

# Efficient Number Theoretic Transform Implementation on GPU for Homomorphic Encryption

Özgün Özerk · Can Elgezen · Ahmet  
Can Mert · Erdinç Öztürk · Erkay Savaş

**Abstract** Lattice-based cryptography forms the mathematical basis for homomorphic encryption, which allows computation directly on encrypted data. Homomorphic encryption enables privacy-preserving applications such as secure cloud computing; yet, its practical applications suffer from the high computational complexity of homomorphic operations. Fast implementations of the homomorphic encryption schemes heavily depend on efficient polynomial arithmetic; multiplication of very large degree polynomials over polynomial rings, in particular. Number theoretic transform (NTT) accelerates polynomial multiplication significantly and therefore, it is the core arithmetic operation in the majority of homomorphic encryption scheme implementations. Therefore, practical homomorphic applications require efficient and fast implementations of NTT in different computing platforms. In this work, we present an efficient and fast implementation of NTT, inverse NTT (INTT) and NTT-based polynomial multiplication operations for GPU platforms. To demonstrate that our GPU implementation can be utilized as an actual accelerator, we experimented with the key generation, the encryption and the decryption operations of the Brakerski/Fan-Vercauteren (BFV) homomorphic encryption scheme implemented in Microsoft's SEAL homomorphic encryption library on GPU, all of which heavily depend on the NTT-based polynomial multiplication. Our GPU implementations improve the performance of these three BFV operations by up to 141.95 $\times$ , 105.17 $\times$  and 90.13 $\times$ , respectively, on Tesla v100 GPU compared to the highly-optimized SEAL library running on an Intel i9-7900X CPU.

**Keywords** Lattice-based Cryptography · Homomorphic Encryption · SEAL · Number Theoretic Transform · Polynomial Multiplication · GPU · CUDA

---

This work is supported by TÜBİTAK under Grant Number 118E725.

Ö. Özerk, C. Elgezen, A. C. Mert, E. Öztürk, E. Savaş  
Faculty of Engineering and Natural Sciences, Sabanci University, Istanbul, Turkey  
E-mail: {ozgunozerk, celgezen, ahmetcanmert, erdinco, erkays}@sabanciuniv.edu

## 1 Introduction

Lattice-based cryptography is conjectured to be secure against attacks from quantum computers and thus supports post-quantum cryptography (PQC). Also, it provides the mathematical basis for fully homomorphic encryption (FHE) schemes, as demonstrated by Gentry in 2009 [20]. FHE allows computation on the encrypted data requiring neither decryption nor secret key, and therefore, enables secure processing of sensitive data. FHE offers a variety of applications ranging from private text classification to secure cloud computing [2].

Since Gentry’s breakthrough, homomorphic encryption has gained tremendous amount of attention and different homomorphic encryption schemes are proposed in the literature such as Brakerski-Gentry-Vaikuntanathan (BGV) [12], Brakerski/Fan-Vercauteren (BFV) [18] and Cheon-Kim-Kim-Song (CKKS) [14]. There are also various efforts for developing their practical implementations. As such, there are different open-source and highly optimized software libraries such as SEAL [39], HELib [24] and PALISADE [34] for homomorphic encryption and computation. The SEAL library is developed by Microsoft Research and it supports the BFV and the CKKS schemes. HELib supports the BGV and the CKKS schemes while PALISADE supports the BGV, the BFV and the CKKS schemes.

Although the potential applications of FHE are of groundbreaking nature, its high algorithmic complexity is a standing impediment for efficient and practical implementations thereof. Among different core arithmetic operations in various FHE schemes, multiplication over polynomial rings is probably the most time-consuming. Therefore, there are different methods in the literature proposed for the efficient implementation of multiplication of two very large degree polynomials over polynomial ring  $\mathbf{R}_{q,n}$ , where  $n$  and  $q$  represent the degree of polynomials in the ring and coefficient modulus, respectively. The Toom-Cook [42] or Karatsuba [25] multiplications have been methods in use for a long time and generally utilized in schemes with polynomial rings, for which NTT is not suitable [31]. NTT-based polynomial multiplication is, on the other hand, highly utilized in lattice-based cryptosystems and it reduces the  $\mathcal{O}(n^2)$  computational complexity of the schoolbook polynomial multiplication to  $\mathcal{O}(n \cdot \log n)$  [15].

Although utilizing NTT improves the performance of polynomial multiplication operation, it is still inefficient for real life applications. Therefore, there are different NTT-based polynomial multiplication implementations proposed in the literature for efficient and practical lattice-based cryptosystems on different platforms: hardware architectures [37, 32, 41, 36], software implementations [40, 24, 1, 39], and implementations on GPUs [16, 3, 6, 4, 28, 22, 27, 23, 44, 5]. There are also hybrid approaches combining NTT-based and Karatsuba multiplication methods for the polynomial multiplication operation [7]. The NTT-based polynomial multiplication operation can be performed for a range of parameters  $n$  and  $q$  in different applications. For example, FHE applications require usually large  $n$  and  $q$  parameters while PQC utilizes smaller

parameters. Therefore, an efficient implementation of NTT-based polynomial multiplication requires flexibility of supporting both FHE and PQC in addition to offering high-performance.

With a similar motivation of the works in the literature, we propose efficient NTT, INTT and NTT-based polynomial multiplication implementations on GPU in this work. The proposed implementations on GPU support a wide range of polynomial rings. The proposed implementations can perform a single NTT and INTT operations in  $39\mu s$  and  $23\mu s$ , respectively, for the largest ring with  $n = 32768$  and  $\log_2(q) = 61$  in Tesla V100 GPU including overhead of kernel calls.

In order to show that the proposed GPU implementations can be useful as actual accelerators in the homomorphic encryption schemes, for proof of concept, the proposed implementations are utilized to implement and accelerate the key generation, the encryption and the decryption operations of the BFV homomorphic encryption scheme on GPU.

For a quick recap, our contributions are listed as follows:

1. We present high-performance and efficient GPU implementations for NTT, INTT and NTT-based polynomial multiplication operations. The proposed GPU implementations support polynomials of degrees ranging from 2048 to 32768 with 30-bit and 61-bit coefficients<sup>1</sup>. We run the implementations on three different GPU platforms, Nvidia GTX 980, Nvidia GTX 1080, Nvidia Tesla V100; and a single NTT operation for polynomials of degree 32768 with 61-bit coefficients is performed in  $73\mu s$ ,  $36\mu s$ ,  $39\mu s$  on Nvidia GTX 980, Nvidia GTX 1080 and Nvidia Tesla V100, respectively.
2. In order to parallelize the NTT and INTT operations on GPU, we modified their algorithmic structures significantly, which consist of three nested *for* loops and removed the dependencies in between as much as possible. A GPU *block* can process a block of 2048 array elements in NTT operation and our implementations benefit from the fast *shared* memory for polynomials of moderate degrees (i.e.,  $n = 8192$  for 30-bit modulus and  $n = 4096$  for 61-bit modulus). For larger polynomial degrees, global memory is used. To support large polynomial degrees in our implementations, we propose a hybrid approach, whereby we make a separate kernel call for each NTT iteration when the NTT block is large; but as soon as the NTT block becomes sufficiently small we switch to a different working mode, in which the rest of the NTT iterations are completed in a single kernel call. The hybrid approach achieves up to  $9\times$  speed up for different polynomial degrees compared to baseline design, which performs  $\log_2(n)$  kernel calls for one NTT operation.
3. The key generation, encryption and decryption operations of the BFV scheme are fully implemented on GPU and, compared to the BFV implementation on Microsoft’s SEAL library running on an Intel i9-7900X CPU; and we observed up to  $60.31\times$ ,  $43.84\times$ ,  $33.89\times$  speed-up values on Nvidia GTX 980;  $56.07\times$ ,  $40.85\times$ ,  $25.05\times$  speed-up values on Nvidia GTX

---

<sup>1</sup> A sample code is available at <https://github.com/SU-CISEC/gpu-ntt>

1080; and  $141.95\times$ ,  $105.17\times$ ,  $90.13\times$  speed-up values on Nvidia Tesla V100, for key generation, encryption and decryption operations, respectively.

4. Since key generation and encryption operations of the BFV scheme require random polynomials from uniform, ternary and discrete Gaussian distributions, we also introduced an implementation of random polynomial sampler for these distributions on GPU. We utilized Salsa20 implementation [21] for pseudo-random number generation for uniform distribution and inverse cumulative standard distribution function `normcdfinvf` in CUDA Math API.

The rest of the paper is organized as follows. Section 2 presents the notation, the background on NTT and polynomial multiplication operations and summarizes prior works in the literature. Section 3 introduces SEAL library. Section 4 presents the proposed GPU implementations. Section 5 presents the results and Section 6 concludes the paper.

## 2 Background

In this section, we present the notation used in the rest of the paper, brief descriptions of NTT, INTT, polynomial multiplication operations, structure of GPUs and prior works in the literature.

### 2.1 Notation

The ring  $\mathbb{Z}_q$  consists of the set of integers  $\{0, 1, \dots, q-1\}$ . Let the polynomial ring  $\mathbf{R}_q = \mathbb{Z}_q[x]/\phi(x)$  represent all the polynomials reduced with the irreducible polynomial  $\phi(x)$  with coefficients in  $\mathbb{Z}_q$ . When  $\phi(x)$  is in the form  $(x^n + 1)$ , the polynomial ring  $\mathbb{Z}_q[x]/(x^n + 1)$  is represented with  $\mathbf{R}_{q,n}$ , which consists of polynomials of degree at most  $(n-1)$  with coefficients in  $\mathbb{Z}_q$ . For example,  $\mathbf{R}_{17,32} = \mathbb{Z}_{17}[x]/(x^{32} + 1)$  represents the polynomials of degree at most 31 with coefficients in  $\mathbb{Z}_{17}$ .

A polynomial  $\mathbf{a}(x) = \sum_{i=0}^{n-1} a_i \cdot x^i$  in  $\mathbf{R}_{q,n}$  can also be represented as a vector over  $\mathbb{Z}_q$ ,  $\mathbf{a} = [a_0, a_1, \dots, a_{n-1}]$ , where  $i$  represents the position of the coefficient such that  $0 \leq i < n$ . Similarly, we use  $\mathbf{a}[i]$  to represent the coefficient of polynomial  $\mathbf{a}$  at position  $i$ . Throughout the paper, we represent an integer and a polynomial with regular lowercase (e.g.  $a$ ) and boldface lowercase (e.g.  $\mathbf{a}$ ), respectively. Vectors in NTT domain are represented with a bar over their symbols. For example,  $\bar{\mathbf{a}}$  represents the NTT domain representation of vector  $\mathbf{a}$ . Let  $\cdot$ ,  $\times$  and  $\odot$  represent integer, polynomial and coefficient-wise vector multiplication, respectively. Let  $(\mathbf{a} \cdot b)$  and  $(\mathbf{a} + b)$  represent that coefficients of polynomial  $\mathbf{a}$  are multiplied and added with integer  $b$ , respectively, if one of the operands is a single integer. Let  $\mathbf{a} \leftarrow \mathbf{R}_{q,n}$  and  $\mathbf{a} \xleftarrow{\$} \mathbf{S}$  represent that the polynomial  $\mathbf{a}$  is sampled uniformly from  $\mathbf{R}_{q,n}$  and from the set  $\mathbf{S}$ , respectively. Let  $\lfloor \cdot \rfloor$ ,  $\lceil \cdot \rceil$ ,  $\lfloor \cdot \rfloor$  and  $[\cdot]_q$  represent round to nearest integer, rounding up, rounding down and the reduction by modulo  $q$  operations, respectively. For

the rest of the paper,  $q$  and  $n$  represent the coefficient modulus and the degree of the polynomial ring, respectively.

## 2.2 Number Theoretic Transform

NTT is defined as discrete Fourier transform (DFT) over the ring  $\mathbf{Z}_q$  and any efficient DFT algorithm can be adopted as an NTT algorithm. An  $n$ -point (pt) NTT operation transforms an  $n$  element vector  $\mathbf{a}$  to another  $n$  element vector  $\bar{\mathbf{a}}$  as defined in Eqn. 1.

$$\bar{\mathbf{a}}_i = \sum_{j=0}^{n-1} a_j \cdot \omega^{ij} \pmod{q} \text{ for } i = 0, 1, \dots, n-1. \quad (1)$$

The NTT calculations involve the constant called twiddle factor,  $\omega \in \mathbb{Z}_q$ , which is also defined as the  $n$ -th root of unity. The twiddle factor satisfies the conditions  $\omega^n \equiv 1 \pmod{q}$  and  $\omega^i \neq 1 \pmod{q} \forall i < n$ , where  $q \equiv 1 \pmod{n}$ .

Similarly, the INTT operation uses almost the same formula as NTT operation as shown in Eqn. 2 except that  $\omega^{-1} \pmod{q}$ , which is the modular inverse of  $\omega$  in  $\mathbb{Z}_q$ , is used instead of  $\omega$  and the resulting coefficients need to be multiplied with  $n^{-1} \pmod{q}$  in  $\mathbb{Z}_q$ .

$$\mathbf{a}_i = \frac{1}{n} \sum_{j=0}^{n-1} \bar{a}_j \cdot \omega^{-ij} \pmod{q} \text{ for } i = 0, 1, \dots, n-1. \quad (2)$$

Applying NTT and INTT operations as in Eqn. 1 and Eqn. 2 leads to high computational complexity. Therefore, there are many efficient and fast implementations of NTT operation in the literature [33, 16, 29, 19], constructed around two very well-known approaches: *decimation in time* (DIT) and *decimation in frequency* (DIF) FFTs. The former and latter FFT operations utilize Cooley-Tukey (CT) and Gentleman-Sande (GS) butterfly structures, respectively [15].

## 2.3 Number theoretic transform based polynomial multiplication

The multiplication of polynomials  $\mathbf{a}(x)$  and  $\mathbf{b}(x)$  can be computed using schoolbook polynomial multiplication as shown in Eqn. 3. When the polynomial multiplication is performed in  $\mathbf{R}_q$ , the resulting polynomial  $\mathbf{c}(x)$  should be reduced by  $\phi(x)$ .

$$\mathbf{c}(x) = \mathbf{a}(x) \times \mathbf{b}(x) = \sum_{i=0}^{n-1} \sum_{j=i}^{n-1} a_i \cdot b_j \cdot x^{i+j} \quad (3)$$

NTT and INTT operations enable efficient implementation of polynomial multiplication operation by converting the schoolbook polynomial multiplication operation into coefficient-wise multiplication operations as shown in

**Algorithm 1** Merged In-place NTT Algorithm

---

**Input:**  $\mathbf{a}(x) \in \mathbf{R}_q^n$  in natural-order  
**Input:**  $\psi_{rev}$  (powers of  $\psi$  stored in bit-reversed order where  $\psi_{rev}[k] = \psi^{br(k)} \pmod{q}$ )  
**Output:**  $\bar{\mathbf{a}}(x) \in \mathbf{R}_q^n$  in bit-reversed order

```

1:  $t = n$ 
2: for ( $m = 1$ ;  $m < n$ ;  $m = 2m$ ) do
3:    $t = t/2$ 
4:   for ( $i = 0$ ;  $i < m$ ;  $i++$ ) do
5:      $j_1 = 2 \cdot i \cdot t$ 
6:      $j_2 = j_1 + t - 1$ 
7:     for ( $j = j_1$ ;  $j \leq j_2$ ;  $j++$ ) do
8:        $U = \mathbf{a}_j$ 
9:        $V = \mathbf{a}_{j+t} \cdot \psi_{rev}[m+i] \pmod{q}$ 
10:       $\mathbf{a}_j = U + V \pmod{q}$ 
11:       $\mathbf{a}_{j+t} = U - V \pmod{q}$ 
12:    end for
13:  end for
14: end for
15: return  $\mathbf{a}$ 

```

---

Eqn. 4 where  $\text{NTT}_{2n}$  and  $\text{INTT}_{2n}$  represent  $2n$ -pt NTT and  $2n$ -pt INTT operations, respectively. However, this requires doubling the sizes of input polynomials with zero-padding and there still should be a separate polynomial reduction operation by  $\phi(x)$ .

$$\mathbf{c}(x) = \text{INTT}_{2n}(\text{NTT}_{2n}(\mathbf{a}(x)) \odot \text{NTT}_{2n}(\mathbf{b}(x))) \pmod{\phi(x)} \quad (4)$$

When the polynomial ring is  $\mathbf{R}_{q,n}$  where  $q \equiv 1 \pmod{2n}$ , a technique called *negative wrapped convolution* is utilized, which eliminates the need for doubling the input sizes and the polynomial reduction operation as shown in Eqns. 5-8.

$$\hat{\mathbf{a}}(x) = [a_0, a_1, \dots, a_{n-1}] \odot [\psi^0, \psi^1, \dots, \psi^{(n-1)}] \quad (5)$$

$$\hat{\mathbf{b}}(x) = [b_0, b_1, \dots, b_{n-1}] \odot [\psi^0, \psi^1, \dots, \psi^{(n-1)}] \quad (6)$$

$$\hat{\mathbf{c}}(x) = \text{INTT}_n(\text{NTT}_n(\hat{\mathbf{a}}(x)) \odot \text{NTT}_n(\hat{\mathbf{b}}(x))) \quad (7)$$

$$\mathbf{c}(x) = [\hat{c}_0, \hat{c}_1, \dots, \hat{c}_{n-1}] \odot [\psi^0, \psi^{-1}, \dots, \psi^{-(n-1)}] \quad (8)$$

However, this requires the coefficients of input and output polynomials to be multiplied with  $[\psi^0, \psi^1, \dots, \psi^{(n-1)}]$  and  $[\psi^0, \psi^{-1}, \dots, \psi^{-(n-1)}]$ , which are usually referred as pre-processing and post-processing, respectively. The constant  $\psi$  is called  $2n$ -th root of unity satisfying the conditions  $\psi^{2n} \equiv 1 \pmod{q}$  and  $\psi^i \neq 1 \pmod{q} \forall i < 2n$ , where  $q \equiv 1 \pmod{2n}$ .

Roy *et al.* [38] merged pre-processing and NTT operations by employing DIT NTT operation utilizing CT butterfly structure, which takes the input in standard order and produces the output in bit-reversed order. The algorithm for merged pre-processing and NTT operations are shown in Algorithm 1, where  $br(k)$  function performs bit-reversal on  $\log_2(n)$ -bit input  $k$ . We refer this operation as NTT for the rest of the paper.

**Algorithm 2** Merged In-place INTT Algorithm

---

**Input:**  $\bar{a}(x)R_q^n$  in bit-reversed order  
**Input:**  $\psi_{rev}^{-1}$  (powers of  $\psi^{-1}$  stored in bit-reversed order where  $\psi_{rev}^{-1}[k] = \psi^{-br(k)} \pmod{q}$ )  
**Output:**  $a(x) \in R_q^n$  in natural-order

```

1:  $t = 1$ 
2: for ( $m = n$ ;  $m > 1$ ;  $m = m/2$ ) do
3:    $j_1 = 0$ 
4:    $h = m/2$ 
5:   for ( $i = 0$ ;  $i < h$ ;  $i++$ ) do
6:      $j_2 = j_1 + t - 1$ 
7:     for ( $j = j_1$ ;  $j \leq j_2$ ;  $j++$ ) do
8:        $U = \bar{a}_j$ 
9:        $V = \bar{a}_{j+t}$ 
10:       $\bar{a}_j = U + V \pmod{q}$ 
11:       $\bar{a}_{j+t} = (U - V) \cdot \psi_{rev}^{-1}[h + i] \pmod{q}$ 
12:    end for
13:  end for
14: end for
15: return  $\bar{a}$ 

```

---

In [35], Pöppelmann *et al.* merged INTT and post-processing operations by employing DIF NTT operation utilizing the GS butterfly, which takes the input in bit-reversed order and produces the output in standard order. The algorithm for merged INTT and post-processing operations are shown in Algorithm 2, and we refer this operation as INTT for the rest of the paper. Using both techniques, pre-processing and post-processing operations can be eliminated at the expense of using two different butterfly structures as shown in Eqn. 9, which formulates the NTT-based polynomial multiplication.

$$\mathbf{c} = \text{INTT}_n(\text{NTT}_n(\mathbf{a}(x)) \odot \text{NTT}_n(\mathbf{b}(x))) \quad (9)$$

For the rest of the paper, we use  $\text{NTT}_n$  and  $\text{INTT}_n$  for representing  $n$ -pt merged NTT and INTT operations, respectively.

## 2.4 Graphical Processing Unit (GPU)

Inner structures of NTT and INTT algorithms offer opportunities for parallelization. An in-depth insight of GPU organization and its fundamental working principals is of significant importance for understanding the techniques and algorithms suitable for GPU implementation introduced in the subsequent sections, which take advantage of inherently parallelizable nature of NTT and INTT algorithms.

While GPUs consist of many more cores than CPUs, GPU cores are significantly slower (less powerful) than CPU cores. Thus, GPUs provide us with a plausible alternative, when it comes to performing overly many, relatively simple operations. On the other hand, when the operation at hand is complex and not possible to be partitioned into smaller, concurrently executable parts, CPU stands always a much better alternative.

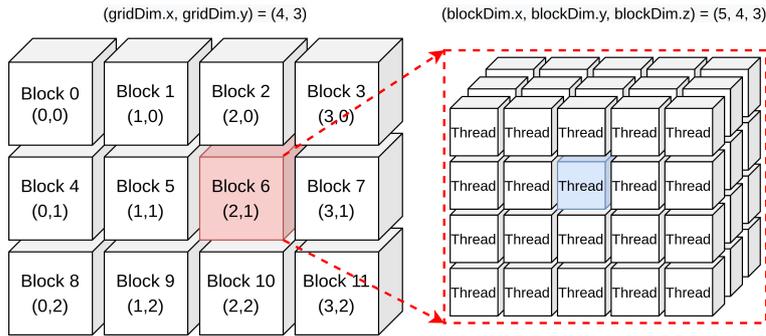


Fig. 1: Structure of a GPU

Kernel in the context of CUDA is a function that is called by the host (CPU) to execute on the device (GPU). Kernels run on streams, each of which can be considered as a single operation sequence. When run on different streams, kernels execute in parallel.

There are three important abstractions in kernels: *grid*, *block*, and *thread*, which are illustrated in Fig. 1. Hierarchically organized, grids consist of blocks, which in turn consist of threads. Each grid can contain a certain maximum number of blocks (actual number varies depending on the GPU model), and these blocks can be organized inside the grid into 1- or 2-dimensional arrays. Every block has an ID determined by the number of dimensions. As shown in Fig. 1, a grid consisting of 12 blocks is organized as an array of 3 rows and 4 columns. The shaded block with ID (2,1) in Fig. 1 indicates that, it is in the first row and the second column of the grid (in other words, it's the 6<sup>th</sup> block). Similar to grids, blocks can contain a certain maximum number of threads (again varies depending on the GPU model); and they can be organized into 1-, 2- or 3-dimensions arrays inside the block. Threads in a block also have indices, which are used to access a specific thread (e.g., the shaded thread in Fig. 1 has the ID derived from its indices (2, 1, 0)) within the block.

Each block is scheduled on a computational unit, referred as *streaming multiprocessor* (SM). After the scheduling, the blocks are then sliced into *warps*, each of which consists of 32 threads. A SM then runs the warps in an arbitrary order determined by the warp-scheduler in the SM.

*Global* memory is the largest memory partition in GPU, accessible by all threads. Among other memory alternatives on GPU, communication with the *global* memory is the most time-consuming; therefore, access to it should be substituted with other memory alternatives, if possible. The data stored in the *constant* memory is cached, and cannot be changed during the execution of operations. The constant memory is significantly faster than the *global* memory. The assignment of each block to SM allows another alternative in the memory hierarchy, the *shared* memory. Unlike *constant* memory, *shared* memory can be modified during execution. It is again faster than the *global* memory, but only accessible by the threads within the same block. In other words, a thread

Table 1: Memory hierarchy inside GPU

Memory Type	A thread can	
	Read	Write
<i>Global Memory</i>	✓	✓
<i>Texture Memory</i>	✓	-
<i>Constant Memory</i>	✓	-
<i>Shared Memory (of its block)</i>	✓	✓
<i>Shared Memory (of another block)</i>	-	-
<i>Registers of other threads (in another warp)</i>	-	-
<i>Registers of other threads (in the same warp)</i>	✓	-
<i>Registers of itself</i>	✓	✓

cannot access other blocks’ *shared* memories. *Registers* are the fastest, yet smallest storage units, which are private to the threads. A thread can neither read or write other thread’s *registers*. The only exception is that threads can read each others’ registers only if they are in the same warp. All these details pertaining to memory access rights are summarized in Table 1.

## 2.5 Prior Works

Efficient implementations of NTT and NTT-based polynomial multiplication operations for various platforms (GPU, FPGA, CPU) and applications (PQC, FHE) have been studied in the literature. GPUs offer a wide variety of optimizations for efficient implementations and they are profitably utilized for accelerating intrinsically parallelizable operations. Moreover, some works in the literature employ more than one GPU by adapting their designs to take advantage of larger degree of parallelism offered by multiple GPUs. For instance, Badawi *et al.* [5] provided the performance evaluation results of their implementation for a variant of the BFV scheme on such multi-GPU design. Another strategy is to take advantage of the memory hierarchy in GPUs. Overall speed of execution can be significantly improved by utilizing faster memory types that are located on the higher part of memory hierarchy (i.e., *shared*) and reducing the communication with global memory [16, 44]. Goey *et al.* [22] proposed a GPU accelerator for NTT operations utilized in FHE applications. Besides very-well known strategies for optimizing memory access in GPUs, they also utilized *warp shuffle*, which allows threads in the same warp to read from each others registers. This enables faster data access than both global and shared memories [22]. Lee *et al.* [27] proposed a method to improve the performance of NTT by eliminating the recursive part of the Nussbaumer algorithm with nega-cyclic convolution, along with the optimizations on the non-coalesced memory access pattern. Lee *et al.* focused on mitigating the warp divergence problem of the NTT operation and they utilized NTT operation in qTESLA, which is a lattice-based post-quantum digital signature scheme [28]. In [23], Gupta *et al.* proposed a GPU accelerator for PQC schemes NewHope and Kyber with an efficient implementation of NTT. In [26], Kim

*et al.* analyzed and improved the performance of NTT operation by introducing on-the-fly twiddle factor generation for FHE schemes. Dai *et al.* [16] proposed cuHE, which is a complete homomorphic encryption library written with CUDA for GPU platforms, creating a solid reference point for future works. Our proposed GPU implementations are compared with these works in Section 5.

### 3 Microsoft’s SEAL Homomorphic Encryption Library [39]

The SEAL homomorphic encryption library is recently developed by Cryptography Research Group at Microsoft Research, which enables fast and efficient homomorphic applications for a variety of applications ranging from private information retrieval to secure neural network inference [8, 13]. SEAL supports two homomorphic encryption schemes, BFV and the CKKS, for implementing homomorphic operations, where the former works with integers while the latter enables homomorphic arithmetic using real numbers. Since, in this work, we utilize the proposed GPU implementation for accelerating the operations of the BFV scheme, details of the CKKS scheme will not be presented in this section.

#### 3.1 BFV Homomorphic Scheme

BFV [18] is a homomorphic encryption scheme proposed by Fan *et al.*, which extends Brakerski’s construction [11]. It is based on ring learning with errors (RLWE) problem [30] and it involves intensive polynomial arithmetic. Let the plaintext and ciphertext spaces be  $\mathbf{R}_{t,n}$  and  $\mathbf{R}_{q,n}$ , respectively, for some integer  $t > 1$ , where neither  $q$  nor  $t$  has to be prime. Suppose  $\Delta = \lfloor q/t \rfloor$  and let  $\chi$  represent discrete Gaussian distribution. Also, while the symbol  $\leftarrow$  represents random sampling from uniform distribution,  $\overset{\$}{\leftarrow} \chi$  stands for sampling from the distribution  $\chi$ . Key generation, encryption and decryption operations described in the textbook-BFV scheme are shown below.

– **Key Generation:**  $\mathbf{s} \leftarrow \mathbf{R}_{2,n}$ ,  $\mathbf{a} \leftarrow \mathbf{R}_{q,n}$  and  $\mathbf{e} \overset{\$}{\leftarrow} \chi$ ,

$$sk = \mathbf{s}, pk = (p_0, p_1) = ([-(\mathbf{a} \times \mathbf{s} + \mathbf{e})]_q, \mathbf{a}).$$

– **Encryption:**  $\mathbf{m} \in \mathbf{R}_{t,n}$ ,  $\mathbf{p}_0, \mathbf{p}_1 \in \mathbf{R}_{q,n}$ ,  $\mathbf{u} \leftarrow \mathbf{R}_{2,n}$  and  $\mathbf{e}_1, \mathbf{e}_2 \overset{\$}{\leftarrow} \chi$ ,

$$ct = (\mathbf{c}_0, \mathbf{c}_1) = ([\mathbf{m} \cdot \Delta + \mathbf{p}_0 \times \mathbf{u} + \mathbf{e}_1]_q, [\mathbf{p}_1 \times \mathbf{u} + \mathbf{e}_2]_q).$$

– **Decryption:**  $\mathbf{c}_0, \mathbf{c}_1 \in \mathbf{R}_{q,n}$  and  $sk \in \mathbf{R}_{2,n}$ ,

$$\mathbf{m} = [\lfloor \frac{t}{q} [\mathbf{c}_0 + \mathbf{c}_1 \times \mathbf{s}]_q \rfloor]_t.$$

### 3.2 Residue Number System (RNS)

For homomorphic computation, an operation is first expressed as a logic or arithmetic circuit. Homomorphic encryption applications are practical only if the multiplicative depth of the circuit, which is to be homomorphically evaluated, is not very high. For example, private information retrieval from a table of 65536 entries [37], which requires a multiplicative depth of at least 4, is reasonably fast when implemented using homomorphic encryption. More complicated homomorphic operations result in larger values of  $n$  and  $q$  due to the increased depth of the circuit.

Efficient arithmetic with large values of the modulus  $q$  is very challenging. Therefore, Residue Number System (RNS) enabling parallelism at algorithmic level for modular integer arithmetic is frequently utilized in the implementations of homomorphic encryption schemes. In RNS, a set of coprime moduli  $q_i$  is used such that

$$q = \prod_{i=0}^{r-1} q_i,$$

where  $r$  is the number of moduli used. Using RNS, arithmetic operations modulo  $q$  (and thus on  $\mathbf{R}_q$ ) can be mapped into operations on smaller  $q_i$  values (or  $\mathbf{R}_{q_i}$ ), which can be performed in parallel. For example, a large 109-bit modulus can be constructed using 3 smaller moduli of sizes 36-bit, 36-bit and 37-bit. The RNS arithmetic requires conversion of a larger integer  $a$  in  $q$  to smaller integers,  $a_i = a \bmod q_i$ , in moduli  $q_i$ . Reconstruction of  $a$  from integers  $a_i$  in moduli  $q_i$ s via Chinese Remainder Theorem (CRT) can be performed as

$$a = \sum_{i=0}^{r-1} a_i \cdot M_i \cdot m_i \pmod{q},$$

where  $M_i = (q/q_i)$  and  $m_i = M_i^{-1} \pmod{q_i}$  for  $i = 0, \dots, r-1$ .

The SEAL library utilizes RNS and implements its arithmetic operations slightly different than the textbook-BFV, where polynomial arithmetic can be performed in parallel for each modulo  $q_i$  [9]. We also use the same approach in our GPU implementation since RNS enables leveraging the parallelism supported in GPU architecture. For the rest of the paper, a polynomial  $\mathbf{a}$  in  $\mathbf{R}_{q,n}$  with  $q = \prod_{i=0}^{r-1} q_i$  will be represented as the array of  $\mathbf{a}^i$  in  $\mathbf{R}_{q_i,n}$  for  $i = 0, 1, \dots, r-1$ .

### 3.3 Implementation of Homomorphic Operations in SEAL Library

The full RNS implementations of key generation, encryption and decryption operations in SEAL are based on Algorithms 3, 4 and 5, respectively. Key generation and encryption operations of SEAL require random polynomials. To this end, SEAL utilizes three different distributions, (i) ternary distribution in  $\mathbf{R}_{2,n}$ , (ii) uniform distribution in  $\mathbf{R}_{q,n}$ , (iii) discrete Gaussian distribution  $\chi$

**Algorithm 3** Implementation of Key Generation Operation in SEAL

---

**Output:**  $\bar{s}^i \in \mathbf{R}_{q,n}$ ,  $\bar{p}_0^i, \bar{p}_1^i \in \mathbf{R}_{q_i,n}$  for  $0 \leq i < r$

- 1:  $\mathbf{s} \leftarrow \mathbf{R}_{2,n}$  ▷ Secret key generation
- 2: **for** ( $i = 0$ ;  $i < r$ ;  $i = i + 1$ ) **do**
- 3:    $\bar{s}^i = \text{NTT}_n(\mathbf{s})$  ▷ Operations in  $\mathbf{R}_{q_i,n}$
- 4: **end for**
- 5:  $\mathbf{e} \xleftarrow{\$} \chi$  ▷ Public key generation
- 6: **for** ( $i = 0$ ;  $i < r$ ;  $i = i + 1$ ) **do**
- 7:    $\bar{\mathbf{a}}^i \leftarrow \mathbf{R}_{q_i,n}$  ▷ Already in NTT domain
- 8:    $\bar{p}_0^i = [-(\bar{\mathbf{a}}^i \odot \bar{s}^i + \text{NTT}_n(\mathbf{e}))]_{q_i}$  ▷ Operations in  $\mathbf{R}_{q_i,n}$
- 9:    $\bar{p}_1^i = \bar{\mathbf{a}}^i$
- 10: **end for**
- 11: **return**  $\bar{\mathbf{s}}, \bar{p}_0, \bar{p}_1$

---

**Algorithm 4** Implementation of Encryption Operation in SEAL

---

**Input:**  $\mathbf{m} \in \mathbf{R}_{t,n}$ ,  $\bar{p}_0^i, \bar{p}_1^i \in \mathbf{R}_{q_i}^n$  for  $0 \leq i < r$

**Output:**  $\mathbf{c}_0^i, \mathbf{c}_1^i \in \mathbf{R}_{q_i,n}$  for  $0 \leq i < (r - 1)$

- 1:  $\mathbf{e}_1, \mathbf{e}_2 \xleftarrow{\$} \chi$
- 2:  $\mathbf{u} \leftarrow \mathbf{R}_{q,n}$
- 3: **for** ( $i = 0$ ;  $i < r$ ;  $i = i + 1$ ) **do**
- 4:    $\bar{\mathbf{u}}^i = \text{NTT}_n(\mathbf{u})$  ▷ Operations in  $\mathbf{R}_{q_i}^n$
- 5:    $\mathbf{c}_0^i = \text{INTT}_n([\bar{p}_0^i \odot \bar{\mathbf{u}}^i + \mathbf{e}_1]_{q_i})$  ▷ Operations in  $\mathbf{R}_{q_i}^n$
- 6:    $\mathbf{c}_1^i = \text{INTT}_n([\bar{p}_1^i \odot \bar{\mathbf{u}}^i + \mathbf{e}_2]_{q_i})$  ▷ Operations in  $\mathbf{R}_{q_i}^n$
- 7: **end for**
- 8: **for** ( $i = 0$ ;  $i < (r - 1)$ ;  $i = i + 1$ ) **do**
- 9:    $\mathbf{t}_0 = [\bar{\mathbf{c}}_0^{r-1} \cdot \lfloor \frac{q_{r-1}}{2} \rfloor]_{q_i}$
- 10:    $\mathbf{t}_1 = [\bar{\mathbf{c}}_1^{r-1} \cdot \lfloor \frac{q_{r-1}}{2} \rfloor]_{q_i}$
- 11:    $\mathbf{c}_0^i = [(\mathbf{c}_0^i - \mathbf{t}_0 + \lfloor \frac{q_{r-1}}{2} \rfloor) \cdot q_{r-1}^{-1}]_{q_i}$
- 12:    $\mathbf{c}_1^i = [(\mathbf{c}_1^i - \mathbf{t}_1 + \lfloor \frac{q_{r-1}}{2} \rfloor) \cdot q_{r-1}^{-1}]_{q_i}$
- 13: **end for**
- 14: **for** ( $j = 0$ ;  $j < n$ ;  $j = j + 1$ ) **do** ▷ Add  $\mathbf{m}$  to  $\mathbf{c}_0$
- 15:    $f = \mathbf{m}[j] \cdot (q \pmod{t}) + \lfloor \frac{t+1}{2} \rfloor$
- 16:   **for** ( $i = 0$ ;  $i < r - 1$ ;  $i = i + 1$ ) **do**
- 17:      $\mathbf{c}_0^i[j] = [\mathbf{c}_0^i[j] + (\mathbf{m}[j] \cdot \frac{q_i}{t} + f)]_{q_i}$  ▷ Operations in  $\mathbf{R}_{q_i}^n$
- 18:   **end for**
- 19: **end for**
- 20: **return**  $\mathbf{c}_0, \mathbf{c}_1$

---

with 0 mean and relatively small standard deviation  $\sigma$ . Besides, all three operations utilize polynomial arithmetic. The key generation operation requires  $2r$  NTT and  $r$  coefficient-wise multiplications of two vectors. The encryption operation requires  $r$  NTT,  $2r$  INTT and  $(6r - 3)$  coefficient-wise multiplications of two vectors. The decryption operation requires  $(r - 1)$  NTT,  $(r - 1)$  INTT and  $(3r - 1)$  coefficient-wise multiplications of two vectors.

In order to find out the most time-consuming operations in key generation, encryption and decryption operations of the SEAL implementation, their timing breakdowns for different parameter sets are obtained on an Intel i9-7900X CPU running at 3.30 GHz  $\times$  20 with 32 GB RAM using GCC version 7.5.0 in Ubuntu 16.04.6 LTS, using SEAL version 3.5 and shown in Table 2.

**Algorithm 5** Implementation of Decryption Operation in SEAL

---

**Input:**  $\bar{s}^i \in \mathbf{R}_{q_i,n}$ ,  $c_0^i, c_1^i \in \mathbf{R}_{q_i,n}$  for  $0 \leq i < (r-1)$ ,  $\gamma$  where  $\gamma > q$  and  $\gcd(\gamma, q) = 1$   
**Output:**  $m \in \mathbf{R}_{t,n}$

```

1: for ( $i = 0$ ;  $i < (r-1)$ ;  $i = i+1$ ) do
2:    $mt^i = \text{INTT}_n(\text{NTT}_n(c_1^i) \odot \bar{s}^i)$                                 ▷ Operations in  $\mathbf{R}_{q_i,n}$ 
3:    $mt^i = [mt^i + c_0^i]_{q_i}$                                           ▷ Operations in  $\mathbf{R}_{q_i,n}$ 
4: end for
5: for ( $i = 0$ ;  $i < (r-1)$ ;  $i = i+1$ ) do                                ▷ Convert from base  $q$  to base  $(t, \gamma)$ 
6:    $mt^i = [mt^i \cdot (t \cdot \gamma)]_{q_i} \cdot q^{-1}$ 
7: end for
8: for ( $j = 0$ ;  $j < 2$ ;  $j = i+1$ ) do                                    ▷  $\text{mod}[t, \gamma]_0 = t$  and  $\text{mod}[t, \gamma]_1 = \gamma$ 
9:   for ( $k = 0$ ;  $k < n$ ;  $k = k+1$ ) do
10:     $acc = 0$ 
11:    for ( $i = 0$ ;  $i < (r-1)$ ;  $i = i+1$ ) do
12:      $acc = (acc + mt^i[k] \cdot q) \text{ mod } [t, \gamma]_j$ 
13:    end for
14:     $mt\gamma_j[k] = acc \text{ mod } [t, \gamma]_j$ 
15:  end for
16: end for
17:  $mt\gamma_0 = [mt\gamma_0 \cdot (-q)^{-1}]_t$                                 ▷ Divide polynomial by  $(-q)$  in  $\mathbb{Z}_t$ 
18:  $mt\gamma_1 = [mt\gamma_1 \cdot (-q)^{-1}]_\gamma$                                 ▷ Divide polynomial by  $(-q)$  in  $\mathbb{Z}_\gamma$ 
19: for ( $i = 0$ ;  $i < n$ ;  $i = i+1$ ) do
20:   if  $mt\gamma_1[i] > \frac{\gamma}{2}$  then
21:     $m[i] = [mt\gamma_1[i] + \gamma - mt\gamma_0[i]]_t$ 
22:   else
23:     $m[i] = [mt\gamma_1[i] - mt\gamma_0[i]]_t$ 
24:   end if
25: end for
26:  $m = [m \cdot \gamma^{-1}]_t$ 
27: return  $m$ 

```

---

Table 2: Timing breakdown of arithmetic operations in key generation, encryption and decryption operations in terms of percentage (%)

Operation	(2048,54,1)			(8192,218,5)			(32768,881,16)		
	Key	Enc	Dec	Key	Enc	Dec	Key	Enc	Dec
PA	10.5	13.8	49.8	25.7	43.2	73.6	32.2	59.1	74.2
$\leftarrow \mathbf{R}_{2,n}$	19.1	12.5	–	6.2	6.4	–	3.9	4.7	–
$\leftarrow^{\$} \chi$	35.8	67.0	–	16.4	37.8	–	7.5	18.8	–
$\leftarrow \mathbf{R}_{q,n}$	23.9	–	–	49.2	–	–	53.9	–	–
Other	10.7	6.7	50.2	2.5	12.6	26.4	2.5	17.4	25.8

PA: Polynomial Arithmetic; Key: Key Generation; Enc: Encryption; Dec: Decryption.

The results are obtained for three parameter sets targeting 128-bit security  $(n, \log_2(q), r) \in \{(2048, 54, 1), (8192, 218, 5), (32768, 881, 16)\}$  with  $t = 1024$ . As shown in Table 2, the percentage of polynomial arithmetic in overall execution times increases as polynomial degree  $n$  and ciphertext modulus  $q$  are increased. For  $n = 32768$  and  $\log_2(q) = 881$ , polynomial arithmetic constitutes 32.2%, 59.1% and 74.2% of the execution time for key generation, encryption and decryption operations, respectively.

## 4 Proposed Implementation

In this section, we will discuss the design of the proposed GPU implementations and the optimizations thereof, in a bottom-up fashion. We will start from efficient modular reduction operation implementation, then describe the NTT implementation and the optimizations. Finally, we will present GPU implementations of key generation, encryption and decryption operations of the BFV scheme.

Our proposed GPU implementations comes in two versions. The first version supports modular arithmetic with modulus up to 30-bit while the second does up to 61-bit modulus. For the rest of the paper, the implementations with 30-bit and 61-bit coefficient modulus  $q$  are referred as 30-bit and 61-bit designs, respectively. The proposed GPU implementations support polynomials of degrees ranging from 2048 to 32768 where polynomial degree is power of two.

### 4.1 Barrett Reduction Algorithm and 128-bit Unsigned Integer Library

---

#### Algorithm 6 Barrett Reduction Algorithm for Modular Multiplication

---

**Input:**  $a, b < q, w = K = \lceil \log_2(q) \rceil, \mu = \lfloor \frac{2^{2K}}{q} \rfloor$   
**Output:**  $c = a \cdot b \bmod q$

```

1:  $t \leftarrow a \cdot b$ 
2:  $s \leftarrow \lfloor \frac{t \cdot \mu}{2^{2w}} \rfloor$   $\triangleright s \leftarrow (t \cdot \mu) \ggg 2w$  (i.e. right shift by  $2w$  bits)
3:  $r \leftarrow s \cdot q$ 
4:  $c \leftarrow t - r$ 
5: if  $c \geq q$  then
6:   return  $c - q$ 
7: else
8:   return  $c$ 
9: end if

```

---

In NTT, INTT and polynomial multiplication operations, there is a substantial number of multiplications modulo  $q$ , which take a considerable amount of time. Therefore, a lightweight and efficient modular reduction algorithm is essential to have high performance implementation. As our modular reduction algorithm, we chose Barrett reduction [10].

The Barrett reduction algorithm replaces the division in reduction operation with multiplication, shift and subtraction operations, which are all significantly faster than the division operation, as shown in Algorithm 6. Barrett reduction calculates  $c = a \cdot b$  modulo  $q$  by approximation. For a modulus of bit-length  $K$ , when the number to be reduced has bit-length of  $L$ , where  $0 \leq L \leq 2K$ , this approximation becomes exact. Keeping that property in mind, it can be observed that the coefficients of the polynomials in  $\mathbf{R}_{q,n}$  are all smaller than the coefficient modulus,  $q$ . When two coefficients are multiplied, the result is guaranteed to have a bit-length that is smaller than or

equal to  $2K$ . Thus, the Barrett reduction algorithm always returns the correct results in our case. For a modulus  $q$ , the algorithm requires a precomputed value,  $\mu$  defined as

$$\mu = \lfloor \frac{2^{2w}}{q} \rfloor, \quad (10)$$

where  $w$  is takes as bit-length of modulus,  $K$ .

For the Barrett algorithm to fit better into GPU design, here, we tweak the original Barrett reduction in Algorithm 6. The original Barrett reduction algorithm multiplies the result of  $a \cdot b$  by  $\mu$  (see Step 2 of Algorithm 6), which can lead to integer variables as long as  $3K$ -bits in bit length during the computation. In our 30-bit and 61-bit designs, this generates an integer as large as 90-bits and 183-bits, respectively. Storing and performing operations on these integers are demanding, especially on GPU resources. Therefore, to reduce the size of the integers obtained after these operations down to a maximum of  $2K$ -bit, we divide the right shift operation by  $2w$  as shown in Step 2 of Algorithm 6 into two separate  $(w - 1)$  and  $(w + 1)$  right shift operations as shown in Steps 2 and 4 of Algorithm 7, respectively. The first right shift operation is performed after the multiplication  $a \cdot b$ , where the result gets right shifted by  $(w - 1)$ . The second right shift operation is performed after the multiplication with  $\mu$ , where the result gets right shifted by  $(w + 1)$ . Using a modulus  $q$  of at most 61-bits, by this tweak it is guaranteed that the intermediate values will be below or equal to  $2K = 122$  bits, which saves memory and improves the performance. In Algorithm 7, we present the modified Barrett reduction algorithm optimized for GPU implementation.

---

**Algorithm 7** Barrett Modular Multiplication Optimized for in GPU

---

**Input:**  $a, b < q, w = K = \lceil \log_2(q) \rceil, \mu = \lfloor \frac{2^{2K}}{q} \rfloor$   
**Output:**  $c = a \cdot b \bmod p$   
1:  $t \leftarrow a \cdot b$   
2:  $x_1 \leftarrow t \gg (w - 2)$   
3:  $x_2 \leftarrow x_1 \cdot \mu$   
4:  $s \leftarrow x_2 \gg (w + 2)$   
5:  $r \leftarrow s \cdot q$   
6:  $c \leftarrow t - r$   
7: **if**  $c \geq q$  **then**  
8:     **return**  $c - q$   
9: **else**  
10:    **return**  $c$   
11: **end if**

---

Next, for our 61-bit design, we need a way to store unsigned integers up to 128-bits. To that end, we develop a 128-bit unsigned integer library and overload all the arithmetic, logic and binary operators, with a focus on speed. A 128-bit unsigned integer is represented by two 64-bit unsigned integers, which are naturally supported by CUDA and C/C++. One integer represents the least significant 64 bits while the other integer represents most significant

64 bits. During the computations, a multiplication of two 64-bit unsigned integers and subtraction of two 128-bit unsigned integers from each other are performed. These operations are coded using the *PTX* inline assembly feature of CUDA for the maximum possible speed. *PTX* enables inserting pseudo assembly code in a higher level language such as C/C++ provides features that are unavailable otherwise such as setting and using the carry bits. Therefore, multiplication and subtraction operations written in *PTX* compile into fewer number of instructions than those written in C/C++. Indeed, integer multiplication operation written using *PTX* shows slightly better performance than Nvidia’s intrinsic function for multiplication `umul` and its derivatives.

#### 4.2 NTT Implementation on GPU

As shown in Algorithm 1, the NTT operation is performed with 3 nested *for* loops. If *for* loops are unrolled, it can be seen that an  $n$ -pt NTT operation actually consists of  $\log_2(n)$  iterations and there are  $(n/2)$  butterfly operations in each iteration. Iterations of the NTT operation for a toy example (i.e.,  $n = 8$ ) are depicted in Fig. 2, where the box with the letter B illustrates the butterfly operation and the box with powers of  $\psi$  shows the particular power of the  $2n$ -th root of unity used in the corresponding iteration. During execution, results are read from and written back to the same input array  $a$ , which is called *in-place* computation. Since  $n = 8$  in Fig. 2, there are  $\log_2(8) = 3$  iterations and 4 butterfly operations in each iteration. Note that the NTT blocks are halved after each iteration. Our INTT algorithm can be visualized to work in a similar way; but the order of the butterfly operations are somewhat reversed.

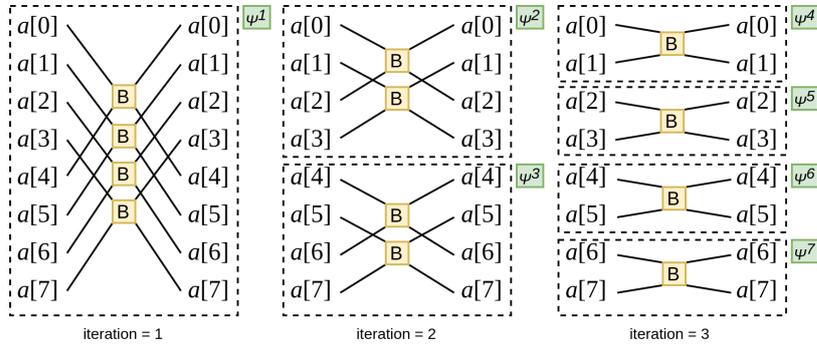


Fig. 2: Visualization of NTT operation for  $n=8$

The challenge in implementing the NTT/INTT algorithms in GPU is to assign the threads efficiently to achieve high utilization. For the best performance, all the threads should be busy, and the workload of each thread should be equivalent. From Fig. 2, for every butterfly operation, B, in each iteration,

---

**Algorithm 8** Scheduling of Array Elements to Threads in GPU for NTT Algorithm
 

---

```

1: for ( $\text{length} = 1$ ;  $\text{length} < n$ ;  $\text{length} = 2 \cdot \text{length}$ ) do
2:    $\text{tid} =$  Global index of thread in GPU
3:    $\text{step} = (n/\text{length})/2$ 
4:    $\text{psi\_step} = \text{tid}/\text{step}$ 
5:    $\text{target\_idx} = (\text{psi\_step} \cdot \text{step} \cdot 2) + (\text{tid} \bmod \text{step})$ 
6:    $\text{step\_group} = \text{length} + \text{psi\_step}$ 
7:    $\text{psi} = \text{psis}[\text{step\_group}]$ 
8:    $U = a[\text{target\_idx}]$ 
9:    $V = a[\text{target\_idx} + \text{step}]$ 
10:   $\triangleright$  Thread assignment is completed at this point
11:   $\triangleright$  Every thread has its corresponding  $U$ ,  $V$  and  $\text{psi}$  values
12: end for

```

---

two elements of the array  $\mathbf{a}$  are used together (see also Steps 10 and 11 of Algorithm 1). Thereby, we elect to use only one GPU thread for two elements of array  $\mathbf{a}$ . Algorithm 8 shows how one thread is scheduled for each butterfly operation with two array elements. From Steps 2-8 of Algorithm 8, we can observe that each thread obtains the necessary values for performing one butterfly operation; namely two elements of the array  $\mathbf{a}$  and the corresponding power of  $\psi$ .

Fig. 3 illustrates how the array elements are assigned to threads in Algorithm 8 for the example when  $n = 8$ . Also, it shows the role of each variable defined in Algorithm 8. In Fig. 3, the array  $\text{psis}$ , which holds the powers of  $\psi$ , and its indices  $\text{psi\_step}$  defined in Algorithm 8 are shown with boxes with solid edges. The boxes with dashed edges indicate the array elements using the same power of  $\psi$ . The array elements in the same dashed box form the *step group*, whose index is kept in the variable  $\text{step\_group}$ . After each iteration, the number of step groups doubles, and the variable  $\text{length}$  keeps track of the number of such boxes. In the step groups, the numbers on the left and on the right stand for the indices of the array elements in Step 8 and Step 9 of Algorithm 8, respectively. The variable  $\text{step}$  is used to determine the second element of the array  $\mathbf{a}$  that is assigned to the same thread. In summary, using the variables  $\text{step}$  and  $\text{length}$ , a thread can easily find out the array elements to process in every iteration independently. For example, in the second iteration, where both  $\text{length}$  and  $\text{step}$  are 2, the first two array elements are assigned to the first two threads; and we use the variable  $\text{step}$  to find the two consecutive elements of the array  $\mathbf{a}$  (here  $\mathbf{a}[2]$  and  $\mathbf{a}[3]$ ) that will be assigned to the first two threads. The elements in the dashed box below are processed in an identical way.

Also, each thread uses the  $\text{step\_group}$  variable to keep track of its step group in each iteration and use it to access the correct element of the array  $\text{psis}$ . For example, in the second iteration in Fig. 3, the first and second step groups consists of array elements with indices 0,1,2,3 and those with indices 4,5,6,7, respectively.

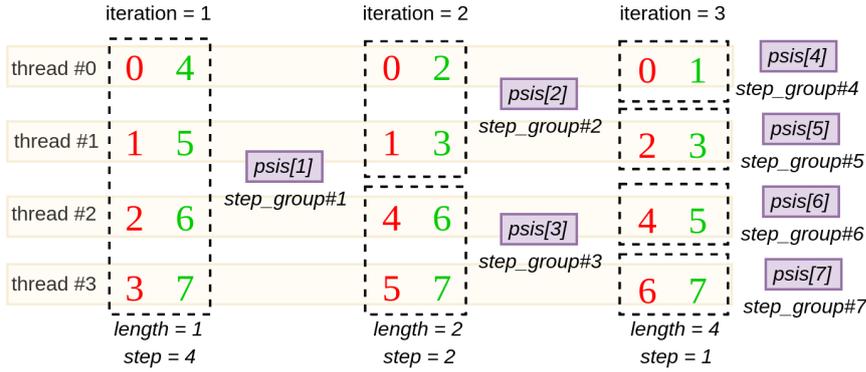


Fig. 3: GPU thread assignment for  $n = 8$

The `target_idx` variable is used by a thread to compute the index of the first array element to be processed by that thread in each iteration (see Step 5 of Algorithm 8 for its computation). The thread first computes the number of array elements in the previous step groups for that iteration ( $\text{psi\_step} \cdot \text{step} \cdot 2$ ), which yields the first array element in its step group. Then, the thread uses its own id to find out the array element that it will process in that iteration.

Note that Algorithm 8 illustrates thread scheduling when the dimension of the input array is at most 2048 as we use one GPU thread per two elements of the array `a` and majority of the current GPU technologies do not support having more than 1024 threads per block to the best of our knowledge. Therefore, all threads in a GPU block performs the same steps in Algorithm 8. Nevertheless, handling an array with more than 2048 elements can be challenging and different approaches are possible. Here, we partition the array into groups of 2048 elements and use multiple GPU blocks to process them in different iterations of NTT computation.

#### 4.2.1 Optimizations of Memory Usage and Kernel Calls

Efficient use of GPU memory hierarchy plays a significant role in the overall performance of our implementations. Coefficients of the input polynomial (array `a`) and different elements of the array `psis` are being constantly reached, rendering the memory access patterns a crucial part of the optimization process. The most basic and direct solution would be utilizing the global memory; however, this also would possibly lead to the worst performance.

**Utilizing Shared Memory:** Since we process the elements of the array `a` in GPU blocks, shared memory offers the best alternative for the threads that are located in the same block. The critical part here is that every element of the array `psis` is being accessed only once per block. Thereby, copying its elements into shared memory would carry no benefit. On the other hand, the elements of `a` are accessed more than once during the computation of NTT/INTT; in

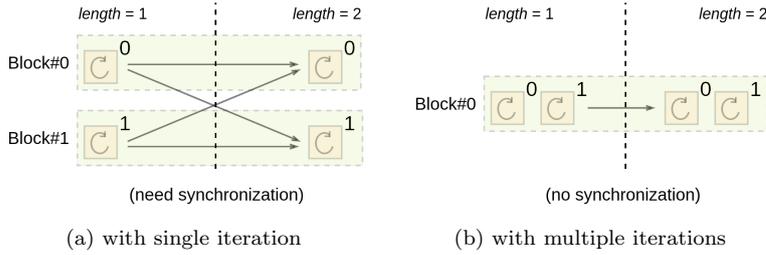


Fig. 4: Single-kernel synchronization for  $n=2^{12}$

fact,  $\log_2(n)$  times, to be precise. Instead of making  $\log_2(n)$  requests to global memory for a single element, we can make only one request to global memory for copying an element of  $\mathbf{a}$  to shared memory, and then access shared memory  $\log_2(n)$  times for it.

#### 4.2.2 Hybrid Approach for Kernel Calls

We also investigate alternative methods for kernel calls. Essentially, we consider two alternatives: i) in **single-kernel approach** we call a kernel once for the computation of NTT/INTT of the whole array (resulting in a total of  $\log_2(n)$  iterations in a single-kernel call); and ii) in **multi-kernel approach** we call a kernel for every iteration for the computation of NTT/INTT operation (resulting in a total of  $\log_2(n)$  kernel calls). Both approaches have their advantages and disadvantages. For smaller sized arrays, single-kernel approach provides better results. However, for larger array sizes (from  $n = 2^{14}$  to  $n = 2^{16}$ , which is determined experimentally on the GPUs used in this study), a multi-kernel approach is advantageous.

In Fig 4, the principle of our single-kernel approach is explained. The boxes with circular arrow symbols represent the operation, in which a GPU block of 1024 threads processes 2048 elements of the array  $\mathbf{a}$ . We have two alternatives for scheduling of array elements to GPU blocks. In the first alternative, every GPU block processes 2048 elements of  $\mathbf{a}$  concurrently as shown in Fig 4a. However, due to the dependencies shown with arrows connecting the boxes in Fig. 4a, a synchronization mechanism is needed for GPU blocks to wait for each other before progressing into the next iteration. Otherwise, race-conditions may occur and lead to wrong results. But, if we assign a group of consecutive  $\text{step} \times 2$  elements of  $\mathbf{a}$  to the same GPU block as shown in Fig. 4b, then no GPU block needs to wait for other GPU blocks processing the other parts of the array  $\mathbf{a}$ . This solution however, exploits limited parallelism offered by the GPU block itself. Multi-kernel approach, on the other hand, can benefit from parallelism beyond the limitation of the single kernel approach.

In the multi-kernel approach, we schedule  $n/2048$  GPU blocks for every kernel call, all of which operate concurrently. Consequently, the multi-kernel

approach can perform better when the array size is sufficiently large. But, every call to kernel incurs certain overhead that results in worse performance than the single-kernel approach for smaller sized arrays. Fully utilizing the potential of parallelism due to multiple kernel calls while reducing the call overhead in an optimal way, we develop the so-called hybrid approach, in which, we start with the multi-kernel approach, but switch to the single-kernel approach when the step groups have fewer number of array elements.

We illustrate the hybrid approach for an array of size  $n = 2^{15}$  in Fig. 5 for 61-bit design. In the leftmost column, we make a kernel call, whereby each of 16 block operates on separate 2048 elements of array  $a$ , in parallel. We make two more kernel calls in the same manner for the second and third iterations of the NTT computation. Until this point, we adopt the multi-kernel approach. However, when the number of elements in a step group is  $2^{12}$ , we switch to single kernel approach. Here, a total of eight GPU blocks process the rest of the iterations on their part of the array in a single kernel call. Note that when the step group size is  $2^{12}$ , the lower and the upper 2048 elements of the array are processed sequentially. At first, this may seem to be rather inefficient. Nevertheless, the alternative would be making an additional kernel call, which brings in the kernel overhead. Then, we can immediately see there is a trade-off between the overhead and the time incurred due to sequential processing of two 2048 elements. We experimentally find out that the kernel call overhead is higher in terms of execution time when the step group size is  $2^{12}$ . In other words, the step group size of  $2^{12}$  is the optimal point to switch from the multi-kernel approach to the single-kernel approach. Note that the switching point can be different for the INTT operation and for the 30-bit design.

In the multi-kernel phase of the hybrid approach, the input array is accessed from the slow global memory, instead of copying it to the faster shared memory. This is due to the fact that the shared memory does not survive across the kernel calls. However, in the single-kernel phase, array elements are copied to the shared memory once, and they are accessed from there for the rest of the iterations as they are in the same block within the same kernel. Table 3 lists the number of NTT iterations that can be performed using shared memory.

The INTT algorithm is slightly different from NTT as it starts from small step groups and merges them after each iteration. Consequently the pattern of kernel calls is the exact opposite of that of the NTT operation; we switch from single to multi-kernel approach. We find the optimal switching point when the step size is 2048 for 61-bit design in INTT computation.

Finally, when they are needed, all CPU-GPU memory transfer operations (`DeviceToHost`, `HostToDevice`) are performed asynchronously to achieve further optimization.

#### 4.2.3 Other Optimizations

Zhang *et al.* proposes a technique to eliminate the final multiplication of coefficients with  $\frac{1}{n}$  in  $\mathbb{Z}_q$  after INTT operation as shown in Eqn. 2 [43]. Instead,

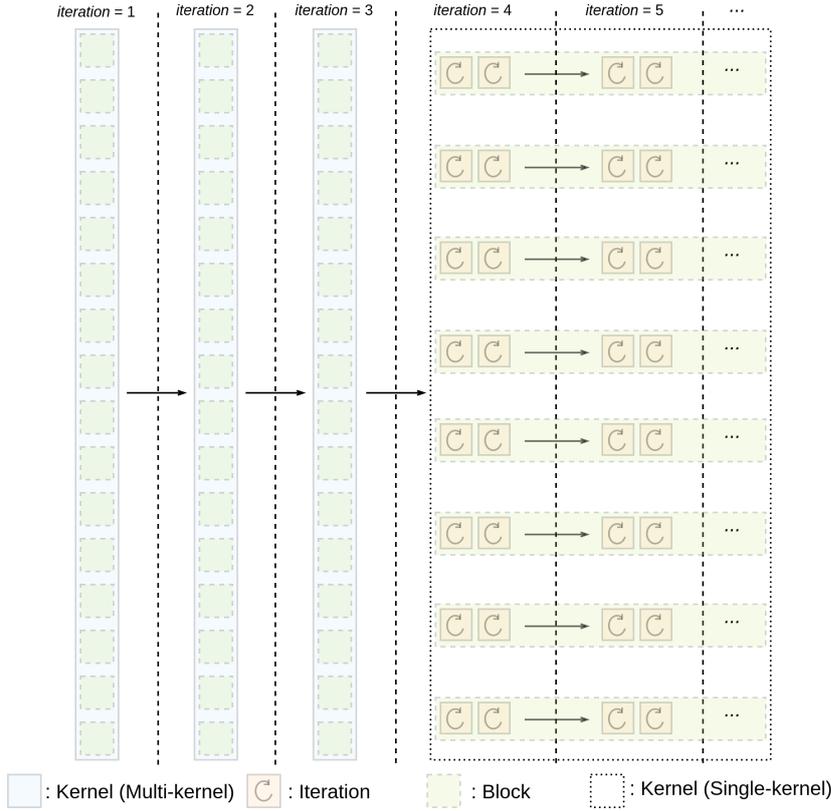


Fig. 5: Hybrid approach internals for  $n = 2^{15}$  and 61-bit moduli

outputs of the GS butterfly operation utilized in INTT can be multiplied with  $\frac{1}{2}$  in  $\mathbb{Z}_q$ , which generates  $\frac{U+V}{2} \pmod{q}$  and  $\frac{(U-V)\psi}{2} \pmod{q}$ . For an odd coefficient modulus  $q$ , division by 2, namely  $\frac{a}{2}$ , can be performed in  $\mathbb{Z}_q$  as shown in Eqn. 11.

$$\frac{a}{2} \pmod{q} = a \cdot \left( \frac{q+1}{2} \right) \quad (11)$$

If  $a$  is an even integer, expression in Eqn. 11 becomes  $(a \gg 1)$ , and if  $a$  is an odd integer, the expression becomes  $(a \gg 1) + \frac{q+1}{2}$  where  $\gg$  represents right shift operation. This optimization replaces modular multiplication operations after INTT with right shift and addition operations, both of which are faster.

#### 4.3 Implementation of the BFV Scheme on GPU

To demonstrate the performance of the proposed NTT and INTT implementations in a practical setting, we implemented the key generation, the encryption

Table 3: Shared memory utilization with hybrid approach

$n$	TI	30-bit				61-bit			
		NTT		INTT		NTT		INTT	
		I(S-)	I(S+)	I(S-)	I(S+)	I(S-)	I(S+)	I(S-)	I(S+)
2048	11	0	11	0	11	0	11	0	11
4096	12	0	12	0	12	0	12	1	11
8192	13	0	13	1	12	1	12	2	11
16384	14	1	13	2	12	2	12	3	11
32768	15	2	13	3	12	3	12	4	11

TI: Total number of iterations

I(S-): Number of iterations performed without shared memory

I(S+): Number of iterations performed with shared memory

and the decryption operations of the BFV scheme on GPU. For these operations, all necessary parameters and powers of  $\psi$  are generated on CPU of the host computer, then sent to GPU prior to any computation. GPU uses the same parameters for the key generation, the encryption and the decryption operations until a new parameter set is sent to GPU. We adopt an approach, referred as *kernel merge strategy*, whereby we find the optimal number of kernel calls that yields the best performance.

#### 4.3.1 Key Generation

Key generation of the BFV scheme is implemented as shown in Algorithm 3. The key generation algorithm takes  $n$ ,  $q_i$  and  $\chi$  as inputs and generates  $r$  secret key polynomials and  $2r$  public key polynomials in  $\mathbf{R}_{q_i, n}$  for  $0 \leq i < r$ . The key generation operation requires randomly generated polynomials from uniform, ternary and discrete Gaussian distributions, which require a random polynomial sampler. Also, it performs NTT and coefficient-wise multiplication operations.

For the generation of random polynomials, we first generate a sequence of cryptographically secure random bytes and, then converted them into random integers from desired distributions (uniform, ternary, discrete Gaussian). We use the Salsa20 stream cipher [17] for generating random bytes by utilizing an existing implementation [21].

For discrete Gaussian distribution, consecutive 4 byte outputs from Salsa20 are interpreted as an unsigned 32-bit integer and converted to a number between 0 and 1. If the converted value is either 0 or 1, we add or subtract the smallest floating point number, defined in CUDA to get a value strictly between 0 and 1. Then, we applied inverse cumulative standard distribution function `normcdfinvf` in CUDA Math API to a value, which results in a random number from discrete Gaussian distribution with mean 0 and standard deviation 1. Since the SEAL library uses a discrete Gaussian distribution with standard deviation 3.2, we multiply the number with 3.2.

For uniform distribution in  $\mathbb{Z}_{q_i}$ , consecutive 8 byte outputs from Salsa20 are interpreted as an unsigned 64-bit integer, converted to a value between 0 and 1, and then multiplied by  $(q_i - 1)$ . For ternary distribution, each byte output from Salsa20 is interpreted as an unsigned 8-bit integer, mapped to a value between 0 and 2, decremented by 1 and finally its fractional part is discarded to obtain either  $-1$ , 0 or 1.

As shown in Algorithm 3, random polynomial generation from uniform distribution, NTT and coefficient-wise multiplication operations are performed for different  $q_i$  values in every iteration of the *for* loops. By following the kernel merge strategy, instead of invoking a separate kernel call for every iteration of the *for* loops, we invoke a single kernel call to handle all the iterations, avoiding the overhead of calling multiple kernels.

#### 4.3.2 Encryption

The encryption of the BFV scheme is implemented as shown in Algorithm 4. The encryption algorithm takes  $n$ ,  $q_i$ ,  $\chi$ , one plaintext polynomial in  $\mathbf{R}_{t,n}$  and  $2r$  public key polynomials in  $\mathbf{R}_{q_i,n}$  as inputs and generates  $2r$  ciphertext polynomials in  $\mathbf{R}_{q_i,n}$  for  $0 \leq i < r$ . The encryption operation requires randomly generated polynomials from ternary and discrete Gaussian distributions. Also, it performs polynomial arithmetic including mainly NTT, INTT and coefficient-wise multiplication operations. As in the key generation implementation, we use the kernel merge strategy. Here, we invoke a single kernel call for generating random polynomials from discrete Gaussian and uniform distributions. Also, we use a single kernel call for each *for* loop.

#### 4.3.3 Decryption

The decryption of the BFV scheme is implemented as shown in Algorithm 5. The decryption algorithm takes  $n$ ,  $q_i$ ,  $\gamma$  parameters and  $r$  secret key polynomials in  $\mathbf{R}_{q_i,n}$  and  $2r$  ciphertext polynomials in  $\mathbf{R}_{q_i,n}$  as inputs and generates one plaintext polynomials in  $\mathbf{R}_{t,n}$  for  $0 \leq i < r$ . The decryption operation does not utilize randomly generated polynomials. It performs polynomial arithmetic including mainly NTT, INTT and coefficient-wise multiplication operations.

In addition to kernel merging strategy, different streams are also utilized here. When not specified explicitly, a kernel is scheduled on the default stream; and thus all operations are being executed sequentially. For concurrent operations, we use multiple streams, which allows to run them in parallel.

## 5 Implementation Results

Performance results of the proposed NTT and INTT implementations are obtained on three different GPUs: Nvidia GTX 980, Nvidia GTX 1080 and Nvidia Tesla V100 with test environments shown in Table 4. Performance results of the SEAL are obtained on an Intel i9-7900X CPU running at 3.30

Table 4: Test environments

Specifications	GPUs		
	GTX 980	GTX 1080	Tesla V100
# of cores	2048	2560	5120
Memory	4GB	8GB	32GB
Frequency (MHz)	1127	1607	1230
Bandwidth	224.4 GB/s	320.3 GB/s	897 GB/s
FP64 Performance	155.6 GFLOPS	277.3 GFLOPS	7.066 TFLOPS
	(1:32)	(1:32)	(1:2)
Compute Capability	5.2	6.1	7.0
CPU	Intel E5-2620 v4	Intel i9 7900X	Intel Gold 5122
RAM	20GB	32GB	62GB

GHz  $\times$  20 with 32 GB RAM using GCC version 7.5.0 in Ubuntu 16.04.6 LTS, using SEAL version 3.5. Performance results of the key generation, the encryption and the decryption implementations of BFV scheme on GPU are obtained only for Nvidia Tesla V100 with test environment shown in Table 4. For obtaining the performance results of GPU implementations, we used `nvvp` profiler utilizing `nvprof`. Note that the time required for initialization steps such as sending necessary parameters and powers of  $\psi$  to GPU from CPU are not included in performance results of NTT/INTT operations.

### 5.1 Performance Results and Comparison with SEAL

In order to evaluate the effects of individual optimizations, we experimented with five different designs for 30-bit implementation with different  $n$  values and obtained performance results on Nvidia GTX 980 GPU: (i) the design **D1** uses single kernel for one NTT/INTT operation, (ii) the design **D2** uses multiple kernels for one NTT/INTT operation, (iii) the design **D3** uses single kernel for one NTT/INTT operation and utilizes shared memory, (iv) the design **D4** uses single kernel for one NTT/INTT operation and utilizes warp shuffle (as utilized in [22]), (v) the design **D5** utilizes the proposed hybrid approach for one NTT/INTT operation presented in Section 4. The performance results of all designs are shown in Table 5.

As shown in Table 5, utilizing shared memory (**D3**) improves the performance of the NTT/INTT implementation with single kernel up to 10%, 25% and 17% for  $n$  values 2048, 4096 and 8192, respectively, compared to **D1**. Due to limited capacity of shared memory on GPU, this optimization can only be applied to the polynomials of degree 8192 or less ( $n \leq 8192$ ) for 30-bit implementation. Applying warp shuffle to the implementation with single kernel (**D4**) also shows some improvements, but not as significant as the shared memory utilization. As shown in Table 5, the warp shuffle optimization improves the performance of NTT/INTT up to 10% for different  $n$  values. This is expected since the warp shuffle mechanism can only be used for the first five iterations of NTT/INTT operation.

Table 5: Performance results with optimizations on GTX 980 in  $\mu s$ 

Operation	Design	$\log_2(n)$					
		11	12	13	14	15	16
NTT	<b>D1</b>	12	21	38	73	149	309
	<b>D2*</b>	95	98	103	109	117	127
	<b>D3</b>	10	16	32	–	–	–
	<b>D4</b>	11	21	37	72	145	303
	<b>D5*</b>	10	16	32	42	49	61
INTT	<b>D1</b>	14	24	40	75	151	315
	<b>D2*</b>	88	102	108	111	114	118
	<b>D3</b>	13	17	33	–	–	–
	<b>D4</b>	13	24	41	75	148	305
	<b>D5*</b>	13	17	21	30	36	46

\*: Including kernel overhead for multiple kernels.

Table 6: Performance results of NTT and INTT implementations in  $\mu s$ 

Operation	Platform	$\log_2(q)$	$\log_2(n)$					
			11	12	13	14	15	16
NTT	GTX 980	30	10	16	32	42	49	61
		55	20	36	43	51	73	–
	GTX 1080	30	7	15	16	23	24	28
		55	11	21	27	33	36	–
	Tesla V100	30	7	11.5	22.5	25.5	27.7	39
		55	12.5	22.5	27	29	39	–
INTT	GTX 980	30	13	17	21	30	36	46
		55	25	31	35	41	52	–
	GTX 1080	30	7	9	12	14	16	22
		55	12	15	18	20	24	–
	Tesla V100	30	7.5	13	14.5	16.3	18.3	20.7
		55	12.5	15.5	18	21	23	–

As shown in Table 5, the design with single kernel utilizing shared memory (**D3**) shows  $9.4\times$ ,  $6.1\times$  and  $3.2\times$  better performance than the design with multiple kernels (**D2**) for  $n$  values 2048, 4096 and 8192, respectively. On the other hand, utilizing multiple kernels per NTT/INTT operation shows better performance in case of  $n > 8192$ . The design with multiple kernels (**D2**) shows  $1.5\times$ ,  $1.5\times$  and  $2.4\times$  better performance than the design with single kernel without shared memory utilization (**D1**) for  $n$  values 16384, 32768 and 65536, respectively. Therefore, as explained in Section 4, we combined both approaches and utilized the hybrid approach with shared memory (**D5**). The design with hybrid approach shows  $1.2\text{-}5\times$  better performance for different  $n$  values compared to the design **D1**. Similar comparison results are applicable to the INTT operation as well. Since the design with hybrid approach (**D5**) provides the best result, we obtained the performance results of 30-bit and 61-bit implementations for NTT and INTT with hybrid approach for three different GPU platforms as shown in Table 6.

Table 7: Performance results of the key generation, the encryption and the decryption operations of the BFV scheme in  $\mu s$ 

$(n, \log_2(q), r)$	Operation	SEAL	TW	Speedup
(4096,109,3)	Key Generation	2020.00	123.86	16.31×
	Encryption	1744.20	85.82	20.32×
	Decryption	363.37	79.46	4.57×
(8192,152,4)	Key Generation	4480.00	135.81	32.99×
	Encryption	4009.76	99.93	40.13×
	Decryption	1012.74	87.46	11.58×
(16384,237,5)	Key Generation	10541.00	176.64	59.68×
	Encryption	9286.84	119.26	77.87×
	Decryption	2688.53	104.13	25.82×
(32768,496,9)	Key Generation	35125.00	273.73	128.32×
	Encryption	29038.60	276.10	105.17×
	Decryption	11365.70	160.05	71.01×
(32768,880,16)	Key Generation	60729.00	427.81	141.95×
	Encryption	48586.70	514.73	94.39×
	Decryption	22215.50	246.48	90.13×

TW: This work.

The key generation, the encryption and the decryption operations of the BFV scheme are implemented on GPU as explained in Section 4. The operations are implemented for five different parameter sets  $(n, \log_2(q), r) \in \{(4096, 109, 3), (8192, 152, 4), (16384, 237, 5), (32768, 496, 9), (32768, 880, 16)\}$  with  $t = 1024$  being used in the SEAL library and targeting various security levels. GPU implementations of the key generation, the encryption and the decryption utilize the proposed efficient NTT and INTT implementations. The performance results of the SEAL library and the proposed GPU implementations for the given set of parameters are presented in Table 7.

As shown in Table 7, the speedup values by GPU implementation compared to the SEAL library become larger as the the ring parameters are increased. For  $(n, \log_2(q), r) = (32768, 880, 16)$ , the GPU implementations of the key generation, the encryption and the decryption operations are improved by 141.95×, 94.39× and 90.13×, respectively, compared to those of the SEAL library running on CPU. There are mainly two reasons for this improvement: (i) the percentage of polynomial arithmetic increases as the ring parameters are increased as shown in Table 2, (ii) as  $r$  is increased, GPU platform leverages parallelism better compared to the CPU platform. For the sake of simplicity, we only present best improvement results for GPU platforms Nvidia GTX 980 and Nvidia GTX 1080. We observe up to 60.31×, 43.85×, 33.89× speed up on Nvidia GTX 980; and 53.08×, 40.85×, 25.06× speed up on Nvidia GTX 1080, for the key generation, the encryption and the decryption operations, respectively.

## 5.2 Comparison with Prior Works

There are plenty of works in the literature targeting acceleration of FHE and PQC schemes on GPU. In [4], Badawi *et al.* accelerates polynomial multiplication operation implemented on NFFlib [1] using GPU with a variant of Stockham NTT algorithm. They present performance results only up to  $n = 8192$  with  $q$  around 400-bit for polynomial multiplication operation. However, their implementation does not utilize RNS; therefore, it is not easy to compare our work with [4]. In [6,5], Badawi *et al.* implements and accelerates some arithmetic operations of the BFV scheme on GPU. However, their work use Discrete Galois Transform (DGT) instead of NTT. In [6], their GPU implementation for the key generation, the encryption and the decryption operations of the BFV scheme achieves  $5\times$ ,  $7\times$  and  $22\times$  speedup on average, respectively, on Tesla P100 for different parameter sets compared to the SEAL library (version 2.3). Our work, on the other hand, improves the performance of the key generation, the encryption and the decryption operations of the BFV scheme by  $75.8\times$ ,  $67.6\times$  and  $40.6\times$  on average, respectively, on Tesla V100 compared to the SEAL library (version 3.5).

Some of the works [28,27,23] focus on PQC schemes utilizing small  $n$  (namely 1024 or less) values and fixed  $q$ . Therefore, they are not applicable to FHE schemes and they are not included in the comparison.

The works in [16,44,3,22,26] focus on accelerating FHE schemes and they also present separate performance results for NTT/INTT operations. We compare our results with these works in Table 8. Compared to the works in [16,44,22,26], our implementations on GTX 1080 and Tesla V100 show better timing performance. The work in [3] outperforms our implementations; however, it uses a fixed special prime as moduli.

The works in [16,3,22,44] utilize a special prime,  $q = 2^{64} - 2^{32} + 1$ , called Solinas prime, which is used as a *carrier* modulus for NTT/INTT operations in  $\mathbf{R}_{q_i,n}$ . This special prime enables following two very useful properties: (i) efficient modular reduction operation using only addition and subtraction, (ii)  $64^{th}$  root of unity of  $q$  is 8, which enables converting multiplication operation with twiddle factor into simple shift operation. Although using this prime as carrier modulus reduces the computational complexity and improves the performance, it comes with the following limitations: (i) size of the input polynomials needs to be doubled (from  $n$  to  $2n$ ) and upper  $n$  coefficients should be zero-padded for NTT-based polynomial multiplication (i.e. 32768-pt NTT operation should be used for a ring with  $n = 16384$ ), (ii) although NTT operation is performed in  $\mathbf{R}_{q,n}$  where  $q$  is the special prime,  $q_i$  of encryption scheme is constrained by  $q_i^2 \cdot n < q$  (i.e.  $q_i$  can be a 24-bit modulus at most for  $q = 2^{64} - 2^{32} + 1$  and  $n = 65536$ ), (iii) the resulting polynomial after INTT operation requires polynomial reduction with  $(x^n + 1)$  for bringing its  $2n$  size back to  $n$ , (iv) multiplying more than two polynomials requires going back-and-forth between NTT and polynomial domains.

In our work, we do not employ a carrier modulus and, thus we are not restricted by the aforementioned limitations. We do not have to double the

Table 8: Comparison table

Work	Platform	$n$	$\log_2(q)$	NTT ( $\mu s$ )	INTT ( $\mu s$ )
[16]*	GTX 690	16384	64 <sup>†</sup>	56	65.3
		32768	64 <sup>†</sup>	71.2	83.6
[16]*, <sup>a</sup>	Tesla K80	16384	64 <sup>†</sup>	12.9	12.5
		32768	64 <sup>†</sup>	19	21.6
[16]*, <sup>b</sup>	GTX 1070	16384	64 <sup>†</sup>	66.8	–
[3]*	Tesla K80	16384	64 <sup>†</sup>	9.6	9.7
		32768	64 <sup>†</sup>	15.3	16.2
[22]*	GTX 1070	16384	64 <sup>†</sup>	57.8	–
[44]*	RTX 2080 Ti	32768	64 <sup>†</sup>	83.3	96
[26]	Titan V	16384	60	44.1	–
		32768	60	84.2	–
This Work	GTX 980	16384	55	51	41
		32768	55	73	52
	GTX 1080	16384	55	33	20
		32768	55	36	24
	Tesla V100	16384	55	29	21
		32768	55	39	23

\*: Uses constant  $q=0xFFFFFFFF00000001$ ; <sup>a</sup>: Results are from [3]; <sup>b</sup>: Results are from [22];

<sup>†</sup>: Actual  $q_i$  is restricted by  $q_i^2 n < 2^{64} - 2^{32} + 1$ .

size of input polynomials, apply polynomial reduction operation or go back-and-forth between NTT and polynomial domains for multiplying more than two polynomials.

## 6 Conclusion and Future Work

In this work, we proposed efficient GPU implementations of NTT and INTT operations to be used in NTT-based polynomial multiplication. The proposed implementations can perform a single NTT and INTT operations in  $39\mu s$  and  $23\mu s$ , respectively, including kernel overhead for the largest ring with  $n = 32768$  and  $\log_2(q) = 61$  on Tesla V100 GPU. The proposed GPU implementations are utilized for accelerating the key generation, the encryption and the decryption operations of the BFV homomorphic encryption scheme, which use NTT and INTT operations frequently. The proposed GPU implementations improve the performance of the key generation, the encryption and the decryption operations by up to  $141.95\times$ ,  $105.17\times$  and  $90.13\times$ , respectively, on Tesla V100 compared to the implementations on SEAL library running on an Intel i9-7900X CPU. The proposed GPU implementations prove to be useful as accelerators for computation-demanding FHE schemes and we also showed that utilizing GPUs as accelerators for FHE libraries such as SEAL is very promising.

As future work, we target accelerating homomorphic multiplication and generating a complete GPU accelerator library for BFV and CKKS schemes, which are required for practical FHE applications. Finally, with minor modifications the proposed NTT and INTT implementations can be used to accelerate PQC schemes with smaller polynomial rings or NTT-unfriendly lattice-based cryptosystems.

## References

1. Aguilar-Melchor, C., Barrier, J., Guelton, S., Guinet, A., Killijian, M.O., Lepoint, T.: Nflib: Ntt-based fast lattice library. In: *Topics in Cryptology - CT-RSA 2016*, pp. 341–356. San Francisco, CA, USA
2. Al Badawi, 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)
3. Al Badawi, A., Veeravalli, B., Aung, K.M.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. IEEE (2018)
4. Al Badawi, A., Veeravalli, B., Aung, K.M.M., Hamadicharef, B.: Accelerating subset sum and lattice based public-key cryptosystems with multi-core cpus and gpus. *Journal of Parallel and Distributed Computing* **119**, 179–190 (2018)
5. Al Badawi, A., Veeravalli, B., Lin, J., Xiao, N., Kazuaki, M., Khin Mi Mi, A.: Multi-gpu design and performance evaluation of homomorphic encryption on gpu clusters. *IEEE Transactions on Parallel and Distributed Systems* **32**(2), 379–391 (2021)
6. 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* pp. 70–95 (2018)
7. Alkim, E., Bilgin, Y.A., Cenk, M.: Compact and simple rlwe based key encapsulation mechanism. In: *International Conference on Cryptology and Information Security in Latin America*, pp. 237–256. Springer (2019)
8. Angel, S., Chen, H., Laine, K., Setty, S.: Pir with compressed queries and amortized query processing. In: *2018 IEEE Symposium on Security and Privacy (SP)*, pp. 962–979. IEEE (2018)
9. Bajard, J.C., 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. NL, Canada
10. Barrett, P.: Implementing the Rivest Shamir and Adleman Public Key Encryption Algorithm on a Standard Digital Signal Processor. *Advances in Cryptology — CRYPTO’86* **263**, 311–323 (1986)
11. Brakerski, Z.: Fully homomorphic encryption without modulus switching from classical gapsvp. In: *Annual Cryptology Conference*, pp. 868–886. Springer (2012)
12. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (leveled) fully homomorphic encryption without bootstrapping. *ACM Transactions on Computation Theory (TOCT)* **6**(3), 1–36 (2014)
13. Brutzkus, A., Elisha, O., Gilad-Bachrach, R.: Low latency privacy preserving inference. In: *International Conference on Machine Learning* (2019)
14. Cheon, J.H., Kim, A., Kim, M., Song, Y.: Homomorphic encryption for arithmetic of approximate numbers. In: *International Conference on the Theory and Application of Cryptology and Information Security*, pp. 409–437. Springer (2017)
15. Chu, E., George, A.: *Inside the FFT black box: serial and parallel fast Fourier transform algorithms*. CRC press (1999)
16. Dai, W., Sunar, B.: cuhe: A homomorphic encryption accelerator library. In: *International Conference on Cryptography and Information Security in the Balkans*, pp. 169–186. Springer (2015)

17. D.J., B.: The salsa20 family of stream ciphers, lecture notes in computer science **4986**, 84–97 (2008). DOI 10.1007/978-3-540-68351-3\_8
18. Fan, J., Vercauteren, F.: Somewhat practical fully homomorphic encryption. Cryptology ePrint Archive, Report 2012/144 (2012)
19. Feng, X., Li, S., Xu, S.: RLWE-Oriented High-Speed Polynomial Multiplier Utilizing Multi-lane Stockham NTT Algorithm. IEEE Transactions on Circuits and Systems II: Express Briefs pp. 1–1 (2019). DOI 10.1109/TCSII.2019.2917621
20. Gentry, C., Boneh, D.: A fully homomorphic encryption scheme, vol. 20. Stanford university Stanford (2009)
21. Ghosh, M.: Salsa20 cuda. [https://github.com/moinakg/salsa20\\_core\\_cuda](https://github.com/moinakg/salsa20_core_cuda)
22. Goey, J.Z., Lee, W.K., Goi, B.M., Yap, W.S.: Accelerating number theoretic transform in gpu platform for fully homomorphic encryption. The Journal of Supercomputing pp. 1–20 (2020)
23. Gupta, N., Jati, A., Chauhan, A.K., Chattopadhyay, A.: Pqc acceleration using gpus: Frodokem, newhope and kyber. IEEE Transactions on Parallel and Distributed Systems pp. 1–1 (2020)
24. Halevi, S., Shoup, V.: Algorithms in helib. In: Advances in Cryptology – CRYPTO 2014, pp. 554–571. Santa Barbara, CA, USA (Aug. 2014)
25. Karatsuba, A.A., Ofman, Y.P.: Multiplication of many-digital numbers by automatic computers. In: Doklady Akademii Nauk, vol. 145, pp. 293–294. Russian Academy of Sciences (1962)
26. Kim, S., Jung, W., Park, J., Ahn, J.H.: Accelerating number theoretic transformations for bootstrappable homomorphic encryption on gpus. 2020 IEEE International Symposium on Workload Characterization (IISWC) (2020). DOI 10.1109/iiswc50251.2020.00033. URL <http://dx.doi.org/10.1109/IISWC50251.2020.00033>
27. Lee, W.K., Akleyek, S., Wong, D.C.K., Yap, W.S., Goi, B.M., Hwang, S.O.: Parallel implementation of nussbaumer algorithm and number theoretic transform on a gpu platform: application to qtesla. The Journal of Supercomputing pp. 1–26 (2020)
28. Lee, W.K., Akleyek, S., Yap, W.S., Goi, B.M.: Accelerating number theoretic transform in gpu platform for qtesla scheme. In: International Conference on Information Security Practice and Experience, pp. 41–55. Springer (2019)
29. Longa, P., Naehrig, M.: Speeding up the number theoretic transform for faster ideal lattice-based cryptography. In: Cryptology and Network Security, pp. 124–139. Milan, Italy (Nov. 2016)
30. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: Advances in Cryptology – EUROCRYPT 2010, pp. 1–23. French Riviera,
31. Mera, J.M.B., Karmakar, A., Verbauwhede, I.: Time-memory trade-off in toom-cook multiplication: an application to module-lattice based cryptography. IACR Transactions on Cryptographic Hardware and Embedded Systems pp. 222–244 (2020)
32. Mert, A.C., Öztürk, E., Savaş, E.: Design and implementation of encryption/decryption architectures for bfv homomorphic encryption scheme. IEEE Transactions on Very Large Scale Integration (VLSI) Systems **28**(2), 353–362 (2019)
33. Pollard, J.M.: The fast Fourier transform in a finite field. Mathematics of computation **25**(114), 365–374 (1971)
34. Polyakov, Y., Rohloff, K., Ryan, G.W.: Palisade lattice cryptography library user manual. Cybersecurity Research Center, New Jersey Institute of Technology (NJIT), Tech. Rep (2017)
35. Pöppelmann, T., Oder, T., Güneysu, T.: High-performance ideal lattice-based cryptography on 8-bit atxmega microcontrollers. In: International Conference on Cryptology and Information Security in Latin America, pp. 346–365. Springer (2015)
36. Riazi, M.S., Laine, K., Pelton, B., Dai, W.: Heax: An architecture for computing on encrypted data. In: Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS ’20, p. 1295–1309. Association for Computing Machinery, New York, NY, USA (2020). DOI 10.1145/3373376.3378523. URL <https://doi.org/10.1145/3373376.3378523>
37. Roy, S.S., Turan, F., Jarvinen, K., Vercauteren, F., Verbauwhede, I.: Fpga-based high-performance parallel architecture for homomorphic computing on encrypted data. Cryptology ePrint Archive, Report 2019/160 (2019)

38. Roy, S.S., Vercauteren, F., Mentens, N., Chen, D.D., Verbauwhe, I.: Compact ring-lwe cryptoprocessor. In: L. Batina, M. Robshaw (eds.) *Cryptographic Hardware and Embedded Systems – CHES 2014*, pp. 371–391. Springer Berlin Heidelberg, Berlin, Heidelberg (2014)
39. Microsoft SEAL (release 3.6). <https://github.com/Microsoft/SEAL> (2020). Microsoft Research, Redmond, WA.
40. Seiler, G.: Faster avx2 optimized ntt multiplication for ring-lwe lattice cryptography. *IACR Cryptology ePrint Archive* **2018**, 39 (2018)
41. Sinha Roy, S., Järvinen, K., Vliegen, J., Vercauteren, F., Verbauwhe, I.: Hepcloud: An fpga-based multicore processor for fv somewhat homomorphic function evaluation. *IEEE Transactions on Computers* **67**(11), 1637–1650 (2018). DOI 10.1109/TC.2018.2816640
42. Toom, A.L.: The complexity of a scheme of functional elements realizing the multiplication of integers. In: *Soviet Mathematics Doklady*, vol. 3, pp. 714–716 (1963)
43. Zhang, N., Yang, B., Chen, C., Yin, S., Wei, S., Liu, L.: Highly Efficient Architecture of NewHope-NIST on FPGA using Low-Complexity NTT/INTT. *IACR Trans. on CHES* **2020**(2), 49–72 (2020). DOI 10.13154/tches.v2020.i2.49-72
44. Zheng, Z.: Encrypted cloud using gpus (2020)