

Efficient Multi-Key Homomorphic Encryption with Packed Ciphertexts with Application to Oblivious Neural Network Inference

Hao Chen¹, Wei Dai¹, Miran Kim², and Yongsoo Song¹

¹ Microsoft Research, Redmond, USA

{haoche, Wei.Dai, Yongsoo.Song}@microsoft.com

² University of Texas, Health Science Center at Houston, USA

Miran.Kim@uth.tmc.edu

Abstract. Homomorphic Encryption (HE) is a cryptosystem which supports computation on encrypted data. López-Alt et al. (STOC 2012) proposed a generalized notion of HE, called Multi-Key Homomorphic Encryption (MKHE), which is capable of performing arithmetic operations on ciphertexts encrypted under different keys.

In this paper, we present multi-key variants of two HE schemes with packed ciphertexts. We present new relinearization algorithms which are simpler and faster than previous method by Chen et al. (TCC 2017). We then generalize the bootstrapping techniques for HE to obtain multi-key fully homomorphic encryption schemes. We provide a proof-of-concept implementation of both MKHE schemes using Microsoft SEAL. For example, when the dimension of base ring is 8192, homomorphic multiplication between multi-key BFV (resp. CKKS) ciphertexts associated with four parties followed by a relinearization takes about 116 (resp. 67) milliseconds.

Our MKHE schemes have a wide range of applications in secure computation between multiple data providers. As a benchmark, we homomorphically classify an image using a pre-trained neural network model, where input data and model are encrypted under different keys. Our implementation takes about 1.8 seconds to evaluate one convolutional layer followed by two fully connected layers on an encrypted image from the MNIST dataset.

Keywords: multi-key homomorphic encryption; packed ciphertext; ring learning with errors; neural networks

1 Introduction

As large amount of data are being generated and used for driving novel scientific discoveries, the effective and responsible utilization of large data remain to be a big challenge. This issue might be alleviated by outsourcing to public cloud service providers with intensive computing resources. However, there still remains a problem in privacy and security of outsourcing data and analysis. In the past few years, significant progresses have been made on cryptographic techniques for *secure computation*. Among the techniques for secure computation, Multi-Party Computation (MPC) and Homomorphic Encryption (HE) have received increasing attention in the past few years due to technical breakthroughs.

The history of MPC dates back three decades ago [49, 5], and since then it has been intensively studied in the theory community. In this approach, two or more parties participate in an interactive protocol to compute a function on their private inputs, where only the output of the function is revealed to the parties. Recent years witnessed a large body of works on improving the practical efficiency of MPC, and state-of-the-art protocols have achieved orders of magnitude improvements on performance (see e.g. [48, 20, 36]). However, these protocols are still inherently inefficient in terms of communication complexity: the number of bits that the parties need to exchange during the protocol is proportional to the *product* between the complexity of the function and the number of parties. Therefore, the high communication complexity remains the main bottleneck of MPC protocols.

Moreover, the aforementioned MPC protocols may not be desirable for cloud-based applications, as all the parties involved need to perform local computation proportional to the complexity of the function.

However, in practical use-cases, we cannot expect the data providers to either perform large amount of work or stay online during the entire protocol execution. Another model was proposed where the data owners secret-share their data with a small number of independent servers, who perform an MPC to generate the computation result [23, 43]. These protocols have good performance and they moved the burden from the data providers to the servers, but their privacy guarantees rely on the assumption that the servers do not collude.

HE refers to a cryptosystem that allows computing on encrypted data without decrypting them, thus enabling securely outsourcing computation in an untrusted cloud. There have been significant technical advances on HE after Gentry’s first construction [24]. For example, one can encrypt multiple plaintext values into a single *packed* ciphertext, and use the single instruction multiple data (SIMD) techniques to perform operations on these values in parallel [47, 26]. Hence, HE schemes with packing techniques [7, 6, 22, 16] have good amortized complexity per plaintext value, and they have been applied to privacy-preserving big data analysis [39, 11, 37]. However, traditional HE schemes only allow computation on ciphertexts decryptable under the same secret key. Therefore, HE does not naturally support secure computation applications involving multiple data providers, each providing its own secret key.

López-Alt et al. [42] proposed a *Multi-Key Homomorphic Encryption* (MKHE) scheme, which is a cryptographic primitive supporting arithmetic operations on ciphertexts which are not necessarily decryptable to the same secret key. In addition to solving the aforementioned issues of HE, MKHE can be also used to design round-efficient MPC protocols with minimal communication cost [44]. In addition, an MPC protocol from MKHE satisfies the *on-the-fly* MPC [42] property, where the circuit to be evaluated can be dynamically decided after the data providers upload their encrypted data.

Despite its versatility, MKHE has been seldom used in practice. Early studies [19, 44, 45] used a multi-key variant of the GSW scheme [28]. These constructions have large ciphertexts and their performance does not scale well with the number of parties. Previous work [8, 9] proposed MKHE schemes with short ciphertexts, with the caveat that one ciphertext encrypts only a single bit. The only existing MKHE scheme with packed ciphertexts [13, 50] is a multi-key variant of the BGV scheme [7]. Note that all the above studies were purely abstract with no implementation given, and it remains an open problem whether an MKHE scheme supporting SIMD operations can be practical.

1.1 Our Contributions

We design multi-key variants of the BFV [6, 22] and CKKS [16] schemes. We propose a new method for generating a relinearization key which is simpler and faster compared to previous technique in [13]. Furthermore, we adapt the state-of-the-art bootstrapping algorithms for these schemes [12, 14, 9] to the multi-key scenario to build Multi-Key Fully Homomorphic Encryptions with packed ciphertexts. Finally, we give a proof-of-concept implementation of our multi-key schemes using Microsoft SEAL [46] and present experimental results. To the best of our knowledge, this is the first practical implementation of MKHE schemes that support packed ciphertexts.

We also present the first viable application of MKHE that securely evaluates a pre-trained convolutional neural network (CNN) model. We build an efficient protocol where a cloud server provides on-line prediction service to a data owner using a classifier from a model provider, while protecting the privacy of both data and model using MKHE. Our scheme with support for the multi-key operations makes it possible to achieve this at a low end-to-end latency, and near-optimal cost for the data and model providers, as shown in Fig. 1. The server can store numerous ciphertexts encrypted under different keys, but the computational cost of a certain task depends only on the number of parties related to the circuit.

1.2 Overview of Our Construction

Let $R = \mathbb{Z}[X]/(X^n+1)$ be the cyclotomic ring with a power-of-two dimension n , and $s_i \in R$ be the secret of the i -th party. The starting point of the construction of a ring-based MKHE scheme is the requirement that the resulting scheme should be able to handle homomorphic computations on ciphertexts under independently generated secret keys. A ciphertext of our MKHE scheme associated to k different parties is of the form $\overline{\text{ct}} = (c_0, c_1, \dots, c_k) \in R_q^{k+1}$ for a modulus q , which is decryptable by the concatenated

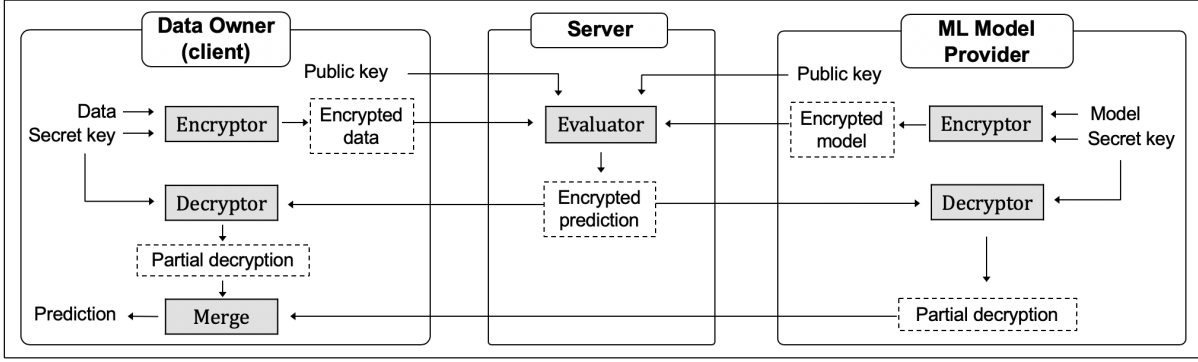


Fig. 1. High-level overview of the application to oblivious neural network inference.

secret $\overline{\mathbf{sk}} = (1, s_1, \dots, s_k)$. In other words, its phase $\mu = \langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \rangle \pmod{q}$ is a randomized encoding of a plaintext message m corresponding to the base scheme.

Homomorphic multiplication of BFV or CKKS consists of two steps: tensor product and relinearization. The tensor product of two input ciphertexts satisfies $\langle \overline{\mathbf{ct}}_1 \otimes \overline{\mathbf{ct}}_2, \overline{\mathbf{sk}} \otimes \overline{\mathbf{sk}} \rangle = \langle \overline{\mathbf{ct}}_1, \overline{\mathbf{sk}} \rangle \cdot \langle \overline{\mathbf{ct}}_2, \overline{\mathbf{sk}} \rangle$, so it is a valid encryption under the tensor squared secret $\overline{\mathbf{sk}} \otimes \overline{\mathbf{sk}}$. In the relinearization step, we aim to transform the extended ciphertext $\overline{\mathbf{ct}} = \overline{\mathbf{ct}}_1 \otimes \overline{\mathbf{ct}}_2 \in R_q^{(k+1)^2}$ into a canonical ciphertext encrypting the same message under $\overline{\mathbf{sk}}$. This step can be understood as a key-switching process which requires a special encryption of $\overline{\mathbf{sk}} \otimes \overline{\mathbf{sk}}$. We note that $\overline{\mathbf{sk}} \otimes \overline{\mathbf{sk}}$ contains entries $s_i s_j$ which depend on secrets of two distinct parties. Hence a relinearization key corresponding to non-linear entries cannot be generated by a single party, different from the traditional HE schemes.

We propose an RLWE-based cryptosystem to achieve this functionality. It looks similar to the ring variant of GSW [28, 21] but our scheme supports some operations between ciphertexts under different keys. Let $\mathbf{g} \in \mathbb{Z}^d$ be an integral vector, called the gadget vector. This scheme assumes the Common Reference String (CRS) model so all parties share a random polynomial vector $\mathbf{a} \in R_q^d$. Each party i generates a special encryption of secret s_i by itself, which is a matrix $\mathbf{D}_i = [\mathbf{d}_{i,0} | \mathbf{d}_{i,1} | \mathbf{d}_{i,2}] \in R_q^{d \times 3}$ satisfying $\mathbf{d}_{i,0} + s_i \cdot \mathbf{d}_{i,1} \approx r_i \cdot \mathbf{g} \pmod{q}$ and $\mathbf{d}_{i,2} \approx r_i \cdot \mathbf{a} + s_i \cdot \mathbf{g} \pmod{q}$ where r_i is a small polynomial sampled from the key distribution. It is published as the evaluation key of the i -th party.

We present two relinearization methods with different advantages. For each pair $1 \leq i, j \leq k$, the first method combines the i -th evaluation key \mathbf{D}_i with the j -th public key $\mathbf{b}_j \approx -s_j \cdot \mathbf{a} \pmod{q}$ to generate $\mathbf{K}_{i,j} \in R_q^{d \times 3}$ such that $\mathbf{K}_{i,j} \cdot (1, s_i, s_j) \approx s_i s_j \cdot \mathbf{g} \pmod{q}$. That is, $\mathbf{K}_{i,j}$ can be used to relinearize one entry $c_{i,j}$ of an extended ciphertext into a triple (c'_0, c'_i, c'_j) such that $c'_0 + c'_i s_i + c'_j s_j \approx c_{i,j} s_i s_j \pmod{q}$. This method can be viewed as a variant of the previous GSW ciphertext extension proposed in [44]. In particular, each row of $\mathbf{K}_{i,j}$ consists of three polynomials in R_q (compared to $O(k)$ dimension of previous work [13, 50]), so that the bit size of a shared relinearization key $\{\mathbf{K}_{i,j}\}_{1 \leq i,j \leq k}$ is $O(dk^2 \cdot n \log q)$ and the complexity of key generation is $O(d^2 k^2)$ polynomial operations modulo q (see Section 3 for details). The relinearization algorithm repeats $O(k^2)$ key-switching operations from $s_i s_j$ to $(1, s_i, s_j)$, so its complexity is $O(dk^2)$ operations in R_q . We note that $\mathbf{K}_{i,j}$ can be pre-computed before multi-key operations, and a generated key can be reused for any computation related to the parties i and j .

Our second approach directly linearizes each of the entries of an extended ciphertext by multiplying the j -th public key \mathbf{b}_j and i -th evaluation key \mathbf{D}_i in a recursive way. The first solution should generate and store a shared relinearization key $\{\mathbf{K}_{i,j}\}_{1 \leq i,j \leq k}$, so its space and time complexity grow quadratically on k . However, the second algorithm allows us to keep only the individual evaluation keys which is linear on k . Furthermore, it significantly reduces the variance of additional noise from relinearization, so that we can use a smaller parameter while keeping the same functionality. The only disadvantage is that, if we exclude the complexity of a shared key generation from the first approach, then the second algorithm entails additional costs (about 1/3 of the complexity of the first relinearization). However, it is ignorable

compared to the overall performance gain from its various advantages. Finally, we adapt the modulus raising technique [27] to the second approach to reduce the noise growth even more.

As an orthogonal issue, the bootstrapping of packed MKHE schemes has not been studied in the literature. We generalize the existing bootstrapping methods for HE schemes [25, 32, 12, 14, 9] to the multi-key setting. The main issue of generalization is that the pipeline of bootstrapping includes some advanced functionalities such as slot permutation. We resolve this issue and provide all necessary operations by applying the multi-key-switching technique in [10] to Galois automorphism.

Finally, we apply the state-of-art optimization techniques for implementing HE schemes [4, 30, 15] to our MKHE schemes for performance improvement. For example, we implement full Residue Number System (RNS) variants of MKHE schemes and use an RNS-friendly decomposition method [4, 30, 15] for relinearization, thereby avoiding expensive high-precision arithmetic.

1.3 Related Works

López-Alt et al. [42] firstly proposed an MKHE scheme based on NTRU. After that, Clear and McGoldrick [19] suggested a multi-key variant of GSW together with ciphertext extension technique to design an MKHE scheme and it was simplified by Mukherjee and Wichs [44]. Peikert and Shiehian [45] developed two multi-hop MKHE schemes based on the same multi-key GSW scheme. However, these schemes could encrypt only a single bit in a huge extended GSW ciphertext.

Brakerski and Perlman [8] suggested an MKHE scheme with short ciphertexts whose length grow linearly on the number of parties involved. Chen, Chillotti and Song [10] improved its efficiency by applying the framework of TFHE [17] with the first implementation of MKHE primitive. However, their scheme does not support the packing technique, thereby having similar inherent (dis)advantages from TFHE.

Chen, Zhang and Wang [13] described a multi-key variant of BGV [7] by adapting the multi-key GSW scheme for generating a relinearization key. Their performance was improved by Zhou et al. [50], however, each key-switching key from $s_i \cdot s_j$ to the ordinary key has $O(k)$ components. In addition, these works did not provide any implementation or empirical result. This study is an extension of these works in the sense that our relinearization method and other optimization techniques can be applied to BGV as well. We also stress that the performance of previous batch MKHE schemes can be improved by observing the sparsity of evaluation keys, but this point was not pointed out in the manuscripts.

2 Background

2.1 Notation

All logarithms are in base two unless otherwise indicated. We denote vectors in bold, e.g. \mathbf{a} , and matrices in upper-case bold, e.g. \mathbf{A} . We denote by $\langle \mathbf{u}, \mathbf{v} \rangle$ the usual dot product of two vectors \mathbf{u}, \mathbf{v} . For a real number r , $\lceil r \rceil$ denotes the nearest integer to r , rounding upwards in case of a tie. We use $x \leftarrow D$ to denote the sampling x according to distribution D . For a finite set S , $U(S)$ denotes the uniform distribution on S . We let λ denote the security parameter throughout the paper: all known valid attacks against the cryptographic scheme under scope should take $\Omega(2^\lambda)$ bit operations.

2.2 Multi-Key Homomorphic Encryption

A multi-key homomorphic encryption is a cryptosystem which allows us to evaluate an arithmetic circuit on ciphertexts, possibly encrypted under different keys.

Let \mathcal{M} be the message space with arithmetic structure. An MKHE scheme **MKHE** consists of five PPT algorithms (**Setup**, **KeyGen**, **Enc**, **Dec**, **Eval**). We assume that each participating party has a reference (index) to its public and secret keys. A multi-key ciphertext implicitly contains an *ordered* set $T = \{id_1, \dots, id_k\}$ of associated references. For example, a fresh ciphertext $\text{ct} \leftarrow \text{MKHE.Enc}(\mu; \text{pk}_{id})$ corresponds to a single-element set $T = \{id\}$ but the size of reference set gets larger as the computation between ciphertexts from different parties progresses.

- **Setup:** $pp \leftarrow \text{MKHE.Setup}(1^\lambda)$. Takes the security parameter as an input and returns the public parameterization. We assume that all the other algorithms implicitly take pp as an input.
- **Key Generation:** $(sk, pk) \leftarrow \text{MKHE.KeyGen}(pp)$. Outputs a pair of secret and public keys.
- **Encryption:** $ct \leftarrow \text{MKHE.Enc}(\mu; pk)$. Encrypts a plaintext $\mu \in \mathcal{M}$ and outputs a ciphertext $ct \in \{0, 1\}^*$.
- **Decryption:** $\mu \leftarrow \text{MKHE.Dec}(\overline{ct}; \{sk_{id}\}_{id \in T})$. Given a ciphertext \overline{ct} with the corresponding sequence of secret keys, outputs a plaintext μ .
- **Homomorphic evaluation:**

$$\overline{ct} \leftarrow \text{MKHE.Eval}(\mathcal{C}, (\overline{ct}_1, \dots, \overline{ct}_\ell), \{pk_{id}\}_{id \in T}).$$

Given a circuit \mathcal{C} , a tuple of multi-key ciphertexts $(\overline{ct}_1, \dots, \overline{ct}_\ell)$ and the corresponding set of public keys $\{pk_{id}\}_{id \in T}$, outputs a ciphertext \overline{ct} . Its reference set is the union $T = T_1 \cup \dots \cup T_\ell$ of reference sets T_j of the input ciphertexts \overline{ct}_j for $1 \leq j \leq \ell$.

Semantic Security. For any two messages $\mu_0, \mu_1 \in \mathcal{M}$, the distributions $\{\text{MKHE.Enc}(\mu_i; pk)\}$ for $i = 0, 1$ should be computationally indistinguishable where $pp \leftarrow \text{MKHE.Setup}(1^\lambda)$ and $(sk, pk) \leftarrow \text{MKHE.KeyGen}(pp)$.

Correctness and Compactness. An MKHE scheme is compact if the size of a ciphertext relevant to k parties is bounded by $\text{poly}(\lambda, k)$ for a fixed polynomial $\text{poly}(\cdot, \cdot)$.

For $1 \leq j \leq \ell$, let \overline{ct}_j be a ciphertext (with reference set T_j) such that $\text{MKHE.Dec}(\overline{ct}_j, \{sk_{id}\}_{id \in T_j}) = \mu_j$. Let $\mathcal{C} : \mathcal{M}^\ell \rightarrow \mathcal{M}$ be a circuit and $\overline{ct} \leftarrow \text{MKHE.Eval}(\mathcal{C}, (\overline{ct}_1, \dots, \overline{ct}_\ell), \{pk_{id}\}_{id \in T})$ for $T = T_1 \cup \dots \cup T_\ell$. Then,

$$\text{MKHE.Dec}(\overline{ct}, \{sk_{id}\}_{id \in T}) = \mathcal{C}(\mu_1, \dots, \mu_\ell) \quad (1)$$

with an overwhelming probability. The equality of (1) can be substituted by approximate equality similar to the CKKS scheme for approximate arithmetic [16].

2.3 Ring Learning with Errors

Throughout the paper, we assume that n is a power-of-two integer and $R = \mathbb{Z}[X]/(X^n + 1)$. We write $R_q = R/(q \cdot R)$ for the residue ring of R modulo an integer q . The Ring Learning with Errors (RLWE) assumption is that given any polynomial number of samples $(a_i, b_i = a_i \cdot s + e_i) \in R_q^2$, where a_i, s are uniformly random in R_q and e_i is drawn from an error distribution, the b_i 's are computationally indistinguishable from uniformly random elements from R_q . In our implementation, we sample the secret s uniformly at random from the set of binary polynomials.

2.4 Gadget Decomposition

Let $\mathbf{g} = (g_i) \in \mathbb{Z}^d$ be a gadget vector and q an integer. The gadget decomposition, denoted by \mathbf{g}^{-1} , is a function from R_q to R^d which transforms an element $a \in R_q$ into a vector $\mathbf{u} = (u_0, \dots, u_{d-1}) \in R^d$ of *small* polynomials such that $a = \sum_{i=0}^{d-1} g_i \cdot u_i \pmod{q}$.

The gadget decomposition technique is widely used in the construction of HE schemes. For example, homomorphic evaluation of a nonlinear circuit is based on the key-switching technique and most of HE schemes exploit various gadget decomposition method to control the noise growth. There have been suggested in the literature various decomposition methods such as bit decomposition [6, 7], base decomposition [21, 17] and RNS-based decomposition [4, 30]. Our implementation exploits an RNS-friendly decomposition for the efficiency.

3 Relinearizing Multi-key Ciphertexts

This section provides a high-level description of our MKHE schemes and explain how to perform the relinearization procedures which are core operations in homomorphic arithmetic.

3.1 Overview of HEs with Packed Ciphertexts

In recent years, there have been remarkable advances in the performance of HE schemes. For example, the ciphertext packing technique allows us to encrypt multiple data in a single ciphertext and perform parallel homomorphic operations in a SIMD manner. Currently the batch HE schemes such as BGV [7], BFV [6, 22] and CKKS [16] are the best-performing schemes in terms of amortized size and timing per plaintext slot. They adapt some DFT-like algorithms to transform a vector of plaintext values into an element of cyclotomic ring.

Let $\mathbf{sk} = (1, s)$ for the secret $s \in R$. A canonical RLWE-based ciphertext is of the form $\mathbf{ct} = (c_0, c_1) \in R_q^2$ such that the inner product $\mu = \langle \mathbf{ct}, \mathbf{sk} \rangle \pmod{q}$, called the *phase*, is a randomized encoding of a plaintext m . For example, the phase of a BFV ciphertext has the form of $\mu = (q/t) \cdot m + e$ for the plaintext modulus t while the phase $\mu = m + e$ of CKKS is an approximate value of the plaintext.

For homomorphic computation, we basically perform arithmetic operations between the phases of given ciphertexts. In particular, homomorphic multiplication of RLWE ciphertexts consists of two steps: tensor product and relinearization. For input ciphertexts \mathbf{ct}_1 and \mathbf{ct}_2 , we first compute their tensor product and return the extended ciphertext $\mathbf{ct} = \mathbf{ct}_1 \otimes \mathbf{ct}_2$ that satisfies $\langle \mathbf{ct}, \mathbf{sk} \otimes \mathbf{sk} \rangle = \langle \mathbf{ct}_1, \mathbf{sk} \rangle \cdot \langle \mathbf{ct}_2, \mathbf{sk} \rangle$. Since $\mathbf{sk} \otimes \mathbf{sk}$ contains the nonlinear entry s^2 , it requires to perform the relinearization procedure which transforms the extended ciphertext to a canonical ciphertext encrypting the same message. Roughly speaking, we publish a relinearization key which is some kind of ciphertext encrypting s^2 under \mathbf{sk} and run the key-switching algorithm for this conversion.

In the multi-key case, a ciphertext related to k different parties is of the form $\overline{\mathbf{ct}} = (c_0, c_1, \dots, c_k) \in R_q^{k+1}$ which is decryptable by the concatenated secret $\overline{\mathbf{sk}} = (1, s_1, \dots, s_k)$, i.e., its phase is computed by $\mu = \langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \rangle = c_0 + \sum_{i=1}^k c_i \cdot s_i$. If we follow the same pipeline for homomorphic operation as in the single-key setting, the tensor product step returns an extended ciphertext corresponding $\overline{\mathbf{sk}} \otimes \overline{\mathbf{sk}}$. Hence, we need to generate a relinearization key which consists of multiple ciphertexts encrypting the entries $s_i \cdot s_j$ of $\overline{\mathbf{sk}} \otimes \overline{\mathbf{sk}}$. Different from the classical HE schemes, it requires some additional computations since the term $s_i \cdot s_j$ depends on two secret keys which are independently generated by different parties. In the following, we will explain how to efficiently generate a relinearization key for multi-key homomorphic multiplication.

3.2 Basic Scheme

In this section, we present a ring-based scheme which will be used to generate some public material for relinearization.

- **Setup**(1^λ): For a given security parameter λ , set the RLWE dimension n , ciphertext modulus q , key distribution χ and error distribution ψ over R . Generate a random vector $\mathbf{a} \leftarrow U(R_q^d)$. Return the public parameter $pp = (n, q, \chi, \psi, \mathbf{a})$.
- **KeyGen**(pp): Sample the secret key $s \leftarrow \chi$. Sample an error vector $\mathbf{e} \leftarrow \psi^d$ and set the public key as $\mathbf{b} = -s \cdot \mathbf{a} + \mathbf{e} \pmod{q}$ in R_q^d .
- **UniEnc**($\mu; s$): For an input plaintext $\mu \in R$, generate a ciphertext $\mathbf{D} = [\mathbf{d}_0 | \mathbf{d}_1 | \mathbf{d}_2] \in R_q^{d \times 3}$ as follows:
 1. Sample $r \leftarrow \chi$.
 2. Sample $\mathbf{d}_1 \leftarrow U(R_q^d)$ and $\mathbf{e}_1 \leftarrow \psi^d$, and set $\mathbf{d}_0 = -s \cdot \mathbf{d}_1 + \mathbf{e}_1 + r \cdot \mathbf{g} \pmod{q}$.
 3. Sample $\mathbf{e}_2 \leftarrow \psi^d$ and set $\mathbf{d}_2 = r \cdot \mathbf{a} + \mathbf{e}_2 + \mu \cdot \mathbf{g} \pmod{q}$.

The public parameter pp contains a randomly generated vector $\mathbf{a} \in R_q^d$, so we are assuming the common reference string model. All parties should take the same public parameter as an input of the key-generation algorithm to support multi-key homomorphic arithmetic. We note that the same assumption was made in all the previous researches on MKHE.

The uni-encryption algorithm is a symmetric encryption which can encrypt a single ring element. An uni-encrypted ciphertext $\mathbf{D} = [\mathbf{d}_0 | \mathbf{d}_1 | \mathbf{d}_2] \leftarrow \text{UniEnc}(\mu; s)$ consists of three vectors in R_q^d so is (3/4)

times as large as an ordinary RGSW ciphertext in $R_q^{2d \times 2}$. For an uni-encrypted ciphertext \mathbf{D} , the first two columns $[\mathbf{d}_0 | \mathbf{d}_1]$ can be viewed as an encryption of r under the secret s while $[\mathbf{d}_2] - \mathbf{a}$ forms an encryption of μ under secret r .

Security. We claim that the uni-encryption scheme is IND-CPA secure under the RLWE assumption. We will show that the distribution

$$\{(\mathbf{a}, \mathbf{b}, \mathbf{D}) : pp = (n, \chi, \psi, \mathbf{a}) \leftarrow \text{Setup}(1^\lambda), (s, \mathbf{b}) \leftarrow \text{KeyGen}(pp), \mathbf{D} \leftarrow \text{UniEnc}(\mu; s)\}$$

is computationally indistinguishable from the uniform distribution over $R_q^d \times R_q^d \times R_q^{d \times 3}$ for an arbitrary $\mu \in R$.

First, we can modify \mathbf{b} and \mathbf{d}_0 so that we sample them independently from the uniform distribution over R_q^d . This step relies on the hardness of RLWE with parameter (n, χ, ψ) and secret s . Second, \mathbf{d}_2 can also be changed into the uniform distribution under the same RLWE assumption with secret r . Since the uniform distribution over $R_q^d \times R_q^d \times R_q^{d \times 3}$ is independent from the plaintext μ , the uni-encryption scheme is semantically secure.

3.3 Relinearization

We revisit the relinearization procedure on extended ciphertexts and present two solutions with different advantages. We recall that the tensor product $\overline{\mathbf{ct}} = \overline{\mathbf{ct}}_1 \otimes \overline{\mathbf{ct}}_2$ of two multi-key ciphertexts $\overline{\mathbf{ct}}_i \in R_q^{k+1}$ encrypted under the concatenated secret $\overline{\mathbf{sk}} = (1, s_1, \dots, s_k)$ can be viewed as a ciphertext corresponding to the tensor squared secret $\overline{\mathbf{sk}} \otimes \overline{\mathbf{sk}}$. Note that $\overline{\mathbf{sk}} \otimes \overline{\mathbf{sk}}$ contains some nonlinear entries $s_i \cdot s_j$ related to two different parties. Therefore, the computing server should be able to transform the extended ciphertext $\overline{\mathbf{ct}} \in R_q^{(k+1) \times (k+1)}$ into a canonical ciphertext by linearization of the non-linear entries $s_i \cdot s_j$.

Our relinearization methods require the same public material (evaluation key) that is generated by individual parties as follows:

- EvkGen(s): For a secret $s \in R$, set the evaluation key $\mathbf{D} \leftarrow \text{UniEnc}(s; s)$.

To be precise, each party i generates its own secret, public, and evaluation keys by running the algorithms $(s_i, \mathbf{b}_i) \leftarrow \text{KeyGen}(pp)$ and $\mathbf{D}_i \leftarrow \text{EvkGen}(s_i)$, then publishes the pair $(\mathbf{b}_i, \mathbf{D}_i)$. In the rest of this section, we present two relinearization algorithms and explain their pros and cons.

We make an additional *circular security* assumption since the evaluation key is an uni-encryption of secret s encrypted by itself. However, we stress that our assumption is no stronger than the same assumption in HE schemes [27, 22, 17, 16] requiring either bootstrapping or relinearization of ciphertexts.

3.3.1 First Method

This solution includes a pre-processing step which generates a shared relinearization key corresponding to the set of involved parties. A shared relinearization key consists of encryptions of $s_i \cdot s_j$ for all pairs $1 \leq i, j \leq k$. Then, we can linearize an extended ciphertext by applying a standard key-switching technique.

This approach is similar to a method proposed in previous researches [13, 50] which also generates a shared evaluation key. However, each element of our shared relinearization key is computed from the public information of at most two parties so consists of three vectors, while previous method based on the multi-key GSW scheme has $O(k)$ dimensional entries.

- Convert($\mathbf{D}_i, \mathbf{b}_j$): It takes as the input a pair of an uni-encryption $\mathbf{D}_i = [\mathbf{d}_{i,0} | \mathbf{d}_{i,1} | \mathbf{d}_{i,2}] \in R_q^{d \times 3}$ and a public key $\mathbf{b}_j \in R_q^d$ generated by (possibly different) parties i and j . Let $\mathbf{k}_{i,j,0}$ and $\mathbf{k}_{i,j,1}$ be the vectors in R_q^d such that $\mathbf{k}_{i,j,0}[\ell] = \langle \mathbf{g}^{-1}(\mathbf{b}_j[\ell]), \mathbf{d}_{i,0} \rangle$ and $\mathbf{k}_{i,j,1}[\ell] = \langle \mathbf{g}^{-1}(\mathbf{b}_j[\ell]), \mathbf{d}_{i,1} \rangle$ for $1 \leq \ell \leq d$, i.e., $[\mathbf{k}_{i,j,0} | \mathbf{k}_{i,j,1}] = \mathbf{M}_j \cdot [\mathbf{d}_{i,0} | \mathbf{d}_{i,1}]$ where $\mathbf{M}_j \in R_q^{d \times d}$ is the matrix whose ℓ -th row is $\mathbf{g}^{-1}(\mathbf{b}_j[\ell]) \in R^d$. Let $\mathbf{k}_{i,j,2} = \mathbf{d}_{i,2}$ and return the ciphertext $\mathbf{K}_{i,j} = [\mathbf{k}_{i,j,0} | \mathbf{k}_{i,j,1} | \mathbf{k}_{i,j,2}] \in R_q^{d \times 3}$.

Algorithm 1 Relinearization method 1

Input: $\overline{\text{ct}} = (c_{i,j})_{0 \leq i,j \leq k}$, $\overline{\text{rlk}} = \{\mathbf{K}_{i,j}\}_{1 \leq i,j \leq k}$.

Output: $\overline{\text{ct}}' = (c'_i)_{0 \leq i \leq k} \in R_q^{k+1}$.

- 1: $c'_0 \leftarrow c_{0,0}$
 - 2: **for** $1 \leq i \leq k$ **do**
 - 3: $c'_i \leftarrow c_{0,i} + c_{i,0} \pmod{q}$
 - 4: **end for**
 - 5: **for** $1 \leq i, j \leq k$ **do**
 - 6: $(c'_0, c'_i, c'_j) \leftarrow (c'_0, c'_i, c'_j) + \mathbf{g}^{-1}(c_{i,j}) \cdot \mathbf{K}_{i,j} \pmod{q}$
 - 7: **end for**
-

$$\left[\begin{array}{c|c} \mathbf{k}_{i,j,0} & \mathbf{k}_{i,j,1} \end{array} \right] = \left[\begin{array}{c} \mathbf{g}^{-1}(\mathbf{b}_j[1]) \\ \vdots \\ \mathbf{g}^{-1}(\mathbf{b}_j[d]) \end{array} \right] \cdot \left[\begin{array}{c|c} \mathbf{d}_{i,0} & \mathbf{d}_{i,1} \end{array} \right], \quad \left[\begin{array}{c} \mathbf{k}_{i,j,2} \end{array} \right] = \left[\begin{array}{c} \mathbf{d}_{i,2} \end{array} \right].$$

• Relin($\overline{\text{ct}}; \{(\mathbf{D}_i, \mathbf{b}_i)\}_{1 \leq i \leq k}$): Given an extended ciphertext $\overline{\text{ct}} = (c_{i,j})_{0 \leq i,j \leq k}$ and k pairs of evaluation/public keys $\{(\mathbf{D}_i, \mathbf{b}_i)\}_{1 \leq i \leq k}$, generate a ciphertext $\overline{\text{ct}}' \in R_q^{k+1}$ as follows:

1. Compute $\mathbf{K}_{i,j} \leftarrow \text{Convert}(\mathbf{D}_i, \mathbf{b}_j)$ for all $1 \leq i, j \leq k$ and set the relinearization key as $\overline{\text{rlk}} = \{\mathbf{K}_{i,j}\}_{1 \leq i,j \leq k} \in (R_q^{d \times 3})^{k^2}$.
2. Run Alg. 1 to relinearize $\overline{\text{ct}}$.

We note that the first step (generation of $\overline{\text{rlk}}$) can be pre-computed on public information $\{(\mathbf{D}_i, \mathbf{b}_i)\}_{1 \leq i \leq k}$ without taking a ciphertext as the input.

Correctness. We first claim that, if \mathbf{D}_i is a uni-encryption of $\mu_i \in R$ encrypted by the i -th party and \mathbf{b}_j is the public key of the j -th party, then the output $\mathbf{K}_{i,j} \leftarrow \text{Convert}(\mathbf{D}_i, \mathbf{b}_j)$ of the conversion algorithm is an encryption of $\mu_i s_j$ with respect to the secret $(1, s_i, s_j)$, i.e., $\mathbf{k}_{i,j,0} + s_i \cdot \mathbf{k}_{i,j,1} + s_j \cdot \mathbf{k}_{i,j,2} \approx \mu_i s_j \cdot \mathbf{g} \pmod{q}$. It is derived from the following formulas:

$$\begin{aligned} \mathbf{k}_{i,j,0} + s_i \cdot \mathbf{k}_{i,j,1} &= \mathbf{M}_j \cdot (\mathbf{d}_0 + s_i \cdot \mathbf{d}_1) \approx \mathbf{M}_j \cdot r_i \mathbf{g} = r_i \mathbf{b}_j \pmod{q}, \\ s_j \cdot \mathbf{k}_{i,j,2} &= s_j \cdot \mathbf{d}_2 \approx r_i s_j \cdot \mathbf{a} + \mu_i s_j \cdot \mathbf{g} \approx -r \cdot \mathbf{b}_j + \mu_i s_j \cdot \mathbf{g} \pmod{q}. \end{aligned}$$

Note that \mathbf{M}_j , s_j and r_i should be small to hold the approximate equalities. We estimate the size of noise in Appendix B.

We now show the correctness of our algorithm. Since the evaluation key \mathbf{D}_i of the i -th party is a uni-encryption of $\mu_i = s_i$, we obtain that $\mathbf{K}_{i,j} \cdot (1, s_i, s_j) \approx s_i s_j \cdot \mathbf{g} \pmod{q}$. From the definition of $\overline{\text{ct}}'$, we get

$$\begin{aligned} \langle \overline{\text{ct}}', \overline{\text{sk}} \rangle &= c'_0 + \sum_{i=1}^k c'_i \cdot s_i \\ &= c_{0,0} + \sum_{i=1}^k (c_{0,i} + c_{i,0}) s_i + \sum_{i,j=1}^k \mathbf{g}^{-1}(c_{i,j}) \cdot \mathbf{K}_{i,j} \cdot (1, s_i, s_j) \pmod{q} \\ &\approx c_{0,0} + \sum_{i=1}^k (c_{0,i} + c_{i,0}) s_i + \sum_{i,j=1}^k c_{i,j} \cdot s_i s_j = \langle \overline{\text{ct}}, \overline{\text{sk}} \otimes \overline{\text{sk}} \rangle \pmod{q}, \end{aligned}$$

as desired.

Algorithm 2 Relinearization method 2

Input: $\overline{\mathbf{ct}} = (c_{i,j})_{0 \leq i,j \leq k}$, $\{(\mathbf{D}_i = [\mathbf{d}_{i,0}|\mathbf{d}_{i,1}|\mathbf{d}_{i,2}], \mathbf{b}_i)\}_{1 \leq i \leq k}$.

Output: $\overline{\mathbf{ct}}' = (c'_i)_{0 \leq i \leq k} \in R_q^{k+1}$.

```
1:  $c'_0 \leftarrow c_{0,0}$ 
2: for  $1 \leq i \leq k$  do
3:    $c'_i \leftarrow c_{0,i} + c_{i,0} \pmod{q}$ 
4: end for
5: for  $1 \leq i, j \leq k$  do
6:    $c'_{i,j} \leftarrow \langle \mathbf{g}^{-1}(c_{i,j}), \mathbf{b}_j \rangle \pmod{q}$ 
7:    $(c'_0, c'_i) \leftarrow (c'_0, c'_i) + \mathbf{g}^{-1}(c'_{i,j}) \cdot [\mathbf{d}_{i,0}|\mathbf{d}_{i,1}] \pmod{q}$ 
8:    $c'_j \leftarrow c'_j + \langle \mathbf{g}^{-1}(c_{i,j}), \mathbf{d}_{i,2} \rangle \pmod{q}$ 
9: end for
```

3.3.2 Second Method

Our second solution does not generate a shared relinearization key different from the previous one. Instead, it directly linearizes each entry $c_{i,j}$ of an extended ciphertext $\overline{\mathbf{ct}} = (c_{i,j})_{0 \leq i,j \leq k}$ by multiplying it to \mathbf{b}_j and \mathbf{D}_i in a recursive way.

• **Relin**($\overline{\mathbf{ct}}; \{(\mathbf{D}_i, \mathbf{b}_i)\}_{1 \leq i \leq k}$): For given extended ciphertext $\overline{\mathbf{ct}} = (c_{i,j})_{0 \leq i,j \leq k}$ and k pairs of evaluation/public keys $\{(\mathbf{D}_i, \mathbf{b}_i)\}_{1 \leq i \leq k}$, generate a ciphertext $\overline{\mathbf{ct}}' \in R_q^{k+1}$ as described in Alg. 2.

We will analyze and compare two relinearization methods in the following section. In short, the second method has advantages in storage and noise growth while the first method could be faster if a shared evaluation key is used repeatedly to relinearize multiple ciphertexts corresponding to the same set of parties. We first show the correctness of the second method.

Correctness. At each iteration of the second for-loop in Alg. 2, we compute $c'_{i,j} = \langle \mathbf{g}^{-1}(c_{i,j}), \mathbf{b}_j \rangle$, then add $\mathbf{g}^{-1}(c'_{i,j}) \cdot [\mathbf{d}_{i,0}|\mathbf{d}_{i,1}]$ and $\langle \mathbf{g}^{-1}(c_{i,j}), \mathbf{d}_{i,2} \rangle$ to (c'_0, c'_i) and c'_j , respectively. We note that

$$\mathbf{g}^{-1}(c'_{i,j}) \cdot [\mathbf{d}_{i,0}|\mathbf{d}_{i,1}] \cdot (1, s_i) \approx r_i \cdot c'_{i,j} \pmod{q},$$

and

$$\langle \mathbf{g}^{-1}(c_{i,j}), \mathbf{d}_{i,2} \rangle \cdot s_j \approx \langle \mathbf{g}^{-1}(c_{i,j}), -r_i \cdot \mathbf{b}_j + s_i s_j \cdot \mathbf{g} \rangle = -r_i \cdot c'_{i,j} + c_{i,j} \cdot s_i s_j \pmod{q}.$$

From the definition of $\overline{\mathbf{ct}}'$, we get

$$\begin{aligned} \langle \overline{\mathbf{ct}}', \overline{\mathbf{sk}} \rangle &= c'_0 + \sum_{i=1}^k c'_i \cdot s_i \\ &= c_{0,0} + \sum_{i=1}^k (c_{0,i} + c_{i,0}) s_i + \sum_{i,j=1}^k \mathbf{g}^{-1}(c'_{i,j}) \cdot [\mathbf{d}_{i,0}|\mathbf{d}_{i,1}] \cdot (1, s_i) + \sum_{i,j=1}^k \langle \mathbf{g}^{-1}(c_{i,j}), \mathbf{d}_{i,2} \rangle \cdot s_j \pmod{q} \\ &\approx c_{0,0} + \sum_{i=1}^k (c_{0,i} + c_{i,0}) s_i + \sum_{i,j=1}^k c_{i,j} \cdot s_i s_j = \langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \otimes \overline{\mathbf{sk}} \rangle \pmod{q}, \end{aligned}$$

as desired.

3.3.3 Performance of Relinearization Algorithms

Suppose that there are k different parties involved in a multi-key computation. For relinearizing an extended ciphertext $\overline{\mathbf{ct}} = (c_{i,j})_{0 \leq i,j \leq k} \in R_q^{(k+1)^2}$, both of our relinearization methods repeat some computations on each $c_{i,j}$ to switch its corresponding secret $s_i \cdot s_j$ into $(1, s_i, s_j)$. So we will focus on a single step (i, j) of each solution to compare their performance.

In our first method, a computing party generates a shared relinearization key $\mathbf{K}_{i,j}$ and uses it to linearize an input extended ciphertext. The generation of $\mathbf{K}_{i,j}$ includes a multiplication between $d \times d$ and $d \times 2$ matrices so its complexity is $2d^2$ polynomial multiplications. However, the computation of $\mathbf{g}^{-1}(c_{i,j}) \cdot \mathbf{K}_{i,j}$ in Step 6 of Alg. 1 requires only $3d$ polynomial multiplications. Meanwhile, the second method does not have any pre-processing but a single iteration of Alg. 2 requires $4d$ polynomial multiplications. As a result, the first method can be up to $(4/3)$ times faster when one performs multiple homomorphic arithmetic on the same set (or its subset) of parties using a pre-computed shared relinearization key, however, the required storage grows quadratically on k compared to the linear memory of the second method.

The second method also has an advantage in noise management, which we will discuss below together with modulus raising technique.

3.3.4 Special Modulus Technique

Noise growth is the main factor determining the parameter size and thereby overall performance of a cryptosystem. In general, we can use a large decomposition degree d to reduce the size of a decomposed vector $\mathbf{g}^{-1}(\cdot)$ as well as key-switching error, but this naive method causes performance degradation. In addition, the benefit of this trade-off between noise growth and computational complexity gets smaller and smaller as d increases. Therefore, this method is not the best option when we should have a small noise.

The special modulus (a.k.a. modulus raising) technique proposed in [27] is one attractive solution to address this noise problem with a smaller overhead. Roughly speaking, it raises the ciphertext modulus from q to pq for an integer p called special modulus, and then computes the key-switching procedure over R_{pq} followed by modulus reduction back to q . The main advantage of this method is that a key-switching error is decreased by a factor of about p due to the modulus reduction. We apply this technique to our relinearization and encryption algorithms. In particular, a special modulus variant of relinearization requires two sequential modulus switching operations (see Appendix A for details).

We recall that for an extended ciphertext $\overline{\mathbf{ct}} \in R_q^{(k+1)^2}$, the goal of relinearization is to generate a ciphertext $\overline{\mathbf{ct}}' \in R_q^{k+1}$ such that $\langle \overline{\mathbf{ct}}', \overline{\mathbf{sk}} \rangle = \langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \otimes \overline{\mathbf{sk}} \rangle + e_{lin}$ for some error e_{lin} , which should be minimized for efficiency. We refer the reader to Appendix B which provides a noise analysis based on the variance of polynomial coefficients, but we present a concise summary in this section.

Let u be a uniform random variable over R_q . We consider its decomposition $\mathbf{g}^{-1}(u)$ and denote by V_g the average of variances of its coefficients. We respectively estimate the variance of a relinearization error from our first and second methods:

$$V_1 \approx k^2 n^2 \sigma^2 \cdot d^2 V_g^2, \quad V_2 \approx k^2 n^2 \sigma^2 \cdot d V_g.$$

In addition, the special modulus variant of the second method achieves a smaller noise whose variance is

$$V_2' = p^{-2} \cdot V_2 + \frac{1}{24}(k^2 + k)n.$$

Compared to the first method, our second solution has significant advantages in practice because we may use an efficient decomposition method with a small d while obtaining the same level of noise growth. Furthermore, its modulus raising variant obtains an even smaller error variance which is not nearly affected by the size of decomposition since V_2' is dominated by the second term (rounding error) when we introduce a special modulus p which can cancel out the term V_2 .

4 Two MKHE Schemes with Packed Ciphertexts

In this section, we present multi-key variants of the BFV [6, 22] and CKKS [16] schemes. They share the following setup and key generation phases but have different algorithms for message encoding and homomorphic operations.

- **MKHE.Setup**(1^λ): Run **Setup**(1^λ) and return the parameter pp .
- **MKHE.KeyGen**(pp): Each party i generates secret, public and evaluation keys by $(s_i, \mathbf{b}_i) \leftarrow \text{KeyGen}(pp)$ and $\mathbf{D}_i \leftarrow \text{EvkGen}(s_i)$, respectively.

Encryption, decryption and homomorphic arithmetic of our MKHE schemes are described in the next subsections. We have a common pre-processing when performing a homomorphic operation between ciphertexts. For given ciphertexts $\overline{\mathbf{ct}}_i \in R_q^{k_i+1}$, we denote $k \geq \max\{k_1, k_2\}$ the number of parties involved in either $\overline{\mathbf{ct}}_1$ or $\overline{\mathbf{ct}}_2$. We rearrange the entries of $\overline{\mathbf{ct}}_i$ and pad zeros in the empty entries to generate some ciphertexts $\overline{\mathbf{ct}}_i^*$ sharing the same secret $\overline{\mathbf{sk}} = (1, s_1, \dots, s_k)$. To be precise, a ciphertext $\overline{\mathbf{ct}}_i = (c_0, c_1, \dots, c_{k_i})$ corresponding to the tuple of parties $(id_1, \dots, id_{k_i}) \in \{1, 2, \dots, k\}^{k_i}$ is converted into the ciphertext $\overline{\mathbf{ct}}_i^* = (c_0^*, c_1^*, \dots, c_k^*) \in R_q^{k+1}$ which is defined as $c_0^* = c_0$ and

$$c_i^* = \begin{cases} c_j & \text{if } i = id_j \text{ for some } 1 \leq j \leq k_i; \\ 0 & \text{otherwise,} \end{cases}$$

for $1 \leq i \leq k$. We remark that

$$\langle \overline{\mathbf{ct}}_i, (1, s_{id_1}, \dots, s_{id_{k_i}}) \rangle = \langle \overline{\mathbf{ct}}_i^*, (1, s_1, \dots, s_k) \rangle.$$

For simplicity, we will assume that this pre-processing is always done before homomorphic arithmetic so that two input ciphertexts are related to the same set of k parties.

Security and Correctness. Our MKHE schemes inherit the semantic security of underlying HE schemes because they have exactly the same encryption algorithms as the ordinary HE schemes. BFV and CKKS both randomize the public key to generate a randomized RLWE sample and add an encoded plaintext to the first component. Hence our MKHE schemes are IND-CPA secure under the RLWE assumption of parameter (n, q, χ, ψ) . We will briefly show the correctness of our schemes in the following sections but we refer the reader to Appendix B for the rigorous proof with noise estimation.

4.1 Multi-Key BFV

The BFV scheme [6, 22] is a scale-invariant HE which supports exact computation on a discrete space with a finite characteristic. We denote by t the plaintext modulus and $\Delta = \lfloor q/t \rfloor$ be the scaling factor of the BFV scheme. The native plaintext space is the set of cyclotomic polynomials R_t , but a plaintext is decoded to a tuple of finite field elements via a ring isomorphism from R_t depending on the relation of t and n [47].

- **MK-BFV.Enc**($m; \mathbf{b}, \mathbf{a}$): This is the standard BFV encryption which takes a polynomial $m \in R_t$ as the input. Let $a = \mathbf{a}[0]$ and $b = \mathbf{b}[0]$. Sample $v \leftarrow \chi$ and $e_0, e_1 \leftarrow \psi$. Return the ciphertext $\mathbf{ct} = (c_0, c_1) \in R_q^2$ where $c_0 = v \cdot b + \Delta \cdot m + e_0 \pmod{q}$ and $c_1 = v \cdot a + e_1 \pmod{q}$.
- **MK-BFV.Dec**($\overline{\mathbf{ct}}; s_1, \dots, s_k$): Let $\overline{\mathbf{ct}} = (c_0, c_1, \dots, c_k) \in R_q^{k+1}$ be a ciphertext associated to k parties and s_1, \dots, s_k be their secret keys. Set $\overline{\mathbf{sk}} = (1, s_1, \dots, s_k)$ and compute $\lfloor (t/q) \cdot \langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \rangle \rfloor \pmod{t}$.
- **MK-BFV.Add**($\overline{\mathbf{ct}}_1, \overline{\mathbf{ct}}_2$): Given two ciphertexts $\overline{\mathbf{ct}}_i \in R_q^{k+1}$, return the ciphertext $\overline{\mathbf{ct}}' = \overline{\mathbf{ct}}_1 + \overline{\mathbf{ct}}_2 \pmod{q}$.
- **MK-BFV.Mult**($\overline{\mathbf{ct}}_1, \overline{\mathbf{ct}}_2; \{\mathbf{D}_i, \mathbf{b}_i\}_{1 \leq i \leq k}$): Given two ciphertexts $\overline{\mathbf{ct}}_i \in R_q^{k+1}$, compute $\overline{\mathbf{ct}} = \lfloor (t/q) \cdot (\overline{\mathbf{ct}}_1 \otimes \overline{\mathbf{ct}}_2) \rfloor \pmod{q} \in R_q^{(k+1)^2}$ and return the ciphertext $\overline{\mathbf{ct}}' \leftarrow \text{Relin}(\overline{\mathbf{ct}}; \{\mathbf{D}_i, \mathbf{b}_i\}_{1 \leq i \leq k})$.

The correctness of our scheme is obtained from the properties of the basic BFV and relinearization algorithm. A multi-key BFV encryption of $m \in R_t$ is a vector $\overline{\mathbf{ct}} = (c_0, c_1, \dots, c_k) \in R_q^{k+1}$ such that $\langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \rangle \approx \Delta \cdot m \pmod{q}$ for the secret $\overline{\mathbf{sk}} = (1, s_1, \dots, s_k)$. So the decryption algorithm can recover m correctly. If $\overline{\mathbf{ct}}_1$ and $\overline{\mathbf{ct}}_2$ are encryptions of m_1 and m_2 with respect to the secret $\overline{\mathbf{sk}} = (1, s_1, \dots, s_k)$, then their (scaled) tensor product $\overline{\mathbf{ct}} = \lfloor (t/q) \cdot (\overline{\mathbf{ct}}_1 \otimes \overline{\mathbf{ct}}_2) \rfloor \pmod{q}$ satisfies $\langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \otimes \overline{\mathbf{sk}} \rangle \approx \Delta \cdot m_1 m_2 \pmod{q}$ similar to the ordinary BFV scheme. The output $\overline{\mathbf{ct}}' \leftarrow \text{Relin}(\overline{\mathbf{ct}}; \text{rlk})$ holds $\langle \overline{\mathbf{ct}}', \overline{\mathbf{sk}} \rangle \approx \langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \otimes \overline{\mathbf{sk}} \rangle \approx \Delta \cdot m_1 m_2 \pmod{q}$.

4.2 Multi-Key CKKS

The CKKS scheme [16] is a leveled HE scheme with support for approximate fixed-point arithmetic. We assume $q = \prod_{i=0}^L p_i$ for some integers p_i to have a chain of ciphertext moduli $q_0 < q_1 < \dots < q_L$ for $q_\ell = \prod_{i=0}^\ell p_i$. The native plaintext is a small polynomial $m \in R$, but one can pack at most $(n/2)$ complex numbers in a single polynomial via DFT. In addition to the basic arithmetic operations, it supports the rescaling algorithm to control the magnitude of encrypted message. For homomorphic operations between ciphertexts at different levels, it requires to transform a high-level ciphertext to have the same level as the other.

- **MK-CKKS.Enc**($m; \mathbf{b}, \mathbf{a}$): Let $m \in R$ be an input plaintext and let $a = \mathbf{a}[0]$ and $b = \mathbf{b}[0]$. Sample $v \leftarrow \chi$ and $e_0, e_1 \leftarrow \psi$. Return the ciphertext $\mathbf{ct} = (c_0, c_1) \in R_q^2$ where $c_0 = v \cdot b + m + e_0 \pmod{q}$ and $c_1 = v \cdot a + e_1 \pmod{q}$.
- **MK-CKKS.Dec**($\overline{\mathbf{ct}}; s_1, \dots, s_k$): Let $\overline{\mathbf{ct}} = (c_0, c_1, \dots, c_k) \in R_{q_\ell}^{k+1}$ be a ciphertext at level ℓ associated to k parties and s_1, \dots, s_k be their secret keys. Set $\overline{\mathbf{sk}} = (1, s_1, \dots, s_k)$ and return $\langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \rangle \pmod{q_\ell}$.
- **MK-CKKS.Add**($\overline{\mathbf{ct}}_1, \overline{\mathbf{ct}}_2$): Given two ciphertexts $\overline{\mathbf{ct}}_i \in R_{q_\ell}^{k+1}$ at level ℓ , return the ciphertext $\overline{\mathbf{ct}}' = \overline{\mathbf{ct}}_1 + \overline{\mathbf{ct}}_2 \pmod{q_\ell}$.
- **MK-CKKS.Mult**($\overline{\mathbf{ct}}_1, \overline{\mathbf{ct}}_2; \{(\mathbf{D}_i, \mathbf{b}_i)\}_{1 \leq i \leq k}$): Given two ciphertexts $\overline{\mathbf{ct}}_i \in R_{q_\ell}^{k+1}$ at level ℓ , compute $\overline{\mathbf{ct}} = \overline{\mathbf{ct}}_1 \otimes \overline{\mathbf{ct}}_2 \pmod{q_\ell} \in R_{q_\ell}^{(k+1)^2}$ and return the ciphertext $\overline{\mathbf{ct}}' \leftarrow \mathbf{Relin}(\overline{\mathbf{ct}}; \{(\mathbf{D}_i, \mathbf{b}_i)\}_{1 \leq i \leq k}) \in R_{q_\ell}^{k+1}$. The relinearization algorithm is defined over modulus $q = q_L$, but we compute the same algorithm modulo q_ℓ for level- ℓ ciphertexts.
- **MK-CKKS.Rescale**($\overline{\mathbf{ct}}$): Given a ciphertext $\overline{\mathbf{ct}} = (c_0, c_1, \dots, c_k) \in R_{q_\ell}^{k+1}$ at level ℓ , compute $c'_i = \lfloor p_\ell^{-1} \cdot c_i \rfloor$ for $0 \leq i \leq k$ and return the ciphertext $\overline{\mathbf{ct}}' = (c'_0, c'_1, \dots, c'_k) \in R_{q_{\ell-1}}^{k+1}$.

A level- ℓ multi-key encryption of a plaintext m with respect to the secret $\overline{\mathbf{sk}} = (1, s_1, \dots, s_k)$ is a vector $\overline{\mathbf{ct}} = (c_0, c_1, \dots, c_k) \in R_{q_\ell}^{k+1}$ satisfying $\langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \rangle \approx m \pmod{q_\ell}$. For basic homomorphic operation, we take as input level- ℓ encryptions of m_1 and m_2 . Then, homomorphic addition (resp. multiplication) returns a ciphertext $\overline{\mathbf{ct}}'$ such that $[\langle \overline{\mathbf{ct}}', \overline{\mathbf{sk}} \rangle]_{q_\ell}$ is approximately equal to $m_1 + m_2$ (resp. $m_1 m_2$). Finally, we show that for a level- ℓ encryption $\overline{\mathbf{ct}}$ of m , the rescaling algorithm returns a ciphertext $\overline{\mathbf{ct}}'$ at level $(\ell - 1)$ encrypting $p_\ell^{-1} \cdot m$ from the equation $[\langle \overline{\mathbf{ct}}', \overline{\mathbf{sk}} \rangle]_{q_{\ell-1}} \approx p_\ell^{-1} \cdot [\langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \rangle]_{q_\ell}$.

4.3 Distributed Decryption

In the classical definition of MKHE primitive, all the secrets of the parties involved are required to decrypt a multi-key ciphertext. In practice, however, it is not reasonable to assume that there is a party holding multiple secret keys. Instead, we can ‘imagine a protocol between several key owners to jointly decrypt a ciphertext. The decryption algorithms of our schemes are (approximate) linear combinations of secrets with known coefficients, and there have been proposed some secure methods for this task. We introduce one simple solution based on the noise flooding technique, but any secure solution achieving the same functionality can be used.

The distributed decryption consists of two algorithms: partial decryption and merge. In the first phase, each party i receives the i -th entry of a ciphertext and decrypts it with a noise. We set the noise distribution ϕ which has a larger variance than the standard error distribution ψ of basic scheme. Then, we merge partially decrypted results with c_0 to recover the message.

- **MKHE.PartDec**(c_i, s_i): Given a polynomial c_i and a secret s_i , sample an error $e_i \leftarrow \phi$ and return $\mu_i = c_i \cdot s_i + e_i \pmod{q}$.
- **MK-BFV.Merge**($c_0, \{\mu_i\}_{1 \leq i \leq k}$): Compute $\mu = c_0 + \sum_{i=1}^k \mu_i \pmod{q}$ and return $m = \lfloor (t/q) \cdot \mu \rfloor$.
- **MK-CKKS.Merge**($c_0, \{\mu_i\}_{1 \leq i \leq k}$): Compute and return $\mu = c_0 + \sum_{i=1}^k \mu_i \pmod{q}$.

For a multi-key ciphertext $\overline{\mathbf{ct}} = (c_0, \dots, c_k)$, both multi-key BFV and CKKS schemes compute $\mu = c_0 + \sum_{i=1}^k \mu_i = \langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \rangle + \sum_{i=1}^k e_i \approx \langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \rangle \pmod{q}$ in the merge phase. Then, BFV extracts the plaintext by cancelling the scaling factor (q/t) .

5 Bootstrapping for two MKHE schemes

There have been several studies on the bootstrapping procedures of the standard (single-key) ring-based HE schemes [25, 32, 12, 14, 9]. Previous work had different goals and solutions depending on the underlying schemes but they are basically following the Gentry's technique [24] – homomorphic evaluation of the decryption circuit. In particular, the BFV and CKKS schemes have a very similar pipeline for bootstrapping which consists of four steps: (1) Modulus Raise, (2) Coeff to Slot, (3) Extraction and (4) Slot to Coeff. The second and last steps are specific linear transformations, which require rotation operations on encrypted vectors.

In the rest of this section, we first explain how to perform the rotation operation on multi-key ciphertexts based on the evaluation of Galois automorphisms. Then, we revisit the bootstrapping procedures for BFV and CKKS to generalize the existing solutions to our MKHE schemes.

5.1 Homomorphic Evaluation of Galois Automorphisms

The Galois group $\mathcal{Gal}(\mathbb{Q}[X]/(X^n + 1))$ of a cyclotomic field consists of the transformation $X \mapsto X^j$ for $j \in \mathbb{Z}_{2n}^*$. We recall that BFV (resp. CKKS) uses the DFT on R_t (resp. R) to pack multiple plaintext values into a single polynomial. As noted in [26], these automorphisms provide special functionalities on packed ciphertext such as rotation of plaintext slots.

The evaluation of an automorphism can be done based on the key-switching technique. In some more details, let $\tau_j : a(X) \mapsto a(X^j)$ be an element of the Galois group. Given an encryption $\overline{\mathbf{ct}} = (c_0, c_1, \dots, c_k) \in R_q^{k+1}$ of m , we denote by $\tau_j(\overline{\mathbf{ct}}) = (\tau_j(c_0), \dots, \tau_j(c_k))$ the ciphertext obtained by taking τ_j to the entries of $\overline{\mathbf{ct}}$. Then $\tau_j(\overline{\mathbf{ct}})$ is a valid encryption of $\tau_j(m)$ corresponding the secret key $\tau_j(\overline{\mathbf{sk}})$. We then perform the key-switching procedure from $\tau_j(\overline{\mathbf{sk}})$ back to $\overline{\mathbf{sk}}$, so as to generate a new ciphertext encrypting the same message under the original secret key $\overline{\mathbf{sk}}$.

In the following, we present two algorithms for the evaluation of the Galois element. The first algorithm generates an evaluation key for the Galois automorphism τ_j . The second algorithm gathers the evaluation keys of multiple parties and evaluates τ_j on a multi-key ciphertext using the multi-key-switching technique proposed in [10].

- **MKHE.GkGen**($j; s$): Generate a random vector $\mathbf{h}_1 \leftarrow U(R_q^d)$ and an error vector $\mathbf{e}' \leftarrow \psi^d$. For an RLWE secret $s \in R$, compute $\mathbf{h}_0 = -s \cdot \mathbf{h}_1 + \mathbf{e}' + \tau_j(s) \cdot \mathbf{g} \pmod{q}$. Return the Galois evaluation key as $\mathbf{gk} = [\mathbf{h}_0 | \mathbf{h}_1] \in R_q^{d \times 2}$.
- **MKHE.EvalGal**($\overline{\mathbf{ct}}; \{\mathbf{gk}_i\}_{1 \leq i \leq k}$): Let $\mathbf{gk}_i = [\mathbf{h}_{i,0} | \mathbf{h}_{i,1}]$ be the Galois evaluation key of the i -th party for $1 \leq i \leq k$. Given a ciphertext $\overline{\mathbf{ct}} = (c_0, \dots, c_k) \in R_q^{k+1}$, compute and return the ciphertext $\overline{\mathbf{ct}}' = (c'_0, \dots, c'_k)$ by

$$\begin{aligned} c'_0 &= \tau_j(c_0) + \sum_{i=1}^k \langle \mathbf{g}^{-1}(\tau_j(c_i)), \mathbf{h}_{i,0} \rangle \pmod{q}, \quad \text{and} \\ c'_i &= \langle \mathbf{g}^{-1}(\tau_j(c_i)), \mathbf{h}_{i,1} \rangle \pmod{q} \quad \text{for } 1 \leq i \leq k. \end{aligned}$$

In the context of CKKS, all the computations are carried out over modulus $q = q_\ell$ for level- ℓ ciphertext. We now show the correctness of our algorithms.

Correctness. From the definition, the output ciphertext $\overline{\mathbf{ct}}' = (c'_0, \dots, c'_k) \leftarrow \text{MKHE.EvalGal}(\overline{\mathbf{ct}}; \{\mathbf{gk}_i\}_{1 \leq i \leq k})$ holds

$$\begin{aligned}
\langle \overline{\mathbf{ct}}', \overline{\mathbf{sk}} \rangle &= c'_0 + \sum_{i=1}^k c'_i \cdot s_i \\
&= \tau_j(c_0) + \sum_{i=1}^k \langle \mathbf{g}^{-1}(\tau_j(c_i)), \mathbf{h}_{i,0} \rangle + \langle \mathbf{g}^{-1}(\tau_j(c_i)), \mathbf{h}_{i,1} \rangle \cdot s_i \pmod{q} \\
&\approx \tau_j(c_0) + \sum_{i=1}^k \langle \mathbf{g}^{-1}(\tau_j(c_i)), \tau_j(s_i) \cdot \mathbf{g} \rangle \pmod{q} \\
&= \langle \tau_j(\overline{\mathbf{ct}}), \tau_j(\overline{\mathbf{sk}}) \rangle = \tau_j(\langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \rangle) \pmod{q},
\end{aligned}$$

as desired. In other words, if the input ciphertext has the phase $\mu(X) = \langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \rangle \pmod{q}$, then the phase of the output ciphertext is approximately equal to $\tau_j(\mu(X)) = \mu(X^j)$.

Besides the rotation of plaintext slots, we can evaluate the Frobenius endomorphism $X \mapsto X^t$ on BFV ciphertexts using the same technique. In the case of CKKS, the map $X \mapsto X^{-1}$ corresponds to the complex conjugation over plaintext slots.

Any linear transformation can be represented as a linear combination of shifted plaintext vectors. We note that previous HE optimization techniques [25, 32, 9] for linear transformations can be directly applied to our MKHE schemes.

5.2 Bootstrapping for Multi-Key BFV

The authors of [12] described a bootstrapping procedure for the single-key BFV scheme, which follows the paradigm of [32], done for BGV scheme. The bootstrapping procedure in [12] takes as input a ciphertext with an arbitrary noise and outputs another ciphertext with a low noise encrypting the same plaintext. Below we present a multi-key variant of [12].

1. The previous work [12] published an encryption of the secret key by itself to raise the modulus. However, we observe that this step can be done by multiplying a constant without extra information. Suppose that the input ciphertext $\overline{\mathbf{ct}}$ encrypts a message m with a plaintext modulus t , i.e., $\langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \rangle = \frac{q}{t}m + e \pmod{q}$ for some error e . Then we perform a modulus-switching down to a divisor q' of q , resulting in $\langle \overline{\mathbf{ct}}', \overline{\mathbf{sk}} \rangle = \frac{q'}{t}m + e' \pmod{q'}$, then multiply the ciphertext with q/q' and get a ciphertext $\overline{\mathbf{ct}}''$ whose phase is $\langle \overline{\mathbf{ct}}'', \overline{\mathbf{sk}} \rangle = \frac{q}{q'} \cdot \left(\frac{q'}{t}m + e' \right) \pmod{q}$. This is a trivial (noise free) encryption of $\mu = \frac{q'}{t}m + e'$ with plaintext modulus q' and ciphertext modulus q .
2. It computes a homomorphic linear transform which produces multiple ciphertexts holding the coefficients $\mu_i \in \mathbb{Z}_t$ of μ in their plaintext slots. We note that this step can be done using the additions, scalar multiplications and multi-key rotations.
3. We homomorphically evaluate a polynomial, called lower digits removal [12], on the multi-key ciphertexts obtained in previous step. It removes the noise e' and leaves the coefficients of m in plaintext slots.
4. The final step is another linear transformation which inverts the second step and outputs an encryption of m .

As a consequence, the output ciphertext has the phase $\frac{q}{t}m + e'' \pmod{q}$ for an error which is smaller than the initial noise e .

5.3 Bootstrapping for Multi-Key CKKS

The authors of [14] presented a bootstrapping procedure for the single-key CKKS scheme and its performance was improved in the follow-up research [9]. The bootstrapping procedure of CKKS aims to refresh a low-level ciphertext and return an encryption of the (almost) same messages in a larger ciphertext modulus. We describe its multi-key version as follows.

1. The first step takes a lowest-level ciphertext $\overline{\text{ct}}$ as an input. Let $\mu = \langle \overline{\text{ct}}, \overline{\text{sk}} \rangle \pmod{q_0}$. Then $\langle \overline{\text{ct}}, \overline{\text{sk}} \rangle = q_0 \cdot I + \mu$ for a small $I \in R$, so $\overline{\text{ct}}$ can be considered as an encryption of $t = q_0 \cdot I + \mu$ in the largest ciphertext modulus q_L .
2. We apply a homomorphic linear transformation to compute one or two ciphertexts encrypting the coefficients of $t(X)$ in their plaintext slots. This step requires multi-key rotation and conjugation described in Section 5.1.
3. We evaluate a polynomial which approximates the reduction modular q_0 function. It removes the I part of t and leaves coefficients of μ in the slots.
4. Finally, we apply the inverse linear transformation of the second step to pack all the coefficients of μ back into a ciphertext.

The output ciphertext $\overline{\text{ct}}'$ encrypts the same plaintext μ in a higher level than the input ciphertext $\overline{\text{ct}}$, i.e., $\langle \overline{\text{ct}}', \overline{\text{sk}} \rangle \approx \mu \pmod{q_\ell}$ for some $0 < \ell < L$.

6 Implementation

We provide a proof-of-concept implementation to show the performance of our MKHE schemes. Our source code is developed in C++ with Microsoft SEAL version 3.2.0 [46] which includes BFV and CKKS implementations. We summarize our optimization techniques, recommended parameter sets, and some experimental results in this section. Finally, we apply the multi-key CKKS scheme to evaluate an encrypted neural network model on encrypted data and report the experimental result to classify handwritten images on the MNIST dataset [40]. All experiments are performed on a ThinkPad P1 laptop: Intel Xeon E-2176M @ 4.00 GHz single-threaded with 32 GB memory, compiled with GNU C++ 7.3.0 (-O2).

6.1 Optimization Techniques

Basic Optimizations. In the relinearization process, we first compute the tensor product of two ciphertexts which corresponds to the tensor squared secret $\overline{\text{sk}} \otimes \overline{\text{sk}}$. It has duplicated entries at (i, j) and (j, i) , so we can reduce its dimension from $(k+1)^2$ down to $\frac{1}{2}k(k+1)$. Both the size of the relinearization key and complexity of the algorithm are almost halved. Furthermore, each of the diagonal entries s_i^2 of $\overline{\text{sk}} \otimes \overline{\text{sk}}$ depends on a single party, so we can include a key-switching key for s_i^2 in the generation of an evaluation key. It increases the size of evaluation keys but reduces the complexity and noise of relinearization.

RNS and NTT. Our schemes are designed on the ring structure R_q , so we need to optimize the basic polynomial arithmetic. There is a well-known technique to use an RNS by taking a ciphertext modulus $q = \prod_{i=0}^L p_i$ which is a product of coprime integers. Based on the ring isomorphism $R_q \rightarrow \prod_{i=0}^L R_{p_i}$, $a \mapsto (a \pmod{p_i})_{0 \leq i \leq L}$, we achieve asymptotic/practical improvements in polynomial arithmetic over R_q . In particular, it has been studied how to design full-RNS variants of BFV and CKKS [4, 30, 15], which do not require any RNS conversions. In addition, each of base prime can be chosen properly so that there exists a $(2n)$ -th root of unity modulo p_i . It allows us to exploit an efficient Number Theoretic Transformation (NTT) modulo p_i . Our implementation adapts these techniques to improve the speed of polynomial arithmetic.

Gadget Decomposition. As mentioned before, the gadget decomposition has a major effect on the performance of homomorphic arithmetic. Bajard et al. [4] observed that the formula $a = \sum_i g_i \cdot [a]_{p_i} \pmod{q}$ where $g_i = \left[\left(\prod_{j \neq i} p_j \right)^{-1} \right]_{p_i} \cdot \left(\prod_{j \neq i} p_j \right)$ can be used to build an RNS-friendly decomposition $a \mapsto ([a]_{p_i})_i$ with the gadget vector $\mathbf{g} = (g_i)_i$. We adapt this decomposition method and take an advantage of an RNS-based implementation by storing ciphertexts in the RNS form.

In [4, 30], the authors further combined this method with the classical digit decomposition method to provide a more fine-grained control of the trade-off between complexity and noise growth. However, we realize that this hybrid method increases the decomposition degree (and thereby space and computational complexity) several times, and the special modulus technique described in Section 3.3.4 provides a much better trade-off. Therefore, the digit decomposition is *not* used in our implementation.

Parameter					Public Key		Evaluation Key	
ID	n	$\lceil \log q \rceil$	$\lceil \log p_i \rceil$	$\# p'_i s$	Size	Gen.	Size	Gen.
I	2^{13}	218	49–60	4	0.75 MB	3 ms	2.25 MB	8 ms
II	2^{14}	438	53–60	8	7 MB	24 ms	21 MB	59 ms
III	2^{15}	881	54–60	16	60 MB	195 ms	180 MB	470 ms

Table 1. Proposed parameter sets. $\log q$ and $\log p_i$ denote the bit lengths of the largest RLWE modulus and individual RNS primes, respectively. $\# p'_i s$ denotes the number of RNS primes. The standard deviation of fresh RLWE samples is $\sigma = 3.2$. Public keys’ and evaluations keys’ generation times and sizes are those of each party. ms = 10^{-3} sec.

6.2 Micro-benchmarks for MKHE Schemes

Table 1 illustrates the selected parameter sets used in experiments. They are default parameter sets in Microsoft SEAL which provide at least 128-bit of security level according to LWE-estimator [3] and HE security standard [2]. Generation time and size of secret keys, and execution time of encryption are the same as those in single-key BFV and CKKS. Decryption and ciphertext addition take $\frac{1}{2}(k+1)$ times longer than the ordinary HE schemes. We remark that generation time and size of public and evaluation keys $\{(\mathbf{b}_i, \mathbf{D}_i)\}_{1 \leq i \leq k}$ do not depend on the number of parties or the scheme because the generation can be executed in a synchronous way.

In our experiments, a homomorphic multiplication is always followed by a relinearization procedure. BFV requires more NTTs than CKKS overall to perform these operations, and is therefore slower, which is confirmed by the timing results.

ID	#Parties	Mult + Relin		EvalGal	
		BFV	CKKS	BFV	CKKS
I	1	20 ms	8 ms	3 ms	4 ms
	2	44 ms	22 ms	7 ms	8 ms
	4	116 ms	67 ms	14 ms	16 ms
	8	365 ms	229 ms	28 ms	31 ms
II	1	110 ms	59 ms	22 ms	24 ms
	2	257 ms	165 ms	47 ms	49 ms
	4	717 ms	521 ms	88 ms	95 ms
	8	2,350 ms	1,845 ms	176 ms	193 ms
III	1	675 ms	465 ms	170 ms	172 ms
	2	1,715 ms	1,364 ms	333 ms	359 ms
	4	5,025 ms	4,287 ms	646 ms	711 ms
	8	17,450 ms	15,159 ms	1,332 ms	1,413 ms

Table 2. Execution time that depends on the number of parties. Multiplication is always followed by a relinearization in MKHE. ms = 10^{-3} sec.

The execution times of multiplications in both MKHE schemes are asymptotically quadratic in the number of parties as discussed in Section 3.3. In practice, they are better than quadratic, as reported in Table 2. This is because both multiplication and relinearization include a notable portion of computation that is linear in the number of parties. The execution times of homomorphic evaluation of Galois automorphisms are almost linear on the number of the parties as described in Section 5.1.

For the single-party scenario in Table 2, we measured performance of our modified Microsoft SEAL [46] with a special modulus. It is infeasible to fairly compare the ordinary BFV and CKKS with their multi-key variants because the performance of a scheme can be analyzed from various perspectives: space/time

complexity, noise growth, functionality, etc. It is provided merely as a reference point, for a more portable estimation of MKHE on different processors.

6.3 Application to Oblivious Neural Network Inference

The authors of [34] proposed a novel framework to test encrypted neural networks on encrypted data in the single-key scenario. We consider the same service paradigm but in a multi-key setting: the data and trained model are encrypted under different keys.

6.3.1 Homomorphic Evaluation of CNN

We present an efficient strategy to evaluate CNN prediction model on the MNIST dataset. Each image is a 28×28 pixel array and will be labeled with 10 possible digits after an arbitrary number of hidden layers. We assume that a neural network is trained with the plaintext dataset in the clear. Table 3 describes our neural network topology which uses one convolution layer and two fully-connected (FC) layers with square activation function. The final step is to apply the softmax activation function for a purpose of probabilistic classification, so it is enough to obtain an index of maximum values of outputs in a prediction phase. Our objective is to predict a single image in an efficient way, thereby achieving a low latency. In Appendix C, we describe the detailed algorithms for encryption and evaluation.

Layer	Description
Convolution	Input image 28×28 , window size 4×4 , stride $(2, 2)$, number of output channels 5
1 st square	Squaring each of the 845 inputs
FC-1	Fully connecting with 845 inputs and 64 outputs
2 nd square	Squaring each of the 64 inputs
FC-2	Fully connecting with 64 inputs and 10 outputs

Table 3. Description of our CNN to the MNIST dataset.

The Convolutional Layer. As noted in [35], strided convolution can be decomposed into a sum of simple convolutions (i.e., the stride parameter = 1). From our choice of the parameters, each of such simple convolutions takes as inputs 14×14 images and 2×2 filters. This representation allows more SIMD parallelism, since we can pack all the inputs into a single ciphertext and perform four simple convolutions in parallel. Once this is done, we can accumulate the results across plaintext slots using rotate-and-sum operations in [31]. Moreover, we can pack multiple channels in a single ciphertext as in [35, Section VI.D], yielding in a fully-packed ciphertext of the convolution result.

The First Square Layer. This step applies the square activation function to all the encrypted output of the convolutional layer in an SIMD manner.

The FC-1 Layer. In general, an FC layer with n_i inputs and n_o outputs can be computed as a matrix-vector multiplication. Let \mathbf{W} and \mathbf{v} be the $n_o \times n_i$ weight matrix and n_i -length vector, respectively. We assume that n_i and n_o are smaller than the number of plaintext slots, and n_o is much lower than n_i in the context of FC layers. Halevi and Shoup [31] presented the diagonal encoding method which puts a square matrix in diagonal order, multiplies each of them with a rotation of the input vector, and then accumulates all the output vectors to obtain the result. Juvekar et al. [35] extended the method to multiply a vector by a rectangular matrix. If the input vector is encrypted in a single ciphertext, the total complexity is n_o homomorphic multiplications, $(n_o - 1)$ rotations of the input ciphertext of \mathbf{v} , and $\log(n_i/n_o)$ rotations for rotate-and-sum algorithm.

We extend their ideas to split the original matrix \mathbf{W} into smaller sized blocks and perform computation on the sub-matrices as shown in Fig. 2. Suppose that the vector \mathbf{v} is split into ℓ many sub-strings with the same length. For simplicity, we consider the first ℓ rows of \mathbf{W} . We first apply the diagonal method

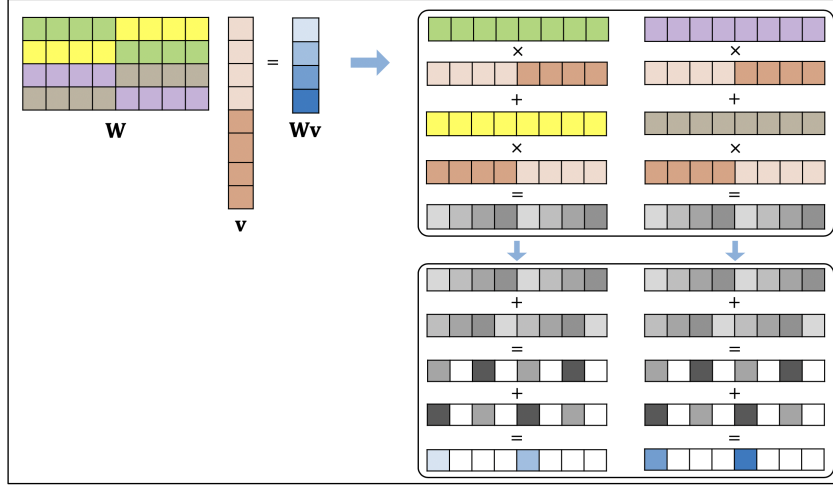


Fig. 2. Our matrix-vector multiplication algorithm ($\ell = 2$).

to arrange the $1 \times (n_i/\ell)$ sized sub-matrices of \mathbf{W} in a way that intermediate numbers are aligned in the same position across multiple slots after homomorphic multiplications. To be precise, the encryptions of diagonal components are multiplied with ℓ rotations of the encrypted vector and all these encryptions are added together similar to the diagonal method. Then, the output ciphertext represents (n_i/ℓ) -sized ℓ chunks, each containing partial sums of ℓ entries of n_i inputs. Finally, we can accumulate these using a rotate-and-sum algorithm with $\log(n_i/\ell)$ rotations. As a consequence, the output ciphertext encrypts the first ℓ many entries of $\mathbf{W}\mathbf{v}$. We repeat this procedure for each ℓ many rows of \mathbf{W} , resulting in (n_o/ℓ) ciphertexts.

When n_i is significantly smaller than the number of plaintext slots n_s , the performance can be improved by packing multiple copies of the input vector into a single ciphertext and performing (n_s/n_i) aforementioned operations in parallel. The computational cost is $(n_o \cdot n_i)/n_s$ homomorphic multiplications, $(\ell - 1)$ rotations of the input ciphertext of \mathbf{v} , and $(n_o \cdot n_i)/(n_s \cdot \ell) \cdot \log(n_i/\ell)$ rotations. We provide additional details in Appendix C.2. As a result, our method provides a trade-off between rotations on the same input ciphertext (which can benefit from the hoisting optimization of [33]) and rotations on distinct ciphertexts (which cannot benefit from hoisting).

As described in Fig. 2, all slots except the ones corresponding to the result components may reveal information about partial sums. We therefore multiply the output ciphertexts by a constant zero-one plaintext vector to remove the information.

The Second Square Layer. This step applies the square activation function to all the output nodes of the first FC layer.

The FC-2 Layer. This step performs a multiplication with small sized weight matrix \mathbf{U} and vector \mathbf{v} . As discussed in [31], it can be considered as the linear combination of \mathbf{U} 's columns using coefficients from \mathbf{v} . Suppose that the column vectors are encrypted in a single ciphertext in such a way that they are aligned with the encrypted vector. We first repeatedly rotate the encryption of the vector to generate a single ciphertext with n_o copies of each entry. Then, we apply pure SIMD multiplication to multiply each column vector by the corresponding scalar of the vector in parallel. Finally, we aggregate all the resulting columns over the slots to generate the final output.

6.3.2 Performance Evaluation

We evaluated our framework to classify encrypted handwritten images of the MNIST dataset. We used the library keras [18] with Tensorflow [1] to train the CNN model from 60,000 images of the dataset.

We employ the special modulus variant of the multi-key CKKS scheme to achieve efficiency of approximate computation. Each layer of the network has a depth of one homomorphic multiplication (except the

first FC layer requiring one more depth for multiplicative masking), so it requires 6 levels for the evaluation of CNN. We chose the parameter Set-II from Table 1 so as to cope with such levels of computations.

The data owner first chooses one among 10,000 test images in MNIST dataset, normalizes it by dividing by the maximum value 255, and encrypts it into a single ciphertext using the public key, which takes 1.75 MB of space. Meanwhile, the model provider generates a relatively large number of ciphertexts for the trained model: four for the multiple channels, eight for the weight matrix of the FC-1 layer, and one of each for the other weight or bias. Therefore, the total size of the output ciphertexts is 18.5 MB and it takes roughly 7 times longer to encrypt the trained model than an image, but it is an one-time process before data outsourcing and so it is a negligible overhead. After the evaluation, the cloud server outputs a single multi-key ciphertext encrypting the prediction result with respect to the extended secret key of the data and model owners. Table 4 shows the timing result for the evaluation of CNN. It takes about 1.8 seconds to classify an encrypted image from the encrypted training model.

Our parameter guarantees at least 32-bit precision after the decimal point. That is, the infinity norm distance between encrypted evaluation and plain computation is bounded by 2^{-32} . Therefore, we had enough space to use the noise flooding technique for decryption. In terms of the accuracy, it achieves about 98.4% on the test set which is the same as the one obtained from the evaluation in the clear.

	Stage	Runtime
Data owner	Image encryption	31 ms
Model provider	Model encryption	236 ms
Cloud server	Convolutional layer	705 ms
	1 st square layer	143 ms
	FC-1 layer	739 ms
	2 nd square layer	75 ms
	FC-2 layer	135 ms
	Total evaluation	1,797 ms

Table 4. Performance breakdown for evaluating an encrypted neural network on encrypted MNIST data, where the two encryptions are under different secret keys.

6.3.3 Comparison with Previous Works

In Table 5, we compare our benchmark result with the state-of-the-art frameworks for oblivious neural network inference: CryptoNets [29], MiniONN [41], Gazelle [35], and E2DM [34]. The first column indicates the framework and the second column denotes the cryptographic primitives used for preserving privacy. The last columns give running time for image classification as well as amortized time per instance if applicable.

Among the aforementioned solutions for private neural network prediction, E2DM relies on a third-party authority holding a secret key of HE, since the data and model are under the same secret key. CryptoNets has good amortized complexity, but it has a high latency for a single prediction. MiniONN and Gazelle have good latency, but they require both parties to be online during the protocol execution, and at least one party performs local work proportional to the complexity of the function being evaluated. Also, the number of rounds for MiniONN and Gazelle scales with the number of layers in the neural network. On the other hand, our solution has a constant number of rounds.

Moreover, our solution allows the parties to outsource homomorphic evaluation to an untrusted server (e.g. a VM in the cloud with large computing power), so both parties only need to pay encryption/decryption cost, and the communication cost only scales with the input/model sizes, but not the complexity of the network itself. This feature is made possible since our scheme supports multi-key operations. Note that the server is only assumed to be semi-honest: we do *not* require non-collusion assumptions since even if the server colludes with one party, they cannot learn the other party’s private inputs due to the IND-CPA security of MKHE. Therefore, we believe that our work presents an interesting point in the design space of oblivious machine learning inference.

Framework	Methodology	Runtime	
		Latency	Amortized
CryptoNets	HE	570 s	0.07 s
MiniONN	HE, MPC	1.28 s	-
Gazelle	HE, MPC	0.03 s	-
E2DM	HE	28.59 s	0.45 s
Ours	MKHE	1.80 s	-

Table 5. MNIST benchmarks of privacy-preserving neural network frameworks.

7 Conclusion

In this paper, we presented practical multi-key variants of the BFV and CKKS schemes and their bootstrapping methods. We provided the first experimental results of MKHE with packed ciphertexts by implementing our schemes. The main technical contribution is to propose new relinearization algorithms achieving better performance compared to prior works [13, 50]. Finally, we showed that our scheme can be applied to secure on-line prediction services by evaluating an encrypted classifier on an encrypted data under two different keys. We implemented our protocol on convolutional neural networks trained on the MNIST dataset and showed that it can achieve a low end-to-end latency by leveraging the optimized homomorphic convolutions and homomorphic matrix-vector multiplications.

References

1. M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems, 2015. <https://www.tensorflow.org>.
2. M. Albrecht, M. Chase, H. Chen, J. Ding, S. Goldwasser, S. Gorbunov, S. Halevi, J. Hoffstein, K. Laine, K. Lauter, S. Lokam, D. Micciancio, D. Moody, T. Morrison, A. Sahai, and V. Vaikuntanathan. Homomorphic encryption security standard. Technical report, HomomorphicEncryption.org, Toronto, Canada, November 2018.
3. M. R. Albrecht, R. Player, and S. Scott. On the concrete hardness of learning with errors. *Journal of Mathematical Cryptology*, 9(3):169–203, 2015.
4. J.-C. Bajard, J. Eynard, M. A. Hasan, and V. Zucca. A full RNS variant of FV like somewhat homomorphic encryption schemes. In *International Conference on Selected Areas in Cryptography*, pages 423–442. Springer, 2016.
5. M. Ben-Or, S. Goldwasser, and A. Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *Proceedings of the twentieth annual ACM symposium on Theory of computing*, pages 1–10. ACM, 1988.
6. Z. Brakerski. Fully homomorphic encryption without modulus switching from classical GapSVP. In R. Safavi-Naini and R. Canetti, editors, *CRYPTO 2012*, volume 7417 of *Lecture Notes in Computer Science*, pages 868–886. Springer, 2012.
7. Z. Brakerski, C. Gentry, and V. Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In *Proc. of ITCS*, pages 309–325. ACM, 2012.
8. Z. Brakerski and R. Perlman. Lattice-based fully dynamic multi-key FHE with short ciphertexts. In *Annual Cryptology Conference*, pages 190–213. Springer, 2016.
9. H. Chen, I. Chillotti, and Y. Song. Improved bootstrapping for approximate homomorphic encryption. Cryptology ePrint Archive, Report 2018/1043, 2018. <https://eprint.iacr.org/2018/1043>, to appear in EUROCRYPT 2019.
10. H. Chen, I. Chillotti, and Y. Song. Multi-key homomorphic encryption from TFHE. Cryptology ePrint Archive, Report 2019/116, 2019. <https://eprint.iacr.org/2019/116>.
11. H. Chen, R. Gilad-Bachrach, K. Han, Z. Huang, A. Jalali, K. Laine, and K. Lauter. Logistic regression over encrypted data from fully homomorphic encryption. *BMC medical genomics*, 11(4):81, 2018.

12. H. Chen and K. Han. Homomorphic lower digits removal and improved FHE bootstrapping. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 315–337. Springer, 2018.
13. L. Chen, Z. Zhang, and X. Wang. Batched multi-hop multi-key FHE from Ring-LWE with compact ciphertext extension. In *Theory of Cryptography Conference*, pages 597–627. Springer, 2017.
14. J. H. Cheon, K. Han, A. Kim, M. Kim, and Y. Song. Bootstrapping for approximate homomorphic encryption. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 360–384. Springer, 2018.
15. J. H. Cheon, K. Han, A. Kim, M. Kim, and Y. Song. A full RNS variant of approximate homomorphic encryption. In *International Conference on Selected Areas in Cryptography*. Springer, 2018.
16. J. H. Cheon, A. Kim, M. Kim, and Y. Song. Homomorphic encryption for arithmetic of approximate numbers. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 409–437. Springer, 2017.
17. I. Chillotti, N. Gama, M. Georgieva, and M. Izabachene. Faster fully homomorphic encryption: Bootstrapping in less than 0.1 seconds. In *Advances in Cryptology – ASIACRYPT 2016*, pages 3–33. Springer, 2016.
18. F. Chollet et al. Keras, 2015. <https://github.com/keras-team/keras>.
19. M. Clear and C. McGoldrick. Multi-identity and multi-key leveled FHE from learning with errors. In *Annual Cryptology Conference*, pages 630–656. Springer, 2015.
20. I. Damgård, M. Keller, E. Larraia, V. Pastro, P. Scholl, and N. P. Smart. Practical covertly secure MPC for dishonest majority—or: breaking the SPDZ limits. In *European Symposium on Research in Computer Security*, pages 1–18. Springer, 2013.
21. L. Ducas and D. Micciancio. FHEW: Bootstrapping homomorphic encryption in less than a second. In *Advances in Cryptology—EUROCRYPT 2015*, pages 617–640. Springer, 2015.
22. J. Fan and F. Vercauteren. Somewhat practical fully homomorphic encryption. *IACR Cryptology ePrint Archive*, 2012:144, 2012.
23. A. Gascón, P. Schoppmann, B. Balle, M. Raykova, J. Doerner, S. Zahur, and D. Evans. Privacy-preserving distributed linear regression on high-dimensional data. *Proceedings on Privacy Enhancing Technologies*, 2017(4):345–364, 2017.
24. C. Gentry. Fully homomorphic encryption using ideal lattices. In *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing*, STOC ’09, pages 169–178. ACM, 2009.
25. C. Gentry, S. Halevi, and N. P. Smart. Better bootstrapping in fully homomorphic encryption. In *Public Key Cryptography—PKC 2012*, pages 1–16. Springer, 2012.
26. C. Gentry, S. Halevi, and N. P. Smart. Fully homomorphic encryption with polylog overhead. In D. Pointcheval and T. Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *Lecture Notes in Computer Science*, pages 465–482. Springer, 2012.
27. C. Gentry, S. Halevi, and N. P. Smart. Homomorphic evaluation of the AES circuit. In R. Safavi-Naini and R. Canetti, editors, *Advances in Cryptology - CRYPTO 2012*, volume 7417 of *Lecture Notes in Computer Science*, pages 850–867. Springer, 2012.
28. C. Gentry, A. Sahai, and B. Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In *Advances in Cryptology—CRYPTO 2013*, pages 75–92. Springer, 2013.
29. R. Gilad-Bachrach, N. Dowlin, K. Laine, K. Lauter, M. Naehrig, and J. Wernsing. CryptoNets: Applying neural networks to encrypted data with high throughput and accuracy. In *International Conference on Machine Learning*, pages 201–210, 2016.
30. S. Halevi, Y. Polyakov, and V. Shoup. An improved RNS variant of the BFV homomorphic encryption scheme. *Cryptology ePrint Archive*, Report 2018/117, 2018. <https://eprint.iacr.org/2018/117>.
31. S. Halevi and V. Shoup. Algorithms in HELib. In *Advances in Cryptology—CRYPTO 2014*, pages 554–571. Springer, 2014.
32. S. Halevi and V. Shoup. Bootstrapping for HELib. In *Advances in Cryptology—EUROCRYPT 2015*, pages 641–670. Springer, 2015.
33. S. Halevi and V. Shoup. Faster homomorphic linear transformations in HELib. In *Annual International Cryptology Conference*, pages 93–120. Springer, 2018.
34. X. Jiang, M. Kim, K. Lauter, and Y. Song. Secure outsourced matrix computation and application to neural networks. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 1209–1222. ACM, 2018.
35. C. Juvekar, V. Vaikuntanathan, and A. Chandrakasan. GAZELLE: A low latency framework for secure neural network inference. In *27th USENIX Security Symposium (USENIX Security 18)*, Baltimore, MD, 2018. USENIX Association.

36. M. Keller, V. Pastro, and D. Rotaru. Overdrive: making SPDZ great again. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 158–189. Springer, 2018.
37. A. Kim, Y. Song, M. Kim, K. Lee, and J. H. Cheon. Logistic regression model training based on the approximate homomorphic encryption. *BMC medical genomics*, 11(4):83, 2018.
38. M. Kim, Y. Song, B. Li, and D. Micciancio. Semi-parallel logistic regression for GWAS on encrypted data. Cryptology ePrint Archive, Report 2019/294, 2019. <https://eprint.iacr.org/2019/294>.
39. M. Kim, Y. Song, S. Wang, Y. Xia, and X. Jiang. Secure logistic regression based on homomorphic encryption: Design and evaluation. *JMIR medical informatics*, 6(2), 2018.
40. Y. LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
41. J. Liu, M. Juuti, Y. Lu, and N. Asokan. Oblivious neural network predictions via minionn transformations. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 619–631. ACM, 2017.
42. A. López-Alt, E. Tromer, and V. Vaikuntanathan. On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 1219–1234. ACM, 2012.
43. P. Mohassel and Y. Zhang. SecureML: A system for scalable privacy-preserving machine learning. In *2017 38th IEEE Symposium on Security and Privacy (SP)*, pages 19–38. IEEE, 2017.
44. P. Mukherjee and D. Wichs. Two round multiparty computation via multi-key FHE. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 735–763. Springer, 2016.
45. C. Peikert and S. Shiehian. Multi-key FHE from LWE, revisited. In *Theory of Cryptography Conference*, pages 217–238. Springer, 2016.
46. Microsoft SEAL (release 3.2). <https://github.com/Microsoft/SEAL>, Feb. 2019. Microsoft Research, Redmond, WA.
47. N. P. Smart and F. Vercauteren. Fully homomorphic SIMD operations. *Designs, codes and cryptography*, 71(1):57–81, 2014. Early verion at <http://eprint.iacr.org/2011/133>.
48. X. Wang, S. Ranellucci, and J. Katz. Global-scale secure multiparty computation. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 39–56. ACM, 2017.
49. A. C.-C. Yao. How to generate and exchange secrets. In *Foundations of Computer Science, 1986., 27th Annual Symposium on*, pages 162–167. IEEE, 1986.
50. T. Zhou, N. Li, X. Yang, Y. Han, and W. Liu. Efficient multi-key FHE with short extended ciphertexts and less public parameters. Cryptology ePrint Archive, Report 2018/1054, 2018. <https://eprint.iacr.org/2018/1054>.

A Special Modulus Variant of Multi-Key CKKS

- **MKHE.Setup**(1^λ): Given a security parameter λ , set the RLWE dimension n , ciphertext modulus q , special modulus p , key distribution χ and error distribution ψ over R . Generate a random vector $\mathbf{a} \leftarrow U(R_{pq}^d)$. Return the public parameter $pp = (n, p, q, \chi, \psi, \mathbf{a})$.
- **UniEnc**($\mu; s$): For an input plaintext $\mu \in R$, generate a ciphertext $\mathbf{D} = [\mathbf{d}_0 | \mathbf{d}_1 | \mathbf{d}_2] \in R_{pq}^{d \times 3}$ as follows:
 1. Sample $r \leftarrow \chi$.
 2. Sample $\mathbf{d}_1 \leftarrow U(R_{pq}^d)$ and $\mathbf{e}_1 \leftarrow \psi^d$, and set $\mathbf{d}_0 = -s \cdot \mathbf{d}_1 + \mathbf{e}_1 + pr \cdot \mathbf{g} \pmod{pq}$.
 3. Sample $\mathbf{e}_2 \leftarrow \psi^d$ and set $\mathbf{d}_2 = r \cdot \mathbf{a} + \mathbf{e}_2 + p\mu \cdot \mathbf{g} \pmod{pq}$ in R_{pq}^d .
- **MKHE.KeyGen**(pp): Each party i samples the secret key $s_i \leftarrow \chi$, an error vector $\mathbf{e}_i \leftarrow \psi^d$ and sets the public key as $\mathbf{b}_i = -s_i \cdot \mathbf{a} + \mathbf{e}_i \pmod{pq}$. Set the evaluation key $\mathbf{D}_i \leftarrow \text{UniEnc}(s_i; s_i)$.
- **Relin**($\overline{\mathbf{ct}}; \{(\mathbf{D}_i, \mathbf{b}_i)\}_{1 \leq i \leq k}$): Given an extended ciphertext $\overline{\mathbf{ct}} = (c_{i,j})_{0 \leq i, j \leq k} \in R_q^{(k+1)^2}$ and k pairs of evaluation and public keys $\{(\mathbf{D}_i, \mathbf{b}_i)\}_{1 \leq i \leq k} \in (R_{pq}^{d \times 3} \times R_{pq}^d)^k$, generate a ciphertext $\overline{\mathbf{ct}}' \in R_q^{k+1}$ as described in Alg. 3.
- **MK-CKKS.Enc**($m; \mathbf{b}_i, \mathbf{a}$): Let $m \in R$ be an input plaintext and let $a = \mathbf{a}[0]$ and $b_i = \mathbf{b}_i[0]$ be the first entries of the common reference string and public key of the i -th party. Sample $v \leftarrow \chi$ and $e_0, e_1 \leftarrow \psi$. Return the ciphertext $\mathbf{ct} = (m, 0) + \lfloor p^{-1} \cdot (c_0, c_1) \rfloor \in R_q^2$ where $(c_0, c_1) = v \cdot (b, a) + (e_0, e_1) \pmod{pq}$.

Algorithm 3 Relinearization method with modulus raising

Input: $\overline{\mathbf{ct}} = (c_{i,j})_{0 \leq i,j \leq k}$, $\{(\mathbf{D}_i = [\mathbf{d}_{i,0} | \mathbf{d}_{i,1} | \mathbf{d}_{i,2}], \mathbf{b}_i)\}_{1 \leq i \leq k}$.

Output: $\overline{\mathbf{ct}}' = (c'_i)_{0 \leq i \leq k} \in R_q^{k+1}$.

```
1:  $(c''_i)_{0 \leq i \leq k} \leftarrow \mathbf{0}$  ▷ Create a temporary vector modulo  $pq$ 
2: for  $1 \leq i, j \leq k$  do
3:    $c''_{i,j} \leftarrow \langle \mathbf{g}^{-1}(c_{i,j}), \mathbf{b}_j \rangle \pmod{pq}$ 
4:    $c'_{i,j} \leftarrow \lfloor p^{-1} \cdot c''_{i,j} \rfloor$  ▷ It is a polynomial modulo  $q$ 
5:    $(c'_0, c'_i) \leftarrow (c'_0, c'_i) + \mathbf{g}^{-1}(c'_{i,j}) \cdot [\mathbf{d}_{i,0} | \mathbf{d}_{i,1}] \pmod{pq}$ 
6:    $c'_j \leftarrow c'_j + \langle \mathbf{g}^{-1}(c_{i,j}), \mathbf{d}_{i,2} \rangle \pmod{pq}$ 
7: end for
8:  $c'_0 \leftarrow c_{0,0} + \lfloor p^{-1} \cdot c''_0 \rfloor \pmod{q}$ 
9: for  $1 \leq i \leq k$  do
10:   $c'_i \leftarrow c_{0,i} + c_{i,0} + \lfloor p^{-1} \cdot c''_i \rfloor \pmod{q}$ 
11: end for
```

- MK-CKKS.Dec, Add, Mult, Rescale: These algorithms are the same as the ones described in Section 4.2.
- MKHE.GkGen($j; s$): Generate a random vector $\mathbf{h}_1 \leftarrow U(R_{pq}^d)$ and an error vector $\mathbf{e}' \leftarrow \psi^d$. Compute $\mathbf{h}_0 = -s \cdot \mathbf{h}_1 + \mathbf{e}' + \tau_j(s) \cdot \mathbf{g} \pmod{pq}$. Return the Galois evaluation key as $\mathbf{gk} = [\mathbf{h}_0 | \mathbf{h}_1] \in R_{pq}^{d \times 2}$.
- MKHE.EvalGal($\overline{\mathbf{ct}}; \{\mathbf{gk}_i\}_{1 \leq i \leq k}$): Let $\mathbf{gk}_i = [\mathbf{h}_{i,0} | \mathbf{h}_{i,1}]$ be the Galois evaluation key of the i -th party for $1 \leq i \leq k$. Given a ciphertext $\overline{\mathbf{ct}} = (c_0, \dots, c_k) \in R_q^{k+1}$, compute

$$\begin{aligned} c''_0 &= \sum_{i=1}^k \langle \mathbf{g}^{-1}(\tau_j(c_i)), \mathbf{h}_{i,0} \rangle \pmod{pq} \quad \text{and} \\ c''_i &= \langle \mathbf{g}^{-1}(\tau_j(c_i)), \mathbf{h}_{i,1} \rangle \pmod{pq} \quad \text{for } 1 \leq i \leq k, \end{aligned}$$

then return the ciphertext $\overline{\mathbf{ct}}' = (c'_0, \dots, c'_k) \in R_q^{k+1}$ where $c'_0 = \tau_j(c_0) + \lfloor p^{-1} \cdot c''_0 \rfloor \pmod{q}$ and $c'_i = \lfloor p^{-1} \cdot c''_i \rfloor \pmod{q}$ for $1 \leq i \leq k$.

B Noise Analysis

We provide an average-case noise estimation on the variances of polynomial coefficients. In the following, we make a heuristic assumption that the coefficients of each polynomial behave like independent zero-mean random variables of the same variance. We denote by $\text{Var}(a) = \text{Var}(a_i)$ the variance of its coefficients for given random variable $a = \sum_i a_i \cdot X^i \in R$ over R . Hence, the product $c = a \cdot b$ of two polynomials will have the variance of $\text{Var}(c) = n \cdot \text{Var}(a) \cdot \text{Var}(b)$. More generally, for a vector $\mathbf{a} \in R^d$ of random variables, we define $\text{Var}(\mathbf{a}) = \frac{1}{d} \sum_{i=1}^d \text{Var}(\mathbf{a}[i])$.

Specifically, we let $V_g = \text{Var}(\mathbf{g}^{-1}(a))$ of a uniform random variable a over R_q to estimate the size of gadget decomposition. Recall that our implementation exploits the RNS-friendly decomposition $R_q \rightarrow \prod_i R_{p_i}$, $a \mapsto ([a]_{p_i})_i$ for distinct word-size primes of the same bit-size so that $d = \lceil \log q / \log p_i \rceil$ and $V_g \approx \frac{1}{12d} \sum_{i=1}^d p_i^2$. Finally, we assume that every ciphertext $\overline{\mathbf{ct}} \in R_q^{k+1}$ behaves as if it is a uniform random variable over R_q^{k+1} .

B.1 Relinearization

We first specify some distributions for detailed analysis. We set the key distribution χ and the distribution ψ as the uniform distribution over the set of binary polynomials and the Gaussian distribution of variance σ^2 , respectively.

Method 1. We first analyze the conversion algorithm $\mathbf{K}_{i,j} \leftarrow \text{Convert}(\mathbf{D}_i, \mathbf{b}_j)$. Let $\mathbf{D}_i = [\mathbf{d}_{i,0} | \mathbf{d}_{i,1} | \mathbf{d}_{i,2}]$ be an uni-encryption of $\mu_i \in R$ encrypted by the secret s_i and (s_j, \mathbf{b}_j) a pair of the secret and public keys of the j -th party, i.e., $\mathbf{b}_j = -s_j \cdot \mathbf{a} + \mathbf{e}_j \pmod{q}$, $\mathbf{d}_{i,0} = -s_i \cdot \mathbf{d}_{i,1} + \mathbf{e}_{i,1} + r_i \cdot \mathbf{g} \pmod{q}$, and $\mathbf{d}_{i,2} = r_i \cdot \mathbf{a} + \mathbf{e}_{i,2} + \mu_i \cdot \mathbf{g} \pmod{q}$ for fresh errors $\mathbf{e}_j, \mathbf{e}_{i,1}$ and $\mathbf{e}_{i,2}$. We observe that

$$\begin{aligned} \mathbf{k}_{i,j,0} + s_i \cdot \mathbf{k}_{i,j,1} &= \mathbf{M}_j \cdot (\mathbf{d}_{i,0} + s_i \cdot \mathbf{d}_{i,1}) = \mathbf{M}_j \mathbf{e}_{i,1} + r_i \mathbf{b}_j \pmod{q}, \\ s_j \cdot \mathbf{k}_{i,j,2} &= r_i s_j \cdot \mathbf{a} + s_j \cdot \mathbf{e}_{i,2} + \mu_i s_j \cdot \mathbf{g} \pmod{q}, \end{aligned}$$

and consequently $\mathbf{k}_{i,j,0} + s_i \cdot \mathbf{k}_{i,j,1} + s_j \cdot \mathbf{k}_{i,j,2} = (\mathbf{M}_j \mathbf{e}_{i,1} + r_i \mathbf{e}_j + s_j \mathbf{e}_{i,2}) + \mu_i s_j \mathbf{g} \pmod{q}$. Therefore, the noise $\mathbf{M}_j \mathbf{e}_{i,1} + r_i \mathbf{e}_j + s_j \mathbf{e}_{i,2} \in R^d$ of the output ciphertext has the variance of

$$V_{conv} = n\sigma^2 \cdot (d \cdot V_g + 1).$$

We now consider an extended ciphertext $\overline{\mathbf{ct}} = (c_{i,j})_{0 \leq i,j \leq k}$ and the output $\overline{\mathbf{ct}}' \leftarrow \text{Relin}(\overline{\mathbf{ct}}; \{(\mathbf{D}_i, \mathbf{b}_i)\}_{1 \leq i \leq k})$ of the relinearization procedure. As shown in Section 3.3.1, it satisfies that

$$\begin{aligned} \langle \overline{\mathbf{ct}}', \overline{\mathbf{sk}} \rangle &= c'_0 + \sum_{i=1}^k c'_i \cdot s_i \\ &= c_{0,0} + \sum_{i=1}^k (c_{0,i} + c_{i,0}) s_i + \sum_{i,j=1}^k \mathbf{g}^{-1}(c_{i,j}) \cdot \mathbf{K}_{i,j} \cdot (1, s_i, s_j) \pmod{q} \\ &= \langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \otimes \overline{\mathbf{sk}} \rangle + \sum_{i,j=1}^k \langle \mathbf{g}^{-1}(c_{i,j}), \mathbf{e}_{i,j} \rangle \pmod{q}, \end{aligned}$$

where $\mathbf{e}_{i,j} = \mathbf{K}_{i,j} \cdot (1, s_i, s_j) - s_i s_j \cdot \mathbf{g} \pmod{q}$ denotes the error of $\mathbf{K}_{i,j}$. Hence, the relinearization error $e_{lin} = \sum_{i,j=1}^k \langle \mathbf{g}^{-1}(c_{i,j}), \mathbf{e}_{i,j} \rangle$ has the variance of

$$V_{lin} = k^2 \cdot nd \cdot V_g \cdot V_{conv} \approx k^2 \cdot n^2 d^2 \sigma^2 \cdot V_g^2.$$

This variance can be reduced to about half by eliminating the duplicated entries of $\overline{\mathbf{sk}} \otimes \overline{\mathbf{sk}}$ as explained in Section 6.1. We also note that the factor k^2 in the formula can be reduced down to $k_1 k_2$ if $\overline{\mathbf{ct}} = \overline{\mathbf{ct}}_1 \otimes \overline{\mathbf{ct}}_2$ is the tensor product of two sparse ciphertexts $\overline{\mathbf{ct}}_i$ corresponding to $k_i \leq k$ secrets.

Method 2. Suppose that $\overline{\mathbf{ct}} = (c_{i,j})_{0 \leq i,j \leq k}$ is an extended ciphertext and let $\overline{\mathbf{ct}}' \leftarrow \text{Relin}(\overline{\mathbf{ct}}; \{(\mathbf{D}_i, \mathbf{b}_i)\}_{1 \leq i \leq k})$ be the output of the relinearization procedure. As noted in Section 3.3.2, we have that

$$\mathbf{g}^{-1}(c'_{i,j}) \cdot [\mathbf{d}_{i,0} | \mathbf{d}_{i,1}] \cdot (1, s_i) = r_i \cdot c'_{i,j} + \langle \mathbf{g}^{-1}(c'_{i,j}), \mathbf{e}_{i,1} \rangle \pmod{q},$$

and

$$\begin{aligned} \langle \mathbf{g}^{-1}(c_{i,j}), \mathbf{d}_{i,2} \rangle \cdot s_j &= \langle \mathbf{g}^{-1}(c_{i,j}), r_i s_j \cdot \mathbf{a} + s_j \cdot \mathbf{e}_{i,2} + s_i s_j \cdot \mathbf{g} \rangle \\ &= \langle \mathbf{g}^{-1}(c_{i,j}), r_i \cdot (-\mathbf{b}_j + \mathbf{e}_j) + s_j \cdot \mathbf{e}_{i,2} + s_i s_j \cdot \mathbf{g} \rangle \\ &= -r_i \cdot c'_{i,j} + c_{i,j} \cdot s_i s_j + e_{i,j} \pmod{q}. \end{aligned}$$

where $e_{i,j} = \langle \mathbf{g}^{-1}(c_{i,j}), r_i \cdot \mathbf{e}_j + s_j \cdot \mathbf{e}_{i,2} \rangle \pmod{q}$. The variance of $e_{i,j}$ is $n^2 d \cdot \sigma^2 \cdot V_g$.

From the equation

$$\langle \overline{\mathbf{ct}}', \overline{\mathbf{sk}} \rangle = \langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \otimes \overline{\mathbf{sk}} \rangle + \sum_{i,j=1}^k (\langle \mathbf{g}^{-1}(c'_{i,j}), \mathbf{e}_{i,1} \rangle + e_{i,j}) \pmod{q},$$

the variance of a relinearization error e_{lin} is obtained by

$$V_{lin} = k^2 \cdot (n^2 + n) d \cdot \sigma^2 \cdot V_g \approx k^2 \cdot n^2 d \cdot \sigma^2 \cdot V_g.$$

Special Modulus Variant of Method 2. In lines 3 ~ 6 of Alg. 3, we add $\mathbf{g}^{-1}(c'_{i,j}) \cdot [\mathbf{d}_{i,0} | \mathbf{d}_{i,1}]$ and $\langle \mathbf{g}^{-1}(c_{i,j}), \mathbf{d}_{i,2} \rangle$ to the temporary ciphertext. We first note that $p \cdot c'_{i,j} = c''_{i,j} - [c''_{i,j}]_p \pmod{pq}$ and

$$\begin{aligned} \mathbf{g}^{-1}(c'_{i,j}) \cdot (\mathbf{d}_{i,0} + s_i \cdot \mathbf{d}_{i,1}) &= pr_i c'_{i,j} + \langle \mathbf{g}^{-1}(c'_{i,j}), \mathbf{e}_{i,1} \rangle \pmod{pq} \\ &= \langle \mathbf{g}^{-1}(c_{i,j}), r_i \mathbf{b}_j \rangle - r_i \cdot [c''_{i,j}]_p + \langle \mathbf{g}^{-1}(c'_{i,j}), \mathbf{e}_{i,1} \rangle \pmod{pq}. \end{aligned}$$

Meanwhile, the other term satisfies that

$$\begin{aligned} \langle \mathbf{g}^{-1}(c_{i,j}), s_j \cdot \mathbf{d}_{i,2} \rangle &= \langle \mathbf{g}^{-1}(c_{i,j}), -r_i \mathbf{b}_j + r_i \mathbf{e}_j + ps_i s_j \mathbf{g} + s_j \mathbf{e}_{i,2} \rangle \pmod{pq} \\ &= pc_{i,j} \cdot s_i s_j - \langle \mathbf{g}^{-1}(c'_{i,j}), r_i \mathbf{b}_j \rangle + \langle \mathbf{g}^{-1}(c_{i,j}), r_i \mathbf{e}_j + s_j \mathbf{e}_{i,2} \rangle \pmod{pq}. \end{aligned}$$

Consequently, the phase of the temporary ciphertext $(c''_{i,j})_{0 \leq i \leq k}$ is increased by $pc_{i,j} \cdot s_i s_j + e_{i,j}$ for the error

$$e_{i,j} = -r_i \cdot [c''_{i,j}]_p + \langle \mathbf{g}^{-1}(c'_{i,j}), \mathbf{e}_{i,1} \rangle + \langle \mathbf{g}^{-1}(c_{i,j}), r_i \mathbf{e}_j + s_j \mathbf{e}_{i,2} \rangle,$$

whose variance is $\text{Var}(e_{i,j}) = (n^2 + n)d \cdot V_g \cdot \sigma^2 + \frac{1}{24}np^2$. We repeat it for all $1 \leq i, j \leq k$, so the temporary ciphertext will satisfy

$$c''_0 + \sum_{i=1}^k c''_i \cdot s_i = e + p \cdot \sum_{i,j=1}^k c_{i,j} \cdot s_i s_j \pmod{pq}$$

for the error $e = \sum_{i,j=1}^k e_{i,j}$ of variance $\text{Var}(e) = k^2 \cdot \text{Var}(e_{i,j})$. After reducing its modulus down to q , we get the output ciphertext $\overline{\mathbf{ct}'} = (c'_i)_{0 \leq i \leq k}$ whose phase is

$$\langle \overline{\mathbf{ct}'}, \overline{\mathbf{sk}} \rangle = \langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \otimes \overline{\mathbf{sk}} \rangle + p^{-1} \cdot e + e_{rd}$$

with an additional rounding error $e_{rd} = -p^{-1} \cdot \langle [(c'_i)_{0 \leq i \leq k}]_p, \overline{\mathbf{sk}} \rangle$ of variance $\frac{1}{24}kn + \frac{1}{12}$. The final relinearization error $e_{lin} = p^{-1} \cdot e + e_{rd}$ has the variance

$$\begin{aligned} V_{lin} &= p^{-2} \cdot k^2(n^2 + n)d\sigma^2 \cdot V_g + \frac{1}{24}k^2n + \frac{1}{24}kn + \frac{1}{12} \\ &\approx p^{-2} \cdot k^2n^2d \cdot \sigma^2 \cdot V_g + \frac{1}{24}(k^2 + k)n. \end{aligned}$$

B.2 Multi-Key BFV

Encryption. Let $\mathbf{ct} = (c_0, c_1) \in R_q^2$ be an encryption of $m \in R_t$ generated by the randomness $r \leftarrow \chi$ and $e_0, e_1 \leftarrow \psi$. Then we have

$$\begin{aligned} c_0 + c_1 \cdot s &= \Delta \cdot m + r \cdot (b + a \cdot s) + (e_0 + e_1 \cdot s) \pmod{q} \\ &= \Delta \cdot m + (r \cdot e + e_0 + e_1 \cdot s) \pmod{q}, \end{aligned}$$

where $e = \mathbf{e}[0]$ is a noise of public key. Therefore, the encryption noise $e_{enc} = r \cdot e + e_0 + e_1 \cdot s$ has the variance of

$$V_{enc} = \sigma^2 \cdot (1 + n) \approx \sigma^2 n.$$

Multiplication. Suppose that $\overline{\mathbf{ct}}_i$ is an encryption of m_i for $i = 1, 2$, i.e., $\langle \overline{\mathbf{ct}}_i, \overline{\mathbf{sk}} \rangle = q \cdot I_i + \Delta \cdot m_i + e_i$ for some I_i and e_i in R . The variance of $I_i = \left\lfloor \frac{1}{q} \langle \overline{\mathbf{ct}}_i, \overline{\mathbf{sk}} \rangle \right\rfloor$ is computed by $\text{Var}(I_i) \approx \frac{1}{12} (1 + \frac{1}{2}kn) \approx \frac{1}{24}kn$ since $\frac{1}{q} \cdot \overline{\mathbf{ct}}_i$ behaves like a uniform random variable over $\frac{1}{q} \cdot R_q^{k+1}$ whose variance is approximately equal to $\frac{1}{12}$.

The tensor product of the input ciphertexts satisfies

$$\begin{aligned}\langle \overline{\mathbf{ct}}_1 \otimes \overline{\mathbf{ct}}_2, \overline{\mathbf{sk}} \otimes \overline{\mathbf{sk}} \rangle &= \langle \overline{\mathbf{ct}}_1, \overline{\mathbf{sk}} \rangle \cdot \langle \overline{\mathbf{ct}}_2, \overline{\mathbf{sk}} \rangle \\ &= \Delta^2 \cdot m_1 m_2 + q \cdot (I_1 e_2 + I_2 e_1) + \Delta \cdot (m_1 e_2 + m_2 e_1) + e_1 e_2 \pmod{q \cdot \Delta},\end{aligned}$$

and consequently the ciphertext $\overline{\mathbf{ct}} = \lfloor (t/q) \cdot \overline{\mathbf{ct}}_1 \otimes \overline{\mathbf{ct}}_2 \rfloor$ has the phase

$$\langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \otimes \overline{\mathbf{sk}} \rangle = \Delta \cdot m_1 m_2 + (t \cdot (I_1 e_2 + I_2 e_1) + (m_1 e_2 + m_2 e_1) + \Delta^{-1} \cdot e_1 e_2 + e_{rd})$$

for the rounding error $e_{rd} = \langle (t/q) \cdot \overline{\mathbf{ct}}_1 \otimes \overline{\mathbf{ct}}_2 - \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \otimes \overline{\mathbf{sk}} \rangle$. Therefore, the multiplication error

$$e_{mul} = (t \cdot (I_1 e_2 + I_2 e_1) + (m_1 e_2 + m_2 e_1) + \Delta^{-1} \cdot e_1 e_2 + e_{rd})$$

is dominated by the first term $t \cdot (I_1 e_2 + I_2 e_1)$ whose variance is

$$V_{mul} = nt^2 \cdot (\text{Var}(I_1) \cdot \text{Var}(e_2) + \text{Var}(I_2) \cdot \text{Var}(e_1)) \approx \frac{1}{24} kn^2 t^2 \cdot (\text{Var}(e_1) + \text{Var}(e_2)).$$

B.3 Multi-Key CKKS

Encryption. The CKKS scheme has the same encryption error as BFV.

Multiplication. For $i = 1, 2$, let $\mu_i = \langle \overline{\mathbf{ct}}_i, \overline{\mathbf{sk}} \rangle$ be the phase of an input ciphertext $\overline{\mathbf{ct}}_i$. Then $\overline{\mathbf{ct}} = \overline{\mathbf{ct}}'_1 \otimes \overline{\mathbf{ct}}'_2$ has the phase $\langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \otimes \overline{\mathbf{sk}} \rangle = \langle \overline{\mathbf{ct}}'_1, \overline{\mathbf{sk}} \rangle \cdot \langle \overline{\mathbf{ct}}'_2, \overline{\mathbf{sk}} \rangle = \mu_1 \cdot \mu_2 \pmod{q}$. Therefore, the output $\overline{\mathbf{ct}}' \leftarrow \text{Relin}(\overline{\mathbf{ct}}; \overline{\mathbf{rk}})$ satisfies that $\langle \overline{\mathbf{ct}}', \overline{\mathbf{sk}} \rangle = \mu_1 \mu_2 + e_{lin}$ for a relinearization error e_{lin} of variance V_{lin} . As explained above, this variance can be reduced down if $\overline{\mathbf{ct}}_i$ has some zero entries.

Rescaling. Let $\overline{\mathbf{ct}}' \leftarrow \text{MK-CKKS.Rescale}(\overline{\mathbf{ct}}) = \lfloor p_\ell^{-1} \cdot \overline{\mathbf{ct}} \rfloor \in R_{q_{\ell-1}}^{k+1}$ for $\overline{\mathbf{ct}} \in R_{q_\ell}^{k+1}$. Then, we have

$$\langle \overline{\mathbf{ct}}', \overline{\mathbf{sk}} \rangle = p_\ell^{-1} \cdot \langle \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \rangle_{q_\ell} + e_{rd} \pmod{q_{\ell-1}}$$

for the rescaling error $e_{rd} = \langle \lfloor p_\ell^{-1} \cdot \overline{\mathbf{ct}} \rfloor - p_\ell^{-1} \cdot \overline{\mathbf{ct}}, \overline{\mathbf{sk}} \rangle$ from rounding. Note that each component of $\lfloor p_\ell^{-1} \cdot \overline{\mathbf{ct}} \rfloor - p_\ell^{-1} \cdot \overline{\mathbf{ct}}$ behaves like a uniform random variable on $p_\ell^{-1} \cdot R_{p_\ell}$ whose variance is $\frac{1}{12}$. Therefore, the rescaling error has the variance of

$$V_{res} = \frac{1}{12} \left(1 + \frac{1}{2} kn \right) \approx \frac{1}{24} kn.$$

C Homomorphic Evaluation of CNN

For a $d_1 \times d$ matrix \mathbf{A}_1 and a $d_2 \times d$ matrix \mathbf{A}_2 , $(\mathbf{A}_1; \mathbf{A}_2)$ denotes the $(d_1 + d_2) \times d$ matrix obtained by concatenating two matrices in a vertical direction. If two matrices \mathbf{A}_1 and \mathbf{A}_2 have the same number of rows, $(\mathbf{A}_1 | \mathbf{A}_2)$ denotes a matrix formed by horizontal concatenation. As defined in [34], there is a row ordering encoding map to transform a vector of dimension $n = d^2$ into a matrix in $\mathbb{R}^{d \times d}$. For a vector $\mathbf{a} = (a_k)_{0 \leq k < n}$, we define the encoding map $\text{mat} : \mathbb{R}^n \rightarrow \mathbb{R}^{d \times d}$ by $\text{mat} : \mathbf{a} \mapsto \mathbf{A} = (a_{d \cdot i + j})_{0 \leq i, j < d}$, i.e., \mathbf{a} is the concatenation of row vectors of \mathbf{A} . Let vec denote its inverse mapping. We use $\mathbf{A}_{i, \ell_1 : \ell_2}$ to denote a vector of length $(\ell_2 - \ell_1)$ formed by extracting from the ℓ_1 -th entry to the $(\ell_2 - 1)$ -th entry of the i -th row of \mathbf{A} . $\text{PadZeros}(\mathbf{A}, v, \text{dir})$ pads a matrix \mathbf{A} with zeros in the direction specified by dir where v is a vector of non-negative integers that specifies both the amount of padding to add and the dimension along which to add it. For example, direction can be specified as one of the following values: L (left), R (right), U (upper), and B (bottom). For simplicity, when \mathbf{A} is given as a vector and we want to pad zeros to the right side of the vector, we will omit the direction in the notation.

Algorithm 4 Encryption of an image

Input: Image $\mathbf{X} \in \mathbb{R}^{28 \times 28}$, public key of the image provider \mathbf{b}_D .

Output: $\text{ct}.X$.

```
1: for  $0 \leq i, j < 2$  do
2:    $\mathbf{X}'_{i,j} \leftarrow (x_{2k+i, 2l+j})_{0 \leq k, l \leq 13} \in \mathbb{R}^{14 \times 14}$ 
3:    $\bar{\mathbf{X}}_{i,j} \leftarrow (\mathbf{X}'_{i,j}; \dots; \mathbf{X}'_{i,j}) \in \mathbb{R}^{70 \times 14}$ 
4:    $\mathbf{v}_{i,j} \leftarrow \text{PadZeros}(\text{vec}(\bar{\mathbf{X}}_{i,j}), 2^{10} - 70 \cdot 14) \in \mathbb{R}^{1024}$ 
5: end for
6:  $\mathbf{v} \leftarrow (\mathbf{v}_{0,0} | \mathbf{v}_{0,1} | \mathbf{v}_{1,0} | \mathbf{v}_{1,1}) \in \mathbb{R}^{4096}$ 
7:  $\bar{\mathbf{v}} \leftarrow (\mathbf{v} | \mathbf{v}) \in \mathbb{R}^{8192}$ 
8:  $\text{ct}.X \leftarrow \text{MK-CKKS.Enc}(\bar{\mathbf{v}}; \mathbf{b}_D, \mathbf{a})$ 
```

Algorithm 5 Encryption of multiple channels

Input: $\{\mathbf{Y}^{(k)} = (y_{i,j}^{(k)}) \in \mathbb{R}^{4 \times 4}\}_{0 \leq k < 5}$, public key of the model provider \mathbf{b}_M .

Output: $\{\text{ct}.Y_l\}_{0 \leq l < 4}$.

```
1: for  $0 \leq i, j < 4$  do
2:    $\mathbf{v}_{i,j} \leftarrow \emptyset$   $\triangleright$  Null string
3:   for  $0 \leq k < 5$  do
4:      $\mathbf{Y}_{k,i,j} \leftarrow (y_{i,j}^{(k)}, \dots, y_{i,j}^{(k)}) \in \mathbb{R}^{13 \times 13}$ 
5:      $\mathbf{Y}'_{k,i,j} \leftarrow \text{PadZeros}(\mathbf{Y}_{k,i,j}, (1, 1), (R, B)) \in \mathbb{R}^{14 \times 14}$ 
6:      $\mathbf{v}_{i,j} \leftarrow (\mathbf{v}_{i,j} | \text{vec}(\mathbf{Y}'_{k,i,j}))$ 
7:   end for
8:    $\mathbf{v}_{i,j} \leftarrow \text{PadZeros}(\mathbf{v}_{i,j}, 2^{10} - 5 \cdot 14 \cdot 14) \in \mathbb{R}^{1024}$ 
9: end for
10: for  $0 \leq i, j < 2$  do
11:    $\mathbf{v}_{2i+j} \leftarrow (\mathbf{v}_{2i,2j} | \mathbf{v}_{2i,2j+1} | \mathbf{v}_{2i+1,2j} | \mathbf{v}_{2i+1,2j+1}) \in \mathbb{R}^{4096}$ 
12:    $\bar{\mathbf{v}}_{2i+j} \leftarrow (\mathbf{v}_{2i+j} | \mathbf{v}_{2i+j}) \in \mathbb{R}^{8192}$ 
13:    $\text{ct}.Y_{2i+j} \leftarrow \text{MK-CKKS.Enc}(\bar{\mathbf{v}}_{2i+j}; \mathbf{b}_M, \mathbf{a})$ 
14: end for
```

C.1 Encryption of Data and Trained Model

Alg. 4 takes as the input an image $\mathbf{X} = (x_{i,j}) \in \mathbb{R}^{28 \times 28}$ and the public key \mathbf{b}_D of the image provider. Alg. 5 takes the multiple channels of the convolutional layer from the trained model as inputs and generates their encryptions using the public key of the model provider \mathbf{b}_M . Alg. 6 takes as the inputs the weights and biases of the FC layers and outputs their encryptions using the public key.

C.2 Homomorphic Evaluation of CNN

We start with a useful aggregation operation across plaintext slots from the literature [31, 38]. This algorithm is referred as **AllSum**, which is parameterized by integers φ and α . See Algorithm 7 for an implementation. We denote by $\text{MKHE.Rot}(\text{ct}; k)$ a multi-key rotation which transforms an input ciphertext ct into a new ciphertext that encrypts a shifted plaintext vector by k of the original message. Given a ciphertext ct representing a plaintext vector $\mathbf{m} = (m_0, \dots, m_{\ell-1}) \in \mathbb{R}^\ell$ for $\ell = \alpha \cdot \varphi$, the **AllSum** algorithm outputs a ciphertext ct_{out} encrypting α copies of the vector $(\sum_{j=0}^{\alpha-1} m_{j \cdot \varphi}, \sum_{j=0}^{\alpha-1} m_{j \cdot \varphi + 1}, \dots, \sum_{j=0}^{\alpha-1} m_{(j+1) \cdot \varphi - 1}) \in \mathbb{R}^\varphi$. This can be implemented using rotations and additions.

Matrix-Vector Multiplication. Assume that n_i is smaller than the number of plaintext slots n_s and the input $n_o \times n_i$ matrix \mathbf{W} is split into $1 \times (n_i/\ell)$ sized sub-matrices, denoted by $\mathbf{w}_{i,j}$ for $0 \leq i < n_o$ and $0 \leq j < \ell$. Then, we can pack $(\ell \cdot n_s)/n_i$ many different sub-matrices into a single ciphertext, and $n_c = (n_s/n_i)$ copies of the input vector \mathbf{v} in a single ciphertext. For example, consider the first $(\ell \cdot n_c)$

Algorithm 6 Encryption of weights & biases of FC layers

Input: $\mathbf{W} \in \mathbb{R}^{64 \times 845}$, $\mathbf{z}^{(1)} \in \mathbb{R}^{64}$, $\mathbf{U} \in \mathbb{R}^{10 \times 64}$, $\mathbf{z}^{(2)} \in \mathbb{R}^{10}$, public key of the model provider \mathbf{b}_M .

Output: $\{\text{ct}.W_k\}_{0 \leq k < 8}$, $\text{ct}.z_1$, $\text{ct}.U$, $\text{ct}.z_2$.

[Encryption of weight of FC-1]

```
1:  $\mathbf{W}' \leftarrow \emptyset$ 
2: for  $0 \leq i < 64$  do
3:    $\mathbf{w} \leftarrow \emptyset$ 
4:   for  $0 \leq j < 5$  do
5:      $\mathbf{W}_j \leftarrow \text{mat}(\mathbf{W}_{i,13^2 \cdot j : 13^2 \cdot (j+1)}) \in \mathbb{R}^{13 \times 13}$ 
6:      $\mathbf{W}'_j \leftarrow \text{PadZeros}(\mathbf{W}_j, (1, 1), (R, B)) \in \mathbb{R}^{14 \times 14}$ 
7:      $\mathbf{w} \leftarrow (\mathbf{w} \parallel \text{vec}(\mathbf{W}'_j))$ 
8:   end for
9:    $\mathbf{w} \leftarrow \text{PadZeros}(\mathbf{w}, 2^{10} - 5 \cdot 14 \cdot 14) \in \mathbb{R}^{1024}$ 
10:   $\mathbf{W}' \leftarrow (\mathbf{W}'; \mathbf{w})$ 
11: end for
12: for  $0 \leq k < 8$  do
13:    $\mathbf{v} \leftarrow \emptyset$ 
14:   for  $0 \leq i < 64$  do
15:      $\mathbf{v} \leftarrow (\mathbf{v} \parallel \mathbf{W}'_{i, 2^7 \cdot ((i+k)\%8) : 2^7 \cdot ((i+k+1)\%8)})$ 
16:   end for
17:    $\text{ct}.W_k \leftarrow \text{MK-CKKS.Enc}(\mathbf{v}; \mathbf{b}_M, \mathbf{a})$ 
18: end for
```

[Encryptions of bias of FC-1/weight & bias of FC-2]

```
19:  $\mathbf{v}_1 \leftarrow \emptyset$ ,  $\mathbf{u} \leftarrow \emptyset$ ,  $\mathbf{v}_2 \leftarrow \emptyset$ 
20: for  $0 \leq i < 64$  do
21:    $\mathbf{v}_1 \leftarrow (\mathbf{v}_1 \parallel \text{PadZeros}(z_i^{(1)}, 2^7 - 1))$ 
22:    $\mathbf{u} \leftarrow (\mathbf{u} \parallel \text{PadZeros}(\mathbf{U}_i^t, 2^7 - 10))$ 
23:    $\mathbf{v}_2 \leftarrow (\mathbf{v}_2 \parallel \text{PadZeros}(\mathbf{z}^{(2)}, 2^7 - 10))$ 
24: end for
25:  $\text{ct}.z_1 \leftarrow \text{MK-CKKS.Enc}(\mathbf{v}_1; \mathbf{b}_M, \mathbf{a})$ 
26:  $\text{ct}.U \leftarrow \text{MK-CKKS.Enc}(\mathbf{u}; \mathbf{b}_M, \mathbf{a})$ 
27:  $\text{ct}.z_2 \leftarrow \text{MK-CKKS.Enc}(\mathbf{v}_2; \mathbf{b}_M, \mathbf{a})$ 
```

$\triangleright z_i^{(1)} : i\text{-th entry of } \mathbf{z}^{(1)}$
 $\triangleright \mathbf{U}_i : i\text{-th column of } \mathbf{U}$

rows of \mathbf{W} . We encode n_c many the first diagonal vectors of the matrix into a plaintext vector, i.e.,

$$\begin{aligned} & ((\mathbf{w}_{0,0} \parallel \mathbf{w}_{1,1} \parallel \dots \parallel \mathbf{w}_{\ell-1,\ell-1}) \parallel (\mathbf{w}_{\ell,0} \parallel \mathbf{w}_{\ell+1,1} \parallel \dots \parallel \mathbf{w}_{2\ell-1,\ell-1}) \parallel \dots \\ & (\mathbf{w}_{\ell \cdot (n_c-1),0} \parallel \mathbf{w}_{\ell \cdot (n_c-1)+1,1} \parallel \dots \parallel \mathbf{w}_{\ell \cdot (n_c-1)+\ell-1,\ell-1})) \in \mathbb{R}^{n_s}. \end{aligned}$$

As such, each extended diagonal vectors are encrypted in a single ciphertext and these ℓ many ciphertexts are multiplied with ℓ rotations of the encrypted vector \mathbf{v} . Then we add them together similar to the original diagonal method and the output ciphertext represents (n_i/ℓ) -sized $(\ell \cdot n_c)$ chunks, each of which contains partial sums of ℓ entries of n_i inputs. Finally, we can accumulate these using a rotate-and-sum algorithm with $\log(n_i/\ell)$ rotations, yielding in a ciphertext that contains the first $(\ell \cdot n_c)$ entries of $\mathbf{W}\mathbf{v} \in \mathbb{R}^{n_o}$. We repeat these procedures $n_o/(\ell \cdot n_c)$ times. In the end, we get $n_o/(\ell \cdot n_c)$ ciphertexts, each containing $(\ell \cdot n_c)$ entries of the result. The computational cost is $\ell \cdot (n_o/(\ell \cdot n_c)) = (n_o \cdot n_i)/n_s$ homomorphic multiplications, $(\ell - 1)$ rotations of the encrypted vector \mathbf{v} , and $(n_o/(\ell \cdot n_c)) \cdot \log(n_i/\ell)$ rotation on distinct ciphertexts.

Evaluation Strategy of CNN. Alg. 8 provides an explicit description of homomorphic evaluation of CNN. We denote by $\text{SMult}(\text{ct}, \mathbf{u})$ the multiplication of ct with a scalar \mathbf{u} . For simplicity, we omit the evaluation/public keys $\{(\mathbf{D}_i, \mathbf{b}_i)\}_{1 \leq i \leq 2}$ in homomorphic multiplication process.

Algorithm 7 AllSum(ct, ψ , α)

Input: ct, input ciphertext, the unit initial amount by which the ciphertext shifts ψ , the number of the repeated doubling $\alpha > 1$

Output: ct_{out}

```
1: ctout  $\leftarrow$  ct
2: for  $i = 0, 1, \dots, \log \alpha - 1$  do
3:   ctout  $\leftarrow$  MKHE.Add(ctout, MKHE.Rot(ctout;  $\psi \cdot 2^i$ ))
4: end for
```

Algorithm 8 Homomorphic evaluation of CNN

Input: ct.X, {ct.Y_l}_{0 ≤ l < 4}, {ct.W_k}_{0 ≤ k < 8}, ct.z₁, ct.U, ct.z₂.

Output: ct_{out}.

[Convolutional layer]

```
1: ct0  $\leftarrow$  MK-CKKS.Mult(ct.X, ct.Y0) ▷ Simple convolutions
2: for  $1 \leq l < 4$  do
3:   ct.Xl  $\leftarrow$  MK-CKKS.Rot(ct.X;  $14i + j$ ) ▷  $i = \lfloor l/2 \rfloor, j = (l \% 2)$ 
4:   ct  $\leftarrow$  MK-CKKS.Mult(ct.Xl, ct.Yl)
5:   ct0  $\leftarrow$  MK-CKKS.Add(ct0, ct)
6: end for
7: ct0  $\leftarrow$  MK-CKKS.Rescale(ct0)
8: ct0  $\leftarrow$  AllSum(ct0, 1024, 4) ▷ Sum over multiple channels
```

[1st square layer]

```
9: ct1  $\leftarrow$  MK-CKKS.Rescale(MK-CKKS.Mult(ct0, ct0))
```

[FC-1 layer]

```
10: ct2  $\leftarrow$  MK-CKKS.Mult(ct1, ct.W0)
11: for  $1 \leq l < 8$  do
12:   ct  $\leftarrow$  MK-CKKS.Rot(ct1;  $2^7 \cdot l$ )
13:   ct  $\leftarrow$  MK-CKKS.Mult(ct, ct.Wl)
14:   ct2  $\leftarrow$  MK-CKKS.Add(ct2, ct)
15: end for
16: ct2  $\leftarrow$  MK-CKKS.Rescale(ct2)
17: ct2  $\leftarrow$  AllSum(ct2, 1, 64)
18: u  $\leftarrow$  PadZeros((1), 127)  $\in \mathbb{R}^{128}$ 
19: u  $\leftarrow$  (u|u|...|u)  $\in \mathbb{R}^{8192}$ 
20: ct2  $\leftarrow$  MK-CKKS.SMult(ct2, u) ▷ Multiplicative masking
21: ct2  $\leftarrow$  MK-CKKS.Rescale(ct2)
22: ct2  $\leftarrow$  MK-CKKS.Add(ct2, ct.z1)
```

[2nd square layer]

```
23: ct3  $\leftarrow$  MK-CKKS.Rescale(MK-CKKS.Mult(ct2, ct2))
```

[FC-2 layer]

```
24: ct4  $\leftarrow$  AllSum(ct3, -1, 16)
25: ct4  $\leftarrow$  MK-CKKS.Rescale(MK-CKKS.Mult(ct4, ct.U))
26: ct4  $\leftarrow$  AllSum(ct4, 128, 64)
27: ctout  $\leftarrow$  MK-CKKS.Add(ct4, ct.z2)
```
