# Hercules: Boosting the Performance of Privacy-Preserving Federated Learning

Guowen Xu<sup>®</sup>, Xingshuo Han<sup>®</sup>, Shengmin Xu<sup>®</sup>, Tianwei Zhang<sup>®</sup>, Hongwei Li<sup>®</sup>, *Senior Member, IEEE*, Xinyi Huang<sup>®</sup>, and Robert H. Deng<sup>®</sup>, *Fellow, IEEE* 

**Abstract**—In this paper, we address the problem of privacy-preserving federated neural network training with *N* users. We present **Hercules**, an efficient and high-precision training framework that can tolerate collusion of up to N - 1 users. **Hercules** follows the POSEIDON framework proposed by Sav et al. (NDSS'21), but makes a qualitative leap in performance with the following contributions: (i) we design a novel parallel homomorphic computation method for matrix operations, which enables fast Single Instruction and Multiple Data (SIMD) operations over ciphertexts. For the multiplication of two  $h \times h$  dimensional matrices, our method reduces the computation complexity from  $O(h^3)$  to O(h). This greatly improves the training efficiency of the neural network since the ciphertext computation is dominated by the convolution operations; (ii) we present an efficient approximation on the sign function based on the composite polynomial approximation. It is used to approximate non-polynomial functions (i.e., ReLU and max), with the optimal asymptotic complexity. Extensive experiments on various benchmark datasets (BCW, ESR, CREDIT, MNIST, SVHN, CIFAR-10 and CIFAR-100) show that compared with POSEIDON, **Hercules** obtains up to 4% increase in model accuracy, and up to  $60 \times$  reduction in the computation and communication cost.

Index Terms—Privacy protection, federated learning, polynomial approximation

# **1** INTRODUCTION

As a promising neural network training mechanism, Federated Learning (FL) has been highly sought after with some attractive features including amortized overhead and mitigation of privacy threats. However, the conventional FL setup has some inherent privacy issues [1], [2]. Consider a scenario where a company (referred to as the cloud server) pays multiple users and requires them to train a target neural network model collaboratively. Although each user is only required to upload the intermediate data (e.g., gradients) instead of the original training data to the server during the training process, a large amount of sensitive information can still be leaked implicitly from these

Manuscript received 18 June 2022; revised 14 October 2022; accepted 28 October 2022. Date of publication 2 November 2022; date of current version 1 September 2023.

(Corresponding author: Xingshuo Han.) Digital Object Identifier no. 10.1109/TDSC.2022.3218793 intermediate values. Previous works have demonstrated many powerful attacks to achieve this, such as attribute inference attacks and gradient reconstruction attacks [3], [4], [5]. On the other hand, the target model is locally distributed to each user according to the FL protocol, which ignores the model privacy and may be impractical in realworld scenarios. Actually, to protect the model privacy, the server must keep users ignorant of the details of the model parameters throughout the training process.

# 1.1 Related Works

Extensive works have been proposed to mitigate the above privacy threats. In general, existing privacy-preserving deep learning solutions mainly rely on the following two lines of technologies: *Differential Privacy* (DP) [6], [7] and *crypto-based multiparty secure computing* (MPC) [8], [9], [10], [11], [12]. Each one has merits and demerits depending on the scenario to which it is applied.

Differential Privacy. DP is usually applied in the training phase [6], [7]. To ensure the indistinguishability between individual samples while maintaining high training accuracy, each user is required to add noise to the gradient or local parameters that meets the preset privacy budget. Abadi et al. [6] propose the first differentially private stochastic gradient descent (SGD) algorithm. They carefully implement gradient clipping, hyperparameter tuning, and moment accountant to obtain a tight estimate of overall privacy loss, both asymptotically and empirically. Yu et al. [7] design a new DP-SGD, which employs a new primitive called zero concentrated differential privacy (zCDP) for privacy accounting, to achieve a rigorous estimation of the privacy loss. In recent years, many variants of the above works have been designed and applied to specific scenarios [13], [14], [15], [16]. Most of them follow the principle that the

Guowen Xu, Xingshuo Han, and Tianwei Zhang are with the School of Computer Science and Engineering, Nanyang Technological University, Singapore 639798. E-mail: guowen.xu@foxmail.com, xingshuo001@e. ntu.edu.sg, tianwei.zhang@ntu.edu.sg.

Shengmin Xu and Xinyi Huang are with the College of Computer and Cyber Security, Fujian Normal University, Fuzhou 350007, China. E-mail: {smxu1989, xyhuang81]@gmail.com.

<sup>•</sup> Hongwei Li is with the School of Computer Science and Engineering, University of Electronic Science and Technology of China, Chengdu 611731, China. E-mail: hongweili@uestc.edu.cn.

Robert H. Deng is with the School of Information Systems, Singapore Management University, Singapore 178902. E-mail: robertdeng@smu.edu.sg.

This work was supported in part by NTU-Desay Research Program under Grant 2018-0980, in part by the Singapore Ministry of Education (MOE) AcRF Tier under Grant 2MOE-T2EP20121-0006, in part by the National Natural Science Foundation of China under Grants 62102090 and 62032005, and in part by the Science Foundation of Fujian Provincial Science and Technology Agency under Grant 2020J02016.

<sup>1545-5971 © 2022</sup> IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

minimum accumulated noise is added to the gradient or local parameters while meeting the preset privacy budget.

DP is cost-effective because each user is only required to add noise that obeys a specific distribution during training. However, it is forced to make a trade-off between training accuracy and privacy, i.e., a strong privacy protection level can be reached at the cost of certain model accuracy drop [17], [18]. This goes against the motivation of this paper, as our goal is to design a highly secure FL training framework without compromising the model accuracy.

Crypto-Based Multiparty Secure Computing. The implementation of this strategy mainly relies on two general techniques, secret sharing [19] and homomorphic encryption (HE) [11]. MPC enables the calculation of arbitrary functions collaboratively by multiple parties without revealing the secret input of each party. To support privacy-preserving neural network training, most existing works [8], [9], [10], [19], [20], [21] rely on splitting the training task into two or more servers, who are usually assumed to be non-colluding. Then, state-of-the-art secret sharing methods, including arithmetic sharing [19], boolean sharing [8], and Yao's garbled circuit [22] are carefully integrated to efficiently implement various mathematical operations under the ciphertext. Mohassel et al. [20] propose SecureML, the first privacy-preserving machine learning framework for generalized linear model regression and neural network training. It lands on the setting of two non-colluding servers, where users securely outsource local data to them. Then, several types of secret sharing methods are mixed and used to complete complex ciphertext operations. Other works, e.g., ABY<sup>3</sup> [8], QUOTIENT [9], BLAZE [23], Trident [24], are also exclusively based on the MPC protocol between multiple noncolluding servers (or a minority of malicious servers) to achieve fast model training and prediction.

It is cost-effective to outsource the training task among multiple users to several non-colluding servers, avoiding the high communication overhead across large-scale users. However, it may be impractical in real scenarios where the setting of multiple servers is not available. Especially in FL scenarios, users are more inclined to keep their datasets locally rather than uploading data to untrusted servers. To alleviate this problem, several works [2], [11], [12], [25] propose to use multi-party homomorphic encryption (a.k.a. threshold homomorphic encryption, as a variant of the standard HE), as the underlying technology to support direct interactions among multiple data owners for distributed learning. For example, Zheng et al. [11] present Helen, a secure distributed learning approach for linear models, where the threshold Paillier scheme [26] is used to protect users' local data. Froelicher et al. [25] reduce the computation overhead of Helen by using the packed plaintext encoding with the SIMD technology [2]. Sav et al. propose POSEIDON [12], the first distributed training framework with multi-party homomorphic encryption. It relies on the multiparty version of the CKKS (MCKKS) cryptosystem [27] to encrypt users' local data. Compared with the standard CKKS, the secret key of MCKKS is securely shared with multiple entities. As a result, each entity still performs the function evaluation under the same public key. However, the decryption of the result requires the participation of all entities. Besides, non-polynomial functions are approximated

# 1.2 Technical Challenges

In this paper, we follow the specifications of POSEIDON to design our FL training framework, because such a technical architecture enables the users' data to be kept locally without incurring additional servers. However, there are still several critical issues that have not been solved well. (1) Computation overhead is the main obstacle hindering the development of HE. It usually requires more computing resources to perform the same machine learning tasks compared to outsourcingbased solutions [8], [9], [10]. Although there are some optimization methods such as parameter quantization and model compression [9], [28], they inevitably degrade the model accuracy. Recently, Zhang et al. [29] design GALA, which employs a novel coding technique for matrix-vector multiplication. In this way, multiple plaintexts are packed into one ciphertext to perform efficient homomorphic SIMD operations without reducing the calculation accuracy. However, GALA is specifically designed for the MPC protocol that uses a mixture of HE and garbled circuits, and its effectiveness is highly dependent on the assistance of the inherent secret sharing strategy. Therefore, it is necessary to design a computation optimization method that is completely suitable for HE, without sacrificing the calculation accuracy. (2) There is a lack of satisfactory approximation mechanisms for non-polynomial functions in HE. HE basically supports homomorphic addition and multiplication. For nonpolynomial functions, especially ReLU, one of the most popular activation functions in hidden layers, we need to approximate them to polynomials for ciphertext evaluation. The common polynomial approximation method, such as the minimax method, aims to find the approximate polynomial with the smallest degree on the objective function under the condition of a given error bound. However, the computation complexity of evaluating these polynomials is enormous, making it quite inefficient to obtain the fitting function with high-precision [30], [31]. Recently, Lu et al. [32] propose PEGASUS, which can efficiently switch back and forth between a packed CKKS ciphertext and FHEW ciphertext [33] without decryption, allowing us to evaluate both polynomial and non-polynomial functions on encrypted data. However, its performance is still far from practical.

# 1.3 Our Contributions

As discussed above, the HE-based FL is more in line with the needs of most real-world applications, compared to other methods. However, it suffers from computing bottlenecks and poor compatibility with non-polynomial functions. To mitigate these limitations, we present Hercules, an efficient, privacy-preserving and high-precision framework for FL. Hercules follows the tone of the state-of-the-art work POSEIDON [12], but makes a qualitative leap in performance. Specifically, we first devise a new method for parallel homomorphic computation of matrix, which supports fast homomorphic SIMD operations, including addition, multiplication, and transposition. Then, instead of fitting the replacement function of ReLU for training in POSEIDON, we design an efficient method based on the composite polynomial approximation. In short, the contributions of Hercules are summarized as follows:

• We design a new method to execute matrix operations in parallel, which can pack multiple plaintexts

as polynomial functions to be efficiently executed by CKKS. tions in parallel, which can pack multiple plaintexts Authorized licensed use limited to: Nanyang Technological University Library. Downloaded on October 06,2023 at 19:53:34 UTC from IEEE Xplore. Restrictions apply.

into a ciphertext to achieve fast homomorphic SIMD operations (Section 3). Our key insight is to minimize the number of plaintext slots that need to be rotated in matrix multiplication through customized permutations. Compared with existing works [12], [34], our solution reduces the computation complexity from  $O(h^3)$  to O(h) for the multiplication of any two  $h \times h$  matrices. It greatly improves the neural network training efficiency since the ciphertext computation is dominated by the convolution operations. We describe the detail of efficiently executing matrix transposition on packed ciphertexts, and packing multiple matrices into one ciphertext, yielding better-amortized performance.

- We present an efficient approximation on the sign function based on the composite polynomial approximation, with optimal asymptotic complexity (Section 4). The core of our solution is to carefully construct a polynomial *g* with a constant degree, and then make the composite polynomial *g* ∘ *g* ∘ *g* ∘ *s* ∘ · · · ∘ *g* infinitely close to the sign function, as the number of *g* increases. In this way, our new algorithm only requires Θ(log (1/δ)) + Θ(log σ) computation complexity to obtain an approximate sign function result of *m* ∈ [−1, −δ] ∪ [δ, 1] within 2<sup>-σ</sup> error. For example, for an encrypted 20-bit integer *m*, we can obtain the result of the sign function within 2<sup>-20</sup> error with an amortized running time of 20.05 milliseconds, which is 33× faster than the state-of-the-art work [35].
- We show that Hercules provides semantic security in the FL scenario consisting of *N* users and a parameter server, and tolerates collusion among up to *N* – 1 passive users (Section 5). This is mainly inherited from the property of the MCKKS.
- We conduct extensive experiments on various benchmark datasets (BCW, ESR, CREDIT, MNIST, SVHN, CIFAR-10 and CIFAR-100) to demonstrate the superiority of **Hercules** in terms of classification accuracy, and overhead of computation and communication (Section 6). Specifically, compared with POSEIDON, we obtain up to 4% increase in model accuracy, and up to 60× reduction in the computation and communication cost.

*Roadmap*. In Section 2, we review some basic concepts used in this paper, and introduce the scenarios and threat models. In Sections 3, 4, and 5, we give the details of **Hercu-les**. Performance evaluation is presented in Section 6. Section 7 concludes the paper.

# 2 PRELIMINARIES

# 2.1 Neural Network Training

A neural network usually consists of an input layer, one or more hidden layers, and an output layer, where hidden layers include convolutional layers, pooling layers, activation function layers, and fully connected layers. The connections between neurons in adjacent layers are parameterized by  $\omega$  (*i.e.*, model parameters), and each neuron is associated with an element-wise activation function  $\varphi$  (such as sigmoid, ReLU, and softmax). Given the training sample set  $(x, y) \in D$ , training a neural network of  $\mathbb{L}$  layers is generally Authorized licensed use limited to: Narvang Technological University Library. Dow divided into two phases: feedforward and backpropagation. Specifically, at the *k*-th iteration, the weights between layers *j* and *j* + 1 are denoted as a matrix  $\omega_i^k$ ; matrix  $M_j$  represents the activation of neurons in the *j*-th layer. Then the input xis sequentially propagated to each layer with operations of linear transformation (i.e,  $E_i^k = \omega_i^k \times M_{i-1}^k$ ) and non-linear transformation (i.e.,  $M_i^k = \varphi(E_i^k)$ ) to obtain the final classification result  $\bar{y} = M_{\rm I}^k$ . With the loss function L which is usually set as  $L=||y-\bar{y}||_2$ , the mini-batch based Stochastic Gradient Descent (SGD) algorithm [12] is exploited to optimize the parameter  $\omega$ . The parameter update rule is  $\omega_i^{k+1} =$  $\omega_i^k - \frac{\eta}{B} \bigtriangledown \omega_i^k$ , where  $\eta$  and  $\mathcal{B}$  indicate the learning rate and the random batch size of input samples, and  $\nabla \omega_i^k = \frac{\partial L}{\partial k}$ . Since the transposition of matrices/vectors is involved in the *backpropagation*, we use  $V^T$  to represent the transposition of variable V. The *feedforward* and *backpropagation* steps are performed iteratively until the neural network meets the given convergence constraint. The detailed implementation is shown in *Algorithm* 1.

# Algorithm 1. Mini-Batch Based SGD Algorithm

| Input: $\omega_1^k, \omega_2^k, \dots, \omega_1^k$ .   |
|--|
| Output: $\omega_1^{k+1}, \omega_2^{k+1}, \ldots, \omega_{\mathbb{L}}^{k+1}$ .                          |
| 1: for $t = 1$ to $\mathcal{B}$ do   |
| 2: $M_0 = X[t]$ $\triangleright$ feedforward   |
| 3: <b>for</b> $j = 1$ to $\mathbb{L}$ <b>do</b>  |
| 4: $E_i^k = \omega_i^k \times M_{i-1}^k$   |
| 5: $\dot{M}_{i}^{k} = \dot{\varphi}(E_{i}^{k})$  |
| 6: end for   |
| 7: $L_{L}^{k} =   y[t] - M_{L}^{k}  _{2}$ $\triangleright$ backpropagation                             |
| 8: $L_{\mathrm{L}}^{\overline{k}} = \varphi'(E_{\mathrm{L}}^{k}) \odot \overline{L}_{\mathrm{L}}^{k1}$ |
| 9: $\nabla \omega_{\mathbb{L}}^k + = (M_{\mathbb{L}-1}^k)^T \times L_{\mathbb{L}}^k$                   |
| 10: <b>for</b> $j = L - 1$ to 1 <b>do</b> _  |
| 11: $L_{i}^{k} = L_{i+1}^{k} \times (\omega_{i+1}^{k})^{T}$  |
| 12: $L_i^k = \varphi'(E_i^k) \odot L_i^k$  |
| 13: $\nabla \omega_i^k + = (M_{i-1}^k)^T \times L_i^k$   |
| 14: end for  |
| 15: end for  |
| 16: <b>for</b> $j = 1$ to L <b>do</b>  |
| 17: $\omega_i^{k+1} = \omega_i^k - \frac{\eta}{B} \bigtriangledown \omega_i^k$                         |
| 18: end for  |

# 2.2 Multiparty Version of CKKS

Hercules relies on the multiparty version of Cheon-Kim-Kim-Song (MCKKS) [12] fully homomorphic encryption to protect users' data as well as the model's parameter privacy. Compared with the standard CKKS, the secret key of MCKKS is securely shared with all entities. As a result, each entity still performs ciphertext evaluation under the same public key, while the decryption of the result requires the participation of all entities. As shown in [12], MCKKS has several attractive properties: (i) it is naturally suitable for floating-point arithmetic circuits, which facilitates the implementation of machine learning; (ii) it flexibly supports collaborative computing among multiple users without revealing the respective share of the secret key; (iii) it supports the function of key-switch, making it possible to convert a ciphertext encrypted under a public key into a

 $(x, y) \in D$ , training a neural network of  $\mathbb{L}$  layers is generally  $1. \varphi'(\cdot)$  and  $\odot$  indicate partial derivative and element-wise product. Authorized licensed use limited to: Nanyang Technological University Library. Downloaded on October 06,2023 at 19:53:34 UTC from IEEE Xplore. Restrictions apply.

- **SecKeyGen** $(1^{\lambda})$ : Given a security parameter  $\lambda$ , 1) output a secret key  $sk_i$  for each user  $i \in [N]$ , where [N] is the shorthand  $\{1, 2, \dots N\}$  and  $\sum_{i=1}^{i=N} sk_i =$ sk.
- **DKeyGen**( $\{sk_i\}$ ): Given the set of secret keys 2)  $\{sk_i\}, i \in [N]$ , output the collective public key pk.
- 3)  $Ecd(\cdot)$ : Given a plaintext *m* (or a plaintext vector **m**) whose dimension does not exceed  $\mathcal{N}/2$ ), output the encoded (packed) plaintext  $\hat{m} \in R_{Q_{C}}$ , with scale  $\Delta$ .
- **Dcd**( $\hat{m}$ ): Given an encoded (packed) plaintext  $\hat{m} \in$ 4)  $R_{Q_{\mathcal{L}_m}}$  with scale  $\Delta_m$ , output the decoding of m (or the plaintext vector **m**).
- $\mathbf{Enc}(pk, \hat{m})$ : Given the collective public key pk, and 5) an encoded (packed) plaintext  $\hat{m} \in R_{Q_{\mathcal{L}}}$ , output the ciphertext  $[\mathbf{c}]_{pk} \in R_{Q_{\mathcal{L}}}^2$  with scale  $\Delta$ .
- **DDec**([**c**]<sub>*pk*</sub>, {*sk<sub>i</sub>*}): Given a ciphertext [**c**]<sub>*pk*</sub>  $\in R^2_{Q_{\mathcal{L}_{\mathbf{c}}}}$  with scale  $\Delta_{\mathbf{c}_{pk}}$ , and the set of secret keys {*sk<sub>i</sub>*}, 6)  $i \in [1, N]$ , output the plaintext  $p \in R_{Q_{\mathcal{L}_{c}}}$  with scale  $\Delta_{\mathbf{c}_{nl}}$
- 7)  $\mathbf{Add}([\mathbf{c}]_{pk}, [\mathbf{c}']_{pk})$ : Given two ciphertexts  $[\mathbf{c}]_{pk}$  and  $[\mathbf{c}']_{pk}$  encrypted with the same public key pk, output  $[\mathbf{c} + \mathbf{c}']_{pk}$  with level  $\min(\mathcal{L}_{\mathbf{c}_{pk}}, \mathcal{L}_{\mathbf{c}'_{pk}})$  and scale  $\max(\Delta_{\mathbf{c}_{pk}}, \Delta_{\mathbf{c}'_{pk}}).$
- $\begin{aligned} \mathbf{Sub}([\mathbf{c}]_{pk}, [\mathbf{c}']_{pk}^{pk}): \text{ Given two ciphertexts } [\mathbf{c}]_{pk} \text{ and } \\ [\mathbf{c}']_{pk'} \text{ output } [\mathbf{c} \mathbf{c}']_{pk} \text{ with level } \min(\mathcal{L}_{\mathbf{c}_{pk}}, \mathcal{L}_{\mathbf{c}'_{pk}}) \end{aligned}$ 8) and scale  $\max(\Delta_{\mathbf{c}_{pk}}, \Delta_{\mathbf{c}'_{pk}})$ .
- $\operatorname{Mul}_{pt}([\mathbf{c}]_{pk}, \hat{m})$ : Given a ciphertext  $[\mathbf{c}]_{pk}$  and an 9) encoded (packed) plaintext  $\hat{m}$ , output  $[\mathbf{c}m]_{pk}$  with level min $(\mathcal{L}_{\mathbf{c}_{pk}}, \mathcal{L}_{\mathbf{c}'_{pk}})$  and scale  $\Delta_{\mathbf{c}_{pk}} \times \Delta_{m}$ .
- $\operatorname{Mul}_{ct}([\mathbf{c}]_{pk}, [\mathbf{c}']_{pk})$ : Given two ciphertexts  $[\mathbf{c}]_{pk}$  and 10) $[\mathbf{c}']_{pk}$ , output  $[\mathbf{cc}']_{pk}$  with level  $\min(\mathcal{L}_{\mathbf{c}_{pk}}, \mathcal{L}_{\mathbf{c}'_{pk}})$  and scale  $\Delta_{\mathbf{c}_{pk}} \times \Delta_{\mathbf{c}'_{pk}}$ .
- $\mathbf{Rot}([\mathbf{c}]_{pk}, k)$ : Given a ciphertexts  $[\mathbf{c}]_{pk}$ , homomor-11) phically rotate  $[\mathbf{c}]_{pk}$  to the right (k > 0) or to the left (k < 0) by k times.
- $\mathbf{RS}([\mathbf{c}]_{pk})$ : Given a ciphertexts  $[\mathbf{c}]_{pk}$ , output  $[\mathbf{c}]_{pk}$ 12)
- with scale  $\Delta_{\mathbf{c}}/q_{\Delta_{\mathbf{c}}}$  and level  $\mathcal{L}_{\mathbf{c}} 1$ . **DKeySwitch**([ $\mathbf{c}$ ]<sub>*pk*</sub>, *pk'*, {*sk*<sub>*i*</sub>}): Given a ciphertexts [ $\mathbf{c}$ ]<sub>*pk*</sub>, another public key *pk'*, and the set of secret 13) keys  $\{sk_i\}, i \in [N]$ , output  $[\mathbf{c}]_{pk'}$ .
- **DBootstrap**([**c**]<sub>*pk*</sub>,  $\mathcal{L}_{\mathbf{c}_{pk}}$ ,  $\Delta_{\mathbf{c}_{pk}}$ ,  $\{sk_i\}$ ): Given a ciphertexts [**c**]<sub>*pk*</sub> with level  $\mathcal{L}_{\mathbf{c}_{pk}}$  and scale  $\Delta_{\mathbf{c}_{pk}}$ , and 14) the set of secret keys  $\{sk_i\}, i \in [N]$ , output  $[\mathbf{c}]_{pk}$ with initial  $\mathcal{L}$  and scale  $\Delta$ .

#### Fig. 1. Cryptographic operations of MCKKS.

ciphertext under another public key without decryption. Such a property facilitates the decryption of ciphertexts collaboratively. We provide a short description of MCKKS and list all the functions required by Hercules in Fig. 1. Informally, given a cyclotomic polynomial ring with a dimension of  $\mathcal{N}$ , the plaintext and ciphertext space of MCKKS is defined as  $R_{Q_{\mathcal{L}}} = \mathbb{Z}_{Q_{\mathcal{L}}}[X]/(X^{\mathcal{N}}+1)$ , where  $Q_{\mathcal{L}} = \prod_{i=1}^{\mathcal{L}} q_i$ , and each  $q_i$  is a unique prime.  $Q_{\mathcal{L}}$  is the ciphertext module under the initial level *L*. In CKKS, a plaintext vector with up to  $\mathcal{N}/2$  values can be encoded into a ciphertext. As shown in Fig. 1, given a plaintext  $m \in R_{Q_{\mathcal{L}}}$  (or a plaintext vector  $\mathbf{m} =$  $(m_1,\ldots,m_n) \in \mathbb{R}^n_{Q_{\mathcal{L}}}$ , with  $n \leq \mathcal{N}/2$  with its encoded (packed) plaintext  $\hat{m}$ , the corresponding ciphertext is denoted as  $[\mathbf{c}]_{pk} = (c_1, c_2) \in R^2_{Q_{\mathcal{L}}}$ . Besides, we use symbols  $\mathcal{L}_{\mathbf{c}_{pk}}, \Delta_{\mathbf{c}_{pk}}, \mathcal{L}, \Delta$  to indicate the current level of  $[\mathbf{c}]_{pk}$ , the current scale of c, the initial level, and the initial scale of a fresh ciphertext, respectively. All functions named starting with **D** (except for  $Dcd(\cdot)$ ) in Fig. 1 need to be executed cooperatively by all the users, while the rest operations can be executed locally by each user with the public key. For more details about MCKKS, please refer to literature [1], [12], [25].

# 2.3 Threat Model and Privacy Requirements

We consider a FL scenario composed of a parameter server and N users for training a neural network model collaboratively. Specifically, the server (also the model owner) first initializes the target model  $\mathcal{M}$  and broadcasts the encrypted model  $[\mathbf{M}]_{pk} = \mathbf{Enc}(pk, \mathcal{M})$  (i.e., encrypting all the model parameters) to all the users.<sup>2</sup> Then, each user  $P_i$  with a dataset  $\{x, y\} \in D_i$  trains  $[\mathbf{M}]_{pk}$  locally using the mini-batch SGD algorithm and then sends the encrypted local gradients to the server. After receiving the gradients from all the users, the server homomorphically aggregates them and broadcasts back the global model parameters. All the participants perform the above process iteratively until the model converges. Since the final trained model is encrypted with the public key pk, for the accessibility of the server to the plaintext model, we rely on the function DKeySwitch (Fig. 1), which enables the conversion of  $[\mathbf{M}]_{pk}$  under the public key pk into  $[\mathbf{M}]_{pk'}$  under the server's public key pk' without decryption (refer to Section 5 for more details). As a result, the server obtains the plaintext model by decrypting  $[\mathbf{M}]_{nk'}$ with its secret key.

In Hercules, we consider a passive-adversary model with collusion of up to N - 1 users<sup>3</sup>. Concretely, the server and each user abide by the agreement and perform the training procedure honestly. However, there are two ways of colluding in Hercules by sharing their own inputs, outputs and observations during the training process for different purposes: (i) collusion among up to N-1 users to derive the training data of other users or the model parameters of the server; (ii) collusion among the server and no more than N-1 users to infer the training data of other users. Given such a threat model, in the training phase, the privacy requirements of Hercules are defined as below:

- Data privacy: No participant (including the server) should learn more information about the input data (e.g., local datasets, intermediate values, local gradients) of other honest users, except for the information that can be inferred from its own inputs and outputs.
- Model privacy: No user should learn more information about the parameters of the model, except for information that can be inferred from its own inputs and outputs.

In Section 5, we will provide (sketch) proofs of these privacy requirements with the real/ideal simulation formalism [42], [43], [44].

<sup>2.</sup> Note that the server knows nothing about the secret key sk corresponding to pk. sk is securely shared with N users and can only be restored with the participation of all the users.

<sup>3.</sup> See Appendix, which can be found on the Computer Society Digital Library at http://doi.ieeecomputersociety.org/10.1109/TDSC.2022.3218 793, for more discussion about malicious adversary model [36], [37], [38], [39], [40], [41]

Setup: Given two ciphertexts  $[\mathbf{A}]_{pk}$  and  $[\mathbf{B}]_{pk}$  that are the encryption forms of two  $(3 \times 3)$ -dimensional matrix matrices A and B (shown below), respectively, we now describe how to efficiently evaluate their homomorphic matrix multiplication.  $A = \begin{bmatrix} a_0 & a_1 & a_2 \\ a_3 & a_4 & a_5 \\ a_6 & a_7 & a_8 \end{bmatrix}; \quad B = \begin{bmatrix} b_0 & b_1 & b_2 \\ b_3 & b_4 & b_5 \\ b_6 & b_7 & b_8 \end{bmatrix}, \text{ where the vector representations of } A \text{ and } B \text{ are } B = \begin{bmatrix} a_0 & a_1 & a_2 \\ a_3 & a_4 & a_5 \\ b_6 & b_7 & b_8 \end{bmatrix},$  $\begin{bmatrix} b_0 & b_1 & b_2 \end{bmatrix}$  $\mathbf{a} = [a_0, a_1, a_2, a_3, a_4, a_5, a_6, a_7, a_8]$  and  $\mathbf{b} = [b_0, b_1, b_2, b_3, b_4, b_5, b_6, b_7, b_8]$ , respectively. **Step 1-1**: From A and B, we first compute  $U^{\mu}$ ,  $U^{\zeta}$ ,  $V = \{V^1, V^2\}$  and  $P = \{P^1, P^2\}$  based on Eqn.(5)-(8) as follows.  $0 \quad 0 \quad 0^{-1}$ 0 0 0 0 0 0 0 0 0  $0 \ 1 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0$  $0 \quad 0 \quad 1$ 0 0 0 0 0 0 0  $0 \; ; V^1 =$ 0 0 - 0  $\begin{array}{ccc} 0 & 0 \\ 0 & 1 \end{array}$ 0 0 0 0 0 0 0 0 0 0  $0 \ 0 \ 0 \ 0 \ 0$  $1 \quad 0 \quad 0 \quad 0 \quad 0$ 0 0 0 0 г0 0 0 г0 0 0 0 0 0 0  $0 \ 0 \ 0 \ 1$ 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 1 0 0 0  $0 \ 0 \ 0 \ 0 \ 1 \ 0$ 0 0  $0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 1 \ 0 \ 0$ 0 0  $0 ; P^1 = 0$ 0 0 0 0 0 0 1 0 0 0  $0 \ 1 \ 0$ 0 0 0 0 0 0  $0 \ 0 \ 1$  $0 \ 1$ -0  $0 \ 1$ 0 0 - 0 - 0 - 0  $0 \ 0 \ 0 \ 0 \ 0$ 0 0  $1 \ 0 \ 0 \ 0$ 0 0 0 0 - 0  $\lceil a_0 \rceil$  $a_1$  $a_2$ We securely compute  $U^{\mu} \cdot \mathbf{a}$ . Based on Eqn.(9), we have  $U^{\mu} \cdot \mathbf{a} = [a_0, a_1, a_2, a_4, a_5, a_3, a_8, a_6, a_7] \stackrel{\iota(\mathbf{a})}{=}$  $a_4 \ a_5$  $a_3$  $|a_8 \ a_6 \ a_7|$ where  $U^{\mu}$  has exactly  $(2 \times 3 - 1) = 5$  nonzero diagonal vectors (based on Eqn.(10) and (11)), denoted by  $\mathbf{u}_{k}^{\mu}$ for  $k \in \mathbb{Z} \cap (-3,3)$ . Specifically,  $\mathbf{u}_{-2}^{\mu} = [0,0,0,0,0,1,0,0]$ ,  $\mathbf{u}_{-1}^{\mu} = [0,0,0,0,0,0,0,0,1,1]$ ,  $\mathbf{u}_{0}^{\mu} = [1,1,1,0,0,0,0,0,0,0,0]$ ,  $\mathbf{u}_{1}^{\mu} = [0,0,0,1,1,0,0,0,0]$ , and  $\mathbf{u}_{2}^{\mu} = [0,0,0,0,0,0,1,0,0]$ . Then, we can get the ciphertext of  $U^{\mu} \cdot \mathbf{a}$ , denoted by  $[\bar{\mathbf{A}}^{(0)}]_{pk}$ , based on Eqn.(12). **Step 1-2**: We securely compute  $U^{\zeta} \cdot \mathbf{b}$ . Based on Eqn.(13), we have  $U^{\zeta} \cdot \mathbf{b} = [b_0, b_4, b_8, b_3, b_7, b_2, b_6, b_1, b_5] \stackrel{\iota(\mathbf{b})}{=} \begin{bmatrix} b_0 & b_4 & b_8 \\ b_3 & b_7 & b_2 \end{bmatrix}$ where  $U^{\zeta}$  has exactly h = 3 nonzero diagonal vectors, denoted by  $\mathbf{u}_{3\cdot k}^{\zeta}$ , for  $0 \leq k < 3$ . Specifically,  $\mathbf{u}_{0}^{\zeta} = [1, 0, 0, 1, 0, 0, 1, 0, 0], \ \mathbf{u}_{3}^{\zeta} = [0, 1, 0, 0, 1, 0, 0, 1, 0], \ \mathbf{u}_{6}^{\zeta} = [0, 0, 1, 0, 0, 1, 0, 0, 1].$  Then, we can get the ciphertext of  $U^{\zeta} \cdot \mathbf{b}$ , denoted by  $[\mathbf{B}^{(0)}]_{pk}$ , based on Eqn.(14). **Step 2**: This step is used to securely perform column and row shifting operations on  $\mu(A)$  and  $\zeta(B)$  respectively. Specifically, for each column shifting matrix  $V^k$ ,  $1 \le k < 3$ , it has two nonzero diagonal vectors  $\mathbf{v}_k$  and  $\mathbf{v}_{k-h}$  (based on Eqn.(15) and (16)). Hence, the nonzero diagonal vectors in  $V^1$  are  $\mathbf{v}_1 = [1, 1, 0, 1, 1, 0, 1, 1, 0]$ and  $\mathbf{v}_{-2} = [0, 0, 1, 0, 0, 1, 0, 0, 1]$ , and the nonzero diagonal vectors in  $V^2$  are  $\mathbf{v}_2 = [1, 0, 0, 1, 0, 0, 1, 0, 0]$  and  $\mathbf{v}_{-1} = [0, 1, 1, 0, 1, 1]$ . Similarly, for each row shifting matrix  $P^k$ , it has only one nonzero diagonal vector  $\mathbf{p}_{3\cdot k}$ . Then the nonzero diagonal vector in  $P^1$  is  $\mathbf{p}_3 = [1, 1, 1, 1, 1, 1, 1]$  and the nonzero diagonal vector in  $P^2$ are  $\mathbf{p}_6 = [1, 1, 1, 1, 1, 1, 1, 1]$ . Based on this, we can obtain the ciphertexts  $[\mathbf{A}^{(1)}]_{pk}$ ,  $[\mathbf{A}^{(2)}]_{pk}$ ,  $[\mathbf{B}^{(1)}]_{pk}$ , and  $[\mathbf{B}^{(2)}]_{pk}$  of the matrix  $\phi^1 \circ \mu(A)$ ,  $\phi^2 \circ \mu(A)$ ,  $\pi^1 \circ \zeta(B)$ , and  $\pi^2 \circ \zeta(B)$ , respectively, where  $\phi^{1} \circ \mu(A) = \begin{bmatrix} a_{1} & a_{2} & a_{0} \\ a_{5} & a_{3} & a_{4} \\ a_{6} & a_{7} & a_{8} \end{bmatrix}; \phi^{2} \circ \mu(A) = \begin{bmatrix} a_{2} & a_{0} & a_{1} \\ a_{3} & a_{4} & a_{5} \\ a_{7} & a_{8} & a_{6} \end{bmatrix}; \pi^{1} \circ \zeta(B) = \begin{bmatrix} b_{3} & b_{7} & b_{2} \\ b_{6} & b_{1} & b_{5} \\ b_{0} & b_{4} & b_{8} \end{bmatrix}; \pi^{2} \circ \zeta(B) = \begin{bmatrix} b_{6} & b_{1} & b_{5} \\ b_{0} & b_{4} & b_{8} \\ b_{3} & b_{7} & b_{2} \end{bmatrix}$ **Step 3**: For  $0 \le k < 3$ , we compute the element-wise multiplication between  $[\mathbf{A}^{(k)}]_{pk}$  and  $[\mathbf{B}^{(k)}]_{pk}$ . Then,  $[\mathbf{AB}]_{pk}$  is obtained as below.

| $\begin{bmatrix} a_6 & a_7 & a_8 \end{bmatrix} \begin{bmatrix} b_6 & b_7 & b_8 \end{bmatrix} \begin{bmatrix} a_8 & a_6 & a_7 \end{bmatrix} \begin{bmatrix} b_6 & b_1 & b_5 \end{bmatrix} \begin{bmatrix} a_6 & a_7 & a_8 \end{bmatrix} \begin{bmatrix} b_0 & b_4 & b_8 \end{bmatrix} \begin{bmatrix} a_7 & a_8 & a_6 \end{bmatrix} \begin{bmatrix} b_3 & b_4 & b_8 \end{bmatrix} \begin{bmatrix} a_7 & a_8 & a_6 \end{bmatrix} \begin{bmatrix} b_8 & b_8 \end{bmatrix} \begin{bmatrix} a_8 & a_8 & a_8 \end{bmatrix} \begin{bmatrix} a_8 & a_$ | $\begin{bmatrix} a_1 & a_2 \\ a_4 & a_5 \\ a_7 & a_8 \end{bmatrix} \cdot \begin{bmatrix} \\ \\ \end{bmatrix}$ | $\begin{bmatrix} b_0 & b_1 & b_2 \\ b_3 & b_4 & b_5 \\ b_6 & b_7 & b_8 \end{bmatrix}$ | $ = \begin{bmatrix} a_0 & a_1 \\ a_4 & a_5 \\ a_8 & a_6 \end{bmatrix} $ | $\begin{bmatrix} a_2\\a_3\\a_7 \end{bmatrix} \odot \begin{bmatrix} b_0 & b_4\\b_3 & b_7\\b_6 & b_1 \end{bmatrix}$ | $ \begin{bmatrix} b_8 \\ b_2 \\ b_5 \end{bmatrix} + \begin{bmatrix} a_1 & a_2 \\ a_5 & a_3 \\ a_6 & a_7 \end{bmatrix} $ | $\begin{bmatrix} a_0 \\ a_4 \\ a_8 \end{bmatrix} \odot \begin{bmatrix} b_3 & b_7 \\ b_6 & b_1 \\ b_0 & b_4 \end{bmatrix}$ | $ \begin{bmatrix} b_2 \\ b_5 \\ b_8 \end{bmatrix} + \begin{bmatrix} a_2 & a_0 \\ a_3 & a_4 \\ a_7 & a_8 \end{bmatrix} $ | $\begin{bmatrix} a_1\\ a_5\\ a_6 \end{bmatrix} \odot \begin{bmatrix} b_6 & b_1\\ b_0 & b_4\\ b_3 & b_7 \end{bmatrix}$ | $b_5 \\ b_8 \\ b_2 \end{bmatrix}$ |
|--|---|---|---|---|---|---|---|---|-----------------------------------|
|--|---|---|---|---|---|---|---|---|-----------------------------------|

Fig. 2. Homomorphic multiplication of two  $3 \times 3$ -dimensional matrices.

# 3 PARALLELIZED MATRIX HOMOMORPHIC OPERATIONS

**Hercules** essentially exploits MCCK as the underlying architecture to implement privacy-preserving federated neural network training. Since the vast majority of the computation of a Authorized licensed use limited to: Nanyang Technological University Library. Dow neural network consists of convolutions (equivalent to matrix operation), **Hercules** is required to handle this type of operation homomorphically very frequently. In this section, we describe our optimization method to perform homomorphic matrix operations in a parallelized manner, thereby substantially improving the computation performance of HE.

#### 3.1 Overview

At a high level, operations between two matrices, including multiplication and transposition, can be decomposed into a series of combinations of linear transformations. To handle homomorphic matrix operations in an SIMD manner, a straightforward way is to directly perform the relevant linear operations under the packed ciphertext (Section 3.2). However, it is computationally intensive and requires  $O(h^3)$  computation complexity for the multiplication of two  $h \times h$ -dimensional matrices (Section 3.3). Existing state-of-the-art methods [34] propose to transform the multiplication of two  $h \times h$ -dimensional matrices into inner products between multiple vectors. It can reduce the complexity from  $O(h^3)$  to  $O(h^2)$ , however, yielding h ciphertexts to represent a matrix (Section 3.6). Compared to existing efforts, our method only needs O(h) complexity and derives one ciphertext. Our key insight is to first formalize the linear transformations corresponding to matrix operations, and then tweak them to minimize redundant operations in the execution process. In the following we present the technical details of our method. To facilitate understanding, Fig. 2 also provides an intuitive example, where the detailed steps of the multiplication of two  $3 \times 3$ -dimensional matrices are described for comprehensibility.

# 3.2 Preliminary Knowledge

We first introduce some useful symbols and concepts. Specifically, all the vectors in this section refer to row vectors, and are represented in bold (e.g., a). As shown in Fig. 1, given a plaintext vector  $\mathbf{m} = (m_1, \ldots, m_n) \in \mathbb{R}^n_{O_r}$ , with  $n \leq n$  $\tilde{N}/2$ , CKKS enables to encode the plaintext vector **m** into an encoded plaintext  $\hat{m} \in R_{Q_{\mathcal{L}}}$ , where each  $m_i$ ,  $i \in [n]$  has a unique position called a plaintext slot in the encoded  $\hat{m}$ . Then,  $\hat{m}$  is encrypted as a ciphertext  $[\mathbf{c}]_{nk}$ . Hence, performing arithmetic operations (including addition and multiplication) on  $[\mathbf{c}]_{vk}$  is equivalent to doing the same operation on every plaintext slot at once.

The ciphertext packing technology is capable of packing multiple plaintexts into one ciphertext and realizing the homomorphic SIMD operation, thereby effectively reducing the space and time complexity of encryption/calculation of a single ciphertext. However, it is incapable of handling the arithmetic circuits when some inputs are in different plaintext slots. To combat that, CKKS provides a rotation function **Rot**( $[\mathbf{c}]_{pk}$ , k). Given a ciphertext  $[\mathbf{c}]_{pk}$  of a plaintext vector  $\mathbf{m} =$  $(m_1, \ldots, m_n) \in R^n_{Q_{\mathcal{L}}}$ ,  $\mathbf{Rot}([\mathbf{c}]_{pk}, k)$  transforms  $[\mathbf{c}]_{pk}$  into an encryption of **R**( $\mathbf{m}, k$ ) := ( $m_k, ..., m_{n-1}, m_0, ..., m_{k-1}$ ). k can be either positive or negative and we have a rotation by  $\mathbf{R}(\mathbf{m},k) = \mathbf{R}(\mathbf{m},n-k).$ 

Based on the above explanation, we adopt a method proposed by Shai et al. [45], which supports arbitrary linear transformations for encrypted vectors. Specifically, an arbitrary linear transformation  $\mathcal{T}: \mathbb{R}^n \to \mathbb{R}^n$  on the plaintext vector can be expressed as  $\mathcal{T} : \mathbf{m} \to U \cdot \mathbf{m}$  using some matrix  $U \in \mathbb{R}^{n \times n}$ . This process can be implemented in ciphertext by the rotation function **Rot** and constant multiplication operation  $Mul_{vt}$ . Concretely, for  $0 \le k < n$ , a *k*-th *diagonal vector* U is defined as  $\mathbf{u}_k = (U_{0,k}, U_{1,k+1}, \dots, U_{n-k-1,n-1}, U_{n-k,0}, \dots, U_{n-1,k-1}) \in$  $R^n$ . Consequently, we have

Hence, given the matrix 
$$U$$
, and a ciphertext  $[c]_{pk}$  of the vector **m**, *Algorithm* 2 shows the details of computing encrypted  $U \cdot \mathbf{m}$ . We observe that *Algorithm* 2 requires  $n$  additions, constant multiplications and rotations. Because the rotation operation is much more intensive than the other two operations, the computation complexity of *Algorithm* 2 is usually regarded as asymptotically  $O(n)$  rotations.

| Algorithm 2. Homomorphic Linear Transformation  |  |
|---|--|
| procedure HE-LinTrans $([\mathbf{c}]_{pk}, U)$  |  |
| 1: $[\mathbf{c}']_{pk} \leftarrow \mathbf{Mul}_{pt}([\mathbf{c}]_{pk}, \mathbf{u}_0)$ |  |

2: for k = 1 to n - 1 do 3:  $[\mathbf{c}']_{pk} \leftarrow \mathbf{Add}([\mathbf{c}']_{pk}, \mathbf{Mul}_{pt}(\mathbf{Rot}([\mathbf{c}]_{pk}, k), \mathbf{u}_k))$ 4: end for 5: return  $[\mathbf{c}']_{nk}$ 

In the following, we first describe how to express the multiplication between two matrices by permutation. Then, we introduce an encoding method that converts a matrix into a vector. Based on this, we describe the details of matrix multiplication on packed ciphertexts.

### 3.3 Permutation for Matrix Multiplication

Given a  $(h \times h)$ -dimensional matrix  $A = (A_{i,j})_{0 \le i,j < h}$ , we describe four permutation operations  $(\mu, \zeta, \phi, \pi)$  on it. For simplicity, we use  $\mathbb{Z} \cap [0, h)$  to denote the representative of  $\mathbb{Z}_h$ ,  $[i]_h$  indicates the reduction of an integer *i* modulo *h* into that interval. Below all indexes are integers modulo *h*.

We first define four permutation operations as below.

$$\mu(A)_{i,j} = A_{i,i+j}; \zeta(A)_{i,j} = A_{i+j,j};$$
  
$$\phi(A)_{i,j} = A_{i,j+1}; \pi(A)_{i,j} = A_{i+1,j}$$

We can see that  $\phi$  and  $\pi$  are actually shifts of the columns and rows of the matrix, respectively. Given two  $(h \times$ h)-dimensional square matrices A and B, the multiplication of A and B can be parsed as

$$A \cdot B = \sum_{k=0}^{h-1} (\phi^k \circ \mu(A)) \odot (\pi^k \circ \zeta(B)).$$
<sup>(2)</sup>

The correctness of Eq. (2) is shown as follows by calculating the components of the matrix index (i, j).

$$\sum_{k=0}^{h-1} (\phi^k \circ \mu(A))_{i,j} \cdot (\pi^k \circ \zeta(B))_{i,j} = \sum_{k=0}^{h-1} \mu(A)_{i,j+k} \cdot \zeta(B)_{i+k,j}$$
$$= \sum_{k=0}^{h-1} A_{i,i+j+k} \cdot B_{i+j+k,j}$$
$$= \sum_{k=0}^{h-1} A_{i,k} \cdot B_{k,j} = (A \cdot B)_{i,j}.$$

Note that while a single  $\mu(A)_{i,j+k} \cdot \zeta(B)_{i+k,j} = A_{i,i+j+k} \cdot$  $B_{i+j+k,j}$  is not equal to  $A_{i,k} \cdot B_{k,j}$ , it is easy to deduce that  $\sum_{k=0}^{h-1} A_{i,i+j+k} \cdot B_{i+j+k,j} = \sum_{k=0}^{h-1} A_{i,k} \cdot B_{k,j} = (A \cdot B)_{i,j}.$  To be precise, given *i* and *j*,  $\sum_{k=0}^{h-1} A_{i,i+j+k} \cdot B_{i+j+k,j} =$  $\sum_{t=(i+j)}^{h-1+i+j} A_{i,t} \cdot B_{t,j}, \text{ where we set } t=i+j+k. \text{ Then, we}$ have  $\sum_{t=(i+j)}^{h-1+i+j} A_{i,t} \cdot B_{t,j} = \sum_{t=0}^{h-1} A_{i,t} \cdot B_{t,j} \text{ since all the}$ indexes are considered as integers modulo h. Therefore,

 $U \cdot \mathbf{m} = \sum_{0 \le k < n} \mathbf{u}_k \odot \mathbf{R}(\mathbf{m}, k).$ (1)  $\sum_{t=0}^{h-1} A_{i,t} \cdot B_{t,j} = \sum_{k=0}^{h-1} A_{i,k} \cdot B_{k,j}.$ Authorized licensed use limited to: Nanyang Technological University Library. Downloaded on October 06,2023 at 19:53:34 UTC from IEEE Xplore. Restrictions apply.

We observe that Eq. (2) consists of permutation and multiplication of element components between matrix entries. Intuitively, we can evaluate it using the operations (shown in *Algorithm* 2) provided by CKKS for packed ciphertexts. However, since the matrix representation U usually has  $n = h^2$  nonzero diagonal vectors, if we directly use *Algorithm* 2 to evaluate  $A \mapsto \phi^k \circ \mu(A)$  and  $B \mapsto \pi^k \circ \zeta(B)$  for  $1 \le k < h$ , each of them requires rotations with the complexity of  $O(h^2)$ . As a result, the total complexity is  $O(h^3)$ . To alleviate this, we design a new method to substantively improve its efficiency.

# 3.4 Matrix Encoding

We introduce an encoding method that converts a matrix into a vector. Given a vector  $\mathbf{a} = (a_k)_{0 \le k < n}$ , where  $n = h^2$ , the encoding map  $\iota : R^n \to R^{h \times h}$  is shown as below.

$$\iota : \mathbf{a} \mapsto A = (a_{h \cdot i+j})_{0 \le i,j < h}.$$

$$\tag{4}$$

This encoding method makes the vector **a** essentially an ordered concatenation of the rows of the matrix *A*. As a result,  $\iota(\cdot)$  is isomorphic of addition, which means that matrix addition operations are equivalent to the same operations between the corresponding original vectors. Therefore, the matrix addition can be calculated homomorphically in the SIMD environment. The constant multiplication operations can also be performed homomorphically. In this paper, we use  $\iota(\cdot)$  to identify two spaces  $R^n$  and  $R^{h \times h}$ . For example, we say that a ciphertext is the encryption of *A* if  $\mathbf{a} = \iota^{-1}(A)$ .

## 3.5 Tweaks of Permutation

From the definition of matrix encoding, permutation on an  $(h \times h)$ -dimensional matrix can be regarded as a linear transformation  $\mathcal{T}: \mathbb{R}^n \to \mathbb{R}^n$ , where  $n = h^2$ . In general, its matrix representation  $U \in \{0,1\}^{n \times n} \subset \mathbb{R}^{n \times n}$  has *n* nonzero diagonal vectors. Therefore, as presented in Sections 3.2 and 3.3, if we directly use Algorithm 2 to evaluate  $A \mapsto \phi^k \circ \mu(A)$  and  $B \mapsto \pi^k \circ \zeta(B)$  for  $1 \leq k < h$ , each of them requires rotations with the complexity of  $O(h^2)$ . The total complexity will be  $O(h^3)$ . To alleviate this, based on Eq. (2) and our matrix encoding map, we provide a tweak method for matrix permutation to reduce the complexity from  $O(h^3)$  to O(h). Specifically, for four permutation operations ( $\mu$ ,  $\zeta$ ,  $\phi$ , and  $\pi$ ) on the matrix, we use  $U^{\mu}$ ,  $U^{\zeta}$ , V and P to indicate the matrix representations corresponding to these permutations, respectively.  $U^{\mu}$ ,  $U^{\zeta}$  for permutations  $\mu$  and  $\zeta$  can be parsed as below (readers can refer to the example in Fig. 2 for ease of understanding).

$$U^{\mu}_{h\cdot i+j,t} = \begin{cases} 1 & \text{if } t = h \cdot i + [i+j]_h; \\ 0 & \text{otherwise}; \end{cases}$$
(5)

$$U_{h\cdot i+j,t}^{\zeta} = \begin{cases} 1 & \text{if } t = h \cdot [i+j]_h + j, \\ 0 & \text{otherwise}; \end{cases}$$
(6)

where  $0 \le i, j < h$  and  $0 \le t < h^2$ . Similarly, for  $1 \le k < h$ , the matrix representations of  $\phi^k$  and  $\pi^k$  (i.e.,  $V^k$  and  $P^k$ ) can be denoted as below.

$$V_{h\cdot i+j,t}^{k} = \begin{cases} 1 & \text{if } t = h \cdot i + [j+k]_{h}; \\ 0 & \text{otherwise:} \end{cases}$$
(7)

$$P_{h\cdot i+j,t}^{k} = \begin{cases} 1 & \text{if } t = h \cdot [i+k]_{h} + j; \\ 0 & \text{otherwise:} \end{cases}$$

$$\tag{8}$$

where  $0 \le i, j < h$  and  $0 \le t < h^2$ . Reviewing Eq. (1), we use the diagonal decomposition of matrix representation to perform multiplication with encrypted vectors. Hence, we can count the number of nonzero diagonal vectors in  $U^{\mu}$ ,  $U^{\zeta}$ , V, and *P* to evaluate the complexity. For simplicity, we use  $\mathbf{u}_t$  to represent the *t*-th diagonal vector of a matrix *U*, and identify  $\mathbf{u}_{h^2-t}$  with  $\mathbf{u}_{-t}$ . For matrix  $U^{\mu}$ , we can observe that it has exactly (2h-1) nonzero diagonal vectors, denoted by  $\mathbf{u}_k^{\mu}$  for  $k \in \mathbb{Z} \cap (-h, h)$ . There are h nonzero diagonal vectors in  $U^{\zeta}$ , because each *t*-th diagonal vector in  $U^{\zeta}$  is nonzero if and only if *t* is divisible by the integer *h*. For each matrix  $V^k$ ,  $1 \le k < k$ *h*, it has only two nonzero diagonal vectors  $\mathbf{v}_k$  and  $\mathbf{v}_{k-h}$ . Similarly, for each matrix  $P^k$ , it has only one nonzero diagonal vector  $\mathbf{p}_{h,k}$ . Therefore, we only need rotation operations of O(h)complexity to perform permutation  $\mu$  and  $\zeta$ , and O(1) complexity for both  $\phi^k$  and  $\pi^k$  where  $1 \le k < h$ .

## 3.6 Homomorphic Matrix Multiplication

Given two ciphertexts  $[\mathbf{A}]_{pk}$  and  $[\mathbf{B}]_{pk}$  that are the encryption forms of two  $(h \times h)$ -dimensional matrix matrices A and B, respectively, we now describe how to efficiently evaluate homomorphic matrix multiplication between them.

Step 1-1. We perform a linear transformation on the ciphertext  $[\mathbf{A}]_{pk}$  under the guidance of permutation  $U^{\mu}$  (Step 1-1 in Fig. 2). As described above,  $U^{\mu}$  has exactly (2h - 1) nonzero diagonal vectors, denoted by  $\mathbf{u}_{k}^{\mu}$  for  $k \in \mathbb{Z} \cap (-h, h)$ . Then such a linear transformation can be denoted as

$$U^{\mu} \cdot \mathbf{a} = \sum_{-h < k < h} (\mathbf{u}_{k}^{\mu} \odot \mathbf{R}(\mathbf{a}, k)),$$
(9)

where  $\mathbf{a} = \iota^{-1}(A) \in \mathbb{R}^n$  is the vector representation of A. If  $k \ge 0$ , the k-th diagonal vector can be computed as

$$\mathbf{u}_{k}^{\mu}[t] = \begin{cases} 1 & \text{if } 0 \le t - h \cdot k < (h - k); \\ 0 & \text{otherwise,} \end{cases}$$
(10)

where  $\mathbf{u}_k^{\mu}[t]$  represents the *t*-th component of  $\mathbf{u}_k^{\mu}$ . Similarly, if k < 0,  $\mathbf{u}_k^{\mu}$  is computed as

$$\mathbf{u}_{k}^{\mu}[t] = \begin{cases} 1 & \text{if } -k \leq t - (h+k) \cdot h < h; \\ 0 & \text{otherwise,} \end{cases}$$
(11)

As a result, Eq. (9) can be securely computed as

$$\sum_{-h < k < h} \mathbf{Mul}_{pt}(\mathbf{Rot}([\mathbf{A}]_{pk}, k), \mathbf{u}_k^{\mu}),$$
(12)

where we get the ciphertext of  $U^{\mu} \cdot \mathbf{a}$ , denoted as  $[\mathbf{A}^{(0)}]_{pk}$ . We observe that the computation cost is about 2 *h* rotations, constant multiplications and additions.

Step 1-2. This step is to perform the linear transformation on the ciphertext  $[\mathbf{B}]_{pk}$  under the guidance of permutation  $U^{\zeta}$  (Step 1-2 in Fig. 2). Since  $U^{\zeta}$  has *h* nonzero diagonal vectors, this process can be denoted as

$$U^{\zeta} \cdot \mathbf{b} = \sum_{0 \le k < h} (\mathbf{u}_{h \cdot k}^{\zeta} \odot \mathbf{R}(\mathbf{b}, h \cdot k)),$$
(13)

where  $\mathbf{b} = \iota^{-1}(B) \in \mathbb{R}^n$ ,  $\mathbf{u}_{h\cdot k}^{\zeta}$  is the  $(h \cdot k)$ -th diagonal vector of the matrix  $U^{\zeta}$ . We observe that for any  $0 \le k < h$ ,  $\mathbf{u}_{h\cdot k}^{\zeta}$  is a non-zero vector because its  $(k + h \cdot i)$ -th element is nonzero for  $0 \le i < h$ , and zero for all other entries. Therefore, Eq. (13) can be securely computed as

$$\sum_{0 \le k \le h} \mathbf{Mul}_{pt}(\mathbf{Rot}([\mathbf{B}]_{pk}, h \cdot k), \mathbf{u}_{h \cdot k}^{\zeta}), \tag{14}$$

where we get the ciphertext of  $U^{\zeta} \cdot \mathbf{b}$ , denoted as  $[\mathbf{B}^{(0)}]_{pk}$ . We observe that the computation cost is about *h* rotations, constant multiplications and additions.

*Step* 2. This step is used to securely perform column and row shifting operations on  $\mu(A)$  and  $\zeta(B)$  respectively (*Step* 2 in Fig. 2). Specifically, for each column shifting matrix  $V^k$ ,  $1 \le k < h$ , it has only two nonzero diagonal vectors  $\mathbf{v}_k$  and  $\mathbf{v}_{k-h}$ , which are computed as

$$\mathbf{v}_{k}[t] = \begin{cases} 1 & \text{if } 0 \le [t]_{h} < (h-k); \\ 0 & \text{otherwise,} \end{cases}$$
(15)

$$\mathbf{v}_{k-h}[t] = \begin{cases} 1 & \text{if } (h-k) \le |t|_h < h; \\ 0 & \text{otherwise.} \end{cases}$$
(16)

By adding two ciphertexts  $\mathbf{Mul}_{pt}(\mathbf{Rot}([\mathbf{A}^{(0)}]_{pk},k),\mathbf{v}_k)$  and  $\mathbf{Mul}_{pt}(\mathbf{Rot}([\mathbf{A}^{(0)}]_{pk},k-h),\mathbf{v}_{k-h})$ , we can obtain the ciphertext  $[\mathbf{A}^{(k)}]_{pk}$  of the matrix  $\phi^k \circ \mu(A)$ . Similarly, for each row shifting matrix  $P^k$ , it has only one nonzero diagonal vector  $\mathbf{p}_{h,k}$ . Then the encryption of  $\pi^k \circ \zeta(B)$  can be computed as  $[\mathbf{B}^{(k)}]_{pk} \leftarrow \mathbf{Rot}([\mathbf{B}^{(0)}]_{pk},h \cdot k)$ . The computation cost of this process is about 3 *h* rotations, 2 *h* constant multiplications and *d* additions.

*Step* 3. For  $0 \le k < h$ , we now compute the element-wise multiplication of  $[\mathbf{A}^{(k)}]_{pk}$  and  $[\mathbf{B}^{(k)}]_{pk}$  (*Step* 3 in Fig. 2). Then, the ciphertext  $[\mathbf{AB}]_{pk}$  of the product of A and B is finally obtained. The computation cost of this process is h homomorphic multiplications and additions. In summary, the entire process of performing homomorphic matrix multiplication is described in *Algorithm* 3.

| Algorithm 3. Homomorphic Matrix Multiplication   |
|--|
| procedure HE-MatMult $([\mathbf{A}]_{pk}, [\mathbf{B}]_{pk})$  |
| 1: $[\mathbf{A}^{(0)}]_{pk} \leftarrow \texttt{HE-LinTrans}([\mathbf{A}]_{pk}, U^{\mu})$   |
| 2: $[\mathbf{B}^{(0)}]_{pk} \leftarrow \texttt{HE-LinTrans}([\mathbf{B}]_{pk}, U^{\zeta})$   |
| 3: for $k = 1$ to $h - 1$ do   |
| 4: $[\mathbf{A}^{(k)}]_{pk} \leftarrow \texttt{HE-LinTrans}([\mathbf{A}^{(0)}]_{pk}, V^k)$   |
| 5: $[\mathbf{B}^{(k)}]_{pk}^{r} \leftarrow \text{HE-LinTrans}([\mathbf{B}^{(0)}]_{pk}^{r}, P^k)$   |
| 6: end for   |
| 7: $[\mathbf{AB}]_{pk} \leftarrow \mathbf{Mul}_{ct}([\mathbf{A}^{(0)}]_{pk}, [\mathbf{B}^{(0)}]_{pk})$                                   |
| 8: for $\vec{k} = 1$ to $h - 1$ do   |
| 9: $[\mathbf{AB}]_{pk} \leftarrow \mathbf{Add}([\mathbf{AB}]_{pk}, \mathbf{Mul}_{ct}([\mathbf{A}^{(k)}]_{pk}, [\mathbf{B}^{(k)}]_{pk}))$ |
| 10: end for  |
| 11: return $[\mathbf{AB}]_{nk}$  |

*Remark 3.1:* In general, the above homomorphic matrix multiplication requires a total of 5 h additions, 5 h constant multiplications and 6 h rotations. We can further reduce the computation complexity by using the baby-step/giant-step algorithm [46], [47] (See Appendix, available in the online supplemental material, for technical details). This algorithm can be exploited to reduce the complexity of Steps 1-1 and 1-2. As a result, Table 1 summarizes the computation complexity required for each step in *Algorithm* 3.

*Remark* 3.2: As described before, the multiplication of *A* and *B* is parsed as  $A \cdot B = \sum_{k=0}^{h-1} (\phi^k \circ \mu(A)) \odot (\pi^k \circ \zeta(B))$ . A simple way to calculate the product is to directly use *Algorithm* 2: we can evaluate  $A \mapsto \phi^k \circ \mu(A)$  and  $B \mapsto \pi^k \circ \zeta(B)$  for  $1 \le k < h$ . However, each of them requires  $O(h^2)$  homomorphic rotation operations, which results in a total complexity of  $O(h^3)$  [48]. Halevi et al. [34] introduce a matrix encoding

TABLE 1 Complexity of Algorithm 3

| Add | $mul_{pt}$   | Rot   | $mul_{ct}$   |
|-----|--|---|--|
| 2h  | 2h   | $3\sqrt{h}$   | -  |
| h   | h  | $2\sqrt{h}$   | -  |
| 2h  | h  | 3h  | -  |
| h   | -  | —   | h  |
| 6h  | 4h   | $3h + 5\sqrt{h}$  | h  |
|     | $\begin{array}{c} Add \\ 2h \\ h \\ 2h \\ h \\ 6h \end{array}$ | $\begin{array}{c c} Add & mul_{pt} \\ \hline 2h & 2h \\ h & h \\ 2h & h \\ h & - \\ 6h & 4h \\ \end{array}$ | $\begin{array}{c ccc} Add & mul_{pt} & Rot \\ \hline 2h & 2h & 3\sqrt{h} \\ h & h & 2\sqrt{h} \\ 2h & h & 3h \\ h & - & - \\ 6h & 4h & 3h + 5\sqrt{h} \\ \hline \end{array}$ |

method based on diagonal decomposition. This method maps each diagonal vector into a separate ciphertext by arranging the matrix diagonally. As a result, it requires h ciphertexts to represent a matrix, and each ciphertext is required to perform matrix-vector multiplication with the complexity of O(h) rotations, resulting in a total computation complexity of  $O(h^2)$ . Compared with these schemes, our strategy only needs a total computation complexity of O(h) rotations to complete the homomorphic multiplication for two  $(h \times h)$ -dimensional matrices. We note that POSEIDON [12] also proposes an "alternating packing (AP) approach" to achieve matrix multiplication with a complexity approximated as  $\max_{i \in [L]} (\omega_i \times$  $\log(h \times \omega_i)$ , where  $\omega_i$  denotes the number of weights between layers i and i + 1. However, the implementation of this method requires to generate a large number of copies of each element in the matrix (depending on the number of neurons in the neural network layer where the matrix is located), resulting in poor parallel computing performance (see Section 6 for more experimental comparison).

*Remark* 3.3: We also give the methods of how to perform matrix transposition, rectangular matrix multiplication (i.e., calculating general matrix forms such as  $R^{t \times h} \times R^{h \times h} \rightarrow R^{t \times h}$  or  $R^{h \times h} \times R^{h \times t} \rightarrow R^{h \times t}$ ) and parallel matrix operations (using the idleness of the plaintext slots) under packed ciphertext. They follow the similar idea of the above homomorphic matrix multiplication. Readers can refer to Appendix, available in the online supplemental material, for more technical details.

# 4 APPROXIMATION FOR SIGN FUNCTION

In this section, we describe how to efficiently estimate the sign function, and then use the estimated function to approximate the formulas commonly used in neural network training, including ReLU and max functions.

# 4.1 Notations

We first introduce some useful symbols. Specifically, all logarithms are base 2 unless otherwise stated.  $\mathbb{Z}$  and  $\mathbb{R}$  represent the integer and real number fields, respectively. For a finite set M, we use U(M) to represent the uniform distribution on M. Given a function g defined in the real number field  $\mathbb{R}$ , and a compact set  $I \subset \mathbb{R}$ , we say that the infinity norm of g on the set I is defined as  $||g||_{\infty,I} := \max_{m \in I} |g(m)|$ , where |g(m)| means the absolute value of g(m). we use  $g^{(k)}$  :  $= g \circ g \circ g \circ \cdots \circ g$  to indicate the k-times composition of g. Besides, the sign function is defined as below.

$$sgn(m) = \begin{cases} 1 & \text{if } m > 0; \\ 0 & \text{if } m = 0; \\ -1 & \text{if } m < 0. \end{cases}$$

Note that sqn(m) is a discontinuous function at the zero point, so the closeness of g(m) and sgn(m) should be carefully considered in the interval near the zero point. That is, we do not consider the small interval  $(-\delta, \delta)$  near the zero point when measuring the difference between g(m) and sgn(m). We will prove that for some  $k_d > 0$ , the infinity norm of  $g_d^{(k)}(m) - sgn(m)$  is small than  $2^{-\sigma}$  over  $[-1, -\delta] \cup$  $[\delta, 1]$  if  $k > k_d$ , where the definition of  $g_d(m)$  will be explained later.

Given  $\sigma > 0$  and  $0 < \delta < 1$ , we define a function  $g_J^{(k)}(m)$ that is  $(\sigma, \delta)$ -close to sgn(m) on [-1, 1] if it satisfies

$$||g_d^{(k)}(m) - sgn(m)||_{\infty, [-1, -\delta] \cup [\delta, 1]} \le 2^{-\sigma}.$$
(17)

Similar to the previous work [35], we assume that the input is limited to a bounded interval [0, 1], since for any  $m \in$  $[a_1, a_2]$ , where  $a_2 > a_1$ , we can scale it down to [0, 1] by mapping  $m \mapsto (m - a_1)/(a_2 - a_1)$ . Hence, for simplicity, the domain of sgn(m) we consider in this part is [-1, 1].

## 4.2 Composite Polynomial Approximation

As mentioned before, we use a composite function to approximate the sign function. This is advantageous, because a composite polynomial function G, namely  $G = g \circ$  $g \cdots \circ g$ , can be calculated with the complexity of  $O(\log(deg(G)))$ , while the computation complexity of calculating any polynomial G is at least  $\Theta(\sqrt{deg(G)})$  [49], where deg(G) indicates the degree of G. To achieve this, our goal is to find such a k that  $q^{(k)}$  is close enough to sgn(x) in the interval  $[-1, -\delta] \cup [\delta, 1]$ .

Our construction of such a function *q* comes from the following key observations: for any  $m_0 \in [-1, 1]$ , let  $m_i$  be the *i*-th composite value of  $q^{(i)}(m_0)$ . Then, we can easily estimate the behavior of  $m_i$  through the graph of g. Based on this, we ensure that as *i* increases,  $m_i$  should be close to 1 when  $m_0 \in (0,1]$ , and close to -1 when  $m_0 \in [-1,0)$ . Besides, we formally identify three properties of *q* as follow. First, g should be an odd function so as to be consistent with the sign function. Second, g(1) = 1 and g(-1) = -1. This setting makes  $g^{(k)}(m)$  point-wise converge to sgn(m), whose value is  $\pm 1$  for all  $x \neq 0$ . In other words, for some  $m \in$ [-1,1],  $g^{(k)}(m)$  converges to a value y when increasing with the value of k, which means  $g(y) = g(\lim_{k\to\infty} g^{(k)}(m)) =$  $\lim_{k\to\infty} g^{(k)}(m) = y$ . Last, to accelerate the convergence of  $g^{(k)}$  to the sign function, a satisfactory g should be more concave in the interval [0,1] and more convex in the interval [-1,0]. Moreover, the derivative g' of g should have multiple roots at 1 and -1 so as to increase the convexity. These properties are summarized as follows:

*Core Properties of g:* 

Prop. I g(-m) = -g(m)(Origin Symmetry) Prop. II g(1) = 1, g(-1) = -1Prop. III  $g'(m) = p(1-m)^d (1+m)^d$ (Convergence to  $\pm 1$ )

for some p > 0(Fast convergence)

Given a fixed  $d \ge 1$ , a polynomial g of degree (2d + 1)that satisfies the above three properties can be uniquely determined. We denote this polynomial as  $g_d$ , where the constant *p* is indicated as  $p_d$ . Then, based on Prop. I and III, we have  $g_d(m) = p_d \int_0^m (1 - t^2)^d dt$ , where the constant  $p_d$  is also determined by Prop. II. To solve this integral formula Lemmas 4.2 and 4.3 to analyze the lower bound of k Authorized licensed use limited to: Nanyang Technological University Library. Downloaded on October 06,2023 at 19:53:34 UTC from IEEE Xplore. Restrictions apply.

 $q_d(m)$ , a common method is to transform the  $(1-t^2)$  part of the integral formula with Trigonometric Substitutions, a typical technique which can convert formula  $\int (1-t^2)^d dt$  to  $\int (\cos t)^{3d} dt$ . As a result, given the following identity

$$\int_{0}^{m} \cos^{n} t dt = \frac{1}{n} \cos^{n-1} m \cdot \sin m + \frac{n-1}{n} \int_{0}^{m} \cos^{n-2} t dt.$$

which holds for any  $n \ge 1$ , we have

$$g_d(m) = \sum_{i=0}^{i=d} \frac{1}{4^i} \cdot {\binom{2i}{i}} \cdot m(1-m^2)^i.$$

Therefore, we can compute  $g_n$  as follows

- $g_1(m) = -\frac{1}{2}m^3 + \frac{3}{2}m.$
- $g_2(m) = \frac{3}{8}m^5 \frac{10}{8}m^3 + \frac{15}{8}m.$
- $g_3(m) = -\frac{5}{16}m^7 + \frac{21}{16}m^5 \frac{35}{16}m^3 + \frac{35}{16}m.$   $g_4(m) = \frac{35}{128}m^9 \frac{180}{128}m^7 + \frac{378}{128}m^5 \frac{420}{128}m^3 + \frac{315}{128}m.$

Since 
$$\binom{2i}{i} = 2 \cdot \binom{2i-1}{i-1}$$
 is divisible by 2 for  $i \ge 1$ ,

each coefficient of  $g_d$  can be represented as  $n/2^2$ for  $n \in$  $\mathbb{Z}$ , which can be inferred by simply using Binomial Theorem for the coefficients in  $g_d(m)$ .

Size of Constant  $p_d$ . The constant  $p_d$  is crucial for  $g_{\mathcal{A}}^{(k)}$  to converge to the sign function. Informally, since the coefficient term of m is exactly  $p_d$ , we can regard  $g_d(m)$  as  $g_d(m) \simeq$  $p_d \cdot m$  for small *m*. Further we have  $1 - g_d(m) \simeq 1 - p_d \cdot m \simeq$  $(1-m)^{p_d}$ . For simplicity, we can obtain  $p_d$  as follows:

$$\sum_{i=0}^{i=d} \frac{1}{4^i} \cdot \binom{2i}{i},$$

which can be simplified with Lemma 4.1.

**Lemma 4.1.** It holds that  $p_d = \sum_{i=0}^{i=d} \frac{1}{4^i} \cdot \begin{pmatrix} 2i \\ i \end{pmatrix} = \frac{2d+1}{4^d} \begin{pmatrix} 2d \\ d \end{pmatrix}$ .

Proof. Please refer to Appendix, available in the online supplemental material.

# 4.3 Analysis on the Convergence of $g_d^{(k)}$

We now analyze the convergence of  $g_d^{(k)}$  to the sign function as k increases. To be precise, we provide a lower bound on k, under which  $g_d^{(k)}$  is  $(\sigma, \delta)$ -close to the sign function. To accomplish this, we first give two lower bounds about 1  $g_d(m)$  as shown below.

- **Lemma 4.2.** It holds that  $0 \le 1 g_d(m) \le (1 m)^{p_d}$  for  $m \in [0, 1].$
- **Lemma 4.3.** It holds that  $0 \le 1 g_d(m) \le 2^d \cdot (1-m)^{d+1}$  for  $m \in [0, 1]$ , where the value of m is close to 1.
- Proof. Please refer to Appendix, available in the online supplemental material.
- **Theorem 4.4.** If  $k \ge \frac{1}{\log p_d} \cdot \log(1/\delta) + \frac{1}{\log(d+1)} \cdot \log(\sigma-1) + O(1)$ , then  $g_d^{(k)}(m)$  is an  $(\sigma, \delta)$ -close polynomial to sgn(x) over [-1,1].
- **Proof.** Here we only consider the case where the input of  $g_d^{(k)}$  is non-negative, since  $g_d^{(k)}$  is an odd function. We use

TABLE 2 Complexity of Each Approximation Method

| Parameter  | MinMax Approx. [50]                                | [35]                               | Ours                  |
|--|--|------------------------------------|-----------------------|
| $\log\left(\frac{1}{\delta}\right) = \Theta(1)$      | $\Theta(\sqrt{\sigma})$                            | $\Theta(\log^2 \sigma)$            | $\Theta(\log \sigma)$ |
| $\log\left(\frac{1}{\delta}\right) = \Theta(\sigma)$ | $\Theta(\sqrt{\sigma} \cdot 2^{\frac{\sigma}{2}})$ | $\Theta(\sigma \cdot \log \sigma)$ | $\Theta(\sigma)$      |
| $\log\left(\frac{1}{\delta}\right) = 2^{\sigma}$     | $\Theta(\sqrt{\sigma}\cdot 2^{2^{\sigma-1}})$      | $\Theta(\sigma\cdot 2^\sigma)$     | $\Theta(2^{\sigma})$  |

when  $g_d^{(k)}$  converges to  $(\sigma, \delta)$ -close polynomial to sgn(x). Note that when the value of m is close to 0, Lemma 4.2 is tighter than Lemma 4.3 but vice verse when the value of x is close to 1. To obtain a tight lower bound on k, we decompose the proof into the following two steps, each of which applies Lemmas 4.2 and 4.3, separately.

Step 1. We consider the case  $m \in [\delta, 1]$  instead of  $[-1, -\delta] \cup [\delta, 1]$ , since  $g_d^{(k)}$  is an odd function. Let  $k_{\delta} =$  $\left\lceil \frac{1}{\log(p_d)} \cdot \log\left(\log\left(\frac{1}{\gamma}\right)/\delta\right) \right\rceil$  for some constant  $0 < \gamma < 1$ . Then, with Lemma 4.2, we have the following inequality for  $m \in [\delta, 1]$ .

$$1 - g_d^{k_{\delta}}(m) \quad \leq (1 - m)^{p_d^{k_{\delta}}} \leq (1 - \delta)^{\log\left(\frac{1}{\gamma}/\delta\right)} < \left(\frac{1}{e}\right)^{\log\left(\frac{1}{\gamma}\right)} < \gamma,$$

where *e* indicates the Euler's constant.

Step 2. Let  $k_{\sigma} = \lceil \frac{1}{\log(d+1)} \cdot \log((\sigma - 1/\log(\frac{1}{2\gamma}))) \rceil$ . With Lemma 4.3, we have the following inequality for  $m \in$  $[\delta, 1].$ 

$$2 \cdot (1 - g_d^{(k_{\delta} + k_{\sigma})}(m)) \le (2 \cdot (1 - g_d^{k_{\delta}}(m)))^{(d+1)^{k_{\sigma}}}$$
$$\le (2\gamma)^{(d+1)^{k_{\sigma}}} \le (2\gamma)^{\sigma - 1/\log(\frac{1}{2\gamma})}$$
$$= 2^{-\sigma + 1}.$$

Therefore, 
$$1 - g_d^{(k)}(m) \le 2^{-\sigma}$$
 for  $m \in [\delta, 1]$ ,  
if  $k \ge k_{\delta} + k_{\sigma}$ .

Comparisons With Existing Works. We compare the computation complexity of our method with existing approximation methods for the sign function, including the traditional Minmax based polynomial approximation method [50] and the latest work [35]. The results are shown in Table 2. Paterson et al. [49] have proven that when the input is within the interval [-1, 1], the minimum degree of a  $(\sigma, \delta)$ -polynomial function to approximate a sign function is  $\Theta(\sigma/\delta)$ . This means at least multiplications with the complexity of  $\Theta(\log(1/\delta)) + \Theta(\log \sigma)$  are required to complete the approximation of the sign function. Hence, our method achieves an optimality in asymptotic computation complexity. Other works, like [35] as one of the most advanced solutions for approximating the sign function, only achieve quasi-optimal computation complexity (see Appendix, available in the online supplemental material, for more experimental comparisons).

### 4.4 Application to Max and Relu Functions

Given two variables a and b, the max function can be expressed as  $\max(a,b) = \frac{a+b}{2} + \frac{|a-b|}{2}$ . The ReLu function  $f(x) = \max(0, x)$  can be considered as a special case of the max function. Specifically, since  $|m| = m \cdot sgn(m)$ , as long as we give the approximate polynomial about |m|, we can directly get the approximate max function. Therefore,  $\max(a, b)$  can be evaluated by computing  $\frac{a+b}{2} + \frac{a-b}{2} \cdot g_d^{(k)}(a-b)$ . mial ring. Authorized licensed use limited to: Nanyang Technological University Library. Downloaded on October 06,2023 at 19:53:34 UTC from IEEE Xplore. Restrictions apply.

The detailed algorithm is shown in Algorithm 4. We also provide the convergence rate to approximate the absolute function |m| with  $m \cdot g_d^{(k)}(m)$  (See Theorem 4.5).

**procedure** AppMax (a, b, d, k)1:  $m \leftarrow a - b, y \leftarrow \frac{a+b}{2}$ 2: for k = 1 to k = n - 1 do 3:  $m \leftarrow g_d(m)$ 4: end for 5:  $y \leftarrow y + \frac{a-b}{2} \cdot m$ 6: return y

Algorithm 5. High-Level of Federated Neural Network Training

**Input:**  $\{x, y\} \in D_i \subseteq D$ , for  $i \in \{1, \ldots, N\}$ **Output:** Encrypted  $\omega_1^H, \omega_2^H, \ldots, \omega_1^H$ 

Prepare:

- 1: The cloud server C and every user  $P_i$  agree on the parameters  $\mathbb{L}$ ,  $h_1, \ldots, h_{\mathbb{L}}$ ,  $\eta$ ,  $\varphi(\cdot)$ , H and  $\mathcal{B}$ . The cloud server  $\mathcal{C}$  generates its secret key and public key  $\{sk', pk'\} \leftarrow$ SecKeyGen $(1^{\lambda})$ .
- 2: Each user  $P_i$  generates  $sk_i \leftarrow$ **SecKeyGen** $(1^{\lambda})$ .
- 3: All users collectively generate  $pk \leftarrow \mathbf{DKeyGen}(\{sk_i\})$ .
- 4: Each user encodes its input as  $\hat{X}_i$ ,  $\hat{Y}_i$ .<sup>4</sup>
- 5: The cloud server C initializes  $[\boldsymbol{\omega}_1^0]_{pk}, [\boldsymbol{\omega}_2^0]_{pk}, \dots, [\boldsymbol{\omega}_{\mathbb{L}}^0]_{pk}$ . Then, *C* broadcasts them to all users. Local Training:
- 6: **for** k = 0 to k = H 1 **do**
- Each user  $P_i$  computes  $[\nabla \omega_{1,i}^k]_{pk}, \ldots, [\nabla \omega_{L,i}^k]_{pk}$  and sends 7: them to the cloud server. Aggregation:
- 8: for j = 1 to j = L do
- 9:
- $\begin{array}{l} \mathcal{C} \text{ computes } [\nabla \omega_{\mathbf{j}}^{\mathbf{k}}]_{pk} = [\sum_{i=1}^{N} \nabla \omega_{\mathbf{j}}^{\mathbf{k}}]_{pk}. \\ \mathcal{C} \text{ computes } [\omega_{\mathbf{j}}^{\mathbf{k}+1}]_{pk} = [\omega_{\mathbf{j}}^{\mathbf{k}} \frac{\eta}{\mathcal{B} \times N} \nabla \omega_{\mathbf{j}}^{\mathbf{k}}]_{pk} \text{ and broadcasts} \end{array}$ 10: them to all users.
- 11: end for

- 12: end for
- **Theorem 4.5.** If  $k \ge \frac{1}{\log p_d} \cdot \log (\sigma 1)$ , then the error of  $m \cdot$  $g_d^{(k)}(m)$  compared with |m| over [-1,1] is bounded by  $2^{\sigma}$ .

**Proof.** This proof can be easily evolved from Theorem 4.4. We omit it for brevity. Π

#### **IMPLEMENTATION OF HERCULES** 5

We now describe the technical details of implementing Hercules, which provides privacy-preserving federated neural network training. In particular, model parameters and users' data are encrypted throughout the execution process. To achieve this, Hercules exploits the MCKKS as the underlying framework and relies on the packed ciphertext technology to accelerate calculations. Besides, approximation methods based on composite polynomials are used to

<sup>4.</sup>  $\hat{X}_i$  and  $\hat{Y}_i$  can be vectors composed of a single training sample, or a matrix composed of multiple samples. This depends on the size of a single sample and the value of the degree N of the cyclotomic polyno-

approximate ReLU and max functions, which facilitate the compatibility of HE with complex operations.

From a high-level view, the implementation of Hercules is composed of three phases: Prepare, Local Training, and Aggregation. As shown in *Algorithm* 5, we use  $[\cdot]_{nk}$  to denote the encrypted value and  $\omega_{i,i}^k$  to represent the weight matrix of the *j*-th layer generated by  $P_i$  at the *k*-th iteration. The global weight matrix is denoted as  $\omega_i^k$  without index *i*. Similarly, the local gradients computed by user  $P_i$  for each layer *j* at the *k*-th iteration is denoted as  $\nabla \omega_{j,i}^k$ .

1. Prepare: The cloud server C needs to agree with all users on the training hyperparameters, including the number  $\mathbb{L}$  of layers in the model, the number  $h_i$  of neurons in each hidden layer  $j, j \in [L]$ , the learning rate  $\eta$ , the number *H* of global iterations, the number  $\mathcal{B}$  of local batches, the activation function  $\varphi(\cdot)$  and its approximation. Then, C generates its own key pair  $\{sk, pk'\}$ , and each user  $P_i$  generates  $sk_i$  for  $i \in [N]$ . Besides, all users collectively generate pk. Finally, C initializes  $[\boldsymbol{\omega}_1^0]_{vk}, [\boldsymbol{\omega}_2^0]_{vk}, \dots, [\boldsymbol{\omega}_L^0]_{vk}$ , and broadcasts them to all users.

2. Local Training: Each user  $P_i$  executes the mini-batch based SGD algorithm locally and obtains the encrypted local gradients  $[\nabla \omega_{1,i}^{\mathbf{k}}]_{pk}, \ldots, [\nabla \omega_{L,i}^{\mathbf{k}}]_{pk}$ , where  $P_i$  is required to execute the forward and backward passes for  $\mathcal{B}$ times to compute and aggregate the local gradients. Then,  $P_i$  sends these local gradients to the cloud server C.

3. Aggregation: After receiving all the local gradients from users, C updates the global model parameters by computing the averaged aggregated gradients. In our system, training is stopped once the number of iterations reaches *H*. Therefore, after the last iteration, all users need to perform an additional ciphertext conversion operation, i.e., the **DKeySwitch** function (shown in Fig. 1), which enables to convert model **M** encrypted under the public key pk into  $[\mathbf{M}]_{pk'}$  under the cloud server's public key pk' without decryption, so that C can access the final model parameters.

Appendix, available in the online supplemental material, presents the details of Hercules implementation, which essentially executes Algorithm 1 under the ciphertext. This helps readers understand how the functions in MCKKS as well as our new matrix parallel multiplication technology are used in FL.

Security of Hercules: We demonstrate that Hercules realizes the data and model privacy protection defined in Section 2.3, even under the collusion of up to N - 1 users. This is inherited from the property of MCKKS [1]. We give the following Theorem 5.1 and provide the security proof (sketch). The core of our proof is that for any adversary, when only the input and output of passive malicious users in Hercules are provided, there exists a simulator with Probabilistic Polynomial Time computation ability, which can simulate the view of the adversary and make the adversary unable to distinguish the real view from the simulated one.

# Theorem 5.1. Hercules realizes the privacy protection of data and model parameters during the FL process, as long as its underlying MCKKS cryptosystem is secure.

**Proof.** Hercules inherits the security attributes of the MCKKS cryptosystem proposed by Mouchet et al. [1]. Compared with the standard CKKS, the multiparty version constructs additional distributed cryptographic functions trained on the corresponding datasets (See Appendices, Authorized licensed use limited to: Nanyang Technological University Library. Downloaded on October 06,2023 at 19:53:34 UTC from IEEE Xplore. Restrictions apply.

including  $DKeyGen(\cdot)$ ,  $DDec(\cdot)$ ,  $DKeySwitch(\cdot)$  and **DBootstrap** $(\cdot)$ . All of them have been proven secure against a passive-adversary model with up to N-1 colluding parties, under the assumption of the underlying NP hard problem (i.e., RLWE problem [51]). Here we give a sketch of the proof with the simulation paradigm of the real/ideal world. Let us assume that a real-world simulator S simulates a computationally bounded adversary composed of N-1 users colluding with each other. Therefore, S can access all the inputs and outputs of these N-1 users. As mentioned earlier, the MCKKS guarantees the indistinguishability of plaintext under chosen plaintext attacks (i.e., CPA-Secure) even if collusion of N - 1 users. This stems from the fact that the secret key used for encryption must be recovered with the participation of all users. Therefore, S can simulate the data sent by honest users by replacing the original plaintext with random messages. Then these random messages are encrypted and sent to the corresponding adversary. Due to the security of CKKS, the simulated view is indistinguishable from the real view to the adversary. Analogously, the same argument proves that Hercules protects the privacy of the training model, because all model parameters are encrypted with CKKS, and the intermediate and final weights are always in ciphertext during the training process. 

#### **PERFORMANCE EVALUATION** 6

We experimentally evaluate the performance of Hercules in terms of classification accuracy, computation communication and storage overhead. We compare Hercules with POSEIDON [12], which is consistent with our scenario and is also bulit on MCCKS.

#### Experimental Configurations 6.1

We implement the multi-party cryptography operations on the Lattigo lattice-based library [52], which provides an optimized version of the MCKKS cryptosystem. All the experiments are performed on 10 Linux servers, each of which is equipped with Intel Xeon E5-2699v3 CPUs, 36 threads on 18 cores and 189 GB RAM. We make use of Onet [53] and build a distributed system where the parties communicate over TCP with secure channels (TLS). We instantiate **Hercules** with the number of users as N = 10and N = 50, respectively. For parameter settings, the dimension of the cyclotomic polynomial ring in CKKS is set as  $\mathcal{N} = 2^{13}$  for the datasets with the dimension of input h < 132 or  $32 \times 32$  images, and  $2^{14}$  for those with inputs h > 32. The number of initial levels  $\mathcal{L} = 6$ . We exploit  $g_4(m)$ described in Section 4.2 as the basic of compound polynomial to approximate the ReLU and max functions, where we require  $\sigma = 20$ ,  $\delta = 2^{-20}$ . For other continuous activation functions, such as sigmoid, we use the traditional MinMax strategy to approximate it, since it has been proven that a small degree polynomial can fit a non-polynomial continuous function well within a small bounded error.

Consistent with POSEIDON [12], we choose 7 public datasets (i.e., BCW [54], MNIST [55], ESR [56], CREDIT [57], SVHN [58], CIFAR-10, and CIFAR-100 [59]) in our experiments, and design 5 different neural network architectures

TABLE 3 Model Accuracy and Training Cost With  ${\cal N}=10$  Users

| Dataset |             | Training time (s) |          |          |        | Communication cost<br>(GB) |        |         |          |          |
|---------|-------------|-------------------|----------|----------|--------|----------------------------|--------|---------|----------|----------|
|         | Distributed | Local             | POSEIDON | Hercules | POSE   | IDON                       | Here   | cules   | POSEIDON | Hercules |
|         |             |                   |          |          | One-GI | Total                      | One-GI | Total   |          |          |
| BCW     | 97.9%       | 93.7%             | 96.2%    | 97.7%    | 0.40   | 39.92                      | 0.11   | 11.09   | 0.59     | 0.59     |
| ESR     | 93.8%       | 90.2%             | 90.3%    | 93.5%    | 0.92   | 553.44                     | 0.29   | 172.95  | 562.51   | 3.52     |
| CREDIT  | 81.7%       | 79.8%             | 80.3%    | 81.4%    | 0.33   | 163.07                     | 0.13   | 62.73   | 7.32     | 2.93     |
| MNIST   | 92.3%       | 87.7%             | 88.4%    | 91.9%    | 44.67  | 4467.25                    | 1.54   | 1540.43 | 703.13   | 17.58    |

available in the online supplemental material, for more details of the datasets and models used in our experiments. Note that we train two models, CIFAR-10-N1 and CIFAR-10-N2, over the CIFAR-10 dataset for comparison).

# 6.2 Model Accuracy

We first discuss the model accuracy on different datasets when the number of users is 10 and 50 respectively. We choose the following three baselines for comparison. (1) Distributed: distributed training in plaintext, which is in the plaintext form corresponding to **Hercules**. The datasets are evenly distributed to all users to perform FL in a plaintext environment. (2) Local: local training in plaintext, i.e., each user only trains the model on the local dataset. (3) POSEI-DON [12]. We reproduce the exact algorithm designed in [12].

All the baselines are trained on the same network architecture and learning hyperparameters. The learning rate is adaptive to different schemes to obtain the best training accuracy.<sup>5</sup>

As shown in Tables 3 and 4, we can obtain the following two observations. (1) Compared with local training, FL improves the accuracy of model training, especially with the participation of large-scale users. This is drawn from the comparison between the second and fourth columns of Table 4. The reason is obvious: the participation of largescale users has enriched the volume of training samples, and a more accurate model can be derived from such a fertile composite dataset. (2) Compared with distributed training, Hercules has negligible loss in accuracy (less than (0.3%) and is obviously better than POSEIDON (1% to 4%) improvement). In POSEIDON, the non-continuous activation function (i.e, ReLU) is converted into a low-degree polynomial using a traditional approximation method based on the least square method. This is computationally efficient but inevitably brings a non-negligible precision loss. However, given a small error bound, our approximation based on the composite polynomial can approximate non-continuous functions with high-degree polynomials, but only requires the computation complexity of  $O(\log(degG))$ , where degG is the degree of the composite polynomial. Therefore, the accuracy loss caused by the conversion of the activation function is very slight in Hercules.

Note that the model accuracy can be further improved by increasing the number of iterations, but we use the same

5. For example, approximating the activation function at a small interval usually requires a small learning rate to avoid divergence. Authorized licensed use limited to: Nanyang Technological University Library, Dowr number of iterations for the convenience of comparison. To achieve the expected training accuracy, model training over CIFAR-100 usually requires a special network architecture (such as ResNet) and layers (batch normalization) due to the diversity of its labels. For the training simplicity, we choose a relatively simple network architecture, which is also the main reason for the relatively low training accuracy under CIFAR-10 and CIFAR-100 [60], [61]. We leave the model training of more complex architectures [62], [63], [64] and tasks as future work (See Appendix, available in the online supplemental material).

#### 6.3 Computation Overhead

We further discuss the performance of **Hercules** in terms of computation overhead. As shown in Tables 3 and 4, when the number of users is 10, the training time of Hercules over BCW, ESR and CREDIT is less than 3 minutes, and the training time over MNIST is also less than 30 minutes. For N =50, to train specific model architectures over SVHN, CIFAR-10-N1, CIFAR-10-N2 and CIFAR-100, the total cost of Hercules is 8.78 hours, 40.73 hours, 33.3 hours and 126.52 hours, respectively. We also give the running time of one global iteration (One-GI), which can be used to estimate the time required to train these architectures under a larger number of global iterations. Obviously, for the same model architecture and number of iterations, the execution time of Hercules is far less than that of POSEIDON. This stems from the fast SIMD operation under our new matrix multiplication coding method (See Appendix, available in the online supplemental material) for the comparison of the microbenchmark costs of Hercules and POSEIDON under various functionalities). Specifically, POSEIDON designs AP to achieve fast SIMD calculations. AP combines row-based and column-based packing, which means that the rows or columns of the matrix are vectorized and packed into a ciphertext. For the multiplication of two  $(h \times h)$ -dimensional matrices, the complexity of the homomorphic rotation operations required by AP is  $\max_{i \in [L]} (\omega_i \times \log (h \times \omega_i))$ , where  $\omega_i$  denotes the number of weights between layers *i* and i + 1. For example, given h = 64,  $\max_{i \in [L]} \omega_i = 64$ , AP roughly needs 768 homomorphic rotation operations to realize the multiplication calculation of two  $(64 \times 64)$ -dimensional matrices. For Hercules, as shown in Table 1, the complexity required for the matrix multiplication is only  $3 \times 64 + 5\sqrt{64} = 232$ , which is roughly one third of the overhead required by POSEIDON. Moreover, by comparing the complexity, we can infer that the homomorphic multiplication of the matrices in Hercules is only linearly related to

Authorized licensed use limited to: Nanyang Technological University Library. Downloaded on October 06,2023 at 19:53:34 UTC from IEEE Xplore. Restrictions apply.

| IABLE 4                               |              |
|---------------------------------------|--------------|
| Model Accuracy and Training Cost With | N = 50 Users |

| Detect      | Accuracy    |       |          |                | Training time (hrs) |        |          |        | Communication cost<br>(GB) |          |
|-------------|-------------|-------|----------|----------------|---------------------|--------|----------|--------|----------------------------|----------|
| Dataset     | Distributed | Local | POSEIDON | EIDON Hercules | POSEIDON            |        | Hercules |        | POSEIDON                   | Hercules |
|             |             |       |          |                | One-GI              | Total  | One-GI   | Total  |                            |          |
| SVHN        | 68.5%       | 35.7% | 67.7%    | 68.4%          | 0.0013              | 24.15  | 0.0005   | 8.78   | 12656.25                   | 474.61   |
| CIFAR-10-N1 | 54.9%       | 26.7% | 51.9%    | 54.4%          | 0.005               | 126.26 | 0.0016   | 40.73  | 61523.44                   | 2050.78  |
| CIFAR-10-N2 | 63.5%       | 28.4% | 60.4%    | 63.3%          | 0.0059              | 98.32  | 0.002    | 33.33  | 59062.5                    | 1968.75  |
| CIFAR-100   | 45.6%       | 28.2% | 41.1%    | 44.2%          | 0.0069              | 363.11 | 0.0024   | 126.52 | 246796.88                  | 8226.56  |

Note that **Hercules** and POSEIDON produce a relatively high total communication overhead compared to Table 3, which stems from the use of a larger number of global iterations over the above datasets (See Appendix, available in the online supplemental material, for hyperparameter settings).



Fig. 3. Running time of one training epoch. (a) Increase the number of users N when the number of samples for each user is  $|D_i| = 200$ . (b) Increase the number of users N when the total sample size is |D| = 2000. (c) Increase the total sample size |D| when N = 10. (d) Increase the sample dimension when N = 10 and  $|D| = 200 \times N$ .

the dimension of the matrix, and is independent of the number of neurons in each layer of the model. On the contrary, the complexity of AP increases linearly with  $\max_{i \in [L]} \omega_i$ . This implies that **Hercules** is more suitable for complex network architectures than POSEIDON.

We further analyze the scalability of Hercules and POSEIDON in terms of the number of users N, the number of samples |D|, and the number of dimensions h for one sample. Here we use a two-layer architecture with 64 neurons in each layer. The local batch size for each user is 10. Fig. 3 shows the experimental results, where we record the execution time of one training epoch, i.e., all the data of each user are processed once. Specifically, Fig. 3a shows the execution time as the number of users grows, where we fix the number of samples held by each user as 200, and the dimension of each sample as 64. We can observe that the execution time of Hercules and POSEIDON shows a slight linear increase with the increase of the number of users. This stems from the fact that most of the operations performed by each user are concentrated locally except for the distributed bootstrapping procedure. Obviously, the percentage of DBootstrap operations over the total operations under ciphertext training is relatively small. We further fix the total number of samples in the system as 2000, and calculate the execution time of each user as the number of users increases. As shown in Fig. 3b, this causes a linear decrease in execution time since the increase in user data reduces the sample volume held by each user. Given the fixed number of users N = 10 and h = 64, Fig. 3c shows that the execution time of each user increases linearly as |D|increases. It is obvious that the increase in |D| implies an increase in the number of samples in each user. Fig. 3d also shows similar results under different sample dimensions.

In general, **Hercules** and POSEIDON show similar relationships in terms of computation cost under different hyperparameters. However, we can observe that the running time of **Hercules** is far less than that of POSEIDON, due to the superiority of our new matrix multiplication method.

# 6.4 Communication Overhead

Tables 3 and 4 show the total communication overhead required by Hercules and POSEIDON over different datasets. We can observe that during the training process, the ciphertext data that each user needs to exchange with other parties in Hercules are much smaller than that of POSEI-DON. This also stems from the superiority of the new matrix multiplication method we design. Specifically, In POSEIDON, AP is used for matrix multiplication to achieve fast SIMD operations. However, as shown in the fourth row of Protocol 3 in [12], this method requires multiple copies and zero padding operations for each row or column of the input matrix, depending on the number of neurons in each hidden layer, the absolute value of the difference between the row or column dimension of the matrix and the number of neurons in the corresponding hidden layer. In fact, AP is an encoding method that trades redundancy in storage for computation acceleration. On the contrary, our method does not require additional element copy except for a small amount of zero padding in the initial stage to facilitate calculations. Therefore, Hercules obviously exhibits smaller communication overhead. For example, given the MNIST dataset, a 3-layer fully connected model with 64 neurons per layer, the communication overhead of each user in POSEIDON is about 703(GB) to complete 1000 global iterations, while Hercules only needs 17.58(GB) per user.



Fig. 4. Evaluation of communication overhead for one training epoch. (a) Increase the number of users N when given  $|D_i| = 200$  for each user i. (b) Increase the number of users N when the total sample size |D| = 2000. (c) Increase the total sample size |D| when given N = 10. (d) Increase the dimension of a single sample when given N = 10 and  $|D| = 200 \times N$ .

We also analyze the scalability of **Hercules** and POSEI-DON in terms of the number of users N, the number of samples |D|, and the sample dimension h. Here we use a twolayer model architecture with 64 neurons in each layer. The local batch size for each user is 10. Fig. 4 shows the experimental results. Similar to the results for computation cost comparison, we can observe that **Hercules** exhibits better scalability compared to POSEIDON under different hyperparameters. In addition, we also show the storage overhead advantage of **Hercules** compared to POSEIDON, and discuss the performance of **Hercules** compared with other advanced MPC-based solutions. More details can be found in Appendix, available in the online supplemental material.

# 7 CONCLUSION

In this paper, we propose **Hercules** for privacy-preserving FL. We design a novel matrix coding technique to accelerate the training performance under ciphertext. Then, we use a novel approximation strategy to improve the compatibility of **Hercules** for processing non-polynomial functions. Experiments on benchmark datasets demonstrate the superiority of **Hercules** compared with existing works. In the future, we will focus on designing more efficient optimization strategies to further reduce the computation overhead of **Hercules**, to make it more suitable for practical applications.

## ACKNOWLEDGMENTS

We thank the anonymous reviewers for their insightful comments!

### REFERENCES

- C. Mouchet, J. Troncoso-Pastoriza, J.-P. Bossuat, and J.-P. Hubaux, "Multiparty homomorphic encryption from ring-learning-witherrors," Cryptology ePrint Arch., 2020.
- H. Chen, W. Dai, M. Kim, and Y. Song, "Efficient multi-key homomorphic encryption with packed ciphertexts with application to oblivious neural network inference," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2019, pp. 395–412.
- [3] J. Geiping, H. Bauermeister, H. Dröge, and M. Moeller, "Inverting gradients-how easy is it to break privacy in federated learning?," 2020, arXiv:2003.14053.
- [4] L. Zhu and S. Han, "Deep leakage from gradients," in *Federated Learning*. Berlin, Germany: Springer, 2020, pp. 17–31.
- [5] D. Chen, N. Yu, Y. Zhang, and M. Fritz, "GAN-Leaks: A taxonomy of membership inference attacks against generative models," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2020, pp. 343–362.

- [6] M. Abadi et al., "Deep learning with differential privacy," in Proc. ACM SIGSAC Conf. Comput. Commun. Secur., 2016, pp. 308–318.
- [7] L. Yu, L. Liu, C. Pu, M. E. Gursoy, and S. Truex, "Differentially private model publishing for deep learning," in *Proc. IEEE Symp. Secur. Privacy*, 2019, pp. 332–349.
- [8] P. Mohassel and P. Rindal, "ABY<sup>3</sup>: A mixed protocol framework for machine learning," in Proc. ACM SIGSAC Conf. Comput. Commun. Secur., 2018, pp. 35–52.
  [9] N. Agrawal, A. S. Shamsabadi, M. J. Kusner, and A. Gascón,
- [9] N. Agrawal, A. S. Shamsabadi, M. J. Kusner, and A. Gascón, "QUOTIENT: Two-party secure neural network training and prediction," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2019, pp. 1231–1247.
- [10] D. Rathee et al., "CrypTFlow2: Practical 2-party secure inference," in Proc. ACM SIGSAC Conf. Comput. Commun. Secur., 2020, pp. 325–342.
- [11] W. Zheng, R. A. Popa, J. E. Gonzalez, and I. Stoica, "Helen: Maliciously secure coopetitive learning for linear models," in *Proc. IEEE Symp. Secur. Privacy*, 2019, pp. 724–738.
- [12] S. Sav et al., "POSEIDON: Privacy-preserving federated neural network learning," in Proc. Netw. Distrib. Syst. Secur. Symp., 2021, pp. 1–16.
- pp. 1–16.
  [13] K. Wei et al., "Federated learning with differential privacy: Algorithms and performance analysis," *IEEE Trans. Inf. Forensics Secur.*, vol. 15, pp. 3454–3469, 2020.
- [14] N. Wu, F. Farokhi, D. Smith, and M. A. Kaafar, "The value of collaboration in convex machine learning with differential privacy," in *Proc. IEEE Symp. Secur. Privacy*, 2020, pp. 304–317.
- [15] E. Bagdasaryan, O. Poursaeed, and V. Shmatikov, "Differential privacy has disparate impact on model accuracy," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2019, pp. 15479–15488.
- [16] D. Wang and J. Xu, "On sparse linear regression in the local differential privacy model," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 6628–6637.
- [17] B. Jayaraman and D. Evans, "Evaluating differentially private machine learning in practice," in *Proc. USENIX Secur. Symp.*, 2019, pp. 1895–1912.
- [18] B. Hitaj, G. Ateniese, and F. Perez-Cruz, "Deep models under the GAN: Information leakage from collaborative deep learning," in Proc. ACM SIGSAC Conf. Comput. Commun. Secur., 2017, pp. 603–618.
- [19] A. Patra, T. Schneider, A. Suresh, and H. Yalame, "ABY2.0: Improved mixed-protocol secure two-party computation," in *Proc. USENIX Secur. Symp.*, 2021, pp. 2165–2182.
- [20] P. Mohassel and Y. Zhang, "SecureML: A system for scalable privacy-preserving machine learning," in *Proc. IEEE Symp. Secur. Pri*vacy, 2017, pp. 19–38.
- [21] J. Liu, M. Juuti, Y. Lu, and N. Asokan, "Oblivious neural network predictions via MiniONN transformations," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2017, pp. 619–631.
- [22] N. Kumar, M. Rathee, N. Chandran, D. Gupta, A. Rastogi, and R. Sharma, "CryptFlow: Secure tensorflow inference," in *Proc. IEEE Symp. Secur. Privacy*, 2020, pp. 336–353.
- [23] A. Patra and A. Suresh, "BLAZE: Blazing fast privacy-preserving machine learning," in *Proc. Netw. Distrib. Syst. Secur. Symp.*, 2020.
  [24] R. Rachuri and A. Suresh, "Trident: Efficient 4PC framework for
- [24] R. Rachuri and A. Suresh, "Trident: Efficient 4PC framework for privacy preserving machine learning," in Proc. Netw. Distrib. Syst. Secur. Symp., 2020.
- [25] D. Froelicher et al., "Scalable privacy-preserving distributed learning," in Proc. Privacy Enhancing Technol. Symp., 2020, pp. 323–347.
- [26] S. D. Galbraith, "Elliptic curve Paillier schemes," J. Cryptol., vol. 15, no. 2, pp. 129–138, 2002.

- [27] J. H. Cheon, A. Kim, M. Kim, and Y. Song, "Homomorphic encryption for arithmetic of approximate numbers," in *Proc. Int. Conf. Theory Appl. Cryptol. Inf. Secur.*, 2017, pp. 409–437.
- [28] C. Źhang, Š. Li, J. Xia, W. Wang, F. Yan, and Y. Liu, "BatchCrypt: Efficient homomorphic encryption for cross-silo federated learning," in Proc. USENIX Annu. Tech. Conf., 2020, pp. 493–506.
- [29] Q. Zhang, C. Xin, and H. Wu, "GALA: Greedy computation for linear algebra in privacy-preserved neural networks," in *Proc. Netw. Distrib. Syst. Secur. Symp.*, 2021, pp. 1–19.
- Netw. Distrib. Syst. Secur. Symp., 2021, pp. 1–19.
  [30] I. Iliashenko and V. Zucca, "Faster homomorphic comparison operations for BGV and BFV," in *Proc. Privacy Enhancing Technol. Symp.*, 2021, pp. 246–264.
- [31] J. H. Cheon, D. Kim, and D. Kim, "Efficient homomorphic comparison methods with optimal complexity," in *Proc. Int. Conf. The*ory Appl. Cryptol. Inf. Secur., 2020, pp. 221–256.
- [32] W.-J. Lu, Z. Huang, C. Hong, Y. Ma, and H. Qu, "PEGASUS: Bridging polynomial and non-polynomial evaluations in homomorphic encryption," in *Proc. IEEE Symp. Secur. Privacy*, 2021, pp. 1057–1073.
- [33] L. Ducas and D. Micciancio, "FHEW: Bootstrapping homomorphic encryption in less than a second," in *Proc. Annu. Int. Conf. Theory Appl. Cryptographic Techn.*, 2015, pp. 617–640.
- [34] S. Halevi and V. Shoup, "Bootstrapping for HElib," in Proc. Annu. Int. Conf. Theory Appl. Cryptographic Techn., 2015, pp. 641–670.
- [35] J. H. Cheon, D. Kim, D. Kim, H. H. Lee, and K. Lee, "Numerical method for comparison on homomorphically encrypted numbers," in Proc. Int. Conf. Theory Appl. Cryptol. Inf. Secur., 2019, pp. 415–445.
- [36] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership inference attacks against machine learning models," in *Proc. IEEE Symp. Secur. Privacy*, 2017, pp. 3–18.
- [37] Y. Zhang, R. Jia, H. Pei, W. Wang, B. Li, and D. Song, "The secret revealer: Generative model-inversion attacks against deep neural networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 253–261.
- [38] M. Juuti, S. Szyller, S. Marchal, and N. Asokan, "PRADA: Protecting against DNN model stealing attacks," in *Proc. IEEE Eur. Symp. Secur. Privacy*, 2019, pp. 512–527.
- [39] J. Jia, A. Salem, M. Backes, Y. Zhang, and N. Z. Gong, "MemGuard: Defending against black-box membership inference attacks via adversarial examples," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2019, pp. 259–274.
- [40] N. Koti, M. Pancholi, A. Patra, and A. Suresh, "SWIFT: Super-fast and robust privacy-preserving machine learning," in *Proc. USE-NIX Secur. Symp.*, 2021, pp. 2651–2668.
- [41] M. S. Riazi, M. Samragh, H. Chen, K. Laine, K. Lauter, and F. Koushanfar, "XONN: XNOR-based oblivious deep neural network inference," in *Proc. USENIX Secur. Symp.*, 2019, pp. 1501–1518.
- [42] R. Canetti, A. Jain, and A. Scafuro, "Practical UC security with a global random oracle," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, 2014, pp. 597–608.
- [43] C. Weng, K. Yang, X. Xie, J. Katz, and X. Wang, "Mystique: Efficient conversions for zero-knowledge proofs with applications to machine learning," in *Proc. USENIX Secur. Symp.*, 2021, pp. 501– 518.
- [44] M. Keller, "MP-SPDZ: A versatile framework for multi-party computation," in Proc. ACM SIGSAC Conf. Comput. Commun. Secur., 2020, pp. 1575–1590.
- [45] S. Halevi and V. Shoup, "Algorithms in HElib," in Proc. Annu. Cryptol. Conf., 2014, pp. 554–571.
- [46] J. S. Coron, D. Lefranc, and G. Poupard, "A new baby-step giantstep algorithm and some applications to cryptanalysis," in *Proc. Int. Workshop Cryptographic Hardware Embedded Syst.*, 2005, pp. 47– 60.
- [47] X. Jiang, M. Kim, K. Lauter, and Y. Song, "Secure outsourced matrix computation and application to neural networks," in *Proc.* ACM SIGSAC Conf. Comput. Commun. Secur., 2018, pp. 1209–1222.
- [48] 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 *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 201–210.
- [49] M. S. Paterson and L. J. Stockmeyer, "On the number of nonscalar multiplications necessary to evaluate polynomials," SIAM J. Comput., vol. 2, no. 1, pp. 60–66, 1973.
- [50] A. Eremenko and P. Yuditskii, "Uniform approximation of SGN X by polynomials and entire functions," J. d'Analyse Mathématique, vol. 101, no. 1, pp. 313–324, 2007.

- [51] M. Rosca, D. Stehlé, and A. Wallet, "On the ring-LWE and polynomial-LWE problems," in *Proc. Annu. Int. Conf. Theory Appl. Cryptographic Techn.*, 2018, pp. 146–173.
  [52] Lattigo v2.1.1, ePFL-LDS, Dec. 2020. [Online]. Available: http://
- [52] Lattigo v2.1.1, ePFL-LDS, Dec. 2020. [Online]. Available: http:// github.com/ldsec/lattigo
- [53] Cothority network library, ePFL-EFPL, Dec. 2021. [Online]. Available: https://github.com/dedis/onet
- [54] D. Lavanya and D. K. U. Rani, "Analysis of feature selection with classification: Breast cancer datasets," *Indian J. Comput. Sci. Eng.*, vol. 2, no. 5, pp. 756–763, 2011.
- [55] L. Deng, "The MNIST database of handwritten digit images for machine learning research [best of the web]," *IEEE Signal Process. Mag.*, vol. 29, no. 6, pp. 141–142, Nov. 2012.
- [56] Epileptic seizure recognition dataset, 2007. [Online]. Available: https://archive.ics.uci.edu/ml/datasets/EpilepticSeizureRecognition
- [57] The default of credit card clients, 2016. [Online]. Available: https://archive.ics.uci.edu/ml/datasets/default+of+credit+card +clients
- [58] The street view house numbers (SVHN) dataset, 2011. [Online]. Available: http://ufldl.stanford.edu/housenumbers/
- [59] CIFAR-10 and CIFAR-100 dataset, 2009. [Online]. Available: https://www.cs.toronto.edu/kriz/cifar.html
- [60] H. Tang, X. Lian, M. Yan, C. Zhang, and J. Liu, "d<sup>2</sup>: Decentralized training over decentralized data," in Proc. Int. Conf. Mach. Learn., 2018, pp. 4848–4856.
- [61] H. Tang, C. Yu, X. Lian, T. Zhang, and J. Liu, "DoubleSqueeze: Parallel stochastic gradient descent with double-pass error-compensated compression," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 6155–6165.
  [62] I. M. Baytas, C. Xiao, X. Zhang, F. Wang, A. K. Jain, and J. Zhou,
- [62] I. M. Baytas, C. Xiao, X. Zhang, F. Wang, A. K. Jain, and J. Zhou, "Patient subtyping via time-aware LSTM networks," in *Proc. ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2017, pp. 65–74.
- [63] S. Li, W. Li, C. Cook, C. Zhu, and Y. Gao, "Independently recurrent neural network (IndRNN): Building a longer and deeper RNN," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 5457–5466.
  [64] R. Fablet, S. Ouala, and C. Herzet, "Bilinear residual neural net-
- [64] Ř. Fablet, S. Ouala, and C. Herzet, "Bilinear residual neural network for the identification and forecasting of dynamical systems," 2017, arXiv:1712.07003.



**Guowen Xu** received the PhD degree from the University of Electronic Science and Technology of China, in 2020. He is currently a research fellow with Nanyang Technological University, Singapore. He has published papers in reputable conferences/ journals, including ACM CCS, NeurIPS, ASIACCS, ACSAC, ESORICS, *IEEE Transactions on Information Forensics and Security*, and *IEEE Transactions on Dependable and Secure Computing*. His research interests include applied cryptography and privacy-preserving deep learning.



Xingshuo Han is currently working toward the PhD degree with the School of Computer Science and Engineering, Nanyang Technological University, Singapore. He has published papers in reputable conferences/journals, including ACM MM, *IEEE Transactions on Intelligent Transportation Systems* and *IEEE Transactions on Dependable and Secure Computing*. His research interests include safety and privacy of deep learning, safety and security of autonomous vehicles, and intelligent transportation systems.



Shengmin Xu is currently an associate professor with the Fujian Provincial Key Laboratory of Network Security and Cryptology, College of Computer and Cyber Security, Fujian Normal University, Fuzhou, China. Previously, he was a senior research engineer with the School of Computing and Information Systems, Singapore Management University. His research interests include cryptography and information security.



**Tianwei Zhang** received the bachelor's degree from Peking University, in 2011, and the PhD degree from Princeton University, in 2017. He is an assistant professor with the School of Computer Science and Engineering, Nanyang Technological University. His research focuses on computer system security. He is particularly interested in security threats and defenses in machine learning systems, autonomous systems, computer architecture, and distributed systems.



Hongwei Li (Senior Member, IEEE) is currently the head and a professor with the Department of Information Security, School of Computer Science and Engineering, University of Electronic Science and Technology of China. His research interests include network security and applied cryptography. He is the distinguished lecturer of IEEE Vehicular Technology Society.



Xinyi Huang is currently an associate professor with the Thrust of Artificial Intelligence, Information Hub, Hong Kong University of Science and Technology (Guangzhou), China. His research interests include cryptography and information security. He is in the editorial board of the International Journal of Information Security and Science China Information Sciences. He has served as the program/general chair or program committee member in more than 120 international conferences.



**Robert H. Deng** (Fellow, IEEE) is AXA chair professor of Cybersecurity, Singapore Management University. His research interests include the areas of data security and privacy, cloud security, and IoT security. He served on many editorial boards and conference committees, including the editorial boards of the *IEEE Security & Privacy Magazine*, *IEEE Transactions on Dependable and Secure Computing*, *IEEE Transactions on Information Forensics and Security*, and Steering Committee chair of the ACM Asia Conference on Computer and Communications Security.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.