

DeTrust-FL: Privacy-Preserving Federated Learning in Decentralized Trust Setting

Runhua Xu, Nathalie Baracaldo, Yi Zhou, Ali Anwar, Swanand Kadhe, and Heiko Ludwig

IBM Research Almaden, San Jose, CA, USA

{runhua, yi.zhou, ali.anwar2, Swanand.Kadhe}@ibm.com, {baracald, hludwig}@us.ibm.com

Abstract—Federated learning has emerged as a privacy-preserving machine learning approach where multiple parties can train a single model without sharing their raw training data. Federated learning typically requires the utilization of multi-party computation techniques to provide strong privacy guarantees by ensuring that an untrusted or curious aggregator cannot obtain *isolated* replies from parties involved in the training process, thereby preventing potential inference attacks. Until recently, it was thought that some of these secure aggregation techniques were sufficient to fully protect against inference attacks coming from a curious aggregator. However, recent research has demonstrated that a curious aggregator can successfully launch a *disaggregation attack* to learn information about model updates of a target party. This paper presents *DeTrust-FL*, an *efficient privacy-preserving federated learning framework* for addressing the lack of transparency that enables isolation attacks, such as disaggregation attacks, during secure aggregation by assuring that parties' model updates are included in the aggregated model in a private and secure manner. *DeTrust-FL* proposes a decentralized trust consensus mechanism and incorporates a recently proposed decentralized functional encryption scheme in which all parties agree on a *participation matrix* before collaboratively generating decryption key fragments, thereby gaining control and trust over the secure aggregation process in a decentralized setting. Our experimental evaluation demonstrates that *DeTrust-FL* outperforms state-of-the-art FE-based secure multi-party aggregation solutions in terms of training time and reduces the volume of data transferred. In contrast to existing approaches, this is achieved without creating any trust dependency on external trusted entities.

Index Terms—federated learning, secure multi-party aggregation, privacy-enhanced computing, decentralized trust, decentralized functional encryption

I. INTRODUCTION

Training machine learning (ML) models that generalize well and have good predictive accuracy requires good quality training data. Traditionally, to train ML models, data is collected in a central place by the entities that run the training process. This data oftentimes includes private information, whose use and transmission have been restricted by regulations including the Health Insurance Portability and Accountability Act (HIPPA), the European General Data Protection Regulation (GDPR), and California Consumer Privacy Act (CCPA). To address privacy concerns and satisfy regulatory compliance, federated learning (FL) [1] has emerged as a promising learning paradigm, where multiple parties can collaboratively train a single machine learning model without sharing any of its training data. In other words, while traditional machine learning systems require the transmission of all training data

to a single place, FL ensures that each data owner or party can maintain possession of their own data. For this reason, FL represents a substantial privacy improvement.

In FL, an entity called *aggregator* is introduced to orchestrate the learning process among multiple parties who own private data. To train a model in FL, the aggregator sends an initial model to all parties, which then continue training it locally using their private data. Then, each party sends to the aggregator model parameters, e.g., gradients, which are known as *model updates* [2]. The aggregator then combines or fuses all model updates received creating a single shared ML model. These steps are repeated multiple *rounds*, and in each of them, the aggregator shares with the parties a new model and parties train and send back model updates. Private data is never transmitted from the party; only model updates are shared with the aggregator.

In FL all training data remains with the party who owns it, providing basic privacy. However, this simple FL setup still leaves room for *inference attacks* that enable adversaries to obtain information about the training data used by each party [3], [4]. The risk of private information leakage may be intolerable in some scenarios, and hence, multiple defenses to prevent these threats have been proposed including differential privacy (DP) and secure multi-party computation (SMC).

SMC approaches protect privacy of the input through cryptographic techniques to ensure that a curious or mistrusted aggregator cannot see individual model updates. Popular SMC approaches include (partial) homomorphic encryption [5], [6], threshold Paillier [7], functional encryption [8], and pairwise masking protocols [9]. In contrast, to mitigate inference attacks that use the final model or model updates, DP injects a carefully tuned amount of noise into model updates using differentially private mechanisms. While DP provides strong privacy guarantees, approaches relying on DP are notorious for producing noisy models with lower accuracy. Other solutions to prevent data leakage cleverly combine DP and SMC techniques to provide strong differential private guarantees while still enabling good model performance [7].

Unfortunately, existing SMC techniques are not sufficient to prevent inference threats for the following reasons. First, the lack transparency in the aggregation process results in an absence of guarantee that all of the received model updates will be included in a certain round of the FL process. Hence, a malicious aggregator can avoid including a target model update(s) in the secure aggregation process leading to multiple

problems. Besides enabling inference attacks, a model produced by ignoring certain model updates may not generalize well for all participants or lead to unfairness in its predictions. In an *isolation attack* [10], [11], a malicious aggregator may try to segregate certain replies by ignoring some of the received model updates to later on infer private information from a target party, thus violating its privacy. These attacks completely overthrow the purpose of SMC. Additionally, secure aggregation techniques are vulnerable to *disaggregation attacks* [12], [13], where a malicious aggregator can take advantage of the multi-round nature of federated learning to construct a system of equations to ultimately solve for the model updates of a target party. In summary, using existing crypto-based secure aggregation systems, parties cannot be sure that all the received model updates have been aggregated as expected and are vulnerable to potential disaggregation leading to inference attacks.

To address the previously described limitations, we propose *DeTrust-FL*, an efficient, scalable, and secure aggregation based privacy-preserving federated learning decentralized trust approach, that solves the dilemma of secure multi-party aggregation efficiency and centralized trust. *DeTrust-FL* still follows the direction of advanced cryptography based approach to maximize simplicity and efficiency in terms of communication and computation, and is built on a recently proposed decentralized functional encryption (FE) [14], [15] design in which each party generates the functional decryption key collaboratively. The aggregator can only obtain the aggregated result from the received encrypted parties' model updates using the functional decryption key. As a result, *DeTrust-FL* overcomes the limitation of previous state-of-the-art cryptography-based secure aggregation solutions.

Our key contributions are summarized as follows:

- We propose *DeTrust-FL*, an efficient and privacy-preserving federated learning (PPFL) in a decentralized trust setting. A successful process of secure aggregation relies on a transparent consensus among parties, in which each party can present its 'decision' or 'opinion' on a participation matrix that reflects its role in each secure aggregation round and its expectation of the proportion of benign parties in the FL training. *DeTrust-FL* utilizes decentralized functional encryption to perform the fundamental secure computing task.
- *DeTrust-FL* addresses the lack of transparency that enables isolation attacks during secure aggregation by assuring that parties' model updates are included in the aggregated model in a private and secure manner. *DeTrust-FL* also mitigates privacy leaks caused by recently demonstrated disaggregation attacks. To this end, we design a distributed consensus trust setting, which at the same time can address potential attacks where a curious aggregator colluding with a small number of parties.
- We describe an implementation of *DeTrust-FL* and demonstrate its applicability to train neural networks. According to our experimental evaluation, conducted in a cloud-based distributed setup, *DeTrust-FL* provides

the state-of-the-art model accuracy, communication and computation efficiency, while protecting the privacy of model updates from each party.

To the best of our knowledge, this is the first PPFL solution that gives each party fine-grained control over how their model updates are included and whether or not their participation matrix leaks private information, hence increasing the party's trust on collaborating via PPFL framework.

II. MOTIVATION AND PRELIMINARIES

The goal of SMC schemes is to protect the privacy of the inputs. In the FL setting, this means ensuring that nothing is revealed by the aggregation process other than the aggregated model. Attacks that compromise the privacy of existing SMC schemes stem from the ability of the aggregator to *i)* analyze the log of aggregated results or *ii)* manipulate the data that is fed to the SMC procedure.

A curious aggregator may launch a *disaggregation attack* by analyzing the outputs of the SMC over the training rounds. A matrix containing these outputs across training rounds and the knowledge of what parties participated in each round is used to build a system of equations that can be used to solve for the outputs of one of the parties. Note that this curious aggregator does not deviate from the training process and simply analyzes the matrix. To prevent this attack, it is necessary to ensure that the resulting matrix does not enable the construction of equations that can be solved to isolate a single input. Existing techniques [12], [13] only work for very large federations and are not appropriate for consortia with a small number of participants.

Isolation attacks consist of Isolation attacks consist of selectively inputting model updates to the SMC process in such a way that certain model updates can be isolated. Isolation attacks violate privacy and may lead to biased models. Techniques to prevent this attack have been proposed in [8] where an *inference prevention module* is introduced to prevent potential malicious behaviors during the secure aggregation procedure; however, this module is deployed by a fully trusted authority which is difficult to embody in real-life scenarios.

Finally, *replay attacks* occur when a malicious aggregator decides to send *old* model updates to the SMC process to force the construction of a vulnerable matrix. Both isolation and replay attacks are active attacks and are very difficult to detect. A potential solution is using specialize trusted hardware that is typically expensive and unavailable to federations. The goal of this paper is to enable SMC while preventing disaggregation, isolation and replay attacks without requiring specialize hardware or a *third party authority* (TPA).

Preliminaries of Decentralized Functional Encryption: Functional encryption (FE) is a family of cryptosystems that allows the computation of functions over encrypted data and as a result the final function value is obtained in plaintext [16]. FE schemes have shown promising efficiency improvements to address the secure multi-party aggregation challenge compared to existing homomorphic encryption-based approaches [8]. In [8], FL systems made use of a functional encryption scheme

[17] that enabled the computation of inner product using a TPA. Unlike the functional encryption scheme used in [8], this paper uses the recently proposed decentralized multi-client functional encryption (DMCFE) schemes [15], [14], [18].

DMCFE introduces two very interesting features: *i*) each party has a partial key and it does not require a TPA to generate functional keys for all the federation, and *ii*) it ensures that a particular functional key can only be used to compute a particular ciphertext by cryptographically marking the ciphertext with a *label*. Let $\mathcal{F} = \{\mathcal{F}_\rho\}_\rho$ be a family (indexed by ρ) of sets \mathcal{F}_ρ of functions $f : \mathcal{X}_{\rho,1} \times \dots \times \mathcal{X}_{\rho,n_\rho} \rightarrow \mathcal{Y}_\rho$ and $\mathcal{L} = \{0,1\}^* \cup \{\perp\}$ be a set of labels. A DMCFE scheme for the function family \mathcal{F} and the label set \mathcal{L} is a tuple of six algorithms $\mathcal{E}_{\text{DMCFE}} = (\text{Setup}, \text{KeyGen}, \text{KeyDerShare}, \text{KeyDerComb}, \text{Enc}, \text{Dec})$:

- $\text{Setup}(1^\lambda, 1^n)$: Takes as input a security parameter λ and the number of parties n , and generates public parameters pp .
- $\text{KeyGen}(\text{pp})$: Takes as input the public parameters pp and outputs n secret keys $\{\text{sk}_i\}_{i \in [n]}$.
- $\text{KeyDerShare}(\text{pp}, \text{sk}_i, f)$: Takes as input the public parameters pp , a secret key sk_i and a function $f \in \mathcal{F}_\rho$, and outputs a partial functional decryption key dk_i .
- $\text{KeyDerComb}(\text{pp}, \{\text{dk}_i\}_{i \in [n]})$: Takes as input the public parameters pp , n partial functional decryption keys $\{\text{dk}_i\}$ and outputs the functional decryption key dk_f .
- $\text{Enc}(\text{pp}, \text{sk}_i, x_i, l)$: Takes as input the public parameters pp , a secret key sk_i , a message $x_i \in \mathcal{X}_{\rho,i}$ to encrypt, a label $l \in \mathcal{L}$, and outputs ciphertext $\text{ct}_{i,l}$.
- $\text{Dec}(\text{pp}, \text{dk}_f, \{\text{ct}_{i,l}\}_{i \in [n]})$: Takes as input the public parameters pp , a functional key dk_f and n ciphertexts under the same label l and outputs a value $y \in \mathcal{Y}_\rho$.

III. THE PROPOSED *DeTrust-FL* FRAMEWORK

Threat Model and Assumptions. The objective of *DeTrust* is to ensure that the FL training process can provide privacy of the *input*, which guarantees the privacy of each party's local model, and the *output*, which ensures the privacy of the aggregated global model. The input privacy is achieved via secure aggregation, which allows the aggregator to learn only the aggregated model, and the output privacy is achieved through differential privacy. We assume that a malicious aggregator may attempt to infer private information of a target party during or after the training process. We focus our attention on disaggregation, isolation, and replay attacks described in Sec. II, and refer to the aggregator as *semi-trusted*. Moreover, we assume that the aggregator may collude with a *subset of parties* of bounded size to acquire private information of the other parties. We assume that parties are honest-but-curious, and do not consider attacks where malicious parties manipulate the model updates to poison the ML model or reduce its accuracy or to create a denial of service attack.

How does DeTrust-FL work? Figure 1 depicts the overview of *DeTrust-FL*. Similar to [5], [6], [8], *DeTrust-FL* adopts advanced cryptography-based secure aggregation to achieve




Fig. 1: Overview of *DeTrust-FL* framework.

privacy-preserving FL coupled with a consensus mechanism that provides transparency to the secure aggregation process.

Before training starts, *DeTrust-FL* performs a decentralized consensus through its decentralized trust consensus (DTC) module. The DTC module enables each party to inspect and negotiate the *participation matrix* that dictates what parties model updates will be input into the SMC process and with what weights during each round of training. Initially, the aggregator proposes a participation matrix to determine which parties will be selected for each training round. Parties engage in a protocol that enables them to agree on a matrix that does not allow for inference. At the end of DTC process, each party agrees on an updated participation matrix and uses it to generate the functional decryption key fragments, which represent its trust on the FL training process. Training can only proceed after the successful completion of the DTC.

At each training round, each party trains a local model, encrypts the model update using the DMCFE cryptosystem [14], and sends the encrypted model updates to the aggregator. As part of the encryption, each party *cryptographically labels* its model update with the current training round. The aggregator then computes the new *global model* by aggregating all encrypted inputs. The aggregator can execute secure aggregation by decrypting the set of encrypted model updates if and only if it receives all functional key fragments from parties for each training round. In addition, for the decryption to be successful, the aggregator can use the cryptographic key only for the current round, as the encrypted model updates are labeled with that round. Once the aggregator performs the decryption, it returns the new *global model* to each party.

Importantly, while our *DeTrust-FL* solution utilizes the same functional encryption family as previous PPFL approach [8], it does not introduce any additional entities (e.g., the online trusted third-party authority suggested in [8]). Thus, *DeTrust-FL* is architecturally equivalent to the standard FL paradigm [1]. Additionally, *DeTrust-FL* supports both *average* and *weighted* fusion methods directly, whereas the majority of secure aggregation approaches support just the former. This is discussed in Section III-D. In the following, we present the notation we use and more details of the private training process.

A. Notation

We use the following notation throughout the rest of the paper. Let \mathcal{A} be the aggregator and $\mathcal{S}_{\mathcal{P}}$ be the set of n parties

that participate in m FL training rounds. Let \mathcal{D}_j denote the dataset of party $\mathcal{P}_j \in \mathcal{S}_{\mathcal{P}}$. Let $\mathcal{M}_G^{(i)}$ and $\mathcal{M}_{\mathcal{P}_j}^{(i)}$ denote the global model and party \mathcal{P}_j 's local model, respectively, in FL round i . We use $\mathbf{V} \in \mathbb{R}^{m \times n}$ to denote the *participation matrix*, where each row $\mathbf{V}_i = (v_{i,\mathcal{P}_0}, \dots, v_{i,\mathcal{P}_n})_{1 \leq i \leq m}$ represents the fusion vector for i -th training round, defining the weights assigned to parties' local models for aggregation. That is, $\mathcal{M}_G^{(i)} = \sum_j v_{i,\mathcal{P}_j} \mathcal{M}_{\mathcal{P}_j}^{(i)}$. The column $\mathbf{V}_{,\mathcal{P}_j}$ represents the fusion weights of \mathcal{P}_j for all m training rounds. Note that the entry $V_{i,j} = 0$ means that during round i the local model of party j will not be added to the model. Let $DP_{SMC}(\epsilon, \mathcal{M}_{\mathcal{P}_j}^{(i)})$ denote the DP mechanism proposed in [7], which reduces the amount of injected noise by leveraging secure aggregation. We use t_g to denote the minimum number of parties that need to be aggregated in a round.

B. Training Process in DeTrust-FL

DeTrust-FL operation relies in two primary steps: *reaching a trust consensus* on a participation matrix that dictates who will participate during a particular round of training and what weights will be used to aggregate participant's model updates, and then, *performing privacy-preserving training* among parties in a decentralized setting.

Decentralized Trust Consensus (DTC). The objective of DTC is to establish an agreed-upon participation matrix \mathbf{V} for m rounds of FL training among n parties (see Sec. III-A for the definition). The *DTC* process begins with the aggregator proposing an initial participation matrix \mathbf{V} . An honest aggregator may propose a matrix that meets a batch partitioning (BP) constraint to prevent inference [13] and the global threshold requirements. A malicious aggregator may provide other input; however, the following steps would prevent the attack.

Each party investigates the proposed \mathbf{V} as follows: (i) does \mathbf{V} violate the BP criterion? (ii) does $\mathbf{V}_{,\mathcal{P}_j}$ accurately reflect fusion weight of party \mathcal{P}_j for each round of FL training and (ii) does \mathbf{V} ensure that a minimum number of parties are included into secure aggregation for all rounds? If any of these inspections fail, party \mathcal{P}_j can propose a new participation vector $\mathbf{V}_{,\mathcal{P}_j}$. Then, \mathcal{A} compiles all received assessments and conducts another round of party inspections until all the parties agree on the updated \mathbf{V} . Finally, each party generates functional decryption key fragments $\text{dk}^{\mathbf{V}} \in \mathbb{Z}^{m \times 1}$ based on the previously negotiated \mathbf{V} using parties' secret key. We will formalize these procedures in Section III-C.

Privacy-Preserving Training. For each training round, the aggregator queries parties with the current global model. Each party verifies it is participating in the training process during the current training round, which we refer to as being *enrolled* in a round of training, according to the previously agreed participation matrix. The enrolled party first trains a local model update using the current global model and its local dataset. In *DeTrust-FL*, we use the same hybrid method adding differential privacy noise as proposed in [7] to ensure both input and output privacy. Following that, the aggregator waits and collects a sufficient number of protected models in accordance with the agreed participation matrix. Then,

Algorithm 1: Privacy-Preserving DeTrust-FL Training

Input: Party set $\mathcal{S}_{\mathcal{P}}$, each party \mathcal{P}_j has its own dataset \mathcal{D}_j .
Output: Final trained global model $\mathcal{M}_G^{(m)}$.

- 1 Each entity initializes cryptographic keys;
- 2 Perform *Decentralized-Trust-Consensus* to build \mathbf{V} and $\text{dk}^{\mathbf{V}}$;
- 3 Aggregator \mathcal{A} initializes model $\mathcal{M}_G^{(0)}$;
- 4 **foreach** round $i \in \{1, \dots, m\}$ **do**
- 5 **foreach** $\mathcal{P}_j \in \mathcal{S}_{\mathcal{P}}$ **do**
- 6 \mathcal{A} queries $\text{msg}_{\text{req}} = (\mathcal{M}_G^{(i-1)}, \text{hyperparameters})$;
- 7 **if** \mathcal{P}_j is enrolled in \mathbf{V}_i **then**
- 8 $\mathcal{M}_{\mathcal{P}_j}^{(i)} \leftarrow \text{train}(\mathcal{D}_j, \mathcal{M}_G^{(i-1)})$;
- 9 $[\mathcal{M}_{\mathcal{P}_j}^{(i)}] \leftarrow \mathcal{E}_{\text{DMCFE}}.\text{Enc}_{\text{sk}_{\mathcal{P}_j}}(DP_{SMC}(\epsilon, \mathcal{M}_{\mathcal{P}_j}^{(i)}))$;
- 10 \mathcal{P}_j replies $[\mathcal{M}_{\mathcal{P}_j}^{(i)}]$ to \mathcal{A} ;
- 11 $\mathcal{S}_{\text{msg}_{\text{resp}}} \leftarrow \mathcal{A}$ waits and collects responses;
- 12 **if** response status $[\mathcal{M}_{\mathcal{P}_j}^{(i)}] \in \mathcal{S}_{\text{msg}_{\text{resp}}}$ satisfies \mathbf{V}_i **then**
- 13 \mathcal{A} recovers dk_i from fragments $\text{dk}_i^{\mathbf{V}} = \{\text{dk}_{i,\mathcal{P}_j}\}$;
- 14 $\mathcal{M}'_G \leftarrow \mathcal{E}_{\text{DMCFE}}.\text{Dec}_{\text{dk}_i}(\{[\mathcal{M}_{\mathcal{P}_j}^{(i)}]\})$;
- 15 \mathcal{A} updates $\mathcal{M}_G^{(i)} \leftarrow \mathcal{M}'_G$;
- 16 **return** final model $\mathcal{M}_G^{(m)}$;

using the collection of key fragments corresponding to current round's fusion weight vector \mathbf{V}_i , the aggregator recovers the functional decryption key $\text{dk}^{\mathbf{V}_i}$ for the round. The recovered $\text{dk}^{\mathbf{V}_i}$ is also bound to a specific round index and cannot be used in any other training round, as the DMCFE cryptosystem [15], [14] guarantees. Finally, the aggregator obtains the new global model by decrypting the collection of encrypted model updates without learning anything from each of the encrypted model updates. Notice that $\text{dk}^{\mathbf{V}_i}$ is generated collaboratively by the parties based on the trusted participation matrix that is negotiated through the DTC procedure.

Algorithm Overview. Algorithm 1 gives an overview of how *DeTrust-FL* works. Each participant creates its cryptographic keys (Line 1). To initiate the learning process, \mathcal{A} executes the *Decentralized-Trust-Consensus* function with all parties to agree on a participation matrix \mathbf{V} and corresponding functional decryption key fragments $\text{dk}^{\mathbf{V}}$ (Line 2). We present the decentralized trust consensus in detail in Section III-C. Then, \mathcal{A} initiates the training process in the same manner as standard FL does by sending a query message to all parties (Lines 4-6). Each enrolled party trains its local model and adds differential private noise to its model (Lines 7-10). When all responses are received, \mathcal{A} needs to check the response status $\mathcal{S}_{\text{msg}_{\text{resp}}}$ to ensure the required relies are satisfied according to the agreed-upon participation of current training round \mathbf{V}_i (Line 12). If the quorum verification passes, \mathcal{A} recovers current round's decryption key $\text{dk}^{\mathbf{V}_i}$ and performs the decryption among the encrypted model collection to obtain the new global model (Lines 13-15). We elaborate on the hybrid methodology used in privacy-preserving training in Section III-D.

C. Decentralized Trust Consensus Process

As we have illustrated, \mathbf{V}_i defines the assigned fusion weight used for each party's local model at i -th round and functional decryption key $\text{dk}_i^{\mathbf{V}}$ is collaboratively generated

Algorithm 2: Decentralized Trust Consensus

Input: $n :=$ num. parties and $m :=$ num. training rounds.
Output: Agreed-upon \mathbf{V} and decryption key fragments $\text{dk}^{\mathbf{V}}$.

- 1 $\mathcal{P}_j \in \mathcal{S}_{\mathcal{P}}$ initializes and sends local trust threshold $t_{\mathcal{P}_j}$;
- 2 \mathcal{A} collects and sets up the agreed-upon trust threshold $t_g = \max(\{t_{\mathcal{P}_j}\}_{\mathcal{P}_j \in \mathcal{S}_{\mathcal{P}}})$;
- 3 \mathcal{A} initializes $\mathbf{V}'^{m \times n}$ s.t. $\vdash \text{BP}(\mathbf{V}') \wedge \forall i, \|\mathbf{V}'_i\|_{\neq 0} \geq t_g$;
- 4 **function** *Decentralized-Trust-Consensus*($t_g, m, \mathcal{S}_{\mathcal{P}}$)
- 5 \mathcal{A} initializes empty $\text{dk}^{\mathbf{V}} \in \mathbf{0}^{m \times n}$ and threshold set \mathcal{S}_t ;
- 6 **do**
- 7 \mathcal{A} queries all parties with \mathbf{V}' ;
- 8 **foreach** $\mathcal{P}_j \in \mathcal{S}_{\mathcal{P}}$ **do**
- 9 $\text{vf}_{\mathcal{P}_j} \leftarrow \text{Party-Inspect-Negotiate}(\mathbf{V}'_{BP}, t_g)$;
- 10 **if** $\text{vf}_{\mathcal{P}_j} = \mathbf{V}_{\mathcal{P}_j}^{\text{suggest}}$ **is suggested then**
- 11 **if** \mathcal{A} **accepts** $\mathbf{V}_{\mathcal{P}_j}^{\text{suggest}}$ **then** $\mathbf{V}'_{\cdot j} \leftarrow \mathbf{V}_{\mathcal{P}_j}^{\text{suggest}}$;
- 12 **else** $\mathbf{V}'_{\cdot j} \leftarrow \mathbf{0}^{m \times 1}$;
- 13 \mathcal{A} updates \mathbf{V}' s.t. $\vdash \text{BP}(\mathbf{V}') \wedge \forall i, \|\mathbf{V}'_i\|_{\neq 0} \geq t_g$;
- 14 **while** all entities do not agree on participation matrix \mathbf{V} ;
- 15 **foreach** $\mathcal{P}_j \in \mathcal{S}_{\mathcal{P}}$ **do**
- 16 **foreach** row $\mathbf{V}_i \in \mathbf{V}$ **do**
- 17 $\text{dk}_{i, \mathcal{P}_j} \leftarrow \mathcal{E}_{\text{DMCFE}}.\text{KeyDerShare}_{\text{sk}_{\mathcal{P}_j}}(\mathbf{V}_i)$;
- 18 \mathcal{P}_j replies to \mathcal{A} with $\{\text{dk}_{i, \mathcal{P}_j}\}_{i \in \{1 \dots m\}}$;
- 19 \mathcal{A} fills column $\text{dk}_{\cdot j}^{\mathbf{V}} \leftarrow \{\text{dk}_{i, \mathcal{P}_j}\}_{i \in \{1 \dots m\}}$;
- 20 **return** \mathbf{V} and $\text{dk}^{\mathbf{V}}$;
- 21 **function** *Party-Inspect-Negotiate*(\mathbf{V}', t_g)
- 22 \mathcal{P}_j initializes $\text{vf}_{\mathcal{P}_j} = \{\}$;
- 23 **if** $t_g < t_{\mathcal{P}_j}$ **then return** *refuse*;
- 24 **if not** $\vdash \text{BP}(\mathbf{V}')$ **then return** *violate-BP*;
- 25 **foreach** row $i \in \mathbf{V}'$ **do**
- 26 **if** $\|\mathbf{V}_i\|_{\neq 0} \geq t_g$ **then**
- 27 **if** $\mathbf{V}_{i, \mathcal{P}_j}$ **is expected then** $\text{vf}_{\mathcal{P}_j} \leftarrow \text{accept}$;
- 28 **else** $\text{vf}_{\mathcal{P}_j} \leftarrow \mathbf{V}_{i, \mathcal{P}_j}^{\text{suggest}}$;
- 29 **else** $\text{vf}_{\mathcal{P}_j} \leftarrow \text{refuse}$;
- 30 **return** $\text{vf}_{\mathcal{P}_j}$

based on the provided \mathbf{V}_i . The aggregator can decrypt collected encrypted local models to obtain the aggregated global model if and only if (i) the aggregator proposed a “proper” \mathbf{V} that is unanimously agreed among parties, and (ii) those parties collaboratively generate $\text{dk}^{\mathbf{V}}$ key. A mutually agreed-upon \mathbf{V} can represent the parties’ unanimity of trust in the secure aggregation process, as well as in the privacy-preserving FL training. Thus, achieving such an agreement is the major objective of decentralized trust consensus setting, as various parties may have varying degrees of trust in a FL collaboration. **Detailed Procedure.** Algorithm 2 demonstrates how parties can agree on \mathbf{V} and then collaboratively generate $\text{dk}^{\mathbf{V}}$ key fragments for secure aggregation. Prior to the DTC process, *DeTrust-FL* ensures that each party is initialized with a threshold $t_{\mathcal{P}_j}$ that defines the individual trust threshold of the benign party proportion among n parties. Next, \mathcal{A} collects parties’ trust threshold and sets up the global trust threshold as $t_g = \max(\{t_{\mathcal{P}_j}\}_{\mathcal{P}_j \in \mathcal{S}_{\mathcal{P}}})$ (Lines 1-2). The DTC process begins with \mathcal{A} providing an initial $\mathbf{V}'^{m \times n} \in \mathbb{R}$ that complies with the Batch Partitioning (BP) criteria [13] and $\forall i, \|\mathbf{V}'_i\|_{\neq 0} \geq t_g$ (Line 3). The BP verification ensures that the participation of the parties in a particular training round do not enable disaggregation attacks. Next, \mathcal{A} repeatedly queries all parties

to obtain inputs on \mathbf{V}' , updates expected threshold t_{\max} , re-proposes \mathbf{V}' , until all parties reach to an agreed \mathbf{V} (Lines 6-14). To effectively re-propose a participation matrix \mathbf{V} , \mathcal{A} acquires suggested fusion weight $\mathbf{V}_{\mathcal{P}_j}^{\text{suggest}}$ from party \mathcal{P}_j to form new \mathbf{V} and ensure that:

$$\begin{cases} \text{comply with BP criteria, i.e., } \vdash \text{BP}(\mathbf{V}) \\ \forall i \in \{1, \dots, m\}, \|\mathbf{V}_i\|_{\neq 0} \geq t_g \text{ i.e. } t_g = \max(\{t_{\mathcal{P}_j}\}) \\ \forall j \in \{1, \dots, n\}, \mathbf{V}_{\cdot j} \text{ is accepted by } \mathcal{P}_j \end{cases}$$

When a party receives a participation matrix query message, each party first verifies the global trust threshold and the BP criteria compliance (Lines 23-24). For each row of compliance \mathbf{V} , \mathcal{P}_j inspects the enrollment status, trust threshold, and expected fusion weight that are associated to itself. If all inspections are successful, the party can respond with an accept signal; otherwise, it may respond with corresponding feedback (Lines 26-30). Finally, all parties agree on a participation matrix \mathbf{V} and generate functional decryption key fragments $\{\text{dk}_{i, \mathcal{P}_j}\}_{i \in \{1 \dots m\}}$ corresponding to that matrix using each party’s private key (Lines 15-20).

D. Contrasting DeTrust with Existing Techniques

Decentralizing Trust. *DeTrust-FL* relies on a novel FE cryptosystem that does not rely on a fully trusted entity overcoming some of the adoption problems of previously proposed FE-based systems including [8], where dk key is generated by a third-party entity, and the inference prevention module is also controlled by it. As a result, one limitation of [8] is that all parties needs to fully trust the third-party entity. *DeTrust-FL* completely removes this limitation. In *DeTrust-FL*, dk key is collaboratively generated by parties. Each party has the ability to verify the secure aggregation process in a fine-grained manner by controlling the partial functional decryption key generation. Without enough dk key fragments, the aggregator cannot decrypt to acquire the aggregated global model. Only if the totality of the parties reach consensus on the fusion weights, the aggregator can perform the secure aggregation process. In such a design, *DeTrust-FL* decentralizes the trust reliance on a centralized crypto dealer.

Transparency over the Secure Aggregation Process. Secure aggregation techniques suffer from lack of transparency with respect to the inputs an aggregator feeds to the SMC procedure. In *DeTrust-FL*, we allow \mathcal{P}_j to (i) define their personalized local trust threshold $t_{\mathcal{P}_j}$ for \mathbf{V} , and (ii) verify the number of parties participating in a particular round of training and the individual weights used to aggregate them $\mathbf{V}_{\cdot \mathcal{P}_j}$ in \mathbf{V} . The former setting represents the party’s confidence in enrolled parties for training process, whereas the local trust threshold $t_{\mathcal{P}_j}$ indicates the party’s expectation on minimum percentage of enrolled benign parties for one training round. For the enrolled training rounds it allows the aggregator and the party to negotiate on the expected fusion weights used in the FL training. Because the matrix is inextricably linked to the functional decryption key fragments, it dictates the results of the secure aggregation. This design removes the assumption of a trusted and honest aggregator.

Hybrid Methodology Compatibility. The secure aggregation of *DeTrust-FL* is achieved using DMCFE technique, in which each party encrypts its local model update using a secret key $sk_{\mathcal{P}_j}$ and the aggregator is issued a functional decryption key dk_i to decrypt the set of encrypted local models in order to obtain the aggregated model directly without learning the protected local models at i -th training round, as illustrated in Line 10 and Line 15 of Algorithm 1. After obtaining the aggregated global model, the aggregation can conduct another round of FL training. The main distinction between our secure aggregation solution and other FE-based secure aggregation solutions [8] is that in *DeTrust-FL*, dk_i is recovered from a set of key fragments generated by parties in accordance with the agreed-upon participation matrix, rather than being issued by a centralized trusted authority [8]. Additionally, our secure aggregation can be used with the hybrid methodology described in [7], in which we first inject a certain amount of differential privacy noise into the locally trained model and then encrypt the noise-injected model achieving higher model performance. As a result, *DeTrust-FL* retains the benefit of noise reduction for the global model through secure aggregation approaches.

Support for Fusion Algorithms. *DeTrust-FL* supports two commonly used algorithms to aggregate model updates: the *average fusion method* that averages aggregated local models and the *weighted fusion method* that fuses local models from each party with assigned party weight, e.g., FedAvg. Note that the latter is not supported in emerging additive HE-based (e.g., Paillier) secure aggregation solutions [5], [6], threshold Paillier [7] based secure aggregation, and pairwise masking based protocol [9]. As explained above, \mathbf{V}_i specifies the fusion weight assigned to each party’s local model in the aggregated global model at the i -th round. In the average fusion method, \mathbf{V}_i is set as an all- $\frac{1}{\|\mathbf{V}_i\|_{\neq 0}}$ vector in case that all parties are enrolled in the FL training. *DeTrust-FL* also supports FedAvg by ensuring that the weights in matrix V reflect the number of samples in each party.

IV. SECURITY AND PRIVACY ANALYSIS

A. Security of Cryptographic Infrastructure

As DMCFE is the underlying cryptographic infrastructure of our secure aggregation approach, the security of DMCFE is critical in the *DeTrust-FL* framework. Our adoption does not change any algorithms of the DMCFE scheme. According to formal security proof presented in [14], the encrypted model in *DeTrust-FL* has ciphertext indistinguishability and is secure against adaptive corruptions under the classical DDH assumption. To minimize redundancy, we will not discuss the correctness and security proofs to DMCFE here, and readers can refer to [14] for more details.

B. Privacy Guarantee

In this section, we give an overview of the privacy guarantees for *DeTrust-FL*.

Privacy of the Aggregated Global Model. Similarly to the comparable privacy-preserving federated learning proposals [8], [7], *DeTrust-FL* uses a differential privacy mechanism that

takes advantage of the secure aggregation procedure to protect the privacy of the output, i.e., the aggregated global model (in each round). Following the theoretical analysis reported in [7], the output model of *DeTrust-FL* also achieves the same differential privacy guarantee.

Privacy of Each Party’s Local Model. In the threat model, we consider a *semi-trusted aggregator* \mathcal{A} and a *limited number of colluding parties*, with disaggregation, isolation, and replay attacks described in Sec. II.

Inference Attack I (Isolation attack without collusion): \mathcal{A} can launch an inference attack targeting a specific party by manipulating the participation matrix for target i -th training round at target party \mathcal{P}_j . In the case of average fusion method, instead of setting up an all-one vector \mathbf{V}_i , \mathcal{A} may set \mathbf{V}_i as $(0, \dots, v_{i,\mathcal{P}_j}, \dots, 0)$, $v_{i,\mathcal{P}_j} = 1$. Thus, $\sum_i \mathbf{V}_i \mathcal{M}_i = \mathcal{M}_{\mathcal{P}_j}$. As a consequence, \mathcal{A} can infer the input of \mathcal{P}_j . Our proposed *inspection* module filters malicious \mathbf{V}_i by counting non-zero elements as demonstrated in Section III-C, which prevents isolation attacks.

Inference Attack I (Isolation Attack with Colluding Parties): To highlight this attack, consider $n = 7$ parties. Now, \mathcal{A} may propose a malicious $\mathbf{V}_i = (1, 1, 1, 0, 0, 0, 0)$ as one row of participation matrix to acquire the functional decryption key fragments. If \mathcal{A} can collude with two parties, then it can infer a benign party’s input. Our proposed *inspection* module ensures that in each round at least t_g parties’ models are included in the secure aggregation. Therefore, as long as the the aggregator colludes with fewer than $t_g - 2$ parties, it cannot infer the local model of any individual party.

Inference Attack III (Disaggregation Attack): \mathcal{A} can launch the *disaggregation attack* over multiple rounds of FL training by exploiting the side channel information (i.e., participation matrix), as demonstrated in [13], [12]. As illustrated in Section III-C, the *DeTrust-FL* design allows each party to examine the participation matrix to determine whether it complies with the Batch Partitioning (BP) criteria, a defense approach demonstrated in [13]. To prevent such a multi-round disaggregation attack, *DeTrust-FL* incorporates this BP criterion into our secure aggregation in a decentralized setting.

Inference Attack IV (Replay Attack): To highlight this attack, suppose $n = 6$ and let $\mathbf{V}_{i_1} = (1, 1, 1, 1, 0, 0)$ and $\mathbf{V}_{i_2} = (1, 1, 1, 1, 0, 0)$ for some $i_2 > i_1$. In other words, in rounds i_1 and i_2 , model updates from parties 1 through 4 will be averaged together. Now, the aggregator can target a specific party, say \mathcal{P}_1 , by launching a replay attack in round i_2 by using ciphertexts from round i_1 for parties \mathcal{P}_2 , \mathcal{P}_3 , and \mathcal{P}_4 . Note that in round i_1 secure aggregation yields $\sum_{j=1}^4 \mathcal{M}_{\mathcal{P}_j}^{(i_1)}$. Now, due to the replay attack in round i_2 , secure aggregation will yield $\mathcal{M}_{\mathcal{P}_1}^{(i_2)} + \sum_{j=2}^4 \mathcal{M}_{\mathcal{P}_j}^{(i_1)}$. From these two quantities, \mathcal{A} can obtain $\mathcal{M}_{\mathcal{P}_1}^{(i_2)} - \mathcal{M}_{\mathcal{P}_1}^{(i_1)}$, which violates the privacy of the local model for \mathcal{P}_1 . In *DeTrust-FL*, encrypted model updates are labeled with a round ensuring that the aggregator needs to use decryption keys only for the current round. Therefore, the aggregator cannot use ciphertexts from previous rounds in any round, which prohibits inference via replay attacks.

Based on the threat model described in Section III, we do not consider the DDoS attacks or malicious parties. Thus, if the *Party-Inspect-Negotiate* verification is passed, the party does not intentionally refuse the training query.

C. Forced Secure Aggregation

Existing secure aggregation solutions that employ homomorphic encryption schemes (e.g., Paillier and threshold Paillier) cannot ensure that parties’ encrypted models are included in the aggregated model. In these solutions, the aggregator is responsible for the fusion of the encrypted model updates. Simultaneously, the parties perform the decryption or partial decryption such that a *curious* aggregator can manipulate the fusion operation by only involving one party’s encrypted model update. This is a form of *isolation attack* [10], [11]. While it may not result in privacy leakage for some secure aggregation schemes, e.g. one based on Paillier encryption, it may result in a biased global model or aggregated model that does not generalize well to the target party’s data.

Unlike existing solutions, *DeTrust-FL* enables the forced secure aggregation functionality, where each party can ensure that its encrypted local model update is included in the aggregated model. *DeTrust-FL* does not involve any ciphertext fusion operation. To acquire the aggregated model, the aggregator performs the decryption over a set of encrypted local models while the decryption key is associated with the fusion weight vector that all parties should agree on. Thus, a party can check its assigned weight in the participation matrix. Suppose that the aggregator manipulates the fusion weight V_i for target party \mathcal{P}_j by sending other parties $V_i^{\text{manipulate}}$ s.t. $v_{i,j}^{\text{manipulate}} = 0$ while sending \mathcal{P}_j the fusion weight V_i^{valid} s.t. $v_{i,j}^{\text{valid}} = 1$ to make \mathcal{P}_j believe that its encrypted model is involved, but actually not. However, the function decryption key fragments $\text{dk}_{\mathcal{P}_j}^{V_i^{\text{valid}}}$ generated by \mathcal{P}_j and $\{\text{dk}_{\mathcal{P}_k}^{V_i^{\text{manipulate}}}\}_{j \neq k}$ generated by others cannot collaboratively recover $\text{dk}_{V_i^{\text{manipulate}}}$, which is ensured by the underlying cryptosystem. As a result, such manipulation is impossible.

V. EXPERIMENTAL EVALUATION

A. Experimental Setup

The implementation of the *DeTrust-FL* framework builds on the IBM Federated Learning (IBM-FL) framework [2]. We implement the underlying cryptosystem, i.e., the decentralized MCFE scheme [14] in Python based on the *gmpy2* library. We also incorporate the techniques for final discrete logarithm computation as used in [8] in our implementation. To integrate our proposed solution with the IBM-FL framework, we implement *DeTrust-FL* privacy-preserving *global fusion handler* and *local training handler* classes for both aggregator and parties that respectively inherited the *IterAvgFusionHandler* class and *LocalTrainingHandler* class of IBM-FL. We adopted and implemented an instance of DMCFE of [14].

To benchmark the performance of *DeTrust-FL*, we train two types of Keras-based Convolutional Neural Network (CNN) models, as described in TABLE I, to classify the publicly

available MNIST dataset of handwritten digits [19] and CIFAR10 dataset of colour images [20]. The total parameters of two CNN models are 1,199,882 and 890,410, respectively. TABLE I also shows dataset assignment for each party in non-iid setting.

In the configuration of the IBM-FL framework, as partially included in TABLE I, we set the following hyperparameters: (i) termination accuracy is 0.99; (ii) total global training rounds is 20 for MNIST dataset and 30 for CIFAR10 dataset; (iii) learning rate is 0.01; (iv) each local training epoch varies according to the evaluated dataset and adopted model. The crypto-based secure aggregation works on the *integer* field, while model parameters are in *floating-point number* format. *DeTrust-FL* has a parameter of *encoding precision* to define the scale factor of the conversion between integers and floating-point numbers, as do in existing solutions [8], [6], [5]. We use the term *precision* to denote the number of digits after the decimal point that are considered for rounding off. The default encoding precision is 4, which means we keep four digits after the decimal point. The rounded off number is then appropriately scaled to convert it to an integer.

As analysed in Section III-D, the secure aggregation is easily integrated with differential privacy (DP) technique in the hybrid methodology. Because the impact and effect of DP-integration are reported in [8], [7], we omit the similar results for the sake of brevity. We compare our *DeTrust-FL* framework to various baselines:

- General FL training without any secure aggregation setting (General-FL);
- FL training using partially additive homomorphic encryption (i.e., Paillier cryptosystem) based secure aggregation (PHE-FL, e.g., [5], [21]);
- FL training using threshold Paillier based secure aggregation (HybridOne [7]);
- FL training using functional encryption based secure aggregation (HybridAlpha [8]).

Note that, while both *DeTrust-FL* and *HybridAlpha* employ functional encryption techniques, *HybridAlpha* relies on a trusted online third-party entity to provide key service for each training round.

All experiments are performed on a cloud instance equipped with Intel(R) Xeon(R) CPU E5-2683 v4 platform with 32 cores and 64GB of RAM. Note that the network latency is not measured in our experiment as the framework is running on the same machine in a multi-process setting.

B. Experimental Results

1) *Comparison of Performance With Various Approaches: Model Accuracy.* We present the model accuracy of *DeTrust-FL* with comparison to four baselines using the MNIST dataset. Each test case is repeated at least three times, and we report the min/max/avg values in our results. As shown in Fig. 2, we report the model accuracy of the global model over 20 global training rounds, where for each training round each party performs 3 local training epochs. We observe that (i) the introduction of secure aggregation does not impact the final

TABLE I: Summary of Keras Model, Dataset Assignment and FL Setting

Type	Model Architecture	Parameters	Parties	Training/Test per Party	FL Rounds	Local Epochs
CNN-MNIST	2xConv2D→MaxPooling→Droupout→Flatten→Dense→Dense	1,199,882	5	500/2000 (non-iid)	20	3
CNN-CIFAR10	2x(2xConv2D→MaxPooling→Droupout)→Flatten→Dense→Dense	890,410	10	5000/1000 (non-iid)	30	20




Fig. 2: Model accuracy, training time and transmission payload comparison for multiple secure aggregation schemes in FL training on evaluating MNIST dataset.

model accuracy and the rate of learning convergence, and (ii) *DeTrust-FL* has comparable model quality with respect to the general FL training without secure aggregation setting, as well as to compared FL frameworks with secure aggregation.

Training Time. We also analyze the impact of the secure aggregation approaches on total training time, as shown in Fig. 2, the second figure. It shows that *DeTrust-FL* has a training time close to *General-FL* and achieves significant training efficiency improvements compared to existing secure aggregation-enabled FL solutions.

DeTrust-FL achieves efficiency comparable to *HybridAlpha*, with both being the most efficient of the crypto-based secure aggregation approaches. However, *HybridAlpha* relies on a centralized and trusted third-party authority to do so while *DeTrust-FL* does not. In the context of 20 global training rounds, *DeTrust-FL* reduces training time by about 70.2% and 80.3% when compared to secure aggregation solutions of *PHE-FL* and *HybirdOne* frameworks, respectively. The underlying cryptosystem and reduced communication interactions contribute to this improvement in training time efficiency, as the *DeTrust-FL* utilizes decentralized multi-input client encryption with more efficient encryption and decryption algorithms, and only requires one round of interaction per secure aggregation.

Communication Interaction and Transmission Overhead. To measure communication efficiency, first we theoretically analyze the communication rounds needed by each approach and then report the experimental results.

Suppose that the FL training is performed over m rounds of global training among n parties \mathcal{P} , a key server \mathcal{K} and an aggregator \mathcal{A} . Notice that, in *DeTrust-FL*, the key server participates only during the setup time. Table II summarizes total number of interactions. It shows that *DeTrust-FL* uses the smallest number of interactions compared to other secure aggregation solutions. Note that we include one round of registration interaction between the aggregator and n parties, as required in the IBM-FL framework.

We also accumulate the total transmission payload for each FL training round, in terms of aggregator’s query and parties’ response, with respect to all approaches. Fig. 2 reports the

comparison of transmission payload during the entire FL training. Compared to *PHE-FL* and *HybridOne* solutions, *DeTrust-FL* reduces the volume of transmission payload by 73.6% and 82.2%, respectively. Our work also generates a transmission payload similar to *HybridAlpha*. However, we reduce the number of interactions by 16.4% in the setting of 20 global training rounds with 5 parties, as shown in Table II.

TABLE II: Comparison of communication interaction

Proposal	$\mathcal{A} \leftrightarrow \mathcal{P}$	$\mathcal{A} \leftrightarrow \mathcal{K}$	$\mathcal{P} \leftrightarrow \mathcal{K}$	Total
General-FL	$mn + n$	0	0	$mn + n$
PHE-FL	$2mn + n$	1	n	$2mn + 2n + 1$
HybridOne	$2mn + n$	1	n	$2mn + 2n + 1$
HybridAlpha	$mn + n$	$m + 1$	n	$mn + m + 2n + 1$
DeTrust-FL	$mn + n$	1	n	$mn + 2n + 1$

Notation: m rounds of global training with one aggregator \mathcal{A} , n parties \mathcal{P} and one key server (or TPA) \mathcal{K} .

2) *Performance of DeTrust-FL:* As demonstrated above, *DeTrust-FL* and *HybridAlpha* solutions outperform all other solutions in terms of accuracy and training time. We now further examine the performance of *DeTrust-FL* in the context of 30 FL training rounds with 10 parties using the CIFAR10 dataset.

We first report the model performance of *DeTrust FL*, as well as *General-FL* and *HybridAlpha*. We evaluate model accuracy and training time in the setting as illustrated in TABLE I with encoding precision as four on floating-point parameters. As shown in the first two figures of Fig. 3, we report the model accuracy curve and training time for CIFAR10 dataset. We observe that both *DeTrust-FL* and *HybridAlpha* can achieve better model accuracy compared to the *General-FL*. Theoretically, we believe that the slight improvement in accuracy is due to the encoding operation that converts the model weights between integers and floating-point numbers, which discards some information and could be considered as a type of pruning. However, we left it as an open question to be investigated further in the future.

As shown in last two Fig. 3, we also illustrate that *DeTrust-FL* inherits the advantage of general FL wherein the global model has a better rate of convergence and model accuracy




Fig. 3: Performance comparison on evaluating CIFAR10 dataset and decomposition of model accuracy in *DeTrust-FL* training on evaluating CIFAR10 dataset and MNIST dataset.




Fig. 4: Impact of number of parties in *DeTrust-FL* training on evaluating MNIST dataset with settings of precision = 4 and 3 local training epochs per training round.




Fig. 5: Impact of encoding precision of floating-point parameters in *DeTrust-FL* training on evaluating MNIST dataset with setting of 5 parties and 3 local training epochs per training round.

compared to those local models, for both MNIST and CIFAR10 dataset.

3) *Impact of Numbers of Parties*: We also evaluate the impact of numbers of parties on *DeTrust-FL* performance. As reported in Fig. 4, more parties enrolled in the FL training can increase the model accuracy. In our experimental setting, each party’s training samples are assigned randomly; thus, the model accuracy improvement is not as significant as if they were more unbalanced, which often occurs in practice. Fig. 4 also shows the training time results over 20 global training rounds. The training time increases along the number of parties. Each party in the IBM FL experiment is simulated by multiple processes in a parallel setting on a single node. Hence, the increased training time is caused by the increased computation time in the decryption phase and functional decryption key recovering phase at the aggregator, rather than the time spent at the parties.

4) *Impact of Encoding Precision on Floating-point Parameters*: When performing crypto-based secure aggregation, the floating-point model parameters should be encoded in integer format. We evaluate the impact of encoding precision on model accuracy and training time in *DeTrust*. As shown in Fig. 5, the encoding precision setting has no significant influence on model accuracy. Furthermore, a higher encoding precision

setting may slightly increase the training time. As mentioned above, since parties are simulated with multiple processes in this experiment there are some fluctuations in the growth curve.

VI. RELATED WORK

Federated learning (FL) as proposed in [1] has emerged as a promising approach to collaboratively train a global model without sharing local data among participants and thereby providing a primary privacy guarantee. The primary FL design is not sufficient for privacy-sensitive applications because it is still vulnerable from inference attacks on the final model and model updates between parties and aggregator, as demonstrated by [3], [4]. To partially address those threats, several privacy-preserving techniques such as differential privacy and secure multi-party computation have been adopted during the learning process.

Emerging hybrid privacy-preserving methodologies have shown promise to providing strong privacy guarantees [8], [7]. The approach in [7] employs the threshold Paillier cryptosystem to perform multi-party secure aggregation, while *HybridAlpha* [8] uses a functional encryption cryptosystem to securely aggregate encrypted local models. Existing secure aggregation approaches mainly follow three directions: *pairwise masking* based approach [9], [22], *homomorphic encryption* based approach [23], [7], [11] and *functional encryption* [8], [24], [25], [26], [27].

Pairwise masking [9], [22] is a lightweight computation approach where each input is masked with one-time pads while the aggregated output cancels out those pads. For example, Bonawitz et al. [9] make use of double pairwise additive masking by integrating approaches such as secret sharing, key agreement, authenticated encryption, and public key infrastructure. Even though those approaches are efficient in computation, they either require peer-to-peer communication or rely on a server forwarding the encrypted shares and masked inputs for each aggregation round, resulting in communication inefficiency. Additionally, they are vulnerable to disaggregation attacks that are recently demonstrated in [12], [13] in the context of multi-round secure aggregation in FL training.

In homomorphic encryption-based approaches, one set of solutions such as [7], [11] rely on additive homomorphic encryption schemes like the Paillier and its variants. Another set of solutions such as PySyft [23] employs the SPDZ protocol that is built on somewhat homomorphic encryption

(SHE) in the form of BGV [28]. Those approaches are efficient in communication but are computationally inefficient. Both pairwise masking and homomorphic encryption-based secure aggregation solutions only support average fusion algorithms and are not applicable to weighted fusion algorithms. In addition, homomorphic encryption-based secure aggregation cannot guarantee secure aggregation to a party without special hardware.

A functional encryption-based solution, *HybridAlpha*, [8] has achieved significant training efficiency improvements by efficient cryptographic computation and simple communication design. Furthermore, it also supports the weighted secure aggregation situation and addresses ensures a minimum number of pre-defined replies are aggregated. However, *HybridAlpha* introduces and requires an online third-party entity that should be fully trusted by all parties, which makes impractical in some FL scenarios.

None of the approaches outlined above achieves computational and communication efficiency without requiring a trusted crypto authority to provide cryptographic key services or a partially trusted aggregator to honestly execute secure aggregation methods and without conducting stealth targeted or disaggregation attacks.

VII. CONCLUSION

Secure aggregation, the underlying privacy-preserving computation primitive, plays an essential role during the privacy-preserving FL training. Existing approaches do not address the dilemma of efficiency requirements and decentralized trust setting. To address these limitations, in this paper we propose *DeTrust-FL*, a novel FL framework that does not require a fully trusted crypto dealer, as well as mitigating partial trust on the aggregator. In addition, our experimental evaluation shows that our approach achieves comparable performance in terms of model accuracy and training time, and reduces the communication interactions by 16.4% compared to the state-of-the-art efficient approach without the prerequisite of a centralized trust setting.

REFERENCES

- [1] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," *arXiv preprint arXiv:1610.05492*, 2016.
- [2] H. Ludwig, N. Baracaldo, G. Thomas, Y. Zhou, A. Anwar, S. Rajamoni, Y. Ong, J. Radhakrishnan, A. Verma, M. Sinn *et al.*, "Ibm federated learning: an enterprise framework white paper v0. 1," *arXiv preprint arXiv:2007.10987*, 2020.
- [3] M. Nasr, R. Shokri, and A. Houmansadr, "Comprehensive privacy analysis of deep learning: Stand-alone and federated learning under passive and active white-box inference attacks," in *IEEE S&P'19*, IEEE. ACM, 2019.
- [4] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership inference attacks against machine learning models," in *IEEE S&P'17*. IEEE, 2017, pp. 3–18.
- [5] C. Liu, S. Chakraborty, and D. Verma, "Secure model fusion for distributed learning using partial homomorphic encryption," in *Policy-Based Autonomic Data Governance*. Springer, 2019, pp. 154–179.
- [6] C. Zhang, S. Li, J. Xia, W. Wang, F. Yan, and Y. Liu, "Batchcrypt: Efficient homomorphic encryption for cross-silo federated learning," in *USENIX ATC'20*, 2020, pp. 493–506.

- [7] S. Truex, N. Baracaldo, A. Anwar, T. Steinke, H. Ludwig, R. Zhang, and Y. Zhou, "A hybrid approach to privacy-preserving federated learning," in *ACM AISec'19*, 2019, pp. 1–11.
- [8] R. Xu, N. Baracaldo, Y. Zhou, A. Anwar, and H. Ludwig, "Hybridalpha: An efficient approach for privacy-preserving federated learning," in *ACM AISec'19*, 2019, pp. 13–23.
- [9] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth, "Practical secure aggregation for privacy-preserving machine learning," in *ACM CCS'17*, ACM. ACM, 2017, pp. 1175–1191.
- [10] B. Chen and H. Yu, "Secure aggregation with malicious node revocation in sensor networks," in *ICDCS'11*. IEEE, 2011, pp. 581–592.
- [11] Y. Liu, Z. Ma, X. Liu, S. Ma, S. Nepal, and R. Deng, "Boosting privately: Privacy-preserving federated extreme boosting for mobile crowdsensing," *arXiv preprint arXiv:1907.10218*, 2019.
- [12] M. Lam, G.-Y. Wei, D. Brooks, V. J. Reddi, and M. Mitzenmacher, "Gradient disaggregation: Breaking privacy in federated learning by reconstructing the user participant matrix," *arXiv preprint arXiv:2106.06089*, 2021.
- [13] J. So, R. E. Ali, B. Guler, J. Jiao, and S. Avestimehr, "Securing secure aggregation: Mitigating multi-round privacy leakage in federated learning," *arXiv preprint arXiv:2106.03328*, 2021.
- [14] M. Abdalla, F. Benhamouda, M. Kohlweiss, and H. Waldner, "Decentralizing inner-product functional encryption," in *PKC'19*. Springer, 2019, pp. 128–157.
- [15] J. Chotard, E. D. Sans, R. Gay, D. H. Phan, and D. Pointcheval, "Decentralized multi-client functional encryption for inner product," in *Asiacrypt'18*. Springer, 2018, pp. 703–732.
- [16] D. Boneh, A. Sahai, and B. Waters, "Functional encryption: Definitions and challenges," in *Theory of Cryptography*, Y. Ishai, Ed. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 253–273.
- [17] M. Abdalla, D. Catalano, D. Fiore, R. Gay, and B. Ursu, "Multi-input functional encryption for inner products: function-hiding realizations and constructions without pairings," in *Crypto'18*, Springer. Springer, 2018, pp. 597–627.
- [18] J. Chotard, E. Dufour-Sans, R. Gay, D. H. Phan, and D. Pointcheval, "Dynamic decentralized functional encryption," *IACR Cryptol. ePrint Arch.*, p. 197, 2020.
- [19] Y. LeCun, C. Cortes, and B. Christopher J.C., "MNIST handwritten digit database," 2010. [Online]. Available: <http://yann.lecun.com/exdb/mnist/>
- [20] A. Krizhevsky, G. Hinton *et al.*, "Learning multiple layers of features from tiny images," 2009.
- [21] S. Sharma, C. Xing, Y. Liu, and Y. Kang, "Secure and efficient federated transfer learning," in *IEEE BigData'19*. IEEE, 2019, pp. 2569–2576.
- [22] S. Kadhe, N. Rajaraman, O. O. Koyluoglu, and K. Ramchandran, "Fast-secagg: Scalable secure aggregation for privacy-preserving federated learning," *arXiv preprint arXiv:2009.11248*, 2020.
- [23] T. Ryffel, A. Trask, M. Dahl, B. Wagner, J. Mancuso, D. Rueckert, and J. Passerat-Palmbach, "A generic framework for privacy preserving deep learning," in *Proceedings of Privacy Preserving Machine Learning Workshop with NeurIPS 2018*, 2018.
- [24] R. Xu and J. Joshi, "Revisiting secure computation using functional encryption: Opportunities and research directions," *arXiv preprint arXiv:2011.06191*, 2020.
- [25] R. Xu, J. B. Joshi, and C. Li, "Cryptonn: Training neural networks over encrypted data," in *IEEE ICDCS*. IEEE, 2019, pp. 1199–1209.
- [26] R. Xu, J. Joshi, and C. Li, "Nn-emd: Efficiently training neural networks using encrypted multi-sourced datasets," *IEEE TDSC*, 2021.
- [27] R. Xu, N. Baracaldo, Y. Zhou, A. Anwar, J. Joshi, and H. Ludwig, "Fedv: Privacy-preserving federated learning over vertically partitioned data," in *ACM AISec'21*, 2021, pp. 181–192.
- [28] Z. Brakerski, C. Gentry, and V. Vaikuntanathan, "(leveled) fully homomorphic encryption without bootstrapping," *ACM TOCT*, vol. 6, no. 3, pp. 1–36, 2014.