

ScionFL: Efficient and Robust Secure Quantized Aggregation

Yaniv Ben-Itzhak
Shay Vargaftik
Avishay Yanai
VMware Research

Benny Pinkas
Bar-Ilan University

Helen Möllering
Thomas Schneider
Hossein Yalame
Technical University of Darmstadt

Ajith Suresh
Technology Innovation Institute (TII)

Oleksandr Tkachenko
DFINITY Foundation

Christian Weinert
Royal Holloway, University of London

Abstract

Secure aggregation is commonly used in federated learning (FL) to alleviate privacy concerns related to the central aggregator seeing all parameter updates in the clear. Unfortunately, most existing secure aggregation schemes ignore two critical orthogonal research directions that aim to (i) significantly reduce client-server communication and (ii) mitigate the impact of malicious clients. However, both of these additional properties are essential to facilitate cross-device FL with thousands or even millions of (mobile) participants.

In this paper, we unite both research directions by introducing ScionFL, the first secure aggregation framework for FL that operates *efficiently* on quantized inputs and simultaneously provides robustness against malicious clients. Our framework leverages (novel) multi-party computation (MPC) techniques and supports multiple linear (1-bit) quantization schemes, including ones that utilize the randomized Hadamard transform and Kashin’s representation.

Our theoretical results are supported by extensive evaluations. We show that with *no overhead for clients* and moderate overhead on the server side compared to transferring and processing quantized updates in plaintext, we obtain comparable accuracy for standard FL benchmarks. Additionally, we demonstrate the robustness of our framework against state-of-the-art poisoning attacks.

1 Introduction

Federated learning (FL) [91] is a paradigm for large-scale distributed machine learning, where in each training round a subset of clients locally updates a global model that is then centrally aggregated. FL quickly gained popularity due to its promises of data privacy, resource efficiency, and ability to handle dynamic participants.

However, in terms of *privacy*, the central aggregator learns the individual client updates in the clear and thus can infer sensitive details about the clients’ private input data [56, 92, 106, 117, 129]. Hence, many secure aggregation schemes, e.g., [26, 53, 96], have been developed, where the aggregator only learns the aggregation result, i.e., the global model (we refer to [68, 100] for a discussion of differential privacy in FL as an orthogonal privacy-enhancing paradigm). Most prominently, in the “SecAgg” protocol [26], clients exchange masks with peers to blind their model updates such that the masks cancel out during aggregation and reveal only the exact result. However, this approach requires an interactive setup between clients and thus is less reliable when dealing with real-world problems such as client dropouts (except for special variants [67, 83, 133]). Moreover, recent research has shown that a *single malicious aggregator* can reconstruct individual training data from clients’ inputs, despite the use of secure aggregation [24, 25, 101].

In terms of *resource efficiency* of plain FL, clients have to send parameter updates (also known as *gradients*), which typically consist of thousands or millions of coordinates with one floating point number per coordinate. For cross-device FL involving mobile clients with limited upload bandwidth, this quickly becomes infeasible when dealing with increasingly large model architectures, where gradients consist of millions of coordinates. Therefore, quantization schemes that exploit the noise resiliency of gradient-based FL methods (e.g., federated averaging [91]) have been proposed to significantly compress client updates (typically replacing the representation of each coordinate by a single bit, instead of a 32-bit floating point number), e.g., [4, 12, 18, 30, 114, 123, 126].

Unfortunately, so far the ML community has worked on optimizing FL for resource efficiency, whereas the security community has worked separately on optimizing secure aggregation for privacy. One of the very few exceptions is a work [35] that combines SecAgg with sketching [111] for compression of gradients. Besides the mentioned disadvantages of SecAgg, their work considers only a single compression method. However, due to the trade-offs between accuracy and computational efficiency provided by different forms of quantization, which become especially relevant for gradients with millions of coordinates, it is important to have a modular approach.

Beyond data privacy threats, FL was also shown to be vulnerable to manipulations by malicious clients who alter their local models/updates, affecting the characteristics of the final global model [10, 20, 51, 115, 116, 132]. This highlights the need for effective *defense measures* capable of thwarting such attacks. While there exists a plethora of such attacks and defenses in the FL literature (summarized in §A.6), we focus on defending against *untargeted poisoning attacks*, in which the attacker attempts to damage the trained model’s performance for a large number of test inputs, typically resulting in a final global model with a high error rate [11, 51, 115]. Unfortunately, existing defenses for such attacks cannot be efficiently translated to secure aggregation in the quantized setting [99, 137].

1.1 Our Contributions

We propose ScionFL, a framework that enables *efficient* and *robust* secure aggregation for FL with a distributed aggregator that operates *directly* on *quantized* parameter updates. Specifically, ScionFL has virtually *no additional communication overhead for clients* compared to the insecure transfer of quantized plaintext updates. Achieving this is non-trivial for prior single-server solutions that are based on masking techniques or homomorphic encryption (HE) [26, 136].

In ScionFL, we rely on outsourced multi-party computation (MPC), where the clients apply computationally efficient secret-sharing

techniques to distribute their sensitive (quantized) parameter updates among multiple servers that together form a distributed aggregator. These servers then use interactive cryptographic MPC protocols [17, 40, 41] to securely compute the aggregation and only reveal the updated global model.

The distributed aggregator model is well established [53, 57, 89, 107] (although prior works cannot efficiently handle quantized inputs) as it has many practical benefits over single-server solutions, including not requiring an interactive setup between clients [96, 99] and the ability to handle dropouts efficiently. Moreover, in recent years, numerous studies have demonstrated that single server aggregation methods are susceptible to privacy attacks when an aggregator is corrupt [24, 25, 54, 101, 118, 131]. The vulnerability of these methods stems from the fact that the aggregator holds complete authority in selecting clients and the data transmitted to and received from them. On the other hand, with our MPC protocols, data privacy can be guaranteed even if all servers except one are compromised or their operators receive subpoena requests.

As MPC protocols typically induce significant overhead in terms of (inter-server) communication, we propose optimizations for secure aggregation that *support any “linear” quantization scheme*, including *1-bit quantization* schemes that use preprocessing in the form of random rotations [123] and Kashin’s representation [30]. Moreover, we study novel *approximate* MPC variants that exploit the noise tolerance of FL and might be of independent interest. We formally prove that our resulting secure aggregation scheme is an *unbiased estimator* of the arithmetic mean and explore trade-offs between efficiency and accuracy.

We implement a combined end-to-end FL evaluation and MPC simulation environment, which we will open-source upon paper publication. Our prototype implementation allows to assess the performance and accuracy of our solution for stochastic quantization schemes, including recent state-of-the-art distributed mean estimation techniques [30, 123]. Our results demonstrate that standard FL benchmarks’ accuracy is barely impacted while our optimizations and approximation can significantly reduce inter-server communication. For example, when training the LeNet architecture for image classification on the MNIST data set [82] using 1-bit stochastic quantization with Kashin’s representation [30], training accuracy is only slightly reduced from 99.04% to 98.71% after 1000 rounds, while inter-server communication drops from 16.14 GB to 0.94 GB compared to naive MPC-based secure aggregation when considering 500 clients per round, an improvement by factor 17.2 \times .

Since clients may act maliciously and try to degrade accuracy with their updates, we design a novel bipartite defense mechanism called ScionFL-Aura to ensure the robustness of our framework. Specifically, we provide protection against state-of-the-art untargeted poisoning attacks [115], combining magnitude clipping and directional filtering based on the gradients’ approximate L2-norms and cosine distances. Notably, ScionFL-Aura is the first defense mechanism to work directly on quantized inputs and thus enables a highly efficient realization in MPC, whereas existing works require expensive MPC conversions of all individual parameters [5, 99].

We summarize our contributions as follows:

(1) First *secure aggregation* framework to consider (1-bit) quantization with almost *no communication overhead for clients* compared to *plaintext quantized FL*.

(2) Novel *optimizations and approximations* to reduce MPC inter-server communication, offering performance/accuracy trade-offs.

(3) End-to-end FL evaluation and MPC simulation environment, demonstrating the efficiency and accuracy of our approach.

(4) First efficient and effective FL (poisoning) defense operating directly on quantized updates.

Notably, while we study FL as our primary application, our secure aggregation protocols have numerous other applications like privacy-preserving aggregate statistics, for which there are currently (less efficient) real-world deployments that also rely on distributed aggregators (e.g., telemetry reporting in Mozilla’s Firefox browser [1, 2]). For these settings, we achieve improvements in communication of up to 4 \times over prior works like Prio+ [2] and the details are provided in §3.5.

1.2 Related Work

We present a summary of the most relevant related works here, with a more comprehensive discussion in §A.

Quantization and Compression in FL: In this work, we focus on quantization to reduce communication in FL. However, an alternative line of work investigates compression techniques for the same purpose. There are three main directions for gradient compression in cross-device FL: (i) gradient sparsification (e.g., [3, 52, 75, 120]), (ii) client-side memory and error-feedback (e.g., [4, 19, 109, 114]), and (iii) entropy encodings (e.g., [4, 123, 125]). Reviews of current state-of-the-art gradient compression techniques and some open challenges can be found in [68].

Compression techniques are less suited for the requirements of secure aggregation for cross-device FL than quantization, e.g., due to computational overhead, incompatibility with secure aggregation, and state-requirements on the client side. We discuss more details in §2.2 and §A.2.

Secure Aggregation & Model Inference Attacks: In conventional FL with a single aggregator, clients share locally trained model updates with a central party to contribute to the training of a global model. However, sharing those updates makes the system vulnerable to data leakage. Attacks exploiting this leakage are called *inference attacks* [23, 81, 98]. Even a semi-honest central server can learn confidential information about the used private training data by analyzing the received local model updates.

A common countermeasure against inference attacks is to use secure aggregation [49, 80] (cf. §A.5). As FL poses specific challenges such as a large number of clients and drop-out tolerance, tailored secure aggregation protocols for FL have been proposed [16, 26, 53, 67, 113, 119]. Those hide individual updates, ensuring that the server has only access to global updates, hence, effectively prohibiting the analysis of individual updates for inference attacks. The first scheme, SecAgg [26], combines secret sharing with authenticated symmetric encryption, but requires 4 communication rounds per training iteration among servers and client. Bell et al. [16] improve upon SecAgg [26] by reducing client communication and computation to poly-logarithmic complexity. However, from a practical point of view also [16] as well as other existing secure aggregation protocols designed for FL still exhibit significant computation and communication overhead: Due to underlying secret sharing or encryption, those protocols typically encode each local

update in 32-bit and add computational overhead caused by the required cryptographic operations. In contrast, ScionFL enables highly communication-efficient secure aggregation thanks to 1-bit quantization and causes almost *no* additional communication overhead on the client side compared to plaintext FL. Fereidonni et al. [53] provide more details by comparing several secure aggregation protocols with respect to efficiency and practicality. Mansouri et al. [89] provide a comprehensive analysis of secure aggregation schemes w.r.t. their suitability for FL.

To the best of our knowledge, only Chen et al. [35] and Beguier et al. [15] have considered both compression and secure aggregation in combination for FL. Specifically, Chen et al. [35] combine SecAgg [26] with sparse random projection and count-sketching [111] for compression. Moreover, they add noise using a distributed discrete Gaussian mechanism to generate a differential private output. Beguier et al. [15] combine arithmetic secret-sharing with Top- k sparsification [120] and 1-bit quantization [18]. As we point out in §A.2, both sketching and sparsification are suboptimal for our envisioned cross-device setting given that they require memory and error-feedback on the client side. In contrast, we focus on a dynamic scenario where clients might drop-out at any time and may contribute only once to the training.

Poisoning Attacks & Defenses: FL was also shown vulnerable to manipulations by malicious clients [10, 11, 51, 115, 138]. Targeted/backdoor attacks aim at manipulating the inference results for specific attacker-chosen inputs [10, 138]. In contrast, untargeted poisoning attacks [11, 51, 115] reduce the overall accuracy of the global model. As untargeted attacks are considered to be more severe (given they are harder to detect, cf. §4), we focus on defending those using Byzantine-robust defenses like [99, 137]. In §A.6, we provide a more detailed overview of poisoning attacks and defenses.

Comparison: We compare ScionFL qualitatively in Tab. 1 to the state-of-the-art related work with respect to aggregation and quantization, as well as robustness against poisoning.

Categories	Reference	Technique	M.P.	Quant.	P.R.	Dist. Servers	No Client Interaction
Aggregation	[38]	MPC	✓	✗	✗	✓	✓
	[16]	Masking	✓	✗	✗	✗	✗
Quantization	[123]	–	✗	✓	✗	✗	✓
	[15]	MPC	✓	✗	✗	✓	✓
	[35]	Masking+DP	✓	✓	✗	✗	✗
Robustness	[99]	MPC	✓	✗	✓	✓	✓
	[137]	–	✗	✗	✓	✗	✓
ScionFL	This	MPC	✓	✓	✓	✓	✓

Table 1: High-level comparison of ScionFL and previous works. Notation: MPC—Secure Multi-Party Computation, DP—Differential Privacy, M.P.—Model Privacy, Quant.—Quantization (refers to the schemes where compressed gradients are communicated by the clients to the aggregator(s)), P.R.—Poisoning Resilience, Dist.—Distributed. Since the body of literature is vast, comparison is made against only a subset representing each category.

2 Problem Statement

We now define the precise problem of secure quantized aggregation for FL, which we address in our work. We first introduce the necessary preliminaries on FL and quantization schemes, formalize the

functionality we want to compute securely, and finally define our threat and system model for common (cross-device) FL scenarios.

2.1 Aggregation for Federated Learning

Google introduced federated learning (FL) as a distributed machine learning (ML) paradigm in 2016 [75, 91]. In FL, N data owners collaboratively train a ML model G with the help of a central aggregator S while keeping their input data *locally private*. In each training iteration t , the following three steps are executed:

- (1) The server S randomly selects n out of N available clients and provides the most recent global model G^t .
- (2) Each selected client C_i , $i \in [n]$, sets its local model $w_i^{t+1} = G^t$ and improves it using its local dataset D_i for E epochs (i.e., local optimization steps):

$$w_i^{t+1} \leftarrow w_i^{t+1} - \eta_{C_i} \frac{\partial L(w_i^{t+1}, B_{i,e})}{\partial w_i^{t+1}}, \quad (1)$$

where L is a loss function, η_{C_i} is the clients’ learning rate, and $B_{i,e} \subseteq D_i$ is a batch drawn from D_i in epoch e , where $e \in [E]$. After finishing the local training, C_i sends its local update w_i^{t+1} to S .

- (3) The server updates to a new global model G^{t+1} by combining all w_i^{t+1} with an aggregation rule f_{agg} :

$$G^{t+1} \leftarrow G^t - \eta_S \cdot f_{agg}(w_1^{t+1}, \dots, w_n^{t+1}), \quad (2)$$

where η_S is the server’s learning rate. The most commonly used aggregation rule, which we also focus on, is FedAvg [91]. It averages the local updates as follows:

$$\text{FedAvg}(w_1^{t+1}, \dots, w_n^{t+1}) = \sum_{i=1}^n \frac{|D_i|}{n} w_i^{t+1} \quad (3)$$

This process is repeated until some stopping criterion is met (e.g., a fixed number of training iterations or a certain accuracy is reached).

2.2 Stochastic Quantization

Quantization is a central building block in FL, where data transmission is often a bottleneck. Thus, compressing the (thousands or millions of) gradients is essential to adhere to client bandwidth constraints, reducing training time, and allowing better inclusion and scalability. We now review the desired properties and constructions of quantization schemes that will play a key role in our system design and additional details are provided in §A.1.

Unbiasedness: A well-known and desired design property of gradient compression techniques is being unbiased. That is, for an estimate \hat{w} of a gradient $w \in \mathbb{R}^d$, being unbiased means that $\mathbb{E}[\hat{w}] = w$. Unbiasedness is desired because it guarantees that the mean squared error (MSE) of the mean’s estimation decays linearly with respect to the number of clients, which can be substantial in FL. In FL and other optimization techniques based on stochastic gradient descent (SGD) and its variants (e.g., [69, 85, 91]), the MSE measure (or normalized MSE, a.k.a. NMSE, cf. §3.4) is indeed the quantity of interest since it affects the convergence rate and often the final accuracy of models. In fact, provable convergence rates for the non-convex compressed SGD-based algorithms have a linear dependence on the NMSE. Accordingly, to keep the estimates unbiased, modern quantization techniques employ stochastic quantization (SQ) and its variants to compress the gradients.

1-bit SQ: Our focus is on the appealing communication budget of a single bit for each gradient entry, resulting in a $32\times$ compression ratio compared to regular 32-bit floating point entries. Using 1-bit quantization has been the focus of many recent works concerning distributed and FL network efficiency (e.g., [18, 64, 114, 124, 126]). These works repeatedly demonstrated that a budget of 1-bit per coordinate is sufficient to achieve model accuracy that is competitive to that of a non-compressed baseline.

In particular, 1-bit SQ (i.e., SQ using two quantization values) can be done as follows: For a vector X with m dimensions, the client sends each coordinate as $\sigma_i = \text{Bernoulli}(\frac{X_i - s_X^{\min}}{s_X^{\max} - s_X^{\min}})$, where $s_X^{\max} = \max(X)$ and $s_X^{\min} = \min(X)$. The coordinate is then reconstructed by the receiver as $\tilde{X}_\sigma^i = s_X^{\min} + \tilde{\sigma}_X^i \cdot (s_X^{\max} - s_X^{\min})$. This simple technique results in an unbiased quantization as desired, i.e., $\mathbb{E}[\tilde{X}_\sigma^i] = \mathbb{E}[s_X^{\min} + \tilde{\sigma}_X^i \cdot (s_X^{\max} - s_X^{\min})] = X_i$.

Linear SQ Techniques: A key requirement of being able to perform secure aggregation *efficiently* is being able to aggregate client gradients in their compressed (i.e., quantized) form. The schemes with this property are called “linear” henceforth. Towards this, clients can use a “global scales” approach, where they securely learn the maximal and the minimal value across all gradients, i.e., each client uses $s^{\max} = \max_c \{s_{X_c}^{\max}\}$ and $s^{\min} = \min_c \{s_{X_c}^{\min}\}$.¹ While this approach is appealing due to its simplicity, it has several drawbacks: (i) it requires a preliminary communication stage, (ii) it reveals the global extreme values (even if they are computed securely across all clients), and (iii) the reconstruction error (i.e., the NMSE) is increased as it now depends on the extreme values across all round participants. Accordingly, we also consider a second approach where each client continues to use its own “local scales”. Since plainly using individual scales is not “linear”, i.e., it does not allow for aggregating the quantized gradients efficiently without decoding them, to realize this approach, we use a known approximation [15] and adapt it to our setting (cf. §3.1).

While in this work we focus on the mentioned vanilla SQ, SQ with random rotation [123] as well as SQ with Kashin’s representation [30, 87, 112], our framework seamlessly supports any “linear” quantization scheme, namely, any quantization technique that allows for aggregation in a compressed form (cf. §C.1).

Other Approaches: We acknowledge more recent advances in gradient quantization (e.g., [13, 45, 125, 126]). However, these techniques are not linear and thus less applicable to our framework.

2.3 MPC for Secure Aggregation

Secure computation in the form of multi-party computation (MPC) allows a set of parties to privately evaluate any efficiently computable function on confidential inputs. This paradigm can be utilized to securely run the FedAvg aggregation algorithm [47, 53, 96, 99]: The set of selected FL clients uses additive secret sharing to distribute their sensitive inputs among a set of MPC servers, which resemble a distributed aggregator. The MPC servers securely add the received shares and reconstruct the public result from the resulting shares. In the next iteration, the public model is distributed to a new set of clients chosen at random, and the process is repeated

until the desired accuracy is attained. A visualization of this outsourced MPC setting for secure aggregation is provided in Fig. 1.




Figure 1: Secure aggregation in FL for n clients using outsourced MPC; $\langle w_i^{t+1} \rangle$ denotes the secret-shares of gradient w_i^{t+1} that client C_i has in round $t + 1$; G^t is the model of the previous round and η_S the server-side learning rate.

In the remainder of this paper, we work towards a secure aggregation for FL using FedAvg on *quantized* inputs. In the following, we give our assumptions in terms of the threat model and refine the description of our system model.

MPC Servers: The MPC servers are assumed to be semi-honest. This means, they follow the protocol specification, but may try to learn additional information from the protocol messages that they receive. This well-established assumption, both in theory [7, 22, 61, 66, 93, 99, 102] and in practice [1, 6, 127], is motivated by the fact that companies who run FL with a secure aggregation scheme try to actively protect their clients’ data but want to make sure that someone who monitors or breaks into their systems cannot get plaintext access to the data that is being processed. Additionally, this assumption is justified as organizations usually cannot afford the reputational and monetary risk when being caught betraying their users’ trust. The protocols that we design provide this security guarantee in a dishonest majority setting, where data is protected even when an adversary \mathcal{A} corrupts all except one of the involved servers.. Furthermore, our framework is easily extendable to provide *malicious privacy*, which ensures input privacy even when the corrupted servers try to actively cheat [107].

Malicious Clients: We anticipate that some clients might behave maliciously to negatively affect the quality of the global model with manipulated updates (i.e., poisoning attacks, cf. §4). This is because there are significantly less incentives for clients to behave honestly. Also, due to the sheer number of clients in cross-device FL, it is hard to verify their reputation. We assume an honest majority of clients as is standard in FL, yet our framework is secure even when malicious clients collude with corrupted servers.

Preprocessing Model: We focus on MPC in the preprocessing model [14, 42–44, 72]. This means, we try to shift as much computation and communication as possible to a *data-independent preprocessing* or *offline* phase that can be executed at an arbitrary time before the actual computation. This gives several advantages, e.g., service providers can exploit cheap spot instances. Furthermore, it guarantees faster results when the *data-dependent online phase* of the protocol is executed.

Shared Randomness: We assume that clients and MPC servers have access to a shared randomness source, e.g., by agreeing on a PRNG seed. Such a configuration has been widely employed in MPC protocols [7, 33, 55, 77, 94] as well as in ML systems [12, 123] to optimize communication.

¹This approach resembles “scaler sharing” in TernGrad [130].

2.4 Secure Quantized Aggregation

To introduce the problem of secure quantized aggregation, without loss of generality, we consider a simple stochastic binary quantization scheme to begin with. In this scheme, a m -dimensional vector of the form $\vec{X} = \{x_1, \dots, x_m\}$ comprising of ℓ -bit values will be quantized to obtain a triple $\vec{X}_\sigma = (\vec{\sigma}_X, s_X^{\min}, s_X^{\max})$. Here, $\vec{\sigma}_X$ is an m -dimensional binary vector with a zero at an index indicating the value s_X^{\min} and a one indicating the value s_X^{\max} . Also, s_X^{\min} and s_X^{\max} correspond to the minimum and maximum values in the vector \vec{X} . With this binary quantization (cf. §2.2), the quantized value at the i th dimension, denoted by \vec{X}_σ^i , can be written as

$$\vec{X}_\sigma^i = s_X^{\min} + \vec{\sigma}_X^i \cdot (s_X^{\max} - s_X^{\min}). \quad (4)$$

Before going into the details of aggregation, we provide some of the basic notation that will be utilized throughout the paper.

Notation: $\vec{Y}_{\alpha \times \beta}$ denotes a matrix of dimension $\alpha \times \beta$ with \vec{Y}_i being the i th row and \vec{Y}^j being the j th column. An element in the i th row and j th column is denoted by Y_i^j . Also, $\text{Agg-R}(\vec{Y}_{\alpha \times \beta})$ returns a row vector corresponding to an aggregate of the rows of \vec{Y} . Likewise, $\text{Agg-C}(\vec{Y}_{\alpha \times \beta})$ returns an aggregate of the columns of \vec{Y} .

Given $\vec{U}_{\alpha \times \beta}$ and $\vec{V}_{\alpha \times 1}$, we use $\vec{U} \circ \vec{V}$ to denote the column-wise Hadamard product. Similarly, $\vec{U} \oplus \vec{V}$ denote the sum of two matrices in a column-wise fashion. Concretely, for $\vec{M}_{\alpha \times \beta} = \vec{U} \circ \vec{V}$ and $\vec{N}_{\alpha \times \beta} = \vec{U} \oplus \vec{V}$, we have

$$M_i^j = U_i^j \cdot V_i^j \text{ and } N_i^j = U_i^j + V_i^j, \text{ where } i \in [\alpha], j \in [\beta].$$

Quantized Aggregation: To perform aggregation on quantized inputs, a set of n clients first locally prepares their quantized triples, $(\vec{\sigma}_X, s_X^{\min}, s_X^{\max})$, corresponding to their locally trained ML model updates (i.e., gradients) and submits these to a parameter server for aggregation. Let m be the dimension of the underlying ML model. The quantized triples can then be consolidated to a matrix triple of the form $(\vec{B}_{n \times m}, \vec{U}_{n \times 1}, \vec{V}_{n \times 1})$. Here, \vec{B} is a binary matrix that corresponds to the $\vec{\sigma}_X$ vector of the clients. Similarly, \vec{U} and \vec{V} correspond to the s_X^{\min} and s_X^{\max} values of the above-mentioned triple. The aggregate of the quantized inputs is defined as

$$\vec{X}_{1 \times m} = \text{Agg-R}(\vec{U}_{n \times 1} \oplus \vec{B}_{n \times m} \circ (\vec{V}_{n \times 1} - \vec{U}_{n \times 1})). \quad (5)$$

Functionality $\mathcal{F}_{\text{SecAgg}}$

$\mathcal{F}_{\text{SecAgg}}$ interacts with all the τ servers in \mathcal{S} and an adversary \mathcal{A} that controls a subset of the servers in \mathcal{S} .

Input: $\mathcal{F}_{\text{SecAgg}}$ receives $\langle \cdot \rangle$ -shares of the matrix triple $(\vec{B}_{n \times m}, \vec{U}_{n \times 1}, \vec{V}_{n \times 1})$ from the honest servers in \mathcal{S} , while the adversary \mathcal{A} provides the $\langle \cdot \rangle$ -shares on behalf of the corrupt servers. Here $B_i^j \in \{0, 1\}$ and $U_i^j, V_i^j \in \mathbb{Z}_{2^\ell}$.

Computation: $\mathcal{F}_{\text{SecAgg}}$ reconstructs $(\vec{B}, \vec{U}, \vec{V})$ from its $\langle \cdot \rangle$ -shares.

- Set $\vec{S}_{n \times 1} = \vec{V}_{n \times 1} - \vec{U}_{n \times 1}$ and compute $\vec{W}_{n \times m} = \vec{B}_{n \times m} \circ \vec{S}_{n \times 1}$.
- Compute $\vec{X}_{n \times m} = \vec{U}_{n \times 1} \oplus \vec{W}_{n \times m}$.
- Compute $\vec{Y}_{1 \times m} = \text{Agg-R}(\vec{X}_{n \times m})$, i.e., $\vec{Y}^j = \sum_{i=1}^n X_i^j$ for $j \in [m]$.

Output: $\mathcal{F}_{\text{SecAgg}}$ computes $\langle \cdot \rangle$ -shares of \vec{Y} and sends the respective shares to the servers in \mathcal{S} .

Figure 2: Ideal functionality for semi-honest secure quantized aggregation for linear stochastic binary quantization.

Ideal Functionality: To perform secure aggregation of quantized updates using MPC, we model the aggregation as an ideal functionality $\mathcal{F}_{\text{SecAgg}}$, which is provided in Fig. 2. We consider a set of τ servers to which the clients secret-share their quantized updates. The goal is to compute the aggregate of the inputs as shown in Eq. 5. Let $\langle \cdot \rangle$ denote the underlying secret sharing scheme. Looking ahead, we will use linear secret sharing techniques for MPC, in which linear operations such as addition and subtraction are local. As a result, we will concentrate on efficiently computing the column-wise Hadamard product.

3 Our Framework: ScionFL

We now detail our framework “ScionFL” from an MPC standpoint. We present the underlying sharing semantics, client interaction with MPC servers, and methods for performing secure aggregation of client updates. Our generic constructions utilize MPC in a black-box fashion [17, 40, 41, 44, 110], however, the full MPC protocols, including inner product computation, multiplication, and bit-to-arithmetic conversion, are detailed in §B.1.

Masked Evaluation: In our MPC protocols, we use the masked evaluation technique [17, 34, 59, 77, 90, 122], which enables costly *data-independent* calculations to be completed in a preprocessing phase, thus enabling a fast and efficient *data-dependent* on-line phase (cf. §2.3). In this model, the secret-share of every element $v \in \mathbb{Z}_{2^\ell}$, denoted by $\langle v \rangle$, is associated with two values: a random mask $\lambda_v \in \mathbb{Z}_{2^\ell}$ and a masked value $m_v \in \mathbb{Z}_{2^\ell}$ such that $v = m_v + \lambda_v$. While λ_v is split and distributed as q shares as per the underlying MPC scheme (cf. §B.1), m_v is given to all MPC servers.² Since λ_v is random and independent of v , all operations involving λ_v values alone can be computed during the preprocessing phase and thereby leading to a fast online phase.

Client Interaction: Before going into the details of aggregation among the τ MPC servers, we discuss input sharing and the reconstruction of the aggregated vector for clients.

To generate the $\langle \cdot \rangle$ -shares of a value $v \in \mathbb{Z}_{2^\ell}$ owned by client C , it first non-interactively computes the additive shares of the mask λ_v using the shared randomness setup discussed in §2.3. The masked value is then computed as $m_v = v - \lambda_v$ and sent to a single designated MPC server, say S_1 . The input sharing is completed when S_1 sends m_v to all remaining MPC servers.³ For Boolean values (i.e., in \mathbb{Z}_2) the procedure is similar except that addition/subtraction is replaced with XOR and multiplication with AND. We use $\langle \cdot \rangle^B$ to denote the secret sharing over the domain \mathbb{Z}_2 .

After the servers have received all inputs in $\langle \cdot \rangle$ -shared form, they instantiate the $\mathcal{F}_{\text{SecAgg}}$ functionality specified in Fig. 2 and obtain the aggregated vectors in $\langle \cdot \rangle$ -shared form. Recall from §2.1 that the values to be aggregated in our case correspond to FL gradients, and the aggregated result can also be made public. As a result, the servers reconstruct the aggregated result towards a chosen server, say S_1 , which updates the global model according to Eq. (3). In the next iteration, S_1 distributes the updated global model to a fresh set of clients, and the process is repeated.

²Due to differences in the underlying setting, there may be minor differences in how the values m_v and λ_v are distributed among the servers.

³If *malicious privacy* is desired, C can send a hash digest of all the m_v values to the remaining MPC servers, who can verify the correctness of messages from S_1 .

From a client’s perspective, it is only interacting with a single server (apart from a one-time shared-randomness setup), as is the case in the privacy-free variant with a single parameter server [91].

3.1 MPC-based Aggregation

We now discuss three approaches for instantiating $\mathcal{F}_{\text{SecAgg}}$ using MPC protocols that operate on secret-shared values. Recall from Eq. (5) that the MPC servers for the aggregation of quantized values possess $\langle \cdot \rangle$ -shares of matrices $\tilde{\mathbf{U}}_{n \times 1}$ and $\tilde{\mathbf{V}}_{n \times 1}$ along with the $\langle \cdot \rangle^{\mathbf{B}}$ -shares of $\tilde{\mathbf{B}}_{n \times m}$.

Approach I: A naive instantiation of $\mathcal{F}_{\text{SecAgg}}$ would be to have the servers convert binary shares of the matrix $\tilde{\mathbf{B}}$ to their arithmetic shares first, as in Prio+ [2]. This is possible in MPC with a bit-to-arithmetic conversion protocol [2, 77, 94, 104], denoted by Π_{BitA} , which computes the arithmetic shares of $b \in \mathbb{Z}_2$ from its Boolean sharing. Once the arithmetic shares are generated, the servers use an instance of an inner-product protocol, denoted by Π_{IP} , on each of the m columns of matrix $\tilde{\mathbf{B}}$ with the column vector $(\tilde{\mathbf{V}}_{n \times 1} - \tilde{\mathbf{U}}_{n \times 1})$ computed locally to obtain the row aggregate. They complete the protocol by adding a row aggregate of $\tilde{\mathbf{U}}$ to each column of the matrix computed in the previous step. The formal protocol $\Pi_{\text{SecAgg}}^{\text{I}}$ is given in Fig. 3.




Figure 3: Secure aggregation – Approach I.

As shown in Fig. 3, $\Pi_{\text{SecAgg}}^{\text{I}}$ invokes Π_{BitA} for each bit in matrix $\tilde{\mathbf{B}}$, resulting in $n \cdot m$ invocations. Using the masked evaluation technique [78, 122], the online communication of the inner product protocol Π_{IP} can be made independent of the dimension n , which in our case corresponds to the number of clients.

Approach II: We use the bit injection protocol [76, 94, 104], denoted by Π_{BI} , which computes the arithmetic sharing of $b \cdot s$ given the Boolean sharing of $b \in \mathbb{Z}_2$ and the arithmetic sharing of $s \in \mathbb{Z}_\ell$. Given $\langle \tilde{\mathbf{M}}_{\alpha \times 1} \rangle^{\mathbf{B}}$ and $\langle \tilde{\mathbf{N}}_{\alpha \times 1} \rangle$, the high-level idea is to effectively combine the Π_{BitA} and Π_{IP} protocol to a slightly modified instance of Π_{BI} that directly computes the sum [78, 122] (denoted by $\sum_{i=1}^{\alpha} \mathbf{M}_i^1 \cdot \mathbf{N}_i^1$) instead of the individual positions. This can be achieved following Eq. (16) and the details appear in Fig. 17 in §B.1. One significant advantage of this technique is that the overall online communication is no longer dependent on the number of clients n . $\Pi_{\text{SecAgg}}^{\text{II}}$ denotes the resulting protocol and the formal details are given in Fig. 4.




Figure 4: Secure aggregation – Approach II.

Approach III using SepAgg: In a closely related work [15], the authors combine the SIGNSGD compression technique of [18] with additive secret sharing for FL. Unlike our work, which investigates secure aggregation using various linear quantization algorithms in a cross-device environment, they aim for a cross-silo setting in which clients distribute arithmetically shared values to servers rather than single bits. In terms of client-server communication, this indicates a non-optimal communication overhead of factor $\log_2 2n$, where n is the number of involved participants each round.

However, the authors of [15] introduce an interesting approximation called “SepAgg” for computing the averaged inner product between bits and scales:

$$\frac{1}{n} \sum_i \tilde{\sigma}_X^i \cdot \tilde{s}_X^i \approx \frac{1}{n^2} \left(\sum_i \tilde{\sigma}_X^i \right) \left(\sum_i \tilde{s}_X^i \right). \quad (6)$$

We adopt the SepAgg method to our setting to compute $\text{Agg-R}(\tilde{\mathbf{B}} \circ (\tilde{\mathbf{V}} - \tilde{\mathbf{U}}))$ in Eq. (5). In particular, we aggregate the matrices $\tilde{\mathbf{B}}$ and $(\tilde{\mathbf{V}} - \tilde{\mathbf{U}})$ independently and then perform one secure multiplication per coordinate, with the other operations being linear and free in our MPC protocol. As a result, we can utilize linear quantization schemes with local scales at the cost of global scales (ignoring the overhead for global scales to securely determine s_X^{min} and s_X^{max} across all participants). The formal protocol $\Pi_{\text{SecAgg}}^{\text{III}}$ utilizing SepAgg appears in Fig. 5.




Figure 5: Secure aggregation – Approach III (SepAgg [15]).

Accuracy Evaluation: We next provide strong empirical evidence that applying SepAgg for SQ with preprocessing preserves the linear NMSE decay with respect to the number of clients (i.e., unbiased estimates). For this, we simulate the aggregation of random vectors \tilde{v}_i with dimension d drawn from a $(0, 1)$ -log-normal distribution.⁴ Then, we measure the normalized mean square error (NMSE) when comparing the averaged aggregation result agg computed on secret-shared and quantized inputs to the plain averaged aggregation $agg_{\text{orig}} = \sum_i \tilde{v}_i / n$. Concretely, we measure

$$\text{NMSE} = \frac{\|agg_{\text{orig}} - agg\|_2^2}{\sum_i \|\tilde{v}_i\|_2^2 / n}, \quad (7)$$

where agg is computed using various linear quantization schemes for (i) a regular dot product between converted bits and scales and (ii) with SepAgg. Our results shown in Fig. 6 are the average of 10 trials for each experiment with $q = 3$ shares (representing a three-server dishonest majority setting using the masked evaluation technique). While for smaller dimensions we observe a visible effect for SQ without preprocessing, there is only a minor difference for the other two quantization schemes, and sometimes the NMSE for SepAgg is even smaller than for the exact computation.

⁴We use this distribution for preliminary measurements as it was commonly observed in neural network gradients, e.g., [37].

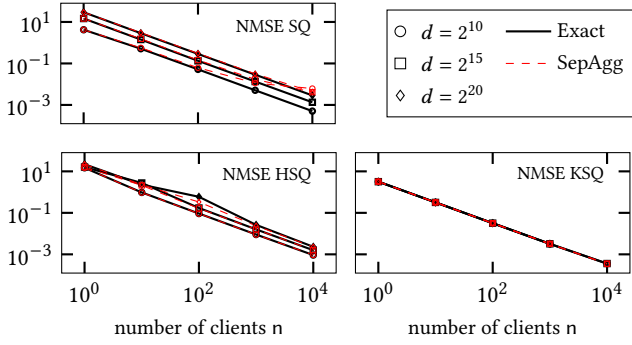


Figure 6: NMSE comparison between exact and SepAgg-based aggregation for vanilla SQ, SQ using the randomized Hadamard transform (HSQ), and SQ using Kashin's representation (KSQ) for various vector dimensions d and number of clients n .

Formally proving that applying SepAgg after different preprocessing techniques (e.g., random rotation and Kashin's representation) results in an unbiased aggregation is a significant theoretical challenge. We leave this intriguing endeavor for future work.

Communication Costs: Tab. 2 provides the theoretical communication costs for our approaches when aggregating n quantized single-dimension vectors. Clearly, our Approach-III (cf. Fig. 5) is the most efficient, with the multiplication-related cost being completely independent of the number of clients n due to SepAgg [15]. The concrete communication costs are provided in §3.5.

The communication costs are primarily determined by the cost of BitA^{pre} . In our masked evaluation technique (cf. §B.1), this relates to the conversion of a random secret-shared bit from Boolean sharing to its additive sharing form [41, 110]. Thus, we propose approximate variants for improving the cost of this operation.

Approach	Offline	Online
Approach-I	$n \cdot \text{BitA}^{\text{pre}} + n \cdot \text{Mult}^{\text{pre}}$	$n \cdot \text{BitA}^{\text{on}} + \text{Mult}^{\text{on}}$
Approach-II	$n \cdot \text{BitA}^{\text{pre}} + n \cdot \text{Mult}^{\text{pre}}$	Mult^{on}
Approach-III	$n \cdot \text{BitA}^{\text{pre}} + \text{Mult}^{\text{pre}}$	Mult^{on}

Table 2: Communication costs for aggregating quantized vectors with a single dimension for n clients. Protocols Π_{BitA} and Π_{Mult} are treated as black-boxes, and their costs are represented as BitA and Mult , respectively. The superscript pre in the costs denotes preprocessing and on denotes the online phase.

3.2 Approximate Bit Conversion in MPC

We reduce our communication and computation costs using a novel *approximate* bit conversion method. For this, consider a bit b represented using two Boolean shares $b_1, b_2 \in \{0, 1\}$, such that $b = b_1 \oplus b_2$. Note that when embedding b_1 and b_2 in a larger field/ring⁵, it holds that $b = b_1 + b_2 - 2b_1b_2$. Similarly, for $b = b_1 \oplus b_2 \oplus b_3$, $b = b_1 + b_2 + b_3 - 2b_1b_2 - 2b_1b_3 - 2b_2b_3 + 4b_1b_2b_3$ holds true. In fact, this concept can be generalized to an arbitrary number of shares, denoted by q , as discussed below.

For $b = \bigoplus_{i=1}^q b_i$, let $Q = \{b_i\}_{i \in [q]}$ denote the set of all q shares of b , and \tilde{b}_i the arithmetic equivalent of the share b_i . Let 2^Q be the powerset of Q and $Q^{|c|}$ the set of all size- c subsets in 2^Q , that is,

⁵The bit (either 0 or 1) is treated as a ring element in \mathbb{Z}_{2^ℓ} in our protocols.

$2^Q = \sum_{i=0}^q Q^{|i|}$. The arithmetic equivalent of b , denoted by \tilde{b} , is given as

$$\begin{aligned} \tilde{b} &= \sum_{\{b_e\} \in Q^{|1|}} \tilde{b}_e - 2 \cdot \sum_{\{b_{e_1}, b_{e_2}\} \in Q^{|2|}} \tilde{b}_{e_1} \tilde{b}_{e_2} + 4 \cdot \sum_{\{b_{e_1}, b_{e_2}, b_{e_3}\} \in Q^{|3|}} \tilde{b}_{e_1} \tilde{b}_{e_2} \tilde{b}_{e_3} \dots + (-2)^{q-1} \cdot \prod_{\{b_{e_1}, \dots, b_{e_q}\} \in Q^{|q|}} \tilde{b}_e \\ &= \sum_{k=1}^q (-2)^{k-1} \sum_{\{b_{e_1}, \dots, b_{e_k}\} \in Q^{|k|}} \tilde{b}_{e_1} \tilde{b}_{e_2} \dots \tilde{b}_{e_k} \end{aligned} \quad (8)$$

Note that the Eq. (8) can be viewed as sum of three terms: Sum (term_s), Middle (term_m), and Product (term_p) as shown in Eq. (9) below. (Note that $Q^{|q|} = Q$).

$$\tilde{b} = \underbrace{\sum_{\{b_e\} \in Q^{|1|}} \tilde{b}_e}_{\text{Sum Term: } \text{term}_s} + \underbrace{\sum_{k=2}^{q-1} (-2)^{k-1} \sum_{\{b_{e_1}, \dots, b_{e_k}\} \in Q^{|k|}} \tilde{b}_{e_1} \tilde{b}_{e_2} \dots \tilde{b}_{e_k}}_{\text{Middle Term: } \text{term}_m} + \underbrace{(-2)^{q-1} \prod_{b_e \in Q} \tilde{b}_e}_{\text{Product Term: } \text{term}_p} \quad (9)$$

Our Approach. Performing this conversion in MPC requires many additions and multiplications. While linear operations like additions can be calculated for “free” in most MPC protocols, non-linear operations such as multiplications require some form of communication between the MPC servers. Hence, computing the middle term is costly, especially when a large number of shares is involved.

To approximate \tilde{b} in Eq. (9), we replace only the middle term term_m with its expected value $\mathbb{E}[\text{term}_m]$ such that the approximate value of \tilde{b} , denoted by \hat{b} , retains $\mathbb{E}[\hat{b}] = b$. The expectation of term_s and term_p in Eq. (9) is first calculated, and $\mathbb{E}[\text{term}_m]$ is inferred using the fact that $\mathbb{E}[\hat{b}] = b$. This analysis is summarised in Lem. 3.1.

Lemma 3.1 (Expected Values). *Given a bit $b = \bigoplus_{i=1}^q b_i$ and $b = \text{term}_s + \text{term}_m + \text{term}_p$ with*

$$\begin{aligned} \text{term}_s &= \sum_{\{b_e\} \in Q^{|1|}} \tilde{b}_e, \quad \text{term}_m = \sum_{k=2}^{q-1} (-2)^{k-1} \sum_{\{b_{e_1}, \dots, b_{e_k}\} \in Q^{|k|}} \tilde{b}_{e_1} \tilde{b}_{e_2} \dots \tilde{b}_{e_k}, \\ \text{term}_p &= (-2)^{q-1} \prod_{b_e \in Q} \tilde{b}_e, \end{aligned}$$

we have $\mathbb{E}[\text{term}_s \mid b] = q/2$, $\mathbb{E}[\text{term}_m \mid b] = (q-1) \bmod 2 - q/2$, and $\mathbb{E}[\text{term}_p \mid b] = b - (q-1) \bmod 2$.

PROOF. For the analysis, we use the truth table of b , denoted by T_b , which has 2^q rows. Half of the rows in T_b correspond to $b = 0$, while the other half correspond to $b = 1$. The truth table for three shares ($q = 3$) is given in Tab. 3 as a reference.

Sum Term (term_s): For each row of the form (b_1, \dots, b_q) in T_b , term_s equals $\tilde{b}_1 + \dots + \tilde{b}_q$, which can be interpreted as the number of \tilde{b}_i 's selected out of the q possible. Furthermore, there are a total of $\binom{q}{k}$ rows with sums equal to k , with k being odd corresponding to the row for $b = 1$ and k being even corresponding to the row for $b = 0$. As a result, given $b = 0$, the expectation of the sum term can be calculated as the product of $1/2^{q-1}$ (corresponding to rows in T_b with $b = 0$) and the sum of terms of the form $k \cdot \binom{q}{k}$ with k being even. Using Lem. B.1 in §B, we get

$$\mathbb{E}[\text{term}_s \mid (b = 0)] = \frac{1}{2^{q-1}} \cdot \sum_{k=0}^{\lfloor q/2 \rfloor} 2k \binom{q}{2k} = \frac{1}{2^{q-1}} \cdot q \cdot 2^{q-2} = q/2.$$

Similarly, we obtain $\mathbb{E}[\text{term}_s \mid (b = 1)] = q/2$. To summarize, we have $\mathbb{E}[\text{term}_s \mid b] = q/2$.

b	b ₁	b ₂	b ₃	term _s	term _m	term _p	\tilde{b}
0	0	0	0	0	0	0	0
1	0	0	1	1	0	0	1
1	0	1	0	1	0	0	1
0	0	1	1	2	-2	0	0
1	1	0	0	1	0	0	1
0	1	0	1	2	-2	0	0
0	1	1	0	2	-2	0	0
1	1	1	1	3	-6	4	1

Table 3: Truth table for $b = b_1 \oplus b_2 \oplus b_3$. The rows corresponding to $b = 0$ are highlighted. \tilde{b} denotes the arithmetic equivalent of b .

Product Term (term_p): The product of all the q shares will be 1 only if all the shares are 1, otherwise it will be 0. Moreover, all shares of b being 1 correspond to $b = 1$ if q is odd, and $b = 0$ otherwise. Now, when q is odd then $\text{term}_p = (-2)^{q-1}$ with probability $\frac{1}{2^{q-1}}$ (when all the shares of b are 1, given that at least one share is 1 as we are in the case $b = 1$), and 0 otherwise. In this case, we can write

$$\begin{aligned} \mathbb{E}[\text{term}_p \mid (b = 0 \wedge q \text{ is odd})] &= \frac{1}{2^{q-1}} \cdot 0 = 0 \\ \mathbb{E}[\text{term}_p \mid (b = 1 \wedge q \text{ is odd})] &= \frac{1}{2^{q-1}} \cdot (-2)^{q-1} = 1 \end{aligned}$$

Similarly, the case for even q can be written as

$$\begin{aligned} \mathbb{E}[\text{term}_p \mid (b = 0 \wedge q \text{ is even})] &= \frac{1}{2^{q-1}} \cdot (-2)^{q-1} = -1 \\ \mathbb{E}[\text{term}_p \mid (b = 1 \wedge q \text{ is even})] &= \frac{1}{2^{q-1}} \cdot 0 = 0 \end{aligned}$$

The above observation can be summarized as $\mathbb{E}[\text{term}_p \mid b] = b - (q-1) \bmod 2$.

Middle Term (term_m): Given $\mathbb{E}[b] = b$, $\mathbb{E}[\text{term}_s \mid b]$ and $\mathbb{E}[\text{term}_p \mid b]$, the expectation of term_m can be calculated as

$$\begin{aligned} \mathbb{E}[\text{term}_m \mid b] &= \mathbb{E}[b] - \mathbb{E}[\text{term}_s \mid b] - \mathbb{E}[\text{term}_p \mid b] \\ &= b - q/2 - (b - (q-1) \bmod 2) = (q-1) \bmod 2 - q/2. \end{aligned}$$

This concludes the proof of Lem. 3.1. \square

Our Approximation. With the aforementioned observations, the approximate arithmetic equivalent of b , denoted by \hat{b} , is as follows:

$$\hat{b} = \underbrace{\sum_{b_e \in \mathcal{Q}} \tilde{b}_e}_{\text{term}_s} + \underbrace{\left((q-1) \bmod 2 - \frac{q}{2} \right)}_{\text{term}_m^a} + \underbrace{(-2)^{q-1} \prod_{b_e \in \mathcal{Q}} \tilde{b}_e}_{\text{term}_p} \quad (10)$$

While term_s is kept because it only involves linear operations on the shares of b (which are free in MPC for any linear secret sharing scheme), we observe that term_p is required to keep the expected values for $b = 0$ and $b = 1$ different. This is evident from Lem. 3.1 where $\mathbb{E}[\text{term}_p \mid b]$ is the only term that depends on b .

In general, if a term that depends on all the q shares of b is missing from the approximation, we get $\mathbb{E}[b = 0] = \mathbb{E}[b = 1]$. The intuition is that only such a term can differentiate between $b = 0$ and $b = 1$, while all other terms will be symmetrically distributed. For instance, consider $q = 3$ and let $\tilde{b} = c_1 b_1 + c_2 b_2 + c_3 b_3 + c_4 b_1 b_2 + c_5 b_2 b_3 + c_6 b_1 b_3 + c_7$ for some random combiners $c_i \in \mathbb{Z}_{2^t}$ and $i \in [7]$. Using the truth table T_b given in Tab. 3, it is easy to verify that

$$\mathbb{E}[\tilde{b} = 0] = \mathbb{E}[\tilde{b} = 1] = \frac{1}{4} \cdot (2c_1 + 2c_2 + 2c_3 + c_4 + c_5 + c_6 + 4c_7) \quad (11)$$

This argument can be generalized to any value of q .

Claim 3.2. *The approximate arithmetic equivalent \hat{b} in Eq. (10) preserves the expectation of the exact bit b in Eq. (8), i.e., $\mathbb{E}[\hat{b} = 0] = 0$ and $\mathbb{E}[\hat{b} = 1] = 1$.*

PROOF. The proof is straightforward as we replace the middle term (term_m) in Eq. (8) with its expected value term_m^a . \square

Efficiency Analysis: We measure the efficiency gains achieved by our approximation method, discussed in §3.2, by counting the number of *cross terms*⁶ that must be computed securely using MPC. Cross terms are terms that compute the product of two or more shares. While the exact amount of computation and communication varies depending on the MPC protocol and setting (e.g., honest vs. dishonest majority or semi-honest vs. malicious security), we believe cross terms can provide a protocol-independent and realistic assessment of scalability.⁷

Computation	#cross-terms	
	Exact (\tilde{b})	Approximate (\hat{b})
Bit-to-Arithmetic	$2^q - q - 1$	1
Bit Injection	$2^q + q^2 - 2q - 1$	$q^2 - q + 1$

Table 4: Efficiency analysis via approximate bit conversion with respect to the #cross-terms involved.

Tab. 4 provides details regarding the number of cross terms involved in obtaining the arithmetic equivalent of $b = \bigoplus_{i=1}^q b_i$. The gains increase significantly with a higher number of shares q due to the exponential growth in the number of cross terms for the exact computation. Tab. 4 also provides details for a bit injection operation in which the product of a Boolean bit b and a scale value s is securely computed. Given $s = \sum_{i=1}^q s_i$, the value $b \cdot s$ can be computed by first computing either \tilde{b} or \hat{b} (depending on whether an exact or approximate value is required) and then multiplying by s .

3.3 Secure Bit Aggregation with Global Scales

Here, we consider the secure bit aggregation problem in the context of “global scales”, as discussed in §2.2. In this case, all the clients use the same set of scales for quantization, denoted by s_G^{\min} and s_G^{\max} . Therefore, it is sufficient to compute

$$\vec{X}_{1 \times m} = s_G^{\min} \oplus \text{Agg-R}(\vec{B}_{n \times m}) \circ (s_G^{\max} - s_G^{\min}) \quad (12)$$

as the aggregation result. Interestingly, when $s_G^{\min} = 0$ and $s_G^{\max} = 1$, this can also be viewed as an instance of privacy-preserving aggregate statistics computation, as demonstrated in the works of Prio [38] and Prio+ [2].

As shown in Eq. (12), the computation becomes simpler in the case of global scales since all clients utilize the same set of public scales, denoted by s^{\min} and s^{\max} , to compute their quantized vector that corresponds to the rows of \vec{B} . Hence, we just need to compute the column-wise aggregate of the \vec{B} matrix and use protocol $\Pi_{\text{BitA}}^{\text{sum}}$ (Fig. 16 in §B.1) to do so. The resulting protocol $\Pi_{\text{SecAgg}}^{\text{Global}}$ appears in Fig. 7.

⁶Terms for which interaction among MPC servers is necessary.

⁷We acknowledge that the analysis cannot provide an exact comparison, owing to the presence of the product term in the approximation. e.g., depending on the underlying MPC setup, the product term (term_p) may require more communication than the middle terms (term_m), and therefore the effect of approximation may be minimized.

Protocol $\Pi_{\text{SecAgg}}^{\text{Global}}$ ($(\langle \tilde{\mathbf{B}}_{n \times m} \rangle^B, s^{\min}, s^{\max})$)

1. Compute $\langle \tilde{\mathbf{W}}^j \rangle = \Pi_{\text{BitA}}^{\text{sum}}(\langle \tilde{\mathbf{B}}^j \rangle^B)$, for each $j \in [m]$.
2. Locally compute $\langle \tilde{\mathbf{Y}}_{1 \times m} \rangle = s^{\min} \oplus (\langle \tilde{\mathbf{W}}_{1 \times m} \rangle \cdot (s^{\max} - s^{\min}))$.

Figure 7: Secure aggregation – Global Scales.

3.4 Accuracy Evaluation

In §3.2, we showed that our approximate bit conversion preserves the expectation of the exact bits. However, we also want to understand the concrete accuracy impact on the aggregation result due to the increased variance. For this, we run a simulation similar to the one described in §3.1. Here, we compare the NMSE computed as in Eq. (7) for an aggregation *agg* when using various linear quantization schemes with global scales (i) with an exact bit-to-arithmetic conversion and (ii) with our approximation enabled. Our results shown in Fig. 8 are the average of 10 trials for each experiment with $q = 3$ shares. Consistently, we observe that applying our approximation increases the NMSE by about three orders of magnitude for stochastic quantization without rotation, and by about one and a half orders of magnitude for rotation-based algorithms.

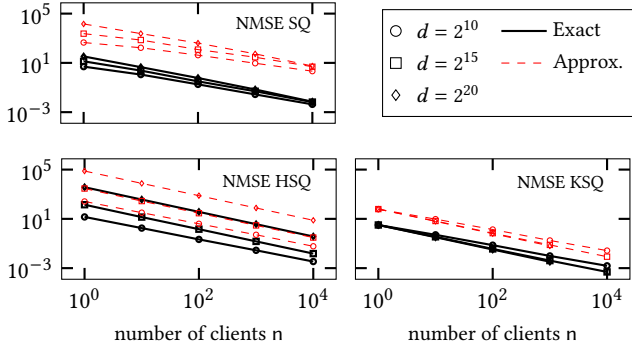


Figure 8: NMSE comparison between exact and approximation-based aggregation for vanilla SQ, SQ using the randomized Hadamard transform (HSQ), and SQ using Kashin’s representation (KSQ) for global scales with $q = 3$ shares and various vector dimensions d and number of clients n .

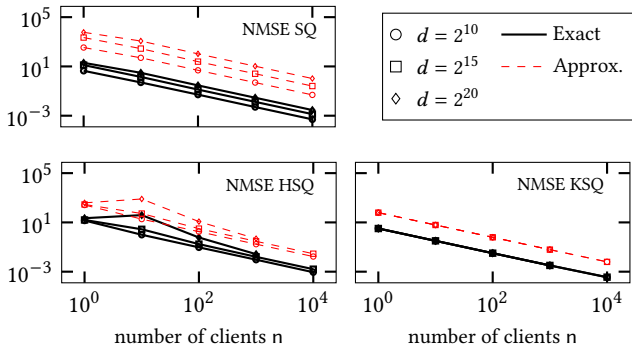


Figure 9: NMSE comparison between exact and approximation-based aggregation for SQ, Hadamard SQ (HSQ), and Kashin SQ (KSQ) for local scales with $q = 3$ shares, various vector dimensions d , and number of clients n .

In Fig. 9, we provide results considering local scales. In contrast to global scales, we can observe that for stochastic quantization without rotation the effect on the NMSE is reduced from three to one order of magnitude. Also, for rotation-based algorithms there are significant concrete improvements. Furthermore, as shown in §3.6, the error is still so small that the impact on the accuracy in common FL settings is negligible.

3.5 Detailed Communication Costs

Here, we provide more insights into the concrete communication costs for our secure aggregation protocols described in §3.1.

n	Method	Exact		Approx.	
		Offline	Online	Offline	Online
20	Approach-I	644.50	1.70	620.27	1.70
	Approach-II	644.50	0.59	620.27	0.59
	Approach-III	89.77	0.59	65.54	0.59
100	Approach-I	3222.51	6.12	3101.36	6.12
	Approach-II	3222.51	0.59	3101.36	0.59
	Approach-III	332.08	0.59	210.93	0.59
500	Approach-I	16112.56	28.24	15506.80	28.24
	Approach-II	16112.56	0.59	15506.80	0.59
	Approach-III	1543.62	0.59	937.85	0.59

Table 5: Inter-server communication per round in MiB for our MNIST/LeNet benchmark for different numbers of clients n per round. Training is done using 1-bit SQ with Kashin’s representation (KSQ). We compare Approach-I (cf. Fig. 3 in §3.1), Approach-II (cf. Fig. 4 in §3.1), and Approach-III (cf. Fig. 5 in §3.1). Additionally, we distinguish between using an exact bit-to-arithmetic conversion and our approximation (cf. §3.2).

In Tab. 5 we provide the detailed communication costs for the secure aggregation approaches discussed in §3.1 when training the LeNet architecture for image classification on the MNIST data set [82] using 1-bit SQ with Kashin’s representation [30]. We instantiate the OT instances required in the preprocessing phase, as discussed in §B.1, with silent OT [39], following Prio+ [2]. Here, we can observe the significant impact of including SepAgg [15] in practice with performance improvements between Approach-II and Approach-III of up to 16.6× in the offline phase.

Approach	$n = 10^2$	$n = 10^3$	$n = 10^4$	$n = 10^5$
Prio+ [2]	9.45	94.50	945.04	9450.44
Approach-III (Exact)	3.94	39.42	394.17	3941.66
Approach-III (Approx.)	2.37	23.75	237.45	2374.53

Table 6: Total communication in MiB of our Approach-III (cf. Fig. 5 in §3.1) compared to Prio+ [2] to calculate the sum of bits for different numbers of clients n and dimension $m = 1000$. For our Approach-III, we distinguish between using an exact bit-to-arithmetic conversion as in Prio+ [2] and our approximation (cf. §3.2).

In Tab. 6, we compare the aggregation of bits (i.e., when not considering quantized inputs that require scale multiplication and hence without SepAgg [15] being applicable) to Prio+ [2]. For a fair comparison, we translate the approach in Prio+ [2] to our three party dishonest-majority setting. As we can see, even for exact bit-to-arithmetic conversion, we improve over Prio+ by factor 2.4× for $n = 10^5$. When applying our approximate bit-to-arithmetic conversion (cf. §3.2), this improvement increases to a factor of 4×.

3.6 Performance Evaluation

We implemented an extensive end-to-end FL evaluation and MPC simulation framework. We first describe our implementation, then the parameters used for our accuracy evaluation, and finally the obtained results.

Implementation: Our implementation is written in Python based on PyTorch. It supports multi-GPU acceleration, also for our MPC simulation. We used a subset of this framework for measuring the accuracy of SepAgg (cf. §3.1) and our approximate bit conversion (cf. §3.4), and we will describe extensions in §4.2 to incorporate evaluations of poisoning attacks and defenses.

Our framework provides a command-line interface to run FL training tasks and observe the resulting training as well as test accuracy. Upon execution, the framework distributes training data among the specified number of virtual clients that locally perform training. The server(s) perform aggregation using FedAvg. When the MPC simulation is enabled, the clients’ input will be secret-shared before aggregation and the protocol described in §3.1 will be executed locally. Note that our goal is not to assess the run-time performance of the MPC protocol but rather precisely measure the impact on accuracy. Our implementation supports all exact and approximate secure aggregation variants described in this paper.

Parameters: We evaluate the accuracy on the following standard FL tasks for image classification: training (i) LeNet on MNIST [82] for 1000 rounds and (ii) ResNet9 on CIFAR-10 [79] for 8000 rounds. For all tasks, we set a client batch size of 8, a learning rate of 0.05, and perform 5 local client train steps per round. For MNIST, we run training using $N \in \{200, 1000, 5000\}$ clients and choose 10% of the clients at random per round. Due to the memory constraints of our system (that simulates all clients at once), we restrict training for CIFAR10 to $N = 1000$ clients and select $n = 40$ per round. As we observed a significant loss in accuracy for plain SQ in our accuracy evaluation for approximate bit conversion as well as SepAgg (cf. §3.4), we focus our evaluation on more accurate linear quantization schemes, i.e., HSQ and KSQ. For the MPC simulation of our approximate secure aggregation following Approach-III (cf. Fig. 5), we choose a three-server dishonest majority setting.

Results: The results for the MNIST/LeNet training are given in Fig. 10. Validation accuracy for our approximate version converges to almost the same final accuracy as the insecure exact aggregation. Specifically, in the final round of training, the difference between the two is diminished to 0.77% and 0.33% for HSQ and KSQ for $N = 5000$, respectively. Similar observations apply to CIFAR10/ResNet9 in Fig. 11. However, here the difference between the exact and approximate version for KSQ is higher with 3.14% in the final round. This gap is expected due to the significantly lower number of clients per round, for which our approximate bit conversion and SepAgg technique result in a comparatively high NMSE over the baseline (cf. Figs. 6 and 8). We expect this effect to vanish for a real cross-device setting with thousands of participants per round (due to the demonstrated linear decay of the NMSE when increasing n), which we unfortunately cannot simulate with complex model architectures due to hardware limitations. Additionally, one may use a *hybrid approach*, where training uses the approximate version for initial rounds until a baseline accuracy is reached, whereas secure

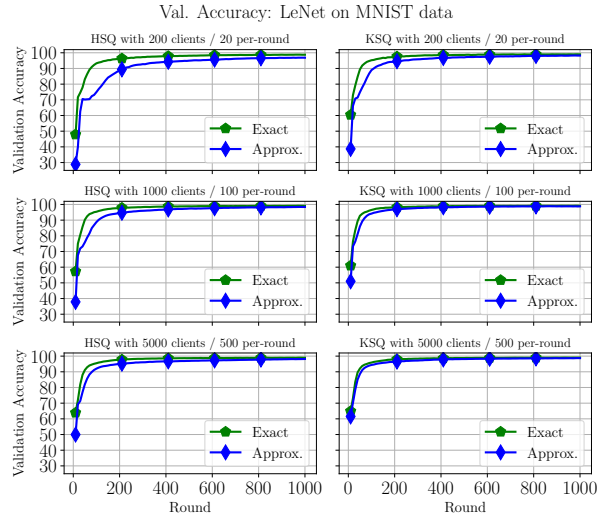


Figure 10: Validation accuracy for training LeNet on the MNIST data set for $n \in \{200, 1000, 5000\}$ clients when selecting 10% of the clients at random per round (n) for SQ with Hadamard (HSQ, left) and with Kashin’s representation (KSQ, right); “Exact” denotes the insecure baseline, “Approx” the simulation of our MPC-based approximate secure aggregation including SepAgg (cf. Fig. 5).

exact training (potentially including only the SepAgg [15] approximation but not our approximate bit-to-arithmetic conversion) is used for fine tuning up to the desired target accuracy.




Figure 11: Validation accuracy for training ResNet9 on the CIFAR10 data set for $N = 1000$ clients with random $n = 40$ selected per round for quantization techniques and protocols as in Fig. 10.

In Tab. 7, we additionally compare the exact inter-server MPC communication cost for a naive MPC implementation of the exact computation to our optimized approximate version including SepAgg. As we can see, we improve the offline communication by factor $\approx 15\times$. For the online communication, we can see a wide range of improvement factors from $2.9\times$ to $48\times$ for MNIST with $n = 500$. This highlights the positive impact when utilizing the SepAgg approach for aggregating an increasingly large number of rows non-interactively. Note that there are slight differences in communication overhead for HSQ and KSQ. This is because for an efficient GPU-friendly implementation of the randomized Hadamard transform, which we use for both rotating the gradients in HSQ and for calculating Kashin’s coefficients in KSQ, we require that the gradients’ size are a power of 2. In §B.3, we detail how we can minimize the resulting overhead by dividing the gradients into chunks, and we also give the exact number of bits per gradient that we assume in our calculations for each algorithm.

Benchmark	n	Method	Naive Exact (cf. Fig. 3)		Our Approx. (cf. Fig. 5)	
			Offline	Online	Offline	Online
MNIST/ LeNet	20	HSQ	572.89	1.51	58.26	0.52
		KSQ	644.50	1.70	65.54	0.59
MNIST/ LeNet	100	HSQ	2864.46	5.44	187.49	0.52
		KSQ	3222.51	6.12	210.93	0.59
MNIST/ LeNet	500	HSQ	14322.28	25.10	833.64	0.52
		KSQ	16112.56	28.24	937.85	0.59
CIFAR10/ ResNet9	40	HSQ	87079.45	189.27	6883.13	39.85
		KSQ	100828.84	219.15	7969.94	46.14

Table 7: Inter-server communication per round for our benchmarks for different numbers of clients n in MiB.

4 Defending Untargeted Poisoning Attacks

Our defense called ScionFL-Aura is designed to mitigate untargeted poisoning attacks in the context of secure quantized aggregation. These attacks pose a significant threat to the deployment of FL for two reasons: (i) Untargeted attacks are particularly difficult to detect because, ignorant of the attack, service providers are unaware that they could have achieved a greater accuracy. (ii) Even a minor drop in accuracy can cause enormous (competitive) damage [116].

Most proposed untargeted poisoning attacks on FL use the (unrealistic) assumption that the adversary \mathcal{A} is aware of either the aggregation rule [51] or all benign updates [11]. However, the *Min-Max* attack proposed by [115] defies this assumption and constitutes the state-of-the-art attack. This attack prevents the manipulations from being detected by allowing the adversary to compute representative benign updates using some clean training data; the attacker can then limit the maximum distance of the manipulated update to any other update by the maximum distances of any two benign updates. This ensures that the malicious gradients are sufficiently similar to the set of benign gradients. We refer to [115, §IV] for more specifics on the attack.

In addition to removing assumptions about the adversary’s knowledge, [115] empirically shows that the *Min-Max* attack outperforms the former state-of-the-art poisoning attack [11] for almost all tested datasets. However, since all benchmarks in [11, 115] were performed on FL schemes without quantization, the impact of the *Min-Max* attack on quantized FL schemes is unclear. Hence, we first test the attack’s effectiveness in our framework using the open-sourced code⁸ as baseline. As we discuss in §4.2, we observe that the attacks are effective even in the context of quantization.

4.1 Our Defense: ScionFL-Aura

From an intuitive standpoint, the adversary in an untargeted poisoning attack seeks to manipulate the global update with malicious updates to deviate it as much as possible from the result of an ideal benign training while evading potentially deployed detection mechanisms. This baseline observation was also used by earlier works to propose defense mechanisms [99, 115], however, those cannot be combined trivially with ScionFL without having to de-quantize all updates and running expensive secure computation machinery.

We now outline the general design of ScionFL-Aura and show its effectiveness against the *Min-Max* attack [115]. In §C.2, we

describe how to efficiently instantiate it in an MPC-friendly manner to reduce communication overhead.

Algorithm 1 Our Defense: ScionFL-Aura

```

1: procedure SCIONFL-AURA( $\{\vec{\sigma}_{X_i}, s_{X_i}^{min}, s_{X_i}^{max}\}_{i \in [n]}$ )
   // Gradient Aggregation including poisoned ones.
2:  $\vec{X}^{agg} \leftarrow \text{AGGREGATE}(\{\vec{\sigma}_{X_i}, s_{X_i}^{min}, s_{X_i}^{max}\}_{i \in [n]})$ 
   //  $L_2$ -norm Computation
3:  $L_2^{avg} \leftarrow 0$ 
4: for  $k \leftarrow 1$  to  $n$  do
5:    $L_2^k \leftarrow \text{L2-NORMQ}(\vec{\sigma}_{X_k}, s_{X_k}^{min}, s_{X_k}^{max})$ 
6:    $L_2^{avg} \leftarrow L_2^{avg} + L_2^k$ 
7: end for
   //  $L_2$ -norm based Scaling
8:  $L_2^{avg} \leftarrow L_2^{avg}/n$  // Average of  $L_2$ -norms
9: for  $k \leftarrow 1$  to  $n$  do
10:  if  $L_2^k > \mu_{th} \cdot L_2^{avg}$  then
11:     $s_{X_k}^{min} \leftarrow s_{X_k}^{min} \cdot (\mu_{th} \cdot L_2^{avg})/L_2^k$ 
12:     $s_{X_k}^{max} \leftarrow s_{X_k}^{max} \cdot (\mu_{th} \cdot L_2^{avg})/L_2^k$ 
13:  end if
14: end for
   // Cosine-distance based Filtering
15: for  $k \leftarrow 1$  to  $n$  do
16:    $\theta^k \leftarrow \text{COSINE}((\vec{\sigma}_{X_k}, s_{X_k}^{min}, s_{X_k}^{max}), \vec{X}^{agg})$ 
17: end for
18:  $\mathcal{X} \leftarrow \text{TOP-K}(\vec{\theta}, \psi)$  // Returns  $k$  for which  $\theta^k > \psi$ 
   // Aggregation of filtered updates
19:  $\vec{X}^{aggd} \leftarrow \text{AGGREGATE}(\{\vec{\sigma}_{X_i}, s_{X_i}^{min}, s_{X_i}^{max}\}_{i \in [n], i \notin \mathcal{X}})$ 
20: end procedure

```

Approach: ScionFL-Aura uses a hybrid approach by combining ideas from existing FL defenses based on the L_2 -norm [10, 99, 121] and cosine similarity [31, 99]. Several works like [99] compute these metrics for each pair of clients, resulting in expensive computation. In contrast, we first aggregate all updates, including the poisoned ones, to produce the vector \vec{X}^{agg} , which we then utilize as the reference. At a high level, L_2 -norm based scaling of the gradient vectors is used at first to bound the impact of malicious contributions that are potentially overlooked (i.e., not filtered) in later stages. In a second step, local updates that significantly deviate from the average update direction are considered to be manipulated and, thus, excluded. Concretely, ScionFL-Aura consists of the following steps:

(1) *L_2 -norm-based Scaling.* In this step, the L_2 -norm of each gradient vector is compared against a public threshold multiplied with the average of the L_2 -norms. Let μ_{th} denote the threshold and L_2^{avg} denote the average of the L_2 -norms across all clients. If $L_2^k > \mu_{th} \cdot L_2^{avg}$ for a gradient vector \vec{X} , the vector is scaled⁹ by a factor of $\mu_{th} \cdot L_2^{avg}/L_2^k$. This ensures that no gradient has an L_2 -norm greater than $\mu_{th} \cdot L_2^{avg}$.

(2) *Cosine-distance-based Filtering.* This step involves computing the cosine distance of each gradient from the reference vector \vec{X}^{agg} discussed above. Following that, another round of aggregation is performed on the updated vectors, but without the top- ψ vectors with the highest cosine distances, which are considered malicious.

⁸<https://github.com/vrt1shjwkr/NDSS21-Model-Poisoning>

⁹Scaling a quantized vector requires simply scaling the scales (cf. §2.4).

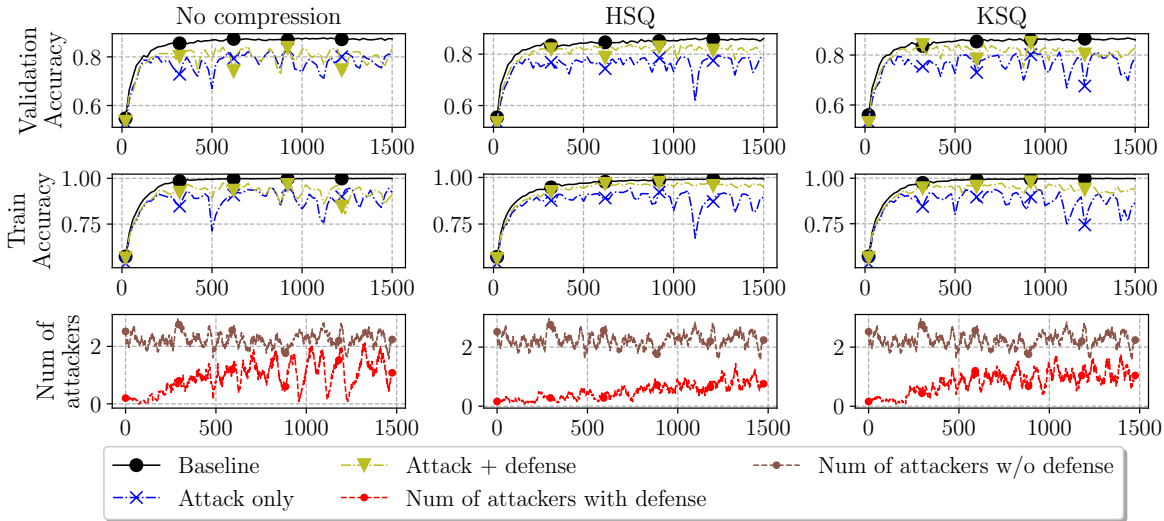


Figure 12: Effect of *Min-Max* attack [115] on training ResNet9 with CIFAR10 with and without our defense ScionFL-Aura assuming 20% of $N = 50$ clients are corrupted. The number of attackers included in the global update varies even without defense due to random client selection.

In this case, ψ is either a known bound (i.e., defined in advance by the service provider) or an accepted percentile determined based on an assumed attacker ratio that follows a normal distribution.

Alg. 1 provides the formal details of ScionFL-Aura, including support for quantized aggregation. Note that we use an optimizer with momentum for FedAvg which ensures that even if the majority of clients picked at random in a training round happens to be malicious, the optimization is still based on benign contributions from the previous round. §C.2 provides details on the sub-protocols utilized in our defense algorithm given in Alg. 1.

4.2 Effectiveness Evaluation

To analyze the effectiveness of ScionFL-Aura, we test it against the *Min-Max* attack [115].

Setup: Training involves $N = 50$ clients of which 20% (as in [115]) are corrupted.¹⁰ Per training iteration, a random subset of $n = 10$ clients is chosen to train the global model. Each client C runs its local training for 10 iterations with batches of $B = 128$ samples and a learning rate of $\eta_C = 0.1$. The defense threshold μ_{th} is set to 3 and the momentum is 0.9.

Experimental Results: Our results when training ResNet9 on CIFAR10 (i) without an attack, (ii) under attack without defense, and (iii) under attack with ScionFL-Aura in place are given in Fig. 12. We compare the attack’s effect when no compression is in place as well as when applying SQ with the randomized Hadamard transform (HSQ) or with Kashin’s representation (KSQ). We also provide similar results for training VGG11 in §C.2. As shown in Fig. 12, our re-implementation of the *Min-Max* attack substantially reduces the validation accuracy by up to 20% when no defense is in place. This

is in line with [115], where the authors report an accuracy degradation between 10.1% and 42.1% for CIFAR10, depending on the model architecture and the aggregation scheme. Furthermore, our experiments show that quantization does not significantly change the impact of the attack. When ScionFL-Aura is enabled, we can remove more than half of the malicious updates in each training iteration compared to when no defense is in place. In fact, quantization supports our defense as the additional noise added to synchronized malicious updates overturns the attacker’s ability of staying just below the detection threshold. As a result, compared to unprotected training, the validation accuracy decreases by at most 7.7% for HSQ and 10.7% for KSQ.

5 Conclusion

In this paper, we proposed ScionFL, the first secure aggregation framework for FL that can efficiently process quantized updates with an optimal client-server communication as low as 1-bit per coordinate. Together with our novel defense against state-of-the-art (untargeted) poisoning attacks, this is the first work to unite security, efficiency, and robustness in a single and coherent framework. As part of future work, we plan to investigate how to efficiently extend our MPC protocols to guarantee also correctness and not “only” privacy when dealing with malicious MPC servers.

Acknowledgments

This project received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 850990 PSOTI). It was co-funded by the Deutsche Forschungsgemeinschaft (DFG) within SFB 1119 CROSSING/236615297 and GRK 2050 Privacy & Trust/251805230, by the German Federal Ministry of Education and Research and the Hessen State Ministry for Higher Education, Research and the Arts within ATHENE, and by the Private AI Collaborative Research Institute, funded by Intel, Avast, and VMware.

¹⁰[116] points out that assuming more than 1% of corrupted clients is unrealistic for most scenarios. However, in our experiments the attack failed to notably reduce the accuracy with such a low corruption level. Thus, we tested against 20% of corrupted clients as in the original attack paper [115].

References

- [1] Josh Aas and Tim Geoghegan. [n. d.]. *Introducing ISRG Prio Services for Privacy Respecting Metrics*. <https://www.abetterinternet.org/post/introducing-prio-services/>.
- [2] Surya Addanki, Kevin Garbe, Eli Jaffe, Rafail Ostrovsky, and Antigoni Polychroniadou. 2022. Prio+: Privacy Preserving Aggregate Statistics via Boolean Shares. In *SCN*.
- [3] Alham Fikri Aji and Kenneth Heafield. 2017. Sparse Communication for Distributed Gradient Descent. In *EMNLP*.
- [4] Dan Alistarh, Torsten Hoefler, Mikael Johansson, Nikola Konstantinov, Sarit Khirirat, and Cédric Renggli. 2018. The Convergence of Sparsified Gradient Methods. In *NeurIPS*.
- [5] Sébastien Andreina, Giorgia Azzurra Marson, Helen Möllering, and Ghassan Karame. 2021. BaFFLe: Backdoor Detection via Feedback-based Federated Learning. In *IEEE ICDCS*.
- [6] Apple and Google. [n. d.]. Exposure Notification Privacy-preserving Analytics (ENPA) White Paper. https://covid19-static.cdn-apple.com/applications/covid19/current/static/contact-tracing/pdf/ENPA_White_Paper.pdf.
- [7] Toshinori Araki, Jun Furukawa, Yehuda Lindell, Ariel Nof, and Kazuma Ohara. 2016. High-Throughput Semi-Honest Secure Three-Party Computation with an Honest Majority. In *ACM CCS*.
- [8] Gilad Asharov, Shai Halevi, Yehuda Lindell, and Tal Rabin. 2018. Privacy-Preserving Search of Similar Patients in Genomic Data. *PETS 2018* (2018).
- [9] Gilad Asharov, Yehuda Lindell, Thomas Schneider, and Michael Zohner. 2013. More efficient oblivious transfer and extensions for faster secure computation. In *ACM CCS*.
- [10] Eugene Bagdasaryan, Andreas Veit, Yiqing Hua, Deborah Estrin, and Vitaly Shmatikov. 2020. How To Backdoor Federated Learning. In *AISTATS*.
- [11] Gilad Baruch, Moran Baruch, and Yoav Goldberg. 2019. A Little Is Enough: Circumventing Defenses For Distributed Learning. In *NeurIPS*.
- [12] Ran Ben Basat, Michael Mitzenmacher, and Shay Vargaftik. 2021. How to Send a Real Number Using a Single Bit (And Some Shared Randomness). In *ICALP*.
- [13] Ran Ben Basat, Shay Vargaftik, Amit Portnoy, Gil Einziger, Yaniv Ben-Itzhak, and Michael Mitzenmacher. 2022. QUIC-FL: Quick Unbiased Compression for Federated Learning. (2022). <https://arxiv.org/abs/2205.13341>.
- [14] Carsten Baum, Ivan Damgård, Tomas Toft, and Rasmus Winther Zakarias. 2016. Better Preprocessing for Secure Multiparty Computation. In *ACNS*.
- [15] Constance Beguier, Mathieu Andreux, and Eric W. Tramel. 2020. Efficient Sparse Secure Aggregation for Federated Learning. (2020). <https://arxiv.org/abs/2007.14861>.
- [16] James Henry Bell, Kallista A. Bonawitz, Adrià Gascón, Tancrede Lepoint, and Mariana Raykova. 2020. Secure Single-Server Aggregation with (Poly)Logarithmic Overhead. In *ACM CCS*.
- [17] Aner Ben-Efraim, Michael Nielsen, and Eran Omri. 2019. Turbospeedz: Double Your Online SPDZ! Improving SPDZ Using Function Dependent Preprocessing. In *ACNS*.
- [18] Jeremy Bernstein, Yu-Xiang Wang, Kamyar Azizzadenesheli, and Animashree Anandkumar. 2018. SIGNSGD: Compressed Optimisation for Non-Convex Problems. In *ICML*.
- [19] Aleksandr Beznosikov, Samuel Horváth, Peter Richtárik, and Mher Safaryan. 2020. On Biased Compression for Distributed Learning. (2020). <https://arxiv.org/abs/2002.12410>.
- [20] Arjun Nitin Bhagoji, Supriyo Chakraborty, Prateek Mittal, and Seraphin B. Calo. 2019. Analyzing Federated Learning through an Adversarial Lens. In *ICML*.
- [21] Peva Blanchard, El Mahdi El Mhamdi, Rachid Guerraoui, and Julien Stainer. 2017. Machine Learning with Adversaries: Byzantine Tolerant Gradient Descent. In *NeurIPS*.
- [22] Fabian Boemer, Rosario Cammarota, Daniel Demmler, Thomas Schneider, and Hossein Yalame. 2020. MP2ML: a mixed-protocol machine learning framework for private inference. In *ARES*.
- [23] Franziska Boenisch, Adam Dziedzic, Roi Schuster, Ali Shahin Shamsabadi, Iliia Shumailov, and Nicolas Papernot. 2021. When the Curious Abandon Honesty: Federated Learning Is Not Private. (2021). <https://arxiv.org/abs/2112.02918>.
- [24] Franziska Boenisch, Adam Dziedzic, Roi Schuster, Ali Shahin Shamsabadi, Iliia Shumailov, and Nicolas Papernot. 2022. All You Need Is Matplotlib. <http://www.cleverhans.io/2022/04/17/fl-privacy.html>.
- [25] Franziska Boenisch, Adam Dziedzic, Roi Schuster, Ali Shahin Shamsabadi, Iliia Shumailov, and Nicolas Papernot. 2023. Is Federated Learning a Practical PET Yet? *CoRR* (2023).
- [26] Kallista A. Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H. Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. 2017. Practical Secure Aggregation for Privacy-Preserving Machine Learning. In *ACM CCS*.
- [27] Elette Boyle, Geoffroy Couteau, Niv Gilboa, Yuval Ishai, Lisa Kohl, and Peter Scholl. 2019. Efficient Pseudorandom Correlation Generators: Silent OT Extension and More. In *CRYPTO*.
- [28] Lennart Braun, Daniel Demmler, Thomas Schneider, and Oleksandr Tkachenko. 2022. MOTION - A Framework for Mixed-Protocol Multi-Party Computation. *ACM Trans. Priv. Secur.* (2022).
- [29] Megha Byali, Harsh Chaudhari, Arpita Patra, and Ajith Suresh. 2020. FLASH: Fast and Robust Framework for Privacy-preserving Machine Learning. *PETS* (2020).
- [30] Sebastian Caldas, Jakub Konečný, H. Brendan McMahan, and Ameet Talwalkar. 2018. Expanding the Reach of Federated Learning by Reducing Client Resource Requirements. (2018). <http://arxiv.org/abs/1812.07210>.
- [31] Xiaoyu Cao, Minghong Fang, Jia Liu, and Neil Zhenqiang Gong. 2021. FLTrust: Byzantine-robust Federated Learning via Trust Bootstrapping. In *NDSS*.
- [32] Octavian Catrina and Amitabh Saxena. 2010. Secure Computation with Fixed-Point Numbers. In *FC*.
- [33] Harsh Chaudhari, Ashish Choudhury, Arpita Patra, and Ajith Suresh. 2019. ASTRA: High Throughput 3PC over Rings with Application to Secure Prediction. In *ACM Conference on Cloud Computing Security Workshop, CCSW@CCS*.
- [34] Harsh Chaudhari, Rahul Rachuri, and Ajith Suresh. 2020. Trident: Efficient 4PC Framework for Privacy Preserving Machine Learning. In *NDSS*.
- [35] Wei-Ning Chen, Christopher A. Choquette-Choo, Peter Kairouz, and Ananda Theertha Suresh. 2022. The Fundamental Price of Secure Aggregation in Differentially Private Federated Learning. In *ICML*.
- [36] Jung Hee Cheon, Andrey Kim, Miran Kim, and Yong Soo Song. 2017. Homomorphic Encryption for Arithmetic of Approximate Numbers. In *ASIACRYPT*. Springer.
- [37] Brian Chmiel, Liad Ben-Uri, Moran Shkolnik, Elad Hoffer, Ron Banner, and Daniel Soudry. 2021. Neural gradients are near-lognormal: improved quantized and sparse training. In *ICLR*.
- [38] Henry Corrigan-Gibbs and Dan Boneh. 2017. Prio: Private, Robust, and Scalable Computation of Aggregate Statistics. In *USENIX NSDI*.
- [39] Geoffroy Couteau, Peter Rindal, and Srinivasan Raghuraman. 2021. Silver: Silent VOLE and Oblivious Transfer from Hardness of Decoding Structured LDPC Codes. In *CRYPTO*.
- [40] Ronald Cramer, Ivan Damgård, Daniel Escudero, Peter Scholl, and Chaoping Xing. 2018. SPDZ^k: Efficient MPC mod 2^k for Dishonest Majority. In *CRYPTO*.
- [41] Ivan Damgård, Daniel Escudero, Tore Kasper Frederiksen, Marcel Keller, Peter Scholl, and Nikolaj Volgushev. 2019. New Primitives for Actively-Secure MPC over Rings with Applications to Private Machine Learning. In *IEEE S&P*.
- [42] Ivan Damgård, Marcel Keller, Enrique Larraia, Valerio Pastro, Peter Scholl, and Nigel P. Smart. 2013. Practical Covertly Secure MPC for Dishonest Majority - Or: Breaking the SPDZ Limits. In *ESORICS*.
- [43] Ivan Damgård, Claudio Orlandi, and Mark Simkin. 2018. Yet Another Compiler for Active Security or: Efficient MPC Over Arbitrary Rings. In *CRYPTO*.
- [44] Ivan Damgård, Valerio Pastro, Nigel Smart, and Sarah Zakarias. 2012. Multiparty computation from somewhat homomorphic encryption. In *CRYPTO*.
- [45] Peter Davies, Vijaykrishna Gurunathan, Niusha Moshrefi, Saleh Ashkboos, and Dan Alistarh. 2021. New Bounds For Distributed Mean Estimation and Variance Reduction. In *ICLR*.
- [46] Daniel Demmler, Thomas Schneider, and Michael Zohner. 2015. ABY - A Framework for Efficient Mixed-Protocol Secure Two-Party Computation. In *NDSS*.
- [47] Ye Dong, Xiaojun Chen, Kaiyun Li, Dakui Wang, and Shuai Zeng. 2021. FLOD: Oblivious Defender for Private Byzantine-Robust Federated Learning with Dishonest-Majority. In *ESORICS*.
- [48] Tariq Elahi, George Danezis, and Ian Goldberg. 2014. PrivEx: Private Collection of Traffic Statistics for Anonymous Communication Networks. In *ACM CCS*.
- [49] Zekeriya Erkin, Juan Ramón Troncoso-Pastoriza, Reginald L. Lagendijk, and Fernando Pérez-González. 2013. Privacy-Preserving Data Aggregation in Smart Metering Systems: An Overview. *IEEE Signal Process. Mag.* (2013).
- [50] Daniel Escudero, Satrajit Ghosh, Marcel Keller, Rahul Rachuri, and Peter Scholl. 2020. Improved Primitives for MPC over Mixed Arithmetic-Binary Circuits. In *CRYPTO*.
- [51] Minghong Fang, Xiaoyu Cao, Jinyuan Jia, and Neil Zhenqiang Gong. 2020. Local Model Poisoning Attacks to Byzantine-Robust Federated Learning. In *USENIX Security*.
- [52] Jiawei Fei, Chen-Yu Ho, Atal Narayan Sahu, Marco Canini, and Amedeo Sapio. 2021. Efficient sparse collective communication and its application to accelerate distributed deep learning. In *ACM SIGCOMM Conference*.
- [53] Hossein Fereidooni, Samuel Marchal, Markus Miettinen, Azalia Mirhoseini, Helen Möllering, Thien Duc Nguyen, Phillip Rieger, Ahmad-Reza Sadeghi, Thomas Schneider, Hossein Yalame, and Shaza Zeitouni. 2021. SAFELearn: Secure Aggregation for private Federated Learning. In *IEEE S&P Workshops*.
- [54] Liam H. Fowl, Jonas Geiping, Wojciech Czaja, Micah Goldblum, and Tom Goldstein. 2022. Robbing the Fed: Directly Obtaining Private Data in Federated Learning with Modified Models. In *ICLR*.
- [55] Jun Furukawa, Yehuda Lindell, Ariel Nof, and Or Weinstein. 2017. High-Throughput Secure Three-Party Computation for Malicious Adversaries and an Honest Majority. In *EUROCRYPT*.

- [56] Karan Ganju, Qi Wang, Wei Yang, Carl A. Gunter, and Nikita Borisov. 2018. Property Inference Attacks on Fully Connected Neural Networks using Permutation Invariant Representations. In *ACM CCS*.
- [57] Till Gehlhar, Felix Marx, Thomas Schneider, Tobias Wehrle, Ajith Suresh, and Hossein Yalame. 2023. SafeFL: MPC-friendly framework for Private and Robust Federated Learning. In *IEEE S&P Workshops*.
- [58] Oded Goldreich, Silvio Micali, and Avi Wigderson. 1987. How to Play any Mental Game or A Completeness Theorem for Protocols with Honest Majority. In *ACM STOC*.
- [59] S. Dov Gordon, Samuel Ranellucci, and Xiao Wang. 2018. Secure Computation with Low Communication from Cross-Checking. In *ASIACRYPT*.
- [60] Vipul Goyal, Hanjun Li, Rafail Ostrovsky, Antigoni Polychroniadou, and Yifan Song. 2021. ATLAS: Efficient and Scalable MPC in the Honest Majority Setting. In *CRYPTO*.
- [61] Aditya Hegde, Helen Möllering, Thomas Schneider, and Hossein Yalame. 2021. SoK: Efficient Privacy-preserving Clustering. *PETS (2021)*.
- [62] Ehsan Hesamifard, Hassan Takabi, Mehdi Ghasemi, and Rebecca N. Wright. 2018. Privacy-preserving Machine Learning as a Service. *PETS (2018)*.
- [63] Nikita Ivkin, Daniel Rothchild, Enayat Ullah, Vladimir Braverman, Ion Stoica, and Raman Arora. 2019. Communication-efficient Distributed SGD with Sketching. In *NeurIPS*.
- [64] Richeng Jin, Yufan Huang, Xiaofan He, Huaiyu Dai, and Tianfu Wu. 2020. Stochastic-Sign SGD for Federated Learning with Theoretical Guarantees. (2020). <https://arxiv.org/abs/2002.10940>.
- [65] Marc Joye and Benoît Libert. 2013. A Scalable Scheme for Privacy-Preserving Aggregation of Time-Series Data. In *FC*.
- [66] Chiraag Juvekar, Vinod Vaikuntanathan, and Anantha Chandrakasan. 2018. GAZELLE: A low latency framework for secure neural network inference. In *USENIX Security*.
- [67] Swanand Kadhe, Nived Rajaraman, Onur Ozan Koyluoglu, and Kannan Ramchandran. 2020. FastSecAgg: Scalable Secure Aggregation for Privacy-Preserving Federated Learning. (2020). <https://arxiv.org/abs/2009.11248>.
- [68] Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, and et al. 2021. Advances and Open Problems in Federated Learning. *Found. Trends Mach. Learn.* (2021).
- [69] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank J. Reddi, Sebastian U. Stich, and Ananda Theertha Suresh. 2020. SCAFFOLD: Stochastic Controlled Averaging for Federated Learning. In *ICML*.
- [70] Marcel Keller, Emmanuela Orsini, and Peter Scholl. 2016. MASCOT: Faster Malicious Arithmetic Secure Computation with Oblivious Transfer. In *ACM SIGSAC*.
- [71] Marcel Keller, Valerio Pastro, and Dragos Rotaru. 2018. Overdrive: Making SPDZ Great Again. In *EUROCRYPT*.
- [72] Marcel Keller, Peter Scholl, and Nigel P. Smart. 2013. An architecture for practical actively secure MPC with dishonest majority. In *ACM CCS*.
- [73] Duhyeong Kim, Yongha Son, Dongwoo Kim, Andrey Kim, Seungwan Hong, and Jung Hee Cheon. 2019. Privacy-preserving Approximate GWAS computation based on Homomorphic Encryption. (2019). <https://eprint.iacr.org/2019/152>.
- [74] Vladimir Kolesnikov and Thomas Schneider. 2008. Improved Garbled Circuit: Free XOR Gates and Applications. In *ICALP*.
- [75] Jakub Konečný, H. Brendan McMahan, Felix X. Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. 2016. Federated Learning: Strategies for Improving Communication Efficiency. (2016). <http://arxiv.org/abs/1610.05492>.
- [76] Nishat Koti, Mahak Pancholi, Arpita Patra, and Ajith Suresh. 2021. SWIFT: Super-fast and Robust Privacy-Preserving Machine Learning. In *USENIX Security*.
- [77] Nishat Koti, Shravani Patil, Arpita Patra, and Ajith Suresh. 2022. MPClan: Protocol Suite for Privacy-Conscious Computations. <https://eprint.iacr.org/2022/675>.
- [78] Nishat Koti, Arpita Patra, Rahul Rachuri, and Ajith Suresh. 2022. Tetrad: Actively Secure 4PC for Secure Training and Inference. In *NDSS*.
- [79] Alex Krizhevsky. 2009. *Learning multiple layers of features from tiny images*. Technical Report.
- [80] Klaus Kursawe, George Danezis, and Markulf Kohlweiss. 2011. Privacy-Friendly Aggregation for the Smart-Grid. In *PETS*.
- [81] Maximilian Lam, Gu-Yeon Wei, David Brooks, Vijay Janapa Reddi, and Michael Mitzenmacher. 2021. Gradient Disaggregation: Breaking Privacy in Federated Learning by Reconstructing the User Participant Matrix. In *ICML*.
- [82] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* (1998).
- [83] Kwing Hei Li, Pedro Porto Buarque de Gusmão, Daniel J. Beutel, and Nicholas D. Lane. 2021. Secure aggregation for federated learning in flower. In *ACM International Workshop on Distributed Machine Learning*.
- [84] Liping Li, Wei Xu, Tianyi Chen, Georgios B. Giannakis, and Qing Ling. 2019. RSA: Byzantine-Robust Stochastic Aggregation Methods for Distributed Learning from Heterogeneous Datasets. In *AAAI*.
- [85] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. 2020. Federated Optimization in Heterogeneous Networks. In *MLSys*.
- [86] Yehuda Lindell, Benny Pinkas, Nigel P. Smart, and Avishay Yanai. 2015. Efficient Constant Round Multi-party Computation Combining BMR and SPDZ. In *CRYPTO*.
- [87] Yuri Lyubarskii and Roman Vershynin. 2010. Uncertainty principles and vector quantization. *IEEE Trans. Inf. Theory* (2010).
- [88] Eleftheria Makri, Dragos Rotaru, Frederik Vercauteren, and Sameer Wagh. 2021. Rabbit: Efficient Comparison for Secure Multi-Party Computation. In *FC*.
- [89] Mohamad Mansouri, Melek Onen, Wafa Ben Jaballah, and Mauro Conti. 2023. SoK: Secure aggregation based on cryptographic schemes for federated Learning. In *PETS*.
- [90] Sahar Mazloom, Phi Hung Le, Samuel Ranellucci, and S. Dov Gordon. 2020. Secure parallel computation on national scale volumes of data. In *USENIX Security*.
- [91] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *AISTATS*.
- [92] Luca Melis, Congzheng Song, Emiliano De Cristofaro, and Vitaly Shmatikov. 2019. Exploiting Unintended Feature Leakage in Collaborative Learning. In *IEEE S&P*.
- [93] Pratyush Mishra, Ryan Lehmkuhl, Akshayaram Srinivasan, Wenting Zheng, and Raluca Ada Popa. 2020. Delphi: A cryptographic inference service for neural networks. In *USENIX Security*.
- [94] Payman Mohassel and Peter Rindal. 2018. ABY³: A Mixed Protocol Framework for Machine Learning. In *ACM CCS*.
- [95] Payman Mohassel and Yupeng Zhang. 2017. SecureML: A System for Scalable Privacy-Preserving Machine Learning. In *IEEE S&P*.
- [96] Arup Mondal, Yash More, Prashanthi Ramchandran, Priyam Panda, Harpreet Virk, and Debayan Gupta. 2022. Scotch: An Efficient Secure Computation Framework for Secure Aggregation. (2022). <https://arxiv.org/abs/2010.07730>.
- [97] Jean-Pierre Münch, Thomas Schneider, and Hossein Yalame. 2021. VASA: Vector AES Instructions for Security Applications. In *ACM ACSAC*.
- [98] Milad Nasr, Reza Shokri, and Amir Houmansadr. 2019. Comprehensive Privacy Analysis of Deep Learning: Passive and Active White-box Inference Attacks against Centralized and Federated Learning. In *IEEE S&P*.
- [99] Thien Duc Nguyen, Phillip Rieger, Huili Chen, Hossein Yalame, Helen Möllering, Hossein Fereidooni, Samuel Marchal, Markus Miettinen, Azalia Mirhoseini, Farihan Koushanfar, Ahmad-Reza Sadeghi, Thomas Schneider, and Shaza Zeitouni. 2022. FLAME: Taming Backdoors in Federated Learning. In *USENIX Security*.
- [100] Ahmed El Ouaadri and Ahmed Abdelhadi. 2022. Differential Privacy for Deep and Federated Learning: A Survey. *IEEE Access* (2022).
- [101] Dario Pasquini, Danilo Francati, and Giuseppe Ateniese. 2022. Eluding Secure Aggregation in Federated Learning via Model Inconsistency. In *CCS*.
- [102] Arpita Patra, Thomas Schneider, Ajith Suresh, and Hossein Yalame. 2021. ABY2.0: Improved Mixed-Protocol Secure Two-Party Computation. In *USENIX Security*.
- [103] Arpita Patra, Thomas Schneider, Ajith Suresh, and Hossein Yalame. 2021. SynCirc: Efficient Synthesis of Depth-Optimized Circuits for Secure Computation. In *IEEE HOST*.
- [104] Arpita Patra and Ajith Suresh. 2020. BLAZE: Blazing Fast Privacy-Preserving Machine Learning. In *NDSS*.
- [105] Raluca A. Popa, Andrew J. Blumberg, Hari Balakrishnan, and Frank H. Li. 2011. Privacy and accountability for location-based aggregate statistics. In *ACM CCS*.
- [106] Apostolos Pyrgelis, Carmela Troncoso, and Emiliano De Cristofaro. 2018. Knock Knock, Who's There? Membership Inference on Aggregate Location Data. In *NDSS*.
- [107] Mayank Rathee, Conghao Shen, Sameer Wagh, and Raluca Ada Popa. 2023. ELSA: Secure Aggregation for Federated Learning with Malicious Actors. In *IEEE S&P Workshops*.
- [108] M. Sadegh Riazi, Mohammad Samragh, Hao Chen, Kim Laine, Kristin E. Lauter, and Farinaz Koushanfar. 2019. XONN: XNOR-based Oblivious Deep Neural Network Inference. In *USENIX Security*.
- [109] Peter Richtárik, Igor Sokolov, and Ilyas Fatkhullin. 2021. EF21: A New, Simpler, Theoretically Better, and Practically Faster Error Feedback. In *NeurIPS*.
- [110] Dragos Rotaru and Tim Wood. 2019. MArBled Circuits: Mixing Arithmetic and Boolean Circuits with Active Security. In *INDOCRYPT*.
- [111] Daniel Rothchild, Ashwinee Panda, Enayat Ullah, Nikita Ivkin, Ion Stoica, Vladimir Braverman, Joseph Gonzalez, and Raman Arora. 2020. FetchSGD: Communication-Efficient Federated Learning with Sketching. In *ICML*.
- [112] Mher Safaryan, Egor Shulgin, and Peter Richtárik. 2020. Uncertainty Principle for Communication Compression in Distributed and Federated Learning and the Search for an Optimal Compressor. (2020). <https://arxiv.org/abs/2002.08958>.
- [113] Sinem Sav, Apostolos Pyrgelis, Juan R Troncoso-Pastoriza, David Froelicher, Jean-Philippe Bossuat, Joao Sa Sousa, and Jean-Pierre Hubaux. 2021. POSEIDON: Privacy-preserving federated neural network learning. In *NDSS*.
- [114] Frank Seide, Hao Fu, Jasha Droppo, Gang Li, and Dong Yu. 2014. 1-bit stochastic gradient descent and its application to data-parallel distributed training of speech DNNs. In *INTERSPEECH*.

- [115] Virat Shejwalkar and Amir Houmansadr. 2021. Manipulating the Byzantine: Optimizing Model Poisoning Attacks and Defenses for Federated Learning. In *NDSS*.
- [116] Virat Shejwalkar, Amir Houmansadr, Peter Kairouz, and Daniel Ramage. 2022. Back to the Drawing Board: A Critical Evaluation of Poisoning Attacks on Production Federated Learning. In *IEEE S&P*.
- [117] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. 2017. Membership Inference Attacks Against Machine Learning Models. In *IEEE S&P*.
- [118] Jinhyun So, Ramy E. Ali, Basak Guler, Jiantao Jiao, and Salman Avestimehr. 2021. Securing Secure Aggregation: Mitigating Multi-Round Privacy Leakage in Federated Learning. *AAAI* (2021).
- [119] Jinhyun So, Basak Guler, and Amir Salman Avestimehr. 2021. Turbo-Aggregate: Breaking the Quadratic Aggregation Barrier in Secure Federated Learning. *IEEE J. Sel. Areas Inf. Theory* (2021).
- [120] Sebastian U. Stich, Jean-Baptiste Cordonnier, and Martin Jaggi. 2018. Sparsified SGD with Memory. In *NeurIPS*.
- [121] Ziteng Sun, Peter Kairouz, Ananda Theertha Suresh, and H Brendan McMahan. 2019. Can you really backdoor federated learning?. In *NeurIPS FL Workshop*.
- [122] Ajith Suresh. 2021. MPCLeague: Robust MPC Platform for Privacy-Preserving Machine Learning. PhD Thesis. <https://arxiv.org/abs/2112.13338>.
- [123] Ananda Theertha Suresh, Felix X. Yu, Sanjiv Kumar, and H. Brendan McMahan. 2017. Distributed Mean Estimation with Limited Communication. In *ICML*.
- [124] Hanlin Tang, Shaoduo Gan, Ammar Ahmad Awan, Samyam Rajbhandari, Conglong Li, Xiangru Lian, Ji Liu, Ce Zhang, and Yuxiong He. 2021. 1-bit Adam: Communication Efficient Large-Scale Training with Adam’s Convergence Speed. In *ICML*.
- [125] Shay Vargafik, Ran Ben Basat, Amit Portnoy, Gal Mendelson, Yaniv Ben-Itzhak, and Michael Mitzenmacher. 2022. EDEN: Communication-Efficient and Robust Distributed Mean Estimation for Federated Learning. In *ICML*.
- [126] Shay Vargafik, Ran Ben-Basat, Amit Portnoy, Gal Mendelson, Yaniv Ben-Itzhak, and Michael Mitzenmacher. 2021. DRIVE: One-bit Distributed Mean Estimation. In *NeurIPS*.
- [127] Tanya Verma and Sudheesh Singanamalla. 2020. Improving DNS Privacy with Oblivious DoH in 1.1.1.1. <https://blog.cloudflare.com/oblivious-dns/>.
- [128] Jianyu Wang, Zachary Charles, Zheng Xu, Gauri Joshi, H. Brendan McMahan, and et al. 2021. A Field Guide to Federated Optimization. (2021). <https://arxiv.org/abs/2107.06917>.
- [129] Lixu Wang, Shichao Xu, Xiao Wang, and Qi Zhu. 2019. Eavesdrop the Composition Proportion of Training Labels in Federated Learning. (2019). <http://arxiv.org/abs/1910.06044>.
- [130] Wei Wen, Cong Xu, Feng Yan, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. 2017. TernGrad: Ternary Gradients to Reduce Communication in Distributed Deep Learning. In *NeurIPS*.
- [131] Yuxin Wen, Jonas Geiping, Liam Fowl, Micah Goldblum, and Tom Goldstein. 2022. Fishing for User Data in Large-Batch Federated Learning via Gradient Magnification. In *ICML*.
- [132] Huang Xiao, Battista Biggio, Gavin Brown, Giorgio Fumera, Claudia Eckert, and Fabio Roli. 2015. Is Feature Selection Secure against Training Data Poisoning?. In *ICML*.
- [133] Chien-Sheng Yang, Jinhyun So, Chaoyang He, Songze Li, Qian Yu, and Salman Avestimehr. 2021. LightSecAgg: Rethinking Secure Aggregation in Federated Learning. (2021). <https://arxiv.org/abs/2109.14236>.
- [134] Andrew Chi-Chih Yao. 1982. Protocols for Secure Computations (Extended Abstract). In *FOCS*.
- [135] Dong Yin, Yudong Chen, Kannan Ramchandran, and Peter L. Bartlett. 2018. Byzantine-Robust Distributed Learning: Towards Optimal Statistical Rates. In *ICML*.
- [136] Chengliang Zhang, Suyi Li, Junzhe Xia, Wei Wang, Feng Yan, and Yang Liu. 2020. BatchCrypt: Efficient Homomorphic Encryption for Cross-Silo Federated Learning. In *USENIX ATC*.
- [137] Zaixi Zhang, Xiaoyu Cao, Jinyuan Jia, and Neil Zhenqiang Gong. 2022. FL-Detector: Defending Federated Learning Against Model Poisoning Attacks via Detecting Malicious Clients. In *KDD*.
- [138] Zhengming Zhang, Ashwinee Panda, Linyue Song, Yaoqing Yang, Michael W. Mahoney, Prateek Mittal, Kannan Ramchandran, and Joseph Gonzalez. 2022. Neurotoxin: Durable Backdoors in Federated Learning. In *ICML*.

A Related Work & Background Information

A.1 Stochastic Quantization

This section provides additional details regarding stochastic quantization schemes discussed in §2.2.

Preprocessing via Random Rotations: To deal with possible limitations of vanilla SQ, recent state-of-the-art works suggest

to *randomly rotate* the input vector prior to SQ [123]. That is, the clients and the aggregator draw rotation matrices according to some known distribution; the clients then send the quantization of the rotated vectors while the aggregator applies the inverse rotation on the estimated rotated vector. Intuitively, the coordinates of a randomly rotated vector are identically distributed, and thus the expected difference between the coordinates is smaller, allowing for a more accurate quantization. For n clients and a gradient with m coordinates, this approach achieves a NMSE¹¹ of $O(\frac{\log m}{n})$ using $O(m)$ bits, which asymptotically improves over the $O(\frac{m}{n})$ NMSE bound of vanilla SQ. The computational complexity, on the other hand, is increased from $O(m)$ to $O(m \log m)$ when utilizing the randomized Hadamard transform for rotations.

Preprocessing via Kashin’s Representation: The rotation approach was recently improved using Kashin’s representation [30, 87, 112]. Roughly speaking, it allows representing an m -dimensional vector using a slightly larger vector with $\lambda \cdot m$ smaller coefficients ($\lambda > 1$). It can be shown that applying SQ to the Kashin coefficients allows for an optimal NMSE of $O(\frac{1}{n})$ using $O(\lambda \cdot m)$ bits. Compared with [123], Kashin’s representation yields a lower NMSE bound by a factor of $\log m$ at the cost of increasing the computational complexity by the same factor [13, 30].

A.2 Additional Compression Techniques

In this work, we focus on quantization as a means to reduce bandwidth. We nevertheless briefly overview some additional techniques considered for FL gradient compression.

Sparsification: Some works like [3, 52, 75, 120] consider sparsifying the gradients. Quantization can also be applied to these sparsified gradients as it reduces the number of bits used per entry, while sparsification reduces the number of entries.

Client-side Memory-based Techniques: Some compression techniques, including Top- k [120] and sketching [63], rely on client-side memory and error-feedback [4, 19, 109, 114] to ensure convergence. We consider the cross-device FL setup where clients are stateless (e.g., a client may appear only once during a training procedure). Therefore, client-side-memory-based techniques are mostly designed for the cross-silo FL setup and are less applicable to cross-device FL.

Entropy Encodings: Some techniques use entropy encoding such as arithmetic encoding and Huffman encoding (e.g., [4, 123, 125]). While such techniques are appealing in their bandwidth-to-accuracy trade-offs, it is unclear how to allow for an efficient secure aggregation as gradients must be decoded before being averaged. Also, such techniques usually incur a higher computational overhead at the clients than fixed-length representations. An additional review of current state-of-the-art gradient compression techniques and some open challenges can be found in [68, 75, 128].

A.3 Secure Multi-party Computation

The field of secure multi-party computation (MPC) started with the seminal work of Yao [134] in 1982. It enables to securely compute arbitrary functions on private inputs without leaking anything beyond what can be inferred from the output. Since then,

¹¹The normalized MSE is the mean’s estimate MSE normalized by the mean clients’ gradient squared norms

the field of MPC has seen a variety of advancements of used primitives effectively improving communication and computation efficiency, e.g., [9, 27, 50, 74, 97]. Also, tailored efficient optimizations for varying number of computation parties have been explored, e.g., [33, 34, 46, 78, 102, 103]. Moreover, MPC research considers different assumptions regarding adversarial behavior such as the well-known semi-honest [28, 46] and malicious security model [41, 70–72]) as well as numbers of corrupted computation parties (e.g., honest majority [58, 60] or dishonest majority/full threshold security [28, 40, 46, 72, 102]). Beyond running the computation among several non-colluding parties, another well-established system model (which we use in our work) is outsourcing, where the data owners secret-share their private input data among a set of non-colluding computing parties which then run the private computation on their behalf [1, 6, 127].

A.4 Approximate Secure Computation

To improve efficiency of MPC, few works already considered approximations of the exact computation. Such approximations in MPC include using integer or fixed-point instead of floating-point operations (too many works to cite), approximations in genomic computation [8], and in privacy-preserving machine learning such as for division [32], activation functions [29, 62, 95], and completely changing the classifier to be MPC-friendly [108]. Also, for FHE, approximations are used such as in the approximate HE scheme CKKS [36], which is implemented in the HEAAN library¹² and was used for approximate genomic computations in [73]. In this work, we propose for the first time to use approximations to substantially improve efficiency of FL when combined with MPC and give detailed evaluations on the errors introduced thereby.

A.5 Secure Aggregation

Performing secure aggregation without revealing anything about the aggregated input values beyond what can be inferred from the output was already investigated more than 10 years ago, for example, in the context of smart metering, e.g., [49, 80]. It has come a long way since then, resulting in practical solutions for real-world applications nowadays.

For example, Prio [38] introduces secure protocols for aggregate statistics such as sum, mean, variance, standard deviation, min/max, and frequency. It uses additive arithmetic secret sharing, offers full-threshold security among a small set of servers running the secure computation, and validates inputs to protect against malicious clients. Prio+ [2] optimizes client computation and communication compared to [38] with a Boolean secret sharing-based client input validation and an additional conversion from Boolean to arithmetic sharing. Similar to our work, it has a multi-server setup to jointly compute statistical functions on private inputs. Compared to Prio+ [2], we optimize the naive bit-to-arithmetic conversion presented in [2] for our $\mathcal{F}_{\text{SecAgg}}$ protocol (cf. §3.1), resulting in reduced communication cost of 2.4× with exact results and 4× with our novel approximating variant for $n = 10^5$, where n is the number of clients. Popa et al. [105] specifically focus on secure location-based aggregation statistics, Joye et al. [65] on time-series data, and PrivEx [48] on traffic data in anonymous communication.

So et al. [118] point out that differences among securely aggregated updates across multiple training iterations can also leak information about the contribution of individual clients. Most existing secure aggregation schemes are executed on one training iteration, i.e., they cannot protect against multi-round attacks. An exception is POSEIDON [113] which runs FL fully under encryption, but at the cost of significant computational overhead on clients’ and server’s side. Instead, So et al. [118] propose to organize clients in batches that can only be chosen together for a training iteration. This approach is orthogonal and fully compatible with ScionFL.

A.6 Poisoning Attacks & Defenses

Poisoning attacks can be categorized into untargeted and targeted attacks based on the goals of the attacker [51]. In the former case, the attacker aims to corrupt the global model so that it reduces or even destroys the performance of the trained model for a large number of test inputs, yielding a final global model with a high error rate [11, 51, 115]. In the latter case, the attacker aims to activate attacker-defined triggers that cause a victim model to do targeted misclassifications, which can then be activated in the inference phase [20, 121]. Notably, other classification results without the trigger behave normally and main task accuracy remains high. The second class of attacks is sometimes also referred to as *backdoor attacks* [10]. As discussed in §4, we consider only untargeted poisoning following the argument in [116]: This class of attacks is particularly challenging as service providers may not notice they are under attack given they do not know which accuracy is achievable in a fresh training of a new model. Also, even small accuracy reductions can lead to serious economical losses.

Below, we detail three state-of-the-art untargeted poisoning attacks, LIE [11], Fang [51], and Shejwalkar et al. [115], which are most relevant to our work.

- *Little is Enough (LIE) attack* [11]: In LIE [11], malicious clients manipulate their local updates by adding noise drawn from the normal distribution to “clean” updates they created following the normal training process to cause a disorientation. LIE assumes independent and identically distributed (iid) data and was tested against various robust aggregations such as trimmed-mean [135].

- *Fang et al.* [51]: The authors of [51] formulate their untargeted poisoning attack as an optimization problem where the manipulated updates aim at maximally disorienting the global model from the benign direction. However, they assume the adversary to either know or guess the deployed (robust) aggregation mechanism. Additionally, the attack was shown to be ineffective for iid as well as severely unbalanced non-iid training datasets [115].

- *Shejwalkar and Houmansadr* [115]: The attacks of [115] follow a similar idea as [51]: they maximize the distances between benign and malicious updates while using the evasion of outlier-based detection mechanisms as a boundary. Concretely, they formalize the following “Min-Max” optimization problem:

$$\operatorname{argmax}_{\gamma} \max_{i \in [n]} \|\nabla^m - \nabla_i\|_2 \leq \max_{i, j \in [n]} \|\nabla_i - \nabla_j\|_2 \quad (13)$$

$$\nabla^m = f_{\text{avg}}(\nabla_{\{i \in [n]\}}) + \gamma \nabla^P, \quad (14)$$

where $f_{\text{avg}}(\nabla_{\{i \in [n]\}})$ is the average gradient and $\gamma \nabla^P$ is the adversary’s perturbation vector, i.e., either the inverse unit vector

¹²<https://github.com/snucrypto/HEAAN>

of the (simulated) benign gradients, the inverse average standard deviations, or the average gradient with flipped sign of all updates. For details, we refer to §IV in [115].

Note that although the authors of [115] suggest several flavours of their attack based on different levels of adversarial knowledge, we compare to their Min-Max attack as it (i) does not make the unrealistic assumption that an adversary knows defenses in place and (ii) it is more destructive than LIE [11] for almost all datasets [115]. We do not consider Fang et al. [51]’s attack as it requires the guess of the robust aggregation rule, i.e., defense mechanism, which is unrealistic in a real-world deployment. Taking those considerations into account, we evaluate the robustness of ScionFL-Aura against the state-of-the-art Min-Max attack of [115] in §4.2.

Poisoning Defenses: Simple parameter-wise averaging is very sensitive to outliers and, thus, can easily hamper accuracy. Therefore, Byzantine-robust defenses aim to make FL robust against (untargeted) attacks. To do so, Krum [21] selects only one local update, namely the one with the closest $n - m - 2$ local updates as update for the global model, where n is the number of clients and m the number of anticipated malicious clients. Multi-krum [21] extends this idea to a selection of c (instead of just one) updates. Median [135] is another coordinate-wise aggregation selecting the coordinate-wise median of each update parameter. A straight-forward idea to assess (to some extent) if a specific gradient is malicious is to use an auxiliary dataset (rootset) at the aggregator to validate the performance of the updated global model [31, 47, 84]. FLTrust [31] and FLOD [47] use the ReLU-clipped cosine-similarity/Hamming distance between each received update and the aggregator-computed baseline update based on the auxiliary dataset. FLDetector [137] detects malicious clients by checking their model updates’ consistency based on historical model updates. RSA [84] uses an L_1 -norm-based regularization, which is also comparing to the aggregator-computed baseline update. The recently proposed Divider and Conquer (DnC) aggregation [115] combines dimensionality reduction using random sampling with an outlier-based filtering.

The so far discussed poisoning defenses are not compatible with secure aggregation protocols in a straight-forward manner or lead to an intolerable overhead. Only two works, namely FLAME [99] and BaFFLe [5] simultaneously consider both threats. Concretely, FLAME [99] uses a density-based clustering to remove updates with significantly different cosine distances (i.e., different directions) combined with clipping (for more subtle manipulations). BaFFLe [5] introduces a feedback loop enabling a subset of clients to evaluate each global model update, while being compatible with arbitrary secure aggregation schemes.

In our work, we combine (1) communication efficiency due to quantization, (2) data privacy due to secure aggregation, and (3) robustness due to a novel poisoning defense.

B Preliminaries

This section provides relevant details regarding the primitives used in this work. We begin with providing the necessary MPC background and protocols. The protocols are presented in a generic manner because our approach is not restricted to any specific MPC setting. Hence, some of the sub-protocols are treated as black-boxes

that can be instantiated using any efficient protocols in the underlying MPC setting. Since we consider dishonest majority setting to work with, we utilize the (semi-honest variant of) primitives from [17, 40, 41, 44] in a black-box manner.

B.1 MPC Protocols

In this section, we go over the details of the underlying MPC protocols used in our scheme. We consider three MPC servers, $\mathcal{S} = \{S_1, S_2, S_3\}$, to which the clients delegate the aggregation computation, as shown in Fig. 1. All the operations are carried out in either an ℓ -bit ring, \mathbb{Z}_{2^ℓ} , or a binary ring, \mathbb{Z}_2 . Before we go into the protocols, we provide additional details regarding the masked evaluation scheme [17, 86, 122] discussed in §3, starting with the sharing semantics.

Sharing Semantics: We use two different sharing schemes:

- (1) $[\cdot]$ -sharing. A value $v \in \mathbb{Z}_{2^\ell}$ is said to be $[\cdot]$ -shared among MPC servers in \mathcal{S} , if each server S_i , for $i \in [3]$, holds $v_i \in \mathbb{Z}_{2^\ell}$ such that $v_1 + v_2 + v_3 = v$.
- (2) $\langle \cdot \rangle$ -sharing. In this sharing, every $v \in \mathbb{Z}_{2^\ell}$ is associated with two values: a random mask $\lambda_v \in \mathbb{Z}_{2^\ell}$ and a masked value $m_v \in \mathbb{Z}_{2^\ell}$, such that $v = m_v + \lambda_v$. Here, the share of an MPC server is defined as a tuple of the form $(m_v, [\lambda_v])$.

Handling Decimal Values: While the MPC protocol that we utilize is designed over a ring architecture, the underlying FL algorithms handle decimal numbers. To address this compatibility issue, we employ the well-known Fixed-Point Arithmetic (FPA) technique [32, 94, 95], which encodes a decimal number in ℓ -bits using the 2’s complement representation. The sign bit is represented by the most significant bit, while the f least significant bits are kept for the fractional component. We use $\ell = 32$ bit values with $f = 16$ in this work.

We will now go over the MPC protocols used in our scheme. We assume that the protocols’ inputs are in $\langle \cdot \rangle$ -shared form, and that the output is generated in $\langle \cdot \rangle$ -shared form among the MPC servers.

Inner Product Computation: For simplicity, consider the multiplication of two values $x, y \in \mathbb{Z}_{2^\ell}$ as per the $\langle \cdot \rangle$ -sharing semantics. We have

$$\begin{aligned} z &= xy = (m_x + \lambda_x)(m_y + \lambda_y) \\ &= m_x m_y + m_x \lambda_y + m_y \lambda_x + \lambda_x \lambda_y. \end{aligned}$$

Since the λ values are independent of the underlying secret, the servers can compute $[\cdot]$ -shares of the term $\lambda_x \lambda_y$ during preprocessing [17, 40]. This enables the servers to locally compute $[\cdot]$ -shares of z during the online phase.

In addition to the above observation, since we operate over FPA representation, truncation [32, 95] must be performed in order to keep the result z in FPA format after a multiplication. For this, we use the truncation pair method [94], wherein a tuple of the form $(r, r/2^f)$ is generated in $\langle \cdot \rangle$ -shared form among the servers during preprocessing using the Π_{Tr} protocol [41]. Then, with very high probability, we have

$$z/2^f = (z - r)/2^f + r/2^f.$$

Hence, during the online phase, servers publicly open the value $(z - r)$ and apply the above transformation to obtain the $\langle \cdot \rangle$ -shares of truncated z , completing the protocol.

Now, in the case of the inner-product (Fig. 13), the task can be split into d multiplications and the result obtained accordingly. Furthermore, because the desired result is the sum of the individual multiplication results, servers can sum them and communicate in a single shot, saving communication cost [102].

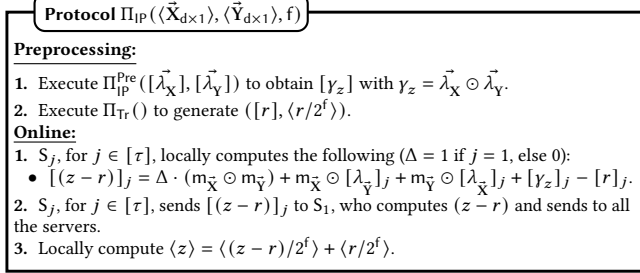


Figure 13: Inner product protocol.

Bit-to-Arithmetic Protocol: Given the Boolean sharing of $b \in \mathbb{Z}_2$, protocol Π_{BitA} computes the arithmetic sharing of the bit b over \mathbb{Z}_{2^ℓ} . As shown in Eq. 15, the arithmetic equivalent \tilde{b} for a bit $b = m_b \oplus \lambda_b$ can be obtained as

$$\tilde{b} = m_b \oplus \lambda_b = M_b + (1 - 2m_b) \cdot \Lambda_b. \quad (15)$$

Here, M_b and Λ_b denote the arithmetic equivalents of m_b and λ_b respectively. In our protocol shown in Fig. 14, MPC servers invoke Π_{BitA}^{Pre} protocol [41, 102] on the Boolean $[-]$ -shares of λ_b in the preprocessing phase to obtain its respective arithmetic shares. This enables the servers to locally compute an additive sharing of \tilde{b} during the online phase, as shown above. The rest of the steps proceed similar to the inner-product protocol and we omit the details.

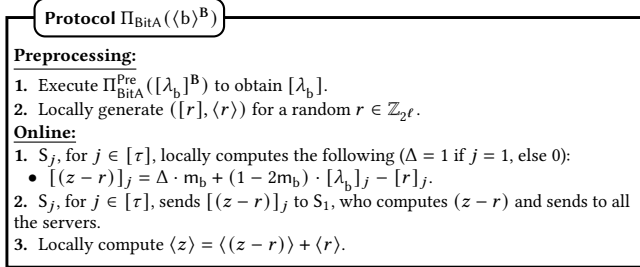


Figure 14: Bit-to-arithmetic conversion protocol.

To instantiate Π_{BitA}^{Pre} , we use SPDZ-style computations [70, 110], where oblivious transfer (OT) instances [9, 27, 39] are used among every pair of servers. Let Π_{OT}^{ij} denote an instance of 1-out-of-2 OT with S_i being the sender and S_j being the receiver. Here, S_i inputs the sender messages (x_0, x_1) while S_j inputs the receiver choice bit $c \in \mathbb{Z}_2$ and obtains x_c as the output, for $x_0, x_1 \in \mathbb{Z}_{2^\ell}$.

To generate the arithmetic sharing of λ_b from its Boolean shares in $[-]$ -shared form, a simple method would be to apply a 3-XOR using a daBit-style approach [110], but would result in 12 executions of 1-out-of-2 OTs. However, as pointed out in Prio+ [2], the cost could be further optimized due to the semi-honest security model being considered in this work rather than the malicious in [110]. Since Prio+ operates over two MPC servers, we extend their optimized daBit-generation protocol (cf. [2, daBitGen_p]) to our setting with three servers.

Given two bits $b_i, b_j \in \mathbb{Z}_2$, the arithmetic share corresponding to their product can be generated using one instance of Π_{OT}^{ij} with $(x_0 = r, x_1 = r + b_i)$ as the OT-sender messages and b_j as the OT-receiver choice bit. With this observation and using Eq. 8, servers can compute $[-]$ -shares corresponding to the bit λ_b using five OT invocations. The formal details appear in Fig. 15.

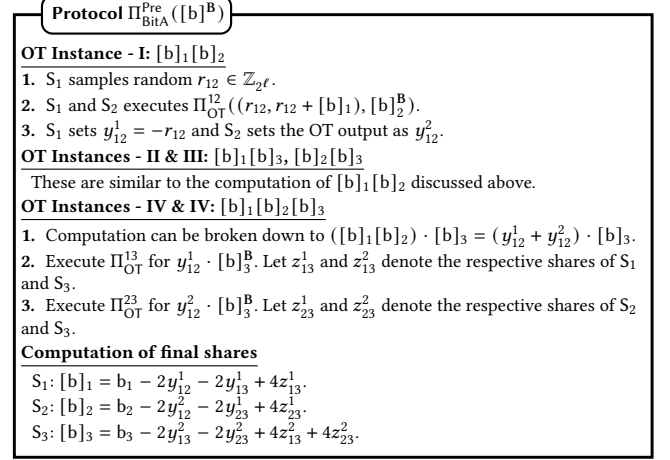


Figure 15: Bit-to-arithmetic preprocessing.

For the case of approximate bit conversion discussed in §3.2, the number of OT instances can be further reduced to three following Eq. 10. Concretely, the conversion involves computation of just $[b]_1 [b]_2 [b]_3$ and hence the OT instances II & III described in Fig. 15 are no longer needed.

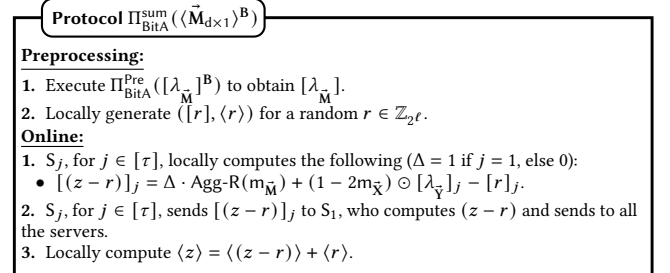


Figure 16: Bit-to-arithmetic sum protocol.

When computing the sum of bits directly, the online communication can be optimized following inner-product protocol and the resulting protocol Π_{BitA}^{Sum} is given in Fig. 16.

Bit Injection Protocol: Given a bit $b = m_b \oplus \lambda_b$ and $s = M_s + \Lambda_s$, the bit injection operation involves computing the value $b \cdot s$ that can be obtained as

$$\begin{aligned} b \cdot s &= (M_b + (1 - 2m_b) \cdot \Lambda_b) \cdot (M_s + \Lambda_s) \\ &= M_b M_s + M_b \Lambda_s + (1 - 2m_b) \cdot (\Lambda_b M_s + \Lambda_b \Lambda_s). \end{aligned} \quad (16)$$

Given a boolean vector $\vec{M}_{d \times 1}$ and an arithmetic vector $\vec{N}_{d \times 1}$ in the secret-shared form, protocol Π_{BI} computes the inner product of the two vectors, defined as $z = \vec{M} \odot \vec{N}$. This protocol is similar to the inner product protocol Π_{IP} presented in Fig. 13, with the main difference being that \vec{M} is a boolean vector.

During the preprocessing, servers first generate the arithmetic shares of $\lambda_{\vec{M}}$ from its boolean shares, similar to the bit-to-arithmetic protocol Π_{BitA} in Fig. 14. In this case, $\Pi_{\text{BI}}^{\text{pre}}$ is same as the $\Pi_{\text{IP}}^{\text{pre}}$ primitive discussed in Fig. 13. The remaining steps are similar to the Π_{IP} in Fig. 13 and we omit the details.

Protocol $\Pi_{\text{BI}}(\langle \vec{M}_{d \times 1} \rangle^B, \langle \vec{N}_{d \times 1} \rangle, f)$

Preprocessing:

1. Execute $\Pi_{\text{BitA}}^{\text{pre}}([\lambda_{\vec{M}}]^B)$ to obtain $[\lambda_{\vec{M}}]$.
2. Execute $\Pi_{\text{BI}}^{\text{pre}}([\lambda_{\vec{M}}], [\lambda_{\vec{N}}])$ to obtain $[\gamma_{\vec{Q}}]$ with $\gamma_{\vec{Q}} = \lambda_{\vec{M}} \circ \lambda_{\vec{N}}$.
3. Execute $\Pi_{\text{Tr}}()$ to generate $([r], \langle r/2^f \rangle)$.

Online:

1. S_j , for $j \in [\tau]$, locally computes the following ($\Delta = 1$ if $j = 1$, else 0):
 - $T_j^1 = \Delta \cdot (m_{\vec{M}} \circ m_{\vec{N}}) + m_{\vec{M}} \circ [\lambda_{\vec{N}}]_j$.
 - $T_j^2 = ((1 - 2m_{\vec{M}}) \circ m_{\vec{N}}) \circ [\lambda_{\vec{M}}]_j + (1 - 2m_{\vec{M}}) \circ [\gamma_{\vec{Q}}]_j$.
 - $[(z - r)]_j = T_j^1 + T_j^2 - [r]_j$.
2. S_j , for $j \in [\tau]$, sends $[(z - r)]_j$ to S_1 , who computes $(z - r)$ and sends to all the servers.
3. Locally compute $\langle z \rangle = \langle (z - r)/2^f \rangle + \langle r/2^f \rangle$.

Figure 17: Bit injection (sum) protocol.

B.2 Binomial Sum

Lemma B.1 (Expected Values). *Given $n, p \in \mathbb{Z}$, we have*

$$(1) \sum_{p=0}^n p \cdot \binom{n}{p} = n \cdot 2^{n-1}.$$

$$(2) \sum_{p=0}^{\lfloor n/2 \rfloor} 2p \cdot \binom{n}{2p} = \sum_{p=0}^{\lfloor n/2 \rfloor} (2p+1) \cdot \binom{n}{2p+1} = n \cdot 2^{n-2}.$$

PROOF. Consider the binomial formula for $(1+y)^n$, given by

$$\sum_{p=0}^n \binom{n}{p} y^p = (1+y)^n \quad (17)$$

Differentiating Eq. (17) with respect to y will give

$$\sum_{p=0}^n \binom{n}{p} p \cdot y^{p-1} = n \cdot (1+y)^{n-1} \quad (18)$$

Substituting $y = 1$ in Eq. (18) gives the first result (1). Similarly, setting $y = -1$ in Eq. (18) gives

$$\sum_{p=0}^n (-1)^{p-1} p \cdot \binom{n}{p} = 0 \quad (19)$$

Combining Eq. (19) with the first result (1) will give the second result (2). \square

B.3 Overhead of HSQ and KSQ Quantization

For being able to use an efficient GPU-friendly implementation of the randomized Hadamard transform, which we use for both rotating the gradients in HSQ and for calculating Kashin’s coefficients in KSQ, we require that the gradients’ size to be a power of 2. A simple solution to meet this requirement is padding. For example, for the LeNet architecture with $\approx 60k$ parameters, we can pad the gradient to $2^{16} = 65536$ entries with a small resulting overhead of $\approx 6.2\%$ (i.e., using ≈ 1.06 bits per coordinate instead of 1). However, a more sophisticated approach is to divide the gradient into decreasing power-of-two-sized chunks and inflate only the

last (smallest) chunk.¹³ For example, for the LeNet architecture, we can decompose it into chunks of size 32768, 16384, 8192, 4096, 512, that sum up to 61952 (with an additional overhead of two floats per chunk) with a resulting overhead of only $\approx 1.44\%$. Also, for Kashin’s representation, we use $\lambda = 1.15$ for each chunk (an extra 15% of space) as used in previous works (e.g., [126]). To summarize, we state these resulting overheads in Tab. 8.

Architecture	n	SQ	HSQ	KSQ
LeNet	61706	61706	62272	73024
ResNet9	4903242	4903242	4915456	5767424
ResNet18	11220132	11220132	11272192	12583040

Table 8: Exact number of bits used for different network architectures and quantization schemes compared to the baseline number of coordinates n .

C ScionFL: Additional Details

This section provides addition details of our FL framework ScionFL presented in §3. We begin with providing additional details regarding the approximate bit conversion discussed in §3.2.

C.1 Multi-bit Quantization Schemes

This section describes how our scheme ScionFL can be extended to support multi-bit linear quantization schemes, in which each coordinate is classified into more than two levels, resulting in each coordinate being represented by more than a single bit.

For instance, consider the quantization in TernGrad [130], where each coordinate is compressed to one of the three levels $\{-1, 0, 1\}$. Here, each coordinate can be represented using two bits, say b_1 and b_2 and the quantized level can be computed as $2b_1 - b_2$.

To use our scheme, each client C_i share the bits separately using the underlying boolean secret sharing scheme, i.e., $\langle b_1 \rangle_i^B$ and $\langle b_2 \rangle_i^B$. MPC servers use our instantiations of $\mathcal{F}_{\text{SecAgg}}$ functionality discussed in §3.1 to aggregate each of the bits and obtain the result in arithmetic sharing format, i.e. $\langle b_1 \rangle$ and $\langle b_2 \rangle$. The final result can be locally computed by the MPC servers as $2\langle b_1 \rangle + \langle b_2 \rangle$, since the underlying MPC protocol used in ScionFL is linear.

C.2 ScionFL-Aura: Additional Details

In this subsection, we provide additional details of our ScionFL-Aura. We begin with discussing some optimizations that could be incorporated to our defense presented in Alg. 1.

MPC-friendly Variant: A naive secure realization of ScionFL-Aura outlined in Alg. 1 utilizing MPC will yield an inefficient solution, particularly over a ring architecture. This is due to some of the algorithm’s non-MPC friendly primitives, which are discussed below.

(1) (*Line 5 in Alg. 1*). The computation of L_2 -norm within $L2\text{-NORMQ}$ (cf. Alg. 3 in §C.2) involves calculating the square root of a ring element, which corresponds to a decimal value. To alleviate this, we ask the clients to submit the L_2 -norm of their gradient vectors and the MPC servers verify them. To be more specific, the provided L_2 -norm is squared and compared to a squared- L_2 -norm computed by the MPC servers via a secure comparison protocol [32, 88].

¹³The size of the last chunk is kept above some threshold, e.g., 2^9 to keep the overhead of the scales small.

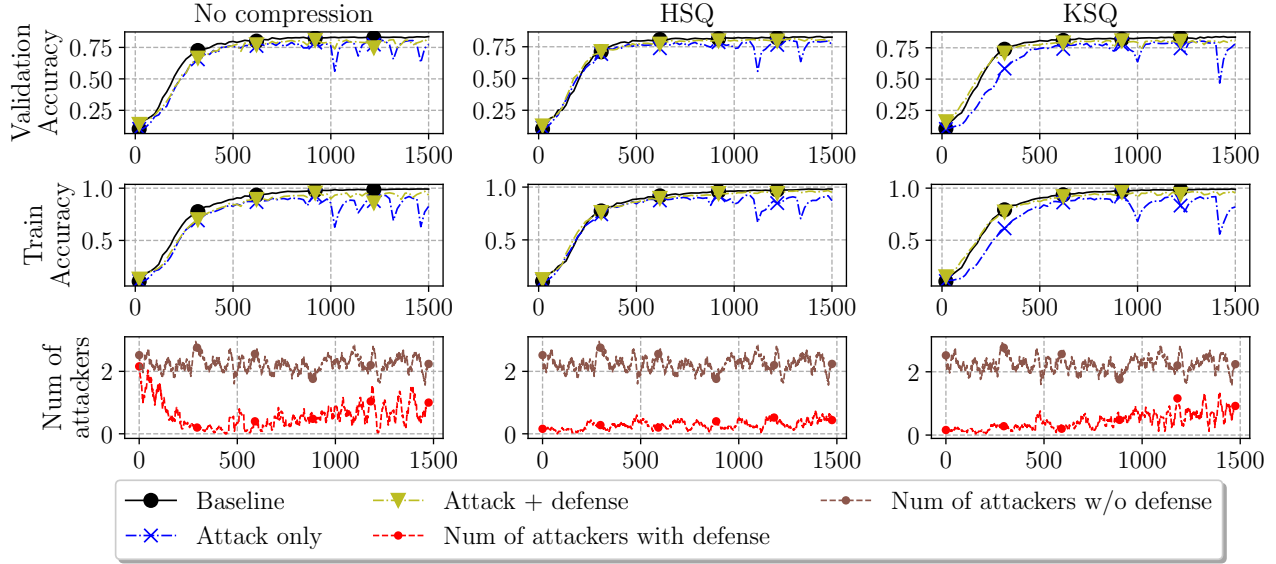


Figure 18: Effect of *Min-Max* attack [115] on training VGG11 with CIFAR10 with and without our defense ScionFL-Aura assuming 20% of $N = 50$ clients being corrupted.

(2) (Lines 11 & 12 in Alg. 1). When using L_2 -norm scaling, the scales of the gradient vector must be bounded if the corresponding L_2 -norm is greater than the limit. In particular, the procedure entails dividing the vector by its L_2 -norm. Because division is expensive in MPC over rings, we ask the client to submit the reciprocal of the L_2 norm as well, similar to the method suggested above. The provided value is validated by multiplying it by the L_2 norm supplied by the client and checking whether the product is a 1.

(3) (Line 16 in Alg. 1). The calculation of the cosine-distance between the gradient vector and the reference \vec{X}^{agg} requires computing the L_2 -norm of \vec{X}^{agg} and dividing by it, as shown in COSINE in Alg. 4. However, cosine distances are used to filter out the top- ψ vectors with the highest cosine distance, as shown in Alg. 1 (Line 18). As a result, we may safely disregard the division by the L_2 -norm of \vec{X}^{agg} when computing the cosine distance for our purpose.

In addition to the aforementioned optimizations, we notice that most of the values computed as part of the \vec{X}^{agg} computation in the AGGREGATE function (Line 2 in Alg. 1) can be reused in the next steps, thus lowering the overhead of the defense scheme over simple aggregation.

Sub-protocols: Here, we provide the details of the sub-protocols used in ScionFL-Aura (cf. Alg. 1 in §4.1).

Alg. 2 computes the aggregation of α quantized vectors. As shown in Eq. 4, the dequantized value of a vector \vec{Y} , given its quantized form $(\vec{\sigma}_Y, s_Y^{\min}, s_Y^{\max})$, can be computed as

$$\vec{Y} = s_Y^{\min} \oplus \vec{\sigma}_Y \circ (s_Y^{\max} - s_Y^{\min}).$$

The above operation essentially places s_Y^{\min} in those positions of the vector \vec{Y} with the corresponding bit in $\vec{\sigma}_Y$ being zero, and the rest with s_Y^{\max} .

Algorithm 2 Quantized Aggregation

```

1: procedure AGGREGATE( $\{\vec{\sigma}_{Y_i}, s_{Y_i}^{\min}, s_{Y_i}^{\max}\}_{i \in \alpha}$ )
2:    $\vec{Z} \leftarrow \vec{0}$ 
3:   for  $k \leftarrow 1$  to  $\alpha$  do
4:      $\vec{Z} \leftarrow \vec{Z} + (s_{Y_k}^{\min} \oplus \vec{\sigma}_{Y_k} \circ (s_{Y_k}^{\max} - s_{Y_k}^{\min}))$ 
5:   end for
6:    $\vec{Z} \leftarrow \vec{Z} / \alpha$ 
7:   return  $\vec{Z}$ 
8: end procedure

```

Algorithm 3 L_2 -Norm Computation (Quantized)

```

1: procedure L2-NORMQ( $\vec{\sigma}_Y, s_Y^{\min}, s_Y^{\max}$ )
2:    $\beta \leftarrow \text{LEN}(\vec{\sigma}_Y)$  // Dimension of  $\vec{\sigma}_Y$ 
3:    $N_O \leftarrow \text{SUM}(\vec{\sigma}_Y)$  // Number of ones in  $\vec{\sigma}_Y$ 
4:    $N_Z \leftarrow \beta - N_O$  // Number of zeros in  $\vec{\sigma}_Y$ 
5:   return  $\sqrt{N_Z \cdot (s_Y^{\min})^2 + N_O \cdot (s_Y^{\max})^2}$ 
6: end procedure

```

Alg. 3 computes the L_2 -norm of a quantized vector. As discussed in §2.4, a quantized vector \vec{Y}_σ consists of a binary vector $\vec{\sigma}_Y$ and the respective min. and max. scales s_Y^{\min}/s_Y^{\max} . In this case, we observe that the squared L_2 -norm can be obtained by first counting the number of zeroes and ones in the vector, denoted by N_Z and N_O respectively, followed by multiplying them with the square of the respective scales and adding the results, i.e. $N_Z \cdot (s_Y^{\min})^2 + N_O \cdot (s_Y^{\max})^2$. Furthermore, computing the number of ones N_O corresponds to the bit-aggregation of the vector \vec{Y} , for which our aggregation methods discussed in §3.1 can be utilized.

Alg. 4 is used to compute the cosine distance between a quantized vector \vec{Y}_σ and a reference vector \vec{S} . The cosine distance is given

Algorithm 4 Cosine Distance Calculation

```
1: procedure COSINE( $(\vec{\sigma}_Y, s_Y^{min}, s_Y^{max}), \vec{S}$ )
2:    $L_2^Y \leftarrow \text{L2-NORMQ}(\vec{\sigma}_Y, s_Y^{min}, s_Y^{max})$ 
3:    $L_2^S \leftarrow \|\vec{S}\|$  // Computes  $L_2$ -norm
4:    $\alpha \leftarrow \text{SUM}(\vec{S})$  // Sum of elements of  $\vec{S}$ 
5:    $\beta \leftarrow \text{INNER-PRODUCT}(\vec{\sigma}_Y, \vec{S})$ 
6:    $\gamma = s_Y^{min} \cdot \alpha + \beta \cdot (s_Y^{max} - s_Y^{min})$ 
7:   return  $\gamma / (L_2^Y \cdot L_2^S)$ 
8: end procedure
```

by $\frac{\vec{Y}_\sigma \odot \vec{S}}{\|\vec{Y}_\sigma\| \cdot \|\vec{S}\|}$, where $\|\cdot\|$ corresponds to the L_2 -norm of the input vector. Using Eq. 4, we can write

$$\begin{aligned} \vec{Y}_\sigma \odot \vec{S} &= (s_Y^{min} \oplus \vec{\sigma}_Y \circ (s_Y^{max} - s_Y^{min})) \odot \vec{S} \\ &= s_Y^{min} \odot \vec{S} + (\vec{\sigma}_Y \odot \vec{S}) \cdot (s_Y^{max} - s_Y^{min}). \end{aligned}$$

Thus, the inner product computation of $\vec{Y}_\sigma \odot \vec{S}$ reduces to computing $\vec{\sigma}_Y \odot \vec{S}$, followed by two multiplications.

Evaluation on VGG11: In addition to our results in §4.2, we evaluate the *Min-Max* attack on VGG11 trained with CIFAR10. The experimental setup is identical to §4.2. The results are shown in Fig. 18.

Similarly as for ResNet9 (cf. Fig. 12), the *Min-Max* attack substantially reduces the validation accuracy when training VGG11: We observe drops of up to 36.8%. However, on average, VGG11 is less impacted by the attack. Concretely, only 15% of the iterations observe a validation accuracy reduction of about 10% or more when using no compression. One third of the training rounds are impacted by about 10% or more when using Kashin’s representation (KSQ) while with the Hadamard transform (HSQ) only very few training rounds showed a significant accuracy reduction. Thus, HSQ seems to be more robust against untargeted poisoning.

With ScionFL-Aura, the accuracy reduction is still smaller for all variants. With HSQ, on average 0.28 malicious updates are included in global updated instead of 2.24 without defense. With respect to the validation accuracy, the difference between having no attack and employing ScionFL-Aura when under attack is less than 4% in almost all training iterations. When using KSQ, a global update includes just 0.44 malicious updates on average, and the attack impact is at least halved in two third of the training iterations.