
PICASO: Secure Aggregation for Federated Learning with Minimal Synchronization

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Preventing private data leakage is crucial in federated learning. Existing secure
2 aggregation (SA) protocols, which are the core protocols for privacy-preserving
3 federated learning, require clients to synchronize at multiple points, meaning
4 they must wait for other clients to send their messages before proceeding. This
5 synchronization ensures that inputs can be aggregated without compromising
6 privacy, while also accounting for client dropouts and message delays.

7 This work presents PICASO, abbreviated from Per Iteration Client At most Syn-
8 chronizes Once, a novel SA protocol minimizing synchronization overhead in
9 privacy preserving federated learning, aligning its communication pattern more
10 closely with that of non-private federated learning. PICASO outperforms previous
11 works like SecAgg, SecAgg+, MicroSecAgg, and Flamingo with server runtime
12 under 1 second for large clients. PICASO demonstrates viability by training various
13 models on different datasets.

14 We also detail extensions to PICASO to achieve various improvements over state-
15 of-the-art algorithms in two key areas - detecting and removing malicious clients,
16 and secure aggregation for heterogeneous datasets. Overall, PICASO presents an ef-
17 ficient, secure, and flexible federated learning solution minimizing synchronization
18 needs.

19 1 Introduction

20 Federated learning (FL) enables collaborative machine learning without sharing client data, by
21 aggregating local model updates at a central coordinator. However, recent works show that training
22 data can still be compromised from model updates alone [51, 82, 71], making secure aggregation
23 (SA) crucial for privacy-preserving FL. SA computes the sum of user inputs while keeping individual
24 inputs private. Like in traditional FL, the server typically selects a set of n clients for each iteration
25 with the server repeating the process until the model converges. Typically, a new set of clients
26 are chosen per iteration. Note that the number of clients n can range from a hundred to tens of
27 millions [58] and similarly the number of model parameters m can scale to millions [13]. The goal
28 is to securely train a global model. Critically, for SA protocols to help with federated learning, the
29 underlying protocol is needed to be robust to client dropouts. Furthermore, SA algorithms typically
30 work over integers or field elements while the weights produced by an ML model are floating point
31 values. Therefore, one often needs to quantize the weights and show that the model produced by the
32 SA protocols still preserve accuracy. Existing SA algorithms broadly rely on Differential Privacy
33 (DP) [44], Homomorphic Encryption (HE) [75, 96], or secure multiparty computation techniques
34 [8, 13, 7, 68, 66, 86, 83, 98, 56, 94, 67, 63].

35 SecAgg [13] introduced a practical solution for privacy-preserving horizontal federated learning. The
36 protocol's core idea involves pairwise masking seeds $s_{u,v}$ shared between clients $u, v \in \mathcal{U}$, where \mathcal{U}
37 is the set of all users. Each client u masks its input using $\sum_{v < u} s_{u,v} - \sum_{v > u} s_{u,v}$. Note that a client

38 $v < u$ would generate its mask by subtracting $s_{u,v}$. Therefore, it is easy to see that the pairwise masks
 39 cancel out in aggregation, i.e., $\sum_{u \in \mathcal{U}} (\sum_{v < u} s_{u,v} - \sum_{v > u} s_{u,v}) = 0$. While a particular user u maybe
 40 offline, the remaining clients would still have used their pairwise mask with u in the aggregation.
 41 Therefore, $s_{u,v}$ for an offline u and any online v needs to be secret shared with other clients to allow
 42 the server to reconstruct $s_{u,v}$ and then remove the mask. Unfortunately, a server could label an online
 43 client u as offline which would give the server $s_{u,v}$, allowing it to unmask an online user’s inputs.
 44 Therefore, a client also uses a self-mask s_u to mask its inputs. This s_u is also secret-shared and is
 45 reconstructed should u be online. As is obvious, SecAgg involves multiple rounds of computation
 46 such as to establish pairwise masks, secret share the masks, sending the masked inputs, and then
 47 reconstructing the sum. Therefore, for n clients and a vector of size m , the protocol requires $O(n^2m)$
 48 computation on the part of the aggregator, $O(mn)$ for each client. Subsequent works have focused
 49 on reducing the complexity through various assumptions and techniques. Here the vector m can be
 50 viewed as the number of parameters in the model, i.e., the inputs to the secure aggregation algorithm.

51 This work aims to address a critical scalability issue with existing SA algorithms. Prior works often,
 52 including SecAgg that was described above, require clients to synchronize their participation with
 53 others, an artifact of techniques where clients mask their inputs but must share masks with a quorum
 54 of clients to facilitate unmasking, if the client was unavailable later. This expensive ritual of sharing
 55 masks induces a bottleneck absent in non-private training, where clients simply train the model and
 56 send updates without additional synchronization.

57 Our main contribution is PICASO, a secure aggregation protocol where each client synchronizes
 58 at most once per training iteration. The key idea is that in iteration ℓ , client i masks its input x_i by
 59 computing $y_i = x_i + \text{GenerateMask}(k_i, \ell)$, where k_i is its *private* key. Client i sends y_i to the server
 60 and $\text{GenerateMask}(k_i, \ell)$ to a separate "collector" party. Our model supports dynamic selection of
 61 collector. They are stateless and run a deterministic computation, simpler than the server’s own
 62 computation. A simple way of choosing a collector would be to use a randomness beacon [37] and
 63 the Algorithm 1 from Flamingo [68, Lines 2-5]. The collector aggregates the masks from all clients
 64 and sends their sum to the server. The server then reconstructs $\sum_i x_i$ using the masked inputs and the
 65 sum of masks. Unlike SecAgg (and its subsequent works), PICASO does not require sending masks
 66 to multiple parties or secret sharing, reducing synchronization overhead. It only needs to synchronize
 67 to identify the collector for that iteration. Looking ahead, the collector acts can be viewed as a single
 68 "decryptor" [68] or "committee member" [63], receiving information from clients, condensing it, and
 69 communicating the result to the server for secure aggregation (Figure 1a). In other words, PICASO
 70 utilizes one intermediate party while SecAgg employed n with subsequent works employing $\log n$
 71 intermediate parties.

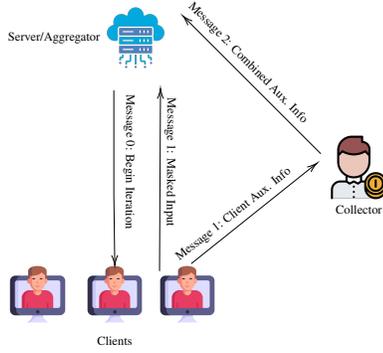
72 Asymptotically, PICASO’s client computation cost is $O(m)$, where m is the input vector length (e.g.,
 73 the number of weights of the model), while the server and collector computation cost is $O(mn)$,
 74 where n is the number of clients (Table 2). PICASO offers several attractive features:

- 75 • Dropout tolerance: Any number of selected clients can opt out without increasing computa-
 76 tional burden on remaining clients or requiring additional interaction.
- 77 • Collusion resistance: Privacy of honest users’ inputs is preserved even if an adversary
 78 corrupts any number of clients and the aggregator.
- 79 • Scalability and dynamism: New clients can join without an expensive setup phase, needing
 80 only public parameters and the aggregator’s iteration key.
- 81 • Enhanced privacy: Input privacy is maintained against both collector and aggregator, pro-
 82 vided they do not collude. The collector can change in each iteration, facilitated by a
 83 randomness beacon which in tern prevents server manipulation.

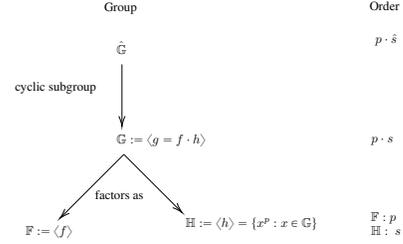
84 We also microbenchmark PICASO, comparing with the state-of-the-art secure aggregation algorithms
 85 to demonstrate competitive performance. For example, PICASO’s server computation time is $< 1s$,
 86 even for large number of clients besting prior work. We also conduct extensive experiments on FL
 87 benchmark datasets to demonstrate that PICASO preserves performance, while guaranteeing privacy.

88 Further, PICASO can easily be extended to offer:

- 89 • a constant-round protocol to detect and remove malicious clients (i.e., sending inconsistent
 90 or incorrect messages), improving on the state-of-the-art ACORN which requires $O(\log n)$
 91 rounds where n is the number of clients.



(a) The PICASO system model operates in iterations. Each iteration begins with the server sending a message to initiate the process (Message 0). In response, clients train the model on their local data, obtain updates, and mask the input. *Concurrently*, clients communicate with both the server and the collector (Message 1): masked input is sent to the server, while auxiliary information is transmitted to the collector. Upon receiving information from all clients, the collector combines these into a single value. Finally, this consolidated data is sent to the server (Message 2), concluding the iteration.



(b) A brief overview of the class group framework we employ. Here, $\hat{\mathbb{G}}$ is group, whose order is $\hat{s} \cdot p$, such that \hat{s} and p are co-prime. Further, s divides \hat{s} and is the order of the group \mathbb{G} , which is generated by g and is denoted as $\mathbb{G} := \langle g \rangle$. Similarly, \mathbb{H} is a subgroup of \mathbb{G} , generated by h whose order is s , while \mathbb{F} has order p and is generated by f . We have $g = f \cdot h$. Further, \hat{s} (and s) is unknown but an upper bound \bar{s} is known. The last property we will rely on is that discrete logarithm is efficient in the subgroup \mathbb{F} .

Figure 1: The backbone of PICASO - the communication system model and the CL framework.

- a secure aggregation protocol supporting heterogeneous datasets via robust stochastic averaging [64], which improves upon DReS-FL [80] as DReS-FL requires the entire dataset to be secret shared among clients, which we avoid.

1.1 Related Work

Secure Aggregation Using Differential Privacy and Homomorphic Encryption. A simple approach to differential privacy (DP) is local DP [38], where clients add noise to their data before sending it to the server. This has been deployed by major tech companies [40, 1]. However, research shows that such data perturbation may reduce accuracy. Our techniques can be composed with DP solutions, using secure aggregation (SA) algorithms to mask noisy local inputs [57]. Meanwhile, BatchCrypt [96] employed homomorphic encryption (HE), building on earlier work. However, it required all clients to use the same key, posing a significant privacy risk.

SA using Multiparty Computation. Secure multiparty computation (MPC) preserves privacy and accuracy by computing over encrypted data. Early works on Private Stream Aggregation [81] focused on secure summing of streaming data. Following SecAggBonawitz et al. [13], Federated Learning protocols with dropout resilience were developed, but multiple interaction rounds increased dropout risk. Subsequent works [8, 7, 68, 66, 86, 83, 98, 56, 94, 67, 63] have focused on reducing the number of intermediate parties to $\log n$, or reusing the masked secret sharing across multiple iterations to reduce round count. This is summarized in Table 1 where we compare various protocols with respect to the following properties: (a) the number of rounds of interaction, (b) whether it can tolerate client dropouts, (c) on whether the aggregate value can be efficiently recovered, (d) public setup for security assumption, and (e) number of intermediate parties needed to help with the aggregation. Multi-server settings [33, 4] face challenges with long inputs due to increased communication and computation demands. Our approach assumes no server-collector collusion, a weaker assumption as the collector changes each iteration and performs less computation.

2 System Model and Relevant Background

We consider a federated learning framework, as shown in Figure 1a. There exist n clients, with each client C_i owning a dataset D_i . The server holds the ML model Θ . In FL, the server first sends Θ to clients, and each client trains its local dataset on D_i to generate the updated weights m_i . Meanwhile,

Table 1: Comparison of various Secure Aggregation Algorithms based on MPC.

	# Rounds	Dropout Resilience	Efficient Aggregation	Public Setup	# Additional Parties
[81]	1	✗	✗	✓	0
[55]	1	✗	✓	✗	0
[9]	1	✗	✗	✗	0
[62]	1	✓	✓	✗	1
SecAgg[13]	3	✓	✓	✓	n
SecAgg+[8]	3	✓	✓	✓	$\log n$
MicroSecAgg[49]	1+2	✓	✗*	✓	$\log n$
LERNA[63]	1+2	✓	✓	✓	$\log n$
Flamingo[68]	1+2	✓	✓	✓	$\log n$
PICASO	1	✓	✓	✓	1

Table 2: Comparison of asymptotic complexity of some secure aggregation protocols. Note that in PICASO, the collector performs $O(mn)$ computation and communication.

	Client		Server	
	Computation	Communication	Computation	Communication
SecAgg[13]	$O(mn + n^2)$	$O(m + n)$	$O(mn^2)$	$O(mn + n^2)$
SecAgg+[8]	$O(m \log n + \log^2 n)$	$O(m + \log n)$	$O(n \log^2 n + mn \log n)$	$O(mn + n \log n)$
SASH[66]	$O(m + n^2)$	$O(m + n)$	$O(m + n^2)$	$O(mn + n^2)$
PICASO	$O(m)$	$O(m)$	$O(mn)$	$O(mn)$

120 the server computes the average of these model updates $\{m_i\}$ to update its global model to Θ' . In the
 121 next iteration Θ' is sent back for the next update. The goal is to use the collector to ensure that the
 122 weights m_i remain secret while still allowing the server to compute the average, and thereby the new
 123 model Θ' .

124 **Threat Model.** Our threat model follows the long line of prior works whereby an adversary can: (a)
 125 corrupt the server *or* the collector, (b) corrupt clients which enables the adversary to choose the client
 126 inputs for an iteration. The goal is to ensure that the honest users' inputs remain private with only
 127 their sum being leaked. Our protocols are described in this setting, like all prior work. Note that prior
 128 works such as Flamingo [68], or LERNA [63] did not guarantee privacy when all the intermediate
 129 parties collude with each other. Similarly, we allow for the collector to corrupt clients and guarantee
 130 the security against a corrupted collector. If the server and collector collude at an iteration, however,
 131 there is no privacy for that iteration.¹ Importantly, our collector can change from every iteration to
 132 iteration and is selected by a randomness beacon using an algorithm similar to how the set is chosen
 133 in Flamingo [68, Algorithm 1]. This is similar to how validators are chosen in some proof of stake
 134 blockchains [45].

135 **Modeling Security.** We model security against both a corrupt server and a corrupt collector. A
 136 corrupt server can adaptively compromise clients and collude with them, issue arbitrary encryption
 137 messages for honest parties in any iteration, and receive the collector's combined information at
 138 each iteration, but cannot corrupt the collector. The adversary selects honest clients H_1, \dots, H_t
 139 and provides two input sets: $\{x_1, \dots, x_t\}$ and $\{x'_1, \dots, x'_t\}$, where $\sum x_i = \sum x'_i$. The challenger
 140 randomly selects and encrypts one set for the target time period τ . The adversary's goal is to determine
 141 which set was chosen with probability significantly exceeding $1/2$.

142 Meanwhile, for privacy against a corrupt collector, the adversary receives individual communication
 143 sent by the clients to the collector. It can corrupt clients and also issue encryption queries, as before.
 144 It cannot corrupt the server but can adaptively issue the above queries. The challenge is the same as
 145 for a corrupt server - to distinguish between honest user inputs. Since it does not receive the final
 146 result, the challenge sets need not have the same sum.

147 **CL Framework.** Cryptographic protocols often use cyclic groups G of prime order q , generated
 148 by g_q , i.e., $G := \{1, g_q, \dots, g_q^{q-1}\}$. The Discrete Logarithm (DL) Assumption [2] states that given
 149 $X \in G$, finding x where $g_q^x = X$ is computationally infeasible.² The Decisional Diffie-Hellman

¹In such a case the server can learn the individual client model updates. To protect against such an attack the best that the clients can do is to add differential private noise to their updates.

²Small x are recoverable, as in [49, 81, 9].

150 (DDH) Assumption [14] posits that given g_q, g_q^x, g_q^y , distinguishing g_q^{xy} from a random element in G
 151 is computationally infeasible.

152 The CL Framework [24, 27–29, 17] introduces the idea of a composite order group, where the order
 153 is unknown, but there is a subgroup of known prime order where the discrete logarithm computation
 154 is easy. This framework utilizes the group where DL is easy to ensure correctness and eventual
 155 message recovery, and the group where DL is hard to achieve security. The framework is summarized
 156 in Figure 1b.

157 The security property relies on a modified DDH assumption (Definition 1) involving indistinguishability
 158 between elements from groups \mathbb{G} and \mathbb{H} within a composite order group. While their orders
 159 are unknown, upper bounds exist. The input space is in \mathbb{F} , and the key space in \mathbb{H} . Distributions \mathcal{D}_G
 160 and \mathcal{D}_H are based on these upper bounds, with \mathcal{D}_G (resp. \mathcal{D}_H) being statistically indistinguishable
 161 from \mathbb{G} 's (resp. \mathbb{H} 's) exponent space. Typically, $\mathcal{D}_H := 0, \dots, B - 1$ where $B = 2^{40} \cdot \bar{s}$ where \bar{s} is
 162 the upperbound of the order of \mathbb{H} , shown by [89] to be 2^{-40} -close to uniform. The security property
 163 relies on a modification of the DDH assumption (see Definition 1), where the indistinguishability is
 164 between elements from two different groups, \mathbb{G} and \mathbb{H} , within the composite order group. However,
 165 the orders of \mathbb{G} and \mathbb{H} are unknown, but there are upper bounds on their orders.

166 **Definition 1** (DDH-f Assumption). *Let $(\hat{\mathbb{G}}, \mathbb{G}, \mathbb{H}, \mathbb{F}, \bar{s})$ be the class group as defined in Figure 1b.*
 167 *Then, the following two distributions are computationally indistinguishable, i.e., there is no “efficient”*
 168 *attacker who can distinguish whether it is the first or the second distribution that a sampled value*
 169 *comes from, with a probability greater than half. Here, $x, y \leftarrow_s \mathcal{D}_H, u \leftarrow_s \mathbb{Z}/p\mathbb{Z}$*

$$\{(h, h^x, h^y, h^{xy})\} \approx_c \{(h, h^x, h^y, h^{xy} \cdot f^u)\}$$

170 We refer the readers to Bouvier *et al.* [16] and Tucker [89] for a detailed exposition on class groups,
 171 techniques, and its extensive use in cryptography. They also survey the utility of CL framework in
 172 building other cryptographic primitives.

173 3 PICASO

174 In this section, we present an efficient privacy-preserving aggregation scheme called PICASO where
 175 each iteration involves a client speaking at most once. We begin by describing an additively homo-
 176 morphic masking algorithm. We then instantiate this in the CL framework. This is a generalization of
 177 the version presented in the introduction. We then present our complete description of PICASO and
 178 formally prove it secure in the random oracle model.

179 3.1 Homomorphic Masking Algorithm

180 Let k_i denote the secret key of Client i . Let k_0 denote the secret key of the aggregator. Further, for
 181 iteration ℓ , let $pk_{i,\ell}$ (resp. $pk_{0,\ell}$) denote the public key of client i (resp. the server) for iteration ℓ .
 182 Then, we require the following properties of our algorithm GenMask:

- 183 • The masking function can be computed using two different ways, i.e., $\text{GenMask}(pk_{i,\ell}, k_0) =$
 184 $\text{GenMask}(pk_{0,\ell}, k_i)$
- 185 • Homomorphic over public key space, i.e., $\prod_{i=1}^n \text{GenMask}(pk_{i,\ell}, k_0) =$
 186 $\text{GenMask}(\prod_{i=1}^n pk_{j,\ell}, k_0)$

187 Further, we require that the generated mask is pseudorandom, i.e., $\text{GenMask}(pk_{0,\ell}, k_i)$ appears
 188 random provided the adversary cannot compute the mask on its own which requires the knowledge of
 189 $(pk_{0,\ell}, k_i)$ or $(pk_{i,\ell}, k_0)$.

190 **Construction 1** (Homomorphic Masking Algorithm). Let $\mathcal{H} : \{0, 1\}^* \rightarrow \mathbb{H}$ be a hash function that
 191 maps strings to the unknown order subgroup of \mathbb{G} .³ Then, for secret key $k_i \leftarrow_s \mathcal{D}_H$ for $i = 0, \dots, n$,
 192 we can define, for iteration ℓ , $pk_{i,\ell} := \mathcal{H}(\ell)^{k_i}$ and the mask value as $\text{mask}_{i,\ell} := \mathcal{H}(\ell)^{k_0 \cdot k_i}$

193 We now show that the construction satisfies the required properties:

³Note that for our purposes we can simply begin by hashing the input to an element in \mathcal{D}_H , and then raising the group generator h to this value. This is because the knowledge of the discrete logarithm is not detrimental. However, [77] present additional methods to hash into a group of unknown order, in a way that the discrete logarithm is unknown.

- 194 • $\text{mask}_{i,\ell} = \mathcal{H}(\ell)^{k_0 \cdot k_i} = \text{pk}_{i,\ell}^{k_0} = \text{pk}_{0,\ell}^{k_i}$
- 195 • $\text{mask}_{i,\ell} \cdot \text{mask}_{j,\ell} = \mathcal{H}(\ell)^{k_0(k_i+k_j)} = (\text{pk}_{i,\ell} \cdot \text{pk}_{j,\ell})^{k_0}$

196 For a particular iteration ℓ , an adversary is either given $\mathcal{H}(\ell), \text{pk}_{0,\ell}, \text{pk}_{i,\ell}, \text{mask}_{i,\ell}$ or
 197 $\mathcal{H}(\ell), \text{pk}_{0,\ell}, \text{pk}_{i,\ell}, U$ where $U \leftarrow_{\mathcal{S}} \mathbb{G}$. This follows from the DDH-f assumption, which we define in
 198 Definition 1. Looking ahead, this pseudorandom mask will be used to mask the client input and
 199 thereby guaranteeing privacy.

200 3.2 Formal Description of PICASO

201 We first informally describe the protocol. At iteration ℓ , the server sends a message identifying
 202 clients who are participating in that round of interaction. In this message, it also includes its
 203 iteration public key $\text{pk}_{0,\ell}$. Client i , with its input $x_{i,\ell}$, first encodes it as $f^{x_{i,\ell}}$. Recall that f is
 204 the generator of the cyclic, prime-order group \mathbb{F} where discrete logarithm is easy, i.e., given this
 205 encoding, there exists an efficient algorithm that outputs $x_{i,\ell}$. Once encoded, it computes the mask
 206 $\text{mask}_{i,\ell} = \text{GenMask}(\text{pk}_{0,\ell}, k_i)$. It sends to the server the masked input $\text{ct}_{i,\ell} = \text{mask}_{i,\ell} \cdot f^{x_{i,\ell}}$.
 207 Meanwhile, it also sends to the collector $\text{pk}_{i,\ell}$.

208 The collector simply multiplies all of the clients' iteration public keys to compute $\text{AUX}_\ell = \prod \text{pk}_{i,\ell}$.
 209 AUX_ℓ is sent to the server. The server does the following: multiplies all of the masked inputs $\prod \text{ct}_{i,\ell}$
 210 and divides it by $\text{GenMask}(\text{AUX}_\ell, k_0)$. It then applies the efficient discrete logarithm to compute the
 211 aggregate. Formally, we present in Construction 2.

212 **Construction 2** (PICASO Protocol for iteration ℓ). The protocol description is as follows:

- 213 • **One-Time Setup Phase:**
 214 Transparent Setup is executed and outputs $\text{pp} = (\widehat{\mathbb{G}}, \mathbb{F}, p, \mathcal{D}_H, \mathcal{D}_G, \bar{s}, g, h, f, \mathcal{H} : \{0, 1\}^* \rightarrow \mathbb{H})$
- 215 • **Begin Iteration:** Server, with key k_0 , computes $\text{pk}_{0,\ell} = \mathcal{H}(\ell)^{k_0}$ and sends to the chosen clients and
 216 collector.
- 217 • **Encryption Phase:** Each online client $C_i \in \mathcal{OL}_\ell$ with key k_i and input $x_{i,\ell}$ does the following:
 - 218 – Compute $\text{mask}_{i,\ell} := \text{GenMask}(\text{pk}_{0,\ell}, k_i)$
 - 219 – Compute masked input $\text{ct}_{i,\ell} = f^{x_{i,\ell}} \cdot \text{mask}_{i,\ell}$
 - 220 – Compute public key $\text{pk}_{i,\ell} = \mathcal{H}(\ell)^{k_i}$
 - 221 – $C_i \rightarrow \text{Server}$: $\text{ct}_{i,\ell}$
 - 222 – $C_i \rightarrow \text{Collector}$: $\text{pk}_{i,\ell}$
- 223 • **Collection Phase:** Collector computes $\text{AUX}_\ell = \prod_{i \in \mathcal{OL}_\ell} \text{pk}_{i,\ell}$
 224 **Collector** \rightarrow **Server**: AUX_ℓ
- 225 • **Aggregation Phase:** Server computes:
 - 226 – Compute $Y_\ell := \prod_{i \in \mathcal{OL}_\ell} \text{ct}_{i,\ell}$
 - 227 – Compute $X_\ell := \text{GenMask}(\text{AUX}_\ell, k_0)$
 - 228 – Compute $\text{Sum}_\ell := Y_\ell / X_\ell$
 - 229 – Take discrete log of Sum_ℓ , which is efficient.

230 We omit the proof due to space constraints. However, the intuition for security comes from the fact
 231 that: (a) the honest user's key is chosen by the honest user and is unknown to the adversary, (b) for
 232 such a random key, the mask generated is indeed pseudorandom under the DDH-f assumption, and
 233 (c) such a pseudorandom mask will blind the honest client's inputs.

234 **Remark 1.** Observe that it is possible that the client's communication to either the server or the
 235 collector is dropped due to network issues. In this situation, the collector's information relayed to the
 236 server does not yield correct aggregate. To handle such situation, the server and the collector can
 237 engage in one additional round of communication, per iteration. In this round, the collector first sends
 238 a list of clients from whom it has received communication. The server respond with the intersection
 239 of the collector's list with its own list of clients. Finally, the collector "collects" only with respect to
 240 this set of clients.

241 **Remark 2.** Note that each $\text{pk}_{i,\ell}$ is pseudorandom, if k_i is unknown. However, masking only with
 242 $\text{pk}_{i,\ell}$ is insufficient for security against the collector. This is because the collector receives the masked
 243 input (masked by $\text{pk}_{i,\ell}$) and $\text{pk}_{i,\ell}$ for ever honest client i . Therefore, the collector can easily recover
 244 the input. This prompts the need for a server's iteration public key, generated as a function of its
 245 secret key k_0 .

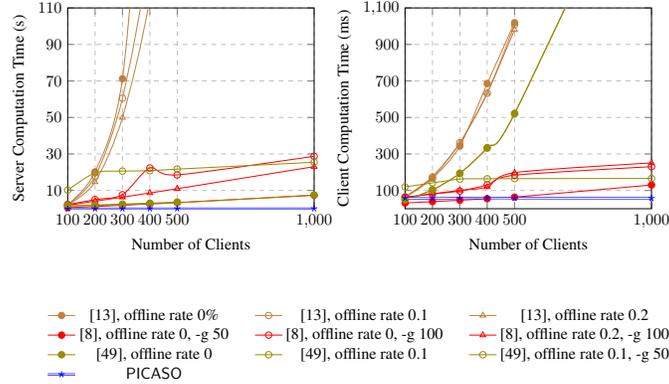


Figure 2: Measure of Server and Client Computation Time as a function of number of clients across various aggregation algorithms.

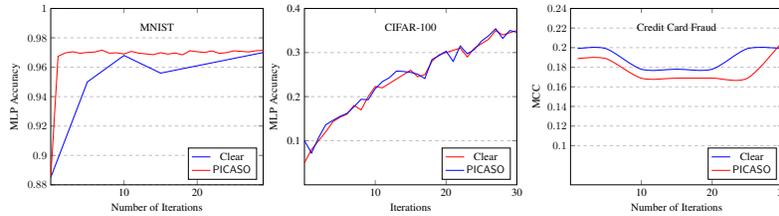


Figure 3: Performance Measurement of PICASO for FL Tasks.

246 4 Experiments

247 In this section, we perform different experimentation to demonstrate the efficiency, accuracy, and
 248 privacy of PICASO. All our experiments were carried out on an Apple M1 Pro CPU with 16 GB of
 249 unified memory, without any multi-threading or related parallelization. More details can be found in
 250 Section B.

251 **Microbenchmarking Secure Aggregation.** We benchmark the client and server computation time
 252 of our protocol against existing state-of-the-art solutions, including [13], [8], MicroSecAgg [49], and
 253 PICASO. Additionally, we compare our results with specific parameter choices from prior work,
 254 such as grouping operations (clients share inputs with 50 or 100 parties) and offline rate (parties can
 255 go offline during the protocol). These settings are not applicable to PICASO. Our reported timing is
 256 taken as a mean of 20 iterations.

257 As shown in Figure 2, our client computation time is significantly better than [13] and [8], and
 258 comparable to MicroSecAgg. However, unlike MicroSecAgg, our protocol does not incur offline
 259 waiting times due to multiple rounds of participation. For instance, when there are 100 clients,
 260 MicroSecAgg requires at least 30ms of offline time, which increases with more clients. Additionally,
 261 MicroSecAgg limits input size to achieve server efficiency, supporting only small model updates or
 262 quantized large model updates. Figure 2 demonstrates that PICASO’s server running time is under 1
 263 second, thanks to a single-round protocol with efficient aggregate recovery. This outperforms all other
 264 protocols. Any additional communication required to capture Remark 1 has a negligible impact on
 265 computation time, as it only involves gathering the list of clients and communicating with the collector.
 266 SASH [66] combines the secure aggregation protocol SecAgg [13] with a seed-homomorphic PRG
 267 to enhance efficiency for encrypting large input vectors. However, their performance is dominated
 268 by SecAgg, which we significantly outperform. Combining SASH with PICASO could achieve
 269 efficient round communication and improved server computation time, optimizing for input size
 270 scaling. Finally, PICASO requires 56 bytes of bandwidth for each of the following: server public
 271 key, masked input, while requiring 32 bytes for client’s iteration public key sent to the collector, and
 272 information sent to the server from the collector.

273 **Benchmarking FL Models.** To demonstrate PICASO’s viability for federated learning, we use
 274 it to train various models. We train a logistic regression model on Kaggle Credit Card Fraud

275 Dataset [76]. Figure 3 shows PICASO’s MCC versus clear learning for varying clients and iterations.
 276 With the accuracy multiplier, PICASO’s MCC is very close to clear learning and even outperforms
 277 sometimes. The highly unbalanced dataset demonstrates PICASO can achieve strong performance
 278 even in challenging real-world scenarios. We then train a vanilla multi-layer perceptron (MLP)
 279 classifier on three datasets: MNIST, CIFAR-100. The details of the datasets, including quantization
 280 and license can be found in Table 4. The MLP accuracy, as a function of the iteration, is plotted
 281 in Figure 3. Our experiments demonstrate that PICASO preserves accuracy, while ensuring the
 282 privacy of client data. Note that vanilla MLP classifiers do not typically offer good performance for
 283 CIFAR datasets, but note that the goal of our experiments was to show that PICASO does not impact
 284 accuracy.

285 5 Extensions to PICASO

286 We defer additional extensions to Section D, due to space constraints.

287 5.1 Robust PICASO

288 PICASO requires clients to send iteration public keys to the collector and masked inputs to the server,
 289 potentially allowing malicious actors to disrupt aggregation by using inconsistent keys. While secure
 290 aggregation has been widely studied, less focus has been on detecting and mitigating malicious
 291 behavior. Prior works in this domain are limited to:

- 292 • ACORN [7]: Offers a constant-round version detecting malicious behavior (aborting on
 293 detection) and a non-constant round version removing malicious inputs.
- 294 • RoFL [67] and ACORN: Use zero-knowledge proofs (e.g., Bulletproofs [19] and improve-
 295 ments [43]) to prevent malicious input injection.

296 The latter requires that the secure aggregation algorithm still proceeds, after having removed malicious
 297 clients. Indeed, PICASO ensures privacy of inputs, even if the server is corrupt and chooses to corrupt
 298 various users. Similarly, against corrupt collector who can corrupt users and inject messages into the
 299 system. In this section, we show how to augment PICASO to detect and remove malicious clients as
 300 described above. In Algorithm 3, we only present the additional proving steps by the client and the
 301 verification steps for the collector. In the construction, \mathcal{C} is representative of the challenge space and
 302 integer A is chosen as a function of the size of \mathcal{C} . We refer the reader to [17, §5.2] for details about
 303 the proof system and its correctness. Here, A is set to be an integer such that the size of challenge
 304 space $C := |\mathcal{C}|$ is negligibly small when compared to A , i.e., C/A is negligible.

305 Signatures can be employed to ensure the collector transmits only client-authenticated information to
 306 the server, mitigating malicious collector behavior. Our protocol can be enhanced with range-proof
 307 techniques from ACORN [7]. Notably, our inputs are encoded in prime-order subgroup \mathbb{F} , which
 308 can be composed with standard Pedersen commitments [74] using a prime order cyclic subgroup G
 309 where the DDH assumption holds.

310 **Construction 3** (Additional Steps in Robust-PICASO). We assume that there is an hash function $\mathbb{H} : \{0, 1\}^* \rightarrow$
 311 \mathcal{C} . Here, $A := 2^{40} \cdot |\mathcal{D}_H| \cdot C$ and $[A] := \{0, \dots, A - 1\}$. We set C to be 2^{128} .

312 **Proof Generation:** Each online client C_i

- 313 • Sample $r_k \leftarrow_{\$} [A], r_x \leftarrow_{\$} \{0, \dots, p - 1\}$
- 314 • Compute $t_1 := \mathcal{H}(\ell)^{r_k}, t_2 := \text{pk}_{0,\ell}^{r_k} \cdot f^{r_x}$
- 315 • Compute $ch := \mathbb{H}(\ell, \text{pk}_{i,\ell}, \text{ct}_{i,\ell}, t_1, t_2, \text{pk}_{0,\ell})$
- 316 • Compute $s_k := r_k + ch \cdot k_i, s_x := r_x + ch \cdot x_{i,\ell} \bmod p$
- 317 • Set $\text{proof}_i := (s_k, s_x, ch)$
- 318 • $C_i \rightarrow \text{Server}: \text{ct}_{i,\ell}$
- 319 • $C_i \rightarrow \text{Collector}: \text{pk}_{i,\ell}, \text{ct}_{i,\ell}, \text{proof}_i$

Proof Verification: Collector does: For client i in $\mathcal{O}\mathcal{L}_\ell$:

- Receive: $(\text{pk}_{i,\ell}, \text{ct}_{i,\ell}, \text{proof}_i) = (s_k, s_x, ch)$
- Compute $t'_2 := \text{pk}_{0,\ell}^{s_k} \cdot f^{s_x} \cdot \text{ct}_{i,\ell}^{-ch}$
- Compute $t'_1 := (\mathcal{H}(\ell))^{s_k} \cdot \mathcal{H}(\ell)^{-ch}$
- Compute $ch' := \mathbb{H}(\ell, \text{pk}_{i,\ell}, \text{ct}_{i,\ell}, t'_1, t'_2, \text{pk}_{0,\ell})$
- **if** $ch \neq ch'$ **then**
 $\mathcal{O}\mathcal{L}_\ell := \mathcal{O}\mathcal{L}_\ell \setminus \{i\}$
 Add $(\text{proof}_i, \text{ct}_{i,\ell}, \text{pk}_{i,\ell})$ to list \mathcal{M}
- Compute $\text{AUX}_\ell := \prod_{i \in \mathcal{O}\mathcal{L}_\ell} \text{pk}_{i,\ell}$
- **Collector** \rightarrow **Server:**
 $\text{AUX}_\ell, \{\text{ct}_{i,\ell}\}_{i \in \mathcal{O}\mathcal{L}_\ell}, \mathcal{M}$

322 5.2 Byzantine-Robust Stochastic Aggregation and Heterogeneous Datasets

323 **Heterogeneity in Data Distribution.** Data-centric methods [97, 72, 54] aim to align local and global
324 distributions while preserving privacy, using techniques like sharing raw, synthesized, or augmented data.
325 However, these approaches may compromise local data privacy [79]. Privacy-preserving machine learning can
326 be achieved through secret sharing schemes such as homomorphic encryption (HE) [42, 46] and multiparty
327 computation (MPC) [70]. However, HE is computationally expensive, and MPC faces scalability issues. Recent
328 frameworks [85] utilize Lagrange coding and polynomial approximations to address these challenges in federated
329 learning settings. RSA [64] is a class of stochastic sub-gradient methods for distributed learning robust to
330 Byzantine workers. It mitigates the effects of incorrect messages due to malicious behavior, communication
331 failures, or uneven data distribution by incorporating a regularization term in the objective function. At each
332 iteration k , clients compute parameter updates based on local data, prior local models, and global parameters.
333 The client and server updates are:

$$334 \begin{aligned} \text{Client: } x_i^{k+1} &= x_i^k - \eta^k \left(\nabla F(x_i^k, \xi_i^k) + \lambda \text{sign}(x_i^k - w^k) \right) \\ \text{Server: } w^{k+1} &= w^k - \eta^k \left(\nabla f_0(w^k) + \lambda \sum_{i \in [n]} \text{sign}(w^k - x_i^k) \right) \end{aligned}$$

335 where η is the learning rate, ξ is a local dataset sample, $F(\cdot, \cdot)$ is the loss function, $f_{\ell_2}(\cdot)$ is the robust
336 regularization term, λ weights the robustness term, sign is element-wise, and $[n]$ is the client set.

337 **Secure Aggregation with RSA.** As pointed out by Franzese *et al.* [41], the only information needed
338 by the server to aggregate is $\text{sign}(w^k - x_i^k)$. In other words, the clients simply need to supply the server
339 with a vector with elements in $\{-1, 1\}$. Furthermore, representing -1 as a 0 yields the following property:
340 $2 \cdot \sum_{i=1}^n v_i - n = \sum_{i=1}^n u_i$ where $u_i \in \{-1, 1\}$ and $v_i = 0$ iff $u_i = -1$. In summary, the server has to perform
341 aggregation over binary vectors. PICASO can be used to perform this securely, with only the client having to
342 prove that the masked input is either 0 or 1. Such a proof is efficient and we describe below. We present the
343 additional steps to be performed by the clients and the server in Construction 4, where the client proving that it
344 has encrypted either a value of 0 or a value of 1. This is an adaptation of Groth and Kohlweiss [48] to the CL
345 Framework. We omit the proof due to space constraints but it follows earlier results from Braun *et al.* [17].

346 **Construction 4** (Secure, Byzantine-Robust Secure Aggregation with PICASO). We assume that there is an
347 hash function $\mathbb{H} : \{0, 1\}^* \rightarrow \mathcal{C}$. Here, $A := 2^{40} \cdot |\mathcal{D}_H| \cdot C$ and $[A] := \{0, \dots, A - 1\}$.

348 **Proof Generation:** Each online client C_i is encrypting $x_{i,\ell} \in \{0, 1\}$ where $\text{ct}_{i,\ell} := \text{pk}_{0,\ell}^{k_i} \cdot f^{x_{i,\ell}}$ **Proof Verification:** Server does: For client i in \mathcal{OL}_ℓ :

- | | |
|--|--|
| <p>350 • Sample</p> <p>351 $r_k, r'_k \leftarrow [A], r_x \leftarrow \{0, \dots, p - 1\}$</p> <p>352 • Compute $t_1 := \text{pk}_{0,\ell}^{r_k} \cdot f^{r_x}, t_2 := \text{pk}_{0,\ell}^{r'_k} \cdot f^{r_x \cdot x_{i,\ell}}$</p> <p>353 • Compute $ch := \mathbb{H}(\ell, \text{pk}_{i,\ell}, \text{ct}_{i,\ell}, t_1, t_2, \text{pk}_{0,\ell})$</p> <p>354 • Compute $s_x := r_x + ch \cdot x_{i,\ell} \bmod p$</p> <p>355 • Compute $s_k := r_k + ch \cdot k_i, s'_k := r'_k + (ch - s_x) \cdot k_i$</p> <p>356 • Set $\text{proof}_i := (s_k, s'_k, s_x, ch)$</p> <p>357 • $C_i \rightarrow \text{Server: } \text{ct}_{i,\ell}, \text{proof}_i$</p> <p>358 • $C_i \rightarrow \text{Collector: } \text{pk}_{i,\ell}$</p> | <p>359 • Receive: $(\text{pk}_{i,\ell}, \text{ct}_{i,\ell}, \text{proof}_i) = (s_k, s'_k, s_x, ch)$</p> <p>360 • Compute $t'_1 := \text{ct}_{i,\ell}^{-ch} \cdot f^{s_x} \cdot \text{pk}_{0,\ell}^{s_k}$</p> <p>361 • Compute $t'_2 := \text{ct}_{i,\ell}^{s_x - ch} \cdot \text{pk}_{0,\ell}^{s'_k}$</p> <p>362 • Compute $ch' := \mathbb{H}(\ell, \text{pk}_{i,\ell}, \text{ct}_{i,\ell}, t'_1, t'_2, \text{pk}_{0,\ell})$</p> <p>363 • if $ch \neq ch'$ then $\mathcal{OL}_\ell := \mathcal{OL}_\ell \setminus \{i\}$</p> |
|--|--|

362 6 Conclusion

363 We present PICASO which is a secure aggregation protocol where client only has to synchronize once. This is an
364 improvement over existing secure aggregation protocols. We also demonstrate that our protocol has a competitive
365 performance over these protocols, and outperforms several of them. Finally, we show that PICASO preserves
366 accuracy, while guaranteeing privacy. Our encryption time increasing proportionally with the length of vector.
367 While this is expected, our use of group exponentiations makes the process slower. A possible direction for future
368 research is to apply the SASH framework [66] with PICASO, which reduces number of group exponentiations.
369 Despite limitations we believe PICASO significantly solves the scalability problem, while ensuring privacy. Its
370 possible extensions improve on several state-of-the-art protocols.

References

- 371
372 [1] Apple machine learning research. URL [https://machinelearning.apple.com/research/](https://machinelearning.apple.com/research/learning-with-privacy-at-scale)
373 [learning-with-privacy-at-scale](https://machinelearning.apple.com/research/learning-with-privacy-at-scale).
- 374 [2] Discrete logarithm. URL en.wikipedia.org/wiki/Discrete_logarithm. Wikipedia.
- 375 [3] D. Abram, I. Damgård, C. Orlandi, and P. Scholl. An algebraic framework for silent preprocessing with
376 trustless setup and active security. In Y. Dodis and T. Shrimpton, editors, *Advances in Cryptology –*
377 *CRYPTO 2022, Part IV*, volume 13510 of *Lecture Notes in Computer Science*, pages 421–452. Springer,
378 Heidelberg, Aug. 2022. doi: 10.1007/978-3-031-15985-5_15.
- 379 [4] S. Addanki, K. Garbe, E. Jaffe, R. Ostrovsky, and A. Polychroniadou. Prio+: Privacy preserving aggregate
380 statistics via boolean shares. In *Security and Cryptography for Networks: 13th International*
381 *Conference, SCN 2022, Amalfi (SA), Italy, September 12–14, 2022, Proceedings*, page 516–539, Berlin,
382 Heidelberg, 2022. Springer-Verlag. ISBN 978-3-031-14790-6. doi: 10.1007/978-3-031-14791-3_23. URL
383 https://doi.org/10.1007/978-3-031-14791-3_23.
- 384 [5] A. Arun, C. Ganesh, S. Lokam, T. Mopuri, and S. Sridhar. Dew: Transparent constant-sized zkSNARKs.
385 Cryptology ePrint Archive, Report 2022/419, 2022. <https://eprint.iacr.org/2022/419>.
- 386 [6] T. Attema, I. Cascudo, R. Cramer, I. B. Damgård, and D. Escudero. Vector commitments over rings and
387 compressed σ -protocols. Cryptology ePrint Archive, Report 2022/181, 2022. <https://eprint.iacr.org/2022/181>.
- 388
- 389 [7] J. Bell, A. Gascón, T. Lepoint, B. Li, S. Meiklejohn, M. Raykova, and C. Yun. ACORN: Input validation
390 for secure aggregation. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 4805–4822,
391 Anaheim, CA, Aug. 2023. USENIX Association. ISBN 978-1-939133-37-3. URL <https://www.usenix.org/conference/usenixsecurity23/presentation/bell>.
- 392
- 393 [8] J. H. Bell, K. A. Bonawitz, A. Gascón, T. Lepoint, and M. Raykova. Secure single-server aggregation
394 with (poly) logarithmic overhead. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and*
395 *Communications Security*, pages 1253–1269, 2020.
- 396 [9] F. Benhamouda, M. Joye, and B. Libert. A new framework for privacy-preserving aggregation of time-
397 series data. *ACM Trans. Inf. Syst. Secur.*, 18(3), mar 2016. ISSN 1094-9224. doi: 10.1145/2873069. URL
398 <https://doi.org/10.1145/2873069>.
- 399 [10] W. Beullens, T. Kleinjung, and F. Vercauteren. CSI-FiSh: Efficient isogeny based signatures through class
400 group computations. In S. D. Galbraith and S. Moriai, editors, *Advances in Cryptology – ASIACRYPT 2019,*
401 *Part I*, volume 11921 of *Lecture Notes in Computer Science*, pages 227–247. Springer, Heidelberg, Dec.
402 2019. doi: 10.1007/978-3-030-34578-5_9.
- 403 [11] J.-F. Biasse. Improvements in the computation of ideal class groups of imaginary quadratic number fields,
404 2010. ISSN 1930-5346. URL [/article/id/0151cba7-9512-448b-83ae-a272f0b836ce](https://arxiv.org/abs/1011.1811).
- 405 [12] J.-F. Biasse, M. J. Jacobson, and A. K. Silverster. Security estimates for quadratic field based cryptosystems.
406 In R. Steinfeld and P. Hawkes, editors, *ACISP 10: 15th Australasian Conference on Information Security*
407 *and Privacy*, volume 6168 of *Lecture Notes in Computer Science*, pages 233–247. Springer, Heidelberg,
408 July 2010.
- 409 [13] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and
410 K. Seth. Practical secure aggregation for privacy-preserving machine learning. In B. M. Thuraisingham,
411 D. Evans, T. Malkin, and D. Xu, editors, *Proceedings of the 2017 ACM SIGSAC Conference on Computer*
412 *and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*, pages
413 1175–1191. ACM, 2017. doi: 10.1145/3133956.3133982. URL [https://doi.org/10.1145/3133956.](https://doi.org/10.1145/3133956.3133982)
414 [3133982](https://doi.org/10.1145/3133956.3133982).
- 415 [14] D. Boneh. The decision diffie-hellman problem. In J. P. Buhler, editor, *Algorithmic Number Theory*, pages
416 48–63, Berlin, Heidelberg, 1998. Springer Berlin Heidelberg. ISBN 978-3-540-69113-6.
- 417 [15] D. Boneh, B. Bünz, and B. Fisch. Batching techniques for accumulators with applications to IOPs and
418 stateless blockchains. In A. Boldyreva and D. Micciancio, editors, *Advances in Cryptology – CRYPTO 2019,*
419 *Part I*, volume 11692 of *Lecture Notes in Computer Science*, pages 561–586. Springer, Heidelberg, Aug.
420 2019. doi: 10.1007/978-3-030-26948-7_20.
- 421 [16] C. Bouvier, G. Castagnos, L. Imbert, and F. Laguillaumie. I want to ride my BICYCL : BICYCL
422 implements cryptography in class groups. *J. Cryptol.*, 36(3):17, 2023. doi: 10.1007/s00145-023-09459-1.
423 URL <https://doi.org/10.1007/s00145-023-09459-1>.

- 424 [17] L. Braun, I. Damgård, and C. Orlandi. Secure multiparty computation from threshold encryption based on
425 class groups. In H. Handschuh and A. Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023,*
426 *Part I*, volume 14081 of *Lecture Notes in Computer Science*, pages 613–645. Springer, Heidelberg, Aug.
427 2023. doi: 10.1007/978-3-031-38557-5_20.
- 428 [18] J. Buchmann and H. C. Williams. A key-exchange system based on imaginary quadratic fields. *Journal of*
429 *Cryptology*, 1(2):107–118, June 1988. doi: 10.1007/BF02351719.
- 430 [19] B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell. Bulletproofs: Short proofs for
431 confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy*, pages 315–334.
432 IEEE Computer Society Press, May 2018. doi: 10.1109/SP.2018.00020.
- 433 [20] B. Bünz, B. Fisch, and A. Szepieniec. Transparent SNARKs from DARK compilers. In A. Canteaut and
434 Y. Ishai, editors, *Advances in Cryptology – EUROCRYPT 2020, Part I*, volume 12105 of *Lecture Notes in*
435 *Computer Science*, pages 677–706. Springer, Heidelberg, May 2020. doi: 10.1007/978-3-030-45721-1_24.
- 436 [21] D. Byrd, M. Hybinette, and T. H. Balch. Abides: Towards high-fidelity multi-agent market simulation.
437 In *Proceedings of the 2020 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation,*
438 SIGSIM-PADS ’20, page 11–22, New York, NY, USA, 2020. Association for Computing Machinery.
439 ISBN 9781450375924. doi: 10.1145/3384441.3395986. URL [https://doi.org/10.1145/3384441.](https://doi.org/10.1145/3384441.3395986)
440 3395986.
- 441 [22] I. Cascudo and B. David. ALBATROSS: Publicly Attestable BATched Randomness based On Secret
442 Sharing. In S. Moriai and H. Wang, editors, *Advances in Cryptology – ASIACRYPT 2020, Part III*, volume
443 12493 of *Lecture Notes in Computer Science*, pages 311–341. Springer, Heidelberg, Dec. 2020. doi:
444 10.1007/978-3-030-64840-4_11.
- 445 [23] G. Castagnos and F. Laguillaumie. On the security of cryptosystems with quadratic decryption: The
446 nicest cryptanalysis. In A. Joux, editor, *Advances in Cryptology – EUROCRYPT 2009*, volume 5479
447 of *Lecture Notes in Computer Science*, pages 260–277. Springer, Heidelberg, Apr. 2009. doi: 10.1007/
448 978-3-642-01001-9_15.
- 449 [24] G. Castagnos and F. Laguillaumie. Linearly homomorphic encryption from DDH. In K. Nyberg, editor,
450 *Topics in Cryptology – CT-RSA 2015*, volume 9048 of *Lecture Notes in Computer Science*, pages 487–505.
451 Springer, Heidelberg, Apr. 2015. doi: 10.1007/978-3-319-16715-2_26.
- 452 [25] G. Castagnos, A. Joux, F. Laguillaumie, and P. Q. Nguyen. Factoring pq^2 with quadratic forms: Nice
453 cryptanalyses. In M. Matsui, editor, *Advances in Cryptology – ASIACRYPT 2009*, volume 5912 of
454 *Lecture Notes in Computer Science*, pages 469–486. Springer, Heidelberg, Dec. 2009. doi: 10.1007/
455 978-3-642-10366-7_28.
- 456 [26] G. Castagnos, L. Imbert, and F. Laguillaumie. Encryption switching protocols revisited: Switching
457 modulo p . In J. Katz and H. Shacham, editors, *Advances in Cryptology – CRYPTO 2017, Part I*, volume
458 10401 of *Lecture Notes in Computer Science*, pages 255–287. Springer, Heidelberg, Aug. 2017. doi:
459 10.1007/978-3-319-63688-7_9.
- 460 [27] G. Castagnos, F. Laguillaumie, and I. Tucker. Practical fully secure unrestricted inner product functional
461 encryption modulo p . In T. Peyrin and S. Galbraith, editors, *Advances in Cryptology – ASIACRYPT 2018,*
462 *Part II*, volume 11273 of *Lecture Notes in Computer Science*, pages 733–764. Springer, Heidelberg, Dec.
463 2018. doi: 10.1007/978-3-030-03329-3_25.
- 464 [28] G. Castagnos, D. Catalano, F. Laguillaumie, F. Savasta, and I. Tucker. Two-party ECDSA from hash proof
465 systems and efficient instantiations. In A. Boldyreva and D. Micciancio, editors, *Advances in Cryptology –*
466 *CRYPTO 2019, Part III*, volume 11694 of *Lecture Notes in Computer Science*, pages 191–221. Springer,
467 Heidelberg, Aug. 2019. doi: 10.1007/978-3-030-26954-8_7.
- 468 [29] G. Castagnos, D. Catalano, F. Laguillaumie, F. Savasta, and I. Tucker. Bandwidth-efficient threshold
469 EC-DNA. In A. Kiayias, M. Kohlweiss, P. Wallden, and V. Zikas, editors, *PKC 2020: 23rd International*
470 *Conference on Theory and Practice of Public Key Cryptography, Part II*, volume 12111 of *Lecture Notes in*
471 *Computer Science*, pages 266–296. Springer, Heidelberg, May 2020. doi: 10.1007/978-3-030-45388-6_10.
- 472 [30] Z. Chai, Y. Chen, L. Zhao, Y. Cheng, and H. Rangwala. Fedat: A communication-efficient federated
473 learning method with asynchronous tiers under non-iid data. *arXiv preprint arXiv:2010.05958*, 2020.
- 474 [31] P. Chaidos and G. Couteau. Efficient designated-verifier non-interactive zero-knowledge proofs of knowl-
475 edge. In J. B. Nielsen and V. Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2018, Part III*,
476 volume 10822 of *Lecture Notes in Computer Science*, pages 193–221. Springer, Heidelberg, Apr. / May
477 2018. doi: 10.1007/978-3-319-78372-7_7.

- 478 [32] Y. Chen, Y. Ning, M. Slawski, and H. Rangwala. Asynchronous online federated learning for edge devices
479 with non-iid data. In *2020 IEEE International Conference on Big Data (Big Data)*, pages 15–24. IEEE,
480 2020.
- 481 [33] H. Corrigan-Gibbs and D. Boneh. Prio: Private, robust, and scalable computation of aggregate statis-
482 tics. In A. Akella and J. Howell, editors, *14th USENIX Symposium on Networked Systems Design
483 and Implementation, NSDI 2017, Boston, MA, USA, March 27-29, 2017*, pages 259–282. USENIX
484 Association, 2017. URL [https://www.usenix.org/conference/nsdi17/technical-sessions/
485 presentation/corrigan-gibbs](https://www.usenix.org/conference/nsdi17/technical-sessions/presentation/corrigan-gibbs).
- 486 [34] G. Couteau, M. Kloof, H. Lin, and M. Reichle. Efficient range proofs with transparent setup from
487 bounded integer commitments. In A. Canteaut and F.-X. Standaert, editors, *Advances in Cryptology
488 – EUROCRYPT 2021, Part III*, volume 12698 of *Lecture Notes in Computer Science*, pages 247–277.
489 Springer, Heidelberg, Oct. 2021. doi: 10.1007/978-3-030-77883-5_9.
- 490 [35] G. Couteau, D. Goudarzi, M. Kloof, and M. Reichle. Sharp: Short relaxed range proofs. In H. Yin,
491 A. Stavrou, C. Cremers, and E. Shi, editors, *ACM CCS 2022: 29th Conference on Computer and Commu-
492 nications Security*, pages 609–622. ACM Press, Nov. 2022. doi: 10.1145/3548606.3560628.
- 493 [36] Y. Deng, S. Ma, X. Zhang, H. Wang, X. Song, and X. Xie. Promise Σ -protocol: How to construct efficient
494 threshold ECDSA from encryptions based on class groups. In M. Tibouchi and H. Wang, editors, *Advances
495 in Cryptology – ASIACRYPT 2021, Part IV*, volume 13093 of *Lecture Notes in Computer Science*, pages
496 557–586. Springer, Heidelberg, Dec. 2021. doi: 10.1007/978-3-030-92068-5_19.
- 497 [37] Drand. Drand/drand: a distributed randomness beacon daemon - go implementation. URL [https://
498 github.com/drand/drand](https://github.com/drand/drand).
- 499 [38] C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Found. Trends Theor.
500 Comput. Sci.*, 9(3–4):211–407, aug 2014. ISSN 1551-305X. doi: 10.1561/04000000042. URL [https://
501 doi.org/10.1561/04000000042](https://doi.org/10.1561/04000000042).
- 502 [39] T. Elgamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE
503 Transactions on Information Theory*, 31(4):469–472, 1985. doi: 10.1109/TIT.1985.1057074.
- 504 [40] Ú. Erlingsson, V. Pihur, and A. Korolova. RAPPOR: Randomized aggregatable privacy-preserving ordinal
505 response. In G.-J. Ahn, M. Yung, and N. Li, editors, *ACM CCS 2014: 21st Conference on Computer and
506 Communications Security*, pages 1054–1067. ACM Press, Nov. 2014. doi: 10.1145/2660267.2660348.
- 507 [41] N. Franzese, A. Dziedzic, C. A. Choquette-Choo, M. R. Thomas, M. A. Kaleem, S. Rabanser, C. Fang,
508 S. Jha, N. Papernot, and X. Wang. Robust and actively secure serverless collaborative learning. In
509 *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL [https://openreview.
510 net/forum?id=SouroWC5Un](https://openreview.net/forum?id=SouroWC5Un).
- 511 [42] C. Gentry. Fully homomorphic encryption using ideal lattices. *Proceedings of the forty-first annual ACM
512 symposium on Theory of computing*, pages 169–178, 2009.
- 513 [43] C. Gentry, S. Halevi, and V. Lyubashevsky. Practical non-interactive publicly verifiable secret sharing
514 with thousands of parties. In O. Dunkelman and S. Dziembowski, editors, *Advances in Cryptology –
515 EUROCRYPT 2022, Part I*, volume 13275 of *Lecture Notes in Computer Science*, pages 458–487. Springer,
516 Heidelberg, May / June 2022. doi: 10.1007/978-3-031-06944-4_16.
- 517 [44] R. C. Geyer, T. Klein, and M. Nabi. Differentially private federated learning: A client level perspective.
518 *CoRR*, abs/1712.07557, 2017. URL <http://arxiv.org/abs/1712.07557>.
- 519 [45] Y. Gilad, R. Hemo, S. Micali, G. Vlachos, and N. Zeldovich. Algorand: Scaling byzantine agreements
520 for cryptocurrencies. In *Proceedings of the 26th Symposium on Operating Systems Principles, SOSP '17*,
521 page 51–68, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450350853.
522 doi: 10.1145/3132747.3132757. URL <https://doi.org/10.1145/3132747.3132757>.
- 523 [46] R. Gilad-Bachrach, N. Dowlin, K. Laine, K. Lauter, M. Naehrig, and J. Wernsing. Cryptonets: Applying
524 neural networks to encrypted data with high throughput and accuracy. In *International conference on
525 machine learning*, pages 201–210. PMLR, 2016.
- 526 [47] N. Glaeser, M. Maffei, G. Malavolta, P. Moreno-Sanchez, E. Tairi, and S. A. Thyagarajan. Foundations of
527 coin mixing services. *Cryptology ePrint Archive, Report 2022/942*, 2022. [https://eprint.iacr.org/
528 2022/942](https://eprint.iacr.org/2022/942).

- 529 [48] J. Groth and M. Kohlweiss. One-out-of-many proofs: Or how to leak a secret and spend a coin. In
530 E. Oswald and M. Fischlin, editors, *Advances in Cryptology – EUROCRYPT 2015, Part II*, volume
531 9057 of *Lecture Notes in Computer Science*, pages 253–280. Springer, Heidelberg, Apr. 2015. doi:
532 10.1007/978-3-662-46803-6_9.
- 533 [49] Y. Guo, A. Polychroniadou, E. Shi, D. Byrd, and T. Balch. Microsecagg: Streamlined single-server
534 secure aggregation. *Proceedings on Privacy Enhancing Technologies*, 2024(4):77–101, 2024. doi:
535 10.56553/popets-2024-0077. URL <https://doi.org/10.56553/popets-2024-0077>.
- 536 [50] J. L. Hafner and K. S. McCurley. A rigorous subexponential algorithm for computation of class groups.
537 *Journal of the American mathematical society*, 2(4):837–850, 1989.
- 538 [51] B. Hitaj, G. Ateniese, and F. Pérez-Cruz. Deep models under the GAN: Information leakage from
539 collaborative deep learning. In B. M. Thuraisingham, D. Evans, T. Malkin, and D. Xu, editors, *ACM*
540 *CCS 2017: 24th Conference on Computer and Communications Security*, pages 603–618. ACM Press,
541 Oct. / Nov. 2017. doi: 10.1145/3133956.3134012.
- 542 [52] D. Hühnlein, M. J. Jacobson Jr., S. Paulus, and T. Takagi. A cryptosystem based on non-maximal imaginary
543 quadratic orders with fast decryption. In K. Nyberg, editor, *Advances in Cryptology – EUROCRYPT’98*,
544 volume 1403 of *Lecture Notes in Computer Science*, pages 294–307. Springer, Heidelberg, May / June
545 1998. doi: 10.1007/BFb0054134.
- 546 [53] M. J. Jacobson. Computing discrete logarithms in quadratic orders. *J. Cryptol.*, 13(4):473–492, jan 2000.
547 ISSN 0933-2790. doi: 10.1007/s001450010013. URL <https://doi.org/10.1007/s001450010013>.
- 548 [54] E. Jeong, S. Oh, H. Kim, J. Park, M. Bennis, and S.-Y. Kang. Communication-efficient on-device
549 machine learning: Federated distillation and augmentation under non-iid private data. *arXiv preprint*
550 *arXiv:1811.11479*, 2018.
- 551 [55] M. Joye and B. Libert. A scalable scheme for privacy-preserving aggregation of time-series data. In A.-R.
552 Sadeghi, editor, *FC 2013: 17th International Conference on Financial Cryptography and Data Security*,
553 volume 7859 of *Lecture Notes in Computer Science*, pages 111–125. Springer, Heidelberg, Apr. 2013. doi:
554 10.1007/978-3-642-39884-1_10.
- 555 [56] S. Kadhe, N. Rajaraman, O. O. Koyluoglu, and K. Ramchandran. Fastsecagg: Scalable secure aggregation
556 for privacy-preserving federated learning. *CoRR*, abs/2009.11248, 2020. URL <https://arxiv.org/abs/2009.11248>.
- 558 [57] P. Kairouz, S. Oh, and P. Viswanath. Extremal mechanisms for local differential privacy. In Z. Ghahra-
559 mani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, editors, *Advances in Neural Information*
560 *Processing Systems*, volume 27. Curran Associates, Inc., 2014. URL [https://proceedings.neurips.](https://proceedings.neurips.cc/paper_files/paper/2014/file/86df7dcfd896fcfa2674f757a2463eba-Paper.pdf)
561 [cc/paper_files/paper/2014/file/86df7dcfd896fcfa2674f757a2463eba-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2014/file/86df7dcfd896fcfa2674f757a2463eba-Paper.pdf).
- 562 [58] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles,
563 G. Cormode, R. Cummings, et al. Advances and open problems in federated learning. *Foundations and*
564 *Trends® in Machine Learning*, 14(1-2):1–210, 2021.
- 565 [59] A. Kate, E. V. Mangipudi, P. Mukherjee, H. Saleem, and S. A. K. Thyagarajan. Non-interactive vss
566 using class groups and application to dkg. *Cryptology ePrint Archive*, Paper 2023/451, 2023. URL
567 <https://eprint.iacr.org/2023/451>. <https://eprint.iacr.org/2023/451>.
- 568 [60] T. Kleinjung. Quadratic sieving. *Math. Comput.*, 85(300):1861–1873, 2016. doi: 10.1090/mcom/3058.
569 URL <https://doi.org/10.1090/mcom/3058>.
- 570 [61] R. W. F. Lai and G. Malavolta. Subvector commitments with application to succinct arguments. In
571 A. Boldyreva and D. Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part I*, volume
572 11692 of *Lecture Notes in Computer Science*, pages 530–560. Springer, Heidelberg, Aug. 2019. doi:
573 10.1007/978-3-030-26948-7_19.
- 574 [62] I. Leontiadis, K. Elkhyaoui, and R. Molva. Private and dynamic time-series data aggregation with trust
575 relaxation. In D. Gritzalis, A. Kiayias, and I. G. Askoxylakis, editors, *CANS 14: 13th International*
576 *Conference on Cryptology and Network Security*, volume 8813 of *Lecture Notes in Computer Science*,
577 pages 305–320. Springer, Heidelberg, Oct. 2014. doi: 10.1007/978-3-319-12280-9_20.
- 578 [63] H. Li, H. Lin, A. Polychroniadou, and S. Tessaro. LERNA: Secure single-server aggregation via key-
579 homomorphic masking. In J. Guo and R. Steinfeld, editors, *Advances in Cryptology – ASIACRYPT 2023,*
580 *Part I*, volume 14438 of *Lecture Notes in Computer Science*, pages 302–334. Springer, Heidelberg, Dec.
581 2023. doi: 10.1007/978-981-99-8721-4_10.

- 582 [64] L. Li, W. Xu, T. Chen, G. B. Giannakis, and Q. Ling. Rsa: Byzantine-robust stochastic aggregation methods
583 for distributed learning from heterogeneous datasets. In *Proceedings of the Thirty-Third AAAI Conference*
584 *on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and*
585 *Ninth AAAI Symposium on Educational Advances in Artificial Intelligence*, AAAI'19/IAAI'19/EAAI'19.
586 AAAI Press, 2019. ISBN 978-1-57735-809-1. doi: 10.1609/aaai.v33i01.33011544. URL <https://doi.org/10.1609/aaai.v33i01.33011544>.
587
- 588 [65] H. Lipmaa. Secure accumulators from euclidean rings without trusted setup. In F. Bao, P. Samarati, and
589 J. Zhou, editors, *ACNS 12: 10th International Conference on Applied Cryptography and Network Security*,
590 volume 7341 of *Lecture Notes in Computer Science*, pages 224–240. Springer, Heidelberg, June 2012. doi:
591 10.1007/978-3-642-31284-7_14.
- 592 [66] Z. Liu, S. Chen, J. Ye, J. Fan, H. Li, and X. Li. SASH: efficient secure aggregation based on SHPRG for
593 federated learning. In J. Cussens and K. Zhang, editors, *Uncertainty in Artificial Intelligence, Proceedings*
594 *of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence, UAI 2022, 1-5 August 2022,*
595 *Eindhoven, The Netherlands*, volume 180 of *Proceedings of Machine Learning Research*, pages 1243–1252.
596 PMLR, 2022. URL <https://proceedings.mlr.press/v180/liu22c.html>.
- 597 [67] H. Lycklama, L. Burkhalter, A. Viand, N. Küchler, and A. Hithnawi. RoFL: Robustness of secure federated
598 learning. In *2023 IEEE Symposium on Security and Privacy*, pages 453–476. IEEE Computer Society
599 Press, May 2023. doi: 10.1109/SP46215.2023.10179400.
- 600 [68] Y. Ma, J. Woods, S. Angel, A. Polychroniadou, and T. Rabin. Flamingo: Multi-round single-server secure
601 aggregation with applications to private federated learning. In *2023 IEEE Symposium on Security and*
602 *Privacy*, pages 477–496. IEEE Computer Society Press, May 2023. doi: 10.1109/SP46215.2023.10179434.
- 603 [69] K. McCurley. Cryptographic key distribution and computation in class groups. *Proceedings of NATO*
604 *ASI Number Theory and applications*, pages 459–479, 1989. URL [https://cir.nii.ac.jp/crid/](https://cir.nii.ac.jp/crid/1570854174866004864)
605 [1570854174866004864](https://cir.nii.ac.jp/crid/1570854174866004864).
- 606 [70] P. Mohassel and Y. Zhang. Secureml: A system for scalable privacy-preserving machine learning. In *2017*
607 *IEEE symposium on security and privacy (SP)*, pages 19–38. IEEE, 2017.
- 608 [71] M. Nasr, R. Shokri, and A. Houmansadr. Comprehensive privacy analysis of deep learning: Passive and
609 active white-box inference attacks against centralized and federated learning. In *2019 IEEE Symposium on*
610 *Security and Privacy*, pages 739–753. IEEE Computer Society Press, May 2019. doi: 10.1109/SP.2019.
611 00065.
- 612 [72] D. C. Nguyen, M. Ding, P. N. Pathirana, A. Seneviratne, J. Li, and H. V. Poor. Federated learning for
613 covid-19 detection with generative adversarial networks in edge cloud computing. *IEEE Internet of Things*
614 *Journal*, 2021.
- 615 [73] J. Nguyen, K. Malik, H. Zhan, A. Yousefpour, M. Rabbat, M. Malek, and D. Huba. Federated learning with
616 buffered asynchronous aggregation. In G. Camps-Valls, F. J. R. Ruiz, and I. Valera, editors, *Proceedings of*
617 *The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of*
618 *Machine Learning Research*, pages 3581–3607. PMLR, 28–30 Mar 2022. URL [https://proceedings.](https://proceedings.mlr.press/v151/nguyen22b.html)
619 [mlr.press/v151/nguyen22b.html](https://proceedings.mlr.press/v151/nguyen22b.html).
- 620 [74] T. P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In J. Feigenbaum,
621 editor, *Advances in Cryptology – CRYPTO'91*, volume 576 of *Lecture Notes in Computer Science*, pages
622 129–140. Springer, Heidelberg, Aug. 1992. doi: 10.1007/3-540-46766-1_9.
- 623 [75] L. T. Phong, Y. Aono, T. Hayashi, L. Wang, and S. Moriai. Privacy-preserving deep learning via additively
624 homomorphic encryption. *IEEE Transactions on Information Forensics and Security*, 13(5):1333–1345,
625 2018. doi: 10.1109/TIFS.2017.2787987.
- 626 [76] A. D. Pozzolo, O. Caelen, R. A. Johnson, and G. Bontempi. Calibrating probability with undersampling for
627 unbalanced classification. In *2015 IEEE Symposium Series on Computational Intelligence*, pages 159–166,
628 2015. doi: 10.1109/SSCI.2015.33.
- 629 [77] I. A. Seres, P. Burcsi, and P. Kutas. How (not) to hash into class groups of imaginary quadratic
630 fields? Cryptology ePrint Archive, Paper 2024/034, 2024. URL <https://eprint.iacr.org/2024/034>.
631 <https://eprint.iacr.org/2024/034>.
- 632 [78] A. Shamir. How to share a secret. *Commun. ACM*, 22(11):612–613, nov 1979. ISSN 0001-0782. doi:
633 10.1145/359168.359176. URL <https://doi.org/10.1145/359168.359176>.
- 634 [79] A. Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.

- 635 [80] J. Shao, Y. Sun, S. Li, and J. Zhang. Dres-fl: dropout-resilient secure federated learning for non-iid
636 clients via secret data sharing. In *Proceedings of the 36th International Conference on Neural Information*
637 *Processing Systems, NIPS '22*, Red Hook, NY, USA, 2024. Curran Associates Inc. ISBN 9781713871088.
- 638 [81] E. Shi, T.-H. H. Chan, E. G. Rieffel, R. Chow, and D. Song. Privacy-preserving aggregation of time-series
639 data. In *ISOC Network and Distributed System Security Symposium – NDSS 2011*. The Internet Society,
640 Feb. 2011.
- 641 [82] R. Shokri, M. Stronati, C. Song, and V. Shmatikov. Membership inference attacks against machine learning
642 models. In *2017 IEEE Symposium on Security and Privacy*, pages 3–18. IEEE Computer Society Press,
643 May 2017. doi: 10.1109/SP.2017.41.
- 644 [83] J. So, R. E. Ali, B. Güler, and A. S. Avestimehr. Secure aggregation for buffered asynchronous federated
645 learning. *CoRR*, abs/2110.02177, 2021. URL <https://arxiv.org/abs/2110.02177>.
- 646 [84] J. So, R. E. Ali, B. Güler, and A. S. Avestimehr. Secure aggregation for buffered asynchronous federated
647 learning, 2021.
- 648 [85] J. So, B. Guler, and A. S. Avestimehr. Codedprivateml: A fast and privacy-preserving framework for
649 distributed machine learning. *IEEE Journal on Selected Areas in Information Theory*, 2(1):441–451, 2021.
- 650 [86] J. So, B. Güler, and A. S. Avestimehr. Turbo-aggregate: Breaking the quadratic aggregation barrier in
651 secure federated learning. *IEEE Journal on Selected Areas in Information Theory*, 2(1):479–489, 2021.
652 doi: 10.1109/JSAIT.2021.3054610.
- 653 [87] J. So, C. He, C.-S. Yang, S. Li, Q. Yu, R. E. Ali, B. Guler, and S. Avestimehr. Lightsecagg: a lightweight
654 and versatile design for secure aggregation in federated learning, 2022.
- 655 [88] S. A. K. Thyagarajan, G. Castagnos, F. Laguillaumie, and G. Malavolta. Efficient CCA timed commitments
656 in class groups. In G. Vigna and E. Shi, editors, *ACM CCS 2021: 28th Conference on Computer and*
657 *Communications Security*, pages 2663–2684. ACM Press, Nov. 2021. doi: 10.1145/3460120.3484773.
- 658 [89] I. Tucker. *Functional encryption and distributed signatures based on projective hash functions, the*
659 *benefit of class groups*. Theses, Université de Lyon, Oct. 2020. URL [https://theses.hal.science/](https://theses.hal.science/tel-03021689)
660 [tel-03021689](https://theses.hal.science/tel-03021689).
- 661 [90] M. van Dijk, N. V. Nguyen, T. N. Nguyen, L. M. Nguyen, Q. Tran-Dinh, and P. H. Nguyen. Asynchronous
662 federated learning with reduced number of rounds and with differential privacy from less aggregated
663 gaussian noise. *arXiv preprint arXiv:2007.09208*, 2020.
- 664 [91] B. Wesolowski. Efficient verifiable delay functions. In Y. Ishai and V. Rijmen, editors, *Advances in*
665 *Cryptology – EUROCRYPT 2019, Part III*, volume 11478 of *Lecture Notes in Computer Science*, pages
666 379–407. Springer, Heidelberg, May 2019. doi: 10.1007/978-3-030-17659-4_13.
- 667 [92] B. Wesolowski. Efficient verifiable delay functions. *Journal of Cryptology*, 33(4):2113–2147, Oct. 2020.
668 doi: 10.1007/s00145-020-09364-x.
- 669 [93] C. Xie, S. Koyejo, and I. Gupta. Asynchronous federated optimization. *arXiv preprint arXiv:1903.03934*,
670 2019.
- 671 [94] C. Yang, J. So, C. He, S. Li, Q. Yu, and S. Avestimehr. Lightsecagg: Rethinking secure aggregation in
672 federated learning. *CoRR*, abs/2109.14236, 2021. URL <https://arxiv.org/abs/2109.14236>.
- 673 [95] T. H. Yuen, H. Cui, and X. Xie. Compact zero-knowledge proofs for threshold ECDSA with trustless
674 setup. In J. Garay, editor, *PKC 2021: 24th International Conference on Theory and Practice of Public*
675 *Key Cryptography, Part I*, volume 12710 of *Lecture Notes in Computer Science*, pages 481–511. Springer,
676 Heidelberg, May 2021. doi: 10.1007/978-3-030-75245-3_18.
- 677 [96] C. Zhang, S. Li, J. Xia, W. Wang, F. Yan, and Y. Liu. BatchCrypt: Efficient homomorphic encryption for
678 Cross-Silo federated learning. In *2020 USENIX Annual Technical Conference (USENIX ATC 20)*, pages
679 493–506. USENIX Association, July 2020. ISBN 978-1-939133-14-4. URL [https://www.usenix.](https://www.usenix.org/conference/atc20/presentation/zhang-chengliang)
680 [org/conference/atc20/presentation/zhang-chengliang](https://www.usenix.org/conference/atc20/presentation/zhang-chengliang).
- 681 [97] L. Zhang, B. Shen, A. Barnawi, Y. Tang, Z. Luo, W. Wang, W. Zhang, and Z. Han. Feddpagan: federated
682 differentially private generative adversarial networks framework for the detection of covid-19 pneumonia.
683 *Information Systems Frontiers*, 23(6):1403–1415, 2021.
- 684 [98] Y. Zhao and H. Sun. Information theoretic secure aggregation with user dropouts. In *2021 IEEE*
685 *International Symposium on Information Theory (ISIT)*, pages 1124–1129, 2021. doi: 10.1109/ISIT45174.
686 2021.9517953.

Table 3: Comparison of the RSA Modulus Size and the Class Group Size for various Security Levels. All the sizes are in bits. The RSA Modulus forms the basis of the DCR-based construction.

Security Level	RSA Size Modulus	Class Group Size
112	2048	1348
128	3072	1872
192	7680	3598
256	15360	5971

687 A Deferred Discussion on Class Groups and Cryptography

688 A.1 Class Group and Cryptography

689 Class Group-based Cryptography has its roots in the late 1980s, with the pioneering work of [18] and [69]. They
690 proposed that the class group of ideals of maximal orders of imaginary quadratic fields might offer better security
691 compared to the multiplicative group of finite fields. One of the key features of class group cryptography is its
692 suitability for protocols involving multiple parties while requiring only a one-time transparent (or public-coin)
693 setup without extensive interaction. This has enabled the construction of verifiable random functions without a
694 trusted setup by [91, 92], accumulators by [65, 15], encryption switching protocols by [26], designated verifier
695 non-interactive zero knowledge proofs of knowledge by [31], SNARKs [61, 20, 5], homomorphic secret sharing,
696 and pseudorandom correlation functions for oblivious transfer by [3], range proofs by [34, 35], and vector
697 commitments by [6]. Subsequent advancements in the computation of the structure of class groups of quadratic
698 imaginary number fields were made by [50], and further improvements were presented in the works of [53],[11],
699 and [60]. However, the computational costs of these algorithms increase with the size of Δ , with the largest
700 current computation involving 512 bits for the size of Δ by [10]. It is worth noting that the subexponential
701 complexity of computing class groups is asymptotically slower than integer factorization. Indeed, in the work of
702 [12], it was conjectured that one needs the size of the discriminant of class group Δ_k to be 1872 bits to achieve
703 128 bits of security. Meanwhile, one needs an RSA modulus of 3072 bits to achieve 128-bit security. Some
704 additional parameters for our class groups are: the size of \bar{s} is 914 bits and the size of the prime p is 128 bits.
705 We present the comparison in Table 3 that details the sizes for various levels of security. Looking ahead, the
706 presence of the RSA modulus and related cryptographic assumptions is the main reason why BatchCrypt [96]
707 and the works of [55] and [62] have inefficient parameter sizes, when compared to class group based protocol.

708 Efforts to improve the efficiency of cryptosystems based on class groups were undertaken by Hühlein *et*
709 *al.* [52]. However, their work was later attacked by [23] and [25]. Despite this setback, there has been a notable
710 resurgence in the use of class groups in various applications over the past decade. [24] made a significant
711 contribution by designing a cryptosystem based on a subgroup of a class group where discrete logarithms are
712 easy to compute. This framework has since become the foundation of several protocols, including projective
713 hash functions used to construct inner-product functional encryption [27], two-party and fully-threshold ECDSA
714 signatures [28, 29, 95, 36], coin-mixing [47], and secure timed commitments [88]. Additionally, the framework
715 has been employed to build secure multiparty computation from threshold encryption [17] and non-interactive
716 verifiable secret sharing [59, 22]. We rely on this framework for our protocol.

717 A.2 Expanded Discussion on CL Framework

718 Broadly, the framework is defined by two functions - CLGen, CLSolve with the former outputting a tuple of
719 public parameters. The elements of this framework are the following:

- 720 • *Input Parameters:* κ_c is the computational security parameter, κ_s is a statistical security parameter, a
721 prime p such that $p > 2^{\kappa_c}$, and uniform randomness ρ that is used by the CLGen algorithm and is
722 made public.
- 723 • *Groups:* $\widehat{\mathbb{G}}$ is a finite multiplicative abelian group, \mathbb{G} is a cyclic subgroup of $\widehat{\mathbb{G}}$, \mathbb{F} is a subgroup of \mathbb{G} ,
724 $\mathbb{H} = \{x^p, x \in \mathbb{G}\}$
- 725 • *Orders:* \mathbb{F} has order p , $\widehat{\mathbb{G}}$ has order $p \cdot \widehat{s}$, \mathbb{G} has order $p \cdot s$ such that s divides \widehat{s} and $\gcd(p, \widehat{s}) =$
726 $1, \gcd(p, s) = 1, \mathbb{H}$ has order s and therefore $\mathbb{G} = \mathbb{F} \times \mathbb{H}$.
- 727 • *Generators:* f is the generator of \mathbb{F} , g is the generator of \mathbb{G} , and h is the generator of \mathbb{H} with the
728 property that $g = f \cdot h$
- 729 • *Upper Bound:* Only an upper bound \bar{s} of \widehat{s} (and s) is provided.
- 730 • *Additional Properties:* Only encodings of $\widehat{\mathbb{G}}$ can be recognized as valid encodings and s, \widehat{s} are
731 unknown.

732 • *Distributions:* \mathcal{D} (resp. \mathcal{D}_p) be a distribution over the set of integers such that the distribution
 733 $\{g^x : x \leftarrow \mathcal{D}\}$ (resp. $\{g_p^x : x \leftarrow \mathcal{D}_p\}$) is at most distance $2^{-\kappa_s}$ from the uniform distribution over \mathbb{G}
 734 (resp. \mathbb{H}).

735 **Remark 3.** The motivations behind these additional distributions are as follows. One can efficiently recognize
 736 valid encodings of elements in $\widehat{\mathbb{G}}$ but not \mathbb{G} . Therefore, a malicious adversary \mathcal{A} can run our constructions by
 737 inputting elements belonging to $\widehat{\mathbb{G}}^p$ (rather than in \mathbb{H}). Unfortunately, this malicious behavior cannot be detected
 738 which allows \mathcal{A} to obtain information on the sampled exponents modulo $\bar{\omega}$ (the group exponent of $\widehat{\mathbb{G}}^p$). By
 739 requiring the statistical closeness of the induced distribution to uniform in the aforementioned groups allows
 740 flexibility in proofs. Note that the assumptions do not depend on the choice of these two distributions. Further,
 741 the order s of \mathbb{H} and group exponent $\bar{\omega}$ of $\widehat{\mathbb{G}}^p$ are unknown and the upper bound \bar{s} is used to instantiate the
 742 aforementioned distribution.

743 We also have the following lemma from Castagnos, Imbert, and Laguillaumie [26] which defines how to sample
 744 from a discrete Gaussian distribution.

745 **Lemma 1.** *Let \mathbb{G} be a cyclic group of order n , generated by g . Consider the random variable X sampled*
 746 *uniformly from \mathbb{G} ; as such it satisfies $\Pr[X = h] = \frac{1}{n}$ for all $h \in \mathbb{G}$. Now consider the random variable Y with*
 747 *values in \mathbb{G} as follows: draw y from the discrete Gaussian distribution $\mathcal{D}_{\mathbb{Z}, \sigma}$ with $\sigma \geq n \sqrt{\frac{\ln(2(1+1/\epsilon))}{\pi}}$ and set*
 748 *$Y := g^y$. Then, it holds that:*

$$\Delta(X, Y) \leq 2\epsilon$$

749 **Definition 2** (Class Group Framework). *The framework is defined by two algorithms (CLGen, CLSolve) such*
 750 *that:*

- 751 • $\text{pp} = (p, \kappa_c, \kappa_s, \bar{s}, f, h, \widehat{\mathbb{G}}, \mathbb{F}, \mathcal{D}_G, \mathcal{D}_H, \rho) \leftarrow \text{CLGen}(1^{\kappa_c}, 1^{\kappa_s}, p; \rho)$
- 752 • *The DL problem is easy in \mathbb{F} , i.e., there exists a deterministic polynomial algorithm CLSolve that*
 753 *solves the discrete logarithm problem in \mathbb{F} :*

$$\Pr \left[x = x' \mid \begin{array}{l} \text{pp} = \leftarrow \text{CLGen}(1^{\kappa_c}, 1^{\kappa_s}, p; \rho) \\ x \leftarrow \mathbb{Z}/p\mathbb{Z}, X = f^x; \\ x' \leftarrow \text{CLSolve}(\text{pp}, X) \end{array} \right] = 1$$

754 A.3 Why CL Framework?

755 Cryptographic protocols are often proven secure, under a hardness assumption. The hardness assumption
 756 guarantees that an adversary cannot break this assumption in polynomial time. A common setting for these
 757 cryptographic protocols is cyclic groups, usually of prime order q . Let G be such a group and because it is cyclic,
 758 there exists a generator g such that $G := \{g^0, \dots, g^{q-1}\}$. Now, consider the simpler setting where we want to
 759 securely sum up integer values. The immediate question is how to we encrypt these integer values when working
 760 over cyclic groups. The solution is to take your input x and map it to an element in the group G as g^x . This was
 761 the idea behind a famous encryption scheme known as ElGamal encryption proposed by Taher ElGamal [39].
 762 Unfortunately, the security of this encryption scheme also required that given a random group element $g' \in G$,
 763 one cannot efficiently recover x' such that $g' = g^{x'}$, i.e., that the discrete logarithm of g' cannot be efficiently
 764 computed. For our use case of the server recovering the sum of the integers, using such an encoding scheme
 765 would require the server to compute the discrete logarithm which is inefficient. However, while computing
 766 the discrete logarithm for a random element is inefficient, the problem becomes simpler if we assume that the
 767 maximum value of the discrete logarithm is bounded. In other words, rather than searching over the entire
 768 set $\{0, \dots, q-1\}$, the problem is simplified to searching over $\{0, \dots, B-1\}$, where $B \ll q$. This is the
 769 assumption made by Shi et al. [81] and Guo et al. [49].

770 The CL Framework also works over cyclic groups but its security does not rely on the inefficiency of computing
 771 the discrete logarithm. Instead, it relies on the inefficiency of computing the order of the group. This gives us the
 772 ability to encode input x in the subgroup F where discrete logarithm is efficient while using the element in \mathbb{H} to
 773 mask the inputs. This gives us the ciphertext $ct := h^k \cdot f^x$ where k is some random key. Note that if the order
 774 of \mathbb{H} (call it s) is known, then the security is lost. This is because I can compute $ct^s = h^{ks} \cdot f^{sx}$ with h^{ks} being
 775 the identity element. Then, sx can be recovered because the discrete logarithm is efficiently computable, with x
 776 recoverable after. Therefore, the benefit of the CL framework is the following: Encode the elements into a cyclic
 777 group, from which the sum can be recovered efficiently without restricting the input sizes.

778 B Experimental Details

779 **ABIDES Framework.** To simulate evaluate real network conditions, we run the ABIDES simulator [21].
 780 ABIDES supports a latency model which consists a base delay and plus a jitter that controls the percentage of
 781 messages that arrive within a given time. We set the base delay to the “global” setting in ABIDES’s default
 782 parameters (the range is 21 microseconds to 53 milliseconds), and use the default parameters for the jitter.

Table 4: Details about Datasets

	Fraud Detection	MNIST	CIFAR-10	CIFAR-100
No. of classes	2	10	10	100
No. of training samples	213k	60,000	50,000	50,000
No. of test samples	71k	10,000	10,000	
Feature Details	30 features	Image Size: 28×28	Image Size: 32×32	Image Size: 32×32
Quantization Multiplier	10^4	2^{16}	2^{16}	2^{16}
License	Open Database License	Creative Commons Attribution-Share Alike 3.0 License	Apache License 2.0	MIT License

783 **Dataset Details.** The license and feature details of our datasets can be found in Table 4.

784 **Architecture Details for MLP Classifier.** The default neural network defined by
 785 sklearn.neural_network.MLPClassifier has the following characteristics. Note that our goal is to show
 786 that the secure aggregation does not impact accuracy when compared to learning in the clear.

- 787 • Architecture:
 - 788 – Multi-layer perceptron (MLP) with one hidden layer
 - 789 – Hidden layer contains 100 neurons by default
- 790 • Activation function:
 - 791 – ReLU (Rectified Linear Unit) for the hidden layer
 - 792 – Softmax for the output layer
- 793 • Optimizer:
 - 794 – Adam (Adaptive Moment Estimation)
- 795 • Learning rate:
 - 796 – Initial learning rate set to 0.001
 - 797 – Uses adaptive learning rate ('constant' schedule)
- 798 • Regularization:
 - 799 – L2 regularization with alpha=0.0001
- 800 • Batch size:
 - 801 – Mini-batch gradient descent with a batch size of 200
- 802 • Maximum iterations:
 - 803 – Default maximum number of iterations is 200
- 804 • Early stopping:
 - 805 – Not enabled by default
- 806 • Initialization:
 - 807 – Xavier initialization for weight initialization

808 **Microbenchmarking Robust-PICASO.** For completeness, we implement Robust-PICASO on the BI-
 809 CYCL library [16] to show its running time. For completeness, we also benchmark the running time of the
 810 encryption and aggregation protocols steps of Robust-PICASO in Figure 4.

811 C Deferred Proofs

812 **Theorem 2.** *Under the DDH – f assumption (see Definition 1), Algorithm 2 is secure in the random oracle*
 813 *model.*

814 The proof of security of this theorem is through a sequence of hybrids. As mentioned in the introduction, the
 815 adversary triggers a challenge phase by presenting a set of honest clients whose inputs it wants to recover:
 816 H_1, \dots, H_t while also presenting two sets of challenge inputs for these clients x_1, \dots, x_t and x'_1, \dots, x'_t . The
 817 challenger chooses to encrypt one of these sets at random and provide it to the attacker. While we defer the
 818 formal proof due to space constraints, we present a description of the hybrids below:

- 819 • Hybrid 0: The challenger outputs 1 if the adversary guesses correctly which of the challenge set was
 820 chosen, else outputs 0.

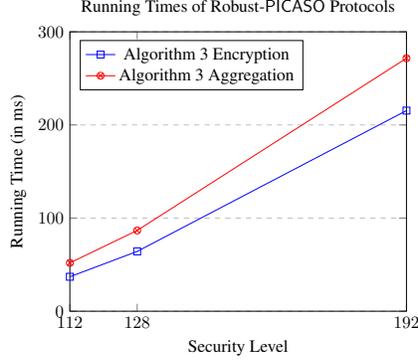


Figure 4: Performance of Robust-PICASO. Here, security level indicates the computational power needed to break the security of the protocol.

- 821 • Hybrid 1: The change is that for every query to the random oracle \mathcal{H} , the challenger tosses a biased
822 coin δ_t . The biasing of the coin is that it takes 1 with probability $1/q_{enc} + 1$ and 0 with the remaining
823 probability. Then let's define an event E : if for target iteration τ , $\delta_t a u = 0$ or for some $t \neq \tau$, for an
824 honest user i , there was an encryption query which produced $\delta_t = 1$. If E happens, then challenger
825 just outputs a random bit. $\Pr[\neg E] = \frac{1}{q_{enc}+1} \cdot \left(\frac{q_{enc}}{q_{enc}+1}\right)^{q_{enc}} \geq \frac{1}{e(q_{enc}+1)}$
- 826 • Hybrid 2: Now, further modify the generation of $\mathcal{H}(t)$. If $\delta_t = 0$, then we sample $w_t \leftarrow \mathcal{D}_H$
827 and setting $\mathcal{H}(t) := h^{w_t}$. Meanwhile, if $\delta_t = 1$, it additionally samples $u_t \leftarrow \mathbb{Z}/p\mathbb{Z}$ setting
828 $\mathcal{H}(t) := h^{w_t} \cdot f^{u_t}$. Note that the view between Hybrids 1 and 2 only happens when $\delta_t = 1$ and this
829 is protected by the DDH-f assumption.

830 Note that if E has not happened, then the challenge epoch has $\delta_\tau = 1$, which implies that $\mathcal{H}(\tau) := h^{w_\tau} \cdot f^{u_\tau}$.
831 In other words, every $ct_{i,\tau}$ will be represented as: $f^{x_{i,\tau} + u_\tau} \cdot \text{mask}_{i,\tau}$. We will then show that this u_τ provides
832 sufficient masking of the inputs. Observe that for the honest parties, the server does not know any information
833 about k_i . Meanwhile, $pk_{i,\tau}$ is also not provided to the adversary. Therefore, at the challenge epoch, we sample
834 individual values $u_{i,\tau}$ for honest i . Each of these $u_{i,\tau}$ masks the inputs $x_{i,\tau}$. Meanwhile, to ensure correctness
835 of decryption, the challenger correctly simulates AUX_τ .

836 D Other Extensions to PICASO

837 D.1 Asynchronous Secure Aggregation

838 While FL algorithms have handled the problem of stragglers by simply considering them as dropouts. However,
839 Asynchronous FL [30, 32, 90, 93] which focused on updating the global model, as soon as the local updates are
840 received, even if the updates are for an old iteration. This was designed to ensure that stragglers in an iteration
841 do not delay the global model update. Increased staleness in local models leads to greater errors when updating
842 the global model [93], which was remedied by their proposed staleness-aware weighted averaging protocol
843 called FedAsync. Unfortunately, Asynchronous FL is not easily composable with existing Secure Aggregation
844 techniques. [73] presented an approach where the local updates are buffered, before updating the global model,
845 with the buffer being stored in a trusted execution environment (TEE) to guarantee privacy. TEEs are known to
846 be expensive. Later, BASecAgg[84] avoided the use of TEE by composing the buffering technique with a secure
847 aggregation protocol, with the aggregation incorporating the staleness function. Unfortunately, BASecAgg
848 required that each client share its mask with every other client, for every iteration which is undesirable for
849 asynchrony.

850 **BASecAgg.** BASecAgg [84] successfully combined techniques of secure aggregation with asynchronous
851 FL. Specifically, the server aggregates the model weights as: $\sum_i \phi(t - t_i) \cdot x_i$ where t_i is the iteration count of
852 client i , with the corresponding updates being x_i . $\phi(t - t_i)$ is a "staleness" function which is 1 if $t = t_i$ and
853 is monotonically decreasing. BASecAgg presented a solution where the aggregation of shares accounted for
854 this staleness function. This is because LightSecAgg [87], unlike other works[13, 8], had the server reconstruct
855 the sum of the masks masks for the online clients which can then be used to unmask. Similarly, PICASO also
856 only helps the server reconstruct the online clients' masks. Therefore, using PICASO as the secure aggregation
857 component of BASecAgg, we can build an asynchronous secure aggregation protocol that avoids the expensive
858 secret sharing costs associated with LightSecAgg.

859 **D.2 Minimizing Trust Assumption**

860 Existing protocols that offer secure aggregation rely on a non-colluding assumption. Typically, this is modeled
861 by allowing the server to collude with up to a threshold t , out of n parties, with the security being completely
862 lost if even $t + 1$ parties are corrupted. It follows that if the server colludes with all n parties, then privacy is lost.
863 These parties are those involved in helping reconstruct the sum, even in the presence of dropouts. In SecAgg,
864 these parties are fellow clients while in LERNA and Flamingo, this corresponds to the committee members.

865 However, PICASO can be easily extended to support a committee of M collectors. In this setting, the server can
866 corrupt to a certain threshold t of collectors and need the help of at least $1 + t$ of collectors to reconstruct the
867 sum.

868 A naive implementation would be to secret-share the current public key among the committee of such collectors.
869 This technique is called secret sharing [78]. More recently, Braun et al. [17] demonstrated how to construct
870 secret sharing techniques, compatible with the CL Framework. However, it is to be noted that this comes at a
871 significant cost:

- 872 • The communication cost and computation cost scales linearly with the size of the committee.
- 873 • The server's computation and communication costs also increase. Additionally, the server has to
874 engage in expensive coordination with the committee. Factoring in networking delays, it is entirely
875 possible that different committee members receive inputs from different subsets of clients. Now, the
876 server has to find the intersection of online clients among this list.

877 **NeurIPS Paper Checklist**

878 The checklist is designed to encourage best practices for responsible machine learning research, addressing
879 issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The**
880 **papers not including the checklist will be desk rejected.** The checklist should follow the references and
881 precede the (optional) supplemental material. The checklist does NOT count towards the page limit.

882 Please read the checklist guidelines carefully for information on how to answer these questions. For each
883 question in the checklist:

- 884 • You should answer [Yes], [No], or [NA].
- 885 • [NA] means either that the question is Not Applicable for that particular paper or the relevant
886 information is Not Available.
- 887 • Please provide a short (1–2 sentence) justification right after your answer (even for NA).

888 **The checklist answers are an integral part of your paper submission.** They are visible to the reviewers, area
889 chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions)
890 with the final version of your paper, and its final version will be published with the paper.

891 The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While
892 "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper
893 justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or
894 "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not
895 grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is
896 often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting
897 evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer
898 [Yes] to a question, in the justification please point to the section(s) where related material for the question can
899 be found.

900 IMPORTANT, please:

- 901 • **Delete this instruction block, but keep the section heading "NeurIPS paper checklist",**
- 902 • **Keep the checklist subsection headings, questions/answers and guidelines below.**
- 903 • **Do not modify the questions and only use the provided macros for your answers.**

904 **1. Claims**

905 Question: Do the main claims made in the abstract and introduction accurately reflect the paper's
906 contributions and scope?

907 Answer: [Yes]

908 Justification: The paper sets out to solve a critical problem in prior work on secure aggregation. In
909 this work, we demonstrate how to reduce the synchronization by employing one additional party. We
910 demonstrate experiments to show competitive performance over prior work. In addition, we also train
911 machine learning models to justify that our protocol can be used for its intended purpose.

912 Guidelines:

- 913 • The answer NA means that the abstract and introduction do not include the claims made in the
914 paper.
- 915 • The abstract and/or introduction should clearly state the claims made, including the contributions
916 made in the paper and important assumptions and limitations. A No or NA answer to this
917 question will not be perceived well by the reviewers.
- 918 • The claims made should match theoretical and experimental results, and reflect how much the
919 results can be expected to generalize to other settings.
- 920 • It is fine to include aspirational goals as motivation as long as it is clear that these goals are not
921 attained by the paper.

922 **2. Limitations**

923 Question: Does the paper discuss the limitations of the work performed by the authors?

924 Answer: [Yes]

925 Justification: We have a conclusion paragraph that draws attention to some of the limitations while
926 identifying how they can be remedied in future work.

927 Guidelines:

- 928 • The answer NA means that the paper has no limitation while the answer No means that the paper
929 has limitations, but those are not discussed in the paper.
- 930 • The authors are encouraged to create a separate "Limitations" section in their paper.
- 931 • The paper should point out any strong assumptions and how robust the results are to violations of
932 these assumptions (e.g., independence assumptions, noiseless settings, model well-specification,
933 asymptotic approximations only holding locally). The authors should reflect on how these
934 assumptions might be violated in practice and what the implications would be.
- 935 • The authors should reflect on the scope of the claims made, e.g., if the approach was only tested
936 on a few datasets or with a few runs. In general, empirical results often depend on implicit
937 assumptions, which should be articulated.
- 938 • The authors should reflect on the factors that influence the performance of the approach. For
939 example, a facial recognition algorithm may perform poorly when image resolution is low or
940 images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide
941 closed captions for online lectures because it fails to handle technical jargon.
- 942 • The authors should discuss the computational efficiency of the proposed algorithms and how
943 they scale with dataset size.
- 944 • If applicable, the authors should discuss possible limitations of their approach to address problems
945 of privacy and fairness.
- 946 • While the authors might fear that complete honesty about limitations might be used by reviewers
947 as grounds for rejection, a worse outcome might be that reviewers discover limitations that
948 aren't acknowledged in the paper. The authors should use their best judgment and recognize
949 that individual actions in favor of transparency play an important role in developing norms that
950 preserve the integrity of the community. Reviewers will be specifically instructed to not penalize
951 honesty concerning limitations.

952 3. Theory Assumptions and Proofs

953 Question: For each theoretical result, does the paper provide the full set of assumptions and a complete
954 (and correct) proof?

955 Answer: [\[Yes\]](#)

956 Justification: The paper introduces all the necessary theoretical framework and assumptions for security
957 of the construction. There's detailed proof deferred to the appendix.

958 Guidelines:

- 959 • The answer NA means that the paper does not include theoretical results.
- 960 • All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- 961 • All assumptions should be clearly stated or referenced in the statement of any theorems.
- 962 • The proofs can either appear in the main paper or the supplemental material, but if they appear in
963 the supplemental material, the authors are encouraged to provide a short proof sketch to provide
964 intuition.
- 965 • Inversely, any informal proof provided in the core of the paper should be complemented by
966 formal proofs provided in appendix or supplemental material.
- 967 • Theorems and Lemmas that the proof relies upon should be properly referenced.

968 4. Experimental Result Reproducibility

969 Question: Does the paper fully disclose all the information needed to reproduce the main experimental
970 results of the paper to the extent that it affects the main claims and/or conclusions of the paper
971 (regardless of whether the code and data are provided or not)?

972 Answer: [\[Yes\]](#)

973 Justification: The protocols are well detailed, including the parameter settings for our classifier. We
974 use publicly available ABIDES framework to simulate real-life networking situations. For our class
975 group operations, we take the BICYCL framework that is open-source and bind it to Python.

976 Guidelines:

- 977 • The answer NA means that the paper does not include experiments.
- 978 • If the paper includes experiments, a No answer to this question will not be perceived well by the
979 reviewers: Making the paper reproducible is important, regardless of whether the code and data
980 are provided or not.
- 981 • If the contribution is a dataset and/or model, the authors should describe the steps taken to make
982 their results reproducible or verifiable.

- 983
- 984
- 985
- 986
- 987
- 988
- 989
- 990
- 991
- 992
- 993
- 994
- 995
- 996
- 997
- 998
- 999
- 1000
- 1001
- 1002
- 1003
- 1004
- 1005
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
 - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

1006 5. Open access to data and code

1007 Question: Does the paper provide open access to the data and code, with sufficient instructions to
1008 faithfully reproduce the main experimental results, as described in supplemental material?

1009 Answer: [NA]

1010 Justification: Unfortunately, there was no support for supplementary material upload. However, we
1011 are happy to furnish the anonymized code for interested reviewers.

1012 Guidelines:

- 1013
- 1014
- 1015
- 1016
- 1017
- 1018
- 1019
- 1020
- 1021
- 1022
- 1023
- 1024
- 1025
- 1026
- 1027
- 1028
- 1029
- 1030
- The answer NA means that paper does not include experiments requiring code.
 - Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
 - While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
 - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
 - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
 - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
 - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
 - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

1031 6. Experimental Setting/Details

1032 Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters,
1033 how they were chosen, type of optimizer, etc.) necessary to understand the results?

1034 Answer: [Yes]

1035 Justification: We detail the quantization metrics along with the choice of datasets with the test-train
1036 splits. Our choice of training algorithms are vanilla versions, with no customized hyperparameters.

1037 Guidelines:

- 1038
- 1039
- 1040
- The answer NA means that the paper does not include experiments.
 - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.

1041 • The full details can be provided either with the code, in appendix, or as supplemental material.

1042 7. Experiment Statistical Significance

1043 Question: Does the paper report error bars suitably and correctly defined or other appropriate informa-
1044 tion about the statistical significance of the experiments?

1045 Answer: [NA]

1046 Justification: Our running time experiments report the mean of 30 iterations, for various choices of
1047 number of clients. Meanwhile, we present the accuracy values, as a function of iterations for various
1048 datasets. At each iteration, the data is randomly split.

1049 Guidelines:

- 1050 • The answer NA means that the paper does not include experiments.
- 1051 • The authors should answer "Yes" if the results are accompanied by error bars, confidence
1052 intervals, or statistical significance tests, at least for the experiments that support the main claims
1053 of the paper.
- 1054 • The factors of variability that the error bars are capturing should be clearly stated (for example,
1055 train/test split, initialization, random drawing of some parameter, or overall run with given
1056 experimental conditions).
- 1057 • The method for calculating the error bars should be explained (closed form formula, call to a
1058 library function, bootstrap, etc.)
- 1059 • The assumptions made should be given (e.g., Normally distributed errors).
- 1060 • It should be clear whether the error bar is the standard deviation or the standard error of the
1061 mean.
- 1062 • It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report
1063 a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is
1064 not verified.
- 1065 • For asymmetric distributions, the authors should be careful not to show in tables or figures
1066 symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- 1067 • If error bars are reported in tables or plots, The authors should explain in the text how they were
1068 calculated and reference the corresponding figures or tables in the text.

1069 8. Experiments Compute Resources

1070 Question: For each experiment, does the paper provide sufficient information on the computer
1071 resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

1072 Answer: [Yes]

1073 Justification: We detail the system settings of the device on which experiments are performed.

1074 Guidelines:

- 1075 • The answer NA means that the paper does not include experiments.
- 1076 • The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud
1077 provider, including relevant memory and storage.
- 1078 • The paper should provide the amount of compute required for each of the individual experimental
1079 runs as well as estimate the total compute.
- 1080 • The paper should disclose whether the full research project required more compute than the
1081 experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into
1082 the paper).

1083 9. Code Of Ethics

1084 Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code
1085 of Ethics <https://neurips.cc/public/EthicsGuidelines>?

1086 Answer: [Yes]

1087 Justification: The paper complies with the code of ethics.

1088 Guidelines:

- 1089 • The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- 1090 • If the authors answer No, they should explain the special circumstances that require a deviation
1091 from the Code of Ethics.
- 1092 • The authors should make sure to preserve anonymity (e.g., if there is a special consideration due
1093 to laws or regulations in their jurisdiction).

1094 10. Broader Impacts

1095 Question: Does the paper discuss both potential positive societal impacts and negative societal impacts
1096 of the work performed?

1097 Answer: [Yes]

1098 Justification: The work focuses on privacy of client-held data. This is surveyed in the introduction and
1099 motivates why privacy-preserving federated learning is important and its positive impact.

1100 Guidelines:

- 1101 • The answer NA means that there is no societal impact of the work performed.
- 1102 • If the authors answer NA or No, they should explain why their work has no societal impact or
1103 why the paper does not address societal impact.
- 1104 • Examples of negative societal impacts include potential malicious or unintended uses (e.g.,
1105 disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deploy-
1106 ment of technologies that could make decisions that unfairly impact specific groups), privacy
1107 considerations, and security considerations.
- 1108 • The conference expects that many papers will be foundational research and not tied to particular
1109 applications, let alone deployments. However, if there is a direct path to any negative applications,
1110 the authors should point it out. For example, it is legitimate to point out that an improvement in
1111 the quality of generative models could be used to generate deepfakes for disinformation. On the
1112 other hand, it is not needed to point out that a generic algorithm for optimizing neural networks
1113 could enable people to train models that generate Deepfakes faster.
- 1114 • The authors should consider possible harms that could arise when the technology is being used
1115 as intended and functioning correctly, harms that could arise when the technology is being used
1116 as intended but gives incorrect results, and harms following from (intentional or unintentional)
1117 misuse of the technology.
- 1118 • If there are negative societal impacts, the authors could also discuss possible mitigation strategies
1119 (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitor-
1120 ing misuse, mechanisms to monitor how a system learns from feedback over time, improving the
1121 efficiency and accessibility of ML).

11. Safeguards

1122 Question: Does the paper describe safeguards that have been put in place for responsible release of
1123 data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or
1124 scraped datasets)?

1125 Answer: [NA]

1126 Justification: We do not use any of the stated models or data sources.

1127 Guidelines:

- 1128 • The answer NA means that the paper poses no such risks.
- 1129 • Released models that have a high risk for misuse or dual-use should be released with necessary
1130 safeguards to allow for controlled use of the model, for example by requiring that users adhere to
1131 usage guidelines or restrictions to access the model or implementing safety filters.
- 1132 • Datasets that have been scraped from the Internet could pose safety risks. The authors should
1133 describe how they avoided releasing unsafe images.
- 1134 • We recognize that providing effective safeguards is challenging, and many papers do not require
1135 this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

1136 Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper,
1137 properly credited and are the license and terms of use explicitly mentioned and properly respected?

1138 Answer: [Yes]

1139 Justification: The experiments have only used publicly available datasets with their license details
1140 specified in a tabular column.

1141 Guidelines:

- 1142 • The answer NA means that the paper does not use existing assets.
- 1143 • The authors should cite the original paper that produced the code package or dataset.
- 1144 • The authors should state which version of the asset is used and, if possible, include a URL.
- 1145 • The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- 1146 • For scraped data from a particular source (e.g., website), the copyright and terms of service of
1147 that source should be provided.

- 1150 • If assets are released, the license, copyright information, and terms of use in the package should
1151 be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for
1152 some datasets. Their licensing guide can help determine the license of a dataset.
- 1153 • For existing datasets that are re-packaged, both the original license and the license of the derived
1154 asset (if it has changed) should be provided.
- 1155 • If this information is not available online, the authors are encouraged to reach out to the asset's
1156 creators.

1157 13. **New Assets**

1158 Question: Are new assets introduced in the paper well documented and is the documentation provided
1159 alongside the assets?

1160 Answer: [NA]

1161 Justification: We are not releasing any new assets.

1162 Guidelines:

- 1163 • The answer NA means that the paper does not release new assets.
- 1164 • Researchers should communicate the details of the dataset/code/model as part of their sub-
1165 missions via structured templates. This includes details about training, license, limitations,
1166 etc.
- 1167 • The paper should discuss whether and how consent was obtained from people whose asset is
1168 used.
- 1169 • At submission time, remember to anonymize your assets (if applicable). You can either create an
1170 anonymized URL or include an anonymized zip file.

1171 14. **Crowdsourcing and Research with Human Subjects**

1172 Question: For crowdsourcing experiments and research with human subjects, does the paper include
1173 the full text of instructions given to participants and screenshots, if applicable, as well as details about
1174 compensation (if any)?

1175 Answer: [NA]

1176 Justification: There were no human subjects involved in this project.

- 1177 • The answer NA means that the paper does not involve crowdsourcing nor research with human
1178 subjects.
- 1179 • Including this information in the supplemental material is fine, but if the main contribution of the
1180 paper involves human subjects, then as much detail as possible should be included in the main
1181 paper.
- 1182 • According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other
1183 labor should be paid at least the minimum wage in the country of the data collector.

1184 15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

1185 Question: Does the paper describe potential risks incurred by study participants, whether such
1186 risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an
1187 equivalent approval/review based on the requirements of your country or institution) were obtained?

1188 Answer: [NA]

1189 Justification: We do not have any research with human subjects that forms a part of this work.

1190 Guidelines:

- 1191 • The answer NA means that the paper does not involve crowdsourcing nor research with human
1192 subjects.
- 1193 • Depending on the country in which research is conducted, IRB approval (or equivalent) may be
1194 required for any human subjects research. If you obtained IRB approval, you should clearly state
1195 this in the paper.
- 1196 • We recognize that the procedures for this may vary significantly between institutions and
1197 locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for
1198 their institution.
- 1199 • For initial submissions, do not include any information that would break anonymity (if applica-
1200 ble), such as the institution conducting the review.