# Securing Secure Aggregation: Mitigating Multi-Round Privacy Leakage in Federated Learning

Abstract

Secure aggregation is a critical component in federated learning (FL), which enables the server to learn the aggregate model of the users without observing their local models. Conventionally, secure aggregation algorithms focus only on ensuring the privacy of individual users in a *single* training round. We contend that such designs can lead to significant privacy leakages over *multiple* training rounds, due to partial user selection/participation at each round of FL. In fact, we show that the conventional random user selection strategies in FL may lead to leaking users' individual models within a number of rounds that is linear in the number of users. To address this challenge, we introduce a secure aggregation framework, Multi-RoundSecAgg, with multi-round privacy guarantees. In particular, we introduce a new metric to quantify the privacy guarantees of FL over multiple training rounds, and develop a structured user selection strategy that guarantees the long-term privacy of each user (over any number of training rounds). Our framework also carefully accounts for the fairness and the average number of participating users at each round. Our experiments on MNIST, CIFAR-10 and CIFAR-100 datasets in the IID and the non-IID settings demonstrate the performance improvement over the baselines, both in terms of privacy protection and test accuracy.

1. Introduction

Federated learning (FL) enables collaborative training of machine learning models over the data collected and stored locally by multiple data-owners (users). The training in FL is typically coordinated by a central server who maintains a global model that is updated locally by the users. The local updates are then aggregated by the server to update the global model. Throughout the training process, the users never share their data with the server, i.e., the data is always kept on device, rather, they only share their local updates. However, as has been shown recently, the local models may still reveal substantial information about the local datasets, and the private training data can be reconstructed from the local models through inference or inversion attacks (see
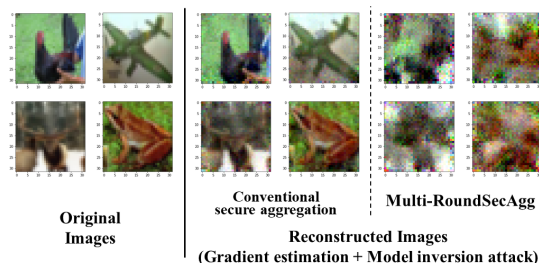
Figure 1: A qualitative comparison of the reconstructed images in two settings is shown. The first setting corresponds to the case that model privacy with random user selection (e.g., `FedAvg` (McMahan et al., 2018)) is protected by conventional secure aggregation schemes as (Bonawitz et al., 2017) at each round. The second setting corresponds to our proposed method, Multi-RoundSecAgg, which ensures the long-term privacy of individual models over any number of rounds, and hence model inversion attack cannot work well. This reconstruction process is described in detail in Appendix G.

e.g., (Fredrikson et al., 2015; Nasr et al., 2019; Zhu & Han, 2020; Geiping et al., 2020)).

To prevent such information leakage, *secure aggregation* protocols are proposed (e.g., (Bonawitz et al., 2017; So et al., 2021b; Kadhe et al., 2020; Zhao & Sun, 2021; Bell et al., 2020; Yang et al., 2021; So et al., 2021a)) to protect the privacy of the local models, both from the server and the other users, while still allowing the server to learn their aggregate. More specifically, the secure aggregation protocols ensure that, at any given round, the server can only learn the aggregate model of the users, and beyond that no further information is revealed about the individual model.

Secure aggregation protocols, however, only ensure the privacy of the individual users in a *single training round*, and do not consider their privacy over multiple training rounds. On the other hand, due to partial user selection (Cho et al., 2020b; Chen et al., 2020; Cho et al., 2020a; Ribero & Vikalo, 2020), the server may be able to reconstruct the individual models of some users using the aggregated models from the previous rounds. In fact, we show that after a sufficient number of rounds, all local models can be recovered with a high accuracy if the server uniformly chooses a random subset of the users to participate at every round. As shown in Fig.1, performing model inversion attack (Geiping et al., 2020) with the recovered local models yields reconstructed images with a similar quality as the original images.

Contributions. As such motivated, we study long-term user privacy in FL. Specifically, our contributions are as follows

1. We introduce a new metric to capture long-term privacy guarantees in FL for the first time. This long-term privacy condition requires that the server cannot reconstruct any individual model using the aggregated models from any number of training rounds. Using this metric, we show that the conventional random user selection schemes can result in leaking the individual user models after a sufficient number of rounds, even if secure aggregation is employed at each round.

2. We propose Multi-RoundSecAgg, a privacy-preserving structured user selection strategy that ensures the long-term privacy of the individual users over any number of training rounds. This strategy also takes into account the fairness of the selection process and the average number of participating users at each round.

3. We demonstrate that Multi-RoundSecAgg creates a trade-off between the long-term privacy guarantee and the average number of participating users. In particular, as the average number of participating users increases, the long-term privacy guarantee becomes weaker.

4. We provide the convergence analysis of Multi-RoundSecAgg, which shows that the long-term privacy guarantee and the average number of participating users control the convergence rate. The convergence rate is maximized when the average number of participating users is maximized. (e.g., the random user selection strategy maximizes the average number of participating users at the expense of not providing long-term privacy guarantees). As we require stronger long-term privacy guarantees, the average number of participating users decreases and a larger number of training rounds is required to achieve the same level of accuracy as the random selection strategy.

5. Finally, our experiments in both IID and non-IID settings on MNIST, CIFAR-10 and CIFAR-100 datasets demonstrate that Multi-RoundSecAgg achieves almost the same test accuracy compared to the random selection scheme while providing better long-term privacy guarantees.

2. Related Work

The underlying principle of the secure aggregation protocol in (Bonawitz et al., 2017) is that each pair of users exchange a pairwise secret key which they can use to mask their local models before sharing them with the server. The pairwise masks cancel out when the server aggregates the masked models, allowing the server to learn the aggregate of the local models. These masks also ensure that the local models are kept private, i.e., no further information is revealed beyond the aggregate of the local models. This protocol,

however, incurs a significant communication cost due to exchanging and reconstructing the pairwise keys.

Recently, several works have developed computation and communication-efficient protocols (So et al., 2021b; Kadhe et al., 2020; Bell et al., 2020; Tang et al., 2021; Choi et al., 2020; Elkordy & Avestimehr, 2020; Yang et al., 2021), which are complementary to and can be combined with our work. Another line of work focused on designing partial user selection strategies to overcome the communication bottleneck in FL while speeding up the convergence by selecting the users based on their local loss (Cho et al., 2020b; Chen et al., 2020; Cho et al., 2020a; Ribero & Vikalo, 2020).

Previous works, either on secure aggregation or on partial user selection, however, do not consider mitigating the potential privacy leakage as a result of partial user participation and the server observing the aggregated models across multiple training rounds. While (Pejó & Biczók, 2020) pointed out to the privacy leakage of secure aggregation, mitigating this leakage has not been considered and our work is the first to address this challenge.

3. System Model

In this section, we first describe the basic federated learning model in Section 3.1. Next, we introduce the multi-round secure aggregation problem for federated learning and define the key metrics to evaluate the performance of a multi-round secure aggregation protocol in Section 3.2.

3.1. Basic Federated Learning Model

We consider a cross-device federated learning setup consisting of a server and N users. User $i \in [N]$ has a local dataset $\mathcal{D}_i$ consisting of $m_i = |\mathcal{D}_i|$ data samples. The users are connected to each other through the server, i.e., all communications between the users goes through the server (McMahan et al., 2017; Bonawitz et al., 2017; Kairouz et al., 2019). The goal is to collaboratively learn a global model $\boldsymbol{x}$ with dimension d, using the local datasets that are generated, stored, and processed locally by the users. The training task can be represented by minimizing a global loss function,

$$\min_{\boldsymbol{x}} L(\boldsymbol{x}) \text{ s.t. } L(\boldsymbol{x}) = \frac{1}{\sum_{i=1}^{N} w_i} \sum_{i=1}^{N} w_i L_i(\boldsymbol{x}), \quad (1)$$

where L_i is the loss function of user i and $w_i \geq 0$ is a weight parameter assigned to user i to specify the relative impact of that user. A common choice for the weight parameters is $w_i = m_i$ (Kairouz et al., 2019). We define the optimal model parameters $\boldsymbol{x}^*$ and $\boldsymbol{x}_i^*$ as $\boldsymbol{x}^* = \arg\min_{\boldsymbol{x} \in \mathbb{R}^d} L(\boldsymbol{x})$ and $\boldsymbol{x}_i^* = \arg\min_{\boldsymbol{x} \in \mathbb{R}^d} L_i(\boldsymbol{x})$.

Federated Averaging with Partial User Participation. To solve (1), the most common algorithm is the *FedAvg*

(federated averaging) algorithm (McMahan et al., 2017). *FedAvg* is an iterative algorithm, where the model training is done by repeatedly iterating over individual local updates. At the beginning of training round t, the server sends the current state of the global model, denoted by $\boldsymbol{x}^{(t)}$, to the users. Each round consists of two phases, local training and aggregation. In the local training phase, user $i \in [N]$ updates the global model by carrying out $E (\geq 1)$ local stochastic gradient descent (SGD) steps and sends the updated local model $\boldsymbol{x}_i^{(t)}$ to the server. One of key features of cross-device FL is partial device participation. Due to various reasons such as unreliable wireless connectivity, or battery issues, at any given round, only a fraction of the users are available to participate in the protocol. We refer to such users as *available* users throughout the paper. In the aggregation phase, the server selects $K \leq N$ users among the available users if this is possible and aggregates their. After receiving the local updates, the server updates the global model as follows

$$\boldsymbol{x}^{(t+1)} = \sum_{i \in \mathcal{S}^{(t)}} w_i' \boldsymbol{x}_i^{(t)} = \mathbf{X}^{(t)^\top} \boldsymbol{p}^{(t)}, \quad (2)$$

where $\mathcal{S}^{(t)}$ is the set of participating users at round t, $w_i' = \frac{w_i}{\sum_{i \in \mathcal{S}^{(t)}} w_i}$, and $\boldsymbol{p}^{(t)} \in \{0, 1\}^N$ is the corresponding characteristic vector. That is, $\boldsymbol{p}^{(t)}$ denotes a participation vector at round t whose i-th entry is 0 when user i is not selected and 1 otherwise. $\mathbf{X}^{(t)}$ denotes the concatenation of the weighted local models at round t, i.e., $\mathbf{X}^{(t)} = \left[w_1' \boldsymbol{x}_1^{(t)}, \ldots, w_N' \boldsymbol{x}_N^{(t)} \right]^\top \in \mathbb{R}^{N \times d}$. Finally, the server broadcasts the updated global model $\boldsymbol{x}^{(t+1)}$ to the users for the next round.

Threat Model. Similar to the prior works on secure aggregation as (Bonawitz et al., 2017; Kadhe et al., 2020; So et al., 2021b), we consider the honest-but-curious server's model. The server follows the protocol honestly in this model, but tries to learn as much as possible about the users. At each round, the privacy of individual model $\boldsymbol{x}_i^{(t)}$ in (2) is protected by secure aggregation such that the server only learns the aggregate of the selected models $\sum_{i \in \mathcal{S}^{(t)}} w_i' \boldsymbol{x}_i^{(t)}$.

3.2. Multi-round Secure Aggregation

Conventional secure aggregation protocols only consider the privacy guarantees over a single training round. While secure aggregation protocols have provable privacy guarantees at any single round, in the sense that no information is leaked beyond the aggregate model at each round, the privacy guarantees do not extend to attacks *that span multiple training rounds*. Specifically, by using the aggregate models and participation information across multiple rounds, an individual model may be reconstructed. For instance, consider the following user participation strategy across three training rounds, $\boldsymbol{p}^{(1)} = [1, 1, 0]^\top$,

$\boldsymbol{p}^{(2)} = [0, 1, 1]^\top$, and $\boldsymbol{p}^{(3)} = [1, 0, 1]^\top$. Assume a scenario where the local updates do not change significantly over time (e.g., models start to converge, or the server fixes the global model over consecutive rounds), i.e., $\boldsymbol{x}_i = \boldsymbol{x}_i^{(t)}$ for all $i \in [3]$ and $t \in [3]$. Then, the server can single out individual model, e.g., $\boldsymbol{x}_1 = (\boldsymbol{x}^{(1)} + \boldsymbol{x}^{(3)} - \boldsymbol{x}^{(2)})/2$. Similarly, the server can single out all individual models $\boldsymbol{x}_i$, even if a secure aggregation protocol is employed at each round.

In this paper, we study secure aggregation protocols with long-term privacy guarantees (which we term *multi-round secure aggregation*) for the cross-device FL setup which has not been studied before. We assume that user $i \in [N]$ drops from the protocol at each round with probability p_i. $\mathcal{U}^{(t)}$ denotes the index set of available users at round t and $\boldsymbol{u}^{(t)} \in \{0, 1\}^N$ is a vector indicating the available users such that $\{\boldsymbol{u}^{(t)}\}_j = \mathbb{1}\{j \in \mathcal{U}^{(t)}\}$, where $\{\boldsymbol{u}\}_j$ is j-th entry of $\boldsymbol{u}$ and $\mathbb{1}\{\cdot\}$ is the indicator function. The server selects K users from $\mathcal{U}^{(t)}$, if $|\mathcal{U}^{(t)}| \geq K$, based on the history of selected users in previous rounds. If $|\mathcal{U}^{(t)}| < K$, the server skips this round. The local models of the selected users are then aggregated via a secure aggregation protocol (i.e., by communicating masked models), at the end of which the server learns the aggregate of the local models of the selected users. Our goal is to design a user selection algorithm $\mathcal{A}^{(t)} : \{0, 1\}^{t \times N} \times \{0, 1\}^N \rightarrow \{0, 1\}^N$,

$$\mathcal{A}^{(t)}\big(\mathbf{P}^{(t)}, \boldsymbol{u}^{(t)}\big) = \boldsymbol{p}^{(t)} \text{ such that } \|\boldsymbol{p}^{(t)}\|_0 \in \{0, K\}, \quad (3)$$

to prevent the potential information leakage over multiple rounds, where $\boldsymbol{p}^{(t)} \in \{0, 1\}^N$ is the participation vector defined in (2), $\|\boldsymbol{x}\|_0$ denotes the L_0-"norm" of a vector $\boldsymbol{x}$ and K denotes the maximum number of selected users. We note that $\mathcal{A}^{(t)}$ can be a random function. $\mathbf{P}^{(t)}$ is a matrix representing the user participation information up to round t, and is termed the *participation matrix*, given by

$$\mathbf{P}^{(t)} = \left[\boldsymbol{p}^{(0)}, \boldsymbol{p}^{(1)}, \ldots, \boldsymbol{p}^{(t-1)} \right]^\top \in \{0, 1\}^{t \times N}. \quad (4)$$

Key Metrics. A multi-round secure aggregation protocol can be represented by $\mathcal{A} = \{\mathcal{A}^{(t)}\}_{t \in [J]}$, where $\mathcal{A}^{(t)}$ is the user selection algorithm at round t defined in (3) and J is the total number of rounds. The inputs of $\mathcal{A}^{(t)}$ are a random vector $\boldsymbol{u}^{(t)}$, which indicates the available users at round t, and the participation matrix $\mathbf{P}^{(t)}$ defined in (4) which can be a random matrix. Given the participation matrix $\mathbf{P}^{(J)}$, we evaluate the performance of the corresponding multi-round secure aggregation protocol through the following metrics.

1. **Multi-round Privacy Guarantee.** The secure aggregation protocols ensure that the server can only learn the sum of the local models of some users in each single round, but they do not consider what the server can learn over the long run.

Our multi-round privacy definition extends the guarantees of the secure aggregation protocols from one round to all training rounds by requiring that the server can only learn a sum of the local models even if the server exploits the aggregate models of all rounds. That is, our multi-round privacy guarantee is a natural extension of the privacy guarantee provided by the secure aggregation protocols considering a single training round.

Specifically, a multi-round privacy guarantee T requires that any non-zero partial sum of the local models that the server can reconstruct, through any linear combination $\mathbf{X}^\top \mathbf{P}^{(J)^\top} \boldsymbol{z}$, where $\boldsymbol{z} \in \mathbb{R}^J \setminus \{\mathbf{0}\}$, must be of the form[1]

$$
\mathbf{X}^\top \mathbf{P}^{(J)^\top} \boldsymbol{z} = \sum_{i\in[n]} a_i \sum_{j\in\mathcal{S}_i} \boldsymbol{x}_j
$$

$$
= a_1 \sum_{j\in\mathcal{S}_1} \boldsymbol{x}_j + a_2 \sum_{j\in\mathcal{S}_2} \boldsymbol{x}_j + \cdots + a_n \sum_{j\in\mathcal{S}_n} \boldsymbol{x}_j, \quad (5)
$$

where $|\mathcal{S}_i| \geq T, a_i \neq 0, \forall i \in [n]$ and $n \in \mathbb{Z}^+$. Here all the sets $\mathcal{S}_i$, the number of sets n, and each a_i could all depend on $\boldsymbol{z}$. In equation (5), we consider the worst-case scenario, where the local models do not change over the rounds. That is, $\mathbf{X}^{(t)} = \mathbf{X}, \forall t \in [J]$. Intuitively, this guarantee ensures that the best that the server can do is to reconstruct a partial sum of T local models which corresponds to the case where $n = 1$. When $T \geq 2$, this condition implies that the server cannot get any user model from the aggregate models of all training rounds (the best it can obtain is the sum of two local models).

Remark 1. (Weaker Privacy Notion). It is worth noting that, a weaker privacy notion would require that $\|\mathbf{P}^{(J)^\top} \boldsymbol{z}\|_0 \geq T$ when $\mathbf{P}^{(J)^\top} \boldsymbol{z} \neq \mathbf{0}$. When $T = 2$, this definition requires that the server cannot reconstruct any individual model (the best it can do is to obtain a linear combination of two local models). This notion, however, allows constructions in the form of $a\boldsymbol{x}_i + b\boldsymbol{x}_j$ for any $a, b \in \mathbb{R} \setminus \{0\}$. When $a \gg b$, however, this is almost the same as recovering $\boldsymbol{x}_i$ perfectly, hence this privacy criterion is weaker than that of (5).

Remark 2. (Multi-round Privacy of Random Selection). In Section 6, we empirically show that a random selection strategy in which K available users are selected uniformly at random at each round does not ensure multi-round privacy even with respect to the weaker definition of Remark 1. Specifically, the local models can be reconstructed within a number of rounds that is linear in N. We also show theoretically in Appendix G that when $\min(N - K, K) \geq cN$, where $c > 0$ is a constant, then the probability that the server can reconstruct all local models after N rounds is at least $1 - 2e^{-c'N}$ for

---

[1]We assume that $w_i = \frac{1}{N}, \forall i \in [N]$ in this paper.

a constant c' that depends on c. Finally, we show that a random selection scheme in which the users are selected in an i.i.d fashion according to $\text{Bern}(\frac{K}{N(1-p)})$ reveals all local models after N rounds with probability that converges to 1 exponentially fast.

Remark 3. (Worst-Case Assumption). In (5), we considered the worst-case assumption where the models do not change over time. When the local models change over rounds, the multi-round privacy guarantee becomes even stronger as the number of unknowns increases.

2. **Aggregation Fairness Gap.** The average aggregation fairness gap quantifies the largest gap between any two users in terms of the expected relative number of rounds each user has participated in training. Formally, the average aggregation fairness gap is defined as follows

$$
F = \max_{i\in[N]} \limsup_{J\to\infty} \frac{1}{J} \mathbb{E}\Big[\sum_{t=0}^{J-1} \mathbb{1}\big\{\{\boldsymbol{p}^{(t)}\}_i = 1\big\} \Big] -
$$
$$
\min_{i\in[N]} \liminf_{J\to\infty} \frac{1}{J} \mathbb{E}\Big[\sum_{t=0}^{J-1} \mathbb{1}\big\{\{\boldsymbol{p}^{(t)}\}_i = 1\big\} \Big], \quad (6)
$$

where $\{\boldsymbol{p}^{(t)}\}_i$ is i-th entry of the vector $\boldsymbol{p}^{(t)}$ and the expectation is over the randomness of the user selection algorithm $\mathcal{A}$ and the user availability. The main intuition behind this definition is that when $F = 0$, all users participate on average on the same number of rounds. This is important to take the different users into consideration equally and our experiments show that the accuracy of the schemes with small F are much higher than the schemes with high F.

3. **Average Aggregation Cardinality.** The aggregation cardinality quantifies the expected number of models to be aggregated per round. Formally, it is defined as

$$
C = \liminf_{J\to\infty} \frac{\mathbb{E}\big[\sum_{t=0}^{J-1} \|\boldsymbol{p}^{(t)}\|_0 \big]}{J}, \quad (7)
$$

where the expectation is over the randomness in $\mathcal{A}$ and the user availability. Intuitively, less number of rounds are needed to converge as more users participate in the training. In fact, as we show in Section 5.2, C directly controls the convergence rate.

3.3. Baseline Schemes

In this subsection, we introduce three baseline schemes for multi-round secure aggregation.
Random Selection. In this scheme, at each round, the server selects K users at random from the set of available users if this is possible.
Random Weighted Selection. This scheme is a modified version of random selection to reduce F when the dropout

probabilities of the users are not equal. Specifically, K users are selected at random from the available users with the minimum frequency of participation in the previous rounds. If less than K users are available, the server skips this round.

User Partitioning (Grouping). In this scheme, the users are partitioned into $G = N/K$ equal-sized groups denoted as $\mathcal{G}_1, \mathcal{G}_2, \cdots, \mathcal{G}_G$. At each round, the server selects one of the groups if none of the users in this group has dropped out. If multiple groups are available, to reduce the aggregation fairness gap, the server selects a group including a user with the minimum frequency of participation in previous rounds. If no group is available, the server skips this round.

4. Proposed Scheme: Multi-RoundSecAgg

In this section, we present Multi-RoundSecAgg, which has two components as follows.

- The first component designs a family of sets of users that satisfy the multi-round privacy requirement. The inputs of the first component are the number of users (N), the desired number of selected at each round (K), and the desired multi-round privacy guarantee (T). The output is a family of sets of K users satisfying the multi-round privacy guarantee T, termed as a *privacy-preserving family*. This family is represented by a matrix $\mathbf{B}$, where the rows are the characteristic vectors of these user sets.

- The second component selects a set from this designed family to satisfy the fairness guarantee. The inputs to the second component are the privacy-preserving family $\mathbf{B}$, the set of available users at round t, $\mathcal{U}^{(t)}$, and the frequency of participation of each user. The output is the set of users that will participate at round t.

We now describe these two components in detail.

Component 1 (Batch Partitioning (BP) of the users to guarantee multi-round privacy). The first component designs a family of R_{BP} sets, where R_{BP} is given in (8), satisfying the multi-round privacy requirement T. We denote the $R_{\mathrm{BP}} \times N$ binary matrix corresponding to these sets by $\mathbf{B} = [\boldsymbol{b}_1, \cdots, \boldsymbol{b}_{R_{\mathrm{BP}}}]^\top$, where $\|\boldsymbol{b}_i\|_0 = K, \forall i \in [R_{\mathrm{BP}}]$. That is, the rows of $\mathbf{B}$ are the characteristic vectors of those sets. The main idea of our scheme is to restrict certain sets of users of size T, denoted as batches, to either participate together or not participate at all. This guarantees a multi-round privacy T as we show in Section 5.

To construct a family of sets with this property, the users are first partitioned into N/T batches. At any given round, either all or none of the users of a particular batch participate in training. The server can choose K/T batches to participate in training, provided that all users in any

given selected batch are available. Since there are $\binom{N/T}{K/T}$ possible sets with this property, then the size of this privacy-preserving family of sets is given by[2]

$$R_{\mathrm{BP}} \overset{\mathrm{def}}{=} \binom{N/T}{K/T}. \tag{8}$$

In the extreme case of $T = 1$, this strategy specializes to random selection where the server can choose any set from the $\binom{N}{K}$ possible sets of K users. In the other extreme case of $T = K$, this strategy specializes to the partitioning strategy where there are N/K possible sets. We next provide an example to illustrate the construction of $\mathbf{B}$.

Example 1 ($N = 8, K = 4, T = 2$). In this example, the users are partitioned into 4 batches as $\mathcal{G}_1 = \{1,2\}, \mathcal{G}_2 = \{3,4\}, \mathcal{G}_3 = \{5,6\}$ and $\mathcal{G}_4 = \{7,8\}$ as given in (9). The server can choose any two batches out of these 4 batches, hence we have $R_{\mathrm{BP}} = \binom{4}{2} = 6$ possible sets. This ensures a multi-round privacy $T = 2$.

$$\mathbf{B} = \begin{pmatrix} 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \end{pmatrix}. \tag{9}$$

Component 2 (Available batch selection to guarantee fairness). At round t, user $i \in [N]$ is available to participate in the protocol with a probability $1 - p_i \in (0, 1]$. The frequency of participation of user i before round t is denoted by $f_i^{(t)} \overset{\mathrm{def}}{=} \sum_{j=0}^{t-1} \mathbb{1}\left\{\{\boldsymbol{p}^{(j)}\}_i = 1\right\}$. Given the set of available users at round t, $\mathcal{U}^{(t)}$, and the frequencies of participation $\boldsymbol{f}^{(t-1)} = (f_1^{(t-1)}, \cdots, f_N^{(t-1)})^\top$, the server selects K users.

To do so, the server first finds the submatrix of $\mathbf{B}$ denoted by $\mathbf{B}^{(t)}$ corresponding to the set of available users at round t, $\mathcal{U}^{(t)}$. Specifically, the i-th row of $\mathbf{B}$ denoted by $\boldsymbol{b}_i^\top$ is included in $\mathbf{B}^{(t)}$ provided that $\mathrm{supp}(\boldsymbol{b}_i) \subseteq \mathcal{U}^{(t)}$. If $\mathbf{B}^{(t)}$ is an empty matrix, then the server skips this round. Otherwise, the server selects a row from $\mathbf{B}^{(t)}$ uniformly at random if $p_i = p, \forall i \in [N]$. If the users have different dropout probabilities, the server selects a row from $\mathbf{B}^{(t)}$ that includes the user with the minimum frequency of participation $\ell_{\min}^{(t-1)} \overset{\mathrm{def}}{=} \arg\min_{i \in \mathcal{U}^{(t)}} f_i^{(t-1)}$. If there are many such rows, then the server selects one of them uniformly at random.

Remark 4. (Necessity of the Second Component). The second component is necessary to guarantee that the aggregation fairness gap goes to zero as we show in Theorem 1 and Section 6.

Overall, the algorithm first designs a privacy-preserving family of sets of K users such that any set can be selected

---

[2]We assume for simplicity that N/T and K/T are integers.
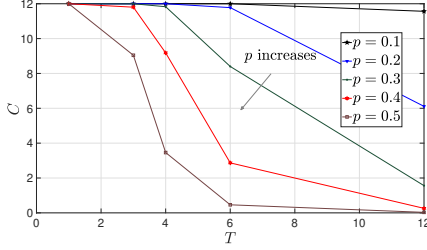
Figure 2: An illustration of the trade-off between the multi-round privacy guarantee T and the average aggregation cardinality C. In this example, $N = 120$ and $K = 12$.

at any given round while ensuring the multi-round privacy guarantee T. Then a specific set is selected from this family at each round to ensure fairness.

We describe the two components of Multi-RoundSecAgg in detail in Algorithm 1 and Algorithm 2 in Appendix D.

5. Theoretical Results

In this section, we provide the theoretical guarantees of Multi-RoundSecAgg in Section 5.1 and the convergence analysis of Multi-RoundSecAgg in Section 5.2.

5.1. Theoretical Guarantees of Multi-RoundSecAgg

In this subsection, we establish the theoretical guarantees of Multi-RoundSecAgg in terms of the multi-round privacy guarantee, the aggregation fairness gap and the average aggregation cardinality.

Theorem 1. *Multi-RoundSecAgg with parameters N, K, T ensures a multi-round privacy guarantee of T, an aggregation fairness gap $F = 0$, and an average aggregation cardinality given by*

$$C = K \left(1 - \sum_{i=N/T-K/T+1}^{N/T} \binom{N/T}{i} q^i (1-q)^{N/T-i} \right),$$

where $q = 1 - (1-p)^T$, when all users have the same dropout probability p.

We provide the proof of Theorem 1 in Appendix A.

Remark 5. (Trade-off between "Multi-round Privacy Guarantee" and "Average Aggregation Cardinality"). Theorem 1 indicates a trade-off between the multi-round privacy and the average aggregation cardinality since as T increases, C decreases which slows down the convergence as we show in Sec. 5.2. We show this trade-off in Fig. 2.

Remark 6. (Necessity of Batch Partitioning (BP)). We show that any strategy that satisfies the privacy guarantee in Equation (5) must have a batch partitioning structure, and for given $N, K, T, K \leq N/2$, the largest number of distinct user sets in any strategy is at most $\binom{N/T}{K/T}$, which is achieved in our design in Section 4. We provide the proof in Appendix C.

Remark 7. (Non-linear Reconstructions of Aggregated Models). The privacy criterion in Equation (5) considers linear reconstructions of the aggregated models. One may also consider more general non-linear reconstructions. The long-term privacy guarantees of batch partitioning hold even under such reconstructions as the users in the same batch always participate together or do not participate at all. Hence, the server cannot separate individual models within the same batch even through non-linear operations.

5.2. Convergence Analysis of Multi-RoundSecAgg

We now provide the convergence guarantees of Multi-RoundSecAgg by first introducing a few common assumptions (Li et al., 2019; Yu et al., 2019).

Assumption 1. $L_1, \ldots, L_N$ in (1) are all ρ-smooth: for all $a, b \in \mathbb{R}^d$ and $i \in [N]$, $L_i(a) \leq L_i(b) + (a-b)^\top \nabla L_i(b) + \frac{\rho}{2}\|a - b\|^2$.

Assumption 2. $L_1, \ldots, L_N$ in (1) are all μ-strongly convex: for all $a, b \in \mathbb{R}^d$ and $i \in [N]$, $L_i(a) \geq L_i(b) + (a - b)^\top \nabla L_i(b) + \frac{\mu}{2}\|a - b\|^2$.

Assumption 3. Let $\xi_i^{(t)}$ be a sample uniformly selected from the local dataset $\mathcal{D}_i$. The variance of the stochastic gradients at each user is bounded, i.e., $\mathbb{E}\|\nabla L_i(x_i^{(t)}, \xi_i^{(t)}) - \nabla L_i(x_i^{(t)})\|^2 \leq \sigma_i^2$ for $i \in [N]$.

Assumption 4. The expected squared norm of the stochastic gradients is uniformly bounded, i.e., $\mathbb{E}\|\nabla L_i(x_i^{(t)}, \xi_i^{(t)})\|^2 \leq G^2$ for all $i \in [N]$.

We now state the convergence guarantees.

Theorem 2. *Consider a FL setup with N users to train a machine learning model from (1). Assume K users are selected by Multi-RoundSecAgg with average aggregation cardinality C defined in (7) to update the global model from (2), and all users have the same dropout rate, hence Multi-RoundSecAgg selects a random set of K users uniformly from the set of available user sets at each round. Then, the following is satisfied*

$$\mathbb{E}[L(x^{(J)})] - L^* \leq \tag{10}$$
$$\frac{\rho}{\gamma + \frac{C}{K}EJ - 1} \left(\frac{2(\alpha + \beta)}{\mu^2} + \frac{\gamma}{2}\mathbb{E}\|x^{(0)} - x^*\|^2 \right),$$

where $\alpha = \frac{1}{N}\sum_{i=1}^N \sigma_i^2 + 6\rho\Gamma + 8(E-1)^2 G^2$, $\beta = \frac{4(N-K)E^2 G^2}{K(N-1)}$, $\Gamma = L^* - \sum_{i=1}^N L_i^*$, and $\gamma = \max\left\{\frac{8\rho}{\mu}, E\right\}$.

We provide the proof of Theorem 2 in Appendix B.

Remark 8. (The average aggregation cardinality controls the convergence rate.) Theorem 2 shows how the average aggregation cardinality affects the convergence. When the average aggregation cardinality is maximized, i.e., $C = K$, the convergence rate in Theorem 2 equals that of the random

selection algorithm provided in Theorem 3 of (Li et al., 2019). In (10), we have the additional term E (number of local epochs) in front of J compared to Theorem 3 of (Li et al., 2019) as we use global round index t instead of using step index of local SGD. As the average aggregation cardinality decreases, a greater number of training rounds is required to achieve the same level of accuracy.

Remark 9. (General Convex and Non-Convex Convergence Rates). Theorem 2 considers the strongly-convex case, but we consider the general convex and the non-convex cases in Appendix H.

Remark 10. (Different Dropout Rates). When the dropout probabilities of the users are not the same, characterizing the theoretical and convergence guarantees of Multi-RoundSecAgg is challenging. This is due to the fact that batch selection based on the frequency of participation breaks the conditional unbiasedness of the user selection, which is required for the convergence guarantee.

6. Experiments

Our experiments consist of two parts. We first numerically demonstrate the performance of Multi-RoundSecAgg compared to the baselines of Section 3.3 in terms of the key metrics of Section 3.2. Next, we implement convolutional neural networks (CNNs) for image classification with MNIST (LeCun et al., 2010), CIFAR-10, and CIFAR-100 (Krizhevsky & Hinton, 2009) to investigate how the key metrics affect the test accuracy.

Setup. We consider a FL setting with $N = 120$ users, where the server aims to choose $K = 12$ users at every round. We study two settings for partitioning the CIFAR-100 dataset across the users.

- **IID Setting.** The 50000 training samples are shuffled and partitioned uniformly across the $N = 120$ users.

- **Non-IID Setting.** We distribute the dataset using a Dirichlet distribution (Hsu et al., 2019), which samples $\mathbf{d}_c \sim \text{Dir}(\beta = 0.5)$ which specifying the prior class distribution over 100 classes, and allocate a portion $d_{c,i}$ of the class c to user i. The parameter β controls the heterogeneity of the distributions at each user, where $\beta \to \infty$ results in IID setting.

We implement a VGG-11 (Simonyan & Zisserman, 2014), which is sufficient for our needs, as our goal is to evaluate various schemes, not to achieve the best accuracy. The hyperparameters are provided in Appendix F.

Modeling dropouts. At each round, user $i \in [N]$ drops from the protocol with probability p_i. p_i is selected from $\{0.1, 0.2, 0.3, 0.4, 0.5\}$ uniformly at random.

Implemented Schemes. For the benchmarks, we implement the three baselines introduced in Sec. 3.3,

referred to as *Random*, *Weighted Random*, and *Partition*. For Multi-RoundSecAgg, we construct three privacy-preserving families of sets with different target multi-round privacy guarantees, $T = 6$, $T = 4$, and $T = 3$ which we refer to as Multi-RoundSecAgg ($T = 6$), Multi-RoundSecAgg ($T = 4$), and Multi-RoundSecAgg ($T = 3$), respectively. One can view the Random and Partition as extreme cases of Multi-RoundSecAgg with $T = 1$ and $T = K$, respectively. Table 1 summarizes the family size R defined in (8).

| Scheme	Family size ($= R$)
Random selection	$\sim 10^{16}$
Weighted random selection	$\sim 10^{16}$
User partition	10
Multi-RoundSecAgg, T=6	190
Multi-RoundSecAgg, T=4	4060
Multi-RoundSecAgg, T=3	91389

Table 1: Family size with $N = 120$, $K = 12$.

Key Metrics. To numerically demonstrate the performance of the six schemes in terms of the key metrics defined in Sec. 3.2, at each round, we measure the following metrics.

- For the multi-round privacy guarantee, we measure the number of models in the partial sum that the server can reconstruct, which is given by $T^{(t)} := \min_{z \in \mathbb{R}^J} \|z^\top \mathbf{P}^{(t)}\|_0$, s.t. $\mathbf{P}^{(t)^\top} z \neq \mathbf{0}$. This corresponds to the weaker privacy definition of Remark 1. We use this weaker privacy definition as the random selection and the random weighted selection strategies provide the worst privacy guarantee even with this weaker definition, as demonstrated later. On the other hand, Multi-RoundSecAgg provides better privacy guarantees with both the strong and the weaker definitions.

- For the aggregation fairness gap, we measure the instantaneous fairness gap, $F^{(t)} := \max_{i \in [N]} F_i^{(t)} - \min_{i \in [N]} F_i^{(t)}$ where $F_i^{(t)} = \frac{1}{t+1} \sum_{l=0}^{t} \mathbb{1}\{\{\boldsymbol{p}^{(l)}\}_i = 1\}$.

- We measure the instantaneous aggregation cardinality as $C^{(t)} := \frac{1}{t+1} \sum_{l=0}^{t} \|\boldsymbol{p}^{(l)}\|_0$.

We demonstrate these key metrics in Figure 3. We make the following key observations.

- Multi-RoundSecAgg achieves better multi-round privacy guarantee than both the random selection and random weighted selection strategies, while user partitioning achieves the best multi-round privacy guarantee, $T = K = 12$. However, the partitioning strategy has the worst aggregation cardinality, which results in the lowest convergence rate as demonstrated later.

- Figure 4 demonstrates the trade-off between the multi-round privacy guarantee T and the average aggregation cardinality C. Interestingly, Multi-RoundSecAgg when $T = 3$ or $T = 4$ achieves better multi-round privacy guarantee than both the random selection and the weighted random selection strategies while achieving

(a) Multi-round privacy guarantee.   (b) Aggregation fairness gap.   (c) Average aggregation cardinality.
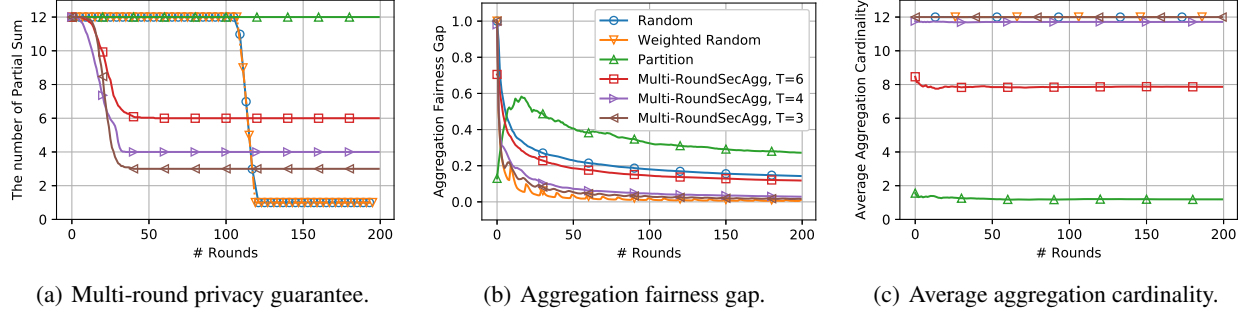
Figure 3: The key metrics with $N = 120$ (number of users), $K = 12$ (number of selected users at each round).
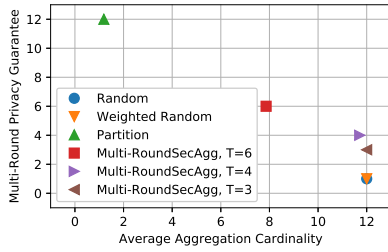
Figure 4: Trade-off between multi-round privacy and average aggregation cardinality with $N = 120$, $K = 12$.

(a) IID data distribution.   (b) Non-IID data distribution.
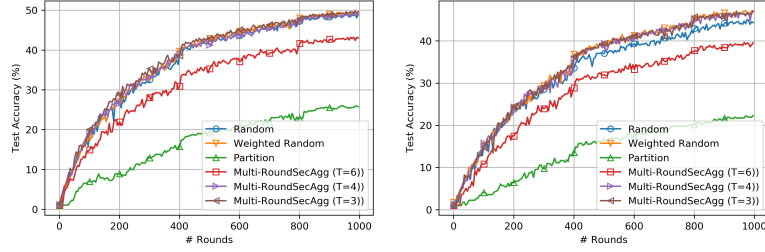
Figure 5: Training rounds versus test accuracy of VGG11 in (Simonyan & Zisserman, 2014) on the CIFAR-100 with $N = 120$ and $K = 12$.

almost the same average aggregation cardinality.

Remark 11. (Multi-round Privacy of Random and Weighted Random). The multi-round privacy guarantees of Random and Weighted Random drop sharply as shown in Fig. 3(a) as the participating matrix $\mathbf{P}^{(t)} \in \{0, 1\}^{t \times N}$ becomes full column rank with high probability when $t \geq N$, and hence the server can reconstruct the individual models by utilizing a pseudo inversion of the matrix $\mathbf{P}^{(t)}$. More precisely, Theorem 3 in Appendix G shows this *thresholding phenomenon*, where the probability that the server can reconstruct individual models after certain number of rounds converges to 1 exponentially fast.

Key Metrics versus Test Accuracy. To investigate how the key metrics affect the test accuracy, we measure the test accuracy of the six schemes in the two settings, the IID and the non-IID settings. Our results are demonstrated in Figure 5. We make the following key observations.

- In the IID setting, the Multi-RoundSecAgg schemes show test accuracies that are comparable to the random selection and random weighted selection schemes while the Multi-RoundSecAgg schemes provide higher levels of privacy. Specifically, the Multi-RoundSecAgg schemes achieve $T = 3, 4, 6$ based on the privacy-preserving family design while the random selection and random weighted selection schemes have $T = 1$, i.e., the server can learn an individual local model.
- In the non-IID setting, Multi-RoundSecAgg not only outperforms the random selection scheme but

also achieves a smaller aggregation fairness gap as demonstrated in Fig. 3(b).
- In both IID and non-IID settings, the user partitioning scheme has the worst test accuracy as its average aggregation cardinality is much smaller than the other schemes as demonstrated in Fig. 3(c).

We also implement additional experiments on MNIST and CIFAR-10 datasets in Appendix E.

7. Conclusion

Partial user participation may breach user privacy in federated learning, even if secure aggregation is employed at every training round. To address this challenge, we introduced the notion of long-term privacy, which ensures that the privacy of individual models are protected over all training rounds. We developed Multi-RoundSecAgg, a structured user selection strategy that guarantees long-term privacy while taking into account the fairness in user selection and average number of participating users, and showed that Multi-RoundSecAgg provides a trade-off between long-term privacy and average number of participating users (hence the convergence rate). Our experiments on the CIFAR-100, CIFAR-10, and MNIST datasets on both the IID and non-IID settings show that Multi-RoundSecAgg achieves comparable accuracy to the random selection strategy (which does not ensure long-term privacy), while ensuring long-term privacy guarantees.

References

Bell, J. H., Bonawitz, K. A., Gascón, A., Lepoint, T., and Raykova, M. Secure single-server aggregation with (poly) logarithmic overhead. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1253–1269, 2020.

Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., and Seth, K. Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1175–1191, 2017.

Chen, W., Horvath, S., and Richtarik, P. Optimal client sampling for federated learning. *arXiv preprint arXiv:2010.13723*, 2020.

Cho, Y. J., Gupta, S., Joshi, G., and Yağan, O. Bandit-based communication-efficient client selection strategies for federated learning. *arXiv preprint arXiv:2012.08009*, 2020a.

Cho, Y. J., Wang, J., and Joshi, G. Client selection in federated learning: Convergence analysis and power-of-choice selection strategies. *arXiv preprint arXiv:2010.01243*, 2020b.

Choi, B., Sohn, J.-y., Han, D.-J., and Moon, J. Communication-computation efficient secure aggregation for federated learning. *arXiv preprint arXiv:2012.05433*, 2020.

Elkordy, A. R. and Avestimehr, A. S. Secure aggregation with heterogeneous quantization in federated learning. *arXiv preprint arXiv:2009.14388*, 2020.

Fredrikson, M., Jha, S., and Ristenpart, T. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pp. 1322–1333, 2015.

Geiping, J., Bauermeister, H., Dröge, H., and Moeller, M. Inverting gradients–how easy is it to break privacy in federated learning? *arXiv preprint arXiv:2003.14053*, 2020.

Hsu, T.-M. H., Qi, H., and Brown, M. Measuring the effects of non-identical data distribution for federated visual classification. *arXiv preprint arXiv:1909.06335*, 2019.

Jain, V., Sah, A., and Sawhney, M. Singularity of discrete random matrices ii. *arXiv preprint arXiv:2010.06554*, 2020.

Kadhe, S., Rajaraman, N., Koyluoglu, O. O., and Ramchandran, K. Fastsecagg: Scalable secure aggregation for privacy-preserving federated learning. *arXiv preprint arXiv:2009.11248*, 2020.

Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., et al. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*, 2019.

Karimireddy, S. P., Kale, S., Mohri, M., Reddi, S., Stich, S., and Suresh, A. T. Scaffold: Stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*, pp. 5132–5143. PMLR, 2020.

Kolesnikov, A., Beyer, L., Zhai, X., Puigcerver, J., Yung, J., Gelly, S., and Houlsby, N. Big transfer (bit): General visual representation learning. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*, pp. 491–507. Springer, 2020.

Krizhevsky, A. and Hinton, G. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.

LeCun, Y., Haffner, P., Bottou, L., and Bengio, Y. Object recognition with gradient-based learning. In *Shape, contour and grouping in computer vision*, pp. 319–345. Springer, 1999.

LeCun, Y., Cortes, C., and Burges, C. MNIST handwritten digit database. *http://yann. lecun. com/exdb/mnist*, 2010.

Li, X., Huang, K., Yang, W., Wang, S., and Zhang, Z. On the convergence of fedavg on non-iid data. In *International Conference on Learning Representations*, 2019.

McMahan, H. B., Moore, E., Ramage, D., Hampson, S., and y Arcas, B. A. Communication-efficient learning of deep networks from decentralized data. In *Int. Conf. on Artificial Int. and Stat. (AISTATS)*, pp. 1273–1282, 2017.

McMahan, H. B., Ramage, D., Talwar, K., and Zhang, L. Learning differentially private recurrent language models. *Int. Conf. on Learning Representations (ICLR)*, 2018.

Nasr, M., Shokri, R., and Houmansadr, A. Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning. In *2019 IEEE symposium on security and privacy (SP)*, pp. 739–753. IEEE, 2019.

Pejó, B. and Biczók, G. Quality inference in federated learning with secure aggregation. *arXiv preprint arXiv:2007.06236*, 2020.

Ribero, M. and Vikalo, H. Communication-efficient federated learning via optimal client sampling. *arXiv preprint arXiv:2007.15197*, 2020.

Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

So, J., Ali, R. E., Güler, B., and Avestimehr, A. S. Secure aggregation for buffered asynchronous federated learning. *arXiv preprint arXiv:2110.02177*, 2021a.

So, J., Güler, B., and Avestimehr, A. S. Turbo-aggregate: Breaking the quadratic aggregation barrier in secure federated learning. *IEEE Journal on Selected Areas in Information Theory*, 2(1):479–489, 2021b.

Stich, S. U. Local sgd converges fast and communicates little. *arXiv preprint arXiv:1805.09767*, 2018.

Tan, M. and Le, Q. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pp. 6105–6114. PMLR, 2019.

Tang, M., Ning, X., Wang, Y., Wang, Y., and Chen, Y. Fedgp: Correlation-based active client selection for heterogeneous federated learning. *arXiv preprint arXiv:2103.13822*, 2021.

Tran, T. The smallest singular value of random combinatorial matrices. *arXiv preprint arXiv:2007.06318*, 2020.

Yang, C.-S., So, J., He, C., Li, S., Yu, Q., and Avestimehr, S. Lightsecagg: Rethinking secure aggregation in federated learning. *arXiv preprint arXiv:2109.14236*, 2021.

Yu, H., Yang, S., and Zhu, S. Parallel restarted sgd with faster convergence and less communication: Demystifying why model averaging works for deep learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 5693–5700, 2019.

Zhao, Y. and Sun, H. Information theoretic secure aggregation with user dropouts. *arXiv preprint arXiv:2101.07750*, 2021.

Zhao, Y., Li, M., Lai, L., Suda, N., Civin, D., and Chandra, V. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*, 2018.

Zhu, L. and Han, S. Deep leakage from gradients. In *Federated Learning*, pp. 17–31. Springer, 2020.

Organization. These appendices are organized as follows.

(A) In Appendix A, we prove Theorem 1.

(B) In Appendix B, we prove Theorem 2.

(C) In Appendix C, we show that batch partitioning is necessary to satisfy the multi-round privacy definition given in (5).

(D) In Appendix D, we provide the two components of Multi-RoundSecAgg which are Algorithm 1 and Algorithm 2.

(E) Appendix E provides additional experiments on the MNIST dataset.

(F) Appendix F provides additional details and the hyperparameters of the experiments of Section 6 and Appendix E.

(G) In Appendix G, we theoretically show that the random selection strategy discussed in Remark 2 that aims to select K available users at each round and the random selection strategy that selects the users in i.i.d fashion both have a multi-round privacy $T = 1$ with high probability. We also empirically demonstrate that the local models can be reconstructed accurately when random selection is used.

(H) Finally, in Appendix H, we consider the convergence rate of the general convex and the non-convex cases.

A. Theoretical Guarantees of Multi-RoundSecAgg: Proof of Theorem 1

In this appendix, we provide the proof of Theorem 1.

Proof. 1. First, we prove that Multi-RoundSecAgg ensures a multi-round privacy of T. We first partition the matrix $\mathbf{B}$ into $R \times T$ matrices as $\mathbf{B} = [\mathbf{B}^{(1)}, \mathbf{B}^{(2)}, \cdots, \mathbf{B}^{(N/T)}]$ and the aggregated models as $\mathbf{X} = [\mathbf{X}^{(1)\top}, \mathbf{X}^{(2)\top}, \cdots, \mathbf{X}^{(N/T)\top}]^\top$. We can then express any linear combination of the aggregated models $\mathbf{X}^\top \mathbf{B}^\top \mathbf{z}$, where $\mathbf{z} \in \mathbb{R}^R \setminus \{\mathbf{0}\}$, as follows

$$\mathbf{X}^\top \mathbf{B}^\top \mathbf{z} = \sum_{i=1}^{N/T} \mathbf{X}^{(i)\top} \mathbf{B}^{(i)\top} \mathbf{z}. \tag{11}$$

Denote the j-th column of $\mathbf{B}^{(i)}$ by $\mathbf{b}_j^{(i)}$ which is either a zero vector or all ones vector due to the batch partitioning structure. That is, $\mathbf{b}_j^{(i)} \in \{\mathbf{0}, \mathbf{1}\}$. Hence, $\mathbf{B}^{(i)\top} \mathbf{z} \in \{\mathbf{0}, a_i.\mathbf{1}\}$ for some $a_i \in \mathbb{R} \setminus \{0\}$. Therefore, we have

$$\mathbf{X}^{(i)\top} \mathbf{B}^{(i)\top} \mathbf{z} = \begin{cases} \mathbf{0} & \mathbf{B}^{(i)\top} \mathbf{z} = 0, \\ a_i \displaystyle\sum_{j=(i-1)T+1}^{iT} \mathbf{x}_j & \text{otherwise}, \end{cases} \tag{12}$$

$\forall i \in [N/T]$, which shows that Multi-RoundSecAgg achieves a multi-round privacy T.

2. Next, we prove that Multi-RoundSecAgg has an aggregation fairness gap $F = 0$.

It is clear that the total number of times user i is being selected up to time J is the same as that of user j who lies in the same batch as user i. This follows since all users in the same batch either participate together or they do not participate at all.

It suffices to show that the *expected* number of selections of user i up to time J is the same as that of user j, where user i and user j are in different batches. The main observation is that our protocol is *symmetric*. Indeed, the only randomness in the system are the user availability randomness and the set selection randomness when there are multiple user sets satisfying the requirements. We note that for any realization of random variables such that the batch of user i is selected at time t, there is a corresponding realization of random variables such that the batch of user j is selected at time t and all other selections remain exactly the same. Hence, $F_i = F_j$ for any $i \neq j$.

3. Finally, we characterize the average aggregation cardinality of Multi-RoundSecAgg. The average aggregation cardinality can be expressed as follows

$$
\begin{aligned}
C &= K \left(1 - \Pr[\text{No row of } \mathbf{B} \text{ is available}]\right) \\
&= K \left(1 - \Pr\left[\text{At least } \frac{N}{T} - \frac{K}{T} + 1 \text{ batches are not available}\right]\right) \\
&= K \left(1 - \sum_{i=N/T-K/T+1}^{N/T} \binom{N/T}{i} q^i (1-q)^{N/T-i}\right),
\end{aligned}
\tag{13}
$$

where q is the probability that a certain batch is not available, which is given by $q = 1 - (1-p)^T$.

$\square$

B. Convergence Analysis of Multi-RoundSecAgg : Proof of Theorem 2

The proof of Theorem 2 is divided into two parts. In the first part, we introduce a new sequence to represent the local updates in each user with respect to step index while we use the global round index t for $\boldsymbol{x}^{(t)}$ in (2). We carefully define the sequence and the step index, and then provide the convergence analysis of the sequence. In the second part, we bridge the newly defined sequence and $\boldsymbol{x}^{(t)}$ in (2), and provide convergence analysis of $\boldsymbol{x}^{(t)}$.

First Part (Convergence analysis of local model updates).

Let $\boldsymbol{w}_i^{(j)}$ be the local model updated by user i at the j-th step. Note that this step index is different from the global round index t in (2) as each user updates the local model by carrying out $E(\geq 1)$ local SGD steps before sending the results to the server. Let $\mathcal{I}_E$ be the set of global synchronization steps, i.e., $\mathcal{I}_E = \{nE | n = 0, 1, 2, \ldots\}$. Importantly, we define the step index j such it increases from nE to $nE + 1$ only when the server does not skip the selection, i.e., there are at least K available users at step $nE + 1$ for $n \in \{0, 1, 2, \ldots\}$. We denote by $\mathcal{H}_{nE}$ the set selected by Multi-RoundSecAgg at step index nE and from the definition, $|\mathcal{H}_{nE}| = K$ for all $n \in \{0, 1, 2, \ldots\}$. Then, the update equation can be described as

$$
\boldsymbol{v}_i^{j+1} = \boldsymbol{w}_i^j - \eta^j \nabla L_i\left(\boldsymbol{w}_i^j, \xi_i^j\right),
\tag{14}
$$

$$
\boldsymbol{w}_i^{j+1} = \begin{cases} \boldsymbol{v}_i^{j+1} & \text{if} \quad j+1 \in \mathcal{I}_E \\ \frac{1}{K} \sum_{k \in \mathcal{H}_{j+1}} \boldsymbol{v}_k^{j+1} & \text{if} \quad j+1 \notin \mathcal{I}_E \end{cases},
\tag{15}
$$

where we introduce an additional variable $\boldsymbol{v}_i^{j+1}$ to represent the immediate result of one step SGD from $\boldsymbol{w}_i^j$. We can view $\boldsymbol{w}_i^{j+1}$ as the model obtained after aggregation step (when $j+1$ is a global synchronization step). Motivated by (Stich, 2018; Li et al., 2019), we define two virtual sequences

$$
\overline{\boldsymbol{v}}^j = \frac{1}{N} \sum_{i=1}^N \boldsymbol{v}_i^j,
\tag{16}
$$

$$
\overline{\boldsymbol{w}}^j = \frac{1}{N} \sum_{i=1}^N \boldsymbol{w}_i^j.
\tag{17}
$$

We can interpret $\overline{\boldsymbol{v}}^{j+1}$ as the result of single step SGD from $\overline{\boldsymbol{w}}^j$. When $j \notin \mathcal{I}_E$, both $\overline{\boldsymbol{v}}^j$ and $\overline{\boldsymbol{w}}^j$ are not accessible. We also define $\overline{\boldsymbol{g}}^j = \frac{1}{N} \sum_{i=1}^N \nabla L_i\left(\boldsymbol{w}_i^j\right)$ and $\boldsymbol{g}^j = \frac{1}{N} \sum_{i=1}^N \nabla L_i\left(\boldsymbol{w}_i^j, \xi_i^j\right)$. Then, $\overline{\boldsymbol{v}}^{j+1} = \overline{\boldsymbol{w}}^j - \eta^j \boldsymbol{g}^j$.

Now, we state our two key lemmas.

Lemma 1 (Unbiased selection). *When $j+1 \in \mathcal{I}_E$, the following is satisfied,*

$$
\mathbb{E}_{\mathcal{H}_{j+1}}[\overline{\boldsymbol{w}}^{j+1}] = \overline{\boldsymbol{v}}^{j+1}.
\tag{18}
$$

Proof. Let $\mathcal{H}_{j+1} = \{i_1, \ldots, i_K\}$. Then, we have

$$\mathbb{E}_{\mathcal{H}_{j+1}}[\overline{\boldsymbol{w}}^{j+1}] = \frac{1}{K} E_{\mathcal{H}_{j+1}} \left[\sum_{k \in \mathcal{H}_{j+1}} \boldsymbol{v}_k^{j+1} \right] = \frac{1}{K} E_{\mathcal{H}_{j+1}} \left[\sum_{k=1}^{K} \boldsymbol{v}_{i_k}^{j+1} \right] = E_{\mathcal{H}_{j+1}}[\boldsymbol{v}_{i_k}^{j+1}]$$

$$= \sum_{k=1}^{N} \frac{1}{N} \boldsymbol{v}_k^{j+1} = \overline{\boldsymbol{v}}^{j+1} \tag{19}$$

where (19) follows as $\Pr[i_k = j] = \frac{1}{N}$ for $i \in [N]$. This is because the sampling probability of each user is identical due to the symmetry in the construction and the fact that all users have the same dropout probability. $\square$

Now, we provide the convergence analysis of the sequence $\overline{\boldsymbol{w}}^j$ defined in (17). We have,

$$\|\overline{\boldsymbol{w}}^{j+1} - \boldsymbol{w}^*\|^2 = \|\overline{\boldsymbol{w}}^{j+1} - \overline{\boldsymbol{v}}^{j+1} + \overline{\boldsymbol{v}}^{j+1} - \boldsymbol{w}^*\|^2$$

$$= \|\overline{\boldsymbol{w}}^{j+1} - \overline{\boldsymbol{v}}^{j+1}\|^2 + \|\overline{\boldsymbol{v}}^{j+1} - \boldsymbol{w}^*\|^2 + 2\left(\overline{\boldsymbol{w}}^{j+1} - \overline{\boldsymbol{v}}^{j+1}\right)^{\top} \left(\overline{\boldsymbol{v}}^{j+1} - \boldsymbol{w}^*\right). \tag{20}$$

When the expectation is taken over $\mathcal{H}_{j+1}$, the last term in (20) becomes zero due to Lemma 1. For the second term in (20), we have

$$\|\overline{\boldsymbol{v}}^{j+1} - \boldsymbol{w}^*\|^2 \leq (1 - \eta^j \mu)\|\overline{\boldsymbol{w}}^j - \boldsymbol{w}^*\|^2 + \alpha(\eta^j)^2, \tag{21}$$

where $\alpha = \frac{1}{N} \sum_{i=1}^{N} \sigma_i^2 + 6\rho\Gamma + 8(E-1)^2 G^2$ and (21) directly follows from Lemma 1, 2, 3 of (Li et al., 2019). The first term in (20) becomes zero if $j + 1 \in \mathcal{I}_E$, and if $j + 1 \notin \mathcal{I}_E$, from Lemma 5 of (Li et al., 2019), it is bounded by

$$\mathbb{E}_{\mathcal{H}_{j+1}}\|\overline{\boldsymbol{w}}^{j+1} - \overline{\boldsymbol{v}}^{j+1}\|^2 \leq \beta(\eta^j)^2, \tag{22}$$

where $\beta = \frac{4(N-K)E^2 G^2}{K(N-1)}$. By combining (20) to (22), we have

$$\mathbb{E}\|\overline{\boldsymbol{w}}^{j+1} - \boldsymbol{w}^*\|^2 \leq (1 - \eta^j \mu)\|\overline{\boldsymbol{w}}^j - \boldsymbol{w}^*\|^2 + (\alpha + \beta)(\eta^j)^2. \tag{23}$$

Then by utilizing the similar induction in (Li et al., 2019), we can show that

$$\mathbb{E}\|\overline{\boldsymbol{w}}^{j+1} - \boldsymbol{w}^*\|^2 \leq \frac{1}{\gamma + t - 1} \left(\frac{4(\alpha + \beta)}{\mu^2} + \gamma \mathbb{E}\|\overline{\boldsymbol{w}}^0 - \boldsymbol{w}^*\|^2 \right), \tag{24}$$

where $\gamma = \max\left\{ \frac{8\rho}{\mu}, E \right\}$. By combining (24) with ρ-smoothness of the global loss function in (1), we have

$$\mathbb{E}[L(\overline{\boldsymbol{w}}^I)] - L^* \leq \frac{\rho}{\gamma + I - 1} \left(\frac{2(\alpha + \beta)}{\mu^2} + \frac{\gamma}{2} \mathbb{E}\|\overline{\boldsymbol{w}}^0 - \boldsymbol{x}^*\|^2 \right). \tag{25}$$

Second Part (Convergence analysis of global model).

Now, we bridge the sequence $\overline{\boldsymbol{w}}^T$ and $\boldsymbol{x}^{(t)}$ in (2) to provide the convergence analysis of $\boldsymbol{x}^{(t)}$. Since we define the step index j such that j increases from nE to $nE + 1$ only when the server does not skip the selection, we have

$$\mathbb{E}[L(\boldsymbol{x}^{(J)})] = \mathbb{E}[L(\overline{\boldsymbol{w}}^{(JE\phi)})] \tag{26}$$

where ϕ is the probability that there are at least K available users at a certain synchronization step, and $\phi = \frac{C}{K}$ due to the fact that $C = K \cdot \Pr[\text{at least one row of } \mathbf{B} \text{ is available}] = K\phi$. By combining (25) and (26), we have that,

$$\mathbb{E}[L(\boldsymbol{x}^{(J)})] - L^* \leq \frac{\rho}{\gamma + \frac{C}{K}EJ - 1} \left(\frac{2(\alpha + \beta)}{\mu^2} + \frac{\gamma}{2} \mathbb{E}\|\boldsymbol{x}^{(0)} - \boldsymbol{x}^*\|^2 \right), \tag{27}$$

which completes the proof.

C. Necessity of Batch Partitioning (BP)

In this appendix, we show that batch partitioning is necessary to satisfy the multi-round privacy guarantee of Equation (5) and our strategy is optimal in the sense that no other strategy can have more distinct user selection sets than our strategy.

Proof. Consider any scheme which selects sets from an $R \times N$ matrix $\mathbf{V} = [\boldsymbol{v}_1, \cdots, \boldsymbol{v}_N]^\top$. Denote the linear coefficients multiplying them by z_i, $i \in [R]$. Then, the i-th element of $\mathbf{V}^\top \boldsymbol{z}$ is given by

$$\{\mathbf{V}^\top \boldsymbol{z}\}_i = \sum_{j \in \mathrm{supp}(\boldsymbol{v}_i)} z_i. \tag{28}$$

We now claim that we can cluster the entries using equivalence of linear functions to groups, where each group must have a size of at least T except for the group corresponding to the zero function. To show this, we choose each $z_i \overset{\text{i.i.d.}}{\sim} U[0,1]$, and the key observation is that if two entries have different linear functions then their final value after this assignment would be different with probability one. Since the scheme satisfies a multi-round privacy T, this implies that for each non-zero linear function of the form of Equation (28), there must be at least T of them. If we group the entries according to the equivalence of linear functions, we get at most N/T groups (ignoring the group of constant zero).

Then, we show that the total number of possible sets R is upper-bounded by $\binom{N/T}{K/T}$. We observe that the total number of non-zero groups we can choose for each vector is at most K/T due to the size of each group, so the total number of distinct vectors satisfying the weight requirement is at most

$$R \leq R_{\max} \overset{\text{def}}{=} \binom{D}{E}, \tag{29}$$

where $D \leq N/T$ is the total number of groups corresponding to the non-zero linear functions, and $E \leq K/T$ is the total number of groups we may select in each round. Next, we have

$$\begin{aligned} R_{\max} &= \binom{D}{E} \\ &\overset{(i)}{\leq} \binom{N/T}{E} \\ &\overset{(ii)}{\leq} \binom{N/T}{K/T} = R_{\mathrm{BP}}, \end{aligned} \tag{30}$$

where (i) follows since $\binom{D}{E}$ is monotonically increasing w.r.t D, and (ii) follows as $\binom{D}{E}$ is monotonically increasing w.r.t E if $E \leq D/2$. □

D. The Two Components of Multi-RoundSecAgg : Algorithms 1 and 2

Algorithm 1 Batch Partitioning Privacy-preserving Family Generation

---

Input: Number of users N, row weight K and the desired multi-round privacy guarantee T.
Output: Privacy-preserving Family $\mathbf{B} \in \{0,1\}^{R_{\mathrm{BP}} \times N}$, where $R_{\mathrm{BP}} = \binom{N/T}{K/T}$
Initialization: $\mathbf{B} = \mathbf{0}_{R_{\mathrm{BP}} \times N}$.

  1: Partition index sets $\{1, 2, \ldots, N\}$ into $\frac{N}{T}$ sets, $\mathcal{G}_1, \ldots, \mathcal{G}_{\frac{N}{T}}$, where $|\mathcal{G}_i| = T$ for all $i \in [\frac{N}{T}]$.
  2: Generate all possible sets each of which is union of $\frac{K}{T}$ sets out of $\frac{N}{T}$ sets $(\mathcal{G}_1, \ldots, \mathcal{G}_{\frac{N}{T}})$ without replacement. Denote the generated sets by $\mathcal{L}_1, \ldots, \mathcal{L}_{R_{\mathrm{BP}}}$.
  3: **for** $i = 1, 2, \ldots, R_{\mathrm{BP}}$ **do**
  4:     **for** $j = 1, 2, \ldots, N$ **do**
 5:   **if** $j \in \mathcal{L}_i$ **then** $\{\boldsymbol{b}_i\}_j = 1$

Algorithm 2 Available Batch Selection

Input: A family of sets $\mathbf{B}$, set of available users $\mathcal{U}^{(t)}$, the frequency of participation vector $\boldsymbol{f}^{(t-1)}$, and the selection mode λ.
$\triangleright \lambda = 0$ when $p_i = p, \forall i \in [N]$ and 1 otherwise
Output: A participation vector $\boldsymbol{p}^{(t)}$.
Initialization: $\mathbf{B}^{(t)} = [\,]$, $\ell_{\min}^{(t-1)} := \arg\min_{i \in \mathcal{U}^{(t)}} f_i^{(t-1)}$.

1: **for** $i = 1, 2, \ldots, R_{\mathrm{BP}}$ **do**
2:     **if** $\mathrm{supp}(\boldsymbol{b}_i) \subseteq \mathcal{U}^{(t)}$ **then** $\mathbf{B}^{(t)} = [\mathbf{B}^{(\mathbf{t})\top}, \boldsymbol{b}_i]^\top$.
3: **if** $\mathbf{B}^{(t)} = [\,]$ **then**
4:     $\boldsymbol{b}_{r(t)}^{(t)} = \mathbf{0}$.
5: **else if** $\lambda = 0$ **then** $\triangleright$ Uniform selection
6:     Select a row from $\mathbf{B}^{(t)}$, $\boldsymbol{b}_{r(t)}^{(t)}$, uniformly at random.
7: **else** $\triangleright$ Fairness-aware selection
8:     Select a row from $\mathbf{B}^{(t)}$, $\boldsymbol{b}_{r(t)}^{(t)}$, uniformly at random from the rows that include $\ell_{\min}^{(t-1)}$.
9: $\boldsymbol{p}^{(t)} = \boldsymbol{b}_{r(t)}^{(t)}$.
10: Update $\boldsymbol{f}^{(t)} = \boldsymbol{f}^{(t-1)} + \boldsymbol{p}^{(t)}$

E. Additional Experiments

MNIST. To further investigate the performance of Multi-RoundSecAgg, we implement a simple CNN (McMahan et al., 2017) with two 5×5 convolution layers, a fully connected layer with ReLU activation, and a final Softmax output layer. This standard model has 1,663,370 parameters and is sufficient for our needs, as our goal is to evaluate various schemes, not to achieve the best accuracy. We study the two settings for partitioning the MNIST dataset across the users.

- **IID Setting.** In this setting, the 60000 training samples are shuffled and partitioned uniformly across the $N = 120$ users, where each user receives 500 samples.

- **Non-IID Setting.** In this setting, we first sort the dataset by the digit labels, partition the sorted dataset into 120 shards of size 500, and assign each of the 120 users one shard. This is similar to the pathological non-IID partitioning setup proposed in (McMahan et al., 2017), where our partition is an extreme case as each user has only one digit label while each user in (McMahan et al., 2017) has two.

(a) IID data distribution.   (b) Non-IID data distribution.
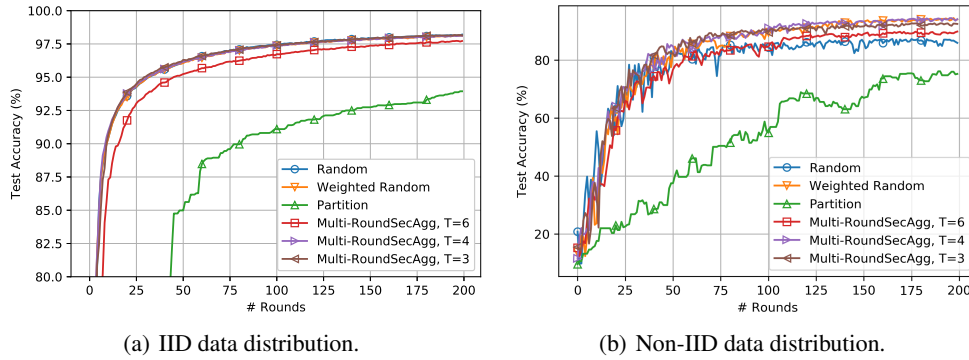
Figure 6: Training rounds versus test accuracy of CNN in (McMahan et al., 2017) on the MNIST with $N = 120$ and $K = 12$.

CIFAR-10 We also consider both IID and Non-IID distribution, and implement LeNet (LeCun et al., 1999) for both setting. While the state-of-the-art models (Kolesnikov et al., 2020; Tan & Le, 2019) achieve 99% accuracy, LeNet is sufficient for our needs, as our goal is to evaluate various schemes, not to achieve the best accuracy.

- **IID Setting.** In this setting, the 50000 training samples are shuffled and partitioned uniformly across the $N = 120$ users, where each user receives 417 or 416 samples.

Table 2: Test accuracy of VGG11 in (Simonyan & Zisserman, 2014) on the CIFAR-100 dataset with $N = 120$ and $K = 12$.

Scheme	IID Setting	Non-IID Setting
Random selection	49.15%	44.32%
Weighted random selection	50.06%	47.11%
User partition	25.73%	22.32%
Multi-RoundSecAgg, T=6	42.89%	39.57%
Multi-RoundSecAgg, T=4	49.43%	46.99%
Multi-RoundSecAgg, T=3	50.22%	47.06%

- **Non-IID dataset.** In this setting, we utilize the *data-sharing strategy* of (Zhao et al., 2018), where the 50000 training samples are divided into a globally shared dataset $\mathcal{G}$ and private dataset $\mathcal{D}$. We set $|\mathcal{G}| = 200$ and $|\mathcal{D}| = 49800$. Then, we sort $\mathcal{D}$ by the labels, partition it into 120 shards of size 415, and assign each of the 120 users one shard. Each user has 200 samples of globally shared data and 415 samples of private dataset with one label.

(a) IID data distribution. (b) Non-IID data distribution.
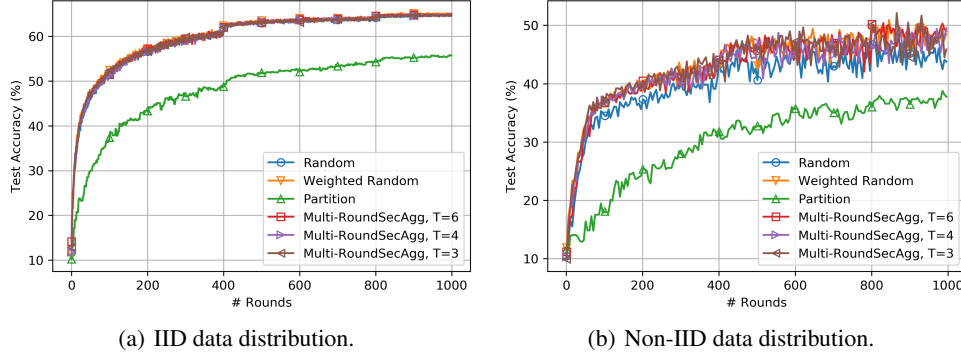
Figure 7: Training rounds versus test accuracy of LeNet in (LeCun et al., 1999) on the CIFAR-10 with $N = 120$ and $K = 12$.

We measure the test accuracy of the six schemes on the MNIST and CIFAR-10 dataset with the two distribution settings, the IID and the Non-IID. Our results are demonstrated in Figure 6 and Figure 7. We make the following key observations, which are similar to the observations on the CIFAR-100 dataset.

- In the IID setting, the Multi-RoundSecAgg schemes show comparable test accuracy to the random selection and random weighted selection schemes while the Multi-RoundSecAgg schemes provide better multi-round privacy guarantee T.

- In the non-IID setting, the Multi-RoundSecAgg schemes outperform the random selection scheme while showing comparable test accuracy to the weighted random selection scheme. This is because Multi-RoundSecAgg schemes have better aggregation fairness gaps as demonstrated in Figure 3(b), which results in better test accuracy in the non-IID setting.

- In both IID and non-IID settings, the user partitioning scheme has the worst test accuracy as its average aggregation cardinality is much smaller than the other schemes.

F. Experiment Details

In this section, we provide more details about the experiments of Section 6 and Appendix E.

We summarize the test accuracy of CIFAR-100, CIFAR-10, and MNIST dataset in Table 2, Table 3 and Table 4, respectively. For all datasets, we run experiments five times with different random seeds and present the average value of the test accuracy in Table 3 and Table 4.

Hyperparameters and computing resources. For a fair comparison between 6 schemes, we find the best learning rate from $\{0.1, 0.03, 0.01, 0.003, 0.001, 0.0003, 0.0001\}$. Given the choice of the best learning rate η, η is decayed to 0.4η every 400 and 800 rounds to train the LeNet on the CIFAR-10 dataset or train VGG11 on the CIFAR-100 dataset while η is not

Table 3: Test accuracy of LeNet in (LeCun et al., 1999) on the CIFAR-10 dataset with $N = 120$ and $K = 12$.

Scheme	IID Setting	Non-IID Setting
Random selection	64.64%	45.20%
Weighted random selection	65.06%	47.89%
User partition	55.70%	37.74%
Multi-RoundSecAgg, T=6	65.01%	46.35%
Multi-RoundSecAgg, T=4	64.95%	47.00%
Multi-RoundSecAgg, T=3	64.80%	47.21%

Table 4: Test accuracy of the CNN in (McMahan et al., 2017) on the MNIST dataset with $N = 120$ and $K = 12$.

Scheme	IID Setting	Non-IID Setting
Random selection	98.21%	85.79%
Weighted random selection	98.10%	94.04%
User partition	93.94%	75.26%
Multi-RoundSecAgg, T=6	97.72%	89.88%
Multi-RoundSecAgg, T=4	98.11%	92.51%
Multi-RoundSecAgg, T=3	98.15%	94.16%

decayed in the CNN on the MNIST dataset. To train the LeNet on the CIFAR-10 dataset or train VGG11 on the CIFAR-100 dataset, we use the mini-batch size of 50 and $E = 1$ local epoch for both IID and Non-IID settings. To train the CNN on the MNIST dataset, we use the mini-batch size of 100 and $E = 1$ local epoch for both IID and Non-IID settings. All experiments are conducted with users equipped with 3.4 GHz 4 cores i-7 Intel CPU and NVIDIA Geforce 1080, and the users communicate amongst each other through Ethernet to transfer the model parameters.

G. Multi-round Privacy Analysis of the Conventional Random User Selection Strategies

In this appendix, we first theoretically study the multi-round privacy of two random user selection strategies, and show that they have a very weak multi-round privacy of $T = 1$ with high probability (for the case where $p_i = p, \forall i \in [N]$). Furthermore, we also provide additional experiments showing that the server can reconstruct the local updates of all users with high accuracy when a random selection strategy is used. In the theoretical analysis, to simplify the problem, we assume that the model of the users have converged and don't change from one round to the next. However, in the experiments, we empirically evaluate the error in approximating the individual models of the users (via least-squares error estimation), and show that the server can approximate individual updates with very small error.

G.1. Theoretical Analysis of the Random Selection Strategies

We start by our theoretical results, where we consider the following two random selection schemes.

1. K-**uniform Random Selection**. In this scheme, at round t, K users are selected uniformly at random from the set of available users $\mathcal{U}^{(t)}$ if $|\mathcal{U}^{(t)}| \geq K$. Otherwise, the server skips this round.

2. **I.I.D Random Selection**. In this scheme, at round t, each user is selected with probability $\frac{K}{N(1-p)}$ independently from the other available users, where $K < N(1-p)$. Hence, the expected number of selected users at each round is K user.

For both schemes, we show that the server can reconstruct all individual models after N rounds in the worst-case scenario (assuming that the models do not change over N rounds). Specifically, we show that the participation matrices in both schemes have full rank with high probability after N rounds. This, in turn, implies that the server can reconstruct all local models after N rounds with high probability in both schemes. We provide our results formally next in Theorem 3.

Theorem 3. *(Random selection schemes have a multi-round privacy guarantee $T = 1$).*

1. *Consider the K-uniform random selection scheme, where $\min(K, N - K) \geq cN$. In this scheme, the server can*

reconstruct all individual models of the N users after N rounds with probability at least

$$1 - 2e^{-c'N}, \tag{31}$$

for some constant $c' > 0$ that depends on c.

2. *Consider the i.i.d random selection scheme, where the users are selected according to $Bern(\frac{K}{N(1-p)})$ distribution and let $t = K/N$. In this scheme, the server can reconstruct the individual models of the N users after N rounds with probability at least*

$$1 - 2N(1-t)^N - (1 + o_N(1))N(N-1)(t^2 + (1-t)^2)^N, \tag{32}$$

which converges to 1 exponentially fast if $t \in (0, 1/2)$ is a fixed constant.

Proof. We first note that if the participation matrix has full rank after N rounds, then the server can reconstruct the model of each individual user. Hence, we analyze the probability of the $N \times N$ participation matrix being full rank. We now consider each scheme separately.

1. In the K-uniform random selection scheme, the probability that the participation matrix after N rounds $\mathbf{P}^{(N)}$ has full rank is lower-bounded as follows (Tran, 2020), when $\min(K, N - K) \geq cN$,

$$\Pr[\mathbf{P}^{(N)} \text{ has full rank}] \geq 1 - 2e^{-c'N},$$

for some constant $c' > 0$ that depends on c. Hence, it follows that the server can reconstruct all individual models with probability at least $1 - 2e^{-c'N}$.

2. In the i.i.d random selection scheme, the probability that the participation matrix after N rounds $\mathbf{P}^{(N)}$ has full rank is lower-bounded as follows (Jain et al., 2020)

$$\Pr[\mathbf{P}^{(N)} \text{ has full rank}] \geq 1 - 2N(1-t)^N - (1 + o_N(1))N(N-1)(t^2 + (1-t)^2)^N,$$

which converges to 1 exponentially fast if $t = K/N \in (0, 1/2)$ is a fixed constant. Hence, it follows that the probability the server can reconstruct all individual models is lower-bounded by the same probability.

$\square$

Remark 12. Our experimental results in Section 6 also show that the multi-round privacy guarantee of the K-uniform random selection scheme goes to 1 after almost N rounds as shown in Fig. 3(a).

G.2. Experimental Results

We now empirically evaluate the error in approximating the individual gradients of the users (via least-squares error estimation), and show that the server can approximate individual gradients of all users with a very small error when K-uniform random selection is used. To do so, we implement a reconstruction algorithm utilizing the least-squares method, and measure the L_2 distance between the true gradients and reconstructed gradients. We consider a FL setting with $N = 40$ users, where the server aims to choose $K = 8$ users at every round, to train the LeNet in (LeCun et al., 1999) on the CIFAR-10 dataset with Non-IID setting, which is the same as the setting in Appendix E.
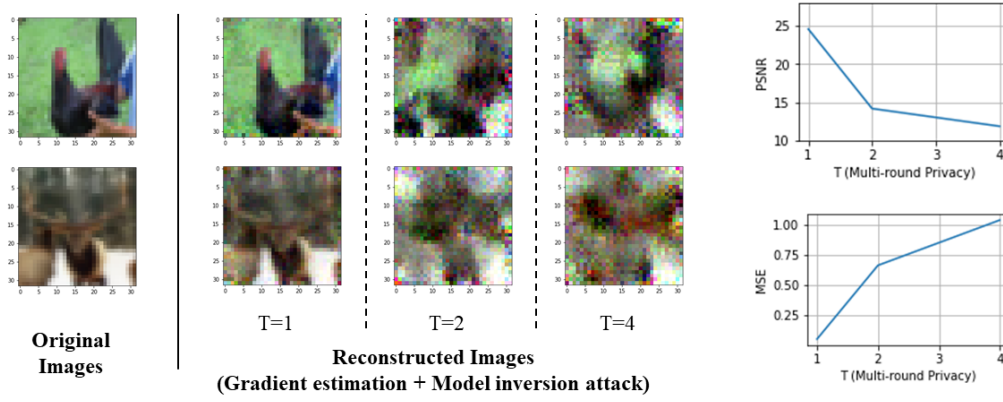
Let $\delta_i^{(t)}$ be the gradient of user i at round t, i.e., $\delta_i^{(t)} = x_i^{(t)} - x^{(t)}$, and $\delta^{(t)}$ be the global update at round t, i.e., $\delta^{(t)} = x^{(t+1)} - x^{(t)} = \mathbf{\Delta}^{(t)\top}_{\text{individual}} p^{(t)}$ where $\mathbf{\Delta}^{(t)}_{\text{individual}} = \left[w_1 \delta_1^{(t)}, \ldots, w_N \delta_N^{(t)} \right]^\top \in \mathbb{R}^{N \times d}$. After a sufficiently large number of rounds t_0, the global model at the server converges and does not change much across the rounds, which results in that local updates also do not change much across the rounds. Then, we have

$$\mathbf{\Delta}^{(t_0;t_1)}_{\text{global}} = \mathbf{P}^{(t_0;t_1)} \mathbf{\Delta}^{(t_0)}_{\text{individual}} + \mathbf{Z}, \tag{33}$$

where $\mathbf{\Delta}^{(t_0;t_1)}_{\text{global}}$ denotes the concatenate of the global updates from round t_0 to round $t_1 - 1$, i.e., $\mathbf{\Delta}^{(t_0;t_1)}_{\text{global}} = \left[\delta^{(t_0)}, \ldots \delta^{(t_1-1)} \right]^\top \in \mathbb{R}^{(t_1-t_0) \times d}$ for $t_1 > t_0$, $\mathbf{P}^{(t_0;t_1)} \in \{0, 1\}^{(t_1-t_0) \times N}$ is the participation matrix from round t_0 to round $t_1 - 1$, and $\mathbf{Z}$ denotes the perturbation (or noise) incurred by the local updates across the rounds.

(a) $K(=8)$-uniform random selection ($T = 1$).

(b) Multi-RoundSecAgg ($T = 2$).
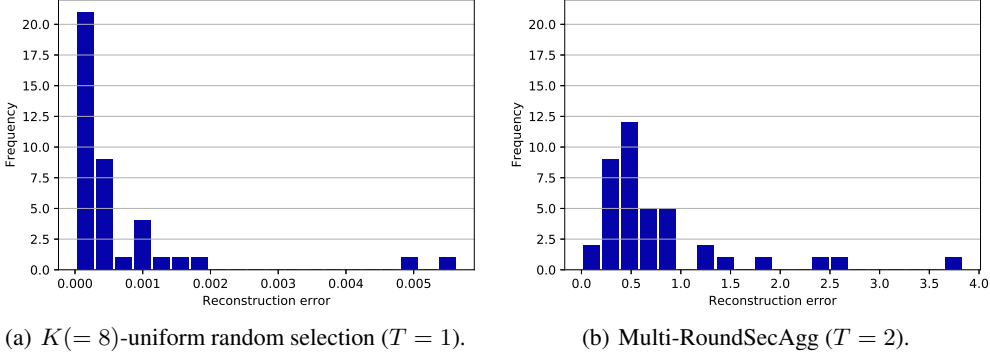
Figure 8: Histogram of the reconstruction error defined in (35) when the $K(=8)$-uniform random selection or Multi-RoundSecAgg ($T = 2$) scheme is used to train the LeNet on the CIFAR-10 dataset. The average reconstruction errors of $K(=8)$-uniform random selection and Multi-RoundSecAgg ($T = 2$) are 6.715×10^{-3} and 0.7829, respectively, which implies that the server can reconstruct all local updates when $K(=8)$-uniform random selection is used while the server cannot reconstruct the local updates when Multi-RoundSecAgg ($T = 2$) is used.

Original Images

T=1   T=2   T=4

Reconstructed Images
(Gradient estimation + Model inversion attack)

Figure 9: Comparison of the reconstructed images using the model inversion attack (Geiping et al., 2020) with different value of multi-round privacy guarantee T (left) and measurement of similarity between the reconstructed images and the original images, where PSNR $= \infty$ and MSE $= 0$ for two identical images (right).

The server can then estimate $\boldsymbol{\Delta}_{\text{individual}}^{(t_0)}$ by utilizing the least-squares method as follows

$$\hat{\boldsymbol{\Delta}}_{\text{individual}}^{(t_0)} = \left(\mathbf{P}^{(t_0;t_1)\top} \mathbf{P}^{(t_0;t_1)} \right)^{-1} \mathbf{P}^{(t_0;t_1)\top} \boldsymbol{\Delta}_{\text{global}}^{(t_0;t_1)}, \tag{34}$$

and we measure the reconstruction error as follows

$$e_i^{(t_0)} = \frac{\|\delta_i^{(t_0)} - \hat{\delta}_i^{(t_0)}\|_2^2}{\|\delta_i^{(t_0)}\|_2^2}, \tag{35}$$

where $\hat{\delta}_i^{(t_0)}$ denotes the reconstructed gradient of user i, which corresponds to i-th row of $\hat{\boldsymbol{\Delta}}_{\text{individual}}^{(t_0)}$ in (34). On the other hand, in Multi-RoundSecAgg with multi-round privacy guarantee $T = 2$, the server cannot estimate the individual gradients by utilizing (34) because $\mathbf{P}^{(t_0;t_1)}$ is not full rank hence the inverse of $\mathbf{P}^{(t_0;t_1)\top} \mathbf{P}^{(t_0;t_1)}$ does not exist. The best that the server can do is to estimate $\sum_{i \in \mathcal{G}_j} \delta_i^{(t_0)}$, where $\mathcal{G}_j$ is the index set of the users in the j-th batch. The server can then estimate $\delta_i^{(t_0)}$ by dividing the estimate of $\sum_{i \in \mathcal{G}_j} \delta_i^{(t_0)}$ by T, where $i \in \mathcal{G}_j$.

Figure 8(a) and Figure 8(b) show the histogram of the reconstruction error of the individual gradients when the K-uniform random selection scheme and Multi-RoundSecAgg ($T = 2$) scheme are used, respectively. We set $t_0 = 1460$ and $t_1 = 1500$ in this experiment. We observe that the K-uniform random selection scheme has much smaller average reconstruction error

$\frac{1}{N}\sum_{i=1}^{N} e_i^{(t_0)} = 6.715 \times 10^{-3}$ than the average reconstruction error of Multi-RoundSecAgg ($T = 2$), which implies that the server can reconstruct all local gradients as the K-uniform random selection scheme has a multi-round privacy guarantee $T = 1$.

Finally, the server can reconstruct the training images by applying model inversion attack (Geiping et al., 2020) to the reconstructed gradient $\hat{\delta}_i^{(t_0)}$. Figure 9 the reconstructed images of random selection scheme ($T = 1$) and Multi-RoundSecAgg ($T = 2, 4$). We measure the reconstruction performance using peak signal-to-noise ratio (PSNR) and mean square error (MSE). Large PSNR and small MSE indicate more similarity between the reconstructed and original images, and hence we can observe that random selection scheme ($T = 1$) leaks much more information about the original image than Multi-RoundSecAgg ($T = 2, 4$).

H. General Convex and Non-Convex Convergence Rates

In this appendix, we discuss extending the convergence proof of (Karimireddy et al., 2020) to our setting. Since we closely follow (Karimireddy et al., 2020), we mainly here focus on the differences. The main difference between the two settings is that the number of participating users in (Karimireddy et al., 2020) is fixed as $|\mathcal{S}^{(t)}| = K$ across all rounds as *dropouts* are not considered, whereas it may change in our case between the rounds based on the availability of the users. Specifically, the server in our case aims to choose K users at each round. If this is not possible due to the dropouts, the server skips this round. That is, $|\mathcal{S}^{(t)}| \in \{0, K\}$ and $\mathbb{E}[|\mathcal{S}^{(t)}|] = C$, where we recall that C is the average aggregation cardinality that depends on the desired multi-round privacy guarantee T.

Next, we recall the setting and the assumptions of (Karimireddy et al., 2020). In (Karimireddy et al., 2020), the problem is formalized as minimizing a global loss function as follows

$$\min_{\boldsymbol{x}} L(\boldsymbol{x}) \text{ s.t. } L(\boldsymbol{x}) = \frac{1}{N}\sum_{i=1}^{N} L_i(\boldsymbol{x}), \tag{36}$$

where L is bounded from below by L^*, the loss function of user i L_i is ρ-smooth and $g_i(\boldsymbol{x}) = \nabla L_i(\boldsymbol{x}, \zeta_i)$ is an unbiased stochastic gradient of L_i with variance bounded by σ^2. Furthermore, following the assumptions of (Karimireddy et al., 2020), we consider the following assumptions.

Assumption 1. (G, B)-Bounded Gradient Dissimilarity (BGD). There exists constants $G \geq 0$ and $B \geq 1$ such that

$$\frac{1}{N}\sum_{i=1}^{N}\|\nabla L_i(\boldsymbol{x})\|^2 \leq G^2 + B^2\|\nabla L_i(\boldsymbol{x})\|^2, \forall \boldsymbol{x}, \tag{37}$$

and when $\{L_i\}$ are convex, this assumption can be relaxed as

$$\frac{1}{N}\sum_{i=1}^{N}\|\nabla L_i(\boldsymbol{x})\|^2 \leq G^2 + 2\rho B^2(L(\boldsymbol{x}) - L^*), \forall \boldsymbol{x}. \tag{38}$$

Assumption 2. δ-Bounded Hessian Dissimilarity (BHD).

$$\|\nabla^2 L_i(\boldsymbol{x}) - \nabla^2 L(\boldsymbol{x})\| \leq \delta, \forall \boldsymbol{x}, \tag{39}$$

and L_i is δ-weakly.

Assumption 3. L_i is μ-convex for $\mu \geq 0$ and satisfies

$$\langle\nabla L_i(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x}\rangle \leq -\left(L_i(\boldsymbol{x}) - L_i(\boldsymbol{y}) + \frac{\mu}{2}\|\boldsymbol{x} - \boldsymbol{y}\|^2\right), \text{for any } i, \boldsymbol{x}, \boldsymbol{y}, \tag{40}$$

where μ can be 0 (general convex case).

Assumption 4. $g_i(\boldsymbol{x}) = \nabla L_i(\boldsymbol{x}; \zeta_i)$ is an unbiased stochastic gradient of L_i with bounded variance. That is, we have

$$\mathbb{E}\left[\|g_i(\boldsymbol{x}) - \nabla L_i(\boldsymbol{x})\|\right] \leq \sigma^2, \text{for any } i, \boldsymbol{x}. \tag{41}$$

Assumption 5. $\{L_i\}$ are ρ-smooth and satisfy

$$\|\nabla L_i(\boldsymbol{x}) - \nabla L_i(\boldsymbol{x})\| \leq \rho\|\boldsymbol{x} - \boldsymbol{y}\|, \text{for any } i, \boldsymbol{x}, \boldsymbol{y}. \tag{42}$$

It is also worth noting that when $\{L_i\}$ are convex and x^* this assumption implies that

$$\frac{1}{2\rho N}\sum_{i=1}^{N}\|\nabla L_i(\boldsymbol{x}) - L_i(\boldsymbol{x}^*)\| \leq L(\boldsymbol{x}) - L^*,\tag{43}$$

and when L_i is twice differentiable this assumption implies that $\|\nabla^2 L_i(\boldsymbol{x})\| \leq \rho$ for any $\boldsymbol{x}$.

We now recall the global and the local updates of the setting considered in (Karimireddy et al., 2020). $\boldsymbol{x}^{(t)}$ is the global model after round t and $\boldsymbol{x}_{i,e}^{(t)}$ is the local model of user i in round t and local step e. In round t, the server selects a subset of users $\mathcal{S}^{(t)}$. Each user then copies the global model $\boldsymbol{x}_{i,0}^{(t)} = \boldsymbol{x}^{(t-1)}$ and performs E local update steps as follows

$$\boldsymbol{x}_{i,e}^{(t)} := \boldsymbol{x}_{i,e-1}^{(t)} - \eta_l g_i(\boldsymbol{x}_{i,e-1}^{(t)}),\tag{44}$$

where η_l is the local step size. The users then send their updates and the server updates the global model as follows

$$\boldsymbol{x}^{(t)} := \boldsymbol{x}^{(t-1)} + \frac{\eta_g}{K}\sum_{i\in\mathcal{S}}(\boldsymbol{x}_{i,E}^{(t)} - \boldsymbol{x}^{(t-1)}),\tag{45}$$

where η_g is the global step size. Finally, the output is given by

$$\bar{\boldsymbol{x}}^{(J)} = \bar{\boldsymbol{x}}^{(t-1)} \text{ with probability } \frac{\theta_t}{\sum_\tau \theta_\tau}\text{for } t \in \{1, \cdots, J+1\}\tag{46}$$

for some weights $\{\theta_t\}$, where J is the total number of rounds. Next, we restate the convergence Theorem of (Karimireddy et al., 2020).

Theorem. *Suppose that $\{L_i\}$ satisfy Assumptions $1, 4$ and 5. Then for each of the following cases there exists weights $\{\theta_t\}$ and local step sizes η_l such that for any global step size $\eta_g \geq 1$, we have*

- **Strongly convex.** *If $\{L_i\}$ satisfy Assumption 3 for $\mu > 0$, $\eta_l \leq \frac{1}{8(1+B^2)\rho E\eta_g}$, $J \geq \frac{8(1+B^2)\rho}{\mu}$, then*

$$\mathbb{E}\left[L(\bar{\boldsymbol{x}}^J)\right] - L(\boldsymbol{x}^*) \leq \tilde{\mathcal{O}}\left(\frac{M^2}{\mu JEK} + \frac{\rho G^2}{\mu^2 J^2} + \mu D^2 \exp\left(-\frac{\mu}{16(1+B^2)\rho}J\right)\right).\tag{47}$$

- **General convex.** *If $\{L_i\}$ satisfy Assumption 3 for $\mu = 0$, $\eta_l \leq \frac{1}{8(1+B^2)\rho E\eta_g}$, $J \geq 1$, then*

$$\mathbb{E}\left[L(\bar{\boldsymbol{x}}^J)\right] - L(\boldsymbol{x}^*) \leq \mathcal{O}\left(\frac{MD}{\sqrt{JEK}} + \frac{D^{4/3}(\rho G^2)^{1/3}}{(J+1)^{2/3}} + \frac{B^2\rho D^2}{J}\right).\tag{48}$$

- **Non-convex.** *If $\{L_i\}$ satisfy Assumption 1 and $\eta_l \leq \frac{1}{8(1+B^2)\rho E\eta_g}$, then*

$$\mathbb{E}\left[\|\nabla L(\bar{\boldsymbol{x}}^J)\|^2\right] \leq \mathcal{O}\left(\frac{\rho M\sqrt{F}}{\sqrt{JEK}} + \frac{F^{2/3}(\rho G^2)^{1/3}}{(J+1)^{2/3}} + \frac{B^2\rho F}{J}\right).\tag{49}$$

where $M^2 = \sigma^2(1 + K/\eta_g^2) + E(1 - K/N)G^2$, $D = \|\boldsymbol{x}^0 - \boldsymbol{x}^*\|$ and $F = L(\boldsymbol{x}^0) - L^*$.

As we discussed, the main difference between our setting and the setting of (Karimireddy et al., 2020) is that the number of selected users in our case in each round is a random variable $|\mathcal{S}^{(t)}| \in \{0, K\}$ with mean equal to the average aggregation cardinality C. We now show the effect of this difference on the key lemma of (Karimireddy et al., 2020). We first recall this key lemma from (Karimireddy et al., 2020) and then derive a simple corollary that extends this lemma to our setting.

Lemma. *(Separating Mean and Variance)[Lemma 4 in (Karimireddy et al., 2020)]. Let $\{\boldsymbol{E}_1, \boldsymbol{E}_2, \cdots, \boldsymbol{E}_w\}$ be random vectors in $\mathbb{R}^d$, which may not be independent. We consider the following two cases.*

- When $\mathbb{E}[\boldsymbol{E}_i] = \boldsymbol{\zeta}_i$ and $\mathbb{E}[\|\boldsymbol{E}_i - \boldsymbol{\zeta}_i\|^2] \leq \sigma^2$, then we have

$$\mathbb{E}[\|\sum_{i=1}^{w} \boldsymbol{E}_i\|^2] \leq \|\sum_{i=1}^{w} \boldsymbol{\zeta}_i\|^2 + w^2\sigma^2. \tag{50}$$

- When $\mathbb{E}[\boldsymbol{E}_i|\boldsymbol{E}_{i-1}, \cdots, \boldsymbol{E}_1] = \boldsymbol{\zeta}_i$ and $\mathbb{E}[\|\boldsymbol{E}_i - \boldsymbol{\zeta}_i\|^2] \leq \sigma^2$, then we have

$$\mathbb{E}[\|\sum_{i=1}^{w} \boldsymbol{E}_i\|^2] \leq 2\|\sum_{i=1}^{w} \boldsymbol{\zeta}_i\|^2 + 2w\sigma^2. \tag{51}$$

In this key lemma, w is constant as the setting (Karimireddy et al., 2020) assumes the number of participating users is fixed as K in every round. That is, this lemma is applied with $w = K$. In our case, however, the number of participating users is a random variable. Hence, we consider this case in the following corollary.

Corollary 1. *Let $\{\boldsymbol{E}_1, \boldsymbol{E}_2, \cdots, \boldsymbol{E}_W\}$ be random vectors in $\mathbb{R}^d$, which may not be independent and $W \in \{0, w\}$ is a random variable that is independent of E_i and $\mathbb{E}[W] = \mu_W$. We consider the following two cases.*

- *When $\mathbb{E}[\boldsymbol{E}_i] = \boldsymbol{\zeta}_i$ and $\mathbb{E}[\|\boldsymbol{E}_i - \boldsymbol{\zeta}_i\|^2] \leq \sigma^2$, then we have*

$$\mathbb{E}[\|\sum_{i=1}^{W} \boldsymbol{E}_i\|^2] \leq \frac{\mu_W}{w}\|\sum_{i=1}^{w} \boldsymbol{\zeta}_i\|^2 + w\mu_W\sigma^2. \tag{52}$$

- *When $\mathbb{E}[\boldsymbol{E}_i|\boldsymbol{E}_{i-1}, \cdots, \boldsymbol{E}_1] = \boldsymbol{\zeta}_i$ and $\mathbb{E}[\|\boldsymbol{E}_i - \boldsymbol{\zeta}_i\|^2] \leq \sigma^2$, then we have*

$$\mathbb{E}[\|\sum_{i=1}^{W} \boldsymbol{E}_i\|^2] \leq 2\frac{\mu_W}{w}\|\sum_{i=1}^{w} \boldsymbol{\zeta}_i\|^2 + 2\mu_W\sigma^2. \tag{53}$$

This is the main difference between our setting and the setting considered in (Karimireddy et al., 2020) and the rest of the proof follows similarly. Similar to Theorem 2, we can see that the average aggregation cardinality C controls the convergence rate and hence there is a trade-off between the multi-round privacy T and the convergence rate.