

# Robust Federated Learning Over the Air: Combating Heavy-Tailed Noise With Median Anchored Clipping

Jiaxing Li<sup>†</sup>, Zihan Chen<sup>‡</sup>, Kai Fong Ernest Chong<sup>‡</sup>, Bikramjit Das<sup>‡</sup>, Tony Q. S. Quek<sup>‡</sup>, Howard H. Yang<sup>†</sup>

<sup>†</sup>ZJU-UIUC Institute, Zhejiang University, Haining 314400, China

<sup>‡</sup>Singapore University of Technology and Design, Singapore 487372

**Abstract**—Leveraging over-the-air computations for model aggregation is an effective approach to cope with the communication bottleneck in federated edge learning. By exploiting the superposition properties of multi-access channels, this approach facilitates an integrated design of communication and computation, thereby enhancing system privacy while reducing implementation costs. However, the inherent electromagnetic interference in radio channels often exhibits heavy-tailed distributions, giving rise to exceptionally strong noise in globally aggregated gradients that can significantly deteriorate the training performance. To address this issue, we propose a novel gradient clipping method, termed Median Anchored Clipping (MAC), to combat the detrimental effects of heavy-tailed noise. We also derive analytical expressions for the convergence rate of model training with analog over-the-air federated learning under MAC, which quantitatively demonstrates the effect of MAC on training performance. Extensive experimental results show that the proposed MAC algorithm effectively mitigates the impact of heavy-tailed noise, hence substantially enhancing system robustness.

**Index Terms**—Analog over-the-air computing, federated learning, gradient clipping, robustness

## I. INTRODUCTION

In lieu of conventional centralized training methods, federated learning (FL) [1]–[4] offers a fresh perspective on collaborative data methodologies, enabling clients to benefit from high-quality model services without compromising the security of their private data. Specifically, clients are not required to upload private local data to the server; rather, they only need to transmit model weights or gradient parameters, effectively averting the exposure of sensitive information. FL establishes a collaborative connection between clients and servers, facilitating the joint learning of a shared global model.

Despite the tremendous success of the FL paradigm, there still exist obvious problems while the executing process. The frequent transmission of model information between clients and server consumes substantial network bandwidth, while the aggregation of a large number of parameters requires extensive computing resources. These factors can result difficulty in accommodating massive access and meeting stringent latency requirements. Moreover, although FL avoids the direct aggregation of user data, model parameter information such as gradients can still potentially compromise user privacy. Inference attackers may exploit shared gradients to reverse-engineer training data, leading to potential indirect breaches of privacy [5].

A viable solution to this problem is by integrating over-the-air (OTA) computation [6]–[9] into the FL system. By leveraging the superposition property of multiple-access channels, OTA computation enables the automatic aggregation of clients' gradients, significantly improving channel utilization while simultaneously reducing computational overhead [10]. Furthermore, as the server receives aggregated gradients rather than individual ones from clients [11], the system's vulnerability to inference attacks is substantially mitigated.

However, the analog channel inherently introduces electromagnetic interference during the transmission [12]–[15]. While such interference enhances privacy protection, it also compromises the reliability of channel transmission, especially when it manifests as impulse interference, rendering the noise exhibiting a heavy-tailed distribution (rather than Gaussian) [16]—this has been consistently demonstrated by both theoretical [17] and empirical evidence [18]. In heavy-tailed distributions, extreme values (i.e. very large or very small values) occur with high probability, which could lead to severe signal distortion, resulting in a gradient explosion in the FL system, and thereby profoundly affecting the training process of OTA FL.

Numerous methods have been proposed to combat the impact of strong channel noise, ranging from channel inversion [19], phase correction [20], [21], to amplitude correction and energy estimation [22]. However, these methods only enhance the channel quality and fail to cope with the heavy tail phenomenon at the algorithmic level. In the field of deep learning, gradient norm clipping (GNC) has been proposed to solve the gradient explosion problem [23], and has been widely discussed as a solution to the gradient heavy-tailed distribution problem [24]–[26], but there is a crucial limitation: Once the data statistical structure of the gradient is altered by noise, GNC struggles to maintain its effectiveness.

To enhance the robustness of OTA FL against heavy-tailed noise, we introduce a novel residual clipping technique named median anchored clipping (MAC). In contrast to GNC, leveraging the stability of the median, this method constrains the magnitude of signals received after centralization, adjusts the proportional relationships among gradients, maximizes gradient retention, and mitigates the impact of heavy-tailed interference on OTA FL. Our main contributions are summarized as follows:

- We propose a novel robust gradient clipping method tailored for OTA FL systems to mitigate the impact of heavy-tailed noise present in the analog channel.
- We derive the convergence rate of the OTA FL gradient descent algorithm with MAC under non-convex conditions.
- We conducted substantial experiments in which the results show that our MAC algorithm effectively mitigates the impact of heavy-tailed noise in analog OTA FL.

## II. SYSTEM MODEL

### A. Setting

We consider the FL system depicted in Fig. 1, which consists of an edge server and  $N$  clients. Every client  $n$  possesses a local dataset  $\mathcal{D}_n$  that contains  $m_n$  data samples  $\{(\mathbf{x}_i, y_i)\}_{i=1}^{m_n}$  where  $\mathbf{x}_i \in \mathbb{R}^d, y_i \in \mathbb{R}$ . We assume that the local datasets are statistically independent from each other. The edge server orchestrates with the clients to learn a statistical model from their datasets while preserving privacy.

More precisely, the clients need to collaboratively find a vector  $\mathbf{w} \in \mathbb{R}^d$  that minimizes the following loss function:

$$f(\mathbf{w}) = \frac{1}{N} \sum_{n=1}^N f_n(\mathbf{w}) \quad (1)$$

where  $f_n(\mathbf{w})$  is the local empirical risk of agent  $n$ . The solution of (1) is commonly known as the empirical risk minimizer, denoted by

$$\mathbf{w}^* = \arg \min f(\mathbf{w}). \quad (2)$$

In this paper, we adopt the OTA FL for model training, which employs analog transmissions for intermediate parameter uploading, harnessing the superposition property of radio signals for fast gradient aggregation.

### B. Federated Model Training Over the Air

The general procedure for analog OTA FL is detailed in [27]. We briefly describe it in this part for completeness. In particular, at the  $k$ -th round of global communication, the edge server broadcasts the global parameter  $\mathbf{w}_k$  to all the clients. Then, each client  $n$  calculates its local gradient  $\nabla f_n(\mathbf{w}_k)$ , modulates this parameter onto the magnitude of a set of common waveforms that are orthogonal to each other, and simultaneously sends the resulting analog signals to the edge server. The edge server passes the received signal to a bank of matched filters, with each branch tuned to one of the waveform bases, and outputs the automatically aggregated (but distorted) gradient. Formally, the global gradient can be written as follows:

$$\mathbf{g}_k = \frac{1}{N} \sum_{n=1}^N h_{n,k} \nabla f_n(\mathbf{w}_k) + \boldsymbol{\xi}_k \quad (3)$$

in which  $h_{n,k}$  represents the channel fading of client  $n$  at the  $k$ -th global iteration, assumed to be a random variable with unit mean and finite variance, independent across the clients, and




Fig. 1: An illustration of the OTA FL training procedure.

varies over time in an i.i.d. manner;  $\boldsymbol{\xi}_k$  results from the electromagnetic interference, modeled as a  $d$ -dimensional random vector where each entry follows an independent symmetrical  $\alpha$ -stable distribution ( $S\alpha S$ ) [28] (with tail index  $\alpha$  and scale parameter  $\tau$ ), accounting for the heavy-tailed distribution of impulse noise.

Consequently, the global parameter is updated as

$$\mathbf{w}_{k+1} = \mathbf{w}_k - \eta \mathbf{g}_k, \quad (4)$$

where  $\eta$  is the learning rate. Then the global parameter will be broadcast to all clients for the next round of computations.

### C. Unstable Training Performance

Normally, the above recursion is executed multiple rounds until convergence (if it occurs), upon which all participating entities have a common model close to  $\mathbf{w}^*$ . However, the spectrum is, by nature, a shared medium, giving rise to potentially strong co-channel interference, which typically manifests as noise during training. A notable feature of analog channel noise is its heavy-tail characteristic [18], which is manifested by frequent occurrence of impulse noise, which could lead to gradient explosion. These impulse noises cause severe distortion in the statistical structure of the gradients, including the proportional relationships between different gradient entries and the distribution of entries with varying numeric magnitudes. Traditional GNC only controls the total norm of the gradient, without addressing the underlying statistical structure. As a result, GNC cannot resolve the noise problem in analog OTA FL.

To that end, the main thrust of the present paper is to develop a scheme to cope with the noise introduced by analog OTA parameter aggregation so as to stabilize the training process and improve the performance of the trained model.

## III. MEDIAN ANCHORED CLIPPING

### A. Proposed Method

As shown in Fig. 2, our MAC algorithm strengthens the robustness of OTA FL against strong communication noise by performing three key steps for aggregated gradients: 1) *centralization*, 2) *clipping*, and 3) *recovery*. The primary concept behind the MAC algorithm is to determine a datum point for a set of gradient entries and, anchoring on this point, recalibrate the magnitudes of entries.

To begin with, we define the vector-median as follows:




Fig. 2: MAC algorithm—illustrated with a multi-dimensional gradient displayed in a three-dimensional view—operates as follows: For an aggregated gradient with multiple entries, the vector-median value is subtracted from each entry, clipping is applied, and the gradient is then restored by adding the vector-median value back.

**Definition 1:** For a vector  $\mathbf{w} \in \mathbb{R}^d$ ,  $\text{med}(\mathbf{w})$  is the median entry of all entries of  $\mathbf{w}$ , i.e., given  $\mathbf{w} = (w_1, w_2, \dots, w_d)^\top$ ,

$$\text{med}(\mathbf{w}) = \text{median}\{w_i, i \in [d]\} \quad (5)$$

where  $[d]$  stands for the set  $\{1, \dots, d\}$ .

1) *Centralization*: Given a globally aggregated gradient  $\mathbf{g}_k$ , we centralize it by subtracting the vector-median from each entry, namely,

$$\mathbf{g}_k \leftarrow \mathbf{g}_k - \text{med}(\mathbf{g}_k) \cdot \mathbf{1} \quad (6)$$

where  $\mathbf{1}$  represents an all-ones vector.

The rationale behind centralizing the global gradient at the median is that this operation minimizes the  $\ell_1$  deviation of the entries (note that due to heavy-tailed noise, the  $\ell_2$  deviation of the entries may be unbounded). As such, during the subsequent clipping procedure, it preserves the original information of entries that are close to the median while eliminates the extreme values introduced by the impulse noise.

2) *Clipping*: Based on the centralized gradient  $\mathbf{g}_k$ , we implement value clipping on each entry, thereby restricting the range of individual entries within a designated threshold  $C$ . More concretely, for a generic entry  $g_{k,i}$ ,  $i \in [d]$ , we have

$$g_{k,i} \leftarrow \text{sgn}(g_{k,i}) \cdot \min(|g_{k,i}|, C) \quad (7)$$

where  $\text{sgn}(\cdot)$  takes the sign of its input variable.

3) *Recovery*: After clipping the centralized gradient, we add back the median to each entry as follows:

$$\tilde{\mathbf{g}}_k \leftarrow \mathbf{g}_k + \text{med}(\mathbf{g}_k) \cdot \mathbf{1}. \quad (8)$$

Ultimately, by employing the MAC algorithm, we obtain a new aggregated global gradient with the detrimental effects of heavy-tailed noise effectively mitigated while retaining the useful information as much as possible. The details of this method are summarized in Algorithm 1.

### B. Convergence Analysis

To facilitate the analysis, we make the following assumptions, which are widely adopted in machine learning research.

### Algorithm 1: OTA FL with MAC algorithm

```

1 Initialize  $\mathbf{w}_0$ 
2 for  $k \in [K]$  (communication round) do
3   for each client  $n \in [N]$  in parallel do
4     Update local model:  $\mathbf{w}_{k,n} = \mathbf{w}_k$ .
5     Local training:  $\nabla f_n(\mathbf{w}_k)$ .
6     Send  $\nabla f_n(\mathbf{w}_k)$  to the server.
7   end
8   Global noisy aggregation:
9    $\mathbf{g}_k = \frac{1}{N} \sum_{n=1}^N h_{k,n} \nabla f_n(\mathbf{w}_k) + \boldsymbol{\xi}_k$ 
10  Server receive  $\mathbf{g}_k$ :
11   $\tilde{\mathbf{g}}_k = \text{MAC}(\mathbf{g}_k, C)$  // Median Anchored Clipping
12   $\mathbf{w}_{k+1} = \mathbf{w}_k - \eta \tilde{\mathbf{g}}_k$  // Server Update
13  Broadcasting  $\mathbf{w}_{k+1}$  to clients.
14 end

15 Function  $\text{MAC}(\mathbf{g}, C)$ :
   Input: Gradient  $\mathbf{g}$ , threshold  $C$ 
   Output: Clipped gradient  $\mathbf{g}$ 
16   $\mathbf{m} = \text{med}(\mathbf{g}) \cdot \mathbf{1}$ 
17   $\mathbf{g} = \mathbf{g} - \mathbf{m}$  // Centralization
18  for  $g_i$  in  $\mathbf{g}$  do
19     $|g_i| = \text{sgn}(g_i) \cdot \min\{|g_i|, C\}$  // Clipping
20  end
21   $\mathbf{g} = \mathbf{g} + \mathbf{m}$  // Recovery
22  return clipped gradient  $\mathbf{g}$ 

```

**Assumption 1:** The objective function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  is lower bounded by a constant  $f(\mathbf{w}^*)$ , i.e., for any  $\mathbf{w} \in \mathbb{R}^d$ , it is satisfied:

$$f(\mathbf{w}) \geq f(\mathbf{w}^*) \quad (9)$$

**Assumption 2:** The objective function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  is  $L$ -smooth, i.e., for any  $\mathbf{w}, \mathbf{v} \in \mathbb{R}^d$ , it is satisfied:

$$f(\mathbf{w}) \leq f(\mathbf{v}) + \langle \nabla f(\mathbf{v}), \mathbf{w} - \mathbf{v} \rangle + \frac{L}{2} \|\mathbf{w} - \mathbf{v}\|^2. \quad (10)$$

**Assumption 3:** The gradients of each client are bounded, i.e., for  $\forall n \in [N]$ , there exists a constant  $G$  that

$$\|\nabla f_n(\mathbf{w})\| \leq G. \quad (11)$$




Fig. 3: Performance comparison between MAC and the baselines, (a) and (d) training ResNet-18 on the CIFAR-10 dataset, (b) and (e) ResNet-34 on the CIFAR-100 dataset, and (c) and (f) training CNN on the FEMNIST dataset.

At this stage, we are ready to present the main theoretical result of this paper as follows.

**Lemma 1:** For any gradient entry, let  $p_C$  denote the probability it remains unclipped at threshold  $C$ , there is

$$p_C \sim 1 - \frac{\tau^\alpha}{C^\alpha}. \quad (12)$$

*Proof:* Please refer to Appendix B.  $\square$

**Theorem 1:** If the learning rate is set as  $\eta \leq \frac{2}{L}$ , and clipping threshold  $C > \sqrt{2}G$ , then Algorithm 1 converges as

$$\begin{aligned} \frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E} [\|\nabla f(\mathbf{w}_k)\|^2] &\leq \frac{2(f(\mathbf{w}_0) - f(\mathbf{w}^*))}{K p_C (2 - \eta L) \eta} \\ &+ \frac{1}{2} \eta^2 dL \left( p_C \left( \frac{\sqrt{2}}{2} C - G \right)^2 + (1 - p_C) C^2 \right) \end{aligned} \quad (13)$$

*Proof:* Please refer to Appendix C.  $\square$

**Remark 1:** The result in (13) demonstrates that regardless of the tail index  $\alpha$  and scale parameter  $\tau$ , running OTA FL in conjunction with MAC consistently achieves a sublinear convergence rate (where the residual error can be reduced by decreasing the learning rate). As such, the robustness of model training is substantially enhanced, making it resilient to the detrimental effects of heavy-tailed communication noise.

**Remark 2:** The effects of noise characteristics (including tail index and scale) and clipping threshold are quantified by  $p_C$ , which is determined within a probabilistic range. As (12) shows, these factors jointly influence the algorithm's convergence rate.

## IV. EXPERIMENTAL RESULTS

### A. Experiment Setup

**System Setting:** We evaluated the effectiveness of our proposed MAC algorithm by utilizing the GNC as a baseline and conducted a comparison of their training performance under identical configurations of OTA FL systems. Additionally, during the experiment, we performed threshold searches on both the MAC and GNC algorithms and conducted tuning to enhance performance, thereby ensuring the fairness of the comparison. In our experiment, Rayleigh fading with a parameter setting of  $\mu = 1$  is utilized to model channel fading. Unless otherwise stated, the following parameters will be used: number of clients  $N = 50$ , learning rate  $\eta = 0.03$ , local epoch  $E = 5$ ,  $S\alpha S$  tail index  $\alpha = 1.5$ ,  $S\alpha S$  scale parameter  $\tau = 0.1$ , and the local batch size is 10.

**Dataset and Models:** We assessed performance on CIFAR-10, CIFAR-100 [29], and FEMNIST (which is processed in a non-i.i.d. manner) [30] using ResNet-18, ResNet-34 [31], and CNN architectures, respectively. As the neural networks have a multi-layer structure, the MAC algorithm is executed parallel to the gradient blocks in a layer-wise manner. In experiments involving non-independent and identically distributed (non-i.i.d.) data on the CIFAR-10 dataset, we use a Dirichlet distribution with a concentration parameter of  $Dir = 0.3$  [32], [33] to characterize data heterogeneity.

### B. Performance Evaluation

In Fig. 3, we present a comprehensive comparative analysis of the performance of OTA FL across diverse model training




Fig. 4: Sensitivity evaluation, (a)–(c) are training ResNet-18 on the CIFAR-10 dataset while non-i.i.d. sampling, respectively correspond different tail index  $\alpha$ , scale parameter  $\tau$  and clipping threshold  $C$ .

tasks. The automatically aggregated but noisy global gradient, is subjected to various post-processing methodologies: MAC, GNC, and without any post-processing (termed as noisy transmission). Furthermore, the ideal scenario is illustrated, wherein channel corruption, encompassing fading and noise, is absent, serving as the upper bound of performance. The figures presented here demonstrate that the training process is severely impacted by heavy-tailed noise, leading to significant fluctuations in the test accuracy curve, a substantial reduction in the achievable maximum accuracy, and an increase in the loss function value, as shown by the comparison between noisy and ideal transmissions. In these circumstances, numerous experiments demonstrate that GNC faces difficulties in sustaining effective performance when the statistical structure of gradients, including the proportional relationships among various gradient dimensions and the distribution of entries with differing magnitudes, is significantly disrupted by impulse noise. Conversely, the application of the MAC algorithm results in a substantial mitigation of these fluctuations, thereby markedly improving test accuracy and training stability, while fortifying the robustness of the OTA FL system.

In Fig. 4, we examine the performance of the MAC by altering system-related parameters. Moreover, distinct from Fig. 3, all experiments depicted in Fig. 4 are executed under non-i.i.d. conditions, thereby illustrating the universal applicability of our MAC algorithm across datasets characterized by varying degrees of data heterogeneity. Specifically, Fig. 4(a) and (b) illustrate a comparative evaluation of the test accuracy of OTA FL MAC across different levels of communication noise. As shown in Fig. 4(a), it is evident that decreasing the tail index from  $\alpha = 1.5$  to  $\alpha = 1.1$  results in the MAC algorithm maintaining consistent stability while achieving high test accuracy. Notably, the condition  $\alpha = 1.1$  represents a scenario with extreme volatility in the channel noise (where  $\alpha = 1$  corresponds to the Cauchy distribution characterized by undefined mean and variance), whilst MAC demonstrates remarkable robustness even under such extreme noise conditions. Fig. 4(b) examines the performance of the MAC algorithm

under conditions of different scale parameter  $\tau$ . It is important to observe that for  $\tau = 0.05, 0.1, 0.2$ , and  $0.3$ , the Signal-to-Noise Ratios (SNRs), defined as the ratio of useful signal power to channel noise (i.e.,  $\text{SNR} = \|\nabla f(\mathbf{w}_k)\|^2 / \|\boldsymbol{\xi}_k\|^2$ ), have average values of up to  $-35$  dB,  $-41$  dB,  $-47$  dB, and  $-50$  dB, respectively, after measurement. The figure clearly corroborates that our MAC algorithm exhibits considerable stability in adverse channel conditions, sustaining its performance even when the SNR reaches  $-50$  dB. In alignment with our analysis, Fig. 4(a) and (b) illustrate that the MAC algorithm substantially mitigates the direct impact of tail index and scale parameters associated with heavy-tailed noise on the convergence rate of OTA FL, thereby enhancing system stability in the extreme case. On the other hand, Fig. 4(c) presents a comparative analysis of the performance of the MAC algorithm across various clipping thresholds. The figure reveals that variations in threshold values yield analogous convergence rates, indicating that the MAC algorithm exhibits insensitivity to the clipping thresholds, demonstrating the robustness and ease of parameterization characteristic to the MAC algorithm.

## V. CONCLUSION

In this study, we introduce a novel algorithm, denoted as median anchored clipping(MAC), which markedly enhances the robustness of over-the-air federated learning(OTA FL) systems against channel noise, frequently characterized by heavy-tailed distributions. The MAC algorithm leverages the median of the aggregated global gradient entries as a datum plane and applies value clipping to truncate the extreme values induced by impulse noise, thereby effectively alleviating the detrimental effects of channel noise while largely preserving the original gradient information. The effectiveness of the MAC algorithm was validated through a convergence analysis conducted under non-convex conditions, alongside a series of empirical experiments that demonstrated the algorithm's consistency across a range of noise conditions. The MAC algorithm is effective, low-complex, and straightforward to implement, rendering it highly applicable in practical scenarios.

## VI. APPENDIX

The appendix provides a detailed account of our convergence analysis, which is organized into three subsections: *A. Prerequisites*, *B. Proof of Lemma 1*, and *C. Proof of Theorem 1*. Appendix A introduces the analytical framework and establishes the foundational concepts required for the subsequent sections. In particular, it outlines a key lemma whose proof is presented in Appendix B. Finally, Appendix C concludes with the convergence proof of the MAC algorithm under non-convex conditions.

### A. Prerequisites

First of all, we denote that

$$\nabla f(\mathbf{w}_k) = \frac{1}{N} \sum_{n=1}^N \nabla f_n(\mathbf{w}_k), \quad (14)$$

and since the channel fading among users is assumed to be independently and identically distributed (i.i.d.) with an expected value of 1, we approximate that:

$$\begin{aligned} \mathbf{g}_k &= \frac{1}{N} \sum_{i=1}^N h_{k,n} \nabla f(\mathbf{w}_k) + \boldsymbol{\xi}_k \\ &\approx \frac{1}{N} \sum_{n=1}^N \nabla f_n(\mathbf{w}_k) + \boldsymbol{\xi}_k \\ &= \nabla f(\mathbf{w}_k) + \boldsymbol{\xi}_k. \end{aligned} \quad (15)$$

In the clipping step of MAC, entries in the aggregated gradient could be divided into two parts, some entries are clipped while others are not. We could represent this phenomenon by a selection matrix, such as  $\mathbf{S}_k = \text{diag}\{s_{k,1}, \dots, s_{k,d}\}$ , in which  $s_{k,i} \in \{0, 1\}$  indicates whether the  $i$ -th entry is clipped (in this case,  $s_{k,i} = 0$ ) or not (in this case,  $s_{k,i} = 1$ ). Consequently, we can express the model update under MAC as follows:

$$\begin{aligned} \mathbf{w}_{k+1} &= \mathbf{w}_k - \eta[\mathbf{S}_k(\nabla f(\mathbf{w}_k) + \boldsymbol{\xi}_k) \\ &\quad + (\mathbf{I} - \mathbf{S}_k)(\text{med}(\nabla f(\mathbf{w}_k) + \boldsymbol{\xi}_k) \cdot \mathbf{1} + \hat{\mathbf{C}}_k)] \\ &= \mathbf{w}_k - \eta[\mathbf{S}_k \nabla f(\mathbf{w}_k) \\ &\quad + (\mathbf{I} - \mathbf{S}_k) \cdot \text{med}(\nabla f(\mathbf{w}_k) + \boldsymbol{\xi}_k) \cdot \mathbf{1}] - \eta \boldsymbol{\zeta}_k \end{aligned} \quad (16)$$

where  $\mathbf{I}$  is identity matrix and  $\boldsymbol{\zeta}_k$  is

$$\boldsymbol{\zeta}_k = \mathbf{S}_k \boldsymbol{\xi}_k + (\mathbf{I} - \mathbf{S}_k) \hat{\mathbf{C}}_k \quad (17)$$

in which  $\hat{\mathbf{C}}_k = (c_{k,1}, \dots, c_{k,d})^\top$  where each  $c_{k,i}$  is given by

$$c_{k,i} = \text{sgn}(g_{k,i} - \text{med}(\nabla f(\mathbf{w}_k) + \boldsymbol{\xi}_k)) \cdot C, \quad i \in [d], \quad (18)$$

where  $\text{sgn}(\cdot)$  takes the sign of its input variable.

Notice that the heavy-tailed noise  $\boldsymbol{\xi}_k$  is zero-mean, then we make an approximation from median to mean:

$$\text{med}(\nabla f(\mathbf{w}_k) + \boldsymbol{\xi}_k) \approx \frac{1}{d} \mathbf{1}^\top \nabla f(\mathbf{w}_k) \quad (19)$$

### B. Proof of Lemma 1

Let us start from the range  $R$  of entries from gradient  $\nabla f(\mathbf{w}_k)$ , which is defined as

$$R = \max(|\nabla f(w_{k,i}) - \nabla f(w_{k,j})|), \forall i, j \in [d]. \quad (20)$$

Then following the Assumption 3 we know that there has

$$R \leq \sqrt{2}G. \quad (21)$$

And we assume that  $C > \sqrt{2}G$ . Then considering the clipping judge condition, if

$$|\nabla f(w_{k,i}) + \xi_{k,i} - \text{med}(\nabla f(\mathbf{w}_k) + \boldsymbol{\xi}_k)| \leq C, \quad (22)$$

then the  $i$ -th entry would not be clipped. If we denote  $\text{med}(\nabla f(\mathbf{w}_k) + \boldsymbol{\xi}_k)$  as  $\nabla f(w_{k,m}) + \xi_{k,m}$ , then there has

$$\begin{aligned} &|\nabla f(w_{k,i}) + \xi_{k,i} - \nabla f(w_{k,m}) - \xi_{k,m}| \\ &\leq |\nabla f(w_{k,i}) - \nabla f(w_{k,m})| + |\xi_{k,i} - \xi_{k,m}| \\ &\leq \sqrt{2}G + |\xi_{k,i} - \xi_{k,m}|. \end{aligned} \quad (23)$$

And we denote  $\xi_{k,c} = \xi_{k,i} - \xi_{k,m}$ , because  $\xi_{k,i}$  and  $\xi_{k,m}$  are independent and identically distributed as  $S\alpha S(\alpha, \tau)$ , thus

$$\xi_{k,c} \sim S\alpha S(\alpha, \sqrt{2}\tau). \quad (24)$$

Then we let

$$\sqrt{2}G + |\xi_{k,c}| \leq C, \quad (25)$$

which is saying that the clipping will not occur if  $|\xi_{k,c}| \leq C - \sqrt{2}G$ , the probability is

$$\mathbb{P}\{|\xi_{k,c}| \leq C - \sqrt{2}G\}. \quad (26)$$

In a similar way, there the clipping will occur if

$$\begin{aligned} C &< |\xi_{k,c}| - |\nabla f(w_{k,i}) - \nabla f(w_{k,m})| \\ &\leq |\nabla f(w_{k,i}) + \xi_{k,i} - \nabla f(w_{k,m}) - \xi_{k,m}|, \end{aligned} \quad (27)$$

or

$$\begin{aligned} C &< |\nabla f(w_{k,i}) - \nabla f(w_{k,m})| - |\xi_{k,c}| \\ &\leq |\nabla f(w_{k,i}) + \xi_{k,i} - \nabla f(w_{k,m}) - \xi_{k,m}|. \end{aligned} \quad (28)$$

That is

$$|\xi_{k,c}| > |\nabla f(w_{k,i}) - \nabla f(w_{k,m})| + C \quad (29)$$

or

$$|\xi_{k,c}| < |\nabla f(w_{k,i}) - \nabla f(w_{k,m})| - C. \quad (30)$$

Because  $C > \sqrt{2}G$ , thus  $|\nabla f(w_{k,i}) - \nabla f(w_{k,m})| - C$  is constant negative,

$$\mathbb{P}\{|\xi_{k,c}| < |\nabla f(w_{k,i}) - \nabla f(w_{k,m})| - C\} = 0. \quad (31)$$

That is when  $|\xi_{k,c}| \geq C + \sqrt{2}G$ , the clipping will occur, the probability is

$$\mathbb{P}\{|\xi_{k,c}| > C + \sqrt{2}G\}. \quad (32)$$




Fig. 5: Clipping decision interval

As illustrated in Fig. 5, for random variable  $|\xi_{k,c}|$ , no clipping occurs if  $|\xi_{k,c}| \leq C - \sqrt{2}G$ . However, if  $|\xi_{k,c}| > C - \sqrt{2}G$ , clipping occurs. In our analysis, we made an approximation wherein this part is entirely subject to clipping. So, at this stage, we get the probability of without clipping which is denoted as  $p_C$ , and

$$\begin{aligned} p_C &= \mathbb{P}\{s_i = 1\} = \mathbb{P}\{|\xi_{k,c}| \leq C - \sqrt{2}G\} \\ &= \mathbb{P}\left\{|\xi_{k,i}| \leq \frac{\sqrt{2}}{2}C - G\right\} \sim 1 - \frac{\tau^\alpha}{C^\alpha} \end{aligned} \quad (33)$$

It is worth noting that our approximation imposes a more stringent clipping mechanism, causing it to occur more frequently. This, in turn, leads to a more relaxed analytical outcome, which is advantageous compared to the actual scenario being analyzed.

### C. Proof of Theorem 1

We have

$$\begin{aligned} \mathbb{E}[f(\mathbf{w}_{k+1}) - f(\mathbf{w}_k)] &\leq \mathbb{E}[\langle \nabla f(\mathbf{w}_k), \mathbf{w}_{k+1} - \mathbf{w}_k \rangle] \\ &\quad + \frac{L}{2} \mathbb{E}[\|\mathbf{w}_{k+1} - \mathbf{w}_k\|^2] \end{aligned} \quad (34)$$

Then

$$\begin{aligned} &\mathbb{E}[\langle \nabla f(\mathbf{w}_k), \mathbf{w}_{k+1} - \mathbf{w}_k \rangle] \\ &= -\eta \left( \mathbb{E}[\langle \nabla f(\mathbf{w}_k), \mathbf{S}_k \nabla f(\mathbf{w}_k) \rangle] \right. \\ &\quad + \mathbb{E}[\langle \nabla f(\mathbf{w}_k), (\mathbf{I} - \mathbf{S}_k) \text{med}(\nabla f(\mathbf{w}_k) + \boldsymbol{\xi}_k) \cdot \mathbf{1} \rangle] \\ &\quad \left. + \mathbb{E}[\langle \nabla f(\mathbf{w}_k), \boldsymbol{\zeta}_k \rangle] \right) \\ &\simeq -\eta \left( \mathbb{E}[\langle \nabla f(\mathbf{w}_k), \mathbf{S}_k \nabla f(\mathbf{w}_k) \rangle] \right. \\ &\quad + \mathbb{E}[\langle \nabla f(\mathbf{w}_k), (\mathbf{I} - \mathbf{S}_k) \frac{1}{d} \mathbf{1}^\top (\nabla f(\mathbf{w}_k) + \boldsymbol{\xi}_k) \cdot \mathbf{1} \rangle] \\ &\quad \left. + \mathbb{E}[\langle \nabla f(\mathbf{w}_k), \boldsymbol{\zeta}_k \rangle] \right). \end{aligned} \quad (35)$$

Notice that

$$\mathbb{E}[\boldsymbol{\zeta}_k] = p_C \mathbb{E}[\boldsymbol{\xi}_k] + (1 - p_C) \mathbb{E}[\hat{\mathbf{C}}_k], \quad (36)$$

and  $\mathbb{E}[\boldsymbol{\xi}_k]$  is at the condition  $|\xi_{k,i}| \leq \frac{\sqrt{2}}{2}C - G$ , where  $i \in [d]$ . For vector  $\hat{\mathbf{C}}_k$ , due to the properties of the median, the following holds probabilistically for  $i \in [d]$ :

$$\mathbb{P}\{c_{k,i} = C\} = \mathbb{P}\{c_{k,i} = -C\} = 0.5, \quad (37)$$

so we have

$$\mathbb{E}[\boldsymbol{\zeta}_k] = \mathbf{0}. \quad (38)$$

Then we get

$$\begin{aligned} &\mathbb{E}[\langle \nabla f(\mathbf{w}_k), \mathbf{w}_{k+1} - \mathbf{w}_k \rangle] \\ &\simeq -\eta \left( p_C \mathbb{E}[\|\nabla f(\mathbf{w}_k)\|^2] + \frac{1}{d} (1 - p_C) \mathbb{E}[\|\mathbf{1}^\top \nabla f(\mathbf{w}_k)\|^2] \right). \end{aligned} \quad (39)$$

Notice that  $\mathbf{S}_k^\top \mathbf{S}_k = \mathbf{S}_k$ , we have

$$\begin{aligned} &\mathbb{E}[\|\mathbf{w}_{k+1} - \mathbf{w}_k\|] \\ &\simeq \eta^2 \mathbb{E} \left[ \left\| \mathbf{S}_k \nabla f(\mathbf{w}_k) + (\mathbf{I} - \mathbf{S}_k) \frac{1}{d} \mathbf{1}^\top \nabla f(\mathbf{w}_k) \cdot \mathbf{1} + \boldsymbol{\zeta}_k \right\|^2 \right] \\ &= \eta^2 \left[ \mathbb{E}[\|\mathbf{S}_k \nabla f(\mathbf{w}_k)\|^2] + \frac{1}{d^2} \mathbb{E}[\|(\mathbf{I} - \mathbf{S}_k) \mathbf{1}^\top \nabla f(\mathbf{w}_k) \cdot \mathbf{1}\|^2] \right. \\ &\quad \left. + \mathbb{E}[\|\boldsymbol{\zeta}_k\|^2] \right] \\ &= \eta^2 \left( p_C \mathbb{E}[\|\nabla f(\mathbf{w}_k)\|^2] + \frac{1}{d} (1 - p_C) \mathbb{E}[\|\mathbf{1}^\top \nabla f(\mathbf{w}_k)\|^2] \right. \\ &\quad \left. + \mathbb{E}[\|\boldsymbol{\zeta}_k\|^2] \right). \end{aligned} \quad (40)$$

Then from inequalities (39), (34) and (40), we get that

$$\begin{aligned} \mathbb{E}[f(\mathbf{w}_{k+1}) - f(\mathbf{w}_k)] &\leq - \left( 1 - \frac{1}{2} \eta L \right) \eta p_C \mathbb{E}[\|\nabla f(\mathbf{w}_k)\|^2] \\ &\quad + \left( \frac{1}{2} \eta L - 1 \right) \eta (1 - p_C) \mathbb{E}[\|\mathbf{1}^\top \nabla f(\mathbf{w}_k)\|^2] \\ &\quad + \frac{1}{2} \eta^2 L \mathbb{E}[\|\boldsymbol{\zeta}_k\|^2]. \end{aligned} \quad (41)$$

Let  $\eta$  be constant, and satisfies that  $\eta \leq \frac{2}{L}$ , then we get

$$\begin{aligned} \mathbb{E}[f(\mathbf{w}_{k+1}) - f(\mathbf{w}_k)] &\leq - \left( 1 - \frac{1}{2} \eta L \right) \eta p_C \mathbb{E}[\|\nabla f(\mathbf{w}_k)\|^2] \\ &\quad + \frac{1}{2} \eta^2 L \mathbb{E}[\|\boldsymbol{\zeta}_k\|^2]. \end{aligned} \quad (42)$$

For  $\mathbb{E}[\|\boldsymbol{\zeta}_k\|^2]$ , notice that  $s_{k,i}^2 = s_{k,i}$ , we have

$$\begin{aligned} \mathbb{E}[\|\boldsymbol{\zeta}_k\|^2] &= \mathbb{E} \left[ \left\| \mathbf{S}_k \boldsymbol{\xi}_k + (\mathbf{I} - \mathbf{S}_k) \hat{\mathbf{C}}_k \right\|^2 \right] \\ &= \mathbb{E} \left[ \left\| \mathbf{S}_k \boldsymbol{\xi}_k \right\|^2 \right] + \mathbb{E} \left[ \left\| (\mathbf{I} - \mathbf{S}_k) \hat{\mathbf{C}}_k \right\|^2 \right] \\ &= \mathbb{E} \left[ \sum_{i=1}^d s_{k,i}^2 \xi_{k,i}^2 \right] + \mathbb{E} \left[ \sum_{i=1}^d (1 - s_{k,i})^2 c_{k,i}^2 \right] \\ &= \mathbb{E} \left[ \sum_{i=1}^d s_{k,i} \xi_{k,i}^2 \right] + \mathbb{E} \left[ \sum_{i=1}^d (1 - s_{k,i}) c_{k,i}^2 \right] \\ &= dp_C \mathbb{E} \left[ \xi_{k,i}^2 \mid |\xi_{k,i}| \leq \frac{\sqrt{2}}{2}C - G \right] + d(1 - p_C) C^2 \\ &\leq dp_C \left( \frac{\sqrt{2}}{2}C - G \right)^2 + d(1 - p_C) C^2. \end{aligned} \quad (43)$$

Then we get that

$$\begin{aligned} \left(1 - \frac{1}{2}\eta L\right) \eta p_C \mathbb{E} [\|\nabla f(\mathbf{w}_k)\|^2] &\leq \mathbb{E}[f(\mathbf{w}_k)] - \mathbb{E}[f(\mathbf{w}_{k+1})] \\ &+ \frac{1}{2}\eta^2 dL \left[ p_C \left( \frac{\sqrt{2}}{2} C - G \right)^2 + (1 - p_C) C^2 \right]. \end{aligned} \quad (44)$$

Summing inequality (44) for all  $k \in \{0, 1, \dots, K-1\}$ , and rearranging the results we get

$$\begin{aligned} \frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E} [\|\nabla f(\mathbf{w}_k)\|^2] &\leq \frac{2(f(\mathbf{w}_0) - \mathbb{E}[f(\mathbf{w}_K)])}{K p_C (2 - \eta L) \eta} \\ &+ \frac{1}{2}\eta^2 dL \left[ p_C \left( \frac{\sqrt{2}}{2} C - G \right)^2 + (1 - p_C) C^2 \right] \end{aligned} \quad (45)$$

which completes the proof.

## VII. ACKNOWLEDGEMENT

The work of Jiaying Li and Howard H. Yang was supported in part by the National Key R&D Program of China under Grant 2024YFE0200700, in part by the open research fund of National Mobile Communications Research Laboratory, and in part by the National Natural Science Foundation of China under Grant 62201504. This work has also been supported in part by the SNS JU project 6G-GOALS under the EU's Horizon program grant agreement No. 101139232.

## REFERENCES

- [1] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. Int. Conf. Artif. Intell. Stat. (AISTATS)*, Fort Lauderdale, USA, Apr. 2017.
- [2] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," *IEEE Signal Process. Mag.*, vol. 37, pp. 50–60, May 2019.
- [3] J. Park, S. Samarakoon, M. Bennis, and M. Debbah, "Wireless network intelligence at the edge," *Proc. IEEE*, vol. 107, no. 11, pp. 2204–2239, Oct. 2019.
- [4] Y.-J. Liu, S. Qin, Y. Sun, and G. Feng, "Resource consumption for supporting federated learning in wireless networks," *IEEE Trans. Wireless Commun.*, vol. 21, no. 11, pp. 9974–9989, June 2022.
- [5] L. Lyu, H. Yu, and Q. Yang, "Threats to federated learning: A survey," *Available as ArXiv:2003.02133*, 2020.
- [6] B. Nazer and M. Gastpar, "Computation over multiple-access channels," *IEEE Trans. Inf. Theory*, vol. 53, no. 10, pp. 3498–3516, Oct. 2007.
- [7] K. Yang, T. Jiang, Y. Shi, and Z. Ding, "Federated learning via over-the-air computation," *IEEE Trans. Wireless Commun.*, vol. 19, no. 3, pp. 2022–2035, Jan. 2020.
- [8] M. Frey, I. Bjelaković, and S. Stańczak, "Over-the-air computation in correlated channels," *IEEE Trans. Signal Process.*, vol. 69, pp. 5739–5755, Aug. 2021.
- [9] B. Xiao, X. Yu, W. Ni, X. Wang, and H. V. Poor, "Over-the-air federated learning: Status quo, open challenges, and future directions," *Fundam. res.*, Feb. 2024.
- [10] G. Zhu, J. Xu, K. Huang, and S. Cui, "Over-the-air computing for wireless data aggregation in massive iot," *IEEE Trans. Wireless Commun.*, vol. 28, no. 4, pp. 57–65, Sept. 2021.
- [11] A. Şahin and R. Yang, "A survey on over-the-air computation," *IEEE Commun. Surv. Tutor.*, Apr. 2023.
- [12] G. Zhu, Y. Wang, and K. Huang, "Broadband analog aggregation for low-latency federated edge learning," *IEEE Trans. Wireless Commun.*, vol. 19, no. 1, pp. 491–506, Oct. 2019.
- [13] M. M. Amiri and D. Gündüz, "Machine learning at the wireless edge: Distributed stochastic gradient descent over-the-air," *IEEE Trans. Signal Process.*, vol. 68, pp. 2155–2169, Mar. 2020.
- [14] H. Guo, A. Liu, and V. K. Lau, "Analog gradient aggregation for federated learning over wireless networks: Customized design and convergence analysis," *IEEE Internet Things J.*, vol. 8, no. 1, pp. 197–210, June 2020.
- [15] Z. Zhang, G. Zhu, R. Wang, V. K. Lau, and K. Huang, "Turning channel noise into an accelerator for over-the-air principal component analysis," *IEEE Trans. Wireless Commun.*, vol. 21, no. 10, pp. 7926–7941, Apr. 2022.
- [16] Z. Chen, H. H. Yang, and T. Q. Quek, "Edge intelligence over the air: Two faces of interference in federated learning," *IEEE Commun. Mag.*, Aug. 2023.
- [17] D. Middleton, "Statistical-physical models of electromagnetic interference," *IEEE Trans. Electromagn. Compat.*, no. 3, pp. 106–127, Aug. 1977.
- [18] L. Clavier, T. Pedersen, I. Larrad, M. Lauridsen, and M. Egan, "Experimental evidence for heavy tailed interference in the iot," *IEEE Commun. Lett.*, vol. 25, no. 3, pp. 692–695, Oct. 2020.
- [19] N. Mital and D. Gündüz, "Bandwidth expansion for over-the-air computation with one-sided csi," in *Proc. IEEE Int. Symp. on Inf. Theory (ISIT)*, Espoo, Finland, June 2022.
- [20] T. Sery and K. Cohen, "On analog gradient descent learning over multiple access fading channels," *IEEE Trans. Signal Process.*, vol. 68, pp. 2897–2911, Apr. 2020.
- [21] J. Lee, Y. Jang, H. Kim, S.-L. Kim, and S.-W. Ko, "Over-the-air consensus for distributed vehicle platooning control (extended version)," *Available as ArXiv:2211.06225*, 2022.
- [22] M. Goldenbaum and S. Stańczak, "Computing the geometric mean over multiple-access channels: Error analysis and comparisons," in *Asilomar Conf. on Signals, Systems, and Computers (ACSSC)*, Pacific Grove, CA, USA, Nov. 2010.
- [23] R. Pascanu, T. Mikolov, and Y. Bengio, "On the difficulty of training recurrent neural networks," in *Proc. Int. Conf. Mach. Learn. (ICML)*, Atlanta GA USA, June 2013.
- [24] P. Seetharaman, G. Wichern, B. Pardo, and J. Le Roux, "Autoclip: Adaptive gradient clipping for source separation networks," in *Proc. IEEE Int. Workshop on Mach. Learn. for Signal Process. (MLSP)*, Espoo, Finland, Sept. 2020.
- [25] J. Qian, Y. Wu, B. Zhuang, S. Wang, and J. Xiao, "Understanding gradient clipping in incremental gradient methods," in *Proc. Int. Conf. Artif. Intell. Stat. (AISTATS)*, Apr. 2021.
- [26] H. Yang, P. Qiu, and J. Liu, "Taming fat-tailed ("heavier-tailed" with potentially infinite variance) noise in federated learning," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, New Orleans, Louisiana, USA, Nov. 2022.
- [27] H. H. Yang, Z. Chen, T. Q. Quek, and H. V. Poor, "Revisiting analog over-the-air machine learning: The blessing and curse of interference," *IEEE J. Sel. Topics Signal Process.*, vol. 16, no. 3, pp. 406–419, Dec. 2021.
- [28] G. Samorodnitsky, M. S. Taqqu, and R. Linde, "Stable non-gaussian random processes: stochastic models with infinite variance," *Bull. London Math. Soc.*, vol. 28, no. 134, pp. 554–555, 1996.
- [29] A. Krizhevsky, "Learning multiple layers of features from tiny images," *Master's thesis, University of Tront*, 2009.
- [30] S. Caldas, S. M. K. Duddu, P. Wu, T. Li, J. Konečný, H. B. McMahan, V. Smith, and A. Talwalkar, "Leaf: A benchmark for federated settings," *Available as ArXiv:1812.01097*, 2018.
- [31] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Las Vegas, NV, USA, June 2016.
- [32] D. A. E. Acar, Y. Zhao, R. M. Navarro, M. Mattina, P. N. Whatmough, and V. Saligrama, "Federated learning based on dynamic regularization," *Available as ArXiv:2111.04263*, 2021.
- [33] T.-M. H. Hsu, H. Qi, and M. Brown, "Measuring the effects of non-identical data distribution for federated visual classification," *Available as ArXiv:1909.06335*, 2019.