Improving Generalization in Federated Learning with Highly Heterogeneous Data via Momentum-Based Stochastic Controlled Weight Averaging

Junkang Liu¹ Yuanyuan Liu^{* 2} Fanhua Shang^{* 1} Hongying Liu^{* 3 4} Jin Liu⁵ Wei Feng¹

Abstract

For federated learning (FL) algorithms such as FedSAM, their generalization capability is crucial for real-word applications. In this paper, we revisit the generalization problem in FL and investigate the impact of data heterogeneity on FL generalization. We find that FedSAM usually performs worse than FedAvg in the case of highly heterogeneous data, and thus propose a novel and effective federated learning algorithm with Stochastic Weight Averaging (called FedSWA), which aims to find flatter minima in the setting of highly heterogeneous data. Moreover, we introduce a new momentum-based stochastic controlled weight averaging FL algorithm (FedMoSWA), which is designed to better align local and global models. Theoretically, we provide both convergence analysis and generalization bounds for FedSWA and FedMoSWA. We also prove that the optimization and generalization errors of FedMoSWA are smaller than those of their counterparts, including FedSAM and its variants. Empirically, experimental results on CIFAR10/100 and Tiny ImageNet demonstrate the superiority of the proposed algorithms compared to their counterparts. Open source code at: https://github. com/junkangLiu0/FedSWA.

1. Introduction

With increasing concerns about data privacy and security, federated learning (FL) has emerged as a potential distributed machine learning paradigm (McMahan et al., 2017).

Proceedings of the 42^{nd} International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).



Figure 1: The training loss surfaces of ResNet-18 trained by FedAvg (a), FedSAM (b), and our FedSWA (c) on CIFAR-100 with Dirichlet-0.1 (high data heterogeneity). The test accuracies of FedAvg, FedSAM, and FedSWA are 45.8%, 40.1%, and 50.3%, respectively, and FedSAM performs the worst and fails in this case. (d) highlights the sharpness minimizing conflicts due to discrepancies between local and global loss landscapes caused by data heterogeneity. (e) and (f) show that FedSWA can obtain global flat minima in the setting of high data heterogeneity. The blue dots denote the client model, and the red dot is the server model in (f). The outer layer is the training loss surface, the inner layer is the test loss surface, and the global model generalizes well at a flat minimum in (f). The double arrow denotes the generalization error gap between clients and the server.

FL allows for collaborative training of a global model across several clients without exchanging raw data, ensuring privacy while exploiting decentralized data. FL algorithms have considerable promise in various areas, including healthcare, finance, and personalized mobile services (Rieke et al., 2020; Antunes et al., 2022; Byrd & Polychroniadou, 2020; Wang et al., 2024; Bian et al., 2024; 2025; Liu et al., 2024; Wu et al., 2025a;b).

FL confronts the issue of data heterogeneity because of the non-independent and identically distributed (Non-IID) data across clients (Zhang et al., 2021; Liu et al., 2024),

¹College of Intelligence and Computing, Tianjin University, Tianjin, China ²School of Artificial Intelligence, Xidian University, Xi'an, China ³Medical School, Tianjin University, Tianjin, China ⁴Peng Cheng Lab, Shenzhen, China ⁵School of Cyber Engineering, Xidian University, Xi'an, China. Correspondence to: Yuanyuan Liu <yyliu@xidian.edu.cn>, Fanhua Shang <fnshang@tju.edu.cn>, Hongying Liu <hyliu2009@tju.edu.cn>.

along with different computational capabilities and network connectivity. These characteristics bring complications, particularly regarding the model's capacity to generalize to unseen data, which is critical to the practical applicability of FL algorithms. The generalization study assesses the model's capacity to generalize from empirical to theoretical optimal risks, providing vital insights into the efficacy of FL algorithms.

The generalization of many deep learning methods has been extensively studied to address the overfitting problem. Finding flat minima in model parameters has been demonstrated to improve generalization performance. The two methods, Stochastic Weight Averaging (SWA) (Izmailov et al., 2018) and Sharpness-Aware Minimization (SAM) (Foret et al., 2021), have received significant attention. In distributed learning, particularly FL, data heterogeneity always enables the global model to converge to a sharp local minimum (Fan et al., 2024), and generalization of sharp local minimum is poor, as shown in Figure 1(a).

To tackle this challenge, several studies have introduced SAM to local client optimization, which aims to enhance generalization by finding a flat minimum through minimizing the local loss after perturbation. Ou et al. (2022) adopted SAM to FL and introduced FedSAM, and they also proposed MoFedSAM, which incorporates local momentum acceleration. Caldarola et al. (2022) applied adaptive SAM optimizer in FL local training and proposed FedASAM. FedGAMMA (Dai et al., 2023) improves FedSAM by using the SCAFFOLD's variance reduction technique (Karimireddy et al., 2020). FedSpeed (Sun et al., 2023) utilized the gradient computed at the SAM-based perturbed weight. Fan et al. (2024) proposed FedLESAM locally estimates the direction of global perturbation on the client side as the difference between global models received in the previous active and current rounds. However, adequate experimental results show that FedSAM is extremely ineffective in the setting of highly heterogeneous data (see in Figure 1(b)). Moreover, Lee & Yoon (2024) also pointed out that SAM in local clients can find a local flat minimum, not a global flat minimum, and FedSAM fails to generalize in the setting of high data heterogeneity. In addition, Kaddour et al. (2022) noted that SAM can be close to sharp directions, which is unfriendly to the federated averaging algorithm.

To address the above shortcomings of SAM-type algorithms as local optimizers, this paper introduces SWA to improve generalization of FL with highly heterogeneous data. SWA improves the model's generalization capability by averaging different model weights at the final stage of training, thereby locating a model lying in a flat region of the loss landscape. Therefore, we propose a new Stochastic Weight Averaging Federated Learning (FedSWA) algorithm, which aims to find flatter minima in the setting of highly heterogeneous data, as shown in Figure 1(e,f). In other words, the goal of FedSWA is to find global flat minima, avoiding the drawbacks of FedSAM, which finds local flat minima instead of global flat minima. Moreover, SWA is computationally more efficient than SAM since SWA does not require additional forward and backward propagations to compute the perturbations in SAM. In this paper, we will revisit SAM and SWA on FL to improve the generalization ability on highly heterogeneous data.

In Figure 1(d), we can observe that FedSAM finds a global flat minima but a high loss surface, and the generalization to the test loss surface does not work very well. This is because FedSAM focuses too much on local flat minima, not global flat minima, causing it to overlook superior global minima. The flatness and low loss surface of FedSAM in local client training does not imply that the global model is flat (Lee & Yoon, 2024). Our FedSWA algorithm aims to find a flat and low global loss surface, which generalizes to the test set much better than FedSAM, as shown in Figure 1(c,e).

While FedSWA enhances the global model's generalization compared to FedSAM, it does not guarantee a consistent flat and low minimum for the local model to agree with the global model in the setting of highly heterogeneous data. To seamlessly integrate the smoothness of both local and global models, we present a more advanced algorithm within our framework, called Federated Learning via Momentum-Based Stochastic Controlled Weight Averaging (FedMoSWA). The key idea of FedMoSWA is inspired by the momentum-based variance reduction (Cutkosky & Orabona, 2020). Intuitively, FedMoSWA estimates the update direction for the server variable m and the client variable c_i . Their difference, $(m - c_i)$, is used to correct the estimation of the client drift of local updates, maintaining the consistency of local updates with global updates. This strategy successfully overcomes data heterogeneity.

Despite that there is considerable research on optimization error and convergence in FL, the analysis of generalization remains significantly challenging. To theoretically characterize generalization performance, researchers have developed a variety of analysis techniques, including uniform convergence methods (Vapnik et al., 1998), operator approximation techniques (Smale & Zhou, 2007), information-theoretic tools (Russo & Zou, 2019), and algorithmic stability analysis (Rogers & Wagner, 1978; Hardt et al., 2016). Among these methods, stability analysis has attracted particular attention due to its theoretical guarantees that are independent of the capacity measurement of the hypothesis function space. It is also widely applicable and sensitive to data distribution. This paper proposes an analytical framework based on uniform stability, which is particularly useful for studying generalization errors by considering the dependence on specific FL algorithms and data heterogeneity.

Table 1: Comparison of generalization and optimization errors in the non-convex setting. σ is the variance of stochastic gradients, σ_g is data heterogeneity, s is the number of participating clients, and T is the number of communication rounds. K is the number of local iterations. $\ell(x, y; \theta)$ is L-Lipschitz continous and β -smooth. F^* is a minimum value of Problem (1) below, $F := F_0 - F^*$, $\bar{c} > \tilde{c} = 1 + (2 + 1/KT)^{K-1}/T \gg 1$, which means the generalization error of FedSWA is smaller than that of FedSAM. The generalization error of MoFedSWA is smaller than that of FedSAM are proved in (Patel et al., 2022). Generalization bounds without client sampling.

Algorithm	Generalization error	Optimization error
FedSAM (Qu et al., 2022)	$\mathcal{O}\left(\frac{L}{mn\beta}e^{1+\frac{1}{T}}(\overline{c}L+\overline{c}\sigma_{g}+\overline{c}\sigma)\right)$	$\mathcal{O}\left(rac{eta F}{\sqrt{TKs}}+rac{\sqrt{K}\sigma_g{}^2}{\sqrt{Ts}}+rac{L^2\sigma^2}{T^{3/2}K}+rac{L^2}{T^2} ight)$
MoFedSAM (Qu et al., 2022)	$\mathcal{O}\left(\frac{L}{mn\beta}e^{1+\frac{1}{T}}(\bar{c}L+\bar{c}\sigma_{g}+\bar{c}\sigma)\right)$	$\mathcal{O}\left(\frac{\beta \widetilde{LF}}{\sqrt{TKs}} + \frac{\beta \sqrt{K} \sigma_g^{\ 2}}{\sqrt{Ts}} + \frac{L^2 \sigma^2}{T^{3/2} K} + \frac{\sqrt{K} \widetilde{L}^2}{T^{3/2} \sqrt{s}}\right)$
FedSWA (ours)	$\mathcal{O}\left(\frac{L}{mn\beta}e^{1+\frac{1}{T}}(\tilde{c}L+\tilde{c}\boldsymbol{\sigma}_{\boldsymbol{g}}+\tilde{c}\boldsymbol{\sigma})\right)$	$\mathcal{O}\left(\frac{\beta\left(\sigma+\sqrt{K}\sigma_{g}\right)\sqrt{F}}{\sqrt{TKs}}+\frac{F^{2/3}\left(\beta\sigma_{g}^{2}\right)^{1/3}}{T^{2/3}}+\frac{\beta F}{T}\right)$
FedMoSWA (ours)	$\mathcal{O}\big(\tfrac{L}{mn\beta}e^{1+\frac{1}{T}}(\tilde{c}L+\boldsymbol{\sigma_g}+\tilde{c}\boldsymbol{\sigma})\big)$	$\mathcal{O}\left(rac{\sigma\sqrt{F}}{\sqrt{TKs}}\left(\sqrt{1+rac{s}{lpha^2}} ight)+rac{eta F}{T}\left(rac{m}{s} ight)^{rac{2}{3}} ight)$

Contributions: We propose a new FL generalization analysis framework and develop two efficient algorithms to improve the generalization and optimization errors. Our main contributions can be summarized as follows:

• New FL algorithm inspired by SWA with better generalization: We revisit the role of SWA and SAM for federated learning and discuss the superiority of SWA with heterogeneous data. To generalize the global model, we first propose a novel algorithm, FedSWA, to improve the generalization of FL, which is better than FedSAM.

• Momentum-based stochastic controlled weight averaging algorithm: To help local models find consistent flat minima aligned with the global model, we develop a new FedMoSWA algorithm. We also propose one momentumbased stochastic controlled method, which is designed to increase FL generalization and overcome data heterogeneity. Our theoretical result shows that FedMoSWA provides theoretically better upper bounds for generalization error than FedSAM. Experimentally, FedMoSWA is better than both FedSAM and his variants.

• Novel FL generalization analysis framework: We propose a new FL generalization error analysis framework, which can account for the effect of data heterogeneity on generalization error. Our theoretical results show that the generalization bound of our FedSWA is $O(\frac{L}{mn\beta}e^{\frac{1}{T}+1}(\tilde{c}L + \tilde{c}\sigma_g + \tilde{c}\sigma))$, which is correlated with data heterogeneity and is superior to that of FedSAM. We prove that the generalization error of FedMoSWA is $O(\frac{L}{mn\beta}e^{\frac{1}{T}+1}(\tilde{c}L + \sigma_g + \tilde{c}\sigma))$, which is better than those of FedSAM and FedSWA, where $\tilde{c} = 1 + (2 + 1/KT)^{K-1}/T \gg 1$.

• **Better practical results:** We also evaluate our algorithms for various federated learning tasks. Experimental results indicate that the proposed algorithms significantly improve the generalization and optimization performance.

2. Related work

• Heterogeneity Issues in FL: In past years, various strategies have been proposed to solve the heterogeneity issues in FL. FedAvgM (Hsu et al., 2019) is a proposed method that introduces momentum terms during global model updating. FedACG (Kim et al., 2024) improves the inter-client interoperability by the server broadcasting a global model with a prospective gradient consistency. SCAFFOLD (Karimireddy et al., 2020) uses SAGA-like control variables to mitigate client-side drift, which can be regarded as adopting the idea of variance reduction at the client side. Our FedMoSWA method is based on momentum-based variance reduction, which is different from SCAFFOLD and other momentum methods.

• Generalization Analysis of FL: Hu et al. (2023) provided a systematic analysis of the generalization error of FL in the two-level framework, which captures the missed participating gap. Sefidgaran et al. (2022) used tools from rate-distortion theory to establish new upper bounds on the generalization error of statistical distributed learning algorithms. Sun et al. (2024) analyzed the generalization performance of federated learning through algorithm stability, but did not propose an improved algorithm to overcome data heterogeneity.

• Uniform Stability Generalization Analysis of Algorithms: Uniform stability (Bousquet & Elisseeff, 2002) is a classical tool to analyze the generalization error of an algorithm. For instance, Hardt et al. (2016) and Zhang et al. (2022) analyzed the generalization of stochastic gradient descent (SGD) (Bottou, 2010) via uniform stability. Yuan et al. (2019) investigated stagewise SGD and showed the advantages of the statewise strategy.

• SWA: The concept of averaging weights can be traced back to early efforts to accelerate the convergence of SGD (Polyak & Juditsky, 1992; Kaddour). SWA is motivated by an observation about SGD's behavior when training neural

networks: although SGD frequently explores regions in the weight space associated with high-performing models, it seldom reaches the central points of this optimal set. By averaging parameter values, SWA guides the solution closer to the centroid of this space (Izmailov et al., 2018).

3. Proposed Algorithm

3.1. Problem Setup

FL aims to optimize global model with the collaboration local clients, i.e., minimizing the following population risk:

$$F(\boldsymbol{\theta}) = \frac{1}{m} \sum_{i=1}^{m} \left(F_i(\boldsymbol{\theta}) := \mathbb{E}_{\zeta_i \sim \mathcal{D}_i} \left[F_i(\boldsymbol{\theta}; \zeta_i) \right] \right). \quad (1)$$

The function F_i represents the loss function on client *i*. $\mathbb{E}_{\zeta_i \sim \mathcal{D}_i}[\cdot]$ denotes the conditional expectation with respect to the sample ζ_i . Generally, the unattainability of the local population risk $F_i(\boldsymbol{\theta})$ forces us to train the model by minimizing the following empirical approximation of population risk $F_{S_i}(\boldsymbol{\theta})$. For the generalization problem, we consider:

$$F_{\mathcal{S}}(\boldsymbol{\theta}) = \frac{1}{m} \sum_{i=1}^{m} F_{\mathcal{S}_i}(\boldsymbol{\theta}) = \frac{1}{mn} \sum_{i=1}^{m} \sum_{j=1}^{n} \ell(\boldsymbol{x}_{i,j}, \boldsymbol{y}_{i,j}; \boldsymbol{\theta}),$$
(2)

where $F_S(\theta)$ and $F_{S_i}(\theta)$ indicate the empirical risks for the global model and the local model of the *i*-th client, respectively. $F_{S_i}(\theta) = \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{x}_{i,j}, \mathbf{y}_{i,j}; \theta)$ represents the empirical risk loss function on client *i*. Given a dataset $S = \{(\mathbf{x}_{i,j}, \mathbf{y}_{i,j})\}_{i=1,j=1}^{i=m,j=n}, S = S_1 \cup \cdots \cup S_i \cup \cdots \cup S_m$, where the dataset $S_i = \{(\mathbf{x}_{i,j}, \mathbf{y}_{i,j})\}_{j=1}^n$, and $(\mathbf{x}_{i,j}, \mathbf{y}_{i,j})$ is drawn from an unknown distribution \mathcal{D}_i . Assume $\{\mathcal{D}_1, \cdots, \mathcal{D}_m\}$ are independently sampled from \mathcal{D} according to P (Hu et al., 2023). *m* is the number of clints, *n* is the number of training samples per client. In this paper, we assume that $F_S(\theta)$ has a minimum value F_S^* , when $\theta = \theta^*$.

3.2. Our FedSWA Algorithm

To enhance the generalization of FL, our FedSWA algorithm uses SWA for local training and weight aggregation. In the client-side update, this paper proposes a local learning rate decay strategy instead of a constant learning rate. In the t-th round of local training,

$$\eta_0^t = \eta_l, \eta_K^t = \rho \eta_l, \eta_k^t = \eta_l \left(1 - \frac{k}{K} \right) + \frac{k}{K} \rho \eta_l, \quad (3)$$

$$\boldsymbol{\theta}_{i,k+1}^{(t)} = \boldsymbol{\theta}_{i,k}^{(t)} - \eta_k^t g_i\left(\boldsymbol{\theta}_{i,k}^{(t)}\right),\tag{4}$$

where $g_i(\boldsymbol{\theta}_{i,k}^t) = \frac{1}{|\mathcal{B}|} \sum_{(\boldsymbol{x},\boldsymbol{y})\in\mathcal{B}} \nabla \ell(\boldsymbol{x}_{i,j},\boldsymbol{y}_{i,j};\boldsymbol{\theta}_{i,k}^t), 0 \leq \rho \leq 1$. Here \mathcal{B} is the sampled minibatch at the (t,k)-th iteration and η_l is the initial local learning rate. k is the k-th local iteration and K is the total number of local iterations.

Algorithm 1 FedSWA, FedMoSWA algorithm.

- Input: λ, ρ, initial server model θ₀, number of clients N, number of communication rounds T, number of local iterations K, local learning rate η_l.
- 2: for t = 0, ..., T do Communicate $(\boldsymbol{\theta}_{t-1})$ to selected clients $i \in [s]$. 3: Communicate (θ_{t-1}, m) to selected clients $i \in [s]$. 4: for $i = 1, \ldots, s$ clients in parallel do 5: for $k = 0, \ldots, K$ local update do 6: Compute mini-batch gradient $g_i(\boldsymbol{\theta}_{i,k}^t)$. 7: $\eta_k^t = \eta_l \left(1 - \frac{k}{K} \right) + \frac{k}{K} \rho \eta_l.$ 8: $\boldsymbol{\theta}_{i,k+1}^{t} \leftarrow \boldsymbol{\theta}_{i,k}^{(t)} - \eta_{k}^{t} \left(g_{i} \left(\boldsymbol{\theta}_{i,k}^{t} \right) \right).$ 9: $\boldsymbol{\theta}_{i,k+1}^t \leftarrow \boldsymbol{\theta}_{i,k}^{(t)} - \eta_k^t \left(g_i \left(\boldsymbol{\theta}_{i,k}^t \right) - \boldsymbol{c_i} + \boldsymbol{m} \right).$ 10: 11: end for Communicate $\left(oldsymbol{ heta}_{i,K}^t
 ight)$ to server. 12: $c_i^+ \leftarrow (i) g_i(\boldsymbol{x}) \text{ or } (ii) c_i - \boldsymbol{m} + \frac{1}{\sum_k \eta_k^t} \left(\boldsymbol{\theta}_{t-1} - \boldsymbol{\theta}_{i,k}^t \right).$ 13: Communicate $(\boldsymbol{\theta}_{i,K}^{t}, \boldsymbol{c}_{i}^{+} - \boldsymbol{m})$ to server, $\boldsymbol{c}_{i} \leftarrow \boldsymbol{c}_{i}^{+}$. 14: 15: end for $\boldsymbol{m} \leftarrow \boldsymbol{m} + \gamma \frac{1}{s} \sum_{i \in [s]} \Delta \boldsymbol{c}_i.$ 16:

17:
$$\boldsymbol{v}_t = \frac{1}{s} \sum_{i=1}^{s} \boldsymbol{\theta}_{i,K}^t, \, \boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} + \alpha \left(\boldsymbol{v}_t - \boldsymbol{\theta}_{t-1} \right).$$

18: end for

 ρ is the coefficient of local learning rate decay, the smaller ρ , the faster the local learning rate decays. Inspired by the LookAhead (Zhang et al., 2019), the server aggregation employs an exponential moving average (EMA) to aggregate the historical weights instead of average aggregation. The specific steps of the algorithm are described in Algorithm 1. In Eq.(3), the local learning rate decreases from η_l to $\rho\eta_l$. After model aggregation, the learning rate is adjusted to the initial learning rate η_l to restart local training. Smith (2017); Gotmare et al. (2019) observed that restoring the initial large learning rate (learning rate restart) helps us to jump out of the worse local minimum and find a better local minimum. We experimentally demonstrate that restarting the learning rate is effective for federated learning (see Table 7).

Compared with FedAvg (McMahan et al., 2017) and Fed-SAM (Qu et al., 2022). FedAvg and FedSAM adopt constant local learning rate and model averaging to aggregate, while our FedSWA uses cyclical learning rate and exponential moving average aggregation which help FedSWA jump out of worse flat minima and look for better flat minima as discussed in (Smith, 2017; Gotmare et al., 2019). The test accuracy of FedSWA is also better than those of FedAvg and FedSAM (see Tables 2 and 3).



Figure 2: (a) and (b) indicate that due to data heterogeneity, the local and global models of FedSWA differ too much to find a less flatter minima, while our FedMoSWA aligns the local and global models and finds a flatter minima.

3.3. Our FedMoSWA Algorithm

Inspired by momentum-based variance reduction (Cutkosky & Orabona, 2020), which achieves variance reduction through the use of a variant of the momentum term, we propose a new FedMoSWA algorithm. Different from our FedSWA, the local update of FedMoSWA becomes

$$\boldsymbol{\theta}_{i,k+1}^{(t)} = \boldsymbol{\theta}_{i,k}^{(t)} - \eta_k^t \left(g_i \left(\boldsymbol{\theta}_{i,k}^{(t)} \right) - \boldsymbol{c_i} + \boldsymbol{m} \right), \quad (5)$$

where c_i is the client control variable, and m is the server control variable. We provide two options to update c_i^+ :

$$\boldsymbol{c}_{i}^{+} \leftarrow \begin{cases} \mathrm{I.} & g_{i}(\boldsymbol{\theta}_{t}), \text{ or} \\ \mathrm{II.} & \boldsymbol{c}_{i} - \boldsymbol{m} + \frac{1}{\sum_{k} \eta_{k}^{t}} (\boldsymbol{\theta}_{t-1} - \boldsymbol{\theta}_{i,K}^{(t)}) \end{cases}$$
 (6)

Depending on the situation, Option I may be more stable than II, but II is less computationally expensive and usually sufficient (we use Option II for our experiments). In the local update in Eq. (5), our proposed FedMoSWA algorithm has global control variable m, which is update by

$$\boldsymbol{m} \leftarrow \boldsymbol{m} + \gamma \frac{1}{s} \sum_{i \in S} \left(\boldsymbol{c}_i^+ - \boldsymbol{m} \right).$$
 (7)

Compared with SCAFFOLD (Karimireddy et al., 2020). The SCAFFOLD algorithm has a global variable, which is updated as: $c = c + \frac{1}{m} \sum_{i \in S} (c_i^+ - c_i)$, but the global control variable *c* has an update delay problem. This is due to the fact that the update of *c* aggregates the old c_i and the new c_i by using the same weights. When the client participation rate is low, the old c_i is too different from *c*, which instead has a bad effect on the update of *c*. Our proposed FedMoSWA algorithm has global variables updated with momentum as $m = m + \gamma \frac{1}{s} \sum_{i \in S} (c_i^+ - m)$, which gives higher weight to the latest uploaded c_i and lower weight to the old c_i , instead of giving the same weight to all c_i as that SCAFFOLD does.

Compared with FedSWA. In Figure 2, we can observe that the local model of our FedMoSWA tends to be more close to

the global optimal than FedSWA. Influenced by the control variables, FedMoSWA has fewer differences between clients, and minima are flatter compared to FedSWA.

Compared with MoFedSAM (Qu et al., 2022). MoFed-SAM is based on FedSAM and presents a momentumbased mechanism to accelerate convergence and improve performance in data heterogeneity. FedMoSWA introduces a momentum-based variance reduction mechanism in FedSWA to accelerate convergence in the setting of high data heterogeneity. Theoretically, our FedMoSWA also outperforms MoFedSAM.

4. FL Generalization via uniform stability

Excess Risk Error Decomposition. One often minimizes the empirical risk $F_{\mathcal{S}}(\theta) = \frac{1}{m} \frac{1}{n} \sum_{i=1}^{m} \sum_{i=1}^{n} \ell(\mathbf{x}_{i,j}, \mathbf{y}_{i,j}; \theta)$ in Eq.(2). In client *i* via a randomized algorithm \mathcal{A}_i , e.g., SGD, to find an estimated optimum $\theta_{\mathcal{A}_i,\mathcal{S}_i} \approx \operatorname{argmin}_{\theta} F_{\mathcal{S}_i}(\theta)$. In server, we use FL optimizer \mathcal{A} , e.g., FedAvg, to find an estimated optimum $\theta_{\mathcal{A},\mathcal{S}} \approx \operatorname{argmin}_{\theta} F_{\mathcal{S}}(\theta)$. However, this empirical solution $\theta_{\mathcal{A},\mathcal{S}}$ differs from the desired optimum $\theta_{\mathcal{D}}^*$ of the population risk $F_{\mathcal{D}}(\theta) = \mathbb{E}_{(\mathbf{x},\mathbf{y})\sim \mathcal{D}}[\ell(\mathbf{x},\mathbf{y};\theta)]$,

$$\boldsymbol{\theta}_{\mathcal{D}}^{\star} \in \operatorname{argmin}_{\boldsymbol{\theta}} F_{\mathcal{D}}(\boldsymbol{\theta}) = \mathbb{E}_{(\boldsymbol{x}, \boldsymbol{y}) \sim \mathcal{D}}[\ell(\boldsymbol{x}, \boldsymbol{y}; \boldsymbol{\theta})].$$
 (8)

To understand the generalization ability of FL algorithms, we want to know: what performance of the estimated optimum $\theta_{\mathcal{A},\mathcal{S}}$ can achieve on the test data $(x, y) \sim \mathcal{D}$ in FL. To answer this question, we analyze the test error $\mathbb{E}_{\mathcal{A},\mathcal{S}} [F_{\mathcal{D}}(\theta_{\mathcal{A},\mathcal{S}})]$ of $\theta_{\mathcal{A},\mathcal{S}}$ via investigating the well-known excess risk error ε_{exc} defined as

$$\varepsilon_{\text{exc}} = \mathbb{E}_{\mathcal{A},\mathcal{S}} \left[F_{\mathcal{D}} \left(\boldsymbol{\theta}_{\mathcal{A},\mathcal{S}} \right) \right] - \mathbb{E}_{\mathcal{A},\mathcal{S}} \left[F_{\mathcal{S}} \left(\boldsymbol{\theta}_{\mathcal{S}}^{*} \right) \right] \\ = \underbrace{\mathbb{E} \left[F_{\mathcal{D}} \left(\boldsymbol{\theta}_{\mathcal{A},\mathcal{S}} \right) - F_{\mathcal{S}} \left(\boldsymbol{\theta}_{\mathcal{A},\mathcal{S}} \right) \right]}_{\varepsilon_{\text{gen}}} + \underbrace{\mathbb{E} \left[F_{\mathcal{S}} \left(\boldsymbol{\theta}_{\mathcal{A},\mathcal{S}} \right) - F_{\mathcal{S}} \left(\boldsymbol{\theta}_{\mathcal{S}}^{*} \right) \right]}_{\varepsilon_{\text{opt}}}$$

where $\boldsymbol{\theta}_{\mathcal{S}}^{\star} \in \operatorname{argmin}_{\boldsymbol{\theta}} F_{\mathcal{S}}(\boldsymbol{\theta})$ is the optimum of $F_{\mathcal{S}}$. The optimization error $\varepsilon_{\text{opt}} = \mathbb{E}_{\mathcal{A},\mathcal{S}} \left[F_{\mathcal{S}} \left(\boldsymbol{\theta}_{\mathcal{A},\mathcal{S}} \right) - F_{\mathcal{S}} \left(\boldsymbol{\theta}_{\mathcal{S}}^{\star} \right) \right]$ denotes the difference between the exact optimum $\boldsymbol{\theta}_{\mathcal{S}}^{\star}$ and the estimated solution $\boldsymbol{\theta}_{\mathcal{A},\mathcal{S}}$, and generalization error $\varepsilon_{\text{gen}} = \mathbb{E}_{\mathcal{A},\mathcal{S}} \left[F_{\mathcal{D}} \left(\boldsymbol{\theta}_{\mathcal{A},\mathcal{S}} \right) - F_{\mathcal{S}} \left(\boldsymbol{\theta}_{\mathcal{A},\mathcal{S}} \right) \right]$ measures the effects of minimizing empirical risk instead of population risk.

Uniform Stability and Generalization. One popular approach to analyze generalization error ε_{gen} of algorithm \mathcal{A} is uniform stability (Hardt et al., 2016; Zhang et al., 2022).

Lemma 4.1. (Uniform Stability and Generalization Error) (Hardt et al., 2016) We say a randomized algorithm \mathcal{A} is ϵ -uniformly stable if for all datasets $S \sim \mathcal{D}$ and $S' \sim \mathcal{D}$ where S and S' differ in at most one sample,

$$\sup_{(\boldsymbol{x},\boldsymbol{y})\sim\mathcal{D}} \mathbb{E}_{\mathcal{A}}\left[\ell\left(\boldsymbol{x},\boldsymbol{y};\boldsymbol{\theta}_{\mathcal{A},\mathcal{S}}\right) - \ell\left(\boldsymbol{x},\boldsymbol{y};\boldsymbol{\theta}_{\mathcal{A},\mathcal{S}'}\right)\right] \leq \epsilon. \quad (9)$$

Moreover, if \mathcal{A} is ϵ -uniformly stable, then its generalization error ε_{gen} , which is defined as $\varepsilon_{gen} = |\mathbb{E}_{\mathcal{A},\mathcal{S}} [F_{\mathcal{D}}(\boldsymbol{\theta}_{\mathcal{A},\mathcal{S}}) - F_{\mathcal{S}}(\boldsymbol{\theta}_{\mathcal{A},\mathcal{S}})]|$ satisfying $\varepsilon_{gen} \leq \epsilon$.

FL generalization error via uniform stability. Here we aim to use uniform stability to upper bound the FL generalization error. Suppose given *n* samples $S_i = \{z_{i,1}, z_{i,2}, \dots, z_{i,n}\}$, where $z_{i,j} = (x_{i,j}, y_{i,j})$ is sampled from an unknown distribution \mathcal{D}_i , and suppose the generated sample set $S_i^{(j)} = \{z'_{i,1}, z'_{i,2}, \dots, z'_{i,n}\} = \{z_{i,1}, z_{i,2}, \dots, z_{i,j-1}, z'_{i,j}, z_{i,j+1} \dots, z_{i,n}\}$ which only differs from the set S_i with the *j*-th sample. $S = \{S_i\}_{i=1}^m, S^{(j_i)} = \{S_1, \dots, S_{i-1}, S_i^{(j)}, S_{i+1}, \dots, S_m\}$. One usually analyzes the stability of an algorithm by replacing one sample in S by another sample from \mathcal{D} , and we get $S^{(j_i)}$. Then based on these two sets, one can train the algorithm on S to obtain different solutions θ of the function $F_S(\theta)$. When using $S^{(j_i)}$, we adopt θ'_t and $\theta^{t,\prime}_{i,k}$ to denote their corresponding versions θ_t and $\theta^t_{i,k}$ (θ_t and $\theta^t_{i,k}$ are trained on S in Algorithm 1).

5. Theoretical Analysis

5.1. Generalization Error Analysis

We analyze generalization based on following assumptions: Assumption 1. The loss function $l(\cdot, z)$ is μ -strongly convex for any sample z, $l(\theta; z) \geq l(\theta'; z) + \langle \nabla l(\theta'; z), \theta - \theta' \rangle + \frac{\mu}{2} ||\theta - \theta'||^2$, for any z, θ, θ' . Assumption 2. The loss function $l(\cdot, z)$ is L Lipschitz continous, $|l(\theta; z) - l(\theta'; z)| \leq L ||\theta - \theta'||$, for any z, θ, θ' . Assumption 3. There exists a constant $\sigma > 0$ such that for any $\theta, i \in [m]$, and $z_i \sim D_i$, $\mathbb{E} \left[||\nabla l(\theta; z_i) - \nabla F_i(\theta)||^2 \right] \leq \sigma^2$.

Assumption 4. The loss function $l(\cdot, z)$ is β -smooth, $\|\nabla l(\theta; z) - \nabla l(\theta'; z)\| \leq \beta \|\theta - \theta'\|$, for any z, θ, θ' .

Assumption 5 (data heterogeneity). Given $i \in [m]$, there exists a constant $\sigma_g > 0$, for any θ we have $\frac{1}{m} \sum_{i=1}^{m} \|\nabla F_i(\theta) - \nabla F(\theta)\| \le \sigma_g$.

Assumption 2, as in (Zhang et al., 2022), links model perturbation to stability. Assumptions 3 and 5 in our study address data heterogeneity and algorithm convergence impacts.

Theorem 5.1 (Generalization Error). Assuming all clients participate in each round with Option I:

Strongly convex: Under Assumptions 1-5, suppose loss $\ell(\boldsymbol{x}, \boldsymbol{y}; \boldsymbol{\theta})$ is μ -strongly convex. By setting $\eta_k^t \leq \frac{1}{\beta KT}$, $\tilde{b} = \frac{K-1}{\beta KT}$

$$1 + \left(\frac{\mu}{(\beta+\mu)K}\right)^{K-1} \frac{1}{T}, \text{ the generalization error satisfies:}$$

$$\begin{aligned} & \text{FedSWA: } \varepsilon_{gen} \leq \frac{1}{mn\beta} e^{-(\beta+\mu)T} \left(bL + b\delta_g + b\delta \right). \end{aligned}$$

$$\begin{aligned} & \text{FedMoSWA: } \varepsilon_{gen} \leq \frac{2L}{mn\beta} e^{1-\frac{\mu}{(\beta+\mu)T}} \left(\tilde{b}L + \sigma_g + \tilde{b}\sigma \right). \end{aligned}$$

Non-convex: Under Assumptions 2-5, assume $\ell(x, y; \theta)$ is β -smooth. Together with $\eta_k^t \leq \frac{1}{\beta KT}$, $\tilde{c} = 1 + (2 + \frac{1}{KT})^{K-1} \frac{1}{T}$, the generalization error satisfies: FedSWA: $\varepsilon_{gen} \leq \frac{2L}{mn\beta} e^{\frac{1}{T}+1} (\tilde{c}L + \tilde{c}\sigma_g + \tilde{c}\sigma)$. FedMoSWA: $\varepsilon_{gen} \leq \frac{2L}{mn\beta} e^{\frac{1}{T}+1} (\tilde{c}L + \sigma_g + \tilde{c}\sigma)$.

Theorem 5.1 shows that for strongly convex and non-convex problems, ε_{opt} has three terms: the first bias term has the Lipschitz continous constant *L*, the second bias term characterizes the effect of data heterogeneity σ_g , and the third term reveals the impact of the stochastic gradient noise σ .

For strongly convex problems, Theorem 5.1 shows that ε_{gen} of FedMoSWA can be upper bounded by $\mathcal{O}(\frac{2L}{mn\beta}e^{1-\frac{\mu}{(\beta+\mu)T}}(\tilde{b}L+\sigma_g+\tilde{b}\sigma))$, which is better than that of FedSWA, $\mathcal{O}(\frac{2L}{mn\beta}e^{1-\frac{\mu}{(\beta+\mu)T}}(\tilde{b}L+\tilde{b}\sigma_g+\tilde{b}\sigma))$ ($\tilde{b}\sigma_g \gg \sigma_g$). Compared to FedSWA, FedMoSWA reduces the effect of data heterogeneity on the generalization error. In addition, we can observe that the generalization ability of the algorithm is affected by the number of clients m and the amount of data mn on the clients. More client participation (m) can effectively enhance FL generalization. Similarly, to achieve smaller generalization error, clients should use large amount of data (mn). Reducing the number of local iterations K improves generalization capabilities. And greater data heterogeneity impairs federated learning generalization.

For the non-convex generalization error, Theorem 5.1 shows that ε_{gen} of FedMoSWA is better than FedSWA's $\mathcal{O}(\frac{2L}{mn\beta}e^{\frac{1}{T}+1}(\tilde{c}L+\tilde{c}\sigma_g+\tilde{c}\sigma))$, and $\tilde{c}\sigma_g \gg \sigma_g$ in Table 1. FedMoSWA and FedSWA generalization error is better than FedSAM in Table 1.

5.2. Optimization Error Analysis

Theorem 5.2 (Optimization Error of FedMoSWA). For β smooth functions $\{F_i\}$, which satisfy Assumptions 6-9, and are the same as in the SCAFFOLD (Karimireddy et al., 2020) algorithm (see the Appendix for details), the output of FedMoSWA has expected error smaller than ϵ .

Strongly convex: $\eta_k^t \leq \min\left(\frac{1}{\beta K \alpha}, \frac{s}{\mu m K \alpha}\right), T \geq \max\left(\frac{\beta}{\mu}, \frac{m}{s}\right)$ then

$$\mathcal{O}\left(\frac{\sigma^2}{\mu TKs}\left(1+\frac{s}{\alpha^2}\right)+\frac{m\mu}{s}D^2\exp\left(-\left\{\frac{s}{m}+\frac{\mu}{\beta}\right\}T\right)\right)$$

Non-convex: $\eta_k^t \leq \frac{1}{K\alpha\beta} \left(\frac{s}{m}\right)^{\frac{2}{3}}$, and $T \geq 1$, then

$$\mathcal{O}\left(\frac{\sigma\sqrt{F}}{\sqrt{TKs}}\left(\sqrt{1+\frac{m}{\alpha^2}}\right) + \frac{\beta F}{T}\left(\frac{m}{s}\right)^{\frac{2}{3}}\right).$$

$$\mathcal{O}^2 := \|\boldsymbol{\theta}^0 - \boldsymbol{\theta}^\star\|^2 + \frac{1}{T}\sum^m \|\boldsymbol{e}^0 - \nabla F_s\|^2$$

Here $D^2 := \left\| \boldsymbol{\theta}^0 - \boldsymbol{\theta}^\star \right\|^2 + \frac{1}{2m\beta^2} \sum_{i=1}^m \left\| \boldsymbol{c}_i^0 - \nabla F_i\left(\boldsymbol{\theta}^\star\right) \right\|^2$, $F := F\left(\boldsymbol{\theta}^0\right) - F\left(\boldsymbol{\theta}^\star\right)$.



Figure 3: Convergence plots on CIFAR100 with ResNet-18.

Theorem 5.2 shows that for strongly convex and non-convex problems, optimization error ε_{opt} has two terms: the first term reveals the impact of the stochastic gradient noise, and the second bias term characterizes the effect of initialization θ_0 . The optimization error ε_{opt} is not affected by the data heterogeneous parameter σ_g and FedMoSWA converges faster as α increases. Convergence is also accelerated when the number of client *s* increases. Its convergence is accelerated when the number of local iteration *K* increases.

For strongly convex problems in Theorem 5.2 shows that ε_{opt} of FedMoSWA is $\mathcal{O}\left(\frac{\sigma^2}{\mu T K s}\left(1+\frac{s}{\alpha^2}\right)+\frac{m\mu}{s}D^2\exp\left(-\left\{\frac{s}{m}+\frac{\mu}{\beta}\right\}T\right)\right)$. For α , it can be observed that with the increase of α , ε_{opt} becomes smaller. The ε_{opt} of FedMoSWA for non-convex problems is $\mathcal{O}\left(\frac{\sigma\sqrt{F}}{\sqrt{TKs}}\left(\sqrt{1+\frac{m}{\alpha^2}}\right)+\frac{\beta F}{T}\left(\frac{m}{s}\right)^{\frac{2}{3}}\right)$, which is not affected by the data heterogeneous parameter σ_g and is better than FedSWA, FedSAM and MoFedSAM, as shown in Table 1. We can prove that the ε_{opt} of FedMoSWA is faster than SCAFFOLD's $\mathcal{O}\left(\frac{\sigma\sqrt{F}}{\sqrt{TKs}}\left(\sqrt{1+m}\right)+\frac{\beta F}{T}\left(\frac{m}{s}\right)^{\frac{2}{3}}\right)$ (Karimireddy et al., 2020).

6. Experiments

6.1. Experimental Settings

Datasets: We evaluate our algorithms on CIFAR10, CI-FAR100 (?), Tiny ImageNet (Le & Yang, 2015). For Non-IID data setup, we simulate the data heterogeneity by sampling the label ratios from Dirichlet distribution (Hsu et al., 2019). Dirichlet-0.1 means very high data heterogeneity and Dirichlet-0.6 implies average data heterogeneity.

Models: To test the robustness of our algorithms, we use various classifiers including LeNet-5 (LeCun et al., 2015), VGG-11 (Simonyan & Zisserman, 2015), ResNet-18 (He et al., 2016), vision Transformer (ViT-Base) (Dosovitskiy et al., 2021).

Methods: We compare FedSWA and FedMoSWA with many SOTA FL methods, including FedAvg (McMahan et al., 2017), FedDyn (Durmus et al., 2021), SCAFFOLD (Karimireddy et al., 2020), FedAvgM (Hsu et al., 2019), FedSAM (Qu et al., 2022), MoFedSAM (Qu et al., 2022), FedLESAM (Fan et al., 2024), FedASAM (Caldarola et al., 2022), and FedACG (Kim et al., 2024).

Hyper-parameter Settings: The number of clients is 100, batch size B = 50, local epoch E = 5, K = 50 and client selection rate C = 0.1. We set the grid search range of client learning rate by $\eta_l \in \{10^{-3}, 3 \times 10^{-3}, ..., 10^{-1}, 3 \times 10^{-1}\}$. The learning rate decay per round of communication is 0.998, and the total is T = 1,000. Specifically, we set $\rho = 0.1$, $\alpha = 1.5$, $\gamma = 0.2$ for FedMoSWA and FedSWA.

6.2. Results on Nonconvex Problems

Results on Convolutional Neural Network: From Table 2 and Figure 3, we have the following observations: (i) FedSWA can improve generalization ability by SWA, and it is better than FedSAM. Our FedMoSWA is better than FedSAM and its variant, MoFedSAM, in terms of generalization. (ii) FedMoSWA further improves the generalization ability by momentum-based stochastic control. Compared to the stochastic control algorithm, SCAFFOLD, our FedMoSWA demonstrates better generalization ability than SCAFFOLD. Compared to the momentum acceleration algorithm FedACG, FedMoSWA shows significantly better generalization combining momentum and variance reduction techniques. Compared to FedSAM, MoFedSAM, SCAFFOLD and FedACG can not achieve desirable flatness. FedMoSWA further minimizes the global sharpness and achieves much flatter loss landscape in Figure 4.

Impact of heterogeneity: From Table 3, we have the following observations: For different data heterogeneities, our FedMoSWA generalizes well. Especially when the degree of data heterogeneity is high (Dirichle-0.1), the final test results of our FedMoSWA are much better than all other algorithms. For example, as Non-IID levels increasing, FedMoSWA achieves higher test accuracy 10.4%, 8.5% and 6.2% and saving communication round than MoFedSAM on the CIFAR-100 dataset.

Impact of α , γ : Figure 5 (a) compares the effects of α on FL aggregation. As the α increases, the acceleration effect of our FedMoSWA becomes more obvious. However, when $\alpha > 1.5$, the performance of the algorithm decreases with increasing α . We can find that the final performance of the algorithm is best when $\alpha = 1.5$. Figure 5 (b) compares the effects of γ in the same settings. When $\gamma = 0.05$, FedMoSWA converges very slowly, and the final algorithm has low accuracy. As the γ increases, the acceleration effect of our FedMoSWA becomes more apparent. The final performance of the algorithm is best when $\gamma = 0.2$.

Impact of ρ : Table 7 and Figure 3 compare the effects of ρ in the local training of FedSWA and MoFedSWA. As ρ decreases, the performance of the algorithms increases. We

		CIFAR10					CIFAR100						
Method	LeN	LeNet-5		VGG-11		ResNet-18		LeNet-5		VGG-11		ResNet-18	
	Acc.(%) 1000R	Rounds 78%	Acc.(%) 1000R	Rounds 80%	Acc.(%) 1000R	Rounds 85%	Acc.(%) 1000R	Rounds 52%	Acc.(%) 1000R	Rounds 53%	Acc.(%) 1000R	Rounds 55%	
FedAvg	$79.6_{\pm 0.2}$	574	$84.0_{\pm 0.3}$	313	$86.0_{\pm 0.2}$	542	$41.2_{\pm 0.2}$	1000+	$48.9_{\pm 0.4}$	1000+	$54.2_{\pm 0.2}$	1000+	
FedDyn	$80.8_{\pm 0.3}$	390	$83.1_{\pm 0.2}$	557	$84.6_{\pm 0.2}$	1000+	$40.8_{\pm 0.5}$	1000+	$45.2_{\pm 0.6}$	1000+	$46.5_{\pm 0.6}$	1000+	
SCAFFOLD	$80.7_{\pm 0.3}$	799	$86.9_{\pm 0.2}$	272	$85.9_{\pm 0.6}$	543	$50.3_{\pm 0.2}$	1000+	$47.9_{\pm 0.2}$	1000+	$54.1_{\pm 0.2}$	1000+	
FedSAM	$79.9_{\pm 0.2}$	821	$85.5_{\pm 0.2}$	227	$83.6_{\pm 0.2}$	1000+	$47.9_{\pm 0.1}$	1000+	$51.7_{\pm 0.3}$	1000 +	$47.8_{\pm 0.1}$	1000+	
MoFedSAM	$82.5_{\pm 0.1}$	483	$87.4_{\pm 0.4}$	345	$87.0_{\pm 0.3}$	569	$49.3_{\pm 0.2}$	1000+	$1.0_{\pm 0.0}$	1000+	$60.1_{\pm 0.2}$	603	
FedLESAM	$79.7_{\pm 0.2}$	496	$83.5_{\pm 0.4}$	169	$89.0_{\pm 0.1}$	480	$41.3_{\pm 0.3}$	1000+	$49.2_{\pm 0.3}$	1000+	$52.1_{\pm 0.1}$	1000+	
FedASAM	$81.0_{\pm 0.2}$	371	$86.3_{\pm 0.3}$	169	$88.3_{\pm 0.2}$	542	$42.1_{\pm 0.3}$	1000+	$53.4_{\pm 0.2}$	908	$49.8_{\pm 0.3}$	1000+	
FedACG	$82.9_{\pm 0.1}$	321	$84.5_{\pm 0.4}$	249	$90.3_{\pm 0.4}$	237	$52.5_{\pm 0.2}$	729	$50.9_{\pm 0.2}$	1000+	$61.7_{\pm 0.4}$	518	
FedSWA	$81.5_{\pm 0.2}$	591	$86.5_{\pm 0.1}$	169	$89.5_{\pm 0.1}$	538	$48.5_{\pm 0.1}$	1000+	$52.5_{\pm 0.4}$	1000 +	$59.8_{\pm 0.3}$	574	
FedMoSWA	$83.8_{\pm 0.1}$	301	$88.1_{\pm 0.2}$	137	$91.2_{\pm 0.1}$	265	$53.8_{\pm0.1}$	622	$61.7_{\pm 0.4}$	389	$67.9_{\pm 0.4}$	330	

Table 2: Comparison of the accuracies at the target rounds (1000R) and the communication round to reach target test accuracy (Acc.%) of each algorithm on Dirichlet-0.6 of the CIFAR10 and CIFAR100 datasets.



Figure 4: Global training loss surfaces for FedSWA, FedMoSWA and other baselines on Dirichlet-0.1 of CIFAR100 dataset with ResNet-18. FedMoSWA could approach a more general and flat loss landscape which high generalization in FL.



Figure 5: Convergence plots for FedMoSWA with different α and γ on CIFAR100 dataset using ResNet-18.

can find that the final performance of the algorithms is best when $\rho = 0.1$. When $\rho = 1$ and $\alpha = 1$, our FedSWA is equal to FedAvg, which means cyclical learning rate works in FL. When $\rho = 1$ and $\alpha = 1$, our FedMoSWA is equal to FedAvg with momentum variance reduction without SWA (65.9%) which means momentum variance reduction works in FL and it is better than SCAFFOLD (52.3%).

Results on Vision Transformer: In Table 3, we use the ViT-Base model on the Tiny ImageNet dataset. We used the

pretrained model from the official website, and conducted many experiments. All the experimental results verify that our FedMoSWA can achieve excellent generalization on both the vision Transformer model and big datasets.

6.3. Ablation Study

We conducted full client participation experiments in 5. Here, FedMo denotes FedMoSWA without SWA (i) but only with Momentum Stochastic Control (ii). FedMo (62.5%) outperforms SCAFFOLD (59.9%) by 2.6%. Even if the client is full participation, our method is still better than SCAFFOLD.

When Dirichlet-0.6, 10% participation clients, FedMo using only momentum stochastic control achieves 65.9%, higher than SCAFFOLD (54.1%), demonstrating that momentum variance reduction mitigates SCAFFOLD's variance reduction delay. By combining SWA and momentum variance reduction, FedMoSWA achieves 67.9% (+13.7%), showing that both SWA (+5.6%) and momentum control (+11.7%) are effective, and momentum control has greater impact. Moreover, FedSWA is a simple algorithm like FedSAM and can be combined with other techniques.

		CIFAR100 (ResNet-18)					Tiny ImageNet (ViT-Base)					
Method	Dirichl	et-0.1	Dirichl	et-0.3	Dirichl	et-0.6	Dirichl	et-0.1	Dirichl	et-0.3	Dirichl	et-0.6
	Acc.(%) 1000R	Rounds 55%	Acc.(%) 1000R	Rounds 55%	Acc.(%) 1000R	Rounds 55%	Acc.(%) 400R	Rounds 70%	Acc.(%) 400R	Rounds 70%	Acc.(%) 400R	Rounds 70%
FedAvg	$45.8_{\pm 0.3}$	1000+	$52.5_{\pm 0.3}$	1000+	$54.2_{\pm 0.2}$	1000+	$70.9_{\pm 0.1}$	258	$71.8_{\pm 0.1}$	223	$72.8_{\pm 0.1}$	208
FedDyn	$45.8_{\pm 0.2}$	1000+	$45.9_{\pm 0.3}$	1000+	$46.5_{\pm 0.2}$	1000+	$67.5_{\pm 0.3}$	400+	$68.2_{\pm 0.3}$	400+	$69.3_{\pm 0.3}$	400+
SCAFFOLD	$44.3_{\pm 0.3}$	1000+	$50.3_{\pm 0.3}$	1000 +	$52.3_{\pm 0.2}$	1000+	$71.6_{\pm 0.1}$	202	$72.5_{\pm 0.1}$	192	$73.1_{\pm 0.2}$	169
FedSAM	$40.1_{\pm 0.4}$	1000+	$49.0_{\pm 0.3}$	1000 +	$51.9_{\pm 0.5}$	1000+	$71.4_{\pm 0.2}$	212	$72.2_{\pm 0.2}$	194	$72.9_{\pm 0.4}$	180
MoFedSAM	$51.5_{\pm 0.2}$	1000+	$57.5_{\pm 0.2}$	770	$60.1_{\pm 0.1}$	603	$71.6_{\pm 0.4}$	229	$72.4_{\pm 0.3}$	214	$72.5_{\pm 0.4}$	209
FedLESAM	$48.7_{\pm 0.2}$	1000+	$53.3_{\pm 0.4}$	1000+	$52.1_{\pm 0.1}$	1000+	$71.9_{\pm 0.3}$	210	$72.1_{\pm 0.2}$	188	$72.5_{\pm 0.3}$	182
FedASAM	$47.7_{\pm 0.3}$	1000+	$46.6_{\pm 0.2}$	1000 +	$49.8_{\pm 0.1}$	1000+	$69.2_{\pm 0.3}$	400+	$71.3_{\pm 0.2}$	234	$72.1_{\pm 0.3}$	196
FedACG	$52.2_{\pm 0.4}$	1000+	$57.7_{\pm 0.2}$	717	$61.7_{\pm 0.4}$	518	$66.2_{\pm 0.2}$	400+	$68.5_{\pm0.1}$	400+	$70.2_{\pm 0.3}$	386
FedSWA	$50.3_{\pm 0.3}$	1000+	$55.5_{\pm 0.4}$	889	$59.8_{\pm 0.3}$	574	$71.9_{\pm 0.3}$	199	$72.6_{\pm 0.2}$	179	$73.2_{\pm 0.2}$	168
FedMoSWA	$61.9_{\pm 0.5}$	577	$66.2_{\pm 0.4}$	468	$67.9_{\pm 0.4}$	330	$73.8_{\pm 0.3}$	161	$74.4_{\pm 0.3}$	152	$74.7_{\pm 0.1}$	144

Table 3: Comparison of each algorithm on the CIFAR100 and Tiny ImageNet datasets with different data heterogeneity.

Table 4: Testing Acc (%) of FedSWA and MoFedSWA with different ρ on CIFAR100 (ResNet-18, Dirichlet-0.6).

ρ	0.1	0.2	0.3	0.4	0.7	0.8	1.0
FedSWA	59.8	58.5	57.9	56.6	55.3	54.6	54.2
FedMoSWA	67.9	67.3	67.0	66.8	66.5	66.3	65.9

Table 5: Full client participation on CIFAR-100 and ResNet-18 with 10 clients and 300 rounds.

Algorithm	Accuracy (%)	Improvement (%)
FedAvg	58.2	-
SCAFFOLD	59.9	+1.7
FedSAM	48.3	-9.9
MoFedSAM	37.9	-20.3
FedLESAM	59.2	+1.0
FedSWA (i)	60.2	+2.0
FedMo (ii)	62.5	+4.3
FedMoSWA (i+ii)	63.2	+5.0

7. Conclusion

This study investigated generalization concerns in FL and proposed two new algorithms, FedSWA and FedMoSWA, for the setting of highly heterogeneous data. In particular, FedMoSWA addresses the essential issue of generalization with highly heterogeneous data. By developing an analytical framework based on algorithmic stability, we theoretically established the superiority of FedSWA and FedMoSWA in terms of generalization error. Our FedMoSWA mitigates the effect of data heterogeneity on the generalization error, though it does not completely eliminate it. Our future works Table 6: 10% client participation on CIFAR-100 and ResNet-18 with 100 clients and 1000 rounds.

Algorithm	Accuracy (%)	Improvement (%)
FedAvg	54.2	-
SCAFFOLD	54.1	-0.1
FedSAM	47.8	-6.4
FedSWA (i)	59.8	+5.6
FedMo (ii)	65.9	+11.7
FedMoSWA (i+ii)	67.9	+13.7

may consider an approach that significantly eliminates the effect of high heterogeneity on the generalization error.

Acknowledgments

This work was supported in the National Natural Science Foundation of China (Nos. 62276182), Peng Cheng Lab Program (No. PCL2023A08), Tianjin Natural Science Foundation (Nos. 24JCYBJC01230, 24JCYBJC01460), and Tianjin Municipal Education Commission Research Plan (No. 2024ZX008).

Impact Statement

The proposed Stochastic Weight Averaging (SWA) and Momentum-Based Stochastic Controlled Weight Averaging (FedMoSWA) algorithms for federated learning aim to address data heterogeneity and improve model generalization. This paper presents one work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here

References

- Antunes, R. S., André da Costa, C., Küderle, A., Yari, I. A., and Eskofier, B. Federated learning for healthcare: Systematic review and architecture proposal. ACM Transactions on Intelligent Systems and Technology (TIST), 13 (4):1–23, 2022.
- Bian, J., Wang, L., Zhang, L., and Xu, J. Lora-fair: Federated lora fine-tuning with aggregation and initialization refinement. arXiv preprint arXiv:2411.14961, 2024.
- Bian, J., Wang, L., Zhang, L., and Xu, J. Fedalt: Federated fine-tuning through adaptive local training with rest-ofworld lora. arXiv preprint arXiv:2503.11880, 2025.
- Bottou, L. Large-scale machine learning with stochastic gradient descent. In Proceedings of COMPSTAT'2010: 19th International Conference on Computational StatisticsParis France, August 22-27, 2010 Keynote, Invited and Contributed Papers, pp. 177–186. Springer, 2010.
- Bousquet, O. and Elisseeff, A. Stability and generalization. Journal of Machine Learning Research, 2:499–526, 2002.
- Bubeck, S. et al. Convex optimization: Algorithms and complexity. *Foundations and Trends*® *in Machine Learning*, 8(3-4):231–357, 2015.
- Byrd, D. and Polychroniadou, A. Differentially private secure multi-party computation for federated learning in financial applications. In *Proceedings of the First ACM International Conference on AI in Finance*, pp. 1–9, 2020.
- Caldarola, D., Caputo, B., and Ciccone, M. Improving generalization in federated learning by seeking flat minima. In *European Conference on Computer Vision*, pp. 654–672. Springer, 2022.
- Cutkosky, A. and Orabona, F. Momentum-based variance reduction in non-convex sgd. In Advances in Neural Information Processing Systems 32, (NeurIPS 2019).Vancouver(CA).8-14 December 2019, 2020.
- Dai, R., Yang, X., Sun, Y., Shen, L., Tian, X., Wang, M., and Zhang, Y. Fedgamma: Federated learning with global sharpness-aware minimization. *IEEE Transactions on Neural Networks and Learning Systems*, 2023.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. An image is worth 16x16 words: Transformers for image recognition at scale. *Proceedings of International Conference on Learning Representations*, 2021.
- Durmus, A. E., Yue, Z., Ramon, M., Matthew, M., Paul, W., and Venkatesh, S. Federated learning based on dynamic regularization. In *International conference on learning representations*, 2021.

- Fan, Z., Hu, S., Yao, J., Niu, G., Zhang, Y., Sugiyama, M., and Wang, Y. Locally estimated global perturbations are better than local perturbations for federated sharpness-aware minimization. *In Proceedings of the* 41st International Conference on Machine Learning, pp. 235:12858-12881, 2024.
- Foret, P., Kleiner, A., Mobahi, H., and Neyshabur, B. Sharpness-aware minimization for efficiently improving generalization. *In Proceedings of International Conference on Learning Representations*, 2021.
- Gotmare, A., Keskar, N. S., Xiong, C., and Socher, R. A closer look at deep learning heuristics: Learning rate restarts, warmup and distillation. *In Proceedings of International Conference on Learning Representations*, 2019.
- Hardt, M., Recht, B., and Singer, Y. Train faster, generalize better: Stability of stochastic gradient descent. In *International Conference on Machine Learning*, pp. 1225–1234. PMLR, 2016.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.
- 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.
- Hu, X., Li, S., and Liu, Y. Generalization bounds for federated learning: Fast rates, unparticipating clients and unbounded losses. In *International Conference on Learning Representations*, 2023.
- Izmailov, P., Podoprikhin, D., Garipov, T., Vetrov, D., and Wilson, A. G. Averaging weights leads to wider optima and better generalization. *In Proceedings of Conference* on Uncertainty in Artificial Intelligence, 2018.
- Kaddour, J. Stop wasting my time! saving days of imagenet and bert training with latest weight averaging. In *Has it Trained Yet? NeurIPS 2022 Workshop*.
- Kaddour, J., Liu, L., Silva, R., and Kusner, M. J. When do flat minima optimizers work? *Advances in Neural Information Processing Systems*, 35:16577–16595, 2022.
- 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.
- Kim, G., Kim, J., and Han, B. Communication-efficient federated learning with accelerated client gradient. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12385–12394, 2024.

- Le, Y. and Yang, X. Tiny imagenet visual recognition challenge. *Technical Report*, 2015.
- LeCun, Y. et al. Lenet-5, convolutional neural networks. URL: http://yann. lecun. com/exdb/lenet, 20(5):14, 2015.
- Lee, T. and Yoon, S. W. Rethinking the flat minima searching in federated learning. In *Forty-first International Conference on Machine Learning*, pp. 27037–27071, 2024.
- Li, T., Sahu, A. K., Zaheer, M., Sanjabi, M., Talwalkar, A., and Smith, V. Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems*, 2:429–450, 2020.
- Liu, J., Shang, F., Liu, Y., Liu, H., Li, Y., and Gong, Y. Fedbcgd: Communication-efficient accelerated block coordinate gradient descent for federated learning. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pp. 2955–2963, 2024.
- McMahan, B., Moore, E., Ramage, D., Hampson, S., and y Arcas, B. A. Communication-efficient learning of deep networks from decentralized data. In *Proceedings of International Conference on Artificial Intelligence and Statistics*, pp. 1273–1282. PMLR, 2017.
- Patel, K. K., Wang, L., Woodworth, B. E., Bullins, B., and Srebro, N. Towards optimal communication complexity in distributed non-convex optimization. *Advances in Neural Information Processing Systems*, 35:13316–13328, 2022.
- Polyak, B. T. and Juditsky, A. B. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control* and Optimization, 30(4):838–855, 1992.
- Qu, Z., Li, X., Duan, R., Liu, Y., Tang, B., and Lu, Z. Generalized federated learning via sharpness aware minimization. In *International Conference on Machine Learning*, pp. 18250–18280. PMLR, 2022.
- Reddi, S., Charles, Z., Zaheer, M., Garrett, Z., Rush, K., Konečný, J., Kumar, S., and McMahan, H. B. Adaptive federated optimization. *In Proceedings of International Conference on Learning Representations*, 2021.
- Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., Bakas, S., Galtier, M. N., Landman, B. A., Maier-Hein, K., et al. The future of digital health with federated learning. *NPJ Digital Medicine*, 3(1):119, 2020.
- Rogers, W. H. and Wagner, T. J. A finite sample distributionfree performance bound for local discrimination rules. *The Annals of Statistics*, pp. 506–514, 1978.
- Russo, D. and Zou, J. How much does your data exploration overfit? controlling bias via information usage.

IEEE Transactions on Information Theory, 66(1):302–323, 2019.

- Sefidgaran, M., Chor, R., and Zaidi, A. Rate-distortion theoretic bounds on generalization error for distributed learning. *Advances in Neural Information Processing Systems*, 35:19687–19702, 2022.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. In 3rd International Conference on Learning Representations (ICLR 2015). Computational and Biological Learning Society, 2015.
- Smale, S. and Zhou, D.-X. Learning theory estimates via integral operators and their approximations. *Constructive Approximation*, 26(2):153–172, 2007.
- Smith, L. N. Cyclical learning rates for training neural networks. In 2017 IEEE Winter Conference on Applications of Computer Vision (WACV), pp. 464–472. IEEE, 2017.
- Sun, Y., Shen, L., Huang, T., Ding, L., and Tao, D. Fedspeed: Larger local interval, less communication round, and higher generalization accuracy. *In Proceedings of International Conference on Learning Representations*, 2023.
- Sun, Z., Niu, X., and Wei, E. Understanding generalization of federated learning via stability: Heterogeneity matters. In *International Conference on Artificial Intelligence and Statistics*, pp. 676–684. PMLR, 2024.
- Vapnik, V. N., Vapnik, V., et al. Statistical learning theory. 1998.
- Wang, L., Bian, J., Zhang, L., Chen, C., and Xu, J. Taming cross-domain representation variance in federated prototype learning with heterogeneous data domains. *Ad*vances in Neural Information Processing Systems, 37: 88348–88372, 2024.
- Wu, W., Chen, Z., Qiu, X., Song, S., Huang, X., Ma, F., and Xiao, J. Llm-enhanced multimodal fusion for cross-domain sequential recommendation. *arXiv preprint arXiv:2506.17966*, 2025a.
- Wu, W., Qiu, X., Song, S., Chen, Z., Huang, X., Ma, F., and Xiao, J. Prompt categories cluster for weakly supervised semantic segmentation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 3198– 3207, 2025b.
- Yuan, Z., Yan, Y., Jin, R., and Yang, T. Stagewise training accelerates convergence of testing error over sgd. Advances in Neural Information Processing Systems, 32, 2019.

- Zhang, M., Lucas, J., Ba, J., and Hinton, G. E. Lookahead optimizer: k steps forward, 1 step back. *Advances in Neural Information Processing Systems*, 32, 2019.
- Zhang, X., Hong, M., Dhople, S., Yin, W., and Liu, Y. Fedpd: A federated learning framework with adaptivity to non-iid data. *IEEE Transactions on Signal Processing*, 69:6055–6070, 2021.
- Zhang, Y., Zhang, W., Bald, S., Pingali, V., Chen, C., and Goswami, M. Stability of sgd: Tightness analysis and improved bounds. In *Proceedings of Conference on Uncertainty in Artificial Intelligence*, pp. 2364–2373. PMLR, 2022.

A. Background

In this section, we provide a concise overview of Sharpness-Aware Minimization (SAM) and Stochastic Weight Averaging (SWA).

A.1. SAM: Overview

SAM seeks to identify a solution θ that lies within a region characterized by uniformly low training loss $\mathcal{L}_{\mathcal{D}}(\theta)$, specifically within a flat minimum. The sharpness of the training loss function is quantified as:

$$\max_{\|\epsilon\|_{p} \le \rho} \mathcal{L}_{\mathcal{D}}(\theta + \epsilon) - \mathcal{L}_{\mathcal{D}}(\theta)$$

Here, ρ is a hyper-parameter that determines the size of the neighborhood, and $p \in [1, \infty)$. SAM focuses on reducing the sharpness of the loss by addressing the following minmax objective:

$$\min_{\theta \in \mathbb{R}^d} \max_{\|\epsilon\|_p \le \rho} \mathcal{L}_{\mathcal{D}}(\theta + \epsilon) + \lambda \|\theta\|_2^2$$

where λ is a hyper-parameter that balances the regularization term. According to (Foret et al., 2021), the optimal choice for p is typically 2. Therefore, we use the ℓ_2 -norm for the maximization over ϵ and simplify the expression by omitting the regularization term. To solve the inner maximization problem $\epsilon^* \triangleq \arg \max_{\|\epsilon\|_2 \le \rho} \mathcal{L}(\theta + \epsilon)$, the authors suggest using a first-order approximation of $\mathcal{L}(\theta + \epsilon)$ around 0:

$$\epsilon^* \approx \arg \max_{\|\epsilon\|_2 \le \rho} \mathcal{L}_{\mathcal{D}}(\theta) + \epsilon^T \nabla_{\theta} \mathcal{L}_{\mathcal{D}}(\theta) = \rho \frac{\nabla_{\theta} \mathcal{L}_{\mathcal{D}}(\theta)}{\|\nabla_{\theta} \mathcal{L}_{\mathcal{D}}(\theta)\|_2} =: \hat{\epsilon}(\theta)$$

This approximation, which is computationally efficient, results in $\hat{\epsilon}(\theta)$ being essentially a scaled gradient of the current parameters θ . The sharpness-aware gradient is then defined as $\nabla_{\theta} \mathcal{L}_{\mathcal{D}}(\theta)|_{\theta + \hat{\epsilon}(\theta)}$ and is used to update the model as follows:

$$\theta_{t+1} \leftarrow \theta_t - \gamma \nabla_{\theta_t} \mathcal{L}_{\mathcal{D}} \left(\theta_t \right) |_{\theta_t + \hat{\epsilon}_t},$$

where γ is the learning rate and $\hat{\epsilon}_t = \hat{\epsilon}(\theta_t)$. This two-step process is repeated iteratively to solve the minmax objective. Essentially, SAM first performs a gradient ascent step to estimate the point $(\theta_t + \hat{\epsilon}_t)$ where the loss is approximately maximized, followed by a gradient descent step at θ_t using the computed gradient.

A.2. Stochastic Weight Averaging: Overview

SWA (Stochastic Weight Averaging) operates by averaging the weights proposed by Stochastic Gradient Descent (SGD) while employing a learning rate schedule to explore regions of the weight space that correspond to high-performing networks. At each step *i* within a cycle of length *c*, the learning rate is gradually reduced from an initial value γ_1 to a final value γ_2 . This learning rate adjustment is defined as:

$$\gamma(i) = (1 - t(i))\gamma_1 + t(i)\gamma_2, \quad t(i) = \frac{1}{c} (mod(i - 1, c) + 1)$$

Here, t(i) represents a linear interpolation factor that depends on the current step i and the cycle length c. If c = 1, the learning rate remains constant at γ_1 throughout the training process. However, for c > 1, the learning rate follows a cyclical schedule, oscillating between γ_1 and γ_2 over the course of the cycle.

Starting from a pre-trained model $f_{\hat{\theta}}$, SWA captures the model weights θ at the end of each cycle and computes their running average. This averaging process is expressed as:

$$\theta_{\text{SWA}} \leftarrow \frac{\theta_{\text{SWA}} \cdot n_{\text{models}} + \theta}{n_{\text{models}} + 1}$$

In this equation, n_{models} tracks the number of completed cycles, and θ_{SWA} represents the averaged weights. The final model, denoted as $f_{\theta_{\text{SWA}}}$, is obtained by aggregating the weights from multiple cycles, effectively smoothing the optimization trajectory and enhancing generalization performance.

By iteratively averaging the weights, SWA encourages the model to converge to a flatter region of the loss landscape, which is often associated with better generalization. This approach leverages the exploration of diverse weight configurations during training, ultimately leading to a more robust and stable model.

B. Implementation of the Experiments

B.1. Datasets

Table 7: A summary of CIFAR10/100, Tiny ImageNet datasets, including number of total images, number of classes, and the size of the images in the datasets.

Datasets	Total Images	Class	Image Size
CIFAR10	60,000	10	$3 \times 32 \times 32$
CIFAR100	60,000	100	$3 \times 32 \times 32$
Tiny ImageNet	100,000	200	$3 \times 224 \times 224$

CIFAR-10/100, Tiny ImageNet are popular benchmark datasets in federated learning. Data samples in CIFAR10 and CIFAR100 are colorful images of different categories with the resolution of 32×32 . There are 10 classes and each class has 6,000 images in CIFAR10. For CIFAR100, there are 100 classes and each class has 600 images. For Tiny ImageNet, there are 200 classes and each class has 500 images. As shown in Table 6, we summarize CIFAR10, CIFAR100, and Tiny ImageNet from the views of number of total images, number of classes, and the size of the images in the datasets.

B.2. More Loss Surface Visualization

Here we show more global loss surface visualizations. As shown in Figure 7-11, we conduct experiments using LeNet-5 and ResNet-18 for CIFAR10 and CIFAR100, and visualize FedAvg, FedSAM, and FedSWA. Among these, our FedSWA achieves a flatter loss landscape.



Figure 6: The training loss surfaces of FedAvg, FedSAM, and our FedSWA on CIFAR-100 with ResNet-18 (Dirichlet-0.1).



Figure 7: The training loss surfaces of FedAvg, FedSAM, and our FedSWA on CIFAR-100 with LeNet-5 (Dirichlet-0.1).



Figure 8: The training loss surfaces of FedAvg, FedSAM, and our FedSWA on CIFAR-10 with LeNet-5 (Dirichlet-0.1).



Figure 9: The training loss surfaces of FedSAM on CIFAR-100 with ResNet-18 (Dirichlet-0.1, Dirichlet-0.3, Dirichlet-0.6).



Figure 10: The training loss surfaces of FedSWA on CIFAR-100 with ResNet-18 (Dirichlet-0.1, Dirichlet-0.3, Dirichlet-0.6).

B.3. Convergence Behavior Curves



Figure 11: Convergence plots for FedSWA, FedMoSWA and other baselines in same settings with Dirichlet-0.6 on CIFAR10 and CIFAR100 using different neural network architectures.

heterogeneity	0.1	0.3	0.6
FedSMOO	46.5	47.8	49.2
FedGAMMA	48.4	51.8	52.6
FedSWA (ours)	50.3	55.5	59.8
FedMoSWA (ours)	61.9	66.2	67.9

Table 8: CIFAR-100 with ResNet-18 under different data heterogeneity levels (0.1, 0.3, 0.6).



Figure 12: Convergence plots for FedSWA, FedMoSWA and baselines with Dirichlet-0.6, 0.3, 0.1 on CIFAR100 using ResNet-18.

C. Basic Assumptions

We analyze generalization based on following assumptions:

Assumption C.1. The loss function $l(\cdot, z)$ is μ -strongly convex for any sample z, $l(\theta; z) \ge l(\theta'; z) + \langle \nabla l(\theta'; z), \theta - \theta' \rangle + \frac{\mu}{2} \|\theta - \theta'\|^2$, for any z, θ, θ' .

Assumption C.2. The loss function $l(\cdot, z)$ is L Lipschitz continuous, $|l(\theta; z) - l(\theta'; z)| \le L ||\theta - \theta'||$, for any z, θ, θ' . We note that under Assumption 2, $\mathbb{E}_{\mathcal{A},\mathcal{S},z} |l(\theta_T; z) - l(\theta'_T; z)| \le L\mathbb{E} ||\theta_T - \theta'_T||$, $\mathbb{E} ||\nabla F(\theta_t)|| \le L$.

Assumption C.3. There exists a constant $\sigma > 0$ such that for any $\theta, i \in [m]$, and $z_i \sim D_i$, $\mathbb{E}\left[\left\| \nabla l(\theta; z_i) - \nabla F_i(\theta) \right\|^2 \right] \leq \sigma^2$.

Assumption C.4. The loss function $l(\cdot, z)$ is β -smooth, $\|\nabla l(\theta; z) - \nabla l(\theta'; z)\| \le \beta \|\theta - \theta'\|$, for any z, θ, θ' .

Assumption C.5. (data heterogeneity). Given $i \in [m]$, there exists a constant $\sigma_g > 0$, for any θ we have $\|\nabla F_i(\theta) - \nabla F(\theta)\| \le \sigma_{g,i}, \frac{1}{m} \sum_{i=1}^m \|\nabla F_i(\theta) - \nabla F(\theta)\| \le \frac{1}{m} \sum_{i=1}^m \sigma_{g,i} = \sigma_g$

D. Main Lemmas

Lemma D.1. Assume that f is β -smooth. Then, the following properties hold. 1. $G_{f,\eta}$ is $(1 + \eta\beta)$ -expansive, $G(v) = G_{f,\eta}(v) = v - \eta \nabla f(v)$ 2. Assume in addition that f is μ -strongly convex. Then, for $\eta \leq \frac{2}{\beta + \mu}$, $G_{f,\eta}$ is $\left(1 - \frac{\eta\beta\mu}{\beta + \mu}\right)$ expansive. Henceforth we will no longer mention which random selection rule we use as the proofs are almost identical for both rules.

Proof. 1. Let $G = G_{f,\eta}$. By triangle inequality and our smoothness assumption,

$$||G(v) - G(w)|| \le ||v - w|| + \eta |\nabla f(w) - \nabla f(v)|| \le ||v - w|| + \eta \beta ||w - v|| = (1 + \eta \beta) ||v - w||.$$
(10)

2. First, note that if f is μ strongly convex, then $\varphi(w) = f(w) - \frac{\mu}{2} ||w||^2$ is convex with $(\beta - \mu)$ -smooth. Hence,

$$\langle \nabla f(v) - \nabla f(w), v - w \rangle \ge \frac{\beta \mu}{\beta + \mu} \|v - w\|^2 + \frac{1}{\beta + \mu} \|\nabla f(v) - \nabla f(w)\|^2$$
(11)

Using this inequality gives

$$\|G_{f,\eta}(v) - G_{f,\eta}(w)\|^{2} = \|v - w\|^{2} - 2\eta \langle \nabla f(v) - \nabla f(w), v - w \rangle + \eta^{2} \|\nabla f(v) - \nabla f(w)\|^{2}$$

$$\leq \left(1 - 2\frac{\eta\beta\mu}{\beta + \mu}\right) \|v - w\|^{2} - \eta \left(\frac{2}{\beta + \mu} - \eta\right) \|\nabla f(v) - \nabla f(w)\|^{2}.$$
 (12)

With our assumption that $\eta \leq \frac{2}{\beta + \mu}$, this implies

$$\|G_{f,\eta}(v) - G_{f,\eta}(w)\| \le \left(1 - 2\frac{\eta\beta\mu}{\beta+\mu}\right)^{1/2} \|v - w\| \le \left(1 - \frac{\eta\beta\mu}{\beta+\mu}\right) \|v - w\|.$$
(13)

The lemma follows by applying the inequality $\sqrt{1-x} \le 1-x/2$ which holds for $x \in [0,1]$.

E. Theoretical Results of FedSWA

E.1. Generalization Analysis for FedSWA under strongly convex setting

Lemma E.1. Suppose Assumptions C.1-C.5 hold. Then for FedSWA with $\eta_k^t \leq 1/\beta KT$, $\alpha = 1$,

$$\mathbb{E}\left\|\theta_{i,k}^{t}-\theta_{t}\right\| \leq (1-b)^{K-1}\tilde{\eta}_{t}\left(\mathbb{E}\left\|\nabla F\left(\theta_{t}\right)\right\|+\sigma_{g,i}+\sigma\right), \forall k=1,\ldots,K,$$

where $\tilde{\eta}_t = \sum_{k=0}^{K-1} \eta_k^t$.

Proof. Considering local update of FedSWA

$$\mathbb{E} \|\theta_{i,k+1} - \theta_t\| = \mathbb{E} \|\theta_{i,k}^t - \eta_k^t g_i\left(\theta_{i,k}^t\right) - \theta_t\| \\
\leq \mathbb{E} \|\theta_{i,k}^t - \theta_t - \eta_k^t\left(g_i\left(\theta_{i,k}^t\right) - g_i\left(\theta_t\right)\right)\| + \eta_k^t \mathbb{E} \|g_i\left(\theta_t\right)\| \\
\stackrel{(a)}{\leq} \left(1 - \frac{\eta_k^t \beta \mu}{\beta + \mu}\right) \mathbb{E} \|\theta_{i,k}^t - \theta_t\| + \eta_k^t \mathbb{E} \|g_i\left(\theta_t\right)\| \\
\leq \left(1 - \frac{\eta_k^t \beta \mu}{\beta + \mu}\right) \mathbb{E} \|\theta_{i,k}^t - \theta_t\| + \eta_k^t \left(\mathbb{E} \|g_i\left(\theta_t\right) - \nabla F_i\left(\theta_t\right)\| + \mathbb{E} \|\nabla F_i\left(\theta_t\right)\|\right) \\
\stackrel{(b)}{\leq} \left(1 - \frac{\eta_k^t \beta \mu}{\beta + \mu}\right) \mathbb{E} \|\theta_{i,k}^t - \theta_t\| + \eta_k^t \left(\mathbb{E} \|\nabla F_i\left(\theta_t\right)\| + \sigma\right),$$
(14)

where (a) follows Lemma D.1; (b) follows Assumption C.3. Unrolling the above and noting $\theta_{i,0} = \theta_t$, with $b = \left(1 - \frac{\eta_K^t \beta \mu}{\beta + \mu}\right)$ for some b > 0 yields

$$\mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{t} \right\| \leq \sum_{l=0}^{k-1} \eta_{l}^{t} \left(\mathbb{E} \left\| \nabla F_{i} \left(\theta_{t} \right) \right\| + \sigma \right) (1-b)^{k-1-l}$$

$$\leq \sum_{l=0}^{K-1} \eta_{l}^{t} \left(\mathbb{E} \left\| \nabla F_{i} \left(\theta_{t} \right) \right\| + \sigma \right) (1-b)^{K-1}$$

$$\leq (1-b)^{K-1} \tilde{\eta}_{t} \left(\mathbb{E} \left\| \nabla F \left(\theta_{t} \right) \right\| + \sigma_{g,i} + \sigma \right),$$
(15)

where the last inequality follows Assumption C.5, $\tilde{\eta}_t = \sum_{k=0}^{K-1} \eta_k^t$.

Lemma E.2. Given Assumptions C.1-C.4 and considering of FedSWA, for $\eta_k^t \leq 1/\beta KT$ we have

$$\mathbb{E}\left\|g_{i}\left(\theta_{i,k}^{t}\right)\right\| \leq \left(1 + (1-b)^{K-1}\beta\tilde{\eta}_{t}\right)\left(\mathbb{E}\left\|\nabla F\left(\theta_{t}\right)\right\| + \sigma_{g,i} + \sigma\right),\$$

where $g_i(\cdot)$ is the sampled gradient of client $i, \tilde{\eta}_t = \sum_{k=0}^{K-1} \eta_k^t$.

Proof. Using Assumptions C.3, C.4, C.5, we obtain

$$\mathbb{E} \left\| g_i\left(\theta_{i,k}^t\right) \right\| \leq \mathbb{E} \left\| g_i\left(\theta_{i,k}^t\right) - \nabla F_i\left(\theta_{i,k}^t\right) \right\| + \mathbb{E} \left\| \nabla F_i\left(\theta_{i,k}^t\right) \right\| \\
\leq^a \mathbb{E} \left\| \nabla F_i\left(\theta_{i,k}^t\right) \right\| + \sigma \\
\leq \mathbb{E} \left\| \nabla F_i\left(\theta_t\right) \right\| + \mathbb{E} \left\| \nabla F_i\left(\theta_{i,k}^t\right) - \nabla F_i\left(\theta_t\right) \right\| + \sigma \\
\leq^b \mathbb{E} \left\| \nabla F\left(\theta_t\right) \right\| + \mathbb{E} \left\| \nabla F_i\left(\theta_t\right) - \nabla F\left(\theta_t\right) \right\| + \beta \mathbb{E} \left\| \theta_{i,k}^t - \theta_t \right\| + \sigma \\
\leq^c \left(1 + (1-b)^{K-1} \beta \tilde{\eta}_t \right) \left(\mathbb{E} \left\| \nabla F\left(\theta_t\right) \right\| + \sigma_{q,i} + \sigma \right),$$
(16)

a is following Assumption C.3, b is is following Assumption C.4, c is is following Assumption C.5 and D.1. \Box

Theorem E.3. (FedSWA). Suppose Assumptions 1-5 hold and consider FedSWA. Let $\{\theta_t\}_{t=0}^T$ and $\{\theta'_t\}_{t=0}^T$ be two datasets S and $S^{(i)}$, respectively. Suppose $\theta_0 = \theta'_0$. Then,

$$\varepsilon_{gen} \leq \frac{2L\tilde{b}}{mn} \frac{1}{\beta} e^{1 - \frac{\mu}{(\beta + \mu)T}} \left(\mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \sigma_{g} + \sigma \right),$$

where $\tilde{\eta}_k^t = \sum_{k=0}^{K-1} \eta_k^t$, $\tilde{b} \leq 1 + (1-b)^{K-1} \beta \tilde{\eta}_t$

Proof. For client *i*, there are two cases to consider. In the first case, SGD selects the index of an sample at local step *k* on which is identical in S and $S^{(i)}$. In this sense, we have

$$\left\|\theta_{i,k+1}^{t} - \theta_{i,k+1}^{t,\prime}\right\| \le \left(1 - \frac{\eta_{k}^{t}\beta\mu}{\beta + \mu}\right) \left\|\theta_{i,k}^{t} - \theta_{i,k}^{t,\prime}\right\|.$$
(17)

And this case happens with probability 1 - 1/n (since only one sample is perturbed for client i).

In the second case, SGD encounters the perturbed sample at local time step k, which happens with probability 1/n. We denote the gradient of this perturbed sample as $g'_i(\cdot)$. Then,

$$\begin{aligned} \left\| \theta_{i,k+1}^{t} - \theta_{i,k+1}^{t,\prime} \right\| &= \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} - \eta_{k}^{t} \left(g_{i} \left(\theta_{i,k}^{t} \right) - g_{i}^{\prime} \left(\theta_{i,k}^{t,\prime} \right) \right) \right\| \\ &\leq \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} - \eta_{k}^{t} \left(g_{i} \left(\theta_{i,k}^{t} \right) - g_{i} \left(\theta_{i,k}^{t,\prime} \right) \right) \right\| + \eta_{k}^{t} \left\| g_{i} \left(\theta_{i,k}^{t} \right) - g_{i}^{\prime} \left(\theta_{i,k}^{t,\prime} \right) \right\| \\ &\leq \left(1 - \frac{\eta_{k}^{t} \beta \mu}{\beta + \mu} \right) \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} \right\| + \eta_{k}^{t} \left\| g_{i} \left(\theta_{i,k}^{t} \right) - g_{i}^{\prime} \left(\theta_{i,k}^{t,\prime} \right) \right\|. \end{aligned}$$
(18)

Combining these two cases we have for client *i*,

$$\mathbb{E}\left\|\theta_{i,k+1}^{t} - \theta_{i,k+1}^{t,\prime}\right\| \leq \left(1 - \frac{\eta_{k}^{t}\beta\mu}{\beta + \mu}\right) \mathbb{E}\left\|\theta_{i,k}^{t} - \theta_{i,k}^{t,\prime}\right\| + \frac{\eta_{k}^{t}}{n} \mathbb{E}\left\|g_{i}\left(\theta_{i,k}^{t}\right) - g_{i}^{\prime}\left(\theta_{i,k}^{t,\prime}\right)\right\| \\ \leq \left(1 - \frac{\eta_{k}^{t}\beta\mu}{\beta + \mu}\right) \mathbb{E}\left\|\theta_{i,k}^{t} - \theta_{i,k}^{t,\prime}\right\| + \frac{2\eta_{k}^{t}}{n} \mathbb{E}\left\|g_{i}\left(\theta_{i,k}^{t}\right)\right\|, \tag{19}$$

where the last inequality follows that $g_i(\cdot)$ and $g'_i(\cdot)$ are sampled from the same distribution.

Iterating the above over t and noting $\theta_0 = \theta'_0$, we conclude the proof. We let \tilde{b} be an upper bound of $1 + (1-b)^{K-1}\beta \tilde{\eta}_t$ since $\tilde{\eta}_t$ is bounded above. Then unrolling 19, and follow E.2,

$$\mathbb{E} \left\| \theta_{i,K}^{t} - \theta_{i,K}^{t,\prime} \right\| \\
\leq \prod_{k=0}^{K-1} \left(1 - \frac{\eta_{k}^{t}\beta\mu}{\beta+\mu} \right) \mathbb{E} \left\| \theta_{t} - \theta_{t}^{\prime} \right\| + \left(\frac{2}{n} \sum_{k=0}^{K-1} \eta_{k}^{t} \tilde{b} \prod_{l=k+1}^{K-1} \left(1 - \frac{\eta_{k}^{t}\beta\mu}{\beta+\mu} \right) \left(\mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \sigma_{g,i} + \sigma \right) \right) \qquad (20)$$

$$\leq^{a} e^{\frac{-\tilde{\eta}_{t}\beta\mu}{\beta+\mu}} \mathbb{E} \left\| \theta_{t} - \theta_{t}^{\prime} \right\| + \frac{2}{n} \tilde{b} \tilde{\eta}_{t} e^{\frac{-\tilde{\eta}_{t}\beta\mu}{\beta+\mu}} \left(\mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \sigma_{g,i} + \sigma \right).$$

a is following $1 - x \le e^{-x}$. With $\theta_{t+1} = \frac{1}{m} \sum_{i=1}^{m} \theta_{i,K}^t$ and $\theta'_{t+1} = \frac{1}{m} \sum_{i=1}^{m} \theta_{i,K}^{t,\prime}$, we have

$$\mathbb{E} \left\| \theta_{t+1} - \theta_{t+1}' \right\| \leq \sum_{i=1}^{m} \frac{1}{m} \mathbb{E} \left\| \theta_{i,K}^{t} - \theta_{i,K}^{t,\prime} \right\| \\
\leq e^{\frac{-\tilde{\eta}_{t}\beta\mu}{\beta+\mu}} \mathbb{E} \left\| \theta_{t} - \theta_{t}' \right\| + \frac{2}{mn} \tilde{b} \tilde{\eta}_{t} e^{\frac{-\tilde{\eta}_{t}\beta\mu}{\beta+\mu}} \left(\mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \sigma_{g} + \sigma \right)$$
(21)

Further, unrolling the above over t and noting $\theta_0 = \theta'_0$, we obtain

$$\mathbb{E} \left\| \theta_{T} - \theta_{T}' \right\| \leq \frac{2\tilde{b}}{mn} \sum_{t=0}^{T-1} \exp\left(\beta \sum_{l=t+1}^{T-1} \tilde{\eta}_{t}\right) \tilde{\eta}_{t} e^{-\frac{\tilde{\eta}_{t}\beta\mu}{\beta+\mu}} \left(\mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \sigma_{g} + \sigma\right)$$

$$\leq \frac{2\tilde{b}}{mn} \frac{1}{\beta} e^{1 - \frac{\mu}{(\beta+\mu)T}} \left(\mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \sigma_{g} + \sigma\right),$$
(22)

With Assumption C.2,

$$\varepsilon_{\text{gen}} \leq \frac{2Lb}{mn} \frac{1}{\beta} e^{1 - \frac{\mu}{(\beta + \mu)T}} \left(\mathbb{E} \left\| \nabla F\left(\theta_t\right) \right\| + \sigma_g + \sigma \right)$$
$$\tilde{b} = 1 + \left(\frac{\mu}{(\beta + \mu)K} \right)^{K-1} \frac{1}{T} \gg 1$$

When the diminishing stepsizes $\eta_k^t \leq 1/\beta KT$ are chosen in the statement of the theorem, we conclude the proof.

E.2. Generalization Analysis for FedSWA under non-convex setting

Lemma E.4. Suppose Assumptions C.2-C.5 hold. Then for FedSWA with $\eta_k^t \leq /\beta KT$,

 $\mathbb{E}\left\|\theta_{i,k}^{t}-\theta_{t}\right\| \leq (1+c)^{K-1}\tilde{\eta}_{t}\left(\mathbb{E}\left\|\nabla F\left(\theta_{t}\right)\right\|+\sigma_{g,i}+\sigma\right), \quad \forall k=1,\ldots,K,$

where $\tilde{\eta}_t = \sum_{k=0}^{K-1} \eta_k^t$.

Proof. Considering local update of FedSWA

$$\mathbb{E} \|\theta_{i,k+1} - \theta_t\| = \mathbb{E} \|\theta_{i,k}^t - \eta_k^t g_i(\theta_{i,k}^t) - \theta_t\|
\leq \mathbb{E} \|\theta_{i,k}^t - \theta_t - \eta_k^t(g_i(\theta_{i,k}^t) - g_i(\theta_t))\| + \eta_k^t \mathbb{E} \|g_i(\theta_t)\|
\leq^a (1 + \beta \eta_k^t) \mathbb{E} \|\theta_{i,k}^t - \theta_t\| + \eta_k^t \mathbb{E} \|g_i(\theta_t)\|
\leq (1 + \beta \eta_k^t) \mathbb{E} \|\theta_{i,k}^t - \theta_t\| + \eta_k^t(\mathbb{E} \|g_i(\theta_t) - \nabla F_i(\theta_t)\| + \mathbb{E} \|\nabla F_i(\theta_t)\|)
\leq^b (1 + \beta \eta_k^t) \mathbb{E} \|\theta_{i,k}^t - \theta_t\| + \eta_k^t(\mathbb{E} \|\nabla F_i(\theta_t)\| + \sigma),$$
(23)

where a is following D.1, and we use Assumptions C.3 in b . With $(1 + \beta \eta_k^t) \le c$, Unrolling the above and noting $\theta_{i,0} = \theta_t$ yields

$$\mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{t} \right\| \leq \sum_{l=0}^{k-1} \eta_{l}^{t} \left(\mathbb{E} \left\| \nabla F_{i} \left(\theta_{t} \right) \right\| + \sigma \right) (1+c)^{k-1-l}$$

$$\leq \sum_{l=0}^{K-1} \eta_{l}^{t} \left(\mathbb{E} \left\| \nabla F_{i} \left(\theta_{t} \right) \right\| + \sigma \right) (1+c)^{K-1}$$

$$\leq (1+c)^{K-1} \tilde{\eta}_{t} \left(\mathbb{E} \left\| \nabla F \left(\theta_{t} \right) \right\| + \sigma_{g,i} + \sigma \right),$$
(24)

where the last inequality follows Assumption C.5, $\tilde{\eta}_t = \sum_{k=0}^{K-1} \eta_k^t$.

Lemma E.5. Given Assumptions C.2-C.5, for $\eta_k^t \leq /\beta KT$, we have

$$\mathbb{E}\left\|g_{i}\left(\theta_{i,k}^{t}\right)\right\| \leq \left(1 + (1+c)^{K-1}\beta\tilde{\eta}_{t}\right)\left(\mathbb{E}\left\|\nabla F\left(\theta_{t}\right)\right\| + \sigma_{g,i} + \sigma\right)$$

where $g_i(\cdot)$ is the sampled gradient of client $i, \tilde{\eta}_t = \sum_{k=0}^{K-1} \eta_k^t$.

Proof. we obtain

$$\mathbb{E} \left\| g_{i}\left(\theta_{i,k}^{t}\right) \right\| \leq \mathbb{E} \left\| g_{i}\left(\theta_{i,k}^{t}\right) - \nabla F_{i}\left(\theta_{i,k}^{t}\right) \right\| + \mathbb{E} \left\| \nabla F_{i}\left(\theta_{i,k}^{t}\right) \right\| \\
\leq^{a} \mathbb{E} \left\| \nabla F_{i}\left(\theta_{i,k}^{t}\right) \right\| + \sigma \\
\leq \mathbb{E} \left\| \nabla F_{i}\left(\theta_{t}\right) \right\| + \mathbb{E} \left\| \nabla F_{i}\left(\theta_{i,k}^{t}\right) - \nabla F_{i}\left(\theta_{t}\right) \right\| + \sigma \\
\leq^{b} \mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \mathbb{E} \left\| \nabla F_{i}\left(\theta_{t}\right) - \nabla F\left(\theta_{t}\right) \right\| + \beta \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{t} \right\| + \sigma \\
\leq^{c} \left(1 + (1+c)^{K-1}\beta\tilde{\alpha}_{i} \right) \left(\mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \sigma_{g,i} + \sigma \right).$$
(25)

a is following Assumption C.3, *b* is is following Assumption C.4, *c* is is following Assumption C.5 and D.1 and E.4. \Box

Theorem E.6. (FedSWA). Suppose Assumptions C.2-C.5 hold and consider FedSWA. Let $k = K, \forall i \in [m] \text{ and } \eta_k^t \leq \frac{1}{\beta KT}$. Then,

$$\varepsilon_{gen} \leq \frac{2L\tilde{c}}{mn} \frac{1}{\beta} e^{1+\frac{1}{T}} \left(\mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \sigma_{g} + \sigma \right),$$

 $\tilde{c} \le 1 + (1+c)^{K-1} \beta \tilde{\eta}_t.$

Proof. For client *i*, there are two cases to consider. In the first case, SGD selects non-perturbed samples in S and $S^{(i)}$, which happens with probability 1 - 1/n. Then, with D.1, we have

$$\left\|\theta_{i,k+1}^{t} - \theta_{i,k+1}^{t,\prime}\right\| \leq \left(1 + \beta \eta_{k}^{t}\right) \left\|\theta_{i,k}^{t} - \theta_{i,k}^{t,\prime}\right\|.$$
(26)

In the second case, SGD encounters the perturbed sample at time step k, which happens with probability 1/n. Then, we have

$$\begin{aligned} \left\| \theta_{i,k+1}^{t} - \theta_{i,k+1}^{t,\prime} \right\| &= \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} - \eta_{k}^{t} \left(g_{i} \left(\theta_{i,k}^{t} \right) - g_{i}^{\prime} \left(\theta_{i,k}^{t,\prime} \right) \right) \right\| \\ &\leq \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} - \eta_{k}^{t} \left(g_{i} \left(\theta_{i,k}^{t} \right) - g_{i} \left(\theta_{i,k}^{t,\prime} \right) \right) \right\| + \eta_{k}^{t} \left\| g_{i} \left(\theta_{i,k}^{t,\prime} \right) - g_{i}^{\prime} \left(\theta_{i,k}^{t,\prime} \right) \right\| . \end{aligned}$$

$$(27)$$

Combining these two cases for client i we have

$$\mathbb{E} \left\| \theta_{i,k+1}^{t} - \theta_{i,k+1}^{t,\prime} \right\| \leq (1 + \beta \eta_{k}^{t}) \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} \right\| + \frac{\eta_{k}^{t}}{n} \mathbb{E} \left\| g_{i} \left(\theta_{i,k}^{t,\prime} \right) - g_{i}^{\prime} \left(\theta_{i,k}^{t,\prime} \right) \right\| \\ \leq (1 + \beta \eta_{k}^{t}) \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} \right\| + \frac{\eta_{k}^{t}}{n} \mathbb{E} \left\| g_{i} \left(\theta_{i,k}^{t} \right) \right\| \\ \leq (1 + \beta \eta_{k}^{t}) \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} \right\| + \frac{2\eta_{k}^{t}}{n} \left(1 + (1 + c)^{K-1} \beta \tilde{\eta}_{t} \right) \left(\sigma + \mathbb{E} \left\| \nabla F \left(\theta_{t} \right) \right\| + \sigma_{g,i} \right) \\ \leq (1 + \beta \eta_{k}^{t}) \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} \right\| + \frac{2\eta_{k}^{t} \tilde{c}}{n} \left(\mathbb{E} \left\| \nabla F \left(\theta_{t} \right) \right\| + \sigma_{g,i} + \sigma \right)$$
(28)

where the last inequality follows Assumption D.1. We let $1 + (1+c)^{K-1}\beta \tilde{\eta}_t \leq \tilde{c}$ and $\tilde{\eta}_t = \sum_{k=0}^{K-1} \eta_k^t$. Then unrolling (28) with E.5 gives

$$\mathbb{E}\left\|\theta_{i,K}^{t}-\theta_{i,K}^{t,\prime}\right\| \leq \prod_{k=0}^{K-1} \left(1+\beta\eta_{k}^{t}\right) \mathbb{E}\left\|\theta_{t}-\theta_{t}^{\prime}\right\| + \left(\frac{2}{n}\sum_{k=0}^{K-1}\eta_{k}^{t}\tilde{c}\prod_{l=k+1}^{K-1}\left(1+\beta\eta_{l}^{t}\right)\left(\mathbb{E}\left\|\nabla F\left(\theta_{t}\right)\right\|+\sigma_{g,i}+\sigma\right)\right) \leq e^{\beta\tilde{\eta}_{t}}\mathbb{E}\left\|\theta_{t}-\theta_{t}^{\prime}\right\| + \frac{2}{n}\tilde{c}\tilde{\eta}_{t}e^{\beta\tilde{\eta}_{t}}\left(\mathbb{E}\left\|\nabla F\left(\theta_{t}\right)\right\|+\sigma_{g,i}+\sigma\right).$$
(29)

With $\theta_{t+1} = \frac{1}{m} \sum_{i=1}^{m} \theta_{i,K}^t$ and $\theta_{t+1}' = \frac{1}{m} \sum_{i=1}^{m} \theta_{i,K}^{t,\prime}$, we have

$$\mathbb{E} \left\| \theta_{t+1} - \theta_{t+1}' \right\| \leq \sum_{i=1}^{m} \frac{1}{m} \mathbb{E} \left\| \theta_{i,K}^{t} - \theta_{i,K}^{t,\prime} \right\|$$

$$\leq e^{\beta \tilde{\eta}_{t}} \mathbb{E} \left\| \theta_{t} - \theta_{t}' \right\| + \frac{2}{mn} \tilde{c} \tilde{\eta}_{t} e^{\beta \tilde{\eta}_{t}} \left(\mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \sigma_{g,i} + \sigma \right)$$
(30)

Further, unrolling the above over t and noting $\theta_0=\theta_0',$ we obtain

$$\mathbb{E} \left\| \theta_T - \theta_T' \right\| \le \frac{2\tilde{c}}{mn} \sum_{t=0}^{T-1} \exp\left(\beta \sum_{l=t+1}^{T-1} \tilde{\eta}_t\right) \tilde{\eta}_t e^{\beta \tilde{\eta}_t} \left(\mathbb{E} \left\| \nabla F\left(\theta_t\right) \right\| + \sigma_g + \sigma\right),\tag{31}$$

We simplify this inequality

$$\mathbb{E} \left\| \theta_{T} - \theta_{T}^{\prime} \right\| \\
\leq \frac{2\tilde{c}}{mn} \frac{1}{\beta} e^{\frac{1}{T} + 1} \left(\mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \sigma_{g} + \sigma \right) \\
\leq \frac{2\tilde{c}}{mn} \frac{1}{\beta} e^{\frac{1}{T} + 1} \left(L + \sigma_{g} + \sigma \right),$$
(32)

With the Assumption C.2,

$$\varepsilon_{\text{gen}} \leq \frac{2L\tilde{c}}{mn} \frac{1}{\beta} e^{1 + \frac{1}{T}} \left(\mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \sigma_{g} + \sigma \right),$$

With

$$\tilde{c} = 1 + \left(2 + \frac{1}{TK}\right)^{K-1} \frac{1}{T} \gg 1$$

When the diminishing stepsizes are chosen in the statement of the theorem, we conclude the proof.

F. Theoretical Results of FedMoSWA

F.1. Generalization Analysis for FedMoSWA under strongly convex setting

where $m{g}_{i,k}^{(t)}$ denotes the stochastic gradient at the point $m{ heta}_{i,k}^{(t)}$.

Lemma F.1. Suppose Assumptions C.1-C.5 hold. Then for FedMoSWA with $\eta_k^t \leq 1/\beta KT$, $\alpha = 1$,

$$\mathbb{E}\left\|\theta_{i,K}^{t} - \theta_{t}\right\| \leq \tilde{\eta}_{t}(1-b)^{K-1}\left(L+\sigma\right)$$

where $\tilde{\eta}_t = \sum_{k=0}^{K-1} \eta_k^t$, $b = \frac{\eta^t \beta \mu}{\beta + \mu}$.

Proof. Considering local update of FedMoSWA, $c_i^t = g_i(\theta_t), m^t = (1 - \gamma)m^{t-1} + \gamma \frac{1}{m} \sum g_i(\theta_t),$

$$\begin{aligned}
\mathbb{E} \left\| \theta_{i,k+1}^{t} - \theta_{t} \right\| \\
&= \mathbb{E} \left\| \theta_{i,k+1}^{t} - \eta_{k}^{t} \left(g_{i} \left(\theta_{i,k+1}^{t} \right) - c_{i}^{t} + \boldsymbol{m}^{t} \right) - \theta_{t} \right\| \\
&= \mathbb{E} \left\| \theta_{i,k+1}^{t} - \eta_{t} - \eta_{k}^{t} g_{i} \left(\theta_{i,k+1}^{t} \right) + \eta_{k}^{t} \left(g_{i} \left(\theta_{t} \right) + \boldsymbol{m}^{t} \right) \right\| \\
&= \mathbb{E} \left\| \theta_{i,k+1}^{t} - \theta_{t} - \eta_{k}^{t} g_{i} \left(\theta_{i,k+1}^{t} \right) + \eta_{k}^{t} g_{i} \left(\theta_{t} \right) \right\| + \eta_{k}^{t} \mathbb{E} \left\| \boldsymbol{m}^{t} \right\| \\
&\leq^{a} \mathbb{E} \left\| \theta_{i,k+1}^{t} - \theta_{t} - \eta_{k}^{t} \left(g_{i} \left(\theta_{i,k+1}^{t} \right) - g_{i} \left(\theta_{t} \right) \right) \right\| + \eta_{k}^{t} \sum_{l=0}^{t} \left(1 - \gamma \right)^{l} \left(\gamma \mathbb{E} \left\| \nabla F \left(\theta_{l} \right) \right\| + \sigma \right) \\
&\leq^{b} \left(1 - \frac{\eta_{k}^{t} \beta \mu}{\beta + \mu} \right) \mathbb{E} \left\| \theta_{i,k+1}^{t} - \theta_{t} \right\| + \eta_{k}^{t} \mathbb{E} \left\| \boldsymbol{m}^{t} \right\| \\
&\leq \left(1 - \frac{\eta_{k}^{t} \beta \mu}{\beta + \mu} \right) \mathbb{E} \left\| \theta_{i,k+1}^{t} - \theta_{t} \right\| + \eta_{k}^{t} \sum_{l=0}^{t} \left(1 - \gamma \right)^{l} \left(\gamma \mathbb{E} \left\| \nabla F \left(\theta_{l} \right) \right\| + \sigma \right) \\
&\leq^{c} \left(1 - \frac{\eta_{k}^{t} \beta \mu}{\beta + \mu} \right) \mathbb{E} \left\| \theta_{i,k+1}^{t} - \theta_{t} \right\| + \eta_{k}^{t} \sum_{l=0}^{t} \left(1 - \gamma \right)^{l} \left(\gamma L + \sigma \right)
\end{aligned}$$
(33)

where, *a* is following Assumption C.3, *b* is following Lemma D.1, c is following Assumption C.2. Below we find an upper bound for $\mathbb{E} \| \boldsymbol{m}^t \|$,

$$\mathbb{E} \|\boldsymbol{m}^{t}\| = \mathbb{E} \left\| (1-\gamma) \, \boldsymbol{m}^{t-1} + \gamma \frac{1}{m} \sum_{i=1}^{m} g_{i}\left(\theta_{t}\right) \right\|$$

$$\leq (1-\gamma) \mathbb{E} \|\boldsymbol{m}^{t-1}\| + \gamma \mathbb{E} \left\| \frac{1}{m} \sum_{i=1}^{m} g_{i}\left(\theta_{t}\right) \right\|$$

$$\leq (1-\gamma) \mathbb{E} \|\boldsymbol{m}^{t-1}\| + \gamma \mathbb{E} \|\nabla F\left(\theta_{t}\right)\| + \gamma \sigma$$

$$\leq^{a} \sum_{l=0}^{t} (1-\gamma)^{l} \left(\gamma \mathbb{E} \|\nabla F\left(\theta_{l}\right)\| + \gamma \sigma\right)$$

$$\leq^{b} L + \sigma$$
(34)

where a follows Assumption C.3; b follows Lemma D.1. Unrolling the (F.1) and noting $\theta_{i,0} = \theta_t$, with $b = \frac{\eta^t \beta \mu}{\beta + \mu}$ for some b > 0, $\tilde{\eta}_t = \sum_{k=0}^{K-1} \eta_k^t$ yields

$$\mathbb{E} \left\| \theta_{i,K}^{t} - \theta_{t} \right\| \\
\leq (1-b)^{K-1} \tilde{\eta}_{t} \mathbb{E} \left\| \boldsymbol{m}^{t} \right\| \\
\leq \sum_{l=0}^{K-1} \eta_{l}^{t} \sum_{p=0}^{t} (1-\gamma)^{p} \left(\gamma \mathbb{E} \left\| \nabla F \left(\theta_{p} \right) \right\| + \gamma \sigma \right) (1-b)^{K-1-l} \\
\leq \sum_{l=0}^{K-1} \eta_{l}^{t} \left(L + \sigma \right) (1-b)^{K-1-l} \\
\leq \tilde{\eta}_{t} (1-b)^{K-1} \left(L + \sigma \right)$$
(35)

Lemma F.2. Given Assumptions C.1-C.5 , for $\eta_k^t \leq 1/\beta KT$ we have

$$\mathbb{E}\left\|g_{i}\left(\theta_{i,k}^{t}\right)\right\| \leq \left(1 + \beta \tilde{\eta}_{t}(1+b)^{K-1}\right)\left(L+\sigma\right) + \sigma_{g,i}$$

where $g_i(\cdot)$ is the sampled gradient of client $i, \tilde{\eta}_t = \sum_{k=0}^{K-1} \eta_k^t$.

Proof. we obtain

$$\mathbb{E} \left\| g_{i}\left(\theta_{i,k}^{t}\right) \right\| \leq \mathbb{E} \left\| g_{i}\left(\theta_{i,k}^{t}\right) - \nabla F_{i}\left(\theta_{i,k}^{t}\right) \right\| + \mathbb{E} \left\| \nabla F_{i}\left(\theta_{i,k}^{t}\right) \right\| \\ \leq^{a} \mathbb{E} \left\| \nabla F_{i}\left(\theta_{i,k}^{t}\right) \right\| + \sigma \\ \leq \mathbb{E} \left\| \nabla F_{i}\left(\theta_{t}\right) \right\| + \mathbb{E} \left\| \nabla F_{i}\left(\theta_{i,k}^{t}\right) - \nabla F_{i}\left(\theta_{t}\right) \right\| + \sigma \\ \leq^{b} \mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \mathbb{E} \left\| \nabla F_{i}\left(\theta_{t}\right) - \nabla F\left(\theta_{t}\right) \right\| + \beta \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{t} \right\| + \sigma \\ \leq^{c} \left(\mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \sigma_{g,i} + \sigma \right) + \sum_{l=0}^{K-1} \eta_{l}^{t} \sum_{p=0}^{l} \left(1 - \gamma \right)^{p} \left(\gamma \mathbb{E} \left\| \nabla F\left(\theta_{p}\right) \right\| + \sigma \right) \left(1 - b \right)^{K-1-l} \\ \leq^{d} \left(1 + \beta \tilde{\eta}_{t} (1 + b)^{K-1} \right) \left(L + \sigma \right) + \sigma_{g,i}$$
(36)

a is following Assumption C.3, *b* is is following Assumption C.4, *c* is is following Assumption C.5 and Lemma D.1 and Lemma F.1, *d* is following $\|\nabla F(\theta_t)\| \leq L$. Similarly, using same techniques we have

$$\mathbb{E} \|g_i(\theta_t)\| \leq \mathbb{E} \|\nabla F_i(\theta_t)\| + \sigma$$

$$\leq \mathbb{E} \|\nabla F_i(\theta_t) - \nabla F(\theta_t)\| + \mathbb{E} \|\nabla F(\theta_t)\| + \sigma$$

$$\leq \sigma_{g,i} + \mathbb{E} \|\nabla F(\theta_t)\| + \sigma.$$
(37)

Theorem F.3. Suppose Assumptions C.1-C.5 hold and consider FedMoSWA (Algorithm 1). Let $\{\theta_t\}_{t=0}^T$ and $\{\theta'_t\}_{t=0}^T$ be two trajectories of the server induced by neighboring datasets S and $S^{(i)}$, respectively. Suppose $\theta_0 = \theta'_0$. Then,

$$\mathbb{E} \|\theta_T - \theta_T'\| \le \frac{2}{mn} \frac{1}{\beta} e^{1 - \frac{\mu}{(\beta + \mu)}} (\tilde{b}L + \sigma_g + \tilde{b}\sigma),$$
$$\varepsilon_{gen} \le \frac{2L}{mn} \frac{1}{\beta} e^{1 - \frac{\mu}{(\beta + \mu)T}} \left(\tilde{b}L + \sigma_g + \tilde{b}\sigma\right)$$

where $\tilde{\eta}_t = \sum_{k=0}^{K-1} \eta_k^t$. \tilde{b} be an upper bound of $1 + (1-b)^{K-1} \beta \tilde{\eta}_t$.

Proof. For client *i*, there are two cases to consider. In the first case, SGD selects the index of an sample at local step *k* on which is identical in S and $S^{(i)}$. In this sense, we have,

$$\begin{aligned} \left\| \theta_{j,k+1}^{t} - \theta_{j,k+1}^{t,\prime} \right\| \\ &\leq \left\| \theta_{j,k+1}^{t} - \theta_{j,k+1}^{t,\prime} - \eta_{k}^{t} \left(g_{j} \left(\theta_{j,k+1}^{t} \right) - g_{j} \left(\theta_{j,k}^{t,\prime} \right) - \boldsymbol{c}_{i}^{t} + \boldsymbol{m}^{t} - \boldsymbol{c}_{i}^{t} + \boldsymbol{m}^{t} \right) \right\| \\ &\leq^{a} \left(1 - \frac{\eta_{k}^{t} \beta \mu}{\beta + \mu} \right) \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} \right\| \end{aligned} \tag{38}$$

where a is following Lemma D.1. And this case happens with probability 1 - 1/n (since only one sample is perturbed for client i).

In the second case, SGD encounters the perturbed sample at local time step k, which happens with probability 1/n. We denote the gradient of this perturbed sample as $g'_i(\cdot)$. Then,

$$\begin{split} \left\| \theta_{i,k+1}^{t} - \theta_{i,k+1}^{t,\prime} \right\| &\leq \left\| \theta_{j,k+1}^{t} - \theta_{j,k+1}^{t,\prime} - \eta_{k}^{t} \left(g_{j} \left(\theta_{j,k+1}^{t} \right) - g_{j}^{\prime} \left(\theta_{j,k}^{t,\prime} \right) - c_{i}^{t} + m^{t} - c_{i}^{t} + m^{t} \right) \right\| \\ &= \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} - \eta_{k}^{t} \left(g_{i} \left(\theta_{i,k}^{t} \right) - g_{i}^{\prime} \left(\theta_{i,k}^{t,\prime} \right) \right) \right\| \\ &\leq \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} - \eta_{k}^{t} \left(g_{i} \left(\theta_{i,k}^{t} \right) - g_{i} \left(\theta_{i,k}^{t,\prime} \right) \right) \right\| + \eta_{k}^{t} \left\| g_{i} \left(\theta_{i,k}^{t,\prime} \right) - g_{i}^{\prime} \left(\theta_{i,k}^{t,\prime} \right) \right\| \\ &\leq \left(1 - \frac{\eta_{k}^{t} \beta \mu}{\beta + \mu} \right) \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} \right\| + \eta_{k}^{t} \left\| g_{i} \left(\theta_{i,k}^{t,\prime} \right) - g_{i}^{\prime} \left(\theta_{i,k}^{t,\prime} \right) \right\|. \end{split}$$

Combining these two cases we have for client *i*,

$$\mathbb{E} \left\| \theta_{i,k+1}^{t} - \theta_{i,k+1}^{t,\prime} \right\| \leq \left(1 - \frac{\eta_{k}^{t}\beta\mu}{\beta+\mu} \right) \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} \right\| + \frac{\eta_{k}^{t}}{n} \mathbb{E} \left\| g_{i} \left(\theta_{i,k}^{t,\prime} \right) - g_{i}^{\prime} \left(\theta_{i,k}^{t,\prime} \right) \right\| \\
\leq^{a} \left(1 - \frac{\eta_{k}^{t}\beta\mu}{\beta+\mu} \right) \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} \right\| + \frac{2\eta_{k}^{t}}{n} \mathbb{E} \left\| g_{i} \left(\theta_{i,k}^{t} \right) \right\| \\
\leq^{b} \left(1 - \frac{\eta_{k}^{t}\beta\mu}{\beta+\mu} \right) \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} \right\| + \frac{2\eta_{k}^{t}}{n} \left(1 + \beta\tilde{\eta}_{t}(1-b)^{K-1} \right) \left(L + \sigma \right) + \frac{2\eta_{k}^{t}}{n} \sigma_{g},$$
(40)

where the *a* inequality follows that $g_i(\cdot)$ and $g'_i(\cdot)$ are sampled from the same distribution, *b* is following Lemma F.2. We let $1 + (1-b)^{K-1}\beta \tilde{\eta}_t \leq \tilde{b}, \tilde{\eta}_t = \sum_{k=0}^{K-1} \eta_k^t$. Then unrolling it we have,

$$\mathbb{E}\left\|\theta_{i,K}^{t}-\theta_{i,K}^{t,\prime}\right\| \leq \prod_{k=0}^{K-1} \left(1-\frac{\eta_{k}^{t}\beta\mu}{\beta+\mu}\right) \mathbb{E}\left\|\theta_{t}-\theta_{t}'\right\| + \left(\frac{2}{n}\sum_{k=0}^{K-1}\eta_{k}^{t}\prod_{l=k+1}^{K-1}\left(1-\frac{\eta_{k}^{t}\beta\mu}{\beta+\mu}\right)\right) \\ \cdot\left(\left(\mathbb{E}\left\|\nabla F\left(\theta_{t}\right)\right\|+\sigma_{g,i}+\sigma\right)+\left(1-b\right)^{K-1}\tilde{\eta}_{t}\mathbb{E}\left\|\boldsymbol{m}^{t}\right\|\right)\right) \\ \leq e^{-\frac{\tilde{\eta}_{t}\beta\mu}{\beta+\mu}}\mathbb{E}\left\|\theta_{t}-\theta_{t}'\right\| + \frac{2}{n}\tilde{\eta}_{t}e^{-\frac{\tilde{\eta}_{t}\beta\mu}{\beta+\mu}}\left(\left(\mathbb{E}\left\|\nabla F\left(\theta_{t}\right)\right\|+\sigma_{g,i}+\sigma\right)+\left(1-b\right)^{K-1}\tilde{\eta}_{t}\mathbb{E}\left\|\boldsymbol{m}^{t}\right\|\right).$$
(41)

Then, with $\theta_{t+1} = \frac{1}{m} \sum_{i=1}^{m} \theta_{i,K}^t$ and $\theta'_{t+1} = \frac{1}{m} \sum_{i=1}^{m} \theta_{i,K}^{t,\prime}$, we have

$$\mathbb{E} \left\| \theta_{t+1} - \theta_{t+1}' \right\| \leq \sum_{i=1}^{m} \frac{1}{m} \mathbb{E} \left\| \theta_{i,K}^{t} - \theta_{i,K}^{t,\prime} \right\|
\leq e^{-\frac{\tilde{\eta}_{t}\beta\mu}{\beta+\mu}} \mathbb{E} \left\| \theta_{t} - \theta_{t}' \right\| + \frac{2}{mn} \tilde{\eta}_{t} e^{-\frac{\tilde{\eta}_{t}\beta\mu}{\beta+\mu}} \left((\mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \sigma_{g,i} + \sigma) + (1-b)^{K-1} \tilde{\eta}_{t} \mathbb{E} \left\| \boldsymbol{m}^{t} \right\| \right)$$
(42)

where we use Lemma F.2 in the last step. Further, unrolling the above over t and noting $\theta_0 = \theta'_0$, $b = \frac{\eta^t \beta \mu}{\beta + \mu} = \frac{\eta^t_k \beta \mu}{\beta + \mu}$ we obtain,

$$\mathbb{E} \|\theta_{T} - \theta_{T}'\| \leq \frac{2}{mn} \sum_{t=0}^{T-1} \exp\left(-\frac{\beta\mu}{\beta+\mu} \sum_{l=t+1}^{T-1} \tilde{\eta}_{t}\right) \tilde{\eta}_{t} e^{-\frac{\tilde{\eta}_{t}\beta\mu}{\beta+\mu}} \left((\mathbb{E} \|\nabla F(\theta_{t})\| + \sigma_{g} + \sigma) + (1-b)^{K-1} \tilde{\eta}_{t} \mathbb{E} \|\boldsymbol{m}^{t}\| \right) \\ \leq \frac{2}{mn} \sum_{t=0}^{T-1} \exp\left(-\frac{\beta\mu}{\beta+\mu} \sum_{l=t+1}^{T-1} \tilde{\eta}_{t}\right) \tilde{\eta}_{t} e^{-\frac{\tilde{\eta}_{t}\beta\mu}{\beta+\mu}} \left(\left(1 + \beta \tilde{\eta}_{t}(1-b)^{K-1}\right) (L+\sigma) + \sigma_{g} \right) \tag{43}$$

When the diminishing stepsizes are chosen in the statement of the theorem, we conclude the proof. \Box

F.2. Generalization Analysis for FedMoSWA under non-convex setting

Lemma F.4. Suppose Assumptions C.1-C.5 hold. Then for FedMoSWA with $\eta_k^t \leq 1/\beta KT$,

$$\mathbb{E}\left\|\theta_{i,k}^t - \theta_t\right\| \le \tilde{\eta}_t (1+c)^{k-1} (L+\sigma)$$

where $\tilde{\eta}_t = \sum_{k=0}^{K-1} \eta_k^t$.

Proof. Considering local update of FedMoSWA, $c_{i}^{t} = g_{i}(\theta_{t}), m^{t} = (1 - \gamma)m^{t-1} + \gamma \frac{1}{m} \sum g_{i}(\theta_{t}),$

$$\mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{t} \right\| \\
= \mathbb{E} \left\| \theta_{i,k}^{t} - \eta_{k}^{t} \left(g_{i} \left(\theta_{i,k}^{t} \right) - c_{i}^{t} + \boldsymbol{m}^{t} \right) - \theta_{t} \right\| \\
= \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{t} - \eta_{k}^{t} g_{i} \left(\theta_{i,k}^{t} \right) + \eta_{k}^{t} \left(g_{i} \left(\theta_{t} \right) + \boldsymbol{m}^{t} \right) \right\| \\
= \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{t} - \eta_{k}^{t} g_{i} \left(\theta_{i,k}^{t} \right) + \eta_{k}^{t} g_{i} \left(\theta_{t} \right) \right\| + \eta_{k}^{t} \mathbb{E} \left\| \boldsymbol{m}^{t} \right\| \\
\leq^{a} \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{t} - \eta_{k}^{t} \left(g_{i} \left(\theta_{i,k}^{t} \right) - g_{i} \left(\theta_{t} \right) \right) \right\| + \eta_{k}^{t} \sum_{l=0}^{t} \left(1 - \gamma \right)^{l} \left(\gamma \mathbb{E} \left\| \nabla F \left(\theta_{l} \right) \right\| + \gamma \sigma \right) \\
\leq^{b} \left(1 + \beta \eta_{k}^{t} \right) \mathbb{E} \left\| \theta_{i,k+1}^{t} - \theta_{t} \right\| + \eta_{k}^{t} \sum_{l=0}^{t} \left(1 - \gamma \right)^{l} \left(\gamma \mathbb{E} \left\| \nabla F \left(\theta_{l} \right) \right\| + \gamma \sigma \right)$$
(44)

where, a is following Assumption C.3, b is following Lemma D.1. then we can get,

$$\mathbb{E} \left\| \theta_{i,k+1}^{t} - \theta_{t} \right\| \leq \left(1 + \beta \eta_{k}^{t} \right) \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{t} \right\| + \eta_{k}^{t} \sum_{l=0}^{t} \left(1 - \gamma \right)^{l} \left(\gamma \mathbb{E} \left\| \nabla F \left(\theta_{l} \right) \right\| + \gamma \sigma \right) \\
\leq \left(1 + \beta \eta_{k}^{t} \right) \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{t} \right\| + \eta_{k}^{t} (L + \sigma) \\
\leq^{c} \sum_{k=0}^{K} \prod_{l=k}^{K} \left(1 + \beta \eta_{l}^{t} \right) \eta_{l}^{t} (L + \sigma) \\
\leq \tilde{\eta}_{t} (1 + c)^{k-1} (L + \sigma)$$
(45)

, c is following Assumption C.2. Below we find an upper bound for $\mathbb{E} \| \boldsymbol{m}^t \|$,

$$\mathbb{E} \|\boldsymbol{m}^{t}\| = \mathbb{E} \left\| (1-\gamma) \, \boldsymbol{m}^{t-1} + \gamma \frac{1}{m} \sum_{i=1}^{m} g_{i}\left(\theta_{t}\right) \right\|$$

$$\leq (1-\gamma) \mathbb{E} \|\boldsymbol{m}^{t-1}\| + \gamma \mathbb{E} \left\| \frac{1}{m} \sum_{i=1}^{m} g_{i}\left(\theta_{t}\right) \right\|$$

$$\leq (1-\gamma) \mathbb{E} \|\boldsymbol{m}^{t-1}\| + \gamma \mathbb{E} \|\nabla F\left(\theta_{t}\right)\| + \gamma \sigma$$

$$\leq \sum_{l=0}^{t} (1-\gamma)^{l} \left(\gamma \mathbb{E} \|\nabla F\left(\theta_{l}\right)\| + \gamma \sigma\right)$$

$$\leq L + \sigma$$

$$(46)$$

where we use Assumptions C.5. With $(1 + \beta \eta_k^t) \le c$, unrolling the above and noting $\theta_{i,0} = \theta_t$ yields

$$\mathbb{E} \left\| \theta_{i,K}^{t} - \theta_{t} \right\| \\
\leq (1+c)^{K-1} \tilde{\eta}_{t} \mathbb{E} \left\| \boldsymbol{m}^{t} \right\| \\
\leq \tilde{\eta}_{t} (1+c)^{K-1} \left(L + \sigma \right)$$
(47)

Lemma F.5. Given Assumptions C.2-C.5 and considering of FedMoSWA, for $\eta_k^t \leq 1/\beta KT$, we have

$$\mathbb{E}\left\|g_{i}\left(\theta_{i,k}^{t}\right)\right\| \leq \left(1 + \beta \tilde{\eta_{t}}(1+c)^{K-1}\right)\left(L+\sigma\right) + \sigma_{g,i}$$

where $g_i(\cdot)$ is the sampled gradient of client $i, \tilde{\eta}_t = \sum_{k=0}^{K-1} \eta_k^t$.

Proof.

$$\mathbb{E} \left\| g_{i}\left(\theta_{i,k}^{t}\right) \right\| \leq \mathbb{E} \left\| g_{i}\left(\theta_{i,k}^{t}\right) - \nabla F_{i}\left(\theta_{i,k}^{t}\right) \right\| + \mathbb{E} \left\| \nabla F_{i}\left(\theta_{i,k}^{t}\right) \right\| \\ \leq^{a} \mathbb{E} \left\| \nabla F_{i}\left(\theta_{i,k}^{t}\right) \right\| + \sigma \\ \leq \mathbb{E} \left\| \nabla F_{i}\left(\theta_{t}\right) \right\| + \mathbb{E} \left\| \nabla F_{i}\left(\theta_{i,k}^{t}\right) - \nabla F_{i}\left(\theta_{t}\right) \right\| + \sigma \\ \leq^{b} \mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \mathbb{E} \left\| \nabla F_{i}\left(\theta_{t}\right) - \nabla F\left(\theta_{t}\right) \right\| + \beta \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{t} \right\| + \sigma \\ \leq^{c} \left(\mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \sigma_{g,i} + \sigma \right) + \sum_{l=0}^{K-1} \eta_{l}^{t} \sum_{p=0}^{l} \left(1 - \gamma \right)^{p} \left(\gamma \mathbb{E} \left\| \nabla F\left(\theta_{p}\right) \right\| + \sigma \right) \left(1 + c \right)^{K-1-l} \\ \leq^{d} \left(1 + \beta \tilde{\eta_{t}} (1 + c)^{K-1} \right) \left(L + \sigma \right) + \sigma_{g,i}$$
(48)

a is following Assumption C.3, *b* is is following Assumption C.4, *c* is is following Assumption C.5 and Lemma D.1 and Lemma F.4, *d* is following $\|\nabla F(\theta_t)\| \leq L$.

Theorem F.6. (*FedMoSWA*). Suppose Assumptions C.2-C.5 hold and consider FedMoSWA (Algorithm 1). Let $k = K, \forall i \in [m]$ and $\eta_k^t \leq \frac{1}{\beta KT}$, $\tilde{c} = (1 + \beta \tilde{\eta}_t (1 + c)^{K-1}) (L + \sigma)$. Then,

$$\varepsilon_{gen} \leq \frac{2L}{mn} \frac{1}{\beta} e^{1 + \frac{1}{T}} \left(\tilde{c}L + \sigma_g + \tilde{c}\sigma \right)$$

where , $\tilde{c} \leq 1+(1+c)^{K-1}\beta\tilde{\eta}_t,$ $\tilde{\eta}_t=\sum_{k=0}^{K-1}\eta_k^t.$

Proof. For client *i*, there are two cases to consider. In the first case, SGD selects non-perturbed samples in S and $S^{(i)}$, which happens with probability 1 - 1/n. Then, we have,

$$\left\|\theta_{i,k+1}^{t} - \theta_{i,k+1}^{t,\prime}\right\| \leq \left(1 + \beta \eta_{k}^{t}\right) \left\|\theta_{i,k}^{t} - \theta_{i,k}^{t,\prime}\right\|.$$

$$\tag{49}$$

In the second case, SGD encounters the perturbed sample at time step k, which happens with probability 1/n. Then, we have

$$\left\| \theta_{i,k+1}^{t} - \theta_{i,k+1}^{t,\prime} \right\| = \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} - \eta_{k}^{t} \left(g_{i} \left(\theta_{i,k}^{t} \right) - g_{i}^{\prime} \left(\theta_{i,k}^{t,\prime} \right) - c_{i}^{t} + \boldsymbol{m}^{t} - c_{i}^{t} + \boldsymbol{m}^{t} \right) \right\|$$

$$\leq \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} - \eta_{k}^{t} \left(g_{i} \left(\theta_{i,k}^{t} \right) - g_{i} \left(\theta_{i,k}^{t,\prime} \right) \right) \right\| + \eta_{k}^{t} \left\| g_{i} \left(\theta_{i,k}^{t,\prime} \right) - g_{i}^{\prime} \left(\theta_{i,k}^{t,\prime} \right) \right\|$$

$$\leq^{a} \left(1 + \beta \eta_{k}^{t} \right) \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} \right\| + \eta_{k}^{t} \left\| g_{i} \left(\theta_{i,k}^{t,\prime} \right) - g_{i}^{\prime} \left(\theta_{i,k}^{t,\prime} \right) - g_{i}^{\prime} \left(\theta_{i,k}^{t,\prime} \right) \right\| .$$

$$(50)$$

where a is following Lemma D.1. Combining these two cases we have for client i,

$$\mathbb{E} \left\| \theta_{i,k+1}^{t} - \theta_{i,k+1}^{t,\prime} \right\| \leq \left(1 + \beta \eta_{k}^{t} \right) \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} \right\| + \frac{\eta_{k}^{t}}{n} \mathbb{E} \left\| g_{i} \left(\theta_{i,k}^{t,\prime} \right) - g_{i}^{\prime} \left(\theta_{i,k}^{t,\prime} \right) \right\| \\
\leq^{a} \left(1 + \beta \eta_{k}^{t} \right) \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} \right\| + \frac{2\eta_{k}^{t}}{n} \mathbb{E} \left\| g_{i} \left(\theta_{i,k}^{t} \right) \right\| \\
\leq^{c} \left(1 + \beta \eta_{k}^{t} \right) \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} \right\| + \frac{2\eta_{k}^{t}}{n} \left(\left(\mathbb{E} \left\| \nabla F \left(\theta_{t} \right) \right\| + \sigma_{g,i} + \sigma \right) + \left(1 + c \right)^{K-1} \tilde{\eta}_{t} \mathbb{E} \left\| \boldsymbol{m}^{t} \right\| \right) \\
\leq^{b} \left(1 + \beta \eta_{k}^{t} \right) \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} \right\| + \left(1 + \beta \tilde{\eta}_{t} (1 + c)^{K-1} \right) (L + \sigma) + \sigma_{g,i},$$
(51)

where the *a* inequality follows that $g_i(\cdot)$ and $g'_i(\cdot)$ are sampled from the same distribution, *b* is following Lemma F.4, *c* is following Lemma F.5. We let $1 + (1 - c)^{K-1} \beta \tilde{\eta}_t \leq \tilde{c}$, $\tilde{\eta}_t = \sum_{k=0}^{K-1} \eta_k^t$. Then unrolling it we have,

$$\mathbb{E}\left\|\theta_{i,K}^{t} - \theta_{i,K}^{t,\prime}\right\| \leq \prod_{k=0}^{K-1} \left(1 + \beta\eta_{k}^{t}\right) \mathbb{E}\left\|\theta_{t} - \theta_{t}^{\prime}\right\| \\
+ \left(\frac{2}{n}\sum_{k=0}^{K-1} \eta_{k}^{t}\prod_{l=k+1}^{K-1} \left(1 + \beta\eta_{k}^{t}\right) \cdot \left(\left(\mathbb{E}\left\|\nabla F\left(\theta_{t}\right)\right\| + \sigma_{g,i} + \sigma\right) + (1 + c)^{K-1}\tilde{\eta}_{t}\mathbb{E}\left\|\boldsymbol{m}^{t}\right\|\right)\right) \\
\leq e^{\beta\tilde{\eta}_{t}}\mathbb{E}\left\|\theta_{t} - \theta_{t}^{\prime}\right\| + \frac{2}{n}\tilde{\eta}_{t}e^{\beta\tilde{\eta}_{t}}\left(\left(\mathbb{E}\left\|\nabla F\left(\theta_{t}\right)\right\| + \sigma_{g,i} + \sigma\right) + (1 + c)^{K-1}\tilde{\eta}_{t}\mathbb{E}\left\|\boldsymbol{m}^{t}\right\|\right).$$
(52)

Then, with $\theta_{t+1} = \frac{1}{m} \sum_{i=1}^{m} \theta_{i,K}^t$ and $\theta'_{t+1} = \frac{1}{m} \sum_{i=1}^{m} \theta_{i,K}^{t,\prime}$, we have

$$\mathbb{E} \left\| \theta_{t+1} - \theta_{t+1}' \right\| \leq \sum_{i=1}^{m} \frac{1}{m} \mathbb{E} \left\| \theta_{i,K}^{t} - \theta_{i,K}^{t,\prime} \right\|
\leq e^{\beta \tilde{\eta}_{t}} \mathbb{E} \left\| \theta_{t} - \theta_{t}' \right\| + \frac{2}{mn} \tilde{\eta}_{t} e^{\beta \tilde{\eta}_{t}} \left((\mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \sigma_{g} + \sigma) + (1+c)^{K-1} \tilde{\eta}_{t} \mathbb{E} \left\| \boldsymbol{m}^{t} \right\| \right)$$
(53)

where we use Lemma F.4 in the last step. Further, unrolling the above over t and noting $\theta_0 = \theta'_0$, we obtain

When the diminishing stepsizes are chosen in the statement of the theorem, we conclude the proof.

G. Convergence Analysis for FedMoSWA under strongly convex setting

For the optimization problem, we consider:

$$\min_{\boldsymbol{\theta} \in \mathbb{R}^d} \left\{ F(\boldsymbol{\theta}) := \frac{1}{m} \sum_{i=1}^m \left(F_i(\boldsymbol{\theta}) := \mathbb{E}_{\zeta_i} \left[F_i(\boldsymbol{\theta}; \zeta_i) \right] \right) \right\}.$$

Before giving our theoretical results, we first present the common assumptions.

Assumption G.1. F_i is μ -strongly-convex for all $i \in [m]$, i.e.,

$$F_i(\boldsymbol{\theta_1}) \ge F_i(\boldsymbol{\theta_2}) + \langle \nabla F_i(\boldsymbol{\theta_2}), \boldsymbol{\theta_1} - \boldsymbol{\theta_2} \rangle + \frac{\mu}{2} \|\boldsymbol{\theta_1} - \boldsymbol{\theta_2}\|^2$$
(55)

for all θ_2, θ_1 in its domain and $i \in [m]$. We allow $\mu = 0$, which corresponds to general convex functions.

Assumption G.2. (Smoothness) The gradient of the loss function is Lipschitz continuous with constant β , for all $\theta_1, \theta_2 \in \mathbb{R}^d$

$$\left\|\nabla F\left(\boldsymbol{\theta}_{1}\right) - \nabla F\left(\boldsymbol{\theta}_{2}\right)\right\| \leq \beta \left\|\boldsymbol{\theta}_{1} - \boldsymbol{\theta}_{2}\right\|$$
(56)

Assumption G.3. Let ζ be a mini-batch drawn uniformly at random from all samples. We assume that the data is distributed so that, for all $\theta \in \mathbb{R}^d$

$$\mathbb{E}_{\zeta|\boldsymbol{\theta}}\left[\nabla F_i(\boldsymbol{\theta};\zeta)\right] = \nabla F_i(\boldsymbol{\theta}). \tag{57}$$

We also can get:

$$\mathbb{E}_{\zeta|\boldsymbol{\theta}}\left[\left\|\nabla F_{i}\left(\boldsymbol{\theta};\zeta_{i}\right)-\nabla F_{i}(\boldsymbol{\theta})\right\|^{2}\right]\leq\sigma^{2}.$$

Assumption G.4. (Bounded heterogeneity) The dissimilarity of $F_i(\theta)$ and $f(\theta)$ is bounded as follows:

$$\frac{1}{m}\sum_{i=1}^{m} \|\nabla F_i(\boldsymbol{\theta}) - \nabla F(\boldsymbol{\theta})\|^2 \le \sigma_g^2.$$
(58)

Assumption G.3 bounds the variance of stochastic gradients, which is common in stochastic optimization analysis (Bubeck et al., 2015). Assumption G.4 bounds the gradient difference between global and local loss functions, which is a widely-used approach to characterize client heterogeneity in federated optimization literature (Li et al., 2020; Reddi et al., 2021).

G.1. Some technical lemmas

Now we cover some technical lemmas which are useful for computations later on. The two lemmas below are useful to unroll recursions and derive convergence rates.

Lemma G.5 (linear convergence rate). For every non-negative sequence $\{d_{r-1}\}_{r\geq 1}$ and any parameters $\mu > 0$, $\eta_{\max} \in (0, 1/\mu]$, $c \geq 0$, $R \geq \frac{1}{2\eta_{\max}\mu}$, there exists a constant step-size $\eta \leq \eta_{\max}$ and weights $w_r := (1 - \mu\eta)^{1-r}$ such that for $W_R := \sum_{r=1}^{R+1} w_r$,

$$\Psi_{R} := \frac{1}{W_{R}} \sum_{r=1}^{R+1} \left(\frac{w_{r}}{\eta} (1 - \mu \eta) d_{r-1} - \frac{w_{r}}{\eta} d_{r} + c \eta w_{r} \right) = \tilde{\mathcal{O}} \left(\mu d_{0} \exp\left(-\mu \eta_{\max} R \right) + \frac{c}{\mu R} \right)$$
(59)

Proof. By substituting the value of w_r , we observe that we end up with a telescoping sum and estimate

$$\Psi_R = \frac{1}{\eta W_R} \sum_{r=1}^{R+1} \left(w_{r-1} d_{r-1} - w_r d_r \right) + \frac{c\eta}{W_R} \sum_{r=1}^{R+1} w_r \le \frac{d_0}{\eta W_R} + c\eta$$

When $R \ge \frac{1}{2\mu\eta}$, $(1 - \mu\eta)^R \le \exp(-\mu\eta R) \le \frac{2}{3}$. For such an R, we can lower bound ηW_R using

$$\eta W_R = \eta (1 - \mu \eta)^{-R} \sum_{r=0}^R (1 - \mu \eta)^r = \eta (1 - \mu \eta)^{-R} \frac{1 - (1 - \mu \eta)^R}{\mu \eta} \ge (1 - \mu \eta)^{-R} \frac{1}{3\mu}.$$

This proves that for all $R \ge \frac{1}{2un}$,

$$\Psi_R \le 3\mu d_0 (1-\mu\eta)^R + c\eta \le 3\mu d_o \exp(-\mu\eta R) + c\eta$$

The lemma now follows by carefully tuning η . Consider the following two cases depending on the magnitude of R and η_{max} :

Suppose $\frac{1}{2\mu R} \leq \eta_{\max} \leq \frac{\log(\max(1,\mu^2 R d_0/c))}{\mu R}$. Then we can choose $\eta = \eta_{\max}$,

$$\Psi_R \le 3\mu d_0 \exp\left[-\mu\eta_{\max}R\right] + c\eta_{\max} \le 3\mu d_0 \exp\left[-\mu\eta_{\max}R\right] + \tilde{\mathcal{O}}\left(\frac{c}{\mu R}\right)$$

Instead if $\eta_{\max} > \frac{\log(\max(1,\mu^2 R d_0/c))}{\mu R}$, we pick $\eta = \frac{\log(\max(1,\mu^2 R d_0/c))}{\mu R}$ to claim that

$$\Psi_R \leq 3\mu d_0 \exp\left[-\log\left(\max\left(1, \mu^2 R d_0/c\right)\right)\right] + \tilde{\mathcal{O}}\left(\frac{c}{\mu R}\right) \leq \tilde{\mathcal{O}}\left(\frac{c}{\mu R}\right)$$

Lemma G.6 (sub-linear convergence rate). For every non-negative sequence $\{d_{r-1}\}_{r\geq 1}$ and any parameters $\eta_{\max} \geq 0$, $c \geq 0$, $R \geq 0$, there exists a constant step-size $\eta \leq \eta_{\max}$ and weights $w_r = 1$ such that,

$$\Psi_R := \frac{1}{R+1} \sum_{r=1}^{R+1} \left(\frac{d_{r-1}}{\eta} - \frac{d_r}{\eta} + c_1 \eta + c_2 \eta^2 \right) \le \frac{d_0}{\eta_{\max}(R+1)} + \frac{2\sqrt{c_1 d_0}}{\sqrt{R+1}} + 2\left(\frac{d_0}{R+1}\right)^{\frac{2}{3}} c_2^{\frac{1}{3}}.$$
 (60)

Proof. Unrolling the sum, we can simplify

$$\Psi_R \le \frac{d_0}{\eta(R+1)} + c_1 \eta + c_2 \eta^2$$
.

Similar to the strongly convex case (Lemma G.5), we distinguish the following cases:

• When $R+1 \leq \frac{d_0}{c_1 \eta_{\max}^2}$, and $R+1 \leq \frac{d_0}{c_2 \eta_{\max}^3}$ we pick $\eta = \eta_{\max}$ to claim

$$\Psi_R \le \frac{d_0}{\eta_{\max}(R+1)} + c_1 \eta_{\max} + c_2 \eta_{\max}^2 \le \frac{d_0}{\eta_{\max}(R+1)} + \frac{\sqrt{c_1 d_0}}{\sqrt{R+1}} + \left(\frac{d_0}{R+1}\right)^{\frac{1}{3}} c_2^{\frac{1}{3}}$$
(61)

• In the other case, we have $\eta_{\max}^2 \ge \frac{d_0}{c_1(R+1)}$ or $\eta_{\max}^3 \ge \frac{d_0}{c_2(R+1)}$. We choose $\eta == \min\left\{\sqrt{\frac{d_0}{c_1(R+1)}}, \sqrt[3]{\frac{d_0}{c_2(R+1)}}\right\}$ to prove

$$\Psi_R \le \frac{d_0}{\eta(R+1)} + c\eta = \frac{2\sqrt{c_1 d_0}}{\sqrt{R+1}} + 2\sqrt[3]{\frac{d_0^2 c_2}{(R+1)^2}}.$$

Next, we state a relaxed triangle inequality true for the squared ℓ_2 norm.

Lemma G.7 (relaxed triangle inequality). Lemma 3 (relaxed triangle inequality). Let $\{v_1, \ldots, v_{\tau}\}$ be τ vectors in \mathbb{R}^d . Then the following are true:

$$\begin{split} I. \ \|\boldsymbol{v}_{i} + \boldsymbol{v}_{j}\|^{2} &\leq (1+a) \|\boldsymbol{v}_{i}\|^{2} + \left(1 + \frac{1}{a}\right) \|\boldsymbol{v}_{j}\|^{2} \text{ for any } a > 0, \text{ and} \\ 2. \ \|\sum_{i=1}^{\tau} \boldsymbol{v}_{i}\|^{2} &\leq \tau \sum_{i=1}^{\tau} \|\boldsymbol{v}_{i}\|^{2}. \end{split}$$

Proof. The proof of the first statement for any a > 0 follows from the identity:

$$\|\boldsymbol{v}_{i} + \boldsymbol{v}_{j}\|^{2} = (1+a) \|\boldsymbol{v}_{i}\|^{2} + \left(1 + \frac{1}{a}\right) \|\boldsymbol{v}_{j}\|^{2} - \left\|\sqrt{a}\boldsymbol{v}_{i} + \frac{1}{\sqrt{a}}\boldsymbol{v}_{j}\right\|^{2}.$$

For the second inequality, we use the convexity of $x \to \|x\|^2$ and Jensen's inequality

$$\left\|rac{1}{ au}\sum_{i=1}^{ au}oldsymbol{v}_i
ight\|^2\leqrac{1}{ au}\sum_{i=1}^{ au}\|oldsymbol{v}_i\|^2$$

Next we state an elementary lemma about expectations of norms of random vectors.

2

Lemma G.8 (separating mean and variance). Let $\{\Xi_1, \ldots, \Xi_\tau\}$ be τ random variables in \mathbb{R}^d which are not necessarily independent. First suppose that their mean is $\mathbb{E}[\Xi_i] = \xi_i$ and variance is bounded as $\mathbb{E}\left[||\Xi_i - \xi_i||^2\right] \leq \sigma^2$. Then, the following holds

$$\mathbb{E}\left[\left\|\sum_{i=1}^{\tau} \Xi_i\right\|^2\right] \le \left\|\sum_{i=1}^{\tau} \xi_i\right\|^2 + \tau^2 \sigma^2$$

Now instead suppose that their conditional mean is $\mathbb{E}[\Xi_i | \Xi_{i-1}, \ldots, \Xi_1] = \xi_i$ i.e. the variables $\{\Xi_i - \xi_i\}$ form a martingale difference sequence, and the variance is bounded by $\mathbb{E}\left[\|\Xi_i - \xi_i\|^2\right] \leq \sigma^2$ as before. Then we can show the tighter bound

$$\mathbb{E}\left[\left\|\sum_{i=1}^{\tau} \Xi_i\right\|^2\right] \le 2\left\|\sum_{i=1}^{\tau} \xi_i\right\|^2 + 2\tau\sigma^2.$$

Proof. For any random variable $X, \mathbb{E}[X^2] = (\mathbb{E}[X - \mathbb{E}[X]])^2 + (\mathbb{E}[X])^2$ implying

$$\mathbb{E}\left[\left\|\sum_{i=1}^{\tau}\Xi_i\right\|^2\right] = \left\|\sum_{i=1}^{\tau}\xi_i\right\|^2 + \mathbb{E}\left[\left\|\sum_{i=1}^{\tau}\Xi_i - \xi_i\right\|^2\right].$$

Expanding the above expression using relaxed triangle inequality (Lemma 3) proves the first claim:

$$\mathbb{E}\left[\left\|\sum_{i=1}^{\tau}\Xi_i - \xi_i\right\|^2\right] \le \tau \sum_{i=1}^{\tau}\mathbb{E}\left[\left\|\Xi_i - \xi_i\right\|^2\right] \le \tau^2 \sigma^2.$$

For the second statement, ξ_i is not deterministic and depends on Ξ_{i-1}, \ldots, Ξ_1 . Hence we have to resort to the cruder relaxed triangle inequality to claim

$$\mathbb{E}\left[\left\|\sum_{i=1}^{\tau}\Xi_{i}\right\|^{2}\right] \leq 2\left\|\sum_{i=1}^{\tau}\xi_{i}\right\|^{2} + 2\mathbb{E}\left[\left\|\sum_{i=1}^{\tau}\Xi_{i} - \xi_{i}\right\|^{2}\right]$$

and then use the tighter expansion of the second term:

$$\mathbb{E}\left[\left\|\sum_{i=1}^{\tau}\Xi_{i}-\xi_{i}\right\|^{2}\right]=\sum_{i,j}\mathbb{E}\left[\left(\Xi_{i}-\xi_{i}\right)^{\top}\left(\Xi_{j}-\xi_{j}\right)\right]=\sum_{i}\mathbb{E}\left[\left\|\Xi_{i}-\xi_{i}\right\|^{2}\right]\leq\tau\sigma^{2}$$

The cross terms in the above expression have zero mean since $\{\Xi_i - \xi_i\}$ form a martingale difference sequence.

G.2. Convergence of FedMoSWA for strongly convex functions

We will first bound the variance of FedMoSWA update in Lemma G.9, then see how sampling of clients effects our control variates in Lemma G.10, and finally bound the amount of client-drift in Lemma G.12. We will then use these three lemmas to prove the progress in a single round in Lemma G.13. Combining this progress with Lemmas G.5 and G.6 gives us the desired rates. Before proceeding with the proof of our lemmas, we need some additional definitions of the various errors we track. As before, we define the effective step-size to be

$$\tilde{\eta} := K\eta_k^t \alpha = K\eta_l \alpha.$$

We define client-drift to be how much the clients move from their starting point:

$$\mathcal{E}_t := rac{1}{Km} \sum_{k=1}^K \sum_{i=1}^m \mathbb{E}\left[\left\| oldsymbol{ heta}_{i,k}^t - oldsymbol{ heta}^{t-1}
ight\|^2
ight].$$

Because we are sampling the clients, not all the client control-variates get updated every round. This leads to some 'lag' which we call control-lag:

$$\mathcal{C}_{t} := \frac{1}{m} \sum_{j=1}^{m} \mathbb{E} \left\| \mathbb{E} \left[\boldsymbol{c}_{i}^{t} \right] - \nabla F_{i} \left(\boldsymbol{\theta}^{\star} \right) \right\|^{2}.$$
$$\mathcal{D}_{t} := \mathbb{E} \left\| \boldsymbol{m}^{t} - \nabla F \left(\boldsymbol{\theta}^{\star} \right) \right\|^{2} = \mathbb{E} \left\| \boldsymbol{m}^{t} \right\|^{2}.$$

Variance of server update. We study how the variance of the server update can be bounded.

Lemma G.9. For updates, we can bound the variance of the server update in any round r and any $\tilde{\eta} := K \eta_k^t \alpha \ge 0$ as follows

$$\mathbb{E}\left[\left\|\boldsymbol{\theta}^{t}-\boldsymbol{\theta}^{t-1}\right\|^{2}\right] \leq 8\beta\tilde{\eta}^{2}\left(\mathbb{E}\left[f\left(\boldsymbol{\theta}^{t-1}\right)\right]-f\left(\boldsymbol{\theta}^{\star}\right)\right)+4\tilde{\eta}^{2}\mathcal{C}_{t-1}+4\tilde{\eta}^{2}\mathcal{D}_{t-1}+4\tilde{\eta}^{2}\beta^{2}\mathcal{E}_{t}+\frac{12\tilde{\eta}^{2}\sigma^{2}}{Ks}\right]$$

Proof. The server update in round t can be written as follows (dropping the superscript t everywhere)

$$\mathbb{E}\|\Delta\boldsymbol{\theta}\|^{2} = \mathbb{E}\left\|-\frac{\tilde{\eta}}{Ks}\sum_{k,i\in\mathcal{S}}\boldsymbol{\theta}_{i,k}\right\|^{2} = \mathbb{E}\left\|\frac{\tilde{\eta}}{Ks}\sum_{k,i\in\mathcal{S}}\left(g_{i}\left(\boldsymbol{\theta}_{i,k-1}\right) + \boldsymbol{m} - \boldsymbol{c}_{i}\right)\right\|^{2},\tag{62}$$

which can then be expanded as

$$\mathbb{E} \|\Delta \theta\|^{2} \leq \mathbb{E} \left\| \frac{\tilde{\eta}}{Ks} \sum_{k,i \in S} \left(g_{i}\left(\theta_{i,k-1}\right) + \boldsymbol{m} - \boldsymbol{c}_{i} \right) \right\|^{2} \\
\leq 4\mathbb{E} \left\| \frac{\tilde{\eta}}{Ks} \sum_{k,i \in S} g_{i}\left(\theta_{i,k-1}\right) - \nabla F_{i}\left(\theta\right) \right\|^{2} + 4\tilde{\eta}^{2}\mathbb{E} \|\boldsymbol{m}\|^{2} + 4\mathbb{E} \left\| \frac{\tilde{\eta}}{Ks} \sum_{k,i \in S} \nabla F_{i}\left(\theta^{*}\right) - \boldsymbol{c}_{i} \right\|^{2} \\
+ 4\mathbb{E} \left\| \frac{\tilde{\eta}}{Ks} \sum_{k,i \in S} \nabla F_{i}\left(\theta\right) - \nabla F_{i}\left(\theta^{*}\right) \right\|^{2} \\
\leq 4\mathbb{E} \left\| \frac{\tilde{\eta}}{Ks} \sum_{k,i \in S} g_{i}\left(\theta_{i,k-1}\right) - \nabla F_{i}\left(\theta\right) \right\|^{2} + 4\tilde{\eta}^{2}\mathbb{E} \|\boldsymbol{m}\|^{2} + 4\mathbb{E} \left\| \frac{\tilde{\eta}}{s} \sum_{i \in S} \nabla F_{i}\left(\theta^{*}\right) - \boldsymbol{c}_{i} \right\|^{2} \\
+ 8\beta\tilde{\eta}^{2}\left(\mathbb{E}[f(\theta)] - f\left(\theta^{*}\right)\right) \\
\leq 4\mathbb{E} \left\| \frac{\tilde{\eta}}{Ks} \sum_{k,i \in S} \nabla F_{i}\left(\theta_{i,k-1}\right) - \nabla F_{i}\left(\theta\right) \right\|^{2} + 4\tilde{\eta}^{2} \|\mathbb{E}[\boldsymbol{m}]\|^{2} + 4 \left\| \frac{\tilde{\eta}}{s} \sum_{i \in S} \nabla F_{i}\left(\theta^{*}\right) - \mathbb{E}\left[\boldsymbol{c}_{i}\right] \right\|^{2} \tag{64} \\
+ 8\beta\tilde{\eta}^{2}\left(\mathbb{E}[f(\theta)] - f\left(\theta^{*}\right)\right) + \frac{12\tilde{\eta}^{2}\sigma^{2}}{Ks}.$$

The inequality before the last used the smoothness of $\{F_i\}$. The variance of $\left(\frac{1}{Ks}\sum_{k,i\in\mathcal{S}}g_i\left(\boldsymbol{\theta}_{i,k-1}\right)\right)$ is bounded by σ^2/Ks . Similarly, \boldsymbol{c}_j for any $j\in[m]$ has variance smaller than σ^2/K and hence the variance of $\left(\frac{1}{m}\sum_{i\in\mathcal{S}}\boldsymbol{c}_i\right)$ is smaller

than σ^2/Ks .

$$\mathbb{E} \|\Delta \boldsymbol{\theta}\|^{2} \leq \frac{4\tilde{\eta}^{2}}{Km} \sum_{k,i} \mathbb{E} \|\nabla F_{i}\left(\boldsymbol{\theta}_{i,k-1}\right) - \nabla F_{i}(\boldsymbol{\theta})\|^{2} + 4\tilde{\eta}^{2} \|\mathbb{E}[\boldsymbol{m}]\|^{2} + \frac{4\tilde{\eta}^{2}}{m} \sum_{i} \|\nabla F_{i}\left(\boldsymbol{\theta}^{\star}\right) - \mathbb{E}\left[\boldsymbol{c}_{i}\right]\|^{2} \\
+ 8\beta\tilde{\eta}^{2}\left(\mathbb{E}[f(\boldsymbol{\theta})] - f\left(\boldsymbol{\theta}^{\star}\right)\right) + \frac{12\tilde{\eta}^{2}\sigma^{2}}{Ks} \\
\leq \underbrace{\frac{4\tilde{\eta}^{2}}{Km} \sum_{k,i} \mathbb{E} \|\nabla F_{i}\left(\boldsymbol{\theta}_{i,k-1}\right) - \nabla F_{i}(\boldsymbol{\theta})\|^{2}}_{\mathcal{T}_{1}} + \underbrace{\frac{4\tilde{\eta}^{2}}{m} \sum_{i} \|\nabla F_{i}\left(\boldsymbol{\theta}^{\star}\right) - \mathbb{E}\left[\boldsymbol{c}_{i}\right]\|^{2} + 4\tilde{\eta}^{2} \|\mathbb{E}[\boldsymbol{m}]\|^{2}}_{\mathcal{T}_{1}} \\
+ 8\beta\tilde{\eta}^{2}\left(\mathbb{E}[f(\boldsymbol{\theta})] - f\left(\boldsymbol{\theta}^{\star}\right)\right) + \frac{12\tilde{\eta}^{2}\sigma^{2}}{Ks}.$$
(65)

Since the gradient of F_i is β -Lipschitz, $\mathcal{T}_1 \leq \frac{\beta^2 4 \tilde{\eta}^2}{Km} \sum_{k,i} \mathbb{E} \|\boldsymbol{\theta}_{i,k-1} - \boldsymbol{\theta}\|^2 = 4 \tilde{\eta}^2 \beta^2 \mathcal{E}$. The definition of the error in the control variate $\mathcal{C}_{t-1} := \frac{1}{m} \sum_{j=1}^m \mathbb{E} \|\mathbb{E}[\boldsymbol{c}_i] - \nabla F_i(\boldsymbol{\theta}^*)\|^2$ completes the proof.

Change in control lag. We have previously related the variance of the server update to the control lag. We now examine how the control-lag grows each round.

Lemma G.10. For updates with the control update and Assumptions G.1–G.2, the following holds true for any $\tilde{\eta} := \eta_k^t \alpha K \in [0, 1/\beta]$:

$$\mathcal{C}_{t} \leq \left(1 - \frac{s}{m}\right) \mathcal{C}_{t-1} + \frac{s}{m} \left(4\beta \left(\mathbb{E}\left[f\left(\boldsymbol{\theta}^{t-1}\right)\right] - f\left(\boldsymbol{\theta}^{\star}\right)\right) + 2\beta^{2} \mathcal{E}_{t}\right).$$

Proof. Recall that after round t, the control update rule implies that c_i^t is set as per

$$\boldsymbol{c}_{i}^{t} = \begin{cases} \boldsymbol{c}_{i}^{t-1} & \text{if } i \notin \mathcal{S}^{t} \text{ i.e. with probability } \left(1 - \frac{s}{m}\right), \\ \frac{1}{K} \sum_{k=1}^{K} g_{i}\left(\boldsymbol{\theta}_{i,k-1}^{t}\right) & \text{with probability } \frac{s}{m}. \end{cases}$$

Taking expectations on both sides yields

$$\mathbb{E}\left[\boldsymbol{c}_{i}^{t}\right] = \left(1 - \frac{s}{m}\right) \mathbb{E}\left[\boldsymbol{c}_{i}^{t-1}\right] + \frac{s}{Km} \sum_{k=1}^{K} \mathbb{E}\left[\nabla F_{i}\left(\boldsymbol{\theta}_{i,k-1}^{t}\right)\right], \forall i \in [m].$$

Plugging the above expression in the definition of C_t we get,

$$\mathcal{C}_{t} = \frac{1}{m} \sum_{i=1}^{m} \left\| \mathbb{E} \left[\boldsymbol{c}_{i}^{t} \right] - \nabla F_{i} \left(\boldsymbol{\theta}^{\star} \right) \right\|^{2}$$

$$= \frac{1}{m} \sum_{i=1}^{m} \left\| \left(1 - \frac{s}{m} \right) \left(\mathbb{E} \left[\boldsymbol{c}_{i}^{t-1} \right] - \nabla F_{i} \left(\boldsymbol{\theta}^{\star} \right) \right) + \frac{s}{m} \left(\frac{1}{K} \sum_{k=1}^{K} \mathbb{E} \left[\nabla F_{i} \left(\boldsymbol{\theta}_{i,k-1}^{t} \right) \right] - \nabla F_{i} \left(\boldsymbol{\theta}^{\star} \right) \right) \right\|^{2} \qquad (66)$$

$$\leq \left(1 - \frac{s}{m} \right) \mathcal{C}_{t-1} + \frac{s}{m^{2}K} \sum_{k=1}^{K} \mathbb{E} \left\| \nabla F_{i} \left(\boldsymbol{\theta}_{i,k-1}^{t} \right) - \nabla F_{i} \left(\boldsymbol{\theta}^{\star} \right) \right\|^{2}.$$

The final step applied Jensen's inequality twice. We can then further simplify using the relaxed triangle inequality as,

$$\mathbb{E}_{t-1} \left[\mathcal{C}_{t} \right] \leq \left(1 - \frac{s}{m} \right) \mathcal{C}_{t-1} + \frac{s}{m^{2}K} \sum_{i,k} \mathbb{E} \left\| \nabla F_{i} \left(\boldsymbol{\theta}_{i,k-1}^{t} \right) - \nabla F_{i} \left(\boldsymbol{\theta}^{\star} \right) \right\|^{2} \\
\leq \left(1 - \frac{s}{m} \right) \mathcal{C}_{t-1} + \frac{2s}{m^{2}} \sum_{i} \mathbb{E} \left\| \nabla F_{i} \left(\boldsymbol{\theta}^{t-1} \right) - \nabla F_{i} \left(\boldsymbol{\theta}^{\star} \right) \right\|^{2} + \frac{2s}{m^{2}K} \sum_{i,k} \mathbb{E} \left\| \nabla F_{i} \left(\boldsymbol{\theta}_{i,k-1}^{t} \right) - \nabla F_{i} \left(\boldsymbol{\theta}^{t-1} \right) \right\|^{2} \\
\leq \left(1 - \frac{s}{m} \right) \mathcal{C}_{t-1} + \frac{2s}{m^{2}} \sum_{i} \mathbb{E} \left\| \nabla F_{i} \left(\boldsymbol{\theta}^{t-1} \right) - \nabla F_{i} \left(\boldsymbol{\theta}^{\star} \right) \right\|^{2} + \frac{2s}{m^{2}K} \beta^{2} \sum_{i,k} \mathbb{E} \left\| \boldsymbol{\theta}_{i,k-1}^{t} - \boldsymbol{\theta}^{t-1} \right\|^{2} \\
\leq \left(1 - \frac{s}{m} \right) \mathcal{C}_{t-1} + \frac{s}{m} \left(4\beta \left(\mathbb{E} \left[f \left(\boldsymbol{\theta}^{t-1} \right) \right] - f \left(\boldsymbol{\theta}^{\star} \right) \right) + \beta^{2} \mathcal{E}_{t} \right).$$
(67)

The last two inequalities follow from smoothness of $\{F_i\}$ and the definition $\mathcal{E}_t = \frac{1}{mK}\beta^2 \sum_{i,k} \mathbb{E} \left\| \boldsymbol{\theta}_{i,k-1}^t - \boldsymbol{\theta}^{t-1} \right\|^2$. \Box

Lemma G.11. Assumptions G.1–G.2, the following holds true for any $\tilde{\eta} := \eta_k^t \alpha K \in [0, 1/\beta]$:

$$\mathcal{D}_{t} \leq (1-\gamma) \mathcal{D}_{t-1} + \gamma \left(4\beta \left(\mathbb{E} \left[f \left(\boldsymbol{\theta}^{t-1} \right) \right] - f \left(\boldsymbol{\theta}^{\star} \right) \right) + 2\beta^{2} \mathcal{E}_{t} \right).$$

Proof. Recall that after round t, the control update rule implies that c_i^t is set as

$$\boldsymbol{m}^{t} = \boldsymbol{m}^{t-1} + \beta \frac{1}{m} \sum_{i=1}^{m} \left(c_{i}^{t} - \boldsymbol{m}^{t-1} \right) = (1 - \beta) \boldsymbol{m}^{t-1} + \beta \frac{1}{m} \sum_{i=1}^{m} c_{i}^{t}$$
(68)

$$\mathbb{E} \left\| \boldsymbol{m}^{t} - \nabla F\left(\boldsymbol{\theta}^{\star}\right) \right\|^{2} = \mathbb{E} \left\| (1-\gamma) \, \boldsymbol{m}^{t-1} + \gamma \frac{1}{m} \sum_{i=1}^{m} \frac{1}{K} \sum_{k=1}^{K} \nabla F_{i}\left(\boldsymbol{\theta}_{i,k-1}^{t}; \zeta\right) - \nabla F\left(\boldsymbol{x}^{\star}\right) \right\|^{2}$$

$$\leq (1-\gamma) \mathbb{E} \left\| \boldsymbol{m}^{t-1} - \nabla F\left(\boldsymbol{\theta}^{\star}\right) \right\|^{2} + \gamma \mathbb{E} \left\| \frac{1}{m} \sum_{i=1}^{m} \frac{1}{K} \sum_{k=1}^{K} \nabla F_{i}\left(\boldsymbol{\theta}_{i,k-1}^{t}; \zeta\right) - \nabla F\left(\boldsymbol{\theta}^{\star}\right) \right\|^{2}$$

$$\leq (1-\gamma) \mathbb{E} \left\| \boldsymbol{m}^{t-1} - \nabla F\left(\boldsymbol{\theta}^{\star}\right) \right\|^{2} + \frac{2}{m} \gamma \mathbb{E} \sum_{i=1}^{m} \left\| \nabla F\left(\boldsymbol{\theta}^{t-1}\right) - \nabla F\left(\boldsymbol{\theta}^{\star}\right) \right\|^{2}$$

$$+ \frac{2}{mK} \gamma \sum_{i,k} \mathbb{E} \left\| \nabla F_{i}\left(\boldsymbol{\theta}_{i}^{t-1,t}\right) - \nabla F_{i}\left(\boldsymbol{\theta}^{t-1}\right) \right\|^{2}$$

$$\leq (1-\gamma) \mathbb{E} \left\| \boldsymbol{m}^{t-1} - \nabla F\left(\boldsymbol{\theta}^{\star}\right) \right\|^{2} + \gamma \left(4\beta \left(\mathbb{E} \left[f\left(\boldsymbol{\theta}^{t-1}\right) \right] - f\left(\boldsymbol{\theta}^{\star}\right) \right) + \beta^{2} \mathcal{E}_{t} \right)$$

Bounding client-drift. We will now bound the final source of error which is the client-drift.

Lemma G.12. Suppose our step-sizes satisfy $\eta_k^t = \eta_l \leq \frac{1}{81\beta K\alpha}$ and F_i satisfies Assumptions G.1–G.2. Then, for any global $\alpha \geq 1$ we can bound the drift as

$$3\beta\tilde{\eta}\mathcal{E}_{t} \leq \frac{\tilde{\eta}^{2}}{3}\mathcal{C}_{t-1} + \frac{\tilde{\eta}^{2}}{3}\mathcal{D}_{t-1} + \frac{\tilde{\eta}}{25\alpha^{2}}\left(\mathbb{E}\left[f\left(\boldsymbol{\theta}^{t-1}\right)\right] - f\left(\boldsymbol{\theta}^{\star}\right)\right) + \frac{\tilde{\eta}^{2}}{K\alpha^{2}}\sigma^{2}.$$

Proof. First, observe that if K = 1, $\mathcal{E}_t = 0$ since $\theta_{i,0} = \theta$ for all $i \in [m]$ and that \mathcal{C}_{t-1} and the right hand side are both positive. For K > 1, we build a recursive bound of the drift. Starting from the definition of the update and then applying the

relaxed triangle inequality, we can expand

$$\frac{1}{m}\mathbb{E}_{t-1}\left[\sum_{i\in\mathcal{S}}\|\boldsymbol{\theta}_{i}-\eta_{l}g_{i}\left(\boldsymbol{\theta}_{i}\right)+\eta_{l}\boldsymbol{m}-\eta_{l}\boldsymbol{c}_{i}-\boldsymbol{\theta}\|^{2}\right] \\
\leq \frac{1}{m}\mathbb{E}_{t-1}\left[\sum_{i\in\mathcal{S}}\|\boldsymbol{\theta}_{i}-\eta_{l}\nabla F_{i}\left(\boldsymbol{\theta}_{i}\right)+\eta_{l}\boldsymbol{m}-\eta_{l}\boldsymbol{c}_{i}-\boldsymbol{\theta}\|^{2}\right]+\eta_{l}^{2}\sigma^{2} \\
\leq \frac{(1+a)}{s}\mathbb{E}_{t-1}\left[\sum_{i\in\mathcal{S}}\underbrace{\|\boldsymbol{\theta}_{i}-\eta_{l}\nabla F_{i}\left(\boldsymbol{\theta}_{i}\right)+\eta_{l}\nabla F_{i}\left(\boldsymbol{\theta}\right)-\boldsymbol{\theta}\|^{2}}_{\mathcal{T}_{2}}\right] \\
+\left(1+\frac{1}{a}\right)\eta_{l}^{2}\underbrace{\mathbb{E}_{t-1}\left[\frac{1}{m}\sum_{i\in\mathcal{S}}\|\boldsymbol{m}-\boldsymbol{c}_{i}+\nabla F_{i}\left(\boldsymbol{\theta}\right)\|^{2}\right]}_{\mathcal{T}_{3}}+\eta_{l}^{2}\sigma^{2}.$$
(70)

The final step follows from the relaxed triangle inequality (Lemma G.7). Applying the contractive mapping Lemma for $\eta_l \leq 1/\beta$ shows

$$\mathcal{T}_{2} = \frac{1}{m} \sum_{i \in \mathcal{S}} \left\| \boldsymbol{\theta}_{i} - \eta_{l} \nabla F_{i}\left(\boldsymbol{\theta}_{i}\right) + \eta_{l} \nabla F_{i}(\boldsymbol{\theta}) - \boldsymbol{\theta} \right\|^{2} \le \left\| \boldsymbol{\theta}_{i} - \boldsymbol{\theta} \right\|^{2}.$$
(71)

Once again using our relaxed triangle inequality to expand the other term \mathcal{T}_3 , we get

$$\begin{aligned} \mathcal{T}_{3} &= \mathbb{E}_{t-1} \left[\frac{1}{m} \sum_{i \in \mathcal{S}} \|\boldsymbol{m} - \boldsymbol{c}_{i} + \nabla F_{i}(\boldsymbol{\theta})\|^{2} \right] \\ &= \frac{1}{m} \sum_{j=1}^{m} \|\boldsymbol{m} - \boldsymbol{c}_{i} + \nabla F_{i}(\boldsymbol{\theta})\|^{2} \\ &= \frac{1}{m} \sum_{j=1}^{m} \|\boldsymbol{m} - \boldsymbol{c}_{i} + \nabla F_{i}(\boldsymbol{\theta}^{\star}) + \nabla F_{i}(\boldsymbol{\theta}) - \nabla F_{i}(\boldsymbol{\theta}^{\star})\|^{2} \\ &\leq 3 \|\boldsymbol{m}\|^{2} + \frac{3}{m} \sum_{j=1}^{m} \|\boldsymbol{c}_{i} - \nabla F_{i}(\boldsymbol{\theta}^{\star})\|^{2} + \frac{3}{m} \sum_{j=1}^{m} \|\nabla F_{i}(\boldsymbol{\theta}) - \nabla F_{i}(\boldsymbol{\theta}^{\star})\|^{2} \\ &\leq 3 \|\boldsymbol{m}\|^{2} + \frac{3}{m} \sum_{j=1}^{m} \|\boldsymbol{c}_{i} - \nabla F_{i}(\boldsymbol{\theta}^{\star})\|^{2} + \frac{3}{m} \sum_{j=1}^{m} \|\nabla F_{i}(\boldsymbol{\theta}) - \nabla F_{i}(\boldsymbol{\theta}^{\star})\|^{2} \\ &\leq 3 \|\boldsymbol{m}\|^{2} + \frac{3}{m} \sum_{j=1}^{m} \|\boldsymbol{c}_{i} - \nabla F_{i}(\boldsymbol{\theta}^{\star})\|^{2} + 6\beta \left(f(\boldsymbol{\theta}) - f(\boldsymbol{\theta}^{\star})\right). \end{aligned}$$

The last step used the smoothness of F_i . Combining the bounds on \mathcal{T}_2 and \mathcal{T}_3 in the original inequality and using $a = \frac{1}{K-1}$ gives,

$$\frac{1}{m}\sum_{i}\mathbb{E}\left\|\boldsymbol{\theta}_{i,k}-\boldsymbol{\theta}\right\|^{2} \leq \frac{\left(1+\frac{1}{K-1}\right)}{m}\sum_{i}\mathbb{E}\left\|\boldsymbol{\theta}_{i,k-1}-\boldsymbol{\theta}\right\|^{2}+\eta_{l}^{2}\sigma^{2} + 6\eta_{l}^{2}K\beta\left(f(\boldsymbol{\theta})-f\left(\boldsymbol{\theta}^{\star}\right)\right)+\frac{3K\eta_{l}^{2}}{m}\sum_{i}\mathbb{E}\left\|\boldsymbol{c}_{i}-\nabla F_{i}\left(\boldsymbol{\theta}^{\star}\right)\right\|^{2}+3K\eta_{l}^{2}\|\boldsymbol{m}\|^{2}$$
(73)

Recall that with the choice of c_i , the variance of c_i is less than $\frac{\sigma^2}{K}$. Separating its mean and variance gives

$$\frac{1}{m}\sum_{i}\mathbb{E}\left\|\boldsymbol{\theta}_{i,k}-\boldsymbol{\theta}\right\|^{2} \leq \left(1+\frac{1}{K-1}\right)\frac{1}{m}\sum_{i}\mathbb{E}\left\|\boldsymbol{\theta}_{i,k-1}-\boldsymbol{\theta}\right\|^{2}+7\eta_{l}^{2}\sigma^{2}+6\eta_{l}^{2}K\beta\left(f(\boldsymbol{\theta})-f\left(\boldsymbol{\theta}^{*}\right)\right)+\frac{3K\eta_{l}^{2}}{m}\sum_{i}\mathbb{E}\left\|\boldsymbol{c}_{i}-\nabla F_{i}\left(\boldsymbol{\theta}^{*}\right)\right\|^{2}+3K\eta_{l}^{2}\|\boldsymbol{m}\|^{2}$$
(74)

Unrolling the recursion, we get the following for any $k \in \{1, \ldots, K\}$

$$\frac{1}{m}\sum_{i}\mathbb{E}\left\|\boldsymbol{\theta}_{i,k}-\boldsymbol{\theta}\right\|^{2} \leq \left(6K\beta\eta_{l}^{2}\left(f(\boldsymbol{\theta})-f\left(\boldsymbol{\theta}^{\star}\right)\right)+3K\eta_{l}^{2}\mathcal{C}_{t-1}+3K\eta_{l}^{2}\mathcal{D}_{t-1}+7\beta\eta_{l}^{2}\sigma^{2}\right)\left(\sum_{k=0}^{k-1}\left(1+\frac{1}{K-1}\right)^{k}\right) \\ \leq \left(6K\beta\eta_{l}^{2}\left(f(\boldsymbol{\theta})-f\left(\boldsymbol{\theta}^{\star}\right)\right)+3K\eta_{l}^{2}\mathcal{C}_{t-1}+3K\eta_{l}^{2}\mathcal{D}_{t-1}+7\beta\eta_{l}^{2}\sigma^{2}\right)\left(K-1\right)\left(\left(1+\frac{1}{K-1}\right)^{K}-1\right) \\ \leq \left(6K\beta\eta_{l}^{2}\left(f(\boldsymbol{\theta})-f\left(\boldsymbol{\theta}^{\star}\right)\right)+3K\eta_{l}^{2}\mathcal{C}_{t-1}+3K\eta_{l}^{2}\mathcal{D}_{t-1}+7\beta\eta_{l}^{2}\sigma^{2}\right)3K \\ \leq 18K^{2}\beta\eta_{l}^{2}\left(f(\boldsymbol{\theta})-f\left(\boldsymbol{\theta}^{\star}\right)\right)+9K^{2}\eta_{l}^{2}\mathcal{C}_{t-1}+9K^{2}\eta_{l}^{2}\mathcal{D}_{t-1}+21K\beta\eta_{l}^{2}\sigma^{2}.$$
(75)

The inequality $(K-1)\left(\left(1+\frac{1}{K-1}\right)^K-1\right) \le 3K$ can be verified for K=2,3 manually. For $K \ge 4$,

$$\left(K-1\right)\left(\left(1+\frac{1}{K-1}\right)^{K}-1\right) < K\left(\exp\left(\frac{K}{K-1}\right)-1\right) \le K\left(\exp\left(\frac{4}{3}\right)-1\right) < 3K.$$
(76)

Again averaging over k and multiplying by 3β yields

$$3\beta \mathcal{E}_{t} \leq 54K^{2}\beta^{2}\eta_{l}^{2}\left(f(\boldsymbol{\theta}) - f\left(\boldsymbol{\theta}^{\star}\right)\right) + 27K^{2}\beta\eta_{l}^{2}\mathcal{C}_{t-1} + 27K^{2}\beta\eta_{l}^{2}\mathcal{D}_{t-1} + 63\beta K\eta_{l}^{2}\sigma^{2}$$

$$= \frac{1}{\alpha^{2}}\left(54\beta^{2}\tilde{\eta}^{2}\left(f(\boldsymbol{\theta}) - f\left(\boldsymbol{\theta}^{\star}\right)\right) + 27\beta\tilde{\eta}^{2}\mathcal{C}_{t-1} + 27\beta\tilde{\eta}^{2}\mathcal{C}_{t-1} + 63\beta\tilde{\eta}^{2}\frac{\sigma^{2}}{K}\right)$$

$$\leq \frac{1}{\alpha^{2}}\left(\frac{1}{25}\left(f(\boldsymbol{\theta}) - f\left(\boldsymbol{\theta}^{\star}\right)\right) + \frac{1}{3}\tilde{\eta}\mathcal{C}_{t-1} + \frac{1}{3}\tilde{\eta}\mathcal{D}_{t-1} + \tilde{\eta}\frac{\sigma^{2}}{K}\right).$$
(77)

The equality follows from the definition $\tilde{\eta} = K \eta_l \alpha$, and the final inequality uses the bound that $\tilde{\eta} \leq \frac{1}{81\beta}$.

Lemma G.13. Suppose assumptions G.1–G.2 are true. Then the following holds for any step-sizes satisfying $\alpha \ge 1, \eta_l \le \min\left(\frac{1}{81\beta K\alpha}, \frac{s}{15\mu m K\alpha}\right)$, and effective step-size $\tilde{\eta} := K\alpha\eta_l$

$$\mathbb{E}\left[\left\|\boldsymbol{\theta}^{t}-\boldsymbol{\theta}^{\star}\right\|^{2}+\frac{9m\tilde{\eta}^{2}}{s}\mathcal{C}_{t}\right] \leq \left(1-\frac{\mu\tilde{\eta}}{2}\right)\left(\mathbb{E}\left\|\boldsymbol{\theta}^{t-1}-\boldsymbol{\theta}^{\star}\right\|^{2}+\frac{9m\tilde{\eta}^{2}}{s}\mathcal{C}_{t-1}\right)$$
(78)

$$-\tilde{\eta}\left(\mathbb{E}\left[f\left(\boldsymbol{\theta}^{t-1}\right)\right] - f\left(\boldsymbol{\theta}^{\star}\right)\right) + \frac{12\tilde{\eta}^{2}}{Ks}\left(1 + \frac{s}{\alpha^{2}}\right)\sigma^{2}$$
(79)

Proof. We can then apply Lemma G.9 to bound the second moment of the server update as

$$\Delta \boldsymbol{\theta} = -\frac{\tilde{\eta}}{Ks} \sum_{k,i \in \mathcal{S}} \left(g_i \left(\boldsymbol{\theta}_{i,k-1} \right) + \boldsymbol{m} - \boldsymbol{c}_i \right), \text{ and } \mathbb{E}[\Delta \boldsymbol{\theta}] = -\frac{\tilde{\eta}}{Km} \sum_{k,i} g_i \left(\boldsymbol{\theta}_{i,k-1} \right).$$

We can bound the second moment of the server update as

$$\mathbb{E}_{t-1} \|\boldsymbol{\theta} + \Delta \boldsymbol{\theta} - \boldsymbol{\theta}^{\star}\|^{2} = \mathbb{E}_{t-1} \|\boldsymbol{\theta} - \boldsymbol{\theta}^{\star}\|^{2} - \frac{2\tilde{\eta}}{Ks} \mathbb{E}_{t-1} \sum_{k,i\in\mathcal{S}} \langle \nabla F_{i}\left(\boldsymbol{\theta}_{i,k-1}\right), \boldsymbol{\theta} - \boldsymbol{\theta}^{\star} \rangle + \mathbb{E}_{t-1} \|\Delta \boldsymbol{\theta}\|^{2}$$

$$\leq \underbrace{\frac{2\tilde{\eta}}{Ks} \mathbb{E}_{t-1} \sum_{k,i\in\mathcal{S}} \langle \nabla F_{i}\left(\boldsymbol{\theta}_{i,k-1}\right), \boldsymbol{\theta}^{\star} - \boldsymbol{\theta} \rangle}_{\mathcal{T}_{4}} + 8\beta\tilde{\eta}^{2} \left(\mathbb{E}\left[f\left(\boldsymbol{\theta}^{t-1}\right)\right] - f\left(\boldsymbol{\theta}^{\star}\right)\right) + 4\tilde{\eta}^{2}\mathcal{C}_{t-1} + 4\tilde{\eta}^{2}\mathcal{D}_{t-1} + 4\tilde{\eta}^{2}\beta^{2}\mathcal{E} + \frac{12\tilde{\eta}^{2}\sigma^{2}}{Ks}.$$
(80)

The term \mathcal{T}_4 can be bounded by using perturbed strong-convexity with $h = F_i$, $\theta = \theta_{i,k-1}$, $\theta = \theta^*$, and $z = \theta$ to get

$$\mathbb{E}\left[\mathcal{T}_{4}\right] = \frac{2\tilde{\eta}}{Ks} \mathbb{E}\sum_{k,i\in\mathcal{S}} \left\langle \nabla F_{i}\left(\boldsymbol{\theta}_{i,k-1}\right), \boldsymbol{\theta}^{\star} - \boldsymbol{\theta} \right\rangle$$

$$\leq \frac{2\tilde{\eta}}{Ks} \mathbb{E}\sum_{k,i\in\mathcal{S}} \left(F_{i}\left(\boldsymbol{\theta}^{\star}\right) - F_{i}(\boldsymbol{\theta}) + \beta \left\|\boldsymbol{\theta}_{i,k-1} - \boldsymbol{\theta}\right\|^{2} - \frac{\mu}{4} \left\|\boldsymbol{\theta} - \boldsymbol{\theta}^{\star}\right\|^{2} \right)$$

$$= -2\tilde{\eta} \mathbb{E} \left(f(\boldsymbol{\theta}) - f\left(\boldsymbol{\theta}^{\star}\right) + \frac{\mu}{4} \left\|\boldsymbol{\theta} - \boldsymbol{\theta}^{\star}\right\|^{2} \right) + 2\beta\tilde{\eta}\mathcal{E}.$$
(81)

Plugging \mathcal{T}_4 back, we can further simplify the expression to get

$$\mathbb{E} \|\boldsymbol{\theta} + \Delta \boldsymbol{\theta} - \boldsymbol{\theta}^{\star}\|^{2} \leq \mathbb{E} \|\boldsymbol{\theta} - \boldsymbol{\theta}^{\star}\|^{2} - 2\tilde{\eta} \left(f(\boldsymbol{\theta}) - f(\boldsymbol{\theta}^{\star}) + \frac{\mu}{4} \|\boldsymbol{\theta} - \boldsymbol{\theta}^{\star}\|^{2} \right) + 2\beta\tilde{\eta}\mathcal{E} + \frac{12\tilde{\eta}^{2}\sigma^{2}}{Ks} + 8\beta\tilde{\eta}^{2} \left(\mathbb{E} \left[f\left(\boldsymbol{\theta}^{t-1}\right) \right] - f\left(\boldsymbol{\theta}^{\star}\right) \right) + 4\tilde{\eta}^{2}\mathcal{C}_{t-1} + 4\tilde{\eta}^{2}\mathcal{D}_{t-1} + 4\tilde{\eta}^{2}\beta^{2}\mathcal{E} = \left(1 - \frac{\mu\tilde{\eta}}{2} \right) \|\boldsymbol{\theta} - \boldsymbol{\theta}^{\star}\|^{2} + \left(8\beta\tilde{\eta}^{2} - 2\tilde{\eta}\right) \left(f(\boldsymbol{\theta}) - f\left(\boldsymbol{\theta}^{\star}\right) \right) + \frac{12\tilde{\eta}^{2}\sigma^{2}}{Ks} + \left(2\beta\tilde{\eta} + 4\beta^{2}\tilde{\eta}^{2}\right)\mathcal{E} + 4\tilde{\eta}^{2}\mathcal{C}_{t-1} + 4\tilde{\eta}^{2}\mathcal{D}_{t-1}$$

$$(82)$$

We use Lemma G.10 can (scaled by $9\tilde{\eta}^2 \frac{m}{s}$) bound the control-lag

$$9\tilde{\eta}^{2}\frac{1}{\gamma}\mathcal{D}_{t} \leq \left(1-\frac{\mu\tilde{\eta}}{2}\right)9\tilde{\eta}^{2}\frac{1}{\gamma}\mathcal{D}_{t-1} + 9\left(\frac{\mu\tilde{\eta}}{2}\frac{1}{\gamma}-1\right)\tilde{\eta}^{2}\mathcal{D}_{t-1} + 9\tilde{\eta}^{2}\left(4\beta\left(\mathbb{E}\left[f\left(\boldsymbol{\theta}^{t-1}\right)\right]-f\left(\boldsymbol{\theta}^{\star}\right)\right)+2\beta^{2}\mathcal{E}\right)$$
(83)

Now recall that Lemma G.12 bounds the client-drift:

$$3\beta\tilde{\eta}\mathcal{E}_{t} \leq \frac{\tilde{\eta}^{2}}{3}\mathcal{C}_{t-1} + \frac{\tilde{\eta}^{2}}{3}\mathcal{D}_{t-1} + \frac{\tilde{\eta}}{25\alpha^{2}}\left(\mathbb{E}\left[f\left(\boldsymbol{\theta}^{t-1}\right)\right] - f\left(\boldsymbol{\theta}^{\star}\right)\right) + \frac{\tilde{\eta}^{2}}{K\alpha^{2}}\sigma^{2}.$$
(84)

Adding all three inequalities together,

$$\mathbb{E} \|\boldsymbol{\theta} + \Delta \boldsymbol{\theta} - \boldsymbol{\theta}^{\star}\|^{2} + \frac{9\tilde{\eta}^{2}m\mathcal{C}_{t}}{S} + 9\tilde{\eta}^{2}\frac{1}{\gamma}\mathcal{D}_{t}$$

$$\leq \left(1 - \frac{\mu\tilde{\eta}}{2}\right) \left(\mathbb{E} \|\boldsymbol{\theta} - \boldsymbol{\theta}^{\star}\|^{2} + \frac{9\tilde{\eta}^{2}m\mathcal{C}_{t-1}}{s}\right) + \left(44\beta\tilde{\eta}^{2} - \frac{49}{25}\tilde{\eta}\right) \left(f(\boldsymbol{\theta}) - f(\boldsymbol{\theta}^{\star})\right)$$

$$+ \frac{12\tilde{\eta}^{2}\sigma^{2}}{Ks} \left(1 + \frac{s}{\alpha^{2}}\right) + \left(22\beta^{2}\tilde{\eta}^{2} - \beta\tilde{\eta}\right)\mathcal{E} + \left(\frac{9\mu\tilde{\eta}m}{2S} - \frac{1}{3}\right)\tilde{\eta}^{2}\mathcal{C}_{t-1} + \left(\frac{9\mu\tilde{\eta}}{2\gamma} - \frac{1}{3}\right)\tilde{\eta}^{2}\mathcal{D}_{t-1}$$
(85)

Finally, the lemma follows from noting that $\tilde{\eta} \leq \frac{1}{81\beta}$ implies $44\beta^2 \tilde{\eta}^2 \leq \frac{24}{25}\tilde{\beta}$ and $\tilde{\eta} \leq \frac{s}{15\mu m}$ implies $\frac{9\mu\tilde{\eta}m}{2s} \leq \frac{1}{3}$. Note that if $c_i^0 = g_i(\theta^0)$, then $\frac{\tilde{\eta}m}{s}C_0$ can be bounded in terms of function sub-optimality F. The final rate for strongly convex follows simply by unrolling the recursive bound in Lemma G.13 using Lemma G.5. Also note that if $c_i^0 = g_i(x^0)$, then $\frac{\tilde{\eta}m}{s}C_0$ can be bounded in terms of function F.

G.3. Convergence Analysis for FedMoSWA under non-convex setting

We now analyze the most general case of FedMoSWA with option II on functions that are potentially non-convex. Just as in the non-convex proof, we will first bound the variance of the server update in Lemma G.14, the change in control lag in Lemma G.16, and finally we bound the client-drift in Lemma G.17. Combining these three together gives us the progress made in one round in Lemma G.18. The final rate is derived from the progress made using Lemma G.6.

Additional notation. Recall that in round r, we update the control variate as

$$\boldsymbol{c}_{i}^{t} = \begin{cases} \frac{1}{K} \sum_{k=1}^{K} g_{i} \left(\boldsymbol{\theta}_{i,k-1}^{t} \right) & \text{ if } i \in \mathcal{S}^{t}, \\ \boldsymbol{c}_{i}^{t-1} & \text{ otherwise} \end{cases}$$

We introduce the following notation to keep track of the 'lag' in the update of the control variate: define a sequence of parameters $\left\{ \alpha_{i,k-1}^{t-1} \right\}$ such that for any $i \in [m]$ and $k \in [K]$ we have $\alpha_{i,k-1}^0 := \theta^0$ and for $t \ge 1$,

$$oldsymbol{lpha}_{i,k-1}^t \coloneqq egin{cases} oldsymbol{ heta}_{i,k-1}^t & ext{if } i \in \mathcal{S}^t, \ oldsymbol{lpha}_{i,k-1}^{t-1} & ext{otherwise} \ . \end{cases}$$

By the update rule for control variates and the definition of $\{\alpha_{i,k-1}^{t-1}\}$ above, the following property always holds:

$$\boldsymbol{c}_{i}^{t} = rac{1}{K} \sum_{k=1}^{K} g_{i} \left(\boldsymbol{\alpha}_{i,k-1}^{t}
ight)$$

We can then define the following Ξ_t to be the error in control variate for round r:

$$\Xi_t := \frac{1}{Km} \sum_{k=1}^{K} \sum_{i=1}^{m} \mathbb{E} \left\| \boldsymbol{\alpha}_{i,k-1}^t - \boldsymbol{\theta}^t \right\|^2.$$

Also recall the closely related definition of client drift caused by local updates:

$$\mathcal{E}_t := \frac{1}{Km} \sum_{k=1}^K \sum_{i=1}^m \mathbb{E}\left[\left\| \boldsymbol{\theta}_{i,k}^t - \boldsymbol{\theta}^{t-1} \right\|^2 \right].$$

Variance of server update. Let us analyze how the control variates effect the variance of the aggregate server update. Define a sequence of parameters $\left\{ \alpha_{i,k-1}^{t-1} \right\}$ such that for any $i \in [m]$ and $k \in [K]$ we have $\alpha_{i,k-1}^0 := \theta^0$ and for $r \ge 1$,

$$\boldsymbol{\theta}_{i,k-1}^t \coloneqq \begin{cases} \mathbf{y}_{i,k-1}^t & p = \beta \\ \boldsymbol{\theta}_{i,k-1}^{t-1} & p = 1 - \beta \end{cases}$$

$$\boldsymbol{m}^{t} = \boldsymbol{m}^{t-1} + \beta \frac{1}{m} \sum_{i=1}^{m} \left(c_{i}^{t} - \boldsymbol{m}^{t-1} \right) = (1-\beta)\boldsymbol{m}^{t-1} + \beta \frac{1}{m} \sum_{i=1}^{m} c_{i}^{t}$$
$$\mathbb{E}[\boldsymbol{m}] = (1-\beta)\mathbb{E}\left[\boldsymbol{m}^{t-1}\right] + \beta \frac{1}{m}\mathbb{E}\left[\sum_{i=1}^{m} c_{i}^{t}\right] = \frac{1}{mK} \sum_{i=1}^{m} \sum_{k} g_{i}\left(\boldsymbol{\theta}_{i,k-1}^{t}\right)$$

We can then define the following Ξ_t to be the error in control variate for round r:

$$\mathcal{G}_{t-1} := \frac{1}{Km} \sum_{k=1}^{K} \sum_{i=1}^{m} \mathbb{E} \left\| \boldsymbol{\theta}_{i,k-1}^{t} - \boldsymbol{\theta}^{t} \right\|^{2}.$$

Variance of server update. Let us analyze how the control variates effect the variance of the aggregate server update. **Lemma G.14.** For updates and assumptions G.3 and G.2, the following holds true for any $\tilde{\eta} := \eta_l \alpha K \in [0, 1/\beta]$:

$$\mathbb{E}\left\|\mathbb{E}_{t-1}\left[\boldsymbol{\theta}^{t}\right]-\boldsymbol{\theta}^{t-1}\right\|^{2} \leq 2\tilde{\eta}^{2}\beta^{2}\mathcal{E}_{t}+2\tilde{\eta}^{2}\mathbb{E}\left\|\nabla F\left(\boldsymbol{\theta}^{t-1}\right)\right\|^{2}, \text{ and}$$
$$\mathbb{E}\left\|\boldsymbol{\theta}^{t}-\boldsymbol{\theta}^{t-1}\right\|^{2} \leq 4\tilde{\eta}^{2}\beta^{2}\mathcal{E}_{t}+4\tilde{\eta}^{2}\beta^{2}\Xi_{t-1}+4\tilde{\eta}^{2}\beta^{2}\mathcal{G}_{t-1}+4\tilde{\eta}^{2}\mathbb{E}\left\|\nabla F\left(\boldsymbol{\theta}^{t-1}\right)\right\|^{2}+\frac{9\tilde{\eta}^{2}\sigma^{2}}{Ks}.$$

Proof. Recall that that the server update satisfies

$$\mathbb{E}[\Delta \boldsymbol{\theta}] = -\frac{\tilde{\eta}}{Km} \sum_{k,i} \mathbb{E}\left[g_i\left(\boldsymbol{\theta}_{i,k-1}\right)\right]$$

From the definition of $\alpha_{i,k-1}^{t-1}$ and dropping the superscript everywhere we have

$$\Delta \boldsymbol{\theta} = -\frac{\tilde{\eta}}{Ks} \sum_{k,i \in \mathcal{S}} \left(g_i \left(\boldsymbol{\theta}_{i,k-1} \right) + \boldsymbol{m} - \boldsymbol{c}_i \right) \text{ where } \boldsymbol{c}_i = \frac{1}{K} \sum_k g_i \left(\boldsymbol{\alpha}_{i,k-1} \right).$$

Taking norm on both sides and separating mean and variance, we proceed as

$$\mathbb{E}\|\Delta\boldsymbol{\theta}\|^{2} = \mathbb{E}\left\|-\frac{\tilde{\eta}}{K_{s}}\sum_{k,i\in\mathcal{S}}\left(g_{i}\left(\boldsymbol{\theta}_{i,k-1}\right) - g_{i}\left(\boldsymbol{\alpha}_{i,k-1}\right) + \boldsymbol{m} - \boldsymbol{c}_{i}\right)\right\|^{2} \\ \leq \mathbb{E}\left\|-\frac{\tilde{\eta}}{K_{s}}\sum_{k,i\in\mathcal{S}}\left(\nabla F_{i}\left(\boldsymbol{\theta}_{i,k-1}\right) + \mathbb{E}[\boldsymbol{m}] - \mathbb{E}\left[\boldsymbol{c}_{i}\right]\right)\right\|^{2} + \frac{9\tilde{\eta}^{2}\sigma^{2}}{K_{s}} \\ \leq \mathbb{E}\left[\frac{\tilde{\eta}^{2}}{K_{s}}\sum_{k,i\in\mathcal{S}}\left\|\nabla F_{i}\left(\boldsymbol{\theta}_{i,k-1}\right) + \mathbb{E}[\boldsymbol{m}] - \mathbb{E}\left[\boldsymbol{c}_{i}\right]\right\|^{2}\right] + \frac{9\tilde{\eta}^{2}\sigma^{2}}{K_{s}} \\ = \frac{\tilde{\eta}^{2}}{K_{m}}\sum_{k,i}\mathbb{E}\left\|\left(\nabla F_{i}\left(\boldsymbol{\theta}_{i,k-1}\right) - \nabla F_{i}\left(\boldsymbol{\theta}\right)\right) + \left(\mathbb{E}[\boldsymbol{m}] - \nabla F(\boldsymbol{\theta})\right) + \nabla F(\boldsymbol{\theta}) - \left(\mathbb{E}\left[\boldsymbol{c}_{i}\right] - \nabla F_{i}(\boldsymbol{\theta})\right)\right\|^{2} + \frac{9\tilde{\eta}^{2}\sigma^{2}}{K_{s}} \\ \leq \frac{4\tilde{\eta}^{2}}{K_{m}}\sum_{k,i}\mathbb{E}\left\|\nabla F_{i}\left(\boldsymbol{\theta}_{i,k-1}\right) - \nabla F_{i}(\boldsymbol{\theta})\right\|^{2} + \frac{4\tilde{\eta}^{2}}{K_{m}}\sum_{k,i}\mathbb{E}\left\|\nabla F_{i}\left(\boldsymbol{\alpha}_{i,k-1}\right) - \nabla F_{i}(\boldsymbol{\theta})\right\|^{2} \\ + \frac{4\tilde{\eta}^{2}}{K_{m}}\sum_{k,i}\mathbb{E}\left\|\nabla F_{i}\left(\boldsymbol{\theta}_{i,k-1}\right) - \nabla F_{i}(\boldsymbol{\theta})\right\|^{2} + 4\tilde{\eta}^{2}\mathbb{E}\|\nabla F(\boldsymbol{\theta})\|^{2} + \frac{9\tilde{\eta}^{2}\sigma^{2}}{K_{s}} \\ \leq 4\tilde{\eta}^{2}\beta^{2}\mathcal{E}_{t} + 4\beta^{2}\tilde{\eta}^{2}\Xi_{t-1} + 4\beta^{2}\tilde{\eta}^{2}\mathcal{G}_{t-1} + 4\tilde{\eta}^{2}\mathbb{E}\|\nabla F(\boldsymbol{\theta})\|^{2} + \frac{9\tilde{\eta}^{2}\sigma^{2}}{K_{s}}. \end{aligned}$$

$$\tag{86}$$

In the first inequality, note that the three random variables $-\frac{1}{Ks}\sum_{k,i\in\mathcal{S}}g_i(\theta_{i,k})$, $\frac{1}{m}\sum_{i\in\mathcal{S}}c_i$, and *m*-may not be independent but each have variance smaller than $\frac{\sigma^2}{Ks}$ and so we can apply Lemma G.8. The rest of the inequalities follow from repeated applications of the relaxed triangle inequality, β -Lipschitzness of F_i , and the definition of Ξ_{t-1} . This proves the second statement. The first statement follows from our expression of $\mathbb{E}_{t-1}[\Delta \theta]$ and similar computations.

Lag in the control variates. We now analyze the 'lag' in the control variates due to us sampling only a small subset of clients each round. Because we cannot rely on convexity anymore but only on the Lipschitzness of the gradients, the control-lag increases faster in the non-convex case.

Lemma G.15. Assumptions G.3, G.2, the following holds true for any $\tilde{\eta} \leq \frac{1}{24\beta} (\gamma)^a$ for $a \in \left[\frac{1}{2}, 1\right]$ where $\tilde{\eta} := \eta_l \alpha K$:

$$\mathcal{G}_{r} \leq \left(1 - \frac{17}{36}\gamma\right)\mathcal{G}_{t-1} + \frac{1}{48\beta^{2}}\left(\gamma\right)^{2a-1} \left\|\nabla F\left(\boldsymbol{\theta}^{t-1}\right)\right\|^{2} + \frac{97}{48}\left(\gamma\right)^{2a-1}\mathcal{E}_{t} + \left(\frac{\gamma}{\beta^{2}}\right)\frac{\sigma^{2}}{32Ks}$$

Proof. The proof except that we cannot rely on convexity. Recall that after round t, the definition of $\theta_{i,k-1}^t$ implies that The proof proceeds similar to that of Lemma G.10 except that we cannot rely on convexity.

$$\mathbb{E}_{\mathcal{S}^t}\left[\boldsymbol{\theta}_{i,k-1}^t\right] = (1-\gamma)\,\boldsymbol{\theta}_{i,k-1}^{t-1} + \gamma \boldsymbol{\theta}_{i,k-1}^t.$$

Plugging the above expression in the definition of Ξ_t we get

$$\Xi_{t} = \frac{1}{Km} \sum_{i,k} \mathbb{E} \left\| \boldsymbol{\theta}_{i,k-1}^{t} - \boldsymbol{\theta}^{t} \right\|^{2}$$

$$= (1 - \gamma) \cdot \underbrace{\frac{1}{Km} \sum_{i} \mathbb{E} \left\| \boldsymbol{\theta}_{i,k-1}^{t-1} - \boldsymbol{\theta}^{t} \right\|^{2}}_{\mathcal{T}_{5}} + \gamma \cdot \underbrace{\frac{1}{Km} \sum_{k,i} \mathbb{E} \left\| \boldsymbol{\theta}_{i,k-1}^{t} - \boldsymbol{\theta}^{t} \right\|^{2}}_{\mathcal{T}_{6}}.$$
(87)

We can expand the second term \mathcal{T}_6 with the relaxed triangle inequality to claim

$$\mathcal{T}_{6} \leq 2\left(\mathcal{E}_{t} + \mathbb{E}\left\|\Delta\boldsymbol{\theta}^{t}\right\|^{2}\right).$$

We will expand the first term \mathcal{T}_5 to claim for a constant $b \ge 0$ to be chosen later

$$\mathcal{T}_{5} = \frac{1}{Km} \sum_{i} \mathbb{E} \left(\left\| \boldsymbol{\theta}_{i,k-1}^{t-1} - \boldsymbol{\theta}^{t-1} \right\|^{2} + \left\| \Delta \boldsymbol{\theta}^{t} \right\|^{2} + \mathbb{E}_{t-1} \left\langle \Delta \boldsymbol{\theta}^{t}, \boldsymbol{\theta}_{i,k-1}^{t-1} - \boldsymbol{\theta}^{t-1} \right\rangle \right) \\ \leq \frac{1}{Km} \sum_{i} \mathbb{E} \left(\left\| \boldsymbol{\theta}_{i,k-1}^{t-1} - \boldsymbol{\theta}^{t-1} \right\|^{2} + \left\| \Delta \boldsymbol{\theta}^{t} \right\|^{2} + \frac{1}{b} \left\| \mathbb{E}_{t-1} \left[\Delta \boldsymbol{\theta}^{t} \right] \right\|^{2} + b \left\| \boldsymbol{\theta}_{i,k-1}^{t-1} - \boldsymbol{\theta}^{t-1} \right\|^{2} \right)$$
(88)

where we used Young's inequality which holds for any $b \ge 0$. Combining the bounds for \mathcal{T}_5 and \mathcal{T}_6 ,

$$\mathcal{G}_{r} \leq (1-\gamma)\left(1+b\right)\mathcal{G}_{t-1} + 2\gamma\mathcal{E}_{t} + 2\mathbb{E}\left\|\Delta\boldsymbol{\theta}^{t}\right\|^{2} + \frac{1}{b}\mathbb{E}\left\|\mathbb{E}_{t-1}\left[\Delta\boldsymbol{\theta}^{t}\right]\right\|^{2}$$

$$\leq \left((1-\gamma)\left(1+b\right) + 16\tilde{\eta}^{2}\beta^{2}\right)\Xi_{t-1} + \left(2\gamma m + 8\tilde{\eta}^{2}\beta^{2} + 2\frac{1}{b}\tilde{\eta}^{2}\beta^{2}\right)\mathcal{E}_{t} + \left(8 + 2\frac{1}{b}\right)\tilde{\eta}^{2}\mathbb{E}\|\nabla F(\boldsymbol{\theta})\|^{2}\right) + \frac{18\tilde{\eta}^{2}\sigma^{2}}{Ks}$$

$$\tag{89}$$

Verify that with choice of $b = \frac{\gamma}{2(1-\gamma)}$, we have $(1-\gamma)(1+b) \le (1-\frac{\gamma}{2})$ and $\frac{1}{b} \le \frac{2}{\gamma}$. Plugging these values along with the bound on the step-size $16\beta^2 \tilde{\eta}^2 \le \frac{1}{36} (\gamma)^{2a} \le \frac{\gamma}{36}$ completes the lemma. Bounding the drift. We will next bound the client drift \mathcal{E}_t .

Lag in the control variates. We now analyze the 'lag' in the control variates due to us sampling only a small subset of clients each round. Because we cannot rely on convexity anymore but only on the Lipschitzness of the gradients, the control-lag increases faster in the non-convex case.

Lemma G.16. Assumptions G.3–G.2, the following holds true for any $\tilde{\eta} \leq \frac{1}{24\beta} \left(\frac{s}{m}\right)^a$ for $a \in \left[\frac{1}{2}, 1\right]$ where $\tilde{\eta} := \eta_l \alpha K$:

$$\Xi_{t} \leq \left(1 - \frac{17s}{36m}\right) \Xi_{t-1} + \frac{1}{48\beta^{2}} \left(\frac{s}{m}\right)^{2a-1} \left\|\nabla F\left(\boldsymbol{\theta}^{t-1}\right)\right\|^{2} + \frac{97}{48} \left(\frac{s}{m}\right)^{2a-1} \mathcal{E}_{t} + \left(\frac{s}{m\beta^{2}}\right) \frac{\sigma^{2}}{32Ks}.$$

Proof. The proof except that we cannot rely on convexity. Recall that after round t, the definition of $\alpha_{i,k-1}^t$ implies that

$$\mathbb{E}_{\mathcal{S}^{t}}\left[\boldsymbol{\alpha}_{i,k-1}^{t}\right] = \left(1 - \frac{s}{m}\right)\boldsymbol{\alpha}_{i,k-1}^{t-1} + \frac{s}{m}\boldsymbol{\theta}_{i,k-1}^{t}.$$

Plugging the above expression in the definition of Ξ_t we get

$$\Xi_{t} = \frac{1}{Km} \sum_{i,k} \mathbb{E} \left\| \boldsymbol{\alpha}_{i,k-1}^{t} - \boldsymbol{\theta}^{t} \right\|^{2}$$
$$= \left(1 - \frac{s}{m}\right) \cdot \underbrace{\frac{1}{Km} \sum_{i} \mathbb{E} \left\| \boldsymbol{\alpha}_{i,k-1}^{t-1} - \boldsymbol{\theta}^{t} \right\|^{2}}_{\mathcal{T}_{5}} + \frac{s}{m} \cdot \underbrace{\frac{1}{Km} \sum_{k,i} \mathbb{E} \left\| \boldsymbol{\theta}_{i,k-1}^{t} - \boldsymbol{\theta}^{t} \right\|^{2}}_{\mathcal{T}_{6}}.$$
(90)

We can expand the second term \mathcal{T}_6 with the relaxed triangle inequality to claim

$$\mathcal{T}_{6} \leq 2\left(\mathcal{E}_{t} + \mathbb{E}\left\|\Delta\boldsymbol{\theta}^{t}\right\|^{2}\right).$$
(91)

We will expand the first term \mathcal{T}_5 to claim for a constant $b \ge 0$ to be chosen later

$$\mathcal{T}_{5} = \frac{1}{Km} \sum_{i} \mathbb{E} \left(\left\| \boldsymbol{\alpha}_{i,k-1}^{t-1} - \boldsymbol{\theta}^{t-1} \right\|^{2} + \left\| \Delta \boldsymbol{\theta}^{t} \right\|^{2} + \mathbb{E}_{t-1} \left\langle \Delta \boldsymbol{\theta}^{t}, \boldsymbol{\alpha}_{i,k-1}^{t-1} - \boldsymbol{\theta}^{t-1} \right\rangle \right) \\ \leq \frac{1}{Km} \sum_{i} \mathbb{E} \left(\left\| \boldsymbol{\alpha}_{i,k-1}^{t-1} - \boldsymbol{\theta}^{t-1} \right\|^{2} + \left\| \Delta \boldsymbol{\theta}^{t} \right\|^{2} + \frac{1}{b} \left\| \mathbb{E}_{t-1} \left[\Delta \boldsymbol{\theta}^{t} \right] \right\|^{2} + b \left\| \boldsymbol{\alpha}_{i,k-1}^{t-1} - \boldsymbol{\theta}^{t-1} \right\|^{2} \right)$$
(92)

where we used Young's inequality which holds for any $b \ge 0$. Combining the bounds for \mathcal{T}_5 and \mathcal{T}_6 ,

$$\begin{aligned} \Xi_t &\leq \left(1 - \frac{s}{m}\right)(1+b)\Xi_{t-1} + 2\frac{s}{m}\mathcal{E}_t + 2\mathbb{E}\left\|\Delta\boldsymbol{\theta}^t\right\|^2 + \frac{1}{b}\mathbb{E}\left\|\mathbb{E}_{t-1}\left[\Delta\boldsymbol{\theta}^t\right]\right\|^2 \\ &\leq \left(\left(1 - \frac{s}{m}\right)(1+b) + 16\tilde{\eta}^2\beta^2\right)\Xi_{t-1} + \left(\frac{2s}{m} + 8\tilde{\eta}^2\beta^2 + 2\frac{1}{b}\tilde{\eta}^2\beta^2\right)\mathcal{E}_t + \left(8 + 2\frac{1}{b}\right)\tilde{\eta}^2\mathbb{E}\|\nabla F(\boldsymbol{\theta})\|^2\right) + \frac{18\tilde{\eta}^2\sigma^2}{Ks} \end{aligned}$$

$$\tag{93}$$

Verify that with choice of $b = \frac{s}{2(m-s)}$, we have $\left(1 - \frac{s}{m}\right)(1+b) \le \left(1 - \frac{s}{2m}\right)$ and $\frac{1}{b} \le \frac{2m}{s}$. Plugging these values along with the bound on the step-size $16\beta^2 \tilde{\eta}^2 \le \frac{1}{36} \left(\frac{s}{m}\right)^{2a} \le \frac{s}{36m}$ completes the lemma. Bounding the drift. We will next bound the client drift \mathcal{E}_t .

Bounding the drift. We will next bound the client drift \mathcal{E}_t . For this, convexity is not crucial and we will recover a very similar result to Lemma G.12 only use the Lipschitzness of the gradient.

Lemma G.17. Suppose our step-sizes satisfy $\eta_l \leq \frac{1}{24\beta K\alpha}$ and F_i satisfies assumptions G.3–G.2. Then, for any global $\alpha \geq 1$ we can bound the drift as

$$\frac{5}{3}\beta^2 \tilde{\eta} \mathcal{E}_t \le \frac{5}{6}\beta^3 \tilde{\eta}^2 \Xi_{t-1} + \frac{5}{6}\beta^3 \tilde{\eta}^2 \mathcal{G}_{t-1} + \frac{\tilde{\eta}}{24\alpha^2} \mathbb{E} \left\| \nabla F\left(\boldsymbol{\theta}^{t-1}\right) \right\|^2 + \frac{\tilde{\eta}^2 \beta}{4K\alpha^2} \sigma^2$$

Proof. First, observe that if K = 1, $\mathcal{E}_t = 0$ since $\theta_{i,0} = \theta$ for all $i \in [m]$ and that Ξ_{t-1} and the right hand side are both positive. Thus the Lemma is trivially true if K = 1 and we will henceforth assume $K \ge 2$. Starting from the update for $i \in [m]$ and $k \in [K]$.

$$\mathbb{E} \|\boldsymbol{\theta}_{i,k} - \boldsymbol{\theta}\|^{2} = \mathbb{E} \|\boldsymbol{\theta}_{i,k-1} - \eta_{l} \left(g_{i} \left(\boldsymbol{\theta}_{i,k-1}\right) + \boldsymbol{m} - \boldsymbol{c}_{i}\right) - \boldsymbol{\theta}\|^{2} \\
\leq \mathbb{E} \|\boldsymbol{\theta}_{i,k-1} - \eta_{l} \left(\nabla F_{i} \left(\boldsymbol{\theta}_{i,k-1}\right) + \boldsymbol{m} - \boldsymbol{c}_{i}\right) - \boldsymbol{\theta}\|^{2} + \eta_{l}^{2} \sigma^{2} \\
\leq \left(1 + \frac{1}{K-1}\right) \mathbb{E} \|\boldsymbol{\theta}_{i,k-1} - \boldsymbol{\theta}\|^{2} + K \eta_{l}^{2} \mathbb{E} \|\nabla F_{i} \left(\boldsymbol{\theta}_{i,k-1}\right) + \boldsymbol{m} - \boldsymbol{c}_{i}\|^{2} + \eta_{l}^{2} \sigma^{2} \\
= \left(1 + \frac{1}{K-1}\right) \mathbb{E} \|\boldsymbol{\theta}_{i,k-1} - \boldsymbol{\theta}\|^{2} + \eta_{l}^{2} \sigma^{2} \\
+ K \eta_{l}^{2} \mathbb{E} \|\nabla F_{i} \left(\boldsymbol{\theta}_{i,k-1}\right) - \nabla F_{i} \left(\boldsymbol{\theta}\right) + \left(\boldsymbol{m} - \nabla F(\boldsymbol{\theta})\right) + \nabla F(\boldsymbol{\theta}) - \left(\boldsymbol{c}_{i} - \nabla F_{i}(\boldsymbol{\theta})\right)\|^{2} \\
\leq \left(1 + \frac{1}{K-1}\right) \mathbb{E} \|\boldsymbol{\theta}_{i,k-1} - \boldsymbol{\theta}\|^{2} + 4K \eta_{l}^{2} \mathbb{E} \|\nabla F_{i} \left(\boldsymbol{\theta}_{i,k-1}\right) - \nabla F_{i} \left(\boldsymbol{\theta}\right)\|^{2} + \eta_{l}^{2} \sigma^{2} \\
+ 4K \eta_{l}^{2} \mathbb{E} \|\boldsymbol{m} - \nabla F(\boldsymbol{\theta})\|^{2} + 4K \eta_{l}^{2} \mathbb{E} \|\nabla F(\boldsymbol{\theta})\|^{2} + 4K \eta_{l}^{2} \mathbb{E} \|\boldsymbol{c}_{i} - \nabla F_{i} \left(\boldsymbol{\theta}\right)\|^{2} \\
\leq \left(1 + \frac{1}{K-1} + 4K \beta^{2} \eta_{l}^{2}\right) \mathbb{E} \|\boldsymbol{\theta}_{i,k-1} - \boldsymbol{\theta}\|^{2} + \eta_{l}^{2} \sigma^{2} + 4K \eta_{l}^{2} \mathbb{E} \|\nabla F(\boldsymbol{\theta})\|^{2} \\
+ 4K \eta_{l}^{2} \mathbb{E} \|\boldsymbol{m} - \nabla F(\boldsymbol{\theta})\|^{2} + 4K \eta_{l}^{2} \mathbb{E} \|\boldsymbol{c}_{i} - \nabla F_{i} \left(\boldsymbol{\theta}\right)\|^{2}
\end{aligned}$$

The inequalities above follow from repeated application of the relaxed triangle inequalities and the β -Lipschitzness of F_i . Averaging the above over *i*, the definition of $\boldsymbol{m} = \frac{1}{m} \sum_i \boldsymbol{c}_i$ and Ξ_{t-1} gives

$$\frac{1}{m}\sum_{i}\mathbb{E}\|\boldsymbol{\theta}_{i,k}-\boldsymbol{\theta}\|^{2} \leq \left(1+\frac{1}{K-1}+4K\beta^{2}\eta_{l}^{2}\right)\frac{1}{m}\sum_{i}\mathbb{E}\|\boldsymbol{\theta}_{i,k-1}-\boldsymbol{\theta}\|^{2} \\
+\eta_{l}^{2}\sigma^{2}+4K\eta_{l}^{2}\mathbb{E}\|\nabla F(\boldsymbol{\theta})\|^{2}+4K\eta_{l}^{2}\beta^{2}\Xi_{t-1}+4K\eta_{l}^{2}\beta^{2}\mathcal{G}_{t-1} \\
\leq \left(\eta_{l}^{2}\sigma^{2}+4K\eta_{l}^{2}\mathbb{E}\|\nabla F(\boldsymbol{\theta})\|^{2}+4K\eta_{l}^{2}\beta^{2}\Xi_{t-1}+4K\eta_{l}^{2}\beta^{2}\mathcal{G}_{t-1}\right)\left(\sum_{k=0}^{k-1}\left(1+\frac{1}{K-1}+4K\beta^{2}\eta_{l}^{2}\right)^{k}\right) \\
= \left(\frac{\tilde{\eta}^{2}\sigma^{2}}{K^{2}\alpha^{2}}+\frac{4\tilde{\eta}^{2}}{K\alpha^{2}}\mathbb{E}\|\nabla F(\boldsymbol{\theta})\|^{2}+\frac{4\tilde{\eta}^{2}\beta^{2}}{K\alpha^{2}}\Xi_{t-1}+\frac{4\tilde{\eta}^{2}\beta^{2}}{K\alpha^{2}}\mathcal{G}_{t-1}\right)\left(\sum_{k=0}^{k-1}\left(1+\frac{1}{K-1}+\frac{4\beta^{2}\tilde{\eta}^{2}}{K\alpha^{2}}\right)^{k}\right) \\
\leq \left(\frac{\tilde{\eta}\sigma^{2}}{24\beta K^{2}\alpha^{2}}+\frac{1}{144\beta^{2}K\alpha^{2}}\mathbb{E}\|\nabla F(\boldsymbol{\theta})\|^{2}+\frac{\tilde{\eta}\beta}{6K\alpha^{2}}\Xi_{t-1}+\frac{\tilde{\eta}\beta}{6K\alpha^{2}}\mathcal{G}_{t-1}\right)3K.$$
(95)

The last inequality used the bound on the step-size $\beta \tilde{\eta} \leq \frac{1}{24}$. Averaging over k and multiplying both sides by $\frac{5}{3}\beta^2 \tilde{\eta}$ yields the lemma statement. Progress made in each round. Given that we can bound all sources of error, we can finally prove the progress made in each round.

Progress made in each round. Given that we can bound all sources of error, we can finally prove the progress made in each round.

Lemma G.18. Suppose the updates satisfy assumptions G.3–G.2. For any effective step-size $\tilde{\eta} := K \alpha \eta_l$ satisfying $\tilde{\eta} \leq \frac{1}{24\beta} \left(\frac{s}{m}\right)^{\frac{2}{3}}$,

$$\left(\mathbb{E}\left[f\left(\boldsymbol{\theta}^{t}\right)\right] + 12\beta^{3}\tilde{\eta}^{2}\frac{m}{s}\Xi_{t}\right) \leq \left(\mathbb{E}\left[f\left(\boldsymbol{\theta}^{t-1}\right)\right] + 12\beta^{3}\tilde{\eta}^{2}\frac{m}{s}\Xi_{t-1}\right) + \frac{5\beta\tilde{\eta}^{2}\sigma^{2}}{Ks}\left(1 + \frac{s}{\alpha^{2}}\right) - \frac{\tilde{\eta}}{14}\mathbb{E}\left\|\nabla F\left(\boldsymbol{\theta}^{t-1}\right)\right\|^{2}.$$

Proof. Starting from the smoothness of f and taking conditional expectation gives

$$\mathbb{E}_{t-1}[f(\boldsymbol{\theta} + \Delta \boldsymbol{\theta})] \le f(\boldsymbol{\theta}) + \langle \nabla F(\boldsymbol{\theta}), \mathbb{E}_{t-1}[\Delta \boldsymbol{\theta}] \rangle + \frac{\beta}{2} \mathbb{E}_{t-1} \|\Delta \boldsymbol{\theta}\|^2.$$
(96)

We as usual dropped the superscript everywhere. Recall that the server update can be written as

$$\Delta \boldsymbol{\theta} = -\frac{\tilde{\eta}}{Ks} \sum_{k,i \in \mathcal{S}} \left(g_i \left(\boldsymbol{\theta}_{i,k-1} \right) + \boldsymbol{m} - \boldsymbol{c}_i \right), \text{ and } \mathbb{E}_{\mathcal{S}}[\Delta \boldsymbol{\theta}] = -\frac{\tilde{\eta}}{Km} \sum_{k,i} g_i \left(\boldsymbol{\theta}_{i,k-1} \right).$$
(97)

Substituting this in the previous inequality to bound $\mathbb{E}\left[\|\Delta \boldsymbol{\theta}\|^2\right]$ gives

$$\mathbb{E}[f(\boldsymbol{\theta} + \Delta\boldsymbol{\theta})] - f(\boldsymbol{\theta})$$

$$\leq -\frac{\tilde{\eta}}{Km} \sum_{k,i} \langle \nabla F(\boldsymbol{\theta}), \mathbb{E} [\nabla F_i(\boldsymbol{\theta}_{i,k-1})] \rangle + \frac{\beta}{2} \mathbb{E} \|\Delta\boldsymbol{\theta}\|^2$$

$$\leq -\frac{\tilde{\eta}}{Km} \sum_{k,i} \langle \nabla F(\boldsymbol{\theta}), \mathbb{E} [\nabla F_i(\boldsymbol{\theta}_{i,k-1})] \rangle +$$

$$2\tilde{\eta}^2 \beta^3 \mathcal{E}_t + 2\tilde{\eta}^2 \beta^3 \Xi_{t-1} + 2\tilde{\eta}^2 \beta^3 \mathcal{G}_{t-1} + 2\beta \tilde{\eta}^2 \mathbb{E} \|\nabla F(\boldsymbol{\theta})\|^2 + \frac{9\beta \tilde{\eta}^2 \sigma^2}{2Ks}$$

$$\leq -\frac{\tilde{\eta}}{2} \|\nabla F(\boldsymbol{\theta})\|^2 + \frac{\tilde{\eta}}{2} \sum_{i,k} \mathbb{E} \left\| \frac{1}{Km} \sum_{i,k} \nabla F_i(\boldsymbol{\theta}_{i,k-1}) - \nabla F(\boldsymbol{\theta}) \right\|^2 +$$

$$2\tilde{\eta}^2 \beta^3 \mathcal{E}_t + 2\tilde{\eta}^2 \beta^3 \Xi_{t-1} + 2\tilde{\eta}^2 \beta^3 \mathcal{G}_{t-1} + 2\beta \tilde{\eta}^2 \mathbb{E} \|\nabla F(\boldsymbol{\theta})\|^2 + \frac{9\beta \tilde{\eta}^2 \sigma^2}{2Ks}$$

$$\leq -\frac{\tilde{\eta}}{2} \|\nabla F(\boldsymbol{\theta})\|^2 + \frac{\tilde{\eta}}{2Km} \sum_{i,k} \mathbb{E} \|\nabla F_i(\boldsymbol{\theta}_{i,k-1}) - \nabla F_i(\boldsymbol{\theta})\|^2 +$$

$$2\tilde{\eta}^2 \beta^3 \mathcal{E}_t + 2\tilde{\eta}^2 \beta^3 \Xi_{t-1} + 2\tilde{\eta}^2 \beta^3 \mathcal{G}_{t-1} + 2\beta \tilde{\eta}^2 \mathbb{E} \|\nabla F(\boldsymbol{\theta})\|^2 + \frac{9\beta \tilde{\eta}^2 \sigma^2}{2Ks}$$

$$\leq -\frac{\tilde{\eta}}{2} \|\nabla F(\boldsymbol{\theta})\|^2 + \frac{\tilde{\eta}}{2Km} \sum_{i,k} \mathbb{E} \|\nabla F_i(\boldsymbol{\theta}_{i,k-1}) - \nabla F_i(\boldsymbol{\theta})\|^2 + \frac{9\beta \tilde{\eta}^2 \sigma^2}{2Ks}$$

$$\leq -\left(\frac{\tilde{\eta}}{2} - 2\beta \tilde{\eta}^2\right) \|\nabla F(\boldsymbol{\theta})\|^2 + \left(\frac{\tilde{\eta}}{2} + 2\beta \tilde{\eta}^2\right) \beta^2 \mathcal{E}_t + 2\beta^3 \tilde{\eta}^2 \mathcal{G}_{t-1} + 2\beta^3 \tilde{\eta}^2 \Xi_{t-1} + \frac{9\beta \tilde{\eta}^2 \sigma^2}{2Ks}.$$

The third inequality follows from the observation that $-ab = \frac{1}{2} ((b-a)^2 - a^2) - \frac{1}{2}b^2 \le \frac{1}{2} ((b-a)^2 - a^2)$ for any $a, b \in \mathbb{R}$, and the last from the β -Lipschitzness of F_i . Now we use Lemma G.16 to Ξ_t as

$$12\beta^{3}\tilde{\eta}^{2}\frac{m}{s}\Xi_{t} \leq 12\beta^{3}\tilde{\eta}^{2}\frac{m}{s}\left(\left(1-\frac{17s}{36m}\right)\Xi_{t-1}+\frac{1}{48\beta^{2}}\left(\frac{s}{m}\right)^{2a-1}\left\|\nabla F\left(\boldsymbol{\theta}^{t-1}\right)\right\|^{2}+\frac{97}{48}\left(\frac{s}{m}\right)^{2a-1}\mathcal{E}_{t}+\left(\frac{s}{m\beta^{2}}\right)\frac{\sigma^{2}}{32Ks}\right)\right.$$
$$=12\beta^{3}\tilde{\eta}^{2}\frac{m}{s}\Xi_{t-1}-\frac{17}{3}\beta^{3}\tilde{\eta}^{2}\Xi_{t-1}+\frac{1}{4}\beta\tilde{\eta}^{2}\left(\frac{m}{s}\right)^{2-2a}\left\|\nabla F(\boldsymbol{\theta})\right\|^{2}+\frac{97}{4}\beta^{3}\tilde{\eta}^{2}\left(\frac{m}{s}\right)^{2-2a}\mathcal{E}_{t}+\frac{3\beta\tilde{\eta}^{2}\sigma^{2}}{8Ks}.$$
$$\tag{99}$$

Now we use Lemma G.17 to bound \mathcal{G}_{t-1} as

$$12\beta^{3}\tilde{\eta}^{2}\frac{1}{\gamma}\mathcal{G}_{r} \leq 12\beta^{3}\tilde{\eta}^{2}\frac{1}{\gamma}\left(\left(1-\frac{17\gamma}{36}\right)\mathcal{G}_{t-1}+\frac{1}{48\beta^{2}}\left(\gamma\right)^{2a-1}\left\|\nabla F\left(\boldsymbol{\theta}^{t-1}\right)\right\|^{2}+\frac{97}{48}\left(\gamma\right)^{2a-1}\mathcal{G}_{t-1}+\left(\frac{\gamma}{\beta^{2}}\right)\frac{\sigma^{2}}{32Ks}\right)\right.$$
$$=12\beta^{3}\tilde{\eta}^{2}\frac{1}{\gamma}\mathcal{G}_{t-1}-\frac{17}{3}\beta^{3}\tilde{\eta}^{2}\mathcal{G}_{t-1}+\frac{1}{4}\beta\tilde{\eta}^{2}\left(\frac{1}{\gamma}\right)^{2-2\alpha}\left\|\nabla F(\boldsymbol{\theta})\right\|^{2}+\frac{97}{4}\beta^{3}\tilde{\eta}^{2}\left(\frac{1}{\gamma}\right)^{2-2a}\mathcal{E}_{t}+\frac{3\beta\tilde{\eta}^{2}\sigma^{2}}{8Ks}.$$
$$(100)$$

Also recall that Lemma G.18 states that

$$\frac{5}{3}\beta^2 \tilde{\eta} \mathcal{E}_t \le \frac{5}{6}\beta^3 \tilde{\eta}^2 \Xi_{t-1} + \frac{5}{6}\beta^3 \tilde{\eta}^2 \mathcal{G}_{t-1} + \frac{\tilde{\eta}}{24a^2} \mathbb{E} \left\| \nabla F\left(\boldsymbol{\theta}^{t-1}\right) \right\|^2 + \frac{\tilde{\eta}^2 \beta}{4K\alpha^2} \sigma^2.$$
(101)

Adding these bounds on Ξ_t and \mathcal{E}_t to that of $\mathbb{E}[f(\theta + \Delta \theta)]$ gives

$$\left(\mathbb{E}[f(\boldsymbol{\theta}+\Delta\boldsymbol{\theta})]+12\beta^{3}\tilde{\eta}^{2}\frac{m}{S}\Xi_{t}+12\beta^{3}\tilde{\eta}^{2}\frac{1}{\gamma}\mathcal{G}_{r}\right) \leq \left(\mathbb{E}[f(\boldsymbol{\theta})]+12\beta^{3}\tilde{\eta}^{2}\frac{m}{s}\Xi_{t-1}\right) + \left(\frac{5}{6}-\frac{17}{3}\right)\beta^{3}\tilde{\eta}^{2}\Xi_{t-1} \\
+ \left(\frac{5}{6}-\frac{17}{3}\right)\beta^{3}\tilde{\eta}^{2}\mathcal{G}_{t-1} - \left(\frac{\tilde{\eta}}{2}-2\beta\tilde{\eta}^{2}-\frac{1}{4}\beta\tilde{\eta}^{2}\left(\frac{m}{s}\right)^{2-2a}\right)\|\nabla F(\boldsymbol{\theta})\|^{2} \\
+ \left(\frac{\tilde{\eta}}{2}-\frac{5\tilde{\eta}}{3}+2\beta\tilde{\eta}^{2}+\frac{97}{4}\beta\tilde{\eta}^{2}\left(\frac{m}{s}\right)^{2-2a}\right)\beta^{2}\mathcal{E}_{t} + \frac{39\beta\tilde{\eta}^{2}\sigma^{2}}{8Ks}\left(1+\frac{s}{\alpha^{2}}\right)$$
(102)

By our choice of $a = \frac{2}{3}$ and plugging in the bound on step-size $\beta \tilde{\eta} \left(\frac{m}{s}\right)^{2-2a} \leq \frac{1}{24}$ proves the lemma. The non-convex rate of convergence now follows by unrolling the recursion and selecting an appropriate step-size $\tilde{\eta}$. Finally note that if we initialize $c_i^0 = g_i \left(\boldsymbol{\theta}^0\right)$ then we have $\Xi_0 = 0$.

The **non-convex rate** of convergence now follows by unrolling the recursion in Lemma G.18 and selecting an appropriate step-size $\tilde{\eta}$ as in Lemma G.6. Finally note that if $c_i^0 = g_i(\theta^0)$ then we have $\Xi_0 = 0$.

H. Effectiveness of local learning rate decay

Theorem H.1 (Effectiveness of local learning rate decay). We define client-drift to be how much the clients move from their starting point. When using a constant learning rate $\eta_k^{(t)} = \eta_l$, we define the client-side drift to be $\mathcal{E}_t := \frac{1}{km} \sum_{k=1}^k \sum_{i=1}^m \mathbb{E}\left[\left\| \mathbf{v}_{i,k}^t - \boldsymbol{\theta}^{t-1} \right\|^2 \right]$, and when using local learning rate decay $\eta_k^{(t)} = \left(\frac{K-2}{K-1} \right)^k$, $\tilde{\mathcal{E}}_t := \frac{1}{km} \sum_{k=1}^k \sum_{i=1}^m \mathbb{E}\left[\left\| \tilde{\mathbf{v}}_{i,k}^t - \tilde{\boldsymbol{\theta}}^{t-1} \right\|^2 \right]$. $\mathcal{E}_t \leq \frac{\tilde{\eta}}{9\beta} \mathcal{C}_{t-1} + \frac{\tilde{\eta}}{9\beta} \mathcal{D}_{t-1} + \frac{1}{75\alpha^2\beta} \left(\mathbb{E}\left[f\left(\mathbf{x}^{t-1} \right) \right] - f\left(\mathbf{x}^* \right) \right) + \frac{\tilde{\eta}}{3K\alpha^2\beta} \sigma^2$ $\tilde{\mathcal{E}}_t \leq \frac{1}{3} \left[\frac{\tilde{\eta}}{9\beta} \mathcal{C}_{t-1} + \frac{\tilde{\eta}}{9\beta} \mathcal{D}_{t-1} + \frac{1}{75\alpha^2\beta} \left(\mathbb{E}\left[f\left(\mathbf{x}^{t-1} \right) \right] - f\left(\mathbf{x}^* \right) \right) + \frac{\tilde{\eta}}{3K\alpha^2\beta} \sigma^2 \right]$.

So local learning rate decay client-side drift $\tilde{\mathcal{E}}_t$ is better than \mathcal{E}_t

I. Generalization Analysis for FedSAM and MoFedSAM under non-convex setting

Lemma I.1. Suppose Assumptions 2-5 hold. Then for FedSAM with $\eta^t \leq /\beta KT$,

$$\mathbb{E} \left\| \theta_{i,k}^t - \theta_t \right\| \le (1 + \overline{\mathbf{c}})^{K-1} \tilde{\eta}_t \left(\mathbb{E} \left\| \nabla F\left(\theta_t\right) \right\| + \sigma_{g,i} + \sigma \right), \quad \forall k = 1, \dots, K,$$

where $\tilde{\eta}_t = \sum_{k=0}^{K-1} \eta^t, \, \overline{c} \le 1 + (1 + c)^{K-1} \beta \tilde{\eta}_t \le 1 + (1 + c)^{K-1} \beta \sum_{k=0}^{K-1} \eta^t$

Proof. Considering local update of FedSAM

$$\mathbb{E} \|\theta_{i,k+1} - \theta_t\| = \mathbb{E} \|\theta_{i,k}^t - \eta^t g_i(\theta_{i,k}^t) - \theta_t\|
\leq \mathbb{E} \|\theta_{i,k}^t - \theta_t - \eta^t (g_i(\theta_{i,k}^t) - g_i(\theta_t))\| + \eta^t \mathbb{E} \|g_i(\theta_t)\|
\leq (1 + \beta \eta^t) \mathbb{E} \|\theta_{i,k}^t - \theta_t\| + \eta^t \mathbb{E} \|g_i(\theta_t)\|
\leq (1 + \beta \eta^t) \mathbb{E} \|\theta_{i,k}^t - \theta_t\| + \eta^t (\mathbb{E} \|g_i(\theta_t) - \nabla F_i(\theta_t)\| + \mathbb{E} \|\nabla F_i(\theta_t)\|)
\leq (1 + \beta \eta^t) \mathbb{E} \|\theta_{i,k}^t - \theta_t\| + \eta^t (\mathbb{E} \|\nabla F_i(\theta_t)\| + \sigma),$$
(103)

where we use Assumptions C.2 and C.3. Unrolling the above and noting $\theta_{i,0} = \theta_t$ yields

$$\mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{t} \right\| \leq \sum_{l=0}^{k-1} \eta^{t} \left(\mathbb{E} \left\| \nabla F_{i} \left(\theta_{t} \right) \right\| + \sigma \right) (1+c)^{k-1-l}$$

$$\leq \sum_{l=0}^{K-1} \eta^{t} \left(\mathbb{E} \left\| \nabla F_{i} \left(\theta_{t} \right) \right\| + \sigma \right) (1+c)^{K-1}$$

$$\leq (1+c)^{K-1} \tilde{\eta}_{t} \left(\mathbb{E} \left\| \nabla F \left(\theta_{t} \right) \right\| + \sigma_{g,i} + \sigma \right),$$
(104)

where the last inequality follows Assumption C.5.

Lemma I.2. Given Assumptions 1 - 3, for $\eta^t \leq /\beta KT$, we have

$$\mathbb{E}\left\|g_{i}\left(\theta_{i,k}^{t}\right)\right\| \leq \left(1 + (1+c)^{K-1}\beta\tilde{\eta}_{t}\right)\left(\mathbb{E}\left\|\nabla F\left(\theta_{t}\right)\right\| + \sigma_{g,i} + \sigma\right)$$

where $g_i(\cdot)$ is the sampled gradient of client $i, \tilde{\eta}_t = \sum_{k=0}^{K-1} \eta^t$.

Proof. we obtain

$$\mathbb{E} \left\| g_{i}\left(\theta_{i,k}^{t}\right) \right\| \leq \mathbb{E} \left\| g_{i}\left(\theta_{i,k}^{t}\right) - \nabla F_{i}\left(\theta_{i,k}^{t}\right) \right\| + \mathbb{E} \left\| \nabla F_{i}\left(\theta_{i,k}^{t}\right) \right\| \\
\leq \mathbb{E} \left\| \nabla F_{i}\left(\theta_{i,k}^{t}\right) \right\| + \sigma \\
\leq \mathbb{E} \left\| \nabla F_{i}\left(\theta_{t}\right) \right\| + \mathbb{E} \left\| \nabla F_{i}\left(\theta_{i,k}^{t}\right) - \nabla F_{i}\left(\theta_{t}\right) \right\| + \sigma \\
\leq \mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \mathbb{E} \left\| \nabla F_{i}\left(\theta_{t}\right) - \nabla F\left(\theta_{t}\right) \right\| + \beta \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{t} \right\| + \sigma \\
\leq \left(1 + (1 + c)^{K-1} \beta \tilde{\alpha}_{i}\right) \left(\mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \sigma_{g,i} + \sigma\right).$$
(105)

Theorem I.3. (FedSAM). Suppose Assumptions 1-3 hold and consider FedSAM. Let $k = K, \forall i \in [m]$ and $\eta^t \leq \frac{1}{\beta KT}$. Then,

$$\varepsilon_{gen} \leq \frac{2Lc}{mn} \frac{1}{\beta} e^{1 + \frac{1}{T}} \left(\mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \sigma_{g} + \sigma \right),$$

Proof. For client *i*, there are two cases to consider. In the first case, SGD selects non-perturbed samples in S and $S^{(i)}$, which happens with probability 1 - 1/n. Then, we have

$$\left\|\theta_{i,k+1}^{t} - \theta_{i,k+1}^{t,\prime}\right\| \le \left(1 + \beta \eta_{k}^{t}\right) \left\|\theta_{i,k}^{t} - \theta_{i,k}^{t,\prime}\right\|.$$

$$(106)$$

In the second case, SGD encounters the perturbed sample at time step k, which happens with probability 1/n. Then, we have

$$\begin{split} \left\| \theta_{i,k+1}^{t} - \theta_{i,k+1}^{t,\prime} \right\| &= \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} - \eta^{t} \left(g_{i} \left(\theta_{i,k}^{t} \right) - g_{i}^{\prime} \left(\theta_{i,k}^{t,\prime} \right) \right) \right\| \\ &\leq \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} - \eta^{t} \left(g_{i} \left(\theta_{i,k}^{t} \right) - g_{i} \left(\theta_{i,k}^{t,\prime} \right) \right) \right\| + \eta^{t} \left\| g_{i} \left(\theta_{i,k}^{t,\prime} \right) - g_{i}^{\prime} \left(\theta_{i,k}^{t,\prime} \right) \right\| \\ &\leq \left(1 + \beta \eta_{k}^{t} \right) \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} \right\| + \eta^{t} \left\| g_{i} \left(\theta_{i,k}^{t,\prime} \right) - g_{i}^{\prime} \left(\theta_{i,k}^{t,\prime} \right) - g_{i}^{\prime} \left(\theta_{i,k}^{t,\prime} \right) \right\|.$$

$$(107)$$

Combining these two cases for client i we have

$$\mathbb{E} \left\| \theta_{i,k+1}^{t} - \theta_{i,k+1}^{t,\prime} \right\| \leq (1 + \beta\eta^{t}) \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} \right\| + \frac{\eta^{t}}{n} \mathbb{E} \left\| g_{i} \left(\theta_{i,k}^{t,\prime} \right) - g_{i}^{\prime} \left(\theta_{i,k}^{t,\prime} \right) \right\| \\
\leq (1 + \beta\eta_{k}^{t}) \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} \right\| + \frac{\eta_{k}^{t}}{n} \mathbb{E} \left\| g_{i} \left(\theta_{i,k}^{t} \right) \right\| \\
\leq (1 + \beta\eta^{t}) \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} \right\| + \frac{2\eta_{k}^{t}}{n} \left(1 + (1 + c)^{K-1}\beta\tilde{\eta}_{t} \right) \left(\sigma + \mathbb{E} \left\| \nabla F \left(\theta_{t} \right) \right\| + \sigma_{g,i} \right) \\
\leq (1 + \beta\eta^{t}) \mathbb{E} \left\| \theta_{i,k}^{t} - \theta_{i,k}^{t,\prime} \right\| + \frac{2\eta^{t}c}{n} \left(\mathbb{E} \left\| \nabla F \left(\theta_{t} \right) \right\| + \sigma_{g,i} + \sigma \right)$$
(108)

We let \overline{c} be an upper bound of $1 + (1+c)^{K-1}\beta \tilde{\eta}_t$ since $\tilde{\eta}_t$ is bounded above. Then unrolling it gives

$$\mathbb{E}\left\|\theta_{i,K}^{t}-\theta_{i,K}^{t,\prime}\right\| \leq \prod_{k=0}^{K-1}\left(1+\beta\eta_{k}^{t}\right)\mathbb{E}\left\|\theta_{t}-\theta_{t}^{\prime}\right\| + \left(\frac{2}{n}\sum_{k=0}^{K-1}\eta_{k}^{t}c\prod_{l=k+1}^{K-1}\left(1+\beta\eta_{l}^{t}\right)\left(\mathbb{E}\left\|\nabla F\left(\theta_{t}\right)\right\|+\sigma_{g,i}+\sigma\right)\right) \\ \leq e^{\beta\tilde{\eta}_{t}}\mathbb{E}\left\|\theta_{t}-\theta_{t}^{\prime}\right\| + \frac{2}{n}c\tilde{\eta}_{t}e^{\beta\tilde{\eta}_{t}}\left(\mathbb{E}\left\|\nabla F\left(\theta_{t}\right)\right\|+\sigma_{g,i}+\sigma\right).$$
(109)

We have

$$\mathbb{E} \left\| \theta_{t+1} - \theta_{t+1}' \right\| \leq \sum_{i=1}^{m} \frac{1}{m} \mathbb{E} \left\| \theta_{i,K}^{t} - \theta_{i,K}^{t,\prime} \right\|$$

$$\leq e^{\beta \tilde{\eta}_{t}} \mathbb{E} \left\| \theta_{t} - \theta_{t}' \right\| + \frac{2}{mn} c \tilde{\eta}_{t} e^{\beta \tilde{\eta}_{t}} \left(\mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \sigma_{g} + \sigma \right)$$
(110)

Further, unrolling the above over t and noting $\theta_0 = \theta_0'$, we obtain

$$\mathbb{E} \left\| \theta_T - \theta_T' \right\| \le \frac{2\tilde{c}}{mn} \sum_{t=0}^{T-1} \exp\left(\beta \sum_{l=t+1}^{T-1} \tilde{\eta}_t\right) \tilde{\eta}_t e^{\beta \tilde{\eta}_t} \left(\mathbb{E} \left\| \nabla F\left(\theta_t\right) \right\| + \sigma_g + \sigma\right),$$
(111)

$$\mathbb{E} \left\| \theta_T - \theta'_T \right\| \le \frac{2\overline{c}}{mn} \frac{1}{\beta} e^{\frac{1}{T} + 1} \left(\mathbb{E} \left\| \nabla F\left(\theta_t\right) \right\| + \sigma_g + \sigma \right), \tag{112}$$

With the Assumption 2,

$$\varepsilon_{\text{gen}} \leq \frac{2L\bar{c}}{mn} \frac{1}{\beta} e^{1+\frac{1}{T}} \left(\mathbb{E} \left\| \nabla F\left(\theta_{t}\right) \right\| + \sigma_{g} + \sigma \right),$$
$$\bar{\mathbf{c}} \leq 1 + (1+c)^{K-1} \beta \tilde{\eta}_{t} \leq 1 + (1+c)^{K-1} \beta \sum_{k=0}^{K-1} \eta^{t}$$
$$\tilde{\mathbf{c}} \leq 1 + (1+c)^{K-1} \beta \tilde{\eta}_{t} \leq 1 + (1+c)^{K-1} \beta \sum_{k=0}^{K-1} \eta_{k}^{t}$$
$$\tilde{\mathbf{c}} \leq \bar{\mathbf{c}}$$

When the diminishing stepsizes are chosen in the statement of the theorem, we conclude the proof. The proof on MoFedSAM can be derived simply from FedSAM.