Faster Non-Convex Federated Learning via Global and Local Momentum (Supplementary Material)

Rudrajit Das[1] **Anish Acharya** [*1] **Abolfazl Hashemi** [*2] **Sujay Sanghavi**[1] **Inderjit S. Dhillon**[1] **Ufuk Topcu**[1]

[1]University of Texas at Austin, USA
[2]Purdue University, West Lafayette, Indiana, USA

Abstract

We propose FedGLOMO, a novel federated learning (FL) algorithm with an iteration complexity of $\mathcal{O}(\epsilon^{-1.5})$ to converge to an ϵ-stationary point (i.e., $\mathbb{E}[\|\nabla f(x)\|^2] \leq \epsilon$) for smooth non-convex functions – under arbitrary client heterogeneity and compressed communication – compared to the $\mathcal{O}(\epsilon^{-2})$ complexity of most prior works. Our key algorithmic idea that enables achieving this improved complexity is based on the observation that the convergence in FL is hampered by two sources of high variance: (i) the global server aggregation step with multiple local updates, exacerbated by client heterogeneity, and (ii) the noise of the local client-level stochastic gradients. The first issue is particularly detrimental to FL algorithms that perform plain averaging at the server. By modeling the server aggregation step as a generalized gradient-type update, we propose a variance-reducing momentum-based global update at the server, which when applied in conjunction with variance-reduced local updates at the clients, enables FedGLOMO to enjoy an improved convergence rate. Our experiments illustrate the intrinsic variance reduction effect of FedGLOMO, which implicitly suppresses client-drift in heterogeneous data distribution settings and promotes communication efficiency.

1 INTRODUCTION

Federated learning (FL) is a new edge-computing approach that advocates training statistical models directly on remote devices by leveraging enhanced local resources on each device (McMahan et al. [2017]). In a standard FL setting, there are n clients, each having its own training data, and a central server that is trying to train a model, parameterized by $\boldsymbol{w} \in \mathbb{R}^d$, using the clients' data. Suppose the data distribution of the i^{th} client is $\mathcal{D}_i$. Then the i^{th} client has an objective function $f_i(\boldsymbol{w})$ which is the expected loss, with respect to some loss function ℓ, over data drawn from $\mathcal{D}_i$, and the goal of the central server is to optimize the average [1] loss $f(\boldsymbol{w})$, over the n clients, i.e.,

$$f(\boldsymbol{w}) := \frac{1}{n}\sum_{i=1}^{n} f_i(\boldsymbol{w}) \ \& \ f_i(\boldsymbol{w}) = \mathbb{E}_{\boldsymbol{x}\sim\mathcal{D}_i}[\ell(\boldsymbol{x}, \boldsymbol{w})]. \quad (1)$$

The setting where the data distributions of all the clients are identical, i.e. $\mathcal{D}_1 = \ldots = \mathcal{D}_n$, is typically known as the "homogeneous" setting. Otherwise, the settings where the data distributions are *not* identical are referred to as the "heterogeneous" settings.

The core algorithmic idea of FL – in the form of FedAvg – was introduced in McMahan et al. [2017]. In FedAvg (summarized in Algorithm 3), a *subset* of the clients perform *multiple* steps of gradient descent based updates on their local data and then communicate back their respective updates to the server, which then averages them to update the global model (hence the name FedAvg). This idea of performing multiple local updates before averaging once reduces the communication cost required for training. Another essential strategy in FL to cut down the communication cost is to have the clients send compressed/quantized messages to the server in every round – this is of particular significance for training deep learning models where the number of model parameters is in millions or more.

In practice however, performing multiple local updates on clients with *heterogeneous* data distributions leads to the so-called phenomenon of "client drift", wherein the individual client updates do not align well (due to over-fitting on the local client data) inhibiting the convergence of FedAvg to the optimum of the average loss over all the clients. In this paper, we identify the high variance associated with the

[*]Equal Contribution

[1]In general this may be a weighted average, but here we only consider uniform weights, i.e., each weight is $1/n$.

Accepted for the 38th Conference on Uncertainty in Artificial Intelligence (UAI 2022).

simple averaging step of FedAvg for the global update to be at the heart of this issue.

Ever since the development of FL, significant attention has been devoted to analyzing FedAvg under different settings, modifying FedAvg using ideas from centralized optimization to accelerate the training or to reduce the communication cost; we discuss these works in Section 2. Compared to centralized optimization, a formidable challenge in the theoretical analysis of FL algorithms is the use of multiple local updates in the clients which is compounded by the *heterogeneous* nature of data distribution among the clients. To limit the extent of client heterogeneity, a standard assumption in FL theory is the *bounded client dissimilarity (BCD) assumption*, i.e.,

$$\frac{1}{n}\sum_{i=1}^{n}\|\nabla f_i(\boldsymbol{w}) - \nabla f(\boldsymbol{w})\|^2 \leq G^2 \;\forall\, \boldsymbol{w}, \tag{2}$$

for some large enough constant $G < \infty$ (e.g., see A1 in Karimireddy et al. [2020]). But this assumption is limiting as it does not allow for *arbitrarily large client heterogeneity*.

Recently, Arjevani et al. [2019] showed that the stochastic first-order complexity of any algorithm in the *centralized setting* to reach an ϵ-stationary point (i.e., $\mathbb{E}[\|\nabla f(\boldsymbol{x})\|^2] \leq \epsilon$) for *smooth non-convex functions* is $\Omega(\epsilon^{-1.5})$. It is well known that vanilla SGD has a suboptimal complexity of $\mathcal{O}(\epsilon^{-2})$ as it cannot mitigate the high variance of the stochastic gradient noise. Recognizing this issue, *variance-reducing* techniques for SGD (Fang et al. [2018], Zhou et al. [2018], Cutkosky and Orabona [2019], Liu et al. [2020]) have been proposed that attain the optimal complexity of $\mathcal{O}(\epsilon^{-1.5})$. Coming to the federated setting, as we discuss in this paper, in addition to the noise in the *local* client-level stochastic gradients, one has to also contend with the high variance associated with the *global* server aggregation step which depends on the client heterogeneity and the number of local update steps. In this case, as we argue in the subsequent sections, applying only local client-level variance-reduction is not enough for improving the iteration complexity of vanilla FedAvg beyond $\mathcal{O}(\epsilon^{-2})$ for smooth, non-convex losses.

To alleviate the issue of variance due to heterogeneity, we propose a novel FL algorithm with *compressed communication* called FedGLOMO (Algorithm 1 and 2) which applies G*lobal* as well as LO*cal variance-reducing* MO*mentum* to the server update and client updates, respectively. We prove that the iteration complexity of FedGLOMO is $\mathcal{O}(\epsilon^{-1.5})$ in the smooth non-convex case, which is better than the $\mathcal{O}(\epsilon^{-2})$ complexity of related works in the FL setting; see Table 1 and Theorem 1. Further, our theory does not use the BCD assumption, i.e. eq. (2), which is a standard assumption in related works. Instead, we propose and use Assumption 4, which is a more realistic and *empirically verified* assumption on the client drift, even allowing for arbitrary client heterogeneity. It is worth mentioning here that for FL, Karimireddy et al. [2020] also propose an algorithm (MimeMVR) which is shown to attain this improved complexity of $\mathcal{O}(\epsilon^{-1.5})$ but *with* the BCD assumption and *no* compressed communication; we talk about this at the end of Section 2.

We summarize our **contributions** next:

(a) We propose FedGLOMO (Alg. 1 and 2), in which we apply a *novel global momentum term at the server* in addition to *local momentum at the clients*. The design of FedGLOMO is motivated by two critical issues that need to be alleviated to accelerate convergence in FL; these are the high variances associated with: (i) the *global* server aggregation step due to heterogeneity of clients when there are multiple local updates, and (ii) the noise of *local* client-level stochastic gradients. Global and local momentum result in *variance reduction* for the global server update and the local client updates, allowing us to tackle (i) and (ii), respectively. This enables FedGLOMO to converge to an ϵ-stationary point (i.e., $\mathbb{E}[\|\nabla f(\boldsymbol{x})\|^2] \leq \epsilon$) for smooth non-convex functions in $\mathcal{O}(\epsilon^{-1.5})$ gradient-based updates, which is better than the $\mathcal{O}(\epsilon^{-2})$ complexity of most related works in the FL setting; see Table 1 and Theorem 1.

(b) Unlike prior work, our theory does not use the limiting bounded client dissimilarity assumption (i.e., eq. (2)). Instead, to tighten our result, we propose and use Assumption 4 – which is a novel assumption on the client drift, even allowing for *arbitrary client heterogeneity* in the worst case. We empirically verify that Assumption 4 holds for FedGLOMO as well as FedAvg. Theoretically, we also show that Assumption 4 holds for *any* FL algorithm in the case of linear regression and also with networks whose training dynamics follow that of a linearized model (a.k.a. the "NTK" regime). Refer to the discussion after Assumption 4 and Remark 2 for details.

(c) FedGLOMO is the *first FL algorithm* achieving $\mathcal{O}(\epsilon^{-1.5})$ complexity while allowing *compressed client-to-server communication*. We emphasize that from the theory perspective, applying compression in FedGLOMO is not trivial and the most obvious approach does not work; see Remark 3.

(d) In Section 6, experiments on CIFAR-10 and Fashion-MNIST (Xiao et al. [2017]) show that in a highly heterogeneous setting of at most two (out of ten) classes per client, FedGLOMO requires only about *one-third* the number of bits used by FedAvg with PyTorch's default momentum applied to the local client updates; see Figure 1. Our experiments also illustrate the variance reduction provided by our scheme which implicitly mitigates client-drift under heterogeneous data distribution and in turn promotes communication-efficiency.

2 RELATED WORK

FedAvg and related methods: Reisizadeh et al. [2020] propose FedPAQ which is basically FedAvg (McMahan et al. [2017]) with quantized client-to-server communication, and establish its convergence for the homogeneous case. Li et al. [2019] establish the convergence of FedAvg for strongly convex functions with heterogeneity (assuming bounded client dissimilarity) but without any compressed communication. Haddadpour et al. [2021] propose FedCOMGATE which incorporates gradient tracking (Pu and Nedić [2020]) and derive results with data heterogeneity and quantized communication. Karimireddy et al. [2019] propose SCAFFOLD which uses control-variates to mitigate the client-drift owing to the heterogeneity of clients. Li et al. [2018] present FedProx which adds a proximal term to control the deviation of the client parameters from the global server parameter in the previous round. Reddi et al. [2020] propose federated versions of commonly used adaptive optimization methods and prove their convergence under heterogeneity. Local SGD (Zinkevich et al. [2010], Stich [2018], Yu et al. [2018], Wang and Joshi [2018], Basu et al. [2019], Stich and Karimireddy [2019], Patel and Dieuleveut [2019], Woodworth et al. [2020], Bayoumi et al. [2020], Liang et al. [2019], Koloskova et al. [2020]) is very similar to FL and is essentially based on the same principle as FedAvg. However, in local SGD, there is usually no data heterogeneity and all the clients participate in each round (known as "full device participation"), both of which do not hold in FL and simplify the derivation of convergence results.

Wang et al. [2019], Huo et al. [2020] present momentum-based updates at the server without any improvement in the convergence rate as compared to momentum-free updates. Qu et al. [2020] present Nesterov accelerated FedAvg for convex objectives. Karimireddy et al. [2020] propose Mime(MVR) which applies momentum at the client-level based on globally computed statistics to control client-drift. Khanduri et al. [2021] propose STEM which applies momentum globally and locally for local SGD; however, their server aggregation step is just plain averaging as they do not have deal with server-side variance reduction, since all the clients participate in local SGD.

Distributed optimization with compression: References Alistarh et al. [2017], Suresh et al. [2017], Reisizadeh et al. [2020], Haddadpour et al. [2021], Tang et al. [2018], Wu et al. [2018], Bernstein et al. [2018], Alistarh et al. [2018], Lin et al. [2017], Stich et al. [2018], Basu et al. [2019], Hashemi et al. [2021], Chen et al. [2020, 2021] aim to minimize the communication bottleneck in distributed optimization by transmitting compressed messages to the central server and establishing their convergence. Horváth et al. [2019], Gorbunov et al. [2021] provide distributed algorithms with improved convergence rates by also applying variance reduction and periodically using full gradients; however, there are no multiple local updates in these works. In Appendix D, we compare our work's complexity against that of Gorbunov et al. [2021]. In this work, we employ the quantization operator of Alistarh et al. [2017].

Complexity for smooth non-convex stochastic optimization: Arjevani et al. [2019] show that the optimal stochastic first-order complexity to reach an ϵ-stationary point (i.e., $\mathbb{E}[\|\nabla f(\boldsymbol{x})\|^2] \leq \epsilon$) is $\mathcal{O}(\frac{\sigma}{\epsilon^{1.5}})$ where σ^2 is the variance of the stochastic gradients. Unfortunately, vanilla SGD is sub-optimal and *variance-reducing* techniques must be applied to attain the optimal complexity; some noteworthy works on variance-reduction for SGD are SVRG (Johnson and Zhang [2013]), SAGA (Defazio et al. [2014]) and SARAH (Nguyen et al. [2017]). SVRG-style algorithms such as SPIDER (Fang et al. [2018]) and SNVRG (Zhou et al. [2018]) attain this optimal complexity by periodically using giant batch sizes. Cutkosky and Orabona [2019] propose STORM which also attains this optimal complexity with adaptive learning rates, but without using any large batches. The key idea of STORM is momentum-based variance reduction, obtained by using the stochastic gradient at the previous point *computed over the same batch* on which the stochastic gradient at the current point is computed. Liu et al. [2020] present a much simpler proof for essentially the same algorithm by employing a constant learning rate and requiring a large batch size only at the first iteration. Our key idea of global and local momentum is STORM-like *variance-reducing* momentum applied to the aggregation step at the server, interpreted as a generalized gradient-type update, and the local client updates, respectively; see Section 4.

Table 1 compares the complexities of the most relevant related works in FL ($r < n$) and local SGD ($r = n$) with ours on smooth non-convex functions. Note that under the more challenging FL setting with partial-device participation, only FedGLOMO and MimeMVR (Karimireddy et al. [2020]) attain the improved iteration complexity of $\mathcal{O}(\epsilon^{-1.5})$ with respect to ϵ. However, unlike Karimireddy et al. [2020], our work does not rely on the bounded client dissimilarity assumption (eq. (2)) and allows for compressed client-to-server communication, in which case maintaining the improved complexity is not trivial; for details, see Remarks 2 and 3, respectively. There are meaningful algorithmic differences between our work and Karimireddy et al. [2020] too. The biggest one is that while we explicitly apply momentum in the server aggregation step (global momentum) as well as in the client updates (local momentum), Karimireddy et al. [2020] only apply *globally computed* momentum in the local client updates. For a detailed discussion of the differences of our work from Karimireddy et al. [2020], see Appendix C. Since Mime is designed to deal with client drift, we empirically compare it against FedGLOMO without compression in a highly heterogeneous setting in Section 6.

3 PRELIMINARIES

Recall the setting and the optimization problem that the server is trying to solve as defined in eq. (1). We assume that the clients have access to unbiased stochastic gradients

Table 1: Number of gradient updates, i.e., T, required to achieve $\mathbb{E}[\|\nabla f(\boldsymbol{w})\|^2] \leq \epsilon$ on smooth non-convex functions. Here, n is the total number of clients and r is the number of clients participating in each round. "Client Participation" asks whether all ($r = n$) or only a subset ($r < n$) of the clients participate in each round. "BCD?" asks if the bounded client dissimilarity assumption (eq. (2)) is used or not. "Compression?" asks whether compressed communication is involved or not.
$*1$: $\alpha \leq n$ is a problem-dependent quantity; in practice, we expect $\alpha \ll n$ as confirmed in our experiments.

Ref.	T	Client Participation	BCD?	Compression?
Koloskova et al. [2020], Wang et al. [2019]	$\mathcal{O}(\frac{1}{n\epsilon^2})$	Full ($r = n$)	Yes	✗
Haddadpour et al. [2021]	$\mathcal{O}(\frac{1}{n\epsilon^2})$	Full ($r = n$)	Yes	✓
Khanduri et al. [2021]	$\mathcal{O}(\frac{1}{n\epsilon^{1.5}})$	Full ($r = n$)	Yes	✗
Karimireddy et al. [2019]	$\mathcal{O}(\frac{1}{r\epsilon^2})$	Partial ($r < n$)	Yes	✗
Karimireddy et al. [2020]	$\mathcal{O}(\frac{1}{\sqrt{r}\epsilon^{1.5}})$	Partial ($r < n$)	Yes	✗
This work (`FedGLOMO`)	$\mathcal{O}\big(\max\big(\sqrt{\frac{\alpha}{n}}, \frac{1}{\sqrt{r}}\big)\frac{1}{\epsilon^{1.5}}\big)^{*1}$	Partial ($r < n$)	**No**	✓

of their individual losses. We denote the stochastic gradient of f_i at $\boldsymbol{w}$ computed over a batch of samples $\mathcal{B}$, by $\widetilde{\nabla} f_i(\boldsymbol{w}; \mathcal{B})$. Also in this paper, K is the number of communication rounds, E is the number of local updates per round or the period, and $T = KE$ is the total number of local updates or the (order-wise) number of gradient-based updates. Further, r is the number of clients that the server accesses in each round, i.e., the global batch size.

Vectors and matrices are written in boldface. For any positive integer m, the set $\{1, \ldots, m\}$ is denoted by $[m]$, and the uniform distribution over the set $\{0, \ldots, m\}$ is denoted by Unif$[0, m]$. $\mathbb{1}(.)$ is the indicator function. Next, we recap smooth functions.

Definition 3.1 (**Smoothness).** A function $g : \Theta \to \mathbb{R}$ is to said to be L-smooth if for all $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \Theta$, $\|\nabla g(\boldsymbol{\theta}) - \nabla g(\boldsymbol{\theta}')\| \leq L\|\boldsymbol{\theta} - \boldsymbol{\theta}'\|$. For all $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \Theta$, we also have: $g(\boldsymbol{\theta}') \leq g(\boldsymbol{\theta}) + \langle \nabla g(\boldsymbol{\theta}), \boldsymbol{\theta}' - \boldsymbol{\theta} \rangle + \frac{L}{2}\|\boldsymbol{\theta}' - \boldsymbol{\theta}\|^2$.

Algorithm 1 `FedGLOMO` - Server Update

1: **Input:** Initial point $\boldsymbol{w}_0$, # of rounds of communication K, period E, learning rates $\{\eta_k\}_{k=0}^{K-1}$ and global batch size r. Q_D is the quantization operator. Set $\boldsymbol{w}_{-1} = \boldsymbol{w}_0$.
2: **for** $k = 0, \ldots, K-1$ **do**
3: Server sends $\boldsymbol{w}_k, \boldsymbol{w}_{k-1}$ to a set $\mathcal{S}_k$ of r clients chosen uniformly at random w/o replacement.
4: **for** client $i \in \mathcal{S}_k$ **do**
5: Set $\boldsymbol{w}_{k,0}^{(i)} = \boldsymbol{w}_k$ and $\widehat{\boldsymbol{w}}_{k-1,0}^{(i)} = \boldsymbol{w}_{k-1}$. Run Algorithm 2 for client i.
6: **end for**
7: **if** $k = 0$ **then**
8: Set $\boldsymbol{u}_k = \frac{1}{r}\sum_{i \in \mathcal{S}_k} Q_D(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)})$.
9: **else**
10: Set $\boldsymbol{u}_k = \frac{\beta_k}{r}\sum_{i \in \mathcal{S}_k} Q_D(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)}) + (1 - \beta_k)\boldsymbol{u}_{k-1} + \frac{(1-\beta_k)}{r}\sum_{i \in \mathcal{S}_k} Q_D((\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)}) - (\boldsymbol{w}_{k-1} - \widehat{\boldsymbol{w}}_{k-1,E}^{(i)}))$. `// (Global Momentum)`
11: **end if**
12: Update $\boldsymbol{w}_{k+1} = \boldsymbol{w}_k - \boldsymbol{u}_k$.
13: **end for**

4 FEDGLOMO: GLOBAL AND LOCAL MOMENTUM-BASED VARIANCE REDUCTION

There are two issues that need to be alleviated for improving the convergence rate in FL: (i) the high variance of simple averaging used in the *global* server aggregation step (of `FedAvg`), when there are multiple local updates, which is exacerbated by heterogeneity of the clients, and (ii) the high variance associated with the noise of *local* client-level stochastic gradients. The key idea of `FedGLOMO` (Algorithm 1 and 2) is to apply *variance-reducing* **global** and **local** momentum to combat (i) and (ii), respectively. We now describe global and local momentum in detail.

Global momentum is applied to the sever aggregation step which is line 10 in Algorithm 1. To understand it better, let us revisit `FedAvg` (summarized in Algorithm 3, although in a slightly different way than usual) and its server aggregation step (line 12) which is just simple averaging. Similar to the update of SGD suffering from high variance, this naive averaging step – which we think of as the average of a batch of generalized stochastic gradients – is characterized by high variance stemming from heterogeneity and multiple local updates. So, this way of server aggregation slows down the convergence rate of `FedAvg` (and other related methods).

In this paper, we re-envision the server aggregation as a generalized gradient-based update by thinking of $(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)})$ as the generalized gradient. Then, we wish to incorporate the style of variance-reducing momentum applied in `STORM` (Cutkosky and Orabona [2019], Liu et al. [2020]) to our gen-

Algorithm 2 `FedGLOMO` - Client Update

1: **for** $\tau = 0, \ldots, E-1$ **do**
2:   **if** $\tau = 0$ **then**
3:     Set $\boldsymbol{v}_{k,\tau}^{(i)} = \nabla f_i(\boldsymbol{w}_{k,\tau}^{(i)})$, $\widehat{\boldsymbol{v}}_{k-1,\tau}^{(i)} = \nabla f_i(\widehat{\boldsymbol{w}}_{k-1,\tau}^{(i)})$.
4:   **else**
5:     Pick a random batch of samples in client i, say $\mathcal{B}_{k,\tau}^{(i)}$. Compute the stochastic gradients of f_i at $\boldsymbol{w}_{k,\tau}^{(i)}$, $\widehat{\boldsymbol{w}}_{k-1,\tau}^{(i)}$, $\boldsymbol{w}_{k,\tau-1}^{(i)}$ and $\widehat{\boldsymbol{w}}_{k-1,\tau-1}^{(i)}$ over $\mathcal{B}_{k,\tau}^{(i)}$ viz. $\widetilde{\nabla} f_i(\boldsymbol{w}_{k,\tau}^{(i)}; \mathcal{B}_{k,\tau}^{(i)})$, $\widetilde{\nabla} f_i(\widehat{\boldsymbol{w}}_{k-1,\tau}^{(i)}; \mathcal{B}_{k,\tau}^{(i)})$, $\widetilde{\nabla} f_i(\boldsymbol{w}_{k,\tau-1}^{(i)}; \mathcal{B}_{k,\tau}^{(i)})$ and $\widetilde{\nabla} f_i(\widehat{\boldsymbol{w}}_{k-1,\tau-1}^{(i)}; \mathcal{B}_{k,\tau}^{(i)})$.
6:     Update: $\boldsymbol{v}_{k,\tau}^{(i)} = \widetilde{\nabla} f_i(\boldsymbol{w}_{k,\tau}^{(i)}; \mathcal{B}_{k,\tau}^{(i)}) + \big(\boldsymbol{v}_{k,\tau-1}^{(i)} - \widetilde{\nabla} f_i(\boldsymbol{w}_{k,\tau-1}^{(i)}; \mathcal{B}_{k,\tau}^{(i)})\big)$ and $\widehat{\boldsymbol{v}}_{k-1,\tau}^{(i)} = \widetilde{\nabla} f_i(\widehat{\boldsymbol{w}}_{k-1,\tau}^{(i)}; \mathcal{B}_{k,\tau}^{(i)}) + \big(\widehat{\boldsymbol{v}}_{k-1,\tau-1}^{(i)} - \widetilde{\nabla} f_i(\widehat{\boldsymbol{w}}_{k-1,\tau-1}^{(i)}; \mathcal{B}_{k,\tau}^{(i)})\big)$. `// (Local Mom.)`
7:   **end if**
8:   Update $\boldsymbol{w}_{k,\tau+1}^{(i)} = \boldsymbol{w}_{k,\tau}^{(i)} - \eta_k \boldsymbol{v}_{k,\tau}^{(i)}$ and $\widehat{\boldsymbol{w}}_{k-1,\tau+1}^{(i)} = \widehat{\boldsymbol{w}}_{k-1,\tau}^{(i)} - \eta_k \widehat{\boldsymbol{v}}_{k-1,\tau}^{(i)}$.
9: **end for**
10: Send $Q_D(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)})$ and $Q_D((\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)}) - (\boldsymbol{w}_{k-1} - \widehat{\boldsymbol{w}}_{k-1,E}^{(i)}))$ to the server.

eralized gradient-based update; note that their method is for stochastic gradients in the case of centralized optimization. To that end, let us briefly recap `STORM`'s update rule. For a function $h(\boldsymbol{z})$, `STORM`'s update for the j^{th} iteration is:

$$\boldsymbol{z}_{j+1} = \boldsymbol{z}_j - \eta_j \boldsymbol{v}_j, \text{ where } \boldsymbol{v}_j = \{\widetilde{\nabla} h(\boldsymbol{z}_j; \xi_j) + (1-\beta_j)(\boldsymbol{v}_{j-1} - \widetilde{\nabla} h(\boldsymbol{z}_{j-1}; \xi_j))\mathbb{1}(j > 0)\}. \quad (3)$$

In eq. (3), ξ_j denotes the source of randomness in the j^{th} iteration and $\beta_j \in [0, 1)$ is the momentum parameter. Note the use of the stochastic gradient at $\boldsymbol{z}_{j-1}$ computed on ξ_j. Coming back to Algorithm 1, the quantity $\boldsymbol{u}_k$ plays the role of $\boldsymbol{v}_j$ in eq. (3). To see this clearly, let us analyze $E_{Q_D}[\boldsymbol{u}_k]$ (see lines 8 and 10 in Algorithm 1). Under Assumption 3, the compression operator Q_D produces an unbiased estimate of the input. Then defining $g(\boldsymbol{w}_k; \mathcal{S}_k) \triangleq \frac{1}{r}\sum_{i \in \mathcal{S}_k}(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)})$ and $\widehat{g}(\boldsymbol{w}_{k-1}; \mathcal{S}_k) \triangleq \frac{1}{r}\sum_{i \in \mathcal{S}_k}(\boldsymbol{w}_{k-1} - \widehat{\boldsymbol{w}}_{k-1,E}^{(i)})$, we have:

$$\mathbb{E}_{Q_D}[\boldsymbol{u}_k] = \{g(\boldsymbol{w}_k; \mathcal{S}_k) + (1-\beta_k)\big(\boldsymbol{u}_{k-1} - \widehat{g}(\boldsymbol{w}_{k-1}; \mathcal{S}_k)\big)\mathbb{1}(k > 0)\}. \quad (4)$$

In eq. (4), $g(\boldsymbol{w}_k; \mathcal{S}_k)$ and $\widehat{g}(\boldsymbol{w}_{k-1}; \mathcal{S}_k)$ play the roles of $\widetilde{\nabla} h(\boldsymbol{z}_j; \xi_j)$ and $\widetilde{\nabla} h(\boldsymbol{z}_{j-1}; \xi_j)$, respectively. With this, one can clearly see that eq. (4) is the analogue of eq. (3) for the global server aggregation in FL. However, this equivalence is not so apparent without looking at the expected value of $\boldsymbol{u}_k$ with respect to Q_D; in fact, the choice of quantities that

Algorithm 3 `FedAvg` McMahan et al. [2017]

Input: Initial point $\boldsymbol{w}_0$, # of communication rounds K, period E, learning rates $\{\eta_k\}_{k=0}^{K-1}$ and global batch size r.
for $k = 0, \ldots, K-1$ **do**
  Server sends $\boldsymbol{w}_k$ to a set $\mathcal{S}_k$ of r clients chosen uniformly at random w/o replacement.
  for client $i \in \mathcal{S}_k$ **do**
    Set $\boldsymbol{w}_{k,0}^{(i)} = \boldsymbol{w}_k$.
    for $\tau = 0, \ldots, E-1$ **do**
      Pick a random batch of samples in client i, $\mathcal{B}_{k,\tau}^{(i)}$. Compute the stochastic gradient of f_i at $\boldsymbol{w}_{k,\tau}^{(i)}$ over $\mathcal{B}_{k,\tau}^{(i)}$, viz. $\widetilde{\nabla} f_i(\boldsymbol{w}_{k,\tau}^{(i)}; \mathcal{B}_{k,\tau}^{(i)})$.
      Update $\boldsymbol{w}_{k,\tau+1}^{(i)} = \boldsymbol{w}_{k,\tau}^{(i)} - \eta_k \widetilde{\nabla} f_i(\boldsymbol{w}_{k,\tau}^{(i)}; \mathcal{B}_{k,\tau}^{(i)})$.
    end for
    Send $(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)})$ to the server.
  end for
  Update $\boldsymbol{w}_{k+1} = \boldsymbol{w}_k - \frac{1}{r}\sum_{i \in \mathcal{S}_k}(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)})$.
end for

are compressed in line 10 of Alg. 2 and used in line 10 of Alg. 1 is crucial for establishing provable guarantees (also see Remark 3).

Now that we understand global momentum, let us move on to **local** momentum. For this see lines 3, 6 and 8 in Algorithm 2; these give us $(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)})$ and $(\boldsymbol{w}_{k-1} - \widehat{\boldsymbol{w}}_{k-1,E}^{(i)})$ after running for E steps. But notice that these lines are the same as eq. (3) with $\beta_j = 0$ and the stochastic gradient at the first iteration replaced by the full gradient. It is worth mentioning here that these local updates are also similar to `SPIDER` which is an SVRG-style update proposed in Fang et al. [2018]. However, recognizing that this is also a special case of the `STORM` update with $\beta_j = 0$, we prefer calling it momentum in order to have a unifying terminology for both the global and local updates.

One might wonder what is the role of global momentum as `SPIDER` can be extended to improve the complexity in distributed optimization *without multiple local updates*. For this, in Appendix F, we consider `FedLOMO` (Algorithm 4 and 5) which is a simpler version of `FedGLOMO` with only local momentum and *no* global momentum (i.e, plain averaging at the server which is equivalent to setting $\beta_k = 1$ in Algorithm 1), and show that it does not achieve $\mathcal{O}(\epsilon^{-1.5})$ complexity under partial-device participation and compression (see Theorem 3). The root cause of this is client heterogeneity which amplifies its effect under *multiple local updates*; without incorporating some form of variance reduction in the server aggregation step, the complexity cannot be improved.

Let us try to provide some intuition as to how incorporating global momentum helps. Suppose we keep $\eta_k = \eta$ and $\beta_k = \beta < 1$ for all k. Theoretically, we get a lower bound for β which is $\mathcal{O}(\eta^2)$. Then with this momentum-based aggregation strategy, the variance reduces by a factor of $\mathcal{O}(\beta/\eta) = \mathcal{O}(\eta)$ as compared to aggregation by plain averaging. (There are some other terms too but these are sufficiently small.) This reduction in the variance by a factor of $\mathcal{O}(\eta)$ is what improves the convergence rate of `FedGLOMO`.

It is true that `FedGLOMO` has to communicate twice the amount of information per round as compared to `FedAvg` or `FedPAQ` (Reisizadeh et al. [2020]) which is just `FedAvg` with compressed communication. One can set the precision of the quantizer sufficiently low to account for the extra per-round communication cost of `FedGLOMO` – we adopt this approach in our experiments. Also, we only assume access to the full client gradient in line 3 of Alg. 2 for simplicity of analysis, but our main result (i.e., Theorem 1) can be readily extended to the case of large enough batch sizes.

5 MAIN RESULT FOR FEDGLOMO

First, we state our assumptions.

Assumption 1 (Smoothness). *$\ell(\boldsymbol{x}, \boldsymbol{w})$ is L-smooth with respect to $\boldsymbol{w}$, for all $\boldsymbol{x}$. Thus, each $f_i(\boldsymbol{w})$ ($i \in [n]$) is L-smooth, and so is $f(\boldsymbol{w})$.*

Assumption 2 (Non-negativity). *Each $f_i(\boldsymbol{w})$ is non-negative and therefore, $f_i^* \triangleq \min f_i(\boldsymbol{w}) \geq 0$.*

Most loss functions used in practice satisfy this anyways and if not, we can just add a constant offset to achieve non-negativity.

Assumption 3 (Quantization). *The quantization operator Q_D in Alg. 1 and 2 is unbiased, i.e., $\mathbb{E}[Q_D(\boldsymbol{x})|\boldsymbol{x}] = \boldsymbol{x}$, and its variance satisfies $\mathbb{E}[\|Q_D(\boldsymbol{x}) - \boldsymbol{x}\|^2|\boldsymbol{x}] \leq q\|\boldsymbol{x}\|^2$ for some $q > 0$. The "qsgd" operator proposed in Section 3.1 of Alistarh et al. [2017] satisfies Assumption 3.*

Assumption 4 (Client Drift/Heterogeneity). *Let $\mathcal{A}$ be an FL algorithm with E local update steps and K communication rounds. Let $\boldsymbol{w}_{k,\tau}^{(i)}$ be the i^{th} client's local parameter at the start of the $(\tau+1)^{st}$ local step of the $(k+1)^{st}$ round of $\mathcal{A}$, for $i \in [n]$ (similar to the notation in Alg. 1, 2, and 3). Define $\widetilde{\boldsymbol{e}}_{k,\tau}^{(i)} \triangleq \nabla f_i(\boldsymbol{w}_{k,\tau}^{(i)}) - \nabla f_i\big(\frac{1}{n}\sum_{j\in[n]} \boldsymbol{w}_{k,\tau}^{(j)}\big)$. Then for some $\alpha \ll n$, the following holds:*

$$\mathbb{E}\Big[\Big\|\sum_{i\in[n]} \widetilde{\boldsymbol{e}}_{k,\tau}^{(i)}\Big\|^2\Big] \leq \alpha \sum_{i\in[n]} \mathbb{E}\Big[\Big\|\widetilde{\boldsymbol{e}}_{k,\tau}^{(i)}\Big\|^2\Big], \qquad (5)$$

$\forall\, \tau \in \{0, \dots, E-1\}$ and $k \in \{0, \dots, K-1\}$. The expectation above is w.r.t. any stochasticity in the local updates.

Equation (5) in the above assumption always holds with $\alpha = n$ for any FL algorithm; this follows from the fact that for any $m > 1$ vectors $\{\boldsymbol{a}_j\}_{j=1}^m$, $\|\sum_{j=1}^m \boldsymbol{a}_j\|^2 \leq m\sum_{j=1}^m \|\boldsymbol{a}_j\|^2$ (this can be obtained by using the Cauchy-Schwarz inequality). However, we empirically observe $\alpha \ll n$ in practice for `FedGLOMO` as well as `FedAvg`; see Section 6 and Appendix H, respectively. The value of α in Assumption 4 is a measure of the amount of client drift induced by the algorithm which also depends on the degree of heterogeneity in the system – as the heterogeneity increases (decreases), we observe α to also increase (decrease).

From Figure 3 (in Section 6), we see that for the highly heterogeneous setting that we consider for our experiments in Section 6, $\alpha < 0.06n$ for most of the trajectory of `FedGLOMO` on both CIFAR-10 and Fashion-MNIST (abbreviated as FMNIST). In the homogeneous case, $\alpha < 0.03n$ and $\alpha < 0.02n$ for most of the trajectory on CIFAR-10 and FMNIST, respectively. We observe a similar trend of α for `FedAvg` in Appendix H. Additionally, we derive a convergence result for `FedAvg` under Assumption 4 and without the bounded client dissimilarity assumption (i.e., eq. (2)) in Appendix H.

Some theoretical motivation for Assumption 4: Let us consider *linear regression* to provide a scenario where $\alpha = 0$ provably for **any** FL algorithm. Suppose in client i, we have feature and label pairs $(\boldsymbol{x}, y) \sim (\mathcal{X}_i, \mathcal{Y}_i)$, where the label

$$y = \langle \boldsymbol{w}_i^*, \boldsymbol{x}\rangle + \xi,$$

with $\xi \sim \mathcal{N}_i$ being independent zero-mean client-dependent random noise. Obviously, the label distribution $\mathcal{Y}_i$ here depends on the feature distribution $\mathcal{X}_i$, noise distribution $\mathcal{N}_i$ and $\boldsymbol{w}_i^*$. We assume that the covariance matrix of the feature vectors is the same across all the clients, i.e., $\mathbb{E}_{\boldsymbol{x}\sim\mathcal{X}_i}[\boldsymbol{x}\boldsymbol{x}^T] = \boldsymbol{Q}$ for all $i \in [n]$; this is possible for e.g., by normalization or whitening of the features. Note that by assuming the same covariance matrix across all the clients, we are *not* assuming that the feature distributions are the same across clients, but even if they are, there is heterogeneity through the different label distributions. Then, with the squared loss, our per-client objective function is:

$$f_i(\boldsymbol{w}) = \mathbb{E}_{(\boldsymbol{x},y)\sim(\mathcal{X}_i,\mathcal{Y}_i)}\Big[\frac{1}{2}(y - \langle \boldsymbol{w}, \boldsymbol{x}\rangle)^2\Big].$$

With the aforementioned conditions, it can be verified that $\nabla f_i(\boldsymbol{w}) = \boldsymbol{Q}(\boldsymbol{w} - \boldsymbol{w}_i^*)$. Thus,

$$\widetilde{\boldsymbol{e}}_{k,\tau}^{(i)} = \boldsymbol{Q}\Big(\boldsymbol{w}_{k,\tau}^{(i)} - \frac{1}{n}\sum_{j\in[n]} \boldsymbol{w}_{k,\tau}^{(j)}\Big),$$

and so $\sum_{i\in[n]} \widetilde{\boldsymbol{e}}_{k,\tau}^{(i)} = \vec{0}$. So, *Assumption 4 holds here with $\alpha = 0$ for any FL algorithm.*

In fact, the above analysis and result (i.e., $\alpha = 0$) can be extended to networks whose training dynamics follow that of a linearized model, which has been shown to be the case for infinite-width networks (see for e.g., Lee et al. [2019] and Jacot et al. [2018]) and has been also used on applications for finite-width networks (for e.g., in Mu et al. [2020]).

We now present the abridged version of the convergence result of FedGLOMO, followed by some important remarks. Its full version and detailed proof are in Appendix A and Appendix G.1, respectively.

Theorem 1 (**Smooth non-convex**). *Let Assumptions 1, 2 and 3 hold. Further, suppose Assumption 4 is true for* FedGLOMO*. In* FedGLOMO*, for each round* k*, set* $\eta_k = \eta = \mathcal{O}(\frac{1}{LEK^{1/3}C^{1/3}})$*, where* $C = \mathcal{O}\big(\max\big(\frac{\alpha}{n}, \frac{E^2(1+q)^2}{r}\big)\big)$*, and* $\beta_k = \mathcal{O}((1+q)\eta^2 L^2 E^4)$*. Suppose we use full-device participation (i.e., the global batch size is* n*)* ***only at*** $k = 0$*. Then,* FedGLOMO *can achieve* $\mathbb{E}_{k^*\sim\text{Unif}[0,K-1]}[\|\nabla f(\boldsymbol{w}_{k^*})\|^2] \leq \epsilon$ *in* $K = \mathcal{O}\big(\max\big(\sqrt{\frac{\alpha}{n}}, \frac{1+q}{\sqrt{r}}\big)\epsilon^{-1.5}\big)$ *rounds of communication and* $E = \mathcal{O}(1)$ *local steps.*

Remark 1 (**Better iteration complexity**). As per Theorem 1, for converging to an ϵ-stationary point, FedGLOMO needs $T = KE$ to be $\mathcal{O}\big(\max\big(\sqrt{\frac{\alpha}{n}}, \frac{1}{\sqrt{r}}\big)\epsilon^{-1.5}\big)$. This iteration complexity is the same as that of MimeMVR (Karimireddy et al. [2020]) *but without using the bounded client dissimilarity assumption*, i.e. eq. (2), (also see the next remark for more details on this) and better than other related works in the federated setting; see Table 1. We underscore the significance of global momentum here by comparing this complexity of FedGLOMO to that of FedLOMO (recall this is a simpler version of FedGLOMO with only local momentum and *no* global momentum, described in Appendix F) under partial-device participation and compression which is $\mathcal{O}\big(\frac{1}{r}\epsilon^{-2}\big)$; see Theorem 3.

Remark 2 (**No requirement of bounded client dissimilarity (BCD) assumption**). Divergent from related works, Theorem 1 *does not use* the commonly used BCD assumption, i.e., eq. (2). This is achieved by utilizing the smoothness and non-negativity of the f_i's, specifically $\frac{1}{n}\sum_{i\in[n]}\|\nabla f_i(\boldsymbol{w})\|^2 \leq \frac{1}{n}\sum_{i\in[n]} 2L(f_i(\boldsymbol{w}) - f_i^*) \leq 2Lf(\boldsymbol{w})$; see the proof outline of Theorem 1 in Appendix A. Instead of the BCD assumption, we use our empirically verified Assumption 4 to provide a tighter (when $\alpha \ll n$) and data-dependent convergence result. Note that Assumption 4 will always hold for some $\alpha \leq n$, regardless of the degree of client heterogeneity. Thus, Theorem 1 allows for *arbitrary client heterogeneity*.

Remark 3 (**Compressed communication**). To our knowledge, FedGLOMO is the *first algorithm* that attains the aforementioned improved iteration complexity for FL on smooth non-convex functions *with compressed communication*. We emphasize that the choice of quantities compressed in line 10 of Algorithm 2 is important. This particular choice enables deriving the improved complexity by first deriving a result analogous to smoothness, i.e., $\|(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)}) - (\boldsymbol{w}_{k-1} - \widehat{\boldsymbol{w}}_{k-1,E}^{(i)})\| \leq \widehat{L}\|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\|$ (see Lemma 9 in Appendix G.1). The straightforward choice of sending $Q_D(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)})$ and $Q_D(\boldsymbol{w}_{k-1} - \widehat{\boldsymbol{w}}_{k-1,E}^{(i)})$ prohibits us from deriving the improved rate, unless we also assume $Q_D(.)$ to be a Lipschitz operator.
In Appendix B, for $r \ll n$, we show that using the quantization scheme of Alistarh et al. [2017] with $s = \sqrt{d}$, FedGLOMO achieves more than a five-fold saving in the *total* communication cost as compared to when there is full-precision communication in FedGLOMO.

Remark 4 (**A limitation**). Even though our iteration complexity of $T = \mathcal{O}(\epsilon^{-1.5})$ is better than that of FedCOMGATE proposed by Haddadpour et al. [2021] (which is $\mathcal{O}(\epsilon^{-2})$), our communication complexity of $K = \mathcal{O}(\epsilon^{-1.5})$ is higher than that theirs which is $K = \mathcal{O}(\epsilon^{-1})$ (albeit under an extra assumption on the quantizer, namely Assumption 5 in their paper). Ideally, we would like to have $E = \mathcal{O}(\epsilon^{-p})$ and $K = \mathcal{O}(\epsilon^{-(1.5-p)})$ for some $p > 0$, in order to reduce FedGLOMO's communication complexity. Exploring whether such a result is obtainable with our proposed style of momentum is an interesting future direction.

6 EXPERIMENTS

To show the efficacy of *global* momentum in FedGLOMO, we compare it against FedLOMO (recall this has only local momentum and no global momentum; see Appendix F) and FedAvg (McMahan et al. [2017]) with the standard momentum available in PyTorch applied to (i) only its local updates, and (ii) both local and global updates – all with compressed client-to-server communication. We denote (i) and (ii) by FedAvg-lm and FedAvg-glm ("lm" and "glm" stand for local momentum, and global + local momentum), respectively. FedAvg *with compression* is referred to as FedPAQ (Reisizadeh et al. [2020]). Similarly, we call FedAvg-lm and FedAvg-glm *with compression*, as FedPAQ-lm and FedPAQ-glm. We also compare against FedCOMGATE (Haddadpour et al. [2021]) which uses gradient tracking to *theoretically* derive a better communication-complexity than us (see Remark 4). For compression, the "qsgd" operator proposed in Alistarh et al. [2017] is used.

We consider the task of classification on CIFAR-10 and Fashion-MNIST (Xiao et al. [2017]) abbreviated as FMNIST henceforth. The model used is a two-layer neural network with ReLU activation in the hidden layers. The size of both the hidden layers is 300/600 for FMNIST/CIFAR-10. We train the models using the categorical cross-entropy loss with ℓ_2-regularization. The weight decay value in PyTorch (to apply ℓ_2-regularization) is set to 1e-4. We consider both homogeneous and heterogeneous data distribution among the clients. Similar to McMahan et al. [2017], for the heterogeneous case, we distribute the data among the clients such that each client can have data from either one or (at most) two classes – note that this is a high degree of heterogeneity. The exact procedure is described in Appendix E. The number of clients (n) in all the experiments is set to 50, with each client having the same number of samples. The global batch-size r is 25, and the number of local updates per round (i.e., E) is 10. All full gradients are replaced by stochastic gradients computed on a (per-client) batch size of 256. The learning rates, momentum parameters of the algorithms, and some other experimental details are in Appendix E.

In Fig. 1, we compare `FedPAQ`-lm, `FedPAQ`-glm, `FedLOMO` and `FedCOMGATE` with 4 (resp., 8) bits per-round against `FedGLOMO` with 2 (resp., 4) bits per-round on FMNIST (resp., CIFAR-10) in the heterogeneous and homogeneous cases. We set the number of per-round bits used by `FedGLOMO` to be half the number used by all other algorithms, so that each one has the same *per-round* communication budget. All plots depict results over 3 independent runs; the shaded regions represent ± 1 standard deviation whereas the solid lines are the respective means. Please see the discussion in the figure caption. These results illustrate the *power of global momentum*.

Next, in the *no-compression heterogeneous* case, we compare against `Mime` (specifically, "`MimeSGDm`") of Karimireddy et al. [2020] which also attains a complexity of $\mathcal{O}(\epsilon^{-1.5})$ but without compressed communication, and is tailored to handle client heterogeneity. Having shown the suboptimality of `FedLOMO` and `FedPAQ`-lm in Fig. 1, we only compare `FedAvg`-glm, `FedGLOMO` without compression and `MimeSGDm` in the heterogeneous case in Fig. 2. The plots in Fig. 2 show that the implicit client-drift controlling ability of our proposed global momentum is on par with the explicit client-drift controlling mechanism of `Mime`. The test error values averaged over the last five rounds for the plots in Figs 1 and 2 are in Tables 2 and 3, respectively.

We also provide some more empirical results on CIFAR-100 in Appendix E.1.

Algo.	**CIFAR-10 Het.**	**FMNIST Het.**
`FedPAQ`-lm	50.26 ± 0.85	16.17 ± 0.53
`FedPAQ`-glm	49.88 ± 1.15	15.87 ± 1.10
`FedLOMO`	53.74 ± 0.17	18.95 ± 0.19
`FedGLOMO`	$\mathbf{46.42 \pm 0.05}$	$\mathbf{13.55 \pm 0.32}$
`FedCOMGATE`	$\mathbf{46.26 \pm 0.25}$	15.32 ± 0.09
Algo.	**CIFAR-10 Hom.**	**FMNIST Hom.**
`FedPAQ`-lm	$\mathbf{45.13 \pm 0.07}$	13.08 ± 0.05
`FedPAQ`-glm	45.70 ± 0.10	11.76 ± 0.06
`FedLOMO`	45.96 ± 0.01	14.22 ± 0.01
`FedGLOMO`	$\mathbf{44.97 \pm 0.05}$	$\mathbf{10.98 \pm 0.05}$
`FedCOMGATE`	45.46 ± 0.03	12.24 ± 0.01

Table 2: Average **test error** % ($\pm$ standard deviation) over the last five rounds for the plots in the *heterogeneous* (*top*) and *homogeneous* (*bottom*) cases in Figure 1.

Verifying Assumption 4 for `FedGLOMO`: For each round k, we compute $\alpha = \max_{\tau \in [E]} \frac{\| \sum_{i \in [n]} \widetilde{\boldsymbol{e}}_{k,\tau}^{(i)} \|^2}{\sum_{i \in [n]} \| \widetilde{\boldsymbol{e}}_{k,\tau}^{(i)} \|^2}$, where $\widetilde{\boldsymbol{e}}_{k,\tau}^{(i)}$ is as defined in Assumption 4, for 4 and 2 bit `FedGLOMO` on CIFAR-10 and FMNIST, respectively. Note that we remove the expectation (w.r.t. the stochastic gradients) while computing α for empirical verification. In Fig. 3, we plot (α/n) over different rounds for the heterogeneous as well as homogeneous case on both datasets; see the discussion in the figure caption.

Algo.	**CIFAR-10 Het.**	**FMNIST Het.**
`FedAvg`-glm	50.26 ± 0.74	16.17 ± 0.53
`MimeSGDm`	46.10 ± 0.13	$\mathbf{13.34 \pm 0.25}$
`FedGLOMO`	$\mathbf{45.41 \pm 0.15}$	$\mathbf{13.48 \pm 0.26}$

Table 3: Average **test error** % ($\pm$ standard deviation) over the last five rounds for the plots in Figure 2.

7 CONCLUSION

We presented `FedGLOMO`, a communication-efficient algorithm for faster federated learning via the application of variance-reducing momentum, both in the aggregation step at the server as well as local client updates. We showed that `FedGLOMO` has better iteration complexity than prior work on smooth non-convex functions with compressed communication. Further, unlike prior work, our result does not use the bounded client dissimilarity assumption, even holding under arbitrary client heterogeneity. We also demonstrate the efficacy of `FedGLOMO` via extensive experiments.

Acknowledgements

This work is supported by NSF grants CCF-1564000, IIS-1546452 and HDR-1934932, AFOSR grant FA9550-19-1-0005, and NASA grant 80NSSC21M0071.

References

Dan Alistarh, Demjan Grubic, Jerry Li, Ryota Tomioka, and Milan Vojnovic. Qsgd: Communication-efficient sgd via gradient quantization and encoding. In *Advances in Neural Information Processing Systems*, pages 1709–1720, 2017.

Dan Alistarh, Torsten Hoefler, Mikael Johansson, Nikola Konstantinov, Sarit Khirirat, and Cédric Renggli. The convergence of sparsified gradient methods. In *Advances in Neural Information Processing Systems*, pages 5973–5983, 2018.

Yossi Arjevani, Yair Carmon, John C Duchi, Dylan J Foster, Nathan Srebro, and Blake Woodworth. Lower bounds for non-convex stochastic optimization. *arXiv preprint arXiv:1912.02365*, 2019.

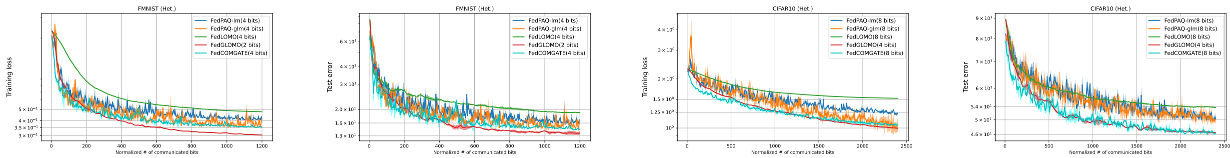

(a) Het. FMNIST train loss (b) Het. FMNIST test err (c) Het. CIFAR10 train loss (d) Het. CIFAR10 test err

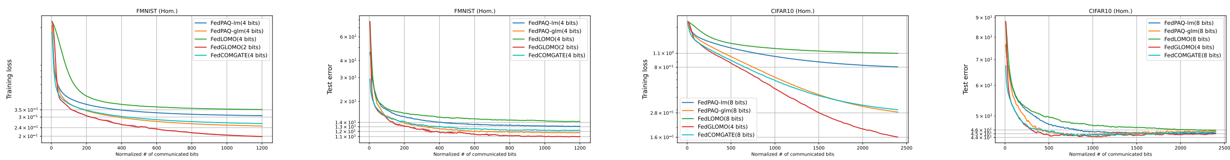

(e) Hom. FMNIST train loss (f) Hom. FMNIST test err (g) Hom. CIFAR10 train loss (h) Hom. CIFAR10 test err

Figure 1: Comparison of `FedPAQ`-lm, `FedPAQ`-glm, `FedLOMO`, `FedGLOMO` and `FedCOMGATE` (Haddadpour et al. [2021]) with the same per-round communication budget on FMNIST and CIFAR-10 in the heterogeneous (top four figs.) and homogeneous (bottom four figs.) settings, respectively. The x-axis is the total number of communicated bits divided by the dimension d and the global batch-size r. `FedGLOMO` is the **fastest** and most **communication-efficient** algorithm in almost all the cases; for e.g., in the heterogeneous case for both datasets, `FedGLOMO` attains the final test error of `FedPAQ`-glm (resp., `FedPAQ`-lm) with less than a **half** (resp., only about a **third**) of the number of bits used by `FedPAQ`-glm (resp., `FedPAQ`-lm). Further, `FedGLOMO` and `FedLOMO` have a smoother trajectory than other algorithms in the heterogeneous case due to variance-reducing momentum. Observe that `FedLOMO` and `FedPAQ`-lm (with only local momentum) are slower than `FedGLOMO` and `FedPAQ`-slm (with both local and global momentum), showing the ineffectiveness of only local momentum and **the power of combining both local and global momentum**. Also, note that `FedGLOMO` performs much better than `FedCOMGATE` in the homogeneous case.

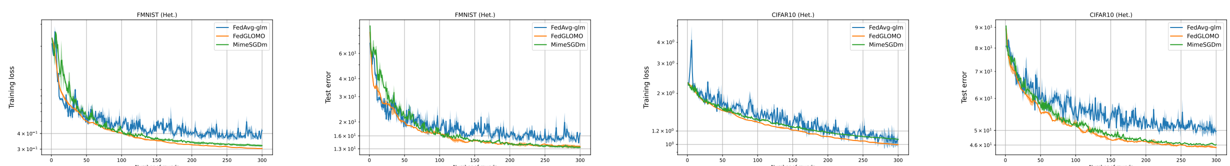

(a) FMNIST train loss (b) FMNIST test err (c) CIFAR-10 train loss (d) CIFAR-10 test err

Figure 2: Comparison of `FedAvg`-glm, `FedGLOMO` (without compression) and `MimeSGDm` on FMNIST and CIFAR-10 in the **heterogeneous** case. On both datasets, `FedAvg`-glm is the slowest while `FedGLOMO` is somewhat faster than `MimeSGDm`. While `Mime` has an explicit client-drift control mechanism, we do not have that in `FedGLOMO`, but still **our proposed global momentum implicitly mitigates client-drift** as well as `Mime`.

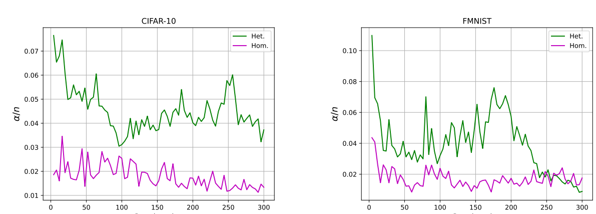

(a) CIFAR-10 (b) FMNIST

Figure 3: Variation of $(\frac{\alpha}{n})$ over different rounds of 4 and 2 bit `FedGLOMO` for CIFAR-10 (Fig. 3a) and FMNIST (Fig. 3b) in the heterogeneous and homogeneous cases. In both cases, notice that $\alpha \ll n$ throughout training. Also, as discussed after the statement of Assumption 4, note that $(\frac{\alpha}{n})$ is higher for the heterogeneous case (except at the end of training for FMNIST). See Figure 4 for the same on `FedAvg`.

Debraj Basu, Deepesh Data, Can Karakus, and Suhas Diggavi. Qsparse-local-sgd: Distributed sgd with quantization, sparsification and local computations. In *Advances in Neural Information Processing Systems*, pages 14695–

14706, 2019.

Ahmed Khaled Ragab Bayoumi, Konstantin Mishchenko, and Peter Richtárik. Tighter theory for local sgd on identical and heterogeneous data. In *International Conference on Artificial Intelligence and Statistics*, pages 4519–4529, 2020.

Jeremy Bernstein, Yu-Xiang Wang, Kamyar Azizzadenesheli, and Anima Anandkumar. signsgd: Compressed optimisation for non-convex problems. *arXiv preprint arXiv:1802.04434*, 2018.

Léon Bottou. Stochastic gradient descent tricks. In *Neural networks: Tricks of the trade*, pages 421–436. Springer, 2012.

Yiuye Chen, Abolfazl Hashemi, and Haris Vikalo. Communication-efficient algorithms for decentralized optimization over directed graphs. *arXiv preprint arXiv:2005.13189*, 2020.

Yiyue Chen, Abolfazl Hashemi, and Haris Vikalo. Communication-efficient variance-reduced decentralized stochastic optimization over time-varying directed graphs. *IEEE Transactions on Automatic Control*, 2021.

Ashok Cutkosky and Francesco Orabona. Momentum-based variance reduction in non-convex sgd. In *Advances in Neural Information Processing Systems*, pages 15236–15245, 2019.

Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. *Advances in neural information processing systems*, 27, 2014.

Cong Fang, Chris Junchi Li, Zhouchen Lin, and Tong Zhang. Spider: Near-optimal non-convex optimization via stochastic path-integrated differential estimator. In *Advances in Neural Information Processing Systems*, pages 689–699, 2018.

Eduard Gorbunov, Konstantin Burlachenko, Zhize Li, and Peter Richtárik. Marina: Faster non-convex distributed learning with compression. *arXiv preprint arXiv:2102.07845*, 2021.

Farzin Haddadpour, Mohammad Mahdi Kamani, Aryan Mokhtari, and Mehrdad Mahdavi. Federated learning with compression: Unified analysis and sharp guarantees. *arXiv preprint arXiv:2007.01154*, 2021.

Abolfazl Hashemi, Anish Acharya, Rudrajit Das, Haris Vikalo, Sujay Sanghavi, and Inderjit S Dhillon. On the benefits of multiple gossip steps in communication-constrained decentralized federated learning. *IEEE Transactions on Parallel and Distributed Systems*, 2021.

Samuel Horváth, Dmitry Kovalev, Konstantin Mishchenko, Sebastian Stich, and Peter Richtárik. Stochastic distributed learning with gradient quantization and variance reduction. *arXiv preprint arXiv:1904.05115*, 2019.

Zhouyuan Huo, Qian Yang, Bin Gu, Lawrence Carin Huang, et al. Faster on-device training using new federated momentum algorithm. *arXiv preprint arXiv:2002.02090*, 2020.

Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.

Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. *Advances in neural information processing systems*, 26, 2013.

Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank J Reddi, Sebastian U Stich, and Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning. *arXiv preprint arXiv:1910.06378*, 2019.

Sai Praneeth Karimireddy, Martin Jaggi, Satyen Kale, Mehryar Mohri, Sashank J Reddi, Sebastian U Stich, and Ananda Theertha Suresh. Mime: Mimicking centralized stochastic algorithms in federated learning. *arXiv preprint arXiv:2008.03606*, 2020.

Prashant Khanduri, Pranay Sharma, Haibo Yang, Mingyi Hong, Jia Liu, Ketan Rajawat, and Pramod Varshney. Stem: A stochastic two-sided momentum algorithm achieving near-optimal sample and communication complexities for federated learning. *Advances in Neural Information Processing Systems*, 34, 2021.

Anastasia Koloskova, Nicolas Loizou, Sadra Boreiri, Martin Jaggi, and Sebastian Stich. A unified theory of decentralized sgd with changing topology and local updates. In *International Conference on Machine Learning*, pages 5381–5393. PMLR, 2020.

Jaehoon Lee, Lechao Xiao, Samuel Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. *Advances in neural information processing systems*, 32, 2019.

Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *arXiv preprint arXiv:1812.06127*, 2018.

Xiang Li, Kaixuan Huang, Wenhao Yang, Shusen Wang, and Zhihua Zhang. On the convergence of fedavg on non-iid data. *arXiv preprint arXiv:1907.02189*, 2019.

Xianfeng Liang, Shuheng Shen, Jingchang Liu, Zhen Pan, Enhong Chen, and Yifei Cheng. Variance reduced local sgd with lower communication complexity. *arXiv preprint arXiv:1912.12844*, 2019.

Yujun Lin, Song Han, Huizi Mao, Yu Wang, and William J Dally. Deep gradient compression: Reducing the communication bandwidth for distributed training. *arXiv preprint arXiv:1712.01887*, 2017.

Deyi Liu, Lam M Nguyen, and Quoc Tran-Dinh. An optimal hybrid variance-reduced algorithm for stochastic composite nonconvex optimization. *arXiv preprint arXiv:2008.09055*, 2020.

Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*, pages 1273–1282. PMLR, 2017.

Fangzhou Mu, Yingyu Liang, and Yin Li. Gradients as features for deep representation learning. *arXiv preprint arXiv:2004.05529*, 2020.

Lam M Nguyen, Jie Liu, Katya Scheinberg, and Martin Takáč. Sarah: A novel method for machine learning problems using stochastic recursive gradient. In *International Conference on Machine Learning*, pages 2613–2621. PMLR, 2017.

Kumar Kshitij Patel and Aymeric Dieuleveut. Communication trade-offs for synchronized distributed sgd with large step size. *arXiv preprint arXiv:1904.11325*, 2019.

Shi Pu and Angelia Nedić. Distributed stochastic gradient tracking methods. *Mathematical Programming*, pages 1–49, 2020.

Zhaonan Qu, Kaixiang Lin, Jayant Kalagnanam, Zhaojian Li, Jiayu Zhou, and Zhengyuan Zhou. Federated learning's blessing: Fedavg has linear speedup. *arXiv preprint arXiv:2007.05690*, 2020.

Sashank Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečnỳ, Sanjiv Kumar, and H Brendan McMahan. Adaptive federated optimization. *arXiv preprint arXiv:2003.00295*, 2020.

Amirhossein Reisizadeh, Aryan Mokhtari, Hamed Hassani, Ali Jadbabaie, and Ramtin Pedarsani. Fedpaq: A communication-efficient federated learning method with periodic averaging and quantization. In *International Conference on Artificial Intelligence and Statistics*, pages 2021–2031, 2020.

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

Sebastian U Stich and Sai Praneeth Karimireddy. The error-feedback framework: Better rates for sgd with delayed gradients and compressed communication. *arXiv preprint arXiv:1909.05350*, 2019.

Sebastian U Stich, Jean-Baptiste Cordonnier, and Martin Jaggi. Sparsified sgd with memory. In *Advances in Neural Information Processing Systems*, pages 4447–4458, 2018.

Ananda Theertha Suresh, X Yu Felix, Sanjiv Kumar, and H Brendan McMahan. Distributed mean estimation with limited communication. In *International Conference on Machine Learning*, pages 3329–3337, 2017.

Hanlin Tang, Shaoduo Gan, Ce Zhang, Tong Zhang, and Ji Liu. Communication compression for decentralized training. In *Advances in Neural Information Processing Systems*, pages 7652–7662, 2018.

Jianyu Wang and Gauri Joshi. Cooperative sgd: A unified framework for the design and analysis of communication-efficient sgd algorithms. *arXiv preprint arXiv:1808.07576*, 2018.

Jianyu Wang, Vinayak Tantia, Nicolas Ballas, and Michael Rabbat. Slowmo: Improving communication-efficient distributed sgd with slow momentum. *arXiv preprint arXiv:1910.00643*, 2019.

Blake Woodworth, Kumar Kshitij Patel, Sebastian U Stich, Zhen Dai, Brian Bullins, H Brendan McMahan, Ohad Shamir, and Nathan Srebro. Is local sgd better than minibatch sgd? *arXiv preprint arXiv:2002.07839*, 2020.

Jiaxiang Wu, Weidong Huang, Junzhou Huang, and Tong Zhang. Error compensated quantized sgd and its applications to large-scale distributed optimization. *arXiv preprint arXiv:1806.08054*, 2018.

Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.

Hao Yu, Sen Yang, and Shenghuo Zhu. Parallel restarted sgd for non-convex optimization with faster convergence and less communication. *arXiv preprint arXiv:1807.06629*, 2 (4):7, 2018.

Dongruo Zhou, Pan Xu, and Quanquan Gu. Stochastic nested variance reduced gradient descent for nonconvex optimization. *Advances in neural information processing systems*, 2018.

Martin Zinkevich, Markus Weimer, Lihong Li, and Alex Smola. Parallelized stochastic gradient descent. *Advances in neural information processing systems*, 23:2595–2603, 2010.

Appendix

Contents

A FULL STATEMENT OF THEOREM 1 AND PROOF OUTLINE

Here we present the full version of Theorem 1 as Theorem 2. The detailed proof of this result can be found in Appendix G.1. We also provide a brief proof outline after the theorem statement.

Theorem 2 (**Expanded version of Theorem 1**). *Let Assumptions 1, 2 and 3 hold. Further, suppose Assumption 4 is true for* `FedGLOMO`*. In* `FedGLOMO`*, for each round k, set:*

$$\eta_k = \eta = \frac{1}{6LEK^{1/3}(\frac{1}{n}(\alpha + \frac{4}{E}) + 800e^2(1+q)(E+1)^2(\frac{q}{n} + \frac{(1+q)(n-r)}{r(n-1)}))^{1/3}} \text{ and}$$

$$\beta_k = \beta = 160e^2(1+q)\eta^2 L^2 E^2 (E+1)^2.$$

Suppose we use full-device participation (i.e., the global batch size is n) ***only at*** *$k = 0$. Then if $\frac{K^{-1/3}}{1200e^2(1+q)\left(\frac{q}{n} + \frac{(1+q)(n-r)}{r(n-1)}\right)} \leq E + 1 \leq \frac{\sqrt{1+q}(n-r)}{3r(n-1)} K$, we have:*

$$\frac{1}{K}\sum_{k=0}^{K-1} \mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] \leq \frac{39Lf(\boldsymbol{w}_0)}{K^{2/3}}\Big(\frac{1}{n}\Big(\alpha + \frac{4}{E}\Big) + 800e^2(1+q)(E+1)^2\Big(\frac{q}{n} + \frac{(1+q)(n-r)}{r(n-1)}\Big)\Big)^{1/3}.$$

Thus, `FedGLOMO` *can achieve $\mathbb{E}_{k^* \sim \text{Unif}[0,K-1]}[\|\nabla f(\boldsymbol{w}_{k^*})\|^2] \leq \epsilon$ in $K = \mathcal{O}\Big(\max\Big(\sqrt{\frac{\alpha}{n}}, (1+q)\sqrt{\frac{(n-r)}{r(n-1)}}\Big)\epsilon^{-1.5}\Big)$ rounds of communication and $E = \mathcal{O}(1)$ local steps.*

Note that the above result is independent of the variance of local stochastic gradients (of the clients). In short, this happens because we use local full gradients at $\tau = 0$ and because the local stochastic gradients are Lipschitz.

PROOF OUTLINE:

Before getting to the proof outline, we would like to mention that the key technical challenge in deriving the improved convergence result with global momentum-based variance reduction is obtaining an analogue of the Lipschitzness of stochastic gradients to the change in local parameters over E local steps. More specifically, for pure stochastic optimization, a key step in proving convergence of momentum-based variance reduction methods is using the Lipschitzness of the stochastic gradients or the update quantities (see Cutkosky and Orabona [2019], Liu et al. [2020]), i.e.,

$$\|\nabla \widetilde{f}(\boldsymbol{x}_t, \xi_t) - \nabla \widetilde{f}(\boldsymbol{x}_{t-1}, \xi_t)\| \leq L\|\boldsymbol{x}_t - \boldsymbol{x}_{t-1}\|.$$

In the FL setting where aggregation is performed at the server, we need an analogue of this at the server, i.e., something like

$$\|(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)}) - (\boldsymbol{w}_{k-1} - \widehat{\boldsymbol{w}}_{k-1,E}^{(i)})\| \leq \widetilde{L}\|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\|.$$

Deriving this result is a part of our contribution and is done in Lemma 9 (in Appendix G.1).

Proof. We set $\eta_k = \eta$ and $\beta_k = \beta \; \forall \; k \in \{0, \ldots, K-1\}$. Then, using Lemma 1 with full global as well as local batch sizes at $k = 0$ (by which $\boldsymbol{u}_0 = \overline{\boldsymbol{\delta}}_0$ in the statement of Lemma 1), we have at any $k' > 0$:

$$\mathbb{E}[f(\boldsymbol{w}_{k'})] \leq f(\boldsymbol{w}_0) - \frac{\eta E}{4}\sum_{k=0}^{k'-1}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \frac{16\eta^3 L^2 E^2(\alpha E + 4)}{n^2}\sum_{k=0}^{k'-1}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]$$
$$+ 160\eta E\beta\Big(\frac{q}{n^2} + \frac{(1+q)}{r(n-1)}\Big(1 - \frac{r}{n}\Big)\Big)\sum_{k=0}^{k'-1}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2], \quad (6)$$

for $4\eta LE^2 \leq 1$ and $\beta \geq \frac{80e^2(1+q)\eta^2 L^2 E^2 (E+1)^2}{(1-4\eta LE)}$.

Also, since the f_i's are L-smooth and non-negative, using Lemma 11, we have that:

$$\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2 \leq \sum_{i\in[n]} 2L(\mathbb{E}[f_i(\boldsymbol{w}_k)] - f_i^*) \leq 2nL\mathbb{E}[f(\boldsymbol{w}_k)] - 2L\sum_{i\in[n]} f_i^* \leq 2nL\mathbb{E}[f(\boldsymbol{w}_k)].$$

This step allows us to circumvent the need for the bounded client dissimilarity assumption. Using this in (6), we get:

$$\mathbb{E}[f(\boldsymbol{w}_{k'})] \leq f(\boldsymbol{w}_0) - \frac{\eta E}{4}\sum_{k=0}^{k'-1}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2]$$
$$+ \underbrace{64\eta LE\Big(\frac{\eta^2 L^2 E(\alpha E + 4)}{n} + 5\beta\Big(\frac{q}{n} + \frac{(1+q)(n-r)}{r(n-1)}\Big)\Big)}_{=\gamma}\sum_{k=0}^{k'-1}\mathbb{E}[f(\boldsymbol{w}_k)]. \quad (7)$$

Unfolding the above recursion and simplifying a bit, we get:

$$\sum_{k=0}^{k'-1}\mathbb{E}[f(\boldsymbol{w}_k)] \leq k' f(\boldsymbol{w}_0) - \frac{\eta E}{4}\sum_{k=0}^{k'-1}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \gamma k'\sum_{k=0}^{k'-1}\mathbb{E}[f(\boldsymbol{w}_k)]. \quad (8)$$

Let us now ensure that $\gamma k' \leq \frac{1}{2}$ for all $k' \in \{1, \ldots, K\}$, so that we can simplify (8) to:

$$\sum_{k=0}^{k'-1}\mathbb{E}[f(\boldsymbol{w}_k)] \leq 2k' f(\boldsymbol{w}_0) - \frac{\eta E}{2}\sum_{k=0}^{k'-1}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] \leq 2k' f(\boldsymbol{w}_0). \quad (9)$$

Now for $8\eta LE^2 \leq 1$, it can be verified that $\beta = 160e^2(1+q)\eta^2 L^2 E^2 (E+1)^2$ is a valid choice. Using this, we get that:

$$\gamma k' \leq \gamma K = \underbrace{64\eta^3 L^3 E^3 K\Big(\frac{1}{n}\Big(\alpha + \frac{4}{E}\Big) + 800e^2(1+q)(E+1)^2\Big(\frac{q}{n} + \frac{(1+q)(n-r)}{r(n-1)}\Big)\Big)}_{\text{(A)}}. \quad (10)$$

Setting

$$\eta = \frac{1}{6LEK^{1/3}(\frac{1}{n}(\alpha + \frac{4}{E}) + 800e^2(1+q)(E+1)^2(\frac{q}{n} + \frac{(1+q)(n-r)}{r(n-1)}))^{1/3}},$$

we have (A) $< \frac{1}{2}$. We also need to ensure that $8\eta LE^2 \leq 1$ and $\beta = 160e^2(1+q)\eta^2 L^2 E^2 (E+1)^2 < 1$. The range of E in the theorem statement is obtained by combining the constraints (on E) that we get from these two requirements.

Finally, using (9) in (7) with $k' = K$, substituting our choice of η and β, and then simplifying a bit more, we get the final convergence result. ■

B REDUCTION IN TOTAL COMMUNICATION COST WHEN $r \ll n$

Here, we derive the claim made in the second paragraph of Remark 3.

We consider the practical regime of $r \ll n$ (as well as $\alpha \ll n$).

First, consider the case where the clients communicate at full precision using 32 bits, i.e., $q = 0$. The number of rounds of communication K_1 needed to reach an ϵ stationary point is:

$$K_1 \approx \Big(\frac{39Lf(\boldsymbol{w}_0)}{\epsilon}\Big)^{1.5}\Big(800e^2(E+1)^2\frac{(n-r)}{r(n-1)}\Big)^{0.5}$$

Since the communication cost per-round is proportional to $r \times (32d)$ bits (recall d is the model dimension), the total communication cost C_1 in this case is:

$$\begin{aligned} C_1 &= 32dr \times K_1 \\ &= 32dr\Big(\frac{39Lf(\boldsymbol{w}_0)}{\epsilon}\Big)^{1.5}\Big(800e^2(E+1)^2\frac{(n-r)}{r(n-1)}\Big)^{0.5}. \end{aligned}$$

Now, let us consider the QSGD quantizer of Alistarh et al. [2017] with $s = \sqrt{d}$ [2]. With this choice, $q = 1$. Here, the number of rounds of communication K_2 needed to reach an ϵ stationary point is:

$$K_2 \approx \Big(\frac{39Lf(\boldsymbol{w}_0)}{\epsilon}\Big)^{1.5}\Big(3200e^2(E+1)^2\frac{(n-r)}{r(n-1)}\Big)^{0.5}.$$

Now using Theorem 3.4 of Alistarh et al. [2017], under the special case of $s = \sqrt{d}$, the communication cost per-round can be reduced to $r \times (2.8d + 32)$ bits. Hence, the total communication cost C_2 in this case is:

$$\begin{aligned} C_2 &\approx (2.8d + 32)r \times K_2 \\ &\approx (2.8d + 32)r\Big(\frac{39Lf(\boldsymbol{w}_0)}{\epsilon}\Big)^{1.5}\Big(3200e^2(E+1)^2\frac{(n-r)}{r(n-1)}\Big)^{0.5}. \end{aligned}$$

Therefore,

$$\frac{C_1}{C_2} \approx \frac{32}{2.8 \times 4^{0.5}} \approx 5.7.$$

C ALGORITHMIC AND THEORETICAL COMPARISON WITH MIME (Karimireddy et al. [2020])

We now discuss the major algorithmic and theoretical differences of our work from Karimireddy et al. [2020].

- Algorithmically, Karimireddy et al. [2020] do not explicitly *apply* any momentum at the server. Instead, they apply globally computed momentum in the local updates of the clients. On the other hand, `FedGLOMO` has an explicit momentum-based update at the server to enable global variance reduction, apart from local momentum applied in the client updates.

[2] Alistarh et al. [2017] use n to denote the dimension

- Unlike `FedGLOMO`, the algorithms of Karimireddy et al. [2020] do not have any quantized/compressed communication. As we discussed in Remark 3, maintaining the improved complexity of $\mathcal{O}(\epsilon^{-1.5})$ with compressed communication is not trivial.
- Even in the absence of any compressed communication, `FedGLOMO` is more communication-efficient than Mime requiring three-fourth / half the number of bits that Mime requires per-round for server to clients as well as clients to server communication / only clients to server communication (which is typically the bottleneck in FL). This is because in Mime, the server needs to send $\boldsymbol{x}$ (sending some other statistics $\boldsymbol{s}$ would require even more bits) and $\boldsymbol{c}$ to the clients, and the clients need to send back $(\boldsymbol{y}_i, \nabla f_i(\boldsymbol{x}))$ to the server (please see their notation). In `FedGLOMO`, the server needs to send $\boldsymbol{w}_k$ and $\boldsymbol{w}_{k-1}$ to the clients, but the clients can just send back $\{(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)}) - (1-\beta_k)(\boldsymbol{w}_{k-1} - \widehat{\boldsymbol{w}}_{k-1,E}^{(i)})\}$ to the server in the absence of any quantization – this can be verified by just removing the quantization operator Q_D and expanding the update rule of $\boldsymbol{u}_k$ (line 10 of Algorithm 1) for $k > 0$.
- Our theory for `FedGLOMO` does not use the bounded client dissimilarity (BCD) assumption, i.e., eq. (2). Instead, we propose and use Assumption 4, which allows for arbitrary client heterogeneity; in the worst case, Assumption 4 will hold with $\alpha = n$. In contrast, the results of MimeMVR use the BCD assumption.
- See the full version of Theorem V (on page 39 of the latest arXiv draft) of Karimireddy et al. [2020] for MimeMVR. Their result is in terms of the gradient of f at the local client parameters and not the actual server parameters, which is not ideal. Our result for `FedGLOMO` is completely in terms of the gradient of f at the server parameters.

D COMPARISON WITH Gorbunov et al. [2021]

As mentioned in Section 2, Gorbunov et al. [2021] also propose algorithms with improved complexity in the *distributed setting without multiple local update steps*. Since our work is under partial-device participation, we compare against their algorithm for the same case, i.e. PP-MARINA. Note that PP-MARINA has a probability p of using *full gradients from all clients* (i.e., full device participation) in each iteration. For a fair comparison against `FedGLOMO`, which uses gradients from all the clients *only in the first round* (see Theorem 1), p should be set to $\frac{1}{K}$ (K being the number of rounds) – in which case, their complexity is $O(\epsilon^{-2})$ which is worse than ours. See Theorem 4.1 in their paper for this.

E EXPERIMENTAL DETAILS AND SOME MORE RESULTS

We first describe the procedure we have used to generate heterogeneous data distribution (among the clients). First, the training data (of both CIFAR-10 and FMNIST) was sorted based on labels and then divided into 100 equal data-shards. Splitting the data into 100 equal shards (after sorting) ensures that each shard contains data from only one class for both CIFAR-10 and FMNIST. Since the number of clients in our experiments is fixed to 50, each client is assigned 2 shards chosen uniformly at random without replacement – this ensures that each client can have data belonging to either just one class or two classes at the most. For the homogeneous case, we distribute the training data uniformly at random among the clients.

In all our experiments, we use the learning rate schedule suggested in Bottou [2012] where we decimate the client learning rate by 1% after every round, i.e., $\eta_k = (0.99)^k \eta_0$, η_0 being the initial learning rate. Note that this learning rate schedule has been used earlier for FL experiments in Haddadpour et al. [2021]. We search the initial learning rates over $\{10^{-3}, 5 \times 10^{-3}, 10^{-2}, 5 \times 10^{-2}, 10^{-1}\}$; the best performance is obtained with an initial learning rate of 10^{-2} in almost all the cases.

For `FedGLOMO`, we use a constant value of $\beta_k = 0.2$. For `FedAvg`-lm, `FedAvg`-glm, `FedPAQ`-lm and `FedPAQ`-glm, the local (client-level) momentum parameter (in the Pytorch optimizer) is set to 0.9 (which is also the default value used). For `FedAvg`-glm and `FedPAQ`-glm, we just implement the server update as a PyTorch optimizer update with momentum, and this momentum parameter is searched over $\{0.9, 0.7, 0.5\}$. For `MimeSGDm`, we search its momentum hyper-parameter over $\{0, 0.9, 0.99\}$ as suggested in Karimireddy et al. [2020].

We make a small modification to `FedGLOMO` in our experiments for the heterogeneous case. Specifically, we modify line 6 (which is the local momentum application step) of Algorithm 2 as follows:

$$\text{Update: } \boldsymbol{v}_{k,\tau}^{(i)} = \widetilde{\nabla} f_i(\boldsymbol{w}_{k,\tau}^{(i)}; \mathcal{B}_{k,\tau}^{(i)}) + 0.8\big(\boldsymbol{v}_{k,\tau-1}^{(i)} - \widetilde{\nabla} f_i(\boldsymbol{w}_{k,\tau-1}^{(i)}; \mathcal{B}_{k,\tau}^{(i)})\big) \text{ and}$$

$$\widehat{\boldsymbol{v}}_{k-1,\tau}^{(i)} = \widetilde{\nabla} f_i(\widehat{\boldsymbol{w}}_{k-1,\tau}^{(i)}; \mathcal{B}_{k,\tau}^{(i)}) + 0.8\big(\widehat{\boldsymbol{v}}_{k-1,\tau-1}^{(i)} - \widetilde{\nabla} f_i(\widehat{\boldsymbol{w}}_{k-1,\tau-1}^{(i)}; \mathcal{B}_{k,\tau}^{(i)})\big).$$

Without applying the above damping factor of 0.8, FedGLOMO seems to diverge – this is probably because we have chosen the number of local updates to be too large.

All experiments are run on a single NVIDIA TITAN Xp GPU.

E.1 RESULTS ON CIFAR-100

We consider the task of classification on CIFAR-100 (100 classes). We use 512-dimensional features extracted from the last layer of a ResNet-34 model pretrained on ImageNet. The architecture is a two-layered ReLU neural network with the size of hidden layers being 300 and 150, respectively. As we have in the main paper, the total number of clients is 50 and 50% of the clients participate in each round (i.e., $r = 0.5n$). Also, we consider heterogeneous and homogeneous (i.i.d.) settings similar to the main paper; in the heterogeneous setting, each client has data from 2 classes. The number of rounds and the number of local steps per round are 50 and 20, respectively. We compare 3 different runs of FedPAQ-lm, FedPAQ-glm, FedLOMO and FedCOMGATE with 8 bits per round against FedGLOMO with 4 bits per round (so that the total number of bits communicated per round of all algorithms is the same). All other details (about learning rates, momentum parameters, etc.) are the same as in the main paper. The test error (error = 100 - accuracy) values averaged over the last five rounds for this case is listed in Table 4.

Algo.	**CIFAR-100 Het.**	**CIFAR-100 Hom.**
FedPAQ-lm	31.36 ± 0.05	31.02 ± 0.12
FedPAQ-glm	31.76 ± 0.09	30.78 ± 0.16
FedLOMO	31.26 ± 0.16	30.73 ± 0.05
FedGLOMO	$\mathbf{30.76 \pm 0.10}$	$\mathbf{30.05 \pm 0.06}$
FedCOMGATE	34.27 ± 0.09	33.59 ± 0.09

Table 4: Average **test error** % ($\pm$ standard deviation) over the last five rounds for CIFAR-100.

In both the heterogeneous and homogeneous settings, notice that FedGLOMO converges to the lowest test error; this is consistent with the results in the main paper on CIFAR-10 and FMNIST.

F FEDLOMO: A SIMPLER VERSION OF FEDGLOMO

Now we consider a simpler version of FedGLOMO, which we call FedLOMO, that applies only local momentum in the client updates and does simple averaging at the server (like FedAvg), i.e., there is no global momentum (and hence the name of this variant does not have a "G"). FedLOMO is summarized in Algorithm 4 and 5. Notice that the momentum application occurs in line 6 of Algorithm 5.

As mentioned in the main paper, FedLOMO does not achieve the optimal convergence rate for smooth non-convex functions due to the absence of global momentum; see Theorem 3 and the subsequent remarks.

Just like the results of FedGLOMO, we do not use the BCD assumption (i.e., eq. (2)) to derive the results of FedLOMO.

Algorithm 4 FedLOMO - Server Update

Input: Initial point $\boldsymbol{w}_0$, # of rounds of communication K, period E, learning rates $\{\eta_k\}_{k=0}^{K-1}$, per-client batch size b, and global batch size r. Q_D is the quantization operator.
for $k = 0, \ldots, K-1$ **do**
    Server chooses a set $\mathcal{S}_k$ of r clients uniformly at random without replacement and sends $\boldsymbol{w}_k$ to them.
    for client $i \in \mathcal{S}_k$ **do**
        Set $\boldsymbol{w}_{k,0}^{(i)} = \boldsymbol{w}_k$ and run Algorithm 5 for client i.
    end for
    Update $\boldsymbol{w}_{k+1} = \boldsymbol{w}_k + \frac{1}{r} \sum_{i \in \mathcal{S}_k} Q_D(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k)$.
end for

Algorithm 5 `FedLOMO` - Client Update

for $\tau = 0, \ldots, E-1$ **do**
  if $\tau = 0$ **then**
    $\boldsymbol{v}_{k,\tau}^{(i)} = \nabla f_i(\boldsymbol{w}_{k,\tau}^{(i)})$.
  else
    Pick a random batch of b samples in client i, say $\mathcal{B}_{k,\tau}^{(i)}$. Compute the stochastic gradients of f_i at $\boldsymbol{w}_{k,\tau}^{(i)}$ and $\boldsymbol{w}_{k,\tau-1}^{(i)}$ over $\mathcal{B}_{k,\tau}^{(i)}$ viz. $\widetilde{\nabla} f_i(\boldsymbol{w}_{k,\tau}^{(i)}; \mathcal{B}_{k,\tau}^{(i)})$ and $\widetilde{\nabla} f_i(\boldsymbol{w}_{k,\tau-1}^{(i)}; \mathcal{B}_{k,\tau}^{(i)})$, respectively.
    Update $\boldsymbol{v}_{k,\tau}^{(i)} = \widetilde{\nabla} f_i(\boldsymbol{w}_{k,\tau}^{(i)}; \mathcal{B}_{k,\tau}^{(i)}) + \big(\boldsymbol{v}_{k,\tau-1}^{(i)} - \widetilde{\nabla} f_i(\boldsymbol{w}_{k,\tau-1}^{(i)}; \mathcal{B}_{k,\tau}^{(i)})\big)$. // `(Local Momentum)`
  end if
  Update $\boldsymbol{w}_{k,\tau+1}^{(i)} = \boldsymbol{w}_{k,\tau}^{(i)} - \eta_k \boldsymbol{v}_{k,\tau}^{(i)}$.
end for
Send $Q_D(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k)$ to the server.

F.1 MAIN RESULT FOR FEDLOMO

Now, we present the convergence result of `FedLOMO` for the smooth non-convex case in Theorem 3. Its proof is in Appendix G.2.

Theorem 3 (Smooth non-convex case for `FedLOMO`). *Let Assumptions 1, 2 and 3 hold. Further, suppose Assumption 4 holds for `FedLOMO`. Define a distribution $\mathbb{P}$ for $k \in \{0, \ldots, K-1\}$ such that $\mathbb{P}(k) = \frac{(1+\zeta)^{(K-1-k)}}{\sum_{k=0}^{K-1}(1+\zeta)^k}$ where ζ will be defined later. Sample k^* from $\mathbb{P}$.*

***1. With compression** ($q > 0$) **and partial-device participation** ($r < n$):*

In `FedLOMO`, set $\eta_k = \frac{1}{8LE\sqrt{BK}}$ for all k, where $B = \frac{q}{n} + \frac{4(1+q)(n-r)}{r(n-1)}$. Then for $K > \frac{1}{64B^3}\big(\frac{1}{n}(\alpha + \frac{4}{E})\big)$:

$$\mathbb{E}[\|\nabla f(\boldsymbol{w}_{k^*})\|^2] \leq \frac{64\sqrt{B}Lf(\boldsymbol{w}_0)}{K^{1/2}} \text{ with } \zeta := \frac{1}{4K} + \frac{1}{16(BK)^{1.5}}\Big(\frac{1}{n}\Big(\alpha + \frac{4}{E}\Big)\Big).$$

So `FedLOMO` needs $K = \mathcal{O}(\frac{1}{r\epsilon^2})$ rounds of communication to achieve $\mathbb{E}[\|\nabla f(\boldsymbol{w}_{k^*})\|^2] \leq \epsilon$, for $\epsilon < \mathcal{O}(\frac{nB^2}{\alpha}) = \mathcal{O}(\frac{n/\alpha}{r^2})$.

***2. No compression** ($q = 0$) **and full-device participation** ($r = n$):*

In `FedLOMO`, set $\eta_k = \frac{1}{4LE}\Big(\frac{n}{(\alpha+\frac{4}{E})K}\Big)^{1/3}$ for all k. Then for $K > \Big(\frac{n}{\alpha+\frac{4}{E}}\Big)$:

$$\mathbb{E}[\|\nabla f(\boldsymbol{w}_{k^*})\|^2] \leq \frac{16Lf(\boldsymbol{w}_0)}{K^{2/3}}\left(\frac{\alpha + \frac{4}{E}}{n}\right)^{1/3}.$$

So `FedLOMO` needs $K = \mathcal{O}(\frac{1}{\epsilon^{1.5}}\sqrt{\frac{\alpha}{n}})$ rounds of communication to achieve $\mathbb{E}[\|\nabla f(\boldsymbol{w}_{k^*})\|^2] \leq \epsilon$, for $\epsilon < \mathcal{O}(\frac{\alpha}{n})$.

We make some remarks to discuss implications of this result and establish connections to some claims made in the main paper.

Remark 5 (Worse iteration complexity than `FedGLOMO` under compression and partial-device participation). Under compression (i.e., $q > 0$) and partial-device participation (i.e., $r < n$), the number of iterations $T = KE$ of `FedLOMO` is $\mathcal{O}(\frac{1}{r\epsilon^2})$ as per the above theorem (since we do not have any constraint on E depending on ϵ). So the iteration complexity of `FedLOMO` is poorer than that of `FedGLOMO` under the same setting. However, it is on a par with the results of Haddadpour et al. [2021], Koloskova et al. [2020], Wang et al. [2019], Karimireddy et al. [2019].

Remark 6 (Same iteration complexity as `FedGLOMO` under no compression and full-device participation). When there is no compression (i.e., $q = 0$) and full-device participation (i.e., $r = n$), the number of iterations $T = KE$ of `FedLOMO` turns out to be $\mathcal{O}(\frac{1}{\epsilon^{1.5}}\sqrt{\frac{\alpha}{n}})$; this is exactly the same as that of `FedGLOMO` under the same setting. Intuitively this makes sense because under full-device participation (and no compression), the global momentum of `FedGLOMO` becomes redundant.

Remark 7 (High variance of simple averaging at the server). At a high level, `FedLOMO` fails to attain the complexity of `FedGLOMO` under partial-device participation and compression because of the high variance of the `FedAvg`-like plain averaging step at the server. The high variance is itself due to the amplified effect of client heterogeneity with multiple local updates. Without the application of some *global variance-reduction* technique (like the one in `FedGLOMO`), the complexity cannot be improved. More precisely, in Theorem 3, B is a constant that is not $\mathcal{O}(\eta LE)$ in general, due to which `FedLOMO` does not achieve the improved convergence rate of $\mathcal{O}(K^{-2/3})$ that `FedGLOMO` attains; see the proof of Theorem 3 in Appendix G.2 for more details. However, in the case of full-device participation and no compression, B is 0 which allows `FedLOMO` to also achieve $\mathcal{O}(K^{-2/3})$ convergence by choosing $\eta = \mathcal{O}(\frac{1}{LEK^{1/3}})$.

G DETAILED PROOFS

G.1 DETAILED PROOF OF THE RESULT OF FEDGLOMO

Some definitions used in the proofs:

$$\boldsymbol{\delta}_k^{(i)} \triangleq \mathbb{E}_{\mathcal{B}_1^{(i)},\ldots,\mathcal{B}_{E-1}^{(i)}}[\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)}] \text{ for any } E-1 \text{ batches } \{\mathcal{B}_1^{(i)},\ldots,\mathcal{B}_{E-1}^{(i)}\} \text{ in client } i, \text{ and } \overline{\boldsymbol{\delta}}_k \triangleq \frac{1}{n}\sum_{i\in[n]}\boldsymbol{\delta}_k^{(i)}.$$

$$g_Q(\boldsymbol{w}_k;\mathcal{S}_k) \triangleq \frac{1}{r}\sum_{i\in\mathcal{S}_k} Q_D(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)})$$

$$\Delta g_Q(\boldsymbol{w}_k,\boldsymbol{w}_{k-1};\mathcal{S}_k) \triangleq \frac{1}{r}\sum_{i\in\mathcal{S}_k} Q_D((\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)}) - (\boldsymbol{w}_{k-1} - \widehat{\boldsymbol{w}}_{k-1,E}^{(i)}))$$

$$g(\boldsymbol{w}_k;\mathcal{S}_k) \triangleq \frac{1}{r}\sum_{i\in\mathcal{S}_k}(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)}) = \mathbb{E}_{Q_D}[g_Q(\boldsymbol{w}_k;\mathcal{S}_k)]$$

$$\widehat{g}(\boldsymbol{w}_{k-1};\mathcal{S}_k) \triangleq \frac{1}{r}\sum_{i\in\mathcal{S}_k}(\boldsymbol{w}_{k-1} - \widehat{\boldsymbol{w}}_{k-1,E}^{(i)})$$

$$\overline{\boldsymbol{w}}_{k,\tau} \triangleq \frac{1}{n}\sum_{i\in[n]}\boldsymbol{w}_{k,\tau}^{(i)} \text{ and } \overline{\boldsymbol{v}}_{k,\tau} \triangleq \frac{1}{n}\sum_{i\in[n]}\boldsymbol{v}_{k,\tau}^{(i)}$$

$$\boldsymbol{e}_{k,\tau}^{(i)} \triangleq \boldsymbol{v}_{k,\tau}^{(i)} - \nabla f_i(\boldsymbol{w}_{k,\tau}^{(i)}) \text{ and } \widetilde{\boldsymbol{e}}_{k,\tau}^{(i)} \triangleq \nabla f_i(\boldsymbol{w}_{k,\tau}^{(i)}) - \nabla f_i(\overline{\boldsymbol{w}}_{k,\tau})$$

Proof of Theorem 2 (recall that this is the full version of Theorem 1):

Proof. We set $\eta_k = \eta$ and $\beta_k = \beta\ \forall\ k \in \{0,\ldots,K-1\}$.
Then using Lemma 1, we have that:

$$\mathbb{E}[f(\boldsymbol{w}_{k'})] \leq f(\boldsymbol{w}_0) - \frac{\eta E}{4}\sum_{k=0}^{k'-1}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \frac{16\eta^3L^2E^2(\alpha E+4)}{n^2}\sum_{k=0}^{k'-1}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]$$
$$+ \frac{5}{4\eta E\beta}\mathbb{E}[\|\boldsymbol{u}_0 - \overline{\boldsymbol{\delta}}_0\|^2] + 160\eta E\beta\Big(\frac{q}{n^2} + \frac{(1+q)}{r(n-1)}\Big(1-\frac{r}{n}\Big)\Big)\sum_{k=0}^{k'-1}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2], \quad (11)$$

for $4\eta LE^2 \leq 1$ and $\beta \geq \frac{80e^2(1+q)\eta^2L^2E^2(E+1)^2}{(1-4\eta LE)}$.

Suppose we use full batch sizes for the local updates as well as the server update at $k=0$ (the latter means $r=n$ only for $k=0$). Then, $\boldsymbol{u}_0 = \overline{\boldsymbol{\delta}}_0$ above. Also, since the f_i's are L-smooth, using Lemma 11, we have that:

$$\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2 \leq \sum_{i\in[n]}2L(\mathbb{E}[f_i(\boldsymbol{w}_k)] - f_i^*) \leq 2nL\mathbb{E}[f(\boldsymbol{w}_k)] - 2L\sum_{i\in[n]}f_i^* \leq 2nL\mathbb{E}[f(\boldsymbol{w}_k)].$$

The last step above follows because the f_i^*'s are non-negative. This trick allows us to circumvent the need for the bounded client dissimilarity assumption.

Using these in (11), we get:

$$\mathbb{E}[f(\boldsymbol{w}_{k'})] \leq f(\boldsymbol{w}_0) - \frac{\eta E}{4}\sum_{k=0}^{k'-1}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \underbrace{64\eta LE\Big(\frac{\eta^2 L^2 E(\alpha E+4)}{n} + 5\beta\Big(\frac{q}{n} + \frac{(1+q)(n-r)}{r(n-1)}\Big)\Big)}_{=\gamma}\sum_{k=0}^{k'-1}\mathbb{E}[f(\boldsymbol{w}_k)]. \quad (12)$$

Using (12) recursively, we get:

$$\sum_{k=0}^{k'-1}\mathbb{E}[f(\boldsymbol{w}_k)] \leq k' f(\boldsymbol{w}_0) - \frac{\eta E}{4}\sum_{k=0}^{k'-2}(k'-1-k)\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \gamma\sum_{k=0}^{k'-2}(k'-1-k)\mathbb{E}[f(\boldsymbol{w}_k)] \quad (13)$$

$$\leq k' f(\boldsymbol{w}_0) - \frac{\eta E}{4}\sum_{k=0}^{k'-1}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \gamma k'\sum_{k=0}^{k'-1}\mathbb{E}[f(\boldsymbol{w}_k)]. \quad (14)$$

Let us now ensure that $\gamma k' \leq \frac{1}{2}$ for all $k' \in \{1, \ldots, K\}$, in which case we can simplify (14) to:

$$\sum_{k=0}^{k'-1}\mathbb{E}[f(\boldsymbol{w}_k)] \leq 2k' f(\boldsymbol{w}_0) - \frac{\eta E}{2}\sum_{k=0}^{k'-1}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] \leq 2k' f(\boldsymbol{w}_0). \quad (15)$$

Now:

$$\gamma k' \leq \gamma K = 64\eta LE\Big(\frac{\eta^2 L^2 E(\alpha E+4)}{n} + 5\beta\Big(\frac{q}{n} + \frac{(1+q)(n-r)}{r(n-1)}\Big)\Big)K. \quad (16)$$

Now if we set $8\eta LE^2 \leq 1$, then it can be verified that $\beta = 160e^2(1+q)\eta^2 L^2 E^2(E+1)^2$ is a valid choice. Using this above, we get that:

$$\gamma k' \leq \gamma K = \underbrace{64\eta^3 L^3 E^3 K\Big(\frac{1}{n}\Big(\alpha + \frac{4}{E}\Big) + 800e^2(1+q)(E+1)^2\Big(\frac{q}{n} + \frac{(1+q)(n-r)}{r(n-1)}\Big)\Big)}_{\text{(A)}}. \quad (17)$$

Setting $\eta = \frac{1}{6LEK^{1/3}(\frac{1}{n}(\alpha+\frac{4}{E})+800e^2(1+q)(E+1)^2(\frac{q}{n}+\frac{(1+q)(n-r)}{r(n-1)}))^{1/3}}$, we have (A) $< \frac{1}{2}$. But we must also have

$$8\eta LE^2 = \frac{4E}{3K^{1/3}\big(\frac{1}{n}\big(\alpha+\frac{4}{E}\big) + 800e^2(1+q)(E+1)^2\big(\frac{q}{n} + \frac{(1+q)(n-r)}{r(n-1)}\big)\big)^{1/3}} \leq 1. \quad (18)$$

This holds for $K^{1/3}(E+1) \geq \frac{1}{1200e^2(1+q)\left(\frac{q}{n}+\frac{(1+q)(n-r)}{r(n-1)}\right)}$.

Further β must be smaller than 1, so

$$\beta = 160e^2(1+q)\eta^2 L^2 E^2(E+1)^2 = \frac{160e^2(1+q)(E+1)^2}{36K^{2/3}\big(\big(\frac{1}{n}\big(\alpha+\frac{4}{E}\big) + 800e^2(1+q)(E+1)^2\big(\frac{q}{n} + \frac{(1+q)(n-r)}{r(n-1)}\big)\big)^{2/3}} < 1. \quad (19)$$

This holds for $E+1 \leq \frac{\sqrt{1+q}(n-r)}{3r(n-1)}K$.

Now using (15) in (12) with $k' = K$ and our choice of $\beta = 160e^2(1+q)\eta^2 L^2 E^2(E+1)^2$ and $\eta = \frac{1}{6LEK^{1/3}(\frac{1}{n}(\alpha+\frac{4}{E})+800e^2(1+q)(E+1)^2(\frac{q}{n}+\frac{(1+q)(n-r)}{r(n-1)}))^{1/3}}$, we get:

$$\mathbb{E}[f(\boldsymbol{w}_K)] \leq f(\boldsymbol{w}_0) - \frac{\eta E}{4}\sum_{k=0}^{K-1}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + 128\eta^3 L^3 E^3\Big(\frac{1}{n}\Big(\alpha + \frac{4}{E}\Big) + 800e^2(1+q)(E+1)^2\Big(\frac{q}{n} + \frac{(1+q)(n-r)}{r(n-1)}\Big)\Big)Kf(\boldsymbol{w}_0). \quad (20)$$

Rearranging the above a bit and using the fact that $f(\boldsymbol{w}_K) \geq 0$, we get:

$$\frac{1}{K}\sum_{k=0}^{K-1}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] \leq \frac{4f(\boldsymbol{w}_0)}{\eta EK} + 512\eta^2 L^3 E^2\Big(\frac{1}{n}\Big(\alpha + \frac{4}{E}\Big) + 800e^2(1+q)(E+1)^2\Big(\frac{q}{n} + \frac{(1+q)(n-r)}{r(n-1)}\Big)\Big)f(\boldsymbol{w}_0). \tag{21}$$

Substituting the value of η above, we get:

$$\frac{1}{K}\sum_{k=0}^{K-1}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] \leq \frac{39Lf(\boldsymbol{w}_0)}{K^{2/3}}\Big(\frac{1}{n}\Big(\alpha + \frac{4}{E}\Big) + 800e^2(1+q)(E+1)^2\Big(\frac{q}{n} + \frac{(1+q)(n-r)}{r(n-1)}\Big)\Big)^{1/3}. \tag{22}$$

This concludes the proof. ■

Lemmas used in the proof of Theorem 2:

Lemma 1. *Suppose $4\eta LE^2 \leq 1$ and $\beta \geq \frac{80e^2(1+q)\eta^2L^2E^2(E+1)^2}{(1-4\eta LE)}$. Then for any $k' \in \{1, \ldots, K\}$, we have:*

$$\mathbb{E}[f(\boldsymbol{w}_{k'})] \leq f(\boldsymbol{w}_0) - \frac{\eta E}{4}\sum_{k=0}^{k'-1}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \frac{16\eta^3L^2E^2(\alpha E+4)}{n^2}\sum_{k=0}^{k'-1}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]$$
$$+ \frac{5}{4\eta E\beta}\mathbb{E}[\|\boldsymbol{u}_0 - \overline{\boldsymbol{\delta}}_0\|^2] + 160\eta E\beta\Big(\frac{q}{n^2} + \frac{(1+q)}{r(n-1)}\Big(1 - \frac{r}{n}\Big)\Big)\sum_{k=0}^{k'-1}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2].$$

Proof. Per the previous definitions:

$$\boldsymbol{u}_k = \beta g_Q(\boldsymbol{w}_k; \mathcal{S}_k) + (1-\beta)\boldsymbol{u}_{k-1} + (1-\beta)\Delta g_Q(\boldsymbol{w}_k, \boldsymbol{w}_{k-1}; \mathcal{S}_k) \tag{23}$$

By L-smoothness of f, we have for $k \geq 1$:

$$\mathbb{E}[f(\boldsymbol{w}_{k+1})] \leq \mathbb{E}[f(\boldsymbol{w}_k)] + \mathbb{E}[\langle\nabla f(\boldsymbol{w}_k), \boldsymbol{w}_{k+1} - \boldsymbol{w}_k\rangle] + \frac{L}{2}\mathbb{E}[\|\underbrace{\boldsymbol{w}_{k+1} - \boldsymbol{w}_k}_{=-\boldsymbol{u}_k}\|^2]$$
$$= \mathbb{E}[f(\boldsymbol{w}_k)] + \underbrace{\mathbb{E}[\langle\nabla f(\boldsymbol{w}_k), -\boldsymbol{u}_k\rangle]}_{\text{(I*)}} + \underbrace{\frac{1}{8\eta E}\mathbb{E}[\|\boldsymbol{u}_k\|^2]}_{\text{(II*)}} - \Big(\frac{1}{8\eta E} - \frac{L}{2}\Big)\mathbb{E}[\|\boldsymbol{w}_{k+1} - \boldsymbol{w}_k\|^2]. \tag{24}$$

Let us analyze (I*) first.

$$\begin{aligned}\mathbb{E}[\langle\nabla f(\boldsymbol{w}_k), -\boldsymbol{u}_k\rangle] &= \mathbb{E}[\langle\nabla f(\boldsymbol{w}_k), -g(\boldsymbol{w}_k; \mathcal{S}_k) - (1-\beta)(\boldsymbol{u}_{k-1} - \widehat{g}(\boldsymbol{w}_{k-1}; \mathcal{S}_k))\rangle] \\ &= \mathbb{E}[\langle\nabla f(\boldsymbol{w}_k), -g(\boldsymbol{w}_k; \mathcal{S}_k)\rangle] - (1-\beta)\mathbb{E}[\langle\nabla f(\boldsymbol{w}_k), \boldsymbol{u}_{k-1} - \widehat{g}(\boldsymbol{w}_{k-1}; \mathcal{S}_k)\rangle] \\ &= \underbrace{\mathbb{E}[\langle\nabla f(\boldsymbol{w}_k), \frac{1}{n}\sum_{i\in[n]}(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k)\rangle]}_{\text{(III*)}} + \underbrace{(1-\beta)\mathbb{E}[\langle-\nabla f(\boldsymbol{w}_k), \boldsymbol{u}_{k-1} - \overline{\boldsymbol{\delta}}_{k-1}\rangle]}_{\text{(IV*)}}\end{aligned} \tag{25}$$

(25) follows by taking expectation with respect to Q_D. (III*) is obtained by taking expectation with respect to $\mathcal{S}_k$ above. (IV*) is obtained by taking expectation with respect to $\{\mathcal{B}_{k,1}^{(i)}, \ldots, \mathcal{B}_{k,E-1}^{(i)}\}_{i=1}^n$ and $\mathcal{S}_k$ above.

From Lemma 2, for $\eta < \frac{1}{L}$ and $E < \frac{1}{4}\min\Big(\frac{1}{\eta L}, \frac{1}{\eta^2L^2} - \frac{1}{\eta L}\Big)$, we can bound (III*) as:

$$\text{(III*)} \leq -\frac{\eta E}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] - \frac{\eta}{2}\Big(1 - \eta^2L^2E^2\Big)\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] + \frac{16\eta^3L^2E^2(\alpha E+4)}{n^2}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]. \tag{26}$$

Note that for $\eta L \ll 1$ (which is going to be the case eventually), we can combine all the above constraints on η and E into $4\eta LE < 1$.

As for (IV*):

$$
\begin{aligned}
\text{(IV*)} &\leq (1-\beta)\mathbb{E}\big[\|\nabla f(\boldsymbol{w}_k)\|\|\boldsymbol{u}_{k-1} - \overline{\boldsymbol{\delta}}_{k-1}\|\big] && (27)\\
&\leq \frac{(1-\beta)}{2}\Big(\frac{\eta E}{2(1-\beta)}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \frac{2(1-\beta)\mathbb{E}[\|\boldsymbol{u}_{k-1} - \overline{\boldsymbol{\delta}}_{k-1}\|^2]}{\eta E}\Big) && (28)\\
&= \frac{\eta E}{4}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \frac{(1-\beta)^2}{\eta E}\mathbb{E}[\|\boldsymbol{u}_{k-1} - \overline{\boldsymbol{\delta}}_{k-1}\|^2]. && (29)
\end{aligned}
$$

(28) above follows by the AM-GM inequality.
Adding (26) and (29), we get:

$$
\begin{aligned}
\text{(I*)} \leq -\frac{\eta E}{4}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] - \frac{\eta}{2}(1-\eta^2L^2E^2)\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] + \frac{(1-\beta)^2}{\eta E}\mathbb{E}[\|\boldsymbol{u}_{k-1} - \overline{\boldsymbol{\delta}}_{k-1}\|^2]\\
+ \frac{16\eta^3L^2E^2(\alpha E+4)}{n^2}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]. \quad (30)
\end{aligned}
$$

Now, let us analyze (II*). We have:

$$
\mathbb{E}[\|\boldsymbol{u}_k\|^2] \leq 2\mathbb{E}[\|\overline{\boldsymbol{\delta}}_k\|^2] + 2\mathbb{E}[\|\boldsymbol{u}_k - \overline{\boldsymbol{\delta}}_k\|^2] \quad (31)
$$

Notice that:

$$
\overline{\boldsymbol{\delta}}_k = \mathbb{E}_{\{\mathcal{B}_{k,1}^{(i)},\ldots,\mathcal{B}_{k,E-1}^{(i)}\}_{i=1}^n}\Big[\frac{1}{n}\sum_{i\in[n]}(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)})\Big] = \mathbb{E}_{\{\mathcal{B}_{k,1}^{(i)},\ldots,\mathcal{B}_{k,E-1}^{(i)}\}_{i=1}^n}[\sum_{\tau=0}^{E-1}\eta\overline{\boldsymbol{v}}_{k,\tau}]. \quad (32)
$$

Thus:

$$
\mathbb{E}[\|\overline{\boldsymbol{\delta}}_k\|^2] \leq \eta^2\mathbb{E}\Big[\Big\|\sum_{\tau=0}^{E-1}\overline{\boldsymbol{v}}_{k,\tau}\Big\|^2\Big] \leq E\eta^2\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2]. \quad (33)
$$

The expectation above is with respect to all the randomness in the algorithm so far.
Using (33) and the result of Lemma 8 in (31) with $2\eta LE^2 \leq 1$, we have that:

$$
\begin{aligned}
\mathbb{E}[\|\boldsymbol{u}_k\|^2] \leq 2E\eta^2\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] + 2\Big\{(1-\beta)^2\mathbb{E}[\|\boldsymbol{u}_{k-1} - \overline{\boldsymbol{\delta}}_{k-1}\|^2] + 2\beta^2\mathbb{E}[\|g_Q(\boldsymbol{w}_k;\mathcal{S}_k) - \overline{\boldsymbol{\delta}}_k\|^2]\\
+ 8e^2(1+q)(1-\beta)^2\eta^2L^2E^2(E+1)^2\mathbb{E}[\|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\|^2]\Big\}. \quad (34)
\end{aligned}
$$

Recalling that (II*) = $\frac{1}{8\eta E}\mathbb{E}[\|\boldsymbol{u}_k\|^2]$, we get:

$$
\begin{aligned}
\text{(II*)} \leq \frac{\eta}{4}\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] + \frac{1}{4\eta E}\Big\{(1-\beta)^2\mathbb{E}[\|\boldsymbol{u}_{k-1} - \overline{\boldsymbol{\delta}}_{k-1}\|^2] + 2\beta^2\mathbb{E}[\|g_Q(\boldsymbol{w}_k;\mathcal{S}_k) - \overline{\boldsymbol{\delta}}_k\|^2]\\
+ 8e^2(1+q)(1-\beta)^2\eta^2L^2E^2(E+1)^2\mathbb{E}[\|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\|^2]\Big\}. \quad (35)
\end{aligned}
$$

Adding (30) and (35):

$$
\begin{aligned}
\text{(I*)} + \text{(II*)} \leq -\frac{\eta E}{4}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] - \frac{\eta}{2}\underbrace{\Big(1-\eta^2L^2E^2-\frac{1}{2}\Big)}_{>0 \text{ for } 4\eta LE \leq 1}\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2]\\
+ \frac{16\eta^3L^2E^2}{n^2}(\alpha E+4)\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2] + \frac{5(1-\beta)^2}{4\eta E}\underbrace{\mathbb{E}[\|\boldsymbol{u}_{k-1} - \overline{\boldsymbol{\delta}}_{k-1}\|^2]}_{\text{from Lemma 8}}\\
+ \frac{\beta^2}{2\eta E}\mathbb{E}[\|g_Q(\boldsymbol{w}_k;\mathcal{S}_k) - \overline{\boldsymbol{\delta}}_k\|^2] + 2e^2(1+q)(1-\beta)^2\eta L^2E(E+1)^2\mathbb{E}[\|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\|^2]. \quad (36)
\end{aligned}
$$

Therefore, using Lemma 8 recursively, we get:

$$\begin{aligned}
\text{(I*)} + \text{(II*)} &\leq -\frac{\eta E}{4}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \frac{16\eta^3 L^2 E^2(\alpha E + 4)}{n^2} \sum_{i\in[n]} \mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2] \\
&\quad + \frac{5(1-\beta)^{2k}}{4\eta E}\mathbb{E}[\|\boldsymbol{u}_0 - \overline{\boldsymbol{\delta}}_0\|^2] + \frac{5\beta^2}{2\eta E}\sum_{l=1}^{k}(1-\beta)^{2(k-l)} \underbrace{\mathbb{E}[\|g_Q(\boldsymbol{w}_l;\mathcal{S}_l) - \overline{\boldsymbol{\delta}}_l\|^2]}_{\text{(V*)}} \\
&\quad + 10e^2(1+q)\eta L^2 E(E+1)^2 \sum_{l=1}^{k}(1-\beta)^{2(k-l+1)}\mathbb{E}[\|\boldsymbol{w}_l - \boldsymbol{w}_{l-1}\|^2]. \quad (37)
\end{aligned}$$

Using Lemma 10, we get:

$$\text{(V*)} \leq 4\eta^2 E\Big(\frac{q}{n^2} + \frac{(1+q)}{r(n-1)}\Big(1 - \frac{r}{n}\Big)\Big) \sum_{i\in[n]}\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{v}_{k,\tau}^{(i)}\|^2]. \quad (38)$$

Putting this back in (37), we get:

$$\begin{aligned}
\text{(I*)} + \text{(II*)} &\leq -\frac{\eta E}{4}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \frac{16\eta^3 L^2 E^2(\alpha E + 4)}{n^2} \sum_{i\in[n]} \mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2] \\
&\quad + \frac{5(1-\beta)^{2k}}{4\eta E}\mathbb{E}[\|\boldsymbol{u}_0 - \overline{\boldsymbol{\delta}}_0\|^2] + 10\eta\beta^2\Big(\frac{q}{n^2} + \frac{(1+q)}{r(n-1)}\Big(1 - \frac{r}{n}\Big)\Big)\sum_{l=1}^{k}(1-\beta)^{2(k-l)}\sum_{i\in[n]}\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{v}_{l,\tau}^{(i)}\|^2] \\
&\quad + 10e^2(1+q)\eta L^2 E(E+1)^2 \sum_{l=1}^{k}(1-\beta)^{2(k-l+1)}\mathbb{E}[\|\boldsymbol{w}_l - \boldsymbol{w}_{l-1}\|^2]. \quad (39)
\end{aligned}$$

Next, using (39) in (24), we get that:

$$\begin{aligned}
\mathbb{E}[f(\boldsymbol{w}_{k+1})] &\leq \mathbb{E}[f(\boldsymbol{w}_k)] - \frac{\eta E}{4}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \frac{16\eta^3 L^2 E^2(\alpha E + 4)}{n^2} \sum_{i\in[n]} \mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2] \\
&\quad + \frac{5(1-\beta)^{2k}}{4\eta E}\mathbb{E}[\|\boldsymbol{u}_0 - \overline{\boldsymbol{\delta}}_0\|^2] + 10\eta\beta^2\Big(\frac{q}{n^2} + \frac{(1+q)}{r(n-1)}\Big(1 - \frac{r}{n}\Big)\Big)\sum_{l=1}^{k}(1-\beta)^{2(k-l)}\sum_{i\in[n]}\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{v}_{l,\tau}^{(i)}\|^2] \\
&\quad + 10e^2(1+q)\eta L^2 E(E+1)^2 \sum_{l=1}^{k}(1-\beta)^{2(k-l+1)}\mathbb{E}[\|\boldsymbol{w}_l - \boldsymbol{w}_{l-1}\|^2] - \Big(\frac{1}{8\eta E} - \frac{L}{2}\Big)\mathbb{E}[\|\boldsymbol{w}_{k+1} - \boldsymbol{w}_k\|^2]. \quad (40)
\end{aligned}$$

Summing the above from $k = 0$ through $(k'-1)$ for any $k' \in \{1, \dots, K\}$, we get:

$$\begin{aligned}
\mathbb{E}[f(\boldsymbol{w}_{k'})] &\leq f(\boldsymbol{w}_0) - \frac{\eta E}{4}\sum_{k=0}^{k'-1}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \frac{16\eta^3 L^2 E^2(\alpha E + 4)}{n^2}\sum_{k=0}^{k'-1}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2] \\
&\quad + \sum_{l=0}^{\infty}\frac{5(1-\beta)^{2l}}{4\eta E}\mathbb{E}[\|\boldsymbol{u}_0 - \overline{\boldsymbol{\delta}}_0\|^2] + 10\eta\beta^2\Big(\frac{q}{n^2} + \frac{(1+q)}{r(n-1)}\Big(1 - \frac{r}{n}\Big)\Big)\sum_{k=0}^{k'-1}\sum_{i\in[n]}\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{v}_{k,\tau}^{(i)}\|^2]\sum_{l=0}^{\infty}(1-\beta)^{2l} \\
&\quad + 10e^2(1+q)\eta L^2 E(E+1)^2(1-\beta)^2\sum_{k=1}^{k'-1}\mathbb{E}[\|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\|^2]\sum_{l=0}^{\infty}(1-\beta)^{2l} - \Big(\frac{1}{8\eta E} - \frac{L}{2}\Big)\sum_{k=0}^{k'-1}\mathbb{E}[\|\boldsymbol{w}_{k+1} - \boldsymbol{w}_k\|^2].
\end{aligned}$$

(41)

Simplifying the above by noting that $\sum_{l=0}^{\infty}(1-\beta)^{2l} \leq \sum_{l=0}^{\infty}(1-\beta)^{l} = 1/\beta$, we get:

$$\begin{aligned}
\mathbb{E}[f(\boldsymbol{w}_{k'})] &\leq f(\boldsymbol{w}_0) - \frac{\eta E}{4}\sum_{k=0}^{k'-1}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \frac{16\eta^3 L^2 E^2(\alpha E + 4)}{n^2}\sum_{k=0}^{k'-1}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2] \\
&+ \frac{5}{4\eta E\beta}\mathbb{E}[\|\boldsymbol{u}_0 - \overline{\boldsymbol{\delta}}_0\|^2] + 10\eta\beta\Big(\frac{q}{n^2} + \frac{(1+q)}{r(n-1)}\Big(1-\frac{r}{n}\Big)\Big)\sum_{k=0}^{k'-1}\sum_{i\in[n]}\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{v}_{k,\tau}^{(i)}\|^2] \\
&+ \underbrace{\frac{10e^2(1+q)\eta L^2 E(E+1)^2}{\beta}\sum_{k=1}^{k'-1}\mathbb{E}[\|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\|^2] - \frac{(1-4\eta LE)}{8\eta E}\sum_{k=0}^{k'-1}\mathbb{E}[\|\boldsymbol{w}_{k+1} - \boldsymbol{w}_k\|^2]}_{\text{(VI*) – want this to be } \leq 0}
\end{aligned} \tag{42}$$

We want (VI*) to be ≤ 0. For this, we must have:

$$\beta \geq \frac{80e^2(1+q)\eta^2 L^2 E^2(E+1)^2}{(1-4\eta LE)}. \tag{43}$$

Note that the denominator above is positive since we already have a constraint of $4\eta LE \leq 1$.

With β satisfying the above constraint, and using the result of Lemma 5 for $\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{v}_{k,\tau}^{(i)}\|^2]$, we get:

$$\begin{aligned}
\mathbb{E}[f(\boldsymbol{w}_{k'})] &\leq f(\boldsymbol{w}_0) - \frac{\eta E}{4}\sum_{k=0}^{k'-1}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \frac{16\eta^3 L^2 E^2(\alpha E + 4)}{n^2}\sum_{k=0}^{k'-1}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2] \\
&+ \frac{5}{4\eta E\beta}\mathbb{E}[\|\boldsymbol{u}_0 - \overline{\boldsymbol{\delta}}_0\|^2] + 160\eta E\beta\Big(\frac{q}{n^2} + \frac{(1+q)}{r(n-1)}\Big(1-\frac{r}{n}\Big)\Big)\sum_{k=0}^{k'-1}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2].
\end{aligned} \tag{44}$$

Finally, note that we have two constraints namely: $4\eta LE \leq 1$ and $2\eta LE^2 \leq 1$. We can merge these constraints into $4\eta LE^2 \leq 1$ for $E \geq 1$ (which is the case).

This gives us the desired result. ■

Lemma 2. *For $\eta < \frac{1}{L}$ and $E < \frac{1}{4}min\Big(\frac{1}{\eta L}, \frac{1}{\eta^2 L^2} - \frac{1}{\eta L}\Big)$, (III*) in the proof of Lemma 1 can be bounded as:*

$$\begin{aligned}
\textit{(III*)} &= \mathbb{E}[\langle\nabla f(\boldsymbol{w}_k), \frac{1}{n}\sum_{i\in[n]}(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k)\rangle] \leq -\frac{\eta E}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] - \frac{\eta}{2}\Big(1-\eta^2 L^2 E^2\Big)\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] \\
&+ \frac{16\eta^3 L^2 E^2(\alpha E + 4)}{n^2}\sum_{i\in[n]}\|\nabla f_i(\boldsymbol{w}_k)\|^2.
\end{aligned}$$

Proof. (III*) $= \mathbb{E}[\langle \nabla f(\boldsymbol{w}_k), \frac{1}{n}\sum_{i\in[n]}(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k)\rangle]$. Then:

$$\begin{aligned}
\text{(III*)} &= \mathbb{E}[\langle \nabla f(\boldsymbol{w}_k), -\frac{1}{n}\sum_{i\in[n]}\sum_{\tau=0}^{E-1}\eta \boldsymbol{v}_{k,\tau}^{(i)}\rangle] \\
&= -\eta \sum_{\tau=0}^{E-1}\mathbb{E}[\langle \nabla f(\boldsymbol{w}_k), \underbrace{\frac{1}{n}\sum_{i\in[n]}\boldsymbol{v}_{k,\tau}^{(i)}}_{=\overline{\boldsymbol{v}}_{k,\tau}}\rangle] \\
&= \sum_{\tau=0}^{E-1}\Big\{ -\frac{\eta}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] - \frac{\eta}{2}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] + \frac{\eta}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k) - \overline{\boldsymbol{v}}_{k,\tau}\|^2]\Big\} \qquad (45) \\
&= \sum_{\tau=0}^{E-1}\Big\{ -\frac{\eta}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] - \frac{\eta}{2}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] + \frac{\eta}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k) - \nabla f(\overline{\boldsymbol{w}}_{k,\tau}) + \nabla f(\overline{\boldsymbol{w}}_{k,\tau}) - \overline{\boldsymbol{v}}_{k,\tau}\|^2]\Big\} \\
&\leq \sum_{\tau=0}^{E-1}\Big\{ -\frac{\eta}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] - \frac{\eta}{2}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] + \eta\mathbb{E}[\underbrace{\|\nabla f(\boldsymbol{w}_k) - \nabla f(\overline{\boldsymbol{w}}_{k,\tau})\|}_{\leq L\|\boldsymbol{w}_k - \overline{\boldsymbol{w}}_{k,\tau}\|}]^2 + \eta\mathbb{E}[\|\nabla f(\overline{\boldsymbol{w}}_{k,\tau}) - \overline{\boldsymbol{v}}_{k,\tau}\|^2]\Big\} \qquad (46) \\
&\leq \sum_{\tau=0}^{E-1}\Big\{ -\frac{\eta}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] - \frac{\eta}{2}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] + \eta L^2\mathbb{E}[\|\boldsymbol{w}_k - \overline{\boldsymbol{w}}_{k,\tau}\|]^2 + \eta\mathbb{E}[\|\nabla f(\overline{\boldsymbol{w}}_{k,\tau}) - \overline{\boldsymbol{v}}_{k,\tau}\|^2]\Big\} \qquad (47)
\end{aligned}$$

(45) above follows by using the fact that for any two vectors $\boldsymbol{a}$ and $\boldsymbol{b}$, $\langle \boldsymbol{a}, \boldsymbol{b}\rangle = \frac{1}{2}(\|\boldsymbol{a}\|^2 + \|\boldsymbol{b}\|^2 - \|\boldsymbol{a}-\boldsymbol{b}\|^2)$. Also, (46) follows from the fact that for any two vectors $\boldsymbol{a}$ and $\boldsymbol{b}$, $\|\boldsymbol{a}+\boldsymbol{b}\|^2 \leq 2\|\boldsymbol{a}\|^2 + 2\|\boldsymbol{b}\|^2$.
Per definitions, observe that:

$$\overline{\boldsymbol{w}}_{k,\tau+1} = \overline{\boldsymbol{w}}_{k,\tau} - \eta\overline{\boldsymbol{v}}_{k,\tau}. \qquad (48)$$

From this, we have that $\boldsymbol{w}_k - \overline{\boldsymbol{w}}_{k,\tau} = \eta\sum_{t=0}^{\tau-1}\overline{\boldsymbol{v}}_{k,t}$. Hence, $\|\boldsymbol{w}_k - \overline{\boldsymbol{w}}_{k,\tau}\|^2 = \eta^2\|\sum_{t=0}^{\tau-1}\overline{\boldsymbol{v}}_{k,t}\|^2 \leq \eta^2\tau\sum_{t=0}^{\tau-1}\|\overline{\boldsymbol{v}}_{k,t}\|^2$ – this follows from the fact that for any $p > 1$ vectors $\{\boldsymbol{u}_1, \ldots, \boldsymbol{u}_p\}$, $\|\sum_{i=1}^{p}\boldsymbol{u}_i\|^2 \leq p\sum_{i=1}^{p}\|\boldsymbol{u}_i\|^2$. Using all this in (47), we get:

$$\begin{aligned}
\text{(III*)} &\leq -\frac{\eta E}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \sum_{\tau=0}^{E-1}\Big\{ -\frac{\eta}{2}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] + \eta^3L^2\tau\sum_{t=0}^{\tau-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,t}\|^2] + \eta\mathbb{E}[\|\nabla f(\overline{\boldsymbol{w}}_{k,\tau}) - \overline{\boldsymbol{v}}_{k,\tau}\|^2]\Big\} \\
&\leq -\frac{\eta E}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] - \frac{\eta}{2}\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] + \frac{\eta^3L^2E^2}{2}\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] + \eta\underbrace{\sum_{\tau=0}^{E-1}\mathbb{E}[\|\nabla f(\overline{\boldsymbol{w}}_{k,\tau}) - \overline{\boldsymbol{v}}_{k,\tau}\|^2]}_{\text{from Lemma 3}} \qquad (49)
\end{aligned}$$

Using Lemma 3 to bound the last term above gives us:

$$\text{(III*)} \leq -\frac{\eta E}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] - \frac{\eta}{2}\Big(1 - \eta^2L^2E^2\Big)\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] + \frac{16\eta^3L^2E^2(\alpha E + 4)}{n^2}\sum_{i\in[n]}\|\nabla f_i(\boldsymbol{w}_k)\|^2. \qquad (50)$$

This gives us the desired result. ■

Lemma 3. *For $\eta < \frac{1}{L}$ and $E < \frac{1}{4}min\Big(\frac{1}{\eta L}, \frac{1}{\eta^2L^2} - \frac{1}{\eta L}\Big)$, we have:*

$$\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau} - \nabla f(\overline{\boldsymbol{w}}_{k,\tau})\|^2] \leq \frac{16\eta^2L^2E^2(\alpha E + 4)}{n^2}\sum_{i\in[n]}\|\nabla f_i(\boldsymbol{w}_k)\|^2,$$

where the expectation is with respect to the randomness due to $\{\mathcal{B}_{k,1}^{(i)}, \ldots, \mathcal{B}_{k,E-1}^{(i)}\}_{i=1}^{n}$.

Proof. Let $\overline{\boldsymbol{e}}_{k,\tau} = \overline{\boldsymbol{v}}_{k,\tau} - \nabla f(\overline{\boldsymbol{w}}_{k,\tau})$. Then:

$$\begin{aligned}
\|\overline{\boldsymbol{e}}_{k,\tau}\|^2 &= \|\overline{\boldsymbol{v}}_{k,\tau} - \nabla f(\overline{\boldsymbol{w}}_{k,\tau})\|^2 \\
&= \Big\| \frac{1}{n} \sum_{i \in [n]} (\boldsymbol{v}_{k,\tau}^{(i)} - \nabla f_i(\overline{\boldsymbol{w}}_{k,\tau})) \Big\|^2 \\
&= \Big\| \frac{1}{n} \sum_{i \in [n]} (\boldsymbol{e}_{k,\tau}^{(i)} + \widetilde{\boldsymbol{e}}_{k,\tau}^{(i)}) \Big\|^2 \\
&\leq \frac{2}{n^2} \Big\| \sum_{i \in [n]} \boldsymbol{e}_{k,\tau}^{(i)} \Big\|^2 + \frac{2}{n^2} \Big\| \sum_{i \in [n]} \widetilde{\boldsymbol{e}}_{k,\tau}^{(i)} \Big\|^2
\end{aligned} \tag{51}$$

So:

$$\mathbb{E}[\|\overline{\boldsymbol{e}}_{k,\tau}\|^2] \leq \frac{2}{n^2} \mathbb{E}\Big[\Big\| \sum_{i \in [n]} \boldsymbol{e}_{k,\tau}^{(i)} \Big\|^2\Big] + \frac{2}{n^2} \mathbb{E}\Big[\Big\| \sum_{i \in [n]} \widetilde{\boldsymbol{e}}_{k,\tau}^{(i)} \Big\|^2\Big] \tag{52}$$

But:

$$\mathbb{E}\Big[\Big\| \sum_{i \in [n]} \boldsymbol{e}_{k,\tau}^{(i)} \Big\|^2\Big] = \sum_{i \in [n]} \mathbb{E}\Big[\Big\|\boldsymbol{e}_{k,\tau}^{(i)}\Big\|^2\Big] + \sum_{i \neq j : i,j \in [n]} \langle \mathbb{E}[\boldsymbol{e}_{k,\tau}^{(i)}], \mathbb{E}[\boldsymbol{e}_{k,\tau}^{(j)}] \rangle$$

In the cross-term above, we can take expectations individually as $\{\mathcal{B}_{k,1}^{(i)}, \ldots, \mathcal{B}_{k,E-1}^{(i)}\}$ and $\{\mathcal{B}_{k,1}^{(j)}, \ldots, \mathcal{B}_{k,E-1}^{(j)}\}$ are independent for $i \neq j$. Next, from Lemma 4, $\mathbb{E}[\boldsymbol{e}_{k,\tau}^{(i)}] = \vec{0} \; \forall \, i, k, \tau$. Hence:

$$\mathbb{E}\Big[\Big\| \sum_{i \in [n]} \boldsymbol{e}_{k,\tau}^{(i)} \Big\|^2\Big] = \sum_{i \in [n]} \mathbb{E}\Big[\Big\|\boldsymbol{e}_{k,\tau}^{(i)}\Big\|^2\Big].$$

Using the above result and Assumption 4 in (52), we get that:

$$\mathbb{E}[\|\overline{\boldsymbol{e}}_{k,\tau}\|^2] \leq \frac{2}{n^2} \sum_{i \in [n]} \mathbb{E}[\|\boldsymbol{e}_{k,\tau}^{(i)}\|^2] + \frac{2\alpha}{n^2} \sum_{i \in [n]} \mathbb{E}[\|\widetilde{\boldsymbol{e}}_{k,\tau}^{(i)}\|^2]. \tag{53}$$

Now:

$$\begin{aligned}
\mathbb{E}\Big[\Big\|\widetilde{\boldsymbol{e}}_{k,\tau}^{(i)}\Big\|^2\Big] &= \mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,\tau}^{(i)}) - \nabla f_i(\overline{\boldsymbol{w}}_{k,\tau})\|^2] \\
&= L^2 \mathbb{E}[\|\boldsymbol{w}_{k,\tau}^{(i)} - \overline{\boldsymbol{w}}_{k,\tau}\|^2] \\
&\leq L^2 \mathbb{E}[\|(\boldsymbol{w}_{k,0}^{(i)} - \eta \sum_{t=0}^{\tau-1} \boldsymbol{v}_{k,t}^{(i)}) - (\overline{\boldsymbol{w}}_{k,0} - \eta \sum_{t=0}^{\tau-1} \overline{\boldsymbol{v}}_{k,t})\|^2]
\end{aligned}$$

But since $\boldsymbol{w}_{k,0}^{(i)} = \boldsymbol{w}_k \; \forall \, i$, we have $\overline{\boldsymbol{w}}_{k,0} = \boldsymbol{w}_k$. Hence:

$$\begin{aligned}
\mathbb{E}\Big[\Big\|\widetilde{\boldsymbol{e}}_{k,\tau}^{(i)}\Big\|^2\Big] &= \eta^2 L^2 \mathbb{E}[\| \sum_{t=0}^{\tau-1} \overline{\boldsymbol{v}}_{k,t} - \sum_{t=0}^{\tau-1} \boldsymbol{v}_{k,t}^{(i)} \|^2] \\
&\leq \eta^2 L^2 \tau \sum_{t=0}^{\tau-1} \mathbb{E}[\|\overline{\boldsymbol{v}}_{k,t} - \boldsymbol{v}_{k,t}^{(i)}\|^2] \\
&= \eta^2 L^2 \tau \sum_{t=0}^{\tau-1} \mathbb{E}[\|\overline{\boldsymbol{v}}_{k,t}\|^2 + \|\boldsymbol{v}_{k,t}^{(i)}\|^2 - 2\langle \overline{\boldsymbol{v}}_{k,t}, \boldsymbol{v}_{k,t}^{(i)} \rangle]
\end{aligned}$$

Substituting the above in (53), we get:

$$\mathbb{E}[\|\overline{\boldsymbol{e}}_{k,\tau}\|^2] \leq \frac{2}{n^2}\sum_{i\in[n]}\mathbb{E}[\|\boldsymbol{e}^{(i)}_{k,\tau}\|^2] + \frac{2\alpha}{n^2}\sum_{i\in[n]}\eta^2L^2\tau\sum_{t=0}^{\tau-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,t}\|^2 + \|\boldsymbol{v}^{(i)}_{k,t}\|^2 - 2\langle\overline{\boldsymbol{v}}_{k,t},\boldsymbol{v}^{(i)}_{k,t}\rangle]$$

$$= \frac{2}{n^2}\sum_{i\in[n]}\mathbb{E}[\|\boldsymbol{e}^{(i)}_{k,\tau}\|^2] + \frac{2\alpha\eta^2L^2\tau}{n^2}\sum_{t=0}^{\tau-1}\{n\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,t}\|^2] + \sum_{i\in[n]}\mathbb{E}[\|\boldsymbol{v}^{(i)}_{k,t}\|^2] - 2\langle\overline{\boldsymbol{v}}_{k,t},\sum_{i\in[n]}\boldsymbol{v}^{(i)}_{k,t}\rangle\} \tag{54}$$

$$= \frac{2}{n^2}\sum_{i\in[n]}\mathbb{E}[\|\boldsymbol{e}^{(i)}_{k,\tau}\|^2] + \frac{2\alpha\eta^2L^2\tau}{n^2}\sum_{t=0}^{\tau-1}\sum_{i\in[n]}\Big(\mathbb{E}[\|\boldsymbol{v}^{(i)}_{k,t}\|^2] - \mathbb{E}[\|\overline{\boldsymbol{v}}_{k,t}\|^2]\Big). \tag{55}$$

$$\leq \frac{2}{n^2}\sum_{i\in[n]}\mathbb{E}[\|\boldsymbol{e}^{(i)}_{k,\tau}\|^2] + \frac{2\alpha\eta^2L^2\tau}{n^2}\sum_{t=0}^{\tau-1}\sum_{i\in[n]}\mathbb{E}[\|\boldsymbol{v}^{(i)}_{k,t}\|^2]. \tag{56}$$

To get (55) from (54), we use the fact $\sum_{i\in[n]}\boldsymbol{v}^{(i)}_{k,t} = n\overline{\boldsymbol{v}}_{k,t}$. Now summing up (56) from $\tau = 0$ through to $E-1$, we get:

$$\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{e}}_{k,\tau}\|^2] \leq \frac{2}{n^2}\sum_{i\in[n]}\underbrace{\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{e}^{(i)}_{k,\tau}\|^2]}_{\text{from Lemma 7}} + \frac{2\alpha\eta^2L^2E^2}{2n^2}\sum_{i\in[n]}\underbrace{\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{v}^{(i)}_{k,\tau}\|^2]}_{\text{from Lemma 5}}. \tag{57}$$

Now using Lemma 7 and Lemma 5 above with $\eta < \frac{1}{L}$ and $E < \frac{1}{4}\min\Big(\frac{1}{\eta L}, \frac{1}{\eta^2L^2} - \frac{1}{\eta L}\Big)$, we get:

$$\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{e}}_{k,\tau}\|^2] \leq \frac{2}{n^2}\sum_{i\in[n]}32E^2\eta^2L^2\|\nabla f_i(\boldsymbol{w}_k)\|^2 + \frac{\alpha\eta^2L^2E^2}{n^2}\sum_{i\in[n]}16E\|\nabla f_i(\boldsymbol{w}_k)\|^2.$$

This gives us the desired result. ■

Lemma 4. $\mathbb{E}_{\mathcal{B}^{(i)}_{k,1},\ldots,\mathcal{B}^{(i)}_{k,\tau}}[\boldsymbol{e}^{(i)}_{k,\tau}] = \vec{0}\;\forall\; k \in \{0,\ldots,K-1\}, \tau \in \{1,\ldots,E-1\}$.

Proof. Note that:

$$\boldsymbol{e}^{(i)}_{k,0} = \boldsymbol{v}^{(i)}_{k,0} - \nabla f_i(\boldsymbol{w}^{(i)}_{k,0}) = \vec{0}.$$

For $\tau > 0$:

$$\begin{aligned}
&\mathbb{E}_{\mathcal{B}^{(i)}_{k,1},\ldots,\mathcal{B}^{(i)}_{k,\tau}}[\boldsymbol{e}^{(i)}_{k,\tau}] = \mathbb{E}_{\mathcal{B}^{(i)}_{k,1},\ldots,\mathcal{B}^{(i)}_{k,\tau}}[\boldsymbol{v}^{(i)}_{k,\tau} - \nabla f_i(\boldsymbol{w}^{(i)}_{k,\tau})]\\
&= \mathbb{E}_{\mathcal{B}^{(i)}_{k,1},\ldots,\mathcal{B}^{(i)}_{k,\tau}}[\widetilde{\nabla} f_i(\boldsymbol{w}^{(i)}_{k,\tau};\mathcal{B}^{(i)}_{k,\tau}) + (\boldsymbol{v}^{(i)}_{k,\tau-1} - \widetilde{\nabla} f_i(\boldsymbol{w}^{(i)}_{k,\tau-1};\mathcal{B}^{(i)}_{k,\tau})) - \nabla f_i(\boldsymbol{w}^{(i)}_{k,\tau})]\\
&= \mathbb{E}_{\mathcal{B}^{(i)}_{k,1},\ldots,\mathcal{B}^{(i)}_{k,\tau-1}}[\mathbb{E}_{\mathcal{B}^{(i)}_{k,\tau}}[\widetilde{\nabla} f_i(\boldsymbol{w}^{(i)}_{k,\tau};\mathcal{B}^{(i)}_{k,\tau}) + (\boldsymbol{v}^{(i)}_{k,\tau-1} - \widetilde{\nabla} f_i(\boldsymbol{w}^{(i)}_{k,\tau-1};\mathcal{B}^{(i)}_{k,\tau})) - \nabla f_i(\boldsymbol{w}^{(i)}_{k,\tau})|\mathcal{B}^{(i)}_{k,1},\ldots,\mathcal{B}^{(i)}_{k,\tau-1}]]\\
&= \mathbb{E}_{\mathcal{B}^{(i)}_{k,1},\ldots,\mathcal{B}^{(i)}_{k,\tau-1}}[(\boldsymbol{v}^{(i)}_{k,\tau-1} - \nabla f_i(\boldsymbol{w}^{(i)}_{k,\tau-1}))]\\
&= \mathbb{E}_{\mathcal{B}^{(i)}_{k,1},\ldots,\mathcal{B}^{(i)}_{k,\tau-1}}[\boldsymbol{e}^{(i)}_{k,\tau-1}].
\end{aligned}$$

Doing this recursively, we get:

$$\mathbb{E}_{\mathcal{B}^{(i)}_{k,1},\ldots,\mathcal{B}^{(i)}_{k,\tau}}[\boldsymbol{e}^{(i)}_{k,\tau}] = \boldsymbol{e}^{(i)}_{k,0} = \vec{0}. \tag{58}$$

Note that this result also holds if we use full gradients at $\tau = 0$. ■

Lemma 5. *For $\eta < \frac{1}{L}$ and $E < \frac{1}{4}min\Big(\frac{1}{\eta L}, \frac{1}{\eta^2L^2} - \frac{1}{\eta L}\Big)$, we have:*

$$\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{v}^{(i)}_{k,\tau}\|^2] \leq 16E\|\nabla f_i(\boldsymbol{w}_k)\|^2.$$

Note that in this lemma, the expectation is with respect to the randomness only due to $\{\mathcal{B}_{k,1}^{(i)}, \ldots, \mathcal{B}_{k,E-1}^{(i)}\}_{i=1}^{n}$.

Proof. First, recall that $\boldsymbol{e}_{k,\tau}^{(i)} = \boldsymbol{v}_{k,\tau}^{(i)} - \nabla f_i(\boldsymbol{w}_{k,\tau}^{(i)})$. Note that $\boldsymbol{e}_{k,0}^{(i)} = \vec{0}$, as we are using clients' full gradients at $\tau = 0$. We have:

$$\mathbb{E}[\|\boldsymbol{v}_{k,\tau}^{(i)}\|^2] \leq 2\mathbb{E}[\|\boldsymbol{e}_{k,\tau}^{(i)}\|^2] + 2\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,\tau}^{(i)})\|^2]. \tag{59}$$

Using Lemma 2.1 of Liu et al. [2020] with $\beta = 0$, we have:

$$\begin{aligned}\mathbb{E}[\|\boldsymbol{e}_{k,\tau}^{(i)}\|^2] &\leq \mathbb{E}[\|\boldsymbol{e}_{k,0}^{(i)}\|^2] + 2L^2 \sum_{t=0}^{\tau-1} \mathbb{E}[\|\boldsymbol{w}_{k,t+1}^{(i)} - \boldsymbol{w}_{k,t}^{(i)}\|^2] \\ &\leq 2L^2 \sum_{t=0}^{\tau-1} \mathbb{E}[\|\boldsymbol{w}_{k,t+1}^{(i)} - \boldsymbol{w}_{k,t}^{(i)}\|^2]. \end{aligned}\tag{60}$$

The last step follows because $\boldsymbol{e}_{k,0}^{(i)} = \vec{0}$.

Summing the above from $\tau = 0$ through to $E - 1$, we get:

$$\begin{aligned}\sum_{\tau=0}^{E-1} \mathbb{E}[\|\boldsymbol{e}_{k,\tau}^{(i)}\|^2] &\leq 2L^2 \sum_{\tau=0}^{E-1}\sum_{t=0}^{\tau-1} \mathbb{E}[\|\boldsymbol{w}_{k,t+1}^{(i)} - \boldsymbol{w}_{k,t}^{(i)}\|^2] \\ &\leq 2EL^2 \sum_{\tau=0}^{E-2} \mathbb{E}[\|\boldsymbol{w}_{k,\tau+1}^{(i)} - \boldsymbol{w}_{k,\tau}^{(i)}\|^2]. \end{aligned}\tag{61}$$

Next, re-arranging equation (11) in Lemma 2.2 of Liu et al. [2020] (observe that in our case, $G_\eta(.)$ is simply the gradient), we get:

$$\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,\tau}^{(i)})\|^2] \leq \frac{2}{\eta}\mathbb{E}[f_i(\boldsymbol{w}_{k,\tau}^{(i)}) - f_i(\boldsymbol{w}_{k,\tau+1}^{(i)})] - \frac{1}{\eta^2}(1-\eta L)\mathbb{E}[\|\boldsymbol{w}_{k,\tau+1}^{(i)} - \boldsymbol{w}_{k,\tau}^{(i)}\|^2] + \mathbb{E}[\|\boldsymbol{e}_{k,\tau}^{(i)}\|^2] \tag{62}$$

Summing (62) from $\tau = 0$ to $E - 1$ and using (61), we get:

$$\begin{aligned}\sum_{\tau=0}^{E-1} \mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,\tau}^{(i)})\|^2] \leq \frac{2}{\eta}(f_i(\boldsymbol{w}_k) - \mathbb{E}[f_i(\boldsymbol{w}_{k,E}^{(i)})]) - \frac{(1-\eta L)}{\eta^2} \sum_{\tau=0}^{E-1} \mathbb{E}[\|\boldsymbol{w}_{k,\tau+1}^{(i)} - \boldsymbol{w}_{k,\tau}^{(i)}\|^2] \\ + 2EL^2 \sum_{\tau=0}^{E-2} \mathbb{E}[\|\boldsymbol{w}_{k,\tau+1}^{(i)} - \boldsymbol{w}_{k,\tau}^{(i)}\|^2]. \end{aligned}\tag{63}$$

Next, summing (61) and (63) gives us:

$$\begin{aligned}\sum_{\tau=0}^{E-1} \{\mathbb{E}[\|\boldsymbol{e}_{k,\tau}^{(i)}\|^2] + \mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,\tau}^{(i)})\|^2]\} \leq \frac{2}{\eta}(f_i(\boldsymbol{w}_k) - \mathbb{E}[f_i(\boldsymbol{w}_{k,E}^{(i)})]) \\ - \underbrace{\left(\frac{1-\eta L}{\eta^2}\right)}_{>0 \text{ for } \eta < \frac{1}{L}} \mathbb{E}[\|\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_{k,E-1}^{(i)}\|^2] - \underbrace{\left(\frac{(1-\eta L)}{\eta^2} - 4EL^2\right)}_{>0 \text{ for } E < \frac{(1-\eta L)}{4\eta^2 L^2}} \sum_{\tau=0}^{E-2} \mathbb{E}[\|\boldsymbol{w}_{k,\tau+1}^{(i)} - \boldsymbol{w}_{k,\tau}^{(i)}\|^2]. \end{aligned}\tag{64}$$

So if we have $\eta < \frac{1}{L}$ and $E < \frac{1}{4}(\frac{1}{\eta^2 L^2} - \frac{1}{\eta L})$, we get:

$$\sum_{\tau=0}^{E-1} \{\mathbb{E}[\|\boldsymbol{e}_{k,\tau}^{(i)}\|^2] + \mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,\tau}^{(i)})\|^2]\} \leq \frac{2}{\eta}(f_i(\boldsymbol{w}_k) - \mathbb{E}[f_i(\boldsymbol{w}_{k,E}^{(i)})]). \tag{65}$$

Now from Lemma 6, for $E < \frac{1}{4}\min\left(\frac{1}{\eta L}, \frac{1}{\eta^2 L^2} - \frac{1}{\eta L}\right)$, we have that:

$$f_i(\boldsymbol{w}_k) - \mathbb{E}[f_i(\boldsymbol{w}_{k,E}^{(i)})] \leq 4\eta E\|\nabla f_i(\boldsymbol{w}_k)\|^2. \tag{66}$$

Putting (66) in (65) and then using it (59) gives us the desired result. ∎

Lemma 6. *For $\eta < \frac{1}{L}$ and $E < \frac{1}{4} min\Big(\frac{1}{\eta L}, \frac{1}{\eta^2 L^2} - \frac{1}{\eta L}\Big)$, we have:*

$$f_i(\boldsymbol{w}_k) - \mathbb{E}[f_i(\boldsymbol{w}_{k,E}^{(i)})] \leq 4\eta E \|\nabla f_i(\boldsymbol{w}_k)\|^2.$$

The expectation above is with respect to the randomness only due to $\{\mathcal{B}_{k,1}^{(i)}, \ldots, \mathcal{B}_{k,E-1}^{(i)}\}_{i=1}^n$.

Proof. By L-smoothness of each f_i, we have:

$$f_i(\boldsymbol{w}_{k,E}^{(i)}) \geq f_i(\boldsymbol{w}_k) + \langle \nabla f_i(\boldsymbol{w}_k), \boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k \rangle - \frac{L}{2}\|\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k\|^2$$

$$\begin{aligned} \implies f_i(\boldsymbol{w}_k) - f_i(\boldsymbol{w}_{k,E}^{(i)}) &\leq \langle \nabla f_i(\boldsymbol{w}_k), \boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)} \rangle + \frac{L}{2}\|\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k\|^2 \\ &\leq \underbrace{\frac{\alpha'}{2}\|\nabla f_i(\boldsymbol{w}_k)\|^2 + \frac{1}{2\alpha'}\|\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k\|^2}_{\text{follows by Young's inequality}} + \frac{L}{2}\|\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k\|^2 \text{ for } \alpha' > 0. \end{aligned}$$

Recall that $\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k = \eta \sum_{\tau=0}^{E-1} \boldsymbol{v}_{k,\tau}^{(i)}$. Hence taking expectation above with $\alpha' = 2\eta E$, we get that:

$$f_i(\boldsymbol{w}_k) - \mathbb{E}[f_i(\boldsymbol{w}_{k,E}^{(i)})] \leq \eta E \|\nabla f_i(\boldsymbol{w}_k)\|^2 + \eta^2 E \Big(\frac{1}{4\eta E} + \frac{L}{2}\Big) \sum_{\tau=0}^{E-1} \mathbb{E}[\|\boldsymbol{v}_{k,\tau}^{(i)}\|^2] \tag{67}$$

$$\leq \eta E \|\nabla f_i(\boldsymbol{w}_k)\|^2 + \frac{3\eta}{8} \sum_{\tau=0}^{E-1} \mathbb{E}[\|\boldsymbol{v}_{k,\tau}^{(i)}\|^2]. \tag{68}$$

(68) follows from the fact that $\eta L E < \frac{1}{4}$. Next, from the proof of Lemma 5, for $E < \frac{(1-\eta L)}{4\eta^2 L^2}$:

$$\sum_{\tau=0}^{E-1} \mathbb{E}[\|\boldsymbol{v}_{k,\tau}^{(i)}\|^2] \leq \frac{2}{\eta}(f_i(\boldsymbol{w}_k) - \mathbb{E}[f_i(\boldsymbol{w}_{k,E}^{(i)})]).$$

Putting this in (68), we get:

$$f_i(\boldsymbol{w}_k) - \mathbb{E}[f_i(\boldsymbol{w}_{k,E}^{(i)})] \leq \eta E \|\nabla f_i(\boldsymbol{w}_k)\|^2 + \frac{3}{4}(f_i(\boldsymbol{w}_k) - \mathbb{E}[f_i(\boldsymbol{w}_{k,E}^{(i)})].$$

$$\implies f_i(\boldsymbol{w}_k) - \mathbb{E}[f_i(\boldsymbol{w}_{k,E}^{(i)})] \leq 4\eta E \|\nabla f_i(\boldsymbol{w}_k)\|^2. \tag{69}$$

■

Lemma 7. *For $\eta < \frac{1}{L}$ and $E < \frac{1}{4} min\Big(\frac{1}{\eta L}, \frac{1}{\eta^2 L^2} - \frac{1}{\eta L}\Big)$, we have:*

$$\sum_{\tau=0}^{E-1} \mathbb{E}[\|\boldsymbol{e}_{k,\tau}^{(i)}\|^2] \leq 32 E^2 \eta^2 L^2 \|\nabla f_i(\boldsymbol{w}_k)\|^2.$$

The expectation above is with respect to the randomness only due to $\{\mathcal{B}_{k,1}^{(i)}, \ldots, \mathcal{B}_{k,E-1}^{(i)}\}_{i=1}^n$.

Proof. Note that in Lemma 5, we have already bounded $\sum_{\tau=0}^{E-1} \mathbb{E}[\|\boldsymbol{e}_{k,\tau}^{(i)}\|^2]$ (see (61)) – but here we expand it more for use in Lemma 3.
First, from (61), we have:

$$\sum_{\tau=0}^{E-1} \mathbb{E}[\|\boldsymbol{e}_{k,\tau}^{(i)}\|^2] \leq 2EL^2 \sum_{\tau=0}^{E-2} \mathbb{E}[\|\boldsymbol{w}_{k,\tau+1}^{(i)} - \boldsymbol{w}_{k,\tau}^{(i)}\|^2].$$

Next, using the fact that $\boldsymbol{w}_{k,\tau+1}^{(i)} = \boldsymbol{w}_{k,\tau}^{(i)} - \eta \boldsymbol{v}_{k,\tau}^{(i)}$, we get:

$$\sum_{\tau=0}^{E-1} \mathbb{E}[\|\boldsymbol{e}_{k,\tau}^{(i)}\|^2] \leq 2E\eta^2 L^2 \sum_{\tau=0}^{E-2} \mathbb{E}[\|\boldsymbol{v}_{k,\tau}^{(i)}\|^2] \leq 2E\eta^2 L^2 \underbrace{\sum_{\tau=0}^{E-1} \mathbb{E}[\|\boldsymbol{v}_{k,\tau}^{(i)}\|^2]}_{\text{from Lemma 5}} \leq 2E\eta^2 L^2 (16E\|\nabla f_i(\boldsymbol{w}_k)\|^2).$$

This gives us the desired result. ■

Lemma 8. *Suppose* $2\eta L E^2 \leq 1$. *Then:*

$$\mathbb{E}[\|\boldsymbol{u}_k - \overline{\boldsymbol{\delta}}_k\|^2] \leq (1-\beta)^2 \mathbb{E}[\|\boldsymbol{u}_{k-1} - \overline{\boldsymbol{\delta}}_{k-1}\|^2] + 2\beta^2 \mathbb{E}[\|g_Q(\boldsymbol{w}_k; \mathcal{S}_k) - \overline{\boldsymbol{\delta}}_k\|^2] \\ + 8e^2(1+q)(1-\beta)^2 \eta^2 L^2 E^2 (E+1)^2 \mathbb{E}[\|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\|^2].$$

Proof. First, note that for each $i \in [n]$, $\mathbb{E}_{\mathcal{B}_{k,1}^{(i)},\ldots,\mathcal{B}_{k,E-1}^{(i)}}[\boldsymbol{w}_k - \widehat{\boldsymbol{w}}_{k,E}^{(i)}] = \boldsymbol{\delta}_k^{(i)}$. So:

$$\mathbb{E}_{\mathcal{S}_k, \{\mathcal{B}_{k,1}^{(i)},\ldots,\mathcal{B}_{k,E-1}^{(i)}\}_{i=1}^n}[g(\boldsymbol{w}_k; \mathcal{S}_k)] = \overline{\boldsymbol{\delta}}_k. \tag{70}$$

Similarly, for each $i \in [n]$, $\mathbb{E}_{\mathcal{B}_{k,1}^{(i)},\ldots,\mathcal{B}_{k,E-1}^{(i)}}[\boldsymbol{w}_{k-1} - \widehat{\boldsymbol{w}}_{k-1,E}^{(i)}] = \boldsymbol{\delta}_{k-1}^{(i)}$. Hence:

$$\mathbb{E}_{\mathcal{S}_k, \{\mathcal{B}_{k,1}^{(i)},\ldots,\mathcal{B}_{k,E-1}^{(i)}\}_{i=1}^n}[\widehat{g}(\boldsymbol{w}_{k-1}; \mathcal{S}_k)] = \overline{\boldsymbol{\delta}}_{k-1}. \tag{71}$$

We have:

$$\begin{aligned}
\mathbb{E}[\|\boldsymbol{u}_k - \overline{\boldsymbol{\delta}}_k\|^2] &= \mathbb{E}[\|\beta g_Q(\boldsymbol{w}_k; \mathcal{S}_k) + (1-\beta)\boldsymbol{u}_{k-1} + (1-\beta)\Delta g_Q(\boldsymbol{w}_k, \boldsymbol{w}_{k-1}; \mathcal{S}_k) - \overline{\boldsymbol{\delta}}_k\|^2] \\
&= \mathbb{E}[\|(1-\beta)(\boldsymbol{u}_{k-1} - \overline{\boldsymbol{\delta}}_{k-1}) + \beta g_Q(\boldsymbol{w}_k; \mathcal{S}_k) - \overline{\boldsymbol{\delta}}_k + (1-\beta)(\overline{\boldsymbol{\delta}}_{k-1} + \Delta g_Q(\boldsymbol{w}_k, \boldsymbol{w}_{k-1}; \mathcal{S}_k))\|^2] \\
&= (1-\beta)^2 \mathbb{E}[\|\boldsymbol{u}_{k-1} - \overline{\boldsymbol{\delta}}_{k-1}\|^2] + \mathbb{E}[\|\beta g_Q(\boldsymbol{w}_k; \mathcal{S}_k) - \overline{\boldsymbol{\delta}}_k + (1-\beta)(\overline{\boldsymbol{\delta}}_{k-1} + \Delta g_Q(\boldsymbol{w}_k, \boldsymbol{w}_{k-1}; \mathcal{S}_k))\|^2]
\end{aligned} \tag{72}$$

The cross-term in (72) vanishes by taking expectation with respect to Q_D and $\mathcal{S}_k$. Next:

$$\begin{aligned}
&\mathbb{E}[\|\beta g_Q(\boldsymbol{w}_k; \mathcal{S}_k) - \overline{\boldsymbol{\delta}}_k + (1-\beta)(\overline{\boldsymbol{\delta}}_{k-1} + \Delta g_Q(\boldsymbol{w}_k, \boldsymbol{w}_{k-1}; \mathcal{S}_k))\|^2] \\
&= \mathbb{E}[\|\beta(g_Q(\boldsymbol{w}_k; \mathcal{S}_k) - \overline{\boldsymbol{\delta}}_k) + (1-\beta)(\overline{\boldsymbol{\delta}}_{k-1} + \Delta g_Q(\boldsymbol{w}_k, \boldsymbol{w}_{k-1}; \mathcal{S}_k)) - \overline{\boldsymbol{\delta}}_k)\|^2] \\
&\leq 2\beta^2 \mathbb{E}[\|g_Q(\boldsymbol{w}_k; \mathcal{S}_k) - \overline{\boldsymbol{\delta}}_k\|^2] + 2(1-\beta)^2 \mathbb{E}[\|\overline{\boldsymbol{\delta}}_{k-1} + \Delta g_Q(\boldsymbol{w}_k, \boldsymbol{w}_{k-1}; \mathcal{S}_k) - \overline{\boldsymbol{\delta}}_k\|^2]
\end{aligned} \tag{73}$$

Next, note that:

$$\begin{aligned}
&\mathbb{E}[\|\overline{\boldsymbol{\delta}}_{k-1} + \Delta g_Q(\boldsymbol{w}_k, \boldsymbol{w}_{k-1}; \mathcal{S}_k) - \overline{\boldsymbol{\delta}}_k\|^2] \\
&= \mathbb{E}[\|\Delta g_Q(\boldsymbol{w}_k, \boldsymbol{w}_{k-1}; \mathcal{S}_k)\|^2] + \mathbb{E}[\|\overline{\boldsymbol{\delta}}_k - \overline{\boldsymbol{\delta}}_{k-1}\|^2] - 2\mathbb{E}[\langle \Delta g_Q(\boldsymbol{w}_k, \boldsymbol{w}_{k-1}; \mathcal{S}_k), \overline{\boldsymbol{\delta}}_k - \overline{\boldsymbol{\delta}}_{k-1} \rangle] \\
&= \mathbb{E}[\|\Delta g_Q(\boldsymbol{w}_k, \boldsymbol{w}_{k-1}; \mathcal{S}_k)\|^2] + \mathbb{E}[\|\overline{\boldsymbol{\delta}}_k - \overline{\boldsymbol{\delta}}_{k-1}\|^2] - 2\mathbb{E}[\|\overline{\boldsymbol{\delta}}_k - \overline{\boldsymbol{\delta}}_{k-1}\|^2] \qquad (74) \\
&\leq \mathbb{E}[\|\Delta g_Q(\boldsymbol{w}_k, \boldsymbol{w}_{k-1}; \mathcal{S}_k)\|^2]. \qquad (75)
\end{aligned}$$

(74) follows by first taking expectation with respect to Q_D and then using (70) and (71).

Further:

$$\begin{aligned}
\mathbb{E}[\|\Delta g_Q(\boldsymbol{w}_k, \boldsymbol{w}_{k-1}; \mathcal{S}_k)\|^2] &= \mathbb{E}\Big[\Big\|\frac{1}{r} \sum_{i \in \mathcal{S}_k} Q_D((\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)}) - (\boldsymbol{w}_{k-1} - \widehat{\boldsymbol{w}}_{k-1,E}^{(i)}))\Big\|^2\Big] \\
&\leq \mathbb{E}_{\mathcal{S}_k}\Big[\frac{r}{r^2} \sum_{i \in \mathcal{S}_k} \mathbb{E}\Big[\|Q_D((\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)}) - (\boldsymbol{w}_{k-1} - \widehat{\boldsymbol{w}}_{k-1,E}^{(i)}))\|^2\Big]\Big] \\
&\leq \mathbb{E}_{\mathcal{S}_k}\Big[\frac{1}{r} \sum_{i \in \mathcal{S}_k} (1+q)\mathbb{E}\Big[\|(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)}) - (\boldsymbol{w}_{k-1} - \widehat{\boldsymbol{w}}_{k-1,E}^{(i)})\|^2\Big]\Big] \qquad (76) \\
&= \frac{1}{n} \sum_{i \in [n]} (1+q)\mathbb{E}\Big[\|(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)}) - (\boldsymbol{w}_{k-1} - \widehat{\boldsymbol{w}}_{k-1,E}^{(i)})\|^2\Big]. \qquad (77)
\end{aligned}$$

(76) follows from Assumption 3 on the variance of Q_D. Further, using Lemma 9, we get

$$\mathbb{E}[\|(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)}) - (\boldsymbol{w}_{k-1} - \widehat{\boldsymbol{w}}_{k-1,E}^{(i)})\|^2] \leq 4e^2\eta^2 L^2 E^2 (E+1)^2 \mathbb{E}[\|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\|^2], \tag{78}$$

for $2\eta LE^2 \leq 1$.

Using this in (77):

$$\mathbb{E}[\|\Delta g_Q(\boldsymbol{w}_k, \boldsymbol{w}_{k-1}; \mathcal{S}_k)\|^2] \leq 4e^2(1+q)\eta^2 L^2 E^2 (E+1)^2 \mathbb{E}[\|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\|^2]. \tag{79}$$

Now using (79) in (75) and then using it in (73), we get:

$$\begin{aligned}\mathbb{E}[\|\beta g_Q(\boldsymbol{w}_k; \mathcal{S}_k) - \overline{\boldsymbol{\delta}}_k + (1-\beta)(\overline{\boldsymbol{\delta}}_{k-1} + \Delta g_Q(\boldsymbol{w}_k, \boldsymbol{w}_{k-1}; \mathcal{S}_k))\|^2] \\ \leq 2\beta^2 \mathbb{E}[\|g_Q(\boldsymbol{w}_k; \mathcal{S}_k) - \overline{\boldsymbol{\delta}}_k\|^2] + 8e^2(1+q)(1-\beta)^2\eta^2 L^2 E^2(E+1)^2 \mathbb{E}[\|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\|^2]. \end{aligned} \tag{80}$$

Finally, putting (80) back in (72) gives us the desired result. ■

Lemma 9. *Suppose $2\eta LE^2 \leq 1$. Then $\forall\ k \geq 0$ and $i \in [n]$, we have:*

$$\mathbb{E}[\|(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)}) - (\boldsymbol{w}_{k-1} - \widehat{\boldsymbol{w}}_{k-1,E}^{(i)})\|] \leq 2e(\eta LE(E+1))\|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\|.$$

Proof. We have for any $i \in [n]$:

$$\begin{aligned}\|(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)}) - (\boldsymbol{w}_{k-1} - \widehat{\boldsymbol{w}}_{k-1,E}^{(i)})\| &= \Big\| \sum_{\tau=0}^{E-1} \eta \boldsymbol{v}_{k,\tau}^{(i)} - \sum_{\tau=0}^{E-1} \eta \widehat{\boldsymbol{v}}_{k-1,\tau}^{(i)} \Big\| \\ &\leq \sum_{\tau=0}^{E-1} \eta \|\boldsymbol{v}_{k,\tau}^{(i)} - \widehat{\boldsymbol{v}}_{k-1,\tau}^{(i)}\|. \end{aligned} \tag{81}$$

The last step follows by the triangle inequality.
Next, we have:

$$\begin{aligned}\|\boldsymbol{v}_{k,\tau}^{(i)} - \widehat{\boldsymbol{v}}_{k-1,\tau}^{(i)}\| = \|\{\widetilde{\nabla} f_i(\boldsymbol{w}_{k,\tau}^{(i)}; \mathcal{B}_{k,\tau}^{(i)}) + (\boldsymbol{v}_{k,\tau-1}^{(i)} - \widetilde{\nabla} f_i(\boldsymbol{w}_{k,\tau-1}^{(i)}; \mathcal{B}_{k,\tau}^{(i)}))\} \\ - \{\widetilde{\nabla} f_i(\widehat{\boldsymbol{w}}_{k-1,\tau}^{(i)}; \mathcal{B}_{k,\tau}^{(i)}) + (\widehat{\boldsymbol{v}}_{k-1,\tau-1}^{(i)} - \widetilde{\nabla} f_i(\widehat{\boldsymbol{w}}_{k-1,\tau-1}^{(i)}; \mathcal{B}_{k,\tau}^{(i)}))\}\|\end{aligned}$$

Note that $\mathcal{B}_{k,\tau}^{(i)}$ can be the full batch too.

Re-arranging the above, using the triangle inequality and the smoothness of the stochastic gradients, we get:

$$\|\boldsymbol{v}_{k,\tau}^{(i)} - \widehat{\boldsymbol{v}}_{k-1,\tau}^{(i)}\| \leq L\|\boldsymbol{w}_{k,\tau}^{(i)} - \widehat{\boldsymbol{w}}_{k-1,\tau}^{(i)}\| + \|\boldsymbol{v}_{k,\tau-1}^{(i)} - \widehat{\boldsymbol{v}}_{k-1,\tau-1}^{(i)}\| + L\|\boldsymbol{w}_{k,\tau-1}^{(i)} - \widehat{\boldsymbol{w}}_{k-1,\tau-1}^{(i)}\|. \tag{82}$$

Unfolding the above recursion, we get:

$$\|\boldsymbol{v}_{k,\tau}^{(i)} - \widehat{\boldsymbol{v}}_{k-1,\tau}^{(i)}\| \leq 2L \sum_{t=0}^{\tau} \|\boldsymbol{w}_{k,t}^{(i)} - \widehat{\boldsymbol{w}}_{k-1,t}^{(i)}\|. \tag{83}$$

Just as a sanity check for (83), observe that $\|\boldsymbol{v}_{k,0}^{(i)} - \widehat{\boldsymbol{v}}_{k-1,0}^{(i)}\| = \|\nabla f_i(\boldsymbol{w}_k) - \nabla f_i(\boldsymbol{w}_{k-1})\| \leq L\|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\|$. Next:

$$\begin{aligned}\|\boldsymbol{w}_{k,\tau+1}^{(i)} - \widehat{\boldsymbol{w}}_{k-1,\tau+1}^{(i)}\| &= \|\boldsymbol{w}_{k,\tau}^{(i)} - \widehat{\boldsymbol{w}}_{k-1,\tau}^{(i)} - \eta(\boldsymbol{v}_{k,\tau}^{(i)} - \widehat{\boldsymbol{v}}_{k-1,\tau}^{(i)})\| \\ &\leq \|\boldsymbol{w}_{k,\tau}^{(i)} - \widehat{\boldsymbol{w}}_{k-1,\tau}^{(i)}\| + \eta\|\boldsymbol{v}_{k,\tau}^{(i)} - \widehat{\boldsymbol{v}}_{k-1,\tau}^{(i)}\| \\ &\leq \|\boldsymbol{w}_{k,\tau}^{(i)} - \widehat{\boldsymbol{w}}_{k-1,\tau}^{(i)}\| + 2\eta L \sum_{t=0}^{\tau} \|\boldsymbol{w}_{k,t}^{(i)} - \widehat{\boldsymbol{w}}_{k-1,t}^{(i)}\|.\end{aligned}$$

The last step follows by using (83). Thus:

$$\|\boldsymbol{w}_{k,\tau}^{(i)} - \widehat{\boldsymbol{w}}_{k-1,\tau}^{(i)}\| \leq \|\boldsymbol{w}_{k,\tau-1}^{(i)} - \widehat{\boldsymbol{w}}_{k-1,\tau-1}^{(i)}\| + 2\eta L \sum_{t=0}^{\tau-1} \|\boldsymbol{w}_{k,t}^{(i)} - \widehat{\boldsymbol{w}}_{k-1,t}^{(i)}\|. \tag{84}$$

Based on (84), we claim that:

$$\|\boldsymbol{w}_{k,\tau}^{(i)} - \widehat{\boldsymbol{w}}_{k-1,\tau}^{(i)}\| \leq (1+2\eta LE)^\tau \|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\|. \tag{85}$$

We prove this by induction. Let us first examine the base case of $\tau = 1$. We have:

$$\begin{aligned}
\|\boldsymbol{w}_{k,1}^{(i)} - \widehat{\boldsymbol{w}}_{k-1,1}^{(i)}\| &= \|\boldsymbol{w}_k - \boldsymbol{w}_{k-1} - \eta(\boldsymbol{v}_{k,0}^{(i)} - \widehat{\boldsymbol{v}}_{k-1,0}^{(i)})\| \\
&= \|\boldsymbol{w}_k - \boldsymbol{w}_{k-1} - \eta(\nabla f_i(\boldsymbol{w}_k) - \nabla f_i(\boldsymbol{w}_{k-1}))\| \\
&\leq \|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\| + \eta L \|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\| \\
&\leq (1+2\eta LE)^1 \|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\|.
\end{aligned}$$

For ease of notation, let us define $d_k \triangleq \|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\|$. Now suppose the claim is true for $\tau \leq t$. Then using (84), we have for $\tau = t+1$:

$$\begin{aligned}
\|\boldsymbol{w}_{k,t+1}^{(i)} - \widehat{\boldsymbol{w}}_{k-1,t+1}^{(i)}\| &\leq \Big\{(1+2\eta LE)^t + 2\eta L \sum_{t_2=0}^{t} (1+2\eta LE)^{t_2}\Big\} d_k \\
&\leq \Big\{(1+2\eta LE)^t + 2\eta L(t+1)(1+2\eta LE)^t\Big\} d_k \\
&\leq (1+2\eta LE)^t (1+2\eta L(t+1)) d_k \leq (1+2\eta LE)^{t+1} d_k.
\end{aligned} \tag{86}$$

This proves our claim.
Now, using our claim, i.e., (85) in (83), we get:

$$\|\boldsymbol{v}_{k,\tau}^{(i)} - \widehat{\boldsymbol{v}}_{k-1,\tau}^{(i)}\| \leq 2L \sum_{t=0}^{\tau} (1+2\eta LE)^t \|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\| \leq 2L(\tau+1)(1+2\eta LE)^\tau \|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\|. \tag{87}$$

Note that this bound is independent of i.
Finally, using (87) in (81), we get:

$$\begin{aligned}
\|(\boldsymbol{w}_k - \boldsymbol{w}_{k,E}^{(i)}) - (\boldsymbol{w}_{k-1} - \widehat{\boldsymbol{w}}_{k-1,E}^{(i)})\| &\leq \sum_{\tau=0}^{E-1} \eta \|\boldsymbol{v}_{k,\tau}^{(i)} - \widehat{\boldsymbol{v}}_{k-1,\tau}^{(i)}\| \\
&\leq \sum_{\tau=0}^{E-1} 2\eta L(\tau+1)(1+2\eta LE)^\tau \|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\| \\
&\leq 2\eta LE(E+1)(1+2\eta LE)^E \|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\| \\
&\leq 2\eta LE(E+1)e^{2\eta LE^2} \|\boldsymbol{w}_k - \boldsymbol{w}_{k-1}\|.
\end{aligned} \tag{88}$$

The last step follows from the fact that $1 + z \leq e^z \; \forall \; z$.

Finally, setting $2\eta LE^2 \leq 1$ gives us the desired result. ■

Lemma 10. *(V*) in the proof of Lemma 1 can be bounded as:*

$$\text{(V*)} \leq 4\eta^2 E\Big(\frac{q}{n^2} + \frac{(1+q)}{r(n-1)}\Big(1 - \frac{r}{n}\Big)\Big) \sum_{i\in[n]} \sum_{\tau=0}^{E-1} \mathbb{E}[\|\boldsymbol{v}_{k,\tau}^{(i)}\|^2].$$

Proof. We have (V*) = $\mathbb{E}[\|g_Q(\boldsymbol{w}_l; \mathcal{S}_l) - \overline{\boldsymbol{\delta}}_l\|^2]$. Note that:

$$\begin{aligned}
\mathbb{E}[\|g_Q(\boldsymbol{w}_l; \mathcal{S}_l) - \overline{\boldsymbol{\delta}}_l\|^2] \leq \eta^2 \underbrace{\mathbb{E}\Big[\Big\|\frac{1}{r}\sum_{i\in\mathcal{S}_l} \frac{Q_D(\boldsymbol{w}_l - \boldsymbol{w}_{l,E}^{(i)})}{\eta} - \frac{1}{n}\sum_{i\in[n]} \frac{Q_D(\boldsymbol{w}_l - \boldsymbol{w}_{l,E}^{(i)})}{\eta}\Big\|^2\Big]}_{\text{(A)}} \\
+ \eta^2 \underbrace{\mathbb{E}\Big[\Big\|\frac{1}{n}\sum_{i\in[n]} \Big\{\frac{Q_D(\boldsymbol{w}_l - \boldsymbol{w}_{l,E}^{(i)})}{\eta} - \frac{(\boldsymbol{w}_l - \boldsymbol{w}_{l,E}^{(i)})}{\eta}\Big\}\Big\|^2\Big]}_{\text{(B)}}
\end{aligned} \tag{89}$$

In (A), we take expectation with respect to $\mathcal{S}_k$ and $Q_D(.)$ – for that, we use Lemma 4 of Reisizadeh et al. [2020]. Note that $\mathbf{x}_{k,\tau}^{(i)} - \mathbf{x}_k$ in their lemma corresponds to $(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k)$ in our case. Specifically, using eqn. (59) and (60) in Reisizadeh et al. [2020] (they also have Assumption 3), we get:

$$\begin{aligned}
\text{(A)} &\leq \frac{1}{r(n-1)}\Big(1-\frac{r}{n}\Big)4(1+q)\sum_{i\in[n]}\mathbb{E}[\|\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k\|^2] \\
&= \frac{1}{r(n-1)}\Big(1-\frac{r}{n}\Big)4(1+q)\sum_{i\in[n]}\mathbb{E}[\|\sum_{\tau=0}^{E-1}\eta\boldsymbol{v}_{k,\tau}^{(i)}\|^2] \\
&\leq \frac{\eta^2}{r(n-1)}\Big(1-\frac{r}{n}\Big)4(1+q)E\sum_{i\in[n]}\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{v}_{k,\tau}^{(i)}\|^2] \qquad (90)
\end{aligned}$$

Next, we deal with (B). We have:

$$\begin{aligned}
\text{(B)} &= \mathbb{E}\Big[\mathbb{E}_{Q_D}\Big[\Big\|\frac{1}{n}\sum_{i\in[n]}\Big\{Q_D\Big(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k\Big) - \Big(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k\Big)\Big\}\Big\|^2\Big]\Big] \\
&\leq \frac{q}{n^2}\sum_{i\in[n]}\mathbb{E}\Big[\Big\|\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k\Big\|^2\Big] \\
&\leq \frac{qE\eta^2}{n^2}\sum_{i\in[n]}\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{v}_{k,\tau}^{(i)}\|^2]. \qquad (91)
\end{aligned}$$

Now using (90) and (91) in (89), we get:

$$\begin{aligned}
\text{(V*)} &\leq \frac{\eta^2}{r(n-1)}\Big(1-\frac{r}{n}\Big)4(1+q)E\sum_{i\in[n]}\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{v}_{k,\tau}^{(i)}\|^2] + \frac{\eta^2 qE}{n^2}\sum_{i\in[n]}\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{v}_{k,\tau}^{(i)}\|^2] \\
&\leq 4\eta^2 E\Big(\frac{q}{n^2} + \frac{(1+q)}{r(n-1)}\Big(1-\frac{r}{n}\Big)\Big)\sum_{i\in[n]}\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{v}_{k,\tau}^{(i)}\|^2]. \qquad (92)
\end{aligned}$$

This gives us the desired result. ■

Lemma 11. *For any L-smooth function $h(\boldsymbol{x})$, we have $\forall\ \boldsymbol{x}$:*

$$\|\nabla h(\boldsymbol{x})\|^2 \leq 2L(h(\boldsymbol{x}) - h^*) \text{ where } h^* = \min_{\boldsymbol{x}} h(\boldsymbol{x}).$$

Proof. For any y, we have that:

$$h^* \leq h(\boldsymbol{y}) \leq \underbrace{h(\boldsymbol{x}) + \langle\nabla h(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\rangle + \frac{L}{2}\|\boldsymbol{y}-\boldsymbol{x}\|^2}_{:= h_2(\boldsymbol{y})} \qquad (93)$$

Setting $\nabla h_2(\boldsymbol{y}) = \vec{0}$, we get that $\widehat{\boldsymbol{y}} = \boldsymbol{x} - \frac{1}{L}\nabla h(\boldsymbol{x})$ is the minimizer of $h_2(\boldsymbol{y})$ (which is a quadratic with respect to $\boldsymbol{y}$). Plugging this back in (93) gives us:

$$h^* \leq h(\boldsymbol{x}) + \Big\langle\nabla h(\boldsymbol{x}), -\frac{1}{L}\nabla h(\boldsymbol{x})\Big\rangle + \frac{L}{2}\Big\| - \frac{1}{L}\nabla h(\boldsymbol{x})\Big\|^2 = h(\boldsymbol{x}) - \frac{1}{2L}\|\nabla h(\boldsymbol{x})\|^2. \qquad (94)$$

This gives us the desired result. ■

G.2 DETAILED PROOF OF THE RESULT OF FEDLOMO

Let us redefine the quantities needed to prove the results of FedLOMO.

$$\overline{\boldsymbol{w}}_{k,\tau} \triangleq \frac{1}{n}\sum_{i\in[n]} \boldsymbol{w}_{k,\tau}^{(i)} \text{ and } \overline{\boldsymbol{v}}_{k,\tau} \triangleq \frac{1}{n}\sum_{i\in[n]} \boldsymbol{v}_{k,\tau}^{(i)}$$

$$\boldsymbol{e}_{k,\tau}^{(i)} \triangleq \boldsymbol{v}_{k,\tau}^{(i)} - \nabla f_i(\boldsymbol{w}_{k,\tau}^{(i)}) \text{ and } \widetilde{\boldsymbol{e}}_{k,\tau}^{(i)} \triangleq \nabla f_i(\boldsymbol{w}_{k,\tau}^{(i)}) - \nabla f_i(\overline{\boldsymbol{w}}_{k,\tau})$$

Proof of Theorem 3:

Proof. Let us set $\eta_k = \eta$.
Using Lemma 12, with $\eta < \frac{1}{L}$ and $E < \frac{1}{4}\min\Big(\frac{1}{\eta L}, \frac{1}{\eta^2 L^2} - \frac{1}{\eta L}\Big)$:

$$\begin{aligned}\mathbb{E}[f(\boldsymbol{w}_{k+1})] \leq \mathbb{E}[f(\boldsymbol{w}_k)] - \frac{\eta E}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] - \frac{\eta}{2}(1 - \eta^2 L^2 E^2 - \eta L E)\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] \\ + 16\eta L E^2\Big\{\frac{\eta^2 L(\alpha E + 4)}{n^2} + \frac{\eta}{2}\Big(\frac{q}{n^2} + \frac{4(1+q)}{r(n-1)}\Big(1 - \frac{r}{n}\Big)\Big)\Big\}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]. \quad (95)\end{aligned}$$

Note here that for $\eta < \frac{1}{2L}$, $\frac{1}{\eta L} < \frac{1}{\eta^2 L^2} - \frac{1}{\eta L}$ and so $E < \frac{1}{4\eta L}$ or $\eta L E < \frac{1}{4}$. Since $E > 1$, we are just left with $\eta L E < \frac{1}{4}$.

Next, we circumvent the need for the bounded client dissimilarity assumption by using the fact that each f_i is L-smooth and so $\|\nabla f_i(\boldsymbol{w}_k)\|^2 \leq 2L(f_i(\boldsymbol{w}_k) - f_i^*)$ using Lemma 11. Hence:

$$\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2] \leq 2L\sum_{i\in[n]}\mathbb{E}[(f_i(\boldsymbol{w}_k) - f_i^*)] = 2nL\mathbb{E}[(f(\boldsymbol{w}_k) - f^* + \Delta^*)], \quad (96)$$

where $\Delta^* := f^* - \frac{1}{n}\sum_{i=1}^n f_i^*$. Using all this in (95), we get:

$$\begin{aligned}\mathbb{E}[f(\boldsymbol{w}_{k+1})] \leq \mathbb{E}[f(\boldsymbol{w}_k)] - \frac{\eta E}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] - \frac{\eta}{2}\underbrace{(1 - \eta^2 L^2 E^2 - \eta L E)}_{> 0 \text{ for } \eta L E < \frac{1}{4}}\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] \\ + 32\eta L^2 E^2\Big\{\eta^2 L\Big(\frac{\alpha E + 4}{n}\Big) + \frac{\eta}{2}\underbrace{\Big(\frac{q}{n} + \frac{4(1+q)(n-r)}{r(n-1)}\Big)}_{:=B}\Big\}\mathbb{E}[(f(\boldsymbol{w}_k) - f^* + \Delta^*)]. \quad (97)\end{aligned}$$

Note that $(1 - \eta^2 L^2 E^2 - \eta L E) > \frac{11}{16}$ for $\eta L E < \frac{1}{4}$. Further, $-f^* + \Delta^* = -f^* + f^* - \frac{1}{n}\sum_{i=1}^n f_i^* = -\frac{1}{n}\sum_{i=1}^n f_i^*$; hence, we can ignore the corresponding term when the f_i^*'s are non-negative. Re-writing the above equation, we get:

$$\begin{aligned}\mathbb{E}[f(\boldsymbol{w}_{k+1})] &\leq \mathbb{E}[f(\boldsymbol{w}_k)] - \frac{\eta E}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + 32\eta L^2 E^2\Big\{\eta^2 L\Big(\frac{\alpha E + 4}{n}\Big) + \frac{\eta B}{2}\Big\}\mathbb{E}[f(\boldsymbol{w}_k)] \\ &\leq \mathbb{E}[f(\boldsymbol{w}_k)]\Big\{1 + \underbrace{\Big(\frac{32\eta^3 L^3 E^3}{n}\Big(\alpha + \frac{4}{E}\Big) + 16B\eta^2 L^2 E^2\Big)}_{=\zeta}\Big\} - \frac{\eta E}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2]. \quad (98)\end{aligned}$$

Let us denote $\frac{32\eta^3 L^3 E^3}{n}\Big(\alpha + \frac{4}{E}\Big) + 16B\eta^2 L^2 E^2$ as ζ for brevity.
Unfolding the above recursion from $k = 0$ through $K - 1$, we get:

$$\mathbb{E}[f(\boldsymbol{w}_K)] \leq f(\boldsymbol{w}_0)(1+\zeta)^K - \frac{\eta E}{2}\sum_{k=0}^{K-1}(1+\zeta)^{(K-1-k)}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2]. \quad (99)$$

Re-arranging the above, we get:

$$\sum_{k=0}^{K-1} p_k\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] \leq \frac{2}{\eta E}\frac{f(\boldsymbol{w}_0)(1+\zeta)^K}{\sum_{k=0}^{K-1}(1+\zeta)^k}, \text{ where } p_k = \frac{(1+\zeta)^{(K-1-k)}}{\sum_{k=0}^{K-1}(1+\zeta)^k}. \quad (100)$$

Notice that p_k defines a distribution over k. Hence, the LHS is $\mathbb{E}_{k\sim\mathbb{P}(k)}[\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2]]$ with $\mathbb{P}(k) = p_k$. Incorporating this and simplifying further, we get:

$$\mathbb{E}_{k\sim\mathbb{P}(k)}[\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2]] \leq \frac{2}{\eta E}\Big\{\frac{f(\boldsymbol{w}_0)\zeta}{1-(1+\zeta)^{-K}}\Big\}, \text{ where } \mathbb{P}(k) = \frac{(1+\zeta)^{(K-1-k)}}{\sum_{k=0}^{K-1}(1+\zeta)^k}. \tag{101}$$

Also note that: $(1+\zeta)^{-K} < 1 - \zeta K + \zeta^2\frac{K(K+1)}{2} < 1 - \zeta K + \zeta^2 K^2$. Hence, $1-(1+\zeta)^{-K} > \zeta K(1-\zeta K)$. Using this in (101), we have for $\zeta K < 1$:

$$\mathbb{E}_{k\sim\mathbb{P}(k)}[\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2]] \leq \frac{2f(\boldsymbol{w}_0)}{\underbrace{\eta E K(1-\zeta K)}_{=d(\eta)}}, \text{ where } \mathbb{P}(k) = \frac{(1+\zeta)^{(K-1-k)}}{\sum_{k=0}^{K-1}(1+\zeta)^k}. \tag{102}$$

Plugging in the value of ζ in (102), the denominator, $d(\eta) = \eta E K\Big(1 - 16\eta^2L^2E^2\Big(\frac{2\eta LE}{n}\Big(\alpha+\frac{4}{E}\Big)+B\Big)K\Big)$.

Case 1, $q \neq 0$ (compression) and $r < n$ (partial-device participation):
Before going ahead, we would like to highlight that the reason `FedLOMO` does not achieve the optimal rate here is because B is a constant that is not $\mathcal{O}(\eta LE)$.

Let us choose $\eta = \frac{1}{8LE\sqrt{BK}}$. Note that:

$$\eta LE \leq \frac{1}{4} \text{ for } K \geq \frac{1}{4B}. \tag{103}$$

Thus, for sufficiently large K, this choice of η is valid. Also:

$$\zeta K = \frac{1}{4} + \frac{1}{16B^{1.5}\sqrt{K}}\Big(\frac{1}{n}\Big(\alpha+\frac{4}{E}\Big)\Big) < \frac{3}{4} \text{ for } K > \frac{1}{64B^3}\Big(\frac{1}{n}\Big(\alpha+\frac{4}{E}\Big)\Big). \tag{104}$$

So for $K \geq \frac{1}{64B^3}(\frac{1}{n}(\alpha+\frac{4}{E}))$,

$$d(\eta) = \eta E K(1-\zeta K) \geq \frac{\sqrt{K}}{8L\sqrt{B}}(1-\frac{3}{4}) = \frac{\sqrt{K}}{32L\sqrt{B}}. \tag{105}$$

Plugging this in (102), we get:

$$\mathbb{E}_{k\sim\mathbb{P}(k)}[\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2]] \leq \frac{64\sqrt{B}Lf(\boldsymbol{w}_0)}{K^{1/2}}, \text{ where } \mathbb{P}(k) = \frac{(1+\zeta)^{(K-1-k)}}{\sum_{k=0}^{K-1}(1+\zeta)^k} \text{ for } k \in \{0,\ldots,K-1\},$$
$$\zeta = \frac{1}{4K} + \frac{1}{16B^{1.5}K^{1.5}}\Big(\frac{1}{n}\Big(\alpha+\frac{4}{E}\Big)\Big) \text{ and } B = \frac{q}{n} + \frac{4(1+q)(n-r)}{r(n-1)}. \tag{106}$$

Case 2, $q = 0$ (no compression) and $r = n$ (full-device participation):
Here, $B = 0$ and $\zeta = \frac{32\eta^3L^3E^3}{n}\Big(\alpha+\frac{4}{E}\Big)$.

Let us choose $\eta = \frac{1}{4LE}\Big(\frac{n}{(\alpha+\frac{4}{E})K}\Big)^{1/3}$. Now note that $\eta LE \leq \frac{1}{4}$ for $K > \Big(\frac{n}{\alpha+\frac{4}{E}}\Big)$. Also, $\zeta K = \frac{1}{2}$ with our choice of η. Plugging this in (102), we get:

$$\mathbb{E}_{k\sim\mathbb{P}(k)}[\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2]] \leq \frac{16Lf(\boldsymbol{w}_0)}{K^{2/3}}\left(\frac{\alpha+\frac{4}{E}}{n}\right)^{1/3},$$
$$\text{where } \mathbb{P}(k) = \frac{(1+\zeta)^{(K-1-k)}}{\sum_{k=0}^{K-1}(1+\zeta)^k} \text{ for } k \in \{0,\ldots,K-1\} \text{ and } \zeta = \frac{1}{2K}.$$

This concludes the proof. ■

Key lemma used in the proof of Theorem 3:

Lemma 12. *For $\eta_k = \eta$ where $\eta < \frac{1}{L}$ and $E < \frac{1}{4}min\Big(\frac{1}{\eta L}, \frac{1}{\eta^2 L^2} - \frac{1}{\eta L}\Big)$ in* `FedLOMO`, *we have:*

$$\mathbb{E}[f(\boldsymbol{w}_{k+1})] \leq \mathbb{E}[f(\boldsymbol{w}_k)] - \frac{\eta E}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] - \frac{\eta}{2}(1 - \eta^2 L^2 E^2 - \eta L E)\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2]$$
$$+ 16\eta L E^2\Big\{\frac{\eta^2 L(\alpha E + 4)}{n^2} + \frac{\eta}{2}\Big(\frac{q}{n^2} + \frac{4(1+q)}{r(n-1)}\Big(1 - \frac{r}{n}\Big)\Big)\Big\}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2].$$

Proof. By the L-smoothness of f, we have:

$$\mathbb{E}[f(\boldsymbol{w}_{k+1})] \leq \mathbb{E}[f(\boldsymbol{w}_k)] + \underbrace{\mathbb{E}\Big[\Big\langle\nabla f(\boldsymbol{w}_k), \frac{1}{r}\sum_{i\in\mathcal{S}_k} Q_D(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k)\Big\rangle\Big]}_{\text{(I)}} + \underbrace{\frac{L}{2}\mathbb{E}\Big[\Big\|\frac{1}{r}\sum_{i\in\mathcal{S}_k} Q_D(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k)\Big\|^2\Big]}_{\text{(II)}} \tag{107}$$

Let us analyze (I) first – taking expectation with respect to $\mathcal{S}_k$ and $Q_D(.)$ (recall that $Q_D(.)$ is unbiased from Assumption 3), we get:

$$\text{(I)} = \mathbb{E}[\langle\nabla f(\boldsymbol{w}_k), \frac{1}{n}\sum_{i\in[n]}(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k)\rangle]$$

But this is the same as (III*) in the proof of Lemma 1; using Lemma 2 and Assumption 4, we get:

$$\text{(I)} \leq -\frac{\eta E}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] - \frac{\eta}{2}(1 - \eta^2 L^2 E^2)\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] + \frac{16\eta^3 L^2 E^2(\alpha E + 4)}{n^2}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2], \tag{108}$$

when $\eta < \frac{1}{L}$ and $E < \frac{1}{4}\min\Big(\frac{1}{\eta L}, \frac{1}{\eta^2 L^2} - \frac{1}{\eta L}\Big)$.

Let us now analyze (II). Recall that:

$$\text{(II)} = \frac{L}{2}\mathbb{E}\Big[\Big\|\frac{1}{r}\sum_{i\in\mathcal{S}_k} Q_D(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k)\Big\|^2\Big].$$

Observe that:

$$\mathbb{E}_{\mathcal{S}_k}\Big[\frac{1}{r}\sum_{i\in\mathcal{S}_k} Q_D(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k)\Big] = \frac{1}{n}\sum_{i\in[n]} Q_D(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k).$$

Hence:

$$\text{(II)} = \frac{L}{2}\Big\{\underbrace{\mathbb{E}\Big[\Big\|\frac{1}{n}\sum_{i\in[n]} Q_D(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k)\Big\|^2\Big]}_{\text{(III)}}$$
$$+ \underbrace{\mathbb{E}\Big[\Big\|\frac{1}{r}\sum_{i\in\mathcal{S}_k} Q_D(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k) - \frac{1}{n}\sum_{i\in[n]} Q_D(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k)\Big\|^2\Big]}_{\text{(IV)}}\Big\}. \tag{109}$$

Note that in (III), the expectation is without $\mathcal{S}_k$. In (IV), we take expectation with respect to $\mathcal{S}_k$ and $Q_D(.)$ – for that, we use Lemma 4 of Reisizadeh et al. [2020]. Note that $\mathbf{x}_{k,\tau}^{(i)} - \mathbf{x}_k$ in their lemma corresponds to $(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k)$ in our case. Specifically, using eqn. (59) and (60) in Reisizadeh et al. [2020] (they also have Assumption 3), we get:

$$\text{(IV)} \leq \frac{1}{r(n-1)}\Big(1 - \frac{r}{n}\Big)4(1+q)\sum_{i\in[n]}\mathbb{E}[\|\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k\|^2]$$
$$= \frac{1}{r(n-1)}\Big(1 - \frac{r}{n}\Big)4(1+q)\sum_{i\in[n]}\mathbb{E}[\|\sum_{\tau=0}^{E-1}\eta\boldsymbol{v}_{k,\tau}^{(i)}\|^2]$$
$$\leq \frac{\eta^2}{r(n-1)}\Big(1 - \frac{r}{n}\Big)4(1+q)E\sum_{i\in[n]}\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{v}_{k,\tau}^{(i)}\|^2] \tag{110}$$

Next, we deal with (III). Noting that $\mathbb{E}_{Q_D}\Big[\frac{1}{n}\sum_{i\in[n]} Q_D(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k)\Big] = (\overline{\boldsymbol{w}}_{k,E} - \boldsymbol{w}_k)$, we get:

$$\begin{aligned}
\text{(III)} &= \mathbb{E}[\|\overline{\boldsymbol{w}}_{k,E} - \boldsymbol{w}_k\|^2] + \mathbb{E}\Big[\mathbb{E}_{Q_D}\Big[\Big\|\frac{1}{n}\sum_{i\in[n]}\Big\{Q_D\Big(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k\Big) - \Big(\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k\Big)\Big\}\Big\|^2\Big]\Big] \\
&\leq \mathbb{E}\Big[\Big\|\sum_{\tau=0}^{E-1}\eta\overline{\boldsymbol{v}}_{k,\tau}\Big\|^2\Big] + \frac{q}{n^2}\sum_{i\in[n]}\mathbb{E}\Big[\Big\|\boldsymbol{w}_{k,E}^{(i)} - \boldsymbol{w}_k\Big\|^2\Big] \\
&\leq \eta^2 E\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] + \frac{qE\eta^2}{n^2}\sum_{i\in[n]}\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{v}_{k,\tau}^{(i)}\|^2]
\end{aligned} \tag{111}$$

Now, using (110) and (111) in (109) gives us:

$$\begin{aligned}
\text{(II)} &\leq \frac{LE\eta^2}{2}\Big\{\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] + \Big(\frac{q}{n^2} + \frac{4(1+q)}{r(n-1)}\Big(1-\frac{r}{n}\Big)\Big)\sum_{i\in[n]}\underbrace{\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{v}_{k,\tau}^{(i)}\|^2}_{\text{from Lemma 5}}\Big\} \\
&\leq \frac{LE\eta^2}{2}\Big\{\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] + \Big(\frac{q}{n^2} + \frac{4(1+q)}{r(n-1)}\Big(1-\frac{r}{n}\Big)\Big)16E\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]\Big\}.
\end{aligned} \tag{112}$$

Therefore, using (108) and (112) in (107), we get:

$$\begin{aligned}
&\mathbb{E}[f(\boldsymbol{w}_{k+1})] \leq \mathbb{E}[f(\boldsymbol{w}_k)] \\
&\quad - \frac{\eta E}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] - \frac{\eta}{2}(1-\eta^2L^2E^2)\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] + \frac{16\eta^3L^2E^2(\alpha E+4)}{n^2}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2] \\
&\qquad + \frac{LE\eta^2}{2}\Big\{\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] + \Big(\frac{q}{n^2} + \frac{4(1+q)}{r(n-1)}\Big(1-\frac{r}{n}\Big)\Big)16E\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]\Big\}
\end{aligned}$$

$$\begin{aligned}
\implies \mathbb{E}[f(\boldsymbol{w}_{k+1})] &\leq \mathbb{E}[f(\boldsymbol{w}_k)] - \frac{\eta E}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] - \frac{\eta}{2}(1-\eta^2L^2E^2-\eta LE)\sum_{\tau=0}^{E-1}\mathbb{E}[\|\overline{\boldsymbol{v}}_{k,\tau}\|^2] \\
&+ 16\eta LE^2\Big\{\frac{\eta^2L(\alpha E+4)}{n^2} + \frac{\eta}{2}\Big(\frac{q}{n^2} + \frac{4(1+q)}{r(n-1)}\Big(1-\frac{r}{n}\Big)\Big)\Big\}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]
\end{aligned} \tag{113}$$

This completes the proof. ■

H CONVERGENCE OF FEDAVG UNDER ASSUMPTION 4

Here, we provide a convergence result for `FedAvg` (Algorithm 3) in the absence of the bounded client dissimilarity assumption (i.e. eq. (2)) and instead assuming that Assumption 4 holds for `FedAvg`.

Before presenting the convergence result, we show empirical proof that Assumption 4 holds for `FedAvg`. For this, we compute and plot α (as we did in Section 6) for 8 and 4 bit `FedAvg` on CIFAR-10 and FMNIST, respectively; the results are in Figure 4.

Theorem 4 (Smooth non-convex case for `FedAvg`). *Let Assumptions 1 and 2 hold. Further, suppose Assumption 4 holds for `FedAvg` (Algorithm 3). Let σ^2 be the maximum variance of the local (client-level) stochastic gradients. In `FedAvg`, set $\eta_k = \frac{1}{LE}\sqrt{\frac{r}{K}}$ for all k. Define a distribution $\mathbb{P}$ for $k \in \{0,\dots,K-1\}$ such that $\mathbb{P}(k) = \frac{(1+\zeta)^{(K-1-k)}}{\sum_{k=0}^{K-1}(1+\zeta)^k}$ where $\zeta := \eta^2L^2E^2\Big(\frac{(n-r)}{6r(n-1)} + \frac{8\alpha\eta LE}{9n}\Big)$. Sample k^* from $\mathbb{P}$. Then for $K \geq \max\Big(\frac{64r^3}{9}(\frac{\alpha}{n})^2, 4r\Big)$:*

$$\mathbb{E}[\|\nabla f(\boldsymbol{w}_{k^*})\|^2] \leq \frac{4Lf(\boldsymbol{w}_0)}{\sqrt{rK}} + \frac{\sigma^2}{\sqrt{rK}}\Big(\frac{1}{E} + \frac{(n-r)}{3(n-1)}\Big) + \frac{8\sigma^2 r}{9K}\Big(\frac{\alpha}{n}\Big) + \frac{\sigma^2}{EK}\Big(\frac{r}{n}\Big). \tag{114}$$

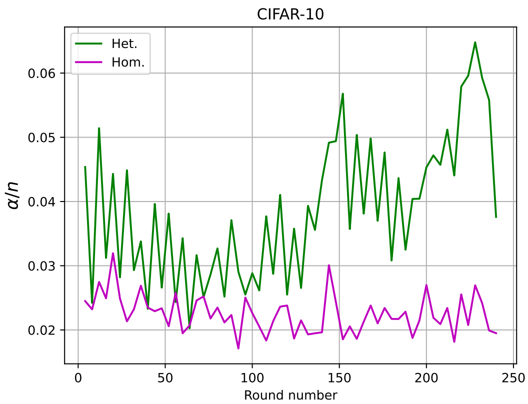

(a) CIFAR-10

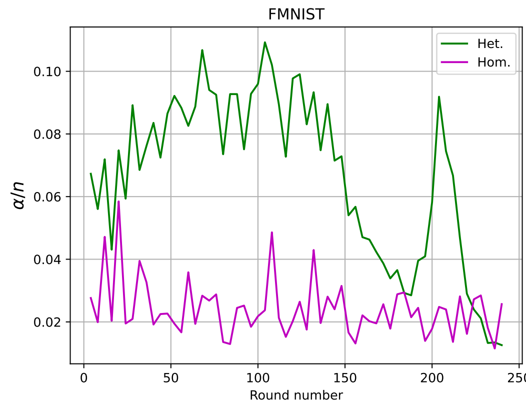

(b) FMNIST

Figure 4: Variation of $\left(\frac{\alpha}{n}\right)$ over different rounds of 8 and 4 bit `FedAvg` for CIFAR-10 (Fig. 4a) and FMNIST (Fig. 4b) in the heterogeneous and homogeneous cases. In both cases, notice that $\alpha \ll n$ throughout training. Also, as expected, observe that $\left(\frac{\alpha}{n}\right)$ is higher for the heterogeneous case (except towards the end of training for FMNIST).

So `FedAvg` needs $K = \mathcal{O}(\frac{1}{r\epsilon^2})$ rounds of communication to achieve $\mathbb{E}[\|\nabla f(\boldsymbol{w}_{k^*})\|^2] \leq \epsilon$, for $\epsilon < \mathcal{O}\big(\max\big(\frac{1}{r}, \frac{n/\alpha}{r^2}\big)\big)$. Note that if we plug in $\alpha = n$ in eq. (114), then we get a convergence result for `FedAvg` without making use of Assumption 4.

Thus, we recover the same complexity for `FedAvg`/Local SGD (which is basically `FedAvg` with full-device participation) as Karimireddy et al. [2019], Koloskova et al. [2020] – but without the bounded client dissimilarity assumption.

Note that the iteration complexity of `FedAvg` is $\mathcal{O}(\epsilon^{-2})$, even with $r = n$. In contrast, note that the iteration complexity of `FedLOMO` improves to $\mathcal{O}(\epsilon^{-1.5})$ with $r = n$ (and no compression) as per Theorem 3.

In the convergence result of Theorem 4, we see that α only shows up in the non-dominant term. So unlike `FedGLOMO`, Assumption 4 holding with $\alpha \ll n$ does not improve the *order-wise* convergence rate/complexity of `FedAvg`.

Proof. Using Lemma 13, for $\eta_k LE \leq \frac{1}{2}$, we can bound the per-round progress as:

$$\mathbb{E}[f(\boldsymbol{w}_{k+1})] \leq \mathbb{E}[f(\boldsymbol{w}_k)] - \frac{\eta_k E}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \eta_k^2 LE^2\Big(\frac{(n-r)}{6r(n-1)} + \frac{8\alpha\eta_k LE}{9n}\Big)\Big(\frac{1}{n}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]\Big) + \frac{\eta_k^2 LE}{2}\Big(\frac{\eta_k LE}{n}\Big(1 + \frac{8\alpha E}{9}\Big) + \frac{1}{r} + \frac{(n-r)E}{3r(n-1)}\Big)\sigma^2. \quad (115)$$

Now applying our earlier trick of using the L-smoothness and non-negativity of the f_i's, we get:

$$\sum_{i\in[n]}\|\nabla f_i(\boldsymbol{w}_k)\|^2 \leq \sum_{i\in[n]} 2L(f_i(\boldsymbol{w}_k) - f_i^*) \leq 2nLf(\boldsymbol{w}_k) - 2L\sum_{i\in[n]} f_i^* \leq 2nLf(\boldsymbol{w}_k).$$

Putting this in eq. (115), we get for a constant learning rate of $\eta_k = \eta$:

$$\mathbb{E}[f(\boldsymbol{w}_{k+1})] \leq \Big(1 + \eta^2 L^2 E^2\Big(\frac{(n-r)}{6r(n-1)} + \frac{8\alpha\eta LE}{9n}\Big)\Big)\mathbb{E}[f(\boldsymbol{w}_k)] - \frac{\eta E}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \frac{\eta^2 LE}{2}\Big(\frac{\eta LE}{n}\Big(1 + \frac{8\alpha E}{9}\Big) + \frac{1}{r} + \frac{(n-r)E}{3r(n-1)}\Big)\sigma^2. \quad (116)$$

For ease of notation, define $\zeta := \eta^2 L^2 E^2\Big(\frac{(n-r)}{6r(n-1)} + \frac{8\alpha\eta LE}{9n}\Big)$ and $\zeta_2 := \Big(\frac{\eta LE}{n}\Big(1 + \frac{8\alpha E}{9}\Big) + \frac{1}{r} + \frac{(n-r)E}{3r(n-1)}\Big)$. Then, unfolding the recursion of eq. (116) from $k = 0$ through to $k = K - 1$, we get:

$$\mathbb{E}[f(\boldsymbol{w}_K)] \leq (1 + \zeta)^K f(\boldsymbol{w}_0) - \frac{\eta E}{2}\sum_{k=0}^{K-1}(1 + \zeta)^{(K-1-k)}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \frac{\eta^2 LE}{2}\zeta_2\sigma^2\sum_{k=0}^{K-1}(1 + \zeta)^{(K-1-k)}. \quad (117)$$

Let us define $p_k := \frac{(1+\zeta)^{(K-1-k)}}{\sum_{k'=0}^{K-1}(1+\zeta)^{(K-1-k')}}$. Then, re-arranging eq. (117) and using the fact that $\mathbb{E}[f(\boldsymbol{w}_K)] \geq 0$, we get:

$$\sum_{k=0}^{K-1} p_k \mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] \leq \frac{2(1+\zeta)^K f(\boldsymbol{w}_0)}{\eta E \sum_{k'=0}^{K-1}(1+\zeta)^{k'}} + \eta L \zeta_2 \sigma^2 \tag{118}$$

$$= \frac{2\zeta f(\boldsymbol{w}_0)}{\eta E(1-(1+\zeta)^{-K})} + \eta LE\Big(\frac{\eta L}{n}\Big(1+\frac{8\alpha E}{9}\Big) + \frac{1}{rE} + \frac{(n-r)}{3r(n-1)}\Big)\sigma^2, \tag{119}$$

where the last step follows by using the fact that $\sum_{k'=0}^{K-1}(1+\zeta)^{k'} = \frac{(1+\zeta)^K - 1}{\zeta}$ and plugging in the value of ζ_2. Now as we did in the proof of Theorem 3:

$$(1+\zeta)^{-K} < 1 - \zeta K + \zeta^2 \frac{K(K+1)}{2} < 1 - \zeta K + \zeta^2 K^2 \implies 1-(1+\zeta)^{-K} > \zeta K(1-\zeta K).$$

Plugging this in eq. (119), we have for $\zeta K < 1$:

$$\sum_{k=0}^{K-1} p_k \mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] \leq \frac{2f(\boldsymbol{w}_0)}{\eta EK(1-\zeta K)} + \eta LE\Big(\frac{\eta L}{n}\Big(1+\frac{8\alpha E}{9}\Big) + \frac{1}{rE} + \frac{(n-r)}{3r(n-1)}\Big)\sigma^2. \tag{120}$$

In this case, note that the optimal step size will be $\eta = \mathcal{O}(\frac{1}{LE\sqrt{K}})$, even for $r = n$. This is in contrast to `FedLOMO` for which the optimal step size is $\eta = \mathcal{O}(\frac{1}{LEK^{1/3}})$ for $r = n$.

So let us pick $\eta = \frac{1}{LE}\sqrt{\frac{r}{K}}$. Note that we need to have $\eta LE \leq \frac{1}{2}$; this happens for $K \geq 4r$. Further, let us ensure $\zeta K < \frac{1}{2}$; this happens for $K \geq \frac{64r^3}{9}(\frac{\alpha}{n})^2$. Thus, we should have $K \geq \max\Big(\frac{64r^3}{9}(\frac{\alpha}{n})^2, 4r\Big)$. Putting $\eta = \frac{1}{LE}\sqrt{\frac{r}{K}}$ in eq. (120) and also using $1 - \zeta K \geq \frac{1}{2}$, we get:

$$\sum_{k=0}^{K-1} p_k \mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] \leq \frac{4Lf(\boldsymbol{w}_0)}{\sqrt{rK}} + \frac{\sigma^2}{\sqrt{rK}}\Big(\frac{1}{E} + \frac{(n-r)}{3(n-1)}\Big) + \frac{8\sigma^2 r}{9K}\Big(\frac{\alpha}{n}\Big) + \frac{\sigma^2}{EK}\Big(\frac{r}{n}\Big). \tag{121}$$

This finishes the proof. ■

Lemma 13. *For $\eta_k LE \leq \frac{1}{2}$, we have:*

$$\mathbb{E}[f(\boldsymbol{w}_{k+1})] \leq \mathbb{E}[f(\boldsymbol{w}_k)] - \frac{\eta_k E}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \eta_k^2 LE^2\Big(\frac{(n-r)}{6r(n-1)} + \frac{8\alpha\eta_k LE}{9n}\Big)\Big(\frac{1}{n}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]\Big)$$
$$+ \frac{\eta_k^2 LE}{2}\Big(\frac{\eta_k LE}{n}\Big(1+\frac{8\alpha E}{9}\Big) + \frac{1}{r} + \frac{(n-r)E}{3r(n-1)}\Big)\sigma^2.$$

Proof. Define

$$\widehat{\boldsymbol{u}}_{k,\tau}^{(i)} := \nabla \widetilde{f}_i(\boldsymbol{w}_{k,\tau}^{(i)}; \mathcal{B}_{k,\tau}^{(i)}),\ \widehat{\boldsymbol{u}}_{k,\tau} := \frac{1}{n}\sum_{i\in[n]}\widehat{\boldsymbol{u}}_{k,\tau}^{(i)},\ \boldsymbol{u}_{k,\tau} := \frac{1}{n}\sum_{i\in[n]}\nabla f_i(\boldsymbol{w}_{k,\tau}^{(i)}),$$

$$\overline{\boldsymbol{w}}_{k,\tau} := \frac{1}{n}\sum_{i\in[n]}\boldsymbol{w}_{k,\tau}^{(i)} \text{ and } \widetilde{\boldsymbol{e}}_{k,\tau}^{(i)} = \nabla f_i(\boldsymbol{w}_{k,\tau}^{(i)}) - \nabla f_i(\overline{\boldsymbol{w}}_{k,\tau}).$$

Then:

$$\boldsymbol{w}_{k+1} = \boldsymbol{w}_k - \eta_k \sum_{\tau=0}^{E-1}\Big(\frac{1}{r}\sum_{i\in\mathcal{S}_k}\widehat{\boldsymbol{u}}_{k,\tau}^{(i)}\Big). \tag{122}$$

$$\overline{\boldsymbol{w}}_{k,\tau} = \boldsymbol{w}_k - \eta_k \sum_{t=0}^{\tau-1}\widehat{\boldsymbol{u}}_{k,t}. \tag{123}$$

$$\mathbb{E}_{\{\mathcal{B}_{k,\tau}^{(i)}\}_{i=1}^n}[\widehat{\boldsymbol{u}}_{k,\tau}] = \boldsymbol{u}_{k,\tau}. \tag{124}$$

$$\mathbb{E}\Big[\Big\|\sum_{t=0}^{\tau-1}\widehat{\boldsymbol{u}}_{k,t}\Big\|^2\Big] \leq \tau\sum_{t=0}^{\tau-1}\mathbb{E}[\|\boldsymbol{u}_{k,t}\|^2] + \frac{\tau\sigma^2}{n}. \tag{125}$$

$$\mathbb{E}\Big[\Big\|\sum_{t=0}^{\tau-1}\widehat{\boldsymbol{u}}_{k,t}^{(i)}\Big\|^2\Big] \leq \tau\sum_{t=0}^{\tau-1}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,t}^{(i)})\|^2] + \tau\sigma^2. \tag{126}$$

Recall that σ^2 is the maximum variance of the local (client-level) stochastic gradients. In eq. (125), the expectation is w.r.t. $\{\mathcal{B}_{k,t}^{(i)}\}_{i=1,t=0}^{n,\tau-1}$ and it follows due to the independence of the noise in each local update of each client. Similarly, eq. (126), the expectation is w.r.t. $\{\mathcal{B}_{k,t}^{(i)}\}_{t=0}^{\tau-1}$ and it follows due to the independence of the noise in each local update.

Next, using the L-smoothness of f and eq. (122), we get

$$\mathbb{E}[f(\boldsymbol{w}_{k+1})] \leq \mathbb{E}[f(\boldsymbol{w}_k)] - \mathbb{E}\Big[\Big\langle\nabla f(\boldsymbol{w}_k), \eta_k\sum_{\tau=0}^{E-1}\Big(\frac{1}{r}\sum_{i\in\mathcal{S}_k}\widehat{\boldsymbol{u}}_{k,\tau}^{(i)}\Big)\Big\rangle\Big] + \frac{L}{2}\mathbb{E}\Big[\Big\|\eta_k\sum_{\tau=0}^{E-1}\Big(\frac{1}{r}\sum_{i\in\mathcal{S}_k}\widehat{\boldsymbol{u}}_{k,\tau}^{(i)}\Big)\Big\|^2\Big] \tag{127}$$

$$= \mathbb{E}[f(\boldsymbol{w}_k)] - \mathbb{E}[\langle\nabla f(\boldsymbol{w}_k), \sum_{\tau=0}^{E-1}\eta_k\widehat{\boldsymbol{u}}_{k,\tau}\rangle] + \frac{\eta_k^2L}{2}\Big\{\frac{n(r-1)}{r(n-1)}\mathbb{E}[\|\sum_{\tau=0}^{E-1}\widehat{\boldsymbol{u}}_{k,\tau}\|^2] + \frac{(n-r)}{r(n-1)}\Big(\frac{1}{n}\sum_{i\in[n]}\mathbb{E}[\|\sum_{\tau=0}^{E-1}\widehat{\boldsymbol{u}}_{k,\tau}^{(i)}\|^2]\Big)\Big\} \tag{128}$$

$$\leq \mathbb{E}[f(\boldsymbol{w}_k)] - \eta_k\mathbb{E}[\langle\nabla f(\boldsymbol{w}_k), \sum_{\tau=0}^{E-1}\boldsymbol{u}_{k,\tau}\rangle] + \frac{\eta_k^2LE}{2}\Big\{\frac{n(r-1)}{r(n-1)}\Big(\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{u}_{k,\tau}\|^2] + \frac{\sigma^2}{n}\Big) \tag{129}$$

$$+ \frac{(n-r)}{r(n-1)}\Big(\frac{1}{n}\sum_{i\in[n]}\sum_{\tau=0}^{E-1}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,\tau}^{(i)})\|^2] + \sigma^2\Big)\Big\}$$

Note that eq. (128) follows by taking expectation w.r.t. $\mathcal{S}_k$ in eq. (127), while eq. (129) follows from eq. (124), eq. (125) and eq. (126).

For any 2 vectors $\boldsymbol{a}$ and $\boldsymbol{b}$, we have that $\langle\boldsymbol{a},\boldsymbol{b}\rangle = \frac{1}{2}(\|\boldsymbol{a}\|^2 + \|\boldsymbol{b}\|^2 - \|\boldsymbol{a}-\boldsymbol{b}\|^2)$. Using this:

$$\langle\nabla f(\boldsymbol{w}_k), \sum_{\tau=0}^{E-1}\boldsymbol{u}_{k,\tau}\rangle = \sum_{\tau=0}^{E-1}\langle\nabla f(\boldsymbol{w}_k), \boldsymbol{u}_{k,\tau}\rangle = \frac{1}{2}\sum_{\tau=0}^{E-1}(\|\nabla f(\boldsymbol{w}_k)\|^2 + \|\boldsymbol{u}_{k,\tau}\|^2 - \|\nabla f(\boldsymbol{w}_k) - \boldsymbol{u}_{k,\tau}\|^2). \tag{130}$$

Putting this in eq. (129), we get:

$$\mathbb{E}[f(\boldsymbol{w}_{k+1})] \leq \mathbb{E}[f(\boldsymbol{w}_k)] - \frac{\eta_kE}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] - \frac{\eta_k}{2}\Big(1 - \eta_kLE\frac{n(r-1)}{r(n-1)}\Big)\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{u}_{k,\tau}\|^2]$$
$$+ \frac{\eta_k}{2}\underbrace{\sum_{\tau=0}^{E-1}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k) - \boldsymbol{u}_{k,\tau}\|^2]}_{\text{(A)}} + \frac{\eta_k^2LE}{2r}\sigma^2 + \frac{(n-r)}{r(n-1)}\frac{\eta_k^2LE}{2}\underbrace{\Big(\frac{1}{n}\sum_{i\in[n]}\sum_{\tau=0}^{E-1}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,\tau}^{(i)})\|^2]\Big)}_{\text{(B)}}. \tag{131}$$

We upper bound (A) and (B) using Lemma 14 and Lemma 15, respectively. Plugging in these bounds, we get:

$$\mathbb{E}[f(\boldsymbol{w}_{k+1})] \leq \mathbb{E}[f(\boldsymbol{w}_k)] - \frac{\eta_kE}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] - \frac{\eta_k}{2}\underbrace{\Big(1 - \eta_kLE\frac{n(r-1)}{r(n-1)} - \eta_k^2L^2E^2\Big)}_{\text{(C)}}\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{u}_{k,\tau}\|^2]$$
$$+ \eta_k^2LE^2\Big(\frac{(n-r)}{6r(n-1)} + \frac{8\alpha\eta_kLE}{9n}\Big)\Big(\frac{1}{n}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]\Big) + \frac{\eta_k^2LE}{2}\Big(\frac{\eta_kLE}{n}\Big(1 + \frac{8\alpha E}{9}\Big) + \frac{1}{r} + \frac{(n-r)E}{3r(n-1)}\Big)\sigma^2, \tag{132}$$

for $\eta_k LE \leq \frac{1}{2}$. Note that (C) ≥ 0 for $\eta_k LE \leq \frac{1}{2}$. Thus, for $\eta_k LE \leq \frac{1}{2}$, we have:

$$\mathbb{E}[f(\boldsymbol{w}_{k+1})] \leq \mathbb{E}[f(\boldsymbol{w}_k)] - \frac{\eta_k E}{2}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)\|^2] + \eta_k^2 LE^2\Big(\frac{(n-r)}{6r(n-1)} + \frac{8\alpha\eta_k LE}{9n}\Big)\Big(\frac{1}{n}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]\Big) + \frac{\eta_k^2 LE}{2}\Big(\frac{\eta_k LE}{n}\Big(1+\frac{8\alpha E}{9}\Big) + \frac{1}{r} + \frac{(n-r)E}{3r(n-1)}\Big)\sigma^2. \quad (133)$$

■

Lemma 14. *For $\eta_k LE \leq \frac{1}{2}$:*

$$\sum_{\tau=0}^{E-1}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k) - \boldsymbol{u}_{k,\tau}\|^2] \leq \eta_k^2 L^2 E^2 \sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{u}_{k,\tau}\|^2] + \frac{16\alpha\eta_k^2 L^2 E^3}{9n^2}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2] + \frac{\eta_k^2 L^2 E^2}{n}\Big(1+\frac{8\alpha E}{9}\Big)\sigma^2.$$

Proof. We have:

$$\begin{aligned}
\mathbb{E}[\|\nabla f(\boldsymbol{w}_k) - \boldsymbol{u}_{k,\tau}\|^2] &= \mathbb{E}[\|\nabla f(\boldsymbol{w}_k) - \nabla f(\overline{\boldsymbol{w}}_{k,\tau}) + \nabla f(\overline{\boldsymbol{w}}_{k,\tau}) - \boldsymbol{u}_{k,\tau}\|^2] &(134)\\
&\leq 2\mathbb{E}[\|\nabla f(\boldsymbol{w}_k) - \nabla f(\overline{\boldsymbol{w}}_{k,\tau})\|^2] + 2\mathbb{E}[\|\nabla f(\overline{\boldsymbol{w}}_{k,\tau}) - \boldsymbol{u}_{k,\tau}\|^2] &(135)\\
&\leq 2L^2\mathbb{E}[\|\boldsymbol{w}_k - \overline{\boldsymbol{w}}_{k,\tau}\|^2] + 2\mathbb{E}\Big[\Big\|\frac{1}{n}\sum_{i\in[n]}\underbrace{(\nabla f_i(\overline{\boldsymbol{w}}_{k,\tau}) - \nabla f_i(\boldsymbol{w}_{k,\tau}^{(i)}))}_{=-\widetilde{\boldsymbol{e}}_{k,\tau}^{(i)}}\Big\|^2\Big] &(136)\\
&\leq 2\eta_k^2 L^2\mathbb{E}\Big[\Big\|\sum_{t=0}^{\tau-1}\widehat{\boldsymbol{u}}_{k,t}\Big\|^2\Big] + \frac{2\alpha}{n^2}\sum_{i\in[n]}\mathbb{E}[\|\widetilde{\boldsymbol{e}}_{k,\tau}^{(i)}\|^2] &(137)\\
&\leq 2\eta_k^2 L^2\Big(\tau\sum_{t=0}^{\tau-1}\mathbb{E}[\|\boldsymbol{u}_{k,t}\|^2] + \frac{\tau\sigma^2}{n}\Big) + \frac{2\alpha L^2}{n^2}\sum_{i\in[n]}\mathbb{E}[\|\boldsymbol{w}_{k,\tau}^{(i)} - \overline{\boldsymbol{w}}_{k,\tau}\|^2]. &(138)
\end{aligned}$$

Equation (136) follows from the L-smoothness of f and the definition of $\boldsymbol{u}_{k,\tau}$. Equation (137) follows from eq. (123) and Assumption 4. Equation (138) follows from eq. (125) and the L-smoothness of f_i.

But:

$$\begin{aligned}
\sum_{i\in[n]}\mathbb{E}[\|\boldsymbol{w}_{k,\tau}^{(i)} - \overline{\boldsymbol{w}}_{k,\tau}\|^2] &= \sum_{i\in[n]}\mathbb{E}[\|(\boldsymbol{w}_{k,0}^{(i)} - \eta_k\sum_{t=0}^{\tau-1}\widehat{\boldsymbol{u}}_{k,t}^{(i)}) - (\overline{\boldsymbol{w}}_{k,0} - \eta_k\sum_{t=0}^{\tau-1}\widehat{\boldsymbol{u}}_{k,t})\|^2] &(139)\\
&= \eta_k^2\sum_{i\in[n]}\mathbb{E}[\|\sum_{t=0}^{\tau-1}\widehat{\boldsymbol{u}}_{k,t} - \sum_{t=0}^{\tau-1}\widehat{\boldsymbol{u}}_{k,t}^{(i)}\|^2] &(140)\\
&\leq \eta_k^2\tau\sum_{i\in[n]}\sum_{t=0}^{\tau-1}\mathbb{E}[\|\widehat{\boldsymbol{u}}_{k,t} - \widehat{\boldsymbol{u}}_{k,t}^{(i)}\|^2] &(141)\\
&= \eta_k^2\tau\sum_{t=0}^{\tau-1}\sum_{i\in[n]}\mathbb{E}[\|\widehat{\boldsymbol{u}}_{k,t}\|^2 + \|\widehat{\boldsymbol{u}}_{k,t}^{(i)}\|^2 - 2\langle\widehat{\boldsymbol{u}}_{k,t}, \widehat{\boldsymbol{u}}_{k,t}^{(i)}\rangle] &(142)
\end{aligned}$$

Equation (140) follows because $\boldsymbol{w}_{k,0}^{(i)} = \boldsymbol{w}_k\ \forall\ i \in [n]$, due to which $\overline{\boldsymbol{w}}_{k,0} = \boldsymbol{w}_k$. Next, using the fact that $\widehat{\boldsymbol{u}}_{k,\tau} =$

$\frac{1}{n}\sum_{i\in[n]}\widehat{\boldsymbol{u}}_{k,\tau}^{(i)}$, we can simplify eq. (142) to:

$$\sum_{i\in[n]}\mathbb{E}[\|\boldsymbol{w}_{k,\tau}^{(i)}-\overline{\boldsymbol{w}}_{k,\tau}\|^2] \leq \eta_k^2\tau\sum_{t=0}^{\tau-1}\sum_{i\in[n]}(\mathbb{E}[\|\widehat{\boldsymbol{u}}_{k,\tau}^{(i)}\|^2]-\mathbb{E}[\|\widehat{\boldsymbol{u}}_{k,t}\|^2]) \tag{143}$$

$$\leq \eta_k^2\tau\sum_{t=0}^{\tau-1}\sum_{i\in[n]}\mathbb{E}[\|\widehat{\boldsymbol{u}}_{k,\tau}^{(i)}\|^2] \tag{144}$$

$$\leq \eta_k^2\tau\sum_{t=0}^{\tau-1}\sum_{i\in[n]}(\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,t}^{(i)})\|^2]+\sigma^2). \tag{145}$$

Next, using Lemma 15 for $\eta_k LE \leq \frac{1}{2}$ in eq. (145), we get:

$$\sum_{i\in[n]}\mathbb{E}[\|\boldsymbol{w}_{k,\tau}^{(i)}-\overline{\boldsymbol{w}}_{k,\tau}\|^2] \leq \frac{4\eta_k^2\tau^2}{3}\sum_{i\in[n]}(2\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]+\sigma^2). \tag{146}$$

Plugging eq. (146) back in eq. (138), we get:

$$\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)-\boldsymbol{u}_{k,\tau}\|^2] \leq 2\eta_k^2L^2\Big(\tau\sum_{t=0}^{\tau-1}\mathbb{E}[\|\boldsymbol{u}_{k,t}\|^2]+\frac{\tau\sigma^2}{n}\Big)+\frac{8\alpha\eta_k^2L^2\tau^2}{3n^2}\sum_{i\in[n]}(2\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]+\sigma^2) \tag{147}$$

$$= 2\eta_k^2L^2\tau\sum_{t=0}^{\tau-1}\mathbb{E}[\|\boldsymbol{u}_{k,t}\|^2]+\frac{16\alpha\eta_k^2L^2\tau^2}{3n^2}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]+\frac{\eta_k^2L^2\tau\sigma^2}{n}\Big(2+\frac{8\alpha}{3}\tau\Big). \tag{148}$$

Summing up eq. (148) for $\tau\in\{0,\ldots,E-1\}$, we get:

$$\sum_{\tau=0}^{E-1}\mathbb{E}[\|\nabla f(\boldsymbol{w}_k)-\boldsymbol{u}_{k,\tau}\|^2] \leq \eta_k^2L^2E^2\sum_{\tau=0}^{E-1}\mathbb{E}[\|\boldsymbol{u}_{k,\tau}\|^2]+\frac{16\alpha\eta_k^2L^2E^3}{9n^2}\sum_{i\in[n]}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]+\frac{\eta_k^2L^2E^2}{n}\Big(1+\frac{8\alpha E}{9}\Big)\sigma^2. \tag{149}$$

■

Lemma 15. *For $\eta_k LE \leq \frac{1}{2}$, we have:*

$$\sum_{t=0}^{\tau-1}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,t}^{(i)})\|^2] \leq \frac{\tau}{3}(8\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]+\sigma^2).$$

Proof.

$$\begin{aligned}\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,t}^{(i)})\|^2] &= \mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,t}^{(i)})-\nabla f_i(\boldsymbol{w}_k)+\nabla f_i(\boldsymbol{w}_k)\|^2]\\ &\leq 2\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]+2\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,t}^{(i)})-\nabla f_i(\boldsymbol{w}_k)\|^2]\\ &\leq 2\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]+2L^2\mathbb{E}[\|\boldsymbol{w}_{k,t}^{(i)}-\boldsymbol{w}_k\|^2]. \end{aligned} \tag{150}$$

But:

$$\mathbb{E}[\|\boldsymbol{w}_k-\boldsymbol{w}_{k,t}^{(i)}\|^2] = \mathbb{E}\Big[\Big\|\eta_k\sum_{t'=0}^{t-1}\nabla\widetilde{f}_i(\boldsymbol{w}_{k,t'}^{(i)};\mathcal{B}_{k,t'}^{(i)})\Big\|^2\Big] \leq \eta_k^2t\sum_{t'=0}^{t-1}\mathbb{E}\Big[\|\nabla\widetilde{f}_i(\boldsymbol{w}_{k,t'}^{(i)};\mathcal{B}_{k,t'}^{(i)})\|^2\Big] \leq \eta_k^2t\sum_{t'=0}^{t-1}(\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,t'}^{(i)})\|^2]+\sigma^2). \tag{151}$$

Putting this back in eq. (150), we get:

$$\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,t}^{(i)})\|^2] \leq 2\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]+2\eta_k^2L^2t\sum_{t'=0}^{t-1}(\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,t'}^{(i)})\|^2]+\sigma^2). \tag{152}$$

Now summing up eq. (152) for all $t \in \{0, \dots, \tau - 1\}$, we get:

$$\sum_{t=0}^{\tau-1} \mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,t}^{(i)})\|^2] \leq 2\tau(\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2]) + 2\eta_k^2 L^2 \sum_{t=0}^{\tau-1} \tau \sum_{t'=0}^{t-1} (\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,t'}^{(i)})\|^2] + \sigma^2)$$

$$\leq 2\tau(\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2) + \eta_k^2 L^2 \tau^2 \sum_{t=0}^{\tau-1} (\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,t}^{(i)})\|^2] + \sigma^2). \tag{153}$$

Let us set $\eta_k LE \leq 1/2$. Then:

$$\sum_{t=0}^{\tau-1} \mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,t}^{(i)})\|^2] \leq 2\tau(\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2) + \frac{1}{4} \sum_{t=0}^{\tau-1} \mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,t}^{(i)})\|^2] + \frac{\sigma^2 \tau}{4}.$$

Simplifying, we get:

$$\sum_{t=0}^{\tau-1} \mathbb{E}[\|\nabla f_i(\boldsymbol{w}_{k,t}^{(i)})\|^2] \leq \frac{\tau}{3}(8\mathbb{E}[\|\nabla f_i(\boldsymbol{w}_k)\|^2] + \sigma^2). \tag{154}$$

■