

Demystifying Why Local Aggregation Helps: Convergence Analysis of Hierarchical SGD

Jiayi Wang¹, Shiqiang Wang², Rong-Rong Chen¹, Mingyue Ji¹

¹ Department of Electrical & Computer Engineering, University of Utah, Salt Lake City, UT, USA

² IBM T. J. Watson Research Center, Yorktown Heights, NY, USA

jiayi.wang@utah.edu, wangshiq@us.ibm.com, rchen@ece.utah.edu, mingyue.ji@utah.edu

Abstract

Hierarchical SGD (H-SGD) has emerged as a new distributed SGD algorithm for multi-level communication networks. In H-SGD, before each global aggregation, workers send their updated local models to local servers for aggregations. Despite recent research efforts, the effect of local aggregation on global convergence still lacks theoretical understanding. In this work, we first introduce a new notion of “upward” and “downward” divergences. We then use it to conduct a novel analysis to obtain a worst-case convergence upper bound for two-level H-SGD with non-IID data, non-convex objective function, and stochastic gradient. By extending this result to the case with random grouping, we observe that this convergence upper bound of H-SGD is between the upper bounds of two single-level local SGD settings, with the number of local iterations equal to the local and global update periods in H-SGD, respectively. We refer to this as the “sandwich behavior”. Furthermore, we extend our analytical approach based on “upward” and “downward” divergences to study the convergence for the general case of H-SGD with more than two levels, where the “sandwich behavior” still holds. Our theoretical results provide key insights of why local aggregation can be beneficial in improving the convergence of H-SGD.

1 Introduction

Stochastic gradient descent (SGD) is a widely used optimization technique in machine learning applications. In distributed SGD, all n workers collaboratively learn a global model \mathbf{w} by minimizing the empirical loss with their local data:

$$\min_{\mathbf{w} \in \mathbb{R}^d} f(\mathbf{w}) := \frac{1}{n} \sum_{j=1}^n F_j(\mathbf{w}), \quad (1)$$

where $F_j(\cdot)$ is the local loss function of worker j . Traditionally, workers send their models or gradients to the central server after one local iteration, which is inefficient due to frequent model aggregations. This may cost high communication latency in practice. Recently, a form of distributed SGD was proposed to reduce communication cost by allowing multiple local iterations during one communication round (McMahan et al. 2017). This is referred to as local SGD (Stich 2019; Lin et al. 2020). However, since data on workers can

be non-IID, to avoid model divergence, the number of local iterations P cannot be too large, which limits its advantage of reducing communication cost.

In practical scenarios, networks often have a hierarchical structure in nature, such as edge computing systems (Li, Ota, and Dong 2018) and software defined networks (SDN) (Kim and Feamster 2013). In these networks, workers often directly communicate with their local server rather than global server. Motivated by this hierarchical structure, a few works (Castiglia, Das, and Patterson 2021; Liu et al. 2020) proposed hierarchical SGD (H-SGD). In H-SGD, workers are partitioned into N groups, where each group has a local server. To reduce communication cost, workers send their models to local servers to do several local aggregations before communicating with global server. They perform I local iterations (referred to as local period) between local aggregations, and G (referred to as global period) local iterations ($G > I$) between global aggregations. In this paper, we call the connection between local servers and workers as “downward” network and the connection between local servers and the global server as “upward” network.

Recently, there have been a few works analyzing the convergence behavior of H-SGD. A structure where the upward network is a peer-to-peer network while downward network is a server-worker network was considered by Castiglia, Das, and Patterson (2021). However, it only considers IID data. The work by Liu et al. (2020) considers non-IID data but it uses full-batch (non-stochastic) gradient descent and the convergence bound is an exponential function of G . In both works, a comprehensive comparison between H-SGD and local SGD is missing and none of them analyzes the effect of local aggregation on overcoming data heterogeneity, which is the key merit of H-SGD. Therefore, a more general and tighter analysis for H-SGD with non-convex objective function, non-IID data and stochastic gradient descent is needed. The effect of local aggregation on overcoming data heterogeneity lacks theoretical understanding.

In this paper, to provide a better theoretical understanding to H-SGD, we devise a novel characterization of the data heterogeneity of H-SGD by “upward” divergence and “downward” divergence, respectively. Furthermore, we show that the global divergence can be partitioned into upward and downward divergences. H-SGD can be seen as performing distributed SGD both within a group and across groups.

Table 1: A summary of convergence bounds in the literature.

Paper	Convergence Bound	Non-IID	SGD	Type	Assumption
Yu, Jin, and Yang (2019)	$O\left(\frac{1+\sigma^2}{\sqrt{nT}} + \frac{n}{T}(P\sigma^2 + P^2\tilde{\epsilon}^2)\right)$	✓	✓	Local SGD	$N = 1$
Liu et al. (2020)	$O\left(\frac{1+B^G\tilde{\epsilon}^2}{\sqrt{nT}}\right)$	✓	✗	H-SGD	$\sigma^2 = 0$
Castiglia, Das, and Patterson (2021)	$O\left(\frac{1+\sigma^2}{\sqrt{nT}} + \frac{n}{T}\frac{G^2}{I}\sigma^2\right)$	✗	✓	H-SGD	$\tilde{\epsilon}^2 = 0$
Ours	$O\left(\frac{1+\sigma^2}{\sqrt{nT}} + \frac{(N-1)(G\sigma^2 + G^2\tilde{\epsilon}^2) + (n-N)(I\sigma^2 + I^2\tilde{\epsilon}^2)}{T}\right)$	✓	✓	H-SGD	None

¹ P : aggregation period in local SGD; G : global aggregation period in H-SGD; I : local aggregation period in H-SGD; N : number of groups in H-SGD; $\tilde{\epsilon}^2$: global divergence (Assumption 2); σ^2 : stochastic gradient noise

² B is a constant and $B > 2$. n is the number of nodes and N is the number of groups. T is the total number of local iterations.

³ Our bound can reduce to the IID case by setting $\tilde{\epsilon}^2 = 0$.

Within a group, workers perform multiple local iterations (on each worker) before local aggregation (within each group). Across groups, each group performs multiple local aggregations before global aggregation. With this characterization, we conduct a novel convergence analysis for H-SGD with non-IID data, non-convex objective function and stochastic gradients. Furthermore, we show that although data can be highly non-IID, local aggregation can help global convergence even when grouping is random. Then by a detailed comparison with local SGD, we show that our convergence upper bound for H-SGD lies in between the convergence upper bounds of two single-level local SGD settings with aggregation periods of I and G , respectively. This is referred to as the “sandwich” behavior, which reveals the fundamental impact of local aggregation.

Our convergence analysis shows that better convergence of H-SGD can be achieved with local aggregation when the number of groups, together with the global and local periods, are chosen appropriately to control the combined effect of upward and downward divergences. In general, we show that since local aggregation is more frequent, grouping strategies with a smaller upward divergence can strengthen the benefit of local aggregation. To reduce the communication cost while maintaining similar or better convergence, it can be beneficial to increase the global period G and decrease the local period I . We also extend our results to multi-level cases where there are multiple levels of local servers, where our characterization with upward and downward divergences can be applied to each level, from which we derive the convergence bound for general multi-level cases.

A comparison of our result with existing results is shown in Table 1. When setting $N = 1$ and $P = I = G$, our result recovers the well-known result for single-level local SGD by Yu, Jin, and Yang (2019). We can also see that choosing $I < G = P$ for H-SGD gives a smaller convergence upper bound, which shows the benefit of the hierarchical structure. The result by Liu et al. (2020) only considers full gradient descent. In this case, the stochastic noise $\sigma^2 = 0$. Even when setting $\sigma^2 = 0$, our result is tighter. The result by Castiglia, Das, and Patterson (2021) only considers IID case where the global divergence $\tilde{\epsilon}^2 = 0$. Setting $\tilde{\epsilon}^2 = 0$, our result is still tighter than theirs since $I < G$. It can be seen that our result

is the most general and tightest.

Main Contributions.

- We introduce the new notion of “upward” and “downward” divergences to characterize data heterogeneity of H-SGD. We show that it can be extended to multi-level cases.
- We derive a general convergence bound for two-level H-SGD with non-IID data, non-convex objective functions and stochastic gradient descent.
- We provide a novel convergence analysis for random grouping and show how local aggregation helps global convergence by a “sandwich” behavior.
- We extend our analysis to multi-level H-SGD and the result shows similar properties as the two-level case. To our knowledge, this is the first analysis which can be extended to multi-level cases.

We also conduct experiments on CIFAR-10, FEMNIST, and CelebA datasets. The results of experiments validate our theoretical results.

2 Related Works

Traditional distributed SGD is introduced by Zinkevich et al. (2010), which can be seen as a special case of local SGD with only one local iteration during one communication round. There have been a large volume of works analyzing the convergence of local SGD, for convex objective functions (Li et al. 2020; Wang et al. 2019), non-convex objective functions (Haddadpour, Farzin et al. 2019; Yu, Yang, and Zhu 2019), and their variants (Karimireddy et al. 2020; Li et al. 2019; Reddi et al. 2020; Wang and Joshi 2019; Yu, Jin, and Yang 2019). Local SGD can be regarded as a special case of H-SGD with only one level. Our theoretical results for H-SGD recover results for both local SGD and traditional distributed SGD. There have been a few works analyzing the convergence of H-SGD, including Castiglia, Das, and Patterson (2021); Zhou and Cong (2019) for IID data, and Liu et al. (2020) for non-IID data with full-batch gradient descent, as described earlier. There are also works on system design for H-SGD without theoretical guarantees (Abad et al. 2020; Luo et al. 2020). In addition, there are works on decentralized SGD (Wang and Joshi 2019; Bellet, Kermarrec, and

Lavoie 2021), where workers exchange their models based on a doubly stochastic mixing matrix. However, the analysis for decentralized SGD cannot be applied to H-SGD, since the second largest eigenvalue of the mixing matrix in the case of H-SGD is one and the decentralized SGD analysis requires this second largest eigenvalue to be strictly less than one.

There are also some works focusing on practical aspects such as model compression and sparsification (Han, Wang, and Leung 2020; Jiang and Agrawal 2018; Jiang et al. 2020; Konecny et al. 2016) and partial worker participation (Bonawitz et al. 2019; Chen et al. 2020). These algorithms and techniques are orthogonal to our work and may be applied together with H-SGD.

3 H-SGD Setup

In two-level H-SGD, all workers are partitioned into N groups $\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_N$. The number of workers in each group is denoted by $n_i := |\mathcal{V}_i|$ ($i = 1, 2, \dots, N$). Then we have $n = \sum_{i=1}^N n_i$. With this grouping, the objective function (1) is equivalent to

$$\min_{\mathbf{w} \in \mathbb{R}^d} f(\mathbf{w}) := \sum_{i=1}^N \frac{n_i}{n} f_i(\mathbf{w}), \quad (2)$$

where $f_i(\cdot)$ is the averaged loss function of workers in group i that is defined as follows:

$$f_i(\mathbf{w}) := \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} F_j(\mathbf{w}). \quad (3)$$

During each local iteration t , each worker j updates its own model using SGD:

$$\mathbf{w}_j^{t+1} = \mathbf{w}_j^t - \gamma \mathbf{g}(\mathbf{w}_j^t, \zeta_j^t), \quad (4)$$

where γ is the learning rate, $\mathbf{g}(\mathbf{w}_j^t, \zeta_j^t)$ is the stochastic gradient of $F_j(\mathbf{w})$, and ζ_j^t represents random data samples from the local dataset \mathcal{D}_j at worker j . We assume that $\mathbb{E}_{\zeta_j^t \sim \mathcal{D}_j} [\mathbf{g}(\mathbf{w}_j^t, \zeta_j^t)] = \nabla F_j(\mathbf{w}_j^t)$.

During one communication round, local models are first averaged within group i ($i = 1, 2, \dots, N$) after every I_i local iterations. In particular, at local iteration $t \in \{I_i, 2I_i, 3I_i, \dots\}$, we compute $\bar{\mathbf{w}}_i^t := \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \mathbf{w}_j^t$. This can be done at a *local server* (e.g., a computational component in close proximity) for group i . After several rounds of “intra-group” aggregations, models from the N groups are averaged globally. Let a global aggregation be performed for every G local iterations. Then, at local iteration t , we have $\bar{\mathbf{w}}^t := \frac{1}{N} \sum_{i=1}^N \bar{\mathbf{w}}_i^t$ for $t = G, 2G, 3G, \dots$. Note that we assume that all workers perform synchronous updates and let G be a common multiple of $\{I_1, \dots, I_N\}$. Therefore, distributed SGD is conducted both at the local level within a group and the global level across groups. We summarize the algorithm in Algorithm 1, where $a \mid b$ (or $a \nmid b$) denotes that a divides (or does not divide) b , i.e., b is (or is not) an integer multiple of a .

In order to understand the fundamental convergence behavior of H-SGD, we will mainly focus on the two-level model introduced above. We will extend our analysis to more than two levels in a later section.

Algorithm 1: Hierarchical SGD (H-SGD)

Input: $\gamma, \bar{\mathbf{w}}^0, G, \{\mathcal{V}_i : i \in \{1, 2, \dots, N\}\}, \{I_i : i \in \{1, 2, \dots, N\}\}$
Output: Global aggregated model $\bar{\mathbf{w}}^T$
for $t = 0$ **to** $T - 1$ **do**
 for Each group $i \in \{1, 2, \dots, N\}$, **in parallel do**
 for Each worker $j \in \mathcal{V}_i$, **in parallel do**
 Compute $\mathbf{g}(\mathbf{w}_j^t, \zeta_j^t)$;
 $\mathbf{w}_j^{t+1} \leftarrow \mathbf{w}_j^t - \gamma \mathbf{g}(\mathbf{w}_j^t, \zeta_j^t)$;
 if $I_i \mid t + 1$ **then**
 Local aggregate: $\bar{\mathbf{w}}_i^{t+1} \leftarrow \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \mathbf{w}_j^{t+1}$;
 if $G \nmid t + 1$ **then**
 Distribute: $\mathbf{w}_j^{t+1} \leftarrow \bar{\mathbf{w}}_i^{t+1}, \forall j \in \mathcal{V}_i$;
 if $G \mid t + 1$ **then**
 Global aggregate: $\bar{\mathbf{w}}^{t+1} \leftarrow \frac{1}{N} \sum_{i=1}^N \bar{\mathbf{w}}_i^{t+1}$;
 Distribute: $\mathbf{w}_j^{t+1} \leftarrow \bar{\mathbf{w}}^{t+1}, \forall j \in \mathcal{V}$;

4 Convergence Analysis for Two-level H-SGD

We begin with a description of the assumptions made in our convergence analysis. Then, we present our results for general two-level H-SGD with fixed groupings. The number of workers and the number of local period during one communication round can vary among groups.

4.1 Assumptions

We make the following minimal set of assumptions that are common in the literature.

Assumption 1. For H-SGD, we assume the following.

a) **Lipschitz gradient**

$$\|\nabla F_j(\mathbf{w}) - \nabla F_j(\mathbf{w}')\| \leq L \|\mathbf{w} - \mathbf{w}'\|, \forall j, \mathbf{w}, \mathbf{w}' \quad (5)$$

where L is some positive constant.

b) **Bounded variance**

$$\mathbb{E}_{\zeta_j^t \sim \mathcal{D}_j} \|\mathbf{g}(\mathbf{w}; \zeta_j^t) - \nabla F_j(\mathbf{w})\|^2 \leq \sigma^2, \forall j, \mathbf{w} \quad (6)$$

where \mathcal{D}_j is the dataset at worker $j \in \mathcal{V}$.

c) **Bounded upward divergence (H-SGD)**

$$\frac{n_i}{n} \sum_{i=1}^N \|\nabla f_i(\mathbf{w}) - \nabla f(\mathbf{w})\|^2 \leq \epsilon^2, \forall \mathbf{w} \quad (7)$$

d) **Bounded downward divergence (H-SGD)**

$$\frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \|\nabla F_j(\mathbf{w}) - \nabla f_i(\mathbf{w})\|^2 \leq \epsilon_i^2, \forall i, \mathbf{w} \quad (8)$$

Note that the Lipschitz gradient assumption also applies to the group objective $f_i(\mathbf{w})$ and overall objective $f(\mathbf{w})$. For example, $\|\nabla f_i(\mathbf{w}) - \nabla f_i(\mathbf{w}')\| = \|\frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}) - \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}')\| \leq \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \|\nabla F_j(\mathbf{w}) - \nabla F_j(\mathbf{w}')\| \leq L \|\mathbf{w} - \mathbf{w}'\|$. Global divergence is often used to describe the data heterogeneity of single-level local SGD (Yu, Jin, and Yang 2019), which is as the following.

Assumption 2. Bounded global divergence

$$\frac{1}{n} \sum_{j=1}^n \|\nabla F_j(\mathbf{w}) - \nabla f(\mathbf{w})\|^2 \leq \tilde{\epsilon}^2, \forall \mathbf{w}. \quad (9)$$

Note that when $N = 1$, H-SGD reduces to local SGD and downward divergence becomes the global divergence while upward divergence becomes zero. When $N = n$, H-SGD also reduces to local SGD but upward divergence becomes the global divergence while downward divergence becomes zero. Here, we will use global divergence to show our results for H-SGD can encompass the results for local SGD.

Note for the global divergence, it is easy to show that

$$\begin{aligned} \frac{1}{n} \sum_{j=1}^n \|\nabla F_j(\mathbf{w}) - \nabla f(\mathbf{w})\|^2 &= \sum_{i=1}^N \frac{n_i}{n} \|\nabla f_i(\mathbf{w}) - \nabla f(\mathbf{w})\|^2 \\ &+ \sum_{i=1}^N \frac{n_i}{n} \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \|\nabla F_j(\mathbf{w}) - \nabla f_i(\mathbf{w})\|^2. \end{aligned} \quad (10)$$

In (10), we see that the upward and downward divergences are in fact a partition of the global divergence, which implies that upward and downward divergences do not increase at the same time. It is the hierarchical structure that makes this partition possible. Later we will show that this partition can exactly explain the benefits of local aggregation. We will discuss more about the relationship between upward/downward and global divergences in a later section.

4.2 Convergence Analysis

Technical Challenge. The main challenge of the analysis is that workers only perform local iterations before global aggregation in local SGD, whereas in H-SGD, workers not only perform local iterations but also do aggregations with other workers in the same group. If we directly apply the analysis for local SGD in upward part, the effect of local aggregations would be neglected and the resulted bound would be loose. To address this, our key idea is to construct the local parameter drift $\|\mathbf{w}_j - \bar{\mathbf{w}}_i\|$, $j \in \mathcal{V}_i$ in the analysis of global parameter drift $\|\bar{\mathbf{w}}_i - \bar{\mathbf{w}}\|$, so that the analysis for the downward part can be incorporated in the analysis for the upward part.

Theorem 1. Consider the problem in (2). For any fixed worker grouping that satisfies Assumption 1, if the learning rate in Algorithm 1 satisfies $\gamma < \frac{1}{2\sqrt{6}GL}$, then for any $T \geq 1$, we have

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\bar{\mathbf{w}}^t)\|^2 \leq \frac{2(f^0 - f^*)}{\gamma T} + \gamma L \frac{1}{n} \sigma^2 \quad (11a)$$

$$+ 2C\gamma^2 G \frac{N-1}{n} \sigma^2 + 3C\gamma^2 G^2 \epsilon^2 \quad (11b)$$

$$+ 2C\gamma^2 \sigma^2 \sum_{i=1}^N \frac{(n_i-1)I_i}{n} + 3C\gamma^2 \sum_{i=1}^N \frac{n_i}{n} I_i^2 \epsilon_i^2, \quad (11c)$$

where $C = 40/3$.

Remark 1. Note that the bound (11a)–(11c) can be partitioned into three parts. The terms in (11a) are the original SGD part. If we set $N = n_i = 1$ (ϵ and ϵ_i become zero), then only (11a) will remain, which is the same as convergence bound for non-convex SGD in (Bottou, Curtis, and Nocedal 2018). The terms in (11b) are the upward part, which consists of noise and upward divergence associated with G . The terms in (11c) are the downward part, which has a similar form as (11b) associated with I_i . Divergence plays a more important role than noise since the coefficients in front of the SGD noise σ^2 include G and I_i , while the coefficients in front of the divergences ϵ and ϵ_i include G^2 and I_i^2 . Since $G > I_i, \forall i$, the upward part has a stronger influence on convergence.

Remark 2. Note that all the divergences of H-SGD can be written as $O(\gamma^2 G^2 \epsilon^2 + \gamma^2 \sum_{i=1}^N \frac{n_i}{n} I_i^2 \epsilon_i^2)$ while the corresponding part in local SGD is $O(\gamma^2 G^2 \epsilon^2)$ (Yu, Jin, and Yang 2019). This exactly shows how local aggregation overcomes data heterogeneity: global divergence is partitioned into two parts and local aggregation weakens the effects of the downward part. This also brings us some insights on how to group workers. The grouping strategy should have the smallest upward divergence since this can make full use of benefits of local aggregation. We will validate this property in the experiments.

Remark 3. Let $\gamma = \sqrt{\frac{n}{T}}$ with $T \geq \frac{1}{24G^2L^2\sqrt{n}}$, when T is sufficiently large, $\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\bar{\mathbf{w}}^t)\|^2 = O\left(\sqrt{\frac{1}{nT}}\right) + O\left(\frac{1}{T}\right)$, which achieves a linear speedup in n .

In the following corollary, we show that our results can encompass local SGD with local period P .

Corollary 1. (Degenerate to local SGD) Let $N = 1$ and $\gamma \leq \frac{1}{2\sqrt{6}PL}$, from Theorem 1, we obtain

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\bar{\mathbf{w}}^t)\|^2 &\leq \frac{2(f^0 - f^*)}{\gamma T} + \frac{\gamma L \sigma^2}{n} \\ &+ 2C\gamma^2 L^2 \sigma^2 \left(1 - \frac{1}{n}\right) P + 3C\gamma^2 L^2 P^2 \tilde{\epsilon}^2 \\ &= O\left(\frac{1}{\gamma T}\right) + O\left(\frac{\gamma \sigma^2}{n}\right) \\ &+ O\left(\gamma^2 P \sigma^2 \left(1 - \frac{1}{n}\right)\right) + O\left(\gamma^2 P^2 \tilde{\epsilon}^2\right). \end{aligned} \quad (12)$$

While (12) is similar to the bound by Yu, Jin, and Yang (2019), we note that the third term $O\left(\gamma^2 P \sigma^2 \left(1 - \frac{1}{n}\right)\right)$ in the last equality in (12) has an additional term of $\left(1 - \frac{1}{n}\right)$ compared to the bound by Yu, Jin, and Yang (2019). This term can potentially make our bound tighter. Another important observation is that the techniques used to obtain this term is the key to make our H-SGD bound in Theorem 1 encompass original SGD cases.

4.3 H-SGD with Random Grouping

We now state our convergence results for H-SGD with random grouping. We will show that the convergence upper bound of H-SGD with random grouping takes a value that is between the convergence upper bounds of two single-level

local SGD settings with local and global periods of I and G , respectively. This will provide insights on when and why local aggregation helps.

For worker grouping, we consider all possible grouping strategies with the constraint that $n_i = n/N, \forall i$. Then, we uniformly select one grouping strategy at random. Let the random variable S denote the uniformly random grouping strategy. This means that each realization of S corresponds to one grouping realization. First, we introduce two key lemmas for our convergence analysis.

Lemma 1. *Using the uniformly random grouping strategy S , for any \mathbf{w} , the average upward divergence is*

$$\mathbb{E}_S \left[\frac{1}{N} \sum_{i=1}^N \|\nabla f(\mathbf{w}) - \nabla f_i(\mathbf{w})\|^2 \right] \leq \left(\frac{N-1}{n-1} \right) \tilde{\epsilon}^2, \quad (13)$$

where $\tilde{\epsilon}$ is given in (9).

Lemma 2. *Using the uniformly random grouping strategy S , for any \mathbf{w} , the average downward divergence is*

$$\mathbb{E}_S \left[\frac{1}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \|\nabla f_i(\mathbf{w}) - \nabla F_k(\mathbf{w})\|^2 \right] \leq \left(1 - \frac{N-1}{n-1} \right) \tilde{\epsilon}^2. \quad (14)$$

Similar to (10), Lemma 1 and Lemma 2 show that the sum of upper bounds for upward and downward divergences is equal to the global divergence. Furthermore, when the number of groups increases, upward divergence becomes larger while downward divergence becomes smaller. When grouping is random, N is the key parameter that determines how global divergence is partitioned. For simplicity, we let each group have the same local period $I_i = I, \forall i$. Then, we can obtain the following theorem.

Theorem 2. *Using the uniformly random grouping strategy S , let $\gamma \leq \frac{1}{2\sqrt{6}GL}$, then we have*

$$\begin{aligned} \mathbb{E}_S \left[\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\bar{\mathbf{w}}^t)\|^2 \right] &\leq \frac{2(f^0 - f^*)}{\gamma T} + \frac{\gamma L \sigma^2}{n} \\ &+ 2C\gamma^2 L^2 \left[\left(\frac{N-1}{n} \right) G + \left(1 - \frac{N}{n} \right) I \right] \sigma^2 \\ &+ 3C\gamma^2 L^2 \left[\left(\frac{N-1}{n-1} \right) G^2 + \left(1 - \frac{N-1}{n-1} \right) I^2 \right] \tilde{\epsilon}^2, \end{aligned} \quad (15)$$

where $C = 40/3$.

Remark 4. Note that both multiplicative factors of the noise term with σ^2 and the divergence term with $\tilde{\epsilon}^2$ are composed of two parts, where the upward part is “modulated” by G while the downward part is “modulated” by I . As N becomes larger, the upward part has a stronger influence on the convergence bound since both $\frac{N-1}{n}$ and $\frac{N-1}{n-1}$ become larger. Note that the convergence upper bound of H-SGD can be sandwiched by the convergence upper bounds of two local SGD. To see this, we consider the following three scenarios: 1) local SGD with aggregation period $P = G$, 2) local SGD with aggregation period $P = I$, and 3) H-SGD with local aggregation period I and global aggregation period G . We

let them all start with the same $\bar{\mathbf{w}}^0$ and choose learning rate γ such that $\gamma \leq \frac{1}{2\sqrt{6}GL}$. We can see that all these three cases have the same first two terms in (15). However, for the third and fourth terms in (15), we have

$$\left(1 - \frac{1}{n} \right) I \leq \left(\frac{N-1}{n} \right) G + \left(1 - \frac{N}{n} \right) I \leq \left(1 - \frac{1}{n} \right) G, \quad (16)$$

$$I^2 \leq \left(\frac{N-1}{n-1} \right) G^2 + \left(1 - \frac{N-1}{n-1} \right) I^2 \leq G^2. \quad (17)$$

Equations (16) and (17) show that the convergence upper bound of H-SGD has a value between the convergence upper bounds of local SGD with aggregation periods of $P = I$ and $P = G$, respectively. The grouping approach can explicitly characterize how much the convergence bound moves towards the best case, i.e., local SGD with $P = I$. However, this case incurs the highest global communication cost, so grouping can adjust the trade-off between convergence and communication cost.

Remark 5. Even with a large G , if we make I sufficiently small, the convergence bound of H-SGD can be improved. To see this, consider the last two terms of (15). Suppose $G = mI, m = 1, 2, \dots$. For a non-trivial worker grouping, i.e., $1 < N < n$, if we increase G to $G' = lG, 1 < l < \sqrt{\frac{1}{m^2} \frac{n-N}{N}} + 1$ and decrease I to $I' = qI, q \leq \sqrt{1 - m^2(l^2 - 1) \frac{N}{n-N}}$, then one can show that the bound (15) using G' and I' can be lower than that using G and I . A similar behavior can also be seen for fixed grouping, as shown empirically in Figure 3b in the experiments.

5 H-SGD with More Than Two Levels

In the two-level H-SGD setting discussed in the previous sections, there is only one level of local servers between the global server and workers. In this section, we extend our analysis to H-SGD with more than one level of local servers. Specifically, as shown in Figure 1, we consider a total of $M \geq 2$ levels. For $\ell = 1, 2, \dots, M-1$, each server at level $\ell-1$ is connected to N_ℓ servers at level ℓ , where we assume that the global server is at a “dummy” level $\ell = 0$ for convenience. For $\ell = M$, each server at level $M-1$ directly connects to N_M workers. As a result, we have a total of $n = \prod_{\ell=1}^M N_\ell$ workers. With this notation, a sequence of indices $(k_1, k_2, \dots, k_\ell)$ denotes a “path” from the global server at level 0 to a local server/worker at level ℓ , where this path traverses the k_1 -th server at level 1, the k_2 -th server at level 2 connected to the k_1 -th server at level 1, and so on. Hence, the index sequence $(k_1, k_2, \dots, k_\ell)$ uniquely determines servers (and workers if $\ell = M$) down to level ℓ .

We assume in the multi-level case that each server at level $\ell-1$ connects to the same number (N_ℓ) of servers/workers in the next level ℓ , for ease of presentation. Our results can be extended to the general case with different group sizes.

Let $F_{k_1 \dots k_M}(\mathbf{w})$ denote the objective function of worker $k_1 \dots k_M$ and $f_{k_1 \dots k_\ell}(\mathbf{w})$ denote the averaged objective func-




Figure 1: Three-level example with $N_1 = N_2 = 2, N_3 = 3$.

tion of workers connected to server $k_1 \dots k_\ell$,

$$f_{k_1 \dots k_\ell}(\mathbf{w}) := \frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \dots \sum_{k_M=1}^{N_M} F_{k_1 \dots k_M}(\mathbf{w}). \quad (18)$$

Then the global objective function can be rewritten as

$$\begin{aligned} f(\mathbf{w}) &= \frac{1}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_M=1}^{N_M} F_{k_1 \dots k_M}(\mathbf{w}) \\ &= \frac{1}{\prod_{j=1}^\ell N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_\ell=1}^{N_\ell} f_{k_1 \dots k_\ell}(\mathbf{w}). \end{aligned} \quad (19)$$

Let P_ℓ ($\ell = 1, 2, \dots, M$) denote the period (expressed as the number of local iterations) that the parameters at servers at level ℓ are aggregated by their parent server at level $\ell - 1$. We require that $P_1 > P_2 > \dots > P_M$ and P_ℓ is an integer multiple of $P_{\ell-1}$. The algorithm is in Appendix D.

In the following, we provide results for the random grouping case. Results for the fixed grouping case can be found in Appendix D. We apply the divergence partition idea to each level, so we extend Lemma 1 and Lemma 2 to obtain the following lemma.

Lemma 3. *Using the uniformly random grouping strategy \mathcal{S} , the ℓ -th level averaged upward and downward divergences are given by*

$$\mathbb{E}_{\mathcal{S}} \left[\frac{1}{n_\ell} \sum_{k_1=1}^{N_1} \dots \sum_{k_\ell=1}^{N_\ell} \|\nabla f_{k_1 \dots k_\ell}(\mathbf{w}) - \nabla f(\mathbf{w})\|^2 \right] \leq \left(\frac{n_\ell - 1}{n - 1} \right) \tilde{\epsilon}, \quad (20)$$

$$\begin{aligned} &\mathbb{E}_{\mathcal{S}} \left[\frac{n_\ell}{n} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \dots \sum_{k_M=1}^{N_M} \|\nabla f_{k_1 \dots k_\ell}(\mathbf{w}) - \nabla F_{k_1 \dots k_M}(\mathbf{w})\|^2 \right] \\ &\leq \left(1 - \frac{n_\ell - 1}{n - 1} \right) \tilde{\epsilon}^2 \end{aligned} \quad (21)$$

respectively, $\forall \mathbf{w}, \forall k_1, \dots, k_\ell$ in (21), where $n_\ell = \prod_{j=1}^\ell N_j$ and $\tilde{\epsilon}$ is the global divergence.

When the number of servers in the ℓ -th-level n_ℓ increases, the ℓ -th-level's upward divergence becomes larger while its downward divergence becomes smaller. The sum of the ℓ -th-level upward and downward divergences are upper bounded by the global divergence. Based on this lemma, we derive the convergence bound for multi-level case as the following.

Theorem 3. *Consider uniform random grouping strategy \mathcal{S} , if $\gamma \leq \frac{1}{2\sqrt{6}P_1L}$, then for multi-level case,*

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}_{\mathcal{S}} \left\| \nabla f(\bar{\mathbf{w}}^t) \right\|^2 &\leq \frac{2}{\gamma T} (f^0 - f^*) + \frac{\gamma L \sigma^2}{n} \\ &\quad + C\gamma^2 L^2 \frac{1}{M-1} \cdot \sum_{\ell=1}^{M-1} [2A_1(\ell)\sigma^2 + 3A_2(\ell)\tilde{\epsilon}^2] \end{aligned} \quad (22)$$

where, $A_1(\ell) := P_1 \left(\frac{1}{\prod_{j=\ell}^M N_j} - \frac{1}{n} \right) + P_\ell \left(1 - \frac{1}{\prod_{j=\ell}^M N_j} \right)$, $A_2(\ell) := P_1^2 \left(\frac{n_\ell - 1}{n - 1} \right) + P_\ell^2 \left(1 - \frac{n_\ell - 1}{n - 1} \right)$.

Remark 6. Note that when $M = 2$, this bound reduces to the bound given in Theorem 2. The effect of the ℓ -th level can be represented as $2A_1(\ell)\sigma^2 + 3A_2(\ell)\tilde{\epsilon}^2$, where each $A_i(\ell), i = 1, 2$ is composed of an upward part “modulated” by P_1 and a downward part “modulated” by P_ℓ . Similar to (16) and (17), we can obtain the sandwich behavior as:

$$\left(1 - \frac{1}{n} \right) P_M \leq \frac{1}{M-1} \sum_{\ell=1}^{M-1} A_1(\ell) \leq \left(1 - \frac{1}{n} \right) P_1, \quad (23)$$

$$P_M^2 \leq \frac{1}{M-1} \sum_{\ell=1}^{M-1} A_2(\ell) \leq P_1^2. \quad (24)$$

It can be seen that the convergence upper bound of H-SGD with more than two levels also takes a value that is between the convergence upper bounds of local SGD with local periods P_1 and P_M , respectively. Compared to the two-level case, this provides greater freedom to choose parameters in $A_1(\ell)$ and $A_2(\ell)$, for $\ell = 1, \dots, M-1$.

6 Experiments

In this section, we validate our theoretical results with experiments on training the VGG-11 model over CIFAR-10 (Krizhevsky, Hinton et al. 2009) and CelebA (Liu et al. 2015), and training a convolutional neural network (CNN) over FEMNIST (Cohen et al. 2017), all with non-IID data partitioning across workers. The communication time is emulated by measuring the round-trip time of transmitting the model between a device (in a home) and local (near) / global (far) Amazon EC2 instances. For the computation time, we measured the averaged computation time needed for VGG-11 during one iteration, where we ran SGD with VGG-11 on a single GPU for 100 times and then computed the averaged computation time, which is approximately 4 ms per iteration on each worker. The experiments presented in this section are for *two-level* H-SGD. Additional details of the setup and experiments with *more levels* and *partial worker participation* can be found in Appendix E.

In all plots, curves with P (denoting the aggregation period) are for local SGD and curves with G (global period), I (local period), and N (number of groups) are for H-SGD. Figure 2 shows that for all the datasets and models evaluated in the experiments, H-SGD can achieve a better accuracy than local SGD given the same communication time. In other words, for any given accuracy, H-SGD can achieve this accuracy using shorter or similar communication time. In Table 2,

Table 2: Total time (s) needed to achieve 50% test accuracy for VGG-11 with CIFAR-10.

Case	$P = 5$	$P = 10$	$P = 50$	$G = 50, I = 5$	$G = 50, I = 10$
Time (s)	673.5	690.2	944.3	160.3	381.1




Figure 2: Test accuracy v.s. communication time ($N = 2$).




Figure 3: Results with CIFAR-10. Test accuracy v.s. local iterations. By default, $N = 2$.

we provide the total time (including both communication and computation) needed for VGG-11 with CIFAR-10 to achieve 50% test accuracy (since 50% is nearly the best accuracy for $P = 50$). We can see H-SGD cases ($G = 50, I = 5$ and $G = 50, I = 10$) perform much better than single-level local SGD cases. This is because communication time plays a much more important role and H-SGD can achieve 50% accuracy within less amount of time in total.

To show the “sandwich” behavior, we examine the test accuracy as a function of the total number of local iterations for the non-IID scenario in Figure 3a. We can observe that the performance of H-SGD with global period G and local period I is between that of local SGD with $P = G$ and with $P = I$. We can also observe that as N becomes larger, the performance becomes worse since the upward divergence becomes larger, which is consistent with our analysis in Remark 4. Comparing Figures 3a and 3b, we see that decreasing I while increasing G can even improve the performance, which is consistent with Theorem 2. For example, $G = 200, I = 2$ gives a better performance than $G = 50, I = 5$, while $G = 50, I = 5$ is similar to $P = 10$. This shows that by allowing more local aggregations, H-SGD can reduce the number of global aggregations by 95% ($G = 200$ vs. $P = 10$) while maintaining a similar performance to local SGD. In Figure 3c, we show the effects

of grouping corresponding to our analysis for Theorem 1. Group-IID is a grouping strategy which makes upward divergence nearly zero while group-non-IID is a grouping strategy with a large upward divergence. We can see that group-IID performs as well as group-non-IID after reducing I by half.

7 Conclusion

We have studied H-SGD with multi-level model aggregations. In particular, we have provided a thorough theoretical analysis on the convergence of H-SGD over non-IID data, under both fixed and random worker grouping. We have successfully answered the important question on how local aggregation affects convergence of H-SGD. Based on our novel analysis of the local and global divergences, we established explicit comparisons of the convergence rate for H-SGD and local SGD. Our theoretical analysis provides valuable insights into the design of practical H-SGD systems, including the choice of global and local aggregation periods. Different grouping strategies are considered to best utilize the divergences to reduce communication costs while accelerating learning convergence. In addition, we have extended the analysis approach to H-SGD with arbitrary number of levels. Future work could theoretically analyze the effect of partial worker participation in H-SGD.

Acknowledgment

This research was partly sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

References

- Abad, M. S. H.; Ozfatura, E.; Gunduz, D.; and Ercetin, O. 2020. Hierarchical federated learning across heterogeneous cellular networks. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 8866–8870. IEEE.
- Bellet, A.; Kermarrec, A.-M.; and Lavoie, E. 2021. D-Cliques: Compensating NonIIDness in Decentralized Federated Learning with Topology. *arXiv preprint arXiv:2104.07365*.
- Bonawitz, K.; Eichner, H.; Grieskamp, W.; Huba, D.; Ingerman, A.; Ivanov, V.; Kiddon, C.; Konečný, J.; Mazzocchi, S.; McMahan, H. B.; et al. 2019. Towards Federated Learning at Scale: System Design. In *SysML 2019*.
- Bottou, L.; Curtis, F. E.; and Nocedal, J. 2018. Optimization methods for large-scale machine learning. *SAIM Review*, 60(2): 223–311.
- Castiglia, T.; Das, A.; and Patterson, S. 2021. Multi-Level Local SGD: Distributed SGD for Heterogeneous Hierarchical Networks. In *International Conference on Learning Representations*.
- Chen, C.; Chen, Z.; Zhou, Y.; and Kailkhura, B. 2020. Fed-Cluster: Boosting the Convergence of Federated Learning via Cluster-Cycling. *arXiv preprint arXiv:2009.10748*.
- Cohen, G.; Afshar, S.; Tapson, J.; and Van Schaik, A. 2017. EMNIST: Extending MNIST to handwritten letters. In *2017 International Joint Conference on Neural Networks (IJCNN)*, 2921–2926. IEEE.
- Haddadpour, Farzin; et al. 2019. Local SGD with Periodic Averaging: Tighter Analysis and Adaptive Synchronization. In *Advances in Neural Information Processing Systems*.
- Han, P.; Wang, S.; and Leung, K. K. 2020. Adaptive Gradient Sparsification for Efficient Federated Learning: An Online Learning Approach. In *IEEE International Conference on Distributed Computing Systems (ICDCS)*.
- Jiang, P.; and Agrawal, G. 2018. A Linear Speedup Analysis of Distributed Deep Learning with Sparse and Quantized Communication. In *Advances in Neural Information Processing Systems*.
- Jiang, Y.; Wang, S.; Valls, V.; Ko, B. J.; Lee, W.-H.; Leung, K. K.; and Tassiulas, L. 2020. Model Pruning Enables Efficient Federated Learning on Edge Devices. In *NeurIPS Workshop on Scalability, Privacy, and Security in Federated Learning (SpicyFL)*.
- Karimireddy, S. P.; Kale, S.; Mohri, M.; Reddi, S.; Stich, S.; and Suresh, A. T. 2020. SCAFFOLD: Stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*, 5132–5143. PMLR.
- Kim, H.; and Feamster, N. 2013. Improving network management with software defined networking. *IEEE Communications Magazine*, 51(2): 114–119.
- Konecny, J.; McMahan, H. B.; Yu, F. X.; Richtarik, P.; Suresh, A. T.; and Bacon, D. 2016. Federated Learning: Strategies for Improving Communication Efficiency. In *NeurIPS Workshop on Private Multi-Party Machine Learning*.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Li, H.; Ota, K.; and Dong, M. 2018. Learning IoT in Edge: Deep Learning for the Internet of Things with Edge Computing. *IEEE Network*, 32(1): 96–101.
- Li, T.; Sahu, A. K.; Talwalkar, A.; and Smith, V. 2019. Federated learning: Challenges, methods, and future directions. *arXiv preprint arXiv:1908.07873*.
- Li, X.; Huang, K.; Yang, W.; Wang, S.; and Zhang, Z. 2020. On the Convergence of FedAvg on Non-IID Data. In *International Conference on Learning Representations*.
- Lin, T.; Stich, S. U.; Patel, K. K.; and Jaggi, M. 2020. Don't Use Large Mini-batches, Use Local SGD. In *International Conference on Learning Representations*.
- Liu, L.; Zhang, J.; Song, S.; and Letaief, K. B. 2020. Client-edge-cloud hierarchical federated learning. In *ICC 2020-2020 IEEE International Conference on Communications (ICC)*, 1–6. IEEE.
- Liu, Z.; Luo, P.; Wang, X.; and Tang, X. 2015. Deep Learning Face Attributes in the Wild. In *Proceedings of International Conference on Computer Vision (ICCV)*.
- Luo, S.; Chen, X.; Wu, Q.; Zhou, Z.; and Yu, S. 2020. HFEL: Joint Edge Association and Resource Allocation for Cost-Efficient Hierarchical Federated Edge Learning. *IEEE Transactions on Wireless Communications*, 19(10): 6535–6548.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proc. International Conference on Artificial Intelligence and Statistics (AISTATS)*, 1273–1282.
- Reddi, S.; Charles, Z.; Zaheer, M.; Garrett, Z.; Rush, K.; Konečný, J.; Kumar, S.; and McMahan, H. B. 2020. Adaptive Federated Optimization. *arXiv:2003.00295*.
- Stich, S. U. 2019. Local SGD Converges Fast and Communicates Little. In *International Conference on Learning Representations*.
- Wang, J.; and Joshi, G. 2019. Cooperative SGD: A unified Framework for the Design and Analysis of Communication-Efficient SGD Algorithms. In *ICML 2019 Workshop on Coding Theory for Large-Scale ML*.
- Wang, S.; Tuor, T.; Salonidis, T.; Leung, K. K.; Makaya, C.; He, T.; and Chan, K. 2019. Adaptive Federated Learning in Resource Constrained Edge Computing Systems. *IEEE Journal on Selected Areas in Communications*, 37(6): 1205–1221.

- Yu, H.; Jin, R.; and Yang, S. 2019. On the Linear Speedup Analysis of Communication Efficient Momentum SGD for Distributed Non-Convex Optimization. In *ICML*, 7184–7193.
- Yu, H.; Yang, S.; and Zhu, S. 2019. Parallel Restarted SGD with Faster Convergence and Less Communication: Demystifying Why Model Averaging Works for Deep Learning. In *AAAI*.
- Zhou, F.; and Cong, G. 2019. A distributed hierarchical SGD algorithm with sparse global reduction. *arXiv preprint arXiv:1903.05133*.
- Zinkevich, M.; Weimer, M.; Li, L.; and Smola, A. J. 2010. Parallelized stochastic gradient descent. In *NeurIPS*, 2595–2603.

Appendix

A	Additional Discussions	11
A.1	Differences between H-SGD and Decentralized SGD	11
A.2	Justification of the Bounded Gradient Divergence Assumptions	11
B	Proofs for the Two-level Fixed Grouping Case	13
B.1	Proof of Theorem 1	13
B.2	Bounding Upward Parameter MSE	15
B.3	Bounding Downward Parameter MSE	18
C	Proofs for the Two-level Random Grouping Case	21
C.1	Proof of Lemma 1	21
C.2	Proof of Lemma 2	22
C.3	Proof of Theorem 2	23
D	Extension to the Multi-Level Case ($M \geq 2$)	24
D.1	Proof of Theorem D.1	24
D.2	Proof of Lemma 3	30
D.3	Proof of Theorem 3	30
E	Additional Details and Results of Experiments	31
E.1	Additional Results for Two-level Case	32
E.2	Multi-level Experiments with ($M = 3$)	34

A Additional Discussions

In this section, we provide some more discussions on the differences between decentralized SGD and H-SGD and also on the justification of the bounded gradient divergence assumptions (Assumption 1 for H-SGD and Assumption 2 for local SGD).

A.1 Differences between H-SGD and Decentralized SGD

In decentralized SGD, a doubly stochastic mixing is used to describe how workers exchange models with their neighbors. A necessary assumption for this mixing matrix is that its secondary largest absolute eigenvalue ξ is strictly less than 1. We can see this in classic works for decentralized SGD (Wang and Joshi 2019; Lian et al. 2017). This is because in their proofs, $1 - \xi$ is present in the denominator.

In H-SGD, local aggregation can also be described by a matrix \mathbf{A} . However, \mathbf{A} does not satisfy the assumption for the mixing matrix in decentralized SGD. To see this, consider the following non-trivial example for H-SGD. Suppose there are $K = \frac{n}{N}$, $1 < N < n$ nodes in each group. Let

$$\mathbf{A} = \begin{bmatrix} \mathbf{D} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{D} & \cdots & \mathbf{0} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{D} \end{bmatrix},$$

where

$$\mathbf{D} = \begin{bmatrix} \frac{1}{K} & \frac{1}{K} & \cdots & \frac{1}{K} \\ \frac{1}{K} & \frac{1}{K} & \cdots & \frac{1}{K} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{1}{K} & \frac{1}{K} & \cdots & \frac{1}{K} \end{bmatrix}.$$

Then we have $\det(\mathbf{A} - \lambda \mathbf{I}) = \det(\mathbf{D} - \lambda \mathbf{I})^N$. Since \mathbf{D} is a rank-1 matrix with eigenvalue 1, there will be N eigenvalues with value 1, which does not satisfy the assumption for the mixing matrix in decentralized SGD.

The above shows that the convergence result of H-SGD cannot be obtained from the approach based on the mixing matrix. Our result provides useful insights on how grouping and local/global periods affect the convergence, which cannot be obtained using standard approaches for decentralized SGD or local SGD.

A.2 Justification of the Bounded Gradient Divergence Assumptions

Bounded gradient divergence assumption is widely used and well accepted in literature for theoretical analysis for local SGD (Haddadpour, Farzin et al. 2019; Karimireddy et al. 2020; Li et al. 2020; Stich 2019; Wang et al. 2019; Yu, Yang, and Zhu 2019; Yu, Jin, and Yang 2019) to describe the degree of data heterogeneity. In this paper, upward and downward divergence are essential for the convergence analysis of H-SGD. First, with a fixed global divergence, there can be different kinds of data distributions in a group. We use upward and downward divergence to capture these characteristics and to show how they can affect the convergence of H-SGD. The standard global divergence assumption cannot capture the heterogeneous characteristics of data distribution within each group and across different groups.

Upward and downward divergence can also provide us insights on how to group workers. If we group workers with similar data distribution, then the downward divergence can be very small and the benefit of local aggregation can be limited, so it is better to group workers with diverse data. This is one implication of Theorem 1.

A classic example that satisfies this bounded gradient divergence assumption is the cross-entropy loss with softmax. Consider the following model:

$$\mathbf{z} = \mathbf{W}\mathbf{x},$$

where \mathbf{x} is the input data and \mathbf{W} is a matrix representing the model parameters. After softmax, we have

$$s_j = \frac{e^{z_j}}{\sum_i e^{z_i}}, \forall j,$$

where z_j is the j -th component of \mathbf{z} . The cross-entropy loss function is given by

$$L(\mathbf{y}, \mathbf{s}) = - \sum_i y_i \log(s_i),$$

where \mathbf{y} is the one-hot encoded ground-truth label. Then we have

$$\frac{\partial L}{\partial \mathbf{z}} = \mathbf{y} - \mathbf{s}.$$

Because $0 \leq s_i \leq 1, 0 \leq y_i \leq 1, \forall i$, we have $|y_i - s_i| \leq 1$. That is, the derivative of L with respect to \mathbf{z} is bounded. Since \mathbf{z} is a linear function of \mathbf{W} , the derivative of \mathbf{z} with respect to \mathbf{W} is also bounded as long as the data sample \mathbf{x} is bounded.

Then, the gradient of the loss L with respect to the model parameter \mathbf{W} is the product of $\frac{\partial L}{\partial \mathbf{z}}$ and $\frac{\partial \mathbf{z}}{\partial \mathbf{W}}$, which is also bounded and depends on the input data \mathbf{x} . This boundedness of derivatives extends to neural network models with sigmoid or ReLU activation functions, since these functions also have bounded derivatives. Hence, the divergence, which is defined as the norm of the difference between gradients, is also bounded.

B Proofs for the Two-level Fixed Grouping Case

We present the proofs of Theorem 1 in this section. Throughout the proof, we use the following inequalities frequently:

$$\left\| \sum_{i=1}^M p_i \mathbf{x}_i \right\|^2 \leq \sum_{i=1}^M p_i \|\mathbf{x}_i\|^2, \quad (\text{B.1})$$

and

$$\sum_{i=1}^M p_i \|\mathbf{x}_i - \bar{\mathbf{x}}\|^2 = \sum_{i=1}^M p_i \|\mathbf{x}_i\|^2 - \|\bar{\mathbf{x}}\|^2 \leq \sum_{i=1}^M p_i \|\mathbf{x}_i\|^2, \quad (\text{B.2})$$

where $\bar{\mathbf{x}} := \sum_{i=1}^M p_i \mathbf{x}_i$, $0 \leq p_i \leq 1$ and $\sum_{i=1}^M p_i = 1$ and (B.1) is due to Jensen's Inequality.

For the ease of notations, we define $\mathbf{g}_j(\mathbf{w}) := \mathbf{g}(\mathbf{w}, \zeta_j)$ for any \mathbf{w} .

B.1 Proof of Theorem 1

Although the averaged global model $\bar{\mathbf{w}}^t$ is not observable in the system at $t \neq aG$, $a = 0, 1, 2, \dots$, here we use $\bar{\mathbf{w}}^t$ for analysis.

$$\begin{aligned} \mathbb{E}f(\bar{\mathbf{w}}^{t+1}) &= \mathbb{E}f\left[\bar{\mathbf{w}}^t - \gamma \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^t)\right] \\ &\stackrel{(a)}{\leq} \mathbb{E}f(\bar{\mathbf{w}}^t) - \gamma \mathbb{E}\left\langle \nabla f(\bar{\mathbf{w}}^t), \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^t) \right\rangle + \frac{\gamma^2 L}{2} \mathbb{E}\left\| \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^t) \right\|^2, \end{aligned} \quad (\text{B.3})$$

where (a) is a proposition of Lipschitz smooth, which is shown in (4.3) of (Bottou, Curtis, and Nocedal 2018). For the inner product term

$$\begin{aligned} & - \gamma \mathbb{E}\left\langle \nabla f(\bar{\mathbf{w}}^t), \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^t) \right\rangle \\ & \stackrel{(a)}{=} - \gamma \mathbb{E}\left\{ \mathbb{E}\left[\left\langle \nabla f(\bar{\mathbf{w}}^t), \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^t) \right\rangle \middle| \{\mathbf{w}_j^t : \forall j \in \mathcal{V}\} \right] \right\} \\ & \stackrel{(b)}{=} - \gamma \mathbb{E}\left\langle \nabla f(\bar{\mathbf{w}}^t), \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^t) \right\rangle \\ & = \frac{\gamma}{2} \left(\mathbb{E}\left\| \nabla f(\bar{\mathbf{w}}^t) - \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^t) \right\|^2 - \mathbb{E}\|\nabla f(\bar{\mathbf{w}}^t)\|^2 - \mathbb{E}\left\| \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^t) \right\|^2 \right). \end{aligned} \quad (\text{B.4})$$

Recalling that $\mathbf{g}_j(\mathbf{w}_j^t) = \mathbf{g}(\mathbf{w}_j^t, \zeta_j^t)$, we note that the expectation operator \mathbb{E} in (B.3) is over all random samples $\{\zeta_j^\tau : \tau = 0, 1, \dots, t, j = 1, \dots, n\}$, where ζ_j^t denotes the random samples used for SGD in iteration t . In step (a) of (B.4), a conditional expectation is taken for given $\{\mathbf{w}_j^t, j = 1, \dots, n\}$. We note that $\mathbb{E}[\mathbf{g}_j(\mathbf{w}_j^t) | \{\mathbf{w}_j^t : \forall j \in \mathcal{V}\}] = \nabla F_j(\mathbf{w}_j^t)$ due to the unbiased gradient assumption, from which (b) follows. This way of replacing the random gradient $\mathbf{g}_j(\mathbf{w}_j^t)$ by its unbiased average $\nabla F_j(\mathbf{w}_j^t)$ through the conditional expectation is also used later in the proof, where we may not write out the conditional expectation step for compactness.

For the last term of (B.3), we have

$$\begin{aligned} & \mathbb{E}\left\| \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^t) \right\|^2 \\ & \stackrel{(a)}{=} \mathbb{E}\left\| \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} (\mathbf{g}_j(\mathbf{w}_j^t) - \nabla F_j(\mathbf{w}_j^t)) \right\|^2 + \mathbb{E}\left\| \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^t) \right\|^2 \\ & \leq \frac{1}{n} \sigma^2 + \mathbb{E}\left\| \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^t) \right\|^2, \end{aligned} \quad (\text{B.5})$$

where (a) holds due to $\mathbb{E}\|\mathbf{x}\|^2 = \mathbb{E}\|\mathbf{x} - \mathbb{E}[\mathbf{x}]\|^2 + \|\mathbb{E}[\mathbf{x}]\|^2$.

Substitute (B.4) and (B.5) into (B.3), we have

$$\begin{aligned} \mathbb{E}f(\bar{\mathbf{w}}^{t+1}) &\leq \mathbb{E}f(\bar{\mathbf{w}}^t) + \frac{\gamma^2 L}{2n} \sigma^2 - \frac{\gamma}{2} \mathbb{E}\|\nabla f(\bar{\mathbf{w}}^t)\|^2 \\ &\quad - \left(\frac{\gamma}{2} - \frac{\gamma^2 L}{2}\right) \mathbb{E}\left\|\frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^t)\right\|^2 + \frac{\gamma}{2} \mathbb{E}\left\|\nabla f(\bar{\mathbf{w}}^t) - \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^t)\right\|^2. \end{aligned} \quad (\text{B.6})$$

Suppose $\gamma \leq \frac{1}{L}$, that is, $\frac{\gamma}{2} - \frac{\gamma^2 L}{2} \geq 0$. We can obtain

$$\begin{aligned} \mathbb{E}f(\bar{\mathbf{w}}^{t+1}) &\leq \mathbb{E}f(\bar{\mathbf{w}}^t) + \frac{\gamma^2 L}{2n} \sigma^2 - \frac{\gamma}{2} \mathbb{E}\|\nabla f(\bar{\mathbf{w}}^t)\|^2 \\ &\quad + \frac{\gamma}{2} \mathbb{E}\left\|\nabla f(\bar{\mathbf{w}}^t) - \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^t)\right\|^2. \end{aligned} \quad (\text{B.7})$$

Now we compute the upper bound of the last term of inequality (B.7)

$$\begin{aligned} &\mathbb{E}\left\|\nabla f(\bar{\mathbf{w}}^t) - \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^t)\right\|^2 \\ &\leq \mathbb{E}\left\|\nabla f(\bar{\mathbf{w}}^t) - \sum_{i=1}^N \frac{n_i}{n} \nabla f_i(\bar{\mathbf{w}}_i^t) + \sum_{i=1}^N \frac{n_i}{n} \nabla f_i(\bar{\mathbf{w}}_i^t) - \sum_{i=1}^N \frac{n_i}{n} \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^t)\right\|^2 \\ &\leq 2\mathbb{E}\left\|\nabla f(\bar{\mathbf{w}}^t) - \sum_{i=1}^N \frac{n_i}{n} \nabla f_i(\bar{\mathbf{w}}_i^t)\right\|^2 + 2\mathbb{E}\left\|\sum_{i=1}^N \frac{n_i}{n} \nabla f_i(\bar{\mathbf{w}}_i^t) - \sum_{i=1}^N \frac{n_i}{n} \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^t)\right\|^2 \\ &\stackrel{(a)}{\leq} 2L^2 \sum_{l=1}^N \frac{n_l}{n} \mathbb{E}\|\bar{\mathbf{w}}^t - \bar{\mathbf{w}}_l^t\|^2 + 2L^2 \frac{1}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E}\|\bar{\mathbf{w}}_i^t - \mathbf{w}_k^t\|^2, \end{aligned} \quad (\text{B.8})$$

where (a) is due to Lipschitz gradient Assumption 1(a) in the main paper and (B.1).

Substituting (B.8) to (B.7) and rearranging the order, we have

$$\begin{aligned} \frac{\gamma}{2} \mathbb{E}\|\nabla f(\bar{\mathbf{w}}^t)\|^2 &\leq \mathbb{E}f(\bar{\mathbf{w}}^t) - \mathbb{E}f(\bar{\mathbf{w}}^{t+1}) + \frac{\gamma^2 L}{2n} \sigma^2 + \frac{\gamma}{2} \left[2L^2 \sum_{l=1}^N \frac{n_l}{n} \mathbb{E}\|\bar{\mathbf{w}}^t - \bar{\mathbf{w}}_l^t\|^2 \right. \\ &\quad \left. + 2L^2 \frac{1}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E}\|\bar{\mathbf{w}}_i^t - \mathbf{w}_k^t\|^2 \right]. \end{aligned} \quad (\text{B.9})$$

Dividing both sides by $\frac{\gamma}{2}$ and taking the average over time, we have

$$\begin{aligned} &\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}\|\nabla f(\bar{\mathbf{w}}^t)\|^2 \\ &\leq \frac{2}{\gamma T} \left[f(\bar{\mathbf{w}}^0) - \mathbb{E}f(\bar{\mathbf{w}}^T) \right] + \frac{\gamma L \sigma^2}{n} + 2L^2 \frac{1}{T} \sum_{t=0}^{T-1} \sum_{l=1}^N \frac{n_l}{n} \mathbb{E}\|\bar{\mathbf{w}}^t - \bar{\mathbf{w}}_l^t\|^2 \\ &\quad + 2L^2 \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E}\|\bar{\mathbf{w}}_i^t - \mathbf{w}_k^t\|^2 \\ &\leq \frac{2}{\gamma T} \left[f(\bar{\mathbf{w}}^0) - f^* \right] + \frac{\gamma L \sigma^2}{n} + 2L^2 \frac{1}{T} \sum_{t=0}^{T-1} \sum_{l=1}^N \frac{n_l}{n} \mathbb{E}\|\bar{\mathbf{w}}^t - \bar{\mathbf{w}}_l^t\|^2 \\ &\quad + 2L^2 \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E}\|\bar{\mathbf{w}}_i^t - \mathbf{w}_k^t\|^2, \end{aligned} \quad (\text{B.10})$$

where $f^* = \min_{\mathbf{w}} f(\mathbf{w})$.

Note that (B.10) is the key equation that we will use to prove Theorem 1. It involves two major terms. The first is $\frac{1}{T} \sum_{t=0}^{T-1} \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \|\bar{\mathbf{w}}^t - \bar{\mathbf{w}}_l^t\|^2$, which represents the upward parameter mean-square-error (MSE) between the globally aggregated parameter $\bar{\mathbf{w}}^t$ and the locally aggregated parameters $\{\bar{\mathbf{w}}_l^t, l = 1, \dots, N\}$ from N groups. The second term is $\frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E} \|\bar{\mathbf{w}}_i^t - \mathbf{w}_k^t\|^2$, which represents the downward parameter MSE between each local parameter \mathbf{w}_j^t and its locally aggregated parameter $\bar{\mathbf{w}}_i^t$ for each worker j that belongs to group i . Next, we will analyze these two terms separately.

B.2 Bounding Upward Parameter MSE $\frac{1}{T} \sum_{t=0}^{T-1} \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \|\bar{\mathbf{w}}^t - \bar{\mathbf{w}}_l^t\|^2$ in (B.8)

In this section, we will bound $\frac{1}{T} \sum_{t=0}^{T-1} \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \|\bar{\mathbf{w}}^t - \bar{\mathbf{w}}_l^t\|^2$ in terms of global divergence $\epsilon_{\mathbf{w}}^2$ and σ^2 . Note that when $t = aG, a = 0, 1, 2, \dots$, we have $\bar{\mathbf{w}}^t = \bar{\mathbf{w}}_l^t, l = 1, 2, \dots, N$. Suppose $t = aG + c, c = 0, 1, \dots$. When $1 \leq c \leq G-1$, we have

$$\begin{aligned}
& \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \|\bar{\mathbf{w}}^t - \bar{\mathbf{w}}_l^t\|^2 \\
&= \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \left\| \left(\bar{\mathbf{w}}^{t-1} - \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^{t-1}) \right) - \left(\bar{\mathbf{w}}_l^{t-1} - \frac{\gamma}{n_l} \sum_{m \in \mathcal{V}_l} \mathbf{g}_m(\mathbf{w}_m^{t-1}) \right) \right\|^2 \\
&= \frac{\gamma^2}{N} \sum_{l=1}^N \mathbb{E} \left\| \sum_{\tau=aG}^{aG+c-1} \left(\frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \mathbf{g}_m(\mathbf{w}_m^\tau) - \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^\tau) \right) \right\|^2 \\
&= \frac{\gamma^2}{N} \sum_{l=1}^N \mathbb{E} \left\| \sum_{\tau=aG}^{aG+c-1} \left(\frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \mathbf{g}_m(\mathbf{w}_m^\tau) - \nabla f_l(\bar{\mathbf{w}}_l^\tau) + \nabla f_l(\bar{\mathbf{w}}_l^\tau) - \sum_{i=1}^N \frac{n_i}{n} \nabla f_i(\bar{\mathbf{w}}_i^\tau) + \sum_{i=1}^N \frac{n_i}{n} \nabla f_i(\bar{\mathbf{w}}_i^\tau) \right. \right. \\
&\quad \left. \left. - \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^\tau) \right) \right\|^2 \\
&\stackrel{(a)}{\leq} 2\gamma^2 \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \left\| \sum_{\tau=aG}^{aG+c-1} \left(\left(\frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \mathbf{g}_m(\mathbf{w}_m^\tau) - \nabla f_l(\bar{\mathbf{w}}_l^\tau) \right) \right. \right. \\
&\quad \left. \left. - \left(\frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^\tau) - \sum_{i=1}^N \frac{n_i}{n} \nabla f_i(\bar{\mathbf{w}}_i^\tau) \right) \right) \right\|^2 + 2\gamma^2 \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \left\| \sum_{\tau=aG}^{aG+c-1} \left(\nabla f_l(\bar{\mathbf{w}}_l^\tau) - \sum_{i=1}^N \frac{n_i}{n} \nabla f_i(\bar{\mathbf{w}}_i^\tau) \right) \right\|^2, \quad (\text{B.11})
\end{aligned}$$

where (a) is due to (B.1). Next, we will derive an upper bound on each of the two terms in (B.11) separately.

For the first term of (B.11), we have

$$\begin{aligned}
& 2\gamma^2 \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \left\| \sum_{\tau=aG}^{aG+c-1} \left(\left(\frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \mathbf{g}_m(\mathbf{w}_m^\tau) - \nabla f_l(\bar{\mathbf{w}}_l^\tau) \right) - \left(\frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^\tau) - \sum_{i=1}^N \frac{n_i}{n} \nabla f_i(\bar{\mathbf{w}}_i^\tau) \right) \right) \right\|^2 \\
&= 2\gamma^2 \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \left\| \sum_{\tau=aG}^{aG+c-1} \left(\left(\frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \mathbf{g}_m(\mathbf{w}_m^\tau) - \frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \nabla F_m(\mathbf{w}_m^\tau) + \frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \nabla F_m(\mathbf{w}_m^\tau) - \nabla f_l(\bar{\mathbf{w}}_l^\tau) \right) \right. \right. \\
&\quad \left. \left. - \left(\frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^\tau) - \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^\tau) + \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^\tau) - \sum_{i=1}^N \frac{n_i}{n} \nabla f_i(\bar{\mathbf{w}}_i^\tau) \right) \right) \right\|^2 \\
&= 2\gamma^2 \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \left\| \sum_{\tau=aG}^{aG+c-1} \left(\left(\frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \mathbf{g}_m(\mathbf{w}_m^\tau) - \frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \nabla F_m(\mathbf{w}_m^\tau) \right) - \left(\frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^\tau) - \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^\tau) \right) \right. \right. \\
&\quad \left. \left. + \left(\frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \nabla F_m(\mathbf{w}_m^\tau) - \nabla f_l(\bar{\mathbf{w}}_l^\tau) \right) - \left(\frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^\tau) - \sum_{i=1}^N \frac{n_i}{n} \nabla f_i(\bar{\mathbf{w}}_i^\tau) \right) \right) \right\|^2 \\
&\stackrel{(a)}{\leq} 4\gamma^2 \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \left\| \sum_{\tau=aG}^{aG+c-1} \left(\left(\frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \mathbf{g}_m(\mathbf{w}_m^\tau) - \frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \nabla F_m(\mathbf{w}_m^\tau) \right) - \left(\frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^\tau) - \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^\tau) \right) \right) \right\|^2
\end{aligned}$$

$$+ 4\gamma^2 \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \left\| \sum_{\tau=aG}^{aG+c-1} \left(\frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \nabla F_m(\mathbf{w}_m^\tau) - \nabla f_l(\bar{\mathbf{w}}_l^\tau) \right) \right\|^2, \quad (\text{B.12})$$

where (a) follows from (B.1) with $M = 2$ and (B.2).

For the first term of (B.12), we have

$$\begin{aligned} & 4\gamma^2 \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \left\| \sum_{\tau=aG}^{aG+c-1} \left(\left(\frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \mathbf{g}_m(\mathbf{w}_m^\tau) - \frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \nabla F_m(\mathbf{w}_m^\tau) \right) - \left(\frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^\tau) - \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^\tau) \right) \right) \right\|^2 \\ & \stackrel{(a)}{=} 4\gamma^2 \left(\sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \left\| \sum_{\tau=aG}^{aG+c-1} \left(\frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \mathbf{g}_m(\mathbf{w}_m^\tau) - \frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \nabla F_m(\mathbf{w}_m^\tau) \right) \right\|^2 \right. \\ & \quad \left. - \mathbb{E} \left\| \sum_{\tau=aG}^{aG+c-1} \left(\frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^\tau) - \frac{1}{n} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^\tau) \right) \right\|^2 \right) \\ & \stackrel{(b)}{=} 4\gamma^2 \sum_{\tau=aG}^{aG+c-1} \left(\sum_{l=1}^N \frac{n_l}{n} \frac{1}{n_l^2} \sum_{m \in \mathcal{V}_l} \mathbb{E} \left\| \mathbf{g}_m(\mathbf{w}_m^\tau) - \nabla F_m(\mathbf{w}_m^\tau) \right\|^2 - \frac{1}{n^2} \sum_{i=1}^N \sum_{j \in \mathcal{V}_i} \mathbb{E} \left\| \mathbf{g}_j(\mathbf{w}_j^\tau) - \nabla F_j(\mathbf{w}_j^\tau) \right\|^2 \right) \\ & \stackrel{(c)}{=} 4\gamma^2 \sum_{\tau=aG}^{aG+c-1} \frac{1}{n} \sum_{i=1}^N \left(\frac{1}{n_i} - \frac{1}{n} \right) \sum_{j \in \mathcal{V}_i} \mathbb{E} \left\| \mathbf{g}_j(\mathbf{w}_j^\tau) - \nabla F_j(\mathbf{w}_j^\tau) \right\|^2 \\ & \leq 4\gamma^2 G \frac{N-1}{n} \sigma^2 \quad (\text{use Assumption 1b}), \end{aligned} \quad (\text{B.13})$$

where (a) is due to (B.2), (b) and (c) hold due to the independence of SGD noise over different time, groups and nodes, and $\mathbb{E}[\mathbf{g}_i(\mathbf{w}_i^t)] = \mathbb{E}[\mathbb{E}[\mathbf{g}_i(\mathbf{w}_i^t) | \mathbf{w}_i^t]] = \mathbb{E}[\nabla F_i(\mathbf{w}_i^t)]$, $\forall i, t$. This implies that $\mathbb{E}[\langle \mathbf{g}_i(\mathbf{w}_i^t) - \nabla F_i(\mathbf{w}_i^t), \mathbf{g}_j(\mathbf{w}_j^{t'}) - \nabla F_j(\mathbf{w}_j^{t'}) \rangle] = 0$, $i \neq j$.

For the second term of (B.12), we have

$$\begin{aligned} & 4\gamma^2 \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \left\| \sum_{\tau=aG}^{aG+c-1} \left(\frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \nabla F_m(\mathbf{w}_m^\tau) - \nabla f_l(\bar{\mathbf{w}}_l^\tau) \right) \right\|^2 \\ & \leq 4\gamma^2 G \sum_{l=1}^N \frac{n_l}{n} \sum_{\tau=aG}^{aG+c-1} \mathbb{E} \left\| \frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \left(\nabla F_m(\mathbf{w}_m^\tau) - \nabla F_m(\bar{\mathbf{w}}_l^\tau) \right) \right\|^2 \quad (\text{use (B.1) with } M = G) \\ & \leq 4\gamma^2 GL^2 \sum_{\tau=aG}^{aG+c-1} \sum_{l=1}^N \frac{n_l}{n} \frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \mathbb{E} \|\mathbf{w}_m^\tau - \bar{\mathbf{w}}_l^\tau\|^2 \quad (\text{use (B.1) with } M = N_\ell \text{ and Assumption 1a}). \end{aligned} \quad (\text{B.14})$$

For the second term of (B.11), we have

$$\begin{aligned} & 2\gamma^2 \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \left\| \sum_{\tau=aG}^{aG+c-1} \left(\nabla f_l(\bar{\mathbf{w}}_l^\tau) - \sum_{i=1}^N \frac{n_i}{n} \nabla f_i(\bar{\mathbf{w}}_i^\tau) \right) \right\|^2 \\ & \leq 2\gamma^2 \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \left\| \sum_{\tau=aG}^{aG+c-1} \left(\nabla f_l(\bar{\mathbf{w}}_l^\tau) - \nabla f(\bar{\mathbf{w}}_l^\tau) + \nabla f(\bar{\mathbf{w}}_l^\tau) - \nabla f(\bar{\mathbf{w}}^\tau) + \nabla f(\bar{\mathbf{w}}^\tau) - \sum_{i=1}^N \frac{n_i}{n} \nabla f_i(\bar{\mathbf{w}}_i^\tau) \right) \right\|^2 \\ & \stackrel{(a)}{\leq} 6\gamma^2 G^2 \epsilon^2 + 6\gamma^2 L^2 G \sum_{\tau=aG}^{aG+c-1} \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \|\bar{\mathbf{w}}_l^\tau - \bar{\mathbf{w}}^\tau\|^2 + 6\gamma^2 L^2 G \sum_{\tau=aG}^{aG+c-1} \sum_{i=1}^N \frac{n_i}{n} \mathbb{E} \|\bar{\mathbf{w}}^\tau - \bar{\mathbf{w}}_i^\tau\|^2 \\ & = 6\gamma^2 G^2 \epsilon^2 + 12\gamma^2 L^2 G \sum_{\tau=aG}^{aG+c-1} \sum_{i=1}^N \frac{n_i}{n} \mathbb{E} \|\bar{\mathbf{w}}^\tau - \bar{\mathbf{w}}_i^\tau\|^2, \end{aligned} \quad (\text{B.15})$$

where (a) uses (B.1), and Assumptions 1c and 1a.

Now back to inequalities (B.11), with results of inequality (B.12), (B.13), (B.14), (B.15), we can obtain

$$\sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \|\bar{\mathbf{w}}^t - \bar{\mathbf{w}}_l^t\|^2$$

$$\begin{aligned}
&\leq 4\gamma^2 G \frac{N-1}{n} \sigma^2 + 6\gamma^2 G^2 \epsilon^2 + 4\gamma^2 GL^2 \sum_{\tau=aG}^{t-1} \sum_{l=1}^N \frac{n_l}{n} \frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \mathbb{E} \|\mathbf{w}_m^\tau - \bar{\mathbf{w}}_l^\tau\|^2 \\
&\quad + 12\gamma^2 L^2 G \sum_{\tau=aG}^{t-1} \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \|\bar{\mathbf{w}}_l^\tau - \bar{\mathbf{w}}^\tau\|^2 \\
&\stackrel{(a)}{\leq} 4\gamma^2 G \frac{N-1}{n} \sigma^2 + 6\gamma^2 G^2 \epsilon^2 + 4\gamma^2 GL^2 \sum_{\tau=aG}^{aG+G-1} \sum_{l=1}^N \frac{n_l}{n} \frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \mathbb{E} \|\mathbf{w}_m^\tau - \bar{\mathbf{w}}_l^\tau\|^2 \\
&\quad + 12\gamma^2 L^2 G \sum_{\tau=aG}^{aG+G-1} \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \|\bar{\mathbf{w}}_l^\tau - \bar{\mathbf{w}}^\tau\|^2, \tag{B.16}
\end{aligned}$$

where (a) is due to $1 \leq c \leq G-1$. Since $\|\bar{\mathbf{w}}^t - \bar{\mathbf{w}}_l^t\|^2 = 0$, (a) also holds when $t = aG, a = 0, 1, 2, \dots$. Suppose $T-1 = dG, d = 0, 1, 2, \dots$. Take the average over time on both sides of (B.16), we have

$$\begin{aligned}
&\frac{1}{T} \sum_{t=0}^{T-1} \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \|\bar{\mathbf{w}}^t - \bar{\mathbf{w}}_l^t\|^2 \\
&\stackrel{(a)}{=} \frac{1}{T} \sum_{t=0}^{T-2} \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \|\bar{\mathbf{w}}^t - \bar{\mathbf{w}}_l^t\|^2 \\
&\leq 4\gamma^2 G \frac{N-1}{n} \sigma^2 + 6\gamma^2 G^2 \epsilon^2 + 4\gamma^2 GL^2 \frac{1}{T} \sum_{a=0}^{d-1} \left(G \sum_{\tau=aG}^{aG+G-1} \sum_{l=1}^N \frac{n_l}{n} \frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \mathbb{E} \|\mathbf{w}_m^\tau - \bar{\mathbf{w}}_l^\tau\|^2 \right) \\
&\quad + 12\gamma^2 L^2 G \frac{1}{T} \sum_{a=0}^{d-1} \left(G \sum_{\tau=aG}^{aG+G-1} \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \|\bar{\mathbf{w}}_l^\tau - \bar{\mathbf{w}}^\tau\|^2 \right) \\
&= 4\gamma^2 G \frac{N-1}{n} \sigma^2 + 6\gamma^2 G^2 \epsilon^2 + 4\gamma^2 G^2 L^2 \frac{1}{T} \sum_{t=0}^{T-2} \sum_{l=1}^N \frac{n_l}{n} \frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \mathbb{E} \|\mathbf{w}_m^t - \bar{\mathbf{w}}_l^t\|^2 \\
&\quad + 12\gamma^2 L^2 G^2 \frac{1}{T} \sum_{t=0}^{T-2} \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \|\bar{\mathbf{w}}_l^t - \bar{\mathbf{w}}^t\|^2 \\
&\stackrel{(b)}{=} 4\gamma^2 G \frac{N-1}{n} \sigma^2 + 6\gamma^2 G^2 \epsilon^2 + 4\gamma^2 G^2 L^2 \frac{1}{T} \sum_{t=0}^{T-1} \sum_{l=1}^N \frac{n_l}{n} \frac{1}{n_l} \sum_{m \in \mathcal{V}_l} \mathbb{E} \|\mathbf{w}_m^t - \bar{\mathbf{w}}_l^t\|^2 \\
&\quad + 12\gamma^2 L^2 G^2 \frac{1}{T} \sum_{t=0}^{T-1} \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \|\bar{\mathbf{w}}_l^t - \bar{\mathbf{w}}^t\|^2, \tag{B.17}
\end{aligned}$$

where (a) and (b) hold due to $\|\mathbf{w}_m^{T-1} - \bar{\mathbf{w}}_l^{T-1}\|^2 = 0$ and $\|\bar{\mathbf{w}}_l^{T-1} - \bar{\mathbf{w}}^{T-1}\|^2 = 0$. Suppose $1 - 12\gamma^2 L^2 G^2 > 0$, move the last term of (B.17) to the left, we can obtain

$$\begin{aligned}
&(1 - 12\gamma^2 L^2 G^2) \frac{1}{T} \sum_{t=0}^{T-1} \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \|\bar{\mathbf{w}}^t - \bar{\mathbf{w}}_l^t\|^2 \\
&\leq 4\gamma^2 G \frac{N-1}{n} \sigma^2 + 6\gamma^2 G^2 \epsilon^2 + \frac{1}{T} \sum_{t=0}^{T-1} 4\gamma^2 G^2 L^2 \frac{1}{n} \sum_{l=1}^N \sum_{m \in \mathcal{V}_l} \mathbb{E} \|\mathbf{w}_m^t - \bar{\mathbf{w}}_l^t\|^2, \tag{B.18}
\end{aligned}$$

then

$$\begin{aligned}
&\frac{1}{T} \sum_{t=0}^{T-1} \sum_{l=1}^N \frac{n_l}{n} \mathbb{E} \|\bar{\mathbf{w}}^t - \bar{\mathbf{w}}_l^t\|^2 \\
&\leq \frac{4\gamma^2 G \frac{N-1}{n} \sigma^2}{1 - 12\gamma^2 L^2 G^2} + \frac{6\gamma^2 G^2 \epsilon^2}{1 - 12\gamma^2 L^2 G^2}
\end{aligned}$$

$$+ \frac{4\gamma^2 G^2 L^2}{1 - 12\gamma^2 L^2 G^2} \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{n} \sum_{l=1}^N \sum_{m \in \mathcal{V}_l} \mathbb{E} \|\mathbf{w}_m^t - \bar{\mathbf{w}}_l^t\|^2. \quad (\text{B.19})$$

The importance of (B.19) is that it presents an upper bound on the upward parameter MSE in terms of the downward parameter MSE. Next, we will proceed to analyze the downward parameter MSE.

B.3 Bounding Downward Parameter MSE $\frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E} \|\bar{\mathbf{w}}_i^t - \mathbf{w}_k^t\|^2$ in (B.8)

Now we consider the local aggregation. Suppose $t = b_i I_i + c_i$ for $i = 1, 2, 3, \dots, N$ where $b_i = 0, 1, \dots, \frac{G}{I_i} - 1$ and $c_i = 0, 1, 2, \dots, I_i - 1$. When $c_i = 0$, we have $\bar{\mathbf{w}}_i^t = \mathbf{w}_k^t, k \in \mathcal{V}_i, \forall i$. When $c_i \geq 1$, we have

$$\begin{aligned} & \frac{1}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E} \|\bar{\mathbf{w}}_i^t - \mathbf{w}_k^t\|^2 \\ &= \frac{1}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E} \left\| \left(\bar{\mathbf{w}}_i^{t-1} - \gamma \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^{t-1}) \right) - \left(\mathbf{w}_k^{t-1} - \gamma \mathbf{g}_k(\mathbf{w}_k^{t-1}) \right) \right\|^2 \\ &= \frac{\gamma^2}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E} \left\| \sum_{\tau=b_i I_i}^{t-1} \left(\mathbf{g}_k(\mathbf{w}_k^\tau) - \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^\tau) \right) \right\|^2 \\ &= \frac{\gamma^2}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E} \left\| \sum_{\tau=b_i I_i}^{t-1} \left(\mathbf{g}_k(\mathbf{w}_k^\tau) - \nabla F_k(\mathbf{w}_k^\tau) + \nabla F_k(\mathbf{w}_k^\tau) - \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^\tau) + \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^\tau) - \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^\tau) \right) \right\|^2 \\ &= \frac{\gamma^2}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E} \left\| \sum_{\tau=b_i I_i}^{t-1} \left(\left(\mathbf{g}_k(\mathbf{w}_k^\tau) - \nabla F_k(\mathbf{w}_k^\tau) + \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^\tau) - \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^\tau) \right) + \left(\nabla F_k(\mathbf{w}_k^\tau) - \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^\tau) \right) \right) \right\|^2 \\ &\leq \frac{2\gamma^2}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E} \left\| \sum_{\tau=b_i I_i}^{t-1} \left(\left(\mathbf{g}_k(\mathbf{w}_k^\tau) - \nabla F_k(\mathbf{w}_k^\tau) \right) - \left(\frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^\tau) - \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^\tau) \right) \right) \right\|^2 \end{aligned} \quad (\text{B.20})$$

$$+ \frac{2\gamma^2}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E} \left\| \sum_{\tau=b_i I_i}^{t-1} \left(\nabla F_k(\mathbf{w}_k^\tau) - \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^\tau) \right) \right\|^2. \quad (\text{B.21})$$

For (B.20), we have

$$\begin{aligned} & \frac{2\gamma^2}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E} \left\| \sum_{\tau=b_i I_i}^{t-1} \left(\left(\mathbf{g}_k(\mathbf{w}_k^\tau) - \nabla F_k(\mathbf{w}_k^\tau) \right) - \left(\frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^\tau) - \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^\tau) \right) \right) \right\|^2 \\ &= \frac{2\gamma^2}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E} \left\| \sum_{\tau=b_i I_i}^{t-1} \left(\mathbf{g}_k(\mathbf{w}_k^\tau) - \nabla F_k(\mathbf{w}_k^\tau) \right) \right\|^2 - 2\gamma^2 \sum_{i=1}^N \frac{n_i}{n} \mathbb{E} \left\| \sum_{\tau=b_i I_i}^{t-1} \left(\frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \mathbf{g}_j(\mathbf{w}_j^\tau) - \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^\tau) \right) \right\|^2 \\ &= \frac{2\gamma^2}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \sum_{\tau=b_i I_i}^{t-1} \mathbb{E} \left\| \mathbf{g}_k(\mathbf{w}_k^\tau) - \nabla F_k(\mathbf{w}_k^\tau) \right\|^2 - 2\gamma^2 \sum_{i=1}^N \frac{n_i}{n} \frac{1}{n_i^2} \sum_{j \in \mathcal{V}_i} \sum_{\tau=b_i I_i}^{t-1} \mathbb{E} \left\| \mathbf{g}_j(\mathbf{w}_j^\tau) - \nabla F_j(\mathbf{w}_j^\tau) \right\|^2 \\ &= \frac{2\gamma^2}{n} \sum_{i=1}^N \left(1 - \frac{1}{n_i} \right) \sum_{k \in \mathcal{V}_i} \sum_{\tau=b_i I_i}^{t-1} \mathbb{E} \left\| \mathbf{g}_k(\mathbf{w}_k^\tau) - \nabla F_k(\mathbf{w}_k^\tau) \right\|^2 \\ &\leq 2\gamma^2 \sigma^2 \sum_{i=1}^N \frac{(n_i - 1)I_i}{n}. \end{aligned} \quad (\text{B.22})$$

For (B.21), we have

$$\frac{2\gamma^2}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E} \left\| \sum_{\tau=b_i I_i}^{t-1} \left(\nabla F_k(\mathbf{w}_k^\tau) - \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^\tau) \right) \right\|^2$$

$$\begin{aligned}
&= \frac{2\gamma^2}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E} \left\| \sum_{\tau=b_i I_i}^{t-1} \left(\nabla F_k(\mathbf{w}_k^\tau) - \nabla F_k(\bar{\mathbf{w}}_i^\tau) + \nabla F_k(\bar{\mathbf{w}}_i^\tau) - \nabla f_i(\bar{\mathbf{w}}_i^\tau) + \nabla f_i(\bar{\mathbf{w}}_i^\tau) - \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^\tau) \right) \right\|^2 \\
&\leq \frac{6\gamma^2}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E} \left\| \sum_{\tau=b_i I_i}^{t-1} \left(\nabla F_k(\mathbf{w}_k^\tau) - \nabla F_k(\bar{\mathbf{w}}_i^\tau) \right) \right\|^2 + \frac{6\gamma^2}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E} \left\| \sum_{\tau=b_i I_i}^{t-1} \left(\nabla F_k(\bar{\mathbf{w}}_i^\tau) - \nabla f_i(\bar{\mathbf{w}}_i^\tau) \right) \right\|^2 \\
&\quad + 6\gamma^2 \sum_{i=1}^N \frac{n_i}{n} \mathbb{E} \left\| \sum_{\tau=b_i I_i}^{t-1} \left(\nabla f_i(\bar{\mathbf{w}}_i^\tau) - \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} \nabla F_j(\mathbf{w}_j^\tau) \right) \right\|^2 \\
&\stackrel{(a)}{\leq} \frac{12\gamma^2}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E} \left\| \sum_{\tau=b_i I_i}^{t-1} \left(\nabla F_k(\mathbf{w}_k^\tau) - \nabla F_k(\bar{\mathbf{w}}_i^\tau) \right) \right\|^2 + 6\gamma^2 \sum_{i=1}^N \frac{n_i}{n} I_i^2 \epsilon_i^2 \\
&\leq \frac{12\gamma^2 L^2}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \sum_{\tau=b_i I_i}^{t-1} I_i \mathbb{E} \left\| \mathbf{w}_k^\tau - \bar{\mathbf{w}}_i^\tau \right\|^2 + 6\gamma^2 \sum_{i=1}^N \frac{n_i}{n} I_i^2 \epsilon_i^2,
\end{aligned} \tag{B.23}$$

where (a) is due to (B.1).

Substituting (B.22) and (B.23) back to (B.20) and (B.21), we have

$$\begin{aligned}
&\frac{1}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E} \left\| \bar{\mathbf{w}}_i^t - \mathbf{w}_k^t \right\|^2 \\
&\leq 2\gamma^2 \sigma^2 \sum_{i=1}^N \frac{(n_i - 1)I_i}{n} + 6\gamma^2 \sum_{i=1}^N \frac{n_i}{n} I_i^2 \epsilon_i^2 + \frac{12\gamma^2 L^2}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \sum_{\tau=b_i I_i}^{t-1} I_i \mathbb{E} \left\| \mathbf{w}_k^\tau - \bar{\mathbf{w}}_i^\tau \right\|^2 \\
&\stackrel{(a)}{\leq} 2\gamma^2 \sigma^2 \sum_{i=1}^N \frac{(n_i - 1)I_i}{n} + 6\gamma^2 \sum_{i=1}^N \frac{n_i}{n} I_i^2 \epsilon_i^2 + \frac{12\gamma^2 L^2}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \sum_{\tau=b_i I_i}^{b_i I_i + I_i - 1} I_i \mathbb{E} \left\| \mathbf{w}_k^\tau - \bar{\mathbf{w}}_i^\tau \right\|^2
\end{aligned} \tag{B.24}$$

where (a) holds due to $c_i \leq I_i - 1$. (B.24) also holds when $t = b_i I_i$ since $\frac{1}{n_i} \sum_{k \in \mathcal{V}_i} \mathbb{E} \left\| \bar{\mathbf{w}}_i^{b_i I_i} - \mathbf{w}_k^{b_i I_i} \right\|^2 = 0$.

Since $T - 1 = dG$, $d = 0, 1, \dots, T - 1$ is also an integral multiple of I_i . That is, $T - 1 = f_i I_i$. Taking the average over time yields

$$\begin{aligned}
&\frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E} \left\| \bar{\mathbf{w}}_i^t - \mathbf{w}_k^t \right\|^2 \\
&\leq 2\gamma^2 \sigma^2 \sum_{i=1}^N \frac{(n_i - 1)I_i}{n} + 6\gamma^2 \sum_{i=1}^N \frac{n_i}{n} I_i^2 \epsilon_i^2 + \frac{12\gamma^2 L^2}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \sum_{b_i=0}^{f_i-1} \sum_{\tau=b_i I_i}^{b_i I_i + I_i - 1} I_i \mathbb{E} \left\| \mathbf{w}_k^\tau - \bar{\mathbf{w}}_i^\tau \right\|^2 \\
&\leq 2\gamma^2 \sigma^2 \sum_{i=1}^N \frac{(n_i - 1)I_i}{n} + 6\gamma^2 \sum_{i=1}^N \frac{n_i}{n} I_i^2 \epsilon_i^2 + \frac{1}{T} \sum_{\tau=0}^{T-1} \frac{12\gamma^2 L^2}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} I_i^2 \mathbb{E} \left\| \mathbf{w}_k^\tau - \bar{\mathbf{w}}_i^\tau \right\|^2 \\
&\leq 2\gamma^2 \sigma^2 \sum_{i=1}^N \frac{(n_i - 1)I_i}{n} + 6\gamma^2 \sum_{i=1}^N \frac{n_i}{n} I_i^2 \epsilon_i^2 + \frac{1}{T} \sum_{\tau=0}^{T-1} \frac{12\gamma^2 L^2 I_{max}^2}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E} \left\| \mathbf{w}_k^\tau - \bar{\mathbf{w}}_i^\tau \right\|^2,
\end{aligned} \tag{B.25}$$

where $I_{max} = \max_i I_i$. Suppose $1 - 12\gamma^2 L^2 I_{max}^2 > 0$, we have

$$\begin{aligned}
&\frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{n} \sum_{i=1}^N \sum_{k \in \mathcal{V}_i} \mathbb{E} \left\| \bar{\mathbf{w}}_i^t - \mathbf{w}_k^t \right\|^2 \\
&\leq \frac{2\gamma^2 \sigma^2 \sum_{i=1}^N \frac{(n_i-1)I_i}{n}}{1 - 12\gamma^2 L^2 I_{max}^2} + \frac{6\gamma^2 \sum_{i=1}^N \frac{n_i}{n} I_i^2 \epsilon_i^2}{1 - 12\gamma^2 L^2 I_{max}^2}.
\end{aligned} \tag{B.26}$$

Obtaining the Final Result Substituting (B.26) into (B.19), we obtain an upper bound on the upward parameter MSE as

$$\frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{N} \sum_{l=1}^N \mathbb{E} \left\| \bar{\mathbf{w}}^t - \bar{\mathbf{w}}_l^t \right\|^2$$

$$\begin{aligned}
&\leq \frac{4\gamma^2 G^{\frac{N-1}{n}} \sigma^2}{1 - 12\gamma^2 L^2 G^2} + \frac{6\gamma^2 G^2 \epsilon^2}{1 - 12\gamma^2 L^2 G^2} \\
&\quad + \frac{4\gamma^2 G^2 L^2}{1 - 12\gamma^2 L^2 G^2} \left(\frac{2\gamma^2 \sigma^2 \sum_{i=1}^N \frac{(n_i-1)I_i}{n}}{1 - 12\gamma^2 L^2 I_{max}^2} + \frac{6\gamma^2 \sum_{i=1}^N \frac{n_i}{n} I_i^2 \epsilon_i^2}{1 - 12\gamma^2 L^2 I_{max}^2} \right). \tag{B.27}
\end{aligned}$$

Taking the average over time and substitute (B.26) and (B.27) to (B.10), we obtain the final convergence upper bound:

$$\begin{aligned}
&\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\bar{\mathbf{w}}^t)\|^2 \\
&\leq \frac{2}{\gamma T} \left(\mathbb{E} f(\bar{\mathbf{w}}^0) - f^* \right) + \frac{\gamma L \sigma^2}{n} \\
&\quad + \frac{8\gamma^2 G^{\frac{N-1}{n}} \sigma^2}{1 - 12\gamma^2 L^2 G^2} + \frac{12\gamma^2 G^2 \epsilon^2}{1 - 12\gamma^2 L^2 G^2} \\
&\quad + \left(1 + \frac{8\gamma^2 G^2 L^2}{1 - 12\gamma^2 L^2 G^2} \right) \left(\frac{4\gamma^2 \sigma^2 \sum_{i=1}^N \frac{(n_i-1)I_i}{n}}{1 - 12\gamma^2 L^2 I_{max}^2} + \frac{12\gamma^2 \sum_{i=1}^N \frac{n_i}{n} I_i^2 \epsilon_i^2}{1 - 12\gamma^2 L^2 I_{max}^2} \right). \tag{B.28}
\end{aligned}$$

When $0 < \gamma \leq \frac{1}{2\sqrt{6}GL}$, we have

$$0 < \frac{8\gamma^2 G^2 L^2}{1 - 12\gamma^2 L^2 G^2} \leq \frac{2}{3}, \tag{B.29}$$

and

$$0 < \frac{1}{1 - 12\gamma^2 L^2 I_{max}^2} \leq \frac{1}{1 - 12\gamma^2 L^2 G^2} \leq 2. \tag{B.30}$$

By substituting (B.29) and (B.30) into (B.27) of the main paper and setting $C = \frac{40}{3}$, we obtain the upper bound in Theorem 1. Note that the constant in front of the term $\sum_{i=1}^N \frac{n_i}{n} \left(1 - \frac{1}{n_i}\right) I_i$ is intentionally increased by a factor of 2 in order to better match the constants given in Theorem 1.

C Proofs for the Two-level Random Grouping Case

In the following, we will present the proofs of Lemmas 1 and 2 and Theorem 2. As discussed in Section 4.3 of the main paper, we consider uniformly random grouping here where the groups $\{\mathcal{V}_i : i = 1, 2, \dots, N\}$ are random sets of equal size, i.e., $n_i = |\mathcal{V}_i| = \frac{n}{N}$ for all i . For simplicity, we use $K := \frac{n}{N} = n_i$ to denote the size of each group. The uniformly random partitioning of workers into groups is equivalent to the following problem. Assume that there are n bins, where the first K bins contain workers in group $i = 1$, the next K bins contain workers in group $i = 2$, and so on. We randomly and uniformly assign the n workers to the n bins, such that each bin contains one (and only one) worker and each worker is assigned to one (and only one) bin. A random realization \mathbf{S} corresponds to a mapping between workers and bins.

Let k_j denote the worker assigned to bin j , where k_j is a random variable due to random assignment. Also let $[K] := \{1, 2, \dots, K\}$ and $[n] := \{1, 2, \dots, n\} = \mathcal{V}$ (the last equality is by definition of \mathcal{V}). We proceed with the proofs in the following.

We define

$$\epsilon_{\mathbf{w}}^2 := \mathbb{E}_{\mathbf{S}} \left(\|\nabla f(\mathbf{w}) - \nabla F_{k_j}(\mathbf{w})\|^2 \right) \stackrel{(a)}{=} \frac{1}{n} \sum_{k=1}^n \|\nabla f(\mathbf{w}) - \nabla F_k(\mathbf{w})\|^2 \stackrel{(b)}{\leq} \tilde{\epsilon}^2. \quad (\text{C.1})$$

where (a) is because for any given bin j , the probability of any worker being assigned to it is $\frac{1}{n}$, and (b) is due to Assumption 2 in the main paper.

C.1 Proof of Lemma 1

We have

$$\begin{aligned} & \mathbb{E}_{\mathbf{S}} \left[\sum_{i=1}^N \frac{n_i}{n} \left\| \nabla f(\mathbf{w}) - \frac{1}{K} \sum_{k \in \mathcal{V}_i} \nabla F_k(\mathbf{w}) \right\|^2 \right] \\ &= \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\mathbf{S}} \left\| \nabla f(\mathbf{w}) - \frac{1}{K} \sum_{k \in \mathcal{V}_i} \nabla F_k(\mathbf{w}) \right\|^2 \\ &\stackrel{(a)}{=} \mathbb{E}_{\mathbf{S}} \left\| \nabla f(\mathbf{w}) - \frac{1}{K} \sum_{k \in \mathcal{V}_1} \nabla F_k(\mathbf{w}) \right\|^2 \\ &\stackrel{(b)}{=} \mathbb{E}_{\mathbf{S}} \left\| \nabla f(\mathbf{w}) - \frac{1}{K} \sum_{j=1}^K \nabla F_{k_j}(\mathbf{w}) \right\|^2 \\ &= \frac{1}{K^2} \mathbb{E}_{\mathbf{S}} \left\| \sum_{j=1}^K (\nabla f(\mathbf{w}) - \nabla F_{k_j}(\mathbf{w})) \right\|^2 \\ &= \frac{1}{K^2} \mathbb{E}_{\mathbf{S}} \left(\sum_{j=1}^K \|\nabla f(\mathbf{w}) - \nabla F_{k_j}(\mathbf{w})\|^2 \right) + \frac{1}{K^2} \mathbb{E}_{\mathbf{S}} \left(\sum_{\substack{j, j' \in [K] \\ j \neq j'}} \left\langle \nabla f(\mathbf{w}) - \nabla F_{k_j}(\mathbf{w}), \nabla f(\mathbf{w}) - \nabla F_{k_{j'}}(\mathbf{w}) \right\rangle \right), \end{aligned} \quad (\text{C.2})$$

where we note that \mathcal{V}_i is a random set of workers in group i and k_j is a random worker index assigned to the j -th bin, both of which capture the outcome of random grouping. Equality (a) is due to symmetry, so that the expectation involving any group i is equal to the expectation involving the first group $i = 1$. Equality (b) is from the fact that the first group \mathcal{V}_1 includes workers in bins $j = 1, 2, \dots, K$, hence $\sum_{k \in \mathcal{V}_1} \nabla F_k(\mathbf{w}) = \sum_{j=1}^K \nabla F_{k_j}(\mathbf{w})$. We analyze the two terms in (C.2) in the following.

For the first term in (C.2), we have

$$\frac{1}{K^2} \mathbb{E}_{\mathbf{S}} \left(\sum_{j=1}^K \|\nabla f(\mathbf{w}) - \nabla F_{k_j}(\mathbf{w})\|^2 \right) = \frac{1}{K^2} \sum_{j=1}^K \mathbb{E}_{\mathbf{S}} \left(\|\nabla f(\mathbf{w}) - \nabla F_{k_j}(\mathbf{w})\|^2 \right) = \frac{\epsilon_{\mathbf{w}}^2}{K}, \quad (\text{C.3})$$

where the last inequality is from (C.1).

For the second term in (C.2), we have

$$\begin{aligned} & \frac{1}{K^2} \mathbb{E}_{\mathbf{S}} \left(\sum_{\substack{j, j' \in [K] \\ j \neq j'}} \left\langle \nabla f(\mathbf{w}) - \nabla F_{k_j}(\mathbf{w}), \nabla f(\mathbf{w}) - \nabla F_{k_{j'}}(\mathbf{w}) \right\rangle \right) \\ &= \frac{1}{K^2} \sum_{\substack{j, j' \in [K] \\ j \neq j'}} \mathbb{E}_{\mathbf{S}} \left\langle \nabla f(\mathbf{w}) - \nabla F_{k_j}(\mathbf{w}), \nabla f(\mathbf{w}) - \nabla F_{k_{j'}}(\mathbf{w}) \right\rangle \end{aligned}$$

$$\begin{aligned}
&\stackrel{(a)}{=} \frac{1}{K^2} \sum_{\substack{j,j' \in [K] \\ j \neq j'}} \sum_{\substack{k,k' \in [n] \\ k \neq k'}} \frac{1}{n(n-1)} \left\langle \nabla f(\mathbf{w}) - \nabla F_k(\mathbf{w}), \nabla f(\mathbf{w}) - \nabla F_{k'}(\mathbf{w}) \right\rangle \\
&= \frac{1}{K^2} \sum_{\substack{j,j' \in [K] \\ j \neq j'}} \frac{1}{n(n-1)} \left[\left\| \sum_{k=1}^n \left(\nabla f(\mathbf{w}) - \nabla F_k(\mathbf{w}) \right) \right\|^2 - \sum_{k=1}^n \left\| \nabla f(\mathbf{w}) - \nabla F_k(\mathbf{w}) \right\|^2 \right] \\
&\stackrel{(b)}{=} -\frac{1}{K^2} \sum_{\substack{j,j' \in [K] \\ j \neq j'}} \frac{1}{n(n-1)} \sum_{k=1}^n \left\| \nabla f(\mathbf{w}) - \nabla F_k(\mathbf{w}) \right\|^2 \\
&= -\frac{K-1}{K} \cdot \frac{1}{n-1} \cdot \frac{1}{n} \sum_{k=1}^n \left\| \nabla f(\mathbf{w}) - \nabla F_k(\mathbf{w}) \right\|^2 \\
&\stackrel{(c)}{=} -\frac{K-1}{K(n-1)} \epsilon_{\mathbf{w}}^2,
\end{aligned} \tag{C.4}$$

where (a) is because the probability of assigning worker k to bin j and at the same time a different worker k' to a different bin j' is $\frac{1}{n(n-1)}$, (b) is because $\sum_{k=1}^n \left(\nabla f(\mathbf{w}) - \nabla F_k(\mathbf{w}) \right) = 0$ according to the definition in (2) and the fact that n_i is equal for all i , and (c) is from (C.1).

Substituting (C.3) and (C.4) into (C.2), we obtain

$$\begin{aligned}
&\mathbb{E}_{\mathcal{S}} \left[\sum_{i=1}^N \frac{n_i}{n} \left\| \nabla f(\mathbf{w}) - \frac{1}{K} \sum_{k \in \mathcal{V}_i} \nabla F_k(\mathbf{w}) \right\|^2 \right] \\
&= \frac{1}{K} \left[\epsilon_{\mathbf{w}}^2 - \frac{K-1}{n-1} \epsilon_{\mathbf{w}}^2 \right] \\
&= \frac{1}{K} \left(\frac{n-K}{n-1} \right) \epsilon_{\mathbf{w}}^2 = \left(\frac{N-1}{n-1} \right) \epsilon_{\mathbf{w}}^2 \stackrel{(a)}{\leq} \left(\frac{N-1}{n-1} \right) \tilde{\epsilon}^2,
\end{aligned} \tag{C.5}$$

where (a) is from (C.1). This proves the lemma.

C.2 Proof of Lemma 2

We have

$$\begin{aligned}
&\mathbb{E}_{\mathcal{S}} \left[\frac{1}{K} \sum_{k \in \mathcal{V}_i} \left\| \nabla F_k(\mathbf{w}) - \frac{1}{K} \sum_{k' \in \mathcal{V}_i} \nabla F_{k'}(\mathbf{w}) \right\|^2 \right] \\
&\stackrel{(a)}{=} \mathbb{E}_{\mathcal{S}} \left[\frac{1}{K} \sum_{k \in \mathcal{V}_1} \left\| \nabla F_k(\mathbf{w}) - \frac{1}{K} \sum_{k' \in \mathcal{V}_1} \nabla F_{k'}(\mathbf{w}) \right\|^2 \right] \\
&\stackrel{(b)}{=} \mathbb{E}_{\mathcal{S}} \left[\frac{1}{K} \sum_{j=1}^K \left\| \nabla F_{k_j}(\mathbf{w}) - \frac{1}{K} \sum_{j'=1}^K \nabla F_{k_{j'}}(\mathbf{w}) \right\|^2 \right] \\
&= \frac{1}{K} \sum_{j=1}^K \mathbb{E}_{\mathcal{S}} \left\| \nabla F_{k_j}(\mathbf{w}) - \nabla f(\mathbf{w}) + \nabla f(\mathbf{w}) - \frac{1}{K} \sum_{j'=1}^K \nabla F_{k_{j'}}(\mathbf{w}) \right\|^2 \\
&= \frac{1}{K} \sum_{j=1}^K \mathbb{E}_{\mathcal{S}} \left[\left\| \nabla F_{k_j}(\mathbf{w}) - \nabla f(\mathbf{w}) \right\|^2 \right] + \frac{1}{K} \sum_{j=1}^K \mathbb{E}_{\mathcal{S}} \left\| \nabla f(\mathbf{w}) - \frac{1}{K} \sum_{j'=1}^K \nabla F_{k_{j'}}(\mathbf{w}) \right\|^2 \\
&\quad + \frac{2}{K} \sum_{j=1}^K \mathbb{E}_{\mathcal{S}} \left\langle \nabla F_{k_j}(\mathbf{w}) - \nabla f(\mathbf{w}), \nabla f(\mathbf{w}) - \frac{1}{K} \sum_{j'=1}^K \nabla F_{k_{j'}}(\mathbf{w}) \right\rangle \\
&= \frac{1}{K} \sum_{j=1}^K \mathbb{E}_{\mathcal{S}} \left[\left\| \nabla F_{k_j}(\mathbf{w}) - \nabla f(\mathbf{w}) \right\|^2 \right] + \mathbb{E}_{\mathcal{S}} \left\| \nabla f(\mathbf{w}) - \frac{1}{K} \sum_{j'=1}^K \nabla F_{k_{j'}}(\mathbf{w}) \right\|^2
\end{aligned}$$

$$\begin{aligned}
& + 2\mathbb{E}_S \left\langle \frac{1}{K} \sum_{j=1}^K \nabla F_{k_j}(\mathbf{w}) - \nabla f(\mathbf{w}), \nabla f(\mathbf{w}) - \frac{1}{K} \sum_{j'=1}^K \nabla F_{k_{j'}}(\mathbf{w}) \right\rangle \\
& = \frac{1}{K} \sum_{j=1}^K \mathbb{E}_S \left[\left\| \nabla F_{k_j}(\mathbf{w}) - \nabla f(\mathbf{w}) \right\|^2 \right] - \mathbb{E}_S \left\| \nabla f(\mathbf{w}) - \frac{1}{K} \sum_{j'=1}^K \nabla F_{k_{j'}}(\mathbf{w}) \right\|^2 \\
& \stackrel{(c)}{=} \epsilon_{\mathbf{w}}^2 - \left(\frac{N-1}{n-1} \right) \epsilon_{\mathbf{w}}^2 \\
& \stackrel{(d)}{\leq} \left(1 - \frac{N-1}{n-1} \right) \tilde{\epsilon}^2,
\end{aligned} \tag{C.6}$$

where (a) is due to symmetry of groups so that the expectation over \mathcal{V}_i is equal to the expectation over \mathcal{V}_1 , (b) is because the workers in group $i = 1$ are in the first K bins, the first term in (c) is from (C.1), the second term in (c) is from the intermediate result in (C.5) of Lemma 1's proof and also the symmetry of groups, and (d) is from (C.1) and $1 - \frac{N-1}{n-1} \geq 0$.

C.3 Proof of Theorem 2

From Lemma 1 and Lemma 2, we have the expected global divergence and local divergence:

$$\mathbb{E}_S \left[\sum_{i=1}^N \frac{n_i}{n} \left\| \nabla f_i(\mathbf{w}) - \nabla f(\mathbf{w}) \right\|^2 \right] \leq \left(\frac{N-1}{n-1} \right) \tilde{\epsilon}^2, \forall \mathbf{w}, \tag{C.7}$$

$$\mathbb{E}_S \left[\frac{1}{n_i} \sum_{k \in \mathcal{V}_i} \left\| \nabla F_k(\mathbf{w}) - \nabla f_i(\mathbf{w}) \right\|^2 \right] \leq \left(1 - \frac{N-1}{n-1} \right) \tilde{\epsilon}^2, \forall \mathbf{w}, \forall i \tag{C.8}$$

where we recall that $n_i = K = \frac{n}{N}$ and $f_i(\mathbf{w}) = \frac{1}{n_i} \sum_{j \in \mathcal{V}_i} F_j(\mathbf{w})$ by definition.

The proof of Theorem 2 follows similar steps as that of Theorem 1, with the following differences: 1) we take an additional expectation over S (i.e., all possible groupings) and consider the expectations currently in Theorem 1, and their proofs as being conditioned on a given grouping realization S ; 2) we substitute (C.7) and (C.8) into (B.15) and (B.23) to replace the original upper bounds ϵ^2 and ϵ_i^2 with the expected divergence upper bounds $\left(\frac{N-1}{n-1} \right) \tilde{\epsilon}^2$ and $\left(1 - \frac{N-1}{n-1} \right) \tilde{\epsilon}^2$, respectively. Due to the similarity, we do not write the full proof of Theorem 2, and refer readers to the proofs of Theorem 1 while noting the above minor changes.

D Extension to the Multi-Level Case ($M \geq 2$)

The multi-level SGD algorithm is shown in Algorithm D.1.

Algorithm D.1: Multi-level H-SGD

Input: $\gamma, \bar{\mathbf{w}}^0, \{P_i : i \in \{1, 2, \dots, M\}\}$

Output: Global aggregated model $\bar{\mathbf{w}}^T$

for $t \leftarrow 0$ **to** $T - 1$ **do**

for Each worker $(k_1, \dots, k_M) \in \mathcal{V}$, *in parallel* **do**

Compute $\mathbf{g}_{k_1 \dots k_M}^t(\mathbf{w}_{k_1 \dots k_M}^t, \zeta_{k_1 \dots k_M}^t)$;

$\mathbf{w}_{k_1 \dots k_M}^{t+1} \leftarrow \mathbf{w}_{k_1 \dots k_M}^t - \gamma \mathbf{g}_{k_1 \dots k_M}^t(\mathbf{w}_{k_1 \dots k_M}^t, \zeta_{k_1 \dots k_M}^t)$;

for level $i \leftarrow 1$ **to** M **do**

if $P_i \mid t + 1$ **then**

i -th level aggregate: $\bar{\mathbf{w}}_{k_1 \dots k_i}^{t+1} \leftarrow \frac{1}{N_i \dots N_M} \sum_{k_i=1}^{N_i} \dots \sum_{k_M=1}^{N_M} \mathbf{w}_{k_1 \dots k_M}^{t+1}, \forall k_1, \dots, k_M$;

Distribute: $\mathbf{w}_{k_1 \dots k_M}^{t+1} \leftarrow \bar{\mathbf{w}}_{k_1 \dots k_i}^{t+1}, \forall k_1, \dots, k_M$;

break; // Lower level aggregation not needed when aggregated at higher level

First, we formally state an assumption for multi-level upward divergence and downward divergence as follows.

Assumption 3. For any server at level ℓ , the ℓ -th level upward and downward divergences are bounded by

$$\frac{1}{n_\ell} \sum_{k_1=1}^{N_1} \dots \sum_{k_\ell=1}^{N_\ell} \|\nabla f_{k_1 \dots k_\ell}(\mathbf{w}) - \nabla f(\mathbf{w})\|^2 \leq \hat{\epsilon}_\ell^2, \quad (\text{D.1})$$

$$\frac{n_\ell}{n} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \dots \sum_{k_M=1}^{N_M} \|\nabla f_{k_1 \dots k_\ell}(\mathbf{w}) - \nabla F_{k_1 \dots k_M}(\mathbf{w})\|^2 \leq \hat{\epsilon}_\ell'^2 \quad (\text{D.2})$$

respectively, $\forall \mathbf{w}, \forall k_1, \dots, k_\ell$ in (21), where $n_\ell = \prod_{j=1}^\ell N_j$.

Next, we provide the theorem for multi-level H-SGD under the fixed grouping case and its proof. We will extend the proof to the uniform random grouping case in Sections D.2 and D.3.

Theorem D.1. For any fixed worker grouping that satisfies Assumption 3, if the learning rate in Algorithm D.1 satisfies $\gamma < \frac{1}{2\sqrt{3}P_1L}$, then for any $T \geq 1$, we have

$$\begin{aligned} & \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left\| \nabla f(\bar{\mathbf{w}}^t) \right\|^2 \\ & \leq \frac{2}{\gamma T} \left(f^0 - f^* \right) + \frac{\gamma L \sigma^2}{n} \\ & \quad + \frac{1}{M-1} \sum_{l=1}^{M-1} \left[\frac{8\gamma^2 L^2 P_1 \left(1 - \frac{1}{\prod_{j=1}^l N_j} \right) \frac{1}{\prod_{j=l+1}^M N_j} \sigma^2}{1 - 12\gamma^2 L^2 P_1^2} + \frac{12\gamma^2 L^2 P_1^2 \hat{\epsilon}_l^2}{1 - 12\gamma^2 L^2 P_1^2} \right. \\ & \quad \left. + \left(1 + \frac{8\gamma^2 P_1^2 L^2}{1 - 12\gamma^2 L^2 P_1^2} \right) \left(\frac{4\gamma^2 L^2 P_l \left(1 - \frac{1}{\prod_{j=l+1}^M N_j} \right) \sigma^2}{1 - 12\gamma^2 L^2 P_l^2} + \frac{12\gamma^2 L^2 P_l^2 \hat{\epsilon}_l'^2}{1 - 12\gamma^2 L^2 P_l^2} \right) \right]. \end{aligned}$$

D.1 Proof of Theorem D.1

Similar to (B.3)–(B.7), when $\gamma \leq \frac{1}{L}$, that is, $\frac{\gamma}{2} - \frac{\gamma^2 L}{2} \geq 0$, we can obtain

$$\mathbb{E} f(\bar{\mathbf{w}}^{t+1}) \leq \mathbb{E} f(\bar{\mathbf{w}}^t) + \frac{\gamma^2 L \sigma^2}{2n} - \frac{\gamma}{2} \mathbb{E} \|\nabla f(\bar{\mathbf{w}}^t)\|^2 + \frac{\gamma}{2} \mathbb{E} \left\| \nabla f(\bar{\mathbf{w}}^t) - \frac{1}{n} \sum_{i=1}^n \nabla F_i(\mathbf{w}_i^t) \right\|^2.$$

Next, we will upper bound the term $\mathbb{E} \left\| \nabla f(\bar{\mathbf{w}}^t) - \frac{1}{n} \sum_{i=1}^n \nabla F_i(\mathbf{w}_i^t) \right\|^2$ for the multi-level case. In order to bound this term, we divide it into upward part as in (D.3) and downward part in (D.4). The upward part means the difference between the global

gradient and i -th ($1 \leq i \leq M-1$) level averaged gradient. The downward means the difference between the gradients of workers and i -th ($1 \leq i \leq M-1$) level averaged gradient, which is similar to (B.8). Then each part is divided into $M-1$ terms corresponding to $M-1$ levels between the global server and workers. The details are as follows.

$$\begin{aligned}
& \mathbb{E} \left\| \nabla f(\bar{\mathbf{w}}^t) - \frac{1}{n} \sum_{i=1}^n \nabla F_i(\mathbf{w}_j^t) \right\|^2 \\
&= \mathbb{E} \left\| \nabla f(\bar{\mathbf{w}}^t) - \frac{1}{M-1} \left[\frac{1}{N_1} \sum_{k_1=1}^{N_1} \nabla f_{k_1}(\bar{\mathbf{w}}_{k_1}^t) - \frac{1}{N_1} \sum_{k_1=1}^{N_1} \nabla f_{k_1}(\bar{\mathbf{w}}_{k_1}^t) + \frac{1}{N_1 N_2} \sum_{k_1=1}^{N_1} \sum_{k_2=1}^{N_2} \nabla f_{k_1 k_2}(\bar{\mathbf{w}}_{k_1 k_2}^t) \right. \right. \\
&\quad \left. \left. - \frac{1}{N_1 N_2} \sum_{k_1=1}^{N_1} \sum_{k_2=1}^{N_2} \nabla f_{k_1 k_2}(\bar{\mathbf{w}}_{k_1 k_2}^t) \right. \right. \\
&\quad \left. \left. + \dots \right. \right. \\
&\quad \left. \left. + \frac{1}{\prod_{j=1}^{M-1} N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_{M-1}=1}^{N_{M-1}} \nabla f_{k_1 \dots k_{M-1}}(\bar{\mathbf{w}}_{k_1 \dots k_{M-1}}^t) - \frac{1}{\prod_{j=1}^{M-1} N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_{M-1}=1}^{N_{M-1}} \nabla f_{k_1 \dots k_{M-1}}(\bar{\mathbf{w}}_{k_1 \dots k_{M-1}}^t) \right] \right. \\
&\quad \left. \left. - \frac{1}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_M=1}^{N_M} \nabla F_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^t) \right\|^2 \right. \\
&\stackrel{(a)}{\leq} 2\mathbb{E} \left\| \frac{1}{M-1} \left[\nabla f(\bar{\mathbf{w}}^t) - \frac{1}{N_1} \sum_{k_1=1}^{N_1} \nabla f_{k_1}(\bar{\mathbf{w}}_{k_1}^t) + \nabla f(\bar{\mathbf{w}}^t) - \frac{1}{N_1 N_2} \sum_{k_1=1}^{N_1} \sum_{k_2=1}^{N_2} \nabla f_{k_1 k_2}(\bar{\mathbf{w}}_{k_1 k_2}^t) + \dots \right. \right. \\
&\quad \left. \left. + \nabla f(\bar{\mathbf{w}}^t) - \frac{1}{\prod_{j=1}^{M-1} N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_{M-1}=1}^{N_{M-1}} \nabla f_{k_1 \dots k_{M-1}}(\bar{\mathbf{w}}_{k_1 \dots k_{M-1}}^t) \right] \right\|^2 \tag{D.3}
\end{aligned}$$

$$\begin{aligned}
&+ 2\mathbb{E} \left\| \frac{1}{M-1} \left[\frac{1}{N_1} \sum_{k_1=1}^{N_1} \nabla f_{k_1}(\bar{\mathbf{w}}_{k_1}^t) - \frac{1}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_M=1}^{N_M} \nabla F_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^t) \right. \right. \\
&+ \frac{1}{N_1 N_2} \sum_{k_1=1}^{N_1} \sum_{k_2=1}^{N_2} \nabla f_{k_1 k_2}(\bar{\mathbf{w}}_{k_1 k_2}^t) - \frac{1}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_M=1}^{N_M} \nabla F_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^t) + \dots \\
&\left. \left. + \frac{1}{\prod_{j=1}^{M-1} N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_{M-1}=1}^{N_{M-1}} \nabla f_{k_1 \dots k_{M-1}}(\bar{\mathbf{w}}_{k_1 \dots k_{M-1}}^t) - \frac{1}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_M=1}^{N_M} \nabla F_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^t) \right] \right\|^2 \tag{D.4}
\end{aligned}$$

$$\begin{aligned}
&\stackrel{(b)}{\leq} \frac{1}{M-1} \left[2\mathbb{E} \left\| \nabla f(\bar{\mathbf{w}}^t) - \frac{1}{N_1} \sum_{k_1=1}^{N_1} \nabla f_{k_1}(\bar{\mathbf{w}}_{k_1}^t) \right\|^2 \right. \\
&+ 2\mathbb{E} \left\| \frac{1}{N_1} \sum_{k_1=1}^{N_1} \nabla f_{k_1}(\bar{\mathbf{w}}_{k_1}^t) - \frac{1}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_M=1}^{N_M} \nabla F_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^t) \right\|^2 \Big] \\
&+ \frac{1}{M-1} \left[2\mathbb{E} \left\| \nabla f(\bar{\mathbf{w}}^t) - \frac{1}{N_1 N_2} \sum_{k_1=1}^{N_1} \sum_{k_2=1}^{N_2} \nabla f_{k_1 k_2}(\bar{\mathbf{w}}_{k_1 k_2}^t) \right\|^2 \right. \\
&+ 2\mathbb{E} \left\| \frac{1}{N_1 N_2} \sum_{k_1=1}^{N_1} \sum_{k_2=1}^{N_2} \nabla f_{k_1 k_2}(\bar{\mathbf{w}}_{k_1 k_2}^t) - \frac{1}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_M=1}^{N_M} \nabla F_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^t) \right\|^2 \Big] \\
&+ \dots \\
&+ \frac{1}{M-1} \left[2\mathbb{E} \left\| \nabla f(\bar{\mathbf{w}}^t) - \frac{1}{\prod_{j=1}^{M-1} N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_{M-1}=1}^{N_{M-1}} \nabla f_{k_1 \dots k_{M-1}}(\bar{\mathbf{w}}_{k_1 \dots k_{M-1}}^t) \right\|^2 \right. \\
&+ 2\mathbb{E} \left\| \frac{1}{\prod_{j=1}^{M-1} N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_{M-1}=1}^{N_{M-1}} \nabla f_{k_1 \dots k_{M-1}}(\bar{\mathbf{w}}_{k_1 \dots k_{M-1}}^t) - \frac{1}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_M=1}^{N_M} \nabla F_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^t) \right\|^2 \Big], \tag{D.5}
\end{aligned}$$

where (a) and (b) are due to (B.1).

For any level ℓ , $1 \leq \ell \leq M-1$, we have

$$\begin{aligned} & \mathbb{E} \left\| \nabla f(\bar{\mathbf{w}}^t) - \frac{1}{\prod_{j=1}^{\ell} N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_{\ell}=1}^{N_{\ell}} \nabla f_{k_1 \dots k_{M-1}}(\bar{\mathbf{w}}_{k_1 \dots k_{M-1}}^t) \right\|^2 \\ & \stackrel{(a)}{\leq} L^2 \frac{1}{\prod_{j=1}^{\ell} N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_{\ell}=1}^{N_{\ell}} \mathbb{E} \|\bar{\mathbf{w}}^t - \bar{\mathbf{w}}_{k_1 \dots k_{\ell}}^t\|^2 \end{aligned} \quad (\text{D.6})$$

and

$$\begin{aligned} & \mathbb{E} \left\| \frac{1}{\prod_{j=1}^{\ell} N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_{\ell}=1}^{N_{\ell}} \nabla f_{k_1 \dots k_{M-1}}(\bar{\mathbf{w}}_{k_1 \dots k_{M-1}}^t) - \frac{1}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_M=1}^{N_M} \nabla F_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^t) \right\|^2 \\ & \stackrel{(b)}{\leq} L^2 \frac{1}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_{\ell}=1}^{N_{\ell}} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k_M=1}^{N_M} \mathbb{E} \|\bar{\mathbf{w}}_{k_1 \dots k_{\ell}}^t - \mathbf{w}_{k_1 \dots k_{\ell} \dots k_M}^t\|^2, \end{aligned} \quad (\text{D.7})$$

where (a) and (b) are due to (B.1) and Lipschitz gradient assumption. Similar to Sections B.2 and B.3, we will bound averaged (D.6) and (D.7) over time, denoted as the upward parameter MSE and downward parameter MSE in multi-level case, respectively.

Bounding Upward Parameter MSE In this section, we will provide a general way to bound the upward parameter MSE for level ℓ .

$$\begin{aligned} & \frac{1}{\prod_{j=1}^{\ell} N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_{\ell}=1}^{N_{\ell}} \mathbb{E} \|\bar{\mathbf{w}}^t - \bar{\mathbf{w}}_{k_1 \dots k_{\ell}}^t\|^2 \\ & = \frac{\gamma^2}{\prod_{j=1}^{\ell} N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_{\ell}=1}^{N_{\ell}} \mathbb{E} \left\| \sum_{\tau=aP_1}^{aP_1+P_1-1} \left(\frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k_M=1}^{N_M} \mathbf{g}_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^{\tau}) \right. \right. \\ & \quad \left. \left. - \frac{1}{\prod_{j=1}^M N_j} \sum_{k'_1=1}^{N_1} \cdots \sum_{k'_M=1}^{N_M} \mathbf{g}_{k'_1 \dots k'_M}(\mathbf{w}_{k'_1 \dots k'_M}^{\tau}) \right) \right\|^2 \\ & \stackrel{(a)}{\leq} \frac{2\gamma^2}{\prod_{j=1}^{\ell} N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_{\ell}=1}^{N_{\ell}} \mathbb{E} \left\| \sum_{\tau=aP_1}^{aP_1+P_1-1} \left[\left(\frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k_M=1}^{N_M} \mathbf{g}_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^{\tau}) - \nabla f_{k_1 \dots k_{\ell}}(\bar{\mathbf{w}}_{k_1 \dots k_{\ell}}^{\tau}) \right) \right. \right. \\ & \quad \left. \left. - \left(\frac{1}{\prod_{j=1}^M N_j} \sum_{k'_1=1}^{N_1} \cdots \sum_{k'_M=1}^{N_M} \mathbf{g}_{k'_1 \dots k'_M}(\mathbf{w}_{k'_1 \dots k'_M}^{\tau}) - \frac{1}{\prod_{j=1}^M N_j} \sum_{k'_1=1}^{N_1} \cdots \sum_{k'_M=1}^{N_M} \nabla f_{k'_1 \dots k'_M}(\bar{\mathbf{w}}_{k'_1 \dots k'_M}^{\tau}) \right) \right] \right\|^2 \\ & \quad + \frac{2\gamma^2}{\prod_{j=1}^{\ell} N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_{\ell}=1}^{N_{\ell}} \mathbb{E} \left\| \sum_{\tau=aP_1}^{aP_1+P_1-1} \left(\frac{1}{\prod_{j=1}^M N_j} \sum_{k'_1=1}^{N_1} \cdots \sum_{k'_M=1}^{N_M} \nabla f_{k'_1 \dots k'_M}(\bar{\mathbf{w}}_{k'_1 \dots k'_M}^{\tau}) - \nabla f_{k_1 \dots k_{\ell}}(\bar{\mathbf{w}}_{k_1 \dots k_{\ell}}^{\tau}) \right) \right\|^2, \end{aligned} \quad (\text{D.8})$$

where (a) is due to (B.1). Similar to (B.13), for the first term of (D.8), we have

$$\begin{aligned} & \frac{2\gamma^2}{\prod_{j=1}^{\ell} N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_{\ell}=1}^{N_{\ell}} \mathbb{E} \left\| \sum_{\tau=aP_1}^{aP_1+P_1-1} \left[\left(\frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k_M=1}^{N_M} \mathbf{g}_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^{\tau}) - \nabla f_{k_1 \dots k_{\ell}}(\bar{\mathbf{w}}_{k_1 \dots k_{\ell}}^{\tau}) \right) \right. \right. \\ & \quad \left. \left. - \left(\frac{1}{\prod_{j=1}^M N_j} \sum_{k'_1=1}^{N_1} \cdots \sum_{k'_M=1}^{N_M} \mathbf{g}_{k'_1 \dots k'_M}(\mathbf{w}_{k'_1 \dots k'_M}^{\tau}) - \frac{1}{\prod_{j=1}^M N_j} \sum_{k'_1=1}^{N_1} \cdots \sum_{k'_M=1}^{N_M} \nabla f_{k'_1 \dots k'_M}(\bar{\mathbf{w}}_{k'_1 \dots k'_M}^{\tau}) \right) \right] \right\|^2 \\ & \stackrel{(a)}{\leq} \frac{4\gamma^2}{\prod_{j=1}^{\ell} N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_{\ell}=1}^{N_{\ell}} \mathbb{E} \left\| \sum_{\tau=aP_1}^{aP_1+P_1-1} \left[\left(\frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k_M=1}^{N_M} \mathbf{g}_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^{\tau}) \right. \right. \right. \\ & \quad \left. \left. - \frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k_M=1}^{N_M} \nabla F_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^{\tau}) \right) \right] \right\|^2 \end{aligned}$$

$$\begin{aligned}
& - \left(\frac{1}{\prod_{j=1}^M N_j} \sum_{k'_1=1}^{N_1} \cdots \sum_{k'_M=1}^{N_M} \mathbf{g}_{k'_1 \dots k'_M}(\mathbf{w}_{k'_1 \dots k'_M}^\tau) - \frac{1}{\prod_{j=1}^M N_j} \sum_{k'_1=1}^{N_1} \cdots \sum_{k'_M=1}^{N_M} \nabla F_{k'_1 \dots k'_M}(\mathbf{w}_{k'_1 \dots k'_M}^\tau) \right) \Big\| \Big\|^2 \\
& + \frac{4\gamma^2}{\prod_{j=1}^\ell N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_\ell=1}^{N_\ell} \mathbb{E} \left\| \sum_{\tau=aP_1}^{aP_1+P_1-1} \left(\frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k_M=1}^{N_M} \nabla F_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^\tau) - \nabla f_{k_1 \dots k_\ell}(\bar{\mathbf{w}}_{k_1 \dots k_\ell}^\tau) \right) \right\|^2,
\end{aligned} \tag{D.9}$$

where (a) is due to (B.1). For the first term of (D.9), we have

$$\begin{aligned}
& \frac{4\gamma^2}{\prod_{j=1}^\ell N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_\ell=1}^{N_\ell} \mathbb{E} \left\| \sum_{\tau=aP_1}^{aP_1+P_1-1} \left[\left(\frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k_M=1}^{N_M} \mathbf{g}_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^\tau) \right. \right. \right. \\
& \left. \left. \left. - \frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k_M=1}^{N_M} \nabla F_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^\tau) \right) \right] \right\|^2 \\
& - \left(\frac{1}{\prod_{j=1}^M N_j} \sum_{k'_1=1}^{N_1} \cdots \sum_{k'_M=1}^{N_M} \mathbf{g}_{k'_1 \dots k'_M}(\mathbf{w}_{k'_1 \dots k'_M}^\tau) - \frac{1}{\prod_{j=1}^M N_j} \sum_{k'_1=1}^{N_1} \cdots \sum_{k'_M=1}^{N_M} \nabla F_{k'_1 \dots k'_M}(\mathbf{w}_{k'_1 \dots k'_M}^\tau) \right) \Big\| \Big\|^2 \\
& = \frac{4\gamma^2}{\prod_{j=1}^\ell N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_\ell=1}^{N_\ell} \mathbb{E} \left\| \sum_{\tau=aP_1}^{aP_1+P_1-1} \left(\frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k_M=1}^{N_M} \mathbf{g}_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^\tau) \right. \right. \\
& \left. \left. - \frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k_M=1}^{N_M} \nabla F_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^\tau) \right) \right\|^2 \\
& - 4\gamma^2 \mathbb{E} \left\| \sum_{\tau=aP_1}^{aP_1+P_1-1} \left(\frac{1}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_M=1}^{N_M} \mathbf{g}_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^\tau) - \frac{1}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_M=1}^{N_M} \nabla F_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^\tau) \right) \right\|^2 \\
& = 4\gamma^2 \left(1 - \frac{1}{\prod_{j=1}^\ell N_j} \right) \mathbb{E} \left\| \sum_{\tau=aP_1}^{aP_1+P_1-1} \left(\frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k_M=1}^{N_M} \mathbf{g}_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^\tau) \right. \right. \\
& \left. \left. - \frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k_M=1}^{N_M} \nabla F_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^\tau) \right) \right\|^2 \\
& \stackrel{(a)}{\leq} 4P_1\gamma^2 \left(1 - \frac{1}{\prod_{j=1}^\ell N_j} \right) \frac{1}{\prod_{j=\ell+1}^M N_j} \sigma^2,
\end{aligned} \tag{D.10}$$

where (a) is due to (B.1) and bounded variance in Assumption 1. For the second term of (D.9), we have

$$\begin{aligned}
& \frac{4\gamma^2}{\prod_{j=1}^\ell N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_\ell=1}^{N_\ell} \mathbb{E} \left\| \sum_{\tau=aP_1}^{aP_1+P_1-1} \left(\frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k_M=1}^{N_M} \nabla F_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^\tau) - \nabla f_{k_1 \dots k_\ell}(\bar{\mathbf{w}}_{k_1 \dots k_\ell}^\tau) \right) \right\|^2 \\
& \leq \frac{4\gamma^2 L^2 P_1}{\prod_{j=1}^M N_j} \sum_{\tau=aP_1}^{aP_1+P_1-1} \sum_{k_1=1}^{N_1} \cdots \sum_{k_M=1}^{N_M} \mathbb{E} \|\mathbf{w}_{k_1 \dots k_M}^\tau - \bar{\mathbf{w}}_{k_1 \dots k_\ell}^\tau\|^2,
\end{aligned} \tag{D.11}$$

Similar to (B.15), for the second term of (D.8), we have

$$\begin{aligned}
& \frac{2\gamma^2}{\prod_{j=1}^\ell N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_\ell=1}^{N_\ell} \mathbb{E} \left\| \sum_{\tau=aP_1}^{aP_1+P_1-1} \left(\frac{1}{\prod_{j=1}^M N_j} \sum_{k'_1=1}^{N_1} \cdots \sum_{k'_M=1}^{N_M} \nabla f_{k'_1 \dots k'_M}(\bar{\mathbf{w}}_{k'_1 \dots k'_M}^\tau) - \nabla f_{k_1 \dots k_\ell}(\bar{\mathbf{w}}_{k_1 \dots k_\ell}^\tau) \right) \right\|^2 \\
& = \frac{\gamma^2}{\prod_{j=1}^\ell N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_\ell=1}^{N_\ell} \mathbb{E} \left\| \sum_{\tau=aP_1}^{aP_1+P_1-1} \left(\frac{1}{\prod_{j=1}^\ell N_j} \sum_{k'_1=1}^{N_1} \cdots \sum_{k'_\ell=1}^{N_\ell} \nabla f_{k'_1 \dots k'_\ell}(\bar{\mathbf{w}}_{k'_1 \dots k'_\ell}^\tau) - \nabla f(\bar{\mathbf{w}}^\tau) + \nabla f(\bar{\mathbf{w}}^\tau) \right. \right. \\
& \left. \left. - \nabla f(\bar{\mathbf{w}}_{k_1 \dots k_\ell}^\tau) + \nabla f(\bar{\mathbf{w}}_{k_1 \dots k_\ell}^\tau) - \nabla f_{k_1 \dots k_\ell}(\bar{\mathbf{w}}_{k_1 \dots k_\ell}^\tau) \right) \right\|^2
\end{aligned}$$

$$\stackrel{(a)}{\leq} 6P_1^2\gamma^2\hat{\epsilon}_\ell^2 + \frac{12P_1\gamma^2L^2}{\prod_{j=1}^\ell N_j} \sum_{\tau=aP_1}^{aP_1+P_1-1} \sum_{k_1=1}^{N_1} \dots \sum_{k_\ell=1}^{N_\ell} \mathbb{E} \|\bar{\mathbf{w}}^\tau - \bar{\mathbf{w}}_{k_1\dots k_\ell}^\tau\|^2, \quad (\text{D.12})$$

where (a) is due to bounded divergence in Assumption 1 and (B.1).

Similar to (B.17)–(B.19), the upper bound for upward parameter MSE can be derived as follows.

$$\begin{aligned} & \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{\prod_{j=1}^\ell N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_\ell=1}^{N_\ell} \mathbb{E} \|\bar{\mathbf{w}}^t - \bar{\mathbf{w}}_{k_1\dots k_\ell}^t\|^2 \\ & \leq \frac{4P_1\gamma^2 \left(1 - \frac{1}{\prod_{j=1}^\ell N_j}\right) \frac{1}{\prod_{j=\ell+1}^M N_j} \sigma^2}{1 - 12\gamma^2 L^2 P_1^2} + \frac{6P_1^2\gamma^2\hat{\epsilon}_\ell^2}{1 - 12\gamma^2 L^2 P_1^2} \\ & \quad + \frac{4\gamma^2 L^2 P_1^2}{1 - 12\gamma^2 L^2 P_1^2} \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_M=1}^{N_M} \mathbb{E} \|\mathbf{w}_{k_1\dots k_M}^\tau - \bar{\mathbf{w}}_{k_1\dots k_\ell}^\tau\|^2. \end{aligned} \quad (\text{D.13})$$

Bounding Downward Parameter MSE In this section, we will upper bound the downward parameter MSE for level ℓ .

$$\begin{aligned} & \frac{1}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_M=1}^{N_M} \mathbb{E} \|\bar{\mathbf{w}}_{k_1\dots k_\ell}^t - \mathbf{w}_{k_1\dots k_M}^t\|^2 \\ & = \frac{\gamma^2}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_M=1}^{N_M} \mathbb{E} \left\| \sum_{\tau=aP_\ell}^{aP_\ell+P_\ell-1} \left(\mathbf{g}_{k_1\dots k_M}(\mathbf{w}_{k_1\dots k_M}^\tau) \right. \right. \\ & \quad \left. \left. - \frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k'_{\ell+1}=1}^{N_{\ell+1}} \dots \sum_{k'_M=1}^{N_M} \mathbf{g}_{k_1\dots k'_{\ell+1}\dots k'_M}(\mathbf{w}_{k_1\dots k'_{\ell+1}\dots k'_M}^\tau) \right) \right\|^2 \\ & \stackrel{(a)}{\leq} \frac{2\gamma^2}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_M=1}^{N_M} \mathbb{E} \left\| \sum_{\tau=aP_\ell}^{aP_\ell+P_\ell-1} \left(\mathbf{g}_{k_1\dots k_M}(\mathbf{w}_{k_1\dots k_M}^\tau) - \nabla F_{k_1\dots k_M}(\mathbf{w}_{k_1\dots k_M}^\tau) \right) \right. \\ & \quad \left. + \frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k'_{\ell+1}=1}^{N_{\ell+1}} \dots \sum_{k'_M=1}^{N_M} \nabla F_{k_1\dots k'_{\ell+1}\dots k'_M}(\mathbf{w}_{k_1\dots k'_{\ell+1}\dots k'_M}^\tau) \right. \\ & \quad \left. - \frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k'_{\ell+1}=1}^{N_{\ell+1}} \dots \sum_{k'_M=1}^{N_M} \mathbf{g}_{k_1\dots k'_{\ell+1}\dots k'_M}(\mathbf{w}_{k_1\dots k'_{\ell+1}\dots k'_M}^\tau) \right\|^2 \\ & \quad + \frac{2\gamma^2}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_M=1}^{N_M} \mathbb{E} \left\| \sum_{\tau=aP_\ell}^{aP_\ell+P_\ell-1} \left(\nabla F_{k_1\dots k_M}(\mathbf{w}_{k_1\dots k_M}^\tau) \right. \right. \\ & \quad \left. \left. - \frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k'_{\ell+1}=1}^{N_{\ell+1}} \dots \sum_{k'_M=1}^{N_M} \nabla F_{k_1\dots k'_{\ell+1}\dots k'_M}(\mathbf{w}_{k_1\dots k'_{\ell+1}\dots k'_M}^\tau) \right) \right\|^2 \end{aligned} \quad (\text{D.14})$$

where (a) is due to (B.1). Similar to (B.22), for the first term of (D.14), we have

$$\begin{aligned} & \frac{2\gamma^2}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \dots \sum_{k_M=1}^{N_M} \mathbb{E} \left\| \sum_{\tau=aP_\ell}^{aP_\ell+P_\ell-1} \left(\mathbf{g}_{k_1\dots k_M}(\mathbf{w}_{k_1\dots k_M}^\tau) - \nabla F_{k_1\dots k_M}(\mathbf{w}_{k_1\dots k_M}^\tau) \right) \right. \\ & \quad \left. + \frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k'_{\ell+1}=1}^{N_{\ell+1}} \dots \sum_{k'_M=1}^{N_M} \nabla F_{k_1\dots k'_{\ell+1}\dots k'_M}(\mathbf{w}_{k_1\dots k'_{\ell+1}\dots k'_M}^\tau) \right. \\ & \quad \left. - \frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k'_{\ell+1}=1}^{N_{\ell+1}} \dots \sum_{k'_M=1}^{N_M} \mathbf{g}_{k_1\dots k'_{\ell+1}\dots k'_M}(\mathbf{w}_{k_1\dots k'_{\ell+1}\dots k'_M}^\tau) \right\|^2 \end{aligned}$$

$$\begin{aligned}
&= \frac{2\gamma^2}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_M=1}^{N_M} \mathbb{E} \left\| \sum_{\tau=aP_\ell}^{aP_\ell+P_\ell-1} \left(\mathbf{g}_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^\tau) - \nabla F_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^\tau) \right) \right\|^2 \\
&\quad - \frac{\gamma^2}{\prod_{j=1}^\ell N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_\ell=1}^{N_\ell} \mathbb{E} \left\| \sum_{\tau=aP_\ell}^{aP_\ell+P_\ell-1} \left(\frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k_M=1}^{N_M} \nabla F_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^\tau) \right. \right. \\
&\quad \left. \left. - \frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k_M=1}^{N_M} \mathbf{g}_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^\tau) \right) \right\|^2 \\
&\stackrel{(a)}{\leq} 2\gamma^2 P_\ell \left(1 - \frac{1}{\prod_{j=\ell+1}^M N_j} \right) \sigma^2
\end{aligned} \tag{D.15}$$

where (a) is due to bounded variance in Assumption 2 and (B.1). Similar to (B.23), for the second term of (D.14), we have

$$\begin{aligned}
&\frac{2\gamma^2}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_M=1}^{N_M} \mathbb{E} \left\| \sum_{\tau=aP_\ell}^{aP_\ell+P_\ell-1} \left(\nabla F_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^\tau) \right. \right. \\
&\quad \left. \left. - \frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k'_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k'_M=1}^{N_M} \nabla F_{k_1 \dots k'_{\ell+1} \dots k'_M}(\mathbf{w}_{k_1 \dots k'_{\ell+1} \dots k'_M}^\tau) \right) \right\|^2 \\
&= \frac{2\gamma^2}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_M=1}^{N_M} \mathbb{E} \left\| \sum_{\tau=aP_\ell}^{aP_\ell+P_\ell-1} \left(\nabla F_{k_1 \dots k_M}(\mathbf{w}_{k_1 \dots k_M}^\tau) - \nabla F_{k_1 \dots k_M}(\bar{\mathbf{w}}_{k_1 \dots k_M}^\tau) \right. \right. \\
&\quad \left. \left. + \nabla F_{k_1 \dots k_M}(\bar{\mathbf{w}}_{k_1 \dots k_M}^\tau) - \nabla f_{k_1 \dots k_\ell}(\bar{\mathbf{w}}_{k_1 \dots k_\ell}^\tau) + \nabla f_{k_1 \dots k_\ell}(\bar{\mathbf{w}}_{k_1 \dots k_\ell}^\tau) \right. \right. \\
&\quad \left. \left. - \frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k'_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k'_M=1}^{N_M} \nabla F_{k_1 \dots k'_{\ell+1} \dots k'_M}(\mathbf{w}_{k_1 \dots k'_{\ell+1} \dots k'_M}^\tau) \right) \right\|^2 \\
&\stackrel{(a)}{\leq} 6\gamma^2 P_\ell^2 \hat{\epsilon}_\ell'^2 + \frac{12\gamma^2 L^2 P_\ell}{\prod_{j=1}^M N_j} \sum_{\tau=aP_\ell}^{aP_\ell+P_\ell-1} \sum_{k_1=1}^{N_1} \cdots \sum_{k_M=1}^{N_M} \mathbb{E} \|\mathbf{w}_{k_1 \dots k_M}^\tau - \bar{\mathbf{w}}_{k_1 \dots k_M}^\tau\|^2
\end{aligned} \tag{D.16}$$

where (a) is due to bounded divergence in Assumption 1 and (B.1). Combining the above two inequalities, applying the same technique in (B.24), we can obtain

$$\begin{aligned}
&\frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{\prod_{j=1}^M N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_M=1}^{N_M} \mathbb{E} \|\bar{\mathbf{w}}_{k_1 \dots k_M}^t - \mathbf{w}_{k_1 \dots k_M}^t\|^2 \\
&\leq \frac{2\gamma^2 P_\ell \left(1 - \frac{1}{\prod_{j=\ell+1}^M N_j} \right) \sigma^2}{1 - 12\gamma^2 L^2 P_\ell^2} + \frac{6\gamma^2 P_\ell^2 \hat{\epsilon}_\ell'^2}{1 - 12\gamma^2 L^2 P_\ell^2}.
\end{aligned} \tag{D.17}$$

Obtaining the Final Result. Substituting (D.13) and (D.17) into (D.5) and combining it with (B.7), since $\gamma < \frac{1}{\sqrt{12}LP_1}$, we have

$$\begin{aligned}
&\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left\| \nabla f(\bar{\mathbf{w}}^t) \right\|^2 \\
&\leq \frac{2}{\gamma T} \left(f^0 - f^* \right) + \frac{\gamma L \sigma^2}{n} \\
&\quad + \frac{1}{M-1} \sum_{l=1}^{M-1} \left[\frac{8\gamma^2 L^2 P_1 \left(1 - \frac{1}{\prod_{j=1}^\ell N_j} \right) \frac{1}{\prod_{j=\ell+1}^M N_j} \sigma^2}{1 - 12\gamma^2 L^2 P_1^2} + \frac{12\gamma^2 L^2 P_1^2 \hat{\epsilon}_\ell'^2}{1 - 12\gamma^2 L^2 P_1^2} \right. \\
&\quad \left. + \left(1 + \frac{8\gamma^2 P_1^2 L^2}{1 - 12\gamma^2 L^2 P_1^2} \right) \left(\frac{4\gamma^2 L^2 P_\ell \left(1 - \frac{1}{\prod_{j=\ell+1}^M N_j} \right) \sigma^2}{1 - 12\gamma^2 L^2 P_\ell^2} + \frac{12\gamma^2 L^2 P_\ell^2 \hat{\epsilon}_\ell'^2}{1 - 12\gamma^2 L^2 P_\ell^2} \right) \right].
\end{aligned} \tag{D.18}$$

D.2 Proof of Lemma 3

Note that by directly substituting K as $K = \prod_{j=\ell+1}^M N_j$ in the proofs Lemma 1 and Lemma 2, we extend Lemma 1 and Lemma 2 to the ℓ -th level case in Lemma 3 as follows:

$$\begin{aligned} & \mathbb{E}_S \left[\frac{1}{\prod_{j=1}^{\ell} N_j} \sum_{k_1=1}^{N_1} \cdots \sum_{k_{\ell}=1}^{N_{\ell}} \left\| \nabla f(\mathbf{w}) - \frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k_M=1}^{N_M} \nabla F_{k_1 \dots k_M}(\mathbf{w}) \right\|^2 \right] \\ & \leq \left(\frac{\prod_{j=1}^{\ell} N_j - 1}{n - 1} \right) \tilde{\epsilon} \end{aligned} \quad (\text{D.19})$$

$$\begin{aligned} & \mathbb{E}_S \left[\frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k_M=1}^{N_M} \left\| \frac{1}{\prod_{j=\ell+1}^M N_j} \sum_{k'_{\ell+1}=1}^{N_{\ell+1}} \cdots \sum_{k'_M=1}^{N_M} \nabla F_{k_1 \dots k'_{\ell+1} \dots k'_M}(\mathbf{w}) - \nabla F_{k_1 \dots k_M}(\mathbf{w}) \right\|^2 \right] \\ & \leq \left(1 - \frac{\prod_{j=1}^{\ell} N_j - 1}{n - 1} \right) \tilde{\epsilon}^2. \end{aligned} \quad (\text{D.20})$$

D.3 Proof of Theorem 3

From (D.18) to Theorem 3, we first take expectation over S for both sides. Then, similar to Section C.3, substitute $\frac{8\gamma^2 P_1^2 L^2}{1-12\gamma^2 L^2 P_1^2} \leq \frac{2}{3}$ and $\frac{1}{1-12\gamma^2 L^2 P_1^2} \leq 2$ to (D.18) and set $C = \frac{40}{3}$. Then applying (D.19) and (D.20) instead of $\hat{\epsilon}_{\ell}^2$ and $\hat{\epsilon}'_{\ell}^2$ in (D.18), respectively, we obtain the final result.

E Additional Details and Results of Experiments

In this section, we will provide additional details of our experiments and more results are shown for both two-level and multi-level (with $M = 3$) cases.

Environment and Parameters. All our experiments are implemented in PyTorch and run on a server with four NVIDIA 2080Ti GPUs. Without specific tuning, we set the learning rate $\gamma = 0.02$ for VGG-11 and $\gamma = 0.1$ for CNN. We use an SGD mini-batch size of 20. Each experiment took around 6 hours. We run each experiment 10 times then plot their average.

For each experiment, we plot both results of training loss and test accuracy. The training loss and test accuracy are computed over the whole training set and the whole test set, respectively. The communication time is emulated by measuring the round-trip time of transmitting the model between a worker device (in a home) and local (near) / global (far) Amazon EC2 instances, as summarized in Table E.1.

Datasets. In our experiments, we study the performance of our algorithm on three real datasets: CIFAR-10, FEMNIST, and CelebA. The details are as follows.

- **CIFAR-10.** The CIFAR-10 dataset (Krizhevsky, Hinton et al. 2009) consists of 60000 32×32 color images in 10 classes, with 6,000 images per class. There are 50,000 training images and 10,000 test images. There are the same number of data points on each worker. This dataset (<https://www.cs.toronto.edu/~kriz/cifar.html>) has a citation requirement which we have satisfied.
- **FEMNIST.** FEMNIST is a federated version of the EMNIST dataset proposed by LEAF (Caldas et al. 2019), which consists of 28×28 gray images in 62 classes. We use 156 workers as training data which consists of 34,659 data points and the test dataset consists of 4,973 data points. This corresponds to 5% of the entire dataset (which is the one of the standard settings of LEAF), while ignoring workers that have less than 100 data samples. This dataset is from LEAF (<https://github.com/TalwalkarLab/leaf>) which has a BSD 2-Clause license.
- **CelebA.** The CelebA dataset (Liu et al. 2015) consists of 32×32 color images in binary classes. We use the data partition for CelebA proposed by LEAF (Caldas et al. 2019). We use 472 workers as training data which consists of 8,855 data points and the test dataset consists of 1,159 data points. This corresponds to 5% of the entire dataset (which is the one of the standard settings of LEAF), while ignoring workers that have less than 5 data samples. This dataset is also from LEAF (<https://github.com/TalwalkarLab/leaf>) which has a BSD 2-Clause license.

The reason that we use 5% of the FEMNIST and CelebA datasets is because processing and storing the entire dataset requires excessive computation and disk space. We list out the statistics of the number of data points on each worker of FEMNIST and CelebA datasets in Table E.2.

Table E.1: Per-round communication time (ms)

Model	Local (near) EC2	Global (far) EC2
CNN	0.29 ± 0.14	4.53 ± 0.66
VGG-11	27.81 ± 4.61	291.82 ± 15.11

Table E.2: Statistics of the number of data points

Dataset	Number of workers	Total (all workers)	Number of data points Mean (each worker)	Standard deviation (each worker)
FEMNIST	156	34,659	222	92
CelebA	472	8,855	19	7

Models. We train VGG-11 models on CIFAR-10 and CelebA datasets. For FEMNIST dataset, we train a 9-layer CNN model. The structure of the CNN is $5 \times 5 \times 32$ Convolutional $\rightarrow 2 \times 2$ MaxPool \rightarrow Local Response Normalization $\rightarrow 5 \times 5 \times 32$ Convolutional \rightarrow Local Response Normalization $\rightarrow 2 \times 2$ MaxPool $\rightarrow 1568 \times 256$ Fully connected $\rightarrow 256 \times 62$ Fully connected \rightarrow Softmax.

Data Partitioning. For experiments with FEMNIST and CelebA, we use the data partition proposed by LEAF. We manually partition CIFAR-10 in a non-IID manner. The assigned label for each worker is different. For H-SGD with $N = 2$, we partition 10 workers into two groups. Group 1 contains only data with labels $\{0, 1, 2, 3, 4\}$ and Group 2 contains only data with labels $\{5, 6, 7, 8, 9\}$. In Group-IID/Group-non-IID experiments, we consider a different way of partitioning. For the “*group-IID*” case, the 5 workers in each group have labels $\{0, 1\}, \{2, 3\}, \{4, 5\}, \{6, 7\}, \{8, 9\}$, respectively. The basic idea is to make the upward divergence nearly zero in this case. In the “*group-non-IID*” case, workers in the first group have labels $\{0, 1\}, \{2, 3\}, \{4, 5\}, \{0, 1\}, \{2, 3\}$, respectively, while workers in the second group have labels

$\{4, 5\}$, $\{6, 7\}$, $\{8, 9\}$, $\{6, 7\}$, $\{8, 9\}$, respectively. Experimental results for CIFAR-10 with more workers and more groups are also shown in Figures E.4–E.6 later.

E.1 Additional Results for Two-level Case

In this section, we will provide all the training loss results and results based on iterations for FEMNIST and CelebA datasets of experiments in Section 6, as shown in Figures E.1, E.2, and E.3. It can be observed that the behavior of the training loss results are consistent with that of test accuracy shown in Section 6. We summarize the final test accuracy in Table E.3.

Table E.3: Final test accuracy (%) for two-level cases.

CIFAR-10	$P = 5$	$P = 10$	$P = 50$	$G = 50, I = 5$	$G = 50, I = 10$
	82.3 ± 0.5	77.9 ± 1.2	50.5 ± 1.8	78.5 ± 0.4	69.6 ± 1.5
FEMNIST	$P = 10$	$P = 50$	$P = 100$	$G = 100, I = 10$	$G = 100, I = 20$
	83.5 ± 0.3	80.6 ± 1.1	79.6 ± 1.6	82.6 ± 0.6	81.7 ± 1.1
CelebA	$P = 10$	$P = 50$	$G = 50, I = 10$		
	96.9 ± 0.2	86.9 ± 0.5	96.8 ± 0.3		




Figure E.1: Results with FEMNIST; (a), (b) training loss v.s. communication time and iterations, (c) test accuracy v.s. iterations.




Figure E.2: Results with CelebA; (a), (b) training loss v.s. communication time and iterations, (c) test accuracy v.s. iterations.

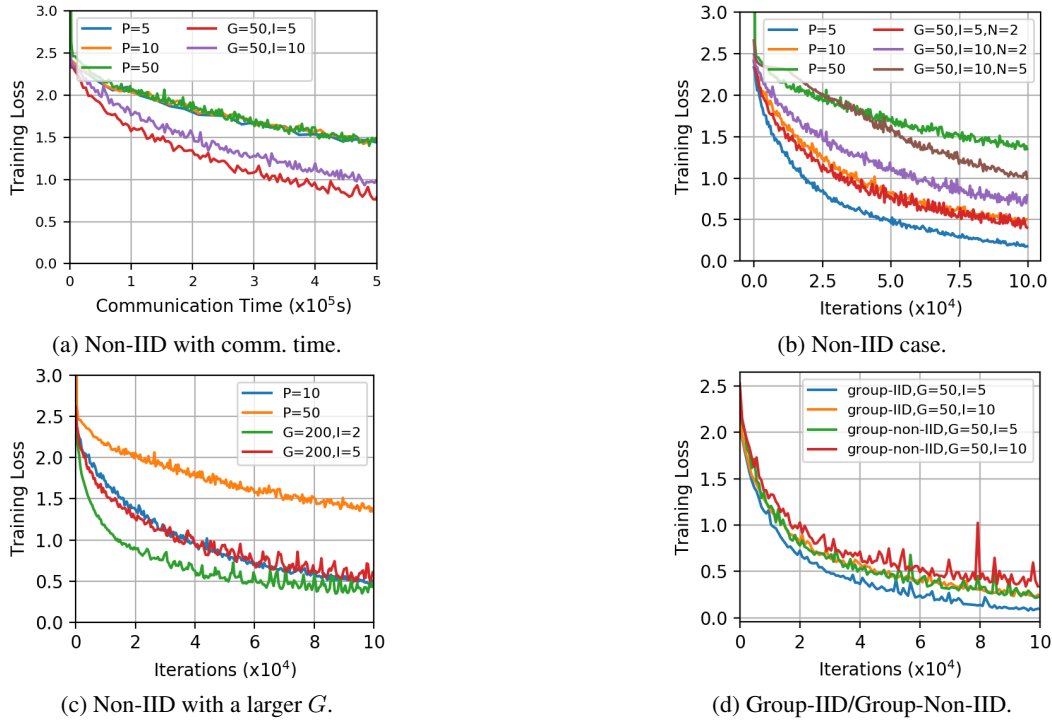


Figure E.3: Results with CIFAR-10. Training loss v.s. (a) communication time, (b)–(d) local iterations. By default, $N = 2$.

Further, we provide experiments with CIFAR-10 dataset for partial worker participation in two-level HF-SGD with random worker sampling in each round, as shown in Figures E.4–E.6. We separate data into 50 workers in a non-IID manner. There are the same number of workers in each group. For each HG-SGD round, we uniformly sample 20% of workers in each group. The results show that the same insights as described in Section 6 of the main paper can be observed here as well.

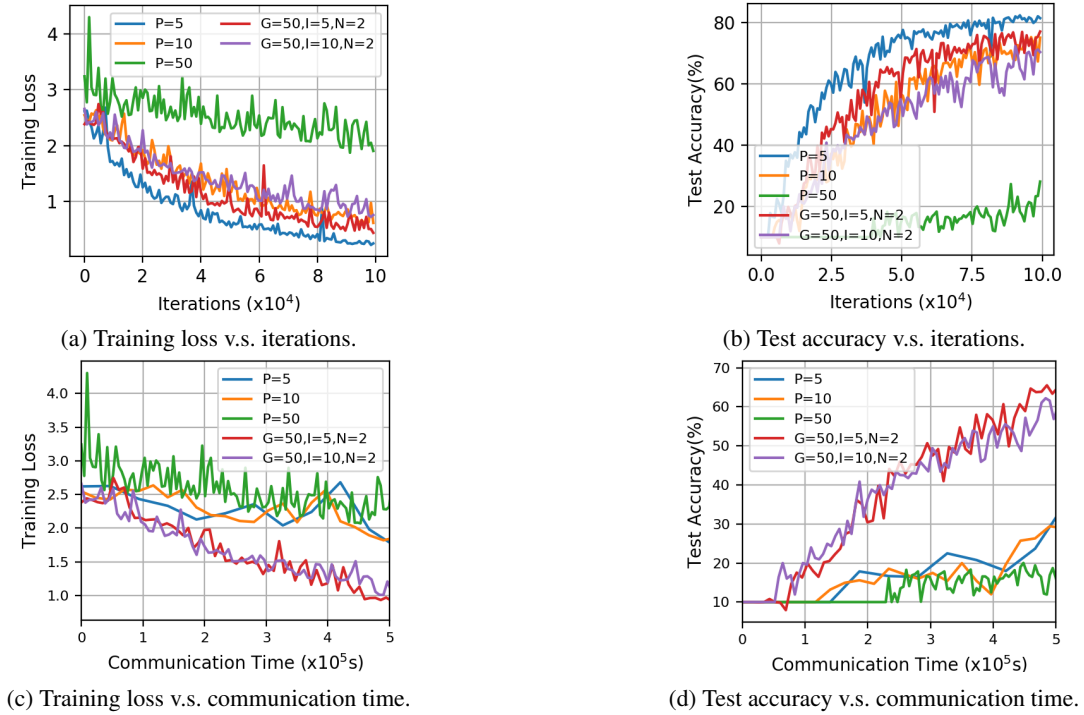


Figure E.4: Comparison between local SGD and H-SGD.

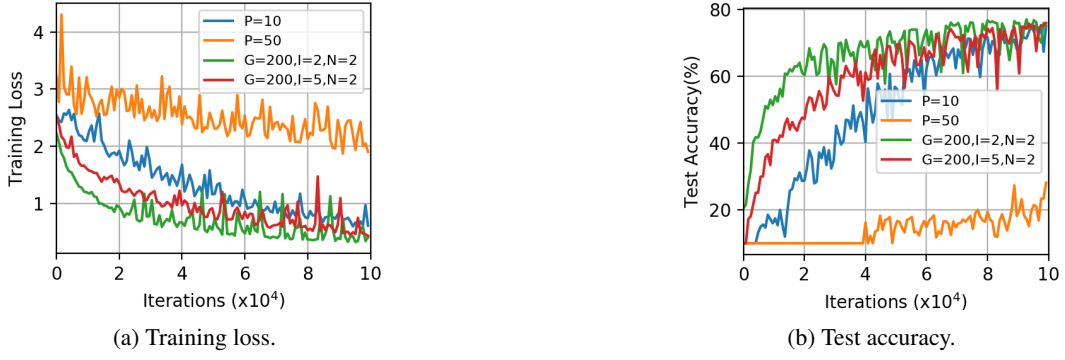


Figure E.5: Comparison between SF-SGD and HF-SGD with a large G .




Figure E.6: Performance comparison for different group sizes.

E.2 Multi-level Experiments with ($M = 3$)

In order to show the advantage of the multi-level HF-SGD and demonstrate the effectiveness of our analysis shown in Theorem 3, in this section, we provide experimental results of the 3-level case using the CIFAR-10 dataset. We construct the 3-level system shown in Figure E.7, where for example, server (1, 1) in level 2 computes the local average of SGD from workers (1, 1, 1) and (1, 1, 2), and server (1) in level 1 computes the average of SGD from servers (1, 1) and (1, 2) in level 2. Here, the CIFAR-10 dataset is partitioned into 10 workers such that each worker has data of one class. For example, the worker (1, 1, 1) has the data of label 0.




Figure E.7: The 3-level structure we used in experiments. The number inside each circle representing the worker is the label available at this worker.

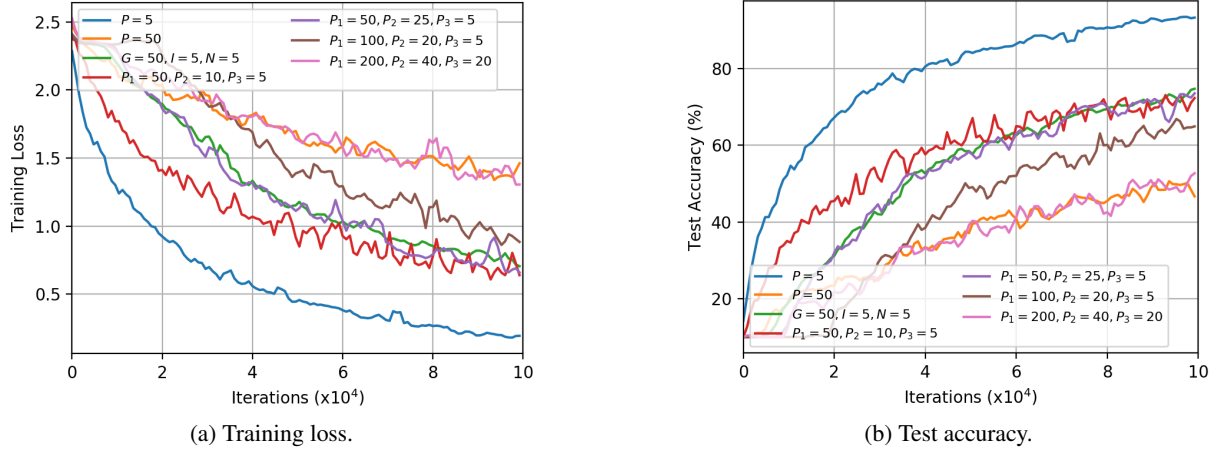


Figure E.8: 3-level experiments v.s. iterations. For 3-level case, we set $N_1 = 2, N_2 = 5$ by default.

From Figure E.8, first, we can observe that second level aggregation helps. For example, compare the 2-level case with $G = 50, I = 5$ (green curve) to the 3-level cases with $P_1 = 50, P_2 = 10, P_3 = 5$ (red curve) and $P_1 = 50, P_2 = 25, P_3 = 5$ (purple curve), more second level aggregation (smaller P_2) can improve the performance. Note that the two-level case with G, I is equivalent to three-level case with $P_1 = P_2 = G, P_3 = I$. Second, compared with single-layer case, 3-level cases also shows the same “sandwich” behavior as in 2-level cases, e.g., the red curve ($P_1 = 50, P_2 = 10, P_3 = 5$) is between the blue curve ($P = 5$) and the orange curve ($P = 50$) in Figure E.8. Third, it can be seen that although the case with $P_1 = 100, P_2 = 20, P_3 = 5$ has a larger global period than single-layer case $P = 50$, its performance (brown curve) is still better than single-layer case (orange curve). Similar behavior can be observed even if we increase P_1, P_2, P_3 at the same time for the case with $P_1 = 200, P_2 = 40, P_3 = 20$ (magenta curve). This implies that, by doing aggregations at different levels with a larger global period, we can reduce the communication time while keeping the same performance.

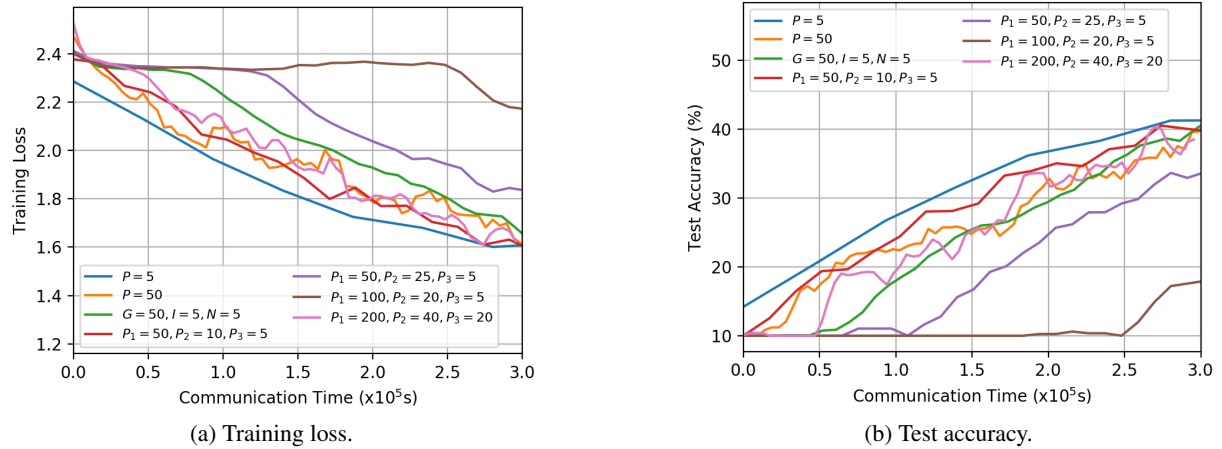


Figure E.9: 3-level experiments v.s. communication time. For 3-level case, we set $N_1 = 2, N_2 = 5$ by default.

In Figure E.9, we plot the training loss/test accuracy of 3-level cases v.s. communication time. In this experiment, we use the communication time measured for VGG-11 model shown in Table E.1. Here we use the local EC2 time in Table E.1 as communication time needed for aggregation by second-level servers. The global EC2 time in Table E.1 is used as communication time needed for aggregation by the global server. We set the ratio of the communication times between the first-level servers to the workers and second-level to the workers to be 2 : 1.¹ We can see from the figure that the cases of $P_1 = 50, P_2 = 10, P_3 = 5$ and $P_1 = 200, P_2 = 40, P_3 = 20$ can achieve better performance compared to others under the same total communication time.

¹Note that if the communication time of the second-level server to the worker is too long, then this additional level of aggregation is not needed.

Additional References for Appendix

- Caldas, S.; Wu, P.; Li, T.; Konečný, J.; McMahan, H. B.; Smith, V.; and Talwalkar, A. 2019. LEAF: A Benchmark for Federated Settings. In *NeurIPS Workshop on Federated Learning for Data Privacy and Confidentiality*.
- Lian, X.; Zhang, C.; Zhang, H.; Hsieh, C.-J.; Zhang, W.; and Liu, J. 2017. Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent. In *NeurIPS*, 5330–5340.