
Federated Learning under Periodic Client Participation and Heterogeneous Data: A New Communication-Efficient Algorithm and Analysis

Michael Crawshaw

Department of Computer Science
George Mason University
Fairfax, VA 22030
mcrawsha@gmu.edu

Mingrui Liu

Department of Computer Science
George Mason University
Fairfax, VA 22030
mingrui.liu@gmu.edu

Abstract

In federated learning, it is common to assume that clients are always available to participate in training, which may not be feasible with user devices in practice. Recent works analyze federated learning under more realistic participation patterns, such as cyclic client availability or arbitrary participation. However, all such works either require strong assumptions (e.g., all clients participate almost surely within a bounded window), do not achieve linear speedup and reduced communication rounds, or are not applicable in the general non-convex setting. In this work, we focus on nonconvex optimization and consider participation patterns in which the chance of participation over a fixed window of rounds is equal among all clients, which includes cyclic client availability as a special case. Under this setting, we propose a new algorithm, named Amplified SCAFFOLD, and prove that it achieves linear speedup, reduced communication, and resilience to data heterogeneity simultaneously. In particular, for cyclic participation, our algorithm is proved to enjoy $\mathcal{O}(\epsilon^{-2})$ communication rounds to find an ϵ -stationary point in the non-convex stochastic setting. In contrast, the prior work under the same setting requires $\mathcal{O}(\kappa^2\epsilon^{-4})$ communication rounds, where κ denotes the data heterogeneity. Therefore, our algorithm significantly reduces communication rounds due to better dependency in terms of ϵ and κ . Our analysis relies on a fine-grained treatment of the nested dependence between client participation and errors in the control variates, which results in tighter guarantees than previous work. We also provide experimental results with (1) synthetic data and (2) real-world data with a large number of clients ($N = 250$), demonstrating the effectiveness of our algorithm under periodic client participation.

1 Introduction

Federated learning (FL) [27, 21, 16, 53] is a distributed learning paradigm that emphasizes client privacy [28, 41, 29], limited communication [19, 25], and data heterogeneity across clients [17, 55]. FL has attracted attention in recent years due to the ability to leverage data and compute from user devices while respecting privacy [49, 9]. For large-scale FL, it is common to limit the number of simultaneously participating devices, and many works do so by assuming that a random subset of clients can be sampled independently at each round [27, 3, 50, 22, 38]. However, this pattern of client participation is not always practical. If clients are user devices like mobile phones, they may not have 24/7 availability due to low battery or bad internet connection [15, 31]. In particular, if client availability is correlated with geographical location (e.g. mobile phones charging at night), then

Table 1: Communication and computation complexity of various methods to find an ϵ -stationary point for L -smooth, non-convex objectives. N : number of clients, κ : data heterogeneity $\sup_{\mathbf{x}} \|\nabla f_i(\mathbf{x}) - \nabla f(\mathbf{x})\| \leq \kappa$. S : number of participating clients per round, \bar{K} : number of groups for cyclic participation. See Section 3.2 for a description of each participation pattern. We say that an algorithm exhibits reduced communication if its dependence in terms of ϵ is strictly smaller than $\mathcal{O}(\epsilon^{-4})$. Derivation of complexities for Amplified FedAvg can be found in Appendix C.

Setting	Communication Complexity (R)	Iteration Complexity (RI)	Reduced Communication	Unaffected by Heterogeneity
i.i.d. Participation (S)				
FedAvg [17]	$\frac{\Delta\kappa^2}{S\epsilon^4} \left(1 - \frac{S}{N}\right) + \frac{\sqrt{L}\kappa}{\epsilon^3}$	$\frac{\Delta L\sigma^2}{S\epsilon^4}$	✗	✗
SCAFFOLD [17]	$\frac{\Delta L}{\epsilon^2} \left(\frac{N}{S}\right)^{2/3}$	$\frac{\Delta L\sigma^2}{S\epsilon^4}$	✓	✓
Amplified FedAvg [39]	$\frac{\Delta LN\kappa^2}{S\epsilon^4} \left(1 - \frac{S}{N}\right) + \frac{\Delta L + \kappa^2}{\epsilon^2}$	$\frac{\Delta L\sigma^2}{S\epsilon^4}$	✗	✗
Amplified SCAFFOLD (ours)	$\frac{\Delta L N}{\epsilon^2 S}$	$\frac{\Delta L\sigma^2}{S\epsilon^4}$	✓	✓
Regularized Participation (P, ρ)				
Amplified FedAvg [39]	$\frac{\Delta LP + \kappa^2 + \sigma^2}{\epsilon^2}$	$\frac{\Delta L\rho^2\sigma^2}{\epsilon^4}$	✓	✗
Amplified SCAFFOLD (ours)	$\frac{\Delta LP}{\epsilon^2}$	$\frac{\Delta L\rho^2\sigma^2}{\epsilon^4}$	✓	✓
Cyclic Participation (\bar{K}, S)				
Amplified FedAvg [39]	$\frac{\Delta LN\kappa^2}{S\epsilon^4} \left(1 - \frac{S\bar{K}}{N}\right) + \frac{\Delta L\bar{K} + \kappa^2}{\epsilon^2}$	$\frac{\Delta L\sigma^2}{S\epsilon^4}$	✗	✗
Amplified SCAFFOLD (ours)	$\frac{\Delta L\bar{K} N}{\epsilon^2 S}$	$\frac{\Delta L\sigma^2}{S\epsilon^4}$	✓	✓

client availability follows a cyclic pattern [54]. Therefore, it remains an important open question to design federated optimization algorithms with provable efficiency under non-i.i.d client participation.

Several works have investigated optimization in FL under non-i.i.d. client participation [10, 2, 8, 39, 46]. However, to the best of our knowledge, no existing algorithm in a non-i.i.d. participation setting provably exhibits reduced communication cost, linear speedup with respect to the number of clients, and resilience to client data heterogeneity for general non-convex optimization.

In this work, we consider FL under an arbitrary participation framework [39], where client participation during each round is a random variable with potentially unknown distribution. We focus on client participation patterns that are periodic, in the sense that all clients are expected to participate with equal frequency over a window of multiple training rounds.

For this setting, we propose Amplified SCAFFOLD, an optimization algorithm for FL under periodic client participation. Amplified SCAFFOLD utilizes (a) amplified updates across participation periods and (b) control variates computed across entire participation periods, to eliminate the effect of data heterogeneity even under non-i.i.d. participation. We show that Amplified SCAFFOLD exhibits significantly reduced communication cost, linear speedup, and is unaffected by client data heterogeneity. To the best of our knowledge, this is the first result demonstrating reduced communication or resilience to data heterogeneity without assuming i.i.d. participation. The complexity of Amplified SCAFFOLD is compared against baselines in Table 1. For cyclic participation, Amplified SCAFFOLD improves the previous best communication cost from $\mathcal{O}(\kappa^2\epsilon^{-4})$ to $\mathcal{O}(\epsilon^{-2})$.

The main challenges of achieving these properties are (1) simultaneously handling randomness from stochastic gradients and non-i.i.d. participation; and (2) controlling the error of control variates under non-i.i.d. participation. Previous work in this setting [39] performs an in-expectation analysis, by taking expectation only over randomness from stochastic gradients; this avoids (1) but cannot leverage properties of the participation pattern to reduce communication. We present a tighter analysis that addresses (1) by taking expectation over both client participation and the stochastic gradients throughout the analysis, and carefully treating the trajectory variables which depend on both sources of randomness. We address (2) by recursively bounding the control variate errors, which involves a non-uniform average of non-uniform averages of error terms resulting from non-i.i.d. participation. We show that this nested non-uniform average can be bounded using mild regularity conditions on the participation pattern.

Our contributions are summarized below.

- We introduce Amplified SCAFFOLD, an optimization algorithm for federated learning under non-i.i.d. client participation. Our convergence analysis demonstrates its computational and communication efficiency: Amplified SCAFFOLD exhibits reduced communication, linear speedup, and is unaffected by data heterogeneity. These guarantees are achieved with a tighter analysis than used in previous work [39], with a fine-grained treatment of the two sources of randomness: client participation and stochastic gradients. In the case of cyclic participation, we reduce the previous best communication cost of $\mathcal{O}(\kappa^2 \epsilon^{-4})$ to $\mathcal{O}(\epsilon^{-2})$.
- Experimental results show that Amplified SCAFFOLD converges faster than baselines on both synthetic and real-world problems under realistic non-i.i.d. client participation patterns. We also include an ablation study which demonstrates the robustness of our algorithm to changes in data heterogeneity, the number of participating clients per round, and the number of client groups in cyclic participation.

The paper is outlined as follows. We discuss related work in Section 2 and Section 3 provides a formal specification of the optimization problem. Amplified SCAFFOLD is introduced and theoretically analyzed in Section 4, and we provide experiments in Section 5. We conclude with Section 6.

2 Related Work

Federated Optimization. FedAvg [27] characterizes partial client participation and local updates in each round. FedAvg was analyzed in the full participation setting [35] [37] [50] [51] [43] [42] [18] [12]. Other federated optimization algorithms aim to improve communication efficiency [32] [52] and tackle data heterogeneity [22] [17]. The analysis of FL optimization algorithms typically either assumes full client participation or partial client participation where clients are sampled uniformly randomly [47] [38] [17] [23] [43] [1]. [30] provides lower bounds for distributed stochastic, smooth optimization with intermittent communication and non-convex objectives, both in the full and partial participation settings. They also include algorithms employing variance reduction which match (or closely match) lower bounds in the full and partial participation settings. However, none of the works above are applicable for general participation patterns such as periodic participation.

Client Participation. Cyclic data sampling was considered for stochastic convex optimization in [10], where they propose “pluralistic” solutions instead of learning a single model for all clients. There is a recent line of work considering various participation patterns, including client selection [11] [4] [33] biased participation [34] [6] [7], independent participation across rounds [17] [22] [23], unbiased participation [36] [13], bounded rounds of unavailability [46] [14] [48], asynchronous participation [24] [2], cyclic participation [8], and arbitrary participation [39] [40]. However, none of these works enjoy linear speedup, reduced communication rounds, and resilience to data heterogeneity under the general setting of non-convex objectives and periodic client participation. Concurrent work [44] considers non-stationary client participation under the condition that for every client and every round, the probability of participation is bounded away from zero; for this setting, they propose an algorithm with linear speedup and dependence on gradient dissimilarity for non-convex, Lipschitz objectives.

See Appendix F for a detailed discussion comparing our results with a small number of closely related baselines.

3 Problem Setup

We consider a federated learning problem with N clients, with the overall objective

$$\min_{\mathbf{x} \in \mathbb{R}^d} \left\{ f(\mathbf{x}) := \frac{1}{N} \sum_{i=1}^N f_i(\mathbf{x}) \right\},$$

where each $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$ is the local objective of one client. We consider the stochastic optimization problem, so that $f_i(\mathbf{x}) = \mathbb{E}_{\xi \sim \mathcal{D}_i} [F(\mathbf{x}; \xi)]$, and the optimization algorithm can access $F_i(\mathbf{x}; \xi)$ and $\nabla F_i(\mathbf{x}; \xi)$ for individual values of ξ . We make the following assumptions about the objectives:

Assumption 1. **(a)** $f(\mathbf{x}_0) - \min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x}) \leq \Delta$. **(b)** Each f_i is L -smooth, i.e., $\|\nabla f_i(\mathbf{x}) - \nabla f_i(\mathbf{y})\| \leq L \|\mathbf{x} - \mathbf{y}\|$ for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^d$. **(c)** The stochastic gradient has variance σ^2 , i.e., $\mathbb{E}_{\xi \sim \mathcal{D}_i} [\|\nabla F_i(\mathbf{x}; \xi) - \nabla f_i(\mathbf{x})\|^2] \leq \sigma^2$ for all $\mathbf{x} \in \mathbb{R}^d$.

Since each f_i may be non-convex, we consider the problem of finding an ϵ -stationary point of f , that is, a point $\mathbf{x} \in \mathbb{R}^d$ such that $\|\nabla f(\mathbf{x})\| \leq \epsilon$.

3.1 Participation Framework

We consider a federated learning framework consisting of R rounds. For any round $r \in \{0, \dots, R-1\}$ and client $i \in [N]$, the availability of client i at round r is a random variable q_r^i , following the arbitrary participation framework of [39]. If $q_r^i = 0$, then client i may not participate during round i . For example, under the conventional i.i.d. sampling of clients, at each round r a subset of clients $\mathcal{S}_r \subset [N]$ is sampled uniformly without replacement, and the weights are set as $q_r^i = \frac{\mathbb{1}\{i \in \mathcal{S}_r\}}{S}$.

For some $P \in \mathbb{N}$, let \mathcal{Q}_{r_0} be the filtration generated by $\{q_r^i : r_0 \leq r < r_0 + P, i \in [N]\}$, let \mathcal{Q} be the filtration generated by $\mathcal{Q}_0, \dots, \mathcal{Q}_{R-P}$, and let \mathcal{G} be the filtration generated by $\{\xi_{r,k}^i : 0 \leq r < R, 0 \leq k < I, i \in [N]\}$, where $\xi_{r,k}^i$ is the random sampling of the stochastic gradient of round r , step k , client i . We make the following assumptions about the participation distribution.

Assumption 2. For all $r \in \{0, \dots, R-1\}$: (a) $\sum_{i=1}^N q_r^i = 1$ and $\sum_{i=1}^N (q_r^i)^2 \leq \rho^2$. (b) The distribution of $\{q_r^i\}$ is unbiased across clients over every window of P rounds, i.e., $\mathbb{E}_{\mathcal{Q}_{r_0}}[\frac{1}{P} \sum_{r=mP}^{(m+1)P-1} q_r^i] = 1/N$ for every $m < R/P$ and $i \in [N]$. (c) Each client has a non-zero probability of being sampled over every window of P rounds, i.e., $\mathbb{P}_{\mathcal{Q}_{r_0}}(\frac{1}{P} \sum_{r=mP}^{(m+1)P-1} q_r^i > 0) > p_{\text{sample}}$ for every $m < R/P$ and $i \in [N]$. (d) \mathcal{Q} and \mathcal{G} are independent.

For each round, Assumption 2(a) enforces that the participation weights q_r^i are normalized to sum to 1, and characterizes the spread of participation weights across clients with the constant ρ^2 . Assumption 2(b) requires that the set of rounds can be partitioned into windows of length P within which clients are expected to participate with equal frequency. Lastly, Assumption 2(c) enforces that within each window, for each client the probability of being sampled is nonzero. Conventional i.i.d client sampling satisfies Assumption 2 with $\rho = S^{-1/2}$, $P = 1$, and $p_{\text{sample}} = S/N$.

An important difference from conventional i.i.d. participation is that here, client participation is not necessarily independent across rounds. Accordingly, we emphasize that the expectation and probability in Assumptions 2(b)-(c) are taken only over \mathcal{Q}_{r_0} . Therefore, the mean participation weight in Assumption 2(b) may itself be a random variable if client participation at some rounds is dependent on the outcome of participation in previous rounds. Similarly, the sampling probability in Assumption 2(c) may be a random variable. For the participation patterns considered in the next section, Assumption 2 is satisfied even when client sampling is not independent across rounds.

3.2 Specific Participation Patterns

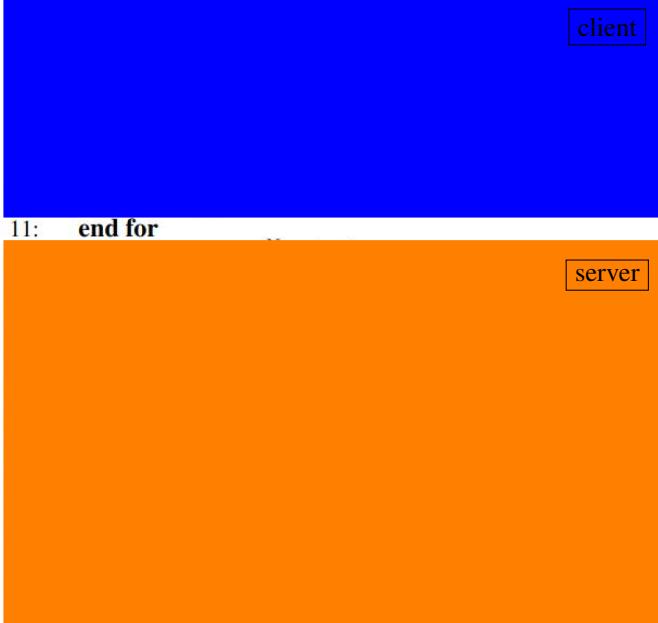
Regularized Participation We say that client participation is *regularized* [39] if $\bar{q}_{r_0}^i = \frac{1}{N}$ almost surely for all r_0 and i , where $\bar{q}_{r_0}^i$ is defined on Line 18 of Algorithm 1 as the participation of client i averaged over rounds $r_0, \dots, r_0 + P - 1$. In this case, Assumption 2 is satisfied with $p_{\text{sample}} = 1$, while P and ρ^2 are parameters of the participation pattern. Regularized participation is a relatively strong constraint, since every client must participate within each window, which may not be practical. However, it is flexible in that there is no constraint on how clients participate within each window. Regularized participation was also considered for strongly convex objectives [26].

Cyclic Participation Following the CyCP framework [8], N clients are partitioned into \bar{K} equally sized subsets, and at round r only clients in group $(r \bmod \bar{K})$ may participate. S clients are sampled without replacement from group $(r \bmod \bar{K})$, for whom the participation weight is $q_r^i = 1/S$. All other clients are assigned $q_r^i = 0$. Cyclic participation satisfies Assumption 2 with $P = \bar{K}$, $\rho = S^{-1/2}$, and $p_{\text{sample}} = \bar{S}\bar{K}/N$. Notice that i.i.d. client sampling is the special case where $\bar{K} = 1$.

Cyclic participation can model a situation where each client group is available at a different time of day. For example, if client devices are mobile phones, then clients are available for participation at night, when phones are charging, likely to have internet connection, and otherwise idle. If devices are spread across the globe, then client groups are naturally formed by time zones. Cyclic participation is less stringent than regularized participation since not all clients are required to participate within each window. FedAvg was analyzed under cyclic participation for PL objectives [8], although this analysis is not applicable in our setting, which uses general non-convex objectives.

Algorithm 1 Amplified SCAFFOLD

```
1: Initialize  $\bar{x}_0, \mathbf{G}_0^i, \mathbf{G}_0 \leftarrow \frac{1}{N} \sum_{i=1}^N \mathbf{G}_0^i, r_0 \leftarrow 0, \mathbf{u} \leftarrow 0$ 
2: for  $r = 0, 1, \dots, R - 1$  do
3:   for  $i \in [N]$  do
```



```
11:   end for
```

```
24:   end if
25: end for
```

4 Algorithm and Analysis

In this section, we present Amplified SCAFFOLD, our algorithm to solve the FL problem described in Section 3. Pseudocode for Amplified SCAFFOLD is shown in Algorithm 1. The main components of the Amplified SCAFFOLD algorithm are (1) amplified updates and (2) long-range control variates.

4.1 Algorithm Overview

To deal with the non-stationarity of client availability, Amplified SCAFFOLD performs amplified updates based on information accumulated over a window of P rounds. In Algorithm 1, the variable \mathbf{u} holds a weighted average of local updates to client models, weighted by client participation. Every P rounds, the global model is updated in the direction \mathbf{u} scaled by the amplification factor γ . Informally, the direction \mathbf{u} includes information from all clients with equal representation, according to Assumption 2(b). Similar amplified updates are used in Amplified FedAvg [39].

Control variates for heterogeneous federated learning were first introduced by SCAFFOLD [17]. However, SCAFFOLD-style control variates are updated every time a client participates, which may not be appropriate under periodic availability. For example, under non-i.i.d. participation control variates for different clients would be updated with different frequencies, so that some clients may have consistently less accurate control variates than others. Informally, this may lead to a bias in which some clients' objective is underweighted relative to others. To deal with this issue, Amplified SCAFFOLD updates control variates based on information accumulated over a window of P rounds, which enforces equal representation of all clients in expectation, according to Assumption 2(b).

Comparison with [17, 39] Although the two algorithmic components of Amplified SCAFFOLD individually appear in previous work [17, 39], we emphasize that our complexity results cannot be achieved by simply combining the analyses of these two works. The analysis of [39] requires ϵ^{-4} communication cost due to their treatment of the randomness in client participation. Here, we present

a tighter analysis with ϵ^{-2} cost from a more fine-grained treatment of the two sources of randomness (stochastic gradients and client sampling). See Section 4.4 for more details on our approach.

4.2 Main Results

Let $\hat{\mathbf{x}} = \bar{\mathbf{x}}_{mP}$, where m is sampled uniformly from $\{0, \dots, R/P-1\}$, and let $r_0 \in \{0, P, \dots, R/P\}$. Denote $w_{r_0}^i = \frac{1}{N} \sum_{j=1}^N \frac{\mathbb{1}\{\bar{q}_{r_0}^j > 0\}}{P\bar{q}_{r_0}^j} \sum_{s=r_0}^{r_0+P-1} q_s^i q_s^j$ and $v_{r_0}^i = \bar{q}_{r_0}^i - \frac{1}{N}$. Informally, $w_{r_0}^i$ represents the "non-uniformity" of the client sampling distribution. We also consider a variable $\Lambda_{r_0}^i$ that depends only on the client sampling distribution and characterizes the sample size from which $\mathbf{G}_{r_0}^i$ is computed. See Appendix A.1 for further discussion of these quantities. We consider convergence under the following conditions, which are satisfied by several participation patterns of interest.

$$\mathbb{E}[w_{r_0}^i] \leq \frac{P^2}{N} \quad \text{for all } r_0 \text{ and } i, \quad (1)$$

$$\mathbb{E} \left[\sum_{i=1}^N (v_{r_0}^i)^2 \Lambda_{r_0}^i \right] \leq \rho^2 \quad \text{for all } r_0 \text{ and } i, \quad (2)$$

Theorem 1. Suppose that Assumptions 1 and 2 hold, and that Equation 1 and Equation 2 hold. If $\gamma\eta \leq \frac{p_{\text{sample}}}{60LIP}$ and $\eta \leq \frac{\sqrt{p_{\text{sample}}}}{60LIP}$, then Algorithm 1 satisfies

$$\mathbb{E}[\|\nabla f(\hat{\mathbf{x}})\|^2] \leq \mathcal{O} \left(\frac{\Delta}{\gamma\eta IR} + (\gamma\eta L\rho^2 + \eta^2 L^2 IP) \sigma^2 \right).$$

Corollary 1. For any $\epsilon > 0$ and $I \geq 1$, there exist choices of γ and η such that $\mathbb{E}[\|\nabla f(\hat{\mathbf{x}})\|^2] \leq \mathcal{O}(\epsilon^2)$ as long as $R \geq \mathcal{O} \left(\frac{\Delta L \rho^2 \sigma^2}{I\epsilon^4} + \frac{\Delta LP}{p_{\text{sample}} \epsilon^2} \right)$.

The complexity of Amplified SCAFFOLD has several important properties:

Reduced Communication By choosing $I = \Theta(\Delta\rho^2\sigma^2p_{\text{sample}}P^{-1}\epsilon^{-2})$, Amplified SCAFFOLD has communication complexity $R = \mathcal{O}(LPp_{\text{sample}}^{-1}\epsilon^{-2})$, which improves upon the ϵ^{-4} complexity of parallel SGD. We are not aware of any existing work that achieves this communication reduction for non-convex federated optimization with periodic participation.

Unaffected by Heterogeneity The iterations RI and the number of communications R are unaffected by heterogeneity, which is not achieved for periodic participation by any existing work [39][8].

Linear Speedup The number of iterations $RI = \mathcal{O}(\Delta L \rho^2 \sigma^2 \epsilon^{-4})$ will exhibit linear speedup in the number of clients through the term ρ^2 , depending on the client participation pattern.

4.3 Application to Participation Patterns

The results from Section 4.2 apply under any participation pattern that satisfies Assumption 2, Equation 1, and Equation 2, and below we discuss the participation patterns discussed in Section 3.2. The complexity of Amplified SCAFFOLD for each participation pattern are shown in Table 1, and these results can be obtained by plugging ρ^2 , P , and p_{sample} into Corollary 1, together with a choice of I as described in Section 4.2. The derivations of each result below are given in Appendix B.

Regularized Participation Recall that regularized participation satisfies Assumption 2 with $p_{\text{sample}} = 1$, and P , ρ^2 are parameters of the participation pattern. Also, under regularized participation, $w_{r_0}^i = \bar{q}_{r_0}^i = 1/N$ almost surely, so that $\mathbb{E}[w_{r_0}^i] = 1/N \leq P^2/N$ and $v_{r_0}^i = 0$. Therefore Equation 1 and Equation 2 are satisfied. Plugging $p_{\text{sample}} = 1$ into Corollary 1 yields

$$R = \mathcal{O}(LP\epsilon^{-2}), \quad RI = \mathcal{O}(\Delta L \rho^2 \sigma^2 \epsilon^{-4}).$$

In this setting, our algorithm exhibits reduced communication and resilience to heterogeneity. To our knowledge, the only existing algorithm with theoretical guarantees for non-convex problems under regularized participation is Amplified FedAvg [39]. However, as seen in Table 1, the communication complexity of Amplified FedAvg has order ϵ^{-4} in terms of ϵ and suffers from a κ^2 dependence.

Cyclic Participation Recall that cyclic participation satisfies Assumption 2 with $P = \bar{K}$, $\rho = S^{-1/2}$, and $p_{\text{sample}} = S/N$. Also, $\mathbb{E}[w_{r_0}^i] = S/N^2 \leq P^2/N$ and $\mathbb{E}[(v_{r_0}^i)^2 \Lambda_{r_0}^i] = \rho^2$, so that Equation 1 and Equation 2 are satisfied. Based on the above parameter values, the resulting complexities are

$$R = \mathcal{O}\left(\frac{L\bar{K}}{\epsilon^2} \left(\frac{N}{S}\right)\right), \quad RI = \mathcal{O}\left(\frac{\Delta L\sigma^2}{S\epsilon^4}\right).$$

Again, Amplified SCAFFOLD achieves reduced communication, linear speedup, and resilience to heterogeneity. Amplified FedAvg [39] is the only existing algorithm with theoretical guarantees in this setting, but it fails to achieve resilience to heterogeneity or reduce communication cost outside of the trivial case of full participation ($\bar{K} = 1, S = N$). Also, even for the setting of PL-functions, the convergence rate of FedAvg under cyclic participation from [8] does not demonstrate an improvement with respect to the number of local steps. See Appendix F for further discussion of their results.

Recall that i.i.d. participation is a special case of cyclic participation with $\bar{K} = 1$. In this case, Amplified FedAvg fails to recover the reduced communication usually achieved under i.i.d. participation, such as by SCAFFOLD [17]. In fact, Amplified FedAvg fails to recover the communication cost of FedAvg under i.i.d. participation, requiring an additional factor of LN . The larger communication cost of Amplified FedAvg is a result of its convergence analysis, which does not leverage the property of unbiased participation (Assumption 2(b)) during the analysis, and requires $P = \mathcal{O}(\epsilon^{-2})$ in order to converge (see Appendix C for more details). In contrast, Amplified SCAFFOLD succeeds in recovering the results of SCAFFOLD under i.i.d. participation, with only a slightly worse dependence of R on N/S . This difference in the order of $\frac{N}{S}$ is due to a potential small issue in the analysis of SCAFFOLD, which we intentionally avoided by accepting a slightly worse dependence on $\frac{N}{S}$. We provide a detailed discussion of the N/S dependence in Appendix F.

4.4 Proof Sketch

The main challenges for demonstrating convergence are (1) simultaneously handling randomness from stochastic gradients and non-i.i.d. client sampling, and (2) controlling error of control variates under non-i.i.d. client sampling. Previous work [39] subverts (1) by conditioning on \mathcal{Q} throughout the entire analysis. However, this eliminates the possibility of utilizing the condition $\mathbb{E}[q_{r_0}^i] = 1/N$, and ultimately incurs a dependence on the data heterogeneity (see the term $\tilde{\delta}^2(P)$ in Theorem 3.1 of [39]). Instead, we take expectation over both sources of randomness throughout the analysis, which requires a careful treatment of each iterate's dependence, and enables communication reduction. For (2), previous analysis of federated algorithms with control variates [17] recursively bounds the error of control variates between consecutive rounds. However, this recursion crucially depends on i.i.d. client participation. We extend this analysis to our setting, establishing a recursion over the control variate error between consecutive windows of P rounds. Establishing this recursion under non-i.i.d. participation involves a non-uniform average of non-uniform averages of error terms, which we handle by invoking the regularity conditions stated in Equation 1 and Equation 2.

Using smoothness of f , the objective function decrease $f(\bar{\mathbf{x}}_{r_0+P}) - f(\bar{\mathbf{x}}_{r_0})$ is upper bounded by $\langle \nabla f(\bar{\mathbf{x}}_{r_0}), \bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0} \rangle + \frac{L}{2} \|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2$. Letting $\bar{\mathbf{x}}_{r,k} = \sum_{i=1}^N q_r^i \mathbf{x}_{r,k}^i$ be a weighted average of local models, the sum of the previous inner product and quadratic terms can be bounded by $-\gamma\eta IP \|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2$, plus standard noise terms, the additional "drift" terms

$$\tilde{D}_{r,k} = \sum_{i=1}^N q_r^i \|\mathbf{x}_{r,k}^i - \bar{\mathbf{x}}_{r,k}\|^2 \quad \tilde{M}_{r,k} = \|\bar{\mathbf{x}}_{r,k} - \bar{\mathbf{x}}_{r_0}\|^2,$$

and control variate errors $C_{r_0}^i = \|\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \mathbf{G}_{r_0}^i\|^2$. $\tilde{D}_{r,k}$ captures the distance between local client models, while $\tilde{M}_{r,k}$ captures the distance from local models to the previous global model $\bar{\mathbf{x}}_{r_0}$.

Taking conditional expectation $D_{r,k} = \mathbb{E}[\tilde{D}_{r,k} | \mathcal{Q}]$ and $M_{r,k} = \mathbb{E}[\tilde{M}_{r,k} | \mathcal{Q}]$, Lemma 1 bounds the drift terms by establishing and unrolling a mutually recurrent relation between $D_{r,k}$ and $M_{r,k}$. The resulting bound involves a non-uniform average over the control variate errors: $\sum_{i=1}^N q_r^i \mathbb{E}[C_{r_0}^i | \mathcal{Q}]$.

Denoting the average control variate error $C_{r_0} = \frac{1}{N} \sum_{i=1}^N C_{r_0}^i$, we want to bound $\mathbb{E}[C_{r_0+P}]$ in terms of $\mathbb{E}[C_{r_0}]$. C_{r_0+P} can be decomposed into drift terms, but the result is a non-uniform average:

$$\sum_{s=r_0}^{r_0+P-1} \sum_{k=0}^{I-1} q_s^i (D_{s,k} + M_{s,k}).$$

Since the bound for each $M_{s,k} + D_{s,k}$ from Lemma 1 involves a non-uniform average over $C_{r_0}^i$, the resulting bound of C_{r_0+P} involves a non-uniform average of non-uniform averages of $C_{r_0}^i$, instead of the uniform average C_{r_0} . The regularity conditions in Equation 1 and Equation 2 allow us to bound this nested non-uniform average by a uniform average, which finishes the recursion.

Putting everything together, we obtain the descent inequality

$$\mathbb{E}[\tilde{f}_{r_0+P}] \leq \mathbb{E}[\tilde{f}_{r_0}] - \gamma \eta I P \mathbb{E}[\|\nabla f(\bar{x}_{r_0})\|^2] + \gamma \eta I P (\gamma \eta L \rho^2 + \eta^2 L^2 I P) \sigma^2,$$

where $\tilde{f}_{r_0} := f(\bar{x}_{r_0+P}) + \Phi(r_0 + P)$ and Φ is a potential function that depends on the control variate errors. Theorem 1 is then obtained by averaging over r_0 and isolating the gradient.

5 Experiments

We experimentally validate our algorithm for non-i.i.d. client participation under three settings: minimizing a synthetic function, logistic regression for Fashion-MNIST¹ [45], and training a CNN for CIFAR-10 [20]. We also include an ablation study on Fashion-MNIST, to investigate how each algorithm is affected by changes in data heterogeneity, the number of participating clients, and the number of client groups in cyclic participation.

5.1 Setup

All of our experiments utilize a non-i.i.d. client participation pattern similar to cyclic participation (discussed in Section 3.2). We partition the total set of N clients into \bar{K} equally sized subsets, and at each training round only a single client group is available for participation. In our experiments, the available group does not change every round; instead, each group is available for g rounds at a time. Under this pattern, Assumption 2 is satisfied with $P = g\bar{K}$. We refer to g as the availability time.

We evaluate five algorithms: FedAvg [27], FedProx [22], SCAFFOLD [17], Amplified FedAvg [39], and Amplified SCAFFOLD (ours). We tune each algorithm's parameters by grid search, including learning rate η , amplification rate γ , and FedProx's μ . The search ranges and tuned values can be found in Appendix D. All experiments were run on a single node with eight NVIDIA A6000 GPUs. Code is available at the following repository: <https://github.com/MingruiLiu-ML-Lab/FL-under-Periodic-Participation>

Synthetic We evaluate each algorithm's convergence on a difficult objective based on a lower bound for FedAvg [43]. The objective maps \mathbb{R}^4 to \mathbb{R} , is convex, and is parameterized by a smoothness L , stochastic gradient variance σ^2 , and heterogeneity κ , so that it satisfies Assumption 1 by construction. The complete definition of the objective can be found in Appendix D. Since there are only two distinct local objectives, we set the number of clients $N = 2$ and the number of sampled clients $S = 1$, and the number of groups $\bar{K} = 2$. All other settings can be found in Appendix D.

Fashion-MNIST and CIFAR-10 We evaluate each algorithm for training an image classifier, using logistic regression for Fashion-MNIST and a two-layer CNN for CIFAR-10. To simulate heterogeneous data in federated learning, we use a common protocol [17, 39], to partition each dataset into client datasets according to a data similarity parameter s . This protocol is detailed in Appendix D. Following [39], we set the number of clients $N = 250$, data similarity $s = 5\%$, and the number of sampled clients per round $S = 10$. For client participation, we set the number of groups $\bar{K} = 5$, so that each group contains clients that have majority label from two different classes. We run all baselines with 5 different random seeds and report the mean results with error bars in Section 5.2 (the radius of each error bar is 1 standard deviation). All other settings can be found in Appendix D.

Additional experimental results are provided in Appendix E where we compare against extra baselines (FedAdam [32], FedYogi [32], FedAvg-M [5], and Amplified FedAvg with FedProx regularization), and evaluate training under another non-i.i.d. client participation pattern.

¹Fashion-MNIST is licensed under the MIT License.

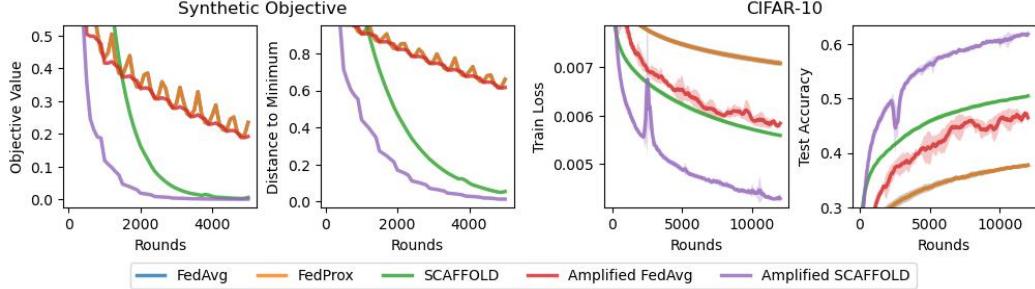


Figure 1: Results for synthetic objective and CIFAR-10. Left: Amplified SCAFFOLD and SCAFFOLD both converge to the global minimum, but Amplified SCAFFOLD converges significantly faster. Right: Amplified SCAFFOLD converges to the best solution by a significant margin. Note that in both cases, the curves for FedAvg and FedProx are nearly overlapping.

5.2 Main Results

Results for the synthetic experiment and CIFAR-10 are shown in Figure 1, and results for Fashion-MNIST are shown in Figure 2. We make the following observations:

Amplified SCAFFOLD converges the fastest. In all three settings, Amplified SCAFFOLD reaches the best overall solution among all algorithms (by all metrics) and requires the fewest communication rounds. In the synthetic experiment, Amplified SCAFFOLD requires 800 communication rounds to reach an objective value of 0.2, while SCAFFOLD requires 1900 rounds, and both FedAvg and Amplified FedAvg require 4800 rounds to reach the same objective value.

Amplified FedAvg is comparable to FedAvg. Amplified FedAvg shows slight improvement over FedAvg for the synthetic experiment and for Fashion-MNIST. Only for CIFAR-10 is Amplified FedAvg significantly faster than FedAvg, but there it also exhibits a reduction in stability. The underwhelming experimental performance of Amplified FedAvg corroborates our discussion from Section 4.3; Amplified FedAvg requires many communication rounds and suffers from data heterogeneity.

Contrary to our findings, the original evaluation of Amplified FedAvg [39] showed a significant improvement over FedAvg. One explanation is that the original evaluation employed pretraining using FedAvg, so that each algorithm was evaluated only for fine-tuning. Our experiments suggest that Amplified FedAvg may have limited improvement over FedAvg when training from scratch.

SCAFFOLD beats Amplified FedAvg. Despite a lack of theoretical guarantees under non-i.i.d. participation, SCAFFOLD outperforms Amplified FedAvg in all settings. This suggests that SCAFFOLD may have reasonable performance under some non-i.i.d. participation patterns. For the synthetic objective and CIFAR-10, SCAFFOLD is still significantly slower than Amplified SCAFFOLD.

5.3 Ablation Study

To understand how each algorithm’s performance is affected by data heterogeneity, the number of participating clients, and the number of client groups, we perform an ablation study on Fashion-MNIST. First, we fix the data similarity $s = 5\%$ and number of groups $\bar{K} = 5$ while varying the number of participating clients (S) over $\{5, 15, 20, 25\}$. Next, we fix $S = 10, \bar{K} = 5$ while varying the similarity s over $\{2.5\%, 10\%, 33\%, 100\%\}$. Lastly, we fix $s = 5\%, S = 10$ while varying the number of client groups $\bar{K} \in \{2, 4, 6, 8\}$. In each of these 12 scenarios, we evaluate all five algorithms using the same settings as detailed in Section 3. We train with three random seeds for each algorithm, and report the average results in Figure 2(right).

Amplified SCAFFOLD reaches the best solution in all settings. Similarly to Section 5.2, Amplified SCAFFOLD consistently reaches the best solution in terms of both training loss and testing accuracy. While our theoretical results provide guarantees for optimization, these experiments show that Amplified SCAFFOLD also exhibits superior generalization in a variety of settings.

Robustness to data heterogeneity. When changing from completely homogeneous data ($s = 100\%$) to extremely heterogeneous data ($s = 2.5\%$), the test accuracy of Amplified SCAFFOLD exhibits a

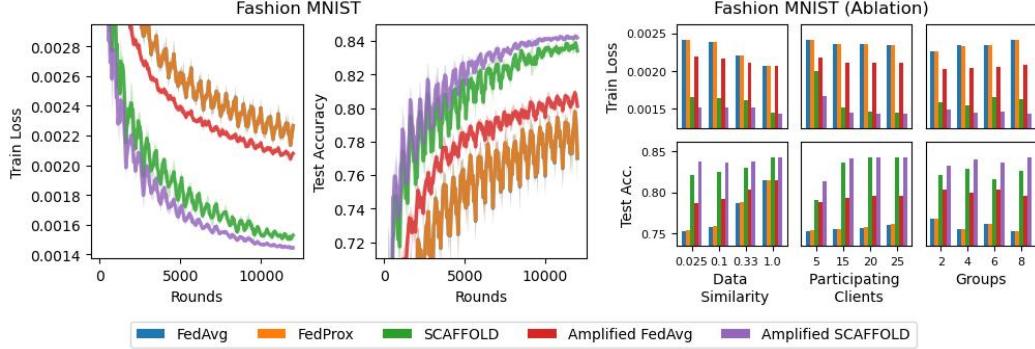


Figure 2: Results for Fashion MNIST and ablation study. Left: Amplified SCAFFOLD reaches the best solution, but SCAFFOLD is competitive. Other baselines are much slower. Right: Amplified SCAFFOLD is robust to changes in data heterogeneity, number of participating clients, and number of client groups.

very small decrease from 84.6% to 84.45%, so that our algorithm behaves nearly identically with homogeneous data as with extremely heterogeneous data. All baselines suffer a larger decrease when transitioning from homogeneous data to heterogeneous data.

Robustness to number of participating clients. The number of participating clients has a smaller effect on performance than data heterogeneity, but some degradation happens in the extreme case $S = 5$. In particular, SCAFFOLD has competitive performance with large $S \geq 15$, but its test accuracy drops off significantly compared to Amplified SCAFFOLD in the case $S = 5$.

Robustness to number of client groups. As \bar{K} increases, FedAvg and Amplified FedAvg get worse, while SCAFFOLD and Amplified SCAFFOLD maintain performance. It makes intuitive sense for an algorithm to degrade as \bar{K} increases, since a larger \bar{K} means that the participation is in some sense “further” from i.i.d. participation. Still, Amplified SCAFFOLD (and SCAFFOLD) are able to maintain performance even as \bar{K} increases. While the worst-case communication complexity of Amplified SCAFFOLD (listed in Table 1) actually increases with \bar{K} , these experiments demonstrate that in practice, Amplified SCAFFOLD can maintain performance as \bar{K} increases.

6 Conclusion

We propose Amplified SCAFFOLD, an optimization algorithm for federated learning under periodic client participation, and prove that it exhibits reduced communication cost, linear speedup, and is unaffected by data heterogeneity. We also show that Amplified SCAFFOLD experimentally outperforms baselines on standard benchmarks under non-i.i.d. client participation, and that the performance of our algorithm is robust to changes in data heterogeneity and the number of participating clients.

Limitations While our analysis covers a general class of participation patterns, it may not cover some participation patterns that appear in practice. Our framework requires that all clients have an equal chance of participation across well-defined windows of time that are known to the algorithm implementer, which may not always hold. One such practical situation is where clients may freely join or leave the federated learning process during training. Extending our algorithm and guarantees for this situation would require a reformulation of the optimization problem, and possibly additional assumptions about the participation structure. We leave such analysis for future work.

Acknowledgements

We would like to thank the anonymous reviewers for their helpful comments. This work is supported by the Institute for Digital Innovation fellowship from George Mason University, a ORIEI seed funding, an IDIA P3 fellowship from George Mason University, a Cisco Faculty Research Award, and NSF award #2436217, #2425687. Experiments were partially run on Hopper, a research computing cluster provided by the Office of Research Computing at George Mason University (URL: <https://orc.gmu.edu>).

References

- [1] Durmus Alp Emre Acar, Yue Zhao, Ramon Matas, Matthew Mattina, Paul Whatmough, and Venkatesh Saligrama. Federated learning based on dynamic regularization. In *International Conference on Learning Representations*, 2021.
- [2] Dmitrii Avdiukhin and Shiva Kasiviswanathan. Federated learning under arbitrary communication patterns. In *International Conference on Machine Learning*, pages 425–435. PMLR, 2021.
- [3] Keith Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingerman, Vladimir Ivanov, Chloe Kiddon, Jakub Konečný, Stefano Mazzocchi, Brendan McMahan, et al. Towards federated learning at scale: System design. *Proceedings of machine learning and systems*, 1: 374–388, 2019.
- [4] Wenlin Chen, Samuel Horvath, and Peter Richtarik. Optimal client sampling for federated learning. *arXiv preprint arXiv:2010.13723*, 2020.
- [5] Ziheng Cheng, Xinneng Huang, Pengfei Wu, and Kun Yuan. Momentum benefits non-iid federated learning simply and provably. In *The Twelfth International Conference on Learning Representations*, 2023.
- [6] Yae Jee Cho, Jianyu Wang, and Gauri Joshi. Client selection in federated learning: Convergence analysis and power-of-choice selection strategies. *arXiv preprint arXiv:2010.01243*, 2020.
- [7] Yae Jee Cho, Jianyu Wang, and Gauri Joshi. Towards understanding biased client selection in federated learning. In *International Conference on Artificial Intelligence and Statistics*, pages 10351–10375. PMLR, 2022.
- [8] Yae Jee Cho, Pranay Sharma, Gauri Joshi, Zheng Xu, Satyen Kale, and Tong Zhang. On the convergence of federated averaging with cyclic client participation. *arXiv preprint arXiv:2302.03109*, 2023.
- [9] Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. *Advances in Neural Information Processing Systems*, 30, 2017.
- [10] Hubert Eichner, Tomer Koren, Brendan McMahan, Nathan Srebro, and Kunal Talwar. Semi-cyclic stochastic gradient descent. In *International Conference on Machine Learning*, pages 1764–1773. PMLR, 2019.
- [11] Yann Fraboni, Richard Vidal, Laetitia Kameni, and Marco Lorenzi. Clustered sampling: Low-variance and improved representativity for clients selection in federated learning. In *International Conference on Machine Learning*, pages 3407–3416. PMLR, 2021.
- [12] Margalit R Glasgow, Honglin Yuan, and Tengyu Ma. Sharp bounds for federated averaging (local sgd) and continuous perspective. In *International Conference on Artificial Intelligence and Statistics*, pages 9050–9090. PMLR, 2022.
- [13] Michał Grudzień, Grigory Malinovsky, and Peter Richtárik. Improving accelerated federated learning with compression and importance sampling. In *Federated Learning and Analytics in Practice: Algorithms, Systems, Applications, and Opportunities*, 2023.
- [14] Xinran Gu, Kaixuan Huang, Jingzhao Zhang, and Longbo Huang. Fast federated learning in the presence of arbitrary device unavailability. *Advances in Neural Information Processing Systems*, 34:12052–12064, 2021.
- [15] Dzmitry Huba, John Nguyen, Kshitiz Malik, Ruiyu Zhu, Mike Rabbat, Ashkan Yousefpour, Carole-Jean Wu, Hongyuan Zhan, Pavel Ustinov, Harish Srinivas, et al. Papaya: Practical, private, and scalable federated learning. *Proceedings of Machine Learning and Systems*, 4: 814–832, 2022.
- [16] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *Foundations and trends® in machine learning*, 14(1–2):1–210, 2021.

[17] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning. In *International conference on machine learning*, pages 5132–5143. PMLR, 2020.

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

[19] Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*, 2016.

[20] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.

[21] Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020.

[22] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *Proceedings of Machine learning and systems*, 2:429–450, 2020.

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

[24] Xiangru Lian, Wei Zhang, Ce Zhang, and Ji Liu. Asynchronous decentralized parallel stochastic gradient descent. In *International Conference on Machine Learning*, pages 3043–3052. PMLR, 2018.

[25] Wei Yang Bryan Lim, Nguyen Cong Luong, Dinh Thai Hoang, Yutao Jiao, Ying-Chang Liang, Qiang Yang, Dusit Niyato, and Chunyan Miao. Federated learning in mobile edge networks: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 22(3):2031–2063, 2020.

[26] Grigory Malinovsky, Samuel Horváth, Konstantin Burlachenko, and Peter Richtárik. Federated learning with regularized client participation. *arXiv preprint arXiv:2302.03662*, 2023.

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

[28] H Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. Learning differentially private recurrent language models. *arXiv preprint arXiv:1710.06963*, 2017.

[29] Viraaji Mothukuri, Reza M Parizi, Seyedamin Pouriyeh, Yan Huang, Ali Dehghantanha, and Gautam Srivastava. A survey on security and privacy of federated learning. *Future Generation Computer Systems*, 115:619–640, 2021.

[30] Kumar Kshitij Patel, Lingxiao Wang, Blake E Woodworth, Brian Bullins, and Nati Srebro. Towards optimal communication complexity in distributed non-convex optimization. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 13316–13328. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper_files/paper/2022/file/56bd21259e28ebdc4d7e1503733bf421-Paper-Conference.pdf.

[31] Matthias Paulik, Matt Seigel, Henry Mason, Dominic Telaar, Joris Kluivers, Rogier van Dalen, Chi Wai Lau, Luke Carlson, Filip Granqvist, Chris Vandevelde, et al. Federated evaluation and tuning for on-device personalization: System design & applications. *arXiv preprint arXiv:2102.08503*, 2021.

[32] Sashank Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konecny, Sanjiv Kumar, and H Brendan McMahan. Adaptive federated optimization. *ICLR*, 2021.

[33] Elsa Rizk, Stefan Vlaski, and Ali H Sayed. Federated learning under importance sampling. *IEEE Transactions on Signal Processing*, 70:5381–5396, 2022.

[34] Yichen Ruan, Xiaoxi Zhang, Shu-Che Liang, and Carlee Joe-Wong. Towards flexible device participation in federated learning. In *International Conference on Artificial Intelligence and Statistics*, pages 3403–3411. PMLR, 2021.

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

[36] Alexander Tyurin, Lukang Sun, Konstantin Pavlovich Burlachenko, and Peter Richtárik. Sharper rates and flexible framework for nonconvex sgd with client and data sampling. *Transactions on Machine Learning Research*, 2022.

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

[38] Jianyu Wang, Zachary Charles, Zheng Xu, Gauri Joshi, H Brendan McMahan, Maruan Al-Shedivat, Galen Andrew, Salman Avestimehr, Katharine Daly, Deepesh Data, et al. A field guide to federated optimization. *arXiv preprint arXiv:2107.06917*, 2021.

[39] Shiqiang Wang and Mingyue Ji. A unified analysis of federated learning with arbitrary client participation. *Advances in Neural Information Processing Systems*, 35:19124–19137, 2022.

[40] Shiqiang Wang and Mingyue Ji. A lightweight method for tackling unknown participation probabilities in federated averaging. *arXiv preprint arXiv:2306.03401*, 2023.

[41] Kang Wei, Jun Li, Ming Ding, Chuan Ma, Howard H. Yang, Farhad Farokhi, Shi Jin, Tony Q. S. Quek, and H. Vincent Poor. Federated learning with differential privacy: Algorithms and performance analysis. *IEEE Transactions on Information Forensics and Security*, 15:3454–3469, 2020. doi: 10.1109/TIFS.2020.2988575.

[42] Blake Woodworth, Kumar Kshitij Patel, Sebastian Stich, Zhen Dai, Brian Bullins, Brendan McMahan, Ohad Shamir, and Nathan Srebro. Is local sgd better than minibatch sgd? In *International Conference on Machine Learning*, pages 10334–10343. PMLR, 2020.

[43] Blake E Woodworth, Kumar Kshitij Patel, and Nati Srebro. Minibatch vs local sgd for heterogeneous distributed learning. *Advances in Neural Information Processing Systems*, 33:6281–6292, 2020.

[44] Ming Xiang, Stratis Ioannidis, Edmund Yeh, Carlee Joe-Wong, and Lili Su. Efficient federated learning against heterogeneous and non-stationary client unavailability. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.

[45] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. 2017.

[46] Yikai Yan, Chaoyue Niu, Yucheng Ding, Zhenzhe Zheng, Fan Wu, Guihai Chen, Shaojie Tang, and Zhihua Wu. Distributed non-convex optimization with sublinear speedup under intermittent client availability. *arXiv preprint arXiv:2002.07399*, 2020.

[47] Haibo Yang, Minghong Fang, and Jia Liu. Achieving linear speedup with partial worker participation in non-iid federated learning. In *International Conference on Learning Representations*, 2020.

[48] Haibo Yang, Xin Zhang, Prashant Khanduri, and Jia Liu. Anarchic federated learning. In *International Conference on Machine Learning*, pages 25331–25363. PMLR, 2022.

[49] Timothy Yang, Galen Andrew, Hubert Eichner, Haicheng Sun, Wei Li, Nicholas Kong, Daniel Ramage, and Françoise Beaufays. Applied federated learning: Improving google keyboard query suggestions. *arXiv preprint arXiv:1812.02903*, 2018.

[50] Hao Yu, Rong Jin, and Sen Yang. On the linear speedup analysis of communication efficient momentum SGD for distributed non-convex optimization. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, pages 7184–7193, 2019.

- [51] Hao Yu, Sen Yang, and Shenghuo Zhu. Parallel restarted sgd with faster convergence and less communication: Demystifying why model averaging works for deep learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5693–5700, 2019.
- [52] Honglin Yuan and Tengyu Ma. Federated accelerated stochastic gradient descent. *Advances in Neural Information Processing Systems*, 33:5332–5344, 2020.
- [53] Chen Zhang, Yu Xie, Hang Bai, Bin Yu, Weihong Li, and Yuan Gao. A survey on federated learning. *Knowledge-Based Systems*, 216:106775, 2021.
- [54] Chen Zhu, Zheng Xu, Mingqing Chen, Jakub Konečný, Andrew Hard, and Tom Goldstein. Diurnal or nocturnal? federated learning of multi-branch networks from periodically shifting distributions. In *International Conference on Learning Representations*, 2021.
- [55] Hangyu Zhu, Jinjin Xu, Shiqing Liu, and Yaochu Jin. Federated learning on non-iid data: A survey. *Neurocomputing*, 465:371–390, 2021.

Contents

1 Introduction	1
2 Related Work	3
3 Problem Setup	3
3.1 Participation Framework	4
3.2 Specific Participation Patterns	4
4 Algorithm and Analysis	5
4.1 Algorithm Overview	5
4.2 Main Results	6
4.3 Application to Participation Patterns	6
4.4 Proof Sketch	7
5 Experiments	8
5.1 Setup	8
5.2 Main Results	9
5.3 Ablation Study	9
6 Conclusion	10
A Proof of Theorem 1	16
A.1 Preliminary Definitions	16
A.2 Proofs	17
B Proofs for Specific Participation Patterns	45
C Complexity of Amplified FedAvg	48
D Experiment Details	49
D.1 Client Sampling Parameters	49
D.2 Heterogeneity Protocol	50
D.3 Hyperparameter Tuning	50
D.4 Synthetic Objective	50
D.5 CNN Architecture	51
E Additional Experimental Results	51
E.1 Additional Baselines	51
E.2 CIFAR-10 with Stochastic Client Availability	52
F Extended Comparison with Baselines	52

A Proof of Theorem 1

A.1 Preliminary Definitions

Let \mathcal{G}_r denote the filtration generated by $\{\xi_{r,k}^i : 0 \leq k \leq I-1, i \in [N]\}$, that is, by the randomness in the stochastic gradients during round r . Also, let \mathcal{Q}_{r_0} denote the filtration generated by $\{q_r^i : r_0 \leq r < r_0 + P, i \in [N]\}$, that is, the randomness in client sampling between rounds r_0 and $r_0 + P - 1$ (inclusive). Also, let \mathcal{G} denote the filtration generated by $\mathcal{G}_0 \cup \mathcal{G}_1 \cup \dots \cup \mathcal{G}_{R-1}$ and \mathcal{Q} denote the filtration generated by $\mathcal{Q}_0 \cup \mathcal{Q}_P \cup \dots \cup \mathcal{Q}_{R-P}$. Similarly, let $\mathcal{G}_{:r_0}$ denote the filtration generated by $\mathcal{G}_0 \cup \mathcal{G}_1 \cup \dots \cup \mathcal{G}_{r_0-1}$ and $\mathcal{Q}_{:r_0}$ denote the filtration generated by $\mathcal{Q}_0 \cup \mathcal{Q}_P \cup \dots \cup \mathcal{Q}_{r_0-P}$. Lastly, denote $\mathbb{E}_{r_0}[\cdot] = \mathbb{E}[\cdot | \mathcal{Q}_{:r_0}, \mathcal{G}_{:r_0}]$.

In order to analyze the control variate errors (i.e., $\|\nabla f_i(\bar{x}_{r_0}) - \mathbf{G}_{r_0}^i\|$), we introduce notation to refer to the iterates whose stochastic gradients were used to construct $\mathbf{G}_{r_0}^i$. These iterates are exactly the iterates during the most recent window before r_0 in which client i was sampled, i.e., where $\frac{1}{P} \sum_{s=m}^{(m+1)P-1} q_s^i > 0$. For each $r \in \{r_0, r_0 + 1, \dots, r_0 + P - 1\}$, $k \in \{0, \dots, I-1\}$, and $i \in [N]$, let

$$t_i(r) = \begin{cases} r & \text{if } \bar{q}_{r_0}^i > 0 \\ t_i(r-P) & \text{if } \bar{q}_{r_0}^i = 0 \end{cases}$$

with the initialization $t_i(r) = 0$ for all $r \in \{-P, \dots, -1\}$. Then denote

$$\begin{aligned} \mathbf{y}_{r,k}^i &= \mathbf{x}_{t_i(r),k}^i \\ z_r^i &= q_{t_i(r)}^i \\ \zeta_{r,k}^i &= \xi_{t_i(r),k}^i. \end{aligned}$$

Also, denote $\bar{z}_{r_0}^i = \frac{1}{P} \sum_{r=r_0}^{r_0+P-1} z_r^i$. Then we can rewrite

$$\mathbf{G}_{r_0}^i = \frac{1}{P \bar{z}_{r_0-P}^i I} \sum_{s=r_0-P}^{r_0-1} z_s^i \sum_{k=0}^{I-1} \nabla F_i(\mathbf{y}_{s,k}^i; \zeta_{s,k}^i).$$

Define

$$\begin{aligned} \bar{x}_{r,k} &= \sum_{i=1}^N q_r^i \mathbf{x}_{r,k}^i \\ D_{r,k} &= \sum_{i=1}^N q_r^i \mathbb{E} \left[\|\mathbf{x}_{r,k}^i - \bar{x}_{r,k}\|^2 \middle| \mathcal{Q} \right] \\ M_{r,k} &= \mathbb{E} \left[\|\bar{x}_{r,k} - \bar{x}_{r_0}\|^2 \middle| \mathcal{Q} \right] \\ S_{r_0}^i &= \frac{1}{IP} \sum_{s=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} \frac{z_s^i}{\bar{z}_{r_0-P}^i} \mathbb{E} \left[\|\mathbf{y}_{s,k}^i - \bar{x}_{r_0}\|^2 \middle| \mathcal{Q} \right] \\ w_{r_0}^i &= \frac{1}{N} \sum_{j=1}^N \frac{\mathbf{1}\{\bar{q}_{r_0}^j > 0\}}{P \bar{q}_{r_0}^j} \sum_{s=r_0}^{r_0+P-1} q_s^i q_s^j \\ v_{r_0}^i &= \bar{q}_{r_0}^i - \frac{1}{N} \\ \Lambda_{r_0}^i &= \frac{\frac{1}{P} \sum_{s=r_0-P}^{r_0-1} (z_s^i)^2}{\left(\frac{1}{P} \sum_{s=r_0-P}^{r_0-1} z_s^i \right)^2}. \end{aligned}$$

As discussed in the main body, $\bar{x}_{r,k}$ is a weighted average over local client models, weighted according to client participation. $D_{r,k}$ and $M_{r,k}$ are the drift terms described in the proof sketch of Section 4.4. In the proof sketch, we also informally discuss the control variate error $C_{r_0}^i = \|\nabla f_i(\bar{x}_{r_0}) - \mathbf{G}_{r_0}^i\|^2$.

This error is closely related to the term $S_{r_0}^i$ defined above, since $\mathbb{E} \left[\left\| \nabla f_i(\bar{\mathbf{x}}_{r_0}) - \mathbb{E}[\mathbf{G}_{r_0}^i | \mathcal{Q}] \right\|^2 \middle| \mathcal{Q} \right] = L^2 S_{r_0}^i$.

The quantities $w_{r_0}^i, v_{r_0}^i$ and $\Lambda_{r_0}^i$ arise in our proof from the use of control variates under non-i.i.d. client participation. As discussed in the proof sketch (Section 4.4), bounding the errors introduced by control variates involves a non-uniform average of non-uniform averages of error terms. This nested non-uniform average can be bounded by a uniform average as long as the term $w_{r_0}^i$ is not too large, which is exactly the requirement stated in Equation 1. On the other hand, the term $\Lambda_{r_0}^i$ arises while bounding the variance of the update $\|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|$, and it represents the sample size for which the correction $\mathbf{G}_{r_0}^i$ is computed. When the client sampling distribution is closer to uniform, then each client may participate frequently during each window of P rounds, and the effective sample size used to compute each control variate is large. In this case, the variance of the update $\|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|$ will be smaller, since larger sample size implies smaller variance. For general client sampling distributions, $\Lambda_{r_0}^i$ describes the reduction of variance of one component of the update (i.e. $\mathbf{G}_{r_0}^i$) due to the sampling of stochastic gradients during the previous window of P rounds. In our analysis, the variance of the update $\|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|$ can be bounded in terms of the variance of $\mathbf{G}_{r_0}^i$, which in turn depends on $\Lambda_{r_0}^i$. The condition in Equation 2 essentially enforces that clients participate sufficiently uniformly, so that the variance of $\mathbf{G}_{r_0}^i$ can be bounded. We note that both Equation 1 and Equation 2 are satisfied by i.i.d. participation, regularized participation, and cyclic participation.

Finally, for ease of exposition, one clarification must be made about the conditional expectation $\mathbb{E}[\cdot | \mathcal{Q}]$ as it appears in, for example, $M_{r,k}$. The filtration \mathcal{Q} contains randomness from every round from 0 to $R-1$, but the expression of which we are taking expectation may only depend on randomness for a smaller subset rounds. For example, in $M_{r,k}$, the expression $\|\bar{\mathbf{x}}_{r,k} - \bar{\mathbf{x}}_{r_0}\|^2$ only depends on the randomness up to round $r \leq r_0 + P$, so that

$$\mathbb{E} \left[\|\bar{\mathbf{x}}_{r,k} - \bar{\mathbf{x}}_{r_0}\|^2 \middle| \mathcal{Q} \right] = \mathbb{E} \left[\|\bar{\mathbf{x}}_{r,k} - \bar{\mathbf{x}}_{r_0}\|^2 \middle| \mathcal{Q}_{:r_0+P} \right].$$

In several places throughout the proof, we will replace $\mathbb{E}[\cdot | \mathcal{Q}]$ by $\mathbb{E}[\cdot | \mathcal{Q}_{:r_0+P}]$ in similar situations.

A.2 Proofs

Lemma 1. *If $\eta \leq \frac{1}{60LIP}$ and $\mathbb{E}_{r_0}[\bar{q}_{r_0}^i] = \frac{1}{N}$, then*

$$\begin{aligned} D_{r,k} + M_{r,k} &\leq 108\eta^2 I^2 P^2 \mathbb{E} \left[\left\| \nabla f(\bar{\mathbf{x}}_{r_0}) \right\|^2 \middle| \mathcal{Q} \right] + (75\eta^2 I + 65\eta^2 IP\rho^2) \sigma^2 \\ &\quad + 109\eta^2 L^2 I^2 P^2 \sum_{i=1}^N \left(\bar{q}_{r_0}^i + \frac{1}{N} \right) S_{r_0}^i + 36\eta^2 L^2 I^2 \sum_{i=1}^N q_r^i S_{r_0}^i, \end{aligned} \quad (3)$$

and

$$\begin{aligned} \mathbb{E} \left[\sum_{r=r_0}^{r_0+P-1} \sum_{k=0}^{I-1} (D_{r,k} + M_{r,k}) \right] &\leq 108\eta^2 I^3 P^3 \mathbb{E} \left[\left\| \nabla f(\bar{\mathbf{x}}_{r_0}) \right\|^2 \right] + (75\eta^2 I^2 P + 65\eta^2 I^2 P^2 \rho^2) \sigma^2 \\ &\quad + 254\eta^2 L^2 I^3 P^3 \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0}^i]. \end{aligned} \quad (4)$$

Proof. Denote

$$\begin{aligned} \bar{\mathbf{G}}_{r_0}^i &= \frac{1}{P\bar{z}_{r_0-P}^i} \sum_{s=r_0-P}^{r_0-1} z_s^i \sum_{k=0}^{I-1} \nabla f_i(\mathbf{y}_{s,k}^i) \\ \bar{\mathbf{G}}_{r_0} &= \frac{1}{N} \sum_{i=1}^N \bar{\mathbf{G}}_{r_0}^i. \end{aligned}$$

Also, denote $\mathbf{g}_{r,k}^i = \nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) - \mathbf{G}_{r_0}^i + \mathbf{G}_{r_0}$ and $\bar{\mathbf{g}}_{r,k}^i = \nabla f_i(\mathbf{x}_{r,k}^i) - \bar{\mathbf{G}}_{r_0}^i + \bar{\mathbf{G}}_{r_0}$. $\bar{\mathbf{G}}_{r_0}^i$ is the analogue of $\bar{\mathbf{G}}_{r_0}$ that depends on deterministic gradients $\nabla f_i(\mathbf{y}_{s,k}^i)$ instead of stochastic gradients.

Variance of updates We first compute the errors $\mathbb{E} \left[\left\| \mathbf{g}_{r,k}^i - \bar{\mathbf{g}}_{r,k}^i \right\|^2 \middle| \mathcal{Q} \right]$ and $\mathbb{E} \left[\left\| \sum_{i=1}^N q_r^i \left(\mathbf{g}_{r,k}^i - \bar{\mathbf{g}}_{r,k}^i \right) \right\|^2 \middle| \mathcal{Q} \right]$, which will be needed in several places.

$$\begin{aligned}
\mathbb{E} \left[\left\| \mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i \right\|^2 \middle| \mathcal{Q} \right] &= \mathbb{E} \left[\left\| \frac{1}{P \bar{z}_{r_0-P}^i I} \sum_{s=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} z_s^i (\nabla F_i(\mathbf{y}_{s,k}^i; \zeta_{s,k}^i) - \nabla f_i(\mathbf{y}_{s,k}^i)) \right\|^2 \middle| \mathcal{Q} \right] \\
&\stackrel{(i)}{\leq} \frac{1}{P \bar{z}_{r_0-P}^i I} \sum_{s=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} z_s^i \mathbb{E} \left[\left\| \nabla F_i(\mathbf{y}_{s,k}^i; \zeta_{s,k}^i) - \nabla f_i(\mathbf{y}_{s,k}^i) \right\|^2 \middle| \mathcal{Q} \right] \\
&\leq \frac{1}{P \bar{z}_{r_0-P}^i I} \sum_{s=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} z_s^i \sigma^2 \\
&= \sigma^2,
\end{aligned} \tag{5}$$

where (i) uses Jensen's inequality. Therefore

$$\begin{aligned}
\mathbb{E} \left[\left\| \mathbf{g}_{r,k}^i - \bar{\mathbf{g}}_{r,k}^i \right\|^2 \middle| \mathcal{Q} \right] &= \mathbb{E} \left[\left\| (\nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) - \nabla f_i(\mathbf{x}_{r,k}^i)) - (\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i) + (\mathbf{G}_{r_0} - \bar{\mathbf{G}}_{r_0}) \right\|^2 \middle| \mathcal{Q} \right] \\
&\leq 3\mathbb{E} \left[\left\| \nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) - \nabla f_i(\mathbf{x}_{r,k}^i) \right\|^2 \middle| \mathcal{Q} \right] + 3\mathbb{E} \left[\left\| \mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i \right\|^2 \middle| \mathcal{Q} \right] \\
&\quad + 3\mathbb{E} \left[\left\| \mathbf{G}_{r_0} - \bar{\mathbf{G}}_{r_0} \right\|^2 \middle| \mathcal{Q} \right] \\
&= 3\mathbb{E} \left[\left\| \nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) - \nabla f_i(\mathbf{x}_{r,k}^i) \right\|^2 \middle| \mathcal{Q} \right] + 3\mathbb{E} \left[\left\| \mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i \right\|^2 \middle| \mathcal{Q} \right] \\
&\quad + 3\mathbb{E} \left[\left\| \frac{1}{N} \sum_{j=1}^N (\mathbf{G}_{r_0}^j - \bar{\mathbf{G}}_{r_0}^j) \right\|^2 \middle| \mathcal{Q} \right] \\
&= 3\mathbb{E} \left[\left\| \nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) - \nabla f_i(\mathbf{x}_{r,k}^i) \right\|^2 \middle| \mathcal{Q} \right] + 3\mathbb{E} \left[\left\| \mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i \right\|^2 \middle| \mathcal{Q} \right] \\
&\quad + 3 \frac{1}{N} \sum_{j=1}^N \mathbb{E} \left[\left\| \mathbf{G}_{r_0}^j - \bar{\mathbf{G}}_{r_0}^j \right\|^2 \middle| \mathcal{Q} \right] \\
&\leq 9\sigma^2,
\end{aligned} \tag{6}$$

where the last line uses Equation 5. Also

$$\begin{aligned}
& \mathbb{E} \left[\left\| \sum_{i=1}^N q_r^i (\mathbf{g}_{r,k}^i - \bar{\mathbf{g}}_{r,k}^i) \right\|^2 \middle| \mathcal{Q} \right] \\
&= \mathbb{E} \left[\left\| \sum_{i=1}^N q_r^i ((\nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) - \nabla f_i(\mathbf{x}_{r,k}^i)) - (\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i) + (\mathbf{G}_{r_0} - \bar{\mathbf{G}}_{r_0})) \right\|^2 \middle| \mathcal{Q} \right] \\
&\leq 3\mathbb{E} \left[\left\| \sum_{i=1}^N q_r^i (\nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) - \nabla f_i(\mathbf{x}_{r,k}^i)) \right\|^2 \middle| \mathcal{Q} \right] + 3\mathbb{E} \left[\left\| \sum_{i=1}^N q_r^i (\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i) \right\|^2 \middle| \mathcal{Q} \right] \\
&\quad + 3\mathbb{E} \left[\left\| \sum_{i=1}^N q_r^i (\mathbf{G}_{r_0} - \bar{\mathbf{G}}_{r_0}) \right\|^2 \middle| \mathcal{Q} \right] \\
&= 3\mathbb{E} \left[\left\| \sum_{i=1}^N q_r^i (\nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) - \nabla f_i(\mathbf{x}_{r,k}^i)) \right\|^2 \middle| \mathcal{Q} \right] + 3\mathbb{E} \left[\left\| \sum_{i=1}^N q_r^i (\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i) \right\|^2 \middle| \mathcal{Q} \right] \\
&\quad + 3\mathbb{E} \left[\left\| \sum_{i=1}^N \frac{1}{N} (\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i) \right\|^2 \middle| \mathcal{Q} \right] \\
&\stackrel{(i)}{=} 3 \sum_{i=1}^N \mathbb{E} \left[\|q_r^i (\nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) - \nabla f_i(\mathbf{x}_{r,k}^i))\|^2 \middle| \mathcal{Q} \right] + 3 \sum_{i=1}^N \mathbb{E} \left[\|q_r^i (\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i)\|^2 \middle| \mathcal{Q} \right] \\
&\quad + 3 \sum_{i=1}^N \mathbb{E} \left[\left\| \frac{1}{N} (\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i) \right\|^2 \middle| \mathcal{Q} \right] \\
&= 3 \sum_{i=1}^N (q_r^i)^2 \mathbb{E} \left[\|\nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) - \nabla f_i(\mathbf{x}_{r,k}^i)\|^2 \middle| \mathcal{Q} \right] + 3 \sum_{i=1}^N (q_r^i)^2 \mathbb{E} \left[\|\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i\|^2 \middle| \mathcal{Q} \right] \\
&\quad + 3 \frac{1}{N^2} \sum_{i=1}^N \mathbb{E} \left[\|\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i\|^2 \middle| \mathcal{Q} \right] \\
&\leq 6\sigma^2 \sum_{i=1}^N (q_r^i)^2 + 3 \frac{\sigma^2}{N} \\
&\leq \sigma^2 \left(6\rho^2 + 3 \frac{1}{N} \right) \\
&\stackrel{(ii)}{\leq} 9\rho^2\sigma^2,
\end{aligned} \tag{7}$$

where (i) uses the fact that stochastic gradient noise is independent across each client, and (ii) uses $\rho^2 \geq \frac{1}{N}$.

One-step recursive bound for $D_{r,k}$ For any $k \geq 0$,

$$\begin{aligned}
D_{r,k+1} &= \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \mathbf{x}_{r,k+1}^i - \sum_{j=1}^N q_r^j \mathbf{x}_{r,k+1}^j \right\|^2 \middle| \mathcal{Q} \right] \\
&= \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \mathbf{x}_{r,k}^i - \eta \mathbf{g}_{r,k}^i - \sum_{j=1}^N q_r^j (\mathbf{x}_{r,k}^j - \eta \mathbf{g}_{r,k}^j) \right\|^2 \middle| \mathcal{Q} \right] \\
&= \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \mathbf{x}_{r,k}^i - \sum_{j=1}^N q_r^j \mathbf{x}_{r,k}^j - \eta \left(\mathbf{g}_{r,k}^i - \sum_{j=1}^N q_r^j \mathbf{g}_{r,k}^j \right) \right\|^2 \middle| \mathcal{Q} \right] \\
&= \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \mathbf{x}_{r,k}^i - \sum_{j=1}^N q_r^j \mathbf{x}_{r,k}^j - \eta \left(\bar{\mathbf{g}}_{r,k}^i - \sum_{j=1}^N q_r^j \bar{\mathbf{g}}_{r,k}^j \right) \right\|^2 \middle| \mathcal{Q} \right] \\
&\quad + \eta^2 \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| (\mathbf{g}_{r,k}^i - \bar{\mathbf{g}}_{r,k}^i) + \sum_{j=1}^N q_r^j (\mathbf{g}_{r,k}^j - \bar{\mathbf{g}}_{r,k}^j) \right\|^2 \middle| \mathcal{Q} \right] \\
&\stackrel{(i)}{=} \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \mathbf{x}_{r,k}^i - \sum_{j=1}^N q_r^j \mathbf{x}_{r,k}^j - \eta \left(\bar{\mathbf{g}}_{r,k}^i - \sum_{j=1}^N q_r^j \bar{\mathbf{g}}_{r,k}^j \right) \right\|^2 \middle| \mathcal{Q} \right] + 36\eta^2\sigma^2 \\
&\stackrel{(ii)}{\leq} \left(1 + \frac{1}{\lambda_1} \right) \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \mathbf{x}_{r,k}^i - \sum_{j=1}^N q_r^j \mathbf{x}_{r,k}^j \right\|^2 \middle| \mathcal{Q} \right] \\
&\quad + \eta^2 (1 + \lambda_1) \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \bar{\mathbf{g}}_{r,k}^i - \sum_{j=1}^N q_r^j \bar{\mathbf{g}}_{r,k}^j \right\|^2 \middle| \mathcal{Q} \right] + 36\eta^2\sigma^2 \\
&= \left(1 + \frac{1}{\lambda_1} \right) D_{r,k} + \eta^2 (1 + \lambda_1) \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \bar{\mathbf{g}}_{r,k}^i - \sum_{j=1}^N q_r^j \bar{\mathbf{g}}_{r,k}^j \right\|^2 \middle| \mathcal{Q} \right] + 36\eta^2\sigma^2, \quad (8)
\end{aligned}$$

where (i) uses

$$\begin{aligned}
&\mathbb{E} \left[\left\| (\mathbf{g}_{r,k}^i - \bar{\mathbf{g}}_{r,k}^i) + \sum_{j=1}^N q_r^j (\mathbf{g}_{r,k}^j - \bar{\mathbf{g}}_{r,k}^j) \right\|^2 \middle| \mathcal{Q} \right] \\
&\leq 2\mathbb{E} \left[\left\| \mathbf{g}_{r,k}^i - \bar{\mathbf{g}}_{r,k}^i \right\|^2 \middle| \mathcal{Q} \right] + 2\mathbb{E} \left[\left\| \sum_{j=1}^N q_r^j (\mathbf{g}_{r,k}^j - \bar{\mathbf{g}}_{r,k}^j) \right\|^2 \middle| \mathcal{Q} \right] \\
&\leq 2\mathbb{E} \left[\left\| \mathbf{g}_{r,k}^i - \bar{\mathbf{g}}_{r,k}^i \right\|^2 \middle| \mathcal{Q} \right] + 2 \sum_{j=1}^N q_r^j \mathbb{E} \left[\left\| \mathbf{g}_{r,k}^j - \bar{\mathbf{g}}_{r,k}^j \right\|^2 \middle| \mathcal{Q} \right] \\
&\stackrel{(a)}{\leq} 36\sigma^2,
\end{aligned}$$

(a) uses Equation 6, and (ii) uses Young's inequality. Focusing on the second term of Equation 8:

$$\begin{aligned}
& \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \bar{\mathbf{g}}_{r,k}^i - \sum_{j=1}^N q_r^j \bar{\mathbf{g}}_{r,k}^j \right\|^2 \middle| \mathcal{Q} \right] = \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \sum_{j=1}^N q_r^j (\bar{\mathbf{g}}_{r,k}^i - \bar{\mathbf{g}}_{r,k}^j) \right\|^2 \middle| \mathcal{Q} \right] \\
& \stackrel{(i)}{\leq} \sum_{i=1}^N q_r^i \sum_{j=1}^N q_r^j \mathbb{E} \left[\left\| \bar{\mathbf{g}}_{r,k}^i - \bar{\mathbf{g}}_{r,k}^j \right\|^2 \middle| \mathcal{Q} \right] \\
& \leq \sum_{i=1}^N q_r^i \sum_{j=1}^N q_r^j \mathbb{E} \left[\left\| \nabla f_i(\mathbf{x}_{r,k}^i) - \bar{\mathbf{G}}_{r_0}^i - \nabla f_j(\mathbf{x}_{r,k}^j) + \bar{\mathbf{G}}_{r_0}^j \right\|^2 \middle| \mathcal{Q} \right] \\
& \leq \sum_{i=1}^N q_r^i \sum_{j=1}^N q_r^j \left(2\mathbb{E} \left[\left\| \nabla f_i(\mathbf{x}_{r,k}^i) - \bar{\mathbf{G}}_{r_0}^i \right\|^2 \middle| \mathcal{Q} \right] + 2\mathbb{E} \left[\left\| \nabla f_j(\mathbf{x}_{r,k}^j) - \bar{\mathbf{G}}_{r_0}^j \right\|^2 \middle| \mathcal{Q} \right] \right) \\
& = 4 \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \nabla f_i(\mathbf{x}_{r,k}^i) - \bar{\mathbf{G}}_{r_0}^i \right\|^2 \middle| \mathcal{Q} \right],
\end{aligned}$$

where (i) uses Jensen's inequality. Using the decomposition

$$\nabla f_i(\mathbf{x}_{r,k}^i) - \bar{\mathbf{G}}_{r_0}^i = (\nabla f_i(\mathbf{x}_{r,k}^i) - \nabla f_i(\bar{\mathbf{x}}_{r,k})) + (\nabla f_i(\bar{\mathbf{x}}_{r,k}) - \nabla f_i(\bar{\mathbf{x}}_{r_0})) + (\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}^i),$$

we have

$$\begin{aligned}
& \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \bar{\mathbf{g}}_{r,k}^i - \sum_{j=1}^N q_r^j \bar{\mathbf{g}}_{r,k}^j \right\|^2 \middle| \mathcal{Q} \right] \\
& \leq 4 \sum_{i=1}^N q_r^i \left(3\mathbb{E} \left[\left\| \nabla f_i(\mathbf{x}_{r,k}^i) - \nabla f_i(\bar{\mathbf{x}}_{r,k}) \right\|^2 \middle| \mathcal{Q} \right] + 3\mathbb{E} \left[\left\| \nabla f_i(\bar{\mathbf{x}}_{r,k}) - \nabla f_i(\bar{\mathbf{x}}_{r_0}) \right\|^2 \middle| \mathcal{Q} \right] \right. \\
& \quad \left. + 3\mathbb{E} \left[\left\| \nabla f_i(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}^i \right\|^2 \middle| \mathcal{Q} \right] \right) \\
& \leq 12L^2 \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \mathbf{x}_{r,k}^i - \bar{\mathbf{x}}_{r,k} \right\|^2 \middle| \mathcal{Q} \right] + 12L^2 \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \bar{\mathbf{x}}_{r,k} - \bar{\mathbf{x}}_{r_0} \right\|^2 \middle| \mathcal{Q} \right] \\
& \quad + 12 \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \nabla f_i(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}^i \right\|^2 \middle| \mathcal{Q} \right] \\
& = 12L^2 D_{r,k} + 12L^2 M_{r,k} + 12 \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \nabla f_i(\bar{\mathbf{x}}_{r_0}) - \frac{1}{P\bar{z}_{r_0-P}^i} \sum_{s=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} z_s^i \nabla f_i(\mathbf{y}_{s,k}^i) \right\|^2 \middle| \mathcal{Q} \right] \\
& = 12L^2 D_{r,k} + 12L^2 M_{r,k} + 12 \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \frac{1}{P\bar{z}_{r_0-P}^i} \sum_{s=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} z_s^i (\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \nabla f_i(\mathbf{y}_{s,k}^i)) \right\|^2 \middle| \mathcal{Q} \right] \\
& \leq 12L^2 D_{r,k} + 12L^2 M_{r,k} + 12 \sum_{i=1}^N q_r^i \left(\frac{1}{P\bar{z}_{r_0-P}^i} \sum_{s=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} z_s^i \mathbb{E} \left[\left\| \nabla f_i(\bar{\mathbf{x}}_{r_0}) - \nabla f_i(\mathbf{y}_{s,k}^i) \right\|^2 \middle| \mathcal{Q} \right] \right) \\
& \leq 12L^2 D_{r,k} + 12L^2 M_{r,k} + 12L^2 \sum_{i=1}^N q_r^i \left(\frac{1}{P\bar{z}_{r_0-P}^i} \sum_{s=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} z_s^i \mathbb{E} \left[\left\| \bar{\mathbf{x}}_{r_0} - \mathbf{y}_{s,k}^i \right\|^2 \middle| \mathcal{Q} \right] \right) \\
& \leq 12L^2 D_{r,k} + 12L^2 M_{r,k} + 12L^2 \sum_{i=1}^N q_r^i S_{r_0}^i.
\end{aligned}$$

Plugging back into Equation 8,

$$\begin{aligned}
D_{r,k+1} &\leq \left(1 + \frac{1}{\lambda_1}\right) D_{r,k} + 12\eta^2 L^2(1 + \lambda_1) \left(D_{r,k} + M_{r,k} + \sum_{i=1}^N q_r^i S_{r_0}^i\right) + 36\eta^2 \sigma^2 \\
&\leq \left(1 + \frac{1}{\lambda_1} + 12\eta^2 L^2(1 + \lambda_1)\right) D_{r,k} + 12\eta^2 L^2(1 + \lambda_1) \left(M_{r,k} + \sum_{i=1}^N q_r^i S_{r_0}^i\right) + 36\eta^2 \sigma^2.
\end{aligned}$$

Now choose $\lambda_1 = 2I$, so that $1 \leq \frac{\lambda_1}{2}$ and

$$12\eta^2 L^2(1 + \lambda_1) \leq 18\eta^2 L^2 \lambda_1 \leq 36\eta^2 L^2 I \leq \frac{1}{2I},$$

where we used the condition $\eta \leq \frac{1}{60LIP}$. This yields the following recursive bound on $D_{r,k+1}$:

$$D_{r,k+1} \leq \left(1 + \frac{1}{I}\right) D_{r,k} + 18\eta^2 L^2 I M_{r,k} + 18\eta^2 L^2 I \sum_{i=1}^N q_r^i S_{r_0}^i + 36\eta^2 \sigma^2. \quad (9)$$

One-step recursive bound for $M_{r,k}$ For any $k \geq 0$,

$$\begin{aligned}
M_{r,k+1} &= \mathbb{E} \left[\left\| \sum_{i=1}^N q_r^i \mathbf{x}_{r,k+1}^i - \bar{\mathbf{x}}_{r_0} \right\|^2 \middle| \mathcal{Q} \right] \\
&= \mathbb{E} \left[\left\| \sum_{i=1}^N q_r^i \mathbf{x}_{r,k}^i - \bar{\mathbf{x}}_{r_0} - \eta \sum_{i=1}^N q_r^i \mathbf{g}_{r,k}^i \right\|^2 \middle| \mathcal{Q} \right] \\
&\leq \mathbb{E} \left[\left\| \sum_{i=1}^N q_r^i \mathbf{x}_{r,k}^i - \bar{\mathbf{x}}_{r_0} - \eta \sum_{i=1}^N q_r^i \bar{\mathbf{g}}_{r,k}^i \right\|^2 \middle| \mathcal{Q} \right] + \eta^2 \mathbb{E} \left[\left\| \sum_{i=1}^N q_r^i (\mathbf{g}_{r,k}^i - \bar{\mathbf{g}}_{r,k}^i) \right\|^2 \middle| \mathcal{Q} \right] \\
&\stackrel{(i)}{\leq} \mathbb{E} \left[\left\| \sum_{i=1}^N q_r^i \mathbf{x}_{r,k}^i - \bar{\mathbf{x}}_{r_0} - \eta \sum_{i=1}^N q_r^i \bar{\mathbf{g}}_{r,k}^i \right\|^2 \middle| \mathcal{Q} \right] + 9\eta^2 \rho^2 \sigma^2 \\
&\stackrel{(ii)}{\leq} \left(1 + \frac{1}{\lambda_2}\right) \mathbb{E} \left[\left\| \sum_{i=1}^N q_r^i \mathbf{x}_{r,k}^i - \bar{\mathbf{x}}_{r_0} \right\|^2 \middle| \mathcal{Q} \right] + \eta^2 (1 + \lambda_2) \mathbb{E} \left[\left\| \sum_{i=1}^N q_r^i \bar{\mathbf{g}}_{r,k}^i \right\|^2 \middle| \mathcal{Q} \right] + 9\eta^2 \rho^2 \sigma^2 \\
&\leq \left(1 + \frac{1}{\lambda_2}\right) M_{r,k} + \eta^2 (1 + \lambda_2) \mathbb{E} \left[\left\| \sum_{i=1}^N q_r^i \bar{\mathbf{g}}_{r,k}^i \right\|^2 \middle| \mathcal{Q} \right] + 9\eta^2 \rho^2 \sigma^2 \\
&\leq \left(1 + \frac{1}{\lambda_2}\right) M_{r,k} + \eta^2 (1 + \lambda_2) \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \bar{\mathbf{g}}_{r,k}^i \right\|^2 \middle| \mathcal{Q} \right] + 9\eta^2 \rho^2 \sigma^2,
\end{aligned} \quad (10)$$

where (i) uses Equation 7 and (ii) uses Young's inequality. Focusing on the second term in Equation 10

$$\begin{aligned}
& \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \bar{\mathbf{g}}_{r,k}^i \right\|^2 \middle| \mathcal{Q} \right] \\
&= \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \nabla f_i(\mathbf{x}_{r,k}^i) - \bar{\mathbf{G}}_{r_0}^i + \bar{\mathbf{G}}_{r_0} \right\|^2 \middle| \mathcal{Q} \right] \\
&= \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| (\nabla f_i(\mathbf{x}_{r,k}^i) - \nabla f_i(\bar{\mathbf{x}}_{r,k})) + (\nabla f_i(\bar{\mathbf{x}}_{r,k}) - \nabla f_i(\bar{\mathbf{x}}_{r_0})) \right. \right. \\
&\quad \left. \left. + (\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}^i) + (\bar{\mathbf{G}}_{r_0} - \nabla f(\bar{\mathbf{x}}_{r_0})) + \nabla f(\bar{\mathbf{x}}_{r_0}) \right\|^2 \middle| \mathcal{Q} \right] \\
&\leq 5 \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \nabla f_i(\mathbf{x}_{r,k}^i) - \nabla f_i(\bar{\mathbf{x}}_{r,k}) \right\|^2 \middle| \mathcal{Q} \right] + 5 \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \nabla f_i(\bar{\mathbf{x}}_{r,k}) - \nabla f_i(\bar{\mathbf{x}}_{r_0}) \right\|^2 \middle| \mathcal{Q} \right] \\
&\quad + 5 \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \nabla f_i(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}^i \right\|^2 \middle| \mathcal{Q} \right] + 5 \mathbb{E} \left[\left\| \bar{\mathbf{G}}_{r_0} - \nabla f(\bar{\mathbf{x}}_{r_0}) \right\|^2 \middle| \mathcal{Q} \right] + 5 \mathbb{E} \left[\left\| \nabla f(\bar{\mathbf{x}}_{r_0}) \right\|^2 \middle| \mathcal{Q} \right] \\
&\leq 5L^2 \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \mathbf{x}_{r,k}^i - \bar{\mathbf{x}}_{r,k} \right\|^2 \middle| \mathcal{Q} \right] + 5L^2 \sum_{i=1}^N q_r^i \mathbb{E} \left[\left\| \bar{\mathbf{x}}_{r,k} - \bar{\mathbf{x}}_{r_0} \right\|^2 \middle| \mathcal{Q} \right] \\
&\quad + 5 \sum_{i=1}^N \left(q_r^i + \frac{1}{N} \right) \mathbb{E} \left[\left\| \nabla f_i(\bar{\mathbf{x}}_{r_0}) - \frac{1}{P \bar{z}_{r_0-P}^i I} \sum_{s=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} z_s^i \nabla f_i(\mathbf{y}_{s,k}^i) \right\|^2 \middle| \mathcal{Q} \right] + 5 \mathbb{E} \left[\left\| \nabla f(\bar{\mathbf{x}}_{r_0}) \right\|^2 \middle| \mathcal{Q} \right] \\
&\leq 5L^2 D_{r,k} + 5L^2 M_{r,k} \\
&\quad + 5 \sum_{i=1}^N \left(q_r^i + \frac{1}{N} \right) \mathbb{E} \left[\left\| \frac{1}{P \bar{z}_{r_0-P}^i I} \sum_{s=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} z_s^i (\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \nabla f_i(\mathbf{y}_{s,k}^i)) \right\|^2 \middle| \mathcal{Q} \right] + 5 \mathbb{E} \left[\left\| \nabla f(\bar{\mathbf{x}}_{r_0}) \right\|^2 \middle| \mathcal{Q} \right] \\
&\leq 5L^2 D_{r,k} + 5L^2 M_{r,k} \\
&\quad + 5 \sum_{i=1}^N \left(q_r^i + \frac{1}{N} \right) \frac{1}{P \bar{z}_{r_0-P}^i I} \sum_{s=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} z_s^i \mathbb{E} \left[\left\| \nabla f_i(\bar{\mathbf{x}}_{r_0}) - \nabla f_i(\mathbf{y}_{s,k}^i) \right\|^2 \middle| \mathcal{Q} \right] + 5 \mathbb{E} \left[\left\| \nabla f(\bar{\mathbf{x}}_{r_0}) \right\|^2 \middle| \mathcal{Q} \right] \\
&\leq 5L^2 D_{r,k} + 5L^2 M_{r,k} + 5L^2 \sum_{i=1}^N \left(q_r^i + \frac{1}{N} \right) S_{r_0}^i + 5 \mathbb{E} \left[\left\| \nabla f(\bar{\mathbf{x}}_{r_0}) \right\|^2 \middle| \mathcal{Q} \right].
\end{aligned}$$

Plugging back into Equation 10

$$\begin{aligned}
M_{r,k+1} &\leq \left(1 + \frac{1}{\lambda_2} \right) M_{r,k} + 5\eta^2 L^2 (1 + \lambda_2) \left(D_{r,k} + M_{r,k} + \sum_{i=1}^N \left(q_r^i + \frac{1}{N} \right) S_{r_0}^i \right) \\
&\quad + 5\eta^2 (1 + \lambda_2) \mathbb{E} \left[\left\| \nabla f(\bar{\mathbf{x}}_{r_0}) \right\|^2 \middle| \mathcal{Q} \right] + 9\eta^2 \rho^2 \sigma^2 \\
&\leq \left(1 + \frac{1}{\lambda_2} + 5\eta^2 L^2 (1 + \lambda_2) \right) M_{r,k} + 5\eta^2 L^2 (1 + \lambda_2) \left(D_{r,k} + \sum_{i=1}^N \left(q_r^i + \frac{1}{N} \right) S_{r_0}^i \right) \\
&\quad + 5\eta^2 (1 + \lambda_2) \mathbb{E} \left[\left\| \nabla f(\bar{\mathbf{x}}_{r_0}) \right\|^2 \middle| \mathcal{Q} \right] + 9\eta^2 \rho^2 \sigma^2.
\end{aligned}$$

Now choose $\lambda_2 = 2IP$, so that $1 \leq \frac{\lambda_2}{2}$ and

$$5\eta^2 L^2 (1 + \lambda_2) \leq \frac{15}{2} \eta^2 L^2 \lambda_2 \leq 15\eta^2 L^2 IP \leq 15L^2 IP \frac{1}{72L^2 I^2 P^2} \leq \frac{1}{2IP},$$

where we used the condition $\eta \leq \frac{1}{60LIP}$. This yields the following recursive bound on $M_{r,k+1}$:

$$\begin{aligned} M_{r,k+1} &\leq \left(1 + \frac{1}{IP}\right) M_{r,k} + 15\eta^2 L^2 IP D_{r,k} + 15\eta^2 L^2 IP \sum_{i=1}^N \left(q_r^i + \frac{1}{N}\right) S_{r_0}^i \\ &\quad + 15\eta^2 IP \mathbb{E} \left[\|\nabla f(\bar{x}_{r_0})\|^2 \middle| \mathcal{Q} \right] + 9\eta^2 \rho^2 \sigma^2. \end{aligned} \quad (11)$$

Note that the same argument can be used to bound $M_{r,0}$ in terms of $M_{r-1,I-1}$ with the same recurrence relation.

Unrolling the recursion For $0 \leq t \leq IP$, let $r_t = r_0 + \lfloor t/I \rfloor$ and $k_t = t - \lfloor t/I \rfloor * I$. Then let $d_t = D_{r_t,k_t}$ and $m_t = M_{r_t,k_t}$. Also, let

$$\begin{aligned} q_1 &= 18\eta^2 L^2 I \\ q_2 &= 15\eta^2 L^2 IP \\ a_r &= 18\eta^2 L^2 I \sum_{i=1}^N q_r^i S_{r_0}^i + 36\eta^2 \sigma^2 \\ b_r &= 15\eta^2 L^2 IP \sum_{i=1}^N \left(q_r^i + \frac{1}{N}\right) S_{r_0}^i + 15\eta^2 IP \mathbb{E} \left[\|\nabla f(\bar{x}_{r_0})\|^2 \middle| \mathcal{Q} \right] + 9\eta^2 \rho^2 \sigma^2. \end{aligned}$$

Then Equation 9 and Equation 11 imply the following mutually recursive relation over d_t and m_t :

$$\begin{aligned} d_t &= \begin{cases} 0 & k_t = 0 \\ (1 + \frac{1}{I}) d_{t-1} + q_1 m_{t-1} + a_{r_{t-1}} & \text{otherwise} \end{cases} \\ m_t &= \left(1 + \frac{1}{IP}\right) m_{t-1} + q_2 d_{t-1} + b_{r_{t-1}}, \end{aligned} \quad (12)$$

where $d_0 = m_0 = 0$. We will unroll this recurrence, showing a bound for m_t and d_t in terms of the following quantities. For any $0 \leq t \leq IP - 1$ and $0 \leq s \leq r_t$, let $j(s,t) = \min\{I, t - sI\}$ and $\ell(s,t) = \max\{0, t - (s+1)I\}$. For any s, t , if $k_{t+1} = 0$, define

$$\alpha_{s,t+1} = 0.$$

Otherwise, define

$$\begin{aligned} \alpha_{s,t+1} &= \alpha_{s,t} \left(1 + \frac{1}{I}\right) + q_1 P \left(\left(1 + \frac{1}{IP}\right)^{j(s,t)} - 1 \right) \left(1 + \frac{1}{IP}\right)^{\ell(s,t)} \\ &\quad + q_1 q_2 \sum_{i=0}^{t-sI-2} \left(1 + \frac{1}{IP}\right)^i \alpha_{s,t-1-i}. \end{aligned}$$

Also, define

$$\begin{aligned} \beta_{s,t+1} &= \beta_{s,t} \left(1 + \frac{1}{IP}\right) + \mathbb{1}\{s = r_t\} \frac{q_2}{P} \left(\left(1 + \frac{1}{I}\right)^{k_t} - 1 \right) + q_1 q_2 \sum_{i=0}^{k_t-2} \left(1 + \frac{1}{I}\right)^i \beta_{s,t-1-i} \\ \psi_{s,t+1} &= \frac{1}{P} \sum_{i=0}^{t-sI-1} \left(1 + \frac{1}{IP}\right)^i \alpha_{s,t-i} \\ \phi_{s,t+1} &= P \sum_{i=0}^{k_t-1} \left(1 + \frac{1}{I}\right)^i \beta_{s,t-i}, \end{aligned}$$

under the initial conditions

$$\begin{aligned} \phi_{s,kI} &= 0 \text{ for all } s \leq P \text{ and } s \leq k \leq P \\ \alpha_{s,kI} &= 0 \text{ for all } s \leq P \text{ and } s \leq k \leq P \\ \psi_{s,sI} &= 0 \text{ for all } s \leq P \\ \beta_{s,sI} &= 0 \text{ for all } s \leq P. \end{aligned}$$

Now we can unroll the recurrence in Equation 12 by proving the following statements by induction on t :

$$d_t \leq \left(\left(1 + \frac{1}{I} \right)^{k_t} - 1 \right) a_{r_t} I + q_1 I \sum_{s=0}^{r_t} \phi_{s,t} a_s + I \sum_{s=0}^{r_t} \alpha_{s,t} b_s \quad (13)$$

$$m_t \leq IP \sum_{s=0}^{r_t} \left(\left(\left(1 + \frac{1}{IP} \right)^{j(s,t)} - 1 \right) \left(1 + \frac{1}{IP} \right)^{\ell(s,t)} + q_2 \psi_{s,t} \right) b_s + IP \sum_{s=0}^{r_t} \beta_{s,t} a_s. \quad (14)$$

Equation 13 and Equation 14 hold for the base case $t = 0$, since $d_0 = m_0 = 0$. Now suppose that Equation 13 and Equation 14 hold for some $t \leq IP$, and we will show that they hold for $t + 1$. We consider two cases: $k_{t+1} \neq 0$ and $k_{t+1} = 0$.

In the first case, $r_{t+1} = r_t$, i.e., step $t + 1$ is in the same round as step t , and $k_{t+1} = k_t + 1$. Using Equation 12 together with the inductive hypothesis:

$$\begin{aligned} d_{t+1} &\leq \left(1 + \frac{1}{I} \right) \left(\left(\left(1 + \frac{1}{I} \right)^{k_t} - 1 \right) a_{r_t} I + q_1 I \sum_{s=0}^{r_t} \phi_{s,t} a_s + I \sum_{s=0}^{r_t} \alpha_{s,t} b_s \right) \\ &\quad + q_1 \left(IP \sum_{s=0}^{r_t} \left(\left(\left(1 + \frac{1}{IP} \right)^{j(s,t)} - 1 \right) \left(1 + \frac{1}{IP} \right)^{\ell(s,t)} + q_2 \psi_{s,t} \right) b_s + IP \sum_{s=0}^{r_t} \beta_{s,t} a_s \right) + a_{r_t} \\ &\leq \left(\left(1 + \frac{1}{I} \right)^{k_t+1} - \left(1 + \frac{1}{I} \right) + \frac{1}{I} \right) a_{r_t} I + q_1 I \sum_{s=0}^{r_t} \left(\left(1 + \frac{1}{I} \right) \phi_{s,t} + P \beta_{s,t} \right) a_s \\ &\quad + I \sum_{s=0}^{r_t} \left(\left(1 + \frac{1}{I} \right) \alpha_{s,t} + q_1 P \left(\left(1 + \frac{1}{IP} \right)^{j(s,t)} - 1 \right) \left(1 + \frac{1}{IP} \right)^{\ell(s,t)} + q_1 q_2 P \psi_{s,t} \right) b_s \\ &\stackrel{(i)}{\leq} \left(\left(1 + \frac{1}{I} \right)^{k_t+1} - 1 \right) a_{r_t} I + q_1 I \sum_{s=0}^{r_t} \phi_{s,t+1} a_s + I \sum_{s=0}^{r_t} \alpha_{s,t+1} b_s \\ &\stackrel{(ii)}{\leq} \left(\left(1 + \frac{1}{I} \right)^{k_{t+1}} - 1 \right) a_{r_{t+1}} I + q_1 I \sum_{s=0}^{r_{t+1}} \phi_{s,t+1} a_s + I \sum_{s=0}^{r_{t+1}} \alpha_{s,t+1} b_s, \end{aligned}$$

where (i) uses the fact that $\phi_{s,t+1} = (1 + \frac{1}{I}) \phi_{s,t} + P \beta_{s,t}$ and the definitions of $\alpha_{s,t+1}$, $\psi_{s,t}$, and (ii) uses $k_{t+1} = k_t + 1$ and $r_{t+1} = r_t$. Similarly,

$$\begin{aligned}
m_{t+1} &\leq \left(1 + \frac{1}{IP}\right) \left(IP \sum_{s=0}^{r_t} \left(\left(\left(1 + \frac{1}{IP}\right)^{j(s,t)} - 1 \right) \left(1 + \frac{1}{IP}\right)^{\ell(s,t)} + q_2 \psi_{s,t} \right) b_s + IP \sum_{s=0}^{r_t} \beta_{s,t} a_s\right) \\
&\quad + q_2 \left(\left(\left(1 + \frac{1}{I}\right)^{k_t} - 1 \right) a_{r_t} I + q_1 I \sum_{s=0}^{r_t} \phi_{s,t} a_s + I \sum_{s=0}^{r_t} \alpha_{s,t} b_s \right) + b_{r_t} \\
&\leq IP \sum_{s=0}^{r_t} \left(\left(\left(1 + \frac{1}{IP}\right)^{j(s,t)} - 1 \right) \left(1 + \frac{1}{IP}\right)^{\ell(s,t)+1} \right. \\
&\quad \left. + \mathbb{1}\{s = r_t\} \frac{1}{IP} + q_2 \left(\left(1 + \frac{1}{IP}\right) \psi_{s,t} + \frac{1}{P} \alpha_{s,t} \right) \right) b_s \\
&\quad + IP \sum_{s=0}^{r_t} \left(\left(1 + \frac{1}{IP}\right) \beta_{s,t} + \mathbb{1}\{s = r_t\} \frac{q_2}{P} \left(\left(1 + \frac{1}{I}\right)^{k_t} - 1 \right) + \frac{q_1 q_2}{P} \phi_{s,t} \right) a_s \\
&\stackrel{(i)}{\leq} IP \sum_{s=0}^{r_t} \left(\left(\left(1 + \frac{1}{IP}\right)^{j(s,t+1)} - 1 \right) \left(1 + \frac{1}{IP}\right)^{\ell(s,t+1)} + q_2 \left(\left(1 + \frac{1}{IP}\right) \psi_{s,t} + \frac{1}{P} \alpha_{s,t} \right) \right) b_s \\
&\quad + IP \sum_{s=0}^{r_t} \left(\left(1 + \frac{1}{IP}\right) \beta_{s,t} + \mathbb{1}\{s = r_t\} \frac{q_2}{P} \left(\left(1 + \frac{1}{I}\right)^{k_t} - 1 \right) + \frac{q_1 q_2}{P} \phi_{s,t} \right) a_s \\
&\stackrel{(ii)}{\leq} IP \sum_{s=0}^{r_t} \left(\left(\left(1 + \frac{1}{IP}\right)^{j(s,t+1)} - 1 \right) \left(1 + \frac{1}{IP}\right)^{\ell(s,t+1)} + q_2 \psi_{s,t+1} \right) b_s + IP \sum_{s=0}^{r_t} \beta_{s,t+1} a_s \\
&\stackrel{(iii)}{\leq} IP \sum_{s=0}^{r_{t+1}} \left(\left(\left(1 + \frac{1}{IP}\right)^{j(s,t+1)} - 1 \right) \left(1 + \frac{1}{IP}\right)^{\ell(s,t+1)} + q_2 \psi_{s,t+1} \right) b_s + IP \sum_{s=0}^{r_{t+1}} \beta_{s,t+1} a_s,
\end{aligned}$$

where (i) uses the fact that for $s < r_t$:

$$\begin{aligned}
&\left(\left(1 + \frac{1}{IP}\right)^{j(s,t)} - 1 \right) \left(1 + \frac{1}{IP}\right)^{\ell(s,t)+1} + \mathbb{1}\{s = r_t\} \frac{1}{IP} \\
&= \left(\left(1 + \frac{1}{IP}\right)^I - 1 \right) \left(1 + \frac{1}{IP}\right)^{t+1-(s+1)I} \\
&= \left(\left(1 + \frac{1}{IP}\right)^{j(s,t+1)} - 1 \right) \left(1 + \frac{1}{IP}\right)^{\ell(s,t+1)}
\end{aligned}$$

and for $s = r_t$:

$$\begin{aligned}
&\left(\left(1 + \frac{1}{IP}\right)^{j(s,t)} - 1 \right) \left(1 + \frac{1}{IP}\right)^{\ell(s,t)+1} + \mathbb{1}\{s = r_t\} \frac{1}{IP} \\
&= \left(\left(1 + \frac{1}{IP}\right)^{t-sI} - 1 \right) \left(1 + \frac{1}{IP}\right) + \frac{1}{IP} \\
&= \left(\left(1 + \frac{1}{IP}\right)^{t+1-sI} - 1 \right) \\
&= \left(\left(1 + \frac{1}{IP}\right)^{j(s,t+1)} - 1 \right) \left(1 + \frac{1}{IP}\right)^{\ell(s,t+1)},
\end{aligned}$$

(ii) uses the fact that $\psi_{s,t+1} = (1 + \frac{1}{IP}) \psi_{s,t} + \frac{1}{P} \alpha_{s,t}$ and the definitions of $\beta_{s,t+1}$, $\phi_{s,t}$, and (iii) uses $k_{t+1} = k_t + 1$ and $r_{t+1} = r_t$. This completes the inductive step for the first case ($k_{t+1} \neq 0$).

In the second case (i.e., $k_{t+1} = 0$), Equation 13 must hold for $t + 1$, since $d_{t+1} = 0$. So it only remains to show Equation 14 holds for $t + 1$. Note that $r_t = r_{t+1} - 1$. As in the first case, we can

use Equation 12 together with the inductive hypothesis to obtain:

$$\begin{aligned} m_{t+1} &\leq IP \sum_{s=0}^{r_t} \left(\left(\left(1 + \frac{1}{IP} \right)^{j(s,t+1)} - 1 \right) \left(1 + \frac{1}{IP} \right)^{\ell(s,t+1)} + q_2 \psi_{s,t+1} \right) b_s + IP \sum_{s=0}^{r_t} \phi_{s,t+1} a_s \\ &\leq IP \sum_{s=0}^{r_{t+1}} \left(\left(\left(1 + \frac{1}{IP} \right)^{j(s,t+1)} - 1 \right) \left(1 + \frac{1}{IP} \right)^{\ell(s,t+1)} + q_2 \psi_{s,t+1} \right) b_s + IP \sum_{s=0}^{r_{t+1}} \phi_{s,t+1} a_s, \end{aligned}$$

where the second line uses the fact that the r_{t+1} -st element of both sums is 0, since $j(r_{t+1}, t+1) = 0$, $\psi_{r_{t+1}, t+1} = 0$, and $\phi_{r_{t+1}, t+1} = 0$. This completes the inductive step for both cases, and proves Equation 13 and Equation 14.

We can now bound $\alpha_{s,t}$ and $\beta_{s,t}$ separately by induction on t . First, for any $s \leq P-1$ and t with $sI \leq t \leq IP$, we claim that

$$\alpha_{s,t} \leq 36q_1 IP, \quad (15)$$

which we will show by induction on t . Let $s \leq P-1$ be given. For the base case, Equation 15 holds when $t = sI$ since $\alpha_{s,sI} = 0$. Now suppose that it holds for all $t' \leq t$. Then

$$\begin{aligned} &\left(\left(1 + \frac{1}{IP} \right)^{m_1} - 1 \right) \left(1 + \frac{1}{IP} \right)^{m_2} \\ &\leq \left(\left(1 + \frac{1}{IP} \right)^{t-sI} - 1 \right) \left(1 + \frac{1}{IP} \right)^0 \\ &\leq \left(1 + \frac{1}{IP} \right)^{t-sI} - 1, \end{aligned} \quad (16)$$

and

$$\begin{aligned} q_1 q_2 \sum_{i=0}^{t-sI-2} \left(1 + \frac{1}{IP} \right)^i \alpha_{s,t-1-i} &\leq 36q_1^2 q_2 IP \sum_{i=0}^{t-sI-2} \left(1 + \frac{1}{IP} \right)^i \\ &= 36q_1^2 q_2 I^2 P^2 \left(\left(1 + \frac{1}{IP} \right)^{t-sI-1} - 1 \right) \\ &\leq q_1 P \left(\left(1 + \frac{1}{IP} \right)^{t-sI} - 1 \right), \end{aligned} \quad (17)$$

where the last line uses the definition of q_1 and q_2 together with the condition $\eta \leq \frac{1}{60LIP}$:

$$36q_1 q_2 I^2 P = 36 \cdot 270 \eta^4 L^4 I^4 P^2 \leq \frac{36 \cdot 270}{60^4} L^4 I^4 P^2 \frac{1}{L^4 I^4 P^4} \leq \frac{1}{P^2} \leq 1.$$

Plugging Equation 16 and Equation 17 into the definition of $\alpha_{s,t+1}$ yields

$$\begin{aligned} \alpha_{s,t+1} &\leq \alpha_{s,t} \left(1 + \frac{1}{I} \right) + 2q_1 P \left(\left(1 + \frac{1}{IP} \right)^{t-sI} - 1 \right) \\ &\stackrel{(i)}{\leq} 2q_1 P \sum_{i=0}^{k_t} \left(1 + \frac{1}{I} \right)^i \left(\left(1 + \frac{1}{IP} \right)^{t-sI-i} - 1 \right) \\ &\leq 2q_1 P \left(1 + \frac{1}{I} \right)^I \sum_{i=0}^{k_t} \left(\left(1 + \frac{1}{IP} \right)^{t-sI-i} - 1 \right) \\ &\leq 6q_1 P \sum_{i=0}^{k_t} \left(1 + \frac{1}{IP} \right)^{t-sI-i} \leq 6q_1 P \left(1 + \frac{1}{IP} \right)^{(r_t-s)I} \sum_{i=0}^{k_t} \left(1 + \frac{1}{IP} \right)^i \\ &\leq 18q_1 P \sum_{i=0}^{I-1} \left(1 + \frac{1}{IP} \right)^i \stackrel{(ii)}{\leq} 18q_1 \sum_{i=0}^{IP-1} \left(1 + \frac{1}{IP} \right)^i \\ &\leq 18q_1 IP \left(\left(1 + \frac{1}{IP} \right)^{IP} - 1 \right) \leq 36q_1 IP, \end{aligned}$$

where (i) unrolls the recurrence on the first line until $\alpha_{s,r_t I} = 0$, (ii) uses $P \sum_{i=0}^{I-1} (1 + \frac{1}{IP})^i \leq \sum_{i=0}^{IP-1} (1 + \frac{1}{IP})^i$ since $(1 + \frac{1}{IP})^i$ is increasing with i , and we repeatedly used $(1 + \frac{1}{x})^x < e \leq 3$. This completes the induction and proves Equation 15.

We will similarly prove the statement

$$\beta_{s,t} \leq 12 \frac{q_2 I}{P}. \quad (18)$$

for all $s \leq P - 1$ and t with $sI \leq t \leq IP$ by induction on t . Let $s \leq P - 1$ be given. For the base case, Equation 15 holds when $t = sI$ since $\beta_{s,sI} = 0$. Now suppose that it holds for all $t' \leq t$. Then

$$\begin{aligned} q_1 q_2 \sum_{i=0}^{k_t-2} \left(1 + \frac{1}{I}\right)^i \beta_{s,t-1-i} &\leq 12 \frac{q_1 q_2^2 I}{P} \sum_{i=0}^{k_t-2} \left(1 + \frac{1}{I}\right)^i \\ &= 12 \frac{q_1 q_2^2 I^2}{P} \left(\left(1 + \frac{1}{I}\right)^{k_t-1} - 1 \right) \\ &\leq 12 \frac{q_1 q_2^2 I^2}{P} \left(\left(1 + \frac{1}{I}\right)^{k_t} - 1 \right) \end{aligned} \quad (19)$$

$$\leq \frac{q_2}{P} \left(\left(1 + \frac{1}{I}\right)^{k_t} - 1 \right), \quad (20)$$

where the last line uses the definition of q_1 and q_2 together with the condition $\eta \leq \frac{1}{60LIP}$:

$$12q_1 q_2 I^2 \leq 12 \cdot 270 \eta^4 L^4 I^4 P \leq \frac{12 \cdot 270}{60^4} L^4 I^4 P \frac{1}{L^4 I^4 P^4} \leq \frac{1}{P^3} \leq 1.$$

We now consider two cases: $t < (s+1)I$ and $t \geq (s+1)I$. In the first case, we have $s = r_t$, and plugging Equation 20 into the definition of $\beta_{s,t+1}$ yields

$$\begin{aligned} \beta_{s,t+1} &\leq \beta_{s,t} \left(1 + \frac{1}{IP}\right) + 2 \frac{q_2}{P} \left(\left(1 + \frac{1}{I}\right)^{k_t} - 1 \right) \\ &\stackrel{(i)}{\leq} 2 \frac{q_2}{P} \sum_{i=0}^{k_t} \left(1 + \frac{1}{IP}\right)^i \left(\left(1 + \frac{1}{I}\right)^{k_t-i} - 1 \right) \\ &\leq 2 \frac{q_2}{P} \left(1 + \frac{1}{IP}\right)^{k_t} \sum_{i=0}^{k_t} \left(\left(1 + \frac{1}{I}\right)^{k_t-i} - 1 \right) \\ &\leq 2 \frac{q_2}{P} \left(1 + \frac{1}{IP}\right)^I \sum_{i=0}^{k_t} \left(1 + \frac{1}{I}\right)^i \\ &\leq 6 \frac{q_2 I}{P} \left(\left(1 + \frac{1}{I}\right)^{k_t+1} - 1 \right) \leq 6 \frac{q_2 I}{P} \left(\left(1 + \frac{1}{I}\right)^I - 1 \right) \\ &\leq 12 \frac{q_2 I}{P}, \end{aligned}$$

where (i) unrolls the recurrence on the first line until $\beta_{s,sI} = 0$ and we repeatedly used $(1 + \frac{1}{x})^x < e \leq 3$. In the second case, we have $s > r_t$, so plugging Equation 19 into the definition of $\beta_{s,t+1}$

yields

$$\begin{aligned}
\beta_{s,t+1} &\leq \beta_{s,t} \left(1 + \frac{1}{IP}\right) + 12 \frac{q_1 q_2^2 I^2}{P} \left(\left(1 + \frac{1}{I}\right)^{k_t} - 1 \right) \\
&\leq \beta_{s,t} \left(1 + \frac{1}{IP}\right) + 12 \frac{q_1 q_2^2 I^2}{P} \left(\left(1 + \frac{1}{I}\right)^I - 1 \right) \\
&\leq \beta_{s,t} \left(1 + \frac{1}{IP}\right) + 24 \frac{q_1 q_2^2 I^2}{P} \\
&\stackrel{(i)}{\leq} 24 \frac{q_1 q_2^2 I^2}{P} \sum_{i=0}^{t-sI} \left(1 + \frac{1}{IP}\right)^i \\
&= 24 q_1 q_2^2 I^3 \left(\left(1 + \frac{1}{IP}\right)^{t-sI+1} - 1 \right) \leq 24 q_1 q_2^2 I^3 \left(\left(1 + \frac{1}{IP}\right)^{IP} - 1 \right) \\
&\leq 48 q_1 q_2^2 I^3 \leq 12 \frac{q_2 I}{P},
\end{aligned}$$

where (i) unrolls the recurrence on the previous line until $\beta_{s,sI} = 0$ and the last inequality uses the definition of q_1 and q_2 together with the condition $\eta \leq \frac{1}{60LIP}$:

$$48 q_1 q_2 I^2 = 48 \cdot 270 \eta^4 L^4 I^4 P \leq \frac{48 \cdot 270}{60^4} L^4 I^4 P \frac{1}{L^4 I^4 P^4} \leq \frac{12}{P^3} \leq \frac{12}{P}.$$

This completes the induction in both cases and proves Equation 18.

We can then use Equation 15 and Equation 18 to yield bounds for the remaining terms of Equation 13 and Equation 14 as follows:

$$\begin{aligned}
q_2 \psi_{s,t} &= \frac{q_2}{P} \sum_{i=0}^{t-sI-1} \left(1 + \frac{1}{IP}\right)^i \alpha_{s,t-1-i} \leq 36 q_1 q_2 I \sum_{i=0}^{t-sI-1} \left(1 + \frac{1}{IP}\right)^i \\
&\leq 36 q_1 q_2 I \sum_{i=0}^{IP-1} \left(1 + \frac{1}{IP}\right)^i \leq 36 q_1 q_2 I^2 P \left(\left(1 + \frac{1}{IP}\right)^{IP} - 1 \right) \\
&\leq 72 q_1 q_2 I^2 P \leq 72 \cdot 270 \eta^4 L^4 I^4 P^2 \leq \frac{72 \cdot 270}{60^4} L^4 I^4 P^2 \frac{1}{L^4 I^4 P^4} \leq \frac{1}{P^2},
\end{aligned}$$

and

$$\begin{aligned}
q_1 \phi_{s,t} &= q_1 P \sum_{i=0}^{k_t-1} \left(1 + \frac{1}{I}\right)^i \beta_{s,t-i} \leq 12 q_1 q_2 I \sum_{i=0}^{k_t-1} \left(1 + \frac{1}{I}\right)^i \\
&= 12 q_1 q_2 I^2 \left(\left(1 + \frac{1}{I}\right)^{k_t} - 1 \right) \leq 12 q_1 q_2 I^2 \left(\left(1 + \frac{1}{I}\right)^I - 1 \right) \leq 24 q_1 q_2 I^2.
\end{aligned}$$

Finally, we can plug these into Equation 13 to yield

$$\begin{aligned}
d_t &\leq \left(\left(1 + \frac{1}{I}\right)^{k_t} - 1 \right) a_{r_t} I + 24 q_1 q_2 I^3 \sum_{s=0}^{r_t} a_s + 36 q_1 I^2 P \sum_{s=0}^{r_t} b_s \\
&\leq 2 a_{r_t} I + 24 q_1 q_2 I^3 \sum_{s=0}^{r_t} a_s + 36 q_1 I^2 P \sum_{s=0}^{r_t} b_s,
\end{aligned}$$

and into Equation 14

$$\begin{aligned}
m_t &\leq IP \sum_{s=0}^{r_t} \left(\left(\left(1 + \frac{1}{IP} \right)^{m_1} - 1 \right) \left(1 + \frac{1}{IP} \right)^{m_2} + \frac{1}{P^2} \right) b_s + 12q_2 I^2 \sum_{s=0}^{r_t} a_s \\
&\stackrel{(i)}{\leq} IP \sum_{s=0}^{r_t} \left(\frac{6}{P} + \frac{1}{P^2} \right) b_s + 12q_2 I^2 \sum_{s=0}^{r_t} a_s \\
&\leq 7I \sum_{s=0}^{r_t} b_s + 12q_2 I^2 \sum_{s=0}^{r_t} a_s,
\end{aligned}$$

where (i) uses

$$\begin{aligned}
\left(1 + \frac{1}{IP} \right)^{m_1} - 1 &\leq \left(1 + \frac{1}{IP} \right)^I - 1 = IP \sum_{i=0}^{I-1} \left(1 + \frac{1}{IP} \right)^i \\
&\leq I \sum_{i=0}^{IP-1} \left(1 + \frac{1}{IP} \right)^i = \frac{1}{P} \left(\left(1 + \frac{1}{IP} \right)^{IP} - 1 \right) \leq \frac{2}{P},
\end{aligned}$$

and

$$\left(1 + \frac{1}{IP} \right)^{m_2} \leq \left(1 + \frac{1}{IP} \right)^{IP} \leq 3.$$

Or, in the original notation:

$$\begin{aligned}
D_{r,k} &\leq 2I \left(18\eta^2 L^2 I \sum_{i=1}^N q_{r_t}^i S_{r_0}^i + 36\eta^2 \sigma^2 \right) + 6480\eta^4 L^4 I^5 P \sum_{s=r_0}^r \left(18\eta^2 L^2 I \sum_{i=1}^N q_s^i S_{r_0}^i + 36\eta^2 \sigma^2 \right) \\
&\quad + 648\eta^2 L^2 I^3 P \sum_{s=r_0}^r \left(15\eta^2 L^2 IP \sum_{i=1}^N \left(q_s^i + \frac{1}{N} \right) S_{r_0}^i + 15\eta^2 IP \mathbb{E} \left[\|\nabla f(\bar{x}_{r_0})\|^2 \middle| \mathcal{Q} \right] + 9\eta^2 \rho^2 \sigma^2 \right) \\
&\leq 36\eta^2 L^2 I^2 \sum_{i=1}^N q_r^i S_{r_0}^i + 18 \cdot 6480\eta^6 L^6 I^6 P \sum_{s=r_0}^r \sum_{i=1}^N q_s^i S_{r_0}^i \\
&\quad + 9720\eta^4 L^4 I^4 P^2 \sum_{s=r_0}^r \sum_{i=1}^N \left(q_s^i + \frac{1}{N} \right) S_{r_0}^i + 9720(r-r_0)\eta^4 L^2 I^4 P^3 \mathbb{E} \left[\|\nabla f(\bar{x}_{r_0})\|^2 \middle| \mathcal{Q} \right] \\
&\quad + (72\eta^2 I + 36 \cdot 6480(r-r_0)\eta^6 L^4 I^5 P + 5832(r-r_0)\eta^4 L^2 I^3 P \rho^2) \sigma^2 \\
&\stackrel{(i)}{\leq} 36\eta^2 L^2 I^2 \sum_{i=1}^N q_r^i S_{r_0}^i + 18 \cdot 6480\eta^6 L^6 I^6 P \sum_{s=r_0}^{r_0+P-1} \sum_{i=1}^N q_s^i S_{r_0}^i \\
&\quad + 9720\eta^4 L^4 I^4 P^2 \sum_{s=r_0}^{r_0+P-1} \sum_{i=1}^N \left(q_s^i + \frac{1}{N} \right) S_{r_0}^i \\
&\quad + (72\eta^2 I + 36 \cdot 6480\eta^6 L^4 I^5 P^2 + 5832\eta^4 L^2 I^3 P^2 \rho^2) \sigma^2 + 9720\eta^4 L^2 I^4 P^3 \mathbb{E} \left[\|\nabla f(\bar{x}_{r_0})\|^2 \middle| \mathcal{Q} \right] \\
&= 36\eta^2 L^2 I^2 \sum_{i=1}^N q_r^i S_{r_0}^i + 9720\eta^4 L^4 I^4 P^3 \frac{1}{N} \sum_{i=1}^N S_{r_0}^i \\
&\quad + (18 \cdot 6480\eta^6 L^6 I^6 P^2 + 9720\eta^4 L^4 I^4 P^3) \sum_{i=1}^N \bar{q}_{r_0}^i S_{r_0}^i \\
&\quad + (72\eta^2 I + 36 \cdot 6480\eta^6 L^4 I^5 P^2 + 5832\eta^4 L^2 I^3 P^2 \rho^2) \sigma^2 + 9720\eta^4 L^2 I^4 P^3 \mathbb{E} \left[\|\nabla f(\bar{x}_{r_0})\|^2 \middle| \mathcal{Q} \right] \\
&\stackrel{(ii)}{\leq} 36\eta^2 L^2 I^2 \sum_{i=1}^N q_r^i S_{r_0}^i + 9720\eta^4 L^4 I^4 P^3 \frac{1}{N} \sum_{i=1}^N S_{r_0}^i + 9753\eta^4 L^4 I^4 P^3 \sum_{i=1}^N \bar{q}_{r_0}^i S_{r_0}^i \\
&\quad + (73\eta^2 I + 5832\eta^4 L^2 I^3 P^2 \rho^2) \sigma^2 + 9720\eta^4 L^2 I^4 P^3 \mathbb{E} \left[\|\nabla f(\bar{x}_{r_0})\|^2 \middle| \mathcal{Q} \right], \tag{21}
\end{aligned}$$

where (i) uses the fact that the interval $\{r_0, \dots, r\}$ is contained in the interval $\{r_0, \dots, r_0 + P - 1\}$ and (ii) uses the condition $\eta \leq \frac{1}{60LIP}$ to simplify non-dominating terms. Also

$$\begin{aligned}
M_{r,k} &\leq 7I \sum_{r=r_0}^r \left(15\eta^2 L^2 IP \sum_{i=1}^N \left(q_s^i + \frac{1}{N} \right) S_{r_0}^i + 15\eta^2 IP \mathbb{E} \left[\|\nabla f(\bar{x}_{r_0})\|^2 \middle| \mathcal{Q} \right] + 9\eta^2 \rho^2 \sigma^2 \right) \\
&\quad + 180\eta^2 L^2 I^3 P \sum_{r=r_0}^r \left(18\eta^2 L^2 I \sum_{i=1}^N q_s^i S_{r_0}^i + 36\eta^2 \sigma^2 \right) \\
&\leq 105\eta^2 L^2 I^2 P \sum_{r=r_0}^r \sum_{i=1}^N \left(q_s^i + \frac{1}{N} \right) S_{r_0}^i + 3240\eta^4 L^4 I^4 P \sum_{r=r_0}^r \sum_{i=1}^N q_s^i S_{r_0}^i \\
&\quad + (63(r-r_0)\eta^2 I\rho^2 + 6480(r-r_0)\eta^4 L^2 I^3 P) \sigma^2 + 105(r-r_0)\eta^2 I^2 P \mathbb{E} \left[\|\nabla f(\bar{x}_{r_0})\|^2 \middle| \mathcal{Q} \right] \\
&\stackrel{(i)}{\leq} 105\eta^2 L^2 I^2 P \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N \left(q_s^i + \frac{1}{N} \right) S_{r_0}^i + 3240\eta^4 L^4 I^4 P \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N q_s^i S_{r_0}^i \\
&\quad + (63\eta^2 IP\rho^2 + 6480\eta^4 L^2 I^3 P^2) \sigma^2 + 105\eta^2 I^2 P^2 \mathbb{E} \left[\|\nabla f(\bar{x}_{r_0})\|^2 \middle| \mathcal{Q} \right] \\
&\leq 105\eta^2 L^2 I^2 P^2 \frac{1}{N} \sum_{i=1}^N S_{r_0}^i + (105\eta^2 L^2 I^2 P^2 + 3240\eta^4 L^4 I^4 P^2) \sum_{i=1}^N \bar{q}_{r_0}^i S_{r_0}^i \\
&\quad + (63\eta^2 IP\rho^2 + 6480\eta^4 L^2 I^3 P^2) \sigma^2 + 105\eta^2 I^2 P^2 \mathbb{E} \left[\|\nabla f(\bar{x}_{r_0})\|^2 \middle| \mathcal{Q} \right] \\
&\stackrel{(ii)}{\leq} 105\eta^2 L^2 I^2 P^2 \frac{1}{N} \sum_{i=1}^N S_{r_0}^i + 106\eta^2 L^2 I^2 P^2 \sum_{i=1}^N \bar{q}_{r_0}^i S_{r_0}^i \\
&\quad + (63\eta^2 IP\rho^2 + 6480\eta^4 L^2 I^3 P^2) \sigma^2 + 105\eta^2 I^2 P^2 \mathbb{E} \left[\|\nabla f(\bar{x}_{r_0})\|^2 \middle| \mathcal{Q} \right], \tag{22}
\end{aligned}$$

where (i) and (ii) use the same operations as in Equation 21. Summing Equation 22 and Equation 21,

$$\begin{aligned}
D_{r,k} + M_{r,k} &\leq 36\eta^2 L^2 I^2 \sum_{i=1}^N \bar{q}_r^i S_{r_0}^i + (9720\eta^4 L^4 I^4 P^3 + 105\eta^2 L^2 I^2 P^2) \frac{1}{N} \sum_{i=1}^N S_{r_0}^i \\
&\quad + (9753\eta^4 L^4 I^4 P^3 + 106\eta^2 L^2 I^2 P^2) \sum_{i=1}^N \bar{q}_{r_0}^i S_{r_0}^i \\
&\quad + (73\eta^2 I + 5832\eta^4 L^2 I^3 P^2 \rho^2 + 63\eta^2 IP\rho^2 + 6480\eta^4 L^2 I^3 P^2) \sigma^2 \\
&\quad + (9720\eta^4 L^2 I^4 P^3 + 105\eta^2 I^2 P^2) \mathbb{E} \left[\|\nabla f(\bar{x}_{r_0})\|^2 \middle| \mathcal{Q} \right] \\
&\leq 36\eta^2 L^2 I^2 \sum_{i=1}^N \bar{q}_r^i S_{r_0}^i + 109\eta^2 L^2 I^2 P^2 \sum_{i=1}^N \left(\bar{q}_{r_0}^i + \frac{1}{N} \right) S_{r_0}^i \\
&\quad + (75\eta^2 I + 65\eta^2 IP\rho^2) \sigma^2 + 108\eta^2 I^2 P^2 \mathbb{E} \left[\|\nabla f(\bar{x}_{r_0})\|^2 \middle| \mathcal{Q} \right],
\end{aligned}$$

where the last inequality uses $\eta \leq \frac{1}{60LIP}$. This proves Equation 3. Equation 4 follows by summing over $r \in \{r_0, \dots, r_0 + P - 1\}$ and $k \in \{0, \dots, I - 1\}$, taking total expectation, and applying the condition $\mathbb{E}_{r_0}[\bar{q}_{r_0}^i] = \frac{1}{N}$. \square

Lemma 2. If $\eta \leq \frac{1}{60LIP}$, $\gamma\eta \leq \frac{1}{60LIP}$, $\mathbb{E}_{r_0}[\bar{q}_{r_0}^i] = \frac{1}{N}$, and $\mathbb{E} \left[\sum_{i=1}^N (v_{r_0}^i)^2 \Lambda_{r_0}^i \right] \leq \rho^2$, then

$$\begin{aligned}
\mathbb{E} \left[\|\bar{x}_{r_0+P} - \bar{x}_{r_0}\|^2 \right] &\leq 6\gamma^2 \eta^2 I^2 P^2 \mathbb{E} \left[\|\nabla f(\bar{x}_{r_0})\|^2 \right] + \gamma\eta IP (5\gamma\eta\rho^2 + 7\eta^2 LI) \sigma^2 \\
&\quad + 11\gamma^2 \eta^2 L^2 I^2 P^2 \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0}^i]. \tag{23}
\end{aligned}$$

Proof. By the algorithm definition,

$$\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0} = -\gamma\eta \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N \sum_{k=0}^{I-1} q_r^i (\nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) - \mathbf{G}_{r_0}^i + \mathbf{G}_{r_0}) = -\gamma\eta \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N \sum_{k=0}^{I-1} q_r^i \mathbf{g}_{r,k}^i.$$

To obtain the variance of the update $\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}$,

$$\begin{aligned} & \mathbb{E} \left[\left\| \sum_{r=r_0}^{r_0+P} \sum_{i=1}^N \sum_{k=0}^{I-1} q_r^i (\mathbf{g}_{r,k}^i - \bar{\mathbf{g}}_{r,k}^i) \right\|^2 \right] \\ &= \mathbb{E} \left[\left\| \sum_{r=r_0}^{r_0+P} \sum_{i=1}^N \sum_{k=0}^{I-1} q_r^i ((\nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) - \nabla f_i(\mathbf{x}_{r,k}^i)) - (\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i) + (\mathbf{G}_{r_0} - \bar{\mathbf{G}}_{r_0})) \right\|^2 \right] \\ &\leq 2 \mathbb{E} \underbrace{\left[\left\| \sum_{r=r_0}^{r_0+P} \sum_{i=1}^N \sum_{k=0}^{I-1} q_r^i (\nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) - \nabla f_i(\mathbf{x}_{r,k}^i)) \right\|^2 \right]}_{A_1} \\ &\quad + 2 \mathbb{E} \underbrace{\left[\left\| \sum_{r=r_0}^{r_0+P} \sum_{i=1}^N \sum_{k=0}^{I-1} q_r^i (-\mathbf{G}_{r_0}^i + \bar{\mathbf{G}}_{r_0}^i + \mathbf{G}_{r_0} - \bar{\mathbf{G}}_{r_0}) \right\|^2 \right]}_{A_2}. \end{aligned} \tag{24}$$

We can bound the two terms A_1 and A_2 separately as follows. For A_1 ,

$$\begin{aligned} A_1 &= 2\mathbb{E} \left[\left\| \sum_{r=r_0}^{r_0+P} \sum_{i=1}^N \sum_{k=0}^{I-1} q_r^i (\nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) - \nabla f_i(\mathbf{x}_{r,k}^i)) \right\|^2 \right] \\ &= 2\mathbb{E} \left[\mathbb{E} \left[\left\| \sum_{r=r_0}^{r_0+P} \sum_{i=1}^N \sum_{k=0}^{I-1} q_r^i (\nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) - \nabla f_i(\mathbf{x}_{r,k}^i)) \right\|^2 \middle| \mathcal{Q} \right] \right] \\ &\stackrel{(i)}{=} 2\mathbb{E} \left[\sum_{r=r_0}^{r_0+P} \sum_{i=1}^N \sum_{k=0}^{I-1} \mathbb{E} \left[\|q_r^i (\nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) - \nabla f_i(\mathbf{x}_{r,k}^i))\|^2 \middle| \mathcal{Q} \right] \right] \\ &= 2 \sum_{r=r_0}^{r_0+P} \sum_{i=1}^N \sum_{k=0}^{I-1} \mathbb{E} \left[(q_r^i)^2 \mathbb{E} \left[\|\nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) - \nabla f_i(\mathbf{x}_{r,k}^i)\|^2 \middle| \mathcal{Q} \right] \right] \\ &\leq 2\sigma^2 \sum_{r=r_0}^{r_0+P} \sum_{i=1}^N \sum_{k=0}^{I-1} \mathbb{E} \left[(q_r^i)^2 \right] \\ &\leq 2IP\rho^2\sigma^2, \end{aligned} \tag{25}$$

where (i) uses the fact that for each i , $\left\{ q_r^i (\nabla F_i(\mathbf{x}_{r,k}^i) - \nabla f_i(\mathbf{x}_{r,k}^i)) \right\}_{r,k}$ is a martingale difference sequence with respect to \mathcal{G} (when conditioned on \mathcal{Q}) and that stochastic gradient noise is independent

across clients. For A_2 ,

$$\begin{aligned}
& \sum_{r=r_0}^{r_0+P} \sum_{i=1}^N \sum_{k=0}^{I-1} q_r^i (-\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i + \mathbf{G}_{r_0} - \bar{\mathbf{G}}_{r_0}) \\
&= - \sum_{r=r_0}^{r_0+P} \sum_{i=1}^N \sum_{k=0}^{I-1} q_r^i (\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i) + \sum_{r=r_0}^{r_0+P} \sum_{i=1}^N \sum_{k=0}^{I-1} q_r^i (\mathbf{G}_{r_0} - \bar{\mathbf{G}}_{r_0}) \\
&= -I \sum_{r=r_0}^{r_0+P} \sum_{i=1}^N q_r^i (\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i) + IP (\mathbf{G}_{r_0} - \bar{\mathbf{G}}_{r_0}) \\
&= -IP \sum_{i=1}^N \left(\frac{1}{P} \sum_{r=r_0}^{r_0+P} q_r^i \right) (\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i) + IP \sum_{i=1}^N \frac{1}{N} (\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i) \\
&= -IP \sum_{i=1}^N \left(\bar{q}_{r_0}^i - \frac{1}{N} \right) (\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i),
\end{aligned}$$

so

$$\begin{aligned}
A_2 &= 2I^2 P^2 \mathbb{E} \left[\left\| \sum_{i=1}^N \left(\bar{q}_{r_0}^i - \frac{1}{N} \right) (\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i) \right\|^2 \right] \\
&= 2I^2 P^2 \mathbb{E} \left[\mathbb{E} \left[\left\| \sum_{i=1}^N \left(\bar{q}_{r_0}^i - \frac{1}{N} \right) (\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i) \right\|^2 \middle| \mathcal{Q} \right] \right] \\
&\stackrel{(i)}{=} 2I^2 P^2 \mathbb{E} \left[\sum_{i=1}^N \mathbb{E} \left[\left\| \left(\bar{q}_{r_0}^i - \frac{1}{N} \right) (\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i) \right\|^2 \middle| \mathcal{Q} \right] \right] \\
&= 2I^2 P^2 \mathbb{E} \left[\sum_{i=1}^N \left(\bar{q}_{r_0}^i - \frac{1}{N} \right)^2 \mathbb{E} \left[\|\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i\|^2 \middle| \mathcal{Q} \right] \right] \\
&= 2I^2 P^2 \mathbb{E} \left[\sum_{i=1}^N \left(\bar{q}_{r_0}^i - \frac{1}{N} \right)^2 \mathbb{E} \left[\left\| \frac{1}{P \bar{z}_{r_0-P}^i I} \sum_{s=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} z_s^i (\nabla F_i(\mathbf{y}_{s,k}^i; \zeta_{s,k}^i) - \nabla f_i(\mathbf{y}_{s,k}^i)) \right\|^2 \middle| \mathcal{Q} \right] \right] \\
&= 2I^2 P^2 \mathbb{E} \left[\sum_{i=1}^N \left(\bar{q}_{r_0}^i - \frac{1}{N} \right)^2 \sum_{s=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} \mathbb{E} \left[\left\| \frac{1}{P \bar{z}_{r_0-P}^i I} z_s^i (\nabla F_i(\mathbf{y}_{s,k}^i; \zeta_{s,k}^i) - \nabla f_i(\mathbf{y}_{s,k}^i)) \right\|^2 \middle| \mathcal{Q} \right] \right] \\
&= 2I^2 P^2 \mathbb{E} \left[\sum_{i=1}^N \left(\bar{q}_{r_0}^i - \frac{1}{N} \right)^2 \sum_{s=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} \frac{1}{P^2 (\bar{z}_{r_0-P}^i)^2 I^2} (z_s^i)^2 \mathbb{E} \left[\|\nabla F_i(\mathbf{y}_{s,k}^i; \zeta_{s,k}^i) - \nabla f_i(\mathbf{y}_{s,k}^i)\|^2 \middle| \mathcal{Q} \right] \right] \\
&= 2I^2 P^2 \frac{\sigma^2}{I} \mathbb{E} \left[\sum_{i=1}^N \left(\bar{q}_{r_0}^i - \frac{1}{N} \right)^2 \sum_{s=r_0-P}^{r_0-1} \frac{(z_s^i)^2}{P^2 (\bar{z}_{r_0-P}^i)^2} \right] \\
&= 2IP \sigma^2 \mathbb{E} \left[\sum_{i=1}^N \left(\bar{q}_{r_0}^i - \frac{1}{N} \right)^2 \frac{\frac{1}{P} \sum_{s=r_0-P}^{r_0-1} (z_s^i)^2}{\left(\frac{1}{P} \sum_{s=r_0-P}^{r_0-1} z_s^i \right)^2} \right] \leq 2IP \rho^2 \sigma^2,
\end{aligned} \tag{26}$$

where (i) uses the fact that, conditioned on \mathcal{Q} , the variables $\mathbf{G}_{r_0}^i - \bar{\mathbf{G}}_{r_0}^i$ depend only on stochastic gradient noise, which is independent across clients, and the last inequality uses the condition $\mathbb{E} \left[\sum_{i=1}^N (v_{r_0}^i)^2 \Lambda_{r_0}^i \right] \leq \rho^2$. Plugging Equation 25 and Equation 26 into Equation 24 yields

$$\mathbb{E} \left[\left\| \sum_{r=r_0}^{r_0+P} \sum_{i=1}^N \sum_{k=0}^{I-1} q_r^i (\mathbf{g}_{r,k}^i - \bar{\mathbf{g}}_{r,k}^i) \right\|^2 \right] \leq 4IP \rho^2 \sigma^2. \tag{27}$$

Therefore

$$\begin{aligned}
& \mathbb{E} \left[\|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2 \right] \\
&= \gamma^2 \eta^2 \mathbb{E} \left[\left\| \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N \sum_{k=0}^{I-1} q_r^i \mathbf{g}_{r,k}^i \right\|^2 \right] \\
&\leq \gamma^2 \eta^2 \mathbb{E} \left[\left\| \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N \sum_{k=0}^{I-1} q_r^i \bar{\mathbf{g}}_{r,k}^i \right\|^2 \right] + \gamma^2 \eta^2 \mathbb{E} \left[\left\| \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N \sum_{k=0}^{I-1} q_r^i (\mathbf{g}_{r,k}^i - \bar{\mathbf{g}}_{r,k}^i) \right\|^2 \right] \\
&\stackrel{(i)}{\leq} \gamma^2 \eta^2 \mathbb{E} \left[\left\| \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N \sum_{k=0}^{I-1} q_r^i \bar{\mathbf{g}}_{r,k}^i \right\|^2 \right] + 4\gamma^2 \eta^2 IP \rho^2 \sigma^2 \\
&\leq \gamma^2 \eta^2 IP \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N \sum_{k=0}^{I-1} \mathbb{E} \left[q_r^i \|\nabla f_i(\mathbf{x}_{r,k}^i) - \bar{\mathbf{G}}_{r_0}^i + \bar{\mathbf{G}}_{r_0}\|^2 \right] + 4\gamma^2 \eta^2 IP \rho^2 \sigma^2 \\
&\stackrel{(ii)}{\leq} 5\gamma^2 \eta^2 IP \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N \sum_{k=0}^{I-1} \mathbb{E} \left[q_r^i (\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 + \|\nabla f_i(\mathbf{x}_{r,k}^i) - \nabla f_i(\bar{\mathbf{x}}_{r,k})\|^2 \right. \\
&\quad \left. + \|\nabla f_i(\bar{\mathbf{x}}_{r,k}) - \nabla f_i(\bar{\mathbf{x}}_{r_0})\|^2 + \|\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}^i\|^2 + \|\nabla f(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}\|^2) \right] + 4\gamma^2 \eta^2 IP \rho^2 \sigma^2 \\
&\leq 5\gamma^2 \eta^2 I^2 P^2 \mathbb{E} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \right] \\
&\quad + 5\gamma^2 \eta^2 L^2 IP \sum_{r=r_0}^{r_0+P-1} \sum_{k=0}^{I-1} \sum_{i=1}^N \left(\mathbb{E} \left[q_r^i \|\mathbf{x}_{r,k}^i - \bar{\mathbf{x}}_{r,k}\|^2 \right] + \mathbb{E} \left[q_r^i \|\bar{\mathbf{x}}_{r,k} - \bar{\mathbf{x}}_{r_0}\|^2 \right] \right) \\
&\quad + 5\gamma^2 \eta^2 I^2 P^2 \mathbb{E} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}\|^2 + \sum_{i=1}^N \left(\frac{1}{P} \sum_{r=r_0}^{r_0+P-1} q_r^i \right) \|\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}^i\|^2 \right] \\
&\quad + 4\gamma^2 \eta^2 IP \rho^2 \sigma^2 \\
&\stackrel{(iii)}{\leq} 5\gamma^2 \eta^2 I^2 P^2 \mathbb{E} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \right] + 5\gamma^2 \eta^2 L^2 IP \mathbb{E} \left[\sum_{r=r_0}^{r_0+P-1} \sum_{k=0}^{I-1} D_{r,k} + M_{r,k} \right] \\
&\quad + 10\gamma^2 \eta^2 I^2 P^2 \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left[\|\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}^i\|^2 \right] + 4\gamma^2 \eta^2 IP \rho^2 \sigma^2, \tag{28}
\end{aligned}$$

where (i) uses Equation 27, (ii) uses the decomposition

$$\begin{aligned}
\nabla f_i(\mathbf{x}_{r,k}^i) - \bar{\mathbf{G}}_{r_0}^i + \bar{\mathbf{G}}_{r_0} &= \nabla f(\bar{\mathbf{x}}_{r_0}) + (\nabla f_i(\mathbf{x}_{r,k}^i) - \nabla f_i(\bar{\mathbf{x}}_{r,k})) + (\nabla f_i(\bar{\mathbf{x}}_{r,k}) - \nabla f_i(\bar{\mathbf{x}}_{r_0})) \\
&\quad + (\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}^i) + (\nabla f(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}),
\end{aligned}$$

and (iii) uses $\mathbf{G}_{r_0} = \frac{1}{N} \sum_{i=1}^N \mathbf{G}_{r_0}^i$ and $\mathbb{E}_{r_0} [\bar{q}_{r_0}^i] = \frac{1}{N}$ to obtain

$$\begin{aligned}
& \mathbb{E} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}\|^2 + \sum_{i=1}^N \left(\frac{1}{P} \sum_{r=r_0}^{r_0+P-1} q_r^i \right) \|\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}^i\|^2 \right] \\
&= \mathbb{E} \left[\mathbb{E}_{r_0} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}\|^2 + \sum_{i=1}^N \left(\frac{1}{P} \sum_{r=r_0}^{r_0+P-1} q_r^i \right) \|\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}^i\|^2 \right] \right] \\
&= \mathbb{E} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}\|^2 + \sum_{i=1}^N \mathbb{E}_{r_0} \left[\frac{1}{P} \sum_{r=r_0}^{r_0+P-1} q_r^i \right] \|\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}^i\|^2 \right] \\
&= \mathbb{E} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}\|^2 + \frac{1}{N} \sum_{i=1}^N \|\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}^i\|^2 \right] \\
&= \mathbb{E} \left[\left\| \frac{1}{N} \sum_{i=1}^N (\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}^i) \right\|^2 + \frac{1}{N} \sum_{i=1}^N \|\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}^i\|^2 \right] \\
&\leq \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \|\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}^i\|^2 + \frac{1}{N} \sum_{i=1}^N \|\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}^i\|^2 \right].
\end{aligned}$$

The remaining term in Equation 28 can be simplified as:

$$\begin{aligned}
\mathbb{E} [\|\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \bar{\mathbf{G}}_{r_0}^i\|^2] &\leq \mathbb{E} \left[\left\| \nabla f_i(\bar{\mathbf{x}}_{r_0}) - \frac{1}{IP} \sum_{r=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} \frac{z_r^i}{\bar{z}_{r_0-P}^i} \nabla f_i(\mathbf{y}_{r,k}^i) \right\|^2 \right] \\
&\leq \mathbb{E} \left[\left\| \frac{1}{IP} \sum_{r=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} \frac{z_r^i}{\bar{z}_{r_0-P}^i} (\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \nabla f_i(\mathbf{y}_{r,k}^i)) \right\|^2 \right] \\
&\leq \frac{1}{IP} \sum_{r=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} \mathbb{E} \left[\frac{z_r^i}{\bar{z}_{r_0-P}^i} \|\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \nabla f_i(\mathbf{y}_{r,k}^i)\|^2 \right] \\
&\leq \frac{L^2}{IP} \sum_{r=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} \mathbb{E} \left[\frac{z_r^i}{\bar{z}_{r_0-P}^i} \|\bar{\mathbf{x}}_{r_0} - \mathbf{y}_{r,k}^i\|^2 \right] \\
&= L^2 \mathbb{E} \left[\frac{1}{IP} \sum_{r=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} \frac{z_r^i}{\bar{z}_{r_0-P}^i} \mathbb{E} [\|\bar{\mathbf{x}}_{r_0} - \mathbf{y}_{r,k}^i\|^2 | \mathcal{Q}] \right] \\
&= L^2 \mathbb{E} [S_{r_0}^i].
\end{aligned}$$

Plugging back to Equation 28 yields

$$\begin{aligned}
& \mathbb{E} \left[\|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2 \right] \\
& \leq 5\gamma^2\eta^2I^2P^2\mathbb{E} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \right] + 4\gamma^2\eta^2IP\rho^2\sigma^2 \\
& \quad + 5\gamma^2\eta^2L^2IP\mathbb{E} \left[\sum_{r=r_0}^{r_0+P-1} \sum_{k=0}^{I-1} (D_{r,k} + M_{r,k}) \right] + 10\gamma^2\eta^2L^2I^2P^2 \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0}^i] \\
& \stackrel{(i)}{\leq} 5\gamma^2\eta^2I^2P^2\mathbb{E} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \right] + 4\gamma^2\eta^2IP\rho^2\sigma^2 \\
& \quad + 5\gamma^2\eta^2L^2IP \left(108\eta^2I^3P^3\mathbb{E} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \right] + (75\eta^2I^2P + 65\eta^2I^2P^2\rho^2) \sigma^2 \right. \\
& \quad \left. + 254\eta^2L^2I^3P^3 \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0}^i] \right) + 10\gamma^2\eta^2L^2I^2P^2 \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0}^i] \\
& \leq 5\gamma^2\eta^2I^2P^2 (1 + 108\eta^2L^2I^2P^2) \mathbb{E} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \right] \\
& \quad + \gamma\eta IP (4\gamma\eta\rho^2 + 375\gamma\eta^3L^2I^2P + 325\gamma\eta^3L^2I^2P^2\rho^2) \sigma^2 \\
& \quad + 10\gamma^2\eta^2L^2I^2P^2 (1 + 127\eta^2L^2I^2P^2) \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0}^i] \\
& \stackrel{(ii)}{\leq} 6\gamma^2\eta^2I^2P^2\mathbb{E} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \right] + \gamma\eta IP (5\gamma\eta\rho^2 + 7\eta^2LI) \sigma^2 + 11\gamma^2\eta^2L^2I^2P^2 \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0}^i].
\end{aligned}$$

where (i) uses Equation 4 from Lemma 1 and (ii) uses $\eta \leq \frac{1}{60LIP}$ and $\gamma\eta \leq \frac{1}{60LIP}$. This proves Equation 23. \square

Lemma 3. Suppose that $P_{\mathcal{Q}_{r_0}}(\bar{q}_{r_0}^i > 0) \geq p_{sample}$, $\mathbb{E} [w_{r_0}^i | \mathcal{Q}_{r_0}] \leq \frac{P^2}{N}$, $\mathbb{E}_{r_0}[\bar{q}_{r_0}^i] = \frac{1}{N}$, and $\mathbb{E} \left[\sum_{i=1}^N (v_{r_0}^i)^2 \Lambda_{r_0}^i \right] \leq \rho^2$. If $\eta \leq \frac{\sqrt{p_{sample}}}{60LIP}$ and $\gamma\eta \leq \frac{p_{sample}}{60LIP}$, then

$$\begin{aligned}
\frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0+P}^i] & \leq \left(324\eta^2I^2P^2 + \frac{20}{p_{sample}}\gamma^2\eta^2I^2P^2 \right) \mathbb{E} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \right] \\
& \quad + \left(226\eta^2I + 195\eta^2IP\rho^2 + \frac{17}{p_{sample}}\gamma^2\eta^2IP\rho^2 \right) \sigma^2 \\
& \quad + \left(1 - \frac{1}{2}p_{sample} \right) \frac{1}{N} \sum_{j=1}^N \mathbb{E} [S_{r_0}^j].
\end{aligned}$$

Proof. Let $1 \leq i \leq N$. We can consider the value $S_{r_0+P}^i$ under two cases: $\bar{q}_{r_0}^i > 0$ and the complement. Let $A_{r_0}^i = \{\bar{q}_{r_0}^i > 0\}$. Denote

$$\begin{aligned}
B_1^i & = \mathbb{1} \{ A_{r_0}^i \} \frac{1}{IP\bar{q}_{r_0}^i} \sum_{s=r_0}^{r_0+P-1} \sum_{k=0}^{I-1} q_s^i \mathbb{E} \left[\|\mathbf{x}_{s,k}^i - \bar{\mathbf{x}}_{r_0+P}\|^2 \middle| \mathcal{Q} \right] \\
B_2^i & = \mathbb{1} \{ \bar{A}_{r_0}^i \} \frac{1}{IP\bar{z}_{r_0}^i} \sum_{s=r_0}^{r_0+P-1} \sum_{k=0}^{I-1} z_s^i \mathbb{E} \left[\|\mathbf{y}_{s,k}^i - \bar{\mathbf{x}}_{r_0+P}\|^2 \middle| \mathcal{Q} \right].
\end{aligned}$$

Then $S_{r_0+P}^i = B_1^i + B_2^i$, and we can consider the two cases separately.

Notice that

$$\|\mathbf{x}_{s,k}^i - \bar{\mathbf{x}}_{r_0+P}\|^2 \leq 3\|\mathbf{x}_{s,k}^i - \bar{\mathbf{x}}_{s,k}\|^2 + 3\|\bar{\mathbf{x}}_{s,k} - \bar{\mathbf{x}}_{r_0}\|^2 + 3\|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2.$$

Therefore

$$\begin{aligned}
B_1^i &= \mathbb{1}\{A_{r_0}^i\} \frac{1}{IP\bar{q}_{r_0}^i} \sum_{s=r_0}^{r_0+P-1} \sum_{k=0}^{I-1} q_s^i \mathbb{E} \left[\|\mathbf{x}_{s,k}^i - \bar{\mathbf{x}}_{r_0+P}\|^2 \middle| \mathcal{Q} \right] \\
&\leq \mathbb{1}\{A_{r_0}^i\} \frac{3}{IP\bar{q}_{r_0}^i} \sum_{s=r_0}^{r_0+P-1} \sum_{k=0}^{I-1} q_s^i \left(\mathbb{E} [\|\mathbf{x}_{s,k}^i - \bar{\mathbf{x}}_{s,k}\|^2 \middle| \mathcal{Q}] + \mathbb{E} [\|\bar{\mathbf{x}}_{s,k} - \bar{\mathbf{x}}_{r_0}\|^2 \middle| \mathcal{Q}] \right. \\
&\quad \left. + \mathbb{E} [\|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2 \middle| \mathcal{Q}] \right) \\
&\leq \mathbb{1}\{A_{r_0}^i\} \frac{3}{IP\bar{q}_{r_0}^i} \sum_{s=r_0}^{r_0+P-1} \sum_{k=0}^{I-1} q_s^i (D_{s,k} + M_{s,k}) + 3\mathbb{1}\{A_{r_0}^i\} \mathbb{E} [\|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2 \middle| \mathcal{Q}].
\end{aligned}$$

Using Equation 3 from Lemma 1

$$\begin{aligned}
&\mathbb{1}\{A_{r_0}^i\} \frac{3}{IP\bar{q}_{r_0}^i} \sum_{s=r_0}^{r_0+P-1} \sum_{k=0}^{I-1} q_s^i (D_{s,k} + M_{s,k}) \\
&\leq \mathbb{1}\{A_{r_0}^i\} \frac{3}{IP\bar{q}_{r_0}^i} \sum_{s=r_0}^{r_0+P-1} \sum_{k=0}^{I-1} q_s^i \left(108\eta^2 I^2 P^2 \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \middle| \mathcal{Q}] + (75\eta^2 I + 65\eta^2 IP\rho^2) \sigma^2 \right. \\
&\quad \left. + 109\eta^2 L^2 I^2 P^2 \sum_{j=1}^N \left(\bar{q}_{r_0}^j + \frac{1}{N} \right) S_{r_0}^j + 36\eta^2 L^2 I^2 \sum_{j=1}^N q_s^j S_{r_0}^j \right) \\
&\leq 324\eta^2 I^2 P^2 \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \middle| \mathcal{Q}] + (225\eta^2 I + 195\eta^2 IP\rho^2) \sigma^2 \\
&\quad + 327\eta^2 L^2 I^2 P^2 \sum_{j=1}^N \left(\bar{q}_{r_0}^j + \frac{1}{N} \right) S_{r_0}^j + \mathbb{1}\{A_{r_0}^i\} \frac{3}{P\bar{q}_{r_0}^i} \sum_{s=r_0}^{r_0+P-1} q_s^i \left(36\eta^2 L^2 I^2 \sum_{j=1}^N q_s^j S_{r_0}^j \right) \\
&= 324\eta^2 I^2 P^2 \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \middle| \mathcal{Q}] + (225\eta^2 I + 195\eta^2 IP\rho^2) \sigma^2 \\
&\quad + 327\eta^2 L^2 I^2 P^2 \sum_{j=1}^N \left(\bar{q}_{r_0}^j + \frac{1}{N} \right) S_{r_0}^j + 108\eta^2 L^2 I^2 \sum_{j=1}^N \left(\frac{\mathbb{1}\{A_{r_0}^i\}}{P\bar{q}_{r_0}^i} \sum_{s=r_0}^{r_0+P-1} q_s^i q_s^j \right) S_{r_0}^j.
\end{aligned}$$

Denote $w_{r_0}^{i,j} = \frac{\mathbb{1}\{A_{r_0}^i\}}{P\bar{q}_{r_0}^i} \sum_{s=r_0}^{r_0+P-1} q_s^i q_s^j$, so that $w_{r_0}^i = \frac{1}{N} \sum_{j=1}^N w_{r_0}^{i,j}$. Then

$$\begin{aligned}
B_1^i &\leq 324\eta^2 I^2 P^2 \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \middle| \mathcal{Q}] + (225\eta^2 I + 195\eta^2 IP\rho^2) \sigma^2 \\
&\quad + 327\eta^2 L^2 I^2 P^2 \sum_{j=1}^N \left(\bar{q}_{r_0}^j + \frac{1}{N} \right) S_{r_0}^j + 108\eta^2 L^2 I^2 \sum_{j=1}^N \left(\frac{\mathbb{1}\{A_{r_0}^i\}}{P\bar{q}_{r_0}^i} \sum_{s=r_0}^{r_0+P-1} q_s^i q_s^j \right) S_{r_0}^j \\
&\quad + 3\mathbb{1}\{A_{r_0}^i\} \mathbb{E} [\|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2 \middle| \mathcal{Q}] \\
\mathbb{E} [B_1^i \mid \mathcal{Q}_{:r_0}] &\stackrel{(i)}{\leq} 324\eta^2 I^2 P^2 \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \mid \mathcal{Q}_{:r_0}] + (225\eta^2 I + 195\eta^2 IP\rho^2) \sigma^2 \\
&\quad + 654\eta^2 L^2 I^2 P^2 \frac{1}{N} \sum_{j=1}^N S_{r_0}^j + 108\eta^2 L^2 I^2 \sum_{j=1}^N \mathbb{E} \left[\left. \frac{\mathbb{1}\{A_{r_0}^i\}}{P\bar{q}_{r_0}^i} \sum_{s=r_0}^{r_0+P-1} q_s^i q_s^j \right| \mathcal{Q}_{:r_0} \right] S_{r_0}^j \\
&\quad + 3\mathbb{E} [\mathbb{1}\{A_{r_0}^i\} \|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2 \mid \mathcal{Q}_{:r_0}] \\
\mathbb{E} [B_1^i] &\stackrel{(ii)}{\leq} 324\eta^2 I^2 P^2 \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2] + (225\eta^2 I + 195\eta^2 IP\rho^2) \sigma^2 \\
&\quad + 654\eta^2 L^2 I^2 P^2 \frac{1}{N} \sum_{j=1}^N \mathbb{E} [S_{r_0}^j] + 108\eta^2 L^2 I^2 \sum_{j=1}^N \mathbb{E} [\mathbb{E} [w_{r_0}^{i,j} \mid \mathcal{Q}_{:r_0}] S_{r_0}^j] \\
&\quad + 3\mathbb{E} [\mathbb{1}\{A_{r_0}^i\} \|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2], \tag{29}
\end{aligned}$$

where (i) uses $\mathbb{E}_{r_0}[\bar{q}_{r_0}^i] = \frac{1}{N}$ and the tower property $\mathbb{E}[\mathbb{E}[\cdot | \mathcal{Q}] | \mathcal{Q}_{:r_0}] = \mathbb{E}[\cdot | \mathcal{Q}_{:r_0}]$, and (ii) uses the tower property $\mathbb{E}[\mathbb{E}[\cdot | \mathcal{Q}_{:r_0}]] = \mathbb{E}[\cdot]$.

Now consider B_2^i . Under $\bar{A}_{r_0}^i$, $z_r^i = z_{r-P}^i$ and $\mathbf{y}_{r,k}^i = \mathbf{y}_{r-P,k}^i$ for all $r \in \{r_0, \dots, r_0 + P - 1\}$. Therefore

$$\begin{aligned}
B_2^i &= \mathbb{1}\{\bar{A}_{r_0}^i\} \frac{1}{IP\bar{z}_{r_0}^i} \sum_{s=r_0}^{r_0+P-1} \sum_{k=0}^{I-1} z_s^i \mathbb{E}\left[\|\mathbf{y}_{s,k}^i - \bar{\mathbf{x}}_{r_0+P}\|^2 \middle| \mathcal{Q}\right] \\
&= \mathbb{1}\{\bar{A}_{r_0}^i\} \frac{1}{IP\bar{z}_{r_0-P}^i} \sum_{s=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} z_s^i \mathbb{E}\left[\|\mathbf{y}_{s,k}^i - \bar{\mathbf{x}}_{r_0+P}\|^2 \middle| \mathcal{Q}\right] \\
&\stackrel{(i)}{\leq} \mathbb{1}\{\bar{A}_{r_0}^i\} \frac{1}{IP\bar{z}_{r_0-P}^i} \sum_{s=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} z_s^i \left((1+\lambda) \mathbb{E}\left[\|\mathbf{y}_{s,k}^i - \bar{\mathbf{x}}_{r_0}\|^2 \middle| \mathcal{Q}\right] \right. \\
&\quad \left. + \left(1 + \frac{1}{\lambda}\right) \mathbb{E}\left[\|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2 \middle| \mathcal{Q}\right] \right) \\
&= \mathbb{1}\{\bar{A}_{r_0}^i\} (1+\lambda) \frac{1}{IP\bar{z}_{r_0-P}^i} \sum_{s=r_0-P}^{r_0-1} \sum_{k=0}^{I-1} z_s^i \mathbb{E}\left[\|\mathbf{y}_{s,k}^i - \bar{\mathbf{x}}_{r_0}\|^2 \middle| \mathcal{Q}\right] \\
&\quad + \mathbb{1}\{\bar{A}_{r_0}^i\} \left(1 + \frac{1}{\lambda}\right) \mathbb{E}\left[\|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2 \middle| \mathcal{Q}\right] \\
&= \mathbb{1}\{\bar{A}_{r_0}^i\} (1+\lambda) S_{r_0}^i + \mathbb{1}\{\bar{A}_{r_0}^i\} \left(1 + \frac{1}{\lambda}\right) \mathbb{E}\left[\|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2 \middle| \mathcal{Q}\right]
\end{aligned}$$

where (i) uses Young's inequality with an arbitrary $\lambda > 0$. Taking conditional expectation $\mathbb{E}[\cdot | \mathcal{Q}_{:r_0}]$ followed by total expectation yields

$$\begin{aligned}
\mathbb{E}[B_2^i | \mathcal{Q}_{:r_0}] &\leq (1 - p_{\text{sample}}) (1 + \lambda) S_{r_0}^i + \left(1 + \frac{1}{\lambda}\right) \mathbb{E}\left[\mathbb{1}\{\bar{A}_{r_0}^i\} \|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2 \middle| \mathcal{Q}_{:r_0}\right] \\
\mathbb{E}[B_2^i] &\leq (1 - p_{\text{sample}}) (1 + \lambda) \mathbb{E}[S_{r_0}^i] + \left(1 + \frac{1}{\lambda}\right) \mathbb{E}\left[\mathbb{1}\{\bar{A}_{r_0}^i\} \|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2\right]. \quad (30)
\end{aligned}$$

Adding Equation 29 and Equation 30 yields

$$\begin{aligned}
\mathbb{E}[S_{r_0+P}^i] &= \mathbb{E}[B_1^i] + \mathbb{E}[B_2^i] \\
&\leq 324\eta^2 I^2 P^2 \mathbb{E}\left[\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2\right] + (225\eta^2 I + 195\eta^2 IP\rho^2) \sigma^2 \\
&\quad + 654\eta^2 L^2 I^2 P^2 \frac{1}{N} \sum_{j=1}^N \mathbb{E}[S_{r_0}^j] + 108\eta^2 L^2 I^2 \sum_{j=1}^N \mathbb{E}\left[\mathbb{E}[w_{r_0}^{i,j} | \mathcal{Q}_{:r_0}] S_{r_0}^j\right] \\
&\quad + (1 - p_{\text{sample}}) (1 + \lambda) \mathbb{E}[S_{r_0}^i] + \left(3 + \frac{1}{\lambda}\right) \mathbb{E}\left[\|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2\right].
\end{aligned}$$

Averaging over i ,

$$\begin{aligned}
& \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0+P}^i] \\
& \leq 324\eta^2 I^2 P^2 \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2] + (225\eta^2 I + 195\eta^2 I P \rho^2) \sigma^2 + 654\eta^2 L^2 I^2 P^2 \frac{1}{N} \sum_{j=1}^N \mathbb{E} [S_{r_0}^j] \\
& \quad + 108\eta^2 L^2 I^2 \sum_{j=1}^N \mathbb{E} \left[\mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N w_{r_0}^{i,j} \mid \mathcal{Q}_{:r_0} \right] S_{r_0}^j \right] \\
& \quad + (1 - p_{\text{sample}}) (1 + \lambda) \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0}^i] + \left(3 + \frac{1}{\lambda} \right) \mathbb{E} [\|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2] \\
& \leq 324\eta^2 I^2 P^2 \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2] + (225\eta^2 I + 195\eta^2 I P \rho^2) \sigma^2 \\
& \quad + ((1 - p_{\text{sample}}) (1 + \lambda) + 654\eta^2 L^2 I^2 P^2) \frac{1}{N} \sum_{j=1}^N \mathbb{E} [S_{r_0}^j] \\
& \quad + 108\eta^2 L^2 I^2 \sum_{j=1}^N \mathbb{E} [\mathbb{E} [w_{r_0}^j \mid \mathcal{Q}_{:r_0}] S_{r_0}^j] + \left(3 + \frac{1}{\lambda} \right) \mathbb{E} [\|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2] \\
& \leq 324\eta^2 I^2 P^2 \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2] + (225\eta^2 I + 195\eta^2 I P \rho^2) \sigma^2 \\
& \quad + ((1 - p_{\text{sample}}) (1 + \lambda) + 762\eta^2 L^2 I^2 P^2) \frac{1}{N} \sum_{j=1}^N \mathbb{E} [S_{r_0}^j] + \left(3 + \frac{1}{\lambda} \right) \mathbb{E} [\|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2],
\end{aligned}$$

where the last inequality uses the condition $\mathbb{E} [w_{r_0}^i \mid \mathcal{Q}_{:r_0}] \leq \frac{P^2}{N}$. We can then apply Lemma 2 and choose $\lambda = \frac{3p_{\text{sample}}}{10(1-p_{\text{sample}})}$ to obtain

$$\begin{aligned}
& \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0+P}^i] \\
& \leq \left(324\eta^2 I^2 P^2 + \left(3 + \frac{1}{\lambda} \right) 6\gamma^2 \eta^2 I^2 P^2 \right) \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2] \\
& \quad + \left(225\eta^2 I + 195\eta^2 I P \rho^2 + \left(3 + \frac{1}{\lambda} \right) 5\gamma^2 \eta^2 I P \rho^2 + \left(3 + \frac{1}{\lambda} \right) 7\gamma \eta^3 L I^2 P \right) \sigma^2 \\
& \quad + \left((1 - p_{\text{sample}}) (1 + \lambda) + 762\eta^2 L^2 I^2 P^2 + \left(3 + \frac{1}{\lambda} \right) 11\gamma^2 \eta^2 L^2 I^2 P^2 \right) \frac{1}{N} \sum_{j=1}^N \mathbb{E} [S_{r_0}^j] \\
& \leq \left(324\eta^2 I^2 P^2 + \frac{20}{p_{\text{sample}}} \gamma^2 \eta^2 I^2 P^2 \right) \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2] \\
& \quad + \left(225\eta^2 I + 195\eta^2 I P \rho^2 + \frac{17}{p_{\text{sample}}} \gamma^2 \eta^2 I P \rho^2 + \frac{23}{p_{\text{sample}}} \gamma \eta^3 L I^2 P \right) \sigma^2 \\
& \quad + \left(1 - \frac{7}{10} p_{\text{sample}} + 762\eta^2 L^2 I^2 P^2 + \frac{110}{3p_{\text{sample}}} \gamma^2 \eta^2 L^2 I^2 P^2 \right) \frac{1}{N} \sum_{j=1}^N \mathbb{E} [S_{r_0}^j] \\
& \leq \left(324\eta^2 I^2 P^2 + \frac{20}{p_{\text{sample}}} \gamma^2 \eta^2 I^2 P^2 \right) \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2] \\
& \quad + \left(226\eta^2 I + 195\eta^2 I P \rho^2 + \frac{17}{p_{\text{sample}}} \gamma^2 \eta^2 I P \rho^2 \right) \sigma^2 + \left(1 - \frac{1}{2} p_{\text{sample}} \right) \frac{1}{N} \sum_{j=1}^N \mathbb{E} [S_{r_0}^j],
\end{aligned}$$

where the last inequality uses the conditions $\eta \leq \frac{\sqrt{p_{\text{sample}}}}{60LIP}$ and $\gamma\eta \leq \frac{p_{\text{sample}}}{60LIP}$, and we used that $3 + \frac{1}{\lambda} \leq \frac{10}{3p_{\text{sample}}}$. \square

Theorem 2 (Theorem 1 restated). *Suppose Assumptions 1 and 2 hold, and $\mathbb{E}[w_{r_0}^i | \mathcal{Q}_{:r_0}] \leq \frac{P^2}{N}$, and $\mathbb{E} \left[\sum_{i=1}^N (v_{r_0}^i)^2 \Lambda_{r_0}^i \right] \leq \rho^2$. If*

$$\gamma\eta \leq \frac{p_{\text{sample}}}{60LIP}, \quad \eta \leq \frac{\sqrt{p_{\text{sample}}}}{60LIP},$$

then Algorithm 1 satisfies

$$\frac{P}{R} \sum_{r_0 \in \{0, P, \dots, R-P\}} \mathbb{E}[\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2] \leq \frac{5\Delta}{\gamma\eta IR} + (20\gamma\eta L\rho^2 + 5785\eta^2 L^2 IP) \sigma^2.$$

Proof. Recall that

$$\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0} = -\gamma\eta \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N q_r^i \sum_{k=0}^{I-1} \nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) + \gamma\eta IP \sum_{i=1}^N \left(\bar{q}_{r_0}^i - \frac{1}{N} \right) \mathbf{G}_{r_0}^i \quad (31)$$

Using the quadratic upper bound for smooth functions and taking total expectation,

$$\begin{aligned} \mathbb{E}[f(\bar{\mathbf{x}}_{r_0+P}) - f(\bar{\mathbf{x}}_{r_0})] &\leq -\gamma\eta \mathbb{E} \left[\left\langle \nabla f(\bar{\mathbf{x}}_{r_0}), \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N q_r^i \sum_{k=0}^{I-1} \nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) \right\rangle \right] \\ &\quad + \gamma\eta IP \mathbb{E} \left[\left\langle \nabla f(\bar{\mathbf{x}}_{r_0}), \sum_{i=1}^N \left(\bar{q}_{r_0}^i - \frac{1}{N} \right) \mathbf{G}_{r_0}^i \right\rangle \right] \\ &\quad + \frac{L}{2} \mathbb{E} [\|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2] \\ &\stackrel{(i)}{\leq} -\gamma\eta \mathbb{E} \left[\left\langle \nabla f(\bar{\mathbf{x}}_{r_0}), \mathbb{E}_{r_0} \left[\sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N q_r^i \sum_{k=0}^{I-1} \nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) \right] \right\rangle \right] \\ &\quad + \gamma\eta IP \mathbb{E} \left[\left\langle \nabla f(\bar{\mathbf{x}}_{r_0}), \sum_{i=1}^N \left(\mathbb{E}_{r_0} [\bar{q}_{r_0}^i] - \frac{1}{N} \right) \mathbf{G}_{r_0}^i \right\rangle \right] \\ &\quad + \frac{L}{2} \mathbb{E} [\|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2] \\ &\stackrel{(ii)}{\leq} -\gamma\eta \mathbb{E} \left[\left\langle \nabla f(\bar{\mathbf{x}}_{r_0}), \mathbb{E}_{r_0} \left[\sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N q_r^i \sum_{k=0}^{I-1} \nabla f_i(\mathbf{x}_{r,k}^i) \right] \right\rangle \right] \\ &\quad + \frac{L}{2} \mathbb{E} [\|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2], \end{aligned} \quad (32)$$

where (i) uses the law of total expectation $\mathbb{E}[\cdot] = \mathbb{E}[\mathbb{E}_{r_0}[\cdot]]$, and (ii) uses the condition $\mathbb{E}_{r_0}[\bar{q}_{r_0}^i] = \frac{1}{N}$ together with

$$\begin{aligned}
\mathbb{E}_{r_0} \left[\sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N q_r^i \sum_{k=0}^{I-1} \nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) \right] &= \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N \sum_{k=0}^{I-1} \mathbb{E}_{r_0} [q_r^i \nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i)] \\
&= \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N \sum_{k=0}^{I-1} \mathbb{E}_{r_0} [\mathbb{E} [q_r^i \nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) | \mathbf{x}_{r,k}^i]] \\
&= \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N \sum_{k=0}^{I-1} \mathbb{E}_{r_0} [q_r^i \mathbb{E} [\nabla F_i(\mathbf{x}_{r,k}^i; \xi_{r,k}^i) | \mathbf{x}_{r,k}^i]] \\
&= \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N \sum_{k=0}^{I-1} \mathbb{E}_{r_0} [q_r^i \nabla f_i(\mathbf{x}_{r,k}^i)] \\
&= \mathbb{E}_{r_0} \left[\sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N q_r^i \sum_{k=0}^{I-1} \nabla f_i(\mathbf{x}_{r,k}^i) \right].
\end{aligned}$$

Focusing on the inner product term of Equation 32:

$$\begin{aligned}
& -\gamma\eta \mathbb{E} \left[\left\langle \nabla f(\bar{\mathbf{x}}_{r_0}), \mathbb{E}_{r_0} \left[\sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N q_r^i \sum_{k=0}^{I-1} \nabla f_i(\mathbf{x}_{r,k}^i) \right] \right\rangle \right] \\
&= -\frac{\gamma\eta}{IP} \mathbb{E} \left[\left\langle IP \nabla f(\bar{\mathbf{x}}_{r_0}), \mathbb{E}_{r_0} \left[\sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N q_r^i \sum_{k=0}^{I-1} \nabla f_i(\mathbf{x}_{r,k}^i) \right] \right\rangle \right] \\
&\leq -\frac{\gamma\eta IP}{2} \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2] + \frac{\gamma\eta}{2IP} \mathbb{E} \left[\left\| \mathbb{E}_{r_0} \left[\sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N q_r^i \sum_{k=0}^{I-1} \nabla f_i(\mathbf{x}_{r,k}^i) \right] - IP \nabla f(\bar{\mathbf{x}}_{r_0}) \right\|^2 \right], \tag{33}
\end{aligned}$$

where we used $-\langle a, b \rangle = \frac{1}{2}\|b - a\|^2 - \frac{1}{2}\|a\|^2 - \frac{1}{2}\|b\|^2 \leq \frac{1}{2}\|b - a\|^2 - \frac{1}{2}\|a\|^2$. Also

$$\begin{aligned}
& \mathbb{E} \left[\left\| \mathbb{E}_{r_0} \left[\sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N q_r^i \sum_{k=0}^{I-1} \nabla f_i(\mathbf{x}_{r,k}^i) \right] - IP \nabla f(\bar{\mathbf{x}}_{r_0}) \right\|^2 \right] \\
&= \mathbb{E} \left[\left\| \mathbb{E}_{r_0} \left[\sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N q_r^i \sum_{k=0}^{I-1} (\nabla f_i(\mathbf{x}_{r,k}^i) - \nabla f(\bar{\mathbf{x}}_{r_0})) \right] \right\|^2 \right] \\
&\stackrel{(i)}{\leq} 3 \mathbb{E} \left[\left\| \mathbb{E}_{r_0} \left[\sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N q_r^i \sum_{k=0}^{I-1} (\nabla f_i(\mathbf{x}_{r,k}^i) - \nabla f_i(\bar{\mathbf{x}}_{r,k})) \right] \right\|^2 \right] \\
&\quad + 3 \mathbb{E} \left[\left\| \mathbb{E}_{r_0} \left[\sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N q_r^i \sum_{k=0}^{I-1} (\nabla f_i(\bar{\mathbf{x}}_{r,k}) - \nabla f_i(\bar{\mathbf{x}}_{r_0})) \right] \right\|^2 \right] \\
&\quad + 3 \mathbb{E} \left[\left\| \mathbb{E}_{r_0} \left[\sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N q_r^i \sum_{k=0}^{I-1} (\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \nabla f(\bar{\mathbf{x}}_{r_0})) \right] \right\|^2 \right] \\
&\stackrel{(ii)}{\leq} 3 \mathbb{E} \left[\left\| \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N q_r^i \sum_{k=0}^{I-1} (\nabla f_i(\mathbf{x}_{r,k}^i) - \nabla f_i(\bar{\mathbf{x}}_{r,k})) \right\|^2 \right] \\
&\quad + 3 \mathbb{E} \left[\left\| \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N q_r^i \sum_{k=0}^{I-1} (\nabla f_i(\bar{\mathbf{x}}_{r,k}) - \nabla f_i(\bar{\mathbf{x}}_{r_0})) \right\|^2 \right] \\
&\quad + 3 \mathbb{E} \left[\left\| I \sum_{i=1}^N \mathbb{E}_{r_0} \left[\sum_{r=r_0}^{r_0+P-1} q_r^i \right] (\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \nabla f(\bar{\mathbf{x}}_{r_0})) \right\|^2 \right] \\
&\stackrel{(iii)}{\leq} 3IP \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N \sum_{k=0}^{I-1} \mathbb{E} \left[q_r^i \|\nabla f_i(\mathbf{x}_{r,k}^i) - \nabla f_i(\bar{\mathbf{x}}_{r,k})\|^2 \right] \\
&\quad + 3IP \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N \sum_{k=0}^{I-1} \mathbb{E} \left[q_r^i \|\nabla f_i(\bar{\mathbf{x}}_{r,k}) - \nabla f_i(\bar{\mathbf{x}}_{r_0})\|^2 \right] \\
&\quad + 3I^2 P^2 \mathbb{E} \left[\left\| \sum_{i=1}^N \mathbb{E}_{r_0} \left[\frac{1}{P} \sum_{r=r_0}^{r_0+P-1} q_r^i \right] (\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \nabla f(\bar{\mathbf{x}}_{r_0})) \right\|^2 \right] \\
&\stackrel{(iv)}{\leq} 3L^2 IP \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N \sum_{k=0}^{I-1} \mathbb{E} \left[q_r^i \|\mathbf{x}_{r,k}^i - \bar{\mathbf{x}}_{r,k}\|^2 \right] + 3L^2 IP \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N \sum_{k=0}^{I-1} \mathbb{E} \left[q_r^i \|\bar{\mathbf{x}}_{r,k} - \bar{\mathbf{x}}_{r_0}\|^2 \right] \\
&= 3L^2 IP \sum_{r=r_0}^{r_0+P-1} \sum_{k=0}^{I-1} \mathbb{E} \left[\sum_{i=1}^N q_r^i \mathbb{E} \left[\|\mathbf{x}_{r,k}^i - \bar{\mathbf{x}}_{r,k}\|^2 \mid \mathcal{Q} \right] \right] \\
&\quad + 3L^2 IP \sum_{r=r_0}^{r_0+P-1} \sum_{k=0}^{I-1} \mathbb{E} \left[\sum_{i=1}^N q_r^i \mathbb{E} \left[\|\bar{\mathbf{x}}_{r,k} - \bar{\mathbf{x}}_{r_0}\|^2 \mid \mathcal{Q} \right] \right] \\
&\leq 3L^2 IP \mathbb{E} \left[\sum_{r=r_0}^{r_0+P-1} \sum_{k=0}^{I-1} (D_{r,k} + M_{r,k}) \right],
\end{aligned}$$

where (i) uses the decomposition

$$\nabla f_i(\mathbf{x}_{r,k}^i) - \nabla f(\bar{\mathbf{x}}_{r_0}) = (\nabla f_i(\mathbf{x}_{r,k}^i) - \nabla f_i(\bar{\mathbf{x}}_{r,k})) + (\nabla f_i(\bar{\mathbf{x}}_{r,k}) - \nabla f_i(\bar{\mathbf{x}}_{r_0})) + (\nabla f_i(\bar{\mathbf{x}}_{r_0}) - \nabla f(\bar{\mathbf{x}}_{r_0})),$$

(ii) and (iii) use Jensen's inequality, and (iv) uses smoothness of each f_i along with the condition $\mathbb{E}_{r_0} [q_{r_0}^i] = \frac{1}{N}$. Plugging into Equation 33 yields

$$\begin{aligned}
& -\gamma\eta\mathbb{E} \left[\left\langle \nabla f(\bar{\mathbf{x}}_{r_0}), \mathbb{E}_{r_0} \left[\sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N q_r^i \sum_{k=0}^{I-1} \nabla f_i(\mathbf{x}_{r,k}^i) \right] \right\rangle \right] \\
& \leq -\frac{\gamma\eta IP}{2} \mathbb{E} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \right] + \frac{3}{2}\gamma\eta L^2 \mathbb{E} \left[\sum_{r=r_0}^{r_0+P-1} \sum_{k=0}^{I-1} (D_{r,k} + M_{r,k}) \right] \\
& \stackrel{(i)}{\leq} -\frac{\gamma\eta IP}{2} \mathbb{E} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \right] + \frac{3}{2}\gamma\eta L^2 \left(108\eta^2 I^3 P^3 \mathbb{E} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \right] \right. \\
& \quad \left. + (75\eta^2 I^2 P + 65\eta^2 I^2 P^2 \rho^2) \sigma^2 + 254\eta^2 L^2 I^3 P^3 \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0}^i] \right) \\
& \leq \gamma\eta IP \left(-\frac{1}{2} + 162\eta^2 L^2 I^2 P^2 \right) \mathbb{E} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \right] + \gamma\eta IP (113\eta^2 L^2 I + 98\eta^2 L^2 IP \rho^2) \sigma^2 \\
& \quad + 381\gamma\eta^3 L^4 I^3 P^3 \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0}^i],
\end{aligned}$$

where (i) uses Equation 4 from Lemma 1.

Plugging into Equation 32 and applying Lemma 2 yields

$$\begin{aligned}
& \mathbb{E}[f(\bar{\mathbf{x}}_{r_0+P}) - f(\bar{\mathbf{x}}_{r_0})] \\
& \leq \gamma\eta IP \left(-\frac{1}{2} + 162\eta^2 L^2 I^2 P^2 \right) \mathbb{E} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \right] + \gamma\eta IP (113\eta^2 L^2 I + 98\eta^2 L^2 IP \rho^2) \sigma^2 \\
& \quad + 381\gamma\eta^3 L^4 I^3 P^3 \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0}^i] + \frac{1}{2} L \mathbb{E} [\|\bar{\mathbf{x}}_{r_0+P} - \bar{\mathbf{x}}_{r_0}\|^2] \\
& \leq \gamma\eta IP \left(-\frac{1}{2} + 162\eta^2 L^2 I^2 P^2 \right) \mathbb{E} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \right] + \gamma\eta IP (113\eta^2 L^2 I + 98\eta^2 L^2 IP \rho^2) \sigma^2 \\
& \quad + 381\gamma\eta^3 L^4 I^3 P^3 \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0}^i] + \frac{1}{2} L \left(6\gamma^2 \eta^2 I^2 P^2 \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2] \right. \\
& \quad \left. + \gamma\eta IP (5\gamma\eta\rho^2 + 7\eta^2 LI) \sigma^2 + 11\gamma^2 \eta^2 L^2 I^2 P^2 \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0}^i] \right) \tag{34} \\
& \leq \gamma\eta IP \left(-\frac{1}{2} + 162\eta^2 L^2 I^2 P^2 + 3\gamma\eta LIP \right) \mathbb{E} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \right] \\
& \quad + \gamma\eta IP \left(\frac{5}{2}\gamma\eta L\rho^2 + 117\eta^2 L^2 I + 98\eta^2 L^2 IP \rho^2 \right) \sigma^2 \\
& \quad + (381\gamma\eta^3 L^4 I^3 P^3 + 6\gamma^2 \eta^2 L^3 I^2 P^2) \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0}^i] \\
& \stackrel{(i)}{\leq} \gamma\eta IP \left(-\frac{1}{2} + 162\eta^2 L^2 I^2 P^2 + 3\gamma\eta LIP \right) \mathbb{E} \left[\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2 \right] \\
& \quad + \gamma\eta IP \left(\frac{5}{2}\gamma\eta L\rho^2 + 117\eta^2 L^2 I + 98\eta^2 L^2 IP \rho^2 \right) \sigma^2 + p_{\text{sample}} \gamma\eta L^2 IP \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0}^i], \tag{35}
\end{aligned}$$

where (i) uses $\eta \leq \frac{\sqrt{p_{\text{sample}}}}{60LIP}$ and $\gamma\eta \leq \frac{p_{\text{sample}}}{60LIP}$.

Using Lemma 3,

$$\begin{aligned}
& 2\gamma\eta L^2 IP \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0+P}^i] \\
& \leq 2\gamma\eta L^2 IP \left(\left(324\eta^2 I^2 P^2 + \frac{20}{p_{\text{sample}}} \gamma^2 \eta^2 I^2 P^2 \right) \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2] \right. \\
& \quad \left. + \left(226\eta^2 I + 195\eta^2 IP\rho^2 + \frac{17}{p_{\text{sample}}} \gamma^2 \eta^2 IP\rho^2 \right) \sigma^2 + \left(1 - \frac{1}{2} p_{\text{sample}} \right) \frac{1}{N} \sum_{j=1}^N \mathbb{E} [S_{r_0}^j] \right) \\
& \leq \gamma\eta IP \left(648\eta^2 L^2 I^2 P^2 + \frac{40}{p_{\text{sample}}} \gamma^2 \eta^2 L^2 I^2 P^2 \right) \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2] \\
& \quad + \gamma\eta IP \left(552\eta^2 L^2 I + 390\eta^2 L^2 IP\rho^2 + \frac{34}{p_{\text{sample}}} \gamma^2 \eta^2 L^2 IP\rho^2 \right) \sigma^2 \\
& \quad + 2\gamma\eta L^2 IP \frac{1}{N} \sum_{j=1}^N \mathbb{E} [S_{r_0}^j] - p_{\text{sample}} \gamma\eta L^2 IP \frac{1}{N} \sum_{j=1}^N \mathbb{E} [S_{r_0}^j] \\
& \leq \gamma\eta IP \left(648\eta^2 L^2 I^2 P^2 + \frac{40}{p_{\text{sample}}} \gamma^2 \eta^2 L^2 I^2 P^2 \right) \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2] \\
& \quad + \gamma\eta IP \left(552\eta^2 L^2 I + 390\eta^2 L^2 IP\rho^2 + \gamma\eta L\rho^2 \right) \sigma^2 \\
& \quad + 2\gamma\eta L^2 IP \frac{1}{N} \sum_{j=1}^N \mathbb{E} [S_{r_0}^j] - p_{\text{sample}} \gamma\eta L^2 IP \frac{1}{N} \sum_{j=1}^N \mathbb{E} [S_{r_0}^j],
\end{aligned}$$

where the last inequality uses $\gamma\eta \leq \frac{p_{\text{sample}}}{60LIP}$. Adding this to Equation 35 and rearranging,

$$\begin{aligned}
& \mathbb{E}[f(\bar{\mathbf{x}}_{r_0+P})] + 2\gamma\eta L^2 IP \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0+P}^i] \\
& \leq \left(\mathbb{E}[f(\bar{\mathbf{x}}_{r_0})] + 2\gamma\eta L^2 IP \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0}^i] \right) \\
& \quad \gamma\eta IP \left(-\frac{1}{2} + 810\eta^2 L^2 I^2 P^2 + 3\gamma\eta LIP + \frac{40}{p_{\text{sample}}} \eta^2 L^2 I^2 P^2 \right) \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2] \\
& \quad + \gamma\eta IP \left(\frac{7}{2} \gamma\eta L\rho^2 + 669\eta^2 L^2 I + 488\eta^2 L^2 IP\rho^2 \right) \sigma^2 \\
& \leq \left(\mathbb{E}[f(\bar{\mathbf{x}}_{r_0})] + 2\gamma\eta L^2 IP \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_{r_0}^i] \right) \\
& \quad - \frac{1}{5} \gamma\eta IP \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2] + \gamma\eta IP (4\gamma\eta L\rho^2 + 669\eta^2 L^2 I + 488\eta^2 L^2 IP\rho^2) \sigma^2,
\end{aligned}$$

where we used $\gamma\eta \leq \frac{1}{60LIP}$ and $\eta \leq \frac{1}{60LIP}$. Finally, we can average over $r_0 \in \{0, P, 2P, \dots, R-P\}$ and rearrange to obtain

$$\begin{aligned}
\frac{P}{R} \sum_{r_0 \in \{0, P, \dots, R-P\}} \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_{r_0})\|^2] & \leq 5 \frac{\mathbb{E}[f(\bar{\mathbf{x}}_0) - f(\bar{\mathbf{x}}_R)]}{\gamma\eta IR} + 10 \frac{L^2 P}{R} \frac{1}{N} \sum_{i=1}^N \mathbb{E} [S_0^i - S_R^i] \\
& \quad + (20\gamma\eta L\rho^2 + 3345\eta^2 L^2 I + 2440\eta^2 L^2 IP\rho^2) \sigma^2 \\
& \stackrel{(i)}{\leq} \frac{5\Delta}{\gamma\eta IR} + (20\gamma\eta L\rho^2 + 3345\eta^2 L^2 I + 2440\eta^2 L^2 IP\rho^2) \sigma^2 \\
& \stackrel{(ii)}{\leq} \frac{5\Delta}{\gamma\eta IR} + (20\gamma\eta L\rho^2 + 5785\eta^2 L^2 IP) \sigma^2,
\end{aligned}$$

where (i) uses $\frac{1}{N} \sum_{i=1}^N \mathbb{E}[S_R^i] \geq 0$ and $\frac{1}{N} \sum_{i=1}^N \mathbb{E}[S_0^i] = 0$ by initialization, and (ii) uses $P \geq 1$ and $\rho \leq 1$. \square

Corollary 2 (Corollary 1 restated). *Let $\epsilon > 0$ and $I \geq 1$. Suppose Assumptions 1 and 2 hold, and that $\mathbb{E}[w_{r_0}^i | \mathcal{Q}_{:r_0}] \leq \frac{P^2}{N}$, and $\mathbb{E} \left[\sum_{i=1}^N (v_{r_0}^i)^2 \Lambda_{r_0}^i \right] \leq \rho^2$. If*

$$\begin{aligned} R &\geq \frac{\Delta L \rho^2 \sigma^2}{I \epsilon^4} + \frac{300 \Delta L P}{p_{\text{sample}} \epsilon^2} \\ \eta &\leq \frac{1}{264} \min \left\{ \frac{\Delta^{1/4} \rho^{1/2}}{L^{3/4} I^{3/4} R^{1/4} P^{1/2} \sigma^{1/2}}, \frac{\rho \sqrt{p_{\text{sample}}}}{LIP} \right\} \\ \gamma &= \frac{1}{\eta} \min \left\{ \frac{\sqrt{\Delta}}{60 \sqrt{LIR} \rho \sigma}, \frac{p_{\text{sample}}}{60 LIP} \right\}, \end{aligned}$$

then Algorithm 1 satisfies

$$\frac{P}{R} \sum_{r_0 \in \{0, P, \dots, R-P\}} \mathbb{E}[\|\nabla f(\bar{x}_{r_0})\|^2] \leq 302\epsilon^2.$$

Proof. First, $\eta \leq \frac{\rho \sqrt{p_{\text{sample}}}}{264 LIP} \leq \frac{\sqrt{p_{\text{sample}}}}{60 LIP}$ and $\gamma \eta \leq \frac{p_{\text{sample}}}{60 LIP}$ together with Assumptions 1 and 2 imply that the conditions of Theorem 2 are satisfied. Therefore

$$\frac{P}{R} \sum_{r_0 \in \{0, P, \dots, R-P\}} \mathbb{E}[\|\nabla f(\bar{x}_{r_0})\|^2] \leq \frac{5\Delta}{\gamma \eta IR} + (20\gamma \eta L \rho^2 + 5785\eta^2 L^2 IP) \sigma^2.$$

From our choice of η ,

$$\begin{aligned} 5785\eta^2 L^2 IP &\leq \frac{5785}{264^2} L^2 IP \min \left\{ \frac{\Delta^{1/2} \rho}{L^{3/2} I^{3/2} R^{1/2} P \sigma}, \frac{\rho^2 p_{\text{sample}}}{L^2 I^2 P^2} \right\} \\ &\leq 5L\rho^2 \min \left\{ \frac{\Delta^{1/2}}{60 L^{1/2} I^{1/2} R^{1/2} \rho \sigma}, \frac{p_{\text{sample}}}{60 LIP} \right\} \\ &= 5\gamma \eta L \rho^2. \end{aligned}$$

Therefore

$$\begin{aligned} \frac{P}{R} \sum_{r_0 \in \{0, P, \dots, R-P\}} \mathbb{E}[\|\nabla f(\bar{x}_{r_0})\|^2] &\leq \frac{5\Delta}{\gamma \eta IR} + 25\gamma \eta L \rho^2 \sigma^2 \\ &\leq \frac{5\Delta}{IR} \left(\frac{60\sqrt{LIR} \rho \sigma}{\sqrt{\Delta}} + \frac{60 LIP}{p_{\text{sample}}} \right) + 25L\rho^2 \sigma^2 \frac{\sqrt{\Delta}}{60\sqrt{LIR} \rho \sigma} \\ &\leq \frac{301\sqrt{\Delta L} \rho \sigma}{\sqrt{IR}} + \frac{300 \Delta L P}{p_{\text{sample}} R}, \end{aligned}$$

where we used our choice of $\gamma \eta$. Finally, from our choice of R ,

$$\begin{aligned} \frac{P}{R} \sum_{r_0 \in \{0, P, \dots, R-P\}} \mathbb{E}[\|\nabla f(\bar{x}_{r_0})\|^2] &\leq \frac{301\sqrt{\Delta L} \rho \sigma}{\sqrt{I}} \frac{\sqrt{I} \epsilon^2}{\sqrt{\Delta L} \rho \sigma} + \frac{300 \Delta L P}{p_{\text{sample}}} \frac{p_{\text{sample}} \epsilon^2}{300 \Delta L P} \\ &\leq 302\epsilon^2. \end{aligned}$$

□

B Proofs for Specific Participation Patterns

Here we derive the convergence rates for the specific participation patterns discussed in Section 4.3. In order to apply Corollary 1, the conditions in Assumption 2, Equation 1 and Equation 2 must be satisfied. Since we already showed that Assumption 2 is satisfied by regularized participation and cyclic participation (in Section 3.2), it only remains to show that Equation 1 and Equation 2 is satisfied, and plug the appropriate values of ρ^2 , P , and p_{sample} into Corollary 1.

Corollary 3 (Regularized Participation). *Under regularized participation, for any $\epsilon > 0$, there exist choices of η, γ, I , and R such that $\mathbb{E}[\|\nabla f(\hat{\mathbf{x}})\|^2] \leq O(\epsilon^2)$ and*

$$R = \mathcal{O}\left(\frac{LP}{\epsilon^2}\right), \quad RI = \mathcal{O}\left(\frac{\Delta L \rho^2 \sigma^2}{\epsilon^4}\right).$$

Proof. We first show that Equation 1 and Equation 2 are satisfied under regularized participation. Let $r_0 \in \{0, \dots, R - P\}$. Using the fact that $\bar{q}_{r_0}^i = \frac{1}{N}$ almost surely:

$$\begin{aligned} \mathbb{E}[w_{r_0}^i | \mathcal{Q}_{:r_0}] &= \mathbb{E}\left[\frac{1}{N} \sum_{j=1}^N \frac{\mathbb{1}\{\bar{q}_{r_0}^j > 0\}}{P \bar{q}_{r_0}^j} \sum_{s=r_0}^{r_0+P-1} q_s^i q_s^j \middle| \mathcal{Q}_{:r_0}\right] = \frac{1}{P} \mathbb{E}\left[\sum_{j=1}^N \sum_{s=r_0}^{r_0+P-1} q_s^i q_s^j \middle| \mathcal{Q}_{:r_0}\right] \\ &= \frac{1}{P} \mathbb{E}\left[\sum_{s=r_0}^{r_0+P-1} q_s^i \sum_{j=1}^N q_s^j \middle| \mathcal{Q}_{:r_0}\right] = \frac{1}{P} \mathbb{E}\left[\sum_{s=r_0}^{r_0+P-1} q_s^i \middle| \mathcal{Q}_{:r_0}\right] \\ &= \mathbb{E}[q_{r_0}^i | \mathcal{Q}_{:r_0}] = \frac{1}{N} \leq \frac{P^2}{N}, \end{aligned}$$

where we used the fact that $\sum_{i=1}^N q_r^i = 1$. This shows that Equation 1 is satisfied. Also, $\bar{q}_{r_0}^i = \frac{1}{N}$ means that $v_{r_0}^i = 0$, so that $\mathbb{E}[(v_{r_0}^i)^2 \Lambda_{r_0}^i] = 0 \leq \rho^2$, and Equation 2 is satisfied. Therefore we can apply Corollary 1. The complexity result follows by plugging in $p_{\text{sample}} = 1$ and choosing $I = \mathcal{O}\left(\frac{\Delta \rho^2 \sigma^2}{P \epsilon^2}\right)$. \square

Corollary 4 (Cyclic Participation). *Under cyclic participation with \bar{K} groups and S participating clients in each round, for any $\epsilon > 0$, there exist choices of η, γ, P, I , and R such that $\mathbb{E}[\|\nabla f(\hat{\mathbf{x}})\|^2] \leq O(\epsilon^2)$ and*

$$R = \mathcal{O}\left(\frac{L\bar{K}}{\epsilon^2} \left(\frac{N}{S}\right)\right), \quad RI = \mathcal{O}\left(\frac{\Delta L \sigma^2}{S \epsilon^4}\right).$$

Proof. Again, we show that Equation 1 and Equation 2 are satisfied under cyclic participation. Let $r_0 \in \{0, \dots, R - P\}$. Denoting $g(i) = \lfloor \frac{\bar{K}}{N}(i - 1) \rfloor$ and $r(i) = r_0 + g(i)$, it holds that, over the rounds $r \in \{r_0, \dots, r_0 + P - 1\}$, client i belongs to group $g(i)$ and is available only during round $r(i)$. Also, let $\Pi(i)$ denote the set of indices of clients in group $g(i)$. To simplify $w_{r_0}^i$, recall that for client i in group $(r \bmod \bar{K})$:

$$q_r^i = \begin{cases} \frac{1}{S} & \text{with probability } \frac{S}{N} \\ 0 & \text{with probability } 1 - \frac{S}{N} \end{cases},$$

and $q_r^i = 0$ for all clients i not in group $(r \bmod \bar{K})$. So

$$\begin{aligned}
\mathbb{E} [w_{r_0}^i | \mathcal{Q}_{:r_0}] &= \mathbb{E} \left[\frac{1}{N} \sum_{j=1}^N \frac{\mathbb{1}\{\bar{q}_{r_0}^j > 0\}}{P \bar{q}_{r_0}^j} \sum_{s=r_0}^{r_0+P-1} q_s^i q_s^j \middle| \mathcal{Q}_{:r_0} \right] \\
&= \frac{1}{NP} \mathbb{E} \left[\sum_{s=r_0}^{r_0+P-1} q_s^i \sum_{j=1}^N \frac{\mathbb{1}\{\bar{q}_{r_0}^j > 0\}}{\bar{q}_{r_0}^j} q_s^j \middle| \mathcal{Q}_{:r_0} \right] \\
&\stackrel{(i)}{=} \frac{1}{NP} \mathbb{E} \left[q_{r(i)}^i \sum_{j=1}^N \frac{\mathbb{1}\{\bar{q}_{r_0}^j > 0\}}{\bar{q}_{r_0}^j} q_{r(i)}^j \middle| \mathcal{Q}_{:r_0} \right] \\
&\stackrel{(ii)}{=} \frac{1}{NP} \mathbb{E} \left[q_{r(i)}^i \sum_{j \in \Pi(i)} \frac{\mathbb{1}\{\bar{q}_{r_0}^j > 0\}}{\bar{q}_{r_0}^j} q_{r(i)}^j \middle| \mathcal{Q}_{:r_0} \right] \\
&\stackrel{(iii)}{=} \frac{1}{NP} \mathbb{E} \left[q_{r(i)}^i \sum_{j \in \Pi(i)} \frac{\mathbb{1}\{\bar{q}_{r_0}^j > 0\}}{1/SP} \frac{1}{S} \middle| \mathcal{Q}_{:r_0} \right] \\
&= \frac{1}{N} \mathbb{E} \left[q_{r(i)}^i \sum_{j \in \Pi(i)} \mathbb{1}\{\bar{q}_{r_0}^j > 0\} \middle| \mathcal{Q}_{:r_0} \right] \\
&\stackrel{(iv)}{=} \frac{S}{N} \mathbb{E} \left[q_{r(i)}^i \middle| \mathcal{Q}_{:r_0} \right] \\
&= \frac{S}{N} \left(\frac{S}{N} \frac{1}{S} + \left(1 - \frac{S}{N}\right) \cdot 0 \right) \\
&= \frac{S}{N^2} \leq \frac{P^2}{N},
\end{aligned}$$

where (i) uses $q_r^i = 0$ for all $r \neq r(i)$, (ii) uses $q_{r(i)}^j = 0$ for all $j \notin \Pi(i)$, (iii) uses the fact that $\bar{q}_{r_0}^j > 0$ implies $q_{r(i)}^j = \frac{1}{S}$ and $\bar{q}_{r_0}^j = \frac{1}{SP}$, and (iv) uses the fact that S clients are sampled in each round. This shows that Equation 1 is satisfied.

To see that Equation 2 is satisfied:

$$\begin{aligned}
\mathbb{E} \left[\sum_{i=1}^N (v_{r_0}^i)^2 \Lambda_{r_0}^i \right] &\stackrel{(i)}{=} \mathbb{E} \left[\sum_{i=1}^N \mathbb{E} \left[\left(\bar{q}_{r_0}^i - \frac{1}{N} \right)^2 \middle| \mathcal{Q}_{:r_0} \right] \Lambda_{r_0}^i \right] \\
&\stackrel{(ii)}{=} \frac{1}{SNP} \left(1 - \frac{S\bar{K}}{N} \right) \sum_{i=1}^N \mathbb{E} [\Lambda_{r_0}^i] \\
&= \frac{1}{SNP} \left(1 - \frac{S\bar{K}}{N} \right) \sum_{i=1}^N \left(\frac{1}{P} \sum_{s=r_0-P}^{r_0-1} (z_s^i)^2 \middle/ \left(\frac{1}{P} \sum_{s=r_0-P}^{r_0-1} z_s^i \right)^2 \right) \\
&\stackrel{(iii)}{=} \frac{1}{SP} \left(1 - \frac{S\bar{K}}{N} \right) \frac{1}{P} \frac{1}{S^2} \middle/ \left(\frac{1}{P} \frac{1}{S} \right)^2 \\
&= \frac{1}{S} \left(1 - \frac{S\bar{K}}{N} \right) \leq \frac{1}{S} = \rho^2,
\end{aligned}$$

where (i) uses the tower property, (ii) uses the variance of $\bar{q}_{r_0}^i$ as computed in Appendix C, and (iii) uses the fact that $\bar{z}_{r_0}^i > 0$ by construction of $\bar{z}_{r_0}^i$, and in this case $z_r^i = \frac{1}{S}$ for exactly one $r \in \{r_0 - P, \dots, r_0 - 1\}$ and $z_r^i = 0$ for all other r . Therefore Equation 2 is satisfied.

This shows we can apply Corollary 1. The complexity result follows by plugging in $p_{\text{sample}} = \frac{S}{N}$, $P = \bar{K}$, $\rho = \frac{1}{\sqrt{S}}$, and choosing $I = \mathcal{O} \left(\frac{\Delta \sigma^2}{\bar{K} N \epsilon^2} \right)$. \square

C Complexity of Amplified FedAvg

In this section we derive the computation and communication complexity required for Amplified FedAvg [39] to find an ϵ -stationary point under various client participation patterns, which we listed in Table 1 of the main paper. Table 1 compares the complexity under i.i.d. participation, regularized participation, and cyclic participation. Since i.i.d. participation is a special case of cyclic participation with $K = 1$ groups, here we only consider regularized and cyclic participation, and the result for i.i.d. participation follows.

Many works in federated learning characterize data heterogeneity by assuming that there exists a constant κ such that

$$\|\nabla f_i(\mathbf{x}) - \nabla f(\mathbf{x})\| \leq \kappa,$$

for all \mathbf{x} . The previous analysis of Amplified FedAvg [39] instead assumes an upper bound $\tilde{\delta}(P)$ on a weighted heterogeneity that depends on client sampling. Specifically, they assume that there exists $\tilde{\delta}(P)$ such that

$$\left\| \frac{1}{P} \sum_{r=r_0}^{r_0+P-1} \sum_{i=1}^N q_r^i (\nabla f_i(\mathbf{x}) - \nabla f(\mathbf{x})) \right\|^2 \leq \tilde{\delta}^2(P),$$

for all \mathbf{x} and r_0 . We restate their Corollary 3.2 for convenience:

Corollary 5 (Corollary 3.2 [39] informally restated). *There exist parameter choices such that Amplified FedAvg satisfies*

$$\min_r \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_r)\|^2] \leq \mathcal{O} \left(\frac{\sqrt{\Delta L} \rho \sigma}{\sqrt{R I}} + \frac{\Delta L P + \kappa^2}{R} + \frac{\sigma^2}{R I P} + \mathbb{E} [\tilde{\delta}^2(P)] \right). \quad (36)$$

As pointed out in their Section 4.2, we can interpret $\tilde{\delta}(P)$ in terms of the conventional heterogeneity constant κ as:

$$\tilde{\delta}^2(P) \leq N \kappa^2 \sum_{i=1}^N \left(\bar{q}_{r_0}^i - \frac{1}{N} \right)^2.$$

Our Assumption 2(b) implies that $\mathbb{E}_{\mathcal{Q}_{r_0}} [\bar{q}_{r_0}^i] = \frac{1}{N}$. If we choose $v \geq 0$ such that $\text{Var}[\bar{q}_{r_0}^i] \leq v^2$, then we can take expectation of the above to obtain

$$\mathbb{E} [\tilde{\delta}^2(P)] \leq N^2 \kappa^2 v^2.$$

This ‘‘conversion’’ of $\tilde{\delta}(P)$ to κ and v will allow us to compare their complexity to that of algorithms that use the conventional heterogeneity assumption by computing v for each participation pattern.

Regularized Participation In this case, $\bar{q}_{r_0}^i = \frac{1}{N}$ almost surely, so that $v = 0$ and accordingly $\tilde{\delta}^2(P) = 0$. Therefore

$$\min_r \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_r)\|^2] \leq \mathcal{O} \left(\frac{\sqrt{\Delta L} \rho \sigma}{\sqrt{R I}} + \frac{\Delta L P + \kappa^2}{R} + \frac{\sigma^2}{R I P} \right).$$

Therefore the choices

$$R \geq \mathcal{O} \left(\frac{\Delta L P + \kappa^2}{\epsilon^2} \right), \quad R I \geq \mathcal{O} \left(\frac{\Delta L \rho^2 \sigma^2}{\epsilon^4} + \frac{\sigma^2}{P \epsilon^2} \right),$$

imply

$$\min_r \mathbb{E} [\|\nabla f(\bar{\mathbf{x}}_r)\|^2] \leq \mathcal{O} (\epsilon^2).$$

Cyclic Participation We can compute v in terms of the parameters of the participation pattern: number of groups \bar{K} and number of participating clients S in each round. Although we chose $P = \bar{K}$ for Amplified SCAFFOLD, we can satisfy Assumption(b) [2] by choosing $P = m\bar{K}$ for any $m \in \mathbb{N}$, and indeed to achieve $\mathbb{E}[\tilde{\delta}^2(P)] \leq \epsilon^2$ we must choose $P = m\bar{K}$ with m depending on ϵ .

Let $A(i)$ denote the set of rounds in $\{r_0, \dots, r_0 + P - 1\}$ during which client i is available. Note that $A(i)$ has size m , since $P = m\bar{K}$. Then

$$\bar{q}_{r_0}^i - \frac{1}{N} = \frac{1}{P} \sum_{r=r_0}^{r_0+P-1} q_r^i - \frac{1}{N} \stackrel{(i)}{=} \frac{1}{P} \sum_{r \in A(i)} q_r^i - \frac{1}{N} = \frac{1}{m\bar{K}} \sum_{r \in A(i)} q_r^i - \frac{1}{b\bar{K}} = \frac{1}{\bar{K}} \left(\frac{1}{m} \sum_{r \in A(i)} q_r^i - \frac{1}{b} \right),$$

where (i) uses the fact that $q_r^i = 0$ for all $r \notin A(i)$. Therefore

$$\begin{aligned} v^2 &= \frac{1}{\bar{K}^2} \mathbb{E} \left[\left(\frac{1}{m} \sum_{r \in A(i)} q_r^i - \frac{1}{b} \right)^2 \right] \\ &\stackrel{(i)}{=} \frac{1}{\bar{K}^2} \frac{1}{m^2} \sum_{r \in A(i)} \mathbb{E} \left[\left(q_r^i - \frac{1}{b} \right)^2 \right] \\ &\stackrel{(ii)}{=} \frac{1}{\bar{K}^2} \frac{1}{m^2} \sum_{r \in A(i)} \frac{1}{S^2} \frac{S}{b} \left(1 - \frac{S}{b} \right) \\ &= \frac{1}{Sm\bar{K}^2b} \left(1 - \frac{S}{b} \right) \\ &= \frac{1}{SNP} \left(1 - \frac{S\bar{K}}{N} \right), \end{aligned}$$

where (i) uses the fact that $\{q_r^i\}_{r \in A(i)}$ are independent and (ii) uses the fact that q_r^i equals $\frac{1}{S}$ times a Bernoulli variable for $r \in A(i)$.

With a bound for v^2 , we can bound the remaining term in Equation 36 as follows:

$$\mathbb{E}[\tilde{\delta}^2(P)] \leq N^2 \kappa^2 v^2 \leq \frac{\kappa^2}{P} \frac{N}{S} \left(1 - \frac{S\bar{K}}{N} \right).$$

Therefore, Amplified FedAvg under cyclic participation can find an ϵ -stationary point with the choices

$$\begin{aligned} P &\geq \max \left\{ \bar{K}, \frac{\kappa^2}{\epsilon^2} \frac{N}{S} \left(1 - \frac{S\bar{K}}{N} \right) \right\} \\ R &\geq \mathcal{O} \left(\frac{\Delta L \bar{K} + \kappa^2}{\epsilon^2} + \frac{\Delta L N \kappa^2}{S \epsilon^4} \left(1 - \frac{S\bar{K}}{N} \right) \right) \\ RI &\geq \mathcal{O} \left(\frac{\Delta L \sigma^2}{S \epsilon^4} \right). \end{aligned}$$

D Experiment Details

Here we discuss experimental details deferred from the main body: client sampling parameters, heterogeneity protocol, hyperparameter tuning, definition of the synthetic objective, and specification of the CNN architecture used for the image classification experiments.

D.1 Client Sampling Parameters

For the synthetic objective, we set the number of groups $\bar{K} = 2$ and the availability time $g = 240$. We set the communication interval $I = 10$ and train for $R = 5000$ rounds.

For Fashion-MNIST, we set the communication interval $I = 30$ and train a logistic regression model for $R = 2000$ rounds, with availability time $g = 4$.

For CIFAR-10, we set the communication interval $I = 5$ and train a two-layer CNN for $R = 12000$ rounds, with availability time $g = 10$.

Table 2: Hyperparameter search ranges and final values.

	γ values	$\gamma\eta$ values	μ values
Synthetic			
FedAvg	{1}	{ 10^{-6} , 10^{-5} , 10^{-4} , 10^{-3} }	{0}
FedProx	{1}	{ 10^{-6} , 10^{-5} , 10^{-4} , 10^{-3} }	{0.01, 0.1, 1.0, 10.0}
SCAFFOLD	{1}	{ 10^{-6} , 10^{-5} , 10^{-4} , 10^{-3} }	{0}
Amplified FedAvg	{1.25, 1.5, 2, 3 }	{ 10^{-6} , 10^{-5} , 10^{-4} , 10^{-3} }	{0}
Amplified SCAFFOLD	{1.25, 1.5 , 2, 3}	{ 10^{-6} , 10^{-5} , 10^{-4} , 10^{-3} }	{0}
Fashion-MNIST			
FedAvg	{1}	{ 10^{-5} , 10^{-4} , 10^{-3} , 10^{-2} }	{0}
FedProx	{1}	{ 10^{-5} , 10^{-4} , 10^{-3} , 10^{-2} }	{0.01, 0.1, 1.0, 10.0 }
SCAFFOLD	{1}	{ 10^{-5} , 10^{-4} , 10^{-3} , 10^{-2} }	{0}
Amplified FedAvg	{1.25, 1.5, 2, 3}	{ 10^{-5} , 10^{-4} , 10^{-3} , 10^{-2} }	{0}
Amplified SCAFFOLD	{1.25, 1.5 , 2, 3}	{ 10^{-5} , 10^{-4} , 10^{-3} , 10^{-2} }	{0}
CIFAR-10			
FedAvg	{1}	{ 10^{-6} , 10^{-5} , 10^{-4} , 10^{-3} }	{0}
FedProx	{1}	{ 10^{-6} , 10^{-5} , 10^{-4} , 10^{-3} }	{0.01, 0.1, 1.0, 10.0}
SCAFFOLD	{1}	{ 10^{-6} , 10^{-5} , 10^{-4} , 10^{-3} }	{0}
Amplified FedAvg	{1.25, 1.5, 2, 3 }	{ 10^{-6} , 10^{-5} , 10^{-4} , 10^{-3} }	{0}
Amplified SCAFFOLD	{1.25, 1.5 , 2, 3}	{ 10^{-6} , 10^{-5} , 10^{-4} , 10^{-3} }	{0}

D.2 Heterogeneity Protocol

The following protocol is commonly used in the literature [17][39] to convert a dataset into a collection of heterogeneous local datasets according to a data similarity parameter s , where $s = 0\%$ creates maximal data heterogeneity across clients, and $s = 100\%$ means that data is allocated to each client uniformly at random.

A single, non-federated dataset (e.g. CIFAR-10) is partitioned into two subsets: $s\%$ of the samples are allocated to an i.i.d. pool and randomly shuffled, and the remaining $(100 - s)\%$ of the samples are allocated to a non-i.i.d. pool and are sorted by label. Samples are allocated to each client so that $s\%$ of each local dataset comes from the i.i.d. pool, and the remaining $(100 - s)\%$ comes from the non-i.i.d. pool, so that with a small s , the majority of each local dataset consists of a small number of labels.

D.3 Hyperparameter Tuning

For all three experiments in the main body (synthetic objective, Fashion-MNIST, and CIFAR-10), each of the four baselines are individually tuned with grid search. For algorithms that use amplified updates (Amplified FedAvg and Amplified SCAFFOLD), we tune the amplification rate γ by searching over a fixed range of values. For the other algorithms (FedAvg and SCAFFOLD), we indicate the lack of amplified updates by setting $\gamma = 1$. We tune η by allowing $\gamma\eta$ over a fixed range of values. For FedProx's μ parameter, we also search over a fixed range of values. The search range and final values for each parameter are written in Table 2, along with the final values adopted for each algorithm.

D.4 Synthetic Objective

For the synthetic experiment, we use a difficult objective from a lower bound analysis of FedAvg [43]. As defined in the lower bound analysis, the objective is parameterized by H , κ , σ , c , μ , and L .

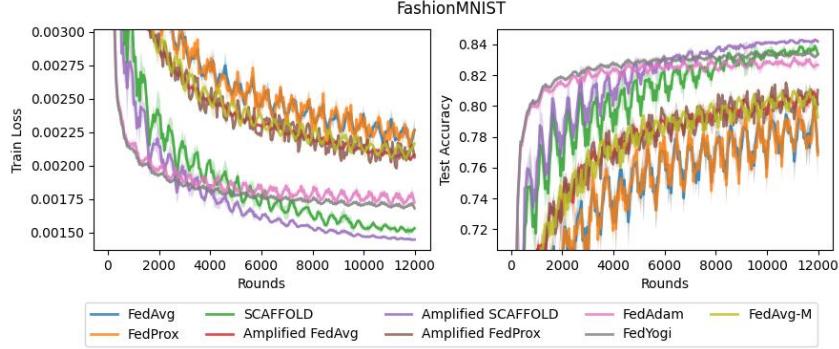


Figure 3: FashionMNIST with additional baselines. Amplified SCAFFOLD maintains the best performance.

The objective maps \mathbb{R}^4 to \mathbb{R} , and there are only two clients with corresponding local objectives

$$f_1(\mathbf{x}) = \frac{\mu}{2} (x_1 - c)^2 + \frac{H}{2} \left(x_2 - \frac{\sqrt{\mu}c}{\sqrt{H}} \right)^2 + \frac{H}{8} (x_3^2 + [x_3]_+^2) + \kappa x_4$$

$$f_2(\mathbf{x}) = \frac{\mu}{2} (x_1 - c)^2 + \frac{H}{2} \left(x_2 - \frac{\sqrt{\mu}c}{\sqrt{H}} \right)^2 + \frac{H}{8} (x_3^2 + [x_3]_+^2) - \kappa x_4,$$

where $[x]_+ := \max\{x, 0\}$. The stochastic gradients for f_1 and f_2 are sampled from the distributions $\nabla f_1(\mathbf{x}) + \xi e_3$ and $\nabla f_2(\mathbf{x}) + \xi e_3$, respectively, where $\xi \sim \mathcal{N}(0, \sigma^2)$ and e_3 denotes the third standard basis vector in \mathbb{R}^4 . We set the parameters of the objective as follows:

$$H = 16, \quad \kappa = 16, \quad \sigma = 1$$

$$c = 1, \quad \mu = 2, \quad L = 2.$$

D.5 CNN Architecture

We use a simple 2-layer CNN for CIFAR-10. The first layer is a convolutional layer with 64 channels, a 5×5 kernel, stride of 2, padding of 2, and a ReLU activation. The second layer is a fully connected layer with no activation.

E Additional Experimental Results

In this section, we provide two additional experimental results. First, in Section E.1, we add four baselines to the experimental settings of the main paper: FedAdam [32], FedYogi [32], FedAvg-M [5], and Amplified FedAvg with FedProx regularization. Second, in Section E.2, we evaluate all nine algorithms (five from the main paper and the four additional baselines) under another non-i.i.d. client participation pattern for the CIFAR-10 dataset, which we refer to as Stochastic Client Availability (SCA).

E.1 Additional Baselines

We evaluate the four additional baselines for the Fashion-MNIST and CIFAR-10 experiments from Section 5, keeping the same experimental setup. We tuned the hyperparameters of all baselines according to the hyperparameter ranges suggested in the original paper of each algorithm, and we allow the same compute budget for tuning each baseline as we did for tuning the algorithms in the original paper, in terms of the total number of hyperparameter combinations evaluated. Also, the results are averaged over five random seeds. The results are shown in Figures 3 and 4.

For FashionMNIST, FedAdam and FedYogi reach moderate training loss quickly, but are soon overtaken by Amplified SCAFFOLD and later by SCAFFOLD. FedAvg-M exhibits a minor advantage over FedAvg, but performs about the same as Amplified FedAvg. Amplified FedProx (i.e. Amplified FedAvg with FedProx regularization) performs nearly identically to Amplified FedAvg.

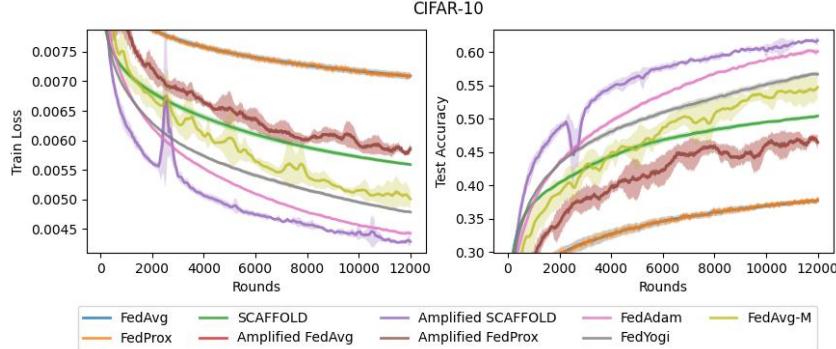


Figure 4: CIFAR-10 with additional baselines. FedAdam is competitive, but Amplified SCAFFOLD maintains superiority.

For CIFAR-10, FedAdam is more competitive, but is still outperformed by Amplified SCAFFOLD. FedYogi and FedAvg-M are further behind, though both still outperform SCAFFOLD. Amplified FedProx is again nearly identical to Amplified FedAvg.

These additional comparisons demonstrate that Amplified SCAFFOLD outperforms strong empirical baselines (FedAdam, FedYogi) under cyclic client participation, reinforcing the empirical validation of our algorithm. This performance is consistent with the fact that Amplified SCAFFOLD has convergence guarantees under periodic participation, while FedAdam and FedYogi were not designed for settings beyond i.i.d. client sampling.

E.2 CIFAR-10 with Stochastic Client Availability

Here, we include an evaluation under another non-i.i.d. participation pattern, which we refer to as Stochastic Cyclic Availability (SCA). SCA models device availability which is both periodic and unreliable. Similarly to cyclic participation, the set of clients is divided into groups, and at each round one group is deemed the "active" group, while the others are inactive. Unlike cyclic participation, in SCA not every client in the active group is always available: Instead, when a group becomes active, the clients in that group become available for sampling with probability 80%, while clients in inactive groups have probability 5% to be available for participation. The active group changes every g rounds. This stochastic availability models the real-life situation where a client device can be unavailable at a time of day when it is usually available, or vice versa. In this way, SCA is more flexible than cyclic participation and better captures the unreliability of client devices. Lastly, we reused the remaining settings (g , \bar{K} , I , etc.) and the tuned hyperparameters for each baseline from the CIFAR-10 experiment under cyclic participation. Again, we average each algorithm's performance over five random seeds.

Results for CIFAR-10 under SCA participation are shown in Figure 5. Again, Amplified SCAFFOLD outperforms all baselines under SCA participation. The relative performance of each baseline is similar as under cyclic participation, with FedAdam staying competitive with Amplified SCAFFOLD, followed by FedYogi and FedAvg-M. The remainder of the baselines have significantly worse performance, and again Amplified FedAvg has not benefitted by adding FedProx regularization.

F Extended Comparison with Baselines

Here we include an extended comparison against two relevant prior works.

FedAvg with Cyclic Participation [8] analyzes FedAvg under cyclic participation, but the resulting convergence rate does not benefit from local steps unless regularized participation is satisfied. They analyze FedAvg for L -smooth and μ -PL objectives. Using their notation, K is the number of client groups, κ is the condition number of the objective, γ is the intra-group heterogeneity, M is the total number of clients, N is the number of clients that participate in each round, and T is the number of communication rounds. Then the dominating term in their convergence rate for FedAvg under cyclic

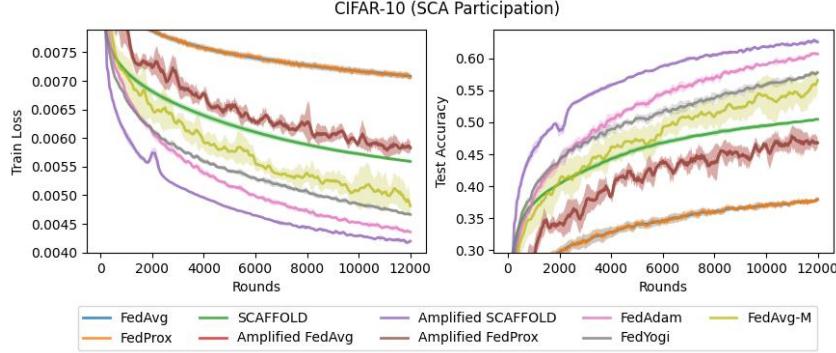


Figure 5: CIFAR-10 under SCA (stochastic cyclic availability). Amplified SCAFFOLD converges fastest.

participation (Theorem 2) is

$$\tilde{\mathcal{O}}\left(\frac{\bar{K}\kappa\gamma^2}{\mu NT}\left(\frac{M/\bar{K}-N}{M/\bar{K}-1}\right)\right).$$

Notice that the number of local steps (denoted τ) does not appear in this dominating term, so there is no way to reduce the communication complexity (compared to parallel SGD) by taking local steps. The only exception is when this term is zero from $N = M/\bar{K}$, which is equivalent to the condition that every client participates within every cycle of availability. Therefore, this result cannot show a benefit from local steps unless the client participation is regularized.

SCAFFOLD In Section 4.3, we mentioned a discrepancy between the complexity of SCAFFOLD vs. Amplified SCAFFOLD in terms of the dependence on N/S . In Table 1, the communication complexity of SCAFFOLD and Amplified SCAFFOLD under i.i.d. participation differs by their dependence on $\frac{N}{S}$. The complexity of Amplified SCAFFOLD is $\mathcal{O}\left(\frac{N}{S}\right)$, while that of SCAFFOLD is $\mathcal{O}\left(\left(\frac{N}{S}\right)^{2/3}\right)$. This difference in the order of $\frac{N}{S}$ is due to a potential small issue in the analysis of SCAFFOLD, which we intentionally avoided by accepting a slightly worse dependence on $\frac{N}{S}$.

This difference stems from an apparent mistake in the original SCAFFOLD analysis. In the proof of Lemma 16 (PMLR version of SCAFFOLD), the second-to-last equation of page 32 is obtained with an incorrect step. Namely, while the current version includes

$$\begin{aligned}\Xi_r &= \frac{1}{KN} \sum_{i,k} \mathbb{E} \|\alpha_{i,k-1}^r - \mathbf{x}^r\|^2 \\ &= \left(1 - \frac{S}{N}\right) \frac{1}{KN} \sum_{i,k} \mathbb{E} \|\alpha_{i,k-1}^{r-1} - \mathbf{x}^r\|^2 + \frac{S}{N} \frac{1}{KN} \sum_{i,k} \mathbb{E} \|\mathbf{y}_{i,k-1}^r - \mathbf{x}^r\|^2,\end{aligned}\quad (37)$$

where the last line is obtained by conditioning on the event $\alpha_{i,k-1}^r = \alpha_{i,k-1}^{r-1}$ and the complement $\alpha_{i,k-1}^{r-1} = \mathbf{y}_{i,k-1}^r$, which have probabilities $1 - \frac{S}{N}$ and $\frac{S}{N}$, respectively. However, this condition is not denoted in Equation 37. A corrected version should be written

$$\begin{aligned}\Xi_r &= \left(1 - \frac{S}{N}\right) \frac{1}{KN} \sum_{i,k} \mathbb{E} \left[\|\alpha_{i,k-1}^{r-1} - \mathbf{x}^r\|^2 \mid \alpha_{i,k-1}^r = \alpha_{i,k-1}^{r-1} \right] \\ &\quad + \frac{S}{N} \frac{1}{KN} \sum_{i,k} \mathbb{E} \left[\|\mathbf{y}_{i,k-1}^r - \mathbf{x}^r\|^2 \mid \alpha_{i,k-1}^{r-1} = \mathbf{y}_{i,k-1}^r \right],\end{aligned}$$

but this conditional expectation is not specified in the proof of SCAFFOLD. For the remainder of the proof of Lemma 16, these terms are treated as total expectation, leading to an inconsistency. Lemma 16 concludes by applying Lemma 15 in order to bound the term $\mathbb{E} \left[\|\mathbb{E}_{r-1} [\Delta \mathbf{x}^r]\|^2 \right]$. However, when we make the correction to include the conditional expectation, we do not have a bound of Ξ_r in terms

of $\mathbb{E} \left[\|\mathbb{E}_{r-1} [\Delta \mathbf{x}^r]\|^2 \right]$, we instead have a bound in terms of $\mathbb{E} \left[\left\| \mathbb{E}_{r-1} \left[\Delta \mathbf{x}^r \mid \alpha_{i,k-1}^{r-1} = y_{i,k-1}^r \right] \right\|^2 \right]$. But this term with a conditional expectation inside the norm can't be bounded with Lemma 15.

A pessimistic solution is to use Jensen's inequality to bound

$$\begin{aligned} \mathbb{E} \left[\left\| \mathbb{E}_{r-1} \left[\Delta \mathbf{x}^r \mid \alpha_{i,k-1}^{r-1} = y_{i,k-1}^r \right] \right\|^2 \right] &\leq \mathbb{E} \left[\mathbb{E}_{r-1} \left[\|\Delta \mathbf{x}^r\|^2 \mid \alpha_{i,k-1}^{r-1} = y_{i,k-1}^r \right] \right] \\ &= \mathbb{E} \left[\|\Delta \mathbf{x}^r\|^2 \right], \end{aligned}$$

where the second line follows from the tower property. This is the step that we perform in our analysis of Amplified SCAFFOLD, and this results in the $\frac{N}{S}$ dependence.

Fixing their analysis to recover the same $(\frac{N}{S})^{2/3}$ dependence may be possible, but we have instead focused on achieving the best known complexity in terms of ϵ, κ, σ , and other problem parameters.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: Every claim made in the abstract is specifically tied to a theoretical convergence property of our proposed algorithm, which are stated in Section 4.2 and proven in A and B.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discuss limitations in the "limitations" paragraph of Section 6.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: All of the theoretical results are stated with the full set of assumptions in Sections 3, 3.1, and 4.2. The proofs are contained in Section A and B.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Experimental details are fully specified in Section 5.1 and Appendix D.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The code for our experiments is included in the supplemental material with instructions for reproduction.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The main training details are specified in Section 5.1. Details such as hyperparameters, tuning protocols, architecture choices, etc. are specified in Section D.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: All of our experiments show results averaged over 3-5 random seeds, and error bars are shown for learning curves. Error bar calculation is specified in Section 5.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer “Yes” if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)

- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: Hardware resources are specified in Section [5.1](#).

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [\[Yes\]](#)

Justification: We have read and conformed to the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [\[NA\]](#)

Justification: Our paper is a theoretical paper on mathematical optimization problems arising in distributed learning. We do not see any direct paths to negative applications beyond those existing for any application of distributed learning or machine learning in general.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our paper does not involve the release of any data or models.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Our paper uses existing datasets (Fashion-MNIST, CIFAR-10), and we cite original sources for both datasets in Section 5.1

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: This paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This paper does not involve crowdsourcing or research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This paper does not involve crowdsourcing or research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.