

# STSyn: Speeding Up Local SGD with Straggler-Tolerant Synchronization

Feng Zhu, Jingjing Zhang, *Member, IEEE* and Xin Wang, *Fellow, IEEE*

**Abstract**—Synchronous local stochastic gradient descent (local SGD) suffers from some workers being idle and random delays due to slow and straggling workers, as it waits for the workers to complete the same amount of local updates. To address this issue, a novel local SGD strategy called STSyn is proposed in this paper. The key point is to wait for the  $K$  fastest workers while keeping all the workers computing continually at each synchronization round, and making full use of any effective (completed) local update of each worker regardless of stragglers. To evaluate the performance of STSyn, an analysis of the average wall-clock time, average number of local updates, and average number of uploading workers per round is provided. The convergence of STSyn is also rigorously established even when the objective function is nonconvex for both homogeneous and heterogeneous data distributions. Experimental results highlight the superiority of STSyn over state-of-the-art schemes, thanks to its straggler-tolerant technique and the inclusion of additional effective local updates at each worker. Furthermore, the impact of system parameters is investigated. By waiting for faster workers and allowing heterogeneous synchronization with different numbers of local updates across workers, STSyn provides substantial improvements both in time and communication efficiency.

**Index Terms**—Local SGD, heterogeneous synchronization, straggler tolerance, distributed learning

## I. INTRODUCTION

AS the size of datasets increases exponentially, it is no longer economic and feasible to leverage all the data to update the model parameters in large-scale machine learning, which magnifies the disadvantage of gradient descent (GD) methods in computational complexity. Instead, stochastic gradient descent (SGD) that uses just one or a mini-batch of samples is proved to converge faster and can achieve significantly reduced computation load [1]–[3]. Therefore, SGD has become the backbone of large-scale machine learning tasks.

To exploit parallelism and further accelerate the training process, the distributed implementation of worker nodes, i.e., distributed learning, has been gaining popularity [4]–[6]. The parameter server (PS) setting is one of the most common scenarios of SGD-based distributed learning frameworks, where in each iteration the PS aggregates the uploaded gradients/models from all the worker nodes to update the global parameter [2], [4], [5], [7]–[9]. In such a setting,

communication cost has become the bottleneck for the performance of the system due to the two-way communication and the larger number of parameters to be updated at each iteration [1], [5], [10]–[12]. Accordingly, a novel scheme justified by the name local SGD is proposed to reduce the communication frequency between the PS and the workers by allowing each worker to perform local updates.

The idea of local SGD is first introduced in [6] where each worker performs a certain number of local updates instead of one, and the PS aggregates the latest model of each worker as the final output. However, the scheme in [6] has been proved to be no better than single-worker algorithms in the worst-case scenario due to the divergence of local optima across workers [13]. To address this issue, [14] proposes the federated averaging (FedAvg) scheme, where each sampled worker runs multiple local updates before uploading the local model to the central server for aggregation, and then the server re-sends the aggregated model back to the workers to repeat the above process. By increasing the communication frequency, FedAvg has been proved to be communication-efficient and enjoy remarkable performance under homogeneous data distribution. The convergence analysis of the latter is widely explored in [15]–[18]. It should also be pointed out that FedAvg is proved to be able to achieve linear speedup concerning the number of workers in this setting. Furthermore, when data is heterogeneously distributed, the convergence of FedAvg is also proved by [19]–[23]. Local SGD has also been proved to achieve better time and communication performance than large mini-batch training with the size of the mini-batch being the same as the total batch size used in a training round of local SGD [24].

**Adaptive Communication Strategies for Local SGD:** It has been proved by [15], [21] that it is theoretically reasonable and advantageous to vary the communication frequency as the training proceeds. The authors of [25] present the convergence analysis of periodic averaging SGD (PASGD), which is FedAvg with no user sampling, and propose an adaptive scheme that gradually decreases the communication frequency between workers and the server to reduce total communication rounds. Meanwhile, [26] also proposes an adaptive strategy named AdaComm based on PASGD where the communication frequency is increased along the training process to achieve a decrease in the true convergence time. Furthermore, [27] and [28] focus on increasing the number of workers that the server waits for per round to enhance communication efficiency. In addition, [29] proposes the STagewise Local SGD (STL-SGD) that increases the communication period along with decreasing learning rate to achieve further improvement.

The authors are with the Key Laboratory for Information Science of Electromagnetic Waves (MoE), Department of Communication Science and Engineering, School of Information Science and Technology, Fudan University, Shanghai 200433, China (e-mail: {20210720072, jingjingzhang, xwang11}@fudan.edu.cn). This work has been supported by the National Natural Science Foundation of China Grant No. 62101134, and by the National Natural Science Foundation of China under Grants No. 62231010 and No. 62071126. (*Corresponding author: Jingjing Zhang.*)

Local SGD presents a possibility to reduce the communication cost of the system; yet, other than communication efficiency, the wall-clock time required for convergence is also an important issue. Due to the discrepancy among computing capabilities of machines, there is considerable variability in the computing time of different machines, and the machines that consume excessively more time per round than others are normally known as stragglers [30]. Numerous researches have been conducted in the development of straggler-tolerant techniques [31], which can be generally categorized into synchronous schemes and asynchronous ones.

**Dealing with Stragglers with Synchronous SGD:** In synchronous parallel SGD, the PS always waits for all the workers to finish their computation. As a result, the whole process could be severely slowed down due to the existence of potential stragglers. One way to mitigate the effect of straggling is to exploit redundancy such as [32]–[36] that allow the PS to wait for only a subset of workers by utilizing storage and computation redundancy. [37] proposes an optimization-based coordinate gradient coding method to mitigate stragglers. Other than redundancy, [38] develops the FlexRR system that attempts to detect and avoid stragglers. [39] introduces a frame-quantization-based method for straggler mitigation in wireless networks. For real-world applications, [40] studies a method for efficiently training spike neural networks (SNN) in federated learning. Another well-known truth is that the convergence analysis of  $M$ -worker synchronous parallel SGD is the same as mini-batch SGD, only with the mini-batch size enlarged  $M$  times. Inspired by this idea, [9] and [41] propose the  $K$ -sync SGD where the PS only waits for the first  $K$  workers to finish their computation. In addition, [41] extends the  $K$ -sync SGD to  $K$ -batch-sync SGD in which the PS updates the parameter using the first  $K$  mini-batches that finish, such that no worker is idle in each training round.

In addition to the approaches mentioned above, there is a growing interest in allowing heterogeneous synchronization to deal with stragglers, where the number of local updates across workers varies in a synchronization round. Heterogeneous synchronization has been studied in several works, including [42]–[44]. In [42], dynamic communication patterns are explored, where workers are manually designated to perform different numbers of local updates under a convex objective global function. FedNova, proposed in [43], balances the heterogeneity of local updates across workers by imposing weights inversely proportional to the number of local updates performed by the workers over the local models. Meanwhile, [44] investigates the impact of incomplete local updates under heterogeneous data distribution by comparing the performances of three schemes. In that work, workers are required to perform a given number of local updates, and by the time of synchronization, those who have finished the task are considered to have “complete” local updates, while the rest are considered to have “incomplete” local updates. Scheme A only utilizes the complete local updates of the workers; Scheme B aggregates both complete and incomplete updates, and Scheme C is the canonical FedAvg scheme. However, all of [42], [43], and [44] do not address the issue of fast workers being idle. Given that, [45] introduces the Anytime Minibatch (AMB) method that

makes full use of the computation of stragglers via averaging over consensus rounds.

**Dealing with Stragglers with Asynchronous SGD:** Apart from confronting stragglers in the synchronous schemes, asynchronous SGD (ASGD) is also proposed to address this issue [46]. In ASGD, the PS updates the parameter immediately after any worker finishes its computation and uploads it to the server. The time used per round of ASGD schemes is much less than that used in synchronous SGD. Nevertheless, since the gradient used to update the parameter may not be the one used to compute it, the trade-off between wall-clock time and the staleness of the gradient becomes the main issue in ASGD, for which some variants of ASGD have been proposed. In [47], the authors develop the Hogwild! scheme which is a lock-free implementation of ASGD and proves its effectiveness both theoretically and experimentally; and [48] improves the performance of ASGD based on Hogwild! with no conflict in parallel execution. In [41], the authors review the  $K$ -async SGD and  $K$ -batch-async SGD first introduced in [9] and [49] respectively, and propose the AdaSync scheme to achieve the best error-runtime trade-off. Another trend of research is to process the delayed gradients in ASGD. To this end, [50] proposes the AdaDelay algorithm that imposes a penalty on delayed gradients, and [51] introduces the DC-ASGD scheme that compensates for the delayed gradients to make full use of all the gradients. [52] introduces a semi-asynchronous FL protocol named SAFA to improve the efficiency of the FL system over model distribution, client selection, and global aggregation. [53] proposes the asynchronous FedBuff method to improve the performance of the system and preserve security at the same time. [31] provided optimal time complexities for parallel optimization methods, under the assumption that the computation time for each stochastic gradient is fixed for each worker. As a followup work of [31], the computation time conforms to some distribution instead of being fixed in [54]. However, both of these two papers assume homogeneous dataset distribution among workers.

### A. Our Contributions

This paper investigates the efficiency of straggler mitigation with heterogeneous synchronization both under homogeneous and heterogeneous data distributions in a distributed PS-based architecture. The main contributions of our work are as follows:

- First, we propose a novel local SGD scheme that allows different numbers of local updates across workers while remaining robust to stragglers. Our approach builds on previous works that use heterogeneous synchronization, but we enhance it with straggler-tolerant techniques and the inclusion of all effective local updates from every worker to improve both time and communication efficiency.
- Second, we provide an analysis of the average runtime, average number of uploading workers, and average number of local updates per round, to justify the improvement in time and communication efficiency of our proposed scheme, named STSyn. We also rigorously establish the

convergence of STSyn and prove that it can achieve a sublinear convergence rate, similar to other local SGD variants, but with a reduced convergence upper bound.

- Finally, we present experimental results to corroborate the superiority of STSyn against state-of-the-art schemes for both homogeneous and heterogeneous data distributions in terms of wall-clock time and communication efficiency. We also study the influence of system hyper-parameters.

The rest of the paper is organized as follows. Section II describes the system model. The development of the proposed STSyn scheme is delineated in Section III. Section IV presents the analysis of STSyn. Numerical results are provided in Section V. Section VI concludes the work.

**Notations:** Boldface lowercase letters and calligraphic letters represent vectors and sets, respectively;  $\mathbb{R}$  denotes the real number fields;  $\mathbb{E}[\cdot]$  denotes the statistical expectation;  $\cup$  denotes the union of sets;  $\mathcal{A} \subseteq \mathcal{B}$  means that set  $\mathcal{A}$  is a subset of set  $\mathcal{B}$ ;  $|\mathcal{A}|$  denotes the number of elements in set  $\mathcal{A}$ ;  $\nabla F$  represents the gradient of the function  $F$ ;  $\|\mathbf{x}\|$  denotes the  $\ell_2$ -norm of vector  $\mathbf{x}$ ;  $\langle \mathbf{x}, \mathbf{y} \rangle$  denotes the inner product of vectors  $\mathbf{x}$  and  $\mathbf{y}$ .

## II. PROBLEM FORMULATION

### A. System Model

We first introduce the general problem framework of distributed SGD with local updates to reduce communication overhead. To implement the local-SGD-based training process, a parameter-server setting with a set of  $M$  workers in set  $\mathcal{M} \triangleq \{1, \dots, M\}$  is considered. Each worker  $m$  maintains a local dataset  $\mathcal{D}_m$ , which is equally allocated from the global training dataset, and we have  $\mathcal{D} = \bigcup_{m \in \mathcal{M}} \mathcal{D}_m$ . Our objective is to minimize the following empirical risk function given as

$$F(\mathbf{w}) = \frac{1}{M} \sum_{m=1}^M F_m(\mathbf{w}), \quad (1)$$

where variable  $\mathbf{w} \in \mathbb{R}^d$  is the  $d$ -dimensional parameter to be optimized; and  $F_m(\mathbf{w})$  is the local loss function of worker  $m$  with  $F_m(\mathbf{w}) \triangleq \mathbb{E}[F_m(\mathbf{w}; \xi_m)]$  where  $\xi_m$  is drawn randomly from the local dataset  $\mathcal{D}_m$ .

To elaborate on the communication and updating protocol, we define  $\mathbf{w}^j$  as the global parameter, and  $\mathbf{w}_m^{j,0}$  as the local parameter that worker  $m \in \mathcal{M}$  has available at training round  $j$  prior to computation.

In local SGD, at each synchronization round  $j$ , the PS first broadcasts the global model parameter  $\mathbf{w}^j$  to a subset of workers  $\mathcal{M}^j \subseteq \mathcal{M}$ . Then by setting the local model as  $\mathbf{w}_m^{j,0} = \mathbf{w}^j$ , each worker  $m$  in  $\mathcal{M}^j$  starts to perform local updates. At the end of round  $j$ , each worker  $m$  is assumed to complete  $U_m^j \geq 0$  local updates, with each given as

$$\mathbf{w}_m^{j,u+1} = \mathbf{w}_m^{j,u} - \frac{\alpha}{B} \sum_{b=1}^B \nabla F_m(\mathbf{w}_m^{j,u}; \xi_{m,b}^{j,u}), \quad (2)$$

for any  $u = 0, \dots, U_m^j - 1$  if  $U_m^j \geq 1$ . Here  $\alpha$  is the stepsize; the mini-batch  $\{\xi_{m,b}^{j,u}\}_{b=1}^B$  of size  $B$  is generated randomly for each local update from dataset  $\mathcal{D}_m$ ; and  $\nabla F_m(\mathbf{w}_m^{j,u}; \xi_{m,b}^{j,u})$  is the local gradient of loss function  $F$  computed with local

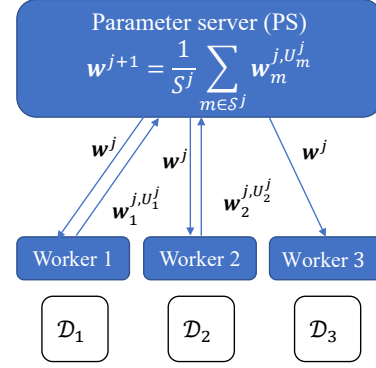


Fig. 1. Illustration of the PS setting with  $M = 3$ ,  $\mathcal{M}^j = \{1, 2, 3\}$  and  $S^j = \{1, 2\}$  at the  $j$ th round.

parameter  $\mathbf{w}_m^{j,u}$  and data sample  $\xi_{m,b}^{j,u}$ . After the computation, the PS then selects a subset  $S^j \subseteq \mathcal{M}^j$  of  $S^j$  workers (i.e.,  $|S^j| = S^j$ ) to upload their local parameters, yielding the global parameter  $\mathbf{w}^{j+1}$  given as

$$\mathbf{w}^{j+1} = \frac{1}{S^j} \sum_{m \in S^j} \mathbf{w}_m^{j,U_m^j}. \quad (3)$$

The next round  $j + 1$  then starts and the training continues until a convergence criterion is satisfied, i.e., the trained model reaches a target test accuracy, with the total number of rounds defined as  $J$ . The whole process is illustrated in Fig. 1.

### B. Preliminaries

Based on the distributed PS architecture under consideration, our proposed scheme stems from a couple of local SGD variants, which we briefly review next.

**FedAvg.** One pioneering work on local updates is the so-called FedAvg. More precisely, at each round, the PS first randomly selects a subset  $\mathcal{M}^j \subseteq \mathcal{M}$  of workers to broadcast the global parameter to. Then each selected worker  $m$  performs the same amount  $U_m^j = U$  of local updates and sends the local parameter back to the PS, i.e., we have  $S^j = \mathcal{M}^j$ . In particular, when all the workers participate in the computation in each round, i.e., with  $S^j = \mathcal{M}^j = \mathcal{M}$ , we can arrive at periodic-averaging SGD (PASGD) [25].

**Adaptive Communication Strategies.** Based on PASGD, [26] proposes an adaptive communication strategy called AdaComm in order to achieve a faster convergence speed in terms of wall-clock runtime. This is achieved by adaptively increasing the communication frequency between the PS and the workers. Note that each worker still performs the same number of local updates in each training round. Similar to AdaComm, FLANP proposed in [27] is also an adaptive strategy that flexibly changes the number of workers the server waits for in each round. Intuitively, in the beginning stage of training, the server only needs samples from fewer workers to gain a coarse direction of optimization. As the training process moves along, more samples/gradients are needed to provide a finer direction. In particular, each worker performs the same number of local updates  $U$  at each round  $j$ .

**FedNova.** This is one of the pioneering works to investigate heterogeneous synchronization due to the discrepancy among

the computing capabilities of machines [43], i.e., the numbers  $\{U_m^j\}_{m \in \mathcal{M}}$  of local updates across workers are not necessarily the same. FedNova aggregates the local models in a weighted manner to balance the heterogeneity of local updates, with the weight inversely proportional to the number of local updates the worker has completed.

### C. Performance Metrics

In addition to training/learning accuracy, the time consumption and communication cost with the distributed learning scheme are also important performance metrics. Indeed, the training accuracy should be evaluated as a function of time and/or communication cost to fairly gauge the time/communication efficiency of a learning scheme. To this end, we specifically define wall-clock runtime and the communication cost as follows.

Firstly, the wall-clock runtime for each worker  $m$  at the local iteration  $u$  ( $u < U_m^j$ ) of round  $j$  is defined as the time elapsed while computing the mini-batch gradient  $\frac{1}{B} \sum_{b=1}^B \nabla F(\mathbf{w}_m^{j,u}; \xi_{m,b}^{j,u})$ . It is denoted by  $T_m^{j,u}$ , and is assumed to conform to exponential distribution with mean  $\mu$ . We also assume that the random variable  $T_m^{j,u}$  is independent across all workers, rounds and local iterations, which is a common assumption since workers are similar machines working independently [4], [32], [41]. Therefore, the total time consumed at round  $j$  can be written as

$$T^j = \max_{m \in \mathcal{S}^j} \left\{ \sum_{u=0}^{U_m^j-1} T_m^{j,u} \right\}. \quad (4)$$

Note that in particular, although we both assume exponential computing time for each worker as in [41], the analysis we perform is completely disparate. [41] analyzes the wall-clock time for synchronous/asynchronous SGD without even considering local SGD, let alone heterogeneous synchronization.

Secondly, note that the global model parameter broadcast from the PS to workers and the local parameters sent from the workers to the PS are of the same size  $d$ . Hence, at each synchronization round  $j$ , we can define the communication cost  $C^j$  as the sum of the number  $|\mathcal{M}^j|$  of workers that download the global model from the PS, and the number  $|\mathcal{S}^j|$  of workers that send the local update to the PS; i.e.,

$$C^j = |\mathcal{M}^j| + |\mathcal{S}^j|. \quad (5)$$

This model is extensively applied in the analysis of communication efficiency, as demonstrated in key studies like [55]–[57].

These two performance metrics will be studied in the theoretical analysis in Section IV and will be used to evaluate the time/communication efficiency of the schemes in the experimental results in Section V.

## III. LOCAL SGD WITH HETEROGENEOUS SYNCHRONIZATION (STSYN)

In this section, we introduce STSyn, a novel strategy that aims to accelerate local SGD both under i.i.d. and non-i.i.d. data with heterogeneous synchronization. The key points are two-fold. First, in each training round, no worker stops

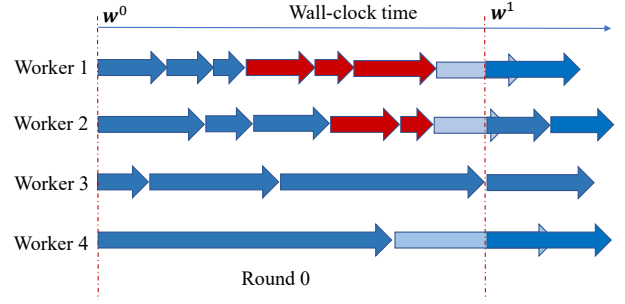


Fig. 2. Illustration of STSyn with  $M = 4$ ,  $U = 3$  and  $K = 3$ . In round 0, workers 1, 2, and 3 are the fastest  $K = 3$  workers that have completed  $U = 3$  local updates; the red arrows represent the additional local updates performed by the fastest  $K - 1 = 2$  workers; and the arrows in light blue represent the straggling updates that are canceled.

computing until the  $K$ th fastest one completes a specific number  $U$  of local updates. Hence, the faster the worker is, the more local updates it computes. Unlike typical local SGD methods where only the selected workers carry out the same number of local updates in each round, STSyn ensures that no worker remains idle, thus taking advantage of both straggler mitigation and training speedup. Second, in the aggregation phase, any worker that completes at least one local update, i.e., with  $U_m^j \geq 1$ , sends its local parameter back to the PS. By increasing the number of workers that participate in the aggregation with each effective local update, STSyn can nicely balance the heterogeneity among workers and accelerate the training process. This also brings the benefit of communication efficiency in the sense that the reduction of synchronization rounds can compensate for the increase in communication cost at each round.

Building on these ideas, we propose the novel STSyn scheme as follows. In particular, at each training round  $j$ , the PS sends the current global model parameter  $\mathbf{w}^j$  to all the workers in  $\mathcal{M}$ , i.e., we have  $\mathcal{M}^j = \mathcal{M}$ . Then each worker  $m$  sets its local model parameter  $\mathbf{w}_m^{j,0} = \mathbf{w}^j$  and begins its computation. After computing the first  $U$  local updates, each worker acknowledges the completion but continues to update locally. Once the PS receives  $K$  ( $K \leq M$ ) acknowledgments, it terminates the computation of round  $j$ . However, for the rest  $M - K$  straggling workers, they might have finished at least one local update, i.e., with  $U_m^j \geq 1$ , which can also be used for the global update. In other words, at round  $j$ , we might have a subset  $\mathcal{S}^j \subseteq \mathcal{M}$  with  $|\mathcal{S}^j| \geq K$  workers that have completed at least one update.

At the end of training round  $j$ , each worker in the subset  $\mathcal{S}^j$  (instead of the set of the fastest  $K$  workers) is required to upload its latest model parameter  $\mathbf{w}_m^{j,U_m^j}$ . Then, the PS updates the global parameter  $\mathbf{w}^j$  through (3). The next round  $j+1$  starts similarly and the training continues. An illustration of the STSyn scheme is shown in Fig. 2. Note that the cost and time elapsed of communicating the acknowledgment are assumed negligible here as they are typically much smaller than those with the local parameters. The complete procedure of STSyn is summarized in Algorithm 1.

*Merits of STSyn:* The compound of the additional local

**Algorithm 1** STSyn

---

**Input:** number of workers  $M$ , constant  $U, K$ , stepsize  $\alpha > 0$ , batch size  $B$

**Server executes:**

Initiate  $\mathbf{w}^0$

**for**  $j = 0, 1, \dots, J - 1$  **do**

Sends the global model  $\mathbf{w}^j$  to all the workers in  $\mathcal{M}$

**for each worker**  $m \in \mathcal{M}$  **do**

Executes  $\text{WorkerUpdate}(m, \mathbf{w}^j)$

**end for**

**if** receives  $K$  acknowledgements **then**

Stops the computation of the round and collects the local models  $\{\mathbf{w}_m^{j, U_m^j}\}_{m \in \mathcal{S}^j}$

**end if**

Updates  $\mathbf{w}^j$  with (3)

**end for**

**WorkerUpdate**( $m, \mathbf{w}^j$ ):

Sets local model  $\mathbf{w}_m^{j, 0} = \mathbf{w}^j$  and  $U_m^j = 0$

**repeat**

Updates its local model via (2) and sets  $U_m^j = U_m^j + 1$

**if**  $U_m^j == U$  **then**

Sends an acknowledgment to the server

**end if**

**stops computing immediately** when noticed (by the server) the end of current round

**if**  $U_m^j \geq 1$  **then**

$\mathcal{S}^j \ni m$

Uploads its latest model  $\mathbf{w}_m^{j, U_m^j}$  to the server

**end if**

---

updates from both the fastest  $K$  workers and the rest  $M - K$  stragglers brings multiplicative benefits; consequently, STSyn can outperform its local SGD counterparts. Intuitively, compared to FedAvg, with the choice  $U_m^j = U$  and  $\mathcal{M}^j = \mathcal{M}$ , STSyn collects more local updates in a training round and also only has to wait for the fastest  $K$  workers, thereby achieving better time and communication efficiency. Compared to FedNova, which relies on weighted heterogeneous synchronization, STSyn presents a distinct advantage without the need for such weights, as will be demonstrated in detail in Section V. In a nutshell, the straggler-tolerant techniques and the additional effective local updates together endow STSyn with the advantages both in wall-clock time and communication efficiency over state-of-the-art schemes.

#### IV. PERFORMANCE ANALYSIS

In this section, we provide the analysis of the average runtime, average number of local updates, and average number of uploading workers for each training round along with numerical illustrations. The benefits from additional effective local updates from the  $K$  fastest workers and the rest  $M - K$  stragglers are highlighted via comparison with PASGD. In addition, convergence of the proposed STSyn is rigorously established.

##### A. Average Wall-Clock Runtime

As specified in Subsection II-C, we assume that the wall-clock times  $\{T_m^{j,u}\}$  are i.i.d. exponential random variables across workers, rounds and local iterations, i.e., we have  $T_m^{j,u} \sim \text{Exp}(\mu)$  [4], [32].

At each round  $j$ , to carry out  $U_m^j = U$  local updates, the runtime  $T_m^j$  of each worker  $m$  is given as  $T_m^j := \sum_{u=0}^{U-1} T_m^{j,u}$ . It then readily follows that  $T_m^j$  is an Erlang random variable with mean  $U\mu$  and variance  $U\mu^2$ . Since the PS only waits for the  $K$ th fastest worker to finish  $U$  local updates, the average wall-clock time  $\bar{T}$  taken per round can be computed as

$$\bar{T} = \mathbb{E}[T_{K:M}^j] = \frac{\mu M!}{(K-1)!(M-K)!} I \quad (6)$$

where  $T_{K:M}^j$  is the  $K$ th order statistic of variables  $\{T_m^j\}_{m \in \mathcal{M}}$ , as analyzed in [58], and we have the integral

$$I = \int_0^\infty x \{Q(x)\}^{K-1} \{1 - Q(x)\}^{M-K} q(x) dx, \quad (7)$$

with  $q(x)$  and  $Q(x)$  being the functions

$$q(x) = \frac{x^{U-1} \exp(-x/\mu)}{\mu^U (U-1)!},$$

$$Q(x) = \frac{\int_0^x t^{U-1} \exp(-t) dt}{(U-1)!}.$$

##### B. Average Number of Local Updates and Uploading Workers

**Average Number of Local Updates per Round.** It is known that the average runtime per round is  $\mathbb{E}[T_{K:M}^j]$  and we have  $T_m^{j,u} \sim \text{Exp}(\mu)$ . As a result, the average number  $\bar{U}$  of local updates that each worker  $m$  performs can be derived as

$$\bar{U} = \mathbb{E}[U_m^j] = \frac{\mathbb{E}[T_{K:M}^j]}{\mu} = \frac{M!}{(K-1)!(M-K)!} I. \quad (8)$$

It is easy to see that  $\bar{U}$  increases as  $K$  since  $\mathbb{E}[T_{K:M}^j]$  does. On one hand, a larger  $K$  indicates more local updates at each worker, which are beneficial for communication efficiency to some extent. On the other hand, the robustness against stragglers is decreased as we need to wait for more workers. Therefore, a trade-off between communication and time efficiency should be addressed, which is further explored through experiments in Section V.

**Average Number of Uploading Workers per Round.** Note that in STSyn, any worker that can finish at least one local update sends back its completed local parameter. To elaborate, for any integer  $s \leq M$ , we define as  $T_{s:M}^{j,0}$  the  $s$ th order statistic of i.i.d. variables  $\{T_m^{j,0}\}_{m=1}^M$ , where  $T_m^{j,0}$  is the runtime of worker  $m$  consumed for the first local update. Then within runtime  $T_{K:M}^j$  at round  $j$ , the number of selected workers  $S^j$  can be given as

$$S^j = \max \left\{ s \mid T_{s:M}^{j,0} \leq T_{K:M}^j \right\}. \quad (9)$$

For exponential distribution, the average of  $T_{s:M}^{j,0}$  can be approximated as [41],

$$\mathbb{E}[T_{s:M}^{j,0}] = \sum_{i=M-s+1}^M \frac{\mu}{i} \approx \mu \log \frac{M}{M-s}. \quad (10)$$

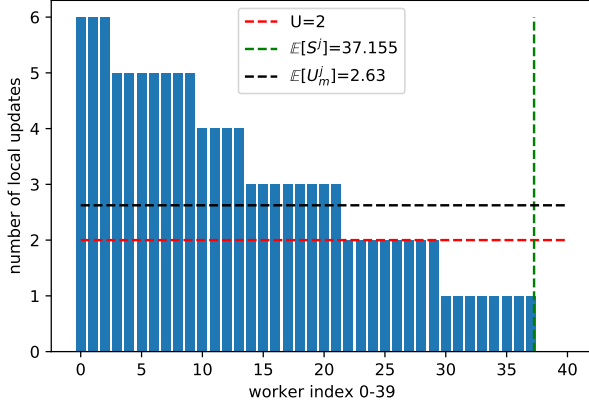


Fig. 3. The numbers of local updates across workers in a training round with  $M = 40$ ,  $K = 30$ , and  $U = 2$ .

By using (10) and (9),  $\mathbb{E}[S^j]$  can be derived as

$$\mathbb{E}[S^j] \approx M - \exp\left(\log M - \frac{M!}{(K-1)!(M-K)!}I\right). \quad (11)$$

It is easy to see that  $\mathbb{E}[S^j]$  is also positively related to  $K$  and  $U$ . That is to say, as the PS waits for more workers to finish more local iterations, a larger number of workers could complete at least one update and be selected to upload their local parameters at each round.

Fig. 3 shows an example of a certain training round of STSyn under the CIFAR-10 dataset using a CNN neural network with  $M = 40$ ,  $K = 30$ , and  $U = 2$ . It can be observed that almost half of the workers have completed more than  $U = 2$  local updates and 7 stragglers have one. These additional local updates are beneficial to the convergence speed, as will be illustrated in the sequel. Over 200 rounds, we also calculate the average number of local updates and that of workers selected to upload per round. As indicated in Fig. 3, the values are 2.63 and 37.155, both of them are consistent with the derived results  $\mathbb{E}[U_m^j] = 2.6352$  and  $\mathbb{E}[S^j] = 37.132$  in (8) and (11), respectively.

**Benefits from Additional Effective Local Updates.** By keeping all the workers computing until the end of each round, STSyn makes full use of any completed (effective) local update. This includes the additional ones performed both by the fastest  $K$  workers and by the rest  $M - K$  stragglers (see Fig. 3), as compared to the typical local SGD methods such as PASGD with  $U_m^j = U$ .

The respective benefits of exploring both types are illustrated in Fig. 4, where we trained a CNN on the CIFAR-10 dataset with  $M = 20$ ,  $K = 10$ , and  $U = 5$ . Data is heterogeneously distributed among workers. In particular, STSyn\_modified1 merely utilizes the additional local updates of the fastest workers, i.e., we have  $U_m^j \geq U$  for the  $K$  fastest workers and  $S^j = K$ ; while STSyn\_modified2 only keeps the additional ones performed by the stragglers, i.e., we have  $U_m^j = U$  for the  $K$  fastest workers and  $S^j \geq K$ , which is equivalent to scheme B in [44]. It can be observed that both STSyn\_modified1 and STSyn\_modified2 require fewer training rounds than the baseline PASGD to reach the same test accuracy. In other words, any type of effective local update

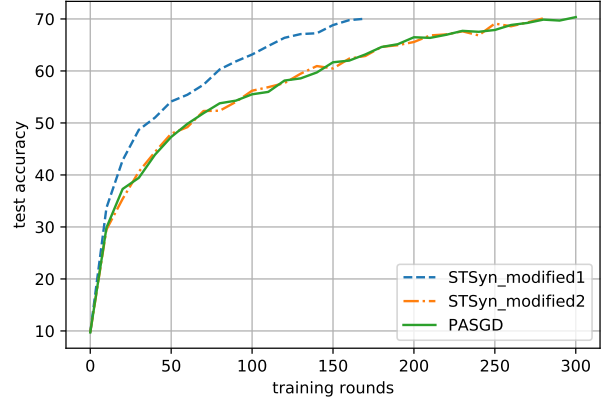


Fig. 4. Illustration of the impact of additional effective local updates on learning accuracy. In contrast to STSyn, STSyn\_modified1 is fixed with  $S^j = K$ ; and STSyn\_modified2 is fixed with  $U_m^j = U$ .

is favorable. It is also evidently seen that the ones of the fastest workers are indeed more useful. Meanwhile, the local updates from the stragglers are also contributive to balance the heterogeneity in synchronization and mitigate the client-drift effect [59]. Furthermore, the joint effect of the two types is magnificent, which will be shown in Section V.

### C. Convergence Analysis

We next rigorously establish the convergence of STSyn for both non-i.i.d. and i.i.d. data distributions. Our convergence analysis is based on the following standard assumptions, which are widely adopted in e.g., [16], [19], [23], [60].

**Assumption 1** (Unbiasedness and Variance Boundedness) *For fixed model parameter vector  $\mathbf{w}$ , the stochastic gradient  $\nabla F_m(\mathbf{w}; \xi)$  is an unbiased estimator of the true gradient  $\nabla F_m(\mathbf{w})$ , i.e., we have*

$$\mathbb{E}_\xi[\nabla F_m(\mathbf{w}; \xi)] = \nabla F_m(\mathbf{w}). \quad (12)$$

*Moreover, there exists two scalars  $\sigma_L > 0$  and  $\sigma_G > 0$  such that*

$$\mathbb{E}_\xi \left[ \|\nabla F_m(\mathbf{w}; \xi) - \nabla F_m(\mathbf{w})\|^2 \right] \leq \sigma_L^2. \quad (13)$$

$$\mathbb{E} \left[ \|\nabla F_m(\mathbf{w}) - \nabla F(\mathbf{w})\|^2 \right] \leq \sigma_G^2. \quad (14)$$

**Assumption 2** (Smoothness and Lower Boundedness) *The local loss function  $F_m$  is  $L$ -smooth, i.e.,*

$$\|\nabla F_m(\mathbf{w}_1) - \nabla F_m(\mathbf{w}_2)\| \leq L \|\mathbf{w}_1 - \mathbf{w}_2\|, \quad (15)$$

$$F_m(\mathbf{w}_1) - F_m(\mathbf{w}_2) \leq \langle \nabla F_m(\mathbf{w}_2), \mathbf{w}_1 - \mathbf{w}_2 \rangle + \frac{L}{2} \|\mathbf{w}_1 - \mathbf{w}_2\|^2, \quad (16)$$

$\forall \mathbf{w}_1, \mathbf{w}_2 \in \mathbb{R}^d$ . We also assume that the objective function  $F$  is bounded below by a scalar  $F^*$ . Based on these assumptions, we can then derive an upper bound for the expectation of the average squared gradient norms  $\frac{1}{J} \mathbb{E} \left[ \sum_{j=0}^{J-1} \|\nabla F(\mathbf{w}^j)\|^2 \right]$  of the objective function  $F$ , which is a common metric used for convergence analysis with general nonconvex objectives.

**Theorem 1:** *Let the constant stepsize  $\alpha$  be chosen such that  $\alpha \leq \min \left\{ \frac{\bar{U}}{\sqrt{30LU_{max}^2}}, \frac{K^2}{LUM^2} \right\}$ . Under Assumptions 1 and 2,*

the expectation of the average squared norm of the gradients with heterogeneous data distribution can be upper bounded by

$$\frac{1}{J} \mathbb{E} \left[ \sum_{j=0}^{J-1} \|\nabla F(\mathbf{w}^j)\|^2 \right] \leq \frac{F(\mathbf{w}^0) - F^*}{c\alpha\bar{U}J} + \frac{1}{c} \left[ \frac{\alpha U_{max} L}{2KB\bar{U}} \sigma_L^2 + \frac{5\alpha^2 U_{max}^3 L^2}{2\bar{U}^2} (\sigma_L^2 + 6U_{max} \sigma_G^2) \right], \quad (17)$$

where  $c$  is a constant and we have bounded  $U_m^j \leq U_{max}$ , with  $U_{max} := \max\{U_m^j\}_{m \in \mathcal{M}, 0 \leq j \leq J-1}$ , as the possibility that  $U_m^j$  becomes unbounded is almost zero; we also used the expression  $\bar{U} = \mathbb{E}[U_m^j]$ , as defined in Subsection IV-A.

*Proof.* The proof is presented in Appendix A.  $\square$

**Corollary 1:** Under the conditions stated in Theorem 1, by selecting  $\alpha = \frac{\sqrt{K}}{\sqrt{\bar{U}JL}}$ , the convergence rate of STSyn for heterogeneous data distribution is

$$\frac{1}{J} \mathbb{E} \left[ \sum_{j=0}^{J-1} \|\nabla F(\mathbf{w}^j)\|^2 \right] = \mathcal{O} \left( \frac{1}{\sqrt{K\bar{U}J}} + \frac{K}{\bar{U}^3 J} \right). \quad (18)$$

Theorem 1 and Corollary 1 state that STSyn can achieve a sublinear convergence rate  $\mathcal{O} \left( \frac{1}{\sqrt{K\bar{U}J}} + \frac{K}{\bar{U}^3 J} \right)$ , providing a theoretical guarantee for the performance of STSyn, which is consistent with the result of local SGD in heterogeneous data case. Since [61] proves the fastest convergence rate of local SGD for nonconvex objective functions under heterogeneous data distribution, it would be a good baseline for comparison. Compared with the convergence upper bound in [61] where each uploading worker performs the same number  $U$  of local updates, here in (17) the role of  $U$  is replaced with  $\bar{U}$  and  $U_{max}$ , which are the average and upper bound of  $U_m^j$ , respectively. Since  $\bar{U}$  is larger than  $U$  as discussed in Section IV-B, the first term of the RHS of (17) is smaller, arriving at a faster convergence rate.

*Remark:* The analysis of the i.i.d. data distribution follows a similar approach to the non-i.i.d. case. By setting  $\sigma_G = 0$  in Assumption 1, we can directly deduce the convergence results for the homogeneous distribution. Consequently, the convergence rate of STSyn under i.i.d. data can be expressed as  $\mathcal{O} \left( \frac{1}{\sqrt{K\bar{U}J}} + \frac{K}{\bar{U}^3 J} \right)$  by selecting  $\alpha = \frac{\sqrt{K}}{\sqrt{\bar{U}JL}}$ .

## V. EXPERIMENTS

In this section, we evaluate the effectiveness of the proposed STSyn. For comparison, we consider state-of-the-art schemes including FedNova [43], PASGD [25] and AdaComm [26] mentioned in Subsection II-B. In addition, we explore the effect of the hyper-parameters  $K$  and  $U$  on the performance of STSyn.

### A. Experimental Setting and Metric Calculation

**Experimental Setting.** We consider training a three-layer CNN on the CIFAR-10 dataset with both homogeneous and heterogeneous data distribution. To be specific, for the homogeneous case, data is first shuffled randomly and then distributed to each worker with equal size in order; for the

heterogeneous case, all training data samples are first sorted according to their labels and then equally (in size) allocated to the workers in order. We set a total of  $M = 20$  workers for the training process. The wall clock times for each worker to compute a minibatch gradient is i.i.d. from an exponential distribution across workers, local iterations, and rounds. The mean of the exponential distribution is set to  $\mu = 10^{-4}$  sec. In particular, each worker maintains a tuple recording the number of local steps it has accomplished and the amount of time elapsed. For STSyn, the time for each round is recorded once the  $K$  fastest worker has finished  $U$  local steps. The stepsize  $\alpha$  is 0.1 and the batch size  $B$  is 100 for each scheme.

**Metric Calculation.** We determine the per-round metrics, i.e., wall-clock time and communication cost, for FedNova, PASGD, AdaComm and the proposed STSyn, as follows. Note that, we generally take one synchronous communication from the workers to the PS as one round.

*Wall-clock time.* In FedNova, the number of local updates performed across workers conform to some distribution and the PS has to wait for all the workers. Hence, the wall-clock time consumed at round  $j$  can be expressed as

$$T_{FedNova}^j = Q_{M:M}^j, \quad (19)$$

where  $Q_m^j := \sum_{u=0}^{U_m^j-1} T_m^{j,u}$  and  $Q_{K:M}^j$  is the  $K$ th order statistic of  $\{Q_m^j\}_{m \in \mathcal{M}}$ .

For PASGD, each worker performs  $U$  local updates per round and the PS waits for the slowest worker to finish. Therefore the per-round wall clock time is

$$T_{PASGD}^j = T_{M:M}^j. \quad (20)$$

For AdaComm, in each round  $j$ , each worker carries out  $\tau^j$  local updates, with  $\tau^j$  being decided adaptively and the PS also waits for all the workers. Then the per-round wall-clock time for AdaComm is

$$T_{AdaComm}^j = V_{M:M}^j, \quad (21)$$

where  $V_m^j := \sum_{u=0}^{\tau^j-1} T_m^{j,u}$  and  $V_{K:M}^j$  is the  $K$ th order statistic of  $\{V_m^j\}_{m \in \mathcal{M}}$ .

The wall-clock time for FLANP at round  $j$  is

$$T_{FLANP}^j = T_{S^j:M}^j. \quad (22)$$

since the server waits for the fastest  $S^j$  workers.

Finally, as shown in Section IV, the wall-clock time for STSyn at round  $j$  is given as

$$T_{STSyn}^j = T_{K:M}^j. \quad (23)$$

*Communication cost.* For FedNova, in each round, all the workers upload their local models. Therefore, the communication cost at round  $j$  can be written as

$$C_{FedNova}^j = 2M. \quad (24)$$

Similarly, the communication cost for PASGD is


$$C_{PASGD}^j = 2M. \quad (25)$$

For AdaComm, the per-round communication cost is


$$C_{AdaComm}^j = 2M, \quad (26)$$

TABLE I  
COMPARISON OF DIFFERENT ALGORITHMS REGARDING DIFFERENT METRICS


	Time at round $j$	Communication at round $j$	Number of local updates at round $j$	Convergence rate
FedNova	$Q_{M:M}^j$	$2M$	$\{U_m^j\}_{m \in \mathcal{M}, 0 \leq j \leq J-1}$	$\mathcal{O}\left(\frac{1}{\sqrt{KUJ}} + \frac{1}{UJ}\right)$ (nonconvex)
PASGD	$T_{M:M}^j$	$2M$	$U$	$\mathcal{O}\left(\frac{1}{\sqrt{KUJ}} + \frac{1}{J}\right)$ (nonconvex)
AdaComm	$V_{M:M}^j$	$2M$	$\{U^j\}_{0 \leq j \leq J-1}$	$\mathcal{O}\left(\frac{1}{\sqrt{KJ}} + \frac{KU}{J}\right)$ (nonconvex)
FLANP	$T_{S^j:M}^j$	$M + S^j$	$U$	$\mathcal{O}\left(\frac{1}{J}\right)$ ( $\mu$ -strongly convex)
STSyn	$T_{K:M}^j$	$M + S^j$	$\{U_m^j\}_{m \in \mathcal{M}, 0 \leq j \leq J-1}$	$\mathcal{O}\left(\frac{1}{\sqrt{KUJ}} + \frac{K}{U^3 J}\right)$ (nonconvex)




(a) i.i.d. distribution




(a) i.i.d. distribution




(a) i.i.d. distribution



(b) non-i.i.d. distribution



(b) non-i.i.d. distribution



(b) non-i.i.d. distribution

Fig. 5. Comparison of test accuracy against wall-clock time for FedNova, AdaComm, PASGD, FLANP, and STSyn with a) i.i.d. distribution and b) non-i.i.d. distribution.

Fig. 6. Comparison of test accuracy versus communication cost for FedNova, AdaComm, PASGD, FLANP, and STSyn with i.i.d. and non-i.i.d. distribution.

Fig. 7. Influence of hyper-parameter  $K$  of STSyn in terms of time-to-accuracy performance with i.i.d. and non-i.i.d. data distribution.

since the PS has to wait for all the workers in  $\mathcal{M}$ .

For FLANP, the communication cost at round  $j$  is

$$C_{FLANP}^j = M + S^j, \quad (27)$$

since all the workers download the global parameter and the number of workers selected to upload each round changes accordingly.

The communication cost for STSyn at round  $j$  is

$$C_{STSyn}^j = M + S^j, \quad (28)$$

since in STSyn all the workers in  $\mathcal{M}$  download the parameter  $\mathbf{w}^j$  from the PS but only the ones in the subset  $S^j$  upload their local parameters. We summarize the results in Table I, together with the metrics of number of local updates and convergence rate, where  $\bar{U}$  is the average number of local updates at each round. It can be observed that STSyn also achieves the best convergence rate supposing  $\bar{U}^2 \geq K$ .


## B. Experimental Results

**STSyn Outperforms State-of-the-Art Schemes.** Fig. 5 and 6 compare the performance of FedNova, AdaComm, PASGD,


and STSyn in terms of test accuracy for both the i.i.d. and non-i.i.d. cases. The local update numbers for FedNova conform to an exponential distribution with a mean of 10, while the initial number  $U_{max}$  of local updates for AdaComm is set to be 10. PASGD uses a fixed number of 10 local updates, and STSyn is set with  $K = 5, U = 10$  in both cases. All schemes stop training once the accuracy reaches 70%. The updated formula for FedNova is provided below for reference.

$$\begin{aligned} \mathbf{w}_m^{j,u+1} &= \mathbf{w}_m^{j,u} - \frac{\alpha}{B} \sum_{b=1}^B \nabla F_m(\mathbf{w}_m^{j,u}; \xi_{m,b}^{j,u}) \\ \mathbf{w}^{j+1} &= \mathbf{w}^j + \frac{1}{M} \sum_{m \in \mathcal{M}} \frac{\sum_{m \in \mathcal{M}} U_m^j}{MU_m^j} (\mathbf{w}_m^{j,U_m^j} - \mathbf{w}^j). \end{aligned} \quad (29)$$


Fig. 5 (a) and (b) provide a comprehensive comparison of the test accuracy of different schemes against wall-clock time under both homogeneous and heterogeneous data distributions. For both cases, it is clear that AdaComm performs better than PASGD by selecting the optimal communication frequency  $U$  in different time intervals. In contrast, the weighted aggregation technique employed in FedNova proves to be ineffective




(a) i.i.d. distribution




(a) i.i.d. distribution




(a) i.i.d. distribution



(b) non-i.i.d. distribution



(b) non-i.i.d. distribution



(b) non-i.i.d. distribution

Fig. 8. Influence of hyper-parameter  $K$  of STSyn in terms of communication-to-accuracy performance with i.i.d. and non-i.i.d. data distribution.

Fig. 9. Influence of hyper-parameter  $U$  of STSyn in terms of time-to-accuracy performance with i.i.d. and non-i.i.d. data distribution.

Fig. 10. Influence of hyper-parameter  $U$  of STSyn in terms of time-to-accuracy performance with i.i.d. and non-i.i.d. data distribution.

in these scenarios, resulting in the poorest performance among the schemes considered. Out of all the schemes, STSyn stands out in both cases, thanks to its incorporation of additional local updates and its robustness against stragglers, a key feature not present in the other three methods. We have to note here that FLANP is the fastest algorithm in the i.i.d. case because of the adaptive scheme but the performance degrades severely when data is non-i.i.d. This is because, in the non-i.i.d. environment, it is far from sufficient to provide a decent direction for optimization with a few workers' gradients due to heterogeneity (We do not equip FLANP with the FedGATE solver as proposed in that paper to make a fair comparison.), which reflects the advantage of STSyn that includes as many workers that have finished at least one update as possible to enable a good start. It is worth noting that the curves corresponding to the i.i.d. distribution appear smoother compared to the other cases. This smoothness is attributed to the smaller variance inherent in the homogeneous data distribution.

Fig. 6 (a) and (b) present the performances of the communication efficiency exhibited by different schemes. In the i.i.d. case, STSyn demonstrates notable performance, thanks to the inclusion of additional effective local updates, particularly those with  $U_m^j \leq U$ . These additional updates effectively balance the synchronization process and contribute to the improved communication efficiency observed in STSyn. As depicted in Fig. 5 (a), FLANP offers a slight edge over STSyn in i.i.d. cases, owing to its adaptive nature. On the other hand, PASGD incurs a higher communication cost compared to both STSyn and FedNova, primarily due to its limited utilization of local updates.

In the more complex non-i.i.d. scenario, all five schemes demonstrate similar levels of communication efficiency, with

the notable exception of FLANP. FLANP performs the worst since it requires a larger number of rounds to achieve the target accuracy. This is largely due to the client-drift phenomenon, which occurs as a result of data heterogeneity among workers, as discussed in prior studies [59], [62]. In other words, the communication efficiency does not consistently increase with the number of local updates.

In summary, the STSyn algorithm stands out as a notably superior solution, consistently surpassing its counterparts in terms of both time and communication efficiency across various data distribution scenarios. This consistent performance highlights the algorithm's robustness and adaptability.

**Effect of Hyper-Parameter  $K$ .** By comparing STSyn schemes with varying values of  $K$ , we can observe the trade-off between wall-clock time and communication cost.

Fig. 7 (a) and (b) depict the performances of test accuracy against wall-clock time for different values of  $K$  under both i.i.d. data and non-i.i.d. data distributions. In the case of i.i.d. data, the wall-clock time remains relatively stable when  $K$  is small. As  $K$  increases, the training process converges faster, leading to a reduction in the number of required rounds. This reduction offsets the increase in average round time caused by including more workers. However, it should be noted that as  $K$  approaches the total number  $M$ , this compensation effect diminishes, resulting in an overall increase in wall-clock time. Conversely, under non-i.i.d. data distributions, the situation differs. As  $K$  increases, the resulting speedup is unable to fully compensate for the reduction in the number of rounds due to the heterogeneity of the data. As a result, the wall-clock time performance degrades as  $K$  increases.

Fig. 8 (a) and (b) demonstrate the communication-to-accuracy performance of STSyn for different values of  $K$




Fig. 11. Comparison of test accuracy against a) wall-clock time and b) communication cost for FedNova, AdaComm, PASGD, FLANP, and STSyn with non-i.i.d. distribution under the weaker assumption.

under both i.i.d. and non-i.i.d. data distributions. In contrast to the observations regarding wall-clock time, it is obvious that a larger value of  $K$ , corresponding to a larger  $S^j$ , results in higher communication efficiency for i.i.d. data. This finding aligns with the analysis provided in Theorem 1, which explains that a larger  $K$  accelerates convergence, and the reduction in the number of training rounds outweighs the increase in per-round communication cost. In the case of non-i.i.d. data, the variation in  $K$  has minimal impact. This outcome primarily arises from the counterbalancing effects of the reduced number of training rounds and the increased per-round communication load.

In summary, when considering an i.i.d. data distribution, a larger value of  $K$  is preferred as it leads to improved communication efficiency while maintaining a relatively high level of time efficiency. On the other hand, for a non-i.i.d. data distribution, a smaller value of  $K$  is a better choice as it allows STSyn to achieve better time efficiency while still maintaining decent communication efficiency.

**Effect of Hyper-Parameter  $U$ .** As discussed in Section IV, there is a positive relationship between the value of  $U$  and the number  $S^j$  of workers selected per round. Consequently, increasing the value of  $U$  results in the selection of a greater number of workers in each round.

Fig. 9 (a) and (b) showcase the time-to-accuracy performances of STSyn with various values of  $U$  for both i.i.d. and non-i.i.d. data distributions. The figures reveal that the impact of  $U$  differs from that of  $K$ . In both cases, the time-to-accuracy performance of STSyn consistently increases as  $U$  decreases, indicating that the reduction in per-round wall-clock time through  $U$  is a more influential factor in terms of time efficiency. These findings highlight that the choice of  $U$  plays a crucial role in optimizing time efficiency in STSyn, regardless of the underlying data distribution.

Fig. 10 (a) and (b) illustrate the communication-to-accuracy performances of STSyn. In the case of i.i.d. data, the communication efficiency increases as the value of  $U$  becomes larger. This observation suggests that, for homogeneous data distribution, having more local updates per round leads to

higher communication efficiency. However, in the case of non-i.i.d. data, increasing  $U$  does not necessarily result in higher system communication efficiency. This is due to the client-drift effect, wherein more local updates cause the workers to diverge towards their own local optima.

Therefore, for homogeneous data distribution, selecting a moderate value of  $U$  provides a preferable choice as it allows STSyn to strike a balance between time and communication efficiency. On the other hand, for heterogeneous data distribution, opting for a smaller value of  $K$  is advisable in order to achieve better time efficiency and mitigate the client-drift effect.

In summary, the optimal selection of  $U$  and  $K$  depends heavily on the specific characteristics of the data distribution. For homogeneous data, a moderate choice of  $K$  and  $U$  strikes the right balance. In contrast, for heterogeneous data, a smaller  $K$  and a moderate  $U$  are preferred to enhance time efficiency and minimize the client-drift effect.

**Additional Experiments under Weaker Assumptions.** We proceed to consider a more realistic scenario where computing times across local iterations are correlated (e.g., a worker experiencing delays is likely to continue doing so). Specifically, we evaluate the case where the means of the exponential distribution vary across workers, and computing times remain consistent across local iterations for each worker within the same round. This consistency could result in greater variance in the number of local updates for STSyn. To manage this, we cap the maximum number of local updates at  $1.5U$ , ensuring that the faster workers stop computing after completing  $1.5U$  local updates.

Fig. 11 illustrates the time and communication efficiencies of STSyn compared to other advanced algorithms under non-i.i.d. data distribution. Compared with Fig. 5 (b) and Fig. 6 (b), it is obvious that STSyn demonstrates greater time efficiency while maintaining similar communication efficiency to other algorithms, except for FLANP. This increased time efficiency is attributed to the weaker assumptions, which significantly amplify the variance in the number of local updates. Consequently, STSyn effectively mitigates the impact

of stragglers, albeit with a slight trade-off in communication efficiency. These findings further underscore the adaptability and superiority of STSyn.

## VI. CONCLUSION

Building on a PS-based distributed learning framework for both i.i.d. and non-i.i.d. data distribution, we proposed STSyn, a novel local-SGD scheme with heterogeneous synchronization, which is both straggler-tolerant and communication-efficient. We derived the average wall-clock time, average number of local updates and average number of uploading workers for the proposed STSyn to substantiate the effectiveness of the additional local updates. The convergence of STSyn under homogeneous and heterogeneous data distributions was also rigorously established. Extensive experimental results corroborated the superior time and communication efficiency of the proposed STSyn over the existing alternatives. The impacts of the hyper-parameters were also explored.

Our work leaves open a number of future research directions. First, it would be interesting to explore the potential advantages of heterogeneous synchronization for asynchronous implementations, where any worker can perform a number of local updates and send the result back to the PS. Second, generalization and analysis of the proposed STSyn scheme to the scenario where the wireless medium is considered, are worthy of further investigation. Lastly, STSyn is a highly integrated algorithm that could be combined with many schemes such as adaptive communication, adaptive worker selection, client-drift mitigation and other techniques to improve the overall performance of the distributed system.

## REFERENCES

- [1] L. Bottou, "Large-scale machine learning with stochastic gradient descent," in *Proc. COMPSTAT*, 2010, pp. 177–186.
- [2] X. Lian, C. Zhang, H. Zhang, C. J. Hsieh, W. Zhang, and J. Liu, "Can decentralized algorithms outperform centralized algorithms? A case study for decentralized parallel stochastic gradient descent," in *Proc. Neural Inf. Process. Syst.*, vol. 30, 2017, pp. 5336–5346.
- [3] Z. Jiang, A. Balu, C. Hegde, and S. Sarkar, "Collaborative deep learning in fixed topology networks," in *Proc. Neural Inf. Process. Syst.*, vol. 30, 2017, pp. 5906–5916.
- [4] M. Li, D. G. Andersen, J. W. Park, A. J. Smola, and B. Y. Su, "Scaling distributed machine learning with the parameter server," in *Proc. USENIX Symp. Oper. Syst. Des. Implement.*, 2014, pp. 583–598.
- [5] J. Dean, G. S. Corrado, R. Monga, C. Kai, and A. Y. Ng, "Large scale distributed deep networks," in *Proc. Neural Inf. Process. Syst.*, vol. 25, 2012, pp. 1223–1231.
- [6] M. Zinkevich, M. Weimer, A. J. Smola, and L. Li, "Parallelized stochastic gradient descent," *Proc. Neural Inf. Process. Syst.*, vol. 23, pp. 2595–2603, 2010.
- [7] F. Sattler, S. Wiedemann, K.-R. Müller, and W. Samek, "Robust and communication-efficient federated learning from Non-iid data," *IEEE Trans. Neural Netw. and Learn. Syst.*, vol. 31, no. 9, pp. 3400–3413, 2019.
- [8] H. Cui, J. Cipar, Q. Ho, J. Kim, S. Lee, A. Kumar, J. Wei, D. Wei, G. R. Ganger, and P. B. Gibbons, "Exploiting bounded staleness to speed up big data analytics," in *Proc. USENIX Annu. Tech. Conf.*, 2014, pp. 37–48.
- [9] S. Gupta, W. Zhang, and F. Wang, "Model accuracy and runtime tradeoff in distributed deep learning: A systematic study," in *Proc. IEEE Int. Conf. Data Mining*, 2016, pp. 171–180.
- [10] P. Goyal, P. Dollár, R. Girshick, P. Noordhuis, L. Wesolowski, A. Kyrola, A. Tulloch, Y. Jia, and K. He, "Accurate, large minibatch SGD: Training imagenet in 1 hour," *arXiv preprint arXiv:1706.02677*, 2017.
- [11] P. Dube, T. Suk, and C. Wang, "AI gauge: Runtime estimation for deep learning in the cloud," in *Proc. IEEE Int. Symp. Comput. Architecture and High Perform. Comput.*, 2019, pp. 160–167.
- [12] V. Smith, C.-K. Chiang, M. Sanjabi, and A. Talwalkar, "Federated multi-task learning," in *Proc. Neural Inf. Process. Syst.*, vol. 30, 2017, pp. 4427–4437.
- [13] Y. Arjevani and O. Shamir, "Communication complexity of distributed convex learning and optimization," in *Proc. Neural Inf. Process. Syst.*, vol. 28, 2015, pp. 1756–1764.
- [14] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. Artif. Intell. and Statist.*, 2017, pp. 1273–1282.
- [15] F. Zhou and G. Cong, "On the convergence properties of a K-step averaging stochastic gradient descent algorithm for nonconvex optimization," in *Proc. Int. Joint Conf. Artif. Intell.*, 2018, pp. 3219–3227.
- [16] S. U. Stich, "Local SGD converges fast and communicates little," in *ICLR 2019-International Conference on Learning Representations*, no. CONF, 2019.
- [17] J. Wang and G. Joshi, "Cooperative SGD: A unified framework for the design and analysis of communication-efficient SGD algorithms," in *Proc. ICML Workshop Coding Theory Mach. Learn.*, 2019.
- [18] B. Woodworth, J. Wang, B. McMahan, and N. Srebro, "Graph oracle models, lower bounds, and gaps for parallel stochastic optimization," *Proc. Neural Inf. Process. Syst.*, vol. 31, pp. 8505–8515, 2018.
- [19] X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang, "On the convergence of FedAvg on Non-IID data," in *Proc. Int. Conf. Learn. Representations*, 2019.
- [20] A. Khaled, K. Mishchenko, and P. Richtárik, "First analysis of local GD on heterogeneous data," *arXiv preprint arXiv:1909.04715*, 2019.
- [21] H. Yu, S. Yang, and S. Zhu, "Parallel restarted SGD with faster convergence and less communication: Demystifying why model averaging works for deep learning," in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, no. 01, 2019, pp. 5693–5700.
- [22] S. Wang, T. Tuor, T. Salonidis, K. K. Leung, C. Makaya, T. He, and K. Chan, "Adaptive federated learning in resource constrained edge computing systems," *IEEE J. Sel. Areas Commun.*, vol. 37, no. 6, pp. 1205–1221, 2019.
- [23] F. Haddadpour and M. Mahdavi, "On the convergence of local descent methods in federated learning," *arXiv preprint arXiv:1910.14425*, 2019.
- [24] L. Tao, S. U. Stich, and M. Jaggi, "Don't use large mini-batches, use local SGD," in *Proc. Int. Conf. Learn. Representations*, 2019.
- [25] F. Haddadpour, M. M. Kamani, M. Mahdavi, and V. R. Cadambe, "Local SGD with periodic averaging: Tighter analysis and adaptive synchronization," in *Proc. Neural Inf. Process. Syst.*, vol. 32, 2019, pp. 11 082–11 094.
- [26] J. Wang and G. Joshi, "Adaptive communication strategies to achieve the best error-runtime trade-off in local-update SGD," in *Proc. Mach. Learn. and Syst.*, vol. 1, 2019, pp. 212–229.
- [27] A. Reiszadeh, I. Tziotis, H. Hassani, A. Mokhtari, and R. Pedarsani, "Straggler-resilient federated learning: Leveraging the interplay between statistical accuracy and system heterogeneity," *IEEE Journal on Selected Areas in Information Theory*, vol. 3, no. 2, pp. 197–205, 2022.
- [28] S. K. Hanna, R. Bitar, P. Parag, V. Dasari, and S. E. Rouayheb, "Adaptive stochastic gradient descent for fast and communication-efficient distributed learning," *arXiv preprint arXiv:2208.03134*, 2022.
- [29] S. Shen, Y. Cheng, J. Liu, and L. Xu, "STL-SGD: Speeding up local SGD with stagewise communication period," in *Proc. AAAI Conf. Artif. Intell.*, vol. 35, no. 11, 2021, pp. 9576–9584.
- [30] J. Dean and L. A. Barroso, "The tail at scale," *Commun. ACM*, vol. 56, no. 2, pp. 74–80, 2013.
- [31] A. Tyurin and P. Richtarik, "Optimal time complexities of parallel stochastic optimization methods under a fixed computation model," in *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [32] R. Tandon, L. Qi, A. G. Dimakis, and N. Karampatziakis, "Gradient coding: Avoiding stragglers in distributed learning," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 3368–3376.
- [33] S. Li, M. A. Maddah-Ali, Q. Yu, and A. S. Avestimehr, "A fundamental tradeoff between computation and communication in distributed computing," *IEEE Trans. Inf. Theory*, vol. 64, no. 1, pp. 109–128, 2017.
- [34] H. Wang, Z. Charles, and D. Papailiopoulos, "Erasurehead: Distributed gradient descent without delays using approximate gradient coding," *arXiv preprint arXiv:1901.09671*, 2019.
- [35] J. Zhang and O. Simeone, "LAGC: Lazily aggregated gradient coding for straggler-tolerant and communication-efficient distributed learning," *IEEE Trans. Neural Netw. and Learn. Syst.*, vol. 32, no. 3, pp. 962–974, 2020.

- [36] E. Ozfatura, D. Gündüz, and S. Ulukus, “Gradient coding with clustering and multi-message communication,” in *Proc. IEEE Data Sci. Workshop*, 2019, pp. 42–46.
- [37] Q. Wang, Y. Cui, C. Li, J. Zou, and H. Xiong, “Optimization-based block coordinate gradient coding for mitigating partial stragglers in distributed learning,” *IEEE Transactions on Signal Processing*, vol. 71, pp. 1023–1038, 2023.
- [38] A. Harlap, H. Cui, W. Dai, J. Wei, G. R. Ganger, P. B. Gibbons, G. A. Gibson, and E. P. Xing, “Addressing the straggler problem for iterative convergent parallel ML,” in *Proc. ACM Symp. Cloud Comput.*, 2016, pp. 98–111.
- [39] K. Son and W. Choi, “Distributed matrix multiplication based on frame quantization for straggler mitigation,” *IEEE Transactions on Signal Processing*, vol. 70, pp. 3058–3073, 2022.
- [40] Y. Venkatesha, Y. Kim, L. Tassiulas, and P. Panda, “Federated learning with spiking neural networks,” *IEEE Transactions on Signal Processing*, vol. 69, pp. 6183–6194, 2021.
- [41] S. Dutta, J. Wang, and G. Joshi, “Slow and stale gradients can win the race: Error-runtime trade-offs in distributed SGD,” in *Proc. Int. Conf. Artif. Intell. and Statist.*, 2018, pp. 803–812.
- [42] E. Rizk, S. Vlaski, and A. H. Sayed, “Dynamic federated learning,” in *IEEE International Workshop on Signal Processing Advances in Wireless Communications*, 2020, pp. 1–5.
- [43] J. Wang, Q. Liu, H. Liang, G. Joshi, and H. V. Poor, “Tackling the objective inconsistency problem in heterogeneous federated optimization,” in *Proc. NeurIPS*, vol. 33, 2020, pp. 7611–7623.
- [44] Y. Ruan, X. Zhang, S.-C. Liang, and C. Joe-Wong, “Towards flexible device participation in federated learning,” in *Proc. Int. Conf. Artif. Intell. and Statist.*, 2021, pp. 3403–3411.
- [45] N. Ferdinand, H. Al-Lawati, S. Draper, and M. Nokleby, “Anytime minibatch: Exploiting stragglers in online distributed optimization,” in *International Conference on Learning Representations*, 2018.
- [46] A. Agarwal and J. C. Duchi, “Distributed delayed stochastic optimization,” in *Proc. Neural Inf. Process. Syst.*, vol. 24, 2011, pp. 873–881.
- [47] B. Recht, C. Re, S. Wright, and F. Niu, “Hogwild!: A lock-free approach to parallelizing stochastic gradient descent,” *Advances in neural information processing systems*, vol. 24, 2011.
- [48] X. Pan, M. Lam, S. Tu, D. Papailiopoulos, C. Zhang, M. I. Jordan, K. Ramchandran, and C. Ré, “Cyclades: Conflict-free asynchronous machine learning,” *Advances in Neural Information Processing Systems*, vol. 29, 2016.
- [49] X. Lian, Y. Huang, Y. Li, and J. Liu, “Asynchronous parallel stochastic gradient for nonconvex optimization,” in *Proc. Neural Inf. Process. Syst.*, vol. 28, 2015, pp. 2737–2745.
- [50] S. Sra, A. W. Yu, M. Li, and A. J. Smola, “Adadelay: Delay adaptive distributed stochastic optimization,” in *Proc. Artif. Intell. and Statist.*, 2016, pp. 957–965.
- [51] S. Zheng, Q. Meng, T. Wang, W. Chen, N. Yu, Z. M. Ma, and T. Y. Liu, “Asynchronous stochastic gradient descent with delay compensation,” in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 4120–4129.
- [52] W. Wu, L. He, W. Lin, R. Mao, C. Maple, and S. Jarvis, “Safa: A semi-asynchronous protocol for fast federated learning with low overhead,” *IEEE Transactions on Computers*, vol. 70, no. 5, pp. 655–668, 2021.
- [53] J. Nguyen, K. Malik, H. Zhan, A. Yousefpour, M. Rabbat, M. Malek, and D. Huba, “Federated learning with buffered asynchronous aggregation,” in *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, ser. Proceedings of Machine Learning Research, G. Camps-Valls, F. J. R. Ruiz, and I. Valera, Eds., vol. 151, 28–30 Mar 2022, pp. 3581–3607.
- [54] A. Tyurin, K. Gruntkowska, and P. Richtárik, “Freya page: First optimal time complexity for large-scale nonconvex finite-sum optimization with heterogeneous asynchronous computations,” *arXiv preprint arXiv:2405.15545*, 2024.
- [55] N. Ivkin, D. Rothchild, E. Ullah, I. Stoica, R. Arora *et al.*, “Communication-efficient distributed SGD with sketching,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [56] D. Rothchild, A. Panda, E. Ullah, N. Ivkin, I. Stoica, V. Braverman, J. Gonzalez, and R. Arora, “Fetchsgd: Communication-efficient federated learning with sketching,” in *International Conference on Machine Learning*, 2020, pp. 8253–8265.
- [57] S. U. Stich, J.-B. Cordonnier, and M. Jaggi, “Sparsified SGD with memory,” *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [58] S. Nadarajah, “Explicit expressions for moments of order statistics,” *Statistics & Probability Letters*, vol. 78, no. 2, pp. 196–205, 2008.
- [59] Y. Zhao *et al.*, “Federated learning with non-IID data,” *arXiv preprint arXiv:1806.00582*, 2018.
- [60] S. J. Reddi, Z. Charles, M. Zaheer, Z. Garrett, K. Rush, J. Konečný, S. Kumar, and H. B. McMahan, “Adaptive federated optimization,” in *Proc. ICLR*, 2020.
- [61] H. Yang, M. Fang, and J. Liu, “Achieving linear speedup with partial worker participation in non-IID federated learning,” in *Proc. ICLR*, 2020.
- [62] S. P. Karimireddy, S. Kale, M. Mohri, S. Reddi, S. Stich, and A. T. Suresh, “Scaffold: Stochastic controlled averaging for federated learning,” in *Proc. Int. Conf. Mach. Learn.* PMLR, 2020, pp. 5132–5143.

## APPENDIX A PROOF OF THEOREM 1

*Proof.* According to (16) in Assumption 2, the one-step difference of the objective function can be written as

$$\begin{aligned}
& F(\mathbf{w}^{j+1}) - F(\mathbf{w}^j) \\
& \stackrel{(a1)}{\leq} \langle \nabla F(\mathbf{w}^j), \mathbf{w}^{j+1} - \mathbf{w}^j \rangle + \frac{L}{2} \|\mathbf{w}^{j+1} - \mathbf{w}^j\|^2 \\
& \stackrel{(a2)}{=} \langle \nabla F(\mathbf{w}^j), \mathbb{E}[g^j + \alpha \bar{U} \nabla F(\mathbf{w}^j) - \alpha \bar{U} \nabla F(\mathbf{w}^j)] \rangle + \frac{L}{2} T_2 \\
& \stackrel{(a3)}{=} -\alpha \bar{U} \|\nabla F(\mathbf{w}^j)\|^2 + \langle \nabla F(\mathbf{w}^j), \mathbb{E}[g^j + \alpha \bar{U} \nabla F(\mathbf{w}^j)] \rangle \\
& \quad + \frac{L}{2} T_2, \\
& \stackrel{(a4)}{=} -\alpha \bar{U} \|\nabla F(\mathbf{w}^j)\|^2 + \underbrace{\langle \nabla F(\mathbf{w}^j), \mathbb{E}[g^j + \alpha \bar{U} \nabla F(\mathbf{w}^j)] \rangle}_{T_1} \\
& \quad + \frac{L}{2} T_2, \tag{30}
\end{aligned}$$

where (a2) holds because of the definitions  $T_2 = \mathbb{E}[\|g^j\|^2]$  and  $g^j = \frac{1}{S^j} \sum_{m \in S^j} g_m^j$  where  $g_m^j = -\alpha \sum_{u=0}^{U_m^j-1} \frac{1}{B} \sum_{b=1}^B \nabla F_m(\mathbf{w}_m^{j,u}; z_{m,b}^{j,u})$ ; (a3) is directly from (a2) and (a4) is due to the definition that  $\bar{g}^j = \frac{1}{M} \sum_{m \in \mathcal{M}} g_m^j$  since the distribution of the computing capabilities of workers is homogeneous.

Next, we will bound the expectation of the first and second terms  $T_1$  and  $T_2$  in (30), respectively. For the first term, we have

$$\begin{aligned}
& \langle \nabla F(\mathbf{w}^j), \mathbb{E}[(\bar{g}^j + \alpha \bar{U} \nabla F(\mathbf{w}^j))] \rangle \\
& \stackrel{(b1)}{=} \left\langle \nabla F(\mathbf{w}^j), \mathbb{E} \left[ -\frac{1}{BM} \sum_{m=1}^M \sum_{u=0}^{U_m^j-1} \sum_{b=1}^B \alpha \nabla F_m(\mathbf{w}_m^{j,u}; z_{m,b}^{j,u}) \right. \right. \\
& \quad \left. \left. + \alpha \bar{U} \nabla F(\mathbf{w}^j) \right] \right\rangle \\
& \stackrel{(b2)}{=} \left\langle \nabla F(\mathbf{w}^j), \mathbb{E} \left[ -\frac{1}{M} \sum_{m=1}^M \sum_{u=0}^{U_m^j-1} \alpha \nabla F_m(\mathbf{w}_m^{j,u}) \right. \right. \\
& \quad \left. \left. + \alpha \bar{U} \frac{1}{M} \sum_{m=1}^M \nabla F(\mathbf{w}^j) \right] \right\rangle \\
& \stackrel{(b3)}{=} \left\langle \sqrt{\alpha \bar{U}} \nabla F(\mathbf{w}^j), -\frac{\sqrt{\alpha}}{\sqrt{\bar{U}}} \mathbb{E} \left[ \sum_{m=1}^M \frac{1}{M} \sum_{u=0}^{U_m^j-1} \Delta_m^{j,u} \right] \right\rangle \\
& \stackrel{(b4)}{\leq} \frac{\alpha \bar{U}}{2} \|\nabla F(\mathbf{w}^j)\|^2 + \frac{\alpha}{2\bar{U}} \mathbb{E} \left[ \left\| \frac{1}{M} \sum_{m=1}^M \sum_{u=0}^{U_m^j-1} \Delta_m^{j,u} \right\|^2 \right]
\end{aligned}$$

$$\begin{aligned}
& -t_{11} \left\| \sum_{m=1}^M \sum_{u=0}^{U_m^j-1} \nabla F_m(\mathbf{w}_m^{j,u}) + (\bar{U} - U_m^j) M \nabla F(\mathbf{w}^j) \right\|^2 \\
& \stackrel{(b5)}{\leq} \frac{\alpha \bar{U}}{2} \|\nabla F(\mathbf{w}^j)\|^2 + \frac{\alpha}{2\bar{U}M} \sum_{m=1}^M U_m^j \sum_{u=0}^{U_m^j-1} \mathbb{E} \left[ \|\Delta_m^{j,u}\|^2 \right] \\
& - t_{11} \mathbb{E} \left\| \sum_{m=1}^M \sum_{u=0}^{U_m^j-1} \nabla F_m(\mathbf{w}_m^{j,u}) \right\|^2 \\
& - t_{11} (\bar{U} - U_m^j)^2 M^2 \|\nabla F(\mathbf{w}^j)\|^2 \\
& - 2t_{11} \mathbb{E} \left[ \left\langle \sum_{m=1}^M \sum_{u=0}^{U_m^j-1} \nabla F_m(\mathbf{w}_m^{j,u}), (\bar{U} - U_m^j) M \nabla F(\mathbf{w}^j) \right\rangle \right] \\
& \stackrel{(b6)}{\leq} \frac{\alpha \bar{U}}{2} \|\nabla F(\mathbf{w}^j)\|^2 + \frac{\alpha L^2}{2\bar{U}M} \sum_{m=1}^M U_m^j \sum_{u=0}^{U_m^j-1} \mathbb{E} \left[ \|\mathbf{w}_m^{j,u} - \mathbf{w}^j\|^2 \right] \\
& - \frac{\alpha}{2\bar{U}M^2} \mathbb{E} \left\| \sum_{m=1}^M \sum_{u=0}^{U_m^j-1} \nabla F_m(\mathbf{w}_m^{j,u}) \right\|^2 \\
& \stackrel{(b7)}{\leq} \left( \frac{\alpha \bar{U}}{2} + \frac{15\alpha^3 L^2 U_{max}^4}{\bar{U}} \right) \|\nabla F(\mathbf{w}^j)\|^2 \\
& + \frac{5\alpha^3 U_{max}^3 L^2}{2\bar{U}} t_{12} - t_{11} \mathbb{E} \left\| \sum_{m=1}^M \sum_{u=0}^{U_m^j-1} \nabla F_m(\mathbf{w}_m^{j,u}) \right\|^2 \\
& \stackrel{(b8)}{=} \left( \frac{\alpha \bar{U}}{2} + \frac{15\alpha^3 L^2 U_{max}^4}{\bar{U}} \right) \|\nabla F(\mathbf{w}^j)\|^2 \\
& + \frac{5\alpha^3 U_{max}^3 L^2}{2\bar{U}} t_{12} - t_{11} \mathbb{E} \left\| \sum_{m \in \mathcal{S}^j} \sum_{u=0}^{U_m^j-1} \nabla F_m(\mathbf{w}_m^{j,u}) \right\|^2 \quad (31)
\end{aligned}$$

where we have defined the terms  $t_{11} = \frac{\alpha}{2\bar{U}M^2}$  and  $t_{12} = (\sigma_L^2 + 6U_{max}\sigma_G^2)$ ; (b1) is from the definition of  $\bar{g}^j$ ; (b2) is due to the unbiasedness assumption; (b3) is from direct computation and we define  $\Delta_m^{j,u} = \nabla F_m(\mathbf{w}_m^{j,u}) - \nabla F(\mathbf{w}^j)$ ; (b4) uses the fact that  $\langle \mathbf{x}, \mathbf{y} \rangle = \frac{1}{2}(\|\mathbf{x}\|^2 + \|\mathbf{y}\|^2 - \|\mathbf{x} - \mathbf{y}\|^2)$ , where  $\mathbf{x} = \sqrt{\alpha\bar{U}}\nabla F(\mathbf{w}^j)$ , and  $\mathbf{y} = -\frac{\sqrt{\alpha}}{\sqrt{\bar{U}}} \sum_{m=1}^M \frac{1}{M} \sum_{u=0}^{U_m^j-1} (\nabla F_m(\mathbf{w}_m^{j,u}) - \nabla F(\mathbf{w}^j))$ ; (b5) comes from Cauchy-Schwartz inequality and unrolling using the perfect square formula; (b6) follows from the  $L$ -smoothness assumption, the fact that  $\frac{\alpha}{2\bar{U}M^2} (\bar{U} - U_m^j)^2 M^2 \|\nabla F(\mathbf{w}^j)\|^2 \geq 0$ , and  $\mathbb{E} \left[ \left\langle \sum_{m=1}^M \sum_{u=0}^{U_m^j-1} \nabla F_m(\mathbf{w}_m^{j,u}), (\bar{U} - U_m^j) M \nabla F(\mathbf{w}^j) \right\rangle \right] = 0$  since  $\mathbb{E}[U_m^j] = \bar{U}$ ; (b7) is from [60, Lemma 3], which proves that

$$\begin{aligned}
& \mathbb{E} \left[ \|\mathbf{w}_m^{j,u} - \mathbf{w}^j\|^2 \right] \leq 5U_{max}\eta^2(\sigma_L^2 + 6U_{max}\sigma_G^2) \\
& + 30U_{max}^2\eta^2 \|\nabla F(\mathbf{w}^j)\|^2, \quad (32)
\end{aligned}$$

under the condition that  $\alpha \leq \frac{1}{8L\bar{U}U_{max}}$ , where  $\sigma_L$  and  $\sigma_G$  are two constants defined in Assumption 1. Considering that  $U_m^j$  varies among workers and  $U_{max}$  is the maximal value, we can naturally arrive at the conclusion and (b8) uses the fact that only the workers in subset  $\mathcal{S}^j$  have completed effective local updates.

For the second term, we have

$$\begin{aligned}
& \mathbb{E}[\|g^j\|^2] \stackrel{(c1)}{=} \mathbb{E} \left[ \left\| \frac{1}{S^j} \sum_{m \in \mathcal{S}^j} g_m^j \right\|^2 \right] \\
& \stackrel{(c2)}{=} \mathbb{E} \left[ \left\| \frac{1}{S^j} \sum_{m \in \mathcal{S}^j} \alpha \sum_{u=0}^{U_m^j-1} \frac{1}{B} \sum_{b=1}^B \nabla F_m(\mathbf{w}_m^{j,u}; z_{m,b}^{j,u}) \right\|^2 \right] \\
& \stackrel{(c3)}{\leq} t_2 \mathbb{E} \left[ \left\| \sum_{m \in \mathcal{S}^j} \sum_{u=0}^{U_m^j-1} \frac{1}{B} \sum_{b=1}^B \nabla F_m(\mathbf{w}_m^{j,u}; z_{m,b}^{j,u}) \right\|^2 \right] \\
& \stackrel{(c4)}{=} t_2 \mathbb{E} \left[ \left\| \sum_{m \in \mathcal{S}^j} \sum_{u=0}^{U_m^j-1} \left( \frac{1}{B} \sum_{b=1}^B \nabla F_m(\mathbf{w}_m^{j,u}; z_{m,b}^{j,u}) - \nabla F_m(\mathbf{w}_m^{j,u}) \right) \right\|^2 \right] \\
& + t_2 \mathbb{E} \left[ \left\| \sum_{m \in \mathcal{S}^j} \sum_{u=0}^{U_m^j-1} \nabla F_m(\mathbf{w}_m^{j,u}) \right\|^2 \right] \\
& \stackrel{(c5)}{\leq} \frac{t_2}{B^2} \sum_{m \in \mathcal{S}^j} \sum_{u=0}^{U_m^j-1} \sum_{b=1}^B \mathbb{E} \left[ \left\| \nabla F_m(\mathbf{w}_m^{j,u}; z_{m,b}^{j,u}) - \nabla F_m(\mathbf{w}_m^{j,u}) \right\|^2 \right] \\
& + \frac{\alpha^2}{(S^j)^2} \mathbb{E} \left[ \left\| \sum_{m \in \mathcal{S}^j} \sum_{u=0}^{U_m^j-1} \nabla F_m(\mathbf{w}_m^{j,u}) \right\|^2 \right] \\
& \stackrel{(c6)}{\leq} \frac{\alpha^2 U_{max}}{S^j B} \sigma_L^2 + \frac{\alpha^2}{(S^j)^2} \mathbb{E} \left[ \left\| \sum_{m \in \mathcal{S}^j} \sum_{u=0}^{U_m^j-1} \nabla F_m(\mathbf{w}_m^{j,u}) \right\|^2 \right], \quad (33)
\end{aligned}$$

where (c1) is due to the definition of  $g^j$ ; (c2) is due to the definition of  $g_m^j$ ; (c3) comes from direct computation and the definition  $t_2 = \frac{\alpha^2}{(S^j)^2}$ ; (c4) uses the fact that  $\mathbb{E}[\|x\|^2] = \mathbb{E}[\|x - \mathbb{E}[x]\|^2] + \|\mathbb{E}[x]\|^2$  and the unbiasedness assumption, where  $x = \sum_{m \in \mathcal{S}^j} \sum_{u=0}^{U_m^j-1} \frac{1}{B} \sum_{b=1}^B \nabla F_m(\mathbf{w}_m^{j,u}; z_{m,b}^{j,u})$  and  $\mathbb{E}[x] = \sum_{m \in \mathcal{S}^j} \sum_{u=0}^{U_m^j-1} \nabla F_m(\mathbf{w}_m^{j,u})$ ; (c5) is due to the fact that  $\mathbb{E}[\|x_1 + \dots + x_n\|^2] = \mathbb{E}[\|x_1\|^2 + \dots + \|x_n\|^2]$  if  $x_i$ 's are independent with zero mean and (c6) uses Assumption 1.

Combining (33) and (31) into (30), we readily have

$$\begin{aligned}
& F(\mathbf{w}^{j+1}) - F(\mathbf{w}^j) \\
& \stackrel{(d1)}{\leq} - \left( \frac{\alpha \bar{U}}{2} - \frac{15\alpha^3 L^2 U_{max}^4}{\bar{U}} \right) \|\nabla F(\mathbf{w}^j)\|^2 + \frac{\alpha^2 U_{max} L}{2S^j B} \sigma_L^2 \\
& + \frac{5\alpha^3 U_{max}^3 L^2}{2\bar{U}} (\sigma_L^2 + 6U_{max}\sigma_G^2) \\
& + \left( \frac{\alpha^2 L}{2(S^j)^2} - \frac{\alpha}{2\bar{U}M^2} \right) \mathbb{E} \left[ \left\| \sum_{m \in \mathcal{S}^j} \sum_{u=0}^{U_m^j-1} \nabla F_m(\mathbf{w}_m^{j,u}) \right\|^2 \right] \\
& \stackrel{(d2)}{\leq} -\alpha \bar{U} \left( \frac{1}{2} - \frac{15\alpha^2 L^2 U_{max}^4}{\bar{U}^2} \right) \|\nabla F(\mathbf{w}^j)\|^2 + \frac{\alpha^2 U_{max} L}{2S^j B} \sigma_L^2 \\
& + \frac{5\alpha^3 U_{max}^3 L^2}{2\bar{U}} (\sigma_L^2 + 6U_{max}\sigma_G^2) \\
& \stackrel{(d3)}{\leq} -\alpha \bar{U} c \|\nabla F(\mathbf{w}^j)\|^2 + \frac{\alpha^2 U_{max} L}{2KB} \sigma_L^2
\end{aligned}$$

$$+ \frac{5\alpha^3 U_{max}^3 L^2}{2\bar{U}} (\sigma_L^2 + 6U_{max}\sigma_G^2) \quad (34)$$

where (d1) is from direct computation; (d2) follows from  $\frac{\alpha^2 L}{2(S^j)^2} - \frac{\alpha}{2UM^2} \leq 0$  if  $\alpha \leq \frac{K^2}{LU^2M^2}$  and (d3) holds because there exists a positive constant  $c$  satisfying  $\frac{1}{2} - \frac{15\alpha^2 L^2 U_{max}^4}{U^2} > c$  if  $\alpha \leq \frac{\bar{U}}{\sqrt{30LU_{max}^2}}$  and the fact that  $S^j \geq K$ .

Summing (34) from 0 to  $J - 1$  and rearranging the terms we have:

$$\begin{aligned} \frac{1}{J} \mathbb{E} \left[ \sum_{j=0}^{J-1} \|\nabla F(\mathbf{w}^j)\|^2 \right] &\leq \frac{F(\mathbf{w}^J) - F^*}{c\alpha\bar{U}J} \\ &+ \frac{1}{c} \left[ \frac{\alpha U_{max} L}{2KB\bar{U}} \sigma_L^2 + \frac{5\alpha^2 U_{max}^3 L^2}{2\bar{U}^2} (\sigma_L^2 + 6U_{max}\sigma_G^2) \right]. \quad (35) \end{aligned}$$

The proof is then complete.