
Anarchic Federated Learning

Haibo Yang¹ Xin Zhang² Prashant Khanduri^{1,3} Jia Liu¹

Abstract

Present-day federated learning (FL) systems deployed over edge networks consists of a large number of workers with high degrees of heterogeneity in data and/or computing capabilities, which call for flexible worker participation in terms of timing, effort, data heterogeneity, etc. To satisfy the need for flexible worker participation, we consider a new FL paradigm called “Anarchic Federated Learning” (AFL) in this paper. In stark contrast to conventional FL models, each worker in AFL has the freedom to choose i) when to participate in FL, and ii) the number of local steps to perform in each round based on its current situation (e.g., battery level, communication channels, privacy concerns). However, such chaotic worker behaviors in AFL impose many new open questions in algorithm design. In particular, it remains unclear whether one could develop convergent AFL training algorithms, and if yes, under what conditions and how fast the achievable convergence speed is. Toward this end, we propose two Anarchic Federated Averaging (AFA) algorithms with two-sided learning rates for both cross-device and cross-silo settings, which are named AFA-CD and AFA-CS, respectively. Somewhat surprisingly, we show that, under mild anarchic assumptions, both AFL algorithms achieve the best known convergence rate as the state-of-the-art algorithms for conventional FL. Moreover, they retain the highly desirable *linear speedup effect* with respect of both the number of workers and local steps in the new AFL paradigm. We validate the proposed algorithms with extensive experiments on real-world datasets.

1. Introduction

Federated Learning (FL) has recently emerged as an important distributed learning framework that leverages numerous workers to collaboratively learn a joint model (Li et al., 2019a; Yang et al., 2019; Kairouz et al., 2019). Since the inception, FL systems have become increasingly powerful and are able to handle various heterogeneity in data, network environments, worker computing capabilities, etc. Furthermore, most of the prevailing FL algorithms (e.g., FedAvg (McMahan et al., 2016) and its variants (Li et al., 2018; Zhang et al., 2020a; Karimireddy et al., 2020b;a; Acar et al., 2021)) enjoy a desirable “*linear speedup effect*,” i.e., the convergence time to a first-order stationary point decreases linearly as the number of workers and local steps increases (Stich, 2018; Yu et al., 2019; Wang & Joshi, 2018; Khaled et al., 2019; Karimireddy et al., 2020b; Yang et al., 2021; Qu et al., 2020).

However, to achieve these salient features, most of the existing FL algorithms have adopted a *server-centric* approach, i.e., the worker behaviors are tightly “dictated” by the server. Such dictation is typically manifested in three aspects: i) determine either all or a subset of workers to participate in each round of FL update; ii) fully control the timing for synchronization and whether to accept/reject information sent from the workers; iii) precisely specify the algorithmic operations (e.g., the number of local steps performed at each worker before communicating with the server). Despite achieving strong performance guarantees, these server-centric FL algorithms often implicitly rely on the following strong assumptions: (1) each worker is available for training upon the server’s request and throughout a complete round; (2) all participating workers are willing to execute the same number of local updates and communicate with the server in a synchronous manner following a common clock. Unfortunately, in edge networks where many FL systems are deployed, these assumptions are restrictive or even problematic. First, many requested edge devices on the worker side may not be available in each round because of, e.g., communication errors or battery outages. Second, the use of synchronous communication and an identical number of local updates across all workers ignores the fact that worker devices in edge-based FL systems are heterogeneous in computation and communication capabilities. As a result, stragglers (i.e., slow workers) could significantly

¹Department of Electrical and Computer Engineering, The Ohio State University, Columbus, OH 43210, USA; ²Department of Statistics, Iowa State University, Ames, IA 50011, USA; ³Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455, USA. Correspondence to: Haibo Yang <yang.5952@osu.edu>, Jia Liu <liu@ece.osu.edu>.

slow down the training process. To mitigate the straggler effect, various robust FL algorithms have been developed. For example, the server in FedAvg (McMahan et al., 2016) can simply ignore and drop the information from the stragglers to speedup learning. However, this may lead to other problems such as wasted computation/energy (Wang et al., 2019), slower convergence (Li et al., 2018), or biased/unfair uses of worker data (Kairouz et al., 2019). Moreover, the synchronous nature of the server-centric approaches implies many networking problems (e.g., interference between workers, periodic traffic spikes, high complexity in maintaining a network-wide common clock).

The above limitations of the current server-centric FL approaches motivate us to propose a new paradigm in FL, which we call **Anarchic Federated Learning** (AFL). In stark contrast to server-centric FL, workers in AFL are completely *free* of the “dictation” from the server. Specifically, each worker has complete freedom to choose when and how long to participate in FL without following any control signals from the server. As a result, the information fed back from workers is inherently asynchronous. Also, each worker can independently determine the number of local update steps to perform in each round based on its current local situation (e.g., battery level, communication channels, privacy concerns). In other words, the amount of local computation at each worker is time-varying, device-dependent, and fully controlled by the worker itself. Clearly, AFL has a much lower server-worker coordination complexity and avoids the aforementioned pitfalls in server-centric FL approaches. However, AFL also introduces significant challenges in algorithmic design on the server-side because the server needs to work much harder to handle the chaotic worker behaviors in AFL (e.g., asynchrony, spatial and temporal heterogeneity in computing). Toward this end, several foundational questions naturally arise: 1) *Is it possible to design algorithms that converge under AFL?* 2) *Under what condition and how fast could the algorithms converge?* 3) *If the answer to 1) is “yes” and 2) can also be resolved, could such algorithms still achieve the desired “linear speedup effect” as in conventional FL?*

In this paper, our goal is to obtain a fundamental understanding to the above questions. Our main contributions and key results are summarized as follows:

- We consider a new FL paradigm called Anarchic Federated Learning (AFL), where the workers are allowed to engage in training *at will* and choose the number of local update steps based on their own time-varying situations (computing resources, energy levels, etc.). This *loose worker-server coupling* significantly simplifies the implementations and renders AFL particularly suitable for FL deployments in edge computing environments. For any AFL algorithms under general worker information arrival processes and non-i.i.d. data across workers, we first establish a fundamental

convergence error lower bound to characterize the effect of worker participation in the AFL system.

- For AFL in the *cross-device* (CD) setting, we study the convergence of an anarchic federated averaging algorithm (AFA-CD), which is a natural counterpart of the popular FedAvg algorithm (McMahan et al., 2016) for server-centric FL. Our analysis reveals that, under bounded maximum delay, AFA-CD converges to an error ball whose size matches the fundamental lower bound, with an $\mathcal{O}(1/\sqrt{mKT})$ convergence rate. Here, m is the number of collected workers in each round of update, K is the local steps and T is the total number of rounds. We note that this convergence rate retains the highly desirable “linear speedup effect” in both worker’s number m and local steps K under AFL.¹ Moreover, under the stronger condition of uniform workers’ participation in AFL, AFA-CD converges to a singleton stationary point at the same convergence rate order that matches the state-of-the-art of server-centric FL.
- For AFL in the *cross-silo* (CS) setting, we study the convergence of a CS version of AFA (AFA-CS), where the special features of CS allow one to leverage historical feedback information and variance reduction techniques. We show that AFA-CS achieves an improved convergence rate of $\mathcal{O}(1/\sqrt{MKT})$, where M is the total number of workers. This suggests that, not only can “linear speedup” be achieved under AFA-CS, the speedup factor also depends on the total number of workers M instead of the number of collected workers m in each round ($M > m$).
- We validate both AFA algorithms with extensive experiments on CV and NLP tasks and explore the effect of the asynchrony and local step number in AFL. We also numerically show that AFA can be integrated with various advanced FL techniques (e.g., FedProx (Li et al., 2018) and SCAFFOLD (Karimireddy et al., 2020b)) to further enhance the AFL performance.

The rest of the paper is organized as follows. In Section 2, we review related work. In Section 3, we introduce AFL and the AFA algorithms, which are followed by their convergence analysis in Section 4. We present the numerical results in Section 5 and conclude the work in Section 6. Due to space limitation, we relegate all proofs and some experiments to the supplementary material.

2. Related Work

1) Server-Centric FL Algorithms: To date, one of the prevailing FL algorithms is Federated Averaging (FedAvg),

¹To attain ϵ -accuracy, it takes $\mathcal{O}(1/\epsilon^2)$ steps for an algorithm with an $\mathcal{O}(1/\sqrt{T})$ convergence rate, while needing $\mathcal{O}(1/m\epsilon^2)$ steps for another algorithm with an $\mathcal{O}(1/\sqrt{mT})$ convergence rate (the hidden constant in Big-O is the same). In this sense, $\mathcal{O}(1/\sqrt{mT})$ implies a *linear speedup* in terms of m .

which was first proposed in (McMahan et al., 2016) as a heuristic to improve communication efficiency and data privacy for FL. Since then, there have been substantial follow-ups of FedAvg that focus on non-i.i.d. (heterogeneous) data (see, e.g., FedProx (Li et al., 2018), FedPD (Zhang et al., 2020a), SCAFFOLD (Karimireddy et al., 2020b), FedNova (Wang et al., 2020), FedDyn (Acar et al., 2021), and MIME (Karimireddy et al., 2020a)), which are closely related to our work. The main idea for these algorithms is to control the “model drift” (due to heterogeneous datasets and the use of multiple local update steps on the worker side). While achieving various degrees of success in handling data heterogeneity, these algorithms are all server-centric and synchronous, which are more restrictive in edge-based settings (see discussions in Section 1).

2) FL with Flexible Worker Participation: To achieve high concurrency and avoid stragglers, researchers have proposed various FL methods with flexible worker participation, which can be roughly categorized into three classes: The first class utilizes different local steps to accommodate worker heterogeneity, while maintaining a synchronous communication between the server and workers (Wang et al., 2020; Ruan et al., 2021; Avdiukhin & Kasiviswanathan, 2021). The second class is based on asynchronous distributed optimization (Zhang et al., 2015; Lian et al., 2018; Niu et al., 2011; Agarwal & Duchi, 2012; Paine et al., 2013; Xie et al., 2019; Zhang et al., 2020b) (with identical local steps) (Nguyen et al., 2021; Avdiukhin & Kasiviswanathan, 2021). Specifically, Xie et al. (2019) proposed an asynchronous FL (FedAsync) method to tackle stragglers and heterogeneous latency, where the server continuously triggers one worker for local training. However, this work did not consider the convergence performance of non-convex problems that are more relevant to learning. Nguyen et al. (2021) utilized buffered asynchronous aggregation and achieved an $O(\frac{1}{\sqrt{TK}})$ convergence rate for non-convex problems, but it was unclear whether a linear speedup in terms of m is achievable. Avdiukhin & Kasiviswanathan (2021) proposed AsyncCommSGD by allowing asynchronous communication, while assuming an identical computation rate across all workers. This work achieved an $O(\frac{1}{\sqrt{mKT}})$ convergence rate for non-convex problems under a bounded gradient assumption, matching the convergence rate of synchronous FedAvg. The third class considers arbitrary device unavailability, though the server and workers still communicate in a synchronous fashion (i.e., following a system-wide common clock). In this class, the algorithms in (Gu et al., 2021) and (Yan et al., 2020) achieve $O(\frac{1}{\sqrt{MKT}})$ and $O(\frac{1}{\sqrt{M^{0.5}T}})$, respectively. However, Gu et al. (2021) required a Lipschitz Hessian assumption, where Yan et al. (2020) needed a bounded stochastic gradient assumption.

3) Key Differences of AFL from Related Works: Compar-

ing to the aforementioned related works, the AFL paradigm in this paper allows both heterogeneity: i) different local steps across workers and ii) asynchronous communications between the server and workers. In other words, the AFL paradigm subsumes all the above settings as special cases. Moreover, AFL fundamentally differs from aforementioned FL algorithms in that the worker’s participation in AFL and its local optimization process are completely *determined by the workers, and not by the sampling requests from the server*. This is more practical since it allows workers to participate in FL under drastically different situations in the network states, charging/idle cycles, etc. Due to the complex couplings between multiple sources of randomness and layers of heterogeneity in spatial and temporal domains in AFL, the training algorithm design for AFL and its theoretical analysis is far from a straightforward combination of existing FL techniques. Interestingly, we show that the AFA algorithms (counterparts of FedAvg under AFL) achieve the same convergence rate without strong assumptions (e.g., Lipschitz Hessian condition in (Gu et al., 2021)).

3. Anarchic Federated Learning

In this section, we first formally define the notion of AFL. Then, we will present the AFA algorithmic framework, which contains two variants called AFA-CD and AFA-CS for cross-device and cross-silo AFL, respectively.

1) Overview of Anarchic Federated Learning: The goal of FL is to solve an optimization problem in the form of $\min_{x \in \mathbb{R}^d} f(x) := \frac{1}{M} \sum_{i=1}^M f_i(x)$, where $f_i(x) \triangleq \mathbb{E}_{\xi_i \sim D_i} [f_i(x, \xi_i)]$ is the local (non-convex) loss function associated with a local data distribution D_i and M is the total number of workers. For the setting with heterogeneous (non-i.i.d.) datasets at the workers, we have $D_i \neq D_j$, if $i \neq j$. In terms of the assumption on the size of workers, FL can be classified as cross-device FL and cross-silo FL (Kairouz et al., 2019; Wang et al., 2021). Cross-device FL is designed for large-scale FL with a massive number of mobile or IoT devices (M is large). As a result, the server can only afford to collect information from a subset of workers in each round of update and is unable to store workers’ information across rounds. In comparison, the number of workers in cross-silo FL is relatively small. Although the server in cross-silo FL may still have to collect information only from a subset of workers in each round, it has enough capacity to store each worker’s most recent information.

The general framework of AFL is illustrated in Algorithm 1. Here, the server is responsible for: 1) collecting the local updates returned from workers, and 2) aggregating the obtained updates once certain conditions are satisfied (e.g., upon collecting $m \in (0, M]$ local updates from workers) to update the global model. Note that these two threads are *concurrent*, so it completely avoids system locking on

Algorithm 1 The General AFL Framework.

At the Server (Concurrently with Workers):

1. (Concurrent Thread) Collect local updates returned from the workers.
2. (Concurrent Thread) Aggregate local update returned from collected workers and update global model following some server-side optimization process.

At Each Worker (Concurrently with Server):

1. Once decided to participate in the training, pull the global model with current timestamp.
 2. Perform (multiple) local update steps following some worker-side optimization process.
 3. Return the result and the associated pulling timestamp to the server, with extra processing if so desired.
-

server’s side. Also, idling is allowed at each worker between each two successive participations in training. Whenever a worker intends to participate in the training, it first pulls the current model parameters from the server. Then, upon finishing multiple local update steps (more on this later) by some worker-side optimization process (e.g., using stochastic gradients or additional information such as variance-reduced and/or momentum adjustments), the worker reports the results to the server (potentially with extra processing if so desired, e.g., compression for communication efficiency).

We remark that AFL is a general computing architecture that subsumes the conventional FL and asynchronous distributed optimization as special cases. From an optimization perspective, the server and the workers may adopt independent optimization processes, thus enabling a much richer set of learning “control knobs” (e.g., separated learning rates, separated batch sizes). Specifically, each worker is able to completely take control of its own optimization process, even using a time-varying number of local update steps and optimizers, which depend on its local dataset and/or its device status (e.g., battery level, privacy preference). More importantly, from the system level, the concurrent processes at worker and server side enable loose worker-server coupling and thus avoiding server-worker interlocking and reducing synchronization overhead.

2) A Convergence Error Lower Bound for AFL: To thoroughly understand AFL, we will first obtain some fundamental insights on the *performance limit of any AFL training algorithms*. Toward this end, we first state several assumptions that are needed for our theoretical analysis throughout the rest of this paper.

Assumption 1. (*L-Lipschitz Continuous Gradient*) *There exists a constant $L > 0$, such that $\|\nabla f_i(\mathbf{x}) - \nabla f_i(\mathbf{y})\| \leq L\|\mathbf{x} - \mathbf{y}\|$, $\forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^d$, and $i \in [M]$.*

Assumption 2. (*Unbiased Local Stochastic Gradient*) *Let ξ^i be a random local data sample at worker i . The local stochastic gradient is unbiased, i.e., $\mathbb{E}[\nabla f_i(\mathbf{x}, \xi^i)] = \nabla f_i(\mathbf{x})$, $\forall i \in [m]$, where the expectation is taken over the local data distribution D_i .*

Assumption 3. (*Bounded Local and Global Variances*) *There exist two constants $\sigma_L \geq 0$ and $\sigma_G \geq 0$, such that the variance of each local stochastic gradient estimator is bounded by $\mathbb{E}[\|\nabla f_i(\mathbf{x}, \xi^i) - \nabla f_i(\mathbf{x})\|^2] \leq \sigma_L^2$, $\forall i \in [M]$, and the global variability of local gradient of the cost function is bounded by $\|\nabla f_i(\mathbf{x}) - \nabla f(\mathbf{x})\|^2 \leq \sigma_G^2$, $\forall i \in [M]$.*

The first two assumptions are standard in the convergence analysis of non-convex optimization (see, e.g., (Ghadimi & Lan, 2013; Bottou et al., 2018)). For Assumption 3, the bounded local variance is also a standard assumption. We utilize a universal bound σ_G to quantify the data heterogeneity among different workers. This assumption is also frequently used in the literature of FL with non-i.i.d. datasets (Reddi et al., 2020; Wang et al., 2019; Yang et al., 2021) as well as in decentralized optimization (Kairouz et al., 2019).

To establish a fundamental convergence error lower bound, we consider the most general case where *no assumption* on the arrival processes of the worker information is made, except that each worker’s participation in FL is independent of each other. In such general worker information arrival processes, we prove the following lower bound of convergence error by constructing a worst-case scenario:

Theorem 1 (Convergence Error Lower Bound for AFL with General Worker Information Arrival Processes). *For any level of heterogeneity characterized by σ_G , there exists loss functions satisfying Assumptions 1- 3 and a specific worker participation process for which the output $\hat{\mathbf{x}}$ of any convergent (and potentially random) FL algorithm satisfies:*

$$\mathbb{E}[\|\nabla f(\hat{\mathbf{x}})\|^2] = \Omega(\sigma_G^2).$$

Remark 1. (Proof in Appendix B.1) The lower bound in Theorem 1 indicates that no algorithms for AFL could converge to a stationary point under general worker information arrival processes, due to the significant system heterogeneity and randomness caused by such general worker information arrivals. The rationale is that there always exist objective value drifts owing to general worker information arrival processes in the worst-case scenario, which further lead to an inevitable error in convergence. We note that this lower bound is different from previous optimization lower bounds in FL (Karimireddy et al., 2020b; Woodworth et al., 2020; Gu et al., 2021). Our lower bound captures objective deviations due to worker participation while previous bounds focus on the optimization error with ideal worker participation (i.e., full worker or uniformly random worker participation). Considering a worst-case scenario in FL by removing such assumption of ideal worker participation, our lower bound

also holds for non-i.i.d. FL including synchronous FedAvg and its variants, thus also providing insights for conventional FL. To ensure convergence to a stationary point, extra assumptions for the worker information arrivals need to be made, e.g., uniformly distributed arrivals (see Theorem 3) and bounded delays (see Theorem 4).

4. The Anarchic Federated Averaging (AFA) Algorithms for AFL

Upon obtaining a basic understanding of the training algorithm performance limit from the convergence error in Theorem 1, in this section, we study convergence conditions and performance of two anarchic federated averaging (AFA) algorithms for cross-device (CD) and cross-silo (CS) settings in Section 4.1 and 4.2, respectively, both of which can be viewed as an extension of FedAvg under AFL.

4.1. The AFA-CD Algorithm for Cross-Device AFL

1) The AFA-CD Algorithm: First, we consider the AFA-CD algorithm for the cross-device AFL setting. As mentioned earlier, cross-device AFL is suitable for cases with a massive number of edge devices. In each round of global model update, only a small subset of workers are used in the training. The server is assumed to have *no historical information* of the workers. As shown in Algorithm 2, AFA-CD closely follows the AFL architecture shown in Algorithm 1. Here, we use the standard stochastic gradient descent (SGD) method as the server- and worker-side optimizer. In each update $t = 0, \dots, T - 1$, the server waits until collecting m local updates $\{\mathbf{G}_i(\mathbf{x}_{t-\tau_{t,i}})\}$ from workers to form a set \mathcal{M}_t with $|\mathcal{M}_t| = m$, where $\tau_{t,i}$ represents the random delay of the local update of worker $i \in \mathcal{M}_t$ (Server Code, Line 1). Once \mathcal{M}_t is formed, the server aggregates all local updates $\mathbf{G}_i(\mathbf{x}_{t-\tau_{t,i}}), i \in \mathcal{M}_t$ and updates global model (Server Code, Line 2). We count each global model update as one communication round for direct comparison with previous FL results. Meanwhile, for each worker, it pulls the current global model parameters with time stamp μ once it decides to participate in training (Worker Code, Line 1). Each worker can then choose a desired number of local update steps $K_{t,i}$ (could be time-varying and device-dependent) to perform SGD updates for $K_{t,i}$ times, and then return the rescaled sum of all stochastic gradients with timestamp μ to the server (Worker Code, Lines 2–3).

2) Convergence Analysis of the AFA-CD Algorithm: We first analyze the convergence of AFA-CD under general worker information arrival processes. We use $f_0 = f(\mathbf{x}_0)$ and f_* to denote the initial and the optimal objective values, respectively. We have the following convergence result for the AFA-CD algorithm (see proof details in Appendix B.2):

Theorem 2 (AFA-CD with General Worker Information Ar-

Algorithm 2 AFA-CD Algorithm for Cross-Device AFL.

At the Server (Concurrently with Workers):

1. In the t -th update round, collect m local updates $\{\mathbf{G}_i(\mathbf{x}_{t-\tau_{t,i}}), i \in \mathcal{M}_t\}$ returned from the workers to form the set \mathcal{M}_t , where $\tau_{t,i}$ represents the random delay of the worker i 's local update, $i \in \mathcal{M}_t$.
2. Aggregate and update: $\mathbf{G}_t = \frac{1}{m} \sum_{i \in \mathcal{M}_t} \mathbf{G}_i(\mathbf{x}_{t-\tau_{t,i}}), \mathbf{x}_{t+1} = \mathbf{x}_t - \eta \mathbf{G}_t$.

At Each Worker (Concurrently with Server):

1. Once decided to participate in the training, retrieve the parameter \mathbf{x}_μ from the server and its timestamp, set the local model: $\mathbf{x}_{\mu,0}^i = \mathbf{x}_\mu$.
 2. Choose a number of local steps $K_{t,i}$, which can be time-varying and device-dependent. Let $\mathbf{x}_{\mu,k+1}^i = \mathbf{x}_{\mu,k}^i - \eta_L \mathbf{g}_{\mu,k}^i$, where $\mathbf{g}_{\mu,k}^i = \nabla f_i(\mathbf{x}_{\mu,k}^i, \xi_{\mu,k}^i), k = 0, \dots, K_{t,i} - 1$.
 3. Sum and rescale the stochastic gradients: $\mathbf{G}_i(\mathbf{x}_\mu) = \frac{1}{K_{t,i}} \sum_{j=0}^{K_{t,i}-1} \mathbf{g}_{\mu,j}^i$. Return $\mathbf{G}_i(\mathbf{x}_\mu)$.
-

rival Processes). Suppose that the resultant maximum delay under AFL is bounded, i.e., $\tau := \max_{t \in [T], i \in \mathcal{M}_t} \{\tau_{t,i}\} < \infty$. Suppose that the server-side and worker-side learning rates η and η_L are chosen as such that the following conditions are satisfied: $6\eta_L^2(2K_{t,i}^2 - 3K_{t,i} + 1)L^2 \leq 1, 180\eta_L^2 K_{t,i}^2 L^2 \tau < 1, \forall t, i$ and $2L\eta\eta_L + 6\tau^2 L^2 \eta^2 \eta_L^2 \leq 1$. Under Assumptions 1–3, the output sequence $\{\mathbf{x}_t\}$ generated by AFA-CD with general worker information arrival processes satisfies:

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\mathbf{x}_t)\|^2 \leq \frac{4(f_0 - f_*)}{\eta\eta_L T} + 4(\alpha_L \sigma_L^2 + \alpha_G \sigma_G^2),$$

where the constants α_L and α_G are defined as:

$$\begin{aligned} \alpha_L &= \frac{L\eta\eta_L}{m} \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{K_t} + \frac{3\tau^2 L^2 \eta^2 \eta_L^2}{m} \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{K_t} \\ &\quad + \frac{15\eta_L^2 L^2}{2} \frac{1}{T} \sum_{t=0}^{T-1} \bar{K}_t, \\ \alpha_G &= \frac{3}{2} + 45L^2 \eta_L^2 \frac{1}{T} \sum_{t=0}^{T-1} \hat{K}_t^2. \end{aligned}$$

Here,

$$\frac{1}{\bar{K}_t} = \frac{1}{m} \sum_{i \in \mathcal{M}_t} \frac{1}{K_{t,i}}, \bar{K}_t = \frac{1}{m} \sum_{i \in \mathcal{M}_t} K_{t,i}, \hat{K}_t^2 = \frac{1}{m} \sum_{i \in \mathcal{M}_t} K_{t,i}^2.$$

The learning rates conditions imply that $\eta\eta_L = \mathcal{O}(\frac{1}{\tau L})$ and $\eta_L^2 \leq \frac{1}{\bar{K}_{t,i}^2}$, which is a natural extension of that in SGD. With

Theorem 2, if we assume a constant local step number and proper learning rates, we immediately have the following convergence rate for AFA-CD, which implies the “linear speedup effect” in both m and K .

Corollary 1 (Linear Speedup to an Error Ball). *Suppose a constant local step K for each worker, by setting $\eta_L = \frac{1}{\sqrt{T}}$, and $\eta = \sqrt{mK}$, the convergence rate of AFA-CD with general worker information arrival processes is:*

$$\mathcal{O}\left(\frac{1}{m^{1/2}K^{1/2}T^{1/2}}\right) + \mathcal{O}\left(\frac{\tau^2}{T}\right) + \mathcal{O}\left(\frac{K^2}{T}\right) + \mathcal{O}(\sigma_G^2).$$

Remark 2. Clearly, due to the chaotic worker behaviors in AFL, one cannot expect that an algorithm for AFL can converge under any arbitrary condition. Theorem 2 and Corollary 1 suggest that, as long as the consequence of the chaotic workers behaviors remains “manageable” in the sense that i) the maximum delay due to asynchrony is bounded and ii) the learning rates used by the workers and server are sufficiently small, then the iterates produced by AFA-CD can converge to a neighborhood around a stationary point. Moreover, if the workers are less “anarchic” in the sense that they know the T -value from the server and are willing to set $\eta_L = \frac{1}{\sqrt{T}}$ accordingly, then the non-vanishing error term $\mathcal{O}(\sigma_G^2)$ in Corollary 1 matches the lower bound in Theorem 1. This implies that the convergence error of AFL-CD is *order-optimal* in this setting.

Remark 3. Recall that the non-vanishing convergence error $\mathcal{O}(\sigma_G^2)$ in Corollary 1 is a consequence of objective function drift under the general worker information arrivals (no assumption on the arrivals of the worker participation in each round of update) and is independent of the choices of learning rates, the number of local update steps, and the number of global update rounds (more discussion in the supplementary material). Also, for a sufficiently large T , the dominant term $\mathcal{O}\left(\frac{1}{m^{1/2}K^{1/2}T^{1/2}}\right)$ implies that AFA-CD achieves the linear speedup in terms of m and K before reaching a constant error neighborhood with size $\mathcal{O}(\sigma_G^2)$.

Given the weak convergence result under general workers’ information arrivals, it is important to understand what extra conditions on the worker information arrivals are needed under AFL in order to achieve stronger convergence performance. Toward this end, we consider a special setting where the arrivals of worker returned information in each round for global update is uniformly distributed among the workers. In this setting, \mathcal{M}_t can be viewed as a subset with size m independently and uniformly sampled from $[M]$ without replacement. It has been empirically found in (McMahan et al., 2016; Li et al., 2019a) that, for FL systems with a massive number of workers, the assumption of uniformly distributed arrivals is a good approximation for worker participation in cross-device FL. In what follows, we show that the convergence performance of AFA-CD in this special

setting can be improved as follows (see proof details in Appendix B.3):

Theorem 3. *Under the same delay condition in Theorem 2 and suppose that the server-side and worker-side learning rates η and η_L are chosen as such that the following relationships hold: $6\eta_L^2(2K_{t,i}^2 - 3K_{t,i} + 1)L^2 \leq 1, \forall t, i$, $L\eta\eta_L + L^2\eta^2\eta_L^2\tau^2 \leq \frac{1}{2M}$, and $120L^2\hat{K}_t^2\eta_L^2\tau < 1, \forall t$. Then, under Assumptions 1–3, the output sequence $\{\mathbf{x}_t\}$ generated by AFA-CD with uniformly distributed worker information arrivals satisfies:*

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\mathbf{x}_t)\|_2^2 \leq \frac{4(f_0 - f_*)}{\eta\eta_L T} + 4(\alpha_L \sigma_L^2 + \alpha_G \sigma_G^2),$$

where α_L and α_G are defined as following:

$$\begin{aligned} \alpha_L &= \frac{L\eta\eta_L}{m} \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{K_t} + \frac{2\tau^2 L^2 \eta^2 \eta_L^2}{m} \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{K_t} \\ &\quad + 5\eta_L^2 L^2 \frac{1}{T} \sum_{t=0}^{T-1} \bar{K}_t, \\ \alpha_G &= 30L^2 \eta_L^2 \frac{1}{T} \sum_{t=0}^{T-1} \hat{K}_t^2, \end{aligned}$$

and other parameters are defined the same as in Theorem 2.

The requirement for learning rates could be easily satisfied as that in Theorem 2. Furthermore, with appropriate server- and worker-side learning rates, we immediately have the following linear speedup convergence result for AFA-CD:

Corollary 2 (Linear Speedup to a Stationary Point). *Suppose a constant local step K , let $\eta_L = \frac{1}{\sqrt{T}}$, and $\eta = \sqrt{mK}$, the convergence rate of AFA-CD with uniformly distributed worker information arrivals is:*

$$\mathcal{O}\left(\frac{1}{m^{1/2}K^{1/2}T^{1/2}}\right) + \mathcal{O}\left(\frac{\tau^2}{T}\right) + \mathcal{O}\left(\frac{K^2}{T}\right)$$

Remark 4. For a sufficiently large T , the linear speedup convergence to a stationary point (rather than a constant error neighborhood) can be achieved under bounded maximum delay τ , i.e., $\mathcal{O}\left(\frac{1}{m^{1/2}K^{1/2}T^{1/2}}\right)$. Note that this rate does not depend on the delay τ after sufficiently many rounds T (i.e., $\tau \leq \min\left\{\frac{T^{1/4}}{m^{1/4}K^{1/4}}, \frac{T^{1/2}}{m^{1/2}K^{5/2}}\right\}$), the negative effect of using outdated information in such an asynchronous setting vanishes asymptotically. Further, for $\sigma_G = 0$ (i.i.d. data) and $K = 1$ (single local update step), AFA-CD can be viewed as an extension of the AsySG-con algorithm (Lian et al., 2015) in asynchronous parallel distributed optimization. It can be readily verified that AFA-CD achieves the same rate as that of the AsySG-con algorithm. Furthermore, AsyncCommSGD (Avdiukhin & Kasiviswanathan, 2021) achieves $\mathcal{O}\left(\frac{1}{\sqrt{mKT}}\right)$ for FL by allowing asynchronous communication assuming an identical

Algorithm 3 The AFA-CS Algorithm for Cross-Silo AFL.

At the Server (Concurrently with Workers):

1. In the t -th update round, collect m local updates.
2. Update worker i 's information in the memory using the returned local update \mathbf{G}_i .
3. Aggregate and update: $\mathbf{G}_t = \frac{1}{M} \sum_{i \in [M]} \mathbf{G}_i$,
 $\mathbf{x}_{t+1} = \mathbf{x}_t - \eta \mathbf{G}_t$.

At Each Worker (Concurrently with Server): Same as AFA-CD Worker Code.

computation rate across workers and bounded gradients. Interestingly, AFA-CD achieves the same convergence rate while allowing flexible worker participation and without such assumptions. Surprisingly, this rate even matches the best known rate for the general non-convex setting in FL (Karimireddy et al., 2020b; Reddi et al., 2020). It is worth noting that Nguyen et al. (2021) proposed the FedBuff algorithm for FL, which is akin to AFA-CD and boosts FL concurrency. However, FedBuff achieves an $O(\frac{1}{\sqrt{TK}})$ convergence rate, which does not achieve the linear speedup in terms of m .

4.2. The AFA-CS Algorithm for Cross-Silo AFL

1) The AFA-CS Algorithm: As mentioned earlier, cross-silo FL is suitable for collaborative learning among a relatively small number of (organizational) workers. Thanks to the relatively small number of workers, each worker's feedback can be stored at the server. As a result, the server could reuse the historical information of each specific worker in each round of global update.

As shown in Algorithm 3, the AFA-CS algorithm also closely follows the AFL architecture as shown in Algorithm 1. In each round of global model update, a subset of workers could participate in the training (Server Code, Line 1). Compared to AFA-CD, the key difference in AFA-CS is in Line 2 of the Server Code, where the server stores the collected local updates $\{\mathbf{G}_i\}$ for each worker $i \in \mathcal{M}_t$ into the memory space at the server (Server Code, Line 2). As a result, whenever a worker i returns a local update to the server upon finishing its local update steps, the server will update the memory space corresponding to worker i to replace the old information with this newly received update from worker i . Similar to AFA-CD, every m new updates in the AFA-CS algorithm trigger the server to aggregate all the $\mathbf{G}_i, i \in [M]$ and update the global model. The Worker Code in AFA-CS is exactly the same as AFA-CD and its description is omitted for brevity.

2) Convergence Analysis of the AFA-CS Algorithm: We divide stochastic gradient returns $\{\mathbf{G}_i\}$ into two groups. One is for those without delay ($\mathbf{G}_i(x_t), i \in \mathcal{M}_t, |\mathcal{M}_t| =$

m') and the other is for those with delay ($\mathbf{G}_i(\mathbf{x}_{t-\tau_{t,i}}), i \in \mathcal{M}_t^c, |\mathcal{M}_t^c| = M - m'$). For cross-silo AFL, the AFA-CS algorithm achieves the following convergence performance (see proof details in Appendix C):

Theorem 4. *Suppose that the resultant maximum delay in the system is bounded, i.e., $\tau := \max_{t \in [T], i \in \mathcal{M}_t^c} \{\tau_{t,i}\} < \infty$. Suppose that the server-side and worker-side learning rates η and η_L are chosen as such that the following relationships hold: $6\eta_L^2(2K_{t,i}^2 - 3K_{t,i} + 1)L^2 \leq 1, \forall t, i, \left(\frac{\eta\eta_L(M-m')^2L^2\tau^2}{M^2} + \frac{L}{2}\right)\eta\eta_L \leq \frac{1}{4}$, and $\frac{30L^2\eta_L^2\tau}{M} \left(\sum_{i \in [M]} K_{t,i}^2\right) \leq \frac{1}{4}$. Then, under Assumptions 1-3, the output sequence $\{\mathbf{x}_t\}$ generated by the AFA-CS algorithm under general worker information arrival processes satisfies:*

$$\frac{1}{T} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|^2 \leq \frac{4f(\mathbf{x}_0) - f(\mathbf{x}_T)}{\eta\eta_L T} + \alpha_L \sigma_L^2 + \alpha_G \sigma_G^2,$$

where the constants α_L and α_G are defined as follows:

$$\begin{aligned} \alpha_L &= \frac{4}{M} \left[5L^2\eta_L^2 \frac{1}{T} \sum_{t=0}^{T-1} \bar{K}_t \right. \\ &\quad \left. + \left(\frac{2\eta^2\eta_L^2(M-m')^2L^2\tau^2}{M^2} + L\eta\eta_L \right) \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{K_t} \right], \\ \alpha_G &= \frac{120L^2\eta_L^2}{M} \frac{1}{T} \sum_{t=0}^{T-1} \hat{K}_t^2, \end{aligned}$$

and other parameters are defined the same as in Theorem 2.

With appropriate learning rates, we immediately have stronger linear speedup convergence:

Corollary 3 (Linear Speedup). *Suppose a constant local step K , and let $\eta_L = \frac{1}{\sqrt{T}}$, and $\eta = \sqrt{MK}$, the convergence rate of the AFA-CS algorithm under general worker information arrival processes is:*

$$\mathcal{O}\left(\frac{1}{M^{1/2}K^{1/2}T^{1/2}}\right) + \mathcal{O}\left(\frac{K^2}{MT}\right) + \mathcal{O}\left(\frac{\tau^2(M-m')^2}{TM^2}\right).$$

Remark 5. Compared to Corollary 1, we can see that, by reusing historical data, AFA-CS can eliminate the non-vanishing $\mathcal{O}(\sigma_G^2)$ error term even under general worker information arrival processes and bounded delay. The bounded delay implicitly requires each workers at least participate in the training process, eliminating the worst-case scenario in Theorem 1. On the other hand, although the server only collects m workers' feedback in each round of global model update, the server leverages all M workers' feedback by reusing historical information. Intuitively, this translates the potential objection function drift originated

from general worker information arrival process into the negative effect of delayed returns $\mathbf{G}(\mathbf{x}_{t-\tau_{t,i}})$ from workers. It can be shown that such a negative effect vanishes asymptotically as the number of communication rounds T gets large and in turn diminishes the convergence error. This also explains the *stronger* linear speedup $\mathcal{O}(1/\sqrt{MT})$. Specifically, even with partial (m) workers participation in each round, AFA-CS achieves a speedup with respect to total number of workers M ($M > m$). From the lower bound in FL (Proposition 6.1 in Gu et al. (2021)), Corollary 3 is tight.

Remark 6. AFA-CS generalizes the lazy aggregation strategy in distributed learning (e.g., LAG (Chen et al., 2018)) by setting $K = 1$ (single local update), $\tau = 0$ (synchronous setting) and $\sigma_L = 0$ (using full gradient descents instead of stochastic gradients) and further improve the rate of LSA (Chen et al., 2020) from $\mathcal{O}(1/\sqrt{T})$ to $\mathcal{O}(1/\sqrt{MT})$. We note that Gu et al. (2021) and Yan et al. (2020) achieved $\mathcal{O}(\frac{1}{\sqrt{MKT}})$ and $\mathcal{O}(\frac{1}{\sqrt{M^{0.5}T}})$ for FL, respectively, by using historical information, which is similar to AFA-CS. However, they both requires additional assumptions. Specifically, Gu et al. (2021) required a Lipschitz Hessian assumption and Yan et al. (2020) needed bounded stochastic gradient assumption. By contrast, AFA-CS achieves the same optimal rate without such assumptions.

5. Numerical results

In this section, we conduct experiments to verify our theoretical results. We use i) logistic regression (LR) on manually partitioned non-i.i.d. MNIST dataset (LeCun et al., 1998), ii) convolutional neural network (CNN) for manually partitioned CIFAR-10 (Krizhevsky, 2009), and iii) recurrent neural network (RNN) on natural non-i.i.d. dataset *Shakespeare* (McMahan et al., 2016). In order to impose data heterogeneity in MNIST and CIFAR-10 data, we distribute the data evenly into each worker in label-based partition following the same process in the literature (e.g., McMahan et al. (2016); Yang et al. (2021); Li et al. (2019c)). Therefore, we can use a parameter p to represent the classes of labels in each worker’s dataset, which signifies data heterogeneity: the smaller the p -value, the more heterogeneous the data across workers (cf. Yang et al. (2021); Li et al. (2019c) for details). Due to space limitation, we relegate the details of models, datasets and hyper-parameters, and further results of CNN and RNN to the appendix.

In Figure 1, we illustrate the test accuracy for LR on MNIST with different p -values. We use the classical FedAvg algorithm (McMahan et al., 2016) for conventional FL with uniform worker sampling as a baseline, since it corresponds to the most ideal scenario where workers are fully cooperative with the server. We examine the learning performance degradation of AFA algorithms (due to anarchic worker behaviors) compared to this ideal baseline. For our AFA-CD




Figure 1. Test accuracy for logistic regression on non-i.i.d. MNIST with different p -values.

and AFA-CS with general worker information arrival processes, the test accuracy is comparable to or nearly the same as that of FedAvg. This confirms our theoretical results and validates the effectiveness of our AFA algorithms. We further evaluate the impacts of various factors in AFL, including asynchrony, heterogeneous computing, worker’s arrival process, and non-i.i.d. datasets, on convergence rate of our proposed AFA algorithms. Note that AFL subsumes FedAvg and many variants when the above hyper-parameters are set as constant. Also, AFL coupled with other FL algorithms such as FedProx (Li et al., 2018) and SCAFFOLD (Karimireddy et al., 2020b) is tested. Our results show that the AFA algorithms are robust against all asynchrony and heterogeneity factors in AFL. Due to space limitation, we refer readers to the appendix for all these experimental results.

6. Conclusions

In this paper, we propose a new paradigm in FL called “Anarchic Federated Learning” (AFL). In stark contrast to conventional FL models where the server and the worker are tightly coupled, AFL has a much lower server-worker coordination complexity, allowing a flexible worker participation. We proposed two Anarchic Federated Averaging algorithms with two-sided learning rates for both cross-device and cross-silo settings, which are named AFA-CD and AFA-CS, respectively. We showed that both algorithms retain the highly desirable linear speedup effect in the new AFL paradigm. Moreover, we showed that our AFL framework works well numerically by employing advance FL algorithms FedProx and SCAFFOLD as the optimizer in worker’s side.

Acknowledgements

This work has been supported in part by NSF grants CA-REER CNS-2110259, CNS-2112471, CNS-2102233, CCF-2110252, and a Google Faculty Research Award.

References

- Acar, D. A. E., Zhao, Y., Navarro, R. M., Mattina, M., Whatmough, P. N., and Saligrama, V. Federated learning based on dynamic regularization. In *International Conference on Learning Representations*, 2021.
- Agarwal, A. and Duchi, J. C. Distributed delayed stochastic optimization. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pp. 5451–5452. IEEE, 2012.
- Avdiukhin, D. and Kasiviswanathan, S. Federated learning under arbitrary communication patterns. In *International Conference on Machine Learning*, pp. 425–435. PMLR, 2021.
- Bottou, L., Curtis, F. E., and Nocedal, J. Optimization methods for large-scale machine learning. *Siam Review*, 60(2):223–311, 2018.
- Charles, Z., Garrett, Z., Huo, Z., Shmulyian, S., and Smith, V. On large-cohort training for federated learning. *Advances in Neural Information Processing Systems*, 34, 2021.
- Chen, T., Giannakis, G. B., Sun, T., and Yin, W. Lag: Lazily aggregated gradient for communication-efficient distributed learning. In *NeurIPS*, 2018.
- Chen, T., Sun, Y., and Yin, W. Lasg: Lazily aggregated stochastic gradients for communication-efficient distributed learning. *arXiv preprint arXiv:2002.11360*, 2020.
- Defazio, A. and Bottou, L. On the ineffectiveness of variance reduced optimization for deep learning. *arXiv preprint arXiv:1812.04529*, 2018.
- Ghadimi, S. and Lan, G. Stochastic first-and zeroth-order methods for nonconvex stochastic programming. *SIAM Journal on Optimization*, 23(4):2341–2368, 2013.
- Gu, X., Huang, K., Zhang, J., and Huang, L. Fast federated learning in the presence of arbitrary device unavailability. *arXiv preprint arXiv:2106.04159*, 2021.
- Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., et al. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*, 2019.
- Karimireddy, S. P., Jaggi, M., Kale, S., Mohri, M., Reddi, S. J., Stich, S. U., and Suresh, A. T. Mime: Mimicking centralized stochastic algorithms in federated learning. *arXiv preprint arXiv:2008.03606*, 2020a.
- Karimireddy, S. P., Kale, S., Mohri, M., Reddi, S., Stich, S., and Suresh, A. T. Scaffold: Stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*, pp. 5132–5143. PMLR, 2020b.
- Khaled, A., Mishchenko, K., and Richtárik, P. First analysis of local gd on heterogeneous data. *arXiv preprint arXiv:1909.04715*, 2019.
- Konecny, J., McMahan, H. B., Ramage, D., and Richtárik, P. Federated optimization: Distributed machine learning for on-device intelligence. *arXiv preprint arXiv:1610.02527*, 2016.
- Krizhevsky, A. Learning multiple layers of features from tiny images. 2009.
- Le Roux, N., Schmidt, M. W., and Bach, F. R. A stochastic gradient method with an exponential convergence rate for finite training sets. In *NIPS*, 2012.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Li, T., Sahu, A. K., Zaheer, M., Sanjabi, M., Talwalkar, A., and Smith, V. Federated optimization in heterogeneous networks. *arXiv preprint arXiv:1812.06127*, 2018.
- Li, T., Sahu, A. K., Talwalkar, A., and Smith, V. Federated learning: Challenges, methods, and future directions. *arXiv preprint arXiv:1908.07873*, 2019a.
- Li, T., Sanjabi, M., Beirami, A., and Smith, V. Fair resource allocation in federated learning. *arXiv preprint arXiv:1905.10497*, 2019b.
- Li, X., Huang, K., Yang, W., Wang, S., and Zhang, Z. On the convergence of fedavg on non-iid data. *arXiv preprint arXiv:1907.02189*, 2019c.
- Lian, X., Huang, Y., Li, Y., and Liu, J. Asynchronous parallel stochastic gradient for nonconvex optimization. *Advances in Neural Information Processing Systems*, 28: 2737–2745, 2015.
- Lian, X., Zhang, W., Zhang, C., and Liu, J. Asynchronous decentralized parallel stochastic gradient descent. In *International Conference on Machine Learning*, pp. 3043–3052. PMLR, 2018.
- McMahan, H. B., Moore, E., Ramage, D., Hampson, S., et al. Communication-efficient learning of deep networks from decentralized data. *arXiv preprint arXiv:1602.05629*, 2016.

- Mohri, M., Sivek, G., and Suresh, A. T. Agnostic federated learning. In *International Conference on Machine Learning*, pp. 4615–4625. PMLR, 2019.
- Nguyen, J., Malik, K., Zhan, H., Yousefpour, A., Rabbat, M., Esmaeili, M. M., and Huba, D. Federated learning with buffered asynchronous aggregation. *arXiv preprint arXiv:2106.06639*, 2021.
- Niu, F., Recht, B., Ré, C., and Wright, S. J. Hogwild!: A lock-free approach to parallelizing stochastic gradient descent. *arXiv preprint arXiv:1106.5730*, 2011.
- Paine, T., Jin, H., Yang, J., Lin, Z., and Huang, T. Gpu asynchronous stochastic gradient descent to speed up neural network training. *arXiv preprint arXiv:1312.6186*, 2013.
- Qu, Z., Lin, K., Kalagnanam, J., Li, Z., Zhou, J., and Zhou, Z. Federated learning’s blessing: Fedavg has linear speedup. *arXiv preprint arXiv:2007.05690*, 2020.
- Reddi, S., Charles, Z., Zaheer, M., Garrett, Z., Rush, K., Konecny, J., Kumar, S., and McMahan, H. B. Adaptive federated optimization. *arXiv preprint arXiv:2003.00295*, 2020.
- Ruan, Y., Zhang, X., Liang, S.-C., and Joe-Wong, C. Towards flexible device participation in federated learning. In *International Conference on Artificial Intelligence and Statistics*, pp. 3403–3411. PMLR, 2021.
- Schmidt, M., Le Roux, N., and Bach, F. Minimizing finite sums with the stochastic average gradient. *Mathematical Programming*, 162(1-2):83–112, 2017.
- Stich, S. U. Local sgd converges fast and communicates little. *arXiv preprint arXiv:1805.09767*, 2018.
- Wang, J. and Joshi, G. Cooperative sgd: A unified framework for the design and analysis of communication-efficient sgd algorithms. *arXiv preprint arXiv:1808.07576*, 2018.
- Wang, J., Liu, Q., Liang, H., Joshi, G., and Poor, H. V. Tackling the objective inconsistency problem in heterogeneous federated optimization. *arXiv preprint arXiv:2007.07481*, 2020.
- Wang, J., Charles, Z., Xu, Z., Joshi, G., McMahan, H. B., Al-Shedivat, M., Andrew, G., Avestimehr, S., Daly, K., Data, D., et al. A field guide to federated optimization. *arXiv preprint arXiv:2107.06917*, 2021.
- Wang, S., Tuor, T., Salonidis, T., Leung, K. K., Makaya, C., He, T., and Chan, K. Adaptive federated learning in resource constrained edge computing systems. *IEEE Journal on Selected Areas in Communications*, 37(6): 1205–1221, 2019.
- Woodworth, B., Patel, K. K., and Srebro, N. Minibatch vs local sgd for heterogeneous distributed learning. *arXiv preprint arXiv:2006.04735*, 2020.
- Xie, C., Koyejo, S., and Gupta, I. Asynchronous federated optimization. *arXiv preprint arXiv:1903.03934*, 2019.
- Yan, Y., Niu, C., Ding, Y., Zheng, Z., Wu, F., Chen, G., Tang, S., and Wu, Z. Distributed non-convex optimization with sublinear speedup under intermittent client availability. *arXiv preprint arXiv:2002.07399*, 2020.
- Yang, H., Fang, M., and Liu, J. Achieving linear speedup with partial worker participation in non- $\{iid\}$ federated learning. In *International Conference on Learning Representations*, 2021.
- Yang, Q., Liu, Y., Chen, T., and Tong, Y. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2):1–19, 2019.
- Yu, H., Jin, R., and Yang, S. On the linear speedup analysis of communication efficient momentum sgd for distributed non-convex optimization. In *International Conference on Machine Learning*, pp. 7184–7193. PMLR, 2019.
- Zhang, S., Choromanska, A. E., and LeCun, Y. Deep learning with elastic averaging sgd. *Advances in neural information processing systems*, 28, 2015.
- Zhang, X., Hong, M., Dhople, S., Yin, W., and Liu, Y. Fedpd: A federated learning framework with optimal rates and adaptivity to non-iid data. *arXiv preprint arXiv:2005.11418*, 2020a.
- Zhang, X., Liu, J., and Zhu, Z. Taming convergence for asynchronous stochastic gradient descent with unbounded delay in non-convex learning. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pp. 3580–3585. IEEE, 2020b.

Appendix

In this supplementary material, we provide the detailed proofs for all theoretical results in this paper. Before presenting the proofs, we introduce some notations that will be used subsequently. We assume there exists M workers in total in the FL systems. In each communication round, we assume a subset \mathcal{M}_t of workers to be used, with $|\mathcal{M}_t| = m$. We use $\mathbf{G}_i(\mathbf{x}_t)$ to represent the local update returned from worker $i, i \in [M]$ given global model parameter $\mathbf{x}_t^0 = \mathbf{x}_t$. Also, we define $\mathbf{G}_i(\mathbf{x}_t) \triangleq \frac{1}{K_{t,i}} \sum_{j=0}^{K_{t,i}-1} \nabla f_i(\mathbf{x}_t^j, \xi_{t,i}^j)$, where \mathbf{x}_t^j represents the trajectory of the local model in the worker. We use Δ_i to denote the average of the full gradients long the trajectory of local updates, i.e., $\Delta_i(\mathbf{x}_t) = \frac{1}{K_{t,i}} \sum_{j=0}^{K_{t,i}-1} \nabla f_i(\mathbf{x}_t^j)$. With the above notations, we are now in a position to present the proofs of the theoretical results in this paper.

A. Proofs of Lemma 1 and Lemma 2

We start with proving two results stated in the following two lemmas, which will be useful in the rest of the proofs.

Lemma 1. $\mathbb{E}[\mathbf{G}_i(\mathbf{x}_t)] = \Delta_i(\mathbf{x}_t)$, $\mathbb{E}[\|\mathbf{G}_i(\mathbf{x}_t) - \Delta_i(\mathbf{x}_t)\|^2] \leq \frac{1}{K_{t,i}} \sigma_L^2, \forall i \in [M]$.

Proof. Taking the expectation of $\mathbf{G}_i(\mathbf{x}_t)$, we have:

$$\begin{aligned} \mathbb{E}[\mathbf{G}_i(\mathbf{x}_t)] &= \mathbb{E}\left[\frac{1}{K_{t,i}} \sum_{j=0}^{K_{t,i}-1} \nabla f_i(\mathbf{x}_t^j, \xi_{t,i}^j)\right] \\ &= \frac{1}{K_{t,i}} \sum_{j=0}^{K_{t,i}-1} \mathbb{E}[\nabla f_i(\mathbf{x}_t^j, \xi_{t,i}^j)] \\ &= \Delta_i(\mathbf{x}_t). \end{aligned}$$

Also, by computing the mean square error between $\mathbf{G}_i(\mathbf{x}_t)$ and $\Delta_i(\mathbf{x}_t)$, we have:

$$\begin{aligned} \mathbb{E}[\|\mathbf{G}_i(\mathbf{x}_t) - \Delta_i(\mathbf{x}_t)\|^2] &= \mathbb{E}\left[\left\|\frac{1}{K_{t,i}} \sum_{j=0}^{K_{t,i}-1} \nabla f_i(\mathbf{x}_t^j, \xi_{t,i}^j) - \sum_{j=0}^{K_{t,i}-1} \nabla f_i(\mathbf{x}_t^j)\right\|^2\right] \\ &= \frac{1}{K_{t,i}^2} \mathbb{E}\left[\left\|\sum_{j=0}^{K_{t,i}-1} \nabla f_i(\mathbf{x}_t^j, \xi_{t,i}^j) - \sum_{j=0}^{K_{t,i}-1} \nabla f_i(\mathbf{x}_t^j)\right\|^2\right] \\ &\leq \frac{1}{K_{t,i}} \sigma_L^2. \end{aligned}$$

Note $\{\nabla f_i(\mathbf{x}_t^j, \xi_{t,i}^j) - \nabla f_i(\mathbf{x}_t^j)\}$ forms a martingale difference sequence. This completes the proof of Lemma 1. \square

Lemma 2. For a fixed set \mathcal{M}_t with cardinality m , $\mathbb{E}\left[\left\|\sum_{i \in \mathcal{M}_t} \mathbf{G}_i(\mathbf{x}_{t-\tau_{t,i}})\right\|^2\right] \leq 2\left\|\sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\right\|^2 + \frac{2m}{K_t} \sigma_L^2$, where $\frac{1}{K_t} = \frac{1}{m} \sum_{i \in \mathcal{M}_t} \frac{1}{K_{t,i}}$.

Proof. By adding and subtracting $\Delta_i(\mathbf{x}_{t-\tau_{t,i}})$, we have:

$$\begin{aligned} \mathbb{E}\left[\left\|\sum_{i \in \mathcal{M}_t} \mathbf{G}_i(\mathbf{x}_{t-\tau_{t,i}})\right\|^2\right] &= 2\mathbb{E}\left[\left\|\sum_{i \in \mathcal{M}_t} \mathbf{G}_i(\mathbf{x}_{t-\tau_{t,i}}) - \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\right\|^2\right] + 2\left\|\sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\right\|^2 \\ &\leq 2 \sum_{i \in \mathcal{M}_t} \frac{1}{K_{t,i}} \sigma_L^2 + 2\left\|\sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\right\|^2 \\ &\leq \frac{2m}{K_t} \sigma_L^2 + 2\left\|\sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\right\|^2. \end{aligned}$$

Here the updates among clients $\{\mathbf{G}_i(\mathbf{x}_{t-\tau_{t,i}}) - \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\}$ are assumed to be independent. \square

B. Proof of the performance of the AFA-CD algorithm

In this section, we provide the proofs of the theoretical results of the AFA-CD algorithm. We consider two cases: i) general worker information arrival processes and ii) uniformly distributed worker information arrivals. As mentioned earlier, for general worker information arrival processes, we do not make any assumptions on the worker information arrival processes except the independence of workers' participation. For uniformly distributed worker information arrivals, \mathcal{M}_t can be viewed as a subset with size m independently and uniformly sampled from $[M]$ without replacement. The similar convergence analysis for independently and uniformly sampling with replacement can be derived in the same way following the techniques in (Yang et al., 2021; Li et al., 2019c).

B.1. Lower Bound for General Worker Information Arrival Processes

Theorem 1 (Convergence Error Lower Bound for AFL with General Worker Information Arrival Processes). *For any level of heterogeneity characterized by σ_G , there exists loss functions satisfying Assumptions 1-3 and a specific worker participation process for which the output $\hat{\mathbf{x}}$ of any convergent (and potentially random) FL algorithm satisfies:*

$$\mathbb{E}[\|\nabla f(\hat{\mathbf{x}})\|^2] = \Omega(\sigma_G^2).$$

Proof. We prove the lower bound by considering a worst-case scenario for simple one-dimensional functions. Let the FL system has two workers with the following loss functions: $f_1(\mathbf{x}) = (\mathbf{x} + G)^2$, $f_2(\mathbf{x}) = (\mathbf{x} - G)^2$, $f(\mathbf{x}) = \frac{1}{2}(f_1(\mathbf{x}) + f_2(\mathbf{x})) = \mathbf{x}^2 + G^2$. It is easy to verify that $\|\nabla f_1(\mathbf{x}) - \nabla f(\mathbf{x})\|^2 \leq 4G^2 = \sigma_G^2$ and $\|\nabla f_2(\mathbf{x}) - \nabla f(\mathbf{x})\|^2 \leq 4G^2 = \sigma_G^2$, where σ_G is the heterogeneity index. We consider a special case for the general worker arrival process when only the first one worker participates in the training as the worst-case scenario, equivalent to optimizing $f_1(\mathbf{x})$ rather than $f(\mathbf{x})$. In such case, any convergent (and potentially random) algorithm would return $\mathbb{E}\hat{\mathbf{x}} = -G$. As a result, $\mathbb{E}\|\nabla f(\hat{\mathbf{x}})\|^2 = \Omega(\sigma_G^2)$. \square

B.2. General Worker Information Arrival Processes

Theorem 2 (AFA-CD with General Worker Information Arrival Processes). *Suppose that the resultant maximum delay under AFL is bounded, i.e., $\tau := \max_{t \in [T], i \in \mathcal{M}_t} \{\tau_{t,i}\} < \infty$. Suppose that the server-side and worker-side learning rates η and η_L are chosen as such that the following conditions are satisfied: $6\eta_L^2(2K_{t,i}^2 - 3K_{t,i} + 1)L^2 \leq 1$, $180\eta_L^2 K_{t,i}^2 L^2 \tau < 1, \forall t, i$ and $2L\eta\eta_L + 6\tau^2 L^2 \eta^2 \eta_L^2 \leq 1$. Under Assumptions 1-3, the output sequence $\{\mathbf{x}_t\}$ generated by AFA-CD with general worker information arrival processes satisfies:*

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}\|\nabla f(\mathbf{x}_t)\|^2 \leq \frac{4(f_0 - f_*)}{\eta\eta_L T} + 4(\alpha_L \sigma_L^2 + \alpha_G \sigma_G^2),$$

where the constants α_L and α_G are defined as:

$$\begin{aligned} \alpha_L &= \frac{L\eta\eta_L}{m} \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{K_t} + \frac{3\tau^2 L^2 \eta^2 \eta_L^2}{m} \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{K_t} \\ &\quad + \frac{15\eta_L^2 L^2}{2} \frac{1}{T} \sum_{t=0}^{T-1} \bar{K}_t, \\ \alpha_G &= \frac{3}{2} + 45L^2 \eta_L^2 \frac{1}{T} \sum_{t=0}^{T-1} \hat{K}_t^2. \end{aligned}$$

Here,

$$\frac{1}{K_t} = \frac{1}{m} \sum_{i \in \mathcal{M}_t} \frac{1}{K_{t,i}}, \bar{K}_t = \frac{1}{m} \sum_{i \in \mathcal{M}_t} K_{t,i}, \hat{K}_t^2 = \frac{1}{m} \sum_{i \in \mathcal{M}_t} K_{t,i}^2.$$

Proof. Due to the L -smoothness assumption, taking expectation of $f(\mathbf{x}_{t+1})$ over the randomness in communication round t , we have:

$$\mathbb{E}[f(\mathbf{x}_{t+1})] \leq f(\mathbf{x}_t) + \underbrace{\langle \nabla f(\mathbf{x}_t), \mathbb{E}[\mathbf{x}_{t+1} - \mathbf{x}_t] \rangle}_{A_1} + \frac{L}{2} \underbrace{\mathbb{E}[\|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2]}_{A_2}.$$

First, we bound the term A_2 as follows:

$$\begin{aligned}
 A_2 &= \mathbb{E} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2 \\
 &= \eta^2 \eta_L^2 \mathbb{E} \left\| \frac{1}{m} \sum_{i \in \mathcal{M}_t} G_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \\
 &\stackrel{(a1)}{\leq} \frac{2\eta^2 \eta_L^2}{m^2} \left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 + \frac{2\eta^2 \eta_L^2}{m K_t} \sigma_L^2,
 \end{aligned}$$

where (a1) is due to Lemma 2. Next, we bound the term A_1 as follows:

$$\begin{aligned}
 A_1 &= \langle \nabla f(\mathbf{x}_t), \mathbb{E}[\mathbf{x}_{t+1} - \mathbf{x}_t] \rangle \\
 &= -\eta \eta_L \langle \nabla f(\mathbf{x}_t), \mathbb{E} \left[\frac{1}{m} \sum_{i \in \mathcal{M}_t} \mathbf{G}_i(\mathbf{x}_{t-\tau_{t,i}}) \right] \rangle \\
 &\stackrel{(a2)}{=} -\frac{1}{2} \eta \eta_L \|\nabla f(\mathbf{x}_t)\|^2 - \frac{1}{2} \eta \eta_L \mathbb{E} \left\| \frac{1}{m} \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 + \underbrace{\frac{1}{2} \eta \eta_L \mathbb{E} \left\| \nabla f(\mathbf{x}_t) - \frac{1}{m} \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2}_{A_3},
 \end{aligned}$$

where (a2) is due to Lemma 1 and the fact that $\langle \mathbf{x}, \mathbf{y} \rangle = \frac{1}{2} (\|\mathbf{x}\|^2 + \|\mathbf{y}\|^2 - \|\mathbf{x} - \mathbf{y}\|^2)$ and Lemma 1. To further bound the term A_3 , we have:

$$\begin{aligned}
 A_3 &= \mathbb{E} \left\| \nabla f(\mathbf{x}_t) - \frac{1}{m} \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \\
 &\leq \frac{1}{m} \sum_{i \in \mathcal{M}_t} \mathbb{E} \|\nabla f(\mathbf{x}_t) - \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\|^2 \\
 &= \frac{1}{m} \sum_{i \in \mathcal{M}_t} \mathbb{E} \|\nabla f(\mathbf{x}_t) - \nabla f(\mathbf{x}_{t-\tau_{t,i}}) + \nabla f(\mathbf{x}_{t-\tau_{t,i}}) - \nabla f_i(\mathbf{x}_{t-\tau_{t,i}}) + \nabla f_i(\mathbf{x}_{t-\tau_{t,i}}) - \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\|^2 \\
 &\stackrel{(a3)}{\leq} \frac{1}{m} \sum_{i \in \mathcal{M}_t} \left[3\mathbb{E} \|\nabla f(\mathbf{x}_t) - \nabla f(\mathbf{x}_{t-\tau_{t,i}})\|^2 + 3\mathbb{E} \|\nabla f(\mathbf{x}_{t-\tau_{t,i}}) - \nabla f_i(\mathbf{x}_{t-\tau_{t,i}})\|^2 \right. \\
 &\quad \left. + 3\mathbb{E} \|\nabla f_i(\mathbf{x}_{t-\tau_{t,i}}) - \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\|^2 \right] \\
 &\stackrel{(a4)}{\leq} \underbrace{\frac{3L^2}{m} \sum_{i \in \mathcal{M}_t} \mathbb{E} \|\mathbf{x}_t - \mathbf{x}_{t-\tau_{t,i}}\|^2}_{A_4} + 3\sigma_G^2 + \frac{3}{m} \sum_{i \in \mathcal{M}_t} \underbrace{\mathbb{E} \|\nabla f_i(\mathbf{x}_{t-\tau_{t,i}}) - \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\|^2}_{A_5},
 \end{aligned}$$

where (a3) follows from the inequality $\|\mathbf{x}_1 + \mathbf{x}_2 + \dots + \mathbf{x}_n\|^2 \leq n \sum_{i=1}^n \|\mathbf{x}_i\|^2$, and (a4) is due to the L-smoothness assumption (Assumption 1) and bounded global variance assumption (Assumption 3).

To further bound the term A_4 , we have:

$$\begin{aligned}
 A_4 &= \frac{1}{m} \sum_{i \in [m]} \mathbb{E} \|\mathbf{x}_t - \mathbf{x}_{t-\tau_{t,i}}\|^2 \\
 &\stackrel{(a5)}{\leq} \mathbb{E} \|\mathbf{x}_t - \mathbf{x}_{t-\tau_{t,u}}\|^2 \\
 &= \mathbb{E} \left\| \sum_{k=t-\tau_{t,u}}^{t-1} \mathbf{x}_{k+1} - \mathbf{x}_k \right\|^2
 \end{aligned}$$

$$\begin{aligned}
 &= \mathbb{E} \left\| \sum_{k=t-\tau_{t,u}}^{t-1} \frac{1}{m} \eta \eta_L \sum_{i \in \mathcal{M}_k} \mathbf{G}_i(\mathbf{x}_{k-\tau_{k,i}}) \right\|^2 \\
 &= \mathbb{E} \left[\frac{\eta^2 \eta_L^2}{m^2} \left\| \sum_{k=t-\tau_{t,u}}^{t-1} \sum_{i \in \mathcal{M}_k} \mathbf{G}_i(\mathbf{x}_{k-\tau_{k,i}}) \right\|^2 \right] \\
 &\stackrel{(a6)}{\leq} \mathbb{E} \left[\frac{\eta^2 \eta_L^2}{m^2} \tau \sum_{k=t-\tau_{t,u}}^{t-1} \left\| \sum_{i \in \mathcal{M}_k} \mathbf{G}_i(\mathbf{x}_{k-\tau_{k,i}}) \right\|^2 \right] \\
 &\stackrel{(a7)}{\leq} \left[\frac{\eta^2 \eta_L^2 \tau}{m^2} \sum_{k=t-\tau_{t,u}}^{t-1} \left(2 \mathbb{E} \left\| \sum_{i \in \mathcal{M}_k} \Delta_i(\mathbf{x}_{k-\tau_{k,i}}) \right\|^2 + \frac{2m}{K_k} \sigma_L^2 \right) \right].
 \end{aligned}$$

In the derivations above, we let $u := \operatorname{argmax}_{i \in [M]} \|\mathbf{x}_t - \mathbf{x}_{t-\tau_{t,i}}\|^2$, which yields (a5). Note also that the maximum delay assumption $\tau \geq \tau_{k,i}, \forall i \in [M]$ implies (a6). Lastly, (a7) follows from Lemma 2.

To further bound the term A_5 , we have:

$$\begin{aligned}
 A_5 &= \mathbb{E} \|\nabla f_i(\mathbf{x}_{t-\tau_{t,i}}) - \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\|^2 \\
 &= \mathbb{E} \left\| \nabla f_i(\mathbf{x}_{t-\tau_{t,i}}) - \frac{1}{K_{t,i}} \sum_{j=0}^{K_{t,i}-1} \nabla f_i(\mathbf{x}_{t-\tau_{t,i}}^j) \right\|^2 \\
 &= \frac{1}{K_{t,i}} \sum_{j=0}^{K_{t,i}-1} \mathbb{E} \left\| \nabla f_i(\mathbf{x}_{t-\tau_{t,i}}) - \nabla f_i(\mathbf{x}_{t-\tau_{t,i}}^j) \right\|^2 \\
 &\stackrel{(a8)}{\leq} \frac{L^2}{K_{t,i}} \sum_{j=0}^{K_{t,i}-1} \underbrace{\mathbb{E} \|\mathbf{x}_{t-\tau_{t,i}} - \mathbf{x}_{t-\tau_{t,i}}^j\|^2}_{A_6} \\
 &\stackrel{(a9)}{\leq} 5K_{t,i} L^2 \eta_L^2 (\sigma_L^2 + 6K_{t,i} \sigma_G^2) + 30K_{t,i}^2 L^2 \eta_L^2 \|\nabla f(\mathbf{x}_{t-\tau_{t,i}})\|^2,
 \end{aligned}$$

where (a8) is due to the L -smoothness assumption (Assumption 1), and (a9) follows from the bound of A_6 shown below. Here, we denote maximum number of local steps of all workers as K , i.e., $K_{t,i} \leq K, \forall t, i$. This definition of K implies (a10).

Now, it remains to bound term A_6 in the derivations above. Note that the bounding proof of A_6 in what follows is the same as Lemma 4 in (Reddi et al., 2020). we restate the proof here in order for this paper to be self-contained. For any worker i in the k -th local step, we have the following results for the norm of parameter changes for one local computation:

$$\begin{aligned}
 A_6 &= \mathbb{E} [\|\mathbf{x}_{t,k}^i - \mathbf{x}_t\|^2] = \mathbb{E} [\|\mathbf{x}_{t,k-1}^i - \mathbf{x}_t - \eta_L g_{t,k-1}^i\|^2] \\
 &\leq \mathbb{E} [\|\mathbf{x}_{t,k-1}^i - \mathbf{x}_t - \eta_L (g_{t,k-1}^i - \nabla f_i(\mathbf{x}_{t,k-1}^i) + \nabla f_i(\mathbf{x}_{t,k-1}^i) - \nabla f_i(\mathbf{x}_t) \\
 &\quad + \nabla f_i(\mathbf{x}_t) - \nabla f(\mathbf{x}_t) + \nabla f(\mathbf{x}_t))\|^2] \\
 &\leq (1 + \frac{1}{2K_{t,i}-1}) \mathbb{E} [\|\mathbf{x}_{t,k-1}^i - \mathbf{x}_t\|^2] + \mathbb{E} [\|\eta_L (g_{t,k-1}^i - \nabla f_i(\mathbf{x}_{t,k-1}^i))\|^2] \\
 &\quad + 6K_{t,i} \mathbb{E} [\|\eta_L (\nabla f_i(\mathbf{x}_{t,k-1}^i) - \nabla f_i(\mathbf{x}_t))\|^2] + 6K_{t,i} \mathbb{E} [\|\eta_L (\nabla f_i(\mathbf{x}_t) - \nabla f(\mathbf{x}_t))\|^2] \\
 &\quad + 6K_{t,i} \|\eta_L \nabla f(\mathbf{x}_t)\|^2 \\
 &\leq (1 + \frac{1}{2K_{t,i}-1}) \mathbb{E} [\|\mathbf{x}_{t,k-1}^i - \mathbf{x}_t\|^2] + \eta_L^2 \sigma_L^2 + 6K_{t,i} \eta_L^2 L^2 \mathbb{E} [\|\mathbf{x}_{t,k-1}^i - \mathbf{x}_t\|^2] \\
 &\quad + 6K_{t,i} \eta_L^2 \sigma_G^2 + 6K_{t,i} \|\eta_L \nabla f(\mathbf{x}_t)\|^2 \\
 &= (1 + \frac{1}{2K_{t,i}-1} + 6K_{t,i} \eta_L^2 L^2) \mathbb{E} [\|\mathbf{x}_{t,k-1}^i - \mathbf{x}_t\|^2] + \eta_L^2 \sigma_L^2 + 6K_{t,i} \eta_L^2 \sigma_G^2 + 6K_{t,i} \|\eta_L \nabla f(\mathbf{x}_t)\|^2
 \end{aligned}$$

$$\stackrel{(a11)}{\leq} \left(1 + \frac{1}{K_{t,i} - 1}\right) \mathbb{E}[\|\mathbf{x}_{t,k-1}^i - \mathbf{x}_t\|^2] + \eta_L^2 \sigma_L^2 + 6K_{t,i} \eta_L^2 \sigma_G^2 + 6K_{t,i} \|\eta_L \nabla f(\mathbf{x}_t)\|^2,$$

where (a11) follows from the fact that $\frac{1}{2K_{t,i} - 1} + 6K_{t,i} \eta_L^2 L^2 \leq \frac{1}{K_{t,i} - 1}$ if $\eta_L^2 \leq \frac{1}{6(2K_{t,i}^2 - 3K_{t,i} + 1)L^2}$.

Unrolling the recursion, we obtain:

$$\begin{aligned} \mathbb{E}[\|\mathbf{x}_{t,k}^i - \mathbf{x}_t\|^2] &\leq \sum_{p=0}^{k-1} \left(1 + \frac{1}{K_{t,i} - 1}\right)^p [\eta_L^2 \sigma_L^2 + 6K_{t,i} \sigma_G^2 + 6K_{t,i} \eta_L^2 \|\eta_L \nabla f(\mathbf{x}_t)\|^2] \\ &\leq (K_{t,i} - 1) \left[\left(1 + \frac{1}{K_{t,i} - 1}\right)^{K_{t,i}} - 1\right] [\eta_L^2 \sigma_L^2 + 6K_{t,i} \eta_L^2 \sigma_G^2 + 6K_{t,i} \|\eta_L \nabla f(\mathbf{x}_t)\|^2] \\ &\leq 5K_{t,i} \eta_L^2 (\sigma_L^2 + 6K_{t,i} \sigma_G^2) + 30K_{t,i}^2 \eta_L^2 \|\nabla f(\mathbf{x}_t)\|^2. \end{aligned} \quad (1)$$

With the above results of the terms A_1 through A_5 , we have:

$$\begin{aligned} \mathbb{E}[f(\mathbf{x}_{t+1})] - f(\mathbf{x}_t) &\leq \underbrace{\langle \nabla f(\mathbf{x}_t), \mathbb{E}[\mathbf{x}_{t+1} - \mathbf{x}_t] \rangle}_{A_1} + \frac{L}{2} \underbrace{\mathbb{E}[\|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2]}_{A_2} \\ &= -\frac{1}{2} \eta_L \|\nabla f(\mathbf{x}_t)\|^2 - \frac{1}{2} \eta_L \mathbb{E} \left\| \frac{1}{m} \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_{t-\tau_t, i}) \right\|^2 + \frac{1}{2} \eta_L \underbrace{\mathbb{E} \left\| \nabla f(\mathbf{x}_t) - \frac{1}{m} \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_{t-\tau_t, i}) \right\|^2}_{A_3} \\ &\quad + \frac{L \eta^2 \eta_L^2}{m^2} \mathbb{E} \left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_{t-\tau_t, i}) \right\|^2 + \frac{L \eta^2 \eta_L^2}{m} \sigma_L^2 \\ &\leq -\frac{1}{2} \eta_L \|\nabla f(\mathbf{x}_t)\|^2 - \frac{1}{2} \eta_L \mathbb{E} \left\| \frac{1}{m} \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_{t-\tau_t, i}) \right\|^2 + \frac{L \eta^2 \eta_L^2}{m^2} \mathbb{E} \left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_{t-\tau_t, i}) \right\|^2 + \frac{L \eta^2 \eta_L^2}{m K_t} \sigma_L^2 \\ &\quad + \frac{3}{2} \eta_L \sigma_G^2 + \frac{3L^2}{2} \eta_L \underbrace{\left[\frac{1}{m} \sum_{i \in \mathcal{M}_t} \mathbb{E} \|\mathbf{x}_t - \mathbf{x}_{t-\tau_t, i}\|^2 \right]}_{A_4} + \frac{3\eta_L}{2m} \sum_{i \in \mathcal{M}_t} \underbrace{\mathbb{E} \|\nabla f_i(\mathbf{x}_{t-\tau_t, i}) - \Delta_i(\mathbf{x}_{t-\tau_t, i})\|^2}_{A_5} \\ &\leq -\frac{1}{2} \eta_L \|\nabla f(\mathbf{x}_t)\|^2 - \frac{\eta_L}{2m^2} \mathbb{E} \left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_{t-\tau_t, i}) \right\|^2 + \frac{L \eta^2 \eta_L^2}{m^2} \mathbb{E} \left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_{t-\tau_t, i}) \right\|^2 + \frac{L \eta^2 \eta_L^2}{m K_t} \sigma_L^2 \\ &\quad + \frac{3}{2} \eta_L \sigma_G^2 + \frac{3L^2}{2} \eta_L \left[\frac{\eta^2 \eta_L^2 \tau}{m^2} \sum_{k=t-\tau_t, u}^{t-1} \left(2 \mathbb{E} \left\| \sum_{i \in \mathcal{M}_k} \Delta_i(\mathbf{x}_{k-\tau_k, i}) \right\|^2 + \frac{2m}{K_k} \sigma_L^2 \right) \right] \\ &\quad + \frac{3\eta_L}{2} \frac{1}{m} \sum_{i \in \mathcal{M}_t} [5K_{t,i} L^2 \eta_L^2 (\sigma_L^2 + 6K_{t,i} \sigma_G^2) + 30L^2 \eta_L^2 K_{t,i}^2 \|\nabla f(\mathbf{x}_{t-\tau_t, i})\|^2] \\ &\leq -\frac{1}{2} \eta_L \|\nabla f(\mathbf{x}_t)\|^2 + 45\eta_L^3 L^2 \frac{1}{m} \sum_{i=1}^m K_{t,i}^2 \|\nabla f(\mathbf{x}_{t-\tau_t, i})\|^2 \\ &\quad + \left[-\frac{\eta_L}{2m^2} + \frac{L \eta^2 \eta_L^2}{m^2} \right] \mathbb{E} \left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_{t-\tau_t, i}) \right\|^2 + \frac{3\tau \eta^3 \eta_L^3}{m^2} \sum_{k=t-\tau_t, u}^{t-1} \mathbb{E} \left\| \sum_{i \in \mathcal{M}_k} \Delta_i(\mathbf{x}_{k-\tau_k, i}) \right\|^2 \\ &\quad + \left[\frac{L \eta^2 \eta_L^2}{m K_t} + \frac{3\tau L^2 \eta^3 \eta_L^3}{m} \sum_{k=t-\tau_t, u}^{t-1} \frac{1}{K_k} + \frac{15\eta_L^3 L^2}{2} \frac{1}{m} \sum_{i \in \mathcal{M}_t} K_{t,i} \right] \sigma_L^2 + \left[\frac{3}{2} \eta_L + 45L^2 \eta_L^3 \frac{1}{m} \sum_{i \in \mathcal{M}_t} K_{t,i}^2 \right] \sigma_G^2. \end{aligned}$$

Summing the above inequality from $t = 0$ to $t = T - 1$ yields:

$$\mathbb{E}f(\mathbf{x}_T) - f(\mathbf{x}_0)$$

$$\begin{aligned}
 &\leq \sum_{t=0}^{T-1} \left[-\frac{1}{2}\eta\eta_L \|\nabla f(\mathbf{x}_t)\|^2 + 45\eta\eta_L^3 L^2 \frac{1}{m} \sum_{i=1}^m K_{t,i}^2 \mathbb{E} \|\nabla f(\mathbf{x}_{t-\tau_{t,i}})\|^2 \right] \\
 &\quad + \sum_{t=0}^{T-1} \left[\left[-\frac{\eta\eta_L}{2m^2} + \frac{L\eta^2\eta_L^2}{m^2} \right] \mathbb{E} \left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 + \frac{3\tau L^2\eta^3\eta_L^3}{m^2} \sum_{k=t-\tau_{t,u}}^{t-1} \mathbb{E} \left\| \sum_{i \in \mathcal{M}_k} \Delta_i(\mathbf{x}_{k-\tau_{k,i}}) \right\|^2 \right] \\
 &\quad + \sum_{t=0}^{T-1} \left[\frac{L\eta^2\eta_L^2}{mK_t} + \frac{3\tau L^2\eta^3\eta_L^3}{m} \sum_{k=t-\tau_{t,u}}^{t-1} \frac{1}{K_k} + \frac{15\eta\eta_L^3 L^2 \frac{1}{m} \sum_{i \in \mathcal{M}_t} K_{t,i}}{2} \right] \sigma_L^2 \\
 &\quad + \sum_{t=0}^{T-1} \left[\frac{3}{2}\eta\eta_L + 45L^2\eta\eta_L^3 \frac{1}{m} \sum_{i \in \mathcal{M}_t} K_{t,i}^2 \right] \sigma_G^2 \\
 &\stackrel{(a12)}{\leq} \sum_{t=0}^{T-1} \left[-\frac{1}{2}\eta\eta_L + 45\eta\eta_L^3 K_{t,max}^2 L^2 \tau \right] \|\nabla f(\mathbf{x}_t)\|^2 \\
 &\quad + \sum_{t=0}^{T-1} \left[-\frac{\eta\eta_L}{2m^2} + \frac{L\eta^2\eta_L^2}{m^2} + \frac{3\tau^2 L^2\eta^3\eta_L^3}{2m^2} \right] \mathbb{E} \left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \\
 &\quad + \sum_{t=0}^{T-1} \left[\frac{L\eta^2\eta_L^2}{mK_t} + \frac{3\tau^2 L^2\eta^3\eta_L^3}{m} \frac{1}{K_t} + \frac{15\eta\eta_L^3 \bar{K}_t L^2}{2} \right] \sigma_L^2 + \sum_{t=0}^{T-1} \left[\frac{3}{2}\eta\eta_L + 45\hat{K}_t^2 L^2\eta\eta_L^3 \right] \sigma_G^2 \\
 &\stackrel{(a13)}{\leq} \sum_{t=0}^{T-1} -\frac{1}{4}\eta\eta_L \|\nabla f(\mathbf{x}_t)\|^2 \\
 &\quad + \eta\eta_L \left[\frac{L\eta\eta_L}{m} \sum_{t=0}^{T-1} \frac{1}{K_t} + \frac{3\tau^2 L^2\eta^2\eta_L^2}{m} \sum_{t=0}^{T-1} \frac{1}{K_t} + \frac{15\eta_L^2 L^2}{2} \sum_{t=0}^{T-1} \bar{K}_t \right] \sigma_L^2 + \sum_{t=0}^{T-1} \eta\eta_L \left[\frac{3}{2} + 45\hat{K}_t^2 L^2\eta_L^2 \right] \sigma_G^2 \\
 &\stackrel{(a14)}{=} \sum_{t=0}^{T-1} -\frac{1}{4}\eta\eta_L \|\nabla f(\mathbf{x}_t)\|^2 + T\eta\eta_L [\alpha_L \sigma_L^2 + \alpha_G \sigma_G^2],
 \end{aligned}$$

where (a12) is due to maximum time delay τ in the system, (a13) holds if $\frac{1}{4} \leq [\frac{1}{2} - 45\eta_L^2 K_{t,max}^2 L^2 \tau]$, i.e., $180\eta_L^2 K_{t,max}^2 L^2 \tau < 1$, and $\left[-\frac{\eta\eta_L}{2m^2} + \frac{L\eta^2\eta_L^2}{m^2} + \frac{3L^2\tau^2\eta^3\eta_L^3}{m^2} \right] \leq 0$, i.e., $2L\eta\eta_L + 6\tau^2 L^2\eta^2\eta_L^2 \leq 1$. Note $\bar{K}_t = \frac{1}{m} \sum_{i \in \mathcal{M}_t} K_{t,i}$, $\hat{K}_t^2 = \frac{1}{m} \sum_{i \in \mathcal{M}_t} K_{t,i}^2$, and $K_{t,max} = \max\{K_{t,i}, i \in [m]\}$. Lastly, (a14) follows from the following definitions:

$$\begin{aligned}
 \alpha_L &= \left[\frac{L\eta\eta_L}{m} \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{K_t} + \frac{3\tau^2 L^2\eta^2\eta_L^2}{m} \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{K_t} + \frac{15\eta_L^2 L^2}{2} \frac{1}{T} \sum_{t=0}^{T-1} \bar{K}_t \right], \\
 \alpha_G &= \left[\frac{3}{2} + 45L^2\eta_L^2 \frac{1}{T} \sum_{t=0}^{T-1} \hat{K}_t^2 \right].
 \end{aligned}$$

Rearranging terms, we have:

$$\frac{1}{T} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|^2 \leq \frac{4(f_0 - f_*)}{\eta\eta_L T} + 4[\alpha_L \sigma_L^2 + \alpha_G \sigma_G^2],$$

and the proof is complete. \square

Corollary 1 (Linear Speedup to an Error Ball). *Suppose a constant local step K for each worker, by setting $\eta_L = \frac{1}{\sqrt{T}}$, and $\eta = \sqrt{m\bar{K}}$, the convergence rate of AFA-CD with general worker information arrival processes is:*

$$\mathcal{O}\left(\frac{1}{m^{1/2} K^{1/2} T^{1/2}}\right) + \mathcal{O}\left(\frac{\tau^2}{T}\right) + \mathcal{O}\left(\frac{K^2}{T}\right) + \mathcal{O}(\sigma_G^2).$$

Proof. Suppose a constant local step K for each worker, and let $\eta_L = \frac{1}{\sqrt{T}}$, and $\eta = \sqrt{mK}$. It then follows that:

$$\alpha_L = \mathcal{O}\left(\frac{1}{m^{1/2}K^{1/2}T^{1/2}}\right) + \mathcal{O}\left(\frac{\tau^2}{T}\right) + \mathcal{O}\left(\frac{K}{T}\right).$$

$$\alpha_G = \mathcal{O}(\sigma_G^2) + \mathcal{O}\left(\frac{K^2}{T}\right).$$

This completes the proof. \square

B.3. Uniformly Distributed Worker Information Arrivals

Now, we consider the special case that the worker information arrivals are uniformly distributed, i.e., the worker in \mathcal{M}_t could be regarded as a uniformly random sample without replacement in $[M]$. As mentioned earlier, this special case acts as a widely-used assumption in FL and could deepen our understanding on the AFA-CD algorithm's performance in large-scale AFL systems.

Theorem 3. *Under the same delay condition in Theorem 2 and suppose that the server-side and worker-side learning rates η and η_L are chosen as such that the following relationships hold: $6\eta_L^2(2K_{t,i}^2 - 3K_{t,i} + 1)L^2 \leq 1, \forall t, i$, $L\eta\eta_L + L^2\eta^2\eta_L^2\tau^2 \leq \frac{1}{2M}$, and $120L^2\hat{K}_t^2\eta_L^2\tau < 1, \forall t$. Then, under Assumptions 1–3, the output sequence $\{\mathbf{x}_t\}$ generated by AFA-CD with uniformly distributed worker information arrivals satisfies:*

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\mathbf{x}_t)\|_2^2 \leq \frac{4(f_0 - f_*)}{\eta\eta_L T} + 4(\alpha_L\sigma_L^2 + \alpha_G\sigma_G^2),$$

where α_L and α_G are defined as following:

$$\begin{aligned} \alpha_L &= \frac{L\eta\eta_L}{m} \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{K_t} + \frac{2\tau^2 L^2 \eta^2 \eta_L^2}{m} \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{K_t} \\ &\quad + 5\eta_L^2 L^2 \frac{1}{T} \sum_{t=0}^{T-1} \bar{K}_t, \\ \alpha_G &= 30L^2 \eta_L^2 \frac{1}{T} \sum_{t=0}^{T-1} \hat{K}_t^2, \end{aligned}$$

and other parameters are defined the same as in Theorem 2.

Proof. The one-step update can be rewritten as: $\mathbf{x}_{t+1} - \mathbf{x}_t = -\eta\eta_L \mathbf{G}_t$. For cross-device FL, $\mathbf{G}_t = \frac{1}{m} \sum_{i \in \mathcal{M}_t} \mathbf{G}_i(\mathbf{x}_{t-\tau_{t,i}})$, where $\tau_{t,i}$ is the delay for client i in terms of the current global communication round t . When $\tau_{t,i} = 0, \forall i \in \mathcal{M}_t$, it degenerates to synchronous FL with partial worker participation.

Due to the L -smoothness in Assumption 1, taking expectation of $f(\mathbf{x}_{t+1})$ over the randomness in communication round t , we have:

$$\mathbb{E}[f(\mathbf{x}_{t+1})] \leq f(\mathbf{x}_t) + \underbrace{\langle \nabla f(\mathbf{x}_t), \mathbb{E}[\mathbf{x}_{t+1} - \mathbf{x}_t] \rangle}_{A_1} + \frac{L}{2} \underbrace{\mathbb{E}[\|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2]}_{A_2}$$

We first bound A_2 as follows:

$$\begin{aligned} A_2 &= \mathbb{E} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2 \\ &= \eta^2 \eta_L^2 \mathbb{E} \left\| \frac{1}{m} \sum_{i \in \mathcal{M}_t} \mathbf{G}_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \end{aligned}$$

$$\begin{aligned}
 &\stackrel{(b1)}{\leq} \frac{2\eta^2\eta_L^2}{m^2} \mathbb{E} \left[\left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 + \frac{m}{K_t} \sigma_L^2 \right] \\
 &\stackrel{(b2)}{\leq} \frac{2\eta^2\eta_L^2}{m^2} \mathbb{E} \left\| \sum_{i=1}^M \mathbb{I}\{i \in \mathcal{M}_t\} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 + \frac{2\eta^2\eta_L^2}{mK_t} \sigma_L^2,
 \end{aligned}$$

where (b1) is due to Lemma 2 and (b2) is due to the uniformly independent information arrival assumption.

To bound the term A_1 , we have:

$$\begin{aligned}
 A_1 &= \langle \nabla f(\mathbf{x}_t), \mathbb{E}[\mathbf{x}_{t+1} - \mathbf{x}_t] \rangle \\
 &= -\eta\eta_L \left\langle \nabla f(\mathbf{x}_t), \mathbb{E} \left[\frac{1}{m} \sum_{i \in \mathcal{M}_t} \mathbf{G}_i(\mathbf{x}_{t-\tau_{t,i}}) \right] \right\rangle \\
 &\stackrel{(b3)}{=} -\eta\eta_L \left\langle \nabla f(\mathbf{x}_t), \frac{1}{M} \sum_{i \in [M]} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\rangle \\
 &\stackrel{(b4)}{=} -\frac{1}{2}\eta\eta_L \|\nabla f(\mathbf{x}_t)\|^2 - \frac{1}{2}\eta\eta_L \left\| \frac{1}{M} \sum_{i \in [M]} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \\
 &\quad + \underbrace{\frac{1}{2}\eta\eta_L \left\| \nabla f(\mathbf{x}_t) - \frac{1}{M} \sum_{i \in [M]} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2}_{A_3},
 \end{aligned}$$

where (b3) is due to the uniformly independent worker information arrival assumption and Lemma 1, (b4) is due to the fact that $\langle \mathbf{x}, \mathbf{y} \rangle = \frac{1}{2}(\|\mathbf{x}\|^2 + \|\mathbf{y}\|^2 - \|\mathbf{x} - \mathbf{y}\|^2)$.

To further bound the term A_3 , we have:

$$\begin{aligned}
 A_3 &= \left\| \nabla f(\mathbf{x}_t) - \frac{1}{M} \sum_{i \in [M]} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \\
 &\stackrel{(b5)}{=} \left\| \frac{1}{M} \sum_{i \in [M]} [\nabla f_i(\mathbf{x}_t) - \Delta_i(\mathbf{x}_{t-\tau_{t,i}})] \right\|^2 \\
 &\leq \frac{1}{M} \sum_{i \in [M]} \|\nabla f_i(\mathbf{x}_t) - \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\|^2 \\
 &= \frac{1}{M} \sum_{i \in [M]} \|\nabla f_i(\mathbf{x}_t) - \nabla f_i(\mathbf{x}_{t-\tau_{t,i}}) + \nabla f_i(\mathbf{x}_{t-\tau_{t,i}}) - \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\|^2 \\
 &\stackrel{(b6)}{\leq} \frac{1}{M} \sum_{i \in [M]} \left[2\|\nabla f_i(\mathbf{x}_t) - \nabla f_i(\mathbf{x}_{t-\tau_{t,i}})\|^2 + 2\|\nabla f_i(\mathbf{x}_{t-\tau_{t,i}}) - \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\|^2 \right] \\
 &\stackrel{(b7)}{\leq} \underbrace{\frac{2L^2}{M} \sum_{i=1}^M \|\mathbf{x}_t - \mathbf{x}_{t-\tau_{t,i}}\|^2}_{A_4} + \frac{2}{M} \sum_{i=1}^M \underbrace{\|\nabla f_i(\mathbf{x}_{t-\tau_{t,i}}) - \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\|^2}_{A_5},
 \end{aligned}$$

where (b5) is due to the fact that $\nabla f(\mathbf{x}) = \frac{1}{M} \sum_{i \in [M]} \nabla f_i(\mathbf{x})$, (b6) follows from the inequality $\|\mathbf{x}_1 + \mathbf{x}_2 + \dots + \mathbf{x}_n\|^2 \leq n \sum_{i=1}^n \|\mathbf{x}_i\|^2$, and (b7) follows from the L -smoothness assumption (Assumption 1).

For A_4 and A_5 , we have the same bounds as in the case of general worker information arrival processes:

$$A_4 = \frac{2L^2}{M} \sum_{i=1}^M \|\mathbf{x}_t - \mathbf{x}_{t-\tau_{t,i}}\|^2$$

$$\begin{aligned}
 &\leq \mathbb{E} \left[\frac{4L^2\eta^2\eta_L^2\tau}{m^2} \sum_{k=t-\tau_{t,u}}^{t-1} \left(\left\| \sum_{i \in \mathcal{M}_k} \Delta_i(\mathbf{x}_{k-\tau_{k,i}}) \right\|^2 + \frac{m}{K_k} \sigma_L^2 \right) \right] \\
 &\leq \frac{4L^2\eta^2\eta_L^2\tau}{m^2} \sum_{k=t-\tau_{t,u}}^{t-1} \left(\left\| \sum_{i=1}^M \mathbb{I}\{i \in \mathcal{M}_k\} \Delta_i(\mathbf{x}_{k-\tau_{k,i}}) \right\|^2 + \frac{m}{K_k} \sigma_L^2 \right)
 \end{aligned}$$

$$\begin{aligned}
 A_5 &= \left\| \nabla f_i(\mathbf{x}_{t-\tau_{t,i}}) - \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \\
 &\leq 5K_{t,i}L^2\eta_L^2(\sigma_L^2 + 6K_{t,i}\sigma_G^2) + 30K_{t,i}^2L^2\eta_L^2\|\nabla f(\mathbf{x}_{t-\tau_{t,i}})\|^2,
 \end{aligned}$$

With the above results of the term A_1 through A_5 , we have:

$$\begin{aligned}
 \mathbb{E}_t[f(\mathbf{x}_{t+1})] - f(\mathbf{x}_t) &\leq \underbrace{\langle \nabla f(\mathbf{x}_t), \mathbb{E}_t[\mathbf{x}_{t+1} - \mathbf{x}_t] \rangle}_{A_1} + \frac{L}{2} \underbrace{\mathbb{E}_t[\|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2]}_{A_2} \\
 &= -\frac{1}{2}\eta\eta_L\|\nabla f(\mathbf{x}_t)\|^2 - \frac{1}{2}\eta\eta_L\left\| \frac{1}{M} \sum_{i \in [M]} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 + \frac{1}{2}\eta\eta_L \underbrace{\left\| \nabla f(\mathbf{x}_t) - \frac{1}{M} \sum_{i \in [M]} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2}_{A_3} \\
 &\quad + \frac{L\eta^2\eta_L^2}{m^2} \mathbb{E} \left\| \sum_{i=1}^M \mathbb{I}\{i \in \mathcal{M}_t\} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 + \frac{L\eta^2\eta_L^2}{mK_t} \sigma_L^2 \\
 &\leq -\frac{1}{2}\eta\eta_L\|\nabla f(\mathbf{x}_t)\|^2 - \frac{1}{2}\eta\eta_L \left\| \frac{1}{M} \sum_{i \in [M]} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 + \frac{L\eta^2\eta_L^2}{m^2} \mathbb{E} \left\| \sum_{i=1}^M \mathbb{I}\{i \in \mathcal{M}_t\} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \\
 &\quad + \frac{1}{2}\eta\eta_L \left[\underbrace{\frac{2L^2}{M} \sum_{i=1}^M \|\mathbf{x}_t - \mathbf{x}_{t-\tau_{t,i}}\|^2}_{A_4} + \frac{2}{M} \sum_{i=1}^M \underbrace{\|\nabla f_i(\mathbf{x}_{t-\tau_{t,i}}) - \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\|^2}_{A_5} \right] + \frac{L\eta^2\eta_L^2}{mK_t} \sigma_L^2 \\
 &\leq -\frac{1}{2}\eta\eta_L\|\nabla f(\mathbf{x}_t)\|^2 - \frac{1}{2}\eta\eta_L \left\| \frac{1}{M} \sum_{i \in [M]} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 + \frac{L\eta^2\eta_L^2}{m^2} \mathbb{E} \left\| \sum_{i=1}^M \mathbb{I}\{i \in \mathcal{M}_t\} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \\
 &\quad + \eta\eta_L L^2 \frac{2\eta^2\eta_L^2\tau}{m^2} \sum_{k=t-\tau_{t,u}}^{t-1} \left(\left\| \sum_{i=1}^M \mathbb{I}\{i \in \mathcal{M}_k\} \Delta_i(\mathbf{x}_{k-\tau_{k,i}}) \right\|^2 + \frac{m}{K_k} \sigma_L^2 \right) \\
 &\quad + \eta\eta_L \frac{1}{M} \sum_{i=1}^M [5K_{t,i}L^2\eta_L^2(\sigma_L^2 + 6K_{t,i}\sigma_G^2) + 30K_{t,i}^2L^2\eta_L^2\|\nabla f(\mathbf{x}_{t-\tau_{t,i}})\|^2] + \frac{L\eta^2\eta_L^2}{mK_t} \sigma_L^2 \\
 &\leq \left[-\frac{1}{2}\eta\eta_L\|\nabla f(\mathbf{x}_t)\|^2 + (30\eta L^2\eta_L^3) \frac{1}{M} \sum_{i=1}^M K_{t,i}^2\|\nabla f(\mathbf{x}_{t-\tau_{t,i}})\|^2 \right] \\
 &\quad + \left[-\frac{\eta\eta_L}{2M^2} \left\| \sum_{i=1}^M \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 + \frac{L\eta^2\eta_L^2}{m^2} \mathbb{E} \left\| \sum_{i=1}^M \mathbb{I}\{i \in \mathcal{M}_t\} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \right. \\
 &\quad \left. + \frac{2L^2\eta^3\eta_L^3\tau}{m^2} \sum_{k=t-\tau_{t,\mu}}^{t-1} \mathbb{E} \left\| \sum_{i=1}^M \mathbb{I}\{i \in \mathcal{M}_k\} \Delta_i(\mathbf{x}_{k-\tau_{k,i}}) \right\|^2 \right] \\
 &\quad + \sigma_L^2 \left[\frac{L\eta^2\eta_L^2}{mK_t} + \frac{2\tau L^2\eta^3\eta_L^3}{m} \sum_{k=t-\tau_{t,\mu}}^{t-1} \frac{1}{K_k} + 5\eta L^2\eta_L^3 \frac{1}{M} \sum_{i=1}^M K_{t,i} \right] + \left[30\eta L^2\eta_L^3 \frac{1}{M} \sum_{i=1}^M K_{t,i}^2 \right] \sigma_G^2.
 \end{aligned}$$

Summing the above inequality from $t = 0$ to $t = T - 1$ yields:

$$\begin{aligned}
 & \mathbb{E}f(\mathbf{x}_T) - f(\mathbf{x}_0) \\
 & \leq \sum_{t=0}^{T-1} \left[-\frac{1}{2}\eta\eta_L\|\nabla f(\mathbf{x}_t)\|^2 + (30\eta L^2\eta_L^3)\frac{1}{M}\sum_{i=1}^M K_{t,i}^2\|\nabla f(\mathbf{x}_{t-\tau_{t,i}})\|^2 \right] \\
 & \quad + \sum_{t=0}^{T-1} \left[-\frac{\eta\eta_L}{2M^2}\left\|\sum_{i=1}^M \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\right\|^2 + \frac{L\eta^2\eta_L^2}{m^2}\mathbb{E}\left\|\sum_{i=1}^M \mathbb{I}\{i \in \mathcal{M}_t\}\Delta_i(\mathbf{x}_{t-\tau_{t,i}})\right\|^2 \right. \\
 & \quad \left. + \frac{2L^2\eta^3\eta_L^3\tau}{m^2}\sum_{k=t-\tau_{t,\mu}}^{t-1}\mathbb{E}\left\|\sum_{i=1}^M \mathbb{I}\{i \in \mathcal{M}_k\}\Delta_i(\mathbf{x}_{k-\tau_{k,i}})\right\|^2 \right] \\
 & \quad + \sum_{t=0}^{T-1} \left[\sigma_L^2 \left(\frac{L\eta^2\eta_L^2}{mK_t} + \frac{2\tau L^2\eta^3\eta_L^3}{m}\sum_{k=t-\tau_{t,\mu}}^{t-1}\frac{1}{K_k} + 5\eta L^2\eta_L^3\frac{1}{M}\sum_{i=1}^M K_{t,i} \right) + \left(30\eta L^2\eta_L^3\frac{1}{M}\sum_{i=1}^M K_{t,i}^2 \right) \sigma_G^2 \right] \\
 & \stackrel{(b8)}{\leq} \sum_{t=0}^{T-1} \left[-\frac{1}{2}\eta\eta_L\|\nabla f(\mathbf{x}_t)\|^2 + (30\eta L^2\eta_L^3\tau)\frac{1}{M}\sum_{i=1}^M K_{t,i}^2\|\nabla f(\mathbf{x}_t)\|^2 \right] \\
 & \quad + \sum_{t=0}^{T-1} \left[-\frac{\eta\eta_L}{2M^2}\left\|\sum_{i=1}^M \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\right\|^2 + \frac{L\eta^2\eta_L^2}{m^2}\mathbb{E}\left\|\sum_{i=1}^M \mathbb{I}\{i \in \mathcal{M}_t\}\Delta_i(\mathbf{x}_{t-\tau_{t,i}})\right\|^2 \right. \\
 & \quad \left. + \frac{2L^2\eta^3\eta_L^3\tau^2}{m^2}\mathbb{E}\left\|\sum_{i \in [M]} \mathbb{I}\{i \in \mathcal{M}_t\}\Delta_i(\mathbf{x}_{t-\tau_{t,i}})\right\|^2 \right] \\
 & \quad + \sum_{t=0}^{T-1} \left[\sigma_L^2 \left(\frac{L\eta^2\eta_L^2}{mK_t} + \frac{2\tau L^2\eta^3\eta_L^3}{m}\sum_{k=t-\tau_{t,\mu}}^{t-1}\frac{1}{K_k} + 5\eta L^2\eta_L^3\frac{1}{M}\sum_{i=1}^M K_{t,i} \right) + \left(30\eta L^2\eta_L^3\frac{1}{M}\sum_{i=1}^M K_{t,i}^2 \right) \sigma_G^2 \right], \\
 & \stackrel{(b9)}{\leq} \sum_{t=0}^{T-1} \left[-\frac{1}{2}\eta\eta_L\|\nabla f(\mathbf{x}_t)\|^2 + (30\eta L^2\eta_L^3\tau)\hat{K}_t^2\|\nabla f(\mathbf{x}_t)\|^2 \right] \\
 & \quad + \sum_{t=0}^{T-1} \left[-\frac{\eta\eta_L}{2M^2}\left\|\sum_{i=1}^M \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\right\|^2 + \frac{L\eta^2\eta_L^2}{m^2}\mathbb{E}\left\|\sum_{i=1}^M \mathbb{I}\{i \in \mathcal{M}_t\}\Delta_i(\mathbf{x}_{t-\tau_{t,i}})\right\|^2 \right. \\
 & \quad \left. + \frac{2L^2\eta^3\eta_L^3\tau^2}{m^2}\mathbb{E}\left\|\sum_{i \in [M]} \mathbb{I}\{i \in \mathcal{M}_t\}\Delta_i(\mathbf{x}_{t-\tau_{t,i}})\right\|^2 \right] \\
 & \quad + \sum_{t=0}^{T-1} \left[\sigma_L^2 \left(\frac{L\eta^2\eta_L^2}{mK_t} + \frac{2\tau^2 L^2\eta^3\eta_L^3}{m}\frac{1}{K_t} + 5\eta L^2\eta_L^3\bar{K}_t \right) + \left(30\eta L^2\eta_L^3\hat{K}_t^2 \right) \sigma_G^2 \right],
 \end{aligned}$$

where (b8) is due to the fact that the delay in the system is less than τ , (b9) follows from that $\hat{K}_t^2 = \frac{1}{M}\sum_{i=1}^M K_{t,i}^2$, $\bar{K}_t = \frac{1}{M}\sum_{i=1}^M K_{t,i}$.

By letting $\mathbf{z}_i = \Delta_i(\mathbf{x}_{t-\tau_{t,i}})$ (omitting the communication round index t for notation simplicity), we have that:

$$\begin{aligned}
 \left\| \sum_{i=1}^M \mathbf{z}_i \right\|^2 &= \sum_{i \in [M]} \|\mathbf{z}_i\|^2 + \sum_{i \neq j} \langle \mathbf{z}_i, \mathbf{z}_j \rangle, \\
 &\stackrel{(b10)}{=} \sum_{i \in [M]} M\|\mathbf{z}_i\|^2 - \frac{1}{2} \sum_{i \neq j} \|\mathbf{z}_i - \mathbf{z}_j\|^2, \\
 \mathbb{E}\left\| \sum_{i=1}^M \mathbb{I}\{i \in \mathcal{M}_t\} \mathbf{z}_i \right\|^2 &= \sum_{i \in [M]} \mathbb{P}\{i \in \mathcal{M}_t\} \|\mathbf{z}_i\|^2 + \sum_{i \neq j} \mathbb{P}\{i, j \in \mathcal{M}_t\} \langle \mathbf{z}_i, \mathbf{z}_j \rangle
 \end{aligned}$$

$$\begin{aligned}
 &\stackrel{(b11)}{=} \frac{m}{M} \sum_{i \in [M]} \|\mathbf{z}_i\|^2 + \frac{m(m-1)}{M(M-1)} \sum_{i \neq j} \langle \mathbf{z}_i, \mathbf{z}_j \rangle \\
 &\stackrel{(b12)}{=} \frac{m^2}{M} \sum_{i \in [M]} \|\mathbf{z}_i\|^2 - \frac{m(m-1)}{2M(M-1)} \sum_{i \neq j} \|\mathbf{z}_i - \mathbf{z}_j\|^2,
 \end{aligned}$$

where (b10) and (b12) are due to the fact that $\langle \mathbf{x}, \mathbf{y} \rangle = \frac{1}{2}[\|\mathbf{x}\|^2 + \|\mathbf{y}\|^2 - \|\mathbf{x} - \mathbf{y}\|^2] \leq \frac{1}{2}[\|\mathbf{x}\|^2 + \|\mathbf{y}\|^2]$, (b11) follows from the fact that $\mathbb{P}\{i \in \mathcal{M}_t\} = \frac{m}{M}$ and $\mathbb{P}\{i, j \in \mathcal{M}_t\} = \frac{m(m-1)}{M(M-1)}$. It then follows that:

$$\begin{aligned}
 &-\frac{\eta\eta_L}{2M^2} \left\| \sum_{i=1}^M \mathbf{z}_i \right\|^2 + \frac{L\eta^2\eta_L^2}{m^2} \mathbb{E} \left\| \sum_{i=1}^M \mathbb{I}\{i \in \mathcal{M}_t\} \mathbf{z}_i \right\|^2 + \frac{L^2\eta^3\eta_L^3\tau^2}{m^2} \mathbb{E} \left\| \sum_{i=1}^M \mathbb{I}\{i \in \mathcal{M}_t\} \mathbf{z}_i \right\|^2 \\
 &\leq \left[-\frac{\eta\eta_L}{2M} + \left(\frac{L\eta^2\eta_L^2}{M} + \frac{L^2\eta^3\eta_L^3\tau^2}{M} \right) \right] \sum_{i=1}^M \|\mathbf{z}_i\|^2 + \frac{\eta\eta_L}{4M^2} \sum_{i \neq j} \|\mathbf{z}_i - \mathbf{z}_j\|^2 \\
 &\leq \left[-\frac{\eta\eta_L}{2M} + \left(\frac{L\eta^2\eta_L^2}{M} + \frac{L^2\eta^3\eta_L^3\tau^2}{M} \right) \right] \sum_{i=1}^M \|\mathbf{z}_i\|^2 + \frac{\eta\eta_L(M-1)}{2M^2} \sum_{i=1}^M \|\mathbf{z}_i\|^2 \\
 &= \left[-\frac{\eta\eta_L}{2M^2} + \left(\frac{L\eta^2\eta_L^2}{M} + \frac{L^2\eta^3\eta_L^3\tau^2}{M} \right) \right] \sum_{i=1}^M \|\mathbf{z}_i\|^2 \\
 &\leq 0
 \end{aligned}$$

The last inequality follows from $L\eta\eta_L + L^2\eta^2\eta_L^2\tau^2 \leq \frac{1}{2M}$.

Using the above results, we finally have:

$$\begin{aligned}
 \mathbb{E}f(\mathbf{x}_T) - f(\mathbf{x}_0) &\leq \sum_{t=0}^{T-1} \left[-\frac{1}{2}\eta\eta_L \|\nabla f(\mathbf{x}_t)\|^2 + (30\eta L^2\eta_L^3\tau) \hat{K}_t^2 \|\nabla f(\mathbf{x}_t)\|^2 \right] \\
 &\quad + \sum_{t=0}^{T-1} \left[\sigma_L^2 \left(\frac{L\eta^2\eta_L^2}{mK_t} + \frac{2\tau^2 L^2\eta^3\eta_L^3}{m} \frac{1}{K_t} + 5\eta L^2\eta_L^3 \bar{K}_t \right) + (30\eta L^2\eta_L^3 \hat{K}_t^2) \sigma_G^2 \right] \\
 &\leq \sum_{t=0}^{T-1} \left[-\frac{1}{4}\eta\eta_L \|\nabla f(\mathbf{x}_t)\|^2 + T\eta\eta_L [\alpha_L \sigma_L^2 + \alpha_G \sigma_G^2] \right]
 \end{aligned}$$

where (b13) follows from the fact that

$$\frac{1}{4} \leq \frac{1}{2} - 30L^2\hat{K}_t^2\eta_L^2\tau$$

if $120L^2\hat{K}_t^2\eta_L^2\tau < 1, \forall t$; α_L and α_G are defined as following:

$$\alpha_L = \frac{1}{T} \sum_{t=0}^{T-1} \left[\left(\frac{L\eta\eta_L}{mK_t} + \frac{2\tau^2 L^2\eta^2\eta_L^2}{mK_t} + 5\bar{K}_t L^2\eta_L^2 \right) \right],$$

and

$$\alpha_G = \frac{1}{T} \sum_{t=0}^{T-1} \left[30\hat{K}_t^2 L^2\eta_L^2 \right].$$

Lastly, by rearranging and telescoping, we have

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla f(\mathbf{x}_t)\|^2 \leq \frac{4(f_0 - f_*)}{\eta\eta_L T} + 4[\alpha_L \sigma_L^2 + \alpha_G \sigma_G^2].$$

This completes the proof. \square

Corollary 2 (Linear Speedup to a Stationary Point). *Suppose a constant local step K , let $\eta_L = \frac{1}{\sqrt{T}}$, and $\eta = \sqrt{mK}$, the convergence rate of AFA-CD with uniformly distributed worker information arrivals is:*

$$\mathcal{O}\left(\frac{1}{m^{1/2}K^{1/2}T^{1/2}}\right) + \mathcal{O}\left(\frac{\tau^2}{T}\right) + \mathcal{O}\left(\frac{K^2}{T}\right)$$

Proof. Suppose a constant local step K , let $\eta_L = \frac{1}{\sqrt{T}}$, and $\eta = \sqrt{mK}$, then it follows that:

$$\alpha_L = \mathcal{O}\left(\frac{1}{m^{1/2}K^{1/2}T^{1/2}}\right) + \mathcal{O}\left(\frac{\tau^2}{T}\right) + \mathcal{O}\left(\frac{K}{T}\right),$$

$$\alpha_G = \mathcal{O}\left(\frac{K^2}{T}\right),$$

and the proof is complete. \square

C. Proof of the performance results of the AFA-CS algorithm

Theorem 4. *Suppose that the resultant maximum delay in the system is bounded, i.e., $\tau := \max_{t \in [T], i \in \mathcal{M}_t^c} \{\tau_{t,i}\} < \infty$. Suppose that the server-side and worker-side learning rates η and η_L are chosen as such that the following relationships hold: $6\eta_L^2(2K_{t,i}^2 - 3K_{t,i} + 1)L^2 \leq 1, \forall t, i$, $\left(\frac{\eta\eta_L(M-m')^2L^2\tau^2}{M^2} + \frac{L}{2}\right)\eta\eta_L \leq \frac{1}{4}$, and $\frac{30L^2\eta_L^2\tau}{M} \left(\sum_{i \in [M]} K_{t,i}^2\right) \leq \frac{1}{4}$. Then, under Assumptions 1-3, the output sequence $\{\mathbf{x}_t\}$ generated by the AFA-CS algorithm under general worker information arrival processes satisfies:*

$$\frac{1}{T} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|^2 \leq \frac{4f(\mathbf{x}_0) - f(\mathbf{x}_T)}{\eta\eta_L T} + \alpha_L \sigma_L^2 + \alpha_G \sigma_G^2,$$

where the constants α_L and α_G are defined as follows:

$$\begin{aligned} \alpha_L &= \frac{4}{M} \left[5L^2\eta_L^2 \frac{1}{T} \sum_{t=0}^{T-1} \bar{K}_t \right. \\ &\quad \left. + \left(\frac{2\eta^2\eta_L^2(M-m')^2L^2\tau^2}{M^2} + L\eta\eta_L \right) \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{\bar{K}_t} \right], \\ \alpha_G &= \frac{120L^2\eta_L^2}{M} \frac{1}{T} \sum_{t=0}^{T-1} \hat{K}_t^2, \end{aligned}$$

and other parameters are defined the same as in Theorem 2.

Proof. We divide the stochastic gradient returns $\{\mathbf{G}_i\}$ into two groups, one is for those without delay ($\mathbf{G}_i(x_t), i \in \mathcal{M}_t, |\mathcal{M}_t| = m'$) and the other is for those with delay ($\mathbf{G}_i(\mathbf{x}_{t-\tau_{t,i}}), i \in \mathcal{M}_t^c, |\mathcal{M}_t^c| = M - m'$).

Then, the update step can be written as follows:

$$\mathbf{x}_{t+1} - \mathbf{x}_t = -\frac{\eta\eta_L}{M} \left[\sum_{i \in \mathcal{M}_t} G_i(\mathbf{x}_t) + \sum_{i \in \mathcal{M}_t^c} G_i(\mathbf{x}_{t-\tau_{t,i}}) \right].$$

Due to the L -smoothness assumption, taking expectation of $f(\mathbf{x}_{t+1})$ over the randomness in communication round t , we have:

$$\mathbb{E}[f(\mathbf{x}_{t+1})] \leq f(\mathbf{x}_t) + \underbrace{\langle \nabla f(\mathbf{x}_t), \mathbb{E}[\mathbf{x}_{t+1} - \mathbf{x}_t] \rangle}_{A_1} + \frac{L}{2} \underbrace{\mathbb{E}[\|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2]}_{A_2}$$

We first bound A_2 as follows:

$$\begin{aligned}
 A_2 &= \mathbb{E}[\|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2] \\
 &= \frac{\eta^2 \eta_L^2}{M^2} \mathbb{E} \left[\left\| \sum_{i \in \mathcal{M}_t} G_i(\mathbf{x}_t) + \sum_{i \in \mathcal{M}_t^c} G_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \right] \\
 &= \frac{\eta^2 \eta_L^2}{M^2} \mathbb{E} \left[\left\| \sum_{i \in \mathcal{M}_t} [G_i(\mathbf{x}_t) - \Delta_i(\mathbf{x}_t)] + \sum_{i \in \mathcal{M}_t^c} [G_i(\mathbf{x}_{t-\tau_{t,i}}) - \Delta_i(\mathbf{x}_{t-\tau_{t,i}})] + \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_t) + \sum_{i \in \mathcal{M}_t^c} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \right] \\
 &\stackrel{(c1)}{\leq} \frac{2\eta^2 \eta_L^2}{M^2} \left(\sum_{i \in \mathcal{M}_t} \frac{1}{K_{t,i}} + \sum_{i \in \mathcal{M}_t^c} \frac{1}{K_{t-\tau_{t,i},i}} \right) \sigma_L^2 + \frac{2\eta^2 \eta_L^2}{M^2} \left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_t) + \sum_{i \in \mathcal{M}_t^c} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \\
 &= \frac{2\eta^2 \eta_L^2}{MK_t} \sigma_L^2 + \frac{2\eta^2 \eta_L^2}{M^2} \left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_t) + \sum_{i \in \mathcal{M}_t^c} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2,
 \end{aligned}$$

where (c1) follows from the similar result in Lemma 1 and $\frac{1}{K_t} = \frac{1}{M} \left(\sum_{i \in \mathcal{M}_t} \frac{1}{K_{t,i}} + \sum_{i \in \mathcal{M}_t^c} \frac{1}{K_{t-\tau_{t,i},i}} \right)$. To bound the term A_1 , we have:

$$\begin{aligned}
 A_1 &= \mathbb{E} \langle \nabla f(\mathbf{x}_t), \mathbf{x}_{t+1} - \mathbf{x}_t \rangle \\
 &= \mathbb{E} \left\langle \nabla f(\mathbf{x}_t), -\frac{\eta \eta_L}{M} \left[\sum_{i \in \mathcal{M}_t} \mathbf{G}_i(\mathbf{x}_t) + \sum_{i \in \mathcal{M}_t^c} G_i(\mathbf{x}_{t-\tau_{t,i}}) \right] \right\rangle \\
 &= -\eta \eta_L \left\langle \nabla f(\mathbf{x}_t), \frac{1}{M} \left[\sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_t) + \sum_{i \in \mathcal{M}_t^c} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right] \right\rangle \\
 &= -\frac{\eta \eta_L}{2} \|\nabla f(\mathbf{x}_t)\|^2 - \frac{\eta \eta_L}{2M^2} \left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_t) + \sum_{i \in \mathcal{M}_t^c} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \\
 &\quad + \frac{\eta \eta_L}{2} \left\| \nabla f(\mathbf{x}_t) - \frac{1}{M} \left[\sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_t) + \sum_{i \in \mathcal{M}_t^c} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right] \right\|^2 \\
 &= -\frac{\eta \eta_L}{2} \|\nabla f(\mathbf{x}_t)\|^2 - \frac{\eta \eta_L}{2M^2} \left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_t) + \sum_{i \in \mathcal{M}_t^c} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \\
 &\quad + \frac{\eta \eta_L}{2M^2} \left\| \sum_{i \in \mathcal{M}_t} [\nabla f(\mathbf{x}_t) - \Delta_i(\mathbf{x}_t)] + \sum_{i \in \mathcal{M}_t^c} [\nabla f(\mathbf{x}_{t-\tau_{t,i}}) - \Delta_i(\mathbf{x}_{t-\tau_{t,i}})] + \sum_{i \in \mathcal{M}_t^c} [\nabla f(\mathbf{x}_t) - \nabla f(\mathbf{x}_{t-\tau_{t,i}})] \right\|^2 \\
 &\leq -\frac{\eta \eta_L}{2} \|\nabla f(\mathbf{x}_t)\|^2 - \frac{\eta \eta_L}{2M^2} \left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_t) + \sum_{i \in \mathcal{M}_t^c} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \\
 &\quad + \frac{\eta \eta_L}{M^2} \left\| \sum_{i \in \mathcal{M}_t} [\nabla f(\mathbf{x}_t) - \Delta_i(\mathbf{x}_t)] + \sum_{i \in \mathcal{M}_t^c} [\nabla f(\mathbf{x}_{t-\tau_{t,i}}) - \Delta_i(\mathbf{x}_{t-\tau_{t,i}})] \right\|^2 \\
 &\quad + \frac{\eta \eta_L}{M^2} \left\| \sum_{i \in \mathcal{M}_t^c} [\nabla f(\mathbf{x}_t) - \nabla f(\mathbf{x}_{t-\tau_{t,i}})] \right\|^2
 \end{aligned}$$

$$\begin{aligned}
 &\leq -\frac{\eta\eta_L}{2} \|\nabla f(\mathbf{x}_t)\|^2 - \frac{\eta\eta_L}{2M^2} \left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_t) + \sum_{i \in \mathcal{M}_t^c} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \\
 &+ \frac{\eta\eta_L}{M} \left[\sum_{i \in \mathcal{M}_t} \|\nabla f(\mathbf{x}_t) - \Delta_i(\mathbf{x}_t)\|^2 + \sum_{i \in \mathcal{M}_t^c} \|\nabla f(\mathbf{x}_{t-\tau_{t,i}}) - \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\|^2 \right] \\
 &+ \frac{\eta\eta_L(M-m')}{M^2} \sum_{i \in \mathcal{M}_t^c} \|\nabla f(\mathbf{x}_t) - \nabla f(\mathbf{x}_{t-\tau_{t,i}})\|^2
 \end{aligned}$$

For each worker i , we have:

$$\begin{aligned}
 \|\nabla f_i(\mathbf{x}_t) - \Delta_i(\mathbf{x}_t)\|^2 &= \left\| \nabla f_i(\mathbf{x}_t) - \frac{1}{K_{t,i}} \sum_{j=0}^{K_{t,i}-1} \nabla f_i(\mathbf{x}_{t,i}^j) \right\|^2 \\
 &= \frac{1}{K_{t,i}} \sum_{j=0}^{K_{t,i}-1} \|\nabla f_i(\mathbf{x}_t) - \nabla f_i(\mathbf{x}_{t,i}^j)\|^2 \\
 &\leq \frac{L^2}{K_{t,i}} \sum_{j=0}^{K_{t,i}-1} \|\mathbf{x}_t - \mathbf{x}_{t,i}^j\|^2 \\
 &\stackrel{(c2)}{\leq} 5K_{t,i}L^2\eta_L^2\sigma_L^2 + 30K_{t,i}^2L^2\eta_L^2\sigma_G^2 + 30K_{t,i}^2L^2\eta_L^2\|\nabla f(\mathbf{x}_t)\|^2,
 \end{aligned}$$

where (c2) follows from the same bound of A6 specified in Eq. (1).

$$\begin{aligned}
 \|\nabla f(\mathbf{x}_t) - \nabla f(\mathbf{x}_{t-\tau_{t,i}})\|^2 &\leq L^2 \|\mathbf{x}_t - \mathbf{x}_{t-\tau_{t,i}}\|^2 \\
 &\leq L^2 \tau_{t,i} \sum_{u=0}^{\tau_{t,i}-1} \|\mathbf{x}_{t-u} - \mathbf{x}_{t-u-1}\|^2.
 \end{aligned}$$

$$\begin{aligned}
 A_1 &\leq -\frac{\eta\eta_L}{2} \|\nabla f(\mathbf{x}_t)\|^2 - \frac{\eta\eta_L}{2M^2} \left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_t) + \sum_{i \in \mathcal{M}_t^c} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \\
 &+ \frac{\eta\eta_L}{M} \left[\sum_{i \in \mathcal{M}_t} \|\nabla f(\mathbf{x}_t) - \Delta_i(\mathbf{x}_t)\|^2 + \sum_{i \in \mathcal{M}_t^c} \|\nabla f(\mathbf{x}_{t-\tau_{t,i}}) - \Delta_i(\mathbf{x}_{t-\tau_{t,i}})\|^2 \right] \\
 &+ \frac{\eta\eta_L(M-m')}{M^2} \sum_{i \in \mathcal{M}_t^c} \|\nabla f(\mathbf{x}_t) - \nabla f(\mathbf{x}_{t-\tau_{t,i}})\|^2 \\
 &\leq -\frac{\eta\eta_L}{2} \|\nabla f(\mathbf{x}_t)\|^2 - \frac{\eta\eta_L}{2M^2} \left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_t) + \sum_{i \in \mathcal{M}_t^c} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \\
 &+ \frac{\eta\eta_L}{M} \left[(5L^2\eta_L^2\sigma_L^2) \left(\sum_{i \in \mathcal{M}_t} K_{t,i} + \sum_{i \in \mathcal{M}_t^c} K_{t-\tau_{t,i},i} \right) + (30L^2\eta_L^2\sigma_G^2) \left(\sum_{i \in \mathcal{M}_t} K_{t,i}^2 + \sum_{i \in \mathcal{M}_t^c} K_{t-\tau_{t,i},i}^2 \right) \right] \\
 &+ \frac{\eta\eta_L}{M} (30L^2\eta_L^2) \left(\sum_{i \in \mathcal{M}_t} K_{t,i}^2 \|\nabla f(\mathbf{x}_t)\|^2 + \sum_{i \in \mathcal{M}_t^c} K_{t-\tau_{t,i},i}^2 \|\nabla f(\mathbf{x}_{t-\tau_{t,i}})\|^2 \right)
 \end{aligned}$$

$$+ \frac{\eta\eta_L(M-m')L^2}{M^2} \sum_{i \in \mathcal{M}_t^c} \left(\tau_{t,i} \sum_{u=0}^{\tau_{t,i}-1} \|\mathbf{x}_{t-u} - \mathbf{x}_{t-u-1}\|^2 \right)$$

Combining A_1 and A_2 , we have:

$$\begin{aligned} \mathbb{E}[f(\mathbf{x}_{t+1})] - f(\mathbf{x}_t) &\leq \underbrace{\langle \nabla f(\mathbf{x}_t), \mathbb{E}[\mathbf{x}_{t+1} - \mathbf{x}_t] \rangle}_{A_1} + \frac{L}{2} \underbrace{\mathbb{E}[\|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2]}_{A_2} \\ &\leq -\frac{\eta\eta_L}{2} \|\nabla f(\mathbf{x}_t)\|^2 - \frac{\eta\eta_L}{2M^2} \left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_t) + \sum_{i \in \mathcal{M}_t^c} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 + \frac{L}{2} \mathbb{E}\|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2 \\ &\quad + \frac{\eta\eta_L}{M} \left[(5L^2\eta_L^2\sigma_L^2) \left(\sum_{i \in \mathcal{M}_t} K_{t,i} + \sum_{i \in \mathcal{M}_t^c} K_{t-\tau_{t,i},i} \right) + (30L^2\eta_L^2\sigma_G^2) \left(\sum_{i \in \mathcal{M}_t} K_{t,i}^2 + \sum_{i \in \mathcal{M}_t^c} K_{t-\tau_{t,i},i}^2 \right) \right] \\ &\quad + \frac{\eta\eta_L}{M} (30L^2\eta_L^2) \left(\sum_{i \in \mathcal{M}_t} K_{t,i}^2 \|\nabla f(\mathbf{x}_t)\|^2 + \sum_{i \in \mathcal{M}_t^c} K_{t-\tau_{t,i},i}^2 \|\nabla f(\mathbf{x}_{t-\tau_{t,i}})\|^2 \right) \\ &\quad + \frac{\eta\eta_L(M-m')L^2}{M^2} \sum_{i \in \mathcal{M}_t^c} \left(\tau_{t,i} \sum_{u=0}^{\tau_{t,i}-1} \|\mathbf{x}_{t-u} - \mathbf{x}_{t-u-1}\|^2 \right). \end{aligned}$$

Summing from $t = 0$ to $T - 1$, we have:

$$\begin{aligned} \mathbb{E}[f(\mathbf{x}_T)] - f(\mathbf{x}_0) &\leq -\frac{\eta\eta_L}{2} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|^2 - \frac{\eta\eta_L}{2M^2} \sum_{t=0}^{T-1} \left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_t) + \sum_{i \in \mathcal{M}_t^c} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 + \frac{L}{2} \sum_{t=0}^{T-1} \mathbb{E}\|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2 \\ &\quad + \frac{\eta\eta_L}{M} \sum_{t=0}^{T-1} \left[(5L^2\eta_L^2\sigma_L^2) \left(\sum_{i \in \mathcal{M}_t} K_{t,i} + \sum_{i \in \mathcal{M}_t^c} K_{t-\tau_{t,i},i} \right) + (30L^2\eta_L^2\sigma_G^2) \left(\sum_{i \in \mathcal{M}_t} K_{t,i}^2 + \sum_{i \in \mathcal{M}_t^c} K_{t-\tau_{t,i},i}^2 \right) \right] \\ &\quad + \frac{\eta\eta_L}{M} (30L^2\eta_L^2) \sum_{t=0}^{T-1} \left(\sum_{i \in \mathcal{M}_t} K_{t,i}^2 \|\nabla f(\mathbf{x}_t)\|^2 + \sum_{i \in \mathcal{M}_t^c} K_{t-\tau_{t,i},i}^2 \|\nabla f(\mathbf{x}_{t-\tau_{t,i}})\|^2 \right) \\ &\quad + \frac{\eta\eta_L(M-m')L^2}{M^2} \sum_{t=0}^{T-1} \sum_{i \in \mathcal{M}_t^c} \left(\tau_{t,i} \sum_{u=0}^{\tau_{t,i}-1} \|\mathbf{x}_{t-u} - \mathbf{x}_{t-u-1}\|^2 \right) \\ &\stackrel{(c3)}{\leq} -\frac{\eta\eta_L}{2} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|^2 - \frac{\eta\eta_L}{2M^2} \sum_{t=0}^{T-1} \left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_t) + \sum_{i \in \mathcal{M}_t^c} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \\ &\quad + \frac{\eta\eta_L}{M} \sum_{t=0}^{T-1} \left[(5L^2\eta_L^2\sigma_L^2) \bar{K}_t + (30L^2\eta_L^2\sigma_G^2) \hat{K}_t \right] + \frac{\eta\eta_L\tau}{M} (30L^2\eta_L^2) \sum_{t=0}^{T-1} \left(\sum_{i \in [M]} K_{t,i}^2 \right) \|\nabla f(\mathbf{x}_t)\|^2 \\ &\quad + \left(\frac{\eta\eta_L(M-m')^2L^2\tau^2}{M^2} + \frac{L}{2} \right) \sum_{t=0}^{T-1} (\|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2) \\ &\leq -\frac{\eta\eta_L}{2} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|^2 - \left[\frac{\eta\eta_L}{2M^2} - \left(\frac{\eta\eta_L(M-m')^2L^2\tau^2}{M^2} + \frac{L}{2} \right) \frac{2\eta^2\eta_L^2}{M^2} \right] \sum_{t=0}^{T-1} \left\| \sum_{i \in \mathcal{M}_t} \Delta_i(\mathbf{x}_t) + \sum_{i \in \mathcal{M}_t^c} \Delta_i(\mathbf{x}_{t-\tau_{t,i}}) \right\|^2 \\ &\quad + \frac{\eta\eta_L}{M} \sum_{t=0}^{T-1} \left[5L^2\eta_L^2\bar{K}_t\sigma_L^2 + 30L^2\eta_L^2\hat{K}_t\sigma_G^2 \right] + \frac{\eta\eta_L\tau}{M} (30L^2\eta_L^2) \sum_{t=0}^{T-1} \left(\sum_{i \in [M]} K_{t,i}^2 \right) \|\nabla f(\mathbf{x}_t)\|^2 \end{aligned}$$

$$\begin{aligned}
 & + \left(\frac{\eta\eta_L(M-m')^2L^2\tau^2}{M^2} + \frac{L}{2} \right) \frac{2\eta^2\eta_L^2}{M} \sigma_L^2 \sum_{t=0}^{T-1} \frac{1}{K_t} \\
 \stackrel{(c4)}{\leq} & - \sum_{t=0}^{T-1} \left[\frac{\eta\eta_L}{2} - \frac{30L^2\eta\eta_L^3\tau}{M} \left(\sum_{i \in [M]} K_{t,i}^2 \right) \right] \|\nabla f(\mathbf{x}_t)\|^2 \\
 & + \frac{\eta\eta_L}{M} \sum_{t=0}^{T-1} \left[\left[5L^2\eta_L^2\bar{K}_t + \left(\frac{\eta\eta_L(M-m')^2L^2\tau^2}{M^2} + \frac{L}{2} \right) 2\eta\eta_L \frac{1}{K_t} \right] \sigma_L^2 + \left[30L^2\eta_L^2\hat{K}_t^2 \right] \sigma_G^2 \right] \\
 \stackrel{(c5)}{\leq} & - \sum_{t=0}^{T-1} \frac{\eta\eta_L}{4} \|\nabla f(\mathbf{x}_t)\|^2 + \frac{\eta\eta_L}{M} \sum_{t=0}^{T-1} \left[\left[5L^2\eta_L^2\bar{K}_t + \left(\frac{\eta\eta_L(M-m')^2L^2\tau^2}{M^2} + \frac{L}{2} \right) 2\eta\eta_L \frac{1}{K_t} \right] \sigma_L^2 + \left[30L^2\eta_L^2\hat{K}_t^2 \right] \sigma_G^2 \right],
 \end{aligned}$$

where (c3) follows from the facts that $\bar{K}_t = \left(\sum_{i \in \mathcal{M}_t} K_{t,i} + \sum_{i \in \mathcal{M}_t^c} K_{t-\tau_t,i} \right)$, $\hat{K}_t^2 = \left(\sum_{i \in \mathcal{M}_t} K_{t,i}^2 + \sum_{i \in \mathcal{M}_t^c} K_{t-\tau_t,i}^2 \right)$, and τ is the maximum delay; (c4) is due to $\left[\frac{\eta\eta_L}{2M^2} - \left(\frac{\eta\eta_L(M-m')^2L^2\tau^2}{M^2} + \frac{L}{2} \right) \frac{2\eta^2\eta_L^2}{M^2} \right] \geq 0$ if $\left(\frac{\eta\eta_L(M-m')^2L^2\tau^2}{M^2} + \frac{L}{2} \right) \eta\eta_L \leq \frac{1}{4}$; and (c5) is due to $\frac{\eta\eta_L}{4} \leq \left[\frac{\eta\eta_L}{2} - \frac{30L^2\eta\eta_L^3\tau}{M} \left(\sum_{i \in [M]} K_{t,i}^2 \right) \right]$ if $\frac{30L^2\eta_L^2\tau}{M} \left(\sum_{i \in [M]} K_{t,i}^2 \right) \leq \frac{1}{4}$.

By rearranging, we have:

$$\frac{1}{T} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{x}_t)\|^2 \leq \frac{4f(\mathbf{x}_0) - f(\mathbf{x}_T)}{\eta\eta_L T} + \alpha_L \sigma_L^2 + \alpha_G \sigma_G^2,$$

where

$$\begin{aligned}
 \alpha_L &= \frac{4}{M} \left[5L^2\eta_L^2 \frac{1}{T} \sum_{t=0}^{T-1} \bar{K}_t + \left(\frac{2\eta^2\eta_L^2(M-m')^2L^2\tau^2}{M^2} + L\eta\eta_L \right) \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{K_t} \right], \\
 \alpha_G &= \frac{120L^2\eta_L^2}{M} \frac{1}{T} \sum_{t=0}^{T-1} \hat{K}_t^2.
 \end{aligned}$$

Note

$$\begin{aligned}
 \frac{1}{K_t} &= \frac{1}{M} \left(\sum_{i \in \mathcal{M}_t} \frac{1}{K_{t,i}} + \sum_{i \in \mathcal{M}_t^c} \frac{1}{K_{t-\tau_t,i}} \right), \\
 \bar{K}_t &= \left(\sum_{i \in \mathcal{M}_t} K_{t,i} + \sum_{i \in \mathcal{M}_t^c} K_{t-\tau_t,i} \right), \\
 \hat{K}_t^2 &= \left(\sum_{i \in \mathcal{M}_t} K_{t,i}^2 + \sum_{i \in \mathcal{M}_t^c} K_{t-\tau_t,i}^2 \right).
 \end{aligned}$$

This completes the proof. \square

Corollary 3 (Linear Speedup). *Suppose a constant local step K , and let $\eta_L = \frac{1}{\sqrt{T}}$, and $\eta = \sqrt{MK}$, the convergence rate of the AFA-CS algorithm under general worker information arrival processes is:*

$$\mathcal{O}\left(\frac{1}{M^{1/2}K^{1/2}T^{1/2}}\right) + \mathcal{O}\left(\frac{K^2}{MT}\right) + \mathcal{O}\left(\frac{\tau^2(M-m')^2}{TM^2}\right).$$

Proof. Let $\eta_L = \frac{1}{\sqrt{T}}$, and $\eta = \sqrt{MK}$. It then follows that:

$$\alpha_L = \mathcal{O}\left(\frac{1}{M^{1/2}K^{1/2}T^{1/2}}\right) + \mathcal{O}\left(\frac{K}{MT}\right) + \mathcal{O}\left(\frac{\tau^2(M-m')^2}{TM^2}\right),$$

$$\alpha_G = \mathcal{O}\left(\frac{K^2}{MT}\right).$$

This completes the proof. □

D. Discussion

Convergence Error: The case with uniformly distributed worker information arrivals under AFL can be viewed as a uniformly independent sampling process from total workers $[M]$ under conventional FL. Also, the case with general worker information arrival processes under AFL can be equivalently mapped to an arbitrarily independent sampling under conventional FL. In each communication round, the surrogate objection function for partial worker participation in FL is $\tilde{f}(x) := \frac{1}{|\mathcal{M}_t|} \sum_{i \in \mathcal{M}_t} f_i(x)$. For uniformly independent sampling, the surrogate object function approximately equals to $f(x) := \frac{1}{M} \sum_{i=1}^M f_i(x)$ in expectation, i.e., $\mathbb{E}[\tilde{f}(x)] = f(x)$. However, the surrogate object function $\tilde{f}(x)$ may deviate from $f(x)$ with arbitrarily independent sampling. More specifically, for uniformly independent sampling, the bound of $\|\nabla f(\mathbf{x}_t) - \nabla \tilde{f}(\mathbf{x}_t)\|^2$ is independent of σ_G (A_3 term in B.3). On the other hand, for arbitrarily independent sampling, $\|\nabla f(\mathbf{x}_t) - \nabla \tilde{f}(\mathbf{x}_t)\|^2 \leq \mathcal{O}(\sigma_G^2)$ (A_3 term in B.2). This deviation may happen in every communication round, so it is non-vanishing even with infinity communication rounds. As a result, such deviation is originated from the arbitrary sampling coupling with non-i.i.d. datasets. In other words, it is irrelevant to the optimization hyper-parameters such as the learning rate, local steps and others, which is different from the objective inconsistency due to different local steps shown in Wang et al. (2020). When we set $\tau = 0$ and $K_{t,i} = K, \forall t, i$, AFA-CD generalizes FedAvg. In such sense, the convergence error also exists in currently synchronous FL algorithms with such arbitrarily independent sampling and non-i.i.d. dataset. Moreover, this sampling process coupling with non-i.i.d. dataset not only results in convergence issue but also potentially induces a new source of bias/unfairness (Mohri et al., 2019; Li et al., 2019b). So how to model the practical worker participation process in practice and in turn tackle these potential bias are worth further exploration.

Variance Reduction: If we view the derivation between local loss function and global loss function as global variance, i.e., $\|\nabla f_i(\mathbf{x}_t) - \nabla f(\mathbf{x}_t)\|^2 \leq \sigma_G^2, \forall i \in [m], \forall t$ as shown in Assumption 3, the AFA-CS algorithm is indeed a variance reduction (VR) method, akin to SAG (Le Roux et al., 2012; Schmidt et al., 2017). SAG maintains an estimate stochastic gradient $v_i, i \in [n]$ for each data point (n is the size of the dataset). In each iteration, SAG only samples one data point (say, j) and update the stochastic gradient on latest model ($v_j = \nabla f_j(x_t)$) stored in the memory space, but then use the average of all stored stochastic gradients as the estimate of a full gradient to update the model ($x_{t+1} = x_t - \eta_t g_t, g_t = \frac{1}{n} \sum_{i=1}^n v_i$). In such way, SAG is able to have a faster convergence rate by reducing the local variance due to the stochastic gradient. AFA-CS algorithm performs in the similar way. The server in the AFA-CS algorithm maintains a parameter for each worker as an estimate of the returned stochastic gradient. In each communication round, the server only receives m updates in the memory space but updates the global model by the average of all the M parameters. As a result, not only can it diminish the convergence error derived from the non-i.i.d. dataset and general worker information arrival processes (arbitrarily independent sampling), but also accelerate the convergence rate with a linear speedup factor M . Previous works have applied VR methods in FL, notably SCAFFOLD (Karimireddy et al., 2020b) and FedSVRG (Konecny et al., 2016). The key difference is that we apply the VR on the server side to control the global variance while previous works focus on the worker side in order to tackle the model drift due to local update steps. Applying VR methods on server and worker side are orthogonal, and thus can be used simultaneously. We believe other variance reduction methods could be similarly extended on the server side in a similar fashion as what we do in AFA-CD. This will be left for future research.

E. Experiments

In this section, we provide the detailed experiment settings as well as extra experimental results that cannot fit in the page limit of the main paper.

E.1. Model and Datasets

We run three models on three different datasets, including i) multinomial logistic regression (LR) on manually partitioned non-i.i.d. MNIST, ii) convolutional neural network (CNN) for manually partitioned non-i.i.d. CIFAR-10, and iii) recurrent neural network (RNN) on natural non-i.i.d. Shakespeare datasets. These dataset are curated from previous FL papers (McMahan et al., 2016; Li et al., 2018) and are now widely used as benchmarks in FL studies (Li et al., 2019c; Yang et al., 2021).

For MNIST and CIFAR-10, each dataset has ten classes of images. To impose statistical heterogeneity, we split the data based on the classes (p) of images each worker contains. We distribute the data to $M = 10$ (or 100) workers such that each worker contains only certain classes with the same number of training/test samples. Specifically, each worker randomly chooses p classes of labels and evenly samples training/testing data points only with these p classes labels from the overall dataset without replacement. For example, for $p = 2$, each worker only has training/testing samples with two classes, which causes heterogeneity among different workers. For $p = 10$, each worker has samples with ten classes, which is nearly i.i.d. case. In this way, we can use the classes (p) in worker’s local dataset to represent the non-i.i.d. degree qualitatively.

The Shakespeare dataset is built from *The Complete Works of William Shakespeare* (McMahan et al., 2016). We use a two-layer LSTM classifier containing 100 hidden units with an embedding layer. The learning task is the next-character prediction, and there are 80 classes of characters in total. The model takes as input a sequence of 80 characters, embeds each of the characters into a learned 8-dimensional space and outputs one character per training sample after two LSTM layers and a densely-connected layer. The dataset and model are taken from LEAF (Li et al., 2018).

For MNIST and CIFAR-10, we use global learning rate $\eta = 1.0$ and local learning rate $\eta_L = 0.1$. For MNIST, the batch size is 64 and the total communication round is 150. For CIFAR-10, the batch size is 500 and the total communication round is 10000. For the Shakespeare dataset, the global learning rate is $\eta = 50$, the local learning rate is $\eta_L = 0.8$, batch size is $b = 10$, and the total communication round is 300. In the following tables and figure captions, we use “ m/M ” to denote that, in each communication round, we randomly choose m workers from $[M]$ to participate in the training.

We emphasize the fact that the goal here is to demonstrate our algorithms give a performance similar to other algorithms. Note that the baseline FedAvg algorithm is server-centric and under a highly coordinated environment (synchrony, uniformly worker sampling, identical local update number, etc.). In comparison, our AFL algorithms work in a far more chaotic environment with asynchrony, arbitrary worker’s arrival process, heterogeneous local steps, non-i.i.d. data, etc. The mere fact that AFL algorithms in a highly chaotic environment are *still* able to provide *comparable performance* to the highly coordinated FedAvg (as shown in our experiments) is surprising. In other words, our goal is to show that AFL algorithms can perform *nearly as well* in a much more chaotic environment, where traditional FL algorithms are not applicable. Furthermore, we use communication round instead of wall-clock time to measure the model performance. With system parameters, the wall-clock time could be easily measured. For example, by using random exponential time model $\lambda = 1$ to simulate the stragglers (Charles et al., 2021), for LR/MNIST with $p = 1$, the AFL/FedAvg ratio of communication rounds to achieve 85% accuracy is 61/46, while the ratio of wall-clock time is 1/2.6, i.e., AFL only takes 1/2.6-fraction of FedAvg’s wall-clock time.

We study the asynchrony and heterogeneity factors in AFL, including asynchrony, heterogeneous computing, worker’s arrival process, and data heterogeneity. To simulate the asynchrony, each participated worker choose one global model from the last recent five models instead of only using the latest global model for synchronous case. To mimic the heterogeneous computing, we simulate two cases: constant and dynamic local steps. For constant local steps, each participated worker performs a fixed c local update steps. In contrast, each worker takes a random local update steps uniformly sampled from $[1, 2 \times c]$ for dynamic local steps. To emulate the effect of various worker’s arrival processes, we use uniform sampling without replacement to simulate the uniformly distributed worker information arrivals, and we use biased sampling with probability $[0.19, 0.19, 0.1, 0.1, 0.1, 0.1, 0.1, 0.1, 0.01, 0.01]$ without replacement for total 10 workers to investigate potential biases with general worker information arrival processes. To study the data heterogeneity, we use the value p as a proxy to represent the non-i.i.d. degree for MNIST and CIFAR-10.

Table 1. CNN Architecture for CIFAR-10.

Layer Type	Size
Convolution + ReLu	$5 \times 5 \times 32$
Max Pooling	2×2
Convolution + ReLu	$5 \times 5 \times 64$
Max Pooling	2×2
Fully Connected + ReLU	1024×512
Fully Connected + ReLU	512×128
Fully Connected	128×10

E.2. Further experimental results

Table 2. Test Accuracy for comparison of asynchrony and local steps.

Models/ Dataset	Non-i.i.d. index (p)	Worker number	Local steps	Synchrony		Asynchrony	
				Constant steps	Dynamic steps	Constant Steps	Dynamic Steps
LR/ MNIST	$p = 1$	5/10	5	0.8916	0.8915	0.8888	0.8868
	$p = 2$	5/10	5	0.8906	0.8981	0.8901	0.8931
	$p = 5$	5/10	5	0.9072	0.9075	0.9059	0.9048
	$p = 10$	5/10	5	0.9114	0.9111	0.9129	0.9143
	$p = 1$	5/10	10	0.8743	0.8786	0.8701	0.8734
	$p = 2$	5/10	10	0.8687	0.8813	0.8661	0.8819
	$p = 5$	5/10	10	0.9016	0.9050	0.9034	0.9065
	$p = 10$	5/10	10	0.9124	0.9135	0.9112	0.9111
	$p = 1$	20/100	5	0.8898	0.8973	0.8909	0.8938
	$p = 2$	20/100	5	0.8968	0.9007	0.8955	0.9000
	$p = 5$	20/100	5	0.9088	0.9088	0.9097	0.9078
	$p = 10$	20/100	5	0.9111	0.9106	0.9126	0.9125
CNN/ CIFAR-10	$p = 1$	5/10	5	0.7474	0.7606	0.7319	0.7350
	$p = 2$	5/10	5	0.7677	0.7944	0.7662	0.777
	$p = 5$	5/10	5	0.7981	0.802	0.8065	0.799
	$p = 10$	5/10	5	0.8081	0.8072	0.8065	0.8119
RNN/ Shakespeare	-	72/143	50	0.4683	0.4831	0.4606	0.4687

Effect of asynchrony, local update steps, and non-i.i.d. level. In table 2, we examine three factors by comparing the top-1 test accuracy: synchrony versus asynchrony, constant steps versus dynamic steps and different levels of non-i.i.d. dataset. The worker sampling process is uniformly random sampling to simulate the uniformly distributed worker information arrivals. The baseline is synchrony with constant steps. When using asynchrony or/and dynamic local steps, the top-1 test accuracy shows no obvious differences. This observation can be observed in all these three tasks. Asynchrony and dynamic local update steps enable each worker to participate flexibly and loosen the coupling between workers and the server. As a result, asynchrony and dynamic local steps introduce extra heterogeneity factors, but the performance of the model is as good as that of the synchronous approaches with constant local steps. Instead, the data heterogeneity is an important factor for the model performance. As the non-i.i.d. level increases (smaller p value), the top-1 test accuracy decreases.

Next, we study convergence speed of the test accuracy for the model training under different settings. Figure 2 illustrates the test accuracy for LR on MNIST with different non-i.i.d. levels. We can see that asynchrony and dynamic local steps result in zigzagging convergence curves, but the final accuracy results have negligible differences. The zigzagging phenomenon is more dramatic as the non-i.i.d. level gets higher. Interestingly, from Figure 3 and Figure 4, we can see that for less non-i.i.d. settings such as $p = 10$ and $p = 5$, the curves of all algorithms are almost identical. Specifically, in Figure 4, the test accuracy curves of the LSTM model oscillates under asynchrony and dynamic local steps. Another observation is that it takes more rounds to converge as the non-i.i.d. level of the datasets increases. This trend can be clearly observed in Figure 3.

Utilizing FedProx and SCAFFOLD as the optimizer on the worker-side. Here, we choose FedProx and SCAFFOLD as two classes of algorithms in existing FL algorithms. FedProx represents these algorithms that modifies the local objective function. Other algorithms belonging to this category includes FedPD (Zhang et al., 2020a) and FedDyn (Acar et al., 2021). In such algorithms, no extra information exchange between worker and server is needed. On the other hand, SCAFFOLD represents VR-based (variance reduction) algorithms. It needs an extra control variate to perform the “variance reduction” step, so extra parameters are required in each communication round. Other algorithms in this class includes FedSVRG (Konecny et al., 2016).

In Table 3, we show the effectiveness of utilizing existing FL algorithms, FedProx and SCAFFOLD, in the AFL framework.




Figure 2. Test accuracy for LR on MNIST with worker number 5/10, local steps 5.

Table 3. Test Accuracy of FedProx and SCAFFOLD.

Models/ Dataset	Non-i.i.d. index (p)	Worker number	Local steps	FedProx	SCAFFOLD	AFL + FedProx	AFL + SCAFFOLD
LR/ MNIST	$p = 1$	5/10	5	0.8893	0.8928	0.8775	0.8946
	$p = 2$	5/10	5	0.8868	0.8970	0.8832	0.8954
	$p = 5$	5/10	5	0.9036	0.9032	0.9004	0.9019
	$p = 10$	5/10	5	0.9075	0.9057	0.9054	0.9022
	$p = 1$	5/10	10	0.8752	0.8789	0.8669	0.8838
	$p = 2$	5/10	10	0.8685	0.8967	0.8789	0.8978
	$p = 5$	5/10	10	0.9019	0.9047	0.8998	0.9029
	$p = 10$	5/10	10	0.9072	0.9071	0.9052	0.9038
CNN/ CIFAR-10	$p = 1$	5/10	5	0.7488	0.1641	0.7415	0.3935
	$p = 2$	5/10	5	0.7728	0.6315	0.7890	0.6971
	$p = 5$	5/10	5	0.7931	0.7828	0.8031	0.7884
	$p = 10$	5/10	5	0.8150	0.8083	0.8143	0.8051
RNN/ Shakespeare	-	72/143	50	0.4690	0.4794	0.4550	0.4515

For FedProx and SCAFFOLD, we examine synchrony and constant local steps settings. When incorporating these two advanced FL algorithms in the AFL framework, we study the effects of asynchrony and dynamic local steps. We set $\mu = 0.1$

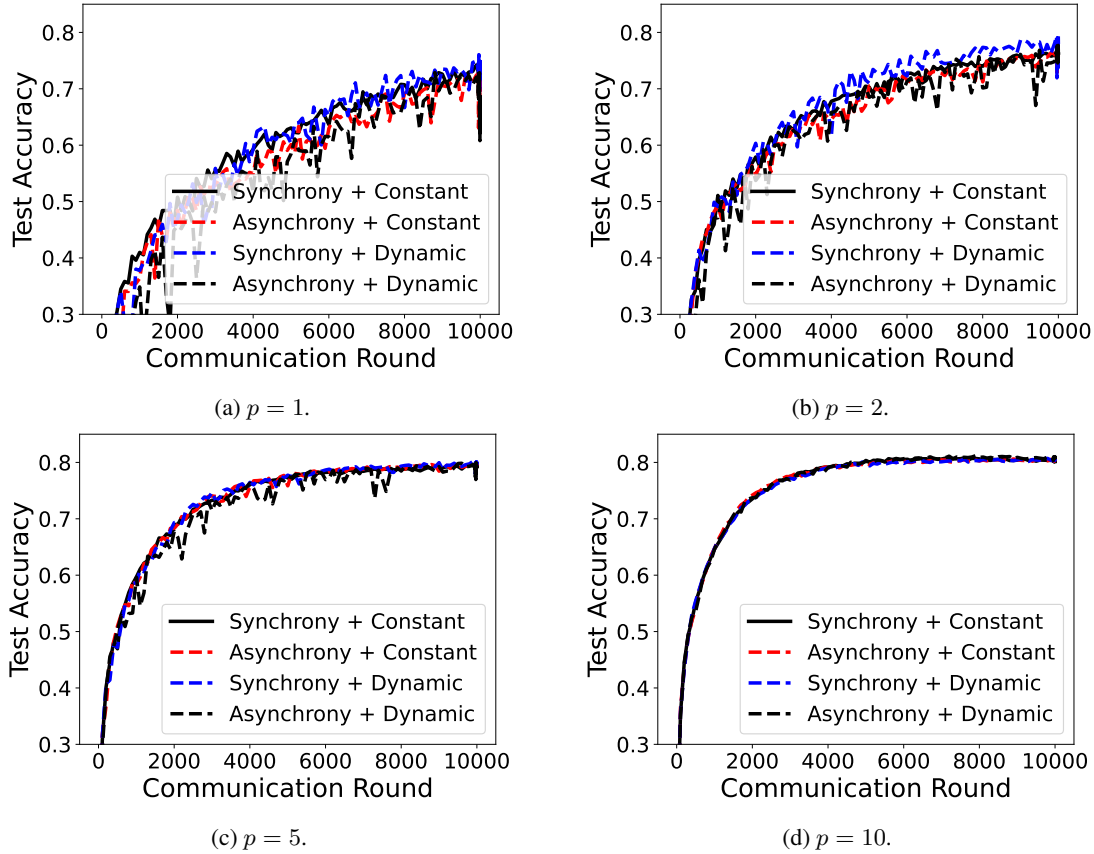


Figure 3. Test accuracy for CNN on CIFAR-10 with worker number 5/10, local steps 5.

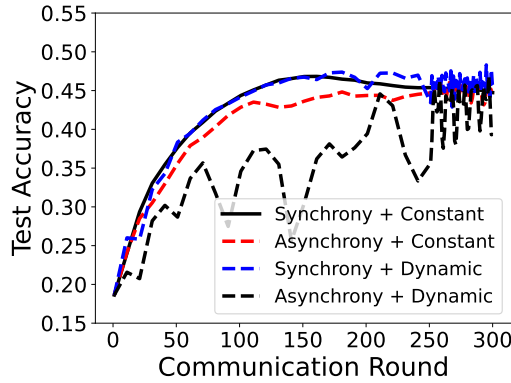


Figure 4. Test accuracy for LSTM on Shakespeare with worker number 72/143, local steps 50.

as default in FedProx algorithm. We can see from Table 3 that FedProx performs as good as FedAvg does (compare with the results in Table 2). Also, there is no performance degradation in AFL framework by utilizing FedProx as the worker’s optimizer. However, while SCAFFOLD performs well for LR on MNIST, it dose not work well for CNN on CIFAR-10, especially in cases with higher non-i.i.d. levels. One possible reason is that the control variates can become stale in partial worker participation and in turn degrade the performance. Previous work also showed similar results (Acar et al., 2021; Reddi et al., 2020). If we view the SCAFFOLD (in synchrony and constant steps setting) as the baseline, no obvious performance degradation happens under AFL with SCAFFOLD being used as the worker’s optimizer.

Effects of different worker information arrival processes. In order to generate different workers’ arrival processes, we use uniform sampling without replacement to simulate the uniformly distributed worker information arrivals and use biased




Figure 5. Test accuracy for LR on MNIST with asynchrony and dynamic local steps.

sampling to simulate the potential bias in general worker information arrival processes. In Figures 5 and 6, we illustrate the effect of the sampling process for LR on MNIST and CNN on CIFAR-10 with asynchrony and dynamic local steps. For highly non-i.i.d. datasets ($p = 1$), the biased sampling process degrades the model performance. This is consistent with the larger convergence error as shown in our theoretical analysis. On the other hand, for other non-i.i.d. cases with $p = 2, 5, 10$, such biased sampling does not lead to significant performance degradation. When applying variance reduction on such biased sampling process by reusing old gradients as shown in AFA-CS, we can see that AFA-CS performs well on MNIST, but not on CIFAR-10. We conjecture that AFA-CS, as a variance reduction method, does not always perform well in practice. This observation is consistent with the previous work (Defazio & Bottou, 2018; Reddi et al., 2020), which also demonstrated the ineffectiveness of variance reduction methods in deep learning and some cases of FL.




Figure 6. Test accuracy for CNN on CIFAR-10 with asynchrony and dynamic local steps.