



HAL
open science

Federated Learning under Heterogeneous and Correlated Client Availability

Angelo Rodio, Francescomaria Faticanti, Othmane Marfoq, Giovanni Neglia,
Emilio Leonardi

► **To cite this version:**

Angelo Rodio, Francescomaria Faticanti, Othmane Marfoq, Giovanni Neglia, Emilio Leonardi. Federated Learning under Heterogeneous and Correlated Client Availability. IEEE INFOCOM 2023 - IEEE Conference on Computer Communications, May 2023, New York City, United States. pp.1-10, 10.1109/INFOCOM53939.2023.10228876 . hal-04364285

HAL Id: hal-04364285

<https://hal.science/hal-04364285>

Submitted on 26 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Federated Learning under Heterogeneous and Correlated Client Availability

Angelo Rodio*, Francescomaria Faticanti*, Othmane Marfoq*[†], Giovanni Neglia*, Emilio Leonardi[‡]

*Inria, Université Côte d’Azur, France. Email: {firstname.lastname}@inria.fr,

[†]Accenture Labs, Sophia-Antipolis, France. Email: {firstname.lastname}@accenture.com,

[‡]Politecnico di Torino, Turin, Italy. Email: {firstname.lastname}@polito.it

Abstract—The enormous amount of data produced by mobile and IoT devices has motivated the development of federated learning (FL), a framework allowing such devices (or clients) to collaboratively train machine learning models without sharing their local data. FL algorithms (like FedAvg) iteratively aggregate model updates computed by clients on their own datasets. Clients may exhibit different levels of participation, often correlated over time and with other clients. This paper presents the first convergence analysis for a FedAvg-like FL algorithm under heterogeneous and correlated client availability. Our analysis highlights how correlation adversely affects the algorithm’s convergence rate and how the aggregation strategy can alleviate this effect at the cost of steering training toward a biased model. Guided by the theoretical analysis, we propose CA-Fed, a new FL algorithm that tries to balance the conflicting goals of maximizing convergence speed and minimizing model bias. To this purpose, CA-Fed dynamically adapts the weight given to each client and may ignore clients with low availability and large correlation. Our experimental results show that CA-Fed achieves higher time-average accuracy and a lower standard deviation than state-of-the-art AdaFed and F3AST, both on synthetic and real datasets.

Index Terms—Federated Learning, Distributed Optimization.

I. INTRODUCTION

The enormous amount of data generated by mobile and IoT devices motivated the emergence of distributed machine learning training paradigms [2], [3]. Federated Learning (FL) [4]–[7] is an emerging framework where geographically distributed devices (or clients) participate in the training of a shared Machine Learning (ML) model without sharing their local data. FL was proposed to reduce the overall cost of collecting a large amount of data as well as to protect potentially sensitive users’ private information. In the original Federated Averaging algorithm (FedAvg) [5], a central server selects a random subset of clients from the set of available clients and broadcasts them the shared model. The sampled clients perform a number of independent Stochastic Gradient Descent (SGD) steps over their local datasets and send their local model updates back to the server. Then, the server aggregates the received client updates to produce a new global model, and a new training round begins. At each iteration of FedAvg, the server typically samples randomly a few hundred devices to participate [8], [9].

This research was supported by the French government through the 3IA Côte d’Azur Investments in the Future project by the National Research Agency (ANR) with reference ANR-19-P3IA-0002, and by Groupe La Poste, sponsor of Inria Foundation, in the framework of FedMalin Inria Challenge.

An extended version of this work, including detailed proofs, is available [1].

In real-world scenarios, the availability/activity of clients is dictated by exogenous factors that are beyond the control of the orchestrating server and hard to predict. For instance, only smartphones that are idle, under charge, and connected to broadband networks are commonly allowed to participate in the training process [5], [10]. These eligibility requirements can make the availability of devices correlated over time and space [8], [11]–[13]. For example, *temporal correlation* may originate from a smartphone being under charge for a few consecutive hours and then ineligible for the rest of the day. Similarly, the activity of a sensor powered by renewable energy may depend on natural phenomena intrinsically correlated over time (e.g., solar light). *Spatial correlation* refers instead to correlation across different clients, which often emerges as a consequence of users’ different geographical distribution. For instance, clients in the same time zone often exhibit similar availability patterns, e.g., due to time-of-day effects.

Temporal correlation in the data sampling procedure is known to negatively affect the performance of ML training even in the centralized setting [14], [15] and can potentially lead to *catastrophic forgetting*: the data used during the final training phases can have a disproportionate effect on the final model, “erasing” the memory of previously learned information [16], [17]. Catastrophic forgetting has also been observed in FL, where clients in the same geographical area have more similar local data distributions and clients’ participation follows a cyclic daily pattern (leading also to spatial correlation) [8], [11], [12], [18]. Despite this evidence, a theoretical study of the convergence of FL algorithms under both temporally and spatially correlated client participation is still missing.

This paper provides the first convergence analysis of FedAvg [5] under heterogeneous and correlated client availability. We assume that the clients’ availability follows an arbitrary finite-state Markov chain: this assumption models a realistic scenario in which the activity of clients is correlated and, at the same time, still allows the analytical tractability of the system. Our theoretical analysis (i) quantifies the negative effect of correlation on the algorithm’s convergence rate through an additional term, which depends on the spectral properties of the Markov chain; (ii) points out a trade-off between two conflicting objectives: slow convergence to the optimal model, or fast convergence to a biased model, i.e., a model that minimizes an objective function different from the initial target. Guided by insights from the theoretical analysis,

we propose **CA-Fed**, an algorithm which dynamically assigns weights to clients and achieves a good trade-off between maximizing convergence speed and minimizing model bias. Interestingly, **CA-Fed** can decide to ignore clients with low availability and high temporal correlation. Our experimental results demonstrate that excluding such clients is a simple, but effective approach to handle the heterogeneous and correlated client availability in FL. Indeed, while **CA-Fed** achieves a comparable maximum accuracy as the state-of-the-art methods **F3AST** [19] and **AdaFed** [20], its test accuracy exhibits higher time-average and smaller variability over time.

The remainder of this paper is organized as follows. Section II describes the problem of correlated client availability in FL and discusses the main related works. Section III provides a convergence analysis of **FedAvg** under heterogeneous and correlated client participation. **CA-Fed**, our correlation-aware FL algorithm, is presented in Section IV. We evaluate **CA-Fed** in Section V, comparing it with state-of-the-art methods on synthetic and real-world data. Section VI concludes the paper.

II. BACKGROUND AND RELATED WORKS

We consider a finite set \mathcal{K} of N clients. Each client $k \in \mathcal{K}$ holds a local dataset D_k . Clients aim to jointly learn the parameters $\mathbf{w} \in W \subseteq \mathbb{R}^d$ of a global ML model (e.g., the weights of a neural network architecture). During training, the quality of the model with parameters \mathbf{w} on a data sample $\xi \in D_k$ is measured by a loss function $f(\mathbf{w}; \xi)$. The clients solve, under the orchestration of a central server, the following optimization problem:

$$\min_{\mathbf{w} \in W \subseteq \mathbb{R}^d} \left[F(\mathbf{w}) := \sum_{k \in \mathcal{K}} \alpha_k F_k(\mathbf{w}) \right], \quad (1)$$

where $F_k(\mathbf{w}) := \frac{1}{|D_k|} \sum_{\xi \in D_k} f(\mathbf{w}; \xi)$ is the average loss computed on client k 's local dataset, and $\alpha = (\alpha_k)_{k \in \mathcal{K}}$ are positive coefficients such that $\sum_k \alpha_k = 1$. They represent the *target importance* assigned by the central server to each client k . Typically $(\alpha_k)_{k \in \mathcal{K}}$ are set proportional to the clients' dataset size $|D_k|$, such that the objective function F in (1) coincides with the average loss computed on the union of the clients' local datasets $D = \cup_{k \in \mathcal{K}} D_k$.

Under proper assumptions, precised in Section III, Problem (1) admits a unique solution. We use \mathbf{w}^* (resp. F^*) to denote the minimizer (resp. the minimum value) of F . Moreover, for $k \in \mathcal{K}$, F_k admits a unique minimizer on W . We use \mathbf{w}_k^* (resp. F_k^*) to denote the minimizer (resp. the minimum value) of F_k .

Problem (1) is commonly solved through iterative algorithms [5], [9] requiring multiple communication rounds between the server and the clients. At round $t > 0$, the server broadcasts the latest estimate of the global model $\mathbf{w}_{t,0}$ to the set of available clients (A_t) . Client $k \in A_t$ updates the global model with its local data through $E \geq 1$ steps of local Stochastic Gradient Descent (SGD):

$$\mathbf{w}_{t,j+1}^k = \mathbf{w}_{t,j}^k - \eta_t \nabla F_k(\mathbf{w}_{t,j}^k, \mathcal{B}_{t,j}^k) \quad j = 0, \dots, E-1, \quad (2)$$

where $\eta_t > 0$ is an appropriately chosen learning rate, referred to as *local learning rate*; $\mathcal{B}_{t,j}^k$ is a random batch sampled from client- k 's local dataset at round t and step j ; $\nabla F_k(\cdot, \mathcal{B}) := \frac{1}{|\mathcal{B}|} \sum_{\xi \in \mathcal{B}} \nabla f(\cdot, \xi)$ is an unbiased estimator of the local gradient ∇F_k . Then, each client sends its local model update $\Delta_t^k := \mathbf{w}_{t,E}^k - \mathbf{w}_{t,0}^k$ to the server. The server computes $\Delta_t := \sum_{k \in A_t} q_k \cdot \Delta_t^k$, a weighted average of the clients' local updates with non-negative *aggregation weights* $\mathbf{q} = (q_k)_{k \in \mathcal{K}}$. The choice of the aggregation weights defines an aggregation strategy (we will discuss different aggregation strategies later). The aggregated update Δ_t can be interpreted as a proxy for $-\nabla F(\mathbf{w}_{t,0})$; the server applies it to the global model:

$$\mathbf{w}_{t+1,0} = \mathbf{Proj}_W(\mathbf{w}_{t,0} + \bar{\eta} \cdot \Delta_t), \quad (3)$$

where $\mathbf{Proj}_W(\cdot)$ denotes the projection over the set W , and $\bar{\eta} > 0$ is an appropriately chosen learning rate, referred to as the *server learning rate*.¹

The aggregate update Δ_t is, in general, a biased estimator of $-\nabla F(\mathbf{w}_{t,0})$, where each client k is taken into account proportionally to its frequency of appearance in the set A_t and to its aggregation weight q_k . Indeed, under proper assumptions specified in Section III, one can show (see Theorem 2) that the update rule described by (2) and (3) converges to the unique minimizer of a biased global objective F_B , which depends both on the clients' availability (i.e., on the sequence $(A_t)_{t>0}$) and on the aggregation strategy (i.e., on $\mathbf{q} = (q_k)_{k \in \mathcal{K}}$):

$$F_B(\mathbf{w}) := \sum_{k=1}^N p_k F_k(\mathbf{w}), \quad \text{with } p_k := \frac{\pi_k q_k}{\sum_{h=1}^N \pi_h q_h}, \quad (4)$$

where $\pi_k := \lim_{t \rightarrow \infty} \mathbb{P}(k \in A_t)$ is the asymptotic availability of client k . The coefficients $\mathbf{p} = (p_k)_{k \in \mathcal{K}}$ can be interpreted as the *biased importance* the server is giving to each client k during training, in general different from the *target importance* α . In what follows, \mathbf{w}_B^* (resp. F_B^*) denotes the minimizer (resp. the minimum value) of F_B .

In some large-scale FL applications, like training Google keyboard next-word prediction models, each client participates in training at most for one round. The orchestrator usually selects a few hundred clients at each round for a few thousand rounds (e.g., see [6, Table 2]), but the available set of clients may include hundreds of millions of Android devices. In this scenario, it is difficult to address the potential bias unless there is some a-priori information about each client's availability. Anyway, FL can be used by service providers with access to a much smaller set of clients (e.g., smartphone users that have installed a specific app). In this case, a client participates multiple times in training: the orchestrating server may keep track of each client's availability and try to compensate for the potentially dangerous heterogeneity in their participation.

Much previous effort on federated learning [5], [18]–[20], [23]–[26] considered this problem and, under different as-

¹The aggregation rule (3) has been considered also in other works, e.g., [9], [21], [22]. In other FL algorithms, the server computes an average of clients' local models. This aggregation rule can be obtained with minor changes to (3).

assumptions on the clients' availability (i.e., on $(A_t)_{t>0}$), designed aggregation strategies that unbiased Δ_t through an appropriate choice of \mathbf{q} . Reference [23] provides the first analysis of FedAvg on non-iid data under clients' partial participation. Their analysis covers both the case when active clients are sampled uniformly at random without replacement from \mathcal{K} and assigned aggregation weights equal to their target importance (as assumed in [5]), and the case when active clients are sampled iid with replacement from \mathcal{K} with probabilities α and assigned equal weights (as assumed in [24]). However, references [5], [23], [24] ignore the variance induced by the clients stochastic availability. The authors of [25] reduce such variance by considering only the clients with important updates, as measured by the value of their norm. References [18] and [26] reduce the aggregation variance through clustered and soft-clustered sampling, respectively.

Some recent works [19], [20], [27] do not actively pursue the optimization of the unbiased objective. Instead, they derive bounds for the convergence error and propose heuristics to minimize those bounds, potentially introducing some bias. Our work follows a similar development: we compare our algorithm with F3AST from [19] and AdaFed from [20].

The novelty of our study is in considering the spatial and temporal correlation in clients' availability dynamics. As discussed in the introduction, such correlations are also introduced by clients' eligibility criteria, e.g., smartphones being under charge and connected to broadband networks. The effect of correlation has been ignored until now, probably due to the additional complexity in studying FL algorithms' convergence. To the best of our knowledge, the only exception is [19], which scratches the issue of spatial correlation by proposing two different algorithms for the case when clients' availabilities are uncorrelated and for the case when they are positively correlated (there is no smooth transition from one algorithm to the other as a function of the degree of correlation).

The effect of temporal correlation on *centralized* stochastic gradient methods has been addressed in [13]–[15], [28]: these works study a variant of stochastic gradient descent where samples are drawn according to a Markov chain. Reference [13] extends its analysis to a FL setting where each client draws samples according to a Markov chain. In contrast, our work does not assume a correlation in the data sampling but rather in the client's availability. Nevertheless, some of our proof techniques are similar to those used in this line of work and, in particular, we rely on some results in [15].

III. ANALYSIS

A. Main assumptions

We consider a time-slotted system where a slot corresponds to one FL communication round. We assume that clients' availability over the timeslots $t \in \mathbb{N}$ follows a discrete-time Markov chain $(A_t)_{t \geq 0}$.²

²In Section III-D we will focus on the case where this chain is the superposition of N independent Markov chains, one for each client.

Assumption 1. *The Markov chain $(A_t)_{t \geq 0}$ on the finite state space $[M]$ is time-homogeneous, irreducible, and aperiodic. It has transition matrix \mathbf{P} and stationary distribution π .*

Markov chains have already been used in the literature to model the dynamics of stochastic networks where some nodes or edges in the graph can switch between active and inactive states [29], [30]. The previous Markovian assumption, while allowing a great degree of flexibility, still guarantees the analytical tractability of the system. The distance dynamics between the current and the stationary distributions of the Markov process can be characterized in terms of the spectral properties of its transition matrix \mathbf{P} [31]. Let $\lambda_2(\mathbf{P})$ denote the second largest eigenvalue of \mathbf{P} in absolute value. Previous works [15] have shown that:

$$\max_{i,j \in [M]} |[\mathbf{P}^t]_{i,j} - \pi_j| \leq C_P \cdot \lambda(\mathbf{P})^t, \quad \text{for } t \geq T_P, \quad (5)$$

where the parameter $\lambda(\mathbf{P}) := (\lambda_2(\mathbf{P}) + 1)/2$, and C_P, T_P are positive constants whose values are reported in [15, Lemma 1].³ Note that $\lambda(\mathbf{P})$ quantifies the correlation of the Markov process $(A_t)_{t \geq 0}$: the closer $\lambda(\mathbf{P})$ is to one, the slower the Markov chain converges to its stationary distribution.

In our analysis, we make the following additional assumptions. Let $\mathbf{w}^*, \mathbf{w}_B^*$ denote the minimizers of F and F_B on W , respectively.

Assumption 2. *The hypothesis class W is convex, compact, and contains in its interior the minimizers $\mathbf{w}^*, \mathbf{w}_B^*, \mathbf{w}_k^*$.*

The following assumptions concern clients' local objective functions $\{F_k\}_{k \in \mathcal{K}}$. Assumptions 3 and 4 are standard in the literature on convex optimization [32, Sections 4.1, 4.2]. Assumption 5 is a standard hypothesis in the analysis of federated optimization algorithms [9, Section 6.1].

Assumption 3 (L-smoothness). *The local functions $\{F_k\}_{k=1}^N$ have L-Lipschitz continuous gradients: $F_k(\mathbf{v}) \leq F_k(\mathbf{w}) + \langle \nabla F_k(\mathbf{w}), \mathbf{v} - \mathbf{w} \rangle + \frac{L}{2} \|\mathbf{v} - \mathbf{w}\|_2^2$, $\forall \mathbf{v}, \mathbf{w} \in W$.*

Assumption 4 (Strong convexity). *The local functions $\{F_k\}_{k=1}^N$ are μ -strongly convex: $F_k(\mathbf{v}) \geq F_k(\mathbf{w}) + \langle \nabla F_k(\mathbf{w}), \mathbf{v} - \mathbf{w} \rangle + \frac{\mu}{2} \|\mathbf{v} - \mathbf{w}\|_2^2$, $\forall \mathbf{v}, \mathbf{w} \in W$.*

Assumption 5 (Bounded variance). *The variance of stochastic gradients in each device is bounded: $\mathbb{E} \|\nabla F_k(\mathbf{w}_{t,j}^k, \xi_{t,j}^k) - \nabla F_k(\mathbf{w}_{t,j}^k)\|^2 \leq \sigma_k^2$, $k = 1, \dots, N$.*

Assumptions 2–5 imply the following properties for the local functions, described by Lemma 1 (proof in Appendix B).

Lemma 1. *Under Assumptions 2–5, there exist constants D , G , and $H > 0$, such that, for $\mathbf{w} \in W$ and $k \in \mathcal{K}$, we have:*

$$\|\nabla F_k(\mathbf{w})\| \leq D, \quad (6)$$

$$\mathbb{E} \|\nabla F_k(\mathbf{w}, \xi)\|^2 \leq G^2, \quad (7)$$

$$|F_k(\mathbf{w}) - F_k(\mathbf{w}_B^*)| \leq H. \quad (8)$$

³Note that (5) holds for different definitions of $\lambda(\mathbf{P})$ as far as $\lambda(\mathbf{P}) \in (\lambda_2(\mathbf{P}), 1)$. The specific choice for $\lambda(\mathbf{P})$ changes the constants C_P and T_P .

Similarly to other works [9], [23], [24], [33], we introduce a metric to quantify the heterogeneity of clients' local datasets:

$$\Gamma := \max_{k \in \mathcal{K}} \{F_k(\mathbf{w}^*) - F_k^*\}. \quad (9)$$

If the local datasets are identical, the local functions $\{F_k\}_{k \in \mathcal{K}}$ coincide among them and with F , \mathbf{w}^* is a minimizer of each local function, and $\Gamma = 0$. In general, Γ is smaller the closer the distributions the local datasets are drawn from.

B. Main theorems

Theorem 1 (proof in Appendix A) decomposes the error of the target global objective as the sum of an optimization error for the biased global objective and a bias error.

Theorem 1 (Decomposing the total error). *Under Assumptions 2–4, the optimization error of the target global objective $\epsilon = F(\mathbf{w}) - F^*$ can be bounded as follows:*

$$\epsilon \leq \underbrace{2\kappa^2(F_B(\mathbf{w}) - F_B^*)}_{:=\epsilon_{\text{opt}}} + \underbrace{2\kappa^4\chi_{\alpha\|\mathbf{p}}^2}_{:=\epsilon_{\text{bias}}}\Gamma, \quad (10)$$

where $\kappa := L/\mu$, and $\chi_{\alpha\|\mathbf{p}}^2 := \sum_{k=1}^N (\alpha_k - p_k)^2/p_k$.

Theorem 2 below proves that the optimization error ϵ_{opt} associated to the biased objective F_B , evaluated on the trajectory determined by scheme (3), asymptotically vanishes. The non-vanishing bias error ϵ_{bias} captures the discrepancy between $F(\mathbf{w})$ and $F_B(\mathbf{w})$. This latter term depends on the chi-square divergence $\chi_{\alpha\|\mathbf{p}}^2$ between the target and biased probability distributions $\alpha = (\alpha_k)_{k \in \mathcal{K}}$ and $\mathbf{p} = (p_k)_{k \in \mathcal{K}}$, and on Γ , that quantifies the degree of heterogeneity of the local functions. When all local functions are identical ($\Gamma = 0$), the bias term ϵ_{bias} also vanishes. For $\Gamma > 0$, the bias error can still be controlled by the aggregation weights assigned to the devices. In particular, the bias term vanishes when $q_k \propto \alpha_k/\pi_k, \forall k \in \mathcal{K}$. Since it asymptotically cancels the bias error, we refer to this choice as *unbiased aggregation strategy*.

However, in practice, FL training is limited to a finite number of iterations T (typically a few hundreds [6], [8]), and the previous asymptotic considerations may not apply. In this regime, the unbiased aggregation strategy can be suboptimal, since the minimization of ϵ_{bias} not necessarily leads to the minimization of the total error $\epsilon \leq \epsilon_{\text{opt}} + \epsilon_{\text{bias}}$. This motivates the analysis of the optimization error ϵ_{opt} .

Theorem 2 (Convergence of the optimization error ϵ_{opt}). *Let Assumptions 1–5 hold and the constants $M, L, D, G, H, \Gamma, \sigma_k, C_P, T_P, \lambda(\mathbf{P})$ be defined as above. Let $Q = \sum_{k \in \mathcal{K}} q_k$. Let the stepsizes satisfy:*

$$\sum_t \eta_t = +\infty, \quad \sum_t \ln(t) \cdot \eta_t^2 < +\infty. \quad (11)$$

Let T denote the total communication rounds. For $T \geq T_P$, the expected optimization error can be bounded as follows:

$$\mathbb{E}[F_B(\bar{\mathbf{w}}_{T,0}) - F_B^*] \leq \frac{\frac{1}{2}\mathbf{q}^\top \Sigma \mathbf{q} + v}{\pi^\top \mathbf{q}} + \psi + \frac{\phi}{\ln(1/\lambda(\mathbf{P}))}, \quad (12)$$

where $\bar{\mathbf{w}}_{T,0} := \frac{\sum_{t=1}^T \eta_t \mathbf{w}_{t,0}}{\sum_{t=1}^T \eta_t}$, and

$$\begin{aligned} \Sigma &= \text{diag}(\sigma_k^2 \pi_k \sum_t \eta_t^2), \\ v &= \frac{2}{E} \|\mathbf{w}_{0,0} - \mathbf{w}^*\|^2 + \frac{1}{4} M Q \sum_t (\eta_t^2 + \frac{1}{t^2}), \\ \psi &= 4L(EQ + 2)\Gamma \sum_t \eta_t^2 + \frac{2}{3}(E-1)(2E-1)G^2 \sum_t \eta_t^2, \\ \mathcal{J}_t &= \min \{ \max \{ \lceil \ln(2C_P H t) / \ln(1/\lambda(\mathbf{P})) \rceil, T_P \}, t \}, \\ \phi &= 2EDGQ \sum_t \ln(2C_P H t) \eta_{t-\mathcal{J}_t}^2. \end{aligned}$$

Theorem 2 (proof in Appendix B) proves convergence of the expected biased objective F_B to its minimum F_B^* under correlated client participation. Our bound (12) captures the effect of correlation through the factor $\ln(1/\lambda(\mathbf{P}))$: a high correlation worsens the convergence rate. In particular, we found that the numerator of (12) has a quadratic-over-linear fractional dependence on \mathbf{q} . Minimizing ϵ_{opt} leads, in general, to a different choice of \mathbf{q} than minimizing ϵ_{bias} .

C. Minimizing the total error $\epsilon \leq \epsilon_{\text{opt}} + \epsilon_{\text{bias}}$

Our analysis points out a trade-off between minimizing ϵ_{opt} or ϵ_{bias} . Our goal is to find the optimal aggregation weights \mathbf{q}^* that minimize the upper bound on total error $\epsilon(\mathbf{q})$ in (10):

$$\begin{aligned} &\underset{\mathbf{q}}{\text{minimize}} && \epsilon_{\text{opt}}(\mathbf{q}) + \epsilon_{\text{bias}}(\mathbf{q}); \\ &\text{subject to} && \mathbf{q} \geq 0, \\ &&& \|\mathbf{q}\|_1 = Q. \end{aligned} \quad (13)$$

In Appendix E we prove that (13) is a convex optimization problem, which can be solved with the method of Lagrange multipliers. However, the solution is not of practical utility because the constants in (10) and (12) (e.g., L, μ, Γ, C_P) are in general problem-dependent and difficult to estimate during training. In particular, Γ poses particular difficulties as it is defined in terms of the minimizer of the target objective F , but the FL algorithm generally minimizes the biased function F_B . Moreover, the bound in (10), similarly to the bound in [33], diverges when setting some q_k equal to 0, but this is simply an artifact of the proof technique. A result of more practical interest is the following (proof in Appendix C):

Theorem 3 (An alternative decomposition of the total error ϵ). *Under the same assumptions of Theorem 1, let $\Gamma' := \max_k \{F_k(\mathbf{w}_B^*) - F_k^*\}$. The following result holds:*

$$\epsilon \leq \underbrace{2\kappa^2(F_B(\mathbf{w}) - F_B^*)}_{:=\epsilon_{\text{opt}}} + \underbrace{8\kappa^4 d_{TV}^2(\alpha, \mathbf{p}) \Gamma'}_{:=\epsilon'_{\text{bias}}}, \quad (14)$$

where $d_{TV}(\alpha, \mathbf{p}) := \frac{1}{2} \sum_{k=1}^N |\alpha_k - p_k|$ is the total variation distance between the probability distributions α and \mathbf{p} .

The new constant Γ' is defined in terms of \mathbf{w}_B^* , and then it is easier to evaluate during training. However, Γ' depends on \mathbf{q} , because it is evaluated at the point of minimum of F_B . This dependence makes the minimization of the right-hand side of (14) more challenging (for example, the corresponding problem is not convex). We study the minimization of the two terms ϵ_{opt} and ϵ'_{bias} separately and learn some insights, which we use to design the new FL algorithm CA-FeD.

D. Minimizing ϵ_{opt}

The minimization of ϵ_{opt} is still a convex optimization problem (Appendix D). In particular, at the optimum, non-negative weights are set accordingly to $q_k^* = a(\lambda^* \pi_k - \theta^*)$ with a , λ^* , and θ^* positive constants (see (29)). It follows that clients with smaller availability get smaller weights in the aggregation. In particular, this suggests that clients with the smallest availability can be excluded from the aggregation, leading to the following guideline:

Guideline A: to speed up the convergence, we can exclude, i.e., set $q_k^ = 0$, the clients with lowest availability π_k .*

This guideline can be justified intuitively: updates from clients with low participation may be too sporadic to allow the FL algorithm to keep track of their local objectives. They act as a noise slowing down the algorithm's convergence. It may be advantageous to exclude these clients from participating.

We observe that the choice of the aggregation weights q does not affect the clients' availability process and, in particular, $\lambda(\mathbf{P})$. However, if the algorithm excludes some clients, it is possible to consider the state space of the Markov chain that only specifies the availability state of the remaining clients, and this Markov chain may have different spectral properties. For the sake of concreteness, we consider here (and in the rest of the paper) the particular case when the availability of each client k evolves according to a Markov chain $(A_t^k)_{t \geq 0}$ with transition probability matrix \mathbf{P}_k and these Markov chains are all independent. In this case, the aggregate process is described by the product Markov chain $(A_t)_{t \geq 0}$ with transition matrix $\mathbf{P} = \bigotimes_{k \in \mathcal{K}} \mathbf{P}_k$ and $\lambda(\mathbf{P}) = \max_{k \in \mathcal{K}} \lambda(\mathbf{P}_k)$, where $\mathbf{P}_i \otimes \mathbf{P}_j$ denotes the Kronecker product between matrices \mathbf{P}_i and \mathbf{P}_j [31, Exercise 12.6]. In this setting, it is possible to redefine the Markov chain $(A_t)_{t \geq 0}$ by taking into account the reduced state space defined by the clients with a non-null aggregation weight, i.e., $\mathbf{P}' = \bigotimes_{k' \in \mathcal{K} | q_{k'} > 0} \mathbf{P}_{k'}$ and $\lambda(\mathbf{P}') = \max_{k' \in \mathcal{K} | q_{k'} > 0} \lambda(\mathbf{P}_{k'})$, which is potentially smaller than the case when all clients participate to the aggregation. These considerations lead to the following guideline:

Guideline B: to speed up the convergence, we can exclude, i.e., set $q_k^ = 0$, the clients with largest $\lambda(\mathbf{P}_k)$.*

Intuition also supports this guideline. Clients with large $\lambda(\mathbf{P}_k)$ tend to be available or unavailable for long periods of time. Due to the well-known catastrophic forgetting problem affecting gradient methods [34], [35], these clients may unfairly steer the algorithm toward their local objective when they appear at the final stages of the training period. Moreover, their participation in the early stages may be useless, as their contribution will be forgotten during their long absence. The FL algorithm may benefit from directly neglecting such clients.

We observe that guideline B strictly applies to this specific setting where clients' dynamics are independent (and there is no spatial correlation). We do not provide a corresponding guideline for the case when clients are spatially correlated

Algorithm 1: CA-Fed (Correlation-Aware FL)

Input : $w_{0,0}$, α , $q^{(0)}$, $\{\eta_t\}_{t=1}^T$, $\bar{\eta}$, E , β , τ

- 1 Initialize $\hat{\mathbf{F}}^{(0)}$, $\hat{\mathbf{F}}^*$, $\hat{\Gamma}^{(0)}$, $\hat{\pi}^{(0)}$, and $\hat{\lambda}^{(0)}$;
- 2 **for** $t = 1, \dots, T$ **do**
- 3 Receive set of active client A_t , loss vector $\mathbf{F}^{(t)}$;
- 4 Update $\hat{\mathbf{F}}^{(t)}$, $\hat{\Gamma}^{(t)}$, $\hat{\pi}^{(t)}$, and $\hat{\lambda}^{(t)}$;
- 5 Initialize $q^{(t)} = \frac{\alpha}{\hat{\pi}^{(t)}}$;
- 6 $q^{(t)} \leftarrow \text{get}(q^{(t)}, \alpha, \hat{\mathbf{F}}^{(t)}, \hat{\mathbf{F}}^*, \hat{\Gamma}^{(t)}, \hat{\pi}^{(t)}, \hat{\lambda}^{(t)})$;
- 7 $q^{(t)} \leftarrow \text{get}(q^{(t)}, \alpha, \hat{\mathbf{F}}^{(t)}, \hat{\mathbf{F}}^*, \hat{\Gamma}^{(t)}, \hat{\pi}^{(t)}, -\hat{\pi}^{(t)})$;
- 8 **for** client $\{k \in A_t; q_k^{(t)} > 0\}$, in parallel **do**
- 9 **for** $j = 0, \dots, E - 1$ **do**
- 10 $w_{t,j+1}^k = w_{t,j}^k - \eta_t \nabla F_k(w_{t,j}^k, \mathcal{B}_{t,j}^k)$;
- 11 $\Delta_t^k \leftarrow w_{t,E}^k - w_{t,0}^k$;
- 12 $w_{t+1,0} \leftarrow \text{Proj}_W(w_{t,0} + \bar{\eta} \sum_{k \in A_t} q_k^{(t)} \cdot \Delta_t^k)$;
- 13 **Function** $\text{get}(q, \alpha, \mathbf{F}, \mathbf{F}^*, \Gamma, \pi, \rho)$:
- 14 Sort \mathcal{K} by descending order in ρ ;
- 15 $\hat{\epsilon} \leftarrow \langle \mathbf{F} - \mathbf{F}^*, \pi \tilde{\odot} q \rangle + d_{TV}^2(\alpha, \pi \tilde{\odot} q) \cdot \Gamma$;
- 16 **for** $k \in \mathcal{K}$ **do**
- 17 $q_k^+ \leftarrow 0$;
- 18 $\hat{\epsilon}^+ \leftarrow \langle \mathbf{F} - \mathbf{F}^*, \pi \tilde{\odot} q^+ \rangle + d_{TV}^2(\alpha, \pi \tilde{\odot} q^+) \cdot \Gamma$;
- 19 **if** $\hat{\epsilon} - \hat{\epsilon}^+ \geq \tau$ **then**
- 20 $\hat{\epsilon} \leftarrow \hat{\epsilon}^+$;
- 21 $q \leftarrow q^+$;
- 22 **return** q

(we leave this task for future research). However, in this more general setting, it is possible to ignore guideline B but still draw on guidelines A and C, or still consider guideline B if the spatially correlated clients can be grouped in clusters, each cluster evolving as an independent Markov chain [1].

E. Minimizing ϵ'_{bias}

The bias error ϵ'_{bias} in (14) vanishes when the total variation distance between the target importance α and the biased importance \mathbf{p} is zero, i.e., when $q_k \propto \alpha_k / \pi_k, \forall k \in \mathcal{K}$. Then, after excluding the clients that contribute the most to the optimization error and particularly slow down the convergence (guidelines A and B), we can assign to the remaining clients an aggregation weight inversely proportional to their availability, such that the bias error ϵ'_{bias} is minimized.

Guideline C: to reduce the bias error, we set $q_k^ \propto \alpha_k / \pi_k$ for the clients that are not excluded by the previous guidelines.*

IV. PROPOSED ALGORITHM

Guidelines A and B in Section III suggest that the minimization of ϵ_{opt} can lead to the exclusion of some available clients from the aggregation step (3), in particular those with low availability and/or high correlation. For the remaining clients, guideline C proposes to set their aggregation weight inversely proportional to their availability to reduce the bias error ϵ'_{bias} . Motivated by these insights, we propose CA-Fed, a client sampling and aggregation strategy that takes into account the problem of correlated client availability in FL, described in

Algorithm 1. CA-Fed learns during training which are the clients to exclude and how to set the aggregation weights of the other clients to achieve a good trade-off between ϵ_{opt} and ϵ'_{bias} . While guidelines A and B indicate which clients to remove, the exact number of clients to remove at round t is identified by minimizing $\epsilon^{(t)}$ as a proxy for the bound in (14):⁴

$$\epsilon^{(t)} := F_B(\mathbf{w}_{t,0}) - F_B^* + d_{TV}^2(\boldsymbol{\alpha}, \mathbf{p})\Gamma'. \quad (15)$$

A. CA-Fed's core steps

At each communication round t , the server sends the current model $\mathbf{w}_{t,0}$ to all active clients and each client k sends back a noisy estimate $F_k^{(t)}$ of the current loss computed on a batch of samples $\mathcal{B}_{t,0}^k$, i.e., $F_k^{(t)} = \frac{1}{|\mathcal{B}_{t,0}^k|} \sum_{\xi \in \mathcal{B}_{t,0}^k} f(\mathbf{w}_{t,0}, \xi)$ (line 3). The server uses these values and the information about the current set of available clients A_t to refine its own estimates of each client's loss ($\hat{\mathbf{F}}^{(t)} = (\hat{F}_k^{(t)})_{k \in \mathcal{K}}$), and each client's loss minimum value ($\hat{\mathbf{F}}^* = (\hat{F}_k^*)_{k \in \mathcal{K}}$), as well as of Γ' , π_k , λ_k , and $\epsilon^{(t)}$, denoted as $\hat{\Gamma}^{(t)}$, $\hat{\pi}_k^{(t)}$, $\hat{\lambda}_k^{(t)}$, and $\hat{\epsilon}^{(t)}$, respectively (possible estimators are described below) (line 4).

The server decides whether excluding clients whose availability pattern exhibits high correlation (high $\hat{\lambda}_k^{(t)}$) (line 6). First, the server considers all clients in descending order of $\hat{\lambda}^{(t)}$ (line 14), and evaluates if, by excluding them (line 17), $\hat{\epsilon}^{(t)}$ appears to be decreasing by more than a threshold $\tau \geq 0$ (line 19). Then, the server considers clients in ascending order of $\hat{\pi}^{(t)}$, and repeats the same procedure to possibly exclude some of the clients with low availability (low $\hat{\pi}_k^{(t)}$) (lines 7).

Once the participating clients (those with $q_k > 0$) have been selected, the server notifies them to proceed updating the current models (lines 9–10) according to (2), while the other available clients stay idle. Finally, model's updates are aggregated according to (3) (line 12).

B. Estimators

We now briefly discuss possible implementation of the estimators $\hat{F}_k^{(t)}$, \hat{F}_k^* , $\hat{\Gamma}^{(t)}$, $\hat{\pi}_k^{(t)}$, and $\hat{\lambda}_k^{(t)}$. Server's estimates for the clients' local losses ($\hat{\mathbf{F}}^{(t)} = (\hat{F}_k^{(t)})_{k \in \mathcal{K}}$) can be obtained from the received active clients' losses ($\mathbf{F}^{(t)} = (F_k^{(t)})_{k \in A_t}$) through an auto-regressive filter with parameter $\beta \in (0, 1]$:

$$\hat{\mathbf{F}}^{(t)} = (\mathbf{1} - \beta \mathbf{1}_{A_t}) \odot \hat{\mathbf{F}}^{(t-1)} + \beta \mathbf{1}_{A_t} \odot \mathbf{F}^{(t)}, \quad (16)$$

where \odot denotes the component-wise multiplication between vectors, and $\mathbf{1}_{A_t}$ is a N -dimensions binary vector whose k -th component equals 1 if and only if k is active at round t , i.e., $k \in A_t$. The server can keep track of the clients' loss minimum values and estimate F_k^* as $\hat{F}_k^* = \min_{s \in [0, t]} \hat{F}_k^{(s)}$. The values of $F_B(\mathbf{w}_{t,0})$, F_B^* , Γ' , and $\epsilon^{(t)}$ can be estimated as follows:

$$\hat{F}_B^{(t)} - \hat{F}_B^* = \langle \hat{\mathbf{F}}^{(t)} - \hat{\mathbf{F}}^*, \hat{\boldsymbol{\pi}}^{(t)} \tilde{\odot} \mathbf{q}^{(t)} \rangle, \quad (17)$$

$$\hat{\Gamma}^{(t)} = \max_{k \in \mathcal{K}} (\hat{F}_k^{(t)} - \hat{F}_k^*), \quad (18)$$

$$\hat{\epsilon}^{(t)} = \hat{F}_B^{(t)} - \hat{F}_B^* + d_{TV}^2(\boldsymbol{\alpha}, \hat{\boldsymbol{\pi}}^{(t)} \tilde{\odot} \mathbf{q}^{(t)}) \cdot \hat{\Gamma}^{(t)}. \quad (19)$$

⁴Following (14), one could introduce a hyper-parameter to weigh the relative importance of the optimization and bias terms in the sum. We have discussed this additional optimization of CA-Fed in the extended version [1].

where $\boldsymbol{\pi} \tilde{\odot} \mathbf{q} \in \mathbb{R}^N$, such that $(\boldsymbol{\pi} \tilde{\odot} \mathbf{q})_k = \frac{\pi_k q_k}{\sum_{h=1}^N \pi_h q_h}$, $k \in \mathcal{K}$.

For $\hat{\pi}_k^{(t)}$, the server can simply keep track of the total number of times client k was available up to time t and compute $\hat{\pi}_k^{(t)}$ using a Bayesian estimator with beta prior, i.e., $\hat{\pi}_k^{(t)} = (\sum_{s \leq t} \mathbb{1}_{k \in A_s} + n_k) / (t + n_k + m_k)$, where n_k and m_k are the initial parameters of the beta prior.

For $\hat{\lambda}_k^{(t)}$, the server can assume the client's availability evolves according to a Markov chain with two states (available and unavailable), track the corresponding number of state transitions, and estimate the transition matrix $\hat{\mathbf{P}}_k^{(t)}$ through a Bayesian estimator similarly to what done for $\hat{\pi}_k^{(t)}$. Finally, $\hat{\lambda}_k^{(t)}$ is obtained computing the eigenvalues of $\hat{\mathbf{P}}_k^{(t)}$.

C. CA-Fed's computation/communication cost

CA-Fed aims to improve training convergence and not to reduce its computation and communication overhead. Nevertheless, excluding some available clients reduces the overall training cost, as we will discuss in this section referring, for the sake of concreteness, to neural networks' training.

The available clients not selected for training are only requested to evaluate their local loss on the current model once on a single batch instead than performing E gradient updates, which would require roughly $2 \times E - 1$ more calculations (because of the forward and backward pass). For the selected clients, there is no extra computation cost as computing the loss corresponds to the forward pass they should, in any case, perform during the first local gradient update.

In terms of communication, the excluded clients only transmit the loss, a single scalar, much smaller than the model update. Conversely, participating clients transmit the local loss and the model update. Still, this additional overhead is negligible and likely fully compensated by the communication savings for the excluded clients.

V. EXPERIMENTAL EVALUATION

A. Experimental Setup

a) Federated system simulator: In our experiments, we simulate the clients' availability dynamics featuring different levels of temporal correlations. We model the activity of each client as a two-state homogeneous Markov process with state space $\mathcal{S} = \{\text{"active"}, \text{"inactive"}\}$. We use $p_{k,s}$ to denote the probability that client $k \in \mathcal{K}$ remains in state $s \in \mathcal{S}$.

In order to simulate the statistical heterogeneity present in the federated learning system, we consider an experimental setting with two disjoint groups of clients G_i , $i = 1, 2$, to which we associate two different data distributions \mathcal{P}_i , $i = 1, 2$, to be precised later. Let $r_i = |G_i|/N$, $i = 1, 2$ denote the fraction of clients in group $i = 1, 2$. In order to simulate the heterogeneity of clients' availability patterns in realistic federated systems, we split the clients of each group in two classes uniformly at random: "more available" clients whose steady-state probability to be active is $\pi_{k,\text{active}} = 1/2 + g$ and "less available" clients with $\pi_{k,\text{active}} = 1/2 - g$, where $g \in$

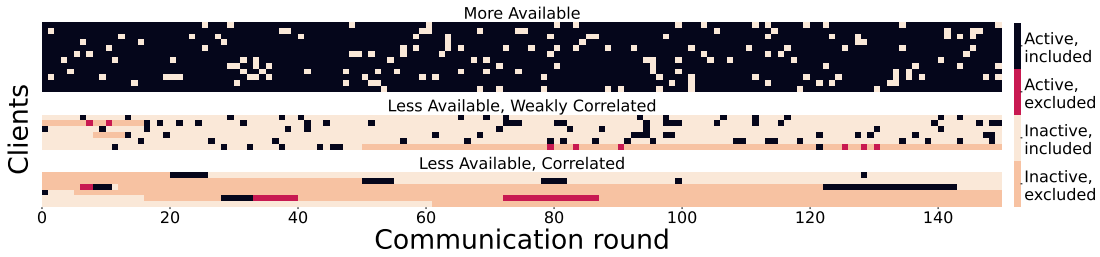


Fig. 1: Clients' activities and CA-Fed's clients selection on the synthetic dataset.





Fig. 2: Importance given to the clients by the different algorithms throughout a whole training process on the synthetic dataset.

$(0, 1/2)$ is a parameter controlling the heterogeneity of clients availability. We furthermore split each class of clients in two sub-classes uniformly at random: “correlated” clients that tend to persist in the same state ($\lambda_k = \nu$ with values of ν close to 1), and “weakly correlated” clients that are almost as likely to keep as to change their state ($\lambda_k \sim \mathcal{N}(0, \varepsilon^2)$, with ε close to 0). In our experiments, we suppose that $r_1 = r_2 = 1/2$, $g = 0.4$, $\nu = 0.9$, and $\varepsilon = 10^{-2}$.

b) Datasets and models: All experiments are performed on a binary classification synthetic dataset (described in Appendix F) and on the real-world MNIST dataset [36], using $N = 24$ clients. For MNIST dataset, we introduce statistical heterogeneity across the two groups of clients (i.e., we make the two distributions \mathcal{P}_1 and \mathcal{P}_2 different), following the same approach in [37]: 1) every client is assigned a random subset of the total training data; 2) the data of clients from the second group is modified by randomly swapping two pairs of labels. We maintain the original training/test data split of MNIST and use 20% of the training dataset as validation dataset. We use a linear classifier with a ridge penalization of parameter 10^{-2} , which is a strongly convex objective function, for both the synthetic and the real-world MNIST datasets.

c) Benchmarks: We compare CA-Fed, defined in Algorithm 1, with the Unbiased aggregation strategy, where all the active clients participate and receive a weight inversely proportional to their availability, and with the state-of-the-art FL algorithms discussed in Section II: F3AST [19] and AdaFed [20]. We tuned the learning rates η , $\bar{\eta}$ via grid search, on the grid $\eta : \{10^{-3}, 10^{-2.5}, 10^{-2}, 10^{-1.5}, 10^{-1}\}$, $\bar{\eta} : \{10^{-2}, 10^{-1.5}, 10^{-1}, 10^{-0.5}, 10^0\}$. For CA-Fed, we used $\tau = 0$, $\beta = 0.2$. We assume all algorithms can access an oracle providing the true availability parameters for each client. In



(a) Synthetic

(b) MNIST

Fig. 3: Test accuracy vs number of communication rounds.

practice, Unbiased, AdaFed, and F3AST rely on the exact knowledge of $\pi_{k,\text{active}}$, and CA-Fed on $\pi_{k,\text{active}}$ and λ_k .⁵

B. Experimental Results

Figure 1 shows the availability of each client during a training run on the synthetic dataset. Clients selected (resp. excluded) by CA-Fed are highlighted in black (resp. red). We observe that excluded clients tend to be those with low average availability or high correlation.

Figure 2 shows the importance p_k (averaged over time) given by different algorithms to each client k during a full training run. We observe that all the algorithms, except Unbiased, depart from the target importance α . As suggested by guidelines A and B, CA-Fed tends to favor the group of “more available” clients, at the expense of the “less available” clients.

Figure 3 shows the time-average accuracy up to round t of the learned model averaged over three different runs. On both datasets, CA-Fed achieves the highest accuracy, which is about a percentage point higher than the second best algorithm (F3AST). Table I shows for each algorithm: the average over three runs of the maximum test accuracy achieved during training, the time-average test accuracy achieved during training, together with its standard deviation within the second half of the training period. While CA-Fed achieves a maximum accuracy which is comparable to the Unbiased baseline and state-of-the-art AdaFed and F3AST, it gets a higher time-average accuracy (1.24 pp on MNIST) in comparison to the second best (F3AST), and a smaller standard deviation ($1.5\times$ on MNIST) in comparison to the second best (F3AST).

⁵The authors have provided public access to their code and data at: <https://github.com/arodio/CA-Fed>.

TABLE I: Maximum and time-average test accuracy, together with their standard deviations, on the Synthetic / MNIST datasets.

	MAXIMUM	TEST ACCURACY	
		TIME-AVERAGE	STANDARD DEVIATION
UNBIASED	78.94 / 64.87	75.32 / 61.39	0.48 / 1.09
F3AST	78.97 / 64.91	75.33 / 61.52	0.40 / 0.94
ADAFED	78.69 / 63.77	74.81 / 60.48	0.59 / 1.37
CA-FED	79.03 / 64.94	76.22 / 62.76	0.28 / 0.61

VI. CONCLUSION

This paper presented the first convergence analysis for a FedAvg-like FL algorithm under heterogeneous and correlated client availability. The analysis quantifies how correlation adversely affects the algorithm's convergence rate and highlights a general bias-versus-convergence-speed trade-off. Guided by the theoretical analysis, we proposed CA-Fed, a new FL algorithm that tries to balance the conflicting goals of maximizing convergence speed and minimizing model bias. Our experimental results demonstrate that adaptively excluding clients with high temporal correlation and low availability is an effective approach to handle the heterogeneous and correlated client availability in FL. Our work represents a first step toward the development of effective FL techniques under temporally and spatially correlated clients' availability. Although our algorithm is mainly designed to handle temporal correlation, it can deal with spatial correlation as well, as discussed in the extended version [1].

APPENDIX

A. Proof of Theorem 1

We bound the optimization error of the target objective as the optimization error of the biased objective plus a bias term:

$$\begin{aligned}
F(\mathbf{w}) - F^* &\leq \frac{1}{2\mu} \|\nabla F(\mathbf{w})\|^2 \leq \frac{L^2}{2\mu} \|\mathbf{w} - \mathbf{w}^*\|^2 \\
&\leq \frac{L^2}{\mu} (\|\mathbf{w} - \mathbf{w}_B^*\|^2 + \|\mathbf{w}_B^* - \mathbf{w}^*\|^2) \\
&\leq \underbrace{\frac{2L^2}{\mu^2} (F_B(\mathbf{w}) - F_B^*)}_{:=\epsilon_{\text{opt}}} + \underbrace{\frac{2L^2}{\mu^2} (F(\mathbf{w}_B^*) - F^*)}_{:=\epsilon_{\text{bias}}},
\end{aligned}$$

where the inequalities follow from Assumptions 3, 4.

We study ϵ_{opt} in Theorem 2. Let us now focus on ϵ_{bias} :

$$\begin{aligned}
\|\nabla F(\mathbf{w}_B^*)\| &\stackrel{(a)}{\leq} \left\| \sum_{k=1}^N (\alpha_k - p_k) \nabla F_k(\mathbf{w}_B^*) \right\| \\
&\stackrel{(b)}{\leq} L \sum_{k=1}^N |\alpha_k - p_k| \|\mathbf{w}_B^* - \mathbf{w}_k^*\| \\
&\stackrel{(c)}{\leq} L \sqrt{\frac{2}{\mu}} \sum_{k=1}^N \frac{|\alpha_k - p_k|}{\sqrt{p_k}} \sqrt{p_k (F_k(\mathbf{w}_B^*) - F_k^*)},
\end{aligned} \tag{20}$$

where (a) uses $\nabla F_B(\mathbf{w}_B^*)=0$; (b) and (c) apply L -smoothness and μ -strong convexity, respectively. By direct calculations:

$$\begin{aligned}
\|\nabla F(\mathbf{w}_B^*)\|^2 &\leq \frac{2L^2}{\mu} \left(\sum_{k=1}^N \frac{|\alpha_k - p_k|}{\sqrt{p_k}} \sqrt{p_k (F_k(\mathbf{w}_B^*) - F_k^*)} \right)^2 \\
&\stackrel{(d)}{\leq} \frac{2L^2}{\mu} \left(\sum_{k=1}^N \frac{(\alpha_k - p_k)^2}{p_k} \right) \left(\sum_{k=1}^N p_k (F_k(\mathbf{w}_B^*) - F_k^*) \right) \\
&\leq \frac{2L^2}{\mu} \chi_{\alpha}^2 \Gamma,
\end{aligned}$$

where (d) uses Cauchy-Schwarz. By strong convexity of F :

$$F(\mathbf{w}_B^*) - F^* \leq \frac{1}{2\mu} \|\nabla F(\mathbf{w}_B^*)\|^2 \leq \frac{L^2}{\mu^2} \chi_{\alpha}^2 \Gamma. \quad \square$$

B. Proof of Theorem 2

1) *Additional notation:* let $\mathbf{w}_{t,j}^k$ be the model parameter computed by device k at the global round t , local iteration j . We define $g_t(A_t) = \sum_{k \in A_t} q_k \sum_{j=0}^{E-1} \nabla F_k(\mathbf{w}_{t,j}^k, \xi_{t,j}^k)$, and $\bar{g}_t(A_t) = \mathbb{E}_{\xi|A_t}[g_t(A_t)]$. The update rule of CA-Fed is:

$$\mathbf{w}_{t+1,0} = \mathbf{Proj}_W(\mathbf{w}_{t,0} - \eta_t g_t(A_t)). \tag{21}$$

2) *Key lemmas and results:* we provide useful lemmas and results to support the proof of the main theorem.

Proof of Lemma 1. The boundedness of W gives a bound on $(\mathbf{w}_{t,0})_{t \geq 0}$ based on (2) and (3). From the convexity of F_k , $D := \sup_{\mathbf{w} \in W, k \in \mathcal{K}} \|\nabla F_k(\mathbf{w})\| < +\infty$. Items (6), (8) follow directly; item (7) follows from (6) and Assumption 5. \square

Lemma 2 (Convergence under heterogeneous client availability). *Let Assumptions 3-5 hold. If $\eta_t \leq \frac{1}{2L(EQ+1)}$, we have:*

$$\begin{aligned}
\sum_t \eta_t \mathbb{E}[\sum_{k \in A_t} q_k (F_k(\mathbf{w}_{t,0}) - F_k(\mathbf{w}_B^*))] &\leq \\
&+ \frac{2}{E} \|\mathbf{w}_{0,0} - \mathbf{w}_B^*\|^2 + 2 \sum_{k=1}^N \pi_k q_k^2 \sigma_k^2 \sum_t \eta_t^2 \\
&+ \frac{2}{3} \sum_{k=1}^N \pi_k q_k (E-1)(2E-1)G^2 \sum_t \eta_t^2 \\
&+ 2L(EQ+2) \sum_{k=1}^N \pi_k q_k \Gamma \sum_t \eta_t^2 := C_1 < +\infty.
\end{aligned}$$

Proof of Lemma 2.

$$\begin{aligned}
\|\mathbf{w}_{t+1,0} - \mathbf{w}_B^*\|^2 &= \|\mathbf{Proj}_W(\mathbf{w}_{t,0} - \eta_t g_t) - \mathbf{Proj}_W(\mathbf{w}_B^*)\|^2 \\
&\leq \|\mathbf{w}_{t,0} - \eta_t g_t - \mathbf{w}_B^* + \eta_t \bar{g}_t - \eta_t \bar{g}_t\|^2 = A_1 + A_2 + A_3,
\end{aligned}$$

where $A_1 = \|\mathbf{w}_{t,0} - \mathbf{w}_B^* - \eta_t \bar{g}_t\|^2$, $A_2 = 2\eta_t \langle \mathbf{w}_{t,0} - \mathbf{w}_B^* - \eta_t \bar{g}_t, \bar{g}_t - g_t \rangle$, $A_3 = \eta_t^2 \|g_t - \bar{g}_t\|^2$. Note $\mathbb{E}[A_2] = 0$. We bound A_1, A_3 using the following key steps in [23]: (1) the variance of g_t is bounded if the variance of the stochastic gradients at each device is bounded; (2) the distance of the local model $\mathbf{w}_{t,E}^k$ from the global model $\mathbf{w}_{t,0}$ is bounded since the expected squared norm of the stochastic gradients is bounded. \square

Lemma 3 (Optimization error after \mathcal{J}_t steps). *Let Assumptions 1, 2, 4 hold, and D, H be defined as in (6), (8). Then:*

$$\begin{aligned}
\sum_t \eta_t \mathbb{E}[\sum_{k \in A_t} q_k (F_k(\mathbf{w}_{t-\mathcal{J}_t,0}) - F_k(\mathbf{w}_{t,0}))] &\leq \\
&\leq EDGQ \sum_t \mathcal{J}_t \eta_t^2 \sum_{k=1}^N \pi_k q_k := \frac{C_3}{\ln(1/\lambda(\mathbf{P}))} < +\infty.
\end{aligned}$$

For the proof of Lemma 3, we introduce the following results:

$$|F_k(\mathbf{v}) - F_k(\mathbf{w})| \leq D \cdot \|\mathbf{v} - \mathbf{w}\|, \quad \forall \mathbf{v}, \mathbf{w} \in W, \tag{22}$$

$$\mathbb{E}_{\mathcal{B}_{t,0}^k, \dots, \mathcal{B}_{t,E-1}^k} \|\mathbf{w}_{t+1,0} - \mathbf{w}_{t,0}\| \leq \eta_t GE (\sum_{k \in A_t} q_k). \tag{23}$$

Equation (22) is proven in [15, Proposition 1]: it is a direct application of the mean value theorem. Equation (23) can be proven using equation (7) and the update rule in (21).

Proof of Lemma 3.

$$\begin{aligned}
& \sum_t \eta_t \mathbb{E}[\sum_{k \in A_t} q_k (F_k(\mathbf{w}_{t-\mathcal{J}_t,0}) - F_k(\mathbf{w}_{t,0}))] \\
& \stackrel{(a)}{\leq} D \sum_t \eta_t \mathbb{E}[\sum_{k \in A_t} q_k \mathbb{E}_B \|\mathbf{w}_{t-\mathcal{J}_t,0} - \mathbf{w}_{t,0}\|] \\
& \stackrel{(b)}{\leq} D \sum_t \eta_t \sum_{d=t-\mathcal{J}_t}^{t-1} \mathbb{E}[\sum_{k \in A_t} q_k \mathbb{E}_B \|\mathbf{w}_{d,0} - \mathbf{w}_{d+1,0}\|] \\
& \stackrel{(c)}{\leq} EDG \sum_t \sum_{d=t-\mathcal{J}_t}^{t-1} \eta_t \eta_d \mathbb{E}[\sum_{k \in A_t} q_k \sum_{k' \in A_d} q_{k'}] \\
& \stackrel{(d)}{\leq} \frac{EDG}{2} \sum_t \sum_{d=t-\mathcal{J}_t}^{t-1} (\eta_t^2 + \eta_d^2) \mathbb{E}[\sum_{k \in A_t} q_k \sum_{k' \in A_d} q_{k'}] \\
& \stackrel{(e)}{\leq} EDGQ \sum_t \mathcal{J}_t \eta_t^2 \sum_{k=1}^N \pi_k q_k := \frac{C_3}{\ln(1/\lambda(\mathbf{P}))},
\end{aligned}$$

where (a) uses (22); (b) uses triangle inequality; (c) uses (23); (d) uses Cauchy–Schwarz; (e) uses $\eta_t < \eta_d \leq \eta_{t-\mathcal{J}_t}$. \square

3) *Core of the proof.* The proof consists in two main steps:

1. $\sum_t \eta_t \sum_{k=1}^N \pi_k q_k \mathbb{E}[F_B(\mathbf{w}_{t-\mathcal{J}_t,0}) - F_B^*] \leq C_2 + \frac{C_3}{\ln(1/\lambda(\mathbf{P}))}$;
2. $\sum_t \eta_t \sum_{k=1}^N \pi_k q_k \mathbb{E}[F_B(\mathbf{w}_{t,0}) - F_B(\mathbf{w}_{t-\mathcal{J}_t,0})] \leq \frac{C_3}{\ln(1/\lambda(\mathbf{P}))}$.

Step 1. Combining Lemma 2 and 3, we get:

$$\sum_t \eta_t \mathbb{E}[\sum_{k \in A_t} q_k (F_k(\mathbf{w}_{t-\mathcal{J}_t,0}) - F_k(\mathbf{w}_B^*))] \leq C_1 + \frac{C_3}{\ln(1/\lambda(\mathbf{P}))}.$$

From the definition of \mathcal{J}_t in Theorem 2 and (5), it follows:

$$|[P^{\mathcal{J}_t}]_{i,j} - \pi_j| \leq C_P \lambda(\mathbf{P})^{\mathcal{J}_t} \leq \frac{1}{2Ht}, \quad \forall i, j \in [M]. \quad (24)$$

Assume $t \geq T_P$. We derive an important lower bound:

$$\begin{aligned}
& \mathbb{E}_{A_t|A_{t-\mathcal{J}_t}}[\sum_{k \in A_t} q_k (F_k(\mathbf{w}_{t-\mathcal{J}_t,0}) - F_k(\mathbf{w}_B^*))] \\
& \stackrel{(a)}{=} \sum_{\mathcal{I}=1}^M \mathbb{P}(A_t=\mathcal{I}|A_{t-\mathcal{J}_t}) \sum_{k \in \mathcal{I}} q_k (F_k(\mathbf{w}_{t-\mathcal{J}_t,0}) - F_k(\mathbf{w}_B^*)) \\
& \stackrel{(b)}{=} \sum_{\mathcal{I}=1}^M [P^{\mathcal{J}_t}]_{A_{t-\mathcal{J}_t}, \mathcal{I}} \sum_{k \in \mathcal{I}} q_k (F_k(\mathbf{w}_{t-\mathcal{J}_t,0}) - F_k(\mathbf{w}_B^*)) \\
& \stackrel{(c)}{\geq} \sum_{\mathcal{I}=1}^M (\pi(\mathcal{I}) - \frac{1}{2Ht}) \sum_{k \in \mathcal{I}} q_k (F_k(\mathbf{w}_{t-\mathcal{J}_t,0}) - F_k(\mathbf{w}_B^*)) \\
& \stackrel{(d)}{\geq} (\sum_{k=1}^N \pi_k q_k) \cdot (F_B(\mathbf{w}_{t-\mathcal{J}_t,0}) - F_B^*) - \frac{1}{2t} MQ, \quad (25)
\end{aligned}$$

where (a) is the definition of the conditional expectation, (b) applies the Markov property, (c) follows from (24), and (d) is due to (8). Taking total expectations:

$$\begin{aligned}
& (\sum_{k=1}^N \pi_k q_k) \sum_t \eta_t \mathbb{E}[F_B(\mathbf{w}_{t-\mathcal{J}_t,0}) - F_B^*] \\
& \leq \sum_t \eta_t \mathbb{E}[\sum_{k \in A_t} q_k (F_k(\mathbf{w}_{t-\mathcal{J}_t,0}) - F_k(\mathbf{w}_B^*))] \\
& \quad + \frac{1}{4} MQ \sum_t (\eta_t^2 + \frac{1}{t^2}) = C_2 + \frac{C_3}{\ln(1/\lambda(\mathbf{P}))}, \quad (26)
\end{aligned}$$

where $C_2 = C_1 + \frac{1}{4} MQ \sum_t (\eta_t^2 + \frac{1}{t^2})$.

Step 2. By direct calculation (similar to Lemma 3):

$$(\sum_{k=1}^N \pi_k q_k) \sum_t \eta_t \mathbb{E}[F_B(\mathbf{w}_{t,0}) - F_B(\mathbf{w}_{t-\mathcal{J}_t,0})] \leq \frac{C_3}{\ln(1/\lambda(\mathbf{P}))}.$$

Summing Step 1 and 2, and applying Jensen's inequality:

$$\begin{aligned}
& (\sum_{t=1}^T \eta_t) (\sum_{k=1}^N \pi_k q_k) \mathbb{E}[F_B(\bar{\mathbf{w}}_{T,0}) - F_B^*] \leq \\
& (\sum_{k=1}^N \pi_k q_k) \sum_{t=1}^T \eta_t \mathbb{E}[F_B(\mathbf{w}_{t,0}) - F_B^*] \leq C_2 + \frac{2C_3}{\ln(1/\lambda(\mathbf{P}))},
\end{aligned}$$

where $\bar{\mathbf{w}}_{T,0} := \frac{\sum_{t=1}^T \eta_t \mathbf{w}_{t,0}}{\sum_{t=1}^T \eta_t}$, and the constants are in (12). \square

C. Proof of Theorem 3

It follows the same lines of Theorem 1, developing (20) as:

$$\begin{aligned}
\|\nabla F(\mathbf{w}_B^*)\| & \leq L \sqrt{\frac{2}{\mu}} \sum_{k=1}^N |\alpha_k - p_k| \sqrt{(F_k(\mathbf{w}_B^*) - F_k^*)} \\
& \leq 2L \sqrt{\frac{2}{\mu}} d_{TV}(\boldsymbol{\alpha}, \mathbf{p}) \sqrt{\Gamma}. \quad \square
\end{aligned}$$

D. Minimizing ϵ_{opt}

Equation 12 defines the following optimization problem:

$$\min_{\mathbf{q}} f(\mathbf{q}) = \frac{\frac{1}{2} \mathbf{q}^T \mathbf{A} \mathbf{q} + B}{\boldsymbol{\pi}^T \mathbf{q}} + C; \quad \text{dom } f = \{\mathbf{q} | \mathbf{q} \geq 0, \boldsymbol{\pi}^T \mathbf{q} > 0, \|\mathbf{q}\|_1 = Q\}.$$

Let us rewrite the problem by adding a variable $s := 1/\boldsymbol{\pi}^T \mathbf{q}$ and then replacing $\mathbf{y} := s\mathbf{q}$. Note that the objective function is the perspective of a convex function, and is therefore convex:

$$\min_{\mathbf{y}, s} f(\mathbf{y}, s) = \frac{1}{2s} \mathbf{y}^T \mathbf{A} \mathbf{y} + Bs + C \quad (27a)$$

$$\text{s.t. } \mathbf{y} \geq 0, s > 0, \boldsymbol{\pi}^T \mathbf{y} = 1, \|\mathbf{y}\|_1 = Qs. \quad (27b)$$

The Lagrangian function \mathcal{L} is as follows:

$$\begin{aligned}
\mathcal{L}(\mathbf{y}, s, \lambda, \theta, \boldsymbol{\mu}) & = \frac{1}{2s} \mathbf{y}^T \mathbf{A} \mathbf{y} + Bs + C + \\
& + \lambda(1 - \boldsymbol{\pi}^T \mathbf{y}) + \theta(\|\mathbf{y}\|_1 - Qs) - \boldsymbol{\mu}^T \mathbf{y}. \quad (28)
\end{aligned}$$

Since the constraint $s > 0$ defines an open set, the set defined by the constraints in (27b) is not closed. However, the solution is never on the boundary $s = 0$ because $\mathcal{L}^* \rightarrow +\infty$ as $s \rightarrow 0^+$, and we can consider $s \geq 0$. The KKT conditions for y_k^* read:

$$\text{if } y_k^* > 0: y_k^* = \frac{s^*}{A[kk]} (\lambda^* \pi_k - \theta^*); \quad y_k^* = 0 \text{ otherwise.} \quad (29)$$

Since $\lambda^* \geq 0$, the clients with smaller π_k may have $q_k^* = 0$.

E. Convexity of $\epsilon_{opt} + \epsilon_{bias}$

In Appendix D, we proved that $\epsilon_{opt}(\mathbf{q})$ is convex. To prove that $\epsilon_{bias}(\mathbf{q})$ is also convex, we need to study the convexity of $\chi_{\boldsymbol{\alpha}|\mathbf{p}}^2 = \sum_{k=1}^N (f_k \circ g_k)(\mathbf{q})$, where $f_k(p_k) = (p_k - \alpha_k)^2/p_k$, and $g_k(\mathbf{q}) = (\pi_k q_k) / \sum_{h=1}^N \pi_h q_h$. We observe that $f_k(p_k)$ is convex, and $g_k(\mathbf{q})$ is a particular case of linear-fractional function [38]. By direct inspection, it can be proved that $(f_k \circ g_k)(\mathbf{q})$ is convex in $\text{dom}(f_k \circ g_k) = \{\mathbf{q} : \|\mathbf{q}\|_1 = Q > 0\}$.

F. Synthetic dataset

Our synthetic datasets has been generated as follows:

- 1) For client $k \in \mathcal{K}$, sample group identity i_k from a Bernoulli distribution of parameter $1/2$;
- 2) Sample model parameters $\mathbf{w}^* \sim \mathcal{N}(0, I_d)$ from the d -dimensional normal distribution;
- 3) For client $k \in \mathcal{K}$ and sample index $j \in \{1, \dots, 150\}$, sample clients input data $\mathbf{x}_k^{(j)} \sim \mathcal{N}(0, I_d)$ from the d -dimensional normal distribution;
- 4) For client $k \in \mathcal{K}$ such that $i_k = 0$ and sample index $j \in \{1, \dots, 150\}$, sample the true labels $y_k^{(j)}$ from a Bernoulli distribution with parameter equal to $\text{sigmoid}(\langle \mathbf{w}^*, \mathbf{x}_k^{(j)} \rangle)$;
- 5) For client $k \in \mathcal{K}$ such that $i_k = 1$ and sample index $j \in \{1, \dots, 150\}$, sample the true labels $y_k^{(j)}$ from a Bernoulli distribution with parameter equal to $0.8 \cdot \text{sigmoid}(\langle \mathbf{w}^*, \mathbf{x}_k^{(j)} \rangle) + 0.2 \cdot (1 - \text{sigmoid}(\langle \mathbf{w}^*, \mathbf{x}_k^{(j)} \rangle))$.

REFERENCES

- [1] A. Rodio, F. Faticanti, O. Marfoq, G. Neglia, and E. Leonardi, "Federated Learning under Heterogeneous and Correlated Client Availability," 2023. [Online]. Available: <https://arxiv.org/abs/2301.04632>
- [2] J. Verbraeken, M. Wolting, J. Katzy, J. Kloppenburg, T. Verbelen, and J. S. Rellermeyer, "A survey on distributed machine learning," *ACM Computing Surveys (CSUR)*, vol. 53, no. 2, pp. 1–33, 2020.
- [3] S. Wang, T. Tuor, T. Salonidis, K. K. Leung, C. Makaya, T. He, and K. Chan, "When edge meets learning: Adaptive control for resource-constrained distributed machine learning," in *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*. IEEE, 2018, pp. 63–71.
- [4] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," *arXiv preprint arXiv:1610.05492*, 2016.
- [5] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial Intelligence and Statistics*. PMLR, 2017, pp. 1273–1282.
- [6] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings *et al.*, "Advances and open problems in federated learning," *Foundations and Trends® in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2021.
- [7] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, "Federated learning: Challenges, methods, and future directions," *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 50–60, 2020.
- [8] H. Eichner, T. Koren, B. McMahan, N. Srebro, and K. Talwar, "Semi-cyclic stochastic gradient descent," in *International Conference on Machine Learning*. PMLR, 2019, pp. 1764–1773.
- [9] J. Wang, Z. Charles, Z. Xu, G. Joshi, H. B. McMahan, M. Al-Shedivat, G. Andrew, S. Avestimehr, K. Daly, D. Data *et al.*, "A field guide to federated optimization," *arXiv preprint arXiv:2107.06917*, 2021.
- [10] K. Bonawitz, H. Eichner, W. Grieskamp, D. Huba, A. Ingerman, V. Ivanov, C. Kiddon, J. Konečný, S. Mazzocchi, B. McMahan *et al.*, "Towards federated learning at scale: System design," *Proceedings of Machine Learning and Systems*, vol. 1, pp. 374–388, 2019.
- [11] Y. Ding, C. Niu, Y. Yan, Z. Zheng, F. Wu, G. Chen, S. Tang, and R. Jia, "Distributed optimization over block-cyclic data," *arXiv preprint arXiv:2002.07454*, 2020.
- [12] C. Zhu, Z. Xu, M. Chen, J. Konečný, A. Hard, and T. Goldstein, "Diurnal or Nocturnal? Federated Learning from Periodically Shifting Distributions," in *NeurIPS 2021 Workshop on Distribution Shifts: Connecting Methods and Applications*, 2021.
- [13] T. T. Doan, "Local stochastic approximation: A unified view of federated learning and distributed multi-task reinforcement learning algorithms," *arXiv preprint arXiv:2006.13460*, 2020.
- [14] T. T. Doan, L. M. Nguyen, N. H. Pham, and J. Romberg, "Convergence rates of accelerated Markov gradient descent with applications in reinforcement learning," *arXiv preprint arXiv:2002.02873*, 2020.
- [15] T. Sun, Y. Sun, and W. Yin, "On Markov chain gradient descent," *Advances in neural information processing systems*, vol. 31, 2018.
- [16] M. McCloskey and N. J. Cohen, "Catastrophic Interference in Connectionist Networks: The Sequential Learning Problem," in *Psychology of Learning and Motivation*, G. H. Bower, Ed. Academic Press, 1989, vol. 24, pp. 109–165.
- [17] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Ramalho, A. Grabska-Barwinska *et al.*, "Overcoming catastrophic forgetting in neural networks," *Proceedings of the National Academy of Sciences*, vol. 114, no. 13, pp. 3521–3526, 2017.
- [18] M. Tang, X. Ning, Y. Wang, J. Sun, Y. Wang, H. Li, and Y. Chen, "FedCor: Correlation-Based Active Client Selection Strategy for Heterogeneous Federated Learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [19] M. Ribero, H. Vikalo, and G. De Veciana, "Federated Learning Under Intermittent Client Availability and Time-Varying Communication Constraints," *arXiv preprint arXiv:2205.06730*, 2022.
- [20] L. Tan, X. Zhang, Y. Zhou, X. Che, M. Hu, X. Chen, and D. Wu, "AdaFed: Optimizing Participation-Aware Federated Learning with Adaptive Aggregation Weights," *IEEE Transactions on Network Science and Engineering*, 2022.
- [21] A. Nichol, J. Achiam, and J. Schulman, "On first-order meta-learning algorithms," *arXiv preprint arXiv:1803.02999*, 2018.
- [22] S. J. Reddi, Z. Charles, M. Zaheer, Z. Garrett, K. Rush, J. Konečný, S. Kumar, and H. B. McMahan, "Adaptive Federated Optimization," in *International Conference on Learning Representations*, 2021.
- [23] X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang, "On the Convergence of FedAvg on Non-IID Data," in *International Conference on Learning Representations*, 2019.
- [24] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," *Proceedings of Machine Learning and Systems*, vol. 2, pp. 429–450, 2020.
- [25] W. Chen, S. Horvath, and P. Richtárik, "Optimal client sampling for federated learning," *arXiv preprint arXiv:2010.13723*, 2020.
- [26] Y. Fraboni, R. Vidal, L. Kameni, and M. Lorenzi, "Clustered sampling: Low-variance and improved representativity for clients selection in federated learning," in *International Conference on Machine Learning*. PMLR, 2021, pp. 3407–3416.
- [27] Y. Jee Cho, J. Wang, and G. Joshi, "Towards Understanding Biased Client Selection in Federated Learning," in *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*. PMLR, 2022, pp. 10 351–10 375.
- [28] T. T. Doan, L. M. Nguyen, N. H. Pham, and J. Romberg, "Finite-time analysis of stochastic gradient descent under Markov randomness," *arXiv preprint arXiv:2003.10973*, 2020.
- [29] A. Meyers and H. Yang, "Markov Chains for Fault-Tolerance Modeling of Stochastic Networks," *IEEE Transactions on Automation Science and Engineering*, 2021.
- [30] H. Olle, P. Yuval, and E. S. Jeffrey, "Dynamical percolation," in *Annales de l'Institut Henri Poincaré (B) Probability and Statistics*, vol. 33, no. 4. Elsevier, 1997, pp. 497–528.
- [31] D. A. Levin and Y. Peres, *Markov chains and mixing times*. American Mathematical Soc., 2017, vol. 107.
- [32] L. Bottou, F. E. Curtis, and J. Nocedal, "Optimization methods for large-scale machine learning," *Siam Review*, vol. 60, no. 2, pp. 223–311, 2018.
- [33] J. Wang, Q. Liu, H. Liang, G. Joshi, and H. V. Poor, "Tackling the objective inconsistency problem in heterogeneous federated optimization," *Advances in neural information processing systems*, 2020.
- [34] I. J. Goodfellow, M. Mirza, D. Xiao, A. Courville, and Y. Bengio, "An empirical investigation of catastrophic forgetting in gradient-based neural networks," *arXiv preprint arXiv:1312.6211*, 2013.
- [35] R. Kemker, M. McClure, A. Abitino, T. Hayes, and C. Kanan, "Measuring catastrophic forgetting in neural networks," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.
- [36] Y. LeCun and C. Cortes, "MNIST handwritten digit database," 2010.
- [37] F. Sattler, K.-R. Müller, and W. Samek, "Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 8, pp. 3710–3722, 2020.
- [38] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.