

# Communication-Efficient Device Scheduling for Federated Learning Using Lyapunov Optimization

Jake B. Perazzone<sup>ID</sup>, Shiqiang Wang<sup>ID</sup>, Senior Member, IEEE, Mingyue Ji<sup>ID</sup>, Member, IEEE, and Kevin Chan<sup>ID</sup>, Senior Member, IEEE

**Abstract**— Federated learning (FL) is a useful tool that enables the training of machine learning models over distributed data without having to collect data centrally. When deploying FL in constrained wireless environments, however, intermittent connectivity of devices, heterogeneous connection quality, and non-i.i.d. data can severely slow convergence. In this paper, we consider FL with arbitrary device participation probabilities for each round and show that by weighing each device’s update by the reciprocal of their per-round participation probability, we can guarantee convergence to a stationary point. Our bound applies to non-convex loss functions and non-i.i.d. datasets and recovers state-of-the-art convergence rates for both full and uniform partial participation, including linear speedup, with only a single-sided learning rate. Then, using the derived convergence bound, we develop a new online client selection and power allocation algorithm that utilizes the Lyapunov drift-plus-penalty framework to opportunistically minimize a function of the convergence bound and the average communication time under a transmit power constraint. We use optimization over manifold techniques to obtain a solution to the minimization problem. Thanks to the Lyapunov framework, one key feature of the algorithm is that knowledge of the channel distribution is not required and only the instantaneous channel state information needs to be known. Using the CIFAR-10 dataset with varying levels of data heterogeneity, we show through simulations that the communication time can be significantly decreased using our algorithm compared to uniformly random participation, especially for heterogeneous channel conditions.

**Index Terms**— Federated learning, Lyapunov stochastic optimization, client selection.

Received 8 May 2024; accepted 3 December 2024; approved by IEEE TRANSACTIONS ON NETWORKING Editor G. Joshi. This work was supported in part by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement W911NF-16-3-0001. The work of Mingyue Ji was supported in part by NSF CAREER under Award 2145835 and in part by NSF under Award 2312227. A preliminary version of this paper titled “Communication-Efficient Device Scheduling for Federated Learning Using Stochastic Optimization” was presented in the IEEE International Conference on Computer Communications (INFOCOM), 2022 [DOI: 10.1109/INFOCOM48880.2022.9796818]. (Corresponding author: Jake B. Perazzone.)

Jake B. Perazzone and Kevin Chan are with the DEVCOM Army Research Laboratory, Adelphi, MD 20783 USA (e-mail: jake.b.perazzone.civ@army.mil; kevin.s.chan.civ@army.mil).

Shiqiang Wang is with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA (e-mail: shiqiang.wang@ieee.org).

Mingyue Ji is with the Electrical and Computer Engineering Department, University of Utah, Salt Lake City, UT 84112 USA (e-mail: mingyueji@ufl.edu).

Digital Object Identifier 10.1109/TON.2025.3539857

## I. INTRODUCTION

FEDERATED learning (FL) enables the training of machine learning (ML) models over decentralized data. Instead of transmitting data from all devices to a centralized location, model training is accomplished through a collaborative procedure in which the participants train on their own locally collected datasets and periodically share their model parameters. This ML technique is immensely powerful for protecting the privacy of the users’ data since raw data never leaves the devices. Furthermore, two other key features of FL, partial device participation and multiple local iterations, help reduce the communication burden by communicating less information less often compared to centralized learning or other distributed training techniques. Accordingly, FL is most advantageous in instances where full data transmission is infeasible or otherwise restricted in some way.

In this paper, we consider the scenario where training is coordinated by a central aggregator that communicates with each device over a wireless network. The aggregator is responsible for choosing devices, accumulating their models, and disseminating the aggregated global model back to the devices. We use device and node interchangeably throughout the paper. A block diagram of the wireless network running FL can be found in Fig. 1. Each node  $n$  is a device that has a unique dataset and an uplink channel to the aggregator with channel gain  $h_n^t$  at time  $t$ . The orchestration of FL over large-scale wireless networks such as this is a challenging task due to the dynamics of the channel. In mobile edge computing (MEC) environments, for example, poor channel quality and intermittent connectivity can completely derail training. In the original FL algorithm, *FedAvg* [2], devices are selected uniformly at random in each round. Although this strategy has been shown to converge [3],[4],[5], in practice, it may lead to poor performance since it is completely agnostic to the communication medium as well as many other practical factors. For example, if devices with poor channels are naively chosen, the model will converge much slower in terms of wall-clock time since it will take longer to communicate the parameters in each aggregation round. On the other hand, if such devices are ignored, convergence will also be negatively affected since those devices’ data will not be used. The end result is the consumption of more network resources than is necessary due

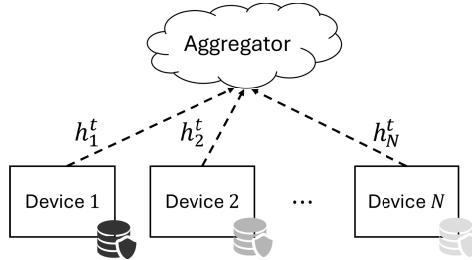


Fig. 1. Block diagram of the uplink communication in federated learning over a wireless network.

to longer training times. Thus, a more intelligent approach to device selection is needed to optimize network resource consumption and minimize wall-clock convergence time.

Before designing such an approach, though, the effect of arbitrary device selection probabilities on convergence must be understood to ensure convergence to a good model. Therefore, in this paper, we first derive a convergence bound for non-convex loss functions with arbitrary device selection probabilities for each round and do not assume uniform bounded gradients as is sometimes done, e.g., [1], [3], [6], and [7]. In fact, we show that we can achieve linear speedup [8] with only a single-sided learning rate. We then use that bound to inform the development of a device selection and power allocation policy that greatly speeds up wall-clock convergence time. More specifically, our upper bound shows that as long as all devices have a non-zero probability of participating in each round, then FL will converge in expectation to a stationary point of the loss function. We then use the knowledge of how the selection probabilities affect the convergence bound to formulate a stochastic optimization problem that determines the optimal transmit powers and selection probabilities for a nonuniform sampling with replacement selection strategy. Optimality in this case means it minimizes a weighted sum of the convergence bound and the time spent communicating model parameters while satisfying constraints on the peak and time average transmit power of each device. The form of the convergence bound and our novel problem formulation allow us to utilize the Lyapunov drift-plus-penalty framework to solve the problem in an online and greedy fashion with optimality and constraint satisfaction guarantees. A key advantage of our new device selection algorithm is that it is able to make decisions according to current channel conditions without knowledge of the underlying channel statistics due to the structure of the Lyapunov framework.

To show the performance of our algorithm, we run numerous experiments on the CIFAR-10 dataset with varying levels of data and channel heterogeneity to demonstrate the saved training time using our developed algorithm. We compare our results to the uniform selection policy of *FedAvg* and show that the time required to reach a target accuracy can be sped up by up to 8.5x. In summary, our main contributions are as follows:

- 1) We derive an upper bound for the convergence of non-convex loss functions using FL with arbitrary selection probabilities and no uniform bounded gradient assumption. We recover state-of-the-art convergence rates, including linear speedup with respect to the number

of devices participating, with only a single-sided learning rate.

- 2) We formulate a novel stochastic optimization problem that minimizes a weighted sum of the convergence bound and the amount of communication time spent on transmitting parameter updates, while satisfying transmit power constraints.
- 3) Using the Lyapunov drift-plus-penalty framework, we derive an optimal online greedy solution that does not require knowledge of the channel statistics. We solve the problem by decoupling device selection from power allocation to obtain a combination of an analytical solution and a Riemannian manifold optimization problem that requires a solver due to a non-convex constraint.
- 4) We provide experimental results that demonstrate a communication savings of up to 8.5 $\times$  compared to traditional uniform selection strategies and examine numerous scenarios involving varying levels of heterogeneity in both the data and the wireless channels.

The rest of the paper is organized as follows. First, we present some related work in Section II before formally presenting our problem formulation in Section III. Then, convergence analysis is provided in IV and the device scheduling policy is developed in Section V. Finally, we present experimental results in Section VI.

## II. RELATED WORKS

Since its introduction via the *FedAvg* algorithm in [2], FL has garnered a lot of attention in both industry and academia. Drawn by the promise of privacy, much work has been done to formally guarantee privacy [9], [10], [11], characterize convergence [3], [4], [5], [6], [12], and minimize communication overhead [13]. For the latter, some strategies and analyses utilize model compression via sparsification [14], [15] and quantization [16], [17], [18] to decrease the communication load while others have focused on optimizing device participation [19], [20], [21], for example. Naively applying these approaches, e.g., biased compression or device selection, however, can lead to poor performance due to model skew which causes the model to drift away from the globally optimal solution. The presence of heterogeneity in the system is the main cause of the skew which inhibits successful implementation of communication-efficient techniques like device selection at scale if not adequately addressed. Heterogeneity can come in the form of differences in network and computational resources as well as from non-identical data distributions across the devices. Overly relying on high-performing nodes during training is one way in which model skew can be unintentionally induced, resulting in a final model that works well for those nodes, but not for the lesser nodes whose data is ignored.

In fact, one of the first works specifically targeting device selection [19] suffers from this problem. In this study, they report poor performance, especially for non-i.i.d. datasets, since they only aggregate models from devices that respond the quickest during each training round. This underlines the importance of accounting for and guaranteeing convergence

in device scheduling algorithms. Some other empirical studies without convergence results include [22], [23]. When dealing with heterogeneity, the challenge is how to quickly and sufficiently learn from all nodes to converge to a good global model without straining resources. Therefore, in this paper, we first focus on ensuring convergence of the global loss function before designing our device selection and power allocation algorithm.

Most theoretical works on the convergence of non-convex loss functions in FL assume uniform device participation via sampling with and without replacement [3],[4],[6],[8] which does not allow for more intelligent device selection. Meanwhile, other papers analyze more sophisticated selection schemes, but often with strict assumptions. Among these is [24], which analyzes the performance of three different scheduling policies but only contains convergence results for simple linear regression tasks. In [20], the authors analyze the convergence of a general, possibly biased, device selection strategy, but only for strongly convex loss functions and uniformly bounded gradients. Importantly, their bound introduces a non-vanishing term to the convergence bound due to selection bias and thus their strategy is not guaranteed to converge to a stationary point of the loss function. Both [25] and [26] also consider convergence, but again only for strongly convex loss functions. Uniquely, though, the work in [26] considers cases where devices may compute incomplete updates, become inactive/unavailable, depart early, or arrive late, but does not provide specific selection strategies. Finally, [27] considers arbitrary participation probabilities for each device, but in their case study, these probabilities are held constant throughout training and are not a design parameter. In fact, the probabilities used in the experiments match the fraction of data samples each node has which follows the traditional *FedAvg* approach. Additionally, in [27], all devices must participate in the first round for convergence. A more recent theoretical treatment of arbitrary participation can be found in [28]. We improve upon these results by considering non-convex loss functions and derive a bound with an easily managed non-vanishing term under the condition that all devices have an arbitrary non-zero probability of participating in each round.

The results most similar to ours can be found in [21]. In the paper, a “clustered” sampling approach is proposed where a multinomial distribution is used to sample devices. This framework allows for the device selection probabilities to change with each draw in a given round. In order to remain unbiased, however, the probabilities must be chosen such that in each round, each node is selected on average proportional to their data ratio. In our approach, by modifying the weights during the global aggregation step, we can maintain unbiased updates while allowing arbitrary participation probabilities at every round without restricting the values. The increased flexibility enables us to choose probabilities depending on the state of the network to increase communication efficiency while still guaranteeing convergence.

As for FL over wireless networks, some works [29], [30] develop frameworks that jointly optimize convergence and communication. Similarly to our approach, they derive a convergence bound and then minimize it by finding the optimal parameter values. For example, in [29], the FL

loss is minimized while meeting the delay and energy consumption requirements via power allocation, user selection, and resource block allocation. Both papers, however, make the unrealistic assumption that the channel remains constant throughout the training process which we do not assume here. In [31], stochastic optimization is used to determine an optimal scheduling and resource block policy that simultaneously minimizes the FL loss function and CSI uncertainties. The loss function considered, though, is simple linear regression and does not readily apply to neural network models. Stochastic optimization is also considered for FL in [32] and [33], but not to design an optimal device selection policy that guarantees convergence of non-convex loss function. In summary, unlike previous work, our approach guarantees convergence of non-convex loss functions with arbitrary participation probabilities and uses the results to develop a device selection and transmit power allocation policy. By leveraging Lyapunov optimization, we provide a partially analytical solution that increases communication efficiency of FL.

### III. PROBLEM FORMULATION

We now explain the FL optimization problem in more detail. Consider a system with  $N$  clients, where each client  $n \in \{1, 2, \dots, N\}$  has a possibly non-convex local objective  $f_n(\mathbf{x})$  with parameter  $\mathbf{x} \in \mathbb{R}^d$ . We would like to solve the following finite-sum problem:

$$\min_{\mathbf{x}} f(\mathbf{x}) := \frac{1}{N} \sum_{n=1}^N f_n(\mathbf{x}). \quad (1)$$

Note that for simplicity, we assume that each node has the same amount of data and thus have uniform weights, i.e.,  $\frac{1}{N}$ , in (1). When this is not the case, it can be easily accounted for by selecting appropriate weights for each local loss function in the sum and carrying it through the analysis.

To solve (1), we follow the typical *FedAvg* [2] FL paradigm, but modify it to allow for arbitrary device participation in each round. In our algorithm, which can be found in Algorithm 1, participation is dictated by  $q_n^t \in (0, 1]$  which is the probability that device  $n$  participates in round  $t$ . To maintain unbiased updates, we weigh each device’s gradient update inversely proportional to the participation probability of that device, as seen in Line 7 of Algorithm 1. Intuitively, this ensures that devices that participate less frequently are weighed more heavily such that they have sufficient influence over the global model when they do participate and vice versa. While knowledge of  $q_n^t$  is required for aggregation, we treat it as a design parameter rather than a property of the environment and thus can conceptually develop an algorithm that chooses its value at each round. We leave cases where  $q_n^t$  is dictated by the environment to future work where it can be estimated. With this in place, we adjust its value for each device and aggregation round to optimize other aspects of the process such as system resource usage. Analyzing the convergence behavior of *FedAvg* with arbitrary probabilities enables the ability to measure the impact of stochastic scheduling on training in order to better design device selection policies.

We now explain Algorithm 1 in further detail. First,  $\mathbb{I}_n^t \in \{0, 1\}$  denotes the random binary variable that represents whether client  $n$  is selected in round  $t$ , where

**Algorithm 1** FedAvg With Client Sampling

---

**Input:**  $\gamma, \mathbf{x}_0, K, T, \{q_n^t\}$   
**Output:**  $\{\mathbf{x}_t\}$

1 **for**  $t \leftarrow 0, \dots, T-1$  **do**

2    Sample  $\mathbb{I}_n^t \sim q_n^t, \forall n$ ;

3    **for**  $n \leftarrow 1, \dots, N$  in parallel **do**

4      $\mathbf{y}_{t,0}^n \leftarrow \mathbf{x}_t$ ;

5     **for**  $i \leftarrow 0, \dots, K-1$  **do**

6        $\mathbf{y}_{t,i+1}^n \leftarrow \mathbf{y}_{t,i}^n - \gamma \mathbf{g}_n(\mathbf{y}_{t,i}^n)$ ;

7      $\mathbf{x}_{t+1} \leftarrow \mathbf{x}_t + \frac{1}{N} \sum_{n=1}^N \frac{\mathbb{I}_n^t}{q_n^t} (\mathbf{y}_{t,K}^n - \mathbf{y}_{t,0}^n)$ ;  
// global parameter update

---

$q_n^t := \Pr\{\mathbb{I}_n^t = 1\}$ . Next, after  $\mathbb{I}_n^t$  is realized and the current global model is broadcasted, each device performs  $K$  rounds of stochastic gradient descent (SGD) on its local dataset where  $\mathbf{g}_n(\mathbf{x})$  denotes the stochastic gradient of  $f_n(\mathbf{x})$  for device  $n$  and  $\gamma > 0$  is the local learning rate. This repeats for  $T$  total aggregation rounds. We use the variable  $y_{t,i}^n$  to represent the intermediate local model updates between global aggregation rounds where  $i$  indexes the local SGD iteration. In the last step, only the selected devices' gradient updates are aggregated. Note that even though this algorithm shows all  $N$  devices performing SGD each round, it is logically equivalent to one in which only the selected clients via  $\mathbb{I}_n^t$  receive the global model, compute gradient updates, and transmit back to the aggregator.

## IV. CONVERGENCE ANALYSIS

In this section, we provide an upper bound on the convergence of (1) using Algorithm 1 for non-convex loss functions. We assume that  $\mathbb{I}_n^t$  and  $\mathbb{I}_{n'}^t$  are independent through time, and that drawing from  $\mathbb{I}_n^t$  is independent of SGD noise, i.e.,  $\mathbb{I}_n^t$  and  $\mathbf{g}_n$  are independent. We also make the following assumptions on the local loss functions.

**Assumption 1** ( $L$ -smoothness).

$$\|\nabla f_n(\mathbf{y}_1) - \nabla f_n(\mathbf{y}_2)\| \leq L\|\mathbf{y}_1 - \mathbf{y}_2\| \quad (2)$$

for any  $\mathbf{y}_1, \mathbf{y}_2$  and some  $L > 0$ .

**Assumption 2** (Unbiased stochastic gradients).

$$\mathbb{E}[\mathbf{g}_n(\mathbf{y})|\mathbf{y}] = \nabla f_n(\mathbf{y}), \quad (3)$$

for any  $\mathbf{y}$ .

**Assumption 3** (Bounded stochastic gradient noise).

$$\mathbb{E}[\|\mathbf{g}_n(\mathbf{y}) - \nabla f_n(\mathbf{y})\|^2] \leq \nu^2, \forall \mathbf{y}, n \quad (4)$$

for some  $\nu > 0$ .

**Assumption 4** (Bounded gradient divergence).

$$\|\nabla f_n(\mathbf{y}) - \nabla f(\mathbf{y})\|^2 \leq \epsilon^2, \forall \mathbf{y}, n \quad (5)$$

for some  $\epsilon > 0$ .

While the first three assumptions are standard in non-convex optimization, the fourth is unique to FL and other distributed SGD techniques, as used in [6], [8], [34], and [35]. It bounds

the differences between the local loss functions across clients due to having non-i.i.d. datasets. The i.i.d. case is recovered when  $\epsilon = 0$ . Now, we state our convergence theorem in Theorem 1.

**Theorem 1.** *Let Assumptions 1–4 hold with  $\gamma, T, K, N$ , and  $q_n^t$  defined as above. Then, if  $\gamma \leq \frac{q_{\min}}{8LK}$ , where we assume the existence of a minimum participation probability  $q_{\min}$  such that such that  $q_{\min} \leq q_n^t$  for all  $n, t$ , Algorithm 1 satisfies*

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[ \|\nabla f(\mathbf{x}_t)\|^2 \right] &\leq \frac{2(\mathbb{E}[f(\mathbf{x}_0)] - \mathbb{E}[f(\mathbf{x}_T)])}{c\gamma KT} \\ &+ \Phi_1 + \frac{\Phi_2}{TN} \sum_{t=0}^{T-1} \frac{1}{N} \sum_{n=1}^N \frac{1}{q_n^t}, \end{aligned} \quad (6)$$

where  $\Phi_1 = \frac{1}{c} 5\gamma^2 K L^2 (\nu^2 + 6K\epsilon^2)$ ,  $\Phi_2 = \frac{2L\gamma\nu^2}{c}$ , and  $c$  is a constant.

*Proof:* The full proof can be found in Appendix A. The proof utilizes the assumption that the client sampling random variable  $\mathbb{I}_n^t$  is independent of the stochastic gradient noise in  $\mathbf{g}_n$  and that  $\mathbb{E}[\mathbb{I}_n^t] = q_n^t$ . Thus, the weighted averaging of  $\frac{1}{q_n^t}$  in Step 7 of Algorithm 1 ensures that the gradient estimate remains unbiased. That is,  $\mathbb{E} \left[ \frac{1}{q_n^t} \mathbf{g}_n(\mathbf{x}) \right] = \nabla f_n(\mathbf{x})$ .

Another key assumption is the existence of a constant  $0 < q_{\min} \leq q_n^t$  for all  $n, t$ . This is reasonable since, otherwise,  $q_n^t \rightarrow 0$  for at least one  $n$ , which would cause the weight in the averaging step to go to infinity, trivially causing divergence. Additionally,  $q_{\min}$  can be viewed as a design parameter, e.g., to control fairness.  $\square$

Our convergence bound consists of three terms. The first vanishes as  $T$  increases and, while the second and third are non-vanishing, they can be managed with appropriate selection of  $\gamma$  (see Corollary 1). The non-vanishing terms are separated into ones that are not amplified by partial participation,  $\Phi_1$ , and ones that are,  $\Phi_2$ . The convergence bound increases when participation probabilities are small. This follows with the intuition that sampling fewer participants in a round leads to a noisier estimate of the full gradient and thus more iterations will be required to converge to a desired loss.

Interestingly, the level of data heterogeneity, as measured by  $\epsilon^2$ , does not amplify the effect of partial participation. This suggests that partial participation affects convergence solely through increasing the variance of the gradient estimate and not by a biasing of the model towards the minima of the more frequent participants. This decoupling is due to the inverse-proportional weighting of participation probabilities in the aggregation step of Algorithm 1 which ensures unbiased gradient estimates. It essentially compensates for infrequent participants by taking larger steps towards their local minima in the low-probability event that they are called upon. Similarly, the upper bound on step size requires that when the minimum participation probability,  $q_{\min}$ , is small, the step size must also be small. This makes sense intuitively since larger steps would also heavily skew the model towards the more frequent participants' local stationary points.

In the bound, the partial participation term can be trivially minimized by setting  $q_n^t = 1$  for all  $n$  and  $t$ , i.e., full parti-

pation, but it is impractical to assume that every device can or will participate in every round. This necessitates a selection strategy that minimizes the time average  $\frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{N} \sum_{n=1}^N \frac{1}{q_n^t}$  under some constraints. Before introducing our approach to this problem, we present two corollaries to further examine convergence behavior and show that linear speedup is achieved.

**Corollary 1.** *If we choose  $\gamma = \min\{\frac{q_{\min}}{8LK}, \frac{\sqrt{N}q_{\min}}{\sqrt{TK}L}\}$ , then we have*

$$\begin{aligned} & \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(\mathbf{x}_t)\|^2] \\ & \leq \mathcal{O} \left( \frac{L}{q_{\min}T} + \frac{L}{\sqrt{TNK}q_{\min}} + \frac{Nq_{\min}\epsilon^2}{TK} + \frac{\sqrt{q_{\min}\nu^2}Q}{\sqrt{TNK}} \right), \end{aligned} \quad (7)$$

where  $Q = \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{N} \sum_{n=1}^N \frac{1}{q_n^t} \leq \frac{1}{q_{\min}}$ .

*Proof:* The order terms can be obtained by plugging in the step size in (6) and noting that  $\max\{x, y\} \leq x + y$ .  $\square$

**Corollary 2.** *In the full participation scenario, i.e.,  $q_n^t = 1 \forall n, t$ , the convergence rate of Algorithm 1 is  $\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(\mathbf{x}_t)\|^2] \leq \mathcal{O} \left( \frac{1}{\sqrt{TNK}} \right)$ .*

*Furthermore, for uniform sampling with and without replacement, the convergence rate of Algorithm 1 is  $\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(\mathbf{x}_t)\|^2] \leq \mathcal{O} \left( \frac{1}{\sqrt{TmK}} \right)$ .*

*Proof:* For the full participation and uniform sampling without replacement scenarios, the corollary can be obtained by noting that  $q_{\min} = q_n^t = 1$  and  $q_{\min} = q_n^t = m/N$  for all  $n, t$ , for each scenario, respectively, and then applying Corollary 1 with  $Q = 1/q_{\min}$ .

For the uniform sampling with replacement scenario, we note that  $q_{\min} = q_n^t = 1 - (1 - 1/N)^m$ . To show the result, we first prove that  $1/q_{\min} = \mathcal{O}(N/m)$  for sampling with replacement and then the result follows from the sampling without replacement strategy.

First, consider the case where  $m \geq N$ ,<sup>1</sup> then we have

$$1 - \left(1 - \frac{1}{N}\right)^m \geq 1 - e^{-m/N} \geq 1 - e^{-1}.$$

Next, consider the case where  $m \leq N$  and let  $r = m/N$  such that  $r \leq 1$ . Observe, then, that the derivative with respect to  $r$  of  $(1 - (1 - r/m)^m)/r$  is always negative and thus is minimized when  $r = 1$ . Then, we have

$$\begin{aligned} \left(1 - \left(1 - \frac{1}{N}\right)^m\right) \frac{N}{m} &= (1 - (1 - r/m)^m)/r \\ &\geq 1 - (1 - 1/m)^m \\ &\geq 1 - e^{-1}, \end{aligned}$$

which gives

$$1 - \left(1 - \frac{1}{N}\right)^m \geq \frac{m}{N} (1 - e^{-1}).$$

<sup>1</sup>In sampling with replacement, since devices can be re-selected each draw, you can continue to sample beyond the number of devices and thus it is possible for  $m \geq N$ . For sampling without replacement,  $m = N$  would result in full participation and no more draws can be made.

Combining the two cases, we have a final bound of

$$1 - \left(1 - \frac{1}{N}\right)^m \geq \frac{\min\{m, N\}}{N} (1 - e^{-1}),$$

and thus,

$$1/q_{\min} = \frac{1}{1 - \left(1 - \frac{1}{N}\right)^m} = \mathcal{O} \left( \frac{N}{m} \right).$$

The result then follows simply via Corollary 1. Finally, we note that if  $m > N$ , then the  $m$  in the order term is replaced with  $N$ .  $\square$

Corollary 2 shows that we recover the state-of-the-art convergence rate in [3] and [4] that indicates a linear speedup in terms of the number of clients for both full *and* partial participation. This is achieved *without the use of a two-sided learning rate*<sup>2</sup> which prior methods require. Linear speedup is desired since the number of iterations required to reach a given loss can be reduced by proportionally increasing the number of clients in a given round. This behavior ensures that our algorithm fully leverages the parallelism of federated learning, even with non-i.i.d. data. For general participation probabilities, the speedup is in terms of  $Nq_{\min}$ . Interestingly, with the choice of step size in Corollary 1, if  $q_{\min} \leq 1/N$ , then speedup is lost.

Now equipped with an understanding of how device selection probability  $q_n^t$  affects FL convergence, we will design an optimization problem that adaptively determines both  $q_n^t$  and transmit power allocations in order to minimize communication overhead while guaranteeing convergence.

## V. COMMUNICATION-EFFICIENT SCHEDULING POLICY

In this section, we propose a novel device selection and transmit power allocation policy that minimizes communication overhead while guaranteeing convergence. The policy is based on solving a stochastic optimization problem that minimizes a time average function consisting of the convergence bound in (6) and the average time spent communicating model updates, while satisfying time average transmission power constraints for each device. This optimization problem is formally stated in Section V-B. The output of the optimization problem is the selection probabilities  $q_n^t$  used in Algorithm 1 and the transmit powers  $P_n^t$  used in each round. The communication of multiple devices' model updates over many rounds may cause a huge bottleneck in FL, especially in heterogeneous environments, which is why we focus on minimizing the time spent during this phase. Unlike previous work, we speed up training in terms of wall-clock time, minimizing the burden on the network, while still theoretically guaranteeing convergence. The natural form of the bound allows us to formulate the problem in terms of time averages. This lends itself perfectly to the application of the Lyapunov drift-plus-penalty framework which specifically deals with this class of problems. A major benefit of the approach is that it does not require knowledge of the exact dynamics or statistics of the channel; only the instantaneous channel state information (CSI) is needed.

<sup>2</sup>Although, our weighted averaging step in (7) may be viewed as a per-device, per-round learning rate that becomes a global learning rate in the uniform sampling scenario.

### A. Communication Model

In our model, we consider a simple wireless star network where all devices have direct uplink and downlink channels to the global aggregator (see Figure 1). Each device has its own channel  $h_n^t$  to the aggregator, but must take turns in transmitting their parameters with transmit power  $P_n^t$  via time-division multiple access (TDMA). For simplicity, we only consider the uplink communication time as the downlink is a broadcast of common information by a highly capable base station which takes much less time. By assuming that we can transmit at capacity, the communication time for device  $n$  in a given round  $t$  is

$$\tau_n^{\text{comm}}(t) = \frac{\ell \mathbb{I}_n^t}{B \log_2 \left( 1 + |h_n^t|^2 \frac{P_n^t}{N_0} \right)}, \quad (8)$$

where  $\ell$  is the size of the model in bits,  $B$  is the bandwidth, and  $\log_2(\cdot)$  denotes the base 2 logarithm. We assume that the aggregator has current CSI in the form of channel gain  $|h_n^t|^2$  and noise power  $N_0$  at each round  $t$  in order to calculate this, but not the underlying distributions.

Next, as is typical in FL [2], [4], [8], we assume that the aggregator is using sampling *with* replacement with  $m$  draws. With this policy, a minimum of one client and a maximum of  $m$  clients are selected in each round. More specifically, the aggregator draws from a multinomial distribution where each device has probability  $\omega_n^t$  of being chosen in each of the  $m$  draws such that  $\sum_{n=1}^N \omega_n^t = 1$  for every round  $t$ . Our goal is to determine  $\omega_n^t$  such that communication time is minimized and convergence is guaranteed. However, since our convergence bound is in terms of arbitrary marginal probabilities  $q_n^t$ , we must relate  $\omega_n^t$  to  $q_n^t$  in order to measure the policy's impact on convergence. For  $m$  draws, the relationship is

$$q_n^t = 1 - (1 - \omega_n^t)^m, \quad (9)$$

where we now require

$$\sum_{n=1}^N \omega_n^t = \sum_{n=1}^N 1 - (1 - q_n^t)^{1/m} = 1 \quad (10)$$

to be satisfied in order to obtain a valid multinomial distribution. In practice, sampling *without* replacement can be used, but expressing  $q_n^t$  in terms of arbitrary  $\omega_n^t$  and  $m$  is not tractable and thus difficult to optimize directly.

### B. Problem Formulation

Similar to some prior works, such as [1], [12], [36], [37], and [38], we minimize the convergence bound (6) as a proxy for the actual convergence loss since the direct effect that parameters have on loss is generally unknown. Simultaneously, we also wish to minimize the time spent communicating the parameters, as in (8), to decrease wall-clock training time. The former is minimized when all devices participate in every round while the latter is minimized when fewer devices participate. Thus, to balance the competing goals, we formulate our objective function as the weighted sum between the two:

$$y_0(t, q_n^t, P_n^t) := \sum_{n=1}^N \left( \frac{1}{N q_n^t} + \lambda \cdot \tau_n^{\text{comm}}(t) \right), \quad (11)$$

where  $\lambda > 0$  is a user-defined, tunable parameter that controls the trade-off. The parameter also subsumes the coefficients in the last term in (6) such that  $L$  and  $\nu$ , for example, can be ignored in the following optimization problem.

Finally, we wish to limit the power expenditure of communication over time to a given time-average threshold,  $\bar{P}_n$ , and instantaneous power  $P_{\max}$ . Thus, we formulate our optimization problem as

$$\begin{aligned} \min_{\{q_n^t\}, \{P_n^t\}} \quad & \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[y_0(t)] \\ \text{s.t.} \quad & \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} P_n^t q_n^t \leq \bar{P}_n, \quad \forall n = 1, \dots, N \\ & 0 \leq P_n^t \leq P_{\max}, \quad n = 1, \dots, N \\ & \sum_{n=1}^N 1 - (1 - q_n^t)^{1/m} = 1 \\ & q_n^t \in (0, 1], \end{aligned} \quad (12)$$

where

$$\begin{aligned} \mathbb{E}[y_0(t)] &= \mathbb{E}[y_0(t, q_n^t, P_n^t)] \\ &= \sum_{n=1}^N \left( \frac{1}{N q_n^t} + \lambda \cdot \frac{\ell q_n^t}{B \log_2 \left( 1 + |h_n^t|^2 \frac{P_n^t}{N_0} \right)} \right). \end{aligned} \quad (13)$$

The first two constraints limit the time average and peak power, respectively, where  $P_n^t q_n^t = \mathbb{E}[P_n^t \mathbb{I}_n^t]$  is the expected power usage in round  $t$  by device  $n$ . The third constraint ensures that a proper multinomial distribution is found for sampling with replacement. In its current state, (12) is difficult to solve as it is not possible to know the behavior of the channel  $h_n^t$  for all  $t$  ahead of time.

### C. Lyapunov Formulation

The novelty of our convergence bound and formulation comes from the fact that both the effect of  $q_n^t$  on convergence and the communication time are in the form of a time average. This allows us to leverage the Lyapunov stochastic optimization framework [39] to reformulate (12) into a form that we can solve greedily at each round with optimality guarantees. By converting our transmission power constraints into a set of *virtual* queues, we can apply the Lyapunov theory to analyze our problem and derive an online solution. The practical implications of stabilizing virtual queues will be explored at the end of this section and its effect will be further illustrated in the experiments of Section VI.

To put our optimization problem into the Lyapunov drift-plus-penalty framework and using standard notation, we turn the first time-average constraint in (12) into a virtual queue  $Z_n^t$  for each client  $n$  such that

$$Z_n^{t+1} = \max[Z_n^t + y_n(t), 0], \quad (14)$$

where

$$y_n(t) = P_n^t q_n^t - \bar{P}_n. \quad (15)$$

Since we have no actual queues, the Lyapunov function is

$$\mathcal{L}(\Theta^t) := \frac{1}{2} \sum_{n=1}^N (Z_n^t)^2, \quad (16)$$

where  $\Theta^t$  represents the current queue states, which in this case, is just  $\{Z_n^t : \forall n\}$ . Next, we define the Lyapunov drift:

$$\Delta^{t+1} = \mathcal{L}^{t+1} - \mathcal{L}^t, \quad (17)$$

where we drop  $\Theta^t$  for simplicity. Finally, we have the Lyapunov drift-plus-penalty function that we aim to minimize:

$$\Delta^t + V\mathbb{E}[y_0(t)|\Theta^t], \quad (18)$$

where  $V > 0$  is another arbitrarily chosen weight that controls the fundamental trade-off between queue stability and optimality of the objective functions (the effect of which is also explored in Section VI).

Now, by utilizing Lemma 4.6 from [39] and assuming that the random event, i.e., channel gain  $|h_n^t|^2$ , is i.i.d.<sup>3</sup> with respect to  $t$ , we can upper bound (18):

$$\begin{aligned} \Delta^t + V\mathbb{E}[y_0(t)|\Theta^t] &\leq C + V\mathbb{E}[y_0(t)|\Theta^t] \\ &\quad + \sum_{n=1}^N Z_n^t \mathbb{E}[y_n(t)|\Theta^t] \end{aligned} \quad (19)$$

where  $C > 0$  is a constant. Next, according to the Min Drift-Plus-Penalty Algorithm, we opportunistically minimize the expectation in the right hand side of (19) at each time step  $t$ :

$$\begin{aligned} \min_{\{\omega_n^t\}, \{P_n^t\}} \quad &f(q_n^t, P_n^t) := V y_0(t) + \sum_{n=1}^N Z_n^t y_n(t) \\ \text{s.t.} \quad &0 \leq P_n^t \leq P_{\max}, \quad \forall n = 1, \dots, N \\ &\sum_{n=1}^N \omega_n^t = 1 \\ &q_n^t = 1 - (1 - \omega_n^t)^m \\ &\omega_n^t \in (0, 1]. \end{aligned} \quad (20)$$

We now have the problem in a form that is decoupled in time allowing for a greedy optimization approach.

#### D. Solving the Problem (20)

First, we utilize the fact that (20) is convex with respect to  $P_n^t$  and that the optimal value for  $P_n^t$  is independent of  $q_n^t$  to produce the following analytical expression in Theorem 2.

**Theorem 2.** *The  $P_n^t$  that optimizes (20) is independent of  $q_n^t$  and is given by either the endpoint, i.e.,  $P_n^{t,\text{opt}} = P_{\max}$ , or by*

$$P_n^{t,\text{opt}} = \frac{N_0}{|h_n^t|^2} \left( \frac{A}{4} W_0 \left( \sqrt{\frac{A}{4}} \right)^{-2} - 1 \right) \quad (21)$$

where  $A = \frac{V\lambda\ell|h_n^t|^2(\log(2))^2}{N_0 B Z_n^t}$  and  $W_0(\cdot)$  is the principal branch of the Lambert  $W$  function.

<sup>3</sup>We make this assumption out of simplicity. Additional analyses can show that the algorithm converges even when this doesn't hold, including for non-ergodic processes [39].

*Proof:* The proof can be found in Appendix B.  $\square$

Optimizing over  $q_n^t$  is trickier because of the additional constraint. We first note that the problem (20) can be simplified to the following form

$$\begin{aligned} \min_{\{\omega_n^t\}} \quad &\sum_{n=1}^N \left( A_n^t (q_n^t)^{-1} + B_n^t q_n^t + C_n^t \right) \\ \text{s.t.} \quad &\sum_{n=1}^N \omega_n^t = 1 \\ &q_n^t = 1 - (1 - \omega_n^t)^m \\ &\omega_n^t \in (0, 1], \end{aligned} \quad (22)$$

where  $A_n^t = \frac{V}{N}$ ,  $B_n^t = \frac{V\lambda\ell}{B \log_2 \left( 1 + |h_n^t|^2 \frac{P_n^{t,\text{opt}}}{N_0} \right)} + Z_n^t P_n^{t,\text{opt}}$ ,  $C_n^t = -Z_n^t \bar{P}_n$ , and  $P_n^{t,\text{opt}}$  is the optimal  $P_n^t$  determined by (21). While the objective function is convex, we have introduced a non-convex equality constraint which makes this problem significantly more difficult to solve. Additionally, the constraint creates a coupling of the individual device selection probabilities and cannot be solved independently.

A natural approach to solving this optimization problem is to use the Lagrangian multiplier method. Unfortunately, however, the gradient of the Lagrangian of (22) does not have a simple closed form solution, so we must resort to numerical solvers. The feasible region of this problem is the probability simplex which is a Riemannian manifold. To solve this problem, we use Manopt [40] which is a toolbox for optimization on manifolds.

#### E. Optimality Guarantees

Theorem 4.8 in [39], Theorem 2, and assuming we obtain the optimal value for  $q_n^t$  from the solver guarantee that this algorithm satisfies

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[y_0(t)] \leq y_o^{\text{opt}} + \frac{C+D}{V}, \quad (23)$$

where  $y_o^{\text{opt}}$  is the minimum of  $y_o$ ,  $D \geq 0$  is a constant whose existence is guaranteed in Theorem 4.8, and  $C \geq 0$  is a constant that bounds the expected difference between the optimal value of (20) and the value given by an approximately optimal decision. Since the subproblem in (22) is non-convex, we may choose a local optimum, so  $C \neq 0$  but is finite. The theorem also guarantees that the time-average transmit power constraint is satisfied as  $t \rightarrow \infty$ . The user-defined parameter  $V$  that traditionally controls the trade-off between the average queue backlog and the gap from optimality now controls the speed of convergence in addition to the optimality gap in (23).

In (21), we can see that when there is a large virtual queue  $Z_n^t$ , the transmission power is decreased in order to satisfy the constraint. In this way, the virtual queue represents how far from the time average constraint we are. As  $V$  is increased, the effect that the current virtual queue has on selection becomes less important and it takes longer to satisfy the average power constraint. This is also explored experimentally in Section VI-F. A large  $\lambda$  favors the minimization of communication time rather than the convergence bound which

naturally leads to more unbalanced  $q_n^t$ , whereas small  $\lambda$  will result in more uniform  $q_n^t$ .

## VI. EXPERIMENTS

In our experiments, we trained a convolutional neural network (CNN) on the CIFAR-10 benchmark [41] with data spread across  $N = 100$  clients. We examine the convergence behavior over wall-clock time for numerous scenarios with varying levels of data and channel heterogeneity, number of clients draws  $m$ , and values for hyperparameter  $\lambda$ . We then compare our device scheduling algorithm with the uniform selection baseline to demonstrate the speedup in convergence ours provides. Since the cost of an iteration/aggregation round is variable, analyzing convergence in terms of wall-clock time instead of iterations gives more insights into how the algorithms will perform in practice. The standard assumption that it is better to converge in fewer iterations does not necessarily hold if the cost of each iteration is large. The wall-clock time consists of both communication time and computation time where communication time is calculated as in (13) and computation time is chosen to be a constant. It is assumed that computation is performed in parallel and each device completes all  $K$  local iterations at the same deterministic time as opposed to communication which occurs sequentially via TDMA.

### A. Setup

The CIFAR-10 dataset consists of 60,000 color images of 10 classes where 10,000 images are reserved for testing. We assign each device 500 samples from the dataset. In order to simulate a non-i.i.d. data distribution among devices, we use the approach presented in [42] where the distribution of the 10 classes among each client is determined by a categorical distribution with a Dirichlet prior. This allows for the level of heterogeneity to be quantified by a single parameter  $\alpha$ . More specifically, for each experiment and client, we draw the probabilities for the categorical distribution via  $\mathbf{c} \sim \text{Dir}(\alpha\mathbf{e})$  where  $\mathbf{e}$  is a uniform probability vector. We can then vary  $\alpha$  to produce different levels of heterogeneity. For example, at the extremes, we have  $\alpha \rightarrow 0$  which results in very non-i.i.d. data where  $\mathbf{c}$  becomes a one-hot vector such that each client is assigned samples from only one class. Meanwhile, on the other hand, when we have  $\alpha \rightarrow \infty$ , we encounter the i.i.d. case where each client gets an even distribution of all classes. Then, after obtaining  $\mathbf{c}$  for each client, we draw from the resulting categorical distribution 500 times to determine the class of each sample. When the class is determined, one sample from that class is uniformly drawn from the dataset. We note that this does allow for the possibility of certain samples to be assigned more than once or not at all, but this approach provides the most randomness and flexibility.

For training, we use the same CNN model as in [1] which contains  $d = 555,178$  parameters. Assuming 32-bit floating point numbers, we accordingly set  $\ell = 32d$ . We also set the minibatch size to 32,  $\gamma = 0.01$ ,  $I = 10$ , and  $B = 22 \times 10^6$  to simulate WiFi bandwidth. The power constraints are set to  $\bar{P}_n = 1$  and  $P_{\max} = 35$  dB and noise power is normalized

to  $N_0 = 1$ . We set  $V = 100$  in our selection algorithm and justify this choice later in Section VI-F. For the channel model, we assume each device experiences Rayleigh fading such that  $|h_n^t|$  is distributed as a Rayleigh random variable. In the first set of experiments, we examine the heterogeneous channel case such that each device experiences a different level of fading. In particular, we assign the Rayleigh parameters in a linear increasing fashion from  $\sigma = 0.1$  to  $\sigma = 10$  for the 100 devices. For the next set of experiments, we examine the homogeneous channel case such that each device has  $\sigma = 1$ . Note that our algorithm only requires knowledge of the instantaneous  $|h_n^t|$  and not the underlying distribution, requiring only that it is i.i.d. across time slots.

To fairly compare to the uniform selection baseline, we set  $q_n^t$  to be the appropriate values for the aggregation weights in Algorithm 1 and we set  $P_n^t$  to maximally satisfy the transmit power constraint in (12). So, for uniform sampling with replacement, we set  $q_n^t = \frac{1}{1-(1-1/N)^m}$  and  $P_n^t = \min\{P_{\max}, \bar{P}_n \cdot q_n^t\}$ . To avoid big outliers that likely would not be chosen by either selection policy in practice, we lower bound the possible values for  $|h_n^t|^2 \geq .001$ . This is very generous to the uniform case as it is agnostic to channel conditions and would greatly suffer from choosing a device in deep fade. Finally, since the total elapsed time varies across runs at each communication round, we perform linear interpolation for each run to generate points at uniform time intervals. This allows for us to more easily average the performance of the 3 runs for each scenario. We then perform a rolling average to smooth the curves to better see the convergence trends.

### B. Exploring the Impact of Computation Time

First, we give an illustrative example of how the optimal number of sampling with replacement draws  $m$  in each round is impacted by increasing computation time. In Fig. 2, we plot the time to reach a target accuracy of 0.775 versus computation time per round for a basic scenario with  $m = 1, 5$ , and 10 draws using our algorithm. In this scenario, there are three clear regions in which a different number of draws reaches the target accuracy quickest. It can be seen that as computation time increases, the optimal number of device draws also increases. In the first region, *scheduling only one device per round is optimal* despite the fact that it takes more iterations to converge than the other cases (depicted later in Fig. 5). However, since the communication phase in each iteration takes less time, due to the aggregator only needing to wait for 1 device in a sequential communication process, the lower quality gradient estimates are overcome by the quick communication. On the other hand, when computation time is comparatively larger than communication, drawing more devices per round is optimal. This is because as computation time increases, it begins to dominate the total time, so communicating additional devices' updates in each round has a smaller comparative impact on total time. Thus, *it is more advantageous to schedule more devices in each round* in order to leverage the parallelism of computation and obtain more informative updates.

We expect that as computation time continues to increase, the optimal number of draws will continue to increase up to a

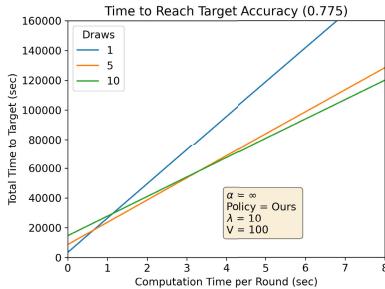


Fig. 2. The optimal number of devices chosen each round depends on computation time.

certain point. It is hard to predict this point as it is related to optimal minibatch sizes. This is a similarly unsolved problem where the generalizability of smaller minibatches generally outweighs the benefits of more accurate gradient estimates obtained from larger batch sizes [43]. Additionally, at a certain point Amdahl's Law must be taken into consideration for very large computation times. Even though our algorithm still technically provides an advantage over uniform, if communication only accounts for a small percentage of the total time, then optimizing communication alone cannot speed up performance significantly.

### C. Best Hyperparameters

Next, in Table I, we identify the best hyperparameters for both our algorithm and the uniform baseline for 3 different levels of data heterogeneity over computation time ranges. We see similar trends for all three levels of data heterogeneity and both selection policies where, again, *higher computation time results in a greater number of draws*. It is also interesting to note the optimal  $\lambda$  values. For large  $\lambda$ , minimizing communication time is favored over minimizing the convergence bound. Here, we see that as computation time increase comparative to communication time, lower  $\lambda$  values become optimal. This is because as the percentage of total time due to computation time increases, communication time becomes less important. Thus, it is better to choose devices in a more uniform manner in order to avoid client drift toward devices with better channels. We will use these optimal hyperparameter ranges in the following results, but we note that we only tested a discrete subset of hyperparameters so they are not necessarily truly optimal.

### D. Heterogeneous Channels

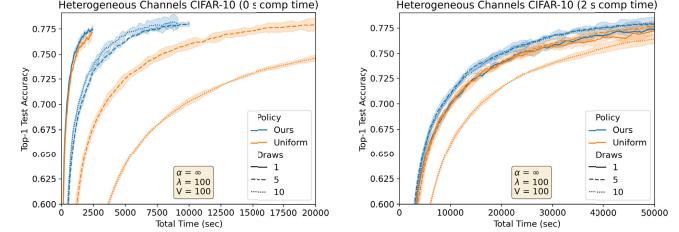
In this section, we investigate the convergence behavior of our algorithm over wall-clock time and compare it to the uniform baseline. We begin with the heterogeneous channels case where some devices have very good channels on average while others do not. In Fig. 3, we show how top-1 accuracy and training loss are affected by different sampling with replacement draws  $m$ . In Figs. 3a and 3c, we show the results for instantaneous computation such that the total time only consists of communication time while in Figs. 3b and 3d, the computation time is set to 2 seconds. The plots clearly show that *our algorithm outperforms uniform*, especially when the

TABLE I  
BEST HYPERPARAMETERS FROM EXPERIMENTS FOR DIFFERENT COMPUTATION TIME RANGES AND LEVELS OF DATA HETEROGENEITY

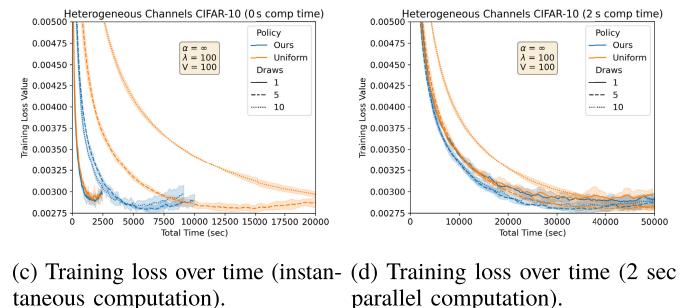
Heterogeneous Channel Gain			
Case	Comp Times (s)	Ours	
		Draws (m)	$\lambda$
$\alpha = \infty$ , target = 0.775	0 – 0.83	1	100
	0.83 – 2.3	10	100
	2.3 – 7.5	5	10
	7.5 – 16.3	10	10
	16.3 – $\infty$	10	1
$\alpha = 1$ , target = 0.73	0 – 0.61	1	100
	0.61 – 3.1	10	100
	3.1 – 15.6	5	1
	15.6 – $\infty$	10	1
$\alpha = 0$ , target = 0.65	0 – 0.26	1	100
	0.26 – 1.8	10	100
	1.8 – $\infty$	10	10

Uniform			
Case	Comp Times (s)	Ours	
		Draws (m)	
$\alpha = \infty$ , target = 0.775	0 – 2	1	
	2 – $\infty$	5	
$\alpha = 1$ , target = 0.73	0 – 0.84	1	
	0.84 – 46	5	
	46 – $\infty$	10	
$\alpha = 0$ , target = 0.65	0 – 0.74	1	
	0.74 – 6.1	5	
	6.1 – $\infty$	10	



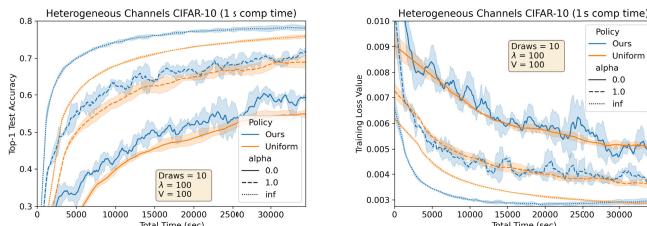
(a) Testing accuracy over time (instantaneous computation). (b) Testing accuracy over time (2 sec parallel computation).



(c) Training loss over time (instantaneous computation). (d) Training loss over time (2 sec parallel computation).

Fig. 3. Comparison of total communication time for uniform selection vs proposed algorithm on CIFAR-10 dataset.

number of draws is large. In fact, for the  $m = 10$  regime, we see a very large speedup of  $8.5 \times$  to reach a target accuracy of 0.775 where our policy takes 5,339 seconds on average and uniform takes 45,748 seconds. The speedup is lesser for the  $m = 1$  case which is about  $1.3 \times$  for the same target. In these first two plots, where only communication time is taken in consideration, scheduling *just one device per round* is quickest. As explained in the previous section, this is a scenario where it is *better to quickly aggregate updates from fewer devices* since the shorter communication time per round outweighs the higher variance of each update.



(a) Testing accuracy over time for varying levels of data heterogeneity.

(b) Training loss over time for varying levels of data heterogeneity.

Fig. 4. Convergence for heterogeneous data.

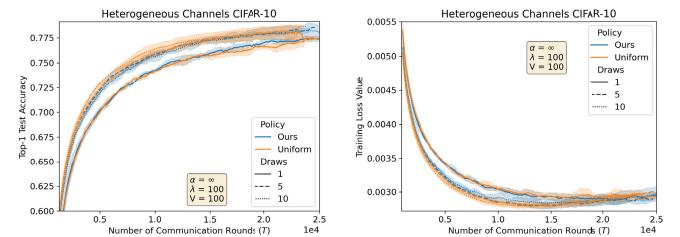
In the second set of plots, we show the same scenario but with parallel computation time of two seconds per round. This causes the optimal number of draws to change to  $m = 5$  with a speedup of  $1.31 \times$  over uniform. The speedup in the  $m = 10$  case remains large, however. Since optimizing communication time only affects the fraction of the total time due to communication, the speedup is less significant when computation times are large.

Next, in Fig. 4, we examine the effect of heterogeneity in the data through parameter  $\alpha$ . As  $\alpha$  increases, the data distribution across devices becomes more homogeneous such that each device has access to samples of more classes. For this experiment, we set the computation time to one second,  $m = 10$ , and  $\lambda = 100$ . We chose a one second computation time since the optimal hyperparameters for all three  $\alpha$  levels is the same as seen in Table I. The results again show that our algorithm outperforms the uniform baseline, but that *the advantage is lesser for higher degrees of heterogeneity* in the data. This intuitively makes sense since a more uniform sampling approach will help avoid client shift in the model as each node will contribute equally on average. Still, though, our algorithm provides an advantage by avoiding nodes with poor instantaneous channels. The general trend of slower convergence for higher non-i.i.d.-ness also holds. The convergence speed for  $\alpha = 0$  is especially bad since each device only has access to samples from one class.

Finally, in Fig. 5, we plot the same results as in Fig. 3 but in terms of communication rounds rather than wall-clock time. This is shown to demonstrate how cases like  $m = 1$  take more iterations to converge and are slowest in the traditional accuracy per iteration/epoch convergence metric. However, since these iterations can occur so quickly with less time spent communicating in each round, these cases are much better in practice. As computation time increases and begins to dominate the total time, though, the trends over total time will begin to match the ones shown here. Although, we reiterate again that simply setting  $m = N = 100$  does not guarantee the best performance when computation is high as the generalization of minibatch sizes comes into play [43].

### E. Homogeneous Channels

In this section, we briefly explore the convergence behavior for homogeneous channels where each device has the same Rayleigh fading parameter of  $\sigma = 1$ . We produce a table



(a) Testing accuracy per communication round ( $T$ ). (b) Training loss per communication round ( $T$ ).

Fig. 5. Convergence over communication rounds/iterations.

TABLE II

BEST HYPERPARAMETERS FROM EXPERIMENTS FOR DIFFERENT COMPUTATION TIME RANGES AND LEVELS OF DATA HETEROGENEITY

Case	Homogeneous Channel Gain		
	Ours		
$\alpha = \infty$ , target = 0.775	Comp Times (s)	Draws ( $m$ )	$\lambda$
	0 – 2.48	1	100
	2.48 – 13.1	5	100
$\alpha = 1$ , target = 0.74	Comp Times (s)	Draws ( $m$ )	$\lambda$
	13.1 – $\infty$	10	1
	0 – 0.69	1	100
$\alpha = \infty$ , target = 0.775	Comp Times (s)	Draws ( $m$ )	
	0.69 – 10.1	10	100
	10.1 – $\infty$	10	1
Uniform			
$\alpha = 1$ , target = 0.73	Comp Times (s)	Draws ( $m$ )	
	0 – 2.5	1	
	2.5 – $\infty$	5	
$\alpha = 1$ , target = 0.73	Comp Times (s)	Draws ( $m$ )	
	0 – 1.3	1	
	1.8 – 18	5	
	18 – $\infty$	10	

of optimal hyperparameters in Table II and plot convergence results in Figure 6. We plot in terms of top-1 testing accuracy and training loss and for two different levels of data heterogeneity. The table contains the same trends as in Table I. In Fig. 6, it can be seen that our algorithm still *outperforms the uniform baseline*, especially for i.i.d. data, but the speedup is not as significant. For example, the time to target accuracy of 0.775 for  $\alpha = \infty$  and  $m = 1$ , achieves a  $1.24 \times$  speedup. For fair comparison, we picked  $m$  such that both selection policies are optimal for the same value according to Table II. For different levels of  $m$ , the speedup is more significant since the probability of choosing a slow device increases with more draws when using the uniform sampling policy. For example, when  $m = 10$ , there is a speed up of  $6.7 \times$  (not pictured here) over uniform as compared to  $8.5 \times$  in the heterogeneous channels case.

Since all devices are equally likely to have good channels in the homogeneous channel case, device selection in our algorithm is more uniform. Therefore, we expect the impact of data heterogeneity to be lessened as compared to the heterogeneous channel case. However, this does not appear to be the case. The reasoning for this is that even though the devices with higher  $\sigma$  are chosen more often in the heterogeneous case (see Section VI-F), all 10 classes are still present among those top choices. So, even with skewed sampling, the algorithm is learning from all classes. If the distribution of classes is dependent on channel quality, however, then the decrease in convergence would be expected to be more pronounced.

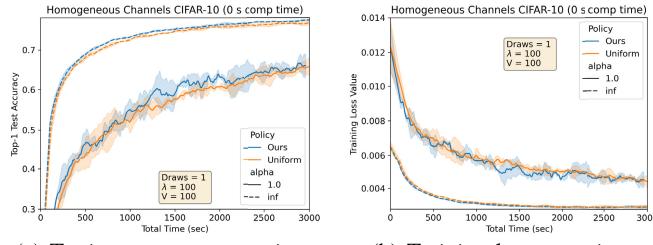
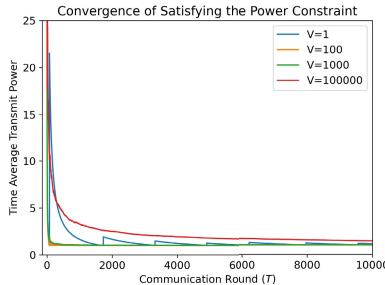


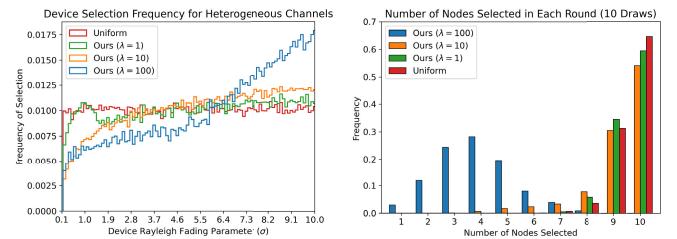
Fig. 6. Convergence for homogeneous channels.

Fig. 7. The convergence of the constraint for different values of  $V$ . The larger the  $V$ , the more rounds it takes until the constraint is satisfied. Here, the constraint is  $\bar{P}_n = 1$  for all  $n$ .

#### F. Additional Convergence and Statistical Insights

In this section, we further explore the behavior of our client sampling algorithm. First, we show how the choice of parameter  $V$  affects convergence of satisfying the time average power constraint. From (23), we know that  $V$  controls the gap in optimality where larger values result in a smaller gap. Normally,  $V$  also conversely controls the queue backlog lengths, but since we only deal with virtual queues, the behavior is slightly different. In Figure 7, we plot the expected time average transmit power  $\frac{1}{T} \sum_{t=0}^{T-1} P_n^t q_n^t$  over the course of multiple communication rounds. We see here that increased  $V$  causes the constraint of  $\bar{P} = 1$  to be satisfied at a slower rate. For example,  $V = 10^5$  does not reach  $\bar{P} = 1$  within a realistic number of rounds. This is important because even though it is guaranteed to converge asymptotically, the FL training will have to terminate in finite time in practice. Similarly for  $V = 1$ , the constraint may be violated at a given finite termination time  $T$  due to its oscillatory behavior despite offering the worst optimality gap. *Our algorithm sacrifices initial constraint violation in finite time in order to make early gains in performance.* The improvements over uniform are not solely attributed to this, however. For the previous experiments, we chose  $V = 100$  since it satisfies the constraint without extreme advantage over uniform.

Last, we examine the device selection statistics of our algorithm versus uniform in Fig. 8. Since we mainly plotted the  $\lambda = 100$  case in the previous sections, we show here how other values affect the device selection process. In Fig. 8a, we show a histogram of the selection frequency of all  $N = 100$  devices in the heterogeneous case for various  $\lambda$  values. The x axis displays each device's Rayleigh fading parameter in increasing order. When  $\lambda$  is larger, minimizing communication time is prioritized over the convergence bound. Thus, for  $\lambda = 100$ , there is a large skew in selecting the devices



(a) The frequency of selection for each device for different  $\lambda$  and policy. (b) The probability of the number of devices chosen per round varying  $\lambda$  and policy.

Fig. 8. Device selection statistics.

with better channels on average. As  $\lambda$  decreases, the selection becomes more uniform where  $\lambda = 1$  and uniform have similar distributions. Next, in Figure 8b, we investigate how  $\lambda$  affects the number of devices chosen per round when using a sampling *with* replacement policy. Since a device can be drawn more than once, the actual number of devices scheduled in a given round may not equal  $m$ , even in the uniform case. Here, you can see that lower  $\lambda$  results in probabilities closer to uniform, whereas when  $\lambda = 100$ , it is much more common to schedule only 3 to 5 devices despite drawing  $m = 10$  times. Through the selection probabilities, *our algorithm is able to adaptively change how many devices are chosen depending on current channel conditions.*

## VII. CONCLUSION

In this paper, we explored the impact of device scheduling over wireless links in federated learning and proposed a novel scheduling algorithm that leverages stochastic optimization theory. Our experiments saw up to 8.5x speedup over the uniform baselines in some cases. The first contribution of our work was developing a convergence bound for non-convex FL loss functions with arbitrary device participation probabilities without a bounded gradient assumption. We showed that it recovers the state-of-the-art convergence rate with linear speedup. We then used the bound to formulate a time average optimization problem with a time average power constraint that can be solved using the Lyapunov drift-plus-penalty framework. The resulting online algorithm can be solved at each communication round without requiring knowledge of channel statistics. Our experiments show impressive performance over the uniform baseline in terms of wall-clock time. We explored many scenarios such varying levels of data and channel heterogeneity as well as varying computation time and hyperparameters. Depending on the scenario, we showed that it is sometimes better to perform quicker, less informative updates than slower, more informative updates as seen in the fantastic performance of the  $m = 1$  regime for low computation time. There are very interesting trade-offs to consider when the cost of each training iteration is variable which is the subject of future research.

## APPENDIX

We use the following inequalities throughout the convergence proof.

**Preliminary (in)equalities.** We will use Jensen's inequality:

$$\left\| \frac{1}{M} \sum_{m=1}^M \mathbf{y}_m \right\|^2 \leq \frac{1}{M} \sum_{m=1}^M \|\mathbf{y}_m\|^2 \quad (24)$$

$$\left\| \sum_{m=1}^M \mathbf{y}_m \right\|^2 \leq M \sum_{m=1}^M \|\mathbf{y}_m\|^2, \quad (25)$$

and the Peter-Paul inequality:

$$\langle \mathbf{y}_1, \mathbf{y}_2 \rangle \leq \frac{\rho \|\mathbf{y}_1\|^2}{2} + \frac{\|\mathbf{y}_2\|^2}{2\rho} \quad (26)$$

for  $\rho > 0$ . We also use

$$\langle \mathbf{x}, \mathbf{y} \rangle = \frac{1}{2} \left( \|\mathbf{x}\|^2 + \|\mathbf{y}\|^2 - \|\mathbf{x} - \mathbf{y}\|^2 \right) \quad (27)$$

#### A. Proof of Theorem 1

First, we note that via Algorithm 1, we have

$$\begin{aligned} \mathbf{x}_{t+1} - \mathbf{x}_t &= \frac{1}{N} \sum_{n=1}^N \frac{\mathbb{I}_n^t}{q_n^t} (\mathbf{y}_{t,K}^n - \mathbf{y}_{t,0}^n) \\ &= -\frac{\gamma}{N} \sum_{n=1}^N \frac{\mathbb{I}_n^t}{q_n^t} \sum_{i=0}^{K-1} \mathbf{g}_n(\mathbf{y}_{t,i}^n). \end{aligned}$$

Let  $\mathbb{E}_t[\cdot] := \mathbb{E}[\cdot | x_t, \{q_n^t\}]$ , which denotes the expectation over the randomness of SGD and client sampling. From  $L$ -smoothness, we have

$$\begin{aligned} &\mathbb{E}_t[f(\mathbf{x}_{t+1})] \\ &\leq f(\mathbf{x}_t) + \langle \nabla f(\mathbf{x}_t), \mathbb{E}_t[\mathbf{x}_{t+1} - \mathbf{x}_t] \rangle \\ &\quad + \frac{L}{2} \mathbb{E}_t \left[ \|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2 \right] \\ &= f(\mathbf{x}_t) - \left\langle \nabla f(\mathbf{x}_t), \mathbb{E}_t \left[ \frac{1}{N} \sum_{n=1}^N \frac{\mathbb{I}_n^t}{q_n^t} \sum_{i=0}^{K-1} \gamma \mathbf{g}_n(\mathbf{y}_{t,i}^n) \right] \right\rangle \\ &\quad + \frac{L\gamma^2}{2N^2} \mathbb{E}_t \left[ \left\| \sum_{n=1}^N \frac{\mathbb{I}_n^t}{q_n^t} \sum_{i=0}^{K-1} \mathbf{g}_n(\mathbf{y}_{t,i}^n) \right\|^2 \right] \\ &\stackrel{(a)}{=} f(\mathbf{x}_t) - \left\langle \nabla f(\mathbf{x}_t), \mathbb{E}_t \left[ \frac{\gamma}{N} \sum_{n=1}^N \sum_{i=0}^{K-1} \nabla f_n(\mathbf{y}_{t,i}^n) \right] \right\rangle \\ &\quad + \frac{L\gamma^2}{2N^2} \mathbb{E}_t \left[ \left\| \sum_{n=1}^N \frac{\mathbb{I}_n^t}{q_n^t} \sum_{i=0}^{K-1} \mathbf{g}_n(\mathbf{y}_{t,i}^n) \right\|^2 \right], \quad (28) \end{aligned}$$

where (a) uses the independence between  $\mathbb{I}_n^t$  and  $\mathbf{g}_n$ , the fact that  $\mathbb{E}[\mathbb{I}_n^t | \mathbf{x}_t] = \mathbb{E}[\mathbb{I}_n^t] = q_n^t$ , and the total expectation  $\mathbb{E}[\mathbf{g}_n(\mathbf{y}_{t,i}^n) | \mathbf{x}_t] = \mathbb{E}[\mathbb{E}[\mathbf{g}_n(\mathbf{y}_{t,i}^n) | \mathbf{y}_{t,i}^n, \mathbf{x}_t] | \mathbf{x}_t] = \mathbb{E}[\nabla f_n(\mathbf{y}_{t,i}^n) | \mathbf{x}_t]$ .

We now bound the terms in (28) separately. We begin with the second term and note that

$$\begin{aligned} &-\left\langle \nabla f(\mathbf{x}_t), \mathbb{E}_t \left[ \frac{\gamma}{N} \sum_{n=1}^N \sum_{i=0}^{K-1} \nabla f_n(\mathbf{y}_{t,i}^n) \right] \right\rangle \\ &= -\left\langle \nabla f(\mathbf{x}_t), \gamma \mathbb{E}_t \left[ \frac{1}{N} \sum_{n=1}^N \sum_{i=0}^{K-1} \nabla f_n(\mathbf{y}_{t,i}^n) \right] \right\rangle \\ &= \left\langle \nabla f(\mathbf{x}_t), -\gamma \mathbb{E}_t \left[ \frac{1}{N} \sum_{n=1}^N \sum_{i=0}^{K-1} \nabla f_n(\mathbf{y}_{t,i}^n) - K \nabla f(\mathbf{x}_t) \right] \right\rangle \end{aligned}$$

$$\begin{aligned} &-\gamma K \|\nabla f(\mathbf{x}_t)\|^2 \\ &= \left\langle \sqrt{\gamma K} \nabla f(\mathbf{x}_t), -\frac{\sqrt{\gamma}}{N\sqrt{K}} \mathbb{E}_t \left[ \sum_{n=1}^N \sum_{i=0}^{K-1} \left( \nabla f_n(\mathbf{y}_{t,i}^n) - \nabla f_n(\mathbf{x}_t) \right) \right] \right\rangle \\ &\quad - \gamma K \|\nabla f(\mathbf{x}_t)\|^2 \\ &\stackrel{(a)}{=} \frac{\gamma}{2N^2 K} \mathbb{E}_t \left[ \left\| \sum_{n=1}^N \sum_{i=0}^{K-1} (\nabla f_n(\mathbf{y}_{t,i}^n) - \nabla f_n(\mathbf{x}_t)) \right\|^2 \right] \\ &\quad - \frac{\gamma}{2N^2 K} \mathbb{E}_t \left[ \left\| \sum_{n=1}^N \sum_{i=0}^{K-1} \nabla f_n(\mathbf{y}_{t,i}^n) \right\|^2 \right] \\ &\quad - \frac{\gamma K}{2} \|\nabla f(\mathbf{x}_t)\|^2 \\ &\stackrel{(b)}{\leq} \frac{\gamma K L^2}{2} \left( 5K\gamma^2 (\nu^2 + 6K\epsilon^2) + 30K^2\gamma^2 \|\nabla f(\mathbf{x}_t)\|^2 \right) \\ &\quad - \frac{\gamma}{2N^2 K} \mathbb{E}_t \left[ \left\| \sum_{n=1}^N \sum_{i=0}^{K-1} \nabla f_n(\mathbf{y}_{t,i}^n) \right\|^2 \right] \\ &\quad - \frac{\gamma K}{2} \|\nabla f(\mathbf{x}_t)\|^2 \\ &= \frac{5\gamma^3 K^2 L^2}{2} (\nu^2 + 6K\epsilon^2) \\ &\quad + \left( 15\gamma^3 K^3 L^2 - \frac{\gamma K}{2} \right) \|\nabla f(\mathbf{x}_t)\|^2 \\ &\quad - \frac{\gamma}{2N^2 K} \mathbb{E}_t \left[ \left\| \sum_{n=1}^N \sum_{i=0}^{K-1} \nabla f_n(\mathbf{y}_{t,i}^n) \right\|^2 \right], \end{aligned}$$

where (a) is due to (27) and (b) is from

$$\begin{aligned} &\mathbb{E}_t \left[ \left\| \sum_{n=1}^N \sum_{i=0}^{K-1} (\nabla f_n(\mathbf{y}_{t,i}^n) - \nabla f_n(\mathbf{x}_t)) \right\|^2 \right] \\ &\leq NK \sum_{n=1}^N \sum_{i=0}^{K-1} \mathbb{E}_t \left[ \|(\nabla f_n(\mathbf{y}_{t,i}^n) - \nabla f_n(\mathbf{x}_t))\|^2 \right] \\ &\leq NKL^2 \sum_{n=1}^N \sum_{i=0}^{K-1} \mathbb{E}_t \left[ \|\mathbf{x}_t - \mathbf{y}_{t,i}^n\|^2 \right] \\ &\leq N^2 K^2 L^2 \left( 5K\gamma^2 (\nu^2 + 6K\epsilon^2) \right. \\ &\quad \left. + 30K^2\gamma^2 \|\nabla f(\mathbf{x}_t)\|^2 \right) \quad (29) \end{aligned}$$

which utilizes Lemma A.1 for the final inequality.

Next, we bound the third term in (28) as

$$\begin{aligned} &\mathbb{E}_t \left[ \left\| \sum_{n=1}^N \frac{\mathbb{I}_n^t}{q_n^t} \sum_{i=0}^{K-1} \mathbf{g}_n(\mathbf{y}_{t,i}^n) \right\|^2 \right] \\ &= \mathbb{E}_t \left[ \left\| \sum_{n=1}^N \sum_{i=0}^{K-1} \frac{\mathbb{I}_n^t}{q_n^t} (\mathbf{g}_n(\mathbf{y}_{t,i}^n) - \nabla f_n(\mathbf{y}_{t,i}^n) + \nabla f_n(\mathbf{y}_{t,i}^n)) \right\|^2 \right] \\ &\leq 2\mathbb{E}_t \left[ \left\| \sum_{n=1}^N \sum_{i=0}^{K-1} \frac{\mathbb{I}_n^t}{q_n^t} \nabla f_n(\mathbf{y}_{t,i}^n) \right\|^2 \right] \\ &\quad + 2\mathbb{E}_t \left[ \left\| \sum_{n=1}^N \sum_{i=0}^{K-1} \frac{\mathbb{I}_n^t}{q_n^t} (\mathbf{g}_n(\mathbf{y}_{t,i}^n) - \nabla f_n(\mathbf{y}_{t,i}^n)) \right\|^2 \right] \end{aligned}$$

$$\begin{aligned}
&\stackrel{(a)}{=} 2\mathbb{E}_t \left[ \left\| \sum_{n=1}^N \sum_{i=0}^{K-1} \frac{\mathbf{I}_n^t}{q_n^t} \nabla f_n(\mathbf{y}_{t,i}^n) \right\|^2 \right] \\
&\quad + 2 \sum_{n=1}^N \sum_{i=0}^{K-1} \mathbb{E}_t \left[ \left\| \frac{\mathbf{I}_n^t}{q_n^t} (\mathbf{g}_n(\mathbf{y}_{t,i}^n) - \nabla f_n(\mathbf{y}_{t,i}^n)) \right\|^2 \right] \\
&= 2\mathbb{E}_t \left[ \left\| \sum_{n=1}^N \sum_{i=0}^{K-1} \frac{\mathbf{I}_n^t}{q_n^t} \nabla f_n(\mathbf{y}_{t,i}^n) \right\|^2 \right] \\
&\quad + 2 \sum_{n=1}^N \sum_{i=0}^{K-1} \frac{1}{q_n^t} \mathbb{E}_t \left[ \left\| \mathbf{g}_n(\mathbf{y}_{t,i}^n) - \nabla f_n(\mathbf{y}_{t,i}^n) \right\|^2 \right] \\
&\leq 2\mathbb{E}_t \left[ \left\| \sum_{n=1}^N \sum_{i=0}^{K-1} \frac{\mathbf{I}_n^t}{q_n^t} \nabla f_n(\mathbf{y}_{t,i}^n) \right\|^2 \right] + 2NKQ_t\nu^2,
\end{aligned}$$

where (a) is due to the fact that  $\{\mathbf{g}_n(\mathbf{y}_{t,i}^n) - \nabla f_n(\mathbf{y}_{t,i}^n)\}$  is independent over  $n$  and is a martingale difference sequence over  $i$  such that  $\mathbb{E}_t \left[ \left\| \sum \mathbf{x} \right\|^2 \right] = \sum \mathbb{E}_t \left[ \left\| \mathbf{x} \right\|^2 \right]$ , [4, Lemma 4]. The final inequality is from Assumption 3 and we define  $Q_t = \frac{1}{N} \sum_{n=1}^N \frac{1}{q_n^t}$ .

Before returning to (28), we note a few more inequalities. Let  $\mathbf{z}_n = \sum_{i=0}^{K-1} \nabla f_n(\mathbf{y}_{t,i}^n)$  and suppose that there exists a  $q_{\min} \leq q_n^t$  for all  $n, t$ . Then, we have

$$\begin{aligned}
\mathbb{E}_t \left[ \left\| \sum_{n=1}^N \frac{\mathbf{I}_n^t}{q_n^t} \mathbf{z}_n \right\|^2 \right] &= \sum_{i,j}^N \frac{1}{q_i^t q_j^t} \mathbb{E}_t [\mathbf{I}_i^t \mathbf{I}_j^t] \mathbb{E}_t [\langle \mathbf{z}_i, \mathbf{z}_j \rangle] \\
&\leq \sum_{i,j}^N \frac{1}{\sqrt{q_i^t q_j^t}} \mathbb{E}_t [\langle \mathbf{z}_i, \mathbf{z}_j \rangle] \\
&\leq \sum_{i,j}^N \frac{1}{q_{\min}} \mathbb{E}_t [\langle \mathbf{z}_i, \mathbf{z}_j \rangle] \\
&= \frac{1}{q_{\min}} \mathbb{E}_t \left[ \left\| \sum_{n=1}^N \mathbf{z}_n \right\|^2 \right] \tag{30}
\end{aligned}$$

where we use the fact that sampling and gradient estimation are independent, and we use Cauchy-Schwarz inequality such that

$$\begin{aligned}
\mathbb{E}_t [\mathbf{I}_i^t \mathbf{I}_j^t] &\leq \sqrt{\mathbb{E}_t [(\mathbf{I}_i^t)^2] \mathbb{E}_t [(\mathbf{I}_j^t)^2]} \\
&= \sqrt{\mathbb{E}_t [\mathbf{I}_i^t] \mathbb{E}_t [\mathbf{I}_j^t]} \\
&= \sqrt{q_i^t q_j^t}.
\end{aligned}$$

Returning to (28) and applying the previous bounds, we have

$$\begin{aligned}
\mathbb{E}_t [f(\mathbf{x}_{t+1})] &\leq f(\mathbf{x}_t) - \left\langle \nabla f(\mathbf{x}_t), \mathbb{E}_t \left[ \frac{\gamma}{N} \sum_{n=1}^N \sum_{i=0}^{K-1} \nabla f_n(\mathbf{y}_{t,i}^n) \right] \right\rangle \\
&\quad + \frac{L\gamma^2}{2N^2} \mathbb{E}_t \left[ \left\| \sum_{n=1}^N \frac{\mathbf{I}_n^t}{q_n^t} \sum_{i=0}^{K-1} \mathbf{g}_n(\mathbf{y}_{t,i}^n) \right\|^2 \right] \\
&\leq f(\mathbf{x}_t) + \frac{5\gamma^3 K^2 L^2}{2} (\nu^2 + 6K\epsilon^2)
\end{aligned}$$

$$\begin{aligned}
&\quad - \frac{\gamma}{2N^2 K} \mathbb{E}_t \left[ \left\| \sum_{n=1}^N \sum_{i=0}^{K-1} \nabla f_n(\mathbf{y}_{t,i}^n) \right\|^2 \right] \\
&\quad + \left( 15\gamma^3 K^3 L^2 - \frac{\gamma K}{2} \right) \|\nabla f(\mathbf{x}_t)\|^2 + \frac{L\gamma^2 K Q_t}{N} \nu^2 \\
&\quad + \frac{L\gamma^2}{N^2} \left( \mathbb{E}_t \left[ \left\| \sum_{n=1}^N \sum_{i=0}^{K-1} \frac{\mathbf{I}_n^t}{q_n^t} \nabla f_n(\mathbf{y}_{t,i}^n) \right\|^2 \right] \right) \\
&\leq f(\mathbf{x}_t) + \frac{5\gamma^3 K^2 L^2}{2} (\nu^2 + 6K\epsilon^2) + \frac{L\gamma^2 K Q_t}{N} \nu^2 \\
&\quad + \left( 15\gamma^3 K^3 L^2 - \frac{\gamma K}{2} \right) \|\nabla f(\mathbf{x}_t)\|^2, \tag{31}
\end{aligned}$$

where the last inequality uses (30) such that

$$\begin{aligned}
&\frac{L\gamma^2}{N^2} \mathbb{E}_t \left[ \left\| \sum_{n=1}^N \frac{\mathbf{I}_n^t}{q_n^t} \mathbf{z}_n \right\|^2 \right] - \frac{\gamma}{2N^2 K} \mathbb{E}_t \left[ \left\| \sum_{n=1}^N \mathbf{z}_n \right\|^2 \right] \\
&\leq \frac{\gamma}{N^2} \left( \frac{L\gamma}{q_{\min}} - \frac{1}{2K} \right) \mathbb{E}_t \left[ \left\| \sum_{n=1}^N \mathbf{z}_n \right\|^2 \right] \\
&= 0, \tag{32}
\end{aligned}$$

if  $\gamma < \frac{q_{\min}}{2LK}$ .

Continuing from (31), we have

$$\begin{aligned}
\mathbb{E}_t [f(\mathbf{x}_{t+1})] &\leq f(\mathbf{x}_t) + \frac{5\gamma^3 K^2 L^2}{2} (\nu^2 + 6K\epsilon^2) + \frac{L\gamma^2 K Q_t}{N} \nu^2 \\
&\quad + \left( 15\gamma^3 K^3 L^2 - \frac{\gamma K}{2} \right) \|\nabla f(\mathbf{x}_t)\|^2, \\
&\leq f(\mathbf{x}_t) + \frac{5\gamma^3 K^2 L^2}{2} (\nu^2 + 6K\epsilon^2) + \frac{L\gamma^2 K Q_t}{N} \nu^2 \\
&\quad - c \frac{\gamma K}{2} \|\nabla f(\mathbf{x}_t)\|^2, \tag{33}
\end{aligned}$$

where the final inequality holds since there exists a constant  $c > 0$  such that  $(1 - 30\gamma^2 K^2 L^2) > c > 0$  if we assume that  $\gamma < \frac{1}{6LK}$ .

Finally, taking total expectation and then rearranging and summing  $t$  from 0 to  $T-1$ , we have

$$\begin{aligned}
\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(\mathbf{x}_t)\|^2] &\leq \frac{2(\mathbb{E}[f(\mathbf{x}_0)] - \mathbb{E}[f(\mathbf{x}_T)])}{c\gamma KT} \\
&\quad + \Phi_1 + \frac{\Phi_2}{TN} \sum_{t=0}^{T-1} Q_t, \tag{34}
\end{aligned}$$

where  $\Phi_1 = \frac{1}{c} 5\gamma^2 K L^2 (\nu^2 + 6K\epsilon^2)$  and  $\Phi_2 = \frac{2L\gamma\nu^2}{c}$ .

**Lemma A.1** ([35, Lemma 3]). *If  $\gamma \leq \frac{1}{8LK}$ , then*

$$\begin{aligned}
\mathbb{E}_t [\|\mathbf{y}_{t,i}^n - \mathbf{x}_t\|^2] &\leq 5K\gamma^2 (\nu^2 + 6K\epsilon^2) + 30K^2\gamma^2 \|\nabla f(\mathbf{x}_t)\|^2. \tag{35}
\end{aligned}$$

*Proof:* This lemma is a slight variation of Lemma 3 in [35] where instead of taking the arithmetic average over all  $N$  clients, we instead bound uniformly across all  $n$ . This change is straightforward, so we omit details for space.  $\square$

## B. Proof of Theorem 2

Since the objective function is an independent sum over  $n$  and has independent boundary constraints for  $P_n^t$ , we can find the minimizing values  $P_n^t$  by finding the roots of the gradient of the objective function and ensuring that they are within the upper and lower bounds. If no roots are within that set, one of the end points will minimize the function, so we only need to check those points.

To find the roots, we compute the gradient of the objective function for each  $n$  in (20)

$$\nabla f(q_n^t, P_n^t) = \left[ \begin{array}{l} -\frac{V}{N(q_n^t)^2} + \frac{V\lambda\ell}{B \log_2 \left( 1 + |h_n^t|^2 \frac{P_n^t}{N_0} \right)} + Z_n^t P_n^t \\ \frac{-V\lambda\ell |h_n^t|^2}{N_0 B \left( 1 + |h_n^t|^2 \frac{P_n^t}{N_0} \right) \left( \log_2 \left( 1 + |h_n^t|^2 \frac{P_n^t}{N_0} \right) \right)^2} q_n^t + Z_n^t q_n^t \end{array} \right]. \quad (36)$$

We first look at the partial derivative with respect to  $P_n^t$  and note that setting it equal to zero and dividing by  $q_n^t$  gives

$$0 = \frac{-V\lambda\ell |h_n^t|^2 / (N_0 B)}{\left( 1 + |h_n^t|^2 \frac{P_n^t}{N_0} \right) \left( \log_2 \left( 1 + |h_n^t|^2 \frac{P_n^t}{N_0} \right) \right)^2} + Z_n^t$$

which *does not depend* on  $q_n^t$ . Next, let  $A = \frac{V\lambda\ell |h_n^t|^2 (\log(2))^2}{N_0 B Z_n^t}$  and  $x = 1 + |h_n^t|^2 \frac{P_n^t}{N_0}$ , then we have something in the form of

$$A = x (\log(x))^2 = x (\log(1/x))^2.$$

By dividing both sides by 1/4, letting  $x' = \sqrt{\frac{A}{4}} \frac{1}{\sqrt{x}}$ , and rearranging, we have

$$\sqrt{A/4} = x' e^{x'}$$

that has a known solution of  $x' = W_k \left( \sqrt{\frac{A}{4}} \right)$  where  $W_k(\cdot)$  is the Lambert  $W$  function which solves  $w \exp w = z$  for  $w$ .

To get the critical point for  $P_n^t$ , we unwrap and substitute  $P_n^t = \frac{N_0}{|h_n^t|^2} (x - 1)$ , to get

$$P_n^{\text{opt}} = \frac{N_0}{|h_n^t|^2} \left( \frac{A}{4} W_k \left( \sqrt{\frac{A}{4}} \right)^{-2} - 1 \right) \quad (37)$$

which has a single root at  $k = 0$  since  $\sqrt{\frac{A}{4}} \geq 0$ .

## ACKNOWLEDGMENT

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

## REFERENCES

- [1] J. Perazzone, S. Wang, M. Ji, and K. S. Chan, "Communication-efficient device scheduling for federated learning using stochastic optimization," in *Proc. IEEE Int. Conf. Comput. Commun. (INFOCOM)*, May 2022, pp. 1449–1458.
- [2] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. Y. Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. 20th Int. Conf. Artif. Intell. Statist.*, vol. 54, 2017, pp. 1273–1282.
- [3] X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang, "On the convergence of FedAvg on non-IID data," in *Proc. Int. Conf. Learn. Represent.*, Jan. 2020, pp. 1–26.
- [4] S. P. Karimireddy, S. Kale, M. Mohri, S. Reddi, S. Stich, and A. T. Suresh, "SCAFFOLD: Stochastic controlled averaging for federated learning," in *Proc. 37th Int. Conf. Mach. Learn.*, vol. 119, Jul. 2020, pp. 5132–5143.
- [5] A. Mitra, R. H. Jaafar, G. J. Pappas, and H. Hassani, "Linear convergence in federated learning: Tackling client heterogeneity and sparse gradients," in *Proc. Adv. Neural Inf. Process. Syst.*, Jan. 2021, pp. 14606–14619.
- [6] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," in *Proc. 3rd Mach. Learn. Syst. Conf.*, vol. 2, 2020, pp. 429–450.
- [7] H. Yu, S. Yang, and S. Zhu, "Parallel restarted SGD with faster convergence and less communication: Demystifying why model averaging works for deep learning," in *Proc. Conf. Artif. Intell. (AAAI)*, 2019, vol. 33, no. 1, pp. 5693–5700.
- [8] H. Yang, M. Fang, and J. Liu, "Achieving linear speedup with partial worker participation in non-IID federated learning," in *Proc. ICLR*, Jan. 2021, pp. 1–23.
- [9] V. Mothukuri, R. M. Parizi, S. Pouriyeh, Y. Huang, A. Dehghanianha, and G. Srivastava, "A survey on security and privacy of federated learning," *Future Gener. Comput. Syst.*, vol. 115, pp. 619–640, Feb. 2021.
- [10] R. C. Geyer, T. Klein, and M. Nabi, "Differentially private federated learning: A client level perspective," 2017, *arXiv:1712.07557*.
- [11] S. Truex et al., "A hybrid approach to privacy-preserving federated learning," in *Proc. 12th ACM Workshop Artif. Intell. Secur.*, Nov. 2019, pp. 1–11.
- [12] S. Wang et al., "Adaptive federated learning in resource constrained edge computing systems," *IEEE J. Sel. Areas Commun.*, vol. 37, no. 6, pp. 1205–1221, Jun. 2019.
- [13] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," in *Proc. NIPS Workshop Private Multi-Party Mach. Learn.*, 2016, pp. 1–10.
- [14] P. Han, S. Wang, and K. K. Leung, "Adaptive gradient sparsification for efficient federated learning: An online learning approach," in *Proc. IEEE 40th Int. Conf. Distrib. Comput. Syst. (ICDCS)*, Nov. 2020, pp. 300–310.
- [15] F. Sattler, S. Wiedemann, K.-R. Müller, and W. Samek, "Robust and communication-efficient federated learning from non-IID data," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 31, no. 9, pp. 3400–3413, Sep. 2019.
- [16] D. Alistarh, D. Grubic, J. Li, R. Tomioka, and M. Vojnovic, "QSGD: Communication-efficient SGD via gradient quantization and encoding," in *Proc. Adv. NIPS*, vol. 30, Dec. 2017, pp. 1709–1720.
- [17] A. Albayoni, M. Safaryan, L. Condat, and P. Richtárik, "Optimal gradient compression for distributed and federated learning," 2020, *arXiv:2010.03246*.
- [18] A. Reisizadeh, A. Mokhtari, H. Hassani, A. Jadbabaie, and R. Pedarsani, "FedPAQ: A communication-efficient federated learning method with periodic averaging and quantization," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2020, pp. 2021–2031.
- [19] T. Nishio and R. Yonetani, "Client selection for federated learning with heterogeneous resources in mobile edge," in *Proc. IEEE Int. Conf. Commun.*, May 2019, pp. 1–7.
- [20] Y. J. Cho, J. Wang, and G. Joshi, "Towards understanding biased client selection in federated learning," in *Proc. Int. Conf. Artif. Intell. Statist.*, 2022, pp. 10351–10375.
- [21] Y. Fraboni, R. Vidal, L. Kameni, and M. Lorenzi, "Clustered sampling: Low-variance and improved representativity for clients selection in federated learning," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 3407–3416.
- [22] M. Ribeiro and H. Vikalo, "Communication-efficient federated learning via optimal client sampling," 2020, *arXiv:2007.15197*.
- [23] J. Goetz, K. Malik, D. Bui, S. Moon, H. Liu, and A. Kumar, "Active federated learning," 2019, *arXiv:1909.12641*.

- [24] H. H. Yang, Z. Liu, T. Q. S. Quek, and H. V. Poor, "Scheduling policies for federated learning in wireless networks," *IEEE Trans. Commun.*, vol. 68, no. 1, pp. 317–333, Jan. 2019.
- [25] J. Ren, Y. He, D. Wen, G. Yu, K. Huang, and D. Guo, "Scheduling for cellular federated edge learning with importance and channel awareness," *IEEE Trans. Wireless Commun.*, vol. 19, no. 11, pp. 7690–7703, Nov. 2020.
- [26] Y. Ruan, X. Zhang, S.-C. Liang, and C. Joe-Wong, "Towards flexible device participation in federated learning," in *Proc. Int. Conf. Artif. Intell. Statist. (AISTATS)*, vol. 130, Apr. 2021, pp. 3403–3411.
- [27] X. Gu, K. Huang, J. Zhang, and L. Huang, "Fast federated learning in the presence of arbitrary device unavailability," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 34, Jan. 2021, pp. 12052–12064.
- [28] S. Wang and M. Ji, "A unified analysis of federated learning with arbitrary client participation," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 35, Jan. 2022, pp. 19124–19137.
- [29] M. Chen, Z. Yang, W. Saad, C. Yin, H. V. Poor, and S. Cui, "A joint learning and communications framework for federated learning over wireless networks," *IEEE Trans. Wireless Commun.*, vol. 20, no. 1, pp. 269–283, Jan. 2020.
- [30] Z. Yang, M. Chen, W. Saad, C. S. Hong, and M. Shikh-Bahaei, "Energy efficient federated learning over wireless communication networks," *IEEE Trans. Wireless Commun.*, vol. 20, no. 3, pp. 1935–1949, Mar. 2021.
- [31] M. M. Wadu, S. Samarakoon, and M. Bennis, "Federated learning under channel uncertainty: Joint client scheduling and resource allocation," in *Proc. IEEE Wireless Commun. Netw. Conf.*, May 2020, pp. 1–6.
- [32] T. Huang, W. Lin, W. Wu, L. He, K. Li, and A. Y. Zomaya, "An efficiency-boosting client selection scheme for federated learning with fairness guarantee," *IEEE Trans. Parallel Distrib. Syst.*, vol. 32, no. 7, pp. 1552–1564, Jul. 2021.
- [33] Z. Zhou, S. Yang, L. Pu, and S. Yu, "CEFL: Online admission control, data scheduling, and accuracy tuning for cost-efficient federated learning across edge nodes," *IEEE Internet Things J.*, vol. 7, no. 10, pp. 9341–9356, Oct. 2020.
- [34] X. Lian, C. Zhang, H. Zhang, C. Hsieh, W. Zhang, and J. Liu, "Can decentralized algorithms outperform centralized algorithms? A case study for decentralized parallel stochastic gradient descent," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, Dec. 2017, pp. 5330–5340.
- [35] S. J. Reddi et al., "Adaptive federated optimization," in *Proc. Int. Conf. Learn. Represent.*, Jan. 2020, pp. 1–38.
- [36] S. Wang, J. Perazzone, M. Ji, and K. S. Chan, "Federated learning with flexible control," in *Proc. IEEE Conf. Comput. Commun.*, May 2023, pp. 1–10.
- [37] W. Shi, S. Zhou, Z. Niu, M. Jiang, and L. Geng, "Joint device scheduling and resource allocation for latency constrained wireless federated learning," *IEEE Trans. Wireless Commun.*, vol. 20, no. 1, pp. 453–467, Sep. 2020.
- [38] L. Cui, X. Su, Y. Zhou, and J. Liu, "Optimal rate adaption in federated learning with compressed communications," in *Proc. IEEE Conf. Comput. Commun.*, May 2022, pp. 1459–1468.
- [39] M. J. Neely, "Stochastic network optimization with application to communication and queueing systems," *Synth. Lectures Commun. Netw.*, vol. 3, no. 1, pp. 1–211, Sep. 2010.
- [40] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre, "Manopt, a MATLAB toolbox for optimization on manifolds," *J. Mach. Learn. Res.*, vol. 15, pp. 1455–1459, Aug. 2014. [Online]. Available: <https://www.manopt.org>
- [41] A. Krizhevsky and G. Hinton, "Learning multiple layers of features from tiny images," Univ. Toronto, Toronto, ON, Canada, Tech. Rep., 2009.
- [42] T.-M. H. Hsu, H. Qi, and M. Brown, "Measuring the effects of non-identical data distribution for federated visual classification," 2019, *arXiv:1909.06335*.
- [43] N. S. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, and P. Tang, "On large-batch training for deep learning: Generalization gap and sharp minima," in *Proc. Int. Conf. Learn. Represent.*, Sep. 2016, pp. 1–16.

**Jake B. Perazzone** received the B.S. degree in electrical engineering from The College of New Jersey, Ewing, NJ, USA, and the M.S. and Ph.D. degrees in electrical engineering from Lehigh University, Bethlehem, PA, USA. During this work, he was a Research Scientist with DEVCOM Army Research Laboratory, Adelphi, MD, USA. He is currently a Senior Machine Learning Engineer with DeepSig, Inc., Arlington, VA, USA. His research interests include machine learning for wireless as well as wireless security.

**Shiqiang Wang** (Senior Member, IEEE) received the Ph.D. degree from Imperial College London, U.K., in 2015. He is a Staff Research Scientist at IBM T. J. Watson Research Center, NY, USA. His research focuses on the intersection of artificial intelligence (AI), distributed computing, networking, and optimization. Currently emphasizing quality and efficiency aspects related to distributed data and models, with a broad range of applications including distributed data analytics, efficient model training and inference, edge-based artificial intelligence (Edge AI), and large language models (LLMs). He has made foundational contributions to edge computing and federated learning that generated both academic and industrial impact. He serves as an Associate Editor of the *IEEE TRANSACTIONS ON MOBILE COMPUTING* and *IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS*. He has also been actively organizing workshops at the intersection of edge computing and machine learning, and regularly participates in technical program committees (TPCs) of prominent conferences and review panels of research grants. He received the IEEE Communications Society (ComSoc) Leonard G. Abraham Prize in 2021, IEEE ComSoc Best Young Professional Award in Industry in 2021, IBM Outstanding Technical Achievement Awards (OTAA) in 2019, 2021, 2022, and 2023, multiple Invention Achievement Awards from IBM since 2016, Best Paper Finalist of the IEEE International Conference on Image Processing (ICIP) 2019, and Best Student Paper Award of the Network and Information Sciences International Technology Alliance (NIS-ITA) in 2015.

**Mingyue Ji** (Member, IEEE) received the Ph.D. degree from the Ming Hsieh Department of Electrical and Computer Engineering, University of Southern California, in 2015, where he received the USC Annenberg Fellowship from 2010 to 2014. He subsequently was a Staff II System Design Scientist with Broadcom Inc., from 2015 to 2016. He is currently an Associate Professor with the Department of Electrical and Computer Engineering, University of Florida. His research interests span a broad spectrum, including cloud and edge computing, distributed machine learning, and 5G and beyond wireless communications, networking, and sensing. His research activities cover fundamental theory study, algorithm design and analysis, and practical system implementation and experimentation. He received the NSF CAREER Award in 2022, the IEEE Communications Society Leonard G. Abraham Prize for the Best IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS (JSAC) Paper in 2019, the Best Paper Awards at 2021 IEEE GLOBECOM Conference, 2015 IEEE ICC Conference and 2024 IEEE ISICN Conference, the Best Student Paper Award at 2010 IEEE European Wireless Conference, the 2022 Outstanding ECE Teaching Award, and the 2023 Outstanding ECE Research Award at the University of Utah. He has been serving as Associate Editors for *IEEE TRANSACTIONS ON INFORMATION THEORY* since 2022 and *IEEE TRANSACTIONS ON COMMUNICATIONS* since 2020.

**Kevin Chan** (Senior Member, IEEE) received the B.S. degree in electrical and computer engineering and engineering and public policy from Carnegie Mellon University, Pittsburgh, PA, USA, and the M.S. and Ph.D. degrees in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, GA, USA. He is currently the Lead Research Electronics Engineer and Network Science Team Lead in the Networks, Cyber, and Computational Sciences Division at the U.S. Army Combat Capabilities Development Command, Army Research Laboratory, Adelphi, MD, USA. He is actively involved in research on network science, edge computing, and cybersecurity. He received the 2021 IEEE Communications Society Leonard G. Abraham Prize and multiple best paper awards. He also recently served as Co-Editor of the *IEEE Communications Magazine–Military Communications and Networks Series*.