

Client Selection and Bandwidth Allocation in Wireless Federated Learning Networks: A Long-Term Perspective

Jie Xu, Heqiang Wang

Abstract

This paper studies federated learning (FL) in a classic wireless network, where learning clients share a common wireless link to a coordinating server to perform federated model training using their local data. In such wireless federated learning networks (WFLNs), optimizing the learning performance depends crucially on how clients are selected and how bandwidth is allocated among the selected clients in every learning round, as both radio and client energy resources are limited. While existing works have made some attempts to allocate the limited wireless resources to optimize FL, they focus on the problem in individual learning rounds, overlooking an inherent yet critical feature of federated learning. This paper brings a new long-term perspective to resource allocation in WFLNs, realizing that learning rounds are not only temporally interdependent but also have varying significance towards the final learning outcome. To this end, we first design data-driven experiments to show that different temporal client selection patterns lead to considerably different learning performance. With the obtained insights, we formulate a stochastic optimization problem for joint client selection and bandwidth allocation under long-term client energy constraints, and develop a new algorithm that utilizes only currently available wireless channel information but can achieve long-term performance guarantee. Further experiments show that our algorithm results in the desired temporal client selection pattern, is adaptive to changing network environments and far outperforms benchmarks that ignore the long-term effect of FL.

I. INTRODUCTION

Mobile devices nowadays generate a massive amount of data each day. This rich data has the potential to power a wide range of machine learning (ML)-based applications, such as learning the activities of smart phone users, predicting health events from wearable devices or adapting

J. Xu and H. Wang are with the Department of Electrical and Computer Engineering, University of Miami, FL, USA

to pedestrian behavior in autonomous vehicles. Due to the growing storage and computational power of mobile devices as well as privacy concerns associated with uploading personal data, it is increasingly attractive to store and process data directly on each mobile device. The aim of “federated learning” (FL) [1] is to enable mobile devices to collaboratively learn a shared ML model with the coordination of a central server while keeping all the training data on device, thereby decoupling the ability to do ML from the need to upload/store the data in the cloud.

This paper focuses on FL in a classic wireless network setting where the clients, e.g., mobile devices, share a common wireless link to the server. We call this system a *wireless federated learning network* (WFLN). The network operates for a number of learning rounds as follows: in each round, the clients download the current ML model from the server, improve it by learning from their local data, and then upload the individual model updates to the server via the wireless link; the server then aggregates the local updates to improve the shared model. Similar to a traditional throughput-oriented wireless network, the limited wireless network resources require the WFLN to determine in each round which clients access the wireless channel to upload the model updates and how much bandwidth is allocated to each client. However, due to the specific application in consideration, namely FL, the resource allocation objective and consequently the outcome can be very different from, e.g., throughput maximization.

Optimizing WFLNs faces unique challenges compared to optimizing either FL or the traditional wireless networks. On the one hand, the wireless network sets resource constraints on performing FL as the finite wireless bandwidth limits the number of clients that can be selected in each round, and the selection must be adaptive to the highly variable wireless channel conditions. On the other hand, FL is likely to change the way wireless networks should be optimized as model training is a complex *long-term* process where decisions across rounds are interdependent and collectively decide the final training performance. Further, since mobile devices often have finite energy budgets due to, e.g., a finite battery, the number of rounds each individual mobile device can participate during the entire course of FL is also limited. An extremely crucial yet largely overlooked question is: does learning in different rounds contribute the same or differently to the final learning outcome and hence should the wireless resources be allocated discrepantly across rounds? Without a good understanding of its answer, conventional wireless network optimization approaches that treat each time slot independently and equally may lead to considerably suboptimal FL performance.

This paper aims to formalize this fundamental problem of client selection and bandwidth

allocation in WFLNs and derive critical knowledge to enable the efficient operation of these networks. We study how resources (i.e., bandwidth and energy) should be allocated among clients in each learning round as well as *across* rounds given finite client energy budgets in a volatile network environment. Our main contributions are summarized as follows.

(1) While existing works [2], [3] have shown that including more clients in FL generally improves the learning performance, there is little understanding of how this improvement depends on the learning rounds. For a fixed total number of selected clients during the entire course of FL, should client selection be uniform across rounds or more biased toward the early/late FL rounds? Although analytical characterization seems extremely difficult, we show in two representative ML tasks, i.e., image classification and text generation, that selecting more clients in the later FL rounds not only achieves higher accuracy and lower training loss but also is more robust than selecting more clients in the early FL rounds. This finding, to our best knowledge, is the first that relates the temporal client selection pattern to the final FL performance.

(2) With the understanding of a desired temporal client selection pattern, we formulate a *long-term* client selection and bandwidth allocation problem for a finite number of FL rounds under finite energy constraints of individual clients. Because wireless channel conditions vary over time but future conditions are unpredictable, we leverage the Lyapunov technique [4] to convert the long-term problem into a sequence of per-round problems via a virtual energy deficit queue for each client. A new online optimization algorithm called OCEAN is proposed, which in each FL round solves a finite number of convex optimization problems using only currently available wireless information, and hence the algorithm is practical and has low complexity.

(3) We prove that OCEAN achieves the FL performance of the desired client selection pattern within a bounded gap while approximately satisfying the energy constraints of the clients. Specifically, OCEAN demonstrates an $[O(1/V), O(\sqrt{V})]$ learning-energy tradeoff where V is an algorithm parameter. In addition, we investigate the structure of the client selection and bandwidth allocation outcome. Our findings are two-fold: in each round, clients are selected according to a priority metric, which is the ratio of the client's current energy deficit queue length and its current wireless channel state. However, among the selected ones, more bandwidth is allocated to clients with a lower priority (i.e., worse channel and larger energy deficit queue length). This is in stark contrast to a traditional throughput-oriented wireless network where more bandwidth is allocated to clients with a better channel condition in order to maximize throughput.

II. RELATED WORK

Since the proposal of FL [1], [5], a lot of research effort has been devoted to tackling various challenges in this new distributed machine learning framework, including developing new optimization and model aggregation algorithms [6]–[8], handling non-i.i.d. and unbalanced datasets [9]–[11], and preserving model privacy [12]–[16] etc. Among these challenges, improving the communication efficiency of FL has been a key challenge due to the tension between uploading a large amount of data for model aggregation and the limited network resource to support this transmission. In this regard, a strand of literature focuses on modifying the FL algorithm itself to reduce the communication burden on the network, e.g., updating clients with significant training improvement [17], compressing the gradient vectors via quantization [18], or accelerating training using sparse or structured updates [1], [19]. Hierarchical FL networks [20] have also been proposed where multiple edge servers perform partial model aggregation first, whose outputs are further aggregated by a cloud server. Recognizing the unique physical property of wireless transmission, [2], [21] propose analog model aggregation over the air, provided that a very stringent synchronization is available.

As wireless networks are the envisioned main deployment scenario of FL, how to optimally allocate the limited bandwidth and energy resources for FL has also received much attention. Many existing works [22]–[25] study the inherent trade-off between local model update and global model aggregation, e.g., to adapt the frequency of global aggregation [22] or to optimize uplink transmission power/rate and the local update CPU frequency [24], [25]. In all these works, all clients participate in every FL round. Although both empirical studies [2], [3] and theoretical analysis [26] show that including more clients improves the FL convergence speed, the limited bandwidth of wireless networks cannot support many clients to upload their local updates at the same time. For FL at scale, client scheduling policies, which select only a subset of clients in every round, are necessary. In [27], the convergence performance of FL under three basic scheduling policies, namely random, round-robin and proportional fair, is analyzed. Different types of joint bandwidth allocation and client scheduling policies, e.g., [3], [28]–[32], have been proposed to either minimize the learning loss or the training time. However, their optimization problems are formulated by considering individual FL rounds separately or treating every FL round equally, and hence the same network resources are allocated across learning rounds. Our paper differs from these works in that we explicitly consider the varying significance of FL

rounds and study a long-term bandwidth allocation and client selection problem under long-term energy constraints and with uncertain wireless channel information.

III. IMPACT OF TEMPORAL CLIENT SELECTION PATTERN

Existing works [2], [3] have shown that the FL performance (in terms of training loss and prediction accuracy) can be improved by selecting more clients in each round. However, selecting more clients is not always possible if each client is subject to a long-term energy constraint due to, e.g., a finite battery: selecting more clients in early learning rounds depletes the battery of the clients and hence fewer clients can be selected in later learning rounds. Hence, even with the same average number of selected clients, the temporal pattern can be considerably different, yet there is little understanding of how the temporal pattern affects the final FL outcome. In this section, we design two experiments to show that the temporal client pattern indeed has a considerable impact on the final FL performance.

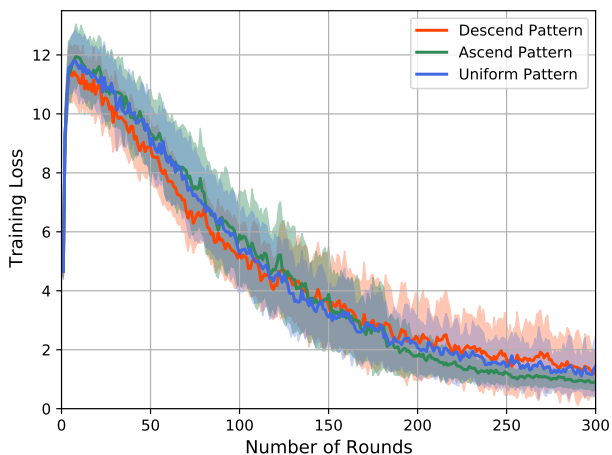


Fig. 1: Training Loss (MNIST Dataset)




Fig. 2: Accuracy (MNIST Dataset)

A. Image Classification on the MNIST Dataset

Our first experiment is conducted using the TensorFlow Federated (TFF) framework [33] on the MNIST dataset for image classification. A deep neural network (DNN) classifier is trained on 10 clients (index from 1 through 10) using the FedAvg algorithm [1] over a total of 300 rounds. Three temporal selection patterns are investigated: **Uniform** – in each round, 5 clients are randomly selected to upload their model parameters; **Ascend** – the number of selected clients




Fig. 3: Training Loss (Shakespeare Dataset)

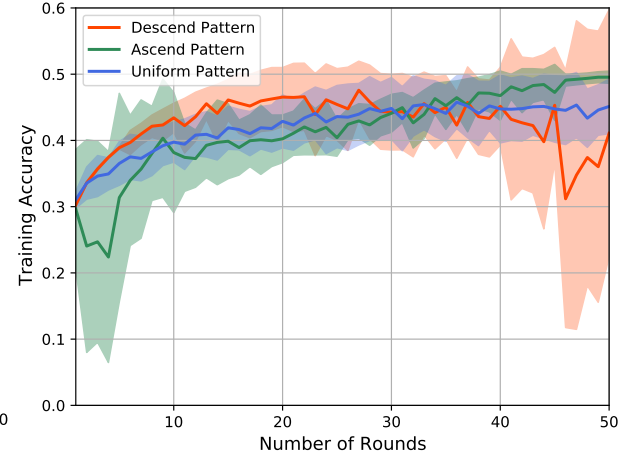


Fig. 4: Accuracy (Shakespeare Dataset)

gradually increases from 1 to 10 over 300 rounds with an average number of 5 clients selected per round; **Descend** – the number of selected clients gradually decreases from 10 to 1 over 300 rounds with an average number of 5 clients selected per round.

Figures 1 and 2 illustrate the training loss and the prediction accuracy, respectively, over 300 rounds for the three temporal patterns. Each curve is generated by averaging over 60 runs and the standard deviation of these curves are also shown in the figures. As can be seen, although the average number of selected clients is the same, different temporal patterns result in different training loss and prediction accuracy by the end of the 300 rounds. In particular, **Ascend** results in the best performance compared to **Uniform** and **Descend**. There is a good reason behind this result: early learning rounds are “easy” rounds where the learning performance is less sensitive to the number of selected clients. Hence, even if **Ascend** selects fewer clients in the early rounds, learning speed is minimally affected. However, the later learning rounds are the more “difficult” rounds, and to push accuracy even higher requires more clients to update the shared model using their data. In fact, not only **Ascend** wins in training loss and accuracy, but it is also much more robust as the standard deviation is much smaller. This is again because more clients participate in model updating towards the end of learning, which can smooth out abrupt changes in the learned model of individual clients.

B. Text Generation on the Shakespeare Dataset

To verify that the above findings are generalizable, we conduct a similar experiment on a text generation task. We utilize the decentralized text generation dataset based on *The Complete Works of Shakespeare* provided in the TensorFlow Federated tutorial [33]. A Recurrent Neural Network with eager execution is pre-trained on the text from Charles Dickens' *A Tale of Two Cities* and *A Christmas Carol* as the initial model, and FL is used to fine-tune this model for the Shakespeare dataset. As can be seen in Figures 3 and 4, although the task and the dataset are very different, similar observations can be made as before: the **Ascend** selection pattern significantly outperforms **Descend** and **Uniform** in terms of training loss, accuracy and robustness.

We note that the exact optimal temporal selection pattern seems impossible to analytically characterize, which would also change across different learning tasks, models, datasets and algorithms. However, as the general ascending trend leads to considerable performance improvement, it offers valuable guidance for designing client selection schemes across rounds.

IV. WIRELESS FEDERATED LEARNING NETWORK MODEL

With the insights obtained in Section III, we now move on to optimize the WFLN. Consider a WFLN with one server and K clients, indexed by the set $\mathcal{K} = \{1, \dots, K\}$. Each participating client $k \in \mathcal{K}$ has a local dataset \mathcal{D}_k . In the supervised learning case, \mathcal{D}_k defines the collection of data samples given as a set of input-output pairs $\{x_i, y_i\}_{i=1}^{\mathcal{D}_k}$, where $x_i \in \mathbb{R}^d$ is a d -dimensional input feature vector, and $y_i \in \mathbb{R}$ is the ground-truth output label. This data can be generated through the usage of the client via mobile applications and can be employed for various ML tasks, e.g., user activity prediction or health event prediction.

FL iterates between two steps: 1) the server updates the global model by aggregating local models transmitted over a multi-access channel by the clients; 2) the clients update their local models using the global model broadcasted by the server. We call each iteration a *learning round*. WFLN has to decide in each round which clients upload their local model updates depending on their wireless channel condition and remaining battery to maximize the learning performance. We use $a_k^t \in \{1, 0\}$ to denote whether or not client k is selected in round t , and $\mathbf{a}^t = (a_1^t, \dots, a_K^t)$ collects the overall client selection decisions.

A. Client Energy Consumption

For a selected client k in round t (i.e., $a_k^t = 1$), it incurs energy consumption due to uploading the local updates to the edge server via the wireless channel. We consider a specific wireless multi-access scheme, i.e., orthogonal frequency-division multiple access (OFDMA) for local model uploading with a total bandwidth B . Let $b_k^t \in [0, 1]$ be the bandwidth allocation ratio for client k in round t , and hence its allocated bandwidth is $b_k^t B$. Let $\mathbf{b}^t = (b_1^t, \dots, b_K^t)$. Bandwidth allocation must satisfy $\sum_{k \in \mathcal{K}} b_k^t = 1, \forall t$. Clearly, if $a_k^t = 0$, namely client k is not selected in round t , then it is the best not to allocate any bandwidth to this client, i.e., $b_k^t = 0$. On the other hand, if $a_k^t = 1$, then we require that at least a minimum bandwidth b_{min} is allocated to client k , i.e., $b_k^t \geq b_{min}$. This is because practical systems cannot assign an arbitrarily small bandwidth to an individual client. In addition, a close-to-zero bandwidth allocation will require an extremely high transmit power and hence result in an extremely high energy consumption to achieve a target transmission rate. To make the problem feasible, we assume $b_{min} \leq 1/K$.

Let p_k^t denote the transmission power (in Watt/Hz) of client k in round t . The achievable rate (in bit/s), denoted by r_k^t , can be written according to the Shannon's formula as

$$r_k^t = b_k^t B \log_2 \left(1 + \frac{p_k^t (h_k^t)^2}{N_0} \right) \quad (1)$$

where N_0 is the variance of the complex white Gaussian channel noise and h_k^t is the channel state of client k in round t . Let L denote the data size of the adopted machine learning model (in bit), then the time needed to upload the local model update to the edge server is $\tau_k = L/r_k^t$. For a target upload time deadline $\bar{\tau}$, the required transmission power can be derived using (1) and hence the transmission energy consumption of client k is

$$E(a_k^t, b_k^t | h_k^t) = \frac{\bar{\tau} N_0 B b_k^t}{(h_k^t)^2} \left(2^{\frac{L}{\bar{\tau} B b_k^t}} - 1 \right) a_k^t \quad (2)$$

B. System Learning Performance

Existing works and our empirical study show that in order to accelerate learning, it is desirable for the WFLN to select as many clients as possible in each round. However, client selection is constrained by finite radio and battery resources. Therefore, the WFLN must judiciously select clients to perform federated learning in each round, without quickly draining clients' battery and

causing insufficient model updates in later communication rounds. To this end, we introduce the following metric to describe the FL performance in round t :

$$U^t(\mathbf{a}^t) = \eta^t \sum_{k=1}^K a_k^t \quad (3)$$

where η^t is a temporal weight to capture the varying significance of selecting more clients in different learning rounds. As suggested by Section III, an increasing sequence of η^t often results in better FL performance as more clients are likely to be selected in later rounds of learning.

We note, however, although the above metric will facilitate our subsequent resource allocation, it does not exactly characterize the FL speed or accuracy, which is extremely difficult, if not impossible, to model due to the complex and non-convex nature of many ML algorithms.

C. Problem Formulation

As we emphasize the long-term performance and final outcome of FL, the goal is to maximize the weighted sum of selected clients defined in (3) for a total number of T learning rounds while satisfying the long-term energy budget constraints of individual clients, through joint client selection \mathbf{a}^t and bandwidth allocation \mathbf{b}^t in every round $t = 0, \dots, T - 1$. Although the performance metric defined in (3) is artificial, we will relate it to the actual FL performance (i.e., training loss and accuracy) in experiments. Formally, the problem that we aim to solve is

$$\mathbf{P1} \quad \max_{\mathbf{a}^0, \mathbf{b}^0, \dots, \mathbf{a}^{T-1}, \mathbf{b}^{T-1}} \sum_{t=0}^{T-1} U^t(\mathbf{a}^t) \quad (4)$$

$$\text{s.t.} \quad \sum_{t=0}^{T-1} E(a_k^t, b_k^t | h_k^t) \leq H_k, \forall k \quad (5)$$

$$b_{min} \leq b_k^t \leq 1, \forall k, \forall t, \quad \sum_{k=1}^K b_k^t = 1, \forall t \quad (6)$$

$$a_k^t \in \{0, 1\}, \forall k, \forall t \quad (7)$$

Constraint (5) requires that the total energy consumption over the T rounds for each client k does not exceed an energy budget H_k (e.g., battery capacity or energy limit set by the client). Constraint (6) is the feasibility condition on the bandwidth allocation. Constraint (7) is the feasibility condition on the client selection.

So far we have formulated a long-term optimization problem for client selection and bandwidth allocation in WFLNs. However, several challenges impede the derivation of the optimal solution

to **P1**. The first is the lack of future information: optimally solving **P1** requires complete offline information (i.e., channel conditions) over the entire FL period (i.e., T learning rounds) that is very difficult to accurately predict in advance. Furthermore, **P1** belongs to mixed-integer nonlinear programming and is difficult to solve, even if the long-term future information is accurately known a priori. Thus, these challenges demand an online approach that can efficiently make joint client selection and bandwidth allocation decisions without foreseeing the far future.

D. Offline Benchmark: R -Round Lookahead Algorithm

Before we propose the online algorithm, we first introduce an offline algorithm with R -round lookahead information (i.e., the channel information in the next R learning rounds are *assumed* to be known) as a benchmark. Specifically, we divide the entire FL period into $M \geq 1$ frames, each having $R \geq 1$ learning rounds such that $T = MR$, and present the following problem formulation:

$$\mathbf{P2} : \max_{\mathbf{a}^0, \mathbf{b}^0, \dots, \mathbf{a}^{T-1}, \mathbf{b}^{T-1}} \sum_{t=mR}^{(m+1)R-1} U^t(\mathbf{a}^t) \quad (8)$$

$$\text{s.t.} \quad \sum_{t=mR}^{(m+1)R-1} E(a_k^t, b_k^t | h_k^t) \leq H_k/M, \forall k \quad (9)$$

Constraints (6), (7)

Essentially, **P2** defines a family of offline algorithms parameterized by the lookahead window size R . Clearly, there exists at least one sequence of joint client selection and bandwidth allocation decisions that satisfies all constraints of **P2** (e.g., no client is selected in any round in each frame). We denote the optimal learning performance for the m -th frame by U_m^* , for $m = 0, \dots, M-1$, considering all the decisions that satisfy the constraints and have perfect information over the frame. Thus, the optimal long-term learning performance achieved by the oracle's optimal R -round lookahead algorithm is given by $\sum_{m=0}^{M-1} U_m^*$.

We note that because of the assumed lookahead information, the R -round lookahead algorithms are impractical (unless $R = 1$). The purpose of introducing these algorithms is only to use them as a benchmark for our practical online algorithm to be proposed in the next section.

V. ONLINE CLIENT SELECTION AND BANDWIDTH ALLOCATION

In this section, we develop the One Client selection and bandwidth allocation N algorithm, called OCEAN, and then characterize its structural properties. We also prove that it is efficient

Algorithm 1 OCEAN

- 1: **Input:** $q_k(0) = 0, \forall k$ and R
 - 2: **for** $t = 1, 2, \dots, T$ **do**
 - 3: **if** $t = mR, \forall m = 1, \dots, M - 1$ **then**
 - 4: $q_k(t) \leftarrow 0, \forall k$ and $V \leftarrow V_m$
 - 5: **end if**
 - 6: Observe the current channel state $h_k^t, \forall k$
 - 7: Solve **P3**
 - 8: Update energy queue according to (10)
 - 9: **end for**
-

compared to the optimal offline algorithm with R -round lookahead information.

A. The OCEAN Algorithm

A major challenge of directly solving **P1** is that the long-term energy constraint of the clients couples the client selection and bandwidth allocation decisions across different learning rounds: selecting more clients in the current round reduces the bandwidth allocated to each individual client, thereby increasing the energy consumption of these clients; furthermore, more energy consumption in the current round potentially reduces the energy budget available for future FL rounds, and yet the decisions have to be made without foreseeing the future. To address this challenge, we leverage the Lyapunov technique and construct a virtual energy deficit queue $q_k(t)$ for each client k to guide the client selection and bandwidth allocation decisions to follow the long-term energy constraint. The virtual energy queue of client k starts with $q_k(0) = 0, \forall k$, and is updated at the end of each round t as follows

$$q_k(t+1) = [E(a_k^t, b_k^t | h_k^t) - H_k/T + q_k(t)]^+ \quad (10)$$

where $[\cdot]^+ = \max\{\cdot, 0\}$. Hence, $q_k(t)$ is the queue length indicating the deviation of the current energy consumption of client k from its long-term energy constraint H_k . Let $\mathbf{q}(t) = (q_1(t), q_2(t), \dots, q_K(t))$ collect the energy deficit queues for all clients.

We now present OCEAN in Algorithm 1. OCEAN is purely online and requires only the currently available channel state information as inputs (i.e. $h_k(t), \forall k$). We use V_0, V_1, \dots, V_{M-1} to denote a sequence of positive control parameters to dynamically adjust the tradeoff between

maximizing the number of selected clients and minimizing energy consumption over the M frames, each having R communication rounds. The importance of the control parameters will be revisited in Section V.C. In every round t , we aim to solve the following per-round problem:

$$\mathbf{P3} \quad \max_{\mathbf{a}^t, \mathbf{b}^t} V \cdot U^t(\mathbf{a}^t) - \sum_{k=1}^K q_k(t) E(a_k^t, b_k^t | h_k^t) \quad (11)$$

$$\text{s.t.} \quad (6), (7) \quad (12)$$

By considering the additional term $\sum_{k=1}^K q_k(t) E(a_k^t, b_k^t | h_k^t)$, the system takes into account the energy deficit of the clients during the current round's client selection and bandwidth allocation. As a consequence, when $q_k(t)$ is larger, minimizing the energy deficit is more critical. Thus, our algorithm works following the philosophy of "if violate the energy constraint, then use less energy", and the energy deficit queue maintained without foreseeing the future guides the system towards meeting the energy constraints of the clients. OCEAN decomposes the long-term optimization problem into a series of per-round problems **P3**. For a more rigorous derivation of this decomposition, please refer to the proof of Theorem 1. Now, to complete OCEAN, it remains to solve **P3**, which however is still very difficult.

B. Solving the Per-Round Problem

The per-round problem **P3** is a difficult mixed-integer problem. To see more clearly how the objective function depends on \mathbf{a}^t and \mathbf{b}^t , we write out and rearrange it as follows

$$\max_{\mathbf{a}^t, \mathbf{b}^t} \sum_{k=1}^K \left[V \eta^t - q_k(t-1) \cdot \frac{\bar{\tau} N_0 B b_k^t}{(h_k^t)^2} \left(2^{\frac{L}{\bar{\tau} B b_k^t}} - 1 \right) \right] a_k^t \quad (13)$$

Notice that a_k^t is a binary integer variable and b_k^t is a continuous variable in $[b_{\min}, 1]$. In general, mixed-integer problems are difficult to solve and often there is no polynomial-time optimal algorithm. Fortunately, our problem **P3** exhibits a special structure and we are able to exploit this structure to develop an algorithm that returns the optimal solution by solving at most K convex optimization problems. To simplify the notations, we drop the index t in this subsection.

Our algorithm to solve **P3**, called OCEAN-P, incrementally adds clients into the selection set S based on a metric $\rho_k \triangleq \frac{q_k(t)}{(h_k^t)^2}$, which we call the *selection priority* (the lower value, the higher priority). Initially, all clients with $\rho_k = 0$ (which also means $q_k(t) = 0$ as $(h_k^t)^2$ is always positive) are added into S . We denote this initial set by S^0 . Then, clients with $\rho_k > 0$ are

added into S one by one in the ascending order of ρ_k , and for each possible selection set, the corresponding bandwidth allocation is computed by solving the following optimization problem

$$\mathbf{P4} \quad \max_{\{b_k^t\}_{k \in S-S^0}} \sum_{k \in S-S^0} \left(V\eta^t - \rho_k N_0 \tilde{\tau} B b_k \left(2^{\frac{L}{\tilde{\tau} B b_k}} - 1 \right) \right) \quad (14)$$

$$\text{s.t.} \quad \sum_{k \in S-S^0} b_k^t = 1 - |S^0| \cdot b_{min} \quad (15)$$

$$b_k \geq b_{min}, \forall k \in S - S^0 \quad (16)$$

Let $\mathbf{b}^*(S)$ be the optimal bandwidth allocation for a given selection set S , and $W^*(S)$ be the optimal value. Clearly, for the initial set S^0 , $W^*(S^0) = \eta^t |S^0|$ as $\rho_k = 0, \forall k \in S^0$. The number of selection sets that can possibly emerge following the above set expanding rule, which are collected in \mathcal{S} , is at most K . Finally, the implemented optimal selection is $S^* = \arg \max_{S \in \mathcal{S}} W^*(S)$ and the implemented optimal bandwidth allocation is $\mathbf{b}^* = \mathbf{b}^*(S^*)$.

Because there are K clients and in every iteration, one more client is added into S , we ensure that the algorithm only needs to solve at most K optimization problems **P4** to return the optimal solution. In fact, we can reduce the number of times for solving **P4** by adding a termination condition: if for some S , its optimal bandwidth allocation $\mathbf{b}^*(S)$ results in $\eta^t - \rho_k N_0 \tilde{\tau} B b_k \left(2^{\frac{L}{\tilde{\tau} B b_k}} - 1 \right) < 0$ for the last added client k , then the algorithm stops adding more clients into the selection set. This termination condition can significantly reduce the number of convex optimization problems to be solved when K is large. The pseudocode of OCEAN-P is given in Algorithm 2. Next, we first prove that **P3** is a convex optimization problem and then prove the optimality of OCEAN-P.

Lemma 1. *The function $f(x) = x(2^{\frac{\beta}{x}} - 1)$ where $\beta > 0$ is decreasing and convex in $x \in (0, \infty)$ and is increasing and concave in $x \in (-\infty, 0)$.*

Proof. See Appendix A. □

Lemma 1 readily proves that **P4** is a convex optimization problem as $b_k^t \in [b_{min}, 1)$. Convex optimization problems are extensively studied in the literature and many efficient algorithms [34] and mature software tools (such as CVX [35] and SciPy [36]) exist. Next, we prove that our algorithm returns the optimal solution by solving **P4** at most K times.

Theorem 1. *OCEAN-P returns the optimal solution to the per-round problem **P3** by solving at most K convex optimization problems.*

Algorithm 2 OCEAN-P

- 1: **Input:** $q_k(0) = 0, \forall k$
 - 2: Rank the clients according to ρ . Hence we have $\rho_1 \leq \rho_2 \leq \dots \leq \rho_K$
 - 3: Set $S^0 = \{k : \rho_k = 0\}$, $S = S^0$, and $\mathcal{S} = \{S^0\}$.
 - 4: **for** $k = |S^0| + 1, \dots, K$ **do**
 - 5: Update $S = S \cup \{k\}$
 - 6: Solve **P4** and obtain $\mathbf{b}^*(S)$ and $W^*(S)$
 - 7: **if** $V\eta^t - \rho_k N_0 \tilde{\tau} B b_k^t \left(2^{\frac{L}{\tilde{\tau} B b_k^t}} - 1\right) < 0$ **then**
 - 8: Stop iteration
 - 9: **else**
 - 10: Add S to \mathcal{S} , i.e. $\mathcal{S} = \mathcal{S} \cup \{S\}$
 - 11: **end if**
 - 12: **end for**
 - 13: Find $S^* = \arg \max_{S \in \mathcal{S}} W^*(S)$
 - 14: Return \mathbf{a}^* where $a_k^* = \mathbf{1}\{k \in S^*\}, \forall k$ and $\mathbf{b}^* = \mathbf{b}^*(S^*)$
-

Proof. See Appendix B. □

C. Structural Results and Performance Analysis

In this subsection, we first investigate the structure of the optimal solution produced by OCEAN-P in every round, and then characterize the performance of OCEAN.

In Theorem 1, we have already proven a thresholding result on the client selection, namely only clients whose selection priority ρ_k is below a threshold are selected to participate in a FL round. Proposition 1 characterizes how bandwidth is allocated among the selected clients and their incurred energy consumption.

Proposition 1. *In any learning round t , the allocated bandwidth $b_k^{t,*}$ of a selected client k and its weighted energy consumption $q_k(t)E_k^t(b_k^{t,*})$ are non-decreasing with ρ_k^t .*

Proof. See Appendix C. □

Theorem 1 and Proposition 1 together show that a client with a smaller energy deficit $q_k(t)$ and a better channel condition $(h_k^t)^2$ (and hence a smaller ρ_k^t) is more likely to be selected to

participate in the current FL round; however, among the selected clients, a client with a smaller ρ_k^t is allocated with less bandwidth. This is because although allocating more bandwidth to client k with a smaller ρ_k^t reduces the energy consumption and deficit of this client, it reduces the bandwidth that can be allocated to clients with larger ρ , which leads to even higher increased energy consumption and deficit of those clients. Moreover, in the optimal solution, the overall effect of energy deficit and consumption, namely $q_k(t-1)E_k^t(b_k^*)$, is still increasing in ρ_k^t .

With the optimality of OCEAN-P, we then prove the performance guarantee of OCEAN.

Theorem 2. *For any $R \in \mathbb{Z}^+$ and $M \in \mathbb{Z}^+$ such that $T = MR$, when comparing OCEAN with the R -round lookahead algorithm, the following statements hold:*

(a) *The energy constraint of every client k is approximately satisfied with a bounded deviation:*

$$\sum_{t=0}^T E_k(a_k^t, b_k^t | h_k^t) \leq H_k + \sum_{m=0}^{M-1} \sqrt{\frac{2(V_m \eta^t K + C_1)}{R}}, \forall k \quad (17)$$

where $C_1 \triangleq K(E^{\max} - H^{\min}/T)^2/2$.

(b) *The federated learning performance satisfies:*

$$\sum_{t=0}^{T-1} U(\mathbf{a}^t) \geq \sum_{m=0}^{M-1} U_m^* - C_2 \sum_{m=0}^{M-1} \frac{1}{V_m} \quad (18)$$

where $C_2 \triangleq C_1 R + \frac{R(R-1)K}{2}(E^{\max})^2$ and U_m^* is the optimal value achieved by the R -round lookahead algorithm in frame m .

Proof. See Appendix D. □

Theorem 2 shows that, given a fixed value of R and M , OCEAN is $O(1/V)$ -optimal with respect to the FL performance against the optimal R -lookahead policy, while the energy consumption is guaranteed to be approximately satisfied with a bounded factor $O(\sqrt{V})$. Thus, OCEAN demonstrates an $[O(1/V), O(\sqrt{V})]$ learning-energy tradeoff. Note that when $R = T$, the T -lookahead benchmark has complete future information of the entire T rounds. Even in this case, the $[O(1/V), O(\sqrt{V})]$ tradeoff still holds.

VI. SIMULATION RESULTS

In this section, we simulate a WFLN to evaluate the performance of OCEAN.

Federated Dataset. To simulate FL, we leverage the TensorFlow Federated (TFF) framework and the MNIST dataset for hand-written digit classification. Each client's local dataset is keyed

by the original writer of the digits. Since each writer has a unique style, this dataset exhibits the kind of non-i.i.d. behavior expected of federated datasets. We use the first 10 clients in the MNIST dataset to conduct our simulation with each client having 100 training data samples. Since the hand-written digit classification is a relatively easy image classification task, we follow TFF's tutorial to construct a simple three-layer neural network with the first layer being input, the second containing 10 neurons and the third performing the softmax operation. This neural network's model size is $L = 3.4 \times 10^5$ bits. FedAvg [1] is used as the learning algorithm.

Wireless Network. To simulate the wireless network, we consider an OFDMA system where the total bandwidth $B = 10$ MHz. Each client's wireless channel gain is modelled as independent free-space fading with average path loss 36dB. The variance of the complex white Gaussian channel noise is set as $N_0 = 10^{-12}$ W. To ensure timely model update, we set the target uploading time in each round to be $\bar{\tau} = 300$ ms. The minimal bandwidth b_{min} is set as 2×10^5 Hz. For each client k , the energy budget is set as $H_k = 0.15$ J. The network runs for $T = 300$ rounds.

A. Benchmarks

We compare the performance of OCEAN with the following three benchmark algorithms.

- **Select-All:** All 10 clients are selected in every learning round. Bandwidth is allocated to minimize the total energy consumption while satisfying the upload deadline requirement.
- **Static Myopic Optimal (SMO):** In every learning round, SMO uses only currently available information independently across rounds (which is equivalent to the 1-Round Lookahead algorithm) to solve

$$\max_{\mathbf{a}^t, \mathbf{b}^t} \sum_k a_k^t \quad (19)$$

$$\text{s.t. } E(a_k^t, b_k^t | h_k^t) \leq H_k/T, \forall k \quad (20)$$

Constraints (6), (7)

This problem is easy to solve: for each client k , first compute the required bandwidth $b_k^\dagger \geq b_{min}$ so that using H_k/T energy can meet the upload time target $\bar{\tau}$; then rank b_k^\dagger in the ascending order and select clients until the total required bandwidth exceeds B . SMO mimics existing approaches (e.g., [3]) that solves bandwidth allocation and client selection independently across learning rounds.

- **Adaptive Myopic Optimal (AMO):** SMO has a clear deficiency which can result in energy under-utilization: when a client is not selected in a round, its energy H_k/T is wasted and will not be used in future rounds. To address this issue, we also consider a modified version of SMO, which recycles previously unused energy budget for future rounds. In particular, the energy budget for client k in round t is modified to $(H_k - \sum_{\tau=0}^{t-1} E_k^\tau)/(T - t)$.

For OCEAN, we let $R = T$ and hence the sequence V_1, \dots, V_M becomes a single scalar V . Moreover, we implement three variants using different temporal importance sequences η_t : Ascending (OCEAN-a); Descending (OCEAN-d); and Uniform (OCEAN-u).

B. Performance Comparison

Figure 5 shows the number of selected clients in every round for different approaches, which is obtained by averaging over 10 runs. As the name suggests, Select-All selects all 10 clients in every round, resulting in the ideal optimal client selection for FL. SMO selects much fewer clients due to the hard energy budget allocation in every round. Many clients do not get to upload their local model updates due to the bad channel state that they are experiencing. AMO starts with selecting few clients due to the same reason as SMO. However, as time goes on, energy budget not used in the previous rounds accumulates. This allows the client to transmit at the desired rate using a higher transmission power in later rounds, especially in those towards the very end, thereby countering the effects of bad channel states. As a (fortunate) by-product, AMO also achieves an ascending pattern of client selection. Our proposed algorithm, OCEAN-a, is able to select many more clients than SMO because it uses energy as needed without imposing a hard per-round energy constraint. Compared to AMO, it is able to fine-tune the temporal pattern of client selection by using different sequences of temporal weights η . As can be seen in Figure 6, OCEAN-a results in an increasing number of selected clients, OCEAN-d results in a decreasing number of selected clients, while OCEAN-u keeps the number of selected clients almost the same across rounds.

Figure 7 shows the actual energy consumption of individual clients by the end of 300 learning rounds for different approaches in a particular run. Because Select-All completely ignores the energy budgets of the clients, it results in a very large energy consumption, far exceeding the energy budgets. On the other hand, SMO does not fully utilize the client's energy budget because in many learning rounds the client is not selected. Both AMO and OCEAN-a incur a total energy consumption close to the given energy budget (i.e. 0.15) for individual clients.




Fig. 5: Temporal Client Selection Patterns of OCEAN and Benchmarks Fig. 6: Temporal Client Selection Patterns of OCEAN Variants




Fig. 7: Per-Client Energy Consumption Comparison

As our ultimate goal is to improve the FL performance, we show the training loss and accuracy for different approaches in Figures 8 and 9. Select-All, as expected, results in the best FL performance, with the smallest training loss, the highest accuracy and the fastest convergence among all approaches. Due to the insufficient selection of clients in the course of learning, SMO’s learning performance is considerably inferior to all other approaches. Thanks to the fortunate by-product of AMO, AMO’s FL performance is comparable to OCEAN-a in this specific setting, which is close to the ideal case Select-All. However, we will show in the next set of experiments that AMO’s “luck” does not extend to other more complex network environments.




Fig. 8: Training Loss of OCEAN-a and Benchmarks




Fig. 9: Accuracy of OCEAN-a and Benchmarks

C. Adaptability to Varying Network Condition

Although the performance of AMO seems comparable to OCEAN-a in the last experiment, it is achieved in a relatively easy network, where the wireless channel is relatively stable. In this set of experiments, we simulate more challenging network environments where the wireless channel can vary considerably due to, e.g., client mobility. In particular, we simulate two scenarios. In Scenario 1, the average path loss gradually increases from 32 dB to 45 dB, mimicking a scenario where clients move away from the server over time. In Scenario 2, the average path loss gradually decreases from 45 dB to 32 dB, mimicking a scenario where clients move towards the server over time.

Scenario 1. Figure 10 shows the number of selected clients over 300 rounds for OCEAN-a and AMO and Figure 11 shows the their FL accuracy. In the early rounds when the wireless channel is good, AMO selects some clients. However, as the channel gain degrades, AMO is not able to adapt to this change as the pre-allocated energy budget (even if the unused budget from the previous rounds is incorporated) cannot support even a single client to finish uploading the local model before the deadline $\bar{\tau}$. Only in the rounds towards the very end does the energy budget become sufficient and hence, some clients again are selected to upload their local model updates. Because of the long idle period in the middle when no clients are selected, the learning performance of AMO is significantly worse than OCEAN.

Scenario 2. Figure 12 shows the number of selected clients over 300 rounds for OCEAN-a and

AMO and Figure 13 shows the their federated learning accuracy. In this scenario, the channel state in the early rounds is bad and hence, hardly any client can be selected to upload its local model update due to insufficient energy budget in AMO. As the channel state improves, AMO starts to select some clients but it becomes too late to do so.

In both scenarios, OCEAN is able to adapt its client selection decision because of its soft per-round energy budget allocation, yet the total consumed energy is still made close to the total energy budget. The per-client total energy consumption of OCEAN-a is shown in Figure 14 for the two considered scenarios.




Fig. 10: Client Selection of Scenario 1




Fig. 11: Accuracy of Scenario 1




Fig. 12: Client Selection of Scenario 2

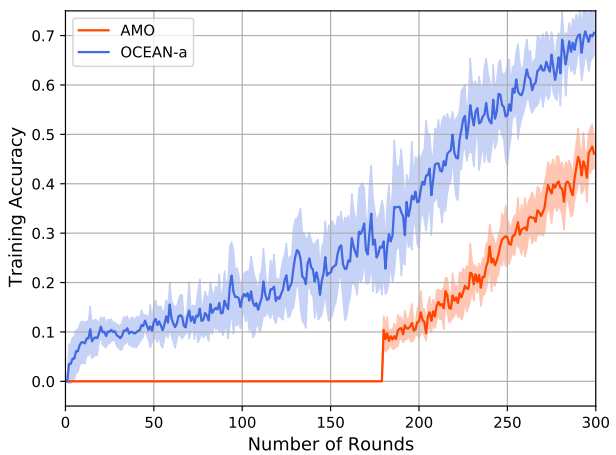


Fig. 13: Accuracy of Scenario 2

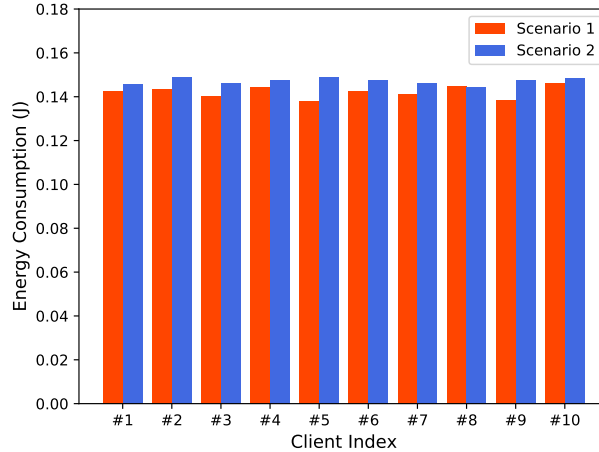


Fig. 14: Energy Consumption of OCEAN-a for the Two Scenarios

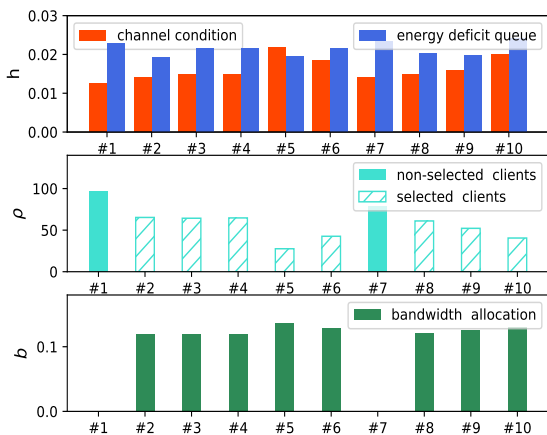


Fig. 15: Client Selection and Bandwidth Allocation Outcomes

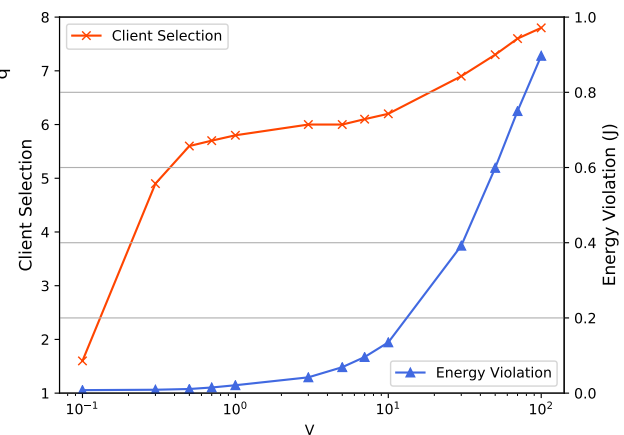


Fig. 16: Tradeoff Between Learning and Energy

D. Features of OCEAN

1) *Client Selection and Bandwidth Allocation Outcomes*: To have a deeper understanding of how OCEAN works, we illustrate, in one specific round, how clients are selected and bandwidth is allocated depending on the clients' channel condition and energy deficit queue in that round. In Figure 15, the upper subplot shows the current channel condition and energy deficit queue for each client. The middle subplot shows the computed selection priority ρ , with shaded bars indicating the selected clients. The bottom subplot shows the bandwidth allocation among the selected clients. As can be seen, a better channel condition and a larger deficit queue result in a

higher priority (i.e., a smaller value of ρ). However, among the selected clients, more bandwidth is allocated to clients of a lower priority (i.e., a larger value of ρ).

2) *Learning - Energy Tradeoff*: Finally, we show the impact of the algorithm parameter V on achieving different learning v.s. energy tradeoff of OCEAN. Figure 16 shows the number of selected clients, the learning accuracy and the per-client energy consumption violation as a function of V . As can be seen, a larger V emphasizes more on the learning performance, resulting in more selected clients and higher accuracy. On the other hand, a smaller V emphasizes more on the energy consumption, resulting in a smaller violation (if any) on the total energy budget.

VII. CONCLUSION

Resource allocation in wireless networks is an old topic, but it also faces constantly changing new challenges as new applications emerge. With FL being the trending new wireless network application, the old mindset of resource allocation for traditional applications such as file downloading or video streaming must be changed. This paper identifies a key property of FL, namely the temporal dependency and varying significance of learning rounds, that may significantly reshape how wireless resources should be allocated for optimized network and learning performance, yet is largely overlooked in the literature. While our formulation and algorithm have shown superior performance in real-world FL experiments, there are several future research directions that may extend the impact of this work. For example, we showed that an ascending client selection pattern is generally desired, but it is still not clear what the optimal pattern is. Moreover, client heterogeneity in terms of the computing power and local data size/distribution can be incorporated into the model to further enhance the understanding of resource allocation in more complex WFLNs.

APPENDIX A

PROOF OF LEMMA 1

To prove the monotonicity and convexity of $f(x)$, we investigate the first and second order derivatives, respectively. The first-order derivative is

$$f'(x) = 2^{\frac{\beta}{x}} - 1 + x \ln 2 \cdot 2^{\frac{\beta}{x}} \left(-\frac{\beta}{x^2}\right) = 2^{\frac{\beta}{x}} \left(1 - \ln 2 \cdot \frac{\beta}{x}\right) - 1 \quad (21)$$

The second-order derivative is

$$f''(x) = \ln 2 \cdot 2^{\frac{\beta}{x}} \left(-\frac{\beta}{x^2}\right) \left(1 - \ln 2 \cdot \frac{\beta}{x}\right) + 2^{\frac{\beta}{x}} \ln 2 \cdot \frac{\beta}{x^2} = (\ln 2)^2 2^{\frac{\beta}{x}} \frac{\beta^2}{x^3} \quad (22)$$

Therefore, for $x \in (0, \infty)$, $f''(x) > 0$ and hence $f'(x)$ is an increasing function. We also know that $\lim_{x \rightarrow \infty} f'(x) = 0$ and thus, $f'(x) < 0$ for $x \in (0, \infty)$. This proves that $f(x)$ is decreasing on $(0, \infty)$. Similarly, for $x \in (-\infty, 0)$, $f''(x) < 0$ and hence $f'(x)$ is a decreasing function. Because $\lim_{x \rightarrow -\infty} f'(x) = 0$, $f'(x) > 0$ for $x \in (-\infty, 0)$. This proves that $f(x)$ is increasing on $(-\infty, 0)$. Moreover, $f(x)$ is convex on $(0, \infty)$ as $f''(x) > 0$; it is concave on $(-\infty, 0)$ as $f''(x) < 0$.

APPENDIX B

PROOF OF THEOREM 1

The key is to prove that the optimal solution must have the following thresholding structure: there exists k^* so that $a_k^* = 1, \forall k \leq k^*$ and $a_k^* = 0, \forall k > k^*$. To prove this, suppose that in the optimal solution $(\mathbf{a}^*, \mathbf{b}^*)$, there exist $k_1 < k_2$ so that $a_{k_1}^* = 0$ and $a_{k_2}^* = 1$. Let us consider a different solution $(\tilde{\mathbf{a}}, \tilde{\mathbf{b}})$ which is obtained by swapping the decisions for k_1 and k_2 in $(\mathbf{a}^*, \mathbf{b}^*)$. Specifically,

$$\tilde{a}_{k_1} = a_{k_2}^* = 1, \tilde{b}_{k_1} = b_{k_2}^* \quad (23)$$

$$\tilde{a}_{k_2} = a_{k_1}^* = 0, \tilde{b}_{k_2} = b_{k_1}^* \quad (24)$$

Since the decisions for other clients remain the same, the difference in the objective function value is

$$\sum_{k=1}^K (\eta^t - \rho_k N_0 \tilde{\tau} B f(\tilde{b}_k)) \tilde{a}_k - \sum_{k=1}^K (\eta - \rho_k N_0 \tilde{\tau} B f(b_k^*)) a_k^* \quad (25)$$

$$= (\eta - \rho_{k_1} N_0 \tilde{\tau} B f(\tilde{b}_{k_1})) - (\eta - \rho_{k_2} N_0 \tilde{\tau} B f(b_{k_2}^*)) = (\rho_{k_2} - \rho_{k_1}) N_0 \tilde{\tau} B f(b_{k_2}^*) > 0 \quad (26)$$

where $f(b) = b(2^{\frac{L}{\tilde{\tau} B b}} - 1)$. This is a contradiction to the optimality of $(\mathbf{a}^*, \mathbf{b}^*)$. Therefore, the optimal solution must have the aforementioned thresholding structure.

Next, we prove that the termination condition is correct. Let i^* be the first client with $\eta - \rho_i N_0 \tilde{\tau} B f(b_{i^*}^*[i^*]) < 0$ where we use $b_k^*[i]$ to denote the optimal bandwidth allocation for the selection set $S = \{1, \dots, i\}$. Clearly, when only clients $k \in \{1, \dots, i^* - 1\}$ are selected, we obtain

a higher utility because

$$\sum_{k=1}^{i^*-1} (\eta - \rho_k N_0 \tilde{\tau} B f(b_k^*[i^* - 1])) - \sum_{k=1}^{i^*} (\eta - \rho_k N_0 \tilde{\tau} B f(b_k^*[i^*])) \quad (27)$$

$$= \sum_{k=1}^{i^*-1} (\eta - \rho_k N_0 \tilde{\tau} B f(b_k^*[i^* - 1])) - \sum_{k=1}^{i^*-1} (\eta - \rho_k N_0 \tilde{\tau} B f(b_k^*[i^*])) \quad (28)$$

$$- (\eta - \rho_{i^*} N_0 \tilde{\tau} B f(b_{i^*}^*[i^*])) > 0 \quad (29)$$

Therefore, client i^* must not be in the optimal selection set. Because of the thresholding structure, we know that clients $k = i^* + 1, \dots, K$ also must not be in the optimal selection set. This proves the correctness of the termination condition.

APPENDIX C

PROOF OF PROPOSITION 1

It suffices to consider the following optimization problem with two clients

$$\min_{b_1, b_2} \rho_1 f(b_1) + \rho_2 f(b_2) \quad (30)$$

$$\text{s.t. } b_1 + b_2 = \delta, \quad b_1, b_2 \geq b_{\min} \quad (31)$$

where δ is any constant in $(2b_{\min}, 1]$. Let $\rho_1 < \rho_2$. Suppose the optimal bandwidth allocation satisfies $b_1^* > b_2^*$, then by Lemma 1, we know $f(b_1^*) < f(b_2^*)$. Let us construct a different bandwidth allocation solution $\tilde{\mathbf{b}}$ where $\tilde{b}_1 = b_2^*$ and $\tilde{b}_2 = b_1^*$. In other words, the bandwidth allocation decisions are swapped. This solution also satisfies all constraints. We compare the respective objective values and have

$$\rho_1 f(b_1^*) + \rho_2 f(b_2^*) - (\rho_1 f(\tilde{b}_1) + \rho_2 f(\tilde{b}_2)) = (\rho_1 - \rho_2)(f(b_1^*) - f(b_2^*)) > 0 \quad (32)$$

This contradicts the optimality of \mathbf{b}^* . Therefore, we must have $b_1^* \leq b_2^*$.

To prove $\rho_1 f(b_1^*) \leq \rho_2 f(b_2^*)$, let $b_2 = \delta - b_1$ and ignore the constraint that $b_1, b_2 \geq b_{\min}$ for now. The first-order condition requires

$$\rho_1 df(b_1)/db_1 + \rho_2 df(b_2)/db_2 \cdot db_2/db_1 = 0 \quad (33)$$

This leads to

$$\rho_1 f'(b_1^*) = \rho_2 f'(b_2^*) \quad (34)$$

Because $f'(x) < 0$, we can instead prove $f(b_1^*)/f'(b_1^*) \geq f(b_2^*)/f'(b_2^*)$. Let us define $g_1(x) \triangleq f(x)/f'(x)$. Since we have proven $b_1^* \leq b_2^*$ in the above, we only need to prove that $g_1(x)$ is a non-increasing function in $x > 0$. To this end, consider the first order derivative of $g_1(x)$,

$$g_1'(x) = \frac{(f'(x))^2 - f(x)f''(x)}{(f'(x))^2} \quad (35)$$

Let $g_2(x) \triangleq (f'(x))^2 - f(x)f''(x)$. We have to prove $g_2(x) \leq 0$ for $x > 0$.

$$g_2(x) = \left(2^{\frac{\beta}{x}} \left(1 - \ln 2 \cdot \frac{\beta}{x}\right) - 1\right)^2 - x(2^{\frac{\beta}{x}} - 1) \cdot (\ln 2)^2 2^{\frac{\beta}{x}} \frac{\beta^2}{x^3} \quad (36)$$

$$= (2^{\frac{\beta}{x}} - 1)^2 - 2(2^{\frac{\beta}{x}} - 1)2^{\frac{\beta}{x}} \ln 2 \cdot \frac{\beta}{x} + (\ln 2)^2 2^{\frac{\beta}{x}} \frac{\beta^2}{x^2} \quad (37)$$

To simplify notations, we use a change of variable by letting $y = \beta/x$. Then proving $g_2(y) \leq 0$ for $y > 0$ is equivalent to proving $g_2(x) \leq 0$ for $x > 0$. We rewrite $g_2(y)$ below:

$$g_2(y) \triangleq (2^y - 1)^2 - 2(2^y - 1)2^y \ln 2 \cdot y + (\ln 2)^2 2^y y^2 \quad (38)$$

Clearly, $g_2(0) = 0$. In order to prove $g_2(y) \leq 0$, we prove $g_2(y)$ is decreasing in $y \geq 0$.

$$g_2'(y) = -(\ln 2)^2 y 2^y (4(2^y - 1) - \ln 2 \cdot y) \quad (39)$$

It is easy to verify that $g_3(y) \triangleq 4(2^y - 1) - \ln 2 \cdot y$ is increasing in $y > 0$ and $g_3(0) = 0$, which means $g_3(y) \geq 0$ for $y \geq 0$. Hence, $g_2'(y) < 0$ for $y > 0$. This concludes the proof for $\rho_1 f(b_1^*) \leq \rho_2 f(b_2^*)$ by ignoring the constraint $b_1, b_2 \geq b_{min}$.

When the constraint $b_1, b_2 \geq b_{min}$ is considered, there are two cases. In the first case, $b_{min} \leq b_1^* \leq b_2^*$. In this case, the constraint is automatically satisfied and hence, our above conclusion holds. In the second case, $b_1^* \leq b_{min} \leq b_2^*$. In this case, the optimal allocation is modified to $\tilde{b}_1^* = b_{min} \geq b_1^*$ and $\tilde{b}_2^* = 1 - b_{min} \leq b_2^*$. Since $f(x)$ is a decreasing function, $\rho_1 f(\tilde{b}_1^*) \leq \rho_1 f(b_1^*) \leq \rho_2 f(b_2^*) \leq \rho_2 f(\tilde{b}_2^*)$. This completes the proof.

APPENDIX D

PROOF OF THEOREM 2

We define the quadratic Lyapunov function $L(\mathbf{q}(t)) \triangleq \frac{1}{2} \sum_{k=1}^K q_k^2(t)$. Let $\Delta_1(t)$ be the 1-round Lyapunov drift yielded by some control decisions over one round: $\Delta_1(t) \triangleq L(\mathbf{q}(t+1)) - L(\mathbf{q}(t))$. Similarly, let $\Delta_R(t)$ be the R -round Lyapunov drift: $\Delta_R(t) \triangleq L(\mathbf{q}(t+R)) - L(\mathbf{q}(t))$. Based on the queue dynamics, we have

$$\frac{1}{2} \sum_k^K q_k^2(t+1) \leq \frac{1}{2} \sum_{k=1}^K [E_k(a_k^t, b_k^t | h_k^t) - H_k/T + q_k(t)]^2 \quad (40)$$

Then, it can be easily show that

$$\Delta_1(t) \leq C_1 + \sum_{k=1}^K q_k(t) \cdot [E_k(a_k^t, b_k^t | h_k^t) - H_k/T] \quad (41)$$

where C_1 is a constant satisfying $C_1 \geq \frac{1}{2} \sum_{k=1}^K (E^{\max} - H^{\min}/T)^2, \forall t$, which is finite due to the boundedness of the channel condition h_k^t and the minimum bandwidth allocation requirement b_{min} . Next, it is straightforward that $\forall m$ and $\forall t = mR, \dots, (m+1)R - 1$

$$V_m \cdot U(\mathbf{a}^t) - \Delta_1(t) \geq V_m \cdot U(\mathbf{a}^t) - \sum_{k=1}^K q_k(t) \cdot [E_k(a_k^t, b_k^t | h_k^t) - H_k] - C_1 \quad (42)$$

As we can see, by solving **P3**, OCEAN-P actually maximizes a lower bound of $V_m \cdot U(\mathbf{a}^t) - \Delta_1(t)$.

Let $\hat{\mathbf{a}}^0, \hat{\mathbf{b}}^0, \dots, \hat{\mathbf{a}}^{T-1}, \hat{\mathbf{b}}^{T-1}$ be the sequence of decisions derived by the online algorithm.

(a) Consider a specific sequence of decisions where $\tilde{a}_k^t = 0, \forall t, k$. Clearly, in this case, $U(\tilde{\mathbf{a}}^t) = 0$ and $E_k(\tilde{a}_k^t, \tilde{b}_k^t | h_k^t) - H_k/K = -H_k/K$. Because $\hat{\mathbf{a}}^t, \hat{\mathbf{b}}^t$ maximizes the right-hand side of (42), we have

$$V_m \cdot U(\hat{\mathbf{a}}^t) - \Delta_1(t) \geq V_m \cdot 0 + \sum_{k=1}^K \tilde{q}_k(t) H_k - C_1 \geq -C_1 \quad (43)$$

Therefore

$$\Delta_1(t) \leq V_m \cdot U(\hat{\mathbf{a}}^t) + C_1 \leq V_m \eta^t K + C_1 \quad (44)$$

As enforced by the online algorithm,

$$\Delta_R(mR) = \frac{1}{2} \sum_{k=1}^K (q_k^2(mR+R) - q_k^2(mR)) = \frac{1}{2} \sum_{k=1}^K q_k^2(mR+R) \quad (45)$$

is the R -round drift calculated after the m -th rest but before the $(m+1)$ -th reset of the energy deficit queue (so $q_k(mR) = 0$). Thus, before the $(m+1)$ -th reset of the energy deficit queue, we have

$$\sum_{k=1}^K q_k^2(mR+R) = 2\Delta_R(mR) = 2 \sum_{t=mR}^{mR+R-1} \Delta_1(t) \leq 2R(V_m \eta^t K + C_1) \quad (46)$$

Therefore, $\forall k$,

$$q_k(mR+R) \leq \sqrt{2R(V_m \eta^t K + C_1)} \quad (47)$$

On the other hand, according to the queue dynamics (10), we have

$$q_k(t+1) - q_k(t) \geq E_k(a_k^t, b_k^t | h_k^t) - H_k/T \quad (48)$$

Summing both sides over the rounds in the m -th frame, namely $t = mR, \dots, (m+1)R - 1$, and dividing by R , we have

$$\frac{1}{R} \sum_{t=mR}^{(m+1)R-1} (E_k(a_k^t, b_k^t | h_k^t) - H_k/T) \leq \frac{q_k((m+1)R) - q_k(mR)}{R} = \frac{q_k((m+1)R)}{R} \quad (49)$$

Plugging (47) into (49), we have

$$\frac{1}{R} \sum_{t=mR}^{(m+1)R-1} (E_k(\hat{a}_k^t, \hat{b}_k^t | h_k^t) - H_k/T) \leq \sqrt{\frac{2(V_m \eta^t K + C_1)}{R}} \quad (50)$$

Considering all M frames, we obtain

$$\sum_{t=0}^T E_k(\hat{a}_k^t, \hat{b}_k^t | h_k^t) \leq H_k + \sum_{m=0}^{M-1} \sqrt{\frac{2(V_m \eta^t K + C_1)}{R}}, \forall k \quad (51)$$

(b) Consider the R -round weighted learning utility minus drift:

$$V_m \sum_{t=mR}^{mR+R-1} U(\mathbf{a}^t) - \Delta_R(mR) \quad (52)$$

$$\geq V_m \sum_{t=mR}^{mR+R-1} U(\mathbf{a}^t) - C_1 R - \sum_{t=mR}^{mR+R-1} \sum_{k=1}^K q_k(t) \cdot [E_k(a_k^t, b_k^t | h_k^t) - H_k/T] \quad (53)$$

Because $\hat{\mathbf{a}}^0, \hat{\mathbf{b}}^0, \dots, \hat{\mathbf{a}}^{T-1}, \hat{\mathbf{b}}^{T-1}$ explicitly maximizes the right-hand side of the above equation, the following must also hold

$$V_m \sum_{t=mR}^{mR+R-1} U(\hat{\mathbf{a}}^t) - \Delta_R(mR) \quad (54)$$

$$\geq V_m \sum_{t=mR}^{mR+R-1} U(\mathbf{a}^{*,t}) - C_1 R - \sum_{t=mR}^{mR+R-1} \sum_{k=1}^K q_k(t) \cdot [E_k(a_k^{*,t}, b_k^{*,t} | h_k^t) - H_k/T] \quad (55)$$

$$\geq V_m \sum_{t=mR}^{mR+R-1} U(\mathbf{a}^{*,t}) - C_1 R - \sum_{t=mR}^{mR+R-1} \sum_{k=1}^K (t - mR) E^{\max} \cdot [E_k(a_k^{*,t}, b_k^{*,t} | h_k^t) - H_k/T] \quad (56)$$

$$- \sum_{k=1}^K q_k(mR) \sum_{t=mR}^{mR+R-1} [E_k(a_k^{*,t}, b_k^{*,t} | h_k^t) - H_k/T] \quad (57)$$

$$\geq V_m \sum_{t=mR}^{mR+R-1} U(\mathbf{a}^{*,t}) - \sum_{k=1}^K q_k(mR) \sum_{t=mR}^{mR+R-1} [E_k(a_k^{*,t}, b_k^{*,t} | h_k^t) - H_k/T] - \left(C_1 R + \frac{R(R-1)K}{2} (E^{\max})^2 \right) \quad (58)$$

$$\geq V_m \sum_{t=mR}^{mR+R-1} U(\mathbf{a}^{*,t}) - \left(C_1 R + \frac{R(R-1)K}{2} (E^{\max})^2 \right) \quad (59)$$

where in $\hat{\mathbf{a}}^{*,0}, \hat{\mathbf{b}}^{*,0}, \dots, \hat{\mathbf{a}}^{*,T-1}, \hat{\mathbf{b}}^{*,T-1}$ is the optimal decision that solves the R -round lookahead problems. Notice that $q_k(t)$ in the above equation is still derived by OCEAN-P. The first inequality holds because OCEAN-P maximizes the lower bound. The last inequality holds because $q_k(mR)$ is reset to zero as enforced by OCEAN-P.

Noticing $\Delta_R(mR) \geq 0$ and dividing both sides by V_m , we have

$$\sum_{t=mR}^{mR+R-1} U(\hat{\mathbf{a}}^t) \geq \sum_{t=mR}^{mR+R-1} U(\mathbf{a}^{*,t}) - \frac{1}{V_m} \left(C_1 R + \frac{R(R-1)K}{2} (E^{\max})^2 \right) \quad (60)$$

By summing over $m = 0, \dots, M-1$, we have

$$\sum_{t=0}^{T-1} U(\hat{\mathbf{a}}^t) \geq \sum_{m=0}^{M-1} U_m^* - C_2 \sum_{m=0}^{M-1} \frac{1}{V_m} \quad (61)$$

where $C_2 \triangleq C_1 R + \frac{R(R-1)K}{2} (E^{\max})^2$.

REFERENCES

- [1] J. Konecny, H. B. McMahan, F. X. Yu, P. Richtarik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," *arXiv preprint arXiv:1610.05492*, 2016.
- [2] K. Yang, T. Jiang, Y. Shi, and Z. Ding, "Federated learning via over-the-air computation," *IEEE Transactions on Wireless Communications*, vol. 19, no. 3, pp. 2022–2035, 2020.
- [3] Q. Zeng, Y. Du, K. K. Leung, and K. Huang, "Energy-efficient radio resource allocation for federated edge learning," *arXiv preprint arXiv:1907.06040*, 2019.
- [4] M. J. Neely, "Stochastic network optimization with application to communication and queueing systems," *Synthesis Lectures on Communication Networks*, vol. 3, no. 1, pp. 1–211, 2010.
- [5] J. Konecny, H. B. McMahan, D. Ramage, and P. Richtarik, "Federated optimization: Distributed machine learning for on-device intelligence," *arXiv preprint arXiv:1610.02527*, 2016.
- [6] S. P. Karimireddy, S. Kale, M. Mohri, S. J. Reddi, S. U. Stich, and A. T. Suresh, "Scaffold: Stochastic controlled averaging for on-device federated learning," *arXiv preprint arXiv:1910.06378*, 2019.
- [7] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," *arXiv preprint arXiv:1812.06127*, 2018.
- [8] F. Haddadpour and M. Mahdavi, "On the convergence of local descent methods in federated learning," *arXiv preprint arXiv:1910.14425*, 2019.
- [9] Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, and V. Chandra, "Federated learning with non-iid data," *arXiv preprint arXiv:1806.00582*, 2018.
- [10] V. Smith, C.-K. Chiang, M. Sanjabi, and A. S. Talwalkar, "Federated multi-task learning," in *Advances in Neural Information Processing Systems*, 2017, pp. 4424–4434.
- [11] L. Corinzia and J. M. Buhmann, "Variational federated multi-task learning," *arXiv preprint arXiv:1906.06268*, 2019.
- [12] A. Bhowmick, J. Duchi, J. Freudiger, G. Kapoor, and R. Rogers, "Protection against reconstruction and its applications in private federated learning," *arXiv preprint arXiv:1812.00984*, 2018.
- [13] R. C. Geyer, T. Klein, and M. Nabi, "Differentially private federated learning: A client level perspective," *arXiv preprint arXiv:1712.07557*, 2017.

- [14] S. Truex, N. Baracaldo, A. Anwar, T. Steinke, H. Ludwig, R. Zhang, and Y. Zhou, “A hybrid approach to privacy-preserving federated learning,” in *Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security*, 2019, pp. 1–11.
- [15] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth, “Practical secure aggregation for privacy-preserving machine learning,” in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 2017, pp. 1175–1191.
- [16] M. Nasr, R. Shokri, and A. Houmansadr, “Comprehensive privacy analysis of deep learning: Stand-alone and federated learning under passive and active white-box inference attacks,” *arXiv preprint arXiv:1812.00910*, 2018.
- [17] T. Chen, G. Giannakis, T. Sun, and W. Yin, “Lag: Lazily aggregated gradient for communication-efficient distributed learning,” in *Advances in Neural Information Processing Systems*, 2018, pp. 5050–5060.
- [18] Y. Lin, S. Han, H. Mao, Y. Wang, and W. J. Dally, “Deep gradient compression: Reducing the communication bandwidth for distributed training,” *arXiv preprint arXiv:1712.01887*, 2017.
- [19] A. F. Aji and K. Heafield, “Sparse communication for distributed gradient descent,” *arXiv preprint arXiv:1704.05021*, 2017.
- [20] L. Liu, J. Zhang, S. Song, and K. B. Letaief, “Edge-assisted hierarchical federated learning with non-iid data,” *arXiv preprint arXiv:1905.06641*, 2019.
- [21] M. M. Amiri and D. Gunduz, “Federated learning over wireless fading channels,” *IEEE Transactions on Wireless Communications*, 2020.
- [22] S. Wang, T. Tuor, T. Salonidis, K. K. Leung, C. Makaya, T. He, and K. Chan, “Adaptive federated learning in resource constrained edge computing systems,” *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 6, pp. 1205–1221, 2019.
- [23] N. H. Tran, W. Bao, A. Zomaya, N. M. NH, and C. S. Hong, “Federated learning over wireless networks: Optimization model design and analysis,” in *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*. IEEE, 2019, pp. 1387–1395.
- [24] X. Mo and J. Xu, “Energy-efficient federated edge learning with joint communication and computation design,” *arXiv preprint arXiv:2003.00199*, 2020.
- [25] Y. Zhan, P. Li, and S. Guo, “Experience-driven computational resource allocation of federated learning by deep reinforcement learning,” in *Proc. of IPDPS*, 2020.
- [26] S. U. Stich, “Local sgd converges fast and communicates little,” *arXiv preprint arXiv:1805.09767*, 2018.
- [27] H. H. Yang, Z. Liu, T. Q. Quek, and H. V. Poor, “Scheduling policies for federated learning in wireless networks,” *IEEE Transactions on Communications*, 2019.
- [28] 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,” *arXiv preprint arXiv:1909.07972*, 2019.
- [29] W. Shi, S. Zhou, and Z. Niu, “Device scheduling with fast convergence for wireless federated learning,” *arXiv preprint arXiv:1911.00856*, 2019.
- [30] T. Nishio and R. Yonetani, “Client selection for federated learning with heterogeneous resources in mobile edge,” in *ICC 2019-2019 IEEE International Conference on Communications (ICC)*. IEEE, 2019, pp. 1–7.
- [31] Z. Yang, M. Chen, W. Saad, C. S. Hong, and M. Shikh-Bahaei, “Energy efficient federated learning over wireless communication networks,” *arXiv preprint arXiv:1911.02417*, 2019.
- [32] M. Chen, H. V. Poor, W. Saad, and S. Cui, “Convergence time optimization for federated learning over wireless networks,” *arXiv preprint arXiv:2001.07845*, 2020.
- [33] M. Reneer, “Tensorflow federated,” <https://www.tensorflow.org/federated>, 2020.
- [34] S. Boyd, S. P. Boyd, and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.

- [35] M. Grant and S. Boyd, “Cvx: Matlab software for disciplined convex programming, version 2.1,” 2014.
- [36] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright *et al.*, “Scipy 1.0: fundamental algorithms for scientific computing in python,” *Nature methods*, pp. 1–12, 2020.