Lyapunov-based Optimization of Edge Resources for Energy-Efficient Adaptive Federated Learning

This paper was downloaded from TechRxiv (<https://www.techrxiv.org>).

LICENSE

CC BY 4.0

SUBMISSION DATE / POSTED DATE

04-10-2021 / 12-10-2021

CITATION

Battiloro, Claudio; Di Lorenzo, Paolo; Merluzzi, Mattia; Barbarossa, Sergio (2021): Lyapunov-based Optimization of Edge Resources for Energy-Efficient Adaptive Federated Learning. TechRxiv. Preprint. <https://doi.org/10.36227/techrxiv.16732651.v1>

DOI

[10.36227/techrxiv.16732651.v1](https://doi.org/10.36227/techrxiv.16732651.v1)

Lyapunov-based Optimization of Edge Resources for Energy-Efficient Adaptive Federated Learning

Claudio Battiloro, Paolo Di Lorenzo, Mattia Merluzzi, and Sergio Barbarossa

Abstract

The aim of this paper is to propose a novel dynamic resource allocation strategy for energy-efficient adaptive federated learning at the wireless network edge, with latency and learning performance guarantees. We consider a set of devices collecting local data and uploading processed information to an edge server, which runs stochastic gradient-based algorithms to perform continuous learning and adaptation. Hinging on Lyapunov stochastic optimization tools, we dynamically optimize radio parameters (e.g., set of transmitting devices, transmit powers, bits, and rates) and computation resources (e.g., CPU cycles at devices and at server) in order to strike the best trade-off between power, latency, and performance of the federated learning task. The framework admits both a model-based implementation, where the learning performance metrics are available in closed-form, and a data-driven approach, which works with online estimates of the learning performance of interest. The method is then customized to the case of federated least mean squares (LMS) estimation, and federated training of deep convolutional neural networks. Numerical results illustrate the effectiveness of our strategy to perform energy-efficient, low-latency, adaptive federated learning at the wireless network edge.

Index Terms

Federated learning, Lyapunov stochastic optimization, edge machine learning, resource allocation.

I. INTRODUCTION

In the last few years, with the advent of 5G (and beyond) systems, communication networks are evolving from a pure communication framework to service enablers in several different sectors (including *verticals*), such as Industry 4.0, Internet of Things (IoT), autonomous driving, remote surgery, etc. [1]–[5]. As key enablers of this vision, machine learning (ML) and artificial

Battiloro, Di Lorenzo, and Barbarossa are with the Department of Information Engineering, Electronics, and Telecommunications, Sapienza University of Rome, Via Eudossiana 18, 00184, Rome, Italy. Merluzzi is with Univ. Grenoble Alpes, CEA, Leti, F-38000 Grenoble, France. E-mail: {claudio.battiloro, paolo.dilorenzo, sergio.barbarossa}@uniroma1.it, mattia.merluzzi@cea.fr. This work was partly funded by the H2020 EU-Taiwan Project 5G-CONNI nr. AMD-861459-3, and partly funded by European Commission through the H2020 project Hexa-X (Grant Agreement no. 101015956) for the work of Merluzzi.

intelligence will be largely exploited in future wireless communication networks, in order to build an effective complex system able to learn and dynamically adapt to the evolving network landscape [6]. Indeed, the advent of a new breed of intelligent devices and high-stake applications foreseen in beyond 5G have sparked a huge interest in distributed, low-latency and reliable ML, calling for a novel system design coined *edge machine learning*, in which: (i) training data is unevenly distributed over a large number of edge devices (including phones, cameras, vehicles, and drones); (ii) every edge device has access to a tiny fraction of the data and training is carried out collectively and distributively; (iii) the inference process is performed on the edge devices, requiring not only high learning accuracy and reliability, but also a very short response time necessary for autonomous decision making in highly dynamic environments. However, differently from cloud-based ML that has virtually infinite computing resources, edge ML is a nascent research field whose system design is entangled with communication and on-device resource constraints (e.g., energy and computing power). Moreover, the process of decentralized training involves a large number of devices that are interconnected over wireless links, hindering learning and adaptation due to communications under poor channel conditions. As such, enabling edge ML introduces novel research problems in terms of jointly optimizing inference, training, communication, computation, and control under end-to-end latency, reliability, and learning performance requirements [7]–[10].

Related works. Training ML models at the edge mainly relies on *federated learning* (FL) [11]–[19]. These learning architectures perform (variants of) parallel stochastic gradient descent (SGD) across multiple edge devices, whose intermediate results are aggregated by an Edge Server (ES). FL has several benefits of data privacy, and is empowered by a large amount of device participants with modern powerful processors and low-delay mobile-edge networks. The work in [11] provides a comprehensive survey on FL algorithms and introduced various challenges, problems, and solutions for enhancing FL effectiveness. In [12], the authors develop two update methods to reduce the uplink communication costs for FL. The work in [15] presents a practical update method for a deep FL algorithm and conduces an extensive empirical evaluation for different FL models. The authors in [18] study FL and the problem of joint power and resource allocation for ultra-reliable low latency communication in vehicular networks. The work in [19] develops a new approach to minimize the computing and transmission delay for FL algorithms.

Other works on FL explicitly focus on the optimization of radio resource allocation [20]–[33]. In [20], the authors propose a control algorithm that determines the best trade-off between local

update and global parameter aggregation to minimize the loss function under a given resource budget. In [21], the authors propose energy-efficient strategies for bandwidth allocation and scheduling to enable latency-constrained FL. The work in [22] proposes a joint learning and wireless resource allocation framework to optimize the FL performance. In [23], the authors characterize how the computation and communication latencies of edge devices affect the trade-offs between energy consumption, learning time, and accuracy of the FL task. The work in [25] studies the relationship between batch size and learning rate to alleviate the negative impact of synchronization barrier through adaptive batch size during model training in the FL paradigm. In [26], the authors model the interaction between a global server and the participating devices for federated learning via a Stackelberg game to motivate the participation of the devices in the federated learning process. In [27], the authors provide an optimization problem whose goal is to minimize the total energy consumption of the system under a latency constraint; to solve the problem, an iterative algorithm is proposed where, at every step, closed-form solutions for time allocation, bandwidth allocation, power control, computation frequency, and learning accuracy are derived. The work in [28] propose a federated deep-reinforcement-learning-based cooperative edge caching framework which enables base stations (BSs) to cooperatively learn a shared predictive model by considering the first-round training parameters of the BSs as the initial input of the local training, and then uploads near-optimal local parameters to the BSs to participate in the next round of global training. In [29], the authors propose adapting federated averaging to use a distributed form of Adam optimization along with a compression technique. The work in [30] proposes a FL approach with adaptive and distributed parameter pruning, which adapts the model size during FL to reduce both communication and computation overhead and minimize the overall training time, while maintaining a similar accuracy as the original model. In [31], the authors consider two transmission protocols for edge devices to upload model parameters to edge server, based on non orthogonal multiple access and time division multiple access, respectively. Under both protocols, they minimize the total energy consumption at all edge devices over a particular finite training duration subject to a given training accuracy, by jointly optimizing the transmission power and rates at the edge devices for uploading model parameters and their central processing unit frequencies for local update. In [32], the authors propose a joint device scheduling and resource allocation policy to maximize the model accuracy within a given total training time budget for latency constrained wireless FL. The work in [34] analyzes how to design dynamic FL in mobile edge networks that optimally chooses the number

of selected clients and the number of local iterations in each training round to minimize the total cost while ensuring convergence. Finally, reference [33] proposes a dynamic user selection scheme to minimize the FL convergence time.

Contributions. The goal of this paper is to introduce a novel dynamic optimization framework for *adaptive federated learning*, which jointly encompasses communication, computation, and learning aspects of the problem. Differently from previous works that mainly focused on a static learning task (where FL is carried out up to convergence and then the learning process stops), here we consider adaptive FL strategies, with the aim of endowing wireless networks with continuous learning, adaptation, and tracking capabilities [35]. Hinging on Lyapunov stochastic optimization [36], we develop a dynamic resource allocation strategy that works at the same time-scale of the gradient-based algorithm, while optimizing on the fly radio parameters (e.g., set of transmitting devices, transmit powers, bits and rates) and computation resources (e.g., CPU cycles at devices and at edge server) in order to strike the best trade-off between energy, latency, and performance (e.g., convergence rate, accuracy or mean-squared error) of the adaptive FL task. Particular emphasis is devoted to the definition and the online control of proper performance metrics in both a model based scenario, where the metrics are known in closed-form, and in a data-driven case, where performance must be inferred online from data. In both cases, the method is able to adaptively minimize the average power needed for the FL task, while ensuring guaranteed latency and learning performance. Finally, the proposed strategy is customized to adaptive federated Least Mean Squares (LMS) estimation and deep convolutional neural network training. Part of this work was presented in the preliminary conference paper [37], which is here largely extended in terms of theory, models, and numerical results. Due to the upcoming convergence of communication and learning in future beyond 5G networks, we claim it is fundamental to merge these aspects with a mathematically formal analysis of both communication and computation power consumption and delays, as well as learning performance metrics in terms of accuracy of the learning task, rate of convergence and adaptation. This mathematical analysis, also exploited in the proposed online solution, represents the main contribution of our work with respect to the current state of the art previously presented.

Notation. Scalar, column vector, and matrix variables are respectively indicated by plain letters a (A), bold lowercase letters $\mathbf{a}$, and bold uppercase letters $\mathbf{A}$. $\mathbb{I}(\cdot)$ denotes the indicator function; a_{ij} is the (i, j) -th element of $\mathbf{A}$, $\mathbf{I}$ is the identity matrix, and $\mathbf{1}_N$ ($\mathbf{0}_N$) is the $N \times 1$ vector of all ones (zeros). $\text{diag}\{\mathbf{a}\}$ denotes a diagonal matrix having vector $\mathbf{a}$ on its main diagonal. $\mathbb{E}\{\cdot\}$

denotes the expectation operator, $\text{Tr}\{\cdot\}$ denotes the matrix trace operator, and $\lambda_{\min}\{\cdot\}$ represents the minimum eigenvalue. Other specific notation is defined along the paper.

In this section, we present the mathematical model used to describe the FL algorithm and its performance, the overall latency, and the system power consumption. Let us consider a scenario with N edge devices and an AP equipped with an edge server, as illustrated in Fig. 1. The devices are cooperatively performing a training task aimed at learning a weight vector $\mathbf{w} \in \mathbf{R}^m$. To this aim, at each time t , the devices collect labelled data (i.e., input/output pairs) given by $(\mathbf{x}_{i,t}, y_{i,t}) \in \mathbf{R}^d \times \mathbf{R}$, for all $i = 1, \dots, N$, and $t \geq 0$. Then, assuming that device i has a local loss function $J_i(\mathbf{w}; \mathbf{x}_{i,t}, y_{i,t})$, whose structure depends on the specific learning task, the goal of FL can be mathematically cast as:

where the expectation is carried out over the data distribution. Now, at each time t , letting $\mathbf{w}_t$ be the instantaneous guess for the weight vector $\mathbf{w}$, we proceed by optimizing problem (1) using an adaptive stochastic gradient descent procedure [35]. In particular, let us denote

for all i, t , to shorten the notation. Then, the adopted SGD recursion reads as:

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \mu \sum_{i \in S_t} \mathbf{g}_{i,t}(\mathbf{w}_t), \quad (2)$$

where $t \geq 0$, $\mu > 0$ is a (sufficiently small) step-size parameter; and $\mathcal{S}_t$ is the set of nodes that participate to the optimization at time t . In the considered FL scenario, the algorithm in (2) can be implemented in two ways. The straightforward one requires that, at each time t , the edge devices belonging to $\mathcal{S}_t$ (to be determined for all t) compute in parallel the gradients of the local cost functions (i.e., $\mathbf{g}_{i,t}(\mathbf{w}_t)$) and upload them to the AP. Then, the edge server aggregates the local information to compute the new estimate $\mathbf{w}_{t+1}$, which is finally fed back to the devices. An example of the data exchange required for this implementation is illustrated in Fig. 1. The second implementation of algorithm (2) requires instead that, at each time t , multiple edge devices compute one step of gradient-based algorithms on the current model using local data, and then the server takes a weighted average of the resulting models. In particular, device i evaluates a local estimate $\psi_{i,t}$ given by:

$$\psi_{i,t} = \mathbf{w}_t - \mu |\mathcal{S}_t| \mathbf{g}_{i,t}(\mathbf{w}_t), \quad (3)$$

for all $i \in \mathcal{S}_t$. Then, the edge server aggregates the local information in (3) to compute the new estimates $\mathbf{w}_{t+1}$ as:

$$\mathbf{w}_{t+1} = \frac{1}{|\mathcal{S}_t|} \sum_{i \in \mathcal{S}_t} \psi_{i,t}. \quad (4)$$

It is straightforward to see that the combination of (3) and (4) is equivalent to the direct SGD implementation in (2). The latter implementation is known as federated averaging (FedAvg), and is typically preferred for privacy reasons [15], since transmitting the gradients might reveal information about the data. However, the first implementation is less sensible to errors introduced by quantization effects, which can be controlled through the choice of the step-size μ . For this reason, following the idea of [10], we will act on the source encoder of each transmitting device, dynamically adapting the quantization level in order to strike the best trade-off between power, latency, and learning performance. Furthermore, adding quantization noise to gradients naturally induces a differential privacy behavior [38], [39]. Thus, in the sequel, we will consider only the direct SGD implementation in (2), thus modeling its energy consumption, latency, and learning performance. Similar expressions hold for the FedAvg implementation in (3)-(4).

A. Dithered quantization of uplink data

We assume that each device i is endowed with a dynamic uniform quantizer that uses $b_{i,t}$ bits to transmit data to the AP at time t . The quantizer used by device i at time t is defined by the

following vector mapping $\mathbf{q}(\cdot; b_{i,t})$ as:

$$\mathbf{q}(\mathbf{z}; b_{i,t}) = [l_1 \Delta_{i,t}, \dots, l_n \Delta_{i,t}]^T = \mathbf{z} + \mathbf{e}_{i,t}(\mathbf{z}), \quad (5)$$

where the entries of $\mathbf{z}$, the dynamic quantization step $\Delta_{i,t} > 0$, and the error $\mathbf{e}_{i,t}$ satisfy

$$\begin{cases} \Delta_{i,t} = \frac{A}{2^{b_{i,t}} - 1}, \quad \forall i, t, \\ (l_n - 1/2)\Delta_{i,t} \leq z_n \leq (l_n + 1/2)\Delta_{i,t}, \quad \forall n, i, t, \\ -\frac{\Delta_{i,t}}{2} \mathbf{1}_m \leq \mathbf{e}_{i,t}(\mathbf{z}) \leq \frac{\Delta_{i,t}}{2} \mathbf{1}_m, \quad \forall \mathbf{z}, i, t, \end{cases} \quad (6)$$

with A denoting the size of input data variation. Conditioned on the input, the quantization error $\mathbf{e}_{i,t}(\mathbf{z})$ in (5) is deterministic. This induces a correlation among the quantization errors at different times, which may affect the convergence properties of the iterative algorithm (2). To avoid undesired error correlations, we introduce dithering [40], [41]. In particular, the dither added to randomize the quantization effects satisfies a special condition, namely the Schuchman condition, as in subtractively dithered systems [42]. Then, adding to each component z_n a dither noise $\nu_{n,i,t}$ of i.i.d. uniformly distributed random variables on $[-\Delta_{i,t}/2, \Delta_{i,t}/2]$ independent of the input sequence, the resultant error $\epsilon_{n,i,t}$ becomes

$$\epsilon_{n,i,t} = q(z_n + \nu_{n,i,t}; b_{i,t}) - (z_n + \nu_{n,i,t}), \quad (7)$$

which is now a uniformly distributed random variable on $[-\Delta_{i,t}/2, \Delta_{i,t}/2]$, independent of the input sequence. To implement (2), the devices transmit dithered quantized versions of $\mathbf{g}_{i,t}(\mathbf{w}_t)$ at time t , which from (7) read as:

$$\mathbf{q}(\mathbf{g}_{i,t}(\mathbf{w}_t) + \boldsymbol{\nu}_{i,t}; b_{i,t}) = \mathbf{g}_{i,t}(\mathbf{w}_t) + \underbrace{\boldsymbol{\nu}_{i,t} + \boldsymbol{\epsilon}_{i,t}}_{\mathbf{v}_{i,t}(b_{i,t})}, \quad (8)$$

where $\mathbf{v}_{i,t}(b_{i,t})$ is an additive noise term with zero mean and covariance $\frac{\Delta_{i,t}^2}{6} \mathbf{I}$. The dithered gradients in (8) are encoded into $m \cdot b_{i,t}$ bits, and transmitted by device i to the AP at time t .

B. Performance of adaptive federated learning

In this section, we derive expressions for the learning performance of the proposed FL strategy. Using the dithered quantized gradients (8) in (2), the SGD recursion is given by:

$$\begin{aligned} \mathbf{w}_{t+1} &= \mathbf{w}_t - \mu \sum_{i \in \mathcal{S}_t} \mathbf{q}(\mathbf{g}_{i,t}(\mathbf{w}_t) + \boldsymbol{\nu}_{i,t}; b_{i,t}) \\ &= \mathbf{w}_t - \mu \sum_{i \in \mathcal{S}_t} [\mathbf{g}_{i,t}(\mathbf{w}_t) + \mathbf{v}_{i,t}(b_{i,t})] = \mathbf{w}_t - \mu \sum_{i \in \mathcal{S}_t} \left[\mathbb{E}\{\mathbf{g}_{i,t}(\mathbf{w}_t)\} + \mathbf{s}_{i,t}(\mathbf{w}_t, b_{i,t}) \right], \end{aligned} \quad (9)$$

where $s_{i,t}(\mathbf{w}_t, b_{i,t})$ is a gradient noise process defined as:

$$\mathbf{s}_{i,t}(\mathbf{w}_t, b_{i,t}) = \mathbf{g}_{i,t}(\mathbf{w}_t) - \mathbb{E}\{\mathbf{g}_{i,t}(\mathbf{w}_t)\} + v_{i,t}(b_{i,t}), \quad (10)$$

for all $i \in \mathcal{S}_t$ and $t \geq 0$, which depends on both data and dithered quantization statistics. Then, from (9), the aim of the analysis is to find expressions for the algorithm's performance in terms of learning accuracy and convergence rate. Let us denote by G_t the performance metric for the learning problem at time slot t . Depending on the specific task, G_t might represent different metrics as, e.g., prediction error or classification accuracy. Also, let us denote by α_t the convergence rate at time slot t . Now, performing a mean-square analysis for the recursion in (10) can be a formidable task for general cost functions. Thus, in the sequel, we consider two different assumptions on the global loss function in (1).

1) Strongly convex objective. In this case, we assume that (1) has a favorable structure that helps making the problem mathematically tractable. In particular, letting

$$J(\mathbf{w}) = \sum_{i=1}^N \mathbb{E}\{J_i(\mathbf{w}; \mathbf{x}_i, y_i)\},$$

we consider the following assumption [35].

Assumption 1: The aggregate cost $J(\mathbf{w})$ is twice differentiable, ν -strongly convex, and its gradient is δ -Lipschitz.

This is the case of many important learning paradigms, spanning from least-mean squares adaptation, support vector machines, logistic regression, and so on [35]. As we will see in the sequel, Assumption 1 is useful to give closed form expressions for the steady-state performance and learning rate of our federated learning strategy (cf. (13) and (15)), which are in turn important to control the resource allocation of our system in a very efficient manner using Lyapunov optimization (cf. Sec. III). Furthermore, following similar arguments as in [35, Ch. 5], we consider the following assumption on the gradient noise process affecting (9).

Assumption 2: The gradient noise process in (10) is zero-mean and satisfies

$$\mathbb{E}\{\|\mathbf{s}_{i,t}(\mathbf{w}_t, b_{i,t})\|^2 | \mathcal{F}_{t-1}\} \leq \beta^2 \|\mathbf{w}_t - \mathbf{w}^o\|^2 + \sigma_s^2, \quad (11)$$

for all $i = 1, \dots, N$ and $t \geq 0$, and some β^2 and σ_s^2 , where $\mathcal{F}_{t-1}$ is the filtration of the random process $\mathbf{w}_t$ up to time $t-1$, and $\mathbf{w}^o$ is the global minimum of $J(\mathbf{w})$ (which is guaranteed to exist under Assumption 1). Assumption 2 ensures boundedness of the second-order moments of the gradient noise process in (10), and is instrumental to derive the mean-square performance of the

algorithm in (9). From (9) and (10), it is clear that the performance of the SGD algorithm depends on the temporal behavior of the set $\mathcal{S}_t$ (i.e., which devices are transmitting) and $\{b_{i,t}\}_{i \in \mathcal{S}_t}$ (i.e., how many bits encode the information transmitted by each device). Of course, the two variables are related to each other as:

$$i \in \mathcal{S}_t \iff b_{i,t} \neq 0, \quad (12)$$

i.e., a device belongs to the set of transmitting nodes if the number of uploaded bits is different from zero, and viceversa. Thus, the set $\mathcal{S}_t$ is fully known given the quantization bits $\{b_{i,t}\}_{i=1}^N$, which represent the variables to be selected by the resource allocation algorithm. In particular, for a fixed transmission scheme (i.e., $\mathcal{S}_t = \mathcal{S} \iff b_{i,t} = b_i$ for all t), under Assumptions 1 and 2, for any $0 < \mu < 2\nu/(\delta^2 + \beta^2)$, the Mean-square Deviation (MSD) and the Excess Risk (ER) can be expressed as [35, Ch. 5]:

$$\text{MSD} = \lim_{t \rightarrow \infty} \mathbb{E}\{\|\mathbf{w}_t - \mathbf{w}^o\|^2\} = \frac{\mu}{2} \text{Tr} \left\{ \left(\sum_{i \in \mathcal{S}} \mathbf{H}_i \right)^{-1} \left(\sum_{i \in \mathcal{S}} \mathbf{R}_i(b_i) \right) \right\}, \quad (13)$$

$$\text{ER} = \lim_{t \rightarrow \infty} \mathbb{E}\{J(\mathbf{w}_t) - J(\mathbf{w}^o)\} = \frac{\mu}{4} \text{Tr} \left\{ \sum_{i \in \mathcal{S}} \mathbf{R}_i(b_i) \right\} \quad (14)$$

where $\mathbf{H}_i = \nabla_{\mathbf{w}}^2 \mathbb{E}\{J_i(\mathbf{w}^o; \mathbf{x}_i, y_i)\}$ is the local Hessian, and

$$\mathbf{R}_i(b_i) = \mathbb{E}\{\mathbf{s}_i(\mathbf{w}^o, b_i) \mathbf{s}_i(\mathbf{w}^o, b_i)^T | \mathcal{F}_{t-1}\}$$

is the covariance matrix of the local gradient noise, both evaluated at $\mathbf{w}^o$, for all $i = 1, \dots, N$. Also, we can approximate the convergence rate (i.e., the rate at which the error variance $\mathbb{E}\{\|\mathbf{w}_t - \mathbf{w}^o\|^2\}$ approaches its steady-state region) of the algorithm in (9) by [35, Ch. 5]:

$$\alpha = 1 - 2\mu \lambda_{\min} \left(\sum_{i \in \mathcal{S}} \mathbf{H}_i \right). \quad (15)$$

Thus, hinging on (13) and (15), a possible way to define approximate learning performance metrics of (9) at time t is:

$$G_t = \text{MSD}(\mathcal{S}_t, \{b_{i,t}\}_{i=1}^N) \quad (16)$$

$$\alpha_t = \alpha(\mathcal{S}_t, \{b_{i,t}\}_{i=1}^N). \quad (17)$$

Alternatively, the learning accuracy G_t can be defined with respect to the ER in (14). Of course, the expressions (16) and (17) represent only (instantaneous) approximations of the learning performance achieved by (9), considering the values of $\{b_{i,t}\}_{i \in \mathcal{S}_t}$ at a given time-slot t . However,

as we will show in the numerical results, an accurate performance prediction can still be achieved, thanks to the fast adaptation capabilities of algorithm (9) in both training and steady-state phases.

2) Non-convex objective. The non-convex scenario is typical of many FL tasks involving, e.g., (deep) neural network training. In such a case, we almost never have theoretical performance metrics, which enable reliable prediction of the algorithm's accuracy, following a model-based approach. Thus, we follow an alternative data-driven approach, which involves an online estimation of learning performance, i.e., accuracy and convergence rate, in order to drive the dynamic resource allocation. In Sec. III, we will first derive the resource allocation framework assuming to have access to some closed form expression of the performance metrics G_t and α_t , as in (13) and (15). Then, in Sec. IV, we will extend the proposed framework to handle the case where closed form expressions for G_t and α_t are not generally available and must be inferred online from the data, which is typical of non-convex scenarios. Specific details of the proposed data-driven approach will be provided in Sec. IV.

C. Latency of SGD iterations

The latency necessary to perform one SGD iteration at time t , together with the convergence rate in (17), quantify the time the FL algorithm in (9) needs to learn and adapt. In particular, the overall latency of one iteration is composed of four main sources of delay, which vary over time due to availability of resources (radio and computation) and wireless channel states.

(i) The *local processing time* at the i -th device, necessary to compute the local gradient $g_{i,t}(\mathbf{w}_t)$, is given by:

$$L_{i,t}^l = \frac{N_i^l}{f_{i,t}^l}, \quad i \in \mathcal{S}_t, \quad (18)$$

where N_i^l is the number of CPU cycles necessary to perform this task, and $f_{i,t}^l$ is the CPU frequency of device i at time t .

(ii) The *uplink communication time*, necessary to upload the local gradients to the edge server. Since the i -th device adopts a dithered quantization scheme that encodes local gradients into $m \cdot b_{i,t}$ bits at time t , this latency term reads as:

$$L_{i,t}^u = \frac{m \cdot b_{i,t}}{R_{i,t}^u}, \quad i \in \mathcal{S}_t, \quad (19)$$

where $R_{i,t}^u$ is the uplink data rate. In principle, if device i belongs to $\mathcal{S}_t$, it is going to compute and transmit at time t , incurring in the delays (18) and (19).

(iii) The *remote processing time* at the edge server, necessary to produce the global estimate in (9) is given by

$$L_t^r = \frac{O |\mathcal{S}_t|}{f_t^r}, \quad (20)$$

where O is the number of CPU cycles necessary to perform one summation (between m -dimensional vectors), the cardinality $|\mathcal{S}_t|$ denotes the number of transmitting nodes, and f_t^r is the CPU frequency of the edge server.

(iv) A *downlink communication time*, say $L_{i,t}^d$, $i = 1, \dots, N$, necessary to send the global estimate $\mathbf{w}_{t+1}$ back to the devices. Here, we assume that the number of bits used to encode downlink data is a fixed value, which is chosen sufficiently large to have a negligible impact on the performance of the algorithm in (9). Also, since our interest is mainly focused on uplink communications, the downlink communication time $L_{i,t}^d$ is assumed to be given and ensured by the AP at any slot; thus, it will not be optimized over time.

Finally, to evaluate the overall latency of each SGD iteration, we need to consider the maximum among communication delays of all transmitting (and receiving) devices in the overall latency, plus the computation time at the edge server, which at a given time t reads as:

$$L_t = \max_{i \in \mathcal{S}_t} \{L_{i,t}^l + L_{i,t}^u\} + L_t^r + \max_{i=1, \dots, N} \{L_{i,t}^d\}. \quad (21)$$

As we will show in the sequel, our aim is to keep the average value of L_t in (21), i.e., the average time of SGD iterations, below a given threshold.

D. Power consumption

In this paragraph, we evaluate the power consumption of the proposed federated learning strategy. We consider two sources of power consumption for each device: local computation and transmission; then, we take into account the power spent for computation at the edge server. At time t , the power spent by device i for local computation is:

$$p_{i,t}^c = \kappa_l (f_{i,t}^l)^3, \quad (22)$$

where κ_l is the effective switched capacitance of the processor [43]. Moreover, given the uplink data rate $R_{i,t}^u$, the power spent for uplink transmission is computed by inverting the Shannon formula, thus obtaining:

$$p_{i,t}^u = \frac{B_{i,t}^u N_0}{h_{i,t}^u} \left[\exp \left(\frac{R_{i,t}^u \ln 2}{B_{i,t}^u} \right) - 1 \right], \quad (23)$$

where $B_{i,t}^u$ is the bandwidth assigned to device i at time t , $h_{i,t}^u$ is the uplink channel power gain, and N_0 is the noise power spectral density. At each time t , we assume that the AP allocates orthogonal frequency channels with pre-allocated bandwidth $B_{i,t}^u$ to all transmitting users (i.e., those belonging to $\mathcal{S}_t$). On the other side, at time t , the power spent by the edge server for computation of (9) is given by:

$$p_{s,t}^c = \kappa_r (f_t^r)^3, \quad (24)$$

where f_t^r and κ_r are the CPU frequency and the effective switched capacitance of the ES processor, respectively.

In this paper, our goal is to minimize the long-term average of the system power consumption, given by the sum of the devices and ES powers, which reads as:

$$p_t^{\text{tot}} = \sum_{i \in \mathcal{S}_t} (p_{i,t}^u + p_{i,t}^c) + p_{s,t}^c. \quad (25)$$

In the following section, we will formulate the proposed dynamic strategy for wireless network edge optimization, aimed at performing energy-efficient FL with guaranteed latency and learning performance requirements.

III. DYNAMIC OPTIMIZATION OF WIRELESS EDGE RESOURCES

We can now formulate the problem of dynamic resource allocation for FL. The aim is to find the optimal joint dynamic resource allocation of radio (i.e., the set $\mathcal{S}_t$ of transmitting devices, uplink data rates $\{R_{i,t}^u\}_{i \in \mathcal{S}_t}$, quantization bits $\{b_{i,t}\}_{i \in \mathcal{S}_t}$) and computation (i.e., CPU cycles at devices $\{f_{i,t}^l\}_{i \in \mathcal{S}_t}$ and at the server f_t^r) resources to minimize the long-term average system power consumption in (25), with constraints on the average learning performance in (16)-(17), and the average latency in (21). Then, the dynamic resource allocation problem can be mathematically cast as:

$$\begin{aligned} & \min_{\Psi_t} \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbb{E} \{p_{\tau}^{\text{tot}}\} \\ & \text{subject to} \\ & (a) \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbb{E} \{L_{\tau}\} \leq \bar{L}; \quad (b) \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbb{E} \{G_{\tau}\} \leq \bar{G}; \quad (c) \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{\tau=0}^{t-1} \mathbb{E} \{\alpha_{\tau}\} = \bar{\alpha}; \\ & \left. \begin{aligned} & b_{i,t} \in \mathcal{B}_i \quad \forall i \in \mathcal{S}_t, t; \quad R_i^{\min} \leq R_{i,t}^u \leq R_{i,t}^{\max} \quad \forall i \in \mathcal{S}_t, t; \\ & f_i^{\min} \leq f_{i,t}^l \leq f_i^{\max} \quad \forall i \in \mathcal{S}_t, t; \quad f^{r,\min} \leq f_t^r \leq f^{r,\max}, \quad \forall t, \end{aligned} \right\} \mathcal{X}_t \end{aligned} \quad (26)$$

where $\Psi_t = [\{b_{i,t}\}_{i \in \mathcal{S}_t}, \{R_{i,t}^u\}_{i \in \mathcal{S}_t}, \{f_{i,t}^l\}_{i \in \mathcal{S}_t}, f_t^r]$, and the expectations are taken with respect to the random channel states, whose statistics are supposed to be unknown. The constraints of (26) have the following meaning: (a) the average latency of SGD iterations does not exceed a predefined value $\bar{L}^1$; (b) the average performance metric G_t does not exceed a predefined value $\bar{G}$; of course, if G_t represents an accuracy metric, the sign of the constraint should be reversed, i.e., the average accuracy must be greater than a certain target value; (c) the average learning rate is constrained to be equal to $\bar{\alpha}$; finally, the constraints in $\mathcal{X}_t$ impose that $\{b_{i,t}\}_{i \in \mathcal{S}_t}$ can take values only from a finite set $\mathcal{B}_i$ of discrete quantization bits, and impose instantaneous bounds (e.g., budget constraints, minimum rates and CPU frequencies) on the resource variables $\{R_{i,t}^u\}_{i \in \mathcal{S}_t}, \{f_{i,t}^l\}_{i \in \mathcal{S}_t}, f_t^r$. In the sequel, we introduce a dynamic algorithmic framework to solve the long-term optimization problem (26).

A. Algorithmic solution via stochastic optimization

We now introduce a method to transform (26) into a stability problem, building on the tools of stochastic Lyapunov optimization [36]. In particular, to deal with the long-term constraints (a)-(c), we introduce three *virtual queues*. The first one, used to impose (a), evolves as:

$$Z_{t+1} = \max \{0, Z_t + \epsilon_z (L_t - \bar{L})\}, \quad (27)$$

where ϵ_z is a positive step-size used to control the convergence speed of the algorithm. The second virtual queue, used to impose constraint (b), read as:

$$Q_{t+1} = \max \{0, Q_t + \epsilon_q (G_t - \bar{G})\}, \quad (28)$$

with $\epsilon_q > 0$. Finally, the virtual queue associated with constraint (c) is given by:

$$Y_{t+1} = Y_t + \epsilon_y (\alpha_t - \bar{\alpha}), \quad (29)$$

with $\epsilon_y > 0$. Note that virtual queue Y_t reads slightly differently from the others, due to the fact that it is used to impose an equality constraint [36]. Interestingly, ensuring the mean-rate stability of the virtual queues in (27)-(29) is equivalent to satisfy the three corresponding constraints [36]. To this aim, we first define the Lyapunov function as:

$$\mathcal{U}_t = \mathcal{U}(\Phi_t) = \frac{1}{2} (Z_t^2 + Q_t^2 + Y_t^2), \quad (30)$$

¹More sophisticated probabilistic constraints can be used as shown in [44].

where $\Phi_t = [Z_t, Q_t, Y_t]$. The Lyapunov function is a measure of the congestion state of the virtual queues, and is fundamental to define the *drift-plus-penalty* function [36]:

$$\Delta_t^p = \mathbb{E} \left\{ \mathcal{U}_{t+1} - \mathcal{U}_t + V \cdot p_t^{\text{tot}} \middle| \Phi_t \right\}. \quad (31)$$

The drift-plus-penalty function in (31) is the conditional expected change of $\mathcal{U}_t$ over successive slots, with a penalty factor that weights the objective function of (26), with a weighting parameter V . Now, if Δ_t^p is lower than a finite constant for all t , the virtual queues are stable and the optimal solution of (26) is asymptotically reached as V increases [36, Th. 4.8]. In practical scenarios with finite V values, the higher is V , the more importance is given to the objective function, rather than the virtual queue backlogs, thus pushing the solution toward optimality while still guaranteeing stability of the system. Thus, we proceed by minimizing an upper-bound of the drift-plus-penalty in (31), which reads as (the derivations are omitted due to the lack of space):

$$\Delta_t^p \leq \zeta + \mathbb{E} \left\{ Z_t \left(\tilde{L}_t - \bar{L} \right) + Q_t \left(G_t - \bar{G} \right) + Y_t \left(\alpha_t - \bar{\alpha} \right) + V \cdot p_t^{\text{tot}} \middle| \Phi_t \right\}, \quad (32)$$

where G_t , α_t and p_t^{tot} are defined in (16), (17) and (25), respectively; whereas, ζ is a positive finite constant, and $\tilde{L}_t$ is an upperbound of (21) given by:

$$\tilde{L}_t = \sum_{i \in \mathcal{S}_t} \left(\frac{N_i^l}{f_{i,t}^l} + \frac{m \cdot b_{i,t}}{R_{i,t}^u} \right) + \frac{O|\mathcal{S}_t|}{f_t^r} + \max_{i=1, \dots, N} \left\{ L_{i,t}^d \right\}. \quad (33)$$

Finally, using stochastic approximation arguments [36], we optimize (32) removing the expectation, thus obtaining the following deterministic problem at each time-slot t (where all the constant terms of (32) are omitted):

$$\begin{aligned} \min_{\Psi_t} \quad & Z_t \tilde{L}_t + Q_t G_t + Y_t \alpha_t + V \cdot p_t^{\text{tot}} \\ \text{subject to} \quad & \Psi_t \in \mathcal{X}_t \end{aligned} \quad (34)$$

where $\mathcal{X}_t$ is the instantaneous feasible set of (26). Because of the structure of $\mathcal{X}_t$, (34) is a mixed-integer nonlinear optimization problem, which might be very complicated to solve. However, for any given $\{b_{i,t}\}_{i=1}^N$ at time-slot t , it is easy to see that (34) is separable into three sub-problems that admit closed form solutions for the optimal uplink data rates, the optimal CPU clock frequency of devices, and the optimal CPU clock frequency of the edge server, respectively. In the sequel, we present the formulation and the solution of the three sub-problems.

B. Uplink radio resource allocation

The uplink radio resource allocation sub-problem aims at optimizing the transmission rates $\{R_{i,t}^u\}_{i \in \mathcal{S}_t}$ of each transmitting device at time-slot t , once the quantization bits $\{b_{i,t}\}_{i=1}^N$ have been fixed. Of course, if $b_{i,t} = 0$, we have $R_{i,t}^u = 0$. Instead, for any $b_{i,t} > 0$, from (34), (33) and (25), we obtain:

$$\begin{aligned} \min_{\{R_{i,t}^u\}_{i \in \mathcal{S}_t}} \quad & \sum_{i \in \mathcal{S}_t} \left(Z_t \frac{m \cdot b_{i,t}}{R_{i,t}^u} + V p_{i,t}^u \right) \\ \text{subject to} \quad & R_i^{\min} \leq R_{i,t}^u \leq R_{i,t}^{\max} \quad \forall i \in \mathcal{S}_t, \end{aligned} \quad (35)$$

where $R_{i,t}^{\max}$ is the maximum rate achievable using the maximum transmitted power, and $R_i^{\min}$ represents the minimum rate that a user should use in the case of transmission. Of course, in (35), there is an intrinsic admission control condition embedded in terms of feasibility. In particular, denoting the set of nodes for which problem (35) is feasible by:

$$\mathcal{A}_t = \{i \in \{1, \dots, N\} : R_{i,t}^{\max} > R_i^{\min}\}, \quad (36)$$

it clearly holds that the set $\mathcal{S}_t$ of transmitting nodes must be selected as a subset of $\mathcal{A}_t$ ($\mathcal{S}_t \subseteq \mathcal{A}_t$). If problem (35) is feasible, it is also strictly convex with respect to the rates $\{R_{i,t}^u\}_{i \in \mathcal{S}_t}$, and admits a unique closed-form solution. In particular, the Lagrangian function of (35) writes as:

$$\mathcal{L}_i^R = \sum_{i \in \mathcal{S}_t} \left[\left(Z_t \frac{m \cdot b_{i,t}}{R_{i,t}^u} + V p_{i,t}^u \right) - \delta_i (R_{i,t}^u - R_i^{\min}) + \xi_i (R_{i,t}^u - R_{i,t}^{\max}) \right], \quad (37)$$

where δ_i and ξ_i are the Lagrange multipliers associated with the constraints of (35) over the variable $R_{i,t}^u$, for $i \in \mathcal{S}_t$. Then, the Karush-Kuhn-Tucker (KKT) conditions of the strictly convex problem (35) are given by:

$$i) \quad \frac{\partial \mathcal{L}_i^R}{\partial R_{i,t}^u} = -Z_t \frac{m \cdot b_{i,t}}{(R_{i,t}^u)^2} + \frac{V N_0}{h_{i,t}^u} \exp \left(\frac{R_{i,t}^u \ln(2)}{B_{i,t}^u} \right) - \delta_i + \xi_i = 0; \quad (38)$$

$$ii) \quad R_{i,t}^u \geq R_i^{\min}, \quad \delta_i \geq 0, \quad \delta_i (R_{i,t}^u - R_i^{\min}) = 0; \quad (39)$$

$$iii) \quad R_{i,t}^u \leq R_{i,t}^{\max}, \quad \xi_i \geq 0, \quad \xi_i (R_{i,t}^u - R_{i,t}^{\max}) = 0; \quad (40)$$

for all $i \in \mathcal{S}_t$, where we used (23) for $p_{i,t}^u$. Now, exploiting the principal branch of the Lambert function $W(\cdot)$ [45], the KKT conditions (38)-(40) can be solved in closed form as:

$$R_{i,t}^u = \left[\frac{2B_{i,t}^u}{\ln(2)} W \left(\frac{\ln(2)}{B_{i,t}^u} \sqrt{\frac{Z_t m \cdot b_{i,t} h_{i,t}^u}{2V N_0}} \right) \right]_{R_i^{\min}}^{R_{i,t}^{\max}} \quad (41)$$

for all $i \in \mathcal{S}_t$, if problem (35) is feasible.

C. Local computing resources allocation

The local computing resource allocation problem aims at optimizing the CPU frequencies $\{f_{i,t}^l\}_{i \in \mathcal{S}_t}$ of the transmitting (and computing) devices. From (34), (33) and (25), for a given time t , it is easy to see how the local computing resource allocation problem decouples over the computing devices and over the iterations within a slot. Thus, we obtain the following sub-problem at each device:

$$\begin{aligned} \min_{f_{i,t}^l} \quad & Z_t \frac{N_i^l}{f_{i,t}^l} + V \kappa_l (f_{i,t}^l)^3 \\ \text{subject to} \quad & f_i^{\min} \leq f_{i,t}^l \leq f_i^{\max}, \end{aligned} \quad (42)$$

for all $i \in \mathcal{S}_t$. Problem (42) is strictly convex and enjoys a simple closed form solution. Indeed, solving the KKT conditions, it is immediate to obtain:

$$f_{i,t}^l = \left[\left(\frac{Z_t N_i^l}{3 \kappa_l V} \right)^{\frac{1}{4}} \right]_{f_i^{\min}}^{f_i^{\max}} \quad \text{for all } i \in \mathcal{S}_t. \quad (43)$$

Note that (43) contains implicitly the admission control condition of the radio resource problem (35), i.e., a local computation is needed only if data can be subsequently uploaded to the AP. Thus, if $i \notin \mathcal{S}_t$, then $R_{i,t}^u = 0$ and also $f_{i,t}^l = 0$.

D. Remote CPU frequencies allocation

The remote computing resource allocation problem aims at optimizing the CPU frequency f_t^r of the ES. From (34), (33) and (25), for a given $\mathcal{S}_t$ at time t , the remote computing resource allocation problem writes as:

$$\begin{aligned} \min_{f_t^r} \quad & Z_t \frac{O |\mathcal{S}_t|}{f_t^r} + V \kappa_r (f_t^r)^3 \\ \text{subject to} \quad & f^{r,\min} \leq f_t^r \leq f^{r,\max}. \end{aligned} \quad (44)$$

Problem (44) is a strictly convex optimization problem, which enjoys a simple closed form solution. Then, solving the KKT conditions, it is straightforward to see that the optimal remote CPU cycle frequency is given by

$$f_t^r = \left[\left(\frac{Z_t O |\mathcal{S}_t|}{3 \kappa_r V} \right)^{\frac{1}{4}} \right]_{f^{r,\min}}^{f^{r,\max}}. \quad (45)$$

E. Online selection of transmitting users, quantization bits, and wireless edge resources

In Secs. III.B, III.C, and III.D, we have derived the optimal allocation of edge resources as a function of the quantization bits $\{b_{i,t}\}_{i=1}^N$. Thus, exploiting (41), (43), and (45), the objective of (34), say $\tilde{\Delta}_t^p$, can now be expressed as a function of only $\{b_{i,t}\}_{i=1}^N$, i.e., $\tilde{\Delta}_t^p = \tilde{\Delta}_t^p(\{b_{i,t}\}_{i=1}^N)$. In principle, to find the optimal solution of (34) at each time-slot t , one should compute the optimal allocation of edge resources for all possible combinations of $\{b_{i,t}\}_{i=1}^N$, evaluate the corresponding objective function $\tilde{\Delta}_t^p$ in (34), and then select the one that yields the lowest value. This approach faces a main challenge: Even if for the single choice of $\{b_{i,t}\}_{i=1}^N$ the resource allocation is efficient (cf. (41), (43), and (45)), the overall search procedure has still a complexity that grows exponentially with the number N of devices and the maximum cardinality of the set of quantization bits (i.e., $\max_i |\mathcal{B}_i|$). For this reason, in the sequel we act some simplifications to reduce the complexity of (34), while still achieving good performance. In particular, instead of performing an exhaustive search over all possible combinations, we use an iterative greedy approach that, starting from the empty set of transmitting nodes, iteratively adds the most convenient devices, selecting jointly the best number of quantization bits and the associated edge resources in (41), (43), and (45). The method keeps adding devices from the admissible set $\mathcal{A}_t$ in (36) until the resulting value of the objective $\tilde{\Delta}_t^p$ in (34) decreases, and stops when there is no more incentive (in terms of reduction of the objective function) in letting other nodes to transmit any bit of information. Of course, if $\mathcal{A}_t$ is empty, the t -th iteration of the FL algorithm in (9) does not take place. Such greedy method drastically reduces the complexity of the procedure, which becomes polynomial in N and $\max_i |\mathcal{B}_i|$. Then, once the resource have been selected, the federated learning algorithm is updated as in (9), and, finally, the virtual queues Z_t , Q_t and Y_t are updated as in (27), (28) and (29), respectively. The overall dynamic optimization procedure for adaptive federated learning is illustrated in Algorithm 1.

Remark: Of course, there are no guarantees that the proposed greedy procedure in Algorithm 1 finds the optimal solution of (34), inevitably representing only an approximation of it. However, in the context of stochastic Lyapunov optimization, our approach can be interpreted as a C -additive approximation [36, p. 59], which entails inexact solutions (with bounded error) of the drift-plus-penalty method in (34) at each iteration t . In our case, since the objective and the feasible set of (34) are both bounded for all t , the proposed greedy approach clearly leads to a valid C -approximation. In Sec. IV, we will numerically assess the performance of the proposed

Algorithm 1 Model-based Dynamic Resource Optimization for Adaptive Federated Learning

for $t \geq 0$ **do**
S1. Observe $Z_t, Q_t, Y_t, \{h_{i,t}^u\}_{i=1}^N$ and compute $\mathcal{A}_t$ in (36).

if $\mathcal{A}_t = \emptyset$ **then**

 | **continue**
else
S2. Find $\mathcal{S}_t$, the quantization bits $\{b_{i,t}\}_{i \in \mathcal{S}_t}$, and the edge resources through the greedy procedure:

 Set $\mathcal{S}_t = \emptyset, \{b_{i,t}\}_{i=1}^N = 0, \Delta_c = \infty, Flag = 1$
while $Flag$ **do**
S3. $[\Delta_j^*, (j^*, b_j^*)] = \arg \min_{j \in \mathcal{A}_t \setminus \mathcal{S}_t, b_j \in \mathcal{B}_i} \tilde{\Delta}_t^p(\mathcal{S}_t \cup \{j\}, b_j)$
if $\Delta_j^* < \Delta_c$ **then**

 | $\mathcal{S}_t \leftarrow \mathcal{S}_t \cup \{j^*\}; \quad b_{j^*,t} = b_j^*; \quad \Delta_c = \Delta_j^*;$

 | $R_{j^*,t}^u$ as in (41); $f_{j^*,t}^l$ as in (43); f_t^r as in (45);

else

 | $Flag = 0$
end
end
S4. Update the federated learning algorithm as in (9);

S5. Update the virtual queues Z_t, Q_t and Y_t as in (27), (28) and (29), respectively;

end
end

dynamic resource allocation strategy for adaptive federated learning at the wireless network edge.

IV. DATA-DRIVEN RESOURCE ALLOCATION FOR ADAPTIVE FEDERATED LEARNING

In the previous section, we have proposed a model-based algorithm for dynamic resource allocation to enable FL at the wireless edge, which exploits closed-form expressions for the learning performance and convergence rate metrics (e.g., (13) and (15)). Indeed, some convex learning tasks admit closed form expressions for different learning metrics, allowing us to use Algorithm 1. However, in several other cases (e.g., non-convex learning tasks such as deep

neural network training), we do not have knowledge of such performance metrics expressions. Therefore, in this section, we extend the previous strategy incorporating an online mechanism that estimates the learning performance and the convergence rate in a fully data-driven fashion, in order to drive the dynamic resource allocation based on Lyapunov optimization. One of the nicest features of the proposed data-driven approach is that it does not necessarily rely on SGD recursions as in (9), but it works also with more sophisticated gradient-based algorithms such as, e.g., Adam, Adagrad, SGD with momentum, etc. Now, let us assume that the agents collect and process batches of data of size B_t at time-slot t . For simplicity, we assume that the batch size B_t is the same for all devices, and that can be selected from a set $\mathcal{C}$ of discrete values at each time slot t . Then, assuming that N_i^l in (18) is the number of CPU cycles to compute the local gradient from one data unit, if we have batches of B_t data, the local processing time in (18) will be simply multiplied by a factor B_t . Furthermore, considering a general gradient-based algorithm (e.g., Adam), the remote processing time is obtained as in (20) simply generalizing the meaning of the constant O that, starting from the received local gradient, represents the number of CPU cycles necessary to perform the single step of the gradient-based algorithm for each device. Of course, the overall remote processing time is still proportional to the number of transmitting devices. To estimate online the learning performance, we assume that either the ES is provided with a validation set $\mathcal{T}$, or the agents can sense an additional batch $\mathcal{T}$ of data at each time-slot if needed. Then, two task-dependent functions $\hat{G}_t$ and $\hat{\alpha}_t$ are introduced to measure online the learning performance G_t and convergence rate α_t , respectively. As an example, let us consider a classification task, whose validation (or batch) accuracy and its moving average with length $2K$ can be used to estimate learning performance and convergence rate as:

$$\hat{G}_t = \frac{1}{|\mathcal{T}|} \sum_{y \in \mathcal{T}} \mathbb{I}(\hat{y}_t = y), \quad (46)$$

$$\hat{\alpha}_t = \frac{1}{K} \sum_{\tau=t-K}^{K-1} (\hat{G}_\tau - \hat{G}_{\tau-1}), \quad (47)$$

where $\hat{y}_t$ is the prediction for data unit y at time-slot t^2 . Clearly, different metrics can be used based on the task and its complexity, e.g. the ratio of gradients norms at successive time-slots can be exploited for tasks whose learning accuracy/error could be difficult or inefficient to estimate

²Note that, if we use the accuracy metric in (46), the long term constraints (b) in (26) must have opposite sign, i.e., the average performance must be greater than or equal to a target threshold $\bar{G}$.

[23]. Then, we propose to exploit the performance estimates in (46)-(47) to drive the Lyapunov-based resource allocation. In particular, we introduce two new virtual queues updates, which are specific of this data-driven approach. The first virtual queue reads as:

$$\hat{Q}_{t+1} = \max \left\{ 0, \hat{Q}_t + \epsilon_q \left(\bar{G} - \hat{G}_t \right) \right\}, \quad (48)$$

and has the goal to drive the estimated learning performance $\hat{G}_t$ above the target $\bar{G}$. The second queue aims at controlling the convergence rate of the FL algorithm, and is updated as:

$$\hat{Y}_{t+1} = \left[\hat{Y}_t + \epsilon_{y,t} (\hat{\alpha}_t - \bar{\alpha}) \right] \cdot \mathbb{I}(\hat{G}_t \leq \bar{G}), \quad (49)$$

where $\epsilon_{y,t}$ is an adaptive (i.e., time-dependent) step-size. The queue evolution defined in (49) is motivated by the fact that, if the distribution of the data is stationary, there is no need to overshoot the convergence rate after the target level of learning performance is reached. Thus, when $\hat{G}_t \geq \bar{G}$ (i.e., the estimated learning performance is greater than the target), the queue Y_t is set to zero and has no more impact on the convergence rate and the resource allocation. Moreover, the dynamic step-size $\epsilon_{y,t}$ is chosen to adapt the update speed of the queue Y_t depending on the distance of $\hat{G}_t$ from the target $\bar{G}$. A possible choice is $\epsilon_{y,t} = \epsilon_y \left| \bar{G} - \hat{G}_t \right|$. The rationale avoids the queue to unnecessarily impact the resource allocation when the learning performance is approaching $\bar{G}$, because the target convergence rate $\bar{\alpha}_t$ is no more achievable at that point. At the same time, non-stationary behaviors can be detected observing a sharp deterioration in $\hat{G}_t$, which reactivates the virtual queue Y_t , thus boosting again the learning process with the desired convergence rate. Moreover, an adaptive step-size $\epsilon_{z,t}$ can also be exploited also for the latency queue in (27) to speed up the adaptability of the method.

Using $\hat{G}_t$ and $\hat{\alpha}_t$ is useful for the virtual queues update in (48) and (49), but they are not explicitly (i.e., mathematically) related to the number of quantization bits and to the batch size, which must be optimized and adapted to drive the learning performance and the convergence rate. Thus, the control action might still not be easily implementable, due to the lack of closed-form expressions for the performance metrics. One possible solution to this issue builds on the following assumptions [10], which are consistently verified both from a theoretical and a numerical point of view (practical examples follow in Sec. V).

Assumption 3: G_t is a monotone non-decreasing function of the quantization bits $\{b_{i,t}\}_{i \in \mathcal{S}_t}$.

Assumption 4: α_t is a monotone non-decreasing function of the quantization bits $\{b_{i,t}\}_{i \in \mathcal{S}_t}$ and of the batch size B_t .

Assumption 3 hinges on the fact that a finer representation of the data generally leads to better learning performance. At the same time, Assumption 4 exploits the fact that, increasing the batch size and the number of quantization bits, the (stochastic) gradient estimates in (2) get better, thus improving the overall convergence rate. Then, under Assumption 3 and Assumption 4, we propose to exploit two surrogate functions, say $\tilde{G}_t$ and $\tilde{\alpha}_t$, which respectively approximate the non-decreasing behavior of G_t and α_t with respect to the quantization bits $\{b_{i,t}\}_{i \in \mathcal{S}_t}$ and the batch size B_t . Of course, there are several possible surrogates that we can exploit, but the best choice depends on the specific performance metric that we need to approximate (e.g., classification accuracy, mean-squared error, etc.). Examples will be given in Sec. V.B, for the case of deep neural network training. The rationale underlying the choice of the surrogates comes again from the concept of C -additive approximation [36] of the drift-plus-penalty method in (34), which makes possible to use inexact updates of the algorithm at each iteration, provided that the approximation error can be bounded within a finite error C . Then, at a given time slot t , exploiting (48), (49), and the surrogate functions $\tilde{G}_t$ and $\tilde{\alpha}_t$ in (34), we solve the following deterministic problem:

$$\min_{\Gamma_t \in \mathcal{Z}_t} Z_t \tilde{L}_t - \hat{Q}_t \tilde{G}_t - \hat{Y}_t \tilde{\alpha}_t + V \cdot p_t^{\text{tot}} \quad (50)$$

where $\Gamma_t = [\{b_{i,t}\}_{i \in \mathcal{S}_t}, \{R_{i,t}^u\}_{i \in \mathcal{S}_t}, \{f_{i,t}^l\}_{i \in \mathcal{S}_t}, f_t^r, B_t]$, and $\mathcal{Z}_t = \mathcal{X}_t \cup \{B_t \in \mathcal{C}\}$, with $\mathcal{C}$ denoting the discrete set of possible choices for the batch size B_t . We solve problem (50) as in the previous case, slightly modifying Algorithm 1, as illustrated in Algorithm 2. Essentially, Algorithm 2 adds a further selection step for the batch size B_t to the greedy procedure of Algorithm 1. This is done with a small additive complexity, since the number of selectable batch sizes is assumed to be small (e.g., 3 or 4 possibilities). Finally, the main steps of the proposed data-driven approach are the same as in Algorithm 1, with the difference that the virtual queues Q_t and Y_t are replaced by $\hat{Q}_t$ and $\hat{Y}_t$ in (48)-(49). This data-driven strategy will be numerically assessed in Sec. V.

V. NUMERICAL RESULTS

In this section, we assess the performance of the proposed method, considering both model-based and data-driven scenarios. In particular, we will exploit the model-based approach for a least-mean squares regression task (in sec. V.A), and the data-driven approach for a classification task aimed at training a deep convolutional neural network (in sec. V.B). We consider a scenario with $N = 9$ devices and one AP equipped with an edge server, as illustrated in Fig. 1. We set

Algorithm 2 Data-driven Dynamic Resource Optimization for Adaptive Federated Learning

for $t \geq 0$ **do**

 S1. Observe $Z_t, \hat{Q}_t, \hat{Y}_t, \{h_{i,t}^u\}_{i=1}^N$ and compute $\mathcal{A}_t$ in (36).

if $\mathcal{A}_t = \emptyset$ **then**

 | **continue**
else

 S2. Find $\mathcal{S}_t$, the quantization bits $\{b_{i,t}\}_{i \in \mathcal{S}_t}$, the batch size B_t , and the edge resources through the greedy procedure:

 Set $\mathcal{S}_t = \emptyset, \{b_{i,t}\}_{i=1}^N = 0, \Delta_c = \infty, Flag = 1$
for $B \in \mathcal{C}$ **do**

 | **while** $Flag$ **do**

 | | S2. $[\Delta_j^*, (j^*, b_j^*)] = \arg \min_{j \in \mathcal{A}_t \setminus \mathcal{S}_t, b_j \in \mathcal{B}_i} \tilde{\Delta}_t^p(\mathcal{S}_t \cup \{j\}, b_j, B)$

 | | **if** $\Delta_j^* < \Delta_c$ **then**

 | | | $\mathcal{S}_t \leftarrow \mathcal{S}_t \cup \{j^*\}; b_{j^*,t} = b_j^*; \Delta_c = \Delta_j^*;$

 | | | $R_{j^*,t}^u$ as in (41); $f_{j^*,t}^l$ as in (43); f_t^r as in (45);

 | | | $B_t = B$

 | | **else**

 | | | $Flag = 0$

 | | **end**

 | **end**
end

S3. Update the federated learning algorithm as in (9), or according to a tailored gradient-based optimizer;

 S4. Update the virtual queues $Z_t, \hat{Q}_t$ and $\hat{Y}_t$ as in (27), (48) and (49), respectively;

end
end

the radio and computation parameters as follows: $N_0 = -174$ dBm/Hz, $A = 25$, $f_r^{\max} = 3.3$ GHz, $f_i^{\max} = 2.5$ GHz, $R_{i,t}^{\max} = B_{i,t}^u \log_2(1 + h_{i,t}^u P_i^{\max} / (N_0 B_{i,t}^u))$ where $P_i^{\max} = 150$ mW for all i ; $B_{i,u}^t$ is assigned equally splitting the overall bandwidth among the devices transmitting at time t . Moreover, $\mathcal{B}_i = [2, 4, 8]$, and $R_{i,t}^d = R_{i,t}^{\max}$ for all i , and t . The channels are generated using the

Fig. 2: MSD vs t , for different values of $\overline{\text{MSD}}$.

ABG model [46], with a carrier frequency of 23 GHz and adding a Rayleigh fading with unit variance. The LMS results are obtained with Matlab, using a PC with an Intel Core i7-7700HQ CPU at a frequency of 2.80 GHz. The data-driven results are obtained using Python and the JAX framework, exploiting an NVIDIA Tesla K80 GPU.

A. Federated Least-Mean Squares

For this learning task, the input data $\mathbf{x}_{i,t} \in \mathbb{R}^{20}$ is related to the corresponding output via a linear model $y_{i,t} = \mathbf{x}_{i,t}^T \mathbf{w}^0 + v_{i,t}$, at each time instant t . In this context, the SGD algorithm in (2) boils down into a federated LMS adaptive algorithm aimed at learning (and tracking over time) the vector $\mathbf{w}^0$ [35]. The devices locations are chosen at random such that the distance of each device from the AP is sampled from a uniform distribution in the interval $[70, 130]$, for all $i = 1, \dots, N$. We assume that input data $\mathbf{x}_{i,t}$ are zero-mean random vectors with covariance matrix $\sigma_{x,i}^2 \mathbf{I}_{20}$ with $\sigma_{x,i}^2 = 1$ for all i . The observation noise $v_{i,t}$ is Gaussian, zero-mean with variance $\sigma_{v,i}^2$ uniformly selected in the interval $[0, 2 \times 10^{-3}]$, for all $i = 1, \dots, N$, independent from the data and among devices. Also, the overall bandwidth is 100 KHz, $O = 2 \cdot 10^3$, $N_i^l = 5 \cdot 10^6$. The step-size μ is set to 0.015. The learning performance is measured in terms of MSD as in (13), and the convergence rate is given by (15).

As a first result, in Fig. 2, we illustrate the learning curve of the FL algorithm in (2), obtained for different values of learning performance constraints $\overline{G} = \overline{\text{MSD}}$, while fixing the learning rate

Fig. 3: (a) Latency vs power consumption, (b)-(c) Power consumption vs V for $\overline{\text{MSD}} = -40$ dB.

to $\bar{\alpha} = 0.99$ and the latency constraint $\bar{L} = 20$ ms. The results are averaged over 50 independent simulations, setting empirically ϵ_z , ϵ_q , ϵ_y to obtain the fastest convergence. From Fig. 2, we can notice how the proposed optimization method is able to guarantee the prescribed performance in terms of learning rate $\bar{\alpha}$ and steady-state accuracy $\overline{\text{MSD}}$.

Then, in Fig. 3 (a) we illustrate the trade-off curve between average latency and TX power consumption (i.e., the sum of powers transmitted by all users) achieved by the proposed method, considering different values of $\overline{\text{MSD}}$ and fixed $\bar{\alpha} = 0.99$. Each point in the curves of Fig. 3 (a) represents a different value of V , whose magnitude grows from right to left. From Fig. 3 (a), increasing V , the method reduces the transmission power up to a limit value (i.e., the optimum) that still enables to guarantee the target latency constraint. As expected, the trade-off gets worse imposing a stricter requirement on learning performance, due to the larger power (and number of bits) necessary to obtain the target performance. Also, from a computation point of view, in Figs. 3 (b)-(c) we illustrate the average remote and local processing power consumption vs V , fixing $\overline{\text{MSD}} = -40$ dB, $\bar{\alpha} = 0.99$ and $\bar{L} = 20$ ms. As expected, the proposed method is able to reduce all the single contributions of the overall power consumption as V increases.

Finally, in Fig. 4, we show the histogram of quantization bits usage for different values of $\overline{\text{MSD}}$, fixing $\bar{\alpha} = 0.99$ and $\bar{L} = 20$ ms. From Fig. 4, we notice how the method requires on average more quantization bits to obtain a stricter requirement on learning performance, due to the finer required representation of transmitted data.

Fig. 4: Histogram of quantization bits, for different $\overline{\text{MSD}}$.

B. Federated Deep Neural Network Training

In this section, we consider a learning task aimed at training a classifier based on deep convolutional neural networks (CNN). We exploit a CNN architecture made of four convolutional layers with 32, 32, 10 and 10 filters, respectively, with final flatten and dense layers; SAME padding, ReLu non-linearities and Batch Normalization are applied after each convolutional layer, and a final Softmax non-linearity is applied after the flatten and dense layers. The filters dimensions are 5×5 for the first two convolutional layers, and 3×3 for the last two layers.

To train the CNN, we use the MNIST dataset [47], which is made of 28×28 grayscale images of handwritten digits divided in 10 classes. The training data is composed of 6×10^4 examples, while the test set is made of 10^4 elements. The loss is the well-known cross-entropy, and the model is trained using a federated ADAM optimizer [48]. The devices locations are random and the distance of each device from the AP is sampled from a uniform distribution in the interval $[20, 80]$, for all $i = 1, \dots, N$. Also, the overall bandwidth is 100MHz, $O = 10^6$, and $N_i^l = 10^7$. The ADAM step-size is set to 0.008, with forgetting factors $\beta_1 = 0.9$, and $\beta_2 = 0.99$.

For this experiment, we use the performance estimate $\hat{G}_t$ in (46), and $\hat{\alpha}_t$ in (47) with $K = 10$. Also, as a surrogate function for the accuracy metric, we exploit:

$$\tilde{G}_t = \sum_{i \in \mathcal{S}_t} \sigma(b_{i,t} - \text{Median}\{\mathcal{B}_i\}), \quad (51)$$

Fig. 5: Accuracy vs t , for different values of $\bar{\alpha}$.

where $\sigma(\cdot)$ is the logistic sigmoid function, and $\text{Median}\{\cdot\}$ represents the median value. Clearly, (51) satisfies Assumption 3. Regarding the convergence rate, we use instead the surrogate $\tilde{\alpha}_t = B_t \sum_{i \in \mathcal{S}_t} b_{i,t}$, which of course satisfies Assumption 4. Moreover, $\mathcal{C} = [1, 3, 7, 14]$ for Fig. 5 and $\mathcal{C} = [1, 3, 7]$ for the others.

As a first result, in Fig. 5, we illustrate the temporal behavior of the estimated accuracy of the FL algorithm, obtained for different values of the learning rate $\bar{\alpha}$, while fixing the accuracy to $\bar{G} = 0.8$ and the latency constraint $\bar{L} = 50$ ms. As we can notice from Fig. 5, the proposed data-driven method is able to strike the desired learning performance, while controlling the convergence rate. Similarly, in Fig. 6 (a), we illustrate the temporal behavior of the estimated accuracy, obtained for different values of the performance constraint $\bar{G}$, while fixing the convergence rate to $\bar{\alpha} = 0.2$ and the latency constraint $\bar{L} = 50$ ms. Then, in Fig. 6 (b) and 6 (c) we show the temporal evolution of the overall latency and the overall uplink transmission power consumption, respectively, corresponding to Fig. 6 for $\bar{G} = 0.8$ and $\bar{G} = 0.9$. As we can notice from Fig. 6, the proposed method keeps the latency around the requirement $\bar{L}$, while driving the accuracy on the target value $\bar{G}$ during the steady state phase. Interestingly, from Fig. 6 (c), we notice how a significant power saving can be achieved at steady-state if the accuracy requirement is not very strict (i.e., for $\bar{G} = 0.8$), thanks to the beneficial impact of the adaptive step-size in (49). Furthermore, the results empirically confirm the choice of the adopted surrogate functions, and

Fig. 6: (a) Accuracy vs t , (b) Latency vs t and (c) Power consumption vs t , for different $\bar{G}$.

Fig. 7: Accuracy vs t in a non-stationary scenario.

the consistence of Assumptions 3 and 4.

Finally, in Fig. 7 we illustrate the temporal behavior of the estimated accuracy in a non-stationary scenario, in order to highlight the adaptation capabilities of the proposed method. In particular, the MNIST dataset is divided into two sub-dataset of 5 classes each; then, the architecture is trained for the first 180 time-slot with one of the two sub-datasets and for the remaining time-slots with the other one (the last dense layer is obviously reduced to a 5 dimensional

output). Equivalently, this introduces a non-stationary behavior of the data distribution. Then, at time slot 401, we change the accuracy requirement from $\bar{G} = 0.8$ to $\bar{G} = 0.9$, introducing a further level of non-stationarity. The results are averaged over 10 independent simulations. From Fig. 7, we can notice that our dynamic strategy is able to react promptly to both changes in the data distribution and in the accuracy requirement, exhibiting powerful learning and adaptation capabilities in a fully data-driven fashion.

VI. CONCLUSIONS

In this paper, we have proposed a dynamic resource allocation strategy enabling adaptive federated learning at the wireless network edge. The strategy dynamically minimizes the power expenditure of the system, while guaranteeing target learning performance and latency constraints. The proposed method builds on stochastic Lyapunov optimization, which leads to low-complexity procedures for the resource allocation at each time slot, without requiring a-priori knowledge of wireless channel statistics. The approach is valid considering both a model-based approach, where performance metrics can be evaluated in closed-form, or exploiting a data-driven approach, where performance are necessarily estimated online from streaming data. The proposed strategy allows the exploration of a new trade-off of communication networks, including power expenditure, latency, and learning performance.

REFERENCES

- [1] S. Ahmadi, *5G NR: Architecture, Technology, Implementation, and Operation of 3GPP New Radio Standards*, Elsevier Science, 2019.
- [2] E. Calvanese Strinati et al., “6G: The next frontier: From holographic messaging to artificial intelligence using subterahertz and visible light communication,” *IEEE Vehicular Technology Magazine*, vol. 14, no. 3, pp. 42–50, September 2019.
- [3] A. Ndikumana et al., “Joint communication, computation, caching, and control in big data multi-access edge computing,” *IEEE Trans. on Mobile Computing*, vol. 19, no. 6, pp. 1359–1374, June 2020.
- [4] Sergio Barbarossa et al., “Communicating while computing: Distributed mobile cloud computing over 5G heterogeneous networks,” *IEEE Signal Processing Magazine*, vol. 31, no. 6, pp. 45–55, November 2014.
- [5] E. Calvanese Strinati et al., “Wireless environment as a service enabled by reconfigurable intelligent surfaces: The RISE-6G perspective,” *Proc. of EUCNC 6G Summit*, Porto, Portugal, 8-11 June 2021.
- [6] M. Polese et al., “Machine learning at the edge: A data-driven architecture with applications to 5g cellular networks,” *IEEE Trans. on Mobile Computing, Early Access*, 2020.
- [7] J. Park, S. Samarakoon, M. Bennis, and M. Debbah, “Wireless network intelligence at the edge,” *Proc. of the IEEE*, vol. 107, no. 11, pp. 2204–2239, November 2019.
- [8] J. Chen and X. Ran, “Deep learning with edge computing: A review,” *Proc. of the IEEE*, vol. 107, no. 8, pp. 1655–1674, 2019.

- [9] X. Wang et al., “Convergence of edge computing and deep learning: A comprehensive survey,” *IEEE Communications Surveys Tutorials*, vol. 22, no. 2, pp. 869–904, Secondquarter 2020.
- [10] M. Merluzzi, P. Di Lorenzo, and S. Barbarossa, “Wireless Edge Machine Learning: Resource Allocation and Trade-Offs,” *IEEE Access*, vol. 9, pp. 45377–45398, 2021.
- [11] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, “Federated learning: Challenges, methods, and future directions,” *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 50–60, May 2020.
- [12] J. Konečný, H. B. McMahan, D. Ramage, and P. Richtárik, “Federated optimization: Distributed machine learning for on-device intelligence,” *arXiv preprint arXiv:1610.02527*, 2016.
- [13] P. H. Jin, Q. Yuan, F. Iandola, and K. Keutzer, “How to scale distributed deep learning?,” *arXiv preprint arXiv:1611.04581*, 2016.
- [14] T. Li et al., “Federated optimization in heterogeneous networks,” *arXiv preprint arXiv:1812.06127*, 2018.
- [15] H. B. McMahan, Eider Moore, D. Ramage, S. Hampson, and B. A. Y. Arcas, “Communication-efficient learning of deep networks from decentralized data,” in *Proc. of AISTATS*, Lauderdale, FL, USA, 20-22 April 2017.
- [16] M. S. H. Abad, E. Ozfatura, D. Gunduz, and O. Ercetin, “Hierarchical federated learning across heterogeneous cellular networks,” in *Proc. of IEEE ICASSP*, Barcelona, Spain, 4-8 May 2020, pp. 8866–8870.
- [17] J. Konečný, B. McMahan, and D. Ramage, “Federated optimization: Distributed optimization beyond the datacenter,” *arXiv preprint arXiv:1511.03575*, 2015.
- [18] S. Samarakoon, M. Bennis, W. Saad, and M. Debbah, “Distributed federated learning for ultra-reliable low-latency vehicular communications,” *IEEE Trans. on Communications*, vol. 68, no. 2, pp. 1146–1159, Febraury 2020.
- [19] S. Ha, J. Zhang, O. Simeone, and J. Kang, “Coded federated computing in wireless networks with straggling devices and imperfect csi,” in *Proc. of IEEE ISIT*, Paris, France, 7-12 July 2019, pp. 2649–2653.
- [20] S. Wang et al., “Adaptive federated learning in resource constrained edge computing systems,” *IEEE Journal on Selected Areas in Communications*, vol. 37, no. 6, pp. 1205–1221, June 2019.
- [21] Q. Zeng, K. Du, Y. and Huang, and K. K. Leung, “Energy-efficient radio resource allocation for federated edge learning,” in *Proc. of IEEE ICC Workshops*, Dublin, Ireland, 7-11 June 2020, pp. 1–6.
- [22] M. Chen et al., “A joint learning and communications framework for federated learning over wireless networks,” *IEEE Trans. on Wireless Communications*, vol. 20, no. 1, pp. 269–283, January 2021.
- [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 *Proc. of IEEE INFOCOM*, Paris, France, 18-20 September 2019, pp. 1387–1395.
- [24] H. Yang, Z. Liu, T. Q. S. Quek, and H. V. Poor, “Scheduling policies for federated learning in wireless networks,” *IEEE Trans. on Communications*, vol. 68, no. 1, pp. 317–333, January 2019.
- [25] Z. Ma et al., “Adaptive batch size for federated learning in resource-constrained edge computing,” *IEEE Trans. on Mobile Computing, Early Access*, 2021.
- [26] L. U. Khan et al., “Federated learning for edge networks: Resource optimization and incentive mechanism,” *IEEE Communications Magazine*, vol. 58, no. 10, pp. 88–93, October 2020.
- [27] Z. Yang, M. Chen, W. Saad, C. S. Hong, and M. Shikh-Bahaei, “Energy efficient federated learning over wireless communication networks,” *IEEE Trans. on Wireless Communications*, vol. 20, no. 3, pp. 1935–1949, March 2021.
- [28] X. Wang, C. Wang, X. Li, V. C. M. Leung, and T. Taleb, “Federated deep reinforcement learning for internet of things with decentralized cooperative edge caching,” *IEEE IoT Journal*, vol. 7, no. 10, pp. 9441–9455, October 2020.
- [29] J. Mills and G. Hu, J. and Min, “Communication-efficient federated learning for wireless edge intelligence in iot,” *IEEE Internet of Things Journal*, vol. 7, no. 7, pp. 5986–5994, July 2020.
- [30] Y. Jiang et al., “Model pruning enables efficient federated learning on edge devices,” *arXiv preprint:1909.12326*, 2020.

- [31] X. Mo and J. Xu, “Energy-efficient federated edge learning with joint communication and computation design,” *Journal of Communications and Information Networks*, vol. 6, no. 2, pp. 110–124, June 2021.
- [32] 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. on Wireless Communications*, vol. 20, no. 1, pp. 453–467, Jan 2021.
- [33] M. Chen, H. V. Poor, W. Saad, and S. Cui, “Convergence time optimization for federated learning over wireless networks,” *IEEE Trans. on Wireless Communications*, vol. 20, no. 4, pp. 2457–2471, April 2021.
- [34] Bing Luo, Xiang Li, Shiqiang Wang, Jianwei Huang, and Leandros Tassiulas, “Cost-effective federated learning in mobile edge networks,” 2021.
- [35] A. H. Sayed, “Adaptation, learning, and optimization over networks,” *Foundations and Trends in Machine Learning*, vol. 7, pp. 311–801, 2014.
- [36] M. J. Neely, “Stochastic network optimization with application to communication and queueing systems,” *Morgan and Claypool*, 2010.
- [37] P. Di Lorenzo, C. Battiloro, and M. Merluzzi, “Dynamic resource optimization for adaptive federated learning at the wireless network edge,” in *Proc. of IEEE ICASSP*, Toronto, Canada, 6–11 June 2021, pp. 4910–4914.
- [38] N. Agarwal, A. T. Suresh, F. Yu, S. Kumar, and H. B. McMahan, “Cpsgd: Communication-efficient and differentially-private distributed sgd,” in *Proc. of NIPS*, Montréal, Canada, 3–8 December 2018, p. 7575–7586.
- [39] Wei-Ting Chang and Ravi Tandon, “Communication efficient federated learning over multiple access channels,” *arXiv preprint arXiv:2109.05411*, 2020.
- [40] S. P. Lipshitz, R. A. Wannamaker, and J. Vanderkooy, “Quantization and dither: A theoretical survey,” *Journal of the audio engineering society*, vol. 40, no. 5, pp. 355–375, May 1992.
- [41] R. A. Wannamaker, S. P. Lipshitz, J. Vanderkooy, and J. N. Wright, “A theory of nonsubtractive dither,” *IEEE Trans. on Signal Processing*, vol. 48, no. 2, pp. 499–516, February 2000.
- [42] L. Schuchman, “Dither signals and their effect on quantization noise,” *IEEE Trans. on Comm. Technology*, vol. 12, no. 4, pp. 162–165, December 1964.
- [43] T. Burd and R. Brodersen, “Processor design for portable systems,” *J. VLSI Sig. Process. Syst.*, vol. 13, no. 2–3, pp. 203–221, Aug. 1996.
- [44] M. Merluzzi, P. Di Lorenzo, S. Barbarossa, and V. Frascolla, “Dynamic computation offloading in multi-access edge computing via ultra-reliable and low-latency communications,” *IEEE Trans. on Signal and Information Processing over Networks*, vol. 6, pp. 342–356, 2020.
- [45] P. Brito, F. Fabião, and A. Staubyn, “Euler, lambert, and the lambert w-function today,” *The Mathematical Scientist*, vol. 33, January 2008.
- [46] G. R. MacCartney, S. Deng, S. Sun, and T. S. Rappaport, “Millimeter-wave human blockage at 73 ghz with a simple double knife-edge diffraction model and extension for directional antennas,” in *Proc. of IEEE VTC-Fall*, Montréal, 18–21 September 2016, pp. 1–6.
- [47] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, November 1998.
- [48] D. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *Proc. of ICLR*, 12 Banff, Canada, 14–16 April 2014.