Constrained Over-the-Air Model Updating for Wireless Online Federated Learning with Delayed Information

Juncheng Wang*, Yituo Liu*, Ben Liang[†], and Min Dong[‡]

*Department of Computer Science, Hong Kong Baptist University, Hong Kong [†]Department of Electrical and Computer Engineering, University of Toronto, Canada [‡]Department of Electrical, Computer and Software Engineering, Ontario Tech University, Canada

Abstract-We study online federated learning over a wireless network, where the central server updates an online global model sequence to minimize the time-varving loss of multiple local devices over time. The server updates the global model through over-the-air model-difference aggregation from the local devices over a noisy multiple-access fading channel. We consider the practical scenario where information on both the local loss functions and the channel states is delayed, and each local device is under a time-varying power constraint. We propose Constrained Over-the-air Model Updating with Delayed infOrmation (COMUDO), where a new lower-and-upper-bounded virtual queue is introduced to counter the delayed information and control the hard constraint violation. We show that its local model updates can be efficiently computed in closed-form expressions. Furthermore, through a new Lyapunov drift analysis, we show that COMUDO provides bounds on the dynamic regret, static regret, and hard constraint violation. Simulation results on image classification tasks under practical wireless network settings show substantial accuracy gain of COMUDO over state-of-the-art approaches, especially in the low-power region.

I. INTRODUCTION

With the explosion of data at wireless edge devices, alongside their increasing computational capacities and privacy concerns, *federated learning* (FL) [1] has been recognized as a vital framework to support edge intelligence. Migrating machine learning from the cloud to the edge, however, requires *communication-efficient* FL, to reduce the communication overhead caused by transmitting the machine-learning models between the edge devices and the parameter server [2].

Existing works on wireless FL adopt two broad approaches for communication efficiency. The first approach [3]-[7] uses conventional orthogonal multiple access in digital communication, with error-control coding and possible model quantization [8] and sparsification [9]. This approach can still lead to high communication overhead and latency especially for large networks. In contrast, the *over-the-air* (OTA) computation approach [10]-[22] uses a multiple access channel in analog communication, for multiple devices to concurrently transmit their local models. A noisy global model can be directly recovered from the superimposed analog signals at the server.

Due to the analog communication nature, OTA computation is highly sensitive to the channel state during signal transmission and the corresponding transmit power setting. In particular, [21] observed that the optimal model update in each training round of OTA FL should consider not only the steepness of the loss-function gradient, but also whether the

This work was supported in part by the Hong Kong Research Grants Council (RGC) Early Career Scheme (ECS) under grant 22200324 and by the Natural Sciences and Engineering Research Council (NSERC) of Canada. updated model can be efficiently transmitted from the edge devices to the server. However, all existing works on OTA FL with transmit power control [14]-[21] assume that the channel state is known in each model training round. In practice, the channel state in wireless communication fluctuates over time, and the computation time in each training iteration is often longer than the coherence period of the channel. As a result, the channel state for aggregating the updated local models at the end of each training round is different from the channel state used in computing the model update at the beginning of the training round. In other words, the available channel state information is *delayed*.

Furthermore, most existing OTA FL works assume the loss functions are fixed or independent and identically distributed (i.i.d.) over time [10]-[21], which does not allow streaming data that can possibly change arbitrarily in many online applications [23]-[25]. OTA FL with delayed loss function information was recently considered in [22], which was the first work to provide bounds on the static regret with respect to a fixed offline solution. However, [22] still required current *i.e.*, non-delayed, channel state information in each training round. Furthermore, the optimization constraint considered in [22] is fixed, so it cannot accommodate transmit power control, which requires a *time-varying* constraint due to the time-varying channel states in OTA FL.

In this work, we consider both delayed loss function information and delayed channel state information during model updates. Our goal is to develop an online algorithm that provides strong performance guarantees, in terms of dynamic and static regret bounds. To achieve this goal, we must address several challenges: 1) Analog OTA requires careful transmit signal control to counter the channel noise that propagates in the entire learning process. 2) The delays in the loss functions and the channel states require an online design to adapt to the system variations that are unknown a priori. 3) The learning performance and system constraints such as power are tightly coupled over time, necessitating a holistic analysis of the incurred performance loss and constraint violation.

Our contributions can be summarized as follows:

• We formulate an online OTA FL problem with timevarying loss functions, channel states, and power constraints. Both the loss function information and the channel state information are delayed, and as a result the constraint function information is also delayed. Our optimization formulation further accommodates hard constraints [26], which do not allow any compensated violation over time or devices and thus are stronger than the soft constraints in [18]-[22].

- We propose an effective algorithm named Constrained Over-the-air Model Updating with Delayed infOrmation (COMUDO) to solve this problem. COMUDO introduces a new lower-and-upper-bounded virtual queue, which eliminates the need for Slater's condition and enforces a minimum constraint penalty, to strictly control the power violation. The resulting closed-form local updates can adapt to unknown variations of both the local loss functions and channel states under individual power limits.
- We establish a connection between the bounds on the virtual queue and the hard constraint violation through a new Lyapunov drift analysis. We show that COMUDO provides *O*(*T*^{1+max{μ,ω}}) dynamic regret, *O*(*T*^{1+ω}) static regret, and *O*(*T^ν*) hard constraint violation for general convex loss functions. Here, *T* is the number of training rounds, and μ, ν, and ω respectively represent how fast the dynamic online benchmark, power constraints, and channel noise fluctuate over time.
- We experiment on canonical datasets over typical wireless settings to study the performance of COMUDO for both convex and non-convex loss functions. Our simulation results show that COMUDO significantly improves the learning performance over state-of-the-art benchmarks. The performance gain is more substantial in the lowpower region.

II. RELATED WORK

A. Communication-Efficient Error-Free Distributed Learning

Much of the machine learning literature on communicationefficient distributed learning, *e.g.*, [1], [8], [9], overlooks the possibility to further mitigate the communication cost through wireless transmission design. Prior works on distributed learning over wireless networks mainly focus on *error-free* digital communication (see [3] and references therein). For example, adaptive model aggregation [4], joint optimization [5], temporal model similarity [6], and model pruning [7] have been proposed to improve the communication efficiency. These works all adopt the conventional orthogonal multiple access to transmit the model parameters device by device to the server, resulting in possible high communication overhead and latency. In this work, we focus instead on OTA FL.

B. Over-the-Air Federated Learning

To mitigate the communication overhead and reduce the communication latency, OTA computation has been adopted in FL recently. Various solutions have been proposed in the literature, including model truncation [10], device selection [11], aggregation error minimization [12], and truncated channel inversion [13]. The above works, however, do not explicitly consider transmit power control at the local devices.

All existing works on OTA FL that consider transmit power control [14]-[21] assume the channel state for OTA computation remains *unchanged* after model training and the loss functions are *fixed* or *i.i.d.* over time. For example, joint transmit power control and receiver beamforming were considered in [16] with perfect channels and fixed losses. Joint uplink and downlink beamforming were considered in [17] with perfect channels and fixed losses. Long-term transmit power control was achieved by scaling the power factor using channel inversion [18]. Based on i.i.d. channels, the optimality gap was minimized through power allocation in [19]. Under long-term power constraints, a regularized channel inversion scheme was proposed to improve convergence in [20]. Online model update and analog aggregation were jointly considered in [21] with current information of the i.i.d. data and channel.

OTA FL with delayed loss function information was considered in [22]. The modified saddle-point-type algorithm provided $\mathcal{O}(\sqrt{T})$ static regret and $\mathcal{O}(T^{\frac{3}{4}})$ soft constraint violation. However, [22] still assumes non-delayed channel state information as in [14]-[21]. Furthermore, the interaction among devices is only through the optimization constraints, which eliminates the need to consider analog model aggregation and simplifies the bounding analysis. Additionally, [22] only allows *fixed* optimization constraints, which cannot accommodate power control in OTA FL, since the transmit power in OTA generally depends on the time-varying channel state. In contrast, our proposed online algorithm is designed to tolerate both delayed loss function information and delayed channel state information during model updates, while providing bounds on the dynamic regret, static regret, and hard constraint violation for time-varying power constraints.

C. Constrained Online Optimization

To handle the time-varying and delayed loss functions and power constraints, we borrow some techniques from Lyapunov optimization [27] and online convex optimization (OCO) [28]. Early OCO works focused on fixed constraints. For example, the online gradient descent algorithm in [29] achieved $O(\sqrt{T})$ static regret to a fixed offline decision. Dynamic regret to a time-varying decision sequence was analyzed in [29], [30]. Later works considered long-term constraints [25], [31]-[34]. These OCO algorithms, however, are *centralized*.

Almost all existing works on distributed OCO either assume fixed constraints or consider the *soft* constraint violation [35]-[39], which allows the violation to be compensated over time and over the local learners. For example, dynamic regret bound and soft constraint violation bound were provided in [38], [39]. In contrast, with the goal of limiting instantaneous constraint violation, [26], [40], [41] considered a stronger notion of hard constraint violation, which does not allow any compensated violation over time. However, the proposed online algorithms in [26], [40], [41] require centralized implementation.

Distributed OCO with hard constraint violation was more recently considered in [42]. Assuming *error-free* communication without any cost, the distributed saddle-point-type algorithm in [42] achieved $\mathcal{O}(\sqrt{T})$ static regret, and $\mathcal{O}(T^{\frac{3}{4}})$ hard constraint violation even when the constraints are fixed. In contrast, for OTA FL considered in our work, we must consider nonideal communication with channel noise and the need for power control. Furthermore, through a novel lower-and-upperbounded queue together with a new Lyapunov drift analysis, our proposed algorithm provides a *dynamic* regret bound and a hard constraint violation bound that smoothly approach to the optimal $\mathcal{O}(\sqrt{T})$ regret and $\mathcal{O}(1)$ violation, respectively, as the system variations diminish.

III. ONLINE OVER-THE-AIR FEDERATED LEARNING

A. OTA FL with Delayed Information

We consider an online FL system that operates over rounds t = 1, ..., T. There are a total of N local devices in the system coordinated by a central server. Each device n experiences a local loss function $f_t^n(\mathbf{x}) : \mathbb{R}^d \to \mathbb{R}$ at round t. At the *beginning* of each round t, each device n computes a local model $\mathbf{x}_t^n \in \mathbb{R}^d$. Only at the *end* of each round t, device n receives feedback information on the current local loss function $f_t^n(\mathbf{x})$ [22]. Therefore, when device n selects its local model \mathbf{x}_t^n at round t, it only has some delayed information — usually the gradient $\nabla f_{t-1}^n(\mathbf{x})$ — on the previous local loss function $f_t^n(\mathbf{x}) = f^n(\mathbf{x})$. This model includes the standard offline FL scenario $f_t^n(\mathbf{x}) = f^n(\mathbf{x}), \forall t$, considered in [10]-[20] as a special case.

Given the delayed loss function information, FL aims at generating a *global* model sequence $\{\mathbf{x}_t\}$ at the server, to minimize the accumulated losses of the N local devices:

$$\min_{\{\mathbf{x}_t\}} \quad \frac{1}{N} \sum_{t=1}^T \sum_{n=1}^N f_t^n(\mathbf{x}_t). \tag{1}$$

Each device *n* updates its local model \mathbf{x}_t^n via online local gradient descent $\mathbf{x}_t^n = \mathbf{x}_{t-1} - \alpha \nabla f_{t-1}^n(\mathbf{x}_{t-1})$, where $\alpha > 0$ is the step size. At the end of each round *t*, each device *n* sends its model difference $\mathbf{x}_t^n - \mathbf{x}_{t-1} = -\alpha \nabla f_{t-1}^n(\mathbf{x}_{t-1})$ (or local gradient) to the server, which then performs model-difference aggregation to update its global model $\mathbf{x}_t = \mathbf{x}_{t-1} + \frac{1}{N} \sum_{n=1}^{N} (\mathbf{x}_t^n - \mathbf{x}_{t-1})$.¹ We adopt the above model-difference aggregation in this work since it usually incurs less transmit power than directly aggregating the models [21].

We assume OTA computation for efficient global model update at the server. The wireless communication channel between the server and devices is modeled as a noisy multipleaccess fading channel. Let $\mathbf{h}_t^n = [h_t^n[1], \ldots, h_t^n[d]]^T \in \mathbb{C}^d$ denote the channel state information between the server and device n at round t, with $0 < h_{\text{LB}} \leq |h_t^n[i]|, \forall t, \forall n, \forall i$. In practical communication networks, instantaneous channel state information is usually unavailable [43]. Furthermore, the computation of \mathbf{x}_t^n can be lengthy and the channel state may have already changed when \mathbf{x}_t^n is ready for transmission. We therefore consider *delayed* channel state information at all the local devices, *i.e.*, when device n computes \mathbf{x}_t^n , it only has the previous channel state information \mathbf{h}_{t-1}^n . In Fig. 1, we illustrate online OTA FL with delays on *both* the local loss function information and the channel state information.

Wireless channels are naturally fading and noisy. Therefore, the devices need to properly design their transmit signals, so that the global model \mathbf{x}_t can be recovered at the server. Let $\hat{\mathbf{x}}_t$ be the recovered *noisy* global model. Due to the delay caused by computing \mathbf{x}_t^n , the analog signals that carry the information of \mathbf{x}_t^n are transmitted to the server at the *end* of each round t. Adopting the standard OTA technique [10], [13]-





Fig. 1. An illustration of OTA FL with delayed information. The solid arrows indicate signal transmission, and the dashed arrows indicate information flow. At the beginning of round t, when device n updates its local model \mathbf{x}_t^n , it only has the delayed local loss function information $f_{t-1}^n(\mathbf{x})$ and the delayed local channel state information \mathbf{h}_{t-1}^n . At the end of round t, only after updating \mathbf{x}_t^n , the information of $f_t^n(\mathbf{x})$ and \mathbf{h}_t^n becomes available to device n.

[15], [18], [21], we set the analog signal vector $\mathbf{s}_t^n \in \mathbb{C}^d$ to be transmitted by device n as

$$\mathbf{s}_t^n = \mathbf{b}_t^n \circ \left(\mathbf{x}_t^n - \hat{\mathbf{x}}_{t-1} \right)$$
(2)

where $\mathbf{b}_t^n \triangleq \left[\frac{\lambda h_t^n[1]^*}{|h_t^n[1]|^2}, \dots, \frac{\lambda h_t^n[d]^*}{|h_t^n[d]|^2}\right]^H \in \mathbb{C}^d$ is the channel inversion vector of \mathbf{h}_t^n , with $\lambda > 0$ being a scaling factor and $\mathbf{x} \circ \mathbf{y}$ denotes the entry-wise product.

All devices concurrently transmit each entry of $\{\mathbf{s}_t^n\}$ over an orthogonal channel. The analog signals $\{\mathbf{s}_t^n\}$ are summed over the wireless multiple-access channel. The server receives the aggregated signal vector $\mathbf{y}_t = \sum_{n=1}^N \mathbf{h}_t^n \circ \mathbf{s}_t^n + \mathbf{z}_t$, where $\mathbf{z}_t \in \mathbb{C}^d$ is the additive channel noise with $z_{\text{UB}} \geq ||\mathbf{z}_t||_F, \forall t$. The server post-processes \mathbf{y}_t to recover a noisy global model

$$\hat{\mathbf{x}}_t = \hat{\mathbf{x}}_{t-1} + \frac{\Re\{\mathbf{y}_t\}}{N\lambda} = \mathbf{x}_t + \mathbf{n}_t$$
(3)

where $\Re\{\mathbf{x}\}\$ returns the real values of vector \mathbf{x} and $\mathbf{n}_t = \frac{\Re\{\mathbf{z}_t\}}{N\lambda} \in \mathbb{R}^d$ is a *residual* noise vector. The server then broadcasts $\hat{\mathbf{x}}_t$ back to the devices.

Remark 1. We can extend our model in multiple fronts without major technical changes. For example, we can use multi-antenna receiver beamforming to reduce the residual noise \mathbf{n}_t [16]. We can consider noisy downlink channels and incorporate the downlink residual noise into \mathbf{n}_t [17]. We can use the imaginary part of \mathbf{s}_t^n to reduce the number of channels [18]. Some lossy methods, *e.g.*, truncated channel inversion [10], model sparsification [18], and regularized channel inversion [20] will incur additional noise that can be incorporated as part of \mathbf{n}_t in (3), and our derived performance bounds later in Section V still hold.

B. Online Optimization Problem Formulation

With both delayed loss function and channel state information, our optimization objective is to minimize the accumulated loss incurred by the sequence of noisy global models at the server, under time-varying power constraints on the local models at each device. Our formulated online OTA FL problem is

$$\mathbf{P}: \min_{\{\mathbf{x}_t^n \in \mathcal{X}\}} \sum_{t=1}^T \sum_{n=1}^N f_t^n(\hat{\mathbf{x}}_t)$$

s.t. $g_t^n(\mathbf{x}_t^n) \le 0, \quad \forall t, \forall n$ (4)

where

$$g_t^n(\mathbf{x}) \triangleq \|\mathbf{b}_t^n \circ (\mathbf{x} - \hat{\mathbf{x}}_{t-1})\|^2 - P^n, \quad \forall t, \forall n$$
 (5)

with P^n being the power limit of device n, and $\mathcal{X} = \{\mathbf{x} : -x_{\text{UB}}\mathbf{1} \leq \mathbf{x} \leq x_{\text{UB}}\mathbf{1}\}$ is a feasible set with $x_{\text{UB}} > 0$ being the model value upper bound. Note that since \mathbf{b}_t^n in (5) is a function of \mathbf{h}_t^n , the constraint function is also time-varying, and its information is also delayed when \mathbf{x}_t^n is computed.

Since the local loss function $f_t^n(\mathbf{x})$, channel state \mathbf{h}_t^n , and power constraint $g_t^n(\mathbf{x})$ are unknown a priori and can possibly change arbitrarily over time with unknown distributions, obtaining an optimal solution to **P** is impossible, since it needs central calculation using all the information of the local loss functions $\{f_t^n(\mathbf{x})\}$, channel states $\{\mathbf{h}_t^n\}$, and power constraints $\{g_t^n(\mathbf{x})\}$ in hindsight.

In this work, our goal is to compute locally a model sequence $\{\mathbf{x}_t^n\}$ to **P** that can provide sublinear *dynamic regret*

$$\operatorname{Reg}_{d}(T) \triangleq \frac{1}{N} \sum_{t=1}^{T} \sum_{n=1}^{N} \left[f_{t}^{n}(\hat{\mathbf{x}}_{t}) - f_{t}^{n}(\mathbf{x}_{t}^{\circ}) \right]$$
(6)

where $\mathbf{x}_{t}^{\circ} \in \arg\min_{\mathbf{x}\in\mathcal{X}} \left\{ \frac{1}{N} \sum_{n=1}^{N} f_{t}^{n}(\mathbf{x}) |\| \mathbf{b}_{t}^{n} \circ (\mathbf{x} - \hat{\mathbf{x}}_{t}) \|^{2} \leq P^{n}, \forall n \right\}$ is the centralized *dynamic online* benchmark computed using all the current information. When we compare the online model sequence with the centralized *fixed offline* benchmark $\mathbf{x}^{\circ} \in \arg\min_{\mathbf{x}\in\mathcal{X}} \{ \frac{1}{N} \sum_{t=1}^{T} \sum_{n=1}^{N} f_{t}^{n}(\mathbf{x}) |\| \mathbf{b}_{t}^{n} \circ (\mathbf{x} - \hat{\mathbf{x}}_{t}) \|^{2} \leq P^{n}, \forall t, \forall n \}$, the resulting regret

$$\operatorname{Reg}_{s}(T) = \frac{1}{N} \sum_{t=1}^{T} \sum_{n=1}^{N} [f_{t}^{n}(\hat{\mathbf{x}}_{t}) - f_{t}^{n}(\mathbf{x}^{\circ})]$$
(7)

is commonly referred to as the *static regret*. In this work, we provide performance guarantees on both the dynamic regret and the static regret.

Furthermore, since the information of \mathbf{h}_t^n and $g_t^n(\mathbf{x})$ is only available to device n at the end of round t, it is impossible to compute a local model \mathbf{x}_t^n at the beginning of round t, such that the power constraints (4) are strictly satisfied. We therefore need to allow some violations on (4) and introduce a *hard constraint violation* measure to quantify the amount of *instantaneous* violations of the constraints

$$\operatorname{Vio}_{h}(T) \triangleq \frac{1}{N} \sum_{t=1}^{T} \sum_{n=1}^{N} \left[g_{t}^{n}(\mathbf{x}_{t}^{n}) \right]_{+}$$

$$\tag{8}$$

where $[x]_{+} = \max\{x, 0\}$. The hard constraint violation $\operatorname{Vio}_{h}(T)$ in (8), also referred to as cumulative constraint violation in [26], [40]-[42], does not allow the violation at one time or one device to be compensated by any other time or device. It is a *stronger* notion of the standard soft constraint violation, *i.e.*, $\frac{1}{N} \sum_{t=1}^{T} \sum_{n=1}^{N} g_t^n(\mathbf{x}_t^n)$, which quantifies the amount of *compensated* violations over time and devices [21], [22], [35]-[39].

A constrained online algorithm is desired that simultaneously achieves $\lim_{T\to\infty} \frac{\operatorname{Reg}_d(T)}{T} = 0$, $\lim_{T\to\infty} \frac{\operatorname{Reg}_s(T)}{T} = 0$, and $\lim_{T\to\infty} \frac{\operatorname{Vio}_h(T)}{T} = 0$, which implies that the online model sequence computed locally using the delayed information is asymptotically no worse than either the dynamic online benchmark or the fixed offline benchmark, and at the same time, the individual power constraints (4) are satisfied in the time-averaged manner.

IV. CONSTRAINED OVER-THE-AIR MODEL UPDATING WITH DELAYED INFORMATION (COMUDO)

We present the COMUDO algorithm for solving **P**. CO-MUDO introduces a new lower-and-upper-bounded virtual queue, which will be shown in Section V to provide improved hard constraint violation bounds. Furthermore, COMUDO yields closed-form local model updates that are tolerant to both delayed loss function information and delayed channel state information.

A. Preliminaries

We show in the following lemma that **P** satisfies several properties: 1) the residual noise \mathbf{n}_t defined below (3) is bounded; 2) the feasible set \mathcal{X} defined below **P** is bounded; and 3) the power constraint function $g_t^n(\mathbf{x})$ in (5) is bounded. These properties will be useful for our performance analysis. The proof is omitted due to the page limit.

Lemma 1. P has the following properties:

$$\|\mathbf{n}_t\| \le E, \quad \forall t, \tag{9}$$

$$\|\mathbf{x} - \mathbf{y}\| \le R, \quad \forall \mathbf{x}, \mathbf{y} \in \mathcal{X},$$
 (10)

$$|g_t^n(\mathbf{x})| \le G, \quad \forall \mathbf{x} \in \mathcal{X}, \forall t, \forall n, \tag{11}$$

where $E = \frac{z_{\text{UB}}}{N\lambda}, R = 2\sqrt{d}x_{\text{UB}}$, and $G = \max_n \{\max\{P^n, \frac{d(R+E)^2\lambda^2}{h_{\text{LB}}^2} - P^n\}\}.$

B. Lower-and-Upper-Bounded Virtual Queue

COMUDO maintains a virtual queue Q_t^n , which is initialized as $Q_1^n = V$ and updated at the end of each round t > 1, to control the power violation at each local device n:

$$Q_t^n = \max\left\{ (1 - \eta) Q_{t-1}^n + [\gamma g_t^n(\mathbf{x}_t^n)]_+, V \right\}$$
(12)

where $V \in (0, \frac{\gamma G}{\eta})$ is any minimum virtual queue length, $\eta \in (0, 1)$ is a penalty factor on Q_{t-1}^n , and $\gamma > 0$ is a weight on $g_t^n(\mathbf{x}_t^n)$. Note that the information of $g_t^n(\mathbf{x})$ is available at the end of round t and is used to update the virtual queue Q_t^n . This is in contrast to the model update of \mathbf{x}_t^n at the beginning of round t, where only the delayed information of $g_{t-1}^n(\mathbf{x})$ is available.

Remark 2. Through a virtual queue lower bound analysis, the centralized online algorithm in [41] shows matched or improved performance bounds over the saddle-point-type algorithms [40], [42]. In this work, we introduce a novel lower-and-upper-bounded virtual queue in (12), together with a new Lyapunov drift analysis in Section V-B, to bound the hard power constraint violation for OTA FL.

In the following lemma, we show that without requiring Slater's condition, the additional penalty factor η in (12) can be seen as a *virtual* Slater's constant, leading to a virtual queue upper bound that is *not* inversely scaled by the *actual* Slater's constant P^n . Furthermore, V in (12) provides a non-zero virtual queue *lower* bound, which prevents COMUDO from incurring an overly large constraint violation.

Algorithm 1 COMUDO at Local Device

- 1: Initialization: $\hat{\mathbf{x}}_0 = \mathbf{0}$, $\hat{\mathbf{x}}_1 \in \mathcal{X}$ and $Q_1^n = V$.
- 2: At the beginning of t, use $\nabla f_{t-1}^n(\hat{\mathbf{x}}_{t-1})$ and \mathbf{h}_{t-1}^n , do:
- 3: **if** $\tilde{g}_{t-1}^n(-\alpha \nabla f_{t-1}^n(\hat{\mathbf{x}}_{t-1})) > 0$ then
- 4: Update local model \mathbf{x}_t^n via (15).
- 5: else Update local model \mathbf{x}_t^n via (16).
- 6: At the end of t, receive $\nabla f_t^n(\hat{\mathbf{x}}_t)$ and \mathbf{h}_t^n , do:
- 7: Transmit analog signal \mathbf{s}_t^n in (2) to the server.
- 8: Update local virtual queue Q_t^n via (12).

Lemma 2. The virtual queue updated via (12) is bounded by

$$V \le Q_t^n \le \frac{\gamma G}{\eta}, \quad \forall t, \forall n.$$
(13)

Proof: From (12), we readily have $Q_t^n \geq V, \forall t$. We now prove $Q_t^n \leq \frac{\gamma G}{\eta}, \forall t$ by induction. From initialization, $Q_1^n = V < \frac{\gamma G}{\eta}$. Then, suppose $Q_{t-1}^n \leq \frac{\gamma G}{\eta}$ at round t-1 for some t > 1. We now prove that $Q_t^n \leq \frac{\gamma G}{\eta}$ at round t. From (12), the triangle inequality, and the bound on $|g_t^n(\mathbf{x})|$ in (11), Q_t^n is upper bounded by $Q_t^n \leq (1-\eta)Q_{t-1}^n + [\gamma g_t^n(\mathbf{x}_t^n)]_+ \leq (1-\eta)\frac{\gamma G}{\eta} + \gamma G = \frac{\gamma G}{\eta}$, which proves (13).

Lemma 2 shows that the virtual queue upper bound $\frac{\gamma G}{\eta}$ is reversely proportional to the virtual Slater's constant η , instead of the actual Slater's constant P^n for the constraint function (5). This virtual queue upper bound, however, can no longer be directly used to bound the hard constraint violation Vio_h(T). In Section V-B, we propose a new Lyapunov drift analysis, which uses both the lower and upper bounds of the virtual queue in (13) to reconnect Q_t^n and Vio_h(T).

C. COMUDO Algorithm

We decompose **P** into a set of per-slot per-device optimization problems $\{\mathbf{P}_t^n\}$, given by

$$\mathbf{P}_{t}^{n}: \min_{\mathbf{x}\in\mathcal{X}} \langle \nabla f_{t-1}^{n}(\hat{\mathbf{x}}_{t-1}), \mathbf{x} - \hat{\mathbf{x}}_{t-1} \rangle + \frac{1}{2\alpha} \|\mathbf{x} - \hat{\mathbf{x}}_{t-1}\|^{2} + Q_{t-1}^{n} [\gamma \tilde{g}_{t-1}^{n}(\mathbf{x})]_{+}.$$
(14)

where $\langle \mathbf{x}, \mathbf{y} \rangle$ denotes the inner product operation and $\tilde{g}_{t-1}^n(\mathbf{x}) \triangleq \|\mathbf{b}_{t-1}^n \circ (\mathbf{x} - \hat{\mathbf{x}}_{t-1})\|^2 - P^{n,2}$. Note that \mathbf{P}_t^n uses the delayed local gradient $\nabla f_{t-1}^n(\hat{\mathbf{x}}_{t-1})$ and the delayed channel state \mathbf{h}_{t-1}^n . The local gradient $\nabla f_{t-1}^n(\hat{\mathbf{x}}_{t-1})$ at device *n* is taken on $\hat{\mathbf{x}}_{t-1}$ to minimize the accumulated global loss via noisy local gradient descent. Furthermore, each power constraint (4) has been converted to minimizing $Q_{t-1}^n[\gamma \tilde{g}_{t-1}^n(\mathbf{x})]_+$ as part of the objective function in \mathbf{P}_t^n for controlling the hard constraint violation.

We now show that the optimal solution to \mathbf{P}_t^n is in a closed form. The objective function of \mathbf{P}_t^n has a gradient $\nabla f_{t-1}^n(\hat{\mathbf{x}}_{t-1}) + \frac{1}{\alpha}(\mathbf{x} - \hat{\mathbf{x}}_{t-1}) + \gamma Q_{t-1}^n \theta_{t-1}^n \circ (\mathbf{x} - \hat{\mathbf{x}}_{t-1}) \mathbb{1}\{\tilde{g}_{t-1}^n(\mathbf{x}) > 0\}$, where $\theta_{t-1}^n \circ (\mathbf{x} - \hat{\mathbf{x}}_{t-1})$ is the gradient of constraint function $\tilde{g}_{t-1}^n(\mathbf{x})$, with the *i*-th entry of $\theta_{t-1}^n \in \mathbb{R}^d$ being $\theta_{t-1}^n[i] = \frac{2\lambda^2}{|h_{t-1}^n[i]|^2}$, and $\mathbb{1}\{x > 0\}$ is an indicator function. We set the above gradient to zeros for solving \mathbf{x} . Then, to satisfy the fixed constraints in the

Algorithm 2 COMUDO at Central Server

- 1: Initialization: $\alpha > 0$, $\eta \in (0, 1)$, $\gamma > 0$, and $V \in (0, \frac{\gamma G}{\eta})$. 2: At the end of t, do:
- 3: Receive analog signal \mathbf{y}_t over the air.
- 4: Recover a noisy global model $\hat{\mathbf{x}}_t$ in (3).
- 5: Broadcast $\hat{\mathbf{x}}_t$ to all the devices.

feasible set \mathcal{X} , we project the solution onto \mathcal{X} . The resulting optimal solution to \mathbf{P}_t^n depends on the power constraint violation $\tilde{g}_{t-1}^n(-\alpha \nabla f_{t-1}^n(\hat{\mathbf{x}}_{t-1}))$ caused by transmitting the local gradient. Specifically, if $\tilde{g}_{t-1}^n(-\alpha \nabla f_{t-1}^n(\hat{\mathbf{x}}_{t-1})) > 0$, *i.e.*, the amount of power required to transmit $-\alpha \nabla f_{t-1}^n(\hat{\mathbf{x}}_{t-1})$ exceeds the power limit P^n , we update \mathbf{x}_t^n via

$$\mathbf{x}_{t}^{n} = \left[\hat{\mathbf{x}}_{t-1} - \frac{\alpha}{\mathbf{1} + \alpha \gamma Q_{t-1}^{n} \boldsymbol{\theta}_{t-1}^{n}} \circ \nabla f_{t-1}^{n} (\hat{\mathbf{x}}_{t-1})\right]_{-x_{\text{UB}}\mathbf{1}}^{x_{\text{UB}}\mathbf{1}} (15)$$

where $[\mathbf{x}]_{\mathbf{y}}^{\mathbf{z}} = \min\{\mathbf{z}, \max\{\mathbf{x}, \mathbf{y}\}\}$. Otherwise, we update \mathbf{x}_{t}^{n} via projected local gradient descent

$$\mathbf{x}_t^n = \left[\hat{\mathbf{x}}_{t-1} - \alpha \nabla f_{t-1}^n(\hat{\mathbf{x}}_{t-1})\right]_{-x_{\text{UB}}\mathbf{1}}^{x_{\text{UB}}\mathbf{1}}.$$
 (16)

Note that $\nabla f_{t-1}^n(\hat{\mathbf{x}}_{t-1})$ is entry-wise scaled by a factor $\frac{\alpha}{1+\alpha\gamma Q_{t-1}^n\theta_{t-1}^n[i]}$ in (15) that is determined by the ratio of each channel strength $|h_{t-1}^n[i]|^2$ and the virtual queue length Q_{t-1}^n . When the channels are strong and the virtual queue, which measures the hard power constraint violation, is small, *i.e.*, $\frac{Q_{t-1}^n}{|h_{t-1}^n[i]|^2}$ is close to 0, the local model update (15) becomes (16) to greedily minimize the loss. When the virtual queue length dominates the channel strength, (15) becomes $\hat{\mathbf{x}}_{t-1}$ to save the transmit power. Therefore, the update of \mathbf{x}_t^n is both computation- and communication-aware.

Algorithms 1 and 2 summarize the COMUDO algorithm for the local device and the central server, respectively. In Section V-C, we will further discuss how to choose the algorithm parameters, α , η , γ , and V for COMUDO to yield the best convergence rates on $\text{Reg}_{d}(T)$, $\text{Reg}_{s}(T)$, and $\text{Vio}_{h}(T)$.

Remark 3. The computational complexity of COMUDO is mainly determined by the solutions to \mathbf{P}_t^n in (15) or (16). Note that both solutions have *closed-form* expressions, and contain a single local gradient computation $\nabla f_{t-1}^n(\hat{\mathbf{x}}_{t-1})$. This makes COMUDO highly efficient, with a computational complexity comparable to the standard FL algorithm.

V. PERFORMANCE BOUNDS OF COMUDO

In this section, we provide upper bounds on the dynamic regret (6), static regret (7), and hard constraint violation (8) for COMUDO. We take into account the joint impact of the lower-and-upper-bounded virtual queue, residual noise, and information delay on the performance of COMUDO.

To facilitate our analysis, we assume the loss functions are convex as in [22]. We show numerically in Section VI that COMUDO also performs well for non-convex loss functions.

Assumption 1. Function $f_t^n(\mathbf{x})$ is convex and its gradient is upper bounded, *i.e.*, there exists some D > 0, *s.t.*,

$$\|\nabla f_t^n(\mathbf{x})\| \le D, \quad \forall \mathbf{x} \in \mathbb{R}^d, \forall t, \forall n.$$
(17)

²We replace $\hat{\mathbf{x}}_{t-2}$ in the definition of $g_{t-1}^n(\mathbf{x})$ (5) with the more up-to-date information of $\hat{\mathbf{x}}_{t-1}$ available at the beginning of round *t*.

A. Bounding the Regret

The following lemma provides a per-slot per-device performance upper bound for COMUDO.

Lemma 3. The performance of each device n yielded by COMUDO is bounded for each round t by

$$\begin{aligned} & [f_{t-1}^{n}(\hat{\mathbf{x}}_{t-1}) - f_{t-1}^{n}(\mathbf{x}_{t-1}^{\circ})] + Q_{t-1}^{n}[\gamma \tilde{g}_{t-1}^{n}(\mathbf{x}_{t}^{n})]_{+} \\ & \leq \frac{R}{\alpha} \|\mathbf{x}_{t-1}^{\circ} - \mathbf{x}_{t}^{\circ}\| + \frac{2R}{\alpha} \|\mathbf{n}_{t}\| + \frac{1}{2\alpha} \|\mathbf{n}_{t}\|^{2} \\ & + \frac{\alpha D^{2}}{2} + \frac{1}{2\alpha} (\phi_{t} + \psi_{t}^{n}) \end{aligned}$$
(18)

where we define $\phi_t \triangleq \|\mathbf{x}_{t-1}^\circ - \hat{\mathbf{x}}_{t-1}\|^2 - \|\mathbf{x}_t^\circ - \hat{\mathbf{x}}_t\|^2$ and $\psi_t^n \triangleq \|\mathbf{x}_{t-1}^\circ - \mathbf{x}_t\|^2 - \|\mathbf{x}_{t-1}^\circ - \mathbf{x}_t^n\|^2$.

Proof: \mathbf{P}_t^n is $\frac{1}{\alpha}$ -strongly convex over \mathcal{X} , with \mathbf{x}_t^n being its optimal solution. Applying Lemma 2.8 in [28] (optimality condition of strongly convex function) to \mathbf{P}_t^n , we have

$$\langle \nabla f_{t-1}^{n}(\hat{\mathbf{x}}_{t-1}), \mathbf{x}_{t}^{n} - \hat{\mathbf{x}}_{t-1} \rangle + \frac{1}{2\alpha} \| \mathbf{x}_{t}^{n} - \hat{\mathbf{x}}_{t-1} \|^{2} + Q_{t-1}^{n} [\gamma \tilde{g}_{t-1}^{n}(\mathbf{x}_{t}^{n})]_{+} \leq \langle \nabla f_{t-1}^{n}(\hat{\mathbf{x}}_{t-1}), \mathbf{x}_{t-1}^{\circ} - \hat{\mathbf{x}}_{t-1} \rangle + \frac{1}{2\alpha} \| \mathbf{x}_{t-1}^{\circ} - \hat{\mathbf{x}}_{t-1} \|^{2} + Q_{t-1}^{n} [\gamma \tilde{g}_{t-1}^{n}(\mathbf{x}_{t-1}^{\circ})]_{+} - \frac{1}{2\alpha} \| \mathbf{x}_{t-1}^{\circ} - \mathbf{x}_{t}^{n} \|^{2}.$$
(19)

We now bound the right-hand side (RHS) of (19). Since $f_{t-1}^n(\mathbf{x})$ is convex, we have $\langle \nabla f_{t-1}^n(\hat{\mathbf{x}}_{t-1}), \mathbf{x}_{t-1}^\circ - \hat{\mathbf{x}}_{t-1} \rangle \leq f_{t-1}^n(\mathbf{x}_{t-1}) - f_{t-1}^n(\hat{\mathbf{x}}_{t-1})$. From the definition of the dynamic benchmark, we directly have $Q_{t-1}^n[\gamma \tilde{g}_{t-1}^n(\mathbf{x}_{t-1}^\circ)]_+ = 0$. For the rest two terms on the RHS of (19), we have

$$\begin{aligned} \|\mathbf{x}_{t-1}^{\circ} - \hat{\mathbf{x}}_{t-1}\|^{2} - \|\mathbf{x}_{t-1}^{\circ} - \mathbf{x}_{t}^{n}\|^{2} \\ &= \|\mathbf{x}_{t-1}^{\circ} - \hat{\mathbf{x}}_{t-1}\|^{2} - \|(\mathbf{x}_{t-1}^{\circ} - \mathbf{x}_{t}^{\circ}) + (\mathbf{x}_{t}^{\circ} - \hat{\mathbf{x}}_{t})\|^{2} \\ &+ (\|\mathbf{x}_{t-1}^{\circ} - \hat{\mathbf{x}}_{t}\|^{2} - \|\mathbf{x}_{t-1}^{\circ} - \mathbf{x}_{t}^{n}\|^{2}) \\ &\stackrel{(a)}{\leq} \phi_{t} - \|\mathbf{x}_{t-1}^{\circ} - \mathbf{x}_{t}^{\circ}\|^{2} + 2\|\mathbf{x}_{t}^{\circ} - \hat{\mathbf{x}}_{t}\|\|\mathbf{x}_{t-1}^{\circ} - \mathbf{x}_{t}^{\circ}\| \\ &+ (\|\mathbf{x}_{t-1}^{\circ} - \hat{\mathbf{x}}_{t}\|^{2} - \|\mathbf{x}_{t-1}^{\circ} - \mathbf{x}_{t}^{n}\|^{2}) \\ &\stackrel{(b)}{\leq} \phi_{t} + 2R\|\mathbf{x}_{t-1}^{\circ} - \mathbf{x}_{t}^{\circ}\| + 2R\|\mathbf{n}_{t}\| \\ &+ (\|(\mathbf{x}_{t-1}^{\circ} - \mathbf{x}_{t}) - \mathbf{n}_{t}\|^{2} - \|\mathbf{x}_{t-1}^{\circ} - \mathbf{x}_{t}^{n}\|^{2}) \\ &\stackrel{(c)}{\leq} \phi_{t} + 2R\|\mathbf{x}_{t-1}^{\circ} - \mathbf{x}_{t}^{\circ}\| + \psi_{t}^{n} + \|\mathbf{n}_{t}\|^{2} + 4R\|\mathbf{n}_{t}\| \end{aligned} \tag{20}$$

where (a) is because of $\|\mathbf{x} + \mathbf{y}\|^2 \ge \|\mathbf{x}\|^2 + \|\mathbf{y}\|^2 - 2\|\mathbf{x}\|\|\mathbf{y}\|$; (b) is due to $\hat{\mathbf{x}}_t$ being defined in (3) and \mathcal{X} being bounded in (10), such that $\|\mathbf{x}_t^\circ - (\mathbf{x}_t + \mathbf{n}_t)\|\|\mathbf{x}_{t-1}^\circ - \mathbf{x}_t^\circ\| \le (\|\mathbf{x}_t^\circ - \mathbf{x}_t\| + \|\mathbf{n}_t\|)\|\mathbf{x}_{t-1}^\circ - \mathbf{x}_t^\circ\| \le R(\|\mathbf{x}_{t-1}^\circ - \mathbf{x}_t^\circ\| + \|\mathbf{n}_t\|)$; and (c) follows from inequality $\|\mathbf{x} - \mathbf{y}\|^2 \le \|\mathbf{x}\|^2 + \|\mathbf{y}\|^2 + 2\|\mathbf{x}\|\|\mathbf{y}\|$ and the bound on \mathcal{X} in (10).

Substituting the above three inequalities into the RHS of (19), rearranging terms, and noting that $-\langle \nabla f_{t-1}^n(\hat{\mathbf{x}}_{t-1}), \mathbf{x}_t^n - \hat{\mathbf{x}}_{t-1} \rangle - \frac{1}{2\alpha} \|\mathbf{x}_t^n - \hat{\mathbf{x}}_{t-1}\|^2 \leq \frac{\alpha \|\nabla f_{t-1}^n(\hat{\mathbf{x}}_{t-1})\|^2}{2} \leq \frac{\alpha D^2}{2}$ by completing the square, we have proved (18).

Using the results in Lemma 3, we are now ready to derive a dynamic regret bound for COMUDO. Note that the virtual queue length is always strictly positive and the hard constraint violation is non-negative. Therefore, their product $Q_{t-1}^n[\gamma \tilde{g}_{t-1}^n(\mathbf{x}_t^n)]_+$ in (18) remains non-negative. Different

from the soft-constrained algorithm analysis in [22], [25], [31]-[39], this unique property allows us to bound the dynamic regret without worrying about the constraint violation.

Theorem 1. The dynamic regret of COMUDO has the following upper bound:

$$\operatorname{Reg}_{d}(T) \leq \frac{R}{\alpha} \sum_{t=2}^{T} \|\mathbf{x}_{t-1}^{\circ} - \mathbf{x}_{t}^{\circ}\| + \frac{2R}{\alpha} \sum_{t=2}^{T} \|\mathbf{n}_{t}\| + \frac{1}{2\alpha} \sum_{t=2}^{T} \|\mathbf{n}_{t}\|^{2} + \frac{\alpha D^{2}T}{2} + \frac{R^{2}}{2\alpha} + D(R+E).$$
(21)

Proof: Summing (18) over n = 1, ..., N and t = 2, ..., T, noting that $Q_{t-1}^n[\gamma \tilde{g}_{t-1}^n(\mathbf{x}_t^n)]_+ \ge 0$, and dividing both sides by N, we have

$$\frac{1}{N} \sum_{t=1}^{T-1} \sum_{n=1}^{N} \left[f_t^n(\hat{\mathbf{x}}_t) - f_t^n(\mathbf{x}_t^\circ) \right] \\
\leq \frac{R}{\alpha} \sum_{t=2}^{T} \|\mathbf{x}_{t-1}^\circ - \mathbf{x}_t^\circ\| + \frac{2R}{\alpha} \sum_{t=2}^{T} \|\mathbf{n}_t\| + \frac{1}{2\alpha} \sum_{t=2}^{T} \|\mathbf{n}_t\|^2 \\
+ \frac{\alpha D^2 T}{2} + \frac{1}{2\alpha} \sum_{t=2}^{T} \phi_t + \frac{1}{2\alpha N} \sum_{t=2}^{T} \sum_{n=1}^{N} \psi_t^n.$$
(22)

From $\hat{\mathbf{x}}_1 \in \mathcal{X}$ by initialization and the bound on \mathcal{X} in (10), we can show that $\sum_{t=2}^{T} \phi_t = \|\mathbf{x}_1^{\circ} - \hat{\mathbf{x}}_1\|^2 - \|\mathbf{x}_T^{\circ} - \hat{\mathbf{x}}_T\|^2 \leq \|\mathbf{x}_1^{\circ} - \hat{\mathbf{x}}_1\|^2 \leq R^2$. Applying the property of separate convexity, we have $\sum_{n=1}^{N} \psi_t^n = \sum_{n=1}^{N} (\|\mathbf{x}_{t-1}^{\circ} - \frac{1}{N} \sum_{m=1}^{N} \mathbf{x}_t^m\|^2 - \|\mathbf{x}_{t-1}^{\circ} - \mathbf{x}_t^m\|^2) \leq \sum_{n=1}^{N} (\frac{1}{N} \sum_{m=1}^{N} \|\mathbf{x}_{t-1}^{\circ} - \mathbf{x}_t^m\|^2 - \|\mathbf{x}_{t-1}^{\circ} - \mathbf{x}_t^n\|) = 0$. Substituting the two inequalities into the RHS of (22), and from $f_T^n(\hat{\mathbf{x}}_T) - f_T^n(\mathbf{x}_T^{\circ}) \leq \langle \nabla f_T^n(\hat{\mathbf{x}}_t), \hat{\mathbf{x}}_T - \mathbf{x}_T^{\circ} \rangle \leq \|\nabla f_T^n(\hat{\mathbf{x}}_T)\|(\|\mathbf{x}_T - \mathbf{x}_T^{\circ}\| + \|\mathbf{n}_T\|) \leq D(R+E)$, we have (21).

From the analysis of Theorem 1, we can readily derive a static regret bound for COMUDO. The proof follows by replacing the dynamic online benchmark $\{\mathbf{x}_t^o\}$ with the fixed offline benchmark \mathbf{x}^o in the proof of Theorem 1.

Theorem 2. The static regret of COMUDO has the following upper bound:

$$\operatorname{Reg}_{s}(T) \leq \frac{2R}{\alpha} \sum_{t=2}^{T} \|\mathbf{n}_{t}\| + \frac{1}{2\alpha} \sum_{t=2}^{T} \|\mathbf{n}_{t}\|^{2} + \frac{\alpha D^{2}T}{2} + \frac{R^{2}}{2\alpha} + D(R+E).$$
(23)

B. Bounding the Hard Constraint Violation

We bound the hard constraint violation yielded by CO-MUDO. To relate the lower-and-upper-bounded virtual queue to the hard constraint violation, we define a Lyapunov drift

$$\Theta_{t-1}^{n} \triangleq \frac{1}{2} (Q_{t}^{n} - V)^{2} - \frac{1}{2} (Q_{t-1}^{n} - V)^{2}, \quad \forall t, \forall n.$$
 (24)

We bound Θ_{t-1}^n produced by COMUDO in the following lemma, using *both* the virtual queue lower bound and upper bound in (13).³

³Existing virtual queue techniques for constrained online optimization can be divided into two groups: Algorithms like [21], [25], [27], [32], [34] provide soft constraint violation bounds by constructing a virtual queue that admits an upper bound only. The algorithm in [41] bounds the hard constraint violation by constructing a virtual queue that enforces a non-zero lower bound. In contrast, our virtual queue construction allows us to incorporate both lower and upper bounds into the Lyapunov drift analysis, leading to improved performance bounds.

Lemma 4. The Lyapunov drift yielded by COMUDO has the following upper bound for any t and n:

$$\Theta_{t-1}^{n} \leq Q_{t-1}^{n} [\gamma \tilde{g}_{t-1}^{n} (\mathbf{x}_{t}^{n})]_{+} - V [\gamma g_{t}^{n} (\mathbf{x}_{t}^{n})]_{+} + \frac{\gamma^{2} G}{\eta} \left| \tilde{g}_{t-1}^{n} (\mathbf{x}_{t}^{n}) - g_{t}^{n} (\mathbf{x}_{t}^{n}) \right| + 2\gamma^{2} G^{2}.$$
(25)

Proof: We can show that

$$\begin{aligned} (Q_t^n - V)^2 &\leq \left[(Q_{t-1}^n - V) + ([\gamma g_t^n(\mathbf{x}_t^n)]_+ - \eta Q_{t-1}^n) \right]^2 \\ &= (Q_{t-1}^n - V)^2 + \left([\gamma g_t^n(\mathbf{x}_t^n)]_+ - \eta Q_{t-1}^n)^2 - 2V [\gamma g_t^n(\mathbf{x}_t^n)]_+ \\ &+ 2Q_{t-1}^n [\gamma g_t^n(\mathbf{x}_t^n)]_+ - 2\eta (Q_{t-1}^n - V) Q_{t-1}^n \end{aligned}$$
(26)

where the first inequality is due to $|\max\{x, y\} - y| \le |x - y|$ for any $x, y \ge 0$.

Using the constraint function upper bound in (11), the virtual queue upper bound in (13), and the triangle inequality, the second term on the RHS of (26) has an upper bound $|[\gamma g_t^n(\mathbf{x}_t^n)]_+ - \eta Q_{t-1}^n| \leq |\gamma g_t^n(\mathbf{x}_t^n)| + \eta Q_{t-1}^n \leq \gamma G + \eta \frac{\gamma G}{\eta} = 2\gamma G$. From $Q_t^n \leq \frac{\gamma G}{\eta}$ and $|[x]_+ - [y]_+| \leq |x - y|$, the fourth term on the RHS of (26) can be bounded as $Q_{t-1}^n[\gamma g_t^n(\mathbf{x}_t^n)]_+ \leq Q_{t-1}^n[\gamma \tilde{g}_{t-1}^n(\mathbf{x}_t^n)]_+ + \frac{\gamma^2 G}{\eta}|\tilde{g}_{t-1}^n(\mathbf{x}_t^n) - g_t^n(\mathbf{x}_t^n)|$. For the last term on the RHS of (26), using the lower bound on the virtual queue $Q_t^n \geq V$ in (13), we directly have $-(Q_{t-1}^n - V)Q_{t-1}^n \leq 0$. Substituting the above three inequalities into the RHS of (26), rearranging terms, and dividing by 2, we complete the proof.

Using the results in Lemma 4, we bound the hard constraint violation (8) of COMUDO in the following theorem.

Theorem 3. The hard constraint violation yielded by CO-MUDO is upper bounded as

$$\begin{aligned} \operatorname{Vio}_{h}(T) &\leq \frac{\gamma G}{\eta V N} \sum_{t=2}^{T} \sum_{n=1}^{N} \left| \tilde{g}_{t-1}^{n}(\mathbf{x}_{t}^{n}) - g_{t}^{n}(\mathbf{x}_{t}^{n}) \right| \\ &+ \frac{R}{\alpha \gamma V} \sum_{t=2}^{T} \| \mathbf{x}_{t-1}^{\circ} - \mathbf{x}_{t}^{\circ} \| + \frac{2R}{\alpha \gamma V} \sum_{t=2}^{T} \| \mathbf{n}_{t} \| + \frac{1}{2\alpha \gamma V} \sum_{t=2}^{T} \| \mathbf{n}_{t} \|^{2} \\ &+ \frac{D(R+E)T}{\gamma V} + \frac{\alpha D^{2}T}{2\gamma V} + \frac{R^{2}}{2\alpha \gamma V} + \frac{2\gamma G^{2}T}{V} + G. \end{aligned}$$

$$(27)$$

Proof: Substituting (25) of Lemma 4 into (18) of Lemma 3, summing the resulting inequality over n = 1, ..., N and t = 2, ..., T, dividing both sides by N, and then rearranging terms, we have

$$\frac{V}{N} \sum_{t=2}^{T} \sum_{n=1}^{N} [\gamma g_{t}^{n}(\mathbf{x}_{t}^{n})]_{+} \leq \frac{1}{N} \sum_{t=1}^{T-1} \sum_{n=1}^{N} [f_{t}^{n}(\mathbf{x}_{t}^{\circ}) - f_{t}^{n}(\hat{\mathbf{x}}_{t})]
- \frac{1}{N} \sum_{t=2}^{T} \sum_{n=1}^{N} \Theta_{t-1}^{n} + \frac{\gamma^{2}G}{\eta N} \sum_{t=2}^{T} \sum_{n=1}^{N} \left| \tilde{g}_{t-1}^{n}(\mathbf{x}_{t}^{n}) - g_{t}^{n}(\mathbf{x}_{t}^{n}) \right|
+ \frac{R}{\alpha} \sum_{t=2}^{T} \|\mathbf{x}_{t-1}^{\circ} - \mathbf{x}_{t}^{\circ}\| + \frac{2R}{\alpha} \sum_{t=2}^{T} \|\mathbf{n}_{t}\| + \frac{1}{2\alpha} \sum_{t=2}^{T} \|\mathbf{n}_{t}\|^{2}
+ \frac{\alpha D^{2}T}{2} + \frac{1}{2\alpha} \sum_{t=2}^{T} \phi_{t} + \frac{1}{2\alpha N} \sum_{t=2}^{T} \sum_{n=1}^{N} \psi_{t}^{n} + 2\gamma^{2}G^{2}T.(28)$$

We can show that the first term on the RHS of (28) is upper bounded by $f_t^n(\mathbf{x}_t^\circ) - f_t^n(\hat{\mathbf{x}}_t) \le D(R+E), \forall t$. From the initialization of the virtual queue $Q_1^n = V$, the second term on the RHS of (28) satisfies $-\sum_{t=2}^{T} \Theta_{t-1}^n = \frac{1}{2} \sum_{t=2}^{T} [(Q_{t-1}^n - V)^2 - (Q_t^n - V)^2] = \frac{1}{2} [(Q_1^n - V)^2 - (Q_T^n - V)^2] \le \frac{1}{2} (Q_1^n - V)^2 = 0$. Substituting the above two inequalities, and the derived bounds on $\sum_{t=2}^{T} \phi_t$ and $\sum_{n=1}^{N} \psi_t^n$ into (28), dividing by γV , and noting that $|g_1^n(\mathbf{x}_1^n - \hat{\mathbf{x}}_0)| \le G$, we have (27).

C. Reaching Sublinear Regret and Hard Constraint Violation

As is generally the case for online optimization, problem \mathbf{P} may be impossible to solve exactly when both the loss functions and the channel states are fully dynamic and their information is delayed. However, COMUDO can simultaneously achieve sublinear dynamic regret (6), sublinear static regret (7), and sublinear hard constraint violation (8), when the underlying system gradually stabilizes over time. This suggests that it can closely track the system environmental dynamics even when they continuously fluctuate over time.

To capture the system environmental dynamics, we introduce two commonly used system variation measures [25], [29], [30], [33], [34], [38]-[41]. The first one is the accumulated difference of the dynamic online benchmarks (or the path length) $\sum_{t=2}^{T} \|\mathbf{x}_{t-1}^{\circ} - \mathbf{x}_{t}^{\circ}\| = \mathcal{O}(T^{\mu})$. The second one is the accumulated difference of adjacent constraint functions $\frac{1}{N} \sum_{t=2}^{T} \sum_{n=1}^{N} \max_{\mathbf{x} \in \mathcal{X}} |\tilde{g}_{t-1}^{n}(\mathbf{x}) - g_{t}^{n}(\mathbf{x})| = \mathcal{O}(T^{\nu})$. Here, $\mu, \nu \in [0, 1]$ represent the time variability. We set the power regularization factor $\lambda = \mathcal{O}(T^{1-\omega})$, such that the accumulated noise satisfies $\sum_{t=2}^{T} \|\mathbf{n}_{t}\| = O(T^{\omega})$ and $\sum_{t=2}^{T} \|\mathbf{n}_{t}\|^{2} = O(T^{\omega})$ for any $\omega \in [0, 1]$.

The following corollary provides convergence rates for the dynamic regret and the hard constraint violation, in terms of the above system variation measures. This is achieved by properly selecting α , η , γ , and V in the bounds provided by Theorems 1 and 3.

Corollary 1. Dynamic regret and hard constraint violation: Let $\alpha = T^{\frac{\max\{\mu,\omega\}-1}{2}}$, $\eta = T^{\nu-1}$, $\gamma > \frac{1}{G}$, and $V = T^{1-\nu}$ in COMUDO, then we have

$$\operatorname{Reg}_{d}(T) = \mathcal{O}(T^{\frac{1+\max\{\mu,\omega\}}{2}}),$$
(29)

$$\operatorname{Vio}_{\mathrm{h}}(T) = \mathcal{O}(T^{\nu}). \tag{30}$$

We can see from Corollary 1 that a sufficient condition for COMUDO to reach $\operatorname{Reg}_{d}(T) = o(T)$ and $\operatorname{Vio}_{h}(T) = o(T)$ is: $\mu, \nu, \omega < 1$, *i.e.*, the accumulated system variations grow sublinearly over time. In this case, both $\operatorname{Reg}_{d}(T)$ and $\operatorname{Vio}_{h}(T)$ are guaranteed to converge. Note that even for error-free online optimization [25], [29], [30], [33], [34], [38]-[41], sublinear system variation is necessary for an online algorithm to achieve sublinear dynamic performance bounds.

The following corollary provides convergence rates for the static regret and the hard constraint violation, by substituting the corresponding values of α , η , γ , and V into the bounds provided by Theorem 2 and 3.

Corollary 2. Static regret and hard constraint violation: Let $\alpha = T^{\frac{\omega-1}{2}}$, $\eta = T^{\nu-1}$, $\gamma > \frac{1}{G}$, and $V = T^{1-\nu}$ in COMUDO, then we have

$$\operatorname{Reg}_{s}(T) = \mathcal{O}(T^{\frac{1+\omega}{2}}), \tag{31}$$

$$\operatorname{Vio}_{\mathsf{h}}(T) = \mathcal{O}(T^{\nu}). \tag{32}$$

Remark 4. The modified saddle-point-type algorithm proposed in [22] for OTA FL with fixed long-term constraints achieved $\mathcal{O}(\sqrt{T})$ static regret and $O(T^{\frac{3}{4}})$ soft constraint violation when only the information of the loss functions is delayed. No dynamic regret bound is given in [22]. In contrast, with delayed information of the loss functions, channel states, and power constraints, COMUDO achieves $\mathcal{O}(T^{\frac{1+\max\{\mu,\omega\}}{2}})$ dynamic regret that smoothly approaches a static regret of $\mathcal{O}(T^{\frac{1+\omega}{2}})$ as the dynamics of the online benchmark decrease. Meanwhile, COMUDO provides $\mathcal{O}(T^{\nu})$ hard constraint violation that smoothly approaches $\mathcal{O}(1)$ as the channel fluctuations reduce. Note that the learning objective in [22] is for each device to minimize its own accumulated local loss instead of the global loss. This eliminates the need of analog model aggregation and thus the resulting $\mathcal{O}(\sqrt{T})$ static regret bound in [22] is not impacted by the channel noise, *i.e.*, the parameter ω . The $\mathcal{O}(T^{\frac{1+\omega}{2}})$ static regret of COMUDO also approaches $\mathcal{O}(\sqrt{T})$ as the variation of channel noise diminishes.

VI. SIMULATION RESULTS

We conduct numerical experiments on OTA FL with longterm transmit power control [18]-[22] for standard image classification tasks. We showcase the performance advantages of COMUDO compared with the state-of-the-art approaches, for both convex and non-convex losses under realistic wireless network conditions. This complements the theoretical bounding analysis of COMUDO in Section V.

A. Simulation Setup

We investigate an OTA FL system comprised of a central server and N = 10 local devices. The noisy multiple-access fading channel consists of C = 1000 orthogonal subchannels divided in frequency. Each subchannel occupies $B_W = 15 \text{ kHz}$ bandwidth. We set the noise figure to $N_F = 10$ dB and the noise power spectral density to $N_0 = -174$ dBm/Hz. We model the wireless fading channel between the central server and each local device n as a Gauss-Markov process [44], where the channel state evolves according to $\mathbf{h}_{t+1}^n = \kappa \mathbf{h}_t^n + \mathbf{r}_t^n$. Here, κ represents channel correlation, $\mathbf{h}_t^n \sim \mathcal{CN}(\mathbf{0}, \xi^n \mathbf{I})$ with ξ^n representing the path-loss and shadowing effects, $\mathbf{r}_{t}^{n} \sim \mathcal{CN}(\mathbf{0}, (1-\kappa^{2})\xi^{n}\mathbf{I})$ is independent of \mathbf{h}_{t}^{n} . We consider pedestrian speed and set $\kappa = 0.997$. We consider urban macrocells and set $\xi^{n}[dB] = -31.54 - 37 \log_{10}(\rho^{n}) - s^{n}$ [45], where ρ^n is the distance to the server, and $s^n \sim \mathcal{CN}(0, \sigma^2)$ represents shadowing. We set $\rho^n = 500$ m, $\forall n$ and $\sigma^2 = 8$ dB. We assume the devices share the same power limit P.

We conduct classification on the MNIST [46], Fashion-MNIST [47], and CIFAR-10 [48] datasets for model training and evaluation. We consider a setting with *streaming* data and *non-i.i.d.* distribution, where each local device n can only access to data samples of label n and gathers a dataset of 20 random data samples at the end of each round t.

We implement the following benchmark schemes:

• *Idealized FL:* The idealized error-free FL algorithm [1] performs the standard local gradient descent model update and global model averaging at each round. It serves as a performance upper bound for COMUDO.

- *OTA-LPC:* The long-term power control (LPC) scheme [18] sets the transmit power around a predefined threshold by adjusting the power-scaling factor at each round. We set the threshold as the average power limit as in [18].
- *OTA-RCI:* The regularized channel inversion (RCI) strategy [20] (structurally similar to the one derived in [19]) is the state-of-the-art long-term power control approach for OTA FL over homogeneous and static channels. We extend it to model-difference aggregation over heterogeneous and time-varying channels.
- *OMUAA*: The online model updating with analog aggregation (OMUAA) algorithm in [21] is the current best solution for OTA FL with soft long-term power constraints, based on the current information of the i.i.d. data and channel at each round.
- *OTA-MSP:* The modified saddle-point (MSP) algorithm in [22] is the state-of-the-art approach for OTA FL with fixed soft long-term constraints and delayed loss function information. We extend it to model-difference aggregation and replace its fixed constraints with timevarying power constraints in (5).

Among the above benchmarks, only OTA-MSP is designed to use delayed information on the loss functions. For fair comparison with COMUDO, we use delayed loss function and channel state information in all benchmark solutions. Since OTA-LPC, OTA-RCI, OMUAA, and OTA-MSP consider soft power violation, they are run with a time-averaged power limit P at each device.

B. Convex Logistic Regression

For the experiment on convex loss functions, we employ multinomial logistic regression with cross-entropy loss, using the MNIST dataset. The total number of model parameters is 7840. For fair comparison among different schemes, we first find that setting the learning rate $\alpha = 0.01$ in Idealized FL achieves the best learning performance, and then we use the same α for COMUDO, OTA-LPC, OTA-RCI, OMUAA and OTA-MSP. We set $x_{\text{UB}} = 10$, $\lambda = 2 \times 10^{-6}$, $\eta = 1 \times 10^{-3}$, $\gamma = 1.2 \times 10^{-2}$, and V = 20 in COMUDO. In our presented results, we have optimized the hyperparameters of all other schemes.

Fig. 2 shows the averaged test accuracy and the averaged normalized hard power violation $\frac{1}{NT} \sum_{t=1}^{T} \sum_{n=1}^{N} \frac{[P_t^n - P]_+}{P}$ in dB, where P_t^n is the transmit power consumed by device n at time t. For a fair comparison of the learning performance, all schemes except Idealized FL are fine tuned to consume nearly the same average transmit power at the end of round T, *i.e.*, $\frac{1}{TN} \sum_{t=1}^{T} \sum_{n=1}^{N} P_t^n \approx P$. We set P = 16 dBm and T = 500. We observe that OTA-LPC achieves the lowest 79.4% average test accuracy. The reason is that the OTA-LPC approach only adjusts the power scaling factor to satisfy the power constraint, resulting in higher scaled noise during the model training process compared to the other schemes. OTA-RCI achieves 83.5% average test accuracy, which is slightly better than the 81.3% accuracy yielded by OTA-MSP. This is because the subchannel-wise inversion design provides OTA-RCI more direct power control capabilities compared with



Fig. 2. Averaged test accuracy and normalized hard power violation for convex logistic regression on MNIST.



Fig. 3. The impact of transmit power limit P on the averaged test accuracy and training loss.

the single gradient-descent-based dual update used in OTA-MSP. The virtual queue mechanism in COMUDO serves as a power-violation-aware direct regularization on the model difference. Together with its effectiveness in controlling the hard power constraint violation, COMUDO achieves a higher 86.1% average test accuracy than the 84.2% accuracy of OMUAA, while incurring much less hard power variation in both the initial and final training stages.

Fig. 3 presents a comparison of the averaged test accuracy and training loss performance between COMUDO, OTA-LPC, OTA-RCI, OMUAA, and OTA-MSP under different transmit power limit P. We simply change γ in COMUDO to meet different power limits. We have also optimized the hyperparameters of all other schemes to meet different P values. We observe that as P decreases, the accuracies of all other four schemes drop more significantly than COMUDO. For example, COMUDO can achieve $\sim 80\%$ test accuracy while OMUAA achieves around $\sim 70\%$ test accuracy, and the other three schemes can only achieve $\sim 60\%$ test accuracy when P = 8 dBm. This is because the virtual-queue-based hard power violation control in COMUDO dynamically regularizes the model difference to generate smaller noise during model training, compared with adjusting the power scaling factor in OTA-LPC, the channel inversion design in OTA-RCI, the soft power violation control in OMUAA, and the gradient-descentbased dual update in OTA-MSP.

C. Non-Convex Neural Network Training

While the theoretical performance bounds of COMUDO derived in Section V are established for convex loss functions, we also evaluate the practical performance of COMUDO in non-convex convolutional neural network training on the MNIST, Fashion-MNIST, and CIFAR-10 datasets. For the MNIST classification task, we employ a neural network architecture consisting of a convolutional layer with 10 filters of size 7×7 , followed by a ReLU activation layer, and a fully connected layer with a softmax output. The total number of model parameters in this architecture is 48,910. To ensure fair comparison, all the schemes are evaluated using the same $\alpha = 0.02$ learning rate. We set $\gamma = 2 \times 10^{-3}$, $\eta = 1 \times 10^{-3}$,



Fig. 4. Averaged test accuracy and normalized hard power violation for nonconvex neural network training on MNIST.

TABLE I Comparison on Averaged Test Accuracy

Datasets	Idealized-FL	COMUDO	OTA-LPC	OTA-RCI	OMUAA	OTA-MSP
MNIST	90.35	87.72	74.06	80.68	81.27	79.02
Fashion-MNIST	80.56	78.00	71.55	73.23	73.82	72.35
CIFAR-10	63.29	59.51	53.96	55.38	56.57	54.08

and V = 1 in COMUDO and have optimized the hyperparameters of all the other schemes.

Fig. 4 shows the averaged test accuracy and normalized hard power violation for neural network training on MNIST with P = 16 dBm and T = 2000. In the context of non-convex neural network training with noise, gradient descent based algorithms such as COMUDO typically converge to points in the vicinity of a local minimum. Around these local minima, the learning performance of COMUDO is close to Idealized-FL and substantially outperforms the rest of the four schemes.

We further consider fashion-MNIST and CIFAR-10 classification in Table I. For fashion-MNIST, we use a convolutional network consisting of two convolutional layers succeeded by two fully connected layers. The total number of model parameters is 144,370. We use the same power limit P and total training rounds T as for MNIST. For CIFAR-10, we use three convolutional layers followed by three fully connected layers with a total of 2,456,842 model parameters. The detailed neural network settings are available in our released code. Due to the complexity of CIFAR-10 classification, we set P = 23dBm and T = 4000 to reach satisfactory performance. From Table I, we can see that COMUDO significantly outperforms OTA-LPC, OTA-RCI, OMUAA, and OTA-MSP under different datasets and neural network settings.

VII. CONCLUSIONS

We have examined the problem of online OTA FL with time-varying loss functions and power constraints under delays on both the loss function and channel state information. We introduce an effective algorithm named COMUDO that minimizes the accumulated time-varying loss at the server, while adhering to individual power constraints at the local devices. COMUDO provides closed-form local model updates that are resilient to delays on the loss functions, channel states, and power constraints. Through a novel lower-andupper-bounded virtual queue design, COMUDO achieves improved regret and hard constraint violation bounds compared with the current best results, which is proven using a new Lyapunov drift analysis technique. Our simulation results on canonical image classification datasets show that COMUDO significantly outperforms the state-of-the-art methods under practical wireless network conditions.

The authors have provided public access to their code and data at https://github.com/yituo-liu/INFOCOM2025-COMUDO.

References

- H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. Y. Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. Intel. Conf. Artif. Intell. Statist. (AISTATS)*, 2017.
- [2] M. I. Jordan, J. D. Lee, and Y. Yang, "Communication-efficient distributed statistical inference," *J. Amer. Stat. Assoc.*, vol. 19, pp. 2322– 2358, Feb. 2018.
- [3] R. R. Gajjala, S. Banchhor, A. M. Abdelmoniem, A. Dutta, M. Canini, and P. Kalnis, "Huffman coding based encoding techniques for fast distributed deep learning," in *Proc. Workshop on Distrib. Mach. Learn.*, 2020.
- [4] S. Wang, T. Tuor, T. Salonidis, K. K. Leung, C. Makaya, T. He, and K. Chan, "When edge meets learning: Adaptive control for resourceconstrained distributed machine learning," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, 2018.
- [5] N. H. Tran, W. Bao, A. Zomaya, M. N. H. Nguyen, and C. S. Hong, "Federated learning over wireless networks: Optimization model design and analysis," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, 2019.
- [6] J. Wang, B. Liang, M. Dong, G. Boudreau, and A. Afana, "Online distributed optimization with efficient communication via temporal similarity," in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, 2023.
- [7] Z. Jiang, Y. Xu, H. Xu, Z. Wang, J. Liu, Q. Chen, and C. Qiao, "Computation and communication efficient federated learning with adaptive model pruning," *IEEE Trans. Mobile Comput.*, vol. 23, pp. 2003–2021, Mar. 2024.
- [8] D. Alistarh, D. Grubic, J. Li, R. Tomioka, and M. Vojnovic, "QSGD: Communication-efficient SGD via gradient quantization and encoding," in *Proc. Adv. Neural Info. Proc. Sys. (NeurIPS)*, 2017.
- [9] A. F. Aji and K. Heafield, "Sparse communication for distributed gradient descent," in *Proc. Conf. Empirical Methods Natural Lang. Process. (EMNLP)*, 2017.
- [10] G. Zhu, Y. Wang, and K. Huang, "Broadband analog aggregation for low-latency federated edge learning," *IEEE Trans. Wireless Commun.*, vol. 19, pp. 491–506, Jan. 2020.
- [11] K. Yang, T. Jiang, Y. Shi, and Z. Ding, "Federated learning via overthe-air computation," *IEEE Trans. Wireless Commun.*, vol. 19, pp. 2022– 2035, Mar. 2020.
- [12] Y. Chen, G. Zhu, and J. Xu, "Over-the-air computation with imperfect channel state information," in *Proc. IEEE Intel. Workshop on Signal Process. Advances in Wireless Commun. (SPAWC)*, 2022.
- [13] J. Yao, Z. Yang, W. Xu, D. Niyato, and X. You, "Imperfect CSI: A key factor of uncertainty to over-the-air federated learning," *IEEE Wireless Commun. Lett.*, vol. 12, no. 12, pp. 2273–2277, 2023.
- [14] D. Liu and O. Simeone, "Privacy for free: Wireless federated learning via uncoded transmission with adaptive power control," *IEEE J. Sel. Areas Commun.*, vol. 39, pp. 170–185, Jan. 2021.
- [15] N. Zhang and M. Tao, "Gradient statistics aware power control for overthe-air federated learning," *IEEE Trans. Wireless Commun.*, pp. 5115 – 5128, Mar. 2021.
- [16] F. Moradi Kalarde, B. Liang, M. Dong, Y. A. Eldemerdash Ahmed, and H. T. Cheng, "Power minimization in federated learning with over-theair aggregation and receiver beamforming," in *Proc. ACM Int. Conf. Modeling, Anal. Simulation Wireless Mobile Syst. (MSWiM)*, 2023.
- [17] C. Zhang, M. Dong, B. Liang, A. Afana, and Y. Ahmed, "Joint downlinkuplink beamforming for wireless multi-antenna federated learning," in *Proc. Int. Symp. Model. Optim. Mobile Ad Hoc Wireless Netw. (WiOpt)*, 2023.
- [18] M. M. Amiri and D. Gunduz, "Federated learning over wireless fading channels," *IEEE Trans. Wireless Commun.*, vol. 19, pp. 3546–3557, May 2020.
- [19] X. Yu, B. Xiao, W. Ni, and X. Wang, "Optimal adaptive power control for over-the-air federated edge learning under fading channels," *IEEE Trans. Commun.*, vol. 71, pp. 5199–5213, Sep. 2023.
- [20] X. Cao, G. Zhu, J. Xu, Z. Wang, and S. Cui, "Optimized power control design for over-the-air federated edge learning," *IEEE J. Sel. Areas Commun.*, vol. 40, pp. 342–358, Jan. 2022.
- [21] J. Wang, B. Liang, M. Dong, G. Boudreau, and H. Abou-Zeid, "Joint online optimization of model training and analog aggregation for wireless edge learning," *IEEE/ACM Trans. Netw.*, vol. 32, pp. 1212–1228, Apr. 2024.
- [22] X. Cao and T. Başar, "Distributed constrained online convex optimization over multiple access fading channels," *IEEE Trans. Signal Process.*, vol. 70, pp. 3468–3483, Jun. 2022.

- [23] C. Gutterman, K. Guo, S. Arora, T. Gilliland, X. Wang, L. Wu, E. Katz-Bassett, and G. Zussman, "Requet: Real-time QoE metric detection for encrypted YouTube traffic," ACM Trans. Multimedia Comput. Commun. Appl., vol. 16, no. 2s, pp. 1–28, 2020.
- [24] W.-K. Hsu, J. Xu, X. Lin, and M. R. Bell, "Integrated online learning and adaptive control in queueing systems with uncertain payoffs," *Oper. Res.*, vol. 70, no. 2, pp. 1166–1181, 2022.
- [25] J. Wang, M. Dong, B. Liang, G. Boudreau, and H. Abou-Zeid, "Delaytolerant OCO with long-term constraints: Algorithm and its application to network resource allocation," *IEEE/ACM Trans. Netw.*, vol. 31, no. 1, pp. 147–163, 2023.
- [26] J. Yuan and A. Lamperski, "Online convex optimization for cumulative constraints," in *Proc. Adv. Neural Info. Proc. Sys. (NeurIPS)*, 2018.
- [27] M. J. Neely, Stochastic Network Optimization with Application on Communication and Queueing Systems. Springer Nature, 2022.
- [28] S. Shalev-Shwartz, "Online learning and online convex optimization," *Found. Trends Mach. Learn.*, vol. 4, pp. 107–194, Feb. 2012.
- [29] M. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in Proc. Intel. Conf. Mach. Learn. (ICML), 2003.
- [30] O. Besbes, Y. Gur, and A. Zeevi, "Non-stationary stochastic optimization," Oper. Res., vol. 63, pp. 1227–1244, Sep. 2015.
- [31] M. Mahdavi, R. Jin, and T. Yang, "Trading regret for efficiency: Online convex optimization with long term constraints," *J. Mach. Learn. Res.*, vol. 13, pp. 2503–2528, Sep. 2012.
- [32] H. Yu and M. J. Neely, "A low complexity algorithm with O(√T) regret and O(1) constraint violations for online convex optimization with long term constraints," J. Mach. Learn. Res., vol. 21, pp. 1–24, Feb. 2020.
- [33] T. Chen, Q. Ling, and G. B. Giannakis, "An online convex optimization approach to proactive network resource allocation," *IEEE Trans. Signal Process.*, vol. 65, pp. 6350–6364, Dec. 2017.
- [34] X. Cao, J. Zhang, and H. V. Poor, "A virtual-queue-based algorithm for constrained online convex optimization with applications to data center resource allocation," *IEEE J. Sel. Topics Signal Process.*, vol. 12, pp. 703–716, Aug. 2018.
- [35] S. Lee and M. M. Zavlanos, "Distributed primal-dual methods for online constrained optimization," in Proc. Amer. Control Conf. (ACC), 2016.
- [36] S. Paternain, S. Lee, M. M. Zavlanos, and A. Ribeiro, "Distributed constrained online learning," *IEEE Trans. Signal Process.*, vol. 68, pp. 3486–3499, Jun. 2020.
- [37] D. Yuan, A. Proutiere, and G. Shi, "Distributed online optimization with long-term constraints," *IEEE Trans. Automat. Contr.*, vol. 67, pp. 1089– 1104, Mar. 2022.
- [38] X. Yi, X. Li, L. Xie, and K. H. Johansson, "Distributed online convex optimization with time-varying coupled inequality constraints," *IEEE Trans. Signal Process.*, vol. 68, pp. 731–746, Jan. 2020.
- [39] P. Sharma, P. Khanduri, L. Shen, D. J. Bucci, and P. K. Varshney, "On distributed online convex optimization with sublinear dynamic regret and fit," in *Proc. Asilomar Conf. Signal Sys. Comput. (ASILOMARSSC)*, 2021.
- [40] X. Yi, X. Li, T. Yang, L. Xie, T. Chai, and K. Johansson, "Regret and cumulative constraint violation analysis for online convex optimization with long term constraints," in *Proc. Intel. Conf. Mach. Learn. (ICML)*, 2021.
- [41] H. Guo, X. Liu, H. Wei, and L. Ying, "Online convex optimization with hard constraints: Towards the best of two worlds and beyond," in *Proc. Adv. Neural Info. Proc. Sys. (NeurIPS)*, 2022.
- [42] X. Yi, X. Li, T. Yang, L. Xie, T. Chai, and K. H. Johansson, "Regret and cumulative constraint violation analysis for distributed online constrained convex optimization," *IEEE Trans. Automat. Control*, vol. 68, pp. 2875–2890, May 2023.
- [43] H. Viswanathan, "Capacity of Markov channels with receiver CSI and delayed feedback," *IEEE Trans. Inf. Theory*, vol. 45, pp. 761–771, Mar. 1999.
- [44] I. Abou-Faycal, M. Medard, and U. Madhow, "Binary adaptive coded pilot symbol assisted modulation over rayleigh fading channels without feedback," *IEEE Trans. Commun.*, vol. 53, pp. 1036–1046, Jun. 2005.
- [45] H. Holma and A. Toskala, WCDMA for UMTS HSPA evolution and LTE. John Wiely & Sons, 2010.
- [46] Y. LeCun, C. Cortes, and C. Burges, "The MNIST database," 1998.
- [47] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-mnist: A novel image dataset for benchmarking machine learning algorithms," 2017.
- [48] A. Krizhevsky, G. Hinton *et al.*, "Learning multiple layers of features from tiny images," 2009.