

Harnessing Wireless Channels for Scalable and Privacy-Preserving Federated Learning

Anis Elgabli¹, Jihong Park², *Member, IEEE*, Chaouki Ben Issaid¹, and Mehdi Bennis¹, *Fellow, IEEE*

Abstract—Wireless connectivity is instrumental in enabling scalable federated learning (FL), yet wireless channels bring challenges for model training, in which channel randomness perturbs each worker’s model update while multiple workers’ updates incur significant interference under limited bandwidth. To address these challenges, in this work we formulate a novel constrained optimization problem, and propose an FL framework harnessing wireless channel perturbations and interference for improving privacy, bandwidth-efficiency, and scalability. The resultant algorithm is coined *analog federated ADMM (A-FADMM)* based on analog transmissions and the alternating direction method of multipliers (ADMM). In A-FADMM, all workers upload their model updates to the parameter server (PS) using a single channel via analog transmissions, during which all models are perturbed and aggregated over-the-air. This not only saves communication bandwidth, but also hides each worker’s exact model update trajectory from any eavesdropper including the honest-but-curious PS, thereby preserving data privacy against model inversion attacks. We formally prove the convergence and privacy guarantees of A-FADMM for convex functions under time-varying channels, and numerically show the effectiveness of A-FADMM under noisy channels and stochastic non-convex functions, in terms of convergence speed and scalability, as well as communication bandwidth and energy efficiency.

Index Terms—Analog federated ADMM, digital federated ADMM, distributed machine learning, privacy, time-varying channels.

I. INTRODUCTION

WIRELESS connectivity has a great potential to scale up federated learning (FL) [1]–[3] by cutting the wires between workers and their parameter server (PS). To this end, recent works have studied FL over wireless links, and improved the communication and energy efficiencies by optimizing worker scheduling, transmission power, and computing operations [4]–[16]. However, wirelessly connected workers may *interfere* with each other during their over-the-air

Manuscript received October 11, 2020; revised February 16, 2021; accepted April 29, 2021. Date of publication May 10, 2021; date of current version August 16, 2021. This work is supported by Academy of Finland 6G Flagship (grant no. 318927) and project SMARTER, projects EU-ICT IntelliIoT (grant agreement no. 957218) and EU-CHISTERA LearningEdge, Infotech-NOOR. The associate editor coordinating the review of this article and approving it for publication was R. F. Schaefer. (*Corresponding author: Chaouki Ben Issaid.*)

Anis Elgabli, Chaouki Ben Issaid, and Mehdi Bennis are with the Centre of Wireless Communications, University of Oulu, 90014 Oulu, Finland (e-mail: anis.elgabli@oulu.fi; chaouki.benissaid@oulu.fi; mehdi.bennis@oulu.fi).

Jihong Park is with the School of Information Technology, Deakin University, Geelong, VIC 3220, Australia (e-mail: jihong.park@deakin.edu.au).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TCOMM.2021.3078783>.

Digital Object Identifier 10.1109/TCOMM.2021.3078783

transmissions, while competing over limited bandwidth. Most existing works avoid such interference by allocating dedicated channels to different workers [4]–[8], which is not scalable and requires significant amounts of bandwidth to support a large number of workers. Alternatively, taking a cue from FL operations, several recent works have proposed a method harnessing interference without separate channel allocation [12]–[16] as we review next.

As illustrated in Fig. 1(a), FL aims to minimize $\sum_{n=1}^N f_n(\Theta)$ assuming N workers, by periodically uploading the local model θ_n (or local gradient $\nabla f_n(\Theta)$) of each worker and downloading a global model Θ from the PS. Under digital transmissions, i.e., *digital FL*, (i) the PS first receives each θ_n through a *separate channel per each worker*, and (ii) combines them into a global model $\Theta = \frac{1}{N} \sum_{n=1}^N \theta_n$. The first step is however vulnerable to model inversion and reconstruction attacks [17], [18] by an honest-but-curious PS. Since the entire model update trajectory is observable, the PS can infer the training samples, violating data privacy. Furthermore, it is not communication-efficient because workers have to be assigned orthogonal channels in order for the PS to decode their models. However, the PS only needs $\sum_{n=1}^N \theta_n$ rather than individual local models, motivating the need for analog over-the-air aggregation schemes as described next.

Unlike digital signal transmission of bit streams, each analog signal directly represents an element $\theta_{n,i}$ of θ_n by its amplitude, allowing signal superposition. Exploiting this property, each worker in *analog FL* transmits an analog signal $\theta_{n,i}$ over a *shared channel among all workers*, through which all $\theta_{n,i}$ ’s are superpositioned over the air while hiding each private local model in the crowd. Consequently, the PS receives $\sum_{n=1}^N h_{n,i} \theta_{n,i}$ that is *perturbed* by complex fading channel $h_{n,i}$. Due to the perturbed models, the convergence and accuracy of analog FL depend significantly on the channel characteristics. To obviate this problem, it is common to cancel out the perturbation via a *channel inversion* method dividing $\theta_{n,i}$ by $h_{n,i}$ before transmissions, as illustrated in Fig. 1(b). With channel inversion, transmissions are only allowed when $|h_{n,i}|^2 \geq \varepsilon$, in order to avoid excessive transmit power due to the inversion [12]–[14]. The choice of ε is heuristic, hindering the convergence analysis of analog FL. Moreover, this approach does not guarantee privacy. For example, when only one worker has a good gain in one channel, it reveals its local model updates to the PS, compromising privacy. Last but not least, the rule of transmitting only when $|h_{n,i}|^2 \geq \varepsilon$ totally ignores the power of the transmitted symbol itself.

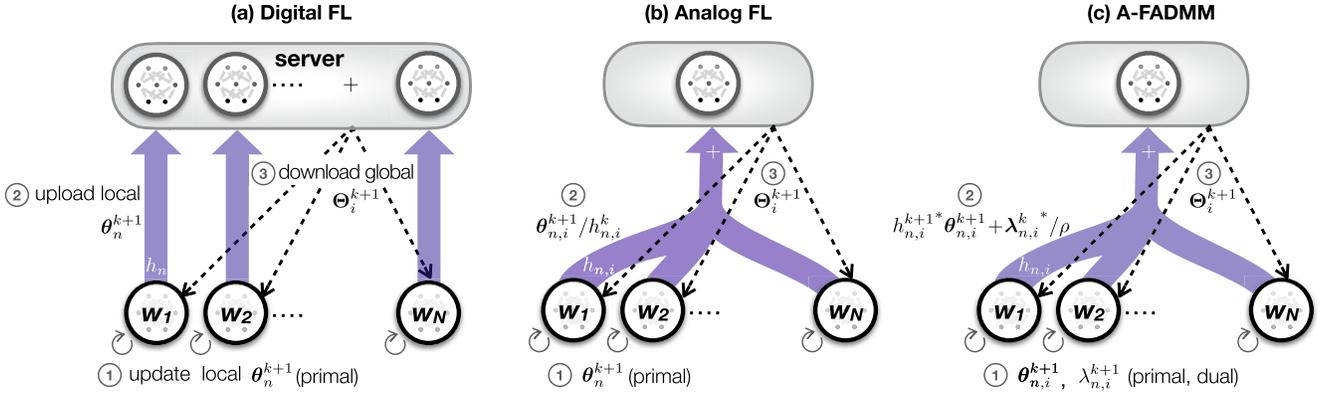


Fig. 1. Schematic illustrations of: (a) digital federated learning (FL), (b) analog FL with channel inversion, and (c) *analog-federated ADMM (A-FADMM)* without channel inversion.

Note that we are sending analog signals, and hence these limitations mandate a non-channel inversion method with formal convergence and privacy guarantees.

In this paper, we propose A-FADMM, a novel federated learning framework rooted in the alternating direction method of multipliers (ADMM) and analog over-the-air aggregation without channel inversion. Compared to the existing analog FL algorithms [12]–[14] based on first-order methods such as GD and SGD, A-FADMM is a *second-order* method providing faster convergence [19], [20]. Furthermore, A-FADMM does not apply channel inversion, so the PS receives the aggregate of perturbed model updates, thereby preserving privacy even when a single worker is transmitting over a given channel. This is done by integrating channel perturbations into the problem formulation, which may hamper the convergence particularly under time-varying channels. A-FADMM thus carefully updates the model parameters so that the time-varying channel does not hinder its convergence. Our major contributions are summarized as follows.

- On the theory front, this is the first work on analog transmission based distributed ADMM (primal-dual method) ensuring convergence while preserving privacy, under convex functions and time-varying channel dynamics. Existing works on analog FL focus on first-order primal methods, without proving convergence nor privacy guarantees [12]–[14].
- On the algorithmic front, our proposed A-FADMM is the first analog FL algorithm overcoming channel perturbations without channel inversion while ensuring convergence over time-varying channels.
- We numerically show that A-FADMM converges faster with comparable accuracy, compared to its digital transmission counterpart D-FADMM. Our simulations clarify under which conditions A-FADMM is preferable to D-FADMM, in terms of energy-efficiency, low-latency, and scalability.
- To further support the feasibility of A-FADMM, we elaborate on how to cope with constrained transmit power. Moreover, to corroborate the feasibility under stochastic and non-convex functions, we provide simulation results for the stochastic version of A-FADMM (SA-FADMM)

based on a deep neural network (DNN) in an image classification task.

The rest of the paper is structured as follows. In Section II, the proposed Analog Federated ADMM (A-FADMM) is described. The convergence of the proposed algorithm is studied in Section III, and the privacy analysis is provided in Section IV. The effectiveness of A-FADMM is numerically corroborated in Section V, in terms of accuracy and communication efficiency for linear regression and image classification using DNNs, followed by our conclusion in Section VI.

II. ANALOG FEDERATED ADMM

A-FADMM aims to aggregate multiple workers' updates at the PS without competition on the available bandwidth via analog transmissions. In this section, we describe A-FADMM operations based on a novel problem formulation, and explain how A-FADMM copes with the nuisances incurred by analog transmissions, in terms of time-varying channel fading, noise, and transmit power limitation.

A. Problem Formulation

The original problem of FL is to minimize $\frac{1}{N} \sum_{n=1}^N f_n(\Theta)$ with N workers, by locally minimizing $f_n(\theta_n)$ at each worker and globally averaging their model parameters θ_n at the PS. This boils down to the average consensus problem **(P1)** below.

$$\begin{aligned}
 \text{(P1)} \quad & \min_{\Theta, \{\theta_n\}_{n=1}^N} \sum_{n=1}^N f_n(\theta_n) \\
 & \text{s.t. } \theta_n = \Theta, \quad \forall n
 \end{aligned} \tag{1}$$

Primal-dual methods can solve **(P1)**, among which ADMM is one popular approach [19]–[21]. To implement this using digital transmissions, workers transmit their local model updates to the PS through *orthogonal channels*, wherein each local model is the primal variable θ_n while its dual variable λ_n is locally updated. Next, we explain the steps of the standard ADMM technique [19]–[21] in solving **(P1)**.

The augmented Lagrangian of **(P1)** is written as

$$\begin{aligned} \mathcal{L}_\rho(\Theta, \{\theta_n\}_{n=1}^N, \{\lambda_n\}_{n=1}^N) \\ = \sum_{n=1}^N f_n(\theta_n) + \sum_{n=1}^N \langle \lambda_n, \theta_n - \Theta \rangle \\ + \frac{\rho}{2} \sum_{n=1}^N \|\theta_n - \Theta\|_2^2, \end{aligned} \quad (2)$$

where $\rho > 0$ is a constant penalty for the disagreement between θ_n and Θ . At iteration $k+1$, each worker updates its primal variable by solving the following problem

$$\theta_n^{k+1} = \arg \min_{\theta_n} \left\{ f_n(\theta_n) + \langle \lambda_n^k, \theta_n - \Theta^k \rangle + \frac{\rho}{2} \|\theta_n - \Theta^k\|_2^2 \right\}. \quad (3)$$

Based on all workers' primal variable updates $\{\theta_n^{k+1}\}_{n=1}^N$ and previous dual variables $\{\lambda_n^k\}_{n=1}^N$, PS updates the global model Θ^{k+1} as follows

$$\Theta^{k+1} = \frac{1}{N} \sum_{n=1}^N (\theta_n^{k+1} + \frac{1}{\rho} \lambda_n^k). \quad (4)$$

Finally, given the updated global model Θ^{k+1} , each worker updates the dual variable λ_n^{k+1} as follows

$$\lambda_n^{k+1} = \lambda_n^k + \rho(\theta_n^{k+1} - \Theta^{k+1}). \quad (5)$$

To implement ADMM operations, each worker uploads θ_n^{k+1} to the PS, and then downloads Θ^{k+1} from PS, followed by locally updating the dual variable λ^{k+1} . Under digital transmission, the entire bandwidth is orthogonally allocated to each worker, while each update uploading or downloading corresponds to exchanging a fixed number of bits, e.g., 32 bits per model's element. To cope with channel fading and noise, adaptive modulation and error-correction coding are used in digital transmission.

By contrast, using analog transmissions, N workers transmit the i -th element of their updates using the same i -th subcarrier (channel). The benefit of analog transmissions is to aggregate all workers' updates over-the-air in one channel use, but at the cost of perturbations by channel fading $h_{n,i}$ which are assumed to follow an independent and identically distributed (IID) complex Gaussian distribution. These fading perturbations are often cancelled by multiplying $1/h_{n,i}$ before transmission, i.e., channel inversion [12]–[14]. Alternatively, we avoid channel inversion by reformulating **(P1)** into **(P2)** below, where the subscript i denotes the i -th element.

$$\text{(P2)} \quad \min_{\Theta, \{\theta_n\}_{n=1}^N} \sum_{n=1}^N f_n(\theta_n) \quad (6)$$

$$\text{s.t. } h_{n,i} \theta_{n,i} = h_{n,i} \Theta_i, \quad \forall n, i \quad (7)$$

where the subscript i represents the i -th element in the vector. Note that, in **(P2)**, we assume that the channel is constant and the system is noise free. We will relax these assumptions later in the section. In **(P2)**, constraint (1) is recast as its equivalent constraint (7) that allows A-FADMM to be updated directly using perturbed updates, i.e., constraints (1) and (7) are the

exact same constraint. Multiplying both sides of the equation with h does not change the constraint. We are utilizing analog transmission without resorting to channel inversion. Therefore, the parameter server will receive an aggregated and channel perturbed version of the local models instead of the aggregate models. Hence, to account for the channel perturbed values of the models in the updating step of ADMM and ensure convergence, we introduce the channel in the constraint and carry out the analysis. With that we, are able to derive an updating formula of the primal and dual variables relying on the aggregated and channel perturbed versions of the local models. The details are explained in the next section.

B. Primal, Dual, and Global Model Updates

The Lagrangian of **(P2)** is written as follows

$$\begin{aligned} \mathcal{L}_\rho(\theta_n, \lambda) = \sum_{n=1}^N f_n(\theta_n) + \sum_{i=1}^d \sum_{n=1}^N \lambda_{n,i}^* h_{n,i} (\theta_{n,i} - \Theta_i) \\ + \frac{\rho}{2} \sum_{i=1}^d \sum_{n=1}^N |h_{n,i}|^2 (\theta_{n,i} - \Theta_i)^2, \end{aligned} \quad (8)$$

where $\lambda_{n,i}^*$ is the conjugate of the complex dual variable $\lambda_{n,i}$, d is the cardinality of θ_n (i.e., model size), and $\rho > 0$ is a constant penalty for the local and global model disagreement. At iteration $k+1$, each worker updates its primal variable θ_n^{k+1} so as to minimize $\mathcal{L}_\rho(\theta_n, \Theta^k, \lambda^k)$. Hence, θ_n^{k+1} , $\forall i = 1, \dots, d$, should satisfy the following equation

$$0 \in \partial_i f_n(\theta_n^{k+1}) + \lambda_{n,i}^k h_{n,i} + \rho |h_{n,i}|^2 (\theta_{n,i}^{k+1} - \Theta_i^k), \quad (9)$$

where $\partial_i f_n(\cdot)$ denotes the i -th element in the sub-gradient vector of $f_n(\cdot)$.

Next, PS updates the i -th element Θ_i^{k+1} of the global model that minimizes $\mathcal{L}_\rho(\theta_n^{k+1}, \Theta, \lambda^k)$. By taking the derivative of $\mathcal{L}_\rho(\theta_n^{k+1}, \Theta, \lambda^k)$ with respect to Θ_i and equating to zero, Θ_i^{k+1} is given by

$$\Theta_i^{k+1} = \frac{1}{\sum_{n=1}^N |h_{n,i}|^2} \sum_{n=1}^N \left(|h_{n,i}|^2 \theta_{n,i}^{k+1} + h_{n,i} \lambda_{n,i}^k / \rho \right). \quad (10)$$

Finally, the dual variables are updated at each worker as follows

$$\lambda_{n,i}^{k+1} = \lambda_{n,i}^k + \rho h_{n,i} (\theta_{n,i}^{k+1} - \Theta_i^{k+1}). \quad (11)$$

Next, we will discuss how to implement the aforementioned update rules under time-varying channel fading, noise, and transmit power limitation.

C. Time-Varying Channel

The primal-dual update rules in (9) and (11) do not ensure the non-increase of the optimality gap when $h_{n,i}^{k+1} \neq h_{n,i}^k$. In this case, instead of updating $\theta_{n,i}^{k+1}$ using (9), we choose $\theta_{n,i}^{k+1} = \theta_{n,i}^k$, and find $\lambda_{n,i}^k$ that satisfies (9); in other words, the primal update problem (9) is flipped to the dual update problem. In doing so, A-FADMM copes with the channel changes reflected in the dual variables, and ensures that the primal-dual variables are still optimal for the given channel $h_{n,i}^{k+1}$.

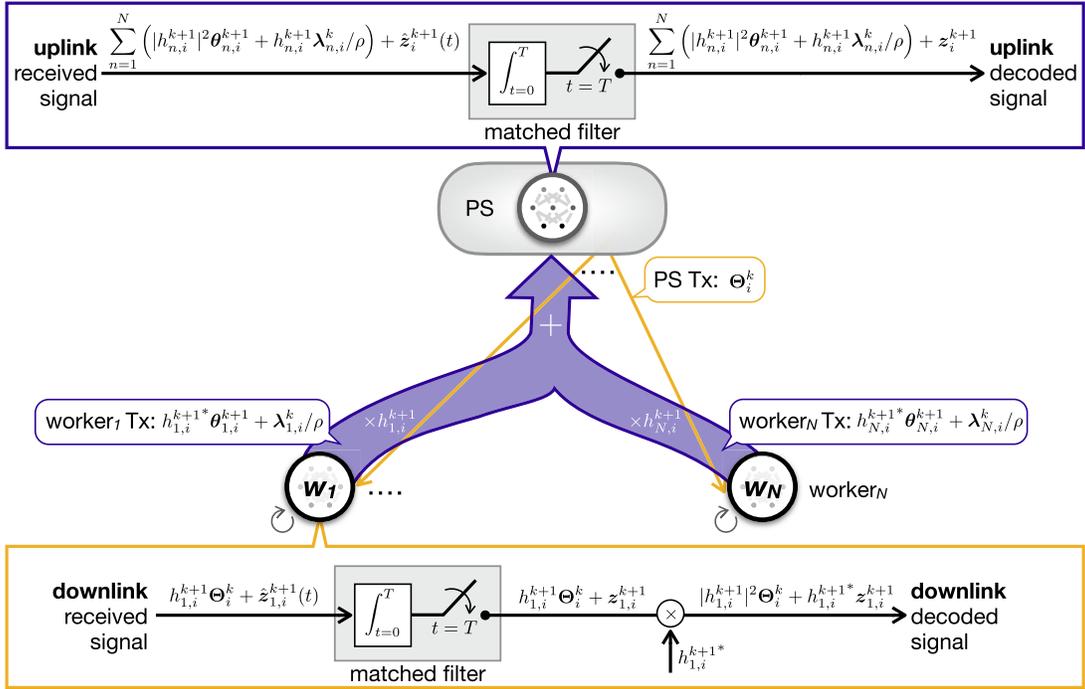


Fig. 2. An illustration of uplink and downlink communication in A-FADMM under channel noise.

Algorithm 1 Analog Federated ADMM (A-FADMM)

- 1: **Input:** $N, f_n(\theta_n) \forall n, \rho, K$, **Output:** $\theta_n, \forall n$
 - 2: **Initialization:** $\theta_n^{(0)}, \Theta_n^{(0)}, \lambda_n^{(0)}, \forall n$
 - 3: **A-FADMM:**
 - 4: **while** $k \leq K$ **do**
 - 5: **All workers** ($n \in \{1, \dots, N\}$): **in Parallel**
 - 6: **for** $i = 1, \dots, d$ **do**
 - 7: **if** $h_{n,i}^{k+1} = h_{n,i}^k$ **then**
 - 8: Find $\theta_{n,i}^{k+1}$ that satisfies (14)
 - 9: **else**
 - 10: $\theta_{n,i}^{k+1} = \theta_{n,i}^k$
 - 11: Find $(\lambda_{n,i}^k)^*$ that satisfies (14)
 - 12: **end if**
 - 13: **end for**
 - 14: Send $(h_{n,i}^{k+1})^* \theta_{n,i}^{k+1} + (\lambda_{n,i}^k)^* / \rho, \forall i = 1, \dots, d$ to the parameter server
 - 15: **Parameter Server:**
 - 16: Find Θ_i^{k+1} that satisfies (13)
 - 17: Broadcast $\Theta_i^{k+1}, \forall i = 1, \dots, d$ to all workers
 - 18: **All workers** ($n \in \{1, \dots, N\}$): **in Parallel**
 - 19: Update $\lambda_{n,i}^{k+1}$ locally via (15)
 - 20: $k \leftarrow k + 1$
 - 21: **end while**
-

D. Uploading and Downloading Information

We assume that every worker knows its individual channel $h_{n,i}^{k+1}$, while the PS knows the aggregate channel $\sum_{n=1}^N |h_{n,i}^{k+1}|^2$ using pilot signals [22]. Then, to update the i -th element of the global model Θ_i , each worker uploads $h_{n,i}^{k+1*} \theta_{n,i}^{k+1} + \lambda_{n,i}^k / \rho$, where $h_{n,i}^k$ is the conjugate of the complex chan-

nel $h_{n,i}$. Hence, after channel perturbation, the PS receives $\sum_{n=1}^N (|h_{n,i}^{k+1}|^2 \theta_{n,i}^{k+1} + h_{n,i}^{k+1} \lambda_{n,i}^k / \rho)$ in (10). By downloading Θ_i , each worker locally updates the primal and dual variables using (9) and (11), respectively.

E. Noisy Channel

In practical systems, the received signal is not only perturbed by channel fading but also distorted by additive white Gaussian noise (AWGN). Under digital transmissions, the noise can be alleviated using digital modulation and error correction coding schemes [23], [24]. By contrast, A-FADMM conveys uncoded information using analog transmissions. Therefore, the received information is perturbed by multiplicative fading and distorted by additive noise. A-FADMM directly utilizes the fading perturbed updates, yet still corrects channel noise using matched filtering (i.e., correlator receiver) as follows.

In the uplink of iteration $k + 1$, as illustrated in Fig. 2, each worker uploads its update $h_{n,i}^{k+1*} \theta_{n,i}^{k+1} + \lambda_{n,i}^k / \rho$ to the PS over the i -th subcarrier for T seconds. Propagating through the wireless channel, each update is perturbed by fading (i.e., multiplying by $h_{n,i}^{k+1}$), aggregated across all workers, and distorted by channel noise (i.e., adding $\hat{z}_i^{k+1}(t) \sim \mathcal{CN}(0, N_o)$). Consequently, the PS receives $\sum_{n=1}^N (|h_{n,i}^{k+1}|^2 \theta_{n,i}^{k+1} + h_{n,i}^{k+1} \lambda_{n,i}^k / \rho) + \hat{z}_{n,i}^{k+1}(t)$ at every instant $t \in [0, T]$, where the AWGN $\hat{z}_{n,i}^{k+1}(t) \sim \mathcal{CN}(0, N_o)$. The matched filter (i.e., correlator receiver) at PS integrates the received signals during T , and takes a sample

at $t = T$, resulting in

$$\begin{aligned} \frac{1}{T} \int_{t=0}^T \left\{ \left(\sum_{n=1}^N |h_{n,i}^{k+1}|^2 \boldsymbol{\theta}_{n,i}^{k+1} + \boldsymbol{\lambda}_{n,i}^k * h_{n,i}^{k+1} / \rho \right) + \hat{z}_i^{k+1}(t) \right\} dt \\ = \sum_{n=1}^N |h_{n,i}^{k+1}|^2 \boldsymbol{\theta}_{n,i}^{k+1} + \boldsymbol{\lambda}_{n,i}^k * h_{n,i}^{k+1} / \rho + \mathbf{z}_{n,i}^{k+1}, \end{aligned} \quad (12)$$

where the resultant noise $\mathbf{z}_{n,i}^{k+1} \sim \mathcal{CN}(0, N_0/T)$ whose variance is reduced from N_0 to N_0/T . Accordingly, the global model update is given by

$$\boldsymbol{\Theta}_i^{k+1} = \frac{\sum_{n=1}^N \left(|h_{n,i}^{k+1}|^2 \boldsymbol{\theta}_{n,i}^{k+1} + h_{n,i}^{k+1} \boldsymbol{\lambda}_{n,i}^k * / \rho + \mathbf{Re}\{\mathbf{z}_{n,i}^{k+1}\} \right)}{\sum_{n=1}^N |h_{n,i}^{k+1}|^2}, \quad (13)$$

where $\mathbf{Re}\{\mathbf{z}_{n,i}^{k+1}\}$ is because $\boldsymbol{\Theta}_i^{k+1}$ is real. Note that $\boldsymbol{\lambda}_{n,i}^k$ is complex, but the term $h_{n,i}^{k+1} \boldsymbol{\lambda}_{n,i}^k *$ is still real.

Likewise, in the downlink, the matched filter at each worker provides $h_{n,i}^{k+1} \boldsymbol{\Theta}_i^{k+1} + \mathbf{z}_{n,i}^{k+1}$. To make this output fit with the primal and dual updates, the output is multiplied by $h_{n,i}^{k+1*}$, and $|h_{n,i}^{k+1}|^2 \boldsymbol{\Theta}_i^{k+1} + h_{n,i}^{k+1*} \mathbf{z}_{n,i}^{k+1}$ is used for the following primal update rule

$$\mathbf{0} \in \partial_i f_n(\boldsymbol{\theta}_n^{k+1}) + (\boldsymbol{\lambda}_{n,i}^k) * h_{n,i}^{k+1} + \rho |h_{n,i}^{k+1}|^2 (\boldsymbol{\theta}_n^{k+1} - \boldsymbol{\Theta}_i^{k+1}) - \rho \mathbf{Re}\{h_{n,i}^{k+1*} \mathbf{z}_{n,i}^{k+1}\}, \quad (14)$$

and the dual update rule is given by

$$\boldsymbol{\lambda}_{n,i}^{k+1} = \boldsymbol{\lambda}_{n,i}^k + \rho h_{n,i}^{k+1} (\boldsymbol{\theta}_n^{k+1} - \boldsymbol{\Theta}_i^{k+1}) - \rho \mathbf{Re}\{\mathbf{z}_{n,i}^{k+1}\}. \quad (15)$$

The aforementioned operations of A-FADMM are summarized in Algorithm 1 in the previous page.

F. Power Control

Another practical concern is each worker's transmit power limitation. In order not to violate the maximum power budget P , before transmission each worker calculates its local power scaling factor α_n^{k+1} such that $(\alpha_n^{k+1})^2 \sum_{i=1}^d |h_{n,i}^{k+1*} \boldsymbol{\theta}_{n,i}^{k+1} + \boldsymbol{\lambda}_{n,i}^k * / \rho|^2 = P$, and sends α_n^{k+1} to the PS. Then, the PS determines $\alpha^{k+1} = \min\{\alpha_1^{k+1}, \dots, \alpha_N^{k+1}\}$ that is downloaded by every worker. Finally, each worker transmits $\alpha^{k+1} (h_{n,i}^{k+1*} \boldsymbol{\theta}_{n,i}^{k+1} + \boldsymbol{\lambda}_{n,i}^k * / \rho)$ to the PS, and after matched filtering and dividing by α^{k+1} , the PS obtains $\sum_{n=1}^N (|h_{n,i}^{k+1}|^2 \boldsymbol{\theta}_{n,i}^{k+1} + h_{n,i}^{k+1} \boldsymbol{\lambda}_{n,i}^k * / \rho) + \mathbf{z}_{n,i}^{k+1} / \alpha^{k+1}$ for the global model update. Note that α^{k+1} and α_n^{k+1} are scalar values that can be exchanged with negligible communication overhead, e.g., through separate control signaling channels [25].

III. CONVERGENCE ANALYSIS

In this section, we prove the optimality and convergence of A-FADMM for convex functions under noise-free but time-varying channels. Before stating the main results of the paper, we make the following assumptions

- *Assumption 1:* The local functions f_n are convex, proper and closed,
- *Assumption 2:* The unaugmented Lagrangian \mathcal{L}_0 has a saddle point.

The necessary and sufficient optimality conditions are the primal and dual feasibility given by

$$\boldsymbol{\theta}_n^* = \boldsymbol{\Theta}^* \quad \forall n \quad \text{and} \quad (16)$$

$$\mathbf{0} \in \partial_i f_n(\boldsymbol{\theta}_n^*) + \boldsymbol{\mu}_{n,i}^* \quad \forall n, \quad (17)$$

where the superscript \star denotes the value at the convergence point. The term $\boldsymbol{\mu}_{n,i} = \boldsymbol{\lambda}_{n,i} * h_{n,i}$ is the dual variable combined with channel fading. According to (11), $\boldsymbol{\mu}_{n,i}$ is updated as follows

$$\boldsymbol{\mu}_{n,i}^{k+1} = \boldsymbol{\mu}_{n,i}^k + \rho |h_{n,i}^{k+1}|^2 \mathbf{r}_{n,i}^{k+1}, \quad (18)$$

where $\mathbf{r}_{n,i}^{k+1} = \boldsymbol{\theta}_n^{k+1} - \boldsymbol{\Theta}_i^{k+1}$ is the n -th worker's primal residual. Applying the modified dual update rule (18) to the primal update rule (9), we obtain

$$\mathbf{0} \in \partial_i f_n(\boldsymbol{\theta}_n^{k+1}) + \boldsymbol{\mu}_{n,i}^{k+1} + \mathbf{S}_{n,i}^{k+1}, \quad (19)$$

where $\mathbf{S}_{n,i}^{k+1} = \rho |h_{n,i}^{k+1}|^2 (\boldsymbol{\Theta}_i^{k+1} - \boldsymbol{\Theta}_i^k)$ is the n -th worker's dual residual. Now, we are in position to introduce our first result, Lemma 1.

Lemma 1: For the iterates $\boldsymbol{\theta}_n^{k+1}$ the optimality gap of A-FADMM, is upper and lower bounded as follows.

$$\begin{aligned} (\text{Lower bound}) & \sum_{n=1}^N \left[f_n(\boldsymbol{\theta}_n^{k+1}) - f_n(\boldsymbol{\theta}_n^*) \right] \\ & \geq \sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^* \mathbf{r}_{n,i}^{k+1} \\ (\text{Upper bound}) & \sum_{n=1}^N \left[f_n(\boldsymbol{\theta}_n^{k+1}) - f_n(\boldsymbol{\theta}_n^*) \right] \\ & \leq - \sum_{n=1}^N \sum_{i=1}^d \left[\boldsymbol{\mu}_{n,i}^{k+1} \mathbf{r}_{n,i}^{k+1} - \mathbf{S}_{n,i}^{k+1} (\boldsymbol{\theta}_n^{k+1} - \boldsymbol{\theta}_n^{k+1}) \right]. \end{aligned} \quad (20)$$

The detailed proof is provided in Appendix VII-A. The main idea for the proof is to utilize the optimality of the updates in (9) and (10). We derive the upper bound for the objective function optimality gap in terms of the primal and dual residuals as stated in (21). To get the lower bound in (20) in terms of the primal residual, the definition of the Lagrangian (8) is used at $\rho = 0$. The result in Lemma 1 is used to derive the main results in Theorem 1 as presented next.

Theorem 1: When $f_n(\boldsymbol{\theta}_n)$ is closed, proper, and convex $\forall n$ and the Lagrangian \mathcal{L}_0 has a saddle point, under a time-varying channel, A-FADMM satisfies the following statements. Then, the optimality gap is non-increasing, i.e.,

$$\begin{aligned} \sum_{n=1}^N \sum_{i=1}^d \left\{ \frac{1}{\rho |h_{n,i}^{k+1}|^2} \left[(\boldsymbol{\mu}_{n,i}^{k+1} - \boldsymbol{\mu}_{n,i}^*)^2 - (\boldsymbol{\mu}_{n,i}^k - \boldsymbol{\mu}_{n,i}^*)^2 \right] \right. \\ \left. + \rho |h_{n,i}^{k+1}|^2 \left[(\boldsymbol{\Theta}_i^{k+1} - \boldsymbol{\Theta}_i^*)^2 - (\boldsymbol{\Theta}_i^k - \boldsymbol{\Theta}_i^*)^2 \right] \right\} \leq 0. \end{aligned} \quad (22)$$

The detailed proof of Theorem 1 is provided in Appendix VII-B. For the time-invariant scenario, we have the following corollary.

Corollary 1: For A-FADMM under a time-invariant channel where $h_{n,i}^{k+1} = h_{n,i}^k \quad \forall k$, it holds that

- The optimality gap converges to zero as $k \rightarrow \infty$, i.e.,

$$\lim_{k \rightarrow \infty} \sum_{n=1}^N f_n(\theta_n^k) = \sum_{n=1}^N f_n(\theta^*) \quad (23)$$

- Both primal and dual residuals converge to zero as $k \rightarrow \infty$, i.e.,

$$\lim_{k \rightarrow \infty} \mathbf{r}_{n,i}^k = \lim_{k \rightarrow \infty} \mathbf{S}_{n,i}^k = \mathbf{0} \quad (24)$$

The proof can be found in Appendix VII-C. The key idea is to define the Lyapunov function $V^k = \sum_{n=1}^N \sum_{i=1}^d \left[\frac{1}{\rho |h_{n,i}|^2} (\mu_{n,i}^k - \mu_{n,i}^*)^2 + \rho |h_{n,i}|^2 (\Theta_i^k - \Theta_i^*)^2 \right]$ and show that the difference between V^{k+1} and V^k monotonically decreases with k . This property enables proving that both primal and dual residuals converge to zero. Next, we apply Lemma 1, and prove the optimality gap goes to zero.

IV. PRIVACY ANALYSIS

Revealing the local model update trajectory is vulnerable to model inversion and membership attacks [17], [18]. These attacks can reconstruct a data sample via a generative model such as a generative adversarial network (GAN) [26], and identify which worker used the sample during its local model training. Against such an adversarial inverse problem, we aim to preserve privacy defined as follows.

Definition 1 [27]: A mechanism $M : M(X) \rightarrow Y$ is defined to be privacy preserving if the input X cannot be uniquely derived from the output Y .

We treat X as local models to be protected, and consider Y as the known information at an eavesdropper such as PS or another worker. Under digital transmissions, PS receives every local model θ_n^{k+1} , always violating privacy. Under analog FL wherein PS receives $\sum_{n=1}^N \theta_n^{k+1}$ after channel inversion, for certain iterations when only one worker sends the local model to PS, privacy is violated.

In sharp contrast, PS in A-FADMM receives $\sum_{n=1}^N (|h_{n,i}^{k+1}|^2 \theta_n^{k+1} + h_{n,i}^{k+1} \lambda_{n,i}^k / \rho)$. This does not violate privacy since the reception is the aggregate of fading-perturbed and dual-variable-distorted local models while $h_{n,i}^*$, $\lambda_{n,i}^*$, and N are unknown at PS. Consequently, according to Definition 1, A-FADMM preserves the privacy of every local update in a local model and its gradient, against PS or any eavesdropper knowing the global model trajectory.

In the literature of FL, privacy has often been analyzed through the lens of differential privacy (DP) [3], [28], [29]. A well-known DP mechanism is to distort the local or global model updates by injecting an additive random noise with zero mean and a constant variance. Compared to such a DP mechanism, A-FADMM distorts the local model updates by adding a random variable $h_{n,i}^k \lambda_{n,i}^k / \rho$ whose statistical properties are non-trivial and time-varying according to the locally computed dual variable updates $\lambda_{n,i}^k$. Moreover, A-FADMM additionally perturbs the local model updates by multiplying the random fading gain, i.e., $|h_{n,i}^{k+1}|^2 \theta_n^{k+1}$. Furthermore, the PS receives the aggregated sum of all perturbed signals of all workers which hides each individual's updated model.

Both additive random distortion with time-varying statistics and multiplicative random perturbation of A-FADMM are

ill-suited for DP analysis. Therefore, it is infeasible to fairly compare the existing DP based mechanisms and that of A-FADMM. Instead, following the privacy analysis in [27], we aim to show that any eavesdropper cannot uniquely infer a local model under A-FADMM, as stated in the following theorems.

Theorem 2: Unless $\Theta_i^{k+1} = \theta_{n,i}^{k+1}$ (i.e., before convergence), at every $k + 1$, A-FADMM preserves the privacy of each local model update $\theta_{n,i}^{k+1}$ and gradient update $\partial f_n(\theta_n^{k+1}) \forall n, i$.

Proof: Intuitively, we show that the inverse problem of an eavesdropper is to solve a set of equations at every iteration, in which the number of unknowns is larger than the number of equations. Therefore, each worker's local model or gradient cannot be uniquely derived. In fact, since $\theta_{n,i}^0$ and $\lambda_{n,i}^0 \forall n \in \{1, \dots, N\}, \forall i \in \{1, \dots, d\}$, are initiated randomly, then their values cannot be revealed by the eavesdropper. For simplicity, we assume that $f_n(\theta_n^{k+1})$ is differentiable and the system is noise free. The eavesdropper needs to solve either of the following two equations to derive $\theta_{n,i}^1$

$$\begin{cases} \theta_{n,i}^1 = \frac{\rho |h_{n,i}^1|^2 \Theta_i^0 - \nabla_i f_n(\theta_n^1) - (\lambda_{n,i}^0)^* h_{n,i}^1}{\rho |h_{n,i}^1|^2}, & \text{if } h_{n,i}^1 = h_{n,i}^0 \\ \theta_{n,i}^1 = \theta_{n,i}^0, & \text{if } h_{n,i}^1 \neq h_{n,i}^0 \end{cases} \quad (25)$$

Then, we can write

$$\begin{aligned} \theta_{n,i}^1 = & \left(\rho \sum_{n=1}^N |h_{n,i}^1|^2 \Theta_i^1 - \rho \sum_{m=1, m \neq n}^N |h_{m,i}^1|^2 \theta_{m,i}^1 \right. \\ & \left. - \sum_{n=1}^N (\lambda_{n,i}^0)^* h_{n,i}^1 \right) / \left(\rho |h_{n,i}^1|^2 \right). \end{aligned} \quad (26)$$

Note that the eavesdropper knows Θ_i^0 and Θ_i^1 . However, the values of $h_{n,i}^1$, $\lambda_{n,i}^0$, $\nabla_i f_n(\theta_n^1)$, $\rho \sum_{m=1, m \neq n}^N |h_{m,i}^1|^2 \theta_{m,i}^1$, and $\theta_{n,i}^0$ are unknown. Hence, even at the absence of the noise at the receiver, the eavesdropper cannot have a unique solution for $\theta_{n,i}^1$ and/or $\nabla_i f_n(\theta_n^1)$ since the number of variables $V = 5$ is greater than the number of equations $E = 2$. Writing the same equations for iteration k

$$\begin{cases} \theta_{n,i}^k = \frac{\rho |h_{n,i}^k|^2 \Theta_i^{k-1} - \nabla_i f_n(\theta_n^k) - (\lambda_{n,i}^{k-1})^* h_{n,i}^k}{\rho |h_{n,i}^k|^2}, & \text{if } h_{n,i}^k = h_{n,i}^{k-1} \\ \theta_{n,i}^k = \theta_{n,i}^{k-1}, & \text{if } h_{n,i}^k \neq h_{n,i}^{k-1} \end{cases} \quad (27)$$

$$\begin{aligned} \theta_{n,i}^k = & \left(\rho \sum_{n=1}^N |h_{n,i}^k|^2 \Theta_i^k - \rho \sum_{m=1, m \neq n}^N |h_{m,i}^k|^2 \theta_{m,i}^k \right. \\ & \left. - \sum_{n=1}^N (\lambda_{n,i}^{k-1})^* h_{n,i}^k \right) / \left(\rho |h_{n,i}^k|^2 \right). \end{aligned} \quad (28)$$

as well as for iteration $k + 1$

$$\begin{cases} \theta_{n,i}^{k+1} = \frac{\rho |h_{n,i}^{k+1}|^2 \Theta_i^k - \nabla_i f_n(\theta_n^{k+1}) - (\lambda_{n,i}^k)^* h_{n,i}^{k+1}}{\rho |h_{n,i}^{k+1}|^2}, \\ \quad \text{if } h_{n,i}^{k+1} = h_{n,i}^k, \\ \theta_{n,i}^{k+1} = \theta_{n,i}^k, \\ \quad \text{if } h_{n,i}^{k+1} \neq h_{n,i}^k \end{cases} \quad (29)$$

$$\begin{aligned} \theta_{n,i}^{k+1} = & \left(\rho \sum_{n=1}^N |h_{n,i}^{k+1}|^2 \Theta_i^{k+1} - \rho \sum_{m=1, m \neq n}^N |h_{m,i}^{k+1}|^2 \theta_{m,i}^{k+1} \right. \\ & \left. - \sum_{n=1}^N (\lambda_{n,i}^k)^* h_{n,i}^{k+1} \right) / \left(\rho |h_{n,i}^{k+1}|^2 \right). \end{aligned} \quad (30)$$

Similarly, the eavesdropper knows Θ_i^k and Θ_i^{k+1} . However, $h_{n,i}^{k+1}$, $\lambda_{n,i}^k$, $\nabla_i f_n(\theta_n^{k+1})$, $\rho \sum_{m=1, m \neq n}^N |h_{m,i}^{k+1}|^2 \theta_{m,i}^{k+1}$, and $\theta_{n,i}^k$ are unknown. We clearly see that if the algorithm has not converged to the optimal solution yet at iteration $k + 1$, i.e., $\theta_{n,i}^{k+1} \neq \Theta_i^{k+1}$, then there is no unique inversion of $\theta_{n,i}^{k+1}$ since the number of variables is more than the number of equations. This finalizes the proof. ■

Theorem 3: When $\Theta_i^{k+1} = \theta_{n,i}^{k+1}$ (i.e., at convergence), A-FADMM preserves the privacy of the local model trajectory $\{\theta_{n,i}^0, \dots, \theta_{n,i}^k\}$ and gradient trajectory $\{\partial f_n(\theta_n^0), \dots, \partial f_n(\theta_n^{k+1})\} \forall n, i$.

Proof: In brief, we show that after A-FADMM convergence when all local models become identical and known to an eavesdropper, this information cannot be used to derive a unique trajectory of each worker's local model and gradient updates. When $\theta_{n,i}^{k+1} = \Theta_i^{k+1}, \forall n, i$, we know from (30) that the following terms can be found at the PS: $\sum_{n=1}^N |h_{n,i}^{k+1}|^2 \Theta_i^{k+1}$, $\rho \sum_{m=1, m \neq n}^N |h_{m,i}^{k+1}|^2 \theta_{m,i}^{k+1}$. However, the terms $(\lambda_{n,i}^k)^* h_{n,i}^{k+1}$ and $\rho |h_{n,i}^{k+1}|^2$ cannot be found, and these two terms are needed to retrieve a unique solution for $\nabla_i f_n(\theta_n^{k+1})$ using (29). Hence, $\nabla_i f_n(\theta_n^{k+1})$ cannot be uniquely derived. From (27)-(28), we clearly see that knowing $\theta_{n,i}^{k+1}$, Θ_i^k , and Θ_i^{k-1} are not enough to find a unique solution for $\nabla_i f_n(\theta_n^k)$ and $\theta_{n,i}^k$ since all other terms in the two equations including $h_{n,i}^k$ and $\lambda_{n,i}^{k-1}$ are also unknown. Therefore, the individual model at the convergence point do not release any unique information about the updating steps of the model and the function gradient trajectory, which concludes the proof. ■

V. EXPERIMENTS

To validate our theoretical foundations, we numerically evaluate the performance of A-FADMM in convex (linear regression) and non-convex (image classification using DNNs) problems.

A. Simulation Settings

For linear regression, we use the California Housing dataset [30] consisting of 20000 samples with 6 features, i.e., model size $d = 6$. At iteration k , the loss is given as $|\sum_{n=1}^N [f(\theta_n^k) - f(\theta^*)]|$. For image classification, we use the MNIST dataset [31] comprising 60000 training and 10000 test samples, each of which represents a hand-written 0-9 digit

image. In this case, we consider a 3-layer fully connected multi-layer perceptron (MLP) comprising an input layer with 784 neurons, two hidden layers with 128 and 64 neurons, respectively, and an output layer with 10 neurons, resulting in the model size $d = 109184$. We use the rectified linear unit (ReLU) activation function, softmax output, and cross entropy loss.

By default, we consider $N = 100$ workers with SNR = 40dB, each of which stores the same number of training samples equally divided and allocated from the training dataset. These workers are supported using 10 and 4096 subcarriers for linear regression and DNNs, respectively. Following the LTE cellular standards [25], each subcarrier provides 15KHz bandwidth during 1ms. Each channel realization is coherent during 10 iterations, and is randomly generated by a Rayleigh fading distribution with zero mean and unit variance for every 10 iterations.

To focus primarily on the uplink bandwidth bottleneck in the simulations, analog transmissions are utilized only for the uplink, while digital transmissions are considered in the downlink where the PS broadcasting the global updates without any bandwidth competition. Consequently, in the resultant A-FADMM implementation under channel noise, the global model update after the analog uplink reception follows (13) as in Algorithm 1, whereas the primal and dual updates (originally given as (14) and (15) in Algorithm 1) after the digital downlink reception use the following rules:

$$0 \in \partial_i f_n(\theta_n^{k+1}) + (\lambda_{n,i}^k)^* h_{n,i}^{k+1} + \rho |h_{n,i}^{k+1}|^2 (\theta_{n,i}^{k+1} - \Theta_i^k) \quad (31)$$

$$\lambda_{n,i}^{k+1} = \lambda_{n,i}^k + \rho h_{n,i}^{k+1} (\theta_{n,i}^{k+1} - \Theta_i^{k+1}). \quad (32)$$

These noise-free primal and dual update rules are implemented as follows. In the digital downlink, each worker decodes Θ_i^{k+1} , and manually perturbs it as $|h_{n,i}^{k+1}|^2 \Theta_i^{k+1}$ that is used for updating primal and dual variables via (31) and (32).

In A-FADMM, the i -th element of the models of all workers are uploaded using the i -th sub-carrier. In linear regression, the model size is less than the number of available subcarriers, i.e., $d = 6 < 10$, and hence A-FADMM requires only one time slot (one upload) to upload all workers' models at each iteration. In image classification where $d = 109184$, it requires $\lceil 109184/4096 \rceil = 27$ time slots to uploads all workers' models per iteration.

In D-FADMM, the number of uploading time slots depends not only on the number of subcarriers but also on the channel gain of each subcarrier. To be precise, following the LTE cellular standards [25], each subcarrier provides $W_i = 15$ KHz bandwidth during 1ms. Each channel realization is coherent during 10 iterations, and is randomly generated by a Rayleigh fading distribution with zero mean and unit variance for every 10 iterations. When each model element consumes 32 bits, the n -th worker requires the uploading time slots \hat{T}_n that is the minimum T_n satisfying the following condition $\int_{t=1}^{T_n} \sum_{i=1}^{4096/N} R_{n,i}(t) dt \geq 32 d$, where $R_{n,i}(t) = W_i \log_2(1 + P|h_{n,i}(t)| / (W_i N_0))$ follows from the Shannon formula. Since each worker has independent channel realizations, to upload all workers' models to PS, it requires $\hat{T} = \max\{\hat{T}_1, \hat{T}_2, \dots, \hat{T}_N\}$

time slots. Based on Algorithm 1, D-FADMM is implemented by replacing (13) for Θ_i^{k+1} , (14) for $\theta_{n,i}^{k+1}$, and (15) for $\lambda_{n,i}^{k+1}$ with (9), (10), and (11), respectively.

In (8), we choose the penalty constant $\rho = 0.5$ yielding fast convergence for both digital and analog implementations from our observations. To run the experiments, we use Matlab for linear regression and TensorFlow for image classification, operated in a MacBook Air computer (1.8 GHz Intel Core i5 CPU, 8 GB 1,600 MHz DDR3 RAM). For each plot, we run 5 simulations, and report mean values (solid curves) and standard deviations (shaded areas, omitted for negligible values). Finally, we compare A-FADMM with the following benchmark algorithms.

- **D-FADMM** is the digital communication version of A-FADMM, wherein the total bandwidth is equally divided and allocated to each worker whose model element consumes 32 bits. i.e., the value of each element in the model vector is transmitted using 32 bits. Following A-FADMM, we use $\rho = 0.5$.
- **A-GD** is the analog communication versions of the distributed gradient descent algorithm (GD) with channel inversion that allows the n -th worker to upload its update only when the channel gain $|h_{n,i}^k| \geq \epsilon$. We use $\epsilon = 10^{-6}$ for communication and 10^{-4} learning rate for GD operations. We observed that A-GD diverges for a larger learning rate.

For the image classification task, we use the following baselines and hyperparameters.

- **A-SFADMM** is the stochastic version of A-FADMM using DNNs. For the ADMM problem, we use $\rho = 0.5$. For the local problem at each global iteration, each worker selects a mini-batch of size 100 samples at random, and uses the Adam optimizer with 0.01 learning rate to update its local model. Per global iteration, we consider 20 local iterations. For different choices of local iterations and learning rates, we study their impact on convergence speed and accuracy in Figures 6 and 7.
- **D-SFADMM** is the stochastic version of D-FADMM which are utilized in the classification problem using DNN. Following A-SFADMM, we use $\rho = 0.5$, Adam optimizer with 0.01 learning rate, mini-batch size 100, and 20 local iterations per global iteration.
- **A-SGD** is the stochastic version of A-GD with channel inversion (i.e., analog FL). Following A-GD, we use $\epsilon = 10^{-6}$ for communication. For SGD operations, we use mini-batch size 100, and choose the learning rate 0.005. Note that from our observations, A-SGD incurs high oscillation under the learning rate 0.01 used in A-SFADMM and D-SFADMM.

For both linear regression and image classification task, the notation **10x** implies an algorithm with 10x more subcarriers (bandwidth) than the default setting. For example, compared to A-SFADMM using 4096 subcarriers, A-SFADMM-10x utilizes 40960 subcarriers at each iteration. Accordingly, given the MLP model size $d = 109184$, A-SFADMM-10x requires $\lceil 109184/40960 \rceil = 3$ time slots

for uploading all workers' models, which is 9x less than A-SFADMM requiring $\lceil 109184/4096 \rceil = 27$ time slots.

B. Communication Efficiency

In linear regression, as observed in Fig. 3(a), A-FADMM requires the lowest communication rounds until achieving a target loss 10^{-4} . Even with 10x more subcarriers, D-FADMM fails to reach the same speed due to the orthogonal subcarrier allocation to each worker under limited bandwidth. However, if one aims to achieve very low loss below 10^{-4} , A-FADMM suffers from noisy reception, and D-FADMM may thus be a better choice, as long as very large bandwidth and/or long uploading time are available. In image classification, Fig.4(a) shows that A-SFADMM achieves the highest accuracy the minimum number of communication rounds. In fact, it is even more communication-efficient than D-SFADMM with 10x more subcarriers (D-SFADMM-10x).

For both tasks, analog FL (i.e., A-GD and A-SGD) struggles with intermittent uploads due to the truncated channel inversion (transmitting only when $|h_{n,i}| \geq \epsilon$). This yields too many communication rounds in linear regression (A-GD) and high variance in image classification (A-SGD), highlighting the importance of non-channel inversion methods used in A-FADMM and A-SFADMM.

C. Energy Efficiency vs. Accuracy

In this experiment, we assume that there are sufficient subcarriers to upload every update in one time slot, and focus on wireless communication energy consumption that often exceeds computing energy [32]. We measure the loss or accuracy when the total channel uses $\sum_{i=1}^j M_i$ at time slot j reaches a target maximum number of channel uses, where M_i is the number of subcarriers used in time slot i .

With linear regression task and 100k maximum number of channel uses, Fig. 3(b) shows that A-FADMM always achieves order-of-magnitude lower loss than D-FADMM, even at very low -10 dB SNR, i.e., low transmit power. With 200k channel uses, D-FADMM outperforms A-FADMM, but only at high SNR exceeding 20dB. This advocates that A-FADMM is more energy-efficient and bandwidth-efficient.

In image classification, as shown by Fig. 4(b), A-SFADMM not only outperforms D-SFADMM, but also achieves the maximum test accuracy even when the SNR is as low as -10 dB and the maximum number of channel uses is 30M. By contrast, D-SFADMM with 40dB SNR and 60M channel uses achieves maximum accuracy that is still lower than A-SFADMM's.

D. Scalability

We investigate the scalability of A-FADMM and A-SFADMM, by counting the number of channel uses until reaching a target loss or accuracy. We vary the number of contributing workers, and we assume that the noise power spectral density is fixed as 10^{-9} W/Hz. In linear regression, we clearly see from Fig.3(c) that A-FADMM does not require more channel uses for more workers to achieve a target

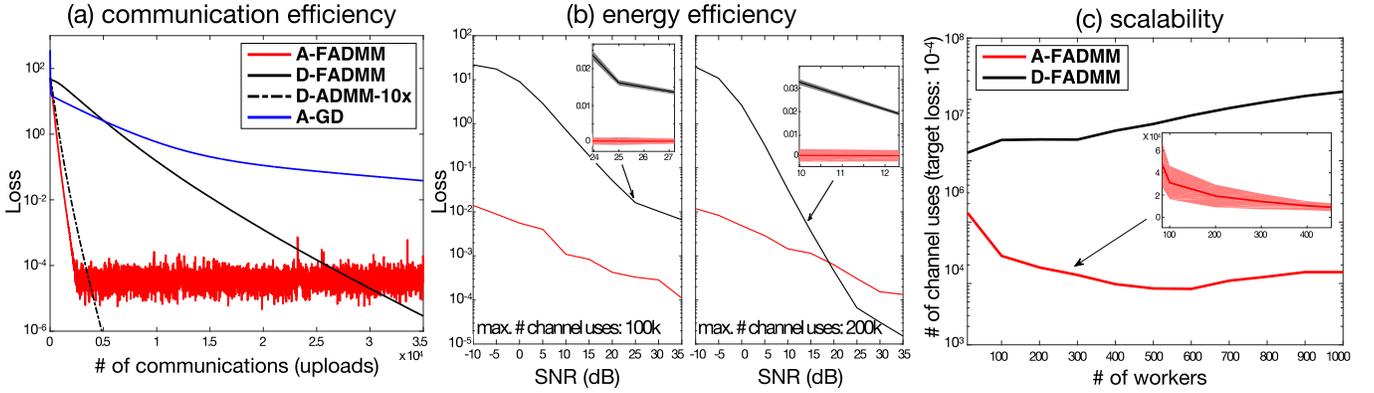


Fig. 3. Linear regression results showing: (a) communication efficiency (loss w.r.t. # of uploads); (b) energy efficiency (loss w.r.t. SNR); and (c) scalability (# of channel uses w.r.t # of workers).

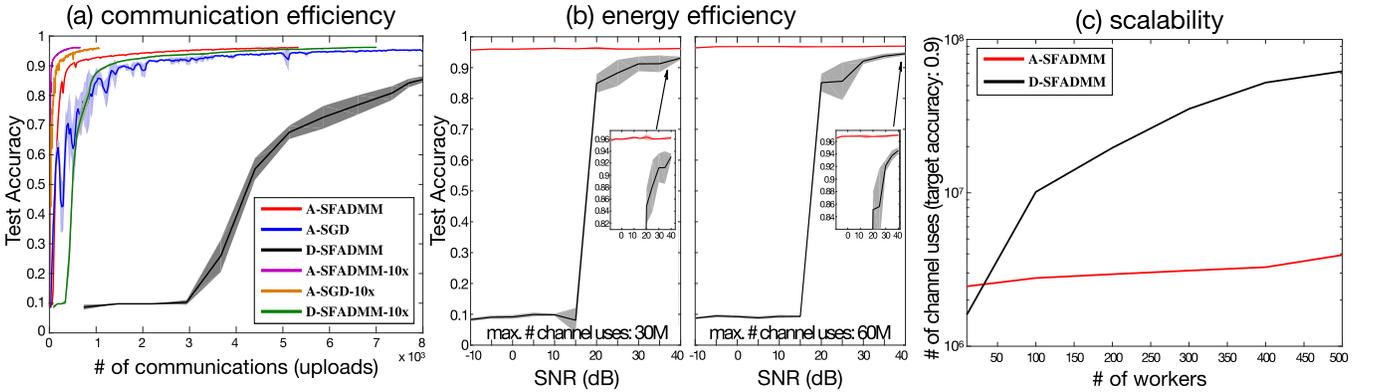


Fig. 4. Image classification results showing: (a) communication efficiency (test accuracy w.r.t. # of uploads); (b) energy efficiency (test accuracy w.r.t. SNR); and (c) scalability (# of channel uses w.r.t # of workers).

loss 10^{-4} . By contrast, D-FADMM necessitates the channel uses linearly proportional to the number of workers due to the orthogonal bandwidth allocation to every worker. It is worth mentioning that even with only $N = 10$ workers, A-FADMM requires order of magnitude less channel uses than D-FADMM. Similar trends are observed in Fig.4(b) for image classification, only except for the cases below $N = 10$ workers.

E. Sensitivity Analysis

In this subsection, we study the impact of hyperparameters on the convergence speed and accuracy of A-FADMM and D-FADMM as well as their stochastic versions. All the training, communication, and simulation environments are identical to the settings in Sec. V, except for the hyperparameters: disagreement penalty weight ρ , learning rate, and the number of local iterations as elaborated next.

a) Impact of ρ : The penalty weight ρ adjusts the degree of disagreement between local and global models in both linear regression and classification tasks. In linear regression, Fig. 5(a) shows that a larger ρ leads to faster convergence with diminishing returns for both A-FADMM and D-FADMM. Our choice $\rho = 0.5$ in Sec. V is thus a value yielding sufficiently fast convergence.

In image classification, on the other hand, Fig. 5(b) shows that a smaller ρ is slower at the beginning, but reaches the highest test accuracy faster. For small ρ , the penalty of disagreeing with other workers is not large. Therefore, every worker is likely to be biased towards its local optima. Since each worker has only a fraction of the global dataset, the convergence speed is fast, but the accuracy cannot outperform the global model averaged across all workers. For large ρ , workers tend to strictly reduce the local model disagreement from the beginning. This yields a faster jump to a high accuracy level at the early phase. However, keeping large ρ slows down the updating step by pushing all workers towards minimizing the disagreement in their model updates at every iteration. Given these observations, our choice $\rho = 0.5$ in Sec. V is a value yielding sufficiently fast convergence to the highest accuracy. To obviate the accuracy reduction at the beginning while keeping fast convergence speed, studying time-varying ρ (e.g., decreasing ρ from a large value with the number of iterations) could be an interesting topic for future study.

b) Impact of the Number of Local Iterations: Our image classification relies on DNNs, and thus cannot be solved in a closed form expression. Instead, at every global iteration k , several local iterations are performed, updating each local model. Ideally, each worker needs to iterate until convergence

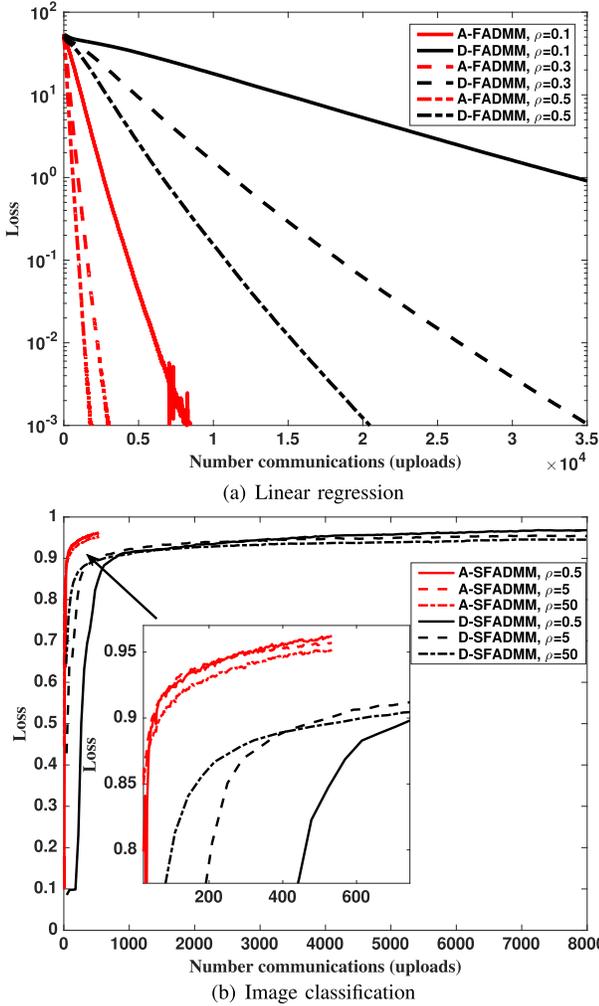


Fig. 5. Impact of the *disagreement penalty weight* ρ in (a) linear regression and (b) image classification using DNNs.

before sharing the model update, which may however consume too much time. Alternatively, following the standard FL settings [1]–[3], we run the local training algorithm (i.e., Adam for A-SFADMM and D-SFADMM and SGD for A-SGD) for a few iterations before uploading each model. The number of local iterations is critical in ensuring convergence and achieving high accuracy. As shown in Fig. 6, with 5 local iterations both A-SFADMM and D-SFADMM suffer from low accuracy, while A-FADMM even struggles with oscillation. With 20 local iterations, we observe that both A-SFADMM and D-SFADMM achieve not only convergence but also the highest accuracy. Optimizing the number of local iterations is intertwined with learning rate, mini-batch size, and communication channels. This interesting-but-challenging problem is deferred to future work.

c) Impact of Learning Rates: In Sec. V, we use the local optimizer’s (Adam or SGD) learning rate 0.01. Here, we additionally test the learning rate 0.001 under $\rho = 5$. As shown in Fig. 7, for A-SFADMM, the learning rate change does not affect the convergence speed and accuracy significantly. By contrast, for D-SFADMM, the learning rate 0.01 leads to

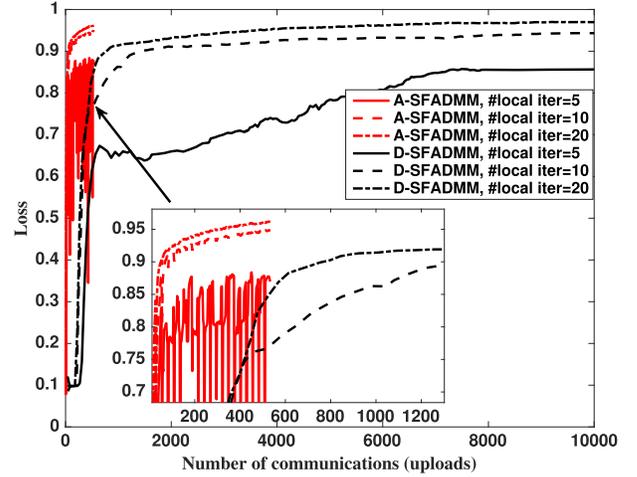


Fig. 6. Impact of the *number of local iterations* in image classification using DNNs.

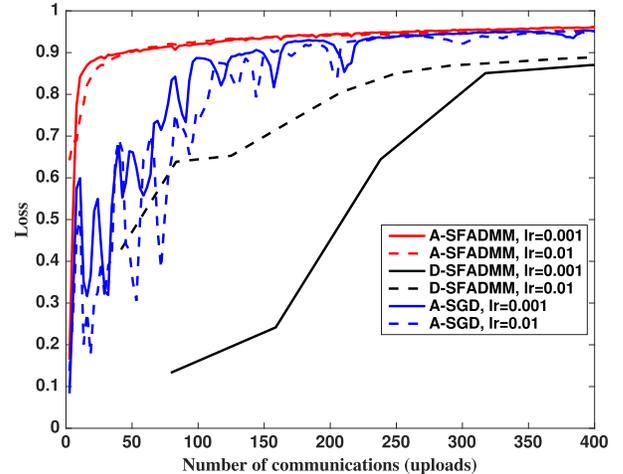


Fig. 7. Impact of the *learning rate* in image classification using DNNs.

faster convergence, while for A-SGD, the learning rate 0.001 yields less oscillation.

VI. CONCLUSION

In this article, we proposed A-FADMM, and proved its theoretical convergence and privacy guarantees, while validating its effectiveness in convex and non-convex problems. To further improve the applicability, we conclude this article by addressing several practical issues and possible extensions.

- **Asynchronous Transmissions:** Analog over-the-air aggregation is sensitive to asynchronous signal transmissions as both early and delayed arrivals incur additional noise at reception. To alleviate this problem, it is possible to turn the less communication rounds of A-FADMM into longer transmission time of each worker, increasing the signal overlapping duration compared to the out-of-synch duration.
- **Large Models:** To convey large models using analog signals, model compression methods should be re-designed. Applying compressive sensing techniques is promising,

in which a sparsified update is encoded by multiplying a random matrix before transmission [33], and the received update is decoded using the approximate message passing (AMP) algorithm [34].

- **Decentralized Architecture:** Workers have limited transmit energy, and hence faraway workers are difficult to reach PS [35], [36], hindering the wide-area coverage of A-FADMM. It could be therefore interesting to study the decentralized version of A-FADMM in which every worker communicates only with neighbors while taking into account their time-varying network topologies.

VII. APPENDICES

A. Proof of Lemma 1

To prove the statement of the lemma, we will proceed by proving the following two statements

- (i) The upper bound on the optimality gap is given as

$$\sum_{n=1}^N [f_n(\boldsymbol{\theta}_n^{k+1}) - f_n(\boldsymbol{\theta}_n^*)] \leq - \sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^{k+1} \mathbf{r}_{n,i}^{k+1} + \sum_{n=1}^N \sum_{i=1}^d \mathbf{S}_{n,i}^{k+1} (\boldsymbol{\theta}_{n,i}^* - \boldsymbol{\theta}_{n,i}^{k+1}), \quad (33)$$

where $\boldsymbol{\mu}_{n,i} = \boldsymbol{\lambda}_{n,i}^* h_{n,i}$.

- (ii) The lower bound on the optimality gap is given as

$$\sum_{n=1}^N [f_n(\boldsymbol{\theta}_n^{k+1}) - f_n(\boldsymbol{\theta}^*)] \geq - \sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^* \mathbf{r}_{n,i}^{k+1}. \quad (34)$$

Proof of statement (i): We note that $f_n(\boldsymbol{\theta}_n)$ for all n is closed, proper, and convex, hence \mathcal{L}_ρ is sub-differentiable. Since $\boldsymbol{\theta}_n^{k+1}$ minimizes $\mathcal{L}_\rho(\boldsymbol{\theta}_n, \boldsymbol{\Theta}^k, \boldsymbol{\lambda}_n^k)$, the following must hold true at each iteration $k+1$

$$0 \in \partial_i f_n(\boldsymbol{\theta}_n^{k+1}) + \boldsymbol{\mu}_{n,i}^k + \rho |h_{n,i}^{k+1}|^2 \boldsymbol{\theta}_{n,i}^{k+1} - \rho |h_{n,i}^{k+1}|^2 \boldsymbol{\Theta}_i^k. \quad (35)$$

Note that when $h_{n,i}^{k+1} \neq h_{n,i}^k$, we choose $\boldsymbol{\theta}_{n,i}^{k+1} = \boldsymbol{\theta}_{n,i}^k$, and under this choice, $\boldsymbol{\theta}_{n,i}^{k+1}$ is still the minimizer of $\mathcal{L}_\rho(\boldsymbol{\theta}_n, \boldsymbol{\Theta}^k, \boldsymbol{\lambda}_n^k)$ since $\boldsymbol{\lambda}_{n,i}^k$ should have been calculated to satisfy (35) given $\boldsymbol{\theta}_{n,i}^k$ when there is change in the channel.

Adding and subtracting the term $\rho |h_{n,i}^{k+1}|^2 \boldsymbol{\Theta}_i^{k+1}$ and re-arranging the terms, we can write

$$0 \in \partial_i f_n(\boldsymbol{\theta}_n^{k+1}) + \boldsymbol{\mu}_{n,i}^k + \rho |h_{n,i}^{k+1}|^2 (\boldsymbol{\theta}_{n,i}^{k+1} - \boldsymbol{\Theta}_i^{k+1}) + \rho |h_{n,i}^{k+1}|^2 (\boldsymbol{\Theta}_i^{k+1} - \boldsymbol{\Theta}_i^k). \quad (36)$$

Using the definitions of $\mathbf{r}_{n,i}^{k+1}$ and $\mathbf{S}_{n,i}^{k+1}$ as well as the update of $\boldsymbol{\mu}_{n,i}^{k+1}$ given in Eq. (18), we obtain

$$0 \in \partial_i f_n(\boldsymbol{\theta}_n^{k+1}) + \boldsymbol{\mu}_{n,i}^{k+1} + \mathbf{S}_{n,i}^{k+1}. \quad (37)$$

The result in (37) implies that $\boldsymbol{\theta}_n^{k+1}$ minimizes the following convex objective function

$$f_n(\boldsymbol{\theta}_n) + \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^{k+1} \boldsymbol{\theta}_{n,i} + \sum_{i=1}^d \mathbf{S}_{n,i}^{k+1} \boldsymbol{\theta}_{n,i}. \quad (38)$$

Next, since $\boldsymbol{\theta}_n^{k+1}$ is the minimizer of (38), then, it holds that

$$f_n(\boldsymbol{\theta}_n^{k+1}) + \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^{k+1} \boldsymbol{\theta}_{n,i}^{k+1} + \sum_{i=1}^d \mathbf{S}_{n,i}^{k+1} \boldsymbol{\theta}_{n,i}^{k+1} \leq f_n(\boldsymbol{\theta}_n^*) + \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^{k+1} \boldsymbol{\theta}_{n,i}^* + \sum_{i=1}^d \mathbf{S}_{n,i}^{k+1} \boldsymbol{\theta}_{n,i}^*, \quad (39)$$

where $\boldsymbol{\theta}^*$ is the optimal value of the problem in (6)-(7). Summing over all workers yields

$$\sum_{n=1}^N f_n(\boldsymbol{\theta}_n^{k+1}) + \sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^{k+1} \boldsymbol{\theta}_{n,i}^{k+1} + \sum_{n=1}^N \sum_{i=1}^d \mathbf{S}_{n,i}^{k+1} \boldsymbol{\theta}_{n,i}^{k+1} \leq \sum_{n=1}^N f_n(\boldsymbol{\theta}_n^*) + \sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^{k+1} \boldsymbol{\theta}_{n,i}^* + \sum_{n=1}^N \sum_{i=1}^d \mathbf{S}_{n,i}^{k+1} \boldsymbol{\theta}_{n,i}^* \quad (40)$$

Similarly, $\boldsymbol{\Theta}_i^{k+1}$ satisfies

$$0 = - \sum_{n=1}^N \boldsymbol{\mu}_{n,i}^k + \rho \sum_{n=1}^N |h_{n,i}^{k+1}|^2 \boldsymbol{\Theta}_i^{k+1} - \rho \sum_{n=1}^N |h_{n,i}^{k+1}|^2 \boldsymbol{\theta}_{n,i}^{k+1}. \quad (41)$$

Using the update of $\boldsymbol{\mu}_{n,i}^{k+1}$, we deduce that $\boldsymbol{\Theta}_i^{k+1}$ minimizes $-\sum_{n=1}^N \boldsymbol{\mu}_{n,i}^{k+1} \boldsymbol{\Theta}_i$, and therefore, we can write

$$- \sum_{n=1}^N \boldsymbol{\mu}_{n,i}^{k+1} \boldsymbol{\Theta}_i^{k+1} \leq - \sum_{n=1}^N \boldsymbol{\mu}_{n,i}^{k+1} \boldsymbol{\Theta}_i^*. \quad (42)$$

Summing over all i yields

$$- \sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^{k+1} \boldsymbol{\Theta}_i^{k+1} \leq - \sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^{k+1} \boldsymbol{\Theta}_i^*. \quad (43)$$

Adding (40) and (43), we get

$$\begin{aligned} & \sum_{n=1}^N f_n(\boldsymbol{\theta}_n^{k+1}) + \sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^{k+1} \boldsymbol{\theta}_{n,i}^{k+1} \\ & + \sum_{n=1}^N \sum_{i=1}^d \mathbf{S}_{n,i}^{k+1} \boldsymbol{\theta}_{n,i}^{k+1} - \sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^{k+1} \boldsymbol{\Theta}_i^{k+1} \\ & \leq \sum_{n=1}^N f_n(\boldsymbol{\theta}_n^*) + \sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^{k+1} \boldsymbol{\theta}_{n,i}^* + \sum_{n=1}^N \sum_{i=1}^d \mathbf{S}_{n,i}^{k+1} \boldsymbol{\theta}_{n,i}^* \\ & - \sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^{k+1} \boldsymbol{\Theta}_i^*. \end{aligned} \quad (44)$$

After rearranging the terms, we get

$$\begin{aligned} & \sum_{n=1}^N [f_n(\boldsymbol{\theta}_n^{k+1}) - f_n(\boldsymbol{\theta}_n^*)] \\ & \leq - \sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^{k+1} (\boldsymbol{\theta}_{n,i}^{k+1} - \boldsymbol{\Theta}_i^{k+1}) \\ & + \sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^{k+1} (\boldsymbol{\theta}_{n,i}^* - \boldsymbol{\Theta}_i^*) \\ & + \sum_{n=1}^N \sum_{i=1}^d \mathbf{S}_{n,i}^{k+1} (\boldsymbol{\theta}_{n,i}^* - \boldsymbol{\theta}_{n,i}^{k+1}). \end{aligned} \quad (45)$$

Using $\mathbf{r}_{n,i}^{k+1} = \boldsymbol{\theta}_{n,i}^{k+1} - \boldsymbol{\Theta}_i^{k+1}$, and $\mathbf{r}_{n,i}^* = \boldsymbol{\theta}_{n,i}^* - \boldsymbol{\Theta}_i^* = 0$ gives

$$\begin{aligned} \sum_{n=1}^N \left[f_n(\boldsymbol{\theta}_n^{k+1}) - f_n(\boldsymbol{\theta}_n^*) \right] &\leq - \sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^{k+1} \mathbf{r}_{n,i}^{k+1} \\ &\quad + \sum_{n=1}^N \sum_{i=1}^d \mathbf{S}_{n,i}^{k+1} (\boldsymbol{\theta}_{n,i}^* - \boldsymbol{\theta}_{n,i}^{k+1}). \end{aligned} \quad (46)$$

and hence we have proved the statement (i).

Proof of statement (ii):

We note that for a saddle point $(\boldsymbol{\Theta}^*, \boldsymbol{\theta}^*, \{\boldsymbol{\lambda}_n^*\}_n)$ of $\mathcal{L}_0(\boldsymbol{\Theta}^*, \{\boldsymbol{\theta}_n\}_n, \{\boldsymbol{\lambda}_n\}_n)$, it holds that, for all n , we have

$$\mathcal{L}_0(\boldsymbol{\Theta}^*, \boldsymbol{\theta}^*, \{\boldsymbol{\lambda}_n^*\}_n) \leq \mathcal{L}_0(\boldsymbol{\Theta}^{k+1}, \{\boldsymbol{\theta}_n^{k+1}\}_n, \{\boldsymbol{\lambda}_n^*\}_n). \quad (47)$$

Substituting the expression for the Lagrangian from (8) on the both sides of (47), we get

$$\begin{aligned} \sum_{n=1}^N f_n(\boldsymbol{\theta}^*) + \sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^* (\boldsymbol{\theta}_{n,i}^* - \boldsymbol{\Theta}_i^*) \\ \leq \sum_{n=1}^N f_n(\boldsymbol{\theta}_n^{k+1}) + \sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^* (\boldsymbol{\theta}_{n,i}^{k+1} - \boldsymbol{\Theta}_i^{k+1}). \end{aligned} \quad (48)$$

Using $\mathbf{r}_{n,i}^{k+1} = \boldsymbol{\theta}_{n,i}^{k+1} - \boldsymbol{\Theta}_i^{k+1}$, and $\mathbf{r}_{n,i}^* = \boldsymbol{\theta}_{n,i}^* - \boldsymbol{\Theta}_i^* = 0$ gives

$$\sum_{n=1}^N \left[f_n(\boldsymbol{\theta}_n^{k+1}) - f_n(\boldsymbol{\theta}^*) \right] \geq - \sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^* \mathbf{r}_{n,i}^{k+1}. \quad (49)$$

which proves the statement (ii).

Finally, combining the statements (i) and (ii) completes the proof.

B. Proof of Theorem 1

The proof relies on using the lower and upper bounds derived in Lemma 1 to show the decrease in the optimality gap. To this end, we start by multiplying both Eqs. (46) and (49) by 2, and then add them up to get

$$\begin{aligned} 2 \sum_{n=1}^N \sum_{i=1}^d (\boldsymbol{\mu}_{n,i}^{k+1} - \boldsymbol{\mu}_{n,i}^*) \mathbf{r}_{n,i}^{k+1} \\ + 2 \sum_{n=1}^N \sum_{i=1}^d \mathbf{S}_{n,i}^{k+1} (\boldsymbol{\theta}_{n,i}^{k+1} - \boldsymbol{\theta}_{n,i}^*) \leq 0 \end{aligned} \quad (50)$$

Since $\boldsymbol{\mu}_{n,i}^{k+1} = \boldsymbol{\mu}_{n,i}^k + \rho |h_{n,i}^{k+1}|^2 \mathbf{r}_{n,i}^{k+1}$, then the first term can be re-written as

$$\begin{aligned} 2 \sum_{n=1}^N \sum_{i=1}^d (\boldsymbol{\mu}_{n,i}^{k+1} - \boldsymbol{\mu}_{n,i}^*) \mathbf{r}_{n,i}^{k+1} \\ = 2 \sum_{n=1}^N \sum_{i=1}^d (\boldsymbol{\mu}_{n,i}^k - \boldsymbol{\mu}_{n,i}^*) \mathbf{r}_{n,i}^{k+1} \\ + 2\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\mathbf{r}_{n,i}^{k+1})^2. \end{aligned} \quad (51)$$

Since $\mathbf{r}_{n,i}^{k+1} = \frac{1}{\rho |h_{n,i}^{k+1}|^2} (\boldsymbol{\mu}_{n,i}^{k+1} - \boldsymbol{\mu}_{n,i}^k)$, we can write

$$\begin{aligned} 2 \sum_{n=1}^N \sum_{i=1}^d (\boldsymbol{\mu}_{n,i}^{k+1} - \boldsymbol{\mu}_{n,i}^*) \mathbf{r}_{n,i}^{k+1} \\ = \frac{2}{\rho} \sum_{n=1}^N \sum_{i=1}^d \frac{1}{|h_{n,i}^{k+1}|^2} (\boldsymbol{\mu}_{n,i}^k - \boldsymbol{\mu}_{n,i}^*) (\boldsymbol{\mu}_{n,i}^{k+1} - \boldsymbol{\mu}_{n,i}^k) \\ + 2\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\mathbf{r}_{n,i}^{k+1})^2. \end{aligned} \quad (52)$$

Using the fact that $\boldsymbol{\mu}_{n,i}^{k+1} - \boldsymbol{\mu}_{n,i}^k = \boldsymbol{\mu}_{n,i}^{k+1} - \boldsymbol{\mu}_{n,i}^* + \boldsymbol{\mu}_{n,i}^* - \boldsymbol{\mu}_{n,i}^k$, we get

$$\begin{aligned} \frac{1}{|h_{n,i}^{k+1}|^2} (\boldsymbol{\mu}_{n,i}^k - \boldsymbol{\mu}_{n,i}^*) (\boldsymbol{\mu}_{n,i}^{k+1} - \boldsymbol{\mu}_{n,i}^k) = \frac{1}{|h_{n,i}^{k+1}|^2} \\ (\boldsymbol{\mu}_{n,i}^k - \boldsymbol{\mu}_{n,i}^*) (\boldsymbol{\mu}_{n,i}^{k+1} - \boldsymbol{\mu}_{n,i}^*) - \frac{(\boldsymbol{\mu}_{n,i}^k - \boldsymbol{\mu}_{n,i}^*)^2}{|h_{n,i}^{k+1}|^2}. \end{aligned} \quad (53)$$

Now, let's re-write $\rho |h_{n,i}^{k+1}|^2 (\mathbf{r}_{n,i}^{k+1})^2$, using $\mathbf{r}_{n,i}^{k+1} = \frac{1}{\rho |h_{n,i}^{k+1}|^2} (\boldsymbol{\mu}_{n,i}^{k+1} - \boldsymbol{\mu}_{n,i}^k)$, as

$$\begin{aligned} \rho |h_{n,i}^{k+1}|^2 (\mathbf{r}_{n,i}^{k+1})^2 &= \frac{1}{\rho |h_{n,i}^{k+1}|^2} (\boldsymbol{\mu}_{n,i}^{k+1} - \boldsymbol{\mu}_{n,i}^k)^2 \\ &= \frac{1}{\rho |h_{n,i}^{k+1}|^2} (\boldsymbol{\mu}_{n,i}^{k+1} - \boldsymbol{\mu}_{n,i}^*)^2 \\ &\quad + \frac{1}{\rho |h_{n,i}^{k+1}|^2} (\boldsymbol{\mu}_{n,i}^k - \boldsymbol{\mu}_{n,i}^*)^2 - \frac{2}{\rho |h_{n,i}^{k+1}|^2} \\ &\quad \times (\boldsymbol{\mu}_{n,i}^k - \boldsymbol{\mu}_{n,i}^*) (\boldsymbol{\mu}_{n,i}^{k+1} - \boldsymbol{\mu}_{n,i}^*), \end{aligned} \quad (54)$$

where we have used that $\boldsymbol{\mu}_{n,i}^{k+1} - \boldsymbol{\mu}_{n,i}^k = \boldsymbol{\mu}_{n,i}^{k+1} - \boldsymbol{\mu}_{n,i}^* + \boldsymbol{\mu}_{n,i}^* - \boldsymbol{\mu}_{n,i}^k$. Going back to Eq. (51), we can write

$$\begin{aligned} 2 \sum_{n=1}^N \sum_{i=1}^d (\boldsymbol{\mu}_{n,i}^{k+1} - \boldsymbol{\mu}_{n,i}^*) \mathbf{r}_{n,i}^{k+1} \\ = \frac{1}{\rho} \sum_{n=1}^N \sum_{i=1}^d \frac{1}{|h_{n,i}^{k+1}|^2} (\boldsymbol{\mu}_{n,i}^{k+1} - \boldsymbol{\mu}_{n,i}^*)^2 \\ - \frac{1}{\rho} \sum_{n=1}^N \sum_{i=1}^d \frac{1}{|h_{n,i}^{k+1}|^2} (\boldsymbol{\mu}_{n,i}^k - \boldsymbol{\mu}_{n,i}^*)^2 \\ + \rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\mathbf{r}_{n,i}^{k+1})^2. \end{aligned} \quad (55)$$

Now, let's examine the second term of Eq. (50)

$$\begin{aligned} 2 \sum_{n=1}^N \sum_{i=1}^d \mathbf{S}_{n,i}^{k+1} (\boldsymbol{\theta}_{n,i}^{k+1} - \boldsymbol{\theta}_{n,i}^*) \\ = 2\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\boldsymbol{\Theta}_i^{k+1} - \boldsymbol{\Theta}_i^k) (\boldsymbol{\theta}_{n,i}^{k+1} - \boldsymbol{\theta}_{n,i}^*) \\ = 2\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\boldsymbol{\Theta}_i^{k+1} - \boldsymbol{\Theta}_i^k) \mathbf{r}_{n,i}^{k+1} \\ + 2\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\boldsymbol{\Theta}_i^{k+1} - \boldsymbol{\Theta}_i^k) (\boldsymbol{\Theta}_i^{k+1} - \boldsymbol{\theta}_{n,i}^*). \end{aligned} \quad (56)$$

Using $\Theta_i^{k+1} - \theta_{n,i}^* = \Theta_i^{k+1} - \Theta_i^k + \Theta_i^k - \theta_{n,i}^*$, we can write

$$\begin{aligned} & 2 \sum_{n=1}^N \sum_{i=1}^d \mathbf{S}_{n,i}^{k+1} (\theta_{n,i}^{k+1} - \theta_{n,i}^*) \\ &= 2\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^{k+1} - \Theta_i^k) \mathbf{r}_{n,i}^{k+1} \\ & \quad + 2\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^{k+1} - \Theta_i^k)^2 \\ & \quad + 2\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^{k+1} - \Theta_i^k) (\Theta_i^k - \theta_{n,i}^*). \end{aligned} \quad (57)$$

Since $\Theta_i^{k+1} - \Theta_i^k = \Theta_i^{k+1} - \theta_{n,i}^* + \theta_{n,i}^* - \Theta_i^k$, then we get

$$\begin{aligned} & 2 \sum_{n=1}^N \sum_{i=1}^d \mathbf{S}_{n,i}^{k+1} (\theta_{n,i}^{k+1} - \theta_{n,i}^*) \\ &= 2\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^{k+1} - \Theta_i^k) \mathbf{r}_{n,i}^{k+1} \\ & \quad + 2\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^{k+1} - \Theta_i^k)^2 \\ & \quad - 2\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^k - \theta_{n,i}^*)^2 \\ & \quad + 2\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^{k+1} - \theta_{n,i}^*) (\Theta_i^k - \theta_{n,i}^*), \end{aligned} \quad (58)$$

Now, let's focus the second term of Eq. (58). Using $\Theta_i^{k+1} - \Theta_i^k = \Theta_i^{k+1} - \theta_{n,i}^* + \theta_{n,i}^* - \Theta_i^k$, we can write

$$\begin{aligned} & \rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^{k+1} - \Theta_i^k)^2 \\ &= \rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^{k+1} - \theta_{n,i}^*)^2 \\ & \quad + \rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^k - \theta_{n,i}^*)^2 \\ & \quad - 2\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^{k+1} - \theta_{n,i}^*) (\Theta_i^k - \theta_{n,i}^*). \end{aligned} \quad (59)$$

Replacing the last equation into Eq. (58), we can write

$$\begin{aligned} & 2 \sum_{n=1}^N \sum_{i=1}^d \mathbf{S}_{n,i}^{k+1} (\theta_{n,i}^{k+1} - \theta_{n,i}^*) \\ &= 2\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^{k+1} - \Theta_i^k) \mathbf{r}_{n,i}^{k+1} \\ & \quad + \rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^{k+1} - \Theta_i^k)^2 \\ & \quad + \rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^{k+1} - \theta_{n,i}^*)^2 \\ & \quad - \rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^k - \theta_{n,i}^*)^2. \end{aligned} \quad (60)$$

Using Eqs. (55) and (60) in (50), we get

$$\begin{aligned} & \frac{1}{\rho} \sum_{n=1}^N \sum_{i=1}^d \frac{(\mu_{n,i}^{k+1} - \mu_{n,i}^*)^2}{|h_{n,i}^{k+1}|^2} \\ & \quad - \frac{1}{\rho} \sum_{n=1}^N \sum_{i=1}^d \frac{(\mu_{n,i}^k - \mu_{n,i}^*)^2}{|h_{n,i}^{k+1}|^2} \\ & \quad + \rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^{k+1} - \Theta_i^k)^2 \\ & \quad + \rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\mathbf{r}_{n,i}^{k+1})^2 \\ & \quad + \rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^{k+1} - \Theta_i^*)^2 \\ & \quad - \rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^k - \Theta_i^*)^2 \\ & \quad + 2\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^{k+1} - \Theta_i^k) \mathbf{r}_{n,i}^{k+1} \leq 0 \end{aligned} \quad (61)$$

Defining the sequence that measures the difference in the optimality gap between iterations $k+1$ and k as

$$\begin{aligned} W^{k+1} &= \frac{1}{\rho} \sum_{n=1}^N \sum_{i=1}^d \frac{(\mu_{n,i}^{k+1} - \mu_{n,i}^*)^2}{|h_{n,i}^{k+1}|^2} \\ & \quad - \frac{1}{\rho} \sum_{n=1}^N \sum_{i=1}^d \frac{(\mu_{n,i}^k - \mu_{n,i}^*)^2}{|h_{n,i}^{k+1}|^2} \\ & \quad + \rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^{k+1} - \Theta_i^*)^2 \\ & \quad - \rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 (\Theta_i^k - \Theta_i^*)^2, \end{aligned} \quad (62)$$

then, Eq. (61) can be re-written as

$$W^{k+1} \leq -\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}^{k+1}|^2 \left(\mathbf{r}_{n,i}^{k+1} + \Theta_i^{k+1} - \Theta_i^k \right)^2. \quad (63)$$

From (63), we note that W^{k+1} is negative since ρ is a positive constant ($\rho > 0$) and the sum square term is positive. Hence, the optimality gap at iteration $k+1$ is non-increasing, which completes the proof.

C. Proof of Corollary 1

Using the derivations made in Theorem 1, we can further show that, in the case of static channel, i.e. $h_{n,i}^{k+1} = h_{n,i}$, both the primal and dual residuals converge to zero, i.e. $\lim_{k \rightarrow \infty} \mathbf{r}_{n,i}^{k+1} = 0$ and $\lim_{k \rightarrow \infty} \mathbf{S}_{n,i}^{k+1} = 0$. To this end, we start defining the Lyapunov function

$$V^k = \frac{1}{\rho} \sum_{n=1}^N \sum_{i=1}^d \frac{(\boldsymbol{\mu}_{n,i}^k - \boldsymbol{\mu}_{n,i}^*)^2}{|h_{n,i}|^2} + \rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}|^2 (\boldsymbol{\Theta}_i^k - \boldsymbol{\Theta}_i^*)^2. \quad (64)$$

We can re-write (63) as

$$\begin{aligned} V^{k+1} - V^k &\leq -\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}|^2 (\mathbf{r}_{n,i}^{k+1})^2 \\ &\quad - \rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}|^2 \left(\boldsymbol{\Theta}_i^{k+1} - \boldsymbol{\Theta}_i^k \right)^2 \\ &\quad - 2\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}|^2 \mathbf{r}_{n,i}^{k+1} \left(\boldsymbol{\Theta}_i^{k+1} - \boldsymbol{\Theta}_i^k \right) \end{aligned} \quad (65)$$

Since $\boldsymbol{\Theta}_i^{k+1}$ minimizes $-\sum_{n=1}^N \boldsymbol{\mu}_{n,i}^{k+1} \boldsymbol{\Theta}_i$, and $\boldsymbol{\Theta}_i^k$ minimizes $-\sum_{n=1}^N \boldsymbol{\mu}_{n,i}^k \boldsymbol{\Theta}_i$, then, after summing over i in both sides of each equation, we can write

$$-\sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^{k+1} \boldsymbol{\Theta}_i^{k+1} \leq -\sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^{k+1} \boldsymbol{\Theta}_i^k, \quad (66)$$

$$-\sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^k \boldsymbol{\Theta}_i^k \leq -\sum_{n=1}^N \sum_{i=1}^d \boldsymbol{\mu}_{n,i}^k \boldsymbol{\Theta}_i^{k+1}. \quad (67)$$

Adding Eqs. (66) and (67), we get

$$\sum_{n=1}^N \sum_{i=1}^d (\boldsymbol{\mu}_{n,i}^{k+1} - \boldsymbol{\mu}_{n,i}^k) (\boldsymbol{\Theta}_i^{k+1} - \boldsymbol{\Theta}_i^k) \geq 0. \quad (68)$$

Since $\boldsymbol{\mu}_{n,i}^{k+1} = \boldsymbol{\mu}_{n,i}^k + \rho |h_{n,i}|^2 \mathbf{r}_{n,i}^{k+1}$, then we get

$$2\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}|^2 (\boldsymbol{\Theta}_i^{k+1} - \boldsymbol{\Theta}_i^k) \mathbf{r}_{n,i}^{k+1} \geq 0. \quad (69)$$

Thus, using Eq. (65), and summing over the iterations from $k = 1, \dots, K$, we get

$$\begin{aligned} \sum_{k=0}^K \left[\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}|^2 (\mathbf{r}_{n,i}^{k+1})^2 \right. \\ \left. + \rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}|^2 (\boldsymbol{\Theta}_i^{k+1} - \boldsymbol{\Theta}_i^k)^2 \right. \\ \left. + 2\rho \sum_{n=1}^N \sum_{i=1}^d |h_{n,i}|^2 (\boldsymbol{\Theta}_i^{k+1} - \boldsymbol{\Theta}_i^k) \mathbf{r}_{n,i}^{k+1} \right] \leq V^0. \quad (70) \end{aligned}$$

Taking the limit as $K \rightarrow \infty$, and using the fact that the terms of the series on the left hand-side are positive, we obtain that the primal and dual residuals goes to zero as $k \rightarrow \infty$, i.e.

$\lim_{k \rightarrow \infty} \mathbf{r}_{n,i}^{k+1} = \mathbf{0}$ and $\lim_{k \rightarrow \infty} \mathbf{S}_{n,i}^{k+1} = \mathbf{0}$. Using the upper and lower bounds, (20) and (21), derived in Lemma 1 and the fact that both the primal and dual residuals goes to zero as $k \rightarrow \infty$, we get that the optimal gap also goes to zero as $k \rightarrow \infty$, i.e., $\lim_{k \rightarrow \infty} \sum_{n=1}^N f_n(\boldsymbol{\theta}_n^k) = \sum_{n=1}^N f_n(\boldsymbol{\theta}_n^*)$, finalizing the proof.

REFERENCES

- [1] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. Y. Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. Artif. Intell. Statist.*, Fort Lauderdale, FL, USA, Apr. 2017, pp. 1273–1282.
- [2] J. Konecny, H. B. McMahan, F. X. Yu, P. Richtarik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," in *Proc. NIPS Workshop (PMPML)*, Barcelona, Spain, Dec. 2016. [Online]. Available: <https://arxiv.org/abs/1610.05492>
- [3] P. Kairouz *et al.*, "Advances and open problems in federated learning," 2019, *arXiv:1912.04977*. [Online]. Available: <http://arxiv.org/abs/1912.04977>
- [4] J. Park, S. Samarakoon, M. Bennis, and M. Debbah, "Wireless network intelligence at the edge," *Proc. IEEE*, vol. 107, no. 11, pp. 2204–2239, Nov. 2019.
- [5] T. Nishio and R. Yonetani, "Client selection for federated learning with heterogeneous resources in mobile edge," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Shanghai, China, May 2019, pp. 1–7. [Online]. Available: <http://arxiv.org/abs/1804.08333>
- [6] S. Wang *et al.*, "Adaptive federated learning in resource constrained edge computing systems," *IEEE J. Sel. Areas Commun.*, vol. 37, no. 6, pp. 1205–1221, Jun. 2019.
- [7] H. H. Yang, Z. Liu, T. Q. S. Quek, and H. V. Poor, "Scheduling policies for federated learning in wireless networks," *IEEE Trans. Commun.*, vol. 68, no. 1, pp. 317–333, Jan. 2020.
- [8] M. Chen, Z. Yang, W. Saad, C. Yin, H. V. Poor, and S. Cui, "A joint learning and communications framework for federated learning over wireless networks," *IEEE Trans. Wireless Commun.*, vol. 20, no. 1, pp. 269–283, Jan. 2021.
- [9] Y. Sun, S. Zhou, and D. Gündüz, "Energy-aware analog aggregation for federated learning with redundant data," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2020, pp. 1–7.
- [10] Z. Yang, M. Chen, W. Saad, C. S. Hong, and M. Shikh-Bahaei, "Energy efficient federated learning over wireless communication networks," *IEEE Trans. Wireless Commun.*, vol. 20, no. 3, pp. 1935–1949, Mar. 2021.
- [11] 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*, Apr. 2019, pp. 1387–1395.
- [12] M. M. Amiri and D. Gündüz, "Over-the-air machine learning at the wireless edge," in *Proc. IEEE 20th Int. Workshop Signal Process. Adv. Wireless Commun. (SPWAC)*, Cannes, France, Jul. 2019, pp. 1–5.
- [13] G. Zhu, Y. Wang, and K. Huang, "Broadband analog aggregation for low-latency federated edge learning," *IEEE Trans. Wireless Commun.*, vol. 19, no. 1, pp. 491–506, Jan. 2020.
- [14] T. Sery and K. Cohen, "On analog gradient descent learning over multiple access fading channels," *IEEE Trans. Signal Process.*, vol. 68, pp. 2897–2911, 2020.
- [15] G. Zhu, Y. Du, D. Gunduz, and K. Huang, "One-bit over-the-air aggregation for communication-efficient federated edge learning: Design and convergence analysis," *IEEE Trans. Wireless Commun.*, vol. 20, no. 3, pp. 2120–2135, Mar. 2021.
- [16] J. Park *et al.*, "Communication-efficient and distributed learning over wireless networks: Principles and applications," *Proc. IEEE*, vol. 109, no. 5, pp. 796–819, May 2021.
- [17] M. Fredrikson, S. Jha, and T. Ristenpart, "Model inversion attacks that exploit confidence information and basic countermeasures," in *Proc. 22nd ACM SIGSAC Conf. Comput. Commun. Secur.* New York, NY, USA: Association for Computing Machinery, Oct. 2015, pp. 1322–1333, doi: [10.1145/2810103.2813677](https://doi.org/10.1145/2810103.2813677).
- [18] B. Hitaj, G. Ateniese, and F. Perez-Cruz, "Deep models under the GAN: Information leakage from collaborative deep learning," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.* New York, NY, USA: Association for Computing Machinery, Oct. 2017, pp. 603–618, doi: [10.1145/3133956.3134012](https://doi.org/10.1145/3133956.3134012).
- [19] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, Jan. 2011.
- [20] W. Deng, M.-J. Lai, Z. Peng, and W. Yin, "Parallel multi-block ADMM with $\mathcal{O}(1/k)$ convergence," *J. Sci. Comput.*, vol. 71, no. 2, pp. 712–736, May 2017.

- [21] R. Glowinski and A. Marroco, "Sur l'approximation, par éléments finis d'ordre un, *et al.* résolution, par pénalisation-dualité d'une classe de problèmes de dirichlet non linéaires," *ESAIM, Math. Model. Numer. Anal.-Modélisation Mathématique et Analyse Numérique*, vol. 9, no. R2, pp. 41–76, 1975.
- [22] M. K. Ozdemir and H. Arslan, "Channel estimation for wireless OFDM systems," *IEEE Commun. Surveys Tuts.*, vol. 9, no. 2, pp. 18–48, 2nd Quart., 2007.
- [23] Q. Liu, S. Zhou, and G. B. Giannakis, "Cross-layer combining of adaptive modulation and coding with truncated ARQ over wireless links," *IEEE Trans. Wireless Commun.*, vol. 3, no. 5, pp. 1746–1755, Sep. 2004.
- [24] S. Catreux, V. Erceg, D. Gesbert, and R. W. Heath, "Adaptive modulation and MIMO coding for broadband wireless data networks," *IEEE Commun. Mag.*, vol. 40, no. 6, pp. 108–115, Jun. 2002.
- [25] *5G NR Physical Channels and Modulation*, document 3GPP TS 38.211, Jun. 2017.
- [26] I. Goodfellow *et al.*, "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 27, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, Eds. Red Hook, NY, USA: Curran Associates, 2014. [Online]. Available: <https://proceedings.neurips.cc/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf>
- [27] C. Zhang, M. Ahmad, and Y. Wang, "ADMM based privacy-preserving decentralized optimization," *IEEE Trans. Inf. Forensics Secur.*, vol. 14, no. 3, pp. 565–580, Mar. 2019.
- [28] K. Wei *et al.*, "Federated learning with differential privacy: Algorithms and performance analysis," *IEEE Trans. Inf. Forensics Secur.*, vol. 15, pp. 3454–3469, 2020.
- [29] Y. Zhao *et al.*, "Local differential privacy based federated learning for Internet of Things," *IEEE Internet Things J.*, early access, Nov. 10, 2021, doi: [10.1109/JIOT.2020.3037194](https://doi.org/10.1109/JIOT.2020.3037194).
- [30] L. Torgo. (2014). *Regression Datasets*. [Online]. Available: <https://www.dcc.fc.up.pt/~ltorgo/Regression/DataSets.html>
- [31] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
- [32] M. S. Elbamby *et al.*, "Wireless edge computing with latency and reliability guarantees," *Proc. IEEE*, vol. 107, no. 8, pp. 1717–1737, Aug. 2019.
- [33] M. M. Amiri and D. Gündüz, "Federated learning over wireless fading channels," *IEEE Trans. Wireless Commun.*, vol. 19, no. 5, pp. 3546–3557, May 2020.
- [34] D. L. Donoho, A. Maleki, and A. Montanari, "Message-passing algorithms for compressed sensing," *Proc. Nat. Acad. Sci. USA*, vol. 106, no. 45, pp. 18914–18919, Nov. 2009. [Online]. Available: <https://www.pnas.org/content/106/45/18914>
- [35] A. Elgabli, J. Park, A. S. Bedi, M. Bennis, and V. Aggarwal, "GADMM: Fast and communication efficient framework for distributed machine learning," *J. Mach. Learn. Res.*, vol. 21, no. 76, pp. 1–39, 2020.
- [36] C. B. Issaid, A. Elgabli, J. Park, M. Bennis, and M. Debbah, "Communication efficient distributed learning with censored, quantized, and generalized group ADMM," 2020, *arXiv:2009.06459*. [Online]. Available: <http://arxiv.org/abs/2009.06459>



Anis Elgabli received the B.Sc. degree in electrical and electronic engineering from the University of Tripoli, Libya, in 2004, the M.Eng. degree from UKM, Malaysia, in 2007, and the M.Sc. and Ph.D. degrees from the Department of Electrical and Computer Engineering, Purdue University, IN, USA, in 2015 and 2018, respectively. He is currently a Post-Doctoral Researcher with the Centre for Wireless Communications, University of Oulu. His main research interests include heterogeneous networks, radio resource management, vehicular communication,

video streaming, and distributed machine learning. He was a recipient of the Best Paper Award in HotSpot Workshop, 2018 (Infocom 2018) and the most JUFO points in 2020 at the Center of Wireless Communication, University of Oulu.



Jihong Park (Member, IEEE) received the B.S. and Ph.D. degrees from Yonsei University, Seoul, South Korea, in 2009 and 2016, respectively. He was a Post-Doctoral Researcher with Aalborg University, Denmark, from 2016 to 2017; and the University of Oulu, Finland, from 2018 to 2019. He is currently a Lecturer (an Assistant Professor) with the School of IT, Deakin University, Australia. His current research interests include communication-efficient distributed machine learning, distributed control, and distributed ledger technology, as well as their applications for beyond 5G/6G communication systems. He has served as a Conference/Workshop Program Committee Member for IEEE GLOBECOM, ICC, and WCNC, as well as NeurIPS, ICML, and IJCAI. He received the IEEE GLOBECOM Student Travel Grant in 2014, the IEEE Seoul Section Student Paper Contest Bronze Prize in 2014, and the 6th IDIS-ETNEWS (The Electronic Times) Paper Contest Award sponsored by the Ministry of Science, ICT, and Future Planning of Korea. He is an Associate Editor of *Frontiers in Data Science for Communications*, a Review Editor of *Frontiers in Aerial and Space Networks*, and a Guest Editor of the Special Issue on "Millimeter Wave Communications and Networking in 5G and Beyond" of *Telecom (MDPI)*.



Chaouki Ben Issaid received the Diplôme d'Ingénieur degree with majors in economics and financial engineering from the Ecole Polytechnique de Tunisie (EPT) in 2013 and the master's degree in applied mathematics and computational science (AMCS) and the Ph.D. degree in statistics from KAUST in 2015 and 2019, respectively. He is currently a Post-Doctoral Fellow with the Centre for Wireless Communications, University of Oulu. His current research interest includes communication-efficient distributed

machine learning.



Mehdi Bennis (Fellow, IEEE) is currently an Associate Professor with the Centre for Wireless Communications, University of Oulu, Finland, an Academy of Finland Research Fellow, and the Head of the Intelligent Connectivity and Networks/Systems Group (ICON). He has coauthored one book and published more than 200 research articles in international conferences, journals, and book chapters. His main research interests include radio resource management, heterogeneous networks, game theory, and machine learning in 5G networks and beyond.

He was a recipient of several prestigious awards, including the 2015 Fred W. Ellersick Prize from the IEEE Communications Society, the 2016 Best Tutorial Prize from the IEEE Communications Society, the 2017 EURASIP Best Paper Award for *Journal of Wireless Communications and Networks*, the all-University of Oulu Award for research, and the 2019 IEEE ComSoc Radio Communications Committee Early Achievement Award. He is also an Editor of IEEE TRANSACTIONS ON COMMUNICATIONS.