



Liu, D. and Simeone, O. (2022) Wireless federated Langevin Monte Carlo: repurposing channel noise for Bayesian sampling and privacy. *IEEE Transactions on Wireless Communications*, (doi: 10.1109/TWC.2022.3215663).

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<https://eprints.gla.ac.uk/284750/>

Deposited on: 4 November 2022

Enlighten – Research publications by members of the University of Glasgow  
<https://eprints.gla.ac.uk>

# Wireless Federated Langevin Monte Carlo: Repurposing Channel Noise for Bayesian Sampling and Privacy

Dongzhu Liu and Osvaldo Simeone

**Abstract**—Most works on federated learning (FL) focus on the most common frequentist formulation of learning whereby the goal is minimizing the global empirical loss. Frequentist learning, however, is known to be problematic in the regime of limited data as it fails to quantify epistemic uncertainty in prediction. Bayesian learning provides a principled solution to this problem by shifting the optimization domain to the space of distribution in the model parameters. This paper proposes a novel mechanism for the efficient implementation of Bayesian learning in wireless systems. Specifically, we focus on a standard gradient-based Markov Chain Monte Carlo (MCMC) method, namely Langevin Monte Carlo (LMC), and we introduce a novel protocol, termed Wireless Federated LMC (WFLMC), that is able to repurpose channel noise for the double role of seed randomness for MCMC sampling and of privacy preservation. To this end, based on the analysis of the Wasserstein distance between sample distribution and global posterior distribution under privacy and power constraints, we introduce a power allocation strategy as the solution of a convex program. The analysis identifies distinct operating regimes in which the performance of the system is power-limited, privacy-limited, or limited by the requirement of MCMC sampling. Both analytical and simulation results demonstrate that, if the channel noise is properly accounted for under suitable conditions, it can be fully repurposed for both MCMC sampling and privacy preservation, obtaining the same performance as in an ideal communication setting that is not subject to privacy constraints.

## I. INTRODUCTION

Federated learning (FL) protocols aim at coordinating multiple devices to collaboratively train a target model in a manner that approximates centralized learning at the cloud, while avoiding the direct exchange of data [1], [2]. Most prior works on wireless FL consider a *frequentist* formulation whose goal is minimizing the empirical loss over the vector of model parameters [3]–[7]. Significant attention has been devoted to uncoded transmission schemes coupled with non-orthogonal multiple access (NOMA), which leverage the superposition property of wireless channels to enable efficient over-the-air aggregation at the server [5], [6]. Furthermore, FL protocols inevitably leak some information about local data via communication. Formal privacy requirements can be met by introducing randomness to the disclosed statistics [8]. When implementing uncoded transmission, noise in wireless channels was accordingly shown to serve as a privacy-preserving mechanism [6], [9].

D. Liu is with the School of Computing Science, University of Glasgow, UK (email: dongzhu.liu@glasgow.ac.uk). O. Simeone is with Department of Engineering of Kings College London, UK (email: osvaldo.simeone@kcl.ac.uk). The work of O. Simeone was supported by European Research Council (ERC) under the European Unions Horizon 2020 Research and Innovation Programme (Grant Agreement No. 725731) and an Open Fellowship by the EPSRC.

Frequentist learning is effective in the regime of large data sets when accuracy is the main concern, but it fails to quantify *epistemic uncertainty* due to the availability of limited data [10], [11]. *Bayesian learning* provides an alternative learning framework in which optimization is done over the distribution of model parameters rather than over a single model parameter vector as in frequentist learning. Practical Bayesian learning methods include *variational inference* (VI), which constrains the model distribution to a parameter family, and *Monte Carlo* (MC) sampling, which draws samples approximately generated from the optimal model distribution [12].

This paper represents the first work on Bayesian FL in wireless networks. We specifically adopt *Langevin Monte Carlo* (LMC), a gradient-based *Markov Chain Monte Carlo* (MCMC) method that adds Gaussian noise to *gradient descent* (GD) updates. LMC is a fundamental building block of computationally efficient Bayesian inference and learning strategies. Unlike simpler random-walk MCMC methods, LMC leverages first-order information about the probabilistic model, striking a useful trade-off between complexity and performance [12], [13]. LMC can be generalized and improved in various directions, such as by accounting also for second-order information [14], [15].

The key contribution of this paper is not that of introducing a new Bayesian learning algorithm. Rather, we introduce a new mechanism for the efficient, and private, implementation of LMC over wireless channels. The approach is based on the idea that channel noise can be repurposed for the double role of seed randomness for the implementation of MC sampling and of privacy-preservation. Our analytical and experimental results provide insights about operating regimes in which channel noise can effectively serve both functions. It is envisaged that the proposed novel method of exploiting channel noise for MC sampling could also be applied and optimized for more sophisticated MCMC solutions such as Hamiltonian Monte Carlo [15].

### A. Related Work

1) *Frequentist and Bayesian FL*: FL protocols alternate between local computing and communication steps. In frequentist FL protocols, devices exchange model parameter vectors, which may be first quantized and compressed [16]. Wireless channels can be treated as bit pipes that limit the communication rates at each round [17], [18], or they can be modelled at the level of baseband signals. In the former case, resource allocation has been addressed to optimize the learning performance [19]. In the latter case, over-the-air aggregation can be leveraged via uncoded transmission – an approach

known as AirComp [5], [6], [20]. AirComp can be combined with sparsification and compression to reduce the communication overhead [21]. Bayesian post-processing estimation methods have been proposed to improve the test accuracy of federated learning, e.g., by exploiting the temporal structure of the received signals [21], by incorporating information about channel distribution and local prior [22], or by addressing data heterogeneity via knowledge distillation [23], [24]. Note that the schemes in [21], [22] do not implement Bayesian learning in the sense explained above of optimizing over a distribution in the model parameter space; while [23], [24] consider ideal communication and require either unlabelled data at the server [23] or additional communication overhead [24].

As discussed, Bayesian learning is, in practice, implemented by approximate methods – either VI or MC sampling. Both have been investigated only to a very limited extent for FL, even in the presence of ideal communication. VI-based methods are proposed in [25], [26] for noiseless communications based on parametric and particle-based representations of the model parameter distribution. Gradient-based MC methods are instead investigated in [27], [28], again under ideal communications.

2) *Private FL*: Differential privacy (DP) is a strong measure of information leakage that relates to the sensitivity of the disclosed statistics on individual data points in the training data set. In FL, a standard model is to assume the edge server to be “honest-but-curious”, requiring the implementation of DP-preserving mechanisms such as noise addition, subsampling random mini-batches, and random quantization of the gradients [29], [30]. Wireless FL can repurpose channel noise so as to ensure DP guarantees by controlling the signal-to-noise (SNR) ratio via transmit power optimization [9]. Furthermore, the superposition property of NOMA not only achieves efficient aggregation, but also amplifies the role of the channel noise as a privacy mechanism by protecting multiple devices’ transmissions simultaneously [31]. To enhance the convergence rate under the DP constraints, reference [6] proposes an optimized adaptive power control strategy that increases the effective SNR over the iterations. We note that the presence of channel noise can also benefit learning by accelerating the convergence for non-convex models [3], [32], or improving the generalization capability of convex models [4].

3) *Private Bayesian learning*: For Bayesian learning, the inherent randomness induced MC sampling automatically satisfies some level of DP requirements. Specifically, producing a single sample from the exact (or approximate) posterior distribution implements a differentially private strategy known as the exponential mechanism [33]. This result can be extended to multiple samples in gradient-based MCMC under proper conditions, such as small learning rate [33] or large scale model [34]. All prior work on private Bayesian learning is limited to centralized settings, and no prior result appears to have studied privacy in the context of Bayesian FL.

## B. Contributions and Organization

In this paper, we introduce a federated implementation of LMC in wireless systems whereby power allocation is optimized to control the SNR level so as to meet the requirement

of both MC sampling and DP. The main contributions and findings of the paper are summarized as follows.

- **Introducing Wireless Federated Langevin Monte Carlo (WFLMC)**: We first introduce Wireless Federated Langevin Monte Carlo (WFLMC), a novel iterative Bayesian learning protocol that relies on power control to repurpose channel noise for the double role of MC sampling via LMC and privacy preservation. WFLMC is based on uncoded transmission and NOMA, and goes beyond existing frequentist AirComp strategies by quantifying epistemic uncertainty through Bayesian learning.

- **Analyzing WFLMC**: Unlike frequentist learning, the goal of MC sampling in Bayesian learning is to ensure that the distribution of the produced samples is close to the global posterior distribution. Accordingly, we measure the learning performance via the 2-Wasserstein distance between the two distributions as in [35]. We provide analytical bounds on the 2-Wasserstein distance that comprise the contribution of the discretization error incurred by LMC, as well as of the gradient error due to channel noise and scheduling. We also present a DP privacy analysis of WFLMC that provides insights into the impact of the channel noise on the privacy loss.

- **Optimized power allocation and scheduling**: Building the analytical results, we formulate the optimization of the power allocation and scheduling policy as the minimization of the 2-Wasserstein distance under DP and power constraints. The resulting optimization is shown to be a convex program, and a closed-form solution is provided under simplifying assumptions. The analysis identifies distinct operating regimes in which the performance is power-limited, DP-limited, or LMC-limited. The three regimes are determined by the relative values of transmitted power, privacy level, and learning rate. The analytical results demonstrate that in the LMC-limited regime channel noise can be fully repurposed for both MC sampling and privacy preservation, obtaining the same performance as in an ideal communication setting that is not subject to DP constraints. For the general case, we formulate a min-max problem that can be converted into a convex problem.

- **Experiments**: We provide extensive numerical results to demonstrate the joint role of channel noise for MC sampling and privacy.

In closing this section, we would like to emphasize the relationship of this work with our previous papers [6] and [36]. As mentioned, in [6], we considered frequentist FL on a wireless channel, and analyzed the problem of optimal power allocation for an AirComp-based strategy that leverages channel noise as a privacy mechanism. The problem formulation has a minor overlap with the setting studied here, which focuses on Bayesian learning. In fact, Bayesian learning requires the analysis of performance metrics based on distributions in the model parameter space [35], and it cannot rely on the standard tools for the convergence of gradient-based schemes used in [6]. In contrast, reference [36] introduces a one-shot Bayesian protocol for a wireless data center setting in which the server has access to the global data set. In this system, the global data set is divided up among the workers to benefit from computational parallelism, but the server uses its access to the global data set during training. Specifically, paper [36]

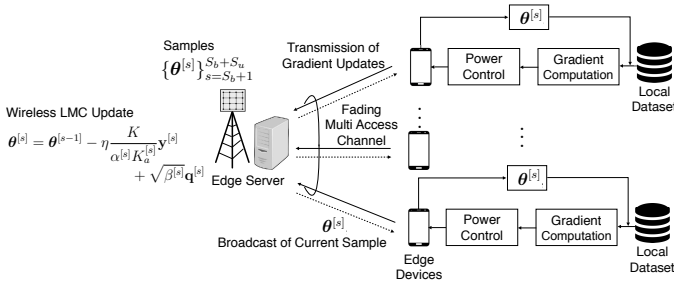


Figure 1. Differentially private federated Bayesian learning system based on Langevin Monte Carlo (LMC).

proposes a novel VI-based strategy that builds on consensus MC [37] by accounting for the presence of fading and channel noise. The contribution of [36] is distinct from the current manuscript for a number of reasons. First, in the current work, we study for the first time iterative, rather than one-shot, Bayesian learning protocols. Second, we concentrate on a federated setting in which the server does not have access to the global data set. Third, no privacy constraints are assumed in [36]. And, fourth, unlike [36], this work provides an analysis of the optimal power allocation strategy and draws theoretical conclusions on the capacity of the channel noise to serve the double role of seed randomness for MC sampling and privacy protection.

**Organization:** The remainder of the paper is organized as follows. Section II introduces the system model. Section III proposes the design of WFLMC. Section IV presents convergence and privacy analysis of WFLMC, while the optimal power allocation and scheduling are provided in Section V, followed by numerical results in Section VI and conclusions in Section VII.

## II. SYSTEM MODEL

As shown in Fig. 1, we consider a wireless federated edge learning system comprising a single-antenna edge server and  $K$  edge devices connected through it via a shared wireless channel. Each device  $k$ , equipped with a single antenna, has its own local dataset  $\mathcal{D}_k$  encompassing  $N_k$  data samples  $\mathcal{D}_k = \{\mathbf{d}_{k,n}\}_{n=1}^{N_k}$ . For supervised learning applications, each data sample  $\mathbf{d}_n = (\mathbf{u}_n, v_n)$  is in turn partitioned into a covariate vector  $\mathbf{u}_n$  and a label  $v_n$ ; while, for unsupervised learning applications such as generative modeling, it consists of a single vector  $\mathbf{d}_n$ . The global data set is denoted as  $\mathcal{D} = \{\mathcal{D}_k\}_{k=1}^K$ . The goal of the system is to carry out Bayesian learning via gradient-based Monte Carlo (MC) sampling. Accordingly, through communication with devices, the server wishes to obtain a number of random samples of the model parameter vector  $\boldsymbol{\theta} \in \mathbb{R}^m$  that are approximately distributed according to the global posterior distribution  $p(\boldsymbol{\theta}|\mathcal{D})$ . Unlike [36], which studied one-shot protocols, the gradient-based MC methods studied in this paper are iterative, in a manner similar to standard federated learning protocols such as FedAvg.

### A. Langevin Monte Carlo

The machine learning model adopted by the system is defined by a likelihood function  $p(\mathbf{d}|\boldsymbol{\theta})$  as well as by a prior

distribution  $p(\boldsymbol{\theta})$ . Accordingly, the likelihood of the data at device  $k$  is

$$p(\mathcal{D}_k|\boldsymbol{\theta}) = \prod_{n=1}^{N_k} p(\mathbf{d}_{n,k}|\boldsymbol{\theta}), \quad (1)$$

where the likelihoods  $p(\mathbf{d}_{n,k}|\boldsymbol{\theta})$  may be different across devices. The goal of Bayesian learning is to compute the global posterior

$$\text{(Global Posterior)} \quad p(\boldsymbol{\theta}|\mathcal{D}) \propto p(\boldsymbol{\theta}) \prod_{k=1}^K p(\mathcal{D}_k|\boldsymbol{\theta}). \quad (2)$$

The global posterior can be expressed in terms of the local posteriors at each device  $k$ , i.e.,

$$\text{(Local Posterior)} \quad p(\boldsymbol{\theta}|\mathcal{D}_k) \propto p(\boldsymbol{\theta})^{1/K} p(\mathcal{D}_k|\boldsymbol{\theta}) \quad (3)$$

since we have the equality

$$p(\boldsymbol{\theta}|\mathcal{D}) \propto \prod_{k=1}^K p(\boldsymbol{\theta}|\mathcal{D}_k). \quad (4)$$

In contrast, frequentist learning is concerned with the optimization of the **global cost function**

$$f(\boldsymbol{\theta}) = -\log p(\boldsymbol{\theta}|\mathcal{D}) = -\sum_{k=1}^K \log p(\boldsymbol{\theta}|\mathcal{D}_k), \quad (5)$$

which is the log-loss evaluated from the unnormalized posterior (2). Minimizing the global cost function (5) yields the maximum a posterior (MAP) solution of frequentist learning. The global cost function (5) can be expressed in terms of the **local cost functions**

$$f_k(\boldsymbol{\theta}) = -\log p(\mathcal{D}_k|\boldsymbol{\theta}) - \frac{1}{K} \log p(\boldsymbol{\theta}), \quad (6)$$

since we have the equality  $f(\boldsymbol{\theta}) = \sum_{k=1}^K f_k(\boldsymbol{\theta})$ .

Directly computing the global posterior distribution (2) is generally of prohibitive complexity. To address this problem, Monte Carlo (MC) methods represent the global posterior distribution (2) in terms of samples approximately distributed from it. Specifically, Markov Chain MC (MCMC) techniques produce a sequence of samples  $\boldsymbol{\theta}^{[s]}$  with  $s = 1, 2, \dots$  with the key property that as  $s$  grows large, the marginal distribution of sample  $\boldsymbol{\theta}^{[s]}$  tends to the desired posterior distribution. As discussed in Sec. I, in this paper, we specifically adopt Langevin MC (LMC), a fundamental MCMC technique that has been widely studied as a means to strike a practical balance between complexity and efficiency (see, e.g., [12], [13]). LMC is a gradient-based MCMC sampling scheme building on the global cost function (5). As we elaborate on next, the LMC update is derived as an approximation of a continuous-time differential process that has the desired property of asymptotically producing samples drawn from the global posterior (2).

To start, we first introduce the continuous-time Langevin diffusion process (LDP)  $\{\bar{\boldsymbol{\theta}}^{(t)} : t \in \mathbb{R}^+\}$  follows the stochastic differential equation

$$\text{(LDP)} \quad d\bar{\boldsymbol{\theta}}^{(t)} = -\nabla f(\bar{\boldsymbol{\theta}}^{(t)})dt + \sqrt{2}d\mathbf{B}^{(t)}, \quad (7)$$

where  $\mathbf{B}^{(t)}$  represents Brownian motion. The Langevin diffusion process (7) has the invariant stable distribution  $p^*(\boldsymbol{\theta}) = p(\boldsymbol{\theta}|\mathcal{D}) \propto \exp[-f(\boldsymbol{\theta})]$ , which corresponds to the desired global posterior (2) [38].

The integral of the stochastic differential equation (7) in the range of  $t \in [s\eta, (s+1)\eta]$  for integers  $s = 1, 2, \dots$  and duration  $\eta$  yields

$$\bar{\boldsymbol{\theta}}^{((s+1)\eta)} - \bar{\boldsymbol{\theta}}^{(s\eta)} = - \int_{s\eta}^{(s+1)\eta} \nabla f(\bar{\boldsymbol{\theta}}^{(t)}) dt + \sqrt{2\eta} \boldsymbol{\xi}^{[s+1]}, \quad (8)$$

where  $\{\boldsymbol{\xi}^{[s+1]}\}$  is a sequence of identical and independent (i.i.d.) random vectors following the Gaussian distribution  $\mathcal{N}(0, \mathbf{I}_m)$ . This is because we have the equality  $\sqrt{2} \int_{s\eta}^{s\eta+\eta} d\mathbf{B}^{(t)} = \sqrt{2}(\mathbf{B}^{(s\eta+\eta)} - \mathbf{B}^{(s\eta)}) \sim \mathcal{N}(0, 2\eta)$ .

In the limit of a sufficiently small  $\eta$  such that the following approximation holds

$$\int_{s\eta}^{(s+1)\eta} \nabla f(\bar{\boldsymbol{\theta}}^{(t)}) dt \approx \eta \nabla f(\boldsymbol{\theta}^{[s]}), \quad (9)$$

the discretization (8) results in Langevin MC (LMC), a gradient-based Markov chain MC (MCMC) sampling scheme that proceeds according to the iterative update rule

$$\text{(LMC)} \quad \boldsymbol{\theta}^{[s+1]} = \boldsymbol{\theta}^{[s]} - \eta \nabla f(\boldsymbol{\theta}^{[s]}) + \sqrt{2\eta} \boldsymbol{\xi}^{[s+1]}, \quad (10)$$

where we have  $\boldsymbol{\theta}^{[s]} = \bar{\boldsymbol{\theta}}^{(s\eta)}$  and  $\eta$  as step size. The error due to approximation (9) was studied in [35], [39], and it will be further discussed in Sec. IV-A.

In order to obtain samples  $\boldsymbol{\theta}^{[s]}$  approximately drawn from the global posterior distribution, LMC discards the samples produced in the first  $S_b$  iterations (10), also known as *burn-in period*. The remaining  $S_u$  samples  $\boldsymbol{\theta}^{[s]}$  with  $s = S_b + 1, S_b + 2, \dots, S_b + S_u$ , are retained and used for downstream applications such as ensemble prediction.

We consider LMC in this work due to its simplicity and scalability, as LMC updates amount to adding Gaussian noise to gradient descent updates. Other gradient-based MCMC methods, such as kinetic MC [14] and Hamilton MC [15], which modify higher-order updates may also be studied in a manner similar to this paper and are left for future research.

## B. Learning Protocol

The goal of the system under study is to implement LMC (10) in the described federated setting with  $K$  devices, while satisfying formal differential privacy (DP) guarantees to be detailed in Sec. II-D. The protocols is organized in iterations  $s = 1, 2, \dots, S_b + S_u$  with  $S_b$  denoting the burn-in period, across which the server maintains sample iterates  $\boldsymbol{\theta}^{[s]}$ . The selection of burn-in period should avoid the regime in which the samples are too dependent on the Markov chain's initialization. A formal test for this purpose is the Gelman-Rubin diagnostic, which is based by using multiple Markov chains. The diagnostic compares the estimated between-chains and within-chain variances for each model parameter, and chooses the burin-in period after which difference between theses variances are sufficiently small.

At each  $s$ -th communication round, the edge server broadcasts the current sample  $\boldsymbol{\theta}^{[s]}$  to all edge devices via the downlink channel. We assume that downlink communication is ideal, so that each device receives the sample  $\boldsymbol{\theta}^{[s]}$  without distortion. This assumption is practically well justified when the edge server communicates through a base station with less stringent power constraint than the devices and the use of the whole downlink bandwidth for broadcasting. It is commonly made in many related papers, such as [5]–[7].

By using the received vector  $\boldsymbol{\theta}^{[s]}$  and the local dataset  $\mathcal{D}_k$ , each device computes the gradient of the local cost function (6) as

$$\text{(Local gradient)} \\ \nabla f_k(\boldsymbol{\theta}^{[s]}) = - \sum_{n=1}^{N_k} \nabla \log p(\mathbf{d}_n | \boldsymbol{\theta}^{[s]}) - \frac{1}{K} \nabla \log p(\boldsymbol{\theta}^{[s]}), \quad (11)$$

which is transmitted over the wireless shared channel to the edge server. The goal is to enable the edge server to approximate the update term in (10), namely

$$-\eta \nabla f(\boldsymbol{\theta}^{[s]}) + \sqrt{2\eta} \boldsymbol{\xi}^{[s+1]} = -\eta \sum_{k=1}^K \nabla f_k(\boldsymbol{\theta}^{[s]}) + \sqrt{2\eta} \boldsymbol{\xi}^{[s+1]}. \quad (12)$$

As we will see, channel noise can be repurposed to contribute to the additive random term  $\boldsymbol{\xi}^{[s+1]}$  in the LMC update (10). The steps in (11) and (10) are iterated across multiple communication rounds until a convergence condition is met. As a result, the server obtains a sequence of global model parameter vectors  $\boldsymbol{\theta}^{[s]}$ , with  $s = 1, 2, \dots, S_b + S_u$ .

## C. Communication Model

The devices communicate via the uplink to the edge server on the shared wireless channel. The proposed approach leverages analog transmission in order to: (i) benefit from over-the-air computing as in many prior works [5], [6], [20]; and (ii) to repurpose channel noise for MC sampling and as a privacy mechanism. We assume a block flat-fading channel, where the channel coefficients remain constant within a communication block, and they vary in a potentially correlated way over successive blocks. Each block contains  $m$  channel uses, allowing the uncoded transmission of a gradient vector via non-orthogonal multiple access (NOMA) as in [6], [20].

We assume symbol-level synchronization among the subset of devices that are scheduled in each block, enabling over-the-air computing. This can be achieved by using standard protocols such as the timing advance procedure in LTE and 5G NR [40]. In the  $s$ -th communication round, the corresponding received signal is

$$\mathbf{y}^{[s]} = \sum_{k=1}^K h_k^{[s]} \mathbf{x}_k^{[s]} + \mathbf{z}^{[s]}, \quad (13)$$

where  $h_k^{[s]}$  is the channel gain for device  $k$  in round  $s$ ,  $\mathbf{x}_k^{[s]} \in \mathbb{R}^m$  is an uncoded function of the local gradient

$\nabla f_k(\boldsymbol{\theta}^{[s]})$ , and  $\mathbf{z}^{[s]}$  is channel noise i.i.d. according to distribution  $\mathcal{N}(0, N_0 \mathbf{I}_m)$ . The transmit power constraint of a device is given as

$$\|\mathbf{x}_k^{[s]}\|^2 \leq P, \quad (14)$$

accounting for per-block power constraints.

#### D. Performance Metrics

In this paper, we aim at designing a wireless federated learning protocol to implement Bayesian learning via Langevin MC under DP constraints. In this subsection, we formalize the performance criteria of interest and elaborate on the role of channel noise in achieving them.

1) *Approximation Error*: Denoting as  $p^{[s]}(\boldsymbol{\theta})$  the distribution of the sample  $\boldsymbol{\theta}^{[s]}$  at the  $s$ -th iteration, the quality of the sample is measured by the 2-Wasserstein distance between  $p^{[s]}(\boldsymbol{\theta})$  and the target global posterior  $p(\boldsymbol{\theta}|\mathcal{D})$  in (2). This is defined as [35], [39]

$$W_2(p^{[s]}(\boldsymbol{\theta}), p(\boldsymbol{\theta}|\mathcal{D})) = \left( \inf_{p(\boldsymbol{\theta}, \boldsymbol{\theta}')} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|_2^2 p(\boldsymbol{\theta}, \boldsymbol{\theta}') d\boldsymbol{\theta} d\boldsymbol{\theta}' \right)^{1/2}, \quad (15)$$

where joint distribution  $p(\boldsymbol{\theta}, \boldsymbol{\theta}')$  is constrained to have marginals  $p^{[s]}(\boldsymbol{\theta})$  and  $p(\boldsymbol{\theta}'|\mathcal{D})$ . The Wasserstein distance is a standard measure of discrepancy between two distributions, and it is routinely used for the analysis of MC algorithms (see, e.g., [35], [39]). It has some useful properties with respect to other measures such as the Kullback-Leibler divergence and total variation distance. For instance, it is well defined and informative even when the two distributions have disjoint supports [41].

2) *Differential Privacy*: We consider a ‘‘honest-but-curious’’ edge server that may attempt to infer information about local data sets from the received signals  $\mathbf{y}$ . We impose the standard  $(\epsilon, \delta)$ -DP metric with some  $\epsilon > 0$  and  $\delta \in [0, 1)$ , for each device  $k$ . This amounts to the inequalities

$$P(\{\mathbf{y}^{[s]}\}_{s=1}^S | \{\cup_{i \neq k} \mathcal{D}_i\} \cup \mathcal{D}'_k) \leq \exp(\epsilon) P(\{\mathbf{y}^{[s]}\}_{s=1}^S | \{\cup_{i \neq k} \mathcal{D}_i\} \cup \mathcal{D}''_k) + \delta, \quad (16)$$

for all  $k = 1, \dots, K$ , where  $S$  is the number of communication rounds, and  $P(\mathbf{y}^{[s]}|\mathcal{D})$  represents the distribution of the received signal (13) conditioned on the global data set  $\mathcal{D}$ . Condition (16) must hold for any two possible neighboring data sets  $\mathcal{D}'_k$  and  $\mathcal{D}''_k$  differing only by one sample, i.e.,  $\|\mathcal{D}'_k - \mathcal{D}''_k\|_1 = 1$ , and for all data sets  $\{\mathcal{D}_i\}_{i \neq k}$  of other devices.

3) *On the Role of Channel Noise*: Without the additive random term in  $\boldsymbol{\xi}^{[s+1]}$  in (10), LMC coincides with standard gradient descent (GD) for frequentist learning, which was studied in [6] in a federated setting under the DP constraints (16). For convex models, the additive channel noise on the uplink channel (13) is harmful to the convergence rate of GD with no privacy constraints [6], although it can improve the generalization performance [4]. In [6], [9], [31], it was shown that channel noise can be repurposed to ensure the DP

constraints (16) for values of  $(\epsilon, \delta)$  that depend on the SNR level.

In this paper, we observe that, in Bayesian learning via MC, by (10), noise can potentially contribute not only to privacy but also to MC sampling, without necessarily compromising the learning performance. This idea was first introduced in [36] for one-shot MC sampling methods, and is studied here for the first time for iterative schemes. Specifically, we investigate the joint role of channel noise as a contributor to MC sampling and as a privacy-inducing mechanism. This interplay between MC sampling and DP was analyzed under ideal communication for centralized learning in [33], and the impact of channel noise in distributed settings is studied here for the first time.

#### E. Assumptions on the Log-Likelihood

Finally, we list several standard assumptions we make on the global cost function  $f(\boldsymbol{\theta})$  in (5) and on its gradient.

**Assumption 1** (Smoothness). The global cost function  $f(\boldsymbol{\theta})$  is smooth with constant  $L > 0$ , that is, it is continuously differentiable and the gradient  $\nabla f(\boldsymbol{\theta})$  is Lipschitz continuous with constant  $L$ , i.e.,

$$\|\nabla f(\boldsymbol{\theta}) - \nabla f(\boldsymbol{\theta}')\| \leq L \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|, \quad \text{for all } \boldsymbol{\theta}, \boldsymbol{\theta}' \in \mathbb{R}^m. \quad (17)$$

**Assumption 2** (Strong Convexity). The global cost function  $f(\boldsymbol{\theta})$  is strongly convex, i.e., the following inequality holds for some constant  $\mu > 0$

$$[\nabla f(\boldsymbol{\theta}) - \nabla f(\boldsymbol{\theta}')]^T (\boldsymbol{\theta} - \boldsymbol{\theta}') \geq \mu \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|^2, \quad \text{for all } \boldsymbol{\theta}, \boldsymbol{\theta}' \in \mathbb{R}^m. \quad (18)$$

Assumptions 1 and 2 imply the following inequality [42, Lemma 3.11]

$$[\nabla f(\boldsymbol{\theta}) - \nabla f(\boldsymbol{\theta}')]^T (\boldsymbol{\theta} - \boldsymbol{\theta}') \geq \frac{\mu L}{\mu + L} \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|^2 + \frac{1}{\mu + L} \|\nabla f(\boldsymbol{\theta}) - \nabla f(\boldsymbol{\theta}')\|^2, \quad \text{for all } \boldsymbol{\theta}, \boldsymbol{\theta}' \in \mathbb{R}^m. \quad (19)$$

**Assumption 3** (Bounded Local Gradient). The local gradient is bounded as

$$\|\nabla f_k(\boldsymbol{\theta}^{[s]})\| \leq \ell, \quad \text{for all } k, s, \quad (20)$$

and some constant  $\ell > 0$ .

This last assumption is essential to ensure DP requirements [33], [43]. One can choose the parameter of  $\ell$  as the maximum value of  $\|\nabla f_k(\boldsymbol{\theta}^{[s]})\|$ , or, in practice, clipping the gradient as  $\overline{\nabla f_k(\boldsymbol{\theta}^{[s]})} = \min\{1, \ell / \|\nabla f_k(\boldsymbol{\theta}^{[s]})\|\} \nabla f_k(\boldsymbol{\theta}^{[s]})$ . We will not account for the impact of clipping in the analysis.

### III. WIRELESS FEDERATED LANGEVIN MONTE CARLO (WFLMC)

In this section, we introduce wireless federated LMC (WFLMC). Specifically, we first present signal design and scheduling protocol, and then detail the rationale behind the proposed approach. The following sections will focus on the analysis of WFLMC.

### A. Signal Design and Scheduling Protocol

In each communication block, all devices transmit their local gradients simultaneously by using uncoded transmission of the form

$$\text{(Transmitted Signal)} \quad \mathbf{x}_k^{[s]} = \alpha_k^{[s]} \nabla f_k(\boldsymbol{\theta}^{[s-1]}), \quad (21)$$

for some power control parameter  $\alpha_k^{[s]}$ . Specifically, we implement truncated channel inversion as in [5], [7], [32], whereby the power control parameter is selected as

$$\alpha_k^{[s]} = \begin{cases} \alpha^{[s]}/h_k^{[s]}, & \text{if } |h_k^{[s]}| \geq g^{[s]}, \\ 0, & \text{if } |h_k^{[s]}| < g^{[s]}, \end{cases} \quad (22)$$

where gain parameter  $\alpha^{[s]} > 0$  and threshold  $g^{[s]} > 0$  are parameters to be optimized. Accordingly, a device  $k$  transmits only if its channel is large enough, i.e.,  $|h_k^{[s]}| \geq g^{[s]}$ . We note that this channel-aware scheduling aims to avoid deep fading channels. Other alternative scheduling policies may achieve better performance by considering also the importance of local data set, such as [44], [45]. We denote as  $K_a^{[s]}$  the number of transmitting devices at round  $s$ , and  $\mathcal{K}_a^{[s]}$  the corresponding set of transmitting devices.

To estimate the LMC update term (12) in (10), the received signal (13) is scaled and a Gaussian random vector is added. Specifically, *wireless federated LMC (WFLMC) update* is proposed as

$$\text{(WFLMC)} \quad \boldsymbol{\theta}^{[s]} = \boldsymbol{\theta}^{[s-1]} - \eta \frac{K}{\alpha^{[s]} K_a^{[s]}} \mathbf{y}^{[s]} + \sqrt{\beta^{[s]}} \mathbf{q}^{[s]}, \quad (23)$$

where the added noise  $\mathbf{q}^{[s]} \sim \mathcal{N}(0, \mathbf{I}_m)$  is independent of all other variables. The variance  $\beta^{[s]}$  of the noise term  $\mathbf{q}^{[s]}$  is chosen as a function of the learning rate  $\eta$ , channel noise  $N_0$ , number of active users  $K_a^{[s]}$ , and gain  $\alpha^{[s]}$  as

$$\beta^{[s]} = \max \left\{ 0, 2\eta - \frac{\eta^2 N_0 K^2}{(\alpha^{[s]} K_a^{[s]})^2} \right\}. \quad (24)$$

### B. Understanding WFLMC

To see the rationale behind the design (23)–(24), let us plug (13) and (21)–(22) into (23) to rewrite the WLMC update as

$$\begin{aligned} \boldsymbol{\theta}^{[s]} &= \boldsymbol{\theta}^{[s-1]} - \eta \frac{K}{K_a^{[s]}} \sum_{k \in \mathcal{K}_a^{[s]}} \nabla f_k(\boldsymbol{\theta}^{[s-1]}) \\ &\quad - \underbrace{\eta \frac{K}{\alpha^{[s]} K_a^{[s]}} \mathbf{z}^{[s]} + \sqrt{\beta^{[s]}} \mathbf{q}^{[s]}}_{\text{effective noise } \tilde{\mathbf{z}}^{[s]}} \quad (25) \\ &= \boldsymbol{\theta}^{[s-1]} - \underbrace{\eta \left[ \frac{K}{K_a^{[s]}} \sum_{k \in \mathcal{K}_a^{[s]}} \nabla f_k(\boldsymbol{\theta}^{[s-1]}) + \sqrt{\tilde{\beta}^{[s]}} \Delta^{[s]} \right]}_{\widehat{\nabla} f(\boldsymbol{\theta}^{[s-1]})} \\ &\quad + \sqrt{2\eta} \widehat{\boldsymbol{\xi}}^{[s]}, \quad (26) \end{aligned}$$

where  $\widehat{\boldsymbol{\xi}}^{[s]} \sim \mathcal{N}(0, \mathbf{I}_m)$  and  $\Delta^{[s]} \sim \mathcal{N}(0, \mathbf{I}_m)$  are the i.i.d. sequences for  $s = 1, 2, \dots$ , and we define

$$\tilde{\beta}^{[s]} = \max \left\{ 0, \frac{N_0 K^2}{(\alpha^{[s]} K_a^{[s]})^2} - \frac{2}{\eta} \right\}. \quad (27)$$

By (25), the power scaling (22) ensures that the gradients of the active devices sum at the receiver, and the scaling by  $K/(\alpha^{[s]} K_a^{[s]})$  in (23) of the received signal compensates for the resulting multiplier of the sum-gradient. The term  $K/K_a^{[s]} \sum_{k \in \mathcal{K}_a^{[s]}} \nabla f_k(\boldsymbol{\theta}^{[s-1]})$  is an empirical estimate of the gradient  $\sum_{k=1}^K \nabla f_k(\boldsymbol{\theta}^{[s-1]})$  in (10) which is exact if  $K_a^{[s]} = K$ .

Based on the discussion so far, in order for (25) to be an estimate of the LMC update (10), we should ideally ensure that the variance of the effective noise term  $\tilde{\mathbf{z}}^{[s]}$  be equal to  $2\eta$ . However, the power of this term, namely,

$$\tilde{\sigma}_z^2 = \eta^2 K^2 / (\alpha^{[s]} K_a^{[s]})^2 + \beta^{[s]} \quad (28)$$

can only be partially controlled through power gain  $\alpha^{[s]}$  and added noise variance  $\beta^{[s]}$ . In particular, we can always choose the added noise variance  $\beta^{[s]}$  such that the variance  $\tilde{\sigma}_z^2$  is no smaller than  $2\eta$ . This condition is ensured by (24). In particular with (24), the effective noise in (25) can be decomposed into two parts as indicated in (26):

1) *LMC noise*: The term  $\sqrt{2\eta} \widehat{\boldsymbol{\xi}}^{[s]}$  with  $\widehat{\boldsymbol{\xi}}^{[s]} \sim \mathcal{N}(0, \mathbf{I}_m)$  serves the role of LMC noise with variance  $2\eta$ ;

2) *Gradient estimation noise*: The remaining noise, denoted as  $-\eta \sqrt{\tilde{\beta}^{[s]}} \Delta^{[s]}$  with  $\Delta^{[s]} \sim \mathcal{N}(0, \mathbf{I}_m)$ , acts as a perturbation on the gradient estimate with variance  $\tilde{\beta}^{[s]}$  in (27). Note that the variance  $\tilde{\beta}^{[s]}$  of estimate noise is non-zero if the channel noise power  $N_0$  is large.

## IV. CONVERGENCE AND PRIVACY ANALYSIS OF WFLMC

In this section, we focus on the performance analysis of WFLMC in terms of (i) convergence through the 2-Wasserstein distance as defined in (15); and (ii) privacy under the DP criterion (16). In this section, we assume that the sequence of power gain and scheduling threshold parameters  $\{\alpha^{[s]}, g^{[s]}\}$  is fixed, and the results are given for an arbitrary sequence  $\{h_k^{[s]}\}$  of channels.

### A. Convergence Analysis

We now study the distribution  $p^{[s]}(\boldsymbol{\theta})$  of the sample  $\boldsymbol{\theta}^{[s]}$  produced by WFLMC via the updates (23). We recall that the goal of LMC is to produce samples distributed according to the global posterior  $p(\mathcal{D}|\boldsymbol{\theta})$ . As discussed in Sec. II-D, we measure the approximation error via 2-Wasserstein distance (15).

There are three main contributions to the discrepancy between the distribution of  $p^{[s]}(\boldsymbol{\theta})$  and the target posterior  $p(\boldsymbol{\theta}|\mathcal{D})$ :

- 1) the initial discrepancy, which is measured by the 2-Wasserstein distance  $W_2(p^{[0]}(\boldsymbol{\theta}), p(\boldsymbol{\theta}|\mathcal{D}))$  between the initial distribution  $p^{[0]}(\boldsymbol{\theta})$  and the posterior  $p(\boldsymbol{\theta}|\mathcal{D})$ ;
- 2) the gradient error in (26), namely

$$\text{(Gradient error)} \quad N_g^{[s]} = \nabla f(\boldsymbol{\theta}^{[s-1]}) - \widehat{\nabla} f(\boldsymbol{\theta}^{[s-1]}), \quad (29)$$

which is caused by the excess noise  $\sqrt{\tilde{\beta}^{[s]}\Delta^{[s]}}$  and by the fact that only a subset of  $K_a^{[s]}$  devices is active;

- 3) and the discretization error due to the approximation (9), which is given as

(Discretization error)

$$N_d^{[s]} = \int_{(s-1)\eta}^{s\eta} \nabla f(\bar{\theta}^{(t)}) - \nabla f(\bar{\theta}^{((s-1)\eta)}) dt. \quad (30)$$

We now bound the last two terms separately, and then use these results to bound square of the desired 2-Wasserstein distance  $W_2(p^{[s]}(\theta), p(\theta|\mathcal{D}))^2$  in (15).

**Lemma 1** (Upper Bound on the Gradient Error). Under Assumption 3, for any  $s$ -th communication round, the average power of the gradient error is bounded as

$$\mathbb{E} \left[ \|N_g^{[s]}\|^2 \right] \leq 4\ell^2 (K - K_a^{[s]})^2 + \tilde{\beta}^{[s]}, \quad (31)$$

where the expectation is taken with respect to the excess noise  $\sqrt{\tilde{\beta}^{[s]}\Delta^{[s]}}$  and the variance  $\tilde{\beta}^{[s]}$  is defined in (27).

*Proof:* See [46, Appendix A].

The error bound (31) is comprised of two parts. The first term is the estimation error due to scheduling, which is zero if all the devices transmit in communication rounds, i.e., if  $K_a^{[s]} = K$ . The second term is the variance  $\tilde{\beta}^{[s]}$  of the excess noise  $\sqrt{\tilde{\beta}^{[s]}\Delta^{[s]}}$  in (27).

The discretization error is constant for all communication round  $s = 1, 2, \dots$ , and can be bounded by following [35, Lemma 3] as detailed in the next lemma.

**Lemma 2** (Upper Bound on the Discretization Error). For any communication round  $s$ , the discretization error is invariant, and it is upper bounded as

$$\mathbb{E} \left[ \|N_d^{[s]}\|^2 \right] \leq \frac{\eta^4 L^3 m}{3} + \eta^3 L^2 m, \quad (32)$$

where the average is computed with respect to the joint distribution of the LDP (7) and the LMC (10).

*Proof:* The proof is detailed in Appendix A.

Lemma 2 shows that the discretization error grows with model dimension  $m$ , with the learning rate  $\eta$ , and with the smoothness constant  $L$  of the global cost function  $f(\theta)$ . Note that a larger  $L$  implies a less smooth function  $f(\theta)$ .

We now leverage Lemma 1 and Lemma 2 to bound the square of 2-Wasserstein distance.

**Proposition 1** (Bound for 2-Wasserstein Distance). For a learning rate  $0 < \eta \leq 2/L$ , define

$$\gamma = \begin{cases} 1 - \eta\mu, & 0 < \eta \leq 2/(\mu + L), \\ \eta L - 1, & 2/(\mu + L) \leq \eta \leq 2/L. \end{cases} \quad (33)$$

Under Assumptions 1, 2 and 3, after any number  $s'$  of iterations, the 2-Wasserstein distance between the sample

distribution produced by WFLMC and the global posterior is upper bounded as

$$\begin{aligned} W_2(p^{[s']}(\theta), p(\theta|\mathcal{D}))^2 &\leq \left(\frac{1+\gamma}{2}\right)^{2s'} W_2(p^{[0]}(\theta), p(\theta|\mathcal{D}))^2 \\ &+ \sum_{s=1}^{s'} \left(\frac{1+\gamma}{2}\right)^{2(s'-s)} \frac{2(1+\gamma)}{1-\gamma} \left[ \frac{\eta^4 L^3 m}{3} + \eta^3 L^2 m \right. \\ &\quad \left. + 4\eta^2 \ell^2 (K - K_a^{[s]})^2 + \eta^2 \tilde{\beta}^{[s]} \right] \end{aligned} \quad (34)$$

$$\triangleq \tilde{W}_2(p^{[s']}(\theta), p(\theta|\mathcal{D}))^2. \quad (35)$$

*Proof:* The proof follows from the upper bound (see, e.g., [35])

$$W_2(p^{[s]}(\theta), p(\theta|\mathcal{D}))^2 \leq \mathbb{E} \left[ \left\| \theta^{[s]} - \bar{\theta}^{(s\eta)} \right\|^2 \right], \quad (36)$$

where  $\bar{\theta}^{(s\eta)}$  is obtained from the LDP (7), while  $\theta^{[s]}$  is the sample produced by the LMC (10) with the LMC noise  $\sqrt{2\eta}\hat{\xi}^{[s]} = \sqrt{2} \int_{s\eta}^{s\eta+\eta} d\mathbf{B}^{(t)}$ . Accordingly, the expectation is taken with respect to the Brownian motion  $\mathbf{B}^{(t)}$  in (7) and over the initial  $\theta^{(0)} \sim p^{[0]}(\theta)$ . Since, assuming the stationary of the LDP (7), the marginal of the LDP output  $\bar{\theta}^{(s)}$  is the target distribution  $p(\theta|\mathcal{D})$  for all  $t > 0$ , the bound (36) follows from the definition (15) by upper bounding the infimum with the described choice of the joint distribution of  $\theta^{[s]}$  and  $\bar{\theta}^{(s\eta)}$ .

Using (26) and (8), the right hand side of (36) can be computed as

$$\begin{aligned} &\mathbb{E} \left[ \left\| \theta^{[s]} - \bar{\theta}^{(s\eta)} \right\|^2 \right] \\ &= \mathbb{E} \left[ \left\| \theta^{[s-1]} - \bar{\theta}^{((s-1)\eta)} - \eta \left[ \nabla f(\theta^{[s-1]}) - \nabla f(\bar{\theta}^{((s-1)\eta)}) \right] \right. \right. \\ &\quad \left. \left. + \underbrace{\eta \left[ \nabla f(\theta^{[s-1]}) - \widehat{\nabla} f(\theta^{[s-1]}) \right]}_{N_g^{[s]}} \right\|^2 \right] \\ &\quad + \underbrace{\left\| \int_{(s-1)\eta}^{s\eta} \nabla f(\bar{\theta}^{(t)}) - \nabla f(\bar{\theta}^{((s-1)\eta)}) dt \right\|^2}_{N_d^{[s]}}. \end{aligned} \quad (37)$$

The rest of the proof involves applying the geometric inequality  $2ab \leq \tau a^2 + \tau^{-1}b^2$  for any  $\tau > 0$ , and using Lemma 1 and Lemma 2 as detailed in Appendix B.  $\square$

Proposition 1 indicates that the 2-Wasserstein distance depends on the initial discrepancy  $W_2(p^{[0]}(\theta), p(\theta|\mathcal{D}))$ , whose contribution decreases exponentially with  $s'$ ; as well as the sum of contributions across the iteration index  $s = 1, 2, \dots, s'$ , with each  $s$ -th error term weighted down by a factor decreasing exponentially with  $s' - s$ , i.e., as one moves towards earlier iterations. This shows that the disturbances at later communication rounds are more harmful to the approximation accuracy. Furthermore, the contribution of each iteration  $s \geq 1$  is given by the sum of the gradient error bounded in Lemma 1, and the discretization error bounded in Lemma 2.

Another interesting aspect highlighted by the bound (34) concerns the optimal choice of the learning rate  $\eta$ . The upper bound (34) increases with  $\gamma$  in (33), and the minimum value of  $\gamma$  is attained when  $\eta = 2/(\mu + L)$ . However, a large



learning rate  $\eta$  causes the gradient error bound (31) and the discretization error bound (32) to increase by Lemma 1 and Lemma 2. Thus, the optimal learning rate is in the range of  $\eta \in (0, 2/(\mu + L)]$ .

### B. Differential Privacy Analysis

The WFLMC scheme implicitly implements a Gaussian DP mechanism [8], [33], since the channel noise  $\mathbf{z}^{[s]}$  in (25) is added to the disclosed function  $\sum_{k \in \mathcal{K}_a^{[s]}} h_k^{[s]} \alpha_k^{[s]} \nabla f_k(\boldsymbol{\theta}^{[s-1]})$ . Note that the noise  $\mathbf{q}^{[s]}$  in (25) is added by the edge server, and hence it does not contribute to privacy. Furthermore, while only part of the channel noise  $\mathbf{z}^{[s]}$  is useful for LMC (see (26)), the entire variance  $N_0$  contributes to DP.

For the Gaussian mechanism, the privacy level  $(\epsilon, \delta)$  depends on the sensitivity of the disclosed information and on the variance of the added noise [6], [8]. In a manner consistent to the definition (16) of DP, the sensitivity quantifies the maximum change of the disclosed function by replacing a single data point. The sensitivity for device  $k$  is accordingly defined as

(Sensitivity)

$$\chi_k^{[s]} = \max_{\mathcal{D}'_k, \mathcal{D}''_k} \left\| h_k^{[s]} \alpha_k^{[s]} \left[ \nabla f'_k(\boldsymbol{\theta}^{[s-1]}) - \nabla f''_k(\boldsymbol{\theta}^{[s-1]}) \right] \right\|, \quad (38)$$

where  $\nabla f'_k(\boldsymbol{\theta}^{[s-1]})$  and  $\nabla f''_k(\boldsymbol{\theta}^{[s-1]})$  are computed by using data sets  $\mathcal{D}'_k$  and  $\mathcal{D}''_k$  respectively, and we have  $\|\mathcal{D}'_k - \mathcal{D}''_k\|_1 = 1$ . By the triangular inequality and Assumption 3, we plug in the definition of  $\alpha_k^{[s]}$  (22) and have the bound

$$\chi_k^{[s]} \leq \mathbf{1}[|h_k^{[s]}| \geq g^{[s]}] \cdot 2\alpha^{[s]}\ell. \quad (39)$$

Following [6, Lemma 1], one can interpret the ratio  $(\chi_k^{[s]})^2/N_0$  as the privacy loss in each communication rounds. This is formalized in the following proposition.

**Proposition 2** (DP Guarantees). For any given sequence of parameters  $\{\alpha^{[s]}, g^{[s]}\}$  and channels  $\{h_k^{[s]}\}$ , after  $s$  communication rounds, WFLMC guarantees  $(\epsilon, \delta)$ -DP if the following condition is satisfied

$$\begin{aligned} & \sum_{s=1}^{s'} \frac{\mathbf{1}[|h_k^{[s]}| \geq g^{[s]}] \cdot 2(\alpha^{[s]}\ell)^2}{N_0} \\ & \leq \left( \sqrt{\epsilon + [\mathcal{C}^{-1}(1/\delta)]^2} - \mathcal{C}^{-1}(1/\delta) \right)^2 \\ & \triangleq \mathcal{R}_{\text{dp}}(\epsilon, \delta), \text{ for all } k, \end{aligned} \quad (40)$$

where  $\mathcal{C}^{-1}(x)$  is the inverse function of  $\mathcal{C}(x) = \sqrt{\pi}xe^{x^2}$ , and  $\mathbf{1}[\cdot]$  is the indicator function.

*Proof:* The result follows from [6, Lemma 1], although reference [6] did not account for the threshold-based scheduling in (22). The extension is direct by redefining the effective channel gain as  $\alpha^{[s]} \cdot \mathbf{1}[|h_k^{[s]}| \geq g^{[s]}]$ .  $\square$

In accordance to the discussion above, the left-hand side of (40) quantifies the overall privacy loss across  $s'$  rounds. Importantly, as anticipated, by (40) the channel noise power  $N_0$  contributes in full to the DP performance. In contrast, by (26), only a portion of the channel noise contributes, in general, to the LMC update.

## V. OPTIMAL POWER ALLOCATION AND SCHEDULING

In this section, we leverage Proposition 1 and Proposition 2 to address the problem of minimizing the convergence error under the  $(\epsilon, \delta)$ -DP constraint (40) and the power constraints (14) over power gain parameters and thresholds  $\{\alpha^{[s]}, g^{[s]}\}_{s=1}^S$  in (22). We recall that WFLMC carries out  $S_b$  communication rounds for the burn-in period, which are followed by  $S_u$  additional rounds to obtain the samples for use in downstream applications. The learning objective is to maximize the quality of the last  $S_u$  samples under the mentioned privacy and power constraints, which apply for the total of  $S = S_b + S_u$  communication rounds.

Throughout this section, we assume that the sequence of channels  $\{\{h_k^{[s]}\}_{k=1}^K\}_{s=1}^S$  is known in advance in order to enable optimization. This assumption can be relaxed at the cost of additional communication overhead. Extensions to an online approach can be directly obtained by iterative one-step-ahead optimization based on predicted values for the future parameters  $\{\{h_k^{[s']}\}_{k=1}^K\}_{s'=s+1}^S$  as detailed in [6]. The optimization is conducted at the edge server. To control Air-Comp transmission, the edge server broadcasts the optimized  $\{\alpha^{[s]}, g^{[s]}\}$  at each communication round, which consumes a negligible amount of communication resources as compared with model broadcasting.

The joint power allocation and scheduling problem of interest is formulated as the min-max optimization

$$\min_{\{\alpha^{[s]}, g^{[s]}\}_{s=1}^S} \max_{s' \in [S_b+1, S_b+S_u]} \widetilde{W}_2(p^{[s']}, p(\boldsymbol{\theta}|\mathcal{D}))^2 \quad (41a)$$

$$\text{s.t. } \sum_{s=1}^S \frac{\mathbf{1}[|h_k^{[s]}| \geq g^{[s]}] \cdot 2(\alpha^{[s]}\ell)^2}{N_0} \leq \mathcal{R}_{\text{dp}}(\epsilon, \delta) \quad (41b)$$

$$\mathbf{1}[|h_k^{[s]}| \geq g^{[s]}] \cdot \frac{(\alpha^{[s]}\ell)^2}{|h_k^{[s]}|^2} \leq P, \quad \forall k, s = 1, \dots, S. \quad (41c)$$

The maximization in (41a) aims at ensuring that the worst-case 2-Wasserstein distance is minimized across all  $S_u$  samples after the burn-in period, under the DP constraint (41b) and the power constraint (41c). With its focus on the distribution of model parameters, problem (41) is notably distinct from the optimization problem for frequentist learning in [6], which only considers the quality of a single vector of model parameters in terms of training loss.

The problem (41) is non-convex since the objective function is non-differentiable in the thresholds  $\{g^{[s]}\}_{s=1}^S$ . In fact, by (22), the threshold  $g^{[s]}$  affects the objective function through the number  $K_a^{[s]}$  of active users. To make progress, we fix the scalar thresholds  $g^{[s]}$  and optimize over the power gain parameter  $\alpha^{[s]}$ .

### A. Zero Additive Noise is Optimal

We start by simplifying problem (41) through the following observation.

**Lemma 3** (Zero Additive Noise). Without compromising optimality, the variance of additive noise  $\sqrt{\beta^{[s]}}\mathbf{q}^{[s]}$  in (23)

can be set as  $\beta^{[s]} = 0$ , which is equivalent to imposing the following constraint on the power gain parameters

(LMC noise requirement)

$$\alpha^{[s]} \leq \frac{K}{K_a^{[s]}} \sqrt{\frac{\eta N_0}{2}}, \quad \text{for } s = 1, \dots, S. \quad (42)$$

*Proof:* Assume by contradiction that we had  $\beta^{[s]} > 0$  at an optimal solution. By (24), this would imply that the optimal  $\alpha^{[s]}$  satisfies the inequality  $\eta^2 N_0 K^2 / (\alpha^{[s]} K_a^{[s]})^2 - 2\eta < 0$ . But one can always choose the smaller value  $\alpha^{[s]} = \sqrt{\eta N_0 / 2} K / K_a^{[s]}$ , which achieves the same value of the objective function, while reducing the left-hand sides of the privacy and power constraints (41b)-(41c). This yields a contradiction, completing the proof.  $\square$

The inequality (42) distinguishes two distinct regimes of operation of the system. If the equality (42) is active, the gradient estimation noise power  $\tilde{\beta}^{[s]}$  in (27) is zero, and hence the channel noise contributes in full to the LMC updates. In contrast, when the inequality is strict, only a fraction of the channel noise is useful for LMC, and the rest contributes to the gradient estimation noise power  $\tilde{\beta}^{[s]}$ .

### B. Optimization of Power Gain Parameters: Single-Sample and Constant Channels

We now tackle the optimization (41) over power gain parameters  $\{\alpha^{[s]}\}_{s=1}^S$  for the special case  $S_u = 1$ , and with constant channels and scheduling thresholds

$$h_k^{[s]} = h_k, \quad \text{for } s = 1, \dots, S, \quad (43)$$

$$g^{[s]} = g, \quad \text{for } s = 1, \dots, S. \quad (44)$$

The solution in this special case will turn out to be especially insightful, and the more general problem will be studied in the next subsection. Under these simplifying assumptions, and using Lemma 3, the min-max optimization in (41) reduces to the problem

$$\min_{\{\alpha^{[s]}\}_{s=1}^{S_b+1}} \sum_{s=1}^{S_b+1} \left(\frac{1+\gamma}{2}\right)^{-2s} \max \left\{ 0, \frac{\eta^2 N_0 K^2}{(\alpha^{[s]} K_a)^2} - 2\eta \right\} \quad (45a)$$

$$\text{s.t.} \quad \sum_{s=1}^{S_b+1} \frac{2(\alpha^{[s]} \ell)^2}{N_0} \leq \mathcal{R}_{\text{dp}}(\epsilon, \delta), \quad (45b)$$

$$\frac{(\alpha^{[s]} \ell)^2}{|h_k|^2} \leq P, \quad \forall k \in \mathcal{K}_a, s = 1, \dots, S \quad (45c)$$

$$\alpha^{[s]} \leq \frac{K}{K_a} \sqrt{\frac{\eta N_0}{2}}, \quad \forall k \in \mathcal{K}_a, s = 1, \dots, S. \quad (45d)$$

**Theorem 1.** Under Assumptions 1-3, and assuming static channels and thresholds as in (43)–(44), the optimal solutions of problem (45) depends on power  $P$  and learning rate  $\eta$  according to the three regimes illustrated in Fig. 2, which are detailed as follows.

1) *LMC-limited Regime:* If the condition

$$\eta \leq \frac{K_a^2}{\ell^2 K^2} \min \left\{ \frac{\mathcal{R}_{\text{dp}}(\epsilon, \delta)}{S}, \min_{k \in \mathcal{K}_a} \frac{2P|h_k|^2}{N_0} \right\} \quad (46)$$

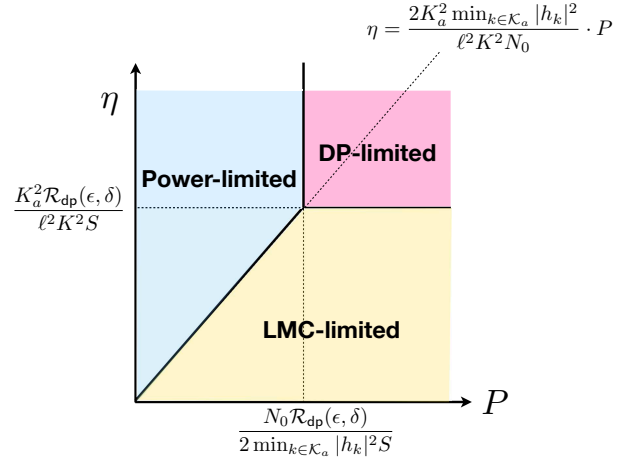


Figure 2. Illustration of the different regimes that describe the optimal power control strategy in Theorem 1.

holds, the optimal power gain parameter is given as

$$\alpha_{\text{opt}}^{[s]} = \frac{K}{K_a} \sqrt{\frac{\eta N_0}{2}}, \quad s = 1, \dots, S. \quad (47)$$

2) *Power-limited Regime:* If the condition

$$P \leq \frac{N_0}{2 \min_{k \in \mathcal{K}_a} |h_k|^2} \min \left\{ \frac{\mathcal{R}_{\text{dp}}(\epsilon, \delta)}{S}, \frac{\ell^2 K^2 \eta}{K_a^2} \right\} \quad (48)$$

holds, the optimal power gain parameter is given as

$$\alpha_{\text{opt}}^{[s]} = \min_{k \in \mathcal{K}_a} \frac{\sqrt{P}|h_k|}{\ell}, \quad s = 1, \dots, S. \quad (49)$$

3) *DP-limited Regime:* Otherwise, we have optimal solution

$$\alpha_{\text{opt}}^{[s]} = \min \left\{ \left(\frac{1+\gamma}{2}\right)^{-s/2} \sqrt{\frac{\eta K N_0^{1/2}}{K_a \lambda^{1/2}}}, \min_{k \in \mathcal{K}_a} \frac{\sqrt{P}|h_k|}{\ell}, \frac{K}{K_a} \sqrt{\frac{\eta N_0}{2}} \right\}, \quad (50)$$

where the value of  $\lambda$  can be obtained by bisection to satisfy the condition  $\sum_{s=1}^{S_b+1} (\alpha_{\text{opt}}^{[s]})^2 = \frac{N_0 \mathcal{R}_{\text{dp}}(\epsilon, \delta)}{2\ell^2}$  for the active devices  $k \in \mathcal{K}_a$ .

*Proof:* The problem is seen to be convex by changing variables  $(\alpha^{[s]})^2 = a^{[s]} \geq 0$  and the solution approach involves applying Lagrange multiplier method and Karush-Kuhn-Tucker (KKT) conditions in a manner similar to [6, Theorem 1]. Details can be found in [46, Appendix D].

The three regimes highlighted in Fig. 2 correspond to settings where each of the corresponding constraints, i.e., (47), (49) or (50), are active. Accordingly, we have the following observations:

- In the LMC-limited regime, the channel noise contributes in full to LMC update, and hence the presence of channel noise and DP does not affect the performance of LMC. In contrast, in the other two regimes, the additional gradient noise causes a loss as compared to a noiseless implementation of LMC.
- In the DP-limited regime, the privacy constraint affects the performance of WFLMC, while in the other regimes,

privacy is obtained “for free”, i.e., as a direct consequence of the presence of channel noise. To optimize the learning performance, the power control parameter should be adaptive across the iterations.

### C. Optimization of Power Gain Parameters: General Case

In this subsection, we consider the general optimization problem (41) with any  $S_u \geq 1$  and under time-varying channels. To this end, we include LMC noise requirement leveraging Lemma 3, and start by rewriting the min-max optimization (41) in the epigraph form

#### (Multi-Sample Opt.)

$$\min_{\{\alpha^{[s]}\}_{s=1}^S} \nu \quad (51a)$$

$$\text{s.t. } \sum_{s=1}^S \frac{\mathbf{1}[|h_k^{[s]}| \geq g^{[s]}] \cdot 2(\alpha^{[s]}\ell)^2}{N_0} \leq \mathcal{R}_{\text{dp}}(\epsilon, \delta), \quad \forall k \quad (51b)$$

$$\left( \frac{\mathbf{1}[|h_k^{[s]}| \geq g^{[s]}] \alpha^{[s]}\ell}{h_k^{[s]}} \right)^2 \leq P, \quad \forall k, s \quad (51c)$$

$$\alpha^{[s]} \leq \frac{K}{K_a^{[s]}} \sqrt{\frac{\eta N_0}{2}}, \quad \forall s \quad (51d)$$

$$\widetilde{W}_2(p^{[s']}, p(\theta|D))^2 \leq \nu, \quad \forall s' = S_b + 1, \dots, S. \quad (51e)$$

Having made the change of variables  $(\alpha^{[s]})^2 = a^{[s]} \geq 0$ , problem (51) can be easily seen to be convex. Therefore, it can be solved using standard numerical tools.

### D. Optimization of Truncated Thresholds

In this subsection, we turn to the problem of optimizing (41a) over the thresholds  $\{g^{[s]}\}_{s=1}^S$ . As mentioned, this problem is characterized by a non-differentiable objective function, and it should be addressed jointly with the optimization of the power gain parameters  $\{\alpha^{[s]}\}_{s=1}^S$ . To make progress, we propose a sub-optimal approach that decouples the two problems by setting the value of  $\alpha^{[s]}$  to ensure equality in the constraint (41c) as

#### (Full Power Transmission)

$$\alpha^{[s]} = \min_{k \in \mathcal{K}_a^{[s]}} \frac{\sqrt{P} h_k^{[s]}}{\ell}, \quad s = 1, \dots, S. \quad (52)$$

Note that the choice in (52) is only made for the purpose of optimizing the thresholds. After the thresholds  $\{g^{[s]}\}_{s=1}^S$  are optimized as explained next, the power gain parameters  $\{\alpha^{[s]}\}_{s=1}^S$  are selected by following the previous subsections.

Using (52), the min-max problem (41a) over  $\{g^{[s]}\}_{s=1}^S$  can be expressed as the sum of gradient error over  $s'$  communication rounds, for  $s' = S_b + 1, \dots, S$

$$\min_{\{g^{[s]}\}_{s=1}^S} \sum_{s=1}^{s'} \frac{2(1+\gamma)^{2(s-s')+1}}{4^{(s'-s)}(1-\gamma)} \left[ 4\eta^2 \ell^2 (K - K_a^{[s]})^2 + \widetilde{\beta}^{[s]} \right], \quad \text{for } s' \in [S_b + 1, S] \quad (53)$$

where we have dropped the constraints since they are assumed to be dealt with by the subsequent optimization over  $\{\alpha^{[s]}\}_{s=1}^S$ . The minimization (53) can be addressed as  $s'$  parallel optimizations, and it is equivalent to focus on  $s' = S$  as

$$\min_{g^{[s]}} \max \left\{ 0, \frac{N_0 K^2 \ell^2}{P (K_a^{[s]})^2 \min_{k \in \mathcal{K}_a^{[s]}} |h_k^{[s]}|^2} - \frac{2}{\eta} \right\} + 4\ell^2 (K - K_a^{[s]})^2, \quad \text{for } s = 1, \dots, S, \quad (54)$$

where the second term is obtained by the definition of  $\widetilde{\beta}^{[s]}$  in (27). The first term is estimation error depends on the bound  $\ell$  of the norm of the local gradient as per Assumption 3. We note that this result can be directly extended to device-dependent bounds  $\ell_k$  for each device  $k$ . With this extension, the threshold  $g^{[s]}$  in the scheduling rule (22) would depend on the properties of the local data set via constants  $\{\ell_k\}_{k=1}^K$ .

The threshold  $g^{[s]}$  affects the objective function (54) through the number  $K_a^{[s]}$  of active users. Decreasing  $K_a^{[s]}$  aggravates the estimation error (first term) due to partial scheduling, while providing a chance to alleviate the excess noise (second term) by silencing the devices with the worst channels. The objective function (54) can take at most  $K$  different values that are attained by setting  $g^{[s]}$  as one of the channel gains  $|h_1^{[s]}|, \dots, |h_K^{[s]}|$ . Therefore, we can limit the search to these values without loss of optimality. It follows that the optimal solution for each threshold  $g^{[s]}$  can be obtained by exhaustive search over the  $K$  values  $\{|h_k^{[s]}\}_{k=1}^K$  by minimizing (54).

## VI. NUMERICAL RESULTS

In this section, we investigate the effectiveness of the proposed scheme, WFLMC, as a mechanism to implement LMC on wireless channels using numerical experiments. We emphasize that our goal is not that of comparing the performance of Bayesian and frequentist techniques. This is a subject that has been extensively explored in the literature (see Sec. I), and we consider it to be beyond the scope of this contribution. Rather, our focus is on evaluating the effectiveness of the specific proposed implementation mechanism of LMC as compared to more conventional solutions. To this end, we consider the following benchmark schemes.

- 1) **WFMLC with equal power allocation:** This reference scheme follows the approach in [31] of dividing up the DP constraint equally across all communication rounds. By Lemma 2, this corresponds to imposing the constraint

$$\frac{2(\alpha^{[s]}\ell)^2}{N_0} < \frac{\mathcal{R}_{\text{dp}}(\epsilon, \delta)}{\sum_{s=1}^S \mathbf{1}[|h_k^{[s]}| \geq g^{[s]}}] \quad (55)$$

for each communication round. Condition (55), along with the power constraint (41c) and LMC noise requirement (42) yield the power scaling gains

$$\alpha^{[s]} = \min \left\{ \frac{1}{\ell} \sqrt{\frac{N_0 \mathcal{R}_{\text{dp}}(\epsilon, \delta)}{2 \sum_{s=1}^S \mathbf{1}[|h_k^{[s]}| \geq g^{[s]}}]}}, \min_{k \in \mathcal{K}_a^{[s]}} \frac{\sqrt{P} |h_k^{[s]}|}{\ell}, \frac{K}{K_a^{[s]}} \sqrt{\frac{\eta N_0}{2}} \right\}, \quad \forall s. \quad (56)$$

- 2) **WFLMC without DP constraint:** In this scheme, the power gain parameter  $\alpha^{[s]}$  is set as

$$\alpha^{[s]} = \min \left\{ \min_{k \in \mathcal{K}_a^{[s]}} \frac{\sqrt{P} |h_k^{[s]}|}{\ell}, \frac{K}{K_a^{[s]}} \sqrt{\frac{\eta N_0}{2}} \right\}, \forall s, \quad (57)$$

which corresponds to the optimal solution without the DP constraint.

We also consider for reference a scenario characterized by ideal communication without fading (i.e.,  $h_k = 1$ ), and channel noise (i.e.,  $N_0 = 0$ ). To guarantee privacy and implement LMC update, noise is added at each devices before transmission as

$$\mathbf{x}_k^{[s]} = \nabla f_k(\boldsymbol{\theta}^{[s-1]}) + \sqrt{\sigma^{[s]}} \mathbf{n}_k^{[s]}, \quad (58)$$

and  $\mathbf{n}_k^{[s]} \sim \mathcal{N}(0, \mathbf{I}_m)$ . The variance  $\sqrt{\sigma^{[s]}}$  depends on the DP constraint. Specifically, following the analysis in WFLMC, we can define the following two additional benchmark schemes.

- 3) **Noiseless federated LMC with DP constraint:** By solving problem (41) without power constraint (41c), the optimal value  $(\alpha^{[s]})_{\text{opt}'}$  is used to set the variance of additive noise as  $\sigma^{[s]} = N_0 / [K(\alpha^{[s]})_{\text{opt}'}^2]$  in order to satisfy the DP constraint.
- 4) **Noiseless federated LMC without DP constraint:** Removing the DP constraints, the implementation of the LMC update (10) requires the variance of update noise in (12) to equal  $2\eta$ , and by (25) we have  $\sigma^{[s]} = 2/(K\eta)$ .

As for the learning model, we consider a Gaussian linear regression with likelihood

$$p(v_n | \boldsymbol{\theta}, \mathbf{u}_n) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(v_n - \boldsymbol{\theta}^\top \mathbf{u}_n)^2}, \quad (59)$$

and the prior  $p(\boldsymbol{\theta})$  is assumed to follow Gaussian distribution  $\mathcal{N}(0, \mathbf{I}_m)$ . Therefore, the posterior  $p(\boldsymbol{\theta} | \mathcal{D})$  is the Gaussian  $\mathcal{N}((\mathbf{U}\mathbf{U}^\top + \mathbf{I})^{-1} \mathbf{U}\mathbf{v}, (\mathbf{U}\mathbf{U}^\top + \mathbf{I})^{-1})$ , where  $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_N]$  is the data matrix and  $\mathbf{v} = [v_1, \dots, v_N]^\top$  is the label vector. The strong convexity parameter  $\mu$  and smoothness parameter  $L$  are computed as the smallest and largest eigenvalues of the data Gramian matrix  $\mathbf{U}\mathbf{U}^\top + \mathbf{I}$ . We will also consider experiments with the MNIST data set at the end of this section.

We consider a synthetic data set  $\{\mathbf{d}_n = (\mathbf{u}_n, v_n)\}_{n=1}^N$  with  $N = 1200$  following the model (59) with covariates  $\mathbf{u}_n \in \mathbb{R}^m$  drawn i.i.d. from Gaussian distribution  $\mathcal{N}(0, \mathbf{I}_m)$  where  $m = 5$ , and ground-truth model parameter  $\boldsymbol{\theta}^* = [0.071, -0.518, 0.9342, 0.7198, 0.4676]^\top$ . The Wasserstein distance between two Gaussian distribution  $p_1 = \mathcal{N}(\mathbf{m}_1, \mathbf{C}_1)$  and  $p_2 = \mathcal{N}(\mathbf{m}_2, \mathbf{C}_2)$  is computed as [47]

$$W_2(p_1, p_2)^2 = \text{Tr} \left( \mathbf{C}_1 + \mathbf{C}_2 - 2(\mathbf{C}_2^{1/2} \mathbf{C}_1 \mathbf{C}_2^{1/2})^{1/2} \right) + \|\mathbf{m}_1 - \mathbf{m}_2\|^2. \quad (60)$$

Unless stated otherwise, we consider homogeneous local data distributions by dividing the data set equally among the  $K = 30$  devices. In this regard, we observe that frequentist federated learning is well known to be significantly affected by the availability of non-i.i.d. data at the devices. This is because the different training objectives at the devices,

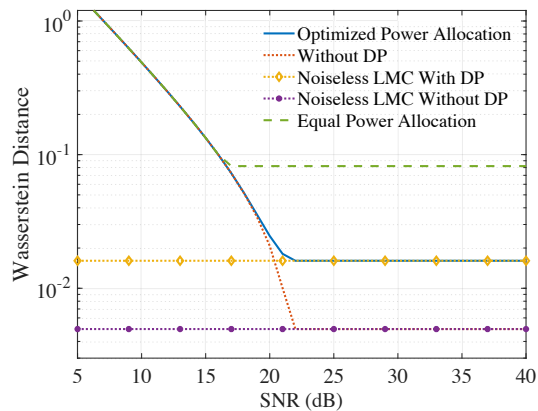

which are entailed by the non-i.i.d. assumption, may cause the local optimization updates to favor different parts of the optimization space [48], [49]. Fundamentally, frequentist federated learning is limited by the fact that there may not be a single model parameter vector that generalizes well for all the non-i.i.d. local distributions. In contrast, Bayesian learning can retain information about all local non-i.i.d. distributions in distinct modes of the global posterior distribution. Therefore, the non-i.i.d. property of the local distributions can be in principled detected and accounted for. We leave a full study of this aspect to future work, and point to [23], [26] for some related observations. That said, in the presence of privacy constraint, the need for clipping causes centralized and federated implementations of LMC to be different, and the performance loss of federated against centralized LMC becomes more pronounced for a smaller clipping threshold in the presence of heterogeneous distribution. Here we set the clipping threshold to  $\ell = 30$ , and leave an investigation of the interplay between data heterogeneity and clipping to future work. Furthermore, the initial sample  $\boldsymbol{\theta}^{[0]}$  is drawn from prior with discrepancy  $W_2(p^{[0]}(\boldsymbol{\theta}), p(\boldsymbol{\theta} | \mathcal{D}))^2 = 6.64$ .

#### A. Single-Sample Case

We now focus on single-sample case ( $S_u = 1$ ) by applying the optimal power allocation detailed in Theorem 1 while scheduling all devices ( $K_a^{[s]} = K$ ). As in Theorem 1, the channels  $h_k^{[s]}$  are constant and set to 0.01 for all devices  $k$ ; the number of communication burn-in periods is set to  $S_b = 50$ ; DP parameters are given as  $\delta = 0.01$  and  $\epsilon = 8$ .

• **Impact of the SNR:** In Fig. 3, we plot the bound (34) on the 2-Wasserstein distance versus SNR, defined as  $P/(mN_0)$ , by varying transmitted power  $P$ . The learning rate is set as  $\eta = 0.4/(\mu + L)$  and  $0.13/(\mu + L)$  in order to illustrate the different regimes defined in Fig. 2. In Fig. 3(a), under the larger learning rate, the performance is limited by the power constraint until SNR = 16.6 dB, after which it becomes limited by DP. In the DP-limited regime, optimized power allocation attains better performance than static power allocation. For SNR > 22 dB, the performance under optimized power allocation becomes equivalent to that of the system with ideal noiseless communication with the DP, showing that all channel noise is repurposed for DP mechanism while only partial is for MC sampling. That is, imposing the DP constraint causes some performance loss due to the need of scaling down the transmission power to decrease the effective SNR. In line with the results in Fig. 2, Fig. 3(b) shows the transition from power-limited regime to LMC-limited regime for the smaller learning rate. In the LMC-limited regime, all the schemes have same performance since the channel noise variance is only determined by the LMC constraint.

• **Impact of the learning rate:** We now further elaborate on the impact of the learning rate  $\eta$  in Fig. 4 by setting SNR = 18 dB and SNR = 30 dB. First, with  $\eta \leq 0.5 \times 10^{-4}$ , by Fig. 2, we are in the LMC-limited regime where all the schemes have same performance. When  $\eta > 0.5 \times 10^{-4}$ , the performance of optimized power allocation is shown to be power-limited for SNR = 18 dB in Fig. 4(a), as it is identical to that without


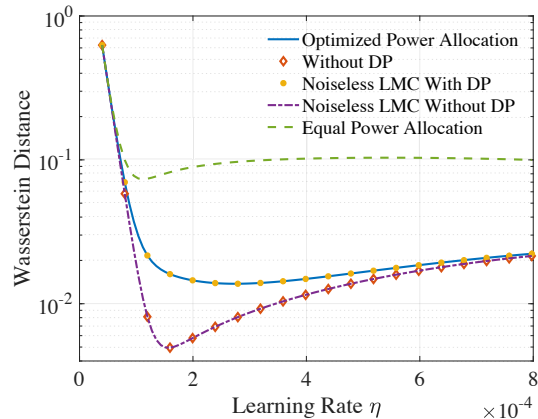
(a)  $\eta = 0.4/(\mu + L)$ .(b)  $\eta = 0.13/(\mu + L)$ .Figure 3. Bound (34) on the 2-Wasserstein distance versus learning rate  $\eta$  for WFLMC and noiseless LMC schemes ( $\epsilon = 8$ ,  $\delta = 0.01$ ,  $S_b = 50$ ,  $S_u = 1$ ).

the DP constraint. In contrast, when  $\text{SNR} = 30$  dB, as seen in Fig. 4(b), the performance of optimized power allocation is limited by DP for  $\eta > 0.5 \times 10^{-4}$ . As for the static power allocation, the performance is always limited by DP for  $\eta > 0.5 \times 10^{-4}$ , and increasing the SNR is not helpful in reducing the approximation error.

• **Impact of the privacy level:** Fig. 5 plots the bound (34) on the 2-Wasserstein distance as a function of the privacy level  $\epsilon$ . The performance of WFLMC is limited by DP until  $\epsilon = 13.5$ , after which it becomes limited by the transmitted power and the DP constraint does not cause a performance loss. In the DP-limited regime, the proposed optimized power allocation outperforms static power allocation. Furthermore, under a stricter DP requirement, i.e., with  $\epsilon \leq 6$ , the performance of optimized power allocation is equivalent to that under ideal communication, and the existence of channel noise does not impair performance.

### B. Multi-Sample Case

We obtain the power control strategy in the multi-sample case by solving the convex problem (51) under the suboptimal scheduling policy in (54). The optimized parameters are then used to implement WLMC. To evaluate the 2-Wasserstein distance, we numerically estimate the mean and covariance

(a)  $\text{SNR} = 18$  dB.(b)  $\text{SNR} = 30$  dB.Figure 4. Bound (34) on the 2-Wasserstein distance versus learning rate  $\eta$  for WFLMC and noiseless LMC schemes ( $\epsilon = 8$ ,  $\delta = 0.01$ ,  $S_b = 50$ ,  $S_u = 1$ ).

matrices of the samples over 100 experiments, and compare with the true posterior via (60). The results present as the worst performance after the burn-in period, as given in the objective (41a). Unless stated otherwise, the channels  $h_k^{[s]}$  are randomly generated following  $\mathcal{CN}(0, 0.01)$  for all  $k$ ; the number of communication round is set to  $S = 100$ ; DP parameters are set to  $\delta = 0.01$  and  $\epsilon = 15$ .

• **Impact of the SNR:** We first study the impact of the  $\text{SNR} = P/(mN_0)$  with number of samples  $S_u = 50$ . The Wasserstein distances of all the wireless LMC schemes are seen to decrease with the SNR until approaching the DP-limited regime for optimized power allocation, and the LMC-limited regime for the scheme without DP constraint. The results emphasize the importance of optimizing power allocation in high SNR regime where the static power allocation is seen to have a significant performance degradation in DP-limited regime. Furthermore, they validate the insights obtained for the analysis under the simplified assumptions considered in Theorem 1.

• **Impact of the Burn-in Period:** In Fig. 7, we investigate the impact of the number of communication periods,  $S_b$ , allocated for the burn-in period where we fix as  $S = 100$  the total number of communication rounds. In this experiment, the SNR is set to 30 dB. Fig. 7 shows that increasing the communication




Figure 5. Bound (34) on the 2-Wasserstein distance versus DP level  $\epsilon$  for WFLMC and noiseless LMC schemes ( $\eta = 0.4/(\mu + L)$ ,  $\delta = 0.01$ , SNR = 20 dB,  $S_b = 50$ ,  $S_u = 1$ ).




Figure 6. Evaluation of the 2-Wasserstein distance (60) versus SNR for WFLMC and noiseless LMC schemes ( $\epsilon = 15$ ,  $\delta = 0.01$ ,  $S = 100$ ,  $S_u = 50$ ).

rounds in the burn-in period helps improve the quality of the samples produced after the burn-in period, as the Wasserstein distances of all the schemes are seen to decrease with the burn-in period. We also note that the enhanced sample quality costs at the cost of sample quantity. The performance of LMC is limited by the bias in discretization error that is determined by the learning rate and can not be diminished by increasing the duration of the burn-in period. Overall, the results emphasize that sample quality can be enhanced by optimizing the power allocation strategy.

### C. MNIST Data Set

We now consider the problem of multinomial logistic regression on the MNIST data set, which is comprised of  $C = 10$  classes representing handwritten digits. We use  $N = 12,000$  data points, and  $N/C$  data points for each class are randomly selected from the MNIST training set. The original data, with dimension 784, is pre-processed by projecting the input images into a subspace of lower dimension 30 via principal component analysis (PCA). This is motivated by the fact that the MNIST images are known to be approximately supported on a manifold with intrinsic dimension lower than 30 [50].




Figure 7. Evaluation of the 2-Wasserstein distance (60) versus the number of communication rounds in burn-in period  $S_b$  for WFLMC and noiseless LMC schemes ( $\epsilon = 15$ ,  $\delta = 0.01$ ,  $S = 100$ , SNR = 30 dB).




Figure 8. Evaluation of expected calibration error versus SNR for WFLMC, noiseless LMC schemes, and SBFL in [22] ( $\epsilon = 50$ ,  $\delta = 0.1$ ,  $S = 500$ ,  $S_u = 50$ ).

For the learning model, the likelihood is considered as

$$p(v_n | \{\theta_c\}_{c=1}^C, \mathbf{u}_n) = \prod_{c=1}^C \left[ \frac{\exp(\theta_c^T \mathbf{u}_n)}{\sum_{c'=1}^C \exp(\theta_{c'}^T \mathbf{u}_n)} \right]^{I_{n,c}}, \quad (60)$$

where we have  $I_{n,c} = \mathbf{1}[v_n = c]$ . The prior of the model parameters  $\theta = \{\theta_c\}_{c=1}^C$  is a produced Gaussian distribution  $\mathcal{N}(0, \mathbf{I}_m)$  for  $m = 300$ . To set the parameters  $\mu$  and  $L$ , we leverage the inequality for the Hessian matrix  $\mathbf{I}_m \preceq \nabla^2 f(\theta) \preceq \frac{1}{2} (\mathbf{I}_c - \frac{1}{C} \mathbf{1}_{c \times c}) \otimes \mathbf{U} \mathbf{U}^T + \mathbf{I}_m$ , where  $\mathbf{1}_{c \times c}$  is a  $c \times c$  matrix with all elements 1, and  $\otimes$  represents Kronecker product. Specifically, we compute the minimum and maximum eigenvalues of the lower bound and upper bound respectively to set  $\mu$  and  $L$ . The channels  $h_k^{[s]}$  are randomly generated following the distribution  $\mathcal{CN}(0, 10^{-4})$  for all  $k$ ; the number of communication round is set to  $S = 600$ ; the number of communication burn-in periods is set to  $S_b = 550$ ; the clipping threshold in (20) is set to  $\ell = 300$ ; and DP parameters are set to  $\delta = 0.1$  and  $\epsilon = 50$ . The results are averaged over 50 experiments.

For tractability, the optimization problem aims at minimizing the averaged Wasserstein distances across the samples after the burn-in period. The problem is convex and can be solved by the method in Sec V-B. We use the test set of MNIST to evaluate the expected calibration error [10]. Furthermore, we

implemented scalable Bayesian federated learning (SBFL) in [22] as another benchmark. We recall that SBFL is a frequentist learning protocol that aggregates the local gradients at the server via a point estimate that incorporates prior knowledge on the gradients. Under the assumption of orthogonal multiple access in [22], the iteration time is set as  $S/K$  for a fair comparison.

Fig. 8 plots the expected calibration error versus SNR with the number of samples  $S_u = 50$ . First, the figure verifies the advantage of Bayesian learning, as all WFLMC schemes are seen to significantly outperform SBFL. Furthermore, by increasing the number of communication rounds, optimized power allocation can outperform the equal power allocation scheme.

## VII. CONCLUSIONS

In this paper, we have proposed a novel Bayesian federated learning (FL) protocol that implements Langevin Monte Carlo (LMC) via uncoded wireless transmission from devices to edge server. The key novel idea is that of repurposing channel noise for both MC sampling in Bayesian FL and privacy preservation. Under simplified assumptions, the analysis has revealed that system operated in different regimes limited by LMC, DP, or power constraints, depending on the values of learning rate, privacy parameters, and transmitted power.

As an extension of the current work, it would be interesting to optimize the learning rate schedule; to study digital implementations of wireless federated LMC as in [7]; to consider multi-hop device-to-device (D2D) network topologies; and to account for non-ideal downlink communication as in [51], [52].

## APPENDIX

### A. Proof of Lemma 2

The vector  $\bar{\theta}^{(t)}$  is the continuous-time Langevin diffusion process in (7) and its distribution is assumed to be invariant, so that we have  $N_d^{[s]} = N_d^{[1]}$  for all  $s$ . The proof then follows from [35, Lemma 3] with the caveat that we bound of  $E[\|N_d^{[1]}\|^2]$  in lieu of  $\sqrt{E[\|N_d^{[1]}\|^2]}$ .

### B. Proof of Proposition 1

For any  $\tau > 0$ , the error (37) is bounded as

$$\begin{aligned} E \left[ \left\| \theta^{[s]} - \bar{\theta}^{(s\eta)} \right\|^2 \right] &\leq (1 + \tau) E \left[ \left\| \theta^{[s-1]} - \bar{\theta}^{((s-1)\eta)} \right\|^2 \right. \\ &\quad \left. - \eta \left[ \nabla f(\theta^{[s-1]}) - \nabla f(\bar{\theta}^{((s-1)\eta)}) \right] \right] \\ &\quad + (1 + \tau^{-1}) \left( 2\eta^2 E[\|N_g^{[s]}\|^2] + 2E[\|N_d^{[s]}\|^2] \right), \end{aligned} \quad (61)$$

where the value of  $\tau$  controls the convergence rate as detailed later. Furthermore, the first term is bounded by

$$\begin{aligned} E \left[ \left\| \theta^{[s-1]} - \bar{\theta}^{((s-1)\eta)} - \eta \left[ \nabla f(\theta^{[s-1]}) - \nabla f(\bar{\theta}^{((s-1)\eta)}) \right] \right\|^2 \right] \\ \leq \gamma^2 E \left[ \left\| \theta^{[s-1]} - \bar{\theta}^{((s-1)\eta)} \right\|^2 \right], \end{aligned} \quad (62)$$

where  $\gamma = 1 - \eta\mu$  if  $0 < \eta \leq 2/(\mu + L)$ ; and  $\gamma = \eta L - 1$  if  $2/(\mu + L) \leq \eta \leq 2/L$ . The proof starts with (19) through [35, Lemma 1]. Specifically, plugging Lemmas 1 and 2 and (62) into (61), and setting  $\tau = \left(\frac{1+\gamma}{2}\right)^2 - 1$  yield

$$\begin{aligned} E \left[ \left\| \theta^{[s]} - \bar{\theta}^{(s\eta)} \right\|^2 \right] &\leq \left( \frac{1+\gamma}{2} \right)^2 E \left[ \left\| \theta^{[s-1]} - \bar{\theta}^{((s-1)\eta)} \right\|^2 \right] \\ &\quad + \frac{2(1+\gamma)^2}{(1-\gamma)(1+3\gamma)} \left[ \frac{\eta^4 L^3 m}{3} + 4\eta^2 \ell^2 (K - K_a^{[s]})^2 \right. \\ &\quad \left. + \eta^3 L^2 m + \max \left\{ 0, \frac{\eta^2 N_0 K^2}{(\alpha^{[s]} K_a^{[s]})^2} - 2\eta \right\} \right]. \end{aligned} \quad (63)$$

The desired result in Proposition 1 is obtained by using the inequality  $\frac{1+\gamma}{1+3\gamma} < 1$ , applying (63) recursively, and the initial point  $\theta^{[0]}$  is set to satisfy  $W_2(p^{[0]}(\theta), p(\theta|\mathcal{D}))^2 = E[\|\theta^{[0]} - \bar{\theta}^{(0)}\|^2]$ .

## REFERENCES

- [1] J. Park, S. Samarakoon, M. Bennis, and M. Debbah, "Wireless network intelligence at the edge," *Proc. IEEE*, vol. 107, no. 11, pp. 2204–2239, 2019.
- [2] G. Zhu, D. Liu, Y. Du, C. You, J. Zhang, and K. Huang, "Toward an intelligent edge: wireless communication meets machine learning," *IEEE Commun. Mag.*, vol. 58, pp. 19–25, Jan. 2020.
- [3] T. Sery, N. Shlezinger, K. Cohen, and Y. C. Eldar, "Over-the-air federated learning from heterogeneous data," *IEEE Trans. Signal Process.*, vol. 69, pp. 3796–3811, June 2021.
- [4] H. H. Yang, Z. Chen, T. Q. Quek, and H. V. Poor, "Revisiting analog over-the-air machine learning: The blessing and curse of interference," [Online]. Available: <https://arxiv.org/pdf/2107.11733.pdf>, 2021.
- [5] 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, Oct. 2019.
- [6] 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, Nov. 2020.
- [7] G. Zhu, Y. Du, D. Gündüz, 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, pp. 2120–2135, March 2021.
- [8] C. Dwork, A. Roth, et al., "The algorithmic foundations of differential privacy," *Foundations and Trends® in Theoretical Computer Science*, vol. 9, no. 3–4, pp. 211–407, 2014.
- [9] Y. Koda, K. Yamamoto, T. Nishio, and M. Morikura, "Differentially private aircomp federated learning with power adaptation harnessing receiver noise," in *Proc. IEEE Glob. Commun. Conf. (GLOBECOM)*, (Virtual), Dec. 2020.
- [10] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," in *Proc. Intl. Conf. Mach. Learning (ICML)*, (Sydney, Australia), pp. 1321–1330, Aug. 2017.
- [11] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," [Online]. Available: <https://arxiv.org/pdf/1612.01474.pdf>, 2016.
- [12] E. Angelino, M. J. Johnson, and R. P. Adams, "Patterns of scalable bayesian inference," [Online]. Available: <https://arxiv.org/pdf/1602.05221.pdf>, 2016.
- [13] Y.-A. Ma, T. Chen, and E. Fox, "A complete recipe for stochastic gradient MCMC," in *Proc. Adv. Neural Info. Proc. Syst. (NIPS)*, vol. 28, (Montreal, Canada), Dec. 2015.
- [14] A. S. Dalalyan and L. Riou-Durand, "On sampling from a log-concave density using kinetic Langevin diffusions," *Bernoulli*, vol. 26, no. 3, pp. 1956–1988, 2020.
- [15] D. Zou and Q. Gu, "On the convergence of Hamiltonian Monte Carlo with stochastic gradients," in *Proc. Conf. Mach. Learning (ICML)*, (Virtual), pp. 13012–13022, July 2021.
- [16] 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. Syst. (NIPS)*, (Long Beach, USA), Dec. 2017.

- [17] S. Wang, T. Tuor, T. Salonidis, K. K. Leung, C. Makaya, T. He, and K. Chan, "Adaptive federated learning in resource constrained edge computing systems," *IEEE J. Sel. Areas Commun.*, vol. 37, no. 6, pp. 1205–1221, 2019.
- [18] N. H. Tran, W. Bao, A. Zomaya, M. N. Nguyen, and C. S. Hong, "Federated learning over wireless networks: Optimization model design and analysis," in *Proc. Intl. Conf. Comput. Commun. (INFOCOM)*, (Paris, France), pp. 1387–1395, IEEE, May 2019.
- [19] 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, 2020.
- [20] 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.
- [21] D. Fan, X. Yuan, and Y.-J. A. Zhang, "Temporal-structure-assisted gradient aggregation for over-the-air federated edge learning," [Online]. Available: <https://arxiv.org/pdf/2103.02270.pdf>, 2021.
- [22] S. Lee, C. Park, S.-N. Hong, Y. C. Eldar, and N. Lee, "Bayesian federated learning over wireless networks," [Online]. Available: <https://arxiv.org/pdf/2012.15486.pdf>, 2020.
- [23] H.-Y. Chen and W.-L. Chao, "Fedbe: Making bayesian model ensemble applicable to federated learning," in *Proc. Intl. Conf. Learning Representations (ICLR)*, (Virtual), May 2021.
- [24] Z. Zhu, J. Hong, and J. Zhou, "Data-free knowledge distillation for heterogeneous federated learning," [Online]. Available: <https://arxiv.org/pdf/2105.10056.pdf>, 2021.
- [25] L. Corinzia and J. M. Buhmann, "Variational federated multi-task learning," [Online]. Available: <https://arxiv.org/pdf/1906.06268.pdf>, 2019.
- [26] R. Kassab and O. Simeone, "Federated generalized bayesian learning via distributed Stein variational gradient descent," *IEEE Trans. Signal Process.*, vol. 70, pp. 2180–2192, April 2022.
- [27] K. el Mekkaoui, D. Mesquita, P. Blomstedt, and S. Kaski, "Federated stochastic gradient Langevin dynamics," in *Proc. 37th Conf. Uncertain. Artif. Intell. (UAI)*, (Virtual), July 2021.
- [28] M. Vono, V. Plassier, A. Durmus, A. Dieuleveut, and E. Moulines, "QLSD: Quantised Langevin stochastic dynamics for Bayesian federated learning," [Online]. Available: <https://arxiv.org/pdf/2106.00797.pdf>, 2021.
- [29] V. Gandikota, R. K. Maity, and A. Mazumdar, "vqSGD: Vector quantized stochastic gradient descent," [Online]. Available: <https://arxiv.org/pdf/1911.07971.pdf>, 2019.
- [30] N. Agarwal, A. T. Suresh, F. X. X. Yu, S. Kumar, and B. McMahan, "cpSGD: Communication-efficient and differentially-private distributed SGD," in *Proc. Adv. Neural Info. Proc. Syst. (NIPS)*, (Montreal, Canada), Dec. 2018.
- [31] M. Seif, R. Tandon, and M. Li, "Wireless federated learning with local differential privacy," in *Proc. IEEE Intl. Symp. Info. Theory (ISIT)*, (Los Angeles, USA), June 2020.
- [32] Z. Zhang, G. Zhu, R. Wang, V. K. Lau, and K. Huang, "Turning channel noise into an accelerator for over-the-air principal component analysis," [Online]. Available: <https://arxiv.org/pdf/2104.10095.pdf>, 2021.
- [33] Y.-X. Wang, S. Fienberg, and A. Smola, "Privacy for free: Posterior sampling and stochastic gradient Monte Carlo," in *Proc. Conf. Mach. Learning (ICML)*, (Lille, France), July 2015.
- [34] B. Li, C. Chen, H. Liu, and L. Carin, "On connecting stochastic gradient MCMC and differential privacy," in *Proc. Intl. Conf. Artif. Intell. Stat. (AISTATS)*, (Naha, Japan), April 2019.
- [35] A. Dalalyan, "Further and stronger analogy between sampling and optimization: Langevin Monte Carlo and gradient descent," in *Proc. Conf. Learning Theory (COLT)*, (Amsterdam, Netherlands), July 2017.
- [36] D. Liu and O. Simeone, "Channel-driven Monte Carlo sampling for Bayesian distributed learning in wireless data centers," *IEEE J. Sel. Areas Commun.*, vol. 40, pp. 562–577, Feb. 2022.
- [37] M. Rabinovich, E. Angelino, and M. I. Jordan, "Variational consensus Monte Carlo," in *Proc. Adv. Neural Info. Proc. Syst. (NIPS)*, (Montreal, Canada), Dec. 2015.
- [38] R. Bhattacharya *et al.*, "Criteria for recurrence and existence of invariant measures for multidimensional diffusions," *The Annals of Probability*, vol. 6, no. 4, pp. 541–553, 1978.
- [39] N. Chatterji, N. Flammarion, Y. Ma, P. Bartlett, and M. Jordan, "On the theory of variance reduction for stochastic gradient Monte Carlo," in *Proc. Intl. Conf. Mach. Learning (ICML)*, (Stockholm, Sweden), pp. 764–773, July 2018.
- [40] A. Mahmood, M. I. Ashraf, M. Gidlund, J. Torsner, and J. Sachs, "Time synchronization in 5G wireless edge: Requirements and solutions for critical-MTC," *IEEE Commun. Mag.*, vol. 57, pp. 45–51, Dec. 2019.
- [41] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *Proc. Intl. Conf. Mach. Learning (ICML)*, (Sydney, Australia), pp. 214–223, Aug. 2017.
- [42] S. Bubeck, "Convex optimization: Algorithms and complexity," *Found. Trends Mach. Learn.*, vol. 8, pp. 231–357, Nov. 2015.
- [43] X. Chen, Z. S. Wu, and M. Hong, "Understanding gradient clipping in private SGD: A geometric perspective," in *Proc. Adv. Neural Info. Proc. Syst. (NIPS)*, (Virtual), Dec. 2020.
- [44] J. Ren, Y. He, D. Wen, G. Yu, K. Huang, and D. Guo, "Scheduling for cellular federated edge learning with importance and channel awareness," *IEEE Trans. Wireless Commun.*, vol. 19, no. 11, pp. 7690–7703, 2020.
- [45] D. Liu, G. Zhu, J. Zhang, and K. Huang, "Data-importance aware user scheduling for communication-efficient edge machine learning," *IEEE Trans. Cogn. Commun. Netw.*, vol. 7, no. 1, pp. 265–278, 2020.
- [46] D. Liu and O. Simeone, "Wireless federated Langevin Monte Carlo: Repurposing channel noise for bayesian sampling and privacy," [Online]. Available: <https://arxiv.org/pdf/2108.07644.pdf>, 2021.
- [47] C. Sun, H. Yan, X. Qiu, and X. Huang, "Gaussian word embedding with a Wasserstein distance loss," [Online]. Available: <https://arxiv.org/pdf/1808.07016.pdf>, 2018.
- [48] J. Wang, Q. Liu, H. Liang, G. Joshi, and H. V. Poor, "Tackling the objective inconsistency problem in heterogeneous federated optimization," in *Proc. Adv. Neural Info. Proc. Syst. (NIPS)*, vol. 33, (Virtual), pp. 7611–7623, Dec. 2020.
- [49] H. Zhang, T. Wu, S. Cheng, and J. Liu, "Fedcos: A scene-adaptive federated optimization enhancement for performance improvement," [Online]. Available: <https://arxiv.org/pdf/2204.03174.pdf>, 2022.
- [50] M. Hein and J.-Y. Audibert, "Intrinsic dimensionality estimation of submanifolds in Rd," in *Proc. Intl. Conf. Mach. Learning (ICML)*, (Bonn, Germany), Aug. 2005.
- [51] X. Wei and C. Shen, "Federated learning over noisy channels: Convergence analysis and design examples," [Online]. Available: <https://arxiv.org/pdf/2101.02198.pdf>, 2021.
- [52] M. M. Amiri, D. Gündüz, S. R. Kulkarni, and H. V. Poor, "Convergence of federated learning over a noisy downlink," *IEEE Trans. Wireless Commun.*, vol. 21, pp. 1422–1437, March 2022.



**Dongzhu Liu** is a Lecturer (Assistant Professor) with the School of Computing Science, University of Glasgow. She received the B.Eng. degree from the University of Electronic Science and Technology of China (UESTC) in 2015, and the Ph.D. degree from The University of Hong Kong in 2019. She was a postdoctoral research associate in the Dept. of Engineering at King's College London from 2019 to 2021. Her research interests include edge intelligence, federated learning, and wireless communications. She was a recipient of Exemplary

Reviewer Award of IEEE Transactions on Wireless Communications.





**Osvaldo Simeone** (Fellow, IEEE) is a Professor of Information Engineering with the Centre for Telecommunications Research at the Department of Engineering of King's College London, where he directs the King's Communications, Learning and Information Processing lab. He received an M.Sc. degree (with honors) and a Ph.D. degree in information engineering from Politecnico di Milano, Milan, Italy, in 2001 and 2005, respectively. From 2006 to 2017, he was a faculty member of the Electrical and Computer Engineering (ECE) Department at

New Jersey Institute of Technology (NJIT), where he was affiliated with the Center for Wireless Information Processing (CWIP). His research interests include information theory, machine learning, wireless communications, neuromorphic computing, and quantum machine learning. Dr Simeone is a co-recipient of the 2022 IEEE Communications Society Outstanding Paper Award, the 2021 IEEE Vehicular Technology Society Jack Neubauer Memorial Award, the 2019 IEEE Communication Society Best Tutorial Paper Award, the 2018 IEEE Signal Processing Best Paper Award, the 2017 JCN Best Paper Award, the 2015 IEEE Communication Society Best Tutorial Paper Award and of the Best Paper Awards of IEEE SPAWC 2007 and IEEE WRECOM 2007. He was awarded an Open Fellowship by the EPSRC in 2022 and a Consolidator grant by the European Research Council (ERC) in 2016. His research has been also supported by the U.S. National Science Foundation, the European Commission, the European Research Council, the Vienna Science and Technology Fund, the European Space Agency, as well as by a number of industrial collaborations including with Intel Labs and InterDigital. He is the Chair of the Signal Processing for Communications and Networking Technical Committee of the IEEE Signal Processing Society and of the UK & Ireland Chapter of the IEEE Information Theory Society. He is currently a Distinguished Lecturer of the IEEE Communications Society, and he was a Distinguished Lecturer of the IEEE Information Theory Society in 2017 and 2018. Dr Simeone is the author of the textbook "Machine Learning for Engineers" published by Cambridge University Press, four monographs, two edited books, and more than 180 research journal and magazine papers. He is a Fellow of the IET, EPSRC, and IEEE.