Federated Learning for Wireless Applications: A Prototype

Varun Laxman Muttepawar[1,*], Arjun Mehra[1,*], Zubair Shaban*, Ranjitha Prasad*, J. Harshan[†]

*Indraprastha Institute of Information Technology Delhi, India; [†]Indian Institute of Technology Delhi, India

Abstract—Wireless embedded edge devices are ubiquitous in our daily lives, enabling them to gather immense data via onboard sensors and mobile applications. This offers an amazing opportunity to train machine learning (ML) models in the realm of wireless devices for decision-making. Training ML models in a wireless setting necessitates transmitting datasets collected at the edge to a cloud parameter server, which is infeasible due to bandwidth constraints, security, and privacy issues. To tackle these challenges, Federated Learning (FL) has emerged as a distributed optimization approach to the decentralization of the model training process. In this work, we present a novel prototype to examine FL's effectiveness over bandwidth-constrained wireless channels. Through a novel design consisting of Zigbee and NI USRP devices, we propose a configuration that allows clients to broadcast synergistically local ML model updates to a central server to obtain a generalized global model. We assess the efficacy of this prototype using metrics such as global model accuracy and time complexity under varying conditions of transmission power, data heterogeneity and local learning.

Index Terms—Software Defined Radios Federated Learning, Distributed Learning, Over the air, MNIST

I. INTRODUCTION

In the era of global connectivity and the Internet of Things (IoT), efficient and privacy-preserving machine learning (ML) solutions are crucial. Federated Learning (FL) [1], an innovative paradigm in ML, emerges as a powerful approach to train models across decentralized devices without compromising on data privacy. Essentially, FL advocates the aggregation of local ML model parameters to obtain a global ML model.

In this dynamic landscape of modern connectivity, where wireless devices permeate our daily lives, the significance of implementing FL on hardware, for example, using software-defined radios (SDRs), is apparent. SDRs serve as the backbone of decentralized networks, facilitating real-time communication. Harnessing the power of FL on such hardware not only optimizes model training but also ensures that privacy-preserving ML becomes a tangible reality. The key objectives of our work are the following:

- Privacy-Preserving ML: Establishing an FL protocol, which is privacy-preserving by design and transmits parameter updates over the wireless channel.
- Real-time ML: Demonstrating the seamless integration of learning with NI-USRPs (National Instruments-Universal Software Radio Peripheral) and Zigbee for real-time update communication.

- Decentralized Model Training: Showcasing the efficiency gains achieved by distributing the learning process across the hardware devices.

Our demo provides a hands-on guide for setting up an FL environment using USRPs and Zigbee devices. The convergence of FL principles with tangible hardware platforms opens new avenues for collaborative and privacy-conscious ML, thereby unlocking the potential of FL on wireless devices. This work also helps to bridge the gap between theoretical advancements [2], [3] and practical implementations in wireless FL.

II. FEDERATED LEARNING & FEDAVG ALGORITHM

Federated learning is a decentralized ML approach where model training occurs locally on the edge devices, and only the model updates, not the raw data, are shared with a central server. In particular, in Federated Averaging (FedAvg), clients share their parameter updates with the server. The server aggregates these updates by computing a weighted average which is assigned as a global ML model. The algorithm is summarized in the following steps:

1) Initialization: The central server initializes and broadcasts a global initialization to the clients.
2) Training & Local updates: Each client computes the local updates by training the model on its local data and transmits them to the central server.
3) Aggregation & Global model: The central server aggregates the local updates using a weighted average and updates the global model.
4) Re-transmission of Global model: The central server then broadcasts the new global model to all the clients, who train on it for the next round.
5) Convergence criteria are monitored, such as reaching an accuracy threshold or a number of iterations.

This collaborative training paradigm enables privacy-preserving communication, making it particularly suitable for scenarios with sensitive distributed data sources, such as mobile devices or edge computing environments. Each edge device is tasked with locally training a model based on its received signal data, capturing unique characteristics of the wireless environment. The goal is to create a global model that not only performs well on the individual signals observed by each edge device but also generalizes effectively across the diverse signal characteristics contributed by different devices. One such work is studied in [4], where the goal is to train models that can identify adversarial attacks in a decentralized fashion by exploring the usage of deep neural networks. The

intent is to create a strong and secure wireless network where numerous nodes work together to detect and eliminate attacks by implementing the principles of FL.

III. DETAILS OF THE DEMONSTRATION

In this section, we discuss the details of the proposed demonstration. Our experimental setup is based on the federated averaging protocol as described in the previous section. Throughout the training process, only model updates (in the form of gradients) are exchanged with a central server, thereby preserving local raw data. Through iterative updates and aggregation, a global model emerges, embodying generalized features of the data despite each client having observed only a subset of the data. We conducted the experiments on a reduced MNIST dataset, where we consider any 4 classes distributed across two clients. We simulate the following scenarios:

- Independent and identically distributed (IID) Setup: We shuffle the dataset and randomly assign half of the dataset to each client. Since we assign the samples randomly, the local data at the clients are inherently IID.
- Non-IID setup: We assign two classes to each of the clients with no overlap, i.e., the client's local dataset is sampled from different distributions and, hence, non-IID.

A. Local Learning

On the clients' side, ML model training involves learning the parameters of a Multi-Layer Perceptron (MLP) with three linear layers, producing approximately $12,000$ parameters per client. The learning rates for the IID and non-IID cases are 10^{-5} and 10^{-4}, respectively. The number of epochs is 10 by default. ReLU activation is applied, and the model is trained using cross-entropy loss with the Adam optimizer.

B. Hardware Setup

Our hardware setup is partitioned into three core components: Clients 1 and 2 and a central server as depicted in Fig. 1. As shown in Fig. 2 each of these segments is equipped with essential hardware components, including NI USRP 2900 and Xbee S2C devices. Here, the NI USRP 2900, mounted with an omnidirectional VERT 2450 antenna, is used for physical layer communication, and Xbee S2C devices are used for network-level synchronization. The USRPs leverage differential modulation to enhance system robustness and eliminate the need for channel estimation. The other key signal specifications encompass various parameters crucial for the communication system. These include an operating frequency of 2.5 GHz, pulse shaping using Root Raised Cosine, a reception gain of 20 dB, an IQ rate of 50k, a symbol rate of 6.25k with 8 samples per symbol, and transmission power of 20dB by default.

C. Synchronization Process

The synchronization process is a challenging aspect of implementing a wireless FL system since the transmission of parameter updates needs to be scheduled meticulously. We realize synchronization by employing Time Division Multiple Access (TDMA) to optimize the scheduling of transmissions.
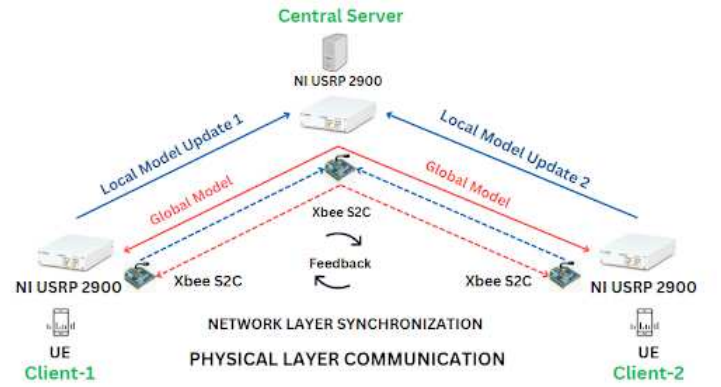

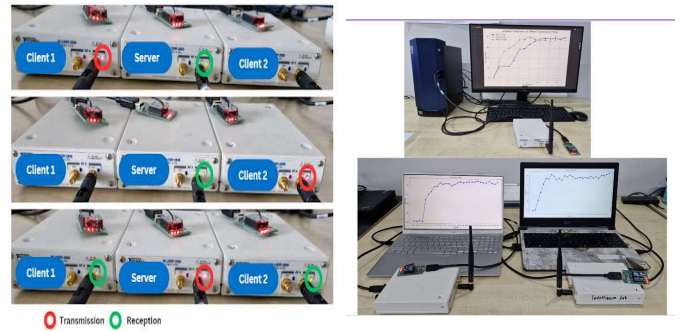Fig. 1: Synergizing Zigbee and USRP for FL Synchronization.

Fig. 2: FL prototype using USRPs and Xbee devices.

It commences with Client 1, transmitting its locally trained model parameters to the server (in $T_{c1,s}$ seconds). Upon receiving this data, the server promptly acknowledges the reception through the established Zigbee network. It is important to note that the server intentionally waits for a response from Client 1 before the server signals Client 2 to initiate the transmission of its locally trained model parameters. Following the reception of data from Client 2 (in $T_{c2,s}$ seconds), the server, once again, acknowledges the successful reception through the Zigbee network. Subsequently, with data from both clients in hand and having received confirmation responses from both clients via the Zigbee network, the server calculates the necessary aggregate of local model parameters to update the global model. Subsequently, the server broadcasts a signal to both clients, instructing them to activate their receptors. The synchronization process concludes one communication round when both clients independently confirm to the server through the Zigbee network that they have received the updated information, at which point the server halts its broadcasting. The order of activation of the SDRs has been depicted in Fig. 2 (left), and the overall time taken for such a TDMA-based scheme is shown in Fig. 3. Note that $T_{c1,c1}$, $T_{c2,c2}$ and $T_{s,s}$ refer to computation time at client1, client 2 and the server, respectively.

D. Client-Server Communication

In our demonstration, USRP is managed through LabVIEW, the XBee devices are controlled using XCTU, and the ML
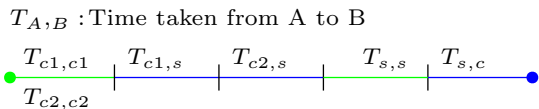
$T_{A,B}$: Time taken from A to B

Fig. 3: Temporal aspects of one communication round.

model parameters are generated via Python. These software components seamlessly work in sync, creating an integrated system for the FL experiment. After obtaining the model updates from the local ML models, it is converted into binary format. This binary data is subsequently organized into packets marked by a unique identifier, effectively serving as its 'jacket number.' These packets consist of 128 bits and are further enhanced by the inclusion of guard bits, sync bits (PN Sequence order = 31), as well as blank frames. A 60-bit cyclic redundancy is added to each packet to bolster recognition at the receiver's end. Once these packets are prepared, the final data packet undergoes differential modulation, transforming it into a complex signal (I+Q). This complex signal then advances through the up-conversion stage, ultimately culminating in the process of final transmission. After down-converting the sampled signal over an acquisition time of 40 milliseconds, we eliminate any DC offset and unwanted frequencies by applying the Hamming window. Subsequently, we employ threshold voltage detection to transform the analog signal into a digital bit stream through demodulation. The synchronized detection is facilitated through convolution using the same PN sequence order. Once the synchronization process is complete, we extract the message bits from the data packets. These bits are then mapped to their corresponding packet numbers embedded into the generated bit stream.

IV. NUMERICAL RESULTS

In this section, we present the outcomes of our demonstration. In Fig. 5 (left), we compare IID and non-IID scenarios comprehensively. While we see that in the IID scenarios, an accuracy of 99% is obtained, it drops to 90% under severe heterogeneity, as expected. From Fig. 5 (right), we see that transmission power impacts the stability, accuracy, and rate of convergence in the non-IID setting. However, we also observe improved convergence with higher transmission power.
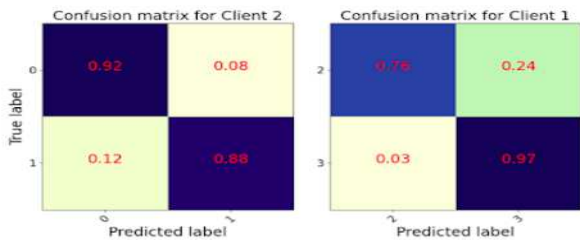

Fig. 4: Confusion Matrix for Cross-accuracy.

In Table. I, we report the best accuracy as a function of the epochs, which measures the amount of local learning and the transmission power. We observe that the convergence is fast in the case of IID data, and higher local learning helps.
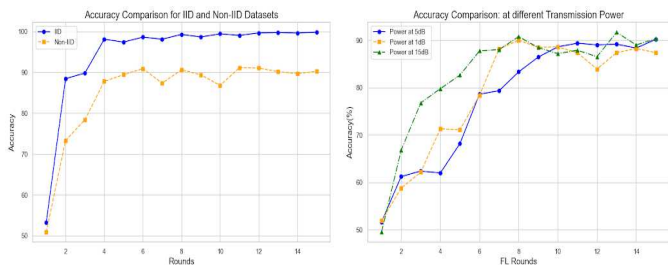

Fig. 5: (Left)Accuracy performance in the IID and non-IID setting & (right) for different transmission powers.

Albeit reduced accuracy gains in the non-IID context, more local learning helps here as well. However, the time taken to achieve convergence is much larger as it takes several more communication rounds as compared to the IID setting.

In order to ensure that we have a generalized global model, we test the cross-accuracy, i.e., since Client 1 has classes 0 and 1, we test if Client 1 can predict well on classes 2 and 3, and analogously, we test Client 2 on classes 0 and 1. In Fig. 4, the confusion matrix shows that the cross-accuracy reported by Client 1 is 87% and Client 2 is 90%, which ascertains that our global ML model is well-generalised.

TABLE I: Best accuracy values for varying epochs and transmission power along with the time taken.

IID	# Epochs		Power(dB)	
	5	10	10	15
Accuracy(%)	99	99	99	99
# Tx. Rounds	5	3	3	3
Time taken(s)	215.43	180.89	180.89	178.71
Non-IID	# Epochs		Power(dB)	
	5	10	10	15
Accuracy(%)	91	90	90	92
# Tx. Rounds	11	8	8	8
Time taken(s)	539.55	507.36	507.36	494.63

V. CONCLUSIONS

In this work, we present a hardware prototype of FL system that utilizes a wireless backbone. The results of our research exemplify the power of FL in establishing a collaborative and privacy-preserving learning environment over wireless channels where multiple clients collectively contribute to a unified, general model. Establishing a network of SDRs and USRPs to enable private ML, while achieving performance similar to centralized architectures is the key takeaway of this prototype. Our approach offers a basic blueprint for implementing FL, which can be extended to delay-specific wireless networks, network with stragglers, etc.

REFERENCES

[1] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Aguera y Arcas. "Communication-efficient learning of deep networks from decentralized data." In AISTATS 2017, pages 1273–12
[2] T. Sery, N. Shlezinger, K. Cohen and Y. C. Eldar, "COTAF: Convergent Over-the-Air Federated Learning," GLOBECOM 2020, pp. 1-6.
[3] Amiri, M., and D. Gündüz. "Federated learning over wireless fading channels." IEEE Trans. on Wireless Communications 19, no. 5 (2020).
[4] Hussain, A., N. Gundapu, S. Drugkar, S. Kiran, J. Harshan, and R. Prasad. "Seeing is Believing: A Federated Learning Based Prototype to Detect Wireless Injection Attacks." arXiv preprint:2311.06564 (2023).