

Demo Abstract: A Hardware Prototype Targeting Federated Learning with User Mobility and Device Heterogeneity

Allen-Jasmin Farcas allen.farcas@utexas.edu The University of Texas at Austin Austin, TX, USA

ABSTRACT

This paper presents a new hardware prototype to explore how centralized and hierarchical federated learning systems are impacted by real-world devices distribution, availability, and heterogeneity. Our results show considerable learning performance degradation and wasted energy during training when users mobility is accounted for. Hence, we provide a prototype that can be used as a design exploration tool to better design, calibrate and evaluate FL systems for real-world deployment.

CCS CONCEPTS

• Computing methodologies → Machine learning; Supervised learning; • Security and privacy;

KEYWORDS

Federated Learning, Edge Devices, Data Heterogeneity, Data Privacy, Communication Cost, Energy Efficiency, Internet-of-Things

ACM Reference Format:

Allen-Jasmin Farcas and Radu Marculescu. 2023. Demo Abstract: A Hardware Prototype Targeting Federated Learning with User Mobility and Device Heterogeneity . In International Conference on Internet-of-Things Design and Implementation (IoTDI '23), May 09–12, 2023, San Antonio, TX, USA. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3576842.3589160

1 INTRODUCTION

Federated Learning (FL) is the de facto solution for large scale deployment of Edge AI applications since it enables distributed learning with data privacy considerations [2]. In FL, edge devices first download a model from the Cloud and then train it using their local data. Finally, all edge devices send their updated models to the Cloud for global aggregation.

Real-world devices have their availability and distribution dictated by the real-world user mobility, since most devices of interest for FL are smartphones or wearables. However, the user mobility is typically ignored in the FL literature as all devices are assumed to be always available and uniformly distributed across the area of deployment. Most edge devices have limited battery life, variable connectivity and bandwidth that can significantly impact the performance of real-world FL systems for Edge AI. Hence, all FL



This work is licensed under a Creative Commons Attribution International 4.0 License.

IoTDI '23, May 09–12, 2023, San Antonio, TX, USA © 2023 Copyright held by the owner/author(s). ACM ISBN 979-8-4007-0037-8/23/05. https://doi.org/10.1145/3576842.3589160 Radu Marculescu radum@utexas.edu The University of Texas at Austin Austin, TX, USA



Figure 1: Hardware prototype containing 36 Raspberry Pi 3B+ (Top) and 36 Odroid MC1 (Bottom) devices, organized in two setups. The inner semicircle in both setups contains Smart Power 2 devices that measure the power consumption for each device while they perform on-device training in a unified FL prototype.

algorithms and solutions targeting Edge AI systems should properly account for the impact of user mobility and device heterogeneity.

To this end, we believe that good analytical models should go hand in hand with *realistic simulations*. By creating a hardware prototype and a software framework, we can properly analyze the impact of mobility and device heterogeneity for both centralized and hierarchical FL systems.

2 APPROACH

Hardware Prototype Considerations The heterogeneous hardware prototype, shown in Fig. 1, contains 36 Raspberry Pi 3B+ and 36 Odroid MC1 devices. We leave as spare devices one Raspberry Pi and one Odroid, hence we consider for our experiments a system with 70 heterogeneous devices. We use TP-Link AX6000 WiFi 6 Router as the main communication point between all devices and the Cloud, with ethernet cable connections for the Odroid devices and Wi-Fi connections for the Raspberry Pi devices. We use a desktop server with a 64-core AMD Threadripper PRO 3995WX CPU, four A6000 GPUs and 512GB RAM as the Cloud in our experiments. **Software Framework for Realistic Federated Learning** Since other centralized or hierarchical FL frameworks [1, 4, 5] cannot



Figure 2: Average measurements per device for Odroid MC1 and Raspberry Pi 3B+ for communication times (a), power (b), temperature (c) and training times (d).

handle a 70 device hardware prototype and user mobility, we build our own FL framework from scratch. We enable our framework to run both centralized FL (CFL) and hierarchical FL (HFL) scenarios with or without mobility of the devices considered. Our solution can scale to N users by extending the algorithm to run sequentially N/70 batches of users directly on our hardware prototype. Additionally, we will open-source our code, thus providing a useful research tool for illustrating the real-world implications in FL.

User Mobility Considerations For user mobility we use the Foursquare dataset [3] for the first 500 hours of May 2020 over a 17.5km×17.5km urban area for deployment. In our scenarios, the edge devices are smartphones associated with people that move around walking or driving. We consider the top 70 devices that appear most often out of 12,866 devices available in the Foursquare dataset and select for HFL 50 Access Points (APs) at random from a total of 37,994 APs. We consider that a device takes less time to communicate compared to the time required for training. Hence, the devices that are present in communication round *i* will start training and will be aggregated the next communication round i + 1, but only if they are still available.

EXPERIMENTAL SETUP AND RESULTS 3

Due to lack of space, we show the real-world system-wide impact of mobility and hardware heterogeneity of current CFL and HFL solutions on MNIST IID. However, our framework can run on multiple datasets (e.g., CIFAR10/100, EMNIST) with both IID and non-IID configurations. We use 500 images per device for MNIST dataset with one local epoch for CFL and one local epoch with five edge aggregations for HFL.

User Mobility Impact on Hardware In existing FL solutions, if devices start training at communication round *i* and are not available at the next communication round i + 1, then they will not be aggregated. Hence, devices may waste energy to train a model that isn't even considered for aggregation. In an ideal scenario, CFL and HFL would not waste any energy since all devices are available all the time and no devices are missing when aggregations are performed. However, as shown in Table 1, when considering realistic user mobility, up to 65% and up to 93% of the total energy consumed is wasted for CFL and HFL, respectively. This shows how

Table 1: Average energy consumption for the	hardware pro-
totype with all 70 devices using Conv5 model	on IID MNIST.

Avg. energy per	MNIST			
device [Joules]	CFL	CFL Mobility	HFL	HFL Mobility
Wasted	0 J	246 J	0 J	469 J
Consumed	399 J	380 J	422 J	501 J

much energy can be wasted when naively deploying current FL solutions in the real world. Hardware heterogeneity also creates a large variation in the average energy consumed per device in all scenarios (i.e., on average 250 J), since the Odroid devices consume more power and take longer to train compared to the Raspberry Pi devices (see Fig. 2(b)).

Another impact of user mobility is shown in Fig. 2 (a), where the average communication times for CFL and HFL vary. It is expected that Odroid MC1 will have lower communication times since they are connected with Ethernet cables to the router, while the Raspberry Pi devices are connected using Wi-Fi. The router becomes a bottleneck when using full device availability in CFL and HFL, hence the larger communication times.

In Fig. 2 (b) and (c), we observe more power and greater temperatures for the MC1 devices, on average, compared to the Raspberry Pi devices. Despite this, in Fig. 2 (d) we see a large training time difference between the two devices, showing how heterogeneous the system is. Using this hardware prototype as a testbed, new FL solutions can be properly evaluated for real-world deployment.

CONCLUSION 4

Ta

In this paper, we have presented a hardware prototype and a software framework to analyze the real-world impact of FL solutions when considering user mobility and availability with device heterogeneity. The framework is flexible enough to include any mobility dataset and can run experiments on GPUs, on real devices or a combination of both using any neural networks on both IID and non-IID datasets.

Our evaluation shows significant increases in wasted energy per device due to the real-world mobility and availability of the devices. Our demonstration consists of a live showcasing of our framework both in GPU simulation and on-device deployment with and without mobility considerations, with real-time power and latency measurements.

ACKNOWLEDGMENTS

This research was supported in part by NSF Grant CCF-2107085 and in part by Cisco Research, Inc.

REFERENCES

- [1] D. J. Beutel et al. 2020. Flower: A friendly federated learning research framework. arXiv preprint arXiv:2007.14390 (2020).
- [2] K. Bonawitz et al. 2019. Towards federated learning at scale: System design. Proceedings of machine learning and systems 1 (2019), 374-388.
- [3] Foursquare. 2023. Independent Location Data & Location Technology Platform. Accessed: 2023-03-03.
- [4] F. Lai et al. 2022. Fedscale: Benchmarking model and system performance of federated learning at scale. In ICML. PMLR, 11814-11827.
- [5] A. Ziller et al. 2021. Pysyft: A library for easy federated learning. Federated Learning Systems: Towards Next-Generation AI (2021), 111-139.