Energy Efficient Federated Learning with Bayesian Optimized Training Pace Control

Hongpeng Guo[1], Haotian Gu[2], Zhe Yang[1], Xiaoyang Wang[1], Eun Kyung Lee[3], Nandhini Chandramoorthy[3], Tamar Eilam[3], Deming Chen[1], Klara Nahrstedt[1]

University of Illinois Urbana-Champaign [1] Berkeley University of California [2] IBM Research [3]

INTRODUCTION

Federated Learning Overview:

- Federated Learning (FL) is a **distributed machine learning** technique where models are trained on local data to **preserve privacy**.
- Federated learning has significant applications in **healthcare**, **material discovery**, **autonomous vehicles** and **smart homes**.
- However, training models on edge devices can consume a lot of energy, making **energy-efficient FL** crucial.

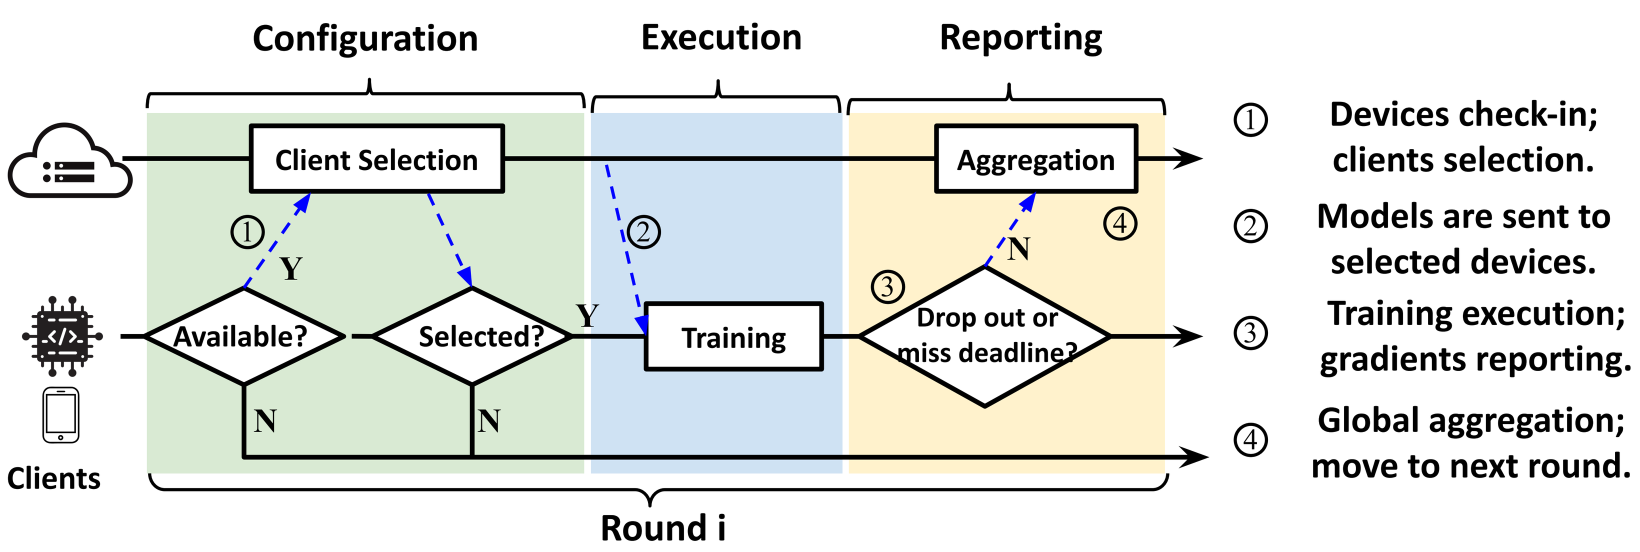

A typical workflow of federated learning.

Dynamic Voltage and Frequency Scaling (DVFS):

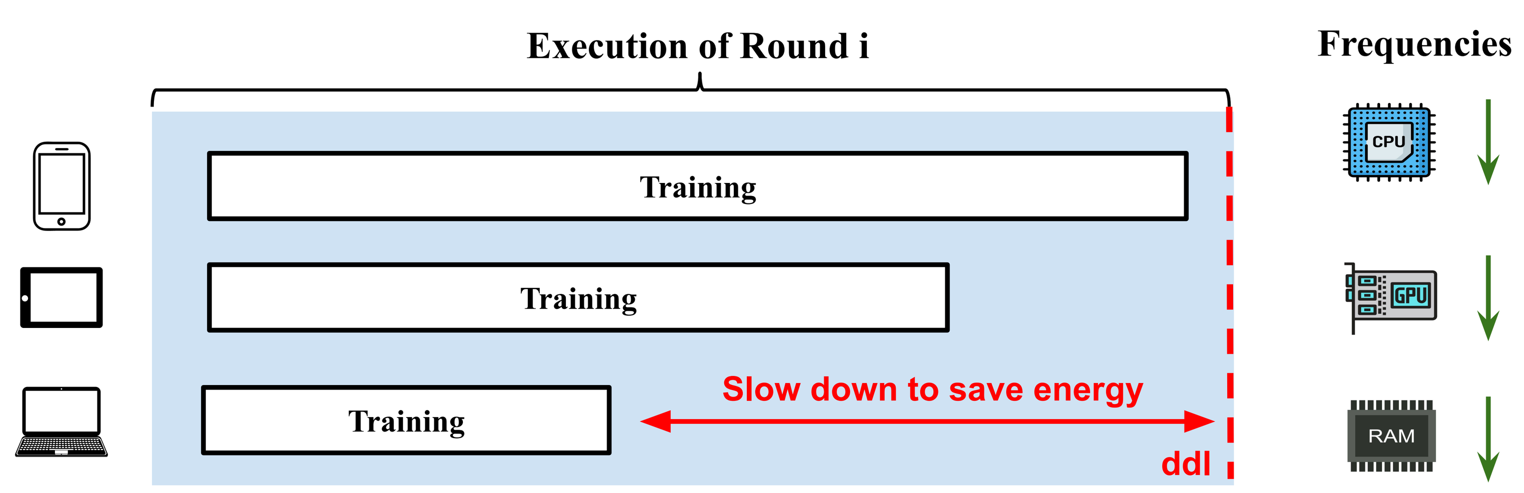

- **DVFS** is a technique that **optimizes energy and performance** according to demand, scaling resources as needed.

In This Work:

- We design a **training pace controller**, named **BoFL**, for Federated Learning clients, which operates by **intelligently modifying hardware configurations in real-time** to optimize energy consumption and performance during the training of machine learning models.

CHALLENGES

How to Select the Best DVFS Configurations?

How to select the **best** DVFS configurations for each round of local model training?

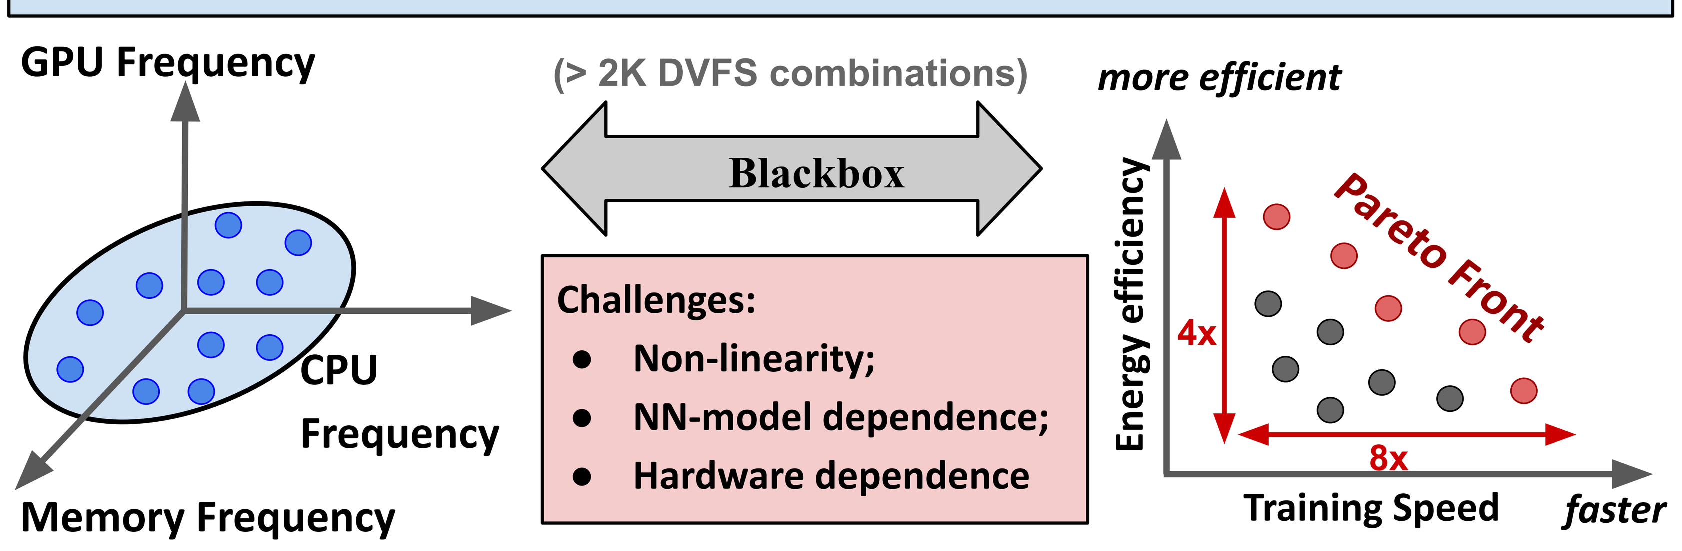

METHODOLOGY

BoFL System Overview:

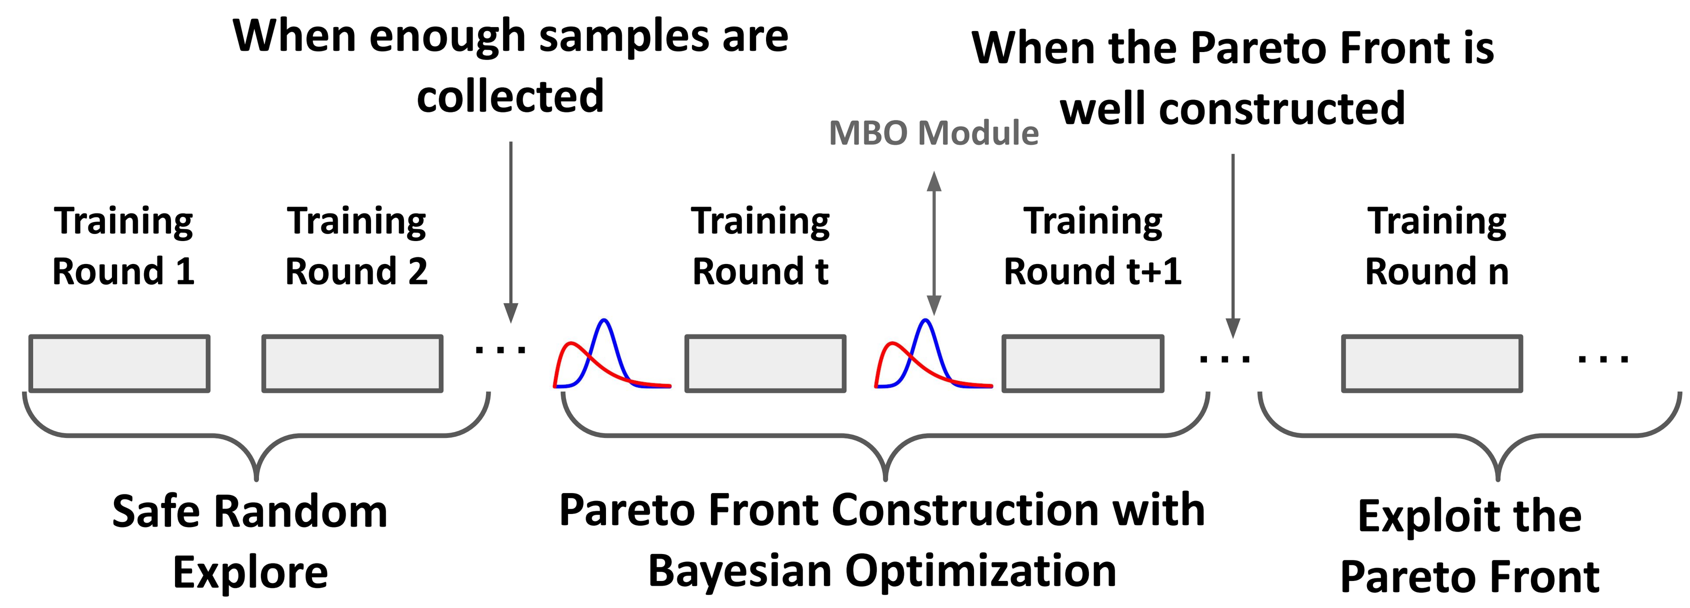

Pareto Construction with Bayesian Optimization:

Bayesian optimization (BO) is a **sample efficient** methodology for optimizing black-box functions that are expensive to evaluate.

BO iteratively selects input values to **balance exploration and exploitation,** finding the global optimum with minimal evaluations.

Multi-Objective Bayesian optimization (MBO)

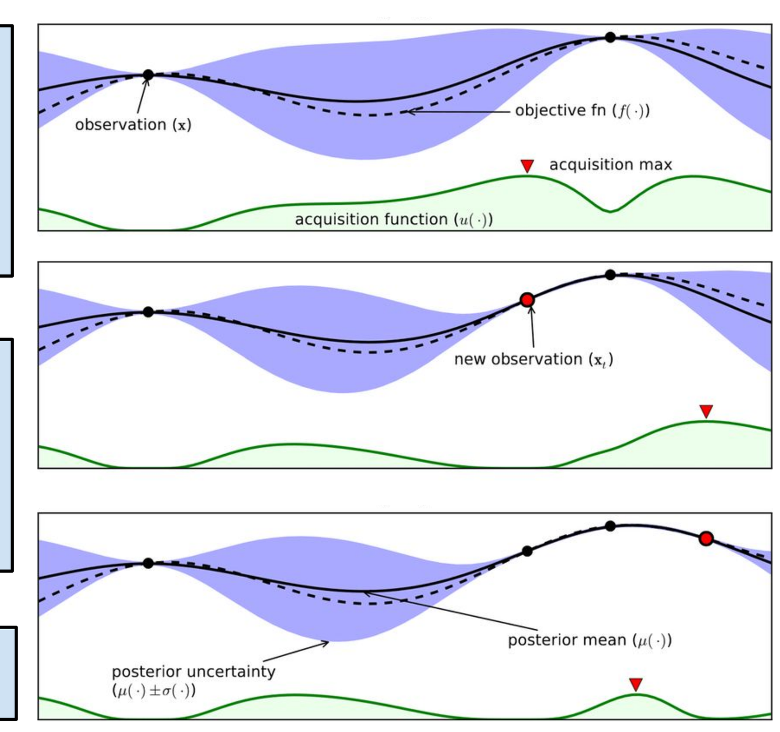

EXPERIMENTAL RESULTS

Testbed Hardware:

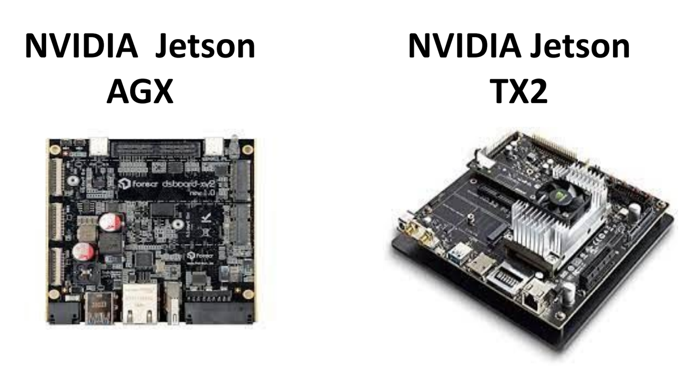

Learning Tasks:

Task Types	Datasets	NN-Model
CV	CIFAR10	Vision-Transformer
CV	ImageNet	ResNet-50
NLP	IMDB	LSTM

Results and Analysis:

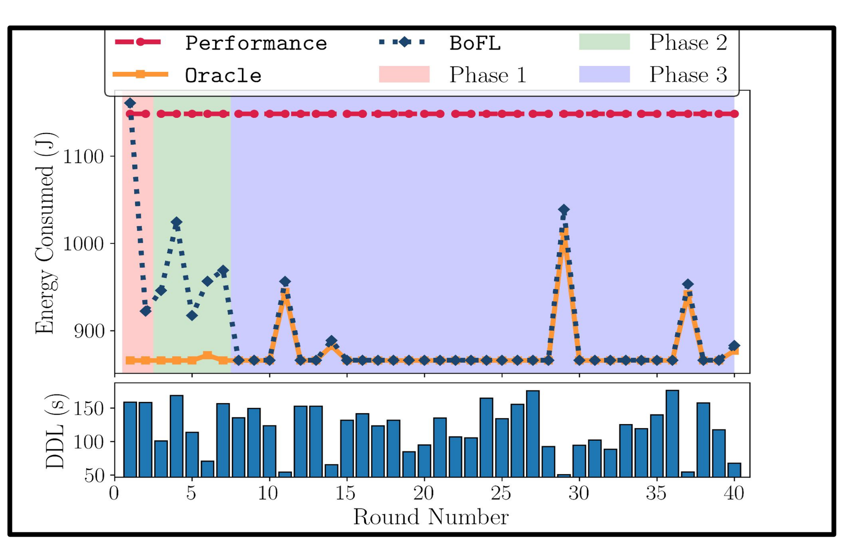

Energy consumption over FL rounds.

- BoFL cuts energy use **22.3%** vs. *Performant*,
- Only **1.48%** energy overhead compared to *Oracle*.

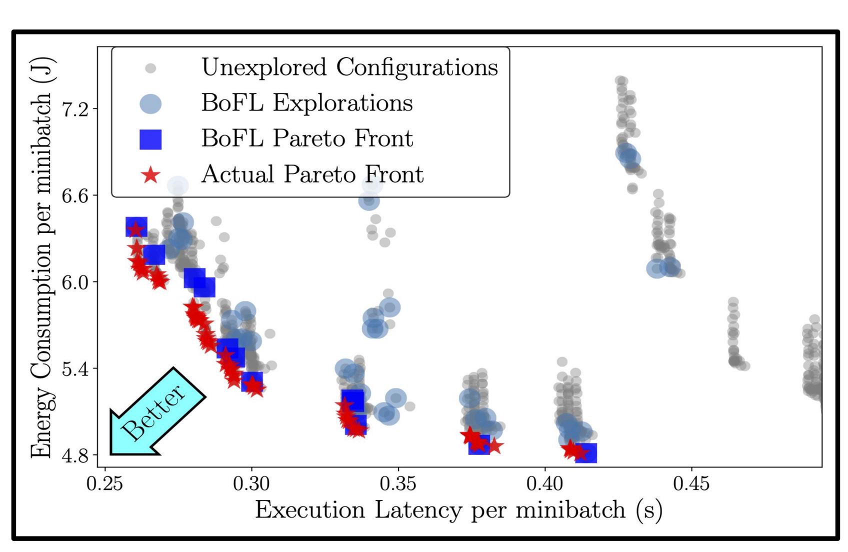

Pareto front searched by BoFL .

- BoFL can successfully find a close approximation to the actual Pareto front over **all three tasks**.

CONCLUSION

BoFL is a **training pace controller** for edge devices that achieves energy efficient federated learning. Experiments show that it can reduce energy consumption by over **20%** compared to the baselines and achieve close to optimal energy efficiency with only **1.2% - 3.4% energy regret**.

ACKNOWLEDGEMENT

This work was supported by **IBM-Illinois Discovery Accelerator Institute**.

More information can be found in our paper "**BoFL: bayesian optimized local training pace control for energy efficient federated learning**"
In the 23rd ACM/IFIP International Middleware Conference. 2022

Contact: hg5@illinois.edu

NAE Regional Meeting 2023