Harnessing Increased Client Participation with Cohort-Parallel Federated Learning

Akash Dhasade[1], Anne-Marie Kermarrec[1], Tuan-Anh Nguyen[2], Rafael Pires[1], Martijn de Vos[1]
[1]EPFL, Switzerland [2]Independent Researcher, France

Motivation

Standard FL focuses on increasing the *cohort* size to induce more parallelism

However, this yields diminishing returns [1]: doesn't satisfactorily harness large client participation

How can we better harness increased client participation in Federated Learning?

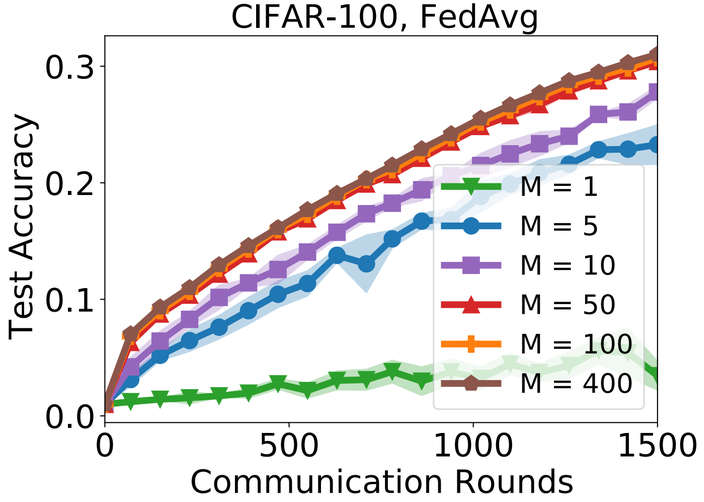

[1] Charles, Zachary, et al. "On large-cohort training for federated learning." In NeurIPS (2021).

CPFL: Cohort-Parallel Federated Learning

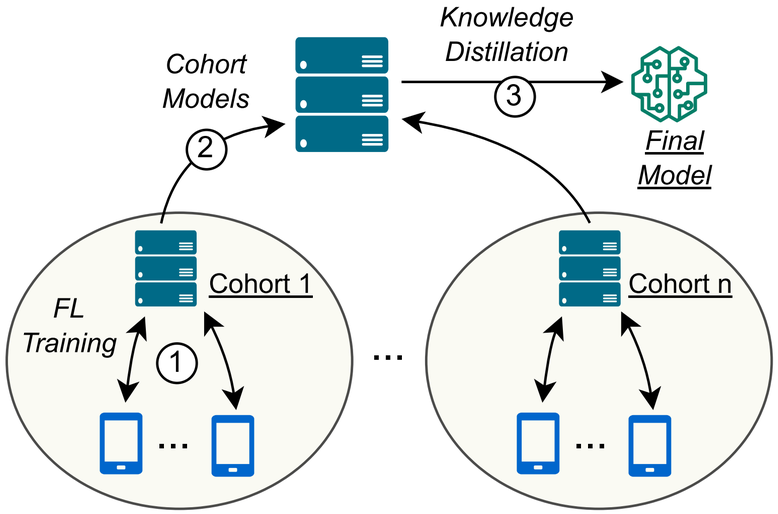

Step 1: Partition the network into **n** random partitions (aka *cohorts*) and perform FL within each cohort

Step 2: Upload trained cohort FL model to the global server

Step 3: Distill cohort FL models into a single global model

Key Insight: Smaller networks better harness individual client updates as well as converge significantly quicker

CPFL significantly lowers

(i) training time
(ii) training resource usage
(iii) communication costs

at a minimal loss in accuracy

CPFL provides flexible means to control training efficiency and scale FL