First Cross-Facility Federated Learning Deployment on DOE Leadership-Class Supercomputers

Argonne NATIONAL LABORATORY

Scientific Achievement

- **First-ever federated learning (FL) deployment across four DOE leadership-class supercomputers**: ALCF-Aurora, ALCF-Polaris, OLCF-Frontier, and NERSC-Perlmutter.
- Developed HPC-aware asynchronous FL algorithms robust to heterogeneous hardware (Intel/AMD/NVIDIA), independent job queue delays, and firewalls.
- Demonstrated secure, cross-silo model training, using the APPFL framework and Globus Compute by using large science LLM.

Significance and Impact

- Opens the door for privacy-preserving scientific AI across DOE facilities.
- Provides a scalable path toward cross-facility foundation model training.
- Demonstrates a new capability for DOE: multi-supercomputer AI workflows running in parallel, coordinated through Globus.
- Advances DOE's AI-for-Science missions by enabling secure, portable, facility-agnostic AI training.

Technical Approach

- Utilized APPFL to coordinate training across facilities.
- Integrated Globus Compute + ProxyStore for model transfer across strict HPC network boundaries.
- Conducted experiments with LLM fine-tuning across all four systems simultaneously; evaluated throughput, communication costs, and queue-time effects

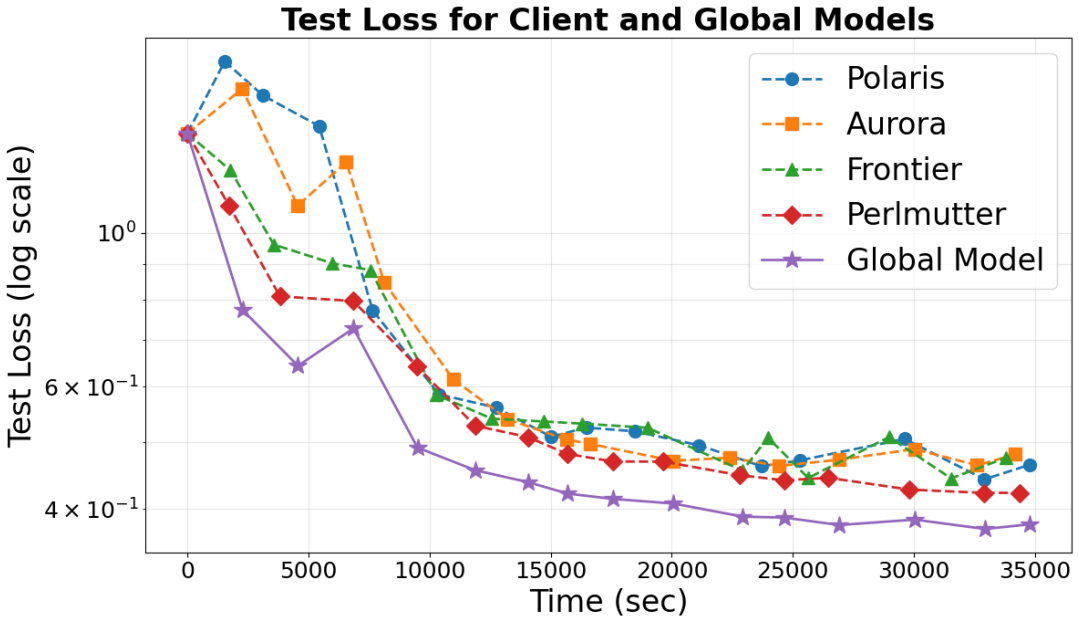

Figure. *Test-loss trajectories from federated training running simultaneously on Aurora, Polaris, Frontier, and Perlmutter. Despite heterogeneous hardware and asynchronous job schedules, all client models converge steadily, while the aggregated global model achieves the fastest and lowest loss—demonstrating robust, cross-facility FL coordination using APPFL and Globus Compute.*

PI(s)/Facility Lead(s): Kibaek Kim
Collaborating Institutions: ALCF, OLCF, NERSC, University of Chicago

ASCR Program: AI for Science
ASCR PM: Hal Finkel