

FedQueue: Queue-Aware Federated Learning for Cross-Facility HPC Training

Scientific Achievement

- Developed **FedQueue**, the first queue-aware federated learning (FL) protocol that predicts and controls HPC batch scheduler delays during cross-facility FL.
- Proved convergence for non-convex objectives at rate $O(1/\sqrt{R})$ under bounded staleness, matching baseline algorithms.
- Deployed across four DOE production HPC facilities (ALCF-Aurora, ALCF-Polaris, OLCF-Frontier, and NERSC-Perlmutter) fine-tuning LLaMA2-7B, achieving **20.5% lower test loss** than baseline algorithms.

Significance and Impact

- Addresses a critical bottleneck in multi-facility FL training: unpredictable scheduler queues significantly degrade existing FL algorithm performances.
- Achieves **34% faster convergence** and up to **67% less model communication** under high queue variance versus baseline algorithms.
- Enables practical, production-ready FL across heterogeneous supercomputers for large-scale AI for Science workflows.

Technical Approach

- **Queue prediction:** EWMA-based online estimator dynamically budgets compute steps per facility per round.
- **Admission control:** Deadline-based cutoffs; late updates are buffered, not discarded.
- **Staleness-aware aggregation:** Harmonic decay weighting and inverse learning rate scaling stabilize training across heterogeneous clients.
- Implemented with APPFL and Globus software suite.




Figure. Test loss vs. wall-clock time across all algorithms during real-world cross-facility deployment. After an initial warm-up, both FedQueue configurations achieve the lowest final loss, with FedQueue-2 being the only algorithm to reach loss below 0.4 within the 17,000-second budget.

PI(s)/Facility Lead(s): Kibaek Kim

Collaborating Institutions: ALCF, OLCF, NERSC, University of Chicago

ASCR Program: AI for Science

ASCR PM: Xujing Davis

Publication for this work: Y. Li, E. Dey, Z. Li, K. Raghavan, R. Madduri, K. Kim. "FedQueue: Queue-Aware Federated Learning for Cross-Facility HPC Training", to appear in ICML 2026