

June 2, 2026

TOWARDS SCALABLE FEDERATED LEARNING FOR SCIENTIFIC COMPUTING

YIJIANG LI

with Emon Dey, Zilinghan Li, Krishnan Raghavan, Kyle Chard, Ian Foster, Todd Munson,
Ravi Madduri, and Kibaek Kim
Argonne National Laboratory



Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.



Data and Compute Problem in Scientific AI

- **Data problem**


- Cannot be centralized: privacy and sheer volume

- **Compute problem**

- Large models are compute-hungry

- **Solution**

- Bring computation to the data: FL on HPC



Challenges of FL on HPC

▪ Waiting in queue



- Highly variable batch queue delay
- Choices between sync and async algorithm

▪ Local training

- Heterogeneous hardware
- Distributed compute configurations
- Local training step budget

▪ Sharing updates

- Multi-GB model transfer per round
- Transfer-speed disparity across sites



Built on APPFL: Cross-HPC-Facility

FL Core (APPFL APPFL)

FL coordination — aggregation, scheduling, privacy

Globus Compute


Remote task dispatch under site-specific security policies

Globus Transfer + ProxyStore

High-throughput model parameter transfer

DeepSpeed + MPI

Distributed local training; ZeRO optimization for large models




Y. Li, Z. Li, K. Chard, I. Foster, T. Munson, R. Madduri, and K. Kim, "Scalable Cross-Facility Federated Learning for Scientific Foundation Models on Multiple Supercomputers." <https://arxiv.org/pdf/2603.19544>

Cross-HPC-Facility in Action

■ Setup

- Fine-tuning of LLaMA-2 7B on SMollnstruct
- Natural non-IID distribution by chemistry subdomain across
 - ALCF-Aurora¹
 - ALCF-Polaris
 - OLCF-Frontier¹
 - NERSC-Perlmutter
- 64 nodes per facility, co-scheduled with reservations




1,700+
GPUs used

3
GPU vendors

1.39→0.37
Global loss drop

1. Aurora and Frontier are both among the world's top-3 supercomputers (TOP500, November 2025)

Built on APPFL: Scheduler-Aware FedQueue




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

FedQueue in Action

■ Setup

- Fine-tuning
 - LLaMA-2 7B on SMollInstruct
 - Same non-IID distribution
 - Local training on two nodes
- Controlled simulation
 - CNN on non-IID MNIST
 - Sweep queue time variance



20%
Improved loss

60%
Faster to target acc.

2x
Fewer local steps


FedQueue-1 and FedQueue-2 are FedQueue with different synchronization horizon.



U.S. DEPARTMENT
of ENERGY

Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

The Same Platform, New Science




Expanding Across DOE and Beyond

■ Broader DOE ecosystem

- Joint demo by Argonne and Sandia
- Expand APPFL deployment to other DOE national labs and commercial clouds

■ Genesis Mission

- Expanding FL capability to train more diverse models and datasets



Challenges and Opportunities

- **Efficiency at scale**

- Computational efficiency for larger and more complex models
- Communication cost and model performance trade-off

- **Dynamic collaboration**

- Real-world FL dynamics with verification, and certifications of model updates
- Alternative domain-specific FL protocols

- **Privacy & Robustness**

- Privacy and robustness definition in real-world FL

Acknowledgments

- This work was supported by the U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research, under Contract DE-AC02-06CH11357.
- An award of computer time was provided by the ASCR Leadership Computing Challenge (ALCC) program. This research used resources of the Argonne Leadership Computing Facility, which is a U.S. Department of Energy Office of Science User Facility operated under contract DE-AC02-06CH11357. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a Department of Energy User Facility using NERSC award ALCC-ERCAP0038201.
- We gratefully acknowledge the computing resources provided on Improv, a high-performance computing cluster operated by the Laboratory Computing Resource Center at Argonne National Laboratory.

Argonne
NATIONAL LABORATORY



U.S. DEPARTMENT
of ENERGY