Workshop on Research Directions in Communication Software for Future Generation DOE Systems and Applications

Overview

The rapid evolution of large-scale computing systems is reshaping the requirements for communication software that underpins scientific computing, data-intensive workloads, and artificial intelligence (AI). Emerging post-exascale architectures are expected to incorporate increasing levels of heterogeneity, deeper memory hierarchies, specialized accelerators, advanced interconnects, and tighter integration of simulation, data analytics, and machine learning workflows. At the same time, AI and foundation-model workloads are becoming first-class drivers of system design, introducing new communication patterns, scalability challenges, and performance requirements.

To address these trends, we are organizing a community workshop at Argonne National Laboratory that will bring together stakeholders from DOE national laboratories, academia, and industry to identify and prioritize future research directions in communication software for DOE systems and applications. Stakeholders include researchers, users, application developers, computing facilities, and vendors.

The workshop will focus on the software ecosystem that enables communication and data movement at scale, including message-passing libraries, partitioned global address space (PGAS) models, collective communication libraries, runtime systems, and communication frameworks supporting both traditional high-performance computing (HPC) and emerging AI workloads. Examples include MPI, OpenSHMEM and related PGAS technologies, communication runtimes, and AI-oriented collective communication libraries (CCLs).

Goals

The primary goals of the workshop are to:

- Assess emerging communication requirements for next-generation HPC, AI, agentic workflows, and converged HPC-AI applications.
- Identify limitations of current communication programming models, runtimes, and libraries in the context of future DOE systems.
- Understand the implications of evolving architectures, including heterogeneous nodes, accelerators, disaggregated memory, advanced networking technologies, and post-exascale system designs.
- Explore opportunities for greater convergence between communication software developed for scientific computing and AI training/inference workloads.
- Define key research challenges in scalability, programmability, performance portability, resilience, security, resource management, elasticity, and energy efficiency.
- Foster collaboration among application developers, communication-library researchers, system software developers, hardware vendors, and facility operators.
- Develop a community-informed research agenda that can guide future investments and collaborations across the DOE ecosystem.

Workshop Format and Key Discussion Topics

The workshop format combines short presentations, invited talks, panels, and breakout sessions to foster interactive discussion. The workshop will consider questions such as:

- What communication abstractions and programming models will be required for future large-scale scientific and AI applications?
- How should communication software evolve to support increasingly heterogeneous and dynamic systems?
- What lessons can be learned from communication libraries developed for large-scale AI training and inference?
- How can communication runtimes better exploit emerging network and accelerator capabilities?
- What new approaches are needed to address resilience, fault tolerance, elasticity, performance variability, and energy constraints at extreme scale?
- How can communication software enable efficient integration of simulation, data analytics, and AI workflows?

Expected Outcomes

The workshop will produce a report documenting the shared understanding of the most pressing research opportunities and challenges facing communication software over the next decade. Discussions will inform a set of community recommendations and research priorities that can help shape future DOE initiatives, foster cross-sector collaboration, and guide the development of communication software for post-exascale systems and next-generation scientific and AI applications.

Organizing Committee

- Rajeev Thakur (Argonne National Laboratory, Chair)
- Meifeng Lin (Brookhaven National Laboratory)
- Damian Rouson (Lawrence Berkeley National Laboratory)
- Ignacio Laguna (Lawrence Livermore National Laboratory)
- John Patchett (Los Alamos National Laboratory)
- David Bernholdt (Oak Ridge National Laboratory)
- Draguna Vrabie (Pacific Northwest National Laboratory)
- Ron Brightwell (Sandia National Laboratories)

Advisors

- Valerie Taylor (Argonne National Laboratory)
- Michael Parks (Oak Ridge National Laboratory)