Summary of Center Accomplishments: 2020 - 2025

Presenter: Prof. Patrick Bridges

CUP-ECS Provided DOE a Full Suite of Communication Abstractions, Models, and Tools

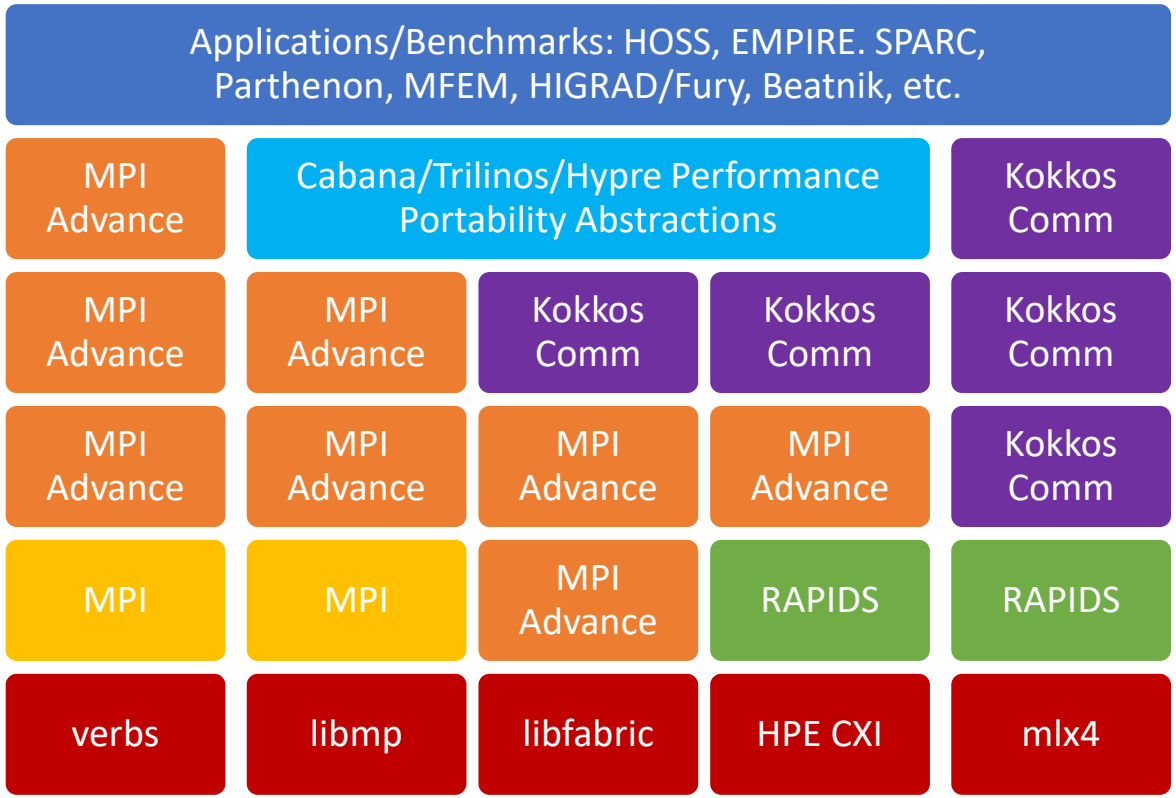

- Researched High-level Application-facing Communication Abstractions and Low-level Primitives necessary to support them
- Development at each level driven by careful assessment, benchmarking, and modeling workflows and tools
- Carefully integrating these abstractions with a wide range of DOE applications and libraries, including multiple production NNSA codes
- Demonstrated performance improvements in full production applications

Trained 20 Students/Postdocs
5 Lab Staff Placements

Current Students

1. Adams, Michael; UNM Ph.D.
2. Avans, Nicole; TN Tech Ph.D.
3. Bacon, Nicholas; UNM Ph.D.
4. Collom, Gerald; UNM Ph.D.
5. Nansamba, Grace; TN Tech Ph.D.
6. Namugwanya, Evelyn; TN Tech Ph.D.
7. Raad, Abdalaziz; UNM B.S.
8. Shipley, Riley; TN Tech Ph.D.
9. Wesley, Jackson; UNM Ph.D.

Lab Placements

1. Dominguez-Trujillo, Jered; UNM M.S. (LANL staff)
2. Gorham, Thomas; UTC M.S. (LLNL Staff)
3. Haskins, Keira; UNM M.S. (Sandia staff)
4. Hooten, Garrett; UTC B.S. (LLNL Staff)
5. Marshall, Ryan; UA Postdoc (LANL Staff)

Other Graduates

1. Broaddus, Justin; UTC M.S. (Industry)
2. McCullough, Jacob; UNM M.S. (GAMA-1)
3. Gekyo, Andrew; UNM B.S. (MPI Ph.D. student)
4. Marts, Pepper; UNM Ph.D. (HPE Staff member)
5. Stewart, Jason; UNM M.S. (UNM Staff)
6. Suggs, Evan; UTC M.S. (TN Tech Staff)
7. Woods, Carson; UTC B.S. (Emory Ph.D. student)

11 Lab Interns Placed

- **Adams, Michael:** Sandia National Laboratories, Daniel Gomez, mentor.
- **Avans, Nicole**: Sandia National Laboratories, Jan Ciesko, mentor.
- **Bacon, Nicholas**: Sandia National Laboratories, Scott Levy, mentor.
- **Broaddus, Justin**: Los Alamos National Laboratory, Galen Shipman, mentor.
- **Collom, Gerald**: Lawrence Livermore National Laboratory, David Boehme and Ruipeng Li, mentors.
- **Dominguez-Trujillo, Jered:** Sandia National Laboratories, Ryan Grant, mentor.
- **Haskins, Keira:** Sandia National Laboratories, Kurt Ferreira, mentor.
- **Marshall, Ryan**: Los Alamos National Laboratory, Jon Reisner, mentor.
- **Nansamba, Grace**: Lawrence Livermore National Laboratory, Olga Pearce, mentor.
- **Namugwanya, Evelyn**: Lawrence Livermore National Laboratory, David Boehme, mentor.
- **Wesley, Jackson:** Sandia National Laboratories, Whit Schonbein, mentor.

Software Products Released/In Progress

- Beatnik Fluid Interface Benchmark: Release 1.1.0 10/31/2024
- Cabana Advection Benchmark: Release 1.0 6/12/2022
- Cabana Communication Abstraction Enhancements: PR Submitted
- ExaMPI: Release with Lassen and Tioga Support Planned
- MPI Advance Communication Abstractions
  - Partitioned Communication Module: Release 1.3.0 3/18/2024
  - Locality-aware Collective Module: Release 1.0 9/19/2022
- RAPIDS: Channel API Release Planned

DOE Software Studied/Enhanced

- Production Application and Libraries
  - Caliper: Pattern annotation API integrated with production code
  - HIGRAD: Initial GPU support prototyped
  - HOSS: Reengineered communication engine, resulting in 25%-40% improved strong scaling. Changes merged with production code
  - Hypre: Added MPI Advance locality aware neighbor collectives and topology object to BoomerAMG solver; integration with production code planned
  - Trilinos: MPI Advance locality aware operations added to Trilinos TPetra
  - xRAGE: Added communication pattern profiling, integrated MPI advance neighbor collectives.
- Research Frameworks and Benchmarks
  - Benchpark: Added Caliper annotation support to a Benchpark branch; merge to main planned
  - Cabana: Added additional communication abstractions, enhanced communication abstractions with MPI Advance neighbor discovery and exchange primitives
  - COMB: Added persistent and partitioned communication backends. Persistent communication backend integrated with production benchmark.
  - AMG2023, Kripke, Laghos: Added Caliper annotations to released LLNL benchmarks

Publication Highlights

42 Peer-Reviewed Publications Accepted or Published

- Jared Dominguez-Trujillo, Keira Haskins, Soheila Jafari Khouzani, Christopher Leap, Satoshi Tashakkori, Quincy Wofford, Trilce Estrada, Patrick G. Bridges, and Patrick M. Widener. Lightweight measurement and analysis of HPC performance variability. In 2020 IEEE/ACM Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS), pages 50–60, 2020.
- Nawrin Sultana, Martin R¨ufenacht, Anthony Skjellum, Purushotham Bangalore, Ignacio Laguna, and Kathryn Mohror. Understanding the use of message passing interface in exascale proxy applications. Concurrency and Computation: Practice and Experience, page e5901, 2020.
- Matthew G. F. Dosanjh, Andrew Worley, Derek Schafer, Prema Soundararajan, Sheikh K. Ghafoor, Anthony Skjellum, Purushotham V. Bangalore, and Ryan E. Grant. Implementation and evaluation of MPI 4.0 partitioned communication libraries. Parallel Comput., 108:102827, 2021.
- W. Pepper Marts, Matthew G. F. Dosanjh, Scott Levy, Whit Schonbein, Ryan E. Grant, and Patrick G. Bridges. Minimod: A modular miniapplication benchmarking framework for HPC. In IEEE International Conference on Cluster Computing, CLUSTER 2021, Portland, OR, USA, September 7-10, 2021, pages 12–22. IEEE, 2021.
- Shelby Lockhart, Amanda Bienz, William Gropp, and Luke Olson. Performance analysis and optimal node-aware communication for enlarged conjugate gradient methods. ACM Trans. Parallel Comput., 10(1), Mar 2023.
- Shelby Lockhart, Amanda Bienz, William D. Gropp, and Luke N. Olson. Characterizing the performance of node-aware strategies for irregular point-to-point communication on heterogeneous architectures. Parallel Computing, 116:103021, 2023.
- Patrick G. Bridges, Anthony Skjellum, Evan D. Suggs, Derek Schafer, and Purushotham V. Bangalore. Understanding GPU triggering apis for MPI+X communication. In Recent Advances in the Message Passing Interface: 31st European MPI Users' Group Meeting, EuroMPI 2024, Perth, WA, Australia, September 25–27, 2024, Proceedings, page 39–55, Berlin, Heidelberg, 2024. Springer-Verlag.
- Gerald Collom, Derek Schafer, Amanda Bienz, Patrick Bridges, and Galen Shipman. Optimizing neighbor collectives with topology objects. In 2024 IEEE International Conference on Cluster Computing (CLUSTER), pages 120–130, 2024.
- Carson Woods, Derek Schafer, Patrick G. Bridges, and Anthony Skjellum. Quantifying and Modeling Irregular MPI Communication. In 2024 IEEE 24th International Symposium on Cluster, Cloud and Internet Computing (CCGrid), pages 525–533, Los Alamitos, CA, USA, May 2024. IEEE Computer Society.

Key DOE Laboratory Collaborators

- Lawrence Livermore National Laboratory: David Boehme, Stephanie Brink, Ignacio Laguna, Ian Lee, Rui Peng Li, Kathryn Mohror, Olga Pearce, Roger Pearce, Peter Pirkelbauer
- Los Alamos National Laboratory: Jered Dominguez-Trujillo, Christoph Junghans, Ryan Marshall, Howard Pritchard, Jon Reisner, Robert Robey, Galen Shipman, Quincy Wofford
- Sandia National Laboratories: Jan Ciesko, Matthew Curry, Matthew Dosanjh, Kurt Ferreira, Ryan Grant, Scott Levy, Reed Milewicz, Carl Pearson, Stephen Olivier, Whit Schonbein, Chris Siefert
- Argonne National Laboratory: Ken Raffenetti, Hui Zheng
- Oak Ridge National Laboratories: Jack Lange, Sam Reeve, Stuart Slattery, Trey White
- Pacific Northwest National Laboratory: Kevin Barker, Jesun Firoz, Sayan Ghosh, Joseph Manzano, Joshua Suetterlein, Stephen Young

Educational Activities

- Tutorial on *MPI Advance - Optimizations and Extensions to MPI* at the NNSA-University Workshop on Exascale Simulation Technologies (NUWEST) workshop
- Tutorial on *MPI Partitioned Communication* at EuroMPI 2023 Symposium
- Developing tutorial on *High Performance Libfabric Programming* for Cray User Group 2026
- Hackathons
  - Weekly online mini-hackathons with students to help with their research
  - Multiple in-person hackathons for students focused on specific research and training problems, including MPI topology objects, libfabric programming, and developing student programming skills
- Course development
  - Created homework assignments focused building new communication primitives for use in multiple courses, tutorials and workshops
  - Offered multiple courses at member institutions on building high-performance communication abstractions

Outreach Activities

- Contributed significant enhancements to the MPI 4.0 and 4.1 standards
- Organized the Workshops on Extreme Scale MPI at SC 2021-2025
- Organized SIAM PP24 mini-symposium on *Realistic Proxy Applications and Datasets for Heterogeneous Architecture Scalable Communication*
- Organized the 4th and 5th Workshop on Compiler-assisted Correctness Checking and Performance Optimization for HPC (C3PO) at ISC 2024 and 2025
- Hosted seminars on advanced communication topics featuring speakers from DOE national laboratories
- UNM PSAAP staff member and student lead and mentored team for SC Student Cluster Competition
- Offered advanced seminars at multiple student organizations

CUP-ECS Provided DOE a Full Suite of Communication Abstractions, Models, and Tools

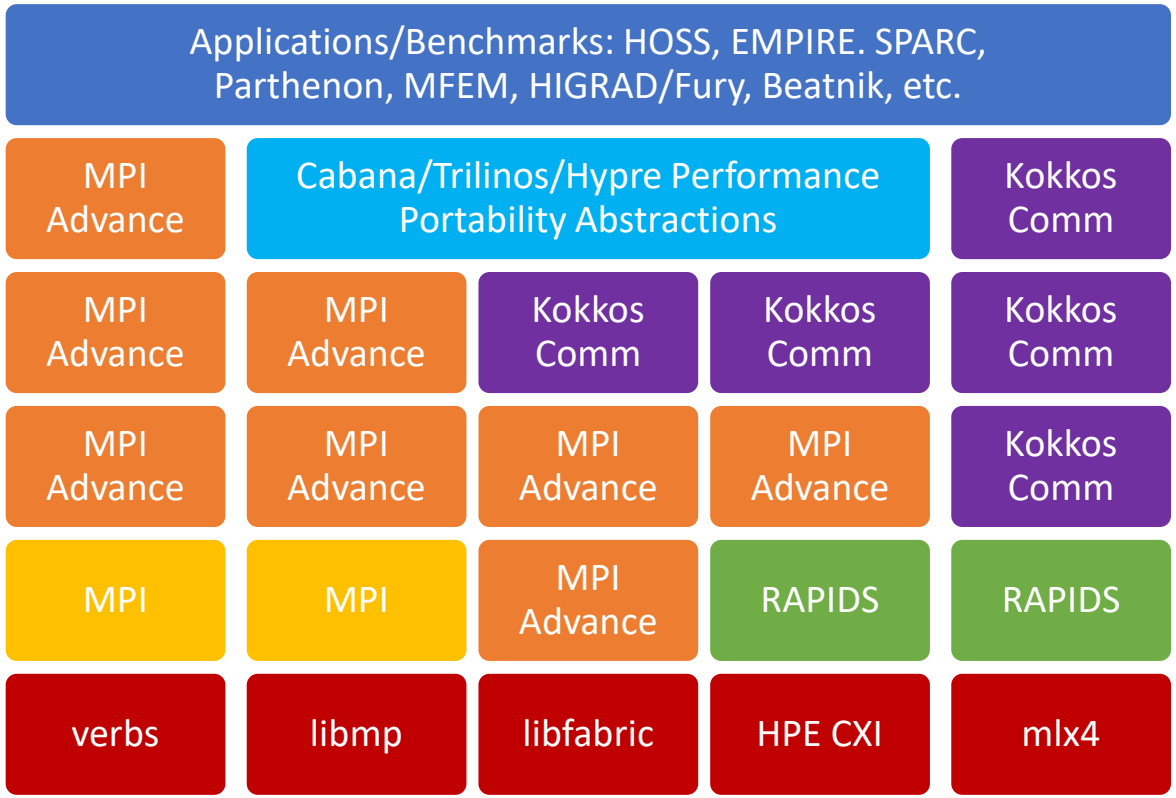

- Researched High-level Application-facing Communication Abstractions and Low-level Primitives necessary to support them
- Development at each level driven by careful assessment, benchmarking, and modeling workflows and tools
- Carefully integrating these abstractions with a wide range of DOE applications and libraries, including multiple production NNSA codes
- Demonstrated performance improvements in full production applications