

# **Year 1 Education Activities and Lab Interactions Annual Review Meeting**

Prof. Patrick G. Bridges

Lead PI and Center Director

August 23, 2021



Center for Understandable, Performant Exascale Communication Systems



# Y1 Education Focused on Student Development

- Recruited many new students to project,
- Education activities in Y1 focused primarily on training them
- Bringing them up to speed on relevant issues
  - Performance monitoring and modeling in HPC systems
  - MPI abstraction, design, and implementation issues
  - Relevant application characteristics
- Building their ties to and knowledge of NNSA lab issues
  - Internships were a key part of this

# Y1 Lab Internships

- PSAAP Internships
  - Jered Dominguez-Trujillo (UNM) - Sandia with Ryan Grant and Matthew Dosanjh
  - Rei Haskins (UNM) – Sandia with Kurt Ferreira and Scott Levy
  - Gerald Collom (UNM) – LLNL with Olga Pearce
  - Tanner Broadus (UTC) – LANL with Galen Shipman
- Other related lab internships
  - Pepper Marts (UNM) – Year-round at Sandia with Matthew Dosanjh
  - Carson Woods (UTC) – Summer 2021 at Sandia with Matthew Curry
- Reverse internships
  - Ryan Goodner (UNM) - Sandia technologist and UNM student hired by the project over summer 2021 to learn HPC and MPI performance issues, reproducibility techniques

# Other Training Activities – Colloquium Series

online colloquium series for students, other project personnel, and interested community members to supplement center-internal talks, discussions and mentoring.

- January 22, Zoom: Colloquium - Kokkos: C++ Performance Portability on HPC Platforms - *Tanner Broaddus, University of Tennessee at Chattanooga*
- February 5, Zoom: Colloquium - Determining Proxy Application Fidelity using ML-Based Similarity Methods - *Jeanine Cook, Sandia National Laboratories*
- February 19, Zoom: Colloquium - Adding Fast GPU Derived Datatype Handing to Existing MPIs - *Carl Pearson, University of Illinois (now Sandia National Laboratories)*
- March 5, Zoom: Colloquium - MPI Partitioned Communication for Highly Concurrent and Heterogeneous Systems - *Ryan Grant, Sandia National Labs*
- March 26, Zoom: Colloquium - DPC++: A modern C++ parallel programming model for Heterogeneous Computing - *Nawrin Sultana, Intel Corporation*
- April 9, Zoom: Colloquium - Recent Development of Multigrid Solvers in HYPRE on Modern Heterogeneous Computing Platforms - *Ruipeng Li, Lawrence Livermore National Laboratory*



# Additional lab interactions and student/education outreach activities

- Bi-weekly coordination meetings with TST lead Jon Reisner
- Partitioned communication research with Sandia staff members Ryan Grant and Matthew Dosanjh
- Performance modeling research with Sandia staff member Patrick Widener
- Periodic group discussions with Olga Pearce, and other LLNL staff on Comb benchmark issues
- Discussions with Sandians Karen Devine and Carl Pearson on potential Trilinos TPetra MPI abstractions
- Discussions with LANL staff members Galen Shipman and Bob Robey on irregular communication issues in adaptive mesh refinement codes
- Discussions with Sandia staff member James Elliot on EMPIRE performance characteristics
- Discussion with Sandia staff member Matthew Curry on use of ReframeHPC as experiment management framework for the center
- “Optimizing MPI for Modern Applications and Systems”, Colloquium talk for U.S. Naval Academy parallel programming students, Patrick Bridges, April 22, 2021.



# **Year 2 Plans and COVID Impact**

## **Annual Review Meeting**

Prof. Patrick G. Bridges

Lead PI and Center Director

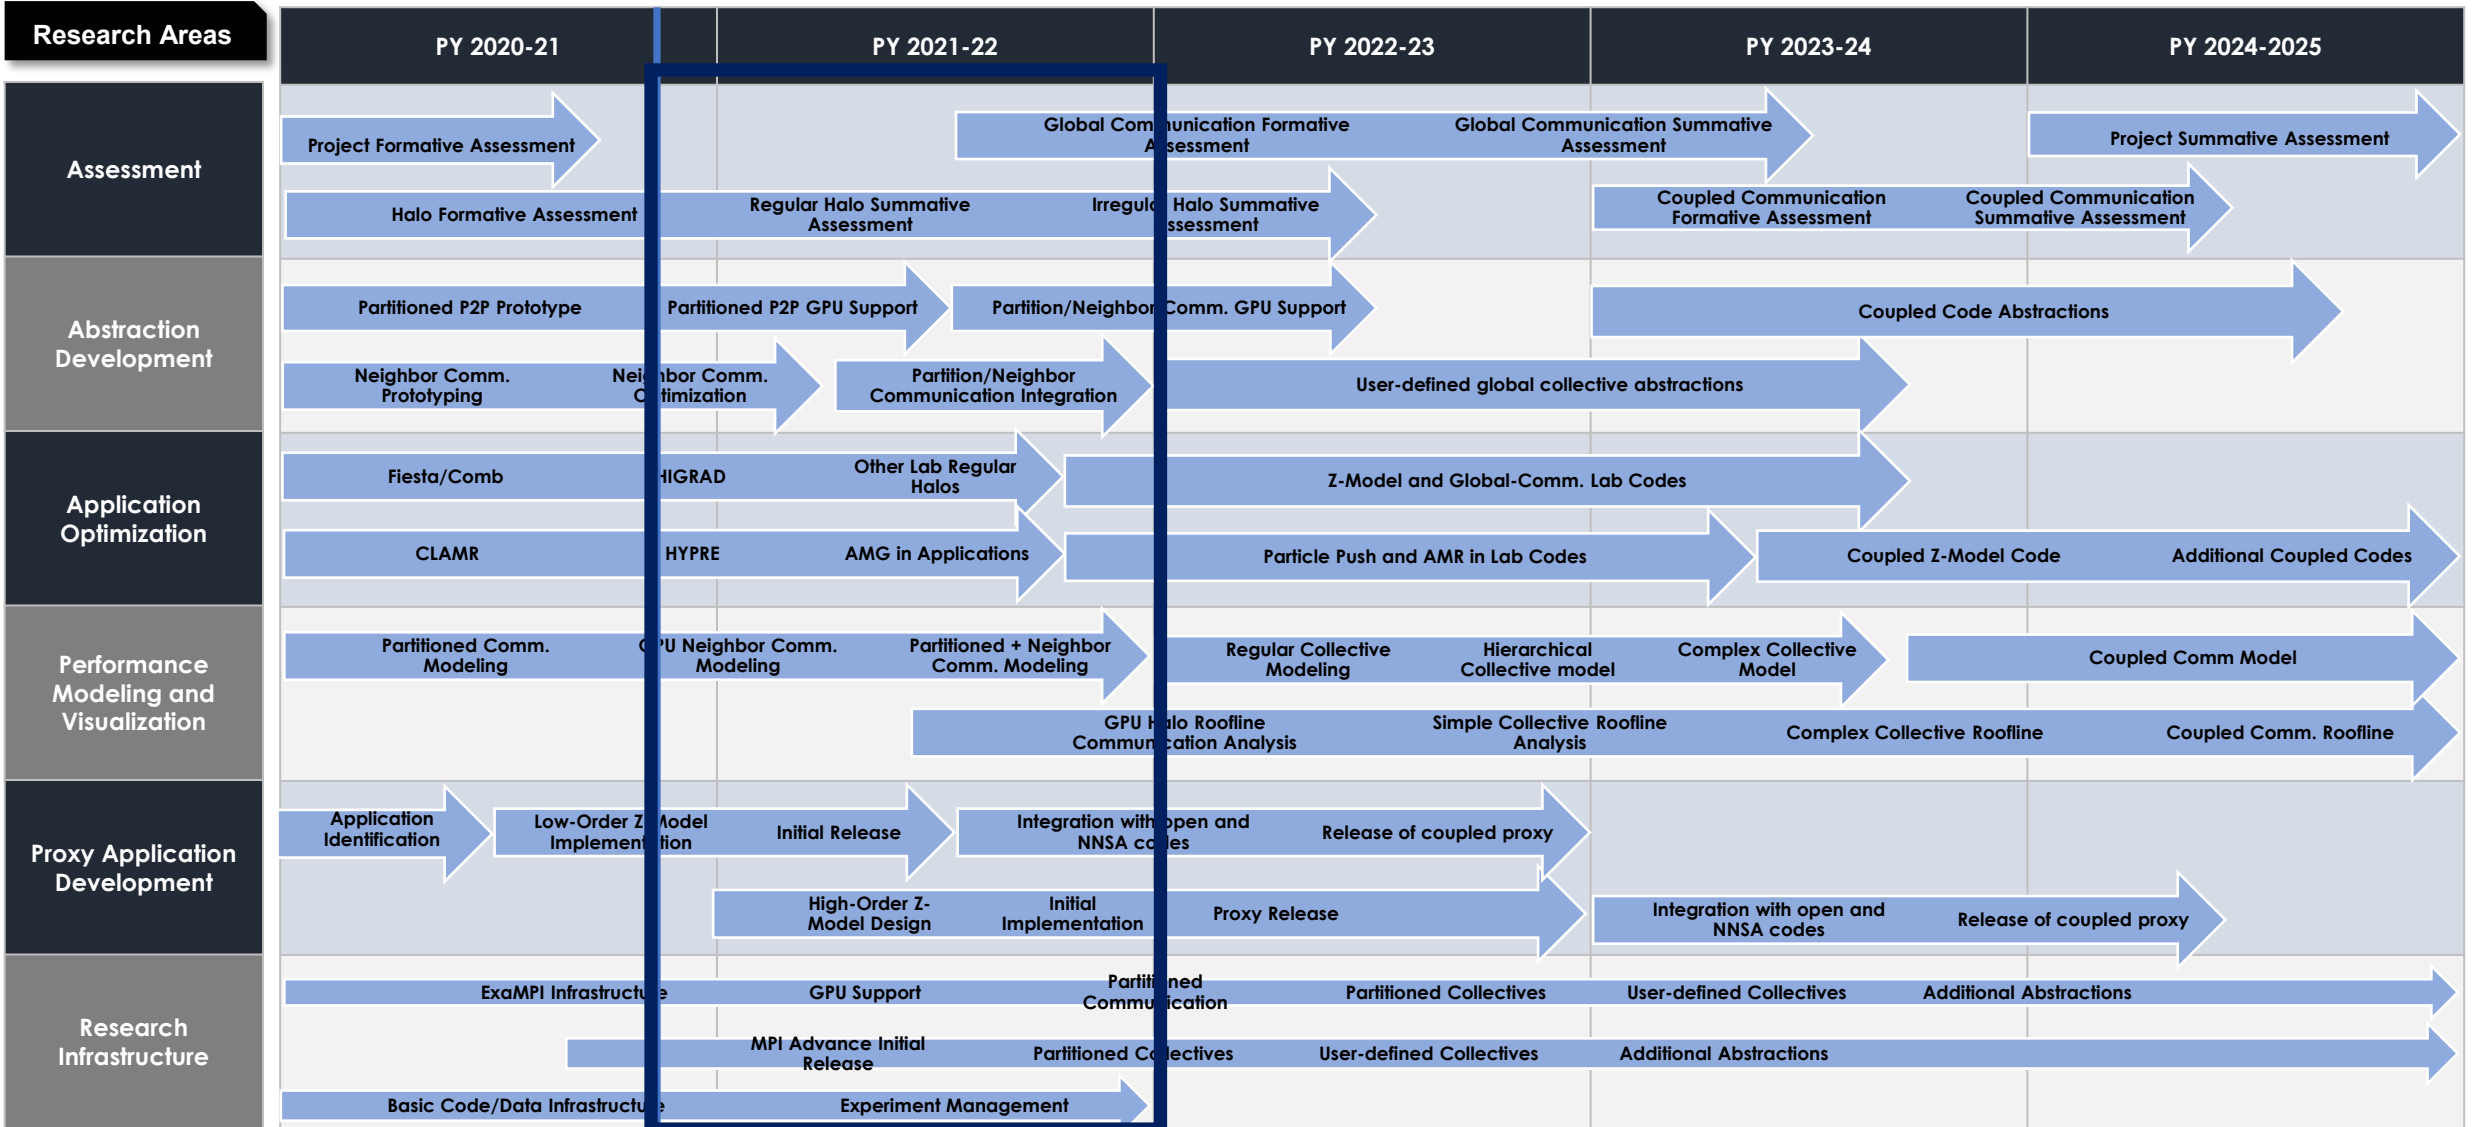
August 23, 2021



Center for Understandable, Performant Exascale Communication Systems



# 5-year Project Roadmap



# Year 2 Regular and Irregular Halos

- Regular Halos
  - Integrate neighbor collectives and (non-GPU) partitioned communication into lab benchmarks applications (HIGRAD and Comb)
  - Perform summative assessment of non-GPU exchanges, begin summative assessment of GPU halo exchanges
  - Begin GPU implementation of partitioned primitives (primarily MPI Advance, infrastructure in ExaMPI)
- Irregular Halos
  - Optimization of current MPI neighbor collectives
  - Research new graph creation and neighbor collective variants to enable better optimizations
  - Assessment in HYPRE AMG solver, potential evaluation of optimized HYPRE for supporting implicit solves in HIGRAD
  - Start integrating these primitives into other irregular halo codes (potential targets – xRage, Trilinos, etc.)
- Examine partitioned neighbor collectives in both cases



# Proxy Development and Research Infrastructure

- Proxy Development
  - Release low-order model as public benchmark using test cases to be specified in collaboration with Shkoller and Reisner
  - Design parallel decomposition/implementation of high-order model - examining FleCSI as framework because of FMM support (in FleCSPH) and Legion
- Research Infrastructure
  - MPI Advance initial release
  - Deploying repository for storing/sharing data artifacts (traces, profiles, etc.)
  - Examining RescaleHPC as possible experiment management infrastructure for specifying and logging performance results



# Key New Year 2 Activities

- Begin assessment of global communication challenges
  - Low-order Z-model FFTs (already working with UTK, CMU)
  - Neighbor discovery in irregular codes
  - Other algorithms with global communication (FMM, sorting, etc.)
- High-order Z-model design, begin implementation
- Examine roofline or other performance visualizations for GPU exchanges (both P2P and neighbor collective)

# COVID Impacts

- This is the first time our team has met in person!
- Coordinating activities across sites has been a real challenge
  - Made aggressive use of collaboration tools like slack
  - Regular (weekly) coordination meetings have helped, created regular meetings for key subgroups
- Internships were virtual
  - Labs worked hard to make them very productive (as you'll see next)
  - Still didn't get full on-site experience with other students, lab personnel, etc.
- Difficult to hire UNM program coordinator

# Questions?



Center for Understandable, Performant Exascale Communication Systems

