Overview of Activities and Accomplishments

PSAAP-III Annual Review

Patrick G. Bridges

September 28, 2023

Updated 5-year Project Roadmap

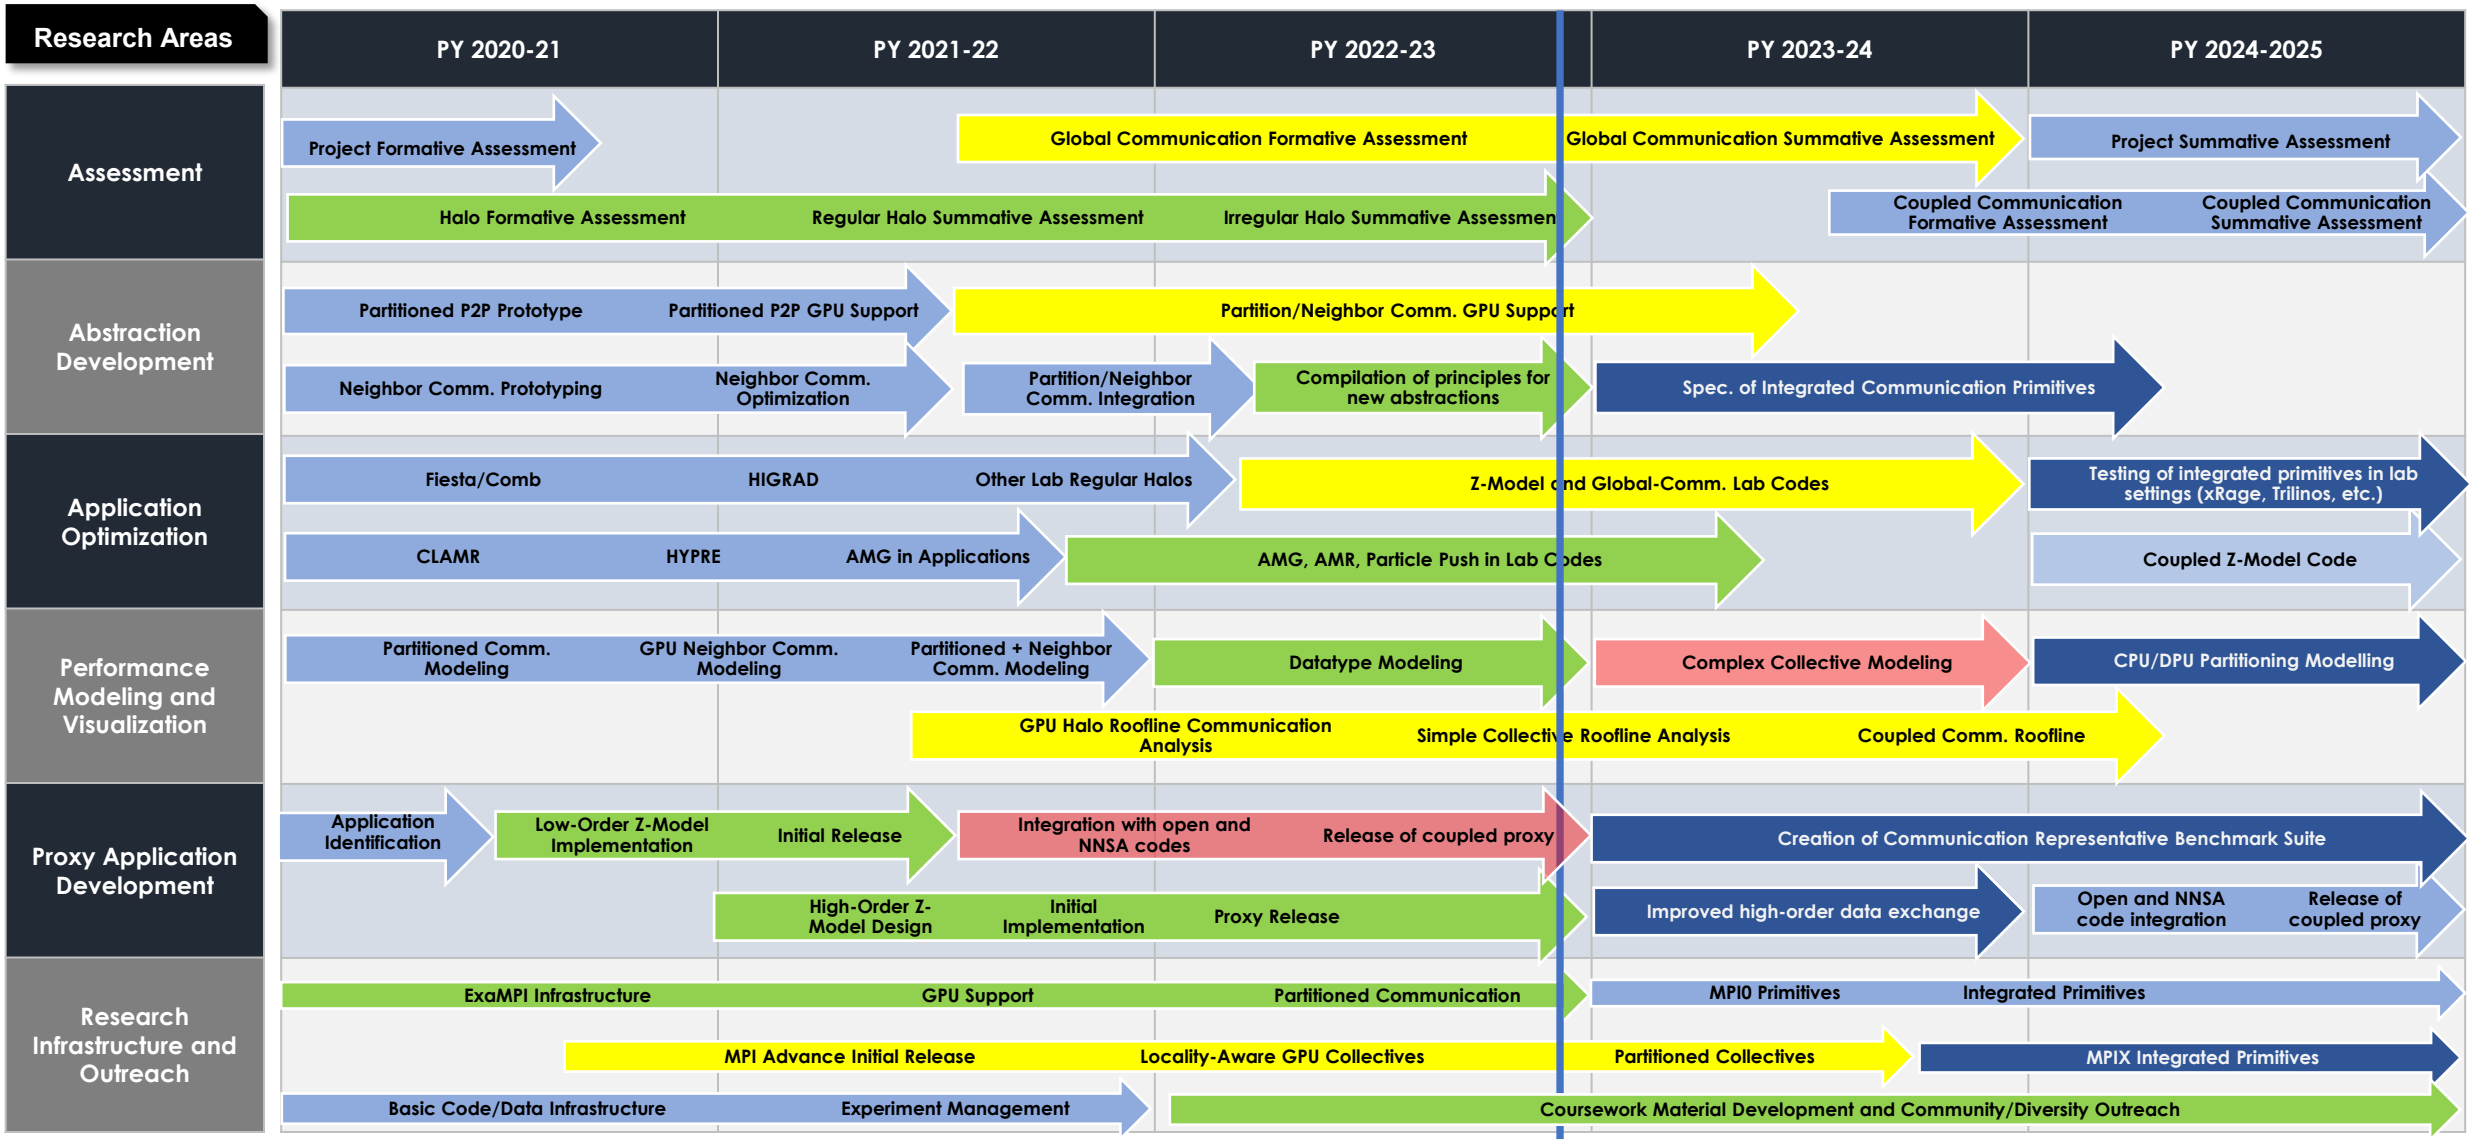

Communication Abstraction Stack View

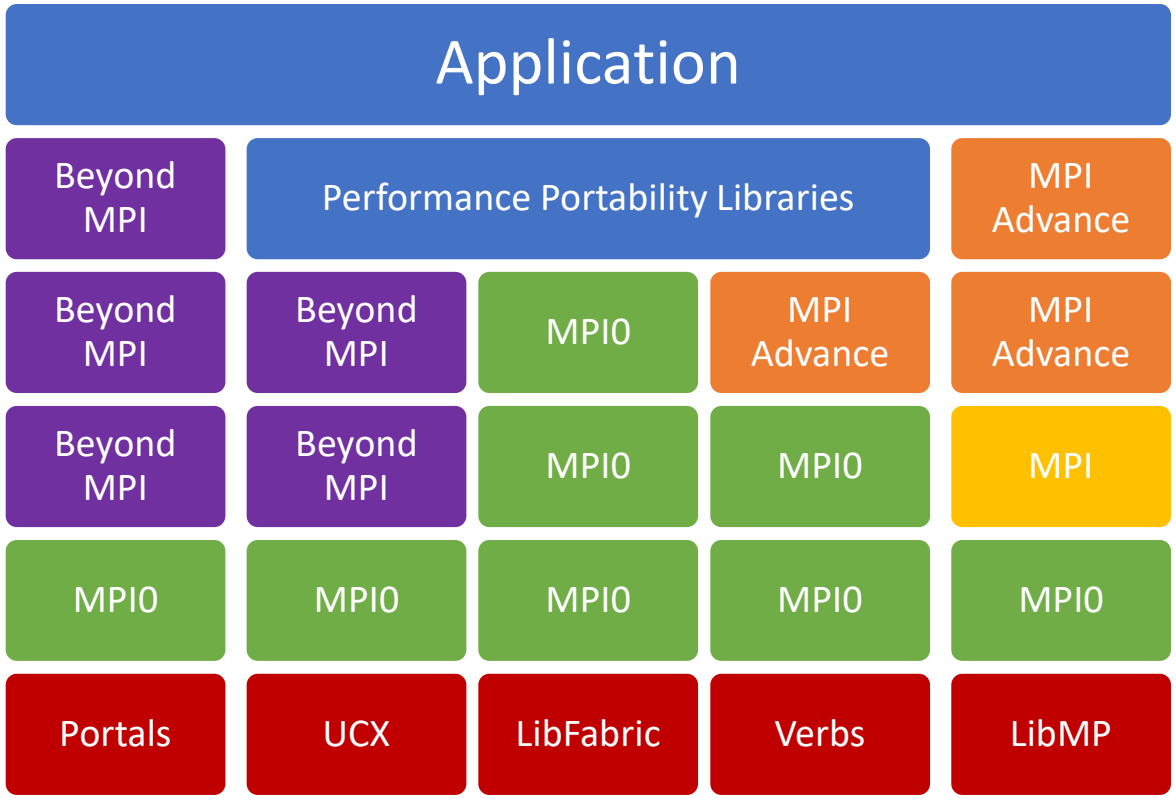

- Goal is to develop new abstractions for applications and libraries
  - Beyond MPI: Application-oriented abstractions that leverage our findings and are unconstrained by legacy MPI
  - MPI Advance: Application-oriented abstractions that extend and push the frontier of existing abstractions
  - MPIO: New low-level communication primitives for use by library writers for building new abstractions
- Development at each level is driven by careful assessment, benchmarking, and modeling
- Through the first two years, focus has mostly been on Advance-level primitives
- Shifting focus to MPIO and Beyond abstractions

GPU and Point to Point Communication

Goal

- Assessment to inform development of MPI0 and Advance abstractions
  - Understand communication requirements and constraints
  - Focus on simple patterns on advanced architectures

Accomplishments

- Pulse benchmark
  - Assess different low-level communication approaches
  - Assess halo communication optimization tradeoffs
- Findings
  - Needs and requirements for low-level portable communication abstractions
  - Needs and requirements for high-level halo-exchange abstractions

Benchmarking and Modeling

Goals

- Assessment of broader communication requirements
  - Key communication abstractions
  - Additional communication patterns
  - Lab communication workloads
- Provide benchmarks and tools to ensure abstractions and optimizations are relevant to lab workloads
- Provide models to predict impact of new abstractions on new architectures

Accomplishments

- Assessment and Modeling
  - Irregular lab communication workloads
  - Datatype abstractions
  - Application scaling models
- Benchmark development
  - Release of Beatnik global communication/ coupled communication benchmark
  - Creation of irregular communication benchmark to replay lab communication patterns
  - Creation of a broader strategy for communication benchmarking

Irregular Communication

Goal: Portable optimizations for codes with irregular communication

Accomplishments

- Integration of MPI Advance into HYPRE and Trilinos
- Optimization of HYPRE using MPI Advance primitives
- Optimization of GPU-based AllToAll communication
- Performance analysis of topology identification algorithms
- Designing abstractions to improve topology identification and topology-based neighbor communication

Abstraction Creation

Goal: Create new communication abstractions across the stack that can be used to improve the performance of lab applications

Accomplishments

- Requirement assessment and creation of initial MPI0 abstractions as a portable, low-level communication abstractions for high-level abstractions
- Identification and evaluation of existing application and library communication primitives that our communication libraries should optimize
- Creation of Kokkos abstractions beyond the standard MPI interface to support first class communication of Kokkos data structures

Research Infrastructure and Outreach

Goals

- Organize abstractions into manageable infrastructure for experimentation
- Outline HPC and MPI fundamentals for class, book development.
- Encourage and increase opportunities for diversity and inclusion

Accomplishments

- Research Infrastructure
  - Addition of more MPI functionality to ExaMPI
  - Integration and spread of MPI Advance
- Education
  - Concretizing assignments, materials for HPC/MPI class
- Outreach
  - Plan to work at Grace Hopper conference
  - Help out in various cluster teams