

Assessment Planning and Next Steps

Purushotham V. Bangalore
Department of Computer Science
University of Alabama at Birmingham



Center for Understandable, Performant Exascale Communication Systems

UAB THE UNIVERSITY OF
ALABAMA AT BIRMINGHAM

Assessment Goals

- Identify applications and input decks that are representative of ongoing and future communication challenges,
- Understand the qualitative and quantitative needs of these applications,
- Assemble public test suites based on this assessment to drive future center and broader community research on optimizing HPC communication systems

Application Characteristics

- Capture communication challenges faced by the broad range of DOE applications on Exascale architectures
- Identify MPI primitives and their common usage, including point-to-point, collective, and one-sided communication
- Represent inter-node communication patterns used by realistic applications, including neighbor halo exchange, static irregular, and dynamic irregular communication
- Express code complexity of full application frameworks and simplified proxy and mini-applications
- Represent relevant programming frameworks and languages
- Cover NNSA-relevant application areas

Application Selection Process

- Examined
 - ECP proxy application suite
 - 100+ open source MPI applications
- Narrowed this down to 40 DOE production applications and proxy or mini-applications
- Further discussion with TST team and NNSA lab collaborators to understand lab application needs
 - Identify production applications with communication challenges
 - Identify proxy or mini-applications that are representative of these production applications

Production & Proxy/mini Apps, I

Application Miniapp	Lab	Availability	Priority	Description
LLNL CFD/Mechanics Apps	LLNL	EC	2	ALE3D, Nike3d, etc.
Comb		Open	1	GPU stencil communication mini-application
EMPIRE	SNL	EC	3	Hybrid PIC Electrodynamics framework
EMPIRELite	SNL	In devel	2	Forthcoming EAR99 version
RefMaxwell	SNL	Open	1	Trilinos AMG solver used in EMPIRE
ExaMPM	ECP-COPA	Open	1	Proxy for EMPIRE particle push
HIGRAD	LANL	EC	1	Shock CFD application
Fiesta	UNM	Open	1	Partial open reimplemention & modernization of key HIGRAD features
MERCURY	LLNL	EC	3	Monte Carlo Radiation Dynamics
Quicksilver	LLNL	Open	2	Simplified Monte Carlo Transport Proxy

Production & Proxy/mini Apps, II

Application Miniapp	Lab	Availability	Priority	Description
PARTISN	LANL	EC	3	Neutron transport; DesignForward traces available
SNAP		Open	2	Proxy of PARTISN compute/communication patterns
SPARC	SNL	EC	3	Reacting and non-reacting hypersonic CFD code
MiniAero		Open	2	Unstructured Navier-Stokes solver mini-app
xRage	LANL	EC	3	Radiation Transport/Hydrodynamics Framework
CLAMR		Open	1	2D Cell-based adaptive mesh refinement mini-app.
EAP Proxy		In devel	2	Forthcoming xRage proxy application

Assessment Phases

Each phase involves modeling, evaluating, and optimizing applications that:

- I. use predominantly point-to-point and nearest-neighbor communication on both CPU and GPU-based systems
- II. use predominantly collective communication primitives
- III. are based on adaptive algorithms and/or use complex communication patterns

Initial Assessment Results

- Preliminary static analysis of a wide range of ECP proxy applications that was conducted prior to CUP-ECS center initiation
- Formative qualitative assessment of the Priority-1 applications identified earlier
- Initial assessment of the ability of MPI nearest neighbor collectives to support a higher-level irregular halo communication primitive that our qualitative assessment showed was common in some DOE applications and frameworks

Outstanding Issues/Challenges

- Feedback on appropriateness of production and proxy/mini applications selected for assessment
- Assistance with identifying unlimited release input decks that expose relevant communication system issues in export controlled applications
- Recommendations on test platforms for qualitative assessment activities and details on upcoming Exascale architectures (e.g., AMD GPUs)

Open Questions

- 1) Is the application coverage sufficient and representative of NNSA application areas?
- 2) Do the proxy/mini applications represent the corresponding production applications with appropriate input decks?
- 3) Do we have input decks available for both production and proxy applications?

Application Miniapp	Lab	Availability	Priority	Description
LLNL CFD/Mechanics Apps	LLNL	EC	2	ALE3D, Nike3d, etc.
Comb		Open	1	GPU stencil communication mini-application
EMPIRE	SNL	EC	3	Hybrid PIC Electrodynamics framework
EMPIRELite	SNL	In devel	2	Forthcoming EAR99 version
RefMaxwell	SNL	Open	1	Trilinos AMG solver used in EMPIRE
ExaMPM	ECP-COPA	Open	1	Proxy for EMPIRE particle push
HIGRAD	LANL	EC	1	Shock CFD application
Fiesta	UNM	Open	1	Partial open reimplementaion & modernization of key HIGRAD features
MERCURY	LLNL	EC	3	Monte Carlo Radiation Dynamics
Quicksilver	LLNL	Open	2	Simplified Monte Carlo Transport Proxy
PARTISN	LANL	EC	3	Neutron transport; DesignForward traces available
SNAP		Open	2	Proxy of PARTISN compute/communication patterns
SPARC	SNL	EC	3	Reacting and non-reacting hypersonic CFD code
MiniAero		Open	2	Unstructured Navier-Stokes solver mini-app
xRage	LANL	EC	3	Radiation Transport/Hydrodynamics Framework
CLAMR		Open	1	2D Cell-based adaptive mesh refinement mini-app.
EAP Proxy		In devel	2	Forthcoming xRage proxy application

Q1: Application Areas Represented

- Computational fluid dynamics
- Electromagnetic plasma simulations
- Atmospheric modeling
- Monte Carlo radiation dynamics
- Neutron transport
- Reacting and non-reacting hypersonic CFD
- Radiation transport/hydrodynamics

Q2: Production vs. Proxy/Mini Apps

Example 1:

EMPIRE	SNL	EC	3	Hybrid PIC Electrodynamics framework
EMPIRELite	SNL	In devel	2	Forthcoming EAR99 version
RefMaxwell	SNL	Open	1	Trilinos AMG solver used in EMPIRE
ExaMPM	ECP-COPA	Open	1	Proxy for EMPIRE particle push

Example 2:

xRage	LANL	EC	3	Radiation Transport/Hydrodynamics Framework
CLAMR		Open	1	2D Cell-based adaptive mesh refinement mini-app.
EAP Proxy		In devel	2	Forthcoming xRage proxy application

Open Questions

- 1) Is the application coverage sufficient and representative of NNSA application areas?
- 2) Do the proxy/mini applications represent the corresponding production applications with appropriate input decks?
- 3) Do we have input decks available for both production and proxy applications?

Questions and Discussion



Center for Understandable, Performant Exascale Communication Systems

UAB THE UNIVERSITY OF
ALABAMA AT BIRMINGHAM