

# Benchmark for AMR Based Unstructured Communication

Savannah Camp  
Dr. Anthony Skjellum  
Dr. Patrick Bridges

This work was supported in part by the U.S. Department of Energy's National Nuclear Security Administration (NNSA) under the Predictive Science Academic Alliance Program (PSAAP-III), Award #DE-NA0003966.



Center for Understandable  
Performant Exascale  
Communication Systems

 THE UNIVERSITY OF TENNESSEE  
CHATTANOOGA

# Background

- CLAMR - Collection of AMR mini-apps
- L7 communication library
- Performance anomaly amongst MPI Implementations
  - MVAPICH2 stuck waiting on requests
  - Introduction of MPI\_Barrier improves MVAPICH2 performance


	Min (s)	Medium (s)	Max (s)
MVAPICH2	0.4611	1.0801	9.3026
OpenMPI	0.2565	0.4076	0.5316
MVAPICH2 w/ Barrier	0.3218	0.4595	0.5910
OpenMPI w/ Barrier	0.3205	0.4522	0.5785

# L7/CLAMR Observations - Tracing



CLAMR-OpenMPI

# L7/CLAMR Observations - Tracing




CLAMR-MVAPICH2



Center for Understandable  
Performant Exascale  
Communication Systems

THE UNIVERSITY OF TENNESSEE  
CHATTANOOGA

# L7/CLAMR Observations - Tracing




CLAMR-MVAPICH2 w/ MPI\_Barrier



Center for Understandable  
Performant Exascale  
Communication Systems

THE UNIVERSITY OF TENNESSEE  
CHATTANOOGA


# L7/CLAMR Observations - Tracing



CLAMR-OpenMPI w/ MPI\_Barrier



# L7/CLAMR Observations - Profiling



CLAMR-OpenMPI Profile w/ Paraprof





# Benchmark

- Work-in-progress benchmark
- Motivated by observed performance anomalies
- How well an MPI implementation performs unstructured communication
- Goals:
  - Observe MPI's performance in unstructured communication
  - Find unexpected performance anomalies



# Use Cases

- Aims to test these different cases:
  - Various ordering of non blocking MPI calls and synchronization primitives (or lack thereof)
  - Different sizes of messages when performing unstructured communication
  - Performance of unstructured communication in different runtime environments (single-node, multi-node, etc)
- Compare performance data between different MPI implementations




# Requirements

- L7 will be used as the facilitating library for performing unstructured communication
- The benchmark will have the option of introducing a synchronization primitive (MPI\_Barrier) within the code before/during/after nonblocking communication
- The benchmark will provide Min/Median/Max times for regions of interest



# Analysis Tools -- Profiling and Tracing

- Caliper
  - Source code annotations
  - Region profiling
- Tau
  - Paraprof - profiling
  - Jumpshot - tracing



# References and Acknowledgements

## References

- Caliper: <https://computing.llnl.gov/projects/caliper>
- CLAMR: <https://github.com/CUP-ECS/CLAMR>
- L7: <https://www.sandia.gov/~rfbarre/l7.v1.ug.ps>
- TAU: <https://www.cs.uoregon.edu/research/tau/home.php>

## Acknowledgements

- Dr. Patrick Bridges
- Justin T. Broaddus
- Dr. Anthony Skjellum
- Dr. Puri Bangalore



Center for Understandable  
Performant Exascale  
Communication Systems

 THE UNIVERSITY OF TENNESSEE  
CHATTANOOGA