Oak Ridge National Laboratory
Computing and Computational Sciences Directorate

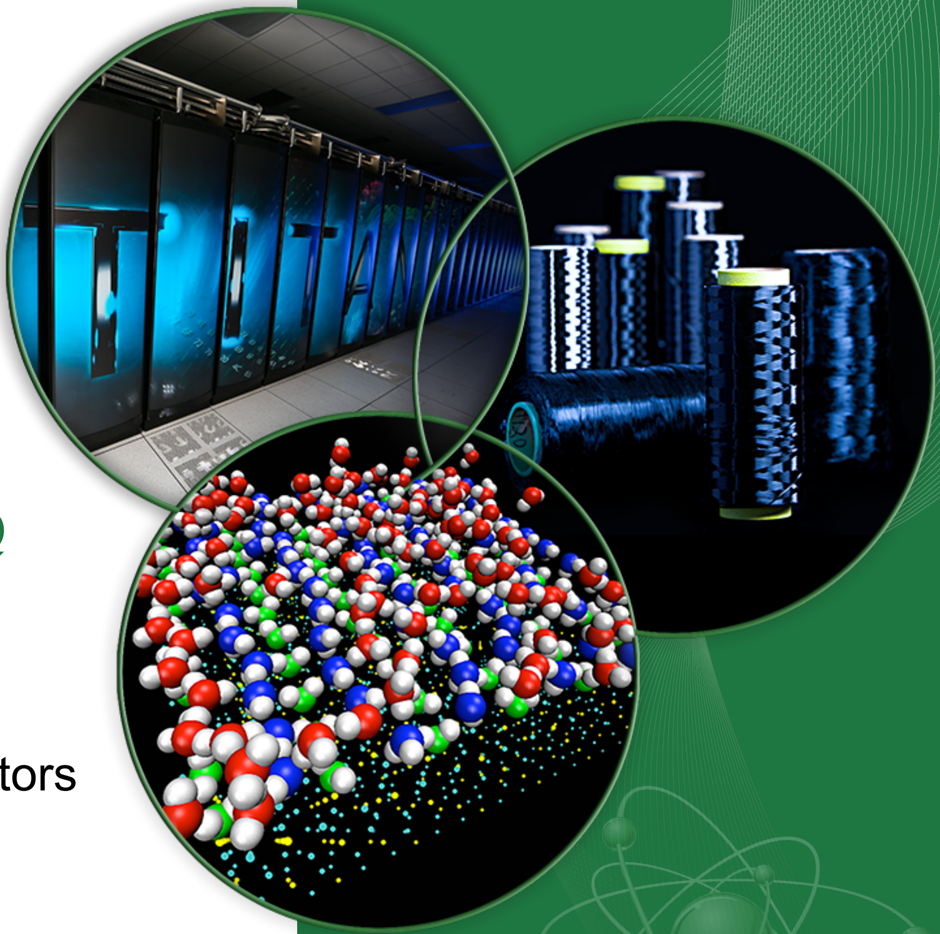

OpenSHMEM Activities @ ORNL

Manjunath Gorentla Venkata
for the Languages Team and Collaborators

Computational Research & Development Programs

Team

- **ORNL**
  - Neena Imam (Director)
  - Manjunath Gorentla Venkata / Manju (Team Lead)
  - Mathew Baker
  - Swen Boehm
  - Swaroop Pophale
  - Thomas Naughton
  - Eric Mitchell
- **DoD**
  - Nick Park
  - Bryant Lam
- **NVIDIA**
  - Sreeram Potluri
  - Anshuman Goswami
  - CJ Newburn
- **Florida State University**
  - Weikuan Yu and Students
- **Paratools**
  - Sameer Shende
  - John Linford

Team Website: **http://www.csm.ornl.gov/openshmem/index.html**

OpenSHMEM : Specification, Research, and Implementation

Overview

- Languages team at ORNL drives the development of the OpenSHMEM programming model and its ecosystem. The effort is focused in four areas:
  - Specification development
  - OpenSHMEM implementation
  - Research and exploratory topics
  - Outreach

Software Artifacts

- OpenSHMEM-X
- ORNL OpenSHMEM Benchmark Suite (OSB)
- OpenUCX : Unified Communication X
- NVSHMEM : OpenSHMEM for GAS Systems
- SHMEMCache: OpenSHMEM based Memcache

Outreach Activities

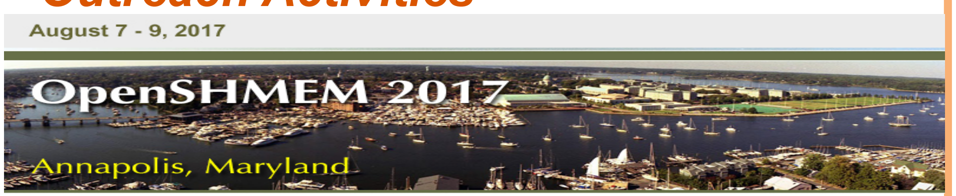

Specification

- **OpenSHMEM 1.4 & Next specification :** Developed OpenSHMEM thread safety proposal, and merged request proposal

R&D

- **OpenSHMEM-X :** Added various capabilities including alltoall and alltoallv collective operations, multithread support, OpenSHMEM contexts, explicit RMA operations, and merged request handles
- OpenSHMEM-X uses **PMIx** for scalability, performance and stability
- **OpenSHMEM Benchmarks:** Developed various OpenSHMEM application kernels and benchmarks including multithreaded Graph 500, latency, bandwidth, SSCA, GUPs, SSSP, and YCSB
- **SHMEMCache** – Memcached implementation, which uses OpenSHMEM for communication
- **NVSHMEM** – OpenSHMEM for GPUs connected via PCIe, NVLINK or InfiniBand
- Improvements to TAU ecosystem for OpenSHMEM programmers with support for Callsites, OTF2 and simplified user interface
- Design, evaluation, and results are available as papers

Computational Research & Development Programs

Team Website: http://www.csm.ornl.gov/openshmem/index.html
Team GitHub: https://github.com/orgs/ornl-languages/

Oak Ridge National Laboratory

Multithreaded OpenSHMEM and Contexts

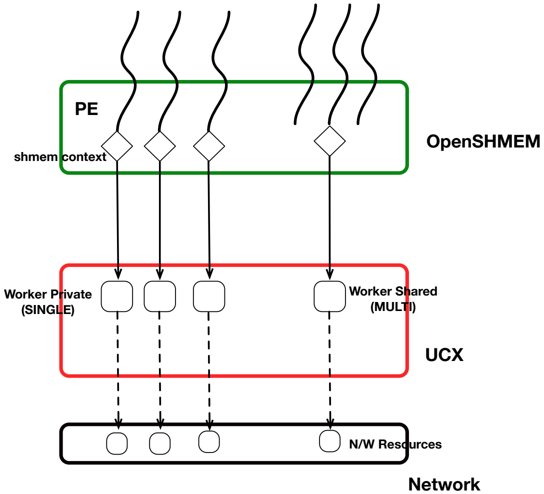

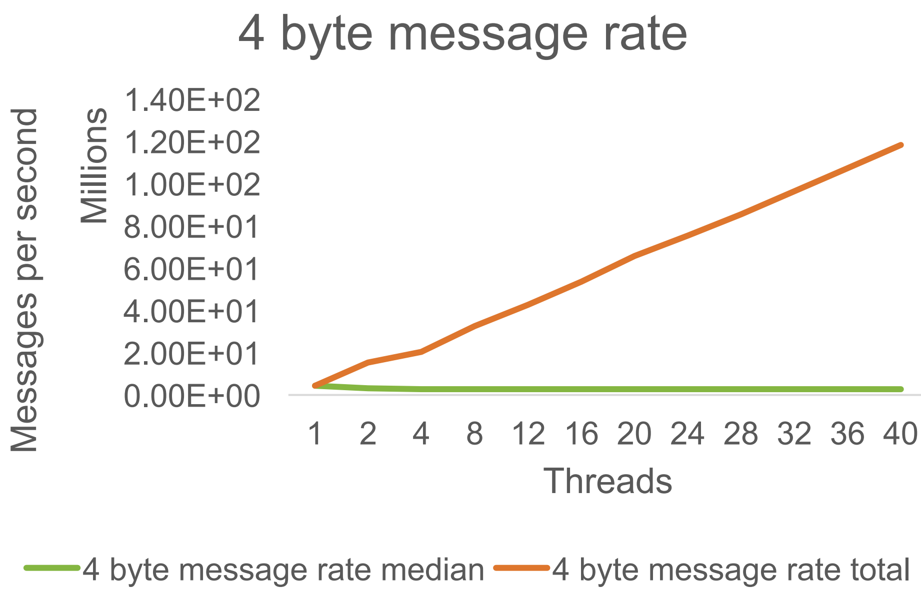

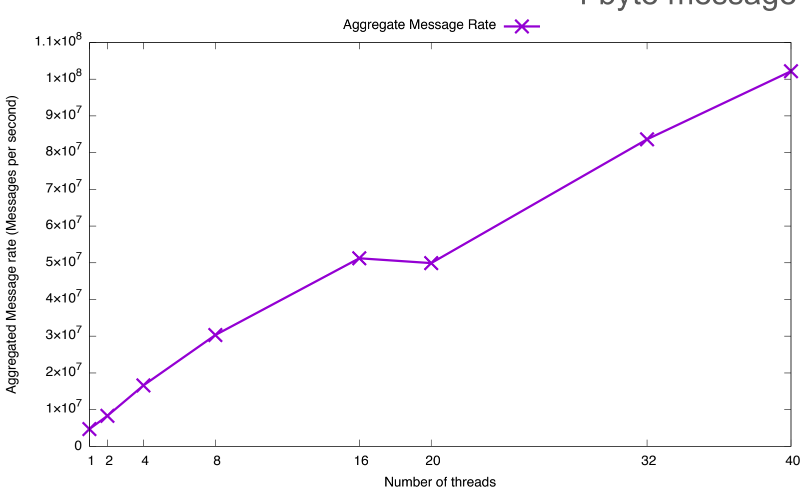

Message rate with Merged Handles

Computational Research & Development Programs

Oak Ridge National Laboratory

OpenSHMEM Benchmarks Developed

Developed OpenSHMEM benchmarks

- Multithreaded Latency and Bandwidth
- Multithreaded SSCA#1
- Multithreaded and Context-based GUPs
- Multithreaded Graph 500
- OpenSHMEM based SSSP (Bellman-Ford and Dijkstra
- OpenSHMEM based YCSB

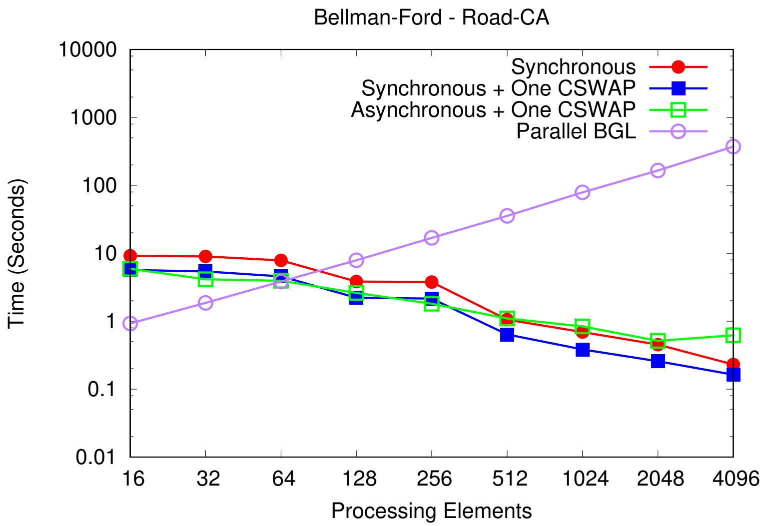

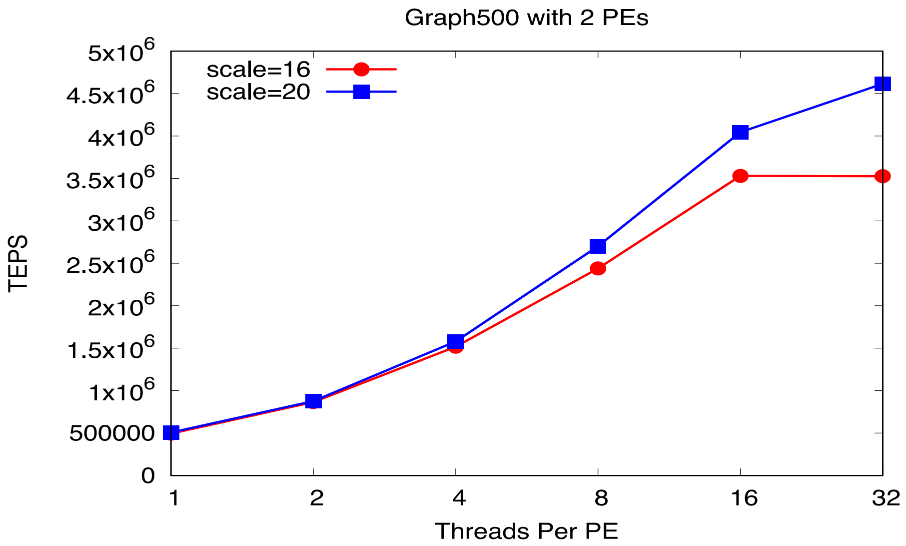

Acknowledgements

- Research sponsored by the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory, managed by UT-Battelle, LLC, for the U. S. Department of Energy.

- Department of Defense