HPC ADVISORY COUNCIL

NAMD
Performance Benchmark and Profiling

April 2017

Note

- **The following research was performed under the HPC Advisory Council activities**
  - Compute resource - HPC Advisory Council Cluster Center

- **The following was done to provide best practices**
  - NAMD performance overview
  - Understanding NAMD communication patterns
  - Ways to increase NAMD productivity

- **For more info please refer to**
  - http://www.ks.uiuc.edu/Research/namd/

NAMD

- A parallel molecular dynamics code that received the 2002 Gordon Bell Award
- Designed for high-performance simulation of large biomolecular systems
  - **Scales to hundreds of processors and millions of atoms**
- Developed by the joint collaboration of the Theoretical and Computational Biophysics Group (TCB) and the Parallel Programming Laboratory (PPL) at the University of Illinois at Urbana-Champaign
- NAMD is distributed free of charge with source code

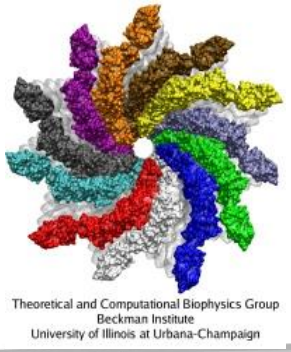

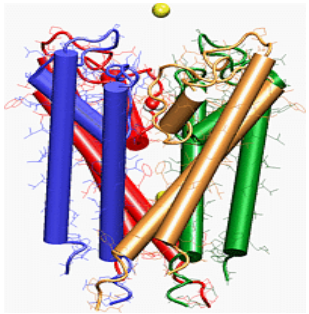

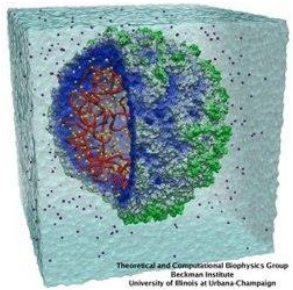

Objectives

- **The presented research was done to provide best practices**
  - NAMD performance benchmarking
    - MPI Library performance comparison
    - Interconnect performance comparison
    - CPUs comparison
    - Optimization tuning
- **The presented results will demonstrate**
  - The scalability of the compute environment/application
  - Considerations for higher productivity and efficiency

Test Cluster Configuration

- **HPE ProLiant DL360 Gen9 128-node (4096-core) "Hercules" cluster**
  - Dual-Socket 16-Core Intel E5-2697A v4 @ 2.60 GHz CPUs
  - Memory: 256GB memory, DDR4 2400 MHz, Memory Snoop Mode in BIOS sets to Home Snoop
  - OS: RHEL 7.2, MLNX_OFED_LINUX-3.4-2.0.0.0 InfiniBand SW stack
- **Mellanox ConnectX-4 EDR 100Gb/s InfiniBand Adapters**
- **Mellanox Switch-IB SB7800 36-port EDR 100Gb/s InfiniBand Switch**
- **Intel® Omni-Path Host Fabric Interface (HFI) 100Gb/s Adapter**
- **Intel® Omni-Path Edge Switch 100 Series**
- **MPI: Intel MPI 2017, Open MPI 2.02**
- **Application: NAMD 2016.2 and 5.1.2**
- **Benchmarks:**
  - Benchmark datasets: lignocellulose3M_rff
  - http://www.prace-i.eu/UEABS/NAMD/1.2/NAMD_TestCaseB.tar.gz

NAMD Performance – MPI Libraries

- **HPC-X demonstrates higher performance than Intel MPI**
  - About 79% better performance seen than Intel MPI OFA provider on NAMD

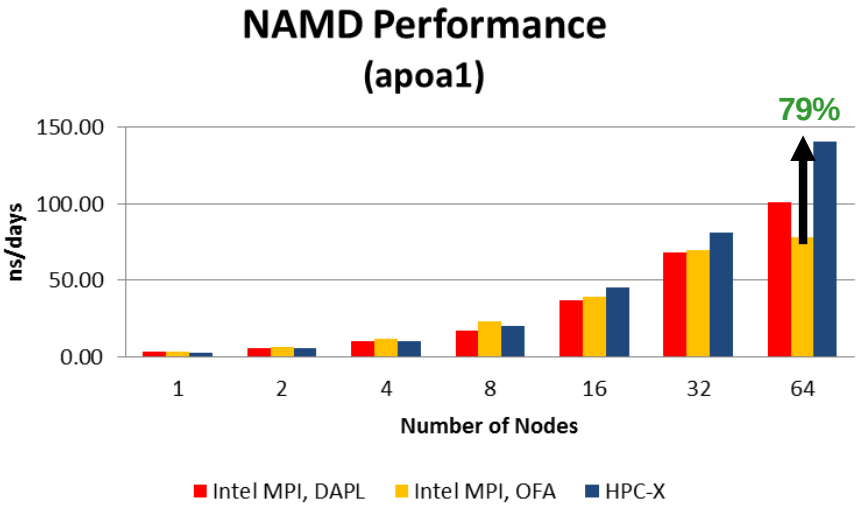

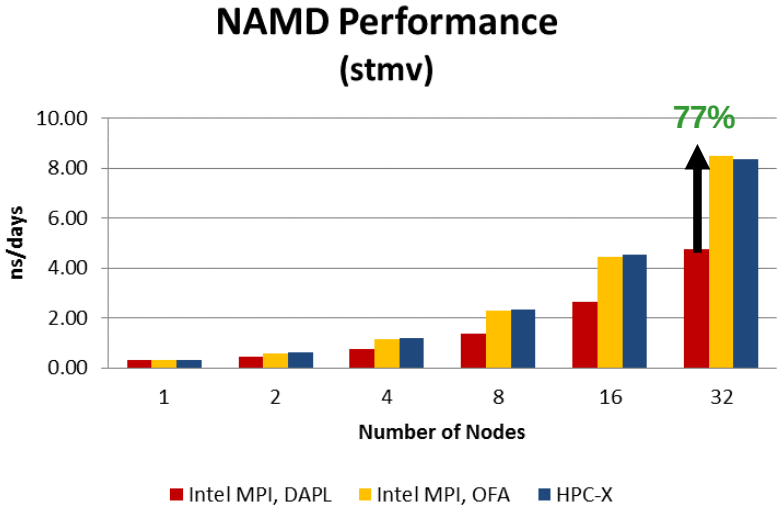

Higher is better

Optimized parameters used

NAMD Performance – Interconnects

- **EDR InfiniBand enables higher scalability than Omni-Path for NAMD**
  - InfiniBand delivers 167% better scaling versus Omni-Path on 64 nodes cluster
  - 32-nodes IB cluster delivers 55% higher performance versus 64-nodes OPA cluster

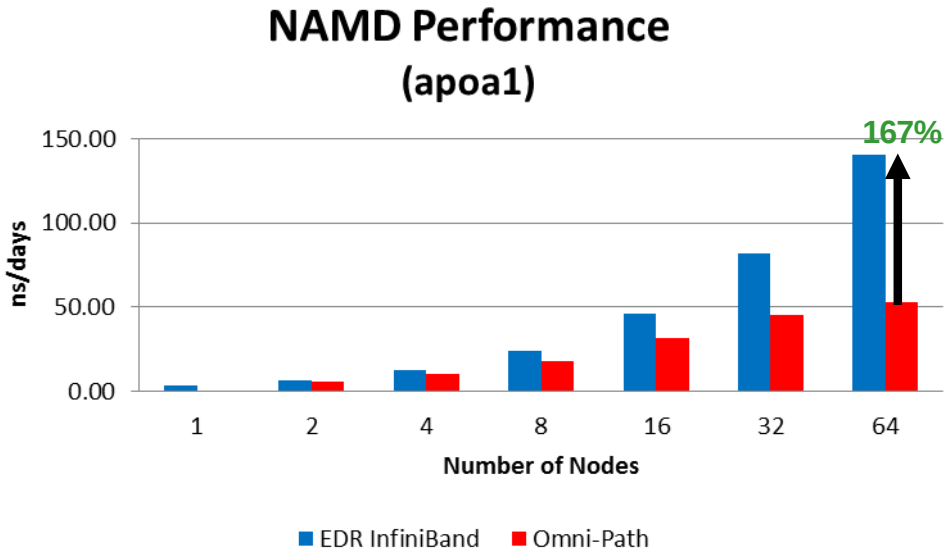

Higher is better

tel MPI

NAMD Profiling – % of MPI Calls

- **For the most time consuming MPI calls (as % of MPI time):**
  - MPI_Iprobe (90%), MPI_Isend (4%), MPI_Test (3%), MPI_Recv (3%)

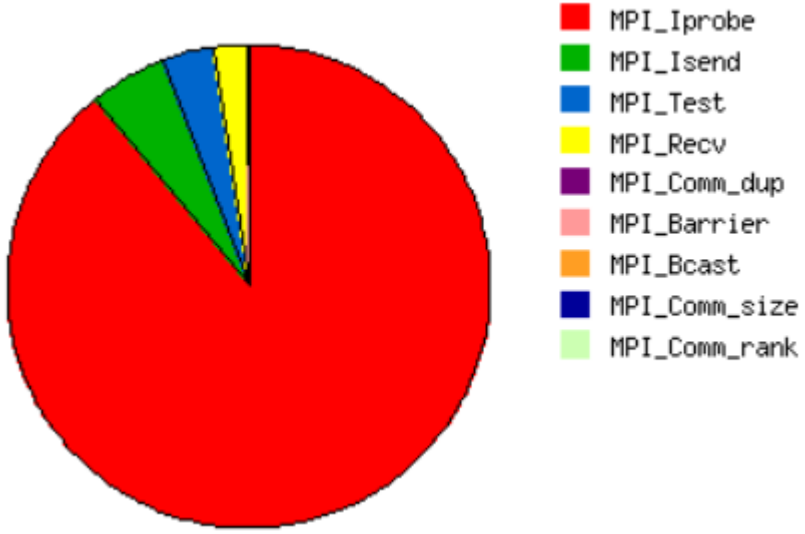

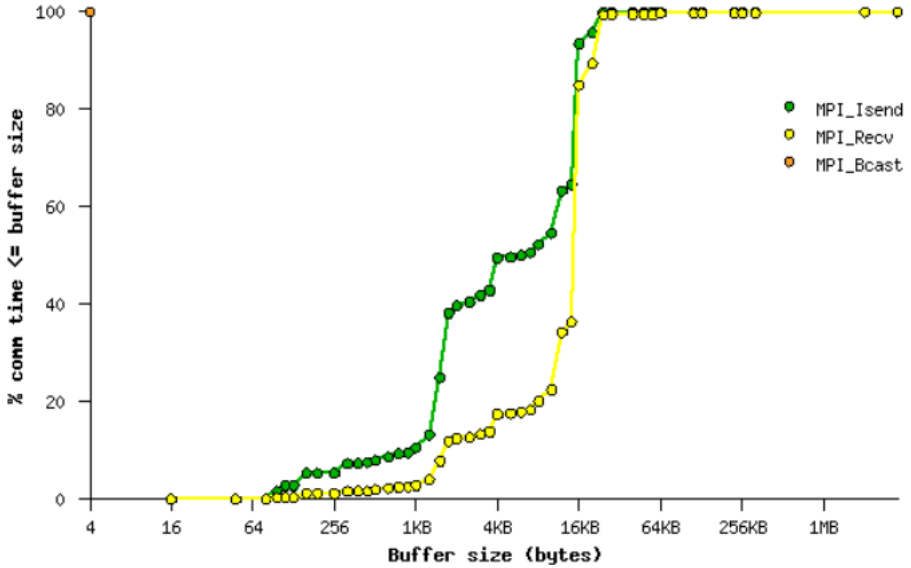

64 Nodes / 2048 Processes

NAMD Summary

- **HPC-X demonstrates higher performance than Intel MPI**
  - About 79% better performance seen than Intel MPI OFA provider on NAMD
- **EDR InfiniBand enables higher scalability than Omni-Path for NAMD**
  - InfiniBand delivers 167% better scaling versus Omni-Path on 64 nodes cluster
  - 32-nodes IB cluster delivers 55% higher performance versus 64-nodes OPA cluster
- **For the most time consuming MPI calls (as % of MPI time):**
  - MPI_Iprobe (90%), MPI_Isend (4%), MPI_Test (3%), MPI_Recv (3%)

Thank You

HPC Advisory Council