

LAMMPS benchmarking on 4 nodes cluster with Intel Xeon Phi 7120P Coprocessors

by Ashish Kumar Singh

This blog explores the application performance analysis of LAMMPS on a cluster of PowerEdge R730 servers with Intel Xeon Phi 7120Ps. All the runs were carried out with Hyper Threading (logical processors) disabled.

LAMMPS (Large Scale Atomic/Molecular Massively Parallel Simulator) is a classical molecular dynamics code, capable of doing simulation for solid-state materials (metals, semi-conductors), soft matter (biomolecules, polymers) and coarse-grained or mesoscopic systems. It can be used to model atoms or more generically as a parallel particle simulator at the atomic, meso or continuum scale.

Test Cluster Configuration:

The test cluster consisted of four PowerEdge R730 servers with two Intel Xeon Phi 7120P co-processors each. Each PowerEdge R730 had two Intel Xeon E5-2695v3 @ 2.3GHz CPU and eight 16GB DIMMS of 2133MHz making it a total of 128GB of memory. Each PowerEdge R730 consisted of one Mellanox FDR Infiniband HCA card in the low-profile x8 PCIe Gen3 slot (Linked with CPU2).

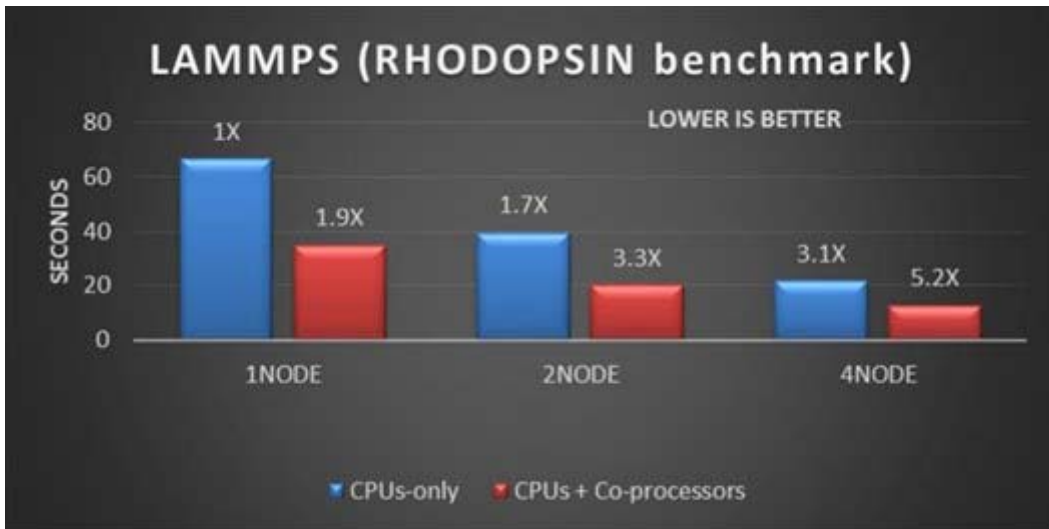
Compute node configuration

Component	Value
<i>Server</i>	<i>R730</i>
<i>Processor</i>	<i>Intel Xeon E5-2695v3 @2.3GHz</i>
<i>Memory</i>	<i>8x 16GB @2133MHz</i>
<i>Infiniband</i>	<i>Mellanox Connect-X3 FDR Adaptor (CX354A)</i>
<i>Cluster Size</i>	<i>4 Servers</i>
<i>OS</i>	<i>RHEL6.5</i>
<i>MPSS version</i>	<i>3.4.1</i>
<i>Coprocessor</i>	<i>2 x 7120P</i>
<i>Intel Compiler</i>	<i>Intel Parallel Studio XE 2015</i>
<i>LAMMPS</i>	<i>Stable version 5-sep-14</i>

The BIOS options selected for this blog were as below:

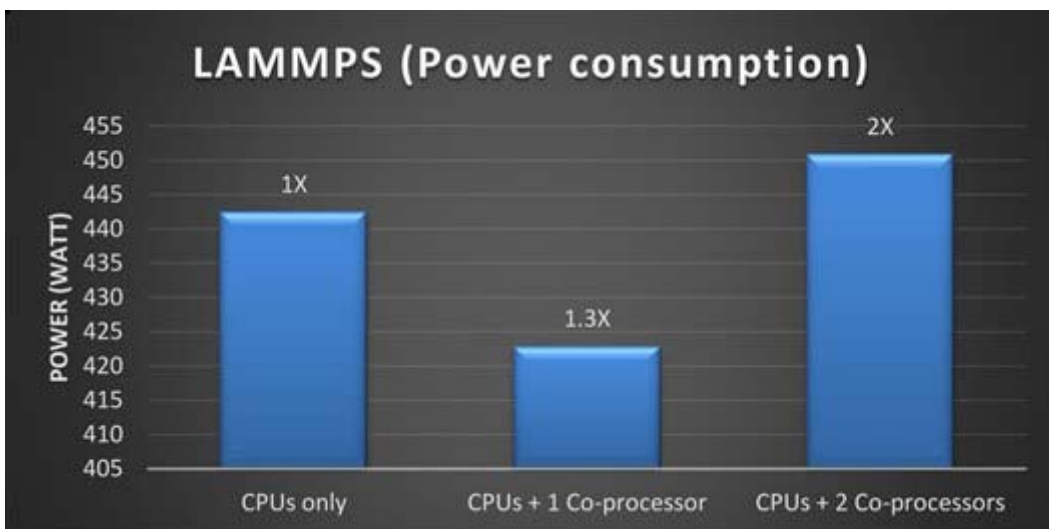
System BIOS Options	Settings
<i>Memory Settings > Snoop Mode</i>	<i>Early Snoop</i>
<i>Processor Settings > Logical Processor</i>	<i>Disabled</i>
<i>Processor Settings > QPI Speed</i>	<i>Maximum Data Rate</i>
<i>Processor Settings > Configurable TDP</i>	<i>Nominal</i>
<i>System Profile Settings > System Profile</i>	<i>Performance</i>

LAMMPS was run for Rhodopsin benchmark. Rhodopsin benchmark simulates the movement of protein in the retina which in turn plays an important role in the perception of light. The protein is solvated lipid bilayer using the CHARMM force field with particle-particle particle-mesh long-range electrostatics and SHAKE constraints. The simulation was performed with 2,048,000 atoms at the temperature of 300K and pressure of 1 atm. The results for single node, two nodes and four nodes are as shown below. On one node with CPU only configuration, the loop-time was 66.5 seconds, while configuration of CPUs and two Intel Xeon Phi 7120Ps had a loop-time of 34.8 seconds. This demonstrated a performance increase of 1.9X. In comparison to CPUs only, CPUs + co-processors from one node to four nodes showed performance increase of 5.2X.



LAMMPS (RHODOPSIN Benchmark)			
	1NODE	2NODE	4NODE
CPUs-only (Seconds)	66.5	39.2	21.5
CPUs + Coprocessors (Seconds)	34.8	20	12.7

The LAMMPS power consumption analysis with RHODOPSIN benchmark is shown below. On single node, the power consumption by a CPU-only configuration was 442.4 watts, while configuration with CPUs and one co-processor consumed around 423W and subsequently configuration with CPUs and two co-processors consumed 450.8W.



Power Consumption (Watt)		
CPUs only	CPUs + 1 Coprocessor	CPUs + 2 Coprocessors
442.4	423	450.8

Performance Per watt		
CPUs only	CPUs + 1 Coprocessor	CPUs + 2 Coprocessors
0.15	0.11	0.077

All the LAMMPS runs on co-processors used the auto-balance mode. The performance per watt demonstrated 2-fold increase with CPUs + 2 co-processors than CPUs only.

Conclusion:

The Intel Xeon Phi 7120Ps cluster with Dell PowerEdge R730 showed sustained performance increase of 2-fold. The power-efficiency was increased by 2X with two Intel Xeon Phi 7120Ps in comparison to CPUs only, resulting in a powerful, energy-efficient HPC platform.