HPC ADVISORY COUNCIL

LS-DYNA
Performance Benchmark and Profiling

October 2017

Note

- **The following research was performed under the HPC Advisory Council activities**
  - Participating vendors: LSTC, Huawei, Mellanox
  - Compute resource - HPC Advisory Council Cluster Center

- **The following was done to provide best practices**
  - LS-DYNA performance overview
  - Understanding LS-DYNA communication patterns
  - Ways to increase LS-DYNA productivity
  - MPI libraries comparisons

- **For more info please refer to**
  - http://www.lstc.com
  - http://www.huawei.com
  - http://www.mellanox.com

LS-DYNA

- **LS-DYNA**
  - A general purpose structural and fluid analysis simulation software package capable of simulating complex real world problems
  - Developed by the Livermore Software Technology Corporation (LSTC)
- **LS-DYNA used by**
  - Automobile
  - Aerospace
  - Construction
  - Military
  - Manufacturing
  - Bioengineering

Objectives

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

Test Cluster Configuration

- **Huawei FusionServer E9000 with FusionServer CH121 V5 16-node (640-core) “Skylake” cluster**
  - Dual-Socket 20-Core Intel Xeon Gold 6138 @ 2.00 GHz CPUs
  - Memory: 192GB memory, DDR4 2666 MHz RDIMMs per node
  - OS: RHEL 7.2, MLNX_OFED_LINUX-4.1-1.0.2.0 InfiniBand SW stack
- **Mellanox ConnectX-5 EDR 100Gb/s InfiniBand Adapters**
- **Mellanox Switch-IB SB7800 36-port EDR 100Gb/s InfiniBand Switch**
- **Compilers: Intel Parallel Studio XE 2018**
- **MPI: Intel MPI 2018, Mellanox HPC-X MPI Toolkit v1.9.7, Platform MPI 9.1.4.3**
- **Application: MPP LS-DYNA R9.1.0, build 113698, single precision**
- **MPI Profiler: IPM (from Mellanox HPC-X)**
- **Benchmarks: TopCrunch benchmarks**
  - Neon Refined Revised (neon_refined_revised), Three Vehicle Collision (3cars), NCAC Minivan Model (Caravan2m-ver10), odb10m (NCAC Taurus model)

Introducing Huawei FusionServer E9000 (CH121) V5

High-Performance 2-Socket Blade Unlocks Supreme Computing Power

FusionServer

HUAWEI

Full-series Intel® Xeon® Scalable Processors, **24** DDR4 DIMMs, **AEP memory supported, 1** PCIe slot, 2 SFF/2 NVMe SSDs/4 **M.2 SSDs high-performance storage, multi-plane network**, LOM supported

LS-DYNA Performance – CPU SKUs and Generation

- **LS-DYNA performance gain by larger core counts and better memory throughput**
  - The “Gold 6140” demonstrates a 50% of performance gain (29% more cores) vs E5-2680v4
  - The “Gold 6148” demonstrates a 61% of performance gain (42% more cores) vs E5-2680v4
  - Base clock are the same on E5-2680 v4 and Gold 6148, while Gold 6140 runs slightly slower
  - Skylake supports 6 memory channels and faster DIMMs which impacts on memory performance

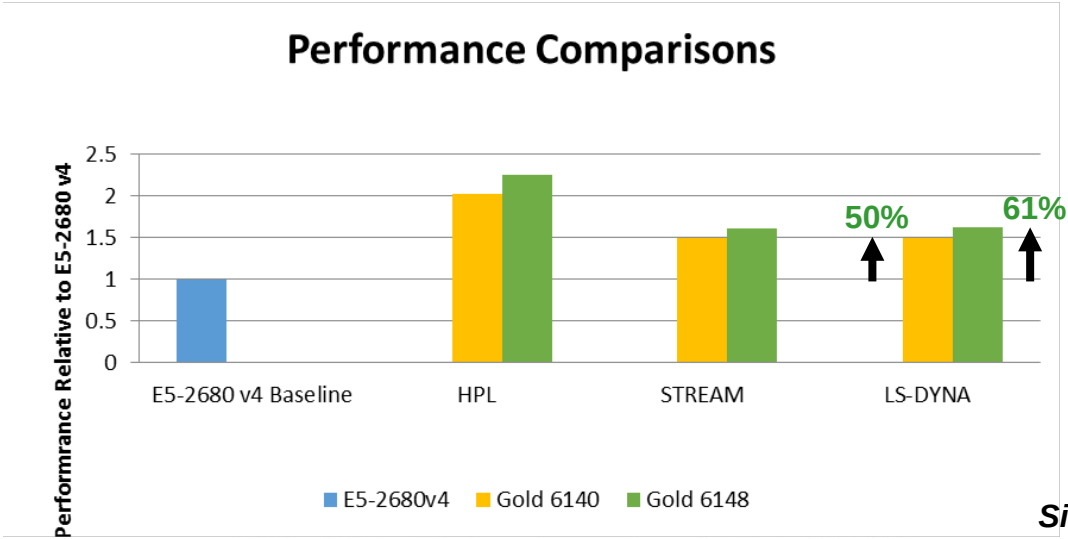

Higher is better

Single Node Performance

LS-DYNA Performance – Memory Speed

- **Memory speed provides some benefits to LS-DYNA performance**
  - Skylake platform supports DIMM speed up to 2666MHz DIMMs
  - 2666MHz DIMMs is theoretically ~11% faster than the 2400MHz DIMMs
  - LS-DYNA reports only about ~2-3% of the improvement on a single node
  - It appears only part of the speed difference is translated into LS-DYNA performance gain

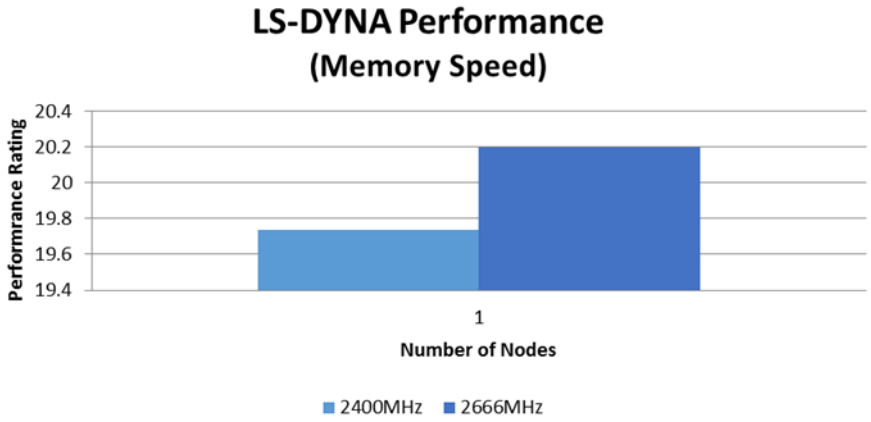

Higher is better

40 MPI Processes / Node

LS-DYNA Performance – Sub-NUMA Clustering

- **Enabling SNC provides some benefits for LS-DYNA**
  - Sub-NUMA Clustering (SNC) is similar to a cluster-on-die (COD) in Haswell/Broadwell generation
  - CPU cores and memory would be split into 2 separate NUMA domains when SNC is enabled
  - SNC generally should demonstrate some benefits for applications that requires good NUMA locality
  - SNC demonstrates a performance gain of ~2-3% on a single node basis

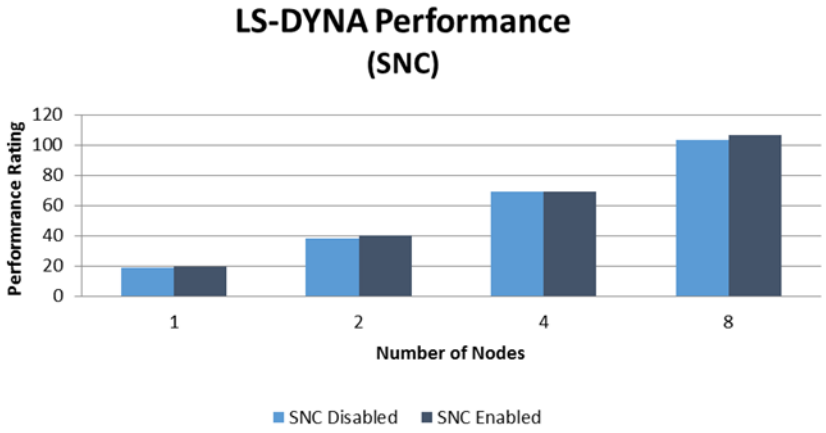

Higher is better

40 MPI Processes / Node

LS-DYNA Performance – CPU Instructions

- **AVX2 outperforms both AVX-512 and SSE2 executables on Skylake CPU**
  - Performance gain of 17% by using AVX2 over AVX-512 executables
  - AVX-512 performs worse compared to AVX2, despite improved vectorization
  - AVX-512 instructions runs at a reduced clock frequency as AVX2 and normal clocks
  - Benefit of AVX2 appears to be larger on bigger dataset (such as car2car)

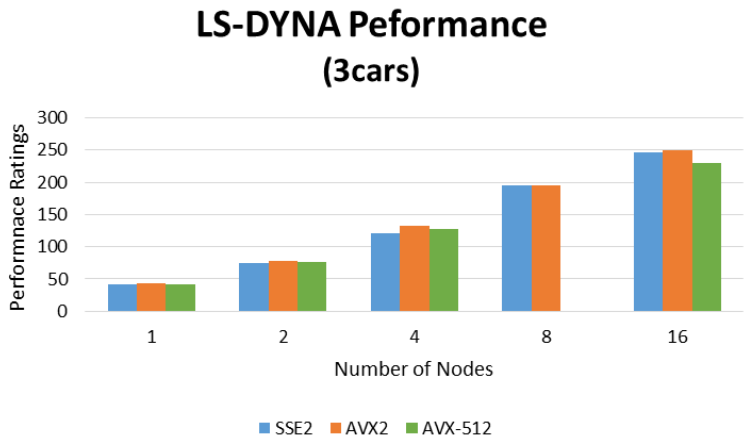

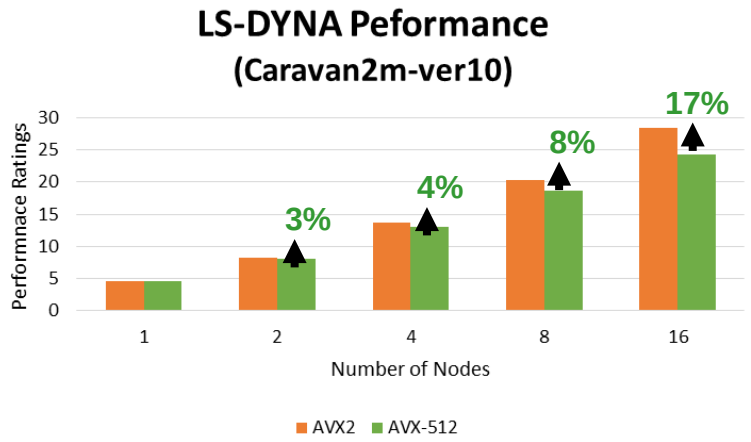

Higher is better

40 MPI Processes / Node

LS-DYNA Performance – CPU Instruction Sets

- **Some variance in performance among different LS-DYNA versions/executables**
  - AVX2 performs better than SSE2 LS-DYNA executables
  - Small variance in performance among different LS-DYNA releases
  - R7.1.3 appeared to perform better on larger datasets

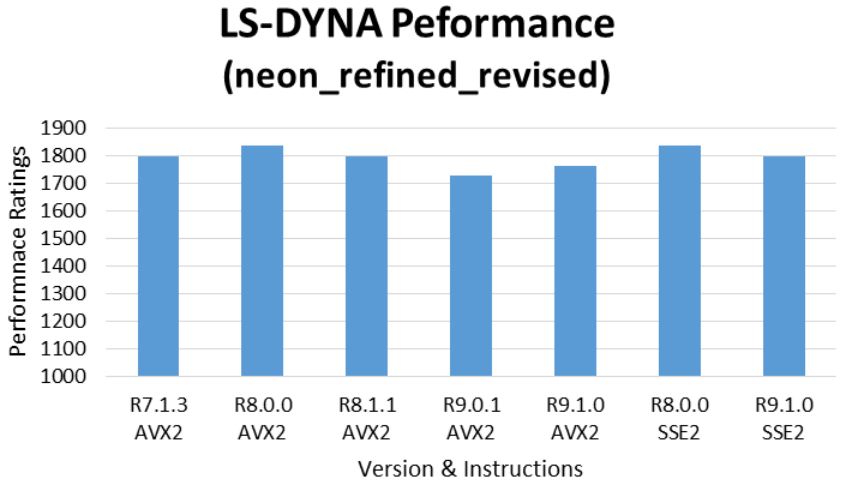

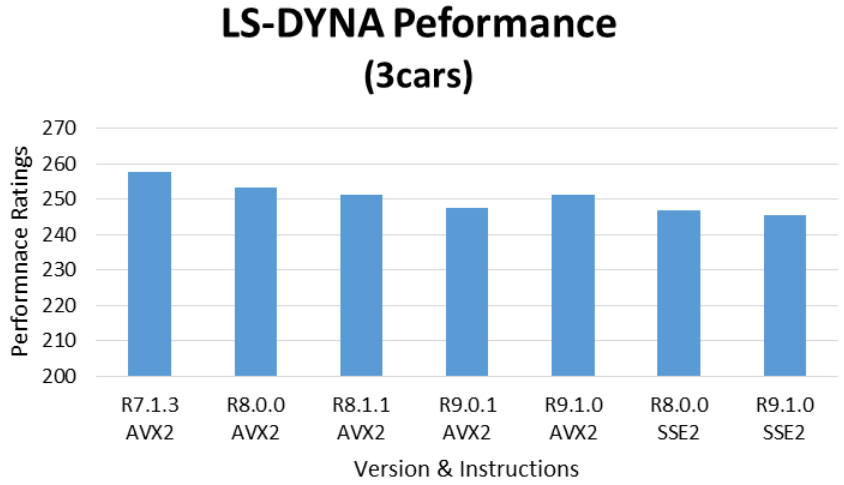

Higher is better

40 MPI Processes / Node

LS-DYNA Performance – MPI Libraries

- **All three MPI implementations shows decent performance at scale**
  - Platform MPI and HPC-X performs similarly, while Intel MPI shows a drop at small dataset at scale

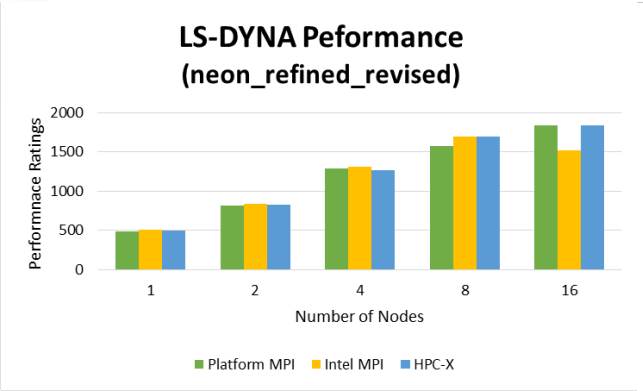

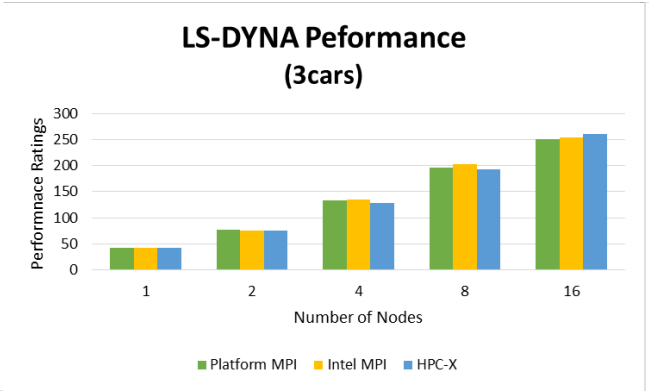

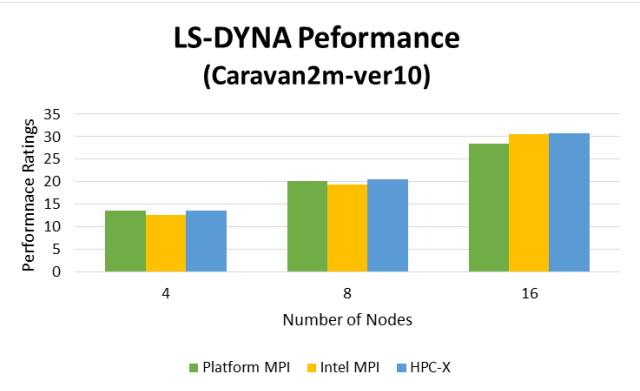

Higher is better

40 MPI Processes / Node

LS-DYNA Performance – System Generations

- **Current Skylake system configuration outperforms prior system generations**
  - Skylake platform outperformed Broadwell by 21%, Haswell by 51%, Ivy Bridge by 89%, Sandy Bridge by 132%, Westmere by 222%, Nehalem by 425%
  - Skylake performs 41% better than Broadwell for the 3cars model on a single-node basis
  - System components used:
    - Skylake: 2-socket 20-core Xeon Gold 6138 2.0GHz, 2666MHz DIMMs, ConnectX-5 EDR InfiniBand
    - Broadwell: 2-socket 14-core Xeon E5-2690v4 2.6GHz, 2400MHz DIMMs, ConnectX-4 EDR InfiniBand
    - Haswell: 2-socket 14-core Xeon E5-2697v3 2.6GHz, 2133MHz DIMMs, ConnectX-4 EDR InfiniBand
    - Ivy Bridge: 2-socket 10-core Xeon E5-2680v2 2.8GHz, 1600MHz DIMMs, Connect-IB FDR InfiniBand
    - Sandy Bridge: 2-socket 8-core Xeon E5-2680 2.7GHz, 1600MHz DIMMs, ConnectX-3 FDR InfiniBand
    - Westmere: 2-socket 6-core Xeon x5670 2.93GHz, 1333MHz DIMMs, ConnectX-2 QDR InfiniBand
    - Nehalem: 2-socket 4-core Xeon x5570 2.93GHz, 1333MHz DIMMs, ConnectX-2 QDR InfiniBand

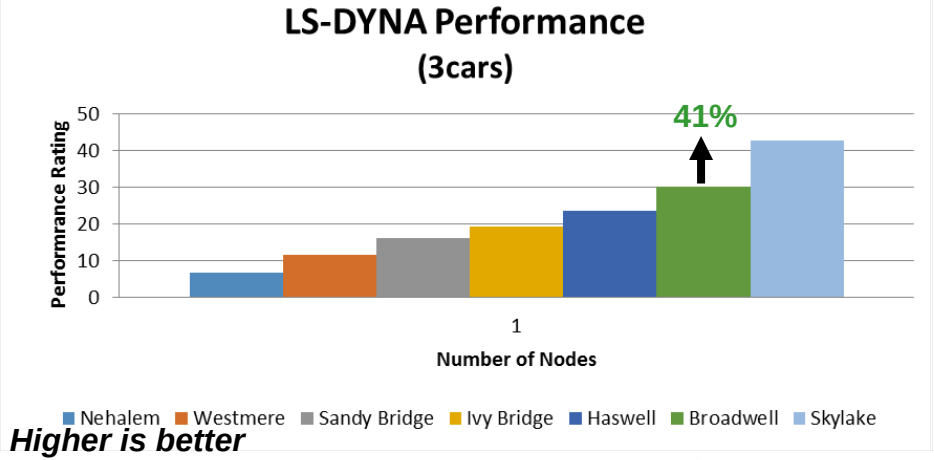

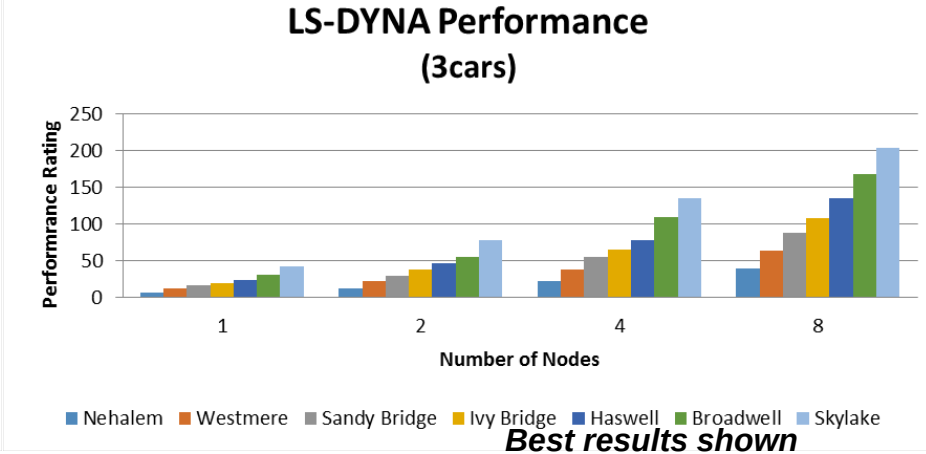

Higher is better

Best results shown

LS-DYNA Summary

- **LS-DYNA is multi-purpose explicit and implicit finite element program**
  - Utilizes both compute, memory and network communications for performance
- **Effect of MPI on performance**
  - Platform MPI and HPC-X performs similarly, Intel MPI shows a drop at small dataset
- **Effect of Skylake generation on performance**
  - Provides substantial performance gain due to the larger core count, support for memory channels
  - Faster 2666MHz DIMM (compares to 2400MHz) translates to increase 2-3% in higher performance
- **Effort of CPU Instructions on performance**
  - AVX-512 performs worse compared to AVX2, despite the improved vectorization
  - AVX-512 instructions runs at a reduced clock frequency as AVX2 and normal clocks
- **Effect of SNC on performance**
  - Enabling Sub-NUMA Clustering provides small advantage (~2-3%) on single node
- **Effect fo LS-DYNA version on performance**
  - Small variance in performance among different LS-DYNA releases; best appeared to be R7.1.3

Thank You

HPC Advisory Council