

LS-DYNA Performance Benchmark and Profiling

August 2010







- **The following research was performed under the HPC Advisory Council activities**
 - Participating vendors: AMD, Dell, Mellanox, LSTC
 - Compute resource - HPC Advisory Council Cluster Center
- **For more info please refer to**
 - www.mellanox.com, www.dell.com/hpc, www.amd.com,
www.lstc.com

- **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



- **LS-DYNA SMP (Shared Memory Processing)**
 - Optimize the power of multiple CPUs within single machine
- **LS-DYNA MPP (Massively Parallel Processing)**
 - The MPP version of LS-DYNA allows to run LS-DYNA solver over High-performance computing cluster
 - Uses message passing (MPI) to obtain parallelism
- **Many companies are switching from SMP to MPP**
 - For cost-effective scaling and performance



- **The presented research was done to provide best practices**
 - LS-DYNA performance benchmarking
 - MPI Library performance comparison
 - Interconnect performance comparison
 - Core/Socket performance comparison

- **The presented results will demonstrate**
 - Balanced system to provide good application scalability

- **Dell™ PowerEdge™ R815 6-node cluster**
- **AMD™ Opteron™ 6174 (code name “Magny-Cours”) 12-cores @ 2.2 GHz CPUs**
- **4 CPUs per server node**
- **Mellanox ConnectX-2 40Gb/s QDR InfiniBand adapter**
- **Mellanox IS3035 36-Port 40Gb/s QDR InfiniBand switch**
- **Memory: 128GB memory per node DDR3 1333**
- **OS: CentOS 5.5, MLNX-OFED 1.5.1 InfiniBand SW stack**
- **MPI: Open MPI 1.4.1, Platform MPI 7.1, MVAPICH-1.2.0**
- **Application: LS-DYNA MPP971_s_R4.2.1**
- **Benchmark Workload**
 - Three Vehicle Collision Test simulation

- **HPC Advisory Council Test-bed System**
- **New 6-node 288 core cluster - featuring Dell PowerEdge™ R815 servers**
 - Replacement system for Dell PowerEdge SC1435 (192 cores) cluster system following 2 years of rigorous benchmarking and product EOL
 - System to be redirected to explore HPC in the Cloud applications
- **Workload profiling and benchmarking**
 - Characterization for HPC and compute intense environments
 - Optimization for scale, sizing and configuration and workload performance
 - Test-bed Benchmarks
 - RFPs
 - Customers/Prospects, etc
 - ISV & Industry standard application characterization
 - Best practices & usage analysis



About Dell PowerEdge™ Platform Advantages

Best of breed technologies and partners

Combination of AMD™ Opteron™ 6100 series platform and Mellanox ConnectX InfiniBand on Dell HPC

Solutions provide the ultimate platform for speed and scale

- Dell PowerEdge R815 system delivers 4 socket performance in dense 2U form factor
- Up to 48 core/32DIMMs per server – 1008 core in 42U enclosure

Integrated stacks designed to deliver the best price/performance/watt

- 2x more memory and processing power in half of the space
- Energy optimized low flow fans, improved power supplies and dual SD modules

Optimized for long-term capital and operating investment protection

- System expansion
- Component upgrades and feature releases




About AMD Opteron™ 6000 Series Platform

Balanced and Scalable - Designed to Support More Cores
Socket compatibility with planned AMD Opteron™ 6200 Series processors.




- Twelve-core and Eight-core, L2 – 512K/core, L3 - 12MB of shared L3 Cache
- AMD CoolCore™ technology, AMD PowerNow!™ technology, Enhanced C1 state, AMD CoolSpeed technology, APML (*in APML enabled platforms*)
- Quad-Channel U/RDDR3 & LV U/RDDR3, ECC, support for on-line spare memory
- Supports up to 3 DIMMs/channel, up to 12 DIMMS per CPU
- Balanced performance with power/thermals
- Quad 16-bit HyperTransport™ 3 technology (HT3) links, up to 6.4 GT/s per link
- AMD SR56x0 chipset with I/O Virtualization and PCIe® 2.0



- **Input Dataset**
 - Three Vehicle Collision Test
- **All MPIs provide same level of performance**
 - Platform MPI is slightly better than others with 192 cores or less

LS-DYNA Benchmark (3 Vehicle Collision)




Higher is better

48-cores per node

- **Core/socket comparison**

- Using 6 cores/socket (24 cores/node) provides higher performance in all tested cases
- Up to 15% higher performance than 48 cores/node

LS-DYNA Benchmark (3 Vehicle Collision)




Higher is better

InfiniBand QDR based

- **AMD “Magny-Cours” enables big performance gain versus previous generations**
 - 87% higher than Barcelona and 51% than Shanghai with 24 CPUs
 - Barcelona /Shanghai with IB DDR and PCIe Gen1, Magny-Cours with IB QDR and PCIe Gen2


LS-DYNA Benchmark (3 Vehicle Collision)



Higher is better

- **13% performance difference measured between 2.2GHz and 1.9GHz CPU frequencies**
 - Performance gap is relative consistent across different cluster sizes
 - Increasing the CPU frequency increases the performance in the same order

LS-DYNA Benchmark (3 Vehicle Collision)

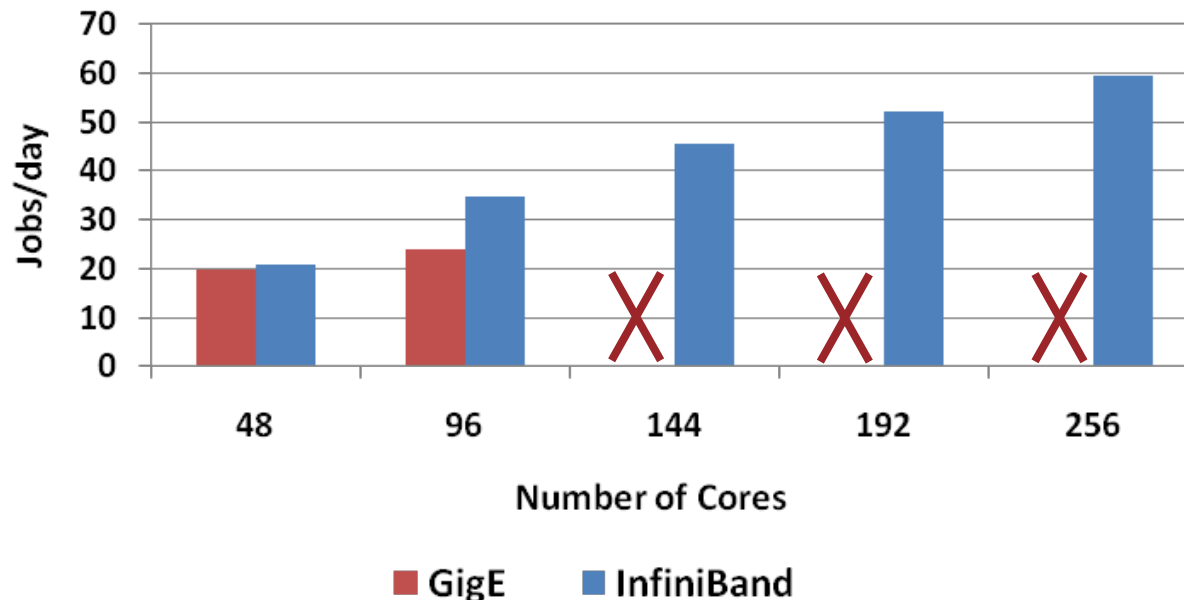


Higher is better

InfiniBand QDR based

- **InfiniBand enables better application performance and scalability**
 - GigE stops scaling after 2 nodes
 - Job runs extremely slow with GigE with 3+ nodes – require more than a day per single job -X
- **Application performance over InfiniBand scales as cluster size increases**

LS-DYNA Benchmark (3 Vehicle Collision)



Higher is better

48-cores per node

- **Application Scalability and efficiency relies on**
 - Balanced components that match each other capabilities
 - CPU, memory, and Interconnect
- **Performance Optimizations**
 - MPI libraries showed comparable performance overall
 - Reducing core utilization/access (from 48 to 24 core) can improve LS-DYNA performance due to higher available memory bandwidth
- **Interconnect Characterization**
 - InfiniBand delivers superior scalability as cluster size increases
 - GigE can't scale beyond 2 nodes
- **CPU Performance Impact**
 - “Magny-Cours” provides much performance versus Barcelona and Shanghai
 - Increasing the CPU frequency increases the performance in the same order

Thank You

HPC Advisory Council



LSTC
Livermore Software
Technology Corp.



All trademarks are property of their respective owners. All information is provided "As-Is" without any kind of warranty. The HPC Advisory Council makes no representation to the accuracy and completeness of the information contained herein. HPC Advisory Council Mellanox undertakes no duty and assumes no obligation to update or correct any information presented herein