

LS-DYNA Performance Benchmarking and Profiling

August 2010



LSTC
Livermore Software
Technology Corp.




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

- **For more info please refer to**
 - www.mellanox.com
 - <http://www.hp.com/go/hpc>
 - 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




- **The presented research was done to provide best practices**
 - LS-DYNA performance benchmarking
 - Interconnect performance comparisons
 - Ways to increase LS-DYNA productivity
 - Power-efficient simulations
- **The presented results will demonstrate**
 - The scalability of the compute environment
 - Considerations for power saving through balanced system configuration


- **HP ProLiant SL2x170z G6 16-node cluster**
 - Six-Core Intel X5670 @ 2.93 GHz CPUs
 - Memory: 24GB per node
 - OS: CentOS5U4, OFED 1.5.1 InfiniBand SW stack
- **Mellanox ConnectX-2 adapters and switches**
- **Fulcrum based 10GigE switch**
- **MPI: Platform MPI 7.1**
- **Application: LS-DYNA MPP971_s_R5.0**
- **Benchmark Workload**
 - Three Vehicle Collision Test simulation

About HP ProLiant SL6000 Scalable System

- **Solution-optimized for extreme scale out**




ProLiant z6000 chassis
Shared infrastructure
– fans, chassis, power



ProLiant SL160z G6 ProLiant SL165z G7
Large memory
-memory-cache apps



ProLiant SL170z G6
Large storage
-Web search and database apps




ProLiant SL2x170z G6
Highly dense
- HPC compute and
web front-end apps

Save on cost and energy -- per node, rack and data center

Mix and match configurations

Deploy with confidence

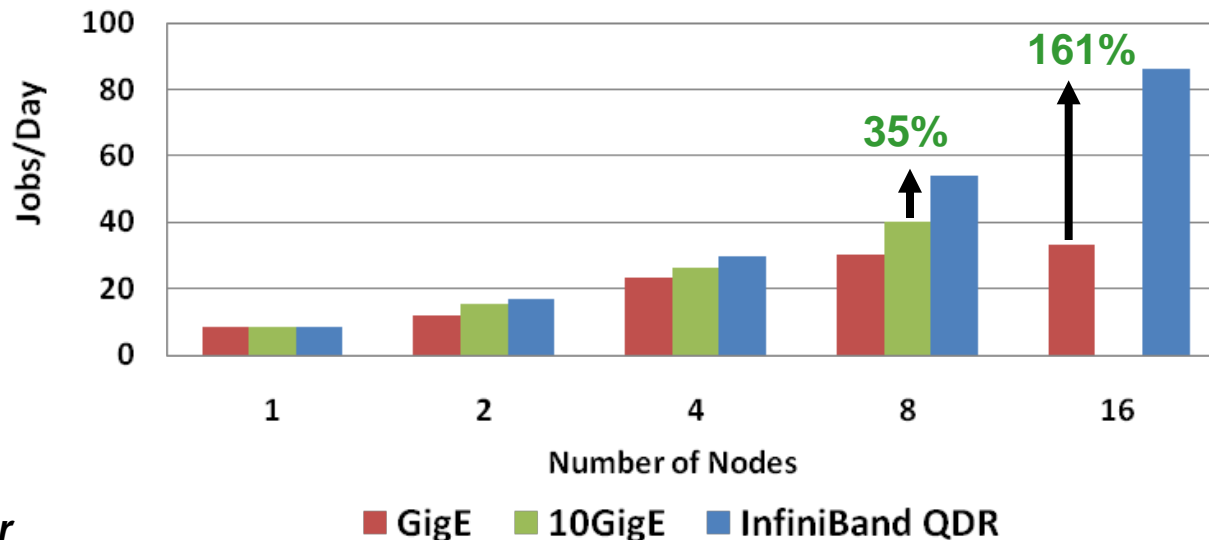


#1
Power
Efficiency*

* SPECpower_ssj2008
www.spec.org
17 June 2010, 13:28

- **Input Dataset**
 - 3 Vehicle Collision (0.8 million elements)
- **InfiniBand QDR enables higher scalability**
 - 161% higher performance than GigE at 16 nodes
 - 35% higher performance than 10GigE 8 nodes
 - Expected more performance gain as cluster size increases
- **InfiniBand reduces electrical energy/job by**
 - 61% or more compared to GigE
 - 26% or more compare to 10GigE

LS-DYNA Benchmark Result (3 Vehicle Collision)



Higher is better

12 Cores/Node


- **Input Dataset**

- 3 Vehicle Collision (0.8 million elements)

- **Core/socket comparison**

- Using 6 cores/socket (12 cores/node) provides higher performance at all node count
- Up to 14% higher performance than 8 cores/node

LS-DYNA Benchmark Result (3 Vehicle Collision)




Higher is better

InfiniBand

- **Input Dataset**
 - 3 Vehicle Collision (0.8 million elements)
- **Performance comparison**
 - 1 job mode: All cores per node used by single job
 - 2 jobs mode: Each job runs over half number of cores per socket
 - Running 2 jobs in parallel delivers 19% higher productivity than single job alone
 - Performance advantage grows as cluster size increases

LS-DYNA Benchmark Result



Higher is better

192 cores/16Nodes

LS-DYNA MPI Profiling – MPI Time

- **Input Dataset**


- 3 Vehicle Collision (0.8 million elements)

- **Runtime Distribution**


- MPI overhead becomes dominated with GigE as node count increases

- InfiniBand enables much smaller MPI communication overhead comparing to GigE

Runtime Distribution (GigE)



Runtime Distribution (InfiniBand QDR)



12 Cores/Node

- **Interconnect comparison shows**
 - InfiniBand delivers superior performance in every cluster size
 - Low latency InfiniBand enables better scalability than 10GigE and GigE
- **Customized job placement increases application productivity**
 - 19% more jobs can be completed by running 2 jobs concurrently
- **Cores/Socket Usage**
 - Fully utilizing all the cores in each node yields the best performance
- **InfiniBand QDR saves power**
 - Reduces power consumption/job by 61% or more compared to GigE and 26% or more compared to 10GigE

Thank You

HPC Advisory Council



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