

ANSYS Performance Benchmark and Profiling




July 2010



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

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

- **Computational Fluid Dynamics (CFD) is a computational technology**
- **CFD brings together a number of different disciplines**
 - Fluid dynamics, mathematical theory of partial differential systems, computational geometry, numerical analysis, Computer science
- **ANSYS FLUENT is a leading CFD application from ANSYS**
 - Widely used in almost every industry sector and manufactured product
- **ANSYS CFX is a high performance, general purpose CFD program**
 - All physical models in the ANSYS CFX solver work in parallel
- **ANSYS Mechanics offer a broad spectrum of capabilities covering**
 - Analysis types, elements, contact, materials, equation solvers and coupled physics




- **The presented research was done to provide best practices**
 - ANSYS Fluent, CFX, and Mechanical performance benchmarking
 - Interconnect performance comparisons
 - Understanding ANSYS communication patterns

- **The presented results will demonstrate**
 - The scalability of the compute environment
 - Considerations for power saving through balanced system configuration


- **HP ProLiant SL170z 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
- **Intel Cluster Ready certified cluster**
- **Mellanox ConnectX2 InfiniBand adapters and switches**
- **MPI: HP-MPI v2.0.7**
- **Application: ANSYS Fluent, CFX, Mechanical, v12.0**
- **Benchmark Workload**
 - **ANSYS Benchmark Suite**

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

ANSYS Fluent Benchmark Results

- **Dataset: Aircraft_2M**
- **InfiniBand enables better application performance and scalability**
 - Up to 135% higher performance than GigE
 - Application performance over InfiniBand scales as cluster size increases




Higher is better

12-cores per node

ANSYS Fluent Benchmark Results

- **Dataset: Truck_14M**
- **InfiniBand enables better application performance and scalability**
 - InfiniBand reduces electrical energy/job by half or more compared to GigE
 - Time to run high priority job is 4x faster with InfiniBand
 - Up to 453% higher performance than GigE
 - GigE stops scaling beyond 14 nodes




Higher is better

12-cores per node

ANSYS CFX Benchmark Results

- **Dataset: Transonic Airfoil**
- **InfiniBand enables better application performance and scalability**
 - Up to 85% higher performance than GigE
 - Application performance over InfiniBand scales as cluster size increases




Higher is better

12-cores per node

ANSYS Mechanical Benchmark Results


- **Dataset: V12cg-2**
 - Engine Block, 6.2 MDOFs
- **InfiniBand enables better application performance and scalability**
 - GigE stops scaling after 4 nodes
 - InfiniBand allows ANSYS Mechanical to scale to 96 cores



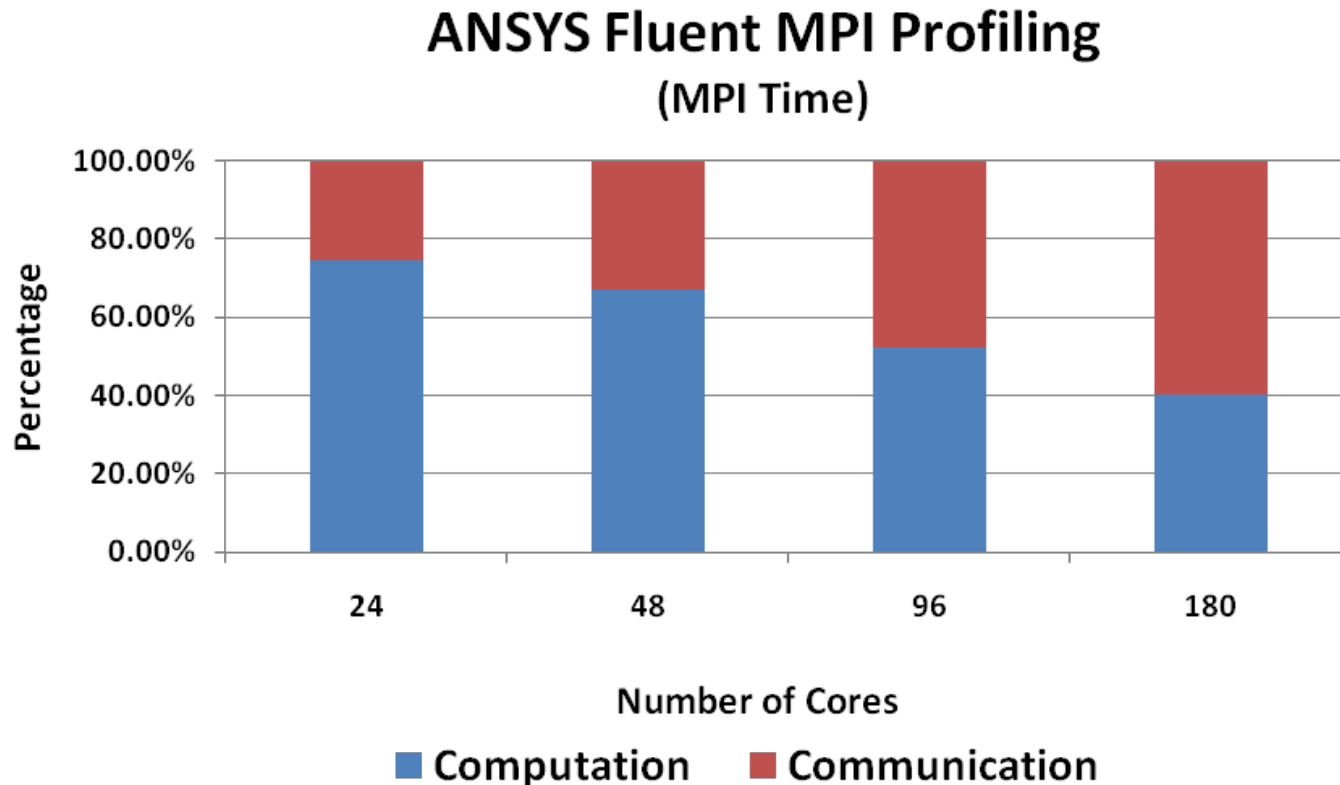
Higher is better

12-cores per node

- **Balanced system – CPU, memory, Interconnect that match each other capabilities - is essential for providing application efficiency**
- **InfiniBand delivers superior performance across all system sizes**
 - Lower power/job
 - Faster execution for high-priority jobs
 - Cost effective way for productive simulations

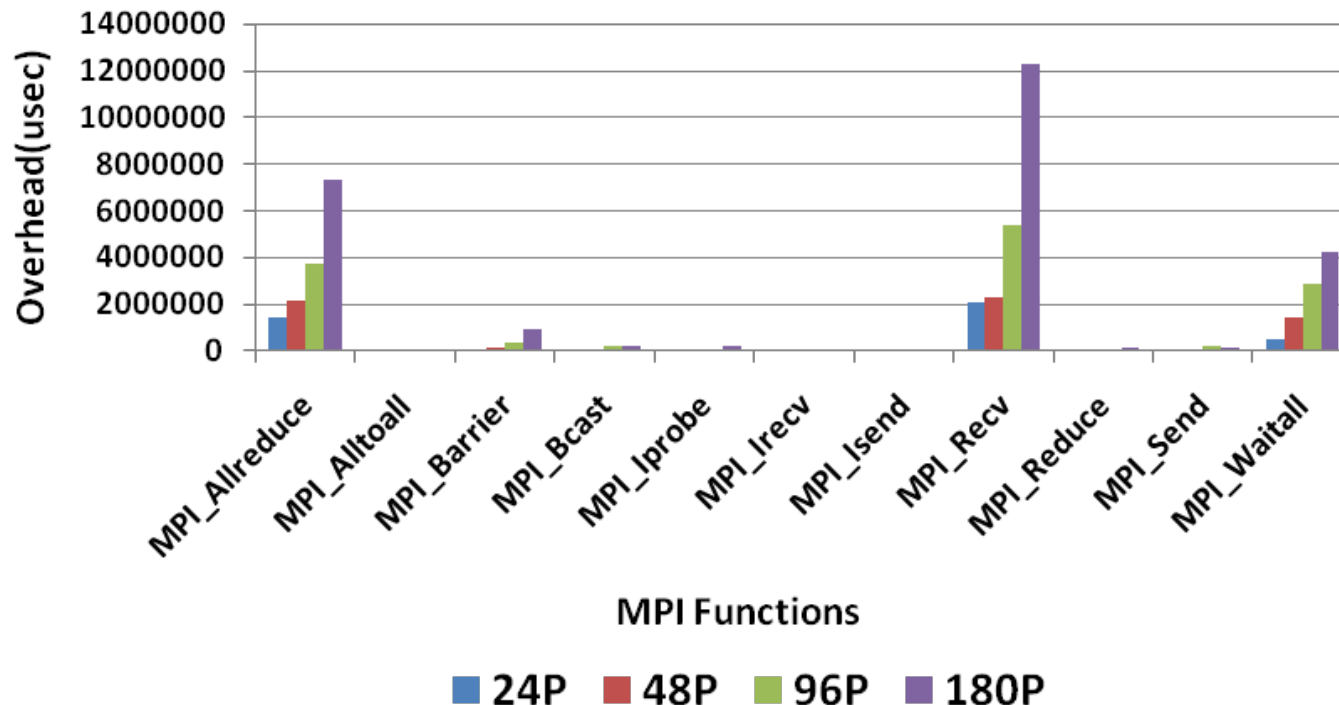


- **Percentage of MPI time increases as cluster size scales**
 - Increases to 60% at 180 cores
 - Truck_14M benchmark

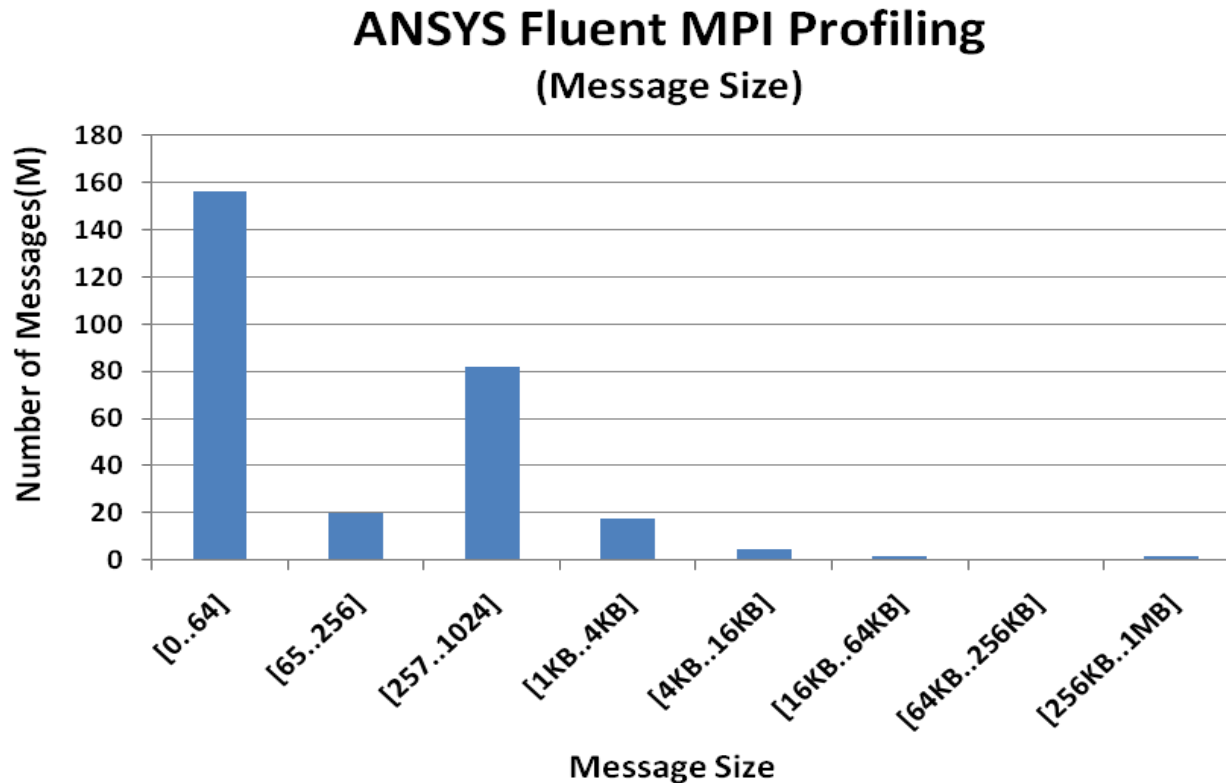


- **MPI_Addreduce, MPI Recv, and MPI_Waitall** are major functions
 - Overhead grows as cluster size increases
 - Truck_14M benchmark

ANSYS Fluent MPI Profiling



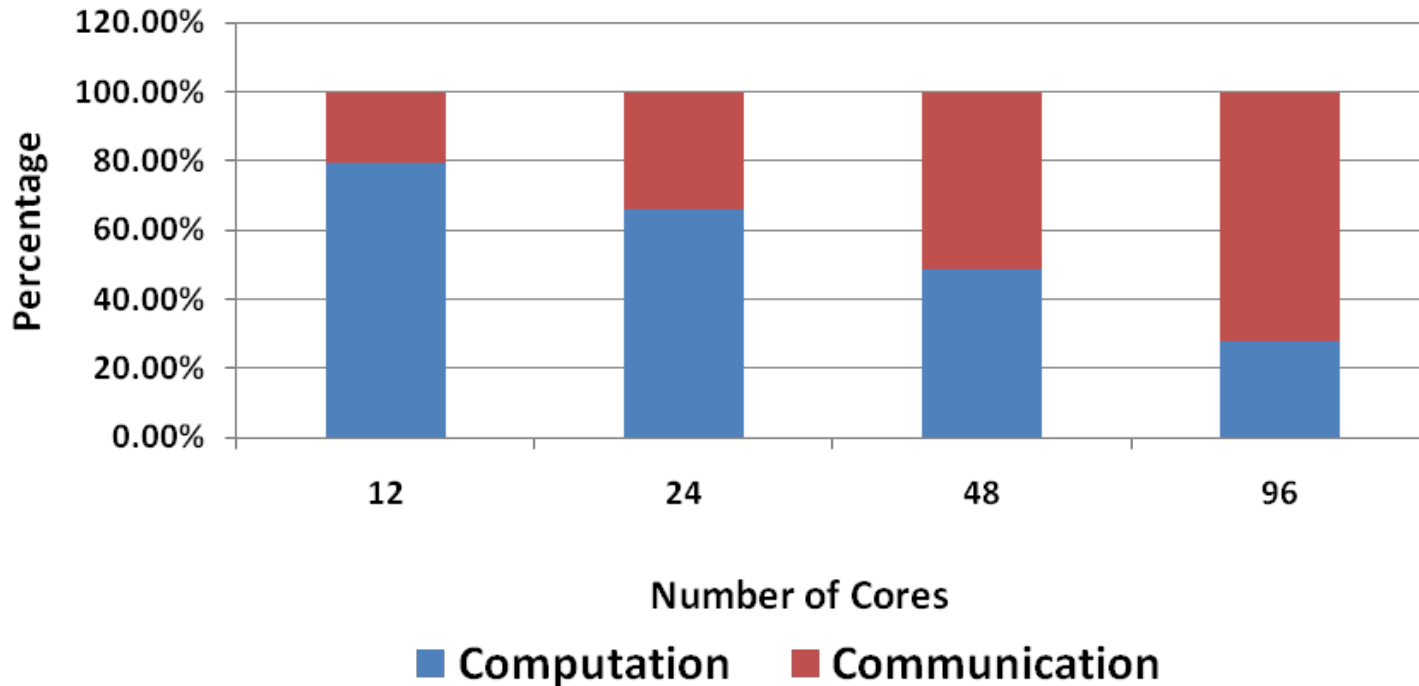
- **Majority messages are small messages**
 - Truck_14M benchmark



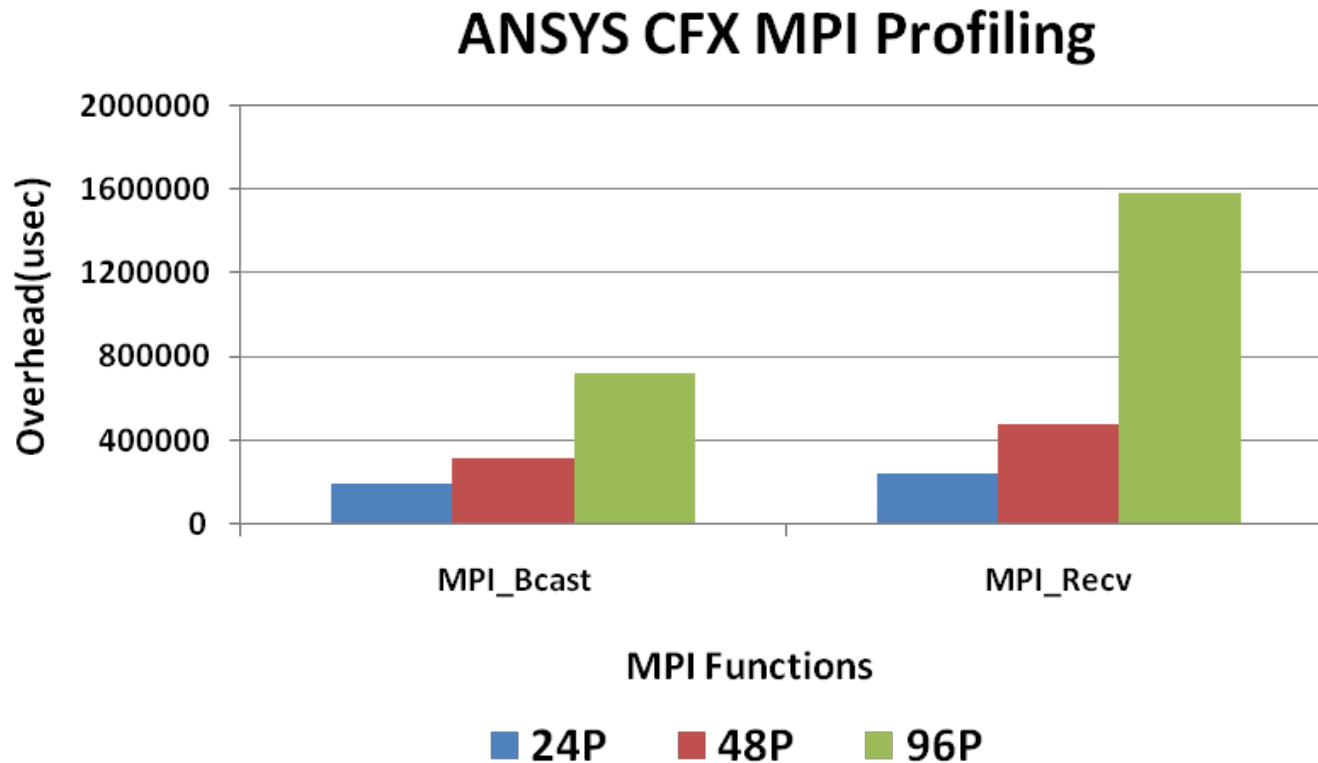
180 Processes

- **Percentage of MPI time increases as cluster size scales**
 - Increases to 76% at 96 cores


ANSYS CFX MPI Profiling (MPI Time)



- **MPI_Bcast, MPI Send/Recv are major functions**
 - Overhead grows as cluster size increases



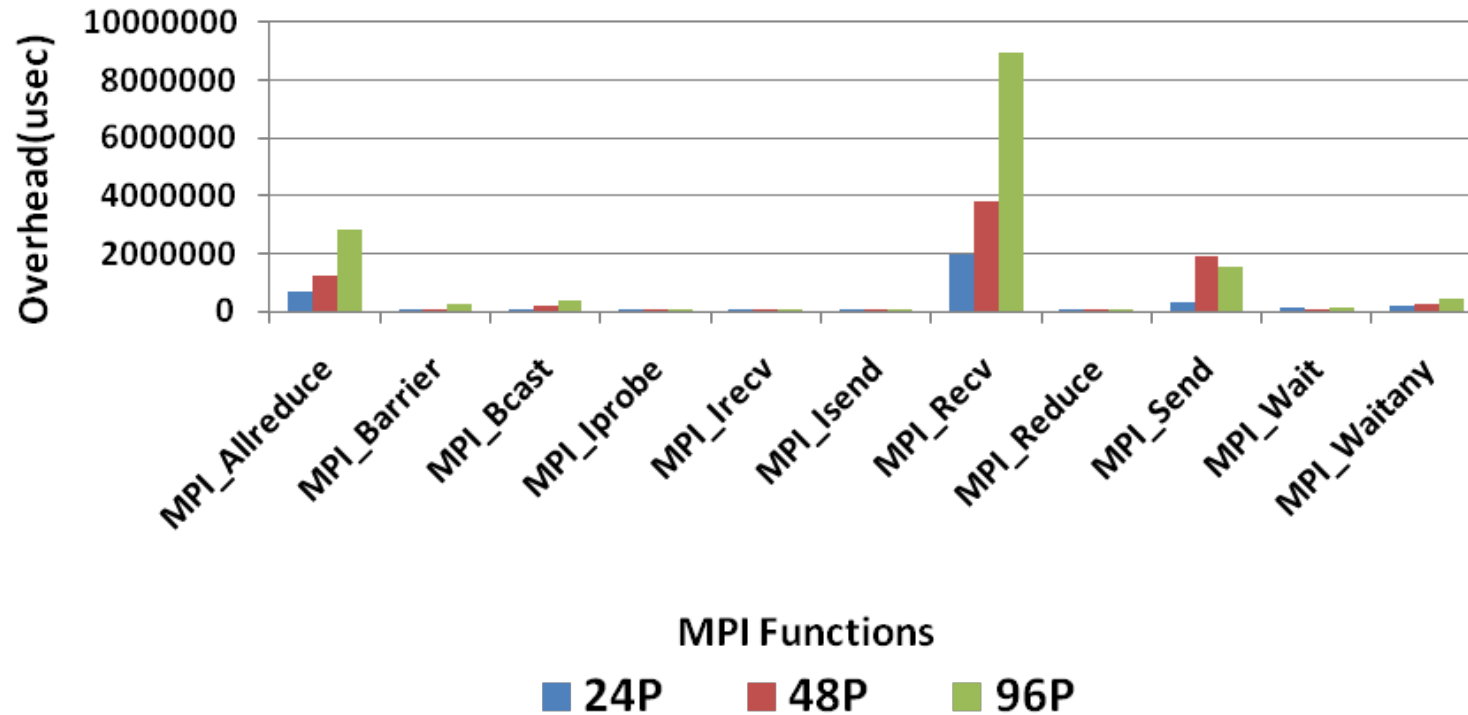
- **Majority messages are small messages**
 - Transonic Airfoil benchmark



96 Processes


- MPI_Allreduce and MPI_Recv are two major functions

ANSYS Mechanical MPI Profiling



- Most messages are small and medium size

ANSYS Mechanical MPI Profiling (Message Size)



96 Processes

- **ANSYS was profiled to identify its communication patterns**
- **Fluent**
 - MPI_Allreduce and MPI_Send/recv generate most overhead
 - Majority messages are small size
 - Interconnect latency is critical to Fluent performance
- **CFX**
 - MPI_Bcast and MPI_Send/recv generate most overhead
 - Majority messages are small size
 - Interconnect latency is critical to CFX performance
- **Mechanical**
 - MPI_Allreduce and MPI_recv create most overhead
 - Both small and medium size messages are used
 - Interconnect latency and bandwidth are critical to Mechanical performance

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