


Improved Performance Speeds Discovery

Optimized HPC
for Life Sciences

Study Shows Up to 94 Percent
Improvement in Time-to-Results



Executive Summary

This performance study, performed jointly by QLogic® and Dell®, focused on the performance gains that scientists will see by upgrading their HPC clusters to the latest processor and interconnect technologies. The study compared performance between generations of Intel® processors using NAMD and GROMACS for simulation purposes. In addition, the study addressed the impact that a high-performance interconnect (InfiniBand®) can have on the overall performance of an HPC cluster. The results of the study showed that significant benefits are provided by HPC technologies in the following areas: computation, communication, and research.

Key Findings:

- *Speed up computation* – Today's processor and interconnect technologies offer a vast improvement in performance compared to previous generations of processors and interconnects. The reduction in time-to-results for a 16-node (128-core) cluster ranged from 78 to 94 percent in all of the NAMD and GROMACS tests.
- *Speed up communication* – InfiniBand has become an industry-standard interconnect for HPC clusters because of its

low latency and high bandwidth. InfiniBand saves researchers valuable time by providing a 50 to 60 percent performance improvement when compared to Gigabit Ethernet.

- *Scale-up research* – The scale or size of the HPC cluster can significantly reduce the time required to complete a simulation. For both NAMD and GROMACS, there is a 12- to 14-fold time improvement in performance when scaling from a single node to 16 nodes.
- *Spend more time on research* – Dell's HPC solution is a modular infrastructure stack. It's built from industry-standard hardware components and best-of-breed partner products. The resulting solution delivers the performance and scalability to satisfy the most demanding application requirements at a fraction of the cost of proprietary systems.

Introduction

The challenge of finding the next great discovery or cure can be accelerated with the right set of tools. Breakthrough research and new discoveries are being powered by High Performance Computing (HPC). In industries such as Life Sciences, speed matters—organizations are searching for life-changing discoveries and are highly dependent upon compute-intensive

tools. Simulation that just a few years ago would have required months to run can now be performed in a matter of minutes using today's HPC systems. The performance of a system determines its speed—speed to discovery, speed to results, and speed to return on investment.

This paper provides key comparative performance data, including computing and interconnect information, for two of the leading molecular dynamic applications. This information can help you select and configure the right computing solution to meet your organization's goals and requirements.


During the past year, Dell and QLogic have partnered together to create multiple HPC test-bed environments designed to better characterize a broad range of workloads and optimize overall performance. The project was performed in QLogic's state-of-the-art NETTrack Developer Center (NDC). The NDC hosts a wide range of systems and incorporates the latest in advanced processors, graphics systems, interconnects, storage, networking, and related hardware that are rigorously tested using leading operating systems, middleware, and applications. As a result, the NDC is useful in helping to characterize the capabilities of various high-performance solutions for compute-intensive workloads. The computing systems at the NDC can be reserved by organizations for a "test drive," so that they can better understand the performance potential of the latest technologies with their specific models and simulations.

Performance Study

Understanding the behavior of an application and how it performs on a given platform is critical to achieving the ultimate goal: fast output and agile analysis of research data and results. During the collaboration on the performance study by Dell and QLogic, the NDC was an instrumental resource in benchmarking and profiling several life science applications. As a result, a greater understanding was achieved of the impact of different processor and interconnect configurations on these applications. The following information describes NAMD and GROMACS, and provides details about test configurations and results:

NAMD


NAMD is a parallel molecular dynamics code designed for high-performance simulation of large biomolecular systems. Based on Charm++ parallel objects, NAMD scales to hundreds of processors on high-end parallel platforms. The code uses the popular molecular graphics program VMD for



simulation setup and trajectory analysis NAMD is distributed free of charge and includes source code. The NIH National Center for Research Resources supports the development of NAMD.

GROMACS

GROningen Machine for Chemical Simulations (GROMACS), a molecular dynamics simulation package, was originally developed at the University of Groningen and is now maintained and extended at various institutions,



including the University of Uppsala, the University of Stockholm, and the Max Planck Institute for Polymer Research. The application is a versatile package primarily designed to simulate the Newtonian equations of motion for systems with hundreds to millions of particles, for example, biochemical molecules such as proteins, lipids, and nucleic acids that have complicated bonded interactions. In addition, GROMACS is extremely fast at calculating non-bonded interactions (which usually dominate simulations), so many groups are also using it for research on non-biological systems (for example, polymers).

Test Configurations

Performance profiles of NAMD and GROMACS simulations were completed using Dell HPC systems, which were configured with multiple interconnects for broad-based analysis of the overall performance and capabilities of the applications. The first test bed was comprised of 128 cores of Intel Xeon® 5472 processors (Harpertown Cluster). The second test bed consisted of 128 cores of the Intel Xeon 5570 processors (Nehalem Cluster). Both systems were used to test generational advances in processor performance as well as data transfer rates and interconnect capabilities between on-board Ethernet and QDR InfiniBand. The following are the configuration details for each cluster:

Cluster Name	Harpertown	Nehalem
Nodes	16	16
Total Core Count	128	128
Server Configuration		
Processor Type	Dual Xeon 5472	Dual Xeon 5570
Processor Speed	3.0 GHz	2.93 GHz
Memory/Node	16GB	24GB
Interconnect		
Giga Ethernet	On-board	On-board
InfiniBand	TrueScale™ DDR	TrueScale QDR

Time-to-Results

The size or scale of the HPC cluster that is used for research can make a significant difference in the time-to-results. Figure 1 shows that there is a consistent improvement in the reduction of time for both NAMD and GROMACS simulations as the size of the cluster is increased.

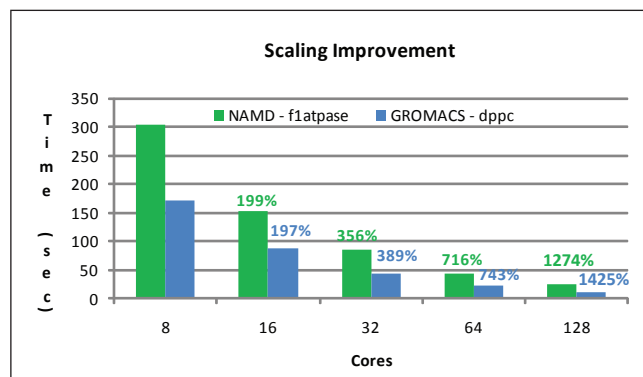


Figure 1 – NAMD and GROMACS Scaling Improvement

The NAMD model tested was f1atpase. ATP synthase is an enzyme that synthesizes adenosine triphosphate (ATP), the common molecular energy unit in cells. The F1-ATPase benchmark is a model of the F1 subunit of ATP synthase containing 327,506 atoms. The GROMACS test consisted of the dppc model, which is a simulation of the phospholipid membrane bilayer, similar to the structure of a cell membrane. The simulation consists of 1,024 dipalmitoylphosphatidylcholine (DPPC) lipids simulated in water totaling 121,856 atoms. Both of these simulations show that a researcher can obtain significantly reduced time-to-results as the cluster is scaled up to yield a potential 12- to 14-fold performance improvement.

Technology Improvement

HPC clusters acquired before April 2009 are probably based on Intel Xeon 5400 or earlier processor technology. If the cluster that you are currently using is based on this generation of processor, and the interconnect is gigabit Ethernet, then substantial performance gains can be achieved by upgrading both the processor and the interconnect technologies.




Figure 2 – NAMD STMV Performance Results

The NAMD Satellite Tobacco Mosaic Virus (STMV) model shows significant performance improvements when both the latest processor and interconnect technologies are used. The STMV benchmark simulates 1,066,628 atoms. It is useful for demonstrating scaling to thousands of processors. STMV is a small, icosahedral plant virus which worsens the symptoms of infection by the Tobacco Mosaic Virus (TMV). The simulation would have taken a single 2006 desktop computer around 35 years to complete. However, Figure 2 shows that a 16-node (128-core) system can complete the simulation in minutes. Use of the latest processor and interconnect technologies can further boost performance and solve the simulation in seconds. There is a 78 percent reduction in time-to-results when both the processor and the interconnect are upgraded.

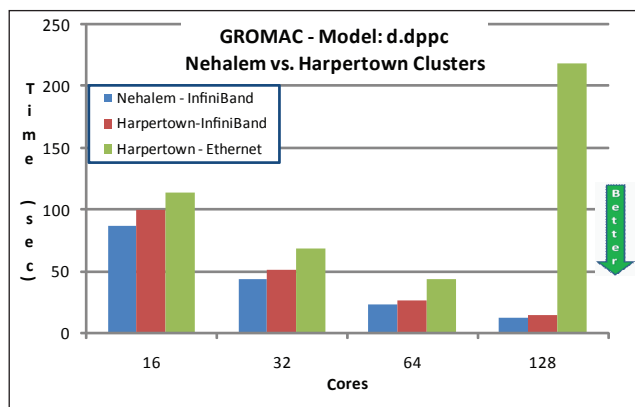


Figure 3 – GROMACS dppc Results

The GROMACS dppc simulation was used again for the next comparison and was run three times. The first run was on the Harpertown cluster with Gigabit Ethernet. The second was on the Harpertown cluster, but this time with InfiniBand. The last run was on the Nehalem cluster with InfiniBand. GROMACS showed a dramatic improvement in performance when both the processor and the interconnect were upgraded. In this test, there was an **outstanding 94 percent reduction in time-to-results** when comparing the

Nehalem-InfiniBand cluster to the Harpertown-Ethernet cluster. There was a 20 percent improvement in performance with the Nehalem cluster compared to the Harpertown cluster when both were using InfiniBand.

Interconnect Impact

The choice of interconnect can make a major difference in both the performance and the effectiveness of the HPC cluster. InfiniBand, with its low latency and high bandwidth, has shown itself to be an optimal interconnect, especially as the size of the cluster is increased. Figure 3 shows the results for the two runs of the NAMD ApoA1 simulation; one with Gigabit Ethernet and the other with InfiniBand. Apolipoprotein A1 (ApoA1) is the major protein component of high-density lipoprotein (HDL) in the bloodstream and plays a specific role in lipid metabolism. The ApoA1 benchmark consists of 92,224 atoms and has been a standard NAMD cross-platform benchmark for years.

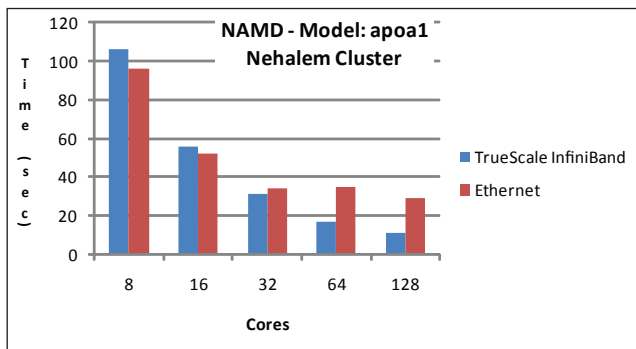


Figure 4 - NAMD ApoA1 Performance Comparison

Figure 4 shows that InfiniBand performance improves as more resources are added to the cluster. For example, going from 32 cores to 128 cores results in a 65 percent reduction in time-to-results, whereas with Ethernet the improvement is only 14 percent.

Conclusion

HPC clusters, based on the latest generation of high-performance processors and interconnect technologies, are an effective scientific tool for accelerating research and discoveries. Whether your organization is upgrading an existing cluster or looking to buy a new cluster, Dell and QLogic can provide an optimized solution which integrates the industry's best-in-breed components.

For Dell HPC information, visit www.dell.com/hpc.

For more information about QLogic's HPC networking solutions, visit www.qlogic.com/Products/HPCNetworking/Pages/HPC_products_landingpage.aspx.

Disclaimer

Reasonable efforts have been made to ensure the validity and accuracy of these performance tests. QLogic Corporation is not liable for any error in this published white paper or the results thereof. Variation in results may be a result of change in configuration or in the environment. QLogic specifically disclaims any warranty, expressed or implied, relating to the test results and their accuracy, analysis, completeness or quality.



Corporate Headquarters QLogic Corporation 26650 Aliso Viejo Parkway Aliso Viejo, CA 92656 949.389.6000

www.qlogic.com

International Offices UK | Ireland | Germany | India | Japan | China | Hong Kong | Singapore | Taiwan

© 2010 QLogic Corporation. Specifications are subject to change without notice. All rights reserved worldwide. QLogic and the QLogic logo are registered trademarks and TrueScale is a trademark of QLogic Corporation. Dell is a registered trademark of Dell Inc. InfiniBand is a registered trademark and service mark of the InfiniBand Trade Association. Intel and Xeon are registered trademarks of Intel Corporation. All other brand and product names are trademarks or registered trademarks of their respective owners. Information supplied by QLogic Corporation is believed to be accurate and reliable. QLogic Corporation assumes no responsibility for any errors in this brochure. QLogic Corporation reserves the right, without notice, to make changes in product design or specifications.