GROMACS 4.6 Pre-Beta Benchmark Report

Benefits of GPU Accelerated Computing

- Faster than CPU only systems in all tests
- Large performance boost with marginal price increase
- Energy usage cut in half or less
- GPUs scale very well within a node and over multiple nodes
- New Tesla K10 GPU board contains 2 of our fastest single precision GPUs yet

Great Scaling in Small Systems

Get up to 3.7x performance compared to CPU-only nodes

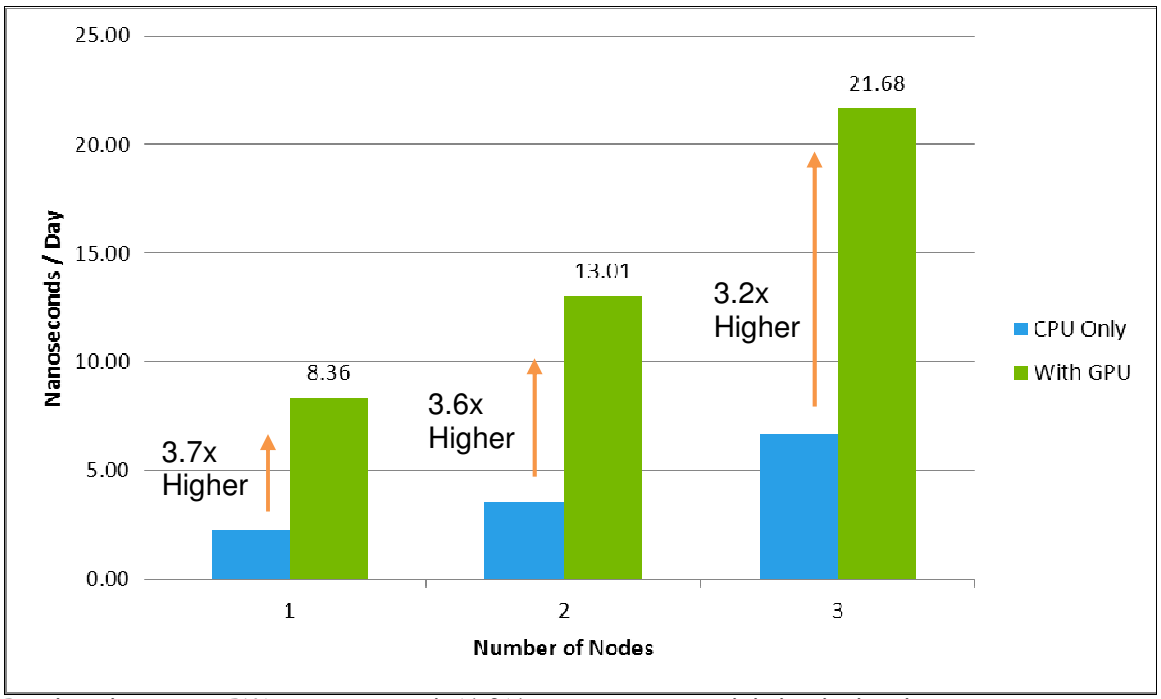

Benchmark systems: RNAse in water with 16,816 atoms in truncated dodecahedron box

Running GROMACS 4.6 pre-beta with CUDA 4.1

Each blue node contains 1x Intel X5550 CPU (95W TDP, 4 Cores per CPU)

Each green node contains 1x Intel X5550 CPU (95W TDP, 4 Cores per CPU) and 1x NVIDIA M2090 (225W TDP per GPU board)

Additional Strong Scaling on Larger System

Up to 128 nodes, NVIDIA GPU-accelerated nodes deliver 2-3x performance when compared to CPU-only nodes

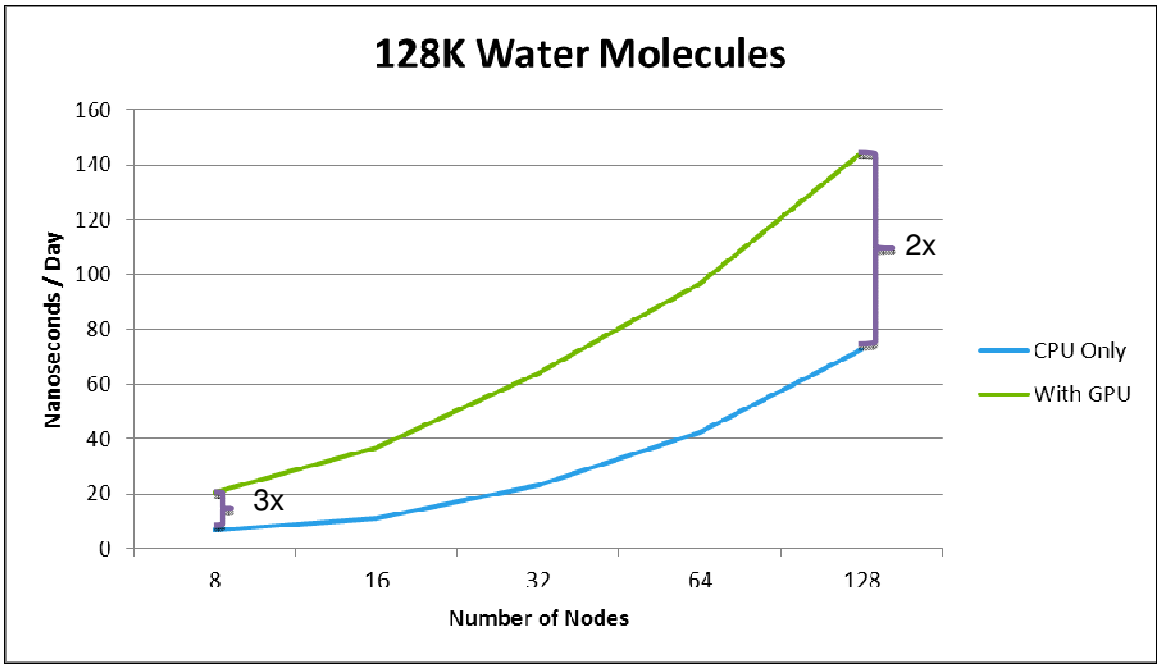

Running GROMACS 4.6 pre-beta with CUDA 4.1

Each blue node contains 1x Intel X5670 (95W TDP, 6 Cores per CPU)

Each green node contains 1x Intel X5670 (95W TDP, 6 Cores per CPU) and 1x NVIDIA M2070 (225W TDP per GPU board)

Replace 3 Nodes with 2 GPUs

Save thousands of dollars and perform 25% faster

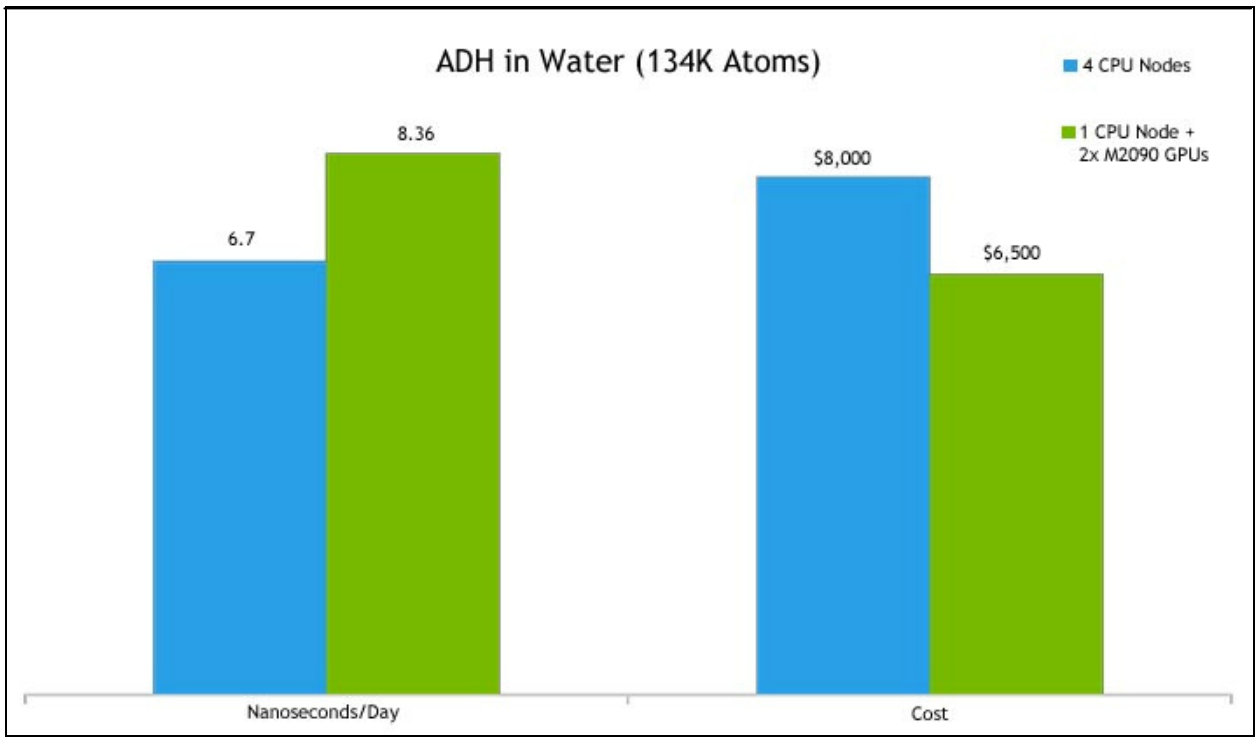

Running GROMACS 4.6 pre-beta with CUDA 4.1

The blue node contains 2x Intel X5550 CPUs (95W TDP, 4 Cores per CPU)

The green node contains 2x Intel X5550 CPUs and 2x NVIDIA M2090 GPU boards as the GPU (225W TDP per GPU board)

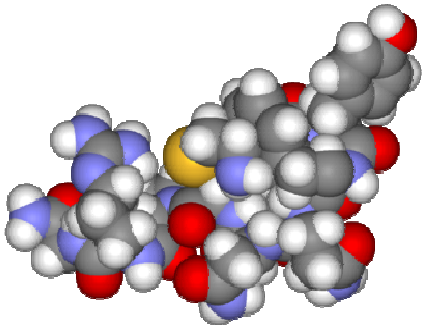

Greener Science

In simulating each nanosecond, the GPU-accelerated system uses 33% less energy

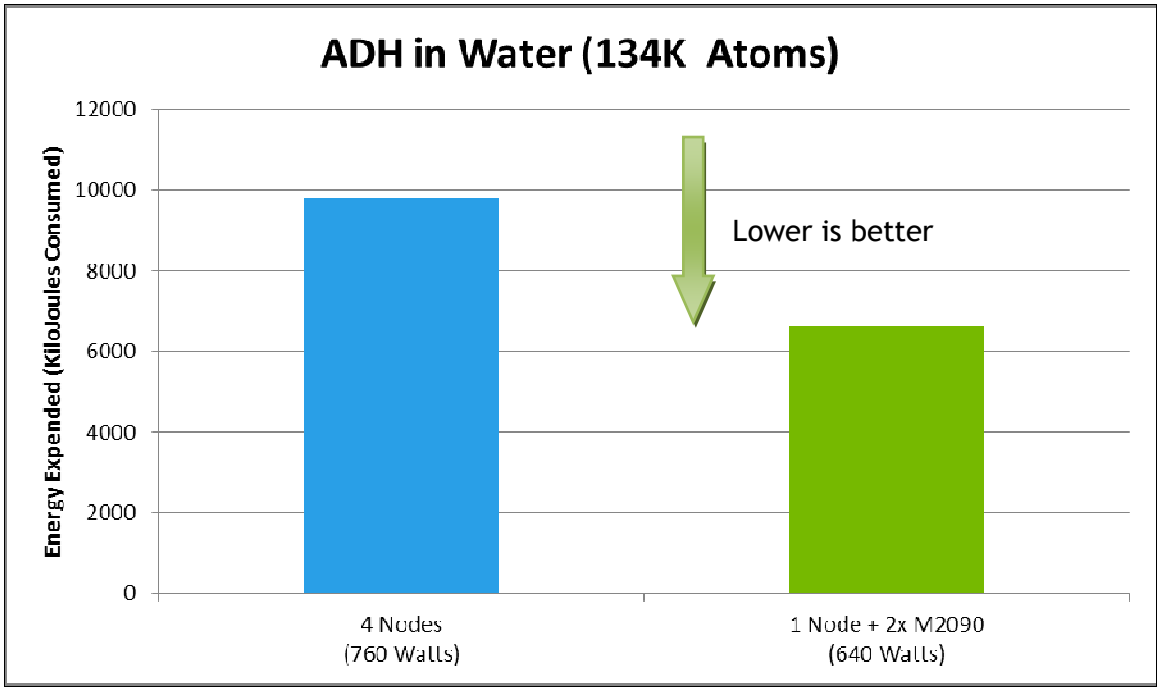

Running GROMACS 4.6 with CUDA 4.1

The blue nodes contain 2x Intel X5550 CPUs (95W TDP, 4 Cores per CPU)

The green node contains 2x Intel X5550 CPUs, 4 Cores per CPU) and 2x NVIDIA M2090s GPU boards (225W TDP per GPU)

Energy Expended = Power x Time

K10 – the Fastest Single Precision GPU

Get 2.6x performance with a K10 GPU Board compared to M2090

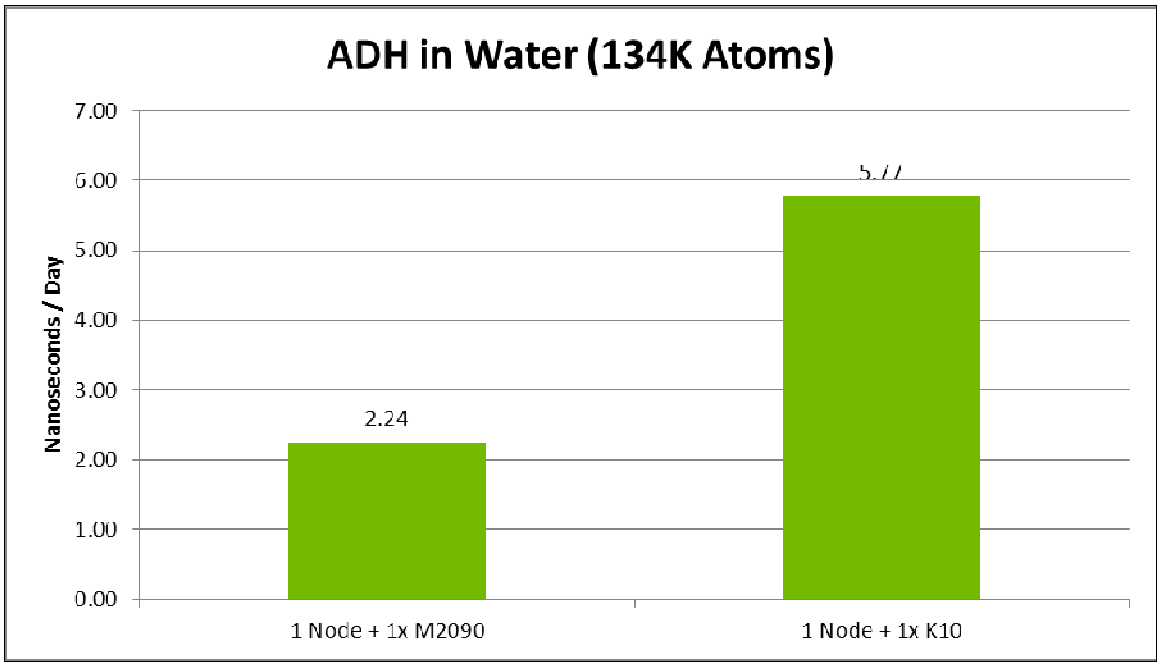

Running GROMACS 4.6 with CUDA 4.1 or CUDA 4.2

Each node contains 2x Intel X5690s for the CPUs (95W TDP, 6 Cores per CPU) and either 1x NVIDIA M2090 or 1x K10 GPU board for the GPUs (225W TDP per GPU board)

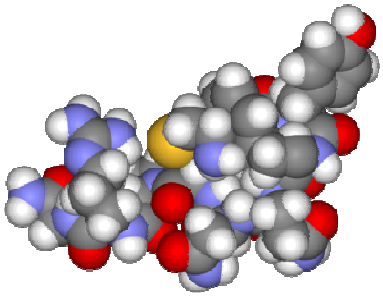

Recommended GPU Node Configuration for Computational Chemistry

Workstation or Single Node Configuration	
# of CPU sockets	2
Cores per CPU socket	4-6+
CPU speed (Ghz)	2.66+
System memory per socket (GB)	32
GPUs	Kepler K10, Fermi M2090, M2075, C2075
# of GPUs per CPU socket	2
GPU memory preference (GB)	6
GPU to CPU connection	PCIe 2.0 or higher
Server storage	500 GB or higher
Network configuration	InfiniBand

Scale to multiple nodes with same single node configuration

Summary/Conclusions

Benefits of GPU Accelerated Computing

- Faster than CPU only systems in all tests
- Large performance boost with small marginal price increase
- Energy usage cut in half or less
- GPUs scale very well from 1 to 4, and beyond
- New Tesla K10 GPU board contains 2 of our fastest single precision GPUs to date

GPU Test Drive- Your Science Apps 5x Faster
Take a Free and Easy Test Drive Today.

Run your computational chemistry simulations 5x faster. Take a free test drive to try GROMACS on a remotely hosted cluster loaded with the latest GPU-accelerated applications and accelerate your results. Simply log on and run your application as usual, no GPU programming expertise required. Try it now and see how you can reduce simulation time from days to hours.

Register for the test drive today at http://www.nvidia.com/object/gpu-test-drive.html#s=gromacs.