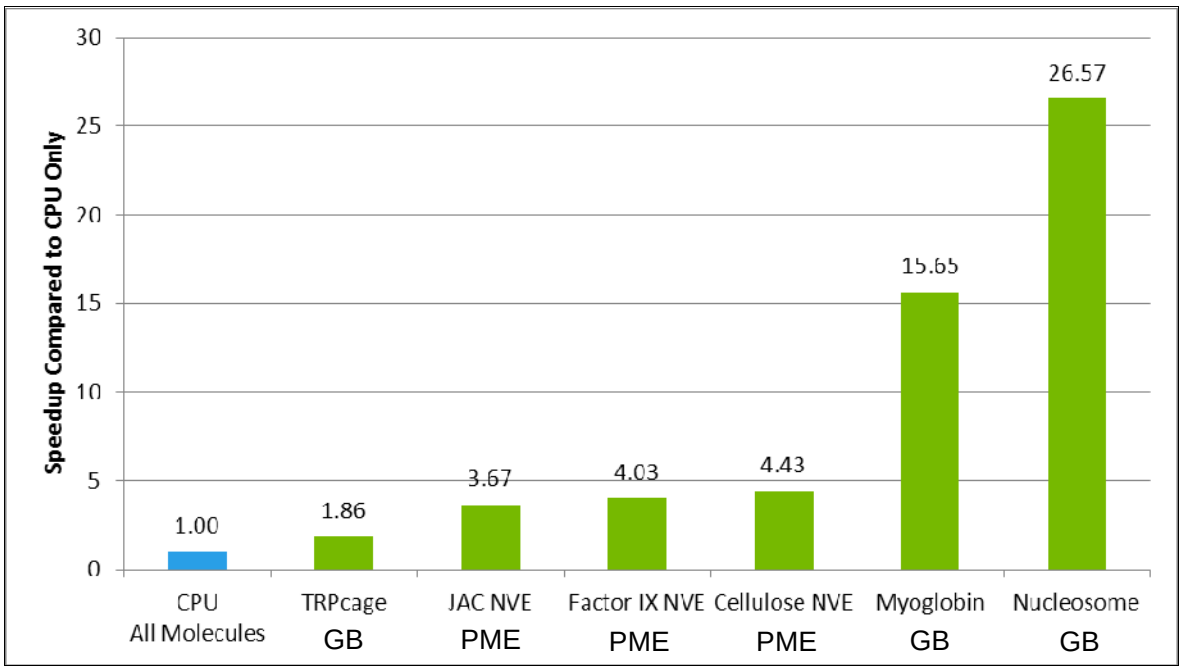
AMBER 12 Benchmark Report

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 to date

Accelerates Simulations of All Sizes

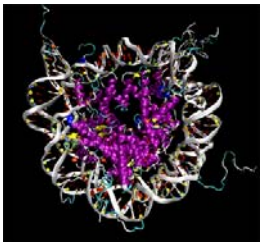
Gain 26x performance by adding just 1x K10 GPU Board when compared to dual CPU performance



Running AMBER 12 with CUDA 4.2 ECC Off

The **blue node** contains 2x Intel X5670 CPUs (6 Cores per CPU)

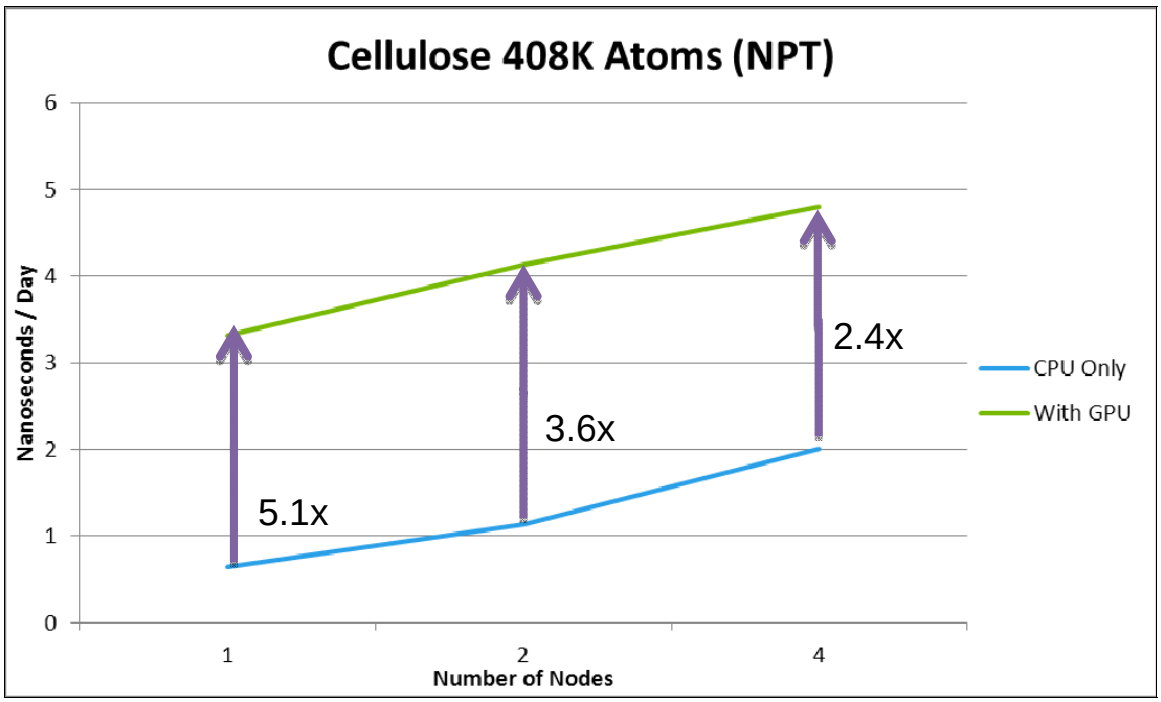
Each **green node** contains 2x Intel X5670 CPUs (6 Cores per CPU) plus 1x NVIDIA K10 GPU board



Nucleosome

Strong Scaling over Nodes

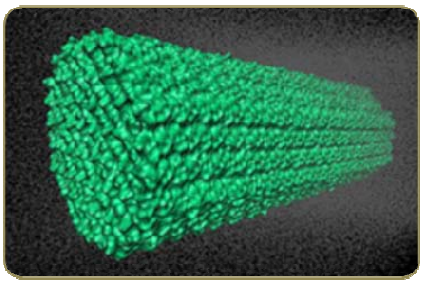
GPUs significantly outperform CPUs while scaling over multiple nodes



Running AMBER 12 with CUDA 4.2 ECC Off

The **blue nodes** contains 2x Intel X5670 CPUs (6 Cores per CPU)

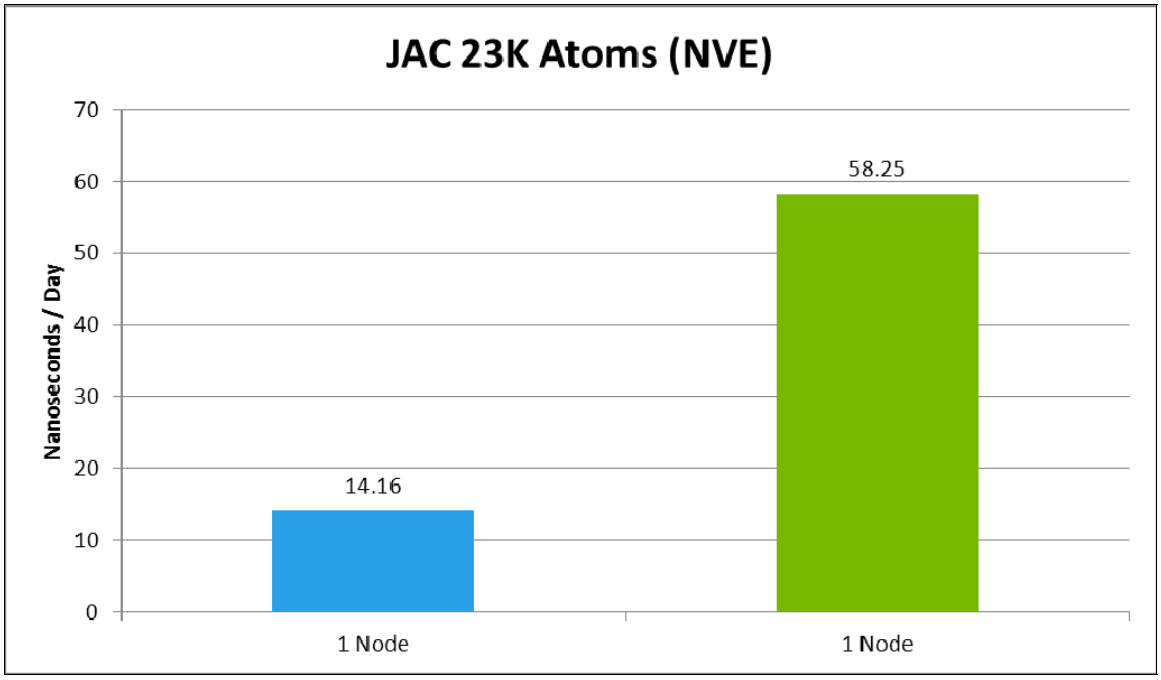
The **green nodes** contains 2x Intel X5670 CPUs (6 Cores per CPU) plus 2x NVIDIA K10 GPU boards



Cellulose

Extreme Performance

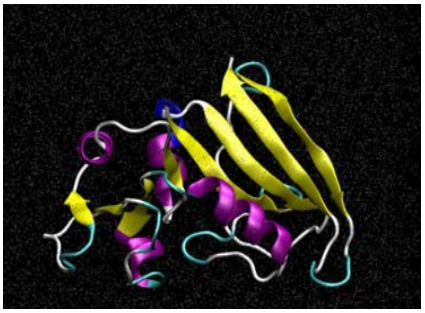
Gain 4X performance by adding just 2x K10 GPU Boards when compared to dual CPU performance



Running AMBER 12 with CUDA 4.2 ECC Off

The **blue node** contains 2x Intel X5670 CPUs (6 Cores per CPU)

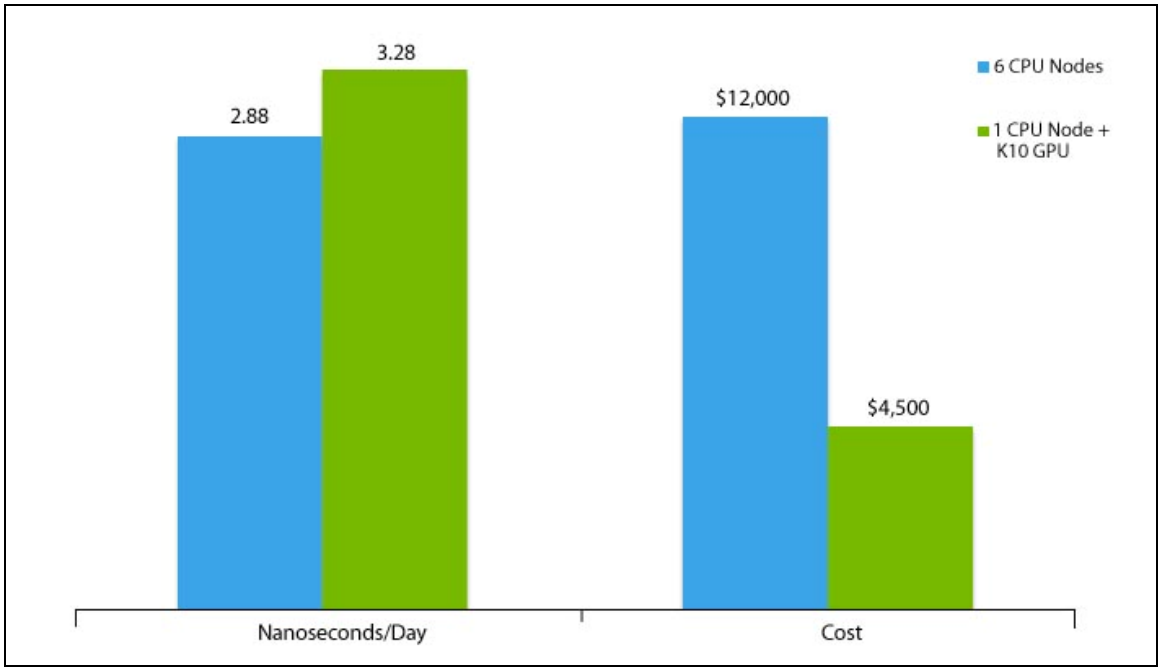
The **green node** contains 2x Intel X5670 CPUs (6 Cores per CPU) plus 2x NVIDIA K10 GPU boards



DHFR

Replace 5 Nodes with 1x K10 GPU Board

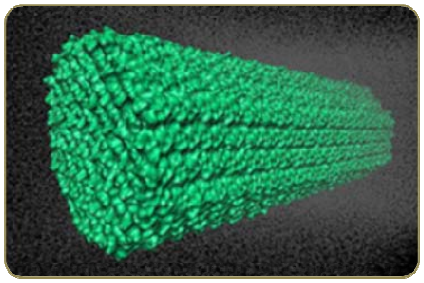
Cut down simulation costs to ⅓ and gain higher performance



Running AMBER 12 with CUDA 4.2 ECC Off

Each **blue node** contains 2x Intel X5670 CPUs

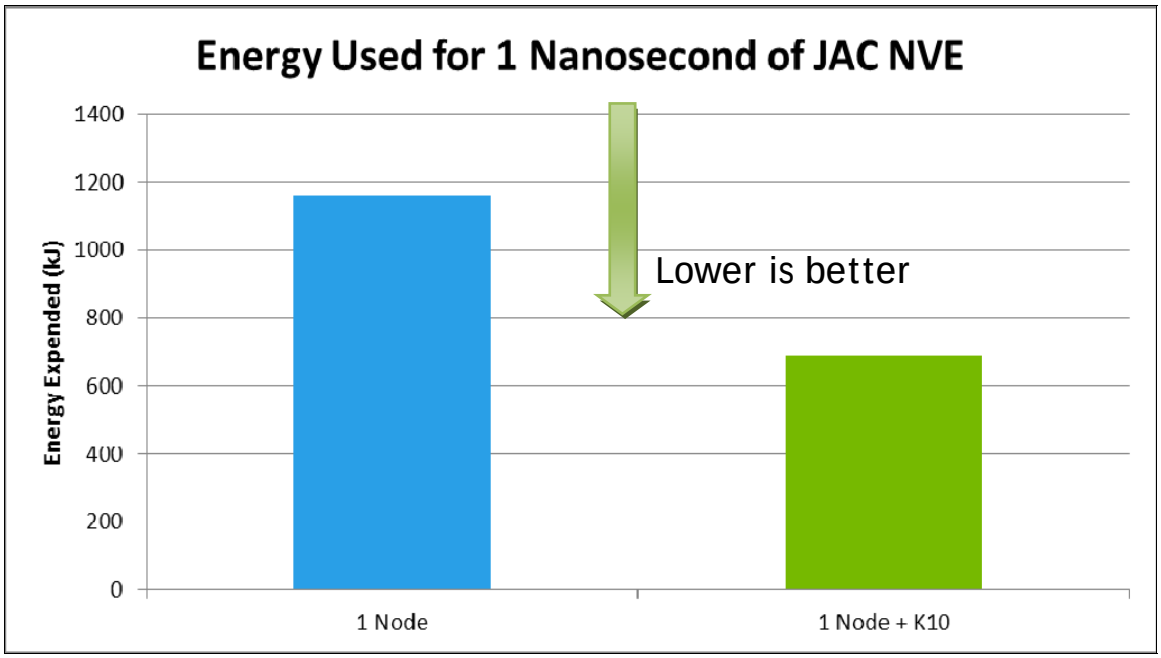
The **green node** contains 2x Intel X5670 CPUs and 1x NVIDIA K10 GPU Board



Cellulose

Greener Science (per Nanosecond of Simulation)

The GPU Accelerated system uses 40% less energy



Running AMBER 12 with CUDA 4.2 ECC Off

The **blue node** contains 2x Intel X5670 CPUs (95W each, 6 Cores per CPU)

The **green node** contains 2x Intel X5670 CPUs (95W each, 6 Cores per CPU) and 1x NVIDIA K10 GPU board (225W per GPU board)

Energy Expended = Power x Time

Recommended GPU Node Configuration for Computational Chemistry

Workstation or Single Node Configuration	
# of CPU sockets	2
Cores per CPU socket	4+
CPU speed (Ghz)	2.66+
System memory per socket (GB)	32
GPUs	Kepler K10, Fermi M2090, M2075, C2075
# of GPUs per CPU socket	1-2 (1 CPU core drives 1 GPU)
GPU memory preference (GB)	6
GPU to CPU connection	PCIe 2.0 or higher
Server storage	500 GB or higher
Network configuration	InfiniBand or better

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 AMBER 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=amber.