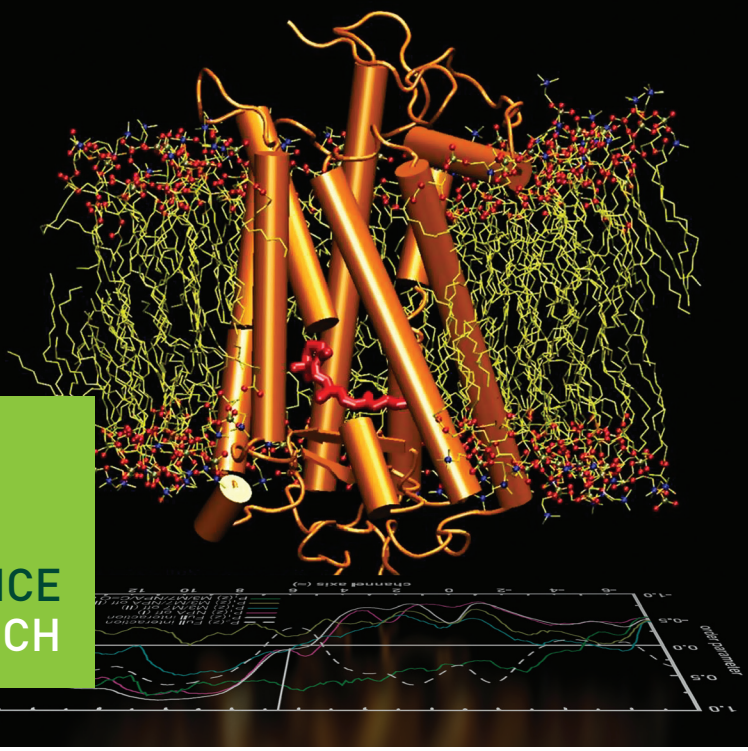

# ENABLING NEW SCIENCE
# TESLA™ BIO WORKBENCH

The NVIDIA® Tesla™ Bio Workbench enables biophysicists and computational chemists to push the boundaries of life sciences research. It turns a standard PC into a "computational laboratory" capable of running complex bioscience codes, in fields such as drug discovery and DNA sequencing, more than 10-20 times faster through the use of NVIDIA Tesla GPUs.

It consists of bioscience applications; a community site for downloading, discussing, and viewing the results of these applications; and GPU-based platforms that enable these applications to run at 1/10$^{th}$ the cost of CPU-only computers.

Complex molecular simulations that had been only possible using supercomputing resources can now be run on an individual workstation, optimizing the scientific workflow and accelerating the pace of research. These simulations can also be scaled up to GPU-based clusters of servers to simulate large molecules and systems that would have otherwise required a supercomputer.

Applications that are accelerated on GPUs include:

- Molecular Dynamics & Quantum Chemistry
  - AMBER, GROMACS, HOOMD, LAMMPS, NAMD, TeraChem (Quantum Chemistry), VMD
- Bio Informatics
  - CUDA-BLASTP, CUDA-EC, CUDA-MEME, CUDASW++ (Smith-Waterman), GPU-HMMER, MUMmerGPU

For more information, visit www.nvidia.com/bio_workbench

## GPU SOLUTIONS

The Tesla Bio Workbench applications can be deployed on GPU-based desktop personal supercomputers or in data center solutions. Built on the revolutionary, massively parallel CUDA™ GPU computing architecture, these solutions are designed to accelerate the pace of computational science and are available today. Visit www.nvidia.com/tesla to find out more about:

### WORKSTATION SOLUTIONS:

TESLA PERSONAL SUPERCOMPUTER
For personal supercomputing at your desk

### DATA CENTER SOLUTIONS:

TESLA GPU COMPUTING CLUSTERS
For computing with large-scale installations

# GPU-BASED MOLECULAR DYNAMICS & QUANTUM CHEMISTRY APPLICATIONS

## AMBER

The explicit solvent and implicit solvent simulations in AMBER have been accelerated using CUDA-enabled GPUs. Paired with a Tesla GPU computing solution built on the CUDA architecture, it enables more than 10x speedup compared to a single quad-core CPU.

Benchmark data for explicit solvent PME and implicit solvent GB Benchmarks. Details can be found on the AMBER 11 NVIDIA benchmark page courtesy of the San Diego Supercomputing Center.

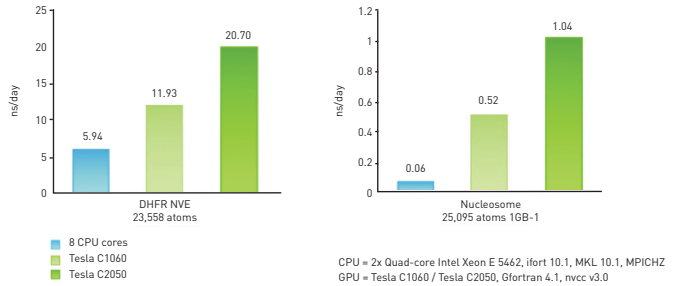


CPU = 2x Quad-core Intel Xeon E 5462, ifort 10.1, MKL 10.1, MPICHZ
GPU = Tesla C1060 / Tesla C2050, Gfortran 4.1, nvcc v3.0

## HOOMD

HOOMD-blue is a general purpose particle dynamics package written from ground-up to take advantage of the revolutionary CUDA architecture in NVIDIA GPUs. It contains a number of different force fields and integrators and its object oriented design enables adding additional ones easily.

Detailed benchmarks are available on the HOOMD-blue page. One Tesla GPU running HOOMD-blue can outperform 32 CPU cores.

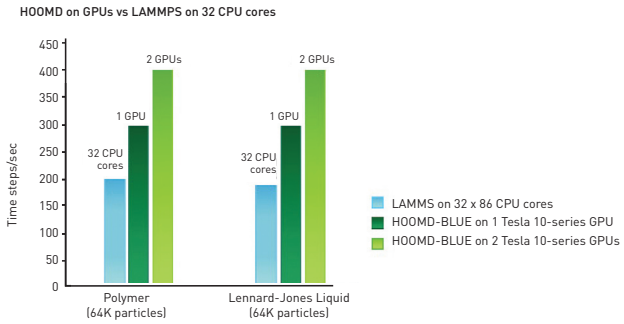


Data courtesy of University of Michigan

## GROMACS

GROMACS is a molecular dynamics package designed primarily for simulation of biochemical molecules like proteins, lipids, and nucleic acids that have a lot complicated bonded interactions. The CUDA port of GROMACS enabling GPU acceleration is now available in beta and supports Particle-Mesh-Ewald (PME), arbitrary forms of non-bonded interactions, and implicit solvent Generalized Born methods.

The CUDA version of GROMACS currently supports a single GPU and produces the following results:

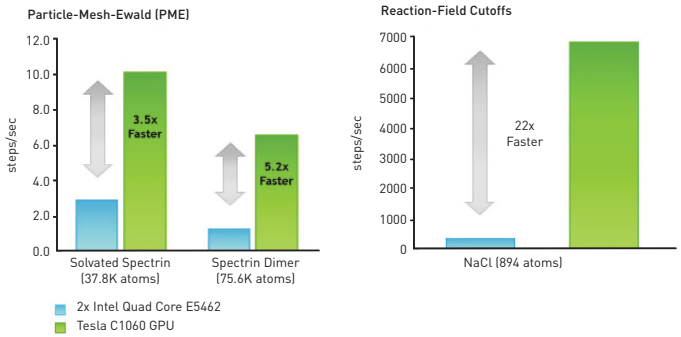


Data courtesy of Stockholm Center for Biomembrane Research

## LAMMPS

LAMMPS is a classical molecular dynamics package written to run well on parallel machines. The CUDA version of LAMMPS is accelerated by moving the force calculations to the GPU.

LAMMPS on the GPU scales very well as seen by the results below. Two Tesla GPUs running GPU-LAMMPS outperforms 24 CPUs.

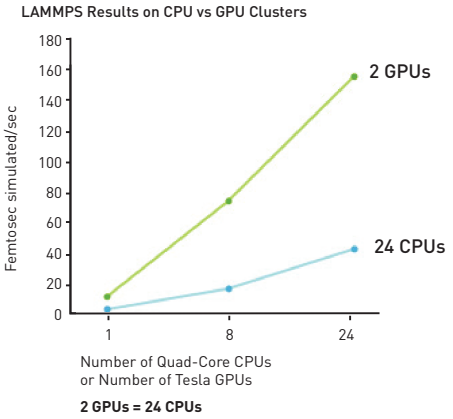


Data courtesy of Scott Hampton & Pratul K. Agarwal, Oak Ridge National Laboratory

# GPU-BASED MOLECULAR DYNAMICS & QUANTUM CHEMISTRY APPLICATIONS

## NAMD

The Team at University of Illinois at Urbana-Champaign (UIUC) has been enabling CUDA-acceleration on NAMD since 2007. They have performed scaling experiments on the NCSA Tesla-based Lincoln cluster and demonstrated that 4 Tesla GPUs can outperform a cluster with 16 quad-core CPUs.

NAMD scales very well on a Tesla GPU cluster as demonstrated by these results.

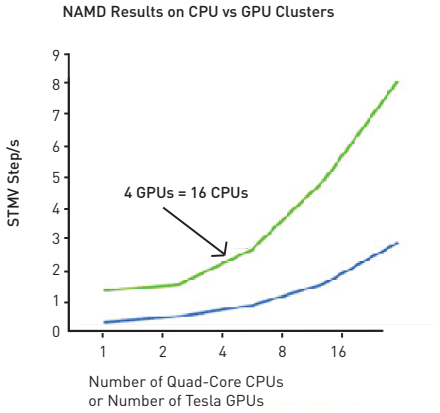


Data courtesy of Theoretical and Computation Bio-physics Group, UIUC

## VMD

VMD is a molecular visualization program for displaying, animating, and analyzing large biomolecular systems using 3-D graphics and built-in scripting. Several key kernels and applications in VMD now take advantage of the massively parallel CUDA architecture of NVIDIA's GPUs. These applications run 20x to 100x faster when using a NVIDIA CUDA GPU compared to running them on a CPU only.

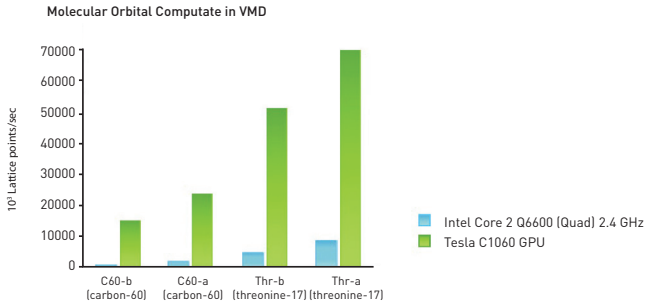


Data courtesy of Theoretical and Computational Bio-physics Group, UIUC

## TERACHEM

TeraChem is the first general purpose quantum chemistry software package designed ground up specifically to take advantage of the massively parallel CUDA architecture of NVIDIA GPUs. TeraChem currently supports density functional theory (DFT) and Hartree-Fock (HF) methods.

TeraChem supports multiple GPUs using p-threads and will soon be MPI enabled. A workstation with 4 Tesla GPUs running TeraChem outperforms 256 quad-core CPUs running GAMESS.

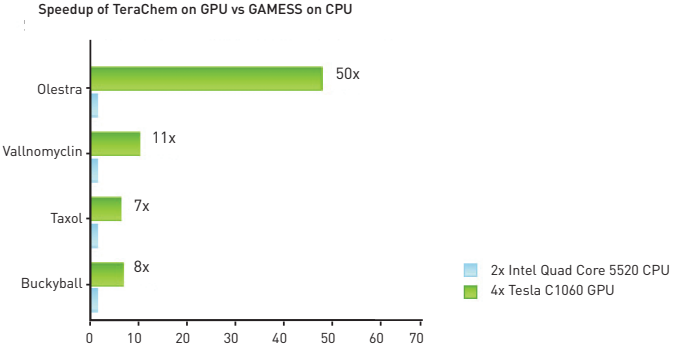


Data courtesy of Petachem

# GPU-BASED BIO-INFORMATICS APPLICATIONS

## CUDA-BLASTP

CUDA-BLASTP is designed to accelerate NCBI BLAST for scanning protein sequence databases by taking advantage of the massively parallel CUDA architecture of NVIDIA Tesla GPUs. CUDA-BLASTP also has a utility to convert FASTA format database into files readable by CUDA-BLASTP.

CUDA-BLASTP running on a workstation with two Tesla C1060 GPUs is 10x faster than NCBI BLAST (2.2.22) running on an Intel i7-920 CPU. This cuts compute time from minutes on CPUs to seconds using GPUs.

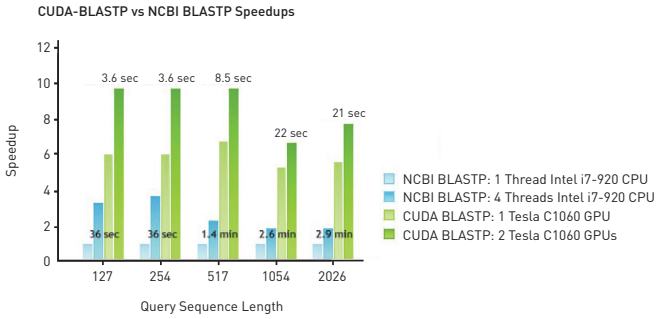


Data courtesy of Nanyang Technological University, Singapore

## CUDA-MEME

CUDA-MEME is a motif discovery software based on MEME (version 3.5.4). It accelerates MEME by taking advantage of the massively parallel CUDA architecture of NVIDIA Tesla GPUs. It supports the OOPS and ZOOPS models in MEME.

CUDA-MEME running on one Tesla C1060 GPU is up to 23x faster than MEME running on a x86 CPU. This cuts compute time from hours on CPUs to minutes using GPUs. The data in the chart below are for the OOPS (one occurrence per sequence) model for 4 datasets.

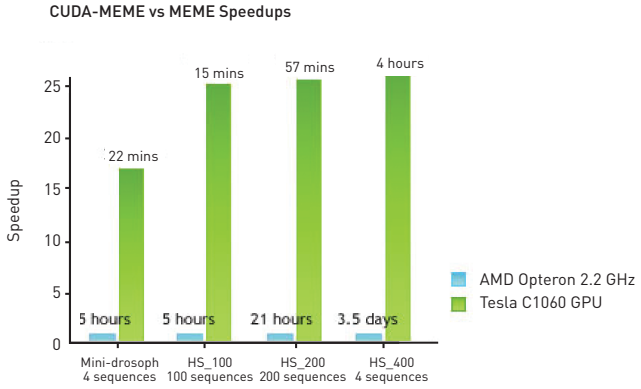


Data courtesy of Nanyang Technological University, Singapore

## CUDA-EC

CUDA-EC is a fast parallel sequence error correction tool for short reads. It corrects sequencing errors in high-throughput short-read (HTSR) data and accelerates HTSR by taking advantage of the massively parallel CUDA architecture of NVIDIA Tesla GPUs. Error correction is a preprocessing step for many DNA fragment assembly tools and is very useful for the new high-throughput sequencing machines.

CUDA-EC running on one Tesla C1060 GPU is up to 20x faster than Euler-SR running on a x86 CPU. This cuts compute time from minutes on CPUs to seconds using GPUs. The data in the chart below are error rates of 1%, 2%, and 3% denoted by A1, A2, A3 and so on for 5 different reference genomes.

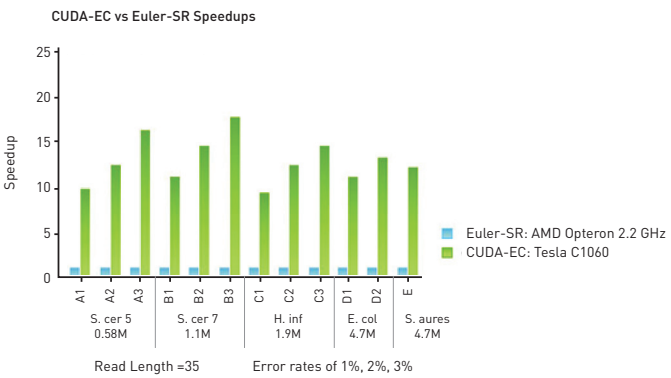


Data courtesy of Nanyang Technological University, Singapore

## CUDASW++ (SMITH-WATERMAN)

CUDASW++ is a bio-informatics software for Smith-Waterman protein database searches that takes advantage of the massively parallel CUDA architecture of NVIDIA Tesla GPUs to perform sequence searches 10x-50x faster than NCBI BLAST. CUDASW++ supports query lengths up to 59K.

CUDASW++ can take advantage of multiple Tesla GPUs and runs between 10x-50x faster than NCBI BLAST, getting up to 30 GCUPs on query lengths of over 5000. Whereas BLAST is a heuristic algorithm, CUDASW++ does optimal local alignment using the Smith-Waterman algorithm.

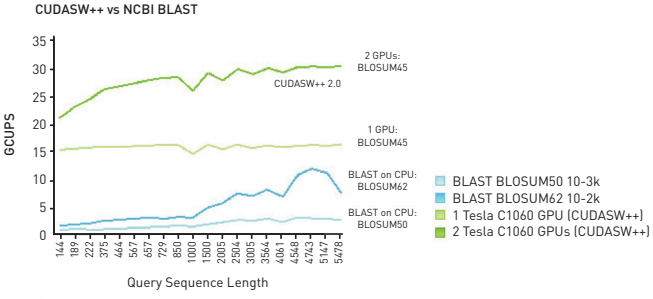


Data courtesy of Nanyang Technological University, Singapore

## GPU-HMMER

GPU-HMMER is a bioinformatics software that does protein sequence alignment using profile HMMs by taking advantage of the massively parallel CUDA architecture of NVIDIA Tesla GPUs. GPU-HMMER is 60-100x faster than HMMER (2.0).

GPU-HMMER accelerates the hmmsearch tool using GPUs and gets speed-ups ranging from 60-100x. GPU-HMMER can take advantage of multiple Tesla GPUs in a workstation to reduce the search from hours on a CPU to minutes using a GPU.

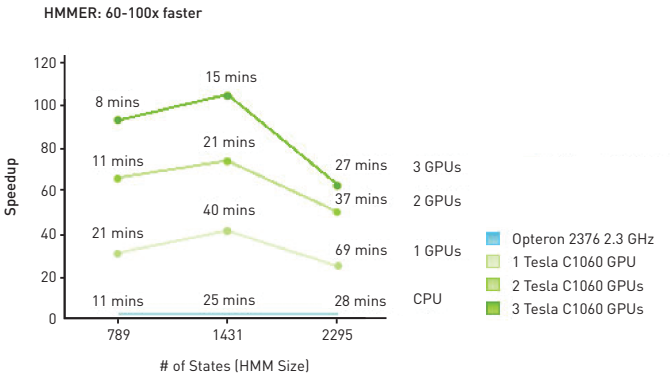


Data courtesy of Scalable Informatics and University of Buffalo

## MUMMERGPU

MUMmerGPU is a bio-informatics software for high-throughput sequence alignment using GPUs. It accelerates the alignment of multiple query sequences against a single reference sequence by taking advantage of the massively parallel CUDA architecture of NVIDIA Tesla GPUs.

MUMmerGPU (version 2.0) provides 3x-4x speedup compared to running MUMmer on CPUs. The data shown in the chart below is based on MUMmerGPU version 1.0 and will soon be updated.

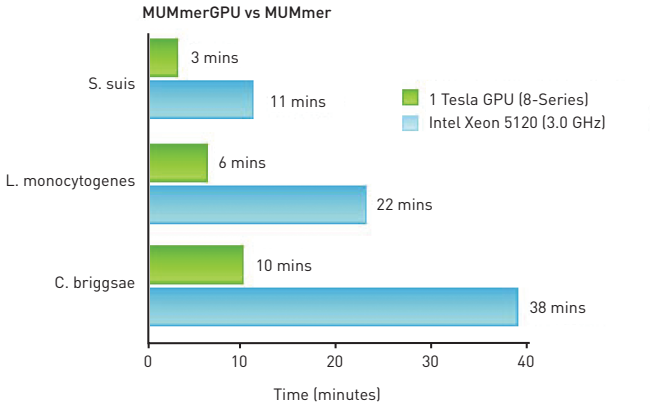


Data courtesy of University of Maryland

To learn more about the NVIDIA Tesla Bio Workbench, go to **www.nvidia.com/bio_workbench.**



NVIDIA

Partner Logo