

# NVIDIA Takes Aim at GPU Acceleration for Bioscience Applications

NVIDIA has announced the Tesla Bio Workbench, a new program designed to bring together the computational components needed to run GPU-accelerated bioscience applications. The rationale is the same one NVIDIA's been touting ever since it got into the high performance computing business: take advantage of the superior performance of the GPU in order to lower the entry point for HPC. In this case, they've assembled a GPU-centric workbench specifically designed for life science researchers and scientists.



In a nutshell, the [Tesla Bio Workbench](#) includes an array of GPU-capable bioscience codes, a community Web site for downloading the codes and providing a forum for exchanging information, and, of course, recommendations for NVIDIA Tesla GPU -equipped workstations and clusters. The strategy is to educate the biotech community that applications and hardware are here and within the reach of more researchers than ever before.

Over the past couple of years, the application set for computational biology codes that are GPU friendly has grown tremendously, thanks mainly to CUDA ports of the CPU versions of the software. This has produced a large number of popular molecular dynamics and quantum chemistry software packages that can now be run on NVIDIA GPUs. These include such codes as AMBER, GROMACS, NAMD, TeraChem, and VMD, among others. A number of bioinformatics codes like CUDA-SW++ (Smith-Waterman), GPU-HMMER, and MUMmerGPU, are also available. All of these can be downloaded via the Tesla Bio Workbench from their respective owner sites. Many of these can be had free of charge, especially if their use is limited to academic research.

The motivation behind all this is NVIDIA's recognition that computational biology is one of the lowest hanging fruits for GPU acceleration. Performance increases on the order of 10X to 100X compared to a CPU are fairly typical for these types of codes. This has not gone unnoticed. "The kind of momentum around GPUs in this domain has been perhaps the biggest and most organic that we've seen," says Sumit Gupta, NVIDIA's senior product manager for the Tesla group. According to him, a lot of biologists have turned to GPUs without any prodding from NVIDIA. The reason for this, he thinks, is that for many small and moderate-sized bio-research projects, the costs and complexity of high performance computing have become a true pain point.

The life sciences sector is already one of the largest markets for high performance computing. In 2008, 29 percent of the supercomputing cycles on TeraGrid were dedicated to bioscience applications, while another 19 percent were running related codes in chemistry and material sciences research. In the commercial realm, HPC demand is being driven by pharmaceutical companies and the emerging genomics industry in their quest for better drugs and treatments. Analyst firm IDC estimates the bioscience vertical is worth well over \$1.5 billion to HPC vendors and expanding at a CAGR of 2.6 percent. By the way, that CAGR figure is post-recession; in 2008 IDC was forecasting a growth rate of 9.3 percent. Nevertheless, the prospects for HPC in this sector are significant.

Drug discovery, in particular, is one area where HPC promises to both lower costs and accelerate the pace of research. Today the physical synthesis of drug compounds and the subsequent testing in high-throughput drug screening is both expensive and time consuming, typically representing a five-year R&D cycle. On modern HPC systems, much of this work can be simulated with molecular dynamics and quantum chemistry codes, in essence, replacing expensive labor and material costs with cheap CPU cycles.

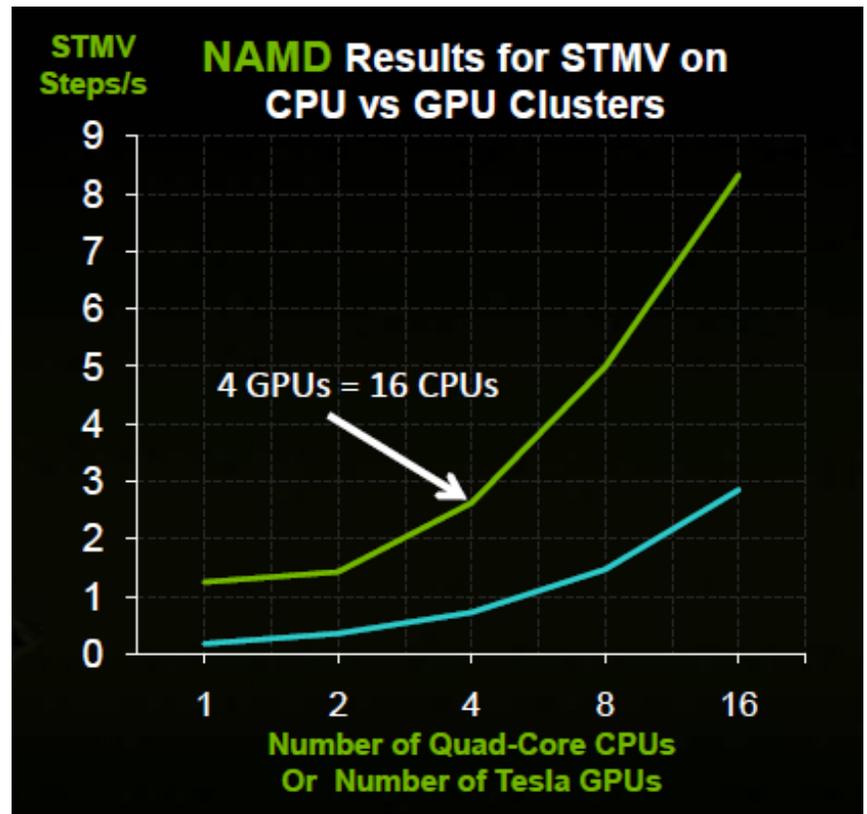
Or GPU cycles, as the case may be. NVIDIA's point with the Tesla Bio Workbench is that GPUs can make computational bioscience a much less expensive proposition than ever before. Because of the data parallel computational capabilities of the modern graphics processor, for many science applications a GPU-equipped

workstation can replace a small CPU cluster, while a moderate-sized GPU cluster can stand in for a high-end supercomputer. This lowers up-front hardware costs, energy use over the life of the system, and datacenter space.

For example, a small simulation of the satellite tobacco mosaic virus (STMV) virus using NAMD, a molecular dynamics code for biomolecular simulations, can be performed on a modern 16-CPU cluster based on quad-core x86 technology. But according to NVIDIA's Gupta, a 4-GPU workstation with a CUDA-version of NAMD will outperform that cluster, and with just a fraction of the power consumption. From the individual researcher's point of view "anything that keeps the job on the workstation is good," says Gupta.

Of course, larger simulations require more computational muscle than a workstation can provide. But since these codes tend to scale very nicely, a GPU cluster is the natural path up. "The key to acceptance here is going to be the fact that it's easy to simulate large molecules," explains Gupta. "You don't have to get time on a supercomputer, because that's too restricting." For a drug company, that means every researcher can have a GPU workstation for their own small experiments and can share a GPU cluster when they need to run a larger problem.

Commercial products resulting from GPU-powered computational biology have yet to appear. At this point the use of these methods for drug discovery at pharmaceutical companies is sporadic. And given the length of clinical trials that must follow the drug design and discovery process, Gupta thinks we probably won't begin to hear of success stories for another five years or so. For NVIDIA, the immediate challenge is to convince the biotech industry that these GPU computational tools and platforms are ready now.



Tags: [bioinformatics](#), [bioscience](#), [Computational Biology](#), [CUDA](#), [GPUs](#), [HPC](#), [life science](#), [NVIDIA](#), [Tesla Bio Workbench](#), [workbench](#)