

```
Drupal.behaviors.print = function(context)
{window.print();window.close();}>
```



# Next-Gen GPUs

July 01, 2011

## Next-Gen GPUs

By Matthew Dublin

The biggest application area for graphics processing units in the life sciences has traditionally been in accelerating molecular modeling and protein simulation software. But with the recent evolution in GPU design — where it's now common for each chip to have multiple processing cores — bioinformatics software developers are looking beyond molecular simulation to other challenging informatics areas.

"Highly parallel GPUs are becoming so ubiquitous that even my MacBook Pro laptop has 48 cores on its GPU. A natural idea is to try to apply that computing power to all the compute-intensive problems out there in bioinformatics," says Michael Schatz, an assistant professor at the Center for Bioinformatics and Computational Biology at the University of Maryland. "Some problems are amazingly well suited to the GPU, and people are reporting 100-fold improvements in run time. This performance gain is truly transformative — a 100-fold improvement condenses a week of work into less than two hours, and changes the scope of problems to consider."

[Click here to subscribe to Genome Technology](#)

In February, a team at Oxford University released STOCHSIMGPU, a tool that uses GPUs to explore stochasticity in biological systems. The open-source implementation is a GPU version of the Gillespie stochastic simulation algorithm, the logarithmic direct method, and the next reaction method, or NRM, and is reportedly 85 times faster than NRM on a CPU. Around the same time, researchers from the Hong Kong University of Science and Technology created GBOOST, a GPU-based tool for identifying gene-gene interactions in genome-wide case-control studies using Boolean operation-based screening and testing, or BOOST. GBOOST has demonstrated performance improvements of up to 40-fold speedups compared the standard BOOST on a CPU when analyzing Wellcome Trust Case-Control Consortium type 2 diabetes genome data.

## **Focus on sequencing**

Bioinformatics developers are also giving more and more attention to how GPUs can help better analyze next-gen sequencing short reads, which are often more error-prone than conventional Sanger shotgun reads. Schatz's group was one of the first to demonstrate how GPUs could be used for high-throughput sequence alignment tasks. In 2007, they released MUMmer-GPU, a GPU implementation of the MUMmer sequence alignment program that achieved 10-fold speedups over the standard CPU version. MUMmerGPU received attention from the bioinformatics community because it demonstrated that, while GPUs are normally thought of as being better with compute-intensive problems versus data-intensive problems, some memory-intensive applications actually run significantly faster on GPUs than on CPUs.

[ pagebreak ]

Developers at Carnegie Mellon University have also accelerated the ubiquitous NCBI Blast algorithm on a GPU. Their GPU-Blast, first released as an open-source tool in November 2010, reportedly achieves speedups of roughly three- to four-times that of the sequential, CPU-based Blast.

More recently, researchers at the Universität Bielefeld in Germany developed a mapping approach for complete short read alignments to microbial genomes using GPUs. Called Semiglobal Alignment of short Reads Using CUDA and NeedleMAN-Wunsch, or SARUMAN, it allows users to compute exact alignments to microbial genomes in parallel on GPUs using Nvidia's CUDA, or compute unified device architecture, technology and returns results in a time frame comparable to or faster than other approaches that use CPUs alone.

"When we started the development of SARUMAN, there were several tools available for the alignment of short reads to genomic reference sequences, but all available approaches were either heuristic and did not guarantee finding all optimal alignments, or they were too slow for the huge amounts of data produced by next-generation short read sequencing systems," says Bielefeld's Jochen Blom. "We decided to use the emerging technique of GPU programming to overcome the discrepancy between speed and accuracy."

The speedup of SARUMAN depends on the length of the sequences to be aligned. For example, for 36 basepair reads, a 25-fold speedup is achieved in comparison to a CPU implementation. For 100 base-pair reads, the speedup is decreased because fewer alignments can be computed in parallel — though a five-fold speedup is still possible. Blom is planning to enhance SARUMAN in the near future, to include support for the widely used Sequence Alignment/Map output format and a paired-end mode for identifying biologically correct mappings among co-optimal results. His team also aims to provide alternative alignment

support for color-space data generated on the SOLiD platform and plans to evaluate how SARUMAN can be efficiently applied to large eukaryotic genomes.

## Correction

Last year, Yongchao Liu, a graduate student at Nanyang Technological University in Singapore, released DecGPU, the first parallel and distributed error correction algorithm for high-throughput short reads that uses both the CUDA and Message Passing Interface, or MPI, programming models. The open source solution employs a "hybrid" computing model that combines GPUs and CPUs with MPI to distribute the computation across multiple compute nodes, enabling users to perform error correction of large-scale high-throughput short reads data sets.

"The rapid evolution of modern, many-core GPU computing architecture has changed the high-performance computing world," Liu says. "Their powerful compute capabilities have been demonstrated to reduce the execution time in a range of bioinformatics applications, such as sequence alignment and motif discovery, so it is viable to use GPU computing to accelerate the error correction process."

In a March *BMC Bioinformatics* paper, Liu and his colleagues describe using DecGPU on simulated and real sequence data sets that exceeded the capacity of existing error correction algorithms, like hSHREC and CUDA-EC. DecGPU can also be combined in an informatics analysis workflow with the popular short read sequence assemblers Velvet and ABySS to improve de novo assembly, and it can also be used with other de Bruijn graph-based assemblers. Liu says that users do not need to be familiar with either CUDA or MPI to take advantage of DecGPU.

[ pagebreak ]

## Backtrack

In May, bioinformatics software developers at the Poznan University of Technology in Poland designed a unique GPU-based method for amino acid sequence alignment using the Needleman-Wunsch and Smith-Waterman algorithms. Unlike MUMerGPU and SARUMAN, which align short read sequences to a reference genome, the team's gpu-pairAlign method produces alignments for every possible sequence pair in given input set. "Before our solution was introduced there were a few implementations of pairwise sequence alignment taking advantage of modern GPUs, but the vast majority of them could only compute the alignment score, whereas the sequence alignment itself was omitted," says Pawel Wojciechowski, an assistant professor at Poznan.

According to Wojciechowski, gpu-pairAlign opens up a wide range of new possibilities for accelerating a -variety of bioinformatics algorithms on GPUs. "The first and the most obvious application of our solution is in the progressive methods for the Multiple Sequence Alignment, or MSA, problem as they require the alignment of every possible pair of input sequences to be computed," he says. "This was one of the motivations that led us to development of this algorithm and we are currently finishing work on G-Coffee, which is an attempt to implement T-Coffee — a well-known algorithm for MSA. -G-Coffee uses in its first step, the pairwise alignment that is described in the paper, and, as such, is immensely fast according to our first tests."

While development environment toolkits have come a long way in the last few years, Schatz says that GPU code development is no place for coding novices. "The development environment has matured considerably," he says. "However, I would still classify it as a platform for very skilled developers."

## Related Stories

- [Big Bioinformatics Hopes](#)

- December 3, 2012 / Genome Technology
- Bioinformatics in Your Pocket  
November 1, 2012 / Genome Technology
- An Accelerated API  
October 1, 2012 / Genome Technology
- Game On  
September 3, 2012 / Genome Technology
- Cloudy Commerce  
July 1, 2012 / Genome Technology

footer