



GPU-Accelerated Computing Reaches Next Generation of Programmers With Python Support of NVIDIA CUDA

Python Productivity With GPU Performance Enables Coming Wave of HPC and Data Analytics Applications

SAN JOSE, CA -- **GTC 2013** -- The growing ranks of programmers using the Python open-source language can now take full advantage of GPU acceleration for their high performance computing (HPC) and big data analytics applications by using the [NVIDIA@ CUDA@ parallel programming model](#), NVIDIA today announced.

Easy to learn and use, Python is among the top 10 programming languages with more than three million users. It enables users to write high-level software code that captures their algorithmic ideas without delving deep into programming details. Python's extensive libraries and advanced features make it ideal for a broad range of HPC science, engineering and big data analytics applications.

Support for NVIDIA CUDA parallel programming comes from NumbaPro, a Python compiler in the new Anaconda Accelerate product from Continuum Analytics.

"Hundreds of thousands of Python programmers will now be able to leverage GPU accelerators to improve performance on their applications," said Travis Oliphant, co-founder and CEO at Continuum Analytics. "With NumbaPro, programmers have the best of both worlds: they can take advantage of the flexibility and high productivity of Python with the high performance of NVIDIA GPUs."

Expanded Access to Accelerated Computing Via LLVM

This new support for GPU-accelerated application development is the result of NVIDIA's [contribution of the CUDA compiler source code](#) into the core and parallel thread execution backend of [LLVM](#), a widely used open source compiler infrastructure.

Continuum Analytics' Python development environment uses LLVM and the NVIDIA CUDA compiler software development kit to deliver GPU-accelerated application capabilities to Python programmers.

The modularity of LLVM makes it easy for language and library designers to add support for GPU acceleration to a wide range of general-purpose languages like Python, as well as to domain-specific programming languages. LLVM's efficient just-in-time compilation capability lets developers compile dynamic languages like Python on the fly for a variety of architectures.

"Our research group typically prototypes and iterates new ideas and algorithms in Python and then rewrites the algorithm in C or C++ once the algorithm is proven effective," said Vijay Pande, professor of Chemistry and of Structural Biology and Computer Science at Stanford University. "CUDA support in Python enables us to write performance code while maintaining the productivity offered by Python."

Anaconda Accelerate is available for Continuum Analytics' [Anaconda](#) Python offering, and as part of the [Wakari](#) browser-based data exploration and code development environment.

About CUDA

CUDA is a [parallel computing platform and programming model](#) developed by NVIDIA. It enables dramatic increases in computing performance by harnessing the power of GPUs. With more than 1.7 million downloads, supporting more than 220 leading engineering, scientific and commercial applications, the CUDA programming model is the most popular way for developers to take advantage of GPU-accelerated computing.

More information about NVIDIA CUDA GPUs is available at the NVIDIA [Tesla@ GPU website](#). To learn more about CUDA or download the latest version, visit the [CUDA website](#).

About NVIDIA

Since 1993, [NVIDIA](#) (NASDAQ: NVDA) has pioneered the art and science of visual computing. The company's technologies are transforming a world of displays into a world of interactive discovery -- for everyone from gamers to scientists, and consumers to enterprise customers. More information at <http://nvidianews.nvidia.com> and <http://blogs.nvidia.com>.

Certain statements in this press release including, but not limited to, statements as to: the impact and benefits of GPU accelerators and the NVIDIA CUDA parallel programming model; and the effects of the company's patents on modern computing are forward-looking statements that are subject to risks and uncertainties that could cause results to be materially different than expectations. Important factors that could cause actual results to differ materially include: global economic conditions; our reliance on third parties to manufacture, assemble, package and test our products; the impact of technological

development and competition; development of new products and technologies or enhancements to our existing product and technologies; market acceptance of our products or our partners' products; design, manufacturing or software defects; changes in consumer preferences or demands; changes in industry standards and interfaces; unexpected loss of performance of our products or technologies when integrated into systems; as well as other factors detailed from time to time in the reports NVIDIA files with the Securities and Exchange Commission, or SEC, including its Form 10-K for the fiscal period ended January 27, 2013. Copies of reports filed with the SEC are posted on the company's website and are available from NVIDIA without charge. These forward-looking statements are not guarantees of future performance and speak only as of the date hereof, and, except as required by law, NVIDIA disclaims any obligation to update these forward-looking statements to reflect future events or circumstances.

© 2013 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, CUDA and Tesla are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated. Features, pricing, availability and specifications are subject to change without notice.

Hector Martinez
Corporate Communications
+1-408-486-3443

hmarinez@nvidia.com

Ken Brown
Corporate Communications
+1-408-486-2626

kebrown@nvidia.com