



NVIDIA, Cray, PGI, CAPS Unveil 'OpenACC' Programming Standard for Parallel Computing

Directives-Based Programming Makes Accelerating Applications Using CPUs and GPUs Dramatically Easier Than Modifying Underlying Code

SEATTLE, WA -- **SC11** -- In an effort to make it easier for programmers to take advantage of parallel computing, NVIDIA, Cray Inc., the Portland Group (PGI), and CAPS enterprise announced today a new parallel-programming standard, known as OpenACC™.

Initially developed by PGI, Cray, and NVIDIA, with support from CAPS, OpenACC is a new open parallel programming standard designed to enable the millions of scientific and technical programmers to easily take advantage of the transformative power of heterogeneous CPU/GPU computing systems.

OpenACC allows parallel programmers to provide simple hints, known as "directives," to the compiler, identifying which areas of code to accelerate, without requiring programmers to modify or adapt the underlying code itself. By exposing parallelism to the compiler, directives allow the compiler to do the detailed work of mapping the computation onto the accelerator.

OpenACC is anticipated to benefit a broad range of programmers working in chemistry, biology, physics, data analytics, weather and climate, intelligence, and many other fields. Existing compilers from Cray, PGI and CAPS are expected to provide initial support for the OpenACC standard beginning in the first quarter of 2012. The OpenACC standard is fully compatible and interoperable with the NVIDIA® CUDA® parallel programming architecture, which is designed to allow detailed control over the accelerator for maximum performance tuning.

Directives provide a common code base that is multi-platform and multi-vendor compatible, offering an ideal way to preserve investment in legacy applications by enabling an easy migration path to accelerated computing. Based on recent data, the majority of developers have reported 2x to 10x increases in application performance in as little as two weeks when using existing directive-based compilers.¹

Quotes

"Compiler directives are integral as we continue the build-out and deployment of the Titan GPU-accelerated supercomputer," said Buddy Bland, Titan project director at Oak Ridge National Laboratory, referring to a system expected to be the world's fastest supercomputer. "Our ultimate goal is to have all Titan supercomputing code run on hybrid CPU/GPU nodes, and OpenACC will enable programmers to develop portable applications that maximize the performance and power efficiency benefits of this architecture."

"OpenACC represents a major development for the scientific community," said Jeffrey Vetter, joint professor in the Computational Science and Engineering School of the College of Computing at Georgia Institute of Technology. "Programming models for open science by definition need to be flexible, open and portable across multiple platforms; OpenACC is well designed to fill this need. It provides a valuable new tool to empower the vast numbers of domain scientists who could benefit from application acceleration, but who may not have the funding or expertise to port their code to emerging architectures."

"I am enthusiastic about the future of accelerator technologies," said Michael Wong, CEO of the OpenMP Architecture Review Board. "The OpenACC announcement highlights the technically impressive initiative undertaken by members of the OpenMP Working Group on Accelerators. I look forward to working with all four companies within the OpenMP organization to merge OpenACC with other ideas to create a common specification which extends OpenMP to support accelerators. We look forward to incorporating accelerator support with the full support of all OpenMP members in a future version of the OpenMP specification."

More information about OpenACC, as well as the OpenACC specification, is available today at www.OpenACC-standard.org. OpenACC is based largely on the PGI and Cray accelerator programming models. NVIDIA, Cray, PGI, and CAPS are members of the OpenMP subcommittee on accelerators, and intend to work within that organization to come to a common standard.

Developers interested in trying directives-based programming can obtain a one-month free trial of the PGI Accelerator Fortran and C compilers at NVIDIA's new [2x in 4 weeks web site](#), or by contacting Cray.

About Cray Inc.

As a global leader in supercomputing, Cray provides highly advanced supercomputers and world-class services and support to government, industry and academia. Cray technology is designed to enable scientists and engineers to achieve

remarkable breakthroughs by accelerating performance, improving efficiency and extending the capabilities of their most demanding applications. Cray's Adaptive Supercomputing vision is focused on delivering innovative next-generation products that integrate diverse processing technologies into a unified architecture, allowing customers to surpass today's limitations and meeting the market's continued demand for realized performance. Go to www.cray.com for more information.

About The Portland Group (PGI)

The Portland Group, a wholly-owned subsidiary of STMicroelectronics (NYSE: STM), is the premier supplier of high-performance parallel Fortran, C, and C++ compilers and tools for workstations, servers, and clusters based on x64 processors from Intel and AMD, and GPUs accelerators from NVIDIA. Further information on The Portland Group products can be found at www.pgroup.com, by calling Sales at (503) 682-2806, or by email to sales@pgroup.com

About CAPS enterprise, The Many-Core Programming Company

CAPS is a major supplier of solutions dedicated to application migration and deployment on many-core processors. CAPS global solutions for many-core lead the developer to performance by providing top-of-the-range technology (HMPP hybrid compiler and wizard), code porting methodology and ecosystem. Its directive-based & multi-target HMPP™ compiler enables developers to safely move to hybrid CPU / GPU model and leverage the computing power of stream processors without the pain associated to GPU programming. HMPP™ is offered within CAPS DevDeck™ package: an ALL-IN-ONE multi-level suite for many-core application definition, porting and optimization with tools, methodology and resources.

About NVIDIA

[NVIDIA](http://www.nvidia.com) (NASDAQ: NVDA) awakened the world to computer graphics when it invented the [GPU](#) in 1999. Today, its [processors](#) power a broad range of products from [smart phones](#) to [supercomputers](#). NVIDIA's [mobile processors](#) are used in [cell phones](#), [tablets](#) and [auto infotainment systems](#). [PC gamers](#) rely on GPUs to enjoy spectacularly immersive worlds. Professionals use them to create visual effects in movies and design everything from golf clubs to jumbo jets. And researchers utilize GPUs to advance the frontiers of science with [high-performance computing](#). The company holds more than 2,100 patents worldwide, including ones covering ideas essential to modern computing. For more information, see www.nvidia.com.

(1) Source -- PGI customers: www.nvidia.com/2xin4weeksstories

Tags / Keywords:

NVIDIA, CUDA, Tegra, GPU, GPU computing, supercomputing, parallel computing, GPGPU, high performance computing, HPC, programmer, directives, compiler, developers, research, scientific computing, Cray, PGI, The Portland Group, CAPS

Certain statements in this press release including, but not limited to statements as to: the effects, benefits and impact of OpenACC, the NVIDIA CUDA architecture and parallel computing; NVIDIA's efforts as part of the OPEN MP subcommittee to create a common standard for accelerators; 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-Q for the fiscal period ended July 31, 2011. 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.

© 2011 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, CUDA and OpenACC 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.

Ken Brown
Corporate Communications
+1-408-486-2626
kebrown@nvidia.com