NVIDIA dresses up CUDA parallel computing platform

January 28 2012, by Nancy Owano

(PhysOrg.com) -- This week's NVIDIA announcement of a dressed up version of its CUDA parallel computing platform is targeted as a good news message for engineers, biologists, chemists, physicists, geophysicists, and other researchers on fast-track computations using GPUs. The new version features an LLVM (low-level virtual machine)-based CUDA compiler, new imaging and signal processing functions added to the NVIDIA Performance Primitives library and a redesigned Visual Profiler with automated performance analysis and expert guidance. NVIDIA says the new enhancements are ways to advance simulations and computational work for these users.

CUDA is a parallel computing platform and programming model that was created by NVIDIA. The company promotes CUDA as the pathway to achieve dramatic increases in computing performance by harnessing the power of the graphics processing unit (GPU). According to the

company, with CUDA, a developer can send C, C++ and Fortran code straight to the GPU; no assembly language is required.

Generally, developers at scientific companies look to GPU computing for speeding up applications for scientific and engineering computing. With this approach, GPU-accelerated applications run the sequential part of their workload on the CPU while accelerating parallel processing on the GPU.

The company notes that a combined team from Harvard Engineering, Harvard Medical School and Brigham & Women's Hospital have used GPUs to simulate blood flow and identify hidden arterial plaque without having to use invasive imaging techniques or exploratory surgery. At NASA, where computer models identify ways to alleviate congestion and keep traffic moving efficiently, a NASA team has made use of GPUs to gain on performance and reduce analysis time.

"When we started creating CUDA, we had a lot of choices for what we could build. The key thing customers said was they didn't want to have to learn a whole new language or API," said Ian Buck, general manager at NVIDIA. "Some of them were hiring gaming developers because they knew GPUs were fast but didn't know how to get to them." He said NVIDIA wanted to provide a solution that could be learned in one session and outperform CPU code.

The revised CUDA parallel computing platform carries three main changes that are supposed to make parallel programming with GPUs easier and faster.

The Visual Profiler with a few clicks is said to deliver an automated performance analysis of the user's application. It highlights problem areas and shows links to suggestions for improvement. This eases application acceleration. Also, NVIDIA is transitioning to new LLVM

based compiler technology, The compiler is based on the LLVM open-source compiler infrastructure, and can deliver an increase in application performance. (LLVM is an umbrella project that hosts and develops a set of close-knit toolchain components such as assemblers, compilers and debuggers. The LLVM project started in 2000 at the University of Illinois at Urbana-Champaign.)

New imaging and signal processing functions are increasing the size of the NVIDIA Performance Primitives (NPP) library. The updated NPP library can be used for image and signal processing algorithms, ranging from basic filtering to advanced workflows.

NVIDIA unveiled CUDA in 2006, announcing CUDA as the world's first solution for general-computing on GPUs. NVIDIA cites some examples on its site of CUDA's user base today. In the consumer market, nearly every major consumer video application has been, or will soon be, accelerated by CUDA, including products from Adobe, Sony , Elemental Technologies, MotionDSP and LoiLo, according to NVIDIA. In scientific research. CUDA accelerates AMBER, a molecular dynamics simulation program used by researchers to speed up new drug discovery.

More information: www.nvidia.com/object/cuda_home_new.html

© 2011 PhysOrg.com