

Software tool helps tap into the power of graphics processing

May 17 2010

Today's computers rely on powerful graphics processing units (GPUs) to create the spectacular graphics in video games. In fact, these GPUs are now more powerful than the traditional central processing units (CPUs) or brains of the computer. As a result, computer developers are trying to tap into the power of these GPUs. Now a research team from North Carolina State University has developed software that could make it easier for traditional software programs to take advantage of the powerful GPUs, essentially increasing complex computing brainpower

Taking advantage of a GPU's processing ability is a big deal, because of the amount of <u>computing power</u> a <u>GPU</u> contains. The CPU from an average computer has about 10 gigaflops of computing power - or 10 billion operations per second. That sounds like a lot until you consider that the GPU from an average modern computer has 1 teraflop of computing power - which is 1 trillion operations per second.

But using a GPU for general computing functions isn't easy. The actual architecture of the GPU itself is designed to process graphics, not other applications. Because GPUs focus on turning data into millions of pixels on a screen, the architecture is designed to have many operations taking place in isolation from each other. The operation telling one pixel what to do is separate from the operations telling other pixels what to do. This hardware design makes graphics processing more efficient, but presents a stumbling block for those who want to use GPUs for more complex computing processes.



A research team from NC state has developed software that could make it easier for traditional software programs to take advantage of GPUs. The research was funded by the National Science Foundation.

"We have developed a <u>software tool</u> that takes computer program A and translates it into <u>computer program</u> B - which ultimately does the same thing program A does, but does it more efficiently on a GPU," says Dr. Huiyang Zhou, an associate professor of electrical and computer engineering at NC State and co-author of a paper describing the research. This sort of translation tool is called a compiler.

Program A, which is the user-provided input, is called a "nad've" version - it doesn't consider GPU optimization, but focuses on providing a clear series of commands that tell the computer what to do. Zhou's compiler software takes the nad've version and translates it into a program that can effectively utilize the GPU's hardware so that the program operates a lot more quickly.

Zhou's research team tested a series of standard programs to determine whether programs translated by their compiler software actually operated more efficiently than code that had been manually optimized for GPU use by leading GPU developers. Their results showed that programs translated by their compiler software ran approximately 30 percent more quickly than those optimized by the GPU developers.

"Tapping into your GPU can turn your personal computer into a supercomputer," Zhou says.

More information: The paper, "A GPGPU Compiler for Memory Optimization and Parallelism Management," was co-authored by Zhou, NC State Ph.D. student Yi Yang, and University of Central Florida Ph.D. students Ping Xiang and Jingfei Kong. The paper will be presented June 7 at the Programming Language Design and Implementation



conference in Toronto.

Provided by North Carolina State University

Citation: Software tool helps tap into the power of graphics processing (2010, May 17) retrieved 1 May 2024 from <u>https://phys.org/news/2010-05-software-tool-power-graphics.html</u>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.