



FOR IMMEDIATE RELEASE

Contacts:

Pico Computing – Mark Hur, Director of Sales and Marketing, (206) 283-2178 x107
mhur@picocomputing.com

Impulse Accelerated Technologies – Brian Durwood, 425.605.9543 x109
brian.durwood@impulseC.com

Pico Computing and Impulse Demonstrate FPGA Cluster Acceleration of C-Language Applications at SC07

*Impulse C™ compiler used to create high-performance, low-power image processing
and scientific demonstrations*

November 7, 2007 – Reno, Nevada. Pico Computing (Seattle, WA) and Impulse Accelerated Technologies (Kirkland, WA) today announced that the two companies will be presenting joint demonstrations at the International Conference for High Performance Computing (SC07) being held November 10 to 16, 2007 in Reno, Nevada.

The demonstrations planned for SC07 include a “Where’s Waldo” image recognition algorithm and an N-body astrophysics simulation. Both of these demonstration applications have been implemented using the Impulse C tools and run on a Pico Computing SuperCluster™ 84-FPGA array. Pico Computing will also be demonstrating an FPGA-accelerated, brute-force WPA cracking algorithm.

“Impulse C allows algorithm developers to more quickly generate software/hardware applications for our SC3 SuperCluster,” said Dr. Robert Trout, President and founder of Pico Computing. “For high-performance computing applications, access to tools like Impulse C is a critical enabler.”

“The SuperCluster is a new and exciting approach to hardware acceleration,” said David Pellerin, CTO of Impulse. “By combining a large number of PCI Express connected FPGAs in a relatively small form-factor chassis, Pico Computing is offering an enormous amount of computing performance per watt of power.”

According to Dr. Trout, the power/performance of the SuperCluster is stunning. A single FPGA module (featuring a relatively modest Virtex-5 LX50 FPGA device) can demonstrate random number generation speed improvements of 12X or better over a standard dual-core processor. In the Pico Computing SC3 SuperCluster, 84 of these modules demonstrate performance comparable to a cluster of 1,000 dual-core processors, while using dramatically less power. In fact, the entire 84-FPGA SuperCluster is capable of operating at-speed using a standard 600W power supply.

Finding Waldo

The “Where’s Waldo” demonstration highlights the potential of FPGAs for acceleration of complex image processing tasks, using a cluster of FPGAs. This demonstration uses the popular children’s book to illustrate the challenges associated with identifying someone or something among a crowd of similar images. The SC3 SuperCluster with 84 FPGAs capitalizes on the parallel nature of the algorithm. Impulse C was used to develop the required FPGA image processing algorithms.

The demonstration application extracts distinctive features of the target image using a Scale Invariant Feature Transform (SIFT) method. The algorithm searches for corresponding features in a video stream, while enforcing the consistency across all feature matches. To ensure precision the algorithm provides a measure of the certainty for each match, for example reporting an 85% chance that Waldo is at a specific location. The success of this demonstration has clear implications for low-power, real-time defense and security applications.

Accelerating Astrophysical Simulations

The goal of the N-body simulation is to model and calculate the gravitational forces between thousands of planets, stars, galaxies, and other objects. N-body simulations are computationally intensive but are regularly used by scientists to understand how galaxies and planets are formed and evolve over time. The computation required is N^2 , with N representing the number of bodies being modeled. The gravitational force on each body is calculated by summing the force between that body and every other body in the system. For example in a simulation of the solar system, the movement of Earth is calculated by summing the gravitational pull of the Sun, other planets, comets, the Earth’s Moon, etc). This is a complex, floating-point problem that is highly parallelizable, and hence a perfect candidate for FPGA acceleration. Impulse C was also used to develop this algorithm, using the streaming, multiple-process features of the Impulse C programming tools.

Cracking WPA Security

WPA is common security algorithm used to secure wireless access points. WPA employs PBKDF2, which runs the SHA1 algorithm thousands of times to convert a password into a key. The key is then used to encrypt the wireless network.

To crack passwords on the WPA network, an authentication session must first be captured. Once this is captured, different passwords can be tried by running through the PBKDF2 (brute-force) function and verifying if the password is correct by verifying against the captured data. This method requires an enormous number of iterations but is highly parallelizable. The result is WPA cracking that is hundreds or thousands of times faster than would be achievable using software-only methods.

About Pico Computing

Pico Computing is a pioneer in rapidly deployable, FPGA-based, reconfigurable and scalable computing systems. Pico Computing platforms are based on Xilinx Virtex-4 and Virtex-5 FPGAs. These unique platforms provide designers and developers the ability to design, develop and debug directly in the CardBus or ExpressCard (PCI Express) slot of a standard laptop or desktop computer. Pico Computing provides all source code (HDL and software) necessary to implement base system-on-chip computing functionality. Pico Computing also provides SuperCluster platforms for use in high performance computing applications. These clusters are based on standard, off-the-shelf computers and can host up to 84 FPGA blades. For more information about Pico Computing FPGA platforms, visit www.PicoComputing.com.

About Impulse

Impulse C allows software application developers to rapidly and cost-effectively move applications originating in C-language to FPGA coprocessors. The Impulse C tools enable partitioning, optimization, and FPGA hardware generation. Interconnections between FPGA coprocessors and host processors are generated automatically for selected platforms including the Pico Computing cards and clusters. For more information about advanced C-to-FPGA design methods and tools, visit www.ImpulseC.com.

###