For further information, contact:
John Melonakos
AccelerEyes
75 5th Street NW, Suite 204, Atlanta, GA 30308
+1 (800) 570-1941
john.melonakos@accelereyes.com

FOR IMMEDIATE RELEASE:

AccelerEyes' Jacket product family supports the latest NVIDIA Fermi GPUs

Atlanta, GA., Jul 13, 2010 – AccelerEyes today announced that its Jacket software platform for MATLAB®, including its new 1.4 release, will support the latest NVIDIA graphics processing units (GPU) based on the Fermi architecture (Tesla 20-series and GeForce GTX 4xx-series). NVIDIA's release of the Fermi architecture brings with it 448 computational cores, increased IEEE-754 floating-point arithmetic precision, error-correcting memory for reliable computation, and enhanced memory caching mechanisms.

AccelerEyes develops Jacket, a software platform that delivers GPU computing power to desktop users of MATLAB and other very high level languages. It enables faster prototyping and problem solving across a range of government, manufacturing, energy, media, biomedical, financial, and scientific research applications. The Jacket software platform enables accelerated double-precision performance for common arithmetic and linear algebra functionality on the new NVIDIA hardware based on the FERMI architecture.

"With this release of Jacket and with the familiar MATLAB environment, domain experts can create highly optimized heterogeneous applications for the latest CPUs from Intel and AMD while also leveraging the latest generation of GPUs from NVIDIA," said John Melonakos, CEO of AccelerEyes. "Efficiently using all available host cores for certain parts of an application while accelerating other portions on GPUs is the key to squeezing maximum performance out of today's GPU-enabled workstations, servers, clusters, and cloud services. With Fermi's improvement in double-precision performance, we expect a big increase in the number and type of applications that benefit from GPU acceleration."

In addition to performance improvements provided by the FERMI hardware:

- AccelerEyes has dramatically improved the performance of Jacket's on-the-fly compilation technology.
- Matrix multiply has seen significant performance improvement ranging from 20-60% speedups.
- Jacket 1.4 provides direct access to the NVIDIA Performance Primitives enabling new Image Processing functionality such as ERODE and DILATE.
- GFOR-support added for various functions such as the direct system solvers with MLDIVIDE and the FFT function.

See the company website and the v1.4 release notes for the full list of enhancements in this release.

Initial customer testing has shown what they are calling "pretty impressive" results. Comparing a 1265 MHz Tesla C1060 running Jacket 1.3 with CUDA 2.3 and a 1025 MHz Fermi C2050 running Jacket 1.4 with CUDA 3.0, the results are as follows:

- Multiply-add (MAD), double precision on GPU, Fermi is 1.4x faster than Tesla for 4000x4000 matrices (2s vs 2.8s)
- 2D 2-point (GRADIENT), single precision on GPU, Fermi is 3x faster than Tesla for 4000x4000 matrices (0.3s vs 0.9s)
- 2D FFT (FFT2), double precision on GPU, Fermi is 2x faster than Tesla for 4000x4000 matrices (0.1s vs 0.2s)

Jacket automatically connects M-code to GPUs through its dynamic compiler technology allowing users to accelerate their applications across any CUDA-capable GPUs no matter which card was used for development – the same Jacket application scales effortlessly across GeForce, Tesla C1060, or the new Tesla C2050 without modification. Jacket eliminates the need to manually port MATLAB code to C, C++, or CUDA to run on GPU-based

workstations and clusters – a productivity step that can take months to years to complete for computationally intensive problems. Support for new NVIDIA GPU platforms in Jacket 1.4 extends across Linux, Windows, and MacOS.

Pricing and availability

Jacket 1.4 is now available for download on the AccelerEyes website. Pricing for a Jacket base license with support for a single GPU is $999US, $800US, and $350US for commercial, government/research, and academic customers, respectively. AccelerEyes provides 12 months of software maintenance and updates with each software license. Volume packages and development bundles are also now available at special price points for a limited time only.

About AccelerEyes

AccelerEyes launched in 2007 to commercialize Jacket, the first software platform to deliver productivity in GPU computing. With advanced language processing and runtime technology to transform CPU applications to high performance GPU codes, Jacket extends from desktop workstation performance to also fully leverage GPU clusters. Based in Atlanta, GA., the privately held company markets Jacket for a range of defense, intelligence, biomedical, financial, research, and academic applications. Additional information is available at www.accelereyes.com.