The World's First Commercially-Available Stream Processor: Architecture, Algorithms and Benchmark Results

Simon McIntosh-Smith
ClearSpeed Technology, Ltd.
Email Address: simon@clearspeed.com

Ron Bell
AWE - Aldermaston
Email Address: ron.bell@awe.co.uk

Abstract:

This briefing describes new applications for ClearSpeed's CS301 device, the first commercially available stream processor. Launched in October 2003, the CS301 is an ultra-high performance next-generation Single-Instruction/Multiple-Data (SIMD) stream processor, delivering 25 GFLOPS and 12.8 GMACS at 1.8 Watts. The CS301's low-power, Multi-Threaded Array Processor (MTAP) architecture scales to hundreds and ultimately thousands of processing elements, each with both floating point and integer hardware, capable of data parallel processing on image and signal processing applications as well as for compression, encryption, search, and general sensor processing applications. The processor is supported by a flexible development environment, including assembly language and C-based language support, as well as a cycle accurate simulator, with plans to develop industry standard API Libraries such as L3 BLAS and FFTW. This new class of stream processor has been shown to provide ten to one hundred times the overall performance of PowerPC or Pentium-based architectures, especially when performing image and signal processing functions, such as FFTs or filters. In general, the architecture has been shown to provide significant throughput, size, and power advantages for embedded processing applications.

AWE Aldermaston has been investigating potential uses for CS301-class processors in its key algorithms and applications. AWE further optimised fast math library routines on the CS301 for SGEMM – a single precision floating point matrix multiply, verifying the CS301's record-breaking math performance. AWE took matrix mutiply from 5 GFLOPS sustained to over 12 GFLOPS sustained on a single CS301. AWE is performing ongoing work exploring the acceleration potential of the CS301 for several in-house and 3rd party scientific codes, such as DL-POLY.

CS301-based accelerator boards having been shipping since January 2003 and multiple algorithms have been ported, with more underway. A dual-processor PCI-based development card is available from ClearSpeed, providing a total of 50 GFLOPS of compute performance, for a total maximum power consumption for the board of 10 Watts. Single systems containing up to 5 boards have been demonstrated for a total compute of 500 GFLOPS and capable of 1 Million FFTs per second (1K complex single precision floating-point). This level of compute density has never before been commercially available, with the CS301 delivering more than 10 GFLOPS per Watt.

In the first half of this briefing ClearSpeed will present performance results from the numerous algorithms and applications that have or are being ported to the CS301 stream processor. These include numerous sizes of FFTs and FIR filters which efficiently utilize the architecture and floating point per PE hardware to gain exceptional performance at very low power dissipation levels. We will include an update on improvements to work previously announced jointly with Lockheed Martin at HPEC03 on pulse compressions for radar (FFT – Complex Multiply – IFFT). The results to be reported are significantly higher than other industry standard processing and DSP platforms. New work on other transforms, such as DCTs, will also be presented. In addition, results from work to develop a Level 3 BLAS (Basic Linear Algebra Subprograms) library will be reported, including performance of certain vector and matrix operations, such as matrix multiplication and matrix inversion, including descriptions of the algorithms required on this high-performance, highly parallel architecture.

In the second half of this briefing AWE Aldermaston will present its work to benchmark the CS301. The briefing will include descriptions of optimizations to a fast matrix multiply algorithm for the MTAP streaming architecture, improving performance from 5 GFLOPS to over 12 GFLOPS on the CS301. AWE will also describe its investigations into using the CS301 to accelerate certain applications used in-house, such as the materials science code DL-POLY.