Numerical libraries for the Epiphany™ architecture

Insight™

The optimized Insight libraries provide numerical signal processing functions for the Epiphany architecture

The functions are optimized for the dual-issue FPU / ALU computational units of the Epiphany architecture that makes full use of the 64 32-bit registers and instruction pipelining to provide maximum performance. The Insight libraries include the most commonly used signal processing algorithms: FFTs, FIR filters, matrix multiplication (real and complex) and many others.

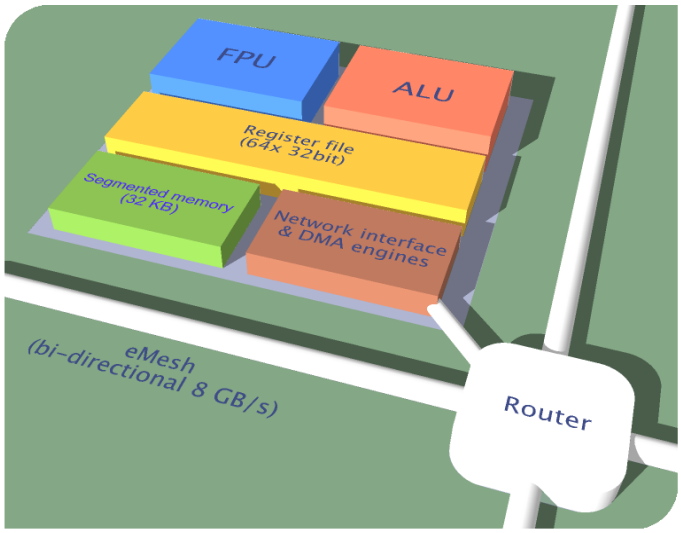

BittWare's Anemone

The Anemone floating-point FPGA co-processor is available from BittWare Inc. with 16 Epiphany cores. The Insight library functions are designed to run over the many-cores of the Anemone co-processor and the numerical functions are designed as the building blocks for larger application designs. Insight includes a utility library that provides the necessary functional glue within the many-core processing environment of the Anemone eMesh. The libraries include single-core examples for each function and multi-core examples to demonstrate the high performances of the architecture.

Performance per watt

BittWare's Anemone is designed as a high-performance floating-point co-processor consuming ultra-low power. Each Epiphany core within the Anemone co-processor is capable of performing two floating-point operations and one memory access each cycle. At a 1 GHz core frequency this equates to a theoretical peak performance of 2 GFLOPS for each core and 32 GFLOPS on the 16-cores of the Anemone co-processor. The Insight matrix-matrix multiplication runs up to 90% of this peak performance.

White papers

A number of white papers are available for download that describe the functional designs of the Insight library and focus on how these algorithms operate over the many-cores of the Anemone's eMesh.
http://www.bittware.com/media/lit/whitepaper.cfm

Paralant

BLAS library

The BLAS (Basic Linear Algebra Subprograms) library routines provide standard building blocks for performing basic vector and matrix operations. The Level 1 BLAS perform scalar, vector and vector-vector operations, the Level 2 BLAS perform matrix-vector operations, and the Level 3 BLAS perform matrix-matrix operations. Because the BLAS are efficient, portable, and widely available, they are commonly used in the development of high quality linear algebra software, LAPACK (Linear Algebra PACKage), for example. BLAS is highly scalable and is the most common method of parallelizing algorithms over many-cores.
http://www.netlib.org/blas

FFT library

The libfft.a library provides both in-place and out-of-place radix-2 and radix-4 FFT functions. The large register files allow for a very efficient radix-4 execution and gives 35% performance improvement over the radix-2 version. The radix-2 is provided for the greater flexibility in available sizes. The out-of-place FFT functions are provided to cater for multi-core implementations where data is optimally read from local memory and written out onto the eMesh for inter-core processing. Examples of single-core and multii-core FFTs are provided.

Function		Size	Cycles	GFLOPS (@1GHz)
Radix-4 FFT	forward	1024	41,807	1.04
	inverse		43,661	1.00
Matrix-matrix multiply	real	40 by 40 by 40	72,290	1.77
	complex	24 by 24 by 32	87,900	1.68
Matrix-vector multiply	real	40 by 40	4,580	0.70
	complex	24 by 24	3,927	1.17
FIR	real	896 by 16-taps	18,525	1.55
	complex		63,429	1.81
Dot-product	real	1024	2,980	0.69
	complex		5,281	1.55

BLAS level-1

sdot	cdotc
sscal	cdotu
saxpy	cscal
scopy	csscal
sswap	caxpy
sasum	ccopy
isamax	cswap
	scasum
	icamax

BLAS level-2

sgemv	cgemv

BLAS level-3

sgemm	cgemm

FFTs

r2fft	4fft
r2ifft	r4ifft

Filter

sfir	cfir

Utility functions

mailbox	barrier

Table 1. Insight library functionality [release 1.1]

The Insight libraries for the Epiphany architecture are produced by Paralant Ltd. [www.paralant.com] and sold by BittWare Inc. [www.bittware.com] for the Anemone FPGA Co-Processor.