

CLEARSPPEED WHITE PAPER: The ClearSpeed Vector Math Library

Abstract

This paper describes the capabilities and performance of the ClearSpeed Vector Math Library (VML).

It shows how the library gives users the ability to maximize performance in code that uses commonly used math functions.

High Performance Math Functions from ClearSpeed

ClearSpeed Vector Math Library

The ClearSpeed Vector Math Library¹ (VML) provides a set of standard math functions similar to those found in the C standard libm library. The difference between the standard library and the VML are:

- the VML functions are specifically designed to maximize performance on ClearSpeed architectures,
- the VML functions provide both poly-scalar – one calculation per processing element (PE) – and poly-vector – four calculations per PE – variants of the functions.

Figure 1 shows the main VML functions.

Usage

The library has been designed to be easy to use within the ClearSpeed development environment using the Cⁿ language, a parallel extension to C.

Consider the following C code:

```
#include <math.h>
...
double x,y;
...
x = sin(y);
```

In Cⁿ it is possible to calculate a sin on each processing element by using a poly double datatype and a VML library call:

```
#include <vmathp.h>
...
poly double x,y;
...
x = cs_sinp(y);
```

To totally exploit the pipeline architecture four values per processing element can be calculated simultaneously:

```
#include <vmathp.h>
...
```

sin	exp
cos	log
sincos	tan
atan	sqrt
sinh	cosh
tanh	atanh
isqrt	reciprocal

Figure 1: Main VML functions

```
__DVECTOR x,y;
...
x = cs_sinp(y);
```

Due to the design of the library, calculating four elements at a time is almost as fast as calculating a single element.

The functions in the library are overloaded based on the input type, so both the poly double and poly double vector version of the function have the same name: cs_sinp.

As can be seen from these code snippets, using the vector math library is as easy as using normal C math libraries.

Performance

Due to the Single Instruction Multiple Data (SIMD) architecture of the processing array in ClearSpeed architectures, the running time of a particular function does not vary depending on the input value to the function.

This makes it possible to accurately measure the performance of each function in terms of processor clock cycles by using the hardware monitoring function of the CSX600 chip.



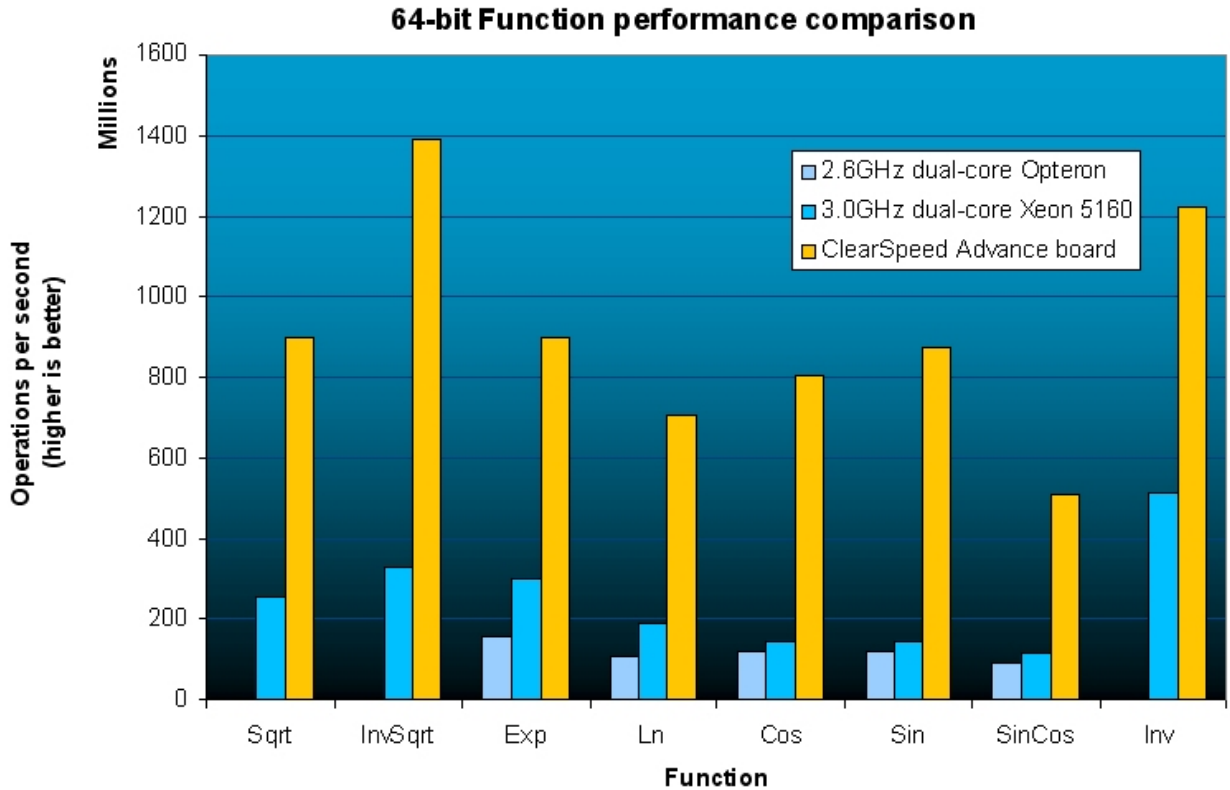


Figure 2 – Performance comparison between the ClearSpeed Advance accelerator, a 2.6 GHz dual core AMD Opteron and a 3.0GHz dual core Intel Xeon 5160

The vector cycle count for a function is how long it takes to calculate four values. So on a CSX600 which has 96 PEs, a call will calculate $96 \times 4 = 384$ values.

Taking this into account, it is possible to determine how many values are calculated per second and compare this against published performance figures from mainstream CPU vendor's high performance math libraries.

Figure 2 shows a comparison between a ClearSpeed Advance™ Board, a dual-core AMD Opteron™ system and a dual-core Intel® Xeon® 5160 (Woodcrest) based system. The operations per second figures are based on published library cycle counts and the clock speeds shown in the graph. Results for the Intel Xeon 5160 and AMD Opteron systems are published in documentation for the Intel Math Kernel Library², (MKL), and

the AMD Core Math Library (ACML™)³ respectively.

Accuracy and Error Checking

Math functions can have errors for two reasons:

- The input value is invalid for the function,
- The returned value is inexact and deviates too far away from the function's actual value.

When using the VML, handling errors in the input are the responsibility of the user. This is to gain as much performance as possible from the library by removing the range checking code.

All performance math libraries provide approximations to the actual value of the function result. The results from libraries

CLEARSPD VECTOR MATH LIBRARY

on different platforms will be close but will not be bit-for-bit comparable.

To check how accurate a platform's library is, an empirical comparison of the results of the library against those results calculated in higher precision will show how close the results are to the correct values.

To test the VML for accuracy, Wolfram *Mathematica*⁴ was used to provide comparisons against higher precision results using 100 significant bits. The input values are generated randomly within the allowable input range of a function. The results of this analysis are shown in Figure 3. This shows both the average and maximum unit in the last place (ULP) error for the VML's double precision functions.

Summary

The ClearSpeed Vector Math Library provides an easy to use, high performance double precision solution for developers requiring standard math functions in ClearSpeed accelerated applications.

	Avg. ULP	Max.ULP
sin	0.36	1.89
cos	0.35	1.96
exp	0.27	1.14
log	0.33	1.00
tan	0.50	3.13
atan	0.89	3.39
sqrt	0.29	1.00
sinh	0.40	2.35
cosh	0.40	2.23
tanh	0.64	3.42
atanh	0.31	1.41
isqrt	0.41	1.88

Figure 3: Accuracy results

¹ Please refer to the *ClearSpeed SDK Reference Manual* and the *ClearSpeed SDK Standard Libraries Reference Manual* for further information on the vector math type.

² <http://www.intel.com/software/products/mkl/data/vml/vmldata.htm>

³ http://developer.amd.com/assets/acml_userguide.pdf

⁴ <http://www.wolfram.com>

ClearSpeed[™]