
Subject: Re: GPULib on my 64-bit WinXP machine
Posted by [Kenneth P. Bowman](#) on Wed, 29 Oct 2008 22:16:42 GMT
[View Forum Message](#) <> [Reply to Message](#)

In article

<ce93a812-3421-47b0-88ad-3772836a72c0@p58g2000hsb.googlegroups.com>,
Vince Hradil <vincehradil@gmail.com> wrote:

>> Do GPUs do IEEE arithmetic (single precision)?

>>

>> Ken Bowman

>

> I believe that is what Peter Messmer said at the user meeting. Of
> course, we're talking specific the CUDA tools on the NVIDIA cards.

>

> More details at: <http://en.wikipedia.org/wiki/CUDA>

>

> including this "limitation"

> Various deviations from the IEEE 754 standard. Denormals and
> signalling NaNs are not supported; only two IEEE rounding modes are
> supported (chop and round-to-nearest even), and those are specified on
> a per-instruction basis rather than in a control word (whether this is
> a limitation is arguable); and the precision of division/square root
> is slightly lower than single precision.

This is like a flashback to the paleocomputing erra of attached array
processors. (Anyone want to admit that they remember FPS, Inc.?)

Also, Crays had their own floating point units (pre-IEEE) with somewhat
different divide/square root units. Maybe NVIDIA is using the same ideas.

Software has gotten so much more powerful now that this is probably
much more transparent. I'll look at the presentation from the UG meeting
when it's posted.

Ken
