Subject: GPULib on my 64-bit WinXP machine
Posted by Vince Hradil on Thu, 23 Oct 2008 16:11:31 GMT
View Forum Message <> Reply to Message

I just wanted to share with the group that I got the GPULib [http://www.txcorp.com/technologies/GPULib/] to run on my 64-bit Windows XP machine with the new NVidia Quadro FX 5600 graphics card (1.5Gb frame buffer [http://www.nvidia.com/object/quadro_fx_5600_4600.html]). Many thanks to Peter Messmer at Tech-X for his help getting the correct executable.

The results are impressive. I ran all the "demos" and the difference is about 21-24 X! I can't wait to try to do some "real" work using this.

The CPU on this machine is a Dual Quad-core Intel Xeon (X5482 @ 3.2GHz), and it has 64Gb RAM, so it's not too shabby itself ;^)

Here's the log from running "bench.pro":

```
IDL> dlm_load, 'gpulib'
% Loaded DLM: GPULIB.
    0.756607     0.756607      1.19352     0.206724    0.0188206
0.756607
    0.756607     0.756607      1.19352     0.206724    0.0188206
0.756607
N iter   =       50
CPU Time =       0.64100003
GPU Time =      0.031000137
Speedup  =        20.677329
IDL> print, !version
{ x86_64 Win32 Windows Microsoft Windows 7.0 Oct 25 2007      64
64}
```

Vince
http://vincehradil.wordpress.com/

Subject: Re: GPULib on my 64-bit WinXP machine
Posted by David Fanning on Tue, 28 Oct 2008 03:04:21 GMT
View Forum Message <> Reply to Message

b_gom@hotmail.com writes:

> I also see that the GTX200 series supports limited double precision
> operations, which might be another trump card.

The point was made at the User's Group meeting that almost

all of the double precision stuff, on any processor, is much
slower than floating point operations. The suggestion was
made to keep everything in floating values if at all possible.

Cheers,

David

--
David Fanning, Ph.D.
Fanning Software Consulting, Inc.
Coyote's Guide to IDL Programming: http://www.dfanning.com/
Sepore ma de ni thui. ("Perhaps thou speakest truth.")

Subject: Re: GPULib on my 64-bit WinXP machine
Posted by Kenneth P. Bowman on Tue, 28 Oct 2008 13:58:32 GMT

View Forum Message <> Reply to Message

In article <MPG.23702f206ddb6fc998a4e0@news.giganews.com>,
David Fanning <news@dfanning.com> wrote:

> b_gom@hotmail.com writes:
>
>> I also see that the GTX200 series supports limited double precision
>> operations, which might be another trump card.
>
> The point was made at the User's Group meeting that almost
> all of the double precision stuff, on any processor, is much
> slower than floating point operations. The suggestion was
> made to keep everything in floating values if at all possible.
>
> Cheers,
>
> David

Do GPUs do IEEE arithmetic (single precision)?

Ken Bowman

Subject: Re: GPULib on my 64-bit WinXP machine
Posted by Vince Hradil on Wed, 29 Oct 2008 02:23:45 GMT

View Forum Message <> Reply to Message

On Oct 28, 8:58 am, "Kenneth P. Bowman" <k-bow...@null.edu> wrote:
> In article <MPG.23702f206ddb6fc998a...@news.giganews.com>,

>   David Fanning <n...@dfanning.com> wrote:
>
>> b_...@hotmail.com writes:
>
>>> I also see that the GTX200 series supports limited double precision
>>> operations, which might be another trump card.
>
>> The point was made at the User's Group meeting that almost
>> all of the double precision stuff, on any processor, is much
>> slower than floating point operations. The suggestion was
>> made to keep everything in floating values if at all possible.
>
>> Cheers,
>
>> David
>
> 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.

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

Subject: Re: GPULib on my 64-bit WinXP machine
Posted by David Fanning on Thu, 30 Oct 2008 04:57:44 GMT

View Forum Message <> Reply to Message

Kenneth P. Bowman writes:

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

Tony Kehoe, who was working for FPS in England, attended the very first IDL class I (or anyone else, probably) ever taught. FPS was to become the first European distributors of IDL. Tony hadn't been in the class 10 minutes before we started laughing about something, and it was pretty much non-stop laughter for the next 10 years, before he died unexpectedly at a very young age.

He would invite me over twice a year to teach an IDL class, but that was completely secondary. Our primary mission was to find somewhere interesting in the UK to walk. We would spend days tramping around, laughing, and talking about IDL. I miss him a lot, and thinking about FPS just makes me miss him more.

Cheers,

David

--
David Fanning, Ph.D.
Fanning Software Consulting, Inc.
Coyote's Guide to IDL Programming: http://www.dfanning.com/
Sepore ma de ni thui. ("Perhaps thou speakest truth.")

Subject: Re: GPULib on my 64-bit WinXP machine
Posted by R.Bauer on Thu, 30 Oct 2008 10:25:07 GMT
View Forum Message <> Reply to Message

David Fanning schrieb:
> b_gom@hotmail.com writes:
>
>> I also see that the GTX200 series supports limited double precision
>> operations, which might be another trump card.
>
> The point was made at the User's Group meeting that almost
> all of the double precision stuff, on any processor, is much
> slower than floating point operations. The suggestion was
> made to keep everything in floating values if at all possible.
>
> Cheers,
>
> David
>

Did one compared AMD to INTEL?

cheers
Reimar

Subject: Re: GPULib on my MacBook Pro
Posted by Brian Larsen on Mon, 03 Nov 2008 23:41:15 GMT
View Forum Message <> Reply to Message

I wanted to add my macbook pro (GeForce 8600M GT) results to this thread. It took a while to get this running, there were 64/32 bit issues and patience (thanks to Nathaniel Sizemore at Tech-X).

```
IDL> @demos/bench/bench
```

```
    0.756607     2.33993    0.196372    0.516154    0.0442747
0.839950
    0.756607     2.33993    0.196372    0.516154    0.0442747
0.839950
CPU Time =      0.79092002
GPU Time =     0.044982195
Speedup  =       17.582957
IDL> print, !version
{ i386 darwin unix Mac OS X 7.0.4 Sep  3 2008      32      64}
```

I am looking forward to playing with this and seeing which of my codes will be much faster.

Cheers,

Brian

 -- --------------
Brian Larsen
Boston University
Center for Space Physics
http://people.bu.edu/balarsen/Home/IDL

Subject: Re: GPULib on my 64-bit WinXP machine
Posted by b_gom on Fri, 07 Nov 2008 23:30:00 GMT

View Forum Message <> Reply to Message

OK, I've run some tests on a Core2 Extreme 3 GHz system (Q6850 processor) with 8 GB ram running WinXP 64-bit, against a GeForce GTX 280 video card (~240 pipelines, 1GB RAM). Running the 'bench.pro' routine included with the GPUlib package, I get the following:

```
% Compiled module: BENCH.
    0.756607    0.756607     1.19352    0.206724    0.0188206
0.756607
    0.756607    0.756607     1.19352    0.206724    0.0188206
0.756607
N iter   =       50
CPU Time =       2.4840000
GPU Time =     0.016000032
Speedup  =       155.24968
```

This seems impressive, but the bench routine is highly artificial; it just runs the gpuLGamma function repeatedly on a static input array. If I include the time to load the array to GPU memory and read it out after each iteration, the speedup drops to a factor of ~12.

For the other test programs, which have more realistic problems, the
speedup ranges from 10 to 100 times.

At any rate, assuming the PC has a fast bus and memory, it appears
that for appropriate types of calculations, you are much better off
buying a $500 video card then spending the premium for multiple or
faster CPUs. Once the bugs get worked out of the library (memory
management is a serious limitation), this should be an extremely
useful package.

Brad

On Oct 23, 9:11 am, Vince Hradil <vincehra...@gmail.com> wrote:

```
>  I just wanted to share with the group that I got
theGPULib[http://www.txcorp.com/technologies/GPULib/] to run on my 64-bit Windows XP
>  machine with the new NVidia Quadro FX 5600 graphics card (1.5Gb frame
>  buffer [http://www.nvidia.com/object/quadro_fx_5600_4600.html]).  Many
>  thanks to Peter Messmer at Tech-X for his help getting the correct
>  executable.
>
>  The results are impressive.  I ran all the "demos" and the difference
>  is about 21-24 X!  I can't wait to try to do some "real" work using
>  this.
>
>  The CPU on this machine is a Dual Quad-core Intel Xeon (X5482 @
>  3.2GHz), and it has 64Gb RAM, so it's not too shabby itself ;^)
>
>  Here's the log from running "bench.pro":
>
>  IDL> dlm_load, 'gpulib'
>  % Loaded DLM:GPULIB.
>      0.756607     0.756607      1.19352     0.206724    0.0188206
>  0.756607
>      0.756607     0.756607      1.19352     0.206724    0.0188206
>  0.756607
>  N iter   =       50
>  CPU Time =       0.64100003
>  GPU Time =      0.031000137
>  Speedup  =        20.677329
>  IDL> print, !version
>  { x86_64 Win32 Windows Microsoft Windows 7.0 Oct 25 2007      64
>  64}
>
>  Vincehttp://vincehradil.wordpress.com/
```

Subject: Re: GPULib on my MacBook Pro
Posted by Marshall Perrin on Sun, 09 Nov 2008 00:39:55 GMT
View Forum Message <> Reply to Message

Brian Larsen <balarsen@gmail.com> wrote:

> I wanted to add my macbook pro (GeForce 8600M GT) results to this
> thread. It took a while to get this running, there were 64/32 bit
> issues and patience (thanks to Nathaniel Sizemore at Tech-X).

Hi Brian,

Is there any chance that you or Nathaniel could post some tips or tricks for replicating your success? I just downloaded GPUlib a little while ago to test it out a bit, so I'd be most grateful if you could share any guidance on the pitfalls to avoid while taking it for a spin. :-) Thanks much,

- Marshall