

---

Subject: Re: GPULib on my 64-bit WinXP machine  
Posted by [b\\_gom](#) on Thu, 23 Oct 2008 17:27:43 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

> 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

Has anyone benchmarked this on a graphics card that doesn't cost more than a high-end PC? It would be interesting to know what kind of performance gain can be achieved, if any, with consumer graphics hardware (i.e. in the \$300 to \$500 range) relative to a normal mid-range PC (~\$1500).

Brad

---

---

Subject: Re: GPULib on my 64-bit WinXP machine  
Posted by [Mort Canty](#) on Thu, 23 Oct 2008 20:57:15 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

b\_gom@hotmail.com schrieb:

>> 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  
>  
> Has anyone benchmarked this on a graphics card that doesn't cost more  
> than a high-end PC? It would be interesting to know what kind of  
> performance gain can be achieved, if any, with consumer graphics  
> hardware (i.e. in the \$300 to \$500 range) relative to a normal mid-  
> range PC (~\$1500).  
>  
> Brad

I knew I was just blogging to myself :-). Have a look at

<http://fwenvi-idl.blogspot.com/>

Cheers

Mort

---

---

Subject: Re: GPULib on my 64-bit WinXP machine  
Posted by [David Fanning](#) on Thu, 23 Oct 2008 21:12:48 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Mort Canty writes:

> I knew I was just blogging to myself :-). Have a look at  
>  
> <http://fwenvi-idl.blogspot.com/>

My goodness, two in one day! I'm going to create a new pointer to IDL blogs on my web page. Anyone else have an IDL blog who wants in on the action? I've been averaging about 2500 hits a day. \*Someone\* might see you here. ;-)

Cheers,

David

P.S. This GPULib thing might really be catching on. I'm hoping to have some time to give it a try tomorrow.

--

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

---

Subject: Re: GPULib on my 64-bit WinXP machine  
Posted by [Michael Galloy](#) on Thu, 23 Oct 2008 21:39:36 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Oct 23, 11:27 am, b\_...@hotmail.com wrote:

>> 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  
>  
> Has anyone benchmarked this on a graphics card that doesn't cost more  
> than a high-end PC? It would be interesting to know what kind of  
> performance gain can be achieved, if any, with consumer graphics  
> hardware (i.e. in the \$300 to \$500 range) relative to a normal mid-  
> range PC (~\$1500).

Running the benchmark demo on a Quadro FX 570, which costs around \$139 - \$250, shows about a 10x speedup. Also see Mort's results at <http://fwenvi-idl.blogspot.com/>, he has a GeForce 8600 GT (about \$100 - \$150).

```
IDL> @bench
% Compiled module: GPUINIT.
% Loaded DLM: GPULIB.
% Compiled module: GPUFLTARR.
% Compiled module: GPUMAKE_ARRAY.
% Compiled module: GPUGETHANDLE.
```

```
% Compiled module: GPUHANDLE__DEFINE.
% Compiled module: GPUPUTARR.
% Compiled module: GPULGAMMA.
    0.756607    2.33993    0.196372    0.516154    0.0442747
0.839950
% Compiled module: GPUGETARR.
    0.756607    2.33993    0.196372    0.516154    0.0442747
0.839950
CPU Time =      0.81534410
GPU Time =      0.075078964
Speedup  =      10.859821
```

```
IDL> err = cudaGetDeviceProperties(prop, 0)
IDL> print, prop.name
Quadro FX 570
```

Mike

--

www.michaelgalloy.com  
Tech-X Corporation  
Associate Research Scientist

---

Subject: Re: GPULib on my 64-bit WinXP machine  
Posted by [Vince Hradil](#) on Thu, 23 Oct 2008 21:56:32 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Oct 23, 3:57 pm, Mort Canty <m.ca...@fz-juelich.de> wrote:  
> b\_...@hotmail.com schrieb:  
>  
>>> 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  
>  
>> Has anyone benchmarked this on a graphics card that doesn't cost more  
>> than a high-end PC? It would be interesting to know what kind of  
>> performance gain can be achieved, if any, with consumer graphics  
>> hardware (i.e. in the \$300 to \$500 range) relative to a normal mid-  
>> range PC (~\$1500).  
>  
>> Brad  
>  
> I knew I was just blogging to myself :-). Have a look at  
>  
> <http://fwenvi-idl.blogspot.com/>  
>  
> Cheers  
>  
> Mort

Oh - I'm glad you responded Mort. I remember reading your blog, but I couldn't find a bookmark for it...

---

---

Subject: Re: GPULib on my 64-bit WinXP machine  
Posted by [russell.grew](#) on Thu, 23 Oct 2008 23:53:43 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

I'm unsure how one goes about linking to older threads, but a thread titled "using GpuLib in IDL" has a few things in it. If you are using google groups, you can easily find it.

I got significant speedup on the spiral benchmark and my video card was only emulating the hardware {i think}

---

---

Subject: Re: GPULib on my 64-bit WinXP machine  
Posted by [b\\_gom](#) on Tue, 28 Oct 2008 02:44:14 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

On Oct 23, 3:39 pm, "mgal...@gmail.com" <mgal...@gmail.com> wrote:

> Running the benchmark demo on a Quadro FX 570, which costs around \$139  
> - \$250, shows about a 10x speedup. Also see Mort's results at <http://fwenvi-idl.blogspot.com/>,  
he has  
> a GeForce 8600 GT (about \$100 - \$150).

I guess what I'm wondering is whether there is a sweet spot in the price range. Are the Quadro 4600\5600 series worth their exorbitant price tags because of their larger memory and 'workstation optimized architecture', or is the cheaper GTX 200 series better because of their larger number of stream processors?

In other words, does the general IDL performance scale directly with the number of processing units times clock speed, assuming there is no bottleneck loading the data into the video ram?

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

Brad

---