Subject: Re: using GpuLib in IDL Posted by Mort Canty on Thu, 05 Jun 2008 17:36:35 GMT

View Forum Message <> Reply to Message

Allard de Wit schrieb:

- > Dear all,
- >
- > After Davids report on the Numerical Analysis Techniques Session, I
- > have been trying to get the GpuLib from TechX working on linux (Fedora
- > 6). I post my findings here, maybe they are useful to others.

>

- 1. First of all you may want to check if your NVIDIA GPU is supported
- > (http://www.nvidia.com/object/cuda_learn_products.html), in the end I
- > found out mine wasn't ;-(.

>

- > 2. Install the CUDA library from http://www.nvidia.com/object/cuda_get.html,
- > I used the CUDA Toolkit version 1.1 for Fedora 7. Add '/usr/local/cuda/
- > bin' to your shell path and '/usr/local/cuda/lib' to your Idconfig
- > path.

>

- > 3. Untar the gpulib package. You will end up with folder
- > 'gpulib-0.3p1' with a couple of folders in there related to MatLab,
- > IDL, Python and some library routines. There is also a folder 'doc'
- > which contains the install-notes. Please ignore these, Peter Messner
- > from TechX told me they are a leftover from another project. Instead
- > relevant installation details are in the 'IDL/doc' folder

>

- > 4. Run the configure script in the gpulib root folder, then run the
- > make command. In my case the make failed at the following statement
- > because of some dependencies to OpenGL:
- > g++ -fPIC -shared -Bsymbolic --warn-once -o gpulib.so gpulib.o ./../
- > vectorOp/gpuVectorOp.o ./../vectorOp/gpuMT.o ./../physicsOp/
- > gpuPhysicsOp.o -L/usr/local/cuda/lib -lcudart -lcublas -lcufft -IGL -
- > IGLU

>

- > According to Peter Messner the dependency on OpenGL wasn't necessary.
- > therefore I run the following command manually (-IGL and -IGLU
- > depencies removed):
- > g++ -fPIC -shared -Bsymbolic --warn-once -o gpulib.so gpulib.o ./../
- vectorOp/gpuVectorOp.o ./../vectorOp/gpuMT.o ./../physicsOp/
- > gpuPhysicsOp.o -L/usr/local/cuda/lib -lcudart -lcublas -lcufft

>

This compiled successfully. >

>

- > 5. Move to the 'IDL' folder and add it to your IDL_DLM_PATH variable.
- > Start IDL and compile 'gpuinit.pro'. Check whether your device is
- > recognized with:
- > IDL> Print, gpuDectectDevice()

>

- > In my case IDL printed -1, showing that no GPU was detected and
- > hardware emulation was enabled. Therefore I cannot show you any
- > performance tests.

>

- > Peter Messner made the following remark on suitable hardware:
- > "Regarding hardware: We have had pretty good experience with the
- > GeForce 8800 GTX or Ulra models. They are a bit on the expensive side
- > (\$400-\$600) but they are worth the money. You would probably also need
- > to upgrade your power supply."

>

> Hope this is useful.

>

> Allard

I was also inspired by David's report to look at GPULib. I'm running IDL 7 on Windows XP with a GeForce 8600GT card (CUDA compatible and cheap). Peter kindly sent me a build for IDL 7 on XP and, after mistakenly installing the CUDA Toolkit 2.0 beta (no joy), I got it up and running in hardware mode pretty quickly. Now I'm keen to start re-coding some of my routines and see what happens. Unfortunately, I have to go on a long weekend vacation and can't start playing around again till Monday. Rats :-(

Cheers anyway,

Mort

Subject: Re: using GpuLib in IDL Posted by russell.grew on Tue, 10 Jun 2008 03:38:18 GMT View Forum Message <> Reply to Message

Dear all,

Just incase anyone is interested, I got GPULib running on Gentoo linux using the ebuild from

http://bugs.gentoo.org/show_bug.cgi?id=167685

to install CUDA.

When configuring GPULib, I added the --disable-MATLAB option and --with-cuda-dir=/opt/cuda

No surprises after that.

The GPULib includes a few benchmarks, one of them, spiral.pro which displays a spiral on screen, unwraps it and coils it up again ran 50

times faster (256 MB nvidia card, a few years old).

Fun.

Subject: Re: using GpuLib in IDL

Posted by russell.grew on Tue, 10 Jun 2008 03:55:59 GMT

View Forum Message <> Reply to Message

Dear all,

Just incase anyone is interested, I got GPULib running on Gentoo linux using the ebuild from

http://bugs.gentoo.org/show_bug.cgi?id=167685

to install CUDA.

When configuring GPULib, I added the --disable-MATLAB option and --with-cuda-dir=/opt/cuda

No surprises after that.

The GPULib includes a few benchmarks, one of them, spiral.pro which displays a spiral on screen, unwraps it and coils it up again ran 50 times faster (256 MB nvidia card, a few years old).

Fun.

Subject: Re: using GpuLib in IDL Posted by russell.grew on Tue, 10 Jun 2008 03:58:51 GMT View Forum Message <> Reply to Message

Also, my card is only a 6800 GT, not listed on

http://www.nvidia.com/object/cuda_learn_products.html

If I do a

print, gpudetectdevice()

I get -1, but don't get any errors if I try and load hardware mode, eg:

gpuinit, /hardware

I am still seeing significant speed advantages. The benchmark, 'gpu_bench' has a speedup of ~1600 [and spiral.pro, ~50 as previously posted].