
Subject: Re: CUDA technology in IDL

Posted by [peter.messmer](#) on Sat, 10 Nov 2007 04:01:03 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Nov 9, 1:12 pm, "pfo...@bcm.tmc.edu" <pf...@bcm.tmc.edu> wrote:

> Greetings,

>

> After many years on the sidelines, I am beginning a project where we
> will be using IDL. The project has the potential to be highly
> parallelized; with that in mind I am curious if CUDA technology is
> supported or has been used with IDL.

>

> <http://developer.nvidia.com/object/cuda.html>

> <http://en.wikipedia.org/wiki/CUDA>

We offer a library that lets you take advantage of a GPU (via CUDA)
from within

IDL. The library offers routines for host-GPU data transfer, kernels
for arithmetic

operations, as well as some GPU optimized IDL intrinsics. For kernels
themselves

(ignoring the data transfer cost) we get speedups of 100x and more,
and for entire

real-world applications we easily get 20x speedup.

E.g.

```
x = findgen(100) * !pi
err = gpuPutArr(x, x_gpu)
err = gpuSin(x_gpu, x_gpu)
err = gpuGetArr(x_gpu, y)
```

allocates a 100 element array x, ships it to the GPU, computes the
sine of it, transfers
it back to the CPU and stores the result in y.

Copying between host and GPU memory is pretty expensive, so the above
sequence pays

only off if your vector is longer than some 1000 elements. However,
once you start to

perform a sequence of operations on the GPU, you can get quite
significant speedup.

For more details, see <http://www.txcorp.com/~messmer/astrogpu.pdf> for
a poster we just

presented at the AstroGPU (www.astrogpu.org) workshop.

I'd be happy to discuss this in more detail,

Cheers,

Peter
