

# GPULib: GPU Computing in IDL (An Update)

Peter Messmer, Paul Mullowney, Mike Galloy, Brian Granger, Dan Karipides, David Fillmore, Nate Sizemore, Keegan Amyx, Dave Wade-Stein, Seth Veitzer

messmer@txcorp.com

**Tech-X Corporation** 5621 Arapahoe Ave., Boulder, CO 80303 www.txcorp.com

#### This work is supported by NASA SBIR Phase-II Grant #NNG06CA13C

IDL User Group Meeting, LASP, Boulder CO, October 16, 2008

#### Why IDL on a GPU?





## **GPUs: Floating-Point Co-Processors**



nultiprocessor')

w (no conv.cache)

lics

ared memory")

#### NVIDIA's CUDA (Compute Unified Device Architecture) ۲

- Latest Generatin (GTX280): Architecture/Programming model no longer focuse
- 128 processing elements, grouped into
- Processors have access to entire
- 2 SIMD processors share a co
- Stream processor: Scalar proce

Support for Double Precision m/cycle, branching, etc.



#### GPUlib: One way to simplify GPU development CECH

- Data objects on GPU represented as structure/object on CPU
  - Contains size information, dimensionality and pointer to GPU memory
- GPULib provides a large set of vector operations
  - Data transfer GPU/CPU, memory management
  - Arithmetic, transcendental, logical functions
  - Support for different types (float, double, complex, dcomplex)
  - Data parallel primitives, reduction, masking (total, where)
  - Array operations (reshaping, interpolation, range selection, **type casting**)
  - NVIDIA's cuBLAS, cuFFT
- Download technology preview http://gpulib.txcorp.com (free for non-commercial use)
- Release at SC'08 (Mid November)

### A GPULib example in IDL





## **GPUlib: Some Vector Operations on GPU**ECH

Memory allocation on GPU

 $y_gpu = gpuFltarr(100, 100)$ 

• Data transfer

gpuPutArr, x, x\_gpu

• Binary operators both plain and affine transform

gpuAdd, x\_gpu, y\_gpu, z\_gpu
gpuExp, a, b, x\_gpu, c, d, z\_gpu

#### • IDL intrinsics

gpuInterpolate, gpuTotal, gpuCongrid, gpuSubArr, gpuReform, gpuWhere, gpuRandomu, gpuFltarr, gpuComplex, gpuMatrix\_multiply ...

IDL structure contains all information about GPU object

type, n\_elements, n\_dimensions, dimensions, handle

#### **Example: Image Deconvolution**



Image is convolved with detector point-spread function:

$$I_{obs}(x, y) = \int I_{true}(x - u, y - v)P(u, v)dudv$$

- Clean image by (complex) division in Fourier space:  $I_{true}(x, y) = FFT^{-1}(FFT(I_{obs}) / FFT(P))$
- Fairly large computational load per CPU-GPU data transfer
- Speedup ranging from 5x 28x for 256x256 3kx3k images

```
gpuFFT, img, gpu_img_fft
gpuDiv, gpu_img_fft, gpu_psf, gpu_img_fft
gpuFFT, gpu_img_fft, gpu_clean, /INVERSE
gpuGetArr, gpu clean, clean
```

#### **GPULib example: Image processing**



*LECH* 

#### **GPULib example: Simulation**





se simulation written in IDL to compute location of scattering maxima (Bragg peaks)

**Data courtesy of** Dr. Matthias Gutmann, **Rutherford Appleton Research Lab, UK** 

### How to write code for GPULib?



- Develop/debug IDL code
- Optimize IDL code via vectorization
  - A well tuned IDL code will benefit most from GPU
- Time it
  - Measure, don't guess
  - Use IDL's profiler
- Successively start moving parts of computation onto GPU
  - gpuPutArr, gpuGetArr for debugging
- Redesign parts of the code
  - reduce CPU<->GPU data transfer
  - increase vector length (  $> 10^4$  elements)
- Create custom kernels
- => talk to us!