Introduction to GPU Programming

Volodymyr (Vlad) Kindratenko

Innovative Systems Laboratory @ NCSA

Institute for Advanced Computing Applications and Technologies (IACAT)

Documentation

- NVIDIA's documentation
  - http://developer.nvidia.com/object/gpucomputing.html
  - Programming Guide ver. 3.0
  - Best Practices Gide ver. 3.0
  - Reference Manual ver. 3.0
- CUDA C SDK Code Samples
  - http://developer.nvidia.com/object/cuda_3_0_downloads.html
- Books
  - David Kirk, Wen-mei W. Hwu, Programming Massively Parallel Processors: A Hands-on Approach, Morgan Kaufmann, 2010
  - Jason Sanders, Edward Kandrot, CUDA by Example: An Introduction to General-Purpose GPU Programming, Addison-Wesley, 2010

Lab Examples

- Exercise 1: Modify fractal code to improve efficiency
  - hint: launch multiple threads per block
- Exercise 2: Modify reduction example to eliminate multiple calls to the kernel
  - hint: use atomic add
- Exercise 3: Modify reduction example to use zero copy
- Exercise 4: Port code in src6 to GPU
  - the code computes volume of a sphere of radius *r* using Monte Carlo integration
  - hint: there is not random number generator function implemented on GPU ☺
- Exercise 5: port tridiagonalization subroutine (**tridiag**) in src7/eigen.c
  - symmetric matrix reduction to tridiagonal form using Givens method