GPGPU computing for image processing and graphics applications

Borko Stosic

Departamento de Estatística e Informática,

Universidade Federal Rural de Pernambuco, Recife, PE, Brazil

The recent boost of GPGPUs represents nothing short of a revolution for scientific computing, as speedups of a 100 times in comparison with traditional CPU simulations on currently easily accessible hardware, are not uncommon. The term "Supercomputing for the masses" coined by Rob Farber [1] synthesizes precisely the current state of the art, and it seems paramount that as many as possible scientists, in diverse areas of knowledge, engage in this novel programming paradigm. This course aims to provide a basis for such a transition for scientists with a solid understanding of the C programming language, by introducing the basic concepts necessary for GPGPU computing, with an emphasis on issues related with image processing, graphics computing, and other scientific applications. First the necessary concepts shall be introduced using concrete examples, and then specific common scientific problems shall be elaborated in detail.

Course outline

Programming model

- Kernels
- Thread hierarchy
- memory hierarchy (global, shared, constant, local, texture)
- Host and Device (CPU e GPU)

Programming interface

- C for CUDA
- OpenCL
- OpenGL interoperability

Performance

- Instruction performance
- Memory bandwidth
- Memory transfer between host and device (CPU e GPU)

- Memory access
- Warp synchronization

Applications

- Examples in image processing
- Examples in computer graphics
- Simulations in Science

Bibliography:

[1] R. Faber, "CUDA, Supercomputing for the Masses", series of articles, Dr. Dobb's Journal, http://drdobbs.com/cpp/231900742

[2] CUDA Programming Guide Version 2.0, http://developer.nvidia.com/nvidia-gpu-computing-documentation

[3] OpenCL Programming Guide, http://www.nvidia.com/content/cudazone/download/opencl/nvidia_opencl_programmingguide.pdf

[4] J. Sanders, and E. Kandrot, "CUDA by Example", Addison Wesley, 2011, ISBN-13: 978-0-13-138768-3

[5] D.B. Kirk and W.W. Hwu, "Programming Massively Parallel Processors: A Hands-on Approach", Morgan Kaufmann Publishers, Elsevier, 2010, ISBN: 978-0-12-381472-2

[6] B. Gaster, L. Howes, D.R. Kaeli, P. Mistry and D. Schaa, "Heterogeneous computing with OpenCL", Morgan Kaufmann Publishers, Elsevier, 2010, ISBN: 978-0-12-387766-6

[7] Wilt, Nicholas. The CUDA handbook: A comprehensive guide to GPU programming. Pearson Education, 2013.