MIC-GPU: High-Performance Computing for Medical Imaging on Programmable Graphics Hardware (GPUs)

Closing Remarks

Klaus Mueller, Ziyi Zheng, Eric Papenhausen

Stony Brook University

Computer Science

Stony Brook, NY

OpenCL

OpenCL: Open Computing Language (based on C)

- support for heterogeneous devices (GPU, CPU, …)
- pick the device best suited for the job
- potential parallelism is key for selection
- recall Amdahl's law

OpenCL Mindset

Platform model:

- a host is connected to one or more OpenCL devices
- a device is divided into one or more compute units (cores)
- compute units are divided into one or more processing elements

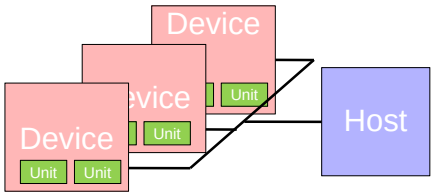

OpenCL Mindset

Execution Model

- host programs execute on the host
- kernels execute on one or more OpenCL devices
- each instance of a kernel is called a *work item*
- work items are organized as *work groups*
- work groups and work items are defined into an *index space*
- index space is created upon kernel submission
- work items can be identified by work group and local work item IDs

→ this is all quite similar to CUDA

CUDA Terminology	OpenCL Terminology
Grid	Index Space
Block	Work Group
Thread	Work Item

Global and Local Dimensions

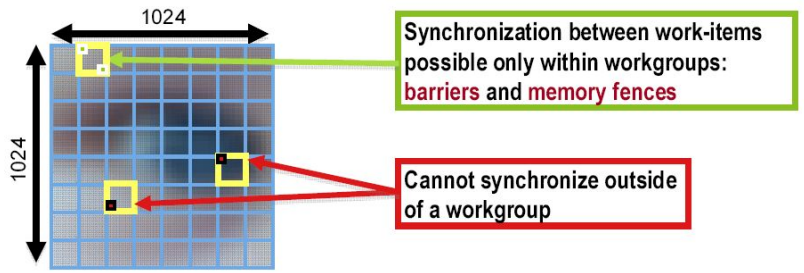

OpenCL Memory Model

Private memory
- per work item

Local memory (16kB)
- shared per work group

Global/constant memory
- not synchronized

Host memory
- on CPU

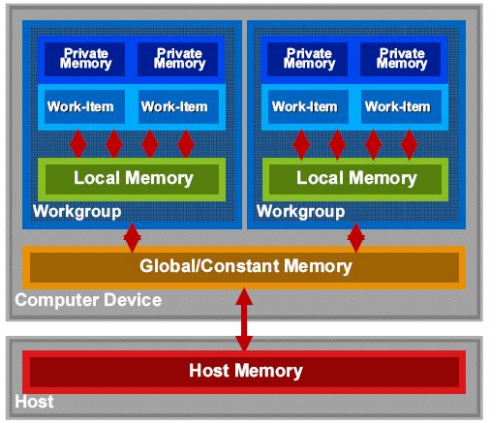

Execution Model

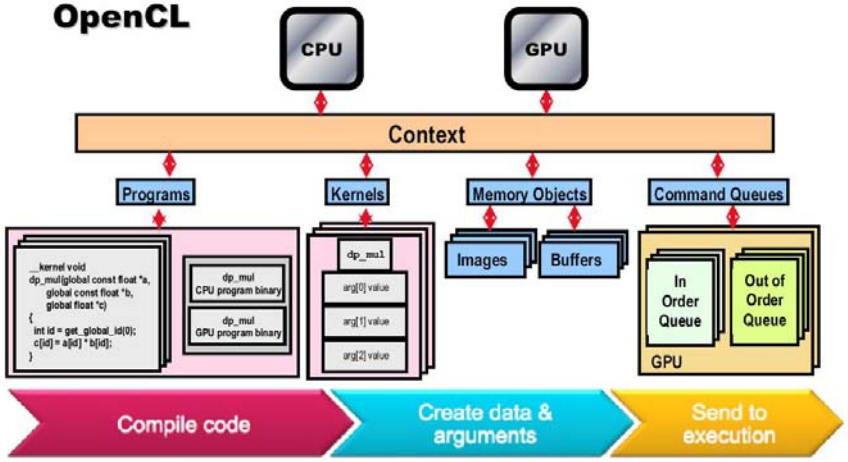

More Information

CUDA
- NVIDIA CUDA Programming Guide (version 2.3)
- NVIDIA CUDA Best Practices Guide (version 2.3)
- go to http://developer.download.nvidia.com

Fermi
- go to http://www.nvidia.com/object/fermi_architecture.html
- also informative: “NVIDIA's 'Fermi' GPU architecture revealed“ by Scott Wasson, available at http://techreport.com/articles.x/17670/1

OpenCL
- Khronos Group OpenCL Overview
- go to http://www.khronos.org/opencl

References

W. Xu, K. Mueller, "Using GPUs to Learn Effective Parameter Settings for GPU-Accelerated Iterative CT Reconstruction Algorithms," GPU Computing Gems Emerald Edition, ed. W. Hwu, pp. 693-708, 2011

E. Papenhausen, Z. Zheng, K. Mueller, "GPU-Accelerated Back-Projection Revisited: Squeezing Performance by Careful Tuning," Workshop on High Performance Image Reconstruction (Fully 3D Image Reconstruction in Radiology and Nuclear Medicine), pp. 19-22, Potsdam, Germany, July 2011.

Z. Zheng, W. Xu, K. Mueller, "Performance Tuning for CUDA-Accelerated Neighborhood Denoising Filters", Workshop on High Performance Image Reconstruction (Fully 3D Image Reconstruction in Radiology and Nuclear Medicine), pp. 52-55, Potsdam, Germany, July 2011.

Z. Zheng, K. Mueller, "A Cache-Aware GPU Memory Scheduling Scheme for CT Reconstruction Back-Projection," IEEE Medical Imaging Conference, Knoxville, TN, October, 2010.

W. Xu, F. Xu, M. Jones, B. Keszthelyi, J. Sedat, D. Agard, K. Mueller, "High-Performance Iterative Electron Tomography Reconstruction with Long-Object Compensation using Graphics Processing Units (GPUs)," Journal of Structural Biology, 171(2):142-153, 2010.

F. Xu, W. Xu, M. Jones, B. Keszthelyi, J. Sedat, D. Agard, K. Mueller, "On the Efficiency of Iterative Ordered Subset Reconstruction Algorithms for Acceleration on GPUs," Computer Methods and Programs in Biomedicine, 98(3):261-270, 2010.

W. Xu, K. Mueller, "A Performance-Driven Study of Regularization Methods for GPU-Accelerated Iterative CT," Workshop on High Performance Image Reconstruction (Fully 3D Image Reconstruction in Radiology and Nuclear Medicine), Beijing, China, September, 2009.

References

W. Xu, K. Mueller, "Accelerating Regularized Iterative CT Reconstruction on Commodity Graphics Hardware (GPU)," (invited paper), IEEE International Symposium on Biomedical Imaging (ISBI), Boston, MA, July, 2009.

F. Xu and K. Mueller, "GPU-Acceleration of Attenuation and Scattering Compensation in Emission Computed Tomography," First Workshop on High Performance Image Reconstruction (Fully 3D Image Reconstruction in Radiology and Nuclear Medicine), pp. 29-32, Lindau, Germany, July 2007.

N. Neophytou, F. Xu, and K. Mueller, " Hardware acceleration vs. algorithmic acceleration: Can GPU-based processing beat complexity optimization for CT?," SPIE Medical Imaging '07, San Diego, February 2007.

F. Xu and K. Mueller, "Real-Time 3D Computed Tomographic Reconstruction Using Commodity Graphics Hardware," Physics in Medicine and Biology, vol. 52, pp. 3405–3419, 2007.

K. Mueller, F. Xu, and N. Neophytou, "Why do GPUs work so well for acceleration of CT?," SPIE Electronic Imaging '07 (Keynote, Computational Imaging V), 64980N, San Jose, January 2007.

K. Mueller and F. Xu, "Practical considerations for GPU-accelerated CT, " IEEE 2006 International Symposium on Biomedical Imaging (ISBI '06), pp. 1184-1187, Arlington, VA, April 6-9, 2006.

F. Xu and K. Mueller, "Accelerating popular tomographic reconstruction algorithms on commodity PC graphics hardware," IEEE Transactions on Nuclear Science, vol. 52, no. 3, pp. 654-663, 2005.

K. Mueller and R. Yagel, "Rapid 3D cone-beam reconstruction with the Algebraic Reconstruction Technique (ART) by using texture mapping hardware," 19(12):1227-1237, IEEE Transactions on Medical Imaging, 2000.

RapidCT GPU Accelerated Tomography

Visit http://www.rapidCT.com

- our papers
- links
- and complete SPIE course notes

Info on our various GPU-accelerated projects

- filtered back projection CT
- iterative CT
- ultrasound transmission tomography CT
- visual CT – monitor ongoing reconstruction
- volume rendering for medical applications
- and much more …