



3-day Course on Image Processing with CUDA

(program could be modified to fit customer's needs)

GPUs are becoming pervasive in scientific computing. Originally served as peripheral accelerators, now they are gradually gaining the main focus in high performance applications development. GPUs are particularly powerful in compute-intensive image processing and may serve in many environments requiring realtime operation, such as video surveillance, traffic analysis or in-vehicle safety systems. Many image processing algorithms could be easily implemented using existing building blocks, such as OpenCV. However, additional performance and flexibility requires deeper knowledge of GPU internals and programming techniques.

Applied Parallel Computing LLC offers a specialized 3-day course on Image Processing with CUDA. The first day is dedicated to the basics of GPU architecture and CUDA programming. The second and the third days of training are dedicated to intensive guided CUDA practice in implementing different types of image filters. Where applicable, CUDA implementations are compared to OpenCV.

The course is delivered on site by one or two NVIDIA-certified CUDA trainers.

Attendees will receive unlimited offline access to all corresponding presentations and code samples in the beginning of each training day.

Day 1: GPU architecture, introduction to CUDA, GPU in realistic applications

09:00-10:00: Introduction to GPU computing

- Types of applications most efficient on accelerators
- Comparing Tesla GPUs, AMD APUs and Intel Xeon Phi
- Key programming principles to achieve high performance on GPU
- Key principles of GPU architecture

10:15-11:45: Introduction to CUDA

- The CUDA programming model and its C++ implementation
- Analogies between MPI+OpenMP for clusters and CUDA for GPUs
- First CUDA program explained
- CUDA compute grid, blocks of threads, grid of blocks, examples
- Realistic CUDA application example (wave propagation)
- Basic optimization techniques

12:00-14:00: GPU and CUDA in practice

- Overview of Eclipse Nsight Edition IDE
- Configuring remote access and file system synchronization for CUDA applications development on remote GPU server
- Monitoring GPU status with *nvidia-smi* utility

- Understanding GPU compute capabilities, *deviceQuery*
- Basic CUDA programming, by example of vector addition
- How to choose efficient CUDA compute grid configuration
- How to measure execution time of CUDA application
- **Hands-on:** “fill-in” exercise on implementing Harris corner detection algorithm in CUDA

14:00-15:30: Lunch

15:30-18:00: GPU architecture, memory hierarchy, CUDA streams

- An overview of Fermi, Kepler and Maxwell GPU architectures
- GPU memory types
- Shared memory
- GPU caches hierarchy and mode switches
- Automatic texture cache (Kepler GK110)
- Unified virtual address space (UVA) in CUDA 6.0
- CUDA streams and asynchronous data transfers

Day 2: Practical image processing with CUDA

09:00-11:30: Computing image histogram in CUDA

11:30-14:00: Histogram-based automatic image contrast enhancement

14:00-15:30: Lunch

15:30-18:00: Background image correction

Day 3: Practical image processing with CUDA

09:00-11:30: Image blurring with 3×3 and 5×5 convolution matrices

11:30-14:00: Image edge detection with 3×3 and 5×5 convolution matrices

14:00-15:30: Lunch

15:30-18:00: Organizing hybrid CPU+GPU workload partitioning with Thread Building Blocks: streamed image processing, by example of background image correction

About Applied Parallel Computing LLC

Applied Parallel Computing LLC is delivering GPU training courses since 2009. Several dozens of courses have been organized all over Europe, both for commercial and academic customers. We work in close partnership with NVIDIA, CUDA Centers of Excellence and Tesla Preferred Partners. In addition to trainings, our company provides GPU porting/optimization services and CUDA certification.