

In-Company Advanced Training

You're designing rockets to get to Mars? Or medical devices? Or self-driving cars? Then you need to know how to port specific algorithms, not what most others find interesting.

After the basic training we offer modules of **4 hours**, discussing advanced subjects. Each subject can be focused on GPU, FPGA, DSP or CPU and using OpenCL, CUDA, OpenMP or OpenCL. Can be combined with a beginner training.

General

- From CUDA to OpenCL - the tricks, tools and optimisation techniques.
- Architecture specific detailed optimizations (or differences across different OpenCL devices)
- Optimizations for host - device interactions (this should include topics such as overlapping data transfers and kernel execution, having multiple command queues or how to work with multiGPU)

Image Processing

- Image Histogram
- Convolutions
- Geometric Scaling
- Point Operations
- Image Segmentation
- Morphological Image Processing

Advanced Data Structures and Parallel Algorithms

- Designing Efficient Data Structures for Parallel Programming
- Parallel Optimization Patterns
- Scan
- Reduce
- Sort
- Graph Traversal Algorithms
- BLAS algorithms

Practical info

These 4-hour blocks build up our inhouse trainings. Costs are ?4000 per half day (one subject). A full training with basics and various advanced subjects costs between ?15,000 and ?30,000.

Trainings are given world-wide.