CUDA Basic

Maximize your hybrid application performances

NVIDIA CUDA

The last generation of NVIDIA GPUs offers a tremendous computing power that can now benefit to not only graphics processing but also to scientific applications.

NVIDIA CUDA is a general purpose parallel computing architecture that leverages the parallel compute engine in NVIDIA graphics processing units (GPUs) to solve many complex computational problems in a fraction of the time required on a CPU.

Duration

2 day training

Objectives

In this 2 day CUDA training, participants will learn the CUDA programming model and tips to reach performance. The training includes hands-on practical labs where participants will be able to progressively program high performance computations.

Prerequisites

- Knowledge in C

Training conditions

The training is limited to 14 persons.

Deliverables

The participants will receive all training materials: courses and practical exercises.

Training content

Day 1

Morning - CUDA Basics
- Introduction to GPU computing
- CUDA architecture and programming model
- CUDA API
- CUDA debugging

Afternoon - CUDA Kernel Performance
- CUDA warps
- Data alignment & coalescing

Day 2

Morning - CUDA kernel performance
- Texture memory & constant memory
- Shared memory

Afternoon - CUDA grids & kernels optimization
- Maximizing occupancy
- Interpreting profiler counters
- CUDA performance tools: Visual Profiler, ...

NVIDIA Partner

CAPS

Innovative software for manycore paradigms

Headquarters

Immeuble CAP Nord
4A Allée Marie Berhaut
35000 Rennes
France

Tel.: +33 (0)2 22 51 16 00

training@caps-entreprise.com
www.caps-entreprise.com

CAPS offers services to help you build optimized applications running on parallel high performance systems. Ranging from training to complete application porting, we can give you all the expertise your problem might require.

CAPS