



Training

HMPP CUDA

Corporate



Benefit from the performance of manycore systems while reducing your development efforts with HMPP Workbench.



**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

Duration

3 day training

Objectives

In this 3 day training, participants will learn the CUDA programming model and tips to reach performance with NVIDIA GPUs. Then, they will learn how to easily take advantage of such architectures and tune their codes with HMPP.

Prerequisites

- Knowledge in C or Fortran
- Knowledge in Linux
- Knowledge in CUDA is a plus

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
- Lab session: getting used to CUDA device
- Lab session: programming a basic addition

Afternoon - CUDA Kernel performance (1/2)

- Using 2D CUDA grid for large computations
- Lab session: programming 2D CUDA grid computation
- CUDA warps
- Data alignment & coalescing
- Lab session: ensuring coalescing

Day 2

Morning - CUDA Kernel performance (2/2)

- Texture memory & constant memory
- Shared memory
- Lab session: use CUDA shared memory as a cache
- Maximizing occupancy
- Interpreting profiler counters
- CUDA performance tools: Visual Profiler

Afternoon - HMPP Basics

- Introduction to parallel hybrid programming
- HMPP overview
- Lab session: HMPP Hello World
- Basis of HMPP programming
- Lab session: offloading a computation into a GPU
- HMPP compilation model
- Lab session: compiling an HMPP application

Day 3

Morning - HMPP Transfers optimization

- Managing data transfers
- Lab session: programming data transfers
- Grouping GPU computations
- Optimizing data movement

Afternoon - Optimizing GPU code generation with HMPP

- Advanced CUDA performance
- Driving the code generation & gridification
- Automatic loop transformations: unrolling, splitting, jamming, ...
- Lab session: optimizing Sgemm code generation

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.

