HMPP Training

Training by CAPS entreprise

19-21
September 2011
(3 days)

Oak Ridge National Laboratory

Agenda

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

Day 2

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

Afternoon - Optimizing code generation with HMPP
- Advanced kernel performance
- Driving the code generation and gridification
- Automatic loop transformations: unrolling, splitting, jamming, …
- Lab session: optimizing Sgemm code generation

Day 3

HMPP advanced
- HMPP C++ API to directly target the HMPP Runtime
- Data mirroring: dynamic management of data identified by their host address
- Computation on a collection of independent data
- Parallel execution of a HMPP collection on multiple GPUs
- Heterogeneous distribution of a parallel loop over multiple GPUs and CPUs

CAPS