nCore
NCT-300

nCore

NCT-300 Programming GPU Processors

Four-day intensive training course covering all aspects of programming GPU processors using advanced methods and techniques

Course Overview

This course covers concepts and approaches related to programming GPU processors using both CUDA and OpenCL. Extensive coverage of GPU hardware, memories, data transport, and performance optimization enable the student to understand the fundamental aspects of GPU programming. In-depth, hands-on laboratories demonstrate how to apply common numerical methods to GPU processors using both the native APIs and open source numerical libraries. This course also covers methods of integrating the Intel TBB threading abstraction layer with GPU software APIs.

Benefits

- Teaches everything necessary to start developing high-performance GPU software on Linux or Windows platforms.
- Covers both CUDA and OpenCL, including open source computational libraries, such as CUBLAS, CUFFT, and CUDPP.
- Explains how to integrate multicore software development techniques with GPUs to increase performance.
- Offers a detailed overview of fundamental concepts, while providing advanced training and practical advice on GPU programming.
- Provides in-depth instruction that increases students knowledge and skills through an online training delivery platform and instructor led, hands-on laboratories.

Course Objectives

- Install GPU libraries and drivers, and compile CUDA/OpenCL programs on Linux and Windows operating systems.
- Understand nVidia GPU hardware and the underlying technical concepts, including SIMD processing and hardware threading architectures.
- Understand the different GPU programming APIs and their appropriate use with various applications.
- Learn single and double precision floating point calculations.
- Recognize the difference between GPU memory types and the advantages and disadvantages of each.
- Effectively orchestrate the tranport of data to and from GPU memory.
- Correctly implement two common types of numerical algorithims - Matrix Multiplication and Reduction.
- Cover performance optimization, including the cudaprof profiling tool, loop unrolling, coalesced memory access, memory bandwdith estimations, and occupancy.
- Participate in labs for commonly used open source computational libraries - CUBLAS, CUFFT, and CUDPP.
- Learn to meld multicore processors and GPUs to take maximum advantage of modern platform performance.
- Learn to integrate the Threading Building Blocks threading abstraction layer with GPU code and migrate TBB primitives to the GPU.
- Discover how to take advantage of multiple GPUs in the same server.
- Cover cudagdb for CUDA debugging, including the use of emulation mode with valgrind.
- Participate in extensive hands-on laboratories with code examples, using both CUDA and OpenCL.

Who should attend

Software architects, developers, team leaders, and managers seeking to develop GPU software. Knowledge of computer architectures and intermediate C++ programming, as well as software development experience are mandatory pre-requisites for this course.

About nCore

nCore HPC is a global provider of professional services and systems focused on high performance, low latency and scalability in embedded computing. nCore delivers state-of-the-art solutions to government agencies, high-technology organizations, defense, research, biomedical and financial companies.

nCore is a working group member of the Multicore Association.

Course Registration

Length: 4 Days **Cost:** $3495

By E-mail: training@ncorehpc.com

THE Multicore ASSOCIATION™

intel Software Partner