

Acceleware 5 day CUDA/OpenCL Training Course Detailed Outline

Lecture 1 - Overview of GPU Computing

- Motivation
- High-level hardware specifications
- Brief history of GPGPU
- CUDA/OpenCL Overview
- *Hands-on Exercise – GPU Memory Management*

Lecture 2 - Data-Parallel Architectures and the GPU Programming Model

- Data-parallelism
- GPU Programming Model
 - GPU Kernels
 - Host vs. device responsibilities
 - CUDA/OpenCL Syntax
 - Thread/Work-item hierarchy
- *Hands-on Exercise – Simple CUDA Kernels*

Lecture 3 - The GPU Memory Model & Thread/Work-item Cooperation

- Task parallelism
- Thread/Work-item Cooperation in GPU Computing
- GPU Memory Model
 - Shared/Local memory
 - Constant memory
 - Global memory
- *Hands-on Exercise – Shared memory and constant memory*

Lecture 4 – Textures

- GPU Texture Hardware Introduction
- Texture Filtering, Texture Caching
- *Hands-on Exercise - Textures*

Lecture 5 - Asynchronous Operations

- Asynchronous vs. Synchronous memory transfers
- Streams/Queues
- Page Locked Memory
- *Hands-on Exercise – Asynchronous operations*

Lecture 6 – Other GPU Features

- Atomics
- Double precision

Lecture 7 - Libraries

- CUFFT, CUBLAS, CUDPP
- *Hands-on Exercise – Experience with CUFFT, CUBLAS*

Lecture 8 - Debugging GPU Programs

- Debugging Tools and Techniques
- *Hands-on Exercise - Debugging*

Lecture 9 – Introduction to Optimizations

- Arithmetic optimizations
- Instruction Cost
- Warps/Wavefronts
- Intrinsic
- *Hands-on Exercise - Arithmetic optimizations*

Lecture 10 - Resource management, Latency, and Occupancy

- Memory Latency
- Occupancy
- Occupancy Calculator
- *Hands-on Exercise - Occupancy Calculator*

Lecture 11 - Memory Performance Optimizations

- Coalescing
- Memory Properties
- Bank Conflicts
- Memory Usage Strategies

Lecture 12 - Profiling applications

- Timers
- Visual Profiler
- *Hands-on Exercise - Profiling*

Lecture 13 - OpenCL/CUDA Runtime/CUDA Driver Overview

- Language overview
- Hardware options overview – NVIDIA and AMD
- Programming comparisons
- *Hands-on Exercise – CUDA driver API*

Case Study

- Finite Difference stencil algorithm
- *Exercise - Finite Difference stencil algorithm or*
- MD5 encryption algorithm
- *Hands-on Exercise - MD5 encryption algorithm*

Lecture 14 - Microsoft HPC Cluster (**Optional**)

- Cluster Overview
- HPC Server Overview
- MPI and MS MPI
- Clustered GPU Programming
- *Hands-on Exercise – GPU and MPI*

Wrap-up

- Course Evaluations
- Course completion certificates

750 Palomar Ave, Sunnyvale, CA 94085

Phone 408.730.2275

Fax 408.730.2274

Web www.colfax-intl.com

© 2010 Colfax International