

## 2 Day CUDA/OpenCL Training Course Outline

### Day 1:

- § Lecture: Overview of GPU Computing
  - § Motivation
  - § High-level hardware specifications
  - § Brief history of GPGPU
  - § CUDA/OpenCL Overview
- § Hands-on-Exercise: Memory Allocation and Memory Transfers
- § Lecture: 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 Kernels
- § Lecture: 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: Asynchronous Operations
  - § Asynchronous vs. Synchronous memory transfers
  - § Streams/Queues
  - § Page Locked Memory

### Day 2:

- § Hands-on-Exercise: Asynchronous Operations
- § Lecture: Introduction to Optimizations
  - § Arithmetic optimizations

- § Instruction Cost
- § Warps/Wavefronts
- § Intrinsic
- § Hands-on-Exercise: Arithmetic Optimization
- § Lecture: Resource Management, Latency and Occupancy
  - § Memory Latency
  - § Occupancy
  - § Occupancy Calculator
- § Hands-on-Exercise: Occupancy Calculator
- § Lecture and Hands-on-Exercise: Memory Optimizations
  - § Coalescing
  - § Memory Properties
  - § Bank Conflicts
  - § Memory Usage Strategies

750 Palomar Ave, Sunnyvale, CA 94085

Phone 408.730.2275

Fax 408.730.2274

Web [www.colfax-intl.com](http://www.colfax-intl.com)

© 2010 Colfax International