Accelerating C, C++, OpenCL, and RTL Applications with the SDAccel Environment

EMBD 2

EMBD-OCLSDA (v1.0)

Course Specification

Course Description

Learn how to develop, debug, and profile new or existing OpenCL™, C/C++, and RTL applications in the SDAccel™ development environment for use on Xilinx FPGAs. Also learn how to run designs on the Alveo™ accelerator card using Nimbix Cloud.

The focus is on learning how to utilize techniques in the SDAccel environment to:

- Reduce latency
- Utilize the massive parallelism inherent to FPGAs
- Optimize throughput
- Pipeline for performance

This course also provides an introduction to targeting the Alveo accelerator card.

Level – EMBD 2

Course Duration – 2 days

Price –

Course Part Number – EMBD-OCLSDA

Who Should Attend? – Anyone who needs to accelerate their software applications using FPGAs.

Prerequisites

- Basic knowledge of Xilinx FPGA architecture
- Comfort with the C/C++ programming language

Software Tools

- SDx™ development environment 2018.3.op

Hardware

- Architecture: Xilinx Kintex® UltraScale™ FPGA

After completing this comprehensive training, you will have the necessary skills to:

- Describe how the FPGA architecture lends itself to parallel computing
- Explain how the SDx development environment helps software developers to focus on applications
- Examine the OpenCL API execution model
- Analyze the OpenCL API memory model
- Create kernels from C, C++, OpenCL, or RTL IP (using the RTL Kernel Wizard)
- Apply host code optimization and kernel optimization techniques
- Move data efficiently between kernel and global memory
- Profile and debug OpenCL API code using the SDAccel development environment

Course Outline

Day 1

- Introduction to the SDAccel Environment and OpenCL Framework {Lecture}
- SDx Tools Overview {Lecture, Lab}
- Makefile Flow {Lecture, Lab}
- Introduction to FPGAs {Lecture}
- Alveo Product Overview {Lecture}
- Alveo Partner Ecosystem Solutions Overview {Lecture}
- Introduction to Nimbix Cloud {Lecture}
- OpenCL Framework Fundamentals 1 {Lecture}
- OpenCL Framework Fundamentals 2 {Lecture, Lab}
- Synchronization {Lecture, Lab}

Day 2

- Introduction to NDRanges {Lecture}
- Working with NDRanges {Lecture, Lab}
- Profiling {Lecture}
- Debugging {Lecture}
- C-Based Kernels {Lecture}
- C-Based Kernel Optimization {Lecture}
- Optimization Methodologies {Lecture}
- Memory Transfer Optimization Techniques {Lecture}
- Kernel Optimization Techniques {Lecture, Lab}
- Using the RTL Kernel Wizard to Reuse Existing IP as Accelerators {Lecture, Lab}

Topic Descriptions

Day 1

- Introduction to the SDAccel Environment and OpenCL Framework – Explains how software engineers and application developers can benefit from the SDAccel development environment and Open Computing Language (OpenCL) framework.
- SDx Tools Overview – Describes the elements of the development flow, such as software emulation, hardware emulation, and system run as well as debugging support for the host code and kernel code.
- Makefile Flow – Introduces the SDAccel environment makefile flow, where the user manages the compilation of host code and kernel(s).
- Alveo Product Overview – Describes the Alveo Data Center accelerator cards and lists the advantages of these cards and the available software solutions stack.
- Alveo Partner Ecosystem Solutions Overview – Describes the partner solutions available in the cloud and on premises for Alveo Data Center accelerator cards.
- Introduction to Nimbix Cloud – Explains the Nimbix Cloud, availability of the Alveo Data Center accelerator cards in the Nimbix Cloud, and how to run the design on the Nimbix Cloud.
- Introduction to FPGAs – Describes fundamental information about FPGAs, which is required to guide the SDAccel tool to the best computational architecture for any algorithm.
- OpenCL Framework Fundamentals 1 – Describes OpenCL framework models such as the Platform model, Execution model, Memory model, and Programming model.
- OpenCL Framework Fundamentals 2 – Describes OpenCL framework components such as the OpenCL platform API, OpenCL run-time API, and OpenCL programming language.
- Synchronization – Describes OpenCL synchronization techniques such as events, barriers, blocking write/read, and the benefit of using out-of-order execution.

Day 2

- Introduction to NDRanges – Explains the basics of NDRange (N dimensional range) and the OpenCL execution model that defines how kernels execute with the NDRange definition.
- Working with NDRanges – Explains the host code and kernel code changes with respect to NDRange. Also explains how NDRange works and the best way to represent the work-group size for the FPGA architecture.

- Profiling – Describes the different reports generated by the tool that help to optimize data transfer and kernel optimization.
- Debugging – Explains the support for debugging host code and kernel code as well as tips to debug the system.
- C-Based Kernels – Describes the trade-offs between C/C++, OpenCL, and RTL applications and the benefits of C-based kernels.
- C-Based Kernel Optimization – Describes techniques for developing a high-performance C kernel.
- Optimization Methodologies – Describes the recommended flow for optimizing an application in the SDAccel environment.
- Memory Transfer Optimization Techniques – Describes the various optimization techniques for data transfer between kernels and global memory.
- Kernel Optimization Techniques {Lecture, Lab} – Apply different techniques such as loop unrolling, pipelining, and DATAFLOW.
- Using the RTL Kernel Wizard to Reuse Existing IP as Accelerators {Lecture, Lab} – Describes how the SDAccel environment provides RTL kernel developers with a framework to integrate their hardware functions into an application running on a host PC connected to an FPGA via a PCIe® interface.

Register Today

Visit the Xilinx Customer Training Center to view schedules and register online.