

FPGA-based Accelerated Cloud Computing with SDAccel

The increasing computational requirements of next-generation Cloud and High-Performance Computing (HPC) applications are pushing the adoption of accelerated computing based on heterogeneous architectures into mainstream, as traditional CPU technology is unable to keep pace.

Xilinx FPGAs are now available, in two different sizes that include up to eight Virtex® UltraScale+ VU9P, on the Amazon Elastic Compute Cloud (EC2) F1 instances, which are designed to accelerate data center workloads, including machine learning inference, data analytics, video processing, and genomics. Furthermore, Amazon Web Services offers the SDAccel™ Development Environment for cloud acceleration, enabling the user to easily and productively develop accelerated algorithms and then efficiently implement and deploy them onto the heterogeneous CPU-FPGA system.

SDAccel completely automates the step of the hardware design flow, offering an easy to use environment for FPGA application design. It offers the possibility to specify a compute kernel using C and C++ for higher-level algorithmic implementation, or using hardware description languages for RTL designs, while using OpenCL APIs to control run-time behavior.

Attendees will use their laptops to connect to the EC2 F1 instances.