

A Simplified Computing Framework for FPGA-Accelerated Workloads

Tuesday, July 10, 2018 4:45 PM (15 minutes)

We introduce SWiF - Simplified Workload-intuitive Framework - a workload-centric, application programming framework designed to simplify the large-scale deployment of FPGAs in end-to-end applications. SWiF intelligently mediates access to shared resources by orchestrating the distribution and scheduling of tasks across a heterogeneous mix of FPGA and CPU resources in order to improve utilization and maintain system requirements. We implemented SWiF atop Intel Accelerator Abstraction Layer (AAL) and deployed the resulting software stack in a datacenter with an Intel-based Xeon+FPGA server running Apache Spark. We demonstrate that by using SWiF's API, developers can flexibly and easily deploy FPGA-enabled applications and frameworks with the addition of only 3 lines of code to existing code base. In particular, we demonstrate that by offloading through SWiF the compression workload of Spark to FPGA, we realize a 3.2X speedup in total job execution, and up to 5X when Spark's Resilient Distributed Datasets (RDDs) are persisted.

Primary author: OJKA, David Nonso (University of Florida (US))

Presenter: OJKA, David Nonso (University of Florida (US))

Session Classification: Posters

Track Classification: Track 5 –Software development