

FaaM: FPGA as a Microservice - A Case study for Data Compression

Thursday 12 July 2018 14:15 (15 minutes)

Field-programmable gate arrays (FPGAs) have largely been used in communication and high-performance computing, and given the recent advances in big data and emerging trends in cloud computing (e.g., serverless [18]), FPGAs are increasingly being introduced into these domains (e.g., Microsoft's datacenters [6] and Amazon Web Services [10]). To address these domains' processing needs, recent research has focused on using FPGAs to accelerate workloads, ranging from analytics and machine learning to databases and network function virtualization. In this paper, we present an ongoing effort to realize a high-performance FPGA-as-a-microservice (FaaS) architecture for the cloud. We discuss some of the technical challenges and propose several solutions for efficiently integrating FPGAs into virtualized environments. Our case study deploying a multithreaded, multi-user compression as a microservice using the FaaS architecture indicate that microservices-based FPGA acceleration can sustain high-performance compared to straightforward implementation with minimal to no communication overhead despite the hardware abstraction.

Primary authors: OJKA, David Nonso (University of Florida (US)); Mr KUNZ, Roland (DELL)

Presenter: Mr KUNZ, Roland (DELL)

Session Classification: T7 - Clouds, virtualization and containers

Track Classification: Track 7 –Clouds, virtualization and containers