

Bump-in-the-Wire FPGAs and HPC in the Cloud

Ahmed Sanaullah, Boston University

Abstract

One of the most important recent datacenter innovations is the use of FPGAs as “bump-in-the-wire” accelerators for network functions. In this project, Red Hat and BU are exploring new forms of workload acceleration, including High Performance Computing applications. A topic of particular interest is leveraging emerging FPGA architectural enhancements to create development environments that make these embedded accelerators easily accessible to application programmers. In this Micro-talk we will, very briefly, introduce FPGAs and their capabilities, describe the architecture of an FPGA-enhanced Cloud, and present the first results showing the possibility of strong scaling in such Clouds. The results show that FPGA-centric Clouds could fill important niches in the space of HPC computational resources. Using FPGA-enhanced nodes to achieve Universal HPC will require rethinking the entire system stack from applications through middleware down to hardware architecture.

Speaker Bios

Ahmed Sanaullah is a 4th year Computer Engineering PhD student at Boston University. His primary research at BU is aimed at scalable and open source High Performance Computing (HPC) for bare-metal public clouds. He has also worked on a number of other topics such as Performance-Portability for HPC, Application aware Machine Learning architectures, and long timescale Molecular Dynamics.