Implications of FPGAs for Floating-Point HPC Systems

Keith D. Underwood and K. Scott Hemmert

FPGAs have experienced a remarkable growth in density, functionality, and clock rate over the last eight years. This has allowed them to become quite capable at floating-point implementations and the algorithms that depend on them. Depending on technology trends, FPGAs may outstrip microprocessors in peak floating-point performance in the near future. This has led to numerous groups investigating floating-point kernels that could be accelerated by FPGA technology. This talk will begin by discussing our work with BLAS operations and FFT operations, and how we expect the performance of those operations to scale over the next few years.

As important as kernels are, they are not applications – HPC systems are not bought to run kernels. Our current research has turned to focusing on issues surrounding how FPGAs can be integrated into HPC systems. We will discuss two types of applications that could leverage BLAS and FFT operations and the implications for library interfaces and system architecture to support that. We will conclude with a brief discussion of what we see as the primary challenges facing the deployment of FPGAs in HPC systems.