SAND2007-1840P

Opportunities for Reconfigurable Computing in HPC

Sandia researchers are engaged in exploring new architectures for use in our high performance computing (HPC) platforms. One notable research effort explores field programmable gate arrays (FPGAs). Our initial findings were striking – over the 6 year period from 1997 to 2003, the *peak double precision floating-point* performance of FPGAs grew at a rate of 4.5X every two years. This was much faster than the 2X every 18 months that microprocessors were experiencing. FPGAs were obtaining performance improvements from transistor density, architecture improvements, and clock rate improvements, where microprocessors were *only* obtaining improvements from clock rate. Our earliest publication of these results has been widely cited and spawned a broad interest in traditional scientific applications that require double precision floating-point arithmetic.

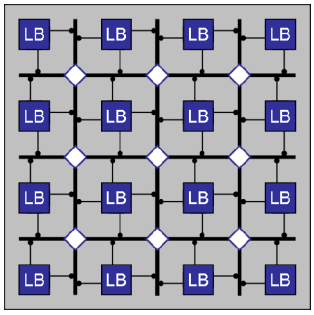

Figure 1: A conceptual diagram of an FPGA

An FPGA is a programmable hardware device that is typically used to replace Application Specific Integrated Circuits (ASICs) for applications that have relatively small hardware requirements or for applications with more complicated hardware requirements, but lower chip volumes. An example of an FPGA is shown in Figure 1. The field of reconfigurable computing uses FPGAs for computing purposes (typically attached to a more conventional microprocessor). Over the past two decades, reconfigurable computing has shown remarkable promise for fields that *do not* use floating-point arithmetic; however, Sandia was the first to demonstrate their viability for applications that require floating-point arithmetic.

Sandia has also been exploring how these devices can be leveraged in applications. Many propose using FPGAs to accelerate library calls (e.g. FFT), but our recent research has demonstrated that traditional library calls would not be sufficient for typical scientific applications. In contrast, a non-blocking interface could be used to deliver more performance than a competing microprocessor (Figure 2). Contrary to popular belief, we also found that the latency between a microprocessor and the attached FPGA is generally irrelevant, but the bandwidth is the dominating factor for scientific computing. The results indicated that it was possible to accelerate important routines, but the high bandwidth connection between the microprocessor and the FPGA would be a very aggressive design point.

While FPGAs have interesting *potential* for future HPC platforms, they still face major challenges. For example, Los Alamos has found that current FPGAs dramatically decrease the reliability of large scale systems due to their inability to tolerate soft errors. In addition, FPGA based systems are notoriously difficult to program and applications for FPGA systems have very low portability. Because of our group's experience with traditional HPC systems (e.g. CPlant and then Red Storm), we have taken an active role in emphasizing the critical nature of reliability, portability, and programmability to the research and vendor communities. As part of an effort to address many of these challenges, Keith Underwood currently serves as the vice president of OpenFPGA – a new organization who's core mission includes creating standards to enable portability in FPGA systems and that provides a key interface between the user and vendor communities.

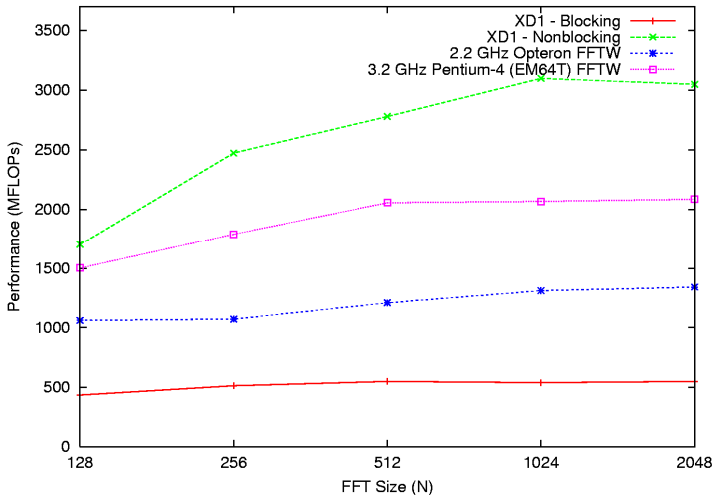

Figure 2: Streaming FFT performance

Contact: Keith D. Underwood, kdunder@sandia.gov

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

National Nuclear Security Administration

Sandia National Laboratories