FPGAs vs. CPUs: Trends in Peak Floating-Point Performance

Keith Underwood

Year of publication: *2004*
Area: *Reconfigurable Computing*

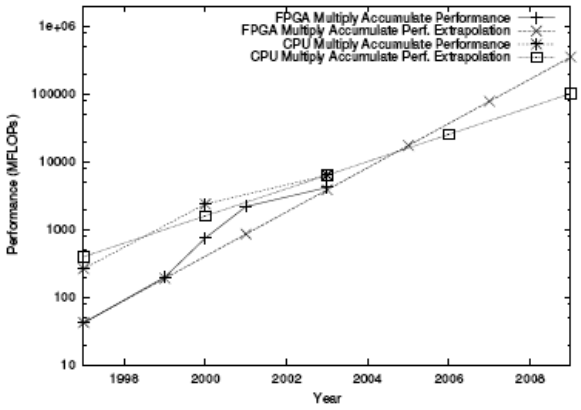

The results presented in this paper showed that there was promise for the future use of FPGAs in supercomputers running applications requiring intensive high-precision, floating-point operations.

Ever since the beginning of FPGA history, there was the vision that FPGAs could be used for computing, which was not the primary intent or market for these devices. At the time this paper was published, computing using fixed-point operations on FPGAs was clearly beneficial and best demonstrated in the area of digital signal processing where FPGAs often showed a significant advantage over programmable DSP processors. However, most scientific computations required, and still require, the use of IEEE floating-point, often in double-precision mode, to achieve computational stability. Until this paper, it was generally believed that FPGAs could not implement enough floating-point functionality to compete with the CPUs of the day.

By measuring the size and maximum clock rate of floating-point adders, multipliers, dividers and multiply accumulate operators over several generations of FPGAs and then dividing the sizes into the capacity of the devices, Underwood generated trend lines for the peak performance of the FPGAs for those operators. For CPUs, trend lines were derived for each operation using the corollary to Moore's Law that predicted performance would double every 18 months. The results showed that the peak capabilities of FPGAs had already crossed, or were about to cross ,the trend lines for CPUs and the lines would continue to diverge.

This is the first paper that tried to quantifiably show that FPGAs could be practical as computing devices in the floating-point domain. By developing a rather coarse model, it was possible to show that the trend lines were favourable towards FPGAs being competitive with CPUs at some point. Reconfigurable computing research still had a future! However, the conclusion of the paper is very much dependent on the meaning of peak performance and how that relates to achievable performance – a problem that is yet to be solved even in the much more mature CPU world.

Paul Chow

DOI: http://doi.acm.org/10.1145/968280.968305