Signal processing at 250 MHz using high-performance FPGA's

Brian von Herzen

Year of publication: *1997*
Area: *Applications*

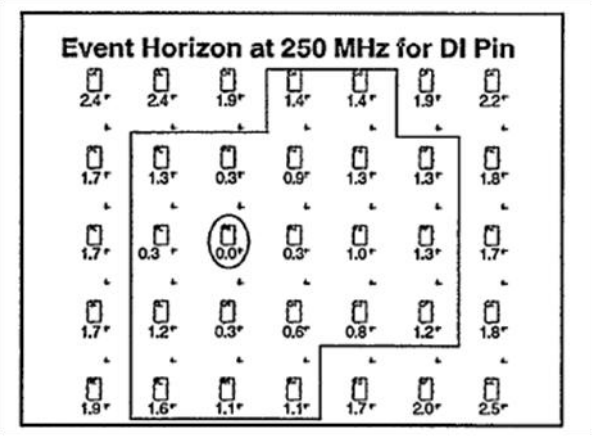

This paper was an inspirational tour de force in FPGA design showing the maximum performance achievable with an FPGA, and how one would go about extracting that performance.

To fully appreciate this paper, you have to remember what designs were like in 1997. Intel's Pentium, often today's standard of high-frequency operation, was only running at 75-100MHz in 0.5µm and 0.6µm technologies. You wouldn't even see a 200MHz Pentium until 0.35 micron. Most FPGA designs ran at 25-40MHz. The general impression was that FPGAs were necessarily slow compared to ASICs and processors.

Von Herzen shows that he can extract useful 250MHz operation out of a 0.6µm (today we might say 600nm) Xilinx XC3100A in 1997 – many FPGA users would be happy to see that performance today out of their 45nm parts! It does require heroic effort in matching the design with the FPGA fabric, including careful placement, pipelining to the layout, and careful planning of exactly how far a signal can travel over the network within a clock cycle (the "Event Horizon" shown in the figure above). Significantly, the paper outlines the design approach. This provides a forward look at what CAD and architectures must do to fully utilize these FPGAs. Later work on clocked FPGAs (e.g. GARP, HSRA [see Tsu, et. al. 1999], CHESS, SFRA, Tabula), interconnect retiming, and streaming compute models build out this vision.

In many ways, this is a model of what you want to see from an FPGA application paper. It demonstrates results that seem beyond the capabilities of FPGAs and details how to use the features of the device to do that – providing lessons that are FPGA-specific and applicable beyond the particular design. Although parts of the paper are dry and read like a lab report, the paper is readable and relevant even today.

André DeHon and Steve Trimberger

DOI: http://doi.acm.org/10.1145/258305.258313