

Measuring and Modeling Variability using Low-Cost FPGAs

Michael Brown, Cyrus Bazeghi, Matthew Guthaus, and Jose Renau
Dept. of Computer Engineering, UCSC

FPGA IMPLEMENTATION

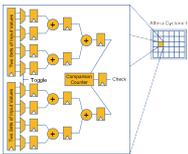
To cover the majority of a chip, we created 80 "builds", each covering roughly 1% of the total chip area available for programmable logic. Each build was comprised of a test block and clock control circuitry. In order to reduce the impact of IR drop in supply circuits, each build with the self-checking circuit was tested separately. This removes the impact that one self-checking circuit may have on an adjacent block.

FPGA SETUP



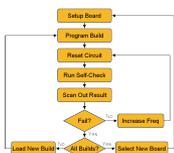
Altera Cyclone II - 60mm² die with external clock generation

TEST CIRCUIT



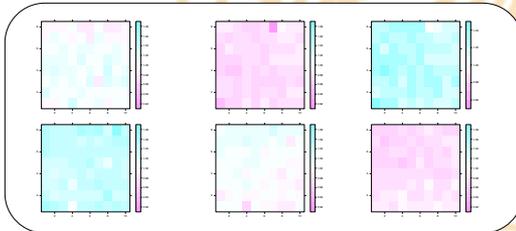
Not all the available block space is used, but we get 75% coverage of physical area over many project builds. Each build is an individual test block. A test block uses 95% of the combinational cells and 58% of the register cells available in its physical region. It is then exported as a hard macro to retain its placement and routing.

TEST PROCEDURE

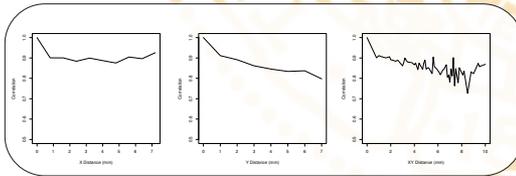


FPGA VARIABILITY

The six figures above below significant WID spatial correlation. We found that there are a couple of outlier points on each die, but it was confirmed that these values are repeatable over a period of time. We also found that there is significant D2D variation.



The bottom figures show plots analyzing the correlation coefficient. The xy plane is quite noisy, but has a general trend of decreasing correlation with distance. We also found that the x- and y-dimensions do not show equal spatial correlation. The x-dimension has a correlation of roughly 0.9 through about 7mm. The y-dimension has decreasing correlation down to about 0.8 at 7mm. Since the FPGA die size is 8mm x 8mm, we are unable to verify correlations beyond this distance at this time.



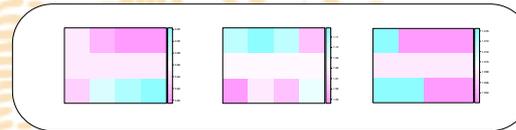
RESULTS

The graphs at right show overall variability distribution estimated for the Pentium D 800 series (near) and the T1 Niagara (far) using the FPGA data without spatial correlation model (Bowman) and with spatial correlation (Proposed). We found that modeling both with and without spatial correlation gave close estimates. This suggested that modeling the spatial correlation for the FMAX model does not have a significant effect on the results.

Even so, we are pleased with our measurement infrastructure and the insights it provides into variation and spatial correlation in multi-core processors. We were able to obtain high resolution variability measurements of real FPGAs and use those measurements to predict variability in off-the-shelf processors.

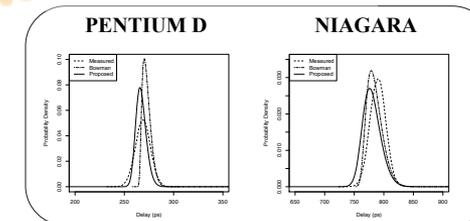
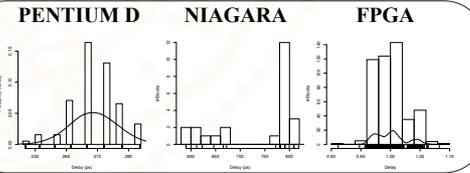
PROCESSOR VARIABILITY

We measured performance for 11 Pentium D 800 series processors (22 cores) and 3 Niagara T1 processors (24 cores). Below are the variation maps for the Niagaras. The right most only has two values because the Niagara will only perform tests as long as the master core is working. In the last Niagara the master core was one of the first to fail, so no data could be obtained after that.



Since the Pentium cores were failing at as much as 45 degrees apart from each other, we plotted the relationship of temp. vs. max frequency and found the relationship to be nearly linear. From this we were able to easily adjust all measured frequencies to be equivalent in temperature. The slight slump of the Pentium distribution on the left side is believed to be because we were unable to obtain any Pentium D 830s, only 820s and 840s. With the additional data available, the slump would likely fill out.

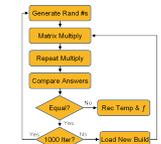
Below for comparison are the adjusted distributions of Intel, Sun, and Altera chips.



PROCESSOR SETUP

The maximum performance of a core is measured as the frequency where a core no longer works properly. In the Sun T1 Niagara cores this is done with a built-in self-test. In the Intel Pentium D cores this is done by turning up the CPU frequency in the BIOS settings. For both processors we record the temperature at which the failure occurred and adjust to the frequencies for the different temperatures. The Pentium system was built to ensure that none of the other components would be the cause of failure. The Niagara system was able to verify this on its own with a service card and its BIST.

PENTIUM TEST PROCEDURE



Tests for the Pentium D were run by setting up a self checking matrix multiply that fit into the cache space. Reducing chip I/O to a minimum would help to eliminate other system components as the cause of failure. We also tried to account for processor binning by the manufacturer by tracking sales records of major distributors.

FMAX MODEL

Estimated processor data was calculated with the FMAX model proposed by Bowman. The model calculates a distribution of the expected maximum frequency of a design based on determined WID and D2D variation, and design elements like the number of critical paths in the design and the nominal delay through a gate.

We attempted to validate the FMAX model as well as expand on it by introducing spatial correlation to the model.

	#Critical Paths	FO4	T _{rand}
Pentium D	1000	12	21.5 ps
Niagara	100	35	21.5 ps