

On-chip Timing Uncertainty Measurements on IBM Microprocessors

R. Franch, P. Restle, N. James¹, W. Huott², J. Friedrich¹, R. Dixon¹, S. Weitzel¹, K. Van Goor³, G. Salem⁴

IBM Research, Yorktown Heights, NY, ¹IBM STG, Austin, TX, ²IBM STG, Poughkeepsie, NY, ³IBM STG, Rochester, MN, ⁴IBM STG, Burlington, VT

Abstract

Timing uncertainty in microprocessors is comprised of several sources including PLL jitter, clock distribution skew and jitter, across chip device variations, and power supply noise. The on-chip measurement macro called SKITTER (SKew+jITTER) was designed to measure timing uncertainty from all combined sources by measuring the number of logic stages that complete in a cycle. This measure of completed delay stages has proven to be a very sensitive monitor of power supply noise, which has emerged as a dominant component of timing uncertainty. This paper describes the Skitter measurement experiences of several IBM microprocessors including PPC970MP, XBOX360TM, CELL Broadband EngineTM, and POWER6TM microprocessors running different workloads.

1. Introduction

Understanding timing uncertainty is important in order to know how to properly budget it in a microprocessor's cycle time. The main components of timing uncertainty have historically been PLL jitter, clock distribution skew and jitter, across chip line width variation (ACLV) and power supply noise effects. Often, it is treated as a simple sum of the estimates for these components and is lumped into the cycle time as a guard band. This guard band is time lost in the cycle since it is time that is taken away from doing real combinational work. Being able to directly measure timing uncertainty from all combined sources not only allows for better budgeting but also allows for learning from current designs and can highlight areas of emphasis for future designs.

The on-chip measurement macro called SKITTER (SKew+jITTER) was designed to measure timing uncertainty from all combined sources [1]. With continued scaling, power supply noise effects have become more important in determining timing uncertainty. Power supply noise results in a complex interaction in the delays of both the clock distribution paths and the logic paths. Since these effects are not independent, it makes more sense to measure the total combined effect instead of the delay variations in the clock and data paths individually. The effects from all combined sources of timing uncertainty are felt by the Skitter circuits. This results in a variation of

the number of delay stages in the Skitter circuit that complete in a cycle, just as it results in a variation of the number of logic stages that complete in path in the microprocessor. This measure of completed delay stages has proven to be a very sensitive monitor of power supply noise, which has emerged as a dominant component of timing uncertainty.

Skitter circuits have been placed in the core and the nest (non-core) regions on several IBM microprocessors. This paper describes several uses of the Skitter circuits. In addition to measuring timing uncertainty, they have been used to monitor the on-chip clock duty cycle measured at the Local Clock Buffers (LCBs) at the end of the clock distribution network. Skitter circuits have also been used to detect power supply noise events and correlate those events to increased switching activity at specific cycles in the instruction stream. In this usage, an instruction stream that causes a power supply noise event is run repetitively and the Skitter circuit is read out cycle by cycle, effectively recovering the on-chip VDD waveform. This has allowed the study of power supply noise events under different instruction streams (workloads). Fixes for the power supply noise can then be evaluated using Skitter.

These measured results have raised the question of how to best test the complicated interaction between the packaging and decoupling schemes and the varied workloads that run on multi-core chips. Some suggestions for testing for worst-case scenarios of power supply noise events between cores are presented.

2. Skitter Circuits and Operating Modes

2.1 Edge Capture and Accumulate Circuit

Skitter contains a latched-tapped delay line of 129 low fan-out inverters, for the best timing resolution, with a nominal delay of 5-8 ps depending on the technology and threshold voltage. The delay line and sampling latches form an edge-capture circuit, shown in Figure 1, which has been tuned to capture rising and falling edges symmetrically, without preference. An edge is detected in the chain when consecutive inverters have the same output (either both 0's or both 1's). The sampling latches take a snapshot of the state of the inverter chain every cycle. They are regular scannable master/slave latches from the standard cell library and the delay chain inverters are made using low-Vt, regular-Vt, or high Vt devices

The workloads that were selected for study are called IDLE, TRASH, and SMOKE. IDLE is a pattern that limits the amount of on-chip switching activity, and this pattern is often used as a reference since it provides a quiet, low power baseline that can be used to compare to other noisy and high-power workloads. TRASH is a pattern that generates random instructions and SMOKE is a workload that tries to maximize power consumption.

The Skitters were operated in single-sample mode and for each workload, 100 measurements were taken and the edge locations were plotted in a histogram as shown in Figure 9. In addition to the 100 single samples, the Skitters were also put into sticky mode for each workload and the variation that was measured in sticky mode is plotted on the same graph as a horizontal line on the x-axis. What was found is that each workload had its own unique “signature” on the Skitter data. In particular, TRASH produced the most variation as indicated by the sticky data (the horizontal line at the base of the histogram), even though the 100 sample histogram had only minor variation. This seems to imply that most of the time, TRASH produces low VDD noise, but if we monitor continuously using sticky mode, on rare occasions TRASH will cause some noise event that will cause a large variation in the edge locations. The sticky mode variation for TRASH was even larger than for SMOKE, the high power workload. The mean slow-down for SMOKE was as expected, since SMOKE will cause the chip to draw more current and result in a larger IR drop. The use of histograms to plot Skitter single-sample data combined with sticky mode data on the same histogram proved to be a useful way to examine the effect of different instruction streams on the timing uncertainty.

3.4 POWER6 Processor

The POWER6 is a dual core PowerPC microprocessor [4] that was built with two power grid designs for comparison. In one design, the core power grids were split from the nest (non-core) region of the chip and from each other. In the other power grid design, all power grids were connected together. Skitter was used to help compare the cycle-by-cycle timing variations from these two options [5]. To evaluate this, Skitter was used in what is called oscilloscope mode.

In oscilloscope mode, the instruction stream that causes a power supply noise event is run repetitively. The Skitter is read out, cycle by cycle, and the bins containing edges are recorded. The bin number is then plotted versus cycle number (or equivalently the bin number is converted to inverter delay using a calibration). This plot effectively reveals the on-chip VDD waveform as a function of the cycle number in the instruction stream. The VDD dips can then be correlated to cycles with large switching activity. Figure 10 shows an example of this on POWER6 with

split power grids. One core (core 0) is running IDLE, the quiet pattern while the other core (core 1) is running TRASH, the random instruction pattern. The Skitter in the noisy core running TRASH measures a large edge

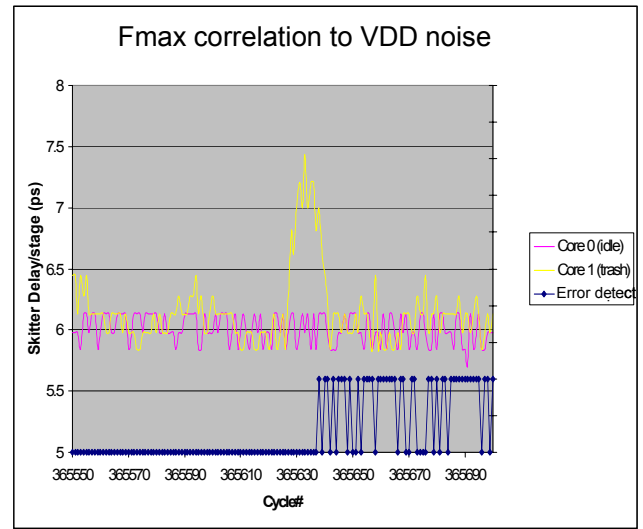


Fig 10:Skitter in o’scope mode detects cycles causing Fmax fail

variation during cycles where the switching activity suddenly jumps -- near cycle # 365630 in the instruction stream. This coincides with the detection of the fail at Fmax. The Skitters in the core running IDLE (core 0) show little edge variation, consistent with low switching activity.

By contrast, the POWER6 design with the connected power grids shows much less timing variation when running the same instruction stream. This is mainly due to the fact in the connected design, the core power grid share the “quiet” decoupling capacitance of the nest, which tends to stabilize the power grid. Figure 11 shows Skitter measurements that compare split to connected power grids. For increased resolution, multiple full-cycle edge locations were added. The noise is significantly reduced in

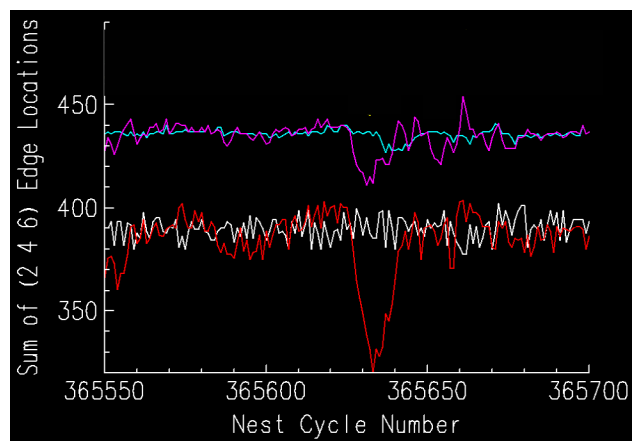


Fig 11: Measured comparison of noise on connected core power grids (top traces) and split power grids (bottom traces)

the connected power grid design.

4. Discussion

Since connected power grids seem to be better from a VDD noise standpoint, this raises the possibility of interaction of noise events from one core to another. For example, what happens if one core has a VDD dip and sends a noise “wave” over to another core. If the wave from the first core should arrive at the same time that second core generates its own VDD dip, then the two events might reinforce at the second core, resulting in a much larger dip at the receiving core.

The propagation speed of such a wave was measured on POWER6 by looking at the response of two Skitters, one in each core. The wave travels 9 mm in 4 ns. The question of how best to test for core to core VDD noise interaction will involve this wave travel time. If the receiving core is 9 mm away, then testing for the worst case noise scenario will involve delaying the starting time of the instruction stream running on the receiving core by the travel delay (4

ns in this case). In this way, the receiving core will be executing the offending cycles just when the VDD wave arrives from the sending core. Figures 12a and 12b are package simulations that show this situation. In Figure 12a, the wave is launched and in Figure 12b it arrives at the receiving core just as the receiving core generates its own VDD noise event. Clearly the worst-case noise occurs when the two VDD dips reinforce each other at the receiving core.

5. Conclusions


The on-chip Skitter circuit combined with a variety of test modes and analysis methods has been found to be valuable in the characterization of several issues on different microprocessors. The sensitivity to process variations, power supply noise, jitter, duty cycle, and skew have made Skitter useful in maximizing performance and reliability. In particular, the sensitivity of Skitter to power supply noise has made it an effective tool for recovering the on-chip VDD waveform during any workload. The measurement capability of Skitter has established it as a standard tool for testing a range of parameters on IBM chip products.

Acknowledgements


The authors would like to acknowledge Tom Strach, Brian Monwai, Steve Wilson, Tim McNamara, Tim Skergan, Steve Jones, John Reilly, Otto Wagner, Otto Torreiter, Andy Maki, Pong-Fei Lu, Tom Bucelot, and Jeff Burns for their contributions to this work. Also to be recognized are a large number of people in various test groups throughout IBM who have generated data and contributed to the growing knowledge base of Skitter measurements.

References

- [1] P. Restle et al. “Timing uncertainty measurements on the Power5 microprocessor”, *ISSCC Digest of Technical Papers*, Feb. 2004 pp 354-355
- [2] D. Pham et al, “The Design and Implementation of a First-Generation CELL Processor”, *ISSCC Digest of Technical Papers*, Feb. 2005, pp. 184-185.
- [3] E. Cohen et al. “A 64B CPU Pair: Dual- and Single-Processor Chips”, *ISSCC Digest of Technical Papers*, Feb 2006 pp 106-107
- [4] J. Friedrich et al., “Design of the POWER6™”, *ISSCC Digest of Technical Papers*, Feb. 2007, pp. 96-97
- [5] N. James et al., “Comparison of Split versus Connected Core Supplies in the POWER6 Microprocessor”, *ISSCC Digest of Technical Papers*, Feb. 2007, pp. 298-299



(a)



(b)

Fig 12: Simulation of VDD noise wave starting at core 1 (a) and the VDD noise wave reinforced at core 2. (b)