An Industrial Case Study of Low Cost Adaptive Voltage Scaling Using Delay Test Patterns

Mahroo Zandrahimi*, Philippe Debaud†, Armand Castillejo†, Zaid Al-Ars*
*Delft University of Technology, The Netherlands
{m.zandrahimi, z.al-ars}@tudelft.nl
†STMicroelectronics, Grenoble, France
{philippe.debaud, armand.castillejo}@st.com

Abstract–In deep sub-micron technologies, the increasing effect of process and environmental variations has lead chip manufacturers to use adaptive voltage scaling techniques in order to adapt operation parameters exclusively to each chip. The increasing effect of process variation is limiting the effectiveness of current chip monitoring approaches, such as on-chip performance monitor boxes (PMBs), which results in yield loss and high design margins, thus high power consumption. This paper proposes an alternative solution for adaptive voltage scaling using delay test patterns, which is able to eliminate the need for PMBs, and thus the long expensive characterization phase of tuning PMBs to each design, while improving the yield as well as power optimization. Results show, using an industrial grade 28nm FD-SOI library developed for low power devices, that delay testing for performance prediction reduces the inaccuracy down to 1.85%.

I. Introduction

Adaptive voltage scaling (AVS) has become a standard approach used by chip manufacturers to ensure low power consumption of their devices [2]. The effectiveness of the AVS approaches depends on appropriately predicting the performance of every manufactured device under specific input voltage values. This device-specific prediction is realized by using on chip performance monitoring boxes (PMBs) integrated on each device, that allows fast performance prediction and voltage pairing during production. Fig. 1 shows an example of a chip, on which various kinds of PMBs are distributed. The figure shows two PMBs created using PMOS and NMOS speedometers that indicate the speed of PMOS and NMOS transistors, while the third shown PMB is a critical path replica designed based on the most used logic cells extracted from the potential critical paths of the design. During production and based on the frequency responses from these monitors, chip performance is estimated, and corresponding voltage is fused for that operation point [1].

However, the correlation process during the characterization stage (i.e., finding the correlation between PMB responses and the actual frequency of the circuit) makes these techniques very expensive, since it should be done for an amount of test chips representative of the process window to make sure (for all manufactured chips) performance prediction based on PMB responses is correlated with application behavior.

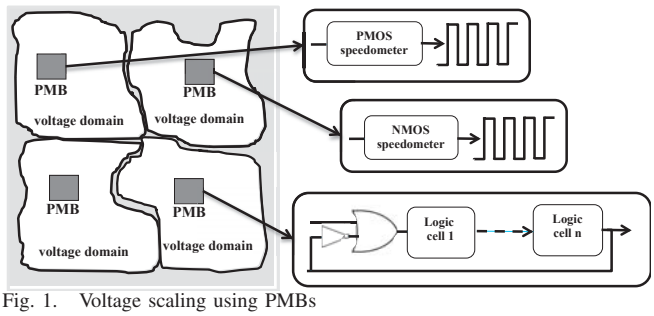

Fig. 1. Voltage scaling using PMBs

In this paper we introduce a cost effective approach for performance prediction during production using small delay defect (SDD) and path delay (PDLY) test patterns, which can be used for general logic as well. We investigate the proposed approach in terms of accuracy and effectiveness using 29 ISCAS'99 benchmarks with an industrial grade 28nm FD-SOI library for 42 different process corners with different characteristics in terms of process and environmental variations as well as aging.

The rest of this paper is organized as follows. Section II proposes the concept of using delay faults for performance prediction. Evaluation of the proposed approach is presented in Section III using simulation results on ISCAS'99 benchmarks. Section IV concludes the paper.

II. AVS using delay testing

In this paper, we propose an innovative test flow for voltage estimation using delay testing during production. Since delay testing covers many path-segments of the circuit [3], [4], it can be a better performance representative than PMBs. Such an approach has a number of unique advantages as compared to PMB-based approaches. Since delay testing is performed to explicitly test for actual chip performance, the expensive effort of correlating PMB responses to chip performance during the characterization stage of manufacturing is not needed anymore, which reduces the cost and time to market dramatically. Moreover, as functional patterns are not used anymore, the delay testing approach could be a solution for general logic, and not only for CPU and GPU components. And last but not least, this approach makes using PMBs redundant, which saves silicon area as well as PMB design time.

The proposed flow performs a binary search to identify the minimum voltage (Vmin), at which the chip can pass all delay test patterns. The following steps are performed for each operation point of the chip: 1. Apply chip setup at nominal values and initialize variables, 2. Set supply voltage to Vmax and wait for stabilization, 3. Apply transition fault at speed test, 4. If the chip fails the test, discard it, otherwise, 5. compute new values and do a binary search to find Vmin. Conversion from Vmin to Fmax might be required depending on either performance estimation is done for yield enhancement or power optimization.

978-3-9819263-0-9/DATE18/©2018 EDAA

TABLE I. ERROR OF SDD AND PDLY VERSUS STA

Benchmark	SDD50	SDD500	PDLY100	PDLY1000	PDLY10000	Benchmark	SDD50	SDD500	PDLY100	PDLY1000	PDLY10000
b01	0.79%	0.79%	0.33%	0.33%	0.33%	b15	2.77%	1.56%	2.46%	0.80%	0.80%
b02	4.33%	4.33%	0.11%	0.11%	0.11%	b15 1	3.60%	1.08%	3.25%	0.57%	0.57%
b03	4.12%	4.12%	4.05%	4.05%	4.05%	b17	3.43%	1.32%	3.69%	2.31%	1.82%
b04	1.70%	1.70%	1.70%	1.70%	1.70%	b17 1	4.37%	3.18%	4.99%	1.89%	1.89%
b05	1.21%	1.21%	1.21%	1.21%	1.21%	b18	5.00%	4.86%	10.54%	0.14%	0.46%
b06	4.36%	4.36%	3.64%	3.64%	3.64%	b18 1	8.13%	6.92%	7.96%	4.02%	4.03%
b07	5.21%	5.21%	2.20%	2.20%	2.20%	b19	11.77%	11.16%	12.35%	5.30%	4.58%
b08	2.84%	2.84%	1.95%	1.95%	1.95%	b19 1	8.83%	8.82%	8.82%	8.76%	8.70%
b09	7.42%	7.42%	7.50%	7.50%	7.50%	b20	8.04%	3.33%	11.69%	1.50%	1.50%
b10	0.18%	0.18%	0.05%	0.05%	0.05%	b20 1	9.75%	7.24%	12.36%	0.43%	0.43%
b11	0.20%	0.20%	0.20%	0.20%	0.20%	b21	7.03%	5.86%	8.56%	0.50%	0.50%
b12	1.75%	1.67%	1.82%	1.82%	1.82%	b21 1	2.47%	2.16%	4.45%	0.33%	0.33%
b13	2.35%	2.35%	2.35%	2.35%	2.35%	b22	6.34%	5.07%	10.29%	0.34%	0.17%
b14	12.16%	6.52%	16.35%	0.23%	0.23%	b22 1	8.44%	5.65%	12.16%	10.90%	0.41%
b14 1	10.15%	3.77%	13.35%	0.22%	0.22%	Average	5.13%	3.96%	5.87%	2.25%	1.85%

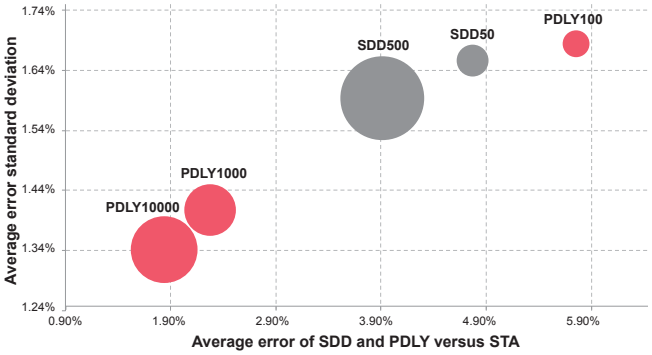

Fig. 2. Average error vs average standard deviation of error for 50 and 500 SDD pattern sets, and 100, 1000 and 10000 PDLY targeted paths; the size of the bubble represents the size of the pattern set used

III. EVALUATION

In this subsection we explore if SDD and PDLY tests correlate with the actual frequency of the circuits. We use low-power 28nm FD-SOI libraries to compare SDD and PDLY maximum frequency versus the critical paths of ISCAS'99 benchmarks using SYNOPSYS tools. We used Design Compiler in topographical mode for physical synthesis, Primetime for static timing analysis (STA), Tetramax for automatic test pattern generation (ATPG), and Vcs for back annotated simulation. Since functional pattens are not available for ISCAS'99 benchmarks, we use STA instead as a reference for comparison versus SDD and PDLY frequencies. This choice can be justified by noting that STA represents worst case circuit performance, which can be used as a lower bound for maximum circuit frequency. On the other hand, any set of functional patterns cannot be complete, since it is very tricky to select an application which reflects the real system performance specially for complex systems. This especially true since identifying the most critical part of the application is not possible in most cases.

We compared the maximum frequency at which each test pattern can be performed for each benchmark versus STA results. Table I presents the differences, what we call error, between each delay test frequency and STA for all benchmarks. We generated the results for 5 pattern sets including 50 and 500 SDD, and 100, 1000, and 10000 PDLY patterns. As it can be seen in this table, with increasing pattern count, the error is reduced. Increasing SDD pattern count from 50 to 500, achieves an error as low as 3.96%. Furthermore, increasing PDLY targeted path count from 100 to 10000 improves the error down to 1.85%. Therefore, depending on the time invested on testing during production, the accuracy of performance prediction can be improved.

The measured error for SDD and PDLY means that in order to make sure the performance prediction is accurate enough, a margin should be added on top of the estimated performance. If the inaccuracy of performance prediction is predictable, it is possible to come up with a safe margin. Fig. 2 illustrates the average standard deviation of the error plotted versus the average error measured using SDD and PDLY for all the circuits in the ISCAS'99 benchmarks. The plotted measurements are represented by the size of the test pattern set (reflected by the size of the circle in the plot). The figure shows that the larger the size of the used test pattern set, the more predictable the performance prediction will be. Therefore, depending on the time invested on testing during production, the accuracy of performance prediction can be improved. PDLY100 shows the worst prediction among all test pattern sets, which means that the size of this pattern set is not enough for performance prediction purposes. More test patterns should be taken into account to increase prediction accuracy as PDLY10000 shows the best performance prediction among all test pattern sets. In general, we can conclude that path delay test patterns are more suitable for performance prediction, however, sufficient pattern set sizes should be taken into account.

IV. CONCLUSIONS

In this paper, we proposed an innovative test flow for adaptive voltage scaling using delay testing during production. Since delay testing is performed to explicitly test for actual chip performance, the expensive effort of correlating PMB responses to chip performance during the characterization stage of manufacturing is not needed anymore, which reduces the cost and time to market dramatically. We compared two approaches (SDD or PDLY) based on their accuracy as performance predictor and how many pattern counts is sufficient for accurate voltage adaptation. We performed simulations using industrial grade 28nm FD-SOI library for 2 SDD pattern sets including 50 and 500 patterns, and 100, 1000, and 10000 PDLY targeted paths. The results show that Increasing SDD pattern count from 50 to 500, achieving an error as low as 3.96%. Furthermore, increasing PDLY targeted path count from 100 to 10000 improves the error down to 1.85%.

REFERENCES

[1] B. Kruseman, A. Majhi, and G. Gronthoud, On Performance Testing with Path Delay Patterns, in VTS, 2007.

[2] N. B. Zain Ali, et al., *Dynamic Voltage Scaling Aware Delay Fault Testing*, in ETS 2006.

[3] M. Sauer, et al., *On the Quality of Test Vectors for Post-Silicon Characterization*, in ETS 2012.

[4] P. Das, et al., *On Generating Vectors for Accurate Post-Silicon Delay Characterization*, in ATS 2011.