Device Degradation and Resilient Computing

Peter Glösekötter
and Ulrich Greveler
University of Applied Sciences
Münster, Germany
Email: gloesek@ieee.org

Gilson I. Wirth
Univ. Federal do Rio Grande do Sul
UFRGS
Porto Alegre, Brazil
Email: wirth@inf.ufrgs.br

Abstract— **In this paper, we study the impact of recent scaling trends on device degradation effects as well as their impact on microarchitecture. A new binning strategy is proposed that takes into account the individual degradation of the processing nodes. The vitality of each processing node is derived from its duty-cycle. By means of random sampling and statistical analysis, the duty-cycle is captured in a resource and time efficient way so that it is feasible to observe more than 100 processing nodes with a single sampling unit. Applications of this kind of degradation handling can primarily be found where conventional reliability verification and life-test acceleration techniques such as burn-in cannot be used anymore due to their complexity and inherent constraints.**

Keywords: resiliency, microarchitecture, device degradation, NBTI, relaxation, duty-cycle, on-chip measurement

I. Introduction

Up to now, speed degradation experienced by products in customer hands is primarily addressed by production burn-in. Anyhow, for high performance devices power consumptions are already very high at normal operation condition, so that the life-test acceleration factor is diminishing [1], [2]. An increased burn-in voltage would increase power and chip temperature drastically. Since traditional stress validation methods have declining viability and backend checks are too late and too costly in the design cycle, it is becoming a serious concern for getting robust products to consumers in the future [3]. Therefore addressing reliability at the design stage becomes increasingly important as scaling can potentially have a negative impact on the circuit reliability [4], [2]. A tight RTL-to-Layout design flow would create the great opportunities to get quality and reliability requirements into the flow, but additional reliability information is required to guarantee the resilient realization during the design cycle.

In this paper, we implement a new binning strategy that takes into account the individual degradation level of the processing nodes. Thanks to this strategy, reliability verification and life-test acceleration techniques become obsolete. The proposed resilient microarchitecture continually detects errors, isolates faults, refreshes processing nodes, and thus adapts the hardware.

The rest of the paper is organized as follows. Section II describes different degradation effects including inherent scaling trends as well as the impact on circuit's integrity. Taking the involved implications into account, a novel circuit architecture is put forward in the subsequent section. The proposed architecture comprises a binning strategy taking into account the stress history of the individual processing node. Over-stressed nodes are consciously decommissioned and then run through a relaxation phase. The vital capacity of each node is checked on a frequent basis as explained in Section IV. Section V reveals an example of the proposed binning strategy. The last Section concludes and summarizes the paper.

II. Impact of Device Degradation

As gate oxide gets thinner than 4 nm, the threshold voltage change caused by negative bias temperature instability (NBTI) for the PMOS transistor has become the dominant factor to limit the life time which is much shorter than that defined by hot-carrier induced degradation (HCI) of the NMOS transistor [5], [3].

The failure criteria are a certain level of parametric shift. For instance, 25 to 100 mV shift in V_{th}, and 8 % to 12 % shift in I_d for the operating life of a common commercial product. For a given process, circuits are designed such that functionality is guaranteed as long as the failure criteria are not reached. Guardbanding is the most popular method of guarding against speed degradation due to hot carrier damage and NBTI. It is performed to ensure product performance specifications are met throughout product life.

Due to the complexity of today's integrated circuits, it is difficult to account for such device degradations. Using static timing analysis tools, the accuracy of pre-silicon max frequency prediction is in the order of 5-10% [6], while the end-of-life-degradation typically is of the order of 2-4% [7], [1]. Further, the circuit behavior and aging crucially depends on stress conditions. Hence, it is doubtful to predict reliably the device degradation at the pre-silicon stage. Apparently, its impact can only be limited by a set of appropriate design rules and suitable guardbands being estimated on the basis of measurement results, since a simple static analysis may provide an extremely pessimistic estimate and consequently, results in over-guardbanding.

Presuming NBTI and HCI would impact over time the circuitry in the same order of magnitude, an inverter's waveform is simply delayed as shown in Fig. 1 a). For the sake of simplicity the graphs in Fig. 1 show idealized rising and falling edges. The first two graphs of Fig. 1 a) show the input and unstressed output waveform at t=0. The third graph shows the

978-1-4244-1684-4/08/$25.00 ©2008 IEEE

Authorized licensed use limited to: UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL. Downloaded on July 30, 2009 at 15:46 from IEEE Xplore. Restrictions apply.

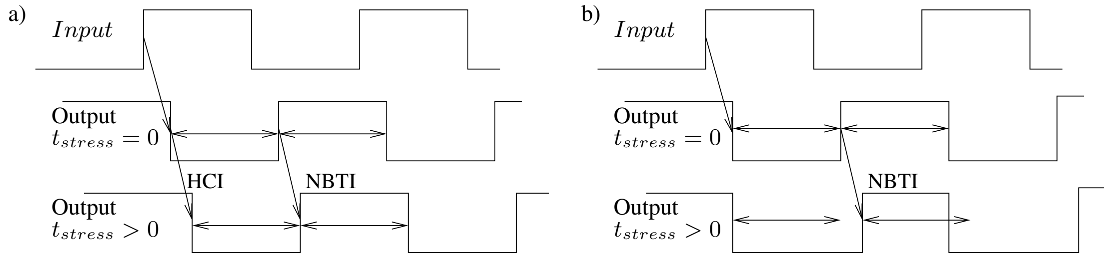

Fig. 1. CMOS inverter: a) Symmetric aging results in unaffected duty-cycle, b) NBTI dominated aging affects duty-cycle.

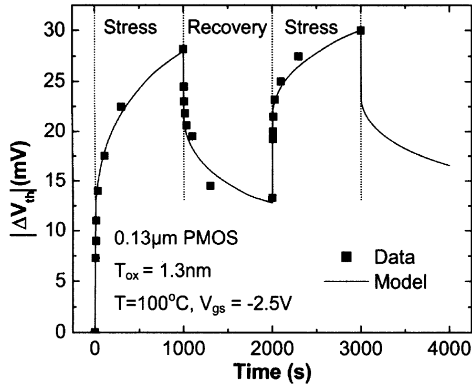

Fig. 2. Stress and recovery (NBTI) impact on ΔV_{th} [3].

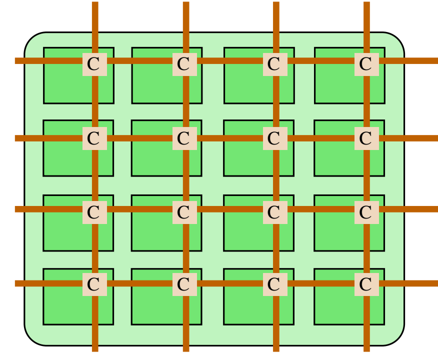

Fig. 3. Proposed resilient many-core Architecture.

output waveform after stress. Both edges rising and falling degrade in a symmetric fashion and obviously, the phase dimensions are preserved.

In contrast to HCI, NBTI is caused during static stress on the oxide even without current flow. As advanced digital systems tend to have longer standby time for lower power consumption, sensitivity of NBTI increases. Taking this trend into account, Fig. 1 b) depicts the impact of a DC stressed inverter. The asymmetry in degradation of the falling and rising edges results in a duty-cycle deformation for common real world applications.

The analysis of NBTI is more complex than that of HCI, since NBTI exhibits both stress and recovery behavior during dynamic circuit operation. PMOS transistors can recover from NBTI stress through the condition of $V_g = V_{DD}$ [8]. The impact of NBTI induced by static stress may be reduced by factor of 2–3. Fig. 2 reveals that V_{th} change under dynamic conditions is dramatically different from that in the static mode.

III. Proposed Architecture

Besides transient errors (Soft Errors) and gradual errors (Variations), time dependent degradations as addressed in Section II will make it impossible to design reliable systems as we know them today. One-time-factory testing is running out of steam and burn-in to catch chip infant-mortality will not be practical anymore [9], [2]. Hence, the hardware test has to be moved into the design cycle in such a way that the resilient realization is assured during design and is maintained during a product's life cycle.

Therefore, the in Fig. 3 proposed many-core architecture has to provide features like [9], [10], [11]:

- Dynamic on-chip testing
- Performance profiling
- Spare hardware
- Binning strategy
- Performance & power management
- Coarse-grain redundancy checking
- Dynamic error detection & reconfiguration

We propose a dynamic self-test and hardware-recovery strategy to meet these requirements and to account for the upcoming impractibility of one-time-factory testing. Each processing node in Fig. 3 is in one of the following four states:

1) busy
2) quiescent
3) test & recovering
4) broken

The *busy* state reflects the normal operation mode of the processing nodes. Each node can stay in this mode for a distinct time depending on its present load and stress-history. Nodes in the *quiescent* state are available for new assignments. Once a node's performance has fallen below a predefined threshold, it has to go through the *test & recovering* phase in order to remove the impact of NBTI dominant stress from PMOS devices. Relaxation can be accomplished though AC

Authorized licensed use limited to: UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL. Downloaded on July 30, 2009 at 15:46 from IEEE Xplore. Restrictions apply.

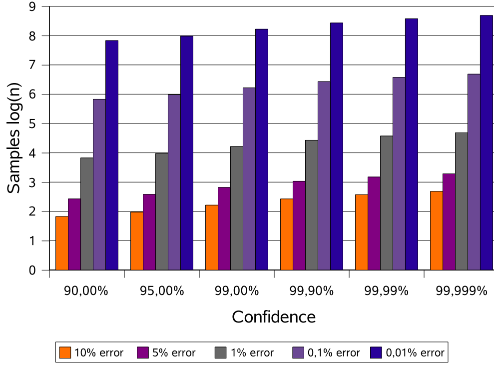

Fig. 4. Sample size versus Confidence Level at $p = 0.5$.

load or the condition of $V_g = V_{DD}$ [8], [12]. If after the $test\,\&\,recovering$ phase the performance degradation remains above a predefined threshold, the node is marked as broken. In this case the node is irretrievable lost due to variations or extreme degradation. The required clearance of the different states heavily depends on the underlying technology and on the usage model. For instance, as shown in Fig. 2 and thoroughly analyzed in [3] a 130 nm technology with $T_{ox} = 1.3\,nm$ exhibits at $T = 100°C$ a V_{th} stress recovery of a factor of 2–3 within time intervals of 10^3 seconds. Predictive models as shown in [3] predict for future technology nodes a reduction of the electric field across gate oxide due to the slow scaling of V_{th} and $T_{ox}(E_{ox} = \frac{V_{GS}-V_{th}}{T_{ox}})$. Although ΔV_{th} due to NBTI is reduced with such trend of scaling, with lower ratio of V_{DD}/V_{th} the transistor performance has increasing sensitivity to V_{th} change.

The appropriate threshold margin for entering the recovery state depends on the usage and availability model. The individual node's degradation is verified on a frequent basis through on-die degradation monitoring as described in Section IV. Adequate on-line learning algorithms provide solid scheduling strategies.

IV. Degradation Monitoring

A. Theoretical Framework

On line degradation monitoring may be performed using different approaches. In [13] a technique to derive the individual node degradation from a worst case inspection of the node in question is presented. In [14] a test structure that can be attached to each node for estimating the performance degradation is described.

To determine the degree of degradation and its impact on the duty-cycle, we pursue a statistical approach and apply the well-known Law of Large Numbers to estimate and the state of the signal to be measured at random instants of time. The idea to repeatedly capture the state of the signal using a random sampling technique was laid out in [15]. The occurrences of edges of the random clock are assumed to be statistically independent (i. e. un-correlated) from the signal of interest. Hence, we can capture all parts of the signal under measurement with an equal probability and use the data for statistical analysis. If the duty-cycle of the signal is p then the probability of capturing logic high or low in a single trial would be equal to p or $q = (1 - p)$, respectively. If X is the number of times logic high is captured at the edge of the random clock in a sample of n trials, from the Law of Large Numbers we have:

$$\lim_{n\to\infty} \frac{X}{n} = p \tag{1}$$

The accuracy and the confidence over the observed result depends on the value of n and should be kept high. $P = X/n$ is the observed proportion of number of ones captured in a sample of n trials. For a very large values of n the duty-cycle is approximated by the following relation [15] with $\pm z_c$ depending on the desired confidence interval:

$$p \approx P \pm \underbrace{z_c\sqrt{\frac{P(1-P)}{n}}}_{\alpha} \tag{2}$$

The observed duty-cycle can be mapped to an actual duty-cycle with a certain confidence level, and the error α can be minimized by calculating the sample size n using:

$$n = \left(\frac{z_c}{\alpha}\right)^2 P(1-P) \tag{3}$$

Fig. 4 shows the relation of the sample size with the desired confidence level and tolerable error level for measurement of a signal expected to have 50 % duty-cycle. Further, it reflects the quadratic relationship of n.

For a measurement error of less than 0.1 % and a confidence level of 99.99 % less than 10^6 samples are needed. Presuming an uncorrelated sampling clock of some moderate MHz, it takes less than $1s$ to measure the duty-cycle. If we further assume stress recovery intervals of 10^3 seconds as shown in Section III and we want to check each processing node 10 times during each degradation interval, then the duty-cycle has to be measured every $100s$. This provides sufficient headroom for monitoring the duty-cycle of about 100 nodes with a single sampling unit. If required, additional sampling units can be applied for monitoring further nodes.

B. Duty-Cycle Measurement

The random sampling unit is simply a gated flip-flop clocked with a random clock (see Fig. 5). However, the practical implementation of such a sampling circuit requires careful handling of metastability issues in case the clock of the latch and the input signal switch simultaneously [15]. The straightforward approach uses cascaded flip-flops to demetastabilize the sampled value of the input signal. The simple random sampling unit in Fig. 5 comprises two counters. *Counter 1* is loaded with the desired sample size (n) and *Counter 2* is reset on the transition of the *Sample* signal. At every rising edge of the clock signal, *Counter 1* is decremented, and *Counter 2* is incremented when the signal is sampled as high. Further sampling is stopped once *Counter 1* decrements to zero. Now, the duty-cycle can be derived from the value of *Counter 2*.

Authorized licensed use limited to: UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL. Downloaded on July 30, 2009 at 15:46 from IEEE Xplore. Restrictions apply.

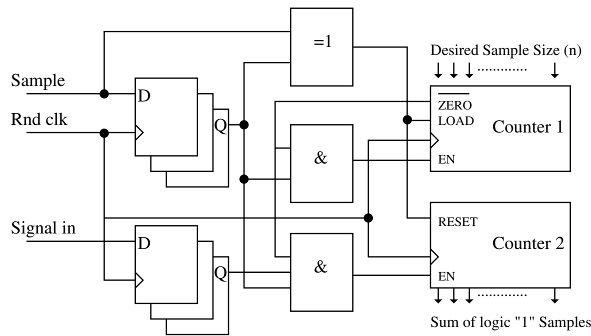

Fig. 5. Random sampling unit.

TABLE I

PROPOSED BINNING STRATEGY FOR AN EXEMPLARY POOL OF 128 NODES.

State	Index	Core-ID	Score
busy or quiescent	001	006	100
	002	117	100
	003	055	050
	004	017	050
	.	.	.
	.	.	.
	.	.	.
	100	114	033
recovering	101	003	025
	102	047	025
	.	.	.
	.	.	.
	.	.	.
	128	022	10

V. BINNING STRATEGY

The purpose of this binning strategy is twofold. First, over-stressed nodes that show a significant duty-cycle shift are decommissioned and run through a $recovering$ phase so that they become again available afterwards. Second, a uniform intra-chip aging should be guaranteed, targeting to maximize product life time. The processing nodes are rated during the duty-cycle measurement according to the number of times each node had to run through the recovering phase.

Let r_i be the number of times node i has entered the recovering phase, then the score of node i is equal to $\frac{100}{r_i+1}$. Exemplarily, table I depicts the proposed binning strategy for a pool of 128 processing nodes. The top 100 nodes (index 001 – 100) are disposable for new jobs and therefore go into to the $busy$ or $quiecent$ state. Nodes below this threshold are assigned to the $recovering$ phase unless they have not reached the $broken$ state before. Hence, the proposed strategy considers the individual degradation behaviors. For instance, a slow degrading node stays longer in the top 100 than the average node and nodes that have entered the $recovering$ phase various times will stay for a longer period in the $recovering$ phase.

This strategy can be augmented with other suitable criteria, e. g. the local temperature of a node's physical neighborhood can be taken into account in order to enhance a decision to activate a node in a rather *cold* neighborhood. Chip temperature may depend on the actual workload. Temperature sensors may be used in various parts of a chip and the temperature readings included into the failure prediction strategy. This way, anticipatory stress factors can be taken into consideration.

VI. CONCLUSION

In this paper, the impact of recent scaling trends on circuit architecture is analyzed in terms of device degradation since the life-test acceleration factor is diminishing and reliability verification is becoming a daunting task at design cycle. A new binning strategy is proposed that comprehensively exploits the underlying hardware architecture by considering the individual stress history of the processing nodes. The node's capabilities are benchmarked through random on-die sampling and statistical analysis of the duty-cycle. Over-stressed nodes are assigned to run through a relaxation phase and put back into the pool of available nodes.

REFERENCES

[1] H. M. et al., "Assessment of a 90nm pmos nbti in the form of products failure rate," in *Int. Conference on Microelectronics Test Structures*. IEEE, April 2005, pp. 89–94.
[2] S. Borkar, "Thousand core chips - a technology perspective," *DAC*, 2007.
[3] R. V. et al., "Modeling and minimization of pmos nbti effect for robust nanometer design," in *DAC*, July 2006, pp. 1047–1052.
[4] Z. L. et al., "Build-in reliability analysis for circuit design in the nanometer technology era," in *Int. Conference on Integrated Circuit Design and Technology*. IEEE, 2004, pp. 205–210.
[5] N. K. et al., "The impact of bias temperature instability for direct-tunneling ultra-thin gate oxide on mosfet scaling," in *VLSI Symp.*, 1999, pp. 73–74.
[6] B. N. Lee, L.-C. Wang, and M. S. Abadir, "Refined statistical static timing analysis through learning spatial delay correlations," in *DAC*, 2006, pp. 149–154.
[7] ITRS, "International technology roadmap for semiconductors," http://www.itrs.net, Tech. Rep., 2005.
[8] S. Rangan, N. Mielke, and E. Yeh, "Universal recovery behavior of negative bias temperature instability," in *IEDM*, 2003, pp. 341–344.
[9] S. Borkar, "Electronics beyond nano-scale cmos," in *DAC*, July 2006, pp. 807–808.
[10] R. I. Bahar, D. Hammerstrom, J. Harlow, W. H. Joyner, C. Lau, D. Marculescu, A. Orailoglu, and M. Pedram, "Architectures for silicon nanoelectronics and beyond," *IEEE Computer*, vol. 40, no. 1, pp. 25–33, January 2007.
[11] D. Sylvester, D. Blaauw, and E. Karl, "Elastic: An adaptive self-handling architecture for unpredictable silicon," in *IEEE Design & Test of Computers*. IEEE CS and IEEE CASS, 2006, pp. 484–490, process Variation and Stochastic Design and Test.
[12] T. Nigam and E. B. Harris, "Lifetime enhancement under high frequency nbti measured on ring oscillators," 44^{th} *Annual International Reliability Physics Symposium*, pp. 289–293, 2006.
[13] B. Paul, K. Kang, H. Kufluoglu, M. Alam, and K. Roy, "Impact of nbti on the temporal performance degradation of digital circuits," *IEEE Electron Device Letters*, vol. 26, no. 8, pp. 560–562, August 2005.
[14] M. Ketchen and M. Bhushan, "Product-representative "at speed" test structures for cmos characterization," *IBM J. Res. & Dev.*, vol. 50, no. 4/5, pp. 451–468, July/September 2006.
[15] R. Bhatti, M. Denneau, and J. Draper, "Duty cycle measurement and correction using a random sampling technique," in *48th Midwest Symposium on Circuits and Systems*, vol. 2, Aug. 2005, pp. 1043 – 1046.

Authorized licensed use limited to: UNIVERSIDADE FEDERAL DO RIO GRANDE DO SUL. Downloaded on July 30, 2009 at 15:46 from IEEE Xplore. Restrictions apply.