

Low Overhead in Situ Aging Monitoring and Proactive Aging Management

Christoph Niemann, Tim Wegner, Dirk Timmermann
Institute of Applied Microelectronics
and Computer Engineering
University of Rostock, Germany
Email: christoph.niemann@uni-rostock.de

Frank Sill Torres
Department of Electronic Engineering,
Federal University of Minas Gerais
Belo Horizonte (MG), Brazil
Email: franksill@ufmg.br

Abstract—Post-Dennard scaling CMOS technologies suffer from considerable degradation due to increasing electrical fields caused by the lack of further reduction of the supply voltage. This aspect of aging is widely disregarded so far and cannot be addressed at design time by adding static margins anymore. Instead, it needs to be counteracted effectively at run time over the entire device lifetime. For this purpose, dynamic run-time approaches for aging management are required, relying on detailed information regarding the current system state. In this paper we propose a novel aging monitoring mechanism providing that crucial information at a marginal resource overhead. The current device degradation is measured via the aging-dependent delay variation, which can be quantified in situ with built-in tests exploiting the strictly monotonic relation between supply voltage and propagation delay. Furthermore, we suggest to utilize the information gained this way for a proactive aging-aware task mapping.

I. INTRODUCTION

Traditionally, integrated circuit reliability has been neglected at higher abstraction layers such as architecture and system layer assuming the underlying hardware to work free of faults within certain environmental conditions and a predefined lifetime. Simplifying the work at higher layers, this approach requires to consider worst-case assumptions for the specification of certain technology parameters (*e. g.* timing margins). However, due to the increasing impact of inherent variations of the fabrication process and varying environmental conditions during system lifetime, this approach proved to be far too pessimistic in many cases. Consequently, the potentially available capability of the target technology is not fully utilized.

Beyond that, aging issues become increasingly dominant within the latest technology generations. This is mainly based on the mismatch of concurrent down-scaling of size and supply voltage (*i. e.* post-Dennard scaling), which results in rising electrical field strengths and thermal power density. Strong electrical fields and high temperatures are major acceleration factors for gate oxide degeneration effects such as Negative Bias Temperature Instability (NBTI) and Hot Carrier Injection (HCI) [1]. The outcome of these aging effects is an additional increase in propagation delay and power consumption. Therefore, a treatment of those issues solely based on margins added at design time becomes insufficient due to increasing performance penalties. Consequently, there is an increasing demand for proactive aging management solutions working at run time that allow for dynamically adjustable margins. Aging

depends on various influences such as fabrication variations, environmental conditions and workloads. Since those are obviously not observable a priori, a suitable aging monitoring is required to provide the necessary information concerning the current system state. The major contributions of this paper are:

- We introduce a novel approach for in situ aging monitoring by measuring the aging-dependent delay shift at lowered supply voltage. This approach is applicable with a marginal overhead utilizing existing Dynamic Voltage Frequency Scaling (DVFS) resources.
- We propose an aging-aware task mapping based on the information gathered by the aging monitoring. It allows either to extend system lifetime or to reduce the aging margins of a chip.
- Our approach enables a better utilization of the potentially available capability of a technology, while it works independently of a specific CMOS technology.

The remainder of this paper is organized as follows. Section II focuses on related work regarding aging monitoring. Section III introduces implications and dependencies of degradation mechanisms in contemporary CMOS technologies. Section IV proposes our novel approach for aging monitoring and prediction. Thereafter, Section V presents an aging-aware task mapping, utilizing the information gathered by the procedure described in Section IV. Section VI focuses on the applied evaluation setup and results. Finally, Section VII concludes the paper.

II. RELATED WORK

A commonly used diagnostic method is Built-In-Test (BIT), a solution in hardware or software that monitors whether the system still matches previously defined requirements. Typically, BIT only monitors whether the system already fails to fulfill the requirements, but is unable to predict failures [2]. Approaches enabling the prediction of the system state mostly rely on aging-induced gradual shifts of the dynamic circuit behavior. Some approaches apply additional dummy circuits that solely measure aging effects. Those dummy circuits experience similar inter-die process variations and environmental influences as the target circuitry, while aspects like intra-die process variations and workload impact are neglected. For example, Simevski et al. [3] propose different types of aging monitors for HCI and NBTI. These monitors are exposed to

TABLE I. COMPARISON OF APPROACHES FOR AGING MONITORING AND PREDICTION (KEY: + GOOD, ○ NEUTRAL, - BAD).

Approach	Prediction	Workload awareness	Intra-die variation awareness	Overhead
BIT [2]	no	+	+	+
Simevski [3]	yes	-	-	○
Agarwal [5]	yes	+	+	-
Li [6]	yes	+	+	-
Our approach	yes	+	+	+

a workload causing maximum degradation, which is far too pessimistic in most cases.

Such issues can be avoided, if the delay shift is measured directly within the target circuitry, as proposed by Agarwal et al. in [4] and [5]. The applied aging monitors are attached to flip-flops in order to detect increased delays nearly violating the setup time restriction. For this purpose, a stability checker inspects the logic signal for transitions within a specific period of time previous to the clock edge. To mark this period, an additional phase-shifted clock is derived from the actual system clock employing a complex, aging-resilient delay circuitry. Avoiding the issues concerning intra-die variations and workload effects on aging, this approach suffers from an immense area penalty due to the required additional logic. Li and Seok [6] propose another aging measurement procedure based on aging-dependent delay shifts. Modified registers and a dedicated feedback network enable the reconfiguration of selected signal paths into ring oscillators (ROs). This approach is facilitated by the dependency of the RO frequency on circuit aging. Since RO frequencies also strongly depend on circuit temperature, the measurements are executed at a specific lowered supply voltage reducing the temperature impact. While this approach seems to support a satisfying precision it suffers from a tremendous area and power overhead caused by the feedback network and the specialized registers. In addition, dedicated maintenance states are required, since system function cannot be provided during the RO measurements. A brief qualitative comparison of the introduced approaches including our concept regarding selected aspects is displayed in Table I. As it can be seen, existing approaches mostly suffer from high resource overhead.

III. CMOS AGING MECHANISMS

NBTI and HCI both increase the threshold voltage and reduce the drain current resulting in a reduced switching speed. Among these two effects the shift of threshold voltage V_{TH} is the predominant one. Its influence on the propagation delay T_P of a gate can be quantified utilizing the alpha-power-law model [7].

$$T_P = \gamma \cdot \frac{W}{L} \frac{C_{load} V_{DD}}{(V_{DD} - V_{TH})^\alpha} \quad (1)$$

W and L are the channel width and length, C_{load} is the load capacity, and α and γ are technology dependent parameters. Typical values are $\alpha = 1...1.5$ for short-channel devices [7]. All aging effects in integrated circuits highly depend on temperature. Even though the extent differs and is more distinct for certain effects (*e.g.* NBTI), the dependency is always an exponential one. Referring to [1] the dependency on temperature can be calculated as:

$$\lambda = \frac{1}{MTTF} = e^{-E_{AA}/k_B T} \quad (2)$$




Fig. 1. Components of the overall timing requirements.

λ is the failure rate, which is directly associated to the mean time to failure (MTTF) by inverse proportionality. E_{AA} is the apparent activation energy depending on the type of degradation process and the used CMOS technology. k_B is the Boltzmann constant and T is the temperature given in Kelvin.

IV. AGING MONITORING

To provide an information base for proactive management decisions, an appropriate aging management has to cover specific requirements. Those are:

- Ability to predict a future system state
- Suitability for a broad spectrum of applications
- Feasibility with negligible resource overhead (*i.e.* minimum chip-area, power and timing slack)
- Easy integration into the regular design process and avoidance of complex design for testability requirements

Aging itself is not a conventionally measurable quantity. It is solely measurable indirectly by measuring aging-affected parameters. As described in Section III, NBTI and HCI result in degradation of timing behavior, gradually consuming the respective timing guard bands (see Figure 1). Hence, the remaining timing margin of an integrated circuit or the decrease of that margin in the course of time can be utilized as an aging indicator. We propose a procedure applying simple BITs in combination with a mechanism artificially slowing down the analyzed circuitry. In detail, this approach evaluates how much a circuit needs to be slowed down, until timing errors start to occur. The degree of tolerable slowdown provides information regarding the remaining timing margin, and thus, the amount of device degradation. Aiming at such a temporary, precisely controllable slowdown a manipulation of V_{DD} is suitable. It can be implemented on minimum overhead by reusing existing DVFS resources, such as a precisely controllable supply voltage.

The influence of a lowered V_{DD} on the propagation delay can be quantified utilizing the alpha-power-law model of Equation 1. Figure 2 illustrates this dependency. The different graphs represent one and the same circuit at different states of aging-induced degradation (degradation is represented as relative shift ΔV_{TH} of the threshold voltage). When the propagation delay of a path exceeds the maximum tolerable delay, the state of the following register will be corrupted, given an appropriate switching activity. By applying suitable tests at incrementally lowered V_{DD} , the value of V_{DD} at which the

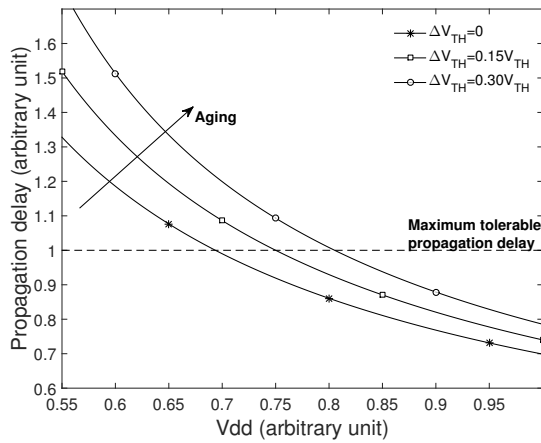


Fig. 2. Aging dependent propagation delay at lowered V_{DD} .

circuit begins to fail can be determined accurately. This value of V_{DD} depends on the aging a circuit has already suffered, as well as on process variations. It is important to remember that aging is a time-dependent mechanism, while manufacturing variations occur only once. Thus, a separation of these two factors can be facilitated by considering the chronological sequence of measured values (*i.e.* aging is measured by comparison of the first failing V_{DD} to the corresponding value of the first test applied to the new circuit).

As BITs we chose simple stored-pattern tests for they are perfectly suitable and realizable at minimum overhead. Such tests check a circuitry by sequentially applying a number of stored input patterns to the inputs and comparing the output with the stored expected results. In case of a mismatch the device fails the test. It is of major importance that the resources required for BIT execution need to work reliable at lowered V_{DD} during the test procedure. This is achieved by applying more severe timing constraints to the BIT infrastructure. The timing margins of the different signal paths of a design are varying in wide ranges. For reasons of efficiency it is desirable to quantify the aging-induced delay shift at preferably low voltage shifts. Therefore, it is essential to choose those paths with the smallest timing margins. This is not necessarily just a single critical path, but rather a group of paths with small timing margins widely covering different parts of the overall circuitry. Hence, the test patterns need to be specially designed for those requirements. This process can be easily automated and integrated into the usual design flow, since the timing behavior of the paths is already estimated during static timing analysis.

During the test procedure, the artificially slowed down circuitry is incapable of providing its normal service. Therefore, an appropriate maintenance state has to be defined. Since aging is an inherently gradual and slow process, the interrupt of regular operation introduced by the proposed method is tolerable due to the infrequent application of the test. Depending on the specific scenario, we assume weekly or even monthly application of the aging monitoring to be sufficient. For this reason, the monitoring procedure can be integrated in power-on routines, already existing maintenance procedures or in idle or power-off modes of the corresponding module. On the other hand, such out-of-service maintenance states permit to deal

with temperature disturbances. After an appropriate idle time (ranging from some hundred milliseconds to a few seconds) an equal temperature can be assumed on the entire surface of the chip [8]. Thus, a single-point measurement of the temperature allows determining and excluding the impact of temperature on the aging calculation.

V. RELIABILITY-AWARE TASK MAPPING IN MANY-CORE SYSTEMS

An unprotected system fails as soon as its first sub-unit fails [9]. Hence, it is desirable to protect the weakest sub-units even at the expense of additionally straining other sub-units to a reasonable extent. This can be achieved by a run-time task mapping, assigning the working loads to the sub-units. In case the system needs to deliver its full performance, this task mapping obviously needs to be optimized with respect to performance. However, typically most of the operating time only partial load is applied to a system, and thus, to its corresponding sub-units. In this case, it is possible to locate the workload in a way, such that weakened resources are less strained, reducing the advancement of degradation and therefore sustaining the system lifetime. We call this approach Reliability Aware Task Mapping (RATeMAP).

A load balancing neglecting the temperature dependency of aging might cause hot spots that result in a disproportional accelerated aging of the affected resources. Therefore, RATeMAP considers the current temperature distribution. This way, a ranking of resources is determined in convenient intervals considering both the measured remaining resilience and the temperature distribution. The time constants of temperature and aging vary widely. The aging measurement itself can be performed infrequently, since aging is a slow, gradual process proceeding over years. In contrast, temperature changes much faster. Consequently, between two aging measurements there are plenty of temperature measurements (e.g. every hundred milliseconds in our implementation). Their results are summed up with the latest aging monitoring results to a score determining a new resource ranking. Temperature is incorporated to this score as Taylor series approximation of Equation 2 to fit its nonlinear influence.

In contrast to HCI degradation, NBTI does not directly depend on switching activity. Consequently, a simple reduction of the workload is insufficient to prevent a weakened structure from further NBTI degradation. Hence, NBTI protection can be achieved by applying an appropriate, RATeMAP-controlled power gating to currently unused resources. While power gating is commonly applied to reduce energy consumption, we extend its use by aging management. Thus, a powerful aging prevention can be achieved on low overhead mitigating both HCI and NBTI.

VI. EVALUATION AND RESULTS

We successfully realized an FPGA-based prototype of RATeMAP and our novel aging measurement. To prove feasibility and effectiveness of our approach we chose a scenario related to digital signal processing. In detail, a management unit maps a workload to four identical Arithmetic Logic Units (ALUs). Whenever the varying amount of workload does not require full system performance, the workload is allocated

according to the RATEMAP resource ranking, that relies on the latest of the infrequently determined aging states of each resource and their current temperature. The latter is frequently measured utilizing one ring oscillator per ALU. The ring oscillators are calibrated according to a dedicated temperature sensor during the free-of-workload maintenance states. Even though we implemented RATEMAP to allocate workloads to ALUs, it is independent from a particular resource or the system type. Hence, it is suitable to manage various kinds of homogeneous resources (*e.g.* task mapping in symmetric multiprocessor systems or networks-on-chip). Furthermore, RATEMAP does not depend on a specific technology. An implemented solution can be migrated to another technology by just adapting selected specific constants concerning the ring oscillator evaluation and the voltage-delay dependency.

We chose the Xilinx ML605 evaluation board as platform for our tests. It utilizes the Texas Instruments PTD08A020W switching regulators for the power supply of the FPGA. These switching regulators can be controlled and monitored via PM-Bus, allowing us to manipulate the core voltage as necessary for our aging monitoring. The Board features a Xilinx Virtex-6 FPGA fabricated in a 40 nm process. The whole system utilizes 54 % of the FPGA's lookup tables. The Microblaze soft core processor serving as management unit and the ring oscillators together occupy only 2.3 % of the utilized lookup tables. Thus, the overhead can be stated negligible. For technical reasons FPGAs cannot operate at arbitrary lowered supply voltages, since the utilized SRAM-based logic requires a certain minimum voltage level to preserve the programmed hardware setup. The Virtex-6 FPGA has a nominal voltage of 1.0V and a recommended minimum operating voltage of 0.75V. However, it satisfies the requirements of our measurement procedure proposed in Section IV. Unfortunately, the application of power gating is prevented by those circumstances.

If no aging management is applied, the described system fails as soon as the first component fails. Hence, in terms of reliability it can be considered as serial system with N components. The reliability R_s of such a system is given by the product of the reliability of its components R_i [9]:

$$R_s(t) = \prod_{i=1}^N R_i(t) \quad (3)$$

To fit the typical bathtub curve, each single component's reliability is modeled by a Weibull distribution. The last section of the bathtub curve is characterized by increasing failure rates due to wearouts. It is the most relevant one in terms of aging. Accordingly, the distribution's shape parameter is chosen as $\beta = 3.5$. The scale parameter λ is arbitrary due to the normalized time scale. The corresponding component's reliability is defined as [9]:

$$R_i(t) = e^{-\lambda t^\beta} \quad (4)$$

RATEMAP equalizes the lifetime of the single components, as the dedicated task mapping slows down the aging of the weakest component. As the system lifetime is determined by the weakest component's lifetime, the system lifetime $R_{RATEMAP}$ increases and asymptotically reaches the lifetime of a non-serial system:

$$R_{RATEMAP}(t) \simeq e^{-\lambda t^\beta} \quad (5)$$




Fig. 3. Reliability of systems composed of multiple resources.

Note that this asymptotic behavior is independent of the number of resources. This results in a considerably higher reliability. For example, at the time an unmanaged 4-component system's reliability reaches 50%, the managed system still offers 84% reliability. The resulting reliability functions of different systems are illustrated in Figure 3.

VII. CONCLUSIONS

Transistor degradation in emerging CMOS technologies demands new measures at system run time. In this paper we present a new method to measure the extent of transistor degradation, in order to allow for a proactive, aging-aware system management at run time. The utilization of existing DVFS resources enables operation at minimum overhead. Furthermore, we propose an appropriate aging-aware task mapping as a proactive measure to cope with emerging aging issues. This mechanism provides the possibility to either extend the system lifetime or to choose smaller static timing margins resulting in an optimized utilization of the technology's capabilities. Finally, the feasibility and efficiency of the approach is proved by implementation of an FPGA-based prototype.

VIII. ACKNOWLEDGEMENTS

We gratefully thank CAPES, CNPq, and FAPEMIG for financial support.

REFERENCES

- [1] JEDEC SOLID STATE TECHNOLOGY ASSOCIATION, *Failure Mechanisms and Models for Semiconductor Devices*, 2011.
- [2] N. M. Vichare and M. G. Pecht, "Prognostics and health management of electronics," *Components and Packaging Technologies, IEEE Transactions on*, vol. 29, no. 1, pp. 222–229, 2006.
- [3] A. Simevski, R. Kraemer, and M. Krstic, "Low-complexity integrated circuit aging monitor," *IEEE DDECS*, pp. 121–126, 2011.
- [4] M. Agarwal, B. C. Paul, M. Zhang, and S. Mitra, "Circuit failure prediction and its application to transistor aging," *IEEE VLSI Test Symposium*, 2007.
- [5] M. Agarwal, V. Balakrishnan, A. Bhuyan, K. Kim, B. C. Paul, W. Wang, B. Yang, Y. Cao, and S. Mitra, "Optimized circuit failure prediction for aging: Practicality and promise," *International Test Conference*, 2008.
- [6] J. Li and M. Seok, "Robust and In-Situ Self-Testing Technique for Monitoring Device Aging Effects in Pipeline Circuits," *DAC*, 2014.
- [7] Neil H. E. Weste; David Money Harris, *CMOS VLSI Design - A Circuits and Systems Perspective*, 4th ed. Boston: Addison-Wesley, 2011.
- [8] T. Wegner, "Modellierung und Steuerung der Temperaturverteilung Network-on-Chip-basierter Systeme," Ph.D. dissertation, University of Rostock, 2014.
- [9] I. Koren and C. Krishna, *Fault Tolerant Systems*. Amsterdam: Elsevier Science & Technology, 2007.