

Aging-Aware Voltage Scaling

Victor M. van Santen*, Hussam Amrouch*, Narendra Parihar†, Souvik Mahapatra† and Jörg Henkel*

*Karlsruhe Institute of Technology, Chair for Embedded Systems (CES), Karlsruhe, Germany

†IIT Bombay, Department of Electrical Engineering, Mumbai, India

{victor.santen; amrouch; henkel}@kit.edu, {np.electro; souvik}@ee.iitb.ac.in

Abstract—As feature sizes of transistors began to approach atomic levels, aging effects have become one of major concerns when it comes to reliability. Recently, aging effects have become a subject to voltage scaling as the latter entered the sub- μs regime. Hence, aging shifted from a sole long-term (as treated by state-of-the-art) to a short and long-term reliability challenge.

This paper interrelates both aging and voltage scaling to explore and quantify for the first time the short-term effects of aging.

We propose “aging-awareness” with respect to voltage scaling which is indispensable to sustain runtime reliability. Otherwise, transient errors, caused by the short-term effects of aging, may occur. Compared to state-of-the-art, our aging-aware voltage scaling optimizes for both short-term and long-term aging effects at marginal guardband overhead.

I. INTRODUCTION

On-chip systems in the current and upcoming technology nodes are thermally constrained [1] due to the continuing scaling that steadily increases on-chip power densities. As a matter of fact, voltage scaling techniques became inevitable in order to fulfill performance constraints while obeying temperature constraints [2]. While, on the one hand, increasing the supply voltage (V_{dd}) allows to boost the CPU performance [3] due to the higher operating frequency, decreasing V_{dd} , on the other hand, helps avoiding critical temperatures.

Ultra-fast voltage scaling: The joint fulfilling of both performance and thermal constraints necessitates to switch the voltage very frequently. However, each V_{dd} switch invokes a performance penalty due to the inoperative phases. This is unavoidable since the power supply would be unstable during switching due to charging/discharging the chip’s capacitances [4]. To increase the efficiency, manufactures started implementing *ultra-fast* voltage regulators where V_{dd} switching moved into the sub-micron regime like the Intel Haswell CPU which switches between voltage levels within less than $1\mu\text{s}$ [4], [5], reducing the performance penalty of voltage scaling.

Aging effects: In the nano-scale era, aging effects are at the forefront of reliability concerns [6] due to their momentous ability to cause hardware failures. During the operation of transistors (i.e. applying/ceasing electric fields) the Bias Temperature Instability (BTI) aging mechanism¹, leads to continuously breaking/annealing Si-H bonds at the Si-SiO_2 interface as well as capturing/emitting charges in the oxide vacancies inside the transistor’s $\text{SiO}_2/\text{high-}\kappa$ dielectric [8]. Overtime, generated defects manifest as a gradual shift in the threshold voltage of a transistor (V_{th}). Aged (i.e. slower) transistors degrade the reliability of on-chip systems as they become less resilient to timing violations manifesting in errors.

Guardband: To sustain reliability during the entire lifetime of an on-chip system, designers employ a guardband, i.e. a

¹We focus solely on BTI as it is responsible for the highest degradation compared to other aging mechanisms [7]. However, our work is applicable to any mechanism featuring recovery, like Hot Carrier Injection.

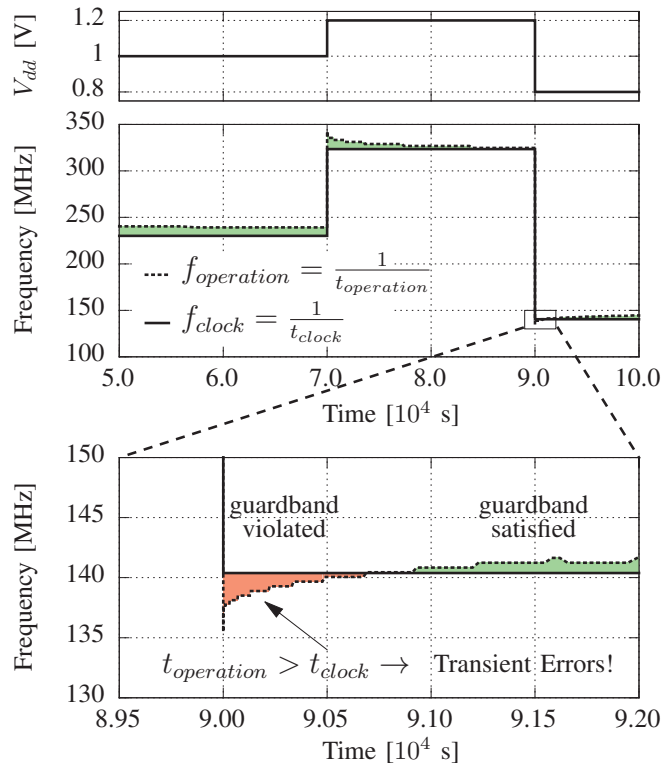


Fig. 1. Aging in conjunction with ultra-fast voltage scaling may lead to transient errors due to the temporary violation of guardband

slack time ($t_{guardband}$) that is added to the nominal delay of chip ($t_{nominal}$), to tolerate the slower operation due to aging.

$$f_{clock} = \frac{1}{t_{clock}}; t_{clock} = t_{nominal} + t_{guardband} \quad (1)$$

$$t_{operation} > t_{clock} \Rightarrow \text{Timing violations}$$

Aging in the scope of voltage scaling: In fact, aging is accelerated/decelerated based on the strength of electric fields and thus based on V_{dd} [8]. Hence, ΔV_{th} indeed follows the tendencies of V_{dd} controlled by the employed voltage scaling technique, i.e. higher $V_{dd} \rightarrow$ higher aging-induced ΔV_{th} , and vice versa. Importantly, switching V_{dd} in an ultra-fast manner opens the door for emerging **transient errors**, as the V_{dd} will be dropped much faster than the speed of aging recovery, as it will be demonstrated in Section II. In practice, such transient errors may appear immediately after switching from high to low V_{dd} level due to the temporary violation of the guardband. In Fig. 1, we show how $t_{operation}$ temporarily grows larger than t_{clock} after switching to a lower V_{dd} level. This is because of the high ΔV_{th} , originating from the previous high V_{dd} level along with the negligible recovery within a transition time of $<1\mu\text{s}$. Recent measurements in [5] through an on-




Fig. 2. Overview of the degradation and recovery of the BTI aging mechanism and its relations with voltage scaling. (a) Aging degradation is determined by the strength of V_{dd} , i.e. higher V_{dd} leads to higher ΔV_{th} . (b) Although the transistor is still on, switching the voltage to a lower level allows an intrinsic recovery to occur contrary to state-of-the-art that assumes recovery only occurs when at 0V. (c) Aging degradation follows the tendencies of voltage scaling. This demonstrates the necessity to *jointly* investigate aging and voltage scaling (as we propose) and not *separately* (as state-of-the-art does)

chip hardware monitor validated the theoretical prediction [9] of a sudden drop in the frequency (see Fig. 1) after the switch from high to low voltage level.

Therefore, aging effects should better be investigated jointly with voltage scaling. Otherwise, reliability may be unsustainable due to the hidden short-term effects of aging.

Our novel contributions within this paper are as follows:

- (1) We explore for the first time the *short-time* effects that aging in conjunction with voltage scaling has on reliability. This is unlike state-of-the-art which treats aging only as a long-term deleterious effect [7], [10].
- (2) To proactively avoid aging-induced transient errors, we propose a technique that adaptively tunes the guardband at runtime towards employing a small, yet sufficient one. Thereby, our technique still maintains the benefits of ultra-fast voltage switching and avoids the high performance loss that incurs from employing non-efficient guardbands.

II. AGING-INDUCED TRANSIENT ERRORS

As soon as a pMOS is turned on, the BTI mechanism occurs and generates defects that shift the V_{th} . The induced ΔV_{th} is determined by the strength of V_{dd} as Fig. 2(a) shows, where ΔV_{th} due to different V_{dd} levels is presented. However, when V_{dd} is switched to a lower level, a partial recovery of the generated defects starts to take place as Fig. 2(b) demonstrates. *State-of-the-art* (e.g., [7], [11]) considers that recovery solely occurs when the pMOS is turned off (i.e. $V_{gs} = 0V$).

However, recent measurements [5] as well as state-of-the-art physics-based BTI modeling [12] demonstrated that an intrinsic recovery occurs as soon as V_{dd} is switched to a lower level proving that recovering aging effects do not necessitate turning the pMOS off. To evaluate that, we employ the state-of-the-art Transient Trap Occupancy Model (TTOM) of BTI [12]. As seen in Fig 2(b), switching V_{dd} from 1.0V down to 0.9V and 0.8V reduces ΔV_{th} by 43% and 59%, respectively. This is in contrast to [13] which shows that voltage scaling has no impact on aging. This is due to employing models that are not capable to capture aging under voltage scaling. Note [13] like others also assumes only long-term effects of aging. Additionally, Fig. 2(c), illustrates how aging degradation follows the tendencies of voltage scaling. All in all, V_{dd} govern aging effects and therefore it is indispensable to investigate them *jointly* with voltage scaling.

In fact, increasing V_{th} results in decreasing the transistor drain current (I_D) which elongates its delay [14]. As a result,

aging increases the delay of the chip's critical path ($t_{operation}$) due to the delay increase of its individual transistor (t_{delay})².

$$t_{operation} = \sum_{i=1}^n t_{delay}(i) : i \in \text{critical path transistors} \quad (2)$$

$$t_{delay} \propto \frac{1}{I_D} \text{ with } I_D \propto (V_{dd} - V_{th} - \Delta V_{th})^2 \quad (3)$$

Susceptibility to aging degradation: Besides its role in governing aging, V_{dd} also determines the susceptibility to the induced degradation, i.e. the impact that ΔV_{th} has on increasing $t_{operation}$. In Fig. 3, we present how the same of aging degradation ($\Delta V_{th} = 10\text{mV}$) leads to a stronger shift in $t_{operation}$ at lower V_{dd} levels. This is consistent with what it can be derived from Eq. 3 where the impact of ΔV_{th} on the t_{delay} magnifies when V_{dd} becomes smaller.

This hints to our key idea of revealing the transient errors induced by aging in conjunction with voltage scaling.

Transient Errors: In state-of-the-art, aging is treated as a long-term problem where degrading the reliability of on-chip systems is in the order of months or even years. This is because aging gradually shifts V_{th} . However, employing ultra-fast voltage scaling changes the situation.

While degradation/recovery of aging still occurs gradually, the impact of aging on reliability becomes sudden in the presence of ultra-fast voltage scaling due to the negligible recovery that is feasible within such tiny transition times (i.e. $<1\mu\text{s}$). Therefore, the high ΔV_{th} , that was induced at the previous high V_{dd} level, will be carried to next low V_{dd} level where a higher susceptibility to aging degradation exists. Such a conjunction between the high aging degradation and the high aging susceptibility may lead to a temporary violation of the employed guardband (i.e. $t_{operation} > t_{clock}$) and thus to executing operations at that time results in transient errors (see Fig. 1). This explains the relevance of aging short-term effects. Despite some works (e.g., [10]) study aging under different V_{dd} levels, such a conjunction between the high aging degradation and the high aging susceptibility was neglected.

III. GUARDBANDS TO SUSTAIN RELIABILITY

Designing the required guardband that sustains reliability (i.e. protects on-chip systems from errors induced by the

²As aging may change which path is critical, works like [15] can be employed to determine the set of potentially critical paths after aging. For simplicity, our method is presented with respect to a single critical path




Fig. 3. SPICE simulations of a ring oscillator with aging modeling from [12] demonstrate that the susceptibility to aging increases as V_{dd} scales down

slower operation of aged transistors) may be either *static* at design-time or *dynamic* at runtime. In both cases, the guardband may be either *optimistically* or *pessimistically* designed.

Optimistic Static Guardband: The designer estimates the aging-induced ΔV_{th} under the worst-case scenario which comes from constantly applying the highest V_{dd} during the entire lifetime (e.g., 10 years). Then, the guardband is designed through calculating the increase in $t_{operation}$ due to the estimated ΔV_{th} , i.e. $t_{guardband} = \Delta t_{operation}$ at the highest ΔV_{th} and the highest V_{dd} . Importantly, such a guardband will be *optimistic* because it does not take into account that V_{dd} may be switched to a lower level causing a conjunction between the high aging degradation (induced at the previous high V_{dd} level) and the high susceptibility (exists at the next low V_{dd} level). In the past, recovery had sufficient time to compensate the higher susceptibility at lower V_{dd} by reducing ΔV_{th} during the voltage switch, but with the introduction of ultra-fast voltage scaling, the OSG approach became unreliable. It may lead to transient errors because such an *optimistic* guardband may temporarily be violated at runtime (see Fig. 1).

Pessimistic Static Guardband: To overcome transient errors, the designer may consider the worst-case scenario in both aging degradation and aging susceptibility. In such a case, the guardband is designed through calculating the $\Delta t_{operation}$ based on the worst-case ΔV_{th} , which is caused by constantly applying the highest V_{dd} along with the worst-case aging susceptibility, which comes from switching to the lowest V_{dd} , i.e. $t_{guardband} = \Delta t_{operation}$ at the highest ΔV_{th} and the lowest V_{dd} . Indeed, the designed guardband is able to overcome all transient errors unlike the previous case. However, such a guardband is *pessimistic* (i.e. larger than what actually be needed at runtime) as it considers the worst-case conjunction, where V_{dd} is always scaled from the highest to the lowest V_{dd} level. Therefore, a considerable performance loss may incur due to the unnecessarily-low operating frequency.

Dynamic Guardbands: To avoid the high performance loss inherent to pessimistic static guardbands, the guardband may *dynamically* be adapted at runtime based on a hardware monitor that provides delay measurements (e.g., [16]). In practice, the on-chip system periodically (at every t_{update}) checks the monitor and adapts the $t_{guardband}$ according to the current delay increase. It is noteworthy that enabling the monitor to get the measurement imposes aging stress on its transistors and hence frequent access leads to rapidly aging the monitor. Therefore, dynamically adapting the guardband based on periodically reading the monitor is a double-edged sword. On the one hand, infrequent reading through employing an *optimistic* t_{update} (e.g., in the order of seconds) avoids aging the monitor. However, it leads to overcoming only the long-term

effects of aging since short-term effects, originating from the ultra-fast voltage scaling, occur in a significantly shorter period of time (i.e. $<1\mu s$). On the other hand, frequently reading the monitor through employing a *pessimistic* t_{update} , which must be smaller than the switching time of V_{dd} (i.e. $<1\mu s$), overcomes both short and long-term effects of aging but it concurrently imposes a severe aging stress on the monitor itself and thus it rapidly ages resulting in a high degree of uncertainty with respect to monitor readings.

IV. OUR PROPOSED A-GEAR TECHNIQUE

To counteract short *and* long-term effects of aging with minimum performance loss, we propose a novel technique that employs an Adaptive Guardband for short- and long-term aging Effects AwaReness (A-GEAR). It is based on an offline (i.e. design-time) analysis, where we investigate the impact that different aging degradations at different V_{dd} levels – which are available within the chip [1] – have on the critical path delay. The analysis results are then used to build an *interpretation table* which interprets the current state of aging degradation to the corresponding guardband that the on-chip system actually needs. This table is employed at runtime to allow an efficient *adaptation* of the guardband (i.e. selecting small, yet sufficient guardbands) based on the existing operation conditions, i.e. the current degradation (ΔV_{th}), the previous and next voltage levels (V_{dd} and V'_{dd}).

A. Aging Effects Investigation

To obtain the current state of aging degradation, we assume the availability of a hardware delay monitor that measures the delay increase at runtime. A wide range of implementations of such monitors has been proposed which, in practice, measure the delay through a ring oscillator and then compare the result with the original/reference delay to capture the delay increase (i.e. $\Delta t_{monitor}$) at any point of time. For instance, state-of-the-art monitor [5] is able to provide its measurements within $<1\mu s$ and for different voltages ($V_{dd} \in [0.8 - 1.4V]$). Authors showed, that their monitor can be implemented through a very simple circuit and hence adds just minor costs/overheads [5].

Once the delay increase ($\Delta t_{monitor}$) is known, the current aging degradation state ΔV_{th} can be estimated³ To achieve that, the hardware monitor circuit (i.e. the ring oscillator) is modeled through a SPICE netlist along with the BSIM4 transistor model [14] on 22nm PTM technology [18]. Then, we employ state-of-the-art physics-based aging modeling [12] that models the impact of BTI on V_{th} and, more importantly, is able to take the voltage dynamics into account. This enables us to accurately consider the joint effect of voltage scaling and aging degradation on reliability. Table I shows an example of such an analysis when an 11-stage ring oscillator is examined.

To the best of our knowledge, the employed aging modeling within this work is the exclusive one that is able to consider the intrinsic aging recovery due to scaling the voltage down (see Fig. 2(b)) from a physical perspective. While the *empirical* aging model [9] is able to consider voltage fluctuations, the aging modeling [12] that we employ is based on modeling the

³To consider the intrinsic variability of BTI [17], the distribution $\Delta V_{th}(\mu, \sigma)$ could be calculated [17] and worst case of ΔV_{th} (e.g. 6σ) selected as the upper bound for degradation. However, our recent model [12] only models the mean $\Delta V_{th}(\mu)$.

underlying *physical processes* behind aging and hence we can model a wide range of voltages, temperatures, etc. with a high degree of certainty. This is indispensable to achieve our goal of exploring the short and long-term effects of aging where we need to accurately investigate aging degradation within very fine-grained time steps (i.e. microseconds) and under highly-fluctuating voltage conditions (i.e. ultra-fast voltage scaling).

TABLE I
EXAMPLE OF THE RESULTING ΔV_{th} DUE TO DIFFERENT ΔT AND V_{dd}

$\Delta t_{monitor} \downarrow V_{dd} \rightarrow$	0.8V	0.9V	1.0V
5%	7mV	9mV	12mV
10%	14mV	17mV	21mV
15%	19mV	23mV	29mV
20%	24mV	30mV	38mV

B. Guardband Estimation

Once, the ΔV_{th} is known, we then estimate the required guardband at the each V_{dd} . This can be achieved through simulating the impact of that particular ΔV_{th} on the chip's critical path delay in SPICE according to different V_{dd} . Table II shows an example of resulting guardband for different conjunctions between aging degradation (ΔV_{th}) and next voltage level (V'_{dd}). It is noteworthy that, the same ΔV_{th} results in different delay increases in the monitor itself and in the critical path (i.e. $\Delta t_{monitor} \neq \Delta t_{operation}$ at the same ΔV_{th}). This is because different circuits may have varied transistors sizes (i.e. different I_D) and therefore the same ΔV_{th} results in different delay increases. For instance, the guardband, at $\Delta V_{th} = 20\text{mV}$ and $V'_{dd} = 0.8\text{V}$, results in 5.3 difference corresponding to a 27% underestimation if we solely rely on the monitor measurement without interpreting it (see Table II). *This illustrates why we cannot directly rely on the monitor to select our guardband unless we interpret its measurement to the corresponding delay increase in the critical path of chip.*

As explained in Section I and motivated in Figs. (1, 3), circuits become more susceptible as V_{dd} scales down. Therefore, guardbands increase if V_{dd} is switched to a lower level and/or ΔV_{th} increases, as it can also be observed in Table II.

TABLE II
THE SAME AGING-INDUCED ΔV_{th} RESULTS IN A DIFFERENT DELAY INCREASE IN THE CRITICAL PATH COMPARED TO THE MONITOR ITSELF

V'_{dd}	0.8V		1.0V	
$\Delta V_{th} \downarrow$	$\Delta t_{monitor}$	$\Delta t_{operation}$	$\Delta t_{monitor}$	$\Delta t_{operation}$
5mV	3.5%	4.5%	2.6%	3.2%
10mV	7.2%	9.1%	4.6%	6.4%
15mV	10.8%	14.0%	6.9%	9.6%
20mV	14.2%	19.5%	9.2%	12.7%

C. Runtime Adaptation to select Guardbands

Based on the aforementioned design-time investigations presented in Tables (I, II), an interpretation table can be extracted to be employed at runtime. Such a table contains the required guardbands that are actually needed to tolerate the delay increase in the critical path according to different operating conditions. Such a table is a two-dimensional ($n \times m$) array, where n is the total number of $\Delta t_{monitor}$ steps and m is the total number of V_{dd} levels. In practice, for each $\Delta t_{monitor}$ step, we calculate the corresponding ΔV_{th} within the hardware monitor transistors. Then, we apply that calculated ΔV_{th} to

TABLE III
EXAMPLE OF THE HARDWARE TABLE, INTERPRETING THE HARDWARE MONITOR DELAY TO A GUARDBAND OF CHIP AT DIFFERENT VOLTAGES.

$\Delta t_{monitor} \downarrow V_{dd} \rightarrow$	0.8V	0.9V	1.0V
5%	6.3%	6.4%	7.3%
10%	12.2%	12.5%	13.6%
15%	18.4%	18.5%	19.0%
20%	24.1%	24.1%	26.3%

the critical path of our on-chip system to estimate the delay increase ($\Delta t_{operation}$). Table III shows an example of the resulting table that will be implemented within the chip to be employed by our runtime algorithm (see Algorithm 1 and further details are in the next section) that present the hardware implementation of our proposed A-GEAR technique. A hardware monitor may have fine-grained $\Delta t_{monitor}$ steps, leading to a large n . To reduce n , we store only the entries, which lead to a different guardbands. As guardbands correspond to the small set of frequency levels within a CPU and hence are coarse-grained in comparison. This allows the feasibility of implementing the table in hardware.

Overcoming Short-term Effects of Aging: Whenever a voltage switch is triggered the responsible control circuit reads, from the hardware monitor, the current delay increase ($\Delta t_{monitor}$) at the requested new voltage level (V'_{dd}). Then, it obtains from the our implemented look-up table the required guardband that sustains a reliable operation based on ($\Delta t_{monitor}, V'_{dd}$). To further optimize our technique, we additionally exploit the intrinsic recovery that is inherent to switching to a lower voltage level (see Fig. II(b) and Section II). As recovery is a exponential function [8], it is worthwhile to adapt the guardband again after a short period of time to exploit the recovery ($t_{recovery}$) which, in turn, enables us to avoid applying a non-efficient guardband (i.e. larger than what the system actually needs).

Overcoming Long-term Effects of Aging: On the other hand, voltage scaling may not be triggered for a prolonged interval of time and hence the employed narrow guardband may become insufficient due to the gradual degradation of aging (i.e. the well-known long-term effect of aging). Therefore, to also counteract long-term effects of aging while employing narrow guardbands we regularly update the guardband, at every t_{update} similar to [16], based on ($\Delta t_{monitor}, V_{dd}$) after rechecking the hardware monitor measurement.

Distinction from existing techniques: Various adaptive guardband techniques have been proposed (e.g., [16]). However, our A-GEAR technique distinguishes itself from others through the following novelties:

- It considers the short *and* long-term effects of aging, instead of solely long-term effects, which prevents transient errors.
- It interprets the aging monitor measurement to the corresponding guardband that the chip's critical path actually needs, instead of directly applying the measurement itself as a guardband, which prevents wrong guardbands.
- It considers the intrinsic recovery of aging in the on-state of the transistor, recently demonstrated [5], [12], which provides efficient guardbands.
- It considers, while adapting the guardband, the impact of voltage scaling on the susceptibility to aging, which allows a correct estimation of guardbands.

Algorithm 1 Algorithm of our hardware A-GEAR technique

Require: Current, new voltages (V_{dd} , V'_{dd}), Timer, Look-up Table

```

1: for every trigger  $\in$  (voltage switch, timer expired) do
2:   if voltage switch then
3:     Read  $\Delta t_{monitor}$  at  $V'_{dd}$  ▷ monitor
4:     Get  $t_{guardband}$  at ( $\Delta t_{monitor}$ ,  $V'_{dd}$ ) ▷ look-up
5:     Set frequency  $f_{clock}$ 
6:     Switch to  $V'_{dd}$ 
7:     if  $V'_{dd} < V_{dd}$  then ▷ intrinsic recovery
8:       Set timer to  $t_{recovery}$ 
9:       Wait until timer expired
10:      Read  $\Delta t_{monitor}$  ▷ monitor
11:      Get  $t_{guardband}$  at ( $\Delta t_{monitor}$ ,  $V_{dd}$ ) ▷ look-up
12:      Set frequency ( $f_{clock}$ )
13:    end if
14:  else if timer expired then
15:    Read  $\Delta t_{monitor}$  ▷ monitor
16:    Get  $t_{guardband}$  at ( $\Delta t_{monitor}$ ,  $V_{dd}$ ) ▷ look-up
17:    Set frequency ( $f_{clock}$ )
18:  end if
19:  Set timer to  $t_{update}$ 
20: end for

```

V. EXPERIMENTAL SETUP

To evaluate our A-GEAR technique and to quantify the short-term effects of aging, we implemented the following:

Thermal Estimation: First, the gem5 simulator [19] extracts the activities of the running application on top of the single-core Alpha CPU⁴. Then, the McPAT simulator [20] provides for a 22nm technology, the V_{dd} levels of the Alpha CPU along with the corresponding maximum frequency and static/dynamic power consumption of the CPU's components at each V_{dd} . Afterwards, the Hotspot thermal simulator [21] estimates the temperature of the CPU's components based on the extracted activity and power. In our experiments, we employed diverse applications from the PARSEC [22] and SPEC2006 [23] benchmark suites exhibiting diverse activities/powers and hence thermal behaviors. In addition, we executed them on top of the Linux OS to consider a more realistic scenario than bare-metal execution.

Dynamic Thermal Management (DTM): We implemented the state-of-the-art DTM technique, namely "Intel Turbo Boost 2.0", from the Intel Haswell CPU [3]. It works as follows [1]: it checks every 1ms whether the critical temperature (e.g., $T_{crit} = 80^\circ\text{C}$) is reached or not. If yes, it decreases frequency by one step (e.g., 133 or 200MHz) and it scales the V_{dd} down to the corresponding V_{dd} level of new frequency. If T_{crit} is not yet reached, the frequency is, instead, increased by one step and the V_{dd} is scaled up to the corresponding V_{dd} level.

Aging Estimation: As explained in Section IV, we estimate aging effects with state-of-the-art BTI aging modeling [12]. Based on the voltage trace, which is resulted from the thermal behavior of the running application and the employed DTM technique, we estimate the corresponding aging degradation trace. The latter enables us to quantify the short and long-term effects of aging jointly with voltage scaling towards capturing when the guardband is violated.

Evaluated Scenarios: For a fair comparison and a more general evaluation, we consider the following four scenarios: (1) *Base*: The unmodified (i.e. nominal) CPU which is not protected against aging (i.e. no guardband is employed).

⁴In many-core system, our A-GEAR needs to be implemented in each core individually to consider different V_{dd} levels per core




Fig. 4. The number of falling edges due to reducing V_{dd} one step (e.g. $0.99\text{V} \rightarrow 0.93\text{V}$) within the resulting voltage traces as on-chip system is susceptible to transient errors only there

- (2) *Optimistic Static Guardband (OSG)*: The CPU is protected against only the long-term effects of aging (see Section III).
- (3) *Pessimistic Static Guardband (PSG)*: The CPU is protected against the short and long-term aging effects (see Section III).
- (4) *A-GEAR*: The CPU is protected against short and long-term effects of aging through adapting the guardband at runtime based on our proposed technique described in Algorithm 1.

It is noteworthy that Base and OSG are unreliable designs as errors due to aging may occur. Whereas, PSG and A-GEAR are reliable designs as they prevent errors due to aging.

VI. EVALUATION, COMPARISON AND ADVANTAGES

Since transient errors due to the short-term effects of aging occur only when V_{dd} is switched to a lower level, we show in Fig. 4 the total number of falling edges after analyzing the resulting voltage trace of each application. The reason behind the variety in voltage traces is that the applications have different thermal behaviors and thus they differently trigger the DTM technique. As a result, different applications exhibit different rates of transient errors that are induced by the short-term effects of aging.

To quantify the latter, we demonstrate in Fig. 5 the total number of occurring transient errors per second in OSG (i.e. not counteracting short-term effects of aging). In such a case, the designed guardband is 17% which is the resulting aging degradation at the end of a 10 years lifetime when the highest V_{dd} (1.2V) is constantly applied. As shown in Fig. 5, designing a guardband – that is unaware of the short-term effects of aging – leads to unreliable behavior due to the high rate of transient errors (on average 94 errors/s). In practice, the $\Delta V_{th} = 44\text{mV}$ that a static guardband of 17% is able to tolerate becomes lower when V_{dd} is switched down and therefore the guardband may temporarily be violated, at the falling edges, resulting in transient errors.

Our A-GEAR technique *adaptively* selects the sufficient guardband that sustains a reliable operation. The distribution of the selected guardbands at runtime for different applications is presented in Fig. 6. As observed, the minority of time is spent within the large guardbands. This is because of the efficient selection of our guardbands due to the exploitation of intrinsic recovery. To evaluate the latter, we demonstrate in Fig. 7 the normalized execution time of each application. Compared to applying the OSG technique that protects the on-chip system against only long-term aging effects, our A-GEAR technique overcomes both short and long-term aging effects and it comes with merely 1% overhead on average.




Fig. 5. Error rate after 1 year operating at $V_{dd} = 1.2V$. Note, employing A-GEAR prevents all errors due to the employment of sufficient guardbands.

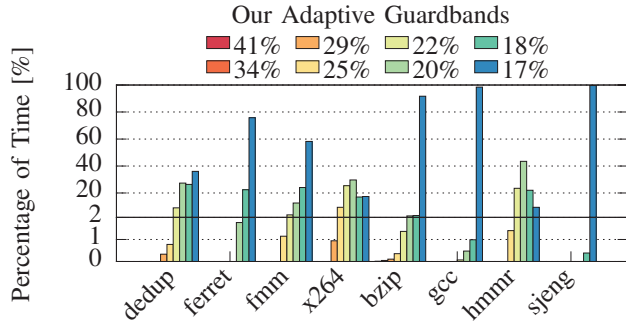


Fig. 6. Percentage of time spent at each guardband that is *adaptively* selected at runtime through our A-GEAR technique

Finally, compared to applying the PSG technique that, similar to ours, is able to overcome shot and long-term aging effects our A-GEAR reduces the overhead by 10% on average and up to 21%.

Monitor Degradation: As explained in Section III, each access to the hardware monitor imposes an aging stress on it. In fact, A-GEAR accesses the monitor only when the V_{dd} scaling is triggered in addition to the regular update at 1s. Our competitor here is dynamic guardband-based techniques (see Section III) when they aim to overcome short and long-term effects of aging. In such a technique, the monitor should very frequently be accessed (i.e. $t_{update} = 1\mu s$) to sustain reliability. Compared to such a case, we mitigate the monitor aging by 4.1x (i.e. we reduce the aging-induced ΔV_{th} in the transistors of the monitor's reference, after a lifetime of 10 years, from 14.03mV to 3.39mV).

VII. CONCLUSION

We demonstrated in this work how voltage scaling techniques, that are widely employed to fulfill performance and temperature constraints, may cause transient errors in conjunction with aging effects. This shows for the first time that designers must counteract the short-term effects of aging in addition to the well-known long-term effects. Our A-GEAR technique adapts the employed guardband at runtime to avoid the considerable performance loss that otherwise is associated with designing guardbands based on state-of-the-art techniques. With merely 1% overhead, it makes on-chip systems resilient to short-term *and* long-term aging effects.

ACKNOWLEDGMENT

This work was supported in parts by the German Research Foundation (DFG) as part of the priority program "Dependable Embedded Systems" (SPP 1500 - spp1500.itec.kit.edu). We thank Christian List for his valuable help in experiments.




Fig. 7. Performance loss evaluation

REFERENCES

- [1] J. Henkel, H. Khdr, S. Pagani, and M. Shafique, "New trends in dark silicon," in *DAC*, 2015.
- [2] J. Henkel, S. Pagani, H. Khdr, F. Kriebel, S. Rehman, and M. Shafique, "Towards Performance and Reliability-Efficient Computing in the Dark Silicon Era," in *DATE*, 2016.
- [3] J. Charles, P. Jassi, N. Ananth, A. Sadat, and A. Fedorova, "Evaluation of the Intel Core i7 Turbo Boost feature," in *IISWC*, 2009.
- [4] E. Burton, G. Schrom, F. Paillet, J. Douglas, W. J. Lambert, K. Radhakrishnan *et al.*, "FIVR Fully integrated voltage regulators on 4th generation Intel® Core SoCs," in *APEC*, 2014.
- [5] S. Satapathy, W. H. Choi, X. Wang, and C. Kim, "A revolving reference odometer circuit for BTI-induced frequency fluctuation measurements under fast DVFS transients," in *IRPS*, 2015.
- [6] J. Henkel, L. Bauer, J. Becker, O. Bringmann, U. Brinkschulte, S. Chakraborty *et al.*, "Design and architectures for dependable embedded systems," in *CODES+ISSS*, 2011.
- [7] H. Amrouch, V. van Santen, T. Ebi, V. Wenzel, and J. Henkel, "Towards interdependencies of aging mechanisms," in *ICCAD*, 2014.
- [8] S. Mahapatra, N. Goel, S. Desai, S. Gupta, B. Jose, S. Mukhopadhyay *et al.*, "A Comparative Study of Different Physics-Based NBTI Models," *T-ED*, 2013.
- [9] C. Zhou, X. Wang, W. Xu, Y. Zhu, V. Reddi, and C. Kim, "Estimation of instantaneous frequency fluctuation in a fast DVFS environment using an empirical BTI stress-relaxation model," in *IRPS*, 2014.
- [10] V. B. Kleeberger, M. Barke, C. Werner, D. Schmitt-Landsiedel, and U. Schlichtmann, "A compact model for NBTI degradation and recovery under use-profile variations and its application to aging analysis of digital integrated circuits," *Microelectronics Reliability*, 2014.
- [11] X. Li, J. Qin, and J. Bernstein, "Compact Modeling of MOSFET Wearout Mechanisms for Circuit-Reliability Simulation," *TDMR*, 2008.
- [12] N. Goel, T. Naphade, and S. Mahapatra, "Combined trap generation and transient trap occupancy model for time evolution of NBTI during DC multi-cycle and AC stress," in *IRPS*, 2015.
- [13] T.-B. Chan, J. Sartori, P. Gupta, and R. Kumar, "On the efficacy of nbtI mitigation techniques," in *DATE*, 2011.
- [14] Y. Chauhan, S. Venugopalan, M. Karim, S. Khandelwal, N. Paydavosi, P. Thakur *et al.*, "BSIM - Industry standard compact MOSFET models," in *ESSCIRC*, 2012.
- [15] J. Chen, S. Wang, and M. Tehranipoor, "Efficient Selection and Analysis of Critical-reliability Paths and Gates," in *GLSVLSI*, 2012.
- [16] C. R. Lefurgy, A. J. Drake, M. S. Floyd, M. S. Allen-Ware, B. Brock, J. A. Tierno *et al.*, "Active Management of Timing Guardband to Save Energy in POWER7," in *MICRO*, 2011.
- [17] A. Kerber and T. Nigam, "Challenges in the characterization and modeling of BTI induced variability in metal gate / High-k CMOS technologies," in *IRPS*, 2013.
- [18] W. Zhao and Y. Cao, "New Generation of Predictive Technology Model for Sub-45 nm Early Design Exploration," *T-ED*, 2006.
- [19] N. Binkert, B. Beckmann, G. Black, S. K. Reinhardt, A. Saidi, A. Basu *et al.*, "The Gem5 Simulator," *SIGARCH Comput. Archit. News*, 2011.
- [20] S. Li, J. H. Ahn, R. D. Strong, J. B. Brockman, D. M. Tullsen, and N. P. Jouppi, "The McPAT Framework for Multicore and Manycore Architectures: Simultaneously Modeling Power, Area, and Timing," *ACM Trans. Archit. Code Optim.*, 2013.
- [21] M. R. Stan, K. Skadron, M. Barcella, W. Huang, K. Sankaranarayanan, and S. Velusamy, "Hotspot: a dynamic compact thermal model at the processorarchitecture level," *Microelectronics Journal*, 2003.
- [22] C. Bienia, S. Kumar, J. P. Singh, and K. Li, "The PARSEC Benchmark Suite: Characterization and Architectural Implications," in *PACT*, 2008, pp. 72–81.
- [23] J. L. Henning, "SPEC CPU2006 Benchmark Descriptions," *SIGARCH Comput. Archit. News*, 2006.