When "things" Get Older: Exploring Circuit Aging in IoT Applications

Xinfei Guo, Vaibhav Verma, Patricia Gonzalez-Guerrero, Mircea R. Stan
ECE Department, University of Virginia, Charlottesville, VA, USA
E-mail: {xguo, vv8dn, lg4er, mircea}@virginia.edu

Abstract— The Internet of Things (IoT) brings a paradigm where humans and "things" are connected. Reliability of these devices becomes extremely critical. Circuit aging has become a limiting factor in technology scaling and a significant challenge in designing IoT systems for reliability-critical applications. As IoT becomes a general-purpose technology which starts to adapt to the advanced process nodes, it is necessary to understand how and on what level aging affects different categories of IoT applications. Since aging is highly dependent on operating conditions and switching activities, this paper classifies the IoT applications based on the aging-related metrics and studies aging using the foundry-provided FinFET aging models. We show that for many IoT applications, aging will indeed add to the already tight design margin. As the expected chip lifetime in IoT devices becomes much longer and the failure tolerant requirements of these applications become much more strict, we conclude that aging needs to be considered in the full design cycle and the IoT lifetime estimation needs to incorporate aging as an important factor. We also present application-specific solutions to mitigate circuit aging in IoT systems.

Keywords— IoT, aging, lifetime, sleep mode, recovery

I. Introduction

The Internet of Things (IoT) has been a powerful enabler to make technology more human centric and real time. IoT is a general-purpose technology by nature ranging from health care, transportation to agriculture and almost all aspects of life [1]. These applications impose common requirements for IoT devices like they should have small form factors in terms of physical dimensions and weights. Since most of these devices are battery-powered or batteryless, they require high energy efficiency and extreme low power consumptions. In addition to these characteristics, IoT devices need to withstand hostile environments such as increased and highly variable temperatures and voltage noise [1]. More importantly, IoT nodes are required to operate reliably for a long lifetime (e.g. decades), which translates to reliability challenges, especially device degradation-induced circuit aging[1]. As the circuit operates, on-chip elements such as transistors and metal wires age gradually when under use, and this can lead to potential permanent failures. Many of the IoT applications (such as automotive or implantable medical devices) require almost zero error during the whole lifetime. Harsh environment such as high temperature accelerate aging. Although the IoT industry has not fully migrated to deeply scaled technologies because of cost and leakage issues but recent advances in technology such as FinFET provide a compelling combination of performance, power, highest integration and ease of design for low-power IoT products [2]. This has pushed the IoT industry to start adapting to these newer technology nodes [3]. As number of on-chip elements scales up, more transistors are susceptible to aging and this leads to the increase of the system failure rate. These advanced technology nodes impose more aging issues than previous generations due to self-heating, reduced oxide thickness, narrower metal and increased current density [4].

[1]This paper focuses on circuit aging. Battery aging and socket (and holder solder) aging are out of the scope of discussion.

Since IoT is a wide concept and circuit aging is a threat to IoT lifetime, it is necessary to understand how current and future IoT systems are impacted by aging. This paper answers this question by conducting extensive circuit-level simulations with foundry-calibrated aging models in advanced FinFET node. Since aging is highly dependent on application behaviors that define the operating voltage, temperature and active time, we perform a survey of existing IoT applications and classify them based on aging-related metrics. By studying aging behaviors in each category, we show that aging can impose immense degradations in performance and design margin for some IoT applications. Our results prove that to meet IoT lifetime requirement, both battery lifetime (energy-efficiency perspective) and chip lifetime (circuit aging perspective) need to be considered together in the full design cycle. As flat guard-band approach could introduce unacceptable energy overheads for IoT systems, several dynamic solutions are also discussed in the paper.

II. Background and Related Work

A. Transistor Aging Basics

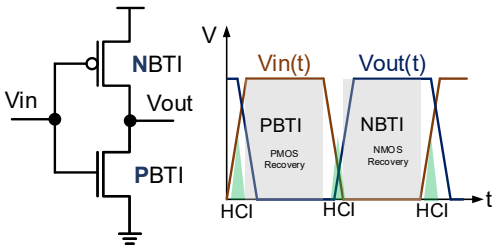

Fig. 1. Transistor Aging: HCI occurs mainly during switching; PBTI happens when NMOS is under stress; NBTI happens when PMOS is under stress. BTI aging partially recovers during OFF states.

The primary cause of aging is the electrical stress across the on chip components such as transistor, dielectrics and interconnects. Different from time zero variation, aging accumulates over time and depends on operating conditions. In general, there are mainly three dominant causes of aging in semiconductor devices. Bias temperature instability (BTI) and Hot carrier injection (HCI) lead to transistor degradations [5] and Electromigration (EM) leads to metal wire resistance increase. In this paper, we will focus on transistor aging, which has been demonstrated as one of the most disruptive reliability threats for modern chips

TABLE I
SUMMARY OF IOT APPLICATIONS SPECIFICATIONS (AGING-RELATED METRICS)
△: Specification for current commercial products (in market); ⋆: Specification for future product (under research)

Applications	Temperature			Core Voltage			Lifetime Requirement			Active Time		
	L	M	H	L	M	H	L	M	H	L	M	H
1 - Implantable/Heathcare		△		⋆	⋆	△			△			△
2 - Consumer Electronics/Wearables	△	△		⋆	⋆	△	△				△	
3 - Automotive			△			△			△		△	
4 - Industrial Processes	△	△	△			△		△				△
5 - Public Transportation	△	△				△		△			△	
6 - Energy Management	△	△	△			△		△			△	
7 - Smart Homes/Buildings/Cities	△	△		⋆	⋆	△		△				△
8 - Retailing/Malls		△				△		△			△	
9 - Agriculture/Environmental	△	△				△	△				△	
10 - Wildlife/Nature Preservation	△	△				△			△		△	

- For temperature, L - Low (≤27℃), M - Medium (27℃ – 100℃), H - High (>100℃);
- For voltage, L - Low (Sub-threshold), M - Medium (Near-threshold), H - High (Nominal Voltage);
- For lifetime requirement, Low (≤3 years), M - Medium (3 years – 10 years), H - High (>10 years);
- Here, "active" means transistors are under stress. In many IoT applications, even when the systems are in "sleep" mode, many circuit blocks (such as accelerometers in a wristband) are still ON and under aging stress. L - Low (Active < 20% of the lifetime), M - Medium (Active 20% - 80% of the lifetime), H - High (Active > 80% of the lifetime).

[6]. Both BTI and HCI impact transistor parameters (e.g. threshold voltage V_{th} and carrier mobility μ) at a level that depends on the operating environment and usage of the circuit. As shown in Fig. 1, NBTI happens to PMOS and is caused by constant electric fields degrading the dielectric. PBTI affects NMOS in a similar way. The aging mechanism associated with HCI shares similarities to BTI i.e. the electrical field shifts the channel free carriers towards the channel dielectric interface, but HCI happens mainly on drain side and primarily occurs during switching. While BTI is partially reversible HCI is an irreversible effect. The parameter shift is also highly temperature dependent since temperature affects the interface trap generation.

B. Previous Work on Aging in IoT Domain

Transistor aging has been explored for a long time in aerospace and safety applications but it didn't gain interests in consumer devices until very recently. For example, in automotive or industrial IoT applications aging happens even when the system is inactive most of the lifetime due to the continuous constant stress across transistors. These devices need to function under all possible scenarios during their expected lifetime [7]. For example, some medical implants will require a reliable operation for more than 50 years [8].

Previously a lot of attention has been paid on battery or package aging for IoT applications [9] while a few studies looked into circuit aging in this domain. [10] introduced a method to obtain multi-threaded switching activity signatures for aging analysis in IoT applications but the focus was on the architectural level framework. Similarly, [11] proposed a solution that leverages the workload dependent reliability analysis for early product failure rate calculations for automotive applications. [8] provided a device level experimental study of BTI aging in ultra-low power applications. [12] proposed a unified model which captures the joint impact of RTN, BTI and PV within a probabilistic reliability estimation for NTV circuits. Most of the previous studies focused on circuit aging in a very specific application or framework. Also previous works used analytical aging model which can lead to inaccurate predictions. The contribution of this paper is that: 1) We study aging impact with foundry-calibrated model instead of predictive models. 2) We investigate circuit aging on a wide spectrum of IoT applications and provide a deep and realistic understanding of how aging affects each IoT category. 3) Several solutions for addressing aging in IoT are also discussed in this paper.

III. Aging as a Reliability Threat for IoT

A. IoT Application Domains

The hardware requirements of the IoT devices are determined by how and where they are deployed [1]. To develop quantitative understanding of these requirements for IoT nodes, we surveyed published SoCs and commercially available IoT products, ranging from agriculture/environmental sensors, automotive, industrial processes to medical implantables, smart cities and consumer electronics. The results are summarized in Table I and are discussed below.

We classify the existing IoT applications into ten groups mainly based on the usage and scale of users [1], [13]. We mainly study aging-related metrics, i.e. voltage, temperature, lifetime requirements and active time which is how long transistors are stressed. IoT chip in all categories are found to operate at super-threshold voltages during active phases of computation for speed purposes even in battery-operated systems. Most of commercial low-power IoT architectures achieve energy efficiency through heavily optimized deep sleep modes or minimization of the unnecessary on-chip components [1]. There is a lot of ongoing research on operating IoT systems completely in near/sub-threshold region to achieve major energy efficiency improvements, especially in applications such as medical devices, sensors and wearables where energy harvesters can be adapted (application 1, 2, 7 in the table). But this comes at the expense of performance and increased sensitivity with respect to variations. In this paper, we focus on IoT chips that operate at nominal voltage only. Application 1 and 2 represent personal IoT where implantable devices usually operate

continuously at human body temperature while consumer electronics such as wearables are exposed to environment. Implantables are always active and require a long lifetime (almost human being lifetime) while wearables have a relatively shorter life cycles (around 3 years) and are inactive most of the lifetime. Applications 3 and 4 are industrial IoT in which automotive sensors monitor the state of the vehicle and mostly reside inside engines. They operate at very high temperature and require a reliable operation throughout car's lifespan of more than 10 years. Similar monitoring strategies are used in industrial environment such as storage warehouse or product lines. IoT devices also enable ubiquitous sensing in city and home infrastructures (applications 5 - 8) and are installed both indoors and outdoors. Thus they experience room or environmental temperatures and have relatively strict lifetime requirement since frequent checking and repairs are not practical. Applications 9 and 10 represent environmental IoT applications. They have similar temperature requirement as city/home-scale IoT. The agriculture sensors usually last for one cycle of crops but other environmental IoTs need to last longer because they are distributed at a very large scale and many of them are not quite accessible physically once they are installed.

B. Simulation Results

In general, process, voltage and temperature (PVT) variations impose additional timing margins that stretch the clock cycle. Aging impacts circuit in a similar way. Aging-induced performance loss requires guardband for margining. In this section, we study impact of aging with the foundry-calibrated models in FinFET technology using Cadence reliability simulator RelXpert (integrated in Virtuoso ADE). Both BTI (including recovery) and HCI mechanisms are included in this model.

B.1 Single Transistor Aging under Different Conditions

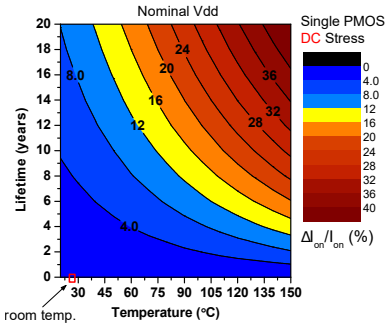

Fig. 2. Single transistor ON current degradation due to aging under **DC nominal** voltage stress.

Transistor aging causes parameter shift, such as increased threshold voltage V_{th} and reduced mobility μ and this leads to reduced current. Two sets of simulations (DC and AC stress) are run at different temperature and lifetime conditions for single transistor under nominal voltage. The ON current degradation (%) is plotted. Fig. 2 shows the DC stress case, in which the PMOS transistor ages continuously during the lifetime without any recovery giving the worst case estimation of NBTI aging. While Fig. 3

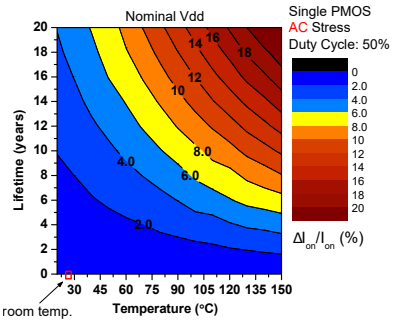

Fig. 3. Single transistor ON current degradation due to aging under **AC nominal** voltage stress with 50% duty cycle. Degradation is about half of the **DC** stress case due to recovery.

shows the results where the transistor is under AC stress with 50% duty cycle allowing BTI recovery period following stress. This includes both BTI and HCI degradations and provides an average aging estimation. As a reference, Monte Carlo simulations show that σ/μ for ON current is around 7%, which indicates the design margin with respect to process variation. The aging-induced degradations are comparable to this, and in high temperature cases, it can be more than 10%. An observation from Fig. 2 and 3 is that AC stress-induced degradation is almost half of the DC case and indicates that the degradation under the same temperature and lifetime condition is almost linearly proportional to the stress time. This assumption will be used in the following sections for estimating how the active time of each application affects aging.

B.2 Aging-induced Timing Failures in IoT Scenarios

Single transistor aging causes the degradation of ON current which will further impose timing error at the circuit level and failures at the system level. Shown in Fig. 4 is a typical data path from the output (Q) of the launch flop to the data input (D) of the capture flop. Aging slows down each unit ($t_{datapath}$ and t_{setup} increase) and this could cause setup time violations. This effect becomes more significant in data paths which have large logic depth. Extra timing margins need to be added to meet the setup timing requirement. To quantitatively study this margin under different operating conditions, we simulate a similar data path consisting of a combination of inverters and buffers. The margin is found by increasing the clock period until the launched data is correctly captured. Fig. 5 plotted the necessary margin vs. temperature and active time (how long the transistor is stressed) at nominal voltage. The result has been normalized to the necessary aging-induced timing margin at 27℃ with an active time of 2 years. This baseline margin is also equal to margin for temperature variation from 27℃ to 110℃ at time zero.

Based on the simulated results in Fig.5, we map the operating conditions of different IoT application listed in I and list the estimated aging margin for each category in Fig.6. As shown in the table, even in the same application category, the IoT devices operate at different temperatures and active time. The error bar in the figure provides an estimated range of the necessary margin which also indicates the estimated aging levels for each application. The

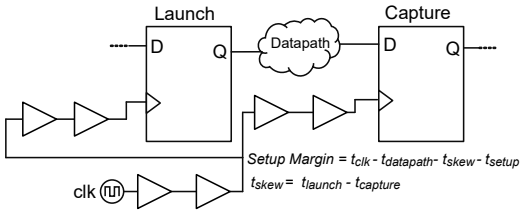

Fig. 4. Simulation setup (an example of datapath): Aging can lead to timing failures such as setup violation by slowing down the datapath. Designers should take extra margins based on aging impact.

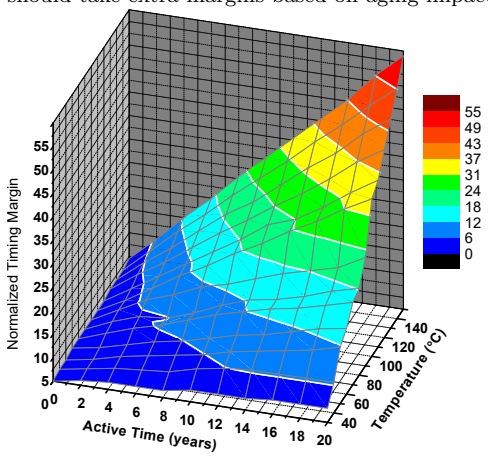

Fig. 5. Normalized timing margin vs. Temperature and Active Time: Margin shown on Y-axis is normalized to the required aging margin for datapath (shown in Fig. 4) for 2 years at room temp (27°C).

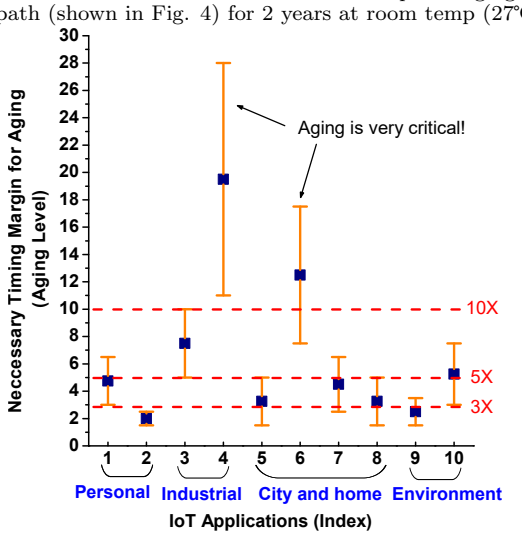

Fig. 6. Estimated aging margin for different IoT applications: X-axis corresponds to IoT application index in Table I, Y-axis shows the normalized design margin and the error bars show design margin range within each category.

top two aging-critical applications are 4 and 6 which correspond to industrial processes IoT and energy management IoT where high temperature and long active time are expected. These applications impose a more than 10× design margin compared to the baseline margin. In the second tier, automotive IoTs (application 3) suffer huge aging due to high operating temperature. Most of the city scale and environmental IoTs (applications 5 - 10) operate under environment temperature but have a relatively long active time and hence they lie in the third tier (3× to 5×). Similarly, implantable devices operating at body temperature need to operate reliably for a long life span, so they are also on the third most critical aging level. As consumer electronics or wearables (Application 2) are usually updated within timescale of 2 years, so they are the least aging-critical, but even so, they need to be margined properly for aging to guarantee the reliable operations spanning their lifetime.

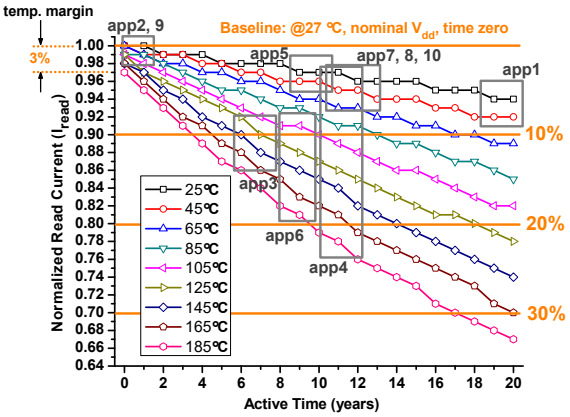

Fig. 7. 6T SRAM read current degradation with aging for different temperatures (nominal voltage).

B.3 Aging Impact on IoT SRAMs

SRAMs act as an external cache for many of today's IoT applications. They usually occupy the largest chunk of SoC and may interact with multiple cores. Hence it becomes imperative to study the impact of aging on SRAMs because the access-time and drive strength degradation may lead to timing failures across the chip. The access-time is directly proportional to the SRAM read current, I_{read}. Fig. 7 shows the I_{read} degradation of a 6T SRAM across different temperatures for different active time. The I_{read} values have been normalized to time-zero I_{read} at 27°C. It shows that many applications will incur more than 5% degradation during their lifetime while some critical applications such as Industrial IoT facing more than 20% degradation in read current. Such a huge loss in SRAM performance due to aging will potentially lead to fatal timing errors and hence should be taken care of during design process. The design process should also aim to appropriately assign timing margin for aging based on target applications.

IV. IoT Lifetime: Battery vs. Chip Lifetime

As battery replacement is not an option both due to the large numbers and inaccessibility of nodes in many IoT applications, the foremost requirement is that they can't rely on constant battery change. Thus the most common ways of defining lifetime of a battery-powered IoT system is by battery lifetime which is the time a node will operate in its normal mode without replacing the battery [1]. It is given by

$$T_{lifetime|battery} \sim E_{battery}/P_{average} \quad (1)$$

As transistor aging could potentially lead to chip failure that might not be recoverable as discussed in the last section, the aging-induced chip lifetime should also be considered to determine the final IoT system lifetime, which is given by

$$T_{lifetime|Final} = min\{T_{lifetime|battery}, T_{lifetime|chip}\} \quad (2)$$

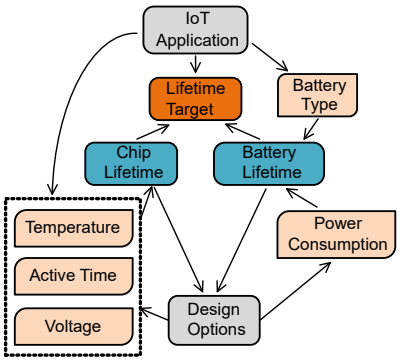

Fig. 8. IoT lifetime: Chip lifetime and Battery lifetime depend on different factors, but they can affect each other indirectly. Two lifetimes together determine the lifetime target of an IoT application.

$$T_{lifetime|chip} \sim F(voltage, temperature, active\ time) \quad (3)$$

Although battery lifetime and chip lifetime depend on different factors, they also impact each other indirectly. Illustrated in Fig. 8 is a suggested flow for closing the IoT lifetime loop as part of the design cycle. The system lifetime target is defined by the applications and specifications, which also constrain the battery size, weight and type. On the right branch, the battery lifetime is determined based on Equation 1. Design knobs such as voltage, power modes and active time can be tuned to achieve lower power consumption while fulfilling the performance requirement. Some of these design knobs are also limited by application itself e.g. implantable devices (Application 1) need to be active continuously and require fast response. On the left branch of the IoT lifetime loop, chip lifetime affected by aging is also highly dependent on knobs such as voltage, temperature and active time. To guarantee that chip lifetime is longer than or equal to the battery lifetime and the overall lifetime target; the design margin needs to be reserved. But as shown in the previous section, this margin can be very large and can translate into wasted energy in the early lifetime, which in turn shortens the battery lifetime. Alternatively, aging can be addressed in run time by adaptive solutions to reduce the design margin, while the additional sensors and circuitry will add to the power budget. To leverage these tradeoffs and meet the expected lifetime, careful design decisions considering both chip and battery lifetime are required. More details are discussed in the following section.

V. Potential Solutions for IoT Circuit Aging

Adding design margin is currently the most common way of addressing aging in the design flow. This is a static solution where all transistors are margin-degraded to a certain amount based on the operating conditions. The difference in performance of aged cell versus the original cell is computed and the ratio (aged/fresh) is used to derate cells in the design. But the large margin in some applications (shown in Section **3-B.2**) can be very conservative and introduce performance penalty in the early lifetime. An alternative solution would be to either recover aging or adapt to it dynamically so that the design margin requirement can be potentially relaxed. This section will briefly discuss several such solutions for IoT applications.

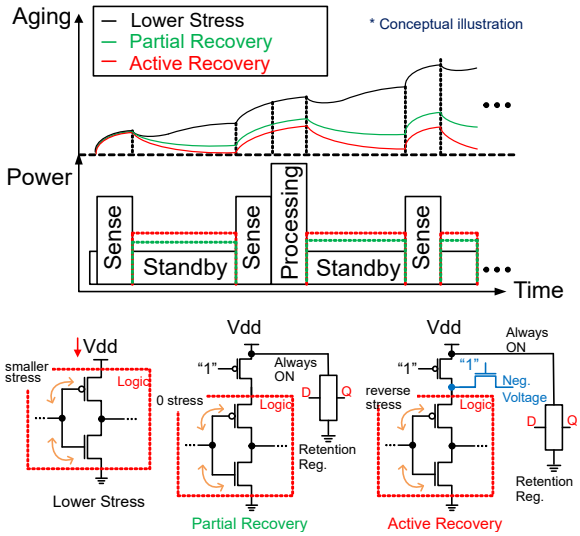

Fig. 9. Power and Aging profile of a typical IoT node: This figure is for illustration only, the height and width are conceptually marked. For aging profile, Y-axis "Aging" corresponds to aging-induced metric change such as ΔV_{th} or timing margin (reduced performance).

A. Lowering the Operating Voltage

Scaling the system operating voltage down into near/sub-threshold regime has been known to be a very effective way of reducing the energy per computation and extending the battery lifetime especially in health care and body sensor IoT applications [1]. Meanwhile, aging has an almost power law dependence on stress voltage [5]. Operation at lower voltages suppresses aging significantly. But the challenges are performance loss and increased sensitivity with respect to variations. The easy solution for performance degradations is to raise the operating voltage back again when necessary to meet the speed requirement, but this in turn will accelerate aging on the entire chip. One potential approach could be to have fine-grain voltage domains which can ensure that voltage boosting is kept small enough so that aging does not introduce large degradations and the impact is only constrained to certain sub-block. Fine-grain voltage domain can also maximize the opportunities to correct variations in paths that are critical due to variations. But this approach certainly leads to significant area overhead and design effort. These tradeoffs need to be leveraged based on the budgets that are defined by IoT applications.

B. IoT Circadian Rhythms: Recover during Standby

Another strategy for energy savings in IoT circuit design is to put the device in "standby" state as long as possible. This is feasible since the devices don't need to be active all the time in many applications shown in Table I. Some aging mechanisms such as BTI are recoverable when the transistor is OFF [5]. Hence these standby periods can be utilized for recovery. Fig. 9 illustrates the power and aging profile of a typical IoT node, where the sensing activity is periodic and triggered by some form of real-time events. One solution to save power and reduce aging is to operate the whole processing unit at the lowest voltage level while maintaining its state in retention mode. Although there is

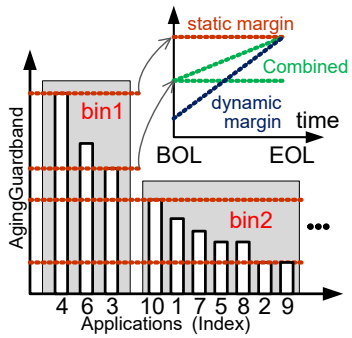

Fig. 10. Conceptual illustration of dynamic margins to enable one chip across multiple IoT applications (BOL - Beginning of lifetime, EOL - End of lifetime).

some recovery when switching from high voltage to lower voltage but aging will still accumulate since transistors are still under stress (at a lower level). An alternative solution would be to turn off certain blocks completely such as fixed function units by power gating and save the state in retention registers. This will result in partial recovery because BTI has a component which cannot be recovered just by switching off the transistors [14]. Recent work has demonstrated that recovery can be activated and accelerated by reversing the bias across transistors [15]. This insight leads to the third option where a negative supply voltage is provided to reverse gate bias the transistor and heal it faster. This approach also helps to recover BTI aging components which are not recoverable at zero bias. This approach enables maximum recovery during standby mode.

However the last two approaches come with additional power and area overheads for the power switches, logic retention, signal isolation and additional floorplanning constraints [1]. The third solution also introduces one more voltage source and domain. But for extremely aging-critical applications such as application 3, 4 and 6 shown in Fig. 6, these overheads are justified to prevent system failures.

C. Dynamic Margins across Multiple IoT Applications

To minimize design costs, circuit designers and chip vendors usually try to use one SoC design across multiple IoT applications. Even within one IoT application category, the operating conditions may change. These variations necessitate run-time compensation of aging such as techniques proposed in [16], where aging events (e.g. delay change) are tracked over operating periods. Once the failure flag (e.g. timing failure) is triggered, adaptive solutions such as dynamic voltage and frequency scaling (DVFS) or error correction are employed to compensate the degradations. But pure dynamic solution can be very limited and costly due to the limited tunability of metrics. A combination of static and dynamic margin methods is a more optimal approach. Fig. 10 illustrates a potential solution where targeted IoT applications are binned based on estimated aging levels. The static margin can be added based on the lowest level in the group, dynamic solutions are also applied for adapting to the worst-case operating conditions. Compared to the purely flat guard-band based approach, the combined static and dynamic margining solution is able to leverage the power-aging tradeoffs while adapting to a wide range of IoT applications.

VI. Conclusions

The diversity of IoT applications and markets leads to a plethora of requirements for IoT reliability. In this paper, we showed that circuit aging introduces new challenges for IoT domain along with energy efficiency. This work demonstrates that circuit aging should be carefully addressed in system design cycle. This paper also presents different static and dynamic solutions to compensate aging in IoT systems. Future work includes the implementation of the dynamic solutions on chip and development of a system level simulation framework for IoT aging to guide the design decisions.

Acknowledgements

This work was funded by NSF grants CCF 1255907, 1619127 and 1543837, by SRC 2410.001, by DARPA under the UPSIDE and PERFECT programs and by C-FAR, one of six SRC STARnet Centers, sponsored by MARCO and DARPA. Xinfei Guo was also supported by an IEEE CASS Pre-Doctoral Scholarship.

References

[1] M. Alioto, "Enabling the Internet of Things: From Integrated Circuits to Integrated Systems." Springer, 2017.
[2] X. Guo *et al.*, "Back to the Future: Digital Circuit Design in the FinFET Era," *Journal of Low Power Electronics*, vol. 13, no. 3, pp. 338–355, 2017.
[3] P. Madoglio *et al.*, "13.6 A 2.4 GHz WLAN digital polar transmitter with synthesized digital-to-time converter in 14nm tri-gate/FinFET technology for IoT and wearable applications," in *ISSCC*. IEEE, 2017, pp. 226–227.
[4] J. H. Stathis *et al.*, "Reliability challenges for the 10nm node and beyond," in *IEDM*. IEEE, 2014, pp. 20–6.
[5] C. Schlünder *et al.*, "On the influence of BTI and HCI on parameter variability," in *Reliability Physics Symposium (IRPS), 2017 IEEE International*. IEEE, 2017, pp. 2E–4.
[6] H. Hong *et al.*, "Lifetime Reliability Enhancement of Microprocessors: Mitigating the Impact of Negative Bias Temperature Instability," *ACM Computing Surveys (CSUR)*, vol. 48, no. 1, p. 9, 2015.
[7] A. Mutschler, "Transistor Aging Intensifies At 10/7nm And Below," https://semiengineering.com/transistor-aging-intensifies-10nm/, 2017, [Online; accessed 13-July-2017].
[8] J. Franco *et al.*, "BTI reliability of ultra-thin EOT MOSFETs for sub-threshold logic," *Microelectronics Reliability*, vol. 52, no. 9, pp. 1932–1935, 2012.
[9] J. Sather, "Battery technologies for iot," in *Enabling the Internet of Things*. Springer, 2017, pp. 409–440.
[10] S. Lerner *et al.*, "Workload-aware ASIC flow for lifetime improvement of multi-core IoT processors," in *ISQED*. IEEE, 2017, pp. 379–384.
[11] A. Sivadasan *et al.*, "Architecture-and workload-dependent digital failure rate," in *IRPS*. IEEE, 2017, pp. CR–8.
[12] V. van Santen *et al.*, "Reliability in Super-and Near-Threshold Computing: A Unified Model of RTN, BTI, and PV," *TCAS-I*, 2017.
[13] J. Gubbi *et al.*, "Internet of Things (IoT): A vision, architectural elements, and future directions," *Future generation computer systems*, vol. 29, no. 7, pp. 1645–1660, 2013.
[14] T. Grasser *et al.*, "The "permanent" component of NBTI revisited: Saturation, degradation-reversal, and annealing," in *IRPS*. IEEE, 2016, pp. 5A–2.
[15] X. Guo and M. R. Stan, "Work hard, sleep well-Avoid irreversible IC wearout with proactive rejuvenation," in *ASP-DAC*. IEEE, 2016, pp. 649–654.
[16] M. Cho *et al.*, "Postsilicon Voltage Guard-Band Reduction in a 22 nm Graphics Execution Core Using Adaptive Voltage Scaling and Dynamic Power Gating," *IEEE JSSC*, vol. 52, no. 1, pp. 50–63, 2017.