A Review of Techniques for Ageing Detection and Monitoring on Embedded Systems

LEANDRO LANZIERI, Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany
GIANLUCA MARTINO, Institute of Embedded Systems, TU Hamburg, Hamburg, Germany
GOERSCHWIN FEY, Institute of Embedded Systems, TU Hamburg, Hamburg, Germany
HOLGER SCHLARB, Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany
THOMAS C. SCHMIDT, Computer Science, HAW Hamburg, Hamburg, Germany
MATTHIAS WÄHLISCH, Computer Science, TU Dresden, Dresden, Germany and Barkhausen Institut, Dresden, Germany

Embedded digital devices are progressively deployed in dependable or safety-critical systems. These devices undergo significant hardware ageing, particularly in harsh environments. This increases their likelihood of failure. It is crucial to understand ageing processes and to detect hardware degradation early for guaranteeing system dependability. In this survey, we review the core ageing mechanisms, and identify and categorize general working principles of ageing detection and monitoring techniques for Commercial-Off-the-Shelf (COTS) components that are prevalent in embedded systems: Field Programmable Gate Arrays (FPGAs), microcontrollers, Systems-on-Chips (SoCs), and their power supplies. From our review, we find that online techniques are more widely applied on FPGAs than on other components, and see a rising trend towards machine learning application for analysing hardware ageing. Based on the reviewed literature, we identify research opportunities and potential directions of interest in the field. With this work, we intend to facilitate future research by systematically presenting all main approaches in a concise way.

CCS Concepts: • **Computer systems organization → Embedded hardware**; • **Hardware → Ageing of circuits and systems**; **Built-in self-test;**

Additional Key Words and Phrases: FPGA, microcontroller, power supplies, ageing monitoring, hardware health

Leandro Lanzieri also with TU Hamburg, Hamburg, Germany and HAW Hamburg, Hamburg, Germany.

We acknowledge the support of DASHH (Data Science in Hamburg — HELMHOLTZ Graduate School for the Structure of Matter) with grant no. HIDSS-0002 as well as the support of the German Federal Ministry of Education and Research with grant C-ray4edge.

Authors' Contact Information: Leandro Lanzieri, Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany; e-mail: leandro.lanzieri@desy.de; Gianluca Martino, Institute of Embedded Systems, TU Hamburg, Hamburg, Germany; e-mail: gianluca.martino@tuhh.de; Goerschwin Fey, Institute of Embedded Systems, TU Hamburg, Hamburg, Germany; e-mail: goerschwin.fey@tuhh.de; Holger Schlarb, Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany; e-mail: holger.schlarb@desy.de; Thomas C. Schmidt, Computer Science, HAW Hamburg, Hamburg, Germany; e-mail: t.schmidt@haw-hamburg.de; Matthias Wählisch, Computer Science, TU Dresden, Dresden, Sachsen, Germany and Barkhausen Institut, Dresden, Sachsen, Germany; e-mail: m.waehlisch@tu-dresden.de.

ACM 0360-0300/2024/10-ART24
https://doi.org/10.1145/3695247

ACM Reference Format:
Leandro Lanzieri, Gianluca Martino, Goerschwin Fey, Holger Schlarb, Thomas C. Schmidt, and Matthias Wählisch. 2024. A Review of Techniques for Ageing Detection and Monitoring on Embedded Systems. *ACM Comput. Surv.* 57, 1, Article 24 (October 2024), 34 pages. https://doi.org/10.1145/3695247

1 Introduction

The ubiquity of embedded devices keeps growing, driven by increasing chip integration levels and computing capacities, as well as by reduced prices and power consumptions. Applications ranging from consumer electronics and **Internet of Things (IoT)** gadgets to space missions, nuclear power plants, and particle accelerators involve embedded systems. Many deployments challenge the lifetime of these devices by making use of unhardened **Commercial Off-the-Shelf (COTS)** components, by extending mission times over the guaranteed life span, or by operating in harsh conditions. Nevertheless, properly performing devices are essential for effective system operation. Performance is crucial for the timely gathering and processing of data, and essential for accurate decision-making and actuating, for example, in industrial environments or in autonomous vehicles.

Microcontrollers, Systems-on-Chip (SoCs), and **Field Programmable Gate Arrays (FPGAs)** are the most widely deployed digital components in embedded applications. Their prevalent use in exceptionally large multiplicities makes them highly relevant, which motivates the focus of this survey (Figure 1).

Microcontrollers and SoCs can perform complex tasks at low cost and low power capabilities. They can even run operating systems with full-featured IPv6 software stacks [8], which makes them a popular central component in embedded scenarios. FPGAs offer a high level of flexibility compared with **Application-Specific Integrated Circuits (ASICs)**. Serving even demands from **Artificial Intelligence (AI)** on the **Adaptive Compute Acceleration Platforms (ACAPs)**, which combine **central processing unit (CPU)** cores, programmable logic, and AI inference acceleration, FPGAs open their way into the edge-computing and high-end edge IoT markets. All systems require power conditioning and delivery with stable parameters over the duration of their deployment. Failures in such power delivery modules are usually fatal for the entire system due to their prominent and critical role. Nonetheless, it has been reported that the reliability of power systems remains inferior to other modules [40, 97, 128] partially due to their complexity and multitude of components.

Hardware ageing is a severe problem for embedded devices [46]. Hence, ageing detection plays a key role in dependable or safety-critical applications. The impact of physical degradation increases for new, highly integrated technologies [59, 78]. A particular issue for many **Integrated Circuits (ICs)** in long-term deployments is transistor degradation, namely, **Bias Temperature Instability (BTI)**, **Hot-Carrier Injection (HCI)**, and **Time-Dependent Dielectric Breakdown (TDDB)**. Their predominant consequence is an increment of the transistor threshold voltage, which reduces the maximum switching speed of **Complementary Metal-Oxide-Semiconductor (CMOS)** logic circuits, which, in turn, potentially affects critical paths in digital applications and the noise margins on digital gates [58].

Embedded systems are composed of a variety of basic components, most of which have been individually analysed [35, 67, 100, 112]. To assess the reliability of embedded systems, a comprehensive approach to detecting failures [122] and monitoring ageing of the entire hardware boards is missing. The present survey closes this gap.

In this article, we concentrate on techniques that detect degradation processes on COTS components without requiring ageing sensors built in by vendors on the chips. Existing surveys

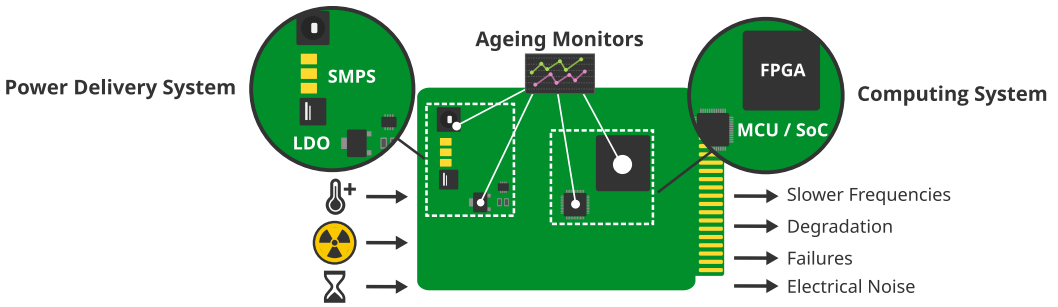

Fig. 1. Ageing monitoring is required to ensure the reliability of embedded systems, which are affected by environmental and operational conditions. This survey covers FPGAs, microcontrollers, and SoCs together with their power supplies as the prevalent system components.

focus on on-chip monitoring solutions [98] or on reconfiguration and ageing monitoring solutions [56]. This excludes the wide variety of industrial and commercial deployments that utilize COTS components without dedicated ageing monitoring [18]. In contrast, our survey places a special focus on techniques that can be applied to COTS components and that require no special silicon sensors to operate. Khoshavi et al. [60] presented and categorized ageing monitoring and mitigation techniques on CMOS devices, focusing on ring-oscillator sensors. Kochte and Wunderlich [65] reviewed literature on self-tests and diagnosis techniques to improve the self-awareness of digital systems. The authors propose a test classification based on the moment when the self-evaluation takes place in a system. We expand this classification work to address ageing detection and monitoring techniques that have not been covered by previous surveys in the field.

The remainder of this article is structured as follows. In Section 2, we describe the ageing mechanisms that affect embedded systems. This includes the degradation of transistors, interconnects, and passive components. Section 3 provides an overview of the studied literature, introduces relevant taxonomies and the proposed classification system, and identifies trends in the techniques. Sections 4 to 6 describe in detail ageing detection and monitoring strategies for FPGAs, microcontrollers and SoCs, and power supplies, respectively. Section 7 presents a discussion on possible research gaps in the field and an outlook with concrete future research directions. The article concludes in Section 8.

2 Ageing Mechanisms in Embedded Systems

A failure is a non-conformance to a defined performance criterion [110]. Dependable systems require reliability, which represents the probability of the system to stay failure free for a given period of time (i.e., to perform within specified limits). Design requirements usually specify failure rates $\lambda(t)$, which are the probability of failing at a certain point in time.

A component failure rate varies during its lifespan and is typically described using the bathtub distribution (Figure 2). The observed failures over time are the sum of three overlapping distributions — early failures, constant random failures, and wear-out failures — which lead to three distinct curve regions. The first region is the start-up, also known as *burn-in*, in which the failure rate decreases over time. Failures in this region are mostly due to manufacturing problems. The middle region is the useful life period, in which failures occur randomly at a constant rate. The last part of the distribution is the wear-out period. The acceleration of degradation mechanisms (e.g., component ageing) causes an increasing failure rate in this region.

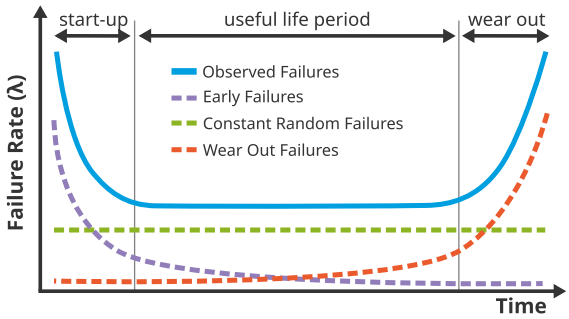

Fig. 2. Bathtub distribution illustrating the typical evolution of the component failure rates over time [110].

Components of embedded systems suffer degradation during their lifetime. Ageing processes are gradual, and their results are usually only noticeable in the long term. **Very Large-Scale Integration (VLSI)** chips are subject to multiple wear-out or degradation processes that shift their parameters away from specification, eventually resulting in system failures [114]. Ageing in chips has two sources: transistor and interconnect degradation. In addition, passive components such as capacitors and inductors age as well [35], causing problems in voltage regulation circuits [17].

In this section, we describe the most relevant ageing mechanisms affecting the core components that compose devices commonly found in embedded systems.

2.1 Ageing of Transistors

Silicon manufacturers encounter more challenges as they explore new miniaturized transistor technologies. Efforts to further reduce device footprints prompted moving from planar to **three-dimensional (3D)** topologies, such as FinFETs and **gate-all-around FETs (GAAFETs)**. This trend continues with vertical GAAFETs featuring vertical nanowires, which are predicted for usage beyond the 5 nm scale [121]. These devices are susceptible to self-heating because the thermal paths to the ambient environment in their complex 3D structures are constrained. An increase in operating temperature accelerates ageing processes and can reduce performance [24]. Additionally, the reduction in size has not been proportionally followed by a lowered operational voltage, causing stronger electric fields applied to the devices.

Next, we describe the main ageing mechanisms that impact **Metal-Oxide-Semiconductor Field-Effect Transistors (MOSFETs)**, the base for all current CMOS technologies, including newer technologies such as FinFETs.

2.1.1 Bias Temperature Instability (BTI). BTI is a silicon ageing mechanism [42, 106], most prevalent on p-channel MOSFETs as **Negative BTI (NBTI)**, which was initially observed on large technologies, such as 40 nm. Since the introduction of high-k/metal transistors, BTI effects are also considerable on n-channel MOSFETs, affected by **Positive BTI (PBTI)** [112]. As devices scaled down, the necessity to boost doping levels accentuated the prominence of the random discrete doping problem over NBTI [104]. Nevertheless, the ability to reduce doping levels due to the introduction of technologies such as FinFETs has resurfaced NBTI as the dominant time-dependent variability-inducing mechanism [3, 71], which is actively under study [58, 92].

NBTI induces an increase in the threshold voltage and a degradation of the carrier mobility, drain current, and transconductance on transistors. Although the mechanism is not fully understood and many models are being debated in research [78], the main cause for NBTI is attributed to the creation of traps in the oxide-substrate (SiO_2/Si) interface and charges in the oxide. The most

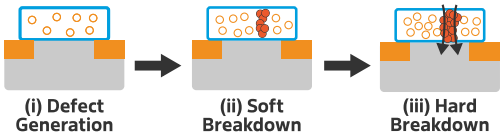

Fig. 3. Cross-section of a transistor with the gate oxide traversing all three stages of dielectric breakdown.

prevalent model is the **Reaction-Diffusion (R–D)** model [53, 90]. It proposes that low-energy SiH bonds on the interface break due to the presence of an applied electric field and high temperature, with a linear relation to the stress time (i.e., reaction). Then, positive charges (H^+ or trapped holes) diffuse into the gate oxide with a time dependence t^n (where usually $n = 1/4$), leaving dangling bonds or interface traps behind [106]. The accumulation of positive charges in the oxide opposes the applied electric field, thus inducing a variation ΔV_{th} in the threshold voltage of the transistor (V_{th}). This reduces the maximum switching speed of the device.

BTI has a recovery phase: once the gate bias is removed, the diffused charges return to the interface and anneal the traps. Although this process reduces ΔV_{th}, it is not completely clear whether the recovery is full: some evidence indicates the possibility of a 100% recovery [99], whereas other studies propose the existence of both permanent and reversible degradation components [32]. As the mechanism is partly governed by a diffusion process, its duration highly relates to the operating temperature. Owing to the recovery phase, degradation depends on the stress duty cycle; thus, continuous stress (DC) is much more severe than non-continuous (AC) [113].

2.1.2 Hot Carrier Injection (HCI). As a lateral electric field between drain and source is applied to a transistor, carriers (electrons or holes) gain kinetic energy, becoming *"hot"* when their energy is significantly larger than the one of the lattice at thermal equilibrium [116]. Due to their high energy, some carriers are able to surpass the potential barrier of the gate oxide, diffusing into the dielectric or causing damage to the interface. This has similar symptoms as BTI, namely, an increase of threshold voltage ΔV_{th}, a decrease of carrier mobility, and a reduction of transconductance in the transistor saturation region [50, 89, 117]. However, unlike BTI, HCI presents no recovery phase when the stress is removed. The impact of HCI is directly proportional to the frequency at which the device is switched, as hot carriers are generated during the transistor transitions. HCI also depends on the device temperature; therefore, its effects are exacerbated by the self-heating process of thermally complex technologies, such as FinFETs [29], for both p-MOS and n-MOS types [29, 63]. This situation is expected to worsen as 3D structures are built taller and narrower [55], making it a clear challenge for technology scaling.

2.1.3 Time-Dependent Dielectric Breakdown (TDDB). TDDB is a gradual and irreversible degradation mechanism in the gate oxide of transistors. Defects accumulated in the SiO_2 layer allow an increase in gate leakage current, thus slowing down the transistor switching frequency. This phenomenon has been studied on planar MOSFETs [100] as well as on newer 3D FinFETs [25, 62].

Nigam et al. [88] describe three stages of the degradation process as illustrated in Figure 3: (i) Defect generation starts when a large electric field is applied to the gate dielectric. As the defect density is still not critical, no gate current leaks. (ii) With continued stress, more defects are generated and start to overlap, producing a soft breakdown. At this point, traps form a resistive path between the gate and the channel, increasing the leakage current and reducing the device switching speed. (iii) Due to the thermal damage caused by the leakage current, more traps form a wider and less resistive path. This increases the gate current even further, which leads to a thermal runaway that completely breaks the dielectric layer. At this point, the transistor structure is destroyed and the device is unusable.

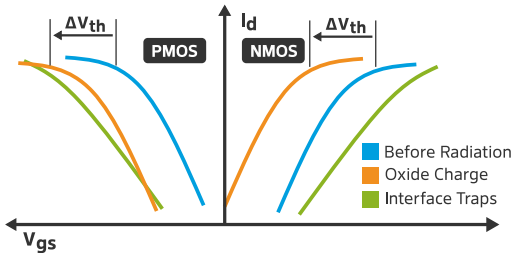

Fig. 4. Illustration of the effect that fixed oxide trapped charges have on the drain current (I_d) vs. gate-source voltage (V_{gs}) characteristic for N-MOS and P-MOS devices [10].

2.1.4 Radiation-Induced Trapped Charges. When high-energy photons or charged particles interact with a material, they can cause ionization. **Total Ionizing Dose (TID)** denotes the total amount of transferred energy from ionizing particles to the material. As a consequence, semiconductors are vulnerable to photon-induced ionization damage via a process that generates two types of trapped charge: (i) oxide-trapped charges and (ii) interface traps [10]. Trapped charges accumulate over time and modify the transistor characteristics. The effect of radiation has been studied in many components commonly found in embedded systems: signal propagation changes within ICs [105], impact on voltage regulators [84], SoCs, and microcontrollers [34, 64, 109].

The trap generation process initiates when an ionizing particle impacts the gate oxide material of the transistor, transferring part of its kinetic energy. A number of electron-hole pairs are generated, proportional to the material activation energy and density [91]. While some pairs are annihilated through recombination, which is a function of the applied electric field [108], other pairs escape this process. A fraction of the remaining holes may fall into deep traps in the oxide or near the Si/SiO_2 interface, forming the trapped positive charges. Other holes react with hydrogen-containing defects in the interface, generating interface traps.

Oxide traps modify the DC characteristics of CMOS circuits, similarly to BTI, as depicted in Figure 4. The most prominent effect is a negative shift in the drain current I_d for a given gate-source voltage V_{gs} on P-MOS and N-MOS. In the first one, the absolute value of V_{th} increases while the drain current is reduced. The latter suffers a reduction of V_{th}, with an increase of the drive current, potentially causing a latch-up (i.e., the transistor always conducts). Interface traps affect the recombination rates of carriers and their mobility through the channel. As the trapping and de-trapping of charges at the interface depend on the applied bias voltage, when these traps build up, they generate an increase in the subthreshold swing of CMOS devices. The consequence is that the I_d vs. V_{gs} response is stretched out, as shown in Figure 4.

Recent research on modern transistor technologies has shown an increased TID tolerance on 28 nm MOSFETs and 16 nm FinFETs, when compared with previous 65 nm nodes [14]. In the case of FinFETs, the TID tolerance was found to strongly depend on the channel length and the number of fins on the transistor [77]. Furthermore, Gorchichko et al. [41] measured an extremely high TID tolerance for GAAFETs, which is partly due to thin gate dielectrics that prevent building up charges and make them good candidates for radiation-enhanced applications.

2.2 Ageing of Interconnects

2.2.1 Electromigration. To keep up with Moore's Law [82], manufacturers constantly increase levels of integration. One way to integrate more devices is to reduce the width of the wires interconnecting them. This optimization comes at a cost: the smaller the conducting area, the higher the current density. Increasing current density entails a higher atomic diffusion phenomenon, called

electromigration. In the presence of high current densities, the strong momentum transfers from electrons to the conductor atoms to cause an atomic diffusion in the direction of the electron flow [118]. The force of the electronic wind causes atoms to deplete "up wind" and accumulate "down wind", until they create an electrical short- or open-circuit, rendering ICs unusable. Black [13] proposed to model the **Median Time to Failure (MTF)** in hours by Equation (1), where A is a constant, J is the current density, ϕ is the activation energy for diffusion, k_B is the Boltzmann constant, and T is the temperature in Kelvin.

$$MTF = AJ^{-2} \exp\left(\frac{\phi}{k_B T}\right) \tag{1}$$

2.3 Ageing of Passive Components

2.3.1 Capacitors. Even though it is well known that electrolytic capacitors wear out with time [35], they are widely used in many embedded applications, from filtering signals to suppressing voltage ripples. As digital circuits lower their operating voltage, they become more susceptible to noise and require functioning capacitors over the entire deployment time. Most of the models representing real capacitors in electrical circuits describe an **Equivalent Series Resistance (ESR)**, which causes energy dissipation as heat. Although ideally the ESR value is negligible, changes have been observed over time when the component is exposed to high temperatures [36]. The vaporization of the capacitor electrolyte through the encapsulation seal increases the ESR. As vaporization depends on the internal temperature and the ESR dissipates heat, it acts in positive feedback degrading the capacitor parameters.

2.3.2 Inductors. Inductors are commonly employed to build analog filters and power supplies. Magnetic cores are usually inserted inside the coils to increase inductance. Depending on their composition and geometry, ferromagnetic cores generate energy losses due to eddy currents in the material induced by the fluctuating magnetic field. Indeed, according to Faraday's law of induction, upon the presence of a changing magnetic flux Φ in an enclosed path, an opposite electromotive force $\mathcal{E}$ is induced. This force will be proportional to the rate of change of the flux, as Equation (2) shows.

$$\mathcal{E} = -\frac{d}{dt}\Phi \tag{2}$$

The electric flux crossing a surface Σ (bounded by a path C) can be expressed as Equation (3), where $\boldsymbol{B}$ is the magnetic field.

$$\Phi = \iint_{\Sigma} \boldsymbol{B} \cdot dA \tag{3}$$

The induced $\mathcal{E}$ relates to the generated electric field $\boldsymbol{E}$ over the closed path C, as shown in Equation (4).

$$\mathcal{E} = \oint_{C} \boldsymbol{E} \cdot dl \tag{4}$$

Finally, Equation (5) shows that the varying magnetic field in the enclosed inductor core will induce an electric field in the ferromagnetic material. Whenever the material of the core is conducting, a current density will appear on it.

$$\oint_{C} \boldsymbol{E} \cdot dl = -\frac{d}{dt} \iint_{\Sigma} \boldsymbol{B} \cdot dA \tag{5}$$

The power dissipation of the eddy currents depends on the induced electric field on the core and its resistivity. Inductor cores can undergo a degradation process when constantly exposed to elevated temperatures for long periods of time, which changes the material resistivity. This increases core losses [17] and contributes to further elevating the temperature.

3 Surveying Ageing Detection Techniques

A defect is an unintended deviation of a component material of physical nature and persistent effects [65]. Defects originate either from manufacturing imperfections or from ageing during service. Defective behaviour can be abstracted to faults when modelled according to the circuit structure. For instance, a transistor threshold voltage degradation (i.e., a physical variation) can be modelled as a delay fault on a combinational circuit. When the system state or its environment triggers a fault, an error occurs. This, in turn, may cause a malfunction or complete failure.

The variety of potential defects of electronic components (see Section 2) evidences the need for health assessment of embedded hardware. Prognostics and remaining-useful-life estimations [52] can leverage this information. The evolution of physical ageing processes can indirectly be assessed during operation by observing deviations of device parameters, either by directly comparing against pre-established profiles or with sophisticated anomaly detection techniques [80]. This opens the door to fully ageing-aware embedded systems that monitor, analyse, and act upon low-level degradation information, either by means of mitigation [61] or even by self-healing [45] techniques.

3.1 Surveying Methodology

The research for the presented literature comprised consulting various notable scientific databases (IEEE, ACM, Scopus, and Google Scholar) while limiting the results with keywords including *ag(e)ing, degradation, failure, monitoring*, and the components of interest: FPGA, microcontroller, SoC, and power supply. We excluded, however, board-level mechanical wear-out effects, such as broken interconnects or the degradation of solder joints. Upon the initial searches, we followed the related work referenced by the results. References were selected with the objective of maximizing the variety of detection and monitoring approaches. We did not include work on the design of special ageing hardware probes for ASICs. Instead, we focused on techniques applied to COTS devices without built-in on-chip ageing sensors. The presented techniques have been implemented in real hardware; therefore, their results are empirical and not exclusively based on simulations. This literature compendium includes methods originally designed to assess the system health status while on service as well as methods that have been employed to study the impact of ageing processes on devices under controlled environments.

We consider that the selected components represent well the minimal building blocks present in most modern embedded systems: a main processing unit, memory, and a power source [81]. While a huge variety of peripherals, such as sensors and actuators, are frequently required on certain embedded system domains (e.g., cyber-physical systems and IoT [46]), they are use-case specific in nature and can vary fundamentally in their functioning principles. We aim to abstract from particular deployment cases for embedded devices. Instead, we aspire for this survey to serve as a comprehensive guide across embedded domains and, thus, focus on the body of work concerning FPGAs, MCUs, SoCs, and power supplies.

3.2 Classification of Tests

Given the wide variety of tests that can be performed on the systems under consideration, possible classifications vary depending on the intended use case. In this work, we adapted the classification system proposed by Kochte and Wunderlich [65] and extended it as presented in Figure 5.

The primary variable is the highest testing frequency that each technique can potentially achieve. The classification distinguishes between infrequent testing that requires a system shutdown or special testing equipment, i.e., *offline testing*, and testing that can be performed autonomously onboard, i.e., *online testing* or self-testing. Online testing can either be *concurrent* or *non-concurrent* depending on whether it runs without influencing the system operation.

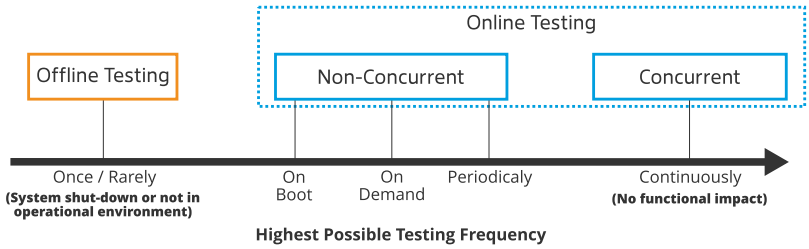

Fig. 5. Test types organized according to their highest testing frequency. We have expanded the classification system proposed by Kochte and Wunderlich [65], which focuses on self-testing systems, in order to apply it to ageing detection techniques.

3.2.1 Offline Testing. Offline testing usually involves automatic testing equipment and extra hardware or tools. These tests are typically required whenever the observed magnitudes need special detection techniques or sensors not present in the system (e.g., radiated electromagnetic emissions). Offline tests tend to be much more expensive and cumbersome to carry out than online tests. For this reason, they are mainly reserved for in-lab evaluations under controlled environments and are rarely deployed in the field. Potential applications of offline tests are acceptance testing for quality assurance (e.g., to detect counterfeits) and post-mortem analysis.

3.2.2 Online Testing. Online tests, or self-tests, leverage the existing capabilities of the devices under test to evaluate their own functionality, thus avoiding extra equipment. Depending on their measurement methodology, online tests are (i) concurrent or (ii) non-concurrent. The former run in parallel with the applications and are useful to monitor systems that cannot afford downtime. The latter are performed outside the device service schedule (e.g., during boot or maintenance periods).

Concurrent Online Testing is a non-intrusive testing approach that requires no suspension of the system service. Such techniques are ideal for ageing detection in safety-critical deployments that operate without interruption for extended periods of time, e.g., nuclear plants or large experimental facilities.

Non-concurrent Online Testing consists of tests that cannot be performed on a running system due to their intrusiveness, e.g., the test requires a reconfiguration in order to be performed. Ageing detection using non-concurrent online tests is possible for systems that can be periodically rebooted or reconfigured.

3.3 Overview of the Results

Considering the criteria described in Section 3.1 and the taxonomy introduced in Section 3.2, relevant literature has been selected, reviewed, and classified. Table 1 summarizes this related work, which is discussed in detail in the following sections.

Approaches are organized by component: FPGA, **Microcontroller Unit (MCU)** and SoC, and **Low Drop-Out (LDO)** and **Switching-Mode Power Supply (SMPS)**. Contributions are further grouped according to their sensing principle. Within each group, references are chronologically ordered so that the latest appear at the end. For each publication, we indicate whether the technique was developed as an online concurrent, online non-concurrent, or offline test. We expose whether techniques were evaluated under aged conditions, whether **machine learning (ML)** techniques were applied, and whether an analytical or empirical model was developed to study the degradation process.

At first glance, Table 1 reveals that most of the literature focusing on FPGAs proposes online test approaches, whereas related work on MCUs, SoCs, or power supplies dominantly relies on

Table 1. Summary of Ageing Detection and Monitoring Techniques Grouped by Component and Sensing Principle

Component	Sensing Principle	Reference	Online		Offline	Ageing	ML	Model
			Conc.	Non-conc.				
FPGA (Section 4)	Shadow Register	Li and Lach [75]	✔					
		Wong et al. [123]		✔				
		Wong et al. [125]		✔				
		Valdes et al. [120]	✔					
		Amouri and Taoori [7]	✔					
		Pfeifer and Pliva [94]	✔					
		Leong et al. [72]	✔					
		Valdes et al. [119]	✔					
		Ghaderi et al. [37]	✔					
		Jiang et al. [54]		✔				
	Ring Oscillator	Ruffoni and Bogliolo [103]			✔			
		Sedcole and Cheung [107]		✔				✔
		Zick and Hayes[129]	✔					✔
		Bruguier et al. [21]			✔			
		Pfeifer and Pliva [95]		✔				
		Pfeifer and Pliva [96]		✔		🌡		
		Amouri et al. [6]			✔	ϟ, 🌡		
		Pfeifer et al. [93]		✔		⧗		
		Naouss and Marc [85]		✔		ϟ		
		Naouss and Marc [87]		✔		ϟ		✔
		Naouss and Marc [86]		✔		ϟ		✔
		Maragos et al. [79]		✔				
		Bender et al. [12]		✔		ϟ, 🌡		✔
		Ahmed et al. [4]			✔	🌡	✔	✔
		Li et al. [76]		✔		ϟ, 🌡	✔	✔
		Sobas and Marc [111]		✔		ϟ, 🌡, ⧗		✔
		Lanzieri et al. [69]		✔		⧗, ☢	✔	
	Transition Probability	Wong et al. [124]		✔				
		Stott et al. [114]		✔				
		Stott et al. [115]		✔		ϟ, 🌡		✔
MCU & SoC (Section 5)	SRAM Pattern	Guo et al. [47]		✔		ϟ, 🌡		
		Guo et al. [48]		✔		ϟ, 🌡		
		Guin et al. [44]		✔		🌡		
		Lanzieri et al. [70]		✔		⧗	✔	
	Time Window	Diggins et al. [31]			✔	☢		
	LDO PSRR	Chowdhury et al. [27]			✔	🌡	✔	
		Acharya et al. [2]			✔		✔	
	EM Charact.	Dawson et al. [30]			✔	🌡		
		Li et al. [74]			✔	🌡		
		Wu et al. [126]			✔	ϟ, 🌡		

(Continued)

Table 1. Continued

Component	Sensing Principle	Reference	Online		Offline	Ageing	ML	Model
			Conc.	Non-conc.				
Power Supply (Section 6)	Equivalent Series Resistance	Lahyani et al. [68]	✔			🌡		✔
		Chen et al. [23]	✔			🌡		
		Boyer et al. [17]			✔	🌡		✔
		Givi et al. [39]	✔					
	EM Charact.	Wu et al. [127]			✔	ϟ	✔	
		Boyer et al. [17]			✔	🌡		✔
		Boyer et al. [16]			✔	🌡		✔
	PSRR	Chowdhury et al. [28]			✔	ϟ, 🌡		
		Chowdhury et al. [26]			✔	🌡	✔	

Within each sensing principle, references are chronologically organized so that newer ones are at the bottom. For each reference, the type of test is indicated (concurrent, non-concurrent, or offline), whether machine learning models have been applied, and whether a model of the degradation processes has been developed. Additionally, the table indicates whether the technique has been evaluated under electrical stress (ϟ), thermal stress (🌡), radiation (☢), natural ageing (⌛), or a combination of them.

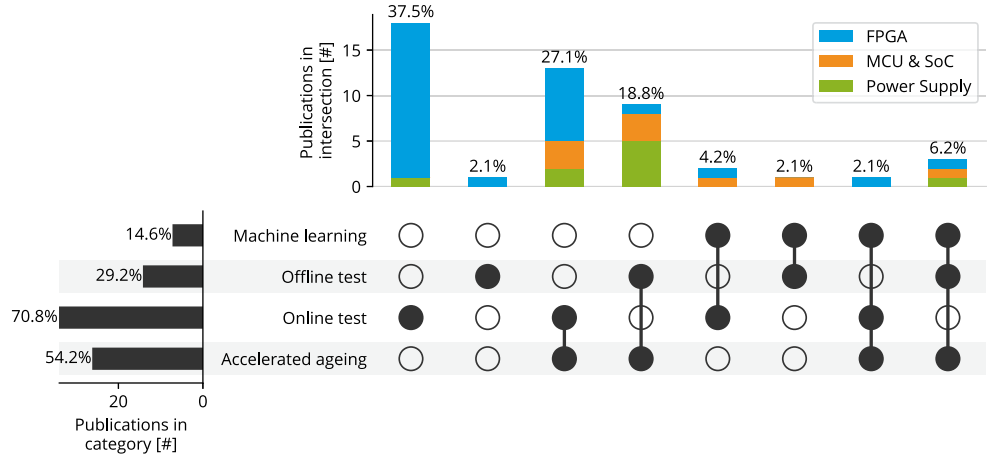

Fig. 6. The UpSet plot visualizes exclusive intersections between publications in different categories. The left bars show the total number of references that belong to each category and the percentage of the total. The upper bars illustrate how many of the publications belong to each exclusive intersection of categories, which are indicated in the matrix below. Each intersection (i.e., each vertical bar) includes papers in the highlighted categories (●) and excludes the ones in the non-highlighted categories (◯).

offline tests using external equipment. While various publications in the field of FPGAs and power supplies develop analytical and empirical models of the ageing mechanisms at play, this is not the case for MCUs and SoCs. In FPGA studies, ring oscillators are the preferred technique for characterizing and modelling. Finally, we can also observe a trend towards the application of ML in recent years, mainly for ageing detection techniques.

By observing the intersections of categories, we can further characterize the state-of-the-art. The UpSet plot [73] in Figure 6 shows the various shares of publications that fit into one or several classification categories. Even though 70% of the techniques are based on online tests, this number reduces to 27% if we restrict to those evaluated under accelerated ageing conditions, of which more

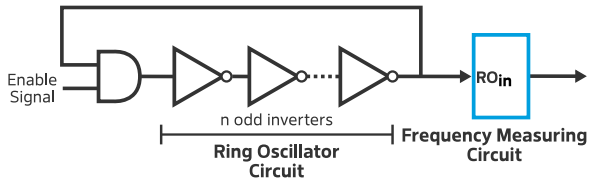

Fig. 7. Generic sensor based on a ring oscillator with an enable signal (see [9]).

than half apply to FPGAs. This differs for techniques describing offline tests, of which more than 85% are tested using accelerated ageing. Notably, more than 57% of the techniques that involve ML models require offline tests to work.

4 Ageing Detection on FPGAs

FPGAs are programmable hardware devices that provide great flexibility, in-field updates, and rapid prototyping of digital designs. The majority of FPGA applications require high-speed processing, including digital filters and control loops with strict low latency requirements. **Static Random Access Memory (SRAM)**–based FPGAs are the most popular in the market and use volatile memory to store their configuration. These devices present a regular structure of configurable logic blocks and programmable switch matrices that allow for a mapping of logic functions.

FPGAs face reliability challenges as their transistor-based circuits are affected by ageing mechanisms. BTI and HCI increase signal propagation delays [6]. With the gradual decay of delays in the signals, critical paths of digital applications may change and even reach values out of specification.

Given their general purpose nature, and unlike ASICs, manufacturers cannot know beforehand what functionality will be synthesized on the FPGA chip, hindering the implementation of precautionary measures to cope with circuit degradation (e.g., guard-bands). This situation moves the issue of monitoring ageing processes to the application developers, who are required to know the tolerances of signal propagation times for their specific use case.

4.1 Ring Oscillator

The flexibility of FPGAs enables the implementation of monitoring techniques directly on the programmable hardware. **Ring Oscillators (ROs)**, which consist of an odd number of inverters connected in a ring, are one of the most common approaches to measuring digital signal propagation time. RO-based sensors are usually composed of an RO and a frequency-measuring circuit (Figure 7). The frequency at the output of an RO depends —at a given temperature and voltage — on the inverter propagation times t_p and the number of gates n, as shown in Equation (6).

$$f = \frac{1}{2 \cdot n \cdot t_p} \tag{6}$$

Zick and Hayes [129] proposed an online concurrent sensing method to measure variations in physical parameters. By using an enhanced RO, an efficient counter, and control logic, the authors developed a compact sensor requiring only 8 **Look-Up Tables (LUTs)**. A residue number system ring counter was implemented, as it requires fewer resources than a binary counter. The temperature sensitivity of the RO was increased to detect hotspots on the chip. RO sensors were placed regularly on a hexagonal tessellation all over the FPGA, together with a softcore, a timer, and a **Universal Asynchronous Receiver/Transmitter (UART)**. The authors measured propagation

times and indirectly estimated a transistor current leakage profile, and localized dynamic power usage and temperature.

Sedcole and Cheung [107] placed ROs in a matrix configuration to measure within-die delay variability. The structure enabled the encoding of regions under test into columns and rows. In such a dense configuration, ROs should only be active for short periods of time to avoid heating neighbouring sensors. A single counter, timer, and unified control logic performed the measurements. Pfeifer and Pliva characterized chip delays during design time [95] and analysed the limitations of 28 nm FPGAs [96] with ROs. To measure the oscillations, they implemented a method based on **Block RAM (BRAM)**, which was later formalized in an approach called "Reliability-on-Chip" [93]. The approach requires a true-dual-port BRAM on the chip to undersample the oscillator outputs and a softcore to read the signal streams for calculating the delays. Li et al. [76] not only studied the process variability of 28 nm FPGAs under NBTI by deploying ROs but also utilized the data extracted from accelerated ageing to train various ML models. By building datasets with information regarding the artificial ageing stress conditions and the RO sensor configurations, they found that the XGBoost model performed best across all conditions.

Lanzieri et al. [69] developed an RO measurement module to perform a large-scale study of the propagation delay on 298 Xilinx Virtex-6 devices, which have been naturally aged and in operation as part of a linear particle accelerator. By comparing delay measurements of used and unused devices, the authors found evidence of effects caused by ageing mechanisms. Moreover, an analysis of the radiation exposure of the devices inside the accelerator showed that slower propagation delays are correlated to higher radiation doses.

The application of ROs as variability sensors extends to more modern node technologies as well. Maragos et al. [79] explored increased intra- and inter-die variability on 16 nm FinFET FPGAs with various ROs controlled by an embedded Cortex-A53 CPU. After an ageing process of 8,000 h, Sobas and Marc [111] measured the degradation of various 16 nm FinFET FPGAs by implementing an RO-based test bench. The authors evaluated the effects of static and dynamic stress, and derived an empirical degradation model for both modes that yields estimations with less than 10% relative error. When comparing their results to similar evaluations on 28 nm MOSFETs, they found that the smaller nodes show better reliability on static stress and that BTI was the predominant ageing mechanism (as opposed to HCI). A similar conclusion was reached by Bender et al. [12] from assessing 16 nm FinFET Xilinx devices with a multi-temperature operational life testing method. By evaluating the impact of different temperatures, voltages, frequencies, and RO sizes, the authors managed to isolate the effects of various degradation mechanisms. Additionally, their results suggest that there is a contribution from the self-heating effect to BTI due to the lower heat dissipation of the transistor fins.

Naouss and Marc [85] proposed a test bench to self-characterize the delay of LUTs. Their design allows stress signals to be injected into the **Circuit Under Test (CUT)** to produce accelerated ageing. This feature was later used to independently study BTI [87] and HCI [86] impact by using different stress signals. They implemented a frequency measuring circuit using three asynchronous counters (one of N bits and two of K bits) and a clock reference. This circuit (Figure 8) allowed for counting the number of cycles in a given period of time as well as the duty cycle of the signal. An enable signal activated the counters and controlled the counting window. Counter A registered the number of cycles of the RO output signal, which was used to calculate the oscillation frequency. Counter B counted the number of clock cycles during the active time of the RO signal, from which the duty cycle was calculated. Measuring the signal duty cycle allowed for studying the impact of ageing on the rising and falling times.

ROs have been employed in offline tests as well [4, 6, 21, 103]. Ahmed et al. [4] used performance degradation to detect recycled FPGAs by exhaustively fingerprinting LUT delays. They

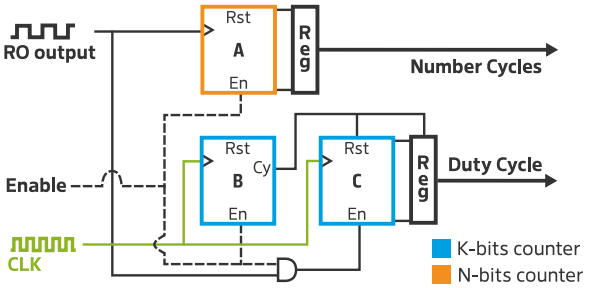

Fig. 8. Measuring circuit of the test bench proposed in [85], which counts the number of cycles of the RO signal and its duty cycle.

synthesized ROs with routes configurable via SRAM. External equipment was used to control the test logic and read out the frequency of each oscillator. Ruffoni and Bogliolo [103] focused on measuring delays of the internal FPGA wires. Two ROs were used, of which one included the wire structure under test. By comparing oscillation frequencies, the delay of the wire was derived. Although the authors employed external equipment, they argued that their method could potentially be operated solely on the device under test. Bruguier et al. [21] proposed a non-invasive method to characterize FPGA performance, analysing the spectra of electromagnetic radiation caused by the ROs. A similar measurement approach was used by Amouri et al. [6] to explore the impact of elevated temperatures and voltages on the performance degradation of FPGAs.

Discussion. RO sensors are relatively easy to implement and very versatile considering that they can provide insights into the effects of BTI and HCI when combined with methods such as multi-temperature operational life testing [12]. Additionally, they can be tuned to indirectly measure other quantities beyond the propagation delay [129]. Although the main sensing principle among approaches is similar, the literature presents multiple approaches to placing the sensors and counting the frequency. A trade-off exists with RO sensors: longer ring chains cover larger areas and require fewer measurement circuits but decrease sensor resolution. Care should be taken to avoid overheating the die or stress power rails by using too few stages in the rings, as this results in unrealistic measurements [95]. A downside of ROs is that they measure delays of the FPGA resources that form the ring and not of the synthesized circuits, which may differ depending on the position and routing. Other techniques measure delays of the existing combinational logic instead, as we describe in the following sections.

4.2 Shadow Register

The usage of **Shadow Registers (SRs)** is a well-studied technique for delay characterization and degradation monitoring. Sensors are usually placed at the end of critical combinational paths in parallel to a destination register for detecting late transitions.

Li and Lach [75], Valdés et al. [120], and Leong et al. [72] proposed to place an SR after the CUT that is clocked by a signal skewed from the destination register (Figure 9). By comparing the latched value on both registers and controlling the phase difference of the clock signals, they determine the delay of the CUT. Li and Lach [75] varied the phase difference in runtime to characterize the FPGA propagation delay and built a histogram. Leong et al. [72] implemented an online concurrent ageing monitoring sensor, which detected when the propagation delay was higher than a predefined threshold. Valdés et al. [120] included an on/off signal to their concurrent sensor that interrupts the clock, enabling authors to differentiate the type of ageing mostly suffered by the sensor itself: static ageing (continuous monitoring) or dynamic ageing (periodic monitoring). The

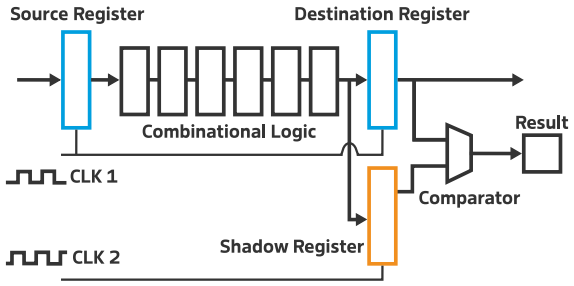

Fig. 9. Sensor based on a shadow register negatively skewed to measure the delay of a combinational logic circuit [75].

sensor functionality was initially tested by operating the circuit under different power supply voltages, which induced a change in its signal propagation delays but did not affect the sensor. The sensor from Valdés et al. was then validated [119] by performing an accelerated ageing process on an FPGA.

The authors reported no significant frequency variations of the clocks on which the sensor reliability depended after the burn-in process. Leong et al. [72] tested sensors by increasing the FPGA frequency and reducing the gap time.

Ghaderi et al. [37] proposed ageing monitors clocked by a single "sensor clock," which set the maximum allowed slack for combinational signals of critical paths. By injecting the CUT signal and a shifted version of it to an XOR, a positive pulse was generated on each transition of the CUT. The XOR was latched by a **Flip-Flop (FF)** and triggered by the sensor clock, thereby detecting invalid transitions whenever the pulse occurred too late.

Pfeifer and Pliva [94] presented an online concurrent delay-fault detection technique for combinatorial circuits. The authors used the D FFs at the input of on-chip BRAMs as SRs, which map the signals to memory rows for later analysis. The interconnect introduced a fixed delay between the destination and the SR to control the sensor sensitivity, and an on-board CPU performed the signal comparison.

Wong et al. [123] presented a self-characterization method, with two registers around a combinational CUT, clocked in counter-phase. An XOR between the CUT output and the SR latched value produced the error signal. Transitions occurring after the first half of the test clock period were invalid. The authors leveraged on-chip clock generation to sweep the test clock frequency until the maximum was found. A non-concurrent circuit for start-up tests was also proposed and optimized in [125], which stored test results of each region on the FPGA RAM.

Amouri and Tahoori [7] implemented an ageing sensor to detect late transitions of combinational paths on a Virtex-5 FPGA. The sensor, illustrated in Figure 10, was composed of two edge-triggered D FFs clocked by the combinational output, with their inputs connected to the principal clock signal. Whenever an invalid change in the combinational signal occurred (i.e., during the active clock cycle), the sensor output was activated. Two FFs were used to detect rising and falling signal transitions. By the addition of independently configurable delay blocks on the combinational output signal and the clock signal, the authors could control the sensor sensitivity (i.e., how late after the rising clock signal a change in the combinational output is detected). When sensitivity is configured to a negative value, the sensor is turned into an early warning monitor, which checks that the signal is stabilized at least by a given time before the clock rises. In comparison with the previously described work, this method bears the great advantage of not requiring extra clock resources for the sensor. On the downside, the approach does not quantitatively measure the propagation delay but rather detects transitions only when slower than a given threshold.

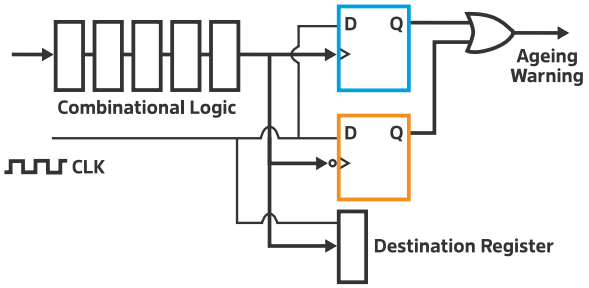

Fig. 10. Sensor implemented in [7] to detect late transitions of combinational logic.

Jiang et al. [54] proposed a similar architecture but connected the inputs Q of both SRs to a shadow clock signal, which had a phase shift relative to the main clock. Given that the frequency of the main clock is known, the authors were able to derive the CUT delay by changing the phase angle between clocks and observing the sensor output.

Discussion. Shadow registers appear more complex to implement and place than ROs, but they provide higher versatility. These sensors enable measuring propagation delays of application-specific circuits (e.g., to characterize chips) as well as implementing late transition detectors. For the detection of transitions within a given time window, this method can verify circuit functionality under different conditions and can even run for continuous monitoring of critical combinational paths to detect degradations caused by BTI and HCI. While ROs either act as probes on unused die areas or temporally replace functioning circuits to test the underlying hardware, SRs are able to run in parallel to application-specific combinational logic. SRs enable at-speed tests, which is a great advantage as they can seamlessly be added concurrently to applications at the cost of additional resource usage.

4.3 Circuit Transition Probability

Propagation delay variations can be detected by observing the **Transition Probability (TP)** of a circuit [114, 115, 124]. Consider a combinatorial digital circuit with an output node z. For each applied input combination, an output value $z(k)$ is produced. The *transition probability* of z, denoted $D(z)$, is the probability of the state changing when the next input stimuli are applied [38] on the following clock cycle. As z can only be zero or one, $D(z)$ is the probability of z experiencing a transition between these states:

$$D(z) = p_z^{10} + p_z^{01}, \tag{7}$$

where p_z^{01} and p_z^{10} indicate the probability of z undergoing the $0 \rightarrow 1$, and $1 \rightarrow 0$ transitions respectively. From [38], this probability can be calculated as the relative number of transitions that occurred in an interval of N clock cycles with $N \rightarrow \infty$.

As an example, p_z^{10} can be defined as

$$p_z^{10} = \lim_{N\to\infty} \frac{1}{N} \sum_{k=1}^{N} z(k)\overline{z(k+1)}. \tag{8}$$

If the probabilities are approximated by observing the transitions during a large number N of clock cycles, we obtain

$$D(z) \approx \frac{1}{N} \sum_{k=1}^{N} \left\{ z(k)\overline{z(k+1)} + \overline{z(k)}z(k+1) \right\}. \tag{9}$$

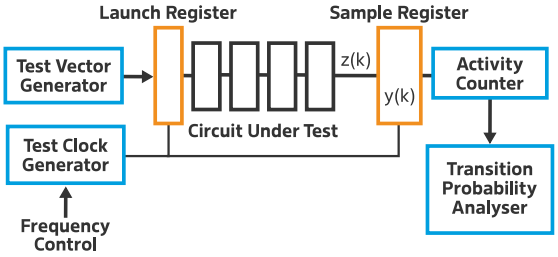

Fig. 11. Schematic of the delay measurement method proposed in [124] based on signal transition probability.

Hence, $D(z)$ can be estimated by the relative amount of rising and falling edges of z over a time interval N.

Ghosh et al. [38] derived a theorem that relates the output value probability with the input value probability on a combinatorial circuit. If the input signals have a probability distribution independent of time (i.e., they form a stationary process), then the output signal z will also have this characteristic. This means for stationary input signals that the TP $D(z)$ does not change in time.

Wong et al. [124] estimated the maximum functioning frequency of an arbitrary circuit by measuring its output TP. With a careful selection of the input signals, they ensured a constant TP of the output under normal operating frequencies. They performed various measurements at increasing frequencies, up to the point at which changes in the TP could be observed. The change indicated that the maximum frequency was reached, and the circuit started to fail. The proposed setup was implemented on a 65 nm Altera Cyclone III FPGA, as shown in Figure 11. The CUT and registers were clocked from a test clock generator. On each clock cycle, the test vector generator injected input vectors, which propagated through the CUT, generating an output $z(k)$. In addition, the sample register captured a sample $y(k)$ from the CUT at a frequency f_{clk}. An asynchronous counter recorded the transitions in $y(k)$ over N clock periods, later used to estimate $D(y)$ in the TP analyser circuit. When f_{clk} is within the operational range and no faults occur on the circuit, then $y(k) = z(k)$ and the transition probabilities $D(y) = D(z)$. If f_{clk} is increased above the CUT propagation time, then y will start to sample values of the previous cycle (i.e., $y(k) = z(k-1)$), thus, changing $D(y)$.

The method proposed by Wong et al. [124] was later applied in the study of circuit degradation under accelerated stress conditions by Stott et al. [114, 115]. The authors implemented multiple CUTs on a pair of Cyclone III FPGAs, which could be measured using the TP method and allowed to be electrically stressed by an input signal. Environmental stress with an ageing acceleration factor of 180 was applied to the chips by means of elevated temperature and core voltage, which sped up the NBTI process. Additionally, the CUT was subjected to electrical stress by controlling its switching activity through the input signals, which triggered NBTI, TDDB, and HCI degradation mechanisms. Their experiments revealed a circuit speed reduction of up to 15% by the end of the test schedule. This stress condition degraded LUTs stronger than interconnects. Moreover, the method was verified against an RO-based (see Section 4.1) frequency measurement.

Discussion. Measuring changes in the TP of a CUT output allows for detecting its maximum operational frequency. Unlike ROs (Section 4.1), this method measures the propagation delay of the FPGA using existing circuits; thus, the evaluation of the impact of BTI and HCI processes on the application is more direct. On the one hand, this technique has the advantage of being implementable with common resources and only requires a controllable clock signal. On the other

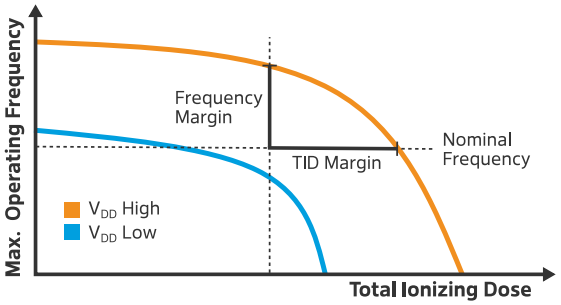

Fig. 12. Illustration of the timing window violation hypothesis proposed by Diggins et al. [31], in which the maximum operating frequency degrades as the total ionizing dose rises.

hand, it requires injecting precise test vectors, which depend on the CUT and need to be stored or consistently generated. As custom inputs are needed, this technique affects the system operation and can only be implemented during a testing period.

5 Ageing on Microcontrollers and SoCs

Microcontrollers and SoCs are digital ASICs. They suffer from similar degradation processes as FPGAs, but their functionality is known from the design process. This knowledge allows designers to study the expected ageing and the tolerances. Guard-bands can be placed to counteract chip degradation (e.g., limiting operating frequencies) and increase product reliability.

Many applications attempt to reduce costs by utilizing COTS components for longer than the guaranteed lifespans, which may raise critical conditions for highly dependable use cases. Despite the wide usage of microcontrollers and SoCs in embedded applications, studies on the impact of hardware ageing on their performance and useful lifetime are missing, in particular for in-field monitoring techniques. Indeed, most work on measuring degradation builds on utilizing external equipment. In this section, we present corresponding work, including techniques to assess the impact of ageing on microcontrollers and SoCs. We focus on COTS devices, which exclude special on-chip ageing sensors. Our motivation for this is that ASICs including dedicated ageing sensors are not commonly found on the market for embedded devices.

5.1 Timing Window Violation

Chip ageing can degrade signal propagation times, which also affects microcontrollers. Many low-end microcontrollers have a dedicated central clock signal and a short execution pipeline. Under such conditions, a violation of the timing windows, in which data is latched, will cause incorrect data to be propagated down the pipe.

Diggins et al. [31] conjectured a timing window violation hypothesis (Figure 12) which explains the relationship between (i) TID degradation, (ii) operating frequency, and (iii) supply voltage. For each voltage, a maximum operating frequency allows for executing software successfully. This frequency decreases as the TID increases [105]. The authors proposed an offline test to measure the impact of TID in the propagation delay of an ATMEGA328P microcontroller, observing violations on timing windows under multiple voltage and frequency conditions while the device was exposed to different values of TID and executing functional tests. Because the operating core frequency controls the length of the timing windows, software tests were executed at different frequencies. For each value of TID, the highest frequency at which the test passed was recorded. The authors reported that by overclocking the microcontroller above its nominal operational frequency, the degradation of the hardware was observable before a failure occurred at the nominal

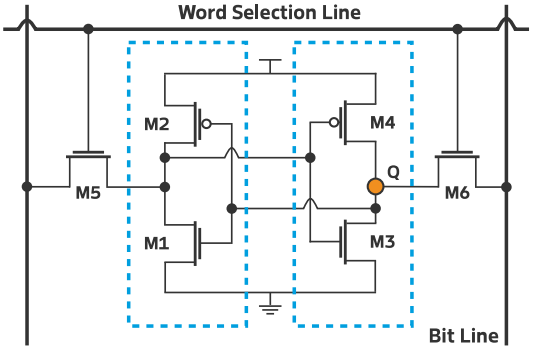

Fig. 13. Schematic of a six-transistor static RAM cell circuit. See [101] for background.

operating point. This observation bears the potential of developing monitoring techniques that activate predictive maintenance or graceful degradation approaches, such as operating at lower frequencies.

Discussion. The timing window violation hypothesis on microcontrollers has so far only been tested on TID-induced degradation. However, we envision the necessity to evaluate other ageing mechanisms that affect signal propagation delays. Although the analysis of timing window violation based on the clock frequency and operating voltage has mostly been tested in laboratories, great potential lies in online non-concurrent tests due to the high clock reconfigurability of many modern microcontrollers [102].

5.2 SRAM Initial Pattern

Many MCUs and FPGAs integrate static RAM (SRAM). A common implementation of an SRAM cell is the six-transistor circuit (Figure 13) owing to its relatively low static power usage. Each cell is composed of two cross-coupled inverters (M1, M2 and M3, M4) and two access transistors (M5 and M6). The cell has two stable states that represent either a logical 0 or 1, depending on the voltage at Q. The word selection line grants access to the cell and connects it to the bit line, which transfers data on read-and-write operations.

When a cell is energized, the initial value is not forced by the bit line; instead, it depends on the mismatch between the threshold voltage of the inverter transistors. The skew of an SRAM cell is its tendency towards a value when powered [49]. A non-skewed cell has a small threshold difference, and its initial value is random and depends on noise. These cells are commonly used as sources of entropy for random number generator systems. Cells with a moderate mismatch have a tendency towards a particular value, but this can be affected by ageing mechanisms that influence the threshold voltage of transistors, such as BTI. Finally, fully skewed cells produce the same initial value with a high probability and can be used to implement SRAM physically unclonable functions.

The **Static Noise Margin (SNM)** of an SRAM cell is the lowest level of voltage noise that flips its logic state. As the SNM depends on the transistor threshold voltage, ageing degrades it. Degraded SNM makes cells more sensitive to electric and thermal noise. A degradation of 15.02% of the SNM has been reported [33] after 10 years of usage owing to BTI effects. The skew of a cell is also impacted by BTI and HCI. Considering a cell that is constantly stressed by retaining a logical 1 on its output, the P-MOS M2 transistor in Figure 13 is active and undergoes a degradation process that modifies its threshold voltage during this time. If, upon startup, the gate threshold voltage of M2 is much larger than that of M4, the cell will be skewed towards 0 because M4 will be activated before M2.

The skew change of SRAM cells over time has been exploited by Guo et al. [47] as a two-phase method to detect recycled SoCs and FPGA chips. This was later refined as a framework [48]. Authors proposed an initial enrolment phase to detect partially skewed cells, which would likely change their startup value when aged. The authors first calculated the probability of starting in 0 and 1 for all SRAM bits, respectively, assuming room and high temperature for a given number of startup cycles. The underlying method foresees that high ambient temperature produces a good prediction for the aged state of the bits. By calculating the difference between the probabilities and comparing it to a pre-selected gap value, they selected the "ageing-sensitive"bits. A chosen threshold value determined the number of flipped bits required to consider a device as used (i.e., aged) during the verification phase. The gap value used to select the ageing-sensitive bits is not universal, and the authors indicated that it should be determined empirically based on measurements of aged devices beforehand. Tests on artificially aged Xilinx Spartan-3 FPGAs showed false accept and reject rates of 0 and 0.03, respectively, when picking the proper parameters.

Guin et al. [44] also proposed a method to detect recycled SoCs and FPGAs based on the assumption that SRAM cells are balanced by design. The authors argued that any bias introduced during manufacturing is random and that the amount of 1s and 0s normally distributes with a centre very close to 50% (i.e., within 1%). Therefore, any significant shift of the mean would be the result of circuit ageing. The technique requires no initial measurement phase or a golden model against which to compare. The approach only requires that the circuit is powered up multiple times to count the amount of 1s. The authors suggested that having a large enough SRAM array (e.g., 64 kbit) and taking sufficient measurements (e.g., 100) would minimize the impact of thermal noise during the assessment. Tests on COTS external SRAM chips showed significant changes in the distribution (up to 14% after ageing) and a slight recovery when stress was removed.

Lanzieri et al. [70] collected and studied SRAM startup patterns from MCUs of 154 naturally aged embedded devices deployed on an IoT testbed. The authors analysed the raw patterns and extracted subtle features, including spatial frequency components, bitwise probability of ones, and bit instability, to unveil correlations with the effective usage times of the devices. The observation of the extracted features revealed not only linear correlations with the usage time, but also discernible patterns in how the firmware interacts with the SRAM, such as initializing variables to 0 and determining their placement within the memory map. Finally, the authors trained and compared ML regression and classification models with the result of the feature engineering process, successfully evaluating the application of the SRAM startup pattern as a ubiquitous and inexpensive on-chip usage monitor.

Discussion. Given the ubiquity of SRAMs in various digital integrated circuits, the use of its initial pattern for indicating transistor degradation is an interesting and deployable method that requires no external equipment. The present work shows that BTI effects are already detectable on SRAMs after a few hours of accelerated ageing as well as in long-term naturally aged deployments. To take further advantage of modifications in SRAM startup patterns and to integrate them into higher-level health assessment systems, it is required to research correlations between the observed changes and the degradation of functional parameters of interest, such as the SNM.

5.3 Electromagnetic (EM) Characteristics

Electronic device ageing affects EM characteristics, such as conducted and radiated emissions [15] as well as susceptibility [74]. By subjecting devices to **EM Compatibility (EMC)** tests, their effective ageing could be determined.

Dawson et al. [30] proposed a method to study the effect of high temperature ageing on the EM emissions of a COTS MCU. The authors proposed to monitor device emissions as a reliable and straightforward measure to monitor ageing, as no electrical contact with the system is required.

An emission test board was built containing a Microchip PIC MCU with pins connected to loaded tracks passing by RF couplers. During tests, the software running on the device toggled the pins while a spectrum analyser captured EM emissions and the output voltage was measured at the pins. The authors reported variations in the emission spectra of artificially aged devices. Some harmonics of the central toggling frequency increased, particularly the high-frequency harmonics. It remains unclear whether a shift in the toggling frequency is the culprit. Additionally, the output voltage on aged devices increased, but there were no conclusions regarding the reasons for the reported behaviour. Although no clear mechanism linking the changes in emissions to device ageing was presented, BTI is a likely candidate due to the temperature-triggered nature of the changes.

EM Robustness (EMR) studies the impact of circuit ageing on EMC [11]. Li, Wu et al. [74, 126] analysed the drift in the EM immunity of an MCU to **Electrical Fast Transients (EFTs)** interference after accelerated ageing. They aged the device in incremental steps by applying high voltage (150% above nominal) and high temperature. The test setup consisted of a probe connecting the device to an EFT signal generator, which injected an interference signal between each ageing step. To assess whether the device was affected by the signal, errors were monitored for the executed software. The authors reported failures, including self-recovering errors, soft errors (reset was needed) and complete damage. Failure rates increased consistently with the ageing time, as well as immunity degradation (observed as a decrement in the maximum tolerated interference voltage for some pins). No further analysis was performed relating internal ageing mechanisms under the observed results. Instead, authors suspected that (i) physical parameters of the functional components of the MCU drifted due to ageing, and (ii) components of the EM interference protection circuitry degraded, thus increasing susceptibility.

Discussion. Although related work reveals a clear impact of circuit degradation in the EMC and EMR of MCUs, the inner dynamics of this phenomenon are still to be explored. As MCUs are complex devices, it is challenging to assess the degradation of every subsystem and their impact on the EM spectrum signature. EM measurements have not been widely developed as ageing monitoring techniques. This may be due to requiring external equipment and limited understanding of the correlation between system degradation of circuits and the changes in EM characteristics. This gap opens an interesting research direction for studying correlations between other ageing indicators for MCUs and SOCs (e.g., SRAM startup patterns) and EMC, which could enable indirect measurements of the latter by means of cheaper and ubiquitous sensors that require no external equipment.

6 Ageing Detection on Power Supplies

All embedded systems require supplies to deliver power with stable parameters over the entire lifespan of an application. Due to their high efficiency (typically more than 90%), **Switching Mode Power Supplies (SMPS)** are a popular choice to provide a stable DC voltage. SMPS are commonly composed of a switching transistor, a diode, a component that accumulates energy (usually an inductor) and a capacitor completing an LC filter for a steady DC output. Depending on their application, they can be built in a step-up (boost) configuration —reducing current and increasing voltage —or in a step-down (buck) configuration (Figure 14) – increasing the delivered current by decreasing the voltage. The controlled switch continuously switches between the ON and OFF states, at which ideally there is no power dissipation. The output voltage is regulated by controlling the duty cycle of the switching signal.

LDO linear voltage regulators (Figure 15) can regulate their output even when the input voltage (power supply) is close to the output voltage. They are commonly found as stand-alone chips in many power delivery networks of embedded systems and integrated into ICs and SoCs. An LDO is

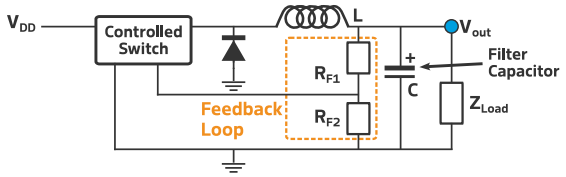

Fig. 14. Schematic of a switching-mode power supply in step-down (or buck) configuration; see also [20].

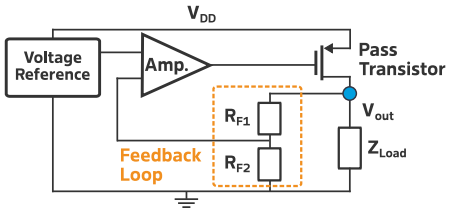

Fig. 15. Schematic of a generic low drop-out voltage regulator. Refer to [19] for background.

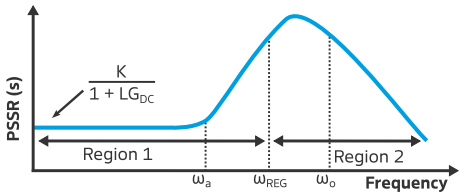

Fig. 16. Characteristic curve of LDO power supply rejection ratio versus frequency [27].

composed of an operational amplifier, a voltage reference, a pass transistor, and a resistive divider, which closes a feedback loop. The amplifier constantly compares the feedback loop voltage with the expected value of the reference. Based on this value, the amplifier controls the gate voltage of the transistor, which acts as a variable resistance and sets the feedback voltage (and, thus, the output voltage) to the desired value. Although simpler than SMPS, LDOs suffer from lower efficiency. As the voltage regulation is performed by the pass transistor in its ohmic region, it actively dissipates energy.

The most popular choices of power supplies in embedded appliances are built with components that undergo ageing as discussed in Section 2. In this section, we present several techniques used to detect and monitor the degradation of power supply performances and shifts in their parameters.

6.1 Power Supply Rejection Ratio

The **Power Supply Rejection Ratio (PSRR)** is an important characteristic of LDO regulators. It indicates its ability to prevent fluctuations in the output voltage in the presence of noise from the input voltage (or power supply voltage). The PSRR is the relation between the output voltage ripple v_o and the input voltage ripple v_i. It is normally measured for the whole operating frequency band of the regulator. A typical PSRR characteristic is illustrated in Figure 16; a lower value means a better rejection of the power supply noise. The PSRR in the frequency domain can be modelled as in Equation (10) [27]. The constant K depends on the feedback resistors, the load, and the internal drain-to-source resistance to the small signal of the transistor. The frequencies ω_a and ω_o are poles of the PSRR transference equation and depend on the resistances and capacitances of the transistor, the amplifier, and the load circuit. A_a and A_o are the loop gains of the amplifier and pass transistor,

respectively.

$$PSRR(s) = \frac{v_o(s)}{v_i(s)} = \frac{K}{\left(1 + \frac{s}{\omega_o}\right)(1 + LG(s))} \tag{10}$$

with the feedback loop transfer function,

$$LG(s) = \frac{A_a A_o}{\left(1 + \frac{s}{\omega_o}\right)\left(1 + \frac{s}{\omega_a}\right)}. \tag{11}$$

The PSSR curve (Figure 16) has two regions: Region 1 at low and middle frequencies and Region 2 at high frequencies. Both regions meet at the unity bandwidth frequency ω_{REG}, where the feedback loop gain is one. Region 1 is governed by the loop gain, which depends on the transconductance of the pass transistor and the amplifier. As this property is affected by transistor ageing mechanisms (e.g., BTI and HCI), the PSRR is impacted by circuit ageing. Region 2 is mainly impacted by parasitic capacitances of the transistor input and the load, whereas the feedback loop gain plays almost no role. HCI changes the gate capacitance of transistors [51], thus also modifying the PSRR in this region.

Considering the effect of circuit ageing on the PSRR, Chowdhury et al. [28] proposed using this parameter to estimate the ageing of LDOs as well as ICs, including such components in their power delivery networks, with the goal of determining whether a device has been recycled. To characterize the impact of ageing in the PSRR, the authors designed their own LDO chips in 65 nm technology. The LDOs were artificially aged at elevated temperature and power supply voltage. To analyse the PSRR of the devices and to record the output spectrum, a spectrum analyser and its tracking generator were employed. After applying AC and DC stress to the LDO, a degradation of the PSRR was observed as well as a shift in the transfer poles. This shows that both NBTI (DC stress) and HCI (AC stress) mechanisms could be detected by measuring the PSRR degradation.

Based on the clear influence of circuit ageing on the PSRR [28], the authors introduced a technique [26, 27] based on ML methods that utilizes this parameter to detect recycled analog and mixed-signal chips and SoCs, respectively. As a first step, an unsupervised **k-Nearest Neighbours (KNN)** model was trained and tested with PSRR data from new and artificially aged devices from one vendor. Next, the authors used a semi-supervised approach, in which they applied the previously trained models and tested them against other vendors in an attempt to avoid the need of golden data from all vendors. Both approaches showed promising results in discriminating used and new LDOs but required golden data and a supervised training phase. Finally, for unsupervised learning, two cases were tested. (i) PSRR data from an unknown device was given to the model together with golden data and, based on the number of clusters returned, the device could be classified as new or aged. (ii) The PSRRs of the unknown device were measured twice, before and after a synthetic ageing process. This last approach had the clear downside of being partially destructive and did not show good results.

Acharya et al. [2] presented and implemented an odometer based on a modified LDO circuit intended for the ageing-based detection of recycled ICs. A parallel feedback path to the regulator (i.e., a reference path) was added, which remained unused throughout the device lifetime. Two signals controlled which path is actively used. A one-label classifier model evaluated the LDO PSRR when using each path to detect degradation in the regulator components. Besides the external measurements, this method has the drawbacks of requiring a custom regulator design (i.e., cannot be applied on COTS components) and employing redundant hardware.

Discussion. Measuring PSRR changes has proven effective for detecting ageing of external LDO regulators. Although Chowdhury et al. [26, 28] claim that the proposed methods are applicable to

ICs, they were only tested on stand-alone LDOs. Measurements of embedded LDOs were solely presented in [27], but the applied models could not detect recycled chips. In addition, the reported results are based on artificially aged devices due to the lack of used chips. Given its direct relation to transistor parameters, PSRR is a sensible parameter to reveal ageing effects of BTI and HCI on LDOs. Thus far, the literature has only focused on short ageing periods, however; a few hours under accelerated ageing correspond to a few days of normal operation. To fully understand the correlation of PSRR with device ageing over longer periods of time, further research is required. Moreover, methods proposed so far heavily rely on external testing equipment, and no strategy has yet been presented to deploy testing in the field.

6.2 Electromagnetic Characteristics

Ageing not only affects the EM characteristics of microcontrollers and SoCs (see Section 5.3) but also those of LDOs [127] and SMPSs [17]. Indeed, as parameters of the voltage regulator shift away from their nominal values due to degradation, so do their EM emissions and susceptibility.

Boyer et al. [17] have studied how thermally triggered ageing mechanisms affect internal passive and switching components of SMPS in step-down or buck configuration, thus modifying its conducted and radiated EM emissions. To this end, they have artificially aged four samples of NCP3163 devices at elevated temperature. Although all samples remained operational after ageing, an average increase between 6 dB and 20 dB in the conducted emissions and 5 dB in the radiated emissions spectra were observed over a large frequency range. The aged inductor was replaced by a new component to determine its impact on the system degradation, which lowered EM emissions according to their initial values. By means of an impedance versus frequency characterization of the inductor before and after thermal stress, the authors encountered a reduction in its impedance and quality factor, which was attributed to an increase in the parasitic parallel capacitor caused by a higher core loss as described in Section 2.3.2.

Boyer et al. [16] also performed related experiments on synchronous buck converters based on the LT3800 controller. Measurements of conducted EM emissions were carried out inside a semi-anechoic chamber. Converter boards and various additional components were aged in an oven at high temperature for 2 weeks. The power iron inductor and aluminium capacitor were the most degraded components, followed by the onboard output capacitor and the power transistors. In addition, the conducted emission measured at the circuit output suffered an increment of up to 15 dB between 4 MHz and 100 MHz.

As analysed in detail by Wu et al. [127], EMI-induced offset is a common failure mode caused by EM noise coupled to the power pin of an LDO. Due to the nonlinearity of the differential transistor pair at the input of the comparator amplifier (see Figure 15), noise injected via the power network induces an offset of the output voltage via a rectification effect. A test bench was built with a signal generator coupled to the voltage source and an oscilloscope monitoring the output and reference voltage. Various devices were aged for a week by electrically stressing them to evaluate the relation between ageing mechanisms and EM susceptibility. The devices still functioned, but their EMI robustness decreased. This indicated that the stress process had indeed triggered the transistor ageing mechanisms in the operational amplifier of the LDO.

Discussion. Component degradation alters EM emissions and affects susceptibility of power supplies. The main culprits are ageing on passive components (via ESR changes on capacitors or quality factor reduction on inductors) and switching components (mainly due to BTI and HCI). Although EM characteristics are indicators of these ageing effects, measurements can be troublesome. EM-based ageing detection methods developed in the literature typically require special setups, procedures, and external equipment, which hinders deployment as an online monitoring technique. Moreover, the usage of expensive dedicated instruments contradicts many typical

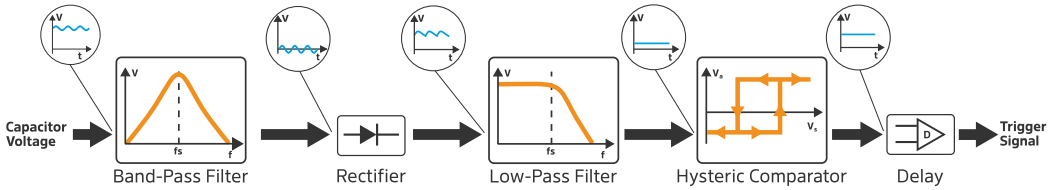

Fig. 17. Block diagram of voltage ripple measurement for electrolytic capacitor failure prediction proposed by [23].

requirements for COTS-based embedded applications, namely, low cost, low power usage, and small size.

6.3 Equivalent Series Resistance (ESR) of Capacitors

Capacitors are widely used as low-pass filters to clean noise from power lines (e.g., to suppress the switching noise introduced by SMPS as visualized Figure 14). Capacitor degradation jeopardizes the availability of the circuit. In fact, electrolytic capacitor failures cause more than 50% of SMPS breakdowns [68]. As digital components require a stable voltage, capacitors are attractive for degradation studies.

The impact of ageing mechanisms on the output voltage ripple of SMPS has been studied in [17]. Several DC-DC switching step-down circuits were subjected to accelerated ageing via increased temperature. The authors reported an increase between 100% and 200% in the output noise after ageing. It was determined that the main culprits for the decay in voltage stability were the filtering electrolytic capacitors at the output. Indeed, the high temperature accelerated the ageing mechanism described in Section 2.3.1. As a consequence, the equivalent serial resistance increased and the capacitance decayed.

As the ESR usually changes more than the capacitance for a given ageing, the former is preferred for monitoring purposes. To prevent failures, it is usually a good practice to replace a component whose ESR has increased by 100% or more [36].

Givi et al. [39] proposed and implemented an online concurrent monitoring system for DC-DC converters. Their technique detects ageing at the power supply output capacitor by indirectly measuring its ESR. The proposed setup required two voltage sensors at the converter: a Rogowski coil sensor on the inductor and a sensor of the output voltage. The authors implemented a monitoring system that derived the capacitor ESR on software from measurements of these two voltages using a DS1104 controller board. By comparing the resistance value with the initial one, the degradation of the component was estimated, and a warning signal could be triggered.

Lahyani et al. [68] argued that only the waveform of the output voltage ripple changes when capacitors degrade. They developed a software-based ESR monitoring method that sampled this ripple and applied a band-pass filter to the signal. They found that measuring the ripple at the power supply switching frequency is a more realistic method than using an average rectified value, and it reduces the load dependence.

On the other hand, Chen et al. [23] also measured ESR by observing the output voltage but proposed to filter out the switching frequency. Following a theoretical analysis, the authors concluded that the amplitude of the output voltage ripple could be accurately modelled as a direct, linear dependence on the ESR. Additionally, they showed that the output voltage ripple barely varies when the load current changes because the capacitor is part of an LC filter. From this, authors derived a method for failure prediction, which measured and processed the capacitor voltage, as illustrated in Figure 17. A band-pass filter eliminates the DC component and the power supply switching

frequency from the signal. After injecting the signal into a rectifier and a low-pass filter, a voltage proportional to the capacitor ESR can be obtained. When compared with a predefined threshold voltage, an early warning trigger signal can be generated. An additional delay circuit serves as mitigation against fake measurements during the circuit startup transient.

Discussion. The literature agrees that the output capacitor ESR serves as a good indicator to assess the overall health of a power supply, mainly due to the high impact of this component in the correct operation of the system. Additionally, its measurement is relatively simple. Given the ubiquity of analog-to-digital converters on most modern microcontrollers of embedded applications, it would be of interest to build in a concurrent online test that performed the signal processing completely digitally without extra circuitry. On the other hand, a drawback of this approach is the high susceptibility of ESR to operating conditions. Although it has been shown [23, 39, 68] that changes in the load do not affect the measurement, ESR still changes significantly with temperature and frequency. Systems also need to measure the environment and keep a record of the initial value under different conditions to account for these changes.

7 Discussion and Future Directions for Research

The body of literature reviewed in this work shows that state-of-the-art research on ageing monitoring predominantly studies degradation effects in a component-wide scope. In reality, components are not deployed to work in isolation. Embedded systems are complex and involve various interconnected, interdependent components. As complex systems, they exhibit composite effects and often present emergent dynamics. To achieve an understanding of system-wide health, further research is needed. Questions, such as 'How do ageing mechanisms of one component interplay with those of others?' demand an answer.

Following this system perspective, it is equally important to recognize the impact of low-level physical conditions on the overall system performance. It is necessary to enhance upper-layer logic with ageing monitoring insights in order to implement more reliable and autonomous systems. Work on degradation mitigation techniques at upper layers has started [61].

The literature represents extensive research towards the implementation and deployment of online tests for FPGAs. Indeed, their re-programmable nature makes them excellent candidates for synthesizing test resources, which can later be removed during operation. Ring oscillators are frequently employed as sensors for variability analysis and ageing monitoring, even on more modern FinFET-based devices. Nevertheless, techniques to assess or monitor degradation on microcontrollers, SoCs, and power supplies have merely focused on offline tests, mainly to understand the impact of ageing processes. To close this research gap, **the Built-In Self-Test (BIST)** and **Self-Built-In Self-Test (SBST)** for physical degradation effects on these components require further investigation, in particular with a focus on ageing monitors that are deployable online.

We envision that research on dynamically adaptive strategies (e.g., dynamic voltage [66] and frequency scaling [102]), augmented with degradation information, can help to advance ageing-aware embedded devices without hindering the overall system performance beyond what is strictly necessary. In an attempt to abstract developers from hardware degradation effects, vendors impose design-time guard bands. Such limits guaranty for embedded components that operate within specifications to account for normal degradation conditions. Nonetheless, in certain deployment environments, degradation may affect components more than expected (e.g., enhanced radiation [31, 69]). Operation outside vendor boundaries can reduce system reliability (e.g., higher clock frequencies, lower core voltage, or deployments longer than guaranteed lifetimes), but it can also boost application performance [57]. Monitoring the device degradation and testing the actual hardware limits may enable these types of operations. Indeed,

by continuously self-assessing the limitations of the components, an ageing-aware dynamic guard band could be determined at which the device still operates safely.

We foresee that adaptive hardware resource management techniques, both at component and system levels, can benefit from degradation monitoring systems. This would enable better informed decisions on resource allocation based on task characteristics and underlying hardware conditions.

We have observed a trend towards applying ML to enhance data produced by performance variability sensing techniques [26, 70, 76], particularly applying traditional models. In parallel, there is a growing research focus on ML-based **predictive maintenance (PdM)** [22]. Still, the field of ML is continuously evolving and introduces new methods, such as probabilistic time series forecasting to PdM [5]. This includes models that can analyse live online data and perform zero-shot inferences without historical data as well as perform online re-training. We consider that research on these ML models in combination with built-in online ageing monitoring systems has strong potential to further enhance the reliability and automation of embedded system deployments with PdM.

In this context, tinyML is under active research as a low-power edge alternative to cloud-based ML solutions [1]. This technology is already being applied in the domains of anomaly detection and PdM, but mostly for sensor data [83]. We envision that tinyML combined with on-device degradation monitoring can enable fully ageing-aware autonomous systems by allowing devices to collect and self-assess the results of built-in tests without requiring third parties or upstream connectivity.

A research program to address the identified gaps in the field could be guided along the following questions.

(1) Is it possible to reuse the working principles of offline tests on SoCs, microcontrollers, etc., to derive in situ online tests of physical device degradation?
(2) Can we identify lightweight, easily measurable indicators that faithfully represent the overall health of a system?
(3) Can we derive reliable system functions for health monitoring and self-assessment from these early indicators?
(4) Is there a versatile lifetime model that predicts system failures based on the indicators available from the self-assessment?
(5) How does one efficiently design a predictive maintenance architecture that can leverage low-level insights from built-in embedded sensing and ML inference based on historical data?

8 Conclusions

In this work, we reviewed the dominant degradation mechanisms that affect the basic components of embedded systems, which include a particular perspective on the trend of transistor miniaturization. We systematized methods for detecting hardware ageing of the elementary building blocks used in COTS embedded systems: FPGAs, SoCs, microcontrollers, and power supplies. This review was motivated by the increasing deployment of such embedded systems in varying contexts of criticality as well as in harsh operational environments. Many deployments are hard to access, either because they are widely distributed in the IoT or because the environmental conditions prevent accessibility (e.g., highly radiated tunnels of particle accelerators). Systems in these environments should operate autonomously and should also be able to monitor and diagnose themselves.

Degradation detection techniques are also useful to ensure component quality and identify counterfeit hardware. Recent chip shortages have increased quality concerns and raised the need for acceptance testing. Counterfeiting is not limited to parts manufactured by different entities; it is estimated that 80% of counterfeits are recycled [43] and potentially under-performing chips [4, 28, 47].

With this overview, we envision fostering future research and development towards self-aware embedded systems that not only can detect the degradation of individual hardware components but also are able to assess their overall health status and predict their remaining lifetime.

References

[1] Youssef Abadade, Anas Temouden, Hatim Bamoumen, Nabil Benamar, Yousra Chtouki, and Abdelhakim Senhaji Hafid. 2023. A comprehensive survey on TinyML. *IEEE Access* 11 (2023), 96892–96922. https://doi.org/10.1109/ACCESS.2023.3294111

[2] Rabin Y. Acharya, Michael Valentin Levin, and Domenic Forte. 2021. LDO-based odometer to combat IC recycling. In *2021 IEEE 34th International System-on-Chip Conference (SOCC'21)*. IEEE, Piscataway, NJ, 206–211. https://doi.org/10.1109/SOCC52499.2021.9739311

[3] Fikru Adamu-Lema, Vihar Georgiev, and Asen Asenov. 2019. Simulation of statistical NBTI degradation in 10nm doped channel pFinFETs. In *2019 International Conference on Simulation of Semiconductor Processes and Devices (SISPAD'19)*. IEEE, Piscataway, NJ, 1–4. https://doi.org/10.1109/SISPAD.2019.8870552

[4] Foisal Ahmed, Michihiro Shintani, and Michiko Inoue. 2021. Accurate recycled FPGA detection using an exhaustive-fingerprinting technique assisted by WID process variation modeling. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 40, 8 (2021), 1626–1639. https://doi.org/10.1109/TCAD.2020.3023684

[5] J. I. Aizpurua, B. G. Stewart, S. D. J. McArthur, M. Penalba, M. Barrenetxea, E. Muxika, and J. V. Ringwood. 2022. Probabilistic forecasting informed failure prognostics framework for improved RUL prediction under uncertainty: A transformer case study. *Reliability Engineering & System Safety* 226 (2022), 108676. https://doi.org/10.1016/j.ress.2022.108676

[6] Abdulazim Amouri, Florent Bruguier, Saman Kiamehr, Pascal Benoit, Lionel Torres, and Mehdi Tahoori. 2014. Aging effects in FPGAs: An experimental analysis. In *2014 24th International Conference on Field Programmable Logic and Applications (FPL'14)*. IEEE, Piscataway, NJ, 1–4. https://doi.org/10.1109/FPL.2014.6927390

[7] Abdulazim Amouri and Mehdi Tahoori. 2011. A low-cost sensor for aging and late transitions detection in modern FPGAs. In *2011 21st International Conference on Field Programmable Logic and Applications*. IEEE, Piscataway, NJ, 329–335. https://doi.org/10.1109/FPL.2011.66

[8] Emmanuel Baccelli, Cenk Gündogan, Oliver Hahm, Peter Kietzmann, Martine Lenders, Hauke Petersen, Kaspar Schleiser, Thomas C. Schmidt, and Matthias Wählisch. 2018. RIOT: An open source operating system for low-end embedded devices in the IoT. *IEEE Internet of Things Journal* 5, 6 (December 2018), 4428–4440.

[9] Mario Barbareschi, Giorgio Di Natale, and Lionel Torres. 2017. Implementation and analysis of ring oscillator circuits on Xilinx FPGAs. In *Hardware Security and Trust: Design and Deployment of Integrated Circuits in a Threatened Environment*, Nicolas Sklavos, Ricardo Chaves, Giorgio Di Natale, and Francesco Regazzoni (Eds.). Springer International Publishing, Cham, 237–251. https://doi.org/10.1007/978-3-319-44318-8_12

[10] Hugh J. Barnaby. 2006. Total-ionizing-dose effects in modern CMOS technologies. *IEEE Transactions on Nuclear Science* 53, 6 (2006), 3103–3121. https://doi.org/10.1109/TNS.2006.885952

[11] Sonia Ben Dhia, Alexandre Boyer, Li Binhong, and A. Ndoye. 2010. Characterisation of electromagnetic compatibility drifts of nanoscale integrated circuit after accelerated life tests. *Electronics Letters* 46 (03 2010), 278–280.

[12] Emmanuel Bender, Joseph B. Bernstein, and Alain Bensoussan. 2020. Reliability prediction of FinFET FPGAs by MTOL. *Microelectronics Reliability* 114 (2020), 113809. https://doi.org/10.1016/j.microrel.2020.113809

[13] James R. Black. 1969. Electromigration–A brief survey and some recent results. *IEEE Transactions on Electron Devices* 16, 4 (1969), 338–347.

[14] Stefano Bonaldo, Teng Ma, Serena Mattiazzo, Andrea Baschirotto, Christian Enz, Daniel M. Fleetwood, Alessandro Paccagnella, and Simone Gerardin. 2022. DC response, low-frequency noise, and TID-induced mechanisms in 16-nm FinFETs for high-energy physics experiments. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1033 (2022), 166727. https://doi.org/10.1016/j.nima.2022.166727

[15] Alexandre Boyer, Sonia Ben Dhia, Binhong Li, Néstor Berbel, and Raul Fernandez-Garcia. 2014. Experimental investigations into the effects of electrical stress on electromagnetic emission from integrated circuits. *IEEE Transactions on Electromagnetic Compatibility* 56, 1 (2014), 44–50.

[16] Alexandre Boyer, M. A. Gonzalez Sentis, C. Ghfiri, and A. Durier. 2017. Study of the thermal aging effect on the conducted emission of a synchronous buck converter. In *2017 11th International Workshop on the Electromagnetic Compatibility of Integrated Circuits (EMCCompo'17)*. IEEE, Piscataway, NJ, 79–84. https://doi.org/10.1109/EMCCompo.2017.7998086

[17] Alexandre Boyer, H. Huang, and Sonia Ben Dhia. 2014. Impact of thermal aging on emission of a buck DC-DC converter. In *2014 International Symposium on Electromagnetic Compatibility, Tokyo*. IEICE, Tokyo, Japan, 77–80. https://doi.org/10.34385/proc.18.13A2-A4

[18] Julien Branlard, Gohar Ayvazyan, Valeri Ayvazyan, Mariusz Grecki, Matthias Hoffmann, Tomasz Jeżyński, Frank Ludwig, Uros Mavrič, Sven Pfeiffer, Holger Schlarb, Christian Schmidt, Henning Weddig, Bin Yang, Paweł Barmuta, Samer Bou Habib, Łukasz Butkowski, Krzysztof Czuba, Maciej Grzegrzółka, Ewa Janas, Jan Piekarski, Igor Rutkowski, Dominik Sikora, Łukasz Zembala, Mateusz Żukociński, Wojciech Cichalewski, Wojciech Jałmużna, Dariusz Makowski, Aleksander Mielczarek, Andrzej Napieralski, Piotr Perek, Adam Piotrowski, Tomasz Poźniak, Konrad Przygoda, Grzegorz Bołtruczyk, Stefan Korolczuk, Maciej Kudła, Jarosław Szewiński, Krzysztof Oliwa, and Wojciech Wierba. 2013. MTCA.4 LLRF system for the European XFEL. In *Proceedings of the 20th International Conference Mixed Design of Integrated Circuits and Systems — MIXDES 2013*. IEEE, Piscataway, NJ, 109–112.

[19] Marty Brown. 2008. An introduction to the linear regulator. In *Power Sources and Supplies*. Newnes, Burlington, 1–12. https://doi.org/10.1016/B978-0-7506-8626-6.00001-6

[20] Martin C. Brown. 2012. *Practical Switching Power Supply Design*. Elsevier, New York.

[21] Florent Bruguier, Pascal Benoit, Philippe Maurine, and Lionel Torres. 2011. A new process characterization method for FPGAs based on electromagnetic analysis. In *2011 21st International Conference on Field Programmable Logic and Applications*. IEEE, Piscataway, NJ, 20–23. https://doi.org/10.1109/FPL.2011.15

[22] Thyago P. Carvalho, Fabrízzio A. A. M. N. Soares, Roberto Vita, Roberto da P. Francisco, João P. Basto, and Symone G. S. Alcalá. 2019. A systematic literature review of machine learning methods applied to predictive maintenance. *Computers & Industrial Engineering* 137 (2019), 106024. https://doi.org/10.1016/j.cie.2019.106024

[23] Yaow-Ming Chen, Ming-Wei Chou, and Hsu-Chin Wu. 2005. Electrolytic capacitor failure prediction of LC filter for switching-mode power converters. In *40th IAS Annual Meeting. Conference Record of the 2005 Industry Applications Conference, 2005*, Vol. 2. IEEE, Piscataway, NJ, 1464–1469. https://doi.org/10.1109/IAS.2005.1518552

[24] Vidya A. Chhabria and Sachin S. Sapatnekar. 2019. Impact of self-heating on performance and reliability in FinFET and GAAFET designs. In *20th International Symposium on Quality Electronic Design (ISQED'19)*. IEEE, Piscataway, NJ, 235–240. https://doi.org/10.1109/ISQED.2019.8697786

[25] Kihyun Choi, Hyun C. Sagong, Minjung Jin, Jiang Hai, Miji Lee, Taeyoung Jeong, Myung Soo Yeo, Hyewon Shim, Da Ahn, Wooyeon Kim, Yongjeung Kim, JuneKyun Park, Hwasung Rhee, and Euncheol Lee. 2020. Reliability on evolutionary FinFET CMOS technology and beyond. In *2020 IEEE International Electron Devices Meeting (IEDM'20)*. IEEE, Piscataway, NJ, 9.3.1–9.3.4. https://doi.org/10.1109/IEDM13553.2020.9371930

[26] Sreeja Chowdhury, Fatemeh Ganji, Troy Bryant, Nima Maghari, and Domenic Forte. 2019. Recycled analog and mixed signal chip detection at zero cost using LDO degradation. In *2019 IEEE International Test Conference (ITC'19)*. IEEE, Piscataway, NJ, 1–10. https://doi.org/10.1109/ITC44170.2019.9000118

[27] Sreeja Chowdhury, Fatemeh Ganji, and Domenic Forte. 2020. Recycled SoC detection using LDO degradation. *SN Computer Science* 1, 6 (26 Sep 2020), 312. https://doi.org/10.1007/s42979-020-00329-2

[28] Sreeja Chowdhury, Haoting Shen, Beomsoo Park, Nima Maghari, and Domenic Forte. 2019. Aging analysis of low dropout regulator for universal recycled IC detection. In *2019 IEEE Computer Society Annual Symposium on VLSI (ISVLSI'19)*. IEEE, Piscataway, NJ, 604–609. https://doi.org/10.1109/ISVLSI.2019.00113

[29] Fong-Min Ciou, Jia-Hong Lin, Po-Hsun Chen, Ting-Chang Chang, Kai-Chun Chang, Jui-Tse Hsu, Yu-Shan Lin, Fu-Yuan Jin, Wei-Chun Hung, Chien-Hung Yeh, Ting-Tzu Kuo, Osbert Cheng, Cheng-Tung Huang, and Yi-Han Ye. 2021. Comparison of the hot carrier degradation of N- and P-Type fin field-effect transistors in 14-nm technology nodes. *IEEE Electron Device Letters* 42, 10 (2021), 1420–1423. https://doi.org/10.1109/LED.2021.3106540

[30] J. F. Dawson, I. D. Flintoft, A. P. Duffy, A. C. Marvin, and M. P. Robinson. 2014. Effect of high temperature ageing on electromagnetic emissions from a PIC microcontroller. In *2014 International Symposium on Electromagnetic Compatibility*. IEEE, Piscataway, NJ, 1139–1143. https://doi.org/10.1109/EMCEurope.2014.6931074

[31] Zachary J. Diggins, Nagabhushan Mahadevan, Daniel Herbison, Gabor Karsai, Brian D. Sierawski, Eric Barth, E. Bryn Pitt, Robert A. Reed, Ronald D. Schrimpf, Robert. A. Weller, Michael L. Alles, and Arthur Witulski. 2014. Total-ionizing-dose induced timing window violations in CMOS microcontrollers. *IEEE Transactions on Nuclear Science* 61, 6 (2014), 2979–2984. https://doi.org/10.1109/TNS.2014.2368125

[32] Maxim Ershov, Supriya Saxena, Hossein Karbasi, S. Winters, and S. Minehane. 2003. Dynamic recovery of negative bias temperature instability in p-type metal-oxide-semiconductor field-effect transistors. *Applied Physics Letters* 83, 8 (2003), 1647–1649.

[33] Rasoul Faraji and Hamid Reza Naji. 2014. Adaptive technique for overcoming performance degradation due to aging on 6T SRAM cells. *IEEE Transactions on Device and Materials Reliability* 14, 4 (2014), 1031–1040.

[34] Tomas Fried, Antonio Di Buono, David Cheneler, Neil Cockbain, Jonathan M. Dodds, Peter R. Green, Barry Lennox, C. James Taylor, and Stephen D. Monk. 2021. Radiation testing of low cost, commercial off the shelf microcontroller board. *Nuclear Engineering and Technology* 53, 10 (2021), 3335–3343. https://doi.org/10.1016/j.net.2021.05.005

[35] Michael L. Gasperi. 1996. Life prediction model for aluminum electrolytic capacitors. In *IAS'96 Conference Record of the 1996 IEEE Industry Applications Conference 31st IAS Annual Meeting*, Vol. 3. IEEE, Piscataway, NJ, 1347–1351.

[36] Michael L. Gasperi. 1997. A method for predicting the expected life of bus capacitors. In *IAS'97 Conference Record of the 1997 IEEE Industry Applications Conference 32nd IAS Annual Meeting*, Vol. 2. IEEE, Piscataway, NJ, 1042–1047.

[37] Zana Ghaderi, Mohammad Ebrahimi, Zainalabedin Navabi, Eli Bozorgzadeh, and Nader Bagherzadeh. 2017. SENSIBle: A highly scalable SENsor DeSIgn for path-based age monitoring in FPGAs. *IEEE Trans. Comput.* 66, 5 (2017), 919–926. https://doi.org/10.1109/TC.2016.2622688

[38] A. Ghosh, S. Devadas, K. Keutzer, and J. White. 1992. Estimation of average switching activity in combinational and sequential circuits. In *[1992] Proceedings 29th ACM/IEEE Design Automation Conference*. IEEE, Piscataway, NJ, 253–259.

[39] Hadi Givi, Ebrahim Farjah, and Teymoor Ghanbari. 2019. A comprehensive monitoring system for online fault diagnosis and aging detection of non-isolated DC-DC converters' components. *IEEE Transactions on Power Electronics* 34, 7 (2019), 6858–6875. https://doi.org/10.1109/TPEL.2018.2875830

[40] Douglas Goodman, James Hofmeister, and Justin Judkins. 2007. Electronic prognostics for switched mode power supplies. *Microelectronics Reliability* 47, 12 (2007), 1902–1906.

[41] Mariia Gorchichko, En Xia Zhang, Pan Wang, Stefano Bonaldo, Ronald D. Schrimpf, Robert A. Reed, Dimitri Linten, Jerome Mitard, and Daniel M. Fleetwood. 2021. Total-ionizing-dose response of highly scaled gate-all-around Si nanowire CMOS transistors. *IEEE Transactions on Nuclear Science* 68, 5 (2021), 687–696. https://doi.org/10.1109/TNS.2021.3066612

[42] Tibor Grasser, Ben Kaczer, Wolfgang Goes, Hans Reisinger, Thomas Aichinger, Philipp Hehenberger, Paul-Jürgen Wagner, Franz Schanovsky, Jacopo Franco, María Toledano Luque, and Michael Nelhiebel. 2011. The paradigm shift in understanding the bias temperature instability: From reaction-diffusion to switching oxide traps. *IEEE Transactions on Electron Devices* 58, 11 (2011), 3652–3666.

[43] Ujjwal Guin, Daniel DiMase, and Mohammad Tehranipoor. 2014. Counterfeit integrated circuits: Detection, avoidance, and the challenges ahead. *Journal of Electronic Testing* 30, 1 (01 Feb 2014), 9–23. https://doi.org/10.1007/s10836-013-5430-8

[44] Ujjwal Guin, Wendong Wang, Charles Harper, and Adit D. Singh. 2019. Detecting recycled SoCs by exploiting aging induced biases in memory cells. In *2019 IEEE International Symposium on Hardware Oriented Security and Trust (HOST'19)*. IEEE, Piscataway, NJ, 72–80. https://doi.org/10.1109/HST.2019.8741032

[45] Xinfei Guo and Mircea R. Stan. 2020. *Future Directions in Self-healing*. Springer International Publishing, Cham, 193–197. https://doi.org/10.1007/978-3-030-20051-0_7

[46] Xinfei Guo, Vaibhav Verma, Patricia Gonzalez-Guerrero, and Mircea R. Stan. 2018. When "things" get older: Exploring circuit aging in IoT applications. In *2018 19th International Symposium on Quality Electronic Design (ISQED'18)*. IEEE, Piscataway, NJ, 296–301. https://doi.org/10.1109/ISQED.2018.8357304

[47] Zimu Guo, Md. Tauhidur Rahman, Mark M. Tehranipoor, and Domenic Forte. 2016. A zero-cost approach to detect recycled SoC chips using embedded SRAM. In *2016 IEEE International Symposium on Hardware Oriented Security and Trust (HOST'16)*. IEEE, Piscataway, NJ, 191–196. https://doi.org/10.1109/HST.2016.7495581

[48] Zimu Guo, Xiaolin Xu, Md. Tauhidur Rahman, Mark M. Tehranipoor, and Domenic Forte. 2018. SCARe: An SRAM-based countermeasure against IC recycling. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 26, 4 (2018), 744–755. https://doi.org/10.1109/TVLSI.2017.2777262

[49] Daniel E. Holcomb, Wayne P. Burleson, and Kevin Fu. 2009. Power-Up SRAM state as an identifying fingerprint and source of true random numbers. *IEEE Trans. Comput.* 58, 9 (2009), 1198–1210.

[50] Chenming Hu, Simon C. Tam, Fu-Chieh Hsu, Ping-Keung Ko, Tung-Yi Chan, and K. W. Terrill. 1985. Hot-Electron-Induced MOSFET degradation —model, monitor, and improvement. *IEEE Journal of Solid-State Circuits* 20, 1 (1985), 295–305.

[51] Yoonjong Huh, Yungkwon Sung, and S. M. Kang. 1998. A study of hot-carrier-induced mismatch drift: A reliability issue for VLSI circuits. *IEEE Journal of Solid-State Circuits* 33, 6 (1998), 921–927.

[52] IEEE. 2017. IEEE standard framework for prognostics and health management of electronic systems. *IEEE Std 1856-2017* 1 (2017), 1–31. https://doi.org/10.1109/IEEESTD.2017.8227036

[53] Kjell O. Jeppson and Christer M. Svensson. 1977. Negative bias stress of MOS devices at high electric fields and degradation of MNOS devices. *Journal of Applied Physics* 48, 5 (1977), 2004–2014.

[54] Weixiong Jiang, Rui Li, Heng Yu, and Yajun Ha. 2020. An accurate FPGA online delay monitor supporting all timing paths. In *2020 IEEE International Symposium on Circuits and Systems (ISCAS'20)*. IEEE, Piscataway, NJ, 1–5. https://doi.org/10.1109/ISCAS45731.2020.9181070

[55] Minjung Jin, Changze Liu, Jinju Kim, Jungin Kim, Seungjin Choo, Yoohwan Kim, Hyewon Shim, Lijie Zhang, Kab-jin Nam, Jongwoo Park, Sangwoo Pae, and Haebum Lee. 2016. Hot carrier reliability characterization in consideration of self-heating in FinFET technology. In *2016 IEEE International Reliability Physics Symposium (IRPS'16)*. IEEE, Piscataway, NJ, 2A–2–1–2A–2–5. https://doi.org/10.1109/IRPS.2016.7574505

[56] Leonardo Rezende Juracy, Matheus Trevisan Moreira, Alexandre de Morais Amory, and Fernando Gehm Moraes. 2020. A Survey of Aging Monitors and Reconfiguration Techniques. arXiv:2007.07829 [cs.AR] https://arxiv.org/abs/2007.07829

[57] Anmol Kaushik, Shivaprasad Chumbalakar, Surya Musunuri, and Anju Pillai. 2022. *Evaluation of Dynamic Frequency Control on an Automotive Microcontroller*. Springer International Publishing, Cham, 313–327. https://doi.org/10.1007/978-981-16-8862-1_21

[58] Usman Khalid, Antonio Mastrandrea, and Mauro Olivieri. 2015. Effect of NBTI/PBTI aging and process variations on write failures in MOSFET and FinFET flip-flops. *Microelectronics Reliability* 55, 12, Part B (2015), 2614–2626. https://doi.org/10.1016/j.microrel.2015.07.050

[59] Seyab Khan, Said Hamdioui, Halil Kukner, Praveen Raghavan, and Francky Catthoor. 2012. BTI impact on logical gates in nano-scale CMOS technology. In *2012 IEEE 15th International Symposium on Design and Diagnostics of Electronic Circuits & Systems (DDECS'12)*. IEEE, Piscataway, NJ, 348–353.

[60] Navid Khoshavi, Rizwan A. Ashraf, Ronald F. DeMara, Saman Kiamehr, Fabian Oboril, and Mehdi B. Tahoori. 2017. Contemporary CMOS aging mitigation techniques: Survey, taxonomy, and methods. *Integration* 59 (2017), 10–22. https://doi.org/10.1016/j.vlsi.2017.03.013

[61] Navid Khoshavi, Rizwan A. Ashraf, Ronald F. DeMara, Saman Kiamehr, Fabian Oboril, and Mehdi B. Tahoori. 2017. Contemporary CMOS aging mitigation techniques: Survey, taxonomy, and methods. *Integration* 59 (2017), 10–22. https://doi.org/10.1016/j.vlsi.2017.03.013

[62] Hyunjin Kim, Minjung Jin, Hyunchul Sagong, Jinju Kim, Ukjin Jung, Minhyuck Choi, Junekyun Park, Sangchul Shin, and Sangwoo Pae. 2018. A systematic study of gate dielectric TDDB in FinFET technology. In *2018 IEEE International Reliability Physics Symposium (IRPS'18)*. IEEE, Piscataway, NJ, 4A.4–1–4A.4–4. https://doi.org/10.1109/IRPS.2018.8353577

[63] Jongsu Kim, Kyushik Hong, Hyewon Shim, HwaSung Rhee, and Hyungcheol Shin. 2020. Comparative analysis of hot carrier degradation (HCD) in 10-nm node nMOS/pMOS FinFET devices. *IEEE Transactions on Electron Devices* 67, 12 (2020), 5396–5402. https://doi.org/10.1109/TED.2020.3031246

[64] R. Kingsbury, F. Schmidt, W. Blackwell, I. Osarentin, R. Legge, K. Cahoy, and D. Sklair. 2013. TID tolerance of popular CubeSat components. In *2013 IEEE Radiation Effects Data Workshop (REDW'13)*. IEEE, Piscataway, NJ, 1–4. https://doi.org/10.1109/REDW.2013.6658220

[65] Michael A. Kochte and Hans-Joachim Wunderlich. 2018. Self-test and diagnosis for self-aware systems. *IEEE Design & Test* 35, 5 (2018), 7–18.

[66] Ulf Kulau, Felix Büsching, and Lars Wolf. 2016. IdealVolting: Reliable undervolting on wireless sensor nodes. *ACM Transactions on Sensor Networks (TOSN)* 12, 2 (2016), 1–38.

[67] Chetan S. Kulkarni, José R. Celaya, Gautam Biswas, and Kai Goebel. 2012. Accelerated aging experiments for capacitor health monitoring and prognostics. In *2012 IEEE AUTOTESTCON Proceedings*. IEEE, Piscataway, NJ, 356–361. https://doi.org/10.1109/AUTEST.2012.6334580

[68] A. Lahyani, P. Venet, G. Grellet, and P.-J. Viverge. 1998. Failure prediction of electrolytic capacitors during operation of a switchmode power supply. *IEEE Transactions on Power Electronics* 13, 6 (1998), 1199–1207. https://doi.org/10.1109/63.728347

[69] Leandro Lanzieri, Lukasz Butkowski, Jiri Kral, Goerschwin Fey, Holger Schlarb, and Thomas C. Schmidt. 2024. Studying the degradation of propagation delay on FPGAs at the European XFEL. In *Proceedings of 27th Euromicro Conference on Digital System Design (DSD'24)*, Paris, France. IEEE, Piscataway, NJ.

[70] Leandro Lanzieri, Peter Kietzmann, Goerschwin Fey, Holger Schlarb, and Thomas C. Schmidt. 2023. Ageing analysis of embedded SRAM on a large-scale testbed using machine learning. In *Proceedings of 26th Euromicro Conference on Digital System Design (DSD'23)*, Durres, AL. IEEE, Piscataway, NJ, 335–342. https://doi.org/10.1109/DSD60849.2023.00054

[71] Kyong T. Lee, Wonchang Kang, Eun-Ae Chung, Gunrae Kim, Hyewon Shim, Hyunwoo Lee, Hyejin Kim, Minhyeok Choe, Nae-In Lee, Anuj Patel, Junekyun Park, and Jongwoo Park. 2013. Technology scaling on High-K & Metal-Gate FinFET BTI reliability. In *2013 IEEE International Reliability Physics Symposium (IRPS'13)*. IEEE, Piscataway, NJ, 2D.1.1–2D.1.4. https://doi.org/10.1109/IRPS.2013.6531956

[72] Carlos Leong, Jorge Semião, Isabel M. C. Teixeira, Marcelino B. Santos, João P. C. Teixeira, María Valdes, Judit Freijedo, Juan J. Rodríguez-Andina, and Fabian Vargas. 2013. Aging monitoring with local sensors in FPGA-based designs. In *2013 23rd International Conference on Field Programmable Logic and Applications*. IEEE, Piscataway, NJ, 1–4. https://doi.org/10.1109/FPL.2013.6645596

[73] Alexander Lex, Nils Gehlenborg, Hendrik Strobelt, Romain Vuillemot, and Hanspeter Pfister. 2014. UpSet: Visualization of intersecting sets. *IEEE Transactions on Visualization and Computer Graphics (InfoVis'14)* 20, 12 (2014), 1983–1992. https://doi.org/10.1109/TVCG.2014.2346248

[74] Chuangwei Li, Jianfei Wu, Yan Huang, and Wei Zhu. 2016. Characterization of change in microcontroller susceptibility during accelerated aging. In *2016 Asia-Pacific International Symposium on Electromagnetic Compatibility (APEMC'16)*, Vol. 01. IEEE, Piscataway, NJ, 751–754. https://doi.org/10.1109/APEMC.2016.7522856

[75] Jie Li and John Lach. 2007. Negative-skewed shadow registers for at-speed delay variation characterization. In *2007 25th International Conference on Computer Design*. IEEE, Piscataway, NJ, 354–359. https://doi.org/10.1109/ICCD.2007.4601924

[76] Zeyu Li, Zhao Huang, Quan Wang, Junjie Wang, and Nan Luo. 2022. Implementation of aging mechanism analysis and prediction for XILINX 7-Series FPGAs with a 28-nm process. *Sensors* 22, 12 (2022), 4439. https://doi.org/10.3390/s22124439

[77] Teng Ma, Stefano Bonaldo, Serena Mattiazzo, Andrea Baschirotto, Christian Enz, Alessandro Paccagnella, and Simone Gerardin. 2022. Influence of fin and finger number on TID degradation of 16-nm Bulk FinFETs irradiated to ultrahigh doses. *IEEE Transactions on Nuclear Science* 69, 3 (2022), 307–313. https://doi.org/10.1109/TNS.2021.3125769

[78] Souvik Mahapatra and Narendra Parihar. 2018. A review of NBTI mechanisms and models. *Microelectronics Reliability* 81 (2018), 127–135. https://doi.org/10.1016/j.microrel.2017.12.027

[79] Konstantinos Maragos, Endri Taka, George Lentaris, Ioannis Stratakos, and Dimitrios Soudris. 2019. Analysis of performance variation in 16nm FinFET FPGA devices. In *2019 29th International Conference on Field Programmable Logic and Applications (FPL'19)*. IEEE, Piscataway, NJ, 38–44. https://doi.org/10.1109/FPL.2019.00016

[80] Gianluca Martino, Arne Gruenhagen, Julien Branlard, Annika Eichler, Goerschwin Fey, and Holger Schlarb. 2021. Comparative evaluation of semi-supervised anomaly detection algorithms on high-integrity digital systems. In *2021 24th Euromicro Conference on Digital System Design (DSD'21)*. IEEE, Piscataway, NJ, 123–130.

[81] Peter Marwedel. 2021. *Embedded System Design: Embedded Systems Foundations of Cyber-physical Systems, and the Internet of Things*. Springer Nature, Cham, Switzerland.

[82] Gordon E. Moore. 1998. Cramming more components onto integrated circuits. *Proc. IEEE* 86, 1 (1998), 82–85.

[83] Alireza Mostafavi and Ali Sadighi. 2021. A novel online machine learning approach for real-time condition monitoring of rotating machines. In *9th RSI International Conference on Robotics and Mechatronics (ICRoM'21)*. IEEE, Piscataway, NJ, 267–273. https://doi.org/10.1109/ICRoM54204.2021.9663495

[84] Matthew Nancekievill, Simon Watson, Peter R. Green, and Barry Lennox. 2016. Radiation tolerance of commercial-off-the-shelf components deployed in an underground nuclear decommissioning embedded system. In *2016 IEEE Radiation Effects Data Workshop (REDW'16)*. IEEE, Piscataway, NJ, 1–5. https://doi.org/10.1109/NSREC.2016.7891730

[85] Mohammad Naouss and François Marc. 2015. Design and implementation of a low cost test bench to assess the reliability of FPGA. *Microelectronics Reliability* 55, 9 (2015), 1341–1345. https://doi.org/10.1016/j.microrel.2015.06.087

[86] Mohammad Naouss and François Marc. 2016. FPGA LUT delay degradation due to HCI: Experiment and simulation results. *Microelectronics Reliability* 64 (2016), 31–35. https://doi.org/10.1016/j.microrel.2016.07.048

[87] Mohammad Naouss and François Marc. 2016. Modelling delay degradation due to NBTI in FPGA Look-up tables. In *2016 26th International Conference on Field Programmable Logic and Applications (FPL'16)*. IEEE, Piscataway, NJ, 1–4. https://doi.org/10.1109/FPL.2016.7577328

[88] Tanya Nigam, K. Y. Yiang, and Amit Marathe. 2017. Moore's Law: Technology scaling and reliability challenges. In *Microelectronics to Nanoelectronics*. CRC Press, Boca Raton, FL.

[89] T. H. Ning, P. W. Cook, R. H. Dennard, C. M. Osburn, S. E. Schuster, and H. N. Yu. 1979. 1-micron MOSFET VLSI technology. IV - Hot-electron design constraints. *IEEE Journal of Solid-State Circuits* 14 (April 1979), 268–275.

[90] Shigeo Ogawa and Noboru Shiono. 1995. Generalized diffusion-reaction model for the low-field charge-buildup instability at the Si-SiO_2 interface. *Phys. Rev. B* 51 (Feb 1995), 4218–4230. Issue 7.

[91] Timothy R. Oldham. 1984. Analysis of damage in MOS devices for several radiation environments. *IEEE Transactions on Nuclear Science* 31, 6 (1984), 1236–1241. https://doi.org/10.1109/TNS.1984.4333489

[92] Narendra Parihar, R. Anandkrishnan, Ankush Chaudhary, and Souvik Mahapatra. 2019. A comparative analysis of NBTI variability and TDDS in GF HKMG planar p-MOSFETs and RMG HKMG p-FinFETs. *IEEE Transactions on Electron Devices* 66, 8 (2019), 3273–3278. https://doi.org/10.1109/TED.2019.2920666

[93] Petr Pfeifer, Ben Kaczer, and Zdenek Pliva. 2014. A reliability lab-on-chip using programmable arrays. In *2014 IEEE International Reliability Physics Symposium*. IEEE, Piscataway, NJ, CA.6.1–CA.6.8. https://doi.org/10.1109/IRPS.2014.6861123

[94] Petr Pfeifer and Zdenek Pliva. 2012. Delay-fault run-time XOR-less aging detection unit using BRAM in modern FPGAs. In *2012 13th Biennial Baltic Electronics Conference*. IEEE, Piscataway, NJ, 81–84. https://doi.org/10.1109/BEC.2012.6376820

[95] Petr Pfeifer and Zdenek Pliva. 2012. On measurement of impact of the metallization and FPGA design to the changes of slice parameters and generation of delay faults. In *22nd International Conference on Field Programmable Logic and Applications (FPL'12)*. IEEE, Piscataway, NJ, 743–746. https://doi.org/10.1109/FPL.2012.6339167

[96] Petr Pfeifer and Zdenek Pliva. 2013. On measurement of parameters of programmable microelectronic nanostructures under accelerating extreme conditions (Xilinx 28nm XC7Z020 Zynq FPGA). In *2013 23rd International Conference on Field Programmable Logic and Applications*. IEEE, Piscataway, NJ, 1–4. https://doi.org/10.1109/FPL.2013.6645584

[97] Karl H. Pflueger. 1997. Power-supply reliability: A practical improvement guide. *EDN* 42, 5 (1997), 151–154.

[98] S. Rahimipour, W. N. Flayyih, I. El-Azhary, S. Shafie, and F. Z. Rokhani. 2012. A survey of on-chip monitors. In *2012 IEEE International Conference on Circuits and Systems (ICCAS'12)*. IEEE, Piscataway, NJ, 243–248. https://doi.org/10.1109/ICCircuitsAndSystems.2012.6408286

[99] S. Rangan, Neal R. Mielke, and E. C. C. Yeh. 2003. Universal recovery behavior of negative bias temperature instability [PMOSFETs]. In *IEEE International Electron Devices Meeting 2003*. IEEE, Piscataway, NJ, 14.3.1–14.3.4.

[100] R. Ranjan, Y. Liu, T. Nigam, A. Kerber, and B. Parameshwaran. 2017. Impact of AC voltage stress on core NMOSFETs TDDB in FinFET and planar technologies. In *2017 IEEE International Reliability Physics Symposium (IRPS'17)*. IEEE, Piscataway, NJ, DG–10.1–DG–10.5. https://doi.org/10.1109/IRPS.2017.7936367

[101] Neetu Rathi, Anil Kumar, Neeraj Gupta, and Sanjay Kumar Singh. 2023. A review of low-power static random access memory (SRAM) designs. In *2023 IEEE Devices for Integrated Circuit (DevIC'23)*. IEEE, Piscataway, NJ, 455–459. https://doi.org/10.1109/DevIC57758.2023.10134887

[102] Michel Rottleuthner, Thomas C. Schmidt, and Matthias Wählisch. 2022. Dynamic clock reconfiguration for the constrained IoT and its application to energy-efficient networking. In *International Conference on Embedded Wireless Systems and Networks (EWSN'22)* (Linz, AT). ACM, New York, NY.

[103] Mattia Ruffoni and Alessandro Bogliolo. 2002. Direct measures of path delays on commercial FPGA chips. In *Proceedings: 6th IEEE Workshop on Signal Propagation on Interconnects*. IEEE, Piscataway, NJ, 157–159. https://doi.org/10.1109/SPI.2002.258304

[104] Nobuyuki Sano, Kazuya Matsuzawa, Mikio Mukai, and Noriaki Nakayama. 2002. On discrete random dopant modeling in drift-diffusion simulations: Physical meaning of atomistic dopants. *Microelectronics Reliability* 42, 2 (2002), 189–199. https://doi.org/10.1016/S0026-2714(01)00138-X

[105] Garrett J. Schlenvogt, Hugh J. Barnaby, Jeff Wilkinson, Scott Morrison, and Larry Tyler. 2013. Simulation of TID effects in a high voltage ring oscillator. *IEEE Transactions on Nuclear Science* 60, 6 (2013), 4547–4554. https://doi.org/10.1109/TNS.2013.2284636

[106] Dieter K. Schroder. 2007. Negative bias temperature instability: What do we understand? *Microelectronics Reliability* 47, 6 (2007), 841–852.

[107] Pete Sedcole and Peter Y. K. K. Cheung. 2006. Within-die delay variability in 90nm FPGAs and beyond. In *2006 IEEE International Conference on Field Programmable Technology*. IEEE, Piscataway, NJ, 97–104. https://doi.org/10.1109/FPT.2006.270300

[108] Marty R. Shaneyfelt, James R. Schwank, Daniel M. Fleetwood, Peter S. Winokur, K. L. Hughes, and Fred W. Sexton. 1990. Field dependence of interface-trap buildup in polysilicon and metal gate MOS devices. *IEEE Transactions on Nuclear Science* 37, 6 (1990), 1632–1640. https://doi.org/10.1109/23.101171

[109] Charu Sharma, Puspalata Rajesh, R. P. Behera, and S. Amirthapandian. 2022. Impact of gamma radiation on 8051 microcontroller performance. *Nuclear Engineering and Technology* 54, 12 (2022), 4422–4430. https://doi.org/10.1016/j.net.2022.08.021

[110] David J. Smith. 2005. 2 - Understanding terms and jargon. In *Reliability, Maintainability and Risk (7th Edition)*, David J. Smith (Ed.). Butterworth-Heinemann, Oxford, 11–23.

[111] Justin Sobas and François Marc. 2024. Degradation measurement and modelling under ageing in a 16 nm FinFET FPGA. *Micromachines* 15, 1 (2024), 19. https://doi.org/10.3390/mi15010019

[112] James H. Stathis, Miaomiao Wang, and Kai Zhao. 2010. Reliability of advanced high-k/metal-gate n-FET devices. *Microelectronics Reliability* 50, 9 (2010), 1199–1202.

[113] James H. Stathis and Sufi Zafar. 2006. The negative bias temperature instability in MOS devices: A review. *Microelectronics Reliability* 46, 2 (2006), 270–286.

[114] Edward Stott, Justin S. J. Wong, and Peter Y. K. Cheung. 2010. Degradation analysis and mitigation in FPGAs. In *2010 International Conference on Field Programmable Logic and Applications*. IEEE, Piscataway, NJ, 428–433. https://doi.org/10.1109/FPL.2010.88

[115] Edward A. Stott, Justin S. J. Wong, Pete Sedcole, and Peter Y. K. Cheung. 2010. Degradation in FPGAs: Measurement and modelling. In *Proceedings of the 18th Annual ACM/SIGDA International Symposium on Field Programmable Gate Arrays* (Monterey, CA) *(FPGA'10)*. ACM, New York, NY, 229–238. https://doi.org/10.1145/1723112.1723152

[116] Alvin W. Strong, Ernest Y. Wu, Rolf-Peter Vollertsen, Jordi Suñe, Giuseppe La Rosa, Stewart E. Rauch, and Timothy D. Sullivan. 2009. Hot carriers. In *Reliability Wearout Mechanisms in Advanced CMOS Technologies*. John Wiley & Sons, Ltd, Hoboken, NJ, Chapter 5, 441–516.

[117] Eiji Takeda and Norio Suzuki. 1983. An empirical model for device degradation due to hot-carrier injection. *IEEE Electron Device Letters* 4, 4 (1983), 111–113.

[118] Carl V. Thompson and James R. Lloyd. 1993. Electromigration and IC interconnects. *MRS Bulletin* 18, 12 (01 Dec 1993), 19–25.

[119] María D. Valdes, Judit Freijedo, María J. Moure Rodríguez, Juan J. Rodríguez-Andina, Jorge Semião, Isabel M. C. Teixeira, João P. C. Teixeira, and Fabian Vargas. 2013. Design and validation of configurable online aging sensors in nanometer-scale FPGAs. *IEEE Transactions on Nanotechnology* 12, 4 (2013), 508–517. https://doi.org/10.1109/TNANO.2013.2253795

[120] María D. Valdes, Judit Freijedo, María J. Moure Rodríguez, Juan J. Rodríguez-Andina, Jorge Semião, Fabian Vargas, Isabel M. C. Teixeira, and João P. C. Teixeira. 2011. Programmable sensor for on-line checking of signal integrity in FPGA-based systems subject to aging effects. In *2011 12th Latin American Test Workshop (LATW'11)*. IEEE, Piscataway, NJ, 1–7. https://doi.org/10.1109/LATW.2011.5985926

[121] Anabela Veloso, Trong Huynh-Bao, Philippe Matagne, Doyoung Jang, Naoto Horiguchi, Julien Ryckaert, and Dan Mocuta. 2019. Nanowire & nanosheet FETs for ultra-scaled, high-density logic and memory applications. In *2019 Joint International EUROSOI Workshop and International Conference on Ultimate Integration on Silicon (EUROSOI-ULIS'19)*. IEEE, Piscataway, NJ, 1–4. https://doi.org/10.1109/EUROSOI-ULIS45800.2019.9041857

[122] Kevin Weiss, Michel Rottleuthner, Thomas C. Schmidt, and Matthias Wählisch. 2021. PHiLIP on the HiL: Automated multi-platform OS testing with external reference devices. *ACM Transactions on Embedded Computing Systems (TECS)* 20, 5s (September 2021), 91:1–91:26. https://doi.org/10.1145/3477040

[123] Justin S. J. Wong, Pete Sedcole, and Peter Y. K. Cheung. 2007. Self-characterization of combinatorial circuit delays in FPGAs. In *2007 International Conference on Field-Programmable Technology*. IEEE, Piscataway, NJ, 17–23. https://doi.org/10.1109/FPT.2007.4439227

[124] Justin S. J. Wong, Pete Sedcole, and Peter Y. K. Cheung. 2008. A transition probability based delay measurement method for arbitrary circuits on FPGAs. In *2008 International Conference on Field-Programmable Technology*. IEEE, Piscataway, NJ, 105–112. https://doi.org/10.1109/FPT.2008.4762372

[125] Justin S. J. Wong, Pete Sedcole, and Peter Y. K. Cheung. 2009. Self-measurement of combinatorial circuit delays in FPGAs. *ACM Trans. Reconfigurable Technol. Syst.* 2, 2, Article 10 (Jun 2009), 22 pages. https://doi.org/10.1145/1534916.1534920

[126] Jianjei Wu, Chuangwei Li, Binhong Li, Wei Zhu, and Hong Y. Wang. 2016. Microcontroller susceptibility variations to EFT burst during accelerated aging. *Microelectronics Reliability* 64 (2016), 210–214. https://doi.org/10.1016/j.microrel.2016.07.099

[127] J. Wu, J. Li, Rongjun Shen, Alexandre Boyer, and Sonia Ben Dhia. 2013. Effect of electrical stresses on the susceptibility of a voltage regulator. In *2013 International Symposium on Electromagnetic Compatibility*. IEEE, Piscataway, NJ, 759–764.

[128] Ying Zeng, Yan-Feng Li, Xiang-Yu Li, and Hong-Zhong Huang. 2019. Tolerance-based reliability and optimization design of switched-mode power supply. *Quality and Reliability Engineering International* 35, 8 (2019), 2774–2784.

[129] Kenneth M. Zick and John P. Hayes. 2010. On-line sensing for healthier FPGA systems. In *Proceedings of the 18th Annual ACM/SIGDA International Symposium on Field Programmable Gate Arrays* (Monterey, California) *(FPGA'10)*. ACM, New York, NY, 239–248. https://doi.org/10.1145/1723112.1723153

Received 7 March 2023; revised 27 August 2024; accepted 30 August 2024