SN Computer Science (2020) 1:312
https://doi.org/10.1007/s42979-020-00329-2

ORIGINAL RESEARCH

Recycled SoC Detection Using LDO Degradation

Sreeja Chowdhury[1,2] · **Fatemeh Ganji**[1,3] · **Domenic Forte**[1]

Received: 1 June 2020 / Accepted: 10 September 2020 / Published online: 26 September 2020

Abstract
Counterfeit electronics form a major roadblock towards a safe and successful economy. An increase in globalization has led to a major increase in the total number of counterfeit products all around the world. While several methods have been designed to detect counterfeits, very few of them have been applied to the system-on-chip (SoC). The influx of a variety of components in SoCs and the conglomeration of different types of properties makes it difficult to detect counterfeit SoCs. In this paper, we aim at detecting recycled counterfeit SoCs by evaluating the degradation of power supply rejection ratio (PSRR) of a low drop-out (LDO) regulator, a principal component of the power supply of the SoC. Since the power supply is a universal component in all SoCs, this method can be considered effective for most SoCs. We apply machine learning (ML) algorithms pertaining to the family of Gaussian mixture models to classify SoCs as recycled or new. Supervised and unsupervised ML algorithms show an accuracy of up to 90% and 74% of recycled detection. We also apply stand-alone LDO PSRR degradation to train the ML algorithm and test on PSRR from embedded LDOs in SoCs. This form of semi-supervised ML performed well for our previous experiments of recycled detection with stand-alone LDOs but was not able to distinguish recycled SoCs from new SoCs, thus increasing the number of false detection.

Keywords Recycled counterfeits · SoC recycled detection · LDO · PSRR · machine learning ML · bayesian ML

Introduction

Counterfeit electronics constitute a significant threat to the global supply chain and jeopardize the root of trust in consumer, military as well as other forms of electronics. A counterfeit electronic component is defined as an electronic part that is (1) an unauthorized copy, (2) does not conform to original component manufacturer's design, model, or performance, (3) is not produced by the original component manufacturer or is produced by unauthorized contractors, (4) is an off-specification, defective, or used original component manufacturer's product sold as "new" or working, or (5) has incorrect or false markings and/or documentation. The taxonomy of counterfeit integrated circuits (ICs) has been described in detail in Ref. [9]. Among the different counterfeit types, Recycled and remarked counterfeits comprise more than 80% of reported counterfeits [8]. Recycled counterfeits are used ICs that are harvested from discarded printed circuit boards (PCBs) and sold as new to consumers without their knowledge.

Several methods have been described in the literature to detect and prevent counterfeit electronics. These methods can be broadly divided into three major categories: (1) Hardware security primitives: This is mostly applicable to new chip designs, where additional security primitives are designed along with existing chip architecture to detect cloned or recycled counterfeits. For cloned counterfeit detection, silicon fingerprints called physical unclonable functions (PUFs) are developed [15]. For recycled detection, odometer or aging sensors called Combatting Die and IC Recycling (CDIRs) components are designed as described

This article is part of the topical collection "Hardware-Assisted Security Solutions for Electronic Systems" guest edited by Himanshu Thapliyal, Saraju P. Mohanty, Wujie Wen and Yiran Chen.

✉ Sreeja Chowdhury
sreeja.chowdhury@ansys.com

Fatemeh Ganji
fganji@wpi.edu

Domenic Forte
dforte@ece.ufl.edu

[1] University of Florida, Gainesville, USA

[2] Present Address: Ansys Inc, San Jose, CA, USA

[3] Present Address: Worcester Polytechnic Institute, Worcester, MA, USA

in Ref. [24]; (2) Targeted electrical testing: This is applicable to detect counterfeit legacy and current ICs: In this case, the addition of new circuits is not an option; thus, these detection methods focus on general electrical tests. These targeted electrical tests evaluate the performance of the ICs and compare the performance with respect to specification sheets or golden data. Most of these approaches (e.g., [25]) used to detect recycled ICs/FPGAs require data from known authentic chips (i.e., golden data), which is often unavailable; thus, it is a drawback. Another essential drawback is that most of these procedures are not automated and applicable to all types of ICs. (3) Physical inspection: These procedures can detect new, active as well as legacy ICs, but require expensive imaging facilities or expert technical guidance to detect discrepancies between the suspect and golden samples. The availability of golden samples is mandatory for most of these procedures. Advanced methods include high-tech imaging procedures involving X-ray tomography, scanning electron microscopy, etc. which are used to detect counterfeits [11]. Due to the wide variety of counterfeit components and their respective parameters, it is difficult to formulate a universal testing technique for detecting recycled counterfeit ICs. In our attempt to devise a ubiquitous, automated method to detect recycled ICs, we have focused on the IC's power delivery network (PDN). Evaluating degradation in electrical properties like PSRR of common PDN elements like low-dropout (LDO) regulators, we have previously investigated a universal strategy to detect recycled counterfeits as described in Ref. [3]. At first, we observed that LDO PSRR is prone to degradation with accelerated transistor aging [4]. The effects of aging degradation on the pass transistor (PT) were depicted from the deviation in current–voltage (*IV*) characteristics of the PMOS PT and also in the overall PSRR. However, the above strategy was only tested for stand-alone LDOs, and the applicability of the former strategy in a complex system on chips (SoCs) is still a question. In this paper, we have expanded our investigation of using LDO degradation towards the detection of recycled SoCs.

A major advantage of using PDNs to detect recycled ICs is that it is available in nearly every IC and SoC. However, this strategy is only applicable to SoCs, which consist of LDOs with output capacitors. To measure the PSRR for any LDO, the subject matter expert (SME) must have access to the LDO's output pin. In most LDOs, the output is coupled to an external capacitor to stabilize the LDO loop. Certain LDO designs may not have an output capacitor; these are known as cap-less LDOs. For a cap-less LDO, reverse engineering, the LDO's output pin may become challenging if embedded within an SoC. Thus, advanced reverse engineering techniques are required to solve the above issue, resulting in extra cost. Thus assuming that we have an SoC consisting of an LDO with an external capacitor, our proposed method can be applied to detect recycled SoCs. In this paper, we observe the degradation of the PSRR of LDOs embedded within SoCs and apply automated machine learning (ML) methods to detect recycled or new SoCs. The supervised and unsupervised algorithms we have used belong to the family of Gaussian mixture models. Our contributions can be summarized as follows:

- We extend the technology of recycled IC detection from stand-alone LDOs to LDOs embedded within SoCs. This extension enables the technique to be useful in most types of SoCs and increases the applicability of the procedure.
- To implement the above, we provide a comparison of the recycled detection in stand-alone LDOs and that in SoCs, providing a clear description of both techniques' pros and cons.
- We also answer the relevant questions about the different challenges that may arise while implementing our process or any other process, in general, to detect recycled SoCs, such as the availability of golden data and whether they are applicable in our technique for recycled SoC detection.
- We implement supervised and unsupervised ML methods to detect recycled SoCs and provide insightful analysis for both techniques. The maximum accuracy of the above techniques is 90% and 74%, respectively. We also implement the detection of recycled SoCs using training data of PSRR degradation from stand-alone LDOs from four different vendors. This type of semi-supervised training was successful in recycled detection for stand-alone LDOs in Ref. [3]. But for SoCs, this method can either detect new or recycled SoCs and is unable to distinguish former from the latter; thus, increasing the risks of false identification. A detailed explanation of the applicability and limitations of the above algorithms are also provided.

We have explained the above contributions in the following sections. The rest of the paper is organized as follows. "Preliminaries" discusses the preliminaries of transistor aging and the generic structure of an LDO with different metrics. "Recycled SoC Detection" describes the proposed approach of recycled SoC detection with a detailed description of the steps and the ML algorithms used. "Experimental Setup and Aging Analysis" describes the experiments performed along with the experimental set up used, and "Recycled Chip Classification Results and Discussion" analyzes the results. "Conclusion and Future Work" concludes the paper with possible future works.

Preliminaries

To explain the proposed technology, it is crucial to provide certain preliminaries on the important concepts used in this paper. In the following subsections, we explain the necessary background on LDOs and transistor aging models.

General Concepts of Transistor Aging

Transistor aging is one of the major causes of reliability issues faced by modern ICs. It is the result of trapped charges and broken bonds at gate dielectric interfaces, which increases threshold voltage (V_{th}) and switching activity, thereby deteriorating transistor performance in scaled modern devices. Bias temperature instability (BTI) results in a positive shift in the absolute value of V_{th} in both PMOS and NMOS. BTI is the condition often referred to as DC stress when the PMOS/NMOS has already pulled up/down, but the gate is still biased in strong inversion. The drain-to-source voltage becomes zero signifying a negligibly small lateral electric field. For PMOS, the condition is called negative BTI (NBTI), whereas, for NMOS, it is positive BTI (PBTI). Hot carrier injection (HCI) occurs when the transistor is switching under strong inversion ($|v_{gs}| \approx V_{dd}$) and the lateral electric field is high ($|v_{ds}| \approx V_{dd}$). During transistor switching, the accelerated carriers drift towards the drain under the influence of the lateral electric field. Channel hot carriers (CHC) are generated when the source-to-drain current flowing through the channel reaches energy above the lattice temperature. These hot carries gain energy and get injected into the gate oxide, forming charge traps. The charge traps cause a shift in the device performance like V_{th}, transconductance, and saturation current of the transistor, as discussed in Ref. [2]. HCI degradation increases by a factor of $t^{1/2}$ (where t is time) and BTI increases as a factor of t^n where $n = 0.1–0.2$. Since the multiplicative constant of HCI is much smaller than that of BTI, BTI overshadows HCI for a short amount of time, as suggested in Ref. [22]. Long term, HCI may cause equal or higher degradation in device parameters than BTI.

Low Dropout (LDO) Regulators

An LDO is a type of linear regulator capable of maintaining an output voltage even when the input is very close to the output (low drop-out). Drop-out voltage is defined as the input-to-output differential voltage, where the regulator fails to regulate the output voltage until the further reduction of the input voltage. The role of an LDO is indispensable in the power supply of any SoC/IC. It provides isolation between the input and output, thus rejecting the noise and ripples (glitches) in the input supply at the output to provide a stable, low noise, fixed output voltage.

As shown in Fig. 1a, the block diagram of an LDO consists of a feedback loop with an error amplifier (EA), a pass transistor (single NMOS or PMOS), and a resistor divider. A bandgap circuit provides a fixed reference voltage to the EA. The pass transistor (PT) acts as a variable resistor controlled by the EA, and the feedback resistor divider circuit level-shifts the output voltage to the EA input. The EA monitors the error between the input and the output voltage and accordingly controls the gate-to-source voltage (v_{gs}) of the PT to regulate the output at a fixed voltage. If the feedback voltage is smaller than the reference voltage, then the gate voltage of the PT is lowered, increasing the v_{gs} as well as the current flowing through the PT, thus increasing the output voltage. If the feedback voltage is higher than the reference voltage, v_{gs} of PT decreases, reducing current and output voltage. The drop-out voltage for a generic LDO, as shown in Fig. 1a, is the drain-to-source voltage drop, which appears across PT. One of the major performance metrics of an LDO is its capability of rejecting the ripples of the input supply at its output. This metric is known as the power supply rejection ratio (PSRR) of the LDO. The ripple can originate from the power supply or from a DC/DC converter or even due to sharing an input supply between different circuit blocks in the system. PSRR is expressed as $\text{PSRR} = 20\log(\frac{v_{out}}{v_{in}})$ where, v_{out} and v_{in} are magnitudes of voltage glitch at output and input, respectively. Apart from PSRR, the quality of voltage regulation provided by an LDO is specified by metrics like transient line regulation and load regulation. Metrics like power efficiency and current efficiency determine the power and current consumption efficiency of the LDO [17].

To observe the effect of accelerated transistor aging on the PSRR and other properties of an LDO, we previously fabricated an LDO in TSMC 65nm process in Ref. [4]. Our experimental results in Fig. 1b–d show the effect of accelerated transistor aging on the PT and the PSRR of the LDO. The effect of DC stress mainly involves operating at a constant supply voltage (V_{DD}), which is increased by 10% over the normal V_{DD} and at a higher temperature of either 85 or 105 °C. But the LDO is providing a constant current at the load. The phenomena of BTI ("General Concepts of Transistor Aging") are mainly observed with DC stress. During AC stress, the operating temperature and V_{DD} is increased like DC stress, but the output current also fluctuates from 0.9 to 1.1 mA. This results in a combined effect of both HCI and BTI on the LDO. The degradation of the PSRR was recorded for both AC and DC stresses and showed an approximate variation of 1.6 dB for DC stress and 2–5 dB for AC stress. We also observed the degradation of DC stress on the I_d/V_{gs} characteristics of the PT of the LDO and observed sufficient degradation. In conclusion, the above experiments provide proof of considerable degradation of LDO PSRR due to accelerated transistor aging, which can be used to detect recycled SoC containing LDOs.

In Ref. [3], we applied the above degradation to detect recycled stand-alone LDOs across four vendors. The PSRR degradation varied from vendor to vendor and also across process variation, but even with small PSRR degradation [for vendor 3 (V3) and vendor 4 (V4)], supervised

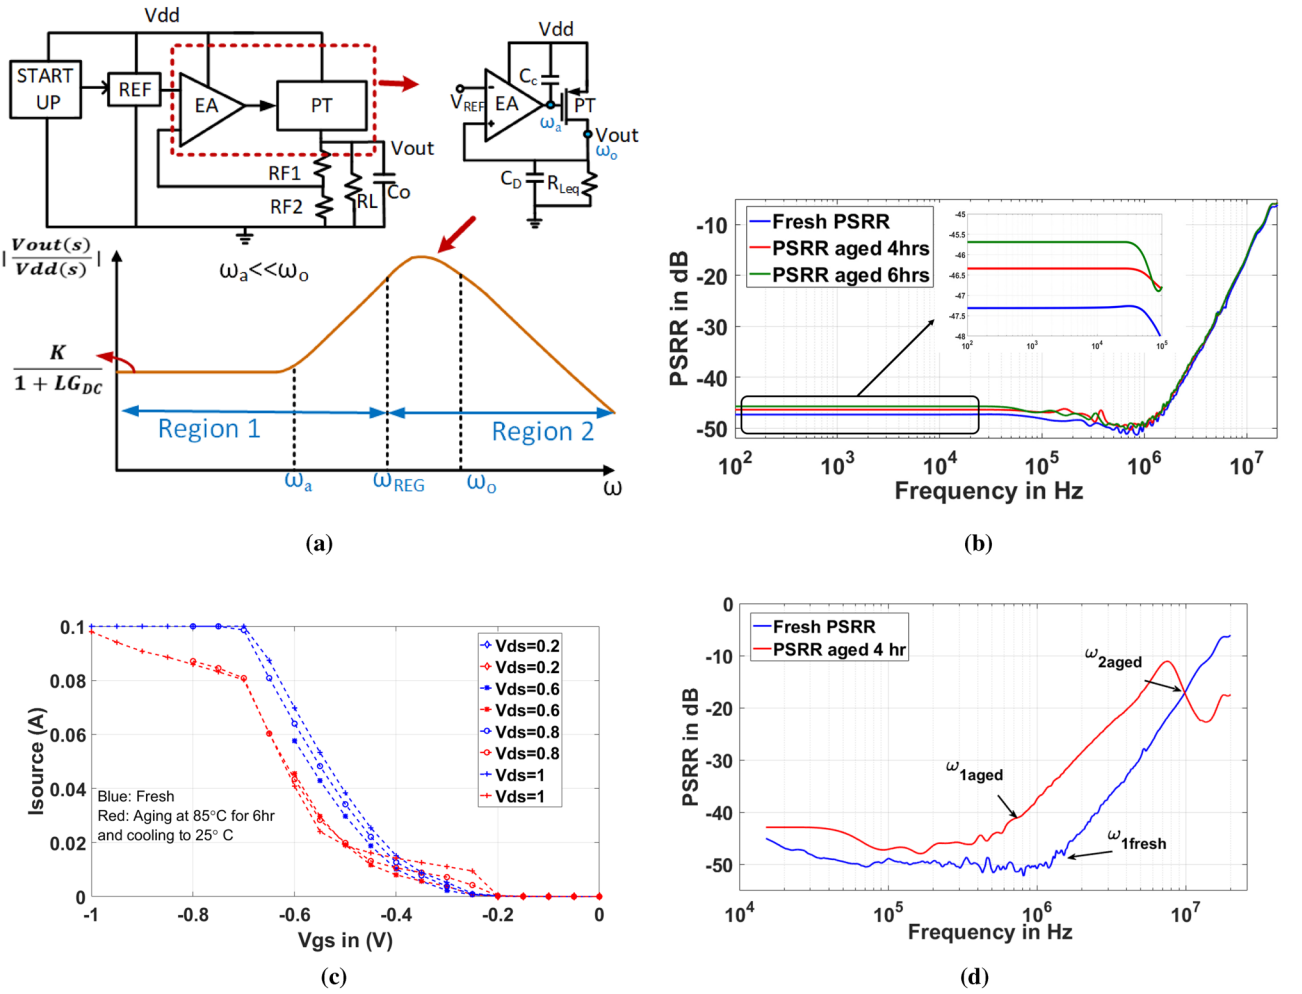

Fig. 1 **a** LDO block diagram with PSR (linear scale) curve [3]; **b** Drain current with respect to gate voltage (I_d/V_{gs}) curve degradation of LDO PT with 6 h of DC stress at 105°C and 10% V_{DD} increase [4]; **c** Degradation of PSRR of LDO under DC stress of 4 h and 6 h at 105 °C [4]; **d** Degradation of LDO PSRR under combined AC and DC stress for four hours at 105 °C [4]

ML algorithms were able to detect recycled LDOs till a maximum accuracy of 90% as reported in Ref. [3]. The primary applicability of the paper targeted universal recycled detection for both digital and analog mixed-signal (AMS) ICs at zero cost with no hardware overhead and minimum measurement equipment. Supervised and unsupervised ML algorithms were used to detect the recycled stand-alone LDOs. For the supervised ML method, the training set involved a specific vendor, and the testing set also pertained to the same vendor. In unsupervised ML, no labels were provided, and thus, the drawback of the requirement of golden data was nullified. The maximum average accuracy for unsupervised ML was 74%. Since the supervised algorithm was more successful than unsupervised applications, we reduced the requirement of golden data using semi-supervised training and improved accuracy. In this case, the training set comprised of one vendor, whereas the testing set consisted of other vendors. Both supervised and semi-supervised applications resulted in detection accuracy greater than 90%. The biggest takeaway from the semi-supervised detection is that there exist certain similarities in PSRR degradation despite design differences existing in the LDO ICs across vendors. This similarity increases the scope of ML algorithms' application to detect recycled LDOs and reduces the requirement of golden data.

It must be noted that the previous work only focuses on the detection of stand-alone LDOs and does not explore the application of the proposed technique to SoCs. While the presence of LDOs in most SoCs/PCBs expands the scope of application of the technique, the aforementioned prior papers do not provide any conclusive results which proves that the same method can be applied to LDOs that are embedded in SoCs. In this paper, we expand the above detection strategy and apply it to LDOs within SoCs and discuss the scope of such detection.

Table 1 Comparison of switching mode power supplies (SMPS), switch capacitor (SC) and LDO power converters [23]

Power converter	SMPS	SC	LDO
Step-up conversion	Possible	Possible	Not possible
Power efficiency	High	Medium	Limited to $\frac{V_{out}}{V_{in}}$
Load regulation	Good	Poor	Good
Physical area	Large	Medium	Small
Applications	Microprocessors, DSPs, SRAMs, hard-discs	EEPROM, DRAM, flash, and mixed-signal	DRAM, SoCs

Recycled SoC Detection

Recycled IC detection has been a targeted research initiative, and there are many methods that have been proposed before to detect IC recycling. Some approaches use statistical methods to detect degradation, as shown in Ref. [16]. Some use low-cost on-chip sensors called CDIRs to detect the transistor aging due to recycling like in Refs. [10, 12], etc. While they are very effective, most of these sensors require additional circuitry adding to hardware design efforts and increased silicon area; hence, to our knowledge, no vendor has adopted them yet. Some other detection methods include side-channel analysis, including power and current analysis [26], require golden data. Compared to the recycled detection of stand-alone ICs, much less work has been done to detect recycled SoCs. In Ref. [13, 14], the authors have proposed a framework to detect recycled SoCs by an aging-sensitive SRAM selection algorithm. The method is applicable to SoCs consisting of embedded SRAMs and can be applied at near zero-cost[1] to most SoCs. While the above holds for most digital SoCs, which contains embedded SRAM, it may not be applicable to purely analog or analog-mixed signal (AMS) SoCs which do not contain embedded SRAM memories. Compared to existing methods, our proposed method involves using a power supply component like an LDO, which is present in most digital, analog-mixed signal SoCs as well as stand-alone ICs and thus can be universally applied. Further, it is free from enrollment steps and does not require any additional silicon area or memory. This makes it applicable even to legacy and commercial-off-the-shelf (COTS) components.

Different Types of Power Supplies in ICs

Increased scaling and requirement of efficient power supply units have revolutionized the power architecture in electronic circuits. Power converter systems in ICs/SoCs mainly consist of a DC-DC converter, as shown in Fig. 2. The power from DC–DC converters is either stepped up (boost) or down (buck) to supply a required amount of current at the load, regulating the output voltage with varying load, line, and pressure–voltage–temperature (PVT) variations. There are two different topologies of power converters used in SoCs: switching topologies, and linear topologies. Switching topologies use passive storage elements like capacitors and inductors to convert and store the power, whereas linear topologies use resistive elements to dissipate the power. Switching mode power supplies (SMPS) consist of an inductor and capacitor (*LC* tank) circuit which stores or dissipates the power and charges the capacitor at the output. Different types of SMPS converters include buck and boost converters, etc. Another type of switching topology includes switch capacitor (SC) power converters, which are often referred to as charge pumps and mainly use a capacitor and a switch for the conversion. Linear power converters are used to convert a DC source from one voltage level to another by dissipating the excess power in the resistive output device. LDO is an example of linear power converters, which can only be used to step down from a given voltage level and fail to step up to a higher voltage level. With increased technology scaling leading to the development of low power designs, LDOs are used in most modern SoCs to step down the off-chip/battery supply voltage to lower levels for the majority of the embedded digital and mixed-signal blocks. A comparison of the different converters is shown in Table 1. It can be seen that LDOs provide proper regulation with small physical area requirements, which makes it suitable for on-chip integration specifically for SoCs. A detailed explanation of the general architecture of the SoC power supply with LDOs is given below.

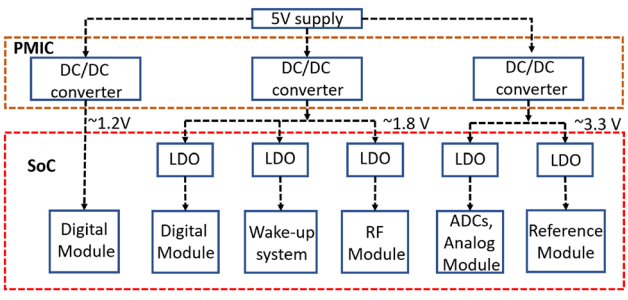

Fig. 2 Block diagram showing general architecture of power supply in a SoC

[1] Enrollment tests are needed as well as non-volatile storage.

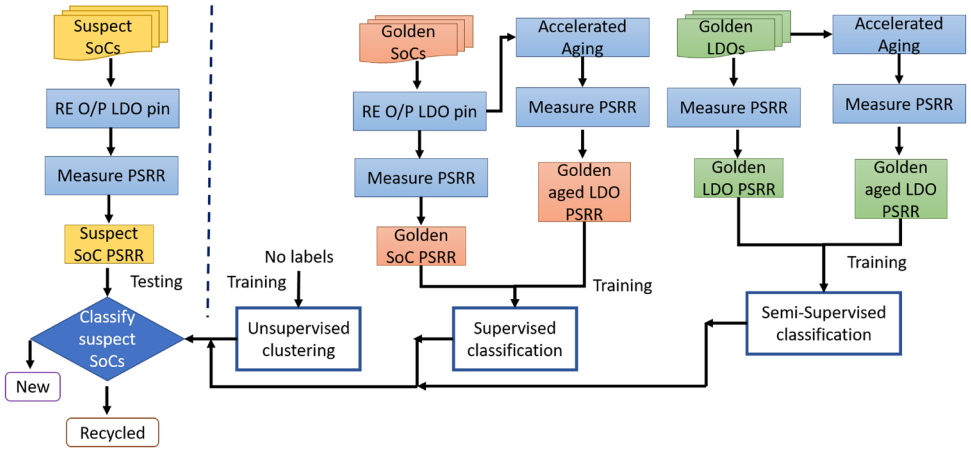

Fig. 3 Flowchart showing the proposed approach for recycled SoC detection

General Architecture for Power Supply in SoCs

A generic block diagram of the power supply architecture in modern SoCs is provided in Fig. 2. The SoC consists of multiple blocks that require different reference or supply voltages due to original specifications. The on-chip voltage supply is significantly lesser than the off-chip supply. It ranges from a fraction of volts for digital modules to a couple of volts for high-precision ADCs, buffers, and analog modules. The supply voltage is provided to the power management integrated circuit (PMIC), which mainly consists of DC–DC converters used to up-convert or down-convert the supply voltage according to the requirement of individual sub-circuits. There can be various sub-circuits associated with an SoC, including digital, analog/analog to digital converters (ADCs), RF modules, etc., as shown in Fig. 2. Each of these modules requires their supply voltage levels, provided by the DC–DC converters and the LDOs. The LDOs perform a major task in these systems of isolating the SoC power from the PMIC. For AMS blocks, suppression of the power supply noise is critical and thus requires LDOs, which can provide a ripple-free regulated output by suppressing the power supply noise appearing at the output of the battery or the DC–DC converter. In digital blocks, there is a substantial amount of switching noise, which should be prevented from getting coupled at the DC–DC converter's output. An LDO also provides this reverse-isolation, which prevents this switching noise from appearing at the converter outputs. Thus, the main functionality of an LDO includes ripple suppression, isolation, and noise regulation, making LDOs an essential component in the power management units (PMUs) of SoCs.

Proposed Methodology

The degradation of electrical parameters like PSRR in Ref. [4] for stand-alone LDOs demonstrates the possibility of degradation over usage/recycling. The successful implementation of recycled detection for stand-alone commercial-off-the-shelf (COTS) LDOs in Ref. [3] discussed in "Low Dropout (LDO) Regulators" and the presence of LDOs in modern SoCs described in "General Architecture for Power Supply in SoCs" motivates the application of the technology to detect recycled SoCs. Apart from the universal applicability of the method, there are also other advantages of the proposed method. It does not require any enrollment tests or other non-volatile storage requirements. The hardware overhead is minimum, and the only cost of application pertains to the experimental setup which consists of basic electrical test components that are easily available in most testing labs. The proposed approach can be divided into the following steps: (1) Identifying the type of SoCs where the proposed technology can be implemented; (2) Reverse engineering the position of the LDO within the SoC and identifying the output of the LDO to measure PSRR; (3) Measuring the PSRR of the sample LDO embedded within the suspect SoC; (4) Identifying the correct set of ML tools developed later on in this paper to classify the SoCs as recycled or new; (5) Identifying correct ML algorithms to develop the ML tools. An elaborate description of the flowchart is shown in Fig. 3.

Identifying Type of SoCs

The essential requirement of our method involves the presence of linear regulators or LDOs within the power supply architecture of the SoC. Since most SoC power converters use LDOs for on-chip integration due to low area requirements, most state-of-art SoCs consist of LDOs. For simple reverse engineering of the LDO output pin, we restrict our proposed approach to only SoCs consisting of embedded LDOs with an output capacitor. A typical LDO regulator requires an external capacitor for better transient operation, improved PSRR, and stability of the LDO. However, the presence of an external capacitor can cost extra area and output pins for on-chip integration in SoCs. Thus, there are LDOs that eliminate such capacitors to save area and

extra pin-outs. These LDO architectures are known as capacitor-less or cap-less. While it may seem that cap-less LDOs are an obvious choice for SoCs, such implementations undermine the LDO performance and result in poor PSRR performance and lesser stability. Due to the advantage of integration, efforts have been made both by researchers and industry to provide practical cap-less LDOs. Nevertheless, cap-less LDOs consist of severe limitations preventing them from usage in practical applications like SoCs. Several reported architectures of cap-less LDOs provide acceptable performance but only for a single parameter like line regulation, load regulation, settling time, etc. rather than multiple parameters. Most of these LDOs can only load control for a narrow range of load capacitance (1–100 pF) and fail to provide regulation at higher load currents (≈ 100 mA, which is typical for commercial LDOs with cap). Cap-less LDOs also suffer from lower performance in terms of PSRR and dynamic performance [21]. Thus, due to the performance restrictions of cap-less LDOs, most SoCs still use LDO architectures with output capacitors where the LDO output pin can be easily reverse-engineered, and the proposed technology can be applied [21].

Inquiries may arise regarding the scalability of the proposed method to state-of-the-art technology nodes like 5 nm or 7 nm. While analyzing transistor aging LDOs we observed the aging degradation in 65nm technology node. It must be understood that for lower technology nodes, the effect of transistor aging is even worse, as suggested in Ref. [19]. Though with technology scaling, the transistor geometries have scaled-down, the supply voltages have not scaled proportionately for performance requirements. Thus, the applied electric field across gates has increased with lower technology nodes, leading to worse aging effects. Thus, the aging behavior analytically should be more prominent for lower technology nodes that would improve our case's detection accuracy. Another concern regarding lower technology nodes would be the availability of an output pin of the LDO. Since LDOs with an output capacitor requires additional hardware overhead, some SoCs in lower (advanced) technology nodes use cap-less LDOs. The technology nodes applicable specifically for AMS and analog ICs are older compared to digital ICs. A detailed explanation of this has been given in Ref. [1], where a comparison has been shown between digital and AMS SoCs. The comparison suggests that AMS SoCs consist of lesser number of transistors (in count of 1000 s) compared to digital SoCs where millions of transistors are present. Also, AMS SoCs require higher supply voltage compared to digital SoCs and are custom designed following older technology nodes. Thus, several AMS, analog, and legacy SoCs use older technology nodes with LDOs containing output capacitors. Since recycling is more prevalent for these SoCs/ICs [1], the application of the proposed method is appropriate. Even for lower technology nodes with cap-less LDOs, reverse engineering methods can be applied to track the embedded LDO's output. This would surely result in excess costs but remains as a possible option. Nonetheless, for SoCs in general, the proposed approach is still an attractive choice as many of them use LDOs with an output capacitor as described in Ref. [21].

Reverse Engineering Output Pin of LDOs

Since we have chosen embedded LDOs with an output capacitor for our proposed method, it is easy to detect the output pin from the SoC specification sheet. Most SoCs containing LDOs with output capacitors contain a dedicated output pin for the external capacitor of the LDO that can be chosen by the user within the specified limits mentioned in the specification sheet. Since the capacitor needs to be attached to the output node of a generic LDO, this pin-out can be considered as the output of the LDO.

Measuring PSRR of Sample LDOs Embedded in Suspect SoC

This step involves the measurement of PSRR of the LDO embedded within suspect SoC, which needs to be identified as recycled (counterfeit) or new (genuine). The experimental setup will be described in later sections. The output pin of the LDO and the V_{DD} supply pin is identified from the previous step of reverse engineering. A small noise signal is coupled to the V_{DD} pin, and the corresponding power spectrum at the LDO output is recorded. The PSRR is calculated by subtracting the input noise spectrum (in dB) from the output power spectrum (also in dB). Likewise, sample PSRR data from suspect SoCs containing LDOs are recorded and are given as input to the ML tool for automated detection.

Identifying Correct Set of ML Tools

In this paper, we have used ML tools to detect SoC recycling as described in our later sections. The technology is provided with three different types of ML approaches that can be used to detect a suspect, SoC. (1) Supervised ML requires golden data of authentic samples of the suspect SoC. The training for this ML tool is executed with

PSRR data from genuine new and aged samples.[2] This method provides an accuracy of up to 90%, as discussed later. Though it provides good accuracy, the stringent requirement of golden data may prove as a drawback for the process. If the subject matter expert (SME) possesses the correct set of golden data, this ML tool can prove to be extremely beneficial. (2) Semi-supervised ML can detect recycled SoCs even if the golden data from the specific category of SoC is unavailable. If the SME possesses PSRR data from new and aged samples of other categories of SoC or even other stand-alone COTS LDOs, semi-supervised ML can effectively detect recycled SoCs. Its training set involves new and aged PSRR data from other SoCs or stand-alone LDOs, thus alleviating stringent requirements of golden data. This method also provides high accuracy of 98%, making it a good choice for the SME, requiring golden data from any other LDO chip belonging to different vendors. But, according to our results, it is observed that the classifier obtained with the above semi-supervised training can either detect a new or a recycled SoCs and not both. Thus, this method has the risk of an increased number of false negatives that limits the applicability of the method; (3) Unsupervised ML can be used when no golden data is available to the SME (worst case scenario). Unsupervised ML requires no label for training and clusters the available PSRR data from suspect samples into new and recycled. Though the requirement of golden data is completely nullified in unsupervised ML, the accuracy obtained is comparatively lower than the other ML approaches. Nonetheless, it provides maximum accuracy of 74% and can prove beneficial in cases where no golden data is available to the SME.

Identifying Correct ML Algorithms to Develop ML Tools

Identification of the correct set of ML tools is an important decision for the SME, as discussed in our previous section. From the point of view of the tool developer, it is also important to understand the specific ML algorithms that must be applied to develop the above set of ML tools discussed. We have used the algorithms belonging to the family of Gaussian mixture models (GMMs) to develop the ML tools that are used to detect recycled SoCs in this paper. A detailed explanation of the algorithms used is discussed below.

Gaussian Mixture Models A dynamic system is dependant on multiple regimes and thus switches its behavior by shifting from one regime to another. To efficiently describe such systems, a mixture of models or components is required. Thus, mixture models are universally applied to describe such dynamic systems. A system state is an unobservable variable that appears within the bounds of the above individual regimes. To estimate the state variable, each of the above regimes is represented by state-space models (SSMs). An SSM is a common approach to analyze structured, sequential data representing a time-series. In ML, one of the foundation steps involves representing the data with the help of a mathematical algorithm.

Here, we will explain the applicability of GMM algorithms to represent the PSRR data. It must be noted that the data is multi-dimensional and is collected over a range of frequency consistently across a predetermined aging time. We collected PSRR data after every hour of accelerated aging over a continuous 8 h range. Thus, our PSRR data can be viewed as a time series. By analyzing from the other dimension of frequency, it can be seen that the data set is ordered and exhibits irregularities due to environmental noise and process variation. Our data set represents a multi-dimensional dynamic system that can be represented as a time series despite being non-temporal naturally. Thus, classification or clustering of the SoC PSRR data can be seen as a sequence labeling problem for a non-temporal data set exhibiting properties of a time-series. In ML literature, such problems are analyzed effectively by time-series analysis that reflects the non-temporal behavior of the data, as can be seen in Refs. [6, 7] etc.

The SSM model described above is a popular method of tacking the above problem. In an SSM, it is assumed that a sequence of measured data y (in a vector form) $\mathbf{y}_1, \mathbf{y}_2, \cdots$ is generated by some hidden state variables $\mathbf{x}_1, \mathbf{x}_2, \ldots$ with joint probability,

$$p(\mathbf{x}_{1:F}, \mathbf{y}_{1:F} \mid \theta) = \prod_{f=1}^{F} p(\mathbf{x}_f \mid \mathbf{x}_{f-1}, \theta)\, p(\mathbf{y}_f \mid \mathbf{x}_f, \theta),$$

where θ is the model parameter, $\mathbf{x}_{1:F}$ and $\mathbf{y}_{1:F}$ are the sequence of F sequences of the hidden state variables and the measurements, respectively. The indices f signify that the PSRR data has been recorded over a specific frequency range in an ordered manner. We have analyzed the SoC PSRR data with linear Gaussian SSMs, composed of multivariate Gaussian variables associated in a linear fashion, according to the following equation [6].

$$\begin{aligned} \mathbf{y}_f &= C\mathbf{x}_f + \mathbf{v}_f \\ \mathbf{x}_f &= A\mathbf{x}_{f-1} + \mathbf{w}_f. \end{aligned} \tag{1}$$

The matrices C and A represent the linear relationship while the vectors $\mathbf{v}$ and $\mathbf{w}$ represent uncertainty. These vectors have a covariance of R and Q and follow Gaussian distributions. Thus, the embedded LDO parameters can be described by

[2] Aged samples could correspond to accelerated aging of SoC/IC samples to represent synthetic recycled samples or to real-time aging of SoC/IC samples. In our experiments, we use the former for simplicity.

the function $\theta = (A, C, Q, R)$. The vectors $\mathbf{v}$ and $\mathbf{w}$ model all the different uncertainties that can be encountered by the system, including the impact of transistor-level aging, environmental noise, measurement noise, and other process variations. Although transistor aging is a physical phenomenon and is difficult to be visualized as an uncertainty, the heterogeneous models of SoCs and the different design components introduce sufficient variability within the aging profile of the LDOs embedded within the SoC. Thus, our ML algorithm models the effect of aging as an uncertainty represented by a Gaussian variable instead of characterizing the gate-level-transistor aging model [5]. It is common, in many works of ML literature, to assume a Gaussian distribution for other unknowns like environmental and measurement noise, etc.

Markov Assumption Another important aspect of the linear Gaussian SSMs discussed before is the dependency of the hidden state variables $\mathbf{x}_f$ on one another or, in other words, the characteristics of first-order Markov dynamics. In our problem, we adopt the frequency from which we start measuring the PSRR data. This can be derived from the PSRR equation of the LDO defined later in Eq. (3), which shows that the PSRR is dependant on the loop gain or *LG*. The component *LG* varies for frequency but can be depicted as a fraction of the *LG* at DC. Thus, if the PSRR is recorded over two successive frequency points, say f_i and f_{i+1}, then the PSRR at frequency f_{i+1} is dependant on the PSRR at the previous frequency f_i.

Parameter optimization To learn the parameters of the linear Gaussian SSM defined above, we employ a *k*-means algorithm for unsupervised clustering. Given a set of n observations of d dimensional vectors $(\mathbf{a}_1, \mathbf{a}_2 ... \mathbf{a}_n)$, the K-means algorithm partitions the set of n observations into k sets $(\mathbf{S}_1, \mathbf{S}_2, ... \mathbf{S}_k)$ where $k < n$ by minimizing the sum of squares (or variance) within each cluster as shown in Eq. (2) below, where μ_i is the mean of the points within the ith set S_i

$$\arg\min_{S} \sum_{i=1}^{k} \sum_{x \in S_i} \|a - \mu_i\|^2 = \arg\min_{S} \sum_{i=1}^{k} |S_i| \sigma^2(S_i). \tag{2}$$

The *k*-means algorithm has much in common with the Expectation–Maximization (EM) algorithm, which is a well-studied approach to learn parameters for linear Gaussian SSMs [6]. The recycled detection of LDOs presents a complex model with several uncertainties, as described above. It may be difficult for the EM model to determine the model size for such a complex model. Therefore, in Ref. [3], we employed a variational Bayesian inference over the parameters of the probabilistic models in conjunction with the EM algorithm. We called this approach the VB algorithm and compared its performance with that of the *k*-means algorithm for stand-alone LDOs. It was concluded that the *k*-means algorithm performed better in terms of accuracy. Since detection of recycled stand-alone LDOs and LDOs embedded in SoCs has considerable similarities in the physical aging procedure and other uncertainties, we only used the *k*-means algorithm in this paper for unsupervised detection.

We also analyze the scenario of recycled SoC detection using supervised and semi-supervised ML tools. In this case, we use the KNN algorithm, which is one of the closest approaches to the VB and *k*-means method. The idea behind this approach is that instances with similar properties will be near to each other in a given data set. Thus, when a new sample is provided to the algorithm, the label of the provided sample should be similar to that of its nearest neighbors. The *k* nearest neighbors are calculated using the Euclidean distance metric, and the label of the new sample is the common label of its nearest neighbors. To accurately classify new samples, the number of neighbors is vital. Thus, for noisy data sets with complex structures and uncertainties, a high value of *k* is selected, whereas, for intimately placed data samples, a smaller *k* needs to be chosen. Nevertheless, when applying KNN algorithm, the value of *k* can be automatically chosen by fitting the best classifier to the data.

Experimental Setup and Aging Analysis

Compared to recycled IC detection, there are many more challenges that need to be answered for recycled SoC detection: (1) Difficulty in automation: In Ref. [3], entire test setup for stand-alone LDO PSRR measurement was easy to automate. Since most LDOs have similar characteristics, it was easier to obtain multiple vendors producing LDOs with similar footprints and related specifications. Whereas SoCs are completely different from one another, and thus measuring PSRR from each type of SoC cannot be automated easily; (2) Increased cost: Compared to stand-alone LDOs, SoCs containing LDOs are much more expensive; thus, the increased cost, lack of automation, and increased time for PSRR measurements served as a bigger challenge in this paper; (3) Lack of samples: The increase in cost and the required time, limited the total number of samples which could be tested. Even allowing more time, the heterogeneous property of SoCs requires different test benches for different vendors and types of SoCs, which also limited the number of samples that can be tested; (4) The reverse engineering of output pin of LDO: In recycled IC detection that only targets individual ICs, output pins are easily available. For recycled SoCs, the output of LDOs needs to be reverse engineered, thus requiring more time and effort than individual ICs.

It is difficult to automate the measurement of SoC PSRR, and also the number of samples recorded is less due to increased cost and time requirement. In this paper, we have explored supervised, semi-supervised, as well as

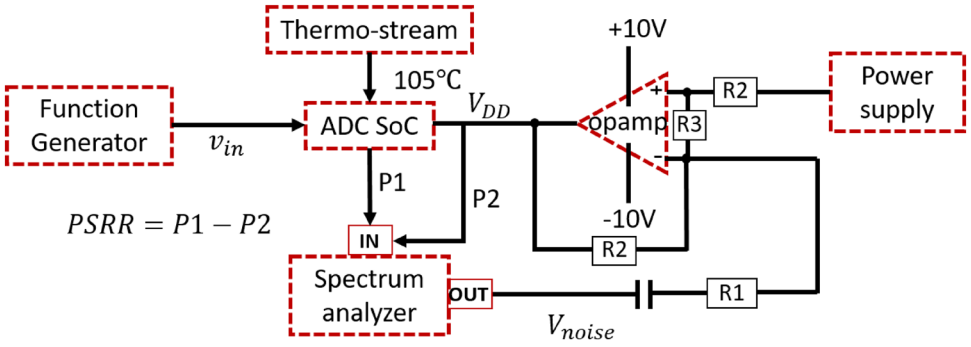

Fig. 4 Block diagram of the experimental setup

unsupervised ML approach for recycled SoC detection. Some of these techniques, like semi-supervised and unsupervised ones, reduces the requirement of a large number of golden samples. For the unsupervised approach, no labels are required, which completely negates the requirement of golden data. But this comes at the cost of detection accuracy, which was approximately a maximum of 86% for stand-alone LDOs in Ref. [3]. For a semi-supervised detection approach for recycled SoCs, the entire training is done with previous recorded PSRR data from COTS stand-alone LDOs, which can be easily automated and obtained at a comparatively lower price, as described in [3]. In Ref. [3], we explored recycled detection on stand-alone LDOs from four different vendors. The supervised ML approach trained on samples from a specific vendor. The testing for the supervised approach was also done on the same vendor from which training samples were obtained. This method provided an accuracy of up to 97%. Another successful attempt of recycled detection performed in Ref. [3] was the semi-supervised approach, where training samples were provided from a specific vendor, but the testing was performed on suspect samples from all other vendors except the above specific vendor. This strategy worked (with maximum accuracy greater than 90%) and was able to detect recycled and new LDOs of a different vendor after being trained on samples from another vendor. The successful implementation of semi-supervised ML for stand-alone LDOs motivated us to observe the same for LDOs embedded in SoCs in this paper. The analysis and results of all the above supervised, unsupervised, and semi-supervised techniques for recycled SoC detection have been documented in "Recycled Chip Classification Results and Discussion". We have also answered the challenge of reverse engineering of output pin by restricting the type of SoCs where this technology can be applied. The increased usage of LDOs with output capacitors for on-chip integration owing to better performance is also an advantage that supports our technology and reduces the complications of reverse engineering. Since recycled samples are mostly unavailable in the market for inspection, we have followed a procedure of accelerated aging to produce a synthetic recycled counterfeit version of the corresponding samples. It must be noted that the amount of accelerated aging can be translated to the correct amount of real-time aging using certain equations shown in Ref. [18].

We have executed the following experiments; (1) Initial Data Collection: At first, we have obtained the initial PSRR data from the LDO embedded within an SoC. In this paper, we have used a Delta-Sigma ADC containing an embedded LDO as a sample SoC. We obtained initial PSRR from four such ADC SoC samples, which constituted the set of the PSRR data for new SoCs. The experimental setup consisted of a 2-channel 24-bit delta-sigma ADC SoC, which is connected to the function generator and the power supply to turn on. The SoC evaluation board came with a software portal that recorded the digitally encoded signal of the analog input signal that was provided to the SoC that was constantly made to run through scripts from the software portal. The initial task was to reverse-engineer the output port of the LDO located within the SoC. Since this type of SoC mainly used LDOs with output capacitors for precise performance and specifications, reverse-engineering the LDO's output pin was easy. The LDO output was provided as one of the pinouts of the SoC, which provided easy access to the PSRR of the LDO. In order to record the PSRR, a tracking spectrum analyzer was used to generate an output noise signal of magnitude 1 dBm, and it was coupled to the V_{DD} of the SoC using a summing amplifier as shown in Ref. [20]. The input power spectrum at the V_{DD} and the output power spectrum at the LDO output was recorded. Subtracting the input power spectrum from the output provided the initial PSRR of the LDO embedded within the SoC. The detailed experimental setup has been shown in Fig. 4. After the initial PSRR data was recorded from four SoCs, we move on to the accelerated aging of the SoCs as our next step.

(2) Accelerated Aging: We executed accelerated aging at a high temperature of 105 °C for eight consecutive hours, and PSRR data was recorded every hour. This comprised the PSRR data set for recycled SoCs. This step involves accelerated aging of the SoCs at a high temperature of 105 ° C while the SoC was always active. The temperature was increased using a thermostream maintaining the SoC at 105 ° C for 8 consecutive hours. In this paper, we implement accelerated

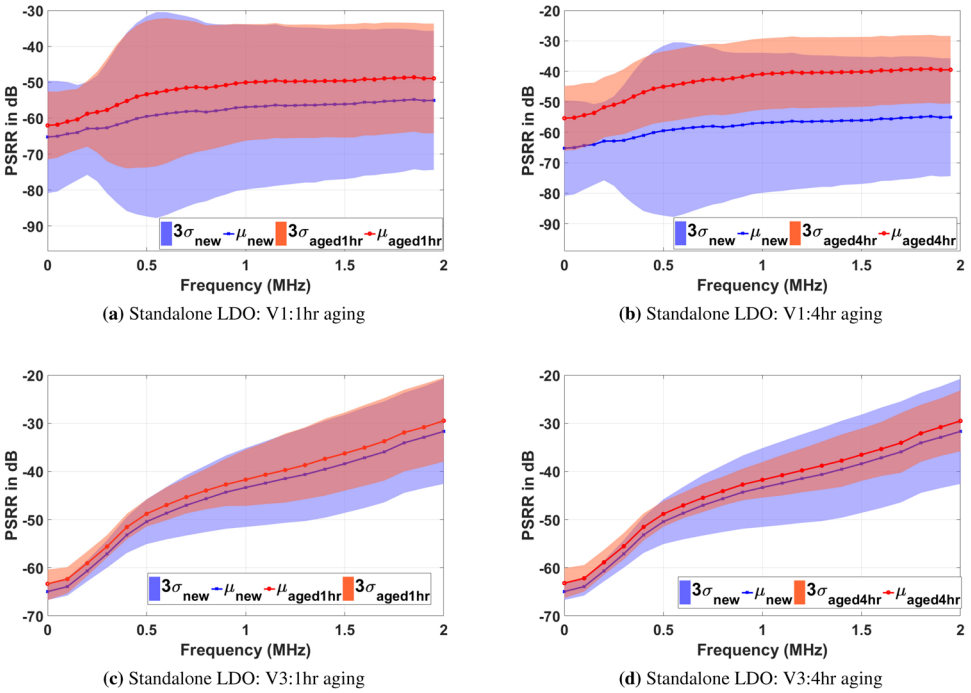

Fig. 5 Silicon data showing mean PSRR degradation of LDO for Vendor 1 (V1) and Vendor 3 (V3) for 1 h (**a**, **b**) and 4 h (**c**, **d**), respectively. This data is provided to solely compare the degradation profile of the LDO embedded in an SoC to that of stand-alone LDOs [3]

(a) Standalone LDO: V1:1hr aging

(b) Standalone LDO: V1:4hr aging

(c) Standalone LDO: V3:1hr aging

(d) Standalone LDO: V3:4hr aging

aging to mimic IC recycling. To emulate the real-time aging, we utilize synthetic aging acceleration. To understand how the artificial aging acceleration correlates to real-time aging, we must consider the accelerating factors like supply voltage and temperature in our case. The calculation of the acceleration factors and the amount of aging time can be used to predict the amount of real-time aging, as shown in paper [3]. According to our calculation, for 65 nm technology node, aging acceleration for 9 h results in approximately 10 days of constant real-time use as shown in the paper [3].

During accelerated aging, the SoC was allowed to continuously operate and the PSRR data was collected every hour to determine the degradation of PSRR across time. The PSRR data collection setup is similar to that of the previous step as shown in Fig. 4. As described in "General Concepts of Transistor Aging", the HCI and BTI effect on the transistor during accelerated aging degrades the performance and other specifications for transistors like threshold voltage (v_{th}), transconductance (g_{m}), etc. This cumulatively degrades the DC PSRR of the LDO as described in Ref. [3]. A simplified equation of the PSRR is given as below:

$$\text{PSR} = \frac{v_{\text{out}}(s)}{v_{\text{dd}}(s)} = \frac{K}{(1+\frac{s}{\omega_o})(1+LG(s))}, \tag{3}$$

$$LG \propto g_{\text{m}}, \qquad g_{\text{m}} \propto v_{\text{eff}}, \qquad v_{\text{eff}} = v_{\text{gs}} - v_{\text{th}}, \tag{4}$$

where k is a constant, ω_o is the pole originating at the output of the LDO in Fig. 1a and LG is the loop gain of the LDO feedback loop. v_{gs} is the gate to source voltage and v_{th} is the threshold voltage of the transistors. We see that the LG is dependant on g_{m}, which again degrades with the deterioration of the v_{th} of the transistors. The PSRR data generated form accelerated aging of the LDO comprises the data set of recycled SoCs' PSRR responses, which were later used in our ML analysis. In Refs. [3, 4], accelerated aging on stand-alone LDOs were observed. Both the papers enlisted the degradation of several parameters and the effect of the same on the LDO's PSRR across hours. In this paper, we have investigated the same structure of an LDO but within an SoC. Thus, the aging response was similar in certain aspects yet dissimilar in other aspects compared to stand-alone LDOs. Compared to stand-alone LDOs (refer: Fig. 5), the amount of noise was more in the SoC (refer: Fig. 6) due to the increased number of modules and other switching activities occurring within the SoC. The spikes, appearing in Fig. 6 is proof of the variance that appeared across certain frequencies in the distribution. As it can be seen in Fig. 6, the variance increased much more between 1–1.5 MHz and 2–2.5 MHz for the PSRR in case of SoCs. But for stand-alone LDOs, the variance was almost similar across the frequency range, and no sudden spikes appeared. This kind of variance was mostly attributed by the process variation among the ICs and not due to frequency dependant noise. Apart from that, the initial degradation was more for stand-alone LDOs as we see that the maximum degradation occurred within the first two-four hours of the accelerated aging (refer Fig. 7a). For SoCs, we found the PSRR degrading consistently across eight hours, as can be seen in Fig. 7b.

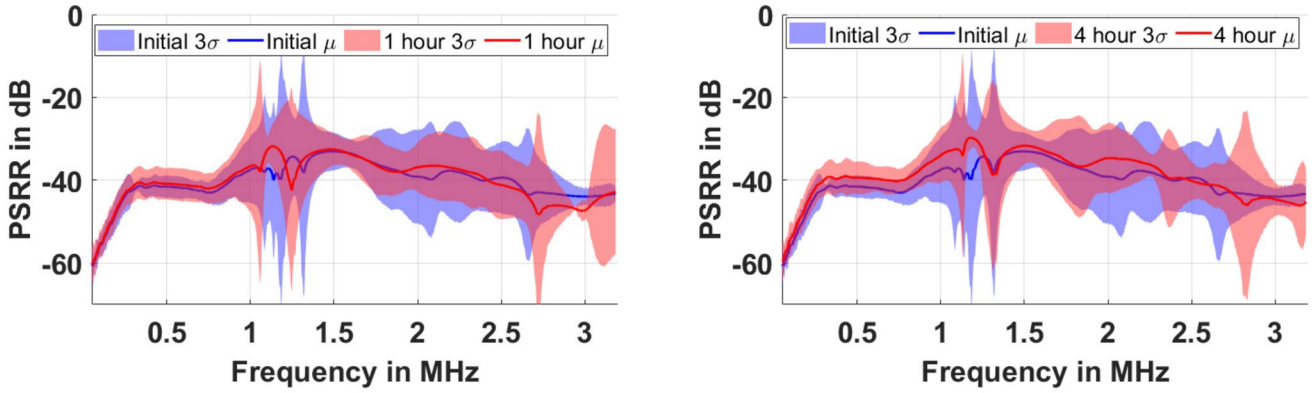

(a) SoC:1hr aging:Original distribution

(b) SoC:4hr aging:Original Distribution

Fig. 6 Silicon data showing PSRR degradation distribution of original SoCs for 1 and 4 h of accelerated aging

Fig. 7 Comparison between the aging profiles of stand-alone LDOs and LDOs embedded in SoC

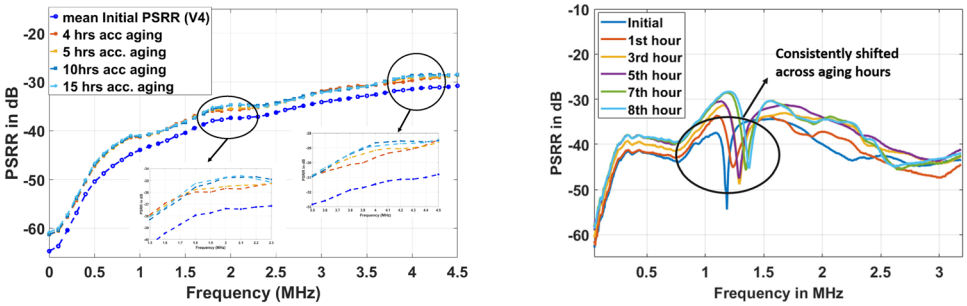

(a) Aging profile:standalone LDO :Vendor 4 (V4) [3]

(b) Aging profile:LDO embedded in SoC

Though this is difficult to explain without understanding the inherent design of the SoC and may vary from one SoC to another. An instinctive conclusion can be made by understanding that apart from the LDOs, other components may degrade within the SoC, affecting the PSRR degradation of the embedded LDO. There were also other consistent changes in the PSRR degradation across time, which were not that prominent for stand-alone LDOs. As can be seen in Fig. 7b, the peak at $\approx$ 1.2 MHz consistently shifted towards right across aging time for all the four SoC samples which we observed. This can be because of sudden changes in the parasitics due to the effect of aging, which was a striking feature we obtained from SoCs and were not that evident from stand-alone LDOs. For stand-alone LDOs, though, there were bumps on the PSRR due to accelerated aging, as seen in Fig. 7a, but the effect was not that consistent across aging hours as we saw in case of the SoCs.

Despite the dissimilarities mentioned above, the inherent aging principle was still similar for both stand-alone and embedded LDOs. As we can see in Figs. 5 and 6, the difference between the mean of new (μ_{new}) and aged (μ_{aged}) PSRR distribution increased with the increase in aging hours for both stand-alone LDOs and LDOs embedded in SoCs. It was seen that the mean difference, $\mu_{\text{new}} - \mu_{\text{aged}}$ for SoCs had an average of $\approx$ 1.5 dB across frequency for four hours of aging, while it was $\approx$ 0.14 dB across frequency for one hour of aging (refer: Fig. 6). For stand-alone LDOs also the difference in the mean PSRR across aging can be seen clearly (refer: Fig. 5). While this difference is higher for certain vendors like V1, V2, etc. it is comparatively smaller for many vendors like V3, V4, etc. Also, the DC PSRR shifted in LDOs embedded in SoC like the stand-alone LDOs following the same principle of aging phenomena as described in Eq. (3). Thus it may be concluded that the initial PSRR was different for each of the above cases, but the difference obtained between the initial PSRR and the resultant PSRR after aging bore similarities in both the cases. In other words, the trend of aging degradation was similar across aging hours, which can be modeled by ML algorithms. This similarity could be of great significance when we use the semi-supervised ML tool in our future analysis. The above similarities in the aging trend motivated us to observe the second type (Case 2) of semi-supervised training, which is explained in "Recycled Chip Classification Results and Discussion". In this scenario, we have used the difference between a new and an aged PSRR to train and test the classifiers. The idea behind such an experiment originates from the identical aging trends of stand-alone LDOs and that of

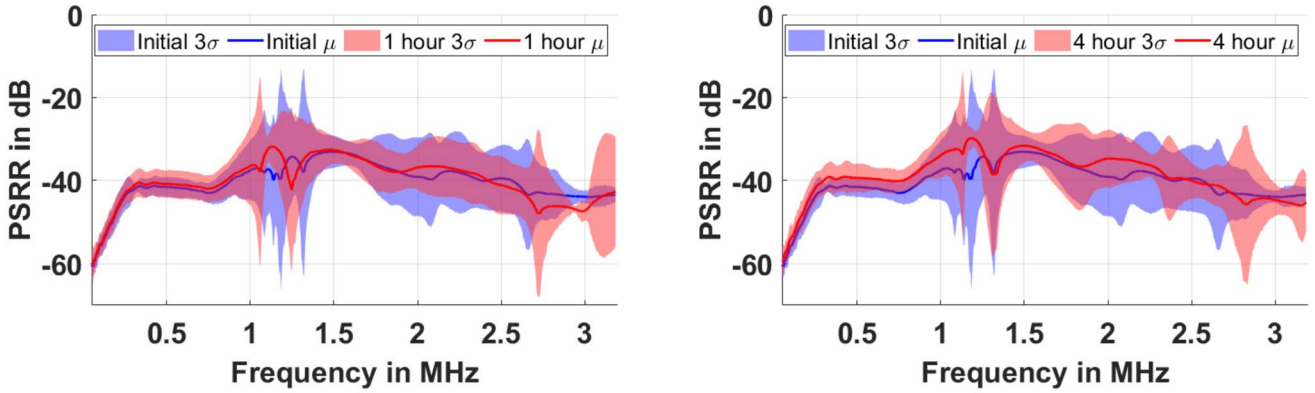

(a) SoC:1hr aging:Augmented Distribution

(b) SoC:4hr aging:Augmented Distribution

Fig. 8 Augmented distribution of the aging profile of 48 chips obtained by applying PCA to original SoC aging distribution shown in Fig. 6

LDOs within SoCs, which can be manifested in the difference of new and aged PSRRs of both LDOs and SoCs. In other words, the initial PSRR or the aged PSRR values of a stand-alone LDO may vary from that of the LDO within an SoC. But, the difference between the initial and aged PSRR is similar due to the similar physical aging phenomenon that impacts both the structures.

After the data collection, a set of ML tools were developed to analyze and aid the detection of recycled SoCs. As described earlier, three different ML tools were developed; namely, (i) Supervised ML tool using k nearest neighborhood (KNN) algorithm. (ii) Unsupervised ML tool using k-means clustering algorithm (iii) Semi-supervised ML tool again using the KNN algorithm. A detailed analysis of the accuracy along with the advantages and disadvantages of the above algorithms are provided in the following section.

Recycled Chip Classification Results and Discussion

We used the data and the experimental setup discussed in "Experimental Setup and Aging Analysis" and conducted ML analysis. Before elaborating on our experimental results, we stress that the constraints linked to the number of available SoC samples, as mentioned earlier, can be overcome by applying data augmentation techniques. In this context, standard, commonly-used data augmentation techniques improve the diversity of data available for training ML models. Thus, we followed a method of data augmentation to strengthen our framework, as explained below.

Data Augmentation

We used a widely applied approach, namely principal component analysis (PCA), to augment our data set. PCA is a linear orthogonal transformation of a data set to a new set of coordinates, where the first coordinate is the greatest variance of the scalar projection of the data set, the second coordinate is the second greatest variance, and so on. PCA analyzes any data set and provides n eigenvalues to represent the principal components of the data set. Out of those, k values were kept similar to that of the original data set, whereas the $n - k$ values were replaced and randomly chosen from a Gaussian distribution to represent the augmented data set. This type of augmentation fitted our needs since we were majorly looking for the same SoC data only with a wider range of process variation than that obtained from the original data set of four SoCs. In this way, a set of 48 augmented PSRR data set was derived from 4 original chips. It must be noted that the augmented distribution was verified with respect to the original distribution, and both appeared similar with respect to statistical properties.

Since we augmented the data set of PSRR from a specific SoC of a particular vendor, we took eigenvalues with the higher variance as the intrinsic aging property, which would remain similar for all the SoCs. Since we had the PSRR data for each hour of aging across a range of 600 frequency points, each separated by 5 KHz, the total number of features was large enough for PCA analysis. To prevent information loss, we randomly altered the lesser varying eigenvalues to reflect the changes that can occur due to process variations. Since our entire data set represented the same unit (dB vs. MHz), there was no requirement of data standardization as mostly required for applying PCA analysis. The distribution of augmented data of 48 SoCs obtained from using PCA to the original four SoCs can be seen in Fig. 8. It can be seen that it is similar to the distribution obtained for the original four SoCs, as seen in Fig. 6.

It must be noted that for SoC recycled detection, we combined the original data from four SoCs and the augmented data from 48 SoCs to generate our final data set for the

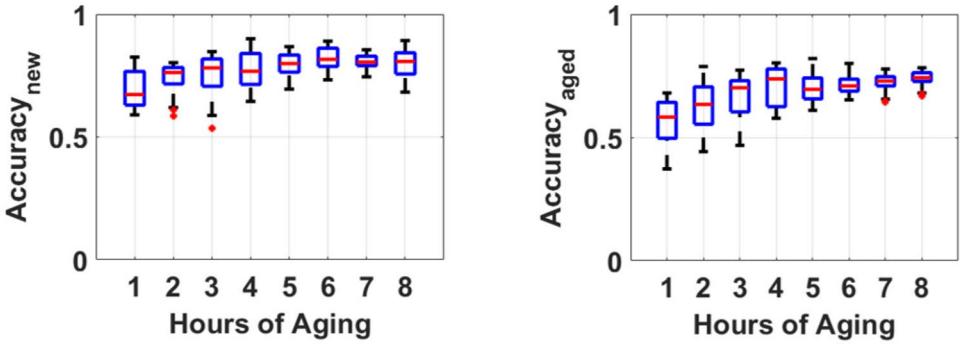

Fig. 9 Supervised classification: detecting accuracy of new and aged SoCs over hours of aging by applying KNN classifiers to our dataset [dot: mean (μ) ; bar: standard deviation (3σ)]

experiments. Thus, we have in total 52 SoCs on which the ML analysis was performed. The sequence of study can be divided into the following categories.

(1) *Supervised classification*: In this type of classification, we provide the PSRR recorded from a subset of the SoCs with their labels, i.e., new or aged, and the algorithm tries to learn the age of the remaining samples of the SoCs and classify them as either new or recycled. Thus training over a small sample set of the SoCs, it is tested whether the tool can generalize the learning for the entire subset of the SoCs.

Results for Supervised Classification

For supervised classification, the labels of PSRR data from a small set of the SoCs, including new and aged ones, are provided to the tool for supervised learning. We ran KNN algorithm package provided by the Matlab software, and the value of k is set as five by the automatic classifier as the best possible parameter. The results for supervised classification are shown in Fig. 9, showing the accuracy of detecting aged and new SoCs, respectively. While calculating the classification accuracy, we use ten-fold cross-validation methods, and the average is reported along with the standard deviation. In simple terms, the results reported in Fig. 9, is the average of the accuracy obtained by running the algorithm ten times for ten sub-sets of the data. In each round of the experiment, the tool is trained over nine parts (sub-sets, or so-called folds) of the data set, whereas the testing is conducted on the remaining one part of the data set. In this way, it is confirmed that the training and suspect data are chosen uniformly across the data set, and there is less bias while computing the accuracy.

The main conclusion from the results is that the model which is extracted from the PSRR of a given set of SoCs can be used to predict the age of the other SoCs up to a maximum accuracy of $\approx$ 90% for new SoCs and $\approx$ 83% for aged/old SoCs. While explaining the results, it must be understood that the average detection accuracy improves with the increase in aging time, as can be seen in the Fig. 9. This trend can be seen for the detection accuracy of both the new and aged (recycled) SoCs. The reason for such an improvement is intuitive and can be obtained from the discussion in "Experimental Setup and Aging Analysis" pertaining to the Fig. 6. As we see, the difference between the mean PSRR of the new chip distribution and the PSRR of the aged chip distribution ($\mu_{\text{new}} - \mu_{\text{aged}}$) increases with the accelerated aging time; it becomes easier for the classifier to distinguish a new chip compared to a recycled one. In conclusion, as the aging time increases, the new chip PSRR distribution separates from the aged chip PSRR, and the classifier can detect the difference even across the process variation and the measurement noise. For only one hour of accelerated aging, the mean difference between the new and aged PSRR is approximately 0.14 dB which can be easily mistaken by the classifier as process variation, but with an increase in the difference, which is about 1.5 dB after four hours of aging, the classifier can distinguish this difference over process variation. The same trend of detection accuracy was seen for supervised classification for stand-alone LDOs in Ref. [3], where we observe the average accuracy increasing with the aging hours for a few vendors (V1 and V2). But for other vendors (V3 and V4) of LDOs in Ref. [3], the detection accuracy was more or less constant across aging hours since, the degradation of PSRR saturated after approximately three hours of accelerated aging. But for the case of LDOs within the SoC observed in this paper, we saw a continuous degradation of PSRR across aging hours till the fifth hour of aging, as seen in Fig. 7b. Thus, the accuracy of detection also saturates after the fifth hour of aging, as seen in Fig. 9.

(2) *Unsupervised classification*: In this setting, no label is provided to the algorithm, i.e., only PSRR data from one SoC (new/aged) is chosen by the SME to be provided to the algorithm. This is considered as the golden sample. While calculating the accuracy, the PSRR from unseen SoC samples is provided to the algorithm. If the unseen LDO and the golden sample are of a similar age, they should be categorized in the same cluster as the tool. If there are differences in the age of the golden and the unknown sample, then the algorithm should be able to assign one of them to the new cluster and the other to the recycled cluster. Two types of

approaches can be followed in this clustering algorithm. In *Case 1*, we only measure the PSRR of the component provided to us. Then we give the measured PSRR values to the tool in a pair-wise manner (i.e., golden and suspect). The suspect PSRR is further determined as new or recycled by the *k*-means algorithm. In *Case 2*, we measure the initial PSRR of the unknown component and then follow a procedure of artificial aging for either 1 or 4 h. These measured PSRR values, along with the initial ones, are provided to the algorithm. This case can be applied only when additional aging can be performed.

Results for Unsupervised Clustering

As discussed earlier, no labels are provided to the tool during unsupervised clustering using *k*-means algorithm and PSRR data from one golden SoC (new or old) along with that of the suspect SoC is provided to the tool. We use the *k*-means function embedded in the Matlab software package for this clustering. During clustering with *k*-means, we apply the Silhouette method to validate the consistency of the clusters, thus increasing our accuracy. We also use the re-sampling technique to find lower local minima of the Euclidean distances between the examples. For this purpose, the centroids, which are determined by running the *k*-means algorithm once, are again utilized while re-running the algorithm. This helps to tackle the noisy samples and provide better accuracy in such cases. The results for *k*-means clustering is presented in Fig. 10. As can be seen, the results are divided into two cases. In this figure, Case 1 and Case 2 refer to what has been explained above. In brief, in Case 1, only the initial PSRR measured from the golden sample, and the suspect sample are provided to the algorithm. But, in Case 2, the SME is allowed to perform artificial aging on both these samples. We have selected one hour and four hours of artificial aging due to the sufficient separation observed between the initial PSRR values and ones collected after one hour and four hours of aging (see Fig. 5). To conduct experiments corresponding to Case 2, the aged PSRR data of both the golden and suspect samples, along with the data provided in Case 1, are fed into the algorithm to improve the learning accuracy further. In this case, we offered the 1-h aged data and also the 4-h aged data separately to the tool. The accuracy for each of the above instances is calculated for components of various ages, as depicted in Fig. 10. In this figure, the *x*-axis shows the minimum age of the SoCs under test. Here if the age equals zero, it refers to the initial data from the sample. Also, the figures show the accuracy of clustering the unseen samples as either new or aged. The maximum accuracy for Case 1 was approximately 74%, while that for Case 2 was approximately 73%.

The results obtained from unsupervised classification provide an accuracy better than random, but the average accuracy is much degraded compared to the supervised classification. The major reason for such degradation of accuracy is due to the fact that the amount of noise and process variation is overpowering the PSRR degradation in this case. Since no labels are provided, it is difficult for the classifier to distinguish the new and recycled SoCs with only a single PSRR (aged/new) as a reference. The high variability across samples which is projected by the standard deviation (3σ) in Fig. 6 is a major roadblock for unsupervised clustering in these data sets. Similar challenges were also seen in the case of the stand-alone LDOs in Ref. [3], where the detection accuracy was also a maximum of 74% for Case 1. But the detection accuracy improved to a maximum of 86% with Case 2. In other words, providing the aged PSRR of the suspect stand-alone LDOs during testing helped the classifier to obtain better accuracy. For, recycled SoC detection, providing the aged PSRR of the LDO within the suspect SoC, while training did not improve the accuracy, rather, in some cases, worsened the detection. This can be understood by the spikes and irregularities that are seen in the data set, which clearly indicates that providing the aged version of the PSRR while testing can sometimes worsen the detection accuracy. In Fig. 5, the noise/process variation profile for stand-alone LDOs was similar across both the new and aged PSRR distribution. There were no irregularities obtained in the distribution with aging that was not initially present in the new chip data set. Thus, when both the new and aged PSRR was provided during testing, the process variation and noise got nullified to a large extent, which helped the algorithm to cluster correctly. But, for an SoC which is continuously active during the aging time, the clock, switching, and other activities caused differences in the noise profile of the new and aged PSRR. Thus, the aged distribution can have certain irregularities that were not initially present in the new SoC distribution. As a result, providing the aged PSRR of the suspect SoC, along with the initial PSRR during testing in Case 2, was not able to reduce the noise or help the algorithm to provide better accuracy.

(3) *Semi-supervised classification*: The classification format is much similar to the supervised classification described above. It uses the same strategy and algorithm, but there is a striking difference in the training set. In this case, we only provide the algorithm the PSRR recorded (new and/or aged) from stand-alone LDOs collected beforehand. The algorithm is then provided with unknown PSRR from the current data set of 52 SoCs for classification as new or recycled. This type of classification is extremely beneficial for our case, where the requirement of golden data is relaxed. Also, the fact that measured PSRR data from stand-alone LDOs can be used to classify PSRR

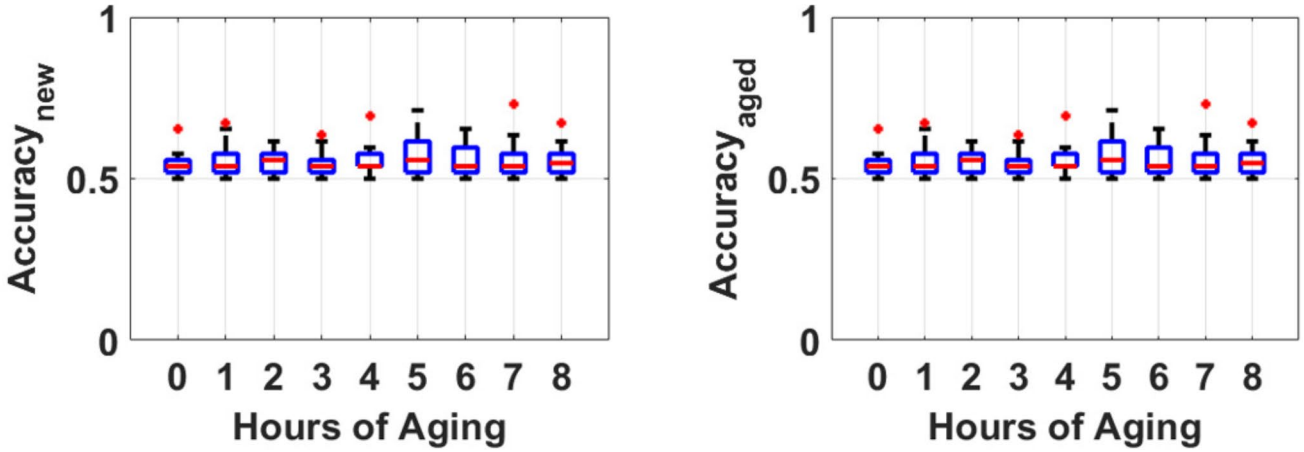

(a) Kmeans: Case 1: Tested on initial SoC PSRR

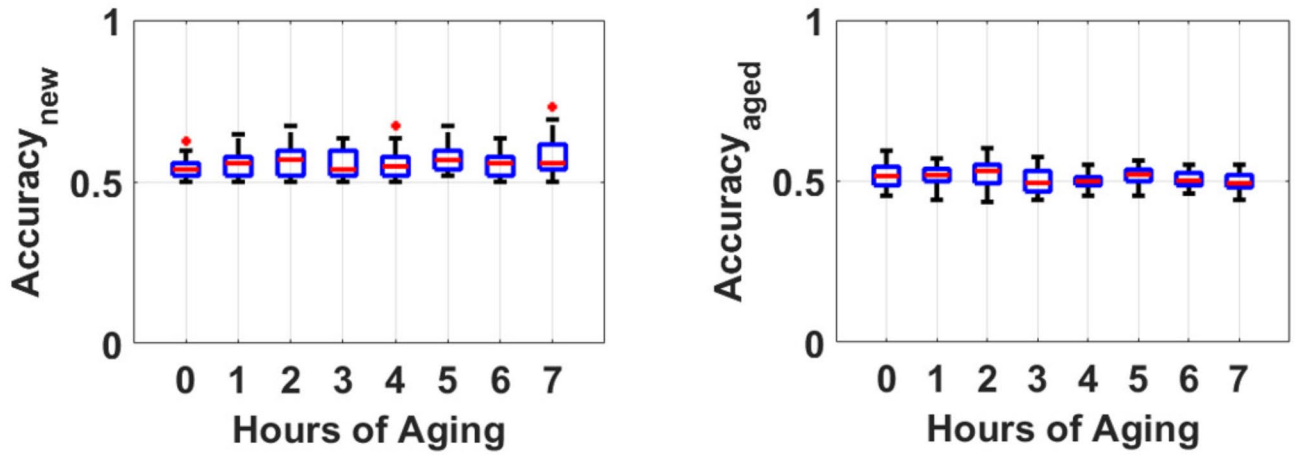

(b) Kmeans: Case 2: Tested on initial and 1 hour aged SoC PSRR

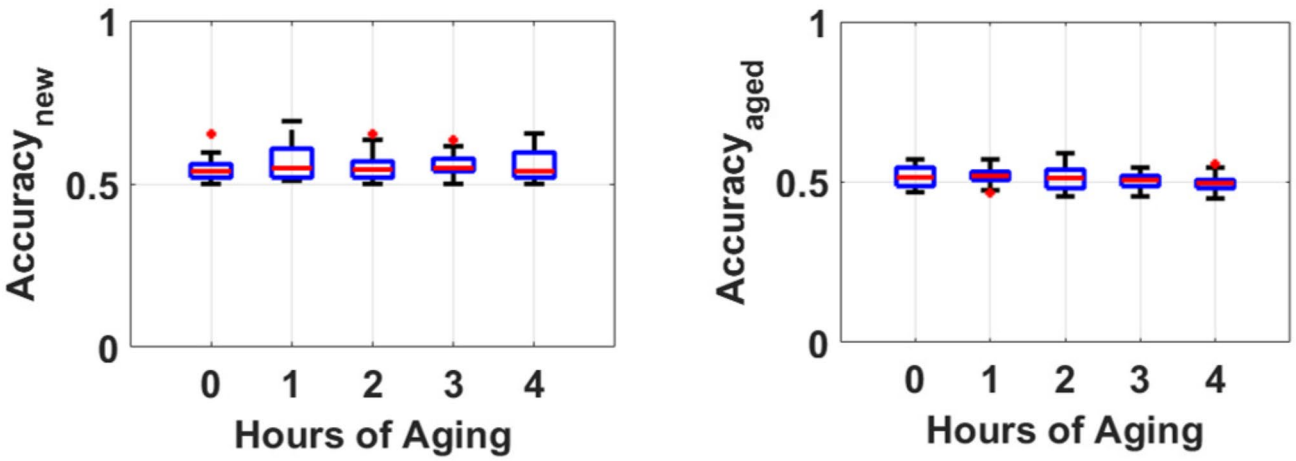

(c) Kmeans: Case 2: Tested on initial and 4 hours aged SoC PSRR

◂**Fig. 10** Unsupervised classification: Detecting accuracy of new and aged SoCs over hours of aging by applying *k*-means clustering for both Case 1: initial PSRR provided for both golden and unseen sample and Case 2: initial and synthetically aged (1 h and 4 h) PSRR provided for samples. [dot:mean (μ) ; bar: standard deviation (3σ)]

data obtained from embedded LDOs in SoC can bolster the detection procedure of recycled SoCs.

Results for Semi-supervised Classification

An obvious question that can be asked is whether the tool constructed above for detecting recycled SoCs using supervised classification can be generalized across other stand-alone LDO ICs available commercially. This is an important aspect since, most of the time, golden data required for supervised classification is not available, and it is easier to procure PSRR data from stand-alone LDOs compared to LDOs embedded in SoCs. Thus, if the supervised ML tool can be generalized to train on other stand-alone LDO PSRR data (aged and/or new), then the entire process can be hugely simplified. For this purpose, we utilized the data we have collected for our previous experiments reported in Ref. [3]. We collected PSRR data from 32 LDOs across four vendors over an aging experiment for over nine hours. The data was recorded hourly after performing aging, similar to our procedure described in "Experimental Setup and Aging Analysis". We utilized this data for training the algorithm. In this regard, we carried out two experiments, referred to Case 1 and Case 2, explained below.

Case 1 First, we extracted classifiers from data sets containing a new and an aged (e.g., aged for one hour) sample of stand-alone LDOs of one vendor at a time. For each vendor, the classifier model obtained was tested to categorize PSRR of embedded LDOs within SoCs that are used in this paper. In doing so, the accuracy of training on each vendor was computed to determine how well the classifier can categorize an SoC as new or aged. Figure 11 illustrates the classification accuracy for both new (aging hour shown as zero) and aged SoCs using model trained on each of the vendors (V1–V4) of stand-alone LDOs. It can be concluded that the classifier obtained can categorize aged (recycled) SoCs correctly up to a $\approx$ 97%, 98%, and 92% after training on V1, V2, and V3, respectively. However, the classification accuracy for detecting a new SoC, in this case, is not high.

In this experiment, the classifiers' inability to classify a new SoC is intuitive as the initial PSRR data is different. Thus, with no aging, the PSRR data obtained from a new SoC can be much different than that of the stand-alone LDO. But, the similarities in the aging trend observed in our prior "Experimental Setup and Aging Analysis" inspired us to provide the difference in the new and aged PSRR of the samples during both training and testing to improve the detection accuracy of new SoCs.

Case 2 Inspired by the above observation, in an attempt to simultaneously improve the classification accuracy for new and aged SoCs, we conducted another experiment (Case 2) as follows. In the training phase, the ML algorithm was given a dataset composed of the differences between the initial PSRR values measured from a new and an aged stand-alone LDOs as well as their corresponding measurements collected after 1 h of synthetic aging. To classify the SoCs, this trained model was applied to the differences between initial PSRR values and ones collected after 1 h of synthetic aging. In this case, as can be seen in Fig. 12, the classification accuracy was improved significantly (up to 96%), when the stand-alone LDOs from V1 was taken into account[3]. Note that this positive result was achieved solely for V1, which was the same manufacturer as that of the SoCs used in our experiments. This further highlights the fact that if stand-alone LDOs and SoCs do not have similarities in terms of the initial PSRR classification of the new SoCs based on stand-alone LDOs may not be feasible. Nevertheless, according to results depicted in Fig. 12, unfortunately, the classification accuracy is not acceptable for the aged SoCs.

Finally, we stress that the classification accuracy for aged SoC is satisfactory in both of the cases (Case 1 and Case 2) when the model is trained using the measurements collected from stand-alone LDOs produced by V2, V3, and V4. For the latter vendor, in Case 2, the average maximum accuracy was improved to $\approx$ 96% from $\approx$ 77% that was reported for Case 1. Nonetheless, a key conclusion that we can draw from our experiments (Case 1 and Case 2) is that the ML models trained to conduct semi-supervised learning can distinguish either the new or the aged SoCs, which limits their applicability in practice.

Summary of Results

In previous sections, we provided the advantages and limitations of the application of ML tools to distinguish recycled SoCs. The most straightforward approach would be to model a classifier with golden new and recycled SoC PSRR data using a supervised ML tool. This provides maximum accuracy of 90% but limits the algorithm with the requirement of golden samples. This limitation can be removed entirely with unsupervised clustering methods, which requires no labels during training, but solely a golden sample (new or aged). However, the accuracy degrades and reduces to the maximum accuracy of 74%, which is lower compared to

[3] Interestingly enough, if the synthetic aging was conducted for 4 h, none of the trained models obtained for our four vendors was useful to classify a new SoC.

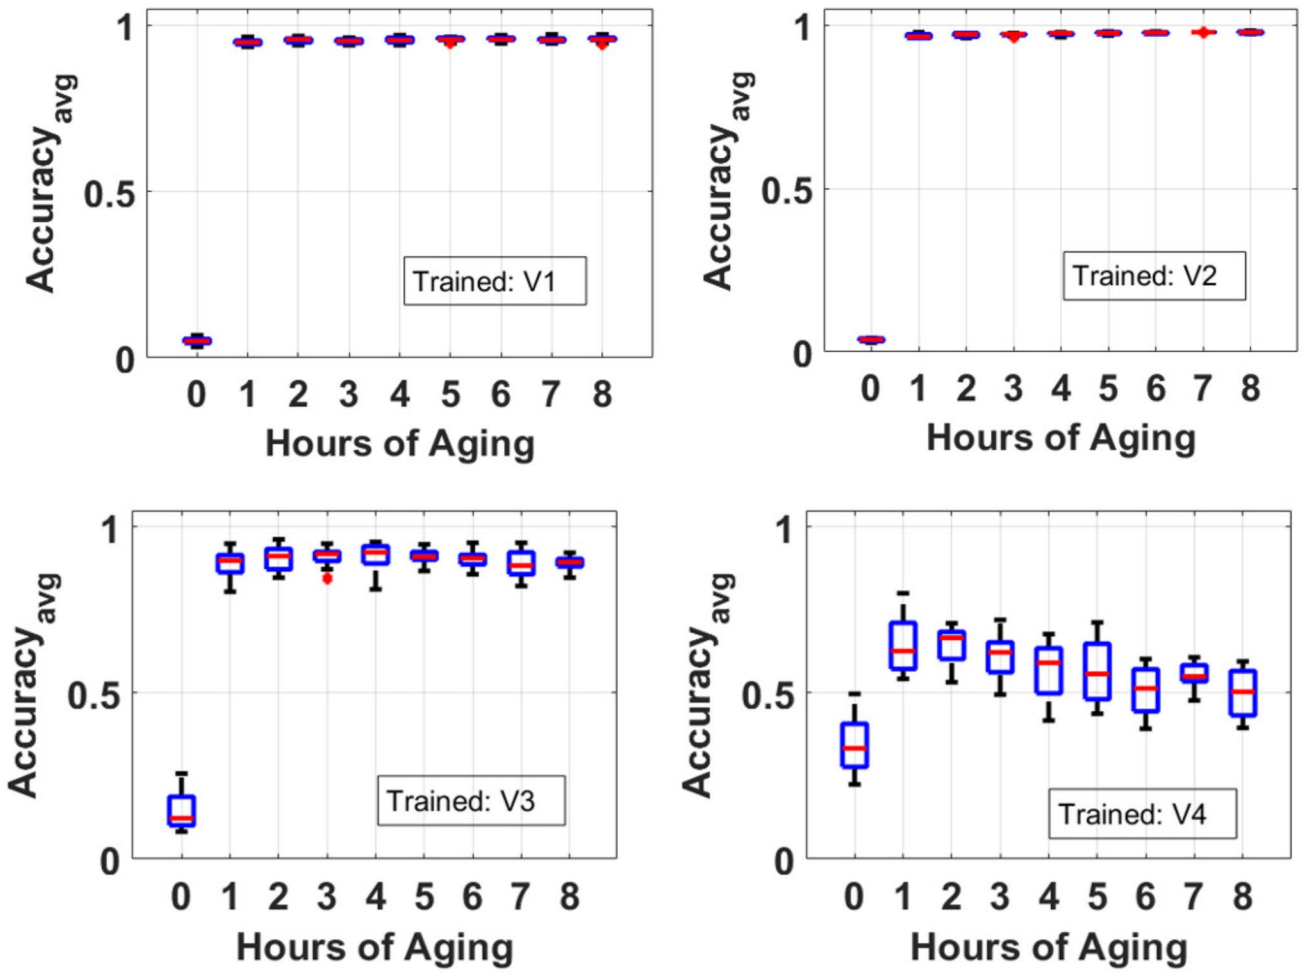

Fig. 11 Semi-supervised classification: case 1: detecting accuracy of unseen SoCs over hours of aging by applying KNN classifiers to data set containing standalone LDOs of four different vendors (V1, V2, V3, V4) [dot:mean (μ) ; bar: standard deviation (3σ)]. Testing was done only on initial SoC PSRR

supervised classification. To offer a trade-off between golden data requirement and accuracy, we propose semi-supervised classification. This type of classification partially removes the constraint of golden data. Here, the classifier is trained on PSRR samples of stand-alone LDOs from four different vendors [3]. Compared to PSRR of embedded LDOs in SoCs, PSRR from stand-alone LDOs are simpler to record and more comfortable to obtain. This improves the accuracy of recycled SoC detection, up to 98%, in our various experiments after training on the data collected from stand-alone LDOs from all the four vendors. Although it is tempting to apply this type of detection, we observe that the trained models cannot simultaneously improve the classification accuracy for new and aged SoCs. Hence, our final conclusion is that the ML models trained to conduct semi-supervised suffer from this limitation, which restricts their applicability in practice.

A detailed comparison of the performance of the proposed recycled SoC detection and recycled LDO detection [3] is provided in Table 2, discussing the advantages and limitations of each procedure. We agree that existing methods for recycled detection rely on golden data, and the proposed method provided also requires some form of golden data, which is a drawback. But the major highlight of the proposed method is its applicability to analog, AMS, and legacy ICs, which are not fulfilled by existing detection methods. The analogy provided in the paper for the advocated method clearly states the requirement of golden data for supervised, semi-supervised, and unsupervised ML techniques. In comparison to the recycled detection using stand-alone LDOs in Ref. [3], the proposed recycled SoC detection suffers many challenges. But, the challenges are justified if we observe the complications that arise with an SoC. The inability of the semi-supervised algorithm to distinguish between new and aged SoCs simultaneously, is a major

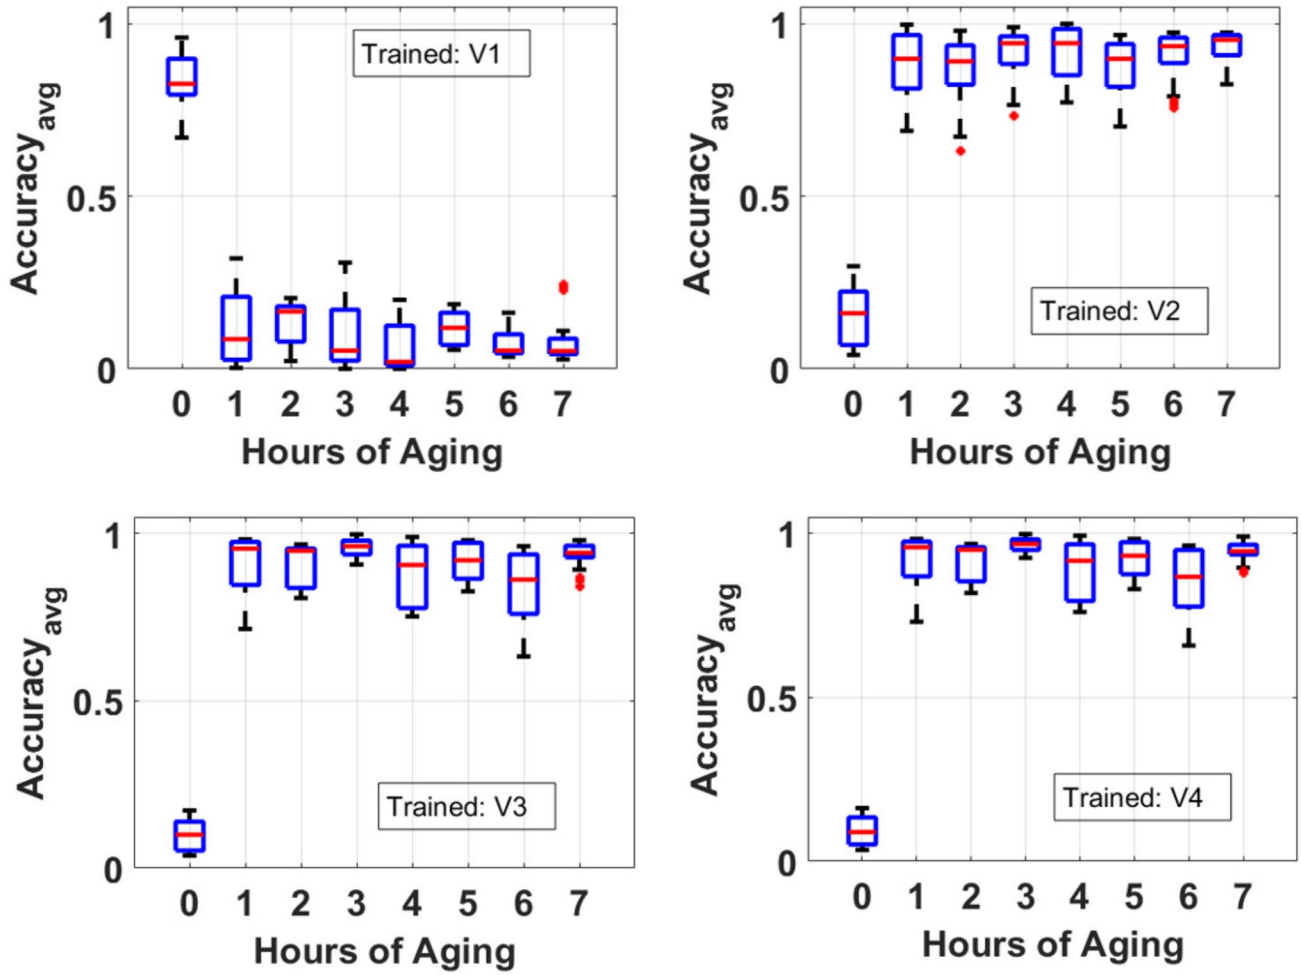

Fig. 12 Semi-supervised classification: case 2: detecting accuracy of unseen SoCs over hours of aging by applying KNN classifiers to data set containing stand-alone LDOs of four different vendors (V1, V2, V3, V4) [dot:mean (μ) ; bar: standard deviation (3σ)]. In the training phase, the initial PSRR values in conjunction with ones collected after 1 h of synthetic aging is given to the algorithm. Testing is performed using the initial SoC PSRR values combined with ones measured from 1 h aged SoCs

challenge we faced. The performance of semi-supervised algorithm could have been improved if different vendors of SoCs with similar specifications were obtained. Though, it is simpler to obtain multiple LDOs from different vendors with similar specifications, footprints, and electrical parameters; it is equally difficult to achieve the same for SoCs, where the performance metrics are unique for every SoC. Nonetheless, the proposed recycled SoC detection approach performs with equivalent accuracy for the supervised and unsupervised techniques, when compared to the stand-alone LDO approach.

Conclusion and Future Work

In this paper, we proposed a universal approach for recycled SoC degradation that relies on the power delivery network available in most SoCs. The degradation of LDO embedded within an SoC can be detected using ML algorithms. Application of supervised and unsupervised ML algorithms provides maximum accuracy of 90% and 74%, respectively. To alleviate the limitation of the golden sample requirement in supervised learning and to improve the accuracy of unsupervised clustering, we employ a semi-supervised learning approach. Though, the semi-supervised method was successful in detecting stand-alone recycled LDOs in our previous works, it failed to detect a new and a recycled SoC simultaneously in the proposed approach. The semi-supervised algorithm was only able to detect either a new or a recycled SoC, which increases

Table 2 Comparison of the maximum accuracy across all the samples. Advantages and limitations of recycled IC detection using stand-alone LDOs [3] and proposed recycled SoC detection using LDOs embedded in SoCs are further summarized

Type and specs	Supervised		Unsupervised				Semi-supervised		
	LDOs	SoCs	LDOs		SoCs		LDOs	SoCs (V*: vendors of stand-alone LDOs)	
$Accuracy_{new}$	≈ 97%	≈ 90%	Case 1: ≈ 74%	Case 2: ≈ 70%	Case 1: ≈ 74%	Case 2: ≈ 73%	≈ 93%	Case 1: ≈ 5–50% across V1–V4	Case 2: ≈ 96% for V1 ≈ 20–30% for V2–V4
$Accuracy_{aged}$	≈ 96%	≈ 83%	Case 1: ≈ 70%	Case 2: ≈ 85%	Case 1: ≈ 74%	Case 2: ≈ 63%	≈ 98%	Case 1: ≈ 96% for V1–V3 ≈ 80% for V4	Case 2: ≈ 32% for V1 > 90% for V2–V3 ≈ 80% for V4
Advantages	Good accuracy for both new and recycled ICs	Good accuracy for both new and recycled SoCs	Golden data not required		Golden data is not required		Reduces golden data requirement	Reduces golden data requirement	
	Accuracy improves with aging time	Accuracy improves with aging time	Accuracy improves with Case 2				Good accuracy for both new and recycled detection		
							Suitable for stand-alone LDOs		
Limitations	golden sample required	golden sample required	Lesser accuracy		Lesser accuracy		Requires some golden data but can be obtained from other vendors	Requires some golden data Unable to distinguish new and aged SoCs simultaneously Increased false detection	

risks of false detection and restricts its applicability for recycled SoC detection. Aging degradation of transistors is helpful for recycled counterfeit detection and has been used extensively for recycled detection in literature. But, it is detrimental for performance and yield. Thus in the future, smart LDO designs can be implemented for new IC/SoC designs that aids recycled detection but do not affect overall chip performance with its aging degradation. In addition, degradation of other metrics of LDOs like transient line/load regulation or efficiency can also be investigated to facilitate recycled detection.

Funding This study was funded by AFOSR award number FA9550-14-1-0351 and NSF award number 1610075.

Compliance with ethical standards

Conflict of interest Prof. Rama Chellappa was in dissertation committee of Prof. Domenic Forte and Prof. Sartaj Sahni is a professor in CISE department in UF.

References

1. Alam M, Chowdhury S, Park B, Munzer D, Maghari N, Tehranipoor M, Forte D. Challenges and opportunities in analog and mixed signal (AMS) integrated circuit (IC) security. J Hardw Syst Secur. 2018;2(1):15–32.
2. Bernstein JB. Chapter 3—failure mechanisms. In: Reliability prediction from burn-in data fit to reliability models. Cambridge: Academic Press; 2014.
3. Chowdhury S, Ganji F, Bryant T, Maghari N, Forte D. Recycled analog and mixed signal chip detection at zero cost using LDO degradation. In: 2019 IEEE international test conference (ITC); 2019.
4. Chowdhury S, Shen H, Park B, Maghari N, Forte D. Aging analysis of low dropout regulator for universal recycled IC detection. In: 2019 IEEE computer society annual symposium on VLSI (ISVLSI); 2019.
5. Dabiri F, Potkonjak M. Hardware aging-based software metering. In: Proceedings of the conference on design, automation and test in Europe, pp. 460–465. European Design and Automation Association; 2009.
6. Ghahramani Z. Unsupervised learning. In: Advanced lectures on machine learning, pp. 72–112. Heidelberg: Springer; 2004.
7. Graves A. Supervised sequence labelling. In: Supervised sequence labelling with recurrent neural networks, pp. 5–13. Heidelberg: Springer; 2012.
8. Guin U, Dimase D, Tehranipoor M. Counterfeit integrated circuits: detection, avoidance, and the challenges ahead. J. Electron. Test 2014;30:9–23. https://doi.org/10.1007/s10836-013-5430-8
9. Guin U, Forte D, Tehranipoor M. Anti-counterfeit techniques: from design to resign. In: 2013 14th Intl. Wkshp. on microprocessor test and verification; 2013
10. Guin U, Forte D, Tehranipoor M. Design of accurate low-cost on-chip structures for protecting integrated circuits against recycling. In: IEEE transactions on very large scale integration (VLSI) systems; 2016
11. Guin U, Huang K, DiMase D, Carulli JM, Tehranipoor M, Makris Y. Counterfeit integrated circuits: a rising threat in the global semiconductor supply chain. Proc. IEEE. 2014;102(8):1207–28.
12. Guin U, Zhang X, Forte D, Tehranipoor M. Low-cost on-chip structures for combating die and IC recycling. In: Proceedings of the 51st annual design automation conference; 2014.
13. Guo Z, Rahman MT, Tehranipoor MM, Forte D. A zero-cost approach to detect recycled SoC chips using embedded SRAM. In: 2016 IEEE international symposium on hardware oriented security and trust (HOST); 2016.
14. Guo Z, Xu X, Rahman MT, Tehranipoor MM, Forte D. SCARe: an SRAM-based countermeasure against IC recycling. IEEE Trans. Very Large Scale Integr. VLSI Syst. 2017;26(4):744–55.
15. Herder C, Yu M, Koushanfar F, Devadas S. Physical unclonable functions and applications: a tutorial. In: Proceedings of the IEEE; 2014.
16. Huang K, Liu Y, Korolija N, Carulli JM, Makris Y. Recycled IC detection based on statistical methods. In: IEEE transactions on computer-aided design of integrated circuits and systems; 2015.
17. Lee BS. Understanding the terms and definitions of LDO voltage regulators. http://www.ti.com/lit/an/slva079/slva079.pdf (1999). Accessed Sept 2020.
18. Maes R, Rozic V, Verbauwhede I, Koeberl P, van der Sluis E, van der Leest V. Experimental evaluation of physically unclonable functions in 65 nm CMOS. In: 2012 Proceedings of the ESSCIRC; 2012.
19. Mutschler AS. Transistor aging intensifies at 10/7 nm and below. https://semiengineering.com/transistor-aging-intensifies-10nm/ (2017). Accessed Sept 2020.
20. Pithadia S, Lester S, Verma A. LDO PSRR measurement simplified. https://www.ti.com/lit/an/slaa414a/slaa414a.pdf?ts=1600199264375 (2017). Accessed Sept 2020.
21. Sanchez-Sinencio E, Rudiak J. The tradeoffs of low dropout (LDO) voltage regulator architectures and the advantages of "Capless" LDOs. https://www.design-reuse.com/articles/45476/capless-low-dropout-ldo-voltage-regulator-architectures.html (2019). Accessed Sept 2020.
22. Sapatnekar SS. What happens when circuits grow old: aging issues in CMOS design. In: 2013 International symposium on VLSI design, automation, and test (VLSI-DAT) (2013)
23. Vaisband IP. Power delivery and management in nanoscale ICs. Ph.D. Thesis, University of Rochester; 2015.
24. Zhang X, Tehranipoor M. Design of on-chip lightweight sensors for effective detection of recycled ICs. In: IEEE transactions on very large scale integration (VLSI) systems; 2014.
25. Zhang X, Xiao K, Tehranipoor M. Path-delay fingerprinting for identification of recovered ICs. In: 2012 IEEE international symposium on defect and fault tolerance in VLSI and nanotechnology systems (DFT); 2010.
26. Zheng Y, Basak A, Bhunia S. CACI: Dynamic current analysis towards robust recycled chip identification. In: 2014 51st ACM/EDAC/IEEE design automation conference (DAC); 2014.