



Reliability improvement of SRAM PUFs based on a detailed experimental study into the stochastic effects of aging

A. Santana-Andreo^a, P. Saraza-Canflanca^b, R. Castro-Lopez^{a,*}, E. Roca^a, F.V. Fernandez^a

^a Instituto de Microelectrónica de Sevilla (IMSE), Universidad de Sevilla – CSIC, Seville, Spain

^b imec, Kapeldreef 75, 3001, Leuven, Belgium

ARTICLE INFO

Keywords:

Aging
Bias Temperature Instability (BTI)
Non-Conductive Hot-Carrier Injection (NC-HCI)
Physical Unclonable Function (PUF)
Reliability
Static Random-Access Memory (SRAM)

ABSTRACT

Physical Unclonable Functions (PUFs) have gained attention as a lightweight hardware security primitive. In particular, the SRAM-based PUF uses the unpredictable power-up value of the cells within an SRAM. Although these values should ideally be always the same within each SRAM to accomplish a correct PUF operation, this is often not the case, especially when factors like circuit aging are considered. While certain studies explore the effects of aging on SRAM PUFs, they often simplify the analysis. For instance, some studies assume that only Bias Temperature Instability (BTI) contributes to circuit degradation while others evaluate the overall degradation without accounting for the stochastic effects of aging on each individual cell. In this work, we first perform a detailed characterization of the nature of aging in SRAM PUFs, demonstrating that the impact of Non-Conductive Hot-Carrier Injection cannot be neglected. We also show that different cells degrade differently, highlighting the importance of accounting for the stochasticity of aging. After that, a method based on the Data Retention Voltage metric to select the cells with the most stable power-up response is introduced. Using these cells to generate the PUF identifier will result in a more stable response, and thus a better PUF performance.

1. Introduction

The field of cybersecurity continues to evolve, introducing new solutions to address the challenges posed by our increasingly interconnected world with a growing number of potential attack vectors. Hardware security, specifically Root-of-Trust (RoT) systems, forms the bedrock of any effective software security protocol. An interesting solution in this context is Physical Unclonable Functions (PUFs). PUFs are utilized to generate unique identifiers for each device, serving as a RoT in encryption schemes. PUFs capitalize on the inherent, unpredictable variations that occur in electronic circuits during fabrication to provide these distinctive identifiers. Typically, PUFs employ straightforward circuitry and generally do not require anti-tamper protection since physically tampering with PUFs is exceedingly challenging without altering the characteristics from which the RoT is derived. Consequently, PUFs are considerably more cost-effective than comparable solutions, such as non-volatile memories, making them especially appealing for resource-constrained applications like the Internet of Things. Numerous silicon PUF implementations have been proposed in the literature [1], including examples like ring oscillator PUFs [2],

arbiter PUFs [3], and SRAM PUFs [4].

Nowadays, SRAM cells can be found in most integrated circuits. SRAM PUFs can repurpose some of these general-purpose SRAM cells to generate the unique chip identifier. This makes SRAM PUFs one of the most practical PUF implementations proposed to date. As a result, they have garnered significant attention in academia [4] and have seen widespread adoption in industry [5]. The SRAM PUF identifier is generated by powering up a set of SRAM cells. Variability, as mentioned earlier, plays a key role in determining whether these SRAM cells power up to a “1” or a “0”. These values collectively constitute the bits of the PUF response. Accessing these values is uncomplicated, involving the simple act of powering up the SRAM array and applying the same read procedure commonly used with general-purpose SRAMs.

Nonetheless, SRAM PUFs come with a notable limitation: their reliability. Naturally, a PUF must consistently provide the same response when utilized as an identifier; otherwise, an erroneous response can lead to an error in the encryption scheme, resulting in implementation failure. The primary approach employed to address this challenge involves the use of Error Correction Codes (ECCs) [6], given their capability to correct erroneous bits and return the exact identifier. Nevertheless,

* Corresponding author.

E-mail addresses: santana@imse-cnm.csic.es (A. Santana-Andreo), pablo.sarazacanjlanca@imec.be (P. Saraza-Canflanca), rafael.castro@csic.es (R. Castro-Lopez), eli@imse-cnm.csic.es (E. Roca), pacov@imse-cnm.csic.es (F.V. Fernandez).

<https://doi.org/10.1016/j.aeue.2024.155147>

Received 16 November 2023; Accepted 13 January 2024

Available online 20 January 2024

1434-8411/© 2024 The Author(s). Published by Elsevier GmbH. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

implementing ECCs carries a significant cost in terms of circuit area and power consumption, with these costs scaling directly with the degree of reliability enhancement needed [7]. This aspect directly contradicts the advantages offered by PUFs, which are known for their cost-effectiveness and suitability for implementation in resource-constrained applications such as devices of the Internet of Things. It also contrasts with the benefits of SRAM PUFs, which can leverage pre-existing circuitry. To confront this issue, a common avenue of research is focused on optimizing ECCs to make them as efficient as possible [6,8,9]. However, an alternative approach that can significantly reduce the need for error correction is to take advantage of the inherent reliability of each SRAM cell. This involves *bit selection*, which means identifying, before deployment, the SRAM cells demonstrating consistent powering up to the same values. Different methods and metrics are employed to carry out the selection [10]–[12].

Irrespective of the chosen bit selection method, it is imperative to maintain PUF reliability under varying operating conditions, such as temperature fluctuations, and throughout the device’s lifespan [13,14]. Unless aided by intensive ECC, the raw SRAM PUF implementation falls notably short in addressing the challenge of aging resilience in modern integration technologies. It is then of paramount importance to possess a precise understanding of the PUF’s evolving reliability, enabling accurate ECC determination to avoid both over-compensation and under-compensation and to predict the reliability of the selected cells as they age.

In acquiring such an understanding there are several challenges. First, the specific nature of the aging phenomena should be considered when developing that understanding. Certain aging effects result in permanent degradation of the transistors within SRAM cells, while others cause a degradation that can be reversed, allowing the devices to recover some of their original characteristics. This recovery can be achieved by making suitable adjustments to voltage and temperature, for example, by powering off the devices. Secondly, it is crucial to avoid using methods or metrics that overlook the inherent heterogeneity in cell reliability. Each cell not only possesses a unique level of reliability right after deployment under standard conditions but also exhibits varying sensitivities to aging degradation [15].

In the literature, most of the works that model PUF reliability tend to follow a homogeneous approach. In this approach, every cell’s behavior is averaged and defined by a global metric [7,12,14,16]. However, some works do propose reliability models that consider cell heterogeneity, although they come with certain limitations. For example, the model proposed in [17] and employed in [15], which is generalized for any PUF, assigns an error rate to each individual response bit but is only applicable under specific nominal conditions, as it does not consider the effects of aging or temperature. The work presented in [18] is applicable for generic PUFs as well and includes a term to accommodate temperature, but not aging. An empirical model specifically made for SRAM PUFs is introduced in [4], where the parameters are obtained by fitting experimental data. However, this model does not account for environmental factors or aging. Conversely, experimental studies that investigate the impact of aging on SRAM PUFs tend to rely on extensive statistical characterizations and the averaging of reliability across numerous cell instances [14,19,20], thus failing to account for the full spectrum of stochastic behaviors that individual cells can exhibit. Consequently, it may lead to incorrect assessments regarding the amount of ECC necessary to ensure the reliability of a PUF implementation. The work in [21] does take into consideration PUF reliability after the application of ECC, offering a more accurate perspective. However, this approach primarily aims at leveraging, rather than mitigating, aging to enhance system reliability or for potential malicious attacks.

In this context, it is crucial to gain a comprehensive understanding of the process for selecting bits (i.e., specific SRAM cells) with long-lasting reliability. To achieve this, one must possess an in-depth comprehension of how exactly cell reliability is influenced by permanent aging-induced

degradation and find out if recovery could be used to regain the initial reliability of the cells. Using this comprehension, an adequate model or metric should be attained that permits the PUF designer to select bits with enduring reliability. Expanding on our work in [22], the goal of this paper is then to provide a detailed insight into the impact of aging on an SRAM PUF by closely examining its evolution at the individual cell level. To do this, we employ our custom chip specifically made to facilitate the characterization of aging phenomena through stress, i.e., applying a supply voltage larger than the nominal value of the technology to accelerate the effects of aging. This accelerated aging study provides us with a clear view of how various aging mechanisms interact and influence cell reliability. Ultimately, our objective is to establish a robust metric that can be used for accurate reliable predictions.

Accordingly, the main contributions of this paper are: (a) a wide statistical study of the reliability of fabricated SRAM cells for PUF application purposes and their sensitivity to aging phenomena; (b) the consideration of the commonly ignored non-conductive hot-carrier degradation to account for some of the observed phenomena, and (c) a bit selection method to use only those SRAM cells that guarantee a more stable PUF performance.

The rest of the paper is structured as follows: Section II explains the basic operation of SRAM PUFs and the expected behavior when subjected to aging according to the literature. Moving on to Section III, the metrics used for reliability characterization are introduced. Section IV details the custom chip and experimental setup employed for this purpose and Section V presents the results obtained from this characterization. In light of these results, Section VI describes a reliability prediction method for SRAM cells. Finally, conclusions are drawn in Section VII.

2. Operation of SRAM PUFs

2.1. SRAM cell operation

The 6-T SRAM cell, shown in Fig. 1, is the standard SRAM implementation and the one used in this work. Among the six transistors in the design, four are interconnected to form a pair of cross-coupled inverters responsible for storing the SRAM value, while the remaining two serve as access transistors that control external read and write operations. This cell exhibits two distinct states, determined by the voltage levels at the internal nodes Q and \bar{Q} :

- Q is low voltage and \bar{Q} is high voltage; associated arbitrarily with state “0.”
- Q is high and \bar{Q} is low; associated arbitrarily with state “1.”




Fig. 1. Schematic of a 6-T SRAM cell.

Upon power-up, the SRAM cell typically settles into one of the two states, i.e., its *preferred value*. This selection hinges primarily on slight manufacturing disparities among the transistors constituting the inverters, particularly concerning the threshold voltage (V_{TH}). Such discrepancies cause one of the two inverters to power up faster than the other, thus determining the cell's final state. Specifically, the PMOS transistors (M_2 & M_4) connect each internal node to the supply line. A higher threshold voltage impedes their ability to drive the node to a high value. Conversely, the NMOS transistors (M_1 & M_3) connect each internal node to the ground, so a higher threshold voltage will favor their node powering up to a high value. Depending on the disparity in threshold voltages between the transistors of the two inverters, some cells exhibit a pronounced *bias* (also known as *skew*) toward powering up to the "1" state while others favor the "0" state. Conversely, certain cells exhibit instability due to closely matched inverters, leading to a power-up behavior characterized by unpredictability and randomness. In these *unstable* cells, upon multiple instances of powering up the cell, the final state may intermittently vary between "1" and "0".

Considering that in an SRAM PUF it is crucial for cells to consistently produce the same bit value (i.e., the power-up state) upon each request, cells displaying a pronounced bias toward a specific power-up state are preferred in this context. Conversely, unstable cells are inappropriate for this application and are better suited for generating random numbers [4].

2.2. Aging in SRAM cells

Aging degradation affects various parameters, including the threshold voltage of transistors, consequently disrupting the initial balance of SRAM cells, and jeopardizing the long-term reliability of SRAM PUFs. One of the most prominent aging mechanisms considered in SRAM PUF degradation is Bias Temperature Instability (BTI) [19]. This phenomenon is gate-activated, meaning that a transistor operating in strong inversion mode will be susceptible to BTI degradation. In the context of an SRAM, this implies that whenever a value is stored, both the PMOS transistor with its gate connected to the low-voltage node and the NMOS transistor with its gate linked to the high-voltage node will undergo NBTI and PBTI degradation, respectively (with "N" and "P" meaning "negative" and "positive"). This is illustrated in Fig. 2. The extent of degradation is contingent on the amount of *stress* experienced by the transistors, which depends on the voltage levels at the transistor terminals, temperature, and the duration for which those voltages are sustained.

Let us consider a specific SRAM cell that, for instance, has a strong tendency to power up to "0" and that stores that value over a period. In this case, the PMOS transistor with its gate connected to the Q node (M_2) and the NMOS transistor with its gate connected to the \bar{Q} node (M_3) will experience NBTI and PBTI degradation, respectively, as long as the voltages at their terminals persist (i.e., as long as the "0" value is kept). As time goes on, the threshold voltage of both transistors will increase

due to BTI, which eventually results in a reduction of the power-up skew. This means that, if the cell is ever powered off and subsequently turned back on, the power-up value may be different from the one the cell had before aging, resulting in reduced reliability over time. Applying the same reasoning to a cell biased to powering up to "1" and storing that value, the threshold voltages of the PMOS transistor M_4 and NMOS transistor M_1 will increase over time thus reducing the cell's reliability. In essence, when an SRAM cell stores its preferred power-up value, the bias of the cell is reduced over time due to BTI, thus degrading its reliability.

Nonetheless, when assessing a cell's reliability over time, it is vital to consider what occurs when the stress is alleviated, such as by discontinuing the supply voltage. This is because BTI-induced degradation comprises two components: one that is permanent or quasi-permanent [23], and another that is recoverable and gradually diminishes with a logarithmic time dependence [24] after the stress is removed. This implies that after the stress subsides, the cell tends to regain some of its original reliability, but it may never fully return to its initial state due to the permanent degradation. Consequently, it is critical to ascertain the role of the recovery process and how it can facilitate the return to an initial level of reliability once the SRAM PUF has been operational for an extended period.

BTI is not the sole aging effect affecting CMOS transistors. Another such effect is Hot Carrier Injection (HCI), which arises when a lateral field is created in the transistor due to a high drain-source voltage. This field accelerates charge carriers, causing high-energy carriers (*hot carriers*) to inject into the gate dielectric. This leads to device characteristics degradation, including the threshold voltage. Unlike BTI, HCI degradation is predominantly permanent, lacking the ability to recover over time. HCI is often modeled as gate- and drain-activated and requires a change of logic state to occur in digital cells, with degradation dependent on the frequency of these state changes. In SRAM applications, these changes coincide with write operations, making HCI a concern in scenarios involving frequent write operations [25].

Another aging phenomenon, typically not explored in SRAM PUFs, is non-conductive HCI degradation (NC-HCI) [26,27], often referred to as *Off-State* degradation. This form of degradation necessitates only a high drain-source voltage when the gate-source voltage is zero. Non-conductive HCI may influence SRAM PUF performance thus leading to a broader spectrum of behaviors as the device ages. For instance, if the cell holds a "1" at node Q , not only will the left NMOS M_1 and the right PMOS M_4 experience BTI, but the left PMOS M_2 and the right NMOS M_3 will also undergo NC-HCI, the latter caused by a high drain-source voltage and a low gate-source voltage, as illustrated in Fig. 2.

An important and common trait shared by BTI, HCI, and NC-HCI is their intrinsic stochastic nature in scaled CMOS technology [28]. This means that each transistor's response to aging varies, even when subjected to identical stress conditions [20,21,29]. Consequently, the effect of these aging mechanisms will differ from one cell to another, resulting in some cells experiencing minimal impact while others might even shift




Fig. 2. Aging mechanisms operating in a 6-T SRAM cell.

their preferred state. Additionally, the recovery process can also vary from one cell to the next.

In summary, the impact that aging has on SRAM cell reliability (with permanent degradations of different magnitudes and recoveries of different rates and magnitudes as well, all taking place at the same time) is not easy to predict. Typically, only BTI has been considered [20]. In this somewhat intricate but comparatively simpler scenario, one could expect that the decline in reliability is either permanent or partially (and possibly even fully) recoverable. However, as it will be shown here, that is not always the case, especially when the less-studied effect of NC-HCI, is considered. Therefore, the complexity of selecting reliable cells right after manufacturing is quite high. However, we will demonstrate here that picking proper metrics and having a detailed picture of BTI and NC-HCI impact can be of use to SRAM PUF designers.

3. Reliability characterization metrics

The main metric used to evaluate the reliability of each cell is the Bit Error Rate (BER). It is determined through multiple power-ups under specific operating conditions (e.g., at a particular temperature and at a specific point in the cell's lifetime) and computing the ratio between the number of erroneous power-ups (e) and the total number of power-ups (n):

$$BER (\%) = \frac{e}{n} \cdot 100\%$$

A bit is deemed erroneous when it differs from the so-called “golden response”, which is derived from the majority response gathered from multiple power-ups performed under nominal conditions when the cell is fresh. This metric directly governs the reliability of the SRAM PUF: lower BER values in SRAM cells indicate a more reliable SRAM PUF and thus a less stringent need for ECC.

Another useful metric is the Data Retention Voltage (DRV). This metric aids in quantifying the skew of each cell to power up to one state or another [16]. Each SRAM cell has a pair of individual DRV values (DRV_1 and DRV_0), one corresponding to the state “1” and another to the state “0”. In essence, the DRV is the supply voltage at which the cell goes from keeping the stored value to switching to the alternate one. Each individual DRV_i is measured by storing value i and then reducing the supply voltage to a certain, non-zero value. Then, it is then raised back to the nominal level while monitoring if the stored state has changed. If no change is observed, the supply voltage is further lowered in incremental steps until either a change is detected, or the entire supply voltage range has been explored.

To fully understand this metric and its potential in evaluating the reliability of an SRAM cell, let us consider several cases that can be identified during a DRV characterization. If a cell has a strong bias towards the initially stored state (e.g., “1”), it will always power up to that value no matter how much the supply voltage is reduced; in this case, its individual DRV value is 0 (e.g., $DRV_1 = 0$). In contrast, when storing the non-preferred value, the individual DRV is typically large (e.g., DRV_0 is large). Most cells do have a strong bias towards either “1” or “0”. As a result, the absolute difference between the DRV_1 and DRV_0 is large. Conversely, cells lacking a pronounced bias toward either power-up state exhibit similar values for DRV_1 and DRV_0 , with both typically being low. This symmetry in values reflects the cell's balanced characteristics and unstable power-up response, which, in turn, indicates a lack of reliability.

To summarize, the BER metric quantifies the SRAM cell's reliability under certain operating conditions while the DRV metric provides a qualitative measure of the bias or skew of the cell towards its preferred value. However, while evaluating the BER metric requires (typically) a large number of power-up cycles, quantifying a cell's DRVs is much faster [11,16]. Furthermore, cells with lower BER values, indicating greater reliability, do not necessarily ensure consistent power-up reliability when aging is considered. Therefore, relying solely on the BER

metric to select the most reliable SRAM cells when they are new may result in unforeseen PUF reliability issues over time.

The DRV metric, on the other hand, has been demonstrated to be useful in pinpointing the most reliable cells [11] under varying operation conditions and aging degradation. This work goes one step forward and uses the underlying correlation between BER and DRV as a predicting tool to identify the most reliable cells not only when a permanent degradation caused by aging is considered but also when recovery is part of the picture. The predictive capability offered by this method may be valuable for PUF designers employing SRAM cells as the fundamental units of entropy. This not only enables them to generate more reliable responses but also provides insights into whether recovery can be leveraged to continue obtaining reliable responses after aging.

4. Array description and experimental setup

To find out what is the correlation between DRV, BER and aging and examine the various scenarios that BTI and NC-HCI may produce, an extensive experimental characterization has been carried out. The integrated circuit used for this purpose fabricated in a CMOS 1.2-V, 65-nm process, contains an array of 832 6-T SRAM cells, each one embedded in a unit cell [30,31]. Regarding the design of each SRAM cell, both the NMOS access (M_5 & M_6) and PMOS pull-up transistors (M_2 & M_4) sizes are $W = 80$ nm and $L = 60$ nm, while the NMOS pull-down transistors (M_1 & M_3) sizes are $W = 160$ nm and $L = 60$ nm, following the standard sizing for SRAM cells. The array has been designed so that, through a Force & Sense architecture, an accurate voltage can be applied at each terminal, a critical aspect required for precise accelerated aging tests due to the high dependence of aging effects on biasing conditions [32]. Its other key features are the independent control of each cell and the ability to perform parallel stress experiments, where a number of cells experience the same stress conditions without having to employ large characterization times. Fig. 3 shows the experimental setup used, which includes a Printed Circuit Board (PCB) specifically designed for this chip, along with a power supply for PCB biasing, a Data Acquisition System (DAQ) from National Instruments to generate the digital control signals, and a Keysight B1500 semiconductor parameter analyzer for precise voltage application and measurement of the SRAM cells. The measurements are automated using a Matlab script that controls the instruments through GPIB and the DAQ through the USB port.

The stress scenario utilized for aging the SRAM PUFs involves storing, in each cell, its preferred power-up state, thus subjecting the transistors to the aging mechanisms described in Section II [29]. Storing the preferred power-up value aims to mitigate the skew toward this




Fig. 3. Experimental set-up used to evaluate the SRAM PUF.

preferred power-up value in the cells. Consequently, a generalized degradation of the reliability of the cells is expected in terms of their power-up value. The experimental procedure goes as follows: First, a comprehensive characterization is conducted, encompassing the measurement of BER and DRV, serving as a reliability benchmark before the aging process begins. Then, an accelerated aging phase is initiated, during which a stress voltage of 2.5 V is applied for 10,000 s at room temperature to the supply line of each cell while they retain their preferred state. This accelerated aging has been devised to emulate the continuous storage of a key in the PUF using the accelerated factor model in [33–35].

The BER and the DRV are reevaluated right after the accelerated aging, and at three subsequent moments (1.5 days, 5 days, and 14 days) to evaluate if and how the recoverable component of the degradation occurs. In between these moments, the SRAM cells are powered off, so no stress whatsoever is applied. The parallelization scheme explained in [30] is employed to measure all cells in the array in a reasonable time while making sure that they endure the same stress conditions. To measure the BER, 400 power-ups in each cell are performed in the first characterization and 200 power-ups afterward. As for the DRV measurement, a step size of 5 mV has been used.

5. Experimental results

5.1. Correlation between BER and DRV

In this section, the results obtained for BER and DRV on the SRAM array described above are presented. A broad picture of these results is given in Table I. It includes the mean BER value of all the array cells before the accelerated aging (AA) is applied and at different times after. Unstable cells (with a BER different than 0 %, i.e., cells with at least one power-up error) are listed in the third column.

Regarding the DRV, the initial distribution of the individual differences, measured before the accelerated aging, is shown in Fig. 4. SRAM cells with positive values of the difference have a more pronounced bias towards a “0” power-up while cells with a negative difference tend to power up more consistently to “1”. Also, the lower the difference, the weaker the bias and therefore a less stable power-up value. The maximum values of DRV_1 and DRV_0 are 260 mV and 240 mV, respectively, while the minimum value is 0 V in both cases. The values of $\max(DRV_1, DRV_0)$ span from 85 mV to 260 mV with an average of 161.09 mV and a standard deviation of 28.85 mV.

The average BER is low when the SRAM cells are fresh. Notably, what adds weight to the DRV as a predictive metric for reliability is the correlation between BER and DRV, as shown in Fig. 5. Clearly, when the difference between DRV_1 and DRV_0 is small (indicating either that both metrics have similar values or one of them is small while the other is zero), the cell has a larger probability of being unstable (i.e., $BER > 0$

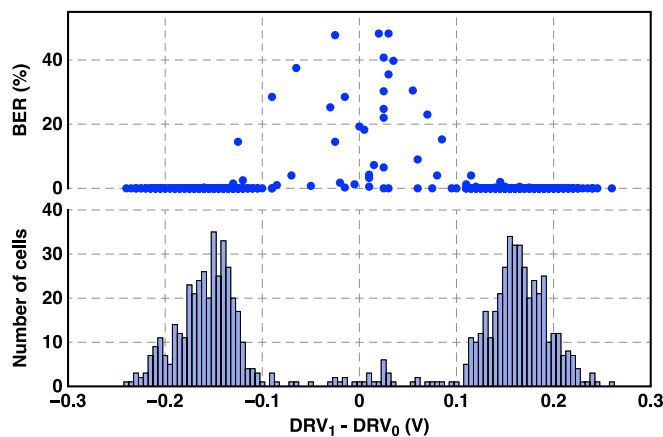


Fig. 5. Correlation between BER and the difference $DRV_1 - DRV_0$ before accelerated aging.

%)

From Table I, it is also clear how the BER increases significantly due to stress, going up by an order of magnitude, indicating that the array, as a PUF, has become much less reliable. This value decreases after 1.5 days, down to 5.50 %. The two remaining measurements, after 5 and 14 days, show similar average BER values, evidencing that most of the recovery occurred in the first 1.5 days due to the logarithmic time nature of the BTI recovery [24]. In any case, it is clear that the fresh state of the array has not been fully recovered, as an average BER of more than 5 % is significantly higher than the initial BER of 0.78 %.

Nonetheless, solely examining the average BER might be deceptive, as it can obscure the intricate nuances of how aging and recovery affect the SRAM PUF. Crucial insights may lie within these intricacies and detailed observations. Indeed, the observation that the average BER seems to worsen 14 days after the accelerated aging with respect to the previous measurement (performed 5 days after the accelerated aging) provides the first clue that looking at the average BER is not sufficient to understand the complete picture. Although the average BER had a tenfold increase after the accelerated aging, the number of actual cells that produced any errors increased only by a factor of 3X. This is because some cells accumulate most of the errors in the array. For instance, several cells switch from always powering up to one value ($BER = 0$ %) to always powering up to the other value, which results in a BER of 100 %, heavily influencing the mean value over the whole array. In addition, some of the initially unstable cells, have become stable to either state (i.e., $BER = 0$ % or $BER = 100$ %) 14 days after. A more detailed picture of the individual evolution of each cell’s BER is shown in Fig. 6. This illustration also underscores the absence of any discernible correlation between the position of a cell and its BER.

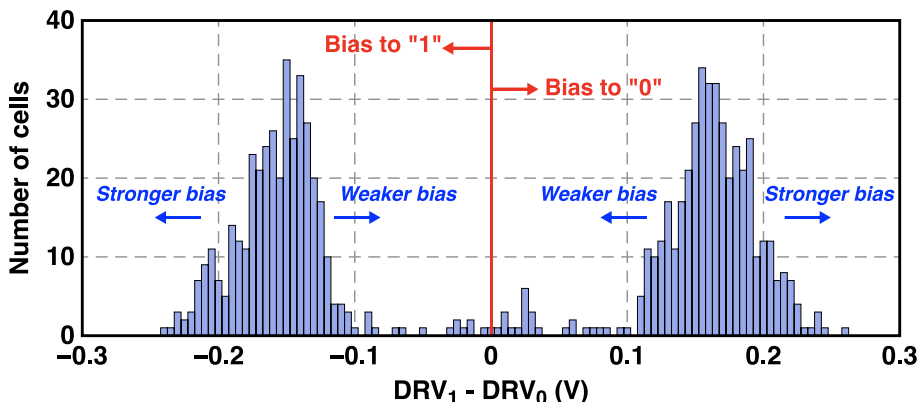


Fig. 4. Difference of individual DRVs before applying stress to create accelerated aging.




Fig. 6. BER values for each cell before AA (left), right after AA (center) and 14 days after AA (right).

As for the cells that would make a reliable SRAM PUF, Fig. 7 illustrates the evolution, in terms of BER, of the 788 cells that are perfectly stable initially (and thus would be ideal for a perfectly reliable PUF if only the BER were considered). From these, 697 cells are undisturbed by the accelerated aging. Right after accelerated aging, 91 cells (~12 %) become unstable (this means that, considering that the accelerated aging emulates the PUF performance at a given time, 91 stable cells would have become unreliable by then, ending up with 697 of the originally stable ones). Note that immediately following the accelerated aging, Table 1 indicates 132 unstable cells, while Fig. 7 indicates 697 stable cells. The total would sum to 829 cells, but there are 832 cells in total. The difference comes from 3 cells that were unstable before the accelerated aging and, as a result, are not included in the representation in Fig. 7, but the accelerated aging made them stable.

Using recovery as a method to regain some reliability, after 14 days, 732 cells of the initial 788 ones are stable, which means that 35 cells have returned to having a BER of 0 % and that ~ 93 % of the initially stable cells can still be used to generate reliable PUF responses. However, Fig. 7 also illustrates an interesting fact: while there are some cells that are permanently damaged (43) and cells that are untouched by aging (695), other cells (50) experience a not-so-straight evolution. For

Table 1

Average BER and number of unstable cells at different points of the characterization.

Measurement	BER (%)	Unstable cells
Before AA	0.78	44
Right after AA	7.93	132
1.5 days after AA	5.50	95
5 days after AA	5.29	98
14 days after AA	5.33	95

instance, there are cells that, after totally recovering at 1.5 days, become unstable again at 5 days (when no stress is applied) to then return to stable at 14 days, while other cells remain unstable to suddenly recover at 14 days. While not a significant portion of the cells, this plethora of behaviors underscores the intricacies of aging and its impact on PUF reliability since the combination of permanent and recoverable aging may yield complex and not easy-to-anticipate behaviors. The DRV metric exhibits complex behavior as well, as depicted in Fig. 8. This figure shows the changes in the DRV difference ($DRV_1 - DRV_0$) for initially stable cells in two scenarios: first, between before and




Fig. 7. Evolution of initially stable cells with accelerated aging and recovery. The numbers indicate the count of cells at each moment and with BER zero (above) or non-zero (below).

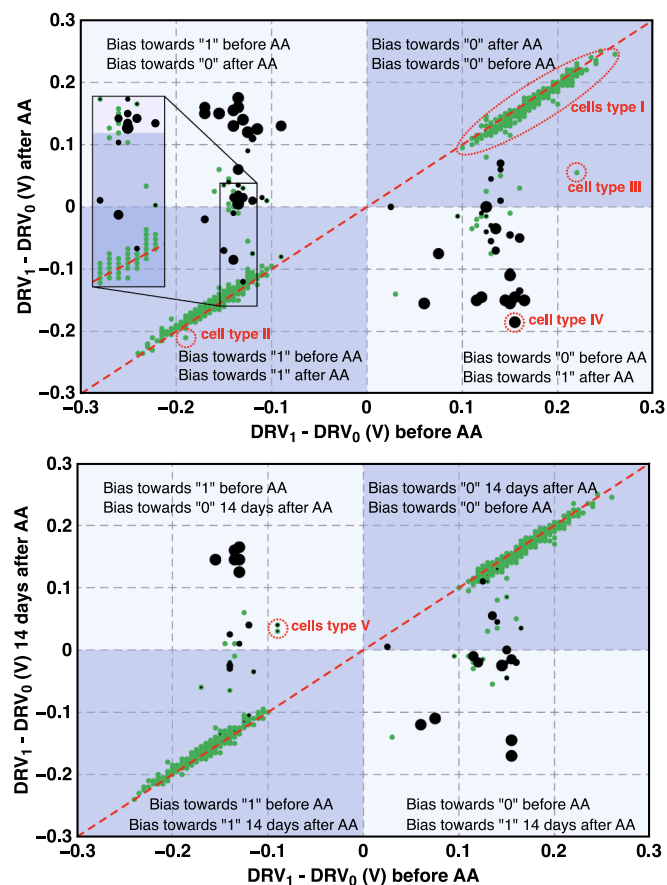


Fig. 8. Change in the difference $DRV_1 - DRV_0$ before and after the AA (top) and before and 14 days after the AA (bottom). The black dots represent unstable cells ($BER > 0\%$) and the size of each black dot is proportional to the BER value.

immediately after accelerated aging (top plot); second, between before and 14 days following accelerated aging (bottom plot). In both plots, the bias in each of the four quadrants, according to the DRV difference, is indicated. As previously noted in Fig. 4, a larger absolute magnitude of the DRV difference corresponds to a stronger bias. Cells with worsened BER after accelerated aging (top plot) or those with incomplete BER recovery after 14 days (bottom plot) are marked with black dots; the size of each black dot is proportional to the BER value.

A significant portion of cells maintains the DRV difference in the same quadrant with a very similar DRV difference value after accelerated aging (e.g., cell type I). Some of these cells even improve the magnitude of difference slightly (e.g., cell type II). A few cells remain in the same quadrant, but the magnitude of the DRV difference decreases, indicating a weakened bias (e.g., cell type III). However, there are cells whose DRV difference changes sign and magnitude (e.g., cells type IV); most of these cells are cells that became unstable ($BER > 0\%$) after the accelerated aging, as indicated by the black dots (top plot), or cells with incomplete BER recovery (bottom plot). Moreover, the further the vertical move to the other quadrant, the larger the BER value (bigger black dots). Comparing the top and bottom plots it is also easy to see that, in general, the recovery diminishes instability, as there are fewer black dots in the bottom plot (note, for instance, that cell type III has fully recovered). In a very low number of cases though, the recovery seems to either be incomplete or worsen the DRV difference (note cells type V indicated in the bottom plot of Fig. 8).

These plots expose three things. First, they underscore the inherent stochasticity of aging phenomena in scaled transistors. The combination of aging stochasticity and process variations, including inherent

stochasticity at time zero, gives rise to unique degradation patterns for each cell. This occurs even in instances where the initial DRV difference values are identical, as evidenced by the zoomed-in section in the top plot of Fig. 8. Here, cells with similar DRV differences before accelerated aging display distinct DRV differences afterward. Second, the different ways the magnitude of the DRV difference changes—whether it improves (cell type II), worsens (type III and IV), or stays the same (type I)—show that basic aging models cannot fully grasp the complex aging process in SRAM cells. Third, recovery helps many cells to regain their original stability but it is not absolutely widespread and there are some instances where the cell's DRV difference worsens or never recovers in full.

5.2. Ultimate causes of individual cell unreliability

Generally, in prior studies, only a permanent NBTI effect was considered [20]. However, such an approach would only predict the existence of cells type I and II and never the evolutions of cells type III, IV or V. This is illustrated with the aging simulation in Fig. 9, where only NBTI was included and the DRV differences, $DRV_1 - DRV_0$, before and after the aging, are plotted. A permanent degradation was considered by adding a voltage source at the gate of the relevant transistor. This source represents NBTI-induced ΔV_{TH} degradation [36]. Although this is a simplified approximation of the actual SRAM cell aging, it offers valuable insight into what is typically expected. In the simulation, the DRV before and after aging of 832 6-T SRAM cells were evaluated (all cells are sized as mentioned in Section IV, process variations were also considered and the preferred value is stored). Please note that the margin between DRV_1 and DRV_0 either worsens or remains the same after the aging of the cells, leading to instability and even a reversal of the preferred value when a sign change occurs. This is evident as there are no data points below (in the left half of the plot) or above (in the right half) the diagonal. This observation contrasts with the experimental results depicted in Fig. 8. The lack of recovery consideration in the simulation, a common characteristic in many prior studies—something that could otherwise explain the data presented in the bottom plot of Fig. 8—suggests that the analysis of SRAM cell aging should involve a more comprehensive set of mechanisms to account for the various scenarios observed experimentally.

This variety emerges if cells are examined on an individual basis. Fig. 10 shows some representative examples from the experimental measurements. Most of the cells behave as Cell A, with a slight decrease in the DRV difference $DRV_1 - DRV_0$ but retaining their preferred value and not having any reliability issues in the BER characterizations over time. Cells like this one should be the preferred cells to form the PUF response as they will be the ones with the largest skew towards its




Fig. 9. Simulated change in the cell's DRV difference due to permanent NBTI only.




Fig. 10. Representative examples of the SRAM cell aging. In the X-axis, the numbers represent the moments where the measurements were taken: “0”, before AA, “1” right after AA, “2” 1.5 days after AA, “3” 5 days after AA, and “4” 14 days after AA.

preferred value. In fact, the 223 cells with the highest values of the difference retain a BER of 0 % at any time. This fact will be used, as explained in the next Section, to set a prediction metric for the SRAM cell reliability. This is confirmed by looking into the behavior of Cell B. This cell is stable before the accelerated aging, but its DRV difference is not among the highest ones, indicating that this cell might be unstable, as it is confirmed after stress where BER increases above 0 % and the DRV difference changes sign. Cells C and D fit the expected behavior considering the recoverable component of the BTI degradation, as their preferred value flips due to aging degradation (there is a sign change in the DRV difference as well) to then recover fully, in terms of BER, and end up with a similar DRV difference to the one they started with. As shown in Fig. 7, there are 37 cells that follow a similar pattern in terms of BER, i.e., having a skew towards a certain value, becoming unstable due

to the stress applied but going back to their original skew with a BER of 0 % (these cells are indicated with a black outline). These cells may not exhibit a reliable PUF response when immediate aging is considered, as their response is unstable immediately following the accelerated aging process. However, it is worth noting that a 14-day recovery period has proven effective in restoring their full reliability.

On the other hand, Cell E goes from always powering up to “1” to always powering up to “0”; note the large change in magnitude if the DRV difference and the sign change. It appears like the degradation caused during the stress has left a strong, permanent mark on the cell that does not recover by leaving the cell powered off for 14 days. Upon scrutinizing Cells C, D, and E, it might appear that permanent degradation and recovery can occur irrespective of the specific value of the cell’s DRV difference before the accelerated aging. However, as detailed

in the next section, there is a threshold above which no permanent reliability degradation persists. This insight can be used to predict reliability over time by evaluating the cells' individual DRVs before deploying the PUF.

Finally, there are some cell behaviors that do not fit the all-recovered or all-degraded pattern. For instance, Cell F remains stable until 5 days after the accelerated aging. This and similar observations can only be explained if a combination of BTI and NC-HCI in PMOS and NMOS transistors are all simultaneously considered. To better understand the different recovery behaviors observed in the cells, an aging experiment was performed on individual transistors of the same size fabricated in the same technology [37].

As previously noted, in the existing literature, the aging-induced degradation of SRAM cells is predominantly linked to BTI in PMOS transistors. This encompasses both a recoverable and a permanent component. To investigate this, 50 PMOS and 50 NMOS transistors underwent an accelerated aging process akin to that experienced by the transistors in Fig. 2, with $|V_{DS}| = 0$ V, $|V_{GS}| = 2.5$ V, for 10 ks at room temperature. Before and after the accelerated aging, the current of transistors is measured with $|V_{DS}| = 0.1$ V, $|V_{GS}| = 0.6$ V. It is noteworthy that no significant degradation was observed in the NMOS transistors. In contrast, the degradation observed in the PMOS transistors was more pronounced. The Cumulative Distribution Functions (CDFs) in Fig. 11 show a notable variance in current degradation (primarily produced by the change in threshold voltage caused by the accelerated aging), clearly illustrating the stochasticity of the phenomena. Additionally, the results demonstrate that recovery from this degradation can be swift, predominantly occurring within the initial 100 s after the accelerated aging period. At 1.5 days following the stress removal, the majority of transistors have recovered, although a few still exhibit a permanent degradation, accounting for approximately 10 % of the initial current after 14 days. This enduring degradation may contribute to the permanent shift in the power-up state observed in certain SRAM cells, but other effects are also at play.

Theoretically, as explained in Section II, if an SRAM cell maintains the power-up value during a certain time, the PMOS transistor in one of the inverters undergoes NBTI degradation (e.g., transistor M_2 in Fig. 1), while the NMOS transistor (M_1) experiences NC-HCI. At the same time, the NMOS transistor of the opposite inverter (M_3) undergoes PBTI and the PMOS (M_4) suffers from NC-HCI. To ascertain the extent to which NC-HCI can degrade the transistors in this technology, the same accelerated aging conditions resulting from the experimental setup described in the previous section ($|V_{GS}| = 0$ V and $|V_{DS}| = 2.5$ V for 10 ks at room temperature) were applied to 50 PMOS and 50 NMOS transistors. As in the previous experiment, the current of transistors is measured with $|V_{DS}| = 0.1$ V, $|V_{GS}| = 0.6$ V before and after the accelerated aging. The results of this experiment, shown in Fig. 12, confirm previous reports [26,27]. PMOS transistors exhibit an enhancement in their baseline

current, implying a reduction in their threshold voltages, which is the opposite effect of NBTI stress. NMOS transistors, however, display a mixed response, with some showing a current degradation and others showing an improvement. In both cases, however, while some recovery exists after 1.5, 5 or 14 days, a significant part of the change is permanent. Also, the current change right after the accelerated aging in PMOS transistors is similar to the one caused by NBTI, albeit with an opposite sign, as shown in Fig. 13. The average current degradation in NMOS due to NC-HCI is approximately 15 % of that caused by the same effect on PMOS transistors. It is also evident that the recovery rate varies between the two phenomena: it is rapid for BTI and slower for NC-HCI.

5.3. Concluding remarks

Given the results in previous subsections, it is apparent that the precise way in which aging-induced degradation and recovery impact SRAM cells can be quite complex and vary significantly from one cell to the next. For instance, let us consider an SRAM cell with a preferred power-up value "1" (i.e., Q node is at 1.2 V) that stores that same value for a given time. Aging can lead to various scenarios when this bit is subsequently requested by the PUF. The threshold voltage of transistor M_4 (PMOS) decreases due to NBTI, weakening its pull-up ability; meanwhile for transistor M_3 (NMOS), an increase or, more likely, a reduction of its threshold voltage may happen because of NC-HCI, thus affecting its pull-down ability. If the pull-down ability becomes stronger, the bias towards powering up to "1" will be diminished, and if it becomes weaker, it depends on the relative magnitude of aging-induced changes in the PMOS and NMOS transistors to determine whether the bias towards powering up to "1" is altered (considering the average magnitude of NC-HCI degradation in NMOS transistors is lower than NBTI, as explained above, it is more likely for the NBTI-aged PMOS to dominate). At the same time, PMOS transistor M_2 is affected by NC-HCI causing its threshold voltage to decrease, further reducing the bias towards powering up to "1" (since node \bar{Q} will power up to "1" instead). This effect compounds the impact of NBTI on transistor M_4 , so aging in both PMOS transistors will reduce the reliability of the cell, while the aging-induced impact in transistor M_3 may have the opposite effect. This explains, for instance, the increase in DRV values in some cells.

Note though that when recovery is considered, the balances can be quite different: in many cases, the larger NBTI degradation disappears while the NC-HCI degradation remains. The rate at which the recovery happens and the amount of permanent damage that remains dictates either an increase or reduction of the bias toward powering up again to the value obtained before the aging. Moreover, in some less common instances, the intensity of the bias may change even during the recovery.

In general, these findings suggest that while they validate the most likely aging trends, the endeavor of formulating an analytical model capable of predicting long-term reliability appears to be a formidable and intricate task. An empirical metric, as detailed in the following section, or even a regression-based machine learning model utilizing experimental data, may offer a more suitable approach for this challenge. It is also important to add a note of caution regarding what is known as *directed aging*, to improve reliability or perform security attacks on the PUFs by forcing false negatives or creating PUF clones [19,20,29]. As evidenced here, the response of all cells to accelerated aging is not as straightforward as previously assumed. Concerning reliability enhancement, some cells may regain their original behaviors following reinforcement, while others may not, depending on the effects of NC-HCI.

The results above have been obtained for a 65-nm CMOS process and, strictly speaking, are only valid for this technology. The relative importance of the wide ensemble of variability phenomena changes with evolving technologies. For instance, the impact of NC-HCI in HKMG technologies is qualitatively and quantitatively different [38], or the PBTI impact becomes comparable to that of NBTI [39]. Therefore,




Fig. 11. Relative current change observed in PMOS transistors due to NBTI.

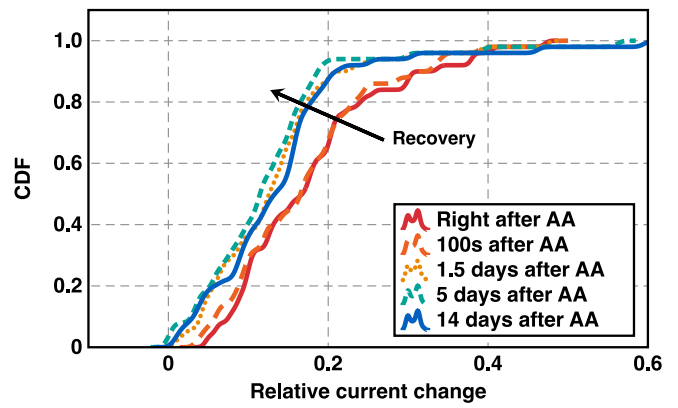



Fig. 12. Relative current change observed in NMOS transistors (left) and PMOS transistors (right) due to NC-HCI.




Fig. 13. Relative current change observed in PMOS transistors caused by NC-HCI and NBTI right after accelerated aging.

although the fact that the PUF reliability is a consequence of the combined impact of different variability phenomena in different devices remains valid, the specific impact for each particular technology should be experimentally studied following a similar methodology to that described in this paper.

6. Reliability prediction using the DRV metric

Fig. 5 shows that a correlation exists between the magnitude DRV_1 - DRV_0 and the cell BER. Building upon this insight, a reliability prediction method using the DRV values has been explored. The aim of this approach is to assist PUF designers using SRAMs in choosing bits with improved and enduring reliability, which can be achieved through the evaluation of the DRV of their cells before deployment.

That is, the method selects the cells that (1) do not show signs of instability reflected in similar and non-zero values of their individual DRVs (DRV_1 and DRV_0), and (2) have the highest DRV values. To do so, this method proceeds as follows: First, with the experimental data for the DRV measurements before any aging has happened, the cells are

separated into two groups: group A, where all cells meet the conditions $\min(DRV_1, DRV_0) = 0$, and group B, where $\min(DRV_1, DRV_0) \neq 0$; Subsequently, cells in group A are sorted based on the value of $\max(DRV_1, DRV_0)$; Then, a specific DRV threshold is selected: only cells in group A with DRV value above said threshold are selected for the PUF to generate an enduring reliable response, while all cells in group A with DRV value below the threshold and all cells in group B are discarded. This method is illustrated in Fig. 14.

Note that the method does not use the BER metric but only the DRV measurements before the accelerated aging. But to find out how accurate is this method in predicting cells that are stable at a particular moment (before, right after the accelerated aging or during recovery), the BER measurements at that moment are needed. In this way, the accuracy is determined by the ratio of the number of cells in group A with $\max(DRV_1, DRV_0)$ above the chosen threshold, to the number of cells in group A, also with $\max(DRV_1, DRV_0)$ above the selected




Fig. 15. Prediction accuracy of the proposed DRV-based method.




Fig. 14. Illustration of the proposed DRV-based method.

threshold, that have a non-zero BER at that specific moment.

Fig. 15 illustrates the accuracy of the proposed method. For each DRV threshold value, the plot showcases the accuracy in predicting stable cells at three moments: before, right after, and 14 days after the accelerated aging. It is evident that, for any of the moments, as the accuracy in selecting reliable cells only improves (i.e., as the DRV threshold value increases along the X-axis), there is a decrease in the number of selected cells. The trade-off is clear here: the more sound the selection of enduring reliable cells is, the lower the number of available cells. Therefore, if a large number of cells is required, there is the risk of including unreliable cells in the selection made by the method. This is depicted in the top subplot of the figure, showing the number of cells in group A with $\max(DRV_1, DRV_0)$ above the threshold. Note also that the method is excellent at selecting only highly reliable cells before aging (with an accuracy above 99.7 %), worsens right after the accelerated aging but gets again better with recovery. In any case, the accuracy is above 91 %. For the sake of comparison, the plot also includes the accuracy when no distinction is made according to $\min(DRV_1, DRV_0)$, that is, no initial group separation is made. Noticeably, the accuracy worsens in all cases with a very minor improvement in the number of cells for every DRV threshold value. Finally, the accuracy prediction for cells that are stable at all moments (i.e., BER = 0 % always) is also shown in the plot. Finally, note from Fig. 15 that there is a DRV threshold value above which all cells have perfect reliability across various measurement moments. This value, 180 mV, allows the designer to use 223 cells out of the 832 available to form a perfectly reliable PUF even when factoring in the impact of aging.

In this sense, and using this method, the PUF designer can select a value of the DRV threshold that presents a convenient trade-off between the reliability of the PUF over time and the number of available bits. For instance, the PUF designer could also use this method to find out how intensive the ECC must be for a 128-bit key generation based on a fuzzy commitment scheme [16]. If a repetition code is used as ECC and a typical value of 10^{-4} % for the Key Error Rate [7] is set (which implies an average BER after applying the ECC of 7.8×10^{-7} %), the resulting required repetition code, as a function of the DRV threshold, is shown in Fig. 16. That is, the lower the selected DRV threshold, the more cells are included (as shown in the top plot of Fig. 15) but the more intensive the repetition code to implement ECC, especially when the ECC must be designed to include the aging-induced degradation of the SRAM cells.

7. Conclusions

In this work, a detailed experimental study of the nature of aging in SRAM PUFs has been presented. The measurements have been performed using a custom SRAM-cell test array that allows the accurate characterization of aging in the cells. It has been shown that, due to the stochastic nature of reliability phenomena in scaled CMOS technologies, different cells will suffer different degradation even under the same stress conditions. Furthermore, it has been demonstrated that the impact of NC-HCI on the SRAM cells needs to be considered together with that of BTI to achieve a comprehensive understanding of the nature of aging in SRAM PUFs. Then, leveraging this knowledge, a new method to select the cells that will have a more reliable power-up value has been developed. This method can improve the performance of SRAM PUFs and thus reduce considerably the overhead associated with ECCs.

CRediT authorship contribution statement

A. Santana-Andreo: Data curation, Formal analysis, Investigation, Software, Validation, Writing – original draft, Writing – review & editing. **P. Saraza-Canflanca:** Conceptualization, Data curation, Formal analysis, Investigation, Software, Validation, Writing – original draft, Writing – review & editing. **R. Castro-Lopez:** Conceptualization, Data curation, Formal analysis, Funding acquisition, Investigation,

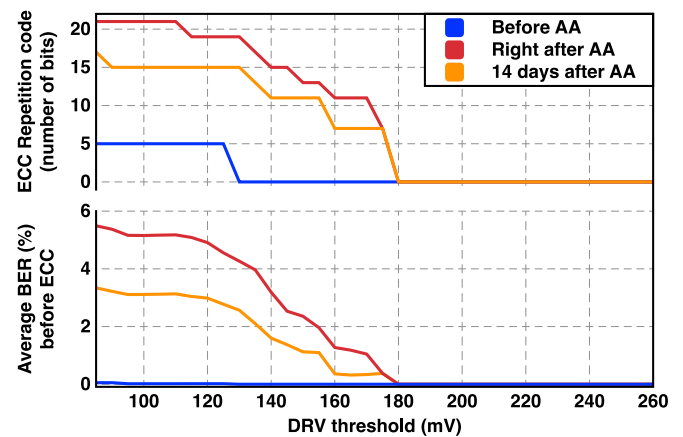


Fig. 16. ECC repetition code before and after the accelerated aging.

Methodology, Project administration, Resources, Software, Supervision, Validation, Visualization, Writing – original draft, Writing – review & editing. **E. Roca:** Conceptualization, Data curation, Formal analysis, Investigation, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **F.V. Fernandez:** Conceptualization, Data curation, Formal analysis, Funding acquisition, Investigation, Methodology, Project administration, Resources, Validation, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgments

This work was supported by grant PID2019-103869RB-C31 funded by MCIN/AEI/10.13039/501100011033 and by grant PID2022-136949OB-C21 funded by MCIN/AEI/10.13039/501100011033 and by “ERDF A way of making Europe”. Andrés Santana Andreo was supported by grant PRE-2020-093167 funded by MCIN/AEI/10.13039/501100011033 and by “ESF Investing in your future”.

References

- [1] McGrath T, Bagci IE, Wang ZM, Roedig U, Young RJ. A PUF taxonomy. *Appl Phys Rev* 2019;6(1):1–25.
- [2] M. C. Martínez-Rodríguez, L. F. Rojas-Muñoz, E. Camacho-Ruiz, S. Sánchez-Solano, and P. Brox, “Efficient RO-PUF for generation of identifiers and keys in resource-constrained embedded systems,” in *Cryptography*, vol. 6, no. 4, vol. 51, pp. 1–20, 2022.
- [3] Machida T, Yamamoto D, Iwamoto M, Sakiyama K. A new arbiter PUF for enhancing unpredictability on FPGA. *Sci World J* 2015;2015:1–13.
- [4] Holcomb DE, Bursleson WP, Fu K. Power-up SRAM state as an identifying fingerprint and source of true random numbers. *IEEE Trans Comput* 2009;58(9):1198–210.
- [5] Intrinsic ID. “The reliability of SRAM PUF,” 2017. [Online]. Available: <https://www.intrinsic-id.com/resources/white-papers/landing-page-white-paper-reliability-sram-puf/>. [Accessed: Nov. 10, 2023].
- [6] Hiller M, Kürzinger L, Sigl G. Review of error correction for PUFs and evaluation on state-of-the-art FPGAs. *J Cryptogr Eng* 2020;10(3):229–47.
- [7] Alioto M. Trends in hardware security: from basics to ASICs. *IEEE Solid-State Circuits Mag* 2019;11(3):56–74.
- [8] Bösch C, Guajardo J, Sadeghi AR, Shokrollahi J, Tuyls P. Efficient helper data key extractor on FPGAs. *Lect Notes Comput Sci* 2008;vol. 5154 LNCS:181–97.
- [9] M. Hiller and G. Sigi, “Increasing the efficiency of syndrome coding for PUFs with helper data compression,” in *Proc. of DATE*, pp. 4–9, 2014.

- [10] Baturone I, Prada-Delgado MA, Eiroa S. Improved generation of identifiers, secret keys, and random numbers from SRAMs. *IEEE Trans Inf Forensics Secur* 2015;10(12):2653–68.
- [11] Saraza-Canflanca P, et al. Improving the reliability of SRAM-based PUFs under varying operation conditions and aging degradation. *Microelectron Reliab* 2021; 118:1–8.
- [12] Wang W, Singh AD, Guin U. A systematic bit selection method for robust SRAM PUFs. *J Electron Test* 2022;38(3):235–46.
- [13] L. Kusters, A. Rikos, and F. M. J. Willems, "Modeling temperature behavior in the helper data for secret-key binding with SRAM PUFs," in *Proc. of CNS*, pp. 1–6, 2020.
- [14] R. Wang, G. Selimis, R. Maes, and S. Goossens, "Long-term continuous assessment of SRAM PUF and source of random numbers," in *Proc. of DATE*, pp. 7–12, 2020.
- [15] Delvaux J, Gu D, Schellekens D, Verbauwheide I. Helper data algorithms for puf-based key generation: overview and analysis. *IEEE Trans Comput Aided Des Integr Circuits Syst* 2015;34(6):889–902.
- [16] Santana-Andreo A, et al. A DRV-based bit selection method for SRAM PUF key generation and its impact on ECCs. *Integration* 2022;85:1–9.
- [17] Skorić B, Tuyls P, Ophey W. Robust key extraction from physical uncloneable functions. *Lect Notes Comput Sci* 2005;3531:407–22.
- [18] R. Maes, "An accurate probabilistic reliability model for silicon PUFs," in *Lecture Notes in Computer Science*, vol. 8086 LNCS, pp. 73-89, 2013.
- [19] R. Maes and V. Van Der Leest, "Countering the effects of silicon aging on SRAM PUFs," in *Proc. of HOST*, pp. 148–153, 2014.
- [20] M. Bhargava, C. Cakir, and K. Mai, "Reliability enhancement of bi-stable PUFs in 65nm bulk CMOS," in *Proc. of HOST*, pp. 25–30, 2012.
- [21] A. Roelke and M. R. Stan, "Attacking an SRAM-based PUF through wearout," in *Proc. of ISVLSI*, pp. 206–211, 2016.
- [22] A. Santana-Andreo, P. Saraza-Canflanca, H. Carrasco-Lopez, R. Castro-Lopez, E. Roca, and F. V. Fernandez, "A detailed, cell-by-cell look into the effects of aging on an SRAM PUF using a specialized test array," in *Proc. of SMACD*, pp. 1-4, 2023.
- [23] Stathis JH, Mahapatra S, Grasser T. Controversial issues in negative bias temperature instability. *Microelectron Reliab* 2018;81:244–51.
- [24] B. Kaczer et al., "Ubiquitous relaxation in BTI stressing—New evaluation and insights," in *Proc. of IRPS*, pp. 20-27, 2008.
- [25] V. M. van Santen et al., "BTI and HCD degradation in a complete 32 × 64 bit SRAM array – including sense amplifiers and write drivers – under processor activity," in *Proc. of IRPS*, pp. 1-7, 2020.
- [26] Lee N-H, Baek D, Kang B. Effect of off-state stress and drain relaxation voltage on degradation of a nanoscale nMOSFET at high temperature. *IEEE Electron Device Lett* 2011;32(7):856–8.
- [27] Lee N-H, Kim H, Kang B. Impact of off-state stress and negative bias temperature instability on degradation of nanoscale pMOSFET. *IEEE Electron Device Lett* 2012; 33(2):137–9.
- [28] M. Simicic, P. Weckx, B. Parvais, P. Roussel, B. Kaczer and G. Gielen, "Understanding the impact of time-dependent random variability on analog ICs: from single transistor measurements to circuit simulations," in *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 27, no. 3, pp. 601-610, March 2019.
- [29] Duan S, Sai G. BTI aging-based physical cloning attack on SRAM PUF and the countermeasure. In *Analog Integrated Circuits and Signal Processing*. 2023.
- [30] P. Saraza-Canflanca et al., "A smart SRAM-Cell array for the experimental study of variability phenomena in CMOS technologies," in *Proc. of IRSP*, pp. P3-1-P3-5, 2022.
- [31] Saraza-Canflanca P, et al. Design considerations of an SRAM array for the statistical validation of time-dependent variability models. In: *In Proc. of SMACD*; 2018. p. 73–6.
- [32] T. Grasser, B. Kaczer, W. Goes, Th. Aichinger, Ph. Hehenberger, and M. Nelhiebel, "A two-stage model for negative bias temperature instability," in *Proc. of IRPS*, pp. 33–44, 2009.
- [33] Mahapatra S, editor. *Recent Advances in PMOS Negative Bias Temperature Instability: Characterization and Modeling of Device Architecture. Material and Process Impact*. Singapore: Springer; 2022.
- [34] R. Maes, V. Rozic, I. Verbauwheide, P. Koeberl, E. van der Sluis and V. van der Leest, "Experimental evaluation of Physically Uncloneable Functions in 65 nm CMOS," in *Proc. of ESSCIRC*, pp. 486–489, 2012.
- [35] *Failure Mechanisms and Models for Semiconductor Devices*, JEDEC JEP122G, 2011.
- [36] A. Lange, F. A. V. Gonzalez, I. Lahbib, and S. Crocoll, "Comparison of modeling approaches for transistor degradation: Model card adaptations vs subcircuits," in *Proc. of ESSDERC*, pp. 186–189, 2019.
- [37] Diaz-Fortuny, et al. A versatile CMOS transistor array IC for the statistical characterization of time-zero variability, RTN, BTI, and HCI. *IEEE J Solid State Circuits* 2019;54(2):476–88.
- [38] A. Spessot et al., "Impact of off state stress on advanced high-K metal gate NMOSFETs," in *Proc. of ESSDERC*, pp. 365-368, 2014.
- [39] Zhang J, Gao R, Duan M, Ji Z, Zhang W, Marsland J. Bias temperature instability of MOSFETs: physical processes, models, and prediction. *Electronics* 2022;11:1420.