

Using Novel Configuration Techniques for Accelerated FPGA Aging

Tanner Gaskin, Hayden Cook, Wesley Stirk, Robert Lucas, Jeffrey Goeders, and Brad Hutchings

Department of Electrical and Computer Engineering
Brigham Young University, Provo, Utah, USA

Abstract—In this work we demonstrate a novel method of accelerating FPGA aging by configuring the FPGA to implement thousands of short circuits, resulting in high on-chip currents and temperatures. Three ring oscillators are placed across the chip and are used to characterize the operating frequency of the FPGA fabric. Over the course of several weeks of running the short circuits, with daily characterization of the FPGA performance, we measured a decrease in FPGA frequency greater than 5%. After aging, the FPGA part was repeatedly characterized during a two week idle period. Results indicated that the slowdown did not change, and the aging appeared to be permanent.

In addition, we demonstrated that this aging could be induced in a non-uniform manner. In our experiments, the short circuits were all placed in the lower two-thirds of the chip, and one of the characterization ring oscillators was placed at the top of the chip, outside of the region with the short circuits. The fabric at this location exhibited a 1.36% slowdown, only one-quarter the slowdown measured in the targeted region.

I. INTRODUCTION

Through normal operation of semiconductor devices, the performance characteristics of transistors gradually decline, resulting in decreased maximum clock speeds. This performance degradation, referred to as *transistor aging*, is the result of several physical mechanisms (*negative bias temperature instability (NBTI)*, *electromigration (EM)*, and more), and is generally a greater concern when scaling to smaller technology nodes [1]. Field-programmable gate arrays (FPGAs) are not immune to this effect, and several studies have measured how FPGA performance is affected by transistor aging [2]–[4].

Understanding FPGA aging is critical. Of highest importance, FPGA design tools must account for aging mechanisms, to ensure the reported f_{max} is achievable over the entire lifetime of a part. However, aging also plays a role in other aspects of FPGA tools. In [4], Dogan et. al. show that FPGA aging models can be used as an effective predictor of recycled parts, and in [5], Maiti et al. demonstrate how aging disrupts the reliability of physical unclonable functions (PUFs). Several works have shown how aging can be rapidly induced, inflicting years of wear-out on parts in a short time period [2], [5]. This is typically accomplished by raising the supply voltage above normal levels, and baking the part in high temperatures.

In this work we induce accelerated FPGA aging by loading an FPGA with a bitstream containing thousands of short-circuits. Once loaded, the FPGA device sinks current in excess

of 7.9 A and self-heats to temperatures over 170 °C. During typical experiments, the FPGA is left in this stressed state for a period of 24 hours. Afterwards, a characterization bitstream is loaded onto the FPGA to determine how much additional delay has been induced by the excessive heat and current. The characterization bitstream uses ring oscillators (ROs) to measure relative increases in delay induced by the accelerated aging process. Aging experiments repetitively load the shorting and characterization bitstreams to measure the accelerated aging process over time. Over the course of several weeks of these aging experiments, we are able to increase circuit delay by over 5%. In addition, the impact appears permanent; after a 16 day idle period, the fabric was still operating up to 5% slower than before the induced aging process.

Our approach is unique in that it does not require modifications to the supply voltage or an external heat source. Potentially, this technique could be used to perform aging remotely, without physical control of the FPGA; however, this would require that the system be equipped with a relatively high-amperage power supply. Of perhaps more interest, our technique can apply aging non-uniformly across the die, something that is not possible with previously published techniques. For example, our experiments demonstrated that the fabric outside of the region targeted with short circuits experienced only one-fourth the slowdown compared to the targeted region. To our knowledge this is the first work that has demonstrated targeted, nonuniform aging on any commodity semiconductor device, not just FPGAs. We believe that as FPGAs continue to be used in more applications, including safety critical IoT devices, and environments where users can deploy FPGA bitstreams on remote machines (Amazon EC3), understanding and planning for these aging techniques is of high importance.

The main contributions and novelty of this work are:

- Demonstrating how FPGA bitstreams containing short-circuits can induce accelerated FPGA aging.
- A high precision experiment that demonstrates up to a 5% slowdown of the FPGA fabric, even after a 16 day recovery period.
- The first work to demonstrate significant non-uniformity in aging; FPGA fabric outside of the targeted region exhibited only one-fourth the slowdown.

In this paper, Section II discusses background material, Section III presents our techniques to induce and measure aging, and Section IV presents our experimental results. Section V

T. Gaskin, H. Cook, W. Stirk, R. Lucas, J. Goeders, and B. Hutchings are also with the NSF Center for Space, High-performance, and Resilient Computing (SHREC)

discusses conclusions and future work.

II. BACKGROUND

A. Aging mechanisms

Prior work has established several aging mechanisms pertinent to VLSI technologies: hot carrier injection (HCI), electromigration (EM), and Negative Bias Temperature Instability (NBTI) [6]–[10]. As further explained in Section IV-F, only two of these aging mechanisms are relevant in this study.

NBTI is a common problem in MOSFET technology, and has become prevalent as technology nodes have gotten smaller [1]. As an electric field is applied across the gate oxide of a MOSFET, the threshold voltage will increase slowly overtime, although more so in pMOS than nMOS. As the threshold voltage increases, the switching speed of the transistors decrease. This effect is more pronounced with higher temperatures and higher voltages [6], [8]–[10].

Electromigration (EM) is the gradual migration of metal atoms over time due to electric currents. As this migration occurs, the metal atoms will slowly accumulate at one end of the channel, thus narrowing the channel or wire. This results in slower transistor switching speeds due to increased resistances. EM is accelerated by both high DC current and high temperatures [6]–[10], both of which are the primary stressors introduced by our aging technique.

B. Related Work

Our work is not the first to introduce short circuits in the FPGA bitstream, nor is it the first to induce FPGA aging; however, it is the first published work to combine the two ideas, resulting in the ability to perform localized aging via configuration.

Beckhoff et al. demonstrate how early partial reconfiguration tools could result in short circuits that increase current and power consumption [11]. Hadzid et al. discuss how short circuits in FPGAs could cause them to operate in an unsafe operating region, potentially exceeding the power supply capability and disrupting operation, or perhaps even damaging the part, though actual damage was not shown [12]. Hutchings et al. show that short circuits can intentionally be inserted into FPGAs and that the relationship between the the number of short-circuits and power consumption is highly linear [13].

Previous aging studies have primarily used a combination of over-voltage and high temperatures to induce NBTI and HCI and accelerate aging [5], [10], [14]. Stott et al. raise the voltage and temperature by 83% and 120% of their nominal values respectively to accelerate aging on an Altera Cyclone III device. They observe a 15% slowdown over 75 days with this method [10]. Maiti et al. demonstrate a slowdown on a Xilinx Spartan 3e device using two stress phases. The first increases voltage and temperature by 25% and 180% respectively, producing a slowdown of approximately 5.0% after 200 hours. The second increases the voltage and temperature by 50% and 220% over nominal respectively, resulting in a total slowdown of approximately 6.7% after an additional 200 hours. In between phases they have a day long recovery

period where the slowdown recovers 0.5%, from 5.0% to 4.5% [5]. Slimani et al. demonstrate aging on a Xilinx Artix-7 part. This is the same FPGA that we use in our experiment. They were able to achieve a 1.8% slowdown by increasing the temperature by 400% of the nominal value for 14 days. They have an eight day recovery period where the slowdown recovers 1.0%, from 1.8% to 0.8%. [14].

As far as we are aware, only one other previous work has discussed using configuration to produce a slowdown. Chakraborty et al. discuss the idea of using ring oscillators to heat up the board and cause a slowdown, but do not validate their model through experimentation [15].

III. INDUCED AGING TECHNIQUE

In the next section we describe our experimental aging process, which consists of performing an initial characterization of the device using ring oscillators, then repeatedly stressing the chip with short circuits, stopping for a short period to re-characterize the device, and repeating. In this section we describe the system infrastructure and CAD flow that enables this experiment.

A. Short Circuits

We accelerate the aging of FPGAs by introducing short circuits into the FPGA fabric. A short circuit, in the context of FPGA designs, is when the output of an FPGA primitive (LUT or FF) driving a logic-1 is connected to the output of a primitive driving a logic-0. Two primitives are selected that connect to the same routing multiplexer, and the two multiplexer inputs are both activated (in legal circuits only one input to a routing mux would ever be active at a time).

In each short circuit of our experiment, we connect a LUT to it's associated FF, as our testing indicated that this combination of primitives created the highest current draw. We configure the LUT as logic-0, and FF as logic-1, although we have determined the reverse produces the same current draw. For each LUT-FF pair, there are several neighboring routing multiplexers that could be used to complete the short circuit; we tested each and chose the one that produced the largest current draw. This configuration results in $380\ \mu\text{A}$ per short. Within a slice (four LUTs, eight FFs), we implement four shorts, each on an associated LUT/FF pair. A diagram of our short circuit is shown in an inset diagram in Figure 1. This circuit is then replicated in as many different tiles as desired. We used a configuration containing 20,798 shorts for the experiments in this paper.

Implementing thousands of short circuits simultaneously in the FPGA fabric causes high current flow, which in turn causes temperatures to rise well past safe thresholds.

B. Ring Oscillators

We used ring oscillators (ROs) to test the intrinsic speed of the FPGA fabric, and measure the effect of these short circuits in causing device degradation. ROs are purely combinational circuits that are built using a ring of an odd number of inverters, causing a pulse to travel around the ring at the

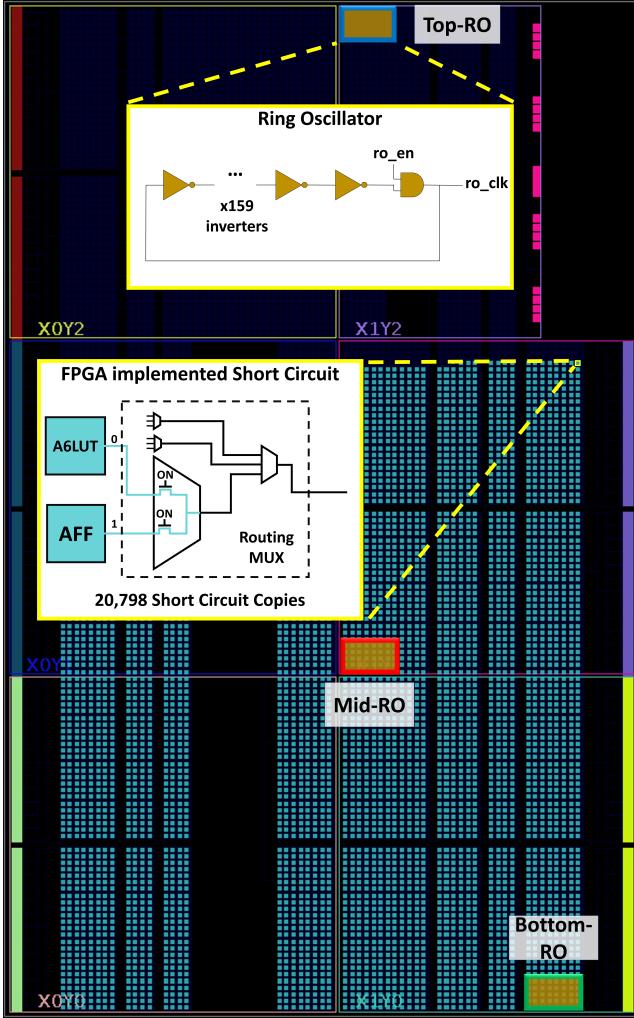


Figure 1: Three ring oscillators are used to measure slowdowns across the chip, during the “characterization” phase of the experiment. 20,798 shorts are used during the “burn” phase of the experiment, covering the lower two-thirds of the part. While both are shown in this figure, they are configured at different times.

maximum speed supported by the FPGA fabric. A logical representation of an RO is included as an inset image in Figure 1. As shown, our ROs contain a single AND-gate that is used to disable and enable the oscillation. A wire is tapped off of the RO, labelled *ro_clk*, that is then fed into the clock input of a large counter, to measure the oscillation frequency of the RO. In our experiment we use three ROs, placed at the top, bottom, and center of the chip to measure the speed of the FPGA fabric before and after introducing the short circuits.

C. System Design and CAD Flow

While Figure 1 shows our short circuits and characterization ROs on the same floorplan, in reality, they are part of two different bitstreams, configured on the FPGA at separate times.

1) *Characterization Bitstream*: The characterization design is responsible for measuring the intrinsic delay of the FPGA fabric before and after short-circuit burns. It includes a static region, restricted to the top-left clock tile, which contains a MicroBlaze, hardware timers, UART, and custom control hardware. It also contains a black box for the entire right-side of the chip, that allows for the insertion of one of our partial bitstreams. There are three partial bitstreams, each one containing one of the three ring oscillators (since they are configured and measured sequentially, not simultaneously). A fourth partial bitstream is blank, and used during the recovery phase at the end of our experiment.

Each RO partial design contains a single RO that is 159 inverters long. The *ro_clk* signal is connected to a counter in the static region, and the *ro_en* signal is controlled from hardware in the static region. This custom hardware enables the RO and counter for an exact number of cycles, measured by a timer in the static region, ensuring that characterization is always performed for the exact same duration of time.

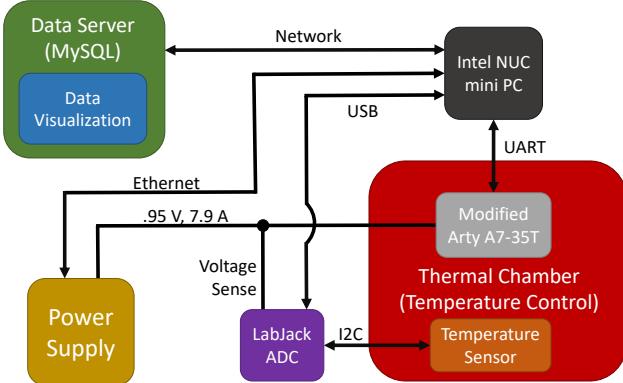
The bitstreams are created using a combination of the Vivado and RapidWright tools. RapidWright is an open-source tool from Xilinx that can manipulate netlists to provide very fine placing and routing control. It interfaces with Vivado via design checkpoints [16]. The static region is designed in Vivado. For each RO, a partial design is created in RapidWright that contains specific placement of the RO inverters. This is then imported into Vivado to be routed with the static region, and a static bitstream and partial bitstream are created. This process is then repeated for the other two ring oscillators, and the blank region, producing the remaining partial bitstreams.

2) *Short Circuit Bitstream*: The short-circuit design does not contain any of the previously described characterization elements, but only 20,798 short circuits, as described previously. The design process of implementing LUTs and FFs with constant outputs, shorted together, is completed using custom scripts in Rapidwright. This allows us to bypass Vivado’s restriction on routing a net with multiple drivers. Once the short circuit design is created in Rapidwright and saved to a design checkpoint, it is then loaded into Vivado. We then turn off all design rule checks (DRCs), which forces Vivado to ignore our shorted nets, and then use Vivado to generate a bitstream that contains all of the short circuits.

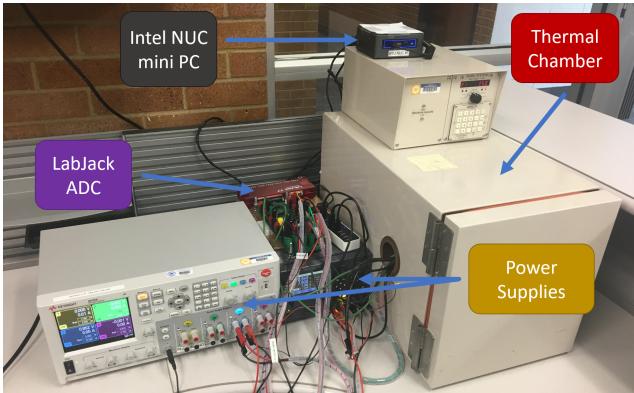
IV. EXPERIMENT AND RESULTS

A. Experimental Setup

Our experiments iteratively sequence through the following steps: (1) load the short-circuit configuration into the FPGA, (2) allow the FPGA to “burn” in this configuration for 24 hours, and (3) load the characterization bitstream to determine increases in circuit delay. The characterization process uses the ring oscillators as described in Section III-B. However, circuit delay and, in turn, RO frequency are not just sensitive to FPGA aging, they are also highly sensitive to both the supply voltage and ambient temperature. To guarantee that our measurements accurately reflected aging, and were not influenced by voltage fluctuations or changes in ambient temperature, we created



(a) A diagram of the experimental setup.



(b) The experiment is kept in a controlled environment by a thermal chamber and a high precision power supply.

Figure 2: Experiment Setup

a test bed that allowed us to have precise control over and detailed measurement/confirmation of both supply voltage and ambient temperature.

Figure 2 illustrates our test setup. In order to control environmental temperature we performed all experiments inside of a thermal chamber. A temperature sensor, located inside the chamber, was monitored using an Intel NUC mini PC and LabJack ADC, and the measurements were recorded in a MySQL database 10 times per second. When set to 35°C this chamber kept the ambient temperature of the board at 35 ± 0.086 °C, as shown in the last column of Table I¹.

The test board was an Arty A7-35T development board with an Artix XC7A35 FPGA. The board was modified such that the *V_{INT}* supply voltage to the FPGA was supplied by a Keysight power supply (E36231A). The external power-supply controls voltage with a much higher degree of accuracy and also supplies higher current than the on-board power regulator. The external power supply is configured to supply 0.95 V to the Arty board for internal logic and used sense lines to compensate for voltage drop over the supply lines.

¹Note that the the thermal chamber only provides a stable ambient temperature and does not play a role in the accelerated aging process.

The NUC PC polled the power supply for voltage and current values over Ethernet, and logged them to the database at a rate of twice per second. As can be seen in Table I, this setup kept the voltage at 0.95 ± 0.00052 V for the duration of the experiment. In addition to the voltage and current reported by the power supply, we also used a 0.1Ω 0.05% shunt resistor, monitored by the NUC via the LabJack, to verify the reported values.

B. Experimental Procedure

An initial characterization consisted of measuring the frequency of the three ROs at the top, middle and bottom of the chip (see Figure 1). Each of the three ROs were characterized three times, six hours apart, to ensure that repeated measurements produced the same frequency results.

After the initial characterization the burn period was started, which consisted of 36 burn cycles. At the beginning of each burn cycle the bitstream containing 20,798 short circuits was configured onto the FPGA. The shorts are distributed across the bottom two-thirds of the chip, as can be seen in Figure 1. After a period of 24 hours, the shorts were removed from the board, and a one hour characterization period was performed. Over the course of the hour, each ring oscillator was programmed onto the board for 20 minutes. The first seven minutes were allotted to allow temperature fluctuations to settle, and then the ring oscillator frequency was measured once per second for 13 minutes.

Following this, another 24 hour burn and one hour characterize cycle would begin. Once 36 of these burn/characterize cycles had taken place, we began a 16 day recovery period, to investigate whether the aging effects were temporary, or long lasting. Over the recovery period we would repeatedly leave the blank bitstream on the FPGA for 24 hours and then perform the same one-hour characterization process.

C. Results

The results of this experiment are summarized in Table I, and the slowdown of the ROs is also plotted in Figure 3. Due to resource constraints, this experiment was only performed on a single Arty A7-35T development board.

a) Slowdown: As can be seen, the maximum slowdown achieved was by the Mid-RO (red), which is located in the middle of the region of short circuits, and slowed by 5.13%. The Bottom-RO (green), also in the short circuit region, but at the corner of the chip, slowed by 4.81%. The Top-RO (blue), located far from the short circuit region, had the least amount of degradation, with a slowdown of only 1.35%.

It can also be seen in Figure 3 that the degradation effects are heavily weighted towards the beginning of the burn period. After the first burn cycle Mid-RO experienced a 1.25% decrease in frequency and by the end of the sixth cycle the slowdown had reached 2.56%, half of the total slowdown seen during the burn period. This early drop-off is also seen in previous FPGA aging work [2], [5], [8], [10], [14].

It should be noted that the test platform allows for these slowdown measurements to be very precise and consistent.

Table I: Experiment results after every five days

Time	Top-RO Slowdown	Mid-RO Slowdown	Bottom-RO Slowdown	Input Current During Burn	Input Voltage During Characterization	Ambient Temp. During Characterization
0 Days	0.00 % \pm 0.020	0.00 % \pm 0.023	0.00 % \pm 0.021	N/A	N/A	N/A
5 Days	0.74 % \pm 0.028	2.39 % \pm 0.020	2.43 % \pm 0.017	$7.94 \pm 7.84 \times 10^{-3}$ A	$0.95 \pm 5.15 \times 10^{-4}$ V	35.053 ± 0.086 °C
9 Days	0.92 % \pm 0.025	3.02 % \pm 0.019	2.95 % \pm 0.021	$7.94 \pm 8.50 \times 10^{-3}$ A	$0.95 \pm 5.15 \times 10^{-4}$ V	35.133 ± 0.051 °C
14 Days	1.00 % \pm 0.028	3.50 % \pm 0.018	3.36 % \pm 0.067	$7.93 \pm 9.48 \times 10^{-3}$ A	$0.95 \pm 5.15 \times 10^{-4}$ V	35.109 ± 0.074 °C
18 Days	1.12 % \pm 0.021	3.91 % \pm 0.024	3.72 % \pm 0.022	$7.92 \pm 7.52 \times 10^{-3}$ A	$0.95 \pm 5.15 \times 10^{-4}$ V	35.116 ± 0.051 °C
23 Days	1.22 % \pm 0.022	4.33 % \pm 0.027	4.08 % \pm 0.021	$7.92 \pm 7.52 \times 10^{-3}$ A	$0.95 \pm 5.15 \times 10^{-4}$ V	35.139 ± 0.051 °C
27 Days	1.21 % \pm 0.023	4.60 % \pm 0.020	4.30 % \pm 0.019	$7.92 \pm 7.84 \times 10^{-3}$ A	$0.95 \pm 5.15 \times 10^{-4}$ V	35.148 ± 0.059 °C
32 Days	1.35 % \pm 0.024	4.93 % \pm 0.018	4.62 % \pm 0.015	$7.91 \pm 7.52 \times 10^{-3}$ A	$0.95 \pm 5.15 \times 10^{-4}$ V	35.115 ± 0.055 °C
36 Days	1.35 % \pm 0.027	5.13 % \pm 0.016	4.81 % \pm 0.019	$7.90 \pm 8.50 \times 10^{-3}$ A	$0.95 \pm 5.15 \times 10^{-4}$ V	35.139 ± 0.082 °C

Recovery Period						
1 Days	1.39 % \pm 0.022	5.15 % \pm 0.017	4.81 % \pm 0.018	$0.05 \pm 1.386 \times 10^{-2}$ A	$0.95 \pm 5.15 \times 10^{-4}$ V	34.950 ± 0.078 °C
8 Days	1.39 % \pm 0.019	5.17 % \pm 0.017	4.81 % \pm 0.020	$0.05 \pm 1.386 \times 10^{-2}$ A	$0.95 \pm 5.15 \times 10^{-4}$ V	34.914 ± 0.082 °C
16 Days	1.39 % \pm 0.016	5.17 % \pm 0.016	4.80 % \pm 0.018	$0.05 \pm 1.387 \times 10^{-2}$ A	$0.95 \pm 5.15 \times 10^{-4}$ V	34.941 ± 0.070 °C

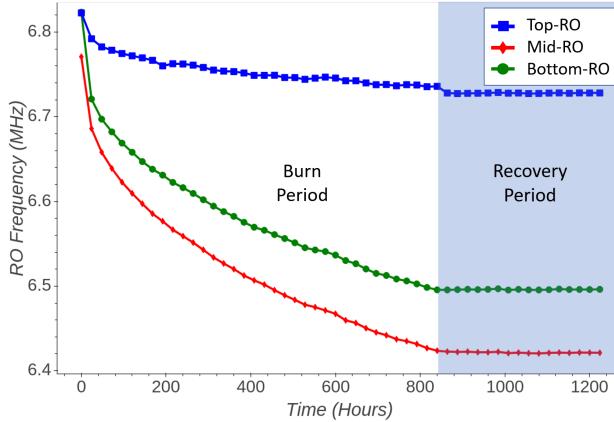


Figure 3: After each cycle the Ring Oscillators were re-characterized. Those results are shown above, which track the slowdown of each RO as time progressed. It is clearly shown that the Top-RO (outside of the burn region), experienced far less slowdown than the two ROs inside the burn region.

During each characterization, the frequency of each RO is measured once per second for 13 minutes, resulting in 780 measurements. A histogram of these measurements at each burn cycle for the Mid-RO is provided in Figure 4. For example, in the initial characterization (shown in the inset of Figure 4), the measured RO frequency ranges from 6.769 to 6.772 MHz, a span of only 0.023%. In the characterization after the final burn cycle, the frequency range had dropped to 6.423 - 6.425 MHz, with a similar distribution and span (0.016%). The narrow ranged, non-overlapping histograms demonstrate that this variance remains small for each characterization cycle, and is indicative of the precision and consistency of the test platform. While Figure 4 illustrates this variance for only the Mid-RO, the variances for each of the three ROs throughout the experiment are listed in Table I.

b) *Temperature*: During the burn period the steady state temperature of the chip, according to the on-chip temperature sensor, was 177.7°C, more than 70°C higher than the

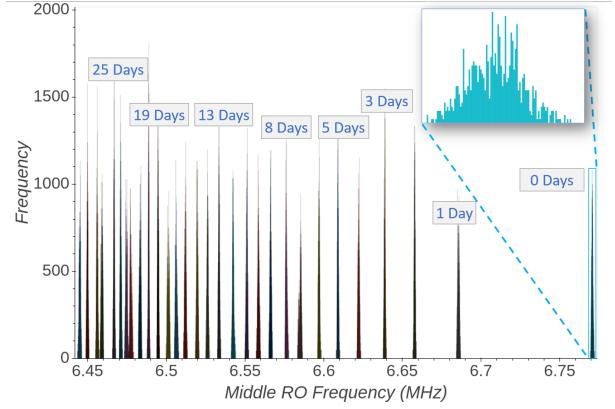


Figure 4: Characterization histograms of Mid-RO.

maximum temperature rating of 100°C for the device on the Arty board [17]. Furthermore, due to the heat spreading capabilities of ICs and the distance between the on-board temperature sensor and the transistors in the burn region, this value only represents a lower bound of the temperature actually experienced by the transistors in the burn region. The actual junction temperature is likely to be much higher.

c) *Current*: The other primary effect of the short circuits is the high current they induce, resulting in a steady state burn current exceeding 7.9 A. An interesting trend in the data is that the steady state current draw decreased over the length of the burn period. As shown in both Table I and Figure 5, the average current dropped from 7.94 A at the beginning of the experiment, to 7.90 A at the end, a 0.04 A drop. While not large, the decrease was a consistent trend and therefore noteworthy. This continuously diminishing current draw is indicative of damage being done to the FPGA fabric. The spikes shown in Figure 5 occur at the beginning of each burn cycle as the temperature settles.

D. Localized Degradation

One of the unique aspects of using short circuits to accelerate aging is the ability to target specific parts of the

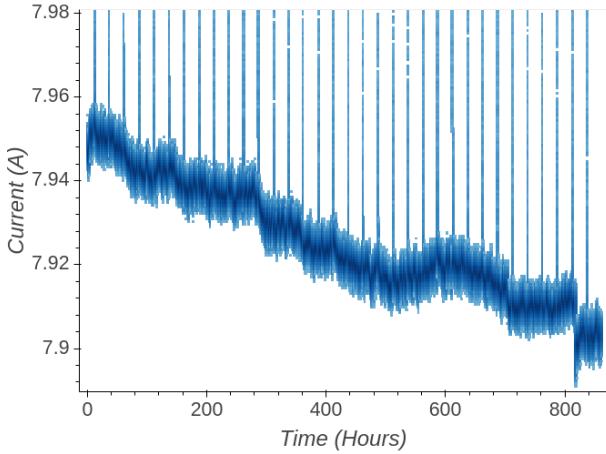


Figure 5: VINT input current during burn period.

fabric, resulting in non-uniform aging. In Figure 3 it can clearly be seen that the two ROs inside the burn region (Mid-RO and Bottom-RO) are heavily degraded due to the short circuits. However, the Top-RO shows a dramatically different aging profile. While the Mid-RO and Bottom-RO slowed down by 5.13% and 4.81% respectively, the Top-RO reached a maximum slowdown of 1.35%, only one-fourth of the slowdown.

Additionally, the slowdown pattern for the Top-RO was more volatile than the slowdown patterns for the other two ROs. The two ROs in the burn region continuously experienced degradation effects for the entire length of the burn period. However, it can be seen that the Top-RO would fluctuate in regards to how much it degraded each cycle. While the overall trend for the Top-RO was towards increased delay, there were several instances where circuit delay decreased during a burn cycle, something that never occurred for the ROs inside the burn region. We leave further analysis of this localized effect to future work.

E. Recovery Period

Another key part of accelerated aging techniques is understanding how permanent the aging effects are. Prior work has shown that up to 56% of the slowdown can be lost during a recovery period [14], showing that recovery periods are a critical part of understanding accelerated aging mechanisms.

We measured the permanency of our aging technique by allowing the FPGA to rest for 16 days, longer than most previous work. Each RO was re-characterized every 24 hours to measure any recovery effects. This method ensured that degradation was not continuing to be accelerated by running the ROs for the entire length of the cycle but also allowed for measurement of recovery effects.

The FPGA did not experience any recovery, as shown in Table I and Figure 3, demonstrating that the aging induced by short circuits is permanent.

F. Aging Effects

While the results of accelerated aging have been clearly demonstrated, determining the cause of the aging is far more difficult. However, there are several well established aging mechanisms that can be discussed as possible contributors to the FPGA degradation. We hypothesize that NBTI and EM are the primary causes of aging in our experiments, as HCI is accelerated by high switching speeds, which doesn't play a role in our experiment [2]–[4], [10].

a) Negative Bias Temperature Inversion: NBTI, as presented in Section II-A, becomes more pronounced in the presence of increased temperature and increased voltage. While we demonstrated in Section IV-A that no external factor could be causing accelerated NBTI degradation, in Section IV-C we show that the short circuits cause on-chip temperatures to exceed maximum ratings by an excess of 70 °C. In [14] it is shown that the temperature only needs to be raised to 125 °C for the Artix 7 part to experience pronounced NBTI effects. As such, we feel that NBTI is at least one of the factors causing the accelerated aging we observed in our experiments.

This claim is further supported by the pronounced initial slowdown. Similar results, where the initial degradation is the most pronounced, are reported in [2], [5], [8], [10], [14], all of which examined the effects of NBTI on FPGAs.

b) Electromigration: EM, as presented in Section II-A, becomes more pronounced in the presence of increased temperature and high DC current. In our experiment we have shown currents in excess of 7.9 A and temperatures in excess of 170 °C, providing the environmental conditions needed for accelerated EM. Furthermore, we show in Section IV-D that the magnitude and pattern of the Top-RO aging varies significantly from the other two ROs, while also showing that the primary difference between the Top-RO and the other two ROs is the presence of high DC current. This leads to the conclusion that EM, the only established aging mechanism accelerated by high DC current, is a primary cause for the differences seen in the Top-RO aging profile.

V. CONCLUSION

This paper introduced a new technique for accelerated aging on FPGAs. By configuring the FPGA with tens of thousands of short circuits, we generated an excessive amount of current and heat, well outside normal operating ranges. This aging technique is novel in that it is independent of environmental conditions, such as input voltage and ambient temperature.

Using this technique we have been able to successfully slowdown the Xilinx Artix-7 fabric by over 5% in 36 days. Compared to previous techniques, this method appears to provide significant aging capability, which aging is shown to be permanent. Most interestingly, this new technique is shown to create nonuniform aging effects in the FPGA fabric, where fabric outside the short circuit region exhibited only one-fourth of the slowdown. These results open the door to a broad spectrum of future work related to additional aging mechanisms, security implications, and more.

REFERENCES

- [1] S. Novak, C. Parker, D. Becher, M. Liu, M. Agostinelli, M. Chahal, P. Packan, P. Nayak, S. Ramey, and S. Natarajan, “Transistor Aging and Reliability in 14nm Tri-Gate Technology”, in *Reliability Physics Symposium*, Apr. 2015, 2F.2.1–2F.2.5.
- [2] A. Amouri, F. Bruguier, S. Kiamehr, P. Benoit, L. Torres, and M. Tahoori, “Aging effects in FPGAs: An Experimental Analysis”, in *Conference on Field Programmable Logic and Applications (FPL)*, Sep. 2014, pp. 1–4.
- [3] Z. Ghaderi and E. Bozorgzadeh, “Aging-aware High-level Physical Planning for Reconfigurable Systems”, in *Asia and South Pacific Design Automation Conference (ASP-DAC)*, Jan. 2016, pp. 631–636.
- [4] H. Dogan, D. Forte, and M. M. Tehranipoor, “Aging Analysis for Recycled FPGA Detection”, in *Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT)*, Oct. 2014, pp. 171–176.
- [5] A. Maiti, L. McDougall, and P. Schaumont, “The Impact of Aging on an FPGA-Based Physical Unclonable Function”, in *Conference on Field Programmable Logic and Applications (FPL)*, Sep. 2011, pp. 151–156.
- [6] N. Weste and D. Harris, “CMOS VLSI design: A circuits and systems perspective”, Jan. 2010.
- [7] JEDEC Solid State Technology Association and others, *JEP122C, Failure Mechanisms and Models for Semiconductor Devices*, Mar. 2006.
- [8] S. Gehrer, S. Leger, and G. Sigl, “Aging Effects on Ring-Oscillator-Based Physical Unclonable Functions on FPGAs”, in *Conference on ReConfigurable Computing and FPGAs (ReConFig)*, Dec. 2015, pp. 1–6.
- [9] S. Gehrer, “Highly Efficient Implementation of Physical Unclonable Functions on FPGAs”, 2017.
- [10] E. A. Stott, J. S. Wong, P. Sedcole, and P. Y. Cheung, “Degradation in FPGAs: Measurement and Modelling”, in *Symposium on Field Programmable Gate Arrays (FPGA)*, ser. FPGA ’10, Monterey, California, USA: Association for Computing Machinery, Feb. 2010, pp. 229–238.
- [11] C. Beckhoff, D. Koch, and J. Torresen, “Short-Circuits on FPGAs caused by Partial Runtime Reconfiguration”, in *Conference on Field Programmable Logic and Applications (FPL)*, Aug. 2010, pp. 596–601.
- [12] I. Hadžić, S. Udani, and J. M. Smith, “FPGA Viruses”, en, in *Conference on Field Programmable Logic and Applications (FPL)*, P. Lysaght, J. Irvine, and R. Hartenstein, Eds., ser. Lecture Notes in Computer Science, Berlin, Heidelberg: Springer, 1999, pp. 291–300.
- [13] B. L. Hutchings, J. Monson, D. Savory, and J. Keeley, “A Power Side-Channel-Based Digital to Analog Converter for Xilinx FPGAs”, in *Symposium on Field-Programmable Gate Arrays (FPGA)*, ser. FPGA ’14, Monterey, California, USA: Association for Computing Machinery, Feb. 2014, pp. 113–116.
- [14] M. Slimani, K. Benkalaia, and L. Naviner, “Analysis of ageing effects on ARTIX7 XILINX FPGA”, *Microelectronics Reliability*, vol. 76-77, Jul. 2017.
- [15] R. S. Chakraborty, I. Saha, A. Palchaudhuri, and G. K. Naik, “Hardware Trojan Insertion by Direct Modification of FPGA Configuration Bitstream”, *IEEE Design Test*, vol. 30, no. 2, pp. 45–54, Apr. 2013.
- [16] C. Lavin and A. Kaviani, “RapidWright: Enabling Custom Crafted Implementations for FPGAs”, in *Symposium on Field-Programmable Custom Computing Machines (FCCM)*, Apr. 2018, pp. 133–140.
- [17] Xilinx, “Artix-7 FPGAs Data Sheet: DC and AC Switching Characteristics (DS181)”, en, p. 64, 2018.