

Fault Characterization Through FPGA Undervolting

Behzad Salami*[†], Osman Unsal*, and Adrian Cristal^{1*†}

*Barcelona Supercomputing Center (BSC), Barcelona, Spain. Emails: {behzad.salami, osman.unsal, and adrian.cristal}@bsc.es

[†]Universitat Politècnica de Catalunya (UPC), Barcelona, Spain.

Abstract—The power and energy efficiency of Field Programmable Gate Arrays (FPGAs) are estimated to be up to 20X less than Application Specific Integrated Circuits (ASICs). What is needed to close this gap is aggressive power/energy savings techniques. A such potentially effective approach is undervolting, which can directly deliver an order of magnitude static and dynamic power savings. However, aggressive undervolting, without accompanying frequency scaling leads to timing related faults, potentially undermining the power savings. Understanding the behavior of these faults and efficiently mitigating them can deliver further power and energy savings in low-voltage designs. In this paper, we conduct a detailed analysis of undervolting FPGA on-chip memories (BRAMs). Through experimental analysis, we find that lowering the supply voltage until a certain conservative level, V_{min} , does not introduce any observable fault. For the studied platforms, we measure this voltage guardband gap to be 39% of the nominal level ($V_{nom} = 1V$, $V_{min} = 0.61V$). Further undervolting corrupts some of data bits stored in BRAMs; however, it also reduces the BRAMs power consumption a further 36.1%. When the voltage is lowered below V_{min} , the rate of these faults exponentially increases to 0.06%, by a fully non-uniform distribution over various BRAMs. This paper comprehensively analyzes the behavior of these faults, in terms of rate, type, location, and environmental temperature.

I. INTRODUCTION

Undervolting is a technique to decrease the supply voltage below the nominal level in order to save power and energy. Unlike Dynamic Voltage and Frequency Scaling (DVFS) [1] [2], the frequency is not scaled down in undervolting, so energy savings can be potentially significant. However, decreasing the voltage while keeping the frequency constant leads to timing related faults, which can cause applications to crash or terminate with wrong results. The severity of these faults depends on the fault rate and location as well as on application characteristics. Therefore, characterization of these undervolting faults and understanding their behavior is critical to mitigate their impact. Although, there have been some previous undervolting works on CPUs [3], Graphic Processor Units (GPUs) [4], and Dynamic RAM (DRAM) memories [5], there is no "deep-dive" undervolting fault characterization studies due to the relatively closed nature of these hardware substrates where the vendors expose few details. In comparison, the relatively open Field Programmable Gate Array (FPGA) architecture makes it possible to conduct and report such detailed studies. However, to the best of our knowledge, such studies have not been thoroughly undertaken for FPGAs. Hence, the main contributions of this paper is extensively characterizing the behavior of faults, when commercial FPGAs are aggressively undervolted. More specifically, we target the

FPGA on-chip Block RAM (BRAM) structures and report the comprehensive experimental findings on undervolting using real hardware to date. For instance, we find that: *first*, a significant voltage guardbanding of 39% exists below the nominal voltage level ($V_{nom} = 1V$, $V_{min} = 0.61V$) before faults manifest, which in turn leads to an order of magnitude power savings. However, the fault rate exponentially increases by further undervolting to a somewhat moderate 0.06% before the FPGA fails. *Second*, the BRAM undervolting fault rate decreases at higher environmental temperatures; thus, experimentally verifying the Inverse Temperature Dependence (ITD) [6]. ITD states that in undervolted nanometer technology nodes, the propagation delay is reduced in higher temperature environments that in turn, leads to a lower fault rate. This is significant since applying thermal stress will reduce the fault rate in low-voltage regions and, in turn, will lower the energy cost of the fault mitigation.

This paper is organized as follows. In Section II, we introduce the experimental setup. The overall behavior of FPGA BRAMs undervolting is described in Section III. Detailed fault characterization is discussed in Section IV. Finally, in Section V, we review the previous work.

II. EXPERIMENTAL METHODOLOGY

We perform our experiments on two FPGAs, i.e., XC7VX485T representing the Virtex7 family on a VC707 board and XC7K325T representing the Kintex7 family on a KC705 board. These FPGAs are respectively equipped with 2060 and 890 BRAMs, distributed over the chip with the size of 16 Kbits each. Each BRAM is a matrix of bitcells with 1024 rows and 16 columns. BRAMs can be either individually accessed or cascaded to build larger memories (with some overheads). This methodology provides flexibility for the FPGA designers to have single-cycle accesses to on-chip memories as per bandwidth or size need. More details of our tested platforms are shown in Table I. Both platforms are fabricated with 28nm technology, and the standard nominal voltage of BRAMs is the same, $V_{nom} = 1V$. However, their difference is that XC7VX485T (Virtex7) is designed for performance while XC7K325T (Kintex7) is optimized for the power consumption. Hence, for a thorough evaluation, we selected these representative FPGAs.

Through the Power Management Bus (PMBUS) standard [7], it is possible to independently and dynamically regulate and monitor the supply voltage of such FPGA components as BRAMs (V_{CCBRAM}) To modify V_{CCBRAM} , we use Texas

TABLE I: Specifications of Tested Platforms.

Hardware Platform (Board)	VC707	KC705
Device Family	Virtex 7	Kintex 7
Chip Model	XC7VX485T	XC7K325T
Number of BRAMs	2060	890
Basic Size of Each BRAM	1024*16-bits	1024*16-bits
Manufacturing Process Technology	28nm	28nm
Nominal V_{CCBRAM} (V_{nom})	1V	1V
Speed Grade	-2	-2

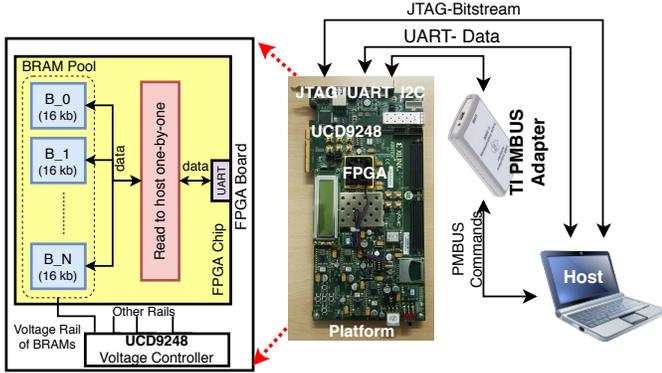


Fig. 1: Undervolting Experimental Setup in FPGA BRAMs.

List 1: Pseudo-code to Study Voltage Scaling on BRAMs, on the Experimental Setup of Fig. 1.

```

1:  $V_{CCBRAM} = V_{min} = 0.61V$ ;
2: while( $V_{CCBRAM} \geq V_{crash}$ ) begin
3:   while(numRun  $\leq$  100) begin
4:     delay(1sec);
5:     Transfer content of BRAMs to the host;
6:     Analyse faulty data (rate and location);
7:     numRun++;
8:   end
9:    $V_{CCBRAM} = V_{CCBRAM} - 10(mV)$ ;
10: end

```

Instrument (TI) PMBUS USB Adapter, and the provided C-based Application Programming Interface (API), which facilitates accessing the on-board voltage controller, with a part number of UCD9248 for our platforms, through the host [8]. The experimental setup is shown in Fig. 1. It is composed of two distinct hardware and software components. The task of the hardware FPGA platform is to access BRAMs, and transmit their contents to the host, using a serial interface. On the other side, the host issues the required PMBUS commands to set a certain voltage to V_{CCBRAM} . Also, it analyzes potentially faulty data retrieved from BRAMs. Note that we verify and validate that the implemented serial interface is entirely reliable in any V_{CCBRAM} level.

On our setup, shown in List 1, first, we initialize V_{CCBRAM} with $V_{min} = 0.61V$. Then, we retrieve contents of BRAMs one-by-one and within each BRAM row-by-row, and transfer them to the host. In the host, we analyze the rate and location

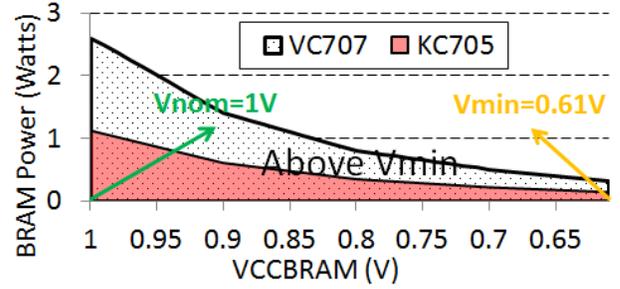


Fig. 2: BRAM Power above $V_{min} = 0.61V$ in the voltage guardband region (no observable fault).

of faults. This process is repeated 100 times for each voltage level to obtain statistically significant results. The reported results in this paper are the median of these 100 tests. After a soft reset, we gradually decrease V_{CCBRAM} by 10mV and repeat the process until the lowest voltage that our design operate, $V_{crash} = 0.54V$. For each voltage level, the fault rate and power consumption of BRAMs are recorded. Finally, to measure the power consumption with acceptable accuracy, we use a power meter, while to extract the power contribution of BRAMs in the nominal voltage level, we use Xilinx Power Estimation (XPE) tool. Note that experiments are performed on the default and fixed internal frequency of BRAMs, i.e., $\sim 500Mhz$ [9].

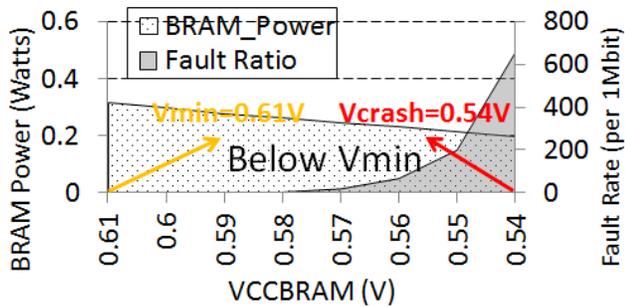
III. OVERALL BEHAVIOR ON BRAM UNDERVOLTING

Our experiments on undervolting BRAMs below nominal level, $V_{nom} = 1V$, demonstrate two thresholds. *First*, a voltage guardband, V_{min} , that separates the fault-free and faulty regions. *Second*, V_{crash} that is the lowest level of the voltage that our design practically operates, below that FPGA fails. In our test environment and for both tested platforms, $V_{nom} = 1V$ due to the factory settings, whereas $V_{min} = 0.61V$ and $V_{crash} = 0.54V$, obtained through our experiments. Also, due to our experiments with various Xilinx FPGAs at 28nm technology, we posit that $V_{crash} = 0.54V$ is strictly set by the factory to prevent device damage in extremely low voltages.

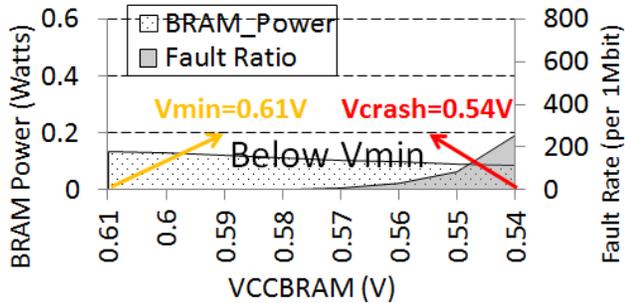
When $V_{CCBRAM} \geq V_{min}$, no observable faults occur. When $V_{CCBRAM} = V_{min} = 0.61V$ we observe significant BRAM power savings over $V_{nom} = 1V$, more than an order of magnitude for both platforms including the sum of static and dynamic power, without incurring any reliability degradation, as shown in Fig. 2. Further lowering V_{CCBRAM} below V_{min} the fault rate exponentially increases, while the power consumption is reduced, as summarized in Fig. 3. As can be seen, both power consumption and reduction are less in KC705 than VC707, which is the consequence of having relatively less BRAMs and also the inherent power optimizations adopted for KC705 by the vendor.

IV. FAULT CHARACTERIZATION

In this section, we characterize the behavior of faults in detail, when V_{CCBRAM} is underscaled from $V_{min} = 0.61V$



(a) VC707



(b) KC705

Fig. 3: BRAM Power and Fault Rate Behavior, undervolting below $V_{min} = 0.61V$ until $V_{crash} = 0.54V$.

to $V_{crash} = 0.54V$. As can be seen in Fig. 3a and 3b, the fault rate in VC707 is relatively more than KC705, by 47.4% on average. This difference can be the consequence of the architectural and technological differences adopted to optimize performance and power in VC707 and KC705, respectively.

A. Fault Stability Over Time

As earlier mentioned, we repeat each test 100 times to get statistically significant results. We do not observe a significant difference among runs, and the standard deviation of the fault rate and locations among these runs is negligible. More details about these 100 runs are summarized in Table. II.

B. Fault Variability Among BRAMs

By statistically analyzing experimental results, we observe that fault rates in voltage regions below $V_{min} = 0.61V$, considerably varies among BRAMs. For instance, experimenting on VC707 at $V_{crash} = 0.54V$, the maximum, minimum, and average fault rate within BRAMs are 2.84%, 0%, and 0.04%, respectively. Also, 38.9% of BRAMs in VC707 and 45.2% in KC705 never experience faults, at the lowest voltage level $V_{crash} = 0.54V$. As a further analysis, we cluster this statistical information in low-, mid-, and high-vulnerable classes of BRAMs, using the k-means clustering algorithm. As can be seen in Table. III, for instance, the vast majority of BRAMs in VC707, 88.6%, is classified as low-vulnerable, with an average fault rate of 0.02%, ~ 3.4 faults within an individual BRAM with the size of 16-Kbits. This significant

TABLE II: Fault Rate Stability Over Time. Fault Rate Analysis of 100 runs at $V_{crash} = 0.54V$ with pattern=16'hFFFF.

Parameter	VC707	KC705
AVERAGE fault rate (per 1 Mbit)	652	254
MINIMUM fault rate (per 1 Mbit)	630	240
MAXIMUM fault rate (per 1 Mbit)	669	264
STANDARD DEVIATION of fault rates	7.3	4.8

TABLE III: Clustering BRAMs due to Vulnerability Features (Reported numbers are at $V_{crash} = 0.54V$).

		VC707	KC705
low-vulnerable	%BRAMs	88.6%	93.4%
	Average Fault Rate (%)	0.02%	0.01%
mid-vulnerable	%BRAMs	9.4%	5.7%
	Average Fault Rate (%)	0.24%	0.17%
high-vulnerable	%BRAMs	1.8%	0.9%
	Average Fault Rate (%)	0.86%	0.74%

fault variability among BRAMs can be due to either the chip-dependent process variation or design tools for place and route. We verify this argument by performing the following test; for our test design, we extract the fault rate of BRAMs with several place-and-route compiles. Repeating the voltage lowering operation on these various bitstreams, we observe almost an identical fault rate in almost identical physical locations of BRAMs. Hence, we conclude that this fault rate variability among BRAMs is the result of the process variation.

C. Impact of the Environmental Temperature

We perform an experiment to study the effect of the environmental temperature on the behavior of faults when V_{CCBRAM} is lowered below V_{min} . Toward this goal, the hardware board is placed inside a chamber that its temperature can be regulated by a heater. We monitor the on-board temperature using PMBus commands. Through experiments, BRAMs fault rates are extracted and shown in Fig. 4 under various on-board temperatures, 50°C (default temperature), 60°C, 70°C, and 80°C. As can be seen, with heating up, the fault rate constantly reduces; for instance, by more than 3X in VC707, with the temperature is increased from 50°C to 80°C. This observation is the consequence of the Inverse Thermal Independence (ITD) property [6]. ITD is a thermal property of digital devices with nano-scale technology nodes; and states that under ultra low-voltage operations, the circuit delay reduces at higher temperatures. The reason is that as the technology node scales down, the supply voltage approaches the threshold voltage. Hence, at low-voltage regimes, increasing the temperature reduces the threshold voltage and allows the device to switch faster. In turn, with the circuit delay decreasing, the number of critical paths, and subsequently, the fault rate reduces. This property is experimentally verified in our case, for commercial FPGAs. Also, as can be seen, the fault rate in VC707 is reduced more aggressive than KC705. A relatively 156% more fault rate in 50°C is reduced to 11.6% less fault rate in 80°C, for VC707 vs. KC705. The architectural and technological difference between these platforms can be the reason since

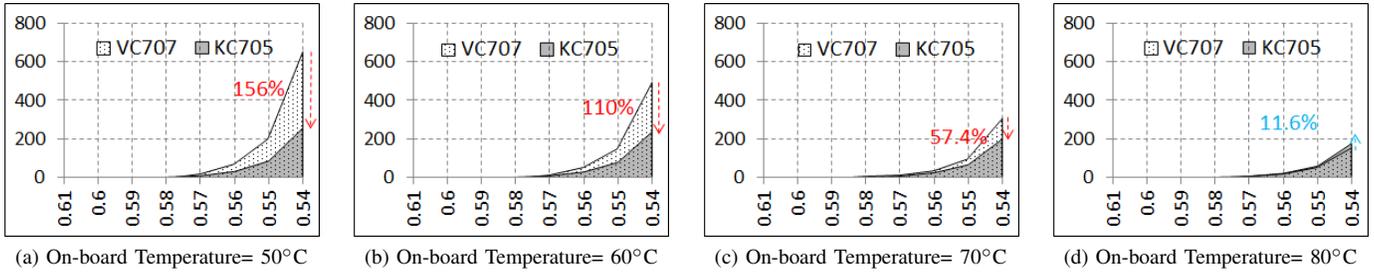


Fig. 4: The Correlation among on-board Temperature, Supply Voltage of BRAMs, Technology (VC707 vs. KC705), and Fault Rate. (x-axis: V_{CCBRAM} from $V_{min} = 0.61V$ to $V_{crash} = 0.54V$, y-axis: the fault rate (per 1Mbit))

their design goal is different, i.e., performance (VC707) vs. power (KC705).

V. RELATED WORK

Most commercial devices are designed with a voltage guardband below the standard minimum nominal supply voltage to ensure the correct functionality in worst environmental and process variation cases. This voltage guardband is fully vendor- and system-dependent; for instance, it is measured to be 20% in GPUs [10] and 16% in DRAMs [5]. We experimentally determine the voltage guardband for Xilinx FPGA BRAMs to be 39%, which in turn, delivers more than an order of magnitude power savings.

Also, tackling with the increased delay in low-voltage regions below V_{min} , accompanying frequency undervolting is a promising solution [1]; however, it can limit the energy savings gain. A more aggressive approach is to allow designs to experience timing faults and in turn, tolerating these faults. Characterizing these faults can allow better power and reliability trade-offs. Among real hardware undervolting evaluations, modern processors are extensively studied [11], [12], [13], [14]; however, there are several recent efforts on other hardware devices, as well, i.e., GPUs [4], ASICs [15], [16], and memory systems, i.e., DRAMs [5] and SRAMs [17]. In parallel, several simulation-based framework [18] or design optimization are also proposed to study undervolting through nano-meter technology parameters [19]; however, it is evident that this approach lacks the exact information of the fault model under very low-voltage operations and their validation on the silicon remains a key question. Our paper studies undervolting for the first time in commercial FPGAs, concentrated on on-chip BRAMs, and extensively evaluates the reliability aspects of BRAMs in low-voltage regions, accounting effects of environmental temperature, as well.

Finally, we are working on effective mitigation techniques. Toward this goal, in addition to evaluating the adaptation of existing techniques, e.g., ECC [20], HTM [21], we are working on customized methods, as well, accounting the observed behavior of faults discussed in this paper.

ACKNOWLEDGMENT

We thank Pradip Bose, Alper Buyuktosunoglu, and Augusto Vega from IBM Watson for their contribution to this work. The

research leading to these results has received funding from the European Union's Horizon 2020 Programme under the LEGaTO Project (www.legato-project.eu), grant agreement n° 780681.

REFERENCES

- [1] Nunez-Yanez, et al. "Energy optimization in commercial FPGAs with voltage, frequency and logic scaling", in *IEEE TC*, 2016.
- [2] G. Semeraro, et al. "Energy-efficient processor design using multiple clock domains with dynamic voltage and frequency scaling", in *HPCA*, 2002.
- [3] R. Bertran, et al. "Very Low Voltage (VLV) Design", in *ICCD*, 2017.
- [4] J. Tan, et al. "Combating the reliability challenge of GPU register file at low supply voltage", in *PACT*, 2016.
- [5] K. K. Chang, et al. "Understanding reduced-voltage operation in modern DRAM devices: Experimental characterization, analysis, and mechanisms", in *Measurement and Analysis of Computing Systems*, 2017.
- [6] Neshatpour, K., et al. "Enhancing Power, Performance, and Energy Efficiency in Chip Multiprocessors Exploiting Inverse Thermal Dependence", in *IEEE TVLSI*, 2018.
- [7] "Power Management Bus (PMBUS)." <http://pmbus.org>
- [8] Texas Instruments (TI), "Fusion Digital Power Designer". http://www.ti.com/tool/FUSION_DIGITAL_POWER_DESIGNER
- [9] Xilinx, "7 Series FPGAs Memory Resources." https://www.xilinx.com/support/documentation/user_guides/ug473_7Series_Memory_Resources.pdf
- [10] J. Leng, et al. "Safe limits on voltage reduction efficiency in GPUs: a direct measurement approach", in *MICRO*, 2015.
- [11] A. B. Kahng, et al. "Designing a processor from the ground up to allow voltage/reliability tradeoffs", in *HPCA*, 2010.
- [12] K. Swaminathan, et al. "Bravo: Balanced reliability-aware voltage optimization", in *HPCA*, 2017.
- [13] G. Papadimitriou, et al. "Harnessing voltage margins for energy efficiency in multicore CPUs", in *MICRO*, 2017.
- [14] A. Bacha, et al. "Dynamic reduction of voltage margins by leveraging on-chip ECC in Itanium II processors", in *ISCA*, 2013.
- [15] G. Tziantzioulis, et al. "b-HiVE: A bit-level history-based error model with value correlation for voltage-scaled integer and floating point units", in *DAC*, 2015.
- [16] P. N. Whatmough, et al. "14.3 a 28nm SOC with a 1.2 ghz 568nj/prediction sparse Deep-Neural-Network engine with > 0.1 timing error rate tolerance for IOT applications", in *ISSCC*, 2017.
- [17] L. Yang, et al. "SRAM voltage scaling for energy-efficient convolutional neural networks", in *ISQED*, 2017.
- [18] J. Zhang, et al. "ThUnderVolt: Enabling Aggressive Voltage Underscaling and Timing Error Resilience for Energy Efficient Deep Neural Network Accelerators", *arXiv:1802.03806*, 2018.
- [19] B. Reagen, et al. "Minerva: Enabling low-power, highly-accurate deep neural network accelerators", in *ACM SIGARCH Computer Architecture News*, 2016.
- [20] G. Yalcin, et al. "Exploiting a fast and simple ECC for scaling supply voltage in level-1 caches", in *IOLTS*, 2014.
- [21] G. Yalcin, et al. "FaultM: error detection and recovery using hardware transactional memory", in *DATE*, 2013.