Impact of Voltage Scaling on Soft Errors Susceptibility of Multicore Server CPUs

Dimitris Agiakatsikas*
Department of Informatics, University of Piraeus
Greece
dagiakatsikas@unipi.gr

George Papadimitriou*
Department of Informatics & Telecomm., University of Athens
Greece
georgepap@di.uoa.gr

Vasileios Karakostas
Department of Informatics & Telecomm., University of Athens
Greece
vkarakos@di.uoa.gr

Dimitris Gizopoulos
Department of Informatics & Telecomm., University of Athens
Greece
dgizop@di.uoa.gr

Mihalis Psarakis
Department of Informatics, University of Piraeus
Greece
mpsarak@unipi.gr

Camille Bélanger-Champagne
Ewart Blackmore
TRIUMF, Vancouver, BC
Canada
cbchampagne@triumf.ca

ABSTRACT

Microprocessor power consumption and dependability are both crucial challenges that designers have to cope with due to shrinking feature sizes and increasing transistor counts in a single chip. These two challenges are mutually destructive: microprocessor reliability deteriorates at lower supply voltages that save power. An important dependability metric for microprocessors is their radiation-induced soft error rate (SER). This work goes beyond state-of-the-art by assessing the trade-offs between voltage scaling and soft error rate (SER) on a microprocessor system executing workloads on real hardware and a full software stack setup. We analyze data from accelerated neutron radiation testing for nominal and reduced microprocessor operating voltages. We perform our experiments on a 64-bit Armv8 multicore microprocessor built on 28 nm process technology. We show that the SER of SRAM arrays can increase up to 40.4% when the device operates at reduced supply voltage levels. To put our findings into context, we also estimate the radiation-induced Failures in Time (FIT) rate of various workloads for all the studied voltage levels. Our results show that the total and the Silent Data Corruptions (SDC) FIT of the microprocessor operating at voltage-scaled conditions can be 6.6× and 16× larger than at the nominal voltage, respectively. Moreover, changes in the microprocessor's clock frequency do not have a noticeable impact on its soft error susceptibility. The findings of this work can aid computer architects in striking a balance between power and dependability, thus, designing more robust and efficient microprocessors.

CCS CONCEPTS

• **Computer systems organization** → **Reliability**; • **Hardware** → **Transient errors and upsets**; **Power and energy**.

KEYWORDS

Microprocessor reliability, error resilience, silent data corruptions, soft errors, energy efficiency, voltage and frequency scaling, power consumption, neutron radiation testing

*Equally contributing first authors.

ACM Reference Format:
Dimitris Agiakatsikas, George Papadimitriou, Vasileios Karakostas, Dimitris Gizopoulos, Mihalis Psarakis, Camille Bélanger-Champagne, Ewart Blackmore. 2023. Impact of Voltage Scaling on Soft Errors Susceptibility of Multicore Server CPUs. In *56th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO '23), October 28-November 1, 2023, Toronto, ON, Canada.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3613424.3614304

1 INTRODUCTION

Device miniaturization drives the integration of more and faster circuits on modern microprocessors. However, this comes at the expense of increased power consumption, elevating energy efficiency as a major design issue. The power consumption of a microprocessor is affected by its supply voltage and operating frequency [39]. Semiconductor vendors set the nominal supply voltage of microprocessors to a pessimistic level to account for 1) static fluctuations due to manufacturing and environmental conditions and 2) dynamic variations of workloads and device degradation [32, 39, 50]. Another major factor that affects power consumption is operating supply voltage [39]. The operating voltage of a microprocessor is strongly affected by the technology node, the static fluctuations (due to manufacturing and environmental conditions), and the dynamic variations (caused by workloads differences and device degradation) [39]. Therefore, to ensure correct execution in all circumstances, the nominal voltage is set to a pessimistic level (i.e., higher than it is actually required) [32, 50].

Exploiting the pessimistic voltage guardbands of microprocessors and unveiling the safe voltage levels beyond the nominal conditions is an effective way to reduce power consumption [5, 49, 57]. Recently, there have been several studies on the limitations of power consumption in modern microprocessors, especially when multi-thread applications are being executed [5, 49, 57]. Reducing the supply voltage can result in considerable power savings without affecting the performance since voltage has a quadratic relationship to the dynamic power (i.e., $P = aCV^2f$, where V is the supply voltage, and f is the clock frequency) [76]. However, with aggressive supply voltage reduction, several circuits may fail due to manufacturing-induced parameter variations [21]. Additionally, voltage overscaling renders circuits more vulnerable to radiation effects [29].

Cosmic radiation is a significant source of errors in modern microprocessor chips, making device reliability issues a critical concern for several decades [9, 12, 13, 15, 52, 54]. Particularly, high-energy neutrons interact with the silicon die, creating a secondary cascade of charged particles. These can create current pulses that invert the values stored either in SRAM arrays (e.g., cache memories) or produced by combinatorial logic; this phenomenon is called a soft error or single-event upset (SEU). Although soft errors are not permanent, their effects may propagate to the system's output and cause malfunctions from minor process crashes to the most severe class of silent data corruptions (SDCs) [55, 56, 69].

Device miniaturization and reduced voltage margins have further increased the vulnerability of integrated circuits to transient radiation effects. This happens because state-of-the-art circuitry carries smaller charges and, thus, can be more easily upset [16, 29, 75]. A lower supply voltage may make the chip more prone to radiation effects because the charge required to upset a node is proportional to the voltage level [16]. Recent studies have shown that voltage scaling techniques beyond nominal conditions reduce power consumption, but also increase failures due to timing violations in the control logic [75]. These failures are typically mitigated by combining voltage overscaling with error recovery mechanisms, such as checkpointing [26]. Semiconductor vendors mitigate soft errors in CPUs with error recovery mechanisms, which introduce overheads and negatively affect power consumption. This can negate the efficacy of supply voltage reduction techniques. Therefore, it is unclear whether energy savings from reduced voltage margins outweigh the overhead of error recovery mechanisms.

Reliability evaluation studies usually consider failures either due to reduced supply voltage levels of the microprocessor chip [49, 57, 75] or due to radiation-induced phenomena [13, 18]. However, radiation-induced effects may severely impact the chip operation at lower voltage levels [22, 27, 79]. A slight increase in the Soft Error Rate (SER) may deteriorate dependability, especially under voltage scaling conditions [14, 82, 84]. Prior research [67] analyzed through simulation how voltage scaling affects the neutron-induced raw SER of microprocessors' SRAM structures (i.e., cache memories). This work goes beyond state-of-the-art *by assessing the trade-offs between voltage scaling and SER on a microprocessor system executing workloads on real hardware and a full software stack setup.* By going beyond simulations and analyzing the effects of upsets in a real-world scenario all the way to the software layer, we can better understand the potential consequences of soft errors that may result in application SDCs or system crashes.

Specifically, we examine the effects of atmospheric-like neutron radiation on a modern Arm-based microprocessor operating at reduced voltage: the Applied Micro's 8-core CPU used in X-Gene 2 server. We present and evaluate data obtained from beam experiments (more than 64 beam hours) conducted at various reduced safe voltage settings beyond the nominal voltage conditions (on real CPU hardware). We show that the upsets per minute for the SRAM arrays of the microprocessor chip are increased by up to 40.4% when the chip operates at reduced voltage conditions. To put our findings into context, we also estimate the radiation-induced Failures in Time (FIT) of various workloads for all tested voltages, assuming operation in New York City (NYC) at sea level. According to our FIT rate results, the probability of SEUs resulting in SDCs when the microprocessor operates at low-voltage conditions is more than 16× larger than nominal voltage conditions, while the total FIT rate is 6.6× larger. Our analysis also reveals that although voltage levels substantially influence the microprocessor's SER, clock frequency changes do not have any impact. This work provides valuable insights that can assist computer architects in making informed decisions in designing efficient microprocessors while reducing the risks of SDCs and system crashes related to low supply voltages.

2 BACKGROUND

2.1 Radiation Effects and Error Rate

The interaction of a galactic cosmic ray with Earth's atmosphere triggers a flux of particles (mainly neutrons), from which some reach the ground [35]. A highly energized neutron strike may perturb the state of one or more transistors, generating bit upset faults (i.e., bit-flips) in memory elements or event transients in combinatorial logic that can be potentially latched in memory elements of a microprocessor system. Such events, in turn, may lead to errors, such as a wrong application output or a system crash. Specifically, a bit upset fault 1) may not affect the application's output (i.e., the fault gets logically masked, or the corrupted data is not used) and 2) may be propagated from the hardware to the software layers, resulting in an SDC (i.e., application output is corrupted without any indication), application crash (i.e., program hang), or system crash (i.e., the device becomes unresponsive or reboots unexpectedly). The error rate of an application executed on a microprocessor depends on its sensitivity to memory and logic radiation effects [8, 46], as well as the probability for the fault to be propagated from the microarchitecture and the software (i.e., the application) layers to the system output [53, 70].

Radiation testing is the most effective way to evaluate essential dependability metrics of microprocessor systems. The dynamic cross-section (DCS) is commonly used to estimate the vulnerability of a microprocessor to radiation-induced events (e.g., memory upsets, SDCs, application crashes, or system crashes) under a certain workload, configuration, and environment conditions. In other words, DCS is a metric that shows the likelihood of a radiation-induced event occurring when highly-energized particles collide with the microprocessor, and it is given from the equation below:

$$\text{Dynamic Cross-Section (DCS)} = \frac{\text{Number of Events}}{\text{Particle Fluence}}, \quad (1)$$

where *Particle Fluence* defines the number of particles passed through cm^2 area of the chip. The larger the DCS, the more susceptible the microprocessor will be to radiation-induced memory upsets and errors. Characterizing the microprocessor's DCS under a certain workload makes it straightforward to calculate the expected rates of memory upsets, SDCs, and application/system crashes for a given radiation environment. For example, the average neutron particle flux in New York City (NYC) at sea level is approximately 13 neutrons per cm^2 per hour (valid for the integrated fluence at neutron energies > 10MeV), which yields the following FIT rate:

$$\text{FIT} = \text{DCS} \times \frac{13 \text{ neutrons}}{\text{cm}^2 \times \text{hour}} \times 10^9 \text{ hours}, \quad (2)$$

that is, the average number of failures that occur within one billion hours of device operation [35].

2.2 SRAM Failures due to Voltage Scaling

SRAM cells can malfunction for several reasons, ranging from environmental conditions, aging, and supply voltage disturbances to process and system variation, leakage, etc. The process change can produce a cell-to-cell shift because depending on which transistor is affected can lead to different types of failure and different voltage thresholds for the cell. This variation is mainly observed when the SRAM dies operate at near-threshold voltage. SRAM failure modes can be summarized as read, write, read stability, and fault retention [86]. A read error occurs when the read discharge time takes longer than the sense amplifier bias, and at the end of a cycle, there is not enough voltage difference between the bit lines. Therefore, the stored value cannot be retrieved. Write failure occurs if the internal node voltage does not reach the desired write-value, resulting in an erroneous stored value. A stable read error occurs when the contents of a memory cell are accidentally reversed during a read. Finally, the hold error occurs when the operating voltage of the CPU is lower than its memory's data hold voltage. A low supply voltage can cause all these failures, creating faulty cells inside the SRAM array.

3 EXPERIMENTAL METHODOLOGY

3.1 Server Platform

We use an 8-core 64-bit out-of-order Armv8-compliant microprocessor (a custom variant of the Cortex A72 core) that offers high-end processing performance and targets server systems. The microprocessor includes a subsystem that features a Power Management processor (PMpro) and a Scalable Lightweight Intelligent Management processor (SLIMpro) to enable flexibility in power management and enhance security. The dedicated SLIMpro processor uses an I^2C interface to communicate with system sensors and peripherals to monitor and configure the system attributes, such as supply voltage and the DRAM refresh rate. It also gathers health status reports, such as soft error events in the microprocessor's L1, L2, and L3 caches. Appropriate drivers are used to access the SLIMpro.

X-Gene 2 has three voltage domains, i.e., the Processor Module Domain (PMD), the System on Chip (SoC) Domain, and the Standby Power Domain, as shown in Figure 1. These domains can be independently regulated to specific voltage levels. This study focuses on the PMD and the SoC domains. The PMD domain includes four dual-core processors interconnected via a Central SWitch (CSW). Each core has private instruction (L1I) and L1D (data) caches, and every core-pair shares a unified L2 cache. The operating voltage of all cores in the PMD can only be changed together (we cannot modify the voltage of a core individually) with a granularity step of 5 mV, beginning from 980 mV. In contrast, we can change the frequency of each dual-core processor in the PMD. The frequency for each dual-core processor ranges from 300 MHz to 2.4 GHz, with a 300 MHz step granularity. The SoC domain contains an L3 cache and DRAM controllers. The voltage of the SoC domain can be independently scaled downwards, with a granularity of 5 mV beginning from 950 mV. Table 1 presents the main characteristics of the microprocessor used in this study.

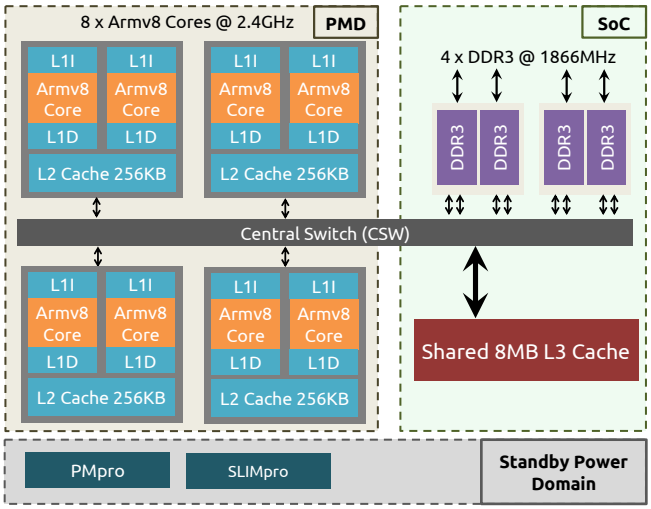

Figure 1: X-Gene 2 microprocessor block diagram.

The X-Gene 2 platform runs CentOS 7.3 with Linux kernel version 4.11. We performed experiments using two different frequencies to provide diverse results: 2.4GHz, the highest available frequency, and 900 MHz.

Note that the Dynamic Voltage and Frequency Scaling (DVFS) of the microprocessor is not enabled during our experiments. DVFS uses nominal voltage levels for each different frequency. We aim to keep the clock frequency at a certain level to analyze the SER of the microprocessor with reduced but safe supply voltages.

Table 1 also provides information regarding the RAS features of X-Gene 2. All levels of cache memories are protected through parity protection (L1 instruction and data caches) or through Single Error Correction Double Error Detection (SECDED) Error Correction Codes (ECC) (L2 and L3 caches). This means that if any upset occurs in these structures (and can be corrected or left uncorrected), it is logged as an upset attributed to a certain cache level. The behavior of the hardware error detection and correction schemes used by the X-Gene 2 is as follows. For SRAM arrays with parity protection only, when the corrupted entry is detected, it is invalidated. The next attempted usage of the entry is canceled and restarted, and

Parameter	X-Gene 2 Server CPU
ISA	Armv8 (AArch64)
Pipeline / CPU Cores	64-bit OoO (4-issue) / 8
Clock Frequency	2.4 GHz
D/I TLBs	20 entries per core (Parity)
Unified L2 TLB	1024 entries per core (Parity)
L1 Instruction Cache	32 KB per core (Parity)
L1 Data Cache	32 KB Write-Through per core (Parity)
L2 Cache	256 KB Write-Back per pair of cores (SECDED)
L3 Cache	8 MB Write-Back Shared (SECDED)
TDP / Technology	35 W / 28 nm
PMD/SoC Nominal Voltage	980 mV / 950 mV

Table 1: X-Gene 2 microprocessor specifications.

a correct new entry is retrieved through the standard cache/TLB-miss handling mechanisms because of the write-through policy of these arrays. This means that single-bit upsets (SBUs) in L1 cache memories are always corrected and do not affect the program's execution. On the other hand, the SECDED protected SRAM arrays (i.e., L2 and L3 caches) detect up to two-bit upsets and correct one SBU per 64-bit words [33]. Thus, SECDED provides a fine-grained reporting mechanism since it can also detect multi-bit upsets (i.e., reported as "uncorrected errors").

3.2 Analyzing a 28 nm microprocessor

The analysis we deliver in this paper uses a 28 nm microprocessor because our special development platform has all the features we need to control the voltage and observe the effects. To our knowledge, no similar platforms for Arm-based systems are implemented in newer technology nodes. Moreover, ever since its introduction in 2011 [1], the 28 nm process technology has been improving and maturing. The development of advanced semiconductor manufacturing has now reached the 5 nm mass production stage [72]. Nonetheless, there is still significant demand for established, mature process technologies, such as 28 nm technology nodes [48, 68].

3.3 Benchmarks

In our analysis, we use a popular parallel benchmark suite for high-performance computing aiming to maximize the utilization of the multicore CPU of our system: the NAS Parallel Benchmarks (NPBs) [7]. NPBs are programs designed to evaluate the performance of parallel supercomputers. We use six benchmarks (CG, EP, FT, IS, LU, and MG) from the NAS Parallel Benchmarks (NPBs) suite [7]. We consider the multicore version of NPBs and perform experiments using all 8 cores. We use the class A of NPBs to have a short execution time (< 5 sec), in order to avoid the accumulation of radiation-induced faults per benchmark run. In other words, we avoid reproducing in the accelerated radiation environment the non-realistic case of multiple radiation-induced events (i.e., multiple soft errors) during the execution of a single benchmark run.

Our preliminary analysis showed that this could be achieved if the execution time of the benchmarks was less than 5 sec. More specifically, previous radiation experiments on a 28nm Multi-Processor System on Chip (MPSoC) showed that its SRAM memory cross-section is in the range of $10^{-15} cm^2$/bit [83]. On the other hand, the flux at TRIUMF Neutron irradiation Facility (TNF) is $\sim 2.5 \times 10^6 \text{neutrons/cm}^2/s$. Therefore, the 28 nm X-Gene 2 microprocessor chip, which is built on a similar 28 nm process node as that MPSoC, is expected to experience one memory upset per 4.8 sec during the radiation tests, assuming 10 MB of on-chip SRAM memory.

3.4 Neutron Beam Experiments

One of the most accurate methods for determining the error rates of devices and applications is accelerated radiation testing. Since the entire chip is irradiated, there is no way to contain faults within a limited set of hardware resources, which is the case in simulation-based fault injection frameworks. The neutron-induced SER of a Design Under Test (DUT) can be measured 1) either by exposing hundreds or thousands of DUTs to natural radiation until adequate soft errors have been detected to give a confident estimate of the SER, or 2) by performing accelerated radiation testing to expose one DUT to high-intensity radiation and collect data in a shorter period to provide statistically significant results [35]. We follow the latter approach in this paper. Our accelerated radiation testing experiments were performed at TRIUMF's TNF in Canada. The TNF beam is designed to study Single Event Effects (SEEs) and thus optimized to extract a neutron spectrum as similar as possible to the atmospheric neutron reference spectrum defined in JEDEC's open standards [35]. Figure 2 shows part of our setup in TRIUMF. To irradiate the X-Gene 2 microprocessor (i.e., our DUT), we mounted the server board on a metallic frame, as shown in the right-side image of Figure 2. We then slid down the frame into the beam zone through the TNF's access channel (as shown in the left-side image of Figure 2).

On the first day of the radiation campaign, we attempted to conduct the tests by directly placing the DUT in the center of the beam path. However, the TNF's beam flux was too intense, and the DUT was experiencing consecutive system boot failures, prohibiting the collection of meaningful results. To overcome this problem and allow the successful execution of the tests, we lowered the beam flux and, therefore, the radiation-induced failure rate of the DUT. Specifically, TRIUMF uses activation foil methods on a yearly basis to estimate the absolute neutron flux in the nominal 5 cm × 12 cm beam spot of the TNF beam [10]. Under typical conditions with a 100 μA proton beam current on the neutron production target, the neutron flux at the test position is in the range of 2×10^6 to 3×10^6 neutrons/cm$^2/s$, for neutrons with energies above $10 MeV$, and cannot be reduced due to other operational constraints of the facility. We raised the DUT up 5-10 cm in the TNF's access channel to position it in the halo of the neutron beam instead of directly in the beam path. This reduced the neutron flux on the DUT to a level where its failure rate allowed the execution of the test.

We performed a relative neutron beam intensity measurement at the newly defined test position using TRIUMF's SRAM-based "golden board" radiation dosimeter [11]. We measured the SEU rate of the dosimeter in the beam center and compared it with six measurements of the SEU rate at the beam halo test position. We moved the DUT up and down the access channel between measurements to account for any mechanical positioning uncertainty at the halo position, which, unlike the nominal test position, is not defined using a built-in mechanical stop. We calculated the flux of neutrons

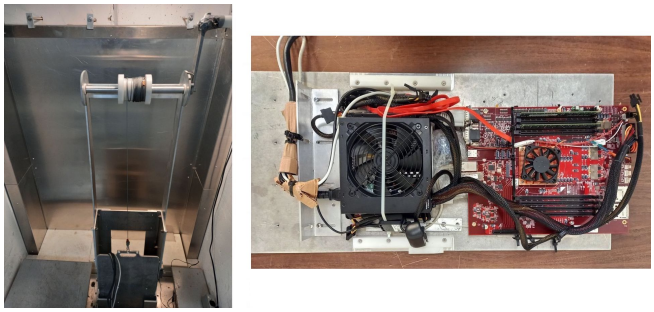

Figure 2: Beam test setup at TRIUMF (left image) and the X-Gene 2 board (right image) ready for beam testing.

with energies ≥ 10 *MeV* at the halo measurement position from the ratio of the SEU rate measurements to be 0.60 ± 0.02%, where the uncertainty quoted is the combination of statistical and position uncertainty obtained from the combination of the six measurements at the halo position. This uncertainty is small compared to the uncertainty on the absolute measurement of the TNF neutron flux, which is on the order of 20% [10]. In a nutshell, we tested the DUT at room temperature, in which the X-Gene2 operated in 40-45°C as verified through periodic temperature and power measurements, with an average neutron flux above 10 MeV of $(2+3)/2 \times 0.6 \times 10^6 = 1.5 \times 10^6 \text{neutrons/cm}^2/s$. Note that our experiments are performed in a temperature-aware manner, as we observed during the offline characterization that the safe V_{min} was not affected up to 50 °C.

We should note that there could be a variation in neutron spectrum in the beam halo compared to the beam center, and a small portion of upsets could be attributed to low-energy thermal neutrons. However, the thermal neutrons account for only about 15% of the ≥10 MeV flux for the upstream production targets in the beamline configuration used.

3.5 Test Sessions

Table 2 summarizes the four test sessions that we performed during the 64-hour radiation campaign at TNF. As shown in the following sections, we tested the X-Gene 2 microprocessor chip under nominal (980 mV) and reduced voltage conditions. As a rule of thumb, statistically significant results can be obtained with radiation tests when the total fluence per test session is greater than 1×10^{11} neutrons/cm^2 [28] or when 100 or more, radiation-induced events get accumulated [65]. Test sessions 1 and 2 (see Table 2) lasted approximately 27 hours, achieving a high fluence ($> 1.5 \times 10^{11}$ neutrons/cm^2). Test session 3 finished shorter than sessions 1 and 2 because we accumulated more than 100 events for errors (SDCs, application and system crashes) and memory upsets. Test session 4 lasted only 165 minutes because the beam time we had reserved was elapsed. Therefore, the statistical significance of that specific experiment is not as high as that of the other beam test session experiments, shown in Table 2.

The fifth row of Table 2 indicates the equivalent period a device needs to be naturally irradiated with neutrons at NYC sea level to receive the same amount of fluence we achieved with the accelerated radiation testing. Rows 6 to 9 show an overview of the SDC, crashes and memory upsets observed in the experiment and the corresponding error rates, which exhibit an increasing soft error vulnerability when the device operates at lower voltages. This trend will be analyzed in detail in the following sections.

To prove the soundness of the experiment, we compare the observed SER with previously published radiation data for 28 nm devices. In [83], the authors studied the atmospheric neutron SEEs on a 28 nm system on chip (SoC) in the nominal voltage and calculated a total SER of 15 FIT per Mbit at Beijing sea level (in NYC, sea level is 13.18) for all the tested memories of the chip. The tenth row of Table 2 shows the SER observed in our experiment, which ranges between 2.08 and 2.45 FIT per Mbit. This SER is slightly lower than the one observed in [83]. Their difference can be attributed to the

Beam test session	1	2	3	4
Voltage Levels (mV)	980	930	920	790
Test duration (minutes)	1651	1618	453	165
Fluence (neutrons/cm^2)	1.49E+11	1.46E+11	4.08E+10	1.48E+10
Years of NYC equivalent radiation	1.30E+06	1.28E+06	3.58E+05	1.30E+05
SDCs and crashes (#)	95	97	141	13
SDCs and crashes rate (per min)	5.75E-02	5.99E-02	3.11E-01	7.87E-02
Memory upsets (#)	1669	1743	506	195
Memory upsets rate (per min)	1.011	1.077	1.117	1.182
Memory SER (FIT per MBit)	2.08	2.22	2.30	2.45

Table 2: Neutron Beam Time Sessions at TRIUMF/TNF.

nature of the benchmarks used in the two experiments. In the aforementioned paper, the authors run a static memory test program that exhaustively tests the memories for upsets, while in our case, the benchmarks neither access the memories regularly nor occupy the entire caches, and thus several upsets are never detected, e.g., the affected memory words are not used or overwritten before read. Thus, it is reasonable that we observe a lower error rate. Note also that, the error margins in the graphs of the accelerated radiation results in this paper assume a confidence level of 95%. The error margins are depicted as error bars, and their absolute numbers are not shown in the figures as labels.

3.6 Test Flow

Prior to the radiation tests, we extensively characterized the 8-core X-Gene 2 processor chip as proposed in [57] and [49] to expose the safe voltage points for each target frequency. The identified safe V_{min} for each frequency allowed a fault-free execution of all benchmarks. *Therefore, any detected errors during the radiation experiments are attributed to neutrons and not to the reduced supply voltage of the processor.* After finding the safe V_{min} for each voltage configuration, we evaluated the impact of the reduced voltage margins on the soft error susceptibility of the microprocessor by performing the following methodology. As shown in Figure 3, a Control-PC located in the control room orchestrates the experiments of the X-Gene 2 located in the beam room. The Control-PC 1) controls, monitors, and collects data from the server, and 2) remotely resets/power cycles (turn ON/OFF) the server during an SDC, application crash, or

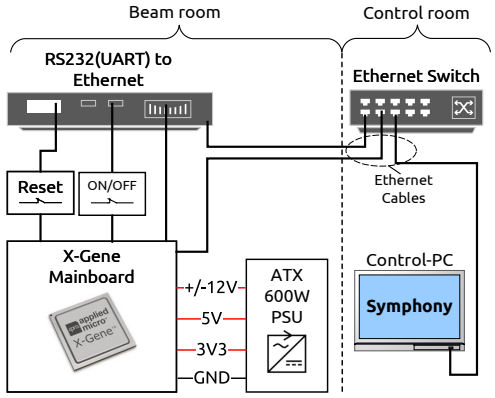

Figure 3: Experimental setup.

system crash; in general, in any unexpected behavior. The connection between the server and the Control-PC is established through Ethernet and RS-232 serial communication links.

During the experiments, SDCs are detected by comparing the application output with its golden reference, i.e., the expected output pre-computed in fault-free conditions. Any mismatch between the experimental and the expected output is marked as an SDC and logged for post-analysis. Additionally, during execution, the X-Gene 2 is connected through a network interface to a control PC to indicate the correct function of the application, as shown in Figure 3. Application and system crashes are detected through response timeouts. Specifically, if after a given period, the DUT is unresponsive, an attempt is made to contact the board and restart the application. If the attempt is successful, the event is logged as an application crash (Linux is still running and responding). When the connection with the server cannot be established, the event is logged as a system crash.

4 SUSCEPTIBILITY ANALYSIS OF SOFT ERRORS

In this section, we first explore the behavior of various multi-threaded workloads to 1) unveil the minimum safe voltage levels so that all workloads operate reliably under reduced supply voltage, and 2) analyze and understand their soft error susceptibility on a low-power microprocessor through accelerated neutron radiation testing.

4.1 Microprocessor Safe Voltage Levels

This part of the paper focuses on a quantitative analysis of the safe V_{min} to expose the potential guardbands of the chip, as well as to quantify what determines the V_{min} of a multicore application in a real hardware experiment. The voltage guardband is the difference between the nominal voltage of the microprocessor and its safe V_{min}. We experimentally obtained the safe V_{min} values of six NAS Parallel benchmarks running on the X-Gene 2 based on the methodology defined in [49, 57]. To maintain a high statistical significance in our experimentation, we ran the entire undervolting experiments hundreds of times for each benchmark and on each frequency. We performed our experiments in two different clock frequencies of the microprocessor chip: 2.4 GHz, the highest available frequency, and 900 MHz. Note that any changes in voltage levels for a given clock frequency do not affect the performance, according to our experiments.

Figure 4 shows the V_{min} characterization results of the NAS Parallel Benchmark suite for 2.4 GHz (left graph) and 900 MHz (right graph) clock frequencies. Note that V_{min} is the lowest safe voltage setting, where executions complete successfully without errors or any abnormal behavior. As shown in both graphs of Figure 4, as the voltage gets reduced (x-axis), the probability of failure (pfail) gets higher (y-axis). Pfail=100% means that all identical executions failed to complete. On the other hand, a pfail=10% means that there are 90% chances for an application to execute correctly in that voltage.

In the left graph of Figure 4, we see that 920 mV is the lowest safe voltage (i.e., V_{min}), in which all benchmarks operate reliably. Beyond this voltage level, the probability of failure increases gradually until 900 mV (i.e., 20 mV lower than the V_{min}), where the probability of failure (pfail) is 100%. However, if we change the frequency to 900 MHz (right graph), we see that the lowest safe voltage level is 790 mV, and the failure probability period beyond this level is shorter (i.e., 10 mV) than the 2.4 GHz clock frequency. Thus, the settings for the PMD domain are 920 mV and 790 mV for 2.4 GHz and 900 MHz clock frequency, respectively. The SoC domain is at the nominal conditions (see Table 3), since the frequency change cannot affect the SoC domain, but only the PMD domain. The V_{min} for all benchmarks is the same for each clock frequency. This is in line with previous studies, which demonstrated that in multicore executions, the workload variation is negligible for the safe V_{min} [49]. In our experiments, we change the voltage levels of both PMD and SoC domains at the same time. The voltage levels are within the safe voltage margins and are shown in Table 3.

4.2 SRAM Upsets Rates

We now present the rate of upsets per minute observed in the SRAM structures of X-Gene 2 during the radiation experiments. Since our study is based on a commercial chip, and thus the observability is extremely limited, we leverage the inherent SRAM protection methods of X-Gene 2 to observe and identify the SRAM upsets. We employ the Linux EDAC (Error Detection And Correction) driver [2], which typically collects the error protection notifications from the hardware and forwards them to the user through the Linux *dmesg* logs. In this way, we report any uncorrected upset error (UE) or corrected upset error (CE) reported by the microprocessor's EDAC mechanisms.

Based on the above mechanisms, we present the rate of upsets in all protected SRAM arrays of X-Gene 2 at different voltage levels (nominal and beyond). Figure 5 shows the upsets per minute, i.e.,

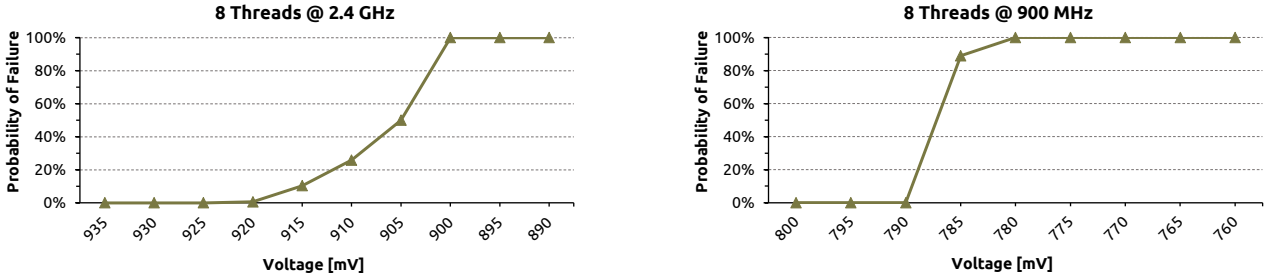

Figure 4: Probability of Failure (pfail) in all voltage levels from nominal level down to the levels of complete failure for different frequencies used in this study.

	Frequency	PMD Voltage	SoC Voltage
Nominal	2.4 GHz	980 mV	950 mV
Safe		930 mV	925 mV
Vmin		920 mV	920 mV
Vmin	900 MHz	790 mV	950 mV

Table 3: Voltage levels used in our experiments.

the average rate of upsets (SBUs and Multiple-Bit Upsets - MBUs), consolidated for all SRAM arrays of X-Gene 2 and all six benchmarks running at 2.4 GHz clock frequency (we will present in the next subsections the upsets rate for 900 MHz clock frequency). By reducing the voltage settings beyond nominal conditions (i.e., the three voltage settings in the x-axis of this graph), we observe that most of the benchmarks show a clear trend of increased upset rates. In Figure 5, we show only the SEUs that occur locally in all protected SRAM arrays, as reported by its hardware detection and correction mechanisms, i.e., not the potential abnormal behavior that the neutron beam experiment can introduce to the running application or the entire system. We present and explain such errors in the following subsections.

We observe that for some benchmarks, the upsets rate of the SRAM arrays may increase by up to 40.4% when reducing the voltage settings of the microprocessor chip. For example, the MG benchmark at the nominal voltage (i.e., 980 mV) experiences 0.94 upsets/minute. At a safe voltage setting beyond the nominal voltage, it experiences 1.02 upsets/minute (i.e., 8.5% increase in the upsets rate), and at the V_{min} (i.e., the lowest safe voltage setting) 1.32 upsets/minute (i.e., 40.4% increase in the upsets rate compared to the nominal voltage). However, there are some cases where the upsets rate decreases in lower voltages, such as in the CG and LU benchmarks. The reason for this upset rate reduction is that all experiments are iteratively executed until we reach 1) approximately 100 SDCs, application crashes, and system crashes in total, or 2) at least 1×10^{11}neutrons/cm^2, as discussed in Section 3.6. CG benchmark, for example, has an increased number of SDCs in lower voltage levels; thus, the upsets rate, reported in Figure 5, is shown lower for lower voltage levels. However, this phenomenon does not change the trend that lower voltage levels increase the SER of caches. The right-most (red-colored) bars of Figure 5 show the total rate of all benchmarks normalized by the execution time for each voltage level, which reveal that the lower the supply voltage, the higher the cache memory upsets rate.

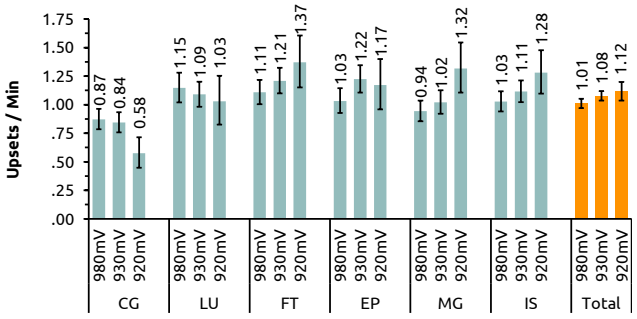

Figure 5: Cache memory upsets per minute for all benchmarks used in this study at 2.4 GHz.

Observation #1: The upsets rate of the SRAM arrays increases by 10.9% on average when reducing the voltage settings to the safe V_{min}.

4.3 SRAM Upsets per Cache Level

In this section, we explore the percentage of SRAM upsets for each cache level and voltage setting independently. Figure 6 shows the upset rate for each cache level at 2.4 GHz per voltage level and for each EDAC recovery attempt (i.e., corrected / uncorrected). We draw two essential observations.

(1) The largest the SRAM structure, the higher the upsets rate. Figure 6 shows that for the ECC corrected upsets, the L2 cache has a higher recovery rate than the L1 cache, and the L3 cache has a higher recovery rate than the L2 cache. On the other hand, uncorrected ECC upsets have only been observed in L3. The reason is primarily attributed to L3's larger memory size; large cache arrays with no memory interleaving schemes are more vulnerable to Multi-Bit Upsets (MBUs) [20]. Note that since the L3 cache is protected through an SECDED scheme (see Table 1), the reported uncorrected errors are at least double bit-flips, which can be detected by SECDED but not corrected.
(2) The lower the supply voltage level, the higher the upsets rate. This trend is illustrated in Figure 6, which shows the upsets rate for each cache level and all benchmarks consolidated.

It is essential to stress that both phenomena also hold for different clock frequency settings, such as the 900 MHz, as shown in Figure 7. However, Figure 7 shows clearly that the impact of voltage reduction on the higher cache levels (TLBs, L1, and L2 cache memories) is more severe in the lower frequency setting (i.e., 900 MHz). Specifically, we see that in the lower clock frequency setting (i.e., Figure 7), the upsets rate of the L1 and the L2 caches is increased by 2.7× and 50%, respectively, compared to the corresponding cache levels in 920 mV at 2.4 GHz (i.e., Figure 6). The reason is that at 790 mV, both the voltage and the clock frequency of the PMD domain (i.e., the CPU cores) are changed (i.e., to 790 mV and 900 MHz), but the voltage and the clock frequency of the SoC domain, which integrates the L3 cache, remain unchanged (see Table 3). Therefore, due to the lower voltage levels of the PMD domain, the upsets percentage is higher for L1 and L2 caches in 790 mV than in 920 mV, while for the L3 cache is lower than in 920 mV.

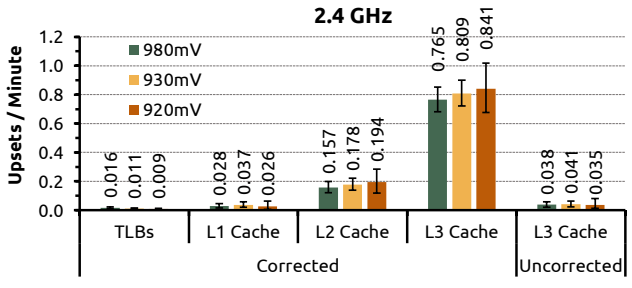

Figure 6: Cache memory upsets per minute for each cache level for 2.4 GHz clock frequency.

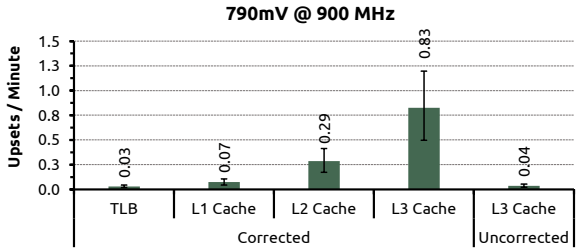

Figure 7: Cache memory upsets per minute for each cache level for 900 MHz clock frequency.

Observation #2: The upsets rate of the SRAM arrays increases as their size becomes larger, regardless of the voltage settings.

Bit-cell failures in SRAM arrays are classified into two categories [22]: 1) persistent bit failures, which strongly depend on a low supply voltage and usually occur due to the Random Dopant Fluctuations (RDF), and 2) non-persistent bit-upsets, where SRAM cells exhibit sporadic failing behavior and primarily occur due to radiation effects. Nevertheless, manufacturing-induced parameter fluctuations become more severe when supply voltage drops beyond the nominal voltage conditions, resulting in the failure of several circuits. These fluctuations limit voltage scaling to the minimum supply voltage (i.e., the safe V_{min}) necessary for the microprocessor chip to operate reliably [22]. For this reason, SRAM bit-cells become more prone to soft errors, especially to multiple-bit upsets during ultra-low voltage conditions. In response, prior works have proposed several solutions and strategies which can enable ultra-low voltage operation of cache memories, primarily due to multiple bit failures [3, 4, 22, 59, 60, 78, 80].

The X-Gene 2 microprocessor is equipped with a 1) parity protection scheme for its Translation Lookaside Buffers (TLBs) and its L1 data and instruction caches and 2) SECDED ECC scheme for its L2 and L3 caches. Our experiments demonstrate that these protection schemes, commonly used in most modern microprocessor chips, are sufficient for preserving the reliability of SRAM structures on a low-voltage operation due to soft errors. This aspect is illustrated in Figure 6 and Figure 7, which show that in all SRAM structures, neither the corrected nor the uncorrected upsets rates exhibit extreme fluctuations at lower voltage levels. This remark aligns with previous studies, which show that only a few cache lines experience single and multiple-bit failures at low voltages [3]. Note that, in the lower frequency setting (i.e., 900 MHz), the upset rate of the parity-protected arrays (i.e., TLB and L1 cache) increases significantly with respect to 920 mV at 2.4 GHz, as discussed earlier. Still, these rates are significantly lower than those of the larger arrays (e.g., L2/L3 cache memories) that are protected through SECDED.

Observation #3: The SRAM upset rates do not exhibit extreme fluctuations at lower voltage levels for a specific frequency.

Design implication #1: Commonly used error mitigation schemes (i.e., parity and SECDED) in microprocessor caches can adequately recover soft errors during increased cache upset rates caused by supply voltage scaling.

4.4 Application and OS Abnormal Behavior

In the previous subsections, we presented the upset rates of the microprocessor's SRAM cache arrays without considering the end-to-end effects on the software layer and the full system operation (i.e., the running application and the entire system, including the operating system). In this section, we explore the radiation effects on the software layer when reducing the voltage levels beyond the nominal conditions. Note that since the voltage levels we examine are all *safe voltage settings for all workloads* (see details in Section 4.1), any cache upset, SDC or application/system crash should occur only due to the neutron radiation effects, which is the focus of our analysis. Figure 8 shows the percentage of abnormal behaviors at the software layer for each voltage setting at the 2.4 GHz operating clock frequency. We observe the following:

(1) In a fixed clock frequency (2.4 GHz in our case), the percentage of system crashes becomes lower at lower voltages. The system crashes initially account for 51.6% of the total abnormal behaviors at the software layer in nominal voltage conditions. However, at 930 mV, the percentage of system crashes is reduced to 37.1% (i.e., 28.1% reduction of system crashes compared to 980 mV). When the voltage is reduced even more to 920 mV, the percentage of system crashes drops to only 5.7% (i.e., 89% reduction of system crashes compared to nominal voltage conditions).
(2) Application crashes (i.e., AppCrash label) follow the same pattern as the system crashes. Specifically, the application crashes are reduced in lower voltages at each target clock frequency. The AppCrash account for 17.9% of the total abnormal behaviors at the software layer in nominal voltage conditions. However, by reducing the voltage to 930 mV, the percentage of AppCrash drops to 7.2% (i.e., 59.8% reduction of AppCrash). When the voltage is reduced even more to 920 mV, the percentage of application crashes accounts for only

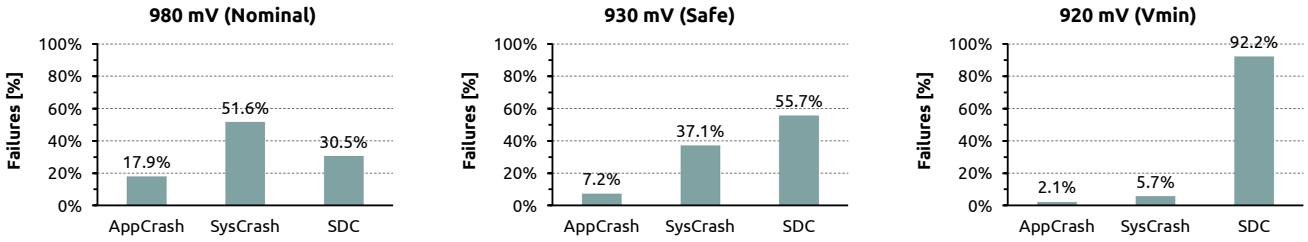

Figure 8: Percentage of abnormal behaviors at the software layer for each voltage setting, independently, for 2.4 GHz clock frequency.

2.1% of the total errors (i.e., 88.3% reduction of application crashes compared to nominal voltage conditions).

(3) In contrast, the percentage of SDCs increases at lower voltages. Specifically, SDCs account for 30.5% of the total abnormal behaviors at the software layer in nominal voltage conditions. The SDCs account for 55.7% (i.e., it is increased by 82.6% compared to 980 mV) when reducing the voltage settings to 930 mV. However, when the voltage is reduced even more to 920 mV, the percentage of SDCs significantly increases to 92.2%. This is a significant observation, shown for the first time, and combines the impact of a low-voltage operation on the soft error susceptibility and the generation of SDCs, which are the most challenging radiation effects.

Observation #4: The probability of an SDC occurrence compared to AppCrash or SysCrash, when the microprocessor operates at low-voltage conditions, is 3× larger than in nominal voltage conditions.

5 POWER CONSUMPTION AND SUSCEPTIBILITY TRADE-OFF

5.1 Power Consumption versus Upsets Rate

This section explores the tradeoff between power consumption and cache upset rate. Figure 9 shows the tradeoff between power consumption and the soft error susceptibility of caches for all voltage levels considered in this study and for all clock frequencies (i.e., 2.4 GHz and 900 MHz). Each bar in the graph presents the total power consumption for both PMD and SoC domains and each voltage setting. The orange line presents the rate of cache upsets per minute for all benchmarks and voltage settings.

This graph shows the correlation between the advantage of low power consumption when reducing the supply voltage beyond nominal conditions and the disadvantage of the high rate of upsets in the microprocessor's caches. Assume, for example, the highest available clock frequency of 2.4 GHz, as shown in Figure 9. The average power consumption of all six benchmarks used in this study is 20.40W at the nominal voltage conditions (i.e., 980 mV). The measured upset rate, which corresponds to the nominal voltage level, is 1.01 upsets per minute. Keeping the clock frequency at its highest levels (i.e., 2.4 GHz) and reducing the supply voltage to 930 mV, we see that the average power consumption for both PMD and SoC domains is 18.63W. This means that for 5.1% voltage reduction (i.e., from 980 mV to 930 mV), the power consumption is reduced by 8.7%. However, the upsets per minute are increased from 1.01 to 1.08, which is a 6.9% increase in soft error susceptibility. The same behavior is also drawn for all voltage levels beyond nominal conditions.

Observation #5: The power consumption can be significantly reduced by exploiting the pessimistic voltage margins of modern microprocessors. However, the susceptibility to soft errors is also increased.

Another major observation is that the reduced clock frequency does not affect significantly the susceptibility of the microprocessor to soft errors. In the rightmost bar of Figure 9, we can see the microprocessor's power consumption when it operates at 790 mV and reduced clock frequency (i.e., 900 MHz). In this bar, it is clear that

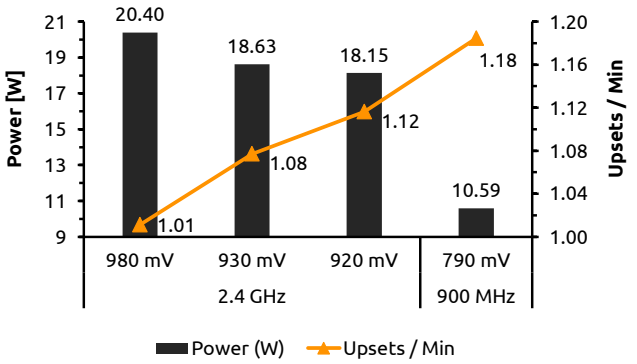

Figure 9: The trade-off between power consumption and the soft error susceptibility for all voltage levels considered in this study and for all clock frequencies.

the power consumption is significantly reduced due to the combination of voltage and clock frequency reduction, compared to the leftmost configurations of this graph, in which the clock frequency remains the same. However, the upsets per minute (i.e., the orange line in Figure 9) is virtually linearly increased compared to all 3 different voltage levels at 2.4 GHz clock frequency. This means that the susceptibility to soft errors in caches when the microprocessor operates beyond the nominal voltage conditions is strongly affected by the voltage levels and is less affected by the frequency changes. Thus, we attribute the increase of the susceptibility more to the voltage reduction from 920 mV to 790 mV, and less to the frequency reduction.

Observation #6: Clock frequency reduction does not affect significantly the susceptibility of the microprocessor to soft errors.

5.2 Power Savings versus Susceptibility

In Section 5.1, we explored the correlation between power consumption and the cache's SER. In this section, we examine the trends between the reduced supply voltage levels' power savings and the caches' susceptibility to soft errors.

Figure 10 shows the percentage of power savings for each voltage level related to the nominal voltage and the maximum clock frequency (i.e., 980 mV and 2.4 GHz) and the increase of the susceptibility of caches to soft errors with respect to the nominal voltage and the maximum clock frequency. Both power savings and susceptibility percentages are average values from all six benchmarks considered in this study. We observe that for both voltage levels at 2.4 GHz (i.e., 930 mV and 920 mV), the power savings curve is increased slower than the susceptibility curve. Specifically, from 930 mV to 920 mV (which accounts for only 1% voltage reduction difference), the power savings are increased by 26.4% (i.e., from 8.7% to 11.0%), while the susceptibility is increased by 58% (i.e., from 6.9% to 10.9%).

However, when we reduce both the supply voltage and the clock frequency (i.e., 790 mV and 900 MHz clock frequency), the power savings are, apparently, extremely high; still, the susceptibility to soft errors follows the same pattern. Specifically, the susceptibility is increased by 54.1% (i.e., from 10.9% to 16.8%). As discussed in Section 5.1, the reason is that the susceptibility of the microprocessor

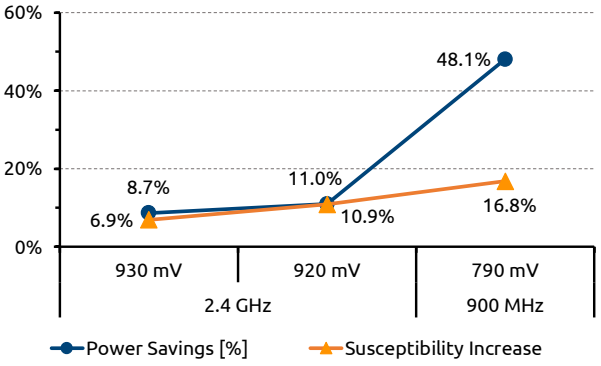

Figure 10: The correlation between power savings and the soft error susceptibility for all voltage levels considered in this study and for all clock frequencies.

to soft errors is strongly affected by the voltage reduction but not by the clock frequency changes. Note that when reducing the voltage levels, the performance is not compromised, but when reducing the clock frequency, the performance reduces as well. This is why in 790 mV at 900 MHz, we have 48.1% power savings; because there is a reduction in both voltage and frequency and thus, in performance.

Observation #7: The power savings increase at a slower pace than the susceptibility at 2.4 GHz. However, at 900 MHz the increase in power savings is higher than the increase in susceptibility.

6 FAILURES IN TIME (FIT) RATES

Apart from the upsets rate for SRAM structures, it is also essential to demonstrate the impact of each individual voltage level on the soft error susceptibility of the entire microprocessor chip. In this section, we present the microprocessor's FIT rates for a NYC flux for each individual voltage level.

6.1 Total Microprocessor FIT Rate

This section presents the total FIT rate of the entire microprocessor at the system level, running the benchmarks on top of Linux for our real measurements. Figure 11 shows the total FIT rates of application and system level errors per voltage level (i.e., 980 mV, 930 mV, and 920 mV) for the target 2.4 GHz clock frequency. The graph illustrates only detected SDC and application/system crash errors, since our setup does not incorporate any error mitigation mechanism at the application and system level to correct errors. The bars showing the total FIT in Figure 11 are the sum of the individual FIT rates of AppCrash, SysCrash, and SDC, independently for each voltage level. Note that the up arrows in Figures 11 and 12 represent numbers that are outside the y-axis values range and in these cases, the error bar is not shown.

We observe that the lower the voltage level, the lower the FIT rates of both application crashes (i.e., AppCrash) and system crashes (i.e., SysCrash). On the other hand, the SDC FIT rate is higher for lower voltages, which is in line with the observations made in Section 4.4. Specifically, the SDC FIT rate increases from 2.54 to 41.43 at the V_{min} (i.e., 16.3× increase). However, the total FIT rates for the entire microprocessor chip are getting higher for lower voltage levels (i.e., 6.6×).

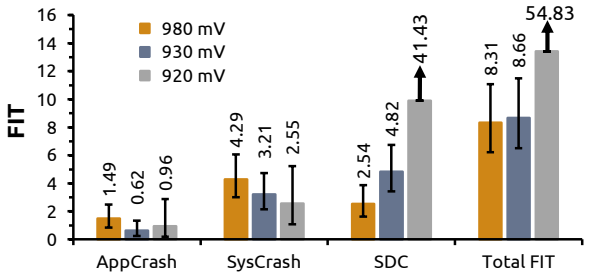

Figure 11: Total FIT rate of the entire CPU chip for all voltage and frequency levels, considered in this study.

The number of AppCrash and SysCrash events observed in the 920 mV test session is relatively low (i.e., 3 AppCrashes and 8 SysCrashes), which means that the level of uncertainty of the statistical results for these failure categories is not negligible. As mentioned in Section 3.4, statistical significance results for each test session in our experiment are achieved when 100 events (SDC, AppCrash, SysCrash) are observed, or the fluence reaches 10^{11} neutrons/cm^2, whichever comes first. In the case of 920 mV, due to the increased number of SDCs (130 SDCs were observed in about 7.5 hours), the test session finished earlier before we collected sufficient AppCrashes and SysCrashes to guarantee statistically significant results. Thus, we can claim that the level of uncertainty is very low for the total FIT rate but not for the AppCrash and SysCrash FIT rates separately. The non-negligible statistical error for the AppCrash and SysCrash events could explain why the individual FIT rates for these categories do not follow the total FIT rate increase at the lower voltage level. Needless to say, we could not have kept the test session running until we had collected a sufficient number of events from all failure categories separately due to limited beam time availability.

Observation #8: The lower the safe voltage is, the higher the FIT rate of the entire microprocessor chip. Specifically, the FIT rate of the entire CPU chip can be as high as 54.83 at the lowest safe voltage level (i.e., 920 mV). Moreover, the SDC FIT rate, at low voltage levels, is extremely higher than the AppCrash and SysCrash FIT rates.

Several efforts have been recently made to enable hardware operation at sub-nominal voltage levels without affecting the clock frequency, taking advantage of reduced voltage margins of microprocessor chips [43, 49]. These studies propose dynamic approaches for setting up, managing, and controlling the reduced voltage levels

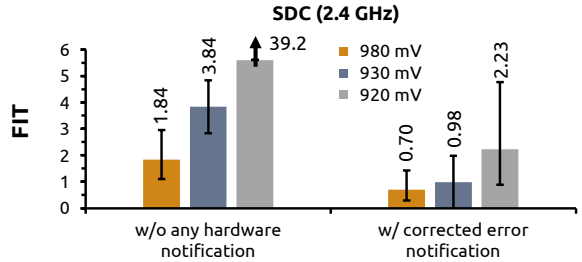

Figure 12: SDC FIT rates for the entire microprocessor chip, with and without hardware error notification for all voltage levels and for the highest clock frequency.

in a context-aware manner, without additional hardware support or modified application binaries. In this context, cloud providers or large datacenter companies can use these methodologies in their already existing infrastructures to reduce their energy footprint without affecting the current performance levels or the high availability demands. However, according to these studies, the proposed operating voltage levels are usually the lowest safe V_{min} (e.g., 920mV in our case). Our study shows that when the microprocessor operates at the lowest safe V_{min} (i.e., 920 mV), the FIT rates are extremely higher than operating it at 10 mV above V_{min} (i.e., at 930 mV). As shown in Figure 11, the total FIT rate at 930mV is slightly increased from the nominal voltage. However, the total FIT rate increases rapidly at the lowest safe V_{min} (i.e., 920 mV).

Design implication #2: Our experimental results suggest that cloud and datacenter providers should operate the microprocessors slightly above the lowest safe V_{min} to reduce the probability of radiation-induced errors in their systems. An optimal supply voltage to save power but not increase SDCs in the case of our target microprocessor is 930 mV rather than V_{min}.

Fault injection is widely used to evaluate hardware vulnerability by injecting faults into certain microprocessor structures and observing their effects on the program's output. For example, simulation-based microarchitecture-level fault injection can be used to measure the Architectural Vulnerability Factor (i.e., AVF) metric [42], which expresses the probability of hardware faults corrupting the program output. There is a direct correlation between the program failure rate, i.e., the FIT or the Mean Time to Failure (MTTF), and the microprocessor's AVF and raw FIT rate per bit [18].

To attribute the FIT of a component to its size, the raw FIT per bit is necessary. Since the size of each hardware structure and the raw FIT per bit (of the certain technology node) are known a priori and the microarchitecture-level simulation experiments provide the AVF, it is feasible to estimate the total FIT of a hardware structure for each different voltage level. As shown in Figure 10, the percentage of susceptibility increase of caches (i.e., SRAM arrays) due to soft errors with respect to the nominal voltage and the maximum clock frequency can be used along with the raw FIT per bit and the AVF, to estimate the total FIT of a certain SRAM structure for different voltage levels.

Design implication #3: The reported cache upset rates can be used in microarchitecture-level fault injection studies to estimate the application FIT rates of different microprocessor designs at scaled supply voltage levels. This can increase design space exploration parameters and enable microarchitectures that balance between power savings and dependability at reduced supply voltage levels.

6.2 Correlation of SDC FIT Rate to the Hardware Errors Notification

SDCs are the most severe class of fault effects in modern CPUs. Errors brought in by transitory defects may silently corrupt the results of the computation performed by the impacted devices. An SDC produced as the result of a single CPU can cascade into a massive problem in modern large-scale infrastructures [25, 34]. Therefore, it is crucial to study the SDC FIT rates for different voltage levels separately. By its definition, an SDC occurs without any indication of the output mismatch in system events or error logs (i.e., faults in unprotected structures or control logic). However, at the system-level setup, it is likely for an output mismatch to happen accompanied by a corrected error notification from a protected SRAM cache structure. These events rarely occurred during our accelerated radiation testing experiments in the following unusual cases: 1) there was a triple-bit upset, and the SECDED scheme recognized it as a single-bit upset, so it was mistakenly reported as a corrected error (and ended up generating a corrupted output), or 2) there were two concurrent events: an error in an unprotected unit (i.e., there is no report for the error occurrence) along with a corrected error in a protected unit (i.e., a single-bit upset in the L2 cache), and despite the correction action the output was corrupted. In both cases, the program output is affected, and therefore we count these as SDCs in our system-level setup.

Figure 12 shows the SDC FIT rates for the entire microprocessor chip, with and without a corrected error notification for all voltage levels and for the highest clock frequency. The hardware error notification is any notification from the hardware for a corrected event (e.g., an ECC-corrected upset in the L2 cache). Therefore, under the case "without error notification" we include the program executions that finish with an output mismatch, but with no indication from hardware or any other resource that there has been any error event (i.e., corrected only).

We observe that in both cases (i.e., w/ and w/o corrected notification), the lower the voltage level is, the higher the SDC FIT rate. However, the SDC FIT rates for all voltage levels are extremely higher in case of no hardware notification than in the case of some indication for a corrected event (rare events). This observation, in conjunction with the previous section that demonstrates the extremely high SDC FIT rate, strengthens the argument that even using sophisticated detection and correction schemes (e.g., SECDED ECC), there is still a small probability for a soft error to result in an SDC. Previous work shows that even using common ECC methods, SDCs are unavoidable, especially in large-scale datacenter infrastructures [45]. The results of our study confirm this phenomenon but also demonstrate that it is exacerbated when the supply voltage decreases, i.e., the probability of a soft error in SRAM structures resulting in an SDC is significantly increased in lower voltage levels.

Observation #9: The SDC FIT rates for all voltage levels are extremely higher in case of no hardware notification than in the case of some indication for a corrected event.

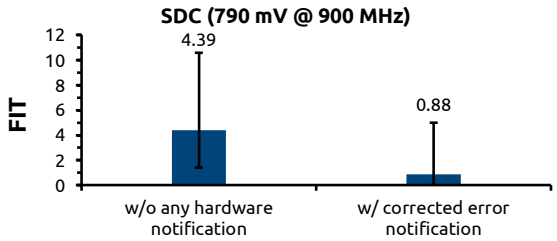

Figure 13: SDC FIT rates for the entire microprocessor chip, with and without hardware error notification, for 900 MHz clock frequency.

Figure 13 shows the SDC FIT rates for the entire microprocessor chip with and without hardware error notification for 900 MHz clock frequency. As we can see, the same behavior also exists in lower clock frequencies, such as 900 MHz. It is clear by Figure 13 that the SDC FIT rate is extremely increased in the case without any hardware notification.

As shown in Figures 12 and 13, there is a high number of SDC FIT rates due to soft errors. More importantly, as shown in these figures, the SDC FIT rate is constantly increased at lower voltage levels. Since SDCs occur without any hardware notification (e.g., from parity or ECC), it is very likely that there are several paths in the core logic (that are not protected by any mechanism), which become more susceptible to soft errors due to low voltage operation. To this end, chip designers can take advantage of this observation to enhance their next-generation designs, for example, by pinpointing the weakest paths or by adding new protection mechanisms.

Design implication #4: SDCs are probably not caused by upsets in SRAM structures when the microprocessor operates at a reduced supply voltage. Thus, computer architects may consider our findings to locate soft errors in those circuit paths causing SDCs due to radiation effects when the microprocessor operates at low supply voltage levels.

7 RELATED WORK

Voltage Margins Characterization & Low-Voltage Operation. Several characterization studies have been presented for off-nominal voltage conditions operation of commercial microprocessor chips with up to 8 cores (e.g., [5, 6, 40, 41, 57, 71, 74, 87]). Bacha *et al.* [5, 6] focused on monitoring the hardware-reported errors in the caches of an Intel Itanium processor running benchmarks in off-nominal voltage conditions. Authors in [36, 37, 51, 58, 74] measured single-core voltage margins in several commercial microprocessor chips to study not only the pessimistic voltage guardbands of the chips but also the core-to-core and chip-to-chip variations for single-core executions. Sasaki *et al.* [64] studied the prevalence of power capping when multiple processes in a multicore microprocessor compete for power, while the power management system attempts to mitigate the contention by slowing down the processor.

Accelerated Beam Experiments. Particle accelerators have been used for many years to measure the reliability of devices and applications [8, 85]. Computing devices' reliability has a strong tradition, motivated mainly by their use in safety-critical applications [23, 47, 67]. Arm Cortex-A9 processors have been exposed to accelerated particles beam and have been subject to fault injection experiments. In [24, 30, 44, 63], the authors presented beam experimental data of Arm Cortex-A9 microprocessors, proposed hardening solutions, and discussed the impact of the operating system in the application and device reliability. Authors in [61, 62] presented results on architecture-level fault injection of the processor core, while [19] includes a microarchitecture-level fault injection on a Cortex-A. Some preliminary studies have proposed a comparison or combination of different reliability evaluation techniques [30, 31, 77]. In [18], a first attempt was made to compare the reliability evaluation of a Cortex-A9 using beam experiments and microarchitectural fault injection. According to the authors, for SDCs, the comparison can be very close, but the difference is significant for caches.

Soft Errors & Supply Voltage Scaling. The authors in [66, 67] evaluated through alpha and neutron accelerated radiation testing the cross-section of 500 nm to 180 nm HP Alpha microprocessors at nominal supply voltage and, in turn, used simulation to extrapolate the SER of their SRAM structures for reduced supply voltages and various clock frequencies. Compared to [66, 67], our study 1) analyzes both the raw SER of SRAM structures and their effects on the application and system layer on real hardware under supply voltage scaling, and 2) targets a modern 28 nm Arm processor.

Chandra and Aitken in [16, 17] performed simulation-based experiments and showed that a 60% reduction in the supply voltage of 65 nm and 45 nm technologies led to a 78% and 81% reduction of Q_{crit}, respectively. Q_{crit} is defined as the minimum charge needed to flip the bit stored in a memory cell. Tonfat *et al.* in [73] performed neutron radiation experiments in a 45 nm Field Programmable Gate Array (FPGA) at different supply voltages to measure its SER. The experimental results showed that an 8% reduction in the supply voltage might result in a 30% higher raw SER. The authors also stated that the application-level failure rate might have a variation of 55% for a voltage variation of 19%. In addition, the authors highlighted that this reliability degradation is expected to increase with device miniaturization and technology scaling. Brendler *et al.* in [14] explored the soft error impact along with voltage scaling through simulations. They showed that the SER is 61% higher at near-threshold conditions, demonstrating the need of strong mitigation strategies to create more reliable circuits.

Wu and Marculescu in [81] proposed a power-aware soft error hardening framework via selective voltage scaling using dual supply voltages for combinatorial logic. They also introduced a heuristic and two refinement techniques for SER reduction. Kastensmidt *et al.* in [38] observed how aging and voltage scaling affect the SER in SRAM-based FPGAs. That work is based on the results of Monte-Carlo electrical simulations and neutron beam experiments. They showed that the error rate could increase by more than twice when taking aging and voltage scaling into account.

All these studies are implemented either on simulations or small FPGA designs. To our knowledge, there is no previous study on microprocessors that provides results through accelerated radiation testing on real hardware with reduced supply voltage.

8 CONCLUSION

In this paper, we presented, for the first time in the literature, a two-knob evaluation of modern server CPUs operation: the impact of a low-voltage operation on the soft error susceptibility of modern multicore microprocessors in full-system level execution. Our analysis on real hardware revealed several important insights and observations, including: 1) the system and application crashes rates decrease while the SDCs rates increase in lower voltage levels, 2) by lowering the supply voltage, the power consumption can be significantly improved, but the susceptibility to soft errors is significantly increased, 3) the clock frequency does not affect the susceptibility of the CPU to soft errors, and 4) the SDCs FIT rate is significantly increased in lower voltage levels. The data collected

from accelerated beam experiments can be used to predict the failure rate of the application code running on the microprocessor under conditions as close as possible to real-world conditions after full system integration. Our analysis can guide the design of future CPUs to achieve good trade-off among performance, energy, and reliability and help architects make better design decisions when incorporating dynamic voltage scaling and fault-tolerance techniques.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their insightful comments and feedback. This work is supported by research gifts from Meta and AMD and by the European Union's Horizon Europe research and innovation programme under grant agreements No 101097224 (REBECCA) and No 101070238 (NEUROPULS). Views and opinions expressed are those of the authors only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them. Beam time at TRIUMF/NIF was awarded by European Union's Horizon 2020 research and innovation programme under grant agreement No 101008126 RADNEXT (1 day) and TRIUMF itself (2 days). TRIUMF receives funding via a contribution agreement with the National Research Council of Canada. The computing equipment used in the experiments was funded by the H2020 Framework Program of the European Union through the UniServer project (grant agreement 688540). The outcome of this work is part of the dissemination for the Hellenic Foundation Research and Innovation (HFRI) project VEMER.

REFERENCES

[1] 2011. TSMC's 28nm Technology. https://www.tsmc.com/english/dedicatedFoundry/technology/logic/l_28nm Accessed: 2023-04-25.

[2] 2017. The Linux Kernel Documentation, Error Detection And Correction (EDAC) Devices. https://www.kernel.org/doc/html/v4.11/driver-api/edac.html. Accessed: 2023-04-25.

[3] Alaa R. Alameldeen, Ilya Wagner, Zeshan Chishti, Wei Wu, Chris Wilkerson, and Shih-Lien Lu. 2011. Energy-Efficient Cache Design Using Variable-Strength Error-Correcting Codes. *SIGARCH Comput. Archit. News* 39, 3 (jun 2011), 461–472. https://doi.org/10.1145/2024723.2000118

[4] Amin Ansari, Shuguang Feng, Shantanu Gupta, and Scott Mahlke. 2011. Archipelago: A polymorphic cache design for enabling robust near-threshold operation. In *2011 IEEE 17th International Symposium on High Performance Computer Architecture*. 539–550. https://doi.org/10.1109/HPCA.2011.5749758

[5] Anys Bacha and Radu Teodorescu. 2013. Dynamic Reduction of Voltage Margins by Leveraging On-Chip ECC in Itanium II Processors. In *Proceedings of the 40th Annual International Symposium on Computer Architecture* (Tel-Aviv, Israel) *(ISCA '13)*. Association for Computing Machinery, New York, NY, USA, 297–307. https://doi.org/10.1145/2485922.2485948

[6] Anys Bacha and Radu Teodorescu. 2014. Using ECC Feedback to Guide Voltage Speculation in Low-Voltage Processors. In *2014 47th Annual IEEE/ACM International Symposium on Microarchitecture*. 306–318. https://doi.org/10.1109/MICRO.2014.54

[7] D. H. Bailey, E. Barszcz, J. T. Barton, D. S. Browning, R. L. Carter, L. Dagum, R. A. Fatoohi, P. O. Frederickson, T. A. Lasinski, R. S. Schreiber, H. D. Simon, V. Venkatakrishnan, and S. K. Weeratunga. 1991. The NAS parallel benchmarks summary and preliminary results. In *Supercomputing '91:Proceedings of the 1991 ACM/IEEE Conference on Supercomputing*. 158–165. https://doi.org/10.1145/125826.125925

[8] R.C. Baumann. 2005. Radiation-induced soft errors in advanced semiconductor technologies. *Device and Materials Reliability, IEEE Transactions on* 5, 3 (Sept 2005), 305–316. https://doi.org/10.1109/TDMR.2005.853449

[9] Majed Valad Beigi, Sudhanva Gurumurthi, and Vilas Sridharan. 2022. Reliability, Availability, and Serviceability Challenges for Heterogeneous System Design. In *2022 IEEE International Reliability Physics Symposium (IRPS)*. 2C.4–1–2C.4–8. https://doi.org/10.1109/IRPS48227.2022.9764554

[10] E.W. Blackmore, P.E. Dodd, and M.R. Shaneyfelt. 2003. Improved capabilities for proton and neutron irradiations at TRIUMF. In *IEEE Radiation Effects Data Workshop*. 149–155. https://doi.org/10.1109/REDW.2003.1281368

[11] Ewart Blackmore, Michael Trinczek, Kai Jiang, Manoj Sachdev, and Derek Wright. 2019. SRAM Dosimeter for Characterizing the TRIUMF Proton and Neutron Beams. *IEEE Transactions on Nuclear Science* 66, 1 (2019), 276–281. https://doi.org/10.1109/TNS.2018.2884148

[12] Pablo Bodmann, George Papadimitriou, Dimitris Gizopoulos, and Paolo Rech. 2021. The Impact of SoC Integration and OS Deployment on the Reliability of Arm Processors. In *2021 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. 223–225. https://doi.org/10.1109/ISPASS51385.2021.00040

[13] Pablo R. Bodmann, George Papadimitriou, Rubens L. Rech Junior, Dimitris Gizopoulos, and Paolo Rech. 2022. Soft Error Effects on Arm Microprocessors: Early Estimations versus Chip Measurements. *IEEE Trans. Comput.* 71, 10 (2022), 2358–2369. https://doi.org/10.1109/TC.2021.3128501

[14] Leonardo H. Brendler, Alexandra L. Zimpeck, Fernanda L. Kastensmidt, Cristina Meinhardt, and Ricardo Reis. 2021. Voltage Scaling Influence on the Soft Error Susceptibility of a FinFET-based Circuit. In *2021 IEEE 12th Latin America Symposium on Circuits and System (LASCAS)*. 1–4. https://doi.org/10.1109/LASCAS51355.2021.9459127

[15] Ramon Canal, Carles Hernandez, Rafa Tornero, Alessandro Cilardo, Giuseppe Massari, Federico Reghenzani, William Fornaciari, Marina Zapater, David Atienza, Ariel Oleksiak, Wojciech Piundefinedtek, and Jaume Abella. 2020. Predictive Reliability and Fault Management in Exascale Systems: State of the Art and Perspectives. *ACM Comput. Surv.* 53, 5, Article 95 (sep 2020), 32 pages. https://doi.org/10.1145/3403956

[16] Vikas Chandra and Robert Aitken. 2008. Impact of Technology and Voltage Scaling on the Soft Error Susceptibility in Nanoscale CMOS. In *2008 IEEE International Symposium on Defect and Fault Tolerance of VLSI Systems*. 114–122. https://doi.org/10.1109/DFT.2008.50

[17] Vikas Chandra and Robert Aitken. 2008. Impact of Technology and Voltage Scaling on the Soft Error Susceptibility in Nanoscale CMOS. In *IEEE International Symposium on Defect and Fault Tolerance of VLSI Systems (DFT)*. 114–122. https://doi.org/10.1109/DFT.2008.50

[18] A. Chatzidimitriou, P. Bodmann, G. Papadimitriou, D. Gizopoulos, and P. Rech. 2019. Demystifying Soft Error Assessment Strategies on ARM CPUs: Microarchitectural Fault Injection vs. Neutron Beam Experiments. In *2019 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 26–38. https://doi.org/10.1109/DSN.2019.00018

[19] Athanasios Chatzidimitriou, Manolis Kaliorakis, Sotiris Tselonis, and Dimitris Gizopoulos. 2017. Performance-aware reliability assessment of heterogeneous chips. In *2017 IEEE 35th VLSI Test Symposium (VTS)*. IEEE. https://doi.org/10.1109/vts.2017.7928940

[20] Athanasios Chatzidimitriou, George Papadimitriou, Christos Gavanas, George Katsoridas, and Dimitris Gizopoulos. 2019. Multi-Bit Upsets Vulnerability Analysis of Modern Microprocessors. In *2019 IEEE International Symposium on Workload Characterization (IISWC)*. 119–130. https://doi.org/10.1109/IISWC47752.2019.9042036

[21] Athanasios Chatzidimitriou, George Papadimitriou, Dimitris Gizopoulos, Shrikanth Ganapathy, and John Kalamatianos. 2019. Assessing the Effects of Low Voltage in Branch Prediction Units. In *2019 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. 127–136. https://doi.org/10.1109/ISPASS.2019.00020

[22] Zeshan Chishti, Alaa R. Alameldeen, Chris Wilkerson, Wei Wu, and Shih-Lien Lu. 2009. Improving cache lifetime reliability at ultra-low voltages. In *2009 42nd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 89–99. https://doi.org/10.1145/1669112.1669126

[23] C. Constantinescu. 2002. Impact of deep submicron technology on dependability of VLSI circuits. In *International Conference on Dependable Systems and Networks (DSN)*. 205–209. https://doi.org/10.1109/DSN.2002.1028901

[24] Ádria Barros de Oliveira, Gennaro Severino Rodrigues, and Fernanda Lima Kastensmidt. 2017. Analyzing Lockstep Dual-core ARM cortex-A9 Soft Error Mitigation in freeRTOS Applications. In *Proceedings of the 30th Symposium on Integrated Circuits and Systems Design: Chip on the Sands (SBCCI '17)*. ACM, New York, NY, USA, 84–89. http://doi.acm.org/10.1145/3109984.3110008

[25] Harish Dattatraya Dixit, Sneha Pendharkar, Matt Beadon, Chris Mason, Tejasvi Chakravarthy, Bharath Muthiah, and Sriram Sankar. 2021. Silent Data Corruptions at Scale. https://arxiv.org/abs/2102.11245

[26] Jack Dongarra, Thomas Herault, and Yves Robert. 2015. *Fault Tolerance Techniques for High-Performance Computing*. Springer International Publishing, Cham, 3–85. https://doi.org/https://doi.org/10.1007/978-3-319-20943-2_1

[27] Henry Duwe, Xun Jian, Daniel Petrisko, and Rakesh Kumar. 2016. Rescuing Uncorrectable Fault Patterns in On-Chip Memories through Error Pattern Transformation. In *2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA)*. 634–644. https://doi.org/10.1109/ISCA.2016.61

[28] European Space Components Coordination (ESCC). 2014. *Single event effects test method and guidelines, ESCC Basic Specification No. 25100.*

[29] Sagi Fisher, Adam Teman, Dmitry Vaysman, Alexander Gertsman, Orly Yadid-Pecht, and Alexander Fish. 2008. Digital subthreshold logic design - motivation

and challenges. In *2008 IEEE 25th Convention of Electrical and Electronics Engineers in Israel*. 702–706. https://doi.org/10.1109/EEEI.2008.4736624

[30] Vinícius Fratin, Daniel Oliveira, Caio Lunardi, Fernando Santos, Gennaro Rodrigues, and Paolo Rech. 2018. Code-Dependent and Architecture-Dependent Reliability Behaviors. In *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. 13–26. https://doi.org/10.1109/DSN.2018.00015

[31] Nisha George, Carl R. Elks, Barry W. Johnson, and John Lach. 2010. Transient fault models and AVF estimation revisited. In *2010 IEEE/IFIP International Conference on Dependable Systems & Networks (DSN)*. IEEE. https://doi.org/10.1109/dsn.2010.5544276

[32] Dimitris Gizopoulos, George Papadimitriou, Athanasios Chatzidimitriou, Vijay Janapa Reddi, Behzad Salami, Osman S. Unsal, Adrian Cristal Kestelman, and Jingwen Leng. 2019. Modern Hardware Margins: CPUs, GPUs, FPGAs Recent System-Level Studies. In *2019 IEEE 25th International Symposium on On-Line Testing and Robust System Design (IOLTS)*. 129–134. https://doi.org/10.1109/IOLTS.2019.8854386

[33] R. W. Hamming. 1950. Error detecting and error correcting codes. *The Bell System Technical Journal* 29, 2 (1950), 147–160. https://doi.org/10.1002/j.1538-7305.1950.tb00463.x

[34] Peter H. Hochschild, Paul Turner, Jeffrey C. Mogul, Rama Govindaraju, Parthasarathy Ranganathan, David E. Culler, and Amin Vahdat. 2021. Cores That Don't Count. In *Proceedings of the Workshop on Hot Topics in Operating Systems* (Ann Arbor, Michigan) *(HotOS '21)*. Association for Computing Machinery, New York, NY, USA, 9–16. https://doi.org/10.1145/3458336.3465297

[35] JEDEC Solid State Technology Association. 2021. *Measurement and Reporting of Alpha Particle and Terrestrial Cosmic Ray-Induced Soft Errors in Semiconductor Devices ('JESD89B - Revision of JESD89A, October 2006)*. https://www.jedec.org/system/files/docs/JESD89B.pdf

[36] Manolis Kaliorakis, Athanasios Chatzidimitriou, George Papadimitriou, and Dimitris Gizopoulos. 2018. Statistical Analysis of Multicore CPUs Operation in Scaled Voltage Conditions. *IEEE Computer Architecture Letters* 17, 2 (2018), 109–112. https://doi.org/10.1109/LCA.2018.2798604

[37] Georgios Karakonstantis, Konstantinos Tovletoglou, Lev Mukhanov, Hans Vandierendonck, Dimitrios S. Nikolopoulos, Peter Lawthers, Panos Koutsovasilis, Manolis Maroudas, Christos D. Antonopoulos, Christos Kalogirou, Nikos Bellas, Spyros Lalis, Srikumar Venugopal, Arnau Prat-Pérez, Alejandro Lampropulos, Marios Kleanthous, Andreas Diavastos, Zacharias Hadjilambrou, Panagiota Nikolaou, Yiannakis Sazeides, Pedro Trancoso, George Papadimitriou, Manolis Kaliorakis, Athanasios Chatzidimitriou, Dimitris Gizopoulos, and Shidhartha Das. 2018. An energy-efficient and error-resilient server ecosystem exceeding conservative scaling limits. In *2018 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. 1099–1104. https://doi.org/10.23919/DATE.2018.8342175

[38] F.L. Kastensmidt, J. Tonfat, T. Both, P. Rech, G. Wirth, R. Reis, F. Bruguier, P. Benoit, L. Torres, and C. Frost. 2014. Voltage scaling and aging effects on soft error rate in SRAM-based FPGAs. *Microelectronics Reliability* 54, 9 (2014), 2344–2348. https://doi.org/10.1016/j.microrel.2014.07.100 SI: ESREF 2014.

[39] Panos Koutsovasilis, Christos D. Antonopoulos, Nikolaos Bellas, Spyros Lalis, George Papadimitriou, Athanasios Chatzidimitriou, and Dimitris Gizopoulos. 2022. The Impact of CPU Voltage Margins on Power-Constrained Execution. *IEEE Transactions on Sustainable Computing* 7, 1 (2022), 221–234. https://doi.org/10.1109/TSUSC.2020.3045195

[40] Charles R. Lefurgy, Alan J. Drake, Michael S. Floyd, Malcolm S. Allen-Ware, Bishop Brock, Jose A. Tierno, and John B. Carter. 2011. Active management of timing guardband to save energy in POWER7. In *2011 44th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 1–11.

[41] Jingwen Leng, Alper Buyuktosunoglu, Ramon Bertran, Pradip Bose, and Vijay Janapa Reddi. 2015. Safe limits on voltage reduction efficiency in GPUs: A direct measurement approach. In *2015 48th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 294–307. https://doi.org/10.1145/2830772.2830811

[42] R. Leveugle, A. Calvez, P. Maistri, and P. Vanhauwaert. 2009. Statistical Fault Injection: Quantified Error and Confidence. In *Proceedings of the Conference on Design, Automation and Test in Europe* (Nice, France) *(DATE '09)*. European Design and Automation Association, Leuven, BEL, 502–506.

[43] Emmanouil Maroudas, Spyros Lalis, Nikolaos Bellas, and Christos D. Antonopoulos. 2021. Exploring the Potential of Context-Aware Dynamic CPU Undervolting. In *Proceedings of the 18th ACM International Conference on Computing Frontiers* (Virtual Event, Italy) *(CF '21)*. Association for Computing Machinery, New York, NY, USA, 73–82. https://doi.org/10.1145/3457388.3458658

[44] Antonio Martínez-Álvarez, Felipe Restrepo-Calle, Sergio Cuenca-Asensi, Leonardo M. Reyneri, Almudena Lindoso, and Luis Entrena. 2016. A Hardware-Software Approach for On-Line Soft Error Mitigation in Interrupt-Driven Applications. *IEEE Trans. Dependable Sec. Comput.* 13, 4 (2016), 502–508. https://doi.org/10.1109/TDSC.2014.2382593

[45] Justin Meza, Qiang Wu, Sanjeev Kumar, and Onur Mutlu. 2015. Revisiting Memory Errors in Large-Scale Production Data Centers: Analysis and Modeling of New Trends from the Field. In *2015 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks*. 415–426. https://doi.org/10.1109/DSN.2015.57

[46] Shubhendu S. Mukherjee, Christopher Weaver, Joel Emer, Steven K. Reinhardt, and Todd Austin. 2003. A Systematic Methodology to Compute the Architectural Vulnerability Factors for a High-Performance Microprocessor. In *Proceedings of the 36th Annual IEEE/ACM International Symposium on Microarchitecture*. IEEE Computer Society, Washington, DC, USA, 29–.

[47] H. T. Nguyen, Y. Yagil, N. Seifert, and M. Reitsma. 2005. Chip-level soft error estimation method. *IEEE Transactions on Device and Materials Reliability* 5, 3 (Sept 2005), 365–381.

[48] Che Pan. 2021. TSMC says Nanjing wafer fab expansion plans on track as second quarter revenue surges 28 per cent. https://www.scmp.com/tech/tech-trends/article/3141240/tsmc-says-nanjing-wafer-fab-expansion-plans-track-second-quarter?module=perpetual_scroll_0&pgtype=article&campaign=3141240 Accessed: 2023-04-25.

[49] George Papadimitriou, Athanasios Chatzidimitriou, and Dimitris Gizopoulos. 2019. Adaptive Voltage/Frequency Scaling and Core Allocation for Balanced Energy and Performance on Multicore CPUs. In *2019 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. 133–146. https://doi.org/10.1109/HPCA.2019.00033

[50] George Papadimitriou, Athanasios Chatzidimitriou, Dimitris Gizopoulos, Vijay Janapa Reddi, Jingwen Leng, Behzad Salami, Osman Sabri Unsal, and Adrian Cristal Kestelman. 2020. Exceeding Conservative Limits: A Consolidated Analysis on Modern Hardware Margins. *IEEE Transactions on Device and Materials Reliability* 20, 2 (2020), 341–350. https://doi.org/10.1109/TDMR.2020.2989813

[51] George Papadimitriou, Athanasios Chatzidimitriou, Manolis Kaliorakis, Yannos Vastakis, and Dimitris Gizopoulos. 2018. Micro-Viruses for Fast System-Level Voltage Margins Characterization in Multicore CPUs. In *2018 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. 54–63. https://doi.org/10.1109/ISPASS.2018.00014

[52] George Papadimitriou and Dimitris Gizopoulos. 2021. Characterizing Soft Error Vulnerability of CPUs Across Compiler Optimizations and Microarchitectures. In *2021 IEEE International Symposium on Workload Characterization (IISWC)*. 113–124. https://doi.org/10.1109/IISWC53511.2021.00021

[53] George Papadimitriou and Dimitris Gizopoulos. 2021. Demystifying the System Vulnerability Stack: Transient Fault Effects Across the Layers. In *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*. 902–915. https://doi.org/10.1109/ISCA52012.2021.00075

[54] George Papadimitriou and Dimitris Gizopoulos. 2022. Anatomy of On-Chip Memory Hardware Fault Effects Across the Layers. *IEEE Transactions on Emerging Topics in Computing* (2022), 1–12. https://doi.org/10.1109/TETC.2022.3205808

[55] George Papadimitriou and Dimitris Gizopoulos. 2023. Silent Data Corruptions: Microarchitectural Perspectives. *IEEE Trans. Comput.* (2023), 1–13. https://doi.org/10.1109/TC.2023.3285094

[56] George Papadimitriou, Dimitris Gizopoulos, Harish Dattatraya Dixit, and Sriram Sankar. 2023. Silent Data Corruptions: The Stealthy Saboteurs of Digital Integrity. In *2023 IEEE 29th International Symposium on On-Line Testing and Robust System Design (IOLTS)*. 1–7. https://doi.org/10.1109/IOLTS59296.2023.10224870

[57] George Papadimitriou, Manolis Kaliorakis, Athanasios Chatzidimitriou, Dimitris Gizopoulos, Peter Lawthers, and Shidhartha Das. 2017. Harnessing Voltage Margins for Energy Efficiency in Multicore CPUs. In *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture* (Cambridge, Massachusetts) *(MICRO-50 '17)*. Association for Computing Machinery, New York, NY, USA, 503–516. https://doi.org/10.1145/3123939.3124537

[58] George Papadimitriou, Manolis Kaliorakis, Athanasios Chatzidimitriou, Charalampos Magdalinos, and Dimitris Gizopoulos. 2017. Voltage margins identification on commercial x86-64 multicore microprocessors. In *2017 IEEE 23rd International Symposium on On-Line Testing and Robust System Design (IOLTS)*. 51–56. https://doi.org/10.1109/IOLTS.2017.8046198

[59] Moinuddin K. Qureshi and Zeshan Chishti. 2013. Operating SECDED-Based Caches at Ultra-Low Voltage with FLAIR. In *Proceedings of the 2013 43rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN) (DSN '13)*. IEEE Computer Society, USA, 1–11. https://doi.org/10.1109/DSN.2013.6575314

[60] David Roberts, Nam Sung Kim, and Trevor Mudge. 2007. On-Chip Cache Device Scaling Limits and Effective Fault Repair Techniques in Future Nanoscale Technology. In *10th Euromicro Conference on Digital System Design Architectures, Methods and Tools (DSD 2007)*. 570–578. https://doi.org/10.1109/DSD.2007.4341526

[61] Gennaro Severino Rodrigues and Fernanda Lima Kastensmidt. 2016. Soft error analysis at sequential and parallel applications in ARM Cortex-A9 dual-core. In *2016 17th Latin-American Test Symposium (LATS)*. IEEE. https://doi.org/10.1109/latw.2016.7483359

[62] Felipe Rosa, Fernanda Kastensmidt, Ricardo Reis, and Luciano Ost. 2015. A fast and scalable fault injection framework to evaluate multi/many-core soft error reliability. In *2015 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFTS)*. IEEE. https://doi.org/10.1109/dft.2015.7315164

[63] T. Santini, L. Carro, F. Rech Wagner, and P. Rech. 2016. Reliability Analysis of Operating Systems and Software Stack for Embedded Systems. *IEEE Transactions*

on Nuclear Science 63, 4 (Aug 2016), 2225–2232. https://doi.org/10.1109/TNS.2015.2513384

[64] Hiroshi Sasaki, Alper Buyuktosunoglu, Augusto Vega, and Pradip Bose. 2016. Characterization and mitigation of power contention across multiprogrammed workloads. In *2016 IEEE International Symposium on Workload Characterization (IISWC)*. 1–10. https://doi.org/10.1109/IISWC.2016.7581266

[65] James R. Schwank, Marty R. Shaneyfelt, and Paul E. Dodd. 2013. Radiation Hardness Assurance Testing of Microelectronic Devices and Integrated Circuits: Test Guideline for Proton and Heavy Ion Single-Event Effects. *IEEE Transactions on Nuclear Science* 60, 3 (2013), 2101–2118. https://doi.org/10.1109/TNS.2013.2261317

[66] N. Seifert, D. Moyer, N. Leland, and R. Hokinson. 2001. Historical trend in alpha-particle induced soft error rates of the Alpha/sup TM/ microprocessor. In *2001 IEEE International Reliability Physics Symposium Proceedings. 39th Annual (Cat. No.00CH37167)*. 259–265. https://doi.org/10.1109/RELPHY.2001.922911

[67] Norbert Seifert, Xiaowei Zhu, and Lloyd W Massengill. 2002. Impact of scaling on soft-error rates in commercial microprocessors. *IEEE Transactions on Nuclear Science* 49, 6 (2002), 3100–3106.

[68] A. Shah. [n. d.]. Chip manufacturers are going back to the future for automotive silicon. https://www.theregister.com/2021/10/19/chip_manufacturer_chips/ Accessed: 2023-04-25.

[69] Adit Singh, Sreejit Chakravarty, George Papadimitriou, and Dimitris Gizopoulos. 2023. Silent Data Errors: Sources, Detection, and Modeling. In *2023 IEEE 41st VLSI Test Symposium (VTS)*. 1–12. https://doi.org/10.1109/VTS56346.2023.10139970

[70] Vilas Sridharan and David R. Kaeli. 2010. Using Hardware Vulnerability Factors to Enhance AVF Analysis. In *Proceedings of the 37th Annual International Symposium on Computer Architecture* (Saint-Malo, France) *(ISCA '10)*. ACM, New York, NY, USA, 461–472. http://doi.acm.org/10.1145/1815961.1816023

[71] Sriram Sundaram, Sriram Samabmurthy, Michael Austin, Aaron Grenat, Michael Golden, Stephen Kosonocky, and Samuel Naffziger. 2016. Adaptive Voltage Frequency Scaling Using Critical Path Accumulator Implemented in 28nm CPU. In *2016 29th International Conference on VLSI Design and 2016 15th International Conference on Embedded Systems (VLSID)*. 565–566. https://doi.org/10.1109/VLSID.2016.106

[72] Md. Nasim Taj. 2022. Intel Chip Manufacturing Technology Roadmap. https://doi.org/10.13140/RG.2.2.11629.46566

[73] Jorge Tonfat, José Rodrigo Azambuja, Gabriel Nazar, Paolo Rech, Fernanda Lima Kastensmidt, Luigi Carro, Ricardo Reis, Juliano Benfica, Fabian Vargas, Eduardo Bezerra, and Christopher Frost. 2014. Measuring the impact of voltage scaling for soft errors in SRAM-based FPGAs from a designer perspective. In *International Mixed-Signals, Sensors, and Systems Test Workshop Proceedings*. 1–6. https://doi.org/10.1109/IMS3TW.2014.6997398

[74] Konstantinos Tovletoglou, Lev Mukhanov, Georgios Karakonstantis, Athanasios Chatzidimitriou, George Papadimitriou, Manolis Kaliorakis, Dimitris Gizopoulos, Zacharias Hadjilambrou, Yiannakis Sazeides, Alejandro Lampropulos, Shidhartha Das, and Phong Vo. 2018. Measuring and Exploiting Guardbands of Server-Grade ARMv8 CPU Cores and DRAMs. In *48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W)*. 6–9. https://doi.org/10.1109/DSN-W.2018.00013

[75] Ioannis Tsiokanos, George Papadimitriou, Dimitris Gizopoulos, and Georgios Karakonstantis. 2021. Boosting Microprocessor Efficiency: Circuit- and Workload-Aware Assessment of Timing Errors. In *2021 IEEE International Symposium on Workload Characterization (IISWC)*. 125–137. https://doi.org/10.1109/IISWC53511.2021.00022

[76] Vasanth Venkatachalam and Michael Franz. 2005. Power Reduction Techniques for Microprocessor Systems. *ACM Comput. Surv.* 37, 3 (sep 2005), 195–237. https://doi.org/10.1145/1108956.1108957

[77] Nicholas J. Wang, Aqeel Mahesri, and Sanjay J. Patel. 2007. Examining ACE analysis reliability estimates using fault-injection. *ACM SIGARCH Computer Architecture News* 35, 2 (Jun 2007), 460. https://doi.org/10.1145/1273440.1250719

[78] Chris Wilkerson, Alaa R. Alameldeen, Zeshan Chishti, Wei Wu, Dinesh Somasekhar, and Shih-lien Lu. 2010. Reducing Cache Power with Low-Cost, Multi-Bit Error-Correcting Codes. *SIGARCH Comput. Archit. News* 38, 3 (jun 2010), 83–93. https://doi.org/10.1145/1816038.1815973

[79] Chris Wilkerson, Hongliang Gao, Alaa R. Alameldeen, Zeshan Chishti, Muhammad Khellah, and Shih-Lien Lu. 2008. Trading off Cache Capacity for Reliability to Enable Low Voltage Operation. In *Proceedings of the 35th Annual International Symposium on Computer Architecture (ISCA '08)*. IEEE Computer Society, USA, 203–214. https://doi.org/10.1109/ISCA.2008.22

[80] Chris Wilkerson, Hongliang Gao, Alaa R. Alameldeen, Zeshan Chishti, Muhammad Khellah, and Shih-Lien Lu. 2008. Trading off Cache Capacity for Reliability to Enable Low Voltage Operation. *SIGARCH Comput. Archit. News* 36, 3 (jun 2008), 203–214. https://doi.org/10.1145/1394608.1382139

[81] Kai-Chiang Wu and Diana Marculescu. 2008. Power-aware soft error hardening via selective voltage scaling. In *2008 IEEE International Conference on Computer Design*. 301–306. https://doi.org/10.1109/ICCD.2008.4751877

[82] Shengqi Yang, Wenping Wang, Tiehan Lu, Wayne Wolf, N. Vijaykrishnan, and Yuan Xie. 2008. Case Study of Reliability-Aware and Low-Power Design. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 16, 7 (2008), 861–873. https://doi.org/10.1109/TVLSI.2008.2000460

[83] Weitao Yang, Yonghong Li, Yang Li, Zhiliang Hu, Fei Xie, Chaohui He, Songlin Wang, Bin Zhou, Huan He, Waseem Khan, and Tianjiao Liang. 2019. Atmospheric neutron single event effect test on Xilinx 28 nm system on chip at CSNS-BL09. *Microelectronics Reliability* 99 (2019), 119–124. https://doi.org/10.1016/j.microrel.2019.05.004

[84] Dakai Zhu, R. Melhem, and D. Mosse. 2004. The effects of energy management on reliability in real-time embedded systems. In *IEEE/ACM International Conference on Computer Aided Design, 2004. ICCAD-2004.* 35–40. https://doi.org/10.1109/ICCAD.2004.1382539

[85] James F Ziegler and Helmut Puchner. 2010. *SER–history, Trends and Challenges: A Guide for Designing with Memory ICs*. Cypress.

[86] Brian Zimmer, Seng Oon Toh, Huy Vo, Yunsup Lee, Olivier Thomas, Krste Asanovic, and Borivoje Nikolic. 2012. SRAM Assist Techniques for Operation in a Wide Voltage Range in 28-nm CMOS. *IEEE Transactions on Circuits and Systems II: Express Briefs* 59, 12 (2012), 853–857. https://doi.org/10.1109/TCSII.2012.2231015

[87] Yazhou Zu, Charles R. Lefurgy, Jingwen Leng, Matthew Halpern, Michael S. Floyd, and Vijay Janapa Reddi. 2015. Adaptive guardband scheduling to improve system-level efficiency of the POWER7+. In *2015 48th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 308–321. https://doi.org/10.1145/2830772.2830824