

Logical Maneuvers: Detecting and Mitigating Adversarial Hardware Faults in Space

Fatemeh Khojasteh Dana*
Worcester Polytechnic Institute
fdana@wpi.edu

Saleh Khalaj Monfared*
Worcester Polytechnic Institute
skmonfared@wpi.edu

Shahin Tajik
Worcester Polytechnic Institute
stajik@wpi.edu

* These authors contributed equally to this work

Abstract—Satellites are highly vulnerable to adversarial glitches or high-energy radiation in space, which could cause faults on the onboard computer. Various radiation- and fault-tolerant methods, such as error correction codes (ECC) and redundancy-based approaches, have been explored over the last decades to mitigate temporary soft errors on software and hardware. However, conventional ECC methods fail to deal with hard errors or permanent faults in the hardware components. This work introduces a detection- and response-based countermeasure to deal with partially damaged processor chips. It recovers the processor chip from permanent faults and enables *continuous operation with available undamaged resources* on the chip. We incorporate digitally-compatible delay-based sensors on the target processor's chip to reliably detect the incoming radiation or glitching attempts on the physical fabric of the chip, even before a fault occurs. Upon detecting a fault in one or more components of the processor's arithmetic logic unit (ALU), our countermeasure employs adaptive software recompilations to resynthesize and substitute the affected instructions with instructions of still functioning components to accomplish the task. Furthermore, if the fault is more widespread and prevents the correct operation of the entire processor, our approach deploys adaptive hardware partial reconfigurations to replace and reroute the failed components to undamaged locations of the chip. To validate our claims, we deploy a high-energy near-infrared (NIR) laser beam on a RISC-V processor implemented on a 28 nm FPGA to emulate radiation and even hard errors by partially damaging the FPGA fabric. We demonstrate that our sensor can confidently detect the radiation and trigger the processor testing and fault recovery mechanisms. Finally, we discuss the overhead imposed by our countermeasure.

I. INTRODUCTION

Satellites are vulnerable to cyberattacks that can destroy them without creating debris, minimizing collateral damage, or taking control of them entirely. Attackers could manipulate satellites in various ways, such as depleting fuel, altering orbits, holding onboard computers for ransom, draining power, or damaging sensors by exposing them to direct sunlight [1]. The shift toward hosting multiple payloads on commercial satellite buses with shared resources (e.g., power source, transponders, data buses, etc.) creates novel cyber-physical vulnerabilities for hosted payloads if one of the payloads is infected and goes rogue, see Fig 1. Such a malicious payload can inject

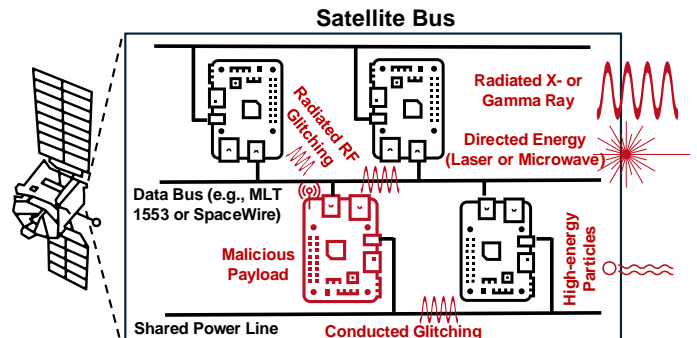


Fig. 1: Fault and radiation threats for hosted payloads on satellites

physical faults into the adjacent payloads' computers through the shared power lines, data buses, or simply vacuum using conducted or radiated electromagnetic glitches. While some proprietary commercial solutions exist for power isolation of hosted payloads to reduce the information leakage [2], it is unclear how effective they are against fault injection attacks. On the other hand, satellite onboard computers are threatened by external high-energy electromagnetic (EM) signals and charged particles caused by directed energy weapons, nuclear detonations, or cosmic rays in space [3]. Such high-energy EM pulses and particles can permanently damage the onboard computer chips.

Radiation hardening (RadHard) and fault-tolerant solutions have been widely researched in the literature to deal with such threats. Error correction codes (ECCs) and redundancy-based methods, such as triple modular redundancy (TMR), are prime examples of detection and recovery schemes from faults. While ECCs are effective against temporary soft errors or transient glitches to detect and correct the corrupted instruction or data, they are ineffective against permanent hard errors on the functional part of the chip as such hardware faults cannot be fixed by the running software. Moreover, while TMR methods can deal with hard errors to some extent using voting, they impose very large overheads on the system, which have constrained size, weight, and power (SWaP) requirements. In addition, due to the lack of sensing capability of these algorithmic countermeasures about the events occurring at the physical layer of a computing chip, they should always be active to respond to possible faults. Interestingly, hard errors are not special to space applications as the shrinking trend of

transistor technology leads to less reliable hardware [4]. There are multiple reports that processors fail randomly due to hard errors in large data centers [5], [6]. While it is straightforward to replace chips on the ground, it is virtually impossible to replace them in a spacecraft in orbit in most cases.

The physical detection of glitching attempts before fault occurrence and the response to it have been researched in the field of hardware security. Deploying delay-based sensors is the primary way for sensing analog disturbances on the chip power delivery network (PDN) caused by voltage [7], EM [8], and laser glitches [9]. Upon detection of such disturbances, the chips can respond by locking the computation [10] or ignoring the processed data [7] for a few prior clock cycles. In contrast to classical hardware security scenarios, where the goal of a fault adversary is to break the confidentiality or violate the integrity of the system, the main goal for fault injection on space assets is to cause safety and availability issues [11]. Hence, while current detection methods could benefit the satellite onboard computers, most conventional response methods are unacceptable as they could make the system unavailable, creating safety issues for a satellite in orbit.

Hence, in this work, we ask the following research questions: *Is it possible to deploy sensors to detect the glitching/radiation attempt in space, recover from a hardware fault, and continue the computation even with a partially damaged processor circuit without imposing a large SWaP overhead and without having redundant processors on standby?*

Our Contribution. In this work, we positively answer the research questions above. We first demonstrate that any glitching attempt or radiation causes voltage fluctuations on the PDN of the processor chip and, consequently, affects the propagation delays of signals (e.g., clock signal) on the chip. We show that such delay variations can be measured reliably by time-to-digital converters (TDCs) [7], [12]. Upon detection of delay variations, a software test is triggered to detect whether a fault has occurred. By detecting a fault in the processor’s arithmetic logic unit (ALU), software recompilation and resynthesis will be performed to replace the affected arithmetic or logic instruction with other available and still functioning instructions to accomplish the task. On the other hand, if a fault makes the entire processor dysfunctional, adaptive hardware partial reconfigurations will be executed to move the damaged parts of the processor to a functional region on the chip.

To show the effectiveness of our approach, we realize a soft RISC-V processor on a common field programmable gate array (FPGA) manufactured with 28 nm technology with TDC sensors. We emulate the radiation on the sensors by employing a near-infrared (NIR) laser beam. To emulate hard errors, we take a step further and create permanent damage on the configuration logic blocks of the FPGA fabric using our laser with higher power. We demonstrate that our software and hardware recovery methods heal the processor and enable continuous operation without any interruption. Finally, we discuss the overhead of our countermeasure in terms of memory, area, and performance overhead.

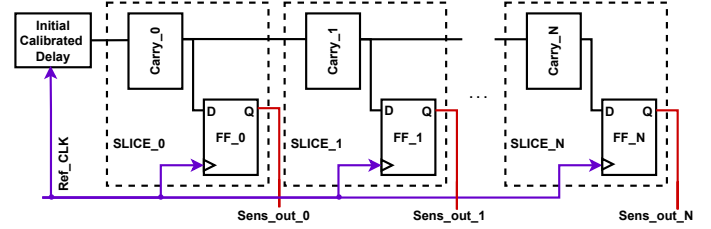


Fig. 2: high-level implementation of a TDC-based fault detection sensor

II. BACKGROUND

A. Hard Errors

The impact of radiation on semiconductor devices can be categorized into soft and hard errors. While soft errors are temporary and can often be corrected through normal operation or a power cycle, hard errors are permanent and result in irreversible damage [13]. The common causes of hard errors are total ionizing dose (TID), single-event burnout (SEB), and electrostatic discharge (ESD). Such hard errors can be modeled as stuck-at-faults, bridge faults, or delay faults. In extreme cases, hard errors cause short or open circuits on transistors connecting VDD to the GND, which makes the circuit nonoperational.

B. Delay-based Sensors

In glitching and radiation attacks, the injected energy causes disturbances on the power line of the chip. Such disturbances can affect the propagation delays of electrical signals in delay-based sensors, such as ring oscillators (ROs) [14] and time-to-digital converters (TDCs) [15], [16]. These sensors have been used to detect voltage, electromagnetic (EM), and laser glitching and radiating attacks on FPGAs, measuring voltage fluctuations on the power delivery network (PDN) caused by injected energy [9]. The main drawbacks of using ROs are their high power consumption and sensitivity to environmental conditions, making them unattractive for space applications requiring low power consumption. Therefore, in this paper, we deploy TDC sensors for attack detection.

We adopt the TDC implementation described by Mahmoud et.al. [17], which is primarily designed for remote side-channel attacks on FPGA. We repurpose the TDC design for fault-sensing applications. A high-level illustration of the employed TDC is shown in Fig. 2. The fundamental building blocks of the sensor comprise an initial delayed signal followed by a tapped delay line and the corresponding output registers. The calibration is carried out offline using adjustable delay elements (e.g., IDELAY [18]) or by including/excluding combinational logic blocks in a chain to ensure that the sensor’s output is in meta-stable status. Ideally, this means that for a N -bit TDC sensor, calibration is achieved if output tends to a value with $HW(Output) = N/2$ during a long enough sampling time, where HW represents the Hamming Weight. The observable delay line is often deployed by fast carry propagation logic units constrained physically to organize a chain of delay. The output of each of these delay units is then propagated to an output register clocked with the original signal. Based on the

propagation delay of the input of each register, the output can be constituted as a binary sensor.

C. Software Instruction Synthesis

The arithmetic logic unit (ALU) is a fundamental building block of various computing machines. The question here is if a computer can accomplish an arithmetic or logical task if one or more components (e.g., Multiplier, XOR, etc.) of the ALU fail. To answer this question, we should understand to what extent it is possible to perform the failed operation/instruction with other available operations/instructions.

Various instruction set architectures (ISA) define how software communicates with the hardware. While there are complex instruction set computers (CISC), such as x86 architecture, in which a single instruction can execute several low-level operations, reduced instruction set computers (RISC) have more simplified instructions in exchange for requiring more instructions to accomplish some complex tasks. In principle, we can reduce the number of instructions to a single instruction to build a One-instruction set computer (OISC) [19], [20] or even a computer without instructions. For instance, it has been shown that the x86 `mov` instruction [21], x86 memory management unit (MMU) page faults [22], or Remote Direct Memory Access (RDMA) offloads are Turing-complete [23], and hence, all instructions can be accomplished by one instruction or operation.

While an OISC is an extreme architecture with a very large overhead in terms of execution time, it is still possible in the case of failure to substitute individual arithmetic or logical instructions by combining and synthesizing other instructions. For example, if the processor's multiplier fails, the processor can still execute the multiplication using add and shift instructions. Therefore, in our hard error scenario, depending on the location of an error, a processor with a partially damaged ALU might still be able to accomplish tasks though in a longer time.

D. Hardware Partial Reconfiguration

Partial Reconfiguration (PR) is a feature of mainstream field programmable gate arrays (FPGAs), allowing for dynamic modifications of a portion of the FPGA circuit without needing a complete system reboot. At the same time, the rest of the system continues to operate uninterrupted, [24]. The use of PR in FPGAs is particularly advantageous in mission-critical applications, like space applications [25], where system downtime is not acceptable.

AMD/Xilinx's Dynamic Function eXchange (DFX) introduces an approach for defining Partially Reconfigurable (PR) regions within the FPGA fabric [26]. However, the Vivado toolchain presents significant limitations, including slow performance for real-time applications and a lack of bitstream relocation support, constraining the reconfigurability [27]. Several notable tools have emerged to address these challenges. For instance, the open-source tool Byteman [27] has substantially improved bitstream manipulation capabilities. It enhances efficiency, speed, and compatibility by supporting the merging of clock, Configurable Logic Block (CLB), and Block RAM data,

along with implementing different merge strategies. These bitstream manipulation tools are crucial, as they provide the ability to generate and deploy partial bitstreams in realtime. Such real-time partial reconfigurations have been recently deployed as countermeasures against side-channel attacks [9], [28], where a processor can choose between multiple partial bitstreams generated in an offline phase and partially reconfigure the circuits on the programmable logic fabric to add noise to physical information leakages. Hence, a housekeeping RadHard processor can use a similar method to replace and reroute defective components of a soft processor in orbit.

III. THREAT MODEL

For our threat model, we assume two types of adversaries, namely, a cyber-physical adversary and a non-kinetic adversaries [29]. We also consider cosmic radiation as an unintentional threat. We assume that there is a RadHard housekeeping microcontroller onboard, which is resilient to radiations and glitches and can control commercial and non-rad-tolerant computer chips. We only consider scenarios where the core of the non-rad-tolerant computer chips is partially damaged, but its communication links to the housekeeping MCU and its JTAG are still functional. Otherwise, our countermeasure will not be useful if the entire chip is fried or it cannot be programmed anymore. Finally, we assume that FPGA scrubbing and ECC are deployed to detect and correct other faults, such as memory contents.

A. Cyber-physical Adversaries

For the cyber-physical attackers, we assume a hosted payload environment in which a satellite bus accommodates multiple payloads by different vendors. The satellite's bus resources, such as power lines, data buses, and physical space, are shared between the payloads. We assume that one of the payloads is infected either by a remote cyber adversary or a supply-chain adversary using malware or hardware Trojans. Therefore, the infected payload can perform physical faults injection attacks, such as voltage glitches/conducted interference through the shared power line or EM glitches/radiated interference through the shared unshielded space inside the bus enclosure, on the adjacent payload's computers or even the computer of the infected system itself, see Fig. 1. The adversary's goal is to affect their functionality or reliability. Examples of such remote fault attacks have been reported for multi-tenant FPGA [12], [30] deployments in the cloud, where the adversary can cause soft and hard errors on the chip.

B. Non-kinetic Adversaries

For the non-kinetic adversaries, we assume that directed energy weapons (from a ground station or another satellite) or nuclear detonation in space are deployed to damage the satellite electronics in orbit. While such attacks are not considered cyber threats, they still can partially damage the onboard electronics and computers, and their effect could be similar to cyber-physical attacks, though on a much larger scale. Directed energy weapons, such as high-energy lasers, can thermally affect the satellite components and cause a current surge in the

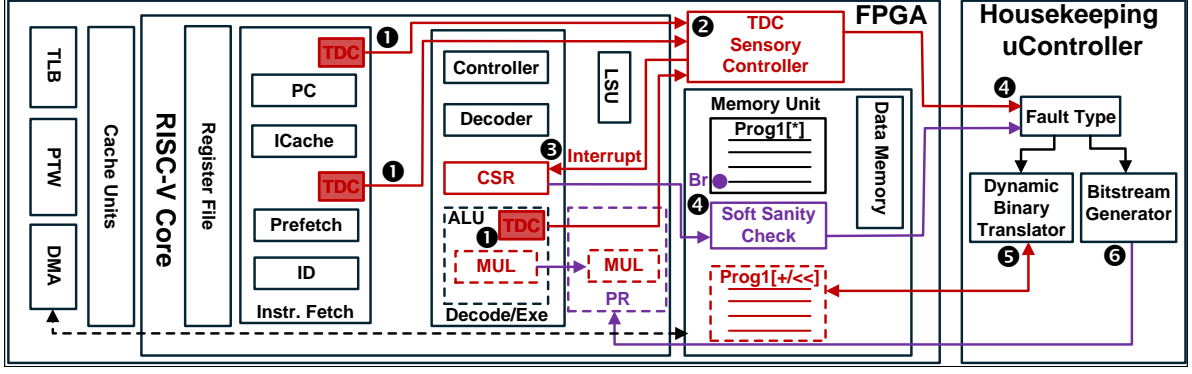


Fig. 3: The Proposed framework

power/data lines of the onboard computer. On the other hand, a nuclear blast in space generates high-energy electromagnetic pulses in the form of X-ray and Gamma radiation, which can ionize the material and create electrical currents that can damage computer chips.

C. Cosmic Radiation

Finally, the source of faults could be natural. Cosmic rays are high-energy particles or particle clusters, mainly protons or atomic nuclei, traveling through space at nearly the speed of light. They originate from the Sun, our galaxy, or even distant galaxies. These high-energy particles can penetrate the enclosures and chip packages and permanently damage the transistors.

IV. APPROACH

This section describes our proposed approach to address hardware-level faults in real-time. Our solution utilizes a hardware/software co-design to mitigate hardware faults in the mission-critical processors deployed in space applications. Concurrently, we seek to incur minimal hardware modification to the existing hardware architectures to preserve the highest level of compatibility. We consider a soft processor on a reconfigurable fabric, such as FPGA, and exploit the reconfigurability to tackle localized permanent hardware damage. Additionally, our main focus is the hardware permanent errors on the silicon as soft errors, such as bitflips in off-core memory units, can be corrected using scrubbing techniques [31], error correction codes [32], re-execution [33], [34], redundancy methods [35], and partial reconfiguration [25].

Figure 3 depicts a high-level representation block diagram of the proposed framework. We consider a soft-core implementation of a RISC-V architecture. We equipped it with additional digital sensors distributed throughout the FPGA design. Highlighted in ①, TDC sensors are placed in different physical locations on the core and are calibrated to detect local disturbance during the execution. Furthermore, we incorporate a hardware sensor controller that constantly checks the output of the sensors. The TDC controller in ② keeps track of environmental conditions for each building block of the processing core. All sensors are tagged and sampled with a relatively higher sampling rate in comparison with the core's clock frequency. The controller verifies the condition of each local

block by comparing the tagged sensor data with its pre-defined threshold that is derived during the calibration. Moreover, we extend the architecture to support an external non-maskable interrupt (NMI) from the sensor controller hardware. If a fault attack is initiated, the environmental disturbance changes the behavior of the TDC sensors, which is detected by the controller in real time. As a result, the core is interrupted with the highest priority of NMI [36] as illustrated in ③. Upon the occurrence of the interrupt, the processor breaks into a pre-defined software check where it verifies the functionality of its different components. As shown in ④, the sanity check software performs a comprehensive test and transmits the test's results to the off-chip Radiation Hardened Housekeeping μ -controller. At the same time, based on the sensory data captured from the TDCs, tags for the potentially faulty building blocks are sent to the Housekeeping system. The fault type detected could be due to several failures diagnosed by the soft sanity check or a sensor abnormality report. Depending on the fault, a software-based response can be taken into account by utilizing a binary translator. As highlighted in ⑤, if the indicated fault type, for instance, is associated with *Multiplier* part of the ALU, a real-time machine-code level translator is invoked to transform the target program to use non-faulty units (e.g., instructions). On the other hand, if multiple execution units are (or about to be) permanently damaged, a full/partial reconfiguration command is sent from the μ controller to the FPGA (e.g., Internal Configuration Access Port ICAP [37]) to re-implement & reconfigure the affected units as indicated in ⑥. Note that the reconfiguration process incurs a relatively higher timing overhead than the binary translation method. However, in scenarios where permanent faults are manifested in multiple logic units, the system cannot be recovered by software, and reconfiguration is the only option for recovery.

V. EXPERIMENTAL SETUP

A. Device Under Test

We used two Genesys 2 development kits [38] for the experiments. These kits contain an AMD/Xilinx Kintex 7 (XC7K325T-2FFG900C) FPGA manufactured with 28 nm technology. The die of the FPGA under test is packaged in a flip-chip package, enabling access to the chip's backside silicon and exposing the FPGA fabric to laser radiation. We did not perform any other modifications (e.g., silicon polishing or

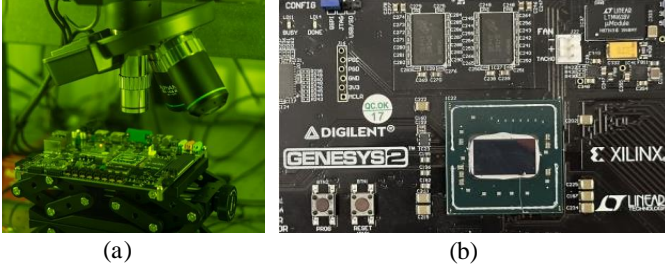


Fig. 4: (a) Device under test under the objective lenses of ALpHANOv setup. (b) The Kintex 7 FPGA chip is in a flip-chip package with a removed heatsink from its backside silicon.

thinning) to the package or board. The FPGA core was supplied by 1.0 V for all the experiments, and the global clock operated at 200 MHz. One kit is used as a reference for all experiments, and the second one was intentionally damaged by the laser beam for hard error emulation.

B. Laser Setup

We used an ALPhANOv S-LMS [39] setup for photon emission microscopy, laser radiation, and fault injection. The microscope consists of a camera system for navigating, an In-GaAs camera for photon mission analysis, and a 1064 nm wavelength laser source. The microscope has been equipped with multiple lenses with 2.5X, 20X (NA=0.6), and 50X (NA=0.7) magnifications on an XYZ stage to focus on a region. The combined setup is controlled using the software and hardware switches to control the XYZ stage and camera options. The software provides an IR view of the die, which can be used for navigation. The peak currents used for radiation and hard error experiments were 1 A and 2.5 A, respectively, while the pulse width was 250 ns with a frequency of 100 kHz. The laser is controlled by the ALPhANOv control software in combination with viewing the die using the camera software. The combination results in a live feed to laser shots on the desired fault region (a screenshot shown in 4). The InGaAs camera was used to verify the FPGA’s damage location by capturing the created short-circuits’ photon emission.

C. Hardware Implementation

A single-core, 32-bit RISC-V is implemented on the FPGA under test. We used the open-source Rocket Chip generator to instantiate the RISC-V Rocket Core [40]. We used an SD card to store and load the executable and linkable format (ELF) file containing bare-metal programs for the processor. We used two parallel UARTs, one for communicating with the running code on the processor and the other one for reading back the sensors’ values. In our experiments, we used a computer to emulate the housekeeping microcontroller unit (MCU) to manage the software recompilations and FPGA’s reconfigurations. As illustrated in Fig. 5 (a), two TDC sensors (represented by blue colors) are positioned in different parts of the FPGA. One of these sensors is placed near the ALU embedded within the processor to maximize the likelihood of detecting errors. Table I provides an overview of the FPGA resource utilization by various modules in the Rocket System

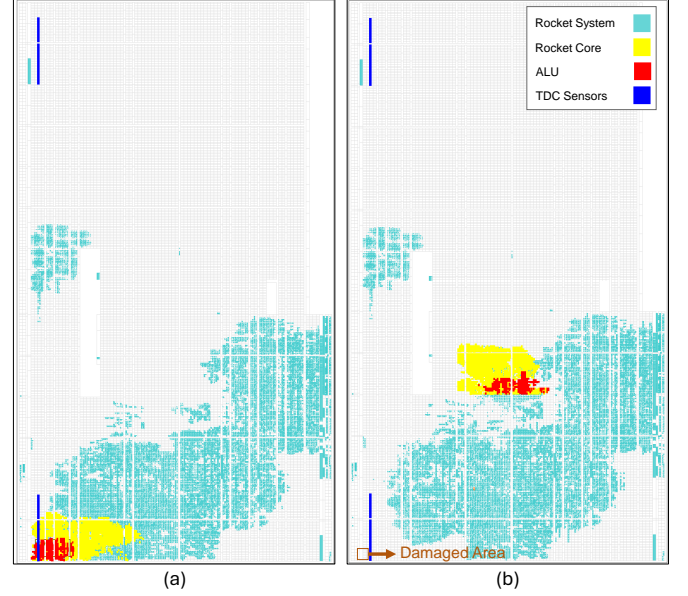


Fig. 5: Two different layouts of the entire Rocket System implementation on Kintex-7 FPGA. The region highlighted in yellow is the core excluding the ALU, and the region highlighted in red is the ALU. (a) The ALU and core are close to the TDC sensor. (b) The ALU and core are replaced and rerouted from the defect region.

in terms of the number of registers and Look-Up Tables (LUTs). There is a hierarchical relationship between the three modules. The Rocket System module encompasses the complete system, including Rocket Core and ALU components. The Rocket Core module contains only the processor core, including the ALU. Peripheral and additional systems (e.g., JTAG and I/Os) are not part of this module. The TDC sensor modules have a minimal resource consumption, and thus, their overheads are negligible. For this implementation, we employed two TDC sensors, each with a width of 128 bits and 32 LUT_5s for the calibration circuitry. The LUT count for ALU reflects 19.4% usage of Rocket Core, and most of its resources are consumed by the multiplier.

TABLE I: Usage Resources

Module	Registers	LUTs
Rocket System	6,350	15,359
Rocket Core	1,557	3,179
ALU	125	617
TDC Sensors	320	64

D. Software Implementation

We selected two benchmarks, namely, a multiply-accumulate (MAC) function along with AND operations and a Reed-Solomon coding scheme. Both benchmarks have been coded using the C programming language. The C codes were compiled into assembly codes. We wrote a translator script to automatically replace and substitute the instructions in the generated output assembly codes. Using the translator, four variations

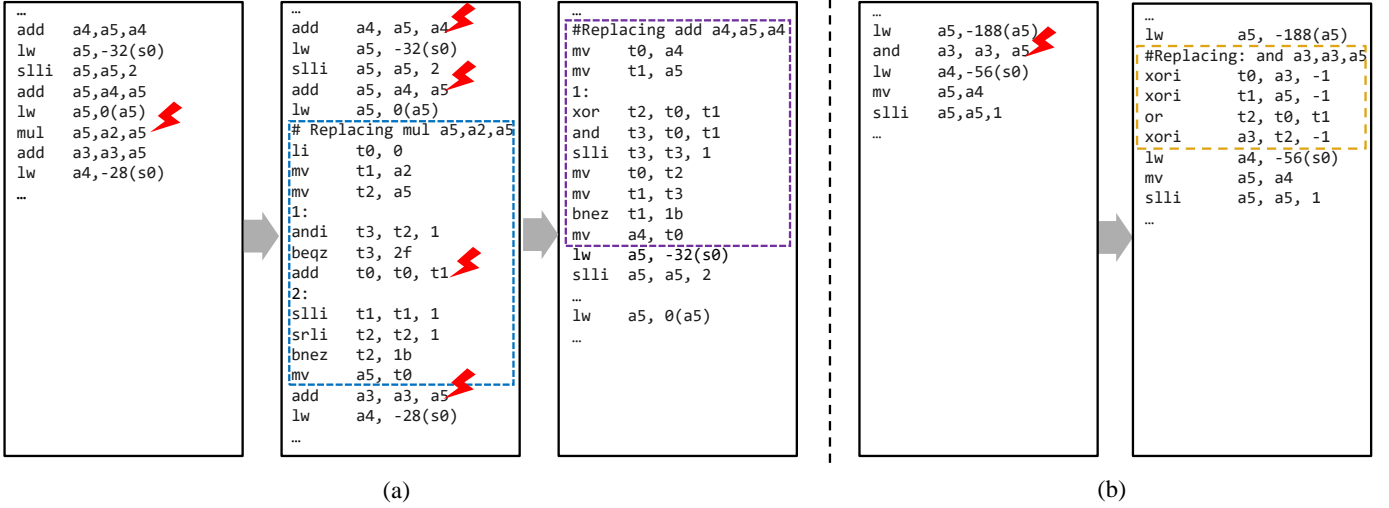


Fig. 6: (a) Translation and resynthesis of arithmetic instructions: Example of substitution of `mul` instructions upon the failure of the multiplier. Further substitution of `add` instruction in the case of addition failure. (b) Translation and resynthesis of logic instructions: Example of substitution of `and` instructions upon the failure of the AND gate in ALU.

of the assembly code for each benchmark were generated. In the first version, the original code is used without any modifications. In the second version, all multiplier operations are replaced with equivalent add-and-shift operations. In the third version, the multiplication and addition operations are substituted with XOR and AND functions. In the last version, to consider the effect of logic failure, the AND operation is replaced with other logic operations.

A script reads the outputs of the TDC sensors via UART. If an error is detected, a diagnostic test routine (soft sanity check) is executed to identify the faulty submodule. This routine determines whether the issue lies with the multiplier, adder, shift register, or AND components. At this stage, a command is sent to the FPGA via the UART interface to execute a different variation of the code based on the identified faulty module.

VI. RESULTS

To evaluate our framework, we consider two fault scenarios that involve hard errors on the target's fabric. Specifically, first, we consider a temporary hardware fault induced by a laser beam and showcase how it is detected and mitigated by the proposed solution. Subsequently, we tackle the scenario where a physical part of the target chip is permanently damaged during the runtime. In the following, we distinctly detail each of these scenarios and describe our mitigation to address them.

A. Radiation Sensing

The radiation sensing system utilizes distributed digital sensors (i.e., TDC) throughout the target IC as described in Sect. IV. We evaluate radiation by statically analyzing digital samples extracted from the sensors and performing a comparative examination between normal and radiating environments emulated via NIR laser exposure.

1) *Normal Operation*: During the *Normal Operation*, the system is set to execute simple bare-metal applications. Simultaneously, two TDC sensors are deployed on the target's fabric as part of our solution in accordance with Fig. 5. As

explained in Sect. IV, the sensory controller keeps track of the output samples of these TDCs. For the sake of illustrious, Fig. 7 demonstrates the samples captured by the controller unit in the system. As anticipated, the HW of the TDC samples tends to show a consistent trend fixed around the value of $128/2 = 64$, where 128 is the width of the sensor. During the sampling process, a slight alteration is observed specifically for *TDC0*, which is physically located near the ALU circuitry. This alteration is possibly due to the local power consumption of the logical unit near the TDC0 sensor. In contrast, *TDC1*'s output depicts a more consistent behavior, possibly due to the fact that it is deployed far away from the execution logic of the core (Please refer to Fig. 5).

2) *Temporary Laser Exposed Operation*: Under the radiation emulation, we expose the backside of the IC with the laser with a setup as described in Sect. V. We limit the peak current usage of the laser to 1.0 A to ensure that neither temporary nor permanent faults are induced in the target IC. Our experiment shows that an average peak current of 1.5 A is required to induce temporary faults (soft errors) on this particular fabric. As a result, here, we merely evaluate our system against radiations that do not cause a fault. This is specifically aligned with scenarios where the target device is about to be exposed to a vulnerable environment in orbit or radiation power is swept by the adversary. In our experiments, we particularly 1) localized the ALU sub-system 2.5X lens, 2) carried out the focusing procedure with the 20X lens, and 3) activated the laser output and performed manual navigation for 5-10 seconds. During this experiment, we captured the TDC outputs from the sensor controller to evaluate the detection mechanism. Furthermore, note that as shown in Fig. 5, the *TDC0* is physically deployed in the vicinity of the ALU. Fig. 8 shows a sampling window of TDC's output during the laser exposure. As the laser stage is moved to expose blocks near the ALU, we can visually identify distinct jumps in the output of *TDC0*, which is located near the target logic. In our framework (see Sect. IV), this fluctuation

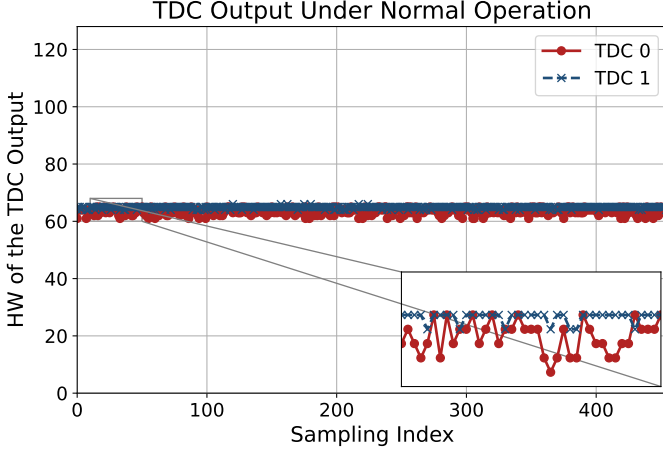


Fig. 7: The TDC outputs represented by HW during a normal operation of the target device

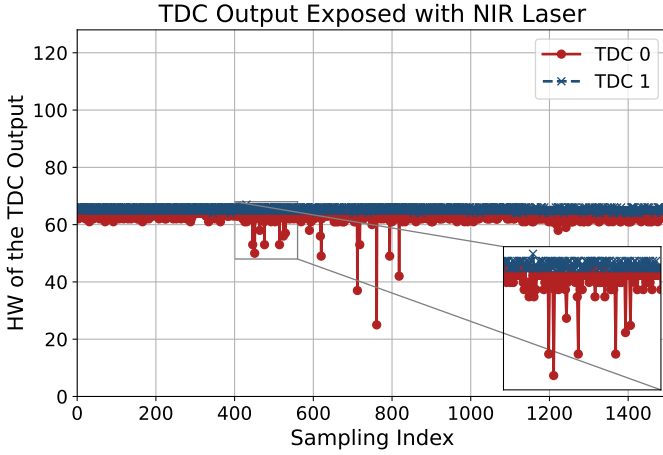


Fig. 8: The TDC outputs represented by HW when the target is exposed with NIR Laser

indicates a potential fault attack and invokes the mitigation mechanism in software and hardware, which are outlined in the following.

B. Instruction Resynthesis

Upon detection of an error in the ALU using the soft sanity check (see Sect. IV), the housekeeping microcontroller executes an algorithm to replace the affected assembly instructions. Here, we present the results of two test scenarios. In the first scenario, one or more arithmetic units, such as the multiplier and adder of the processor, fail. In the second scenario, a logical unit, e.g., the AND gate, fails. Fig. 6(a) shows the basic steps needed to substitute the `mul` instruction with `add` and shifting instructions [41]. The algorithm starts with the loading of the multiplicand number to the `t1` register and a load of the multiplier number to the `t2` register; it initializes the `t0` register to 0 as a product register. The least significant bit of the multiplier register (`t2`, calculated by `anding` it with 1) determines whether the multiplicand is added to the product register. Shifting the multiplicand to the left results in shifting of the intermediate products left. The shift of the multiplier to the right, on the other hand, prepares the next bit of the

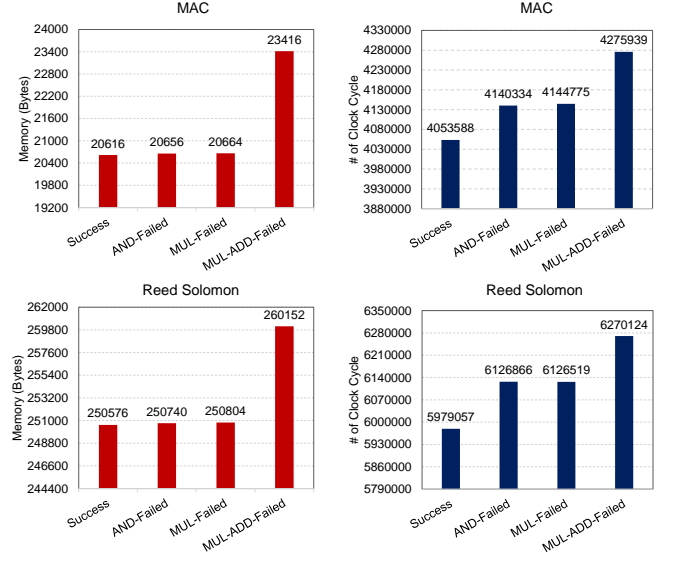


Fig. 9: The comparison of overhead in terms of memory usage and number of clock cycles in various scenarios in which different ALU operations fail.

multiplier to be tested in the next step. It is done in the loop until `t2` is zero. If the adder of the processor also fails, the RISC-V `add` instruction should be replaced with an implementation using bitwise operations for a ripple-carry adder, see Fig. 6(a). The ripple-carry algorithm splits the addition into two separate operations: Computing the sum without carrying using the `xor` operation and Computing the carry using the `and` operation. Instead of handling all bits in parallel, we propagate the carry bit iteratively. The loop guarantees that no carry is left unprocessed, effectively completing the addition.

In the second scenario, where a logical operation like AND fails, the `and` instruction can be replaced by using `not` and `or` instructions based on De Morgan's Laws:

$$(A \wedge B) = \neg(\neg A \vee \neg B)$$

According to the above equation, we can replace an AND operation by inverting the two operands, performing the OR operation on the inverted values, and then inverting the result. In the RISC-V assembly code, inversion is achieved by XORing a register with -1 (0xFFFFFFFF). The assembly language equivalent of an `and` instruction requires the steps of Fig. 6(b).

Fig. 9 illustrates the overhead in calculations in terms of memory and time resources. Specifically, the memory usage of MAC and Reed-Solomon benchmarks demonstrates that `and` and `mul` failures result in an increase of a few bytes of additional memory compared to the original code. In contrast, the failure of `add` and `mul` together require 2800 and 9576 bytes of additional memory in MAC and Reed Solomon, respectively. This significant difference is due to the increased number of `add` instructions in the assembly code. As expected, increasing the number of instructions also affects the number of clock cycles. As illustrated in Fig. 9, both the `mul` failure and the `and` failure take almost the same amount of time to complete

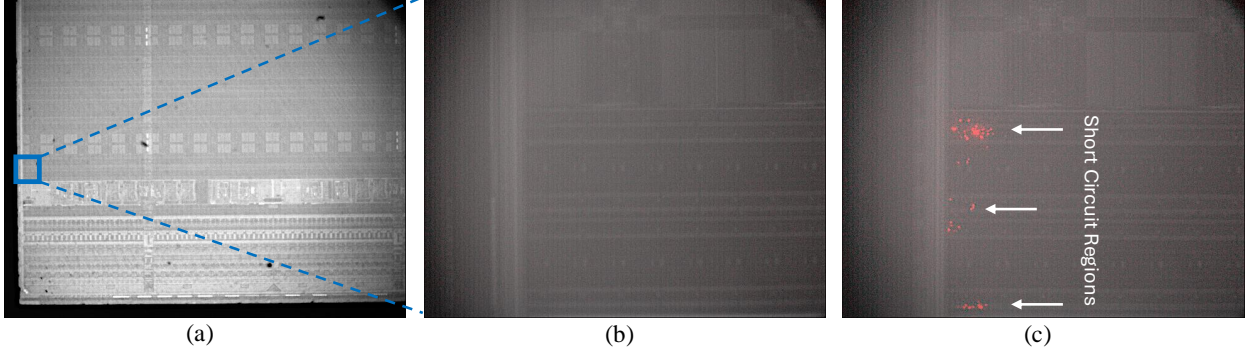


Fig. 10: (a) The reflectance image of one corner of the FPGA with a highlighted region of the injected fault. (b) The zoomed-in photon emission of healthy LUTs of the undamaged FPGA. (c) The zoomed-in photon emission of short-circuited LUTs of the damaged FPGA.

the total operation. However, while the memory requirements for these two scenarios are comparable to the memory usage of the original code, the time required to execute these tasks is substantially larger. This is because `mul` instruction is being replaced with loop-based instructions. Furthermore, the time required for execution of the recovery code when both `add` and `mul` have failed together is longer than for all other forms of the MAC function and Reed-Solomon due to the increased number of `add` instructions. However, because a single addition operation takes less time than a single multiplication, the time ratio between function variations does not directly correspond to memory utilization. Note that we only considered the time overhead imposed on instruction execution. The latency overhead associated with the binary translator or communication with the MCU depends on the onboard computer architecture and deserves a separate study.

C. Hardware Reconfiguration for Extensive HW Damage

As explored in the previous section, if radiation causes faults to the limited number of blocks of ALU, the computation can be detected and consequently recovered with software resynthesis. However, if extensive hardware damage occurs due to a powerful fault injection attack, where, for instance, the routing of the signals is permanently affected, a hardware recovery is necessary. Here, we present a significantly low-overhead reconfiguration technique to recover from major permanent hardware failures compared to existing core redundancy approaches. Fig. 10 demonstrates the images captured from the FPGA. In Fig. 10(a), a zoomed-out layout of the FPGA is depicted with a 2.5X lens. The highlighted blue box indicates the coordinates where the ALU is implemented in accordance with Fig. 5(a). We set the core to perform its normal operation. To emulate a wide-spread hard fault and evaluate our mitigation, we increase the peak current utilization of our NIR laser up to 2.5 A and perform a laser injection focused on the LUTs, realizing the ALU blocks on the IC. As the exposure of the laser is initiated, the core immediately crashes and fails to provide the anticipated output. On the other hand, we observed that the TDC sensory controller indicates a possible fault in our setup. We also observed that the permanent faults caused by the laser resulted in local short circuits on the FPGA. Fig. 10(c) shows

a high photon emission activity due to the short circuits in the damaged FPGA after the high-energy laser exposure. Fig. 10(b) depicts the same physical layout in an undamaged FPGA. Under normal circumstances, the damaged FPGA could not be utilized as an operational processor. However, based on the TDC sensory data captured right before the laser exposure, our framework is able to perform a PR and re-implement the entire core physically positioned near the TDCs in an undamaged location. By considering the reconfiguration time overhead of AMD/Xilinx’s ICAP, we verified the correctness of the core and successful recovery of the system in the new layout shown in Fig. 5(b). Note that reserved LUTs/FFs are required for hardware reconfiguration. The number of reserved LUTs/FFs must match the number of resources needed for the Rocket Core and ALU logic. Moreover, depending on the application and the availability of LUTs, TDCs can be duplicated across all locations in the FPGA. Such duplication assists in more precise localized fault detection while imposing negligible area overhead.

VII. CONCLUSION

In this paper, we presented a detection- and response-based countermeasure for processor chips in space applications to mitigate hard errors caused by adversarial and natural faults. We showed that FPGA-compatible TDC sensors can detect radiation and physical glitching attempts on the processor chip’s PDN and issue a warning for sanity checking. Upon detection of a fault, depending on the level of corruption, software resynthesis for substituting the corrupted instructions or hardware reconfiguration for relocating/rerouting the processor circuit will be carried out to recover the functionality. To assess the effectiveness of our approach, we performed laser radiation and fault injection on a soft RISC-V processor realized on an FPGA. Our results showed that the processor can continue its operation even with the presence of hard errors (e.g., short circuits) to accomplish the arithmetic and logical tasks.

ACKNOWLEDGMENT

This effort was sponsored by NSF Grant CNS-2150123.

REFERENCES

- [1] B. Bailey, "Cybersecurity protections for spacecraft: A threat based approach," *The Aerospace Corporation*, 2021.
- [2] L3Harris, "Hosted Payload Interface Unit (HPIU)," [Online] <https://www.l3harris.com/all-capabilities/hosted-payload-interface-unit-hpiu>, 2019.
- [3] W. Hennigan, "The Warning," *The New York Times*, 2024.
- [4] J. Markoff, "Tiny chips, big headaches," *International New York Times*, pp. NA–NA, 2022.
- [5] H. D. Dixit, S. Pendharkar, M. Beadon, C. Mason, T. Chakravarthy, B. Muthiah, and S. Sankar, "Silent data corruptions at scale," *arXiv preprint arXiv:2102.11245*, 2021.
- [6] P. H. Hochschild, P. Turner, J. C. Mogul, R. Govindaraju, P. Ranganathan, D. E. Culler, and A. Vahdat, "Cores that don't count," in *Proceedings of the Workshop on Hot Topics in Operating Systems*, 2021, pp. 9–16.
- [7] S. Moini, D. Kansagara, D. Holcomb, and R. Tessier, "Fault recovery from multi-tenant fpga voltage attacks," in *Proceedings of the Great Lakes Symposium on VLSI 2023*, 2023, pp. 557–562.
- [8] M. Paquette, B. Marquis, R. Bainbridge, and J. Chapman, "Visualizing electromagnetic fault injection with timing sensors," in *2021 IEEE Physical Assurance and Inspection of Electronics (PAINE)*. IEEE, 2021, pp. 1–8.
- [9] S. K. Monfared, K. Mitard, A. Cannon, D. Forte, and S. Tajik, "LaserEscape: Detecting and Mitigating Optical Probing Attacks," in *2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD)*, 2024.
- [10] B. Yuce, C. Deshpande, M. Ghodrati, A. Bendre, L. Nazhandali, and P. Schaumont, "A secure exception mode for fault-attack-resistant processing," *IEEE Transactions on Dependable and Secure Computing*, vol. 16, no. 3, pp. 388–401, 2018.
- [11] S. Jero, J. Furgala, M. A. Heller, B. Nahill, S. Mergendahl, and R. Skowrya, "Securing the Satellite Software Stack," in *SpaceSec*, 2024.
- [12] M. M. Alam, S. Tajik, F. Ganji, M. Tehranipoor, and D. Forte, "Ramjam: Remote temperature and voltage fault attack on fpgas using memory collisions," in *2019 Workshop on Fault Diagnosis and Tolerance in Cryptography (FDTC)*. IEEE, 2019, pp. 48–55.
- [13] J.-W. Han, M. Meyyappan, and J. Kim, "Single event hard error due to terrestrial radiation," in *2021 IEEE International Reliability Physics Symposium (IRPS)*. IEEE, 2021, pp. 1–6.
- [14] J. Gravelier, J.-M. Dutertre, Y. Teglia, and P. Loubet-Moundi, "High-speed ring oscillator based sensors for remote side-channel attacks on fpgas," in *2019 International conference on ReConfigurable computing and FPGAs (ReConFig)*. IEEE, 2019, pp. 1–8.
- [15] K. M. Zick, M. Srivastav, W. Zhang, and M. French, "Sensing nanosecond-scale voltage attacks and natural transients in fpgas," in *Proceedings of the ACM/SIGDA international symposium on Field programmable gate arrays*, 2013, pp. 101–104.
- [16] M. Zhao and G. E. Suh, "Fpga-based remote power side-channel attacks," in *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2018, pp. 229–244.
- [17] D. G. Mahmoud, O. Glamočanin, F. Regazzoni, and M. Stojilović, "Practical implementations of remote power side-channel and fault-injection attacks on multitenant fpgas," in *Security of FPGA-Accelerated Cloud Computing Environments*. Springer, 2023, pp. 101–135.
- [18] Xilinx. (2017) Zynq-7000 soc (z-7007s, z-7012s, z-7014s, z-7010, z-7015, and z7020): Dc and ac switching characteristics data sheet(ds187). <https://docs.xilinx.com/v/u/en-US/ds187-XC7Z010-XC7Z020-Data-Sheet>.
- [19] F. Mavaddat and B. Parhami, "Urisc: the ultimate reduced instruction set computer," *International Journal of Electrical Engineering Education*, vol. 25, no. 4, pp. 327–334, 1988.
- [20] P. J. Nürnberg, U. K. Wiil, and D. L. Hicks, "A grand unified theory for structural computing," in *Metainformatics: International Symposium, MIS 2003, Graz, Austria, September 17-20, 2003. Revised Papers*. Springer, 2004, pp. 1–16.
- [21] S. Dolan, "mov is Turing-complete," 2013. [Online]. Available: <https://drwho.virtadpt.net/files/mov.pdf>
- [22] J. Bangert, S. Bratus, R. Shapiro, and S. W. Smith, "The {Page-Fault} weird machine: Lessons in instruction-less computation," in *7th USENIX Workshop on Offensive Technologies (WOOT 13)*, 2013.
- [23] W. Reda, M. Canini, D. Kostić, and S. Peter, "{RDMA} is turing complete, we just did not know it yet!" in *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*, 2022, pp. 71–85.
- [24] D. Koch, *Partial reconfiguration on FPGAs: architectures, tools and applications*. Springer Science & Business Media, 2012, vol. 153.
- [25] C. M. Fuchs, N. M. Murillo, A. Plaat, E. van der Kouwe, and T. P. Stefanov, "Dynamic fault tolerance through resource pooling," in *2018 NASA/ESA Conference on Adaptive Hardware and Systems (AHS)*. IEEE, 2018, pp. 9–16.
- [26] Xilinx. (2023, Dec) Xilinx introduction to dynamic function exchange. <https://docs.xilinx.com/r/en-US/ug909-vivado-partial-reconfiguration/Introduction-to-Dynamic-Function-eXchange>.
- [27] K. Manev, J. Powell, K. Matas, and D. Koch, "byteman: A bitstream manipulation framework," in *2022 International Conference on Field-Programmable Technology (ICFPT)*. IEEE, 2022, pp. 1–9.
- [28] S. K. Monfared, D. Forte, and S. Tajik, "Randohm: Mitigating impedance side-channel attacks using randomized circuit configurations," in *2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD)*, 2024.
- [29] C. Swope, K. A. Bingen, M. Young, M. Chang, S. Songer, and J. Tammelleo, "Space threat assessment 2024," 2024.
- [30] D. R. Gnad, F. Oboril, and M. B. Tahoori, "Voltage drop-based fault attacks on fpgas using valid bitstreams," in *2017 27th International Conference on Field Programmable Logic and Applications (FPL)*. IEEE, 2017, pp. 1–7.
- [31] F. Brosser, E. Milh, V. Geijer, and P. Larsson-Edefors, "Assessing scrubbing techniques for xilinx sram-based fpgas in space applications," in *2014 International Conference on Field-Programmable Technology (FPT)*. IEEE, 2014, pp. 296–299.
- [32] C. Spensky, A. Machiry, N. Burrow, H. Okhravi, R. Housley, Z. Gu, H. Jamjoom, C. Kruegel, and G. Vigna, "Glitching demystified: analyzing control-flow-based glitching attacks and defenses," in *2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 2021, pp. 400–412.
- [33] M. De Kruijf, S. Nomura, and K. Sankaralingam, "Relax: An architectural framework for software recovery of hardware faults," *ACM SIGARCH Computer Architecture News*, vol. 38, no. 3, pp. 497–508, 2010.
- [34] X. Ni, E. Meneses, N. Jain, and L. V. Kalé, "Acr: Automatic checkpoint/restart for soft and hard error protection," in *Proceedings of the international conference on high performance computing, networking, storage and analysis*, 2013, pp. 1–12.
- [35] M. J. Wirthlin, A. M. Keller, C. McCloskey, P. Ridd, D. Lee, and J. Draper, "Seu mitigation and validation of the leon3 soft processor using triple modular redundancy for space processing," in *Proceedings of the 2016 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, 2016, pp. 205–214.
- [36] "Ibex risc-v reference guide: Exceptions and interrupts," https://ibex-core.readthedocs.io/en/latest/03_reference/exception_interrupts.html, accessed: 2024-12-10.
- [37] Xilinx. (2023, May) Xilinx 7 series fpgas configurable logic block. <https://www.eng.auburn.edu/~nelson/courses/elec4200/FPGA/ug4747SeriesCLB.pdf>.
- [38] Diligent, "Genesys 2," [Online] <https://diligent.com/reference/programmable-logic/genesys-2/start>, 2023.
- [39] AlphaNov, "Single Laser Fault Injection Microscope - S-LMS," [Online] <https://www.alphanov.com/en/products-services/single-laser-fault-injection>, 2024.
- [40] K. Asanovic, R. Avizienis, J. Bachrach, S. Beamer, D. Biancolin, C. Celio, H. Cook, D. Dabbelt, J. Hauser, A. Izraelevitz *et al.*, "The rocket chip generator," *EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2016-17*, vol. 4, pp. 6–2, 2016.
- [41] Z. F. Baruch, *Structure of computer systems*. UT Pres, 2002.