

RadPC: A Novel Single Event Upset Mitigation Strategy for FPGA-Based Space Computing

Christopher M. Major *

Montana State University, Bozeman, MT, 59715

Advisor: Brock J. LaMeres[†]

Montana State University, Bozeman, MT, 59715

This paper presents a computer architecture with the ability to respond to radiation-induced faults. The architecture presented advances the state-of-the art by implementing a redundant multiprocessor system on a commercial-off-the-shelf field programmable gate array. The system, called *RadPC*, can respond to radiation-induced faults and undergo repairs as necessary. A reliability analysis is conducted using Markov chain models to characterize the architecture's ability to continue performance despite faults. Further analysis regarding expected radiation exposure is provided for runtime context of the architecture.

Nomenclature

p = probability of failure

R = reliability

T = transition matrix

α = rate of context switch, s^{-1}

λ = radiation induced fault rate, s^{-1}

μ = repair rate, s^{-1}

I. Introduction

COMPUTING systems operating on the surface of the Earth run without concern of the harsh radiation environments found in space. However, computing systems for space science and space exploration missions do not have the same luxury. Space computers suffer from the unwanted effects of cosmic energy, unlike the billions of computers and embedded systems in use throughout the world that are effectively shielded from energized particles by the Earth's atmosphere and magnetic field. No such protection is guaranteed for computers outside of the planet's protective radiation shielding [1]. The demand for data processing in space science has become more apparent as the complexity

*Graduate Research Student, Electrical and Computer Engineering Department, 324 Norm Asbjornson Hall

[†]Professor, Electrical and Computer Engineering Department, 316-C Norm Asbjornson Hall. Member AIAA (Corresponding Author)

of exploration systems increases, thus the need for radiation-responsive features has become an area of concern.

There are two classes of space radiation effects on computing platforms: single event effects (SEE) and total ionizing doses (TID) [2]. A single event upset occurs when high-energy particles or heavy ions strike a CMOS device and cause unintended, logic-level transitions – an essentially instantaneous impact on performance. Such transitions can be divided into three sub-categories, depending on their effects within the fabric. A single event transient (SET) occurs when an energized particle changes the voltage of a logic line and therefore changes its logic value. When such a change is stored in a memory device, such as a latching circuit or a D-Flip Flop, a single event upset (SEU) has occurred. Each of these event types can often be corrected with a reset, but if an event causes such malfunction that further counteractive measures such as power cycling the entire system must be taken, the event is classified as a single event functional interrupt (SEFI).

A total ionizing dose occurs when lower-energy particles strike a CMOS device and cause physical damage. This is manifested as trapped charges left inside the insulating layers of a CMOS transistor lock the device into a perpetual on or off state, allowing for continuous and uninterrupted current draw. This results in large power draw by the part, interference with the device's tasks, and the eventual material degradation of the device over time. Where SEEs are momentary, TIDs are gradual and demonstrate cumulative effects that are measured as the amount of energy per unit of mass trapped in the material [3].

In modern integrated circuit design ($< 65\text{nm}$), TID is becoming more and more statistically unlikely to occur due to feature sizes with small process nodes [4]. While the smaller feature sizes are reducing the possibility of damage due to trapped charge, the probability of functionality interruption caused by high-energy particles is drastically increased. Thus, the need to mitigate the damaging effects from SEEs becomes of greater concern than TID in modern computing systems. This need increases in importance considering the trend toward shorter mission lifetimes (i.e., small satellites), where the computers are not in space long enough to accumulate TID damage.

In this paper, we propose and detail a radiation tolerant computer system called RadPC that is designed to respond to the effects of single event upsets. We describe an architecture that utilizes redundant processors configured in N-modular redundancy on a commercial-off-the-shelf (COTS) FPGA. Detection and repair of SEU-induced faults is accomplished through a comprehensive strategy including partial reconfiguration of impacted cores, error correction codes (ECC) for memory that cannot be partially reconfigured, and soft error mitigation monitoring for the configuration memory of the FPGA. We also provide a modeled analysis of the fault-tolerance of RadPC regarding times needed for reconfiguration of the system to flush out single event upsets.

II. Motivation

A. Existing TID Mitigation Strategies

As radiation-based faults are a great concern to spaceflight hardware, there currently exist various methods intended to prevent or mitigate their effects. A basic solution for TID mitigation is the concept of shielding, in which some sort of material serves as a boundary between cosmic radiation and the hardware at risk. Within many ground-based nuclear facilities, shielding is required for device and personnel protection and is implemented with little concern as to the geometry needed for protection.

In the case of spaceflight hardware, however, the mass and volume of a component can greatly impact the cost-effectiveness of the mission to detrimental levels as these metrics directly drive the costs for launch service. The thickness of a shield or the density of a material increases as the radiation dosage levels decrease, thus dramatically increasing the costs needed to transport said hardware into space [5]. Of additional concern is the decreased effectiveness of shielding as additional secondary streams of higher-energy particles can be generated as cosmic rays pass through the shielding hardware [6]. Thus, shielding may succeed in protecting a device from TIDs but is likely to fail to be ineffective in protection from SEEs when considering the cost and mass need to prevent secondary radiation.

Radiation-hardened semiconductor manufacturing modifications are another method of mitigating radiation-induced faults in system hardware [7]. One such technique, known as radiation-hardened-by-design (RHBD), mitigates TIDs by using non-standard layout techniques to reduce the probability of low-energy particles depositing charge within a device's insulating components. This is accomplished through techniques such as TMR implementation, enclosed-layout transistors (ELTs) and guard rings – elements designed to reroute the flow of low-energy particle charge through safer routes so as not to compromise the device's performance [8]. Additional components for error detection and correction in memory are often used in such designs to maintain data integrity [9].

Another technique to mitigate TID is known as radiation-hardened-by-process (RHBP). RHBP is a practice that uses non-standard materials during fabrication to that reduce the susceptibility probability of trapped charges. In doing so, defects are minimized in the substrate and oxide layers of the semiconductors, thus reducing the probability of radiation-induced electron/hole pairs in the material [10].

Systems designed with RHBD and RHBP techniques, however, require larger circuit surface area and therefore experience poorer power efficiency, speed, and performance than standard, commercially manufactured systems. Their dedicated manufacturing processes, combined with the relatively low market demand in comparison with commercial processes, greatly increase the costs of device production. Device performance is also greatly inhibited by these constraints, as RHBP-developed hardware tends to lag commercially developed equivalents by at least a decade [11].

B. Existing SEE Mitigation Strategies

To mitigate SEEs, one commonly used technique is triple modular redundancy (TMR). TMR is an approach in which a logic circuit is triplicated with identical implementations. Each circuit's output feeds into a voter device that determines the correct value of the original circuit through selection by majority. This strategy seeks to mitigate the effects of SEEs by providing the ability to continue operation in the presence of an SEE-induced fault [12]. TMR requires additional area and components for proper implementation, increasing costs, power, and design complexity. Additional overhead is needed for proper voter performance in the system. A block diagram demonstrating the basic structure of a TMR system is shown in Figure 1:




Fig. 1 Block Diagram of a TMR system

Memory in computing systems is uniquely susceptible to SEEs due to the high layout density and large area requirements on the substrate. Memory typically occupies the largest area on the substrate in a computing system. This limits the use of TMR and/or RHBD techniques prohibitive from an area-usage standpoint. The most common techniques to mitigate errors in memory used for data are error correction codes [9].

For memory that is used for program information or configuration data, a technique known as memory scrubbing is often used to mitigate the effects of SEEs. This is accomplished by using a device in hardware or a process in software to periodically iterate through the contents of a memory device and check their values against expected values. Blind scrubbing overwrites the memory cell regardless of whether an error occurred. Readback scrubbing detects if there is any discrepancy, then overwrites the faulty data with good data [13]. Scrubbing has been demonstrated as an effective technique for increasing system reliability and ensuring data integrity. Scrubbing is a common technique for mitigating SEUs in the configuration memory of FPGAs. Faults in the configuration memory of an FPGA represent faults that cannot be repaired using normal reset procedures, thus they are considered a SEFI.

C. Our Contribution

Our work contributes to the field of space computing by providing a cost-effective, radiation tolerant design using a commercial-off-the-shelf FPGA to actively respond to radiation-induced faults in the device. Instead of attempting to impede radiation-induced faults, our design seeks to mitigate their effects as a radiation-tolerant computer system. As advances in semiconductor manufacturing technology have led to decreased gate oxide thicknesses, reducing the possibility of trapped charges from TID at processes of 45 nm or smaller [4], our choice of a modern, commercial FPGA has a lower probability of sustaining damage from low-energy particles. Our strategy implements a redundant, multiprocessor system, a background configuration memory scrubber, and data memory error correction codes to mitigate the effects of SEEs. Our approach has the potential to increase the reliability of space computers while simultaneously reducing the cost, compared to existing systems that exploit RHBD/RHBP techniques.

III. System Design

The architecture of RadPC adapts the strategies of TMR and memory scrubbing by taking advantage of the configuration properties of FPGAs. An FPGA implements digital logic by arranging logic elements according to a design bitstream loaded into its configuration memory. When power is turned on, the device configures all logic, interconnects, inputs, and outputs according to this bitstream. This is known as full configuration. FPGAs provide the flexibility that any design can be implemented and synthesized into an FPGA bitstream. Originally, this capability afforded a unique environment for prototyping. Recently FPGA performance has reached a level that makes implementation in a final system is practical. In our work, we exploit the reconfigurability of FPGAs to dynamically repair and recover from faults.

If correctly enabled by the designer, an FPGA can select a pre-defined subset of the hardware and reconfigure it with a different design. This is known as partial reconfiguration (PR). In PR, a partial bitstream is loaded into a select section of configuration memory to drive this change in logic. This section is often referred to as a “PR region” of the FPGA fabric. Two methods of PR are often used in FPGA design, depending on the state of the device during the process. When the device is shut down or inactive and a partial bitstream is loaded, static partial reconfiguration has occurred. When the device is currently performing its designated task without power cycling or suspending activity outside the PR region, active partial reconfiguration has occurred [14]. Partial reconfiguration has a unique role in our FPGA-based SEE mitigation strategy. When a fault occurs in either the foreground circuitry of the FPGA or in the configuration memory region corresponding to the foreground circuitry, PR will restore that region to its original state. This effectively repairs the fault caused by the SEE [15, 16]. A design approach can be used that divides the system into various PR regions and incorporates PR as part of the fault-mitigation strategy. This does lead to synchronization issues that must be addressed, which will be covered later.

Our design employs a variety of fault-mitigation strategies based on a redundant multiprocessor system implemented within the FPGA fabric. Figure 2, below, demonstrates the general layout of this design.




Fig. 2 Block Diagram of RadPC Architecture

A. Fault Mitigation Strategy

Four parallel processors, denoted as “tiles,” are designated as PR regions and run the identical software. This system extends the capability of TMR to continue operation in the presence of a fault. TMR can withstand one fault in the triad and still produce the intended output. However, when the fault occurs due to an SEU or SEFI, the system needs to stop foreground operation and restore the faulted system. In FPGA-based computer systems, this typically means either reconfiguring the entire FPGA or scrubbing the entire configuration memory array. The significant downtime associated with repairing the faulted member of the triad reduces the reliability of the system for two reasons. First, the time associated with full reconfiguration or full scrubbing of the computer system dramatically reduces the availability of the system. Second, the system is still vulnerable to radiation strikes while performing the repair. If a second strike occurs during repair, the system is put into a state that cannot be recovered from. In our prior work, we discovered that increasing the number of redundant members in the system (N) reduces the possibility of a subsequent radiation strike putting the system into an unrecoverable state [17]. Our prior work also showed that the most gain in reliability was achieved with a 4-member system. Adding more members beyond four did not result in any significant reliability gains because the additional circuit area that was added was also susceptible to radiation strikes so there was diminishing returns. Thus, a 4-modular redundant system is optimal.

Additionally, the use of PR to repair a faulted tile provides a much faster recovery process than a full FPGA reconfiguration. As tasks are completed by each tile, any output results are fed into a voter circuit and analyzed for discrepancy by majority vote. If a tile shows an incorrect value, it is partially reconfigured to flush out the SEU(s) in its PR region. When this occurs, the remaining tiles continue to proceed with the next software task.

After the faulted tile is partially reconfigured, it is ready to be reintroduced into the voting system. In order to synchronize the repaired tile with the other three, a software checkpoint system is used. A checkpoint bus is implemented

between the four tiles that indicates if another tile has been faulted. Code is entered into the software at desired points that will allow the processor to stop execution of the main program and wait for the partially reconfigured tile to synchronize. The locations of the software checkpoints are at the discretion of the developer in order to provide the ability to prioritize critical code execution over synchronizing the repaired PR tile. For mission critical software, the synchronization can wait until there is a break in important execution.

The voter implements additional functionality to serve as the arbitrator of the checkpointing to ensure data integrity and synchronization. The voter observes a register known as the “checkpoint” register, where each tile will set a flag when it has finished executing a given portion of the program. When all four tiles have set their checkpoint flags, the voter receives output data from each of the tiles and checks their values. If these values are all equal, the voter will signal the tiles to continue to the next stage of the program. If these values are not equal, the voter will signal the memory correction component to activate and correct any erroneous memory cells before letting the tiles resume the program.

Each tile accesses an on-FPGA memory block for its data memory. On an FPGA, the user memory is implemented with SRAM technology, which is not impacted by a full or partial reconfiguration. This is advantageous from an operational standpoint because the data is not altered during a PR; however, additional fault mitigation must be included in the system to protect the data memory. In our system, each data memory block employs an ECC system to detect and correct faults that may have occurred in the SRAM cells. Each memory device consists of a data encoder, a memory device, and a data decoder. When a tile writes to this memory device, the data is encoded from an 8-bit format to a 12-bit format. This is accomplished using a Hamming encoder device to add four error correction code bits, which are used in the decoding process to evaluate possible data corruption. When a tile reads from the memory device, the data from the memory device is decoded from a 12-bit format into an 8-bit format. This is accomplished using a Hamming decoder device to check the consistency between the error correction bits and the original data, then attempts to correct any errors present before letting the tile GPIO read the value.

While the ECCs handle correcting any radiation-induced faults in the data memory, they do not have the ability to synchronize the contents with the other three memory blocks in the event of a tile PR. To handle this, a memory correction system is implemented. The encoded data from each tile is stored in a dual port read/write memory unit. One port is dedicated to tile use; the other is used for the separate memory correction component. The memory correction component accesses a tile’s memory block’s second port, bypassing the Hamming encoder and decoder to directly assess the memory contents. It then evaluates the contents of all four memory units simultaneously, comparing them to see which value is inconsistent with the others. If a value is inconsistent, it will be rewritten with the majority value of its neighboring units at that address.

For SEUs that occur outside the FPGA fabric designating a tile, a soft error mitigation (SEM) controller is used as the configuration readback memory scrubber [18]. This is a Xilinx-provided IP, connected to the ICAP primitive and a UART communication port, that uses error correction codes within the Artix-7 configuration memory to check for

Table 1 Summary of Fault Mitigation strategies of RadPC

Fault Condition	Observed By	Action Taken	Fault Severity	Fault Possibility
SEU in Tile foreground circuitry	Voter	PR-effected Tile	Medium	Medium
SEU in non-Tile foreground circuitry	Voter	Full FPGA reconfiguration	High	Low
Single-Bit Error in on-FPGA Data Memory	ECC	Automatic Data Recovery by ECC encoder/decoder	Low	Medium
Multi-Bit Error in on-FPGA Data Memory	Voter	Overwrite corrupt data memory with values from other healthy data memory	Medium	Low
Single-Bit Error in Configuration Memory	SEM	Automatic Correction	Low	High
Double-Bit Adjacent Error in Configuration Memory	SEM	Automatic Correction	Low	Medium
Multi-Bit, Non-Adjacent Errors in Configuration Memory	SEM	Full FPGA reconfiguration	High	Low
Note 1: A fault severity of LOW indicates that the architecture has a strategy to recover from the fault without any noticeable impact on performance. Note 2: A fault severity of MEDIUM indicates that the architecture has a strategy to recover from the fault but will incur a small-time penalty to complete the repair. Note 3: A fault severity of HIGH indicates that the architecture cannot recover from the fault and that the entire FPGA needs to be reconfigured. Existing mission data in the FPGA may be lost.				

single-bit and double-bit adjacent upsets. The controller runs perpetually, independent of the tiles, memory, and voter, to monitor and scrub errors in background operation. When a partial reconfiguration on a tile is performed, the SEM controller is temporarily deactivated to avoid error correction conflicts. Upon completion of the partial reconfiguration, a software reset command is sent to the SEM controller and its monitoring process continues.

To protect the configuration memory itself from faults, a readback memory scrubber can be used. Using the memory's error correction codes and frame address register (FAR), a scrubber can iterate through the memory frame by frame and check the contents for single-bit or adjacent double-bit errors [19]. In the event of an error, the scrubber replaces the faulty bits in memory and resumes operation. If a scrubber is used alongside a PR strategy, an arbitration system is necessary to prevent interference between the two systems. The Xilinx FPGA provides a built-in configuration memory repair system called a soft error mitigation (SEM) block. This is implemented on our system.

B. Unrecoverable Faults

In any system, there will be faults from which the system cannot be recovered. In our system, multiple faults occurring simultaneously can prevent successful recovery. Our design is based on the premise that fault recovery is implemented on high probability faults and lower probability faults represent infrequency events that require either a full FPGA configuration or a full system power cycle.

C. Fault Mitigation Summary

Table 1 summarizes the range of possible faults, fault-mitigation strategies, and probability of risk.

D. Implementation

This architecture was prototyped on a Basys 3 development board, produced by Digilent. This board features an Artix-7 XC7A35T CGP326-1, a low-power, low-cost device available for purchase from Xilinx [20]. The hardware design was developed using the Vivado 2019.2 design suite and the software to be implemented for demonstration was developed using Xilinx's Vitis development suite.

Each tile consists of a Xilinx MicroBlaze softcore processor, 8 KB of program memory, four GPIO devices with two

channels each, and a processor reset system. A single program runs on the MicroBlaze, utilizing the GPIO to access an input register, a memory device, the Voter system, and the Checkpoint register. All these components are contained in a Vivado constraints file and wrapped in a hardware wrapper to allow for instantiation in the top-level VHDL file. The tiles are denoted as TILE_00, TILE_01, TILE_02, and TILE_03.

Each tile has a corresponding partial bitstream which can be used to reset the FPGA fabric in the predefined tile region and bring the processor and its connected components back online. The memory, voter, checkpoint register, and all other external components are not affected by this reconfiguration. Though a reconfigured tile may lag the other tiles in its program's progress, the checkpoint system ensures that the other tiles will not advance in the program until the affected tile can catch up.

IV. Characterization

The reliability of this system depends on the amount of time that the various aspects of recovery take in addition to the predicted rate of incoming faults. This drives the system's availability and mean time to failure (MTTF), but are also dependent on the environment that the computer resides in. The first step is determining the time of the repair strategies in this computer.

Measurements regarding the time needed to perform full configuration and partial reconfigurations of the device were made by measuring the DONE pin of the Basys 3 board using an oscilloscope. According to Xilinx's 7-Series FPGA configuration guides, the DONE pin asserts when a full configuration is completed or when a partial reconfiguration is completed. Thus, this signal was used to evaluate the speeds of both methods.

A full reconfiguration of the system, using the top-level bitstream, takes an average of 1.24 seconds to complete. This test was conducted using the JTAG configuration port to load a 17,536,096 bit configuration bitstream into the FPGA. According to Vivado's bitstream generation reports, 2,793,730 of the 14,663,584 bits of configuration memory, or 19.05%, are classified as "essential bits" that are prioritized in the SEM controller's operation.

The partial reconfiguration of a tile will alter its completion time, depending on the tile to be reconfigured. This is due to the floor planning conducted to allow for the generation of a partial bitstream for each tile, as the XC7A35T has limited fabric space for the positioning of "PBlocks" assigned to various components of the design. Thus, all four tiles differ in position, number of logic elements per PBlock, size, and accessibility to clock and reset lines. Partial reconfiguration of a RadPC tile, therefore, ranges between averages of 217 milliseconds and 434 milliseconds. Specifically, Tile 0 takes an average of 217 milliseconds to be reconfigured, Tile 1 takes an average of 434 milliseconds, Tile 2 takes an average of 404 milliseconds, and Tile 3 takes an average of 217 milliseconds.

Faster configuration times can be obtained by using a different programming interface to the FPGA such as slave serial. For our prototype, the slave serial bus was not connected on the Basys 3 board so it could not be empirically measured. However, an assumption can be made that the ratio of PR to full configuration will remain the same when

moving from JTAG to slave serial. This ratio for the largest PR is 35% (i.e., 434 [ms] / 1.24 [s]). Calculating the amount of time that a slave serial configuration takes, it was determined that a full configuration will take 1.101 seconds and a worst-case partial reconfiguration will take 418 milliseconds [21, 22].

Per Vivado's utilization reports and the available resources on the Artix-7 XC7A35T chip, 12882 Slice Registers out of 41600 (30.97%) and 12068 Slice LUTs out of 20800 (58.02%) are used in the design. 23 Block RAM Tiles out of the available 50 (46%) are likewise used. Vivado's timing summary reports the worst slack for the setup time in the design is 23.790 ns and the worst pulse width slack time is 15.250 ns.

Vivado's power report estimates a total on-chip power of 0.102 watts, with dynamic power accounting for 0.029 watts and device static accounting for the other 0.073 watts. A junction temperature of 25.5 °C was specified in the report. This does not account for I/O power on the Basys board itself, merely representing an estimate of the FPGA device itself.

V. Reliability Analysis

A. Markov Chain Models

To estimate the reliability of the proposed architecture, Markov chain models are used [17]. Markov models are used to model a system using state diagrams, linked together by probabilities describing the likelihood of a transition from one state to another. As such, they assume that, in their most basic form, a transition to another state relies only on the current state of the system [23]. This principle makes them a useful tool for analyzing system reliability. The Markov model will transition to a new state upon a fault in the computer. Since not every radiation strike will cause a fault in the computer system, the fault rate fed into the Markov model is scaled-down to a value that considers the essential bits on the FPGA (i.e., those that represent implemented circuitry) and a de-rating for the type of system implement (i.e., 30% for a microprocessor) [24].

In a Markov chain model, a system is represented by a transition matrix of size (m, n) , where each entry represents the probability of transitioning from state m to state n . In the example of a simple two-state system with states S_0 and S_1 , the transitions can be represented with a 2x2 matrix filled with a probability of each transition in each entry, where λ represents the fault rate of the system. Figure 3, below, demonstrates this model as a state diagram. This model represents a simplex computer system where state S_0 represents an operational system and state S_1 represents a faulted system.

For a general state diagram with s states, a transition matrix T of size $s \times s$ can be created to represent the probability of moving from one state m to another state n - as shown in Eq. 1




Fig. 3 Simple 2-State Markov Chain Model diagram

$$T(m, n) = \begin{bmatrix} t_{1,1} & \dots & t_{1,n} \\ \dots & \dots & \dots \\ t_{m,1} & \dots & t_{m,n} \end{bmatrix} \quad (1)$$

To calculate the probability of being in a state S given a time step k , the following equation must be used.

$$p_s(t = k\Delta t) = \begin{bmatrix} p_0(0) & \dots & p_n(0) \end{bmatrix} \cdot \begin{bmatrix} t_{1,1} & \dots & t_{1,n} \\ \dots & \dots & \dots \\ t_{m,1} & \dots & t_{m,n} \end{bmatrix}^k \quad (2)$$

In the case of the two-state diagram in Figure 3, assume that S_0 represents a state of healthy device operation and S_1 represents a state of faulty or failed device operation. The transition matrix T of this model is:

$$T(m, n) = \begin{bmatrix} 1 - \lambda\Delta t & \lambda\Delta t \\ 0 & 1 \end{bmatrix} \quad (3)$$

The probability of residing in state S after a time step k , p_s , is given by:

$$p_s(t = k\Delta t) = \begin{bmatrix} p_0(0), p_1(0) \end{bmatrix} \cdot \begin{bmatrix} 1 - \lambda\Delta t & \lambda\Delta t \\ 0 & 1 \end{bmatrix}^k \quad (4)$$

The reliability is calculated by the following:

$$R(t) = 1 - p_s(t = k\Delta t) \quad (5)$$

As the performance analysis of RadPC hinges on its reliability, the mean time to failure must be considered. This is defined as the point in time t when the reliability of the system $R(t) = 0.5$.

In the case of RadPC, there are more states to reflect the various capabilities and safeguards in the system. For example, a TMR system can be modeled with three states [17]. State S_0 represents an operational system where none of the three tiles have radiation-induced faults. State S_1 represents the situation where one tile of the triad has been faulted. Note that the system is still operational because two of the three tiles are operational. State S_2 represents a faulted system because two of the three tiles have been faulted and the system can no longer determine the correct output. Figure 4 demonstrates the state diagram used to represent a TMR system as a Markov chain model.




Fig. 4 Markov chain model state diagram for a TMR system

The transition matrix of this system is as follows:

$$T(m, n) = \begin{bmatrix} 1 - 3\lambda\Delta t & 3\lambda\Delta t & 0 \\ 0 & 1 - 2\lambda\Delta t & 2\lambda\Delta t \\ 0 & 0 & 1 \end{bmatrix} \quad (6)$$

The probability of an operation TMR system is given as the product of their individual reliabilities, as follows:

$$R(t) = R^3(t) + 3R^2(t) \cdot (1 - R(t)) \quad (7)$$

$$R(t) = 3R^2(t) - R^3(t) \quad (8)$$

A TMR-based system alone, however, has no option of returning to the initial, fully operational state and therefore reduces the reliability of the system. Thus, the option to repair a tile through partial reconfiguration or soft error mitigation must be represented in the state diagram. This can be referred to as a TMR + Spare system. The repair rate of the tile is represented as a parameter μ , which can be determined as the worst-case number of repairs per second. Based on the performance of the Basys 3 board, the repair rate was found to be $\mu = 437 [ms]$. Figure 5 demonstrates the new state diagram, including the PR/SEM recovery methods:

In order to further increase the reliability of the system, a fourth tile is used in RadPC [17]. This results in a 4-modular redundant (4MR) system, where the additional tile allows two faults to occur in the system without faulting the system. An assumption is made that in the event of two faults, the two damaged tiles will not produce the same, yet




Fig. 5 Markov chain model state diagram for a TMR + PR + SEM system

inaccurate, output. As before, S_0 indicates a healthy and fully operational state. If an operating tile is damaged, the system transitions to S_1 and must undergo repairs. The faulty tile is disregarded in the vote process until it is brought back into operation. This is represented in S_2 . If the faulty tile cannot be repaired, the system can proceed in S_3 until repairs are made. Only when a system has experienced two, unrecoverable faults, is the transition to S_4 made, in which system failure is experienced entirely. Figure 6 demonstrates this behavior as a new state diagram.




Fig. 6 Markov chain model state diagram for a TMR + PR + SEM + Spare Tile System

The difference between a 4MR system and a TMR + Spare system is in the treatment of the "spare tile." In a 4MR system, the spare actively runs the same program as its neighbors and is considered a valid vote in the overall architecture. When an error occurs in any of the tiles, PR is conducted immediately without swapping in new components - as in the previous TMR + Spare system. The voter will then track any faults in the running tiles and perform the repairs as necessary, including memory scrubbing if necessary.

A comparison between architectures was performed in order to verify that the RadPC architecture using four tiles with repair is significantly more reliable than simplex, TMR, and TMR with repair approaches. The graph in Figure 7 demonstrates the exponential reliability curve for the different system architectures. For this comparison, a value of $\mu = 437$ [ms] was used and represents the slowest partial reconfiguration time (i.e., Tile 1). A fault rate of $\lambda = 0.001$ [SEU/ μ s] was used in all systems and was chosen arbitrarily.

The plots indicate the increase in reliability when additional features are implemented alongside a basic TMR design. With a singular, or "simplex" system, or an unprotected TMR system, the reliability decays at an exponential rate and provides little to no assurance of system longevity. With a TMR system protected by a scrubbing mechanism, be that the SEM or PR, the reliability drops at a far slower linear rate. Including a spare tile to be included in the TMR + Spare




Fig. 7 Exponential reliability curve for 4-Tile RadPC system with repair capability

system dramatically decreases the rate of reliability decrease and therefore provides the best system longevity of all possible combinations. Thus, it can be seen that the 4-modular redundant system provides the best reliability of all aforementioned systems. It is left to the discretion of a mission planner to consider the range of radiation environments in which RadPC would be exposed, and to assess whether an environment's expected fault rate would warrant full reconfiguration of the device. [LIST EXAMPLES]

B. CREME96 Analysis

To characterize the radiation environments in which RadPC would be exposed, a numerical modeling tool known as the Cosmic Ray Effects on Micro-Electronics code (CREME96) was used [25]. Given the dimensions of the Artix-7 XC7A35T chip and the essential bit count provided by the Vivado software suite, a series of estimates were given to predict the levels of direct ionization-induced SEEs the system would experience. Table 8, below, displays the results of the ISS-environment and lunar-environment analyses. Accounting for the essential bits in the system and the de-rating for a microprocessor, these fault rates are scaled down by a factor of 0.3. The resulting figures are shown in Table 9.

Generally, the average and peak fault rates are relatively small and do not present a significant risk to the system. A worse-case scenario analysis, however, proves more useful in evaluating the limits of the system. Thus, a script was written in MATLAB to plot the mean-time-to-failure as opposed to fault rate, shown in Figure 10.

The percentage increase in improvement of each of these architectures, in comparison to the simplex architecture, is shown in Table 11.

The plot and table demonstrate that the architecture that best retains its reliability as the fault rate increases is

Essential Bits Analysis (2793730)		Total Bits Analysis (14663584)	
Lunar Environment, Average Conditions	device/day	Lunar Environment, Average Conditions	device/day
Average	1.969	Average	10.336
Peak	5.463	Peak	28.676
Lunar Environment, Inside Earth's Magnetosphere Worst Case Scenarios	device/day	Lunar Environment, Inside Earth's Magnetosphere Worst Case Scenarios	device/day
Worst Week	166696.000	Worst Week	874944.000
Worst Day	910118.000	Worst Day	4776980.000
Worst 5 Minutes	312909000.000	Worst 5 Minutes	1642380000.000
Lunar Environment, Outside Earth's Magnetosphere Worst Case Scenarios	device/day	Lunar Environment, Outside Earth's Magnetosphere Worst Case Scenarios	device/day
Worst Week	166712.000	Worst Week	875200.000
Worst Day	910207.000	Worst Day	4778370.000
Worst 5 Minutes	3473810.000	Worst 5 Minutes	18236700.000
ISS Environment, Stormy Solar Conditions	device/day	ISS Environment, Stormy Solar Conditions	device/day
Average	1179.700	Average	6192.300
Peak	2232.830	Peak	11721.700
ISS Environment, Quiet Solar Conditions	device/day	ISS Environment, Quiet Solar Conditions	device/day
Average	29.046	Average	152.470
Peak	82.695	Peak	434.109
ISS Environment, Worst Case Scenarios	device/day	ISS Environment, Worst Case Scenarios	device/day
Worst Week	9939.990	Worst Week	52182.500
Worst Day	54377.500	Worst Day	285468.000
Worst 5 Minutes	207490.000	Worst 5 Minutes	1089270.000

Fig. 8 CREME96 modeling of ISS-orbit and lunar-orbit radiation environments

the 4-modular redundant system, given the repair rate $\mu = 437$ [ms]. Even in the extreme cases of the ISS-orbit and lunar-orbit environments, it performs significantly better than the simplex and TMR architectures in maintaining reliability. As previous iterations of RadPC have experimented with varying numbers of tiles and concluded that the

Essential Bits Analysis (2793730)		Total Bits Analysis (14663584)	
Lunar Environment, Average Conditions	device/day	Lunar Environment, Average Conditions	device/day
Average	0.591	Average	3.101
Peak	1.639	Peak	8.603
Lunar Environment, Inside Earth's Magnetosphere Worst Case Scenarios	device/day	Lunar Environment, Inside Earth's Magnetosphere Worst Case Scenarios	device/day
Worst Week	50008.800	Worst Week	262483.200
Worst Day	273035.400	Worst Day	1433094.000
Worst 5 Minutes	93872700.000	Worst 5 Minutes	492714000.000
Lunar Environment, Outside Earth's Magnetosphere Worst Case Scenarios	device/day	Lunar Environment, Outside Earth's Magnetosphere Worst Case Scenarios	device/day
Worst Week	50013.600	Worst Week	262560.000
Worst Day	273062.100	Worst Day	1433511.000
Worst 5 Minutes	1042143.000	Worst 5 Minutes	5471010.000
ISS Environment, Stormy Solar Conditions	device/day	ISS Environment, Stormy Solar Conditions	device/day
Average	353.910	Average	1857.690
Peak	669.849	Peak	3516.510
ISS Environment, Quiet Solar Conditions	device/day	ISS Environment, Quiet Solar Conditions	device/day
Average	8.714	Average	45.741
Peak	24.809	Peak	130.233
ISS Environment, Worst Case Scenarios	device/day	ISS Environment, Worst Case Scenarios	device/day
Worst Week	2981.997	Worst Week	15654.750
Worst Day	16313.250	Worst Day	85640.400
Worst 5 Minutes	62247.000	Worst 5 Minutes	326781.000

Fig. 9 CREME96 modeling of ISS-orbit and lunar-orbit radiation environments, 30% microprocessor de-rating

gains of additional redundant tiles diminish shortly after adding one. Thus, it can be concluded that the 4-modular redundant system is the most optimal architecture of those discussed in this paper.




Fig. 10 MMTF versus fault rate for varying architectures

Architecture	% Improvement Over Simplex System	
	ISS-Orbit Worst Case	Lunar-Orbit Worst Case
TMR	4.7%	4.9%
TMR + Repair	327%	117%
4MR	1556%	327%

Fig. 11 MTTF improvement increase of each architecture compared to the Simplex system

VI. Discussion

Several considerations are to be taken into account regarding the performance of the RadPC system. As the software synchronization checkpoints play a vital role in the timing of partial reconfigurations, two approaches to planning checkpoints are to be considered by software developers. The first approach would be to space checkpoints over long periods of time, allowing for multiple partial reconfigurations to complete before checking tile outputs. The second approach would be to shorten the time between checkpoints, thus reducing the amount of time for a faulted tile to continue without repair. Though the former approach would perform well in low-fault conditions, it is recommended that the latter approach be taken so as to quickly assess and correct faults that may occur in a tile, instead of letting errors accumulate between spaced checkpoints.

The time needed to complete a full or a partial reconfiguration depends on the frequency of the external monitoring device performing the bitstream upload. Thus, the μ value of a slave serial configuration can be found by running Vivado's configuration time calculator tool. The slave serial configuration time previously given assumed a 16 MHz serial link from a microcontroller; however, this does not guarantee the fastest possible configuration. Given a controller configuration frequency of 100 MHz, the configuration time drops to 39 [ms] minimum and μ drops to 0.0039. As can be seen in Figure 12, the mean time to failure increases as μ decreases.




Fig. 12 MMTF versus fault rate for varying architectures, best-performing μ for 4MR system

Of additional consideration is the improvement in reliability between the 4MR system with best-performing μ and the simplex system. This architecture exhibits an improvement percentage of upwards to 3567% of the simplex system in ISS-orbit and upwards to 9359% of the simplex system in lunar orbit.

Further analysis of this system will lead to the consideration of single points of failure, specifically the voter and the data memory scrubber. Though these areas of the device are subject to SEUs, like any other component in the FPGA, their relatively small footprints reduce their probability of experiencing damage from faults. The SEM will protect the configuration memory of these elements from damage, but otherwise, an error in their circuitry may require a full reconfiguration. This could be mitigated through implementing these elements on an external, radiation-hardened component, but such a strategy would defeat the purpose of developing a commercial, off-the-shelf design. Rather, the recommended approach would be to establish partially reconfigurable sections of the device dedicated to each component and to employ PR if damage is observed.

Finally, the possibility of tile de-synchronization by a few clock cycles, even after checkpoint synchronization occurs, can become a point of concern in device operation. The choice to attempt synchronization again to constrain these times

remains valid, or to create a voter that considers outputs across a wider range of clock cycles to account for errors.

VII. Conclusion

The computer architecture presented in this paper, named RadPC, demonstrates the ability to respond to faults induced by space radiation. As many current radiation-fault mitigation strategies in space mission computing systems have proven to be infeasible due to cost and complexity, alternate methods to achieve radiation tolerance are necessary. In addressing the impact of SEEs within the circuitry of FPGA-based computers, several methods for mitigating their impact on performance were discussed and implemented within the system. A multiprocessor system inspired and adapted from the common method of TMR was implemented, with the ability to repair impacted tiles and restore functionality. This capability is achieved using partial reconfiguration and soft error mitigation to account for SEE-affected portions of the FPGA fabric. Redundant memory devices, checkpoint flags, and a voter system ensure processor tile synchronization. This system was run through a Markov chain model to characterize its ability to maintain performance despite radiation-induced faults. With expected radiation strikes provided by a model output through CREME96, the design demonstrates the ability to provide a feasible, low-cost, low-power, commercial-off-the-shelf architecture to address radiation-induced faults reliably and expediently.

VIII. Acknowledgements

This material is based upon work supported by the National Aeronautics and Space Administration under Grant/Contract/Agreement No.(s) 80MSFC20C0032, 80NSSC18P2129, 80NSSC20K0107, 80NSSC19M0037, 80NSSC18K1671, 80NSSC17M0037, NNX16AI75A, NNX14AL03A, NNX13AR03A, NNX12AM50G, NNX10AN91A, and NNX10AN32A.

References

- [1] Barth, J. L., Dyer, C. S., and Stassinopoulos, E. G., "Space, atmospheric, and terrestrial radiation environments," *IEEE Transactions on Nuclear Science*, Vol. 50, No. 3, 2003, pp. 466–482.
- [2] Claeys, C., *Radiation Effects In Advanced Semiconductor Materials And Devices*, Springer, 2010.
- [3] Adams, A., and Holmes-Siedle, L., *Handbook Of Radiation Effects*, Oxford University Press, 2004.
- [4] Barnaby, H. J., "Total-Ionizing-Dose Effects in Modern CMOS Technologies," *IEEE Transactions on Nuclear Science*, Vol. 53, No. 6, 2006, pp. 3103–3121.
- [5] Uzel, R., and Özyildirim, A., "A study on the local shielding protection of electronic components in space radiation environment," *2017 8th International Conference on Recent Advances in Space Technologies (RAST)*, 2017, pp. 295–299.

- [6] Cliff, R. A., Danchenko, V., Stassinopoulos, E. G., Sing, M., Brucker, G. J., and Ohanian, R. S., "Prediction and Measurement of Radiation Damage to CMOS Devices on Board Spacecraft," *IEEE Transactions on Nuclear Science*, Vol. 23, No. 6, 1976, pp. 1781–1788.
- [7] Gambes, J. W., and Maki, G. K., "Rad-tolerant flight VLSI from commercial foundries," *Proceedings of the 39th Midwest Symposium on Circuits and Systems*, Vol. 3, 1996, pp. 1227–1230 vol.3.
- [8] Anelli, G., Campbell, M., Delmastro, M., Faccio, F., Floria, S., Giraldo, A., Heijne, E., Jarron, P., Kloukinas, K., Marchioro, A., Moreira, P., and Snoeys, W., "Radiation tolerant VLSI circuits in standard deep submicron CMOS technologies for the LHC experiments: practical design aspects," *IEEE Transactions on Nuclear Science*, Vol. 46, No. 6, 1999, pp. 1690–1696.
- [9] Yu-Lam, D., Lan, J., McMurchie, L., and Sechen, C., "SEE-Hardened-by-Design Area-Efficient SRAMs," *2005 IEEE Aerospace Conference*, 2005, pp. 1–7.
- [10] Makiyara, A., Midorikawa, M., Yamaguchi, T., Iide, Y., Yokose, T., Tsuchiya, Y., Arimitsu, T., Asai, H., Shindou, H., Kuboyama, S., and Matsuda, S., "Hardness-by-design approach for 0.15 μm fully depleted CMOS/SOI digital logic devices with enhanced SEU/SET immunity," *IEEE Transactions on Nuclear Science*, Vol. 52, No. 6, 2005, pp. 2524–2530.
- [11] Keys, A., Adams, J., Frazier, D., Patrick, M., Watson, M., Johnson, M., Cressler, J., and Kolawa, E., "Developments in Radiation-Hardened Electronics Applicable to the Vision for Space Exploration," *AIAA SPACE 2007 Conference & Exposition*, 2012. <https://doi.org/10.2514/6.2007-6269>.
- [12] Sterpone, L., and Violante, M., "Analysis of the robustness of the TMR architecture in SRAM-based FPGAs," *IEEE Transactions on Nuclear Science*, Vol. 52, No. 5, 2005, pp. 1545–1549.
- [13] Garvie, M., and Thompson, A., "Scrubbing away transients and jiggling around the permanent: long survival of FPGA systems through evolutionary self-repair," *Proceedings. 10th IEEE International On-Line Testing Symposium*, 2004, pp. 155–160.
- [14] Kao, C., "Benefits of partial reconfiguration," *Xcell journal*, Vol. 55, 2005, pp. 65–67.
- [15] Gauer, C., LaMeres, B. J., and Racek, D., "Spatial avoidance of hardware faults using FPGA partial reconfiguration of tile-based soft processors," *2010 IEEE Aerospace Conference*, 2010, pp. 1–11.
- [16] Julien, C. R., LaMeres, B. J., and Weber, R. J., "An FPGA-based radiation tolerant SmallSat Computer System," *2017 IEEE Aerospace Conference*, 2017, pp. 1–13.
- [17] Hogan, J. A., Weber, R. J., and LaMeres, B. J., "Reliability analysis of field-programmable gate-array-based space computer architectures," *Journal of Aerospace Information Systems*, Vol. 14, No. 4, 2017, pp. 247–258.
- [18] *Soft Error Mitigation v4.1 Product Guide*, Xilinx, 4 2018. URL https://www.xilinx.com/support/documentation/ip_documentation/sem/v4_1/pg036_sem.pdf.
- [19] Bates, T., and Bridges, C. P., "Single event mitigation for Xilinx 7-series FPGAs," *2018 IEEE Aerospace Conference*, 2018, pp. 1–12.

- [20] *7 Series FPGAs Overview (v2.6)*, Xilinx, 5 2012. URL https://www.xilinx.com/support/documentation/ip_documentation/sem/v4_1/pg036_sem.pdf.
- [21] *Using a Microprocessor to Configure 7 Series FPGAs via Slave Serial or Slave SelectMAP Mode*, Xilinx, 5 2012. URL https://www.xilinx.com/support/documentation/application_notes/xapp583-fpga-configuration.pdf.
- [22] *Vivado Design Suite TCL Command Reference Guide*, Xilinx, 4 2014. URL https://www.xilinx.com/support/documentation/application_notes/xapp583-fpga-configuration.pdf.
- [23] McMurtrey, D., Morgan, K. S., Pratt, B., and Wirthlin, M. J., “Estimating TMR reliability on FPGAs using Markov models,” 2008.
- [24] Weber, R. J., et al., “Reconfigurable hardware accelerators for high performance radiation tolerant computers,” Ph.D. thesis, Montana State University-Bozeman, College of Engineering, 2014.
- [25] Tylka, A. J., Adams, J. H., Boberg, P. R., Brownstein, B., Dietrich, W. F., Flueckiger, E. O., Petersen, E. L., Shea, M. A., Smart, D. F., and Smith, E. C., “CREME96: A Revision of the Cosmic Ray Effects on Micro-Electronics Code,” *IEEE Transactions on Nuclear Science*, Vol. 44, No. 6, 1997, pp. 2150–2160.