Overview of the Upcoming RadPC-Lunar Mission

Chris Major
Montana State University
324 Norm Asbjornson Hall
Bozeman, MT, 59717
christopher.major@student.montana.edu

Brock LaMeres
Montana State University
316 Norm Asbjornson Hall
Bozeman, MT, 59717
lameres@montana.edu

David Klumpar
Montana State University
5th Floor Cobleigh
Bozeman, MT, 59717
klumpar@montana.edu

Larry Springer
Montana State University
5th Floor Cobleigh
Bozeman, MT, 59717
lspringer@physics.montana.edu

John Sample
Montana State University
5th Floor Cobleigh
Bozeman, MT, 59717
john.sample2@montana.edu

Skylar Tamke
Montana State University
P.O. Box 173840
Bozeman, MT, 59717
skylar.tamke@montana.edu

Rubin Meuchel
Montana State University
P.O. Box 173840
Bozeman, MT, 59717
rubin.meuchel@montana.edu

Colter Barney
Montana State University
324 Norm Asbjornson Hall
Bozeman, MT, 59717
colter.barney@student.montana.edu

Annie Bachman
Montana State University
324 Norm Asbjornson Hall
Bozeman, MT, 59717
anniebachman5280@gmail.com

Jake Davis
Montana State University
324 Norm Asbjornson Hall
Bozeman, MT, 59717
jakejdavis3@gmail.com

***Abstract*—This paper describes the mission concept for "RadPC-Lunar", a technology demonstration of a novel computer architecture that can recover from faults caused by ionizing radiation. RadPC-Lunar is a payload that was selected by NASA in 2019 to go to the Moon through its Commercial Lunar Payload Services (CLPS) project as part of the Artemis lunar program. RadPC-Lunar will travel to the lunar surface in 2022-23 onboard a commercial lander and will spend a minimum of 7-days in the Mare Crisium. This demonstration will serve two important purposes in support of future lunar missions. First, it will demonstrate a key technology for computationally intense autonomous lunar activities such as in situ resource utilization, robotic surface operations, and entry/descent/landing maneuverers while providing increased reliability over the state-of-the-art in space computers. Second, it will provide a characterization of the radiation effects environment of the lunar surface by tracking upsets within the computing fabric and correlating them to data from on-board dosimeters. The payload can also provide a unique set of measurements on the ionizing radiation environment as it passes through the Earth's magnetosphere during transit to the Moon.This paper describes the overall mission concept of RadPC-Lunar in addition to the details of the design-of-experiments and the types of data that will be collected. This paper will be of interest to engineers and scientists studying the lunar transit and lunar surface radiation environment and those working with radiation tolerant avionics. The timing of the presentation will allow the RadPC-Lunar team to solicit feedback from the aerospace community that can influence its design-of-experiments prior to completion in order to maximize the return of the mission.**

TABLE OF CONTENTS

978-1-7281-7436-5/21/$31.00 ©2021 IEEE

1. INTRODUCTION

The rapid increase of space exploration missions demonstrates an increased need for computing systems that can sustain operation in hazardous environments. One area of concern that has exhibited detrimental impact to the advancement of such systems is cosmic radiation – specifically, its effects on CMOS devices. Ionizing particles and energy are of little concern to computers and embedded systems on Earth as the planet's atmosphere and magnetic field attenuates these effects to such a degree that they rarely impact a device's performance. This protection is not guaranteed, however, in space [1].

Consideration of mitigating the effects of cosmic radiation varies depending on the category of effect. These classes, capable of causing a fault in a computing system's performance, are single event effects (SEE) and total ionizing doses (TID) [2]. Single event effects occur when high-energy particles or heavy ions strike a device, causing faults by creating unintended, momentary logic-level transitions. These upsets can be further classified into three categories, with increasing risk of functional damage. When the charge deposited by an energized particle changes the voltage level of a logic line, a single event transient (SET) has occurred and the logic value represented by that line has been compromised. If this compromised value is stored within a device, this fault is classified as a single event upset (SEU). If either of these events cause a system malfunction that cannot be corrected by simple measures, such as a device reset, a single event functional interrupt (SEFI) has occurred and more drastic action must be taken. SEEs can be summarized as temporary impacts on system performance, but can lead to system crashes.

In contrast to high-energy radiation effects, a total ionizing

dose is designated by the physical damage inflicted on devices caused by lower-energy particles. Such particles can deposit trapped charges into the insulating layers of CMOS devices, thus biasing transistor gates open to allow perpetual current. As this current can no longer be regulated by the device, the system suffers from abnormal power draws, material degradation, and the inability to properly conduct the intended task. Due to this behavior, TIDs can be summarized as gradual, permanent effects on a system's physical hardware structure [3].

Within the Earth's magnetosphere, many space computer systems have been tested in low-Earth orbit (LEO) and have demonstrated the ability to sustain or reduce damage from cosmic radiation. However, such systems have yet to be tested in other varieties of space radiation environments. One such environment is the lunar surface, since the lunar atmosphere fails to attenuate cosmic radiation due to a weak, nearly nonexistent magnetic field [4]. As continued interest in lunar exploration missions increases, radiation has become of greater concern for not just astronauts, infrastructure, and equipment, but for computing systems as well [5]. Thus, the need for an effective, cost-efficient computational system that can mitigate harsh radiation environments is apparent.

To this end, Montana State University has focused on the development of a reconfigurable system, called "RadPC," that can detect the effects of a SEE-induced fault and respond accordingly to repair damage. A mission-ready version of the system designated "RadPC-Lunar" will be sent to the Moon through NASA's Commercial Lunar Payload Services (CLPS) project as a part of the Artemis lunar program in 2022-23 [13]. The payload will spend a minimum of seven days in the Mare Crisium, where it is expected to experience radiation-induced faults, correct them, and report its status to the lander for transmission back to Earth.

2. Motivation

In response to the need for cost-effective radiation-tolerant space computing systems, Montana State University has developed a computing platform designed using commercial, off the shelf (COTS) components. Current platforms that attempt to address the dangers of radiation-induced faults range from protection through redundancy to resistance through architecture or process; however, many of these methods have proven to be cost-prohibitive for many missions and fail to provide consistent recovery strategies.

A basic approach to attenuating radiation effects is shielding, in which sufficient mass is placed between a computational device and the radiation environment. Though this strategy is already used within nuclear research facilities on Earth, the costs associated with bringing sufficient shielding material into space make this strategy far less desirable for space exploration missions. The addition of mass and volume to a spaceflight payload results in drastically higher costs for a mission and therefore can threaten the feasibility of the launch [6].

Fabrication methods such as Radiation Hardened By Process (RHBP) and Radiation Hardened By Design (RHBD) are often used within space computing systems as they demonstrate ability to reduce the probability of fault-induced charges from TIDs that can damage a system. RHBP decreases the probability of a strike depositing charge within a semiconductor gate by rerouting charge from sensitive components [7]. RHBD achieves a similar end through fabricating components with non-standard materials that reduce the probability of deposited charges in the electron-hole pairs of the semiconductors [8]. Though both techniques have demonstrated effectiveness in space exploration missions, they are expensive to manufacture and test. Additionally, in comparison to their standard manufactured counterparts, computational performance of these systems lags by at least a decade [9].

Given the smaller process nodes of modern integrated circuits, with features less than 65 nanometers, concerns regarding TIDs have been minimized and mitigating SEEs has become a more prominent concern. Triple modular redundant (TMR) systems, for example, are a common way to delay failure by radiation faults from SEEs by putting three identical logic circuits in parallel through a voting mechanism. If one circuit experiences a fault, the system can continue by the reported outputs of the majority [10]. However, this system will still fail when all three circuits are inevitably struck by cosmic radiation, therefore compromising its integrity.

As space exploration missions and programs continue to develop and accelerate in frequency, the need for cost-effective radiation tolerant computing systems has become clearer. Recently, NASA has announced plans for a manned mission to the Moon to be launched in 2024 as a part of the Artemis program [11]. This mission is a part of a multi-stage exploration and research plan, beginning with autonomous landers and equipment and eventually culminating in a continued human presence on the Moon. While existing fault mitigation strategies may assist in these and future missions, it has been shown that there remains a need for more cost-efficient strategies to protect space computing systems in future space exploration missions – especially as an immediate need has arisen through the Artemis program.

3. Mission Concept

Through NASA's Lunar Surface Instrument and Technology Payloads (LSITP) program, RadPC was selected as one of twelve experiments to be deployed on the surface of the Moon [12]. A commercial lander will carry the payload from the Earth to the Moon, coordinated by NASA's Commercial Lander Payload Services (CLPS).

Figure 1 demonstrates the maximum and minimum times that RadPC will be exposed to radiation environments and gather data for analysis. Pre-launch testing will ensure the system remains operational before deployment in space, for an estimated six to nine months. Lunar transit time will depend on the path taken by the payload, ranging from three to thirty-three days, where RadPC-Lunar will remain operational and continue its exposure to various radiation environments. A minimum of seven days will be spent on the lunar surface to gather data, store the data as packets, and transmit the packets back to Earth every eight hours. This data will describe the two key aspects of the mission: the radiation environment experienced by RadPC-Lunar, and the system's response to radiation-induced faults while performing a basic computational task.

4. Experiment Design

The RadPC-Lunar mission consists of a single payload that interfaces with a commercial lunar lander designed specifically for the CLPS mission. This payload is a single enclosure

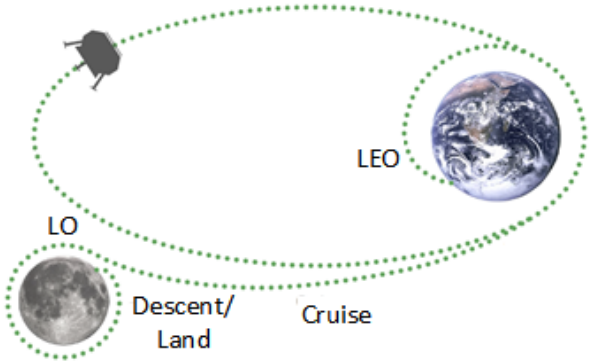

Mission Phase	Estimated Duration	Payload Mode
Pre-Launch	6-9 Months	I & T
Launch	1 Hour	OFF
Lunar Transit	3-33 Days	Operating*
Lunar Descent	10 Min	OFF
Lunar Surface	7 Days	Operating*

* During operation, the payload is powered and acquires experiment data that is to be transmitted back to Earth every 8 hours.

Figure 1. Estimated timeline for RadPC-Lunar mission

with three subsystems: the RadPC architecture, a dosimeter experiment, and a thermal system for temperature regulation. These three subsystems are intended to run independently to reduce the risk of cascading failures in the overall experiment. Though the experiment is expected to run for a week, minimum, on the lunar surface, all subsystems are intended to run during transit time from Earth to the Moon. A diagram of the various components of the RadPC payload is shown in Figure 2, below.

RadPC Experiment

RadPC's architecture is based on the implementation of reconfigurable logic circuitry on a Field Programmable Gate Array (FPGA), a device that can synthesize digital logic from a hardware description language (HDL). The logic to be run on the FPGA fabric is created by referencing a bitstream, stored in the device's configuration memory, to arrange the device's logic elements such that the intended functionality is exhibited. This act is known as a full configuration. Similarly, the act of rearranging only a section of the FPGA fabric's logic elements is known as partial reconfiguration (PR). The FPGA chosen to run this architecture is a Xilinx Artix-7 XC7A35T-FTG256-1 chip, implemented on a custom printed circuit board designed specifically for this mission.

In RadPC's logic design, four softcore processors, called "tiles" in the design, are instantiated in the FPGA fabric. Each tile runs the same program in parallel and sends the outputs to a voter component. This behavior mimics a standard TMR system but uses the fourth tile as an active spare in the case of an irrecoverable fault. Mathematical analysis of this system's performance shows a significant improvement in mean time to failure (MTTF) but also demonstrates rapidly diminishing returns when additional spare tiles are added. Figure 3, below, demonstrates this architecture.

Each tile is given up to 2 kilobytes of data memory with full read/write access. To protect the contents of the memory cells, Error Correction Codes (ECCs) are used. A state machine is used between the tile and the memory on the write line to encode each byte with four bits of ECCs, and a state machine is used between the tile and the read line to decode each twelve-bit word into its original byte using the ECCs. In the case of an error in memory, an external memory scrubber component will iterate through each cell of the data memory, compare values, and ensure all four memory components share the exact same data.

As each tile writes data to its output GPIO, the values are fed into a Voter state machine. This component compares each tile's output values to check for consistency. If any value differs from the majority, it is considered a fault and the tile is marked for repair. Though remarkably unlikely, in the case of an even-split vote between tile outputs, the Voter will mark the system as compromised and require an entire reset of the device.

To handle SEE-induced faults outside of the multiprocessors, a component known as the Soft Error Mitigation (SEM) controller is used. This is an IP core provided by Xilinx to scrub through the configuration memory of the FPGA and correct single-bit or double-bit-adjacent faults. The SEM sends out a status report via a UART connection depending on a given request or a detected fault. This status report can be used to identify where a fault has occurred in the system, if it was successfully corrected, and the device's current state.

Data regarding faulted devices within the FPGA is given to an external microcontroller device that monitors the system and conducts repairs. The microcontroller chosen is an MSP430 variant, commercially available. During RadPC's runtime, the microcontroller observes the voter and SEM UART, then performs repairs when prompted by the FPGA hardware. Data packets are arranged by the device by reporting the board's power statistics, computational output, FPGA health, and runtime status. These are stored in FRAM memory and are sent back to the lander for transmission to Earth at scheduled times in batches.

In the case of a tile failure, the microcontroller will prompt the FPGA to initiate partial reconfiguration of the compromised tile. To ensure that all tiles remain synchronized in program execution, a software checkpointing system is implemented within the Voter that will pause all other tiles mid-runtime while the tile in repair is reconfigured. When this repair is complete and the tile catches up in program execution to the current checkpoint, the Voter signals the tiles to resume program execution and to proceed to the next checkpoint. The results of this process are reported to the microcontroller and this data is included in the packet. Figure 4, below, shows the current revision of the RadPC experiment's printed circuit board.

Dosemiter System

As RadPC demonstrates its performance within a radiation-intense environment, a triplicated series of Teledyne UDOS001-C micro-dosimeters will be included in the payload. These dosimeters will be tuned to various energy levels, reflecting sensitivity to different types of cosmic radiation. As they are struck, a microcontroller will report the change in

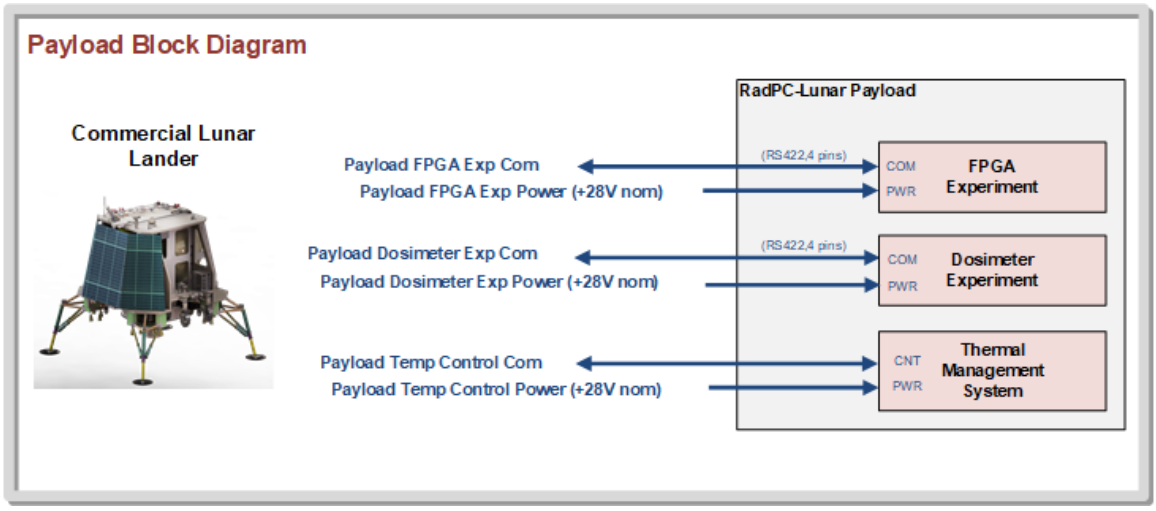

Figure 2. Payload Concept for RadPC-Lunar payload

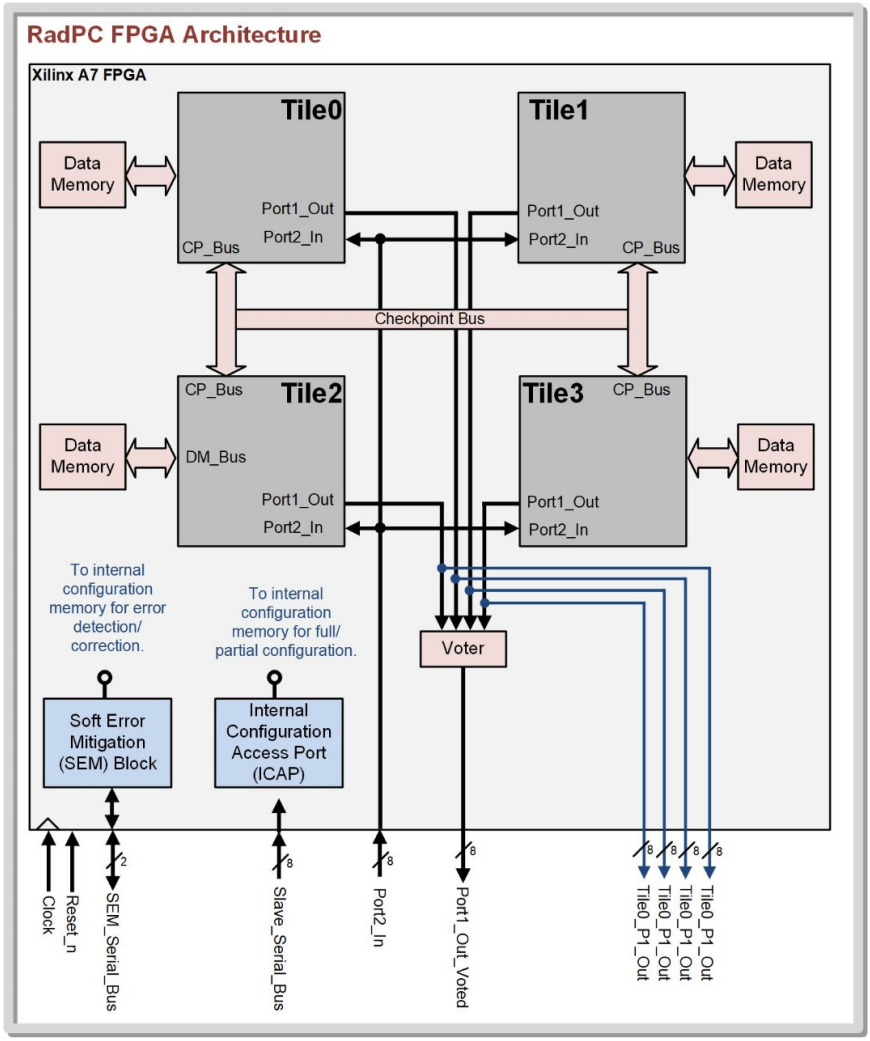

Figure 3. Block diagram of RadPC's 4MR architecture

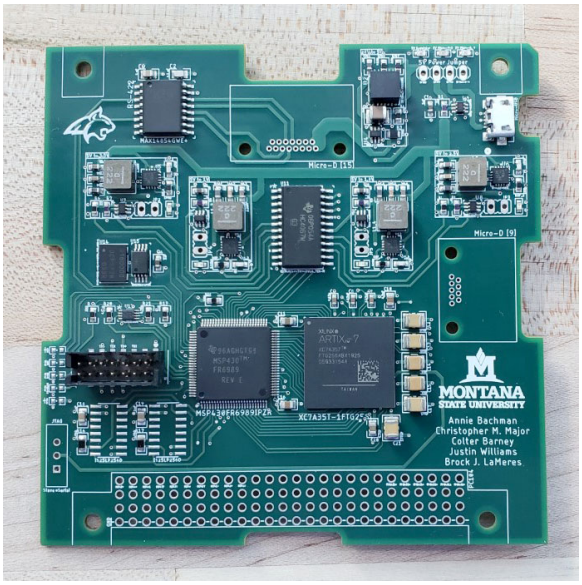

Figure 4. **RadPC experiment printed circuit board, current revision**

voltage across the compromised material and arrange a data packet for transmission by the lander. The dosimeters will allow the payload to collect radiation environment data for comparison with the reports sent from RadPC, giving a better perspective of the radiation events that will occur during the entirety of the mission. Data from the dosimeters is sent separate to the RadPC experiment, for the sake of limiting interdependence between systems.

Telemetry Packet

The data collected by the RadPC experiment microcontroller is arranged into a telemetry packet, which is sent to the lander for transmission back to Earth. The telemetry packet structure is shown in Table 1. After constructing a packet header identifying the packet's origin (Montana State University RadPC-Lunar FPGA), the packet count and time elapsed is given. Power and temperature data are included to ensure proper system operation. The tile input data, output count value, and voter output are given to characterize the system's expected computational performance. A record of all system faults, including those periodically injected to demonstrate recovery strategies, is included with a status report on the SEM controller's current state. The checkpoint and continue signal states give context to the count values and the status of each tile. The status of the memory scrubber is included to ensure data memory integrity. A field for each of the dosimeter values indicates the number of radiation faults. Finally, a CRC check value is added at the end of the packet to ensure the data is complete.

Thermal Management System

The thermal management system is intended to keep the temperatures of RadPC and the dosimeters consistent. As the lunar surface will fluctuate between extremes of cold and heat, a heating element within the payload will activate during cold to raise the temperature to desired levels. When heat needs to be released, the payload will radiate the excess away from the experiments. The consistency in temperatures will allow for smoother performance of the RadPC architecture and dosimeters with no significant power draw from the lander.

Payload

Per CLPS guidelines, the lander will be responsible for providing power and RS422 communication with the payload using a singular interface. The payload enclosure serves three purposes: maintaining the structural integrity of the experiment boards, managing excess heat, and providing a standardized power and communication port. One side will remain exposed to radiate unwanted heat and expose the experiments to radiation. It is expected that the payload will conform to a standard 1.5-U SmallSat format, to be mechanically interfaced with the lander and connect to its power supply and communication system via a D-SUB connector. Figure 5, below, demonstrates the payload-lander interface concept. Figure 6, below, shows a 3D-printed mock-up of the payload enclosure holding the experiment boards.

5. CONCLUSION

The mission concept and payload presented in this paper, named RadPC-Lunar, aims to demonstrate the performance of a cost-effective, commercial off-the-shelf, radiation tolerant space computing system. As space exploration missions continue to arise through the development of upcoming programs, the need for cost-efficient space computing systems that can mitigate faults caused by cosmic radiation continues to perpetuate as an area of concern. Through NASA's Commercial Lunar Payload Services program, an experimental payload developed by Montana State University will be launched in 2022-23 carrying a radiation tolerant space computing platform named RadPC. RadPC will demonstrate its ability to scrub and reconfigure components of its architecture that experience damage from faults induced by space radiation. Additionally, three Teledyne dosimeters will be deployed in a separate experiment within the same payload to characterize the radiation environment to which RadPC will be subjected. Through its time spent in transit to and on the lunar surface, RadPC-Lunar will demonstrate the ability for a computer architecture to continue performance and commit repairs on a system compromised by a harsh radiation environment.

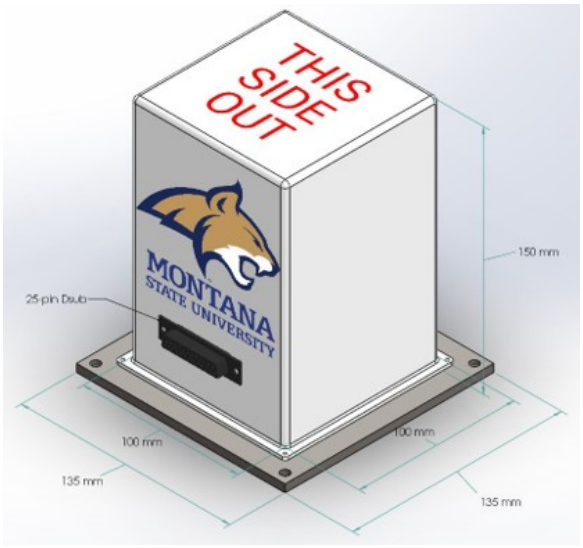

Figure 5. **RadPC-Lunar Payload Enclosure**

Figure 6. **RadPC-Lunar Payload Concept Prototype**

ACKNOWLEDGMENTS

The authors thank NASA Lunar Surface Instrument and Technology Payloads for funding this project.

REFERENCES

[1] J. L. Barth, C. S. Dyer and E. G. Stassinopoulos, "Space, Atmospheric, And Terrestrial Radiation Environments," IEEE Transactions on Nuclear Science, vol. 50, no. 3, pp. 466-482, 2003.

[2] C. Claeys, Radiation Effects In Advanced Semiconductor Materials And Devices, Springer, 2010.

[3] Adams, A. Holmes-Siedle and L., Handbook Of Radiation Effects, Oxford University Press, 2004.

[4] R. Pirich, J. Weir, Leyble, D, S. Chu and M. DiGiuseppe, "Effects of the lunar environment on space vehicle surfaces," in 2010 IEEE Long Island Systems, Applications and Technology Conference, Farmingdale, NY, 2010.

[5] R. K. Tripathi and J. E. Nealy, "Lunar Radiation Risk Assessment and Shielding Design for Ionizing Space Radiation," in 2008 IEEE Aerospace Conference, Big Sky, MT, 2008.

[6] R. Uzel and A. Özyildirim, "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), pp. 295-299, 2017.

[7] G. Anelli, "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, pp. 1690-1696, 1999.

[8] A. Makihara, "Hardness-by-design approach for 0.15 /spl mu/m fully depleted CMOS/SOI digital logic devices with enhanced SEU/SET immunity," IEEE Transactions on Nuclear Science, vol. 52, no. 6, pp. 2524-2530, 2005.

Table 1. Telemetry Packet Structure

Data Category	Metric
Header Data	”Montana State University”
	”RadPC-Lunar FPGA”
Packet Info	PACKET_NUMBER
	PACKETS_SENT_COUNT
Time Running	DAYS
	HOURS
	MINUTES
	SECONDS
Power Data	MAIN_RAIL_V
	MAIN_RAIL_C
	DEVICE_RAIL_V
	DEVICE_RAIL_C
Temperature Data	ARTIX_7_TEMPERATURE
	MSP_430_TEMPERATURE
	CASE_TEMPERATURE
	BOARD_TEMPERATURE
Program Counter	TILE_00
	TILE_01
	TILE_02
	TILE_03
	TILE_INPUT
	VOTER_OUTPUT
Fault Data	TILE_FAULT_COUNT
	SEM_FAULT_COUNT
	TILE_INJECT_COUNT
	SEM_INJECT_COUNT
SEM Data	MONITOR
	CORRECTION_COUNT
	LINEAR_ADDRESS
	FAULTED_WORD
	FAULTED_BIT
	SEM_STATE
	SEM_FLAG
Checkpoint	CHECKPOINT
	CONTINUE_SIGNALS
Memory Scrubber	MEMORY_ADDRESS
	READ/WRITE_SIGNAL
	MEMORY_DATA
	SCRUB_FLAG
Dosimeter Report	DOSE_1_FAULT_COUNT
	DOSE_2_FAULT_COUNT
	DOSE_3_FAULT_COUNT
CC Check	CRC

[9] A. K. a. J. A. a. D. F. a. M. Patrick, M. Watson, M. Johnson, J. Cressler and E. Kolawa, ”Developments in Radiation-Hardened Electronics Applicable to the Vision for Space Exploration,” in AIAA SPACE 2007 Conference & Exposition, Long Beach, California, 2007.

[10] L. Sterpone and M. Violante, ”Analysis of the robustness of the TMR architecture in SRAM-based FPGAs,” IEEE Transactions on Nuclear Science, vol. 52, no. 5, pp. 1545-1549, 2005.

[11] ”NASA Publishes Artemis Plan to Land First Woman, Next Man on Moon in 2024,” NASA, 21 September 2020. [Online]. Available: https://www.nasa.gov/press-release/nasa-publishes-artemis-plan-to-land-first-woman-next-man-on-moon-in-2024/. [Accessed 30 September 2020].

[12] ”NASA Selects 12 New Lunar Science, Technology Investigations,” NASA, 1 July 2019. [Online]. Available: https://www.nasa.gov/press-release/nasa-selects-12-new-lunar-science-technology-investigations. [Accessed 7 October 2020].

[13] ”NASA Enlists Commercial Partners to Fly Payloads to Moon,” NASA, 8 September 2020. [Online]. Available: https://www.nasa.gov/feature/nasa-enlists-commercial-partners-to-fly-payloads-to-moon. [Accessed 7 October 2020].

BIOGRAPHY

Chris Major is a graduate research student studying electrical engineering at Montana State University.

Brock LaMeres is a professor of electrical and computer engineering at Montana State University and the principal investigator of RadPC-Lunar.

David Klumpar is the director of the Space Science Engineering Laboratory at Montana State University.

Larry Springer is the project manager for RadPC-Lunar for the Space Science Engineering Laboratory at Montana State University.

John Sample is a research scientist for the Space Science Engineering Laboratory at Montana State University.

Skylar Tamke is a research engineer at Montana State University.

Rubin Meuchel is a research engineer at Montana State University.

Colter Barney is an undergraduate research student studying computer engineering at Montana State University.

Annie Bachman is an undergraduate research student studying computer engineering at Montana State University.

Jake Davis is an undergraduate research student studying mechanical engineering at Montana State University.