EUROPEAN WORKSHOP ON ON-BOARD DATA PROCESING (OBDP2021), 14-17 JUNE 2021

SYSTEM-LEVEL HARDENING TECHNIQUES USED IN THE COTS-BASED DATA PROCESSING UNIT

Piotr Kuligowski [(1)], Grzegorz Gajoch[(1)], Maciej Nowak [(1)], Wojciech Sładek [(1)]

[(1)] KP Labs Sp. z o.o., 44-100 Gliwice, Poland

ABSTRACT

CubeSat missions, in most cases, utilize commercial off-the-shelf (COTS) components. The COTS components are vulnerable to radiation effects such as single event effects (SEE) or total ionizing dose damage (TID). Those effects decrease overall system reliability and can lead to permanent damage to components. One of the methods of mitigating the risk is system-level hardening.

The most commonly used hardening techniques are fault detection, isolation, and recovery (FDIR) mechanisms implemented in a ruggedized controller that controls state-of-the-art systems on a chip (SoCs), error correction coded (ECC) or triple modular redundant (TMR) memories, and Configuration RAM (CRAM) scrubbing in the SoCs. In some cases, partial or full redundancy of selected components is implemented. These techniques were reviewed then selected techniques were implemented in the KP Labs' Leopard data processing unit (DPU). The proposed solutions may be re-used in other missions to fulfil mission reliability, availability, and safety levels.

The Leopard data processing unit (DPU) is a part of KP Labs' Intuition-1 mission scheduled to be launched in 2023. Intuition-1 will be a 6U-class CubeSat, and it will utilize a specialized hyperspectral camera with spectral resolution in the range of 470-900 nm with 150 spectral bands. The primary purpose of this mission is to technologically demonstrate the reduction of the spatial resolution of hyperspectral images (HSI), hyperspectral band selection, and segmentation of HSI [1] with a neural network-based in-orbit processing hardware that is the Leopard DPU. Implementation of algorithms on-board the satellite will allow to quickly decimate data, reducing the amount of radio air time required to download all the data to Earth.

The Leopard DPU consists of two redundant processing nodes controlled by a shared supervisor. Both elements are built using COTS components. Each processing node utilizes a state-of-the-art Xilinx Zynq Ultrascale+ MPSoC, 16 GB of DDR4 memory with ECC, 4 GB of NAND flash memory, and two 256 GB solid-state drives (SSD). Leopard DPU utilizes Xilinx's Vitis AI development environment platform to support mainstream AI networks such as TensorFlow and Caffe. Zynq's bootloaders and basic Linux image are located on a TMRed QSPI NOR flash memory, placed on a supervisor board. Moreover, the supervisor's board implements basic safety features for the two processing nodes, selects Linux images to be loaded, and multiplexes platform interfaces.

DPUs are powered by highly integrated power management integrated circuits (PMICs) surrounded by multiple sensors. The supervisor monitors currents, voltages, and temperatures to implement FDIR techniques for detecting single event latch-ups and high current events. Single event upsets (SEUs) are corrected on many levels. Zynq's configuration RAM (CRAM) is scrubbed by a soft error mitigation module. DDR4 utilizes an ECC mechanism that detects and corrects SEU errors that occurred in memory.

The Supervisor (SVR) board is composed of a radiation-hardened Vorago's Cortex-M0 microcontroller and accompanying ProASIC3 FPGA. The main tasks of the supervisor are: controlling DPUs, multiplexing platform interfaces such as high-speed X-Band and S-Band links, and routing CubeSat space protocol (CSP) packets.

1. MISSION REQUIREMENTS

To select adequate hardening techniques, system-level and mission requirements should be outlined. The reduction of the spatial resolution of hyperspectral images (HSI) segmentation of HSI with a neural network-based is a quite demanding task from FPGA-resources point of view.

Figure 1 Intuition-1 CubeSat - artist's impression

The estimated required computing power is within the range of 1 to 3 TOPS (tera operations per second, deep neural networks quantized to INT8).
Initially, the data is collected from an on-board hyperspectral camera. The hyperspectral camera collects about 300 frames per second with a 2 megapixel sensor that requires about 4.8 Gbps of frame buffer's bandwidth. Next, the hyperspectral data is collected in a 16 GB frame buffer and then processed with a deep neural network-based data processing system.
The processed or compressed hyperspectral data are then transmitted over X-Band high speed radio link. S-Band radio link is used to update Linux images and FPGA bitstreams that implement deep learning processing units.
Leopard data processing unit is commanded over a CAN interface by an on-board computer (OBC).
Intuition-1 CubeSat will be launched into a sun-synchronous orbit (SSO) which allows to perform imaging but it has additional risks associated with radiation effects, comparing to lower, non-SSO orbits. Originally the mission is scheduled on 1 year of in-orbit operation but it is possible to extend this mission up to 2 years. This imposes a requirement on the total ionization dose absorbed by Leopard. Including all the safety factors, it was assumed that the Leopard DPU should have a resistance to TID at least up to 20kRad.

2. IDENTIFYING THE THREAT

COTS components are susceptible to total ionization dose damage manifested by degrading electrical and/or timing characteristics. Some components may degrade with no visible symptoms and then fail permanently, such as charge pumps used to program Flash memories [2] and phase locked loops (PLLs). Thus proper components preselection process should be performed.
Most of the key components have available TID and/or SEE radiation data. For example, high-current events where identified in the Ultrascale+ family [3][4]. Auxiliary supply voltage (VCCAUX) 1.8V is the most susceptible power rail to high-current events induced by neutron radiation. Internal supply voltage (VCCINT) also indicates similar events, however, they occur less frequently. With no proper current limitation such high-current events may be destructive due to overheating or by exceeding electrical characteristics.
Other critical components are PMICs that deliver supply voltages to the Zynq Ultrascale+ MPSoC and to all the peripherals such as DDR4, NAND Flash and interface buffers. Any voltage transients or output voltage changes will impact Leopard's reliability or it can even lead to component failure. PMICs or simple DCDC converters may respond differently to TID [5]. Component selection should be addressed to select simplest PMICs possible with no digital control loops and signal processing such as IRPS5401 offers. Digital control loops in the power supply pose a risk to the SEU effects on internal configuration registers which may lead an overvoltage. Simpler PMICs that implement analogue loops are preferred such as ADP5052 quad buck converters.
Use of the analogue control loops, however, does not exclude output voltage transients on supply rails. Simple filtering should be done on all critical lines that filters fast transients caused by SET. Additionally, all power supply lines should be equipped with independent protections implemented by an external device.

3. WIDELY USED RADIATION EFFECT MITIGATION TECHNIQUES

There are many radiation effect mitigation techniques used in space industry such as hardening by design (component or circuit level) and system-level hardening [6]. Component-level hardening relies on radiation-hardened components by design. By choosing those components, mission requirements are met. Full, cold redundancy is also common. In low-cost systems and where a high performance is required, it is not possible to use these components. In COTS-based systems system-level hardening is more economic in some cases. Selecting proper system-level radiation effect mitigation techniques, COTS components can meet application radiation requirements. These techniques include system-level analysis in conjunction with mission requirements and then modifying the system-level behaviour to meet environment's requirements.
System-level hardening techniques can be applied to components where both destructive or non-destructive radiation effects have been discovered during radiation tests.
In in some cases appropriate protections may eliminate the destructive events. The most common is to use over-current limits and protections that detect SEL and high-current events. Threat reduction by maintaining safe operating temperature or electrical parameters (such as operating voltages) may be also beneficial reducing radiation effect rates. In extreme cases, where destructive events where not detected during the test campaign and cannot be avoided, cold redundancy allows to switch to a second unit and fulfil mission requirements.
Non-destructive events may be caused by SEE such as SEU cumulated in memories or registers or single event transients. Cumulating SEUs ma lead to SEFI where unit becomes unresponsive or uncontrollable. The most common techniques against SEE are EDACs, memory scrubbers (readback or blind), redundancy with voting such as triple modular redundancy (TMR). Hardening for SEFIs is often implemented with an external component (FPGA or microcontroller, that can be radiation-hardened or radiation-tolerant) which periodically checks the state of the component it monitors.

Although the fact that SEU and SEFI are not destructive, this can have unforeseen consequences that are destructive. In some critical components, such as DCDC converters with digital control loops, SEUs may lead to a serious damage that can propagate throughout the system.

4. SELECTED RADIATION EFFECT HARDENING TECHNIQUES

The selected radiation effect hardening techniques where presented in this chapter from top-level to the lower levels. At the beginning, system assumptions at the highest, system-level were presented. Then individual subcircuits and selected electronic components where discussed.

4.1. Processing node's redundancy

Starting from the most important requirement, which is criticality of this HSI data processing with deep neural networks experiment, architecture with two processing nodes was implemented. This architecture is shown on Figure 2.

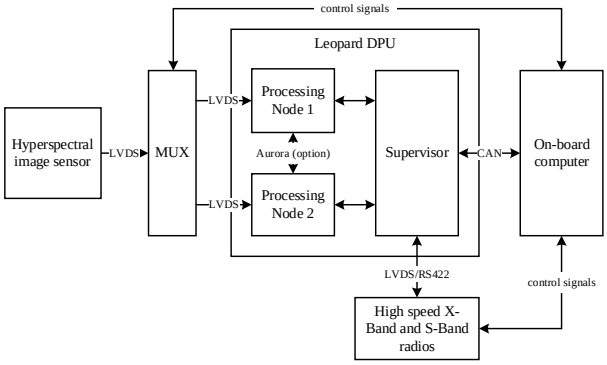

Figure 2 Top-level system architecture

There are two identical processing nodes fitted into one 1U Leopard data processing unit. Each of the processing node is equipped with a Zynq Ultrascale+ MPSoC, 16GB of DDR4 memory and 2x256GB of SSD-based mass memory. Apart from processing nodes, there is a Supervisor which multiplexes bus interfaces and an additional hyperspectral image sensor multiplexer that allows to multiplex one hyperspectral image sensor between the two processing nodes.

The Supervisor board is equipped with a radiation-hardened microcontroller Vorago Cortex-M0. This microcontroller implements EDACs or ECCs on all internal memories and registers what guarantees the most reliable operation possible. It is the only rad-hard component in this system.

It is obvious that this architecture is able to provide cold redundant data processing units. Apart from this type of redundancy, this architecture allows to operate processing nodes in parallel where one processing node is collecting hyperspectral image data, while the second processing node is compressing/processing previously collected images. In this mode of operation in case of any radiation-induced event during imaging process, the second processing node can take control over the image sensor, continuing imaging and allowing to reboot faulty processing node. Both software and hardware allow this type of operation. All these behaviors are controllable by CAN interface and can be scripted in an on-board computer. Intuition-1's OBC software, called Oryx, supports Lua scripting language. This scripting engine has an API to all subsystems and allows to redefine in-orbit all the high-level behaviors.

4.2. Safe Linux images and FSBL in a TMRed memory

NOR Flash memory provides first stage bootloader (FSBL), MPSoC's Cortex-R5 firmware, MPSoC's power management unit (PMU) firmware, U-Boot and safe/minimal Linux images to the processing nodes. To simplify image update process, these NOR Flash memories were located on the Supervisor board. This allows to update these images while processing nodes are disabled.

This memory is critical from mission point of view. Having one NOR Flash only is considered as too risky. Single component failure can lead to loss of the main payload. Triple modular redundant NOR Flash memory has been implemented in the Supervisor board. The Linux image selection process for a processing node was presented on the Figure 3.

Safe Linux image is stored in TMRed NOR Flash memory that can be updated via CAN interface which is used as a main communication interface with Leopard. CAN is used as telemetry and command interface that has a file-oriented protocol implemented. Safe Linux can be utilized to update/modify/repair nominal Linux images (up to 7) on NAND Flash memory that is populated on the processing node. Before turning the processing node on, desired Linux image (either Safe Linux or particular Nominal Linux image and bitstream from NAND) can be selected by command delivered via CAN control interface.

Shared partition on the NAND Flash memory is also responsible for storing multiple FPGA bitstreams that can be loaded on demand by a currently booted Linux image. Multiple bitstreams can be stored for each Linux image.

Any failure occurred in the both NOR Flash or NAND Flash memories does not lead to loss of the payload. While NOR Flash memories are TMRed, one NAND memory is populated for each processing node. In case of any NAND Flash failure, damaged processing node may be booted from TMRed NOR Flash only with no partition for shared data available.

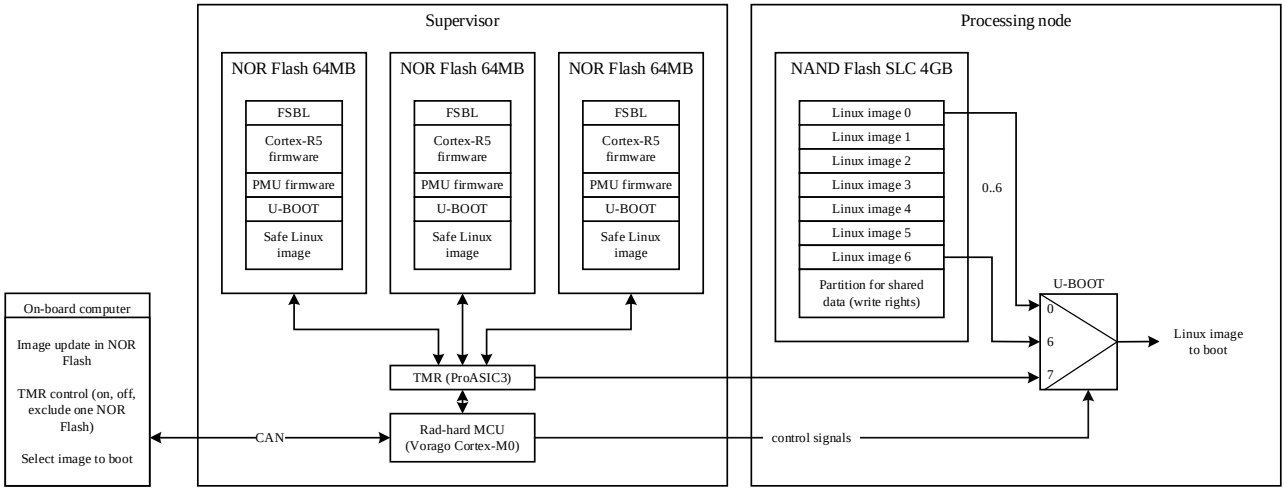

Figure 3 Linux image selection to boot

4.3. Over-current protections

Reconfigurable over-current protections were implemented to especially protect against high-current radiation effects caused by SEE on VCCAUX and VCCINT supply lines. These protections have also been implemented on the remaining power supply lines. Moreover, all supply lines have limited current capabilities that do not exceed MPSoC's electrical characteristics.

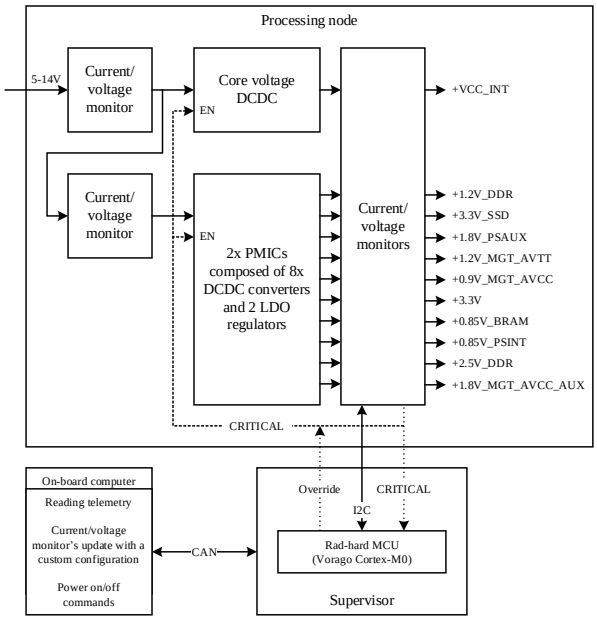

Figure 4 Over-current protection architecture

Over-current protection is composed of multi-channel current-sense and voltage monitors (INA3221 equivalent) placed on the processing node boards independently and a monitoring microcontroller placed on the supervisor. These current-sense monitors expose CRITICAL signals where voltage and current critical levels are configured via I^2C interface by the Supervisor.

Current-sense monitor's configurations are refreshed (along with the reading to verify configuration) every 1s to avoid SEU cumulation. The over-current protection architecture is shown on Figure 4. In case of exceeding any critical level, a processing node is forced to switch off by disabling its PMICs directly with a hard-wired ENABLE signal and an interrupt to the supervisor is generated at same time. Power off forcing can be overwritten by the Supervisor at any time. This allows to choose the reaction method (immediate force to switch off or delayed power cycle) depending on the type of event. This option will be investigated during radiation tests that are planned in 2021.

On the processing node board, the over-current protections were placed just after the two ADP5052 PMICs and a DCDC converter that produce voltages to the MPSoC core and its peripherals. It communicates via I2C interface with the Supervisor board equipped with a radiation-hardened Vorago Cortex-M0 microcontroller what increases reliability of this protection.

4.4. Redundancy of mass memory

Mass memory on a single processing node is composed of two 256GB industrial grade SSD-disks with extended operating temperature range from -40 to 85 °C. These disks use SLC Flash memory chips what reduces SEE cumulation over time due to lower SEU and SEFI cross-sections [8]. However, these disks are limited in size up to 256GB.

In addition, the disks are equipped with over-current protections and current/voltage monitors that report any unusual current consumption.

These disks operate in a RAID1 configuration what reduces the risk of loss of payload data caused by radiation effects.

4.5. Other SEU and SEFI mitigation techniques

There are many other mitigation techniques implemented that are addressed to non-destructive events. It includes MPSoC's configuration memory scrubbing and using of DDR4 with ECC.
The MPSoC was equipped with internal Xilinx's soft error mitigation (SEM) IP-core that performs SEU detection, correction and classification for configuration memory. It uses Readback CRC feature to implement readback scrubbing. Using of readback scrubbing reduces of risk of writing wrong configuration to the configuration memory. Similar behaviours were observed in the previous generation of Xilinx's FPGAs such as Ultrascale and 7000 families [7].
DDR4 that uses 9-chip configuration with ECC delivers SEU mitigation technique to PS' memory that is used both by Linux and an IP-core that interfaces with a hyperspectral image sensor.
SEFI detection is realized with the Supervisor that communicates over SPI interface with MPSoC's Cortex-R5 cores that act as CSP routers. Lack of communication or any undefined symptoms are threated as SEFI condition. Affected processing node is then power cycled to restore operation.

5. CONCLUSIONS

This publication presents one of the existing approaches of selecting system-level hardening techniques suitable for a specific data processing unit that will be used to perform on-board processing of hyperspectral image data on Intuition-1 mission.
Analyses were started by learning about the mission requirements and about criticality of some mission objectives. Next, the existing solutions were reviewed that implement protections against destructive and non-destructive events induced by radiation effects. And then proper techniques were selected such as over-current limits to protect against high-current effects present in the Zynq Ultrascale+ family or booting process was hardened against single point-of-failure on Flash memory caused by radiation effects.
This approach can be applied to harden similar systems against radiation effects to fulfil mission reliability, availability, and safety levels.
Leopard DPU is currently at TRL6 level. Radiation tests are scheduled to be performed in 2021. Test campaign includes both SEE and TID tests that will examine all the system-level mitigation techniques.

REFERENCESS

[1] Nalepa J., Myller M. et al., 2021, "Towards On-Board Hyperspectral Satellite Image Segmentation: Understanding Robustness of Deep Learning through Simulating Acquisition Conditions", MDPI

[2] S. Gerardin, M. Bagatin et al., 2013, "Radiation Effects in Flash Memories", IEEE Transactions on Nuclear Science

[3] David S. Lee, Michael King et al., 2018, "Single-Event Characterization of 16 nm FinFET Xilinx UltraScale+ Devices with Heavy Ion and Neutron Irradiation", IEEE REDW

[4] Jordan D. Anderson, Jennings C. Leavitt, Michael J. Wirthlin, 2018, "Neutron Radiation Beam Results for the Xilinx UltraScale+ MPSoC", IEEE NSREC

[5] Ameel J., Amidei D. et al., 2014, "Radiation-Hard Power Electronics for the ATLAS New Small Wheel", TWEPP

[6] Ladbury R., 2007, "Radiation Hardening at the System Level", IEEE NSREC

[7] David S. Lee, Swift G., Wirthlin M., 2016, "An Analysis of High-Current Events Observed on Xilinx 7-Series and UltraScale Field-Programmable Gate Arrays", SAND2016-6514C

[8] Bagatin M., Gerardin S. et al., 2013, "Proton-induced upsets in SLC and MLC nand flash memories", IEEE Transactions on Nuclear Science