SSC20-VI-07

NASA SpaceCube Intelligent Multi-Purpose System for Enabling Remote Sensing, Communication, and Navigation in Mission Architectures

Cody Brewer, Nicholas Franconi, Robin Ripley, Alessandro Geist, Travis Wise, Sebastian Sabogal,
Gary Crum, Sabrena Heyward, Christopher Wilson
NASA Goddard Space Flight Center
8800 Greenbelt Rd, Greenbelt, MD, 20771; 301-286-1247
cody.g.brewer@nasa.gov

ABSTRACT

New, innovative CubeSat mission concepts demand modern capabilities such as artificial intelligence and autonomy, constellation coordination, fault mitigation, and robotic servicing – all of which require vastly more processing resources than legacy systems are capable of providing. Enabling these domains within a scalable, configurable processing architecture is advantageous because it also allows for the flexibility to address varying mission roles, such as a command and data-handling system, a high-performance application processor extension, a guidance and navigation solution, or an instrument/sensor interface. This paper describes the NASA SpaceCube Intelligent Multi-Purpose System (IMPS), which allows mission developers to mix-and-match 1U (10 cm × 10 cm) CubeSat payloads configured for mission-specific needs. The central enabling component of the system architecture to address these concerns is the SpaceCube v3.0 Mini Processor. This single-board computer features the 20nm Xilinx Kintex UltraScale FPGA combined with a radiation-hardened FPGA monitor, and extensive IO to integrate and interconnect varying cards within the system. To unify the re-usable designs within this architecture, the CubeSat Card Standard was developed to guide design of 1U cards. This standard defines pinout configurations, mechanical, and electrical specifications for 1U CubeSat cards, allowing the backplane and mechanical enclosure to be easily extended. NASA has developed several cards adhering to the standard (System-on-Chip, power card, etc.), which allows the flexibility to configure a payload from a common catalog of cards.

I. INTRODUCTION

Today, advancements in small satellite (SmallSat) technology and miniaturization of sensor technology are enabling NASA to innovate with novel multi-satellite small mission architectures in place of single, monolithic, long-development satellites to achieve key scientific observations. While SmallSats cannot function as an “all-in-one” complete solution to all mission observables and cannot act as a substitute in all cases (i.e. due to limitations imposed by aperture and instrument size/power/precision for specific measurements), they have proven to be valuable contributors to a number of fields. SmallSats benefit from their comparatively lower cost, rapid development, and high launch-opportunity frequency compared to larger flagship-type missions. SmallSats are proving useful for both single spacecraft (i.e. early technology maturation) and constellations (i.e. commercial viability and multi-measurement science) configurations for science, defense, and industry [1]. CubeSat and SmallSat technology advancement is also advantageous for larger spacecraft, since innovation efforts to miniaturize electronics and other components can also be used for larger systems.

Both the relevancy and applicability of small form-factor electronics are becoming rapidly emphasized by the space community through new proposal calls and mission concepts. Additionally, like many other fields, the space research community has become enamored with the perceived capabilities for applications requiring computationally intense operations such as artificial intelligence. Furthermore, concepts for constellations of small spacecraft will need to rely on multi-element autonomy, coordinated fleet navigation, and quality of service and routing for communications. Performing these compute-intensive functions in harsh environments uniquely requires a high-performance onboard computer capable of providing autonomy, robustness, and fault tolerance.

Due to these considerations, NASA especially endeavors to find a balance between the burgeoning “new-space” approach (focusing on all commercial components and short-duration missions or applications with high risk tolerance) and the “traditional-space” approach (focusing on more stringent requirements for harsher environments and longer lasting missions). Frequently, these challenges can be acutely observed in multi-stage proposals where there is an early risk-reduction flight in a more benign low-Earth orbit (LEO), to build confidence for an extended mission in a harsher environment. Frequently, new designers will become captivated with the affordability of CubeSat electronics along with the previous space heritage claimed for a pre-existing CubeSat mission in LEO. However, when attempting to reuse the same design for a harsher environment, more rigor for qualification, development, and testing may be required, and the new environmental restrictions may prohibit the use of many commercial solutions that have only been proven viable in LEO.

This paper describes a payload architecture, SpaceCube Intelligent Multi-Purpose System (IMPS), that harnesses the benefits of the low SWaP-C (size, weight, power, and cost) form factor of CubeSats, while selecting components to meet high-performance requirements for processing and data transfer, and finally combines them with intelligent and novel design practices to improve reliability. To achieve a design that not only is affordable for varying mission environments, but also provides the processing capabilities necessary for onboard computing in a wide range of systems, the NASA Goddard Science Data Processing Branch has developed a CubeSat-sized design that includes multiple CubeSat slices which can meet the needs for a multitude of missions. These interchangeable designs form the structure that allows electronic 1U (10 cm × 10 cm) CubeSat cards to be heavily reused for other missions. This reusability allows for future designs to benefit from extensive heritage, as well as, architecture customizations by a mix-and-match approach from a diverse collection of compatible cards. Targeting reusable design practices and components meets needs for NASA, commercial space, and defense missions.

Throughout numerous mission and instrument formulation experiences, it has been identified that a diverse set of payloads can be realized with the same backbone infrastructure of key reused cards and the simple addition of one or two cards for mission-specific needs. NASA's science missions can greatly benefit from reusable, high-performance computing designs with a supporting infrastructure of cards. Many missions tend to fly technology demonstrations on the International Space Station (ISS) in preparation for future missions. [2] details several opportunities for technology demonstrations provided by the ISS Program Science Office for science research. Additionally, the "Small Satellite Missions for Planetary Science" study [3] led by NASA Glenn along with the National Academies Space Studies Board's "Achieving Science with CubeSats" identified radiation-tolerant flight computers as key needed future technology [1]. Finally, the updated 2020 NASA Taxonomy [4] (formerly NASA Technology Roadmap) illustrates several needs that can be met with crosscutting payload electronics that can be reconfigured for multiple mission classes and science objectives. For upcoming NASA programs, this type of architecture is highly advantageous for developing the technologies required to meet aggressive launch deadlines dictated by the Artemis program. This system is multi-purposed and can provide processing payloads for varying aspects of Artemis. The core technology development can be deployed in a lunar orbit to provide a communications and navigations node as part of LunaNet [5], execute high-performance, finely tuned precision landing algorithms for lunar landers, and be additionally reconfigured to provide mobility guidance capabilities for lunar robots and rovers on the surface. Finally, because these design slices are reusable, they address upcoming concerns described in the National Academies' Review of the Planetary Science Aspects of NASA SMD's Lunar Science and Exploration Initiative [6]. This architecture enables new technological capabilities needed for lunar studies without compromising the needs established in the *Vision and Voyages* planetary decadal study [7].

For defense, in "Outpacing the threat with an agile defense space enterprise," [8] a 2019 report led by the project Thor team of the Aerospace Corporation describes the challenges to United States security, with potential adversaries developing anti-satellite weapons and having much wider, unfettered access to space. In their recommendations, Aerospace Corporation discusses the need for rapid technology development, prototyping, and insertion. This paper describes an extensible architecture that provides a processing baseline, with capabilities to expand and interface new devices into the architecture, and to enable rapid evaluation of new devices monitored and managed by the reliable processing baseline devices. The proposed backplane design adopted by the described architecture allows for modularity and swappable system cards.

Furthermore, the "Air Force Space Command Long-Term Science and Technology Challenges" [9] memorandum describes two critical proposal objectives. Firstly, the document describes using new technologies for space superiority and warfighting in and from the space domain. These new technologies specifically highlight the need for automated and autonomous systems, artificial intelligence, and advanced computer architectures. An example of these applications is provided in Section VI demonstrating a hybrid architecture using two SpaceCube cards. The document also considers "novel and effective ways to support resilience of space systems" specifically highlighting resilient-by-design architectures and dynamically reconfigurable subsystems which are achieved with the fault-tolerant computer architecture of SpaceCube card designs.

II. BACKGROUND

The architecture described in this paper focuses on 1U SpaceCube processing cards. While the SpaceCube designs are highlighted here, the general card standard is described in Section V.

SpaceCube and SpaceCube Design Approach

SpaceCube is a family of Field Programmable Gate Array (FPGA) based on-board science data processing systems developed at the NASA Goddard Space Flight Center (GSFC). The goal of the SpaceCube program is to provide substantial improvements in onboard computing capability while maintaining a high degree of reliability and lowering relative power consumption and cost. In response to the critical and diverse needs of missions and instruments, the Science Data Processing Branch at NASA GSFC pioneered a hybrid-processing approach. This design approach combines radiation-hardened (rad-hard) and commercial components while emphasizing a novel architecture synergizing the best capabilities of CPUs, DSPs, and FPGAs. This division of tasks is conducted with extensive algorithm profiling and partitioning, matched with mission requirements, to best align computational stages with architecture components. This hybrid approach is realized through the SpaceCube family of data processors that have extensive flight heritage for several cards.

In addition to the hybrid architecture design, the SpaceCube approach encompasses several design principles for both reliability and configurability at both card- and box-design levels. A more detailed description of the SpaceCube design principles can be examined in [10]. The summarized key design principles include reliable monitors, quality and intelligent part selection, and modularity.

Challenges for Commercial Processors

While [11] identifies a broad number of commercial vendors in the CubeSat market, there are considerable challenges that must be addressed to incorporate these designs into broad mission types. As shown in [12] many constellation missions focus on Earth observation (EO) and these spacecraft typically reside in radiation benign LEO. Since these missions are typically short duration, they are unlikely to fail due to single-event effects caused by the radiation environment. Due to these mission use-cases, many commercial vendors do not perform any or perform limited radiation testing and parts qualification. Additionally, since radiation effects are not a high priority for many commercial vendors, they do not investigate packaging their systems with fault-tolerant packages or system recommendations for reliability (e.g., scrubbing, triple modular redundancy) that are essential for operating in harsher environments. Researchers at NASA Jet Propulsion Laboratory (JPL) even identified that several commonly used CubeSat processors catastrophically fail to radiation effects, however, due to the low rate of single-event upsets in an equatorial LEO environment, the probability of an event is low [13].

III. APPROACH

This section describes the expected approach for constructing a payload system with the proposed architecture, as well as, a brief list of card slices already available in the format. Section IV details the highly configurable, I/O dense SpaceCube v3.0 Mini that a majority of missions would use to connect to instruments. Section V provides an overview of the details for the card design standard that can be used by the reader to build compatible cards. Finally, Section VI includes example deployment configurations for the architecture.

High-Level Overview

The baseline system architecture includes a SpaceCube v3.0 Mini processor and a backplane (pictured in Figure 1). The CubeSat Card Standard (CS^2), described in Section V, provides pinout configurations, mechanical, and electrical specifications for 1U CubeSat cards, allowing the backplane and mechanical enclosure to be easily extended. NASA has developed several cards adhering to the standard (single-board computers, power cards, routers, etc.), which allows the flexibility to mix-and-match the entire catalog to configure the system. Additionally, included is an I/O card, featuring standard interfaces (such a M.2 connectors), which allows developers to prototype unreliable (or untested) devices interfaced to the reliable architecture with isolation and fault protection.

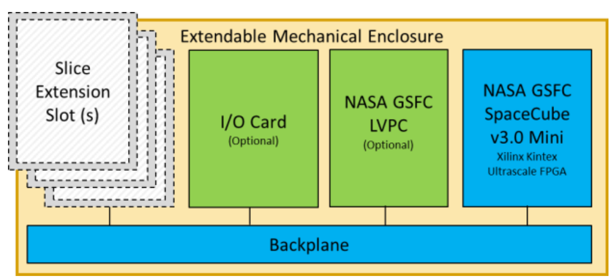

Figure 1: Architecture Diagram

The system design is enabling for current and future NASA missions with varying environments and durations. This paper presents several applications and analysis of the system design that have been incorporated into upcoming missions and proposals in Section VI.

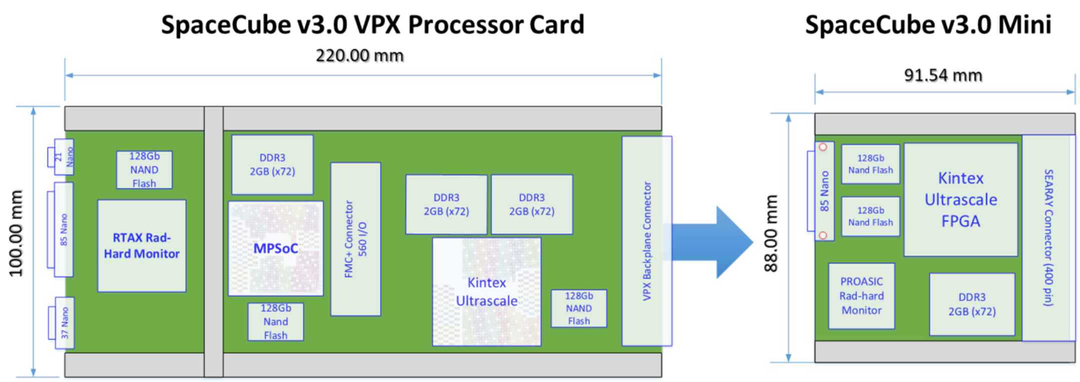

Figure 2: Comparison of SpaceCube v3.0 VPX Main Processor Card to SpaceCube Mini

Compatible Cards

Several cards have already been developed, or are currently in development, for the CS^2 specification (Section V). This section concisely lists some of these designs.

LVPC (Low Voltage Power Converter): This card provides clean and isolated secondary voltages for the processor box, along with switched services for different voltages. This card is used for missions that require the processing box to generate its own secondary voltages from the spacecraft bus power.

SDR (Software-Defined Radio): Under development to provide both remote-sensing and communication applications with a reconfigurable software-defined radio design. This new architecture is optimized for low SWaP-C characteristics and features a scalable design for multi-input multi-output (MIMO).

SpaceCube Mini-Z: This design is an evolution of the popular CSPv1 [19] and features the Xilinx Zynq-7000 SoC (dual-core ARM Cortex-A9, 28-nm FPGA). This card is included on several NASA Goddard CubeSats and has extensive flight heritage.

SpaceCube v3.0 Mini (SCv3M): Next-generation SpaceCube in a CubeSat form-factor with a massive FPGA. This design supports the latest advancements in FGPA design tools and productivity, allowing easy integration of some of the latest Xilinx designs and frameworks, such as the Deep Learning Processor Unit (DPU), Vitis AI, and Vitis High-Level Synthesis. Described in detail in Section IV.

SpaceCube Mini-Z45: Modification of the SpaceCube Mini-Z that upgrades the system to a higher resource capacity FPGA/SoC. This device also includes high-speed multi-gigabit transceivers to connect to sensors or to SCv3M.

Solid State Data Recorder (SSDR): In progress miniaturization of MUSTANG Data Storage Board [14] for CubeSat designs. This design is one of the most frequently requested cards because many missions require extensive storage capacity, typically to make up for limited transmission contacts or large sensor data products.

GPS: Currently under development at NASA Goddard, this design miniaturizes Navigator GPS for CubeSats and is designed to be paired with SCv3M. This effort envisions the NavCube (NASA Goddard's 2016 Innovation of the Year combining SpaceCube v2.0 and the Navigator GPS) card into a 1U CubeSat box design.

IV. SPACECUBE V3.0 MINI (SCV3M)

The SpaceCube v3.0 Mini is a unique, 1U CubeSat-sized single-board computer that features the Xilinx Kintex UltraScale (20-nm FPGA), and currently has no off-the-shelf industry equivalent. This design (displayed in Figure 4) is the latest addition to the SpaceCube family that provides extensive I/O for interconnects and networking, a fault-tolerant architecture, and several multi-gigabit transceivers for high-speed interfaces.

Design Philosophy

Just as its predecessor, the SpaceCube v2.0 Mini (SCv2M [15]), the design methodology for the SpaceCube "Mini" series is to leverage the design of a much larger processor card in the SpaceCube family (within the same Xilinx technology generation), and reduce the core functionality to conform to a CubeSat form-factor design. The SpaceCube v3.0 VPX main processor card (SCv3VPX), described in [10], features a Xilinx Kintex UltraScale (20-nm FPGA) with a Xilinx Zynq UltraScale+ MPSoC (quad-core 64-bit ARM Cortex-A53, dual-core ARM Cortex-R5, 16-nm FinFET+ FPGA) and radiation hardened monitor. Due to the complexity, size of components, and necessary regulators, it is not possible to include both Xilinx devices on a single 1U card. For the SCv3M, the Kintex

Figure 3: SpaceCube v2.0 Mini Design

UltraScale was selected for two primary considerations. First, the Zynq UltraScale+ MPSoC has documented radiation environment considerations, which are mitigated in the SpaceCube v3.0 VPX design with external monitoring and circuitry that would be too constrained by PCB area on a 1U card (MPSoC radiation information available from Xilinx). Secondly, the commercial/industrial-grade Kintex has been tested and performed well in [16] and [17]. Xilinx is also fully committed to supporting the development of the space-grade Kintex UltraScale (XQRKU060)[1] for the space community (same die as the industrial-grade part but with ceramic package and additional screening). Therefore, the Kintex UltraScale combined with a smaller radiation monitor was selected for the architecture design of the SCv3M. An architecture comparison demonstrating the migration from the SpaceCube v3.0 VPX to the Mini card is pictured in Figure 2.

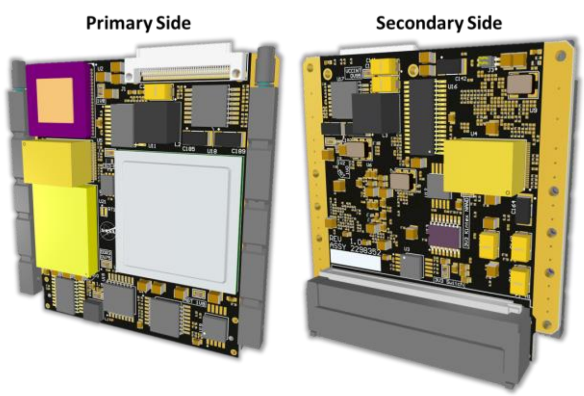

Figure 4: SpaceCube v3.0 Mini 1U Kintex UltraScale CubeSat Single-Board Computer

The SCv3M shares similarities with the VPX main processor card and leverages many features of that design. Some critical examples include design reuse of the interfaces for the Kintex UltraScale (with the DDR3 pinout as the most significant). Just as the SCv3VPX design includes a Microchip (formerly Microsemi) RTAX radiation-hardened monitor (RHM), the SCv3M includes a substantial amount of the same design and reusable code but on a smaller, reconfigurable Microchip RT ProASIC3, more suitable for the condensed CubeSat size. Finally, the most advantageous commonality with the SCv3VPX card is the reuse of components for the bill-of-materials which leverages thorough Electrical, Electronic and Electromechanical (EEE) parts trades, analysis, and circuits, pre-defined and studied for the VPX card.

SpaceCube Mini Heritage and Lessons Learned

The SpaceCube team has extensive experience in building small payload electronics through the development of earlier systems detailed in this section. The SpaceCube v3.0 Mini is a continuation of the "Mini" SpaceCube product line. As described previously, the first of the "Mini" series was the SpaceCube v2.0 Mini, based on the broadly successful SpaceCube v2.0 design [18] that used the Xilinx Virtex-5 family of devices.

The SCv2.0 Mini is incorporated on a couple missions and is extensively described in [15] and pictured in Figure 3. The primary goal of the SCv2.0 Mini was providing a near functional equivalent of the SpaceCube v2.0 in the 1U CubeSat form factor. However, this card was uniquely designed with multiple sections (or cards) interconnected with rigid-flex that allowed the system to be mounted close to the mechanical housing panels to conserve volume, and also folded around a detector lens or smaller payload electronics.

The SpaceCube team identified two major lessons from the development of SCv2.0 Mini. The first lesson was that while the rigid-flex offered the unique capability for folding the cards, it locked the design fabrication to a single vendor for manufacturing. This card also required laser-drilled microvias that when combined

[1]https://www.xilinx.com/news/press/2020/xilinx-lifts-off-with-launch-of-industry-s-first-20nm-space-grade-fpga-for-satellite-and-space-applications.html

with the bookbinder rigid-flex made the card difficult to manufacture. Additionally, this design had also included an Aeroflex rad-hard monitor; however, the logic gate count in the device was limited preventing the inclusion of more robust features.

Following the SCv2.0 Mini design, NASA GSFC would collaborate with the NSF (National Science Foundation) CHREC (Center for High-Performance and Reconfigurable Computing) to develop the popular CSPv1 CubeSat processor card [19], a hybrid system-on-chip design combining dual-core ARM processors with FPGA fabric. The CSPv1 has heritage on numerous missions (STP-H5, STP-H6, etc.) and is proposed on many more. This design would be one of the earliest designs to explore the use of the Samtec SEARAY connector for flexibility and performance. Furthermore, the SpaceCube team would make substantial feature upgrades to the design to produce the SpaceCube Mini-Z card (Figure 5).

Figure 5: SpaceCube Mini-Z Processor Card

Finally, many of the commercially available 1U CubeSat cards have adopted PC104 and similar stacking connectors as part of the growing CubeSat market trend. These types of connectors present a considerable challenge for routing, pin availability, and high-speed signaling across designs, and are more highly detailed in [20]. In contrast to both the earlier SCv2.0 Mini system that connected the modules through rigid-flex and the commercial PC104 designs, the next generation "Mini" series embraced a backplane-style approach. This type of design is favorable because high-speed signals, such as the multi-gigabit transceivers, provided by SCv3M, can be routed more easily to other designs on the backplane. The backplane architecture is scalable and easily extended. The backplane also allows cards to be swapped out during integration and testing, without the complexities and concerns of disassembling a stack of cards, as would be required with PC104. These design considerations have been incorporated into the CS^2 standard highlighted in Section V.

High-Level Architecture Design

The primary component of SCv3M is the Xilinx Kintex UltraScale, however, it also includes a variety of supporting components and resources. Specifically, the board currently supports the -1 or -2 speed grade for the XCKU060-#FFVA1517, with future plans to support the space-grade XQRKU060 part (which was not previously available and recently released as of this publication). The high-level block diagram of the system design is pictured in Figure 6.

For volatile memory resources, the design includes a 2 GB (72-bit wide) DDR3 SDRAM memory module. The extra byte provided by this device is used for ECC (Error-Correcting Code) so the FPGA can respond to and mitigate upsets in the memory module. This memory is typically used to store an initramfs-based operating system (OS) when hosting soft microprocessor cores (e.g., MicroBlaze, RISC-V, etc.) and/or used to buffer dynamic application data, such as images or attached instrument data.

For non-volatile memory, the design includes two 16 GB NAND flash memory modules. One module is connected directly to the Kintex to store OS boot images and/or finalized or intermediate application data products. The other identical module is connected to the RHM, which will be typically used to store configuration files for the Kintex UltraScale, but can also be used for slower transfer long-term data storage. In total, this system provides 32 GB of NAND flash memory, although some portion of the storage would need to be allocated for redundant boot images for fault tolerance.

For external interfaces, the SCv3M provides extensive I/O connections to accommodate the immense volume and speed requirements that may be imposed by high-performance detectors and sensors. Unlike some commercial options, the SCv3M includes 12x multi-gigabit transceivers that can provide up to a maximum transmission rate of 12.5 Gbps (-1 speed grade) or 16.3 Gbps (-2 speed grade). Through the backplane Samtec connector, the design includes 48x LVDS pairs (can also be configured as 96x 1.8V single-ended I/O), 47x 3.3V GPIO, and an assortment of interfaces including SPI, SelectMAP (SMAP), and JTAG. The Kintex is the master of the on-board SPI bus, which is connected to a housekeeping 12-bit, 8-channel ADC (analog-to-digital converter) for current, voltage, and thermistor monitoring. Additional slave devices can be added to this bus over the backplane connector. There is also a discrete IC (integrated circuit) that monitors the Kintex internal temperature diode, which can be read by the Kintex over I2C. Lastly, the optional front-panel 85-pin Nano connector provides 24x LVDS pairs (can also be configured as 48x 1.8V single-ended I/O) and 8x 3.3V GPIO.

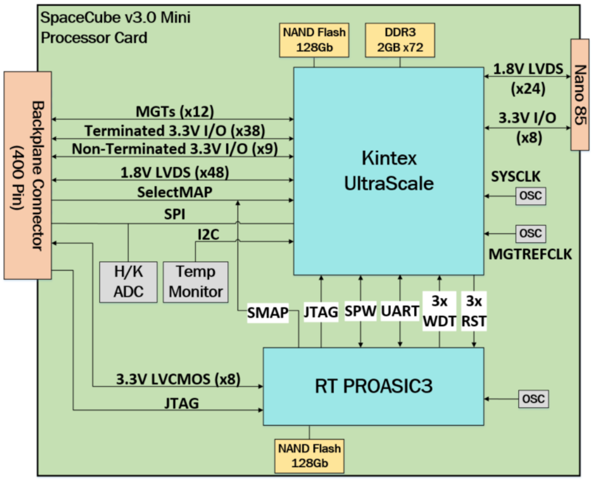

Figure 6: SpaceCube v3.0 Mini Block Diagram

Performance

Selecting the Kintex UltraScale device for this design is significant because it provides orders of magnitude of performance improvement over other 1U CubeSat cards. To compare the Kintex UltraScale to other space computing devices we use a metric known as computational density, measured in gigaoperations per second (GOPS) described in [21]. The purpose of this metric is to develop a means of providing a fair, deterministic measurement to compare the maximum theoretical performance capability of computing devices with different architectures. A comparison of the Kintex UltraScale to other commonly used space processing devices is featured in Figure 7 in logarithmic scale to demonstrate the orders-of-magnitude performance improvement compared to state-of-the-art rad-hard processors.

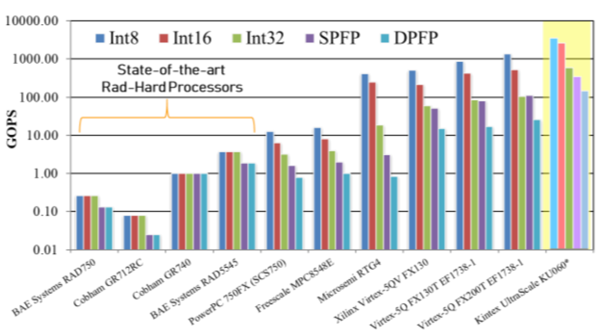

*UltraScale measurements are estimates based off of existing data in **[21],** new metrics are in progress but not currently available

Figure 7: Log Scale Comparison of Giga-Operation Per Second of Space Devices

Fault-Tolerant Board Architecture Design

Reliable operation in varying space environments is challenging due to space radiation effects, which can incur a broad spectrum of damage and errors, from benign bit-flips in unused memory to catastrophic failure of a component. The breadth of these effects and challenges to specific types of EEE components are not described in this paper, however, the field is broad therefore [22] and [23] can be used as a starting survey of paper references that cite authoritative documents in the field. Papers [24] and [25] specifically describe the radiation effects characterization of the Kintex UltraScale. To address these concerns, SCv3M includes both an intelligent fault-tolerant board architecture design and internal FPGA mitigations.

As part of the SpaceCube design methodology, the SCv3M design includes a rad-hard Microchip RT ProASIC3 to provide radiation mitigation and system monitoring for the entire card. The RHM provides the SCv3M with a variety of features. It can configure the Kintex UltraScale, scrub the configuration memory to correct upsets, and monitor the health of the Kintex using watchdog timers. The RHM can be configured to perform simple blind scrubbing (periodic scrubbing) or configuration readback scrubbing for low-latency error detection and correction (via frame-level ECC with global CRC checks). The Kintex configuration files are stored in external non-volatile memories and the RHM also uses error detection (via page-level CRC checks) and multiple copies (typically dozens of configuration files are stored with redundant copies across multiple internal dies of the NAND flash memory) to mitigate against upsets and verify the configuration files in storage. The RHM can also reconstruct a valid configuration file from several corrupted ones in storage if multiple images are corrupted. Internally, the RHM ensures the Kintex programming and boot sequence is completed correctly and will initiate automatic retries of the programming sequence if required.

The RHM also hosts a SpaceWire router that connects externally through the backplane and connects to the Kintex. This interconnect architecture allows the spacecraft to communicate directly to both the RHM and Kintex through the same interface and can be used to issue resets if necessary or change configurations entirely to support in-flight dynamic mission reconfiguration. This allows rapid switching between entirely different functionality for various phases of the mission. Due to the limitations of the RT ProASIC3, the external RHM SpaceWire interface requires LVDS transceivers to be populated externally from the card (typically included on the backplane or I/O card), which would be connected to the RHM's 8x external 3.3V single-ended I/O.

For some more stringent mission classes it has been noted that the RT ProASIC3 has previously experienced some performance degradation at comparatively low total-ionizing dose (TID) levels[2]. To compensate, the software will include a propagation delay derating to account for this increase for timing analysis. However, due to this limitation, the SCv3.0 has a selectable booting configuration. The design can be configured through the SMAP from either the backplane or the onboard RT ProASIC3 supervisor. A companion card or another radiation-hardened processor card could assume the monitoring, booting, and initial configuration of the Kintex device in place of the RT ProASIC3 through the backplane if required. NASA Goddard has developed a Microchip RTG4 1U card that can be used for this purpose in the "MARES" architecture described in [26].

The last fault-tolerance feature included for the SCv3M are several watchdog and reset lines for health and status monitoring. The high-level view of these signals is pictured in Figure 8. The watchdog frequency and reset requirements are configurable in the RHM. These watchdog timers can be used independently to reset different subsets of the Kintex design, including a top-level design reset. It should also be noted that the RHM can also be reset through a reset command issued over the SpaceWire interface (if used), and any of the available FPGA I/O can naturally be configured for this functionality as well.

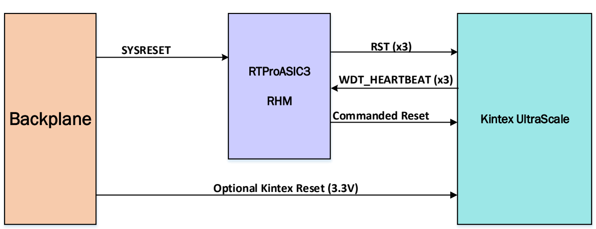

Figure 8: System Watchdog and Resets

Fault-Tolerant FPGA Mitigation

SRAM-based FPGAs can be affected by single-event upsets in a radiation environment that can change their configuration memory. Typically, space FPGA developers will, in addition to scrubbing, employ some variant of redundancy, most commonly Triple Modular Redundancy (TMR). For SCv3M, missions will rely on a softcore processor to perform minor computational tasks. Table 1 provides resource utilizations results for the Xilinx KCU105 development board frequently used as a stand-in board for SCv3M for testbed development. Four different designs are compared in Table 1. Initially provided are the results for a completely unmitigated MicroBlaze design. The second design is one generated using the Xilinx TMR MicroBlaze [27], a built-in Xilinx TMR solution for newer FPGA families that includes the Soft Error Mitigation (SEM) IP core and is included as a part of the Vivado IP integrator facilitating project creation. The third design is the Xilinx MicroBlaze generated using the BL-TMR (BYU-LANL TMR Tool[3]) frequently used by industry and academia for generating TMR designs. Finally, the last design displayed is the BL-TMR for the RISC-V with results provided in [28].

Table 1: Resource Utilization of TMR Designs on KU040

Resource	Unmitigated MicroBlaze	Xilinx TMR MicroBlaze	BL-TMR MicroBlaze	BL-TMR RISC-V3
LUTs	3.29%	9.81%	15.58%	4.48 %
CLB FF	1.63%	4.77%	4.89%	0.6 %
BRAM/ FIFO ECC (36 Kb)	12.50%	37.50%	37.50%	3.0 %
DSP Slices	0.31%	0.94%	0.94%	0.6 %
FMax Ratio to No-FT MicroBlaze	-----	0.95x	0.88x	0.73x

BL-TMR v6.3, MicroBlaze v11, 32-bit 5-stage, FPU, 32 Kb I/D, Vivado 2019.1

Development and Configuration

For ground testing as part of a FlatSat or engineering development unit, two cards were constructed to provide convenient development interfaces to designers, and access to the significant I/O resources provided by the SCv3M. Combined, the three cards constitute the SpaceCube Mini Evaluation Kit pictured in Figure 9.

The first card is the SpaceCube v3.0 Mini Active Evaluation Board. This card provides a number of common interfaces for flight software and FPGA development. It provides all required power to the SCv3M using either bench power supplies or a common wall outlet 12V power brick. An essential design feature of this card is an FMC+ connector that breaks out most of the I/O to be used with other FMC-compatible vendor designs, or custom cards developed to test instruments. Key interfaces of this card include, gigabit Ethernet (RJ45/SGMII), USB JTAG debug, SMAP programming header, 2x SpaceWire ports, 4x RS-422 ports, and the primary FMC+ connector

[2] https://www.microsemi.com/product-directory/rad-tolerant-fpgas/1696-rt-proasic3

[3] http://reliability.ee.byu.edu/edif/

(supporting 11 MGTs, 46x LVDS lines and 22x 3.3V GPIO).

The second card is the Mezzanine "Mezz" Evaluation Card. This board was designed as a breakout board that is compatible with the SCv3M and SCv3VPX, and breaks out a significant amount of I/O through many different interfaces. Interfaces to the SCv3M through this board include 3x SpaceWire ports, 4x RS-422 ports, 2x gigabit Ethernet (RJ45/SGMII), 2x Camera Link (Base or Medium), 1x USB2.0, 1x SATA, 1x CAN bus, 1x MGT over SMA, 7x MGTs over high-speed verSI connector, and a GPIO connector with 9x LVDS and 3x 3.3V GPIO. The USB interface is capable of reading a flash drive, allowing for quick and simple transfer of large files to the SCv3M. The board also hosts a configurable clock generator, which can provide an alternate MGT reference clock, allowing for easy adjustment of MGT speeds to meet varying standards and requirements.

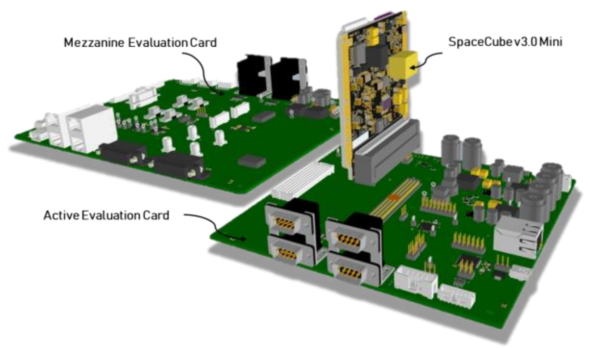

Figure 9: SpaceCube Mini Evaluation Kit

The SCv3M board support package includes a basic FPGA reference design to interface with the peripheral components. A basic Yocto Linux for MicroBlaze is also included as a design reference. Typically, all SpaceCube systems support core Flight System (cFS) as part of the board support package, promoting the rapid deployment philosophies identified by Aerospace Corporation in [8]. Core Flight System[4] is NASA Goddard's open source, flight-software framework licensed under Apache 2.0 and the NASA Open Source Agreement (NOSA). cFS is advantageous because it allows reusable flight software to be re-deployed from mission-to-mission (which has been demonstrated on many SpaceCube missions), dramatically reducing software development costs. Components of cFS have been verified for up to NASA Class A missions.

V. CUBESAT CARD STANDARD (CS^2)

To provide the flexibility and interoperability to mix-and-match varying designs to construct a new system, a standard or template is required for all the designs. For this purpose, the CubeSat Card Standard, also known as CS^2, is defined in this section and is managed by the Embedded Processing Group of the Science Data Processing Branch at NASA Goddard. The standard establishes baseline configurations to develop CubeSat 1U-type cards compatible with several NASA programs, and was developed to address NASA-specific concerns that were not being met by existing standards.

CS^2 establishes the common interface between CubeSat cards, encourages design reuse, and provides a convenient reference to integrate with the numerous cards (and mechanical structures) supported by the SpaceCube family of designs. This section further highlights the respective card connectors, physical card dimensions, and depicts backplane configuration and pinout for flexibility in routing and configuration. Figure 10 identifies defining parts of CS^2.

High-Speed Connectors

A compelling requirement for these systems is they must be capable of supporting high-speed data transfers. In initial market surveys, the Samtec SEARAY connectors were identified to meet future mission needs. These connectors were tested and evaluated on several NASA missions and were previously used on the CSPv1 single-board computer. Characterization reports[5] for these connectors are readily available showing supported rates up to 12.5 GHz or 25 Gbps, which is significantly faster than the rates that can be sustained by devices proposed for this form factor.

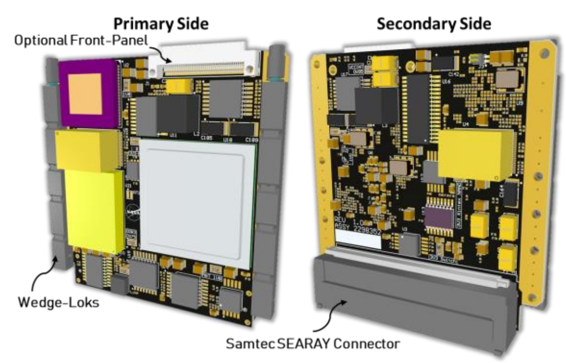

Figure 10: Example Assembled 1U Card with Labelled Components

[4] https://github.com/nasa/cFS

[5] http://suddendocs.samtec.com/testreports/hsc-report-sma_seam-02_seaf-ra_web.pdf

Table 2 provides the compatible card connectors and the corresponding connectors mounted to the backplane. Figure 11 displays models of the varying connector types. Designs should follow manufacturer recommendations for connector overhang to ensure proper engagement with backplane connector.

Table 2: Connectors for 1U Form Factor

Style	Backplane Connector	Card Connector
160 pin	SEAM-40-02.0-L-04-1-A-GP-K-TR	SEAF-40-01-L-04-1-RA-GP-TR
200 pin	SEAM-50-02.0-L-04-1-A-K-TR	SEAF-50-01-L-04-1-RA-TR
400 pin	SEAM-50-02.0-L-08-1-A-K-TR	SEAF-50-01-L-08-1-RA-TR

Figure 11: Model[6] of 400 Pin Card/Backplane Connectors (Left) Backplane Connector (Right) Card Connector

Wedge-Loks

Wedge-Loks are ideal for this system because they are designed to restrain the PCBs through the high vibration and shock environments for spacecraft launch and deployment, while additionally providing a thermal path from the PCB to the chassis wall. Wedge-Loks are preferred over Wedge-Tainers due to improved thermal performance.

The Wedge-Loks can be mounted on either the primary or secondary side of the PCB (see Figure 10), but should be mounted on the opposite side of the Samtec backplane connector. The required Wedge-Lok mount holes should be non-plated through holes to minimize stress and prevent PCB failures. Stitching vias between the top layer and bottom layer chassis copper plane pours can be used under higher thermal loads. The via diameter should be less than 0.025" (0.635 mm), and if possible, offset from directly under the Wedge-Lok to prevent stress failures. Figure 12 displays an example Wedge-Lok. The assembly length dimension compatible with this standard is 2.80" (71.12 mm). The part information is listed in Table 3.

Figure 12: Example Wedge-Lok (courtesy nVent SCHROFF[7])

Table 3: Wedge-Lok Part

Name	Part Number	Description
Series 267	811-2670083	VEN267-2.8ET2LK

While CS^2 describes the specifications for Wedge-Loks, the cards remain compatible with a Wedge-Tainer solution. For the Wedge-Tainer approach, parts listed below in Table 4 have been used.

Table 4: Wedge-Tainer Part

Name	Part Number	Description
Wedge-Tainer Series 340	340L-100S-06EN	Left channel
Wedge-Tainer Series 340	340R-100S-06EN	Right channel

Physical Dimensions

Figure 13 provides the physical dimensions for the PCB card, including mounting hole locations for the Wedge-Loks and 400-pin backplane connector. The areas for Wedge-Loks can include additional holes for conforming to additional board stack-ups.

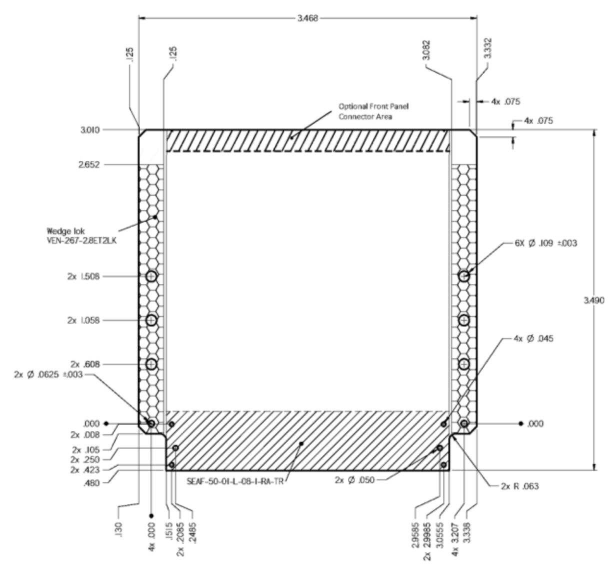

Figure 13: Printed-Circuit Board Dimensions

Figure 14 provides general dimensions for a CS^2-compatible electronics enclosure and highlights

[6]Orange dots indicate polyimide tape for pick and place assembly, models courtesy Samtec (https://www.samtec.com/products/seam)

[7]https://schroff.nvent.com/wcsstore/ExtendedSitesCatalogAssetStore/Attachment/CalmarkBirt cherProductAttachments/Documents/267_DataSheet.pdf

example card spacing. Card pitch is not prescribed by this standard, which allows for cards without tall components to save space by moving closer together on the backplane. The backplane design should ideally consist of only the connectors. This passive design makes the backplane essentially equivalent to harnessing and straightforward to develop.

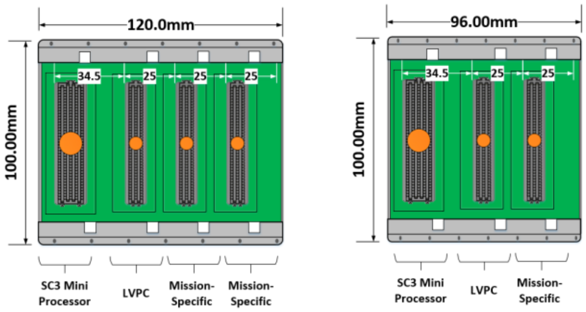

Figure 14: Example Blackplane and Pitch Size

Grounding

Each card has a *main signal ground net* (GND) that is connected to multiple internal copper planes and to signal ground pins on the backplane connector and (optional) front-panel connector. This ground is shared by all cards in the electronics box through their backplane connectors. The backplane has multiple internal copper planes connected to this ground.

Each card has a separate *chassis ground net* for the box chassis (CGND). This connection is established through exposed surface copper under the Wedge-Loks on both edges of the board. The top and bottom layers on the PCB should have exposed copper on the outer edge to provide a thermal and electrical short between the chassis ground (CGND), the Wedge-Lok, and the chassis. This minimizes the thermal resistance as seen from the PCB to the chassis. The front-panel connector, when present, shall have its shell/body connected to CGND.

Separate digital and chassis grounds should be maintained throughout every design unless otherwise required for a specific mission. Each card has optional selective population of a parallel resistor and capacitor in the four corners of the PCB allow for single-point grounding (path between GND and CGND, this should be the only path between the two ground planes) to be maintained and adjusted as needed. A designer must exercise caution when designing an interface board that includes a heatsink to ensure no indirect connection is made between ground and chassis. Similarly, the designer should exercise caution with front-panel connector design to ensure that the shell being in contact with the chassis does not create an inadvertent connection between signal ground and CGND.

Connector Pinouts

There are several configurations of backplane connector I/O for varying pin-connectors. These pinouts were developed as a balance of several design concerns such as signal integrity, I/O density, and ease of routing. Due to the size and complexity of these diagrams they are not suitable for inclusion in this document but can provided upon request.

VI. DEPLOYMENT CONFIGURATIONS

Immediate use of this architecture can be conceived for applications for artificial intelligence (AI), communication and navigation, and finally SmallSat cybersecurity. This section provides brief examples and an initial observation of the configuration for the AI system.

Artificial Intelligence Processing System

A small form-factor dedicated AI processing unit can be strategic for both science and defense applications. Specifically, a modified version of the system described in this section could meet the processing needs described for the Blackjack "Pit Boss" Edge processing node [29]. The most defining challenge for more advanced and capable artificial intelligence on satellites is derived from limitations in the SmallSat platform design. Consequently, these computing restrictions are particularly challenging to machine learning frameworks because a significant amount of progress in deep learning and modern networks has been specifically conducted using graphics processing units (GPUs). Many state-of-the-art network models require high-end GPU devices to run inference, and even more capability is required to train these models. Current space computers would struggle to meet the minimum requirements for complex, deep-learning architectures. Additionally, there are a scarce number of GPUs that have been proven to work in a space environment, while simultaneously meeting the low-power restrictions of SmallSat platforms. Despite these limitations, the proposed SpaceCube system will enable essential AI applications. A proposed configuration, shown in Figure 15, could include a LVPC, a SpaceCube Mini-Z45, and SpaceCube v3.0 Mini.

In [30], the NSF SHREC (Space, High-Performance, and Resilient Computing) Center at the University of Pittsburgh recently examined image classification with convolutional neural networks (CNNs) on their CSP platform using TensorFlow Lite. In detail, they examined CNN architectures designed for mobile applications, including MobileNetV1, MobileNetV2, Inception-ResNetV2, and NASNet Mobile, which were pre-trained on ImageNet. The SpaceCube Mini-Z is an upgrade of that prior platform and is described earlier in this paper.

The design for semantic image segmentation on devices featured in the SpaceCube Mini-Z and SpaceCube v3.0 platforms is presented in [31]. Semantic segmentation is a deep learning algorithm, based on convolutional neural networks (CNNs), that learns to infer dense labels for every pixel of an image. In [31] this application is deployed on a reconfigurable CNN acceleration framework (ReCoN). Semantic segmentation has numerous space applications, from semantic labeling of Earth's features for insights about our changing planet, to monitoring natural disasters, and to gathering intelligence for national security.

The integration of SpaceCube Mini-Z45 and SCv3M interconnected by high-speed interfaces in a CubeSat form-factor can enable two configurations for advanced applications, such as semantic segmentation, necessary for high-performance onboard processing. In one configuration, the SCv3M serves as a co-processing card. The Mini-Z45 can offload massive workloads to the SCv3M for acceleration with minimal communication overhead. In another configuration, the SCv3M serves as a front-end data processor for sensors directly interfaced to this card. In this configuration, the SCv3M can process and convert raw sensor-data into compressed, actionable results or scientific knowledge provided to the Mini-Z45 for downlink to storage. Figure 16 shows the original architecture in [31] deployed on a single SoC device. Figure 17 shows the application deployed onto the AI processor box described in this section.

For an initial demonstration of the combined architecture, two Xilinx development boards (a ZC706 and a KCU105) are used as near hardware equivalent representations of the SpaceCube Mini-Z45 and SCv3M. These boards are connected together with SMA cables to provide an AXI Chip2Chip high-speed interconnect. Semantic Segmentation is then executed on top the ReCoN framework across the combined architecture. Table 5 shows the respective FPGA resource utilization expected for each development board, and Table 6 shows a comparison between the inference performance of the accelerated application. The results demonstrate a massive 1733× speedup over a purely software baseline run on the ARM cores of the ZC706 alone. It should also be emphasized that the performance efficiency will increase with the actual flight designs because the development boards have many peripherals and interfaces that are not included in the SpaceCube designs.

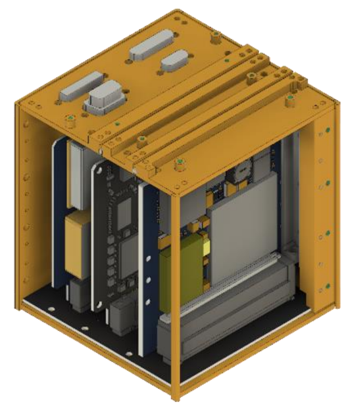

Figure 15: 3-Card 1U CubeSat AI Processing Unit

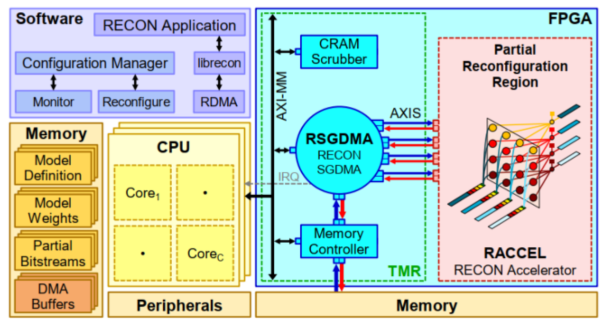

Figure 16: RECON Architecture on Single Zynq SoC or MPSoC

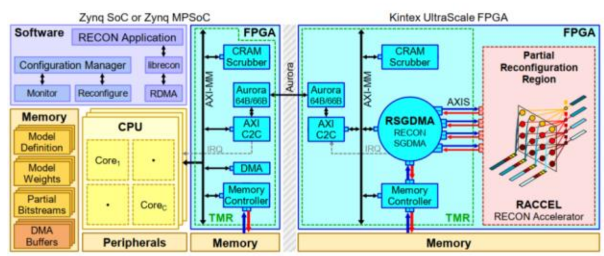

Figure 17: RECON Architecture on Combined Mini-Z45 and SCv3M

Table 5: Resource Utilization of RECON on AI Processing Box Emulator

Resource	SpaceCube Mini-Z45 (ZC706)	SpaceCube v3.0 Mini (KCU105)
LUTs	1.37%	25.59%
FFs	1.19%	19.65%
BRAM (36 Kb)	1.10%	74.67%
DSPs	0.00%	27.66%

Table 6: Performance of RECON on AI Processing Box Emulator

Measurement	SpaceCube Mini-Z45 (ZC706) Fully executed software baseline	AI Processing Box (ZC706/KCU105) Fully accelerated on KCU105
Max Frequency (F_{Max})	800 MHz (PS)	800 MHz (PS) / 215 MHz
Performance (FPS)	0.08	141.74 (**1733×**)
System Power	9.31W	31.58W (9.88+21.7)
Performance/Watt	0.009	4.49 (511×)

ZC706/KCU105; Vivado 2019.2; 515×512 IRRGB (ISPRS Potsdam); INT8 quantization; -O3 optimization; SegNet model (465k weights)

Other Configurations

Two other significant configurations for a reusable payload design are driven by needs in communication and navigation, as well as, SmallSat cybersecurity. While NASA is always focused on incredible science value in missions and experiments, it must also make strides to protect its space systems from cybersecurity threats. With the growing SmallSat industry, it has been noted by defense and research sectors that cybersecurity concerns are often overlooked for spacecraft, and many may be vulnerable to cyberattacks. This emphasis becomes more significant with future plans for constellations enabling capabilities which will feature cross-link communication along with relay links, and NASA has already experienced cybersecurity issues in the past [32]. As described earlier, LunaNet has broad goals for enabling efficient communication and services at the moon. While there are varying aspects of LunaNet that can be addressed by this architecture, [5] does not prescribe a specific solution regarding the proposed hardware for the system. One component for local instrument networking is provided in Figure 18 as a four-card 1U system. Additionally, as a component of addressing cybersecurity concerns, hardware-based cryptography can be implemented in the SpaceCube system. [33] provides several research concepts for enabling system isolation for security inclusive of a SpaceCube-like system with CSPv1s. Several 1U configuration concepts are pictured in Figure 18.

VII. CONCLUSION

The SpaceCube Intelligent Multi-Purpose System is reconfigurable and reusable allowing it to meet future science and defense needs for several mission types. The CS^2 specification allows future developers to design cards to be compatible with the system allowing more combinations of cards to be used to address new mission proposals. This architecture design is advantageous for instruments that can be repurposed to varying science observables without significant changes to the electronics processing cards. This paper has described an architecture that provides high-performance processing, reliability, and the affordability of SWaP-C characteristics intrinsic with a 1U CubeSat form factor, and is immediately relevant for applications in instrument processing, artificial intelligence, communication and navigation, and finally cyber security and encryption.

Acknowledgments

The authors would like to recognize the contributions and support by additional team members of the Science Data Processing Branch Code and Goddard collaborators including Michael Lin, Michael Monaghan, Justin Goodwill, David Wilson, Alan Gibson, Jonathan Boblitt, Luis Santos, Chuck Clagett, Munther Hassouneh, Jason Mitchell, and James Fraction. Additional collaborators include the LunaNet team including Lavida Cooper, Jaime Esper, and Kendall Mauldin. Special thanks to our key sponsors that have supported SpaceCube family development throughout the years including NASA/GSFC Internal Research and Development (IRAD) program and the NASA Earth Science Technology Office (ESTO). The authors also acknowledge support and collaboration with vendors including Javier Valle from Texas Instruments and Dylan Lang at Samtec.

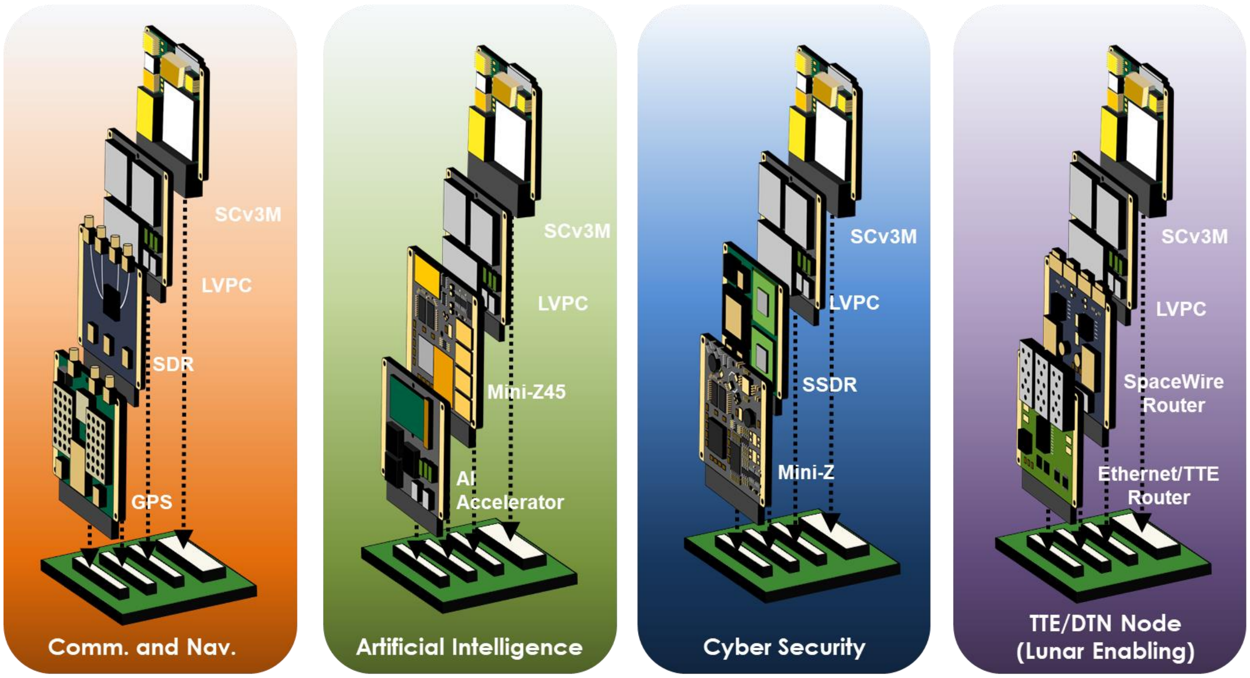

Figure 18 CubeSat 1U Box Configurations for Varying Specializations

References

1. Achieving Science with CubeSats: Thinking Inside the Box, Washington, D.C., USA: National Academy Press, 2016.
2. Hornyak, D., "A Researcher's Guide to: Technology Demonstration, ISS Researcher's Guide Series," NASA ISS Program Science Office, Jul. 6, 2016.
3. Mercer, C., "Small Satellite Missions for Planetary Science," 33rd Annu. AIAA/USU Conf. on Small Satellites, SSC19-WKV-06, Logan, UT, August 3-8, 2019.
4. 2020 NASA Technology Taxonomy, NASA Office of the Chief Technologist, 2020.
5. Israel, D. J., Mauldin, K. D., Roberts, C. J., Mitchell, J. W., Pulkkinen, A. A., Cooper, L. V. D., Jonson, M. A., Christe, S. D., and C. J. Gramling, "LunaNet: a Flexible and Extensible Lunar Exploration Communication and Navigation Infrastructure," IEEE Aerospace, Big Sky, MT, Mar 7 - Mar 14, 2020.
6. Review of the Planetary Science Aspects of NASA SMD's Lunar Science and Exploration Initiative, Washington, D.C., USA: National Academy Press, 2019.
7. Visions and Voyages for Planetary Science in the Decade 2013-2022: A Midterm Review, Washington, D.C., USA: National Academy Press, 2018.
8. Project Thor Team, "Outpacing the Threat with an Agile Defense Space Enterprise," Aerospace Corporation, Sept. 2019.
9. Air Force Space Command, AFSPC Long-Term Science and Technology Challenges, Peterson AFB, CO, USA, 2016.
10. Geist, A., Brewer, C., Davis, M., Franconi, N., Heyward, S., Wise, T., Crum, G., Petrick, D., Ripley, R., Wilson, C., and T. Flatley, "SpaceCube v3.0 NASA Next-Generation High-Performance Processor for Science Applications," 33rd Annual AIAA/USU Conf. on Small Satellites, SSC19-XII-02, Logan, UT, August 3-8, 2019
11. George, A. D. and C. Wilson, "Onboard Processing with Hybrid and Reconfigurable Computing on Small Satellites," Proc. of the IEEE, vol. 106, no. 3, pp. 458-470, Mar 2018.
12. SpaceWorks, "Nano/Microsatellite Market Forecast, 9th Edition," 2019.
13. Guertin, S., "CubeSat and Mobile Processors," NASA Electronics Technology Workshop, June 23-26, 2015.
14. Azarbarzin, A., "MUSTANG Applications Modular Avionics," IEEE Space Computing Conference, 2019 IEEE Space Computing Conference (SCC), Pasadena, CA, USA, 2019.
15. Lin, M., Flatley, T., Geist, A., and D. Petrick, "NASA GSFC Development of the SpaceCube

Mini," 25th Annual AIAA/USU Conf. on Small Satellites, SSC11-X-11, Logan, UT, August 8-11, 2011.

16. Lee, D., Allen, G., Swift, G., Cannon, M., Wirthlin, M., George, J. S., Koga, R., and K. Huey, "Single-Event Characterization of the 20 nm Xilinx Kintex UltraScale Field-Programmable Gate Array under Heavy Ion Irradiation," IEEE Radiation Effects Data Workshop, July 13-17, 2015.
17. Berg, M., Kim, H., Phan, A., Seidleck, C., Label, K., and M. Campola, "Xilinx Kintex-UltraScale Field Programmable Gate Array Single Event Effects (SEE) Heavy-ion Test Report," NASA Electronic Parts and Packaging, 2017.
18. Petrick, D., Gill, N., Hassouneh, M., Stone, R., Winternitz, L., Thomas, L., Davis, M., Sparacino, P., and T. Flatley, "Adapting the SpaceCube v2.0 data processing system for mission-unique application requirements," IEEE Aerospace Conference, Big Sky, MT, June 15-18, 2015.
19. Wilson, C., and A. D. George, "CSP: Hybrid Space Computing," AIAA Journal of Aerospace Information Systems, vol. 15, no. 4, pp. 215–227, Feb 2, 2018. doi: https://doi.org/10.2514/1.I010572
20. Trafford, R., Gorab, N., Smith, T., Fifth, A., Schmalzel, J., Krchnavek, R., and S. Shin, "Wireless Bus Interconnects for Flexible and Reliable CubeSat Signal Integrations," 33rd Annual AIAA/USU Conf. on Small Satellites, SSC19-WKVIII-07, Logan, UT, August 3-8, 2019.
21. Lovelly, T. M. and George, A D., "Comparative Analysis of Present and Future Space-Grade Processors with Device Metrics,"AIAA Journal of Aerospace Information Systems, Vol. 14, No. 3, Mar. 2017, pp. 184-197. doi: 10.2514/1.I010472
22. Wirthlin, M., "High-Reliability FPGA-Based Systems: Space, High-Energy Physics, and Beyond," Proceedings of the IEEE, vol. 103, no. 3, Mar. 2015, pp. 379-389.
23. Campola, M., "Taking SmallSats to the Next Level – Sensible Radiation Requirements and Qualification that Won't Break the Bank," 32nd Annual AIAA/USU Conference on Small Satellites, SSC18-WKV-01, Logan, UT, Aug 4-9, 2018.
24. Lee, D., Allen, G., Swift, G., Cannon, M., Wirthlin, M., George, J. S., Koga, R., and K. Huey, "Single-Event Characterization of the 20 nm Xilinx Kintex UltraScale Field-Programmable Gate Array under Heavy Ion Irradiation," IEEE Radiation Effects Data Workshop, July 13-17, 2015.
25. Berg, M., Kim, H., Phan, A., Seidleck, C., Label, K., and M. Campola, "Xilinx Kintex-UltraScale Field Programmable Gate Array Single Event Effects (SEE) Heavy-ion Test Report," NASA Electronic Parts and Packaging, 2017.
26. Wilson, C., "SpaceCube On-Board Processor Update," GSFC Annual SCaN Technology Review, November 5, 2019.
27. "Microblaze Triple Modular Redundancy (TMR) Subsystem v1.0," PG-286, Xilinx, 2018.
28. Wilson, A. and M. Wirthlin, "Neutron Radiation Testing of Fault Tolerant RISC-V Soft Processors on Xilinx SRAM-based FPGAs," 12th Space Computing Conference, July 30 – August 1, 2019.
29. DARPA, Blackjack Pit Boss Broad Agency Announcement, Arlington, VA, USA: Tactical Technology Office, HR001119S0012, Apr. 1, 2019.
30. Manning, J., Gretok, E., Ramesh, B., Wilson, C., George, A. D., MacKinnon, J., and G. Crum, "Machine-Learning Space Applications on SmallSat Platforms with TensorFlow," 32nd Annual AIAA/USU Conference on Small Satellites, SSC18-WKVII-03, Logan, UT, Aug 4-9, 2018.
31. Sabogal, S., George, A. D., and G. Crum, "ReCoN: Reconfigurable CNN Acceleration for Space Applications A Framework for Hybrid Semantic Segmentation on Hybrid SoCs," 12th Space Computing Conference, July 30 – August 1, 2019.
32. Malik, W. J., "Attack Vectors in Orbit: The Need for IoT and Satellite Security," RSA Conference 2019, San Francisco, CA, March 4-8, 2019.
33. Skowyra, R., "Enabling Trustworthy Remote Recovery with SeL4," SeL4 Summit, Herndon, VA, Nov. 14-16, 2018.