SSC23-P1-32

NASA SpaceCube Next-Generation Artificial-Intelligence Computing for STP-H9-SCENIC on ISS

Alessandro Geist, Gary Crum, Cody Brewer, Dennis Afanasev, Sebastian Sabogal, David Wilson, Justin Goodwill, James Marshall, Noah Perryman, Nick Franconi, Timothy Chase Jr, Travis Wise, Sarah Flanagan, Michael Monaghan, Jeff Hosler, Leyland Young, Judy Kuntz, Lindsey Seo, Joseph Schreck, Tracy Price, Munther Hassouneh, Steven Kenyon, Samuel Price, Christopher Wilson
NASA Goddard Space Flight Center
8800 Greenbelt Rd, Greenbelt, MD, 20771; 301-286-8419
christopher.m.wilson@nasa.gov

Nicole Thai, Madison Hobbs, Andrew Keene, Sammy Lin, Michael Nemerouf,
The Aerospace Corporation
2310 E. El Segundo Blvd, El Segundo, CA 90245; 310-336-5527
nicole.thai@aero.org

Evan T. Kain, Francisco O. Viramontes, Tyler M. Lovelly
Air Force Research Laboratory, U.S. Space Force
3550 Aberdeen Ave SE, Kirtland Air Force Base, NM 87117; 505-846-4746
tyler.lovelly.1@spaceforce.mil

ABSTRACT

Recently, Artificial Intelligence (AI) and Machine Learning (ML) capabilities have seen an exponential increase in interest from academia and industry that can be a disruptive, transformative development for future missions. Specifically, AI/ML concepts for edge computing can be integrated into future missions for autonomous operation, constellation missions, and onboard data analysis. However, using commercial AI software frameworks onboard spacecraft is challenging because traditional radiation-hardened processors and common spacecraft processors cannot provide the necessary onboard processing capability to effectively deploy complex AI models. Advantageously, embedded AI microchips being developed for the mobile market demonstrate remarkable capability and follow similar size, weight, and power constraints that could be imposed on a space-based system. Unfortunately, many of these devices have not been qualified for use in space. Therefore, Space Test Program - Houston 9 - SpaceCube Edge-Node Intelligent Collaboration (STP-H9-SCENIC) will demonstrate inflight, cutting-edge AI applications on multiple space-based devices for next-generation onboard intelligence. SCENIC will characterize several embedded AI devices in a relevant space environment and will provide NASA and DoD with flight heritage data and lessons learned for developers seeking to enable AI/ML on future missions. Finally, SCENIC also includes new CubeSat form-factor GPS and SDR cards for guidance and navigation.

I. INTRODUCTION

Artificial Intelligence (AI) and Machine Learning (ML) are among the fastest growing research areas in both academia and industry. Researchers and companies are demonstrating both substantial and breathtaking applications of AI/ML for terrestrial applications, including autonomous vehicles, social bots, virtual assistance, media manipulator tools, and strategic game systems. Developers seeking to expand their customer bases have begun to create smaller, more power-efficient AI microchips and accelerators targeting the mobile and embedded systems markets. These advances in AI algorithms and custom accelerator electronics can also be harnessed to enable numerous breakthrough capabilities in the space domain, including autonomous swarm/constellation management, reactive health and status monitoring, security analytics, anomaly detection, and responsive onboard data analysis.

While there is exciting potential and many benefits for deploying advanced AI applications in space, using commercial AI software frameworks onboard spacecraft is challenging because traditional radiation-hardened (rad-hard) processors and other common spacecraft processors cannot provide the necessary onboard computing resources and processing capability to effectively deploy complex AI models. Therefore, onboard science data-processing pipelines are substantially restricted to simpler AI/ML approaches. This limitation makes small, low-power AI microchip architectures attractive for space missions where the application-specific design enables both high-performance and power-efficient computing for state-of-the-art AI applications. However, while these embedded AI microchips are cutting-edge embedded solutions, there are few examples of space deployment that demonstrate that these devices are resilient to a harsh space radiation environment.

To demonstrate the capability of inflight cutting-edge AI applications on multiple space-based devices for next-generation onboard intelligence, a new flight experiment, called the Space Test Program-Houston 9-SpaceCube Edge-Node Intelligent Collaboration (STP-H9-SCENIC), was recently deployed on the International Space Station (ISS). The experiment monitors each device's availability and susceptibility due to radiation effects. Using a CubeSat-sized hyperspectral imager, SCENIC collects an extensive image archive of Earth's surface required to train data-driven deep neural networks and perform real-time generation of data products for downlink to the scientific community. This experiment also demonstrates and evaluates NASA's next-generation CubeSat-sized, rad-tolerant, high-performance computer, known as SpaceCube v3.0 Mini, which is supplemented with fault-tolerant computer-architecture design and mitigation strategies. Additionally, SCENIC provides flight validation of new CubeSat-sized guidance and navigation cards to be deployed on future NASA missions. SCENIC is one of eight experiments on the STP-H9 payload, which operates on the International Space Station (ISS) Japanese Experiment Module - Exposed Facility (JEM-EF). STP-H9-SCENIC is developed by NASA Goddard Space Flight Center (GSFC) with collaborators from The Aerospace Corporation, sponsored by the Air Force Research Laboratory (AFRL), and integrated into the STP-H9 payload by the Department of Defense (DoD) STP.

One of the challenges for novel, deep-learning models is having adequate training data. SCENIC is able to provide the science community with a publicly available hyperspectral dataset from a small form-factor hyperspectral imager to train data-driven deep neural networks. This capability provides researchers and scientists with representative datasets to train and develop their own models for future advanced missions. This research also characterizes the performance of embedded AI devices, including the Intel Movidius Myriad X, Google Coral Edge TPU (Tensor Processing Unit), and Xilinx Deep Learning Processor Unit (DPU) in a relevant space environment. SCENIC provides flight heritage data and lessons learned for spacecraft developers seeking to enable future complex missions with modern deep-learning capabilities.

This future capability for AI/ML can be disruptive and enabling for both science and defense mission and instrument concepts. For NASA science objectives, intelligent and autonomous systems have been described in multiple crucial agency-guiding documents such as the decadal surveys. In *Thriving on our Changing Planet*, the Earth science decadal survey [1], state-of-the-art real-time systems using advanced ML techniques are described, as well as improvements to rapid disaster response scenarios such as fires and flood detection. In *Origins, Worlds, and Life*, the decadal survey for planetary science and astrobiology [2], a driving consideration is onboard autonomy for real-time planetary spacecraft control, management, fault detection and recovery, navigation, orbit insertion, landing, and driving, especially for future missions under severe communication or environmental constraints. Finally, the NASA 2020 Technology Taxonomy [3] also includes several primary motivations for SCENIC development including AI/ML, commercial device usage, and improvements in onboard computing. AI/ML is specifically discussed in TX05.5.1 Cognitive Networking, TX10.1 Situational and Self-Awareness, TX10.2 Reasoning and Acting, and finally TX11.4.8 for Edge Computing. Commercial devices and onboard processor improvements are described in TX08.1.2 Electronics, TX02.1.5 High-Performance FPGAs, and TX02.2.8 Use of Commercial-of-the-Shelf (COTS) Technologies.

For defense applications, AI has also been recently highlighted as a key priority. One essential reference is the United States Space Force Long-Term Science and Technology Challenges which emphasizes AI, ML, and autonomy, in an extensive list of use cases both in-situ and ground based, significantly for edge computing, data triage, decision making, analysis, and classification objectives [4]. Other agencies such as the National Geospatial-Intelligence Agency (NGA) have a vested interest in AI applications especially for data analysis with the advent of increased satellite imagery from commercial vendors [5] along with their acquisition of the signature DoD AI/ML tool Project Maven[1]. Finally, DARPA has continued their development of Blackjack[2] which would exhibit a broad host of AI/ML features for low Earth orbit (LEO) resiliency, specifically "mission-level autonomy software and demonstrate autonomous

[1]https://www.c4isrnet.com/intel-geoint/2022/04/27/intelligence-agency-takes-over-project-maven-the-pentagons-signature-ai-scheme/

[2] https://www.darpa.mil/program/blackjack

orbital operations including on-orbit distributed decision processors."

II. BACKGROUND

The following sections describe significant background concepts related to the development and motivations of the SCENIC experiment. This section also describes design concepts and introduces the Department of Defense Space Test Program which enabled this rapid mission-development timeline.

Onboard Computing Limitations

Despite consistent investment into onboard computing performance, electronics remain susceptible to the hazards of the space environment, and testing and qualification for reliable devices and systems is an arduous and gradual process. There is a strident gap in computing performance between traditional rad-hard space-computing devices and the latest commercial embedded devices. For example, the Samsung Galaxy phones have commonly used Qualcomm Snapdragon SoCs (System-on-chip). In a comparison provided in [6], the Snapdragon 855 outperformed the current state-of-the-art rad-hard processors, the Frontgrade GR740 and BAE Systems RAD750 by 75× and 276×, respectively, in a DMIPS benchmark. The new High-Performance Spaceflight Computing (HPSC) processor being developed by Microchip Technology Inc.[3] will be a milestone improvement from the previous legacy rad-hard systems; however, the HPSC will not be available for instruments and missions for years. In the interim between rad-hard developments, NASA and other organizations have adopted hybrid computing approaches combining both commercial and rad-hard components to attain the advantages from both technologies, as described in [7]; however, the computational capability required for modern AI/ML frameworks is significant, even for these hybrid systems. Advantageously, embedded AI microchips being developed for the mobile market demonstrate remarkable capability with architectures optimized for AI/ML workloads that, with qualification and testing, can be adopted for the space domain, with examples demonstrated on SCENIC.

Related Research

There are only sparse examples of current and historic missions that include edge computing devices. This work is not an exhaustive survey of these missions, although several prominent examples are highlighted in this section. There are a growing number of research papers published highlighting AI/ML methods and techniques on ground-development platforms (e.g., small-size development boards); however, some devices have no flight heritage and frequently no corresponding flight-equivalent card. Highlighted examples of onboard embedded devices include the Intel Movidius Myriad 2 on PhiSat-1 [8], the Myriad X on the Unibap iX5-100 in HyTI [9] and WILD RIDE [10], the Google Coral TPU SoM on the CoralReef Payload of The Aerospace Corporation's Slingshot 1[4] and finally, the JPL experiments with the Myriad X inside the ISS [11].

CubeSat Card Specification

The CubeSat Card Specification (CS2) was developed at the NASA Goddard Science Data Processing Branch to establish a common template for design compatibility between CubeSat-sized 1U cards. This specification, described in [12], enables electronics developers to create 1U form-factor cards that can be rapidly integrated into NASA CubeSat systems. This specification provides pinout configurations, mechanical and electrical specifications, and other design recommendations. The specification has been provided to several NASA groups and The Aerospace Corporation, who created their own cards that have been integrated for instruments and payloads, including SCENIC. Currently, NASA Goddard has developed several compliant cards with varying functions (e.g., single-board computers, power cards, analog cards, I/O cards), many featured in the SCENIC hardware architecture, that allow developers to mix-and-match cards within the catalog to build new systems for missions and instruments. Following these specifications, the NASA SpaceCube Intelligent Multi-Purpose System (IMPS) uses the catalog cards expanding from a baseline configuration to provide small, powerful, and reusable processing payloads for NASA missions and instruments. The baseline configuration includes a miniaturized processor card, typically the SpaceCube v3.0 Mini, a power card, and a backplane. From this basic design, the system can be further augmented with additional miniaturized slices to accomplish different objectives for onboard instrument processing, with prominent computing domains including artificial intelligence, communication and navigation, and finally cyber security / encryption. An example design is pictured in Figure 11 for the NavCube-mini stand-alone design.

---

3 https://www.nasa.gov/press-release/nasa-awards-next-generation-spaceflight-computing-processor-contract

4 https://aerospace.org/article/slingshot-platform-fast-tracks-space-systems-using-modularity-and-open-standards

Goddard SpaceCube Family of Processors

SpaceCube is a family of NASA Goddard-developed space processors that established a hybrid-processing approach that provides a blend of the best advantages of both commercial and radiation-hardened technologies to yield unparalleled next-generation systems. These systems commonly combine Xilinx FPGAs with rad-hard supporting power circuity. SCENIC includes the SpaceCube v3.0 Mini [12], the latest generation card.

Space Test Program (STP) and STP-H9

The U.S. DoD STP is a prominent facilitator of space science and space technology providing spaceflight opportunities for technology experiment demonstration payloads [13]. The STP-Houston office is the sole interface to NASA for all DoD payloads on the ISS [14]. STP-H9-SCENIC is one of eight experiments on the STP-H9 payload, which operates on the ISS JEM-EF, pictured in Figure 1.

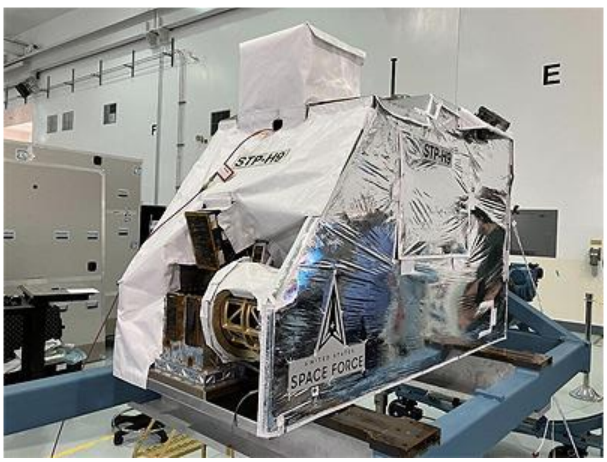

Figure 1: STP-H9 payload inside NASA's Space Station Processing Facility at Kennedy Space Center[5]

III. APPROACH

SCENIC is an ISS-based testbed to evaluate and validate AI/ML technology on a next-generation FPGA platform and custom AI microchip platforms in space. While AI research proposes attractive advantages for future missions, there are very few in-flight mission heritage examples. SCENIC was designed to give science and defense mission developers realizable use-case examples of onboard capability to demonstrate the disparity between ground-based and space-based processing capabilities. SCENIC's primary challenges were a constrained mission budget, a rapid development timeline to meet launch schedule, and finally a software framework that would allow future updates to be completed post-launch due to schedule pressures.

Primary Objectives

SCENIC has several primary experiment objectives for mission success. The first is to demonstrate and evaluate commercial AI microchip technology (i.e., Intel Movidius X, and Google Coral Edge TPU) for radiation characterization in a relevant space environment. There are many novel embedded systems constantly in development; however, only a few systems have any space heritage or example missions. The second is to collect an extensive image archive of hyperspectral Earth observations required to train data-driven deep neural networks and perform real-time generation of data products for downlink to information subscribers. Notably, many of these data-driven models require a significant amount of training data. SCENIC's hyperspectral image (HSI) sensor is CubeSat-sized and would be representative for other missions desiring to incorporate a similar sensor. This dataset would be provided publicly to the academic and research community to train and prototype deep-learning models for future missions. Finally, this mission seeks to demonstrate and evaluate NASA's next-generation CubeSat-sized, rad-tolerant, high-performance computer known as SpaceCube v3.0 Mini including fault-tolerant computer architecture design and mitigation strategies.

Secondary Objectives

SCENIC also includes several secondary experiment objectives. One objective is to upload new mass-less payload experiments (software applications) for additional selected defense applications (e.g., semantic segmentation). SCENIC is also one of the first experiments to adopt several NASA core Flight System (cFS) applications (discussed in Section V) and advance the technology readiness level (TRL) of those applications. Finally, SCENIC also includes new CubeSat form-factor GPS and SDR cards for guidance and navigation for demonstration on ISS, also for TRL increase.

[5] https://spaceflightnow.com/2023/03/13/u-s-military-experiments-hitching-ride-to-space-station-on-spacex-cargo-ship/

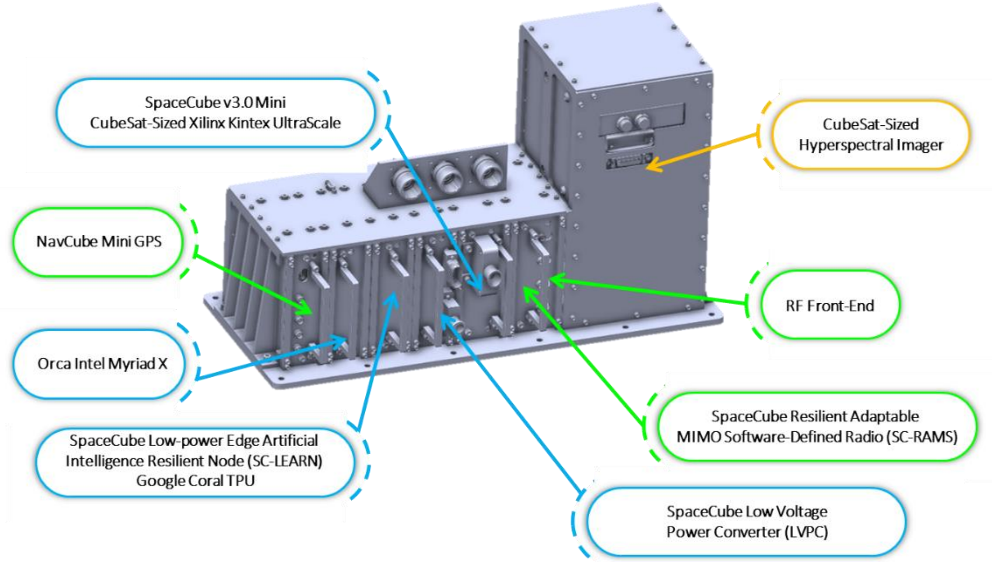

Figure 2: SCENIC Illustrated Electronic Components

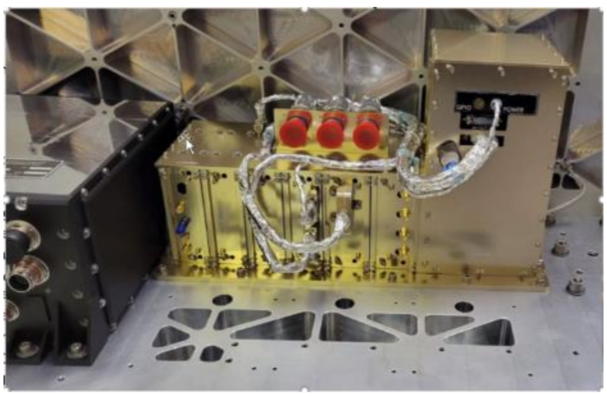

Figure 3: STP-H9-SCENIC installed into STP-H9 NADIR plate

IV. HARDWARE ARCHITECTURE

SCENIC is roughly a 4U CubeSat form-factor (40cm × 10cm × 10cm) avionics payload. It is designed to include a primary backplane card (i.e., a motherboard) supporting card slots in a ~2U card stack with another ~2U reserved for camera electronics, lens, and supporting harnessing. SCENIC is installed within STP-H9 and a circular cutout was positioned in the NADIR plate of STP-H9 for the HSI sensor's field-of-view. There are eight cards in the design described in this section along with the HSI sensor.

The overall hardware architecture is featured in Figure 2 and the post-delivery installation pictures appear in Figure 3. The flight box exposes three external circular connectors used for data, power, and debug interfaces. The debug interface is only used during integration and testing and is capped for flight. SCENIC also includes 5× Global Positioning System (GPS) antennas (two Zenith Antennas and 3×1 Nadir Array) and an external filter/LNA (low-noise amplifier) displayed in Figure 4.

Figure 4: GPS LNA/Filter Assembly (Left) and Antenna (Right)

SpaceCube v3.0 Mini

The SpaceCube v3.0 Mini (SCv3M) processor card is a 1U CubeSat-sized single-board computer described in [12]. The key technology featured in the SpaceCube v3.0 Mini is the AMD-Xilinx Kintex Ultrascale (20-nm FPGA). For memory resources, the design includes 2 GB (72-bit) DDR3 SDRAM memory and 2× 16 GB NAND flash storage. For fault tolerance, the SCv3M incorporates a radiation-hardened watchdog Microchip FPGA and supplementary FPGA-based mitigation techniques. For external interfaces the board provides

12× GTH MGTs, 1.8V LVDS pairs, 3.3V GPIO, and an assortment of RS-422, UART, SelectMAP, and JTAG interfaces. This design is pictured in Figure 5.

The SCv3M instantiates a soft-core MicroBlaze processor and acts as the central communication interface of all the cards within SCENIC to the STP-H9 payload processor. For SCENIC, the 85-pin NANO front-panel connector is used to connect to the HSI sensor using Camera Link. The flight software and FPGA architecture are described in Sections V and VI.

Figure 5: SpaceCube v3.0 Mini 1U Kintex UltraScale CubeSat Single-Board Computer Primary Side

Low-Voltage Power Converter (LVPC)

As described previously, common SpaceCube solution designs include the LVPC to generate the secondary voltage rails for the rest of the system. This card provides clean, isolated secondary voltages (i.e., 3.3V, 5V, 12V) along with up to 8× configurable switched services. On SCENIC, 6× switched services are populated to allow SCv3M to individually control the power state of other cards in the system. The LVPC receives a 28V input from the STP-H9 power system and passes it through an electromagnetic interference (EMI) filter to the secondary voltage regulators.

The LVPC also incorporates 2× 8-channel, 12-bit ADCs (Analog Digital Converter) for monitoring for telemetry through a Serial Peripheral Interface (SPI) interface. Included in the telemetry is the 12V, 5V, and 3.3V voltage and current along with two thermistors.

Figure 6: SpaceCube v3.0 Low-Voltage Power Converter Card Secondary Side

Finally, the LVPC also has a 31-MDM front-panel connector which provides an interface for the RS-422 transceivers and passes through JTAG signals for the SpaceCube v3.0 Mini. This connector is routed from the front-panel of the system up to the circular connectors on the top-body of SCENIC to match the connectors prescribed for STP-H9 payloads. The LVPC is pictured in Figure 6.

SC-LEARN

Goddard has developed a CubeSat Google Coral Edge TPU-based processor card, known as the SpaceCube Low-power Edge Artificial Intelligence Resilient Node (SC-LEARN), which is built to NASA's CubeSat Card Specification (CS2) for integration into SmallSat systems. The 1U SC-LEARN card features three Coral Edge TPU Accelerator Modules. Independently, each of these Coral Edge TPU Accelerator Modules is a multi-chip device that includes the Edge TPU accelerator ASIC, power circuitry, and an internal reference clock. For fault-tolerance, each of the three Coral Edge TPUs include individual load switches controlled by the host processor (SCv3M on SCENIC). Additionally, the supporting circuitry for power, reference clocks, and analog-to-digital conversion is radiation-hardened and designed for flight applications. Extensive details for the SC-LEARN are presented in [15] and a picture is displayed in Figure 7.

Figure 7: SpaceCube Low-power Edge Artificial Intelligence Resilient Node (SC-LEARN) with Thermistors Populated

Planned concept-of-operations for SC-LEARN on SCENIC as one of the primary AI microchips will include uploading and deploying new AI/ML applications for hyperspectral images. This process will involve reconfiguring the onboard FPGA for post-capture processing, reconfiguration of the SC-LEARN's cFS application parameters, and retraining the AI/ML models. As an overview, SCENIC will need to downlink a significant quantity of captures to compose a training dataset from the HSI sensor. Using this dataset, the development team will continuously retrain models for HSI applications on the ground. Finally, the team will convert the models to TensorFlow Lite models compiled for the Edge TPU, upload the TensorFlow Lite models to the SCv3M processor, and deploy on the Edge TPU.

Orca

The Orca payload was developed by the Embedded and Specialized Computing Department (ESCD) at The Aerospace Corporation to demonstrate the capabilities of the Intel Movidius Myriad X Vision Processing Unit (VPU) for AI experimentation in space. ESCD is a group that focuses on evaluation and advancement of accelerated computing technologies in space. The group has prior experience integrating AI solutions-at-the-edge with modules such as the Intel Movidius Myriad 2 (Myriad X predecessor), Intel neuromorphic processors, compute-in-memory chips, and the Google Coral TPU, among others.

The Myriad X features a dedicated on-chip accelerator for executing accurate imaging neural network inferencing applications in real time at-the-edge [16]. These features are significant for evaluation on SCENIC. The developed payload, codenamed Orca, enables batch processing at-the-edge for custom AI model pipelines compatible with ML frameworks such as Keras, Tensorflow, and OpenVINO.

Figure 8: The Aerospace Corporation Orca Myriad X EDU Primary Side

Orca is the custom CubeSat-size card solution to integrate an Intel Myriad X to the NASA CubeSat Card Specification (CS2) to be compatible with other Goddard designs. The Engineering Development Unit (EDU) board is shown in Figure 8. Connectivity to the Intel Myriad X on Orca is through the DepthAI EMB1098 System on Module (SoM) developed by Luxonis. The Myriad X interfaces through the SPI for model inference processing. The NOR flash memory on the DepthAI SoM is exposed through the Quad-SPI (QSPI) for firmware and model pipeline updates. Orca demonstrates a hardware multiplexing technique for switching between the DepthAI SoM SPI and QSPI interfaces to a single Linux SPI device node accessible to the flight software application. An Atmel ATSAMV71Q Automotive Grade ARM Cortex-M7 microcontroller running at 300 MHz with 384 KB of RAM operates as an onboard watchdog to the DepthAI SoM. Firmware on the microcontroller will monitor the voltage and current of Orca in real-time as telemetry to the ground system. The firmware will also control the enable and reset signals for the power circuitry that will boot the DepthAI SoM. The microcontroller is designed for commercial-to-space development transition and has the potential to enable features such as Ethernet communications on Orca to allow future development beyond STP-H9.

SC-RAMS and RF-FE

The SC-RAMS (SpaceCube Resilient Adaptable MIMO Software-Defined Radio) was developed to create a scalable Size, Weight, Power, and Cost (SWaP-C) optimized multi-input multi-output (MIMO) Software-defined Radio (SDR) architecture for Goddard SmallSat and CubeSat missions. The RAMS architecture is designed with a noise optimized power system, clocking configurability for both remote sensing and navigation applications, and MIMO capabilities for communication systems at S-band and X-band.

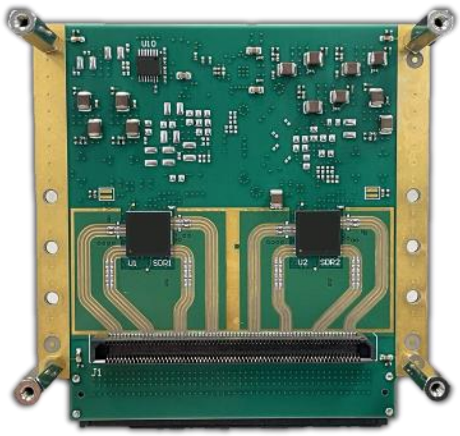

Figure 9: SpaceCube Resilient Adaptable MIMO Software-Defined Radio (SC-RAMS) Primary Side

The proposed SDR was also developed following the CS2 design template that provides the required modularity through simplified backplane development and separation of onboard processor from SDR. Decoupling onboard processor and SDR is advantageous because the SDR is isolated from noise introduced by switching point-of-load regulators and high-speed memories required by processors. The SDR architecture is extensively described in [17] with the card displayed in Figure 9.

For SCENIC, the Radio Frequency – Front End (RF-FE) is designed for a Global Navigation Satellite System Reflectometry (GNSS-R) passive sensing radar experiment. In the SCENIC design, the RF-FE has a connector that mates into SC-RAMS instead of the backplane. SC-RAMS has two Analog Devices AD9361 transceivers configured with four differential receivers configured to amplify L1C and L5 GNSS signals. The RF-FE amplifies reflected L1C and L5 GNSS signals where each GNSS signal requires approximately 55 dB of gain through external low noise amplifiers and internal variable gain amplifiers on each AD9361. The multi-band antennas included on SCENIC are Haigh-Farr wide-band L1/L2/L5 GNSS Antennas with one zenith facing antenna to collect a reference signal and the 3×1 nadir antenna array for reflectometry. The RF-FE is pictured in Figure 10.

Figure 10: Radio Frequency – Front End (RF-FE) Primary Side

NavCube3-mini (NC3m) GNSS Receiver

GSFC has an extensive history building space GPS/GNSS receivers. Starting in the early 2000's GSFC began a program to develop a GPS receiver to specifically target high-altitude applications. This receiver, known as *Navigator*, was designed with high-sensitivity, fast signal acquisition, and an integrated navigation filter, the Goddard Enhanced Onboard Navigation System (GEONS) [18], as key features. Navigator is a mission-enabling technology for the extremely challenging Magnetospheric Multi-Scale Mission (MMS) [19] that was launched in March 2015. It has met or exceeded its performance requirements throughout the mission's duration. The MMS mission has a highly elliptical orbit with current apogee at 185000 km (29 Earth radii), approximately halfway to the moon. Navigator also serves as a critical navigation sensor for the Global Precipitation Measurement Mission (GPM), that has operated in a LEO since February 2014. The heritage Navigator is a rad-hard single frequency receiver that tracks GPS L1 C/A signals. A second-generation version of the Navigator, the NavCube 2.0, was developed using the SpaceCube v2.0 as a flight platform [20]. The NavCube 2.0 introduced a dual-frequency GPS capabilities while achieving modest improvements in SWaP over the heritage Navigator. The NavCube 2.0 was demonstrated on the Space Test Program – Houston 6 (STP-H6) payload in 2019.

Goddard's latest generation GPS/GNSS receiver, the NavCube3-mini (NC3m), uses the SpaceCube v3.0 processor and adds a GNSS RF card in a small form factor resulting in a compact unit that achieves significant further reduction in SWaP compared to previous generations. The NC3m GNSS receiver has performance enhancements over previous generations and is currently capable of tracking GPS L1 C/A and L2C (or L5) signals, with Galileo E1/E5a tracking under development. The standalone version of the NC3m includes a custom backplane and enclosure and has been developed and tested to TRL6 depicted in Figure 11.

The NC3m, under development since 2020, targets onboard navigation and timing as well as science applications in all orbit regimes, including high-altitude applications like geostationary orbit (GEO) and highly elliptical orbit (HEO), and especially lunar applications which are expected to expand rapidly in the near future. After the success of the MMS mission, studies conducted by the NC3m team and others demonstrated that GNSS navigation at the moon can provide excellent on-board real-time performance and potentially provide cost and ground network loading benefits.

For SCENIC, a development version of the NC3m is integrated into the payload. The GNSS RF card has been integrated into the main avionics box and the GNSS software and firmware are hosted on the SCv3M. An external GNSS antenna and front-end electronics are also provided. SCENIC will be the first in-orbit flight demonstration of the NC3m receiver technology. It provides an opportunity to demonstrate basic functions, exercise command and control of the software, act as a testbed for software and firmware updates, and increase the receiver's TRL in preparation for future missions.

Figure 11: NC3m GNSS Receiver Standalone Version

Headwall Hyperspec MV

Headwall Photonics is a leading industry supplier for remote sensing hyperspectral imagers. Their Hyperspec MV hyperspectral system is extremely small, and Goddard commissioned Headwall Photonics to develop a ruggedized light-weight space application version for SCENIC. This design also features a Camera Link connector required for the intended application, where SCv3M could be used in place of the typically included ground-based hardware/software package. While several options were considered, there are no equivalent sensors that meet all the constraining size, weight, power, and connector-type required for the intended payload except this specific model. The sensor includes the key spectral VNIR 400-1000 nm range with 270 spectral bands and 640 spatial pixels (frame resolution width×height of 270×640) essential for science goals.

Backplane

All of the flight cards, except the RF-FE card, insert directly into a single, backplane interconnect board, analogous to a motherboard in a conventional computer system. This backplane, like basic harnessing, encapsulates all interfaces between cards through the backplane connectors and distributes power rails provided by the LVPC. Backplane designs are typically composed of all passive components; however, for SCENIC, some capabilities were included on the backplane that were not integrated onto other individual cards. Significantly, this included an ADC and thermistors around the card edges, designed to monitor varying internal box locations within the experiment. Finally, camera power was also distributed through a set of connectors to the top of the electronic enclosure and routed to the HSI sensor. This backplane design is pictured in Figure 12.

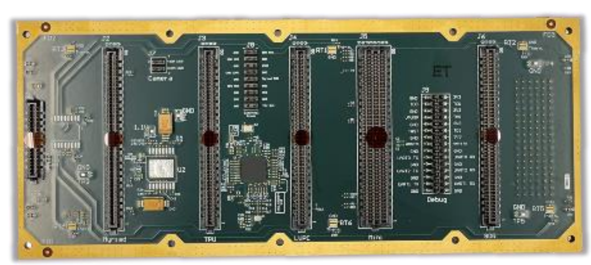

Figure 12: SCENIC EDU Backplane Design

V. SOFTWARE ARCHITECTURE

The SCENIC software team prioritized using open-source software, reusability, and the ability to update code post-launch in order to meet the mission goals. The team and development cycle benefited from extensive previous experience with the open-source tools including cFS and COSMOS, which greatly reduced on-boarding time.

Flight Software

SCENIC's software is based on NASA's open-source flight-software framework, core Flight System[6], version 7.0 (Caelum). NASA cFS provides several services and applications that enables users to add functionality without significant modifications to the framework. Key services include the ability to reload applications while cFS is executing, a publish/subscribe communication architecture (the software bus), and a platform abstraction layer that enables portability for users to develop and test on both flight hardware and development workstations. SCENIC further extended this functionality with Goddard-supported cFS applications to provide file uplink/downlink, application monitoring, and the ability to combine multiple message types into a single packet for downlink.

The SCv3M builds upon cFS by providing a set of applications and libraries specific to its hardware interfaces and onboard subsystems. These applications and libraries may be reused in future SpaceCube-based missions, therefore SCENIC served as a valuable prototyping testbed for future development while also benefiting from the fully assembled and available hardware FlatSat (for other similar project architectures).

This reusability goal presented a significant challenge for developing a flexible code base that could readily share feature updates and bug fixes across multiple missions with diverse hardware. For example, the missions may have varying numbers of ADCs or power switched services to control. This consideration is addressed by ensuring each application incorporated a flexible configuration through either cFS tables, custom header files, or configuration files. The existing cFS build feature of overwriting tables was extended with files of the same name found in the mission `defs` directory to include header files. This modification allowed SCENIC to redefine application message IDs and use custom `#define` directives to change behavior, without necessitating modification of original application source.

SCENIC's approach to reusability is most evident in the Command Ingest (CI) and Telemetry Output (TO) applications. These applications are essential to almost every mission because they provide access to the cFS software bus from external sources and interfaces. The configuration is often changed within a single mission throughout the development and testing cycle. For example, SCENIC uses TCP/IP sockets for testing on developer workstations and a serial interface while developing for the flight hardware. To further exacerbate complexity, the packet format used may also change. Notably, early development used CCSDS (Consultative Committee for Space Data Systems) packets, where the protocol was updated to imitate the ISS interface to SCENIC, and finally modified to communicate with the HOSC (Huntsville Operations Support Center).

JSON-based configuration files and parsing were implemented to facilitate support for multiple configurations and switching between them. The protocols are implemented in C and specified by the configuration files. The CI/TO applications implement modules to support different hardware interfaces such as Ethernet, serial, and SpaceWire. SCENIC has proved the advantages of this approach, which will continue to improve with future mission iterations.

Table 1 shows the list of Goddard, SCv3M, and mission-specific applications. Each application is designed to be operating system (OS) agnostic, with libraries used to implement operating system-specific behavior such as communicating with hardware drivers. These libraries have been omitted for brevity.

Table 1: List of Goddard, SpaceCube v3.0, and mission specific applications

Name	Purpose	Source
CF	File uplink/downlink via CI/TO	GSFC
FM	Basic file operations	GSFC
HK	Combine telemetry packets	GSFC
HS	Monitor application status	GSFC
SCH	Schedules and sends commands	GSFC
Shell	Execute Linux shell commands	GSFC
ADC	Read ADC telemetry	SCv3M
CI	Command Ingest	SCv3M
RHM	Interface with Rad-Hard Monitor	SCv3M
SCM	Record processor resets, startup	SCv3M
SSC	Control LVPC power switches	SCv3M
TO	Telemetry output	SCv3M
CLFG	Interface with Camera	SCENIC
FTDP	Implement STP file transfers	SCENIC
IH	Control image pipeline	SCENIC
NavCube	Interface with NavCube card	SCENIC
ORCA	Interface with Orca card	SCENIC
SDR	Interface with SDR card	SCENIC
TPU	Interface with TPU card	SCENIC

[6] https://github.com/nasa/cfs

SCENIC-Specific Software

SCENIC cFS contains all the applications necessary to communicate with the HSI sensor, Orca Myriad, SC-RAM's SDR, and SC-LEARN's TPUs. The CLFG (Camera Link Frame Grabber) application is responsible for taking captures from the HSI sensor and saving them to SCv3M DDR memory. By default, the CLFG application will capture a 640×480×400 data cube (with 270 of 480 valid hyperspectral bins) consisting of 400 individual frames or lines of 12-bit HSI sensor data from 480 spectral bands and 640 spatial pixels in the along-track dimension. This capture produces a 184 MB raw capture file saved to a predefined location in the Linux file system. Both camera and frame-grabber configuration parameters such as exposure time, number of frames, and timeout are configurable via a cFS configuration table. Images can optionally be automatically saved to the onboard NAND flash (persistent storage), or segmented into 1-MB chunks, and downlinked via CF. CF is programmed to periodically check and downlink files from a predefined path in the Linux filesystem. Applications communicate to each other via the software bus, built on top of the `mqueue` subsystem in Linux.

The TPU application interacts with SC-LEARN and is responsible for ingesting inference commands and outputting inference results. A user may command the application with custom parameters, specifying the `tflite` (TensorFlow Lite) model, raw data, and the bounds of the data to run an inference on. Inference results are automatically downlinked as files via the CF application. The SDR application interacts with SC-RAMS allows for SDR configuration and capturing RF signals, and outputs the signals as files which are also downlinked via CF.

Finally, the Orca Myriad card allows for running inferences on capture data, with a fully configurable system that allows for custom models and bounds on raw data. The flight software for Orca was developed for cFS and provides the connectivity from the cFS app to the microcontroller and the Myriad X through the DepthAI SoM. All system control for the DepthAI SoM is commanded from the cFS Orca app to the microcontroller. This software includes interfacing with peripherals to accurately control the power sequence and configure the SPI node to the DepthAI module for inferencing operations or model pipeline updates. The cFS app will also interface with the microcontroller to periodically report the power consumption recorded for monitoring. From the cFS app, Orca will perform updates from the ground station by writing uploaded firmware and model pipeline updates to the DepthAI SoM NOR flash to execute new model pipelines on the Myriad X. The flight software uses a protocol defined in the DepthAI SPI API to perform SPI transactions with the DepthAI SoM to execute AI inferencing on the Myriad X. Custom model pre- and post-processing procedures for data input and output are supported by dynamically linking libraries that define an application's specifications at software application runtime.

While each hardware application is fully configurable, SCENIC has been setup with an automated inference pipeline for nominal operations and initiated by default during power-on. In this pipeline, the CLFG application captures a HSI data cube and forwards it to the TPU application for inference, and the inference results and the image chunks are then downlinked via CF. This pipeline runs at a periodic timed interval, allowing for automated capture of HSI data cubes and inference results. It is possible to enter a LOS (loss of service) period with the ISS during the downlinking of the HSI chunks. This time period results in incomplete data during live telemetry streaming. However, the data is buffered on the ISS during LOS, downlinked shortly afterwards to the HOSC, and is recorded as replay files. These replay files can be manually requested, allowing for full reconstruction of the HSI cube.

SCENIC PetaLinux

The SCENIC PetaLinux project produces of all necessary software components to run the customized version of cFS on the MicroBlaze processor instantiated within the SCv3M. PetaLinux 2020.2 is the 2020.2 version of the AMD-Xilinx tools for embedded Linux development and is based on Yocto Release 3.0 (codenamed Zeus) and uses the Xilinx fork of the Linux kernel version 5.4. Mission specific kernel configuration fragments, patches, drivers, libraries, and executables were added on top of the base PetaLinux configuration to integrate both SCv3M support logic with flight-card application logic. Notably, this build includes libraries such as TensorFlow Lite and Edge TPU libraries for the SC-LEARN and Linux Industrial Input/Output (`iio`) libraries for the SC-RAMS. Similar modifications were made for acceleration IP, such as from the integration of resources from Xilinx-provided reference projects to run Vitis AI 1.3 applications on the DPU IP through the Xilinx Runtime Library (XRT) and developing custom Linux Userspace I/O (UIO) logic to run the quantization preprocessing acceleration IP to supplement the SC-LEARN.

Ground Station Architecture

SCENIC used COSMOS [21] version 4.5 throughout the entire development cycle. COSMOS is a suite of open-source ground-station applications that provide services, user interfaces, and scripting APIs for command and telemetry streams that are useful during development, testing, and operations. The definitions for packets within these streams are organized into groups called targets. A target was created for each cFS app. At Goddard, COSMOS is multi-purpose and leveraged for three separate use cases.

1. To perform automated testing of newly fabricated SCv3M processor cards
2. To aid in the development and testing of customized cFS applications for mission- or instrument-specific efforts and testing sequences
3. To act as the primary ground station software during integration testing and operations

The SCENIC team has developed a suite of tests in COSMOS to allow automated testing of the SCv3M processor card. These automated tests primarily focus on internal interfaces and component-level diagnostic tests. This feature is invaluable because it enabled rapid functional testing to validate the assembly of physical hardware and detect hardware defects.

COSMOS was used throughout the cFS application development cycle, co-developing the applications with their corresponding COSMOS target command and telemetry definitions. This process allowed the team to develop and test the software in its flight configuration from the start of the project.

Finally, COSMOS is used during integration and testing as well as the actual mission operation control software. COSMOS interacts with the Telescience Resource Kit (TReK) to receive, send, and monitor packets between the ground station computer and SCENIC onboard the ISS. NASA Goddard uses a variety of software suites including ITOS[7] and ASSIST[8]; however, COSMOS was selected due to the other functionality provided across the development cycle.

The COSMOS Command and Telemetry Server serves as the bridge between the user or scripts and the command and telemetry stream. When the CF cFS app downlinks files, the COSMOS CFDP (CCSDS File Delivery Protocol) engine interfaces with the Command and Telemetry Server to reconstruct the files. The COSMOS Command and Telemetry Server is left running on a local ground station. During LOS events, the lost telemetry is played back into the COSMOS Command and Telemetry Server to reconstruct missing files.

VI. FPGA DESIGN

The SCENIC FPGA design on the SCv3M Kintex UltraScale FPGA consists of a MicroBlaze-based system and contains several FPGA IP for interfacing with onboard interfaces and subsystems, such as the rad-hard monitor, and with the other flight cards and sensors. The 32-bit MicroBlaze is configured to run PetaLinux with the "Linux with MMU" configuration and consists of all PetaLinux-required AMD-Xilinx IP, in addition to a local block RAM controller for an embedded first-stage bootloader. This MicroBlaze controls several peripheral IP consisting of standard AMD-Xilinx IP for common interfaces, such as SPI and I2C, and custom developed IP logic for certain application-specific interfaces, such as a Camera Link frame grabber block and SpaceWire, as well as custom card-specific processing IP blocks, such as ADC processing pipelines for the SC-RAMS.

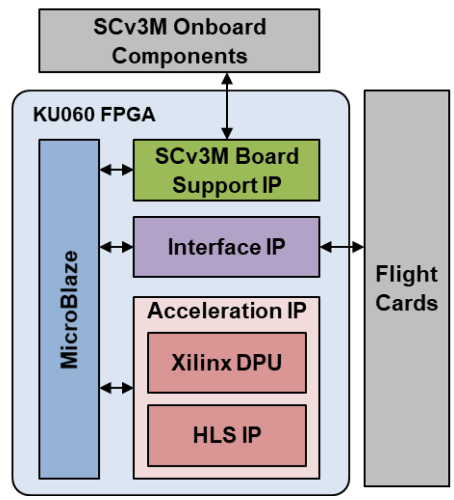

Figure 13: High-Level Diagram of the SCENIC FPGA Design

In addition to interface IP, the SCENIC design also integrates several acceleration IP, such as Deep Learning Processor Unit (DPU) IP, the AMD-Xilinx programmable engine optimized for convolutional neural networks, and a custom high-level synthesis (HLS) IP developed for the SC-LEARN application's preprocessing stages. SCENIC specifically uses the DPUCVDX8G DPU IP configured for the highest performing B4096 DPU architecture along with the options for low RAM usage, high DSP usage, and enabled optional functions (channel augmentation, depth-wise convolution, and leaky rectified linear unit (ReLU) activation function). For the SC-LEARN preprocessing app, the HLS IP was generated from Vitis HLS to accelerate both the unpacking of 12-bit image

[7] https://itos.gsfc.nasa.gov/

[8] https://sed.gsfc.nasa.gov/etd/583/tech/asist

sensor data and the normalization and quantization of that data in a dataflow HLS architecture. The SCENIC FPGA architecture is depicted in Figure 13.

VII. ENVIRONMENTAL TESTING

As an essential step for TRL advancement, SCENIC on STP-H9 has undergone environmental testing including vibration and thermal-vacuum testing. This testing occurred at the experiment level before additional testing conducted at the full payload level (including Electromagnetic interference testing) after the full integration of STP-H9. Several vibration tests were conducted at the integrated level to address varying concerns, notably changing combined load factors when another payload was unable to make delivery. For more detailed environmental testing information, reference [22].

Vibration Testing

SCENIC was required to undergo a modified vibration test from the standard NASA workmanship test. A vibration test is required for experiments and performed to identify latent defects and manufacturing flaws in electrical, electronic, and electromechanical hardware at the component level or sub-assembly level. For STP-H9, the testing spectrum was drawn from the worst-case levels from integrated spectrums from the SpaceX Dragon JEM-EF and NASA Workmanship Low Mass Acceptance.

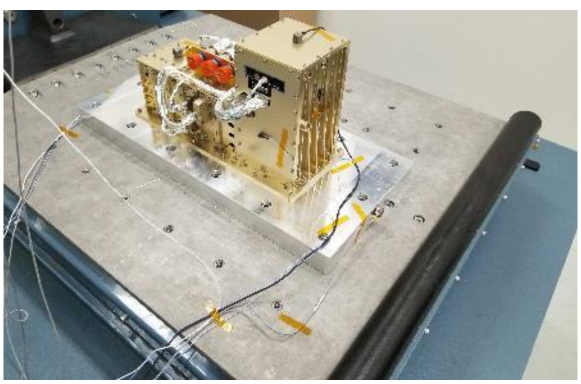

Figure 14: SCENIC at Vibration Testing

Before transport to the testing facility, full functional tests along with Electrical Interface Continuity/Isolation Test (EICIT) were conducted to identify any changes in behavior with comparative analysis before and after the vibration test. The vibration test was performed unpowered, with a sine sweep prior to and after each axis. The results of a sine sweep are necessary because they are compared before and after the Random Vibration Test to verify that there were no changes in frequencies. Any major anomalies or changes would indicate an alteration in the structure and would require further investigation. The workmanship vibration test of SCENIC was performed successfully on all three axes, with no significant changes detected during the sine sweeps or in post-functional tests. SCENIC with vibration adapter plate on vibration table is pictured in Figure 14.

Thermal Vacuum (TVAC) Testing

A TVAC test is performed to verify the continued expected performance of a payload in varying temperature ranges that could occur during the mission duration. The testing temperature range encompasses both the nominal operating conditions and expanded worst-case high and low ranges. The overall profile consisted of multiple cycles in vacuum with a hot operational plateau of 55°C and a cold operational plateau of -10°C, at the SCENIC baseplate. A full-functional performance test was performed at each plateau prior to temperature transitions. Throughout the test, the experiment conducts nominal on-orbit activities throughout the temperature transitions and holds/soaks. The test was performed using minimum and maximum input voltage at various cold/hot cycles in order to capture corner cases, as the specified input voltage could vary during operations. SCENIC performed nominally throughout the TVAC test and initial telemetry from commissioning confirms the anticipated behavior of the experiment. SCENIC in TVAC chamber is displayed in Figure 15.

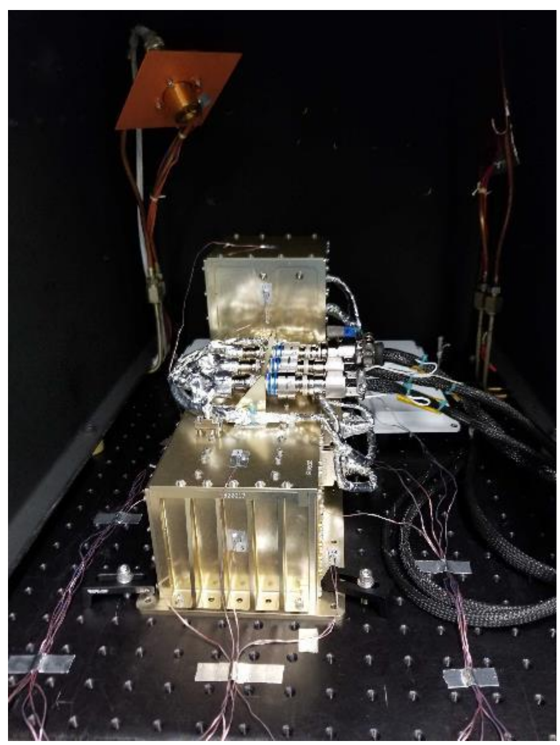

Figure 15: SCENIC in TVAC Chamber

Radiation Testing

SCENIC is a small technology demonstration mission and during development did not have corresponding funding for radiation testing. However, relevant testing for the Myriad X and TPU were conducted for other programs and are discussed here briefly for reference, since these devices are included in the SCENIC architecture as part of the Orca and SC-LEARN card assemblies.

An extensive radiation test campaign was conducted for both the Intel Myriad X and Google Coral Edge TPU devices under the NASA Electronics Parts and Packaging Program (NEPP) and the Goddard Radiation Effects and Analysis Group (REAG). This testing was conducted in support of the DAVINCI (Deep Atmosphere Venus Investigation of Noble gases, Chemistry, and Imaging) instrument Compact Ultraviolet to Visible Imaging Spectrometer (CUVIS) instrument [23]. CUVIS will employ advanced optical systems and innovative artificial intelligence/deep learning (AI/DL) hardware and software for onboard data processing for rapid identification of science results (Edge AI). Several successful radiation test campaigns were conducted [24] including total-ionizing dose at the GSFC Radiation Effects Facility (REF), heavy-ion testing at Brookhaven National Lab NASA Space Radiation Laboratory (NSRL, TPU), heavy-ion testing at Lawrence Berkeley National Laboratory (LBNL, Myriad), and high-energy proton testing at Massachusetts General Hospital (MGH).

VIII. LAUNCH, INSTALLATION, AND OPERATIONS

STP-H9 was launched on SpaceX CRS-27 from LC-39A at NASA Kennedy Space Center, Florida on March 14, 2023, at 20:30 EDT. On March 19th it was removed from the Dragon trunk, installed on the JEM-EF, and powered on. On March 20th all experiments had completed both activation and initial functional and were confirmed to be operating nominally. STP-H9 is now successfully installed and activated on the ISS.

As of May 2023, SCENIC has completed all functional checkouts and is responding as anticipated compared to TVAC data and has begun downloading HSI captures. Unfortunately, commissioning is proceeding slower than expected due to internal staffing conflicts with other projects and notable impediments with NASA Information Technology (IT) resources. The development team is preparing for the first major update of firmware and software for new science applications despite unanticipated delays.

IX. CONCLUSIONS

The Space Test Program - Houston 9 - SpaceCube Edge-Node Intelligent Collaboration (STP-H9-SCENIC) experiment was recently deployed on the ISS to demonstrate artificial-intelligence and machine-learning technologies and capabilities on multiple space-based platforms. Additionally, SCENIC will monitor each device’s availability and susceptibility to radiation effects. The acquisition of Earth observation imagery captured by the onboard CubeSat-sized hyperspectral imager can enable this investigation to provide the scientific community with hyperspectral datasets to facilitate the development of neural network models for future advanced missions with breakthrough capabilities such as responsive, onboard data analysis and autonomous swarm/constellation management. Finally, SCENIC advances the technology readiness level of next-generation NASA technologies for spaceflight including the CubeSat-sized, radiation-tolerant, high-performance computer, SpaceCube v3.0 Mini, and new CubeSat form-factor GPS and SDR cards for guidance and navigation. Currently, SCENIC is in commissioning and will soon advance to deploying newly uploaded AI/ML applications for demonstration.

Acknowledgments

SCENIC is successful thanks to the collaboration and assistance from many friends and colleagues. The authors would like to recognize the contributions and support by additional team members of the Science Data Processing Branch Code 587 and *many* Goddard collaborators including James MacKinnon, Michael Lin, Dong Wu, Alan Gibson, Sergio Maldonado, Hunter Somerville, Dwaine Molock, Sabrena Ball, Benjamin Nold, Darren Midkiff, Zach Boblitt, David Hatke, Luke Winterniz, Bill Bamford, Cindy Goode, Ted Daisy, Caitlin Burth, Christian Billie, Joshua Yacobucci, Gabe Garde, Allen Teasley, Michael Johnson, Christopher Green, Frank Robinson, Colton Cohill, Michael Chapman, Allison Evans, Matt Fritch, Ted Drilling, and Charlie Rogers. The authors also recognize Aerospace collaborators Mikhail Tadjikov and Benjamin Davidson. For Myriad X integration, the authors thank the team of specialists in embedded ML and computer vision at Luxonis including Brandon Gilles, Jon Ngai, Martin Peterlin, et al. Luxonis has been supportive in developing new features within their core embedded SPI framework on the DepthAI SoM to support SCENIC. For accommodation for the SCENIC experiment, we acknowledge the great team at STP-Houston and the Huntsville Operations Support Center. The authors thank supporting organizations including, AMD-Xilinx (specifically Timothy Vales), Apogee Semiconductor (Anton Quiroz), Headwall Photonics (Mike Donahue, David Blair, William Rock, Blair Simon, Ross

Nakatsuji), Carr Astronautics (Jim Carr), and finally Texas Instruments (Joey Christ and Carlos Ventura).

Finally, special thanks to our key sponsors supporting this development, the NASA Earth Science Technology Program (ESTO) and the NASA/GSFC Internal Research and Development (IRAD) program, specifically the Cross-Cutting and Small Satellites Line-of-Businesses.

References

1. National Academies of Sciences, Engineering, and Medicine, "Thriving on Our Changing Planet a Decadal Strategy for Earth Observation from Space," Washington, DC: The National Academies Press. 2018. https://doi.org/10.17226/24938

2. "National Academies of Sciences, Engineering, and Medicine, "Origins, Worlds, and Life: A Decadal Strategy for Planetary Science and Astrobiology 2023-2032," Washington, DC: The National Academies Press. 2022.https://doi.org/10.17226/26522"

3. "2020 NASA Technology Taxonomy," Washington, D.C., USA: NASA Office of the Chief Technologist, 2020.

4. Air Force Space Command, "AFSPC Long-Term Science and Technology Challenges," Colorado Springs, CO: Peterson AFB, 2016.

5. Cardillo, R., "Small Satellite 2017 Keynote Address," 31st Annual AIAA/USU Conference on Small Satellites, Logan, UT, Aug 7, 2017. https://www.nga.mil/MediaRoom/SpeechesRemarks/Pages/Small-Satellites---Big-Data.aspx

6. Swope, J., Mirza, F., Dunkel, E., Towfic, Z., Chien, S., Russell, D., Sauvageau, J., Sheldon, D., Fernandez, M., and C. Knox, "Benchmarking Remote Sensing Image Processing and Analysis on the Snapdragon Processor Onboard the International Space Station," IEEE International Geoscience and Remote Sensing Symposium (IGARSS), Kuala Lumpur, Malaysia, Jul. 17-22, 2022.

7. George, A. D.., and C. Wilson, "Onboard Processing with Hybrid and Reconfigurable Computing on Small Satellites," Proceedings of the IEEE, vol. 106, no. 3, pp. 458-470, Mar. 2018.

8. Giuffrida G., Fanucci, F., Meoni, G., Bati, M., Buckley, L., Dunne, A., van Dijk, C., Esposito, M., and J. Hefele, "The Φ-Sat-1 Mission: The First On-Board Deep Neural Network Demonstrator for Satellite Earth Observation," IEEE Transactions on Geoscience and Remote Sensing, vol. 60, pp. 1-14, 2022. doi: 10.1109/TGRS.2021.3125567

9. Wright, R., "HyTI: high spectral and spatial resolution thermal infrared imaging from a 6U CubeSat," 2021 Earth Science Technology Forum, Virtual, May 6, 2021.

10. "SpaceCloud launched to orbit on June 30,2021 from Cape Canaveral to validate future space data handling," Unibap, [Online]. Available: https://unibap.com/en/our-offer/space/projects/wild-ride-mission/

11. Dunkel, E., Swope, J., Towfic, Z., Chien, S., Russell, D., Sauvageau, J., Sheldon, D., Romero-Canas, J., Espinosa-Aranda, J. L., Buckley, L., Hervas-Martin, E., Fernandez, M., and Knox, C. "Benchmarking Deep Learning Inference of Remote Sensing Imagery on the Qualcomm Snapdragon and Intel Movidius Myriad X Processors Onboard the International Space Station," International Geoscience and Remote Sensing Symposium (IGARSS 2022), Kuala Lumpur, Malaysia, July 2022.

12. Brewer, C., Franconi, N., Ripley, R., Geist, A., Wise, T., Sabogal, S., Crum, G., Heyward, S., and C. Wilson, "NASA SpaceCube Intelligent Multi-Purpose System for Enabling Remote Sensing, Communication, and Navigation in Mission Architectures," 34th Annual AIAA/USU Conference on Small Satellites, SSC20-VI-07, Logan, UT, Aug. 1-6, 2020.

13. Sims, E., "The Department of Defense Space Test Program: Come Fly with Us," Proceedings of IEEE Aerospace Conference, March 2009.

14. McLeroy, J., "Highlights of DoD Research on the ISS," ISS Research and Development Conf., June 26-27, 2012.

15. Goodwill, J., Crum, G., MacKinnon, J. Brewer, C., Monaghan, M., Wise, T, and C. Wilson "NASA SpaceCube Edge TPU SmallSat Card for Autonomous Operations and Onboard Science-Data Analysis," 35th Annual AIAA/USU Conference on Small Satellites, SSC21-VII-08, Virtual Conf., Aug. 7-12, 2021.

16. "Intel Movidius Myriad X VPU Product Brief," Intel, 2018. [Online]. Available: https://www.intel.com/content/www/us/en/products/docs/processors/movidius-vpu/myriad-x-product-brief.html

17. Franconi, N., Sabogal, S., George, A. D., Hassouneh, M., Mitchell, J., and C. Wilson, "A Novel RF Architecture for Simultaneous Communication, Navigation, and Remote Sensing

with Software-Defined Radio," 34th Annual. AIAA/USU Conference on Small Satellites, SSC20-WKVIII-02, Logan, UT, Aug. 1-6, 2020.

18. Gambino, J., and N. Dennehy, "A Multi-Function Guidance, Navigation and Control System for Future Earth and Space Science Missions, 25th Annual AAS Guidance and Control Conference, Breckenridge, CO, Feb. 6-10, 2002.

19. Winternitz, L., Bamford, W., Price, S., Carpenter J., Long, A., and M. Farahmand, "Global Positioning System Navigation Above 76,000 km for NASA's Magnetospheric Multiscale Mission," NAVIGATION: Journal of the Institute of Navigation, vol. 64, no. 2, Jul. 2017.

20. Petrick, D., Geist, A., Albaijes, D., Davis, M., Sparacino, P., Crum, G., Ripley, R., Boblitt, J., and T. Flatley, "SpaceCube v2.0 space flight hybrid reconfigurable data processing system," 014 IEEE Aerospace Conference, Big Sky, MT, USA, 2014, pp. 1-20, doi: 10.1109/AERO.2014.6836226.

21. Melton, R., "Ball Aerospace COSMOS Open Source Command and Control System," 30th Annual AIAA/USU Conf. on Small Satellites, SSC16-IX-03, Logan, UT, August 6-11, 2016.

22. "General Environmental Verification Standard (GEVS) for GSFC Flight Programs and Projects," GSFC-STD-7000, Greenbelt, MD, USA: Goddard Space Flight Center, 2021.

23. Garvin, J., and et al, "Revealing the Mysteries of Venus: The DAVINCI Mission," Planetary Science Journal, vol. 3, no. 5, pp. 1-17, May 2022. doi: 10.3847/PSJ/ac63c2

24. Casey, M., Wyrwas, E., and R. Austin, "Recent Radiation Test Results on COTS AI Edge Processing ASICs," NASA Electronic Parts and Packaging (NEPP) Program 2022 Electronics Technology Workshop, Greenbelt, MD, Jun. 15, 2022.