Commissioning and first flight of a compact intelligent space camera for in-orbit visual operations

Aubrey Dunne, Juan Romero-Cañas, Peter O'Connor, Rosana Rodríguez-Bobada Aranda,
Tom Hendrix, David Rijlaarsdam
Ubotica Technologies Ltd.
Dublin, Ireland
aubrey.dunne@ubotica.com

***Abstract*—The application of advanced AI techniques to support vision-based autonomy has now become commonplace on ground, but this has yet to translate in earnest to operational uses on satellite. This paper presents early results from the first deployment and commissioning of a custom-designed AI-enabled intelligent space camera that brings real-time edge AI capability to visual monitoring applications in-orbit. Camera firmware supports a flexible and reconfigurable processing flow that leverages multiple hardware acceleration blocks in the edge processor to perform vision processing within a low power envelope and a compact 0.4U form factor. Component reliability and radiation robustness have been factored into the design to enable camera operation in both LEO and GEO, making the camera applicable to a range of in-orbit applications, including SSA, RPO, debris detection, visual FDIR and threat detection. The camera has recently flown as a hosted payload on the Apex Aries SN1 productised ESPA-class satellite bus, launched in March 2024. Since launch, the various functions of the camera have been commissioned and tested, and operations are beginning to demonstrate AI-enabled application flows in-orbit. The paper discusses the camera configuration, deployment, and commissioning phases of the mission, and presents commissioning results. Early data on operational performance for experiments conducted to date, including image acquisition of the Earth and the starfield, and on-board processing are presented. Planned future mission operations that aim to demonstrate additional targeted use cases are discussed. Commissioning results prove the feasibility and utility of deploying AI-enabled smart cameras in space.**

***Index Terms*—Intelligent space camera, streaming AI, space domain awareness, resident space objects, rendezvous and proximity operations**

I. INTRODUCTION

Satellite autonomy has the promise of reducing operational latency and costs, whilst simultaneously enabling reactivity and improved duty cycles, by converting satellites into intelligent robotic agents. Multiple approaches to achieving this are being investigated, and this paper explores the aspect of vision-based autonomous operations. Vision is a rich medium that has been shown terrestrially to hugely benefit from the application of AI, in particular of Convolutional Neural Networks (CNNs) and Vision Transformers (ViTs) [1].

Multiple application domains exist for Non-Earth Imaging (NEI) satellites, these being distinct from Earth Observing (EO) use cases by virtue of the imaging targets - for EO, the target is exclusively the Earth, whereas NEI is exclusive of the Earth. Furthermore, NEI is generally considered to exclude astronomical observations. Among NEI's application domains are launch vehicle operations monitoring; Rendezvous and Proximity Operations (RPO); Resident Space Object (RSO) detection and monitoring [2]; Space Domain Awareness (SDA); and visual fault detection. Vision brings deep contextual information in the form of (typically) three-dimensional data to all these applications, and at standard framerates can capture highly unstructured dynamics that are rich in information datapoints.

The deployment of AI in terrestrial vision applications is ubiquitous, ranging from auto-labelling of live broadcasts, to face and smile detection in phone camera applications, to Optical Character Recognition for number plate recognition and document scanning. By appropriate training of the AI models, the interpretation burden can be substantially pushed to the offline pre-deployment (training) stage where latency is not generally an issue and where power, storage and compute resources are in abundance. This burden shift in turn allows the runtime phase of AI deployment to be very efficient even on visual data, enabling realtime interpretation on low-power edge hardware. Since the first deployment of hardware-accelerated AI inference on a cubesat on ϕ -Sat-1 in September 2020 [3], the adoption of deep-learning-based AI on-board satellites has accelerated in pace [4]–[6]. Nevertheless the power and associated thermal constraints of many cubesats and smallsats restricts how such AI applications are deployed, and adoption requires careful planning. Furthermore, effective deployment of AI models on-board satellites requires many additional steps to those required for less resource-constrained terrestrial deployments, including hardware-in-the-loop validation, power analysis, RAM usage analysis, and timing analysis.

Classical deployment of vision sensors on spacecraft involves using radiation hardened or radiation-mitigated Commercial-Off-The-Shelf (COTS) cameras that transform photons to pixels. They consist of an optical frontend (filters, lenses), a sensor (CCD or CMOS), control circuitry, memory, and interfaces for control and image transmission. At a spacecraft subsystem level, the camera is interfaced to a payload data handler, On-Board Computer (OBC), or other processing node in the system that converts the image data to information relevant to the target application on the satellite (e.g., docking or debris detection). In contrast, the camera that is the target of this paper operates on the 'photons in, insights out' principle, achieving this using AI stream processing integrated directly into the camera. By adding the application-specific intelligence into the camera, processing can be truly pushed to the edge and decentralised, improving modularity, reactivity, resilience, and power efficiency.

This work explores the practical realisation of an element of this visual autonomy – that of an AI-enabled intelligent space camera for satellite in-orbit operations [7]. It aims to demonstrate the feasibility and utility of deploying a COTS-based camera with integrated AI capabilities directly on satellite, and presents commissioning and early results for such a deployment. An example pose estimation application executing directly on an Engineering Model of the intelligent space camera is shown in Fig. 1.

Fig. 1. Spacecraft pose estimation executing directly on the Intelligent Space Camera. Camera captures frames of scene on right-hand monitor, and streams position and pose estimates directly to the laptop for display on the left-hand monitor.

The remainder of this paper presents details on the Intelligent Space Camera (Section II), the first space mission into which the camera was integrated (Section III), and camera commissioning results from the mission (Section IV).

II. Intelligent Visual Sensing

A. Camera Overview

The Intelligent Space Camera (ISC) has been designed from the ground up for visual AI workloads. It supports dynamic AI model updates (enabling model hot-swapping for application diversification), and incorporates a reconfigurable fully-hardware Image Signal Processing (ISP) pipeline, hardware AI acceleration, and a hard block for H.265 video encoding. Firmware and software have been developed for efficient image capture, processing, and delivery over Ethernet, as well as boot, control and configuration from the satellite's OBC. A RESTful API enables dynamic configuration of the camera from the host OBC, supporting changes to the AI network, the ISP, the encoder, and the sensor, whilst additionally enabling the querying of telemetry, health and status information. This first flight camera has been configured with a full HD (1920x1200 pixel) area-scan RGB CMOS global-shutter sensor with a 99° FOV (Field of View) optical stack for wide field operations.

B. Hardware and Firmware

The ISC, shown in Fig. 2, has a compact and self-contained custom form factor designed to withstand the harsh environment of space. It comprises an optical stack within a ruggedised mechanical enclosure that houses a processing and control unit with a size envelope of 128x68x47mm (LxWxH) and with a mass of 425g. Designed with automotive-grade components, its features include on-board latch-up protection and non-volatile storage. Data and control are transported over a Gigabit Ethernet interface, and the unit is powered via a 28V bus connection.

All on-camera processing is performed by an Intel Movidius Myriad X Vision Processing Unit [8] - a software-controlled hardware-software System on Chip (SoC) that integrates hardware-accelerated ISP, AI inference, and video compression hard blocks. The Myriad X has been characterised for the ionising radiation environment of space [9], and has spaceflight heritage having flown on several LEO satellites [10], [11].

Camera firmware supports sensor auto-exposure and auto-whitebalance control and a fully hardware configurable ISP pipeline that processes the Full HD frames at up to 120 Frames Per Second (FPS). Dedicated hardware accelerators for Neural Network (NN) computations on Myriad X provide 1 TOPS of AI compute, and these are supported by 16 integrated SHAVE vector engines to enable custom layer processing to support a broad range of NN architectures that are targeted at the device via the OpenVINO toolchain [12]. A lightweight HTTP server in the firmware facilitates a RESTful Application Programming Interface (API).

C. Software

The ISC firmware utilizes a software architecture designed to enable flexible and autonomous AI-driven operations, and which forms the server side of a client-server paradigm. A Manager application (client) resides on the main OBC and acts as the controller for ISC functionalities. The Manager's operations can be customized through JavaScript Object Notation (JSON) configuration files communicated over HTTP, allowing for the specification of operational parameters, such as the uplink and swapping of AI models on the ISC at runtime, without requiring recompilation or software updates. Upon loading the configuration parameters from a designated JSON file, the Manager subsequently initiates execution of lightweight camera controller binaries dedicated to various

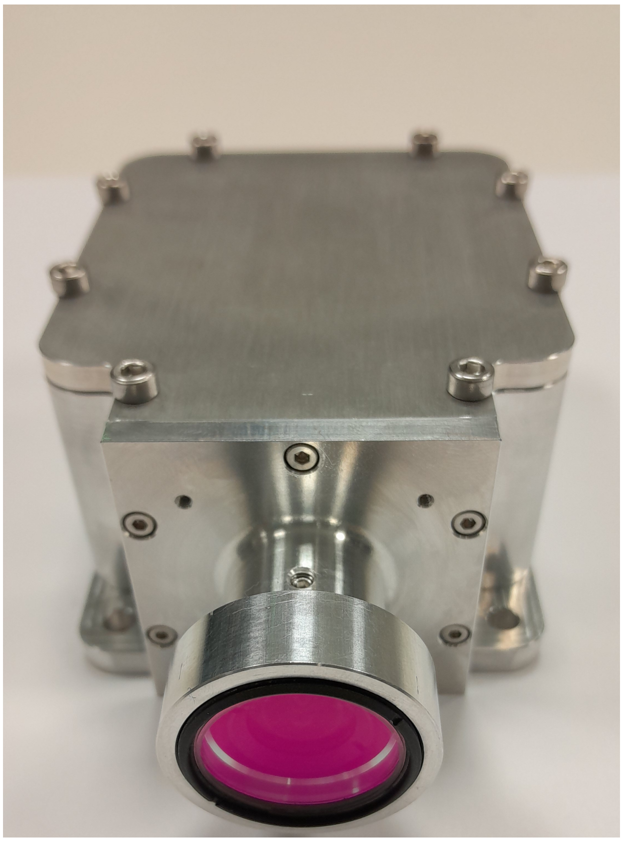

Fig. 2. Intelligent Space Camera (CogniSAT-NEI)

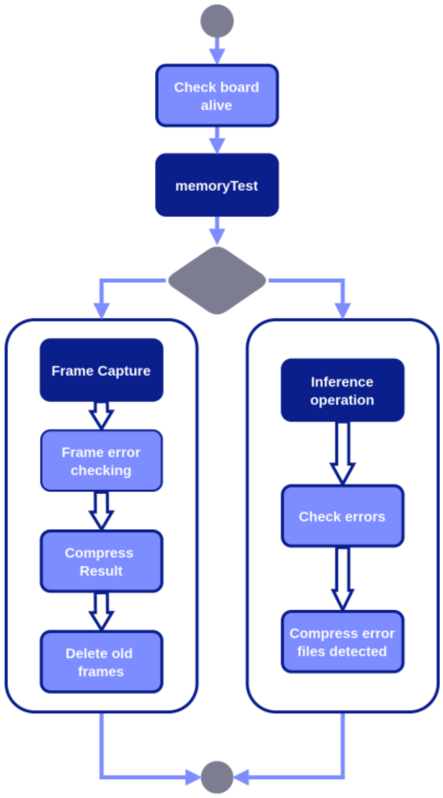

Fig. 3. Manager functional flow

operational procedures. This modular system architecture facilitates enhanced reconfigurability and maintainability.

The Manager orchestrates the execution and validates the successful completion of all operational sequences, while also managing writing to the system's logging subsystem. As depicted in Fig. 3, the high-level functional flow of the Manager comprises two distinct operational pathways: a frame capture path and an inference path. These separate pathways reflect the initial operational modes of the system on-board the spacecraft. Note that for the commissioning phase the frame capture and inference operation paths are mutually exclusive. For camera operations on this first flight, the central management module invokes optimized and compiled applications. These applications are identified as SimpleCam, CogniSatApp, and MemoryTest:

SimpleCam is an application designed to configure the camera for frame capture. Once the custom firmware has booted and is running on the ISC board, frames can be requested from the host side using SimpleCam. This application enables configuration of the ISC board to either process image frames through the ISP, resulting in debayered imagery, or alternatively, to acquire frames in the raw Bayer format. Furthermore, SimpleCam facilitates continuous frame capture, enabling the acquisition of video sequences.

CogniSatApp is an application designed to facilitate the execution of AI inference and custom image processing pipelines on the ISC. This application leverages JSON configuration files to enable the streamlined deployment of novel functionalities with minimal coding requirements, adhering to a low-code or zero-code paradigm. The architecture of CogniSatApp incorporates a generic framework for pre- and post-processing stages, permitting the deployment of a variable number (including zero) of pre-processing operations to enhance flexibility and adaptability.

MemoryTest is a diagnostic application that can be optionally executed prior to each inference or frame capture operation to assess and verify the full functionality of the on-camera Synchronous Dynamic Random-Access Memory (SDRAM).

All three applications generate logs during execution that are available for downlink and on-ground verification. The Manager also supports integrated compression, configurable via JSON, of captured frames. Together with the Manager, the compiled applications enable a wide gamut of applications to be executed on the camera simply by uplinking application-specific JSON files to the spacecraft and providing them as parameters to the Manager invocation.

III. Mission and Integration

Once the ISC design was completed and an initial prototype was produced and successfully functionally tested, flight opportunities were explored for the deployment and in-orbit camera testing and verification. The mission goal was to validate the ISC operation in an orbital LEO environment over a multi-year duration and under operational conditions on a state-of-the-art smallsat bus.

A. *Mission*

The Apex "Call to Adventure" mission was identified as an ideal carrier for the ISC, and integration with the Apex bus began in mid 2023. The mission consists of an Apex Aries satellite (SN1) - a productised ESPA-class satellite bus with a 150kg payload capacity, shown in Fig. 4 - with the camera mounted on the bus exterior such that it has an unobstructed field of view. Aries SN1 is a first-of-its-kind rapid design-to-production spacecraft from Apex (Los Angeles, CA). The satellite was launched in March 2024 on-board SpaceX Transporter 10 from Vandenberg SFB in California.

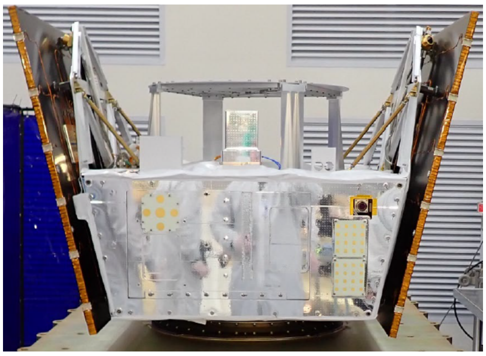

Fig. 4. Apex Aries SN1 satellite (©Apex, used with permission)

B. *Integration*

The ISC is interfaced with the main satellite OBC via an Ethernet connection, and separately to a 28V supply output on the Electrical Power System (EPS) unit. The integration process was conducted by Apex with the support of Ubotica. A set of verification scripts were provided to enable confirmation of the functionality upon integration of the payload board into the spacecraft bus. All software components described in Section II-C were supplied to Apex to facilitate their integration into their build system for the OS on the OBC. The compiled applications underwent verification processes on Apex's dedicated Hardware-In-The-Loop (HITL) environment, where the HITL provides test capability on a (hardware and software) digital twin of the spacecraft.

C. *Commissioning Framework*

Commissioning tests were designed pre-flight to exercise the ISC functionalities as follows:

- **Single Frame Capture** In this test the ISC captures a frame in raw format at 1920 by 1200 resolution and at 12-bit depth, with each frame having a size of approximately 4.6 MB. The raw format corresponds to the data directly from the sensor, without the ISP applied.
- **Long Frame Capture** In this test the ISC continuously captures frames up to a configurable number of total frames. This results in being able to capture multiple frames as an in-orbit video sequence.
- **Orbit Frame Capture** This test verifies intermittent frame capture during a complete orbit, evaluating camera performance during a full orbit temperature cycle, and demonstrating orbital mapping capability.
- **AI Inference** In this test the ISC executes inference using a pre-uploaded model, on pre-uploaded images, with output compared against pre-uploaded verification data. The input data on-board is pre-tiled such that the tile size matches the input tensor size of the ResNet object detection test model. Inference output files are flagged for downlinking only when there is a mismatch between them and the verification data.
- **Long Exposure Capture** This test exercises the ISC exposure time reconfiguration capability. The test was designed to address the ISC use case where the camera is pointed away from Earth to capture the star field, requiring that the exposure setting is changed to accommodate for low illumination.

D. *Compression and Patching*

During commissioning it was required to update the firmware running on the ISC payload. A firmware patching approach was developed that produced an 86KB patch to the original firmware (4.6 MB). The aforementioned process was performed with a binary patching utility, similar to the functionality of `bspatch`. Furthermore, in order to download the captured frames from the satellite in an efficient and bandwidth-saving manner, it was necessary to create previews of the frames so that it could be determined on-ground which were the most appropriate full-resolution images to subsequently downlink. This process involved the Manager extracting the luma channel (Y) from the YUV ISP-processed camera output image, enabling a reduction in size of preview images to approximately 200 KB. Both preview images and full-resolution frames can be compressed by the Manager on the OBC using the 7z library. Through the implementation of these two techniques, it was possible to reduce the downlink bandwidth consumed during commissioning.

IV. Commissioning and Results

A. *Commissioning Process*

The ISC payload commissioning phase commenced in June 2024. The JSON configuration files and FlightJAS command sequences (used to control elements of the satellite operations as they relate to the commissioning tests) were uplinked prior to each commissioniong test, and images, telemetry and log files were made available by the operator after each test. All commissioning operations were validated on-ground on the Apex HITL testbed before file upload and execution on-board the satellite. The FlightJAS command files for the camera controlled the memory testing, timing of frame captures and execution of inference, while the camera pointing mode and powerup time (and therefore ROI) were controlled by the satellite operator based on requests provided by Ubotica. The commissioning tests were performed by the satellite operator with support from Ubotica.

B. Commissioning Results

A total of 681 frames (671 raw, and 10 with on-camera ISP applied) were captured and 20 inference operations were executed during the ISC in-orbit commissioning. Each on-ground test operation was performed on-board successfully. For test validation, the test logs were downloaded and checked, and compressed versions of the images were downloaded and visually inspected, enabling to choose the best uncompressed raw images for download and verification. A commissioning test was performed to validate each ISC function, with test logs, image files, and telemetry data verified against the on-ground pre-flight results from the HITL platform.

Single Frame Capture

A single image was captured in the first commissioning test. This was initially downloaded in compressed JPEG format and checked for validity before the full-size raw file was downloaded and visually confirmed. This very first captured orbital image from the ISC is shown in Fig. 5, in which a cloudy Pacific Ocean and the California coastline are visible (with the Gulf of California on the right of the image). Note that this image, although downloaded in raw format, was subsequently processed through an on-ground ISP. The log file indicated no errors.

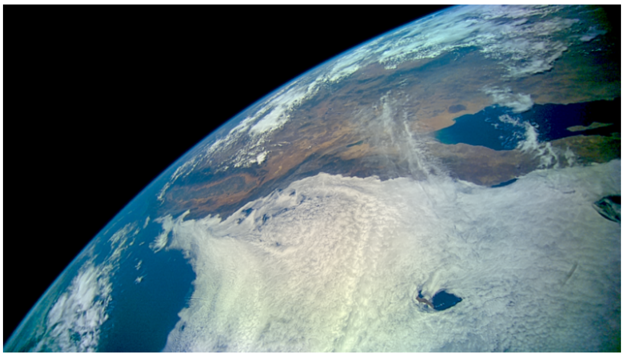

Fig. 5. First image captured by the ISC in orbit

Long Frame Capture

500 frames were captured at a rate of 9FPS in one continuous capture run. A sample of compressed JPEG frames from this sequence was downloaded, grouped at the start and the end of the set. The first frame captured was not correct (blank image). Based on timing information this was deemed likely to be due to the ISP initialisation not being completed for the first frame. The second and subsequent frames were visually checked and confirmed correct. The timing of the first frame capture after initialisation can be controlled during normal operations in order to avoid this issue in the future. The log file indicated no errors.

Orbit Frame Capture

A total of 17 capture periods were executed, with 10 frames captured per period, resulting in 170 frames at 10FPS. Capture scheduling was managed by a FlightJAS file, which managed a single capture period every 6 minutes during the full 96 minute run, as shown in Fig. 6. Log files were used to confirm that the test completed and that all frames were captured without any error being logged. Selective JPEG files were downloaded and confirmed valid, with one image (after on-ground ISP processing) shown in Fig. 7.

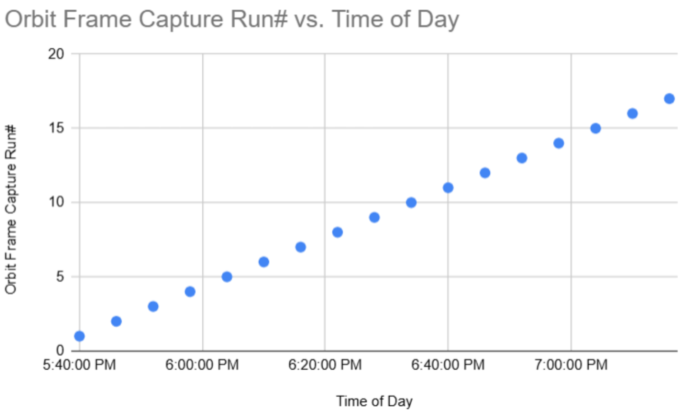

Fig. 6. Orbit Frame Capture Test Timeline

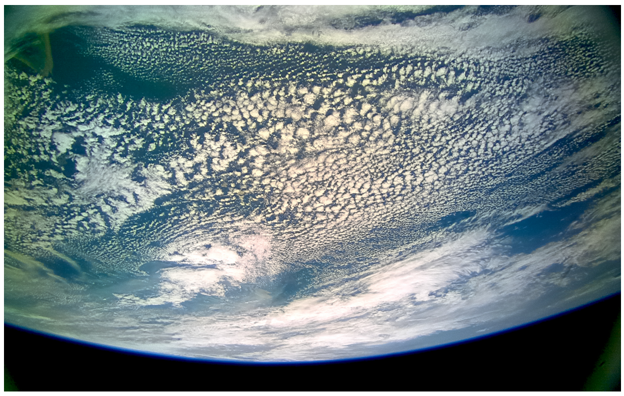

Fig. 7. Image from Orbit Frame Capture test

AI Inference

A total of 20 hardware-accelerated inferneces - 10 synchronous and 10 asynchronous - were executed in a single run on the ISC, utilising a custom UNet segmentation network for vessel detection. Segmentation results (masks) were automatically cross-checked on the OBC against reference gold outputs, with any mismatches written to the log file. No mismatches were reported. The SDRAM MemoryTest was executed before the inference operations, and reported zero bit errors.

TABLE I
COMMISSIONING TEST RESULTS SUMMARY

Test	Runs	Type	Mem Test	Result
Single Frame	1	Frames (raw)	No	Pass
Long Frame	500	Frames (raw)	No	Pass
Orbit Frame	170	Frames (raw)	No	Pass
AI Inference	20	Inference	Yes	Pass
Long Exposure Capture	10	Frames (raw)	Yes	Pass

Long Exposure Capture
A total of 10 long exposure frames were captured with the ISP enabled. A fixed exposure time of 0.5s was set via the JSON configuration prior to frame capture, resulting in a 2FPS framerate. The MemoryTest was executed prior to the frame captures, with zero errors reported. ISC power bus telemetry was captured during this test and is shown in Fig. 8. Telemetry indicates the camera power was 3.64W ($28V \times 0.13A$) during standby and peaked at 4.48W ($28V \times 0.16A$) during frame capture, these values being consistent with on-ground test values.

A summary of the commissioning test results is presented in Table I.

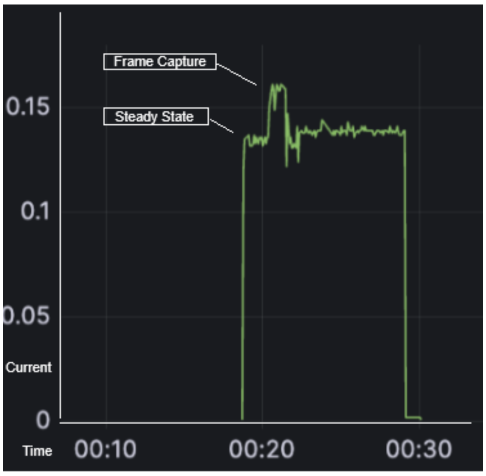

Fig. 8. ISC +28V power bus telemetry taken during Long Exposure Capture test

C. Discussion and Operations

Commissioning results show that the ISC performed to its design specification during the commissioning phase - no malfunctions were detected, and all results are consistent with those seen on HITL during testing (with the exception of a single malformed frame that is thought to be due to frame capture occurring just before the imaging pipeline was fully configured). Although capable of up to 120FPS, commissioning test framerates generally did not surpass 10FPS, which is attributed to write speed limitations on the data storage path of the OBC. Thermal information on the ISC was collected during the AI Inference test only, and due to its rapid execution time ($< 2s$) is of limited value for trend purposes. More extensive temperature logging will be performed in the operations phase.

ISC operations on Aries SN1 will commence in May 2025, where the following operations are planned: horizon-pointing frame and video capture; nadir-pointing frame and video sequence capture; starfield capture. The former two acquisition operations will be executed for both day and night scenes in order to analyse the nightime wide-swathe imaging capability of the camera. Starfield acquisitions will be captured and subsequently processed on-ground by an astrometric processing pipeline in order to attempt to match extracted image features with the known star catalogue.

V. CONCLUSIONS

An AI-enabled designed-for-space camera can have significant utility across a range of space applications, from rendezvous support to in-orbit servicing to debris detection. Realising this utility requires the development and in-orbit testing of such a camera. An Intelligent Space Camera has recently been developed that enables deployment of AI models directly to the camera for in-line streaming AI, and was flown on the Apex Aries bus in March 2024, with commissioning conducted in the following months. Within this phase the camera's primary image capture, processing and inference functionalities were exercised through a series of tests that utilised the on-camera firmware and OBC client software. Commissioning results indicate that the camera is fully functional, with over 600 images captured during commissioning, enabling mission operations with the camera to commence.

REFERENCES

[1] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," 2021.
[2] P. Day, J. Crane, P. Cronk, C. Jimenez, B. Johnson, R. Koeberle, D. Millard, and Z. Werner, "Experimental Results from On-Orbit Edge-Deployed AI Detection of Resident Space Objects Using Computer Vision," in *Advanced Maui Optical and Space Surveillance (AMOS) Technologies Conference*, Sep. 2024.
[3] G. Giuffrida, L. Fanucci, G. Meoni, M. Batič, L. Buckley, A. Dunne, C. van Dijk, M. Esposito, J. Hefele, N. Vercruyssen, G. Furano, M. Pastena, and J. Aschbacher, "The -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, Nov. 2021.
[4] M. Mercè Cañete Muñoz, F. Deconinck, J. B. Ars, N. Melega, C. Pirat, B. J. Soto, A. Dunne, L. Buckley, J. R. Cañas, G. Benelli *et al.*, "Mantis, a 12u smallsat mission taking advantage of super-resolution and artificial intelligence for high-resolution imagery," in *Proceedings of the AIAA/USU Conference on Small Satellites (Small Sat)*, Aug. 2024.
[5] N. Longépé, N. Melega, V. Marchese, A. Paskeviciute, I. Babkina, I. Petrelli, M. Casaburi, D. Peressutti, and N. O. Kadunc, "ϕsat-2 mission overview for the #orbitalai challenge, version 1.0," European Space Agency, Tech. Rep., 2023.
[6] S. Kacker, A. Meredith, K. Cahoy, and G. Labreche, "Machine learning image processing algorithms onboard ops-sat," *Proceedings of the AIAA/USU Conference on Small Satellites, Weekend Session 3: Science/Mission Payloads - Research Academia I, SSC22-WKIV-03*, 2022.

[7] A. Dunne, J. Romero-Cañas, S. Caulfield, S. Romih, and J. L. Espinosa-Aranda, "Intelligent space camera for on-orbit ai-driven visual monitoring applications," in *European Data Handling and Data Processing for Space Conference, Juan-Les-pins, France*, Oct. 2023.
[8] Intel movidius myriad x. [Online]. Available: https://www.intel.com/content/www/us/en/products/sku/125926/intel-movidius-myriad-x-vision-processing-unit-4gb/specifications.html
[9] L. A. Tambara, Ádria Barros de Oliveira, J. Andersson, L. Buckley, and A. Dunne, "Radiation characterization of the cots myriad x edge vision processing unit and use case in space applications," in *RADECS*, Sep. 2024.
[10] D. Rijlaarsdam, T. Hendrix, P. T. T. González, A. Velasco-Mata, L. Buckley, J. P. Miquel, O. A. Casaled, and A. Dunne, "The next era for earth observation spacecraft: An overview of cognisat-6," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 18, pp. 2450–2463, Jan. 2025.
[11] O. Flordal, A. Synodinos, M. Herlitz, H. Magnusson, D. Steenari, K. Förster, M. Castorina, T. George, I. Troxel, S. Reid, C. Brunskill, and F. Bruhn, "Spacecloud cloud computing and in-orbit demonstration," Jun. 2021.
[12] Openvino. [Online]. Available: https://www.intel.com/content/www/us/en/developer/tools/openvino-toolkit/overview.html