

Edge Node: A Multi-User Rendezvous and Proximity Operations On-orbit Testbed

Darren Rowen, Joel Rivera, Susan Lui, Joseph Gangestad, Alex Gallenberger, Binh Dinh,
Phaedrus Leeds, Noah Weiner, Rudd Johnson, Christopher Day, Alexander Christidis
The Aerospace Corporation
2310 E El Segundo Blvd, El Segundo, CA 90245; 310-336-3633
Darren.W.Rowen@aero.org

ABSTRACT

Edge Node is a multi-small satellite free-flying, collaborative testbed for formation flight and rendezvous and proximity operations (RPO) under development by The Aerospace Corporation that is planned for launch in 2026-2027. Edge Node will develop and advance local space situational awareness sensors for CubeSat-scale platforms and facilitate the development and qualification of autonomous on-board RPO software for CubeSats. Edge Node is intended to serve as a multi-user on-orbit facility through which technologies and algorithms can be tested and validated in the authentic operational low Earth orbital environment. Edge Node leverages recent advances in miniature sensor and computing hardware for terrestrial applications (e.g., NVIDIA Jetson TX2 NX, automotive lidar and radar sensors, microbolometer LWIR camera, and polarization image sensors, among others). The mission will enable close-in observations including pose estimation, feature identification, and dynamic motion model characterization. Edge Node is specifically designed to test the limits of autonomous decision-making to support future missions where control with a ground operator in the loop is not possible. Edge Node builds on the prior success of RPO test facilities (e.g., SPHERES and AstroBee) that were confined to operate within the International Space Station (ISS). By operating independently in low-Earth orbit, experimenters will face realistic lighting, complex background scenes, and orbital dynamics that cannot be accurately replicated via terrestrial testbeds or within the confines of the ISS.

The Edge Node mission's computing platform, consisting of a cluster of NVIDIA Jetson TX2 NX modules, builds on a foundation of three prior qualification missions in low Earth orbit in addition to proton 50 MeV radiation testing of the target hardware. A radiation-tolerant 32-bit ARM Cortex-M7 processor oversees the TX2 NX cluster and interfaces with the rest of the spacecraft avionics. Dual terabyte NVMe drives running a ZFS filesystem provide enhanced reliability on top of radiation screened industrial NAND Flash storage drives.

Edge Node will utilize a custom Linux environment utilizing Docker to containerize workloads and provide access to hardware including 256 NVIDIA CUDA GPU cores per TX2 NX for acceleration of workloads. The Robot Operating System (ROS2) will be utilized as an infrastructure layer for both intra-satellite and inter-satellite Remote Procedure Calls and data transport, offering useful communication patterns and will be the primary API for hosted applications from experimenters. An on-board autonomy watchdog ensures safety of flight while enabling advanced algorithm demonstrations (e.g., artificial intelligence / machine learning (AI/ML)).

An earlier risk reduction payload, Edge Node Lite, comprising a subset of the sensing and computing hardware will launch as a hosted payload in early 2025. This risk reduction demonstration mission includes multiple advanced image-processing and machine-learning algorithm demonstrations provided by multiple mission partners.

BACKGROUND

Advancing science of AI/ML for RPO through on-orbit testing is very limited at present. Additionally, small form factor RPO sensors are unproven on-orbit. There is a bow wave of unprecedented, unproven capabilities on the ground that are waiting for opportunities to be demonstrated in space to address autonomous inspection, tracking, maneuver planning, and collision avoidance.

The advancement of autonomy capabilities will allow on-orbit services (OOS)¹ and in-space servicing, assembly, and manufacturing (ISAM) to be achieved at lower cost². As the Consortium for Space Mobility and ISAM Capabilities (COSMIC) Management Entity³, The Aerospace Corporation has collaborated to identify flight testbeds that allow opportunities for test and demonstration of ISAM as a key enabler for advancement of national space priorities^{4,5}.

In addition to providing an in-space algorithm demonstration platform, Edge Node will advance the adoption of compact commercial off the shelf (COTS) computing and sensing technologies not originally designed for space applications. Edge Node will demonstrate sensors originally designed for autonomous driving, computing electronics typically used in terrestrial robotics, and open-source software intended for terrestrial or airborne applications.

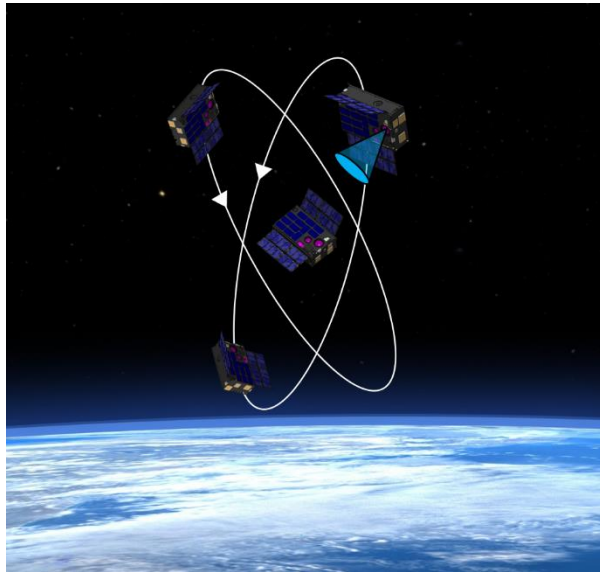


Figure 1: Edge Node Inspection

PROXIMITY OPERATIONS SENSORS

Cepton Lidar

Edge Node will use a Cepton Nova automotive lidar sensor to collect 3D point cloud information. The Cepton Nova weighs approximately 0.5 kg with dimensions of 50 x 50 x 100 mm and achieves a FOV of 120° horizontal and 90° vertical. The main components of the sensor are a near-infrared (NIR) laser source and photodetector that move in sync to scan the FOV. The design is low size, weight & power (SWAP) and rotation free which allows the sensor to operate internally from a 6U CubeSat. Cepton has provided pre-production models of the Nova lidar sensor for the development of Edge Node; the exact specifications of the sensor will vary.

Data from the Cepton Nova is transferred through a 1000BaseT1 (automotive gigabit ethernet) interface. Each point contains a time stamp, xyz coordinates, reflectivity value, segment number and a flag. The segment number refers to one of 48 segments that make up the FOV as shown in Figure 4. The flag labels each point as saturated, noise or a second return point which can be used to distinguish returns from semi-transparent surfaces. The scan pattern of the Cepton Nova produces higher point densities on the edges and the corners of each individual segment and the edges/corners of neighboring segments overlap which results in point density variations across the point cloud. Data volume is approximately 15 MB per second of acquisition for the complete FOV; however, the expected data volume will

True Distance	2.28 m	5.31 m	10.25 m	20.28 m	40.26 m
Number of points in 20 sec acquisition	89,039	43,206	15,273	2,961	1,177
Target distance value from clean points	2.29 ± 0.05	5.27 ± 0.05	10.20 ± 0.05	20.02 ± 0.05	40.01 ± 0.05




Figure 2: The Table shows the number of points in a 20 second Cepton lidar acquisition of the model CubeSat at different distances, the average measured target distance and the standard deviation are also provided. Images show the resulting point cloud at select distances. The color scale indicates the distance with green points being further than blue. The large apertures on the surface are clearly identified by green patches at 2.28 m but only slightly visible at 10.25 m. At 40 m the general shape of the model CubeSat is no longer distinguished.




Figure 3: Top Left: The model CubeSat mounted on a robotic arm at the CAVE lab for dynamic testing. Right: The model CubeSat levitated with a pulley system at the CAST lab. Bottom Left: An example of iterative closest point (ICP) registration with point cloud data from the Cepton Nova.

be much lower for Edge Node operations since most targets will only occupy a sub-region of the point cloud.

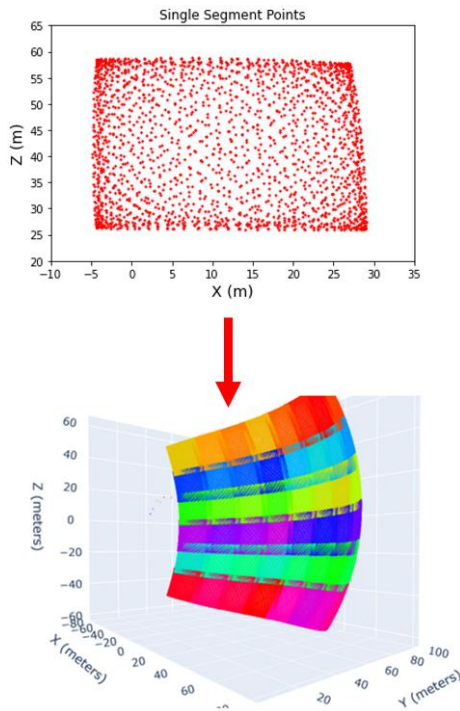


Figure 4: Top: A single segment of data points from the Cepton Nova. Point density is higher on the edges and corners of the segment. Bottom: A full point cloud covering a 120° X 90° FOV with 48 individual segments. Multiple acquisitions have been overlaid to visualize the overlapping regions between the segments.

The performance of the Cepton Nova has been characterized using a model 6U CubeSat. Figure 5 shows the point cloud of a static 6U CubeSat. For this data the center of the CubeSat model is located 2.0 m from the face of the Lidar sensor. At this distance, the solar panels on the edges of the CubeSat model and the 38 mm aperture on the bottom corner of the CubeSat can be resolved; however, flight safety requirements will prevent data collection at this distance. Figure 2 shows the effect of increasing distance on the number of detected points. At a safe distance of 10 m, a 20 second acquisition results in 15,000 points. This is sufficient to partially resolve the 140 x 63.5 mm apertures on the face of the CubeSat model. As the distance increases, surface features are no longer resolved and at 40 m it is no longer possible to distinguish the shape of the model CubeSat.

The Cepton Nova has also been tested with a dynamic target. Dynamic tests were performed at Aerospace’s Collaborative & Autonomous Vehicle Ecosystem (CAVE) Lab and at the Center for Autonomous Systems and Technologies (CAST) at the California Institute of Technology. Both facilities use a network of motion capture cameras to track the position of a dynamic target as it moves on a robotic platform or is suspended from a pulley system as shown in Figure 3. Data from the motion capture system can be used as truth data and compared to data from the Cepton Nova. Point cloud data from dynamic tests is processed using an iterative closest point (ICP) registration algorithm from Open 3D, an open-source library for Python. Figure 3 shows an example of the registration algorithm. Red points show the initial position of the target and blue points show the point cloud after linear displacement of the target. The registration algorithm iteratively determines the best transformation (translation + rotation + scaling) that

must be applied to the red points to achieve the best overlap with the blue points; applying this transformation to the initial point cloud (red) results in the green points. Dynamic data collection and analysis is currently underway. Future data analysis will likely include more advanced registration algorithms and GPU acceleration for real-time, onboard data processing.

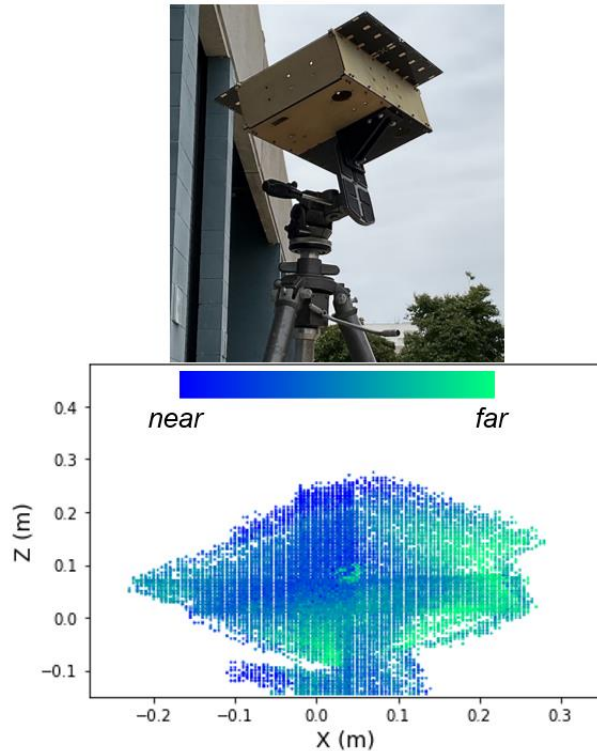


Figure 5: Top: A model of a 6U CubeSat is mounted on an adjustable tripod 2.0 m from the Cepton Nova for static target testing. Bottom: The resulting point cloud. The color scale shows the relative distance from the Cepton Nova. The circular aperture in the bottom face of the model CubeSat produces a patch of green points in the point cloud.

Environmental testing of the Cepton Nova is currently in progress at The Aerospace Corporation. Environmental testing includes ambient pressure thermal cycling, thermal-vacuum, and random vibration testing. The intended concept of operations (CONOPS) will require the sensor to operate for durations of 10 to 90 minutes within its operational range of -30C to 60C. Preliminary results show that this is possible in ambient pressure, but the sensor is triggering its thermal shutoff in vacuum after approximately 10 minutes. Some modifications to

the COTS sensor are in work to improve its thermal performance in vacuum.

mmWave Radar

A COTS radar developed by Texas Instruments (TI) primarily for terrestrial automotive collision avoidance and human detection, was selected for use in object detection and avoidance in space. The goals are detecting an object floating in space at a relative rate of 1 to 3 cm/sec and at a distance ranging from ten (10) to fifty (50) meters. There are several versions available of TI IWR6843 mmWave radar radio board. Because an object in space could be anywhere relative to the CubeSat mmWave sensor, the IWR6843ISK-ODS was selected since it has $\pm 60^\circ$ field-of-view (FoV) in both azimuth (Az) and elevation (El). Although the IWR6843ISK has slightly higher output power (10dBm vs. 8dBm), which is desirable to achieve longer distance detection, it has a more limited FoV of $\pm 15^\circ$ in El. Both the IWR6843ISK and IWR6843ISK-ODS, as seen in Figure 6, have 4 receive (Rx) antennas and 3 transmit (Tx) antennas, allowing the radars to operate in a multiple-input and multiple-output (MIMO) configuration. The layout of these antennas on the board gives the two models different beamwidths and angle resolutions.




Figure 6: TI IWR6843ISK-ODS and IWR6843ISK

TI offers a multitude of demonstration software packages for the user. For initial performance characterization testing, Aerospace used raw Fast Fourier Transform (FFT) output for measuring stationary objects. For dynamic testing and eventual flight applications, the “3D People Counting” firmware was selected due its superior point cloud accuracy on moving objects. This firmware utilizes Capon beam forming to compute a point cloud using the 12 virtual antenna pairs generated from the multiple-input multiple-output (MIMO) processing, as shown in Figure 7.

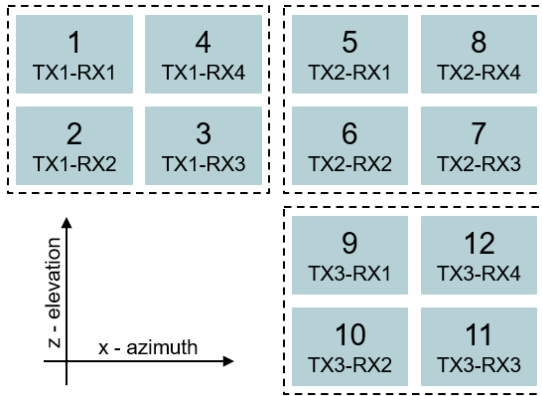


Figure 7: The IWR6843ISK-ODS virtual antenna array produces symmetrical beamwidth and angle resolution in the elevation and azimuth directions⁶.

Material reflectivity has also been a key element in determining suitable materials for use as targets versus test apparatus. For static testing, targets were set on top of Styrofoam stands, which are not detectable at 60 GHz. In general, materials with very low dielectric constants close to air ($\epsilon_r=1.0$) and less than 2.0, worked well as non-detectable targets. Conversely, higher dielectric constant materials worked well as detectable targets, similar to the cubesat exterior metal frame. For dynamic testing, it is advantageous to use a test system that can perform specifically defined and repeatable motions. Typically, such motion systems are large metal robotic structures, which appear as very large targets to the 60 GHz radar.

Testing was performed in a variety of facilities: anechoic chamber, parking lot, Collaborative and Autonomous Vehicle Ecosystem (CAVE) Lab at Aerospace, Center for Autonomous Systems and Technologies (CAST) at The California Institute of Technology (Caltech). Both CAVE and CAST have an independent camera sensor system that provides an independently calibrated source for object location verification. Each test location has different limitations, most commonly distance and interference (Table 1).

Table 1: Range Testing Limits by Test Facility

Test Facility	Max Test Distance (m)
Aerospace Anechoic Chamber	3.5
Aerospace CAVE Lab	6
Caltech CAST Lab	8
Aerospace Parking Lot	42

Initial Phase I testing focused on static testing, which has the advantage of being able to subtract the baseline background to isolate the intended target(s). In all test facilities, an initial baseline static radar measurement is

made with no targets. Next a known target is measured at various distances to confirm the radar is operating as expected. Preliminary static testing was performed in one of the Aerospace anechoic chambers. Although it is a relatively small facility, it serves as a controlled test area to perform testing on different radar radio configurations and settings (Figure 8 and Figure 9). To date, static testing has verified the TI IWR6834ISK radar can detect a 6U Cubesat at 42m (Figure 10 and Figure 11). This addresses the distance task objective. Static testing both in the parking lot and in the CAVE lab verifies the ability to identify multiple targets simultaneously at various distances nearly perpendicular to the radar. Figure 12 - Figure 18 show parking lot, CAVE and CAST lab set-ups along with data plots verifying multiple target locations, which verifies radar distance measurement in 2-dimensions (range and azimuth or elevation).




Figure 8: Initial Testing in Aerospace Anechoic Chamber (A2-2259)

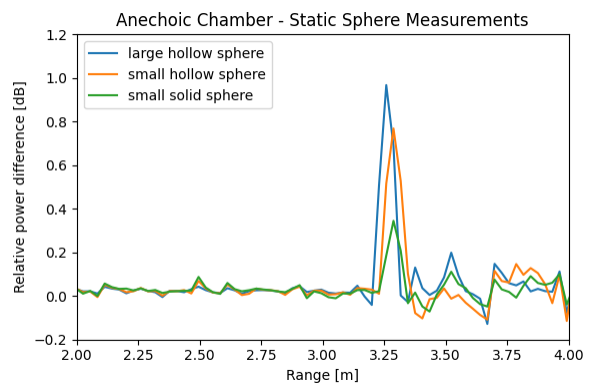


Figure 9: Range versus Target Size in Anechoic Chamber



Figure 10: Parking lot 1D distance test set-up, sphere on Styrofoam stands.

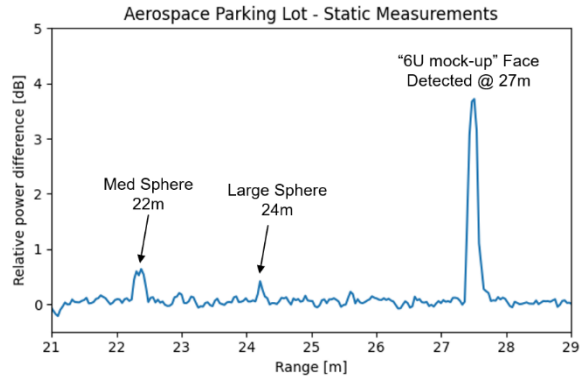


Figure 13: Data plot verifying the radar detected multiple targets in parking lot.

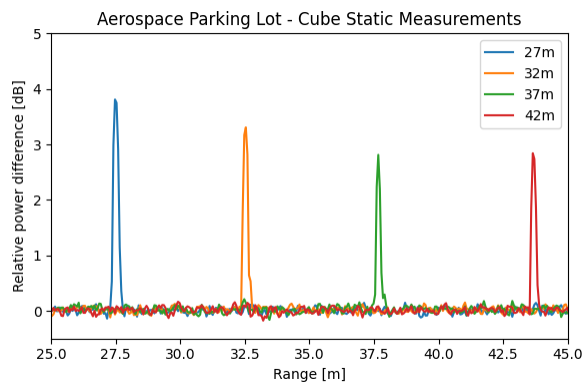


Figure 11: Data plot verifying mmWave radar identified objects at 27m, 32m, 37, and 42m. Mainly 1D target detection.




Figure 14: CAVE lab 2D range/azimuth and multiple target test set-up.




Figure 15: Data plot verifying the radar 2D multiple targets detection in CAVE lab.



Figure 12: Parking lot 2D range/azimuth and multiple target test set-up.



Figure 16: CAST lab 2D radar testing in range and elevation test set-up. Target at various range heights.

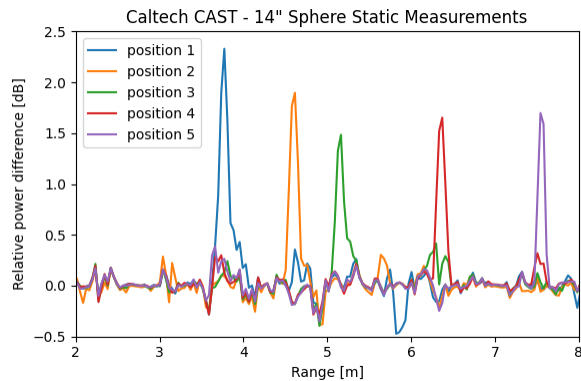


Figure 17: Data plot verifying the radar detected the 14" diameter RCS sphere at various heights in CAST lab.




Figure 18: Variation in returned signal level vs. distance in Figure 17 is believed to be related to variable antenna gain versus angle⁶.

Another aspect of radar detection is the configuration of the target itself. Since the object of interest for this mission is a 6U cubesat, 6U mock-ups were constructed. When the transmitted radar signal is perpendicular to the incident surface, there is maximum reflection. As the angle of incidence diverges from 0° , the amount of reflected signal quickly diminishes. This was verified by testing of a perfect 6U mock-up #1 (Figure 19) with no cubesat features. Figure 20 shows the returned signal response versus angle. To enhance the realism of this response, solar panels were added. For the Edge Node mission, two sets of solar panels will be deployed in fixed 90° and 150° angles, relative to the $\pm X$ 3U panels (Figure 21). Figure 22 shows the improved radar response of the 6U mock-up #2 with added solar panels. Since the solar panels add angles to the $\pm X$ (3U) and $\pm Y$ (6U) panels, they have no effect on the $\pm Z$ (2U) face responses. Mock-up #3 is under construction which will accurately model all the Edge Node 6U features on all panels. The detailed physical features of protruding and recessed structures are expected to further enhance radar detection.




Figure 19: 6U mock-up #1 (featureless faces).




Figure 20: Radar response of featureless 6U mock-up #1 shows reduced response as the angle of incidence increases.




Figure 21: 6U mock-up #2 with deployed solar panels in as-flown configuration.

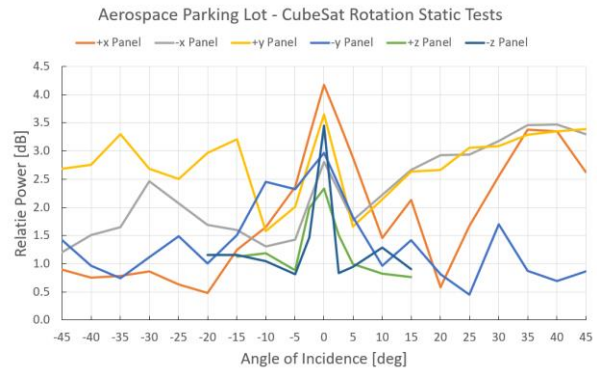


Figure 22: Radar response of mock-up #2. Improved response on 3U and 6U faces versus angle.

Dynamic testing using existing facilities is challenging due to two factors: 1) structures used for controlled motion are equally or in some cases more detectable by the radar than the intended target; 2) large permanent reflecting features in the room, such as 4'x4' electrical panel or 3-story metal chase. As the target moves, the room reflections also change. When operating in space, room reflections will not be an issue. In order to accurately determine performance in space, efforts are being made to minimize interference caused by the room and test fixtures.

When the radar is pointed straight up to the sky, the noise floor of the radar is established. Test facilities and set-ups are then modified by adding absorber material in front of room corners, large metal structures. Figure 23 is a comparison of baseline noise floors for some test facilities used.

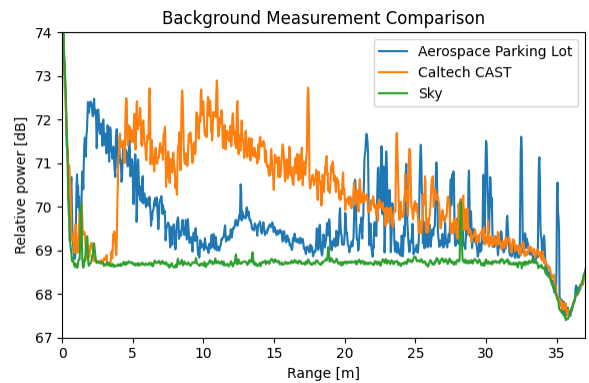


Figure 23: IWR6843ISK-ODS noise floor for test facilities used.

The test aids in the following section are designed to decrease interference from reflections and reduce the measurement noise floor.

In the Aerospace CAVE lab, the ridgeback robot and Kinova robotic arm enable repeatable user-defined motion trajectories to simulate the 6U cubesat motion in space, be it straight up-and-down, back-and-forth, along a diagonal as well as tumbling motion and multi-axis motion. Since the ridgeback is a larger radar target than the cubesat, an “absorber fence” (Figure 24) was designed to “hide” the ridgeback. Prior testing using a kluge of pyramidal absorber sheets verified the absorber effectively masks the ridgeback. Since the CAVE lab is a multi-use facility and cannot be transformed into an anechoic chamber, “absorber walls” (Figure 25) were constructed to be placed on the perimeter of the room to reduce corner and wall reflections and subsequent signal interference.




Figure 24: Absorber fence on top of ridgeback robot. Proof of concept (left). Final product (right).



Figure 25: Portable absorber wall, with pyramidal absorber

Reflections for all surfaces, including the radar mounting plate, affect the measurement. For terrestrial characterization purposes, we also added an absorber plate (Figure 26) in front of the radar board. Figure 27 shows the improvement in noise floor and received signal response due to the absorber. This test also verifies

the geometries of the planned cubesat radar cut-out in the mounting plate, better matching the spacecraft mounting configuration. The cut-out needs to be large enough to not interfere with the radar field of view. Preliminary testing indicates the opening is sufficient with respect to the distance the board is recessed inside the 6U frame. The absorber mounted in front of the radar board is useful for the ground testing campaigns where additional reflections are difficult to avoid. The spacecraft exterior in front of the radar detector will not incorporate this absorber.




Figure 26: mmWave board on test fixture (top) and with absorber face plate installed (bottom).

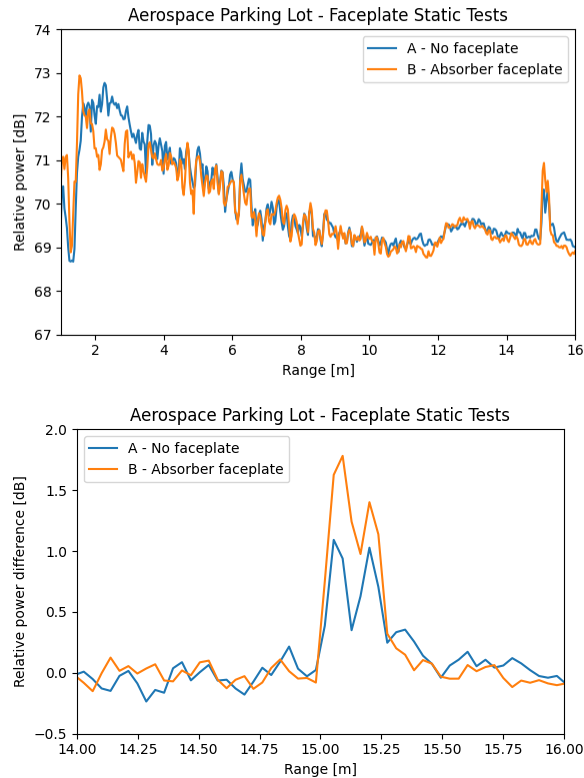


Figure 27: Plots of radar response with and without absorber on faceplate. (Top) From 2m to 4m, there is a reduction in the close-range response with the absorber mounted on the faceplate. At 15m, the received signal amplitude is higher with the absorber on the faceplate (top and bottom). Tests were done with a 21” diameter sphere at 15m in parking lot.

Phase II will focus on dynamic testing. Testing will be performed in the same facilities used for static testing with the addition of the in-house developed absorber test structures. In addition, preparations are in progress to use a drone to move the target.

Laser Rangefinder

Edge Node will use a Jenoptik DLEM 30 range finder for long distance targets. The DLEM 30 uses a 1550 nm laser source with a 0.7 mrad divergence and is capable of measuring targets between 10-14,000 m. The weight of the sensor is 95 g and has dimensions of 97 x 25 x 50 mm. The sensor consumes less than 2W of power and has an operating temperature range of -40° to 80° C. The design of the DLEM 30 is based on the flight proven DLEM SR which has previously flown on AeroCube Optical Communications and Sensor Demonstration satellite and the Seeker-1 program⁷.




Figure 28: Jenoptik DLEM 30 Laser Rangefinder

Longwave Infrared Cameras

Each Edge Node spacecraft will be equipped with two longwave infrared (LWIR) cameras with a spectral range of 8 μm – 14 μm . Both cameras have noise equivalent differential temperature (NEDT) of ≤ 20 mK which extends detection, recognition, and identification (DRI) performance. Each camera can record frames at a rate up to 60 frames per second. The cameras are subject to International Traffic in Arms Regulations (ITAR) control per category 6 (6A003.b.4.b).

The first camera is the Teledyne FLIR BOSON+ model 22640A008–6IAAX which has an 8° Horizontal Field of View (HFOV), a 55 mm focal length lens, and resolution of 640 x 512 pixels. This imager is intended for tracking the other Edge Node spacecraft or targets of opportunity at ranges exceeding 50 meters to aid in maneuver planning for closer proximity approaches. The maximum range where detection and tracking are possible is a characterization objective for the mission. Variation in background scene, lighting, target vehicle orientation, and thermal conditions are expected to statistically influence the maximum tracking range.

The second camera is the Teledyne FLIR BOSON+ model 22640A012–6IAAX which has a 12° Horizontal Field of View (HFOV), a 36 mm focal length lens, and resolution of 640 x 512 pixels. This camera’s lens will be focused to maximize performance at close range of 10 meters. This will aid in close range tracking, feature identification, and pose estimation of the target spacecraft. The lens and sensor resolution will allow detection of features 5 mm and larger at a range of 10 meters. Feature identification is assumed possible so long as the feature is greater than or equal to the size of two pixels projected onto the image plane.

Visible Cameras

Edge Node will host two visible cameras manufactured by Teledyne FLIR: one Blackfly S USB3 color camera, and one Blackfly S USB3 polarized monochromatic camera. The color camera uses a Sony IMX250 monochrome and color imaging sensor. The polar-mono camera uses a Sony IMX250MZR polarization image sensor, which boasts a four-directional on-chip polarizer. Both cameras will be fitted with 12mm/F1.8 lenses. Both are rated at 75 FPS and have a resolution of 5.0 megapixels.

Instead of being attached above the on-chip lens layer of the CMOS sensor, Sony's IMX250MZR puts the polarizer under that layer during the semiconductor process. Each cell of the sensor has its own polarizing filter, with quadrants oriented to 0, 45, 90, or 135 degrees of polarization as shown in Figure 29. In software, each pixel is therefore a matrix of four sub-pixels, which can be indexed using the FLIR-provided Spinnaker software development kit. This enables powerful post-processing techniques, like choosing the darkest pixel (with the smallest degree of polarization) from each polarization quadrant to create a glare-reduced image. Glare reduction can be useful for remote sensing applications that deal with images of water, snow, ice, or human-made reflective surfaces which could include spacecraft. An example of images with glare reduction applied is shown in Figure 30. By analyzing the angle of linear polarization at each pixel, the polarized images can also be used to pick out camouflaged reflective objects from nonreflective backgrounds^{8,9}.

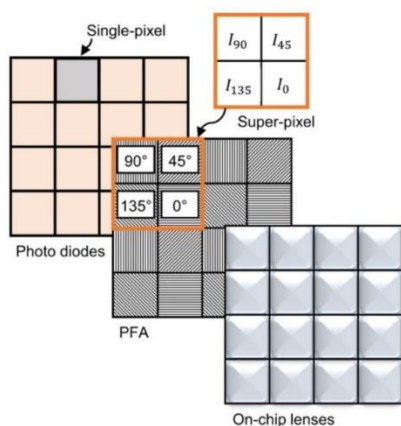


Figure 29: Structure of the Sony IMX250MZR polarizing image sensor, found in one of Edge Node's visible Blackfly cameras. Linear polarizers form the polarizer filter array (PFA) layer, and an on-chip micro-lens covers each polarizer. Sets of four neighboring pixels with different polarizing angles form "super-pixels" [17].

The four so-called Stokes parameters are used to describe the polarization state of a light beam. Only the first three Stokes parameters relate to linear polarization, so only these apply to measurements from the IMX250MZR:

$$S_0 = I = P_0 + P_{90} \quad (1)$$

$$S_1 = Q = P_0 - P_{90} \quad (2)$$

$$S_2 = U = P_{45} - P_{135} \quad (3)$$

S_0 is the intensity of the light beam, calculated by adding the intensities of the vertically and horizontally polarized pixels. S_1 is the difference between the horizontal and vertical components. S_2 is the 45° component. These three parameters can be used to compute degree of linear polarization (DoLP), the proportion of light that is polarized at a given pixel, as:

$$\Pi = \frac{\sqrt{S_1^2 + S_2^2}}{S_0} \quad (4)$$

The Stokes parameters can also be used to compute angle of linear polarization (AoLP), the average polarization angle of the light at a given pixel, as:

$$\theta = \arctan\left(\frac{S_2}{S_1}\right) \quad (5)$$

DoLP processing can help to increase image contrast and distinguish image features. AoLP processing can be used to segment images based on planar surfaces. DoLP and AoLP can also be used to produce useful false color images that highlight irregularities or stresses in material surfaces, a possible application for spacecraft inspection missions¹¹. Examples of these processing techniques are shown in Figure 31.



Figure 30: Using the FLIR Blackfly polarized camera to reduce glare seen on reflective surfaces. The minimum value is selected from each super-pixel, producing the glare-reduced image seen on the right.

Aside from being used for capturing individual color and polarized image frames, the two cameras will be operated in a stereo mode, with a separation of 340.11mm between lens boresights. One camera is designated as the "primary" camera and triggers the "secondary" camera via a GPIO connector cable. Teledyne FLIR provides this synchronization hardware and software, which yields two images whose exposures

begin within microseconds of each other. Each four-quadrant superpixel of the polarized image is collapsed by selecting the subpixel with the smallest degree of polarization. Both images are then converted to monochrome. The OpenCV software package is used along with a real-life 2D checkerboard pattern to perform static calibration of each of the two cameras, which records matrices to correct for lens distortion and map 2D image points to 3D world points. OpenCV is then used to perform stereo rectification, a process which matches keypoints in the two images and computes reprojection matrices to reproject both images onto a common plane parallel to the line between the camera centers. The two reprojected images' epipolar lines align vertically, meaning that the match for a pixel in the left image can be found along the horizontal line with the same y coordinate in the right image as shown in Figure 32. Finally, built-in OpenCV functions are used to match pixels in the two images and display the x disparity of each pixel pair as a relative depth map^{12,13}.

Edge Node aims to test the use of stereo depth imaging for building coarse analysis and awareness of nearby spacecraft, debris, or other objects. Stereo depth image maps could be combined with object classification neural nets and fused with other sensor data like LiDAR or laser rangefinder measurements to build a real-time model of the satellite's surroundings. Compared to using a single visible RGB camera, Edge Node's use of the polarized camera and stereo setup is expected to enable better feature detection when imaging target spacecraft. Stereo depth imaging is also expected to allow for estimating targets' pose and rotation rate.




Figure 31: Images from the FLIR Blackfly polarized camera. Degree of linear polarization (top) and angle of linear polarization (bottom) are computed at each image pixel, allowing for contrast and segmentation in the processed images. In the bottom images, HSV color mapping is used, where the polarization direction is stored in the hue channel and the degree of polarization is stored in the saturation channel.

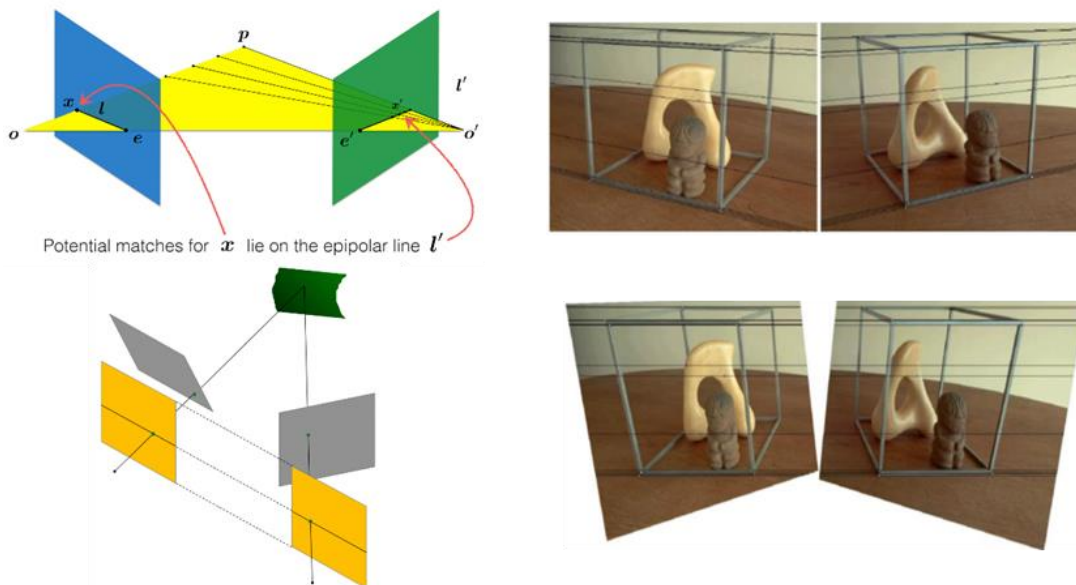


Figure 32: A visualization of stereo rectification, a technique used to enable stereo depth imaging on Edge Node. According to the “epipolar constraint” (top left), potential matches for the left camera’s pixel x must lie on the right camera’s corresponding epipolar line. Reprojecting the two images makes it so that epipolar scanlines occupy matching rows, simplifying the scanning process for stereo pixel matching^{12,13}.

PAYLOAD COMPUTING HARDWARE




Figure 33: Edge Node Backplane Bottom Board, NVIDIA TX2 NX Mounting




Figure 34: Edge Node Backplane Bottom Board, NVMe Mounting




Figure 35: Edge Node Payload Computing Electronics Configuration

EMBEDDED HARDWARE

The computational cores of the Edge Node system are twin TX2 NX supercomputers. While they are not the latest generation, they have a large body of test data for the space environment, software testing experience and documentation. Each TX2 NX contains a heterogenous CPU topology consisting of a Quad-core ARM processor and dual-core NVIDIA Denver 64-bit CPU for conventional processing. Each TX2 NX also contains 256 NVIDIA CUDA cores for sensor data processing and higher-level computational work.

The primary data interfaces between TX2 NXs and other devices on the satellite's data network is through Gigabit Ethernet. Each TX2 NX contains a PHY interface which connects to a 7 port Gigabit Ethernet switch. The SAMV71Q21RT is also a node on the Ethernet network over a RMII interface using its built-in RMII hardware.

A SAMV71Q21RT orchestrates the electrical operation on the Edge Node System. The SAMV71Q21RT is one of the few components which turns on by default when the Edge Node payload receives power, and it also contains a low-level communications path to the spacecraft bus. When directed, the SAMV71Q21RT activates the TX2 NXs and various payloads according to needed mission operations. Individual sensors can be turned on individually or by groups based on target distance and sensor range of use. All power switching is done with soft-start circuits to minimize inrush currents.

The TX2 NX modules are thermally connected to an outer wall of the spacecraft structure which spreads the waste heat to better reject that heat to space. The NVMe drives also have a heat transfer and heat spreading thermal design around them to prevent them from reaching excessive temperatures. The TX2 NX modules combined can create ~30W of waste heat continuously during computationally intensive operations.

EMBEDDED SOFTWARE

Edge Node leverages open-source software (OSS) and commercial off the shelf (COTS) hardware to reduce development costs and enhance users' accessibility. The software stack that Edge Node provides features a customized Linux OS based on Nvidia Jetson Jetpack, a Robot Operating System (ROS) environment with comprehensive tools, middleware, and a suite of containerized applications built on the Docker platform.

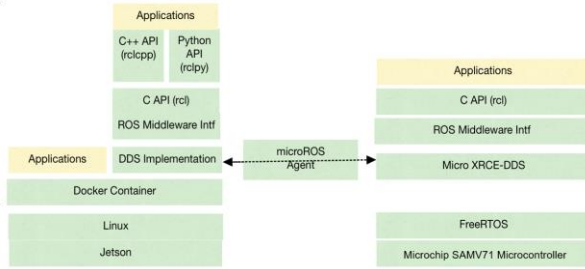


Figure 36: Edge Node Software Stack

Sensor Suite and Communication Interfaces

Edge Node contains multiple sensors for long, medium, and short ranges. WFOV camera, mmWave radar and LIDAR are targeted for short range feature extraction while Laser Range Finder, NFOV and IR cameras are suitable for long and medium range target detection.

Satellite-to-satellite data exchange is facilitated using WIFI and S-band crosslink technology. Such infrastructure and network topology are developed to enable seamless multi-domain ROS communication across several wireless medium, with Quality of Service (QoS) provisions for critical data exchange with respect to priority, real-time performance, and bandwidth.

Edge Node developed a custom carrier board based on radiation tolerant SAMV71 microcontroller for its

payload computing with support for state of health monitoring, and power controls. The system is designed to operate with either or both of its two computing payloads based on processing demands and sensor requirements. Communication between the carrier board and the payload computing is conducted over Ethernet with micro-ROS agent acting as the bridge between bus avionics and Jetson hosts. SatCat5 (mixed-media Ethernet switch) functions as physical layer data exchange hub for bus avionics, designating Ethernet as the primary communication protocol in the CubeSat.

Embedded Linux Environment

Edge Node computing hardware runs Nvidia Jetpack SDK and Linux for Tegra (LAT) board support package. A customization layer provides Edge Node specific hardware support along with tools for building and generating application updates and patches.

Data reliability and integrity is a major concern in space environment. Edge Node addresses that by using ZFS as filesystem for its Western Digital IX SN530 nonvolatile memory express (NVMe) storage drives. The built-in data redundancy and error correction features provide resiliency to data corruption caused by radiation induced bit flips or data loss due to sudden power shutdown. A custom recovery mechanism is also developed in U-Boot bootloader that restores the default ZFS data set to the main storage drive's filesystem on a ground command.

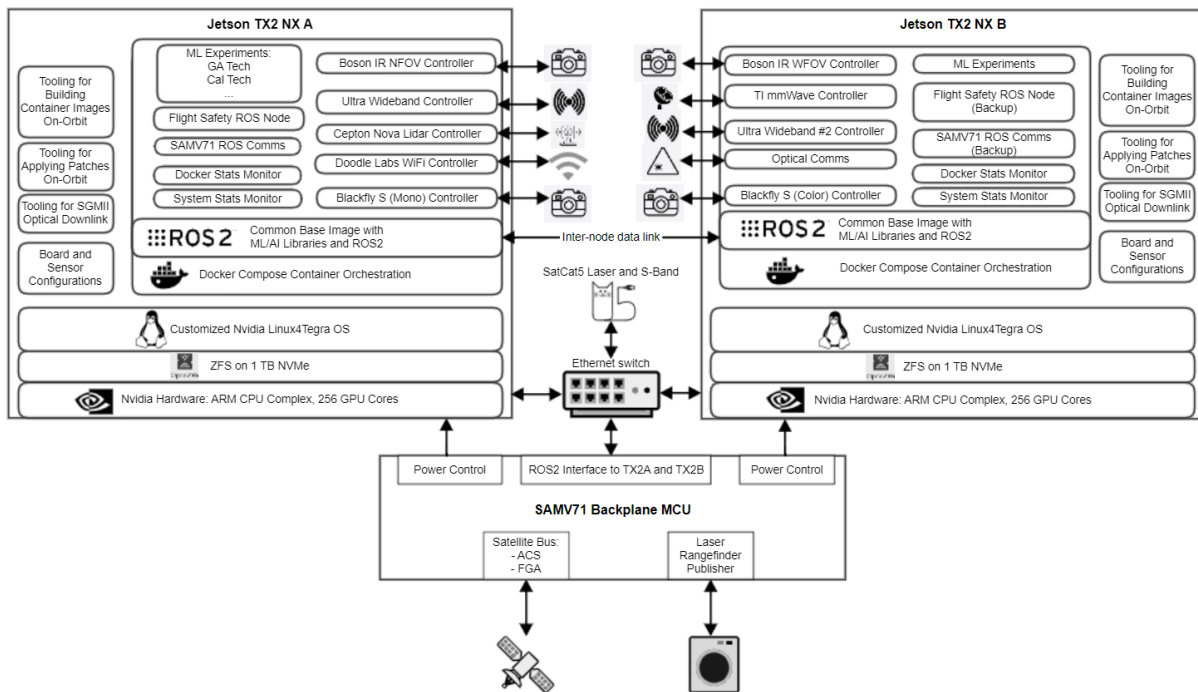


Figure 37: Payload Software Architecture

The Edge Node storage hardware was screened via a testing campaign at Loma Linda Proton Beam facility and shown to have a ten times reduction in single event functional interrupts (SEFI) compared to two alternate NVMe drives tested¹⁴. The radiation testing supports an estimated ~0.5 SEFI per year from a representative low Earth orbit of the International Space Station. The Aerospace Corporation has previously demonstrated the NVIDIA TX2's performance in low Earth orbit via five prior small satellite spacecraft missions.

Computing and Simulating Environment

Robot Operating System (ROS) offers a flexible framework for spacecraft and payload flight software. By interfacing with ROS topics and services, the flight software can efficiently generate and process data through a reliable publisher-subscriber communication model.

Sub-systems are organized as ROS nodes with defined message and service interfaces (published as part of the Edge Node SDK). Spacecraft telemetry, state of health and sensor data are made available to payload computing suite as ROS topics. Additionally, real-time spacecraft statuses and experimental data are captured with “rosviz” tool, facilitating ground telemetry and analysis during flight missions.

Flight software are built, tested, and deployed as Docker containers. The containerization solution allows system services and experimental applications to each have isolated sets of dependencies and filesystem access. A Docker image called “edgenode-base” provides a common base for applications enabling reduced disk usage and isolating software dependencies. The base image also provides NVIDIA platform specific hardware accelerated GPU drivers, high level software, libraries, and frameworks for common applications such as machine learning and artificial intelligence. The integration of ROS and Docker technologies enables payload software development to seamlessly transition from ground testing to in-orbit experimentation.

SPACECRAFT SYSTEMS

Overview




Figure 38: Edge Node Spacecraft Configuration with RPO Sensors




Figure 39: Edge Node Spacecraft Configuration with Internal Components

The Edge Node satellites are two identical 6-unit (6U) CubeSats with stowed dimensions of 36.6 cm x 22.6 cm x 11 cm and a wet mass of approximately 9.9kg. The satellites have two 30cm x 20cm bi-fold wings, with 14 solar cells per side, mounted on opposite edges of their -

X and +X faces. On their +Y face, the vehicles have 12 additional solar cells. Both wing and body solar cells are Spectrolab XTE-SF, with an area of 27cm² which provide an orbit-average solar input power of 10W. During typical day in the life, Edge Node will use ~6W to maintain a +Y face sun tracking, and 15-20W for periods with radio or optical communication to ground stations lasting up to 10 minutes per contact. In eclipse, during communication links, or RPO experiments, the vehicle is powered by two batteries, each comprised of 4 Lithium-Ion Molicel INR-18650-M35A cells in parallel, with a total capacity of 100Whr, which provides sufficient power budget for at least 50 minutes of long-range tracking or close inspection utilizing the RPO sensor suite and dual NVIDIA Jetson TX2 NX modules, which can result in a total system power demand between 50-60W. The vehicle body comprised of six aluminum 6061-T6 machined panels with surface treatments of black anodization or silver Teflon to increase emissivity. Payloads and avionics are connected via copper thermal straps, and for high power components such as the Jetson TX2 NX, mounted on aluminum trays with their heat sinks directly against the vehicle body.

Attitude Determination and Control System

The attitude determination and control system (ADCS) is heritage to multiple prior missions: Optical Communications and Sensor Demonstration (OCSD)¹⁵, AeroCube-10¹⁶, R3¹⁷, Rogue Alpha, Beta¹⁸. The attitude determination system is highly dependent on the SiOnyx detector-based star trackers¹⁹. The ADCS system has demonstrated 0.01 deg of pointing knowledge and payload open-loop precision pointing of about 0.024 degrees while slewing at rates of up to 1 degree per second.

Communication System

Edge Node has two space-to-ground communication paths and two crosslink paths. The primary space-to-ground path is via the SCR-504, a commercial off the shelf (COTS) S-Band radio procured from Innoflight which supports uplink and downlink data rate of up to 1 Mbps and 2 Mbps respectively. The SCR-504 can additionally be used for full duplex 100kbps crosslink at range of up to 100km. The radio has a splitter to patch antennas on both +Z and -Z faces shown in Figure 38 and Figure 39. to increase signal omnidirectionality and enable link closure while tumbling. The second space-to-ground pathway is via laser communication transmitter and is downlink only. The lasercomm system is capable of data rates up to 400Mbps with Forward Error Correction (FEC), transmits at 4W, at a wavelength

of 1064nm and has flight heritage on multiple previous AeroCube missions^{15,17,18,20,21}. The final intersatellite communication pathway is through Doodle Labs RM-1700-21N Mesh Rider WiFi Radio with a single 2.45GHz frequency band²², allowing for password protected authentication and AES-256 encrypted communication across an arbitrary number of satellites at 100-meter range with data rates of ~25Mbps and omnidirectional coverage. Longer ranges may be possible with preferential antenna orientation.

Propulsion System

Edge Node will utilize a steam based propulsion system that is heritage to two prior missions: Optical Communications and Sensor Demonstration (OCSD)¹⁵, and AeroCube-10¹⁶. The propulsion unit shown in Figure 40 for Edge Node has propellant tank dimensions of 300 cm³ with capacity for 190 grams of water. It is estimated that this propulsion unit will be capable of producing greater than 10 m/s of delta-velocity when the usable propellant is fully depleted.




Figure 40: Edge Node Steam Thruster

COLLISION AVOIDANCE

Edge Node must engage in proximity operations closer than 50 meters—and preferably as close as 10 meters—to collect the desired resolved imagery between vehicles with a variety of orientations, backgrounds, and dynamical behavior. The combination of close range and varied relative behavior implies the need for active rendezvous and proximity operations (RPO) behavior such as natural motion circumnavigation (NMC). More passive behavior, such as an in-track “perch” would be inadequate to accomplish all the mission’s goals. However, close-range RPO raises the concern of the vehicles’ possibly colliding in poorly designed operations or due to anomalous behavior. Edge Node is implementing several solutions in its operations to reduce the risk of collision and ensure flight safety.

Both Edge Node vehicles will be capable of monitoring the other's absolute and relative positions by sharing their on-board navigation solution via crosslink. The Edge Node GPS receiver is an Aerospace-developed design that has flown on more than twenty spacecraft since 2011²³. Each fix collected by the receiver has an instantaneous position accuracy of ~ 10 m (1σ). The flight software includes on-board navigation that processes these fixes into an ephemeris that has an accuracy at epoch of <5 m. (This position uncertainty from the GPS-derived ephemeris drives the mission's ~ 10 -meter limit on RPO activities.) During flight, the two vehicles will share their ephemerides via crosslink, so that each vehicle can perform safety processing independently (i.e., formation control and flight safety will not be managed by a single node). An active and healthy crosslink is a prerequisite for the initiation and continuation of RPO activities. If the crosslink goes down, the flight-safety system on board will automatically terminate the RPO and separate the vehicles. Similarly, if either vehicle's ephemeris becomes excessively stale (i.e., a long time has passed since the last GPS fix), the RPO will be terminated.

Edge Node's default RPO behavior will be to maintain passively safe formations via the eccentricity/inclination vector ("e/i") control algorithm of D'Amico²⁴. The e/i control approach transforms the relative states of the two vehicles into a pair of relative-orbital-element vectors that are both intuitive and have a straightforward mapping to the relative RPO behavior. Of particular interest for Edge Node, the relative orientation of these e/i vectors (e.g., perpendicular vs. parallel) controls whether the two vehicles ever cross the in-track direction, which would be a higher risk configuration. If the vectors are parallel (or made to be parallel via propulsion and changing the vehicles' orbital elements), the deputy vehicle's trajectory in an RPO always maintains non-zero radial or crosstrack components. In this case, the deputy orbits around the in-track line of the chief. Or, if there is non-zero relative drift rate, the deputy "corkscrews" around the chief's in-track line. In either case, if the radius of the radial-crosstrack motion is greater than the radial and out-of-plane position uncertainties, an operator can be assured that a collision cannot happen, because the deputy will never be in the same location as the chief's direction of motion (the in-track direction), which is usually the most uncertain. This relative orientation of the vehicles is passively safe: as long as the e/i-vector orientation is maintained, the vehicles cannot collide irrespective of in-track position uncertainties. This passively safe corkscrew behavior meets most of the needs of the Edge Node mission: the rotational corkscrew motion (essentially a drifting NMC) provides dynamic relative behavior and many look angles, and reducing the radius of the corkscrew

allows for precise control of the relative range while maintaining a safe configuration.

Aerospace has previously used this e/i-vector control scheme to accomplish close-range RPO and satellite-to-satellite resolved imagery. The AeroCube-10 mission, which was a pair of 1.5U CubeSats on orbit from 2019-2023, performed a series of close approaches and imaging campaigns in the spring and summer of 2020¹⁶. AeroCube-10 itself was not originally designed for proximity operations and possessed fewer supporting capabilities than Edge Node (e.g., AeroCube-10 had no on-board navigation and all RPO activities were planned and executed open loop), but the mission succeeded in performing several NMCs and close approaches down to ~ 20 meters range. All the AeroCube-10 RPOs were planned with the e/i-vector control approach, which ensured the safety of these vehicles. During multiple NMC and "corkscrew" activities, one AeroCube-10 vehicle captured resolved imagery of the other.

Each Edge Node vehicle will maintain a nominal safety profile by monitoring the relative e/i vector configuration between the vehicles via their shared GPS fixes and on-board navigation solutions. A passively safe configuration is ensured when the e/i vectors are parallel, and the on-board flight safety monitor will track the vectors' orientation to within a pre-specified tolerance. In advance of any RPO maneuver, flight safety will assess whether the maneuver would put the satellites into a potentially unsafe configuration via the e/i vector metric and deny the execution of the maneuvers if so. Furthermore, if an anomaly or other unexpected event triggers a violation of the safety metric, the RPO will be terminated, and maneuvers autonomously executed to return to a safe configuration.

The Edge Node development team anticipates that some collaborators will wish to perform RPO experiments that are not necessarily compatible with the default passively safe configuration. In that situation, the flight safety system can be modified to monitor the projected closest-approach distance, which is a less robust but more geometrically flexible metric. Experiments that require this alternative safety approach will be carefully planned with the operations team in advance.

EDGE NODE LITE

Edge Node Lite is an earlier risk reduction payload comprising a single NVIDIA Jetson TX2 NX and a single LWIR camera, the Teledyne FLIR BOSON+ model 22640A012-6IAAX with a 12° Horizontal Field of View (HFOV). This risk reduction demonstration mission includes multiple advanced image-processing and machine-learning algorithm demonstrations provided by multiple mission partners. The payload

electronics configuration for the objective Edge Node mission was prototyped with the Edge Node Lite payload, shown in Figure 41. The LWIR camera mounting design is shown in Figure 42. The mission hosting Edge Node Lite is manifested to launch in February 2025.




Figure 41: Edge Node Lite Payload Computing Configuration




Figure 42: Edge Node Lite LWIR BOSON+ Mounting Configuration

CONCLUSION

The Edge Node design for an on-orbit RPO testbed has been developed as a community resource to advance

ISAM capabilities via qualification of low-cost compact sensors and advanced onboard autonomous computing. Edge Node will provide a proving ground for inspection, servicing, and assembly concept of operations (CONOPS). It will augment the limited set of in-space demonstration infrastructure to accelerate ISAM technology consistent with national space priorities.

The novel use of Docker containerization and ROS middleware will enable external collaborators to integrate and test their algorithms more rapidly both on the benchtop pre-launch and via on-orbit upload of new algorithm capabilities post launch.

Interested collaborators are encouraged to contact the Edge Node team to explore joint demonstration opportunities.

ACKNOWLEDGEMENTS

The authors would like to thank Professor Soon-Jo Chung (California Institute of Technology) for his laboratory's collaboration on the Edge Node mission.

All trademarks, service marks, and trade names are the property of their respective owners.

References

1. Davis, J.P., J.P. Mayberry, and J.P. Penn, "On-orbit servicing: Inspection, repair, refuel, upgrade, and assembly of satellites in space," https://csp.aerospace.org/sites/default/files/2021-08/Davis-Mayberry-Penn_OOS_04242019.pdf, 2019.
2. Atkinson, D., S. Chien, and E. Mjolsness, "Applications of artificial intelligence for spacecraft autonomy and enhanced science data return," AIAA Space 2000 Conference and Exposition, Long Beach, CA, September 2000.
3. The Consortium for Space Mobility and ISAM Capabilities (COSMIC). <https://cosmicspace.org/>, 2024.
4. National Science & Technology Council, "In-space Servicing, Assembly, and Manufacturing National Strategy," <https://www.whitehouse.gov/wp-content/uploads/2022/04/04-2022-ISAM-National-Strategy-Final.pdf>, 2022.
5. National Science & Technology Council, "National In-space Servicing, Assembly, and Manufacturing Implementation Plan," <https://www.whitehouse.gov/wp-content/uploads/2022/12/NATIONAL-ISAM-IMPLEMENTATION-PLAN.pdf>, 2022.

6. Texas Instruments, "60GHz mmWave Sensor EVMs User's Guide," 2022.
7. Pedrotty, S., J. Sullivan, E. Gambone, and T. Kirven, "Seeker Free-Flying Inspector GNC System Overview," American Astronautical Society Annual Guidance and Control Conference, Breckenridge, CO, February 2019.
8. Teledyne FLIR, "Save Time and Effort with On-Sensor Polarization," <https://www.flir.com/discover/iis/machine-vision/imaging-reflective-surfaces-sonys-first-polarized-sensor/>, 2022.
9. Sony, "Polarization Image Sensor," https://www.sony-semicon.com/files/62/flyer_industry/IMX250_264_253MZR_MYR_Flyer_en.pdf, 2021.
10. Lane, C., D. Rode, T. Rosgen, "Calibration of a polarization image sensor and investigation of influencing factors," Applied Optics, Vol. 61, No. 6, February 2022.
11. The Imaging Source, "Polarization-based Industrial Cameras: Basics and Applications," <https://photonlines-industrie.fr/wp-content/uploads/2020/12/white-paper-camera-polarsens.pdf>, 2020.
12. Kitani K., "Epipolar Geometry," 16-385 Computer Vision, Carnegie Mellon University, <https://www.cs.cmu.edu/~16385/s15/lectures/Lecture18.pdf>.
13. Kitani, K., "Stereo Vision," 16-385 Computer Vision, Carnegie Mellon University, https://www.cs.cmu.edu/~16385/s17/Slides/13.1_Stereo_Rectification.pdf.
14. Leeds, P., D. Rowen, B.T. Dinh, S.C. Davis, "ZFS in Space: Putting an Advanced, Fault-Tolerant Filesystem to Work for CubeSat Resilience", <https://flightsoftware.org/workshop/FSW2024#>, 2024.
15. Rose, T.S., D.W. Rowen, S. LaLumondiere, N.I. Werner, R. Linares, A. Faler, J. Wicker, C.M. Coffman, G.A. Maul, D.H. Chien, A. Utter, R.P. Welle, and S.W. Janson, "Optical Communications Downlink from a 1.5U CubeSat: OCSD Program," Proceedings of the Small Satellite Conference, SSC18-XI-10, 2018.
16. Gangestad, J.W., C.C. Venturini, D.A. Hinkley, and G. Kinum, "A Sat-to-Sat Inspection Demonstration with the AeroCube-10 1.5U CubeSats," Proceedings of the Small Satellite Conference, SSC21-I-11, 2021.
17. Pack, D.W., D.R. Ardila, E. Herman, D.W. Rowen, R.P. Welle, S.J. Wiktorowicz, and B.W. Hattersley, "Two Aerospace Corporation CubeSat Remote Sensing Imagers: CUMULOS and R3", Proceedings of the AIAA/USU Conference on Small Satellites, Next on the Pad, SSC17-III-05, 2017.
18. Pack, D.W., B.S. Hardy, J.R. Santiago, D. Pietrowski, J.C. Mauerhan, P.F. Zittel, D.W. Rowen, C.R. Purcell, P. Thiyanaratnam, L.J. Gelinias, P.K. Su, and J. Gussy, "Flight Operations of Two Rapidly Assembled CubeSats with Commercial Infrared Cameras: The Rogue Alpha, Beta Program," Proceedings of the Small Satellite Conference, SSC21III-03, 2021.
19. Rowen, D.W., A. Utter, G. Kinum, H. Weiher, A. Wu, R. Dolphus, E. Alcid, D. Galindo, C. Beltran, A. Berman, and M. Mellick, 2020. "On-orbit Results from an Ultra-low SWaP Black Silicon Star Tracker," Proceedings of the AIAA/USU Conference on Small Satellites, Advanced Technologies II, SSC20-XI-02, 2020.
20. Johnson, R.C., J.M. Wicker, A. Lin, E.C. Dillingham, P. Le, J. Esquivel, J.C. Mergo, and D. Rowen, "Development and Deployment of Remotely Operable Optical Communication Terminals," Proceedings of the Small Satellite Conference, SSC22-VI-05, 2022.
21. Rose, T.S., D.W. Rowen, S.D. LaLumondiere, N.I. Werner, R. Linares, A.C. Faler, J.M. Wicker, C.M. Coffman, G.A. Maul, D.H. Chen, A.C. Utter, R.P. Welle, and S.W. Janson, "Optical communications downlink from a low earth orbiting 1.5U CubeSat," Optics Express, Vol. 27, No. 17, August 2019.
22. Doodle Labs, RM-1700-21N Mesh Rider Radio, <https://doodlelabs.bitbucket.io/datasheets/nanoOEM/nanoOEM-1700/>, 2024.
23. Gangestad, J.W., B.S. Hardy, and D.A. Hinkley, "Operations, Orbit Determination, and Formation Control of the AeroCube-4 CubeSats," 27th Small Satellite Conference, SSC13-X-4, August 2013.
24. D'Amico, S., and O. Montenbruck, "Proximity Operations of Formation-Flying Spacecraft Using an Eccentricity/Inclination Vector Separation," Journal of Guidance, Control, and Dynamics, Vol. 29, No. 3, May-June 2006.