SSC26-RAII-05

R5 Spacecraft 7, 9, and 10: Proximity Operations, Optical Comm, and Edge Compute

Sam Pedrotty, James Berck, Jack Wisbiski, Kelly Mann, Andrzej Jackowski, Jacob Nunez-Kearny, Robin Onsay, Faheem Quazi
NASA Johnson Space Center
2101 NASA Parkway, Houston, TX 77058; 281-483-2299
samuel.m.pedrotty@nasa.gov

Thomas Kirven, Rob Stoddard, Jordan Fogg
Amentum Technology
2224 Bay Area Blvd., Houston, TX 77058
thomas.c.kirven@nasa.gov

Roger Hunter
NASA Ames Research Center
M/S 213-13, Moffett Field, CA 94035; 650-604-5004
roger.c.hunter@nasa.gov

Anup Parikh, Alex Bongiovanni, Shrusti Jain, Cassidy Kuehl, Kyle Merry
Sandia National Laboratories
1515 Eubank SE, Albuquerque, NM 87123; 505-845-0122
aparikh@sandia.gov

ABSTRACT

The Realizing Rapid, Reduced-cost high-Risk Research (R5) project has designed, assembled, integrated, tested, delivered, and launched three spacecraft with three unique missions in fiscal year 2026. Each spacecraft has applied lessons learned from the prior to enhance the cost and schedule reductions of the R5 process while reducing risk and maximizing spacecraft capabilities. Spacecraft 7 (R5-S7) completed suborbital proximity operations around a Falcon 9 upper stage in November of 2025 using an industrial LiDAR for relative navigation. Spacecraft 10 (R5-S10) launched on Transporter 16 and is intended to build on R5-S7's mission by completing proximity operations in orbit around a Momentus Vigoride vehicle and will also demonstrate using an event camera for celestial navigation. Spacecraft 9 (R5-S9) will have launched on the Transporter 17 mission and is intended to demonstrate a novel optical communication system in addition to an image processing and autonomy demonstration. This paper describes the missions, spacecraft construction, flight results, and associated lessons learned in the hopes that others can leverage them to improve their own cost, schedule, and technical performance.

R5 OVERVIEW AND BACKGROUND

There are many promising novel technologies and capabilities that have the potential to enhance space exploration. To be utilized operationally, these must first be demonstrated in the target environment. Traditional demonstration missions have substantial cost and schedule, creating a bottleneck for the advancement (and therefore infusion) of these technologies and capabilities. The Realizing Rapid, Reduced-cost high-Risk Research (R5) project is primarily intended to provide a demonstration capability far faster and cheaper than the state of the art. R5 is doing this by pioneering new processes and carefully refining mission cost, schedule, and risk with customers. R5 has launched six spacecraft since 2022 and continues to refine its methods with each new vehicle.

R5 has previously detailed the results of its Spacecraft 2 (R5-S2) and Spacecraft 4 (R5-S4) missions.[1] Since then, R5 has designed, assembled, integrated, tested, and launched Spacecraft 7 (R5-S7), Spacecraft 10 (R5-S10), and Spacecraft 9 (R5-S9). Each of these spacecraft had substantially different missions and were developed sequentially across one year. As with every R5 mission, each of these involved many lessons learned across the entire mission lifecycle. Beyond demonstrating technology rapidly and affordably, R5 has other objectives like broadly disseminating its lessons learned (in the hopes that others can leverage them to realize their own cost and schedule savings), advancing technologies in support of in-space servicing, assembly, and manufacturing, and providing hands-on experience to early career personnel.

PROCESS OVERVIEW

The R5 process is continually being adjusted as lessons are learned during the course of each spacecraft life cycle.[1] These adjustment are made in line with the overall philosophy whose high-level tenets are described in this section. Personnel safety is paramount in the R5 process; even as R5 otherwise reevaluates many long-held beliefs and pushes the envelope from a technical perspective. Personnel safety is one of the few areas where there is no compromise.

After separately and independently establishing personnel safety, R5 works closely with customers to help refine mission concepts to clearly delineate what is required and what is desired. By carefully defining these aspects with stakeholders, R5 is able to effectively scope a "minimum viable product" which maximizes mission success by minimizing and focusing scope (and therefore minimizing cost and schedule) and maximizing benefit. The desired elements are kept as aspirational elements, often classified as "full success" or "stretch goals" and are developed/implemented after the "minimum viable product" as time allows. It's common for stakeholders to confuse desired and required products to demonstrate technologies. It may be helpful to carefully walk through Technology Readiness Level (TRL) descriptions with stakeholders in addition to fleshing out their goals and exploring how telemetry products can directly map to their needs to create an understanding of the "minimum viable product."

After the definition of the mission concept, subsystems start component selection and design. While each R5 vehicle starts with a baseline of the prior iteration, there are general guiding principles that are used across the system during design. The first design principle is using relevant data over unrelated experience or poorly validated models. By using stocked, low-cost components, R5 is able to quickly procure and assess components in time to incorporate the resulting conclusions from the assessment into final system design. This has allowed R5 to feature substantially more capability at far lower cost and schedule than prior efforts. A specific result from these types of assessments is the preference to use automotive-grade (AEC-Q200) components where possible.

Another design principle is to design the system for human use. Throughout the assembly, integration, and test (AIT) portions of the spacecraft lifecycle, humans are going to be physically interacting with the systems. By considering this and explicitly accounting for it in design (meaning making the systems robust to human handling), the AIT process can be reduced considerably. The reduced process results in faster AIT as each activity requires less time and further accelerates AIT as it reduces the likelihood of component repair and replacement. This type of consideration is uncommon in most spacecraft design as only the final operating environment is considered which can result in substantial cost and schedule increases associated with increased process on the ground.

While the subsystems consider and account for human handling in design and component selection, they also attempt to minimize the recovery time from hardware failure. Not only is the ease of assembly considered in design, but the ease of disassembly is as well. Component aspects like cost, lead time, and interfaces are all considered. Multiple copies of all components are procured, prepared, and stored to enable rapid replacement in case of failure. The concern of the likelihood of hardware failure is not unique to small projects with leaner process and higher risk tolerance. Even large/expensive missions encounter hardware failure for a variety of reasons (some quite banal), regardless of the amount of process applied.[2]

After assembly (at various levels), integrated testing is performed as soon as possible. This is also done at the lower possible levels of the assembly. For example, the custom power distribution unit (PDU) circuit board has its functionality verified via test before being integrated with other avionics into a completed subassembly. Functional units are continually tested as additional components are added using tests that have been highly automated ensuring rapid, repeatable execution and clear results. This thorough and incremental test philosophy greatly reduces the fault tree when the inevitable fault occurs, greatly accelerating troubleshooting and recovery. By its need for finished flight-like components, hardware testing often occurs late in a project's lifecycle and can therefore be at the highest risk of getting cut due to schedule pressure. R5 requires complete functional testing before flight as hope is not an acceptable verification artifact. R5 ensures this happens by prioritizing testing in the schedule and instead cutting other scope (like analysis) if required.

SPACECRAFT 7

R5-S7 was the R5 project's first proximity operations (ProxOps) mission and Johnson Space Center's (JSC's) first small spacecraft ProxOps mission since Seeker in 2019. R5-S7 (and Seeker before it) was a step along an incremental development effort of an in-space inspection capability.[3] In-space inspection, especially of crewed vehicles, is incredibly challenging and remains an unsolved capability gap.[4] Extravehicular free-flying inspectors provide a complete inspection capability with relatively low mass that missions can bring along with them, making them a superior option to fulfill the gap.[5] JSC has a history of developing these systems, going all

the way back to AERCam Sprint in the 1990s. The Seeker mission first explored what was possible with regard to developing this capability with far less schedule and budget than prior attempts and did so by leveraging the CubeSat ecosystem, some commercial off-the-shelf (COTS) parts, and a leaner process than the standard Class 1 hardware process followed at JSC. Seeker successfully demonstrated this leaner process, the utility of CubeSats for this application, and the suitability of COTS components while also highlighting the challenges of custom/low-volume-manufacture components and launch opportunities for these systems.[6]

Being R5's first ProxOps mission, R5-S7 was intended to demonstrate not only the performance of the integrated vehicle, but the process used to construct the vehicle (including parts selection) and the ability to procure rideshare-like commercial launch services for this type of unique concept of operations (ConOps). R5-S7 used what was at the time the latest iteration of the R5 process, including the evolution of the bus and components after the conclusion of the operations of R5-S2 and R5-S4 and after the completion of assembly, integration, and test (AIT) of R5 Spacecraft 3 (R5-S3) and R5 Spacecraft 5 (R5-S5). To enable the planned incremental development path for inspection systems, regular low-cost launch opportunities were required. NASA's venture-class acquisition of dedicated and rideshare (VADR) launch services contract was used to assess feasibility and interest from commercial entities for this activity. The requirements on the launch services to be provided were explicitly and carefully minimized to ensure the greatest number of responses and the lowest cost. Ultimately, a launch service for R5-S7 was procured through VADR with a similar cost and schedule to typical rideshare manifests. By using VADR and a commercial procurement, it's hoped that this has created a demand signal across the launch industry and has created a repeatable path for similar missions both from government and industry spacecraft operators to rapidly and affordable fly ProxOps demonstration missions. The resulting launch opportunity for R5-S7 was on SpaceX's Transporter 15 mission and had the spacecraft deploy after the upper stage's deorbit burn, minimizing the risk that the ProxOps activity would generate orbital debris. This suborbital release rideshare demonstration also hopefully enables its use for other applications, such as future demonstration of reentry vehicles like the NASA Hypersonic Testing for Entry Platform (HyperSTEP).[7]



Figure 1: R5-S7 front 3/4 view

Spacecraft 7 ConOps

ProxOps were the core part of the ConOps for R5-S7, but their specific definition and limitations were allowed to be driven by the limitations of the launch opportunity to minimize cost and schedule. Even with a suborbital trajectory, a complete ProxOps demonstration can still be accomplished which substantially derisks future orbital demonstration. The full demonstration of the spacecraft's sensors, algorithms, and effectors can be done with a very simple trajectory and one was preferred to minimize the complexity and risks of the ConOps. As such, the project and launch provider baselined the following ConOps:

1. The upper stage completes deorbit burn
2. The upper stage orients itself to deploy R5-S7 in the positive velocity vector direction (+Vbar)
3. 0 seconds after deployment (D+0): R5-S7 is deployed
4. D+31s: R5-S7 completes power on and starts attempting to send telemetry
5. D+46s: R5-S7 acquires and starts tracking the upper stage, keeping it near the center of its LiDAR's field of view
6. D+47s: R5-S7 starts arresting its range rate along its body X axis
7. D+70s: R5-S7 completes arrest and starts drifting while continuing to track the upper stage
8. D<2700s (TBD): R5-S7 burns up in the atmosphere

Spacecraft 7 System

R5-S7 was based on the lessons learned from the prior R5 vehicles and retained the 6UXL form factor and avionics, communication, flight software (FSW),

guidance, navigation, and control (GNC), propulsion, and structures subsystems.

Spacecraft 7 Avionics

The avionics subsystem was developed from the R5-S5 baseline, which itself was evolved from the R5-S2 and R5-S4 configuration.[17] This baseline consists of a mixture of custom and COTS components which encompass all power generation and distribution, communications, sensors, actuators, and harnessing in the spacecraft. The R5 avionics "stack" of custom printed circuit board assemblies (PCBAs) was developed with the intent of modularity and becoming hardware rich, and this was leveraged for the quick development of R5-S7. Each of the three PCBAs was updated during R5-S5's development and those same revisions were able to be adapted to R5-S7 without any updates, drawing from the stock that was already assembled and functionally verified.

Minor changes from the R5-S5 baseline were required for the S7 mission. The primary addition was a LiDAR sensor for ProxOps. A trade study was performed to identify suitable LiDAR candidates and the MicroVision Movia L 60 LiDAR was selected for its combination of low size, weight, and power (SWaP), performance, availability, interfaces, and industrial/automotive qualification. The LiDAR interfaced directly with the flight computer over ethernet for communication and power was provided through the reconfigurable 12V payload power load switch on the R5 custom input/output (IO) interface board. The other major change was the removal of solar panels, since the short mission duration did not require power generation.

Spacecraft 7 Communication

The communication subsystem was identical to the R5-S5 baseline, consisting of an Iridium® 9523-based system and the use of the Iridium Short Burst Data (SBD®) protocol. Like R5-S5, R5-S7 used an NAL A3LA-RS radio paired with their SAF2040B antenna.

Spacecraft 7 Flight Software

R5-S7's FSW was based on R5-S5's, which used NASA's core Flight System (cFS) FSW framework. Beyond the uniqueness of a specific cFS instantiation for a particular set of hardware and ConOps, R5-S7 also used a custom library, the cFS Glue and Unit-Test Library (CGUL). CGUL is intended to simplify development within cFS by abstracting away certain aspects.[8] R5-S7 utilized 16 unique applications beyond the cFS standard compliment, including sensor hardware IO (CAMXIC, CAMXIQ, GPS, LIDAR, IMU), sensor data processing (VNML, MSKY, ST), GNC (NAV, FASTNAV, GUID, CNTRL), communication (TELEM, COMM), and system configuration (AFM, PWR). Relative to other cFS instantiations, R5-S7 also pushed a substantial amount of functionality into standalone libraries to minimize external dependencies.

The Automated Flight Manager (AFM) application, while present on other spacecraft, was reconfigured to enable the unique ConOps of the R5-S7 mission. This critically allowed the spacecraft to complete its mission automatically without any input from the ground.

The telemetry (TELEM) and communication (COMM) applications managed data generated by the spacecraft systems and allowed data to either have its latest state downlinked and the rest discarded or to have all of it saved to disk and deleted only after downlink. The operation of the 9523-based radio was also updated based on the flight experience of R5-S2 and R5-S4 to reduce the rate at which messages were attempted to no faster than every 30 seconds to reduce the chances of trigger traffic management.

In order to onramp software payloads onto the primary flight computer and minimize their verification burden, Docker containers were used to contain and run them. This also ensured there were no dependency conflicts between these software payloads and the primary FSW. These containers communicated over the network and had fixed CPU and RAM allocations, preventing payloads from consuming all of the system resources the primary FSW required to function.

Spacecraft 7 Guidance, Navigation, and Control

The GNC system is instantiated in the primary FSW as four applications (not including sensor IO or preprocessing): the navigation filter (NAV), the filter propagator (FASTNAV), guidance (GUID), and control (CNTRL). GNC also retains ownership of the configuration and high-level function of the AFM as it directly drives how the ConOps (include trajectories) are defined. The AFM was implemented with 13 unique vehicle states that drove the function of the FSW and the associated hardware. Each of these states was able to transition to its successor with a variety of triggers like a timer or application commands.

The R5-S7 navigation system was designed to estimate the position and velocity of the deploying vehicle relative to the R5-S7 body frame. The primary sensor was the MicroVision Movia LiDAR, which provided point cloud data to the navigation application at 5 Hz in all illumination conditions. To simplify processing, the point cloud was condensed into a single relative position

measurement that updated a kinematic relative navigation (RelNav) Extended Kalman Filter (EKF). The RelNav filter propagation utilized gyroscope measurements from the onboard Epson G370 PDT IMU. The navigation system also received bearing measurements to the deploying vehicle via an onboard neural network that processed images from a camera co-boresighted with the LiDAR. For simplicity, these bearing measurements were not incorporated into the RelNav filter but were intended for telemetry downlink. The navigation system also featured an inertial attitude estimation capability instantiated as a multiplicative Extended Kalman Filter (MEKF) based on the Orion Attitude (ATT) MEKF design, but verification of this system was descoped during development as it wasn't required for the primary mission.

To ensure the upper stage was acquired by the LiDAR, a re-acquisition algorithm was implemented. This algorithm ingested the angular rate at GNC initialization and back-propagated the vehicle dynamics using the elapsed time since deployment to determine the delta attitude since deployment. Assuming a +V-bar separation trajectory, the estimated delta attitude and initial LVLH attitude were used to compute the pointing error to the upper stage's LVLH position. Time since deployment was estimated by summing the Linux kernel BOOTTIME and the BIOS boot time, the latter of which was characterized on the ground to be approximately 19 seconds. The total elapsed time from deployment to GNC initialization was approximately 45 seconds.

The guidance software generated attitude, velocity, and position error commands for the control system to support client pointing, self-arrest, and waypoint-seeking modes, respectively. Self-arrest was accomplished by commanding a zero range rate while maintaining client pointing. For waypoint seeking, guidance computed position errors from parameterized offsets relative to the vehicle's current position and provided these errors to the control system for trajectory tracking.

The control system employed a piecewise, bang-bang-inspired attitude control algorithm known as a phase-plane controller. The controller consists of two operating regions: a maneuver region, which commanded vehicle slews toward the desired attitude, and a hold region, which maintained pointing within configurable deadbands. Waypoint seeking was achieved using proportional-derivative (PD) control based on the position and velocity errors generated by guidance.

R5-S7 had two image processing algorithms onboard. Like the R5 spacecraft before it, R5-S7 featured the latest development version of the NASA COTS-based Star Tracker System.[9] As mentioned above, the star tracker system was not verified or used for this flight. R5-S7 also had a convolutional neural network that was used for onboard identification and detection of spacecraft in imagery from the relative navigation camera (a color Ximea XIQ013). This algorithm came from the NASA Pose Bowl challenge.[10]

Spacecraft 7 Propulsion

R5-S7 utilized a cold-gas nitrogen propulsion system for full six degree of freedom (6DOF) control. The subsystem was based on the R5-S2 and R5-S4 propulsion subsystem which was successfully demonstrated on orbit.[1] Several iterative improvements were made based on the AIT of R5-S5, including a more compact high pressure manifold design which added a quick disconnect fitting for the fill port through the top plate, meaning no high pressure tube routing through the vehicle was required. A new Para Plate regulator was used whose interface was originally designed for Mini-AERCam and a new larger Lee Co relief valve was used, as opposed to R5-S2 and R5-S4 which had two smaller relief valves to meet the same flow requirements. The thruster pod low pressure assemblies were made smaller to have a better mounting scheme for tighter tolerancing on the jet locations and to not have the thruster pods be in the critical path of assembly for structures.

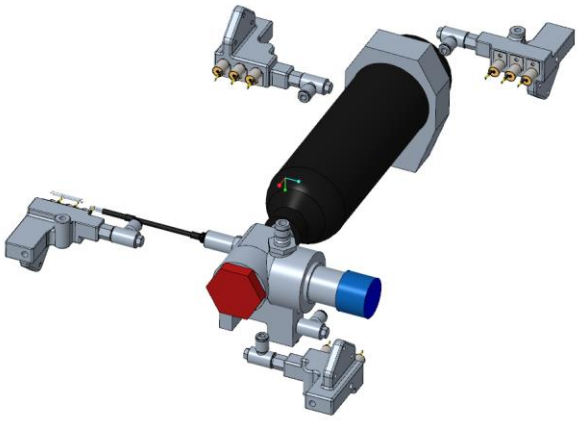


Figure 2: R5 Propulsion subsystem high pressure and low-pressure components shown without tubing

Spacecraft 7 Structures

The physical architecture consisted of the Z+ baseplate, the Z– top plate, and the corner pieces that structurally tied the frame together. Closeout panels attached around the perimeter to enclose the interior. The baseplate and top plate both included a regularly spaced grid of M3 female threads, similar to an optical breadboard, which

enabled rapid internal reconfiguration without redesigning or remanufacturing structural components. Custom 3D-printed Ultem brackets were used to secure internal components and offered rapid turnaround during iterative design cycles.

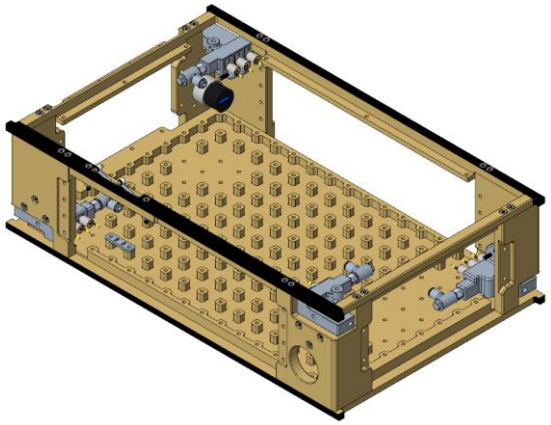


Figure 3: R5-S7 primary structure (no access panels or internal components except thruster pods and GSE)

The R5-S7 design centered on improving the assembly process. Instead of optimizing purely for strength-to-weight or cost, significant emphasis was placed on practical assembly timeline, sequencing, and access. This reduced rework, enabled nonlinear assembly, and forced explicit consideration of connector reachability and tool access for every component. The optical breadboard-like layout also supported a small column height, allowing most harnessing to be passed underneath the components. This meant harness routing could occur later in integration instead of waiting on a full routing model before structure fabrication. The result was reduced harness cost and lead time, along with greater late-stage flexibility. Additionally, the non-structural nature of the closeout panels provided increased flexibility of the material and method by which they could be manufactured.

R5-S7 did not require any external passthroughs beyond those needed for sensor access. Since the vehicle did not recharge on orbit, there were no solar-panel penetrations or charging harness interfaces.

R5-S7 also served as the structural prototyping baseline for subsequent vehicles. Early planning documents reference the start of R5-S7 structural prototype printing and its use in structural fit checks, establishing the architecture that later informed the builds for R5-S10 and future spacecraft.

Fastener selection was intentionally simplified from prior vehicles. M3 × 8 mm socket-head cap screws and M3 × 6 mm flat-head screws were selected as the primary fasteners across the spacecraft and were used wherever possible. Preload torque was used with DOWSIL 3145 as the secondary locking feature. This allowed fastener installation to target a single torque value followed by RTV application, replacing the more time- and work-intensive locking-helicoil approach used previously. The RTV staking also provided a straightforward visual confirmation of fastener lock.

For tolerance control, tight tolerances and precision stackups were avoided wherever possible. The tolerance approach assumed a gaussian or uniform tolerance distribution with a zero mean and used integrated testing (fit checks) to verify the assemblies were still within specification. Assemblies out of tolerance would have been refabricated. This approach has produced assemblies that remain within integrator requirements without requiring any post-machining to date, simplifying fabrication and reducing schedule pressure.

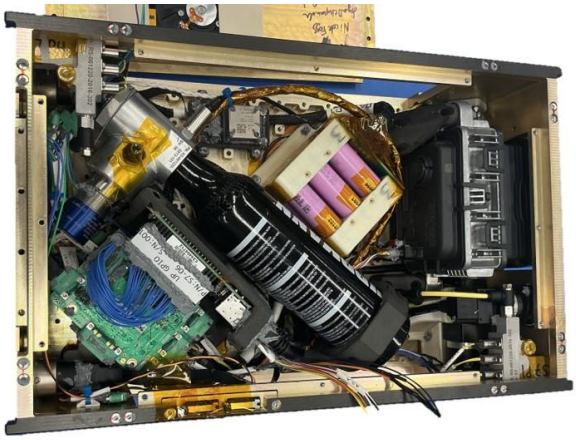


Figure 4: R5-S7 interior prior to lid closure

Spacecraft 7 Mission Results

Given the limitations of the telemetry system and the short mission lifetime, very little spacecraft telemetry was gathered during the R5-S7 mission. The mission artifacts consist of the SpaceX launch livestream (which contains a single frame of the R5-S7 deploy), four telemetry messages, and another video from the upper stage showing R5-S7 loitering nearby. Even with this limited data, the R5-S7 vehicle was proven to have completed its maneuvers as intended.

The R5-S7 telemetry includes the 3D coordinates of a LiDAR point cloud centroid which, when correlated with time, verifies the intended 0.3 m/s deployment of the CubeSat. This LiDAR return also verifies the LiDAR functioned as intended and the spacecraft was able to acquire the upper stage without issue. The

telemetry also includes 42 attitude error measurements recorded at 1/3 Hz (and each only having been written out to telemetry if a new centroid from the LiDAR was processed). Combined with the low attitude rate of the spacecraft, this telemetry shows that approximately 630 successful LiDAR centroids were processed, proving the sensor continued to function past the initial return. The attitude error data itself shows a rapid trend towards 0 and then a fairly steady plateau near 0 for the pitch and yaw axes (roll was not telemetered). Based on simulated results of spacecraft performance, this also strongly suggests the IMU was functioning as expected. Given that the only effectors onboard were the cold-gas jets, the reduction in attitude error suggests that all jets were functioning as expected, even though no explicit firing history or pressure data was downlinked.

Blender was used to reconstruct the physical scene captured by the upper stage camera with R5-S7 loitering near the upper stage. The spacecraft's expected ConOps (including arrest) was modeled and propagated and the resulting trajectory and timeline rendered and compared to the video and both agreed closely, suggesting R5-S7 completed its arrest as expected.

One of the four telemetry messages was a partial image file. Unfortunately, there wasn't enough data to reconstruct the image. However, with the receipt of this message and with prior test results of the system (and from inspection of the FSW), this demonstrates that the relative navigation camera was operating, captured an image, the VNML algorithm was running, identified something in the image as a spacecraft, and the "auto-downlink" capability triggered as intended. If a handful of additional telemetry messages were received, it's highly likely that the full image could be reconstructed. Additional telemetry would have also likely included explicit bearing measurements from the algorithm.

The first SBD® message took 25 minutes to successfully send, with 53 attempts and a reported signal strength of 0 bars. The second SBD® message took 8 minutes to successfully send after the first with 14 attempts and a reported signal strength of 0 bars. The third SBD® message took 30 seconds to successfully send after the second and included no other metadata as it was the image message. The fourth SBD® message took 1 minute to successfully send after the third with 2 attempts and a signal strength of 0 bars. All messages had a size of 265 bytes.

Spacecraft 7 Lessons Learned

Avionics lessons learned include the importance of documenting the full set of system requirements and not just the deltas from the last baseline. An explicit "arm+fire" system was previously specified and implemented in R5 avionics. As the associated avionics components underwent revisions, this requirement was lost and the capability was erroneously removed.

For the communication subsystem, the first telemetry message from R5-S7 was received approximately 25 minutes after deployment, placing it after the majority of the ProxOps. While not necessarily a new lesson learned, Iridium® network connectivity can vary substantially in low earth orbit (LEO) or similar suborbital trajectories. Based on the experience with R5-S2 and R5-S4, these results were known to be possible and the vehicle architecture was developed with this potential outcome in mind. Very few telemetry messages were received, but again this was known before the mission to be possible (and perhaps even likely) as the mission had a unique (and very short) trajectory.

FSW lessons learned include the great utility of hardware-in-the-loop testing to catch bugs as early as possible. From limited FlatSat testing to day-in-the-life (DITL) testing, many system capabilities were able to be verified quickly (or iteratively reworked quickly until finally verified). For the telemetry system, there were multiple places where the telemetry was defined between the FSW and ground software and these would often get out of sync with each other, which took time to manually realign them in order to successfully process telemetry. This was very labor intensive and, while they're out of sync, prevents end-to-end testing. A single source of truth for the telemetry definition would solve this. Another FSW lesson learned was that, depending on the exact issue discovered, debugging memory leaks may be faster than creating a new static implementation. Finally, an assumption on regular and linear telemetry caused an implicit assumption in the telemetry system's instantiation to go unnoticed and greatly complicate temporal reconstruction of the flight telemetry.

GNC lessons learned include the challenges of hardware-in-the-loop testing in a scaled environment, where scaling also drives noise and filter impacts. A bug in the reacquisition logic was identified post-delivery, which could have generated an inaccurate attitude error if the deployer was outside of the LiDAR's FOV upon GNC initialization. Luckily, the vehicle was inside the LiDAR's FOV thanks to low tipoff rate at deploy.

Propulsion lessons learned include the challenges of valve reliability. R5-S7's high pressure isolation valve was built to the same design and specification as prior vehicles, but came from a new lot. Through acceptance testing and integrated subsystem checkouts it was discovered that the isolation valve would close when pressurized in the reverse flow direction, like when

filling the tank. The vehicle was able to be filled through slow mass flow into the tank with a low delta pressure between the pressure source and tank. There was no evidence in integrated testing and in flight that this defect impacted flight performance, but this did highlight the benefit of getting to integrated testing as quickly as possible and showed the difficulty of working with bespoke/low-volume-manufactured hardware.

Structures lessons learned include the benefit of optical breadboard-style architecture. The threaded-grid Z+ baseplate, matching Z– plate pattern, and 3D-printed Ultem brackets allowed rapid internal reconfiguration with minimal redesign. Assembly-first structural design also dramatically reduced rework. Designing around connector/tool access, harness routing beneath the column height, and late-stage harness installation provided substantial schedule flexibility. Taking a nominal-first, machine-only-if-needed tolerance strategy was effective. The approach of machining only parts that came in out of tolerance resulted in fully compliant assemblies with no post-build machining required. All of these combined to enable a complete spacecraft assembly in less than 7 days. Another lessons learned include the sensitivity of the inhibit switches in this location on the rail as the switches were prone to partial engagement during dispenser insertion. RTV staking must be monitored for flushness as RTV protrusion was discovered on rail fasteners just prior to dispenser integration, creating a deployment risk for the low-force springs on that mission. This was addressed with rework in the integration facility, which is always more challenging than in the primary AIT facilities. HDDs are very sensitive to anything that can deform their case and may stop spinning with only a slight deformation.

SPACECRAFT 10

R5-S10 is primarily an orbital ProxOps demo mission and as such directly builds on R5-S7. Some primary differences include the orbital deployment of R5-S10 and its secondary mission afterwards to demonstrate additional payloads. Other differences include the increased complexity of the ProxOps and the associated increase in capability of R5-S10. This mission was done in partnership with Momentus, Inc. codified in a nonreimbursable Space Act Agreement and uses the Vigoride 7 orbit service vehicle (OSV) as the client.

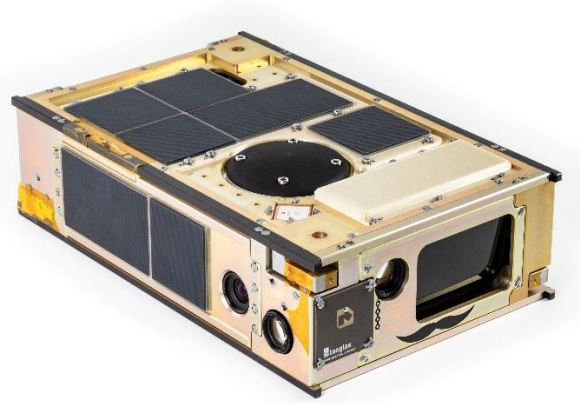


Figure 5: R5-S10 front 3/4 view

Spacecraft 10 ConOps

The R5-S10 ConOps were defined as an incremental improvement from R5-S7 and include arresting a greater departure velocity, the option to perform precise translational maneuvers, the option to perform approaches towards the client, and the demonstration of autonomous departure. Like R5-S7, the ConOps were developed with the mission partner in order to minimize cost, schedule, and risk. This resulted in the following ConOps:

1. Vigoride deploys from the launch vehicle
2. Vigoride performs various payload demonstrations
3. Vigoride orients itself to deploy R5-S10 -Vbar
4. D+0: Vigoride deploys R5-S10
5. D+31s: R5-S10 completes power on
6. D+46s: R5-S10 GNC system completes initialization and starts rotating to acquire Vigoride
7. D+50s: R5-S10 acquires and starts tracking Vigoride, keeping it near the center of its LiDAR's field of view, and starts sending telemetry on TT&C and crosslinks
8. D+51s: R5-S10 begins arresting departure velocity
9. D+80s (TBD): R5-S10 completes the departure arrest and holds position while tracking Vigoride
10. D+500s (TBD): (OPTIONAL) R5-S10 approaches Vigoride
11. D<5000s: R5-S10 departs from the vicinity of Vigoride and starts secondary mission

Spacecraft 10 System

R5-S10 was based on the lessons learned from the prior R5 vehicles and retained the 6UXL form factor and

avionics, communication, FSW, GNC, propulsion, and structures subsystems.

Major differences between R5-S10 and R5-S7 included the addition of a crosslink capability (enabling more data to be downlinked via the client versus the R5-S10 Iridium® system) and enhanced ProxOps capability and verification. The enhanced ProxOps included improved maneuvering capability, deeper verification of collision avoidance, and a "self-deploy" capability.

Another difference between R5-S10 and R5-S7 is the addition of an event camera payload for NASA's Safe and Precise Landing Integrated Capabilities Evolution (SPLICE) project. SPLICE works towards advancing hardware and algorithms to support autonomous precision landing and hazard avoidance on non-terrestrial bodies. SPLICE has been investigating event-based, or neuromorphic, cameras for various applications to support these activities.

Event-based cameras, or simply event cameras, are similar to conventional digital cameras where a photosensitive pixel array behind a lens records information about the light collected from a scene. However, while a conventional camera exposes all pixels for a set amount of time to generate a "frame" that is read out as an image, an event camera asynchronously records data from all pixels continuously and monitors for change in light levels. When a significant change in light intensity is registered at the detector this "event" is logged with a timestamp, pixel location, and polarity to indicate if the change was an increase or decrease in brightness. This seemingly subtle difference in light detection with very similar hardware enables the sensor to achieve very high dynamic range (>120 dB) and temporal resolution (10,000 Hz) in comparison to a conventional camera.[11]

While surveying the use of event cameras for GNC applications, SPLICE identified that while many groups were doing outstanding work on algorithm development for space applications, there remained a dearth of real data collected from modern event camera hardware in the space environment. To help address this gap, SPLICE provided an iDS uEye XCP-E event camera equipped with a Schneider JADE F2.8/25mm 1.2" C-mount lens as a payload for the R5 S10 spacecraft.

The primary application of interest to SPLICE is the use of an event camera as a star tracker. The higher temporal resolution and dynamic range of event cameras is expected to enable this sensor to provide attitude information under high body rates when conventional star trackers would fail due to "smearing" of stars and when there are bright objects such as the Earth or Moon in the field of view of the sensor that would typically result in saturation of the detector, making identification of relatively faint stars extremely challenging with conventional hardware.

The primary objective of the event camera payload is to collect data of star fields with different sensor bias settings. The bias settings of an event camera determine the threshold change of light required to register an event. While optimal bias parameters can be estimated on the ground, in-flight data is valuable as it removes atmospheric scintillation or "twinkle" effects that result in variation of a star's brightness. The bias settings that the event camera payload will be cycled through during the mission are detailed in the table below.

Table 1: Summary of event camera bias parameters to be sampled during S10 operations.

Index	Bias On	Bias Off	Bias Fo	Bias HPF	Bias Refr
1	0	0	0	0	0
2	0	0	55	0	0
3	0	20	-35	0	0
4	0	20	0	0	0
5	0	20	55	0	0
6	0	0	-35	0	0
7	-10	0	0	0	0
8	-10	0	55	0	0
9	-10	20	-35	0	0
10	-10	20	0	0	0
11	-10	20	55	0	0
12	-10	0	-35	0	0
13	-20	0	0	0	0
14	-20	0	55	0	0
15	-20	20	-35	0	0
16	-20	20	0	0	0
17	-20	20	55	0	0
	-20	0	-35	0	0

Event cameras can quickly produce very large data products due to their high temporal sampling. To optimize data value given the bandwidth limitations that are intrinsic to spaceflight, the event camera was configured to record for 5 seconds every 5 minutes with a different bias configuration setting. These files will be transferred from R5-S10 to the OSV over a Wi-Fi cross-link before being downlinked to the ground for analysis.

In addition to the data collect aspect of this mission, the R5-S10 FSW was also integrated with the ability to process event camera data to determine the attitude using imagery of star fields. This was done using the Kitware open source EBS-EKF software which has been previously shown to provide comparable accurate

attitudes to conventional star trackers and also continue functioning at body rates.[12] SPLICE reproduced these results using a panning telescope platform rotating at 4°/s in Houston, TX with no attempt to mitigate light pollution, a sample of which is shown in the below figure.

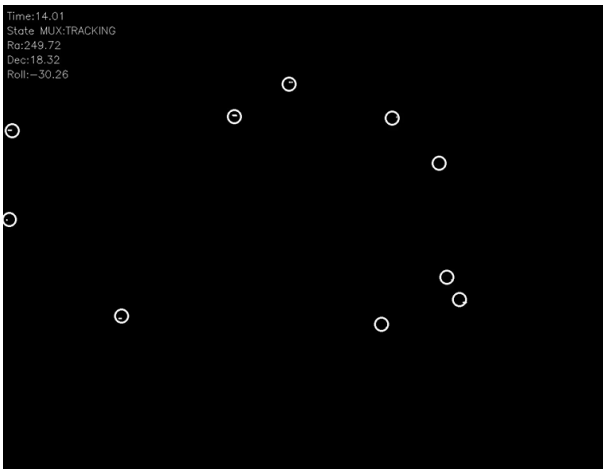


Figure 6: Example of the Kitware EKF-EBS star tracking software performing attitude determination using ground-collected live sky event camera data.

The event camera data collected on orbit from R5-S10 will be processed onboard using the EBS-EKF software after the ProxOps phase of flight is complete. Minimal telemetry to characterize the performance of the software will be downlinked for post-flight analysis.

Spacecraft 10 Avionics

The avionics subsystem was developed from the R5-S7 baseline. With the return to an orbital mission and their associated longer mission lifetime, solar panels were reintroduced, deferring to structures on maximizing use of available surface area. The custom PDU was also updated to a new revision (v4.2) which added a hardware watchdog timer, a standalone voltage rail for the flight computer, and an improved design for manufacturability.

In addition to the LiDAR, two more payloads were introduced – the event camera and WiFi. Both components were owned by other subsystems (GNC and Comm, respectively) and used 5V for power and USB for interfaces. With the desire to save power, both components were load switched. The two additional internal USB 2.0 connections available on a 1mm pitch header on the end of the UP-CHT01 flight computer were broken out by cable to female USB-A connections to provide enough USB ports for all peripherals without introducing additional hardware. With the four primary USB 2.0 connections used by the R5 baseline and the single USB 3.0 connection used for ground support equipment (GSE), all of the USB connections supported by the UP-CHT01 were utilized.

Spacecraft 10 Communication

The Iridium®-based part of the R5-S10 communication subsystem was identical in hardware and configuration to the R5-S7 baseline. R5-S10 also included a crosslink system to support direct communication to the Vigoride OSV that consisted of a Netis WF2190 WiFi dongle and two Taoglas WDMP.2458.A antennas (all COTS).

Spacecraft 10 Flight Software

R5-S10's FSW was developed from the R5-S7 baseline. Notable changes include the addition of a new IO application and docker container to support the event camera payload and its associated image processing software, an application to support data transfer across the WiFi crosslink, and substantial updates to the AFM. The AFM updates enable much more complex behavior in the system's configuration, enabling capabilities like the more advanced ProxOps ConOps and the ability to constantly monitor for various conditions (like propellant and power levels) to trigger additional state transitions.

Spacecraft 10 Guidance, Navigation, and Control

R5-S10's GNC system was developed from the R5-S7 baseline. The sensors were identical to R5-S7, as were the image processing algorithms. Given the similarities, the relative navigation filter was kept the same but had updated tuning parameters based on simulated results of the new ConOps. The target re-acquisition algorithm had some bug fixes implemented and the inertial pointing capability underwent verification.

R5-S10's guidance system was nearly identical to R5-S7's with some tweaks to improve self-arrest from the expected faster 1.0m/s deploy velocity (vs 0.3 m/s on R5-S7). Guidance prioritized nulling only the +X body velocity components while the RelNav filter converged to prevent extraneous Y/Z axis firings. The prioritization was enabled below a parameterized number of LiDAR measurements

The R5-S10 Control system was identical to R5-S7, but with increased maneuver rate to 20 deg/s. Analysis showed faster re-acquisition, hence faster maneuver rate, was needed to ensure self-arrest before the 80 meter range limit imposed by the LiDAR.

The R5-S10 AFM was configured to command GNC apps to arrest velocity upon power on. Upon completion of self-arrest, the CNTRL application triggers AFM into track-and-hold mode, meaning the GNC apps are commanded to hold position relative to the OSV while keeping it centered in the LiDAR FOV. Transition out of track-and-hold requires a command from the ground. If no command is issued, the system will depart upon low power or low prop conditions. If commanded by the ground, the AFM transitions to the first approach mode and the spacecraft then advances to a waypoint 30m from the OSV.

R5-S10 also introduced a "self-deploy" capability over the R5-S7 baseline. This is a special capability where if the spacecraft detects that it's still in the deployer (LiDAR returns < 1 m), GNC initiates open-loop -X thruster firings until separation occurs. The GNC system would then increase separation to more than 10 meters range and hold, awaiting ground command (or until power or propellant departure thresholds were met).

Spacecraft 10 Propulsion

The R5-S10 propulsion subsystem hardware was the same as the baseline of R5-S7. The changes from R5-S7 included increasing the setpoint of the regulator to increase the thrust level of the jets. This was to accommodate the increased propellant usage for self-arrest with the higher deployment speed. Monte Carlo (MC) simulation results from GNC on self-arrest were exported and run on the propulsion system as scripts utilizing the complete propulsion system to see if the available propellant could execute canceling the deployment velocity. The thrust levels were tuned in the MC runs and the actual system would be changed to match the models, then nominal and worst case runs were completed with a full propellant load. A nominal thrust level of 130 mN was chosen as a balance between stopping distance of the vehicle and the propellant remaining to be able to self-deploy.

The propellant amount was increased compared to R5-S7 to help accommodate the initial self-arrest maneuver and was acceptable per the thermal environments provided by the integrator and host-spacecraft.

Spacecraft 10 Structures

The R5-S10 structure was baselined from R5-S7, with the major changes driven by the addition of the WiFi antennas, the event camera payload, and the reincorporation of solar panels. The deployer-facing side of the spacecraft was already heavily populated with sensors, so integrating a forward-facing WiFi antenna required modifying the corner-frame geometry to recover usable surface area on that face. A similar modification was required on the aft side for the second antenna, which needed a double-stack Ultem panel to interface cleanly with the +X reaction control wheel.

Reintroducing solar panels required developing a mounting approach compatible with the available real estate and with the existing structural margins. Based on prior vibration and thermal-vacuum testing of DOWSIL 3145 bond strength, it was selected as the sole attachment method. Solar panels for R5-S10 were bonded to either an Ultem carrier plate that was mechanically fastened to the spacecraft face or, in an alternative configuration, bonded directly to the metal structure. Both approaches were implemented on R5-S10. New pass-throughs in the Z+ and Z- plates were introduced to route the panel harnessing, consistent with earlier structural changes noted in fabrication discussions

Finally, the inhibit switches were relocated from the deployer rail interface to the rail tips to allow the spacecraft to begin its boot sequence as soon as the deployer door opened.

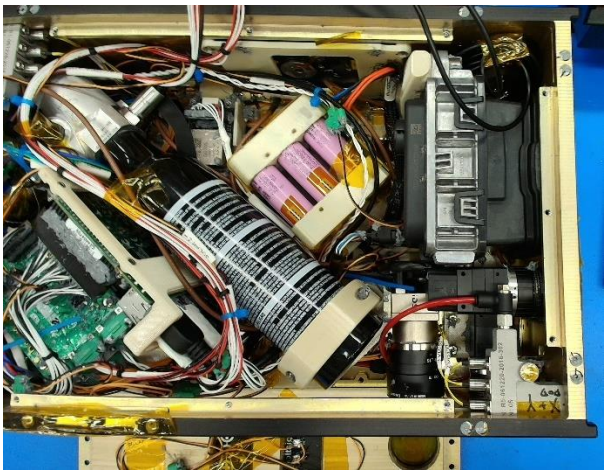


Figure 7: R5-S10 interior prior to lid closure

Spacecraft 10 AIT Results

In the absence of flight results (R5-S10 is currently awaiting deployment from the OSV), results of AIT are included. R5-S10 was subject to random vibration testing and a litany of limited-scope functional tests and far more integrated/involved day-in-the-life tests to verify the system's function and performance. The day-in-the-life tests were scoped and executed to align with the expected on-orbit operations as much as possible and will serve as an expected performance baseline during the mission. This section will focus more on the unique elements of the AIT phase.

R5's high-fidelity, physics-based 6DOF Trick simulation had its MC capability configured to analyze nominal self-arrest through depart. The figures below show a subset of the runs that successfully arrest the deployment velocity, hold position, then depart, while centering the client OSV in the LiDAR FOV. Trajectories from nominal runs post-depart were also assessed over multiple orbits and showed negligible probability of recontact with the OSV.

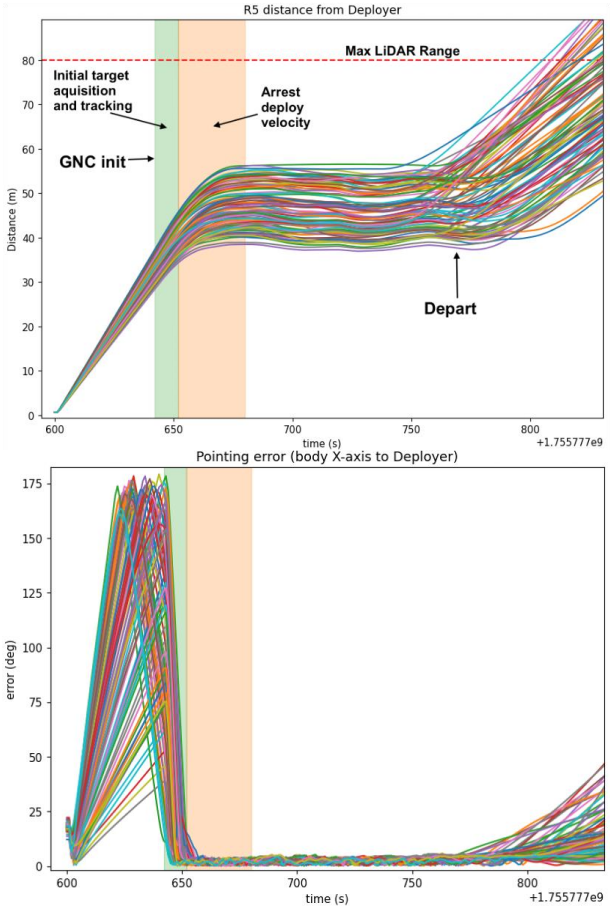


Figure 8: MC results for R5-S10

To further assess risk of recontact with the client OSV, failed thruster scenarios were also analyzed. Twelve MC scenarios were analyzed, one for each failed thruster pair. Results indicated a 1/1200 probability of recontact given any failed thruster pair. The probability of a failed-on thruster pair (most likely a FSW crash while firing thrusters) was not determined quantitatively but qualitatively is low (this qualitative assessment also does not include the recovery capabilities of the R5 bus to such failures, like the watchdog). The overall possibility of a collision is the product of these probabilities: 1/1200 * likelihood of a failed-on thruster. This risk was low enough for stakeholders to accept and proceed with the mission.

Testing on the flight vehicle was performed to verify the polarity of each GNC sensor. LiDAR polarity was verified via pointing at a known structure and correlating independently inspected point clouds with centroids generated by FSW from the same data. IMU polarity was verified through tilt and rotation tests that recorded the IMUs outputs from FSW in the body frame. Star Tracker polarity was verified with live-sky testing: recorded data from FSW was compared to ICRF to body quaternions generated independently using astropy. Thruster mapping and control polarity was verified through a sequence of thruster firings: thruster 1-12 were fired in increasing order while air pulses from the expected thruster in the expected order were identified.

Simulation fidelity work was performed to match the simulated LiDAR and the MicroVision Movia's beam pattern. At ~60 meters range small differences in beam pattern could result in a different number of LiDAR returns. At this range it becomes critical to determine real returns vs noise as there will be fewer than 10 real returns due to beam angular separation at 60m. A persistence-based approach was used to filter the real returns from the noise. This approach worked well when testing with a distant/small object on the ground as shown by the figure below. Between 1-5 points of noise can be seen in the "Total Returns" plot on the left and are filtered out on the right:

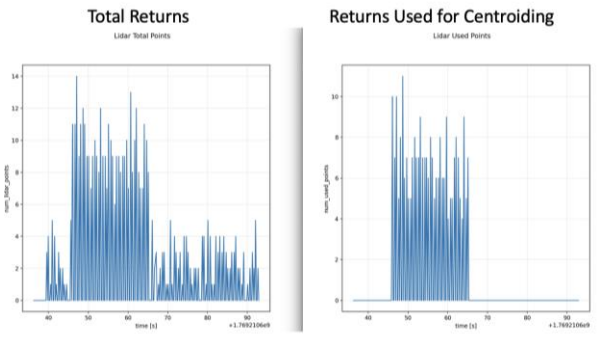


Figure 9: R5-S10 LiDAR point processing logic result

Spacecraft 10 Lessons Learned (so far)

The avionics subsystem encountered many lessons learned. In order to save power during the non-ProxOps mission phases, the avionics subsystem moved the LiDAR to a switched power channel. Late in AIT, it was discovered that this moved the LiDAR initialization time from being in parallel with the flight computer to being serial, which was problematic for GNC. The switch's default state was able to be set to "on" and then managed appropriately in FSW, enabling GNC to function and

realizing the power savings in subsequent mission phases. During AIT, repeated i2c bus failures were traced back to a GSE harness that wasn't backpotted and had a loose conductor shorting the lines. After a random vibration test, some USB devices stopped responding. This was traced to a poor USBA connection, likely due to the flexibility and flowability of the DOWSIL 3145 material used for staking. This was addressed by adding a polycarbonate 3D-printed fixture that ensured the USBA connectors could not move after mating.

For the communication subsystem, the WiFi supported connectivity at ranges greater than expected. Though the results were gathered terrestrially in an environment where reflections or multipath may have enhanced the range, the system was shown to reliably connect at 160 m from the access point when the patch antenna was pointed at the access point and to reliably connect at 140 m from the access point when the patch antenna was pointed 90 degrees off of the access point.

The FSW subsystem was able to bound payload software system utilization with containers, but only within the defined bounds. The disk space usage wasn't limited and a payload ended up consuming far more disk than expected. This was caught in hardware-in-the-loop testing and quickly fixed. If this wasn't caught before flight, this likely would have consumed all space on the disk and prevented the flight computer from rebooting, resulting in loss of mission. It was also observed that when the event camera detected large numbers of events other USB devices would start throwing errors, leading to the hypothesis that large numbers of events may be overwhelming the USB2 controller.

As mentioned above in the avionics section, GNC discovered that the LiDAR's new power configuration and the spacecraft's new deployment velocity were coupling to potentially place the spacecraft beyond the LiDAR's maximum range by the time GNC had initialized. A variety of options were explored and the change of the power switch's default state was the simplest and easiest solution. In order to further reduce the range at which GNC initialized, the deployer was also modified to reduce the deployment velocity. As R5-S10 has a much smaller client at longer ranges than R5-S7, the noise of the LiDAR sensor also became a challenge. This was addressed by updating the onboard software to reject points that matched the pattern of the noise.

Propulsion lessons learned include verifying the construction of all subcomponents. A small metal component on the low pressure side of the system was assumed to be a single component but after vibe separated, revealing it was actually an assembly of two components. While the components remained captive and did not impact the functionality of the system, they were reassembled, staked, and the other components of the system were inspected to ensure there were no additional unanticipated assemblies.

Structures lessons learned include the bond strength of the solar panels to both aluminum and to Ultem with the DOWIL 3145. While this was suggested based on prior test data, the applications held up through all handling and environmental testing. Specifically using an additional panel as an intermediary mount for the solar panels enhanced access to the spacecraft interior as it improved fastener access. Inhibit switch relocation to rail ends was a structural improvement while adding some dispenser installation complexity. The trade was worth it though in that it added a potential capability to our future spacecraft of allowing for immediate activation and potential self-deployment.

SPACECRAFT 9

R5-S9 was built using R5-S10 as a baseline, but is not performing ProxOps. The R5-S9 primary mission is to demonstrate a payload from Sandia National Laboratories (Sandia) that will perform three missions: 1) space observation, 2) maritime observation, and 3) radiation upset measurement. These missions will be executing autonomously using hardware and software resources onboard, with only high-level mission metrics downlinked due to the limited communications capacity. The spacecraft also has secondary mission objectives including the demonstration of a novel LED-based optical communication system developed in partnership with the Aerospace Corporation and the next-generation Iridium® 9704-based radio.[13] The partnership with Sandia is codified in a nonreimbursable Space Act Agreement. Beyond the technical functionality of the spacecraft, this mission was used to demonstrate the deeper collaborations possible through R5 and involved several engineers from Sandia embedding with the R5 team during AIT to get hands-on experience with the R5 process and to provide process feedback directly.



Figure 10: R5-S9 3/4 front view

Spacecraft 9 ConOps

The Sandia payload requirements, bus capabilities, and the R5-S10 AFM architecture drove the definition of the ConOps. R5-S9 will receive commands dictating if the spacecraft should perform primary or secondary missions and the spacecraft will then execute the mission as able. The spacecraft autonomously determines if the missions are feasible and which subset of missions to attempt. These missions have similar ConOps, which take the form of:

1. B+0s: R5-S9 completes power on and starts attempting to send telemetry
2. B+2s: R5-S9 receives command on what mission to execute
3. B+90s: R5-S9 completes state initialization and provides state to the payload
4. B+95s: The payload determines the best course of action and provides pointing target to the bus
5. B+200s: R5-S9 completes slew to pointing target and operates sensors/etc. as required
6. (until power off): R5-S9 continues payload operations and telemetering data

Spacecraft 9 System

R5-S9 was built from an R5-S10 baseline with the LiDAR and WiFi removed and the addition of the Sandia payload hardware, a second Ximea MC023 camera used as a star tracker, an optical communication payload, and an Iridium® 9704-based radio.

The payload hardware included a Basler machine vision imager and ruggedized lens, a custom PCBA with a memory chip for measuring radiological upsets (bit flips), and a modified Nvidia Jetson Orin Nano Super Developer Kit, which was used as the primary payload computer for processing sensor data and communicating with the primary flight computer. Similar machine vision cameras were successfully operated in-flight onboard previous R5 spacecraft, and therefore the operability of the imaging hardware on this payload was assessed to be low risk. The greatest risk and uncertainty associated with this hardware was the Jetson compute speed and operability at the elevated temperatures the payload might experience during parts of its orbit. To estimate payload temperature, a steady-state thermal analysis was performed in an Ansys simulation, based on a CAD model of the mounting panel for the Jetson (including material properties), oriented directly toward the sun, and operating with the maximum specified electrical power draw of the Jetson (25W).

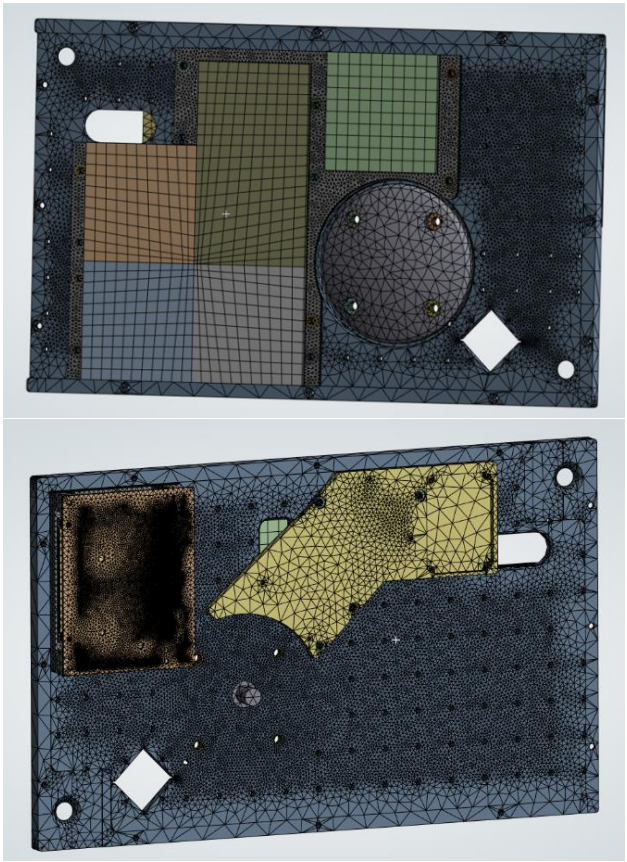


Figure 11: Panel meshes used for thermal analysis. The Jetson is mounted on the copper heatsink in the top-left of the right image.

Assuming a conservative 100% duty cycle, this resulted in an estimated maximum temperature of 111° C. These results suggested that the payload computer would likely thermally throttle, reducing processing speed, and in the worst case, shut down after extended operation in direct sunlight.

To further validate these results, CPU and GPU stress test workloads were executed on the Jetson while running in a thermal-vacuum chamber. The chamber was

set to maintain a constant temperature based on the average temperature R5-S2 and R5-S4 experienced. The mechanical interface and heatsink that would be used to mount the Jetson in R5-S9 was attached for this test, but the Jetson was not attached to the R5-S9 body, and therefore this test did not include the R5-S9 thermal mass. In this test, the Jetson began throttling within 6-8 minutes of execution time, depending on workload and power configuration. To maximize the probability of success, the Jetson was reconfigured to the 7W power mode as default for flight, and mission parameters were modified to reduce the high power processing duration. Also, solar panels directly above the Jetson mounting position were removed and replaced with silver Teflon tape to aid thermal reflectivity.

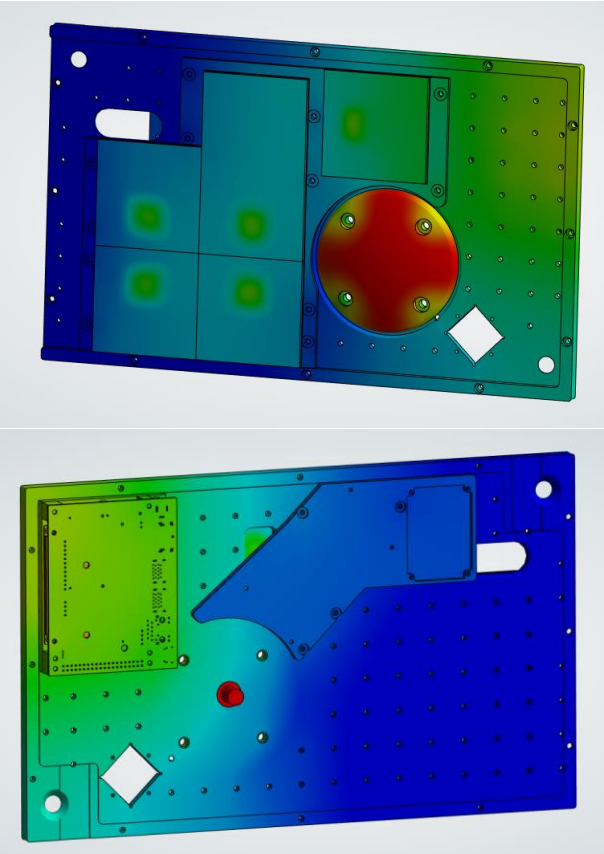


Figure 12: Example thermal analysis results. Solar panels originally planned to be mounted directly above the Jetson were removed and replaced with silver Teflon tape to improve thermal performance

The bulk of the testing was focused on the autonomy software, as much of it was newly developed for this spacecraft. The goal of the space observation mission was to simultaneously detect Earth resident space objects, detect background stars, and calculate a star tracker solution. The software for this mission was tested using simulated imagery and Earth-based astrophotography. The synthetic imagery was generated by injecting small Gaussian blobs or short streaks into base background images with structured noise (colored noise with stray light patterns). After injection, shot noise, sensor artifacts, and digital conversion artifacts are added. The astrophotography imagery was collected using the same imager and a similar lens as the flight payload, but with a variable aperture setting. Long exposure (0.1-1 sec) collections from a stationary mount on clear nights in a suburban area (SQM ~20 mag/arcsec^2) produced sufficiently representative imagery for detection and star solution processing testing. During testing, significant sagittal astigmatism was observed at large aperture settings, and therefore selected a smaller, fixed aperture for the flight unit.

The goal of the maritime observation mission was to detect large ships. No existing public data sets had representative imagery similar to our camera's field of view, ground spatial distance, and video frame rate. Therefore, the maritime observation software was tested with synthetic imagery and modified public imagery. For synthetic images, data was generated using DIRSIG with ocean backgrounds at various sea roughness states, pointing hold and pushbroom control settings, and at various off-nadir angles.[15] For further testing with more realistic imagery, synthetic point targets were added to background images/video. Backgrounds were gathered from the ISS High Definition public live stream and the GOES-16 ABI GeoColor imagery.[16, 16]

End-to-end testing was performed by developing an emulator for the R5 flight computer software. The functionality of the emulator was primarily focused on the interface between the payload and flight computers, orbit and state propagation, and uplink/downlink emulation. The emulation software was deployed on a second Jetson to simulate the network interface and startup connection similar to flight. Testing results in this emulated environment served as a baseline of expected behavior for final integration, verification, and validation with the actual flight hardware and cFS-based flight software.

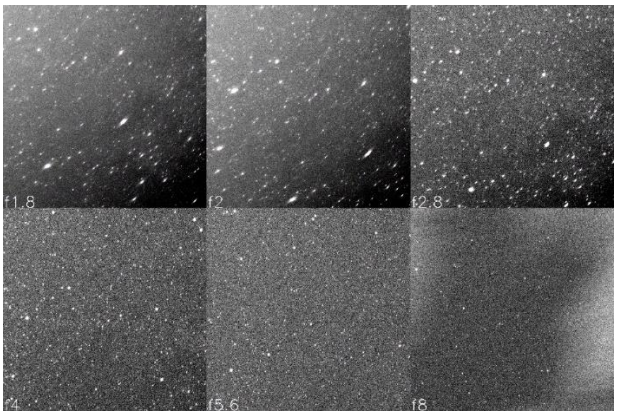


Figure 13: Earth-based test astrophotography image chips of the bottom-right corner of images collected from a camera similar to the payload camera, at various aperture settings. Large apertures (smaller f-number) produce significant sagittal astigmatism.

Spacecraft 9 Avionics

The avionics system was developed from the R5-S10 baseline. To accommodate the power demands of the new payloads, the custom PDU was updated to v5, parallelizing 12 V regulators to achieve 60 W single-rail output, improving power and battery monitoring, and simplifying power selection logic. Additionally, the custom interface PCBA was updated to v3, incorporating several more load switches in 5 V and 12 V banks to continue the practice of switching on payloads only for their operational phases.

Development of R5-S9 included work into the potential for deployable solar array wings (SAWs) which would make the R5 bus power-positive except during high-power operations. To this end, the propulsion controller PCBA was updated (to v2) to include four more solenoid valve actuator circuits which would be used in the deployment mechanism. The solar wiring was planned in parallelized panels with blocking diodes to allow for several potential configurations and to be tolerant against deployment failure. The GPS antenna would also be placed on the furthest deployable panel to improve reception, as found in integrated testing. Ultimately the SAWs were not implemented due to time constraints, but this work will be carried forward on the next R5 iteration.

Spacecraft 9 Communications

The Iridium®-based part of the R5-S9 communication subsystem was identical in hardware and configuration to the R5-S10 baseline. The WiFi system was removed. The R5-S9 communication subsystem also included two new prototype systems for demonstration on this mission; a LEDComm v2B unit developed in partnership with the Aerospace Corporation, and a Ground Control RockBLOCK 9704 with integrated patch antenna (WRL-28896).

The LEDComm unit was configured for 100 kbps to maximize signal to noise ratio for this first demonstration of the system. While it was integrated into the rest of the spacecraft in such a way that it could be used to transmit on-off-keyed (OOK) data, the FSW was only configured to operate the unit in a pseudorandom bit sequence (PRBS). The PRBS should be sufficient to determine the performance of the system. A single ground station at the Aerospace Corporation's El Segundo, CA location is currently active with discussion ongoing about bringing more online.[13]

The Iridium® 9704-based radio shows promise for offering substantially greater data rates over its IMT messaging protocol compared to the baseline R5 TT&C system of the 9523's SBD® protocol, though this has not yet been demonstrated in space. The 9704-based system has been implemented on R5-S9 in such a way so that its connectivity capability can be characterized in space to inform its deployment on future missions.

Spacecraft 9 Flight Software

The R5-S9 FSW was developed from the R5-S10 baseline. With a substantially different mission from R5-S10, there were a number of changes across the FSW system. These include the addition of an IO application for the Sandia payload and a secondary star tracker camera (using the same star tracker image processing application as the first camera), updates to the GPS application and system at large to enable time synchronization from the GPS with or without PPS, and updates to the AFM to match the mission ConOps and enable its unique elements. The optical communication and 9704 radio demos were integrated directly into the AFM to reduce development time instead of getting their own applications and more traditional instantiation within cFS.

New commands were implemented for this mission, including the ability to set the spacecraft time and position from the ground. Additional ground test scripts were also developed to support more flight-like integrated system testing to exercise the interface with the Sandia payload.

Spacecraft 9 Guidance, Navigation, and Control

The R5-S9 GNC system was modified from the R5-S10 system to remove the RelNav filter, LiDAR processing,

and RelNav camera as those capabilities are not needed for the R5-S9 mission. An additional Star Tracker camera was added to improve likelihood of consistent attitude measurements and additional EMI mitigation was applied to the GPS.

R5-S9 guidance implemented a generic attitude pointing method capable of inertial, earth-fixed, or moon pointing upon command. For earth-fixed pointing, guidance generates an attitude error and rate target to maintain pointing to the ground target as the vehicle passes overhead. The pointing method generates the pitch/yaw orientation that fixes the +X body / Sandia camera bore-site on the target, while a secondary selection approach determines the roll around the +X axis. This selection approach uses the earth-moon and sun angles incident to each star tracker camera boresight in order to select a camera and roll orientation. The selected roll orientation minimizes the likelihood of stray-light impact on the selected star tracker that could degrade or prevent attitude measurements.

The control system employed a piecewise, bang-bang-inspired attitude control algorithm known as a phase-plane controller. The controller consists of two operating regions: a maneuver region, which commanded vehicle slews toward the desired attitude, and a hold region, which maintained pointing within configurable deadbands.

A thrusters-only approach to earth-fixed pointing proved intractable due to significant on-time usage throughout the slew during overhead pass. Thus, reaction wheels were employed to maintain the desired attitude and slew rate for Earth-fixed pointing. A PD controller generated wheel speeds to zero attitude error while tracking the rate target from guidance. In simulation it was observed that the selection approach in guidance occasionally caused a 180 degree roll slew due to the relative positional changes of the stray-light sources (sun, moon and earth). When this occurs, CNTRL zeros wheel speeds and allows the phase plane to maneuver before re-engaging the wheels for pointing maintenance.

Spacecraft 9 Propulsion

R5-S9 propulsion subsystem maintained the baseline of the R5-S10 propulsion system in terms of component selection. The thrust level was tuned to the lower baseline demonstrated on R5-S2 and R5-S4 as opposed to R5-S7 and R5-S10 as the former missions were more focused on pointing performance and the latter on translational maneuvers.

Otherwise, the major difference between R5-S10 and R5-S9 was the addition of two valves in the low pressure assembly to support the integration of deployable panels. The design implementation of the panel deployment mechanism used the nitrogen pressure to actuate a small 1/8 piston to release a latch. There was a separate latch for each of the two faces for array deployment, and so two additional thruster valves were used in parallel to control the flow of low pressure gas to enable deployment. The system would have the ability to repeatedly try to deploy the latch mechanism through control of holding open the thruster valves to depressurize the volume behind the piston and reprime to move the opening piston again. The deployment valves were capped off and not used in the final build of the vehicle as the deployables were descoped to meet schedule, but they were present in the vehicle during random vibration testing and did not leak or require to be removed when they were not used, showing the flexibility and resilience of the propulsion system design.

Spacecraft 9 Structures

R5-S9 built upon the R5-S10 baseline and lessons learned. The WiFi and LiDAR were removed and the accuracy of the CAD model overall was improved compared to the R5-S10 model. The Sandia payload, 9704 radio, optical communication system, and additional star tracker were added based on feedback from the payload points of contact and GNC. Deployable solar panels were removed from the baseline configuration after schedule delays prevented their integration before environmental testing. EMI mitigations, like applying foil tape to certain components and relocating the GPS antenna, were also applied to enhance GPS performance.



Figure 14: R5-S9 interior prior to lid closure

Spacecraft 9 AIT Results

In the absence of flight results (R5-S9 is currently awaiting launch), results of AIT are included. R5-S9 was subject to random vibration testing and a litany of limited-scope functional tests and far more integrated/involved day-in-the-life tests to verify the system's function and performance. The day-in-the-life tests were scoped and executed to align with the expected on-orbit operations as much as possible and will serve as an expected performance baseline during the mission. This section will focus more on the unique elements of the AIT phase.

All prior R5 vehicles have had very poor GPS performance. With R5-S9, this was traced back to electromagnetic interference from several avionics sources inside the spacecraft, including the GPS receiver itself. These sources were mitigated with copper foil tape and the GPS antenna was moved further away from the sources. With these mitigations, the GPS regularly and quickly acquired a lock during live-sky terrestrial testing.

During testing it was observed that one of the reaction wheels failed to spin. Simulations with a failed reaction wheel severely degraded pointing performance and increased propellant usage, thus it was decided to disable the wheels for the primary S9 mission. To assess the performance with jets only, a 335 run MC was performed that yielded the pointing error and single thruster on-time shown in the figure below. A 90s time limit imposed on the pointing modes was implemented, therefore some runs do not fully zero the attitude error. This strict time limit was imposed to reduce prop usage from deadbanding, to enable more pointing opportunities.

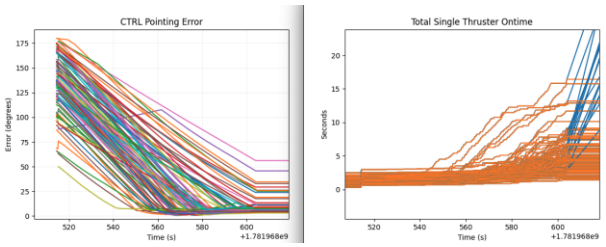


Figure 15: R5-S9 MC results

Testing on the flight vehicle was performed to verify the polarity of each GN&C sensor. IMU polarity was verified through tilt and rotation tests that recorded the IMUs outputs from FSW in the body frame. Upon inspection this data matched expectations. Star Tracker polarity was verified with live-sky testing: recorded data from FSW was compared to ICRF to body quaternions generated independently using astropy. Thruster mapping was verified through a sequence of thruster firings: thruster 1-12 were fired in increasing order while air pulses from the expected thruster in the expected order were identified. Wheel polarity was verified by commanding the wheel demo and verifying the wheels spun in the expected direction in the expected order. Multiple attempts were made to support 3DOF testing (including a turn table and suspending the spacecraft on strings), but none enabled the spacecraft to significantly move under 1G with its effectors. An air bearing system is in work for future spacecraft.

Spacecraft 9 Lessons Learned (so far)

Avionics lessons learned include the generation and mitigation of EMI across the R5-S9 system (detailed above). Avionics also encountered the challenges of working with small connectors (like JSC connectors) and how easily they can be damaged and how hard they are to repair. This resulted in the loss of the ability to use the RelNav camera on this mission, which is acceptable as it's not required for mission success.

The communication subsystem lessons learned include the behavior of the 9704-based radio system. The system shows a variable initialization time that must be accounted for by the FSW in addition to a variable transmit time that appears to be correlated to the size of the IMT message that the radio is attempting to transmit.

The FSW subsystem learned that verification can take much more time than development and the subsystem schedule and priorities should reflect that. Relatedly, FSW noted that final flight configuration verification should be done prior to the final flight configuration verification as at that point no time remains to address any discrepancies found. The FSW subsystem also repeatedly noted that integrated testing highlighted previously unknown bugs (frame discrepancies,

unexpected payload configurations, and many more) and emphasized the value of getting to integrated testing as soon as possible. The FSW team also saw a lot of value from the avionics team being able to independently update test scripts without FSW team assistance, freeing up the FSW team to work more in parallel and highlighting the benefit of keeping the test script architecture as simple and accessible as possible.

GNC lessons learned include the challenges in updating the Trick simulation to support dispersed pointing geometries during MC runs. Additionally, a difference between simulation scenario configuration and the ConOps resulted in an untested AFM transition that ended up needing to be debugged late in AIT. GNC also learned that the R5-S9 reaction wheel control instantiation is sensitive to a single wheel failure as it degrades performance across all axes. Finally, GNC discovered the flight GPS was not outputting PPS which was traced back to functional checkout procedures lagging behind mission requirements (missing PPS checks) and a misconfiguration of the unit due to a procurement miscommunication.

Propulsion lessons learned include the value of utilizing common parts across multiple vehicles and remaining hardware rich. This enabled the quick replacement of a leaking thruster pod.

For the structures subsystem, a thorough 3D printed fit check was replaced by a simplified check of a subset of the components. As a result, there were a number of interferences that weren't caught until final assembly. This further verified the advantage of the structural architecture as the printed flight brackets were quickly modified and re-printed to then install without issue. However, this could have been avoided by a complete 3D printed fit check. This would also have highlighted that the inhibit switches were designed and installed onto the wrong vehicle face. Only because of the symmetry of the vehicle and the lack of ProxOps on this mission, this misconfiguration has no mission impact.

CONCLUSION

The design, AIT, launch, and operations of R5-S7, R5-S10, and R5-S9 have provided feedback on the R5 process and verified its ability to deliver highly performant vehicles within their extraordinarily aggressive cost and schedule targets. The R5 project team hopes that others across government, industry, and academia will be able to leverage aspects of the demonstrated process and lessons learned to enhance and accelerate their work within existing resource constraints. R5 will continue to share details of the evolution of the process and its results and remains highly motivated to partner with other entities to accelerate significant demonstrations and the infusion of the process and lessons learned.

In the next 12 months, the R5 project is planning on launching Spacecraft 6 (R5-S6), Spacecraft 8 (R5-S8), and Spacecraft 11 (R5-S11). These spacecraft will demonstrate a variety of payloads and improvements to the R5 bus. These include autonomous cislunar optical navigation, enhanced optical comm transmitters, modal propellant gauging, improved software containerization, deployable solar array wings, and more.

Acknowledgements

R5-S7, R5-S9, and R5-S10 were funded by the Small Spacecraft Technology program within NASA's Space Technology Mission Directorate and Johnson Space Center's Engineering Directorate. R5 is actively pursuing new public-private partnerships and mission-focused funding partners across the US government and industry. The authors would like to thank our families for their support through the late nights and weekends. In addition, the authors would like to thank the many SpaceX, Maverick, Momentus, SPLICE, Sandia, SEOPS, Iridium®, Launch Services Program, and DynaQual personnel that helped get these vehicles across the finish line and into space. The authors would also like to thank all R5 team members and support personnel past and present for enabling the success of the mission.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. SAND#0000-XXXXX. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

References

1. Berck, J. et. al., 2025. "Launch Fast Learn Fast: Lessons Learned from the R5-S2 and R5-S4 Missions," *Proceedings of the AIAA/USU Conference on Small Satellites*, SSC25-RAIV-01. https://digitalcommons.usu.edu/smallsat/2025/RA-S4-2025/1/
2. NASA Goddard Space Flight Center, 2025. "NOAA-N Prime Mishap Lesson Learned 1580," *NASA Lessons Learned Information System.* https://llis.nasa.gov/lesson/1580

3. Banker, B, and Askew, S., 2019. "Seeler 1.0: Prototype Robotic Free Flying Inspector Mission Overview," *Proceedings of the AIAA/USU Conference on Small Satellites*, SSC19-XI-04. https://digitalcommons.usu.edu/smallsat/2019/all2019/151/

4. Moore, R., 2014. "ISS Inspection Capabilities and Challenges," *NASA Technical Reports Server.* https://ntrs.nasa.gov/api/citations/20150003019/downloads/20150003019.pdf

5. Fredrickson, S., 2012. "Mini AERCam for In-Space Inspection," *NASA Technical Reports Server.* https://ntrs.nasa.gov/api/citations/20120002583/downloads/20120002583.pdf

6. Pedrotty, S., et. al., 2020. "Seeker 1.0 GNC Flight Performance," *43rd Annual AAS Guidance, Navigation, and Control Conference.* https://ntrs.nasa.gov/citations/20200003117

7. Johnson, B., 2025. "Development of a Hypersonic Entry Test Platform," *AIAA SCITECH 2025 Forum.* https://doi.org/10.2514/6.2025-0925

8. Onsay, R. 2025 "CGUL—A Modular Approach to cFS," *2025 NASA cFS Symposium*. https://ntrs.nasa.gov/citations/20250000518

9. NASA., "The Commercial off-the-shelf Camera and Computer-Based Star Tracker," *NASA GitHub.* https://github.com/nasa/COTS-Star-Tracker

10. Berck, J., 2024. "Pose Bowl: Spacecraft Detection and Pose Estimation Challenge," https://www.nasa.gov/directorates/stmd/prizes-challenges-crowdsourcing-program/center-of-excellence-for-collaborative-innovation-coeci/pose-bowl-spacecraft-detection-estimation-challenge/

11. https://www.prophesee.ai/event-based-sensor-imx636-sony-prophesee/

12. https://kitware.github.io/nest-project-page/

13. Hinkley, D. and Pedrotty, S., 2026. "CubeSat LED Transmitter for Optical Communications," *Proceedings of the AIAA/USU Conference on Small Satellites*, SSC26-IV-07. https://digitalcommons.usu.edu/smallsat/2026

14. https://www.rit.edu/dirs/dirsig

15. https://eol.jsc.nasa.gov/esrs/hdev/

16. https://www.goes.noaa.gov/fulldisk.php?sat=G19

17. Knesek, K., Alexandar, M., Wisbiski, J., 2025. "Avionics Design Architecture for Low-Cost CubeSat Missions and Lessons Learned from R5-S2 and R5-S4," *2025 IEEE Aerospace Conference.* https://doi.org/10.1109/AERO63441.2025.11068544