SSC24-S2-05

Demonstrating Rapid Response for Remote Sensing Applications Using Automation and Intelligent Software on GEOStare SV2

Emily Milne, Michael Higgins, Bailey Cortright, Robert Goodloe, Alexander J. Pertica, Valentino Constantinou
Terran Orbital, Ali Mezher
400 Spectrum Center Drive; (949) 344-0181
emily.milne@terranorbital.com

Willem de Vries
Physics Division,
Lawrence Livermore National Laboratory

ABSTRACT

We conducted a campaign to demonstrate rapid responsiveness by tasking a nanosatellite with in-transit plane and ship image captures. By utilizing open-source information, approximate routes can be generated using ANSYS's Systems Tool Kit (STK) Software for these vehicles[1].

GEOStare SV2 was leveraged for this effort by utilizing a new imaging mode for tracking and sweeping between coordinate locations ("nodes"). GEOStare SV2 is a 6U CubeSat in a mid-latitude LEO orbit with a dual camera remote sensing payload. The satellite is operated by Terran Orbital, with the payload designed by Lawrence Livermore National Laboratories (LLNL). Since launching in May 2021, it has taken over 100,000 images of ground and space-based targets, including landmarks, environmental disasters, war zones, satellites, and asteroids at an average cadence of several thousand images per month. The capabilities of GEOStare SV2's attitude determination and control system (ADCS) enable unique imaging modes such as mosaics that can quickly and easily be leveraged by ground operators. For this campaign, mosaics was used to attempt to track and image the position of vehicles of interest. Mosaics is designed to sweep between several nodes at specified timestamps to command the line-of-sight of the camera such that it coincides with a location of interest at a given point in time. This imaging mode allows the spacecraft to capture more information within a given image session by ensuring a moving object remains in frame or by capturing larger portions of the terrain.

A ground scheduling system was used to build command parameters and send them to the spacecraft without an operator in the loop. Ground tools include filters based on the team's experimental findings for improving image integrity, such as weather and exposure times. The images can be downlinked and moved to a desired location without operator intervention. Upon receiving the images, the team is immediately able to view and determine the quality and contents of an image, then schedule new images or deliver them to customers.

CONCEPT

The original concept for this campaign was inspired by Resident Space Object (RSO) tracking. RSO tracking is a hot topic for small satellites, and this campaign sought to perform a similar, operator-in-the-loop version of ground vehicle tracking. The intent was to use publicly available route information for an aircraft, including waypoints, speed, and altitude to determine the exact location of that aircraft at several future times within a few minutes of each other. GEOStare SV2 would then be tasked with imaging locations along the route such that the aircraft remains within frame for every image over a series of imaging captures. As GEOStare SV2 approaches the aircraft, it would begin mosaic imaging with a cadence and number of nodes such that the movement of the aircraft and boresight cause the aircraft to remain in the center of each image.

In practice, it was determined that open-source route information was inaccurate and not suitable for determining the future location of an aircraft. As a result, the focus shifted to slewing around an approximated area of about 50-150km.

This new concept relied heavily on integrating historical route information for the same flight over several days, as well as looking at and adding waypoints in near real-time as they became available on FlightAware[2].

Figure 1: Original (Left) vs New (Right) Aircraft ConOps[1]

The V3 camera covers a ground extent of approximately 18.5x11.6 km, while the V4 camera covers approximately 5.8x3.9 km. Both cameras were used for this campaign. While most ground vehicles are too small to see clearly with the V3, the images from this camera are significantly smaller in data size and faster to downlink. They allowed operators to determine the satellite had slewed between images as expected due to the wider FOV given that most images were taken over the open ocean without any landmarks, and to decide whether to commit ground contact time to the downlink of the larger images. The V4 is also tolerant to much larger slews due to the shorter exposure time necessary to capture an image.

IMPLEMENTATION

Attitude Determination and Control System

Tracking and imaging targets on the Earth's surface using a spacecraft is, at least theoretically, a straightforward problem with a simple solution. To successfully image a target, it is necessary to first compute the line-of-sight (LOS) vector between the spacecraft and the target and then align the spacecraft payload boresight with this vector. This is a trivial calculation to perform, but the validity of the result is inherently dependent on accurate knowledge of the spacecraft's attitude, its orbital position and velocity, and the target's inertial position. Estimating the future location of a moving target using third-party data is handled on the ground, as previously described. GEOStare SV2's ADCS handles estimation and control of the spacecraft state to address the fundamental pointing problem in support of imaging operations.

GEOStare SV2 leverages an extended Kalman filter (EKF) to estimate the spacecraft's inertial attitude. This EKF fuses measurements from onboard star trackers and gyroscopes to provide an accurate prediction of the spacecraft's attitude which is updated at a 10 Hz cadence. This allows the spacecraft to minimize attitude knowledge error as it is attempting to align the desired payload boresight vector with the spacecraft-to-target vector. Attitude determination (AD) is facilitated by calibrating the star trackers and gyroscopes periodically. Star tracker calibration is initiated by aligning the optical payload boresight with the primary star tracker on the vehicle, which shares a common LOS with the payload boresight. The remaining star trackers are then aligned with this primary star tracker. This procedure minimizes high frequency measurement noise in the fused attitude measurements generated by mounting and sensor misalignments. Misalignment between the expected (i.e. configured) and actual star tracker boresight in the vehicle body frame leads to errors in the fused attitude quaternion solutions. Gyroscope calibration is carried out by estimating the gyro misalignment and scale factor errors. Incorporating both the gyroscope and star tracker calibration parameters into the AD EKF allows for significantly better AD performance, which in turn improves image quality. Better pointing knowledge ensures imaging targets stay in frame during imaging, and better pointing stability leads to less noisy images.

To support accurate orbital determination (OD), GEOStare SV2 leverages another EKF. This EKF updates the estimated spacecraft position using GPS measurements and a high-fidelity dynamics model. This dynamics model is based on the EGM2008 geopotential model and an onboard aerodynamic drag estimate. It can also incorporate gravitational disturbances from the Sun and Moon. This high-fidelity model ensures accurate inertial position estimation, even during periods of GPS measurement loss, minimizing biases that can arise from propagation errors and lead to off-centered targets in captured images.

While precise navigation is a crucial component of ground target imaging, control design is an equally important component. GEOStare SV2's control law uses a high bandwidth controller for aggressive tracking of imaging targets. This selection of a control law was made while abiding by a conservative design philosophy, with adequate gain and phase margin accounting for uncertainty in the spacecraft dynamics and disturbance model. Given that access to ground targets is typically brief due to GEOStare SV2's orbital altitude of approximately 550 km, momentum accrual in the reaction wheels while pointing is not significant. This provides the wheels with significant margin for tracking targets, during which time reaction wheel speeds can fluctuate significantly. This margin allows for reaction wheel momentum biasing, whereby the reaction wheel speeds are intentionally kept non-zero to decrease the frequency of zero crossings. Friction internal to the reaction wheels during zero crossings can induce mechanical vibrations that are undesirable during imaging operations.

To improve pointing stability during short imaging sessions, GEOStare SV2's ADCS can solely rely on propagation of gyroscope measurements at the expense of pointing knowledge. Use of gyroscope propagation is a decision that is made on a session-by-session basis, being primarily dependent on the intended session length and required pointing accuracy. In practice, there is a single ADCS parameter that controls this behavior, making it extremely easy for operators to enable this capability. Pointing stability can also be improved by adjusting the behavior of the constraint vector, that is, the vector orthogonal to the pointing vector which controls rotation about the pointing axis. The constraint vector boresight and type can be changed, with constraint types including the Sun direction, the velocity direction, a constant spin rate, or additional options. The constraint vector can also be disabled entirely. All of this constraint vector behavior is controlled by only a handful of parameters that can be changed in a single spacecraft command, making it easy for operators to optimize ADCS performance based on individual session objectives. Individual sessions can also easily utilize different control gainsets if more or less aggressive tracking is desired.

Completing the picture is ADCS's guidance algorithms. While GEOStare SV2's ADCS is equipped with a variety of standard operating modes – inertial pointing, RSO tracking, nadir pointing, stationary ground target tracking, to name several – it was necessary to upgrade ADCS with the capability to track ground targets with more flexibility than previously afforded to track more complex targets. Two such complex targets are outlined. First, it may be desirable to image large area targets that cannot be captured completely within the field of view (FOV) of the imager. In this case, it would be advantageous to sweep the payload boresight of the spacecraft along a defined path between waypoints. By capturing images at a regular cadence, the images could be stitched together in post-processing to form a mosaic depicting the target region of interest. The second case for more complex imaging is moving targets. Planes, boats, trains, rockets, and other vehicles represent potential targets of interest that follow trajectories which are often, at least to some degree, known ahead of time. While GEOStare SV2 is not equipped to receive telemetry from these vehicles in real-time for tracking purposes, it is possible to define a series of waypoints outlining their expected future position and compute tracking profiles based on these waypoints.

In the case of both area targets and moving targets, it is obviously advantageous to treat waypoints as more than a series of individually scheduled static ground targets. Upon further consideration, it also becomes apparent that computing slew profiles based on a constant ground rate tracking is better than using constant spacecraft slew rates for imaging purposes. To visualize why this is true for imaging area targets, one can consider the simplified 2D case of a static spacecraft imaging an area whose center is located directly nadir at a given moment in time. If imaging at a regular cadence using a constant spacecraft slew rate, the distances between the geographical center of each image will grow as the distance from the center of the area target grows. For a camera with a limited FOV, this means that images of areas furthest from the area target center may not overlap with their geographic neighbors. This can be prevented by decreasing the slew rate, but this decreases the potential size of area targets when limited access time is considered. Tracking at a constant ground rate addresses this geometric reality.

Commanding constant ground rate tracking rates to image moving targets is advantageous for a different reason; it more directly accounts for the dynamics of the targets themselves. Planes, trains, and other vehicles travel at a constant velocity for most of their trajectory. Therefore, defining a reference ground velocity rather than a reference spacecraft slew rate maximizes the LOS time within the limited FOV of the spacecraft camera during a limited access period.

Having established the utility of tracking at a constant ground rate, it was necessary to incorporate such capabilities into GEOStare SV2's guidance algorithms. Critically, they were designed from the ground up to be highly flexible with respect to sweeping geometry and easy to use by operators, requiring a minimum number of input parameters. Scheduling this type of tracking can occur extremely close to the desired start of imaging; it is essentially only limited by the time necessary for the spacecraft to slew to the first waypoint.

GEOStare SV2's ADCS provides a robust solution to the fundamental spacecraft pointing problem through the design of its navigation, guidance, and control algorithms. These components ultimately work together to minimize the errors associated with controlling the spacecraft state, reducing the uncertainty inherent in both computing, and tracking the LOS vector. Further, ADCS's guidance algorithms are well equipped to handle the specific imaging cases of interest. Figure 3 depicts a practical demonstration of using these guidance algorithms to image a large area target that could not otherwise be captured in a single image given the camera's FOV.

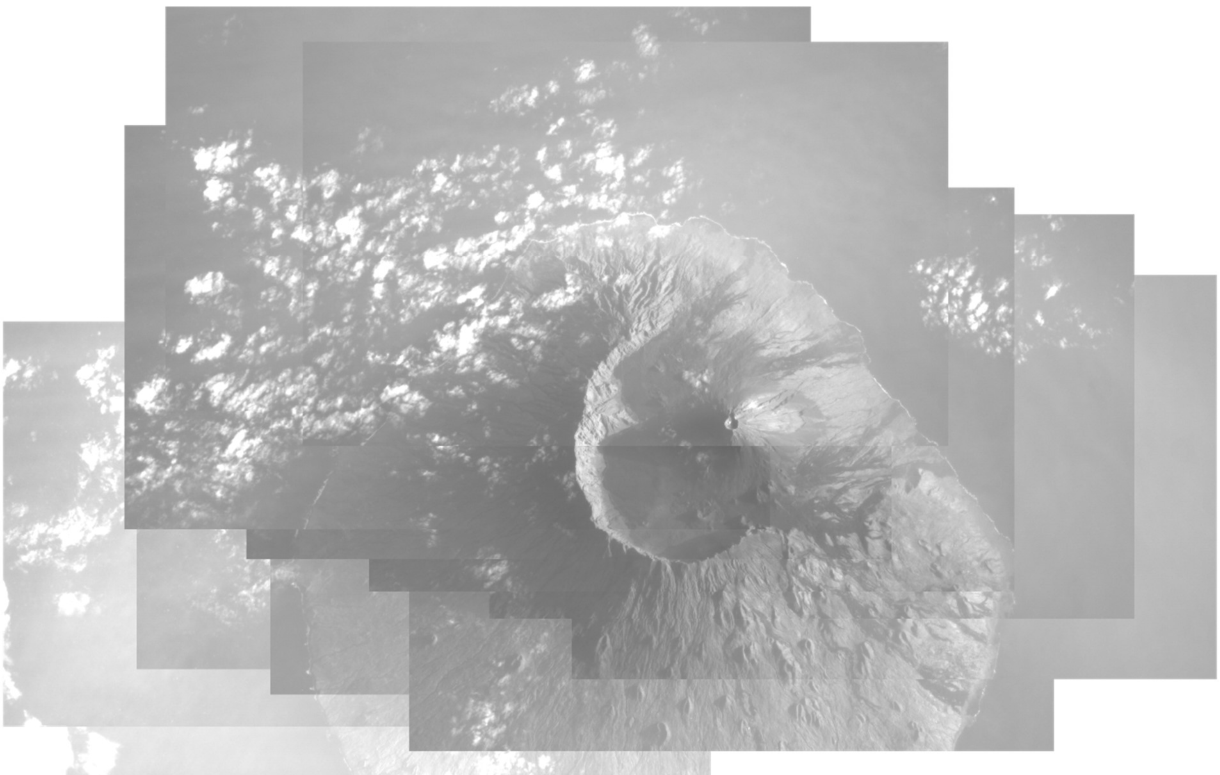

Figure 3: Mosaic of images taken by GEOStare SV2 of Pico do Fogo in Cape Verde on January 24th, 2024

Overall Automation

Images were scheduled using an automated process to turn camera, position, and time parameters into commands. Separate commands are used for scheduling the image capture and mosaics imaging mode. This allowed operators to choose the time of image capture based on access times to the vehicle an hour or two prior to selecting the locations and timestamps that the vehicle was expected to be at, contributing to a more accurate prediction while simplifying the commands that needed to be sent during a contact. For the aircraft and ships, the image capture commands were generally sent about 1-2 hours prior to the last ground contact before the capture, while the mosaic node commands were sent between 30 minutes prior to this contact to 5 minutes into the contact.

Asynchronous commanding software was used for any commands sent outside of passes. The software automatically sends commands to the satellite without operator intervention, freeing them to continue working on route prediction or other tasks.

Aircraft

For this campaign, operators utilized two public databases - FlightRadar24[4] and FlightAware[2] – to obtain a visual of current flights and route information. Operators selected a plane to track using FlightRadar24[4]. The criteria for selecting a flight were as follows:

- The flight needed to be longer than three hours to ensure time to task the satellite with the image capture. Due to the prevalence of flight delays around the world, it was necessary to finalize the imaging locations and times only after departure. This ensured sufficient time for loading data into STK and avoiding a flight phase change from cruise to descent, which would affect altitude and speed[1]. In practice, twenty minutes was often enough to generate these parameters prior to tasking the satellite.
- The access time from the satellite to the plane needed to occur while the plane was in a daylight zone. Areas of eclipse presented a challenge as the exposure time needed to be longer, further limiting acceptable slew rates to avoid blurry images, and preventing the use of the V4, as its necessary exposure times make it ill-suited for imagery in eclipse.
- The satellite is in a mid-latitude orbit, so the flight must not be in polar regions. It also could not be in the South Atlantic Anomaly, where imaging is avoided due to high radiation.

Once the flight was selected in FlightRadar24[4], FlightAware[2] was used to find the waypoints for the flight path as well as to view position and timestamp data in the past and near real-time. Operators inputted the data into STK and checked for access times with the satellite that met these conditions[1]. The scenarios also contained past flight paths for the same flight on different days.

As the selected access time for image capture approached, operators checked FlightAware[2] for velocity and altitude updates, as well as to compare the times at various positions that had since been added for the flight with predictions. Updates were made as necessary, and at the last ground contact opportunity prior to the image capture time, the commands were uploaded to the spacecraft.

This campaign presented several challenges, first and foremost being the level of aviation knowledge necessary to properly predict the flight path. The waypoints specified on the aircraft's flight path were in several difficult-to-understand formats for a team with no aviation experience.

While waypoints provided location information, they did not provide an estimated time of arrival at each waypoint. Looking through old flight data on FlightAware[2] helped provide an approximation of aircraft speed at each phase in its flight path: take-off, climb, cruise, descent, and landing.

Since the satellite was not being tasked until the aircraft had taken off, operators had the opportunity to bias selection of imaging times so that the aircraft was already at cruise when creating waypoints to determine position. This drastically cut down on differences in speed between the time of tasking and the time of imaging, as it was found that speed only varied by about +/- 10 knots during cruise phase, disregarding turbulence and weather. Altitude also remained relatively steady during this phase.

For this campaign, flights taking off or passing through Australia to the East were preferred. This allowed operators to utilize the Terran Orbital ground station located in Peterborough, Australia, for tasking of images right before the image captures. The satellite's orbit is West to East, and if the orbit is carefully selected, the image capture can be uploaded to the satellite within twenty minutes of the image being taken. This lets operators update time estimates for waypoints for as long as possible, increasing the accuracy. Flight ACA8 was eventually chosen due to the consistency of its flight route, convenient access times, and favorable lighting conditions. This daily flight from Hong Kong to Vancouver takes around 11 hours with a cruising altitude of 37,000 ft and speed of 650+ mph. In total, seven days of flight routes were added to STK[1]. Each addition helped narrow down the expected position of the plane. The figure below shows these seven flight paths. While multiple access times were available for the satellite on this route most days, the most frequently used time was around 20:45 UTC. At this point, the plane was closing in on its destination but had not yet begun descending. The black boxes in the area below represent areas of uncertainty for flights around this time. The image capture area the satellite was tasked with was partially dependent on route information up to the point of tasking. That route information would inform which of the smaller boxes was more likely to contain the aircraft. As might be expected, larger boxes generally correspond with lower degrees of uncertainty.

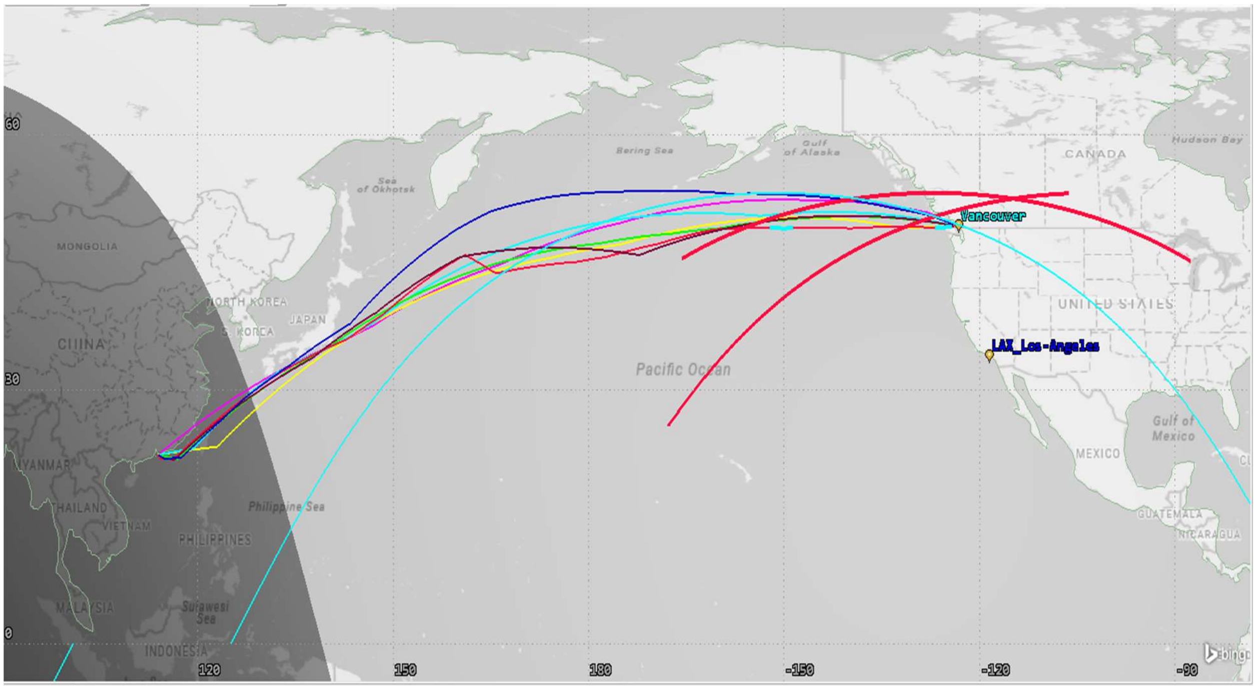

Figure 4: ACA8 Flight Routes April 30-May 8 2024[1,2]

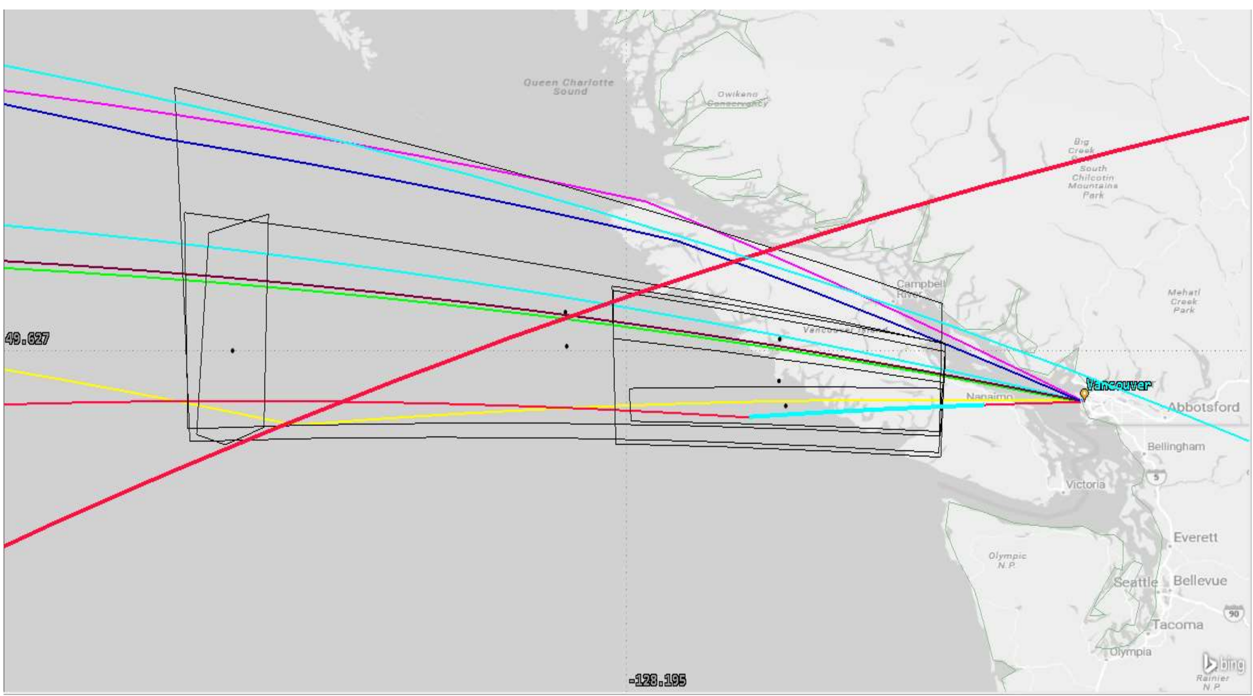

Figure 5: Predicted Locations of ACA8 around 20:45 UTC over several days[1,2]

Ships

Mosaic image sessions were scheduled to capture 3-4 coordinate locations along two busy shipping routes: Strait of Malacca and Strait of Dover. The goal was to capture a ship in transit, using the mosaics imaging type to capture as much area in the strait as possible along its narrowest point. The image below is from a session along the Strait of Malacca that captured three ships with the wide FOV imager. With the native resolution of the V3 camera, some details of the ship are already visible. Post-processing can reveal further information, including additional physical characteristics or location information. Below, the original V3 images is overlayed with the section containing the ship blown-up in both native resolution and utilizing the post-processing techniques developed in-house at Terran Orbital[5]. While not done for this image, image metadata and spacecraft information could be used to obtain the coordinates of the ship as well. For sessions over which the same ship is captured multiple times, a ship velocity could also be generated.

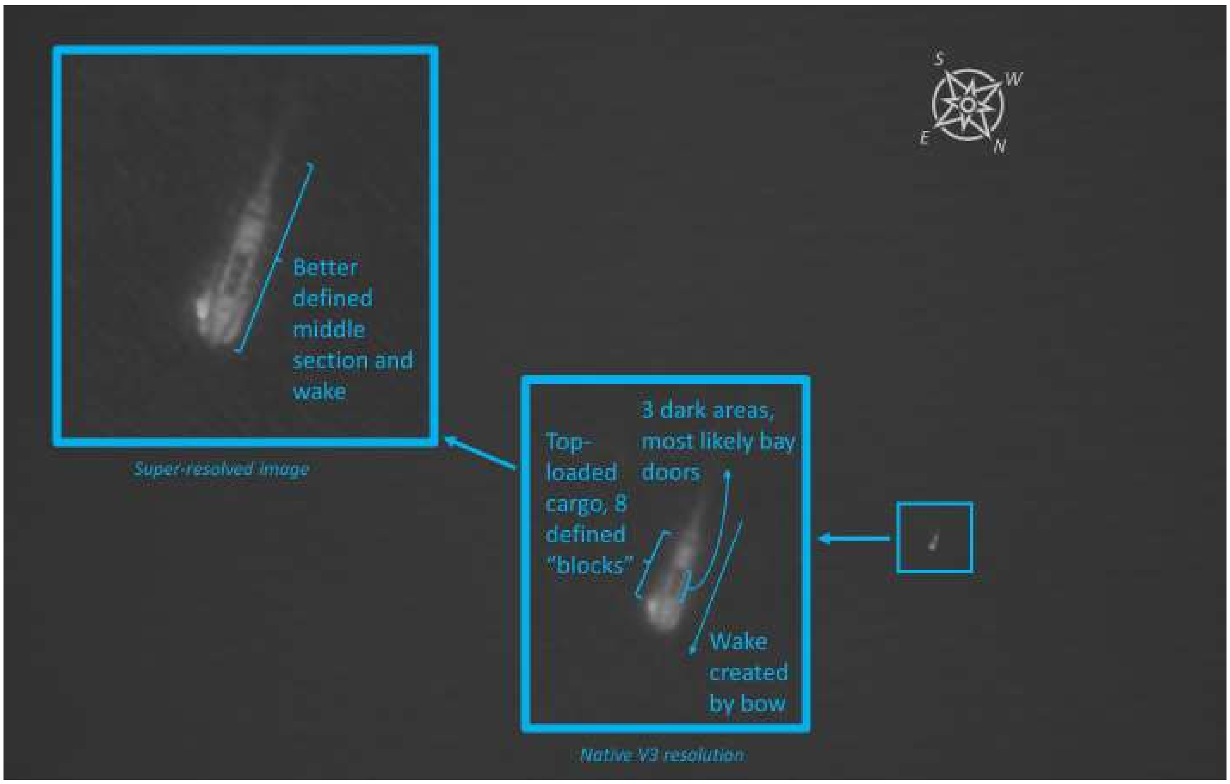

Figure 6: Ship Imaged Along Malacca Strait, 29 Jun 2024 03:36:00 UTC[5]

FUTURE WORK

We will continue improving our models to narrow down the expected area of a vehicle. As the system grows more robust, we may also reattempt to use this imaging mode for tracking by choosing nodes along the route path of a vehicle to follow it in transit. We will also work to rapidly re-task the spacecraft with expected positions for ships captured along a shipping route to image them over subsequent orbits.

While beyond the scope of this campaign, onboard software that can identify, center, and track an object given a reference area would be the ultimate goal. This would allow the satellite to combine the best of both worlds: searching a relatively wide swath of area around the expected vehicle position, and transitioning into tracking once the position has been established. This technique could also be applied to a constellation of satellites to enable rapid re-visiting of a location. This would both make it more likely to successfully image the vehicle and provide more tracking coverage due to a shorter revisit time. To be even more ambitious, satellite crosslinking could be utilized to pass information data between satellites in a constellation to be ingested into the algorithm. If the algorithm can accurately recognize a type of vehicle and is power-positive enough for continuous imaging, there is no need to provide a starting area either – the satellite can identify vehicles of interest as it captures images, ingest the movement over a series of frames into an algorithm, and pass that information along to the next satellite in a constellation with an access time to the vehicle.

Areas for additional ground automation include automating the import of waypoints using a script that can recognize multiple formats to account for the inconsistency in formats within and across data sources, such as the naming and numerical conventions used in aviation. It would also include automatically generating best-guess coordinate sets and timestamps for a vehicle given historical and current route info once a model has been proven. Data could also be automatically pulled in and updated from online data sources.

CONCLUSION

Tracking ground and air-based vehicles proved much more difficult than tracking space-based objects from a scheduling standpoint due to a large degree of uncertainty in location. Accurate waypoints would have greatly increased the likelihood of success of this campaign. However, the campaign demonstrated the robustness of the mosaics imaging mode and introduced the team to flight patterns and modeling unpredictable vehicles. It also demonstrated the rapid tasking capability of the satellite, with some coordinate sets being tasked within 20 minutes of the image capture.

ACKNOWLEDGMENTS

We would like to thank Tyler Hilgendorf and Jacob Beaver from the Mission Operations team at Terran Orbital for contributing their aviation expertise.

REFERENCES

1. “ANSYS’s Systems Tool Kit,” *ANSYS Government Initiatives*. ANSYS, Inc
2. “Flight tracker / flight status,” FlightAware, https://www.flightaware.com/.
3. “Global Ship Tracking Intelligence: AIS MARINE TRAFFIC,” MarineTraffic, https://www.marinetraffic.com/.
4. “Live flight tracker - real-time flight tracker map,” Flightradar24, https://www.flightradar24.com/.
5. V. Constantinou *et al*., "Leveraging Deep Learning for High-Resolution Optical Satellite Imagery From Low-Cost Small Satellite Platforms," in *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 17, pp. 6354-6365, 2024, doi: 10.1109/JSTARS.2024.3365417.