Concept of Operation and Initial Performance Summary of the NorthStar Space-Based Optical SSA System

Daniel O'Connell

NorthStar Earth and Space, Inc.

Jean-Claude Leclerc, Noemi Giammichele, Ryan Comeau, Narendra Gollu, Naron Phou, and Frederic Pelletier

NorthStar Earth and Space, Inc.

1. ABSTRACT

On January 31st, 2024, Rocketlab launched NorthStar's first four satellites with electro-optical payloads into Low Earth Orbit (LEO). This was the initial deployment of the very first space-based constellation dedicated to Space Situational Awareness (SSA), designed and developed by a commercial company solely through private investment. Each of the four 16U platforms hosts a 20 cm aperture payload designed to meet the NorthStar's mission requirements. The launch was a complete success, however, during the commissioning phase, several performance issues with the satellites were encountered that significantly impacted early results. These issues have been dealt with to varying degrees, with some still ongoing. The issues encountered are limited to the sensor performance only, and have not impeded the development and deployment of NorthStar's SSA data processing capabilities or any other aspects of its operations.

The focus of this paper will be on presenting NorthStar's patented Concept of Operations (CONOPS), tailored to the SSA mission, as well as describing the overall system design, focusing on the constellation configuration, an overview of the image processing and how observations are extracted from the raw imagery, and the end-to-end data processing pipeline. Additionally, some of the mitigation techniques developed to address the unique challenges encountered following initial deployment will be presented. The paper will also describe the creation and use of combined observables, including celestial positioning, photometry, and apparent motion.

Results from processing the operational data, relating to the performances of astrometric calibration, image processing, and orbit determination will be described. A performance evaluation will be presented, though it will be preliminary at this time given the sub-optimal sensor output. The paper will elaborate on the choices made by NorthStar over the process and lessons learned to demonstrate the value-add relative to existing SSA systems.

The technical details presented in the paper, together with the unique CONOPS employed by NorthStar's system, introduce a change in thinking in SSA capabilities through the space based SSA. The uniqueness in terms of observational mode/strategy, data collection methodology, and processing pipeline contributes to more frequent observations and unprecedented coverage capability that is not limited by geographic restrictions, ground weather conditions, or orbit regime. This in turn improves the accuracy from astrometric calibration to precision orbit information generation. These are the benefits of a sensor array deployed in space.

2. CONOPS OVERVIEW

NorthStar's CONOPS (CONOPS) was formulated around a set of objectives to develop a transformational SSA data and information service that would deliver both comprehensive coverage of Earth's orbital space in every regime, quality state vectors for Resident Space Objects (RSO's), and value-added analytics products, while significantly adding to the tracked object catalog. From the outset it was recognized that achieving these goals would best be accomplished by a space-based constellation of optical sensors. Hosting a sensor in space, particularly in Low Earth Orbit (LEO), would alleviate any issues with atmospheric phenomena or the geographic restrictions that are inherent

to ground-based systems, and would be much better suited for sweeping through wide areas of space in a short period of time. The limiting factor of the relatively narrow Field of View (FOV) of practical sensors is mitigated by deploying large numbers of sensors in multiple orbital planes, and processing the observations made by the entire set collectively, thus increasing the number of observations dramatically, and reducing the revisit times to individual objects.

Making effective use of such an infrastructure meant that the sensor systems could not be pointed in a track-and-stare mode as standard operations. While observing a single target for an extended period potentially can improve the solution for that target, it does so at the expense of missing everything else that could have been observed in that time frame, and would thus negate much of the advantage of hosting sensors in orbit, failing to produce the desired broad area surveillance capability. Therefore, NorthStar's operations are designed around a "pushbroom" approach, analogous to the method employed by most Earth observation missions that maintain a nadir-pointing attitude to sweep across the surface of the Earth. This methodology is absolutely key to detecting unknown objects with no a priori knowledge of their positions. Ground telescope ops such as the aforementioned tracking-and-staring approaches, are completely inadequate for broad area surveillance because they would take a wildly impractical length of time to complete one observation cycle. Ground-based radar systems have a much wider sensor FOV, but like the telescopes depend on the rotation of the Earth to help them scan a wilder Field of Regard (FOR), which also results in slow overall cycle time. Also, as the maximum range at which a minimum signal can be detected by a radar falls off as the 4^{th} power of the transmit power P_t, and other parameters, meaning radars are typically severely range-limited unless they are extremely expensive, very high-power systems, or are only detecting large, highly reflective objects. Placing optical sensors in space, using a pushbroom pointing profile, and imaging constantly on a non-stop, 24/7 basis, is therefore a far more efficient and practical means of scanning all of near-Earth space effectively in a timely manner.

For NorthStar's SSA application, the (nearly) fixed sensors are side-pointing, approximately normal to the satellite velocity vector and away from the solar vector, to optimize lighting conditions. The orientations are essentially fixed, with the exception of some minor adjustments over the course of an orbit to optimize observing different orbit regimes. With this implementation, visible objects in the FOV of the sensor appear as streaks whose lengths and orientations are a function of the relative dynamics of the objects and the sensor host. Therefore, because the host satellite state vector and orientation are very well known, it is possible to derive detailed information about the state of the observed object from the streak's centroid, length, and orientation within the image frame. Photometry and other characteristics can also be captured to aid in object characterization.

It is well known that the number of observations made of a given object, the spacing of the observations along the object's orbit, and the frequency with which it is observed (revisit time) all directly impact the quality of the resulting orbit solution. Maximizing the observation opportunities requires a high frame rate to produce contiguous image sets. NorthStar's sensors operate at a 1 Hz frame rate. Therefore, once the full constellation of 24 sensors is deployed, more than 2 million images will be produced every single day. The specific orbit parameters of the host satellites obviously have a major impact on the ultimate system performance. Hundreds of simulation runs, using an internally developed simulation tool custom built for the purpose, systematically evaluated different combinations of orbit altitudes, planar combinations, in-plane vehicle phasing, and Line of Sight (LOS) pointing angles, all with the intent of optimizing a set of performance parameters critical to producing a superior catalog and SSA information products. There were of course many other factors that played into the results, which are explained in greater detail in the following sections. An additional consideration was the need to adjust the desired orbit altitudes both during and after the conduct of the trade studies, due to the rapid proliferation of commercial mega-constellations like Starlink occupying nearby bands.

Lastly, the ground data processing is every bit as key to NorthStar operations as the space infrastructure ops. The data processing system operations must be scalable to accommodate processing more than 2 million images per day. Maximizing automation is therefore required to minimize the latency associated with the processing chain while maintaining commercial reliability standards. Furthermore, NorthStar has developed the capability to ingest third-party supplied data, as well as the constellation imagery, even incorporating other sensor types. The data fusion algorithms can produce even higher quality results whenever appropriate data resources are available.

3. SYSTEM DESIGN

3.1 Space Infrastructure Design

The concept of operations serves as the primary driver for the overall system design, encompassing both ground and space segments. Designing a system for space domain awareness utilizing passive optical cameras deployed on multiple satellites necessitates careful consideration of various trade-offs across multiple design aspects. Employing a system-of-systems approach, the overall design was systematically decomposed into distinct systems and subsystems, enabling a comprehensive and modular analysis of the design architecture.

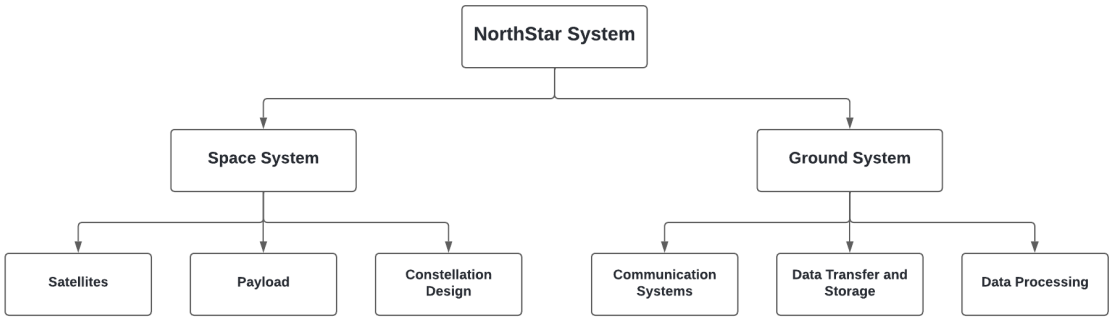

Fig. 1: Top level system hierarchy

A comprehensive simulator was developed to evaluate and optimize the design parameters influencing both the performance of the space system and the data processing of the ground system. The simulator encompasses a range of capabilities, including RSO propagation based on Two-Line Elements (TLE), multi-satellite propagation derived from flight parameters, a maneuver simulation model, a visible camera performance model, a detailed RSO photometric model, and a stray light model that incorporates Earth's ground albedo.

This simulator was utilized to conduct trade-offs among various design parameters, with the objective of optimizing key system capabilities, including the number of detections, the coverage of the RSO catalog, the revisit time of detections, and the minimum detectable object size. To facilitate these trade-offs, a baseline payload design was initially established. Based on the specifications of this baseline payload, the design of the satellite constellation was subsequently undertaken. The requirements for both the camera constellation and the payload were then refined and optimized using the simulator, enabling a comprehensive performance mapping across all target orbital regimes.

3.1.1 Space system: Payload

The operational concept involves utilizing passive cameras operating in a free-running mode, where an RSO will produce streaks within the captured images. The length of these streaks is directly correlated with the apparent relative velocity of the target in the image, as well as the camera's integration time. During this integration period, the light from the RSO is distributed across multiple pixels, which inherently impacts the system's sensitivity. Therefore, the camera aperture must be sufficiently large to optimize detectability in this passive imaging mode. A telescope aperture of 20cm was taken as a baseline.

Given the limited availability of space-qualified, cost-effective detectors, one of the initial decisions involved selecting an appropriate detector size. A detector with a resolution of 1280x1024 pixels was chosen. This choice imposes constraints on both the system's FOV and instantaneous Field of View (iFOV). For a high detection rate, the passive camera system must maximize the FOV. However, to maintain accurate positional determination of the target within the image, the iFOV should be minimized. By ensuring a point spread function of 3x3 pixels, sub-pixel accuracy in target positioning can be achieved. This configuration offers a favorable compromise between maximizing the FOV and minimizing the iFOV while ensuring precise target localization.

A camera performance model was developed to determine optimal system design parameters within the simulation environment. An initial baseline configuration of the satellite constellation was selected based on the best possible target illumination conditions and a predefined number of spacecraft. The camera parameters were subsequently optimized to maximize the detection rate, followed by multiple iterations to refine the constellation deployment parameters effectively.

The simulator is capable of generating high-fidelity images, which were utilized to develop the core components of the processing pipeline. The selection of calibration, detection, and measurement algorithms was directly incorporated into the trade-off analyses. Various noise sources—including detector readout noise, dark current, quantization noise, and both scene and background shot noise—were simulated at the image level to assess their impact on system performance. Considering the performance of these algorithms under a concept of operations that produces streaks for stars and RSOs, a detectability criterion based on the signal-to-noise ratio (SNR) was established. This criterion informed further optimization of the constellation and camera parameters, enhancing the overall system capabilities for the specified CONOPS. A minimal streak length criterion was also derived along the SNR criterion.

3.1.2 Space system: Constellation

The initial constellation design was based on preliminary specifications of the baseline camera payload. The key design parameters included the number of satellites, the number of orbital planes, the number of satellites per orbital plane, and the camera boresight orientation. The primary trade-off criteria for the payload were analyzed in conjunction with the primary design parameters of the constellation, assessing their impact on the constellation's overall performance. Multiple configuration iterations were conducted to optimize system performance across all target orbital regimes.

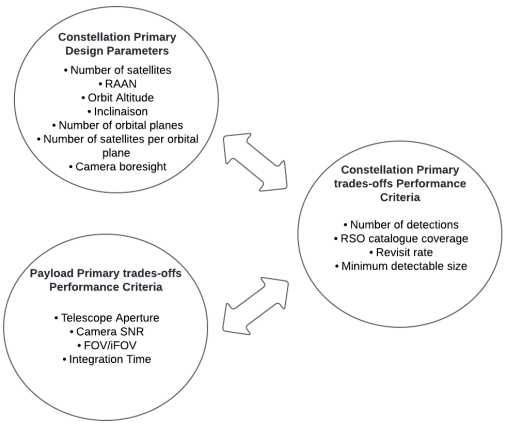

Fig. 2: Design trades dependencies

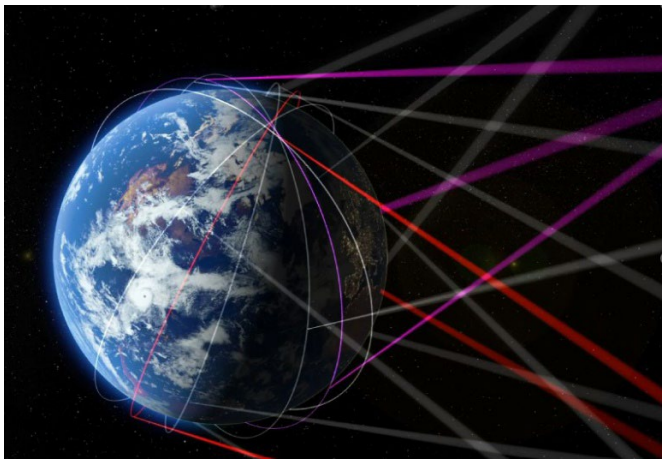

Fig. 3: NorthStar constellation with sensor FOV's depicted

The results of the simulation runs produced the following estimated performance numbers for the 24 satellite constellation. Note that these figures are estimated based on slightly more than 25,000 RSO targets run for a ten day period with the best available constellation configuration data, which is subject to change.

Table 1. Estimated Performance of the 24 Satellite Configuration

	Low LEO	Hi LEO	MEO	GEO	HEO
% Coverage [5m objects]	79.4	99.9	92.9	99.6	98.9
Min Size/Max Mag (Detection) [cm/M_v]	0.3/17.6	0.8/17.7	2.6/17.2	58/16.4	3.8/16.7
Min Size/Max Mag (Custody) [cm/M_v]	10/15.5	15/15.9	10/16.4	70/16.3	5/14.4
Median Revisit [hrs]	7.3	4.3	4.5	2.6	8.2

The final design of the constellation and the payload was optimized with a focus on minimizing cost and mitigating technical risks. The unique advantages inherent in the CONOPS, such as streak generation and the capability to detect both known and unknown RSOs across all orbital regimes, were preserved in the design and effectively integrated into the processing pipeline. The CONOPS guided the development of a system featuring a FOV of 2.8 x 2.2 degrees and an iFOV of 38 microradians. A system aperture of 20 cm was selected as an optimal balance between detectability, system complexity, and cost. The selected constellation design parameters values were: SSO orbit, fixed pointing with respect to orbital plane with a 260 degrees angle from the velocity vector, and 6 planes of 4 spacecraft per plane at LTANs of 0300, 0400, 0500, 0700, 0800, and 0900.

3.1.3 Space system: Satellite Bus

The primary design constraints for the satellite, as dictated by the CONOPS, include downlink capability, storage capacity, power requirements, and pointing stability. The CONOPS establishes conditions where precise pointing knowledge is not critical, as the processing algorithms utilize stars to accurately determine the camera boresight in each image. Since stars are employed for the geometric and distortion calibration of each image, thermal stability and low-frequency jitter in pointing are of lesser concern. The distortions caused by thermal variations and artifacts induced by low-frequency jitter affect both the stars and the RSOs. Consequently, the use of stars for calibration effectively compensates for these effects in the calibration of RSO detections.

The primary requirement focused on minimizing high-frequency jitter in the camera's pointing accuracy. Multiple design iterations were conducted to achieve the necessary performance level within the constraints of the selected 16U satellite bus.

3.2 Ground system: Data Handling / Pipeline Infrastructure

The ground system comprised the following components: ground station locations, data transfer and storage infrastructure, and the data processing pipeline.

The integration of downlink station locations and timing into the simulation allowed for an evaluation of their impact on key performance criteria. A configuration permitting one downlink per orbit per spacecraft was determined to be the optimal compromise, balancing system storage capacity, downlink bandwidth, downlink duration, and overall downlink cost. A trade-off was made regarding data latency, as the complexity, cost, and schedule required for implementing an initial real-time downlink capability were beyond the project's constraints at that time.

The large volume of data generated by the passive camera CONOPS imposed significant demands on data transfer and storage. With a configuration of one downlink per orbit per spacecraft, up to 5,400 images could be downloaded per spacecraft, approximately 15 times per day. The system was designed to be triggered by data availability and to operate fully autonomously, providing 24/7 service. The hardware infrastructure of the processing pipeline was scaled accordingly to handle the substantial volume of data that needed to be processed.

3.2.1 Ground system: processing algorithms

The data processing algorithms were part of the system analysis and trade-off studies. The CONOPS, which utilizes the entire image, ensures the presence of a significant number of stars of varying magnitudes across each image. This facilitates the use of stars for radiometric and geometric/distortion calibration in every frame. The widespread presence of stars around the detected RSOs makes them an efficient tool for correcting system distortions.

The calibration methodology was designed to leverage the stars within the FOV of every image, thereby relaxing the requirements for thermal stability and low-frequency jitter. By developing the calibration in the light flux domain rather than the intensity domain, the calibration process remained effective across different camera integration times, simplifying the optimization of image capture for various RSO orbital regimes when necessary.

The streaks produced by the CONOPS led to a low signal-to-noise ratio for the detection limit. By relying on the streak shape of the object, detections were possible even at a static SNR of 1. A threshold SNR of 2 was chosen as the detection limit criterion, thereby defining the minimum detectable size of RSOs under the chosen payload aperture and constellation design.

The image capture rate resulted in overlapping star fields in subsequent images, providing a unique method for filtering out false positives caused by stars or cosmic rays. A robust and reliable identification process was developed from this continuous imaging CONOPS, allowing for a dependable low SNR detection limit. This approach helped minimize the system aperture, thereby reducing system cost. The identification algorithm utilized multiple measurements from the streak, including streak length (which provides a measure of the object's velocity), streak orientation, streak photometry, and the streak center position. This multifactor identification process resulted in highly reliable identification of all known RSOs across various orbital regimes, also allowing for the use of a lower SNR detection limit criterion.

The RSO position measurement in the image was also designed to accommodate the need for a multi-pixel PSF. Sub-pixel precision was thus achievable, enabling the system to maximize the iFOV and, consequently, the overall FOV, thereby improving the detection rate across all orbital regimes. This approach minimizes the number of spacecrafts required to achieve the desired coverage of the RSO catalog.

The orbit determination algorithm was developed to capitalize on the advantages of the CONOPS in terms of data coverage. Being space-based, the system has the ability to observe RSOs over restricted areas, oceans, and polar regions where ground-based detection systems are typically absent. The complementary data provided by the CONOPS enhances the orbit determination process by supplementing the existing ground-based data. This information was instrumental in optimizing the constellation configuration to maximize value relative to the deployment schedule.

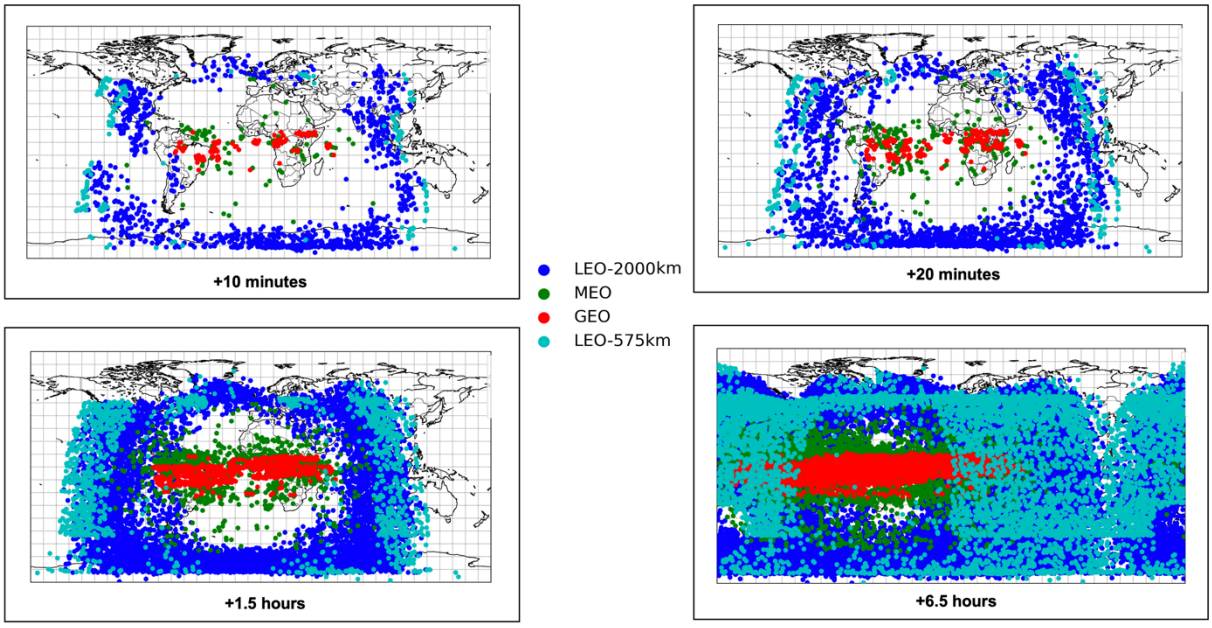

Fig. 4: 2D Projection of a simulation in ECEF with respect to time of orbit class distribution around the Earth in the 24 satellites configuration.

4. ASTROMETRIC CALIBRATION

To reach the best performances out of our system, detailed knowledge of the artefacts that the optical systems introduces into their measurements is critical. To achieve the ultimate limits imposed by the intrinsic noise sources of

the photodetection processes, one must ensure that all deterministic instrument characteristics superimposed on the measurements are known, and that they are properly removed or carefully considered.

The trades studies, driven by the CONOPS, called for a sensor FOV that was sufficiently large to include multiple stars within each image. The calibration methodology involves an initial calibration of the focal plane array using dedicated ground-based measurements to account for the focal plane array artifacts, followed by the use of stars within the image to perform accurate radiometric and geometric calibration for each image. This approach enables continuous correction for sensor performance drifts caused by factors such as thermal instabilities or aging.

Our calibration approach is based on the physical processes involved in the whole sequence of transformations from incident photons through detected photoelectrons, up to measured digital levels. It is a pixel-wise method, meaning that each pixel of the image sensor is independently characterised and calibrated. The method is initially based on a procedure developed for infrared cameras [1], while being adapted to account for the measured specific behaviour of the visible-NIR cameras. However, it goes much further in that it covers all artefacts that the instrument may introduce in its optical measurements. To achieve this goal, the whole calibration aspect is split into four independent parts.

The first part deals with the artefacts added by the sensor, to improve the relative uniformity of the image, referred to as “equalisation”. Its goal is to evaluate all the image sensor specifics that are part of the raw image. These can be removed by appropriate ground data processing. The equalisation block converts the raw image data, in units of Digital Levels (DL’s), into an image representing the photocurrent, i.e., the current inside each detector pixel produced by the detected light flux, in units of electrons per second. Its main role is to spatially “equalise” the raw image by correcting for the main random contributions between the pixels of the image sensor.

The second part relies on a more global instrument characterisation (image sensor plus its dedicated optics), to give an absolute value to the illumination detected by each pixel (radiometry). It carries the information required to convert the photocurrent image into an image of photonic irradiance, in units of photons per second per square meter. Such a physical quantity corresponds to the radiometric way of depicting the calibrated illumination level on the image sensor surface, as produced by the optical system. It thus corresponds to the calibrated image.

These first two parts are the ones that are described in [1]. Pixel-wise on-ground corrections for these 2 categories enable providing images in known absolute units, which in turns enables properly quantifying of the irradiance received from any space object.

The third part is based on the characterisation of the point-spread function of the device (see Fig. 5). The PSF represents the spatial distribution of an imaged point source on the sensor surface produced by the instrument’s optics. The selected model is an elliptic binormal distribution which is represented by a set of three parameters, whose spatial behaviour is characterised across the full image. In NorthStar’s implementation, the PSF is used as an initial seed to facilitate the optimisation of star and RSO models. The PSF information is never used to correct the image since such a process, associated to deconvolution, would be ill-conditioned. This would tend to amplify spatial noise and camera artefacts, so its usage is avoided.

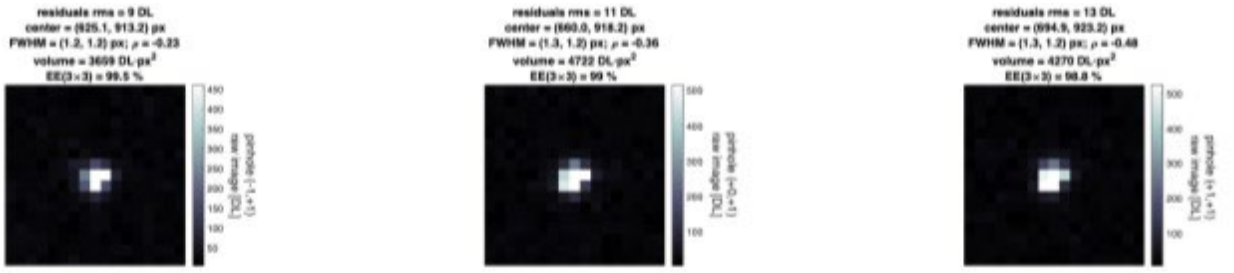

Fig. 5: Example of the PSF characterization done at multiple locations on the focal plane array at room temperature and at atmospheric pressure.

Finally, the last part, the calibration of distortion, plays a critical role since it details how the instrument bends the image position relative to its ideal imaging capability. This is vital relative to the essential task of object localisation.

It contains the information depicting how an object image is shifted (horizontally and vertically) in the image plane relative to its ideal position, as (linearly) predicted by the effective focal length of the optical system. Again, the spatial behaviour of such parameters is characterised across the full image. As with the PSF, the distortion information is never used to correct the image, because such a process would be ill-conditioned relative to spatial noise and camera artefacts, as it would require the fine interpolation of the calibrated image. It is used instead to correct the localisation of objects in the image plane once the optimisation methods have provided an estimated position of the target object in the processed image.

Under nominal space operations, dedicated data sets, such as the ones acquired during the commissioning phase, are regularly analysed to ensure the validity of the active calibration data. Adding to these data sets, the results provided by the performance monitoring of the standard image processing chain are also analysed for the same purpose. Based on all these indicators, updated calibration data sets may be built, and may be selected as the current active calibration data set. Thus, any variations in sensor performance over time are accounted for on a regular basis.

A key aspect of these campaigns is to use the stars as the reference standards. Their known magnitude and locations provide an essential way to ensure the final quality of the calibration data. During the commissioning phase of the satellites, this activity results in producing the first calibration data set of each instrument. Beyond the first calibration data set, an important benefit of the commissioning campaign, derived from the CONOPS, is to provide radiometric comparisons of many distinct measurements of the very same stars. First, some stars are observed at multiple locations on the image sensor plane of each satellite, enabling estimation of the single instrument repeatability. Then, since the same stars are also present in the images taken by different satellites of the constellation, the radiometric repeatability of the whole mission is characterised. We can see in Fig.6 an example of bright stars paths collected over hundreds of images, and used to derive the first set of opto-electric gain map obtained during our commissioning phase, exposed in Fig. 7.

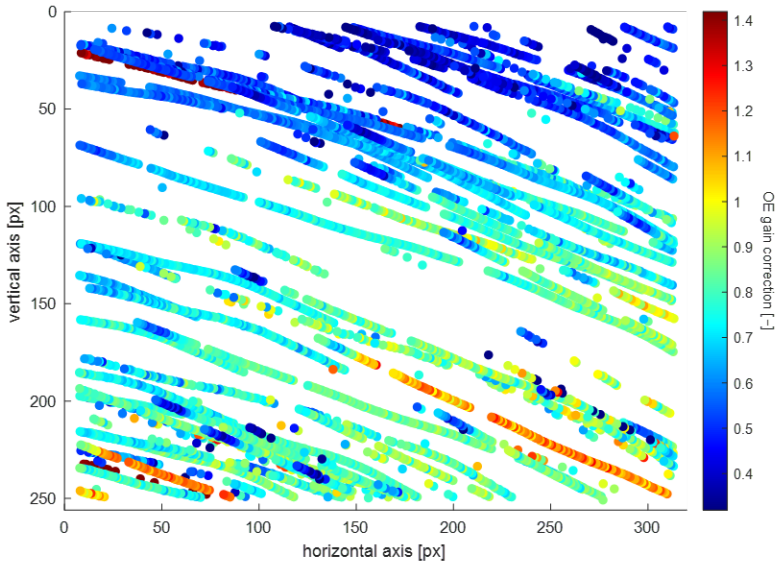

Figure 6: Example of bright stars trajectories in image plane constructed from multiple consecutive frames. The same stars can then be used for multiple locations on the focal plane array.

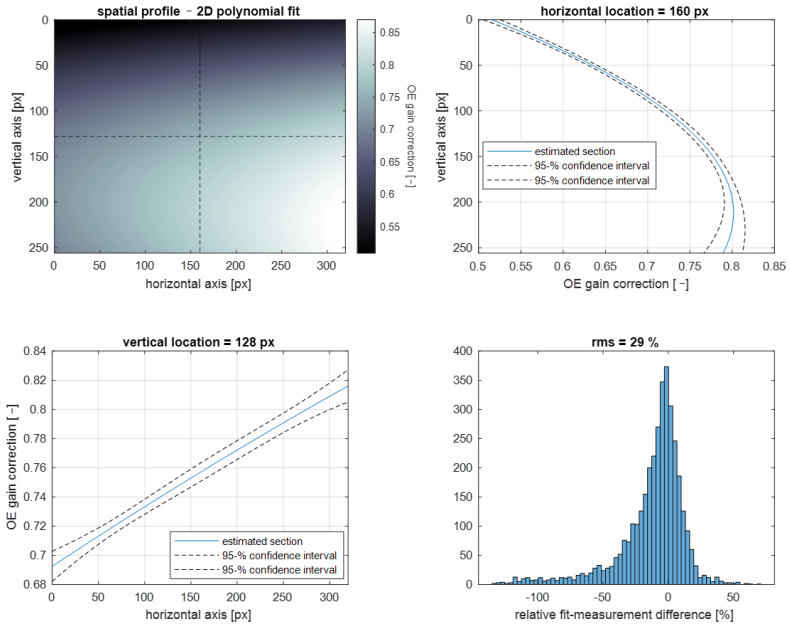

Figure 7: First set of radiometric calibration maps. The absolute opto-electric gain map of the instrument (top left graph) is obtained by comparing, for each star, the expected photonic irradiance, in ph/(s·m²), and the measured one from the calibrated images. The top right and bottom left graphs present a slice of the opto-electric gain correction as a function of each axis, vertical and horizontal respectively. The bottom right graph shows the relative fit-measurement difference (in %). 3644 stars over 325 images were used to produce this first set.

While the main role of the standard processing of the many images produced by the constellation is to detect space objects, it also accumulates different performance indicators that are very powerful measures of the performances of the calibration. Such continuous monitoring of the calibration adds to the calibration quality, and helps account for slow temporal variation in the instruments. Therefore, dedicated campaigns are performed to validate the calibration at regular intervals. Any trend that is estimated during such activity is corrected and an updated calibration data set results. This whole process ensures repeatable calibration performances and provides a unique continuous quality control on every images.

5. IMAGE PROCESSING

5.1 Overview and Product Description

Image processing begins with raw images as the input and yields untagged observations as the output. The process is composed of stages permitting the detection and measurement of known and unknown objects. Known objects are given a NORAD ID before exiting the pipeline. “Known objects” refers to those we search for specifically, and while “unknown objects” are those that show up on image at time t , an are detected by finding unexpected objects which do not resemble stars.

The process follows these broad categories:

1. Calibration and Pre-Processing (Covered in Previous section of this paper)
2. Detection
3. Measurement
4. Association

5. Tagging

The concept of operations enables the extraction of unique product information from the data. In addition to positional measurements such as right ascension and declination, along with their corresponding uncertainties, various other parameters can be derived from the image data. These include, but are not limited to, instantaneous angular velocity inferred from streak length, the orientation of movement within the image plane determined by streak orientation, highly accurate absolute irradiance of the RSO calibrated using stars within the FOV and the absence of atmospheric interference. Moreover, the methodology allows for unique calibration quality control in each image through the stars present in the image plane. A robust identification of the RSO can be achieved by cross-referencing all the above information with time-stamped data from multiple consecutive measurements. The integration of this information supports powerful anomaly detection on multiple parameters within the observation domain, as the CONOPS, deployed through a constellation of satellites, provides numerous observation points to precisely monitor all orbital regimes.

5.2 Pre-processing and Artifact Removal Overview

Calibration is an important part of this process; it removes the optical and sensor imprint on the images (see previous section for details). Noise reduction techniques include an implementation of the background removal model detailed by Lévesque [5] which removes spatially varying background by use of iterations and local statistics across the image. To deal with saturated pixels, we have two categories: 1. Hot pixels (spikes) and, 2. Saturated stars (areas of saturated pixels, see [4] for examples of these stars). During the processing, hot pixels are directly ignored as one pixel isolated on the detector does not affect the processing chain. Saturated stars are either used or ignored depending on the number of saturated pixels contained within. If the number of saturated pixels exceeds a predefined threshold of the total pixels, the entire region of the star is removed and ignored. Below this predefined threshold, the star remains in the processing chain, with the saturated pixels being ignored. Removal of stars is the final step which occurs in the processing chain. Once the boresight of the spacecraft is computed using a star pattern matching algorithm provided by Lang et al. [6], we then remove the stars. Calibration data is used here to ensure that stars are not removed if they're too faint to show up on the detector.

5.3 Detection: Propagation, Event Finder and Matched Filter Construction

The NS CONOPS permits detection in the form of streaks (lines). Statistically, these streaks can be imagined as correlated signals. It's much more likely for a streak to be an object crossing the sensor FOV if it's longer, as it's unlikely for so many pixels to be giving off signal in a straight line by random chance. The matched filter technique given by Levesque [4] provides details on the mathematical implementation of this function, permitting us to make detections much closer to the noise levels, as the match filter technique locates dim correlated signal. In our CONOPS, this process enables both the detection of stars and RSOs in this manner, meaning we can neatly separate the two. For the match filter to function, a known shape is required. This is provided by propagating both our own spacecraft and the catalog of objects outlined in theory by Levesque [2]. The system as we've created it runs this feed-forward loop with complete autonomy (to be discussed in a later section). The propagation, coupled with the event finder, delivers estimated state vectors to the image processing pipeline. This information calculates how the object will appear on the detector by transforming these state vectors on to the image plane. This provides a length in pixels and an orientation. A match filter is applied for each unique pair of length and orientation. GEO spacecraft tend to have identical morphologies, thus a single match filter is required. Fig. 8 below provides visuals of how well the pre-processing and match filter work to extract observations.

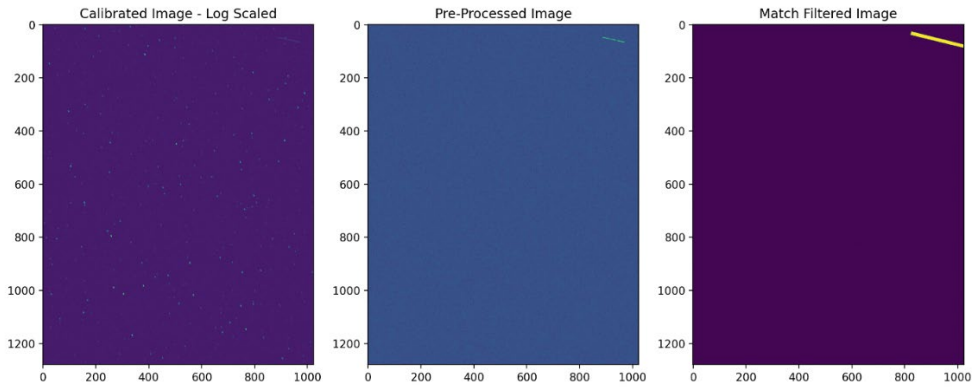

Figure 8: Simulated Images for the purpose of demonstrating. The first on the left is raw image with the calibration data applied. The center image is the result of pre-processing. The key feature here being the stars are removed and the object of interest visible in the upper right corner. The final image on the right shows the zone we identified as the object of interest.

5.4 Measurement

The measurement step yields the precise center-point and morphology of each star and RSO. These precise measurements can be filtered based on the SNR, which is set in the tool. The stars are measured with a least-square curve fitting utility as they are used in the monitoring of calibration data. The RSOs are measured in the same way, allowing, producing precise values for its center-point and shape. These measurements allow the use of differential photometry. Comparing the integrated star signal against the match star catalog stars yields a fitted measure of the RSO integrated signal, which is then converted into a magnitude. Converting the pixel coordinates into Right Ascension (RA) and Declination (Dec) takes place using the results from Astrometry.net [6] astrometric calibration.

5.5 The Association and Tagging of RSO's

Once a collection of observations is established and stored in a database, the objective is to join these observations together. The association process joins individual observations, together which forms a track. The tagging process is applied to tracks, matching them to the NORAD ID. The NORAD ID is not assigned during the match filter process, despite searching for a specific object. This reduces false positives. It's still possible that some artefacts and noise remain in the image The association step is the ultimate reducer of false positives, and it ensures that observations share the same trajectory. Fig. 9 shows the trajectory of an RSO from three Spire images. The purpose of association is to relate these observations into a track. This is accomplished by using both the instantaneous velocity of the streak and the average velocity of a collection. When a streak is long enough (> 7 pixels) in length, its length and orientation are used to project forward in time on the image. Due to the actual system performance of the current on orbit sensors differing from the specified performance, some observations do not have enough length to derive an instantaneous velocity. This case requires the implementation of a Kalman filter to track its average velocity. Fig. 9 shows the motion of an RSO across the detector.

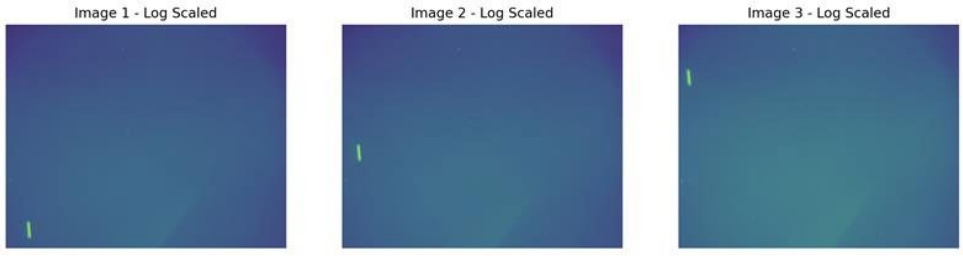

Figure 9: The motion of an object across the detector.

The second key feature of the association is ensuring that the best global solution to the association is found. It's possible to have sets of parallel tracks, meaning similar lengths and orientations, with different trajectories. Thus, the global nearest neighbor's algorithm must be combined with the predict-update steps of the Kalman filter. The following chart [Fig. 10] illustrates the flow of this process.

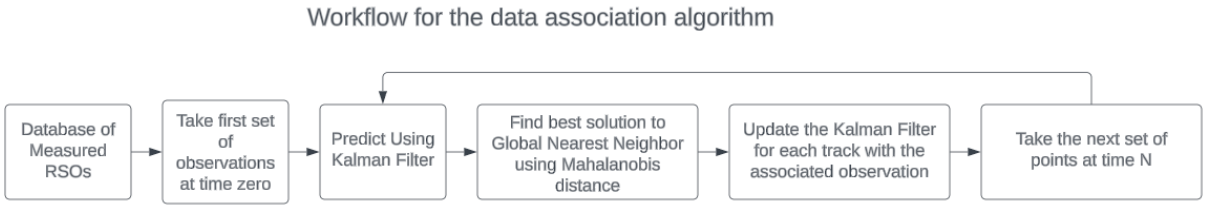

Fig. 10: Illustrates how the data association algorithm joins consecutive observations together, taking observations and forming a track.

5.6 Automation

Autonomous execution from end to end is key, given the throughput of data in the tens to hundreds of thousands of images per day. Therefore, the following systems are implemented to monitor the system performance:

1. Monitoring of the calibration data:
    a. Using the fitted stars to validate when radiometric calibration needs updating
    b. Measuring the PSF of fitted stars to update PSF as a function of position on detector if it's measured as changing over time
    c. Monitoring the dead and hot pixels, tracking their count over time and monitoring the detrimental effects of being in space for prolonged periods.
2. Monitoring statistics on the output data:
    a. The number of observations the pipeline is producing as a function of time,
    b. How many observations are associated to NORAD ID's as a function of time and what the confidence levels for the association and tagging is.
3. Validation of our measures versus those from calibration satellites:
    a. When a calibration satellite ("calsat") passes through the FOV, a dedicated processing thread validates our measures against the calsat to consistently check the accuracy of the system.

The CONOPS helps the automation as many validation steps are completed without specifically executing them. With the wide FOV and a search across all regimes of orbits, calsats regularly pass in front of detectors without needing to task them. Sequential imaging permits automation indirectly, and the association and tagging confidences increase as the number of images increases. Artefact-based false positives are also reduced, since they are unique to a single image. By comparing temporally adjacent images, many false positives are removed.

6. ORBIT DETERMINATION PROCESSING

6.1 OD Overview

Orbit determination involves estimating the state vector of a specific object at a given epoch by processing a set of observations data related to that object. The observations generally are obtained from the image processing pipeline described in the previous sections. NorthStar's current automated orbit determination pipeline is designed to process two distinct types of observational data—radar data (range and range rate) and electro-optical data, to generate a unified state vector. Prior to data ingestion, these varied data types undergo a standardization process, facilitating effective fusion and analysis. While the system is designed for full automation, it retains the flexibility for the manual

intervention of an analyst if required, allowing for adjustments and fine-tuning to address specific scenarios or anomalies. The pipeline triggers as soon as new observations come in.

Because of NorthStar's unique CONOPS, that data ingested from the constellation will heavily involve sets of observations that occur in a very short period of time. This is not uncommon in space-based observations, and is even more pronounced with NorthStar's pushbroom style pointing profiles. Thus, "short arc" and Very Short Arc (VSA) algorithm processing is necessary. The advantage to short arc processing is that information about the orbit can be derived, sufficient enough to produce ephemerides for some amount of propagation, even when there is insufficient data for a standard least squares approach to converge. However, VSA is challenging, and does not result in a full orbit the same way the standard approach can. A particular advantage of NorthStar's CONOPS, though, is that these types of observations will occur much more frequently because of the high revisit rates for most objects (once the full constellation is deployed). The results from different VSA timesteps can then be combined by correlating candidates from one timestep with those of another to generate an estimated orbit.

The first check in the pipeline is to make sure if the observations are tagged to a catalogued object. If the object is not catalogued, an Initial Orbit Determination algorithm (IOD) is triggered. "Initial orbit determination" describes the process of estimating an orbit or state based solely on observations (such as right ascension and declination) and without any a priori orbit information. Presently, we have two different methods implemented. The N-Gooding method, and an IOD using admissible regions. The N-Gooding method [7] is used solely when we have more than 4 observations in a single track. This is often the case for objects in GEO, and in staring mode operations. The N-Gooding is an extension of the classical Gooding method to more than 3 observations. It is purported to be more robust to working efficiently when the initial guess is in the range of 0. The second method addresses scenarios involving very short arcs, as explained above. Short arcs are defined as lasting 2 minutes or less in geostationary orbits and 5-10 seconds in LEO. Very short arcs are frequently encountered in space-based observations. For Electro-Optical (EO) data, the streaklet in the image provides angle and angle rate information, but lacks 2 degrees of freedom, making it difficult to extract usable curvature information due to the short arc length and inherent noise. To address this, admissible regions are employed to constrain the search space for the remaining degrees of freedom [8, 9]. Possible candidates of the admissible region from one VSA timestep are then correlated with candidates from another VSA timestep using either the Initial Value Problem (IVP) or the Boundary Value Problem (BVP). This process results in two pairs of candidates, each with 8 elements, which can be used to construct an orbit with 6 degrees of freedom.

Once a preliminary state vector is established, a batch processing algorithm is initiated to incorporate all available observations for the object. During this phase, the object is classified as a candidate and its state is refined. After this refinement, the object is monitored until it becomes observable again, at which point an Extended Kalman Filter is applied for further state estimation. The sequence of events is show in the Fig. 11 below. After completing the orbit determination processes, notifications are issued based on the detected events. If an anomaly or maneuver is identified, a notification is sent for further analysis. Additionally, a stable state vector is provided for use in other applications, such as conjunction analyses.

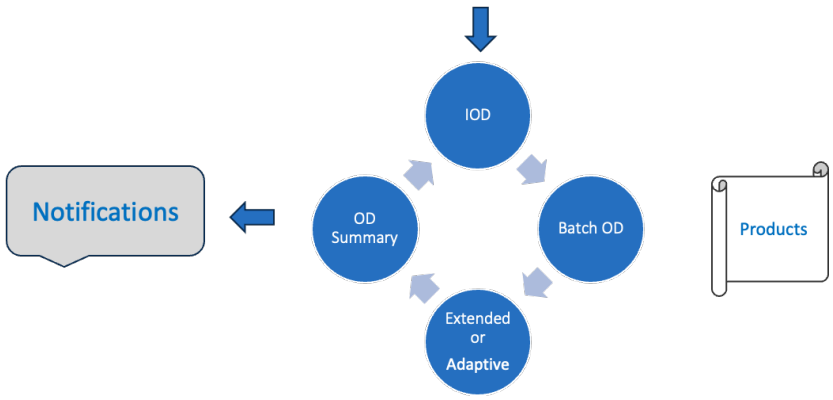

Figure 11: Orbit Determination Automated Pipeline

The block diagram of NorthStar's orbit determination algorithm is presented in Fig. 12 below. The objective is to process the observations received from image processing and estimate an object's trajectory by minimizing the residuals with the computed observables based on dynamic models. The measurement models hold the key part to make sure the computed observables are closely related to the real world. This tool is developed in the Java/Python environment.

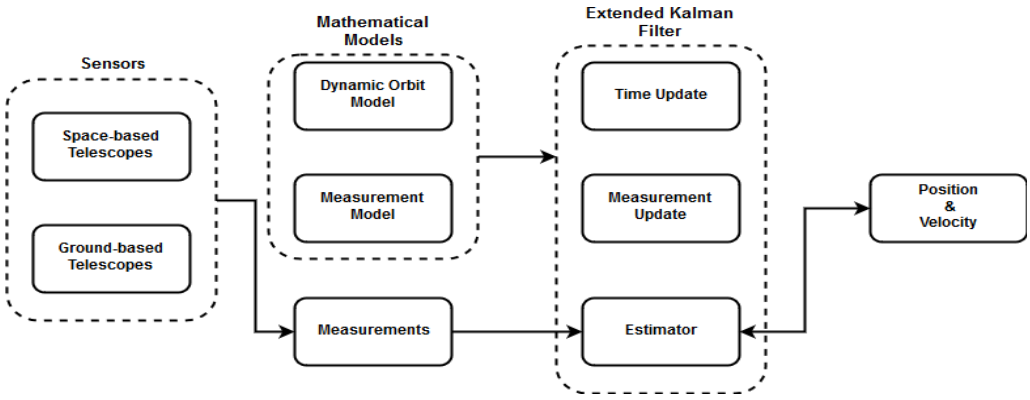

Figure 12: Orbit Determination Block Diagram

6.2 Preliminary Results

Following the description of the orbit determination sequence and algorithm, this section presents preliminary results from data collected on a LEO object using one of our satellites. Images were captured on May 14, 2024, by one of the sensor and processed as detailed in the previous section. In total, four observations were obtained. After processing the images, the orbit determination pipeline was activated, and preliminary analyses were conducted. The state vector generated by NorthStar is initially compared to the Space-track TLE for validation.

Table 2: OD result vs. Space-track TLE comparison

NoradID: 6207

	Initial Spacetrack	Final NorthStar	Delta
UTC Time	2024-05-13T21:48:04.236768Z	2024-05-13T21:48:09.144288Z	
Period	104.6763 [min]	104.6622 [min]	-0.8466 [sec]
Longitude [deg]	-9.8912	-9.4524	0.4387
Long drift rate [deg/day]	4578.8876	4579.5535	0.6659
SMA [km]	7353.4575	7352.6699	-0.7876
Perigee [km]	7338.7373	7337.7364	-1.0009
Apogee [km]	7368.1778	7367.6035	-0.5743
Inclination [deg]	65.8067	65.8212	0.0145
Eccentricity	0.0029	0.0029	0.0000
RAAN [deg]	159.2140	159.2168	0.0028
Argument of perigee [deg]	312.8648	309.5812	-3.2836

While the accuracy levels shown above do not meet expectations, this is to be expected given the limited number of observations and the degraded nature of the observations. It does demonstrate the proper operation of the OD pipeline, even under less-than-ideal conditions.

7. ISSUES & MITIGATION ON THE DEPLOYED FIRST BLOCK

As mentioned in the abstract, the deployed system is suffering from multiple issues that originate from the spacecraft AIT (Assemble, Integration and Test) process. Specifically, all sensors in the first block of satellites suffer from the presence of significant contaminants on the optical surfaces and defocusing caused by misaligned optics, among other issues. These lead to serious degradation of the mission as the SNR is several order of magnitudes lower than anticipated caused primarily by:

1. The amount of straylight in the images, far higher than anticipated
2. The defocus of the system that spreads the signal in a region far larger than designed

To mitigate these unexpected issues, analyses were performed leading to adjustment on the operations, but also on the processing chain. The following sections present the multiple data collection and analyses performed during the commissioning period, as well as the adjustments needed both on the operation and the processing chain.

7.1. **CONOPS Adjustments**

7.1.1 Investigation and qualitative image analyses

Quickly after the first images became available, it had been obvious that something was unexpected. The images captured with the intended integration time were all saturated. A campaign of capture at lower integration time was needed to evaluate whether the sensor could see anything at all. Examples of these captures are presented below.

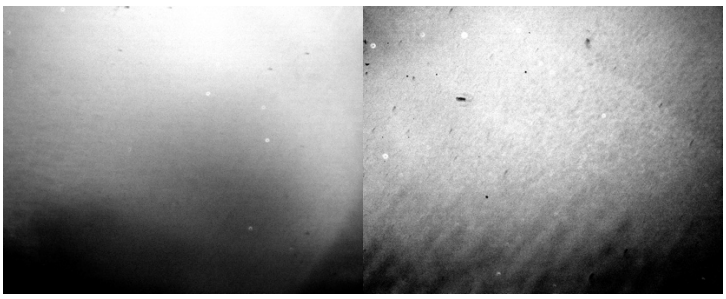

Figure 13: Sample images of sensors showing their non-uniformity due to contamination

These samples show different levels of optical contamination for different spacecrafts. They also show the defocus issue where the PSF appears to have an 85% ensquared energy contained in at least 20x20 pixels versus 3x3 designed. Despite the aforementioned issues, these samples at least show that the sensors can see celestial objects without saturating.

Following this, a campaign to assess the straylight throughout multiple orbits in the Northstar LOF-fixed attitude was commanded. Examples of the data collected are presented below.

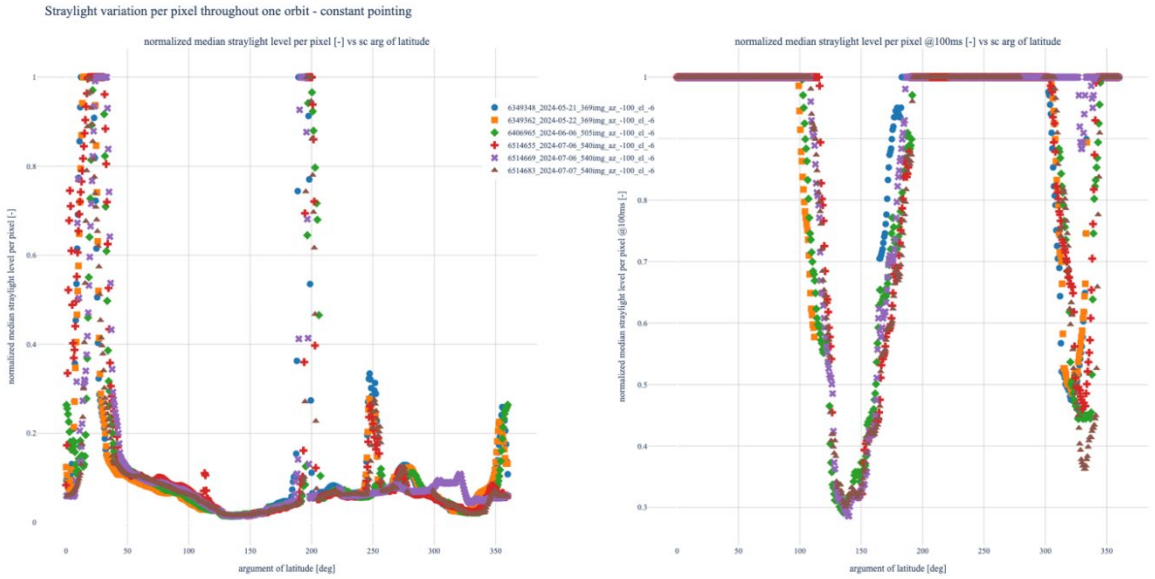

*Figure 14: Straylight variation with the same sensor, with the same pointing, over an orbit at different moments from 2024/05/21 to 2024/07/07. The figure on the left shows the **measured points**, while the one the right is the **extrapolated result** at an integration time meant to optimize star SNR. The raw campaign was performed at 5ms integration time – 200 times lower than the intended for the mission. The data has been normalized to 1 representing the maximum that can be measured (saturation). This shows relatively consistent behaviour with variation that could be explained by a different factor: Earth's longitudinal variation under the spacecraft, pointing imperfections or even a slight difference in the orbital plane. **These results shows that there are at least portions of the orbit (130 deg and 320deg region) where the sensor is not saturating while maximizing SNR of the stars that would be visible at these levels of unwanted light.***

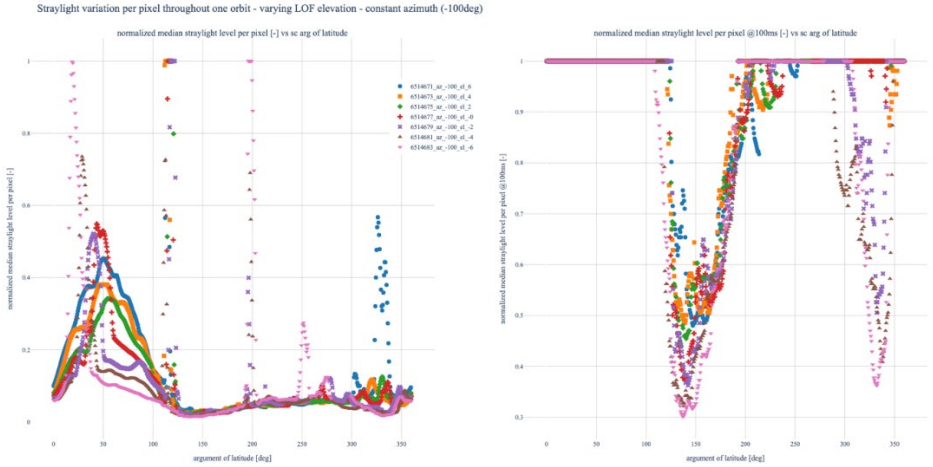

*Figure 15: Straylight variation with the same sensor, varying only the LOF (Local Orbital Frame) elevation angle, the LOF azimuth angle being set at -100 degrees. The figure on the left shows the **measured points**, while the one the right is the **extrapolated result** at an integration time meant to optimize star SNR. This figure shows that similar portions of the orbits can be used across multiple LOF elevation angle for a provided azimuth angle. However the minimum straylight values is achieved in the position where the sensor points the furthest away from Earth (lower pointing pink triangle). However, the furthest data collection (not presented) shows an increase of straylight past that point.*

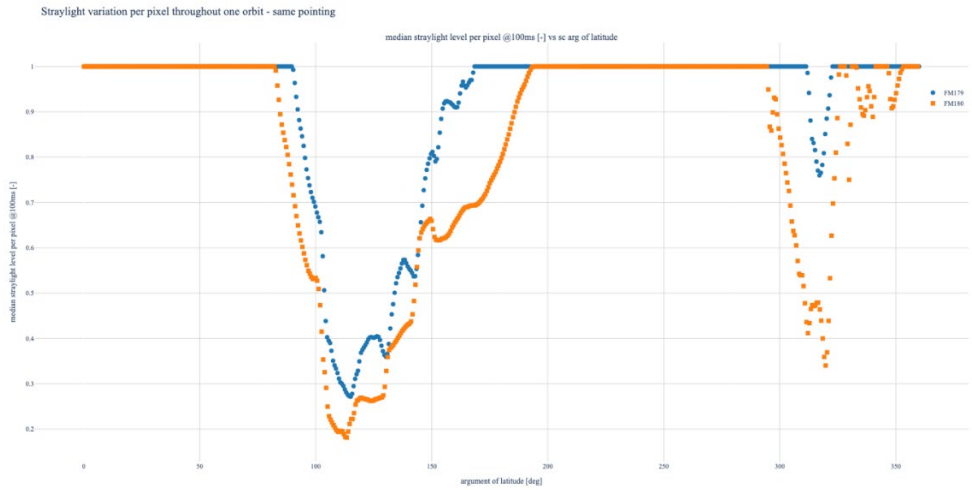

Figure 16: The figure above shows consistency of the straylight variation across an orbit for 2 different spacecrafts on the same orbit that underwent the same AIT process. Many more data collection showed similar trends for both spacecraft. For ease of reading, only 1 orbit for both is presented. The slight improvement for one spacecraft versus the other is most likely explained by different levels of contaminants.

Besides the qualitative analyses, modeling of the payload in details and simulations are on-going to better understand the impact of the contaminants and how they affect the imaging frame. In the following, we present preliminary results of a simulation where effects of the contaminants are modelled by adding a diffuse factor on all optical surfaces of the payload with the baffle having an emissivity of 0.5, as an example of the simulation output.

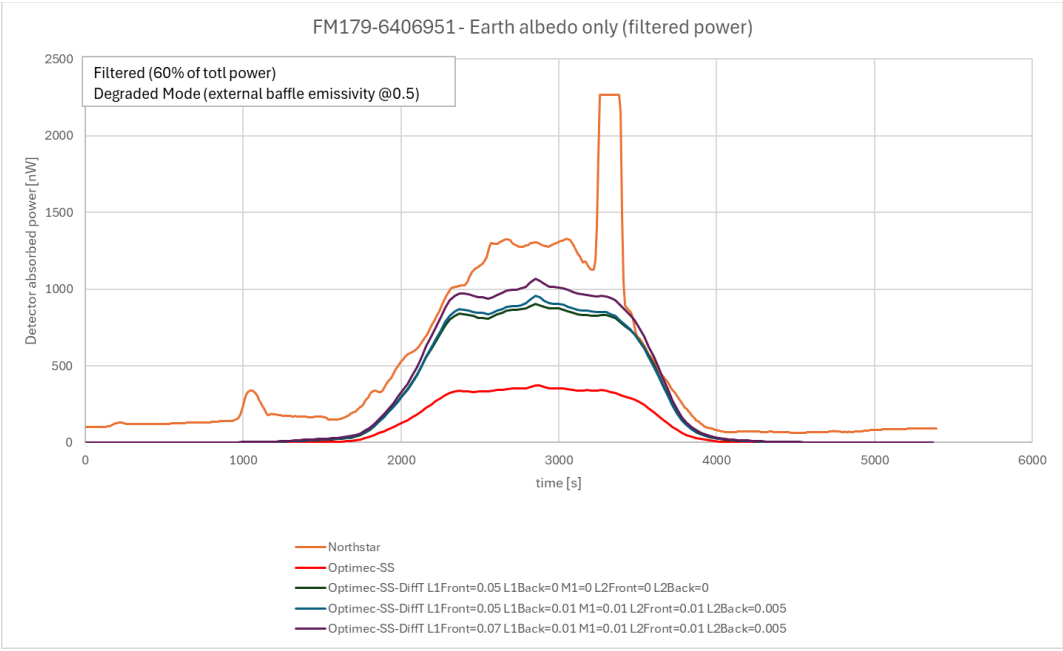

Figure 17: Results showing the simulated power (nW) absorbed by the detector with different diffuse factors using the THERMICA simulation software with only an Earth reflection model accounting for its albedo as source of straylight. The orange line (highest) shows the measured value throughout an orbit. The red line represents the baseline without diffusion effect. These results, consistent with the

data, would tend to point toward a bigger impact of the contaminants on the primary mirror. Additional diffusion on the other surfaces do not seem to change the amplitude of the curve. The spikes on the real data that are not modelled by the simulation are being studied at the time of this paper. Credit: Thomas Garnier, Optimec

Decontamination attempts are also in progress.

While these studies are in progress, the straylight analyses show usability of the sensors in 2 sections of the spacecraft orbits, where:

- The pointing has been selected to minimize the straylight in the current conditions in the Northstar LOF-fixed attitude
- The integration time has been selected to maximize the star SNR in that attitude
- The capture rate is increased to better focus in these regions

However, they come with some operational challenges as the main region of low straylight also coincides with most ground station hits. A ground data pass impacts imaging operations. Despite this, operations have gone relatively well with these adjustments and as shown in the results obtained in the next section.

7.1.2 Preliminary estimation of system performance

The system was designed to detect down to at least stars of apparent magnitude 16, with a more focused system and less straylight impact. A simulation with such content is presented below with a 1 second integration time.

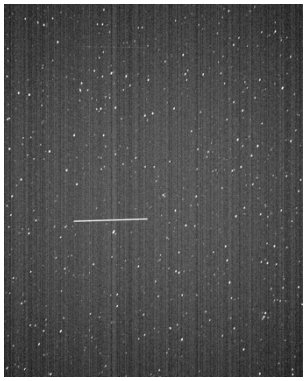

Figure 18: Simulated image of Vm 16 object with nominal camera performance

The deployed system is successful at making detections, but did not meet this specification, even with varied integration times.

7.2 Image Processing Adjustment

7.2.1 Mitigating highly contaminated images

On the processing side, to handle the defocus effect, only minor changes were introduced to keep using simple analytical models of the PSF. The great number of contaminants led to some adjustments. Besides contributing in increasing straylight, they also introduced high spatial frequency artefacts that eventually could be characterized. These artefacts can lead to additional false detections in an automatic system if not handled properly. However, due to NorthStar's LOS-fixed attitude, every light source moves in the frame. That characteristic enables isolating these artefacts using a temporal filter, assuming there are enough successive frames.

Figure 19: Temporal filtering of sequential frames of actual image data showing the result of the background contaminated estimation. The middle image is the estimated contaminated background while the last one is the preprocessed image free of most contamination, containing a RSO (5679 - Thorad Agena D R/B).

7.2.2 Dealing with discontinuous streaks

In the original design, the sensors were meant to capture images with virtually no gaps, except the ones introduced by the electronics of the sensor and attitude change periods. This meant that over several images, an RSO would appear almost as a continuous line. In the adjustment of the current images, this does not hold true, as the combination of the integration time (constrained by straylight) and frame rate (constrained by the bus) leads to more gaps in the streak than nominal RSO signal collection would. The streaklets formed by the RSOs are also much shorter due to the decrease of the integration time, with some RSOs appearing as points, despite their relative velocity. Thus additional robustness in the observation-to-observation association process (streaklet association) was introduced to better handle the problem, through the introduction of Kalman filters to correct the trajectory of the object in the image frame. Such a sequence is displayed in Fig. 20 below, showing the track for an RSO (5679 - Thorad Agena D R/B), crossing the field of view at the vertical line, around the middle of the frame, forming streaklets with gaps between each frame.

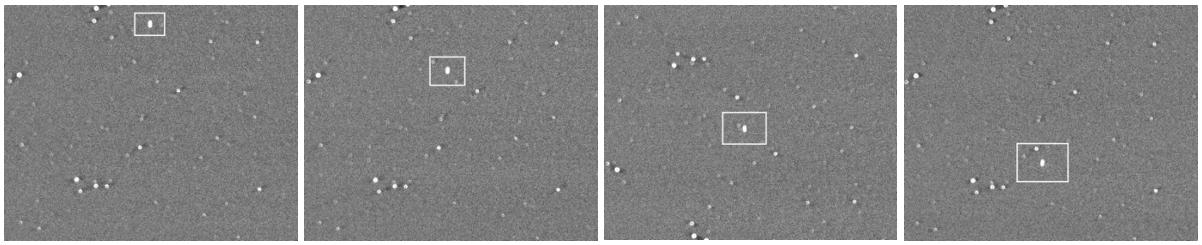

Figure 20: Example of discontinuous streak images

8. CONCLUSION

NorthStar's CONOPS, driven by optical image data supplied by a unique constellation of optical sensors in space, operated in a pushbroom pointing profile, has fully demonstrated the ability to ingest the images, extract observations of RSO's by processing the streaks, and perform orbit determination to produce state vectors in an automated manner. This has been accomplished despite the fact that the initial deployment of four on-orbit sensors suffered severe performance issues that greatly limited the number of observations and degraded the quality of the data relative to specified performance levels. While not yet nearly approaching the product output levels originally intended at this point, NorthStar is delivering its unique data products to customers, will continues to do so. The CONOPS and end-to-end data processing pipeline have been successfully validated. Plans for the replacement of the current on-orbit sensors and for the rollout of the remainder of the constellation are proceeding as rapidly as possible, in order to accelerate the achievement of NorthStar's full performance specifications. As a business, NorthStar is dedicated to continuous improvement, so updates to the planned space infrastructure deployments will always consider technology adjustments or updates, if they will provide added value to product offerings for customers.

9. REFERENCES

[1] Pierre Tremblay, Louis Belhumeur, Martin Chamberland, André Villemaire, Patrick Dubois, Frédérick Marcotte, Charles Belzile, Vincent Farley, Philippe Lagueux, "Pixel-wise real-time advanced calibration method for thermal infrared cameras," Proc. SPIE 7662, Infrared Imaging Systems: Design, Analysis, Modeling, and Testing XXI, 766212 (22 April 2010).

[2] Levesque, M. (2009, September). Automatic reacquisition of satellite positions by detecting their expected streaks in astronomical images. In *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference* (p. E81).

[3] Lévesque, M. P. (2011, September). Detection of artificial satellites in images acquired in track rate mode. In *Proc. AMOS-Tech. Conf., Wailea, Maui, Hawaii, 13–16 September 2011 E* (Vol. 66).

[4] Levesque, M. P., & Buteau, S. (2007). Image processing technique for automatic detection of satellite streaks. *Defense Research and Development Canada Valcartier (Quebec)*.

[5] Lévesque, M. P., & Lelièvre, M. (2007). Evaluation of the iterative method for image background removal in astronomical images. *DRDC Valcartier TN*, *344*.

[6] Lang, D., Hogg, D. W., Mierle, K., Blanton, M., & Roweis, S. (2010). Astrometry. net: Blind astrometric calibration of arbitrary astronomical images. *The astronomical journal*, *139*(5), 1782.

[7] Karimi, R. R., & Mortari, D. (2011). Initial orbit determination using multiple observations. *Celestial Mechanics and Dynamical Astronomy*, *109*, 167-180.

[8] Tommei, Giacomo, Andrea Milani, and Alessandro Rossi. "Orbit determination of space debris: admissible regions." *Celestial Mechanics and Dynamical Astronomy* 97 (2007): 289-304.

[9] Pirovano, Laura. Cataloguing space debris: Methods for optical data association. *Diss. University of Surrey,* 2020.