

NGSatSentry: On-Orbit Detection System for Space Domain Awareness

Nicholas P. Bertrand, Derrick Cheung, Joy Gu, Chris McNamara,
Camille Saidnawey*, Ernie Zenker
Northrop Grumman Corporation

ABSTRACT

A key component of space domain awareness (SDA) involves the detection, tracking, cataloging, and characterization of near-Earth space objects. In addition to general purpose monitoring of various orbits, SDA is crucial for increasing the resiliency of space assets against both accidental and adversarial collisions. While many conventional SDA tracking systems are designed for ground-based operation, there are clear advantages to on-orbit tracking payloads such as improved imaging capabilities due to the absence of atmospheric effects, diversity of vantage point, and improved reaction time. However, in common deployment scenarios, the SDA sensing system is a secondary component on satellites with assorted primary missions (e.g., communication, weather, radar, etc.). Thus, the SDA hardware resides as an auxiliary payload which shares residual resources with the host spacecraft, creating stringent constraints on size, weight, power (SWaP), and limits on downlink utilization. Additionally, in-flight SDA systems require the development of algorithms that account for imaging platform movement and operate efficiently on the resource-constrained flight-capable hardware.

In this work, we introduce NGSatSentry, a low SWaP, hosted SDA payload design. Our discussion focuses on NGSatSentry's suite of on-board space object detection algorithms that adapt ideas from both mature space and ground-based systems and are suitable for deployment in low SWaP flight hardware systems. First, we explore the application of spatial filtering algorithms such as complementary median filtering and subspace projection for suppression of background and stars. Next, we propose a data reduction algorithm that identifies undesired stars in the scene and excludes them from the downlink without the need for an on-board star catalog. We validate our algorithms through experiments on simulated and real ground-based observations of objects in geostationary and other orbits. Finally, we present the results of recent testing on prototype system components in our facility's thermal vacuum chamber and discuss implications on system design.

1. INTRODUCTION

Space-based infrastructure has evolved to become an integral part of a diverse collection of application areas including communications, weather forecasting, navigation, and remote sensing. As our reliance on these technologies continues to grow, space domain awareness (SDA) — the ability to detect, monitor, and characterize resident space objects — provides essential information for mitigating risks to space-based assets such as accidental collisions or adversarial attacks. Although substantial investment in ground-based SDA technologies [1, 2, 3] has resulted in promising advances, there are advantages to space-based SDA systems such as elimination of coverage gaps, improved performance due to atmosphere-free observing conditions, proximity to targets of interest, diversity of observation geometry, and improved reaction time.

One promising path toward a proliferated, space-based SDA sensing network is through the design of a hosted payload with the capability to augment a variety of space vehicles with SDA capabilities. However, the development of an effective hosted SDA payload presents two design challenges: First, the payload must be of low size, weight and power (SWaP) in order to maintain compatibility with a large variety of host vehicles. Second, since communications may be restricted to residual downlink bandwidth, the payload must be able to perform on-board data reduction and tailor its data generation to the available downlink.

*No longer affiliated with Northrop Grumman Corporation.

In this paper, we introduce NGSatSentry, a low-SWaP SDA prototype payload designed to address these challenges. First, we present a suite of on-board algorithms for producing pertinent SDA data from raw, high-resolution, visible-band images. Next, we discuss data reduction approaches that are tunable to match the available downlink allocation. We present a series of simulation results using synthetic and real images of the night sky to evaluate performance and demonstrate potential use cases for the data produced by NGSatSentry. Finally, we present results from recent testing on a commercial off-the-shelf camera system in our facility’s thermal vacuum chamber and discuss implications on system performance.

2. NGSATSENTRY

In this section, we introduce NGSatSentry, a low-SWaP SDA sensor payload intended for use in resource constrained operating scenarios (e.g., in deployment as a hosted payload). Our discussion will focus primarily on the suite of on-board algorithms that NGSatSentry uses to process raw visible-band images and produce pertinent SDA data. The goal of NGSatSentry’s on-board processing is to detect and downlink information about resident space objects (RSOs) in the field of view using imagery from a visible-band focal plane array (FPA). The processing chain is designed to be lightweight such that images can be processed in real-time even on power-constrained on-board processing hardware. The overall processing chain is summarized in Figure 2.1 and we discuss each component in detail in this section.

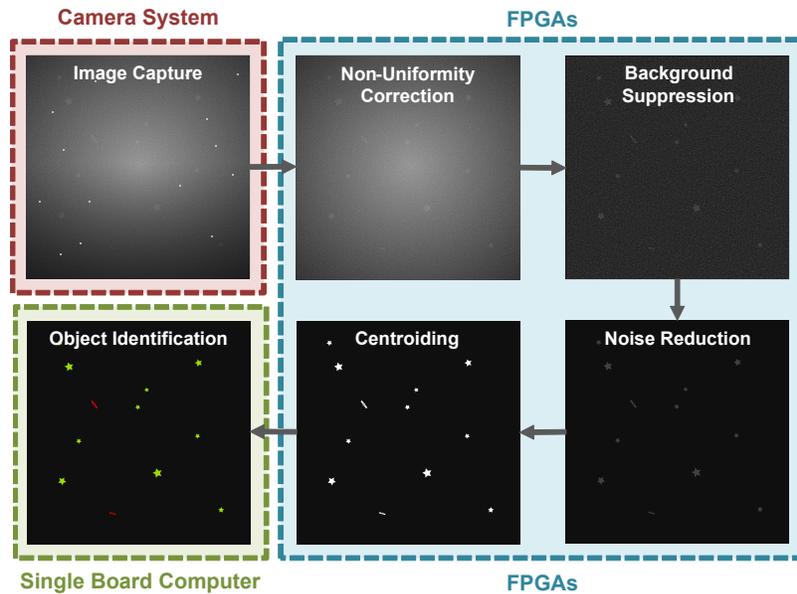


Fig. 2.1: NGSatSentry’s processing pipeline is executed on a combination of a field-programmable gate array (FPGA) and a single board computer (SBC). The FPGA interfaces directly with the camera system to collect raw images and execute image processing algorithms to output a list of exceedance clusters corresponding to objects in the scene. This FPGA implementation of the image processing algorithms is scalable to enable real-time processing of high-resolution images from multiple cameras simultaneously. The SBC then performs data reduction on the resulting exceedance cluster data by identifying objects of interest and excluding other objects such as stars from the downlink.

NGSatSentry’s algorithms have been optimized and implemented on flight-like hardware, proving their feasibility to run in real-time in power constrained on-orbit processing environments. The image processing algorithms, from raw image collection to cluster centroid computation, are implemented using a field-programmable gate array (FPGA). This FPGA implementation is highly scalable and enables real-time processing of high-resolution (e.g., 12 MP) images from multiple cameras simultaneously. The remaining centroid classification and data reduction is performed via an on-board single-board computer (SBC) which also handles communications management and other mission related tasks.

Throughout this section, we use the notation $\mathbf{A} \in \mathbb{R}^{M \times N}$ to denote a matrix with M rows and N columns, and $A[i, j]$ to denote the element with coordinates (i, j) . In each of the image processing stages, we re-use the symbol \mathbf{I} to denote the input image from the previous stage. The notation $\hat{\mathbf{A}}$ is used to denote an estimate of \mathbf{A} , and $\bar{\mathbf{A}} = \frac{1}{N} \sum_{i=1}^N \mathbf{A}_i$ denotes the mean. Finally, we use \odot to indicate element-wise multiplication between matrices or vectors.

2.1 Image Processing and Target Detection

2.1.1 Non-Uniformity Correction

In the first processing stage, raw image data undergoes non-uniformity correction, a standard image processing technique to compensate for variations in response to light across the FPA. While an ideal camera system would produce the same mean intensity in response to a given point source regardless of its position on the focal plane, in practice, this assumption does not hold due to lens design considerations and manufacturing variations in the optics and FPA. We represent the resulting non-uniformities via a linear model

$$\mathbf{I} = \mathbf{G} \odot \mathbf{X} + \mathbf{D}, \quad (2.1)$$

where $\mathbf{X}, \mathbf{G}, \mathbf{D} \in \mathbb{R}^{M \times N}$ denote the true intensity, gain and bias matrices respectively. Provided estimates $\hat{\mathbf{G}}$ and $\hat{\mathbf{D}}$ of the gain and bias matrices, the corrected image $\hat{\mathbf{X}}$ may be produced via the inverse transform

$$\hat{X}[i, j] = \frac{I[i, j] - \hat{D}[i, j]}{\hat{G}[i, j]}. \quad (2.2)$$

Estimation of the gain and bias matrices is performed via a combination of laboratory and on-orbit calibration. The bias matrix is temperature sensitive, and must therefore be updated regularly as thermal conditions change during operation. To perform this update, NGSatSentry shields the front element of the lens to block incoming light, so $\mathbf{X} \approx 0$. Any non-zero detector intensities may then be attributed to imaging noise, so the bias matrix may be estimated by computing the mean over a sequence of images. We estimate the gain matrix by uniformly illuminating the camera and capturing a series of frames $\mathbf{I}_k, k = 1, \dots, K$. The gain for each detector is then computed as the value that normalizes the mean of each detector to the spatial median value:

$$\hat{G}[i, j] = \frac{\text{med}_{ij}(\bar{I}[i, j] - \hat{D}[i, j])}{(\bar{I}[i, j] - \hat{D}[i, j])}, \quad (2.3)$$

where \bar{I} denotes the mean image across the K frames. We use a commercial integrating sphere to ensure uniform illumination of the front lens elements. Fortunately, the gain matrix is relatively robust to changes in temperature and thus does not require frequent updates on-orbit. If further gain calibration is required on-orbit, one may consider using leveraging the known radiance of stars to serve as illumination reference points.

2.1.2 Background Suppression

In the next processing stage, we employ complementary median filtering (CMF) to suppress background illumination which does not share the morphology of typical objects of interest (i.e., streaks or point sources). While this step is especially useful in ground-based test data where light pollution results in illumination gradients, it may also serve to suppress artifacts caused by reflections or stray light due to insufficient solar exclusion in on-orbit data. In this step, we model the image with two additive components: targets of interest denoted by \mathbf{S} and the background denoted by \mathbf{L} . An estimate of the background is produced via median filtering with a kernel \mathcal{K} that is tuned to preserve objects of interest:

$$\hat{\mathbf{L}} = \text{med}_{\mathcal{K}}(\mathbf{I}). \quad (2.4)$$

The background suppressed image is then formed by simply subtracting $\hat{\mathbf{L}}$ from the input image.

Alternatively, the background may be suppressed via subspace projection (SSP) as follows. Suppose the scene is well approximated as $\mathbf{I} = \mathbf{A}\mathbf{W} + \mathbf{B}\mathbf{U}$ where the columns of \mathbf{A} and \mathbf{B} form orthonormal bases

for orthogonal subspaces representing the targets of interest and the background respectively. Provided an estimate of the background subspace basis $\hat{\mathbf{B}}$, we may remove the background via

$$\hat{\mathbf{S}} = \mathbf{I} - \hat{\mathbf{B}}\hat{\mathbf{B}}^T\mathbf{I}. \quad (2.5)$$

Estimation of \mathbf{B} typically requires multiple frames and can be carried out using a variety of methods. One such method for images where the power in the clutter component is larger than that of the target is to simply use the first K singular vectors from a matrix containing a time series of observations, where K is a tunable parameter corresponding to the dimension of the background subspace. While SSP may be advantageous in scenarios with complex clutter environments, we have observed in experiments with ground-based observations that CMF is typically sufficient. Therefore, our results focus mostly on CMF which does not require processing of sequences of multiple frames, making it more suitable for resource constrained processing hardware.

2.1.3 Noise Reduction

Next, matched filtering is employed to increase the signal-to-noise ratio, producing improved detection results in subsequent processing. Here, we model a local region of the image containing a target as

$$\mathbf{I} = \mathbf{S} + \mathbf{N}, \quad (2.6)$$

where \mathbf{S} represents the noiseless target intensities, and \mathbf{N} is a Gaussian distributed random vector. If the target is well-approximated as a point source, then matched filtering with the point response function produces the optimal intensity and location of the target [4]. Furthermore, we have found experimentally using real images of the night sky that this approach improves performance even for non-point source targets such as stars which appear as streaks due to sufficiently long integration durations (see Section 3.2).

2.1.4 Exceedance Cluster Generation

We are now ready to detect objects of interest present in the image. The output of this stage of processing is a collection of *exceedance clusters* $\mathcal{E} = \{E_c\}_{c=1}^C$ corresponding to objects in the scene. Each of the E_c contains a set of tuples $\{(i, j, I[i, j])\}_{p=1}^{P_c}$ indicating the pixel coordinates and intensities for the cluster. Exceedance clusters are generated in three steps: First, we apply a threshold to the input image to produce a binary map defined by

$$M[i, j] = \begin{cases} 1, & \text{if } I[i, j] > \tau \\ 0, & \text{otherwise} \end{cases}, \quad (2.7)$$

where τ is a threshold chosen according to the sensitivity and noise characteristics of the camera. Next, the dilation operator is applied to the resulting binary image to ensure that targets with regions that are close to the threshold value τ remain connected [5]. Finally, the E_c are generated by identifying clusters of connected, non-zero pixels.

2.2 Data Reduction

The remainder of the on-board processing focuses on data reduction and is tunable to the power budget and downlink bandwidth available from the host vehicle. We consider two mechanisms for data reduction: reducing the amount of information per exceedance cluster and reducing total number of exceedance clusters by selectively excluding those that are not of interest from the downlink. The former includes a range of options (ordered from minimal processing/heavy downlink utilization to more processing intensive/downlink friendly): downlink full images; downlink all pixels for each exceedance cluster; downlink only summary information such as cluster centroid coordinates and mean intensity. In scenarios where downlink bandwidth is sufficient, downlinking full pixel intensity information for exceedance clusters may be advantageous, enabling more processing flexibility at the ground segment, such as streak morphology characterization.

Next, we discuss methods for the latter data reduction mechanism: selectively excluding exceedance clusters from the downlink. While a subset of stars may be desired for reference and calibration purposes, most stars

are not targets of interest in the context of SDA. Since stars typically represent the majority of exceedance clusters in a given scene, their rejection from the downlink is a primary mechanism for data reduction. In applications where the payload has access to an on-board star catalog, standard star matching techniques (e.g., [6, 7, 8]) may be used to identify clusters corresponding to stars and exclude them from the downlink. Here we propose an alternative method for scenarios where the payload does not have access to an on-board star catalog. For the purposes of data reduction, we need only distinguish star and non-star exceedance clusters. To accomplish this classification task without access to a star catalog, we consider consecutive frames captured in the same position to estimate cluster velocities and note that star exceedance clusters share approximately the same velocity while non-star clusters tend to differ in velocity.

First, we compute the uncorrected center of mass for the exceedance clusters in each pair of frames via

$$(\tilde{x}_c, \tilde{y}_c) = \frac{1}{P_c} \sum_{(i,j,I[i,j]) \in E_c} I[i,j] \cdot (j, i), \quad (2.8)$$

then apply distortion correction:

$$(x_c, y_c) = (\tilde{x}_c + \Delta x, \tilde{y}_c + \Delta y), \quad (2.9)$$

where $\Delta x, \Delta y$ are distortion correction factors which are discussed in Section 2.3. Next, we convert from image coordinates (i.e., rows and columns) into the corresponding three dimensional line-of-sight vector via

$$\tilde{\ell}_c = \begin{pmatrix} (x_c + \Delta x_c - N_c/2 + 0.5)d \\ (y_c + \Delta y_c - N_r/2 + 0.5)d \\ f \end{pmatrix}, \quad (2.10)$$

where d is the focal plane detector size, f is the lens focal length, and N_c and N_r are the number of columns and rows in the image respectively. Next, we define the normalized line-of-sight vectors $\ell_c = \text{LOS}(x_c, y_c) = \tilde{\ell}_c / \|\tilde{\ell}_c\|_2$. To identify the star exceedance clusters, we consider the line-of-sight vectors $\{\ell_c\}, \{\ell'_c\}$ for two consecutive frames and seek to find the maximum number of cluster correspondences $\mathcal{C} = \{(c_i, c'_i)\}_{i=1}^N$ and a rotation matrix \mathbf{R} such that

$$\|\ell_{c_i} - \mathbf{R}\ell'_{c'_i}\|_2 \leq \eta \quad i = 1, \dots, N, \quad (2.11)$$

where η is a tunable threshold parameter. The problem of determining the correspondence and transformation between two point sets is known as *simultaneous pose and correspondence registration*, and many algorithmic solutions have been developed (e.g., [9, 10]). Here, we propose an adaptation of the well-known Iterative Closest Point algorithm [11] which alternates between i) determining the correspondences for a fixed transformation, and ii) updating the transformation for a fixed set of correspondences. We additionally restrict the correspondences to satisfy (2.11) since not all points are expected to obey the same transformation (i.e., those not associated with stars). The full procedure is described in Algorithm 2.1. On experiments with real data, we terminate the algorithm when the number of correspondences stabilizes, and have observed that the procedure typically converges with fewer than 10 iterations.

Algorithm 2.1: Catalog-Free Star Exceedance Cluster Identification

input : line-of-sight vectors for exceedance clusters from two consecutive frames: $\{\ell_c\}, \{\ell'_{c'}\}$, initial correspondence estimate \mathcal{C} , initial rotation estimate \mathbf{R}

output: index pairs for exceedance clusters that exhibit star-like motion: \mathcal{C}

```
while not converged do
  /* determine correspondences via nearest neighbors */
   $\mathcal{C} \leftarrow \{\}$ ;
  for  $\ell \in \{\ell_c\}$  do
     $\ell^* \leftarrow \operatorname{argmin}_{\ell' \in \{\ell'_{c'}\}} \|\ell - \mathbf{R}\ell'\|_2^2$ ;
    if  $\|\ell - \mathbf{R}\ell^*\|_2 < \eta$  then
      | add indices corresponding to  $\ell$  and  $\ell^*$  to  $\mathcal{C}$ ;
    end
  end
  /* for the registration  $\mathcal{C}$ , determine optimal rotation matrix  $\mathbf{R}$  */
   $\mathbf{L} \leftarrow [\ell_{c_i}], (c_i, c'_i) \in \mathcal{C}$ ; // stack LOS vectors as matrix columns
   $\mathbf{L}' \leftarrow [\ell'_{c'_i}], (c_i, c'_i) \in \mathcal{C}$ ;
   $\mathbf{U}, \mathbf{\Sigma}, \mathbf{V}^T \leftarrow \operatorname{SVD}(\mathbf{L}\mathbf{L}'^T)$ ; // compute singular value decomposition
   $\mathbf{R} \leftarrow \mathbf{U}\mathbf{V}^T$ ;
end
```

2.3 Distortion Correction Profile Calibration

Many SDA related tasks, such as attitude determination, require accurate calculation of line-of-sight vectors to objects in the FOV. Line-of-sight accuracy is impacted by a host of factors, of which two are addressed here: accuracy of the effective focal length estimate and the capability to correct optical distortion. Although the nominal focal length and distortion correction profile may be measured prior to launch, their precise values may differ from the manufacturer specifications and may also change with fluctuations in temperature. In this section, we describe a procedure for updating the distortion correction coefficients and effective focal length during on-orbit operation.

Distortion correction is performed by computing a correction factor in the row and column coordinates of each centroid via an n -degree polynomial

$$\Delta x = \langle \mathbf{f}, \mathbf{p}(x) \rangle, \quad \Delta y = \langle \mathbf{g}, \mathbf{p}(y) \rangle, \quad (2.12)$$

where $\mathbf{p}(x) = [1, x, x^2, x^3, \dots, x^n]$ and \mathbf{f}, \mathbf{g} are the correction coefficients computed via calibration. We can compute the distortion correction coefficients jointly with the focal length via an alternating optimization scheme which seeks to compute values that minimize centroid line-of-sight error compared to a star catalog. The full procedure is described in Algorithm 2.2, and the approach is evaluated on data in Section 3.5.

Algorithm 2.2: Joint Correction Coefficient and Focal Length Estimation

input : star catalog \mathcal{T} , uncorrected centroids $\{(x_c, y_c)\}_{c=1}^C$, initial focal length f and initial correction coefficients \mathbf{f}, \mathbf{g}

output: focal length estimate f , correction coefficients \mathbf{f}, \mathbf{g}

while *not converged* **do**

- $\mathcal{S} \leftarrow \{\}$;
- /* register centroids to stars */*
- for** $c = 1, \dots, C$ **do**

 - $x^* \leftarrow x_c + \langle \mathbf{f}, \mathbf{p}(x_c) \rangle, \quad y^* \leftarrow y_c + \langle \mathbf{g}, \mathbf{p}(y_c) \rangle$; */* compute corrected centroids*
 - $\mathbf{s}^* \leftarrow \operatorname{argmax}_{\mathbf{s} \in \mathcal{T}} \langle \operatorname{LOS}_f(x^*, y^*), \mathbf{s} \rangle$; */* find nearest star*
 - $\theta \leftarrow \cos^{-1}(\langle \operatorname{LOS}_f(x^*, y^*), \mathbf{s} \rangle)$;
 - if** $\theta < \tau_\theta$ **then**

 - | Add centroid, star pair to \mathcal{S} ;

 - end**

- end**
- /* update correction coefficients via least squares */*
- $\mathbf{f}, \mathbf{g} \leftarrow \operatorname{argmin}_{\mathbf{f}, \mathbf{g}} \sum_{i=1}^{|\mathcal{S}|} \|\mathbf{s}_i - (x_i + \langle \mathbf{f}, \mathbf{p}(x_i) \rangle, y_i + \langle \mathbf{g}, \mathbf{p}(y_i) \rangle)\|_2^2$;
- /* update focal length via line search */*
- $f \leftarrow \operatorname{argmin}_f \sum_{i=1}^{|\mathcal{S}|} \cos^{-1} \langle \operatorname{LOS}_f(x_i^*, y_i^*), \mathbf{s} \rangle$

end

3. RESULTS

Next, we discuss experimental results using both simulated and real data to evaluate the efficacy of NGSatSentry’s algorithms, and demonstrate potential use cases for the resulting SDA data. Additionally, we present the results of recent testing of commercial off-the-shelf (COTS) camera systems in our facility’s thermal vacuum chamber, demonstrating the feasibility of low-cost COTS hardware in SDA applications.

3.1 Data Collection and Examples

Images of the night sky were collected using a COTS camera and lens system connected directly to the FPGA. Two camera variants were tested: one with a 3 MP complementary metal-oxide-semiconductor (CMOS) focal plane array (FPA), and one with a similar but larger 12 MP FPA. The lens had a focal length of 50 mm and a 3.6 cm entrance pupil diameter, producing a diagonal field-of-view (FOV) of 10.2° and 20.1° for the 3 MP and 12 MP FPAs respectively. In order to simulate the lighting conditions seen on-orbit as closely as possible using ground-based observations, images were collected at Joshua Tree National Park in California, USA, an area well known for its low levels of light pollution.

An example frame before and after NGSatSentry’s on-board processing is shown in Figure 3.1. NGSatSentry computes centroids for each exceedance cluster, uses these centroids for attitude determination, then registers centroids to stars in the Tycho-2 catalog to identify clusters which do not correspond to stars. To confirm the positions of known (i.e., cataloged) RSOs, we propagate their positions using two-line elements (TLEs) together with the SGP4 orbital model and overlay these predictions on the image. An example region of interest containing four geostationary satellites is shown, demonstrating close agreement between the predicted and observed positions. Another region of interest contains an exceedance cluster with no match in the star or RSO catalogs — such objects would be downlinked for further inspection.

A basic sensitivity analysis is shown in Figure 3.2. We plot a frame-normalized visual magnitude histogram for stars observed over a sequence of images. From this analysis, we may conclude that the sensitivity of our system with this choice of optics is approximately visual magnitude 11.

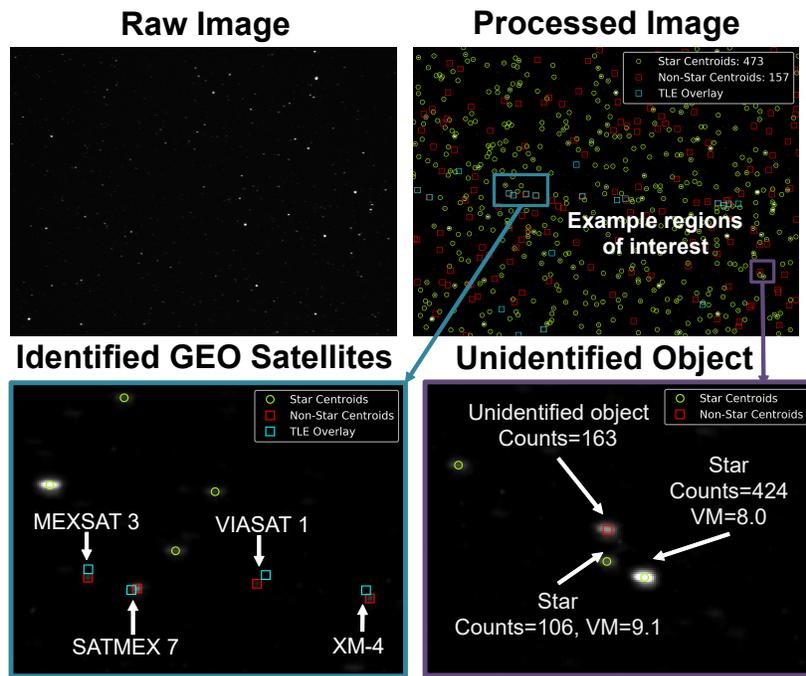


Fig. 3.1: Sample raw and processed images captured in Southern California. (Top-left) To simulate on-orbit viewing conditions of the night sky as closely as possible using ground based measurements, we recorded images in Joshua Tree National Park, a location known for its low light pollution environment. (Top-right) NGSatSentry’s processing performs star matching to identify and reject centroids corresponding to stars. (Bottom-left) By propagating two-line elements from a satellite catalog obtained from `space-track.org`, we see that four of the non-star centroids correspond to a cluster of geostationary satellites. (Bottom-right) Relatively bright unidentified centroid with no corresponding star or TLE object. This is an example of a centroid that would be downlinked for further analysis.

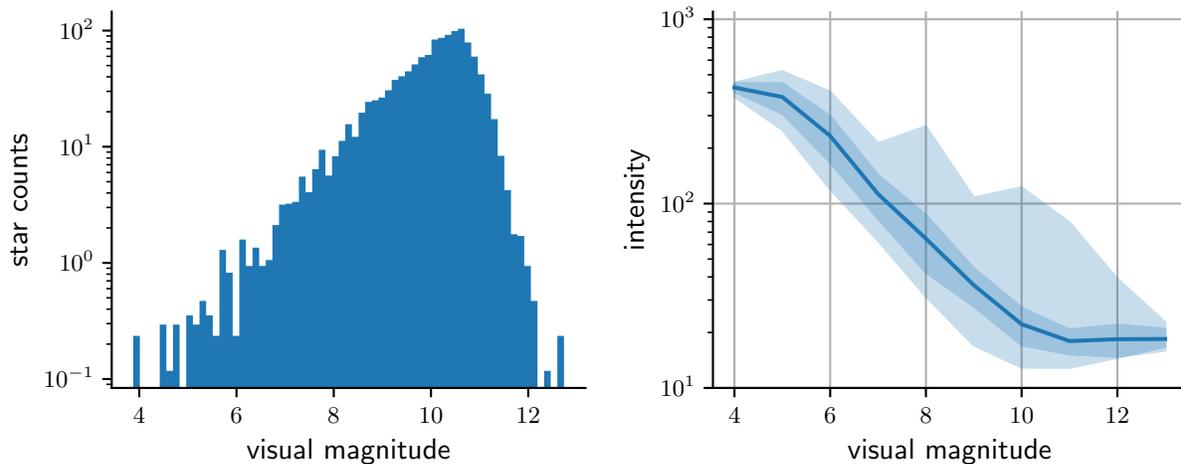


Fig. 3.2: Sensitivity analysis based on real images of the night sky collected in Joshua Tree. (Left) Histogram of star counts as a function of visual magnitude (VM). NGSatSentry is capable of detecting stars up to visual magnitude ~ 12 . (Right) Mean cluster pixel intensity as a function of VM. As expected, the pixel intensity decreases exponentially with increasing VM.

3.2 Component-Level Performance Analysis

The first three blocks in the image processing chain aim to improve detection performance by increasing the signal to noise and clutter ratio. In this section, we evaluate the performance contribution of each block by selectively removing algorithms from the processing chain and comparing performance across various configurations.

The separation of signal from noise is determined by the threshold τ ; pixels that exceed τ are assumed to contain signal and a detection is thus declared. We evaluate performance via the receiver operating characteristic (ROC) curve which characterizes the trade-off between probability of detection and probability of false alarm as a function of τ . Generally, ROC curves show a lower (higher) number of detections and false alarms at a higher (lower) threshold values. The operator has the freedom to adjust threshold that results in the desired trade-off between detections and false alarms. The goal of this analysis is to confirm that our image processing techniques are providing the best trade-off curve (i.e., higher number of detections and lower false alarms for a given threshold value). To this end, we consider ROC curves with each image processing block disabled in turn and compare the resulting performance with that of the full processing chain.

We formulate the detection problem using stars in the Tycho-2 catalog as ground truth. A true positive (i.e., correct detection) corresponds to the event when a centroid is assigned to a nearby star, and a false alarm is declared when an exceedance cluster has no nearby catalog star. We note that in the context of this analysis, RSOs bias the results slightly as they can (correctly) result in exceedance clusters that do not correspond to stars. However, since the stars in a scene vastly outnumber the number of RSOs, the impact of this bias is negligible.

Results for each of the four configurations are shown in Figure 3.3. One noteworthy distinction from traditional ROC curves is that the curve starts with very few detections when the threshold is low. This is a by-product of the clustering process; in the dilation step, nearby clusters are connected together. When the threshold is low, the resulting pixel map is very dense and thus easily connected. In extreme cases, up to 90% of image pixels may be assigned to just a few centroids, resulting in few detections and few false alarms.

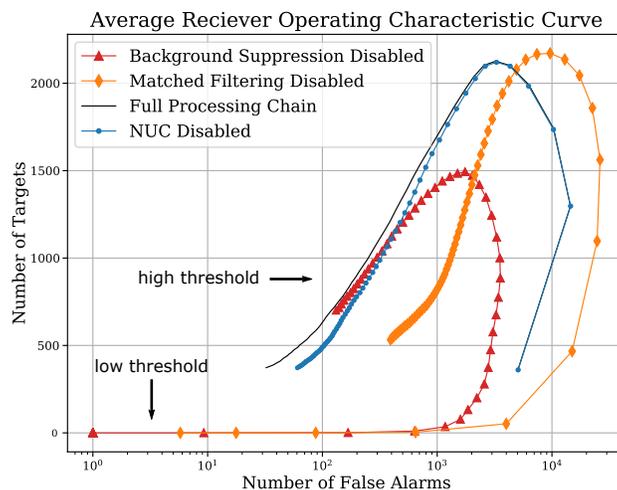


Fig. 3.3: Evaluation of individual processing components. To assess the contributions of each step in the processing chain, we formulate a detection problem using stars in the Tycho-2 catalog as ground truth. Shown are receiver operating characteristic (ROC) curves based on this detection task using data recorded in Joshua Tree, CA. The curves shown represent performance when the specified processing step is removed from the pipeline. Each processing step contributes a measurable performance increase with matched filtering and CMF as the two largest contributors. Note that the morphology of these curves differs from that of traditional ROC curves. This difference arises because the choice of threshold affects the number of exceedance clusters which is effectively the number of observations present in the ROC analysis. In contrast, in the traditional setting, the threshold is applied to a fixed set of estimates.

This analysis suggests that matched filtering and background suppression provide the largest performance increases; the orange curve is further to the right, corresponding to higher false alarms, and the red curve has lower peaks, corresponding to fewer detections. Disabling non-uniformity correction degrades performance to a much lesser degree, though it has higher false alarms at higher threshold values. Furthermore, we expect non-uniformity correction to become more important toward the end of the payload lifecycle as detector performance degrades due to radiation exposure.

3.3 Background Suppression Performance

We now present a comparison of complementary median filtering and subspace projection, the two background suppression methods discussed in Section 2.1. This analysis uses synthetic images to afford control over the relative power in the clutter and target components. Images are generated in three steps: first, a random collection of target locations is generated and the corresponding pixels are set equal to one. Next, we produce the clutter-free target image \mathbf{T} by convolving the resulting image of point sources with a Gaussian point response function. Finally, a radial gradient image \mathbf{B} representing the clutter source is combined additively with the targets to produce the final image:

$$\mathbf{I} = \mathbf{T} + \mathbf{B}. \quad (3.1)$$

We then use CMF and SSP to recover \mathbf{T} given \mathbf{I} and evaluate performance using two metrics: background suppression factor (BSF) and relative mean square error (rMSE). BSF measures the ratio of the power in the background component in the input image to that of the output image and is defined as

$$\text{BSF} = \frac{\sigma_{\text{in}}}{\sigma_{\text{out}}}, \quad (3.2)$$

where σ_{in} and σ_{out} represent background power in the input and output images respectively. Similarly, rMSE is used to measure the relative error between the ground truth target image \mathbf{T} and the estimate produced by each filtering technique $\hat{\mathbf{T}}$, and is defined as

$$\text{rMSE} = \frac{\|\mathbf{T} - \hat{\mathbf{T}}\|_2^2}{\|\mathbf{T}\|_2^2}, \quad (3.3)$$

where $\|\cdot\|_2$ denotes the ℓ_2 -norm.

A comparison of CMF and SSP is shown in Figure 3.4. In terms of BSF, both algorithms perform similarly when the power in the background is small relative to that of the targets. Provided a sufficient number of frames to estimate the background (here we use 10 frames), SSP produces consistent rMSE throughout the range of signal-to-clutter ratio values. The value of the rMSE for SSP is driven by the amount of contamination of the background subspace estimate due to the targets. This quantity may be reduced by using more frames to produce the background subspace estimate. In experiments with real data, we have observed that CMF produces sufficient background suppression performance is thus the default in our algorithm baseline. However, in scenarios with complex backgrounds where i) sufficiently many frames can be captured with backgrounds from the same subspace, and ii) computational resources are sufficient handle these frames, SSP may produce superior results. Finally, we note that the performance of SSP depends on how well the linear model approximates the true signal — in practice, the linear model does not capture some phenomena (e.g., saturation effects) which may reduce SSP performance.

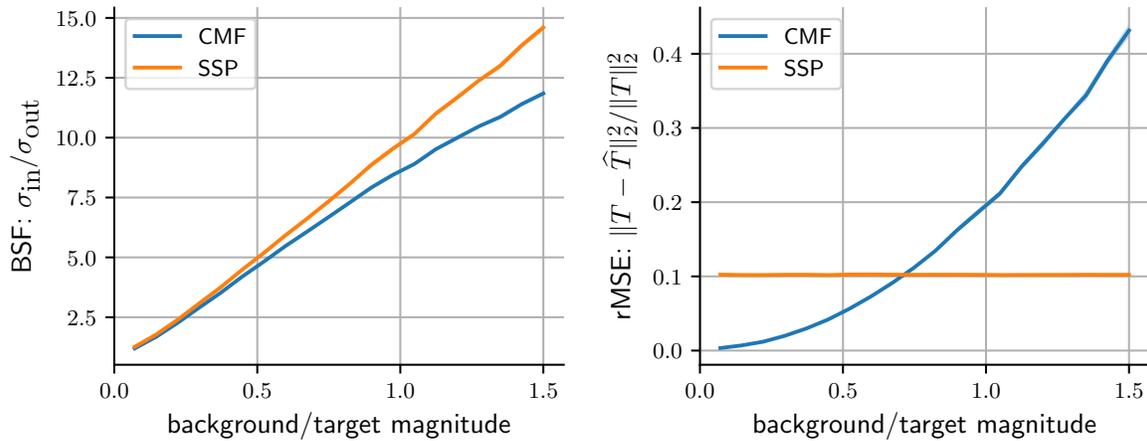


Fig. 3.4: Evaluation of background suppression algorithms on synthetically generated star scenes. The background was generated as a circularly symmetric light source, and stars were modeled as point sources injected additively into the scene. We evaluate background suppression performance via two metrics: background suppression factor (higher is better) and relative mean square error (lower is better). As the power in the background component increases relative to the signal power, the BSF increases accordingly, indicating favorable scaling behavior as the problem becomes more challenging. When the background component is sufficiently small, CMF preserves the target signal better than SSP. As the power in the background component increases, SSP yields superior performance by leveraging background estimates formed over multiple frames.

3.4 Light Curve Analysis

Characterization of space objects is an important component of SDA, and properties like object size, shape, attitude, and surface material can inform the classification of potential threats. Several studies have investigated how to estimate these object properties using light curves, a measure of the brightness of an object across time. Intensity variations in the light curve are influenced by positional angles, surface reflectivity, and object shape and attitude, thus creating an inverse problem to estimate the state parameters causing the changes of magnitude observed in these curves [12]. Techniques like the Unscented Kalman Filter have been studied in estimating spacecraft attitude using light curves [13, 14]. Additional research has shown that light curves can be fused with other data sources like angles to recover information such as spacecraft area and specular reflectivity coefficients [15].

To show that NGSatSentry is capable of providing photometric data used in these types of analyses, light curves were produced using the data collected in Joshua Tree in 2020. The detector counts of four geostationary satellites were trended across 802 frames of data, or approximately 2 hours. Figure 3.5 shows the maximum in-cluster pixel intensity associated with each of the satellites, removing outlier intensities greater than 1.5 times the standard deviation away from the mean. There are still some residual outlier intensities due to stars that cross the same focal plane position on the sensor as the satellite. Overall, the four intensity curves show distinct and unique features for each satellite, some curves have sharp peaks, others are smooth. This type of data could be fed into algorithms like those described above to estimate the object properties that result in these unique features. With ground-based light curves of geostationary satellites generated as a proof of concept, NGSatSentry’s future on-board observations of objects could be a critical asset in the field of object characterization.

Max Intensity Curve for Geostationary Satellites (Outliers Removed)

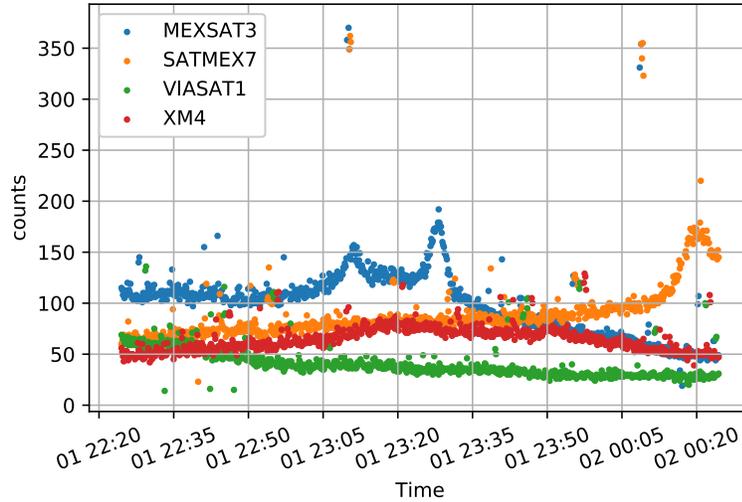


Fig. 3.5: Light curves for geostationary satellites. Four geostationary satellites were registered across 802 frames captured over a span of two hours and maximum in-cluster pixel intensity was computed as a function of time. Such light curves are commonly used for characterizing resident space objects via optical measurements. The fact that NGSatSentry can produce reasonable light curves serves as evidence that its measurements can be used not only for detection and tracking of RSOs, but also for phenomenology characterization.

3.5 Calibration and Astrometric Accuracy

Astrometric accuracy is important for producing useful SDA data, and distortion correction can vastly reduce line-of-sight error, especially in optical systems with wide FOV like the one used in this work. In this section, we evaluate astrometric accuracy before and after using the distortion correction procedure discussed in Section 2.3. Figure 3.6 shows histograms for the line-of-sight errors (presented in right ascension and declination coordinates) computed by registering exceedance cluster centroids to stars in the Tycho-2 star catalog. For images captured using the 3 MP camera, distortion correction results in a reduction in the combined error standard deviation by $0.512''$. The improvement is substantially more apparent in the images captured using 12 MP images, reducing the error standard deviation by $11.555''$. This large margin for improvement is due to the more severe distortion seen near the edges of the FOV which is captured by the larger 12 MP focal plane.

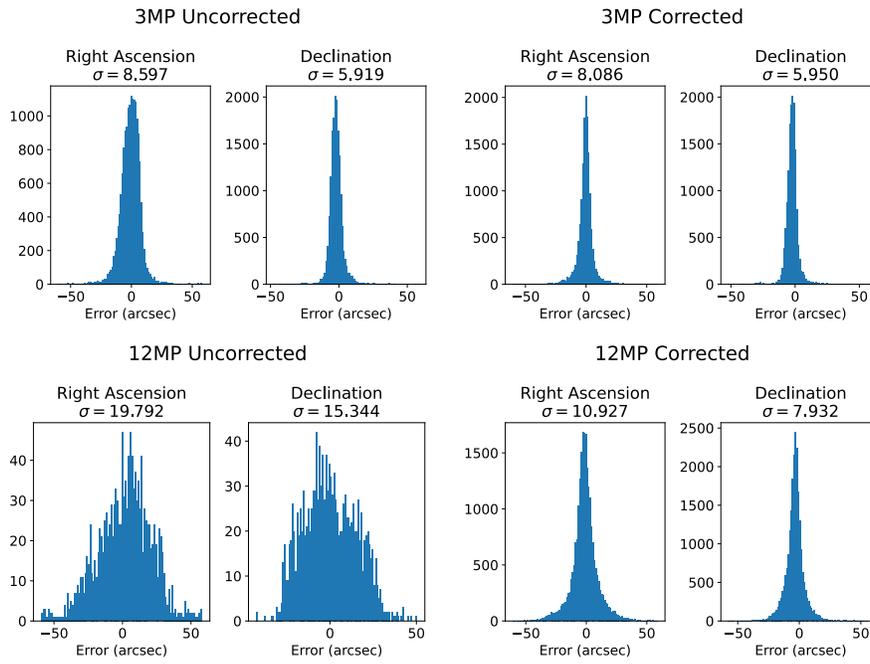


Fig. 3.6: Astrometric accuracy evaluation. Line-of-sight error for exceedance clusters registered to stars in the Tycho-2 catalog. Calibration drastically improves astrometric accuracy, especially when using the larger 12 MP sensor which captures light closer to the edge the field of view where optical distortion is typically more severe.

3.6 Comparison of Data Reduction Methods

Identification of star-induced exceedance clusters to be excluded from the downlink is a critical component in SDA data reduction. Here, we compare the catalog-free star identification technique discussed in Section 2.2 to the mature lost-in-space star registration technique described in [8]. A sequence of 10 images of the night sky containing 11,643 exceedance clusters was processed using each technique. The results are shown in Figure 3.7 which indicates exceedance clusters that would be downlinked (i.e., ones that were not determined to correspond to stars) using the catalog-free method, and exceedance clusters which had nearby stars in the Tycho-2 catalog [16]. The proposed catalog-free star rejection algorithm resulted in a reduction factor of 13.7 in the number of exceedance clusters to be downlinked. The two techniques agreed on the classification of 92.6% of exceedance clusters. Another 5.8% of exceedance clusters were correlated to nearby stars using the catalog registration technique but were not identified as stars by the catalog-free method. Some of these cases were caused by RSOs passing in front of stars which would result in erroneous omission from the downlink using catalog registration techniques alone. In contrast, the catalog-free method leverages multiple frames and correctly classifies these clusters as RSOs. Finally, 1.6% of exceedance clusters were not registered to the star catalog, but were not classified as RSOs using the catalog-free method. Although one may worry that this result suggests that RSOs could be omitted from the downlink, further inspection revealed that the vast majority of such cases are due to one of two reasons: First, closely spaced stars sometimes result in a single cluster with its centroid between the two stars, causing neither one to be registered to the catalog. Second, some exceedance clusters were induced by other distant objects that were not present in the Tycho-2 catalog, but correlated with objects in other astronomical catalogs. Furthermore, in our experiments, none of these “missed” centroids were correlated to cataloged RSOs.

		In Tycho-2	
		Yes	No
Downlinked	Yes	5.8%	1.5%
	No	91.1%	1.6%

Fig. 3.7: Data reduction results for a sequence of night sky images containing over 11,000 exceedance clusters. Clusters determined to exhibit star-like motion via the proposed catalog-free classification algorithm were excluded from the downlink, resulting in a data reduction rate of 13.7. Clusters were also registered against stars in the Tycho-2 star catalog for comparison. The objects that were not present in the star catalog were verified to not be present in the space-track.org TLE catalog of RSOs, and many were present in other astronomical catalogs, suggesting that the algorithm is not excluding potential objects of interest (i.e., RSOs).

3.7 Thermal Vacuum Test Results

A thermal vacuum chamber test was conducted on NGSatSentry’s COTS camera and lens system with the objective of demonstrating the unit’s ability to survive in a somewhat space-like thermal environment while also providing thermal data for optical performance evaluation. Unit testing in the vacuum environment is necessary due to the commercial hardware used as COTS hardware is often not designed to withstand the harsh space environment and can often fail due to poor thermal interface construction for vacuum, lack of resilience to extreme temperatures, or contamination from the use of materials with high outgassing rates. Additionally, due to thermal expansion and contraction, there is an expected impact on the optical performance as the unit undergoes large temperature swings.

The test conducted made use of an aluminum bell jar type chamber with a temperature controlled baseplate and viewing window in the side of the shroud. The viewing window was aligned with the lens so that images could be taken during test. The inside surfaces of the chamber were uncoated, bare aluminum which has a low infrared emissive property. This construction allowed for the unit to be effectively insulated to radiation heat transfer during the test which is important since the shroud of the chamber is unconditioned and therefore close to room temperature. The baseplate temperature was controllable by use of an external chiller that can pump warmed or chilled fluid through pipes embedded in the baseplate. In order to ensure the optical path was aligned with the shroud window, a fixture was designed to support the unit and align with the baseplate bolt pattern. To ensure the unit was well coupled thermally to the baseplate, every bolted interface in the fixture made use of a thermally conductive gap pad filler. Finally, multiple type-T thermocouples were attached to key parts of the fixture and in multiple locations of the unit under test (UUT). The test apparatus is shown in Figure 3.8.

Thermocouples were placed strategically so that the data could be used for multiple purposes such as fitting a thermal model of the device. Additionally, if any large variance in temperature was found across the device surfaces, an appropriate thermocouple could be chosen as the best location to attach thermistors for on-orbit telemetry feedback.

The test began by assembling the test fixture and UUT. The assembly was secured to the baseplate and alignment checked with the shroud window. Next thermocouples were placed on the fixture and UUT with a vacuum compliant room-temperature-vulcanizing (RTV) compound. A layer of aluminum tape was placed over each thermocouple to further reduce radiative coupling with the shroud of the chamber. Next, the chamber was sealed and pumped down to vacuum. After vacuum levels were reached the temperature of the system was elevated to 60 °C to allow for the faster curing time of the RTV compound to the manufacturer’s specifications. After curing the RTV, the temperature was reduced back to ambient and vacuum was broken. The thermocouples were then double checked for adhesion and connected to the data acquisition system. All thermocouples were deemed to be working by a simple touch test and their mapping to the data acquisition system was confirmed. The chamber was re-sealed and pumped back down to vacuum again at which point the UUT was powered on. Initial tests confirming the UUT’s functionality were then performed.

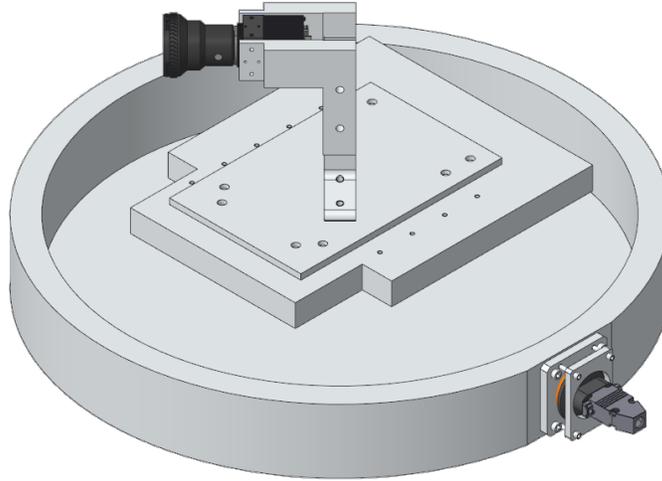


Fig. 3.8: Thermal vacuum chamber test apparatus (shroud not shown).

The set point of the baseplate was then changed and the change in temperature was observed across all thermocouples. Once stability criteria of $1\text{ }^{\circ}\text{C}/\text{h}$ held for 30 minutes and the base of the UUT was within $\pm 3\text{ }^{\circ}\text{C}$ of the targeted temperature, image capture tests were performed. This process was repeated in roughly $10\text{ }^{\circ}\text{C}$ increments between $-24\text{ }^{\circ}\text{C}$ and $61\text{ }^{\circ}\text{C}$ so that optical measurements could be evaluated across a wide temperature range. The temperatures during image capture throughout the test are shown in Table 1.

Location	Temperature ($^{\circ}\text{C}$)							
	T_1	T_2	T_3	T_4	T_5	T_6	T_7	T_8
cold plate	62.2	49.1	25.5	13.9	1.86	-8.00	-19.5	-36.4
camera mount, left	59.5	48.9	29.1	19.0	8.75	0.33	-9.73	-24.5
arm support, bottom	57.9	47.9	29.2	14.8	10.1	2.18	-7.18	-20.8
camera mount, right	57.5	47.4	28.6	19.1	9.3	1.31	-8.25	-22.2
front chassis, top	60.3	49.6	29.6	19.4	9.04	0.53	-9.67	-24.7
mid chassis, top	61.0	50.6	31.2	21.4	11.3	3.03	-6.88	-21.5
lens, outboard	47.4	42.4	30.4	23.7	16.9	12.0	5.24	-4.83
lens, inboard	46.9	42.0	30.2	23.6	17.0	12.2	5.62	-4.19
rear chassis	61.3	50.7	31.5	21.8	11.8	3.68	-6.06	-20.3

Table 1: Thermocouple measurements at thermal equilibrium during thermal vacuum chamber testing.

The test qualitatively demonstrated the hardware’s capability to operate in a vacuum environment and additionally to function to MIL-STD-1540 acceptance temperature limits of $-24\text{ }^{\circ}\text{C}$ to $61\text{ }^{\circ}\text{C}$. This encouraging result demonstrates that this commercial hardware may be a viable option for space-based use cases. Next steps will include a full qualification campaign that will involve increasing the temperature bounds to the qualification levels of $-34\text{ }^{\circ}\text{C}$ to $71\text{ }^{\circ}\text{C}$ and introducing thermal cycling and soaks on multiple units. These additional stressors would confirm quality of workmanship, statistical confidence over a larger sample, and tease out any longer duration issues that could arise from being in a space environment. Finally, the images captured during thermal vacuum testing will be analyzed to predict how noise statistics vary as a function of temperature.

4. SUMMARY

In this paper, we introduced NGSatSentry, a customizable, low size, weight, and power hosted payload for space domain awareness. We discussed NGSatSentry's on-board algorithms that operate in real-time to detect resident space objects and perform data reduction to minimize downlink utilization. These algorithms were implemented on flight-like hardware and through experiments using synthetic and real data, we showed the efficacy of the processing chain and demonstrated that NGSatSentry can detect and characterize geostationary satellites. Finally, we presented results from recent testing on the payload's imaging hardware in our facility's thermal vacuum chamber to demonstrate the feasibility of low-cost, readily available hardware with the potential to enable deployment at scale.

5. REFERENCES

- [1] J. A. Haimerl and G. P. Fonder, "Space fence system overview," in *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*, 2015.
- [2] C. Stokely, E. Stansbery, and R. Goldstein, "Debris flux comparisons from the goldstone radar, haystack radar, and hax radar prior, during, and after the last solar maximum," *Advances in Space Research*, vol. 44, no. 3, pp. 364–370, 2009.
- [3] R. F. Bruck and R. H. Copley, "GEODSS present configuration and potential," in *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*, 2014.
- [4] G. Turin, "An introduction to matched filters," *IRE Transactions on Information Theory*, vol. 6, no. 3, pp. 311–329, 1960.
- [5] R. M. Haralick, S. R. Sternberg, and X. Zhuang, "Image analysis using mathematical morphology," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-9, no. 4, pp. 532–550, 1987.
- [6] D. Mortari, "Search-less algorithm for star pattern recognition," *The Journal of the Astronautical Sciences*, vol. 45, no. 2, pp. 179–194, 1997.
- [7] D. Lang, D. W. Hogg, K. Mierle, M. Blanton, and S. Roweis, "Astrometry.net: Blind astrometric calibration of arbitrary astronomical images," *The Astronomical Journal*, vol. 139, no. 5, pp. 1782–1800, 2010.
- [8] D. Mortari, M. A. Samaan, C. Bruccoleri, and J. L. Junkins, "The pyramid star identification technique," *Navigation*, vol. 51, no. 3, pp. 171–183, 2004.
- [9] B. Guan, J. Zhao, Z. Li, F. Sun, and F. Fraundorfer, "Relative pose estimation with a single affine correspondence," *IEEE Transactions on Cybernetics*, pp. 1–12, 2021.
- [10] D. Campbell, L. Petersson, L. Kneip, and H. Li, "Globally-optimal inlier set maximisation for camera pose and correspondence estimation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, no. 2, pp. 328–342, 2020.
- [11] P. Besl and N. D. McKay, "A method for registration of 3-d shapes," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239–256, 1992.
- [12] Y. Matsushita, R. Arakawa, Y. Yoshimura, and T. Hanada, "Light curve analysis and attitude estimation of space objects focusing on glint," in *The First International Orbital Debris Conference*, vol. 2109, p. 6091, 2019.
- [13] K. A. Rush, M. Yost, L. Smith, and A. Zizzo, "An application of the unscented Kalman filter for spacecraft attitude estimation on real and simulated light curve data," in *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*, 2020.
- [14] C. J. Wetterer and M. Jah, "Attitude determination from light curves," *Journal of Guidance, Control, and Dynamics*, vol. 32, no. 5, pp. 1648–1651, 2009.
- [15] M. Jah and R. A. Madler, "Satellite characterization: angles and light curve data fusion for spacecraft state and parameter estimation," in *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference*, vol. 49, 2007.
- [16] E. Høg, C. Fabricius, V. Makarov, U. Bastian, P. Schwkendiek, A. Wicenc, S. Urban, T. Corbin, and G. Wycoff, "Construction and verification of the Tycho-2 catalogue," *Astronomy and Astrophysics*, vol. 357, pp. 367–386, 2000.