Extremely Accurate Star Tracker for Celestial Navigation

Greg A. Finney,[1] Stephen Fox,[1] Bijan Nemati,[2] Patrick J. Reardon[2]
[1]IERUS Technologies, Inc.,
[2]University of Alabama in Huntsville, Center for Applied Optics

Abstract

IERUS Technologies and the University of Alabama in Huntsville Center for Applied Optics teamed to transition the focal plane metrology technique developed by the NASA Jet Propulsion Lab (JPL) to industry. The interferometric fringe shifting technique enables the location of pixels in a focal plane array to high precision. We developed a laboratory testbed that has shown the potential to determine the effective location of pixels to 1/100th of the pixel pitch. This technique, combined with a precision telescope, was shown via simulation to measure the location of stars on the focal plane to very high accuracy. We developed an initial optical, electrical, and structural design concept that will enable the desired accuracy. Thermal analysis indicated that the anticipated environment would not degrade the accuracy beyond this limit. Optical analysis showed the nominal design would provide near-diffraction limited performance. The system can detect enough stars to calculate its orientation over >99% of the celestial sphere. The projected size is consistent with existing compact star trackers (1300 cc, < 1 kg) which have significantly lower accuracy. In addition to the standard star tracker attitude determination function, this instrument is also capable of detecting earth orbiting satellites with enough precision to provide position data. It thus has the potential to be used for space situational awareness.

In upcoming work, we will complete development of the calibration facility, integrate a calibrated focal plane with a custom designed telescope, and demonstrate the breadboard system. The presentation will provide an overview of the interferometric fringe shifting technique, the design concept for the star tracker, and preliminary feasibility analysis for celestial navigation.

Introduction

A compact star tracker with improved accuracy is of interest both for providing improved attitude determination and for use as an astrometry sensor. There are four primary error sources for a star tracker: detector noise, field distortion error, pointing jitter, and focal plane errors. This effort focused on the most challenging, focal plane errors. The interferometric fringe method developed by National Aeronautics and Space Administration (NASA) Jet Propulsion Laboratory (JPL) enables calibration of focal plane array geometry to 10 milli-pixel (mpix). IERUS Technologies and the University of Alabama in Huntsville (UAH) Center for Applied Optics (CAO) teamed to execute this effort to improve the method and to demonstrate the transition of this technology and feasibility of developing an extremely accurate star tracker in a package with size, weight, and power requirements comparable to commercially available star trackers.

The error budget for the star tracker was focused on the error sources which could be most readily influenced through the star tracker design: detector noise and focal plane errors. Field distortion can be measured and corrected, and pointing jitter is controlled by the platform to which the star tracker is mounted. Detector noise is dominated by shot noise based on the selection of a low-noise focal plane array (FPA) and can be reduced by design of the point

spread function, centroiding algorithm, aperture diameter, and integration time. FPA errors are dominated by the calibration accuracy and the dimensional stability of the FPA. The dimensional stability can be well controlled by temperature control of the FPA. In our error budget, detector noise was limited to 27 mpix and calibration accuracy to 15 mpix. The other error sources, along with a generous reserve, brought the total error budget to 54 mpix. This error budget was based on collection of stars with 10th magnitude stars (V-band).

FPA Calibration Process

The innovation we bring to this task is the precise calibration of the FPA systematic biases as well as the calibration of optical distortion in the optics caused by improper alignment or wavefront errors on the optical elements. The focal plane can be calibrated extensively on the ground prior to launch, and to a lesser extent periodically on orbit. The calibration involves illuminating the focal plane with lasers launched from the tip of a series of fibers. Two fibers illuminating the focal plane create fringes on the pixels. If the fiber ends are attached to a thermally stable metrology block, the fringe spacing can be a stable reference that is the basis of metrology of the pixels. The fringes can be made to move across the focal plane by shifting the phase of the light coming from one of the fibers. With these moving fringes we can calibrate the pixel response functions of all the pixels, the lowest orders of which are familiar as the QE and the position of the pixel.

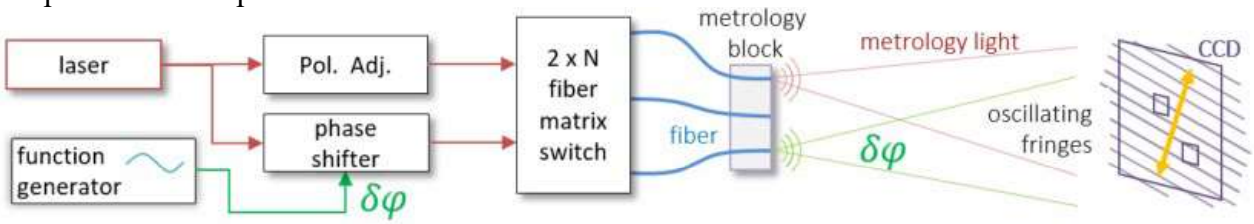

Figure 1: Pixel metrology approach.

The additional detail for the calibration approach is shown in Figure 1. A stabilized He-Ne laser is coupled into a single-mode fiber and split, with one side receiving a phase shift while the other side is only adjusted for polarization mismatches. A small network of 1x2 fiber switches are used to selectively turn the laser output of 4 fibers situated in a 'metrology block' on or off. Pairs of fibers can be selected that have a phase shift applied, causing moving interference fringes on the detector under test. The essence of the technique is to compare the intensity measured by a pixel with that expected for the pixel given its location. The difference between the two values is used to calibrate the pixel location.

The method relies heavily on ideal fringes and ideal input to the camera. Thus, it is important to reject spurious fringes due to any stray light or additional optics. The light path between the fiber tips and the sensor should only be free space. For the test sensor, we chose the large format Sony IMX 455 sensor, inside a ZWO ASI6200MC camera. In this packaging, the sensor comes with a window that must be removed. Before the removal there is clear evidence of 'etalon' effect on the window. After removal there are no such fringes left.

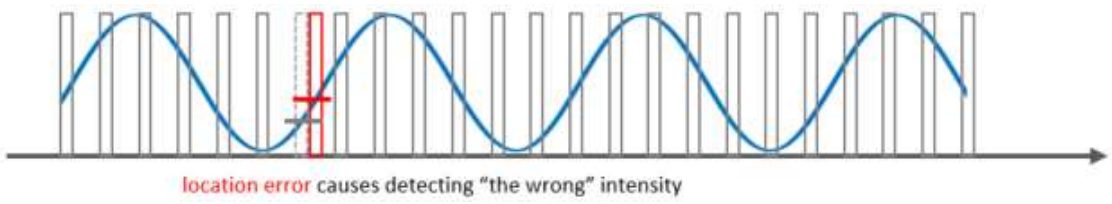

Figure 2: Basic concept of pixel metrology with calibrated fringes. The measured intensity is compared with the expected intensity and the difference is used to calibrate the pixel location.

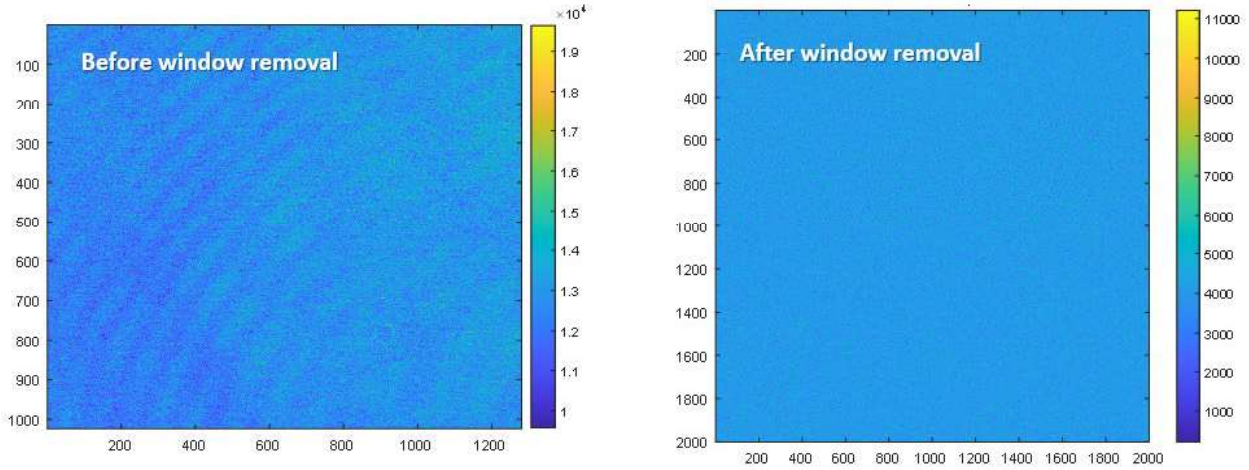

Figure 3: Single fiber illumination gives indication of any stray laser light and spurious fringes. On the left is a frame taken with only a single fiber turned on, before window removal, and on the right is the same after removal. The spurious fringes are gone.

The metrology block and sample fringes are shown in Figure 4 below. The metrology block in the phase I design has 4 fiber tips, labeled A through D. A and B are at one phase and C and D are shifted via the fiber phase shifter. 4 combinations of (A or D) x (B or C) are possible, creating various fringe spacings on the camera. Additionally, the fiber array is mounted on an encoded rotation stage, so that any fringe direction can be generated on the camera. The ability to create different fringe directions is important to the calibration process. In addition to the IMX455, we also processed data provided through our collaboration with JPL, which used an IMX411 sensor.

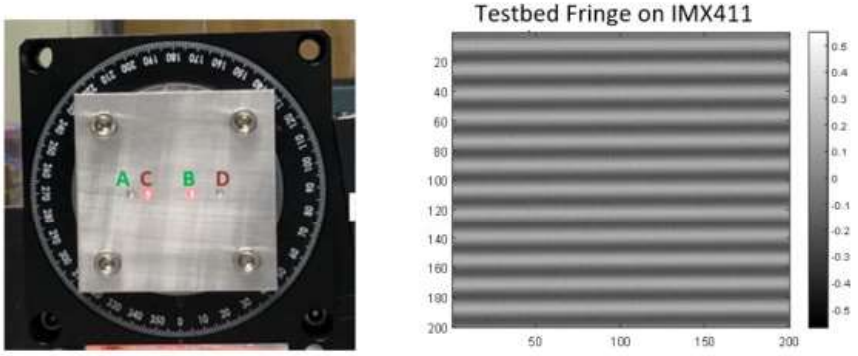

Figure 4: Metrology block (left) and sample fringes from a large format detector (right).

Two fibers that created interference fringes on the camera will not create straight, Young's slit fringes, but hyperbolic ones. Fringes on the detector will have a spacing given by:

$$\Lambda = \frac{\lambda}{a} \cdot d,$$

where λ is the laser wavelength, a is the fiber separation, and d is the distance to the detector.

The left side of Figure 5 shows a comparison of hyperbolic fringes, compared with 'ideal', straight fringes from infinite slits. For a large detector like the Sony IMX 411, which has 150 million pixels, each 3.76 microns wide, by the corner there is a significant difference between the two fringes.

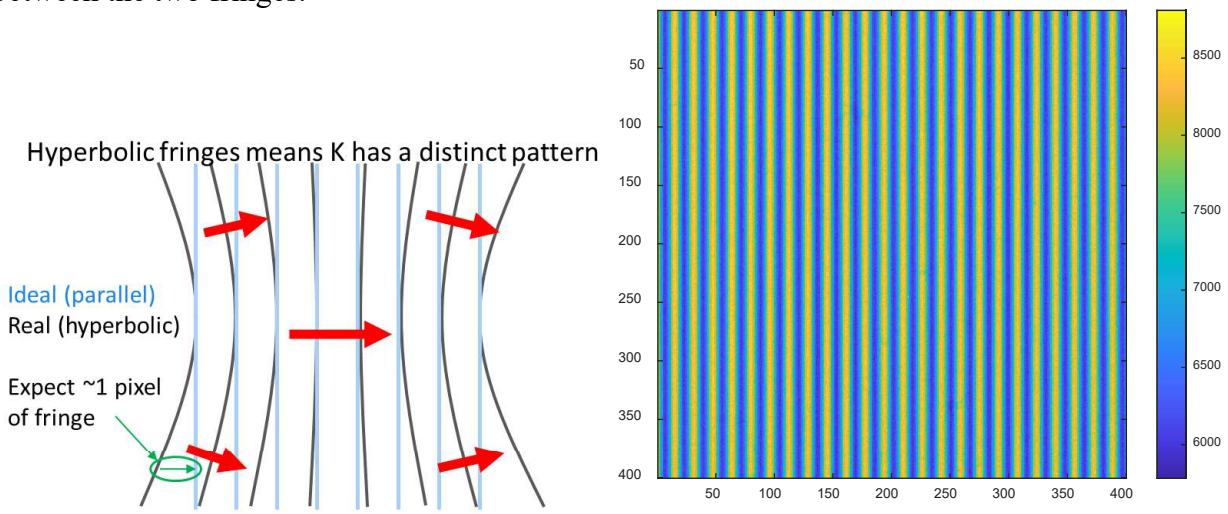

Figure 5: Fringes from a pair of fibers will have a hyperbolic pattern on the detector. One the left, the arrows show the wavenumber vector direction. On the right is a typical 400x400 pixel segment from the IMX411.

The hyperbolicity of the fringes is a low-order effect on the focal plane, needing only a few parameters to describe. We can break up the large format detector into 'tiles' and measure the wavenumber $K = 2\pi/\Lambda$ as a vector over each. We did so for simulated fiber-pair fringes, and for actual data. These are shown in the figure below. We see that the measured wavenumber vectors (top row, where each pixel is a 400x400 area on the detector) is consistent with the expected values (bottom row). For the K_x component, there is a large average value, with the values dropping slightly in the corners. For the K_y component, the average value is near zero, with adjacent corners having opposite signs. Solving for the geometry is possible with this information.

In the initial calibration algorithm, a global nonlinear least squares fit was performed to find the average intensity, fringe modulation, phase, and fringe wavevector components. Our expectation was that the global fit would average out the errors, and the location error for each pixel could be determined. Using the new location values, the process could be repeated, and the location error determined. However, working through the process in simulation, we found that the iterative process did not converge. The issue was found to be the lack of a global coordinate system. The pixel errors were generated in a reference frame determined by the FPA, but that reference frame was not well defined because of those errors.

Figure 8 shows the result of applying this method to simulated data. In the simulation, pixel row and columns errors were randomly generated using a power law distribution with the standard deviation of errors equal to 2% of the pixel pitch (i.e., 20 mpix) and the peak to valley (P-V) error is 100 mpix. After correcting with measurement of the simulated error, the standard deviation of the residual error is 9 mpix, the P-V error is 25 mpix. We expect these errors can be further reduced by an analytical estimate of the gradient error. The final residual for ideal results with all gradient removed shows an error of 27 µpix. Having demonstrated the technique in simulation, in the future we will use the calibration testbed to demonstrate the algorithm on data collected with the IMX455 FPA.

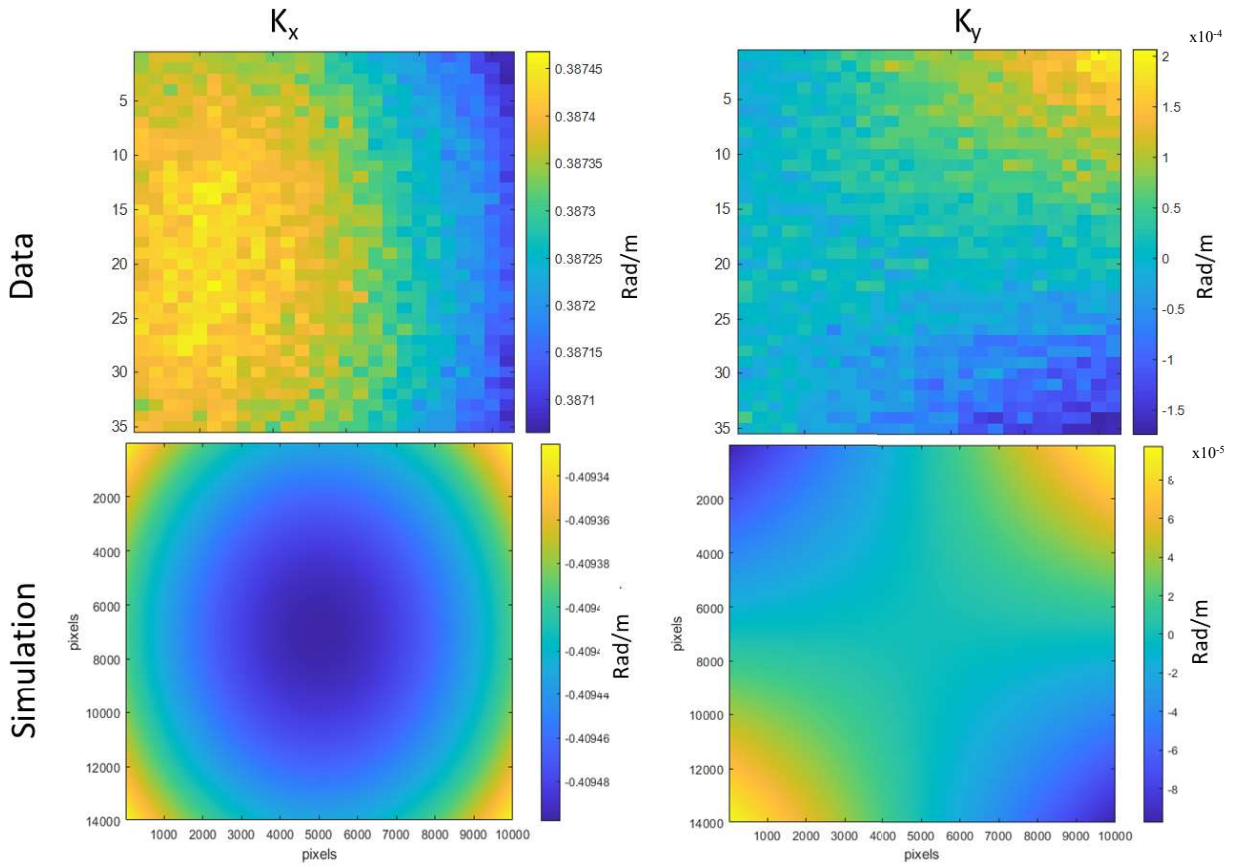

Figure 6. Comparison of simulation and laboratory data

The algorithm has been matured by developing a method for establishing an absolute reference frame. The reference frame is still established by FPA, but in a manner that is explicit and repeatable. This method is illustrated in Figure 7. The process involves definition of three points, labelled 1, 2, and 3. Points 1 and 2 are selected from the same column of the FPA, while point 3 is selected from the same row as Point 1. The rays 12 and 13 will be approximately orthogonal, but the process treats point 3 as a proxy and will define the true orthogonal direction. The ray 12 designates the *x*-axis. It is found by rotating the fringe pattern until the value of $\vec{K} \cdot \vec{r}$ between the two points is approximately a maximum. As the metrology block is rotated in small increments near this orientation, and the fringe phase difference between Points 1 and 2 is computed. The angle of maximum phase difference defines the *x*-axis. In addition, the distance between the points can be accurately found from the fringes. The metrology block then rotated by 90°. Referring again to Figure 7, the true *y*-axis is shown by the solid black line and the ray 13 is shown by the short dashed blue line. In this orientation, the fringes will be running vertically, and are used to compute the distance between Points 1 and 3.

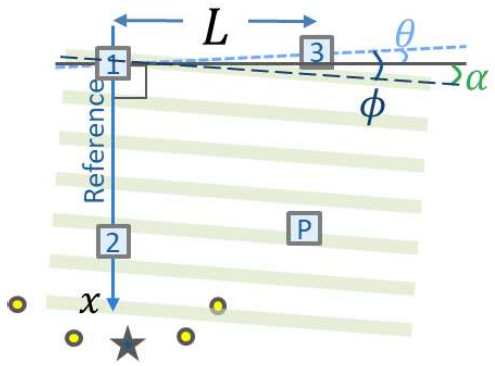

Figure 7. Establishing the absolute reference frame for the FPA

This process is repeated to measure the distance between points 2 and 3. Finally, the stage is returned to the original position (maximizing $\vec{K} \cdot \vec{r}$ along the *x*-axis). However, due to errors in the rotation stage, this rotation will not be exact. This is the situation in Figure 7. The direction of fringe at the nominal position (with error exaggerated) is shown by the long dashed black line. By computing the phase difference between Point 3 and fringe along the direction of the *x*-axis, the angle ϕ can be determined. The angle θ is determined using the measured lengths of the triangle 123. This procedure finally allows calculation of the angle α, defining the *y*-axis.

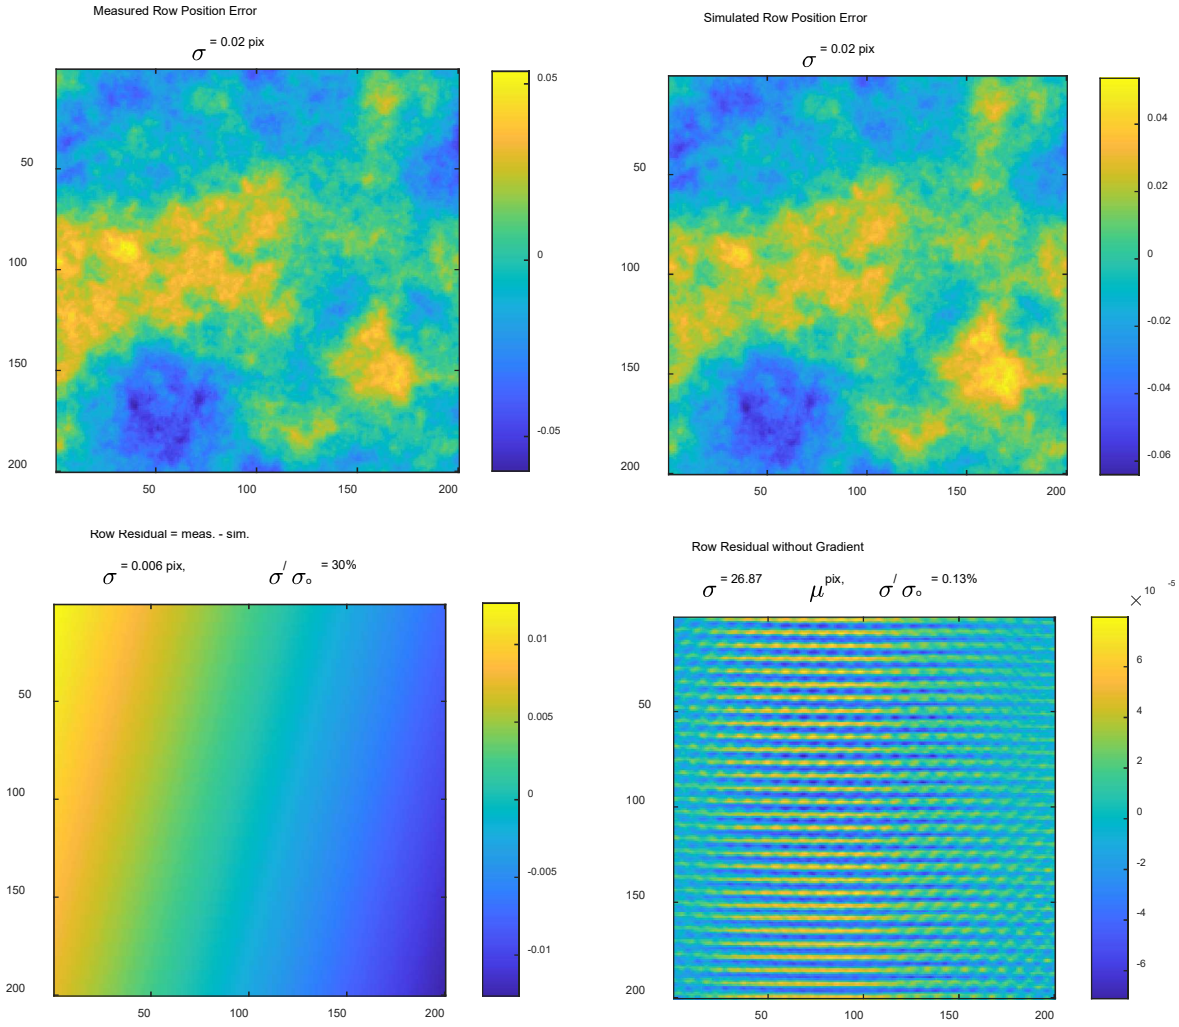

Figure 8. Simulated measurement of row pixel errors using 3-point absolute reference system (upper left) Simulated error, (upper right) measurement of simulated error, (lower left) residual gradient error, (lower right) residual with gradient error removed.

Thus far we have succeeded in setting up a testbed with a large format scientific CMOS detector to be calibrated and have taken test images, as shown in Figure 9.

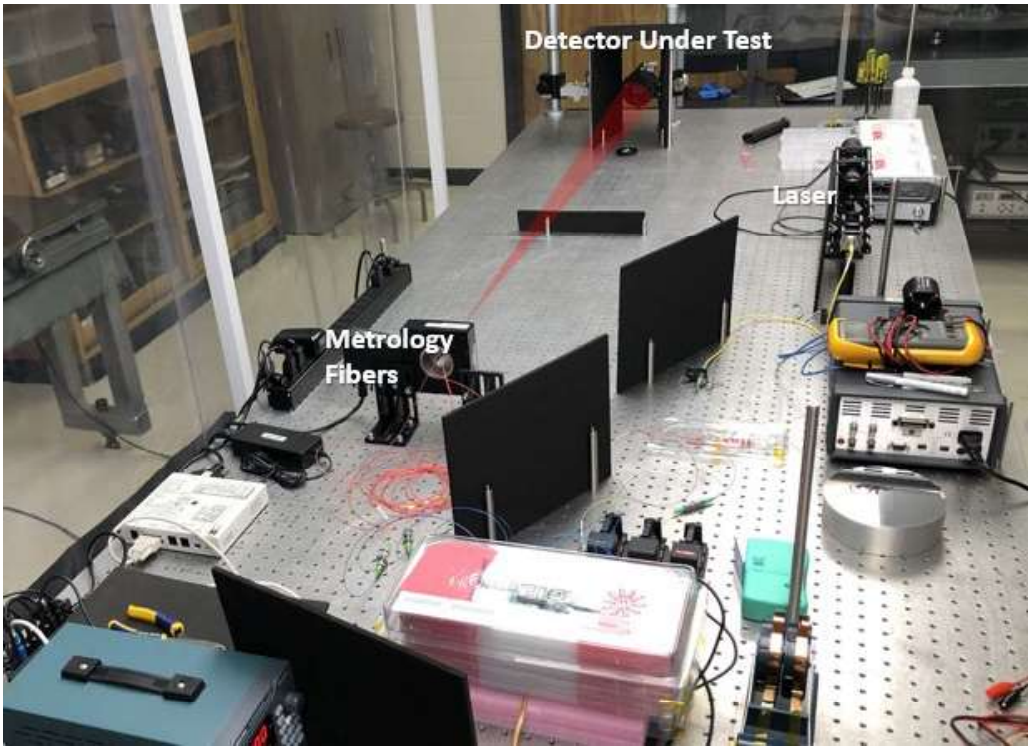

Figure 9. Photograph of testbed setup

Testbed automation has been implemented using a python-based master script that controls:

- MEMS status (fibers on/off)
- Rotation stage angle
- Camera (gain, exposure, # of frames)
- Power supply/phase shifter
- Data organization

A run of the script can produce darks, flats, and fringes at different angles, as well as label and organize the data into an automatically generated folder.

Current challenges we are resolving include overcoming physical error sources. The fibers are very sensitive thermal and vibration transducers. The testbed has been enclosed in a portable soft wall clean room to reduce these environmental effects. In addition, we are developing a more effective baffle system to reduce the stray light effects we are seeing in the data.

Star Tracker Design

The star tracker concept design is shown in Figure 10. Features of the design include:

- f/17 catadioptric design
- 2048x2048, low-noise, back side illuminated CMOS FPA
- Airy disk diameter: 3 pixels
- 450 – 1000 nm band
- Telescope overall length: 310 mm
- Volume: <1300 cm^3
- Mass: <1 kg

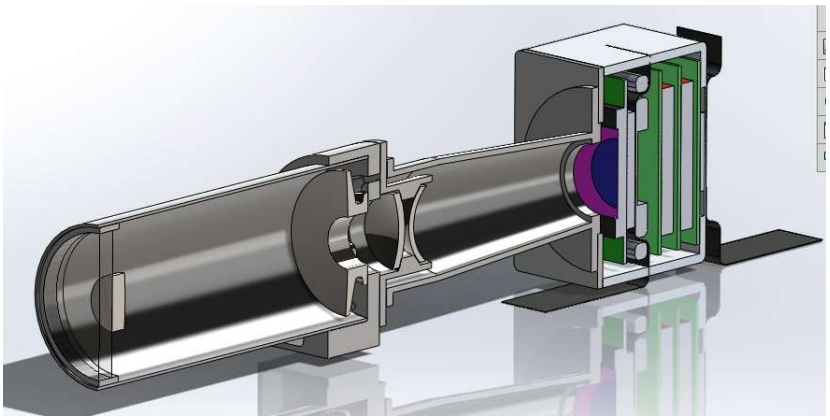

Figure 10. Cutaway view of star tracker

Contemporary star trackers generally have a modest field of view and low sensitivity to enable a lost-in-space capability. Lost-in-space refers to the star tracker's ability to measure its orientation with no external information and using only the onboard processing of the image it collects. A low sensitivity is needed to avoid having too many stars to try to match to its internal catalog, which requires the field of view be large enough to collect sufficient stars in order to calculate its orientation. Star trackers typically operate on stars of about 6th magnitude or brighter (~4,000 stars).

Our high accuracy approach requires a narrow field of view in order to measure the star positions with sufficient precision. The narrow field of view, in turn, requires the sensor to be very sensitive in order to collect sufficient stars in each image, driving toward a large catalog. We are currently targeting 10th magnitude stars as the threshold (~300,000 stars) While the orientation algorithm is still in development, we anticipate that the large number of stars will require the star tracker to have an initial estimate of the boresight direction. For many applications, sacrificing lost-in-space capability for improved accuracy is beneficial. Coarse orientation estimates can be provided either by an inertial measurement unit or by other sensors (e.g., sun / Earth limb sensors).

An analysis was conducted by taking ~40,000 randomly pointed fields of view across the celestial sphere and computing the number of stars in each. For 10th magnitude stars, less than 1% of views contained less than two stars. If the detection threshold is increased to 8th magnitude stars, then almost half of the fields of view contain at least two stars, the fields of view with less than two stars are always within a few degrees of a field with 2 or more. Using brighter stars offers the advantage of a lower integration time, which reduces the impacts of platform motion.

A preliminary thermal analysis was done to assess the impact of thermal loads during operation. We assumed that the star tracker was shielded from direct solar illumination. This is a realistic assumption given the challenges that would be presented by stray light introduced by solar illumination. The initial analysis to check our modeling methodology was a 20°C soak. A transient analysis was performed based on thermal loads from the electronics, conductive cooling through the mount to the vehicle, and radiation cooling to space. In additional to the heat path through the primary mount, some heat spreaders were attached to the housing and the FPA TEC to create a heat path directly to the heat sink. Based on our earliest analysis, we placed an upper bound of 24 W of dissipated by the electronics, so this value was used for the analysis. Later analysis refined this estimate to the much lower value of 5 W, making this analysis very conservative. Likewise, the radiational cooling assumption was very conservative, that is, it will tend to overestimate the thermal gradient. We assumed that one half of the star tracker was

exposed to cold space and emitted blackbody radiation corresponding to a temperature of 320 K and emissivity of 0.5. Using copper heat spreaders to connect the thermal loads to the vehicle resulted in a significant tilt. However, using spreaders made from pyrolytic graphite sheet reduced the thermal deformation to an insignificant level.

Imaging Performance Predictions

For statistical prediction of imaging performance, an array of equal magnitude stars was constructed for use with our sensor modeling tool, the Aero-Optical Prediction Tool (AerOPT). [1] The stars were located far enough apart to ensure there was no overlap of their images. The stars were precisely located in object space so that their image was centered upon one of the 9x9 subsamples that constituted each pixel. In this manner, the image centroid was systematically varied from the pixel center across one quadrant of the pixel in steps of $1/9^{th}$ pixel. This process was repeated to generate 100 frames using independent random draws based on sensor noise characteristics. Examples images in counts are shown in Figure 11. Using this FPA image, the centroid was computed by eliminating saturated pixels (shown in red) and performing a nonlinear least squares fit to the region of interest around each star. Each pixel was weighted based on the uncertainty due to Poisson statistics, propagated from electrons to digital counts on the FPA. The fit was initialized using a center of mass centroid to ensure a good estimate for the fit. Centroid root-mean-square error in the row and column axis for the set of 100 frames was computed to estimate the detector noise error. This analysis was performed for stars with V_{mag} = 3, 8, 10, and 12, as well as a run with zero noise to assess the performance of the fitting algorithm in the absence of noise.

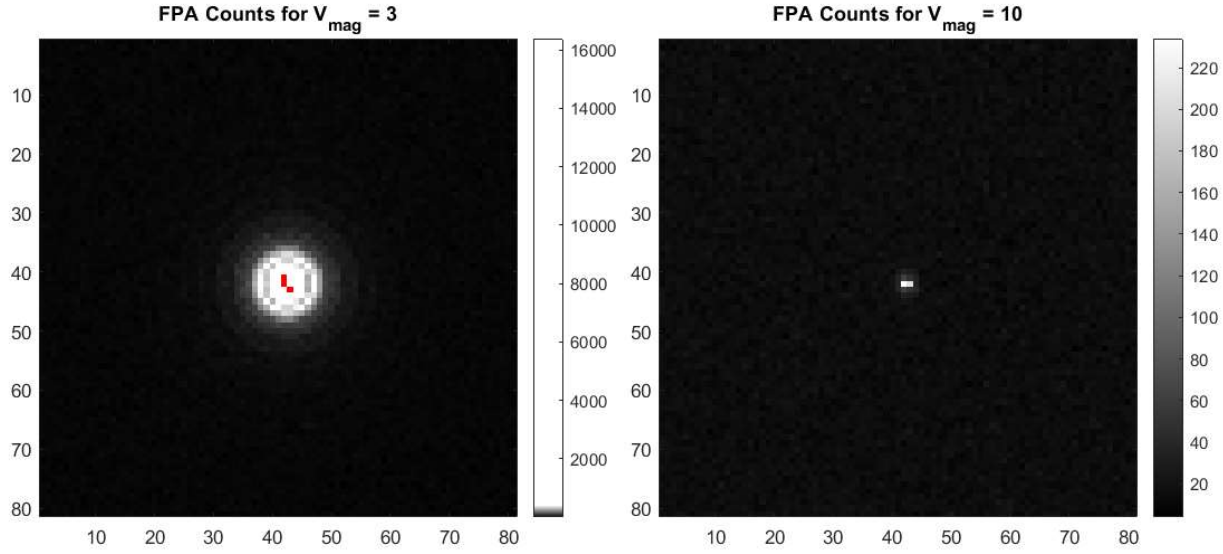

Figure 11. Example results of FPA response in model used for centroid analysis. Red pixels indicate saturation

A summary of these results is presented in Table 1, which shows the error is significantly below the target error of 0.1 pixel. In addition, the centroid error due to photon noise at the target brightness of V_{mag} = 10 is dominated by the nonlinear fit. This observation suggests additional improvement may be possible by improving the fitting function used. As noted above, some slight degradation may result from the asymmetrical point spread function in the presence of thermal deviations. We have also demonstrated in the on-going Phase II that field distortion may be calibrated to within the budget constraints of this design.

Table 1. RMS centroid error summary

V_{mag}	RMS error (pixels)
3	0.033
8	0.030
10	0.038
12	0.086
Noiseless	0.029

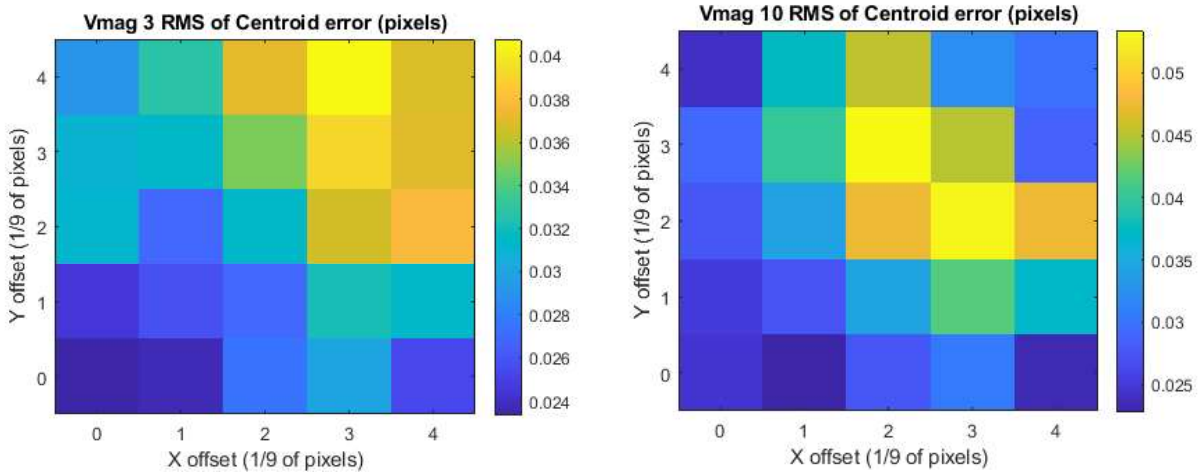

Figure 12. RMS centroid error as a function of centroid location. (left) Vmag = 3 (right) Vmag = 10

A more complex scene containing an actual stellar distribution is being developed for use with AerOPT. For this project, we will generate a physically and radiometrically accurate scene as viewed from the star tracker above the Earth's atmosphere. The selected scene is the open star cluster, the Pleiades (M45). It was chosen because it is readily recognizable, has a combination of bright and dim stars, and includes several challenging features (interstellar dust and galaxies). The scene is designed to cover approximately twice the telescope view of view in both horizontal and vertical directions, which will allow for variations as we test different aspects of the algorithms. The scene includes basic elements and will become more complex as the project and algorithms mature. We will add elements in the following order:

- ✓ Stellar spectral intensity covering B, V, R, and I bands (to $m_V \leq 13$ which yields ~350 stars in the field of regard)
- ✓ Stellar background ($m_V > 13$ averaged over field of view)
- ✓ Zodiacal light
- ✓ Deep sky objects (below noise floor at maximum anticipated integration time)
- ✓ Interstellar dust (below noise floor)

Likewise, the sensor model is currently very basic, and complexity will be added. Characteristics to be modeled will include (check marks indicate current model)

- ✓ Uniform responsivity, dark current, and analog to digital conversion
- ✓ Linear response up to saturation
- ✓ Zero crosstalk
- Non-uniform analog to digital conversion
- Nonlinear response near saturation

- Nonuniform responsivity and dark current
- Nonzero crosstalk
- Camera blooming
- Multiple reflections
- Stray light

Additional features are dependent upon receiving necessary quantitative inputs from the FPA vendor.

Stars are modeled as point sources located at the average distance of the Pleiades (444 ly). Parallax (or any other motion) is currently not modeled, as all the stars are at the same location and the sensor is located at the origin of the ICRF coordinate system. Signature data was taken from the SIMBAD online star catalog. [2]

For comparison, a portion of the resulting AerOPT image at 1 s integration time with a few stars annotated is shown in Figure 13, along with an image from the Digitized Sky Survey (DSS) with the same stars marked. (The AerOPT image is rolled approximately 45° relative to the DSS image.) Finally, Figure 14 shows a comparison of the images of various stars with brightness and contrast scale adjusted to see both bright and dim portions of the image. The brightest star is Alcyone (eta Tau), which has 6 pixels that are saturated. The dimmest star, Cl Melotte 22 1713, has only 86 counts. The background noise, due almost entirely to dark current, has a mean of 10 counts and a standard deviation of 2.7 counts.

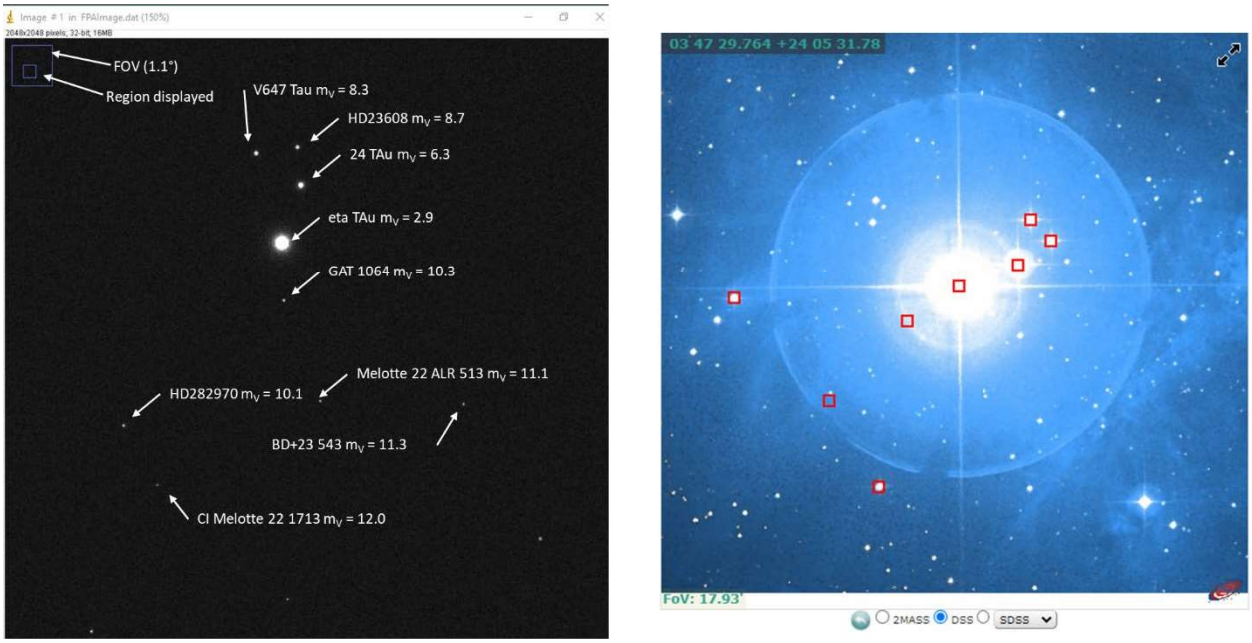

Figure 13. Portion of AerOPT image and DSS image

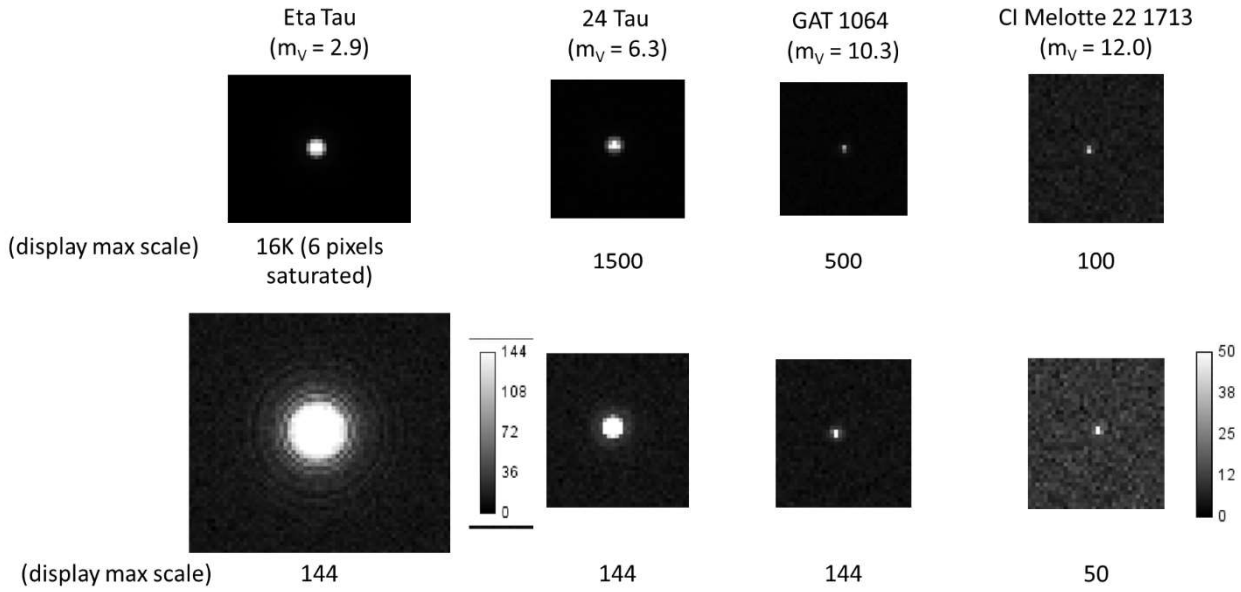

Figure 14. AerOPT image chips showing stars at different contrasts

Application to Celestial Navigation

In addition to functioning as a conventional star tracker to perform attitude determination, the star tracker has the potential to provide celestial navigation by determining angles to objects within the solar system. Because of the finite distances to these objections, collecting multiple angles will enable computing the position with an Earth centered reference frame. We considered two types of objects—natural satellites in orbit around the sun that are located the asteroid belt and artificial satellites in orbit around the Earth. We quickly found using asteroids with the compact star tracker was not feasible. However, using this star tracker for celestial navigation by targeting artificial Earth satellites may be feasible and is discussed below.

The scenario under consideration for this analysis is a temporary loss of Global Positioning System (GPS) signal. This loss could be due to a spacecraft anomaly, jamming, damage, or disruption of the GPS system. In any case, we assume that accurate ephemeris data for satellites is available. Two classes of satellites are being investigated for this task, GPS satellites and geostationary equatorial orbit (GEO) satellites. These satellites are large, spend relatively little time in eclipse, and are visible from a diverse set of orbits. We first investigated the detectability of GPS satellites. The analysis was derived from published information on GPS observations. [2] This reference reported on four GPS Block II/IIA satellites observed over two nights in the B, V, R, and I bands (although the collection matrix was not complete over all of those parameters). The conclusions based on this limited set of observations (along with some additional information from the satellite manufacturer) is that signatures of satellites tend to decrease and become more red over time, the signatures are generally consistent between satellites of the same age, the absolute magnitude of the satellites is modeled well by a linear fit to phase angle in the range of 90° - 166° (phase angle of 180° represents the satellite position at opposition; at 166° the satellites entered eclipse during the observation season). The absolute magnitudes were adjusted for a standard range of 20,000 km from the observer. At their brightest (near 166°), the satellites had I band magnitudes of 8-10 and V band magnitudes of 10-12. This decreased to 13-14 for both bands for a phase of 90°.

For the initial analysis, a moderately bright value of m=12 was used for the V and R bands, which corresponded to roughly a phase angle of 120°. A simple model developed for the telescope and a Teledyne CIS 120 focal plane was used to estimate detectability. This combination predicted the satellites would be difficult but possible to detect with a signal to noise ratio of 8 if the sensor was able to approximately track the GPS satellite. The apparent motion would cause about 20 pixels of image smear, which would push the signal below the noise threshold. While it is possible the signal could be recovered through image processing, that was considered beyond the scope of the present analysis. However, the slew rate is modest enough that open loop tracking is practical and worth further consideration. The signatures of geostationary communications satellites were also reviewed and found to be similar. [3]

To determine how the magnitudes are distributed for a realistic low Earth orbit (LEO) satellite, a Systems Toolkit (STK) model of the current active GPS constellation and six LEO satellites was built. The satellites were distributed in both near equatorial orbits and sun-synchronous orbits. An algorithm was developed to compute the phase angle and range between sensor and GPS satellites. The algorithm was then verified against the STK model.

For simplicity, each band was considered independently. Our simple radiometric model accounts for the bandwidth, spectral irradiance, and quantum efficiency as a function of wavelength for each band. For the two brightest bands, $m_R = 10$ yields a signal to noise ratio of greater than 30, while $m_I = 10$ yields a signal to noise ratio of 16. While neither of these independently can achieve the accuracy desired for attitude determination, either one can be easily detected. Thus, for this preliminary analysis, we assumed that if a satellite is in view with a magnitude of at least 10, it could be detected. The probability of detection becomes the probability of a satellite being in view with the required magnitude. Ideally, for celestial navigation we would make three near-simultaneous measurements. However, if multiple satellites of the necessary brightness are not in view, three observations separated in time can also be used to generate a navigation solution. A lower bound on the error is given by the angular measurement error of the star tracker.

The probability of detection for each of these cases is shown in Figure 15 by varying the apparent magnitude detection threshold. On the left is the probability for each band of detecting at least three satellites simultaneously. The I band gives the highest probability of detection, so the right side shows the probability of detecting at least 1, 2, or 3 satellites in the I band. In the actual sensor, the irradiance from each band would be cumulative, so the actual probability would be higher than shown. Using the threshold of $m_I = 10$ discussed above, it can be seen that three satellites can be detected about 20% of the time, while at least one satellite can almost always be detected. These results are very promising, and additional analysis will be performed in the next phase of the effort.

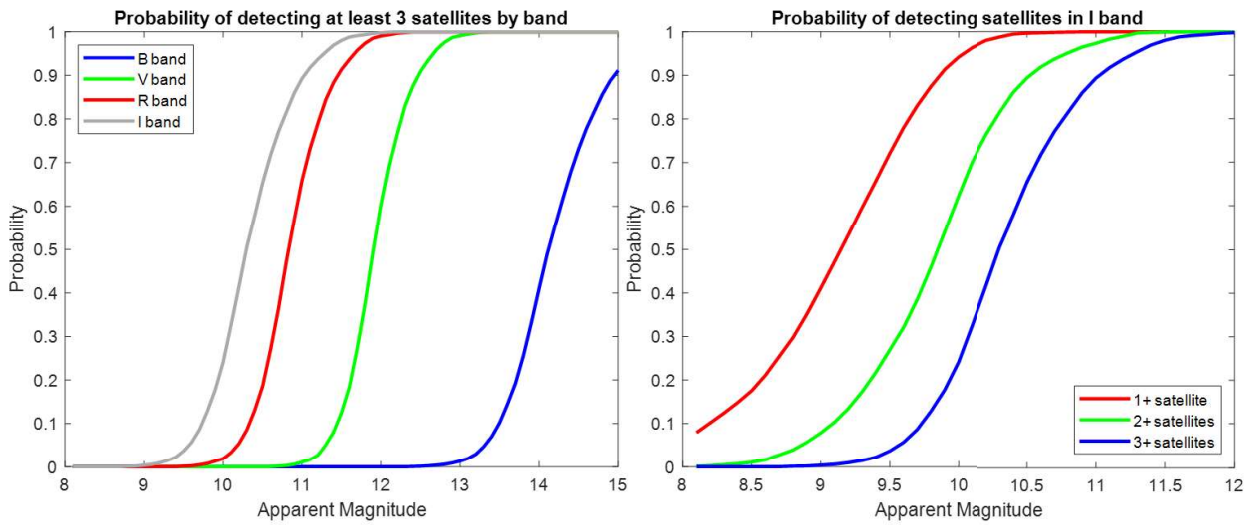

Figure 15. Probability of detection of GPS satellites from the star tracker

Summary

Based on results from the current study, it appears that design of a compact, extremely accurate star tracker is feasible. The optics design is manufacturable and is close to the SWaP goals established. A sensor suitable for use in the star tracker was identified. Performance modeling and systems engineering analysis indicates that the design will enable accurate location of stars down to the 10th magnitude, allowing operation over more than 99% of the celestial sphere. The design is robust against uniform thermal changes. Although the very aggressive thermal environment modeled yields some performance issues, the final design is anticipated to have almost an order of magnitude less heat dissipated. This reduction in dissipation is expected to yield sufficient accuracy for the star tracker mission.

In addition to functioning as a star tracker, this design appears to be able to detect GPS satellites with modest slewing of the platform. Additional analysis is warranted to assess the ability of the system to detect satellites without slewing or in the presence of platform motion. Analysis to provide an estimate of the accuracy that can be achieved with the sensor is also needed.

References

[1] C. M. Persons and G. A. Finney, "Polarimetric Space Situational Awareness using the Aero-Optical Prediction Tool," in *AMOS Technology Conference*, Maui, HI, 2021.

[2] M. Wegner and et. al., "The SIMBAD astronomical database," *Astonomy and Astrophysics Supplement Series,* vol. 143, pp. 9-22, 2000.

[3] H. F. Fliegel, L. F. Warner and F. J. Vrba, "Photometry of Global Positioning System Block II and IIA Satellites on Orbit," *Journal of Spacecraft and Rockets,* vol. 38, p. 609, 2001.

[4] T. E. Payne and S. A. Gregory, "Passive Radiometric Observations of Geosynchronous Satellites," in *IEEE Aerospace Conference Proceedings*, 2004.