Completion and Test of a Compact, Extremely Accurate Star Tracker

Stephen Fox, Greg A. Finney, Aaron M. Mebane
IERUS Technologies, Inc., 2904 Westcorp Blvd, Huntsville, AL 35805
Bijan Nemati
Tellus 1 Scientific, LLC., 57 Wiljoy Cir., Laceys Spring, AL 35754
Kevin Ludwick
University of Alabama in Huntsville, Center for Applied Optics, 301 Sparkman Dr, Huntsville, AL 35899

ABSTRACT

IERUS Technologies, Tellus 1 Scientific, 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 for an Extremely Accurate Star Tracker (EAST) that will enable the desired accuracy. Optical analysis showed the nominal design would provide near-diffraction limited performance. 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 for the host vehicle via celestial navigation. Similarly, it also has the potential to be used for space situational awareness. Last year at this conference we reported our progress on the design and simulation of the EAST concept. This year we report on the results of manufacturing and testing of the EAST system.

1. 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 difficult to measure, focal plane errors, and the dominant error source, optical distortion. 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, Inc., Tellus 1 Scientific, LLC, 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 (EAST) in a package with size, weight, and power requirements comparable to commercially available star trackers. Background on the design and initial simulations have been previously published. [1] A brief summary will be included here to provide context for the reader.

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 its impact on centroid accuracy 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 initial error budget, detector noise was limited to 37 mpix and FPA calibration accuracy to 21 mpix. The dominant error source in our error budget was field distortion errors at 67 mpix. The other error sources, along with a generous reserve, brought the total error budget to 107 mpix. This error budget was based on collection of stars with 10th magnitude stars (V-band).

2. STAR TRACKER FABRICATION

The star tracker concept design is shown in Fig. 1. 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

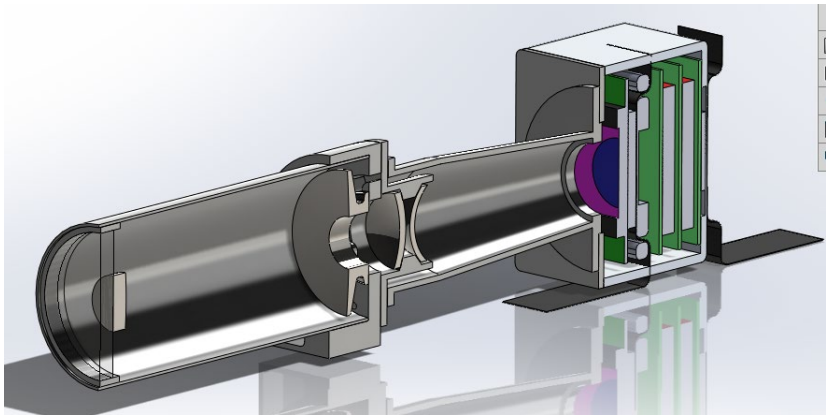

Fig. 1. Cutaway view of star tracker

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. Star trackers typically operate on stars of about 6^{th} magnitude or brighter (~4,000 stars). This consequently requires the field of view be large enough to collect sufficient stars in order to calculate its orientation.

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 10^{th} magnitude stars as the threshold (~300,000 stars), but extended the catalog to 12^{th} magnitude (1.3M stars) to improve matching. Most testing was done using a subset of the catalog pruned to the area of interest, but preliminary testing indicated that matching was successful with the full catalog with longer processing time (~6 s per frame on a laptop). With the pruned catalog, the image could be processed in <100 ms.

2.1 Focal Plane Array

The initial FPA selected for this effort was a 2,048 x 2,048 back side illuminated CMOS sensor. One important feature was this sensor had a radiation hardened version available. For cost efficiency, we planned to use the developer's kit provided by the FPA manufacturer as the hardware interface to the FPA, rather than develop our own FPA circuit board. When we tested the FPA in the developer's kit, we found there was a significant amount of noise. The vendor was able to trace the source of the noise to a power supply in the developer's kit, but they were unable to provide a version with the problem corrected. Early in the project we purchased a commercial astronomy camera that we used as a surrogate while awaiting our primary sensor. This surrogate, which employed a Sony IMX455 sensor, became the sensor used in our prototype. The dimensions of the primary sensor and our surrogate were similar, but the IMX455, which is 9,576 x 6,388, had 3.76 μm pixels instead of the 10 μm pixels of our primary sensor. The higher resolution, along with the lower noise, was expected to yield some improvement in the accuracy of our centroiding results.

2.2 Telescope

IERUS has the capability to fabricate the telescope, but the facilities were dedicated to a production project at the time of telescope manufacture. As a result, the manufacture of the telescope was outsourced. The vendor

manufactured the reflective and transmissive optical elements and integrated them into the telescope. The telescope was not diffraction limited, as shown in Fig. 2. The point spread function shows some symptoms of astigmatism and coma. The width in the narrow direction is close to expected for diffraction limited, but in the long direction is almost a factor of 2 wider than expected from diffraction. The elongated shape was addressed in the algorithm for centroiding the stars and is discussed below.

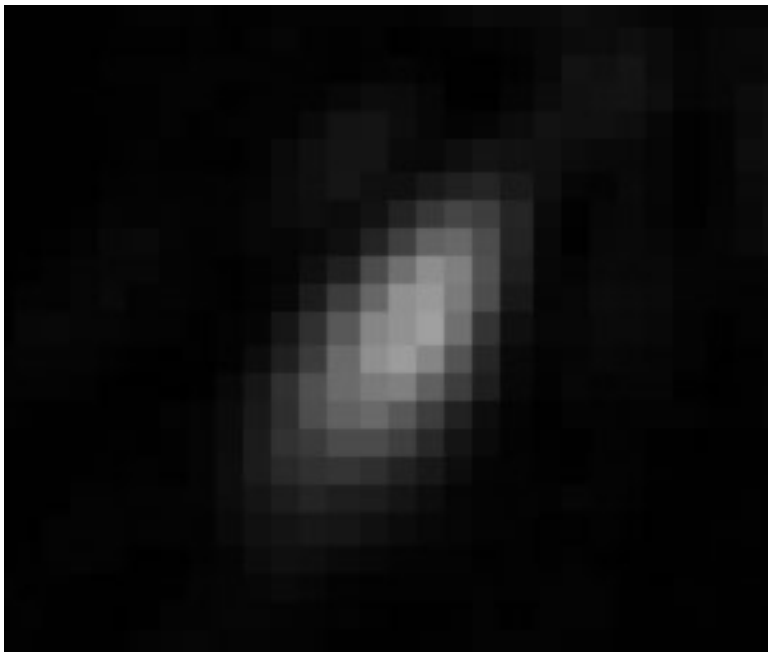

Fig. 2. Off-axis PSF of the completed telescope

2.3 Environmental Enclosure

All calibration was conducted in laboratory conditions, but the customer planned to conduct collections in an outdoor environment with temperatures that could drop well below freezing. Resources available for prototype development did not allow for an athermalized design, so an environmental enclosure was developed that would maintain the telescope at a suitable constant temperature. The enclosure included thermal insulation, patch heaters, temperature sensors, an optical grade window, and electrical interfaces. The design of the enclosure and a photograph of the telescope are shown in Fig. 3.

Fig. 3. Overall thermal enclosure (top) as designed and (bottom) as built

Once concern for the enclosure was deformation of the window under thermal gradient. Although our in-house optical modeling program discussed below can model the effect of window deformation, preliminary analysis showed it was unnecessary to model at that level of fidelity. The maximum temperature gradient that the enclosure supported was 35°C, and with that gradient the estimated radius of curvature of the window is 314 m. Assuming the front and back surfaces are concentric, the change in focal length was <30 nm, so the impact on the PSF is negligible. During field testing, the thermal control panel indicated the temperature remained constant within ±0.5°C, with an atmospheric low temperature during the test of ~5°C.

3. 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. The plate scale can be calibrated periodically on orbit using measurements of star positions. The laboratory 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 quantum efficiency and the position of the pixel.

Some additional detail for the calibration approach is shown in Fig. 4. 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 is 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, as illustrated in Fig. 5. The difference between the two values is used to calibrate the pixel location. The challenge is to determine the error given all the surrounding pixels also have unknown positional error.

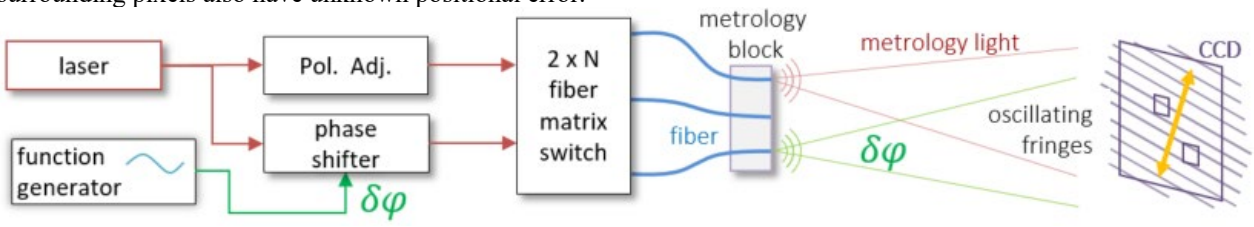

Fig. 4: Pixel metrology approach.

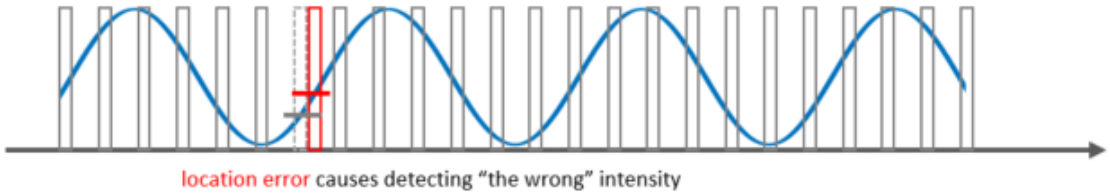

Fig. 5: Basic concept of pixel metrology with calibrated fringes. The measured intensity is compared with the expected intensity and the difference is used to calibrate the pixel location.

The metrology block and sample fringes are shown in Fig. 6 below. The metrology block 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.

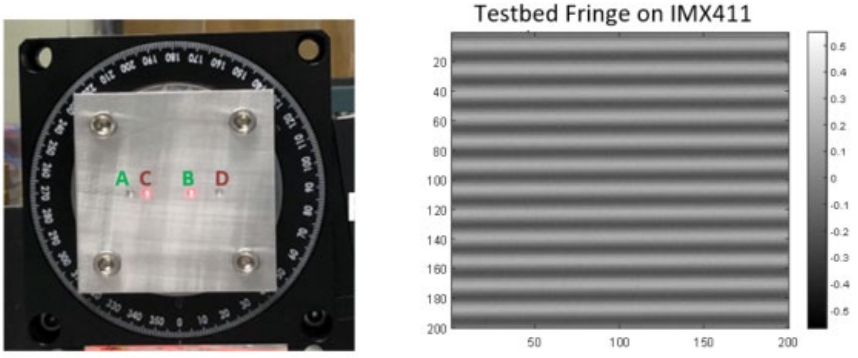

Fig. 6: Metrology block (left) and sample fringes from a large format detector (right).

Using two fibers to create interference fringes on the camera will not create straight, Young's slit fringes, but instead it will create hyperbolic ones. Fringes at the center of 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 Fig. 7 shows a comparison of hyperbolic fringes, compared with 'ideal', straight fringes from infinite slits. For a large detector like the Sony IMX 455, which has 60 million pixels, each 3.76 microns wide, by the corner there is a significant difference between the two fringes. This difference is less than a pixel, but much larger than our goal of 1/100th of a pixel.

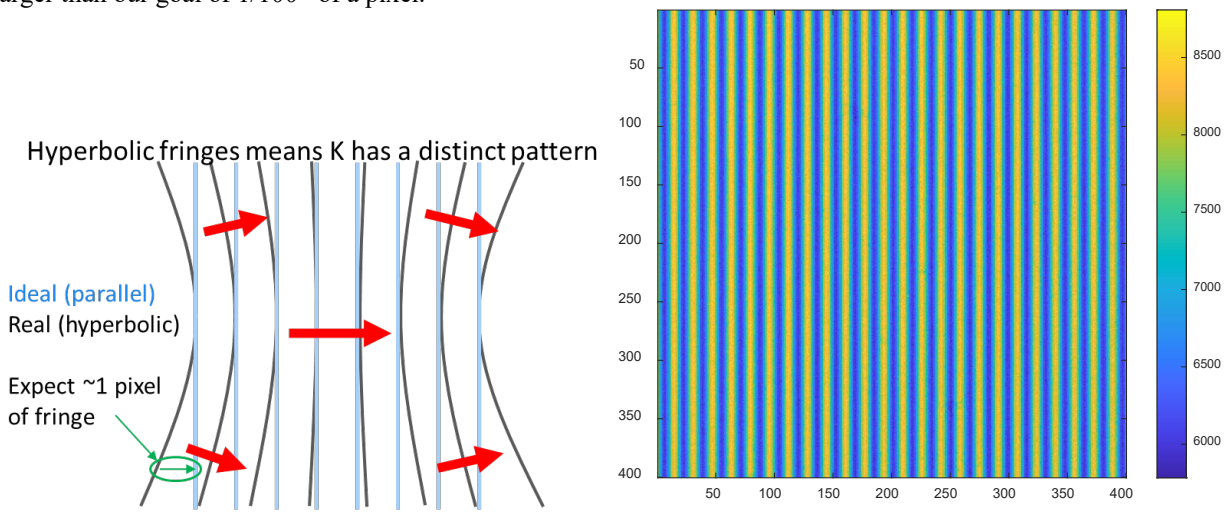

Fig. 7: Fringes from a pair of fibers will have a hyperbolic pattern on the detector. One the left, the arrows show the wavenumber vector direction. On the right is a typical 400x400 pixel segment from the IMX411.

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. Details between the simulation and data vary, including an overall sign difference between the data and simulation for K_x, and the patterns in the data not being centered due to the system alignment. The measured data can be used to solve for these geometric differences.

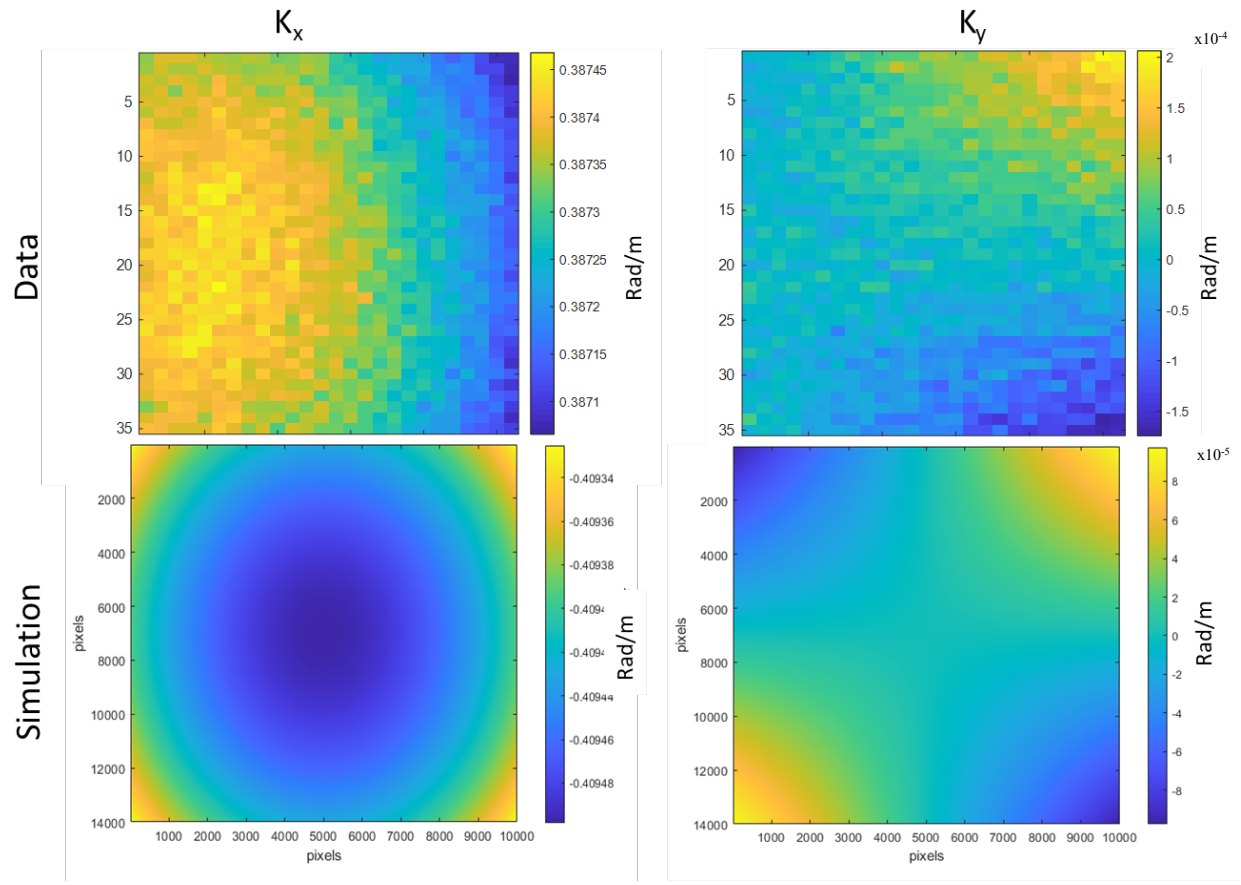

Fig. 8. Comparison of simulation and laboratory data

Measurement challenges we resolved include overcoming physical error sources. The fibers are very sensitive thermal and vibration transducers. The testbed was enclosed in a portable soft wall clean room to reduce these environmental effects. When we reduced the data from early FPA calibration collections, we found there was a high spatial frequency, non-stationary pattern in the pixel displacements overlaid on the low frequency, stationary pattern. Process of elimination indicated this non-stationary component was due to atmospheric turbulence. To mitigate this effect, we allowed the calibration laboratory (within the soft wall cleanroom) to sit undisturbed for 3 days before collecting data. This procedure significantly reduced the non-stationary component. The final pixel position errors that we collected are shown in Fig. 9. This figure shows that the peak-valley pixel position error is 25-30 mpix. The high frequency pattern has been largely eliminated, except for very high spatial frequency pattern that is likely arising from photon noise. It is worth noting that these slowly varying errors would generally have been captured by the field distortion correction algorithm discussed below.

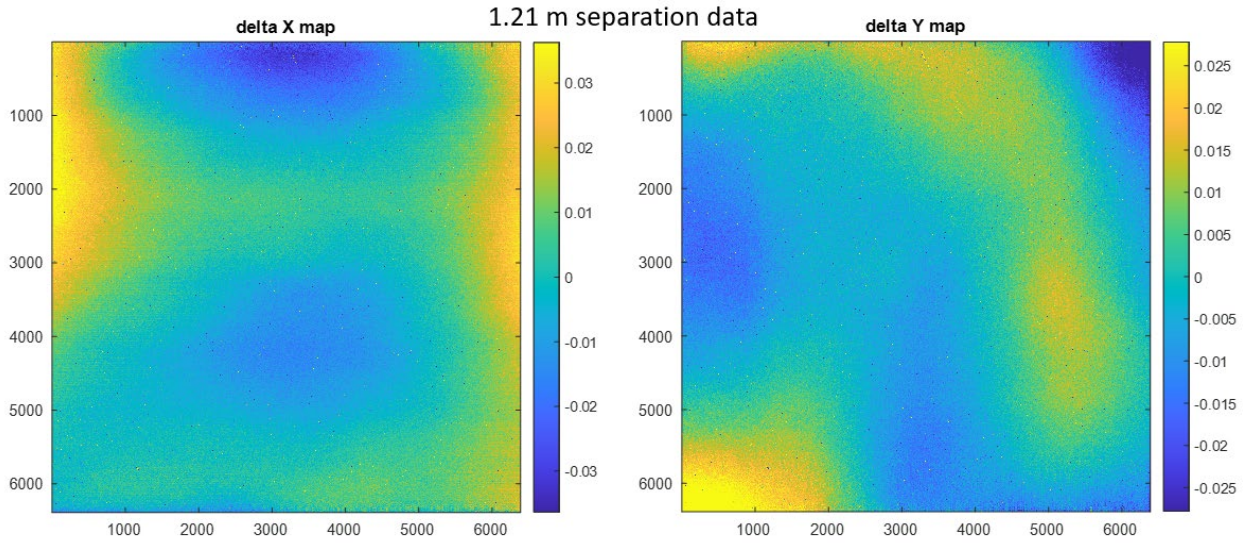

Fig. 9 Final pixel position error maps for the IMX455 in pixel units for x (left) and y (right)

4. FIELD DISTORTION CORRECTION

Field distortion correction was one of the more challenging aspects of the EAST project. The initial design predicted over 50 pixels of distortion at the edge of the field of view, and our error budget was only 67 mpix, requiring almost 1,000X reduction. The analysis of the design confirmed that the distortions would be slowly varying, including predictions from a Monte Carlo analysis over the anticipated manufacturing tolerances. Several approaches for measuring distortion were evaluated, including translation of a point source, imaging a target that was a grid of point sources, and applying pitch/yaw to the telescope using rotation stages. The final approach was to manufacture a diffractive optical element that would produce an array of superimposed collimated beams, with each beam tilted at a different angle. The diffractive element was designed by Tellus 1 and produced by UAH using a lithographic process. Specifications for transmitted wavefront error and spacing error were provided. When imaged, this array produces an array of spots on a regular grid, as shown in Fig. 10. This grid of psuedo-stars is used is used to measure the distortion. The distortion measurements are fit to a polynomial to estimate the distortion correction at any position in the FPA. The telescope and sensor are mounted on rotation stages to provide two angular degrees of freedom and are aligned to ensure the axis of rotation is through the entrance pupil. This ensures that identical portions of the wavefront are sampled at each angle.

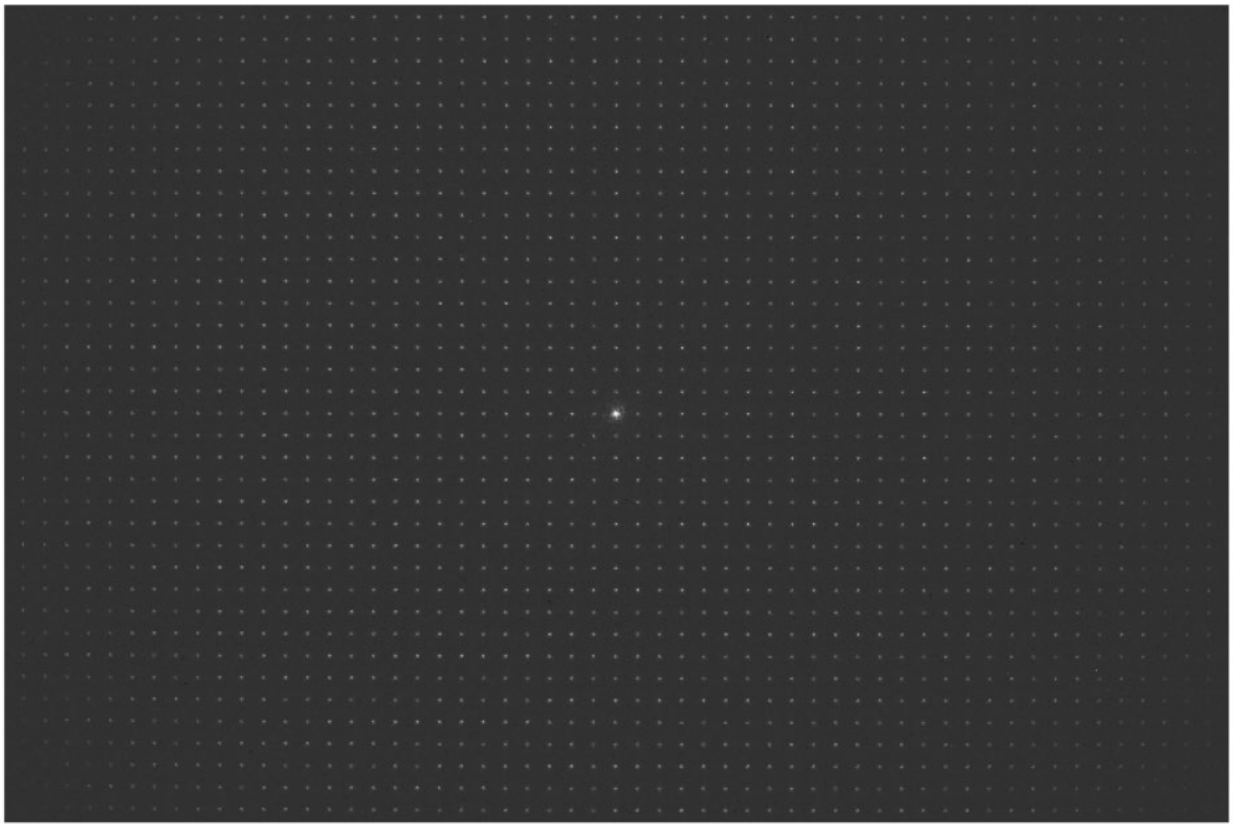

Fig. 10. Logarithmic image of the star array produced by the diffractive element

Fig. 11 shows the distortion pattern from an inferred regular grid and the residual after correction. The centroiding algorithm that seemed to work best overall was an "elliptic" Gaussian fit (i.e., the standard deviation in x and y and the angle of orientation of the Gaussian on the xy plane were fitted as nuisance parameters). The minimization of vibrations in the lab allowed for long exposure times at 0dB gain. Increasing the exposure time and the number of frames used for calibration allowed us to collect more photoelectrons and thus improve the SNR and the calibration.

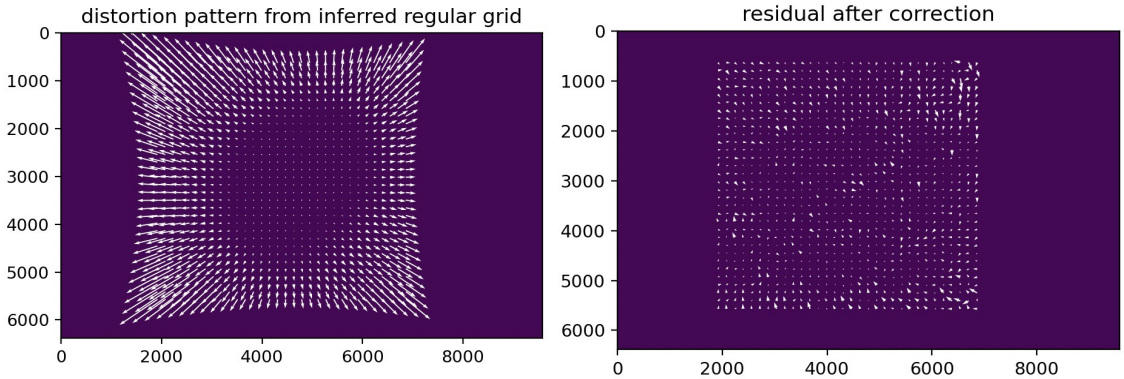

Fig. 11. Quiver plots of distortion before and after correction, with arrows in the right plot magnified by 10X relative to left plot

We utilized new test frames (not involved in the polynomial fit) to estimate the RMS residual error after applying the correction from the polynomial fit to the test frames. The test frames were taken when the telescope was at a different attitude compared to the attitudes when the calibration frames were taken. To determine the required order of the fitting polynomial we tracked RMS error as a function of polynomial order. Fig. 12 shows a sample of this data. After approximately order 18, the RMS error begins to rise again due to fitting of noise artifacts.

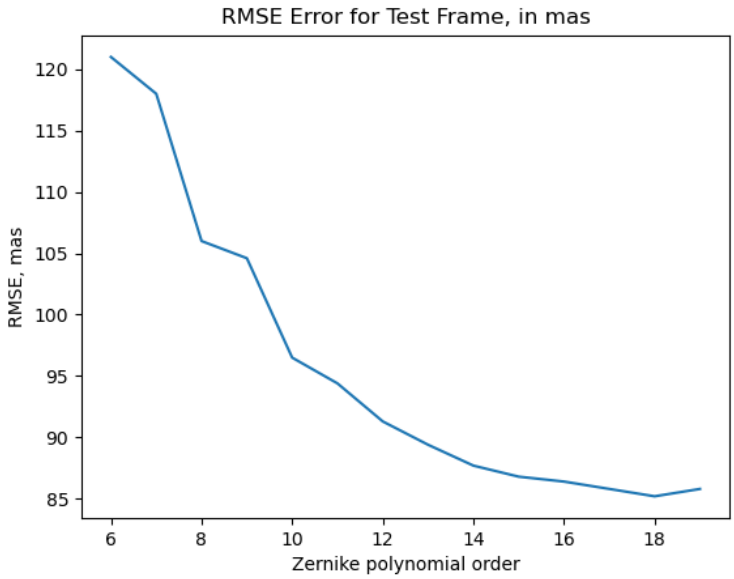

Fig. 12. RMS Error vs. Polynomial Order

In general, our distortion calibration meets the error budget, and we have found it to be a powerful calibration tool even when the FPA calibration is not employed.

5. IMAGING PERFORMANCE PREDICTIONS

A complex scene containing an actual stellar distribution was developed for use with our in-house sensor model, the Aero-Optical Prediction Tool (AerOPT). [2] For this project, we generated a physically accurate scene as viewed from the star tracker above the Earth's atmosphere. The scene radiometric signatures were only approximately correct, as there was not data available for all stars for the full spectrum used by our sensor. 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 field of view in both horizontal and vertical directions, which will allow for variations as we test different aspects of the algorithms. The scene addresses the following elements:

- ✓ 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)
- ✓ Interstellar dust (below noise floor)

The sensor model included effects that are representative of the actual selected sensor, with some simplifying approximations:

- ✓ Uniform responsivity and analog to digital conversion
- ✓ Thermal noise, read noise, and shot noise
- ✓ Linear response up to saturation
- ✓ Field distortion
- ✓ Optical aberrations

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. Our primary star catalog was the GNC v1.1 Star Catalog published by the US Naval Observatory. [3] In some cases, signature data was taken from the SIMBAD online star catalog. [4]

For comparison, a portion of the resulting AerOPT image at 1 s integration time with a few stars annotated is shown in Fig. 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, Fig. 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 top row is scaled to show the full range of counts, while the bottom row is scaled to show the dimmer stars (144 max for the first 3 images and 50 counts max for the fourth image). The brightest star is Alcyone (eta Tau, $m_V = 2.87$), which has 6 pixels that are saturated. The dimmest star annotated, Cl Melotte 22 1713 ($m_V = 11.95$), has only 86 counts. The background, due almost entirely to dark current, has a mean of 10 counts and a standard deviation of 2.7 counts.

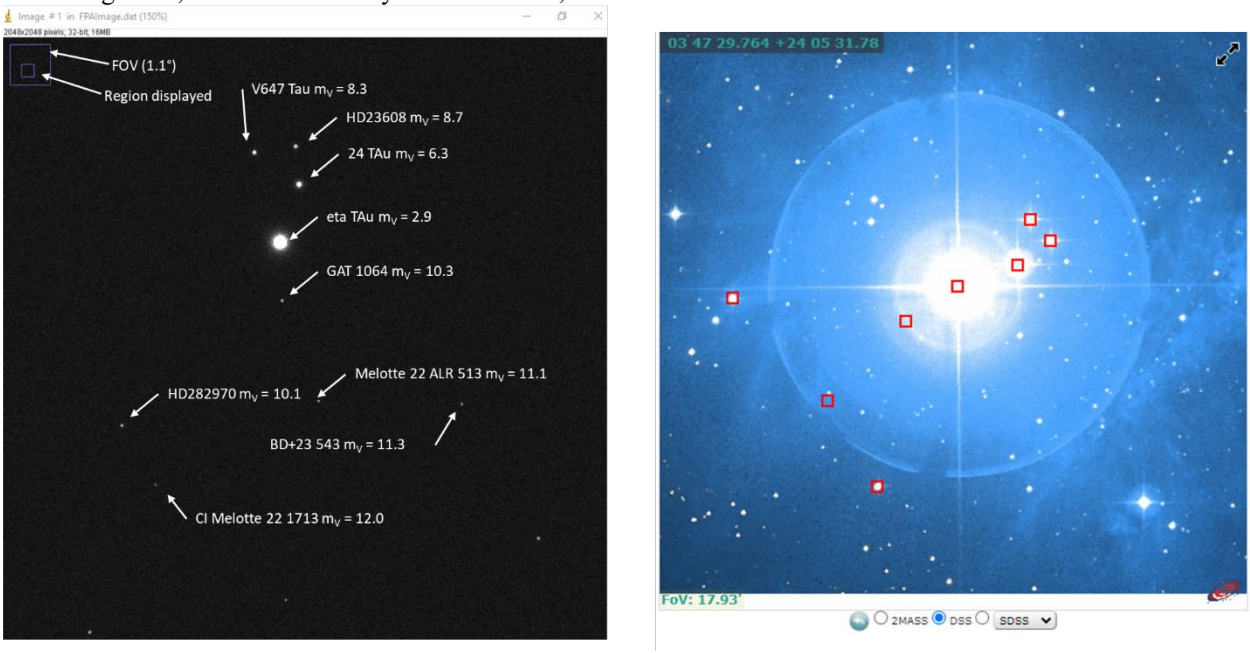

Fig. 13. Portion of AerOPT image and DSS image

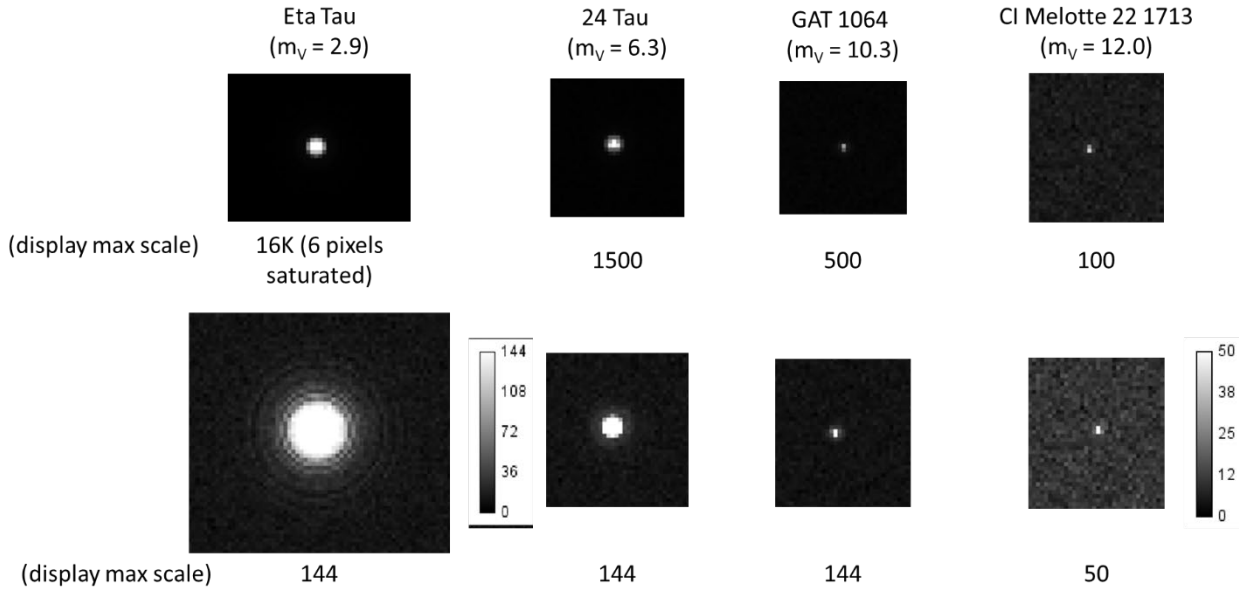

Fig. 14. AerOPT image chips showing stars at different contrasts

The AerOPT scene was used for development of the image processing and orientation solution software while hardware was being designed and fabricated.

6. FIELD TESTING

Field testing was conducted in December 2023 at a rural location outside of Huntsville, AL. Data collections were planned that featured the Pleiades, Polaris, Jupiter, and Orion. When attempting to capture Polaris, an operator error

resulting in collecting images centered on 2 UMi instead of Polaris. Because the orientation solution was not fully functioning during the field test, this error was not recognized until after post processing the data. Other collections were made as planned. An image of the EAST system during the collection and an image of Jupiter are shown in Fig. 15.

Fig. 15. (left) EAST system set up for field collection. (right) Image of Jupiter collected by EAST

It is apparent by comparing the images from the field test with the images of pseudo-stars from the lab tests that focus likely shifted during the final build or during transport. We were not able to perform a focus optimization or realignment in the field. An off axis point spread function collected during lab tests is shown in Fig. 2, while a representative image (near boresight) from the field test is shown in Fig. 16. The PSF from the field collection is larger by approximately a factor of 2 in the long direction and 2.5 in the short direction. This also reduced our achievable SNR by about a factor of 5. Despite the reduced SNR, the software had no problem detecting stars dimmer than magnitude 10. The dimmest star used in our orientation solution was magnitude 11.2, and the software was able to detect additional stars below this value.

The first test we conducted to assess the accuracy of the astrometry was to measure the angular distance between the same star pairs in images taken at two different attitudes. The centroids of the star pairs were corrected for distortion and manually matched, and we used a least-squares fit to solve for the focal plane coordinate transformation between frames. Then we computed the residual difference between the positions of stars in each pair after transforming the coordinates from the second image to match the first image. The per-star RMS error for compared frames of 1-second exposure compared well with the error budget. Several frames were taken near the celestial pole at fairly low elevation (~35°), where refractivity of the atmosphere introduces significant error. In the images taken near zenith, refractivity error is reduced, but sidereal motion introduces error. While the 1-sec exposure time reduced the sidereal motion error, it appears it was not eliminated completely. Given our southern location, we could not collect images that were both near zenith and near the celestial pole.

A challenge we encountered with our processing software was the discrepancy in brightness as reported in the star catalog and as measured by the sensor. Our algorithm used the brightness order as part of the star matching process. It is tolerant to some differences in the brightness order between what is in the catalog and what is in the scene, but the discrepancy we observed resulted in a small number of matches, thus a relatively large error in the boresight solution. Given the initial low-accuracy solution, an iterative approach could have been used to match the stars in the scene and in the catalog, thereby improving the accuracy. However, resources did not allow us to make the

modification to our software. Instead, in post processing the analyst manually updated the brightness order of the stars found in the scene to match the stars in the catalog, which yielded an equivalent result.

As an example of the field test results, the scene containing 2 UMi (right ascension $01^h\ 08^m\ 44.88005^s$, declination 86° 15′ 25.5240″, mV = 4.2) is shown in Fig. 16. The image is in $\log_{10}$(counts), so the maximum corresponds to $10^{4.26} \sim$ 18k counts. The image is masked to a 1° wide field of view to match the field distortion correction. As noted above, the processing software was able to identify the boresight with a relatively large error. From there, out of the set of stars that the software identified, 16 stars were manually matched in brightness order for additional processing. Half of the stars were used for finding the orientation solution, and half were used for testing the accuracy of the solution. To compute the error, the ideal location of the star in the camera frame was computed, and the total angular error between the distortion corrected location and the ideal location was computed. The mean error for the test stars was reduced by almost a factor of 2 by incorporating the correct brightness order, allowing for additional matches. Resources did not allow us to determine the source of the remaining error as part of this effort, although we have identified several areas to investigate.

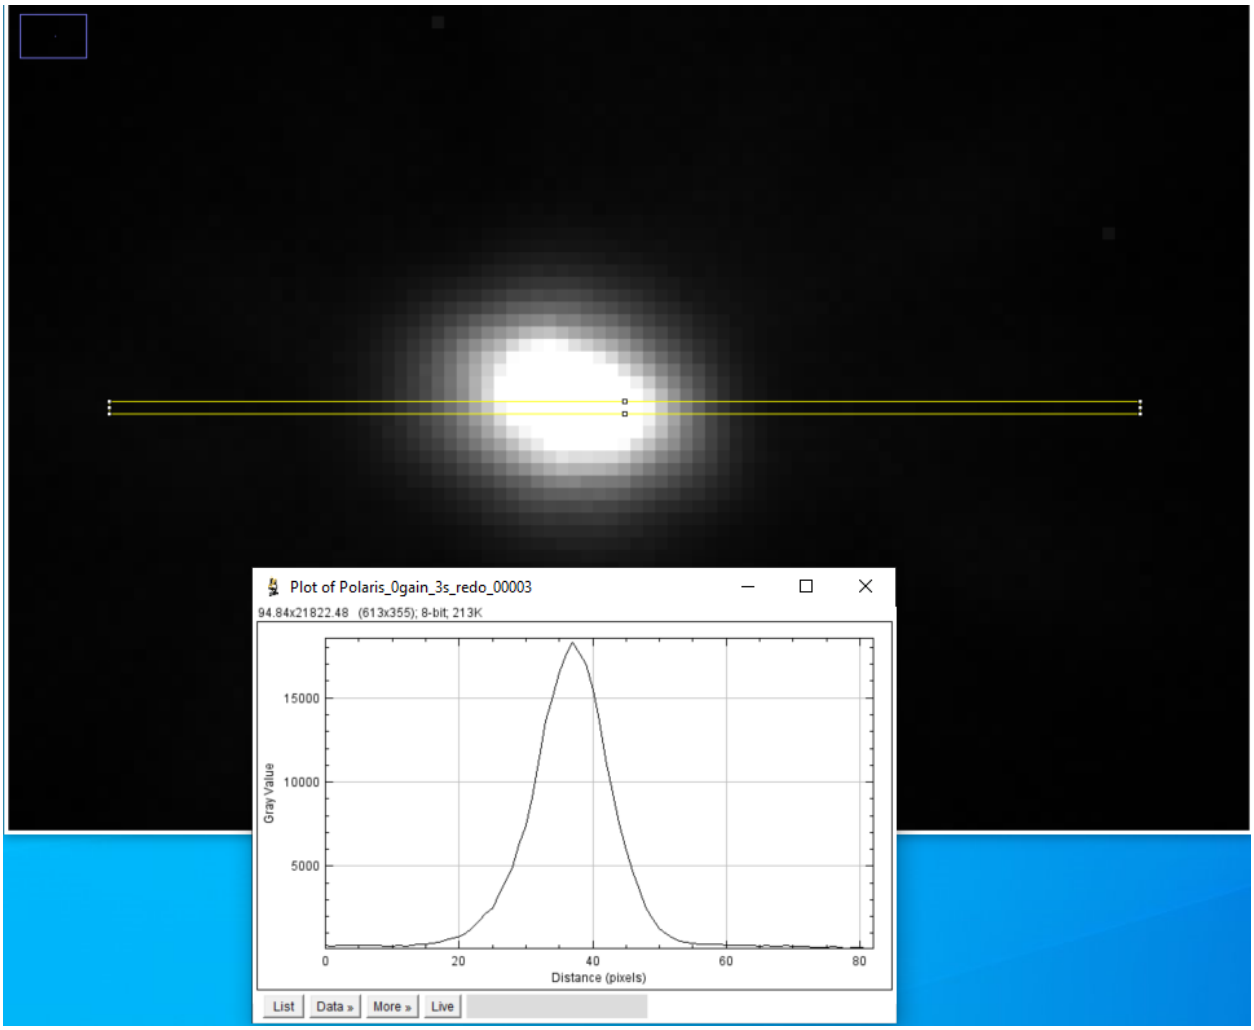

Fig. 16 Image chip of 2 UMi with cross section plot.

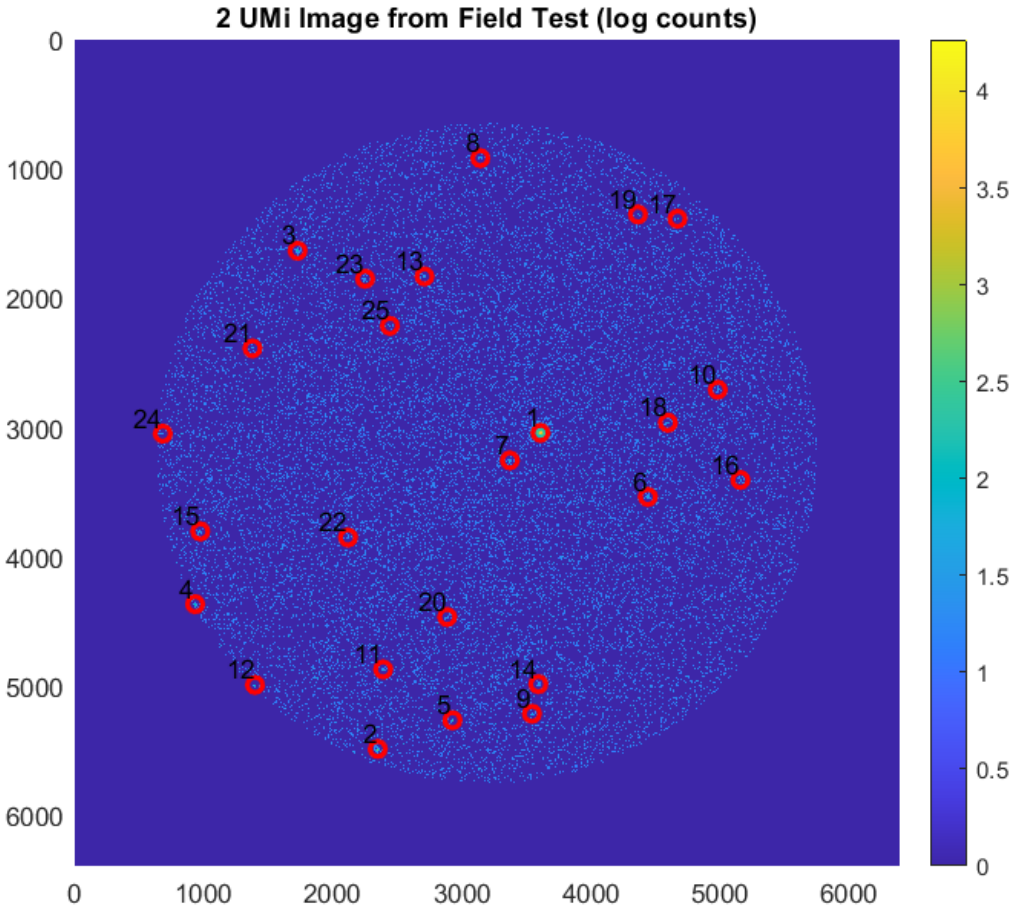

Fig. 17. Image containing 2 UMi from field test (log_{10} counts). Markers show stars identified by the processing software

7. SUMMARY

Based on results from the current study, it appears that design and manufacture of a compact, extremely accurate star tracker is feasible. The optics design was manufactured and is close to the SWaP goals established. A sensor suitable for use in the star tracker was identified, although testing had to be performed with a surrogate. Test results showed the sensor can detect stars down to the 10th magnitude, allowing operation over more than 99% of the celestial sphere. Accuracy during field did not meet the very aggressive goals, but several areas for investigation were identified, including defocus/optical alignment, atmospheric refractivity, sidereal motion, and spectral responsivity. Of these, sidereal motion and atmospheric refractivity are not relevant in the intended space environment but are present due to the field-testing environment.

8. REFERENCES

[1] G. A. Finney, S. Fox, B. Nemati and P. J. Reardon, "Extremely Accurate Star Tracker for Celestial Navigation," in *AMOS Technologies Conference*, Wailei, Maui, HI, 2023.

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

[3] US Naval Observatory, "The International Celestial Reference Frame (ICRF) Overview | USNO," [Online]. Available: https://crf.usno.navy.mil/icrs. [Accessed 23 07 2023].

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