On-orbit Results from an Ultra-low SWaP Black Silicon Star Tracker

Darren Rowen, Alexander Utter, Garrett Kinum, Hannah Weiher, Andy Wu, Richard Dolphus, Eddson Alcid,
Denise Galindo, Carlos Beltran, Ariel Berman, Michael Mellick
The Aerospace Corporation
2310 E El Segundo Blvd, El Segundo, CA 90245; 310-336-3633
Darren.Rowen@aero.org

Martin Pralle
SiOnyx, LLC.
100 Cummings Center, Beverly, MA 01915; 978-922-0684
mupralle@sionyx.com

ABSTRACT

In August 2019, two 1.5U AeroCube-10 satellites built by The Aerospace Corporation were deployed from a Cygnus resupply spacecraft. Each of the satellites has two star trackers which are many times smaller than commercially available alternatives. The significant size reduction is enabled by the SiOnyx XQE-0920 sensor which offers dramatically improved visible and near-infrared sensitivity in an uncooled CMOS platform. This allows the use of a smaller-aperture lens than traditionally used in small form factor star trackers, while maintaining the ability to detect stars of magnitude 5. The reduced volume enables innovative system engineering trades such as forgoing star tracker baffles, and instead flying multiple sensors on the same spacecraft to combat stray light by using the spacecraft body itself as a shield. The additional interior volume made available also enables more capable missions in smaller CubeSat form factors.

On-orbit results are presented showing angular accuracy and solution availability statistics as a function of angular rotation rate. A calibration technique to compensate for optical distortion is also presented, which enables the use of a low-cost COTS lens with a wide field of view. Despite the extremely small volume, the star tracking performance is comparable to units many times larger.

BACKGROUND

Two 10x10x15 cm (1.5U) AeroCube-10 satellites (Figure 1) were deployed from the Cygnus Northrop Grumman 11 (NG-11) resupply spacecraft. The AeroCube-10 spacecraft contained experiments to study the Earth's atmosphere, demonstrate proximity operations, and measure the degradation of solar cells in orbit. To support these experiments, each spacecraft has two miniature star trackers with a high-performance black silicon focal plane made by SiOnyx, LLC. Available volume for the star trackers was extremely limited for this 1.5U CubeSat mission. This motivated the adoption of a much smaller and wider field of view lens than is typically used for star tracking applications.
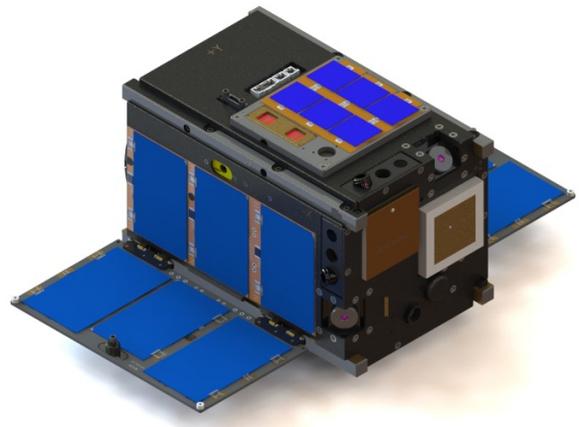

Figure 1: AeroCube-10 Spacecraft CAD Rendering

HARDWARE

SiOnyx XQE-0920 CMOS Image Sensor

The star trackers employed in this mission utilized SiOnyx XQE-0920 silicon CMOS image sensors. The SiOnyx XQE family of enhanced sensors utilize the company's proprietary nano/microtexturing technology (also known as Black Silicon) that modifies absorption characteristics of silicon in a thin layer. Because the SiOnyx XQE-0920 Image Sensor can be manufactured using mostly existing commercial processes, the per unit cost is hundreds of dollars per sensor, much lower than sensors made with specialty processes. A traditional CMOS image sensor uses a physically thin layer of silicon (less than 7µm) as the optically active layer. This results in poor absorption for longer wavelengths and decreased sensitivity for the near infrared (wavelengths between 800 nm and 1200 nm). Silicon is an indirect bandgap semiconductor material that requires the combination of a photon and a phonon to absorb light and generate an electron hole pair. As the wavelength of light approaches the band edge (i.e. longer wavelengths) absorption becomes increasingly dependent on the presence of phonon with a specific momentum. Therefore, the probability for absorption decreases, and photons penetrate deeper into the silicon substrate before being absorbed.[1] The XQE-0920 also uses a thin layer of silicon (~ 7µm) for the optically active layer but traps longer wavelength light within that layer and thereby improves the absorption and quantum efficiency. Increased sensitivity in the near infrared enables enhanced star light imaging by enabling the capture of photons which are invisible to other CMOS sensors.
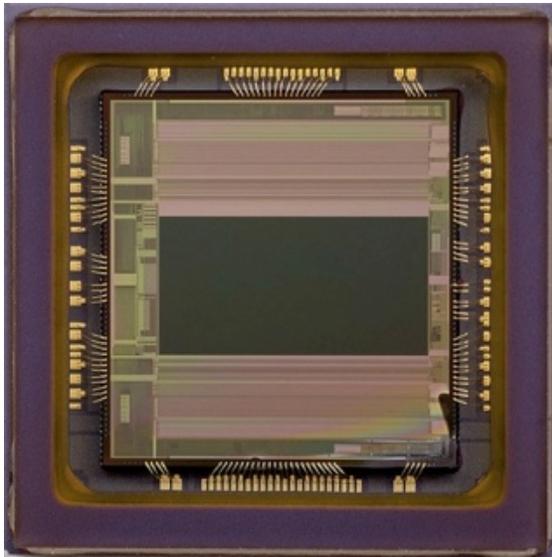

Figure 2: SiOnyx XQE-0920 Image Sensor

The XQE-0920 is a 720p (1280x720 pixels) resolution sensor with 5.6 micron pixel pitch, and ½ inch optical format. The sensor is designed with a columnwise ADC (Analog to Digital Converter) focal plane array architecture for the best possible read noise performance. The pixel design is based on the 4T (four transistor) pixel architecture.[2] The 4T architecture utilizes four transistors in the pixel to control overall device characteristics. Specifically, the addition of the fourth transistor enables correlated double sampling (CDS) of the pixel. CDS reduces the overall read noise of the sensor by subtracting off system noise. The XQE-0920 has a read noise of 1.8 electrons/pixel at maximum column amplifier gain, a room temperature dark current of 24 electrons/pixel/second, linear full well capacity of 24,000 electrons, and a native dynamic range of 72 dB.

In addition to having superior near infrared sensitivity over traditional CMOS image sensors[13], the XQE-0920 has high quantum efficiency over the visible range as well (see Figure 3). The broad spectral response (from 400 nm to 1200 nm) was independently verified by US Army Nightvision Labs and enables detection of very faint starlight.[8]
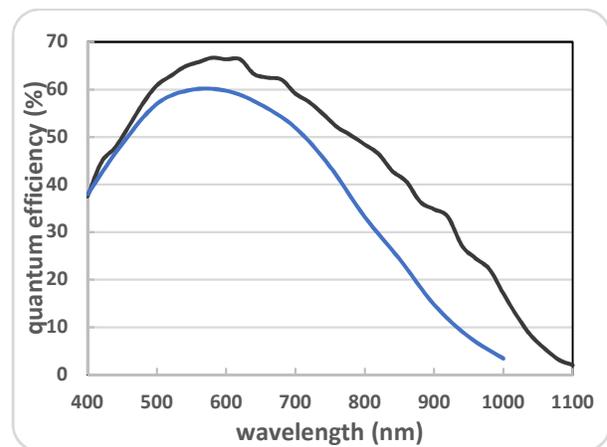

Figure 3: Quantum Efficiency of XQE-0920 (Black) vs sCMOS (Blue)

Silicon-On-Insulator Architecture and Deep Trench Isolation

Broad band spectral sensitivity is important, but there are two other critical device elements of the XQE-0920 for star tracking applications. These elements are the presence of Deep Trench Isolation (DTI) and a silicon-on-oxide (SOI) architecture. Figure 4 illustrates the basic architecture of the pixels in the 0920 sensor.
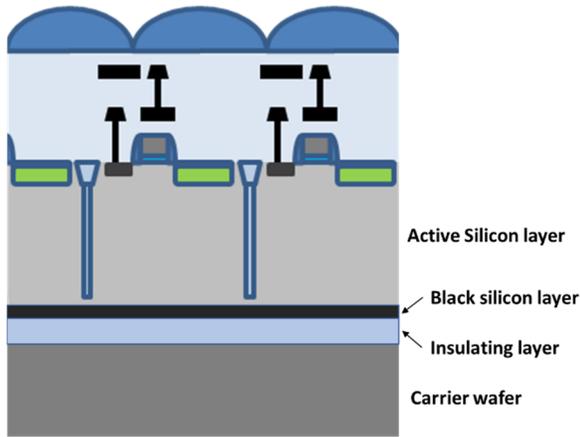
Figure 4: Schematic of the XQE-0920 Pixel Layout

Deep trench isolation is accomplished by placing optical and electrical barriers between each pixel. A physical trench is dug into the active silicon layer and then passivated and filled to avoid electrical or mechanical issues in the final image sensor (Figure 5). This is critical for performance because of the light trapping and extended optical path lengths that lead to higher near infrared sensitivity. Without the DTI the modulation transfer function (MTF), or sharpness, of the image at longer wavelengths is severely degraded. Figure 6 and Figure 7 show the difference in MTF for images taken with 940 nm illumination from a sensor without DTI and one with DTI, respectively. Image sharpness is much improved with DTI, and this provides a variety of benefits for star tracking. One significant benefit is the ability to distinguish two stars that are closely spaced. With a wider field of view, this is even more valuable because the effective angle covered by each pixel is larger and thus covers a larger area of the sky.
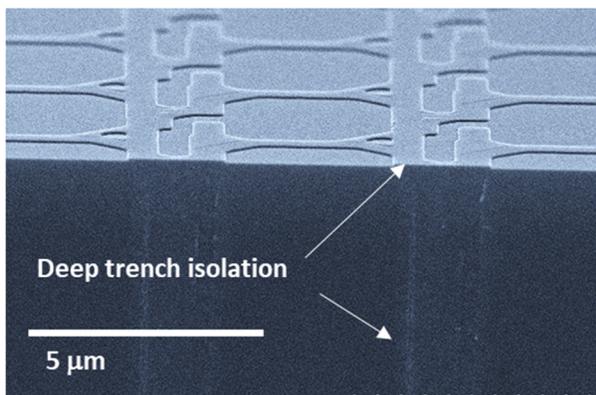

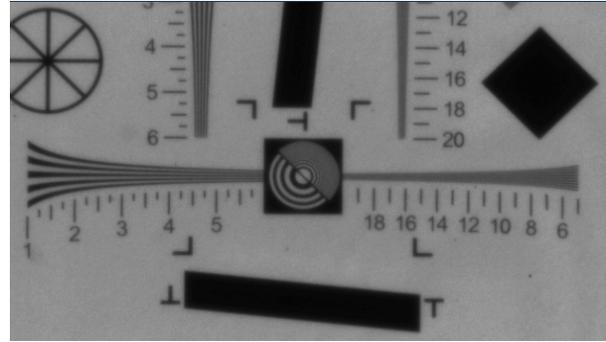**Figure 5: Scanning Electron Micrograph of the XQE-0920 Image Sensor**
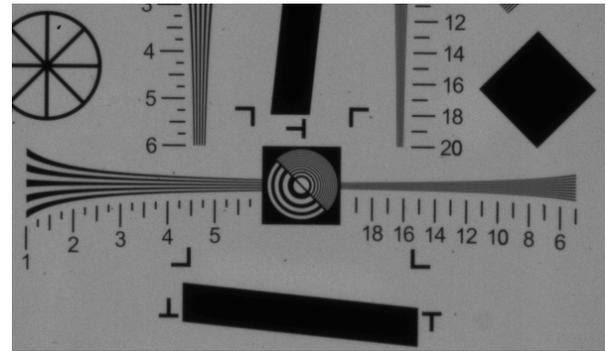

Figure 6: Image Taken Without DTI

Figure 7: Image Taken With DTI

CMOS Image sensors used in space-based applications are susceptible to radiation damage that can degrade performance or destroy the device in question. For an image sensor, the likelihood of radiation interacting with the device is correlated to the thickness of the device layer. In the case of the XQE-0920 architecture, the active silicon layer is isolated from the rest of the silicon wafer by an insulating layer – which is termed a silicon-on-insulator structure. Because of this architecture the XQE-0920 is more radiation tolerant than a traditional CMOS image sensor that utilizes an active silicon epitaxial layer directly grown onto a silicon carrier substrate.

Camera Mechanical Assembly

The star tracker subsystem consists of a printed circuit board (PCB) connected by flat ribbon cables to two camera sub-assemblies. The main camera board, shown in blue in Figure 8, sits in a stack with the other bus and payload electronics. The camera sub-assemblies are in various locations around the perimeter of the spacecraft body, pointing in different directions. The two camera sub-assemblies are shown in Figure 8 in black. Note that there are no light baffles in this design. Each camera sub-assembly (shown in Figure 9) contains a lens, mounting hardware, and an image sensor daughtercard on which the SiOnyx XQE-0920 focal plane is mounted. Up to five image sensors can be connected to the main

camera board using 39-pin fine-pitch flexible flat cables, which provide adequate cable-routing flexibility. The volume for the star trackers is computed for the various sub-components. For the star tracker sub-assembly shown in Figure 9, the volume, computed via enveloping rectangular solid, is 14.74 cm^3. The main camera board volume is 55.51 cm^3; however, this board can control up to five cameras that can be either star trackers or payload cameras. On the AeroCube-10 mission, additional payload cameras exist, so this camera main board services both the star trackers and payload cameras. The main camera board has additional features to output digital data for a laser communication transmitter that is not utilized in the AeroCube-10 configuration. Given the multi-purpose application of the camera main board, it is more valuable to focus on the volume for the camera sub-assembly. In comparison, one of the smallest commercial industry star trackers is the Berlin Space Technologies ST200 which has a volume of 34.2 cm^3, over 2.3 times larger than the AeroCube-10 star tracker's camera sub-assembly.[7] Many other CubeSat scale star trackers are multiple times larger yet.

The commercial off the shelf (COTS) lens used in this design is the Marshall USA V-4406.3-2.0-IRC-LP-5MP which has an M12 mount, 6.3mm focal length, f/2.0 aperture, and provides an effective diagonal field of view 61.1 degrees for the star tracker. With a retail price of $29, this lens was not originally designed with a space application in mind. The lens is the primary driver for the star tracker's camera sub-assembly volume. While larger lenses offer improved solution accuracy, the AeroCube-10 mission was highly volume constrained and would not have been able to accommodate a larger lens.
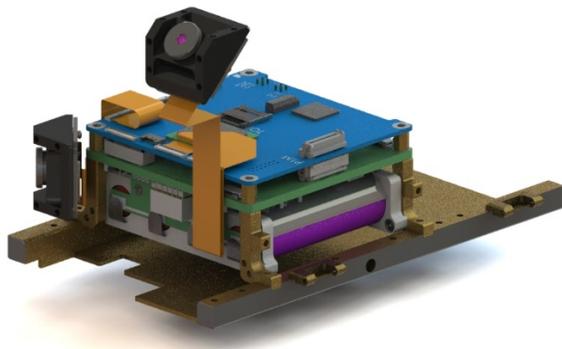

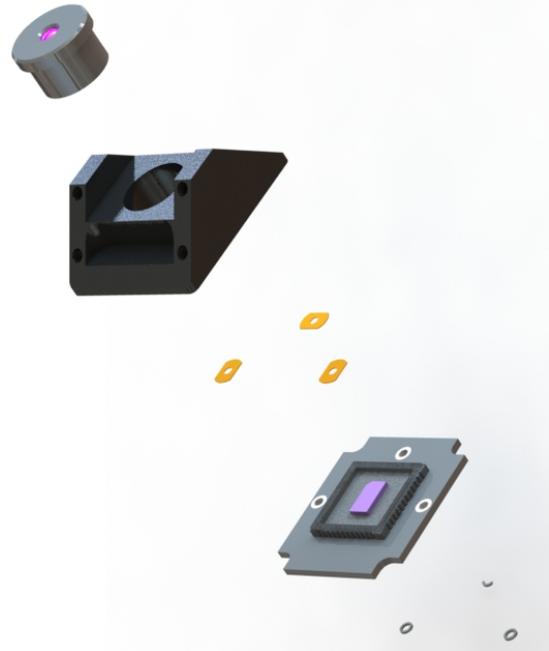**Figure 8: AeroCube-10 Star Trackers, Avionics and Harnessing**

Figure 9: AeroCube-10 Star Tracker Sub-Assembly Exploded View

Camera Board and Daughterboard Electrical Design

The AeroCube Camera system is comprised of a main camera board and a daughterboard for each sensor. The main camera board provides a powerful image acquisition platform with 5 reconfigurable flexible interfaces. The only new hardware required to integrate a new focal plane is a simple PCB, either to carry a bare image sensor or interface to a COTS camera connector. Utilizing the same base board and adding custom sensor modules minimizes the cost, schedule, and risk associated with integrating each new sensor.

The main camera board employs a field programmable gate array (FPGA) for sensor control and configuration, image acquisition, and image processing and storage. The board provides power to the sensor daughtercards and has independently configurable interface voltages for 3 of the 5 camera ports. Attached to this FPGA are 128MB of DDR2 and a microSD card with 8GB of nonvolatile memory to store image and video files.

A Microchip PIC microcontroller on the board performs star-tracking and acts as the system supervisor. It enables or disables the power to the FPGA and acts as a watchdog to reset the FPGA in case of communications timeouts or other errors. It stores the FPGA's configuration file in attached flash memory and can program the FPGA over a parallel interface. The flash memory can be updated with new files from the ground, so in this way, the FPGA code can be altered on-orbit to

fix bugs or add new capabilities. The flash memory also contains the star catalog loaded by the PIC for star tracking.

The SiOnyx daughterboard receives an input clock, a reset signal and an I^2C bus for control from the main camera board. The images are output over a 12-bit data bus with pixel clock, line_valid, and frame_valid signals. All these signals, along with power and ground, are carried over the 39-pin flexible flat cable between the main camera board and daughterboard. The SiOnyx focal plane requires a low-noise analog supply and ground for optimum performance, so the input power is filtered, and the analog and digital rails are kept separate in the daughterboard layout.

EMBEDDED SOFTWARE

Star tracker related processing is performed in three separate locations. Image acquisition and filtering is performed by the FPGA on the main camera board. Filtered image data is then transferred to the low power Star Tracker PIC microprocessor which identifies stars and solves for attitude. The attitude quaternion is then transferred to the Attitude Control System (ACS) PIC microprocessor for spacecraft attitude determination and control. This section details the practical considerations implemented in the embedded processing to account for sensor noise, power requirements, and memory capacity. The section that follows focuses on the algorithms used to process image data into an attitude estimate.

Field Programmable Gate Array

The AeroCube camera's FPGA is a flexible platform supporting multiple CubeSat missions, including multispectral camera payloads[5], star-trackers[6], and general-purpose context cameras. Figure 10 shows a block diagram of the FPGA's internal architecture.
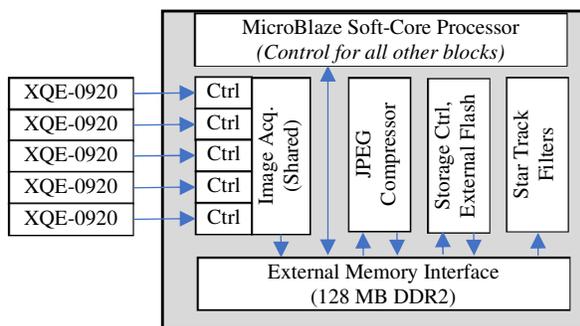

Figure 10: FPGA Internal Block Diagram

In star-tracker mode, the main FPGA functions are to control the focal plane, acquire individual image frames for storage in memory, store captured frames for later diagnostics, and apply a series of image filters to identify the bright pixels of interest.

Control of the focal plane, in the case of the SiOnyx XQE-0920, is through an I^2C interface. The FPGA loads fixed configuration tables and various configuration registers, including those that set exposure-time, analog gain, and other key parameters. Once ready, the focal plane begins streaming video frames.

Image acquisition, also known as a "frame grabber", copies the raw pixel stream to a designated area in external memory. This buffering is required because subsequent steps require multiple passes over the image and cannot be performed in real-time.

The star-tracker filters require multiple steps. The first step, hot-pixel blanking, overwrites the stored value for specific pixels with zero. This is required to reduce false alarms from pixels that are anomalously bright due to manufacturing variation or accumulated radiation damage. The list is specific to each sensor and can be updated on-orbit.

The second step is to apply a series of circular cutout masks, which can remove large unusable image regions caused by lens artifacts, reflections, glare, or excess lens distortion.

The final step is to apply an adaptive threshold, which requires two passes over the image to find the brightest pixels without exceeding available memory in the PIC microcontroller. The final bright-pixel list is then relayed to the Star Tracker PIC for further processing.

Hot Pixel Filtering

Due to manufacturing variance and accumulated radiation damage, some pixels will be consistently brighter than their neighbors under equivalent light exposure. These "hot" pixels can present a challenge to the star tracking algorithm as the worst offenders can occupy several star candidate slots, displacing real star clusters. This results in fewer stars matched to the catalog, and in some cases can result in zero stars matched. One way of addressing these hot pixels is to maintain a list which is used to blank offending pixels. A hot pixel list, updated only a couple of times throughout the mission life, is utilized by the FPGA as a first line of defense; however, this approach fails to filter out new "hot" pixels since the last time the list was updated. The hot pixel list is limited to a maximum of 4,096 pixels. To complement the list-based filter on the FPGA, the Star Tracker PIC can be configured to reject any star cluster comprised of only a single pixel to filter out lone hot pixels that have not yet been added to the removal list.

Power Management

Due to limited power budget, power usage was a major concern when running the attitude control algorithms. Our star tracker, specifically the FPGA responsible for capturing and postprocessing images, consumed a large amount of our power budget. Due to its large percentage of our total power budget, we concluded it was not possible to leave the FPGA on continuously without draining the batteries. The first logical step was to cold boot the FPGA every time it is needed. The problem with that solution is cold booting the FPGA before every image capture would result in a star tracker update frequency orders of magnitude slower than what is required to meet attitude estimation accuracy objectives. To solve this problem, we developed a special low power mode for the FPGA that gave us the benefit of using significantly less power when not in use, while at the same time not having to cold boot the FPGA when needed. The Star Tracker PIC wakes the FPGA from sleep every time a star update is requested. Once the Star Tracker PIC has received the image data, it sends a command to sleep the FPGA. This process of sleeping and waking is done every time a star tracker update is required, unless specified otherwise. The ACS PIC can vary the frequency at which star tracker updates are requested or forego star tracker updates entirely during periods when lower attitude estimation performance is acceptable.

Memory Management

Due to the fact we are using low power microprocessors, data memory is limited. This reduces the amount of data available to determine the vehicle's attitude. The best example of memory constraints directly affecting our control algorithms is the reduction of our star catalog. Our full star catalog contains thousands of stars, which would exceed the memory capacity of the Star Tracker PIC. To solve this issue, we broke up the full catalog into regions with an equal number of stars per region. An initial estimate of the vehicle attitude is used to select which regions of stars to load into data memory as a sub-catalog. The initial attitude estimate can be fairly coarse (inaccuracy of several degrees is tolerated) and is provided by the ACS PIC. The total number of stars that are loaded into the sub-catalog numbers in the hundreds compared to the thousands that exist in the full catalog. The partitioning of the star catalog and the interaction between the Star Tracker PIC and the ACS PIC to enable the sub-catalog approach are further detailed in a previous paper describing the earlier AeroCube-7 vehicles.[6]

ALGORITHMS

The algorithms presented in this section are implemented on the Star Tracker PIC, which receives image data from the FPGA and provides attitude solutions to the ACS PIC for incorporation into an attitude determination Kalman filter. While many of these algorithms have already flown on previous version of this star tracker architecture[6], a notable addition is the optical distortion correction algorithm which enables the use of a low-cost, COTS, wide field of view lens. The methodology used in ground processing to estimate the distortion is also explained.

Pixel Processing and Cluster Calculations

The Star Tracker PIC receives the intensity and location of the brightest pixels from the FPGA and combines these pixels into clusters. A pixel is placed in the same cluster with any of its 8-connected neighbors (touching an edge or corner). The intensity-weighted centroid location and total intensity is calculated for each cluster. A cluster can be rejected if it is comprised of a single pixel (possibly a hot pixel) or if any of its constituent pixels are along the perimeter of the focal plane or circular mask (possibly an incomplete cluster and thus an inaccurate centroid location). The brightest clusters are candidates for star matching.

Star Matching

The cluster centroid locations are converted from pixel coordinates to direction vectors in a sensor-fixed frame using the distortion correction detailed in the next subsection and then rotated into a vehicle body-fixed frame. A star sub-catalog is loaded into memory from a subset of regions of the full catalog. Regions are selected based on an attitude estimate provided by the ACS. A full "lost in space" solution could be found by looping through all of the catalog regions but is not necessary since the ACS PIC runs an attitude determination Kalman filter which is informed by measurements from sun sensors and earth sensors in addition to past star tracker updates.

Star matching is performed by comparing the brightest clusters to the catalog stars two at a time. A pair of direction vector correspondences is required for a 3-axis attitude solution since a single direction vector contains only two pieces of information. The cluster direction vectors are rotated into the same inertial frame as the star catalog direction vectors using the ACS PIC-provided attitude estimate. If the angular separation between the cluster and the catalog star is within a tolerance angle for both candidate correspondences in the pair, then the pair of matches is considered viable. An initial attitude solution is computed from (1) and (2) where s_1 and s_2 are the direction vectors of the two stars in the pair. The

cluster-derived vectors are used for the body-fixed frame representations and the catalog vectors are used for the inertial frame representations.

$$T_{BI} = [\mathbf{x}_B\ \mathbf{y}_B\ \mathbf{z}_B][\mathbf{x}_I\ \mathbf{y}_I\ \mathbf{z}_I]^T \qquad (1)$$

The subscripts B and I denote body-fixed and inertial frame representations, respectively, and the orthonormal unit vectors **x**, **y**, and **z** are constructed as

$$\begin{aligned} \mathbf{x} &= \frac{\mathbf{s}_1 \times \mathbf{s}_2}{|\mathbf{s}_1 \times \mathbf{s}_2|} \\ \mathbf{y} &= \frac{\mathbf{s}_1 + \mathbf{s}_2}{|\mathbf{s}_1 + \mathbf{s}_2|} \\ \mathbf{z} &= \mathbf{x} \times \mathbf{y} \end{aligned} \qquad (2)$$

Using this initial attitude solution T_{BI}, each of the remaining clusters are rotated into the inertial frame and compared to each of the remaining catalog stars. If the angular separation is within the tolerance, that match is deemed viable and added to the solution set.

Lens Distortion Algorithm

Applying lens distortion correction improves the ability to match estimated stars to the catalog. Star matching relies on tight tolerances in relative star positions, which can be difficult to obtain if there is optical distortion. Hence, a calibration routine was created to estimate and correct for these distortion effects. If this type of correction is not applied and the distortion is large enough, a high percentage of frames can get rejected with no valid attitude solution. The lens distortion effects that are estimated and corrected for include tip and tilt, radial, and tangential effects. Note that tip and tilt may also account for tangential corrections, but it was found incorporating another tangential model aided in the correction. The following will discuss how these distortion corrections are applied in the star matching algorithm.

Pixel data collected from the AC10 FPGA sensor feed into the centroiding algorithm which returns the brightest clusters, or possible stars. Using the estimated centroid locations, (u, v) is defined as those locations relative to the image center.

This estimated star location can now be corrected for distortion effects. Starting with tip and tilt corrections represented by the a_1 and a_2 distortion coefficients,

$$U = \frac{u}{a_1 v + a_2 u + 1} \qquad (3)$$

$$V = \frac{v}{a_1 v + a_2 u + 1} \qquad (4)$$

This corrects the star position for the tip and tilt and rotation of the sensor plane[4].

Next radial distortion corrections are applied which, given the shape of the lens, are the dominant source of error. This is modeled by the following polynomial,

$$\delta_{\rho,u} = U(b_1\rho^4 + b_2\rho^2) \qquad (5)$$

$$\delta_{\rho,v} = V(b_1\rho^4 + b_2\rho^2) \qquad (6)$$

where

$$\rho = \sqrt{U^2 + V^2} \qquad (7)$$

and b_1 and b_2 representing the radial distortion coefficients[4]. Higher order terms are considered insignificant in the calculation and are not included.

Finally, tangential distortion corrections are included to account for distortion effects of the lens and sensor not being collinear. It is modeled as

$$\delta_{t,u} = c_1(3U^2 + V^2) + 2c_2 UV \qquad (8)$$

$$\delta_{t,v} = 2c_1 UV + c_2(3V^2 + U^2) \qquad (9)$$

where c_1 and c_2 represent the tangential coefficients. [3]

Condensing these equations, the updated star positions accounting for distortion corrections are

$$u' = U + \delta_{\rho,u} + \delta_{t,u} \qquad (10)$$

$$v' = V + \delta_{\rho,v} + \delta_{t,v}. \qquad (11)$$

The associated unnormalized LOS for these (u', v') coordinates is given by,

$$\mathbf{r}_{est} = \begin{bmatrix} u' \\ v' \\ 1 \end{bmatrix} \qquad (12)$$

and normalizing gives the star unit vector in the tracker frame,

$$s_{est} = \frac{r_{est}}{\|r_{est}\|}. \tag{13}$$

As it goes through the star matching algorithm for attitude processing, it will match these estimated star vectors with the catalog star vectors.

Attitude Computation

Once all potential matches between clusters and catalog stars have been found, the Wahba problem is formulated using all the matches in the solution set as shown in (14), with the same subscript notation used in (1) and (2). The singular value decomposition technique of Markley[9] is used to solve for the attitude solution T_{BI} which is then converted from a direction cosine matrix to a quaternion and sent to the ACS PIC.

$$[s_{B1} \; s_{B2} \; ... \; s_{BM}] = T_{BI}[s_{I1} \; s_{I2} \; ... \; s_{IM}] \tag{14}$$

The measurement matrix for the set of stars (not to be confused with (24) and (25)) is also computed to provide the measurement geometry matrix $H^T H \in \mathbb{R}^{3 \times 3}$ to the ACS PIC as well.

$$H = \begin{bmatrix} skew(s_{B1}) \\ \vdots \\ skew(s_{BM}) \end{bmatrix} \tag{15}$$

where the skew operation accepts a 3D vector and outputs the skew symmetric matrix which when left-multiplying another vector yields the cross product

$$skew\left(\begin{bmatrix} x \\ y \\ z \end{bmatrix}\right) = \begin{bmatrix} 0 & -z & y \\ z & 0 & -x \\ -y & x & 0 \end{bmatrix} \tag{16}$$

The measurement geometry matrix represents the relative accuracy in each axis of the attitude solution accounting for the number and geometric distribution of the stars matched. The ACS PIC uses the inverse of this matrix, scaled by the square of the per-star measurement accuracy, as the measurement covariance matrix R for the star tracker which is used to compute the weight given to the star tracker attitude solution in the attitude estimation Kalman filter.

$$R = \sigma^2 (H^T H)^{-1} \tag{17}$$

Least Squares Estimation for Distortion

To find the optimal distortion coefficients, on-orbit data is post-processed on the ground using the same star matching algorithm that was used onboard AC10 while running the solutions through a least squares estimator.

The distortion estimator algorithm starts by obtaining the residual of the estimated measurement from the predicted star catalog,

$$E = s_{pred} - s_{est} \tag{18}$$

where s_{pred} and s_{est} are found from the star matching algorithm. Initially, without distortion applied, the distortion algorithm will compute E using onboard data for all frames such that there are N stars in the data set. Even though E is a 3-vector for each measurement, there are only two pieces of information as indicated by (12). Define a measurement vector e that consists of two elements from each measurement stacked together,

$$e = \begin{bmatrix} E_{x,1} \\ E_{y,1} \\ E_{x,2} \\ E_{y,2} \\ \vdots \\ E_{x,N} \\ E_{y,N} \end{bmatrix} \tag{19}$$

where $e \in \mathbb{R}^{2N}$.

The goal of the estimation problem is to minimize the norm of this vector e given by,

$$J(x) = \|e(x)\|^2 \tag{20}$$

where $J : \mathbb{R}^m \rightarrow \mathbb{R}$, parameters $x \in \mathbb{R}^m$, and m is the number of estimation states. The intent of the optimization problem[4] is to

$$\underset{x \in \mathbb{R}^m}{minimize} : \quad J(x). \tag{21}$$

The states to be estimated are the lens distortion model coefficients described above which are,

$$x = \begin{bmatrix} a_1 \\ a_2 \\ b_1 \\ b_2 \\ c_1 \\ c_2 \end{bmatrix}. \tag{22}$$

Note that the distortion corrections to be applied are dependent on the lens and not all coefficients may be necessary for each configuration.

A Gauss-Newton method is employed to solve this non-linear least squares problem. After linearizing about some point e_k, J is approximated as,

$$J \simeq \left\| e_k + \frac{\partial e}{\partial x} \delta_x \right\|^2. \tag{23}$$

The solution to this now linearized least squares problem is found by computing the pseudo-inverse of the Jacobian matrix, $H \in \mathbb{R}^{2Nxm}$ defined as

$$H = \frac{\partial e}{\partial x} \tag{24}$$

which becomes

$$\delta_x = -(H^T H)^{-1} H^T e. \tag{25}$$

This is then used to update the state,

$$x_{k+1} = x_k + \delta_x \tag{26}$$

until δ_x drops below a defined threshold and the states have converged.

It is important to perform hot pixel processing on the data before finding the estimated centroid locations, but as an extra precaution, or for outlier data sets, one could remove any 3σ outliers on the error, e, in the subsequent equations before the state update is calculated in equation (26).

Stepping through the distortion coefficient estimation process, start with an initial value for the distortion coefficients, typically zero, run through the star matching algorithm equations (3) - (13), compute an update to the states using equations (18) - (26), and repeat.

ON-ORBIT PERFORMANCE

Thermal and Power

The star tracker subsystem has two components that draw significant power: the focal plane and the image acquisition FPGA. The focal plane draws around 0.5 W and the FPGA draws around 2 W while on. The FPGA also has a sleep mode that drops its power consumption to 0.5 W which is used in between star tracker snapshots. In normal operation, the star tracker is providing a star solution every 3 seconds and the FPGA is on for about 50% of the time bringing the total star tracker average power draw to 1.75 W.

The thermal design for the star tracker components is adequate to dissipate the generated heat and maintain a safe operating temperature for the electronics. The camera FPGA generates a majority of the heat but has good thermal conductivity to its PCB and to the chassis of the spacecraft. Figure 11 shows the temperature of the camera PCB over 80 minutes of typical use. Camera PCB temperatures below 60 Celsius are within the specified operational temperature ranges per the data sheets provided by the component manufacturers.
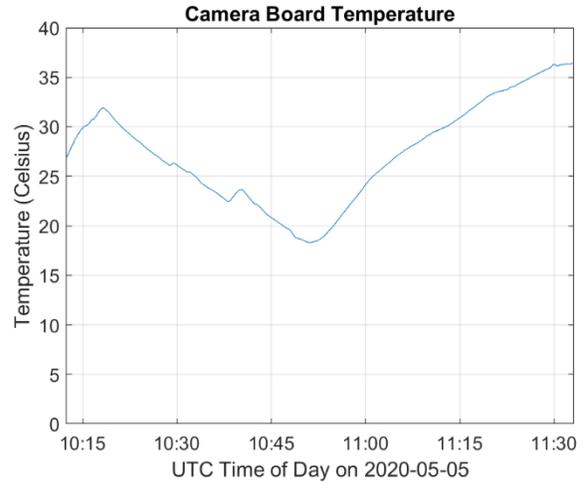

Figure 11: Camera Board PCB Temperature - AeroCube-10B, May 5, 2020

Light Sensitivity

To determine the light sensitivity of the star trackers, we post-processed the on-orbit pixel telemetry on the ground by running them through the same star matching algorithms that are used on-orbit. Due to limited processing power on the vehicles, the star matching algorithms only attempt to match up to 8 of the brightest pixel clusters (computed as the sum of the constituent pixel intensities) to a limited star catalog. However, on the ground we can take advantage of the higher processing power to attempt to match significantly more clusters to a star catalog roughly three times the size of the one on the vehicle; this limit was increased to 30 stars simultaneously. Not only does this provide more star matches overall, but it gives more matches for dimmer stars to better estimate the sensitivity limits of the star tracker. Figure 12 shows all the stars matched with their respective cluster intensities across 872 frames of on-orbit star tracker telemetry taken over a week. The fit line maps the cluster intensity to expected star magnitude at the given exposure time of 0.25 seconds.
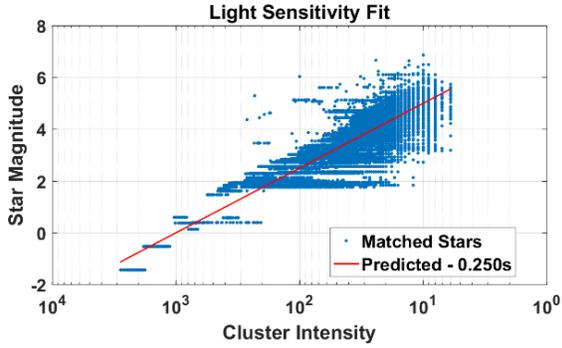
Figure 12: Light Sensitivity Fit of On-orbit Data

The fit equation is of the form given in (29). It is derived from the logarithmic relationship between star magnitude M and brightness B (an increase of 5 in magnitude represents a 100-fold decrease in brightness) (27). Neglecting other effects, cluster intensity I is proportional to brightness B and exposure time dt (28). Rearranging and substituting into (27) yields (29), where c is the fit parameter which captures both the proportionality coefficient k in (28) (the sensitivity of the star tracker) and M_0 in (27). In this case c is 8.9503 when exposure time is given in seconds.

$$M = -2.5 * \log_{10} B + M_0 \qquad (27)$$

$$I = k * B * dt \qquad (28)$$

$$
\begin{aligned}
M &= -2.5 * \log_{10}\left(\frac{I}{k * dt}\right) + M_0 \\
&= -2.5 * \log_{10}\left(\frac{I}{dt}\right) + 2.5 * \log_{10} k + M_0 \quad (29) \\
&= -2.5 * \log_{10}\left(\frac{I}{dt}\right) + c
\end{aligned}
$$

Figure 12 shows that we were able to reliably match stars as dim as magnitude 5.5 based on relative angular position (without needing to compute expected star magnitude from cluster intensity to aid in matching). For comparison, the Berlin Space Technologies ST200 star tracker mentioned previously has an advertised limiting magnitude of 6.0.[7] While we could sense dimmer stars, the catalog choices became hard to distinguish without also discriminating on magnitude, which has a large uncertainty in these measurements (with 0.25 seconds exposure time, cluster intensity is in the single digits for magnitude 5 or dimmer stars, so quantization becomes particularly relevant).

Star Tracker Performance Results

After the distortion estimation algorithm converges to an optimal coefficient set, the error from equation (18) between the estimated stars and catalog stars can be calculated to get a metric for performance. The figure below is a visual representation of the matched stars post-distortion correction.
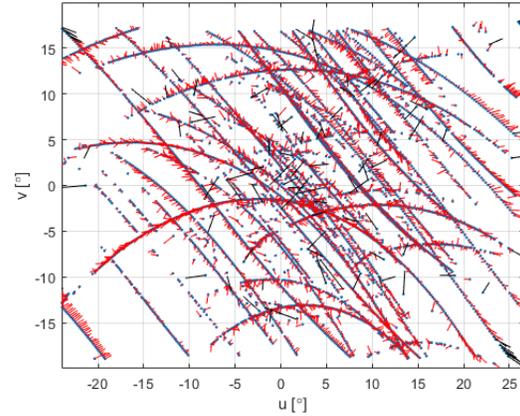

Figure 13: Star Matches Post-Distortion Correction. The blue dots are matched stars on the focal plane and the red lines represent the error.

The (u, v) axes shown in Figure 13 represent the focal plane with the blue dots being the matched star locations found from equations (10) and (11) and the red lines representing the error for each star match. It is important to note that the red lines are scaled up in Figure 13 for visualization purposes. The 3σ rejections, as was discussed in the Least Squares Estimation for Distortion section, are shown by the black lines. Using this visualization along with statistics on e, aid in understanding the expected performance on-orbit.

AC10-A and -B vehicles have two available star trackers of the same design onboard. Although the on-orbit results for both trackers showed improvement after lens distortion correction was applied, a more complete data set was captured with the one of the two trackers on each vehicle. Thus, the following results show the performance of the star tracker with the more complete dataset.

After lens distortion application in the on-orbit star matching algorithm, the percentage of star matches on AC10-A and AC10-B yielded nearly 100% matches amongst a range of missions. It is important to note that the results shown in Table 1 were captured while the vehicles were in eclipse, which minimized stray light in the trackers' field of view during these specific cases.

Table 1: Star Tracker Frame Success Rate (%)

Vehicle	Star Tracker Frame Success Rate (%)
AC10-A	98.6
AC10-B	99.8

The performance of the tracker can be measured by amount of star matches and corresponding residual errors obtained over a range of mission slew rates. Note that a minimum of 4 matching stars is required to achieve a successful attitude solution. As seen in Figure 14, the tracker can successfully match stars over a range of mission rates.
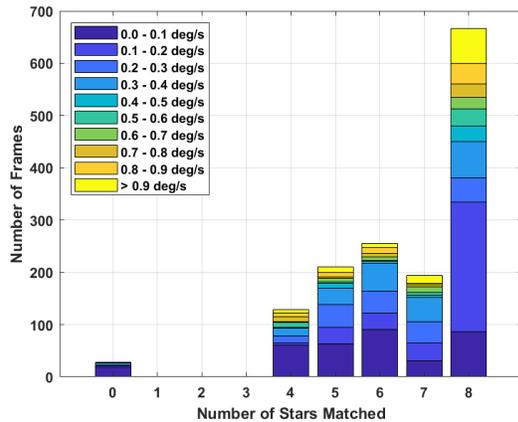

Figure 14: AC10-A Star Matches over a Range of Mission Slew Rates

The performance of individual stars is measured by the residual root mean square (rms) , which is determined by taking the cross product of the matched estimated stars and catalog stars. The corresponding residual error shown in Figure 15 further demonstrates that degradation of star matches remains minimal as rates increase.
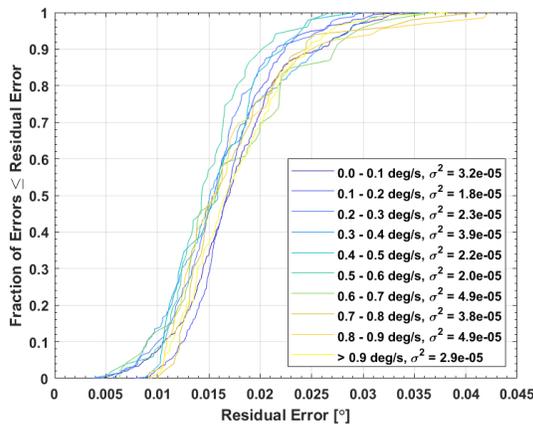

Figure 15: AC10-A Cumulative Distribution of Star Residual Errors

The star matching algorithm provides the tracker's attitude solution to the onboard Kalman filter. Figure 16 shows minimal filter residuals over a range of rates.
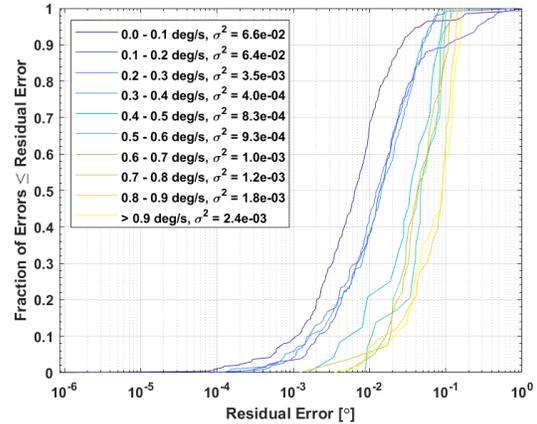

Figure 16: AC10-A Cumulative Distribution of Kalman Filter Residual Errors

Performance Comparison

The star match and Kalman filter update residuals shown in Figure 15 and Figure 16, respectively, are on the order of a few hundredths of a degree which is largely attributable to the 0.05 degree instantaneous field of view (IFOV, the angle viewed by an individual pixel) of the star tracker. Though the WFOV lens results in a relatively large IFOV which limits the accuracy of each star match, the wide field of view does impart the benefit of an overall attitude solution that has more uniform accuracy across all three axes. From (15) and (17) it can be shown that the further apart the stars used to compute the attitude solution are, the better the about-boresight accuracy is.[12] This benefit is further compounded by being able to fly multiple star trackers with boresights pointed in different directions.

Commercially available star trackers from Blue Canyon Technologies[10] and Sinclair Interplanetary[11] offer cross-boresight accuracy of better than 0.002 degrees but the about-boresight uncertainty is greater than 0.01 degrees. While each of these star trackers are an order of magnitude more accurate across the boresight than the AeroCube-10 star tracker, they are also an order of magnitude larger by volume and are roughly comparable in about-boresight accuracy.

CONCLUSION

A miniature star tracker has been developed by The Aerospace Corporation with a high-performance black silicon focal plane made by SiOnyx, LLC. The exceptionally small volume occupied by the camera sub-assembly is multiple times smaller than commercially available CubeSat star trackers. The light sensitivity, quantified by dimmest observable star magnitudes, is comparable to commercial units with significantly larger lenses and apertures. The star tracker has been demonstrated to operate on-orbit at vehicle slew rates exceeding 1.2 degrees per second. The solution accuracy perpendicular to the star tracker boresight is slightly degraded with respect to larger commercial units. For highly volume constrained missions, the reduction in solution accuracy may be worth the gain in available volume for payload components. The wide field of view utilized has the benefit of improved solution accuracy about the star tracker boresight. This new star tracker will help to expand capabilities and potential applications for CubeSat missions 1.5U and below.

ACKNOWLEDGEMENTS

The authors would like to thank John Swigart for his excellent feedback on this paper.

All trademarks, service marks, and trade names are the property of their respective owners.

REFERENCES

1. Sze, S. M. (1981). Physics of Semiconductor Devices. New York: John Wiley & Sons.

2. Mendis, S., Kemeny, S. E., & Fossum, E. R. (1994, March). CMOS Active Pixel Image Sensor. IEEE Transactions on Electron Devices, 41(3), 452-453.

3. Geng, Y., Wang, S., & Chen, B. (2012). Calibration for Star Tracker with Lens Distortion. Proceedings of 2012 IEEE International Conference on Mechatronics and Automation. Chengdu, China.

4. Vaz, B. (2011). Online Calibration for Star Trackers. Master Thesis, Ryerson University.

5. Pack, D. W., Coffman, C., Santiago, J., & Russell, R. (2018). Earth remote sensing results from the CUbesat MULtispectral Observing System, CUMULOS. American Geophysical Union, Fall Meeting.

6. Rowen, D., Utter, A., Dolphus, R., & Alcid, E. (2015). A Miniature, Low-Power Star Tracker for Precision Pointing Nanosatellites. Proceedings of the 38th Annual AAS Rocky Mountain Section Guidance and Control Conference.

7. Berlin Space Technologies GmbH. (2020). Berlin Space Technologies | Star Tracker ST200 and ST400. Retrieved from Berlin Space Technologies: https://www.berlin-space-tech.com/portfolio/star-tracker-st400-and-st200/

8. RDECOM CERDEC Night Vision & Electronic Sensors Directorate (NVESD). (2014). RDECOM CERDEC NVESD Measurements of the Sensitivity, Noise and Resolution of the SiOnyx XQE Sensor.

9. Markley, F. L. (1988, July-September). Attitude Determination Using Vector Observations and the Singular Value Decomposition. The Journal of the Astronautical Sciences, 36(3), 245-258.

10. Blue Canyon Technologies. (2020, May 17). BCT-Datasheet_Star-Trackers.pdf. Retrieved from Blue Canyon Technologies: https://storage.googleapis.com/blue-canyon-tech-news/1/2020/04/BCT-Datasheet_Star-Trackers.pdf

11. Sinclair Interplanetary. (2020, May 17). Second Generation Star Tracker (ST-16RT2). Retrieved from Sinclair Interplanetary: http://www.sinclairinterplanetary.com/startrackers/star%20tracker%20RT%202019a.pdf

12. Dolphus, R. M. (2007, February 20). Autonomous Star Tracker Attitude Error Covariance. The Aerospace Corporation.

13. Coates, C., Fowler, B., & Holst, G. (2009, June 16). sCMOS Scientific CMOS Technology A High-Performance Imaging Breakthrough.