A Ground-Based Prototype of a CMOS Navigational Star Camera for Small Satellite Applications

Brian Shucker
Advisor: Dr. Ernest. D. Fasse

Aerospace and Mechanical Engineering Dept.
University of Arizona, Tucson, AZ 85721
bshucker@cs.arizona.edu

Abstract. Small satellites are now capable of performing missions that require accurate attitude determination and control. However, low size, power, and cost requirements limit the types of attitude sensors that can be used on a small craft, making attitude estimation difficult. In particular, star trackers—often the attitude sensors of choice for larger spacecraft—are not practical for small satellites. This paper describes a miniature navigational star camera based on CMOS sensor technology that is appropriate for small satellite applications. A ground-based prototype version has been built to demonstrate the technology. Hardware design and algorithms for star-pattern recognition and attitude estimation are discussed, along with simulation results and hardware test results.

1 Introduction

In recent years, advances in technology have allowed small satellites to perform missions that would previously have required much larger craft. Some of these missions have required that the small satellite have accurate attitude estimation and even control. This research, while intended for general application, is primarily motivated by one such mission--the UASAT project at the University of Arizona[2,3,4]. The UASAT will detect and image sprites in the upper atmosphere, as well as demonstrate a laser communication system. The latter objective requires UASAT to have 3-axis control with pointing accuracy better than one degree. Relevant statistics for the UASAT are given in Table 1. While the design of the attitude control system is well in hand[2], the problem of estimating the spacecraft's attitude accurately enough is as yet unsolved.

Table 1:UASAT Specifications

Total Mass	68 kg
Dimensions	50 cm × 52 cm
Total Power	<20W continuous
ADCS Subsystem Power	<12W continuous
Orbit	circular: alt 407 km, incl. 51.6°
Required Pointing Accuracy	<1° pitch and yaw, TBD roll

A star sensor of some type is generally required for a spacecraft with accurate 3-axis control[1]. However, this option is not available to the UASAT because of power constraints; the entire spacecraft has a total power budget of about 20W, with only 12W allocated to attitude determination and control. For comparison, the popular Ball Aerospace CT-601 star tracker requires 8-12W operating power[15]. The smaller CT-631 model, intended for the small satellite market, still draws around 8W and requires computational resources to be provided by the host spacecraft[16]. This is clearly impractical for the UASAT attitude determination system. Thus, a new type of low-power star sensor is needed.

Complimentary Metal-Oxide Semiconductor (CMOS) sensor technology makes it possible to develop an extremely small, low-power star camera. CMOS is the technology used in virtually all microprocessors and is thus well-developed and inexpensive. CMOS sensors also use far less power than the CCDs that are currently used in star trackers.

This paper describes a prototype star camera based on the Fillfactory IBIS4 CMOS sensor. The system uses a star-pattern recognition algorithm to compute attitude estimates, rather than locking on to and tracking selected guide stars as with a conventional star tracker[8]. Since small size and limited computational requirements are important, a wide field-of-view (FOV) design was chosen. The wide-FOV approach offers several advantages:

1. Smaller star catalog: Since more of the sky is seen by the sensor at one time, fewer stars must be listed in the onboard star catalog in order to ensure that enough stars are recognized. The catalog is small enough that significant preprocessing results can be stored onboard, which simplifies and accelerates star-pattern recognition.
2. Simpler processing: At this level of accuracy, stars can be effectively viewed as fixed objects in space. Effects such as the proper motion of stars, stellar parallax, and aberration due to the spacecraft's velocity are small and can be ignored.

3. Full-sky search: While a priori attitude knowledge will improve the reliability and speed of star-pattern matching if it is available, the small catalog size makes it possible to search for matches over the entire sky.
4. 3-axis estimate: Since the angular separation between observed stars is fairly large, roll determination is accurate. Thus a single unit can provide a complete attitude estimate.

This design choice also has a few disadvantages:

1. Less accurate estimate: because the wide-FOV design provides low resolution, the positions of stars cannot be determined with great accuracy. However, high accuracy (here meaning better than 0.1 degrees) is not critical in the intended application.
2. Faint signal: commercial off-the-shelf lenses with the required short focal length generally have a very small light-collecting area. Either a custom-built lens or a relatively long integration time is needed to overcome this.

For reference, a comparison of the prototype CMOS navigational star camera and the Ball CT-631 star tracker[16] is given in Table 2.

Table 2: Star Sensor Comparison

	Ball CT-631 Star Tracker	**CMOS Camera**
Mass	2.5kg	< 1kg
Power Consumption	8W	~2W
Field of View	20°	30°
Catalog Size	2000	500
Accuracy	12 arcsec	200 arcsec
Update Rate	5Hz	5 Hz

The purpose of this research is to demonstrate the viability of a CMOS-based star tracker using a wide-FOV pattern recognition approach.

2 Notation

Vectors and set elements will be denoted by lowercase boldface letters. Matrices and sets will be denoted by uppercase boldface letters. Scalars will be denoted by italic letters. The following reference frames are used:

Earth-Centered Inertial Frame (ECI) The ECI frame is centered at the Earth and is fixed with respect to the stars. The z-axis points at the celestial pole. The x-axis points toward the mean equinox, the direction from the Earth to the sun on the first day of spring. The y-axis is chosen such that a right-handed orthonormal triplet is formed.

Camera-Fixed Frame (CAM) The origin of the CAM frame is the center of mass of the spacecraft. The z-axis points along the bore sight of the star camera. The x-axis is perpendicular to the z-axis and is parallel to the direction of the rows of pixels in the sensor. The y-axis is parallel to the direction of the columns of pixels in the sensor.

The attitude of a rigid body may be defined by a reference frame attached to the body[5]. The CAM frame will be used to define the attitude of the star camera with respect to the ECI frame. Denote the frame vectors of the CAM frame in ECI coordinates with $\mathbf{e}_x$, $\mathbf{e}_y$, and $\mathbf{e}_z$. Then the direction cosine matrix

$$\mathbf{A} = [\mathbf{e}_x \;\; \mathbf{e}_y \;\; \mathbf{e}_z]^T \qquad (1)$$

may be interpreted as the camera's attitude with respect to inertial space.

3 Detector Hardware

The detector selected for the star camera is the FillFactory IBIS4 CMOS sensor. The IBIS4 specifications are listed in Table 3[14]. A prototype camera has been constructed using an IBIS4 development board, which includes the sensor, associated electronics, and software for image capture. The prototype optics consist of a single lens of 16mm focal length, giving a field of view of 30.5 by 24.4 degrees with speed f/1.4. In addition, a filter is used to limit the wavelengths of light reaching the detector. This particular filter is a compromise between the standard V and R filters; it allows peak transmission at 590nm, with a full width at half power of 200nm. While the filter reduces the signal strength of stars, it also reduces chromatic aberration and makes it possible to calibrate the sensor on the ground, using wavelengths that are not absorbed by the atmosphere. A standard V filter would serve the same function.

Table 3: IBIS4 Specifications

Pixel Format	1280x1024 SXGA
Pixel Size	7 x 7 um
Frame Rate	7 frames/sec nominal
Spectral Sensitivity Range	400-1000nm (visual and near-IR)

To compute the expected signal strength from a star, we start by computing the stellar flux. The star catalog (see below) contains the magnitude of each star. Also known is the zero magnitude flux for the filter. This value is generally found experimentally by calibrating

against a standard star, usually Vega[9], and is a monochromatic flux measured at the center of the filter passband. To compute the total flux associated with a zero-magnitude star, the monochromatic flux function must be integrated over the passband:

$$F_0 = \int_{l_c-\Delta l}^{l_c+\Delta l} f_o \cdot F(l) \cdot s(l) \cdot f(l) dl \qquad (2)$$

where:

F_0 = total flux-response from zero magnitude star (in A/m^2)

l_c = wavelength at center of filter passband (in nm)

f_0 = zero magnitude monochromatic flux at l_c (in W/ m^2/nm)

Dl = half-width of passband (in nm)

$F(l)$ = stellar flux at l, normalized to $F(l_c)$ (dimensionless)

$s(l)$ = sensor response function evaluated at l (in A/W)

$f(l)$ = filter transmission function evaluated at l (dimensionless)

Note that since the sensor response function is included in the integral, F_0 is the total zero-magnitude flux times the sensor response. The spectral response for the IBIS4 sensor is shown in Fig. 1[14]. A useful approximation of equation 2 is the following:

$$F_0 = f_0 \cdot Dl \cdot s \qquad (3)$$

where s is the mean sensor response over the filter passband.

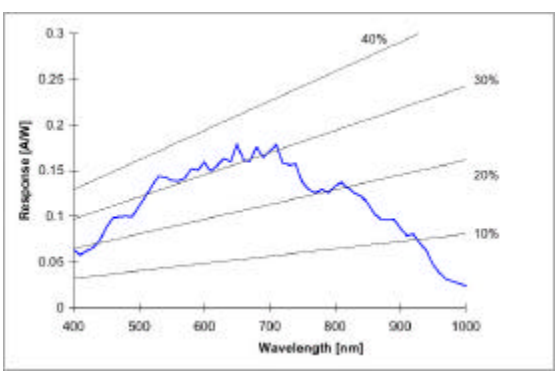

Figure 1: IBIS4 Spectral Response

Once F_0 is known, calculating the flux-response F_m for a star of given magnitude m is straightforward:

$$F_m = F_0 \times 10^{-0.4m} \qquad (4)$$

Note that in practice, F_0 is a parameter that must be measured experimentally to properly calibrate the sensor. F_m is measured for a number of stars with known magnitudes, and these data are used to solve for an empirical value of F_0.

The signal current I_m is the product of the flux-response and the light-collecting area A of the camera lens:

$$I_m = F_m \times A \qquad (5)$$

Once the signal current is known, the resulting accumulated charge on the sensor can be computed by multiplying by the integration time.

$$Q_m = I_m \cdot t \qquad (6)$$

The actual quantity read from the sensor is a voltage that is a function of Q_m. However, since the charge from a given source will be distributed across several pixels, Q_m will be the sum of a number of pixel readings (see Sec. 4).

The most significant source of noise is "dark" thermal noise. See Sec. 9 for a discussion of thermal noise.

4 Image Processing

Once an image is obtained by the sensor, it must be processed in order to extract positions and magnitudes for the observed stars. An initial version of the algorithm for this task has been developed and tested on a small number of real images. The basic algorithm is as follows:

The image is first scanned to find the N brightest pixels, where N is a small constant (~20). If any two of the bright pixels are within a threshold distance of each other, only the brighter one is considered. No pixel below a fixed threshold level is considered, so less than N bright pixels may be found. The radius of a star image, in pixels, is known from the optical properties of the camera. If the bright pixels are surrounded by darkness outside of this radius, then they are likely to be the central pixels of star images.

Each pixel P_i in a star image has center coordinates (x_i, y_i) and is associated with a pixel intensity $I(x_i, y_i)$. The center of the image (x, y) is determined by a weighted average of the pixel locations and is given by the following:

$$x = \frac{\sum_{\forall i} x_i I(x_i, y_i)}{\sum_{\forall i} I(x_i, y_i)} \qquad y = \frac{\sum_{\forall i} y_i I(x_i, y_i)}{\sum_{\forall i} I(x_i, y_i)} \qquad (7)$$

This computation provides sub-pixel resolution, so in theory the center of a star can be found to sub-pixel accuracy. The coordinates are multiplied by the scale factor of the camera (418 microradians/pixel) to obtain measurements in radians, which are then converted into unit vectors in the CAM frame.

The brightness of the image, which is proportional to the quantity Q_m from above, is the sum of the intensities of all the pixels in the image, minus the "dark" intensity I_0.

$$Q_m \propto \sum_{\forall i} I(x_i, y_i) - I_0 \qquad (8)$$

The dark intensity is an experimentally determined quantity that reflects the average intensity of a pixel that is looking at empty space. The proportionality constant may be propagate in the calculations described in Sec. 3 and ultimately subsumed into the quantity F_0 when it is determined empirically.

5 Star Catalog

In order to make an attitude estimate based on the observed stars, the star camera must have a catalog of known star positions available. The catalog consists only of a list with the expected position and magnitude of each star. Position is given as a unit vector in the ECI frame. Information about the proper motion of stars is omitted, because the motion is undetectable at the camera's resolution (as are stellar parallax, aberration due to the camera's velocity, and other such effects). There are several constraints involved in selecting stars to include in the catalog:

1. The camera should be able to detect every star in the catalog.
2. Stars that are not in the catalog should not be detected by the camera, or should be ignored if they are detected.
3. There should be enough stars in the catalog so that, at any time, it is likely that enough of them are in the field of view of the camera in order to make a star pattern match.
4. Within the above constraints, the size of the star catalog should be minimized.

NASA's Goddard Space Flight Center publishes the SKYMAP Master Catalog[12], which contains a great deal of data about astronomical objects (mostly stars) that are useful as attitude references. The SKYMAP catalog contains spectral data as well as the position and magnitude of each star, so it can be used to compute the expected signal from each star as described in Sec. 3 above. The stars can then be sorted by signal strength, and the brightest N stars selected for inclusion in the catalog, with N determined in order to satisfy constraint 3. The camera must then be calibrated to ignore stars below the threshold level that will allow the detection of the Nth brightest star.

NASA has also published several mission-specific star catalogs[13]. For the preliminary catalog, some computation was saved by using the mission catalog

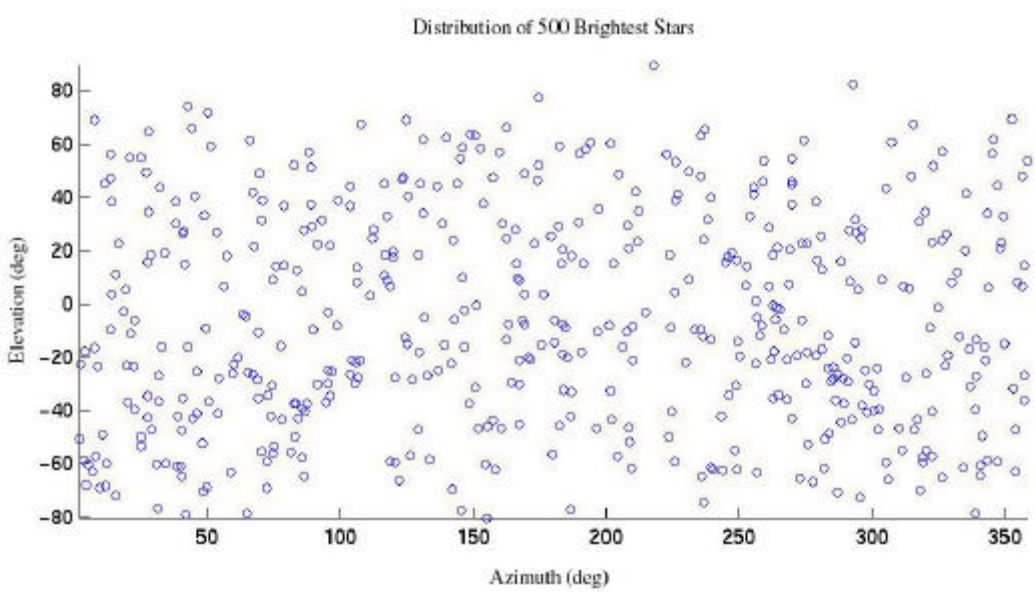

Figure 2

from the Submillimeter Wave Astronomy Satellite (SWAS), rather than the full SKYMAP Master Catalog, as a starting point.

Simulation has shown that a star pattern can be reliably recognized with 3 stars in the camera field of view (see Sec. 8 below). With a camera field-of-view of 30 degrees, a catalog of 500 stars is sufficient to ensure that 3 stars will be in the field-of-view more than 95% of the time (see Sec. 8). The distribution of the brightest 500 stars in spherical coordinates in ECI frame is shown in Fig. 2.

In addition to a list of star unit vectors, the star catalog contains a preprocessed list of the separation between each pair of stars. Each entry in the separations list is a triple (a, i_1, i_2), where i_1 and i_2 are the indices of two stars in the catalog, and a is the cosine of the angle between those two stars. The separations list is sorted in decreasing order by a. Entries with an a value smaller than the cosine of the field-of-view of the camera are not included, because the corresponding stars cannot both be seen by the camera at the same time.

With 500 stars in the catalog, the separations list occupies roughly 100KB of memory, which is not unreasonable. However, the separations list has a size proportional to the square of the number of stars in the catalog, so it may become impractical to store the list if the catalog becomes large. Keeping the separations list small enough to store in memory is one motivator behind constraint 4 above and the wide-FOV design.

6 Star Pattern Matching and Solution

The star-pattern matching algorithm is based in part on known techniques that utilize the observed angular separation between star pairs[6,7], and consists of two phases. In the first phase, candidate matches from the star catalog are selected for each star observed in the camera image. In the second phase, combinations of candidate matches are tested until enough agreeing matches are found to compute an attitude estimate.

Let **S** be the set of M stars in a camera image, and let **s** and **t** be stars in **S**. To chose a candidate match from the star catalog for **s**, let $a_{s,t}$ be the cosine of the angle between **s** and **t**. Compute $a_{s,t}$ for every $\mathbf{t} \in \mathbf{S}$, $\mathbf{t} \neq \mathbf{s}$. For each **t**, look up $a_{s,t}$ in the separations list. When $a_{s,t}$ is found on the separations list, add one "hit" to each of the two star indices (i_1 and i_2) associated with $a_{s,t}$. Ideally, the index that corresponds to **s** is expected to have M-1 hits, because each star pair observed in the image will correspond to a pair of stars in the catalog, one of which is **s**. The other star in the pair will receive a false hit.

The number of false hits is reduced greatly by the use of magnitude information. When assigning hits, the measured magnitude of **s** is compared to the expected magnitude of the star that is being assigned a hit. Hits on stars that do not match the expected magnitude within the measurement accuracy of the sensor are ignored.

Further reduction in the number of false hits may be accomplished if the approximate attitude of the camera is known beforehand. This information may come from other sensors on the spacecraft, or from attitude estimates based on previous star camera images. The camera's expected field-of-view is calculated based on the known attitude, and a window is defined to include this expected field-of-view, enlarged somewhat to account for estimation error and motion of the spacecraft between estimates. Hits on stars outside of this window may be ignored, because those stars cannot be in the camera's field-of-view.

In practice, errors will be introduced in several ways. In particular, the angle to a star cannot be measured with perfect accuracy. If the error in measuring the angle to any star is bounded by e radians and the measured angle between two stars **i** and **j** is $q_{i,j}$, then $q_{i,j}$ may be off by as much as $2e$ in either direction. To compute the possible range of $a_{i,j}$, the following properties of cosines are used:

$$\cos(A-B) = \cos(A)\cos(B) + \sin(A)\sin(B) \qquad (9)$$
$$\cos(A+B) = \cos(A)\cos(B) - \sin(A)\sin(B)$$

for any angles A and B. Thus, that the actual value of $a_{i,j}$ obeys the following:

$$\cos(q_{i,j})\cos(2e) - \sin(q_{i,j})\sin(2e) < a_{i,j} \qquad (10)$$
$$a_{i,j} < \cos(q_{i,j})\cos(2e) + \sin(q_{i,j})\sin(2e)$$

The matching algorithm accounts for this by adding a hit to each pair of stars whose separation lies within the bounds of $a_{i,j}$. Depending on the accuracy of the camera and the size of the star catalog, this may result in a significant number of false hits.

Additional false hits may be created if some light source that is not in the catalog is seen in the image. However, some of these light sources, such as the moon and the planets Jupiter, Venus, and Saturn, will be ignored because they do not have magnitudes similar to stars. Besides false hits, the algorithm is robust against the possibility that stars that are in the catalog may not

be observed. The latter would occur, for example, if the earth were blocking part of the camera's field-of-view. It could also occur if part of the sensor malfunctioned or was miscalibrated. For these reasons, the exact number of hits on the correct star index is difficult to predict, but the correct index is the only one that is likely to accumulate many hits. Thus the selected candidate star is the one whose index has the largest number of hits. Simulation shows that even with a large number of missing or added star images, this algorithm is a reliable method for identifying stars (see Sec. 8).

Phase two of the pattern-matching algorithm uses a simple verification scheme to choose two candidate star matches to use for attitude estimation. For every pair of stars observed in the image, the angular separation between the observed star pair and the corresponding pair of stars in the catalog are compared. If one or both stars have been misidentified, the separations will not match. Also, star pairs that are too close together are discarded, because the attitude estimate will be inaccurate if there is not enough separation between the two reference points used.

Once two stars have been correctly matched, computing an attitude estimate is straightforward. Let **A** be the ECI-to-Camera frame transformation matrix. Let $\mathbf{v}_{ECI,1}$ and $\mathbf{v}_{ECI,2}$ be the unit column vectors for the matched stars in ECI frame (that is, the unit vectors listed in the star catalog) and let $\mathbf{v}_{CAM,1}$ and $\mathbf{v}_{CAM,2}$ be the corresponding unit column vectors measured in the camera frame. Then we know that

$$\mathbf{A}\,\mathbf{v}_{ECI,1} = \mathbf{v}_{CAM,1} \quad \text{and} \quad \mathbf{A}\,\mathbf{v}_{ECI,2} = \mathbf{v}_{CAM,2} \qquad (11)$$

Since **A** represents a rotation from one right-handed orthonormal basis to another, the following property holds for any two vectors u and v:

$$\mathbf{A}(\mathbf{u} \times \mathbf{v}) = \mathbf{A}\mathbf{u} \times \mathbf{A}\mathbf{v} \qquad (12)$$

where $\times$ represents the cross product operation.
Using (11) and (12), we have the following:

$$\begin{aligned}\mathbf{A}(\mathbf{v}_{ECI,1} \times \mathbf{v}_{ECI,2}) &= \mathbf{A}\mathbf{v}_{ECI,1} \times \mathbf{A}\mathbf{v}_{ECI,2} \\ &= \mathbf{v}_{CAM,1} \times \mathbf{v}_{CAM,2} \qquad (13)\end{aligned}$$

Let

$$\begin{aligned}\mathbf{V}_{ECI} &= [\mathbf{v}_{ECI,1} \quad \mathbf{v}_{ECI,2} \quad (\mathbf{v}_{ECI,1} \times \mathbf{v}_{ECI,2})] \text{ and} \\ \mathbf{V}_{CAM} &= [\mathbf{v}_{CAM,1} \quad \mathbf{v}_{CAM,2} \quad (\mathbf{v}_{CAM,1} \times \mathbf{v}_{CAM,2})] \quad (14)\end{aligned}$$

Then we have:

$$\mathbf{A}\,\mathbf{V}_{ECI} = \mathbf{V}_{CAM} \qquad (15)$$

Thus,

$$\mathbf{A} = \mathbf{V}_{CAM}\,(\mathbf{V}_{ECI})^{-1} \qquad (16)$$

The matrix **A** defines the camera's attitude in inertial space.

7 Complexity Analysis

The algorithm described above is designed to run with very limited availability of computational resources, as may be the case for many small satellites and is certainly the case for UASat[4]. In principle, the running time for the star-pattern matching algorithm is independent of the number of stars in the onboard catalog, and is a polynomial function of the number of stars observed in a given image. This latter number may be bounded by a relatively small constant value. For reasons of simplicity, the current implementation does not take full advantage of all available techniques and thus runs in time dependent on the logarithm of the size of the star catalog. In practice, the image processing task takes longer than the pattern-matching task.

The image processing time is dominated by the input/output cost of scanning the image and transferring it into the computer. This time may be considered constant for any image.

$$T_{image} \text{ £ } c_{im} \qquad (17)$$

Note that this represents a worst-case situation; in the typical case, most of the image need not be scanned. This is because a previous image, combined with rate information, may be used to predict the approximate location of stars in the new image. The IBIS4 imager supports subwindow scanning, so only a small subwindow around the expected position of each star must be scanned. Because image processing for the ground-based tests is not done in real time, this optimization has not been implemented for the prototype.

Let N be the number of stars in the onboard star catalog, and let M be the number of stars observed in a particular image. Let M_{max} be an upper bound on M; if there are more than M_{max} stars in the image, only the first M_{max} of them will be considered for pattern matching and the rest will be discarded.

Phase one of the matching algorithm performs a lookup in the star separations list for every pair of stars in the image. The total number of lookups is $M(M-1)$. Since

the separations list is sorted, the current implementation uses a simple binary search to perform each lookup. The time to perform a single lookup depends on the number of star pair separations in the list, which in turn depends on *N*.

In the worst case, the separations list could contain one entry for every pair of stars in the catalog, which would be a total of $N(N-1)/2$ entries. This can be approximated to $N^2/2$ entries. However, only a small fraction k of the entries will have an a value large enough to make the final cut (see above); the exact value of k depends on the star camera's field-of-view and the star distribution. Using a 30° field-of-view and the real distribution of stars in the sky, approximately 2% of the star pair separations are actually listed in the onboard separations list. Thus the lookup time for a binary search is:

$$T_{lookup} = c_1 \times \log_2(kN^2/2) = c_1 \times (\log_2(kN^2) - 1) = c_1 \times (\log_2(0.02N^2) - 1) \text{ with } 30° \text{ FOV} \quad (18)$$

The total time used by phase1 is thus given by

$$T_{phase1} = T_{lookup} \times M(M-1) = c_1 \times (\log_2(kN^2) - 1) \times M(M-1) \; £ \; c_1 \times \log_2(kN^2) \times M^2_{max} \quad (19)$$

As lookups are performed and hits tabulated, the current greatest number of hits and associated index may be stored, so the best-match star is known after all lookups are complete. Thus a second search for the best-match star is not required.

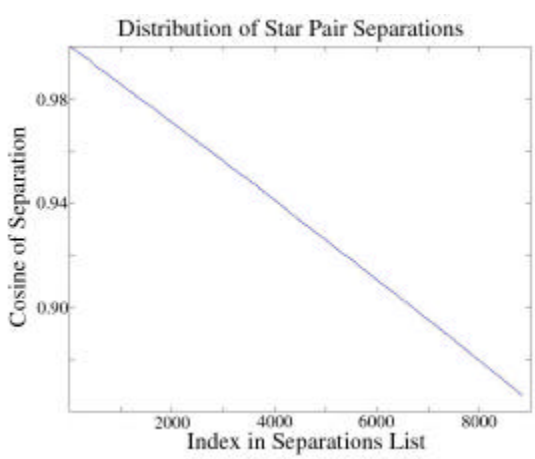

Figure 3

A further optimization makes use of the known distribution of star separations. As seen in Fig. 3, the distribution is nearly linear. Given a separation value to look up, a starting index can be computed by using this linear relation. By starting a search at the computed index, the lookup can be accomplished in constant time, with the constant depending linearly on the maximum deviation of any true index from its corresponding computed index.

Since the star catalog is not very large, the current implementation does not take advantage of either of the above optimizations.

The current implementation of the second phase looks at every pair of matched stars in the worst case, although in the typical case only the first few pairs will be examined before a verified match is found. For the worst case:

$$T_{phase2} = c_2 \times M(M-1)/2 \; £ \; c_2 \times M^2_{max}/2 \quad (20)$$

Once the stars to be used for attitude estimation are selected, solving for the ECI-to-Camera matrix requires two cross product operations, one matrix invert and one matrix multiply, which are all constant-time operations.

$$T_{solve} = c_s \quad (21)$$

The total running time is given by the sum of the four steps:

$$T_{all} = T_{image} + T_{phase1} + T_{phase2} + T_{solve} \; £ \; c_{im} + [c_1 \times \log_2(kN^2) \times M^2_{max}/2] + c_2 \times M^2_{max}/2 + c_s \quad (22)$$

Since c_{im}, c_1, c_2, c_s, k, N, and M_{max} are all constants, the total running time for the complete star camera algorithm can be bounded by a constant. This constant is easily computable and relatively small, although the exact value will depend on the type of processor used.

Source code for the current implementation is available in [9].

8 Simulation Results

A number of simulations have been conducted in order to verify the pattern matching algorithms described above. The simulations account for error in the observed position and magnitude of stars. Measurement errors have components of both a constant and a random nature. Other sources of error are modeled, including the omission of stars that are expected to be observed and the addition of false star images.

The following data are the result of 1000 trial runs in simulation. For each trial, a simulated scene was generated with a random camera orientation. The

actual mission star catalog was used for scene generation. Scene parameters are shown in Table 4.

Table 4 : Simulation Parameters

Maximum position measurement error	0.05 deg (1.7 pixels)
Maximum magnitude measurement error	0.1 magnitudes
Fraction of expected star images omitted	10%
Fraction of scenes with false star images	10%

In addition to simulated measurements of star positions and magnitudes, the software was provided with a previous attitude estimate accurate to within 5 degrees. The results are shown in Table 5.

Table 5: Simulation Results

Stars Seen	# of Trials	Successful Estimates	Failed Pattern Matches	Erroneous Estimates	Success (%)
0	0	0	0	0	N/A
1	5	0	5	0	0
2	22	17	5	0	77
>2	973	973	0	0	100
ALL	1000	990	10	0	99

The median number of stars in a scene was 7, and the max was 19.

In the cases where only one star is observed, the software fails to make an attitude estimate because the task is impossible with only one reference point. When two stars are observed, it may or may not be possible to extract attitude information; one of the observed stars may be a false image, and there may be more than one pair of stars in the catalog that could match the observed pair. In all of these situations, the software reports a lack of sufficient information and does not attempt to compute an attitude estimate.

When three or more stars are observed in a simulated scene, a star pattern match is made in every case.

Thus, the pattern matching algorithm may fail if fewer than three stars are observed, which may be the result of the earth blocking part of the field of view. It will also fail if the sun enters the field of view, because the entire image is likely to be destroyed in that case (however, direct exposure to sunlight will not damage the camera in the short term). Because of these limitations, the UASat project team plans to include at least two, and possibly three, star camera units located at different points on the spacecraft. At least one unit should always be pointing at a clear star field, without the earth or the sun interfering.

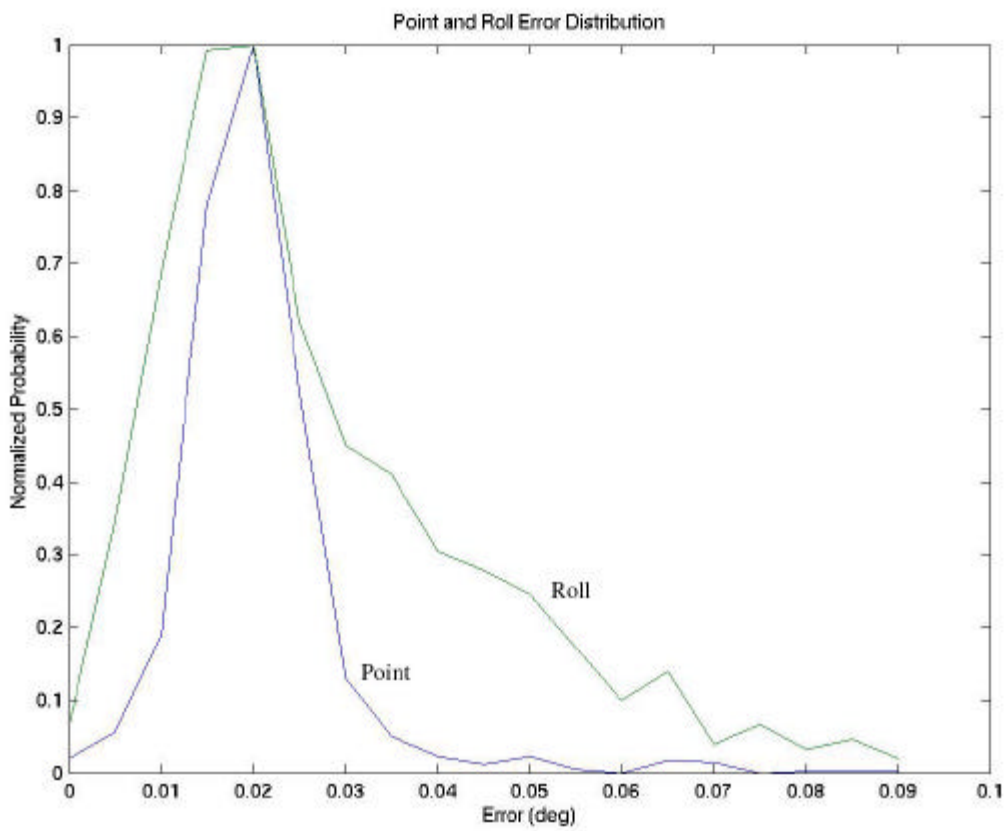

Figure 4

In the 99% of cases where a correct star pattern match is made, the mean accuracy of the computed attitude estimate is 0.023 degrees in pointing (pitch and yaw), which is very close to the mean error in measuring star positions. The mean roll estimate error is 0.039 degrees. The roll estimate error is expected to be somewhat large because the two reference points from which the estimate is made are relatively close together. The distribution of pointing and roll estimate errors is shown in Fig. 4.

9 Hardware Test Results

A series of basic hardware tests was conducted, using the ibis4 development board mentioned above with an inexpensive commercial off-the-shelf video lens and a simple uncooled housing. The camera control software was run on a laptop PC connected to the camera unit, and the images were stored for later processing. Test images were taken without the 590nm filter.

The optical properties of the system, including field of view, pixel scale, and focus, were confirmed in testing. Because the sensor was not actively cooled and it generates some heat while active, the operating temperature was significantly above the ambient temperature. The high temperature resulted in a large amount of random thermal noise in all of the test images. Since there is currently no direct temperature sensor or control on the imager, the exact dependence of noise on temperature has not been experimentally verified. However, the ibis4 specification[14] suggests that a decrease in temperature of 10°C corresponds to a reduction in thermal noise by a factor of 5.

Despite the thermal noise, it was possible to image some of the brighter stars in the sky. Figure 5 shows a raw camera image taken in the direction of Gemini while the half moon was present; Pollux is visible at the upper edge and Procyon at the lower right. The exposure time for this image was 536ms, and the ambient temperature was approximately 25°C. The image processing software successfully located the centers of the stars, and their relative positions agree with the known positions of those stars to within 1/20

Figure 5

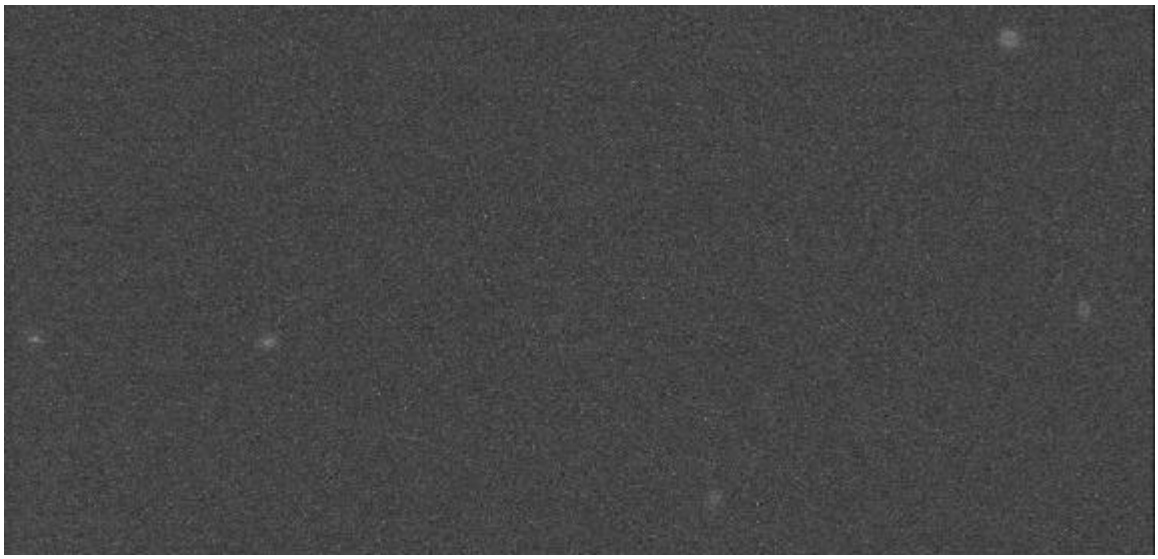

Figure 6

of a degree. However, the noise makes it difficult to determine their brightness. Notice that the moon has oversaturated a section of the image and bloomed out to appear more than four times its actual diameter.

Figure 6 shows a section of an image taken in the direction of Ursa Major (the big dipper); several stars in that constellation are apparent. For this image, the ambient temperature was approximately 10°C and the integration time increased to 4.3s with a lower analog-to-digital converter gain, to bring out the dimmer stars against the background of thermal noise. Even with a low ambient temperature, the sensor still accumulated significant heat from its own operation.

Figure 7 shows a small section of an image containing Arcturus, a magnitude 0 star. This image is from a long exposure (6.4s), so the star has saturated a small section of the sensor. Notice that the blooming is contained in a small area and distributed uniformly around the star, so it is still possible to accurately determine the center. To make individual pixels apparent, the figure is magnified by a factor of 8.

These images, along with other test data, indicate that first magnitude stars can be reliably located with the imager operating at approximately 20°C. Since the 500 brightest stars are all of magnitude 3.5 or brighter, they will be easily detectable if the imager is cooled to 0°C, which is certainly feasible on a spacecraft. However, construction of a thermally controlled prototype was beyond the scope of this project.

The ability of the image processing software to accurately locate stars was verified using the test images. The software also produced reasonable results for the magnitude of each star, but was hampered by the presence of significant noise in the images. Thus, accurate assessment of the magnitude measurements was not possible.

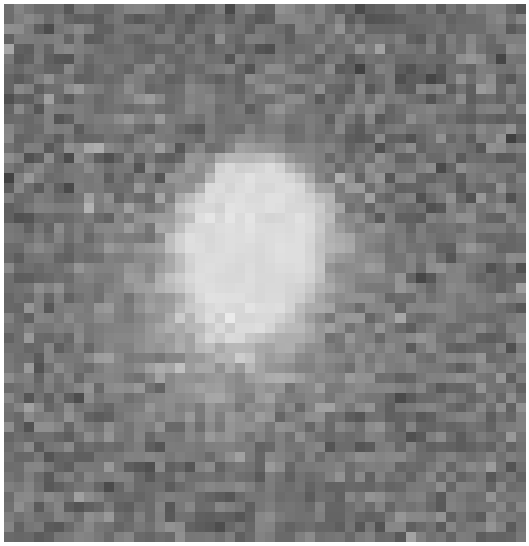

Figure 7

10 Conclusion

Initial results indicate that a CMOS-based navigational star camera is a viable primary attitude sensor for a small satellite. Testing of the prototype will continue in the near future, and work will continue during the 2001-2002 academic year. Future projects include the development of a thermally controlled prototype and integration of the star camera software with existing software for attitude estimation and control of UASat.

11 Acknowledgments

Prototype hardware costs were supported largely by funding from the University of Arizona Student Satellite Program (SSP), a program of the Arizona Space Grant Consortium. The author also thanks professors Ernest Fasse and Uwe Fink for their support, as well as the SSP Guidance, Navigation, and Controls team.

References

1. James R. Wertz and Wiley J Larson (eds). *Space Mission Analysis and Design*, chapter 11. Microcosm Press, Torrance, California. Third edition. 1999.

2. Barry Goeree and Brian Shucker. Geometric Attitude Control of a Small Satellite for Ground Tracking Maneuvers. In *Proceedings of the 13th AIAA/USU Conference on Small Satellites*. Utah State University, August 1999.

3. Barry Goeree and Greg Chatel. Attitude Control of the UASat. Technote GNC-007, Student Satellite Project, University of Arizona. May 1998. Available at http://uasat.arizona.edu/ssp/documents/technotes/GNC-007.pdf (Feb 2001).

4. Igor Ageyev. DCH Flight Hardware Overview. Technote DCH-003, Student Satellite Project, University of Arizona. April 1998. Available at http://uasat.arizona.edu/ssp/documents/technotes/DCH-003.pdf (Feb 2001).

5. M.D. Shuster. A survey of attitude representations. *J. of the Astronautical Sciences*, 41:439-517, 1993.

6. Daniele Mortari. Search-Less Algorithm for Star Pattern Recognition. In *The Journal of the Astronautical Sciences*, Vol. 45, No. 2, April-June 1997.

7. Roberto Lopes, Gustavo Carvalho, and Adenilson Silva. Star Identification for Three-Axis Attitude Estimation of French-Brazilian Scientific Microsatellite. In *Proceedings of the AAS/GSFC International Symposium*, Greenbelt, MD, May 11-15, 1998.

8. Eric Stoneking. A Robust Star Acquisition Algorithm for the Wide-Field Infrared Explorer. In *Proceedings of the AAS/GSFC International Symposium*, Greenbelt, MD, May 11-15, 1998.

9. Brian Shucker. Ground-Based Prototype of CMOS Navigational Star Camera for Small Satellite Applications. Honors Senior Thesis, University of Arizona. May 2001.

10. Michael Zeilik and Stephen Gregory. *Introductory Astronomy and Astrophysics*. Saunders College Publishing, Fort Worth. Fourth Edition. 1998.

11. Myron Kayton and Walter R. Fried, *Avionics Navigation Systems*, Wiley Interscience, New York. Second Edition, 1997.

12. SKY2000 Master Catalog, version 2. Available at http://cheli.gsfc.nasa.gov/dist/attitude/SKYMAP_021201_page.html (April 2001).

13. Submillimeter Wave Astronomy Satellite (SWAS) Mission Run Catalog. Available at http://cheli.gsfc.nasa.gov/dist/attitude/SKYMAP_021201_page.html#SWASSEC (Jan2001).

14. IBIS4 SXGA Image Sensor Datasheet. Fillfactory. Available at http://www.fillfactory.com/htm/cmos/datasheets/ibis4.pdf. (Dec 2000)

15. CT-601 High Accuracy Star Tracker (specification sheet). Ball Aerospace and Technologies Corporation. Available at http://www.ball.com/aerospace/pdf/ct601.pdf (Mar 2001).

16. CT-631 Lightweight Star Tracker (specification sheet). Ball Aerospace and Technologies Corporation. Available at http://www.ball.com/aerospace/pdf/ct631.pdf (Mar 2001).