

Success by 1000 Improvements: Flight Qualification of the ST-16 Star Tracker

Tom Dzamba, John Enright
Department of Aerospace Engineering, Ryerson University
350 Victoria St, Toronto, Ontario, M5B 2K3; 416-979-5000(x4973, 4174)
tdzamba@ryerson.ca, jenright@ryerson.ca

Doug Sinclair, Kofi Amankwah
Sinclair Interplanetary
dns@sinclairinterplanetary.com, kaa@sinclairinterplanetary.com

Ronny Votel
Skybox Imaging Inc.
1061 Terra Bella Ave., Mountain View, CA 94043; 650-316-6640
ronny@skyboximaging.com

Ilija Jovanovic, Geoffrey McVittie
Department of Aerospace Engineering, Ryerson University
350 Victoria St, Toronto, Ontario, M5B 2K3; 416-979-5000(x4973)
ilija.jovanovic@ryerson.ca, gmcvitti@ryerson.ca

ABSTRACT

The first launch of a pair of 90 gram Sinclair Interplanetary ST-16 star trackers was in November, 2013 on-board the Skybox Imaging SkySat-1 satellite. The sensor performance — as captured by the sensors' availability, accuracy, and bad match frequency — fell significantly below expectations. This paper explores the flight qualification campaign undertaken by the sensor developers to bring the sensors back to their intended level of performance. No single fix was sufficient and many small incremental improvements were necessary for success. We discuss the fault diagnosis procedures employed by the team and highlight some of the key improvements to star detection, star measurement, rate estimation, and catalog generation algorithms. Presently the ST-16 sensors on Skysat-1 are reporting availability of around 98% and cross-axis accuracies of roughly 10 arcseconds over an entire orbit in a nominal Earth-observing attitude.

INTRODUCTION

The Skybox Imaging SkySat-1 satellite launched on November 21, 2013 carrying a pair of Sinclair Interplanetary ST-16 Star Trackers. Amongst more prominent objectives, one of SkySat-1's goals was to validate the performance of the ST-16. While the sensors were able to achieve a stellar lock, the availability and accuracy of the ST-16s was far worse than expected: the sensors frequently lost their lock, they saw fewer stars than expected, and the attitude fixes were far noisier and less consistent than observed during ground testing.

Concerned by these results, engineers at Skybox Imaging (SB), Sinclair Interplanetary (SI) and Ryerson University (RU) embarked on an aggressive and comprehen-

sive flight qualification program to understand the causes of these problems and to re-attain our expected performance targets. Just over two months later (February, 2014) we made the last of a sequence of software, catalog and parameter modifications that have met these goals. In this paper we detail the series of individual improvements that helped us remedy the poor performance of the star trackers. We concentrate on technical improvements, but also discuss the collaborative framework between vendors, customer, and researchers that made success possible.

The remainder of this section introduces SkySat-1, the ST-16 star tracker, and describes our approach to diagnosing performance problems. The second section describes

the logistics of collecting data, developing and uploading new sensor software, and the benefits of manufacturer-user relationships. The remaining sections of the paper go onto to describe the various problems we observed, and how we fixed them.

Skybox Imaging and SkySat-1

Skybox Imaging, a venture-backed information and analytics company, recently deployed their first high-resolution imaging satellite, SkySat-1. SkySat-1 was launched on November 21, 2013, out of Yasnny, Russia into a 580 km Sun-synchronous orbit. The satellite has been collecting sub-meter resolution images and video, demonstrating a new high-end platform for monitoring global activity on a daily basis. Skybox will launch and operate several additional satellites in the near future, creating a massive and unprecedented source of data for mining and analytics.

For the imaging products to be useful, the geographic location of each satellite image must be known to high accuracy. The demand for high-quality geolocation knowledge can only be met with precise onboard attitude sensors, namely, star trackers. Historically, volume, mass, and power requirements limited star trackers to large spacecraft. The mechanical and software complexity of early star trackers made them among the most expensive sensors, and arguably the most difficult to calibrate and operate.

Through the utilization of modern, commercial, high-performance electronics, the ST-16 star tracker has achieved levels of performance that have previously been exclusive to its larger predecessors. As a small satellite heavily constrained by mass and volume, SkySat-1 has taken full advantage of the ST-16 capabilities in a miniaturized package. The star trackers have played a major role in making SkySat-1 a world-class earth observation satellite.

The ST-16 Star Tracker

The ST-16 star tracker is a relatively new nanosatellite-class star tracker that became available in 2011. It was developed through a collaboration between Sinclair Interplanetary, the University of Toronto's Space Flight Laboratory, and the Space Avionics and Instrumentation Laboratory at Ryerson University. It is distinguished from other devices by its powerful internal computer, small size and weight, and low power. Although it is able to fit into a nanosatellite, it is also applicable to larger spacecraft (microsatellites). Some key specifications are listed in Table 1, and image of the sensor is shown in Figure 1.

Table 1: Key Parameters of the ST-16 Star Tracker

Accuracy	< 7 arc-sec RMS cross-boresight < 70 arc-sec RMS around boresight
Availability	> 99.9%
Size	59 x 56 x 31.5 mm
Mass	≈ 90 g
Field of View	7.5 deg (half axis)
Exposure Time	100 ms
Catalog	3746 stars

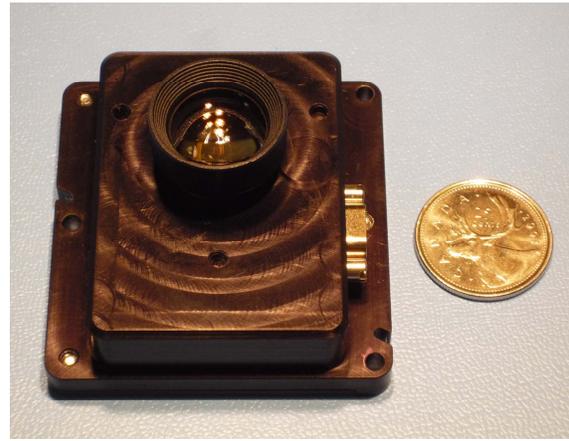


Figure 1: The Sinclair Interplanetary ST-16 Star Tracker.

For more information, please see ^{1,2,3}.

The ST-16 star tracker has had initial market success with over 40 flights units delivered to customers. The two units onboard SkySat-1, designated in this paper as Sensor-A and Sensor-B, were among the first off the production line and together represent the maiden voyage of the ST-16 into space.

Diagnosing Performance Problems: Our Approach

Initial telemetry returns — obtained within the first 24 hours after launch — indicated that something was wrong with both ST-16s onboard SkySat-1. The accuracy and availability of both units was significantly below specification (i.e., compare Table 1 to Table 2). Using an initial batch of collected telemetry and full-frame images, we immediately began to look for the source of problem. Failing to find a single dominant fault, we realized that the poor performance of the star trackers was due to a combination of several smaller faults. Everything was working, but nothing was working well. At this point, we took a step back, and examined the problem from a system-level

Table 2: Measured Performance of SkySat-1 Star Trackers

Operating Scenario	Time Period	Availability (%)		Accuracy (arc seconds)		
		Sensor-A	Sensor-B	Roll Error (1σ)	Pitch Error (1σ)	Yaw Error (1σ)
Tumbling	Launch	34	31	–	–	–
Tumbling	Init. Recal.	80	61	–	–	–
Pointing	Init. Recal.	89	85	21.6	61.5	161.9
Pointing	Final	99	98	10.0	23.0	74.0

perspective.

Diagnosing poor performance relies on an understanding of the nominal processing sequence on the sensor. This sequence has five steps: imaging, star detection, measuring star positions, matching, and finally the attitude solution. Table 3 provides a short explanation of each function.

The overall effectiveness of the sensor processing can be measured using three system-level metrics:

- Accuracy, which describes the uncertainty in the orientation measurement.
- Availability, which describes the fraction of time for which attitude solutions are available.
- Bad Match Frequency, which represents the likelihood that the sensor incorrectly matches the stars in a given scene (this, in turn, will yield a very high attitude error).

These top level metrics are useful for describing the general performance of a star tracker. However, in a case where one, or all, of these metrics is low, they do not illuminate the source of the problem. Breaking down these top-level metrics leads a more descriptive set of performance drivers which align roughly with the basic operations of a star tracker. These are listed in Table 4.

Identifying these lower-level performance drivers and their contributing functions and processes enabled us to trace the performance of subsystems back to top-level performance metrics. This illuminated which metrics need the most revision, and what mechanisms we should utilize to make those revisions.

DATA COLLECTION, SOFTWARE UPDATES, AND KEY ASPECTS OF WORKING AS A TEAM

The debugging and development process for the star trackers runs very differently when the sensor is on-orbit compared to when it is in the lab. Spacecraft activities are scheduled several days in advance, and so data collection

Table 3: Typical Operational Chain of a Star Tracker

	Process	Description
1	Imaging	Optically focus the light onto the detector and digitize the image.
2	Star Detection	Image processing to identify star candidates and isolate their corresponding intensity patterns.
3	Star Vector Computation	Centroid each star image and calculate the corresponding star vector using a camera model.
4	Matching	Identify the observed stars from the onboard star catalog
5	Attitude Solution	Estimate the attitude of the ST from the identified star vectors.

and software updates must be coordinated with the operations team and must respect constraints imposed by other activities.

The very high value of SkySat-1 as an operational orbital asset produces conflicting pressures. On one hand, there is a deeply conservative desire to not endanger the satellite with hasty or untried software loads. On the other hand, there is a desire to rapidly commission the spacecraft to begin producing a valuable payload data product.

Data Collection

Skybox began collecting star tracker data immediately after launch. Within 24 hours, while the spacecraft was still in its initial tumble, we knew that we had a problem.

Three different types of star tracker telemetry were collected. Initially, when we knew very little about what might be going wrong, we downloaded a number of full-frame images from the detector. These images are approximately 10 MB, and so each downlink is a significant burden on the ground segment. Analysis of these images was key to our understanding of changes in the point spread function of the optics, as well as true characterization of the environmental noise.

Next we switched to collection of full telemetry packets

Table 4: Performance Drivers of Star Trackers

Metric	Criteria	Contributing Effects and Processes
Star Detection Reliability	Probability of detection, False detection rate	Focus, Chromatic aberration, Detection routines, Baffle artifacts
Star Measurement Accuracy	Noise equivalent angle, Calibration residual	Centroiding, Camera model, ERS compensation
Matching Robustness	False detection tolerance	Matching acceptance tuning, Acceptance criteria
Catalog Quality	Catalog completeness, Catalog accuracy	Star population, Star positions (proper motion, binary stars), Chromatic aberration, Stellar aberration

from the star tracker. These packets are approximately 2 kB each, and so 5000 can be downloaded for the same cost as a single image. The telemetry packets contain a list of the detected star centroids, along with their matched catalog ID. The on-orbit recalibration process requires a large number of full telemetry packets as input to the optimizer.

Finally we moved to collection of primary telemetry only. These packets are very small, and contain only the quaternion, angular rate, and validity flags. Primary telemetry can be collected at full rate over the course of an entire orbit. By comparing the primary telemetry streams from both star trackers the star tracker error can be determined.

Sensor Software Updates

Following the discovery of a specific problem, the engineering team would then work to find a solution. Depending on the type of the problem, testing the solution would either involve only simulated data, or additionally, post-processing existing on-orbit data (telemetry and images). To minimize risk to the spacecraft, changes to the ST-16 flight software followed a specific quality assurance path before being uploaded to SkySat-1. This path is summarized by the following 6 steps:

1. Develop fix as standalone tool, typically in MATLAB, and test using a limited set of either simulated, or actual, on-orbit data.
2. Implement the fix into a offline PC-based version of the sensor software, and assess performance improvements by reprocessing collected on-orbit data. Reprocessing of telemetry utilizes only a partial implementation of the sensor software (beginning with star centroids and brightnesses). Reprocessing of images utilizes the entire end-to-end sensor software.
3. Assess performance improvements in all scenarios (different orientations, rates, etc.) by reprocessing a

large suite of synthetic images using the PC-based sensor software.

4. Verify performance improvements with hardware-in-the-loop (HITL) tests that utilize an ST-16 engineering model (EM) with a small subset of the synthetic images from step three loaded into memory.
5. Upload sensor software to SkySat-1.

Over the course of this work we updated almost every byte of software and data on the star trackers. There are four different classes of configuration change, ranging from least to most extensive.

The star tracker contains dozens of non-volatile parameters which control its processing. These are ‘knobs’ that were introduced to allow for on-orbit tuning. Parameter modifications are trivial to revert, and changes were made on an almost daily basis with minimal oversight burden.

Each star tracker has a unique calibration structure which contains the geometric model of its camera. Several new calibration structures were generated by the Ryerson team based on modelling of the on-orbit telemetry. The upload of new calibration structures was straightforward.

Several new star catalogs were generated as we came to understand the deficiencies in our original data. The star tracker stores a star table (70 kB), and then a massive explicit enumeration of all of the potential triangles (40 MB). Fortunately in the months prior to launch we had written the software to allow the star tracker to build its own triangle table on-orbit. We were able to upload a new star table, and then send a command to begin the 3-minute process of triangle table generation.

Finally, we performed uploads of new executable code for the star trackers. Unlike the other data uploads, new software revisions required an extensive quality assurance (QA) process at Skybox before they could be uploaded. In the end only two code updates were performed.

After each configuration change was made a standard star tracker test was run. Known as the '4 orbit test', this collected a particular set of data with the spacecraft in a particular set of attitudes. This allowed for apples-to-apples comparisons to determine whether the configuration change was helpful.

Collaborative Relationship

Restoring the star trackers to full function was do-or-die for both Skybox Imaging and Sinclair Interplanetary. Skybox had invested in the spacecraft, and Sinclair in the star tracker product, and neither could afford to fail. While stressful, this unity of purpose was in no small part responsible for timely success. Skybox operations was extremely accommodating in collecting and delivering large quantities of data. Sinclair and Ryerson focused exclusively on this problem for a two month period. In a more relaxed and less motivated environment the necessary advances might not have been made.

IMAGING AND STAR DETECTION

One of the first insights into the poor performance of both sensors was the low reliability of star detection. Both star trackers were detecting significantly fewer stars than expected in each scene. On top of this, Sensor-B was detecting even fewer stars than Sensor-A. Detection performance is directly related to the signal-to-noise ratio (SNR) of star's intensity distribution on the image detector. There are two main drivers that impact the SNR of an imaged star: focus of the sensor optics, and the detection logic used to separate signal from background image noise. Given that detection performance directly impacts matching performance, and attitude accuracy, diagnosing poor detection performance was a priority early in the qualification campaign. This section describes our analysis of sensor focus and the performance of the ST-16's star detection routines.

Focus Evaluation

Analysis of full-frame images showed some stars that appeared to be larger, and as a result, dimmer in each individual pixel, than those which were observed during ground testing. In light of the knowledge that both ST-16s were detecting fewer stars than expected, one of the potential causes we examined was a change in focus. A minor change in focus is expected due to the change in imaging conditions (air vs. vacuum) between the lab (where sensor focus is set), and the ultimate operating environment, space. However, this minor change would not account for the observed changes in star size as seen in the on-orbit images.

A typical method to quantify the focus is to examine the encircled energy from a star source. This is done during the initial focusing and calibration of each sensor. The star source in the lab has a particular color temperature of 2800K, and hence its only representative of a small fraction of the stars within the ST-16's catalog. In addition, the shape and size of a star image vary throughout the sensor field of view (FOV) due to lens aberrations. When comparing these lab results with an on-orbit image that contains stars with vastly different star colors, a specific assessment of sensor focus is difficult. We found that we just didn't have enough images of a star similar to our lab source, spaced through the sensor FOV, to conclude a change in focus.

However, the few on-orbit images we did have allowed us to make relative comparisons of star size between images. By examining the appearance of similar stars, in various parts of the sensor FOV, we were able to conclude that changes in the apparent star size were not due to a bulk change in focus, but rather a varying change in focus dependent on the color temperature of the star we were examining. In more simple terms, we were observing the chromatic aberration of the ST-16's optics.

Due to the fact that we have no control over the physical focus of the ST-16 remotely, we could not change how the stars appeared in the image. However, knowing how these stars looked allowed us to make logical modifications to the ST-16's detection routines which ultimately lead to significant improvements in the detection reliability.

Tuning Detection Parameters

Following directly from the assessment of focus using full images, we revisited the star detection routines used by the ST-16. In the downloaded images from both sensors we noticed that there were many stars that were easily discernible by eye, but that the detection routines failed to pick up. The source of this problem was not singular nor immediately obvious. Stepping through the detection routines, we found that various control parameters which we had set to reduce false detections were much too conservative. Using downloaded images, combined with knowledge of which stars must be in the same scene from various catalogs, we iteratively tuned detection parameters to maximize detection performance.

Before we can describe our tuning approach, we must introduce the star detection logic of the ST-16. This routine can be summarized by describing a moving window, and three threshold parameters:

- **Moving Window.** In order to deal with spatial gradients in the background illumination of a star scene, the ST-16 detection scheme utilizes a moving average filter to estimate the local background level for each pixel. Typical processing utilizes a window 128 pixels in width, and a single pixel in height. The average value of this window is used to define the local background level for the current pixel.
- **Lit Pixel.** This threshold defines the minimum intensity, relative to the local background, of an image pixel that is considered to be lit by star light, as opposed to just sensor noise. Pixels above this threshold are labeled *lit pixels*. The launch configuration utilized a lit pixel threshold of 90 detector counts, out of a possible 4095.
- **Blob Size.** This defines the minimum number of contiguous lit pixels that each candidate star must possess before it can be considered as a valid detection. Configuration at launch required at least 6 contiguous pixels for each valid star detection
- **Integrated Intensity.** This value describes the minimum integrated intensity (summed detector response) of all contiguous lit pixels that compose a candidate star. Candidate stars above this threshold are considered valid detections. Configuration at launch utilized an integrated intensity threshold of 1000 detector counts.

Stars that satisfy these criteria are considered candidate star detections and are passed to the next portion of the processing chain. The ST-16 employs various mechanisms in the matching phase to discard any false stars from this step.

Our approach to tuning detection parameters was to reprocess downloaded images with varying parameter values and examine the resultant star detections and subsequent star matches. Our goals were to maximize star detections, while minimizing false detections. To aid in this process, we assembled custom mini-catalogs for each star scene we had an image for. These mini-catalogs catalogs included additional dim stars on top of the existing stars listed within the onboard ST-16 catalog. By surveying various lit pixel, blob size, and integrated intensity parameters combinations, we created a map of expected star detections given a set of parameter settings. Using this map, we were able to suggest updated detection parameters for use with both sensors. After implementing these changes, we saw the availability of both units increase to above 98% (see last row of Table 2).

Table 5: Tuned Detection Parameters for Sensors A and B

Parameter	Value at Launch	Final Value	
		A	B
Lit Pixel	90	70	60
Blob Size	6	6	6
Integrated Int.	1000	800	500

Using the parameter settings at launch, our analysis of the downloaded images showed the detection performance of Sensor-A exceeds that of Sensor-B. While we were never able to conclude a reason for this discrepancy, we were able to remedy this difference by lower the detection threshold on Sensor-B. Theoretically this does make Sensor-B more susceptible to false detections, but from our analysis of the various on-orbit images, we didn't see any evidence of this in the top-level performance metrics.

Baffles and Stray Light

Even prior to launch, we anticipated that some refinement of the ST-16's stray-light handling logic would be necessary. This subset of the the star detection algorithms is responsible for improving the performance of the sensor when bright objects fall within the FOV. Possible sources of non-stellar light include the moon, the earth, planets, other satellites, and the bright edges of the baffle vanes. These sorts of scenarios are difficult to replicate in ground tests and simulations, so rather than rely on ground tests alone, we developed a flexible processing framework that could be tuned after launch.

Instead of a single block of stray-light-control logic, we formulated a number of simple heuristics designed to minimize the impact of stray light on the processing chain:

1. **Adaptive Threshold Logic.** This algorithm increases the sensor's tolerance to bright backgrounds and large extended sources. In normal processing, pixels bright enough to warrant further attention are flagged on the first pass through the image. This routine excludes affected areas of the detector from further attention based on a simple image contrast calculation.
2. **Region-of-Interest (ROI) Calibration.** When baffle vanes extend into the field of view of the sensor, the areas of the detector behind the baffle vanes will not be in view of the stars. The useful region of the detector is static and can be measured during calibration. Because bright reflections from vane edges can create illumination gradients near the edge of the ROI, our routines require that any candidate

stars must be separated from the boundary by some dark pixels.

3. Large, Bright Object Rejection. Bright areas of the image that had good contrast against the image background could be rejected based on the area of the object.

These first three routines attempt to improve processing effectiveness by reducing the likelihood that parts of the image would be mistakenly detected as stars. Implemented before launch, we were aware that the routines would require some tuning once on-orbit. During the initial qualification we discovered persistent performance problems caused by falsely detected stars. Analysis of the sensor telemetry led to the following additional improvements.

1. Star Cluster Rejection. Even with the above strategies, diffuse stray light would sometimes be interpreted as a cluster of stars inside the FOV. Figure 2 shows the telemetry from an exposure taken with Sensor-B soon after launch. The telemetry shows three 'real' stars, along with a dense line of false detections caused by a bright reflection off a baffle vane edge. These sorts of clusters can also sometimes be found next to bright objects like the moon. Although the star matching algorithms are reasonably tolerant of false stars, clouds of false detections greatly increases the chance of a false match.
2. Tight Matching Tolerances. With improvements in the accuracy of the catalog and centroid determination, the star matching routines could be much stricter about the expected geometry during the star matching process.
3. Strict solution acceptance criteria. One of architectural principles in the ST-16 software was stateless operation: each measurement did not rely on prior measurements. The sensor performance encountered post launch required us to make some small adaptations to the operating concept to help reduce the number of false matches. From the flight data we could see that if the sensor could match four stars, the matching solution was almost always correct. Although most three-star matching solutions were correct, there were an unacceptably large number of incorrect matches when only three stars were identified. To help balance the need for high availability, even when only three stars were detected, and a very low tolerance for false matches. we introduced some very limited reading-to-reading persistent state: any match with four or

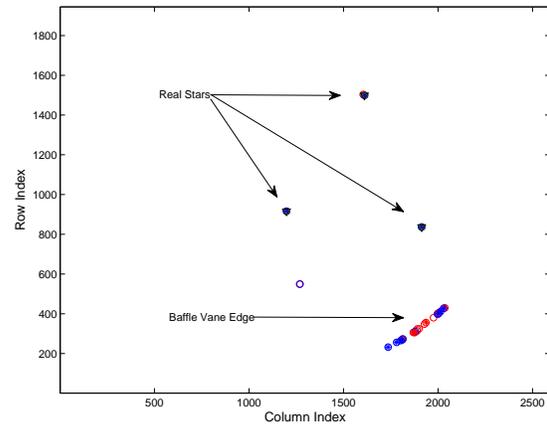


Figure 2: False star detections from edge of baffle vane.

more stars would be considered 'good' and a three-star match would be considered good if the attitude estimate was not too different from an immediately prior good reading.

COMPUTING STAR VECTORS

Despite the fact that most scenes had sufficient star detections, most of which were successfully matched, the attitude accuracy was significantly worse than expected. In addition to generally poor accuracy, many of the returned attitude solutions were bad matches to incorrect star scenes. There are many separate factors contributing to these observed drops in performance. The first issue we discovered was that our measured arc-lengths between stars were different from known values. We traced this symptom to three causes: inaccurate centroiding, improperly set camera model parameters, and a noisy rate estimator. This section describes how these three causes introduced error into our arc-length measurements, and the corresponding fixes we made to restore the performance.

Centroiding Algorithm Improvements

The lowest-level processing that can affect the measured arc-lengths between stars is the determination of the location of the star on the image detector. This routine is commonly referred to as centroiding, because the typical way to compute this is the first-moment of illumination of the observed intensity pattern⁴. During star detection, routines onboard the ST-16 identify and group pixels that belong to the same star image into a *blob*. In an attempt to minimize the contribution of image noise, the centroid was calculated using only pixels within the blob. After careful examination of full-frame images, we discovered

that this decision was negatively impacting our centroid accuracy. When viewing dim stars in various unfavorable imaging conditions (high rate, or around bright bodies), we found that many times, only a portion of the actual star image was identified by the blob. If the portion was asymmetric about the true star centroid, this would result in an erroneous centroid measurement — biased towards the group of selected pixels. To correct this behavior, we developed a window-based centroiding routine that utilizes a circular window to select which pixels are to be used in the centroid estimate.

In ideal lab conditions, the mean centroid accuracy of the ST-16 was measured to be approximately 0.2 pixels or $0.44\mu\text{m}$ ($1-\sigma$). These conditions are set to produce the best possible SNR of the imaged star. The lab star source is tuned to be as bright as possible without saturating the ST-16's image detector. This enables the best separation of the star's intensity pattern from the background image noise. As the brightness of the star decreases, the SNR of the star image drops, and it becomes progressively more difficult to accurately separate the blob from the image background. In ideal imaging conditions (no rate, and no image background gradients), this results in an approximately symmetric loss of pixels about the true centroid. In non-ideal imaging conditions (especially during high rate), a low SNR enables noise contributions to break up the blob. Current detection logic utilizes 4-connectivity to identify pixels within the blob. When the blob is thin, and long, as it would appear at high rate, noise contributions can pull pixels critical to establishing a connection between parts of the blob, below the detection threshold. This separates the blob into two separate candidate stars, ultimately leading to centroid error.

After experimenting with various changes to the ST-16 detection logic, we determined that the approach most resilient to unfavorable imaging conditions was a window-based approach. Following star star detection, an initial centroid estimate is computed using the existing technique. We then revise this estimate by computing the star centroid using all pixels (regardless of whether they are part of the blob) within a circular-shaped window centered on the initial centroid estimate. The ideal radius of this window was empirically determined to be 7 pixels by analyzing full-frame on-orbit images.

On-Orbit Recalibration

Before the matching process can begin, each detected star must be converted into a star vector through the use of a camera model. The camera model used onboard the ST-16 describes the optical characteristics of the sensor with 11 parameters, see Table 6. These parameters are de-

termined in the lab through an offline calibration process that utilizes a single star source, and a 3-axis motorized gimbal. Shortly following launch, the ST-16 was recalibrated on-orbit using collected telemetry. This type of recalibration is typical for star trackers, primarily due to changes in focus as a result of vacuum and having pristine access to real star sources. The role of the ground calibration is simply to enable successful matching of an initial batch of on-orbit telemetry, which is then used for recalibration. Following an initial recalibration process, subsequent analysis highlighted two potential pitfalls of our recalibration approach:

- Our camera model contained some parameters that had minimal impact on the calibration residual, and
- Chromatic aberration from the ST-16's optics caused a star-color-based dependence on focal length.

This subsection discusses our analysis of these pitfalls, and describes our approach to implementing solutions.

Camera Calibration

The initial recalibration of the ST-16s allowed significant improvements in sensor availability and a reduction in false matches, see Table 2. However the overall performance of the sensors — particularly Sensor-A — still fell below expectations. A detailed examination of the effect of the calibration procedure offered some insights into further improvements.

Although the lab calibrations and on-orbit calibrations are parameterized in a similar way, the cost function for the two optimizations are quite different. The lab-calibration uses the residual error between the star vector estimate derived from the sensor, and the true star vector obtained from the test platform kinematics. On-orbit, because we do not have any source of attitude truth, we must adopt a different scheme. The cost-function for the on-orbit calibrations considers the arc-length between pairs of stars. As long as we have a good match for these stars, the true values of these arc-lengths can be calculated from the star catalog and is independent of attitude.

In our lab-calibrations we used an 11-element parameter set that included the scale parameter, g_y . This scale parameter represents a slight variation in the pixel pitch between the x - and y -directions. The original ground calibration yielded very similar results for both sensors (see Table 7), and we assumed that the difference in values was attributable to the inherent calibration and sensor variability. When we recalibrated Sensor-A on-orbit, the

Table 6: Summary of Camera Calibration Parameters

Quantity	Symbol	Number of Parameters	Lab Calibration	On-orbit Calibration
Focal Distance	f	1	Yes	Yes
X/Y Pixel Scale	g_y	1	Yes	No
Optical Centre	m_0, n_0	2	Yes	Yes
Radial Distortion	b_1, b_2	2	Yes	Yes
Axial tilt	a_1, a_2	2	Yes	Yes
Detector Rotation	$\alpha_1, \alpha_2, \alpha_3$	3	Yes	No

optimization found a set of camera parameters that optimized the cost-function value. However, when we looked at the arc-length estimates for a number of star pairs (see Table 8, third column), there was a consistent under-estimation of many of these distances — with the parameters that the optimization could vary, this error could not be corrected. Similar tests with the recalibrated Sensor-B revealed no evidence of the same trend. The calibrated values for g_y represented a 0.5 – 0.8 pixel offset across the whole detector. Although we could envision a rationale for g_y to differ from unity (i.e. perfectly square pixels), there was little justification for individual detectors having different values. When we repeated the on-orbit calibration using the same g_y (from the Sensor-B calibration) in both sensors, the consistent arc-length bias disappeared (arc-length variability remain unchanged). As a consequence, we have removed g_y from the laboratory calibration optimization and simply use a constant value.

The results from this series of tests highlighted a number of potential problems with the laboratory calibration procedures. The original g_y used on Sensor-A allowed the on-orbit recalibration to get stuck in a local minimum. In contrast to the weakly-separable nature of some of the other parameters — e.g., (m_0, n_0) and (a_1, a_2) — a more complex interaction between g_y and several other parameters permitted distinct, steep minima in the optimization. Early sensitivity studies showed that calibrations were extremely sensitive to both f and g_y , but while f had a single optimum value, the non-unique nature of the g_y optimization was overlooked. Night-sky qualification testing of other EM units was generally successful and the noise levels in the night sky tests were generally low. Consequently there was little impetus to closely evaluate the statistics of the arc-length measurement performance for systematic biases. Reviewing many of the calibrations for the EM units used in qualification revealed that not only were incorrect values of g_y often found by the laboratory calibration optimization, but that the arc-length biases are clearly present in the night-sky tests.

Table 7: Comparison of Pixel Scale Values

Sensor	g_y
A	0.99975
B	0.99960

Chromatic Effects

Following from the initial focus assessment using on-orbit images, we believed that our lens possessed significant chromatic aberrations. Satisfied that we had mitigated the effect of these aberrations on star detection by tuning detection parameters, we investigated the impact they had on discrepancies observed in the measured arc-lengths, namely through a change in focal length. Our investigation consisted of separating the collected on-orbit telemetry into two groups based on star color, and then performing separate limited recalibrations with each subset. To simplify the analysis, these recalibrations consisted of only letting the focal length vary, while keeping the remaining camera model parameters constant. Group A consisted of *hot* stars which includes spectral classes: O, B, A, and F. While group B consisted on *cold* stars which includes spectral classes; G, K, M, and C. The result of this investigation showed that the separate recalibrations lead to different determined focal lengths. To illustrate this, Figure 3 shows the residual of the recalibration (defined as the RMS of the arc-length error) as a function of focal length, for both spectral groups. The optimal focal lengths determined by the recalibration procedure for each group lie at the vertex of each parabola.

Figure 3 highlights the dependence of optimal focal length (as defined by the on-orbit recalibration) on star color. Initially, we considered a solution for this dependence that involved implementing a variable focal length, based on matched star color, to revise the computed star vector. However, implementing this solution would require substantial modifications to ST-16 processing routines. Due to time constraints, we chose to utilize the existing recalibrations, not based on spectral class. These corresponding focal lengths are also shown in Figure 3.

Table 8: Sensor-A Recalibration Effect

Pair	Arc-length (deg.)	Mean Arc Error (asec)	Std. Dev. Arc Error (asec)	Mean Arc Error (asec, g_y -fixed)	Std. Dev. Arc Error (asec, g_y -fixed)
1	0.87	-6.53	11.14	-0.21	10.60
2	3.38	-6.79	10.52	3.77	10.20
3	3.43	-8.16	9.48	1.99	9.34
4	0.79	-5.23	9.84	-1.67	9.82
5	3.89	-17.41	11.18	-3.31	11.28
6	3.79	-18.07	9.23	-5.30	9.33
7	2.03	-6.50	12.71	-0.09	12.05
8	5.35	-46.03	7.44	-6.60	7.71
9	4.77	-38.70	11.03	-6.44	10.30
10	3.88	-34.88	9.74	-5.65	9.61

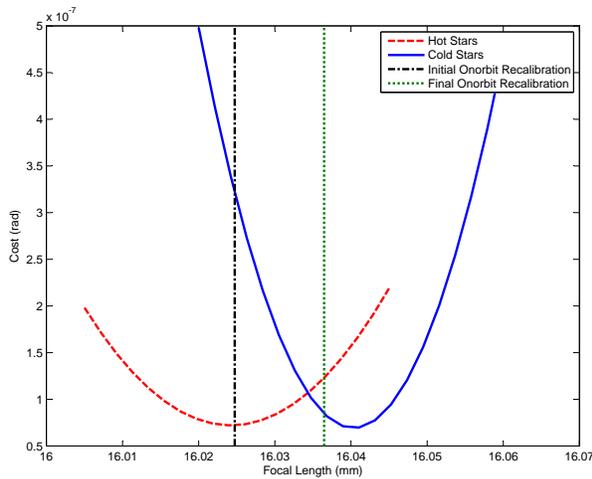


Figure 3: The effect of star temperature on focal length.

A second generation star tracker, titled the ST-16RT, is already in the final stages of development. Leveraging the lessons learned from this analysis of chromatic aberrations, the ST-16RT utilizes a custom lens design with significantly better chromatic performance.

Need for Better Rate Estimates

The detector integrated with the ST-16 uses an electronic rolling shutter (ERS) to control pixel exposure. The primary effect of the ERS is to introduce a time offset between the exposure of different rows on the detector. If uncorrected, this offset skews the observed inter-star geometry and adds error to the attitude estimates. Our algorithm for ERS compensation is described at length in Enright and Dzamba⁵. The effect of the rolling shutter is relatively simple to correct if the angular velocity is known. Thus with a calibrated star tracker we can calculate star vectors from the detector centroid position, (m_i, n_i) , and

angular velocity, ω :

$$\mathbf{b}_i = F(m_i, n_i, \omega) \quad (1)$$

The ST-16 has no internal rate sensor, so our ERS compensation algorithms estimate the angular velocity of the sensor using paired centroids in two successive images:

$$\tilde{\omega} = G(\mathbf{m}_A, \mathbf{n}_A, \mathbf{m}_B, \mathbf{n}_B) \quad (2)$$

This velocity estimate, $\tilde{\omega}$ is used in (1) to remove the effect of the rolling shutter.

The effectiveness of the ERS correction depends on the accuracy of the velocity estimates. The difference between the estimated and the true angular velocity is the angular velocity error: I.e., $\tilde{\omega} = \omega + \delta\omega$. Figure 4 shows how $\delta\omega$ affects the ESOQ2 attitude solution. This plot was generated using a number of simulated scenes across the sky and shows the magnitude of the optimal rotation between the two sets of vectors $\mathbf{b} = F(\mathbf{m}, \mathbf{n}, \omega)$ and $\tilde{\mathbf{b}} = F(\mathbf{m}, \mathbf{n}, \tilde{\omega})$. The exact distribution of stars in each scene and the true value of the angular velocity will cause some spread in the results, but there is an apparently linear relationship between the rate error and the additional attitude error. The attitude solution is much more sensitive to x -axis (cross-boresight) error components than z -axis (about-boresight) components, but since the largest errors in the velocity estimates are about the boresight axis, the $\delta\omega_z$ contributions are usually dominant.

In normal operations the velocity estimates on the ST-16 are derived from pairs of star centroids identified in sequential images. These pairs of images are $\delta t = 0.1$ s apart. As is typical with finite-difference type rate estimates, short sampling intervals effectively amplify any

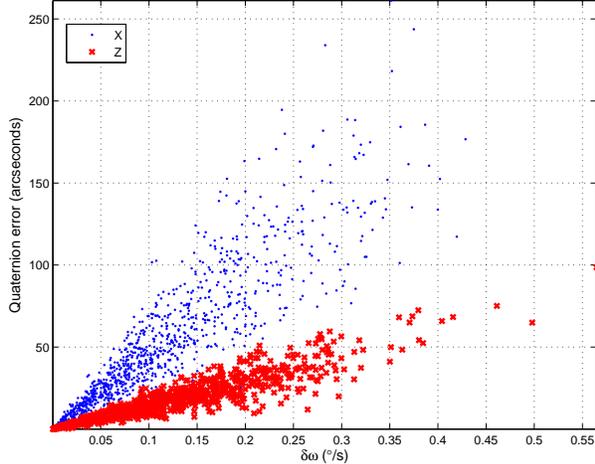


Figure 4: Attitude sensitivity to angular velocity error.

angular errors in the measurements. A one-dimensional approximation to the sensor motion provides some insight into the relationship between angular orientation errors, σ_θ , and the angular velocity error, σ_ω :

$$\sigma_\omega \approx \frac{\sqrt{2}\sigma_\theta}{\delta t} \quad (3)$$

Here the errors are assumed to be zero mean and characterized by their standard deviations.

From this analysis, it is clear that centroid errors can have a compounded effect on the attitude accuracy of the ST-16. Not only do noisy centroid measurements lead to noisy star vector measurements, but these centroid errors also lead to inaccurate rate estimates and additional geometric warping from the ERS effects. Once the star vector errors exceed a certain level, scene-matching becomes more difficult and availability will also suffer.

Initial night-sky testing with the ST-16 did not reveal significant problems with the ERS performance (consider the first row of Table 9). About the x - and y -axes the ERS rate error caused approximately 2 arcseconds of attitude error; about the z -axis the contribution was about 5 arcseconds. These error levels represented sizable but manageable fractions of the error budget on these axes. After launch, the rate errors were significantly higher than the ground tests (the ground tests reported in this table were not made with the same flight unit, but with an engineering model calibrated to about the same level of performance). In this regime, the error contributions were about 7 arcseconds and 25 arcseconds, respectively.

Table 9: Rate Estimator Performance

	δt (s)	Rate ($\times 10^{-3}$ deg/s)		
		σ_x	σ_y	σ_z
Single Reading (night-sky)	0.1	3.89	4.03	39.6
Single Reading (on-orbit)	0.1	17	13	175
Reading-to-Reading (calculated)	0.5	3.40	2.60	35.0
Sparse Dataset (calculated)	40	0.042	0.032	0.438
Sparse Dataset (on-orbit)	40	0.211	0.168	0.614

Although poor centroid accuracy may have been the root cause of the poor rate estimates, the problem was made worse by the short time between observations. As part of the flight qualification campaign, we introduced additional logic into the flight software on the ST-16. Instead of relying solely on the rate estimates derived from a single sensor reading, the sensor would also calculate an estimate based on the reading-to-reading change in orientation. Provided that the two estimates were close, the sensor would prefer to use the estimate derived from reading-to-reading motion. In normal 2 Hz operation, the longer interval between measurements should give a five-fold reduction in noise. Note that this remains a finite difference rate estimate, we have merely increased δt .

We can explore the potential advantages of even longer time intervals when estimating rate; the last few rows of Table 9 show the calculated response and observed noise when we reprocess a sparsely sampled telemetry set with this reading-to-reading algorithm. Data from several orbits are shown in Figure 5. This figure contrasts the single-reading estimates (top) with those derived from the sparse attitude measurements (bottom). Sudden jumps in the ω components indicate invalid reading-to-reading measurements and the use of the single-reading value. Although the noise observed in our orbital data during this last test is good, it is not as low as we would expect from (3). A closer look (Figure 6) at $\tilde{\omega}_x$ reveals that much of the variability that contributes to this discrepancy is actually a periodic variation (once per orbit). This variability in $\tilde{\omega}_z$ is less pronounced than that observed in the other axes, so the predictions are a better match to the observations. The root cause of this periodic term is not fully understood — it could represent a sensor bias or actual motion — but the fact that it does exist explains the increased standard deviation in the rate estimates.

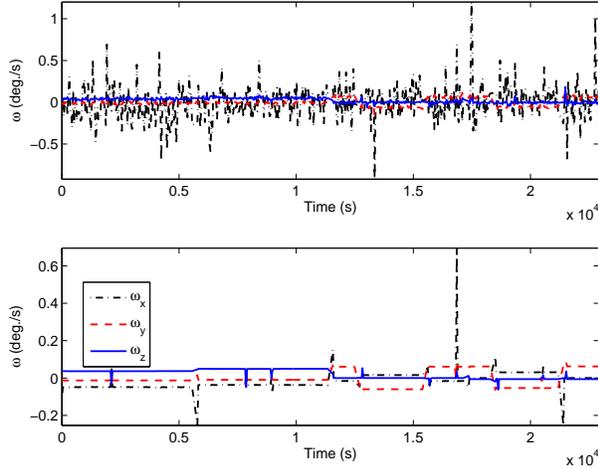


Figure 5: Effect of measurement baseline on $\Delta\omega$ showing single-frame estimates (top) and estimates with $\delta t=40$ s (bottom).

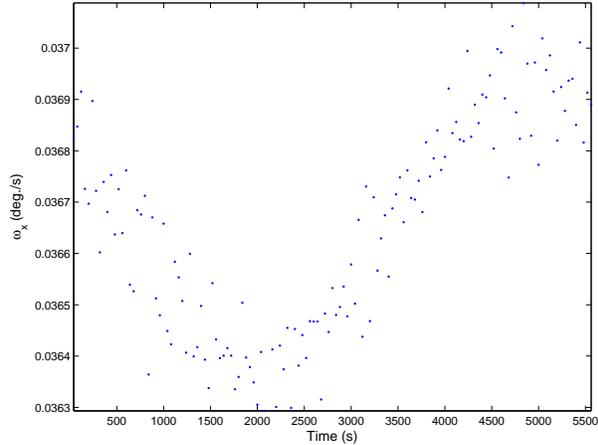


Figure 6: Rate error orbit dependence ($\delta t=40$ s).

The 40 s time difference described in the preceding tests is longer than what is currently in use on SkySat-1. These results are useful, however, in exploring the potential benefits of rate estimates derived over longer time intervals. When angular accelerations are low, the long intervals provide better noise rejection, but during more agile maneuvering the estimates will lag the actual motion. We are currently evaluating a number of strategies to balance noise and responsiveness including adaptive finite difference techniques and recursive estimators.

MATCHING AND THE ATTITUDE SOLUTION

Following the determination of star vectors, the next major operation of a star tracker is to identify these stars by matching the patterns they form to an onboard star cata-

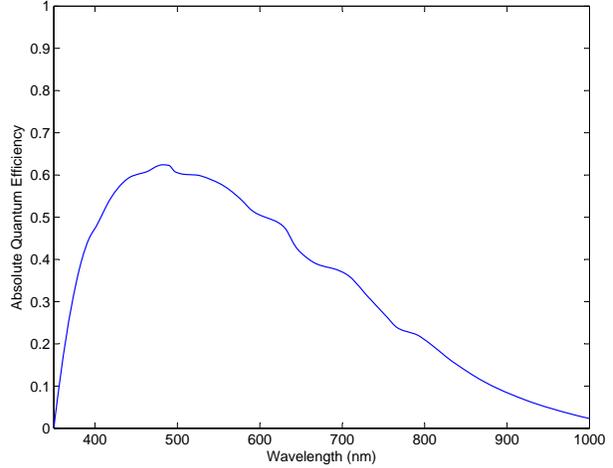


Figure 7: Quantum efficiency of the ST-16 image detector.⁶

log. A successful match allows for the final step of calculating the attitude of the star tracker. Based on analysis of the downloaded full-frame images and telemetry, we discovered that our matching routine was failing to identify certain star scenes, despite these scenes having a sufficient number of star detections. This was due to two problems:

- Some stars we detected were not part of the ST-16 catalog.
- Patterns (triangles) formed by some of these stars were different from their definitions in the catalog.

Since we had already addressed performance issues concerning star detection and star vector calculation, we concluded that these problems must be a result of a poor selection of stars and inaccurate star information. This section describes modifications to the ST-16 star catalog to improve the star matching performance.

Catalog Repopulation

The original star catalog onboard the ST-16 was constructed using stars listed in the Sky2000 stellar database. Stars were chosen based on their visual magnitude, which represents the typical response of the human eye. The spectral response of the ST-16 detector however, covers the entire visual spectrum and portions of the infrared, see Figure 7. Consequently, we realized the use of visual magnitude as a selection criterion misrepresents the expected photon response of the detector.

To solve this issue, we developed a new catalog using the Hipparcos/Tycho-2 (HT2) astronomical database. In addition to position and proper motion information, HT2

contains magnitude data for each star in five spectral bands: U, B, V, R and I. We developed a weighted sum of the five channels that represents how bright a star will appear on the detector. The weighting factors were determined heuristically from orbital telemetry, with guidance from the Next Generation Spectral Library (NGSL) and the detector quantum efficiency (Figure 7).

We selected the brightest stars based on our spectral weighting to populate our new catalog. These are the stars that appear brightest on the ST-16 detector. We then added, or removed stars from the catalog based on star density. This was accomplished by surveying the entire celestial sphere with a large number of possible bore-sight vectors. For each orientation we ensure that there are ten cataloged stars in view. As a consequence of this approach, moderately bright stars in a high density area might not be included, while dim stars in a low density area would be included.

Handling of Difficult Stars

Despite having an improved star catalog, particular stars, while sufficiently bright to be detected, were still not matched. They were being rejected due to errors between their actual and expected star positions. After some more analysis, we traced these errors to the effects of proper motion and the presence of binary star systems.

Following this discovery, we first corrected our star catalog for proper motion to the current year. Secondly, we corrected the positions of star clusters by applying one of the following methods:

- **Apparent Binaries.** Two distinct stars appear joined due to the optics of the star tracker and thus prevent a quality estimate of the component star positions. Since the centroid of the conjoined star cannot be effectively estimated, the apparent binaries are not included in the star catalog. During the matching process, their detection would be treated as an outlier.
- **Equally bright binaries.** A binary star system is composed of two or more stars that are less than 2 magnitudes difference between the stars of the group. Since the stars are sufficiently close, the pair is treated as a single entry in the star catalog. The position of the binary system is approximated to be the magnitude weighted average of the component star positions.
- **Differently bright binaries.** A single star in the group is significantly brighter, more than 2 magnitudes difference, to the next star in the group. The

binary group is again treated as a single entry in the star catalog, with the position of the binary system taken as the position of the brightest star in the group.

CONCLUSIONS

In the preceding sections we have outlined the improvements to sensor processing on the ST-16 that were necessary to bring the sensor performance back up to their intended specifications. These changes include improvements in the logic for star detection, star measurement, rate estimation, and catalog management. Together the algorithmic improvements yielded higher availability, better accuracy, and a much-lower bad match frequency. Although new launches may require a short qualification period to tune calibration and operating parameters, we expect that the core software is stable.

The root causes of some of our post-launch difficulties rest with how we allocated qualification resources between simulation, laboratory, and night-sky tests. Our general approach sought to resolve most of the difficult performance questions in laboratory or simulation studies whereas the night-sky tests provided qualitative confirmation that 'things were working correctly'. For example, methodical laboratory tests were used to help set ideal detection thresholds; night sky tests were used to verify that the sensors could detect stars with the chosen settings but the whole parameter optimization was not actually validated. Overall, the night-sky testing conditions were almost too benign and many anomalies in the test telemetry were blamed on test conditions (e.g, clouds, vibration, etc.) without sufficient analysis. Our testing approach was not unreasonable, but would have benefited from more scrutiny for confirmation bias.

In addition to the many performance issues that have been fixed, there are a number of improvements under development that lay outside the scope of the flight qualification campaign. First we foresee considerable utility in modifying the camera recalibration code so that it can run online, on the sensor. This would drastically reduce the need for our involvement in qualification activities after each sensor launch. Second, investigations into the rolling shutter compensation underlined the criticality of the rate estimates. Recursive filtering could be used to provide better estimates of angular velocity beyond our simple finite difference implementation. Finally, the chromatic aberrations in the ST-16 lens affect the system level performance more than we initially estimated. A follow-on to the ST-16, the ST-16RT, uses a custom lens design with better chromatic performance.

Overall we are pleased with the improvements we were able to make in the ST-16. Good communication and operational integration between all three teams provided the support necessary to successfully diagnose and remedy many of the performance problems. The ST-16 is a successful sensor and we foresee few problems with future launches.

References

1. Dzamba T., Enright J., McVittie G., Sinclair D., Grant C., "Commissioning the S3S Nanosatellite Star Tracker," 61st International Astronautical Congress, vol. 6, pp. 4635–4649, Prague, Czech Republic, 2010.
2. Enright J., Sinclair D., Grant C.C., McVittie G., Dzamba T., "Towards Star Tracker Only Attitude Estimation," 24th Annual AIAA/USU Conference on Small Satellites, vol. SSC10-X-3, 2010.
3. Dzamba T., Enright J., "Optical Trades for Evolving a Small Arcsecond Star Tracker," 2013 IEEE Aerospace Conference, Big Sky, MT, 2013.
4. Liebe C.C., "Accuracy performance of star trackers - a tutorial," Aerospace and Electronic Systems, IEEE Transactions on, vol. 38, no. 2, pp. 587–599, 2002.
5. Enright J., Sinclair D., Dzamba T., "The Things You Can't Ignore: Evolving a Sub-Arcsecond Star Tracker," Proc. of the 26th AIAA/USU Conference on Small, vol. SSC12-X-7, Logan, Utah, 2012.
6. "Monochrome Addendum, Image Sensors: MT9P031I12STM. Aptina Imaging." , 2011, http://www.apgina.com/products/image_sensors/mt9p031i12stm.