

COTS Detectors for Nanosatellite Star Trackers: A Case Study

John Enright
Ryerson University
350 Victoria St. Toronto, Canada; 416-979-5000x4174
jenright@ryerson.ca
and
Doug Sinclair
Sinclair Interplanetary
268 Claremont St. Toronto, Canada; 647-286-3761
dns@sinclairinterplanetary.com
and
K. Christy Fernando
Ryerson University
350 Victoria St. Toronto, Canada; 416-979-5000x4973
cfernando@ryerson.ca

ABSTRACT

In recent years, many imaging sensors for small satellites have adopted CMOS active pixel sensors. The Aptina MT9P031 in particular, has seen many adopters. This paper presents lessons learned from our experience integrating this detector with the Sinclair Interplanetary S3S Star Tracker. We present practical advice for handling and assembly and discuss strategies for maximizing image quality over the sensor lifetime. In particular, we characterize the noise characteristics from several units after high-energy proton testing, and consider a number of approaches to maintaining performance as radiation damage accumulates.

INTRODUCTION

In recent years a number of manufacturers have brought small (< 0.5 kg) star trackers to market. These have enabled, for the first time, 10 arc-second-class pointing in spacecraft under 10 kg. Examination of marketing materials and private conversations with engineers has revealed that the majority of the designers have independently chosen to use the Aptina MT9P031 detector. The MT9P031 is a CMOS active-pixel sensor with five megapixel resolution. It is attractive because of its very low readout noise, but achieving consistent high quality images over the lifetime of the device presents some challenges. We have learned a lot about the operation of this device during the development of the S3S star tracker, and in this paper we share our key findings.

In the remainder of the paper we seek to answer the following questions:

- How do we design and build devices incorporating the MT9P031?
- How do we configure the detector to give the best image quality for star tracking?
- Is the MT9P031 suitable for the radiation environment in low earth orbit (LEO)?

- What adaptation may be necessary to maintain good image quality over the lifetime of the detector?

The remainder of this section gives some context for the study, explaining the history of our star tracker development and summarizing the effects of radiation on CMOS detectors. Subsequent sections seek answers to each of the above questions.

The S3S Star Tracker

Our experience with the MT9P031 stems from our development of the S3S star tracker (Figure 1). Detailed information about this sensor can be found in [1] and [2]. The S3S has been developed over the last three years through a collaboration between Sinclair Interplanetary, the Space Avionics and Instrumentation Laboratory at Ryerson University, and the University of Toronto Space Flight Laboratory. This 90 g sensor has very low requirements for mass, volume, and power and is designed to be compatible with a Cubesat chassis. The S3S provides direct quaternion output at 2 Hz and can obtain attitude estimates of body rates of up to 3 °/s. Internal processing is virtually stateless, and the main sensor processor is rebooted each imaging cycle, reducing the likelihood of radiation induced memory corruption.

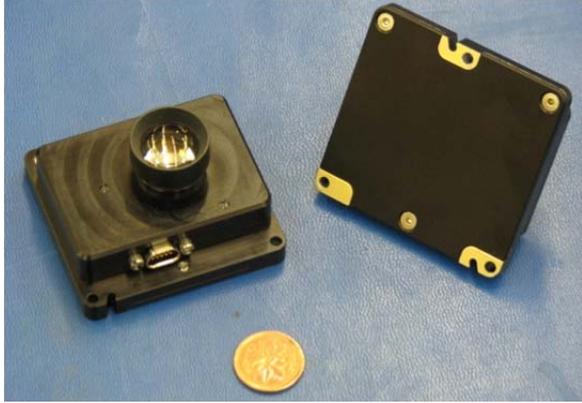


Figure 1: S3S Star Tracker.

We selected the MT9P031 for the S3S for several reasons. As an integrated active-pixel sensor, the detector requires relatively little external support electronics – a feature instrumental in maintaining a compact design. From a star-detection perspective, this detector has very low read noise (2.6 e⁻, RMS) and excellent sensitivity. The designed S3S detection threshold is set at a visual magnitude of 5.75, but in field-testing, we have successfully resolved magnitude 6.2 stars.

Radiation and CMOS Detectors

Most small satellite missions are bound for low Earth orbits – typically at altitudes between 400-1000 km. The dominant type of ionizing particle in this region is high energy protons. Most have energies less than 100 MeV [3]. The flux of the protons is also dependent on the solar cycle, increasing during the solar minimum [4]. Dose also rates vary with shielding and inclination.

Proton interactions with semiconductor devices cause displacement damage in the junctions [5]. Atoms of the silicon lattice structure that are originally excited will recoil off the structure causing additional permanent damage. Due to the displacement the electrical properties of the circuit are changed, increasing the dark current, since the defect has energy within the silicon bandgap. From the radiation testing discussed later in this paper, we can see the effect of the damage on the irradiated sensors as it did have a significant change in the operational characteristics of the detector.

Shielding can provide an effective barrier in some situations, preventing damage caused by electrons and low energy protons. Shielding becomes less effective against high energy particles (e.g., cosmic rays and energetic protons), because as the energy transfer will create secondary particles from interactions with the shielding material. This can cause the rate of SEEs to increase [6].

WORKING WITH THE MT9P031

The detector is a fully integrated active-pixel sensor. It requires far less external support than a more traditional CCD, and this makes it suitable for a miniature star tracker. External connections to the detector are straightforward: the inputs are +1.8 V and +2.8 VDC power, and a reference clock; the outputs are a 12-bit digital bus, an output clock, and line and frame valid signals. An I2C bus is used to set the internal registers.

Even with these simple interfaces, we have found that there is some subtlety in working with this component. The following sections give some tips.

Documentation

A reasonably complete datasheet for the detector is available online [7]. All other documentation requires an NDA with Aptina¹. A prospective user is advised to sign the NDA as soon as possible, and then to request the full datasheet and all available application notes. Earlier revisions of the datasheet contain useful information that has been removed from more recent versions.

In writing this paper we have been careful to abide by the terms of the NDA. In particular, internal register names and addresses are not included here. All discussions we include reference only material available on the public (NDA-free) datasheet.

Soldering

The detector uses a 48-iLCC package. Its small footprint is useful in miniature applications, but care must be taken when designing the board assembly process. Factors to consider are moisture sensitivity, gold removal, soldering, and X-ray inspection.

The package has a moisture sensitivity level (MSL) of 4 per J-STD-020D. Outside of protective packaging its floor life is only 72 hours. After this time it will have absorbed sufficient humidity from the atmosphere that it is in danger of “popcorn” failure during reflow soldering. It will be helpful to have access to a drying oven, a supply of desiccant, moisture barrier bags, and a heat sealer.

The external lands on the component have a 0.5 μm (20 micro-inches) gold finish. This must be removed before soldering to remove the risk of joint

¹ In fact, at the start of the project, even the basic data sheet required the NDA. This policy has apparently changed.

embrittlement from gold intermetallics. We do this by dipping the part in rosin flux and then tinning it in a solder pot filled with Sn60Pb40 alloy at 260 °C. This removes the gold, but leaves uneven bumps of solder on each land. We then hold the part upside down in a vacuum chuck so that a technician can hold a soldering iron in one hand and copper braid in the other to resurface the lands flat.

The component is solderable on the bottom surface only. There is no way to make a joint with a soldering iron – reflow equipment is necessary. We use a Zephyrtronix ZT-7 hot air station with a 10x10 mm nozzle. Various methods of applying solder have been tried. The best results come from abandoning solder paste. Instead, a technician “bumps” the PCB, carefully applying bumps of solder to the lands using a fine tipped soldering iron. After washing the PCB, liquid flux is applied to the footprint, the detector is placed on top, and the reflow operation is performed.

Once the part is soldered it can only be inspected by X-ray, since all of the joints are hidden. If any joint is bad the entire part must be removed. We have had no success reworking joints in situ.

Figure 2 was obtained from a realtime X-ray microscope operating at 80 kV. Within the red ellipse is an open-circuit defect. In the yellow ellipse is a joint with only partial wetting. The circular voids seen in many joints are not defects.

Power Supplies

The internal digital supply (VDD) must be connected to +1.8 V. The internal analog supply (VAA), the pixel supply (VAA_PIX) and the PLL supply (VDD_PIX) must be connected to +2.8 V. The designer can choose whether to connect the IO supply (VDD_IO) to either +1.8 V or +2.8 V. We have always used +1.8 V since it interconnects cleanly with our processor.

We have found that if the +2.8 V supply is not carefully designed, there will be additional noise in the image. The current consumed by the detector is modulated at the pixel, line and frame rates. It is important that the supply source impedance is low across this entire frequency band. Best noise performance is achieved with less than 10 mVp-p ripple.

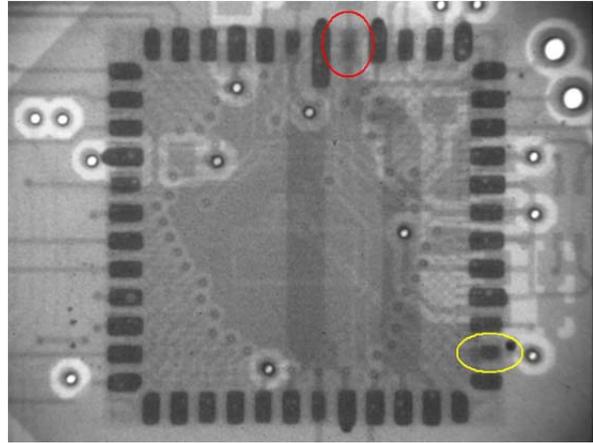


Figure 2: X-ray Showing Solder Defects.

High frequency (pixel rate) impedance can be reduced by decoupling capacitance. We use 0.1 μF at each pin, and see additional benefit from adding 1.0 μF directly under the chip. Medium frequency (line rate) impedance cannot be achieved with practical capacitors, and instead depends on the feedback control performance of the linear or DC/DC converter used to supply +2.8 V. Tuning of compensation networks is worthwhile. Avoid “burst mode” operation in switching supplies.

Power sequencing order is important. The 1.8 V supply must turn on before the 2.8 V supply for correct operation. If the 2.8 V supply is turned on first, we have found that the detector will sink a very large current.

PLL and FIFO

The detector contains a FIFO that permits the parallel output bus to be run at a rate slower than the internal pixel clock, with digital output continuing into the internal horizontal blanking interval. This feature is very useful, as it allows the pixel clock to be maximized (and thus rolling-shutter distortion minimized) in a system with limited digital I/O rate. Sadly this feature is not documented in the Aptina datasheets. By experimentation we have learned how to use it properly.

The phase-lock loop (PLL) must first be enabled. The PLL sets the internal pixel clock as follows:

$$f_{pixclk} = \frac{f_{extclk} \cdot M}{N \cdot P}$$

where M, N and P are multiplier and divider registers that can be set using the I2C bus. When the FIFO is enabled, the digital output clock is given by:

$$f_{fifoclk} = \frac{f_{extclk} \cdot M}{4 \cdot N}$$

where M and N are the same registers referenced above. In essence, the P divider for the FIFO is hard-wired to a value of 4.

Table 1: PLL Settings

f_EXTCLK	16.000 MHz
N	2
M	32
P	3
f_PIXCLK	85.333 MHz
f_FIFOCCLK	64.000 MHz

We use PLL settings shown in Table 1. The output clock is 64 MHz. Even though the output is only rated for 48 MHz with a +1.8 V IO supply we have seen no negative effects from overclocking the device. The noise performance is unaffected.

IMAGE QUALITY AND STAR DETECTION

The S3S sensor is a star tracker, so the particular detector configuration that we use is chosen to emphasize our star tracking performance. In order of importance, our priorities are:

- Maximal sensitivity to light from faint stars. Higher thresholds on visual magnitude (i.e., the ability to reliably detect dim stars), gives better attitude fixes and increases availability
- Low temporal noise. Noise increases centroid error; this degrades the accuracy of our attitude estimates, and increases matching difficulty.
- Locally uniform response. Stars usually distribute their light over several adjacent pixels, so accurate centroid measurements depend on neighbouring pixels behaving similarly to incoming illumination.
- High dynamic range. Once the image of a star saturates, our ability to post-process the star image (e.g., to estimate shape, magnitude, etc.), is severely limited. Maximizing the number of non-saturated stars is a desirable characteristic, but this is comparatively less important than other image quality measures.

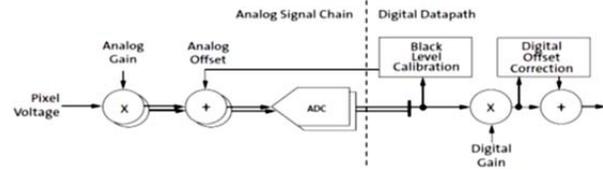


Figure 3: Signal processing path in the MT9P031 [7].

Other performance metrics such as linearity or photometric accuracy may be important to other detector applications, but only weakly impact star tracking performance.

Configuring the detector appropriately to meet these priorities requires an understanding of the internal processing on the MT9P031 (Figure 3). Pixel output voltages are first multiplied by an analog gain. This value is added to an analog offset before reaching the ADC. A digital gain and digital offset is applied after the ADC. Proper selection of the analog offset will guarantee that even small signals lie in the active range of the ADC.

The monochrome MT9P031 sensor that we are using is virtually identical to its colour counterpart. The colour version of this part gets its colour sensitivity from a colour filter array (CFA) (small colour filters in front of each pixel). Even though the spectral response of each pixel is identical, internal processing chains are separated by colour, according to the 2x2 pixel cells of the CFA (Green-1, Red, Blue, Green-2). Gains and offsets are specified separately for each colour. Pixels of each colour are further subdivided into one of two slightly different signal paths. Thus, there are eight distinct colour channels in the device, each with its slightly different response to pixel voltages. Gain and offset adjustments can be made to each colour, but there are no documented adjustments that can be made between the two same-colour channels.

Detector response is dominated by the following factors:

- Illumination and shot noise. These are a consequence of the imaging process.
- Read Noise (from ADC), minimized by operating at high analog gain.
- Quantization noise (from ADC), also minimized with high analog gain.
- Dark Noise. Thermally induced current fluctuations. Increases with temperature and is indistinguishable from photocurrent.

- **Black Level.** Combination of biasing, constant thermal current and other effects. May include components that vary by colour, row, and even individual pixel.

High sensitivity, and low read and quantization noise are managed by the selection of high analog gains. This maximizes the conversion from photons to ADC output units (ADUs). Dark noise can be minimized by keeping the detector cool. These are fairly simple settings. We have found that managing the black level is the trickier part of maintaining good image quality.

Hardware Black Level Correction

If the analog offsets are set to low values, the ADC output will be zero when the detector is not illuminated. However, the detector is only capable of fairly coarse offset adjustments, so excessively low values will cause clipping at the low end of the ADC's range, and the signal from faint stars may be lost. If the offsets are too high, then the response from un-illuminated pixels can become a significant fraction of the full range ADC output. This reduces dynamic range, and does not contribute useful information to star centroid calculations. Furthermore, variation in black levels between pixels can appear as spatial noise, and reduce centroid accuracy.

The detector contains regions of dark rows and columns around the edges of the active array that do not respond to incident light. Examining the values from the dark pixels allows us to estimate the current black level, and even correct the analog offsets. The detector contains logic that can automatically perform these functions (see Figure 3). The *black level calibration* logic uses some of the dark pixels to estimate appropriate values for the analog offsets in order to keep the black levels in a target range. Separate estimates are generated for the four primary colour channels. The *digital offset correction* is a secondary adjustment, calculated for each row and colour, to remove any spatial variations between rows. In this paper, we shall use *black level correction* (BLC), as a general term for any such techniques.

A typical terrestrial user would employ the hardware BLC (HW-BLC) features to ensure uniform dark values across all rows, columns and colour channels. However, during development of the S3S, we were concerned that the HW-BLC performance would be poor in a radiation environment. Single-event proton hits in the dark pixels could result in incorrect dark values. Artificially high dark estimates could, in turn, lead to low values of the analog offsets, clipping the response from dim stars. Some hot pixels can be tolerated in the pixels sampled by the black level calibration, but the row-by-row

corrections provided by the digital offset correction logic samples significantly fewer dark pixels. Hot pixels here would have a much greater impact on this secondary correction. Consequently, when testing the HW-BLC we use only the black level calibration logic and disable the digital offset corrections.

Software Black Level Correction

Concerned about the limitations of HW-BLC, a spacecraft user could disable the HW-BLC and replicate some of its functionality in software (SW-BLC). Here we outline our approach to SW-BLC; later sections compare the performance of these over sensor lifetime.

Using a map of any hot, dark pixels, we can estimate the current black level for each of the colour channels. These *dark offsets* are calculated by computing an average of all the dark pixels that share the same colour channel.

Although the HW-BLC can adjust the analog offsets within a frame, we do not have this capability in the SW-BLC. Instead, we calibrate appropriate values for the analog offsets during the sensor manufacturing process and use these over the life of the sensor. Although the SW-BLC analog offsets do not track black level variations, the dark offsets, will do so.

To calibrate the analog offsets, we measure the mean of the per-pixel temporal noise on each colour channel. This is the contribution from the thermal dark noise. As we reduce the analog offset settings, the temporal noise will suddenly drop to zero. We set the analog offsets just high enough to avoid this clipping effect. Figure 4 is an example of such a calibration. Because two channels share each offset, we must set the offset high enough to prevent both from clipping. This strategy differs from the heuristic used in the HW-BLC, which adjusts the dark level to lie within a target range.

Regardless of whether we use the HW-BLC or the SW-BLC, the dark offsets are always computed before processing the sky image. In the former case, the offsets should be approximately equal to the target dark level used by the hardware calibration logic. In the latter, the offsets estimate the current dark level including the ADC bias. These offset values are subtracted from the active pixel readings when attempting to detect and localize stars in the image.

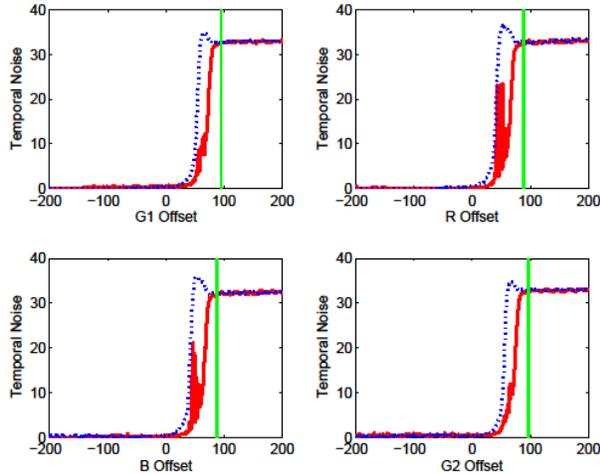


Figure 4: Example of analog offset calibration.

TESTING RESULTS

During development of the S3S, we had some concerns about the general suitability of the MT9P031 detector for orbital applications – it is a commercial part with no deliberate radiation hardening. Part of the qualification process for the star tracker included radiation testing at the TRIUMF proton-beam facility. This section describes our characterization of the detectors after their radiation exposure. We examine both the bulk damage to the detector and comment on the expected effects on star tracking.

Radiation Testing Overview

Two flight-grade S3S electronics assemblies were tested at the TRIUMF accelerator facility in Vancouver, Canada, in April 2010. The facility is equipped with a 105 MeV proton beam, and provides a good analog to the particle environment encountered in LEO. One unit was tested to a total of 9.4 krad and the other to 5 krad. The larger dose was judged to be a conservative estimate of EOL dose for most small satellite missions. A third unit was retained as a control, and was not irradiated.

After the irradiation process, the units were annealed at 125 °C for 48 hours. The annealing process was conducted in order to return functionality to the units as well as help mitigate detector damage that may have occurred. On-orbit this annealing may occur continuously (at moderate temperatures), but some high-temperature reconditioning may be necessary for longer missions. The star tracker can be made a little warmer by configuring the processors for maximum power draw and thus adding electric heat. Peak capability heating is estimated to be 2-3 W, but this must be verified by experiment.

In prior work, Becker, et al [8], also tested the radiation tolerance of the MT9P031. Their tests consisted of a 50 MeV proton beam at dose levels of 500 rad to 5 krad TID. Their findings show clear differences between the pre and post irradiation units and are generally similar to our results. They also report increases in the dark noise as the dosage level increased a rise in the number of hot pixels. Becker proposes using image subtraction to remove the hot pixels, but it is unclear whether this strategy is effective in units (like the S3S), without mechanical shutters. Their tests are all conducted at ambient temperature.

Radiation Testing Results

Performance of the detector degrades with both increasing dose and temperature. Figure 5 shows a histogram of the zero-illumination detector response as the dose increase (these plots only show a single channel). Figure 6 shows the response of the 9 krad unit across temperature. Bulk effects include both an increase in dark noise (peaks broaden), and black level (peaks shift to the right). Although the response at elevated temperature and dose displays some noticeable anomalies, detector behaviour remains quite reasonable at lower temperatures, even at the end of life (EOL).

The distribution curve for the 9 krad unit at 30 °C is very similar to the 0 krad unit at 60 °C, barring a shift to the right, and a consequent, minor reduction in dynamic range. Even the performance at 45 °C may be acceptable. For star tracking, the spread of the distribution – increasing noise and decreasing sensitivity – is more important than the loss in dynamic range.

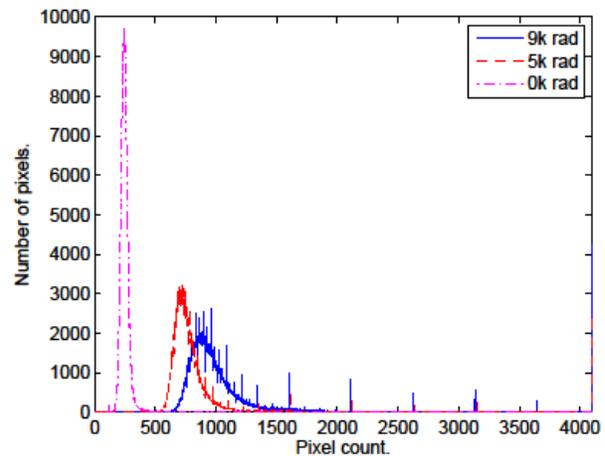


Figure 5: Effect of TID (60°C, SW-BLC).

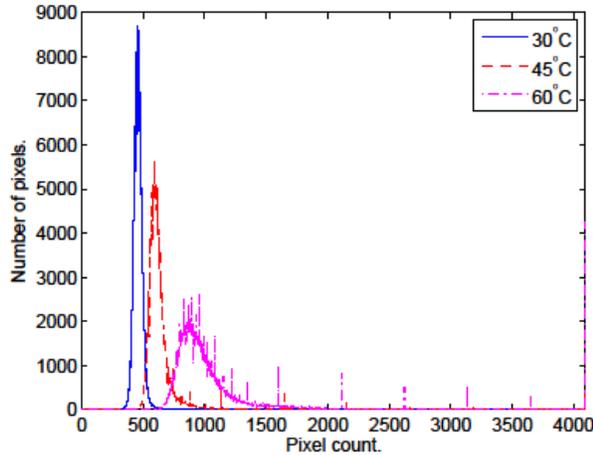


Figure 6: Effect of Increasing Temperature (9 krad, SW-BLC).

Taking a closer look at the 9 krad unit’s behaviour at high temperature gives insight into secondary aging effects that may be occurring (Figure 7). Inspection of the histograms shows a scatter of individual brightness values that are favoured over their neighbours. There appears to be a pattern in their spacing, but they occur in different places for different colour channels.

Table 2 shows the last 10 peaks in all eight colour channels, for the 9 krad unit at 60 °C. When represented in hexadecimal it is clear that there is a pattern to the last eight bits. Colours 2, 3, 6 and 7 favour hexadecimal patterns xFFh. Colour 0 favours x3Eh, colour 1 favours x46h, colour 4 favours x3Fh and colour 5 favours x2Ah.

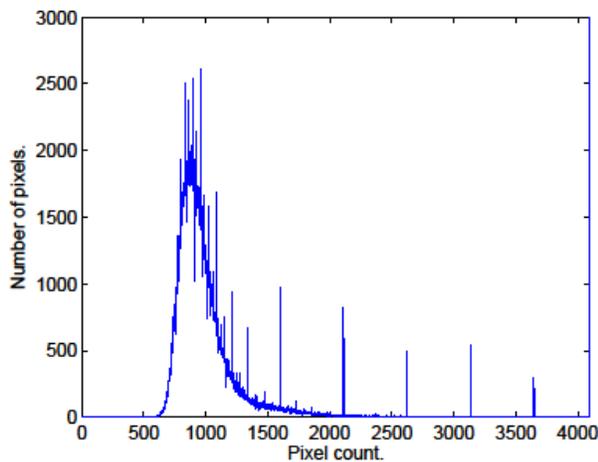


Figure 7: Dark noise distribution near end of life (60°C, 9krad, SW-BLC).

Table 2: Favoured ADC Readings

Ch. 0	Ch. 1	Ch. 2	Ch. 3	Ch. 4	Ch. 5	Ch. 6	Ch. 7
41Eh	3FEh	47Fh	45Fh	43Fh	47Ah	4BFh	4FFh
43Eh	406h	49Fh	47Fh	47Fh	4AAh	4FFh	53Fh
47Eh	446h	4BFh	4BFh	4BFh	4EAh	53Fh	57Fh
4BEh	486h	4FFh	4FFh	4FFh	52Ah	57Fh	5FFh
53Eh	4C6h	57Fh	57Fh	53Fh	5AAh	5FFh	67Fh
63Eh	546h	5FFh	5FFh	63Fh	62Ah	6FFh	6FFh
83Eh	646h	7FFh	7FFh	83Fh	82Ah	7FFh	7FFh
A3Eh	846h	9FFh	9FFh	A3Fh	A2Ah	9FFh	9FFh
C3Eh	A46h	BFFh	BFFh	C3Fh	C2Ah	BFFh	BFFh
E3Eh	C46h	DFh	DFh	E3Fh	E2Ah	DFh	DFh

Patterns of xFFh and x3Fh are strongly suggestive of ADC damage to a successive-approximation type device, as a long sequence of least-significant bits are set to ‘1’. The other favoured codes are less typical of ADC damage, but there is no other credible explanation for the effect.

We do not know how many ADCs are built into the detector. There may be one that samples each pixel in turn, or there may be a number (most likely four) that are pipelined. In any event, it appears that cumulative radiation dose degrades the linearity of the circuits making certain codes more likely than their neighbours. This effect is interesting, but does not appear to seriously interfere with a star tracking function.

In addition to testing the SW-BLC, we can also examine the effectiveness of the HW-BLC near the sensor’s EOL (Figure 8). One immediate conclusion from this figure is the failure of the HW-BLC at high temperature. A large fraction of the curve ends up clipping to zero. The target black level (300 counts in this test), is maintained well at lower temperatures, but at 60 °C, the clipped pixels dominate all the others. Adjusting the target black level may be enough to mitigate this effect, but at time of writing, it is unclear whether this is driven only by the broadening of the distribution, or whether hot dark pixels lead to unusually high dark level estimates and subsequent overcorrection in the analog offsets.

In a test at 0 °C (full results not shown), the SW-BLC showed clipping at the low end of the ADC range, while the HW-BLC performed as expected. This suggests that some adjustment to the BLC parameters will be necessary, regardless of the choice between HW and SW implementations.

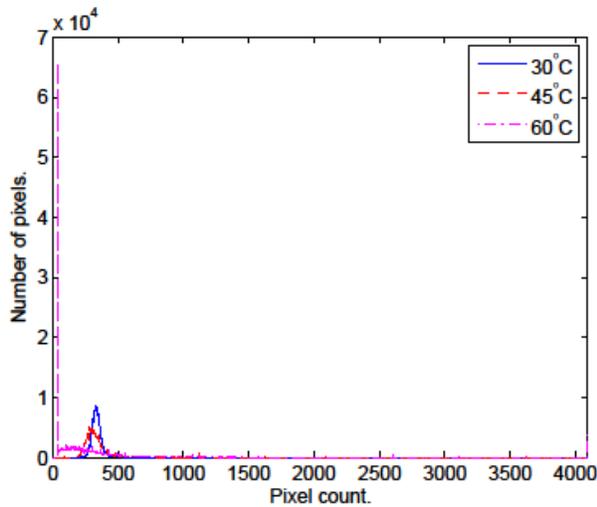


Figure 8: Effect of Temperature (HW-BLC, 9krad).

We can gain insight into the effect of radiation on star detection performance, by looking at cumulative distribution plots (Figure 9-Figure 12). In these figures, we plot the fraction of the image pixels that exceed the measured dark offset by a specified amount. For example the curves at 500 counts on the x-axis represent the fraction of pixels that read 500 counts, or more, above the dark offset. Our star detection routines begin by correcting the measured intensity, by subtracting the measured offset. Consequently, large pixel populations, far from offset value (i.e., large horizontal separation between curves) increase the likelihood of losing faint stars in amidst thermal noise, or falsely identifying noise as stars. These plots also show reductions in dynamic range in the positions of the right endpoints of the curves.

Unsurprisingly, minimizing detector temperature is critical to preserving good imaging characteristics, particularly at EOL (Figure 9). Although some sensitivity in the 9 krad unit is lost at 45 °C, acceptable detection levels are still likely. If we can maintain detector temperatures near 30 °C (Figure 10), we see very little difference in the distributions within a few hundred counts of the measured offset. The longer distribution tails above this value may cause false detections, but are unlikely to significantly impair the detection of faint stars.

At higher temperatures, the potential for detection problems increases with total dose (Figure 11). At 60 °C, the spread in the distribution may make faint detections problematic, even at 5 krad. This figure also suggests that most of the black level elevation (and hence the reduction in dynamic range) occurs from 0 krad to 5 krad.

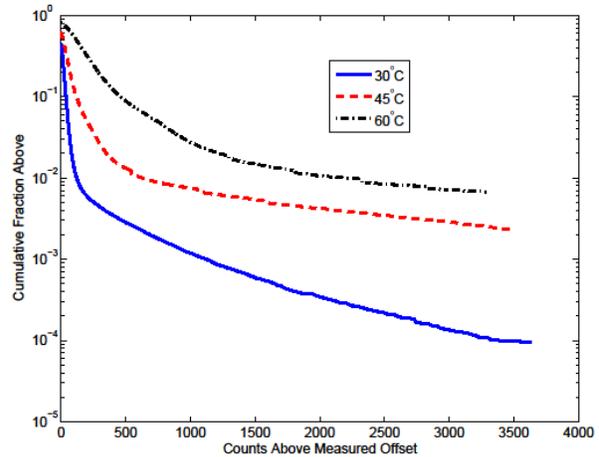


Figure 9: Effect of Temperature (9 krad, SW-BLC).

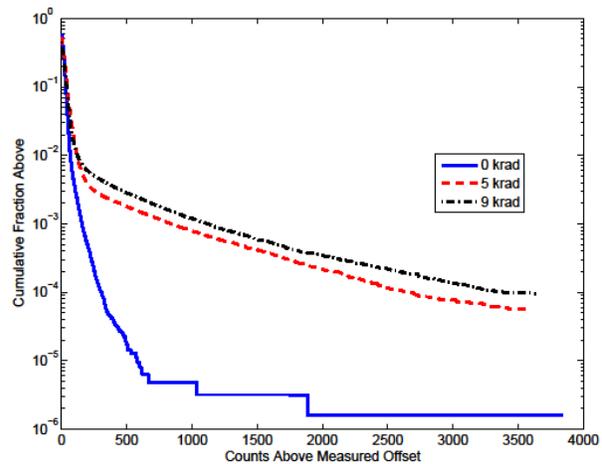


Figure 10: Effect of TID (30 °C, SW-BLC).

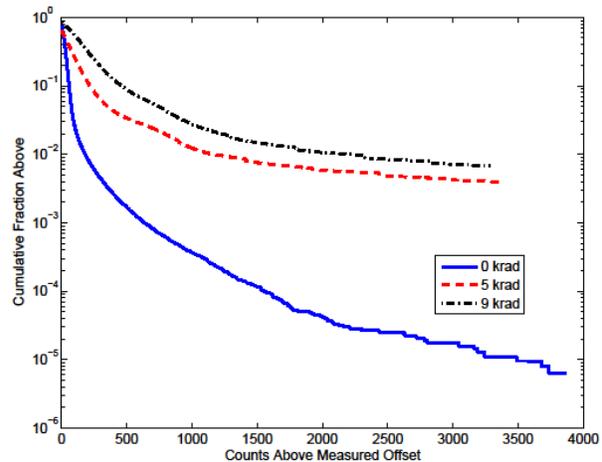


Figure 11: Effect of TID (60 °C, SW-BLC).

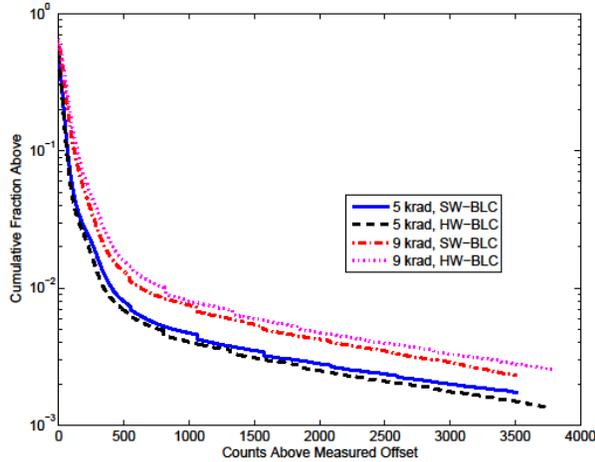


Figure 12: Comparison of HW-BLC and SW-BLC (45 °C).

We can also revisit the comparison between HW-BLC and SW-BLC (Figure 12). As suggested by the earlier histograms, if we avoid the clipping regime, the two schemes perform very similarly. Because the on-chip logic is able to adjust the analog offsets, the HW-BLC preserves more of the detector’s dynamic range.

In addition to increased dark level and dark noise, radiation damage also increases the number of hot pixels present in the detector. Hot pixels are pixels that are always saturated, even under zero-illumination. Table 3 summarizes the total number of hot pixels across the tests. We see that the numbers increase with temperature and even more dramatically by dose. These data suggests that damage from the radiation exposure creates a large number of pixels that become extremely sensitive to temperature. The table also shows that at elevated temperatures, the HW-BLC usually reduces the total count of dark pixels. A close look at the histogram populations reveals that this is not simply a shifting of the high-end distribution – neighbouring bins have only a handful of pixels. This phenomenon may be related to the ADC anomaly discussed above.

LONG-TERM HEALTH MAINTENANCE

As the detector ages and radiation damage accumulates, we can maintain good image quality by detecting and compensating for image defects as they occur. In this

section we compare several strategies for detecting the common types of flaws that we expect to develop in our CMOS detector as the total radiation dose increases. At time of writing, we have not completed our implementation of these algorithms. Instead we present a detailed description of our intended approach and motivation for this investigation.

We approach this problem by considering the types of damaging effects that we would expect over the life of the detector. At one extreme, we have sudden, readily apparent events such as latch-up. While the electronics are designed to mitigate or avoid serious damage, our general philosophy is that such events can be diagnosed quickly (e.g., by current draw or non-responsiveness), and cleared with a reset. This will either have solved the problem, or the sensor will have failed. In any case, we do not need to subsequently monitor the sensor’s health internally.

More interesting from a health-monitoring standpoint is damage that causes a gradual deterioration at the pixel level. These defects appear as outputs that are valid readings, but do not correspond to the intensity of incident light on the detector. Moreover, the combination of large pixel counts, relatively sparse illumination, and higher level processing can mask pixel level damage from immediate detection. As the density of damaged pixels increases, sensor performance will eventually suffer.

As borne out by our experiments, we expect pixel damage to fall into one of the following categories:

- Hot Pixels. This is the most common type of pixel damage. The outputs of these pixels remain saturated regardless of the input illumination.
- Warm Pixels. Similar to hot pixels, this type of damage causes abnormally high dark-level response, but the output is not saturated when the pixel is not illuminated.

Other types of detector damage may occur, but these phenomena become increasingly harder to identify conclusively.

Table 3: Hot Pixel Population

TID	30°C		45°C		60°C	
	SW BLC	HW-BLC	SW-BLC	HW-BLC	SW-BLC	HW-BLC
0 krad	4	4	13	17	52	90
5 krad	380	658	8399	6605	23566	19435
9 krad	609	1063	12791	13581	37095	20398

In the abstract, detecting these defects should not be a particularly difficult problem. The challenge comes from integrating long-term image quality monitoring into the existing S3S electronics and processing architecture: relatively slow external communications precludes readout of full-frame images; our conscious choice to adopt state-less processing makes long-term measurements incompatible with the standard operating concept, and the limited write-lifetime of the internal flash memory, prevents continuous updates of estimated defects.

Our solutions depend on establishing a ‘maintenance mode’ for the sensor. We assume that the host spacecraft can command this mode on occasion in order to update an internal map of faulty pixels. During normal processing the software uses this list and ignores, or interpolates the detector readout from the affected pixels. Some satellite motion would be required to induce intensity changes on the detector from image to image, and we assume the sensor FOV will be free of bright bodies such as the sun, earth, or moon.

We propose three types of processing for the maintenance mode, each with distinct operating concept. We are currently in the process of implementing these algorithms and evaluating their performance. A comparison of these three schemes is shown in Table 4.

Hot Pixel List.

A small amount of state (a few kb) is maintained from one frame to the next, by sending data to the supervisor processor and back to the functional processor. This state information is a list of potential hot or cold pixels that are currently being watched along with some historical information for each one. Pixels will be dropped from the list if they do not maintain the expected behaviour. Once we are confident that a given pixel is faulty, we can update a master list stored in flash. Only a small area of the array can be monitored at one time, but this algorithm requires few changes to the current processing cycle.

Per-Pixel Filtering

A small number of image-sized buffers are kept from cycle to cycle. These calculate running averages,

variance estimates, or other similar metrics for each pixel in the array. Hot pixels can be identified by looking for very low variance, and their locations can be stored in the onboard flash memory.

Motion-Correlated Processing

In this scheme, a number of full-frame image buffers are saved (either into RAM or flash). Periodically, a batch of buffers is processed to estimate image defects. We use angular rate information from the nominal sensor processing to correlate bright areas of the detector from one image to the next (this is akin to optical flow). Image features that do not move with the body motion are assumed to be flaws on the detector. This scheme requires a significant amount of processing, but because it makes better use of the knowledge of the satellite body motion, the algorithm may require less supervision in assessing the reasonableness of its estimates.

CONCLUSIONS

The MT9P031 is quite a capable detector. With proper handling and configuration it can work quite well in space applications. However, thorough testing and careful reading of application notes and datasheets are essential. It is even possible to overclock the detector without compromising performance. Manufacturing devices processes require a reasonable amount of care and successful soldering is probably impossible without reasonably sophisticated tools.

Radiation testing with the detector showed several damage mechanisms as the total dose increased. Some thermal management is crucial to maintaining good performance at EOL, but measures such as active cooling are probably unnecessary. Histogram distributions at 45 °C, suggest that star tracking is still possible, however, in the coming months we will integrate the electronics into sensor chassis and verify this conclusion with functional testing.

Hardware and software BLC algorithms perform quite similarly. This demonstrates that it is fairly simple to replicate the function of the HW-BLC if desired, but also that the SW-BLC is probably acceptable, too. Both schemes require some adjustment to maintain good performance with changing temperature and dose. Scheduling analog offset settings with temperature is

Table 4: Qualitative Comparison of Pixel Defect Detection Schemes

Processing Scheme	Hot Pixels	Warm Pixels	Computational Complexity	Expected Convergence Time
Pixel List	X		Low	Long
Per Pixel Filtering	X		Medium	Medium
Motion-Correlated Processing	X	X	High	Short

probably the most effective approach to managing the SW-BLC.

*Electronic Parts and Packaging (NEPP) Program
Office of Safety and Mission Assurance, 2008.*

Hot pixel populations become significant as the lifetime dose increases. These pixels are not always visible, at lower temperatures, but damage is evident in the large numbers that appear as the temperature rises. Mapping the locations of these pixels is an important part of sensor health management, and effective identification schemes will likely require some cooperation from the host spacecraft. This is an area of ongoing investigation.

Having spent a significant amount of time working with the MT9P031 we feel that we have learned some valuable lessons about this device. These findings should be useful to others who may be considering this part for space applications. For projects considering other detectors, our qualification and characterization studies can also provide helpful guidance.

REFERENCES

- [1] Dzamba and J. Enright, "Commissioning the S3S Nanosatellite Star Tracker," in *61st International Astronautical Congress*, Prague, Czech Republic, 2010, p. IAC-10.B4.4.10.
- [2] J. Enright, D. Sinclair, C. C. Grant, T. Dzamba, and G. McVittie, "Towards Star Tracker Only Attitude Estimation," in *Proc. of the 24th AIAA/USU Conference on Small Satellites*, Logan, Utah, August.
- [3] J. M. Lauenstein and J. L. Barth, "Radiation belt modeling for spacecraft design: model comparisons for common orbits," in *Radiation Effects Data Workshop, 2005. IEEE*, pp. 102-109.
- [4] S. Bourdarie and M. Xapsos, "The Near Earth Space Radiation Environment," in *IEEE Transactions on Nuclear Science Vol. 55, No. 4*, 2008.
- [5] M. Cohen and J. P. David, "Radiation effects on active pixel sensors," in *Radiation and Its Effects on Components and Systems, 1999. RADECS 99. 1999 Fifth European Conference on*, 1999, pp. 450-456.
- [6] K. B. Miller and T. O'Connor, "Radiation Effects on Stellar Reference Sensors," in *Advancement in Astronomical Science, Vol. 121*, 2005.
- [7] Aptina, Inc., *MT9P031 Public Datasheet Rev. E*. <http://www.aplina.com/assets/downloadDocument.do?id=606>, 2010.
- [8] H. N. Becker, M. D. Dolphin, D. O. Thorbourn, J. W. Alexander, and P. M. Salomon, "Commercial Sensor Survey Testing Progress Report," in *NASA*