Irradiation of the TC255 CCD by Fast Neutrons

Kevan S Hashemi and James R Bensinger

Brandeis University

11 September 1998

Abstract: The TC255P is a CCD manufactured by Texas Instruments. We propose to use it in the ATLAS end-cap muon alignment system. To find out if it could operate for ten years in the ATLAS radiation environment, we exposed four TC255Ps to fast neutrons. The doses delivered to the CCDs varied from one to three times the highest expected ten-year ATLAS dose in the end-cap. We observed that the dark current in the CCDs increased dramatically. But we observed no other forms of damage. We believe we understand the means by which neutron bombardment causes an increase in the dark current, and we have an empirical model that describes the progress of the damage with accumulated dose. This model predicts that, once we have made a number of simple changes to the image retrieval procedure, the TC255P will be able to operate for one hundred and thirty years in the ATLAS radiation environment.

Introduction

According to neutron background simulation TP36 [3], there are some parts of the ATLAS forward muon detector that will receive a total dose of 1.1×10^{13} n/cm^2 during the ten-year running time of the experiment. This dose is equivalent to 2.9×10^2 1-MeV n/cm^2.

We are concerned with the radiation-hardness of the TC255, a CCD (charge-coupled device) manufactured by Texas Instruments. We propose to use in the end-cap muon alignment system. It is made by the fabrication process that Texas Instruments and the Jet Propulsion Laboratory developed to make radiation-resistant CCDs for the Gallileo space probe. The CCDs in Gallileo proved resistant to over one hundred kilorads of ionizing radiation. This is far more than the expected ATLAS muon

detector dose of just a few kilorads. We have found no report, however, on the resistance of the CCDs to fast neutrons.

We are fortunate to have a pure source of fast neutrons an hour's drive Brandeis, at University of Massachusetts at Lowell. This produces fast neutrons by bombarding a lithium target with a 20-µA beam of 4 -MeV protons. The energy of the neutrons averages about 1 MeV, and there is only one gamma-ray for every twenty neutrons. We took four TC255s to Lowell and subjected them to between 2×10^{12} and 8×10^{12} 1-MeV n/cm^2.

Apparatus

We mounted four CCD Heads circuits within a few centimeters of the neutron source (see Figure 1). Each CCD Head carried its own TC255. The TC255 has 240 rows and 320 columns, making 76,800 pixels, each of which is 10-µm square. The CCD comes in an 0.4-inch wide, 8-pin plastic DIP package. It has an imaging area and a storage area. Images are captured in the imaging area and then transferred rapidly into the storage area. The storage area is covered by an aluminum light shield so that the image is not exposed to light while it is being transmitted to the CCD Driver. The CCD Driver digitizes the pixel intensities with an 8-bit analog-to-digital converter (ADC). The driver we used in this experiment had an offset of 45 ADC counts. Pixel saturation corresponded to an intensity of 245 counts, so the dynamic range of the pixels was approximately 200 counts. According to the TC255 data sheet, the pixels saturate when they contain approximately 60,000 electrons, so we conclude that one of our ADC counts represents approximately 300 electrons.

Against each CCD we pressed a RASNIK (red alignment system of NIKHEF) mask. An infra-red diode about ten centimeters away from each CCD illuminated each mask, throwing its shadow onto the CCD, and allowing us to record RASNIK patterns during the experiment. We could watch the image quality during the experiment.

Aside from the TC255, each CCD Head carried several resistors, diodes, capacitors, and an op-amp. Two CCD Heads used a radiation-resistant op-amp (Elantec EL2020CN) and two used a low-power op-amp (Elantec EL2044CN).
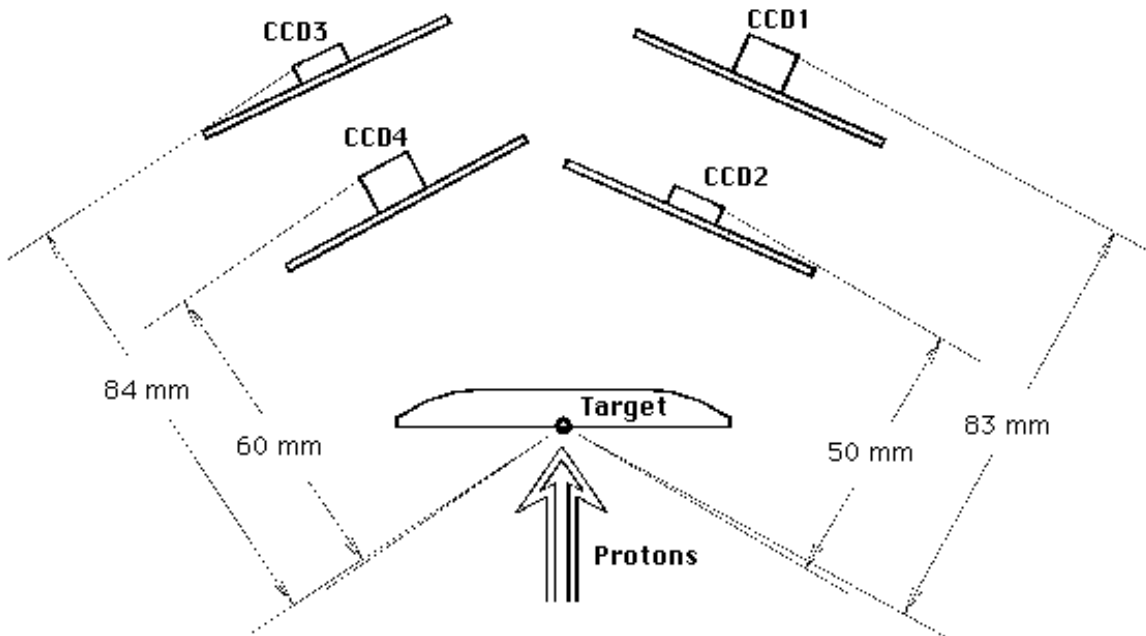
Figure 1: **Arrangement of CCDs in the radiation test. A beam of 4-MeV protons from a van der Graaf accelerator strikes a Li^7 target and produces a near-pure fluence of neutrons.**

We connected the CCD Heads to a nearby CCD Multiplexer [4], and connected the CCD Multiplexer to a CCD Driver [4] in the control room. The cable from the CCD Multiplexer to the CCD Driver was 30 m long. In the control room, we had our VME crate and Macintosh computer. We assigned a name to each CCD Head: 'Head One' to 'Head Four', and readied ourselves for the eight-hour exposure.

Results

Figure 2 shows four images taken from Head One during the course of the experiment. Cumulative damage is strikingly obvious. Our analysis program succeeds with the first three images, but fails on the last. By the time of the last image, Head One had received 3.0×10^{12} 1-MeV n/cm^2, or one worst-case 10-year ATLAS end-cap dose.

To determine whether components other than the CCD were damaged on the CCD Heads, we took Head Two, removed its irradiated CCD, and replaced it temporarily with a fresh CCD. We obtained pristine images with this circuit, no different from the images obtained with a fresh CCD in a fresh CCD Head.

The images show a whitening towards the top. As the damage progresses, the top lines saturate, and the saturation starts to proceed down the image. The final images from Head Two, which received 7.7×10^{12} 1-MeV n/xm^2, are white almost to the bottom. Aside from whitening at the top, the images show no other deterioration. In the parts of the image that are not yet saturated, contrast and sharpness appear undiminished. White specks appear at random in images captured during the irradiation, but they do not appear in images captured after the irradiation. We believe the specks were caused by gamma-rays. As mentioned above, the radiation contained one gamma-ray for every twenty neutrons.
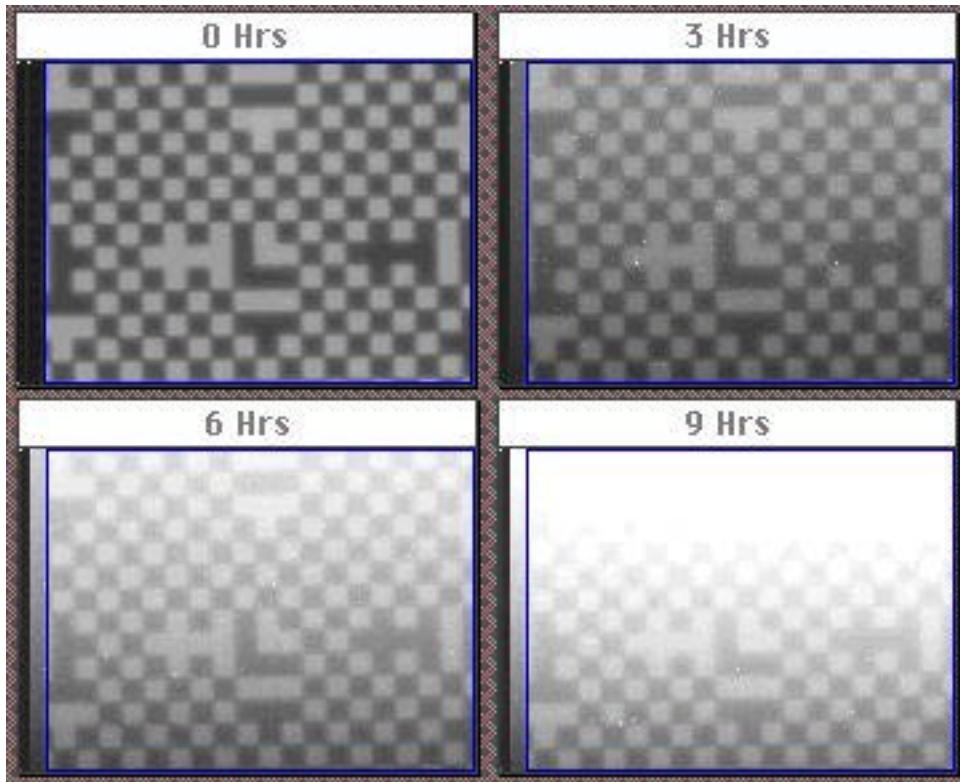

Figure 2: **Four images taken from Head One during the course of the experiment. The time at which each image was captured is given in hours since the start of the experiment. The vertical line on the left of each image is drawn on to mark the extent of the CCD imaging area. The pixels immediately to the left of the line are shielded from light, both in the imaging area and the storage area.**

The asymmetry in the degradation of these images matches an asymmetry in the time spent in the CCD by the image pixels. To capture an image, the driver clears the CCD imaging area (16 ms), exposes it to light (4 ms), moves the image into the CCD's light-shielded storage area

4

(17 ms), and transfers the image from the storage area into its own RAM (174 ms). The final operation takes so much longer than the others because its proceeds one pixel at a time and must traverse a 30-m cable. The other operations proceed one row at a time and are confined to the CCD. An asymmetry arises because the pixels in the bottom row of the image are transferred out of the storage area first. They spend only a few milliseconds waiting in the storage area, while the pixels in the top row wait almost two hundred milliseconds.

Every CCD suffers from 'dark current', which is charge leaking into the pixels from sources other than photon absorption. Dark current is present not only in the imaging area, but also in the storage area of a TC255. In an undamaged TC255, dark current fills the image pixels in four seconds at room temperature. But if the neutron bombardment were to increase the dark current to the point where pixels filled in a fraction of a second, we would see the top of our images getting white. The pixels in the top rows of the image would accumulate ten times as much dark current as those in the bottom rows.

The leftmost fringe of each image contains pixels that never see light. They are shielded both in the imaging area and in the storage area. Consequently, any charge accumulated in these pixels must be due to dark current, which operates in the storage area just as it does in the imaging area. The only difference between the two areas is that the frame store is covered with an aluminum light shield. The intensity within the leftmost fringe of the 6-hr image in Figure 2 varies linearly from 234 counts at the top down to 68 counts at the bottom. The CCD Driver's input offset is 45 counts. With the offset subtracted, the intensity varies from 189 counts down to 23 counts. The time for which the pixels accumulated dark current is 203 ms at the top of the image and 33 ms at the bottom. We conclude that dark current must have increased the intensity of each pixel by close to one count for every millisecond the pixel spent in the CCD. Since one count represents approximately 300 electrons and each pixel is 10 μm square, one count per millisecond is approximately 0.5 μA/cm^2. The TC255 data sheet specifies a dark current of 0.02 μA/cm^2 for an undamaged sensor. We could use 'electrons per ms' or 'μA/cm^2 ' as our measure of dark current, but we prefer 'counts per ms' because it is what we measure.

When the length of the cable between the CCD Driver and the CCD Head is less than a few meters, we can transfer an image from the CCD more rapidly. It takes 48 ms instead of 174 ms. Figure 3 shows two images taken from Head One soon after the irradiation. The first was captured slowly, the second quickly. The top half of the first image is saturated, but there are no saturated pixels in the second image.
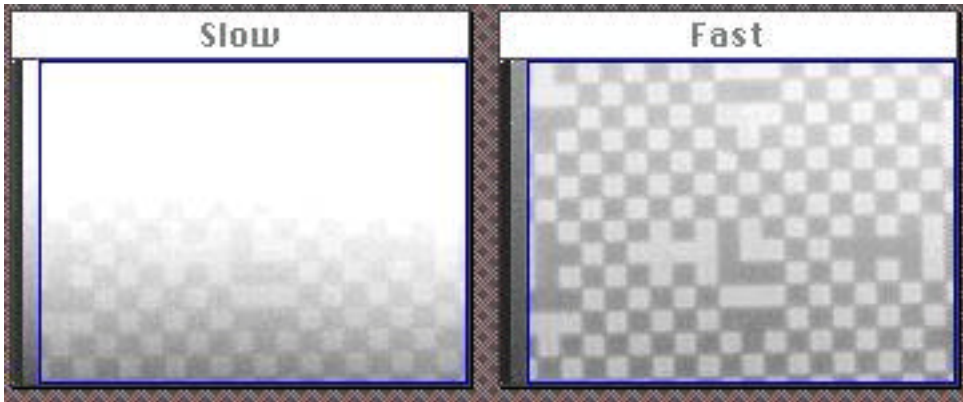
Figure 3: Images captured from Head One at the two different readout speeds. The images were captured within a week of the irradiation.

Figure 4 shows the dark current in our four CCDs as a function of time during the irradiation. The dark currents were calculated using images we captured and saved during the irradiation
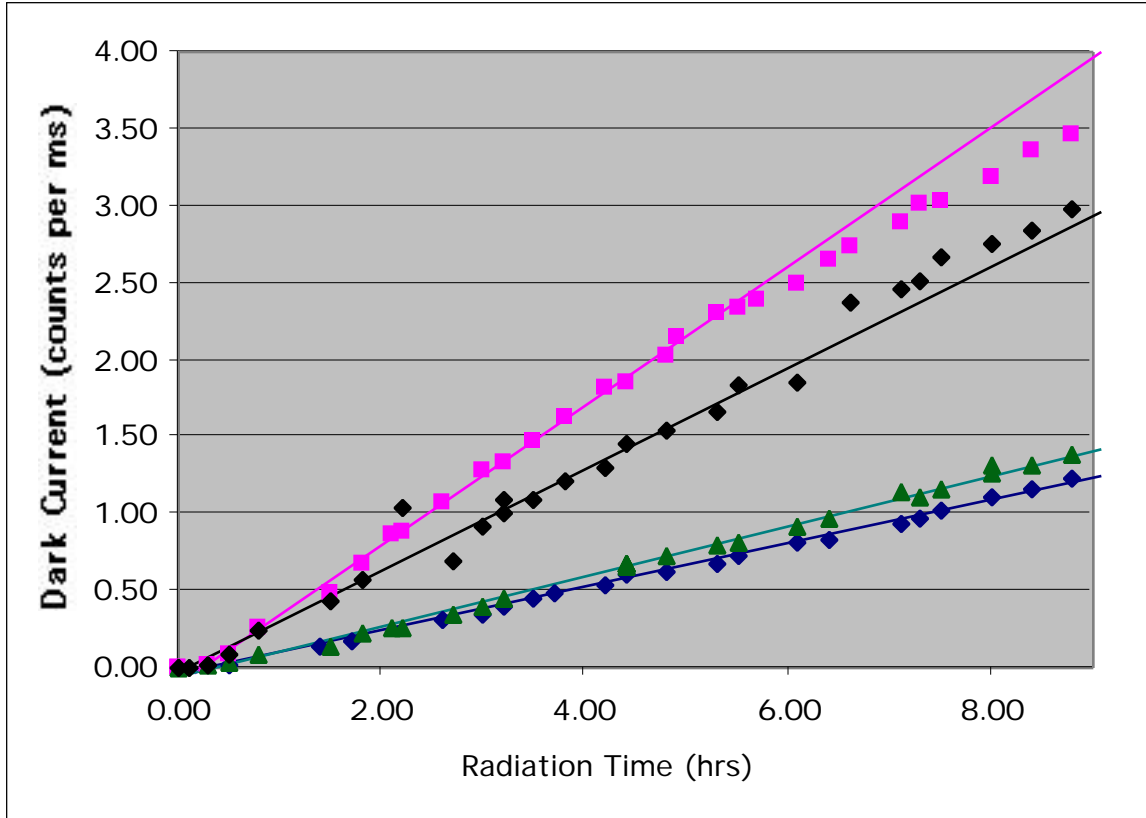
Figure 4: Dark Current (counts per ms) vs. Radiation Time (hours) for our four CCD Heads. One count per millisecond is approximately 300 electrons per millisecond flowing into each pixel, or 0.5 μA/cm².

Table 1 gives the final values of dark current for each CCD, and estimated values for the neutron dose.

Head	Total Dose (10^{12} 1-MeV n cm^{-2})	Dark Current (ADC counts ms^{-1})	Current/Dose (units as specified)
One	3.0	1.6	0.53
Two	7.7	4.6	0.60
Three	2.7	1.8	0.67
Four	5.2	3.9	0.75

Table 1: The total dose and final dark current for the four CCD Heads.

We suspect that our dose estimates are inaccurate, in particular because the dose for Head Three is smaller than for Head One, while the final dark current in Head Three is greater than in Head One. The estimates were calculated with the assumption that the proton beam was well-centered on a 25-mm diameter lithium target. If the beam were 2 mm

left of center, the ratio of dark current to estimated dose would be the same for all CCDs to within a few percent.

Neutron Damage in CCDs
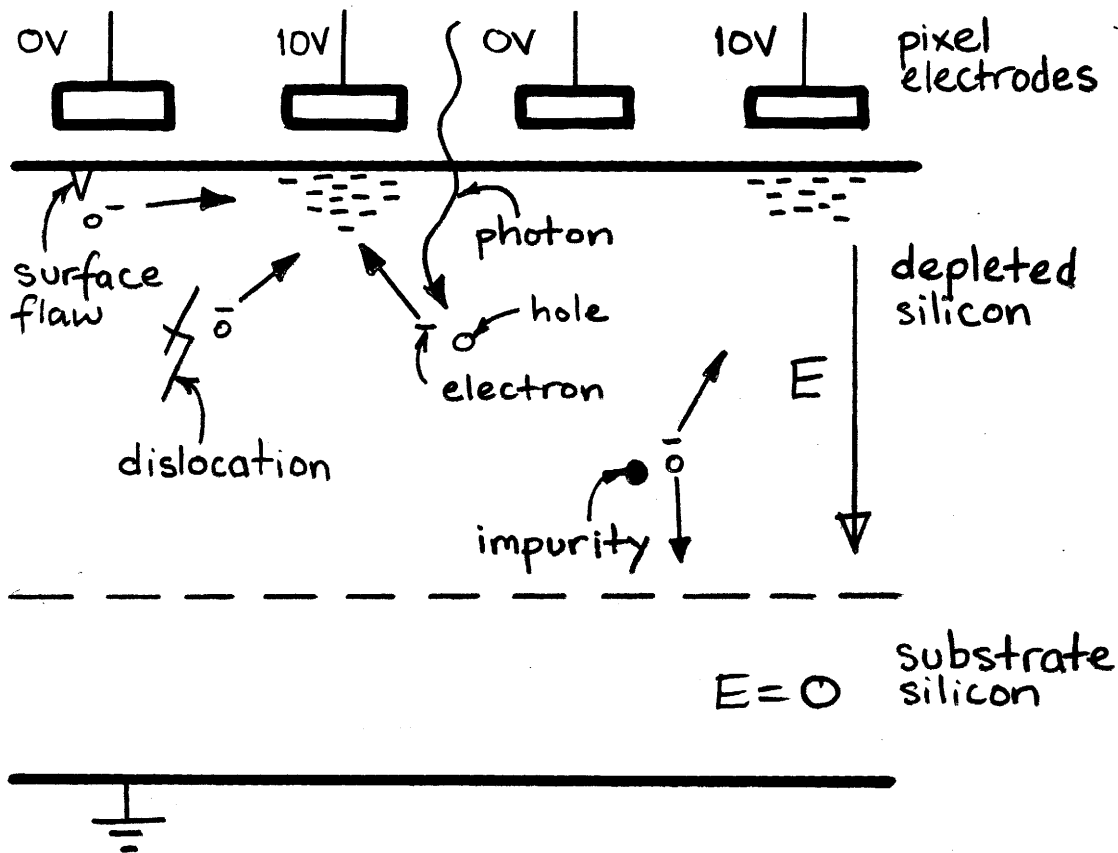
Figure 5 shows a cross section of a CCD such as the TC255.

Figure 5: A cross-section through a CCD chip, showing electrodes, depleted top layer, and conducting substrate. Electron-hole pairs are generated by photons, surface flaws, dislocations, and impurities. The electrons end up beneath the nearest positively charged electrode.

The top two or three microns of the CCD silicon are doped. The doped silicon becomes depleted. It contains an electrical field generated by its depletion charge and by the potential applied to the surface electrodes. This field drives electrons up to the positively charged electrodes, and moves holes down into the substrate. Dark current is made up of electrons from electron-hole pairs generated by thermal excitation at lattice imperfections. Neutrons of sufficient energy bombarding the lattice create new imperfections [2], and so raise the dark current.

The lattice damage caused by neutrons with kinetic energy less than 1 keV is most likely to be the displacement of individual silicon atoms into interstitial locations. Higher energy neutrons are able to cause progressively more severe damage.

Temperature Dependence of Dark Current

Figure 6 is a graph of dark current verses temperature for Head Two, which is the most damaged of the set of four. We pressed the CCD up against a thermistor and cooled it with freezer spray in a light-proof box. As the CCD and thermistor warmed up, we measured their temperature and dark current.
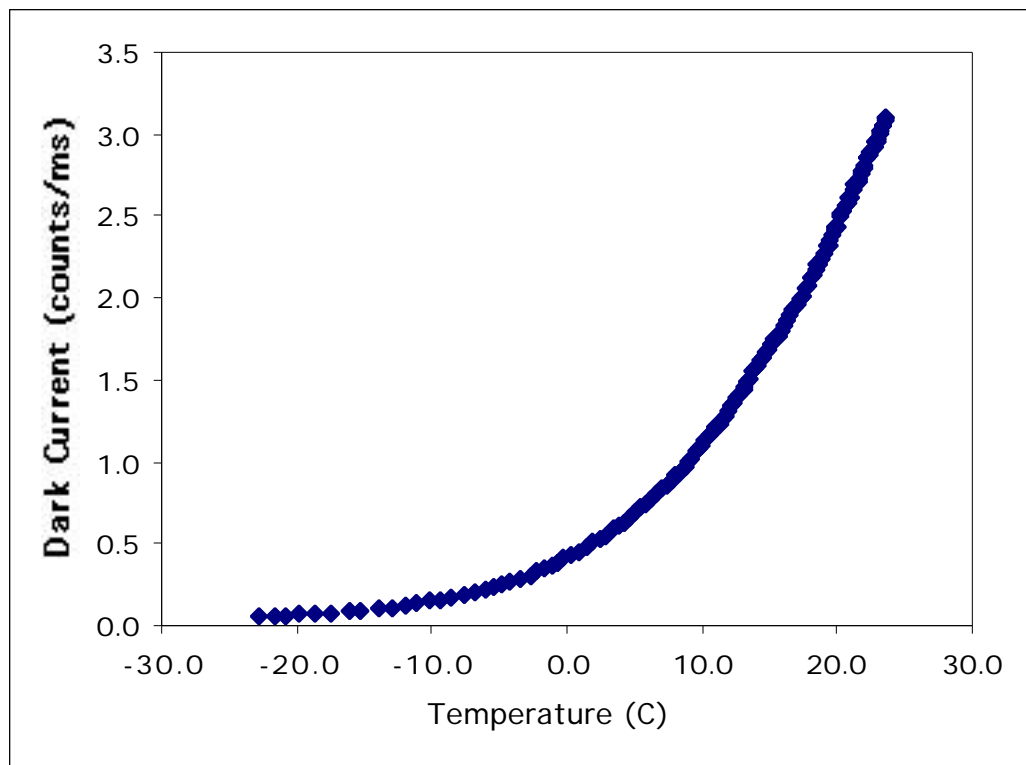

Figure 6: Dark Current vs. Temperature for the CCD from Head Two.

The dark current approximately doubles every 8 °C. Neutron damage has not been a problem for CCDs used in space because the CCDs are so cold that their dark current is negligible despite neutron damage.

Reversal of Damage by Annealing

Neutron damage is completely reversed by annealing at a sufficiently high temperature. The more energetic the neutrons that caused the damage, the higher the temperature required [2]. Most of the damage caused by 1-keV neutrons anneals at room temperature within a few months [2]. Some of the damage caused by 1-MeV neutrons requires annealing at 400 °C [2]. At 400 °C, all damage is annealed within half an hour [1]. We took the CCD of Head Four off its circuit board, and baked it at 90 °C, 150 °C, and 200 °C. At each temperature, we baked it until its dark current stopped decreasing. We measured its dark current by taking it out of the oven, returning it to its socket in Head Four, and letting it cool to room temperature. Room temperature in the laboratory that week was 26 °C. Table 2 gives the dark current before and after each bake.

Temperature (°C)	Initial Dark Current (ADC counts ms^{-1})	Final Dark Current (ADC counts ms^{-1})
90	3.9	1.7
150	1.7	0.68
200	0.68	0.44

Table 2: Dark current after baking at successively higher temperatures.

The remaining CCD Heads were left at room temperature. The dark currents dropped by 50%, 10%, and 36% in CCD Heads One, Two and Three respectively during the 150 days following the irradiation. Most of the changes took place in the first two weeks. We assume that these drops are due to room-temperature annealing. We cannot explain the difference between the drops.

Safety Factor

The 'safety factor' of a device with respect to neutron irradiation is the ratio of the dose at which we expect it to fail to the dose we expect it to receive. Ferrari [3] recommends a safety factor of at least four for electronic components in the end-cap muon detector. If the worst-case dose is 2.9×10^{12} 1-MeV n/cm^2, we should be confident that our electronics can endure 12×10^{12} 1-MeV n/cm^2. To calculate a safety factor for the TC255, we propose a model of the CCD dark current as a function of temperature and 1-MeV neutron dose.

Based upon our graph of dark current verses temperature (Figure 6), we assume an exponential increase in dark current with temperature.

We obtain the exponent by fitting an exponential curve to the graph. Based upon our plots of dark current verses radiation time (Figure 4), we assume the dark current is proportional to the neutron dose. To determine the constant of proportionality in our model, we measured the dark current in Head Two, and found it to be 2.8 counts/ms at 23 C, measured 150 days after the irradiation. We chose Head Two because it experienced the least recovery by room-temperature annealing, and is therefore a pessimistic choice. We obtain the following model.

$$I = D \, e^{\,T}$$

Where

$\quad\quad\quad = 0.082 \, °C^{-1}$
$T \quad\quad = $ temperature in $°C$
$I \quad\quad = $ dark current in counts ms^{-1}
$D \quad\quad = $ neutron dose in 1-MeV n cm^{-2}
$\quad\quad\quad = 0.055$ counts ms^{-1} $(10^{12}$ 1-MeV $n/cm^2)^{-1}$

This model overestimates the dark current in Head One and Head Three by up to 20% at temperatures between 21 °C and 26 °C. We assume this is because Head One and Head Two recovered more through room-temperature annealing than did the Head Two.

To estimate the maximum neutron dose the TC255 can tolerate, we need to know its operating temperature, the time for which dark current will be accumulating in the top pixels, and the maximum amount of dark current charge the pixels can hold and still obtain an accurate image.

Let us begin by estimating how much dark current charge the pixels can hold. Figure 7 shows four images. All are taken from Head Two, and all are strongly intensified to show contrast. In the top-left image, the dark current background is prominent, but you can see traces of a RASNIK image superimposed upon it. The top-right image is a 'background image', captured immediately after the first image but with the light source turned off during the exposure time. The lower-left image is the difference between the first and second images. The result is a clear and readily-analyzed RASNIK pattern.

The lower-right image is the difference between two background images. The pixel intensity along the top row of the image appears to be random. The average intensity along the top row of the original background images corresponds to a pixel charge of 40,000 electrons. We expect the standard deviation of this charge to be 200, or 0.67 counts. We also expect readout noise of 0.5 counts. When we add the readout and counting noises for both images together in quadrature, we obtain an

expected standard deviation for the difference between the two images of 1.2 counts, which is exactly what we observe.

The intensity of our images can vary from the electronic threshold (typically 45 counts) to 200 counts higher (typically 245 counts). At 45 counts, the pixels are empty. At 245 counts, they contain approximately 60,000 electrons. The dynamic range of intensity is therefore 200 counts. The total range of intensity in the lower-left image of Figure 7 is 16 counts. We can generate analyzable RASNIK images even if we use only 10% of our dynamic range for the pattern, leaving the other 90%, or 180 counts, to be filled by dark current. The background-subtraction also works in a BCAM (Boston CCD Angle Monitor). We used Head Two in a BCAM, and obtained 0.2 μm precision locating the center of the spot on the CCD, as compared to 0.1 μm with an undamaged CCD. The loss in precision corresponds do the increase in image noise that is caused by image subtraction and dark current. But a precision of 0.2 μm is adequate for even the most demanding ATLAS applications, so both RASNIK and BCAM instruments can operate with 180 counts of dark current in the top pixels.

We now turn to estimating the time for which the dark current will be accumulated. As we have seen, the effect of dark current upon an image is reduced if we speed up the transfer from the CCD to the CCD Driver. The next version of the CCD Driver will be able to read out a full 344 x 244 pixel image from a TC255 at the end of a 100-m cable in 42 ms. It will read out a 122 x 172 pixel image in 13 ms, and a 61 x 86 pixel image in 4 ms. The lower-resolution images are obtained by skipping pixels in the readout. We tested a 16-m RASNIK and a 16-m BCAM with the 122 x 172 images, and observed no loss of performance. The only allowance we had to make for the lower resolution was in the BCAM, where we had to increase diameter of the light spot on the CCD from 60 μm to 110 μm. When we tried the 61 x 86 pixel images, be found that the standard deviation of measurements taken with both the RASNIK and the BCAM almost doubled.

If we face severe radiation damage, therefore, we can switch to using the 122 x 172 images. The longest exposure required by any of our current instruments is 4 ms, so let us allow for a 4-ms exposure in ATLAS. This means that the pixels in the top of the 122 x 172 images will be at most 17 ms old (4 ms for exposure and 13 ms for readout). The maximum tolerable dark current is therefore 180 counts in 17 ms, or 11 counts/ms.
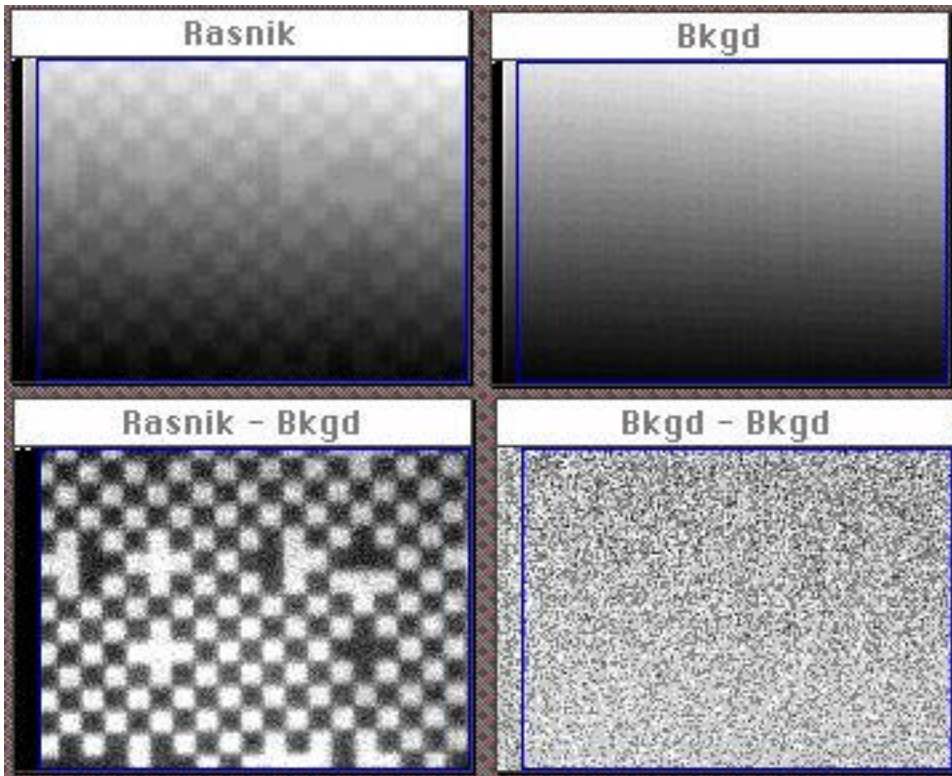
Figure 7: **These four images show how the dark current background in a radiation-damaged image can be subtracted out to leave a clear RASNIK pattern. All the images are intensified to show contrast. The bottom left image is the difference between the top images. The bottom right image is the difference between two background images.**

During the ATLAS experiment, power will be delivered to the CCD Heads only when they capture an image. This will be a small fraction of the time. We assume, therefore, that the operating temperature of the CCD will be equal to the ambient temperature. For ATLAS, the ambient temperature is supposed to be around 20 °C. We will reach the maximum dark current (11 counts/ms) at the operating temperature (20 °C) after a dose of 37×10^{12} 1-MeV n/cm^2, which gives us a safety factor of thirteen in the end-cap.

Conclusion

Our TC255 CCDs were strongly affected by the fast neutron irradiation. The effect of the neutron bombardment was to increase the CCD dark current. The pixels at the top of the image, which are transferred off the CCD last, have more time to fill up with dark current, and are therefore brighter, which makes the effect strikingly obvious in

images recorded from the damaged CCDs. We found that the dark current of a CCD is proportional to the time for which it is irradiated, and increases exponentially with temperature.

The increase in dark current was the only damage we observed in the CCD Heads, and it can be tolerated so long as none of the pixels in the image are allowed to become so full of dark current that they saturate. The dark current of any given CCD is so consistent from one image to the next that it can be removed by subtracting a previously-recorded dark-current-only image. With faster image retrieval, such as we expect from the next version of the CCD Driver, we estimate that the TC255 can endure a dose of up to of 37×10^{12} 1-MeV n/cm^2 and still capture images with adequate contrast and resolution for successful analysis.   Since the maximum dose we expect in the muon detector is 2.9×10^{12} 1-MeV n/cm^2, this gives us a safety factor of thirteen.

References

[1] Kegel at al, *Neutron Irradiation of Insulated Gate Field Effect Transistors*, Journal of Electronic Materials, Vol. 20, No. 11, 1991.

[2] Clinard et al, *Materials Performance in a Radiation Environment*, Materials Research Society Bulletin, April 1997. See www.mrs.org.

[3] Ferrari, *ATLAS Policy on Radiation-Tolerant Electronics Draft Two*, simulation TP36.

[4] Hashemi et al, *Pixel CCD DAQ,* ATLAS Note MUON-No-202.