

Characteristics of the NICMOS Detectors

C. J. Skinner¹, L. E. Bergeron & D. Daou

Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218

Abstract. We present an overview of the properties of the NICMOS flight detectors as measured on-orbit, including the flat-field response, the dark current, the linearity, and the read-noise. We show for the first time all the dependencies of the various components of a NICMOS dark exposure, and show how to generate “synthetic” dark current calibration files. An unexpected time-varying bias known as the “pedestal”, is described, along with some efforts to remedy it. We describe the effects on the detectors of exposure to very bright sources, and finally, we briefly describe the sensitivity of NICMOS to cosmic rays.

1. Introduction

NICMOS underwent its System Level Thermal Vacuum (SLTV) testing at Ball Aerospace during August and September of 1996, and was launched in February 1997. Prior to SLTV, STScI had no data using the NICMOS flight detectors, which are substantially different both physically and in their properties and performance to the NICMOS3 detectors widely used at ground-based telescopes. Thus all of our experience with these detectors has accrued during the last 12 months, and most of it in the last 5 months. During this time our understanding of the detectors has improved at a rapid pace, and the instrumental calibration has been evolving at a similar rate. We describe in the following sections the more important detector properties, as we understand them in September 1997.

2. Detector Properties

2.1. Flat-field Response

We show in Figure 1 the flat-field response for the three NICMOS detectors, measured for the F160W filters. These measurements were made during SLTV. On-orbit measurements of the flat-field response have been made for most of the Camera 1 and 2 filters by now, but many have not yet been reduced and converted into reference files. A detailed analysis of the flight-spare NICMOS detector was presented by Skinner et al., (1995a,b), and the amplitude of the flight detector flat-field response variations, both pixel-to-pixel and on longer scale-lengths, is basically consistent with that presented by Skinner et al. The conclusions of Skinner et al., regarding the effects of the flat-field response on photometric fidelity as a function of wavelength are valid for these detectors.

2.2. Darks

The “dark current” of an array detector is generally thought of as any signal which is accumulated during an exposure with no external illumination. In CCD cameras the only such signal is, in general, electrons generated in the silicon detector at a constant rate and

¹deceased




Figure 1. Flat-field responses for Cameras 1 (left) through 3 (right), on a uniform greyscale.

trapped in the potential wells. This is a continuous feed of electrons into the pixels, unprompted by photons, and the total charge accumulated in an exposure is linearly dependent on exposure time. This signal is not subject to modulation by the detector non-linearity or flat-field response. Thus in a calibration of an image from a CCD, the dark subtraction must be among the first steps. In the case of NICMOS detectors, the darks are considerably more complicated, with contributions from a number of sources. Additionally, it is possible to read the detectors non-destructively, and so it is possible to see the signal accumulating in each pixel during the course of an observation. However, the act of reading the detector in fact has a small effect on the signal. As a result of these complications, a significant “folklore” has been established among ground-based users of NICMOS detectors. One of the elements of this tradition is the belief that to calibrate a NICMOS exposure, it is essential to obtain a “dark” (i.e. unilluminated) exposure matched exactly in number of readouts and timing of all the readouts. We will show here that this is not, in fact, necessary.

The three components of the signal in a NICMOS dark exposure are “amplifier glow”, “shading”, and the real, linear dark current.

A NICMOS detector is a continuous single slab of HgCdTe, pixelated such that each pixel is individually bump-bonded to a single pixel of a CCD which is used as a readout. However, there are four separate readout amplifiers, each of which addresses one quadrant of the detector. Each time the detector is read out, the readout amplifiers, which are situated near the corners of the detector, are turned on. These amplifiers emit IR radiation that is detected by the pixels in the detector - similar to having a small “light bulb” in each corner. This produces a pattern of light that is highest in the corners and decreases towards the center of the detector. This is known as “amp glow”. A typical single readout produces about 20–30 ADUs of amp-glow per pixel in the corners of the detector, and 2–3 DN near the center. Since the readout time of the detector is the same each time (it takes 0.203 seconds to read the whole image), the on-time for the amplifiers is always the same for each readout, and thus the light pattern seen by the array is repeatable. So in a given readout, the amount of signal due to amp-glow in each pixel scales directly with the number of readouts since the last reset:

$$A(i, j) = a(i, j) * n_r \quad (1)$$

where $A(i, j)$ is the observed signal due to amp-glow in a given readout for pixel i, j , $a(i, j)$ is the amp-glow signal per readout (different for each pixel), and n_r is the total number of readouts of the array since the last reset. So in the corners of a full 26-readout MULTI-ACCUM there will be of order 500–800 ADUs due to amp-glow, along with the expected Poisson noise from this signal. It can immediately be seen from the size of this signal that

making excessive numbers of readouts during a MULTIACCUM exposure is harmful to the results, because although the amp-glow signal is highly reproducible and can be subtracted very effectively, the Poisson noise added by it can significantly degrade the S/N of the resulting image, even close to the center of the detector where the amp-glow signal is at its minimum. Amp-glow images for each detector are shown in Figure 2.

The bias level, or “DC offset”, in a given pixel in a NICMOS array, is time-dependent. This is the so-called “shading”, which visually in an uncorrected image looks like a ripple and gradual signal gradient across a given quadrant. The pixels in a given quadrant of a NICMOS detector are read out sequentially. It takes a little over a μ sec to read a single pixel, and so with four readout amplifiers reading in parallel it takes just over 0.2 sec to read the entire 256x256 pixel detector. Considering a quadrant as an array of $i \times j$ pixels, the readout sequence consists of reading sequentially along a detector row i , clocking j from 1 to 128, then moving to row $i + 1$ and clocking j from 1 to 128, and so on. Since the amplifier bias changes pseudo-exponentially with time over the course of the readout, the observed signal, in the absence of any external illumination, varies rather slowly along the rows (i), but rather rapidly along the columns (j). This signal is not accumulated in the pixels each readout, but rather is superimposed on the actual signal at the time of each detector readout. The shading signal is not the same for each readout. Its amplitude, and to some extent its shape, are a function of the time since a pixel was last read out (*not* reset). So readouts with the same DELTATIM (this is the keyword used in NICMOS data to denote the time since the previous readout) will have the same bias signal, for a given pixel. The dependence of this bias on DELTATIM is nearly logarithmic and quite repeatable. It is possible to find a numerical fit to the shading function in DELTATIM for each pixel of each detector. Thus it is (in principle) possible to predict what the bias signal is in any given pixel for any possible readout sequence. Another way to attack the problem is to make an average image of the bias for each of the DELTATIMs in the predefined MULTIACCUM sequences. Then to build a synthetic dark, the bias component can be had by using the bias image for each appropriate DELTATIM in the sequence:

$$B(i, j) = S(i, j, DELTATIM) \quad (2)$$

where $S(i, j)$ is the bias signal in a given pixel as a function of DELTATIM. It is the latter operation which is currently used in generating our synthetic darks.

The linear dark current component is the traditional observed detector dark current when no outside signal is present. This component scales with exposure time only:

$$D(i, j) = T * d(i, j) \quad (3)$$

where $D(i, j)$ is the observed dark current signal in pixel i, j for a given readout, T is the time since the last detector reset, and d is the dark current (in e^-/sec). The NICMOS dark current is extremely small and is very difficult to measure. It is approximately $0.05 e^-/\text{sec}$ for Camera 2, and no larger than about half this value for Cameras 1 and 3 (we have only upper limits for these two).

The challenge in calibrating these dark components for NICMOS detectors is disentangling the three components so that each can be measured independently of the other. Fortunately each component has entirely different sets of dependencies, and this can be used, along with the flexible way the NICMOS detectors can be operated, to achieve our goal. First we note that the amp-glow is dependent only on the number of readouts since the last detector reset. This means that to calibrate the amplifier glow, we need to find images with identical shading components but different number of readouts. Because of the dependence of shading, this turns out to be easy: if we look at any of the STEPxxx MULTIACCUM readout sequences, we find that after a set of logarithmically increasing DELTATIMs, they settle to a constant DELTATIM. For example, in the case of the STEP256 sequence, after an exposure time of 256 sec, readouts occur every 256 sec, and thus the shading is identical




Figure 2. Amplifier glow for Cameras 1 (left) through 3 (right), on a uniform greyscale, and below a plot of row 3 (near the bottom) of each camera.

for each of these linearly spaced readouts. In this linear regime, the difference between the n th and the $(n+1)$ th readouts is simply the signal added by 1 readout's worth of amp-glow (ignoring the dark current). Taking a full set of m linearly spaced readouts yields m lots of amp-glow, and thus obtains the highest S/N for the resulting amp-glow image. Finally, to correct the amp-glow images for the small effects of the linear dark current, we can compare the amp-glow images generated as above using many different STEPxxx readout sequences, and use the difference in DELTATIMs to measure the linear dark current. The resulting amp-glow images for all three cameras are shown in Figure 2.

In order to measure the shading, first we correct each read of a MULTIACCUM exposure for amp-glow, using the amp-glow calibration images obtained above. Having removed this component, shading is now the dominant remaining component. If we plot the amplitude of the shading signal against DELTATIM, we find the two are very strongly correlated. In fact, the shading amplitude is dependent only on DELTATIM, and has no dependence on time since last reset whatsoever. This is illustrated by the fact that if we find any read in the STEP16 pattern whose DELTATIM is 16 seconds, for example, the shading amplitude for the same pixel of the same detector will be identical for a read from the STEP256 pattern for which DELTATIM is 16 seconds - regardless of how many reads have occurred previously in either observation. In Figure 3 we plot the shading amplitude against DELTATIM for a pixel near the center of a quadrant, using observations made using many different MULTIACCUM sequences. Figure 4 shows the shading image as a function of DELTATIM for all three cameras. The detectors are mounted in the cameras such that the readout directions are rotated by multiples of 90 degrees with respect to one another, and so the shading patterns in Cameras 2 and 3 run parallel to one another, while that for Camera 1 is in the orthogonal direction. In Camera 1 the shading generates a bright band along the first row to be read out of each quadrant, parallel to the time axis in Figure 4. In Camera 3 a similar bright band is generated, but now the band runs orthogonal to the time axis of Figure 4. Note that the shading generates a large negative signal. The gradient of the

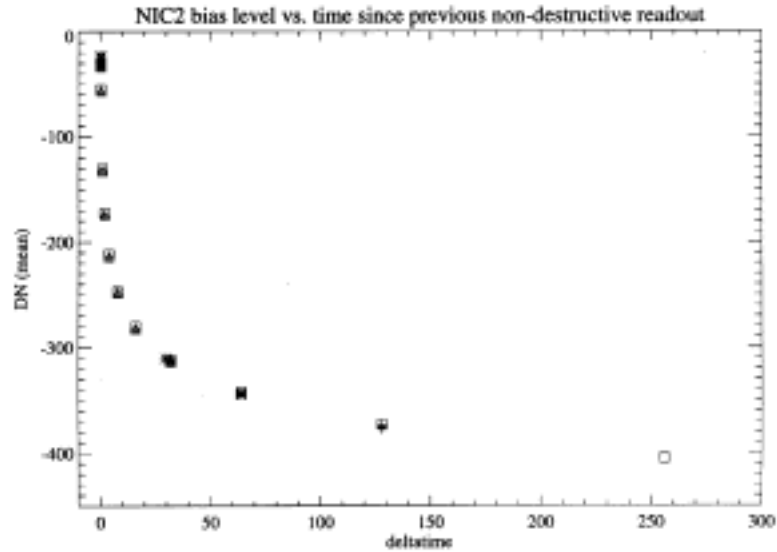


Figure 3. Shading amplitude for a single pixel as a function of DELTATIM. Different symbols represent measurements from different MULTIACCUM sequences: equal DELTATIMs from different sequences coincide, showing that shading is a function *only* of DELTATIM.




Figure 4. Shading images for each camera plotted as a function of DELTATIM (shown in seconds at the top).

shading function is seen to be quite steep in the direction orthogonal to the bright bands seen in Cameras 1 and 3: this direction is the “slow clocking direction”, as the time between pixel readouts in this direction is 128 times a single pixel readout time. In the orthogonal direction, the “fast clocking direction”, it is more difficult to see the gradient in the shading — but it can be seen in Camera 3 by virtue of the quadrant boundaries.

The linear dark current can be measured a number of ways. The technique which has been used so far has been to derive, for a number of MULTIACCUM sequences for which high quality on-orbit darks have been obtained, synthetic dark current images using the amplifier glow and shading contributions as determined above. These are then subtracted from the observations. The residual is seen for Camera 2 to be a very low amplitude signal increasing roughly linearly with time. This is the linear dark current. For Cameras 1 and 3 no residual could be seen above the uncertainties, indicating that the dark current is very small for these two cameras.

The total “dark” signal in any given pixel of any given NICMOS MULTIACCUM readout is just the sum of the 3 components described above:




Figure 5. Histogram of the pixel values in the difference image between an observed and a synthetic STEP8 dark (see text)

$$DARK(i, j) = D(i, j) + A(i, j) + B(i, j) \quad (4)$$

An IDL routine has been developed to make synthetic darks. This routine uses the amplifier glow image displayed in Figure 2, and simply multiplies it by the accumulated number of reads in order to generate $A(i, j)$ for each readout. The shading as a function of DELTATIM has been populated by the technique described above, yielding an array of shading images some of which are displayed in Figure 4. An appropriate image is picked out of this array in order to generate $B(i, j)$ for each readout. Finally, the linear dark current as described above is calculated from the elapsed time per readout. The routine now sums the three contributions for each readout. The calibration database has been populated with MULTIACCUM darks for all sequences which have not yet been observed on-orbit, or for which other effects, such as the pedestal, have contaminated the early on-orbit dark observations. There are plans to tune up the synthetic dark algorithm somewhat (see next section) and eventually release it as an STSDAS tool in the NICMOS package.

Comparisons of on-orbit to synthetic darks show that the differences are relatively small — usually of the order of a few ADUs, with the largest differences in the corners of the detectors. Most of this can be alleviated by better characterizing the amp-glow from on-orbit data. In particular, we are observing some low-level non-linearity in the behavior of the amp-glow — that is, the amplifier glow contribution in observed darks is seen to differ a little from being an exact linear multiple of a single read’s worth of amp-glow. Further investigation of this is required, but we note that the effects are small, and will rarely be significant. As an example, we show here the difference between a synthetic STEP8 dark calibration reference file, and a STEP8 reference file generated as the mean of a set of 8 22-readout STEP8 dark observations obtained as part of the 7116 ERO observation. The observed dark contains a few hundred Cosmic Rays which were not successfully median filtered out when the individual exposures were combined. (Overall we would expect to have received about 2000 Cosmic Ray hits in the set of darks; most of these have been successfully removed by the combination algorithm, with about 10–20% of them remaining, mostly because their amplitude was too low to be detected.) There is no residual shading pattern in the difference image, and very little residual amp-glow (there is about a couple of ADUs in the corners). There is a small gradient across the image from right to left, with an amplitude of about 1 ADU. However, the residual signal is very small. The histogram




Figure 6. Response of pixel 100,100 of Cameras 1 (left) through 3 (right) to a uniform signal, illustrating the linearity behaviour. The dashed line is the theoretical pixel response (if the behaviour were linear), while the solid line is the observed response, plotted only up to the signal level at which the pixel is deemed to be saturated and uncorrectable.

of pixel values in the difference image, which is shown as Figure 5, shows that the modal difference is about 1 ADU between the observed and synthetic dark.

2.3. Linearity

Measurements of NICMOS3 detectors, and their successor the HAWAII detectors, by Rockwell (e.g. Poksheva et al 1993; Kozłowski et al 1994) show non-linearity at both very small and very large total charge accumulations. Measurements of the flight detector non-linearity during SLTV suggested that in fact the detectors were quite linear up to charge accumulations of order 5000 ADUs, and that non-linearity was evident from there up to their saturation level of around 30000 ADUs. A typical linearity curve for a single pixel for each camera, derived from SLTV data, is shown in Figure 6. The linearity curve is somewhat different for different pixels, and the saturation levels show a fairly wide dispersion. From on-orbit measurements we now have indications that at least some pixels may now be saturating at somewhat lower charge accumulations than were determined from SLTV. Whether this is attributable to a problem in the SLTV measurements, to problems in their analysis, or to a real change in the linearity properties since launch, is not currently clear. A new set of measurements of the linearity are planned to take place during October 1997.

2.4. Read Noise

The read noise for the pixels in the NICMOS detectors is highly variable from pixel to pixel, in a basically random fashion. It was measured during SLTV by the usual means, plotting for each pixel the mean detected signal versus its standard deviation. The result is a curve whose gradient yields the gain in e^-/ADU and whose intersection with the signal axis (at zero detected signal) yields the read noise. Images of the read noise for each detector are shown in Figure 7.

3. The Pedestal

This effect was first observed in SLTV data, but not understood. It has been observed on-orbit since launch, and we are now finally beginning to understand it. When a detector is to be unused for some period (e.g. during Earth occultations or filter wheel movements), it enters a mode known as “autoflush”, in which it is reset several times per second to




Figure 7. Images of the read noise for Camera 1 (left) through Camera 3 (right), scaled from 15 electrons (black) to 55 electrons (white).

prevent saturation. When the detector is used after a period in autoflush, its output is observed to be somewhat unstable. It is as though an excess bias is present on emergence from autoflush, which decays on a timescale of many minutes. This excess bias has become known as the “pedestal”. Its amplitude seems to be roughly uniform across a quadrant, but to vary somewhat from quadrant to quadrant. The largest amplitude seen for the pedestal so far is about 20–30ADUs. It seems that the amplitude is usually greater after long periods of autoflush (e.g. Earth occultations), and much smaller after short periods of autoflush (e.g. a spacecraft dither), and the timescale for decay from a large amplitude pedestal can be as long as 30 minutes. An example of the pedestal in a series of STEP8 darks is shown in Figure 8, while in Figure 9 the modes of the reads in another sequence of darks, this time with Camera 1, are shown.

We currently believe that the pedestal may be driven by small changes in the detector temperature. The readout amplifiers are known to be a source of heat. At launch, during autoflush the amplifiers were switched off. During an exposure, the amplifiers must be switched on for readouts, and this led the detector temperatures to rise during exposures. The zero level for the detector output is highly temperature-sensitive, and so the temperature drift causes a zero-level drift which appears as a bias drift, or excess dark signal. Recently the flight software has been changed such that the amplifiers are left switched on during autoflush, and only switched off between readouts during exposures, in an attempt to reduce temperature fluctuations of the detectors. As might be expected, the temperature changes of the detectors are now in the opposite direction, and the pedestal also appears to have reversed its sign. Its amplitude appears to have been reduced, although it has not disappeared. Further work is underway to categorize the new pedestal properties.

4. Persistence

The observed persistence properties are described by Daou & Skinner (1997a,b). Saturation of the NICMOS detectors can cause after-images, or “persistence”, which can linger for up to 30 minutes after the saturated exposure was completed and read out. The count-rates in these persistent images can be significant, and the mechanism and behaviour are not understood. Recent tests have suggested that following a saturated exposure, a few short ACCUM images (minimum exposure time) with many initial and final reads can almost completely eradicate persistence. About 20 to 30 initial and final reads appears to be sufficient.




Figure 8. Images of the pedestal effect in Camera 2. Every readout of this series of eight STEP8 darks is shown sequentially along the x direction, with the first exposure at bottom, and last at top; the last has been subtracted from every one, to show the effect of the decaying pedestal signal.




Figure 9. Modes of a series of darks in Camera 1, differenced exactly as shown in Figure 8, shown as a function of readout number for each exposure in the series. The first exposure shows a strong pedestal, rising to an excess bias of 20 ADUs by the end, and subsequent exposures show decreasing effects.




Figure 10. Part of a Camera 1 image showing a massive CR hit: the readout immediately after the hit is on left, and the previous readout at right.

5. Cosmic Rays

NICMOS detectors are quite sensitive to CRs. For the WFPC2 CCDs, the energy spectrum of detected CRs shows a distinct peak. For NICMOS we observe no such peak; instead the detectors appear to be sensitive to CRs down to energies as low as the read noise (Calzetti, 1997). It is therefore quite possible that a non-negligible fraction of the noise in on-orbit NICMOS images is due to low-energy CR hits. Typically, about 2 CR hits are experienced per second on each detector, if we count hits only down to energies four times the read noise or larger, but the rate is rather variable. Occasionally enormous CR hits are experienced, in which a radius of up to 10 pixels can be instantaneously saturated, with a surrounding ring of what may be secondary emission. An example of such an image is shown in Figure 10. These hits can be distinguished from real sources by the fact that they appear instantaneously between reads, that the entire hit region is saturated without surrounding regions of lower counts, and that there is no spider pattern to the image. We are currently seeing a few of these hits per month. These hits can leave persistent after-images in subsequent NICMOS exposures up to at least half an hour later.

References

- Calzetti, D., 1997, Instrument Science Report NICMOS-97-022 (Baltimore: STScI).
 Daou, D. & Skinner, C. J., 1997, Instrument Science Report NICMOS-97-023 (Baltimore: STScI).
 Daou, D. & Skinner, C. J., 1997, Instrument Science Report NICMOS-97-024 (Baltimore: STScI).
 Kozlowski, L. J., Vural, K., Cabelli, S. C., Chen, C. Y., Cooper, D. E., Bostrup, G. L., Stephenson, D. M., McClevige, W. L., Bailey, R. B., Hodapp, K., Hall, D., Kleinhans, W. E., 1994, SPIE Vol. 2268, 353.
 Poksheva, J. G., Tocci, L. R., Farris, M. C., Tallarico, S. I., 1993, SPIE Vol. 1946, 161.
 Skinner, C. J., 1995, Instrument Science Report NICMOS-006 (Baltimore: STScI).

Skinner, C. J., Mentzell, E., Schneider, G., 1995, Instrument Science Report NICMOS-005 (Baltimore: STScI).