

Improving on the NICMOS Calibration Pipeline Products

C. J. Skinner¹

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

Abstract. We describe a variety of operations which can be carried out using IRAF software to improve upon the quality of the calibrated data when a variety of problems are present in it. The problems covered include such diverse topics as the pedestal, problems with Cosmic Ray rejection, problems with saturated or non-linear pixels, and multiple coronagraphic holes appearing in images.

1. Introduction

We have now had almost a year of experience of processing NICMOS images. The first 6 months consisted only of processing images obtained during SLTV at Ball Aerospace, but even these data contained a number of the problems encountered on-orbit (the “pedestal”, Cosmic Rays), and so gave us early indications of the areas that would later prove the most difficult to deal with in an automated calibration pipeline. For the past 6 months we have been taking data on-orbit, at an ever-increasing rate, and our experience in recognising, and in some cases finding ways to cope with, these difficulties is increasing accordingly. There are some problems for which there is no cure, and for which the major battle is recognising that the problem exists in a given dataset. Other problems can be cured by some extra processing. We describe in this paper our experience with such problems. The list here is by no means exhaustive, and we expect that by the time of the next Calibration Workshop, there will be both a lengthier list and for already known problems better solutions.

2. Pedestal Cures

The pedestal basically acts like a time-dependent extra bias which has been added to the signal. Therefore, in order to fix it, all that is needed in principle is a measurement of the extra bias that has been added. In practise, determining this bias can of course be exceedingly difficult, and in some cases impossible.

The simplest situation is where more than one exposure has been taken back-to-back, at the same pointing. In this case, we should be able to difference the two exposures to obtain a first order estimate of how the pedestal was decaying between the exposures. An example of this is plotted here, where we had two back to back exposures. Figure 1 shows the first exposure, which was badly affected by the pedestal: the pedestal is easily diagnosed by the large inverse flat-field pattern which is stamped on the calibrated image. To estimate its amplitude, we first subtracted the zeroth read from all reads of the first and second exposures, and then differenced the results. This told us how much the signal had changed between the two exposures. However, there is no guarantee that the pedestal had decayed to zero by the beginning of the second exposure. We subtracted this pedestal measurement from the first exposure, and then recalibrated it with CALNICA. The result still yielded an inverse flat-field pattern, but with an amplitude roughly half that present

¹deceased




Figure 1. Left: image in Camera 2 badly affected by the pedestal. Centre: pedestal image, determined by subtracting from the image at left an identical exposure taken immediately afterwards. Right: the same image, with the pedestal image subtracted and then recalibrated using CALNICA – almost all the inverse flat-field signature of the pedestal has been removed, but the S/N is worse.




Figure 2. Top row: the exposure whose last read is displayed in Figure 1, with all 10 reads shown, the zeroth on the right and the final on the left. Bottom row: the pedestal corrected exposure (as described in the text) plotted as for the top row. The removal of the inverse flat-field pedestal signature is dramatic, although the S/N is worse.

in the original calibrated image, confirming that indeed there was still pedestal present in the second exposure. We therefore doubled our pedestal signal measurement, subtracted this from the original raw data, and recalibrated a second time. The result is plotted in Figure 1, and almost all the pedestal signal has now been removed. In Figure 2 we show the whole sequence of 10 reads with the original calibration and the recalibration after pedestal removal. Figure 3 shows the amplitude of the pedestal as a function of time since reset for the duration of the first of the pair of exposures.

The major problem with the above technique is that we are making a pedestal amplitude estimate for each pixel, but the data are intrinsically noisy, and by combining the raw image with our pedestal estimate, we are roughly doubling the noise in the image. We have plotted the pedestal image for the final read in Figure 1, and it appears to be fairly uniform (this is typical for the pedestal, although it does sometimes have more of a quadrant dependence than we found here). Therefore we could obtain a less noisy result by determining an average signal for each quadrant, and subtracting that instead. The type of average to determine is clear: a mean would work very poorly because of the very bright star present in this field, but the mode will work rather well. Determining the quadrant-based modes is a more involved effort, and less amenable to inclusion in an IRAF script, but can be done.




Figure 3. Amplitude of the pedestal signal seen in Figure 2, as a function of time since reset in the first exposure.




Figure 4. Same as Figure 2, but this time instead of subtracting the measured pedestal on a per pixel basis, we have subtracted the mode of the pedestal image for each readout: this yields better S/N.

Using the mode in this manner also works when there are repeated observations, but they have been dithered, so that sources appearing at different locations in the multiple images would spoil the pedestal image calculated as described in the previous paragraph. We show the result of subtracting the mode of the entire pedestal image for each readout in Figure 4, and the result is seen to be very much less noisy for the final few reads. A few of the reads in this image show significant residual pedestal, indicating that the mode is still not estimating the pedestal amplitude perfectly, but the results are clearly better than the previous pixel-by-pixel subtraction. A small further improvement would result from a quadrant based subtraction rather than an entire image subtraction, although in this particular case we find that the amplitude does not vary much from quadrant to quadrant (Figure 1).

The situation becomes much more difficult when there is only one exposure in a given filter. In this situation, the best we can do is to measure the modes for each quadrant, assume that most of this value is due to inverse flat-fielded pedestal signal, and remove this modal value from each quadrant, recalibrate through CALNICA, and then repeat the process iteratively until the modal value has been minimized. An automated procedure of this sort is under development at STScI currently, and may be ready for general release in a few months. However, this use of the mode as an indicator of pedestal will only work provided the field is relatively sparse and Cosmic Rays have been rejected. In a field which is dominated by “real” signal, the mode will not yield a useful estimator of any pedestal, and in this situation there is probably very little that can be done. Anecdotally there has been some suggestion that the “photometrically challenged column” (see Skinner, Bergeron & Daou 1997) may provide a useful barometer of the pedestal, but this has yet to be demonstrated.

3. Multiple Coronagraphic Holes

One of the most frequently received questions at STScI in the NICMOS group regarding data calibration is “Why does my calibrated Camera 2 data contain what looks like three coronagraphic holes?”. If we examine the 300Gb of NICMOS images in the archive, we will find coronagraphic holes at positions which cover a significant fraction of the upper left quadrant of Camera 2. The root cause of the problem is that as the NICMOS dewar has variously expanded or contracted through the course of its 1.5 year life since it was filled with cryogen in 1996, the detectors have been moving in a direction which, while predictable in a gross sense, has many unfathomable fine details. As the Camera 2 detector has moved, the image of the Coronagraphic hole in the FOM has moved across the face of the detector. As a result, exactly where the hole appears in any Camera 2 calibration reference file depends on when the relevant observations were made.

In the case of the flat-field response reference files, some observations were made during SLTV in 1996, some were made during SMOV in the first half of 1997, and some have been made during Cycle 7. Currently, for some filters we have observations made at only one epoch, while for others we have as many as three. Where there are multiple observations, we have been able to remove the hole from the reference file by populating each pixel only with data obtained at a time when that pixel was not looking at the hole. Since some variations in flat-field response between the different epochs have been observed, the results do often yield a visible “patch” at one location or another. If data are available from only one epoch, then it is likely that a hole will be present in the reference file at a location which does not correspond to the current hole location. A hole in the reference file will be flagged in its Data Quality extension, and CALNICA will then flag these pixels.

The hole does not appear in the read noise or the dark current calibration reference files. It does appear in the linearity reference file, however, and in this case, at the time of writing the only suitable calibration data are from SLTV. A new linearity calibration is planned to take place during October, and following this it should be possible to remove the hole from the reference file.

4. Cosmic Ray Problems

A multitude of problems have been experienced with various versions of CALNICA in dealing with CRs, and they are still sometimes problematic. The basic technique adopted in CALNICA since launch has been to calculate a count-rate for each read by determining the total counts received *since the previous read*, and dividing this by the time between the two reads. The count rates for all the reads are then compared, and outliers beyond some limiting threshold are flagged as CR hits. There are, however, a multitude of pitfalls in this process, and we have slowly been discovering some of them and refining the technique.

First, the technique can only ever work on bright sources if the linearity correction is good. On a pixel in which 20,000 ADUs have been accumulated, the S/N is likely to be estimated as 100 or more by CALNICA. If the linearity correction is significantly in error, the countrates may be seen to differ by CALNICA by more than its built-in threshold (5σ), and a CR hit will be flagged. If the saturation value in the linearity calibration reference file is incorrect for the pixel, and the pixel does in fact saturate during the exposure but the saturation is not noticed by CALNICA because of the error in the reference file, then CALNICA may flag some or even all of the reads as CR hits, because all of them may deviate from the mean count-rate it determines. Both of the above problems have occurred frequently. We have reduced the saturation values in the linearity reference files by 10% from the values measured in SLTV, and this appears to have almost completely cured the frequent failures to detect saturation. Incorrect CR identification due to what appear to be incorrect




Figure 5. Left: 2500 second image in Camera 2 badly beset by Cosmic Rays. Right: the final, CR rejected output product from CALNICA: almost every CR was successfully removed, despite rather indifferent dark subtraction and significant pedestal amplitude.

linearity corrections do still sometimes occur. We should soon be obtaining on-orbit linearity calibration measurements, and hopefully this will remedy the problem.

Second, good dark correction is necessary in order for the CR identification algorithm to succeed. The dark current amplitude is quite large for these detectors, as well as highly time variable (see Skinner et al. 1997). Therefore, if it is not correctly subtracted, CALNICA will see a varying count-rate, especially for pixels with only faint (or no) illumination, and again will incorrectly flag CR hits. Situations where this has occurred with NICMOS observations so far include observations badly affected by the pedestal, and observations for which no good dark reference file was available (early in SMOV this included all observations, because no dark calibrations were available at all then; more recently the dark calibration reference files again became rather poor after the Flight Software change on August 22 1997 – see Skinner et al. 1997). Under these circumstances almost all the pixels are sometimes flagged as CR hit, especially early in an exposure when the accumulated source counts are often small and the dark current is changing rapidly.

Given the kinds of calibration problems mentioned in this section so far, effective CR correction will be impossible, and it is advisable to switch off the CR-Identification step in CALNICA (see Bushouse 1997). A new version of CALNICA which is currently under development will contain the ability for the user to set the CR detection threshold manually. Currently the threshold is hardwired in the code.

So long as there are no problems with the reference files, and there is no large instrumental anomaly (such as the pedestal) present, CALNICA usually currently successfully identifies the great majority of CRs in a MULTIACCUM observation. A good example of CALNICA performing well in this regard in the face of adversity is shown in Figure 5. Here we see the final read of a 2500 second exposure in which the dark current was rather poorly corrected; nonetheless, almost all of the CRs seen in the calibrated exposure on the left were successfully removed in the final product on the right. Incidentally, this exposure also shows the presence of three coronagraphic holes, as well as a significant amplitude pedestal and a very prominent “photometrically challenged column”.

5. Calibration Reference File Problems

It is not uncommon for data to be obtained and calibrated at a time when ideal calibration reference files are not available. A good example currently is the lack of any dark calibrations since the August Flight Software change. All data obtained since then have been calibrated in the pipeline using old darks which only partially correct the data. Whenever you receive calibrated data (either as a GO or as an archival researcher), you should check the names of the reference files which were used to calibrate the observations: these can be found by reading the paper products which were sent to you if these were GO data, or can be generated yourself using the paper products software available in IRAF, or you can find the names by reading the latter part of the primary header of the data files (raw or calibrated). The NICMOS pages on the STScI WWW site contain a page listing all the available calibration reference files for NICMOS. Included in this listing is a usage date, which indicates the observing periods during which each reference file is believed to be appropriate. If the file that was used to calibrate your data appears to have been superseded by a more recent file, you should obtain it and recalibrate the data. If you are not sure whether the reference files used were appropriate, contact your Instrument Scientist or the STScI Help Desk.

6. Zero Read Problems

NICMOS has no shutter. In order to make an exposure, the detector is first reset, and then a “zeroth” read made. All subsequent reads have this zeroth read subtracted in order to define the counts per unit time. Unfortunately, the zeroth read occurs 0.203 seconds *after* the reset. For any pixel exposed to a bright source, this means that the total charge accumulated in the potential well will *not* correspond with the counts recorded after zeroth read subtraction, and therefore the subsequent linearity correction and saturation detection steps will not be successful, leading in turn to major CR detection problems.

In the earlier versions of CALNICA, we were not aware of this problem, and bright sources regularly caused major problems in calibration. A new version to be released shortly attempts to correct this problem by using the first few reads to estimate the count-rate for each pixel, and then extrapolating backwards to estimate how many counts should have been recorded in the 0.203 seconds before the zeroth read. If this turns out to be a statistically significant number, it is added onto the signals in all the subsequent steps, after the zeroth read correction. It may seem odd to be subtracting a zeroth read, and then adding back in all the counts that were subtracted! Unfortunately, the zero level for the ADC in these detectors is not stable. From one exposure to the next there is some drift in the zero level, which can approach one thousand ADUs, and is randomly variable from one quadrant to another. It is necessary to subtract the zeroth read in order to correct for this instability. The backwards extrapolation technique outlined above usually works quite well, but in pathological cases the pixel may saturate at the first read, but not be saturated in the zeroth read. To allow for these cases, we have given CALNICA the ability to read a “super zeroth-read” image, which is the average zeroth read level determined from a large number of zeroth read observations. This image can be subtracted from the zeroth read to yield the counts in this read for the pixels which saturated by the later reads. We have found this technique often works quite well, but sometimes still is unable to cope with very bright sources. In such cases it is necessary to individually examine the pixels concerned in the various reads in order to make a determination of the “real” counts: usually there are not very many such pixels in an image, and so the task is not overwhelming.

7. Conclusions

The CALNICA and CALNICB software is in general performing rather well, and most of the problems which need to be corrected before subsequent reprocessing of the data relate to instrumental anomalies (e.g. the pedestal, linearity changes) or changes in instrument operation (e.g. the August 1997 Flight Software change). A few suggestions for how to approach such refinement of the pipeline processing have been suggested here. We expect that over the course of the next year or so our understanding of the best ways to process NICMOS data will continue to improve.

References

- Bushouse, H., 1997, this volume.
Skinner, C.J., Bergeron, L.E. & Daou, D., 1997, this volume.