Instrument Science Report WFC3 2008-39

WFC3 TV3 Testing: IR Channel Nonlinearity Correction

B. Hilbert
2 June 2009

ABSTRACT

Using data taken during WFC3's Thermal Vacuum 3 (TV3) testing campaign, we have calculated the correction coefficients necessary to remove non-linearity effects from IR detector (IR-4, FPA165) data. We find that a set of coefficients for each quadrant of the detector produces a better correction than calculating and applying a separate set of coefficients to each pixel.

Introduction

HgCdTe detectors, such as that in the WFC3-IR channel, are inherently non-linear in their response to incident photons. A description of the physics behind this effect is given in Hilbert (2004a). The goal of this study is to calculate a correction for these effects. Non-linearity correction calculations for previous, non-flight IR detectors are detailed in earlier ISRs (Robberto and Hilbert, 2005, and Hilbert, 2007). We used a similar method in this case. The basic technique involves monitoring the detector's response to incident photons at low signal levels where the measured signal increases linearly with the number of incident photons. By extrapolating to higher signal levels where the detector response is not linear we can quantify the extent of the detector's non-linearity and correct for it.

Data

The data used for these non-linearity calculations were collected during TV3. For this test, the detector was exposed to uniform illumination from the optical stimulus (called "CASTLE") Tungsten lamp. The goal was to expose the detector to enough flux such that most pixels became saturated by the end of the ramp. For our purposes, the saturation level is defined as the signal level at which the detector response becomes non-linear by 5%. This is the maximum signal level to which IR data will be corrected, according to the contract end item (CEI) Specifications. The data consist of 12 ramps, each composed of 15 non-destructive reads of the IR detector, following the STEP50 sample sequence. In this timing pattern, initial reads sample the detector quickly, followed by increasing exposure times between later reads.

Given the low signal threshold at which non-linearity effects were observed in TV2, we wanted to collect as many low-signal reads as possible, to allow for accurate linear fits in the subsequent data analysis. This is the reason behind the use of the STEP50 sample sequence for the data collection, rather than the SPARS sample sequence used in TV2. STEP50 ramps sample the signal more often earlier in the ramps, leading to more samples with which to fit lines.

Analysis

All ramps were run through the initial step of the WFC3 IDL data reduction pipeline (Hilbert, 2004b), using the vertical inboard reference pixels to subtract the bias level from each read. We subtracted a median STEP50 dark current ramp (created from other TV3 data taken close in time to the non-linearity data) from each of the non-linearity ramps. This removed pixel-to-pixel variations in zero level. Normally this data reduction step is accomplished by subtracting the initial read of a given ramp from all subsequent reads of that ramp. However, our goal is to obtain a measure of the detector's non-linearity versus the total signal. If we were to subtract the initial read from the others in a ramp, then the signal in the resulting ramp would not include any photons recorded between the beginning of the ramp and the first read. This would result in a shift of the signal versus non-linearity relationship.

Finally, a bad pixel mask, derived from other TV3 data, was applied to each non-linearity ramp. This mask flagged hot and dead pixels, as well as pixels previously found to have signal in the initial read of a ramp more than 3σ from the mean. In all just over 2% of the active science pixels were masked and ignored in subsequent data analysis.

Non-Linearity Calculation Method

After these initial data reduction steps, non-linearity correction coefficients were calculated in much the same way as with the TV2 data and described in Hilbert (2007).

First, using the 12 input ramps we created a median 15-read ramp, in order to limit the effects of noise on future line and curve-fitting steps.

For each pixel, we created a plot of the measured signal versus time. We then calculated a best-fit line to this data, but only for low signal levels which should be unaffected by non-linearity effects. We used a maximum signal level of 4500 DN for our line-fitting. Figure 1 shows the measured signal versus time, and the corresponding linear fit for one pixel. Our goal, as outlined in the CEI Specifications, is to be able to remove the non-linearity effects up to the point where the measured signal deviates from this linear fit by 5%.

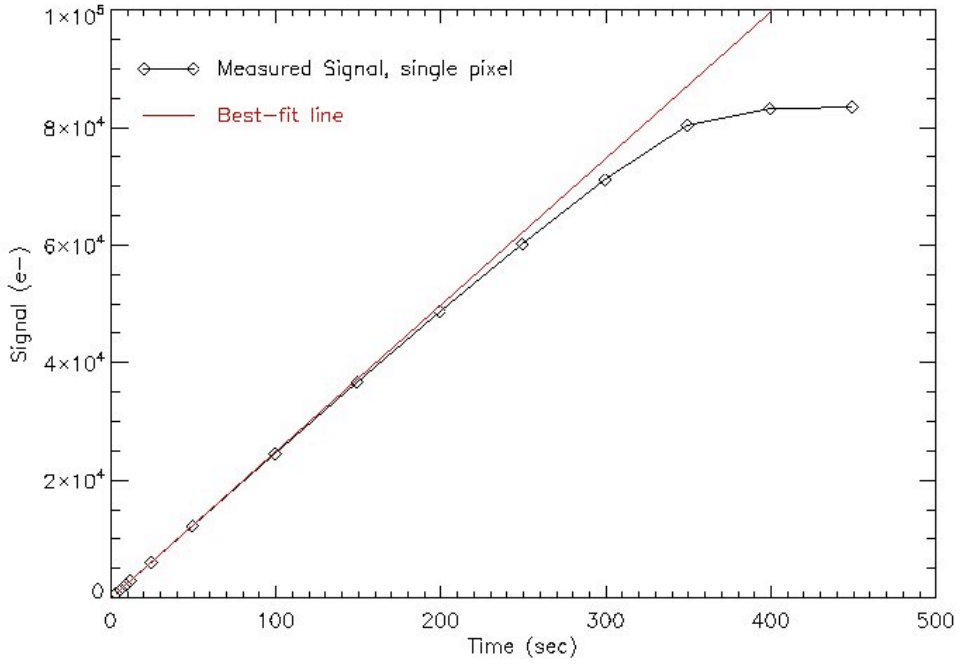

Figure 1: The measured signal (black) and best-fit line (red) for one pixel in the median STEP50 ramp. Points used in the line-fitting included only those up to a signal level of 4500 e-.

Once the linear fit for a given pixel was calculated, we divided the best-fit line by the measured signal for all reads up the ramp. As seen in Figure 2, this ratio (minus 1.0) was plotted against the measured signal up the ramp. At each signal level along the x axis, the red diamonds show the ratio values for a single pixel in the 12 individual STEP50 ramps. The median of these 12 ratio values was computed at each signal level, resulting in a median ratio versus signal curve.

We next fit a simple polynomial, shown in Equation 1, to the median ratio values.

$$y = A + Bx + Cx^2 + Dx^3 \qquad (1)$$

Here, x is the median measured signal, and y is the ratio. After the curve-fitting was complete, we had a set of coefficients, A through D, for each pixel. The blue stars in Figure 2 show the resulting best-fit curve for that pixel. Note that the red diamonds in Figure 2 are the original measurements from which the best-fit, blue stars are derived. At a given signal level, the spread in the red diamonds is several percentage points around the best-fit curve value. This large spread implies that ramp-to-ramp variations in the non-linearity behavior of a pixel will not be completely removed by a non-linearity correction composed of a single polynomial fit, and that the non-linearity of the IR channel will not be corrected to the CEI Specification 4.8.8 of <0.3% for all non-saturated signals.

Using these coefficients, non-linearity corrections could be performed on IR data using Equation 2, where s_{meas} is the measured signal, and s_{corr} is the corrected signal.

$$s_{corr} = s_{meas} * \left(1 + A + Bs_{meas} + Cs_{meas}^{2} + Ds_{meas}^{3}\right) \quad (2)$$

Data are corrected following the same technique as was used to calculate the A through D coefficients. The measured signal, s_{meas}, must include signal measured between the beginning of the ramp and the initial read. Therefore, the non-linearity correction must be applied prior to subtracting the initial read from subsequent reads in a ramp.

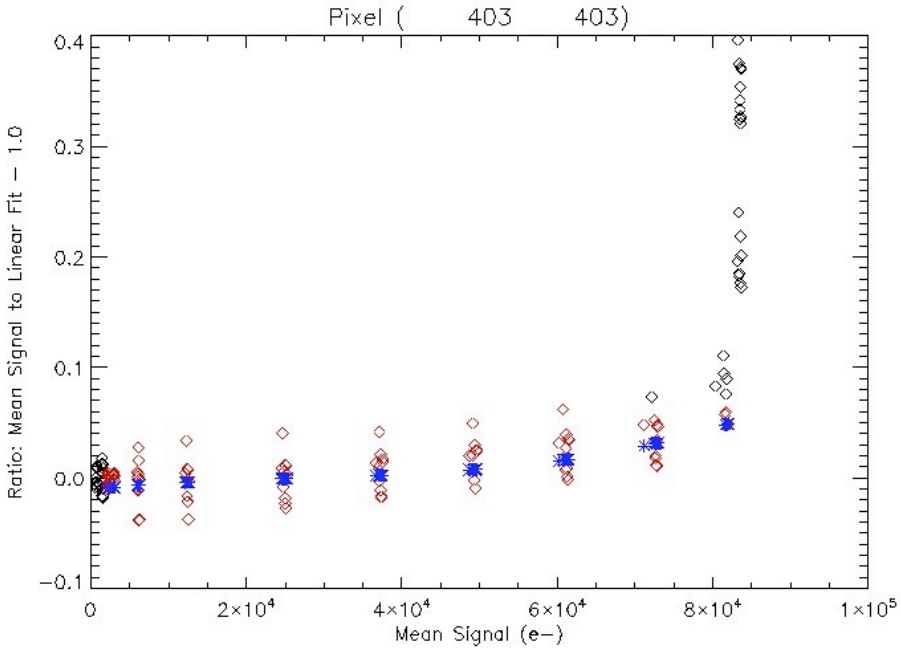

Figure 2: A plot of non-linearity percentage versus signal for a single pixel in 10 ramps. Diamonds show all measurements, red diamonds show measurements used in fitting Equation 1 (after taking the median at each signal level), and blue stars show the results of the curve-fitting. Ratios above 0.07 were not included in the fitting, due to the large change in the curve just above that point.

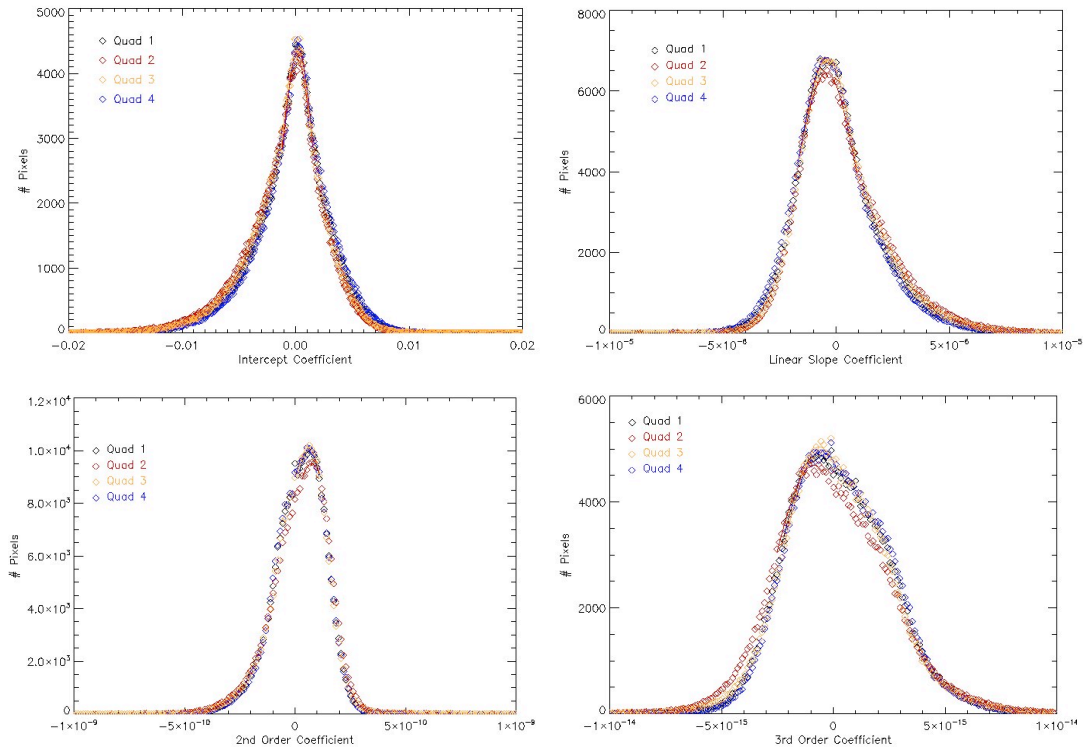

Figure 3: Histograms of the 4 non-linearity coefficients in each of the 4 quadrants.

Figure 3 shows histograms of the four coefficients from Equation 1 for all of the pixels in each of the 4 quadrants of the detector. There are only small differences in the coefficient distributions from quadrant to quadrant. The next question to be answered was whether the non-linearity correction would be better using a separate set of coefficients for each pixel, or by averaging the coefficients together and using a single set for each quadrant. In preparation for this investigation, we calculated quadrant-averaged coefficients from the distributions shown in Figure 3.

Table 1 shows the quadrant-averaged coefficients. The coefficient values listed in this table were calculated through Gaussian fits to the upper portions of the histograms in Figure 3. The peak values of the Gaussian fits were recorded as the mean values and placed in the first column, and the widths from the fits, divided by the square root of the number of pixels in a quadrant, were placed in the third column to show the uncertainty in the estimate of the mean.

The uncertainty values presented in the second column were derived from the uncertainties returned from the polynomial fitting routine (Equation 1). Similar to the coefficients seen in Figure 3, we created a histogram of the uncertainties for all pixels in each quadrant and for each of the four coefficients. The peak values of the Gaussian fits to these histograms are listed in the second column Table 1. These uncertainty values represent the mean uncertainty associated with the polynomial fit to a single pixel on the detector.

	Mean Coefficient	Mean Pixel-to-Pixel Uncertainty	Uncertainty in Estimate of the Mean
Quadrant 1			
A – intercept	2.5×10^{-4}	1.1×10^{-3}	2.7×10^{-6}
B – slope	-4.0×10^{-7}	5.9×10^{-7}	2.5×10^{-9}
C – 2nd Order	6.3×10^{-11}	5.6×10^{-11}	2.1×10^{-13}
D – 3rd Order	-7.3×10^{-16}	1.4×10^{-15}	3.3×10^{-18}
Quadrant 2			
A – intercept	1.3×10^{-4}	1.1×10^{-3}	2.7×10^{-6}
B – slope	-4.2×10^{-7}	6.1×10^{-7}	2.5×10^{-9}
C – 2nd Order	7.5×10^{-11}	5.9×10^{-11}	2.5×10^{-13}
D – 3rd Order	-8.9×10^{-16}	1.4×10^{-15}	3.5×10^{-18}
Quadrant 3			
A – intercept	1.1×10^{-4}	1.1×10^{-3}	2.9×10^{-6}
B – slope	-3.8×10^{-7}	5.8×10^{-7}	2.5×10^{-9}
C – 2nd Order	6.1×10^{-11}	5.5×10^{-11}	2.3×10^{-13}
D – 3rd Order	-6.3×10^{-16}	1.3×10^{-15}	3.1×10^{-18}
Quadrant 4			
A – intercept	2.3×10^{-4}	1.1×10^{-3}	2.9×10^{-6}
B – slope	-4.1×10^{-7}	5.8×10^{-7}	2.5×10^{-9}
C – 2nd Order	5.8×10^{-11}	5.5×10^{-11}	2.1×10^{-13}
D – 3rd Order	-5.3×10^{-16}	1.3×10^{-15}	3.5×10^{-18}

Table 1: Quadrant-wide non-linearity correction coefficients and uncertainties derived using TV3 data.

Comparison of Methods

Upon calculation of correction coefficients for each pixel on the detector, as well as the peak values of the histograms in each quadrant (referred to hereafter as the quadrant average, for brevity), we were faced with the question of which coefficients would produce the best non-linearity correction. Namely, should each pixel be corrected using its own coefficients, or could a quadrant-average be used? The motivation behind using quadrant-averaged coefficients is the reduction in uncertainties through the averaging of 250,000 pixels. The danger behind this method is that pixels with behaviors different from the quadrant-average may not be corrected appropriately.

To compare the efficacy of the individual versus quadrant-averaged correction coefficients, we compared the signal rates across the detector after correction with pixel-

by-pixel versus quadrant-averaged coefficients. We compared the signal rate up the ramp, rather than the absolute signal level, as the differences between the two methods are more pronounced. We calculate signal rates as the signal difference between consecutive reads, divided by the difference in exposure time. In this way, we can examine the changes in signal rate as the ramps progress.

For our comparison of non-linearity correction methods we focused first on individual pixels. The comparison was made on two types of pixels. First, we examined a pixel whose individual coefficients were close to the quadrant-averaged coefficients, where the two corrections should yield similar results. Figure 4 shows the measured signal rate before and after each of the two corrections. By the time the pixel reached saturation (at 300 seconds), the difference between the two corrected signal rates was almost 3%. The correction using the individual pixel coefficients appears to have slightly over corrected the signal rate, with the red line increasing with time between 100 and 250 seconds.

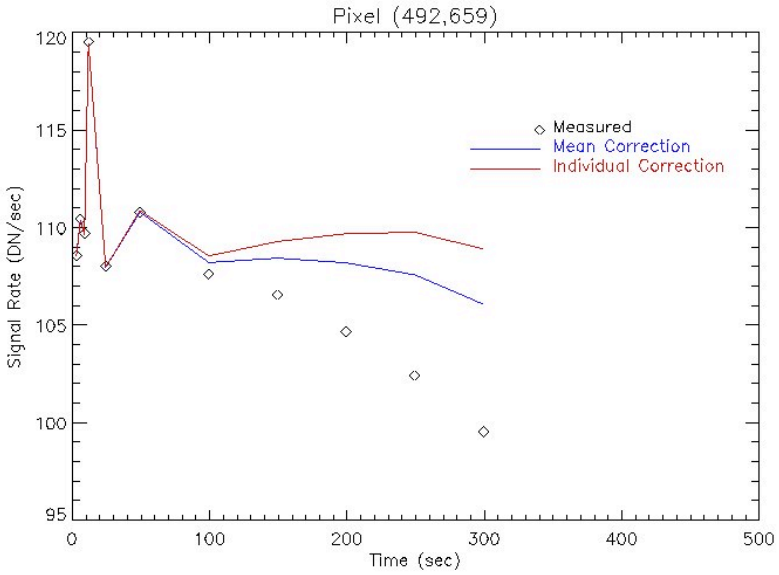

Figure 4: Non-linearity correction results for a single pixel. The correction using the quadrant-averaged coefficients maintains a more constant signal rate for the duration of the ramp.

Figure 5 is similar to Figure 4, but shows the corrected and uncorrected signal rate for a pixel whose individual 2nd and 3rd order coefficients are more than 1σ from the quadrant-wide value mean. In this case, both corrections produced a similar signal rate until 200 seconds into the observation. At this point, as the non-linearity caused the measured signal rate to quickly decrease, the quadrant-mean coefficients partially removed the non-linearity. However, the corrected signal rate still exhibited a decrease with increasing time towards the end of the ramp. Using the individual pixel coefficients,

the corrected signal rate remained at a level within 2% of the maximum rate in the original ramp. This argues in favor of using the non-linearity coefficients calculated on a pixel-by-pixel basis. Roughly 8% of the light sensitive pixels fall into this category, where their individual 2nd and 3rd order coefficients are both more than 1σ from the quadrant-wide mean.

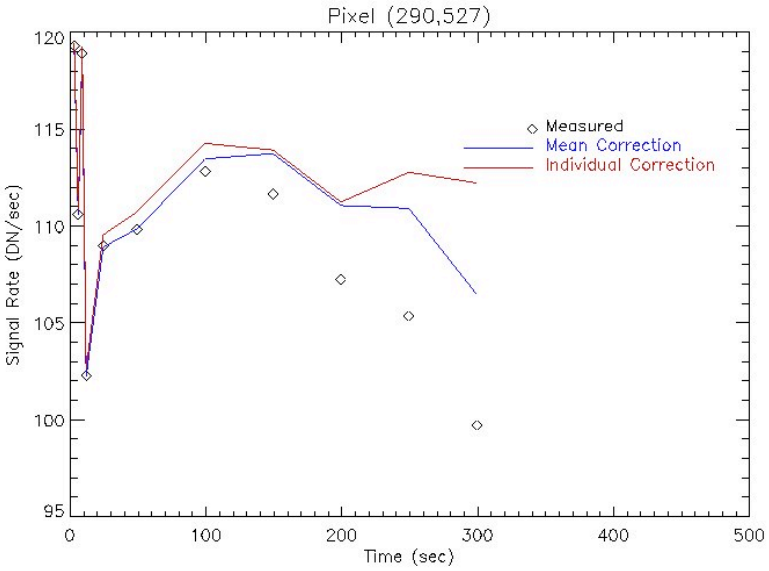

Figure 5: Comparison of correction technique on a pixel whose individual coefficients are far from the quadrant-averaged coefficients.

In the end, most WFC3 users will be interested in the image created by the final signal rate values for all pixels. Therefore, it is useful to take both versions of the corrected ramps, fit a line up the ramp, and compare the resulting best-fit slopes (which give a measure of the ramp average signal rate, rather than the instantaneous rates shown in Figures 4 and 5). For the pixel displayed in Figure 4, the best-fit slopes of the two corrected ramps differed by only 0.8 DN/sec, or 0.7%, while the formal error on each best-fit slope was 0.4 DN/sec. For the pixel shown in Figure 5, the difference in best-fit slopes was 1.0 DN/sec (0.9%), with the same formal error of 0.4 DN/sec. The small magnitude of these errors, combined with the fact that at most 8% of the light-sensitive pixels are worse than the pixel in Figure 5, implies that for most pixels, the use of quadrant-averaged non-linearity correction coefficients will introduce errors smaller than other sources of photometric error.

We next performed a similar comparison using gain data taken during TV3 testing. These data comprise a set of CASTLE flat fields taken with the RAPID sample sequence. As with the non-linearity data, we corrected the gain data using the individual coefficients, as well as the quadrant-averaged coefficients, and then compared the results. We then examined the signal rate in the same pixels as in Figures 4 and 5 after each type

of correction. Figures 6 and 7 show plots similar to Figures 4 and 5. The results are similar to those obtained with the non-linearity data. Figure 6 shows that for the pixel with the individual coefficients very close to the quadrant-wide mean coefficients, the difference between the two corrections is small. Best-fit lines to the two corrected signal ramps reveal rates that differ by only 0.8% (compared to a formal error on each best-fit slope of 0.4%), similar to that in the non-linearity data in Figure 3. The two correction methods also produce very similar results in Figure 7, where the individual correction coefficients are relatively different than the quadrant averages. In this case, best-fit slopes to the two corrected ramps differ by only 0.5% (with a formal error on the fitted slopes of 0.3%). Again, this implies that other sources of photometric errors will dominate over the errors associated with the non-linearity correction in the IR channel.

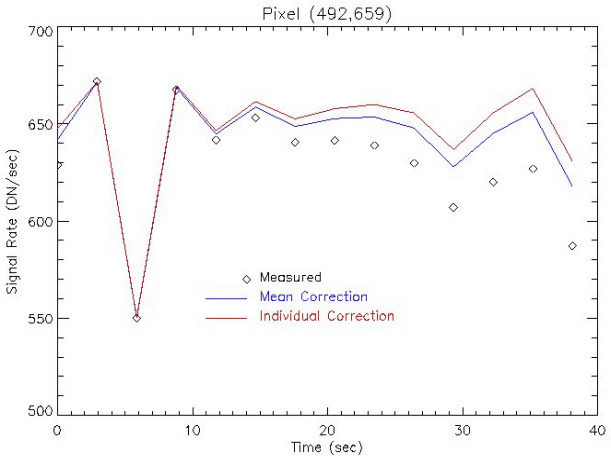

Figure 6: Comparison of the two correction methods performed on the same pixel as Figure 4, using a flat field image from TV3 testing.

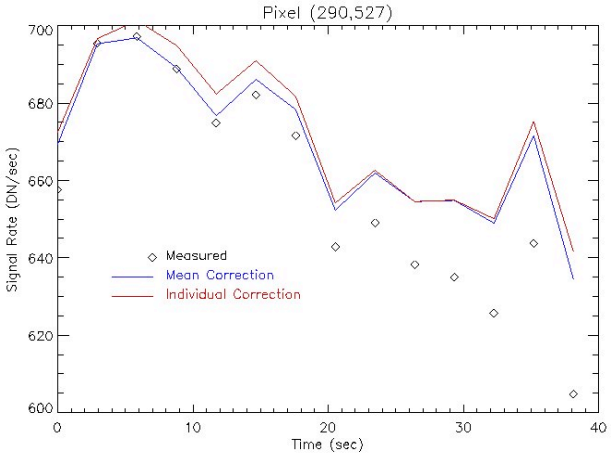

Figure 7: Comparison of the two correction methods performed on the same pixel as Figure 5, using a flat field image from TV3 testing.

We next focused on the entire detector, in order to collect statistics on the overall difference between the two non-linearity correction methods. We used two ramps for our detector-wide comparison of the non-linearity corrections. First, a synthetic ramp was created, with all pixels having identical signal in each read. In this way, we could isolate the effects of the non-linearity correction at a wide range of constant signal levels. Our synthetic ramp was composed of 8 reads, where the signal for all pixels through the reads was 100, 500, 1000, 5000, 10,000, 20,000, 25,000, and 30,000 DN. We applied the non-linearity correction on a pixel-by-pixel basis, and then fit lines up the ramp in order to create a signal rate image. We then repeated this process, but made corrections on a quadrant-by-quadrant basis. To evaluate the difference between the two resulting signal rate images, we took the ratio of one final image to the other, and created histograms of the result. Figure 8 shows that most pixels have similar values in the two final images, but that the peaks of the histograms are not unity, as one would expect. The 1σ widths of the histograms are 1%. Depending on the quadrant, 1.7% to 2.8% of the total light-sensitive pixels have final image values that are more than 2% different between the two methods.

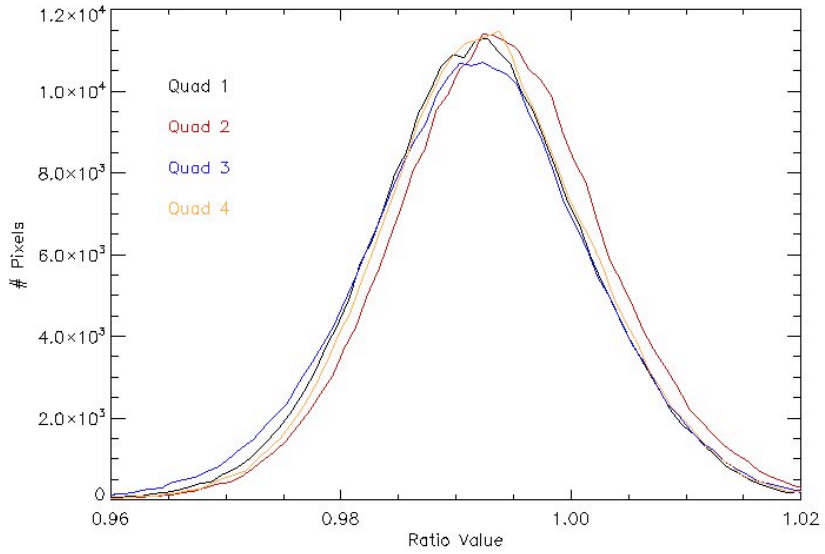

Figure 8: Histograms of the ratio of final images created using quadrant-averaged versus pixel-by-pixel non-linearity correction. Note the peaks of the histograms are shifted from unity.

The next step was to make the comparison described above on a real WFC3 ramp rather than synthetic data. We used two flat field ramps from TV3 testing (taken through the F125W filter), and performed the same corrections and image division described

above. The resulting histograms are shown in Figure 9, and are very similar to those created from the synthetic data. The results are similar for both flat fields. Note that the peaks of the histograms differ from unity by roughly 0.75%, in agreement with the results from the synthetic data presented above. Again, depending on quadrant, 1.6% to 2.4% of the light-sensitive pixels are more than 2% different between the two correction methods.

Uncertainties associated with the line-fitting in the creation of the final images were roughly 0.3%, which is a significant fraction of the offset of the peaks from unity.

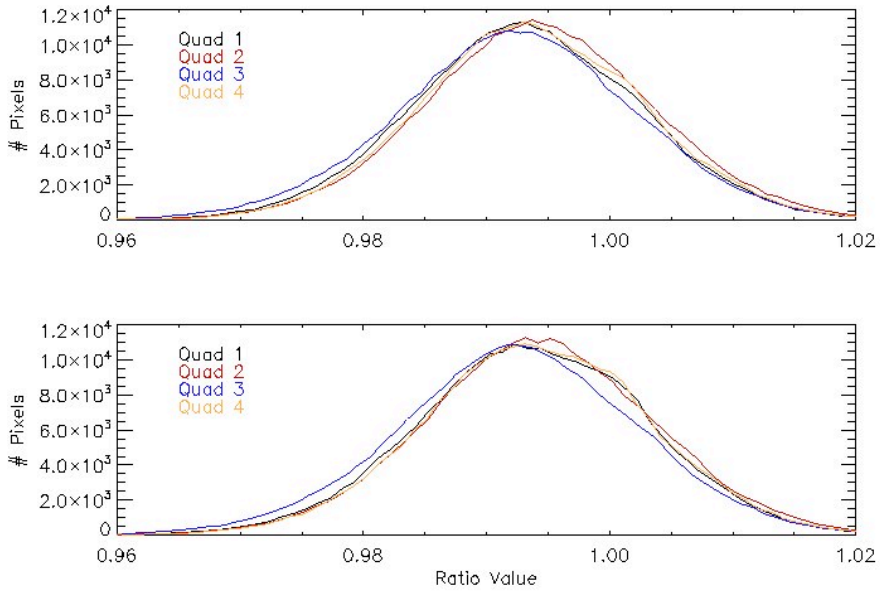

Figure 9: Same as Figure 8, but derived from two flat field ramps taken during TV3 testing. Each panel shows the ratio of the quadrant-by-quadrant corrected final image to the pixel-by-pixel corrected final image for one ramp.

A map of the pixels where the difference between the two corrections is more than 2% (equivalent to 2σ, and shown in Figure 10) reveals a salt and pepper pattern spread across the detector. This implies that for an astronomical point source randomly placed on the detector, the photometry results corrected quad-by-quad versus pixel-by-pixel should differ by less than 2%, as the aperture may contain some pixels with a large difference between the two non-linearity corrections, but will be dominated by those with small differences. At this level of accuracy, the non-linearity correction will be comparable to other sources of photometric error.

The only exception to this is in quadrant 3, where almost all pixels in the "wagon wheel" area of the detector are more than 2% different between the corrections. These pixels are known to exhibit odd behaviors that conspire to increase their uncertainties.

Figure 10: Map of the IR detector. White pixels have a 2% or greater difference between the two non-linearity correction methods. The "wagon wheel" is the feature in the lower right corner.

Saturation Levels

The final task associated with the derivation of the non-linearity correction was the creation of a saturation map of the detector. For each pixel, the saturation map lists the measured signal level at which the uncorrected signal deviates from linearity by 5%. Above this level, a pixel will be labeled as saturated and the measured signal will be ignored. The saturation level for each pixel was calculated through interpolation of the ratio of best-fit line to the measured signal versus measured signal plot, such as in Figure 2. Figure 11 shows a histogram of the calculated saturation levels for all pixels. Due to the large spread in these values, we thought it more appropriate to quote an individual saturation value for each pixel, rather than inserting the quadrant-averaged coefficients into Equation 1 and solving for a 5% non-linearity. The latter situation would lead to cases of low saturation pixels being corrected, rather than marked as saturated, possibly affecting later photometry.

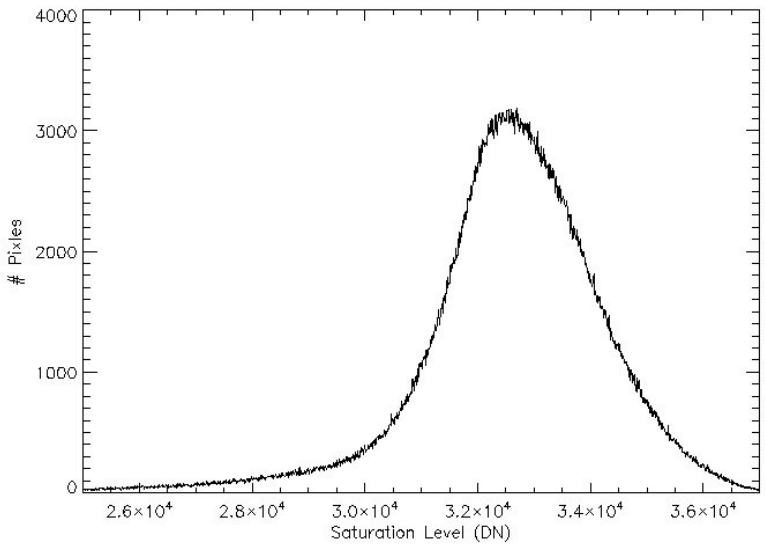

Figure 11: Histogram of the saturation levels of all pixels in the IR detector.

CEI Specification 4.8.8 states that the response of the IR detector shall be linear to <5% (and correctable to <0.3%) up to a level of 70,000e^- (~28,000 DN). Using the 5% non-linearity level as the definition of a pixel's full well, (rather than the 40,000DN specified in CEI Specification 4.8.7), we find that only 2.6% of the light-sensitive pixels have a non-linearity larger than 5% at signal levels below 28,000DN. Another way to state this is that only 2.6% of the light sensitive pixels fall below the CEI Specification for full well.

Conclusions

Using quadrant-averaged coefficients corrects the non-linear response of the IR channel better than using separate coefficients for individual pixels due to decreased noise in the curve-fitting process. For most pixels, differences between the two methods result in final signal rates for a given pixel that are identical to within photometric errors from other sources. This correction will be incorporated into CALWF3, the standard data reduction pipeline for WFC3 data, and applied to all WFC3-IR data ramps.

Given the large spread in non-linearity measured at all signal levels in nominally identical ramps (as seen in the vertical spread of the red diamonds in Figure 2), we cannot produce a set of correction coefficients that will correct the non-linearity of all IR ramps to less than the CEI Spec value of 0.3%. Using a set of multiple ramps, we can only create coefficients capable of correcting a pixel's average non-linear behavior. Therefore, the IR channel does not meet the non-linearity correction portion of the CEI Specification 4.8.8. With this correction method, the signal in the IR detector cannot be corrected to less than 0.3%.

References

Hilbert, B., 2007, WFC3 ISR 2007-25, "***WFC3 TV2 Testing: IR Channel Non-linearity Correction and Unstable Pixel Masking***"

Hilbert, B., 2004a, WFC3 ISR 2004-06, "***Non-Linearity Correction Algorithm for the WFC3 IR* Channel**"

Hilbert, B., 2004b, WFC3 ISR 2004-10, "***Basic IDL Data Reduction Algorithm for WFC3 IR and UVIS Channel***"

Robberto, M. and B. Hilbert, 2005, WFC3 ISR 2005-29, "***WFC3 2004 Thermal Vacuum Campaign: IR channel linearity (flat field illumination - SMS IR04)***"

Appendix

TV3 Data Collection

Non-linearity data during TV3 testing was collected using Science Mission Specification (SMS) IR04S01. Gain data used in the testing of the non-linearity correction were collected using SMS IR02S01.