The behaviour of the WFC3 UVIS and IR Analog-to-Digital Converters

M. Robberto & B. Hilbert
September 22, 2005

ABSTRACT

Using UVIS and IR data from the 2004 thermal vacuum tests, we have searched for anomalous behavior in the WFC3's analog-to-digital converters (ADCs). We find that the ADCs of the CCD channel operate nominally, except for the less significant bits which favor 1's over 0's by up to 5%, depending on the ADC and on the input voltage. The IR channel, clocked at double speed, also shows departures in the least significant bits at ~1% level. These departures from ideal behaviour are statistically significant. We discuss the relevance of these anomalies for the scientific quality of the WFC3 data, simulating images affected by biased ADCs at 1% and 2% level. Whereas the occasional anomalous bits corresponding to signal levels around the cosmic ray filtering threshold may represent a source of salt-pepper noise hard to eliminate, for both the UVIS and the IR channels the added readnoise is negligible and does not affect the ultimate sensitivity of the instrument. The advantages of clocking the CCD detectors at double speed to reduce the detector cross talk largely supersede the disadvantages of the measured extra ADC noise.

1 Introduction

Analog-to-Digital converters (ADCs) are essential components of the readout electronics of both CCDs and IR detectors, as they convert the analog voltage from each pixel into a digital signal. WFC3 uses for both the UVIS and the IR channel the same type of 16-bit ADCs, meaning the analog voltage is converted into 16-bit binary number. The four ADCs of the CCD readout chain (two for each detector, one for each readout channel) are

sampling at about 50KHz, whereas the four ADCs of the IR readout chain (one for each quadrant of the array) are sampling at ~100KHz. In order to eliminate, or significantly reduce, the electronic cross-talk between readout channels of the CCD, it has been proposed to clock the relative ADCs at double speed. However, if the performance of the ADCs degrade with the clock speed, the excellent readout noise performance of the CCD may deteriorate. The purpose of this study is therefore to compare the behavior of the WFC3 ADCs clocked at different speeds.

2 Data

Data files used for this study included images from the IR13 and UV05 thermal vacuum testing procedures. They are flat field images taken with different illumination and exposure times. As seen in Figures 1 (for the CCD) and 4 (for the IR), data with relatively high median signals were most useful, as one can sample several bits of noise. Both the IR and UVIS arrays were read out at the nominal clocking rates. For the IR this is 90.9 KHz, while the UVIS rate is 45.4 KHz.

3 Analysis

The signal of each pixel is originally recorded in 16-bit unsigned integer format. In order to search for anomalous behavior in the ADCs, the flat field data were first converted into binary form, yielding a 16-bit binary number for each pixel. Each of the 16 bits was then examined separately.

Across each quadrant of the detector, the fraction of pixels with a '0' value for a given bit was compared to the fraction with a value of '1'. Ideally, for bits measuring white noise on top of a larger signal, there should be equal numbers of 0's and 1's. Any deviation from a 50/50 split of 0's and 1's in the noise dominated bits would imply a problem in the ADC for that bit.

4 Results

4.1 UVIS

Figure 1, obtained from the file iu05a112r_04274203456_raw.fits, shows the distribution of values for the 16 bits of quadrant 1 of the UVIS detector. The median signal level of the exposure is 15,186 DN. This value is coded in Figure 1: the five most significant bits (0 to 4) are fixed at a constant value. The three bits fixed at 1 contribute values of 2^{13}, 2^{12}, and 2^{11} (8,192, 4,096, and 2,048 = 14,336 DN total) respectively to the final measured signal. The fact that these bits remain fixed with 100% of values either 1 or 0 gives us confidence that the ADCs at these signal levels are well behaving. Bits 5,6 and in some measure also 7 fluctuate in a basically anticorrelated way: when bit 5 is 1, contributing

2^{10}=1024 DN, bit 6 is 0 and vice versa Bits 8 through 15 represent small perturbations (2^7 through 2^0 DN) on top of the measured signal. Our assumption is that these bits sample the noise, dominated by photon shot noise, and ideally should have values of 0 and 1 in equal proportions.

Figure 1, as well as similar plots for other flat fields that we do not present, show that all of the UVIS ADCs behave in the same manner. Bits 7 to 14 are almost exactly at 50% level, whereas bit 15 is biased. In this quadrant it displays a slight preference to favour a value of 1 over 0.

The calculations yielding Figure 1 were repeated on each of the 49 files in the UV05 test, taken with different illumination. In Figure 2 we show, for each of the 49 files, the fraction of 1 in bit 15 versus the median signal level. Quadrants 1 and 2 show a systematic trend: as the median signal increases, the probability of bit 15 having a value of 1 decreases. Low signal files show probabilities as high as 55%, while high signal files have probabilities as low as 48%. The skewed distribution of 1's and 0's appears limited to these 2 quadrants.

The test was repeated on the next less significant bit (14), in order to study the extent of the non-nominal behavior. As shown in Figure 3, bit 14 still shows a marginal preference, about 51%, of 1 over 0..
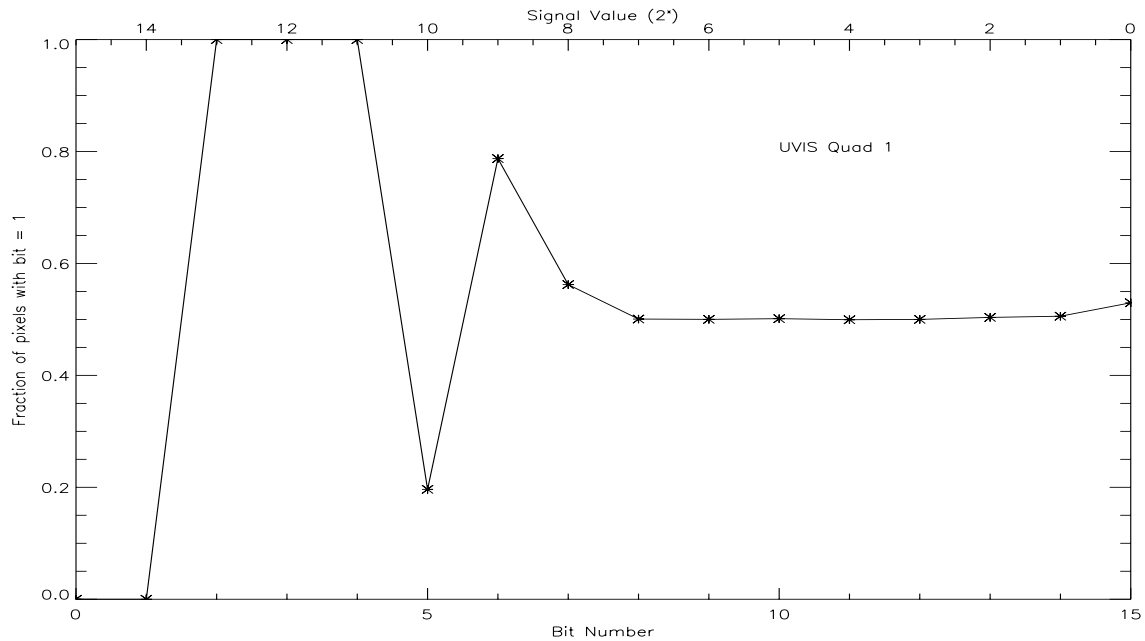

Figure 1: The fraction of bits with a value of 1 for each of the 16 bits in quadrant 1 of a UVIS flatfield. The median signal level in the quadrant is 15,186 DN, explaining the universal bit value of 1 for bits 2 through 4. Bits 8 through 15 represent small numbers and therefore are dominated by noise. As expected, these bit values are split nearly evenly between 0's and 1's
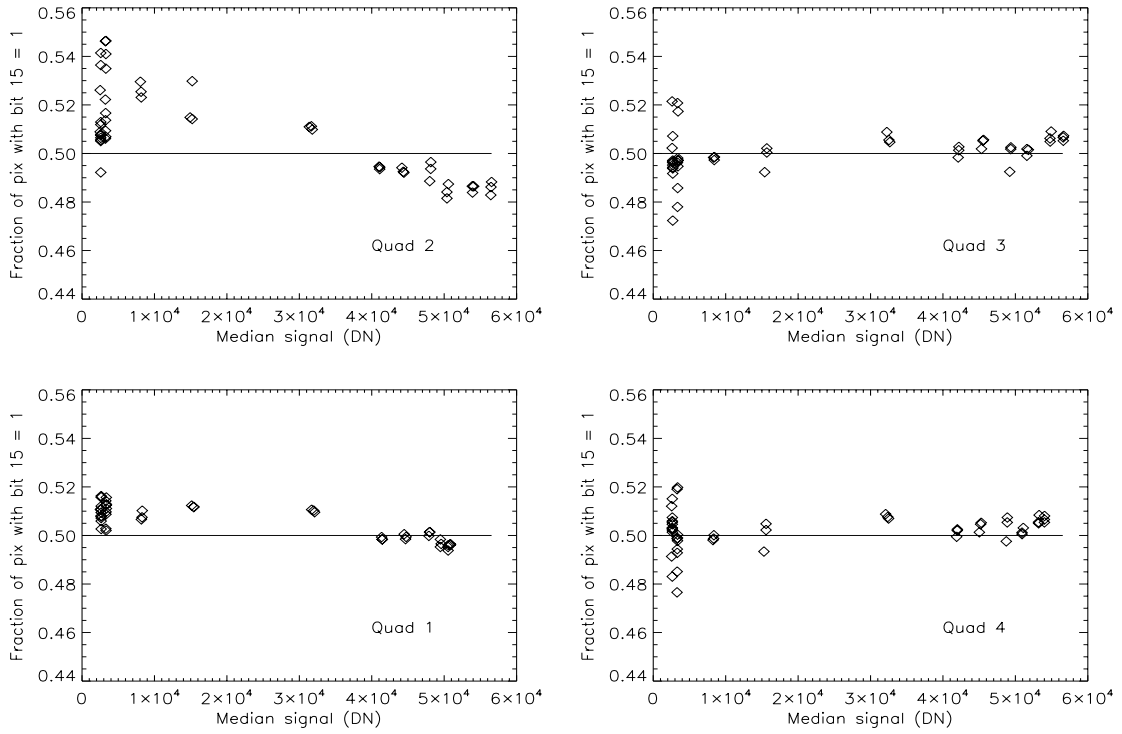
Figure 2: The frequency of 1 in the UVIS bit 15, plotted against median signal level.
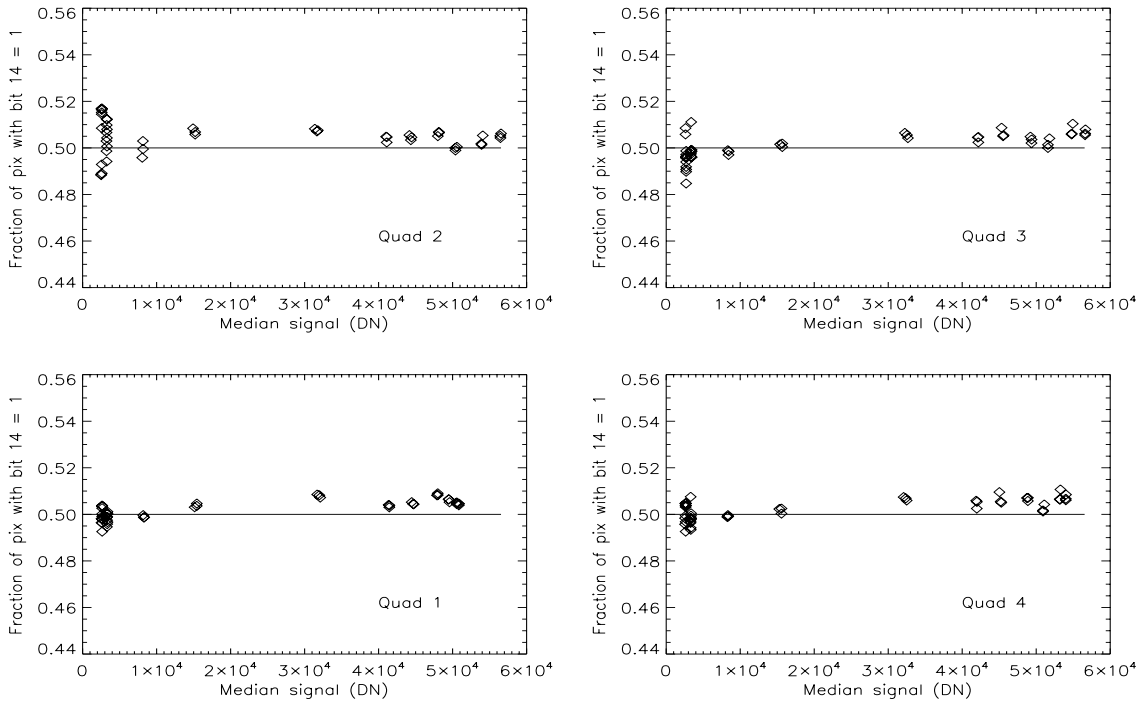

Figure 3: Same as Figure 2 for the UVIS bit 14

4.2 IR

Figure 4, the IR equivalent of Figure 1, shows the frequency of the 16 bits in one quadrant of the IR channel. Here the median signal value was 31,883 DN. The main signal level is set by bits 0 through 4 (signal contributions of 2^{15} through 2^{11} DN), with bits 5 through 15 measuring noise on top of the main signal. As with the UVIS channel, the bits which sample the noise have values that are split evenly between 0's and 1's.

However, tracking the behavior of bit 15 for the different illumination levels of the IR13 test reveals a more complex pattern (Figure 5). Bit 15 takes the value of 1 between 50% and 52% of the times, depending on the input voltage. There is a cyclical variation that seems to extend into the bits 14 and 13 as well, although with a smaller amplitude than in bit 15. This is shown in particular in Figure 6 for bit 13, which seems to favor 1's over 0's by up to 51%. The ADCs start showing the expected behavior, with equal fractions of 1's and 0's above (i.e. for all bits more significant than) bit 13. Bit 12 statistics are shown in Figure 7.
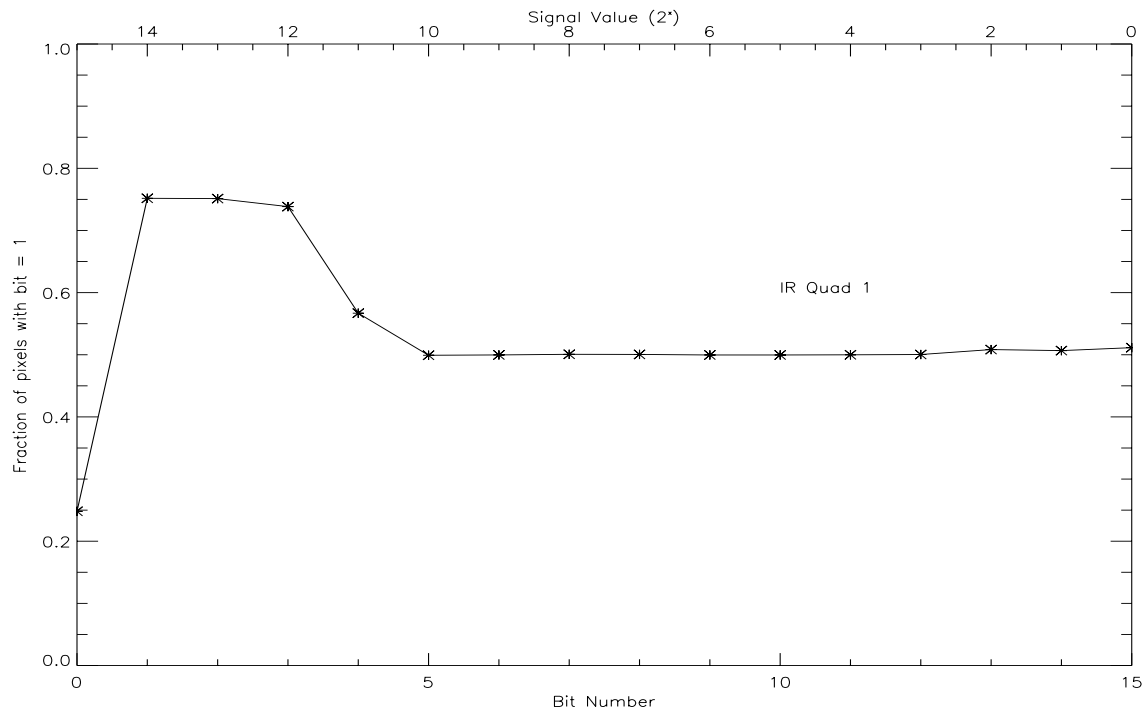

Figure 4: The same plot as shown in Figure 1, for the IR detector. Bits 5 through 15 sample the noise, and have values of 1 in half of the total pixels, as expected.
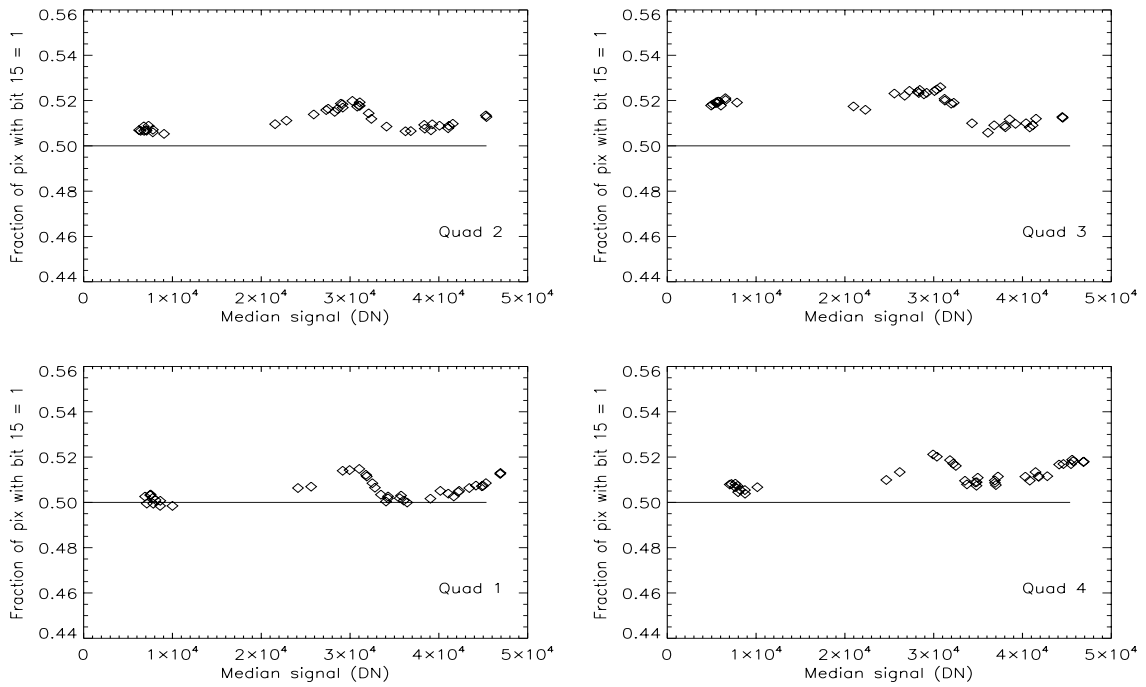
Figure 5: IR bit 15 behaviour. The cyclical values for fraction of 1's is unexpected, and persists in all areas of the detector.
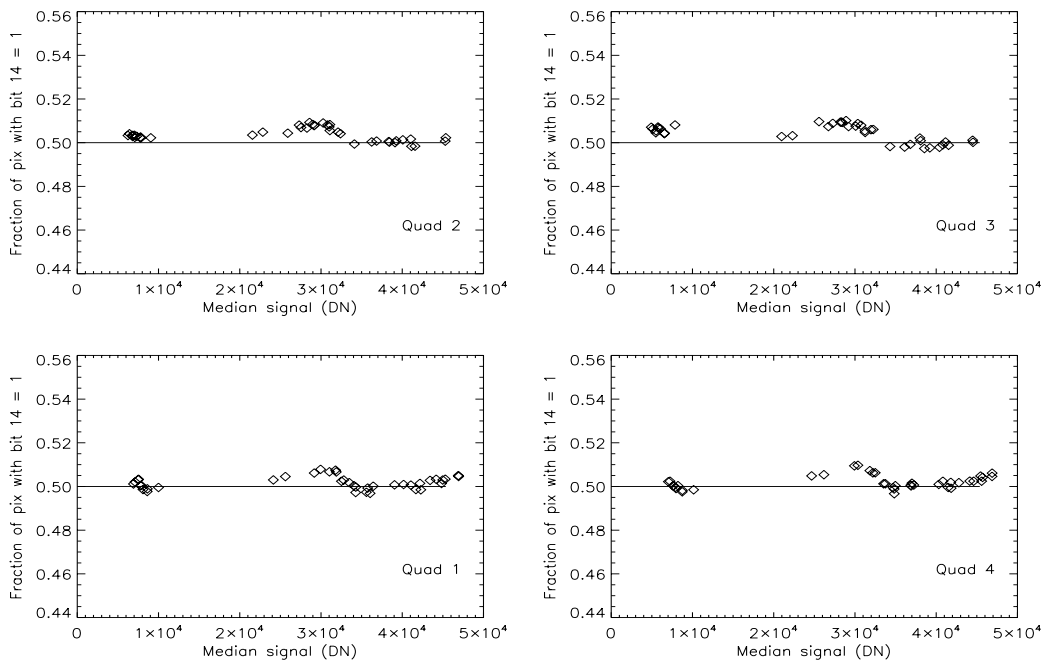

Figure 6: IR bit 14 behaviour. The cyclical variations have a slightly smaller amplitude than in bit 15.
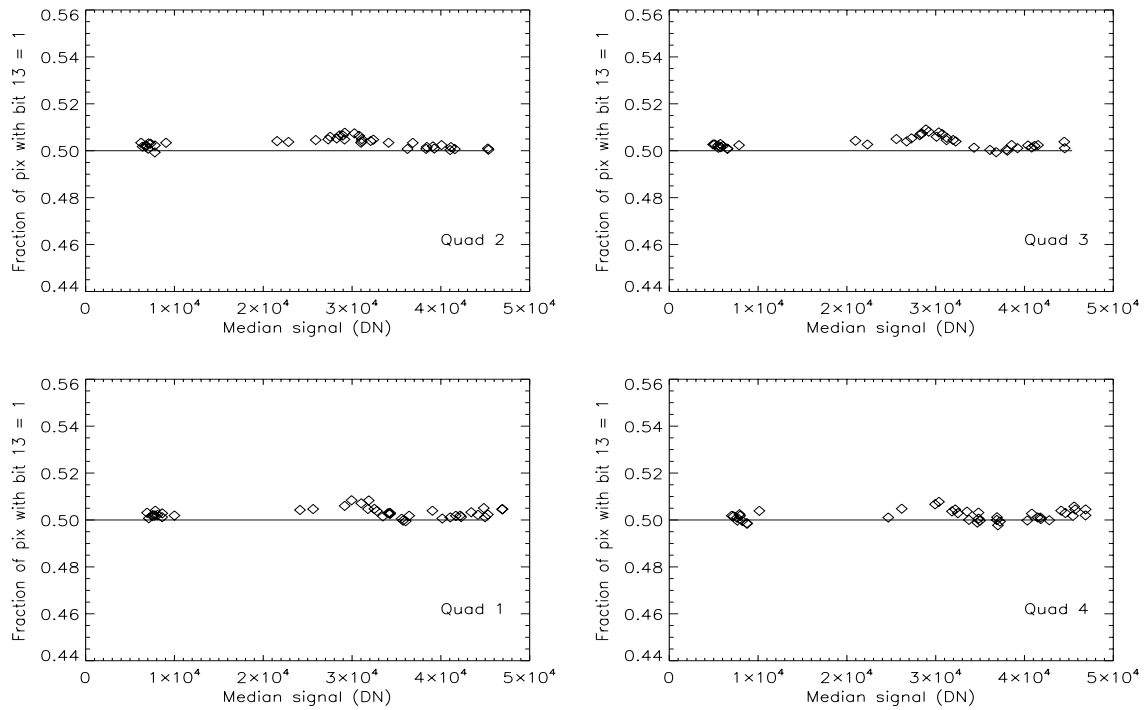
Figure 7: IR bit 13 behaviour. Here, the cyclical variations have a smaller amplitude than in bit 14.
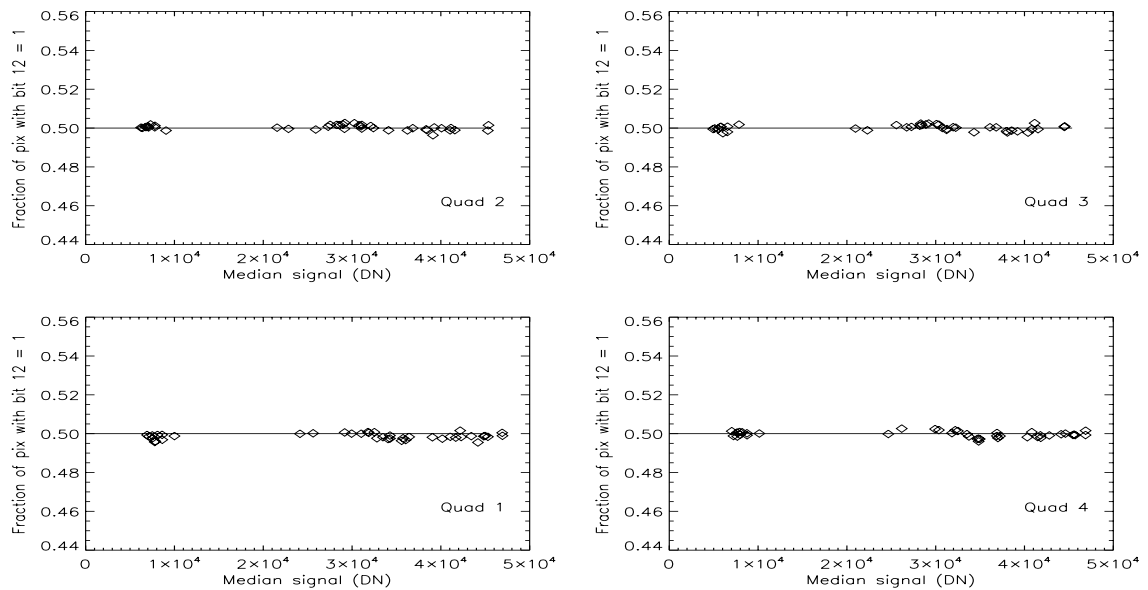

Figure 8: IR bit 12 behaviour. The cyclical variation in the fraction of pixels with values of 1 has damped out, leaving half the pixels with a value of 1, and half with a value of 0, as expected.

5 Statistical properties of an ideal ADC

5.1 ADCs and faked coins

The statistical behaviour of each bit of an ideal ADC can be compared to the problem of tossing a coin a large number of times. If a given bit is sampling pure noise signal, then the frequencies of 1 and 0 must be nearly identical, just as when tossing a coin a large number of times one would expect to obtain an almost equal number of heads and tails. This is under the assumption that the events are randomly distributed and independent, i.e. the ADC is sampling pure white noise or the coin is not flaky[1].

In the case of an ADC, our frames provide us with millions of samples, and therefore the problem can be likened to the classic one of tossing a coin a nearly infinite number of times. One wants to know how close the observed average must be to 0.5 in order to accept the null hypothesis of having a fair coin, or a perfect ADC. For this one can use elementary statistics. Since the single event follows the binomial probability distribution with 0.5 probability, for a large number of events the limit theorem known as the De Moivre-Laplace theorem can be applied. The De Moivre-Laplace theorem states that given a binomial distribution with success probability p, the probability of having a number i of successes in n trials is given by the reduced gaussian distribution of the variable $x=sh$ where:

$$h = \frac{1}{\sqrt{np(1-p)}}$$

For the CCD we have 2Kpix x 2Kpix per channel. A flat field image therefore contains $n=4M$ samples and therefore $h=1K$. For the IR detector a quadrant contains 512x512 pixels (including the reference pixels), i.e. $n=256K$ and $h=256$. The gaussian distribution of s has therefore standard deviation $s=1/h$, narrowly peaked around 0.5: in fact, for the CCD channel the 3.3 sigma level (99.90% confidence) of correct ADC behaviour corresponds to frequencies in the range 0.5+-0.003, whereas for the IR channel the range is 0.5+-0.012. A fraction of 1's outside of this range leads to the nearly certain rejection of the null hypothesis that the particular bit of the ADC is behaving ideally.

5.2 How much faked?

If a pixel is not ideal it will be biased toward another average, giving e.g. a 1 value 51% of the time. One may want to estimate the impact of this bias on the image, possibly in the

---

1.   Persi Diaconis and coworkers at Stanford University have recently demonstrated that real coins are biased and tend to land on the same face they started out on (paper available at http://stat.stanford.edu/~cgates/PERSI/papers/headswithJ.pdf)

perspective of correcting for it. In this case, however, whereas we know that the pixel is not ideal, we do not know a priori its real behaviour: what we have is just a single estimate of its average response. We want to evaluate the distribution of an unknown parameter having in our hands only a sample. This is the type of problems approached by Bayesian techniques. [1]

The Bayesian probability theory is used to calculate the posterior probability density function, where posterior refers to the fact that we already know the outcome of an experiment. In our case, the number i of 1's we have measured is our best estimator of the unknown real value I, and the same for the number $o(=n-i)$ of 0's for the corresponding true value $O(=n-I)$. The theorem states that the probability r of obtaining 1 in a single random read follows the posterior probability density function f given by:

$$f(r|I = i, O = o) = \frac{L(I = i|r, N = i + o)P(r)}{\int_0^1 L(I = i|r, N = i + o)P(r)dr}$$

This theorem therefore states that the probability $f(r/I=i,O=o)$ of obtaining r successes having performed an experiment that has given i successes and o failures is given by the product of the likelihood $L(I-i/r,N=i+o)$ of extracting, during our experiment with $N=i+o$ samples, i successes out of the original distribution of unknown parameter r, times the prior probability P(r) of having that particular r value. Everything is, of course, properly normalized to give total probability equal to 1. We can assume that, in the absence of any observation, $P(r)$ is uniform over the interval $[0,1]$: $P(r)=1$. In what concerns the Likelihood function $L(r)$, it is simply given by the binomial distribution:

$$L(r|N = i + o) = \binom{N}{i} r^i (1 - r)^o$$

We have therefore

$$f(r|I = i, O = o) = \frac{r^i (1 - r)^o}{\int_0^1 r^i (1 - r)^o dr}$$

---

1. Admittedly, this section gives us a pretext to refresh some non-elementary statistical techniques.

The integral in the denominator is the definition of the Beta function B(ι+1,o+1). B is related to the Gamma function Γ by the relation

$$B(i, o) = \frac{\Gamma(i)\Gamma(o)}{\Gamma(i + o)}$$

For any positive integer value of the parameter λ, it is Γ(λ)=(λ−1)!, therefore

$$f(r|I = i, O = o) = \frac{(i + o + 1)!}{i!o!} r^i (1 - r)^o$$

The probability density function can be calculated for asymptotically large values of n,i,o by using the Stirling's formula:

$$\lim_{n \to \infty} n! = n^n e^{-n} \sqrt{2\pi n}$$

and using natural logarithms. It is therefore

$$f(r|I = i, O = o) = (i + o + 1)\ln(i + o + 1) + 1 + \frac{1}{2}\ln(2\pi i + o + 1)$$
$$- i\ln i - \frac{1}{2}\ln(2\pi i) + o\ln o - \frac{1}{2}\ln(2\pi o) + i\ln(r) + o\ln(1 - r)$$

This expression can be integrated between any range of r values, e.g. between 0.509 and 0.510, to get the posterior probability of finding r within that range. For example, if we have for a CCD channel n=2048^2 samples with 2,139,095 times 1 (=51%), the probability of having an ADC with average value between 0.5095 and 0.5105 is given by:

$$\int_{0.5095}^{0.5105} f(r|I = 2139095, O = 2055203) = 0.959$$

corresponding to nearly a 2σ level of confidence.

6 Discussion

The previous section has shown that the range of variability of the average of a random signal sampled by an ideal ADC reading a quadrant of our detectors is extremely tight, of the order of a few parts per thousand. It is clear, therefore, that the ~1% asymmetry shown

by the least significant bits of the ADCs in both the UVIS and IR channels of WFC3 indicate non-ideal behaviour.

Two out of four ADCs of the UVIS channel show a trend with signal level in the least significant bit, i.e. a preference for bit values of 1 at low signal levels shifting to values of 0 at higher signal levels. The last bit of the other two ADCs gives 1 more than 50% of the time at high signal levels, i.e. opposite to the other two, but no clear evidence of a trend with signal. For all four channels the next least significant bit behaves similarly to these last two. All other bits appear to behave properly.

Also the IR channel shows an irregular behavior of the ADCs. All 4 quadrants of the IR channel display an asymmetry in the response of bits 13 through 15, again with a preference of 1's vs. 0's. The amplitude of this variation is as large as ~2% for the last significant bit at intermediate flux levels (Fig.5), and can be traced at 1% level up to bit 13. As with the UVIS channel, the ADCs appear to function nominally in the bits representing signals comparable to the readnoise.

In general, the 2 or 3 least significant bits of the ADCs can be regarded as (slightly) biased estimators of the true values. The worse performance of the ADCs of the IR channel vs. those of the CCD channel at bit 13 may reflect a degradation of the ADC performance at higher clock speed.

Our theoretical analysis shows that the discrepancies are very well measured. One may wonder how these departures from the ideal behaviour may impact the scientific quality of the data. If bit n (starting from the least significant bit) has 51% average response instead of 50%, then there is 1 pixel out of 100 that has value higher by 2^{n-1}. If it is one of the most significant bits, the error produces a spurious isolated hot pixel which should be masked out by cosmic ray subtraction. This, however, is apparently not our case (see Section 4.1), as we are dealing with the least significant bits close to, or below, the noise floor of the detector. To investigate the effect on astronomical data we have run a simulation.

Effect on astronomical images

The behaviour of the CCD detector with ideal ADCs can be treated by assuming a normally distributed population with mean=0 and standard deviation=2. This is the optimal noise level of the UVIS channel, measured at 2DN (3 e$^-$, Hilbert et. al. 2005). We assume zero sky noise and that calibration files do not introduce extra noise. Then we add a 1% or 2% of "poisoned" pixels to generate the fraction of the pixels with a biased bit. They have a value higher by 1 DN (bit 15), 2 DN (bit14) or 4 DN (bit 13), equally distributed among them (i.e. we neglect pixels having more than one biased ADC at the same time). The histogram of the distribution before and after adding the poisoned pixels (Figure 9) clearly shows that the distribution becomes skewed to higher values. The bins to the left of the peak bin (up to 0 DN) lose more pixels than they get from the lower bins, and viceversa for the bins on the right side of the peak. The mean of the distributions, in our test case with 1,000,000 values, increases from 0.000405 for the original distribution to 0.0373 with 1%

of the pixels containing a biased bit. The standard deviation also increases with values of 2.000(0%), 2.049 (1%) and 2.095 (2%). There is therefore a small, negligible deviation from ideal performance. The same applies to the IR channel, where the higher readout noise (approximately 21.5 e⁻, or 8.6 DN at 2.5e⁻/DN gain, Hilbert, 2005) makes the variation even more negligible.
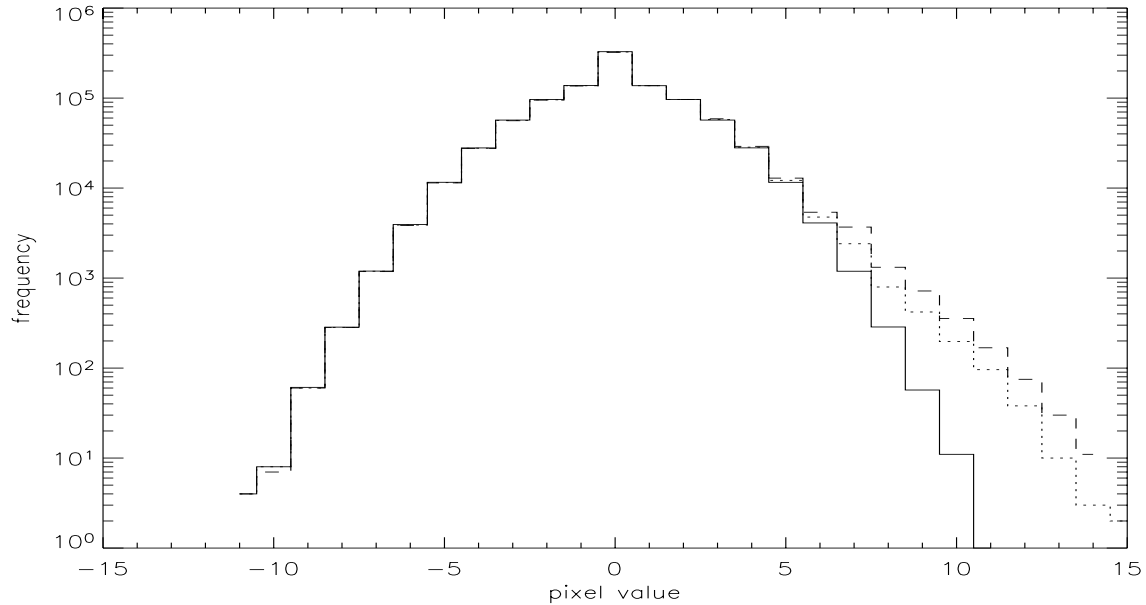

Figure 9: Histograms of the CCD pixels both with, and without digitization noise added.

In general, the bits that represent the highest danger are those that offset the signal by an amount comparable to the cosmic ray rejection threshold. Major outliers (i.e. anomalous response from the most significant bits) can be easily masked out and lower outliers are in the noise. But bits at around 3-5 sigma of the sky level may represent a concern, as they contribute to the noise with a salt-pepper pattern that may be hard to remove if one has only a few images to play with. Given e.g. a noise level of 2 DN, corresponding to the fluctuations of bit 14, one will be mostly concerned by bits 13-12, which control the 3-5 sigma level (~6-10DN) above the noise.

In Figure 10a-c we show three simulated images obtained with the same parameters used for Figure 9.
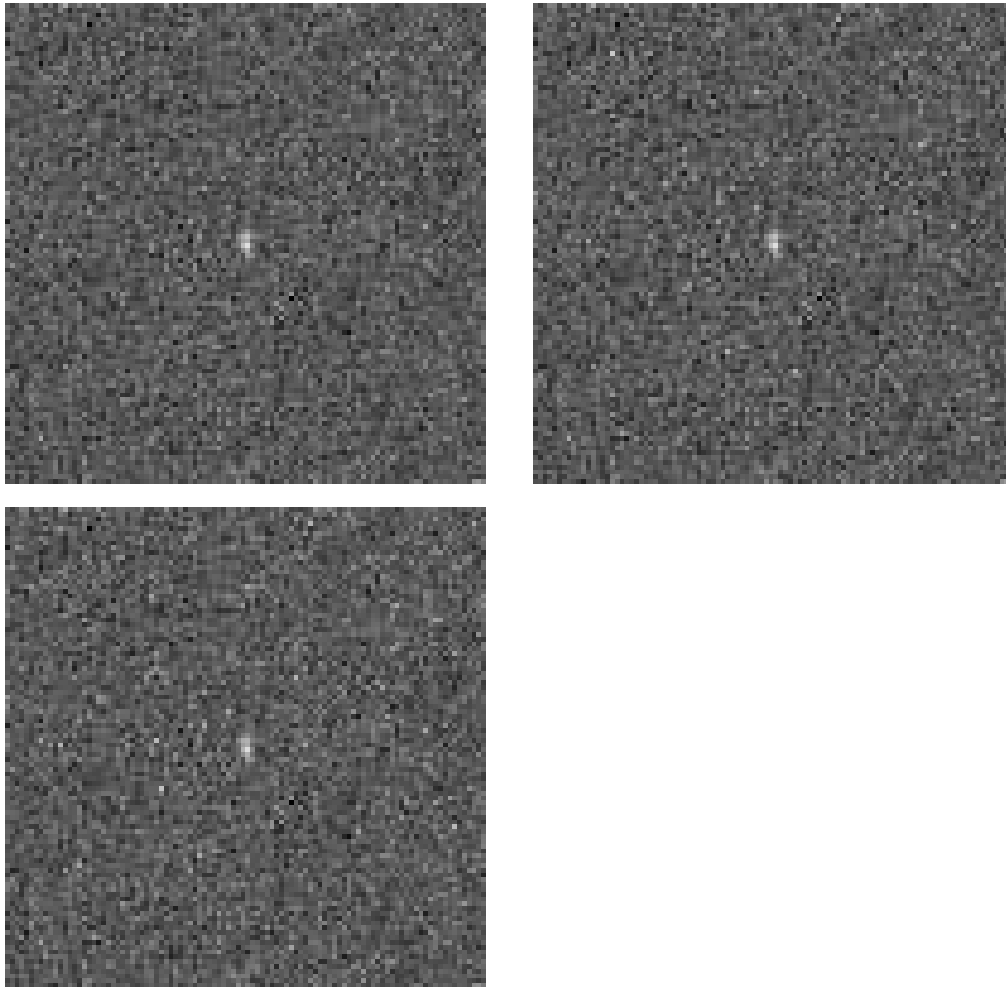
Figure 10: Clockwise from the top left, Figure 10a, with pure 2DN noise, Figure 10b, with 1% digitization noise, and Figure 10c, with 2% digitization noise.

The images show three 100x100 tiles of pure 2DN noise (10a), with 1% of digitization noise in the least 3 significant bits (10b) and 2% (10c). At the center of the field we have added a noiseless artificial source, a 2-d gaussian with FWHM=2 and 4 in the horizontal and vertical directions respectively and a peak value of 10 DN (i.e. 5 sigma above the original noise) to represent a faint elliptical galaxy. The images are presented with the same scale (+/- 10DN) and illustrate the increased salt-pepper noise. Aperture photometry on the source, however, gives the results shown in Table 1.

Biased Pixel Fraction	Flux	Error
0%	126.3	19.1
1%	126.3	19.1
2%	127.3	19.4

Table 1. Photometry results for a source in images affected by digitization noise.

In practice, there is a very small (~1%) degradation of the sensitivity induced by the biased ADC bits if their fraction raises from 1% to 2%. For the most ambitious programs that require the combination of several exposures, like a UDF-like survey, biased bits are more easily averaged out. This does not exclude that, in some situation, it may be convenient to apply a more strict criteria for cosmic ray removal.

In conclusion, the bias of the WFC3 ADCs seems to contribute negligibly to the readout noise. The gain coming with the elimination of the detector cross talk largely compensates for any additional ADC noise found in this study.

References

Hilbert, B. and S. Baggett, **Results of WFC3 Thermal Vacuum Testing - UVIS Readnoise and Dark Current**. Instrument Science Report WFC3-2005-13, http://www.stsci.edu/hst/wfc3/documents/ISRs/WFC3-2005-13.pdf. March 2005.

Hilbert, B. **Results of WFC3 Thermal Vacuum Testing - IR Channel Readnoise**. Instrument Science Report WFC3-2005-15, http://www.stsci.edu/hst/wfc3/documents/ISRs/WFC3-2005-15.pdf. April 2005.