In situ CCD testing.

T. M. C. Abbott

24 May 1995

Abstract

This document describes a number of methods of CCD testing which may be used at the telescope to verify the integrity of a CCD system. The rationale for each test is discussed, and sample test results are presented.

1 How to use this document

This document can be used as a cookbook. Several routes may be followed through the text.

The directions for data collection and for data reduction have been designed to follow independent paths so that the reader may collect the required data without needing to read about data reduction and *vice versa*. These paths are summarised below.

Note that each section begins with a rationale behind each test item. These may be disregarded if the reader wishes only to collect or to reduce data.

For instructions on what **data to collect** for a CCD test, work through these sections:

Section 2.1 ... Introduction (start here)
Section 3.1 ... Bias images
Section 4.1 ... Transfer and Linearity Curves
Section 6.1 ... Low count level flats
Section 7.1 ... Dark Images
Section 8.1 ... CTE
Section 9.1 ... The Shutter Pattern
Section 10.1 ... Bit Biases

For general (ie. for no specific environment) instructions on **how to reduce this data**, work through these sections:

Section 2.2 ... Introduction (start here)
Section 3.2 ... Biases
Section 4.2 ... Transfer curve
Section 5.2 ... Linearity curves
Section 6.2 ... Low count level flats
Section 7.2 ... Dark images
Section 8.2 ... CTE
Section 9.2 ... The Shutter Pattern
Section 10.2 ... Bit Biases

2 Introduction

This document derives from a CCD testing programme designed and implemented at ESO, La Silla. The intent of this programme was to provide regular, comprehensive tests of the CCDs normally offered for observations in order to detect problems which might have compromised astronomical observations.

The philosophy was to test each CCD in the environment in which it is used. A number of images were collected as part of each test which were then reduced and a report written describing the results.

It is often difficult to exactly determine the characteristics of a CCD outside the lab. Therefore, our concern was to draw attention to those problems which are large enough to affect observations.

Since collecting the data is currently a time-consuming task and, in most cases, data reduction must be performed offline, every effort has been made to present methods of extracting useful information from flawed data as well as ideal data.

Each section of this document begins with a description of the rationale behind the procedures, including, where possible, a thorough analysis of what the user might expect to see and how to interpret it.

2.1 Introduction: Data Required

All images should be unbinned and full frame. The images should include bias overscan regions in both axes. If the control system does not have sufficient memory to hold an entire frame, the images should cover the area of the CCD normally used for observations.

All images should be collected in the readout mode normally used for observations and with maximum digital dynamic range.

Image Labels:

In an effort to maintain an upward compatability with the Calibration database for the NTT, the image label keyword used at ESO was set for each image as described in Chapter 4 of the VLT document "EMMI/SUSI calibration plan for an on-line calibration database" reference number OSDH-SPEC-ESO-00000-0002/2.0.

In summary, the label read:

< detector > / < exp-type > : < exp-time >

where < detector > is the CCD number, < exp-type > is "BIAS", "DARK", "FF-det" (for the exposures taken in the sequence of 16 pairs), "FF-lc" (for the low-level exposures), or "FF-foc" (for the exposure taken in "FOCUS" mode) and < exp-time > is the exposure time, in seconds.

This labelling convention is also recommended for use at other observatories.

Go to section 3.1 for the first step in data collection

2.2 Introduction: Data reduction

The detection of possible problems is greatly facilitated by an approach to data reduction which is suspicious that not all is immediately visible. We therefore recommend that although canned routines may form the basis of reduction strategy, the user should be creative in inspecting the data in ways not automatically provided.

Go to section 3.2 for the first step in data reduction

This document is arranged so that following the cookbook (section 1) will ensure that the MIDAS routines are executed in the proper sequence.

3 Bias images

Perhaps the most information about the condition of a CCD can be found by careful examination of a bias frame, or dark exposure of zero integration time. The ideal bias frame is a flat noise image with the amplitude of the noise being the readout noise of the CCD in ADU. The exercise, therefore, is to identify and quantify any unwanted structure, and determine whether or not it is significant when compared to the readout noise of the CCD.

In practise, this is accomplished by collecting 9 bias images. Each can be examined seperately for the presence of various forms of noise:

- Normal readout noise, of Gaussian distribution.
- Coherent pickup noise - interference noise which is the same for every image and which may be removed from science data by subtraction of a mean or median bias image.
- Incoherent pickup noise - interference noise which is different for every exposure and therefore cannot be removed from science data.
- Burst noise - related to interference noise, random noise may be present on short (compared to the time taken to read out an image) timescales.

Bias images can also be used to identify hot pixels and measure the stability of the bias injection.

3.1 Bias images: Data Required

Obtain 9 dark images of zero integration time.

Use the image label "*NN*/BIAS:", where *NN* is the CCD number.

Use the unbinned, full frame with overscan regions in both axes (or, if not possible, use the same region as used for normal observations). Use the "SLOW" readout mode.

Inspect for obvious flaws. Be sure there are no light leaks.

Go to section 4.1 for the next step in data collection

3.2 Bias images: Reduction Procedure

9 Bias images will have been collected in the course of the test.

- Examine each bias image carefully by eye. Make sure that the cut levels chosen to display these do not hide low level effects (ie. if it is dominated by high pixels or cosmic rays). In particular, look for:

- excessive numbers of cosmic rays - more than a few ($\sim$ 5) in a 512x512 pixel subregion may be considered excessive, particularly if such numbers appear in all bias images. Large numbers of cosmic rays may be indicative of a radioactive dewar window.
- wavy patterns indicating unwanted electronic interference. The human eye is very good at spotting low-level patterns in noisy image data. In fact, many astronomers will complain at the slightest hint of an interference pattern. In principle, they are right to do so, since, under ideal conditions, it is perfectly possible to develop CCD control electronics which can generate images that are apparently entirely clear of unwanted patterns. In practise, however, the telescope dome is a far from ideal environment and there are often signals present which can penetrate even the best shielding schemes and contaminate the data.
- Burst noise. On occasion we have observed intereference in CCD images which take the form of a few pixels in a row with significantly greater noise (the cause of the most obvious example of this was never adequately identified, but may have been due to a defective component). In one case, several rows were affected, appearing as a band across the image (in this case, an airconditioning unit was found to have a defective switch which sparked when the unit turned itself on or off).
- Salt-and-pepper noise - described in Massey and Jacoby (1992, ASP conference series vol. 23, p 240), this effect has also been seen at the NTT. Look for neighbouring pairs of high and low pixels.
- Slopes - significant slopes in bias images may indicate unstable bias injection. However, note that many systems will show an exponential trend towards the true bias value starting at the leading end of each row. This is a result of the finite recovery time of the AC-coupled output amplifier after a vertical transfer. The pattern is usually stable and may be removed from a science image by subtraction of a mean bias frame.
- "Hot" regions. Some older CCDs (most notably RCAs) form LEDs in the output amplifiers. This will appear a a brightening towards the corner of an image, and are most easily seen in long darks. Such LEDs are a sign of aging and may signify the impending demise of the CCD.
- Dark spots. A dark region in a bias image is a cause for concern and the bias injection circuitry should be checked. Be sure that there were no light leaks (is there a step at the bias overscan?) and that the dark region is not just a shadow.

- Measure the mean and standard deviation of a suitable region of each bias frame. The mean should be within 1 ADU of the same value for every frame. The standard deviation is a measure of the readout noise in ADU/pixel and should also be substantially the same for each image. For any significantly outlying measurements, check the original bias frame.

- Make a histogram of a subregion of at least one raw bias frame. On a log plot, this should take the form of a parabola, indicating that the noise follows a Normal

distribution. Figure 1 is an example demonstrating an obvious problem.

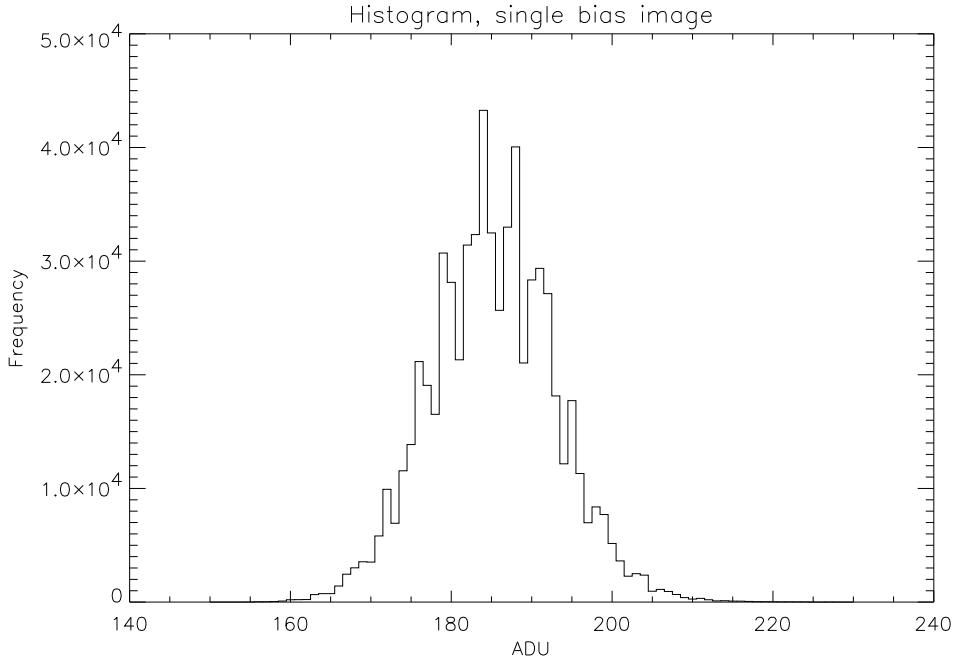

Figure 1: Sample histogram of a single bias image. Note the "picket fence" pattern caused, in this case, by biasing in less significant bits. These data were taken from ESO, La Silla CCD test 09-94-10-23.

- Construct a median stack of the nine bias frames such that the counts in a pixel of the result is the median of the counts in the same pixel in the nine raw frames. Make a note of any hot pixels seen (pixels with signal greater than several times the noise in the image), their amplitude, and how much of the column above that they affect. Figure 2 is an example map of the hot pixels and the columns they affect in ESO CCD#36 (A Tek 2048x2048 pixel CCD).

  Any patterns present in the median image represent coherent interference signals which should be the same for each image and so may be removed from any science data by subtraction of the median image.

- Make a histogram of the median stack, check that the distribution is Normal.

- Plot the mean row and column of the median stacked bias image. These plots can be used to diagnose the presence of low-level light leaks, since a step will be seen between the imaging area and the overscan region and a slope may be seen in the vertical direction. Also, any interference patterns which are identical for each row will be more clearly brought out in the mean row. Figure 3 is an example with significant fixed pattern noise.

- The Fast Fourier Transform or FFT is the best way to quantify noise signals and should always be performed, even if no signals are immediately obvious in the raw

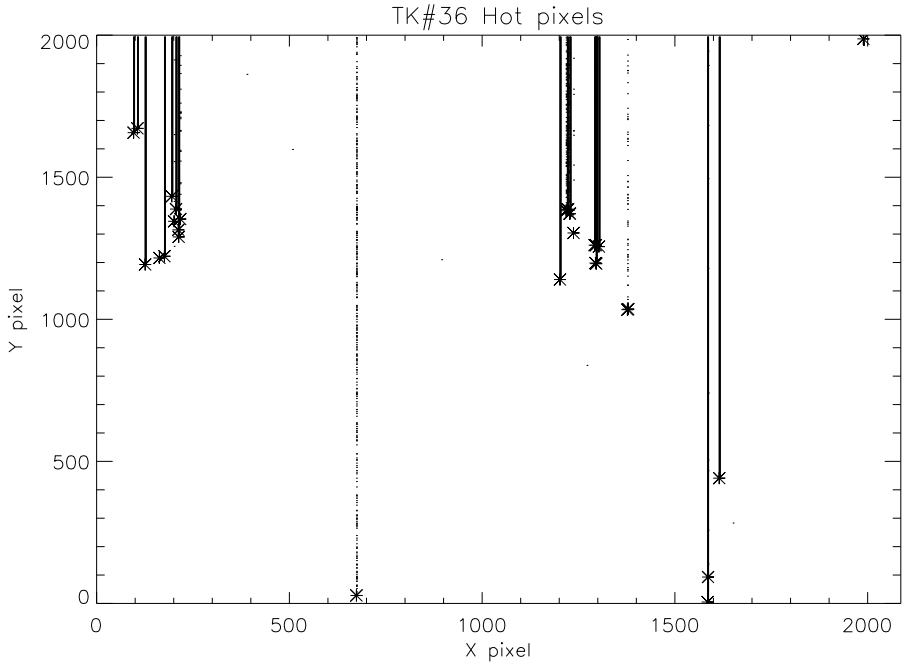

Figure 2: Sample map of hot pixels found in CCD#36, taken from ESO, La SIlla CCD data set 36-94-05-05

bias frames or the median stack (some may be present at too low a level to be seen by eye, or they may be "dithered" such that they are not coherent from row to row). There are several strategies which may be used. (Remember to remove the average signal from the image before computing the FFT. Obvious large scale variations should also be removed by fitting polynomials and subtracting.)

- Most signals of interest are of relatively high frequency, with a wavelength (in pixels) significantly less than the length of a row. These may be mostly cleanly identified by computing the FFT for each row and taking the average for the entire image. Measure the amplitudes of any significant (eg. $>\sim 0.05$ ADU) signals. If the clocking rate is known, the frequencies of these signals may also be measured and the source perhaps tracked down. Figure 4 is an example of such a "mean row" amplitude spectrum.
- A two-dimensional FFT can provide information in the frequency domain in both the horizontal and vertical directions and, in particular, will indicate any modulation of a signal, as measured by the mean row FFT, as the CCD is read out (a stable signal will produce a high column in the 2-d FFT which is the same amplitude at all vertical frequencies, an unstable signal will not).

- Apply the above FFT calculations to both a raw frame and the median stack.
- A good rule of thumb is provided by considering that an interference signal contributes to the effective readout noise of the CCD, but in a non-Gaussian fashion.

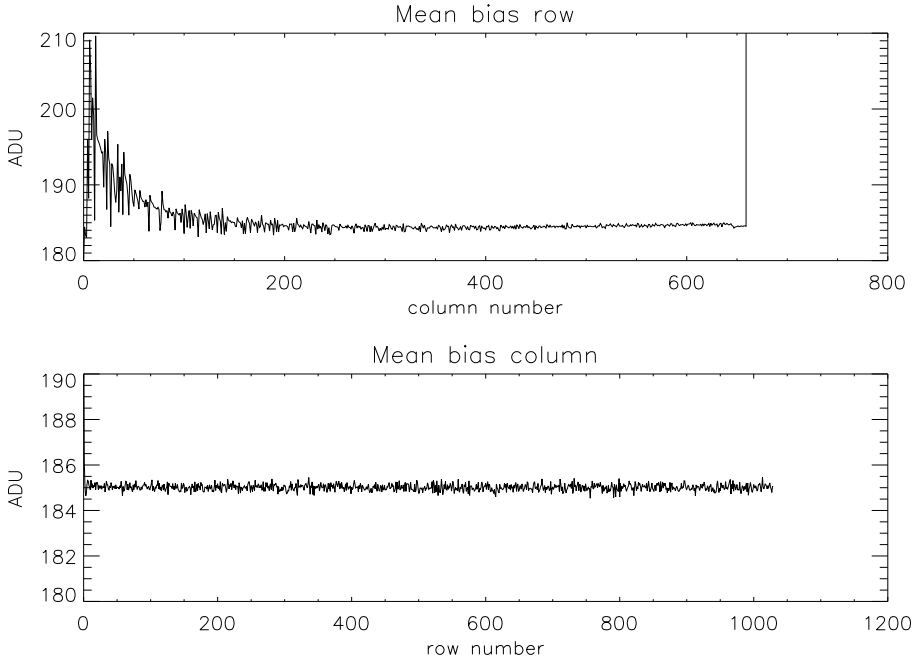

Figure 3: Sample mean row and column of a median stack of 9 bias frames. The slope at the left edge is a fixed pattern and may be subtracted from science data. The mean column was constructed using columns away from this slope. The data set were taken from ESO, La Silla CCD test 09-94-10-23.

Thus any signals present should be small compared to the CCD readout noise (in ADU). If there are many interference signals present, they may interfere with eachother so that their combined effect can be large compared to the readout noise of the CCD.

The next step in data reduction is section 4.2.

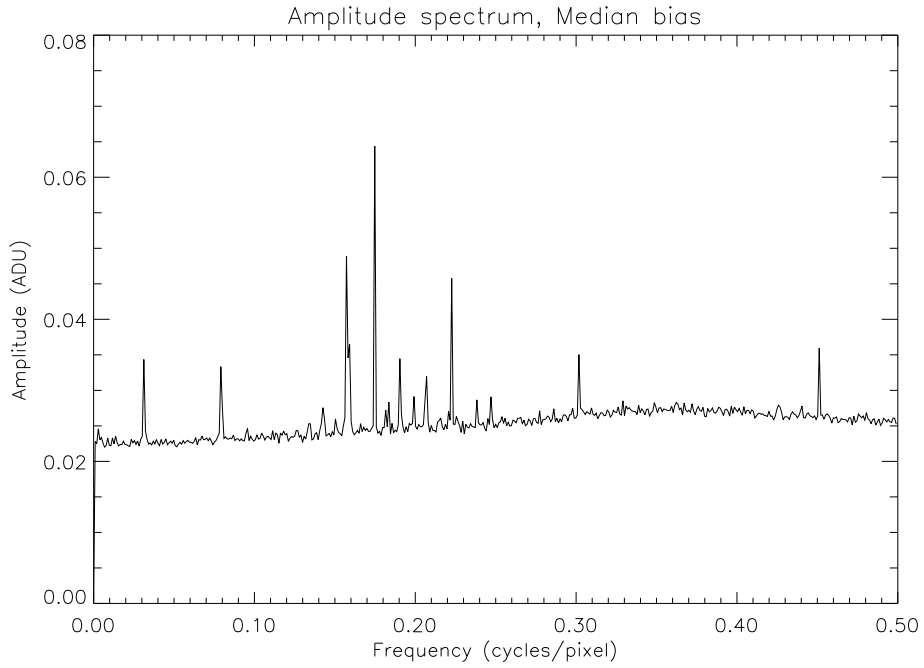

Figure 4: Sample "mean row" amplitude spectrum. Note the many peaks indicating a complex interference pattern. Taken from ESO, La Silla CCD data set 19-94-09-04.

4 The Transfer Curve

The transfer curve is a plot of the square of the noise (variance) in an image after correction for flat-field effects versus the mean counts. It was first developed as a tool for investigation of CCD properties by Janesick (1987) and is used to measure the CCD conversion factor in electrons/ADU.

At illumination levels less than saturation, the variance of the signal in electron units, σ_e^2 is given by:

$$\sigma_e^2 = \sigma_p^2 + \sigma_{ron}^2. \tag{1}$$

where σ_p is the noise in the signal and σ_{ron} is the readout noise in electrons. Assuming Poisson statistics, σ_p^2 is the signal in electrons and is equal to the illumination level in electrons, I_e which is the illumination level in ADU, I_A, multiplied by the CCD conversion factor (or, imprecisely, the gain), g, in electrons/ADU. Similarly, σ_e is the standard deviation of the signal in ADU, σ_A, multiplied by g. Thus:

$$g^2\sigma_A^2 = gI_A + \sigma_{ron}^2 \tag{2}$$

or

$$\sigma_A^2 = \frac{1}{g}I_A + \frac{1}{g^2}\sigma_{ron}^2. \tag{3}$$

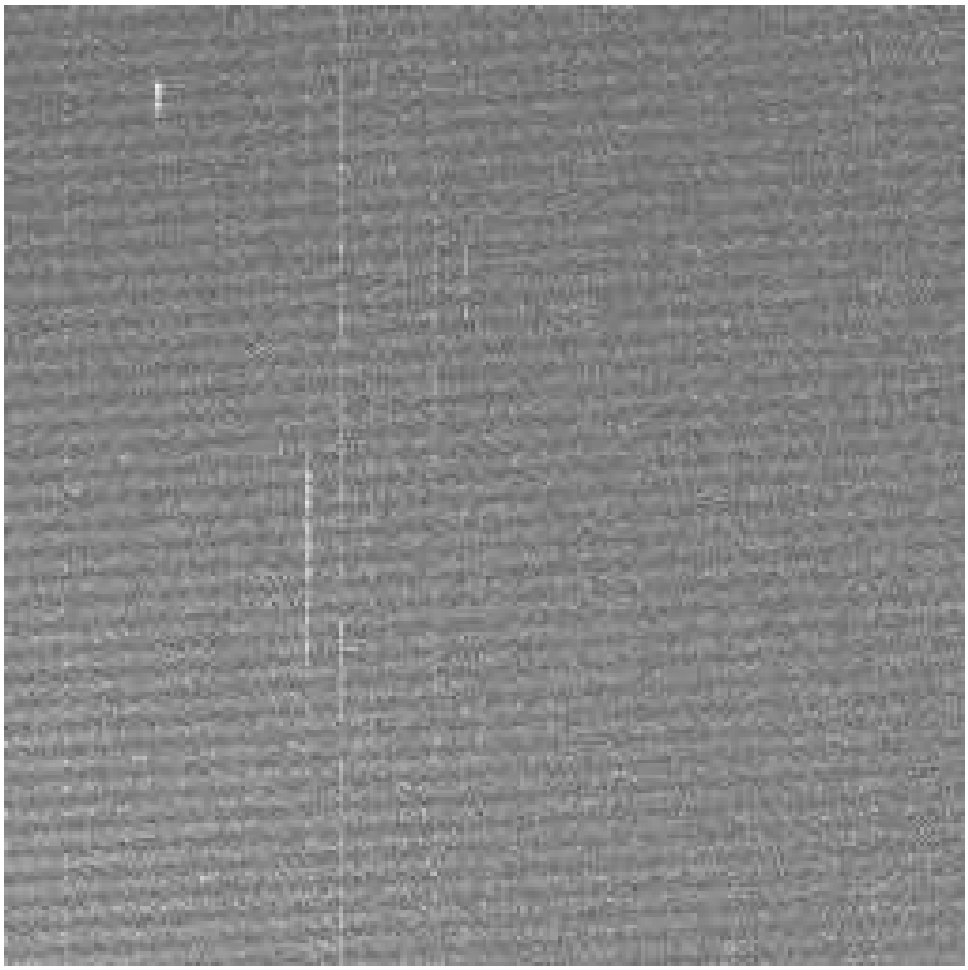

Figure 5: Sample two-dimensional amplitude spectrum of a single bias image taken from ESO, La Silla CCD data set 19-94-09-04. The near horizontal bands do not appear in the amplitude spectrum of the median stack image.

Therefore, the slope of the plot of the variance of the signal in ADU, σ_A^2, versus the mean counts, I_A, is the inverse of the CCD conversion factor.

The readout noise is obtained from the square root of the y-intercept, which is simply the variance of the signal in a bias frame, divided by the conversion factor.

Note: at high count levels, the transfer curve levels off and may turn over. This is because the system has become saturated. Saturation may be defined in either of two ways.

- Digital saturation - the CCD, while still accumulating charge in linear proportion to the illumination level and delivering it to the output amplifier upon readout, is producing signals which are above the dynamic range of the analogue-to-digital converter. Such pixels will be measured to have a content equal to the maximum value delivered by the converter (32768 ADU in a 15-bit converter), but will have contained more charge. The variance of a fully saturated image is 0.

- Analogue saturation - individual potential wells of the CCD are filled to their limit. Any further charge generated by incident light will spill over into neighbouring pixels of greater well depth, primarily in the vertical direction (an effect known as "blooming"). Note that the response of the CCD may still be apparently linear, but the variance no longer increases linearly with signal as the unconstrained, excess charge is smeared across the image. The response will start to level off and the variance to decrease with increasing signal when the majority of potential wells (pixels) are filled with signal electrons.

Normally, the system gain should be set so that digital saturation is encountered at signal levels less than analogue saturation.

4.1 Transfer curves: Data Required

Use the unbinned, full frame with overscan regions in both axes (or, if not possible, use the same region as used for normal observations).

Use the image label "NN/FF-det:TT", where NN is the CCD number and TT is the exposure time in seconds.

Illuminate the CCD with a stable light source (eg. a β light or feedback stabilised LED), such that the count rate is 500-1000 ADU/sec. Collect a sequence of pairs of images such that the exposure times of both images of a pair are the same and the exposure times of the image pairs generate illumination levels ranging from just above bias to near digital saturation. The exposure times should normally be distributed such that the illumination levels are obtained are spread evenly over the dynamic range to be studied. Alternatively, more exposures can be collected at short integration times, but be warned that a mechanical shutter may respond nonlinearly to very short exposure times. For reasons explained in the section on CCD linearity (section 5), the sequence of pairs of flat field images should be obtained in two groups - the first with increasing integration times and the second with decreasing integration times. The exposure times of the two groups may be the same, but it is better that they be different and interleaved. For example, the first group may follow the exposure sequence 2x10 sec, 2x20 sec, 2x30 sec etc and the second 2x25 sec, 2sx15 sec and 2x15 sec.

Here is a sample sequence, as collected for TK#33 on 12/12/94, where the images were concentrated at short exposure times. The first column is a file name, the second is the image label, the third is the exposure time in seconds and the fourth and fifth are the image dimensions in pixels.

```
tk330010.mt            33/FF-DET:1    1.000   581   520
tk330011.mt            33/FF-DET:1    1.000   581   520
tk330012.mt            33/FF-DET:3    3.000   581   520
tk330013.mt            33/FF-DET:3    3.000   581   520
tk330014.mt            33/FF-DET:5    5.000   581   520
tk330015.mt            33/FF-DET:5    5.000   581   520
tk330016.mt            33/FF-DET:9    9.000   581   520
```

```
tk330017.mt            33/FF-DET:9    9.000   581   520
tk330018.mt           33/FF-DET:15   15.000   581   520
tk330019.mt           33/FF-DET:15   15.000   581   520
tk330020.mt           33/FF-DET:25   25.000   581   520
tk330021.mt           33/FF-DET:25   25.000   581   520
tk330022.mt           33/FF-DET:45   45.000   581   520
tk330023.mt           33/FF-DET:45   45.000   581   520
tk330024.mt           33/FF-DET:80   80.000   581   520
tk330025.mt           33/FF-DET:80   80.000   581   520
tk330027.mt           33/FF-DET:90   90.000   581   520
tk330028.mt           33/FF-DET:90   90.000   581   520
tk330029.mt           33/FF-DET:60   60.000   581   520
tk330030.mt           33/FF-DET:60   60.000   581   520
tk330031.mt           33/FF-DET:35   35.000   581   520
tk330032.mt           33/FF-DET:35   35.000   581   520
tk330033.mt           33/FF-DET:20   20.000   581   520
tk330034.mt           33/FF-DET:20   20.000   581   520
tk330035.mt           33/FF-DET:12   12.000   581   520
tk330036.mt           33/FF-DET:12   12.000   581   520
tk330037.mt            33/FF-DET:6    6.000   581   520
tk330038.mt            33/FF-DET:6    6.000   581   520
tk330039.mt            33/FF-DET:4    4.000   581   520
tk330040.mt            33/FF-DET:4    4.000   581   520
tk330041.mt            33/FF-DET:2    2.000   581   520
tk330042.mt            33/FF-DET:2    2.000   581   520
```

The next step in data collection is section 5.1.

4.2 Transfer curve: Reduction Procedure

To generate a transfer curve, the noise must be measured at each exposure level, independent of flat-field variations. To compute this, take each pair of images (both with the same integration time) and subtract one from the other. Select a region which, in a raw image, is clear of nonlinear defects (such as traps or hot pixels). Compute the variance of this region in the difference image and divide by two. Also compute the mean signal in the same region of both the pair of raw images. Perform the same calculation for all the other pairs of flat-fields. When finished, you should have a table with four columns (exposure time of the pair, mean counts in each frame and one-half the variance of the difference frame, referred to as columns 1 through 4 in following discussion) and as many rows as pairs of images obtained. Table 1 is derived from the images listed in section 4.1.

The bias level should be subtracted from the data in columns 2 and 3. This can be done before the above calculation by subtracting the median stack of the 9 bias images as computed in section 3.2 or after, by subtracting a mean bias value from the data in columns 2 and 3 of the table. Note that if there is significant fixed-pattern noise in the bias frame, then a master bias frame must be subtracted before calculation.

The transfer curve is then a plot of column 4 (the variance) along the y-axis versus column 2 (or 3) along the x-axis. Figure 6 is the transfer curve derived from the data in table 1. This plot should be a straight line up to count levels approaching saturation, when the noise turns over and may go to zero in a completely saturated image.

The transfer curve, below saturation, should be a straight line. Fit a straight line to this curve and compute the reciprocal of the slope. This is the conversion factor of the CCD system in electrons/ADU. The readout noise (in electrons/pixel) may be computed by multiplying the standard deviation of a bias image by the conversion factor.

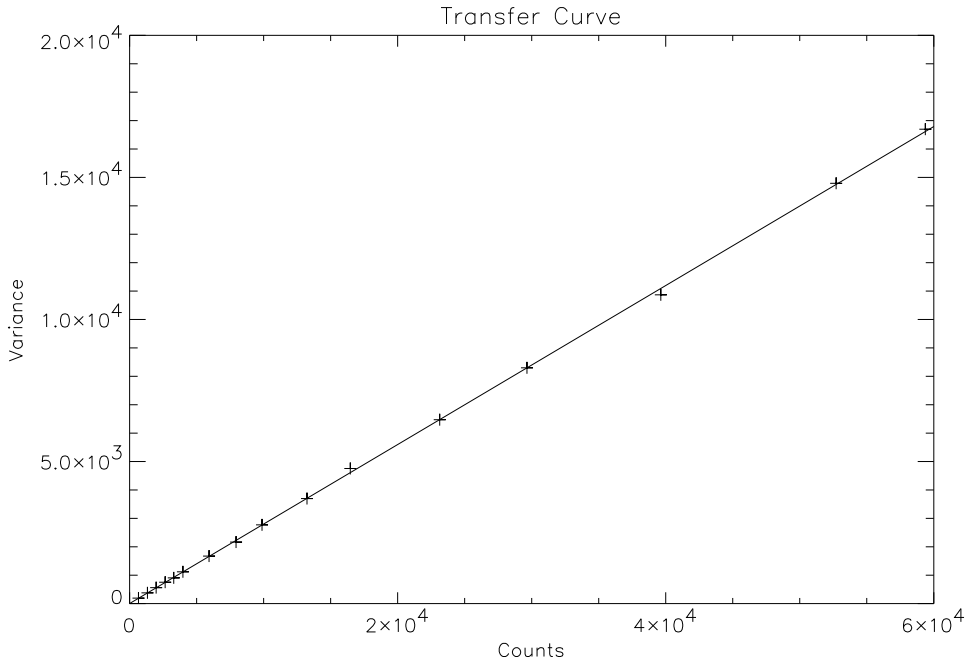

Figure 6: Sample transfer curve for the ESO, La Silla CCD test data set 33-94-12-12. The inverse of the slope of the fitted line gives a conversion factor of 3.57 electrons/ADU. The standard deviation of bias is 2.45 ADU, yielding a readout noise of 8.75 electrons/pixel.

The next step in data reduction is section 5.2.

5 Linearity

In this section, we will consider what is measured by the mean counts in a set of exposures of different exposure times when the CCD is illuminated by a stable light source such as a β light.

If F is the number of signal electrons collected per second by the CCD (which equals the incident flux times the quantum efficiency in the range of interest), g is the CCD conversion factor in electrons/ADU, and T_t is the actual exposure time (ie. not necessarily the *measured* exposure time), then the true mean count level, I_t, is:

$$I_t = gFT_t. \tag{4}$$

However, no CCD is strictly linear and nonlinearities may be modelled by a polynomial dependence of g on FT_t:

$$g = g_0 + g_1 FT_t + g_2(FT_t)^2 + ... = \sum_{i=0}^{n} g_i(FT_t)^i \tag{5}$$

There may also be an error in the bias subtraction such that the observed mean counts, I_o, are in error by an amount b, such that:

$$I_o = I_t + b \tag{6}$$

Finally, the observed exposure time, T_o, is the true exposure time, T_t, less the shutter delay s (which is assumed to be constant for all exposure times):

$$T_t = T_o + s \tag{7}$$

Therefore, the mean counts measured by the experiment are:

$$I_o = F(T_o + s)\sum_{i=0}^{n} g_i[F(T_o + s)]^i + b. \tag{8}$$

Where F, g_i, s and b are unknowns. For now, we shall assume that F is a constant (ie. the light source is constant) and that b is small enough that we can disregard it.

Ideally, we should like to fit Equation 8 to real data and solve for the unknowns. In practise, this is a difficult task, if not impossible. For example, if we assume that the CCD exhibits only first-order nonlinearities ($g_{i\ (i>1)} = 0$), then Equation 8 becomes:

$$I_o = g_1 F^2 T_o^2 + (g_0 F + 2sg_1 F^2)T_o + g_0 Fs + g_1 F^2 s^2. \tag{9}$$

Therefore a fit of a quadratic to a plot of the observed mean counts in an image versus the reported integration time (the "linearity" curve, by its normal definition) will provide

measurements of 3 coefficients. However, the coefficients of T_o in equation 9 include 4 unknowns (g_0, g_1, s and F) and so a complete solution is not possible. Nevertheless, it is possible to extract a measurement of s, since at $I_o = 0$, $T_o = -s$.

g_1 can be estimated as follows - if we fit the equation

$$I_o = a_2 T_o^2 + a_1 T_o + a_0 \tag{10}$$

to the observed linearity curve, then, after some tortuous algebra, the following relation is found to hold:

$$\frac{g_0^2}{g_1} = \frac{a_1^2}{a_2} - 4a_0. \tag{11}$$

Since g_0 is measured by the transfer curve (to first order only, since for completeness, we should apply a similar analysis of nonlinearities in the transfer curve), we can therefore estimate the value of g_1 and thus the magnitude of a first-order nonlinearity (and, if desired, F).

This measurement of g_1 rests on a number of untested assumptions - that s and F are constant, b is small and there are no higher nonlinearities present in the response of the CCD. It is true that the residuals of the fit might provide some clues as to the presence of higher-order nonlinearities and possible variations in the light source (or F) during the test, but even if these are small, then the precision of the measurement of g_1 is still difficult to estimate.

In practise, we are most interested in determining the shutter error as precisely as possible, particularly at short exposure times (where the relative error is large and because the observer normally takes calibration exposures at short exposure times), an estimate of the total amplitude of any nonlinearities in the typical dynamic range used for observations, and some idea of the stability of the light source over the test.

We have converged on a practical strategy which exposes these three factors in a clear, qualitative fashion and, in the right circumstances, produces reliable quantitative measures for the shutter delay and nonlinearity of the CCD system.

The count rate may be derived by dividing equation 8 by the exposure time. Ideally, we should like to divide by the true exposure time, T_t, but in practise we must divide by the reported exposure time, T_o. If the shutter error, s, is non-zero, then the measured count rate will tend to anomolously large values with shorter exposure times (or small values if $s < 0$, which is unusual). Therefore, in order to obtain an estimate of s, we compute the count rates, C, with the formula (ignoring b):

$$C = \frac{I_o}{T_o + s'} = F(T_o + s)/(T_o + s') \sum_{i=0}^{n} g_n [F(T_o + s)]^n \tag{12}$$

where s' is a test value of the shutter delay. When $s' = s$, the count rate is given by:

$$C = F \sum_{i=0}^{n} g_n [F(T_o + s)]^n \tag{13}$$

which, in the linear case ($n = 0$), is a constant ($= g_0 F$), or, in the case with first-order nonlinearity ($n = 1$), a straight line (with respect to $T_o + s$). As a further refinement, divide by the mean count rate, $\sim g_0 F$, over the test to yield:

$$C_f = 1 + \frac{\sum_{i=1}^{n} g_n [F(T_o + s)]^n}{g_0} \tag{14}$$

Thus a plot of this count rate computed with observed exposure times corrected for the shutter delay versus the observed exposure time (or the mean counts in the image) provides a curve whose amplitude of deviation from unity is a measure of the total nonlinearities present in the system over the dynamic range explored by the test, expressed as a fraction of the linear term (and may therefore described as a "fractional" nonlinearity).

For this to be valid, s' must be as close to s as possible. A suitable estimate can be obtained by assuming that any nonlinearities of second order or higher are negligible. In that case, the count rate curve should be a straight line. Therefore, the optimum value of s' may be selected by the straight line fit to C_f versus I_o which has the smallest standard deviation of fit. Note that this computation implicitly gives greater weight to those images of shorter exposure time as the count rates for these are most affected by errors in the reported exposure time.

Typically, any variations in the brightness of the light source used to illuminate the CCD occur on a timescale comparable to the time taken to collect the data, or longer. Such variations may also be comparable to the amplitudes of the nonlinearities in the system. For example, a β light shows a variation in generated flux of $\sim 0.6\%$ per degree centigrade change in temperature. As a result, we must require that the temperature of the housing of the β light be stable to within $\sim 0.1°$C in order to confirm the presence of $\sim 0.06\%$ nonlinearities. This kind of temperature stability is very hard to achieve in an open instrument mounted at Cass. or even in a Coudé room. Therefore, until we can obtain a light source with the required stability, we are forced to express the nonlinearities present in the system as an upper limit based on what we know of the variability of the light source.

If the images are collected in a simple sequence of increasing (or decreasing) exposure time, any variation in illumination will normally (unless the temperature changes quickly) take a form very similar in form to a smooth nonlinearity in the response of the CCD. Instead, if the data are collected in two sequences, say one of increasing exposure time and the other of decreasing exposure times (with different, interleaved exposure times in each case), then a smooth change in illumination level will appear as two distinct bands of data. An example is shown in figure 10 where the light level has decreased over the course of the test.

If, on the other hand, the two sequences follow the same curve, then we may be reasonably sure that the illumination level did not change during the test (unless it followed the exact same change in reverse through the second sequence as through the first, which is unlikely).

In these circumstances, we can obtain a precise measure of the form of any nonlinearities present and use the resulting relation to correct science data for them.

One further point remains to be addressed - how precise does the bias subtraction need to be to avoid contaminating the results? If we assume a linear response from the CCD, C_f becomes:

$$C_f = \frac{I_o}{T_o + s'} = \frac{g_0 F(T_o + s) + b}{T_o + s'} \tag{15}$$

Therefore, we may require:

$$\frac{b}{g_0 F(T_o + s)} << \frac{s}{T_o} \tag{16}$$

or

$$\frac{b}{g_0 F} << s. \tag{17}$$

Typically, $s \sim 0.05$ sec and $g_0 F \sim 1000$ ADU/pixel/sec and thus we require $b << 50$, which is reasonable under almost any circumstances.

5.1 Linearity curves: Data Required

The data set required for the analysis of the linearity of the CCD system is the same as that required to generate the tranfser curve (except that pairs of images at each exposure time are not strictly necessary). Refer to section 4.1 for details.

The next step in data collection is section 6.1.

5.2 Linearity curves: Reduction Procedure

In this section, we will assume that data have been collected and reduced in the manner as described in the section on the transfer curve (section 4.2). Thus we have already obtained a table with four columns such that column 1 contains the integration times of each pair of exposures, columns 2 and 3 contain the (bias-subtracted) mean counts in a suitable region of each of the pair of images and column 4 contains the variance of this region corrected for flat-field effects. In the analysis of linearity, we are concerned only with columns 1 and 2 (columns 1 and 3 should produce identical results).

For completeness, plot the conventional linearity curve - mean counts versus exposure time - and fit a straight line (or quadratic) to those points that are not affected by saturation. Figure 7 is this plot for the data set given in table 1. The exposure-time-axis intercept provides a first-order measurement of the shutter delay (ie. the exposure time at zero counts). The residuals may provide some estimate of the nonlinearities present or variation in the illumination level during the test.

Generate a fifth column in the table containing the measured count rate by dividing column 2 by column 1. Divide column 5 by the mean of its contents. If there is a significant shutter

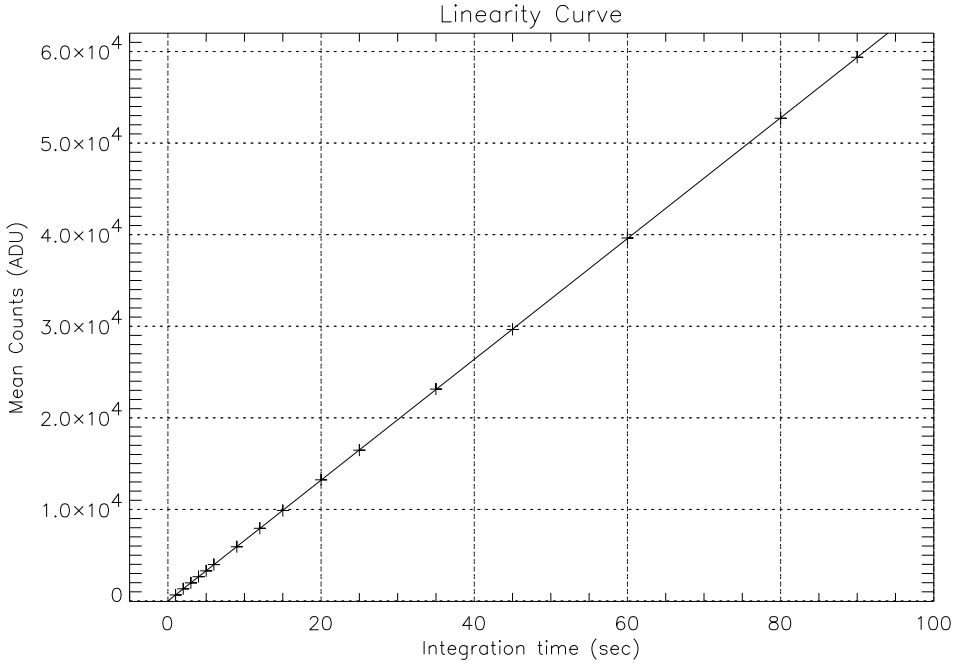

Figure 7: Sample linearity curve for the ESO, La Silla CCD test data set 33-94-12-12.

error and no nonlinearities in the CCD response nor variation in the illumination during the test, then plotting column 5 versus column 2 should produce a horizontal straight line which tends to anomolously small values at small count rates (see figure 8). The shutter error can be deduced by iteratively adjusting column 1 by a range of guesses at this error, recomputing column 5 and fitting a straight line until the best fit to a straight line is found. The input guess is then the shutter error, allowing for first order nonlinearity in the CCD. The amplitude of the points on the curve is the amplitude of any nonlinearities over the dynamic range covered by the test expressed as a fraction of the linear term in the response function. An example is shown in figure 9.

If the light source varied during the test, it will normally have done so smoothly and at the level of $\sim$ 1% or less. Since the images used to construct the count rate curve should have been collected in two sequences - one of increasing and the other of decreasing integration times - the variation should cause the curve to present two distinct arms. In figures 10 and 11 the data were collected with increasing exposure time first, starting with the shortest integration time of the whole test (1 sec, see section 4.1) and the count rate curve is dominated by an increase in the brightness of the light source. It may still be possible to obtain a reasonable estimate of the shutter delay, as in the example used here, but this may require some careful judgement. Note that only only an upper limit to the nonlinearities can be measured.

In an intermediate case, it may be found that the nonlinearities are of the same order as the variations in the light source. In figure 12, although the two sequences are seperated such that the light source varied, it either did so in a reverse sense in the second group as in the first and to a similar degree (an improbable circumstance), or the light source

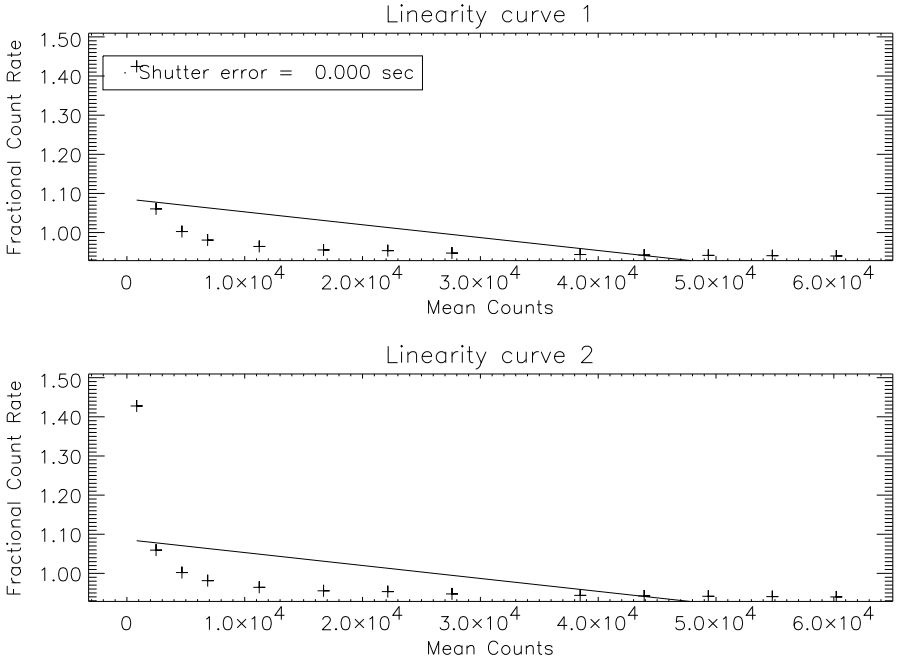

Figure 8: Sample count rate curves, uncorrected for a 0.55 sec shutter error. The data are taken from ESO, La Silla CCD test set 36-94-08-25.

changed steadily through the test, but did so to a small degree and the CCD itself is nonlinear such that it responds with a slightly higher count rate at higher signal levels than at low signal levels.

If a large number of exposures are taken at short exposure time, it may be possible to determine at what exposure time the shutter becomes unreliable. An example is shown in figure 13. Note the increased scatter at short exposure times.

The next step in data reduction is section 6.2.

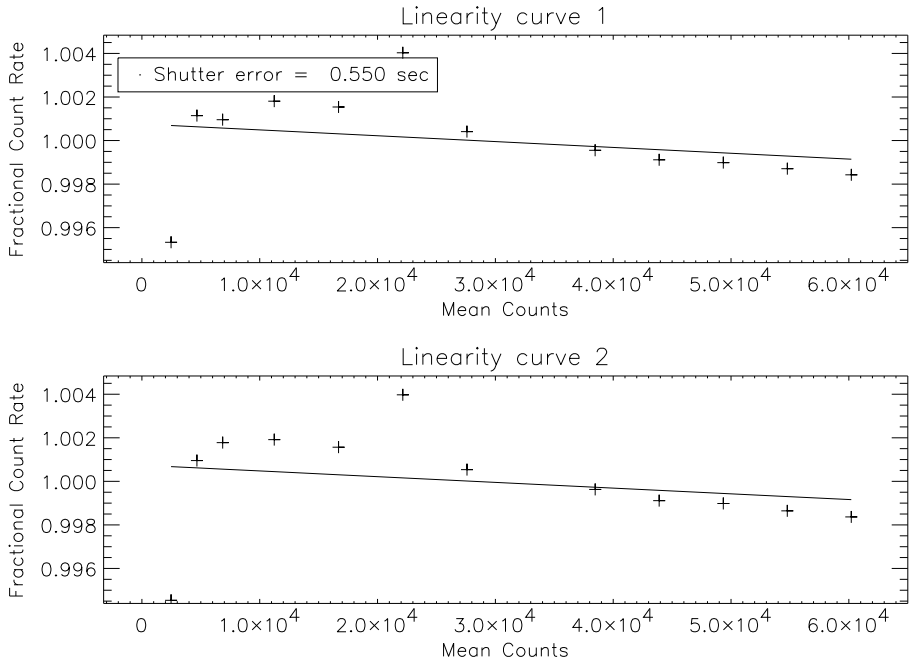

Figure 9: Sample count rate curves, corrected for a 0.55 sec shutter error. The nonlinearities present are at an amplitude of $\sim 0.4\%$. The data are taken from ESO, La Silla CCD test set 36-94-08-25.

6 Low count level flats

A low-noise, flat field image with a small number of counts per pixel can be used to identify traps and dark pixels in the CCD imaging area. A single image of a few hundred ($\sim< 300$) electrons per pixel is typically too noisy to accurately identify all traps and dark pixels, therefore a median stack of several images should be constructed and inspected.

6.1 Low count level flats: data required.

Obtain 9 flat field images with counts $\sim<$ 150 ADU/pixel. Use the stable light source (β light or stabilised LED).

Use the image label "*NN*/FF-LC:*TT*", where *NN* is the CCD number and *TT* is the exposure time in seconds.

Use the unbinned, full frame with overscan regions in both axes (or, if not possible, use the same region as used for normal observations). Use the "SLOW" readout mode.

Inspect for obvious flaws.

Go to section 7.1 for the next step in data collection

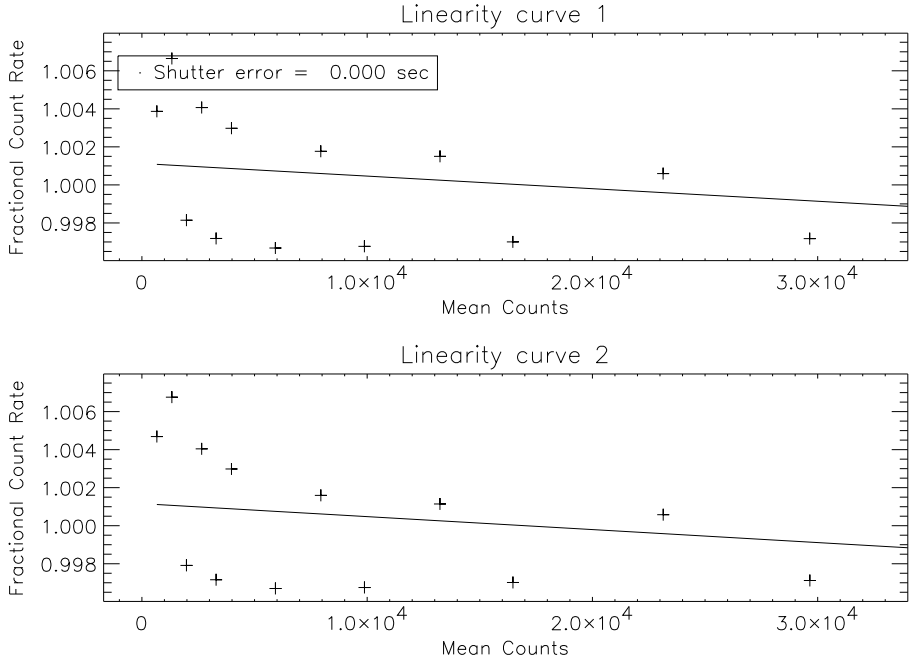

Figure 10: Sample count rate curves for the ESO, La Silla CCD test data set 33-94-12-12, uncorrected for a shutter error of 0.01 sec. Note the two bands of data.

6.2 Low count level flats: Reduction Procedure

9 low count level ($\sim<$ 300 electrons/pixel) will have been collected in the course of the test.

Construct an image in which each pixel is the median of the same pixel in the 9 raw flat fields.

Examine this median stack for traps and dark pixels. A trap appears as a dark pixel with a dark trail in the vertical direction. An example appears in figure 14. Sometimes and entire column is affected. A dark pixel is a pixel with a significantly lower response to light than the surrounding pixels.

Column offsets may also be seen. These are pairs of columns with one higher and the other lower than the surrounding columns. Some examples are shown in figure 15 Construct the mean row of the image and measure the height and depth of the high and low columns over the surrounding columns.

The next step in data reduction is section 7.2.

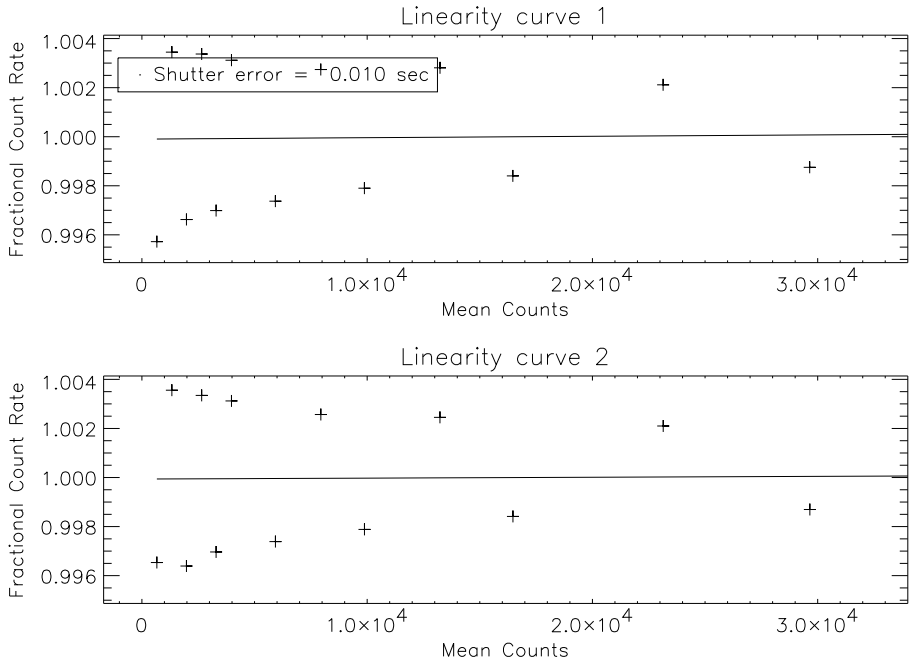

Figure 11: Sample count rate curves for the ESO, La Silla CCD test data set 33-94-12-12, corrected for a shutter error of 0.01 sec. Note the two bands of data.

7 Darks

Long darks can be used to measure the dark current, although with most modern CCDs, the dark current is so low that even the smallest light leaks can cause problems and cosmic rays must be carefully removed. Small variations in bias level can also compromise the data and overscan regions should normally be used.

7.1 darks: data required.

Obtain three half hour integrations with the shutter closed.

Use the image label "*NN*/DARK:*TT*", where *NN* is the CCD number and *TT* is the integration time in seconds.

Use the unbinned, full frame with overscan regions in both axes (or, if not possible, use the same region as used for normal observations). Use the "SLOW" readout mode.

Inspect for obvious flaws.

Go to section 8.1 for the next step in data collection

7.2 Darks: Reduction Procedure

Three one half hour dark integrations will have been collected in the course of the test.

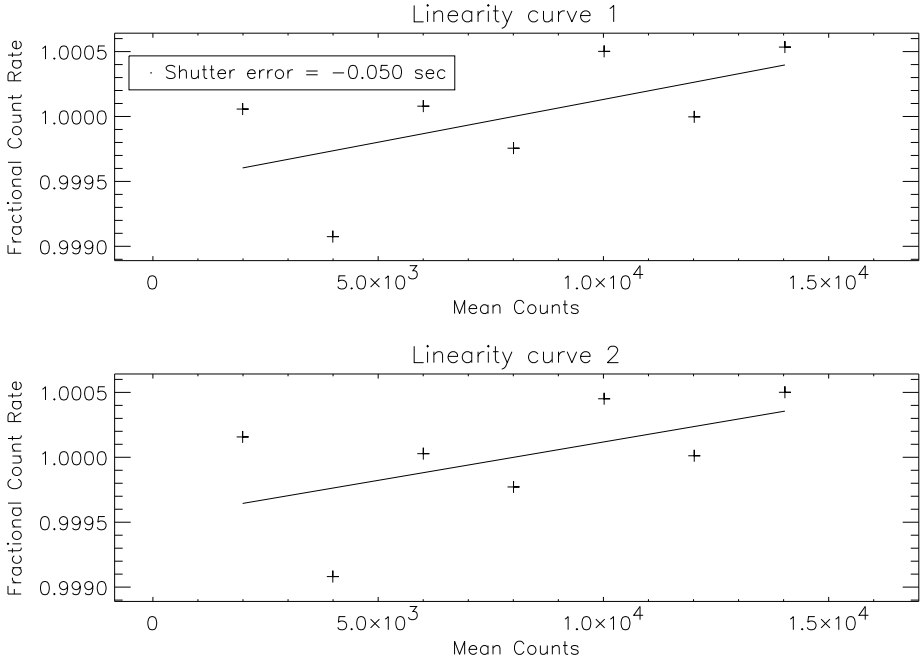

Figure 12: Sample count rate curves for the ESO, La Silla CCD test data set 09-94-10-23, corrected for a shutter error of -0.01 sec. Note the two bands of data combined with a small CCD nonlinearity

In most cases, it is sufficient to add together the three images after median smoothing and measure the counts in the imaging area above the bias level measured in the overscan regions.

The next step in data reduction is section 8.2.

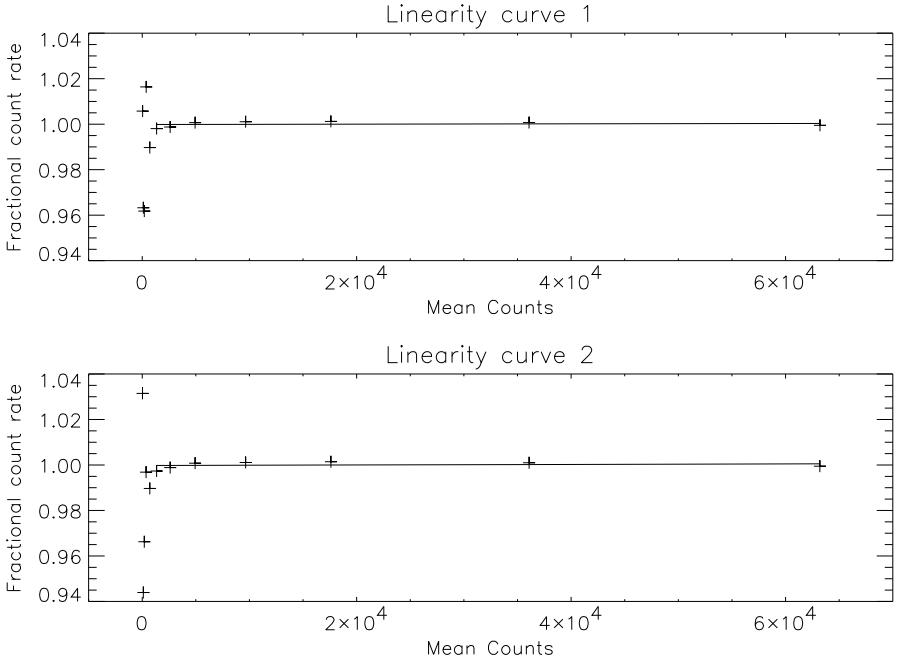

Figure 13: Sample count rate curves for the ESO, La Silla CCD test data set 19-94-05-12, corrected for a shutter error of 0.162 sec. Note the large scatter at very short integration times indicating mechanical nonlinearities in the shutter.

8 Charge Transfer Efficiency

The Charge Transfer Efficiency, or CTE, can be estimated using the Extended Pixel Edge Response (EPER) technique developed by Janesick (1987). In this technique, the bulk CTE along a row (or column) is given by:

$$CTE = 1 - \frac{I_{n+1}}{I_n n} \tag{18}$$

Where I_n are the counts, above bias, in the last pixel of the imaging region of the row; I_{n+1} are the counts, above bias, in the first pixel in the overscan region, and n is the number of pixel transfers along the row which the nth pixel made before arriveing at the output amplifier (ie. including any prescan pixels).

Normally, an average row should be derived by collapsing a well-exposed image in the vertical direction and the calculation applied to that row. The same calculation can be applied to an average column to obtain the CTE in the vertical direction.

Note that this calculation produces an *average* value for the CTE across an image.

The CTE can also be checked by examining cosmic rays in the dark images also obtained as part of the test – significant smearing towards the upper and/or left edges is an indicator of problems.

8.1 CTE: data required.

No special data is necessary for study of CTE (images can be taken from the sequence of pairs of images collected to obtain the transfer curve).

Go to section 9.1 for the next step in data collection

8.2 CTE: Reduction Procedure

Select a well-exposed flat-field image from the sequence of pairs of images used to derive the transfer curve. This image should not be saturated anywhere. Compute the mean row and column of this image (remembering to leave out the vertical overscan region when computing the mean row and *vice versa*). Use equation 18 to deduce the horizontal and vertical CTE.

Examine the cosmic rays that appear in the three dark images obtained as part of the test. Significant smearing of the cosmic rays, increasing towards the corner farthest from the output amplifier is an indication of poor CTE.

The next step in data reduction is section 9.2.

9 The Shutter Pattern

A mechanical shutter takes a finite time to travel from fully closed to fully open. In a short exposure, this can have a significant effect which is not uniform across the CCD (unless the shutter is in the parallel beam). The form of the shutter delay pattern can be deduced by comparing a long exposure flat field with an exposure in which the shutter has been opened and closed many times.

The counts in pixel i in an image where the shutter has been opened and closed once are given by:

$$I_{i,1} = (t_1 + s_i)F_i \tag{19}$$

where t_1 is the *reported* exposure time of the image, s_i is the total shutter delay at pixel i, and F_i is the counts detected per second at that pixel.

If the shutter is opened and closed n times before the CCD is read out, then the counts at pixel i will be:

$$I_{2,i} = (t_2 + ns_i)F_i \tag{20}$$

where t_2 is the reported exposure time for this second image and it has been assumed that the shutter delay is the same for every open-close cycle.

Since F_i is the same for both images, we can deduce the shutter delay at each pixel in the image from the following:

$$s_i = \frac{t_1 I_{2,i} - t2 I_{1,i}}{n I_{1,i} - I_{2,i}} \tag{21}$$

9.1 The Shutter Pattern: data required

A single, well-exposed but unsaturated exposure is required, taken under the same conditions as those for the sequence of pairs of images, but with the shutter opened and closed as many times as possible during the exposure. This can be achieved either by means of the "FOCUS" mode exposure, if one is available, or manually, by executing many PAUSE/CONTINUE cycles in the observing programme. Be sure to note the number of times that the shutter opens and closes *including* at the beginning and end of the exposure.

Use the image label "*NN*/FF-FOC:*TT*", where *NN* is the CCD number and *TT* is the integration time in seconds.

Use the unbinned, full frame with overscan regions in both axes (or, if not possible, use the same region as used for normal observations). Use the "SLOW" readout mode.

Inspect for obvious flaws.

9.2 The Shutter Pattern: Reduction Procedure

Select a well-exposed, but unsaturated flat-field image from the sequence of pairs of images used to derive the transfer curve. Together with the image taken with multiple shutter cycles, apply equation 21 to deduce the shutter delay pattern in seconds. This is best displayed as a contour plot. An example is shown in figure 16

The next step in data reduction is section 10.2.

10 Bit Biases

A perfect analogue-to-digital converter (ADC) generates as many 1's as 0's in each bit perfect white noise is input. In practise, no ADC is perfect and occasionally a bit can become completely stuck. This failure mode can be guarded against by counting the number of 1's and 0's in each bit for a well-exposed and unsaturated image.

10.1 Bit Biases: data required

No additional data is required to test for biasing of bits.

10.2 Bit Biases: Reduction Procedure

Select a well-exposed, but unsaturated flat-field image from the sequence of pairs of images used to derive the transfer curve. For each bit of the digital dynamic range used, count the number of occurences of 1's and 0's in the image. If there are no problems with the ADC, the numbers of 1's and 0's in the lower-significant bits should be almost equal, up to bit-values near the level of illumination of the image. Figure 17 shows an example where there is a slight biasing towards values of 0.

Integration Time	Median Counts 1	Median Counts 2	Variance of Difference
1.00000	663.228	663.951	192.908
2.00000	1330.13	1330.64	378.661
3.00000	1978.34	1978.42	562.158
4.00000	2653.44	2654.10	752.836
5.00000	3294.06	3294.88	904.867
6.00000	3975.84	3976.95	1119.61
9.00000	5926.32	5928.04	1672.88
12.0000	7942.08	7942.90	2164.89
15.0000	9878.07	9880.59	2771.98
20.0000	13233.3	13232.2	3697.87
25.0000	16467.3	16472.0	4754.34
35.0000	23137.3	23143.3	6471.91
45.0000	29646.3	29652.7	8299.01
60.0000	39633.0	39642.3	10867.1
80.0000	52712.9	52729.6	14793.5
90.0000	59370.5	59382.9	16697.7

Table 1: Sample Noise Test table from ESO, La SIlla CCD data set 33-94-12-12.

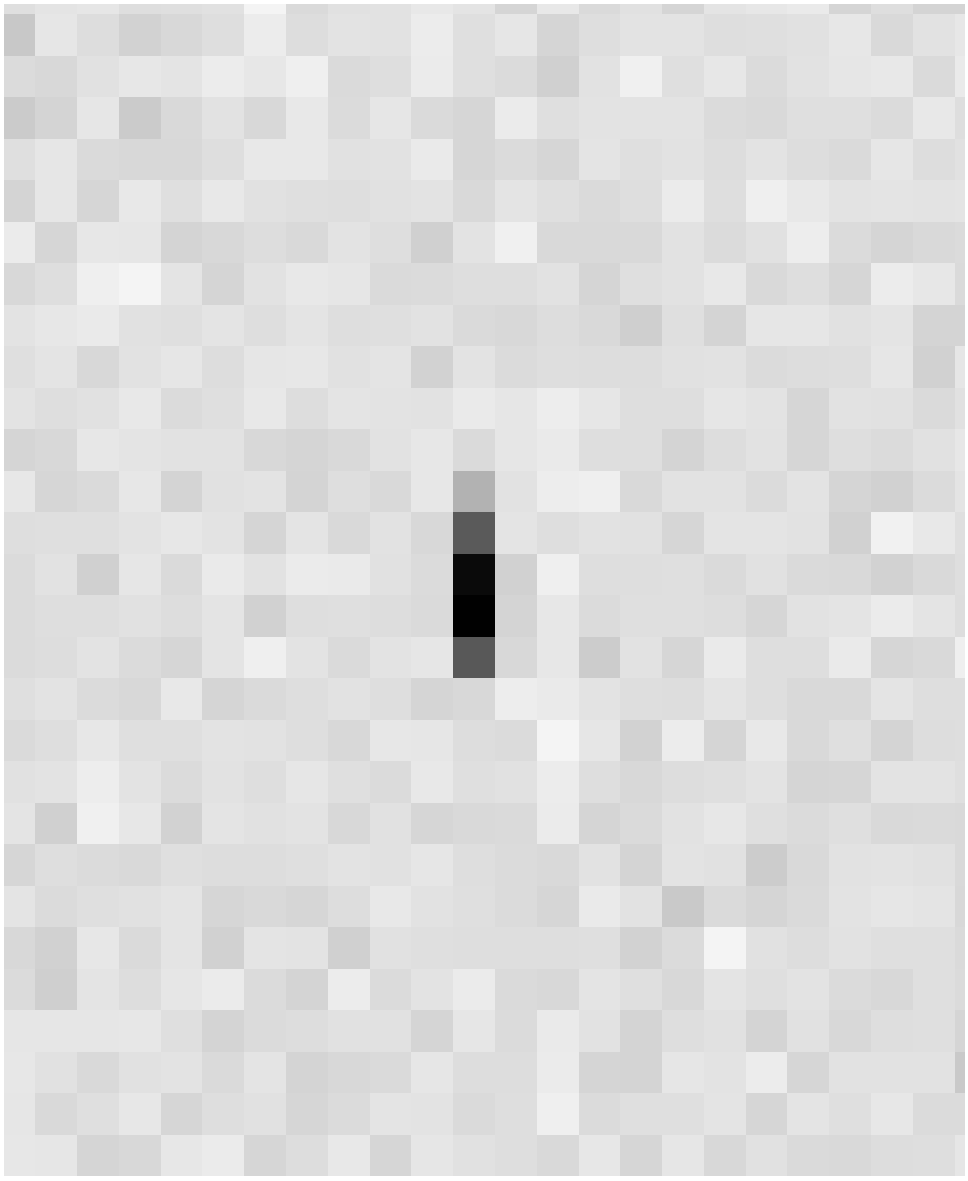

Figure 14: An example of a trap in a low count level image.

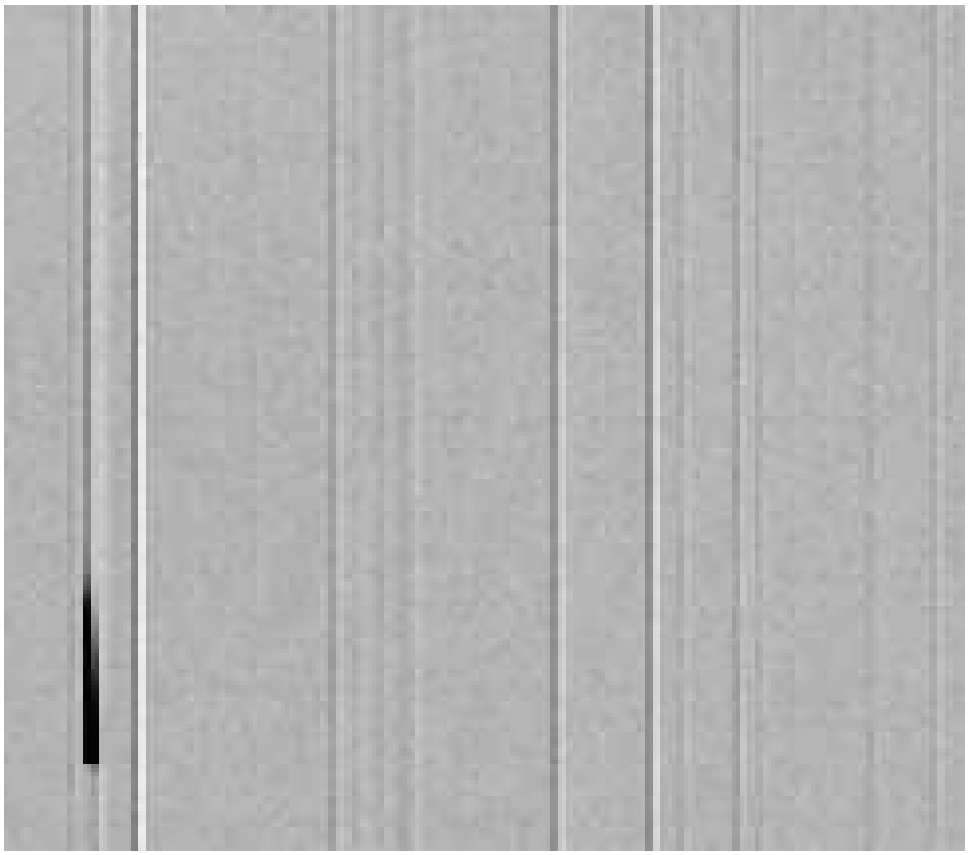

Figure 15: Example column offsets in a low count level flat from ESO CCD RCA#9. A large trap is also shown at bottom left.

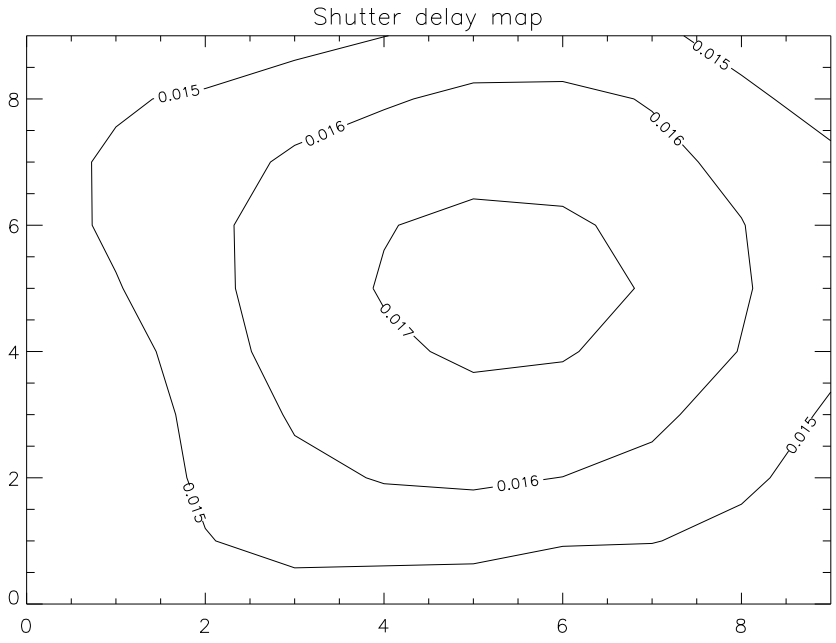

Figure 16: Example shutter delay map, taken from ESO, La Silla CCD test data set 33-94-12-12

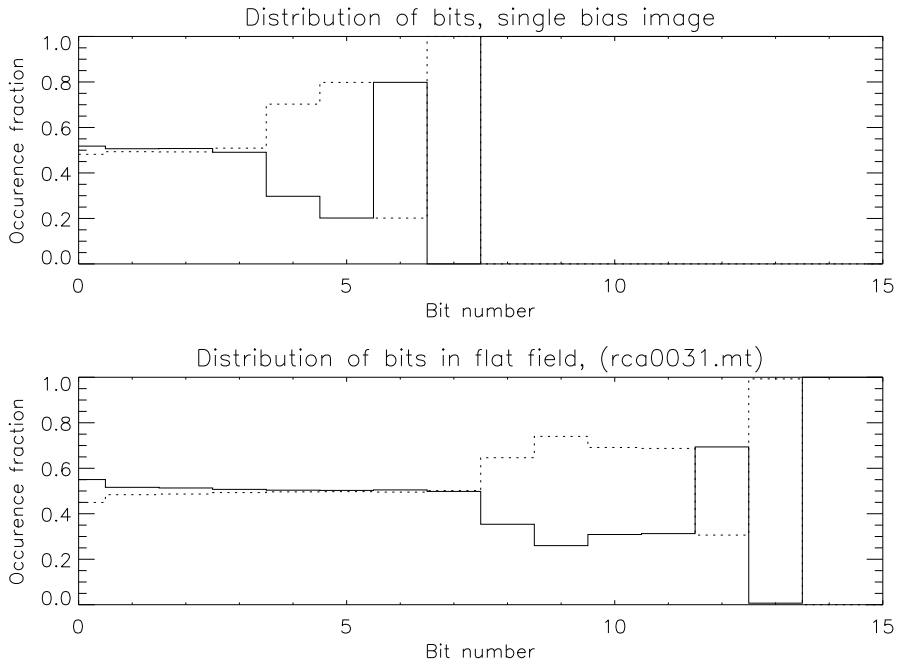

Figure 17: Example of distribution of 1's and 0's in a bias and well-illuminated flat field. The solid line indicates the fraction of bits which have a 0 value, the dotted line is the fraction of 1's. Note that the numbers of 1's and 0's are almost equal for the lower 7 bits in the well-illuminated flat (which are receiving essentially random data), but there is a slight bias towards more frequent 0's. The data were taken from the ESO, La Silla CCD test data set 09-94-10-23.