Chapter 3

Correcting for CCD non-linearities

3.1 Introduction

A charge coupled device (CCD) consists of an array of elements (pixels) arranged on a very thin silicon layer. Incident photons on each pixel are converted to electron-hole pairs. Up to 50 000–500 000 electrons can be stored in each pixel, depending on the CCD. After an exposure is finished, the electrons can be moved around and read-out by a controller which converts the electron charge in each pixel (or binned group of pixels) to digital counts (analogue-to-digital units or ADU). The conversion factor (electrons/ADU) is called the gain and is typically in the range 2–5.

CCDs are not perfectly linear systems as some people often assume. After bias subtraction, the number of ADU counts is not exactly proportional to the number of incident photons. The most probable reason for this non-linearity is that the controller gain is not a constant function of the number of electrons. When the CCD approaches saturation, there will be much larger non-linearities due to the reduced probability of capturing photons in nearly full pixels.

Correcting for non-linearities is important whenever a differential intensity measurement needs to be made. For example, measuring the equivalent width of lines in spectroscopic data or in most photometry. In general, non-linearities will be less important for Doppler shift measurements, except in the case of asymmetric line profiles where the measured wavelength displacement depends on the depth in the line.

3.1.1 Testing for non-linearities

To test for non-linearities, we need a fairly stable light source, for example either a lamp or clear daytime sky. There are two methods that I will discuss here, a bracketed repeat-exposure method (Section 4 from Gilliland et al. 1993) and a ratio method developed by myself.

The bracketed repeat-exposure (BRE) method involves making single and multiple exposures with the CCD. We do not change the exposure time as we cannot assume that the nominal exposure time is accurate, i.e., there may be an constant offset to any selected exposure time. Instead, we keep the exposure time constant and make repeat exposures of the CCD before reading-out. A typical example of the method consists of several read-outs with the no. of repeats = 0,1,2,1,3,1,4,1,5,1,6,1,7,1,8,1,9,1,10,1,0. The first and last read-outs are bias frames and the single exposures (bracketing the multiple exposures) are used to calibrate the change in the lamp's intensity, i.e., to determine the expected counts for the multiple exposures. Only a small region of the CCD is read out, such that the read-out time is only a few seconds, while the exposure time should be about two seconds. It is then possible to plot a curve of relative gain[1] (measured-counts / expected-counts) versus measured-counts for various regions on the CCD (note that all counts are bias-corrected first). For a linear CCD, the gain should be constant. Each plot can be normalised to have the same value of relative gain at a certain number of measured counts.

The ratio method is an indirect method and, as the name suggests, involves calculating the ratio between the measured light level on two regions of the CCD. By varying the exposure time, we can then make a plot of the ratio versus measured-counts (of one of the regions). For a linear CCD, the ratio should be constant. This method is unaffected by uncertainties in the exposure time or in the light level, except in the case of spectra, where changes in the temperature of the lamp will affect the ratio between some regions. To avoid this problem, half of the length of the slit can be covered with a filter to create a light level difference at each wavelength. These ratio measurements can be used to test a non-linearity curve obtained by the BRE method. Another good test involves taking daytime spectra to determine how the equivalent width of lines changes for various light levels. This is similar to the ratio method in that the equivalent width is a measure of the mean depth of a line relative to the continuum. Taking daytime spectra is a good way of testing to see if a particular type of measurement depends significantly on non-linearities.

---

[1] This 'relative gain' is proportional to ADU/electrons (inverse to the normal definition of gain).

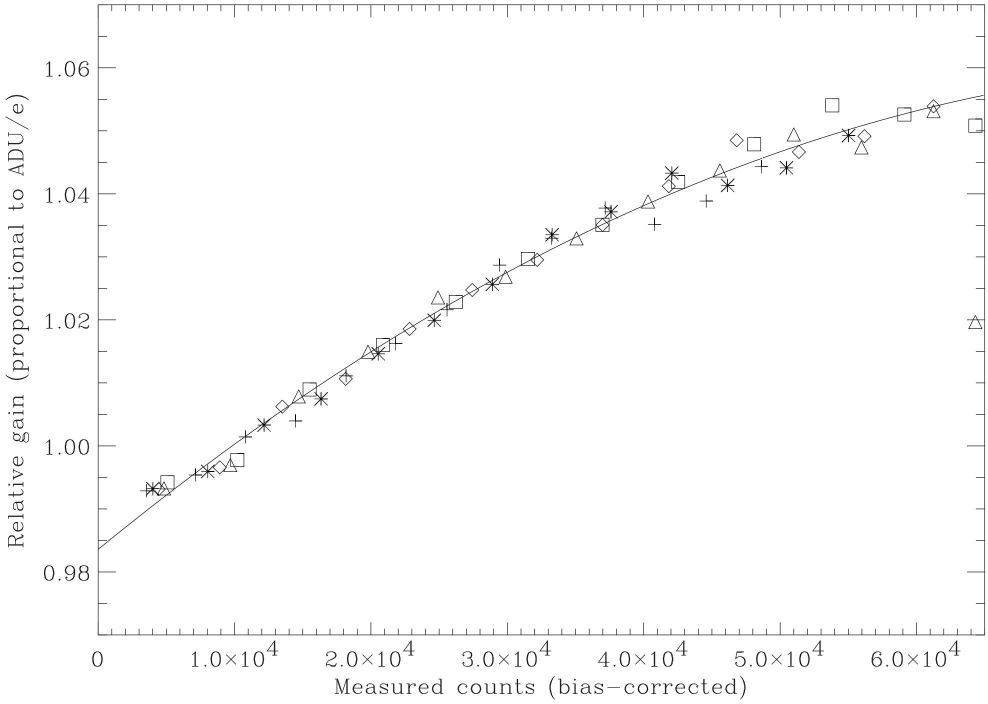

Figure 3.1 BRE measurement of non-linearity from a set of exposures in March 1997

3.2 Measurement of non-linearity using BRE method

BRE measurements of the non-linearity of the CCD system at the 74-inch Telescope at Mt. Stromlo Observatory were made in February and March 1997, along with an application of the ratio method where half of the slit was covered with a filter (Section 3.4). I used the 2Kx2K Tektronix chip (CCD10, serial number 1509BR24-01), with a nominal gain of 2e/ADU, with the B grating (dispersion 0.5Å/pixel) set up to look at the 6000Å to 7000Å wavelength region. A tungsten lamp was used for the light source.

Sixteen sets of repeat exposures were made for the BRE method as described in Section 3.1.1, with the maximum number of repeats ranging from 10 to 16. For each set of repeat exposures, plots of non-linearity (measured-counts / expected-counts versus measured-counts) were made using five different regions on the CCD. A straight line was fitted to each plot across the region from 0 to 30000 measured-counts so that each data set could be normalised to a relative gain of 1.0 at 10000 counts. For each set, a quadratic fit was made using all the five plots. Figure 3.1 shows one of the sixteen sets with the five plots combined, this set has the best quadratic fit. There is a 7% change in the gain between 0 and 64300 counts when the detector becomes digitally saturated (the bias level is $\sim$1200 ADU and the digital saturation before bias correction is 65535 ADU). The large non-linearity in this case may be because the CCD is only of engineering-grade; there are CCDs with a change in gain of less than 0.5% across a similar range.

Next, I combined the best five sets, which were all from March 1997, to produce a final

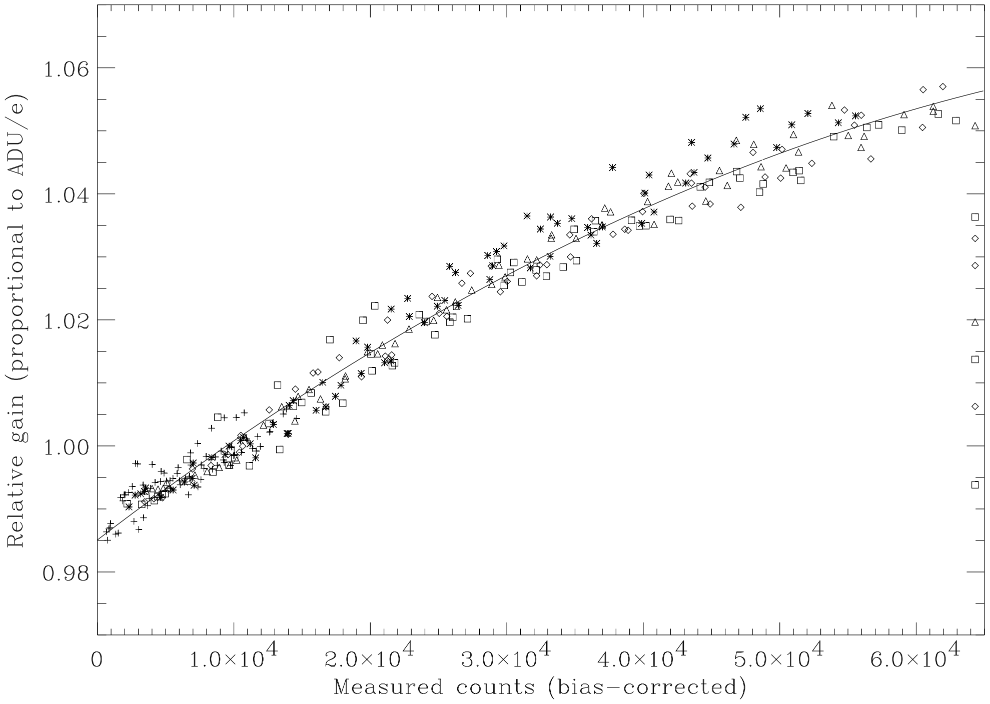

Figure 3.2 BRE measurement of non-linearity from five sets of repeat exposures

plot. For unknown reasons, the BRE measurements of non-linearity from February 1997 were not as good. Figure 3.2 shows the final plot with a quadratic fit to the data points. 3rd and 4th order polynomial fits to the data were only slightly better than a quadratic fit. The coefficients of different fits are shown in Table 3.1. In order to convert data to expected-counts, we divide the measured-counts by the relative gain.

Table 3.1 Coefficients of polynomial fits to non-linearity measurements

Polynomial	Coeff.				
1st order	0.989	1.17E−06			
2nd order	0.985	1.66E−06	−8.65E−12		
3rd order	0.987	1.32E−06	5.63E−12	−1.60E−16	
4th order	0.988	8.22E−07	4.28E−11	−1.12E−15	7.90E−21

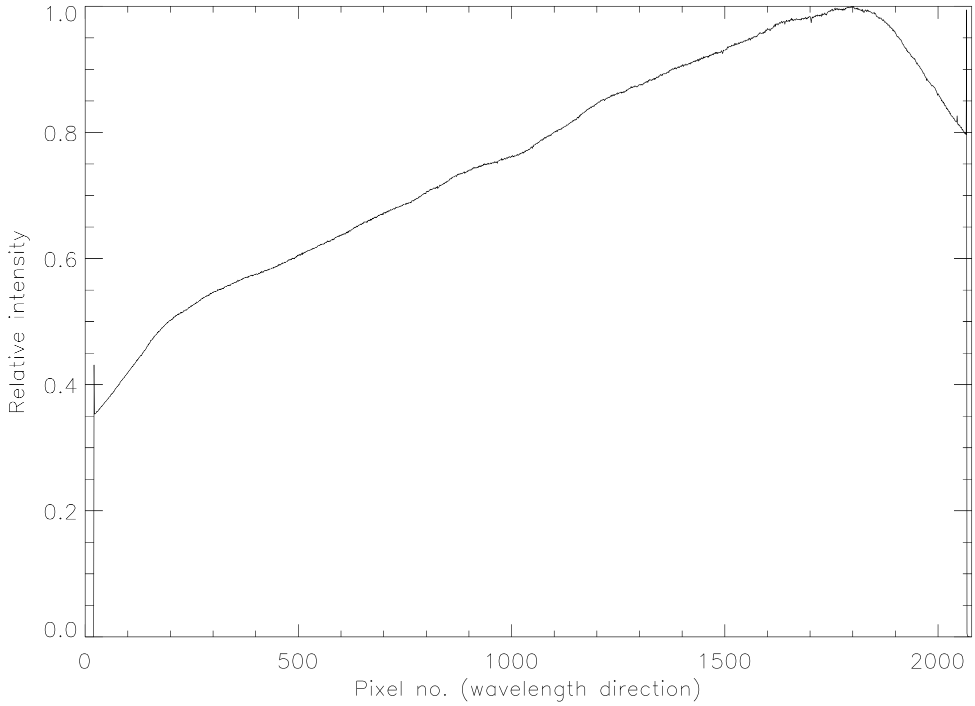

Figure 3.3 Cross-section of flat-field from May 1996

3.3 Stability of non-linearity curve

In this section, I compare non-linearity data taken in May 1996 and February 1997 using the ratio method in order to test the stability of the CCD10 system. The BRE method and the ratio method using a filter are more accurate, but no measurements were made in May 1996. However, flat-field spectra of various exposure times were taken and can be analysed using the ratio method. Figure 3.3 shows a cross-section of the tungsten lamp flat-field images that I used for the non-linearity tests. The aim is to test whether the non-linearity curve measured in 1997 can be applied to the 1996 data.

I measured the ratios between certain regions of the flat-field images and the highest intensity part of the images. For a linear detector and a lamp with a stable temperature, the ratio between two fixed regions in the flat-field spectrum should not depend on the actual light level. Figures 3.4–3.8 show the non-linearity ratio plots for the two years at ratios of approximately 0.450, 0.585, 0.680, 0.810 and 0.915. The scatter in the plots are mainly due to changes in temperature of the flat-field lamp which produces a variable intensity gradient across the spectrum. However, there is clearly a decrease in the measured ratio at higher light levels due to the non-linearity of the CCD. This is consistent with an increase in the relative gain of the CCD at higher light levels, as seen using the BRE method. At a ratio of approximately 0.450 (see Figure 3.4), the decrease in the ratio between light levels of 5000 counts and 60000 counts is about 2%. Note that at near 64300 counts the ratio can go up because the ADU counts are saturated for the higher value. The shapes of

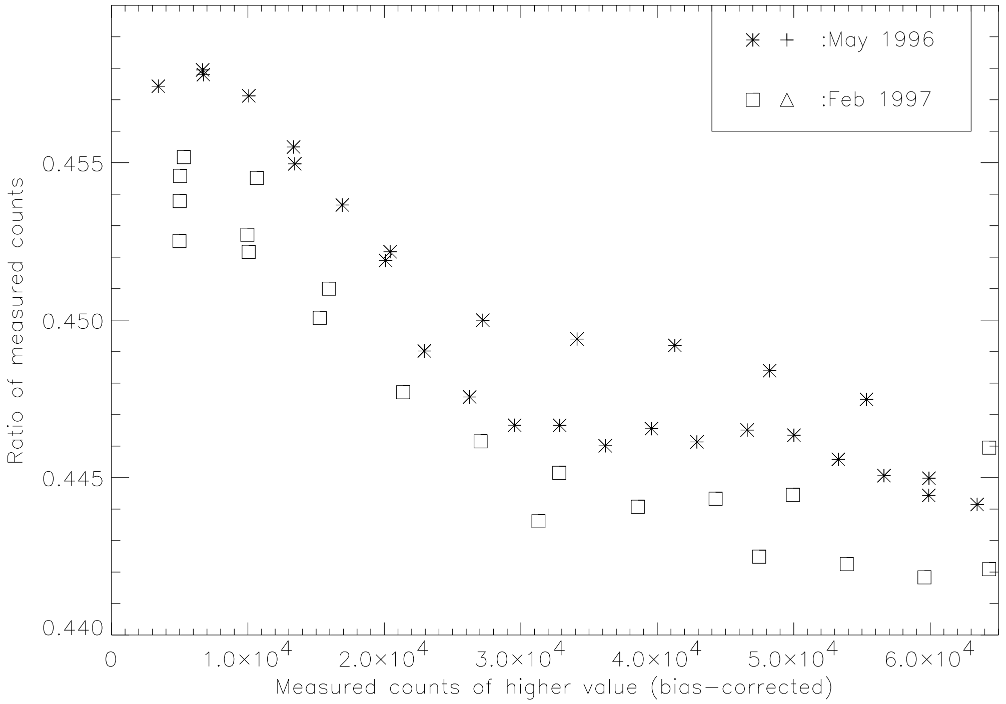

Figure 3.4 Comparison of non-linearity data using ratio method

the ratio curves are approximately the same for May 1996 and for February 1997 which means that I cannot detect any change in the non-linear behaviour of the CCD. This is evident in Figures 3.4–3.6 when the ratio change is large. I assumed that there was no significant difference between the non-linearity from 1996 and from 1997 and, therefore, used the non-linearity curve measured in 1997 to correct for the non-linearities in the 1996 data.

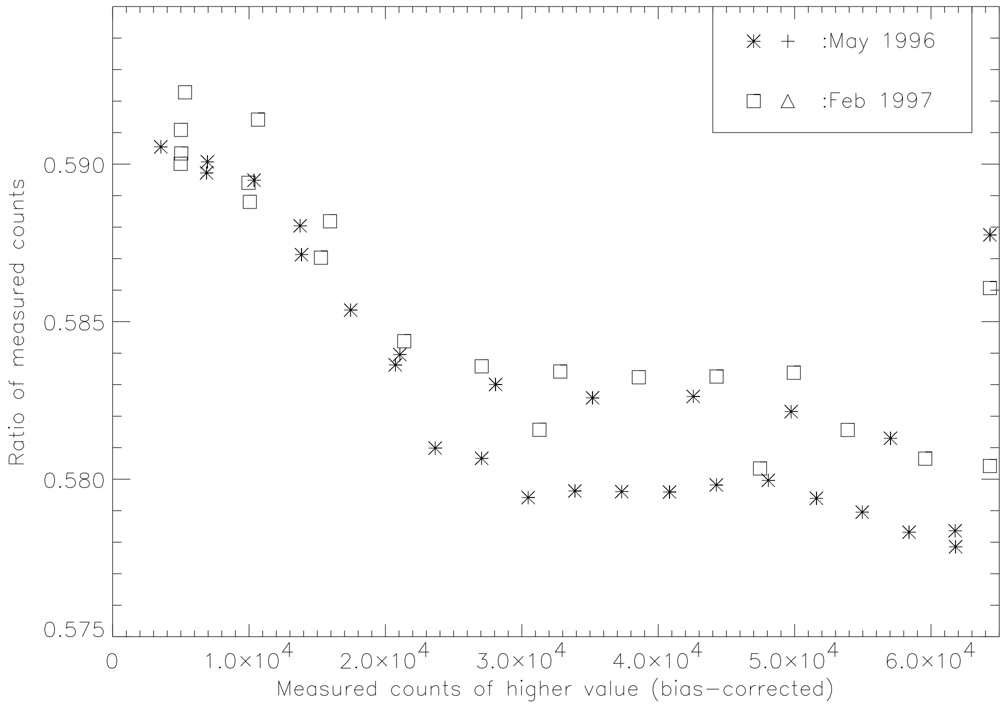

Figure 3.5 Comparison of non-linearity data using ratio method

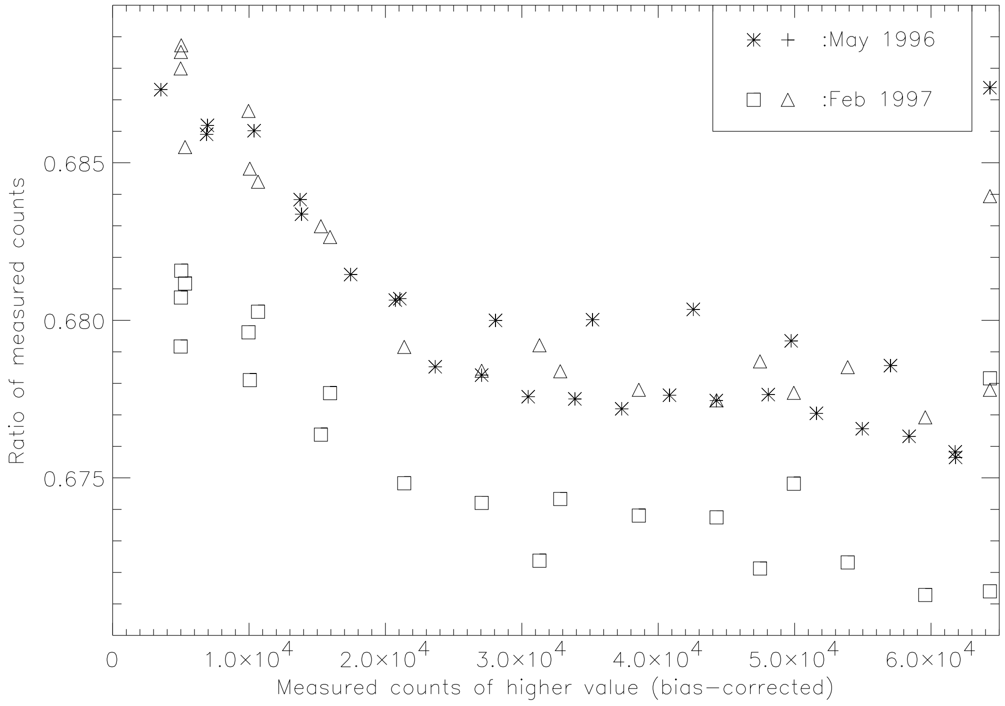

Figure 3.6 Comparison of non-linearity data using ratio method

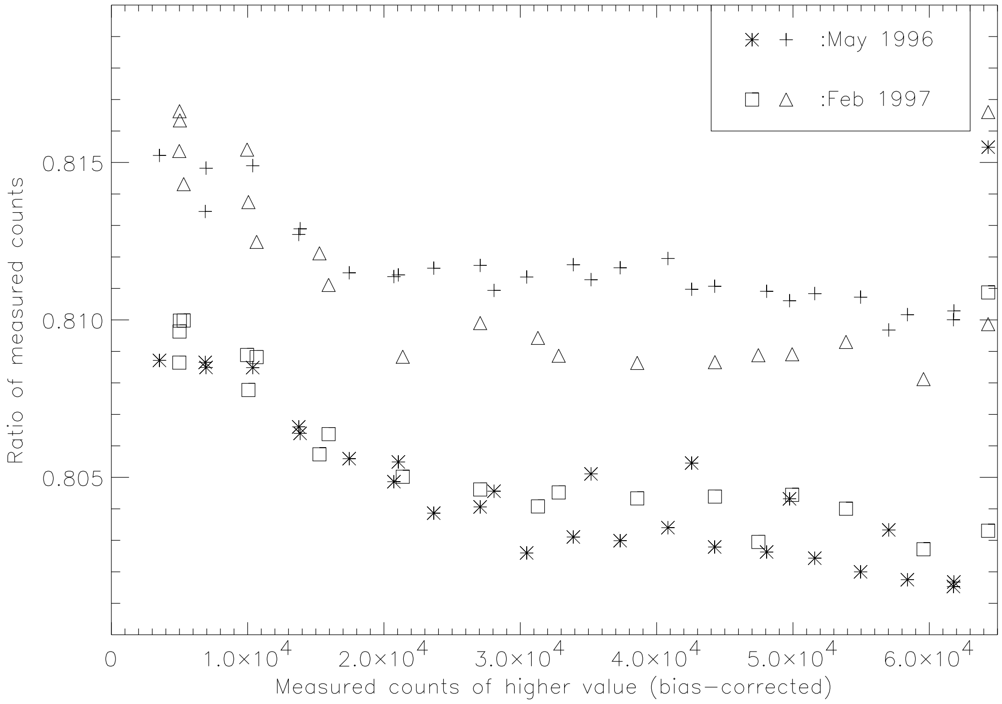

Figure 3.7 Comparison of non-linearity data using ratio method

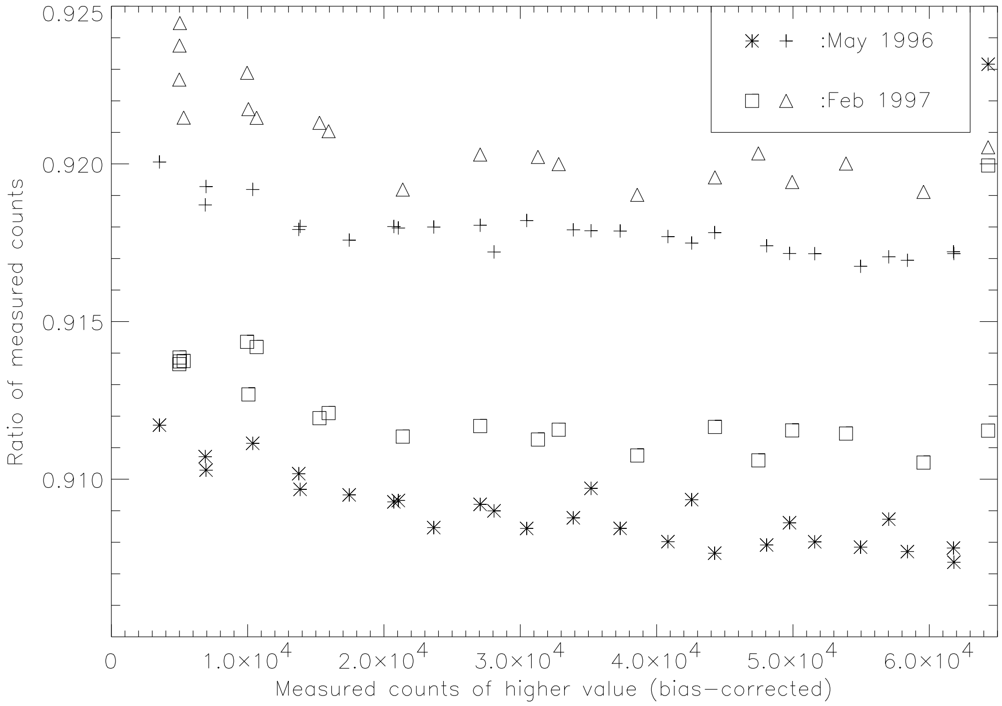

Figure 3.8 Comparison of non-linearity data using ratio method

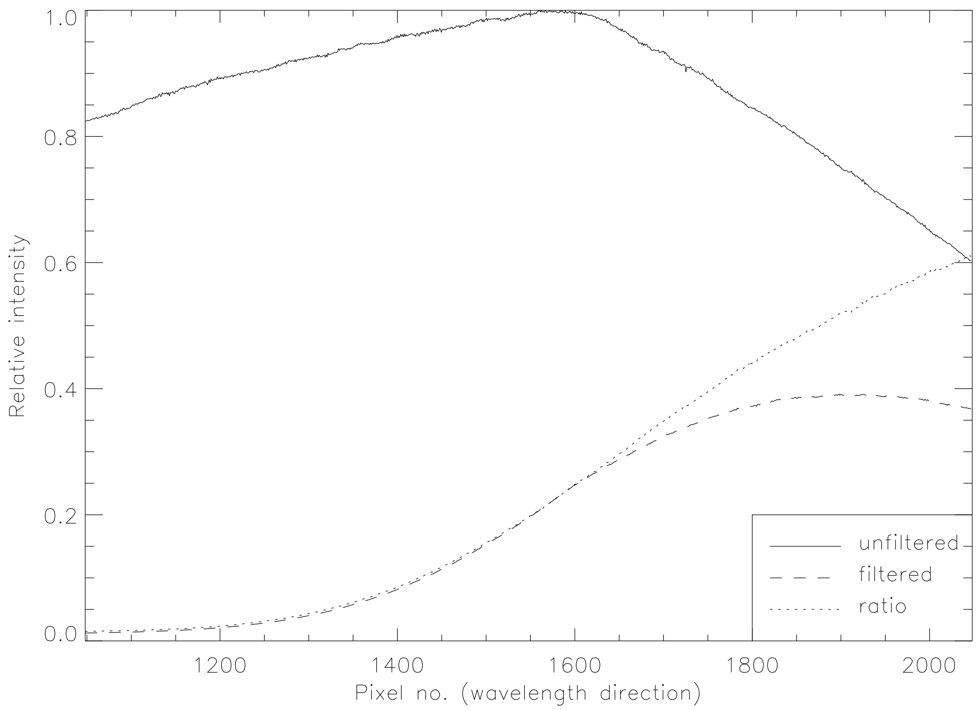

Figure 3.9 Cross-sections of flat-field from February 1997. The solid line represents the direct spectrum of the lamp with vignetting above pixel 1600. The dashed line represents a parallel cross-section where the light has passed through a colour filter. The dotted line represents the transmission profile of the filter.

3.4 Checking the non-linearity curve

Ratio measurements using a colour filter, to cover half of the length of the slit, were made in 1997 to check the validity of the non-linearity curve of CCD10 derived using the BRE method. Figure 3.9 shows two cross-sections of a flat-field image from a filtered and an un-filtered part and the ratio between the two parts. In Figures 3.10–3.14, the ratios between the filtered and un-filtered parts of the image, at a certain wavelength, versus counts of the un-filtered part are plotted. The same measurements were made on the data before and after the 2nd order non-linearity correction given in Table 3.1.

There is a definite improvement in all the ratio changes by a factor of about 5 after the non-linearity correction has been made. However, this ratio method has a very low scatter and seems to be detecting higher order non-linearities (Figures 3.13–3.14) which were not evident using the BRE method. Note that at light levels below 10000 counts, the scatter is much higher, probably because of errors in bias subtraction. I have ignored this region for the purposes of this analysis. From Figures 3.10–3.11, it appears that the non-linearity correction is working quite well but perhaps slightly under-correcting at high light levels. I have tried re-reducing the BRE method by using different normalisations and different sets of exposures but the small discrepancy remains.

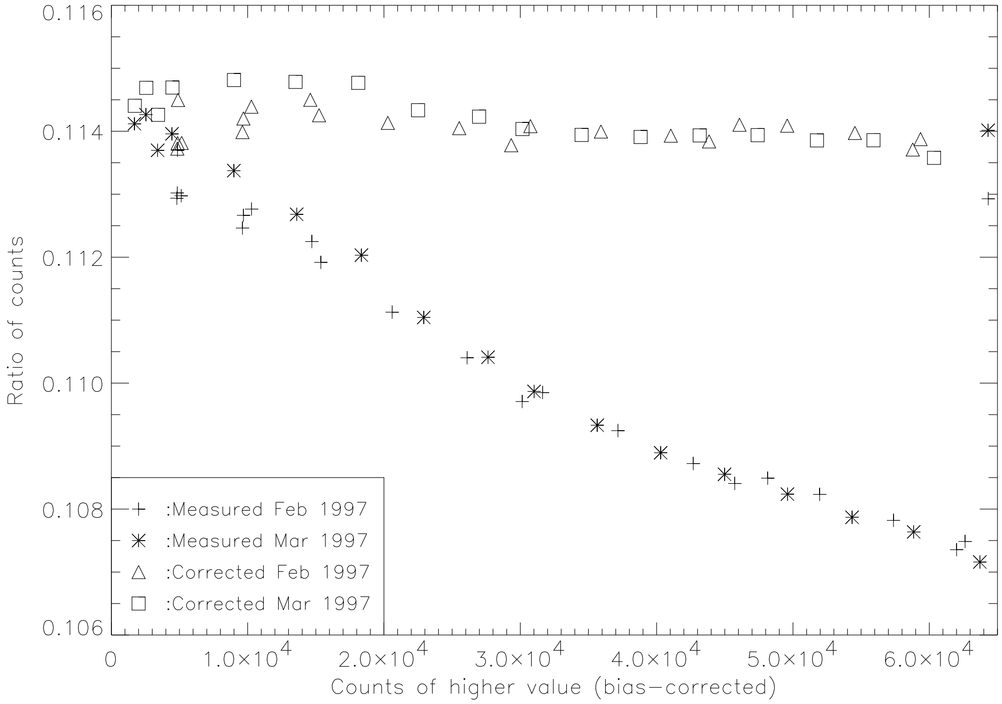

Figure 3.10 Testing of non-linearity curve using ratio method

Structure in the non-linearity curve appears around 18000 and 31000 counts, as is most evident in Figures 3.12–3.14. The obvious change in ratio at 31000 counts in Figure 3.14 is interesting because at this ratio (∼0.58), there is structure in the non-linearity curve around the lower value (18000) as well as the higher value. Using 3rd and 4th order non-linearity corrections from the BRE method only slightly improves the glitches in the ratio curves. This is not surprising as a polynomial fit is not a good representation of what is occurring in the non-linearity curve at these points. It is also likely that the BRE method has smoothed over these high-order changes because of the normalisation and the scatter in the plots. Despite these glitches, the 2nd order fit derived by the BRE method has improved the non-linearity significantly and, therefore, it is valuable to use this non-linearity correction on data obtained with this CCD.

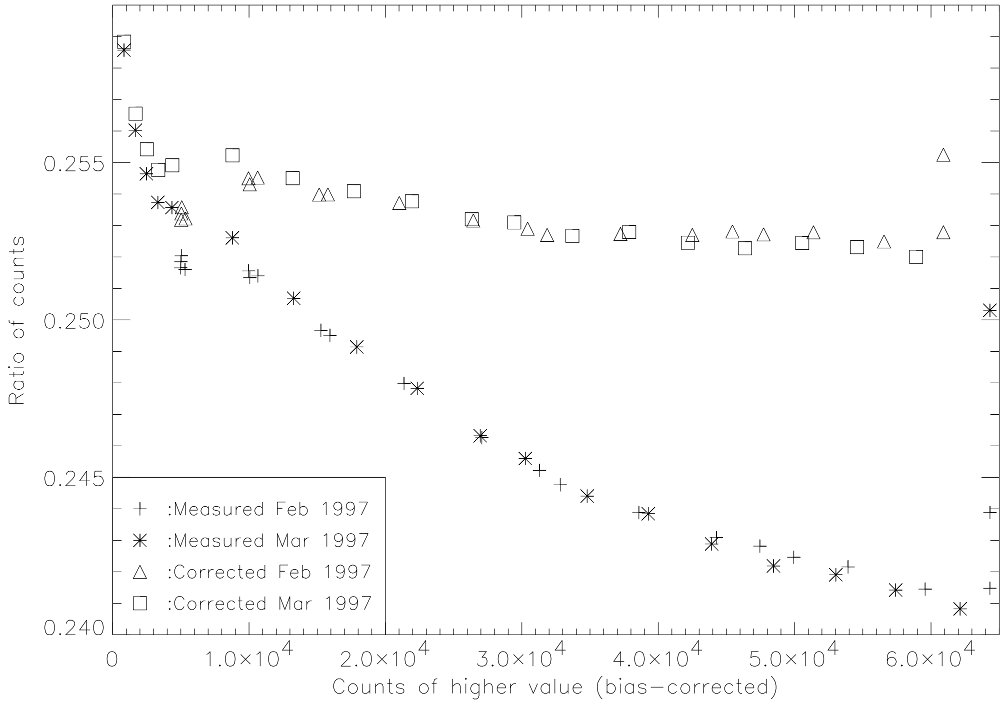

Figure 3.11 Testing of non-linearity curve using ratio method

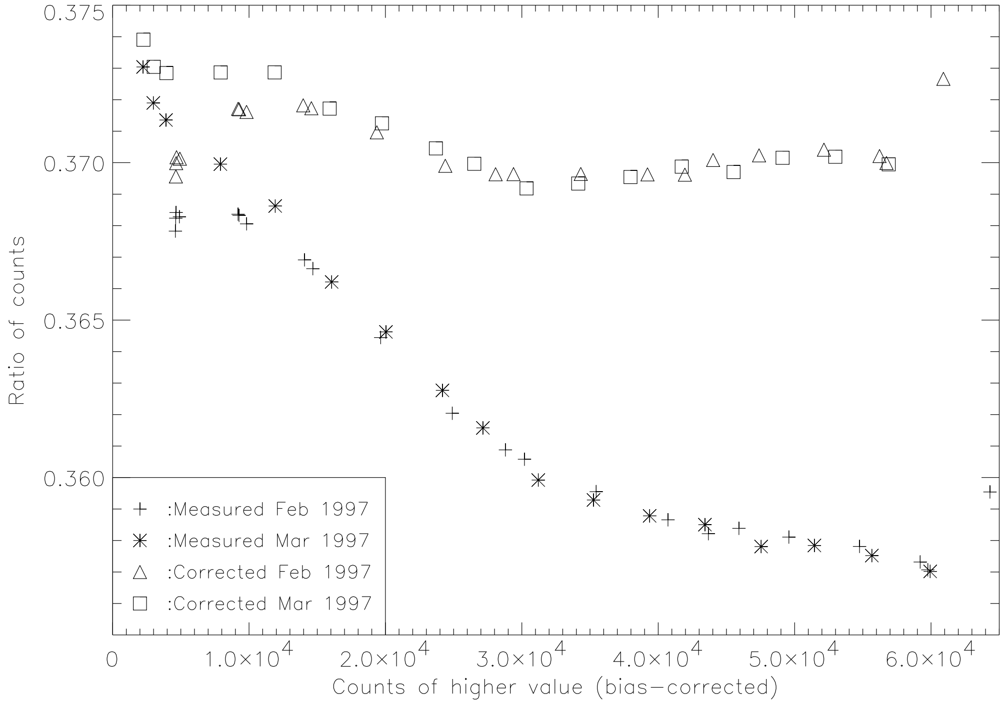

Figure 3.12 Testing of non-linearity curve using ratio method

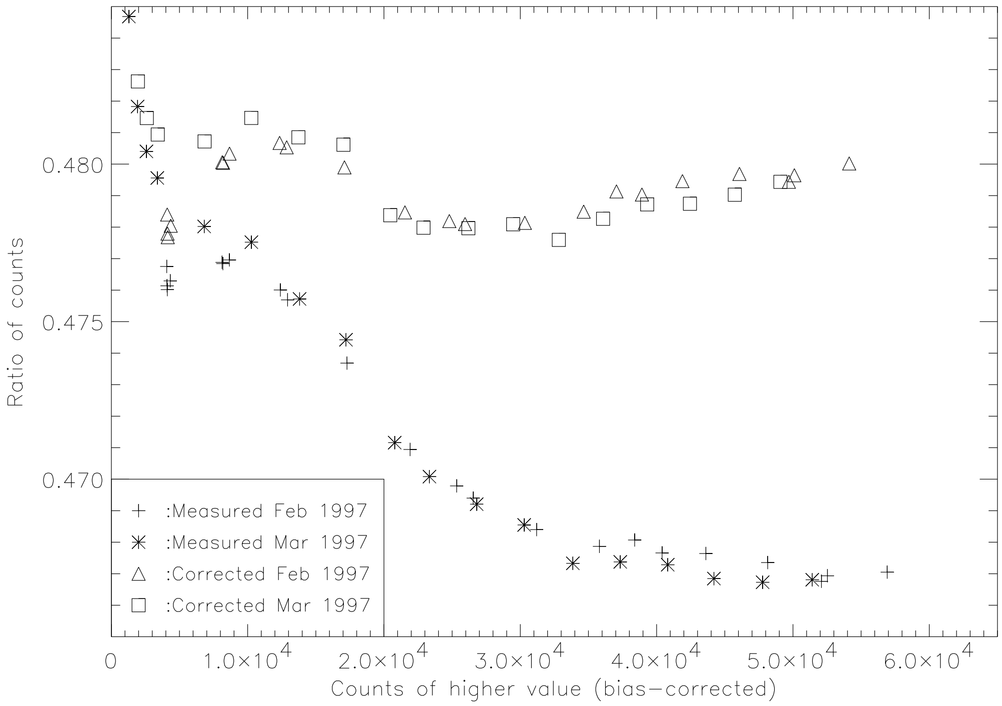

Figure 3.13 Testing of non-linearity curve using ratio method

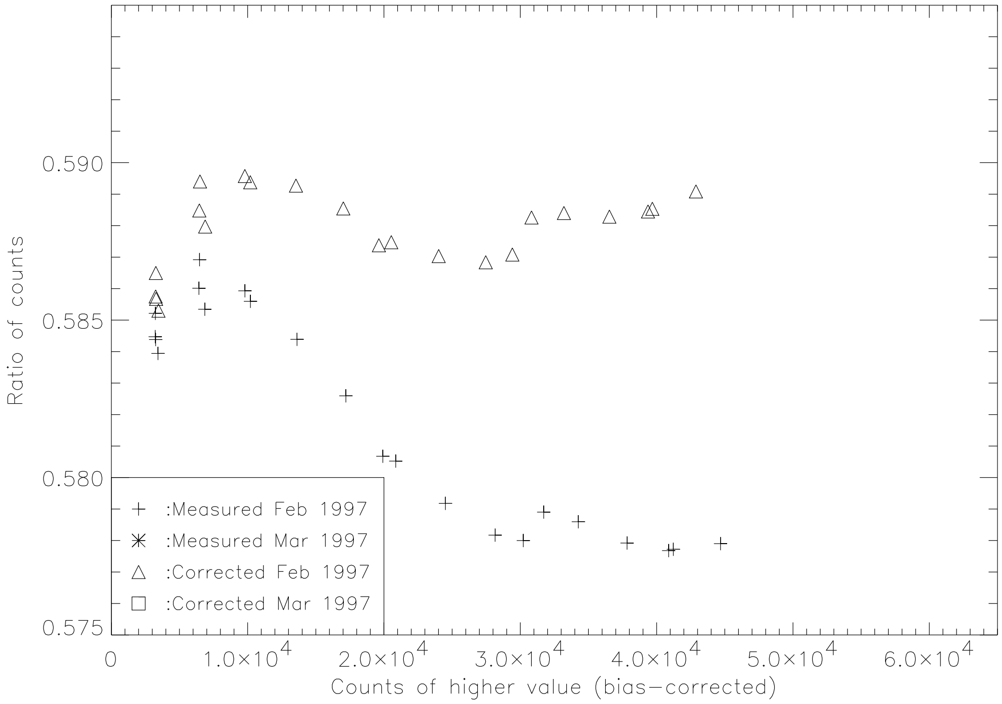

Figure 3.14 Testing of non-linearity curve using ratio method

3.5 Measurement of non-linearity using ratio method

It is possible that the ratio method can provide a better non-linearity curve, for example, by solving an equation of the form

$$R_{12} = \frac{N_1/f(N_1)}{N_2/f(N_2)} = \frac{N_1'/f(N_1')}{N_2'/f(N_2')} = \frac{N_1''/f(N_1'')}{N_2''/f(N_2'')} = \ldots\ldots, \tag{3.1}$$

where N_1 and N_2 are the measured counts of the two regions, at different light levels N N' N'' etc., $f(N)$ is the relative-gain function equivalent to measured-counts / expected-counts and, R_{12} is the ratio of the expected (or true) counts between the two regions (the corrected ratio). A fit can be made to the parameters chosen for the relative-gain function, by minimising the scatter in the corrected ratio between different light levels. Similar equations can be solved for different pairs of regions on the CCD. For low intensity non-linearity tests, it will be necessary to obtain a good overscan region for accurate bias subtraction. A non-linearity measurement of a CCD using the ratio method is described in this section.

In December 1998, I observed at Mt. Stromlo as part of a multi-site campaign to study the roAp star HR 1217 (organised by S. Frandsen et al.). A different CCD was used (2Kx4K SITe chip: CCD17, serial number 6044FCD04-01, gain 2.5 e/ADU, 1x2 binning), from the one that was used in 1996 and 1997 to study α Cir and HR 3831. For this campaign, it was important to measure the non-linearity of this CCD. I used primarily the ratio method and fitted only one parameter, the α parameter (Tinney 1996), to the data. This demonstrates that the ratio method can be used as an accurate and independent method to determine the non-linearity of CCDs.

The spectrum of a flat-field lamp was projected onto the CCD, using the B-grating of the coudé spectrograph in the range 6300–6900Å (dispersion 0.31Å/ pixel). A colour filter (BG 38) was placed directly over the slit, covering about half the length of the slit. This was secured with tape to avoid any movement of the filter during the measurements. The result of placing this filter was that a filtered spectrum and a direct spectrum were both projected onto the CCD. The ratio between the intensities in each spectrum varied between 0.1 and 0.4 depending on the wavelength. If the CCD is a perfectly linear system, the ratio between the intensities (bias-corrected) at each wavelength should remain the same and not depend on the exposure time or changes in the lamp flux.

17 exposures were taken with exposure times varying between 0.1 s and 6 s. The exposure times were first increased and then decreased to check if there was a systematic change with time. An average bias frame was subtracted from all the exposures, to remove spatial structure due to bias, and then overscan regions of the CCD were used to fine tune the bias correction. Six pairs of regions, each pair consisting of a filtered and a direct region, were chosen equally spaced along the dispersion direction (each region consisted of about 300 pixels). The mean intensity was measured for each region and exposure. The measured intensity ratios for each pair of regions are shown in Figures 3.15–3.20.

Tinney (1996) described the AAO CCD non-linearities in terms of an α parameter.

$$N_m = N_t(1 + \alpha N_t) \tag{3.2}$$

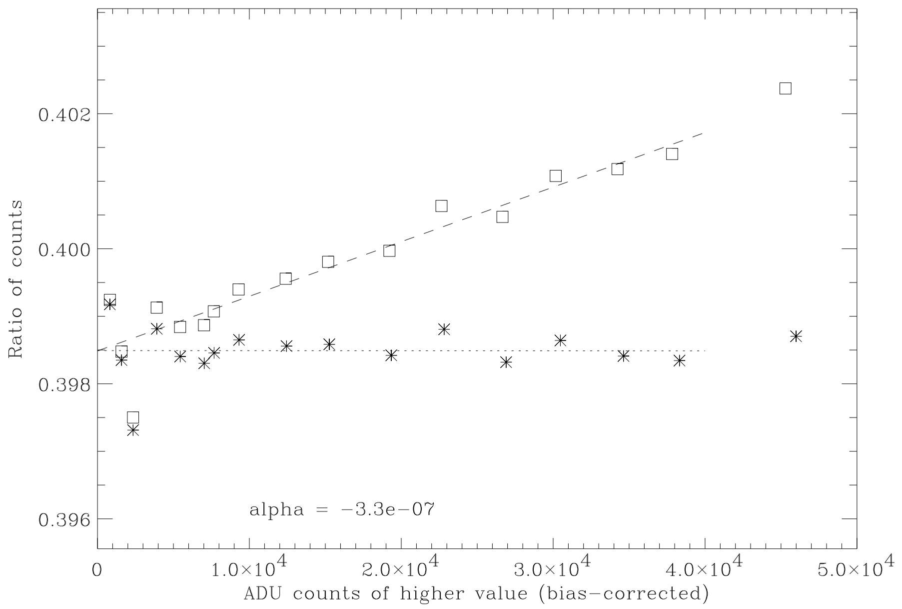

Figure 3.15 Non-linearity measurement using ratio method. The squares represent the measured ratios, while the asterisks represent the ratios after correcting the intensities using the alpha parameter. The lines are best fits to each set of ratios, between 0 and 40 000 counts, with lower weight given to measurements below 10 000. The alpha value has been chosen so that the slope of the best fit to the corrected ratios is zero.

where N_m are the measured counts in ADU above the bias level and N_t are the 'true' counts (normalised so that $N_t = N_m$ for $N_t \to 0$). The non-linearity of CCD17 was assumed to be represented by a single parameter similar to that defined by Tinney, except that the true intensity is defined in terms of the measured counts; $N_t = N_m/(1 + \alpha N_m)$. To first order, it makes no difference to the value of α.

For each set of ratio measurements, α was varied until the best fit for the corrected ratios had a slope of zero. The fit was obtained using the ratios with counts of the higher-value between 0 and 40 000. Lower weight was given to those with counts below 10 000 because of increased noise. Figures 3.15–3.20 show the value of α and the corrected ratios, for each set of measurements. For CCD17, $\alpha = -3.5 \pm 0.2 \times 10^{-7}$ from Figs. 3.15–3.18, where the correction factor is well defined. The α parameter is less well defined from Figs. 3.19–3.20 because of higher noise factors.

The ratio method removes the problems of requiring accurate exposure times and lamp-temperature stability to make accurate non-linearity measurements. Further improvement of the accuracy of the method described in this section could be made by (i) taking more exposures, (ii) increasing and decreasing the exposure time several times, and (iii) interspersing the exposures with bias frames. In the second case, this will reduce problems which might arise from a systematic change in the measurements with time. In the third case, monitoring any changes in bias frames will improve the accuracy of bias subtraction which is critical for measurements with low counts.

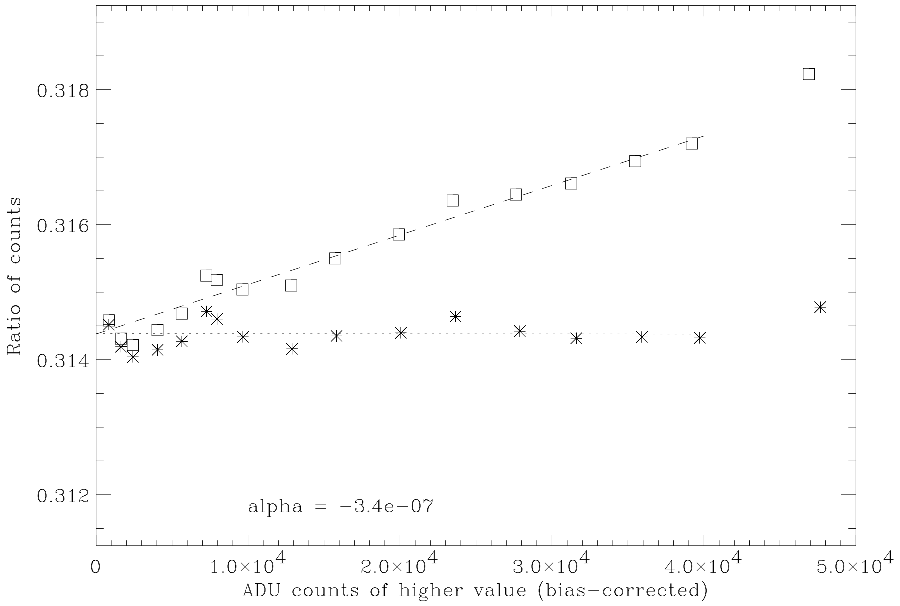

Figure 3.16 Non-linearity measurement using ratio method. See Fig. 3.15 for details.

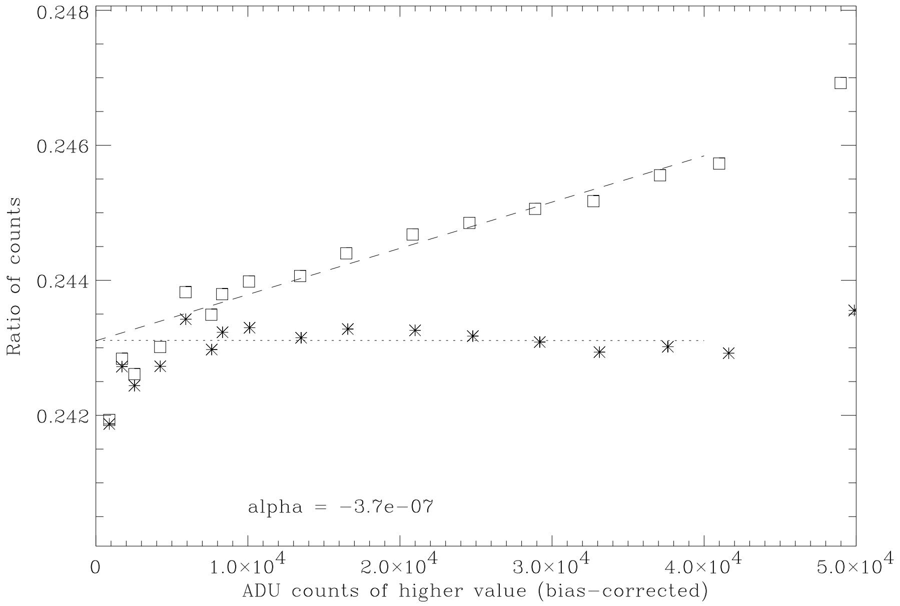

Figure 3.17 Non-linearity measurement using ratio method. See Fig. 3.15 for details.

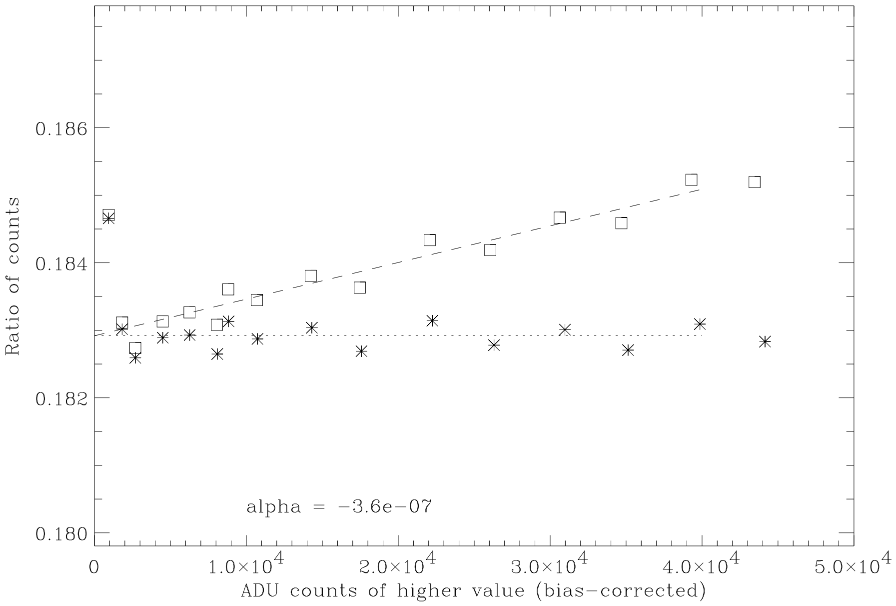

Figure 3.18 Non-linearity measurement using ratio method. See Fig. 3.15 for details.

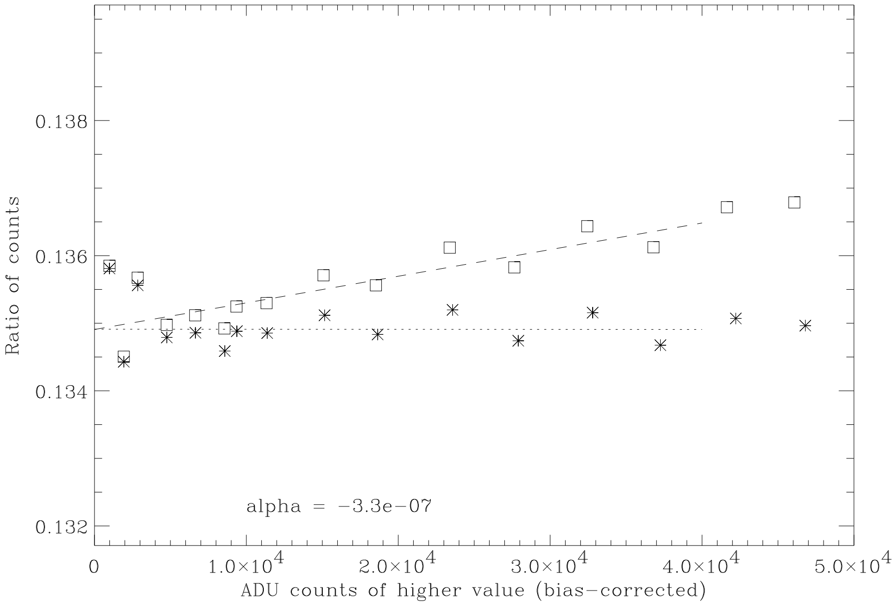

Figure 3.19 Non-linearity measurement using ratio method. See Fig. 3.15 for details.

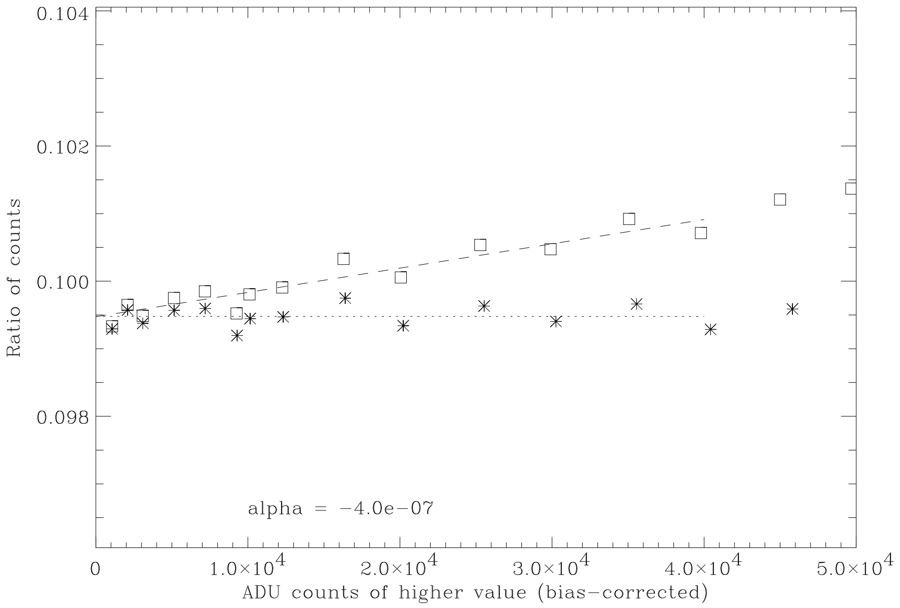

Figure 3.20 Non-linearity measurement using ratio method. See Fig. 3.15 for details.

3.6 Discussion

In Sections 3.2–3.4, I looked at the non-linearity of the Tek CCD10 at Mt. Stromlo (serial number 1509BR24-01, gain 2.0 e/ADU). This is the CCD with which I obtained most of the data for this thesis. In this section, I compare the non-linearity of this CCD with other CCDs.

For CCD10, the non-linearity cannot be accurately quantified using the α parameter due to higher-order effects. The measured value of α varies between 1.17×10^{-6} and 1.29×10^{-6}, depending on the range of counts considered (0–25000 up to 0–60000). Tinney (1996) quotes an α value of -3×10^{-8} for the AAO 1Kx1K Tek #2 chip (commissioned in July 1992) in 'normal' mode (gain 2.7 e/ADU). The linearity for this Tek CCD is a factor of 40 better than for CCD10 at Mt. Stromlo, the difference being due to the grade of the chip. Top science-grade CCDs should have a non-linearity parameter of $|\alpha| < 10^{-7}$ with a gain of about two electrons/ADU. This means that at an ADU level of 50 000, the correction is less than 0.5%. Note that, for a given CCD, the α parameter will be proportional to the gain (e/ADU) if the non-linearity is only a function of the number of electrons and not a function of the number of electrons and the gain.

In 1998, I observed at Mt. Stromlo using the SITe CCD17 (serial number 6044FCD04-01, gain 2.5 e/ADU). The α parameter was measured to be -3.5×10^{-7} in the range 0–40 000 ADU (Section 3.5). The CCD is of moderate science-grade with the non-linearity well characterised by one parameter, which means the non-linearity can be easily and accurately corrected when high-precision measurements need to be made.

The CCD17 measurements were taken using 1x2 binning to match the spectral data being taken during the observing run. The saturation level was determined to be around 57 000 ADU which is approximately double the saturation level when no binning is used. This is because the serial-register pixels of the CCD, where the electrons are combined, have a higher electron capacity. In this sense, the non-linearity is different when the CCD is binned. The α parameter should not depend on the binning because it is related to the controller conversion of electrons to ADU.

CCD10 has a non-linearity characterised by an α parameter of around $+1.2 \times 10^{-6}$, which is significantly non-linear compared to science-grade CCDs. Therefore, in order for this CCD to perform adequately in terms of high precision spectroscopy, it is necessary to apply a non-linearity correction. I used the 2nd order fit derived by the BRE method (Table 3.1) to correct the CCD data taken from Mt. Stromlo, for the results in Chapters 5–7. It would be difficult to make improvements on this correction because the non-linearity for this CCD is not easily quantified by two or three parameters.

In this chapter, I have demonstrated two techniques for measuring or checking the non-linearity of CCDs. The BRE method measures the variation in the intensity of a region on the CCD (using multiple exposures bracketed by single exposures to monitor any changes in the lamp's intensity). The ratio method measures the variation in the ratio between the intensities of two regions on the CCD. This can provide a more accurate non-linearity curve because it is less affected by changes in the lamp's flux and is unaffected by uncertainties in the exposure time.