

HANDBOOK OF CCD ASTRONOMY

STEVE B. HOWELL

*ASTROPHYSICS GROUP, PLANETARY SCIENCE INSTITUTE, TUCSON,
AND DEPARTMENT OF PHYSICS AND ASTRONOMY,
UNIVERSITY OF WYOMING*



CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS

The Edinburgh Building, Cambridge CB2 2RU, UK <http://www.cup.cam.ac.uk>
40 West 20th Street, New York, NY 10011-4211, USA <http://www.cup.org>
10 Stamford Road, Oakleigh, Melbourne 3166, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain

© Cambridge University Press 2000

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2000

Printed in the United States of America

Typeface Computer Modern 10/13 pt. *System* L^AT_EX [TB]

*A catalog record for this book is available from
the British Library.*

Library of Congress Cataloging-in-Publication data is available.

ISBN 0 521 64058 X hardback

ISBN 0 521 64834 3 paperback

Contents

<i>Preface</i>	<i>page</i>	<i>ix</i>
1 Introduction		1
1.1 Nomenclature		2
1.2 Why Use CCDs?		2
2 CCD Manufacturing and Operation		7
2.1 CCD Operation		8
2.2 CCD Types		12
2.3 CCD Coatings		21
2.4 Analog-to-Digital Converters		22
3 Characterization of Charge-Coupled Devices		26
3.1 Quantum Efficiency		26
3.2 Readout Noise		30
3.3 Dark Current		32
3.4 CCD Pixel Size, Pixel Binning, Full Well Capacity, and Windowing		35
3.5 Overscan and Bias		37
3.6 CCD Gain and Dynamic Range		39
3.7 Summary		46
4 CCD Imaging		47
4.1 Image or Plate Scale		47
4.2 Flat Fielding		48
4.3 Calculation of Read Noise and Gain		52
4.4 Signal-to-Noise Ratio		53
4.5 Basic CCD Data Reduction		58
4.6 CCD Imaging		61
5 Photometry and Astrometry		75
5.1 Stellar Photometry from Digital Images		77
5.2 Two-Dimensional Profile Fitting		85

5.3	Aperture Photometry	88
5.4	Absolute versus Differential Photometry	91
5.5	Astrometry	94
5.6	Pixel Sampling	96
6	Spectroscopy with CCDs	100
6.1	Review of Spectrographs	102
6.2	CCD Spectrographs	104
6.3	CCD Spectroscopy	109
6.4	Signal-to-Noise Calculations for Spectroscopy	113
6.5	Data Reduction for CCD Spectroscopy	114
6.6	Extended Object Spectroscopy	118
6.7	Slitless Spectroscopy	121
7	CCDs Used in Space and at Short Wavelengths	128
7.1	CCDs in Space	130
7.2	Radiation Damage in CCDs	136
7.3	CCDs in the UV and EUV (300–3,000 Å) Spectral Range	138
7.4	CCDs in the X-Ray (<500 Å) Spectral Range	140
Appendix A	CCD Reading List	143
Appendix B	CCD Manufacturers and CCD Website Information	146
Appendix C	Some Basics of Image Displays	150
	References	155
	Index	161

3

Characterization of Charge-Coupled Devices

Even casual users of CCDs have run across the terms read noise, dark current, signal-to-noise, and many other possibly mysterious sounding bits of CCD jargon. This chapter will discuss the meanings of the terms used to characterize the properties of CCD detectors. Techniques and methods by which the reader can determine some of these properties on their own and why certain CCDs are better or worse for a particular application are discussed in the following chapters. Within the discussions, mention will be made of older types of CCDs. While these are generally not available or used anymore, there is a certain historical perspective to such a presentation and it will likely provide some amusement for the reader along the way.

One item to keep in mind throughout this chapter and in the rest of the book is that all electrons look alike. When a specific amount of charge is collected within a pixel during an integration, one can no longer know the exact source of each electron (i.e., was it due to a stellar photon or is it an electron generated by thermal motions within the CCD itself). We have to be clever to separate the signal from the noise. There are two notable quotes to keep in mind while reading this text. The first is from an early review article on CCDs by Craig Mackay (1986), who states: “The only uniform CCD is a dead CCD.” The second is from numerous discussions I have had with CCD users and it is: “To understand your signal, you must first understand your noise.”

3.1 Quantum Efficiency

The composition of a CCD is essentially pure silicon. This element is thus ultimately responsible for the response of the detector to various wavelengths of light. The wavelength dependence of silicon can be understood




Fig. 3.1. The photon absorption length in silicon is shown as a function of wavelength in nanometers. From Reicke (1994).

in an instant by glancing at Figure 3.1. Shown here is the length of silicon needed for a photon of a specific wavelength to be absorbed. Absorption length is defined as the distance for which 63% ($1/e$) of the incoming photons will be absorbed. Figure 3.1 clearly shows that, for light outside the range of about 3,500 to over 8,000 Å, the photons (1.) pass right through the silicon, (2.) get absorbed within the thin surface layers or gate structures, or (3.) simply reflect off the CCD surface. At short wavelengths, 70% or more of the photons are reflected, and for very short wavelengths

(as for long wavelengths) the CCD becomes completely transparent. Thus the quantum efficiency of a typical CCD device will approximately mirror the photon absorption curve for silicon. Shortward of $\sim 2,500 \text{ \AA}$ (for thinned devices) or about 25 \AA (for thick devices) the detection probability for photons increases again. However, due to their much higher energy, these photons lead to the production of multiple electron-hole pairs within the silicon and may also produce damage to the CCD itself (see Chapter 7).

CCD quantum efficiencies are therefore very dependent on the thickness of the silicon that intercepts the incoming photons. This relation between absorption probability and CCD thickness is why front-side illuminated devices are more red sensitive (the photons have a higher chance of absorption) and why they have lower overall (blue) QEs (since the gate structures can be close to or even exceed the necessary absorption depths of as small as only a few atomic layers). A few front-side CCDs have been produced with special gate structures that are transparent to incoming blue and UV photons. In thinned devices, the longer wavelength photons are likely to pass right through the CCD without being absorbed at all.

Figure 3.2 provides some examples of typical quantum efficiency curves. Note the large difference in QE between thinned and thick CCDs. Quantum efficiency or QE curves allow one to quickly evaluate the relative collecting power of the device as a function of wavelength. QE curves shown in the literature are generally assumed to be representative of each and every pixel on the device, that is, all pixels of a given device are assumed to work identically and have the same wavelength response. This is almost true, but it is the “almost” that makes flat fielding of a CCD necessary. In addition, the QE curves shown or delivered for a particular device may be only representative of a “typical” device of the same kind, but they may not be for the exact device of interest.

Placing an antireflection (AR) coating on a CCD (both for visible and near-UV light) increases the QE and extends the range of good QE well into the $3,000 \text{ \AA}$ region. The need for AR coatings comes from the fact that silicon, like most metallic substances, is a good reflector of visible light. If you ever have a chance to hold a CCD, you will easily see just how well the surface does indeed reflect visible light. All the QE curves in Figure 3.2 have the overall shape expected based on the absorption properties of silicon as shown in Figure 3.1. Graphical illustrations of QE curves almost always include photon losses due to the gate structures, electron recombination within the bulk silicon itself, surface reflection,

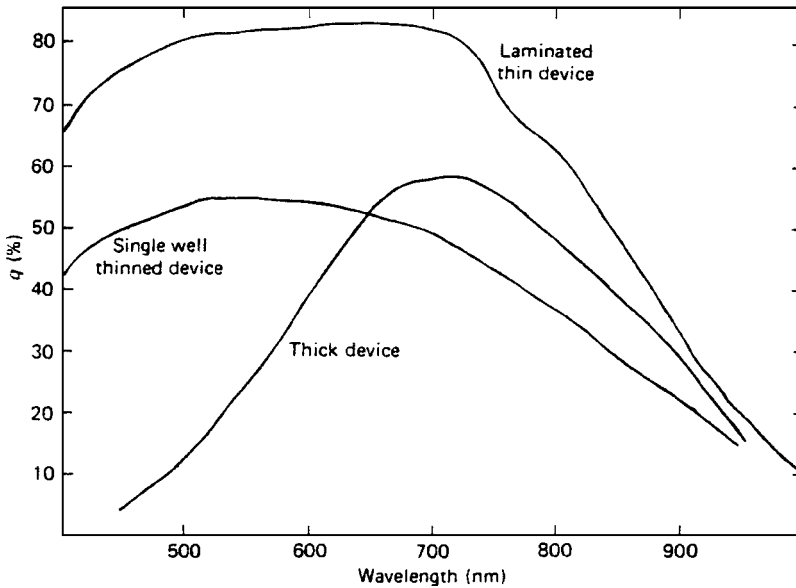
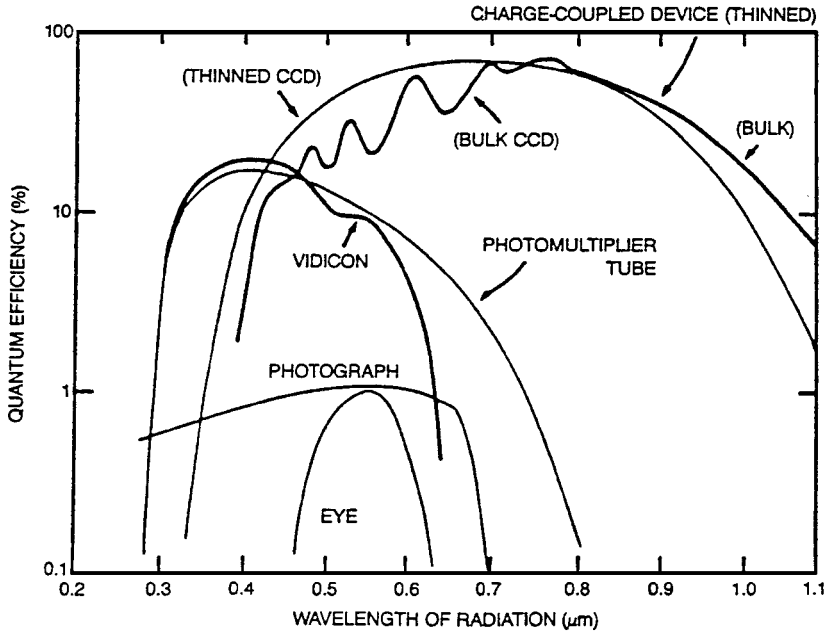


Fig. 3.2. The top panel shows QE curves for various devices indicating why CCDs are a quantum leap above all previous imaging devices. The failure of CCDs at wavelengths shorter than about $3,500 \text{ \AA}$ has been completely eliminated via thinning or coating of the devices. The bottom plot shows representative QE curves for both thick and thinned devices. Laminated and single-well devices are two processes used to produce mechanically stable thinned CCDs.

and, for very long or short wavelengths, losses due to the almost complete lack of absorption by the CCD. Given that all these losses are folded together into the QE value for each wavelength, it should be obvious that changes within the CCD structure itself (such as radiation damage or operating temperature changes) can cause noticeable changes in its quantum efficiency.

Measurement of the quantum efficiency of a CCD is usually performed with the aid of complicated laboratory equipment including well-calibrated photodiodes. Light at each wavelength is used to illuminate both the CCD and the photodiode, and the relative difference in the two readings is recorded. The final result of such an experiment is an absolute QE curve for the CCD (with respect to the calibrated diode) over the range of all measured wavelengths.

To measure a CCD QE curve yourself, a few possibilities exist. You may have access to a setup such as that described above. Measurements can also be made at the telescope itself. One good method of QE measurement for a CCD consists of employing a set of narrow-band filters and a few spectrophotometric standard stars. Performing such a task will provide a good relative QE curve and, if one knows the filter and telescope throughput well, a good absolute QE curve. A detailed reference as to what is involved in the measurement of a spectrophotometric standard star is provided by Tüg et al. (1977). When producing a QE curve by using the above idea, the narrow-band filters provide wavelength selection while the standard stars provide a calibrated light source. A less ambitious QE curve can be produced using typical broad-band (i.e., Johnson) filters, but the final result is not as good because of the large bandpasses and throughput overlap of some of the filters. In between a detailed laboratory setup and the somewhat sparse technique of using filters at the telescope, another method exists. Using an optics bench, a calibrated light source covering the wavelength range of interest, and some good laboratory skills, one can produce a very good QE curve for a CCD and can even turn the exercise into a challenging classroom project.

3.2 Readout Noise

Readout noise, or just read noise, is usually quoted for a CCD in terms of the number of electrons introduced per pixel into your final signal upon readout of the device. Read noise consists of two inseparable components. First is the conversion from an analog signal to a digital number, which is not perfectly repeatable. Each on-chip amplifier and A/D circuit will

produce a statistical distribution of possible answers centered on a mean value.[†] Thus, even for the hypothetical case of reading out the same pixel twice, each time with identical charge, a slightly different answer may be produced. Second, the electronics themselves will introduce spurious electrons into the entire process, yielding unwanted random fluctuations in the output. These two effects combine to produce an additive uncertainty in the final output value for each pixel. The average (one sigma) level of this uncertainty is the read noise and is limited by the electronic properties of the on-chip output amplifier and the output electronics (Djorgovski, 1984).[‡]

The physical size of the on-chip amplifier, the integrated circuit construction, the temperature of the amplifier, and the sensitivity (generally near 1–4 $\mu\text{V}/\text{detected photon}$, i.e., collected photoelectron) all contribute to the read noise for a CCD. In this micro world, the values for electronic noise are highly related to the thermal properties of the amplifier, which in turn determines the sensitivity to each small output voltage. Amazing as it seems, the readout speed, and thus the rate at which currents flow through the on-chip amplifier, can cause thermal swings in the amplifier temperature, which can affect the resulting read noise level. Generally, slower readout speeds produce lower read noise but this reduced readout speed must be weighed against the overall camera duty cycle. Small effects caused by amplifier heating can even occur between the readout of the beginning and end of a single CCD row, as the last charge packets pass through a slightly hotter circuit. Increasing the physical size of the already small microamplifiers can alleviate these small temperature swings, but larger amplifiers have a higher input capacitance, thereby lowering the sensitivity of the amplifier to small voltages.

Additional work on amplifier design, methods of clocking out pixels, and various semiconductor doping schemes can be used to improve the performance of the output electronics. CCD manufacturers invest large efforts into balancing these issues to produce very low read noise devices. Many details of the various aspects of read noise in CCDs are discussed by Janesick & Elliott (1992).

In the output CCD image, read noise is added into every pixel every time the array is read out. This means that a CCD with a read noise of

[†] The distribution of these values is not necessarily Gaussian (Merline & Howell, 1995).

[‡] We note here that the level of the read noise measured, or in fact any noise source within a CCD, can never be smaller than the level of discretization produced by the A/D converter (see Sections 2.4 and 3.6).

20 electrons will, on average, contain 20 extra electrons of charge in each pixel upon readout. High read noise CCDs are thus not very good to use if co-addition of two or more images is necessary. The final resultant image will not be quite as good as one long integration of the same total time, as each co-added image will add in one times the read noise to every pixel in the sum. However, for modern CCDs (see Section 4.4), read noise values are very low and are hardly ever the dominant noise with which one must be concerned. Good read noise values in today's CCDs are in the range of 10 electrons per pixel per read or less. These values are far below read noise levels of ten years ago, which were as high as 50–100 electrons, and even those are well down from values of 500 or more electrons/pixel/read present in the first astronomical CCDs.

In Section 4.3, we will discuss a simple method by which one may determine for oneself the read noise of a given CCD. This determination can be performed with any working CCD system and does not require special equipment, removal of the CCD from the camera dewar, or even removal from the telescope.

3.3 Dark Current

Every material at a temperature much above absolute zero will be subject to thermal noise within. For silicon in a CCD, this means that when the thermal agitation is high enough, electrons will be freed from the valence band and become collected within the potential well of a pixel. When the device is read out, these dark current electrons become part of the signal, indistinguishable from astronomical photons. Thermal generation of electrons in silicon is a strong function of the temperature of the CCD, which is why astronomical use generally demands some form of cooling (McLean, 1997b). Figure 3.3 shows a typical CCD dark current curve, which relates the amount of thermal dark current to the CCD operating temperature. Within the figure the theoretical relation for the rate of thermal electron production is given.

Dark current for a CCD is usually specified as the number of thermal electrons generated per second per pixel or as the actual current generated per area of the device (i.e., picoamps cm^{-2}). At room temperature, the dark current of a typical CCD is near 2.5×10^4 electrons/pixel/second. Typical values for properly cooled devices range from 2 electrons per second per pixel down to very low levels of approximately 0.04 electrons per second for each pixel. Although 2 electrons of thermal noise generated within a pixel every second sounds very low, a




Fig. 3.3. Experimental (symbols) and theoretical (line) results for the dark current generated in a typical three-phase CCD. The rate of dark current, in electrons generated within each pixel every second, is shown as a function of the CCD operating temperature. E_g is the band gap energy for silicon. From Robinson (1988a).

typical 15 minute exposure of a faint astronomical source would include 1,800 additional (thermal) electrons within each CCD pixel upon readout. These additional charges cannot, of course, be uniquely separated from the photons of interest after readout. The dark current produced in a CCD provides an inherent limitation on the noise floor of a CCD. Because dark noise has a Poisson distribution the noise actually introduced by thermal electrons into the signal is proportional to the square root of the dark current (see Section 4.4).

Cooling of CCDs is generally accomplished by one of two methods. The first, and usually the one used for scientific CCDs at major observatories, is via the use of liquid nitrogen (or in some cases liquid air). The CCD and associated electronics (the ones on or very near the actual CCD itself, called the head electronics) are encased in a metal dewar under vacuum. Figure 3.4 shows a typical astronomical CCD dewar




Fig. 3.4. A typical CCD dewar. This is the Mark-II Universal dewar originally produced in 1984 at Kitt Peak National Observatory. The dewar held 1.75 liters of liquid nitrogen providing a CCD operating time of approximately 12 hours between fillings. This dewar could be used in up-looking, down-looking, and side-looking orientations. From Brar (1984).

(Brar, 1984; Florentin-Nielsen, Anderson, & Nielsen, 1995). The liquid nitrogen (LN₂) is placed in the dewar and, although not in direct physical contact with the CCD, cools the device to temperatures of near -100°C . Since LN₂ itself is much colder than this, CCDs are generally kept at a constant temperature ($\pm 0.1^{\circ}\text{C}$) with an on-board heater.

In fact, the consistency of the CCD temperature is very important as the dark current is a strong function of temperature (Figure 3.3) and will vary considerably due to even modest changes in the CCD temperature.

A less expensive and much less complicated cooling technique makes use of thermoelectric cooling methods. These methods are employed in essentially all “off-the-shelf” CCD systems and allow operation at temperatures of -20 to -50°C or so, simply by plugging the cooler into an electrical outlet. Peltier coolers are the best known form of thermoelectric cooling devices and are discussed in Martinez & Klotz (1998). CCD operation and scientific quality imaging at temperatures near -30°C is possible, even at low light levels, owing to advances in CCD design and manufacturing techniques and the use of multipinned phase operation (see Chapter 2). Other methods of cooling CCDs that do not involve LN2 are discussed in McLean (1997a).

The amount of dark current a CCD produces depends primarily on its operating temperature, but there is a secondary dependence upon the bulk properties of the silicon used in the manufacture. Even CCDs produced on the same silicon wafer can have slightly different dark current properties. As with most of the noise properties of a given CCD, custom tailoring the CCD electronics (such as the bias level and the readout rate) can produce much better or much worse overall dark current and noise performance.

3.4 CCD Pixel Size, Pixel Binning, Full Well Capacity, and Windowing

This section is a combination of a few related topics concerning the amount of charge that can be stored within a given pixel during an integration. We have seen that CCD thinning, MPP operation, and small physical pixel size all place limitations on the total number of electrons that can be collected within a pixel. The general rule of thumb is that the physically larger the pixel (both in area and in thickness) the more charge that it can collect and store.

The amount of charge a pixel can hold in routine operation is termed its full well capacity. A Kodak CCD with 9-micron pixels (meaning 9 microns on a side for the projected area, but giving no indication of the thickness of the CCD) operating in MPP mode has a full well capacity per pixel of 85,000 electrons. In contrast, a SiTe CCD with 24-micron pixels can have a full well capacity per pixel of over 350,000 electrons.

When we discussed the method by which a CCD is read out (Chapter 2) it was stated that each row is shifted in turn into the output register and then digitized, and the resulting DN value is sent off to the computer. During this process, each pixel's value is increased on average by one times the read noise. If we could add up the charge within say 4 pixels before they are digitized, we would get a final signal level equal to ~ 4 times each single pixel's value, but only one times the read noise. This process is called on-chip binning and, if selected, occurs prior to readout within the CCD output register (Merline & Howell, 1995; Smith, 1990b). Pixels can be binned (summed) in both vertical and horizontal directions. "On-chip" means that the accumulated charge from each pixel involved in the binning is brought together and summed before the process of A/D conversion occurs. This summing process is done in the output register and is limited by the size of the "pixels" within this register. Generally, the output register pixels can hold five to ten times the charge of a single active pixel. This deeper full well capacity of the output register pixels allows pixel summing to take place.

Older CCD systems that allowed on-chip binning had plug boards mounted on the sides of the dewar. Certain combinations of the plug wires produced different on-chip binning patterns and users could change these to suit their needs. Today, most CCDs have the ability to perform pixel summing as a software option (Leach, 1995). Binning terminology states that normal operation (or "high resolution" as it is called by many low-cost CCDs) is a readout of the CCD in which each pixel is read, digitized, and stored. This is called 1×1 binning or unbinned operation. Binning of 2×2 would mean that an area of 4 adjacent pixels will be binned or summed on-chip within the output register during readout, but before A/D conversion. The result of this binning operation will produce only one "superpixel" value, which is digitized and stored in the final image; the original values in each of the four summed pixels are lost forever. Mainly for spectroscopic operation, binning of 3×1 is commonly used, with the 3 being in the direction perpendicular to the dispersion. Binning of CCD pixels decreases the image resolution, usually increases the final signal-to-noise value of a measurement, and reduces the total readout time and final image size. For example, a 1024×1024 CCD binned 2×2 will have a final image size of only 512×512 pixels and the readout time will be reduced by about a factor of four.

Pixel binning gives flexibility to the user for such applications as (using a high binning factor) quick readout for focus tests, nights with poor seeing, or very low surface brightness observations. Spectroscopic observations with a CCD, high spatial resolution imaging, or bright object observations will benefit from the use of a low binning factor. Binning factors that are very large (say 40×40 pixels) might be of use in some rare cases, but they will be limited by the total amount of charge one can accumulate in a single superpixel of the output register.

A related function available with some CCDs is “windowing.” Windowing allows the user to choose a specific rectangular region (or many regions) within the active area of the CCD to be read out upon completion of the integration. The CCD window is often specified by providing the operating software with a starting row and column number and the total number of x, y pixels to use. For example, using a 2048×2048 CCD to make high speed imaging observations would be difficult, but windowing the CCD to use only the first 512 rows and columns (0, 0, 512, 512) allows for much faster readout and requires far less storage for the image data. Of course, your object of interest must now be positioned within these first 512 rows and columns, and not at the center of the CCD as may be usual. Other applications of CCD windowing would include choosing a cosmetically good subregion of a large CCD or only a rectangular strip to read out from a larger square CCD, when making spectroscopic observations. CCD windowing is independent of any on-chip binning, and one can both window and bin a CCD for even more specific observational needs.

3.5 Overscan and Bias

In an attempt to provide an estimate of the value produced by an empty or unexposed pixel within a CCD, calibration measurements of the bias level can be used.[†] Bias or zero images allow one to measure the zero noise level of a CCD. For an unexposed pixel, the value for zero collected photoelectrons will translate, upon readout and A/D conversion, into a mean value with a small distribution about zero.[‡] To avoid negative

[†] For more on bias frames and their use in the process of CCD image calibration, see Chapter 4.

[‡] Before bias frames, and in fact before any CCD frame is taken, a CCD should undergo a process known as “wiping the array.” This process makes a fast read of the detector, without A/D conversion or data storage, in order to remove any residual dark current or photoelectron collection that may have occurred during idle times between obtaining frames of interest.

numbers in the output image,[†] CCD electronics are set up to provide a positive offset value for each accumulated image. This offset value, the mean “zero” level, is called the bias level. A typical bias level might be a value of 400 ADU (per pixel), which, for a gain of $10\text{ e}^-/\text{ADU}$, equals 4,000 electrons. This value might seem like a large amount to use, but historically temporal drifts in CCD electronics due to age, temperature, or poor stability in the electronics, as well as much higher read noise values, necessitated such levels.

To evaluate the bias or zero noise level and its associated uncertainty, specific calibration processes are used. The two most common ones are: (1.) overscan regions produced with every object frame or (2.) usage of bias frames. Bias frames amount to taking observations without exposure to light (shutter closed), for a total integration time of 0.000 seconds. This type of image is simply a readout of the unexposed CCD pixels through the on-chip electronics, through the A/D converter, and then out to the computer producing a two-dimensional bias or zero image.

Overscan strips, as they are called, are a number of rows or columns (usually 32) or both that are added to and stored with each image frame. These overscan regions are not physical rows or columns on the CCD device itself but additional pseudo-pixels generated by sending additional clock cycles to the CCD output electronics. Both bias frames and overscan regions are techniques that allow one to measure the bias offset level and, more importantly, the uncertainty of this level.

Use of overscan regions to provide a calibration of the zero level generally consists of determining the mean value within the overscan pixels and then subtracting this single number from each pixel within the CCD object image. This process removes the bias level pedestal or zero level from the object image and produces a bias-corrected image. Bias frames provide more information than overscan regions, as they represent any

[†] Representation of negative numbers requires a sign bit to be used. This bit, number 15 in a 16-bit number, is 0 or 1 depending on whether the numeric value is positive or negative. For CCD data, sacrificing this bit for the sign of the number leaves one less bit for data, thus reducing the overall dynamic range. Therefore, most good CCD systems do not make use of a sign bit. One can see the effects of having a sign bit by viewing CCD image data of high numeric value but displayed as a signed integer image. For example, a bright star will be represented as various grey levels, but at the very center (i.e., the brightest pixels) the pixel values may exceed a number that can be represented by 14 bits (plus a sign). Once bit 15 is needed, the signed integer representation will be taken by the display as a negative value and the offending pixels will be displayed as black. This is due to the fact that the very brightest pixel values have made use of the highest bit (the sign bit) and the computer now believes the number is negative and assigns it a black (negative) greyscale value. This type of condition is discussed further in Appendix C.

two-dimensional structure that may exist in the CCD bias level. Two-dimensional (2-D) patterns are not uncommon for the bias structure of a CCD, but these are usually of low level and stable with time. Upon examination of a bias frame, the user may decide that the 2-D structure is nonexistent or of very low importance and may therefore elect to perform a simple subtraction of the mean bias level value from every object frame pixel. Another possibility is to remove the complete 2-D bias pattern from the object frame using a pixel-by-pixel subtraction (i.e., subtract the bias image from each object image). When using bias frames for calibration, it is usually best to work with an average or median frame composed of many (10 or more) individual bias images (Gilliland, 1992). This averaging eliminates cosmic rays,[†] read noise variations, and random fluctuations which will be a part of any single bias frame.

Variations in the mean zero level of a CCD are known to occur over time and are usually slow drifts over many months or longer, not noticeable changes from night to night or image to image. These latter types of changes indicate severe problems with the readout electronics and require correction before the CCD image data can be properly used.

Producing a histogram of a typical averaged bias frame will reveal a Gaussian distribution with the mean level of this distribution being the bias level offset for the CCD. We show an example of such a bias frame histogram in Figure 3.5. The width of the distribution shown in Figure 3.5 is related to the read noise of the CCD (caused by shot noise variations in the CCD electronics (Mortara & Fowler, 1981)) and the device gain by the following expression:

$$\sigma_{\text{ADU}} = \frac{\text{Read Noise}}{\text{Gain}}.$$

3.6 CCD Gain and Dynamic Range

The gain of a CCD is set by the output electronics and determines how the amount of charge collected in each pixel will be assigned to a digital number in the output image. Gain values are usually given in terms of the number of electrons needed to produce one ADU step within the A/D converter. Listed as electrons/Analog-to-Digital Unit (e^-/ADU), common gain values range from 1 (photon counting) to 150 or more.

[†] Cosmic rays are not always cosmic! They can be caused by weakly radioactive materials used in the construction of CCD dewars (Florentin-Nielsen, Anderson, & Nielsen, 1995).




Fig. 3.5. Histogram of a typical bias frame showing the number of pixels vs. each pixel ADU value. The mean bias level offset or pedestal level in this Loral CCD is near 1,017 ADU, and the distribution is very Gaussian in nature with a FWHM value of near 2 ADU. This CCD has a read noise of 10 electrons and a gain of $4.7 \text{ e}^-/\text{ADU}$.

One of the major advantages of a CCD is that it is linear in its response over a large range of data values. Linearity means that there is a simple linear relation between the input value (charge collected within a each pixel) and the output value (digital number stored in the output image).

The largest output number that a CCD can produce is set by the number of bits in the A/D converter. For example, if you have a 14-bit A/D, numbers in the range from zero to 16,383 can be represented.[†] A 16-bit A/D would be able to handle numbers as large as 65,535 ADU.

Figure 3.6 provides a typical example of a linearity curve for a CCD. In this example, we have assumed a 15-bit A/D converter capable of producing output DN values in the range of 0 to 32,767 ADU, a device gain of $4.5 \text{ e}^-/\text{ADU}$, and a pixel full well capacity of 150,000 electrons. The linearity curve shown in Figure 3.6 is typical for a CCD, revealing

[†] The total range of values that a specific number of bits can represent equals $2^{(\text{number of bits})}$, e.g., $2^{14} = 16,384$. CCD output values are zero based, that is, they range from 0 to $2^{(\text{number of bits})} - 1$.

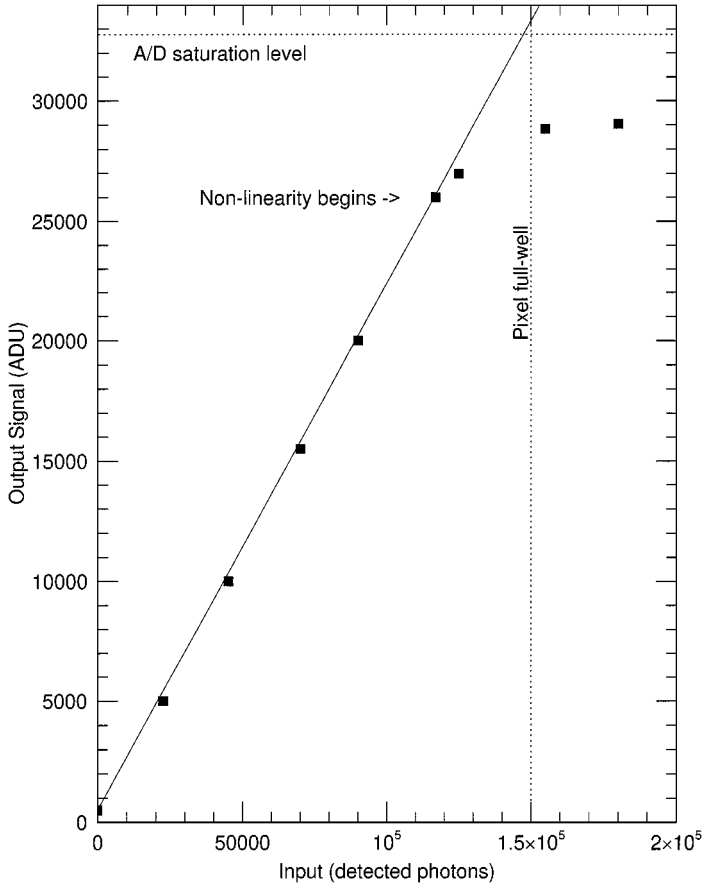


Fig. 3.6. CCD linearity curve for a typical three-phase CCD. We see that the device is linear over the output range from 500 ADU (the offset bias level of the CCD) to 26,000 ADU. The pixel full well capacity is 150,000 electrons and the A/D converter saturation is at 32,767 ADU. In this example, the CCD nonlinearity is the limiting factor of the largest usable output ADU value. The slope of the linearity curve is equal to the gain of the device.

that over most of the range the CCD is indeed linear in its response to incoming photons. Note that the CCD response has the typical small bias offset (i.e., the output value being nonzero even when zero incident photons occur), and the CCD becomes nonlinear at high input values. For this particular CCD, nonlinearity sets in near an input level of 1.17×10^5 photons (26,000 ADU), a number still well within the range of possible output values from the A/D.