

many areas of observational science. Astronomy, physics, chemistry, medical imaging, and remote sensing are just a few of the fields of study in which they play a crucial role. This handbook provides a concise and accessible reference on all practical aspects of using CCDs.

Starting with the electronic workings of these marvels of modern science, this handbook discusses the basic characteristics of CCDs and then gives methods and examples of how to determine these values. While the focus of the book is on the use of CCDs in professional observational astronomy, advanced amateur astronomers and researchers in physics, chemistry, medical imaging, and remote sensing will also find it very valuable. The core of the book focuses on the detection of optical light, but due coverage is given to other wavelengths, particularly high-energy astronomy. Tables of useful and hard-to-find data and key practical equations round the book off and ensure that it provides an ideal introduction to the practical use of CCDs for graduate students, as well as a handy reference for more experienced researchers.

Dr. Steve Howell was raised in the eastern United States. Starting his research career with laser physics work at General Motors Laboratories, he eventually moved west to work at Kitt Peak National Observatory where he supported software and hardware development for CCD instrumentation. Since that time, Howell has worked at Goddard Space Flight Center as a member of the International Halley Watch, developing image processing software for use by NASA and the Shuttle programs, was a team member for the Galileo Spacecraft imager, worked at the Planetary Science Institute as a staff scientist, and was a professor of astrophysics at the University of Wyoming. Currently, Howell is an adjunct faculty member of the Physics and Astronomy Department at Wyoming and Head of the Astrophysics Group of the Planetary Science Institute in Tucson, Arizona.

1.1  
Q 8

12.7.  
. H 6  
2000

## Cambridge Observing Handbooks for Research Astronomers

Today's professional astronomers must be able to adapt to use telescopes and interpret data at all wavelengths. This series is designed to provide them with a collection of concise, self-contained handbooks which covers the basic principles peculiar to observing in a particular spectral region, or to using a special technique or type of instrument. The books can be used as an introduction to the subject and as a handy reference for use at the telescope, or in the office. They also promote an understanding of other disciplines in astronomy and a modern, multi-wavelength, multi-technique approach to research. Although aimed primarily at graduate students and researchers, many titles in the series are of interest to keen amateurs and undergraduate students.

### *Series editors*

Professor Richard Ellis, Institute of Astronomy, University of Cambridge  
Professor John Huchra, Center for Astrophysics, Smithsonian Astrophysical Observatory  
Professor Steve Kahn, Department of Physics, Columbia University, New York  
Professor George Rieke, Steward Observatory, University of Arizona, Tucson  
Dr Peter B. Stetson, Herzberg Institute of Astrophysics, Dominion Astrophysical Observatory, Victoria, British Columbia

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

# Contents

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 CE2 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<sup>A</sup>T<sub>E</sub>X [T<sub>B</sub>]

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

*Library of Congress Cataloging in Publication Data*  
ISBN 0 521 64058 X hardback  
ISBN 0 521 64834 3 paperback

	<i>page</i>
<i>Preface</i>	ix
<b>1 Introduction</b>	<b>1</b>
1.1 Nomenclature	2
1.2 Why Use CCDs?	2
<b>2 CCD Manufacturing and Operation</b>	<b>7</b>
2.1 CCD Operation	8
2.2 CCD Types	12
2.3 CCD Coatings	21
2.4 Analog-to-Digital Converters	22
<b>3 Characterization of Charge-Coupled Devices</b>	<b>26</b>
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
<b>4 CCD Imaging</b>	<b>47</b>
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
<b>5 Photometry and Astrometry</b>	<b>75</b>
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

## Preface

We are all aware of the amazing astronomical images produced with telescopes these days, particularly those displayed as color representations and shown off on websites and in magazines. For those of us who are observers, we deal with our own amazing images produced during each observing run. Just as spectacular are photometric, astrometric, and spectroscopic results generally receiving less fanfare but often of more astrophysical interest. What all of these results have in common is the fact that behind every good optical image lies a good charge-coupled device.

Charge-coupled devices, or CCDs as we know them, are involved in many aspects of everyday life. Examples include video cameras for home use and those set up to automatically trap speeders on British highways, hospital x-ray imagers and high speed oscilloscopes, and digital cameras used as quality control monitors. This book discusses these remarkable semiconductor devices and their many applications in modern day astronomy.

Written as an introduction to CCDs for observers using professional or off-the-shelf CCD cameras as well as a reference guide, this volume is aimed at students, novice users, and all the rest of us who wish to learn more of the details of how a CCD operates. Topics include the various types of CCD; the process of data taking and reduction; photometric, astrometric, and spectroscopic methods; and CCD applications outside of the optical bandpass. The level of presentation is aimed not only at college or professional level readers but also at a more general audience including the ever growing number of highly trained and motivated amateurs and other professionals in technical areas in which CCDs play a role.

Chapters 2 and 3 contain all the fundamental information on CCD operation and characteristics while each remaining chapter can be mastered individually. In a book of this length, many aspects must be treated briefly. However, I have made an effort to provide self-contained detail of the important aspects of CCDs while including numerous references to the detailed professional literature for those desiring a deeper understanding. Additionally, throughout the book, examples related to common observational occurrences as well as footnotes discussing interesting but not mainstream topics are included. Appendices list other reference works of interest, CCD manufacturers, numerous website addresses of interest, and a brief introduction to image displays.

This book started with an idea for a new series of handbooks, including a volume on CCDs. I am happy to thank the editor, Adann Black, and Peter Stetson for allowing me to be involved in this series and write this book. The folks at Cambridge University Press, particularly Adam, have been very helpful, dealing with my many questions during the writing process. Michie Shaw and the staff at TechBooks have helped greatly in the final steps of production. I would like to thank the anonymous readers of an early draft of this book for their comments and for pointing out some important areas and results I had overlooked. Two readers, Peter Stetson and John Hirsch, suggested coverage of material that has led to the inclusion of additional topics in the final version. A number of colleagues have provided information, graphs, references, and support during the writing of this book, all of which I appreciate. I thank my former and current students and postdocs for keeping me on my toes regarding many of the topics herein. Chris Sabbey kindly provided Figure 6.8, and Figures 1.1 and 4.6 were taken by Simon Tulloch and provided by Derek Ives, both of the UK Astronomy Technology Centre.

I would like to acknowledge and thank my parents, Cecil and Barbara, for allowing me to be "scientific" during my childhood. My experimentation, often at the expense of household items, was not always successful but was never discouraged. These experiences deeply planted the seed of scientific fascination in my being. Appreciation is also passed along to my brother Terry, for the many hours we spent together exploring the world around us. Particularly noteworthy were the times we spent watching, analysing, and laughing at "B" sci-fi movies.

During the writing of this volume on CCDs, many opportunities were missed related to spending time with my son Graham and my wife and friend Mary. Both were always supportive of the effort, encouraged its

completion, and have accumulated many IOUs, which I will now have the pleasure of paying off. I appreciate their unflinching love.

I have had fun writing this book and learning even more about CCDs, almost as much fun as I have when I observe with them. I hope that you the reader will find this work of interest as well and enjoy paging through it often. Astronomy has always fascinated humans and if this treatise allows you to obtain a better knowledge of CCDs and with it even more fascination with the universe around us, it will have been a success.

"Since The Beginning Of Time, The Universe Has Called To Awaken Each Of Us. To Understand The Universe Is To Understand Ourselves."

## Introduction

Silicon. This semiconductor material certainly has large implications on our life. Its uses are many, including silicon oil lubricants, implants to change our body's outward appearance, electric circuitry of all kinds, nonstick frying pans, and of course, charge-coupled devices.

Charge-coupled devices (CCDs) and their use in astronomy will be the topic of this book. We will not discuss the use of CCDs in commercial video cameras or many other industrial and scientific applications. As we will see, there are four main methods of employing CCD imagers in astronomical work: imaging, astrometry, photometry, and spectroscopy. Each of these topics will be discussed in turn. Since the intrinsic physical properties of silicon, and thus CCDs, are most useful at optical wavelengths (about 3,000 to 10,000 Å), the majority of our discussion will be concerned with visible light applications. Additional specialty or lesser used techniques and CCD applications outside the optical bandwidth will be mentioned only briefly.

Applications of infrared (IR) arrays – semiconductor devices with some similar properties to CCDs – while important to modern astronomy, will receive small mention here. A complete treatment of IR arrays is given by Ian Glass in his companion book, *Handbook of Infrared Astronomy*, the first book in this series.

Appendix A provides the reader with a detailed list of other books and major works devoted to CCD information. This appendix does not contain the large body of journal articles in existence; these will be selectively referenced throughout the text. Appendix B provides a list of present day CCD manufacturers that produce both amateur and professional grade CCDs and CCD cameras. Finally, Appendix C discusses some basic principles of image display devices. While not directly related to CCDs, computer displays are the medium by which we view

and examine the information collected by CCDs. Proper interpretation of the CCD image is only possible if one understands how it is displayed.

### 1.1 Nomenclature

CCDs are often listed and named by an apparently strange convention. This small section aims to demystify these odd sounding names. CCDs come in various sizes and shapes and are manufactured by a number of companies (see Appendix B).

Figure 1.1 illustrates a few typical CCDs. Present day CCDs generally come in sizes ranging from 512 by 512 "picture elements" or pixels to arrays as large as 4,096 by 4,096 pixels. Often these size designations are written as 512 × 512 or 512<sup>2</sup>. CCDs are also available as rectangular devices of unequal length and width and with nonsquare pixels. For example, CCDs of size 1024 × 3072 pixels are produced for spectroscopic applications. We will see in Chapter 2 that each pixel acts as an electrically isolated portion of the silicon array and is capable of incoming photon collection, storage of the produced photoelectrons, and readout from the CCD array to an associated computer as a digital number.

The names or designations of CCDs are usually a combination of the company name with the CCD size. RCA512, 4K × 2K EEV, or SITE2048 are examples. Instrumentation at observatories often include a CCD as the detector and are specialized to perform a task such as imaging or spectroscopy. Observatories designate these instruments with a name that may or may not include information about the associated CCD detector. The Royal Greenwich Observatory (RGO) on La Palma has the FOS#1A (a 512 × 1024 Loral CCD used in their Faint Object Spectrograph), and the Tek 2K CCD of the National Optical Astronomy Observatories (NOAO) is a 2048 × 2048 pixel array used in their 0.9-m telescope imaging camera. Observatories keep lists of each of their CCDs and instruments with associated information about the CCD specifications. For examples of such information, check out the observatory websites listed in Appendix B.

### 1.2 Why Use CCDs?

Most astronomical detectors in use today at professional observatories, as well as with many amateur telescopes, are CCDs. This fact alone gives the reader an impression that there must be something very special




Fig. 1.1. A selection of recent CCDs that have been placed in operation at the Royal Greenwich Observatory. Clockwise from bottom left they are: SITE-002 (2048 × 4096), Loral 2K3eb (2048 × 2048), EEV CCD42-80 (2048 × 4096), SITE-424 (2048 × 2048), GEC P8603 (385 × 578), EEV 15-11 (1024 × 256), Tek1024 (1024 × 1024), Loral 512FT (512 × 1024), EEV-05-30 (1242 × 1152), EEV CCD42-10 (2048 × 512), Loral-64 (64 × 64), and EEV CCD39-01 (80 × 80).



or useful about CCDs; otherwise why all the fuss? CCDs have revolutionized modern astronomy. They will take their place in astronomical history along with other important discoveries such as the telescope, photographic plates, prisms, and spectroscopy. The contribution to our knowledge of the heavens brought about by CCDs is astounding, even more so when one considers that they have been in use for only about twenty-five years.

First introduced as electronic analogues to magnetic bubble memory (Boyle & Smith, 1976; Ameilo, Tompsett, & Smith, 1970) at Bell Labs, CCDs provided their first astronomical image in 1975 when scientists from the Jet Propulsion Laboratory imaged the planet Uranus at a wavelength of 8,900 Å (Janesick & Blouke, 1987). This observation used the University of Arizona 61-inch telescope atop Mt. Lemmon and a prototype version of a CCD made by Texas Instruments Corporation as part of a development project for NASA spacecraft missions.

During the past quarter century, tremendous progress has been made in the manufacturing process and therefore in the properties of the CCD itself. These improvements have allowed much lower noise properties for CCDs, thereby increasing their overall astronomical usage. We outline below two of the important reasons why CCDs are considered as essentially the perfect imaging device. Details of the manufacturing techniques and properties of CCDs will be presented in Chapters 2 and 3.

### 1.2.1 Noise Properties

The usefulness of a detector is very often determined by the amount of inherent noise within the device itself. We shall see in Chapter 3 how the noise properties of a CCD are determined, but, suffice it to say here, modern CCDs are almost noise free. The original line of photosensitive electronic array detectors, such as television-type imaging detectors, vidicons, silicon intensified targets, and image-dissector scanners, all had very high noise properties. For comparison, silicon intensified target imagers (SITs) had a noise level upon readout of 800 electrons per picture element. Some very good systems of this type could be produced with read noise values of only 200 electrons (Eccles, Sim, & Tritton, 1983). The first CCDs had readout noise levels similar to this latter value, while modern CCDs have noise values of fifteen to only a few electrons per pixel. The large noise levels present in early array detectors not only limited the signal-to-noise ratio obtainable for a given

measurement, they also severely limited the total dynamic range available to the camera.

### 1.2.2 Quantum Efficiency and Bandpass

Quantum efficiency (QE) is the term used to report on the ability of a detector to turn incoming photons into useful output. Generally, it is defined as the ratio of incoming photons to those actually detected or stored in the device. A QE of 100% would be an ideal detector with every incoming photon detected and accounted for in the output. Bandpass is a term that means the total spectral range for which a detector is sensitive to the incoming photons. Our eyes, for example, have a very limited bandpass covering only about 2,000 Å of the optical spectral range, from near 4,500 to 6,500 Å.

One of the great advantages of CCDs compared with earlier detectors is their ability to convert a large percentage of incoming photons into photoelectrons. Photographic plates had an intrinsic quantum efficiency of only about 2% (Kodak IIIa's reached 3%), with "hypersensitized" plates (plates treated to a black-magic process involving heating and exposure to various gases) reaching claimed QEs as high as 10%. Because photographic emulsions were originally more sensitive to UV and blue light, numerous dyes and coatings were developed to both extend their bandpass coverage and allow detection of yellow to red optical photons.

Early solid-state imaging devices and intensified target devices could reach quantum efficiencies as high as 20–40%. These devices relied on television-type tube technology and electron beam scanning for readout of detected photons. Since they used silicon (or similar type materials) as the detector material, their useful spectral bandpass was similar to that of modern CCDs. However, besides the relatively low QE, these early electronic detectors had other drawbacks. For example, image tubes needed high voltage input for electron acceleration and the observed two-dimensional scene was not easily or consistently divided into well-determined  $x, y$  positional output (Walker, 1987). Changes in the voltage levels, collected charge, and telescope position resulted in electric and magnetic field variations leading to positional and flux measurement uncertainties.

Even the earliest CCDs (those manufactured by Fairchild or GEC) easily reached peak QEs of 40%. Typical CCD QE curves today not only peak near 90% but are 60% or more quantum efficient over two

thirds of their total spectral range. The bandpass available in a modern CCD (with a QE of 10% or more) is about 3,000–10,500 Å. Coatings and phosphors deposited on the CCD surface or the use of some form of pre-processor device can extend the bandpass sensitivity limits or increase the QE in specific wavelength ranges (see Chapter 2).

This volume of the Cambridge Observing Handbooks for Research Astronomers will explore the world of CCDs from their inner workings to numerous applications in observational astronomy. Appendices are included to provide ancillary information related to the main text. The chapters will make little assumption as to the readers previous knowledge on the subject, each attempting to be somewhat self-contained. Chapters 4, 5, and 6 deal directly with astronomical applications while Chapters 2 and 3 are of general interest to those wanting an overall understanding of CCDs as detectors. Chapter 7 discusses the use of CCDs at nonoptical wavelengths. In a short treatise such as this, coverage of numerous details and nuances is not possible; thus a detailed reference list to general texts or collections of articles on CCDs is provided in Appendix A. For those wishing to explore a subject at a deeper level, pertinent research articles are provided throughout the text itself.

## CCD Manufacturing and Operation

Before we begin our discussion of the physical and intrinsic characteristics of charge-coupled devices (Chapter 3), we want to spend a brief moment looking into how CCDs are manufactured and some of the basic, important properties of their electrical operation.

The method of storage and information retrieval within a CCD is dependent on the containment and manipulation of electrons (negative charge) and holes (positive charge) produced within the device when exposed to light. The produced photoelectrons are stored in the depletion region of a metal insulator semiconductor (MIS) capacitor, and CCD arrays simply consist of many of these capacitors placed in close proximity. Voltages, which are static during collection, are manipulated during readout in such as way as to cause the stored charges to flow from one capacitor to another, providing the reason for the name of these devices. These charge packets, one for each pixel, are passed through readout electronics that detect and measure each charge in a serial fashion. An estimate of the numerical value of each packet is sent to the next step in this process, which takes the input analog signal and assigns a digital number to be output and stored in computer memory.

Thus, originally designed as a memory storage device, CCDs have swept the market as replacements for video tubes of all kinds owing to their many advantages in weight, power consumption, noise characteristics, linearity, spectral response, and others. We now further explore some of the details glossed over in the above paragraph to provide the reader with a basic knowledge of the tortuous path that the detected photon energy takes from collection to storage. The design of CCD electronics, semiconductor technology, and detailed manufacturing methods are far beyond the level or space constraints of this book. For further information the reader is referred to the excellent discussion

in Janesick & Elliott (1992) and the other technical presentations listed in Appendix A.

### 2.1 CCD Operation

The simplest and very understandable analogy for the operation of a CCD is also one that has been used numerous times for this purpose (Janesick & Blouke, 1987). This is the "water bucket" idea in which buckets represent pixels on the CCD array, and a rainstorm provides the incoming photons (rain drops). Imagine a field covered with buckets aligned neatly in rows and columns throughout the entirety of the area (Figure 2.1). After the rainstorm (CCD integration), each bucket is transferred in turn and metered to determine the amount of water collected. A written record (final CCD image) of the amount of water in each bucket will thus provide a two-dimensional record of the rainfall within the field.

Referring to the actual mechanisms at work within a CCD, we start with the method of charge generation within a pixel: the photoelectric




Fig. 2.1. CCDs can be likened to an array of buckets that are placed in a field and collect water during a rainstorm. After the storm, each bucket is moved along conveyor belts until it reaches a metering station. The water collected in each field bucket is then emptied into the metering bucket within which it can be measured. From Janesick & Blouke (1987).

effect. Incoming photons strike the silicon within a pixel and are easily absorbed if they possess the correct wavelength (energy). Silicon has a band gap energy of 1.14 electron volts (eV), and so it easily absorbs light of energy 1.1 to 4 eV (11,000 to 3,000 Å).<sup>†</sup> Photon absorption causes the silicon to give up a valence electron and move it into the conduction band. Photons of energy 1.1 eV to near 4 or so eV generate single electron-hole pairs, whereas those of higher energy produce multiple pairs (see Section 2.2.7 and Chapter 7). Left to themselves, these conduction band electrons would recombine back into the valence level within approximately 100 microseconds. Silicon has a useful photoelectric effect range of 1.1 to about 10 eV, which covers the near-IR to soft X-ray region (Reicke, 1994). Above and below these limits, the CCD material appears transparent to the incoming photons.

Once electrons have been freed to the conduction band of the silicon, they must be collected and held in place until readout occurs. The details of the actual construction of each pixel within a CCD, that is, the formation of the MIS capacitor with its doped silicon, layers of silicon dioxide, etc., are beyond the scope of this book (Janesick & Elliott, 1992; Eccles, Sim, & Tritton, 1983), but suffice it to say that each pixel has a structure conducive to allowing applied voltages to be placed on sub-pixel sized electrodes called gates. These gate structures provide each pixel with the ability to collect the freed electrons and hold them in a potential well until the end of the exposure. In a typical arrangement, each pixel has associated with it three gates, each of which can be set to a different voltage potential. The voltages are controlled by clock circuits with every third gate connected to the same clock. Figure 2.2 illustrates this clocking scheme for a typical three-phase device.

We note in Figure 2.2 that, when an exposure ends, the clock voltages are manipulated such that the electrons that have been collected and held in each pixel's +10 volt potential well by clock voltage V3 can now be shifted within the device. Note that electrons created anywhere within the pixel during the exposure (where each pixel has a surface area equal to the total area under all three gates) will be forced to migrate toward the deepest potential well. When the exposure is terminated and CCD readout begins, the voltages applied to each gate are cycled (this process is called clocking the device) such that the charge stored within each pixel during the integration is electronically shifted.

<sup>†</sup> The energy of a photon of a given wavelength (in electron volts) is given by  $E(\text{eV}) = 12,407/\lambda(\text{Å})$ .




Fig. 2.2. Schematic voltage operation of a typical three-phase CCD. The clock voltages are shown at three times during the readout process, indicating their clock cycle of 0, 10, and 5 volts. One clock cycle causes the stored charge within a pixel to be transferred to its neighboring pixel. CCD readout continues until all the pixels have had their charge transferred completely out of the array and through the A/D converter. From Walker (1987).

A simple change in the voltage potentials ( $V_3$  goes to +5 volts, while  $V_1$  becomes +10 volts and so on) allows the charge to be shifted in a serial fashion along columns from one CCD pixel to another throughout the array. The transfer of the total charge from location to location within the array is not without losses. As we will see, each charge transfer (one of which occurs for each voltage change or clock cycle) has an associated efficiency. This efficiency value is the percent of charge transferred compared with that which was actually collected. Modern values for the charge transfer efficiency (CTE) are near 0.99999 (i.e., 99.9999% efficient) for each transfer.

Each column in the array is connected in parallel and thus each pixel shift is mimicked throughout the entire array simultaneously. One clock cycle moves each row of pixels up one column, with the top row being shifted off the array into what is called the output shift register or horizontal shift register. This register is simply another row of pixels hidden from view (i.e., not exposed to incident light) and serves as the transition between active rows on the array and the output of the device. Once an entire row is shifted into the output register, and before any further row shifts on the active area occur, each pixel in the output register is shifted out one at a time (in a similar manner as before) into the output electronics. Here, the charge collected within each pixel is measured as a voltage and converted into an output digital number (see Section 2.4). Each pixel's collected charge is sensed and amplified by an output amplifier. CCD output amplifiers are designed to have low noise and are built directly into the silicon circuitry; thus they are often referred to as on-chip amplifiers. These amplifiers must work with extremely small voltages and are rated, as to their sensitivity, in volts per electron. Typical values are in the range of 0.5 to 4 microvolts per electron.

The output voltage from a given pixel is converted to a digital number (DN) and is typically discussed from then on as either counts or ADUs (analog-to-digital units). The amount of voltage needed (i.e., the number of collected electrons or received photons) to produce 1 ADU is termed the gain of the device. We will discuss the gain of a CCD in Chapter 3 and here only mention a few items of interest about it. A typical CCD gain might be 10 electrons/ADU, which means that for every 10 electrons collected within a pixel, the output from that pixel will produce, on average, a count or DN value of 1. For example, with this gain value if a pixel collects 1,000 electrons (photons), the output pixel value stored in the computer would be 100 ADUs. For 1,500 electrons 150 ADUs would be produced and for 17,234 electrons, the output pixel value would be 1,723 ADUs (note, not 1723.4). Digital output values can only be integer numbers and it is clear already that the discrimination between different pixel values can only be as good as the resolution of the gain and digital conversion of the device.

The process of conversion of the output voltage signal into a DN is performed within a device called an analog-to-digital converter (A/D or ADC). We will see later on that there is an intimate connection between the number of digital bits available in the A/D and the value of the gain that can or should be used for the CCD. The output DN's are usually

stored initially in computer memory and then moved to tape or disk for later manipulation.

The process of shifting each entire CCD row into the output register, shifting each pixel along within this register, and finally performing the voltage conversion of each pixel's stored charge by the A/D to produce a DN value is continued until the entire array of pixels has been read out. For large-format, modern CCD arrays, this process may take upwards of a few minutes to complete a single readout of the entire device. Note that for a  $2048 \times 2048$  CCD, the charge collected in the last pixel to be readout has to be transferred over four million times.

The array size of a single CCD, as well as the size of a given pixel on a device, is controlled by the current limitations of manufacturing. How large one can make a good quality, large-scale integrated circuit and how small one can make a MIS capacitor, both of which have demanding requirements for near perfect operation, set the scale of CCD and pixel sizes that are available. The interested reader is referred to Janesick & Elliott (1992) and Robinson (1988a) for further details.

## 2.2 CCD Types

When reading about CCDs, one of the most confusing issues can be the various terms listed in the literature or in commercial documents. Items such as backside illuminated, buried channel, and antiblooming are just a few. This section will provide a brief discussion to some of these terms while Chapter 3 will discuss CCD characteristics in detail. Further information can be found in the references listed in Appendix A. In particular, readers desiring a microscopic look at the electronic structures of a CCD integrated circuit are referred to the many SPIE articles listed. Some terms, such as quantum efficiency and full well capacity, will be used here without proper introduction. This will be rectified in the next chapter.

### 2.2.1 Surface Channel versus Buried Channel CCDs

As discussed above, charge stored within a pixel is moved from pixel to pixel during readout via changes in electrical potential between the biased gates. For surface channel CCDs, this charge movement occurs "on the surface" of the CCD, being transferred between overlapping gates. For example, each pixel in a three-phase CCD has three overlapping gates for each imaging cell or pixel. Each cell has an associated transfer efficiency and maximum possible transfer rate. The efficiency is near to,

but not quite, 100% for each cell (we will discuss this charge-transfer efficiency (CTE) further in Chapter 3). Typical rates of transfer for CCD television cameras (or video cameras) are several megahertz, while low-light level applications with cooled devices (i.e., those occurring in astronomy) use rates nearer to the kilohertz range (Eccles, Sim, & Tritton, 1983).

The major drawback to the surface channel CCD is the presence of trapping states that occur at the boundaries of the gates. These traps, which are caused by imperfections within the silicon lattice or the gate structures, cause a loss of transfer efficiency by trapping some of the electrons from other pixels as they pass by during readout. Traps within CCDs are therefore undesirable for faint light level applications such as astronomical imaging. One method that can be used to eliminate traps (although not in general use anymore because of the advent of buried channel devices; see below) is to raise the entire surface charge level of the CCD above the level needed to fill in any traps. This is accomplished by illuminating the CCD at a low level prior to exposure of the astronomical source. This technique, called a pre-flash or a fat-zero, allows any nonlinearities resulting from traps at low light levels to be avoided while only slightly increasing the overall noise level of the resulting image (Djorgovski, 1984; Tyson & Seitzer, 1988).

Deferred charge is another source of nonlinearity sometimes present at low light levels (Gjililand, 1992; Baum, Thomson, & Kreidl, 1981). Often referred to as charge skimming, charge depletion, or low light level nonlinearity, this particular worry has all but disappeared in modern, low readout noise devices. We will address nonlinearity in CCDs in a more general manner in Section 13.6.

A better solution than those just discussed is to move the charge from pixel to pixel during readout via a channel of semiconductor material that is away from the gate structures. This "buried channel" architecture results from the application of an additional layer of semiconductor material placed under the CCD surface. The buried channel greatly enhances the charge movement through the device by reducing the number of possible trap sites and by decreasing the transfer time between adjacent cells. Higher transfer rates (up to 100 MHz) and high CTE values (up to >99.995%) are easily accomplished with buried channel devices.

The price paid for the addition of the buried channel is that the total charge storage capacity for each pixel is reduced by 3 or 4 times that of a surface channel detector. However, since the operation at very low signal levels is much improved, the overall dynamic range and sensitivity of a buried channel device become much higher.

### 2.2.2 Front-Side and Back-Side Illuminated CCDs

CCDs are manufactured as single large-scale integrated devices. They have a front side, where all the gate structures and surface channel layers are deposited, and a back side, which is simply bulk silicon generally covered with a thin conductive layer of gold. CCDs used as front-side illuminated devices work as the name implies, that is, illumination occurs on the front of the CCD with the photons being absorbed by the silicon after passing directly through the surface gate structures. The device thickness is of order 300 microns from front to back making these chips relatively high in their susceptibility to detection of cosmic rays. Because the photons must first pass through the gate structures before they can be absorbed by the silicon, front-side illuminated CCDs have lower overall quantum efficiencies than the back-side devices (discussed below). However, front-side devices provide a flatter imaging surface and the actual CCD itself is easier to handle and work with. Figure 2.3 provides a schematic view of a single front-side illuminated pixel.

Back-side illuminated devices, also known as thinned devices, are again justly named. The CCD, after manufacture, is physically thinned to about 15 microns by various etching techniques (Lesser, 1994). The device is then mounted on a rigid substrate upside down and illuminated from behind. The incoming photons are now able to be absorbed directly into the bulk silicon pixels without the interference of the gate structures. The advantages in this type of CCD are that the relative quantum efficiency greatly exceeds that of a front-side device and the response of the detector to shorter wavelength light is improved since the photons




Fig. 2.3. Schematic view of a single front-side illuminated CCD pixel. The square labeled "front contact" is a representation of part of the overall gate structure. The letters "p" and "n" refer to regions within the pixel consisting of silicon doped with phosphorus and boron respectively.

no longer need to pass through the pixel gates. Disadvantages of back-side devices are in the areas of shallower pixel well depths (owing to the much smaller amount of material present per pixel), possible nonuniform thinning leading to surface and flat-field nonuniformities, and increased expense incurred by the thinning and mounting process.

In the next chapter, we will explore the idea of quantum efficiency further and provide an example of the large differences present in these two types of CCDs. Back-side devices generally have about twice the quantum efficiency for a given wavelength compared with front-side devices.

### 2.2.3 Interline and Frame Transfer CCDs

Interline transfer CCDs are specially constructed devices in which each column of imaging (active) pixels is paralleled by a light-shielded column of storage (or inactive) pixels. The device is used as a normal CCD but after the exposure ends, each light-sensitive column of data is quickly shifted into its neighboring light-shielded column. The shift occurs in only a few microseconds and so image smear is almost nonexistent and the need for a mechanical shutter is precluded. The shifted image data can then be clocked out of the device while the active columns are again integrating on the source of interest. Interline devices are used in many of the world's high-speed photon counting array detectors owing to their fast shift operation (Timothy, 1988).

The imaging area in an interline device is not continuous as each active column is paralleled by an inactive one; thus there is an immediate reduction in the overall extrinsic quantum efficiency by a factor of two. The electrode and gate structures of the active pixels are a bit different from those in a normal CCD, causing a further reduction in quantum efficiency and a somewhat different overall spectral response.

Frame transfer devices work in a manner similar to that of having two separate CCDs butted together. One half of the device is active (exposed to incoming light) and records the image while the other half is shielded and used as a storage device. The end of the integration sees the image data quickly transferred or shifted to the identical storage half, where it can be read out slowly as a new integration proceeds on the active side. Frame transfer CCDs are used in most commercial video and television cameras for which the active image is read out at video rates (30 frames per second). Astronomical imaging generally can not occur at such high rates because of photon starvation, but frame transfer devices have been used in a number of interesting instruments.




Fig. 2.4. Cartoon view of (top) a frame transfer CCD and (bottom) an inter-line CCD. From Eccles, Sim, & Tritton (1983).

Figure 2.4 provides an illustration of these two types of CCD. We note here that in both of these special purpose CCD types, the movement of the accumulated charge within each active area to the shielded, inactive area can be accomplished at very high rates. These rates are much higher than standard CCD readout rates because no sampling of the pixel charge or A/D conversion to an output number occurs until the final readout from the inactive areas. As we will see, the A/D process takes a finite time and is one cause of noise added to the output signal. Therefore, any on-chip shifting or summing of charge will introduce essentially no additional noise into the data.

### 2.2.4 Antiblooming CCDs

Observations of bright objects can lead to the collection of too many photoelectrons within one or more of the pixels of a CCD, causing the

affected pixels to overflow.<sup>†</sup> This condition will occur when the full well capacity of a pixel (see next chapter) is exceeded. The result of such saturation is termed bleeding and is easily noted on a CCD image as one or more bright pixels with a bright streak leading away from the region, generally in two opposite directions.<sup>‡</sup> Saturation effects can be irritating or, even worse, the streak can bleed into an object of interest rendering its measurement useless. Figure 2.5 shows a CCD image with a saturated star.

To counter pixel saturation behavior, one could simply integrate for a shorter period of time. This will cause the offending object to supply less photons per integration and thus not saturate any CCD pixels. However, to obtain the same exposure level, a number of frames would have to be taken and co-added together later. We will see in the next chapter that this is not always equivalent (in noise terms) to a single long exposure. Therefore, although shorter exposures are a simple approach, they are not always practical. Many times we wish to study a field in which objects of various brightness are interspersed, requiring the need for considerable dynamic range within a single image. Thus, the antiblooming gate was invented (Neely & Janesick, 1993).

Antiblooming gates have added electrical structures manufactured into the CCD itself, allowing saturated pixels to be "drained" off and controlled instead of bleeding into neighboring pixels, destroying their integrity. The offending pixels, in such a case, are still not usable, but the neighboring pixels are. The trade-off that occurs for a CCD with antiblooming gates is the following: About 30% of the active pixel area is lost (as it becomes the drain gate) giving a remaining active pixel area of only 70%. The use of such a gate structure is therefore equivalent to an effective QE reduction. Such additional gate areas also reduce the

<sup>†</sup> Think here of the buckets in the field during a rainstorm and what would happen if too much water fell into any given one of them.

<sup>‡</sup> The reader should be aware of the fact that CCDs can saturate in two different ways. As discussed here, saturation can occur if a given pixel is exposed to too much light and overflows, exceeding its full well capacity. This type of saturation leads to the result discussed here (see Figure 2.5), that is, bleeding of excess charge into neighboring pixels. The other way in which a CCD can saturate is if any given pixel collects more charge than can be represented by the output from the A/D converter (see Sections 2.1 and 3.1), even if the amount collected is below the full well capacity of the pixel. This type of saturation does not lead to charge bleeding but yields such oddities as flat-topped star profiles or constant (exact) pixel values over the region of A/D saturation. As the amount of charge collected by the pixels within a CCD increases, eventually one of these two types of saturation will set the usable upper limit for a given system. We will see that the linearity of a CCD (Sections 3.4 and 3.6) may also limit the maximum usable ADU level of a given observation.




Fig. 2.5. Two equal-length CCD exposures of a bright star (SAO 110456). The normal CCD exposure (a) shows typical bleeding caused by saturation within the CCD. The CCD exposure on the right (b) was made with an antiblooming CCD and clearly shows the much reduced bleeding from the bright star. From Neely & Jamesick (1993).

overall spatial resolution of any measurement (as the drain gates leave a gap between adjacent pixels) and lower the total pixel well depth by about one half. A rule of thumb here is that a typical antiblooming gate CCD requires about twice the integration time to reach the same signal level as a similar device without antiblooming gates.

Antiblooming gates are currently available on certain makes of CCD. When deciding if they are right for you, keep the following in mind: Longer exposures will be required for a given signal level, and if working in a spectral region of already low or reduced quantum efficiency, antiblooming gates will reduce the signal even further for a given

integration time. On the plus side, if bleeding of bright objects will be a detriment to your science, the antiblooming gate option may be a good choice.

### 2.2.5 Multipinned Phase CCDs

A method developed to achieve ultra-low dark current in a CCD is that of multipinned phase operation or MPP mode. CCDs operating in MPP mode are common in inexpensive, "off-the-shelf" devices for which thermoelectric cooling techniques are exploited. The MPP design relies on the inversion of all three clock voltages in connection with additional semiconductor doping (usually boron or phosphorus) of the phase three gate. The effect of this semiconductor architecture is to greatly lower the dark (thermal) current and this can even allow room temperature operation if the light levels are sufficient. Generally, astronomical low light level applications require CCDs to be cooled to near  $-100^{\circ}\text{C}$  before usable images can be obtained. Dark currents in MPP mode CCDs are near 0.01 electrons per second per pixel at temperatures of only  $-40$  to  $-65^{\circ}\text{C}$ , levels that are equal to non-MPP mode operation at  $-100^{\circ}\text{C}$ .

MPP mode CCDs have allowed a new revolution in CCD technology. Items such as digital cameras (CCDs operating at room temperature) are now commonplace. The lower thermal noise, even at room temperature, allows integration and readout (for high light level scenes) to be used to produce digital photographs. While a boon to higher temperature operation, the full well depth of each pixel is reduced by two to three times that of the same device in non-MPP mode. The higher temperature operation must be balanced against the much reduced dynamic range of the modified CCD. New design strategies and higher quality silicon wafer technology are making progress in restoring the full well depth while keeping the advantages of MPP operation. A CCD modified to run in MPP mode has altered electronic design and pixel structures and, therefore, cannot be made to operate normally. Thus, an MPP CCD does not allow one to decide which mode of operation to use nor does it provide the ability to switch between MPP and normal operation.

### 2.2.6 Orthogonal Transfer CCDs

We have seen that a typical CCD has three gates per pixel and that the readout operation is via clocking the pixel charge in a three-phase



fashion. The charge with a given pixel can only be shifted in the vertical direction, that is, from row to row, until it reaches the output register. A new type of CCD, the orthogonal transfer CCD (OTCCD), has been developed that allows each pixel's charge to be shifted both vertically and horizontally on the array (Burke et al., 1994).

The ability to move the charge in both directions within the OTCCD comes about through the use of a complex, four-phase mode of operation. Four gates are used in each pixel, two of which are triangular and make up the central part and two of which are rectangular but smaller in area and are split into pairs surrounding the pixel center. The larger number and size of the gates lead to a lower overall QE for an OTCCD and more detailed intra pixel QE effects are to be expected (Jordan, Deltorn, & Oates, 1994; Jordan, Deltorn, & Oates, 1993).

The first application of an OTCCD was to compensate for image motion during an integration, similar to a tip-tilt correction (Tonry, Burke, & Schechter, 1997). A  $1024 \times 512$  OTCCD device was built to allow one half of the OTCCD to image, quickly read out, and centroid on a bright star, while the other half was used to integrate on a scene of interest. As the bright star center wandered during the integration, the object frame half of the OTCCD was electronically shifted (both vertically and horizontally) many thousand times per 100 seconds in order to move the image slightly ( $\sim 0.5$  arcsec) and follow the bright star centroid position. The final result from the OTCCD was an image with much improved seeing and little loss of information.

The complex gate structure causes small dead spots within each pixel as well as allowing charge to be trapped in numerous pockets within a pixel during charge shifting. These two effects amount to about 3% flux losses, which are probably compensated for in the final image. Advances in gate structure design and impurity control in the device composition have eliminated most of these losses (Tonry, Burke, & Schechter, 1997). While still early in development, the OTCCD promises to be extremely useful for certain applications in astronomical imaging.

### 2.2.7 Superconducting Tunnel Junctions

Although not truly a CCD, the superconducting tunnel junction (STJ) is a very recent semiconductor device addition to the field of astronomical imaging. Using a layer of polycrystalline niobium (Nb) supercooled to a temperature of only a few hundred milli kelvin, a STJ diode acts as a photon detector (Perryman et al., 1994). Used initially as an X-ray

photon counter, the STJ is seeing new applications being developed in the areas of UV, visible, and IR light detection.

In Chapter 7, we will discuss how CCD detection of a high energy photon is accomplished and additionally how this detection and its associated photoelectron production provide information on the energy (wavelength) of the incoming light particle.<sup>†</sup> In general, light of low energy (UV to IR) does not produce multiple free electron pairs within a semiconductor; thus each incident photon (regardless of wavelength) looks like all the rest upon detection. Separation of spectral band must be accomplished by filtering the input beam.

Supercooled Nb STJ devices, however, allow UV, optical, and even IR photons to produce numerous pairs of free electrons, thus allowing a determination of the energy (wavelength) of each incoming particle. Recent application of a STJ device (Peacock et al., 1996) has shown sensitivity to wavelengths of 2,000 to 10,000 Å with a QE of near 50%. Spectral resolution of the incoming light is near 450 Å at present, but future devices made of other superconducting materials such as aluminum promise wavelength coverage from 10 Å to well over 1 micron at spectral resolutions of 80 Å.

The first astronomical test of a STJ occurred during the fall of 1998 at the Observatorio del Roque de los Muchachos on La Palma in the Canary Islands. A single STJ diode, placed at the telescope focus, produced a broad band spectrum of a star by simply imaging it.

Superconducting tunnel junctions are a far way off from being a routine detector in use at an observatory. After all it took CCDs nearly ten years after their introduction to become commonplace. Thus, although not currently available for use, STJs are a remarkable device that will no doubt become an invaluable asset for the future of astronomy.

### 2.3 CCD Coatings

To make use of the properties of CCDs in wavebands for which silicon is not sensitive or to enhance performance at specific wavelengths, CCDs can and have been coated with various materials.<sup>†</sup> These coating materials allow CCDs to become sensitive to photons normally too

<sup>†</sup> In the world of photon counting astronomy, using the number of photoelectrons produced to determine the wavelength (energy) of the incident photon is termed pulse-height spectroscopy.

<sup>†</sup> Nonoptical operation of a CCD can also be accomplished in this manner. See Chapter 7.

blue to allow absorption by the silicon. They generally consist of organic phosphors that down-convert incident UV light into longer wavelength photons, easily detected by the CCD. These vacuum deposited coatings are often used with front-side illuminated, thick CCDs (although not always; the *Hubble Space Telescope* WF/PC I CCDs, which were thinned devices, were phosphor coated) and can be viewed as an inexpensive method of increasing blue response without the cost and complexity of thinning the device.

A coronene phosphor has been commonly used in recent years to convert photons shortward of 3,900 Å to photons at a wavelength near 5,200 Å. The use of such a coating causes a "notch" in the overall QE curve as the CCD QE is falling rapidly near 4,000 Å but coronene does not become responsive until 3,900 Å. Another common phosphor coating, Lumogen, eliminates this QE notch, as it is responsive to wavelengths between 500 and 4,200 Å (see Figure 2.6). An interesting side note is that Lumogen is the commercial phosphorescent dye used in yellow highlighting pens. The QE of a coated CCD may increase in the UV and shortward regions by amounts upwards of 15–20%. Transparent in the visible and near-IR regions, properly deposited coatings can also act as antireflection (AR) coatings for a CCD. More on the use of coated CCDs is contained in Chapter 7.

The main problem encountered when using coated CCDs is that the deposited material tends to evaporate off the CCD when under vacuum conditions (see Chapter 7). A less important problem is the loss of spatial resolution due to the finite emission cone of the light generated in the coating. The specifics of coating materials and the process used to coat CCDs are both beyond the scope of this book. Appendix A and Janesick & Elliott (1992), Lessor (1990), Schaeffer et al. (1990), and Schenpp (1990) provide further readings in this area.

## 2.4 Analog-to-Digital Converters

Analog-to-digital (A/D) converters are not really a subject for this book. However, the output pixel values from a CCD (digital numbers stored in a computer) are completely determined by the method used to sample the charge on a pixel and how that sampling is converted into a digital number or data number to be stored in computer memory. A/D converters for use in CCDs, including linear and logarithmic types, have been reviewed in Opal (1988).




Fig. 2.6. The top plot shows QE curves for a *Hubble Space Telescope* WF/PC prototype CCD before and after being coated with Lumogen. Note the increased UV response of the coated CCD. The bottom plot shows the QE properties of a WF/PC prototype in the far-UV spectral region. Presented are two curves, one for a Coronene coated CCD and one for a Lumogen coated CCD. From Trauger (1990).

The output from each pixel in a CCD must be examined by electronic means in order to determine how much charge has been collected during a given integration and to convert that charge value into a digital number. As each packet is clocked out of the CCD output register, it is passed through an on-chip amplifier circuit built directly into the CCD

substrate. Next, the packet passes into an external low-noise solid state amplifier, to increase the signal strength, followed by input into the integrating amplifier. The output from the integrating amplifier is then sent directly into an A/D converter, which converts the analog voltage into a digital number to be stored in computer memory. The electronic circuitry just described, minus the on-chip amplifier, is often called the CCD head electronics and is usually located in a box very near or attached directly to the CCD dewar itself.

To measure the voltage of each charge packet, in preparation for conversion into a digital number, a difference measurement is performed between the charge packet itself and the sum of the charge packet plus a constant reset voltage. A well-regulated constant voltage source is needed in this step to supply the voltage to the electronics of the inverting plus integrating amplifier circuits. An integrating capacitor, connected across the integrating amplifier, first samples the reset voltage passed through the inverting amplifier for typically 20  $\mu$ s. A single pixel charge packet is then passed through the noninverting amplifier with the output being sampled by the same integrating capacitor for the same time period. Equal time sampling is critical as one has to worry about charge decay in the capacitor. Generally, however, the RC time constant for the capacitor is of order a few seconds, while the sampling is a few tens of microseconds. Both signals sampled by the capacitor contain the reset voltage, but only one has the pixel charge packet. Thus the difference formed by the two opposite signed voltages provides a very good estimate of the actual charge level within the sampled pixel. The value of the reset voltage, whatever it may be, is completely eliminated by this differencing. This technique, for sampling the amount of charge within each pixel, is used in almost all current day CCDs and is called correlated double sampling (CDS). Details of the CDS process and some CDS variations such as the dual slope process and triple correlated sampling are discussed in Jamesick & Elliott (1992), Opal (1988), Joyce (1992), and McLean (1997b).

The assignment of an output digital number to the value of the charge packet from each pixel is the job of the A/D converter. The charge stored in each pixel has an analog value (discrete only at the 1 electron level<sup>†</sup>) and the process of CDS decides how to assign each pixel's charge packet to a specific data number (DN) or analog-to-digital unit (ADU) (Merline

<sup>†</sup> Even though CCDs are called digital devices, the output from each pixel is really an analog voltage that must be measured and digitized by the A/D converter.

& Howell, 1995). As briefly mentioned in Chapter 1, this assignment is dependent on the device gain as follows.

If the gain for a CCD has a value of 10 electrons/ADU, the digitization process tries to divide each pixel's charge packet into units of 10 electrons each, assigning 1 ADU to each 10 electrons it measures. If, after division, there are less than 10 electrons left then the remaining charge is not counted and is lost and therefore unknown to the output pixel value. We will see (Sections 3.4 and Howell & Merline (1991)) that the gain of a device can have a great effect on its ability to provide good photometric information.

The ultimate readout speed of a given CCD depends on how fast the process of pixel examination and A/D conversion can take place. This entire process is called digitization. Modern large-format devices can, and often do, use two or four output amplifiers during readout (one at each device corner) providing a faster overall readout time. However, such readout schemes can introduce systematic offsets between each quadrant of the CCD owing to slight differences in the four output amplifiers and associated electronics. This effect is often seen in images produced by large-format CCDs, where one or more quadrants will show an apparently different background level. The readout speed of a CCD is also related to the number of bits of precision desired (A/D converters with fewer bits of precision work faster than those with more bits) and the acceptable readout noise level (faster pixel readout increases the introduced noise level). Current ultra-low noise A/D converters can distinguish the reference value from the collected charge in a given pixel at the 4 to 5 electron level.

Current CCD controllers (i.e., the hardware that controls the electronic functions of a CCD) are quite complex in design and yet simple in their use. Readout rates for CCDs are a prime concern as larger format chips may require excessively long readout times. Controllers such as those developed by San Diego State University (Leach, 1995) can read pixels out at a maximum rate of 10  $\mu$ s/pixel, but a practical rate of near 50  $\mu$ s/pixel is generally used to keep the read noise to a minimum (<10 electrons/pixel/read). Even at these seemingly fast rates (20 kHz), a 2048  $\times$  2048 CCD, containing over four million pixels, takes over 200 seconds to read out.

## 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

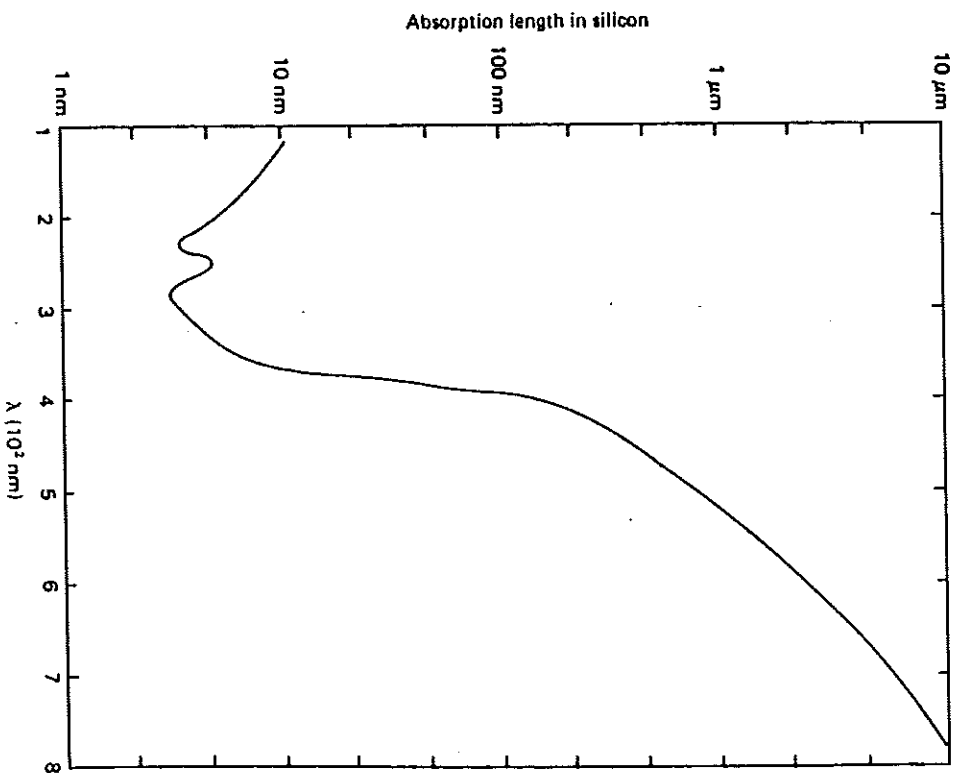


Fig. 3.1. The photon absorption length in silicon is shown as a function of wavelength in nanometers. From Reicke (1994).

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  $25 \text{ \AA}$  (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) QE (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,

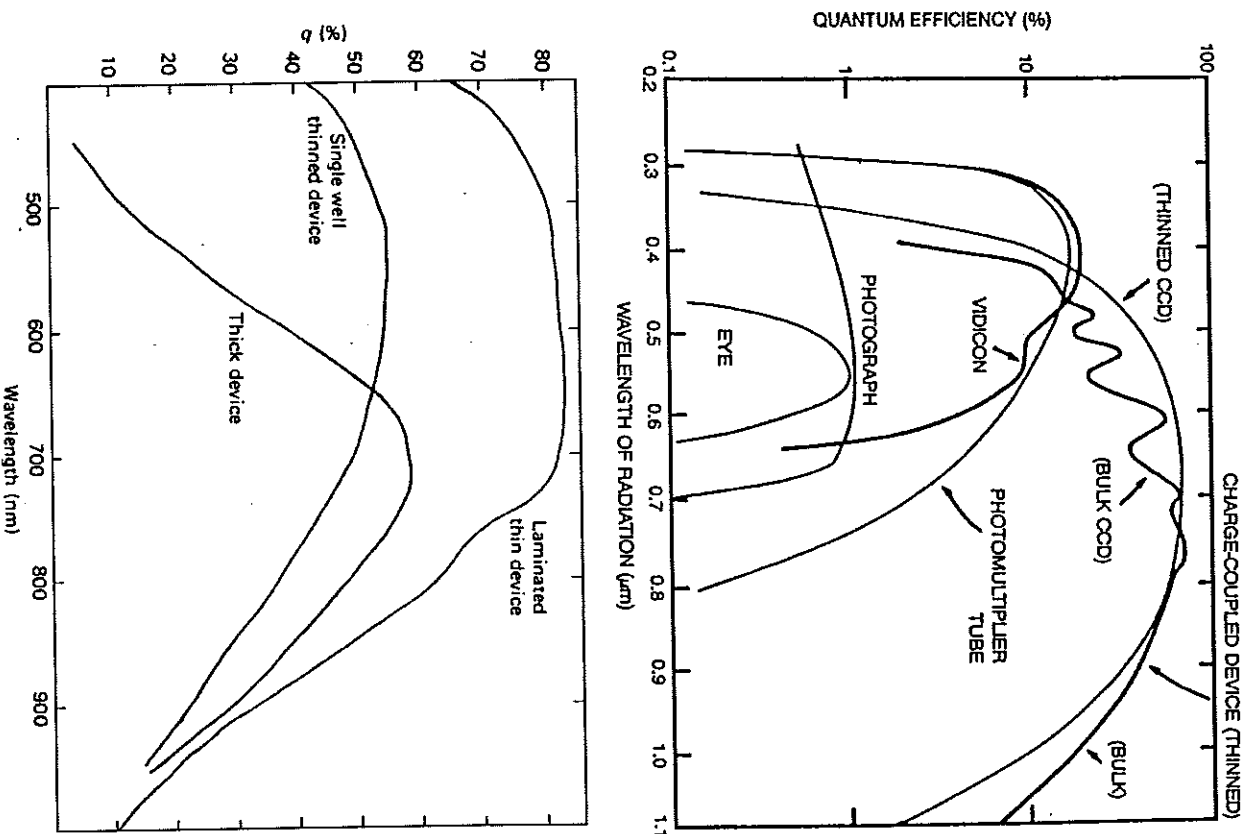


Fig. 3.2. The top panel shows QE curves for various devices indicating why CCDs are a quantum leap above all previous imaging devices. The failure of CCDs at wavelengths shorter than about  $3,500 \text{ \AA}$  has been completely eliminated via thinning or coating of the devices. The bottom plot shows 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 scope 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  $\text{cm}^{-2}$ ). At room temperature, the dark current of a typical CCD is near  $2.5 \times 10^4$  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 (LN<sub>2</sub>) 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 LN<sub>2</sub> 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 LN<sub>2</sub> 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  $\sim 4$  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  $1 \times 1$  binning or unbinned operation. Binning of  $2 \times 2$  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  $3 \times 1$  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  $1024 \times 1024$  CCD binned  $2 \times 2$  will have a final image size of only  $512 \times 512$  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  $40 \times 40$  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  $2048 \times 2048$  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,<sup>†</sup> 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  $10 e^-/\text{ADU}$ , 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

<sup>†</sup> 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 by viewing CCD image data of high numeric value but displayed as a signed integer 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,<sup>†</sup> 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 (Morfara & 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 ( $e^-/\text{ADU}$ ), common gain values range from 1 (photon counting) to 150 or more.

<sup>†</sup> Cosmic rays are not always cosmic! They can be caused by weakly radioactive materials used in the construction of CCD dewars (Florenth-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  $4.7 e^-/ADU$ .

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.<sup>†</sup> 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  $4.5 e^-/ADU$ , and a pixel full well capacity of 150,000 electrons. The linearity curve shown in Figure 3.6 is typical for a CCD, revealing

<sup>†</sup> 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  $1.17 \times 10^5$  photons (26,000 ADU), a number still well within the range of possible output values from the A/D.

The user must be aware of three possible factors that can limit the largest usable output pixel value in a CCD image: the two types of saturation that can occur (A/D saturation and exceeding a pixel's full well capacity; see Sections 2.2.4 and 2.4) and nonlinearity. For the CCD in the example shown in Figure 3.6, A/D saturation would occur at an output value of  $32,767 \cdot 4.5 = 147,451$  input photons. The pixel full well capacity is 150,000 electrons; thus pixel saturation will occur at a value of 33,333 ADU (150,000/4.5). Both full well and A/D saturations would produce noticeable effects in the output data such as bleeding or fat-topped stars. This particular example, however, illustrates the most dangerous type of situation that can occur in a CCD image. The non-linear region, which starts at 26,000 ADU, is entered into before either type of saturation can occur. Thus, the user could have a number of non-linear pixels (for example the peaks of bright stars) and be completely unaware of it. No warning bells will go off and no flags will be set in the output image to alert the user to this problem. The output image will be happy to contain (and the display will be happy to show) these non-linear pixel values and the user, if unaware, may try to use such values in the scientific analysis.

Thus, it is very important to know the linear range of your CCD and to be aware of the fact that some pixel values, even though they are not saturated, may indeed be within the nonlinear range and therefore unusable. Luckily, most professional grade CCDs reach one of the two types of saturation before they enter their nonlinear regime. Be aware, however, that this is almost never the case with low quality, inexpensive CCD systems that tend to use A/Ds with fewer bits, poor quality electronics, or low grade (impure) silicon. Most observatories have linearity curves available for each of their CCDs and some manufacturers include them with your purchase.<sup>†</sup> If uncertain of the linear range of a CCD, it is best to measure it yourself.

One method of obtaining a linearity curve for a CCD is to observe a field of stars covering a range of brightness. Obtain exposures of say 1, 2, 4, 8, 16, etc. seconds, starting with the shortest exposure needed to provide good signal-to-noise ratios (see Section 4.4) for most of the stars and ending when one or more of the stars begins to saturate. Since you have obtained a sequence that doubles the exposure time for each frame, you should also double the number of incident photons collected per star in each observation as well. Plots of the output ADU values

<sup>†</sup> A caution here is that the supplied linearity curve may only be representative of your CCD.

for each star versus the exposure time will provide you with a linearity curve for your CCD.

A common, although not always good, method of determining the value to use for the CCD gain, is to relate the full well capacity of the pixels within the array to the largest number that can be represented by your CCD A/D converter. As an example, we will use typical values for a Loral 512 × 1024 CCD in current operation at the Royal Greenwich Observatory. This CCD has 15-micron pixels and is operated as a back-side illuminated device with a full well capacity of 90,000 electrons per pixel. Using a 16-bit A/D converter (output values from 0 to 65,535) we could choose the gain as follows. Take the total number of electrons a pixel can hold and divide it by the total ADU values that can be represented:  $90,000/65,536 = 1.37$ . Therefore, a gain choice of  $1.4 e^-/ADU$  would allow the entire dynamic range of the detector to be represented by the entire range of output ADU values. This example results in a very reasonable gain setting, thereby allowing the CCD to produce images that will provide good quality output results.

As an example of where this type of strategy would need to be carefully thought out, consider a CCD system designed for a spacecraft mission in which the A/D converter only had 8 bits. A TI CCD was to be used, which had a full well capacity of 100,000 electrons per pixel. To allow imagery to make use of the total dynamic range available to the CCD, a gain value of 350 ( $\sim 100,000/2^8$ )  $e^-/ADU$  was used. This gain value certainly made use of the entire dynamic range of the CCD, allowing images of scenes with both shadow and bright light to be recorded without saturation. However, as we noted before, each gain step is discrete, thereby making each output ADU value uncertain by  $\pm$  the number of electrons within each A/D step. A gain of 350  $e^-/ADU$  means that each output pixel value has an associated uncertainty of  $\pm 1$  ADU, which is equal to, in this case,  $\pm 350$  electrons, a large error if precise measurements of the incident flux are desired. The uncertainty in the final output value of a pixel, which is caused by the discrete steps in the A/D output, is called digitization noise and is discussed in Merline & Howell (1995).

To understand digitization noise let us take, as an example, a CCD that can be operated at two different gain settings. If we imagine the two gain values to be either 5 or 200  $e^-/ADU$  and that a particular pixel collects 26,703 electrons (photons) from a source, we will obtain output values of 5,340 and 133 ADU respectively. Remember, A/D converters output only integer values and so any remainder is lost. In this example,

3 and 103 electrons respectively are lost as the result of the digitization noise of the A/D. More worrisome than this small loss of incident light is the fact that while each ADU step in the gain equals  $5 e^-/ADU$  case can only be incorrect by  $<5$  electrons, the gain equals  $200 e^-/ADU$  case will be uncertain by  $\pm 200$  electrons in each output ADU value. Two hundred electrons/pixel may not seem like much but think about trying to obtain a precise flux measurement for a galaxy that covers thousands of pixels on a CCD image or even a star that may cover tens of pixels. With an error of  $\pm 200$  electrons/pixel multiplied by tens or many more pixels, the value of a galaxy's surface brightness at some location or similarly a stellar magnitude would both be highly uncertain.

The gain of a particular CCD system is set by the electronics and is generally not changeable by the user or there may be but a few choices available as software options. Often, major observatories allow the user to pre-select, during the proposal process, which CCD (and thereby which gain, well depth, etc.) to use during the observing run. Table 3.1 gives an example of the choices one has if applying for time at Kitt Peak National Observatory to perform spectroscopy. Note that the A/D converter in all cases is 16 bits and that one of the CCDs reaches nonlinearity before it reaches A/D saturation. Thus, some care must be taken in deciding which CCD and gain combination you should use to accomplish your science objectives.

How A/D converters actually determine the assignment of the number to output for each pixel and whether the error in this choice is equally distributed within each ADU step is a detailed matter of interest but lies outside the scope of this book. A discussion of how the digitization noise affects the final output results from a CCD measurement is given in Merline & Howell (1995) and a detailed study of ADCs used for CCDs is given in Opal (1983).

In our above discussion of the gain of a CCD, we mentioned the term dynamic range a few times but did not offer a definition. The dynamic range of any device is the total range over which it operates or for which it is sensitive. For audio speakers this number is usually quoted in decibels, and this tradition is used for CCDs as well. The dynamic range of a CCD is given by the expression

$$D(\text{dB}) = 20 \times \log_{10}(\text{full well capacity/readnoise}).$$

A CCD with a full well capacity of 100,000 electrons per pixel and a read noise of 10 electrons would have a dynamic range of 80 dB.

Table 3.1. Spectrographs and CCDs available at Kitt Peak

Instrument	CCD	Scale ("/pixel)	Pixel Size ( $\mu$ )	Gain ( $e^-/ADU$ )	Read Noise ( $e^-$ )	A/D limit ( $e^-$ )	Linearity ( $e^-$ s at 0.1%)
4-m RC Spec	Tek 2K $\times$ 2K	0.69	24	3.1	4.0	201,500	220,000
4-m CryoCam	Loral 1K $\times$ 1K	0.84	15	0.8	15	52,000	150,000
2.1-m GoldCam	Loral 2K $\times$ 2K	0.78	15	1.4	8.5	91,000	80,000
WIYN Hydra	Tek 2K $\times$ 2K	0.78	24	1.7	4.3	110,500	190,000

Table 3.2. Various typical CCD properties

	RCA	Fairchild	TI	EEV	Thomson	Kodak	Loral (Ford)	Tektronix (SITE)
Pixel Format	320 $\times$ 512	380 $\times$ 488	800 $\times$ 800	2048 $\times$ 4096	1024 $\times$ 1024	765 $\times$ 510	3072 $\times$ 1024	1024 $\times$ 1024
Pixel Size ( $\mu$ )	30	18 $\times$ 30	15	13.5	19	9	15	24
Detector Size (mm)	10 $\times$ 15	7 $\times$ 15	12 $\times$ 12	28 $\times$ 55	20 $\times$ 20	7 $\times$ 4.5	46 $\times$ 15	25 $\times$ 25
Full Well Capacity ( $e^-$ )	350,000	>200,000	50,000	180,000	500,000	85,000	>140,000	>170,000
Dark Current ( $e^-/\text{pixel}/\text{hr}$ )	40	>1000	16	<2	22	1800	11	<5
Illumination	Front	Front	Back	Back	Front	Front	Back	Back
Peak QE (%)	70	12	70	85	40	40	90	75
Read Noise ( $e^-$ )	80	>150	15	7	5	12	9	6
CTE	0.99995	0.99975	0.999985	0.99999	0.99996	0.99997	0.99997	0.999999
Operating Temp. ( $^{\circ}\text{C}$ )	-100	-100	-120	-120	-110	-35	-120	-111
Gain ( $e^-/ADU$ )	13.5	50	5	1.1	5	2.3	1.2	1.3
Readout Time (s)	45	65	70	185	45	45	142	300

## 3.7 Summary

This chapter has concentrated on defining the terminology commonly used when discussing CCDs. The brief nature of this book does not allow the many more subtle effects, such as deferred charge, cosmic rays, or pixel traps, to be discussed further nor does it permit any discussion of the finer points of each of the above items. The reader seeking a deeper understanding of the details of CCD terminology (a.k.a., someone with a lot of time on his or her hands) is referred to the references given in this chapter and the detailed reading list in Appendix A. Above all, the reader is encouraged to find some CCD images and a workstation capable of image processing and image manipulation and to spend a few hours of time exploring the details of CCDs for themselves.

As a closing thought for this chapter, Table 3.2 provides a sample of the main properties of various CCDs. (Also see Appendix B.) The table lists CCDs that were used in the early 1980s up to the best available chips today. The sample shown tries to present the reader with an indication of the typical properties exhibited by CCDs. Included are those of different dimension, of different pixel size, having front and back illumination, cooled by LN<sub>2</sub> or thermoelectrically, and those available from different manufacturers. A scan of Table 3.2 allows one to see at a glance many of the corresponding relations we have just discussed, for example, the pixel well depth variation with front (thick) or back (thin) illuminated devices. In 1980, three companies made commercially available CCDs: RCA, GEC, and Fairchild. By 1994, over thirty commercially produced CCDs were on the market as well as at least five CCDs produced only for defense related work. Table 3.2 will serve us later on in the book when we are in need of typical CCD values for examples or calculations.

## CCD Imaging

This chapter will deal with the most basic use of a CCD, that of direct imaging. We will discuss a few more preliminaries such as flat fields, the calculation of gain and read noise for a CCD, and how the signal-to-noise value for a measurement is determined. The chapter then continues by providing a brief primer on the use of calibration frames in standard two-dimensional CCD data image reduction. Finally, we cover some aspects of CCD imaging itself including applications of wide-field imaging with CCD mosaics and CCD drift scanning.

## 4.1 Image or Plate Scale

One of the basic parameters of importance to a CCD user is that of knowing the plate scale of your image. Plate scale is a term that originated when photographic plates were used as the main imaging device and is often given in arcsec/mm. For a CCD user, however, a more convenient unit for the plate scale is arcsec/pixel. Clearly the conversion from one to the other is simple.

The focal ratio of a telescope is given by

$$f / = \frac{\text{focal length of primary mirror}}{\text{primary mirror diameter}},$$

where both values are in the same units and "primary mirror" would be replaced by "primary objective lens" for a refractor. Taking the focal length of the primary ( $f$ ) in mm and the CCD pixel size ( $\mu$ ) in microns, we can calculate the CCD plate scale as

$$P = \frac{206,265 \times \mu}{1,000 \times f} \quad (\text{arcsec/pixel}),$$

where 206,265 is the number of arcseconds in 1 radian and 1,000 is the conversion factor between millimeters and microns.

For a 1-m telescope of  $f/ = 7.5$ , the focal length ( $f$ ) of the primary would be 7,500 mm. If we were to use a Loral CCD with 15-micron pixels as an imager, the above expression would yield an image scale on the CCD of 0.41 arcsec/pixel. This image scale is usually quite a good value for direct imaging applications for which the seeing is near 1 or so arcseconds.

There are times, however, when the above expression for the plate scale of a CCD may not provide an accurate value. This could occur if there are additional optics within the instrument that change the final  $f$ -ratio in some unknown manner. Under these conditions, or simply as an exercise to check the above calculation, one can determine the CCD plate scale observationally. Using a few CCD images of close optical double stars with known separations, measurement of the center positions of the two stars and application of a bit of plane geometry will allow an accurate determination of the pixel-to-pixel spacings, and hence the CCD plate scale.

#### 4.2 Flat Fielding

To CCD experts, the term "flat field" can cause shivers to run up and down their spine. For the novice, it is just another term to add to the lexicon of CCD jargon. If you are in the latter category, don't be put off by these statements but you might want to take a minute and enjoy your thought of "How can a flat field be such a big deal?" In principle, obtaining flat field images and flat fielding a CCD image are conceptually easy to understand, but in practice the reality that CCDs are not perfect imaging devices sets in.

The idea of a flat field image is simple. Within the CCD, each pixel has a slightly different gain or QE value when compared with its neighbors. In order to flatten the relative response for each pixel to the incoming radiation, a flat field image is obtained and used to perform this calibration. Ideally, a flat field image would consist of uniform illumination of every pixel by a light source of identical spectral response to that of your object frames. That is, you want the flat field image to be spectrally and spatially flat. Sounds easy, doesn't it? Once a flat field image is obtained, one then simply divides each object frame by it and voilà: instant removal of pixel-to-pixel variations.

Before talking about the details of the flat fielding process and why it is not so easy, let us look at the various methods devised to obtain flat field

exposures with a CCD. All of these methods involve a light source that is brighter than any astronomical image one would observe. This light source provides a CCD calibration image of high signal-to-noise ratio. For imaging applications, one very common procedure used to obtain a flat field image is to illuminate the inside of the telescope dome (or a screen mounted on the inside of the dome) with a light source, point the telescope at the bright spot on the dome, and take a number of relatively short exposures so as to not saturate the CCD. Since the pixels within the array have different responses to different colors of light, flat field images need to be obtained through each filter that is to be used for your object observations. As with bias frames discussed in the last chapter, five to ten or more flats exposed in each filter should be obtained and averaged together to form a final or master flat field, which can then be used for calibration of the CCD. Other methods of obtaining a CCD flat field image include taking CCD exposures of the dawn or dusk sky or obtaining spatially offset images of the dark night sky; these can then be median filtered to remove any stars that may be present (Gilliland, 1992; Tobin, 1993; Massey & Jacoby, 1992; Tyson, 1990).

To allow the best possible flat field images to be obtained, many observatories have mounted a flat field screen on the inside of each dome and painted this screen with special paints (Massey & Jacoby, 1992) that help to reflect all incident wavelengths of light as uniformly as possible. In addition, most instrument user manuals distributed by observatories discuss the various methods of obtaining flat field exposures that seem to work best for their CCD systems. Illumination of dome flat field screens has been done by many methods, from a normal 35 mm slide projector, to special "hot filament" quartz lamps, to various combinations of lamps of different color temperature and intensity mounted like headlamps on the front of the telescope itself. Flat fields obtained by observation of an illuminated dome or dome screen are referred to as dome flats, while observations of the twilight or night sky are called sky flats.

CCD imaging and photometric applications use dome or sky flats as a means of calibrating out pixel-to-pixel variations. For spectroscopic applications, flat fields are obtained via illumination of the spectrograph slit with a quartz or other high intensity projector lamp housed in an integrating sphere (Wagner, 1992). The output light from the sphere attempts to illuminate the slit, and thus the grating of the spectrograph, in a similar manner to that of the astronomical object of interest. This type of flat field image is called a projector flat. While the main role of a flat field image is to remove pixel-to-pixel variations within the CCD, these calibration images will also compensate for any image vignetting

and for time-varying dust accumulation, which may occur on the dewar window and/or filters within the optical path.

Well, so far so good. So what is the big deal about flat field exposures? The problems associated with flat field images and why they are a topic discussed in hushed tones in back rooms may still not be obvious to the reader. There are two major concerns. One is that uniform illumination of every CCD pixel (spatially flat) to one part in a thousand is often needed but in practice is very hard to achieve. Second, QE variations within the CCD pixels are wavelength dependent. This wavelength dependence means that your flat field image should have the exact wavelength distribution over the bandpass of interest (spectrally flat) as that of each and every object frame you wish to calibrate. Quartz lamps and twilight skies are not very similar at all in color temperature (i.e., spectral shape) to that of a dark nighttime sky filled with stars and galaxies.† Sky flats obtained of the dark nighttime sky would seem to be our savior here, but these types of flat fields require long exposures to get the needed signal-to-noise ratio and multiple exposures with spatial offsets to allow digital filtering (e.g., median) to be applied in order to remove the stars. In addition, the time needed to obtain (nighttime) sky flats is likely not available to the observer who generally receives only a limited stray at the telescope. Thus, whereas very good calibration data lead to very good final results, the fact is that current policies of telescope scheduling often mean that we must somehow compromise the time used for calibration images with that used to collect the astronomical data of interest.

Within the above detailed constraints on a flat field image, it is probably the case that obtaining a perfect, color-corrected flat field is an impossibility. But all is not lost. Many observational projects do not require total perfection of a flat field over all wavelengths or over the entire two-dimensional array. Stellar photometry resulting in differential measurements or on-band/off-band photometry searching for particular emission lines are examples for which one only needs to have good flat field information over small spatial scales on the CCD. However, a project with end results of absolute photometric calibration over large spatial extents (e.g., mapping of the flux distribution within the spiral arms of an extended galaxy) does indeed place stringent limits on flat fielding requirements. For such demanding observational programs, some observers have found that near-perfect flats can be obtained through the

† One good sky region for twilight flats has been determined to be an area 13 degrees east of zenith just after sunset (Chronney & Hasselbacher, 1996).

use of a combination of dome and sky flats. This procedure combines the better color match and low-spatial frequency information from the dark night sky with the higher signal-to-noise, high spatial frequency information of a dome flat. Experimentation to find the best method of flat fielding for a particular telescope, CCD, and filter combination, as well as for the scientific goals of a specific project, is highly recommended.

A summary of the current best wisdom on flat fields depends on who you talk to and what you are trying to accomplish with your observations. The following advice is one person's view.

What does the term "a good flat field" mean? An answer to that question is: A good flat field allows a measurement to be transformed from its instrumental values into numeric results in a standard system that results in an answer that agrees with other measurements made by other observers. For example, if two observers image the same star, they both observe with a CCD using a V filter, and they each end up with the final result of  $V = 14.325$  magnitudes in the Johnson system then, assuming this is an accurate result, one may take this as an indication of the fact that each observer used correct data reduction and analysis procedures (including their flat fielding) for the observations.

The above is one way to answer the question, but it still relies on the fact that observers need to obtain good flat fields. Without them, near perfect agreement of final results is unlikely. While the ideal flat field would uniformly illuminate the CCD such that every pixel would receive equal amounts of light in each color of interest, this perfect image is generally not produced with dome screens, the twilight sky, or projector lamps within spectrographs. This is because good flat field images are all about color terms. That is, the twilight sky is not the same color as the nighttime sky, neither of which are the same color as a dome flat. If you are observing red objects, you need to worry more about matching the red color in your flats; for blue objects you worry about the blue nature of your flats. Issues to consider include the fact that if the Moon is present, the sky is bluer than when the Moon is absent, dome flats are generally reddish owing to their illumination by a quartz lamp of relatively low filament temperature, and so on. Thus, just as in photometric color transformations, the color terms in flat fields are all important. One needs to have a flat field that is good, as described above, plus one that also matches the colors of interest to the observations at hand.

Proper techniques for using flat fields as calibration images will be discussed in Section 4.5. Appendix A offers further reading on this subject and the material presented in Djorgovski (1984), Sterken (1995),



Tyson (1990), and Gudehus (1990) is of particular interest concerning flat fielding techniques.

### 4.3 Calculation of Read Noise and Gain

We have talked about bias frames and flat field images in the text above and now wish to discuss the way in which these two types of calibration data may be used to determine the read noise and gain for a CCD.

Noted above, when we discussed bias frames, was the fact that a histogram of such an image (see Figure 3.5) should produce a Gaussian distribution with a width related to the read noise and the gain of the detector. Furthermore, a similar relation exists for the histogram of a typical flat field image (see Figure 4.1). The mean level in the flat field shown in Figure 4.1 is  $\bar{F} = 6,950$  ADU and its width (assuming it is




Fig. 4.1. Histogram of a typical flat field image. Note the fairly Gaussian shape of the histogram and the slight tail extending to lower values. For this R-band image, the filter and dewar window were extremely dusty leading to numerous out of focus "doughnuts" (see Figure 4.4), each producing lower than average data values.

perfectly Gaussian (Massey & Jacoby, 1992)) will be given by

$$\sigma_{\text{ADU}} = \frac{\sqrt{\bar{F}} \cdot \text{Gain}}{\text{Gain}}$$

We have made the assumption in this formulation that the Poisson noise of the flat field photons themselves is much greater than the read noise. This is not unreasonable at all given the low values of read noise in present day CCDs.

Let us now look at how bias frames and flat field images can be used to determine the important CCD properties of read noise and gain. Using two bias frames and two flat field images, designated 1 and 2, we can proceed as follows. Determine the mean pixel value within each image.† We will call the mean values of the two bias frames  $B_1$  and  $B_2$  and likewise  $\bar{F}_1$  and  $\bar{F}_2$  will be the corresponding values for the two flats. Next, create two difference images ( $B_1 - B_2$  and  $F_1 - F_2$ ) and measure the standard deviation of these image differences:  $\sigma_{B_1-B_2}$  and  $\sigma_{F_1-F_2}$ . Having done that, the gain of your CCD can be determined from the following:

$$\text{Gain} = \frac{(\bar{F}_1 + \bar{F}_2) - (B_1 + B_2)}{\sigma_{F_1-F_2}^2 - \sigma_{B_1-B_2}^2},$$

and the read noise can be obtained from

$$\text{Read Noise} = \frac{\text{Gain} \cdot \sigma_{B_1-B_2}}{\sqrt{2}}.$$

### 4.4 Signal-to-Noise Ratio

Finally we come to one of the most important sections in this book, the calculation of the signal-to-noise (S/N) ratio for observations made with a CCD.

Almost every article written that contains data obtained with a CCD and essentially every observatory user manual about CCDs contain some version of an equation used for calculation of the S/N of a measurement. S/N values quoted in research papers, for example, do indeed give the reader a feel for the level of goodness of the observation (i.e., a S/N of

† Be careful here not to use edge rows or columns, which might have very large or small values owing to CCD readout properties such as amplifier turn on/off (which can cause spikes). Also, do not include overscan regions in the determination of the mean values.

100 is probably good while a  $S/N$  of 3 is not), but rarely do the authors discuss how they performed such a calculation.

The equation for the  $S/N$  of a measurement made with a CCD is given by

$$\frac{S}{N} = \frac{N_*}{\sqrt{N_* + n_{\text{pix}}(N_S + N_D + N_R^2)}}$$

unofficially named the “CCD Equation” (Mortara & Fowler, 1981). Various formulations of this equation have been produced (e.g., Newberry (1991) and Gullixson (1992)), all of which yield the same answers of course, if used properly. The “signal” term in the above equation,  $N_*$ , is the total number of photons<sup>†</sup> (signal) collected from the object of interest.  $N_*$  may be from one pixel (if determining the  $S/N$  of a single pixel as sometimes is done for a background measurement), or  $N_*$  may be from several pixels, such as all of those contained within a stellar profile (if determining the  $S/N$  for the measurement of a star) or  $N_*$  may even be from say a rectangular area of  $X$  by  $Y$  pixels (if determining the  $S/N$  in a portion of the continuum of a spectrum).

The “noise” terms in the above equation are the square roots of  $N_*$ , plus  $n_{\text{pix}}$  (the number of pixels under consideration for the  $S/N$  calculation) times the contributions from  $N_S$  (the total number of photons per pixel from the background or sky),  $N_D$  (the total number of dark current electrons per pixel), and  $N_R^2$  (the total number of electrons per pixel resulting from the read noise.<sup>‡</sup>)

For those interested in more details of each of these noise terms, how they are derived, and why each appears in the CCD Equation, see Merline & Howell (1995). In our short treatise, we will remark on some of the highlights of that paper and present an improved version of the CCD Equation. However, let’s first make sense out of the equation just presented.

For sources of noise that behave under the auspices of Poisson statistics (which includes photon noise from the source itself), we know that for a signal level of  $N$ , the associated 1 sigma error ( $1\sigma$ ) is given by

<sup>†</sup> Throughout this book we have and will continue to use the terms photons and electrons interchangeably when considering the charge collected by a CCD. In optical observations, every photon that is collected within a pixel produces a photoelectron; thus they are indeed equivalent. When talking about observations, it seems logical to talk about star or sky photons, but for dark current or read noise discussions, the number of electrons measured seems more useful.

<sup>‡</sup> Note that this noise source is not a Poisson noise source but a shot noise; therefore it enters into the noise calculation as the value itself, not the square root of the value as Poisson noise sources do.

$\sqrt{N}$ . The above equation for the  $S/N$  of a given CCD measurement of a source can thus be seen to be simply the signal ( $N_*$ ) divided by the summation of a number of Poisson noise terms. The  $n_{\text{pix}}$  term is used to apply each noise term on a per pixel basis to all of the pixels involved in the  $S/N$  measurement and the  $N_R$  term is squared since this noise source behaves as shot noise, rather than being Poisson-like (Mortara & Fowler, 1981). We can also see from the above equation that if the total noise for a given measurement  $\sqrt{N_* + n_{\text{pix}}(N_S + N_D + N_R^2)}$  is dominated by the first noise term,  $N_*$  (i.e., the noise contribution from the source itself), then the CCD Equation becomes

$$\frac{S}{N} = \frac{N_*}{\sqrt{N_*}} = \sqrt{N_*},$$

yielding the expected result for a measurement of a single Poisson behaved value.

This last result is useful as a method of defining what is meant by a “bright” source and a “faint” source. As a working definition, we will use the term bright source to mean a case for which the  $S/N$  errors are dominated by the source itself (i.e.,  $S/N \sim \sqrt{N_*}$ ), and we will take a faint source to be the case in which the other error terms are of equal or greater significance compared with  $N_*$ , and therefore the complete error equation (i.e., the CCD Equation) is needed.

The CCD equation above provides the formulation for a  $S/N$  calculation given typical conditions and a well-behaved CCD. For some CCD observations, particularly those that have high background levels, faint sources of interest, poor spatial sampling, or large gain values, a more complete version of the error analysis is required. We can write the complete CCD Equation (Merline & Howell, 1995) as

$$\frac{S}{N} = \frac{N_*}{\sqrt{N_* + n_{\text{pix}} \left( 1 + \frac{n_{\text{pix}}}{n_B} \right) (N_S + N_D + N_R^2 + G^2 \sigma_f^2)}}$$

This form of the  $S/N$  equation is essentially the same as that given above, but two additional terms have been added. The first term,  $(1 + n_{\text{pix}}/n_B)$ , provides a measure of the noise incurred as a result of any error introduced in the estimation of the background level on the CCD image. The term  $n_B$  is the total number of background pixels used to estimate the mean background (sky) level. One can see that small values of  $n_B$  will introduce the largest error as they will provide a poor estimate of the mean level of the background distribution. Thus, very large values of  $n_B$  are to be preferred but clearly some trade-off must be made between

providing a good estimate of the mean background level and the use of pixels from areas on the CCD image that are far from the source of interest or possibly of a different character.

The second new term added into the complete S/N equation accounts for the error introduced by the digitization noise within the A/D converter. From our discussion of the digitization noise in Chapter 3, we noted that the error introduced by this process can be considerable if the CCD gain has a large value. In this term,  $G^2\sigma_f^2$ ,  $G$  is the gain of the CCD (in electrons/ADU) and  $\sigma_f^2$  is an estimate of the 1 sigma error introduced within the A/D converter<sup>†</sup> and has a value of approximately 0.289 (Merline & Howell, 1995).

In practice for most CCD systems in use and for most observational projects, the two additional terms in the complete S/N equation are often very small error contributors and can be ignored. In the instances for which they become important – for example, cases in which the CCD gain has a high value (e.g., 100 electrons/ADU), the background level can only be estimated with a few pixels (e.g., less than 200), or the CCD data are of poor pixel sampling (see Section 5.6) – ignoring these additional error terms will lead to an overestimation of the S/N value obtained from the CCD data.

Let us work through an example of a S/N calculation given the following conditions. A 300 second observation is made of an astronomical source with a CCD detector attached to a 1-m telescope. The CCD is a Thomson 1024 × 1024 device with 19-micron pixels and it happens that in this example the telescope has a fast  $f$ -ratio such that the plate-scale is 2.6 arcsec/pixel.<sup>‡</sup> From Table 3.2 we find that, for this particular CCD, the read noise is 5 electrons/pixel/read, the dark current is 22 electrons/pixel/hour, and the gain ( $G$ ) is 5 electrons/ADU. Using 200 background pixels surrounding our object of interest from which to estimate the mean background sky level, we take a mean value for  $N_B$  of 620 ADU/pixel. We will further assume here (for simplicity) that the CCD image scale is such that our source of interest falls completely within 1 pixel (good seeing!) and that after background subtraction (see Section 5.1), we find a value for  $N_*$  of 24,013 ADU. Ignoring the two additional minor error terms discussed above (as the gain is very small

and  $n_B = 200$  is quite sufficient in this case), we can write the CCD Equation as

$$\frac{S}{N} = \frac{24,013(\text{ADU}) \cdot G}{\sqrt{24,013(\text{ADU}) \cdot G + (1) \cdot (620(\text{ADU}) \cdot G + 1.8 + 5^2(e^{-}))}}$$

Note that all of the values used in the calculation of the S/N are in electrons, *not* in ADUs. The S/N value calculated for this example is  $\sim 342$ , a very high S/N. With such a good S/N measurement, one might suspect that this is a bright source. If we compare  $\sqrt{N_*}$  with all the remaining error terms, we see that indeed this measurement has its noise properties dominated by the Poisson noise from the source itself and the expression  $S/N \sim \sqrt{N_*} = 346$  works well here.

While the S/N of a measurement is a useful number to know, at times we would prefer to quote a standard error for the measurement as well. Using the fact that  $S/N = 1/\sigma$ , where  $\sigma$  is the standard deviation of the measurement, we can write

$$\sigma_{\text{magnitudes}} = \frac{1.0857\sqrt{N_* + p}}{N_*}$$

In this expression,  $p$  is equal to  $n_{\text{pix}}(1 + n_{\text{pix}}/n_B)(N_S + N_D + N_R^2 + G^2\sigma_f^2)$ , the same assumptions apply concerning the two “extra” error terms, and the value of 1.0857 is the correction term between an error in flux (electrons) and that same error in magnitudes (Howell, 1993). We again see that if the Poisson error of  $N_*$  itself dominates, the term  $p$  can be ignored and this equation reduces to that expected for a  $1\sigma$  error estimate in the limiting case of a bright object.

Additionally, one may be interested in a prediction of the S/N value likely to be obtained for a given CCD system and integration time.  $N_*$  is really  $N \cdot t$ , where  $N$  is the count rate in electrons (photons) per second (for the source of interest) and  $t$  is the CCD integration time. Noting that the integration time is implicit in the other quantities as well, we can write the following (Masse, 1990):

$$\frac{S}{N} = \frac{Nt}{\sqrt{Nt + n_{\text{pix}}(N_S t + N_D t + N_R^2)}}$$

in which we have again ignored the two minor error terms. This equation illustrates a valuable rule of thumb concerning the S/N of an observation:  $S/N \propto \sqrt{t}$ , not to  $t$  itself. Solving the above expression for  $t$  we find

$$t = \frac{-B + (B^2 - 4AC)^{1/2}}{2A}$$

<sup>†</sup> The parameter  $\sigma_f^2$  and its value depend on the actual internal electrical workings of a given A/D converter. We assume here that for a change level that is half way in between two output ADU steps (that is, 1/2 of a gain step), there is an equal chance that it will be assigned to the lower or to the higher ADU value when converted to a digital number. See Merline & Howell (1995) for further details.

<sup>‡</sup> Using the results from Section 4.1, what would the  $f$ -ratio of this telescope be?

where  $A = N^2$ ,  $B = -(S/N)^2 (N + n_{\text{pix}}(N_S + N_D))$ , and  $C = -(S/N)^2 n_{\text{pix}} N_S^2$ . Most instrument guides available at major observatories provide tables that list the count rate expected for an ideal star (usually 10th magnitude and of 0 color index) within each filter and CCD combination in use at each telescope. Similar tables provide the same type of information for the observatory spectrographs as well. The tabulated numeric values, based on actual CCD observations, allow the user, via magnitude, seeing, or filter width, to scale the numbers to a specific observation and predict the  $S/N$  expected as a function of integration time.

#### 4.5 Basic CCD Data Reduction

The process of standard CCD image reduction makes use of a basic set of images that form the core of the calibration and reduction process (Gullixson, 1992). The types of images used are essentially the same (although possibly generated by different means) in imaging, photometric, and spectroscopic applications. This basic set of images consists of three calibration frames – bias, dark, and flat field – and the data frames of the object(s) of interest. Table 4.1 provides a brief description of each image type and Figures 4.2-4.5 show examples of typical bias, dark, flat field, and object CCD images. Note that a CCD dark frame contains not only information on the level and extent of the dark current but also includes bias level information as well.

The use of the basic set of calibration images in the reduction of CCD object frames is as follows. First, subtract a mean bias frame (or dark frame if needed<sup>†</sup>) from your object frame. Then, divide the resulting image by a (bias subtracted) mean flat field image. That's all there is to it! These two simple steps have corrected your object frame for bias level, dark current (if needed), and nonuniformity within each image pixel. During the analysis of your object frames, it is likely that the background or sky contribution to the image will need to be removed or accounted for in some manner. This correction for the background sky

<sup>†</sup> The need for dark frames instead of simply bias frames depends entirely on the level of dark current expected during an integration or the stability of the dark current from integration to integration. The first situation depends on the operating temperature of the CCD. LN2 systems have essentially zero dark current, and thus bias frames are all that is needed. Inexpensive and thermoelectrically cooled CCD systems fall into the category of generally always needing dark frames as part of the calibration process.

Table 4.1. Types of CCD images

CCD Image Type	Image Description
Bias	This type of CCD image has an exposure time of zero seconds. The shutter remains closed and the CCD is simply readout. The purpose of a bias or zero frame is to allow the user to determine the underlying noise level within each data frame. The bias value in a CCD image is usually a low spatial frequency variation throughout the array, caused by the CCD on-chip amplifiers. This variation should remain constant with time. The rms value of the bias level is the CCD read noise. A bias frame contains both the DC offset level (overscan) and the variations on that level. The nature of the bias variations for a given CCD are usually column-wise variations, but may also have small row-wise components as well. Thus, a 2-D, pixel-by-pixel subtraction is often required. A single bias frame will not sample these variations well in a statistical fashion, so an average bias image of 10 or more single bias frames is recommended.

\*\*\*\*\*

Dark CCD dark frames are images taken with the shutter closed but for some time period, usually equal to that of your object frames. That is, if one is planning to dark correct a 45 second exposure, a 45 second dark frame would be obtained. Longer dark frames can often be avoided using the assumption that the dark current increases linearly with time and a simple scaling can be applied. However, this is not always true. Dark frames are a method by which the thermal noise (dark current) in a CCD can be measured. They also can give you information about bad or "hot" pixels that exist as well as provide an estimate of the rate of cosmic ray strikes at your observing site. Observatory class CCD cameras are usually cooled with LN2 to temperatures at which the dark current is essentially zero. Many of these systems therefore do not require the use of a dark exposure CCD frames in the calibration process. Thermoelectrically cooled systems are not cooled to low enough temperatures such that one may ignore the dark current. In addition, these less expensive models often have poor temperature stability allowing the dark current to wander a bit with time. Multiple darks averaged together are the best way to produce the final dark calibration frame. Note that if dark frames are used, the bias level of the CCD is present in them as well, and therefore separate bias frames are not needed.

(Cont.)

Table 4.1. (Continued)

CCD Image Type	Image Description
Flat Field	Flat field exposures are used to correct for pixel-to-pixel variations in the CCD response as well as any nonuniform illumination of the detector itself. Flat fields expose the CCD to light from either a dome screen, the twilight sky, the nighttime sky, or a projector lamp in an attempt to provide a high S/N, uniformly illuminated calibration image. For narrow-band imaging, flats are very helpful in removing fringing which may occur in object frames. Flat field calibration frames are needed for each color, wavelength region, or different instrumental setup used in which object frames are to be taken. A good flat should remain constant to about 1%, with 2% or larger changes being indicators of a possible problem. As with the other calibration frames, at least 5 or more flat fields should be taken and averaged to produce the final flat used for image calibration.

\*\*\*\*\*

**Object**

These are the frames containing the astronomical objects of interest. They are of some exposure length from 1 second or less up to many hours, varying for reasons of type of science, brightness of object, desired temporal sampling, etc. Within each object image pixel is contained contributions from the object and/or sky, read noise, thermally generated electrons, and possibly contributions from cosmic rays. Each pixel responds similarly but not exactly to the incident light, so nonuniformities must be removed. All of the noise and spatial factors are correctable to very low levels via standard CCD reductions as described in the text.

level in your image frames is performed as part of each specific analysis step using "sky" regions in the object frame itself and is not removed or corrected for with some sort of separate "sky" frame. In equational form, the calibration process can be written as

$$\text{Final Reduced Object Frame} = \frac{\text{Raw Object Frame} - \text{Bias Frame}}{\text{Flat Field Frame}}$$

where, again, the flat field image has already been bias subtracted and the bias frame would be replaced by a dark frame when appropriate.




Fig. 4.2. Shown is a typical CCD bias frame. The histogram of this image was shown in Figure 3.5. Note the overall uniform structure of the bias frame.

#### 4.6 CCD Imaging

This section details issues related to the application of using CCDs to produce images of an extended area of the sky. Examples of this type of CCD observation are multicolor photometry of star clusters, galaxy imaging to isolate star forming regions within spiral arm structures, deep wide-field searches for quasars, and extended low surface brightness mapping of diffuse nebulae. Use of the areal nature of a CCD introduces some additional issues related to the calibration procedures and the overall cosmetic appearance as any spurious spatial effects will have implications on the output result. We briefly discuss here a few new items that are of moderate concern in two-dimensional imaging and then move on to the topic of wide-field imaging with CCD mosaics.




Fig. 4.3. Shown is a typical CCD dark frame. This figure shows a dark frame for a Kodak CCD operating in MPP mode and thermoelectrically cooled. Notice the nonuniform dark level across the CCD, being darker (greater ADU values) on the top. Also notice the two prominent partial columns with higher dark counts, which extend from the top toward the middle of the CCD frame. These are likely to be column defects in the CCD that occurred during manufacture, but with proper dark subtraction they are of little consequence. The continuation of the figure shows the histogram of the dark frame. Most of the dark current in this 180 second exposure is uniformly distributed near a mean value of 180 ADU with a secondary maximum near 350 ADU. The secondary maximum represents a small number of CCD pixels that have nearly twice the dark current of the rest, again most likely due to defects in the silicon lattice. As long as these increased dark current pixels remain constant, they are easily removed during image calibration.




Fig. 4.3. (Continued)

#### 4.6.1 CCD Fringing and Cosmetic Effects

We mentioned earlier that observations of monochromatic (or nearly so) light can cause a pattern of fringes to occur on a CCD image. These fringes, which are essentially Newton's rings, are caused by interference between light waves that reflect within the CCD or long wavelength light that passes through the array and reflects back into the array. Fringing may occur for CCD observations in the red part of the optical spectrum, when narrow-band filters are used, or if observations are made of a spectral regime (e.g., the I-band) that contains strong narrow emission lines.

The troubling aspect with fringing in CCD data is that it is often the case that the fringe pattern does not occur in the flat field frames or the pattern is highly variable throughout the night. Without a pattern match between the flats and the image data, fringe removal will not occur during image calibration, and residual fringes will remain in the final object images. One of the major causes of CCD fringing is the night sky emission lines that occur in the Earth's upper atmosphere (Pecker, 1970). These night sky lines are mainly attributed to OH transitions in the atmosphere and are powered by sunlight during the day. Since they are forbidden transitions they have long decay lifetimes and are




Fig. 4.4. Shown is a typical CCD flat field image. This is an R-band flat field image for a  $1,024 \times 1,024$  Loral CCD. The numerous "doughnuts" are out of focus dust specks present on the dewar window and the filter. The varying brightness level and structures are common in flat field images. As seen in the histogram of this image (Figure 4.1) this flat field has a mean level near 6,950 ADU, with an approximate dispersion of (FWHM) 400 ADU.

very narrow spectrally. In addition, owing to upper atmosphere motions, OH concentrations, and their long decay times, these emission lines are highly variable in time and strength, even within a given night. Dealing with fringes that occur in CCD data can be a difficult problem but one for which a number of solutions exist (Wagner, 1992; Broadfoot & Kendall, 1968).

Additionally, cosmetic effects such as bad pixels, hot pixels (LEDS), or dead columns will be present and can mar a CCD image. Not only do these flaws spoil the beauty of the two-dimensional data, they can cause problems during calibration and analysis by hindering software processes




Fig. 4.5. Shown is a typical CCD object frame showing a star field. This image has been properly reduced using bias frame subtraction and division by a flat field image. Note how the background is of a uniform level and distribution; all pixel-to-pixel nonuniformities have been removed in the reduction process. The stars are shown as black in this image and represent R magnitudes of 15th (brightest) to 20th (faintest). The histogram shown in the remainder of the figure is typical for a CCD object frame after reduction. The large grouping of output values on the left (values less than about 125 ADU) are an approximate Gaussian distribution of the background sky. The remaining histogram values (up to 1,500 ADU) are the pixels that contain signal levels above the background (i.e., the pixels within the stars themselves!).

and not allowing correct flux estimates to be made for the pixels that they affect. Procedures for the removal of or correction for these types of problems can be applied during image calibration and reduction. They are highly specialized tasks, depend on the desired output science goals, and generally are specific to a particular CCD, instrument, or type of observation being made. A complete discussion of all of these topics lies




Fig. 4.5. (Continued)

beyond our space limitations but the interested reader will find discussions of such corrections in Gilliland (1992), Gullixson (1992), Djorgovski (1984), Janesick et al. (1987a), and numerous specific instrument manuals and reference papers concerning the finer points of specific CCD related issues (see Appendix A).

#### 4.6.2 Wide-Field CCD Imaging

With the advent of large-footprint CCDs and the construction of CCD mosaic arrays containing many chips, wide-field imaging is becoming one of the major applications in astronomy today. The Wide Field Camera (WFC) on the 2.5-m Isaac Newton Telescope (INT) consists of four  $2K \times 4K$  EEV CCDs plus an additional  $2K \times 2K$  Loral CCD (operating in frame transfer mode for autoguiding), mounted at prime focus. The pixel size for each CCD is 15 microns and a ten minute exposure through a broad-band filter yields a faint limit of near 25th magnitude. Figure 4.6 presents a view of the INT Wide Field Camera instrument.

By far the most ambitious wide-field imaging project is the MEGA-CAM instrument scheduled to begin operation in 2001 on the CFHT




Fig. 4.6. Photograph of the Wide Field Camera mosaic imager. The mosaic bottom is the old version, which used four Loral  $2K \times 2K$  CCDs, while the mosaic top is the current version consisting of four EEV CCDs. In both mosaics, the  $2K \times 2K$  CCD off to the side is used in frame transfer mode as an autoguider.



(Boulade et al., 1998). This imager will contain 36 to 40  $2K \times 4.5K$  EEV CCDs, each with 13.5-micron pixels. The device will image a full 1 square degree field of view at 0.185 arcsec/pixel resolution. As with all shuttered large-area imagers, MEGACAM will use a half-disk spinning shutter that rotates in a constant direction to ensure uniform exposure time across the array. Current plans call for CCD readout at an amazingly fast 20 seconds per chip.

Wide-field CCD mosaic imagers allow a tremendous amount of information (and data) to be collected in one exposure. The WFC, for example, produces over 800 MB per one mosaic exposure. CCD mosaic arrays are pioneering new scientific advances, driving astronomical technology and bringing us back to many projects started years ago with photographic plates, but now pushing them to new depths and precision levels. The efficiency of a large-area survey can be estimated by the quantity

$$\epsilon = \Omega D^2 q,$$

where  $\Omega$  is the solid angle of the field of view,  $D$  is the diameter of the telescope, and  $q$  is the total throughput quantum efficiency of the instrument assuming that the seeing disk is resolved. One can see that the time needed for completion of a survey to a given brightness limit depends inversely on  $\epsilon$ .

Using wide-field CCD imagers leads to the inevitable result that new issues of calibration and data reduction must be developed. For example, when the field of view of a large-area CCD (whether a single large CCD with a wide field of view or an array of chips) approaches  $\sim 0.5$  degrees in size, differential refraction of the images across the field of view of the CCD begins to become important. Color terms therefore propagate across the CCD image and must be corrected for to properly determine and obtain correct photometric information contained in the data.

CCD observations that occur through uniform thin clouds or differential measures essentially independent of clouds are often assumed to be valid, as it is believed that clouds are grey absorbers and that any cloud cover that is present will cover the entire CCD frame. Thus flux corrections (from say previous photometric images of the same field) can be applied to the nonphotometric CCD data, making it usable. Large-field CCD imaging cannot make such claims. A one or more square degree field of view has a high potential of not being uniformly covered by clouds, leading to unknown flux variations across the CCD image.

Observations of large spatial areas using CCD mosaics also necessitate greater effort and expense in producing larger filters, larger dewar

windows, and larger correction optics. Variations of the quality and color dependence of these optical components over the entire field of view are generally noticeable. These optical aberrations will cause point spread function (PSF) changes over the large areas imaged with wide-field CCDs. The use of large-format CCDs or CCD mosaics on Schmidt telescopes is increasing and such an imager provides a good example of the type of PSF changes that occur across the field of view (see Figure 4.7). Coma and chromatic aberrations are easily seen upon detailed inspection of the PSFs, especially near the corners or for very red or blue objects whose peak flux lies outside of the color range for which the optics were designed. Thus, for wide-field applications, such as that represented in Figure 4.7, the typical assumption that all the PSFs will be identical at all locations within the field of view must be abandoned.

A more subtle effect to deal with in wide-field imaging is that of the changing image scale between images taken of astronomical objects and those obtained for calibration purposes. For example, a dome flat field image taken for calibration purposes will not have exactly the same image scale per pixel over the entire CCD image as an object frame taken with the same CCD camera, but of an astronomical scene. Since the image scale (arcsec per pixel) changes over large (astronomical) fields of view, flat field and other calibration corrections become complex as the different types of images have different or changing pixel scales. Therefore, during the reduction process, some type of image interpolation or movement of flux from one pixel scale to another must be performed. This is never a desirable procedure, but for wide-field CCD imaging, flat field and other calibration procedures make it necessary at some level. One way around such an issue is to use sky flats, which will have the same pixel scale if obtained at the same zenith distance as your object frame, but they will also have all of the disadvantages of sky flats discussed above in Section 4.2.

As with previous new advances in CCD imaging, wide-field imaging has issues that must be ironed out. However, this exciting new field of research is still in its infancy and those of you reading this book who are involved in such work are the ones who must help determine the proper data collection and reduction procedures to use.

#### 4.6.3 CCD Drift Scanning and Time-Delay Integration

The standard method of CCD imaging is to point the telescope at a particular place in the sky, track the telescope at the sidereal rate, and integrate with the detector for a specified amount of time. Once the

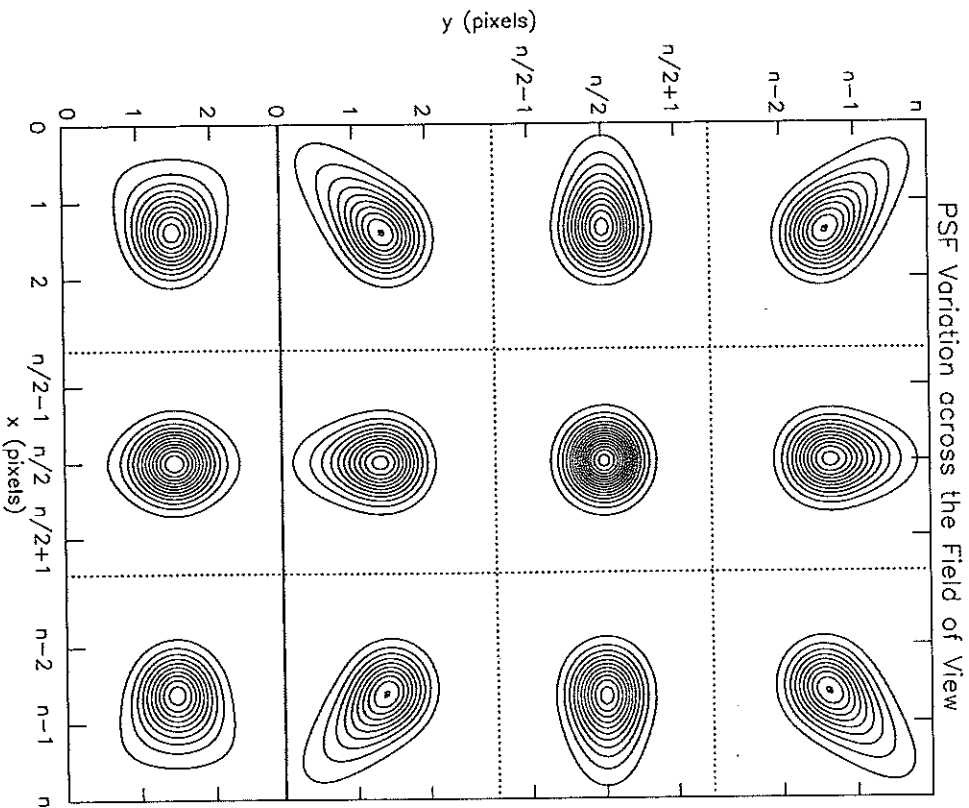


Fig. 4.7. PSF variations for a star imaged at nine locations within the field of view of a large mosaic CCD camera placed at the focal plane of a Schmidt telescope. Only the center of the field has a circular, symmetric PSF while the other positions show extended tails due to optical aberrations and chromatic effects. The three PSFs at the bottom of the figure are column sums of the PSFs vertically above them. From Howell et al. (1996).

desired integration time is obtained, the shutter is closed and the CCD is read out. For telescopes incapable of tracking on the sky or to obtain large areal sky coverage without the need for complex CCD mosaics, the techniques of CCD drift scanning and time-delay integration were developed (McGraw, Angel, & Sargent, 1980; Wright & Mackay, 1981).

Drift scanning consists of reading the exposed CCD at a slow rate while simultaneously mechanically moving the CCD itself to avoid image smear. The readout rate and mechanical movement are chosen to provide the desired exposure time. Each imaged object is thus sampled by every pixel in the column thereby being detected with the mean efficiency of all pixels in the column. Nonuniformities between the pixels in a given column are thus eliminated as each final pixel is, in essence, a sum of many short integrations at each pixel within the column. Cross column efficiency differences are still present but the final image can now be corrected with a one-dimensional flat field. Drift scanning also has the additional advantage of providing an ideal color match to background noise contributions, unavailable with dome flats. Very good flat fielding of a traditional image might reach 0.5 or so percent, while a good drift scanned CCD image can be flattened to near 0.1 percent or better (Tyson & Seitzer, 1988). Drift scanning has even been accomplished with IR arrays (Gorjian, Wright, & Mclean, 1997).

Time-delay integration or TDI is a variant on the drift scanning technique. In TDI, the CCD does not move at all but is read out at exactly the sidereal rate. This type of CCD imaging is necessary if astronomical telescopes such as transit instruments (McGraw, Angel, & Sargent, 1980) or liquid mirror telescopes (Gibson, 1991) are to be used. The same flat fielding advantages apply here as in drift scanning but the integration time per object is limited by the size of the CCD (i.e., the time it takes an object to cross the CCD field of view). For a  $2,048 \times 2,048$  CCD with 0.7 arcsec pixels, the integration time would be only 96 seconds at the celestial equator. Rescanning the same area could be performed and co-added to previous scans as a method of increasing the exposure time, but time sampling suffers.

TDI is mechanically simple, as nothing moves but the electrons in the CCD. This charge movement has been termed electro-optical tracking. Large sky regions can be surveyed, albeit to shallow magnitude limits, very quickly using TDI. Overhead time costs for TDI only consist of the "ramp up" time, that is, the time needed for the first objects to cross the entire field, and the scan time. Using our same  $2048 \times 2048$  CCD as in the example above, we find that a 23 arcsec by 3 degree long strip of the sky at the celestial equator can be scanned in about 2-3 minutes compared with the nearly 25 minutes required if pointed observations of equivalent integration are used.

Although drift scanning and TDI are seemingly great solutions to flat fielding issues and offer the collection of large datasets, drift scanning

requires the CCD to move during the integration with very precise and repeatable steps. This is quite a mechanical challenge and will increase the cost of such an instrument over that of a simple CCD imager. In addition, both techniques suffer two potential drawbacks (Gibson & Hickson, 1992). Images obtained by drift scanning and TDI techniques have elongated PSFs in the east-west direction. This is due to the fact that the rows of the CCD are shifted discretely while the actual image movement is continuous. We note here that objects separated by even small declination differences (i.e., one CCD field of view) do not have the same rate of motion. The resulting images are elongated east-west and are a convolution of the seeing with the CCD pixel sampling.

TDI imagery contains an additional distortion in the north-south direction owing to the curvature of an object's path across the face of the CCD (if imaging away from the celestial equator). This type of distortion is usually avoided in drift scan applications as the telescope and CCD tracking are designed to eliminate this image smearing. This sort of mechanical correction can not be applied to TDI imaging.

These image deformations have been studied in detail (Gibson & Hickson, 1992) and are seen to increase in magnitude for larger format CCDs or declinations farther from the celestial equator. For example, at a declination of  $\pm 30$  degrees, a 1 arcsec per pixel CCD will show an image smear of about 6 pixels. One solution to this large image smear is to continuously reorient the CCD through rotations and translations, such that imaging scans are conducted along great circles on the sky rather than a polar circle or at constant declination. Such a mechanically complex device has been built and used for drift scanning on the 1-m Las Campanas telescope (Zaritsky, Shectman, & Bredthauer, 1996). Another solution is the development of a multilens optical corrector that compensates for the image distortions by tilting and decentering the component lenses (Hickson & Richardson, 1998).

A few telescopes have made good use of the technique of drift scanning or TDI, providing very good astronomical results. Probably the first such project was the Spacewatch telescope (Gehrels et al., 1986) setup to discover and provide astrometry for small bodies within the solar system. Other notable examples are the 2-m transit telescope previously operated on Kitt Peak (McGraw, Angel, & Sargent, 1980) and a 2.7-m liquid mirror telescope currently running at the University of British Columbia (Hickson et al., 1994). This latter telescope contains a rotating mercury mirror and images a 21 arcminute strip of the zenith with an effective integration time of 130 seconds. Using TDI, a typical

integration with this liquid mirror telescope reaches near 21st magnitude in R and continuous readout of the CCD produces about 2 GB of data per night.

Two present-day examples of telescopes employing drift scanning and TDI techniques are the QUEST telescope (Sabby, Coppi, & Oemler, 1998) and the Sloan digital sky survey (Gunn et al., 1998). The QUasar Equatorial Survey Team (QUEST) telescope is a 1-m Schmidt telescope that will provide UB<sub>V</sub> photometry of nearly 4,000 square degrees of the sky to a limiting magnitude of near 19. The focal plane will contain 16 2,048 × 2,048 Loral CCDs arranged in a 4 × 4 array. The telescope is parked and the CCDs are positioned such that the clocking (column) direction is east-west and the readout occurs at the apparent sidereal rate. Each object imaged passes across four CCDs covered, in turn, with a broad-band V, U, B, and V filter. The effective integration time (i.e., crossing time) is 140 seconds, providing nearly simultaneous photometry in U, B, and V.

As we have seen above, a problem with drift scanning is that the paths of objects that drift across the imager are not straight and they can cross the wide field of view with different drift rates. We have discussed a few solutions to these issues, and in the QUEST project (Sabby, Coppi, & Oemler, 1998) we find another. The CCDs are fixed, in groups of four, to long pads lying in the north-south direction. These pads can pivot independently such that they align perpendicular to the direction of the stellar paths. The CCDs are also able to be clocked at different rates, with each being read out at the apparent sidereal rate appropriate for its declination.

The Sloan digital sky survey (SDSS) is a large-format mosaic CCD camera consisting of a photometric array of 30 2,048 × 2,048 STe CCDs and an astrometric array of 24 400 × 2,048 CCDs (see Figure 4.8). The photometric CCDs are arranged in six columns of five CCDs each, providing essentially simultaneous five-color photometry of each image object. The astrometric CCDs are mounted in the focal plane above and below the main array and will be used to provide precise positional information needed for the follow-up multifiber spectroscopy. The SDSS will use a 2.5-m telescope located in New Mexico to image one quarter of the entire sky down to a limiting magnitude of near 23.

TDI scans along great circles will be used by the SDSS to image a region of the sky 2.5 degrees wide. Using six intermediate band filters, covering 3,550 Å to 9,130 Å, scanning at the sidereal rate provides an effective integration per color of 54 seconds with a time delay of 72 seconds




Fig. 4.8. The optical layout within the dewar of the Sloan digital sky survey CCD imager. The right side of the figure labels the CCDs as to their function; 1–15 are photometric CCDs, 16–21 are astrometric CCDs, and 22 (top and bottom) are focus CCDs. The left side gives the dimensions of the array. The labels  $r'$ – $g'$  denote the five separate intermediate band filters, each a single piece of glass covering all six horizontal CCDs. The scan direction is upward causing objects to traverse the array from top to bottom. From Gunn et al. (1998).

between colors caused by CCD crossing time and chip spacing. Complete details of the SDSS, too lengthy for presentation here, can be found in Gunn et al. (1998).

The SDSS has just begun operation and is probably the largest, most complex camera of its kind. Other projects of a similar nature are discussed in Boulaide et al. (1998), Gunn et al. (1998), and Miyazaki et al. (1998).

## Photometry and Astrometry

One of the basic astronomical pursuits throughout history has been to determine the amount and temporal nature of the flux emitted by an object as a function of wavelength. This process, termed photometry, forms one of the fundamental branches of astronomy. Photometry is important for all types of objects from planets to stars to galaxies, each with their own intricacies, procedures, and problems. At times, we may be interested in only a single measurement of the flux of some object, while at other times we could want to obtain temporal measurements on time scales from seconds or less to years or longer.

We start this chapter with a brief discussion of the basic methods of performing photometry when using digital data from 2-D arrays. It will be assumed here that the CCD images being operated on have already been reduced and calibrated as described in detail in the previous chapter. We will see that photometric measurements require that we accomplish only a few steps to provide output flux values. Additional steps are then required to produce light curves or absolute fluxes.

Some photometric output products, such as differential photometry, require fewer additional steps, whereas to obtain the absolute flux for an object, additional CCD frames of photometric standards are needed. These standard star frames are used to correct for the Earth's atmosphere, color terms, and other possible sources of extinction that may be peculiar to a given observing site or a certain time of year (Pecker, 1970).

As an introduction to the level of atmospheric extinction one might expect as a function of observational elevation and wavelength, Table 5.1 lists values of the extinction in magnitudes resulting from the Earth's atmosphere for an observing site at 2,200 m elevation. Note that for observations made at reasonable airmass and redward of 4,000 Å, the