



Scientific Charge-Coupled Devices

James R. Janesick

Scientific Charge-Coupled Devices

James R. Janesick



SPIE PRESS

A Publication of SPIE—The International Society for Optical Engineering
Bellingham, Washington USA

Copyrighted material

**Patent Owner Collect, LLC
Ex. 2059, p. 2**

Library of Congress Cataloging-in-Publication Data

Janesick, James R.

Scientific charge-coupled devices / James R. Janesick.

p. cm. — (SPIE Press monograph ; PM83)

Includes bibliographical references and index.

ISBN 0-8194-3698-4

I. Charge coupled devices. I. Title. II. Series.

TK7871.99.C45 J36 2000

621.36'7-dc21

00-036616

CIP

Published by

SPIE—The International Society for Optical Engineering

P.O. Box 10

Bellingham, Washington 98227-0010

Phone: 360/676-3290

Fax: 360/647-1445

Email: spie@spie.org

WWW: www.spie.org

Copyright © 2001 The Society of Photo-Optical Instrumentation Engineers

All rights reserved. No part of this publication may be reproduced or distributed in any form or by any means without written permission of the publisher.

Printed in the United States of America.

Scientific Charge-Coupled Devices

CHAPTER 1

HISTORY, OPERATION, PERFORMANCE, DESIGN, FABRICATION AND THEORY

1.1 SCIENTIFIC CCD HISTORY

The charge-coupled device (CCD) was invented October 19, 1969, by Willard S. Boyle and George E. Smith at Bell Telephone Laboratories. Figure 1.1(a) shows their original laboratory notebook entry describing the idea. Jack Morton, vice-president of Bell Labs Electronics Technology, was an important motivator behind the CCD. Jack was a strong proponent of the magnetic bubble memory and wondered if an analogous device could be made with a semiconductor. He encouraged Boyle and Smith to look into a potential connection. At the same time, the Picturephone was being developed using diode arrays in silicon.¹ Charge storage on individual diodes had been a problem that prompted some new research. The CCD was born with these two existing technologies in mind. As the story goes, Boyle and Smith brainstormed on the blackboard for approximately a half hour, producing the CCD structure. The first notebook drawing of the device is shown in Fig. 1.1(b), which presents a basic three-phase configuration. The original timing diagram presented also accurately describes today's CCD and is used throughout this book.

A few weeks later, a three-phase device was designed, fabricated and tested: a simple row of nine 100- μm metal plates separated by 3- μm spacings. The first gate electrode was used to inject charge into the second plate. The ninth plate was used to detect charge. Plates 2–8 were clocked to demonstrate the transfer process. The experiment was a success.

The first public announcement, which lasted only five minutes, was made at the New York IEEE convention held March 1970. The ideas were subsequently reported by Boyle and Smith in the April 1970 issue of the *Bell System Technical Journal*.² This paper was important because it introduced potential problems behind the invention. For example, limitations in storage time caused by dark current generation were discussed. Also, they found three important factors that limited charge transfer efficiency: surface interface state trapping, thermal diffusion and fringing fields. The bubble memory analogy was introduced by shifting digital information in the form of charge and implemented with a serial shift register with p-n inputs and outputs. Also very important was the idea that the CCD could be used for imaging using linear and array areas.

The first experimental results were reported in a paper by Amelio, Tompsett, and Smith.³ The first technical paper was presented by Smith at the Device Research Conference in Seattle, June 1970.⁴ This paper presented an 8-bit p-channel CCD with integrated input and output diodes used to inject and read charge. The

1998, 10 AM PDT. The MGS spacecraft was approximately 1,080 km (671 miles) from Phobos at closest approach. This image, about 8.2 km (5.1 miles) wide by 12 km (7.5 miles) tall, shows the full field-of-view. The image as shown here has a scale of 12 m (40 ft) per pixel.

The Mars *Pathfinder* uses the Cassini/Huygens Probe Descent Imaging Spectra Radiometer (DISR) Lockheed CCD mentioned above. Figure 1.6(h) shows a high-resolution image of Mars looking in the direction of Twin Peaks.

Future or new missions include *Deep Space I*, which will visit two asteroids (Braille and Borrelly) and a comet, using a Reticon $1024 \times 1024 \times 9\text{-}\mu\text{m}$ -pixel three-phase CCD. Another exciting mission recently launched by the Shuttle in July 1999 is the Advanced X-ray Astrophysics Facility (AXAF). Now called Chandra, the large x-ray telescope includes the Advanced CCD Imaging Spectrometer (ACIS) instrument. The camera is generating high-resolution x-ray images of active galaxies, supernovas, quasars, etc.; and it simultaneously measures the energy of incident x rays. Several high-performance $1024 \times 1026 \times 24\text{-}\mu\text{m}$ -pixel CCDs fabricated by Lincoln Laboratory are used in Chandra's focal planes. The upcoming Advanced Camera System for Hubble uses two mosaiced SITe $4049 \times 4096 \times 15\text{-}\mu\text{m}$ backside-illuminated three-phase CCDs. The sensitivity of the new camera system is approximately three times better than WF/PC II in terms of QE improvement.^{15,16}

1.2 OPERATION AND PERFORMANCE

1.2.1 OPERATION

As previously mentioned, CCDs were initially conceived to operate as a memory device, specifically as an electronic analog of the magnetic bubble device. In order to function as memory, there must be a physical quantity that represents a bit of information, a means of recognizing the presence or absence of the bit (reading) and a means of creating and destroying the information (writing and erasing). In the CCD, a bit of information is represented by a packet of charges (electrons are assumed in this book, although holes would also work). These charges are stored in the depletion region of a metal-oxide-semiconductor (MOS) capacitor, an important CCD structure that is described below. Charges are moved about in the CCD circuit by placing the MOS capacitors very close to one another and manipulating the voltages on the gates of the capacitors so as to allow the charge to spill from one capacitor to the next. Thus the name "charge-coupled" device. A charge detection amplifier detects the presence of the charge packet, providing an output voltage that can be processed. Charge packets can be created by injecting charge that is from an input diode next to a CCD gate or introduced optically. Like the magnetic bubble device, the CCD is a serial device where charge packets are read one at a time.




Figure 1.7(a) Bucket analogy used to describe CCD operation.

An elegant analogy devised by Jerome Kristian and Morley Blouke, which uses the concept of a “bucket brigade” to describe CCD operation for imaging applications, is presented in Fig. 1.7(a). Their simple model is explained as follows. Determining the brightness distribution in a CCD image can be likened to measuring the rainfall at different points in a field with an array of buckets. Once the rain has stopped, the buckets in each row are moved down vertically across the field on conveyor belts. As the buckets in each column reach the end of the conveyor, they are emptied into another bucket system on a horizontal belt that carries it to a metering station where its contents are measured.

There are many ways to arrange MOS capacitors to form a CCD imager array. Conceptually, the simplest CCD is a three-phase device, the arrangement that Boyle and Smith used for their first CCD. In the three-phase device, the gates are arranged in parallel with every third gate connected to the same clock driver. The basic cell in the CCD, which corresponds to one pixel, consists of a triplet of these gates, each separately connected to phase-1, -2 and -3 clocks making up a pixel register. Figure 1.7(b) shows the layout of a three-phase CCD area array. The image-forming section of the CCD is covered with closely spaced vertical registers, also called parallel registers by some CCD manufacturers. During integration of charge, two vertical phases are biased high (e.g., phase-1 and -2), which results in a high field region that has a higher electrostatic potential relative to the lower biased neighboring gate (phase-3). The dark phases shown represent this bias condition, and it is under these gates where signal electrons collect in a pixel when the CCD is exposed to light. These phases are called collecting phases. The phase




Figure 1.7(b) Primary components that make up a three-phase CCD.

that is biased at lower potential is called the barrier phase since it separates charge being collected in the pixels on either side of this phase.

Vertical registers are separated by potential barriers called channel stops which prevent the spread of the signal charge from one column into another. Similar to our bucket brigade analogy in Fig. 1.7(a), a CCD image is read out by a succession of vertical shifts through the vertical registers. For each vertical shift a line of pixels is transferred into a horizontal register (also called a serial register) which is oriented at right angles to the vertical registers. Before the next line is shifted the charge in the horizontal register is transferred to the output amplifier. The amplifier converts the charge contained in each pixel to a voltage. The device is serially read out line-by-line, pixel-by-pixel, representing the scene of photons incident on the device.

Figure 1.8 schematically shows how charge is transferred in a three-phase CCD register. At time t_1 , the potential for phase-1 is held high, forming collecting wells under these gates. Charge generated in phase-1 will collect under phase-1. Electrons generated in the barrier phases (i.e., phases-2 and -3) will rapidly diffuse into phase-1 because its potential is greater. Therefore, the pixel as shown here is totally photosensitive (i.e., 100% fill factor). At t_2 charge transfer commences which involves the creation of potential wells and barriers by applying the appropriate voltages to the gates at appropriately sequenced times. During t_2 phase-2 clock goes high, forming the same well under phase-2 that exists under phase-1. The signal charge collected now divides between the two wells. Note that a potential barrier between pixels exists under phase-3 since it remains at a low bias potential. At time t_3 phase-1 is returned to ground, which forces charge to transfer to phase-2.




Figure 1.8 Three-phase potential well timing diagram showing signal charge being transferred from left to right.

At t_4 , in the same way, charge is transferred from phase-2 to phase-3. This process continues until charge has been moved one entire pixel in one 3-phase clock cycle from left to right in the illustration. Note also that by interchanging any two of the clocks causes the charge in the 3-phase device to move in the opposite direction, an important characteristic that is beneficial in many CCD applications. For example, if charge is normally clocked in a 1, 2, 3 fashion, reversing phases-1 and -2 will cause the charge to move backwards (reversing phases-2 and -3 or phases-1 and -3 will also result in a backwards motion).

1.2.2 PERFORMANCE FUNCTIONS

Technically speaking, the CCD must perform four primary tasks in generating an image. These performance functions are called (1) charge generation, (2) charge collection, (3) charge transfer and (4) charge measurement. Each performance characteristic and its limit will be fully examined in Chapters 3, 4, 5 and 6, respectively. An introduction to each function is given below.

1.2.2.1 Charge Generation

The first operation—charge generation—is the ability of a CCD to intercept an incoming photon and generate an electric charge. Charge generation efficiency (CGE) is described by a function called quantum efficiency (QE), which is the fraction of incident photons that produces a useful charge in the silicon chip. These

photons are known as interacting photons. An ideal CCD would have 100% QE at all wavelengths, but nature is rarely this kind. QE is highly dependent on wavelength, falling off steeply at the near-infrared (NIR) and blue wavelengths. Charge generation takes place in the silicon body of the CCD. When a photon interacts with this material, it creates one (or for higher-energy photons, several) free electron(s) by a physical process known as the photoelectric effect.

Interaction only occurs if the photon has enough energy to do so. This is because silicon exhibits an energy gap of approximately 1.14 eV, which is bounded by the valence and the conduction energy bands. Photon energy $E(\text{eV})$ can be converted to a wavelength using the equation

$$\lambda = \frac{12,390}{E(\text{eV})}, \quad (1.1)$$

where λ is the wavelength (\AA).

Example 1.1

What is the wavelength of a 1.14-eV photon?

Solution:

From Eq. (1.1),

$$\lambda = 12,390/1.14,$$

$$\lambda = 10,868 \text{ \AA}.$$

Most photons with energies < 1.14 eV pass through the chip; i.e., silicon is transparent in the far infrared. Some of these photons do in fact interact but do not generate signal carriers, a process called free-carrier absorption. Interacting photons with sufficient energy will excite an electron into the conduction band, creating an electron-hole (e-h) pair. The e-h pair created is free to move and diffuse in the silicon lattice structure. For high-quality silicon, the lifetime of the electron is several milliseconds before recombining.

A photon with energy of 1.1 to 3.1 eV (11,263 \AA –4,000 \AA) will generate a single e-h pair. This spectral range covers the near infrared (NIR) and visible spectrum (4,000–7,000 \AA). Energies > 3.1 eV will produce multiple e-h pairs when the energetic conduction band electron collides with other valence electrons.¹⁷ The average number of electrons generated for a photon with energy $E(\text{eV}) > 10$ eV is

$$\eta_i = \frac{E(\text{eV})}{E_{e-h}}, \quad (1.2)$$

where η_i is the ideal quantum yield (electrons/interacting photon) and E_{e-h} is the energy required to generate an electron-hole pair, which for silicon is 3.65 eV/e⁻ at room temperature.

Example 1.2

Find the average number of e-h pairs generated for 10-eV (Lyman-alpha) and 5.9-keV x-ray photons.

Solution:

From Eq. (1.2),

$$\eta_i = (10)/3.65,$$

$$\eta_i = 2.74 \text{ e}^- \text{ at } 10 \text{ eV},$$

$$\eta_i = 1620 \text{ e}^- \text{ at } 5.9 \text{ keV}.$$

The response of the CCD can potentially cover an enormous wavelength range, 1–10,000 Å (1.1 eV to 10 keV). This spectral range includes the near infrared, visible, ultraviolet (UV), extreme ultraviolet (EUV) and soft x-ray regions. CCDs fabricated with other optical materials can extend this range even further.¹⁸ For example, small germanium CCDs were fabricated at Loral. Germanium exhibits a bandgap of 0.66 eV, which extends the IR response out to about 1.6 μm. In addition, the density of germanium is greater than silicon, pushing the x-ray response out to approximately 20 keV. The process employed in fabricating a germanium CCD is very similar to silicon processing, with the exception of the gate insulator. Unfortunately, germanium oxide is not a good insulator (e.g., it is soluble in water). Therefore, deposited SiO₂ was employed as the gate insulator for the germanium CCD. There is a severe lattice mismatch between germanium and SiO₂, which produces a high density of interface states. This leads to very high dark current and large flat-band shift problems. Nevertheless, images were taken when the devices were cooled to very low operating temperatures (< 200 K).

Charge generation in the visible region (4000–7000 Å) typically peaks near 6500 Å. At wavelengths shorter than 6500 Å, the photons begin to be absorbed by the CCD's gate structure that overlies the silicon. The chip's QE drops to only a few percent at 4000 Å. The absorption problem becomes particularly pronounced in the ultraviolet at wavelengths below 3000 Å. For example, at 2500 Å the penetration depth is only 30 Å in silicon, which is a few atomic layers. Most CCDs are inert lumps of metal in this range and regain sensitivity only at wavelengths of about 10 Å, making them useful detectors at x-ray wavelengths.

CCD manufacturers have adopted a variety of techniques to enhance blue and ultraviolet performance. One technique used for the Hubble WF/PC II CCDs was to deposit a phosphor coating directly on the gates. A photon having a wavelength

shorter than 4600 \AA is absorbed by the coating, which re-emits another photon at 5200 \AA where the CCD is more sensitive. This down-converting wavelength technique avoids most of the gate absorption problem. However, only 50% of the 5200 \AA photons are re-emitted in a direction that carries them into the silicon where they are detected, while the other 50% escape the CCD. The short-wavelength quantum efficiency of a phosphor-coated CCD is therefore approximately half that of the same CCD at 5200 \AA .

Phosphor coatings are a cheap way to improve ultraviolet response. A better but much more expensive technique is known as backside illumination, whereby the image is focused on the backside of the silicon in order to bypass the gate structure. This illumination scheme results in the highest QE possible for the CCD. However, to be effective, the thick substrate layer on which the CCD is built must be completely thinned away, because only the electrons generated near the gate structure can be collected efficiently. The membrane thickness for most backside-illuminated CCDs is extremely thin, approximately $12 \mu\text{m}$ thick. Thinning is difficult, and many otherwise good chips are often destroyed in the process.

Reflection is another QE loss problem. At 4000 \AA , silicon reflects about half the light falling on it, which is why most CCDs shine like mirrors. Fortunately, reflection losses can be nearly eliminated by using the same antireflection coating technique found on good camera lenses. A properly coated backside-illuminated chip looks dark, not shiny, and can achieve a peak QE over 90% in the visible range.

1.2.2.2 Charge Collection

The second step, charge collection, is the ability of the CCD to accurately reproduce an image from the electrons generated. Three parameters are included to describe this process: (1) the area or number of pixels contained on the chip, (2) the number of signal electrons that a pixel can hold (i.e., charge capacity of a pixel), and (3) the ability of the "target pixel" to efficiently collect electrons when they are generated. The last parameter is most critical because it defines the spatial resolution performance for the CCD. Spatial resolution is primarily associated with pixel size and charge diffusion characteristics, as explained in Chapter 4.

For most scientific applications, the number of pixels making up an array should be as large as possible. While early CCDs were very small, the number of pixels has grown to overwhelming proportions. For example, Fig. 1.9 shows a 4-in. wafer that holds a $2000 \times 2000 \times 27\text{-}\mu\text{m}$ -pixel CCD designed by Morley Blouke and his group formerly at Tektronix and currently at SITE. This was the first CCD chip to occupy the entire silicon wafer.

Today, 4096×4096 CCDs are fabricated at several CCD manufacturers. A 4096×4096 CCD can compete in number of resolution elements with 35-mm format photographic film. Such CCDs produce digital images with an enormous amount of information. For example, two football fields set side by side, including the side and goal zones, can resolve objects to about 1 inch per pixel. Looking




Figure 1.9 A wafer scale size $2\text{ k} \times 2\text{ k} \times 27\text{-}\mu\text{m}$ -pixel CCD imager built by Tektronix in the early eighties.

skyward, the 4096×4096 sensor can cover 68 arc min of the sky with 1 arc sec resolution (the moon extends about 32 arc min).

Figure 1.10(a) shows a $4000 \times 4000 \times 15\text{-}\mu\text{m}$ -pixel CCD built by Richard Bredthauer and his team at Ford Aerospace. (Ford was subsequently bought out by Loral and later was procured by Lockheed Martin.) Figure 1.10(b) is a dollar bill image taken with the CCD. Only a 600×1000 -pixel portion is presented in the region of the “eye” found on the backside of the bill. The number of pixels displayed is 3.56% of the total 16,777,216 pixels. Figure 1.10(c) is a magnified view showing 300×400 pixels representing 0.715% of the imaging area. Figure 1.10(d) shows a 112×150 image which is less than 0.1% of the pixels. At this magnification individual pixels can be seen.

CCDs have been fabricated even larger than 4000×4000 units. For example, in 1992, JPL initiated a project to fabricate a $9000 \times 7000 \times 12\text{-}\mu\text{m}$ -pixel CCD for astronomical applications. The CCD, built by Philips, occupied a 6-in. wafer and contained over 63,000,000 pixels. A $9000 \times 9000 \times 8.75\text{-}\mu\text{m}$ -pixel CCD was fabricated by Bredthauer in 1996 on a 5-in. wafer containing a whopping 81,000,000 pixels. And a $10,000 \times 10,000 \times 10\text{-}\mu\text{m}$ -pixel is on the drawing board at the time of this writing.




Figure 1.10(a) A $4000 \times 4000 \times 15\text{-}\mu\text{m}$ -pixel CCD built in the early nineties at Loral.




Figure 1.10(b) A 600×1000 dollar bill image taken with a $4000 \times 4000 \times 15\text{-}\mu\text{m}$ CCD.

Are there limits to making even larger detectors? From the manufacturer's standpoint, pixel count is ultimately limited by the production yield. Not every chip fabricated works, and the bigger the chip, the lower the yield and the higher the cost. Large CCDs fail primarily because of electrical shorts in the gates,




Figure 1.10(c) A 300×400 magnified view of Fig. 1.10(b).




Figure 1.10(d) A 112×150 magnified view of Fig. 1.10(c).

a subject discussed below. So, the answer to this question is “no,” assuming resources are available to build such monster devices. For example, a small-pixel $20,000 \times 20,000$ CCD is hypothetically possible today. However, interest and vast amounts of money would be required to build a device this size.

The cost of ultralarge CCDs is considerable because usually only one or two devices can be built per silicon wafer. A lot run typically holds 24 wafers and

possibly only a few good wafers will pass the shorts test. The foundry fabrication cost for a 6-in. silicon run is approximately \$100,000, depending on the manufacturer and the number of options specified for the CCD (i.e., the number of reticles or masks required for lithography). There are also costs associated with design, packaging and testing, which are expensive efforts. A custom CCD might run near \$500,000 before research and development costs are concluded.

Individual off-the-shelf CCDs have price tags of more than \$50,000 for ultralarge backside-illuminated arrays. For example, a class-A fully characterized 2048×2048 thinned chip can cost more than \$70,000. Engineering grades for the same device might go for \$10,000. This works out to be 0.60 and 0.23 cents per pixel, respectively.

Nevertheless, production yields and costs have improved greatly over the years. About 150 production lots were required to make 800×800 arrays for the Hubble WF/PC I and *Galileo* missions. In fact, the *Galileo* sensor was a one-of-a-kind CCD. Of the hundreds of chips fabricated, the sensor that now orbits Jupiter was the only one fabricated that met specifications. The sensor was found hiding in TI's lab, and was in fact an experimental CCD not intended for flight use. In contrast to all of this, only two lots were required to deliver new flight CCDs to Hubble's WF/PC II and satisfy the Cassini project.

Very large CCDs also present many difficulties for the user, primarily in data storage. For example, a $10,000 \times 10,000$ sensor would produce 200 million bytes of information per image, assuming 16-bit encoding. The sensor format would match that of 4096×512 -pixel commercial format CCDs. One night's images at an observatory from such a detector could fill a typical computer's hard drive. More important, however, is that processing time rises exponentially with pixel count. So, to use large CCDs effectively, very fast and expensive computer systems are needed. This requirement has limited the demand for large CCDs.

Also, from the camera designer's standpoint, the problem with large arrays is the amount of time required to read millions of pixels. For example, the WF/PC cameras require several minutes to read out its four 800-pixel-square arrays at 50,000 pixels/sec. At the same rate, a $10,000 \times 10,000$ array would require 33 minutes! Even if this read time were acceptable, long readouts degrade data because of cosmic-ray artifacts and thermally generated dark current that accumulate during the read process. To shorten readout time, large CCDs usually employ several amplifiers working in parallel. For example, the Philips CCD mentioned above uses four amplifiers. Other astronomical CCDs have employed 64 amplifiers for adaptive optics applications. Nevertheless, a price is to be paid for multiple amplifiers. Parallel readout increases the complexity of a camera's electronics and makes image processing and analysis more difficult because each amplifier exhibits its own gain, offset, linearity and noise levels.

A technique used to beat the yield problem and obtain ultralarge CCD focal planes has been to mosaic them.¹⁹⁻²¹ For example, four 4096×4096 CCDs have been designed, fabricated, diced, butted and mosaicked on two edges to obtain a 8192×8192 format. Even larger mosaics have been configured to butt three free

edges. Packaging costs are high using this approach, because very tight tolerances are required in butting the chips to keep the seam between devices to a minimum. A few pixels separation can be achieved.

Well capacity is a measure of how many electrons a pixel can hold, usually 50,000 to a 1,000,000. Again, bigger is better, since the greater the full-well capacity, the greater a CCD's dynamic range and signal-to-noise. Dynamic range is the ratio between the brightest and faintest objects that can be simultaneously discriminated. Large pixels can hold more charge but require more silicon, which in turn degrades the production and increases the price. The smallest pixel that can be manufactured is only a couple of microns in extent, set by lithography design and process rules. Also, charge capacity is an important consideration for very small pixels because full well decreases by pixel area. The advantage of small pixels is higher yield with fewer shorted devices. Production yield is roughly inversely proportional to the area used on the wafer by the CCD. Therefore, yield will be approximately four times greater for a 12- μm pixel than a 24- μm pixel, assuming the same pixel count.

Pixel size is limited due to the need to maintain high CTE with reasonable speed, as will be discussed in Chapter 5. A good compromise between CTE and well capacity for today's astronomical CCD has been the three-phase 15- μm pixel, the size used by WF/PC I and WF/PC II. Today's 15- μm pixel can hold 500,000 e^- , yielding a dynamic range of 250,000 assuming a readout noise of 2 e^- . This means that an exposure could simultaneously record two objects with a brightness difference of 13 1/2 magnitudes, the ratio between Vega (one of the brightest stars in the heavens) and the planet Pluto. Additional phases are incorporated to make super pixels. For example, a 64-phase CCD has been fabricated to obtain pixels 640 μm square. This CCD produced an enormous well capacity ($10^9 e^-$) while simultaneously delivering good CTE characteristics at high pixel rates. A total of 128 clock drivers were required to drive the horizontal and vertical registers.

Ideally, electrons are collected by the target pixel without spreading into neighboring pixels. Diffusion into other pixels makes an image look out of focus and lowers the sensor's modulation transfer function (MTF). Confining charge to a single pixel is also an issue for sensitivity, since it is easier to distinguish a star from the background when images are well defined. Charge leakage between pixels, also called "pixel crosstalk," is caused by "field-free" material found deep in the silicon. Electrons generated farther from the gate feel a weaker electric field and thus have a higher probability of diffusing into neighboring pixels. For frontside-illuminated CCDs this problem is most apparent when imaging in the near IR because long-wavelength photons penetrate deeper into the CCD than do blue photons. CCD manufacturers reduce this effect through the proper choice of silicon material. Backside-illuminated devices, which by design have light falling on the silicon surface farthest from the gates, also experience leakage if not thinned properly. Through careful design and processing, engineers have reduced crosstalk in backside-illuminated CCDs to nearly undetectable levels.

1.2.2.3 Charge Transfer

The third operation, charge transfer, is accomplished by manipulating the voltage on a parallel sequence of gates that form a CCD register, or conveyor belt in our bucket brigade analogy. For quantitative scientific measurements such as those encountered in astronomy, it is important to lose as little of the charge as possible during the transfer process. For example, the photometric accuracy of the Hubble WF/PC II 800-pixel-square array is specified to better than 1% (< 0.01 magnitude). This means that more than 99% of the charge collected in a pixel must survive the transfer process and reach the output amplifier. At the corner of the chip farthest from the amplifier, this means surviving 1600 pixel transfers. This places enormous demands on the CTE process—no member of the bucket brigade can afford to spill a single drop!

The first CCDs fabricated at Bell Laboratories exhibited about 99% efficiency per pixel transfer. Although this may sound good, it was in fact terrible, because the loss is compounded as the number of transfers increases. If 1% of the charge is left behind in each transfer, 63% would be lost after only 100 transfers. Fortunately, great improvements during the last 25 years have raised charge transfer efficiencies (CTE) in modern CCDs to better than 99.9999%. Only one in every million electrons is lost during a typical transfer!

Understanding how charge is lost during transfers has been very important in pushing CTE performance to the limit. Charge is transferred due to fringing fields between phases, thermal diffusion, and self-induced drift, subjects discussed in Chapter 5. The relative importance of each effect is dependent primarily on how much charge is being transferred. Most scientific CCDs move the charge slowly from pixel to pixel, so the time interval needed for diffusion and self-induced drift is not a CTE limitation in slow-scan operation—the changing fields induced by the gates have time to sweep electrons cleanly from one pixel to the next. Some scientific applications, such as Star Tracker cameras, operate at rates of 30 frames per second or greater. CCDs employed in adaptive optics and biological applications operate at thousands of frames per second. For very high speed operation, charge transfer efficiency can suffer because charge diffusion and self-induced drift require a finite time to move electrons completely from phase to phase.

While most pixels in a modern CCD transfer charge smoothly, many chips have small “traps” that prevent achieving 100% efficiency. These traps extract a fixed quantity of charge as a signal packet is clocked past the problem site. The trapped charge is then slowly released into trailing pixels as “deferred charge,” so called since it appears late at the output amplifier. Traps can be created during the design and manufacture of chips, and to some degree, are also found in starting silicon material before the CCD is even made. A single trap could have disastrous effects if it occurred in a pixel near the output amplifier. Nevertheless, traps are avoidable, and through careful design and fabrication, a large chip today will contain just a few pixels with traps that siphon only one hundred electrons from a charge packet.

Energetic particles such as protons, electrons, neutrons, heavy ions, and gamma rays damage the silicon crystal lattice and create traps too. The problem is most pronounced for CCDs that fly on space missions. For example, the 800-pixel-square array currently circling Jupiter on the *Galileo* spacecraft has been severely damaged by energetic protons and neutrons. Engineers had anticipated this problem, however, and compensated for it in the camera design. The spectacular images returned by *Galileo* are evidence that they did their job well. For ground-based astronomy, radiation problems prevail but concern us little. Although constant bombardment by cosmic rays and natural radiation causes damage and CTE loss, even the highest radiation levels would take a 100 years or more to cause a noticeable effect. Since the CCD is only 30 years old, no one has complained.

Bulk traps are due to impurities or lattice defects within the silicon. These typically trap a single electron during charge transfer and set limits on the ultimate CTE performance. They can be reduced by building chips with the best silicon material available. Current limits place CTE at 99.99999%.

1.2.2.4 Charge Measurement

The last major operation to occur during CCD imaging is the detection and measurement of the charge collected in each pixel. This is accomplished by dumping the charge onto a small capacitor connected to an output MOSFET amplifier. This point on the CCD is called the “sense node” or “output diode.” The MOSFET amplifier is the only active element that requires power. The CCD array of pixels itself is passive in that it is made up of MOS capacitors. The output amplifier generates a voltage for each pixel proportional to the signal charge transferred.

CCD designers have worked diligently to make the output capacitor extremely small. The tinier it is, the higher the amplifier's gain and the greater the output signal. For example, a 50 fF capacitance produces a gain of 3.2 μV per electron. Thus, if 1000 electrons are transferred to the sense node, 3.2 mV appear at the transistor's output, which is a respectable signal for digital conversion. Sensitivities as high as 25 μV per electron have been achieved. However, such sensitivities limit the dynamic range of the amplifier and linearity performance. A popular sensitivity for the scientific CCD has been about 2–4 μV per electron.

Besides making the sensitivity as high as possible, engineers make sure there is minimal noise generated by the output MOSFET. This is very important because reducing noise by a factor of two doubles the CCD's sensitivity. The effect is the same as if a telescope's aperture were increased by 41%, assuming signal detection is truly read-noise-limited. Building a quieter camera is easier and cheaper than building a bigger telescope. Although there are many unwanted sources of noise involved with CCD imaging, all can be reduced to zero. For example, there is thermal noise associated with the dark current arising from electrons spontaneously generated within the silicon. Dark current can be controlled and eliminated by cooling the chip. The WF/PC II CCDs are cooled to -90°C for this purpose and generate less than 3.6e^- of dark current per hour per pixel.

The only fundamental source of noise that cannot be completely eliminated comes from the output amplifier. It is generated by the random fluctuations in the current that flows through the transistor. The noise produced is similar to the “snow” sometimes seen on weak television pictures. Early CCDs had about $30e^-$ rms of amplifier noise associated with the signal measured for each pixel. Today’s high-performance CCDs have slightly less than $2e^-$ of noise, a spectacular improvement. The improvement did not come easily, however. Literally millions of dollars have been spent optimizing the CCD amplifier. Design efforts striving for even lower noise amplifiers continue today in pushing the noise level below $1e^-$.

Noise levels this low are achieved not just by careful chip design but also by the design of electronics used to process the CCD signal. The raw output noise that accompanies a CCD signal is typically more than $100e^-$. This noise can be reduced almost a hundredfold by additional off-chip amplification and digital filtering.

1.2.3 PERFORMANCE SPECIFICATIONS

The introductory discussions above acquaint the reader with CCD performance characteristics, which have been grouped into the four basic CCD functions (i.e., charge generation, collection, transfer and measurement). CCD performance is further specified by three major categories associated with the architecture of the CCD: (1) vertical register performance, (2) horizontal register performance and

Table 1.1 Vertical register performance.

Performance Parameter	Specification	Transfer Curve or Image	Conditions
Quantum efficiency 3000 Å 4000 Å 5000 Å 7000 Å 9000 Å	> 0.4 > 0.6 > 0.75 > 0.50 > 0.20	QE Transfer	backside illumination
Charge capacity	> 175,000 e^- > 90,000 e^-	Photon Transfer	partially inverted MPP
Global CTE	> 0.99999	X-ray Transfer	1620 e^- (Fe ⁵⁵)
Local CTE	< three 100 e^- traps	Bar target image	1000 e^- signal
Line transfer time	1 μ sec	Full Well Transfer	at vertical full well
Pixel nonuniformity	< 2%	Photon Transfer	
Global dark current	< 0.5 nA/cm ² < 20 pA/cm ²	Dark Current Transfer	partially inverted MPP (300 K)
Local dark current	< 1 nA/pixel	Dark current image	300 K
Dark current nonuniformity	< 10%	Photon Transfer	
MTF	> 0.45	Modulation Transfer	Nyquist, $\lambda = 4000 \text{ Å}$

Table 1.2 Horizontal register performance.

Performance Parameter	Specification	Transfer Curve	Conditions
Global CTE	> 0.99999	X-ray Transfer	1620 e ⁻ (Fe ⁵⁵)
Charge capacity	> 700,000 e ⁻	Photon Transfer	noninverted
Pixel transfer time	< 10 ns	Full Well Transfer	at vertical full well
Output summing well charge capacity	> 10 ⁶ e ⁻	Photon Transfer	

Table 1.3 Output amplifier performance.

Performance Parameter	Specification	Transfer Curve	Conditions
Read noise	< 3 e ⁻ rms < 30 e ⁻ rms	Photon Transfer	sample time = 4.0 μsec sample time = 0.04 μsec
Sensitivity	> 5 μV/e ⁻	Photon Transfer	at 10 M pixels/sec
Video time constant	τ < 5 nsec		C _L = 5 pF
Nonlinearity	< 1%	Linearity Transfer	over dynamic range

(3) output amplifier performance. As an example of this organization, Tables 1.1–1.3 tabulate major performance parameters for a backside-illuminated 12-μm-pixel three-phase CCD. For those readers new to the CCD, it will be important to return to these tables as these parameters are reviewed in future chapters. Also included is the CCD transfer curve responsible in quantifying the parameter (transfer curves are discussed in Chapter 2) and the critical operating condition for the parameter tested (for example, operating temperature is an important condition in measuring dark current).

1.3 ARCHITECTURE, DESIGN, PHOTOLITHOGRAPHY AND FABRICATION

1.3.1 ARCHITECTURE

While the idea of charge coupling is well defined, it has been implemented by CCD manufacturers in many ways. There are three common scientific CCD architectures. The first is a simple “linear” shift imager, as shown in Fig. 1.8. An image for this detector is recorded by slowly scanning the scene vertically past the imager while rapidly reading out the shift register in the horizontal direction. The second format is an area array employed as a “full-frame” imager, the format shown in Fig. 1.7. For this architecture, the scene is imaged onto the device, integrated, and then, with the shutter closed, the entire array of pixels is read out. Finally, to achieve very high frame rates and provide electronic shuttering, several manufacturers use a “frame transfer” layout. In this case, the full-frame imager is divided into two halves and independently clocked. The upper half of the array contains the “imaging registers,” while the lower half contains the “storage registers.” The scene is