Appendix 3

More Than You Ever Really Wanted to Know About Charge-Coupled Devices

James B. Pawley

INTRODUCTION

The electronic structure of crystalline Si is such that electromagnetic waves having the energy of light photons (1.75–3.0 electron volts) can be absorbed to produce one free or "conduction" electron. If an image is focused onto a Si surface, the number of the photoelectrons (PE) produced at each location over the surface is proportional to the local light intensity. Clearly, all that is needed to create an image sensor is a method for rapidly converting the local PE concentration into an electronic signal. After almost 40 years of NASA and DOD funding, the slow-scan, scientific-grade, charge-coupled device (CCD) camera is now an almost perfect solution to this problem.

Success in modern biological light microscopy depends to an ever-increasing extent on the performance of CCD cameras. Because such cameras differ widely in their capabilities and are also items that most biologists buy separately, rather than as part of a system, some knowledge of their operation may be useful to those practicing biologists who have not yet found it necessary to be particularly interested in "electronics." Although the basics of CCD operation are described in many other chapters (particularly, Chapters 4, 10 and 12) this Appendix describes the operating principles of these devices in greater detail and also discusses the ways that they "don't work as planned." It then covers the operation of the electron-multiplier CCD (EM-CCD), a new variant that reduces the read noise almost to zero, although at the cost of reduced effective quantum efficiency (QE_{eff}).[1] The second section, *How to choose a CCD*, is a review of CCD specifications with comments on the relevance of each in fluorescence microscopy.

PART I: HOW CHARGE-COUPLED DEVICES WORK

The first step is to imagine a rectangular area of the Si surface as being divided into rows and columns, or more usually, lines and pixels. Each pixel is between 4 × 4 µm and about 24 × 24 µm in size and the location of any pixel of the surface can be defined in terms of it being x pixels from the left side, on line *y*.

To construct an actual system like this, start with a smooth Si surface; cover it with a thin, transparent, insulating layer of SiO_2; deposit onto the SiO_2, a pattern of horizontal strips, made out of a transparent conductor called amorphous silicon (or poly-silicon), so that the strips cover the entire image sensor area. Although, viewed from the top, these strips partially overlap each other, they are kept electrically separate from their neighbors by additional layers of SiO_2. Every third stripe is connected together to form three sets of interdigitating strips that we will refer to as Phases 1, 2 and 3 (ϕ_1, ϕ_2, ϕ_3, Fig. A3.1). Taken together, all these phases constitute the vertical register (VR) and, after the assembly has been exposed to a pattern of light, they are used to transfer the photo-induced charge pattern downwards, one line at a time. The pixels along each line are separated from each other by vertical strips of positively doped material injected into the Si. These positive "channel blocks" create fields that prevent charge from diffusing sideways without reducing the active area of the sensor.

Any photon that passes through the stripes and the SiO_2, is absorbed in the Si, producing a PE. If a small positive voltage (~15 volts) is applied to the ϕ_1 electrodes, any PE produced nearby will be attracted to a location just below the nearest ϕ_1 strip (Fig. A3.2). As additional PEs are produced, they form a small cloud of PEs referred to as a charge packet. The number of PEs in the packet is proportional to the local light intensity times the exposure period and the problem now is to convey this packet to some location where its size can be measured, and to do this without changing it or losing track of the location from which it was collected. This will be achieved by using the overlying electrodes to drag the charge packet around in an orderly way until it is deposited at the readout node of the charge amplifier.

Charge Coupling

The dragging mechanism operates in the following way: First ϕ_2 is also made positive so that the cloud diffuses to fill the area underneath both ϕ_1 and ϕ_2. Then ϕ_1 is made zero, forcing the packet to concentrate under ϕ_2 alone (Fig A3.2).

So far, these 3 steps have succeeded in moving the charge packets that were originally under each of the ϕ_1 electrodes downwards by one phase or 1/3 of a "line" in the *x-y* raster. If this sequence is now repeated, but between ϕ_2 and ϕ_3 and then again between ϕ_3 and the ϕ_1 belonging to the next triplet of strips, packets will have moved down by the one entire raster line. PEs created within a particular pixel of each horizontal stripe remain confined by the channel stops as they are transferred to the next line below.

A pixel of the image is therefore defined as the area under a triplet of overlying, vertical charge-transfer electrodes and between two neighboring channel blocks. The pixels on scientific CCDs, are usually square, 4 to 30 µm on a side while those on commercial, video CCDs are likely to be wider than they are high, to conform with the reduced horizontal resolution of commercial video standards. Only square pixels can be conveniently displayed in a truly digital manner. Larger pixels have more leakage current (dark-current), but are also able to store more charge per pixel (see Blooming, below).

[1] This loss can be avoided if the system is used in photon-conting mode.

James B. Pawley • University of Wisconsin, Madison, Wisconsin 53706

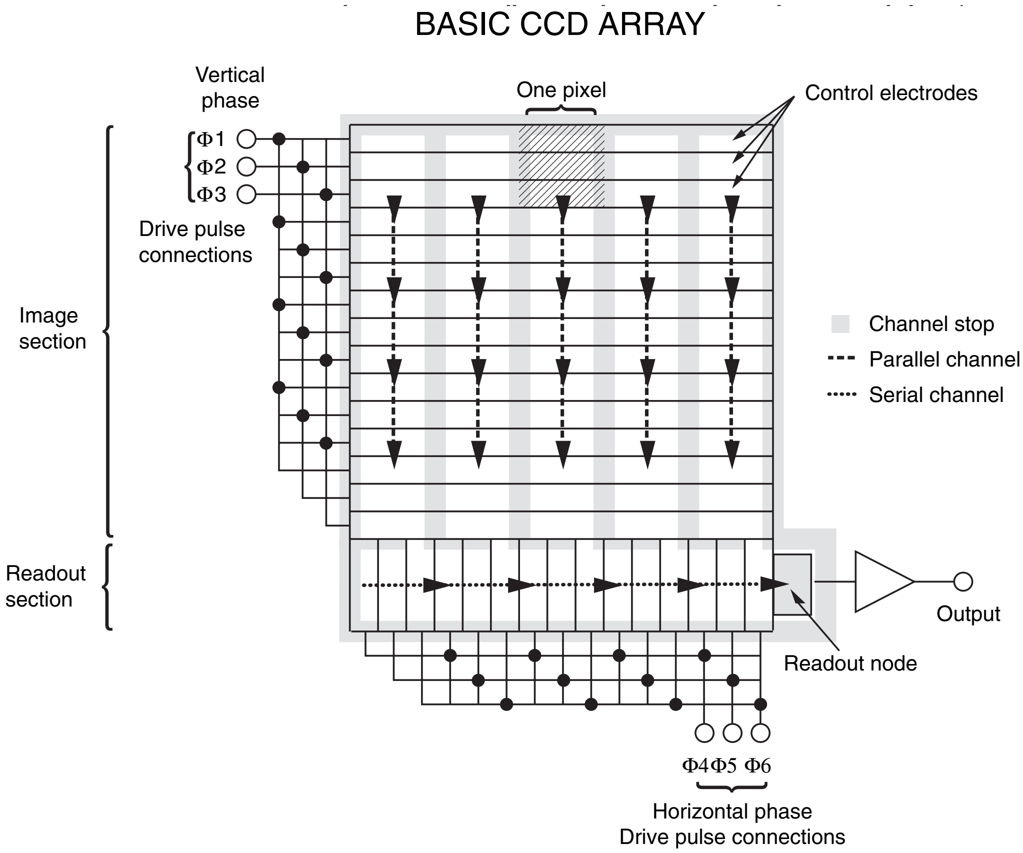

FIGURE A3.1. Layout of CCD array, viewed *en face*.

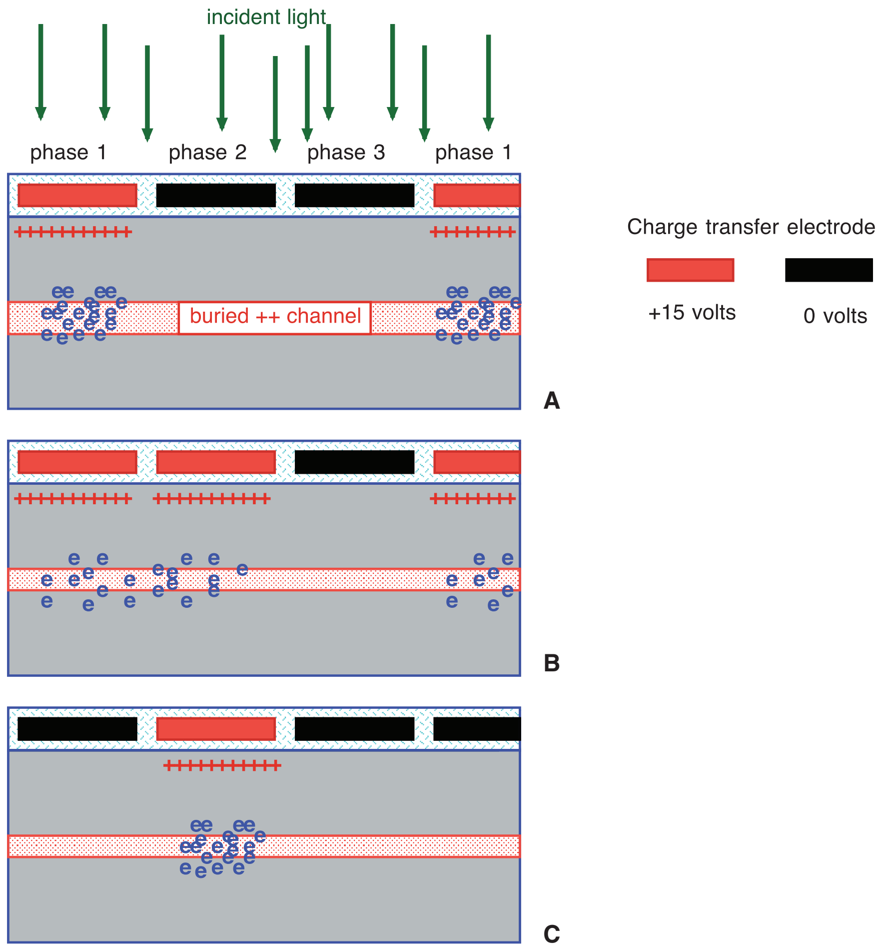

FIGURE A3.2. Charge coupling: Three stages in the process of moving a charge packet initially beneath phase 1 (A), so that it first spreads to be also under phase 2 (B) and finally is confined to entirely under phase 2 (C). These 3 steps must be repeated 3 times before the charge packet has been moved downwards (or in the diagram, to the right) by one line of the CCD array.

At the bottom of the sensor, an entire line of charge packets is simultaneously transferred to the adjacent pixels of the horizontal register (HR, also sometimes called a shift register). Like the VR, the HR is composed of a system of overlying poly-Si electrodes and channel stops. Each column of pixels in the VR is eventually transferred directly into the same specific pixel on the HR. The three phases of the HR (ϕ_4, ϕ_5, ϕ_6) work exactly like those in the VR, except that they must cycle at a much faster rate because the entire HR must be emptied before the next line of packets is transferred down from the bottom line of the VR. In other words, in the time between one complete line-transfer cycle of the VR and the next, the horizontal register must cycle as many times as there are pixels in each line.

At the right-hand end of the HR is a charge amplifier that measures the charge in each packet as it is transferred into it from the last pixel of the HR. The first pixel to be read out is that on the extreme right-hand side of the bottom line. The last pixel will be that on the left side of the top line.[2]

The entire charge-transfer process has the effect of coding position as time. If we digitize the signals from the charge amplifier, and store the resulting numbers in a video memory, we will be able to see a representation of the light intensity pattern striking the sensor on any monitor attached to this video memory. Alternatively, as long as the dimensions of the CCD array match those of some video standard, such as NTSC or PAL, the time sequence of charge-packet readout voltages can be smoothed and, with the addition of synch pulses, turned into an analog video signal. While this latter process is often convenient, it is a poor plan if the analog signal must then be re-digitized. The necessity to digitize twice can reduce the effective horizontal resolution of the CCD sensor by about a factor of 2 and because the process is AC coupled, photometric accuracy is severely compromised.

It is important to understand the relationship between the charge-transfer electrodes and the charge packet. The electrodes do not somehow "connect to" the charge packet, and "conduct" it to the amplifier. Such a process would be subject to resistive losses, charge would be lost and a lot of "wires" would be needed. The charge-coupling process is better thought of in terms of a ball bearing "dragged" over the surface of a loose blanket by moving a cooking pot around underneath the blanket. The weight of the ball and the lip of the pot create a dimple and gravity keeps the ball in the dimple as the pot is moved. The voltage on the charge-transfer electrode creates an electronic "dimple." Changing the voltages on nearby electrodes moves the dimple. In this way, groups of charged particles (electrons) can be pushed around without actually "touching" or losing them.

Readout Methods

There are three distinct methods for reading out the charge pattern of a CCD: full-frame, full-frame transfer and interline transfer (Fig. A3.4). Most early scientific CCDs used the first method, which operates as has just been described. Although **full-frame** readout provides the largest sensitive area for a given area of silicon, the lowest level of readout noise and the greatest photometric accuracy, it also has some disadvantages. One cannot both collect and read out signal at the same time. Unless some sort of shutter is used to prevent light from striking the sensor during vertical transfer, signal will be added to any packets that are transferred past bright features in the image, producing vertical streaking. This problem is more important when the exposure time is short relative to the readout time.

In **frame transfer** readout, at the end of the exposure, the entire charge pattern is rapidly (0.1–3 ms) transferred by charge-coupling to a second 2D **storage** array. The storage array is the same size as the **sensor** array and is located next to it but it is physically masked with evaporated metal to shield it from light. The charge pattern is then read out from the storage array while the sensor array collects a new image. Because vertical transfer can be much faster if the charge packets do not have to be read out, this system reduces streaking by up to 1000× but does not eliminate it and the need for a storage register reduces the fraction of the Si surface area that can be used for sensing by 50%.

In **interline transfer**, the masked storage cells are interlaced between the sensor cells (i.e., each pixel is divided into sense and read areas). After exposure, all charge packets can be moved to the readout array in less than a microsecond. This ability can be used as an electronic shutter to eliminate vertical smearing but, because at least half of the area of each sensor must be masked, and any light striking a masked area is lost, the "fill factor" of the sensor is reduced, proportionately decreasing QE_{eff}. A solution to the "fill-factor" dilemma is to incorporate an array of microlenses, aligned so that there is one above every pixel. With such a system, most of the light striking any pixel will be focused onto the unmasked area.[3] Although microlenses restore the QE_{eff} somewhat, the full-well signal possible is still limited by the smaller sensitive area.

WHAT COULD GO WRONG?

When I first heard the CCD story, it struck me as pretty preposterous! How could you get all the correct voltages (9 different voltage combinations per pixel shift, ~3.6 million for each TV frame, 108 million/s for video rate!) to the right charge-transfer electrodes at the correct times? How could you get all of the charge in a packet to stay together during a transfer? Wouldn't Poisson statistics apply, making even one transfer imprecise and the 2000 transfers needed to read out the top, right pixel of a 1000 × 1000 pixel array impossibly inaccurate? How long would the PEs stay free to be dragged around the lattice? Wouldn't the charge packets decay with time?

In fact, many of these problems did occur, but remedies to most have now been devised. The difference between a $300 commercial CCD camera and a $65,000, top-of-the-line scientific CCD can often be measured in terms of how many of these remedies have been implemented. Therefore, it is worthwhile trying to understand some of them so that one can buy what one needs. The following discussion will define and discuss some of the more important CCD technical specifications.

Quantum Efficiency

Quantum efficiency is the ratio of the number of impinging photons to the number of PEs produced.[4] Any photon with energy in the range of 1–100 eV striking crystalline Si has a very high probability of producing a PE. However, reflections and absorption by the overlying polysilicon electrodes,[5] reduce the QE of

[2] This may seem backwards until one remembers that any image of the real world is usually focused onto the CCD by a single, converging lens, a process that always inverts the image.

[3] This occurs only as long as the initial angle of incidence is near to normal, a condition met when CCDs are used for light microscopy.

[4] In the visible range, each absorbed photon makes only one PE.

[5] Kodak had pioneered the use of charge transfer electrodes made out of In and Sn oxides that scatter less light than do those made of poly-Si.

FIGURE A3.3. Four CCD readout patterns: Full-frame, frame-transfer, interline transfer and gain register (EM-CCD).

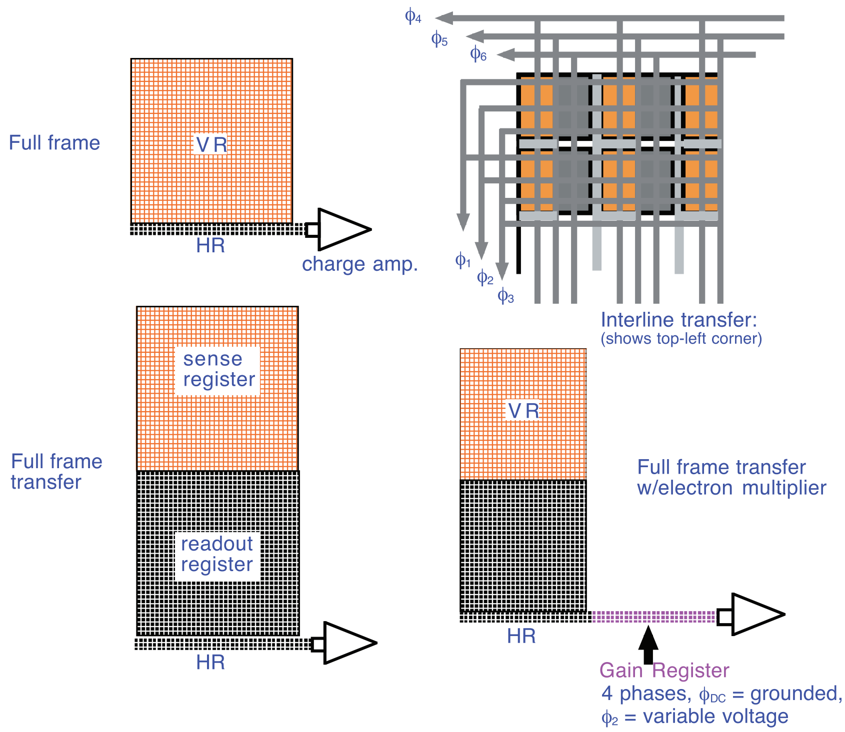

front-illuminated CCDs especially in the blue end of the spectrum. To reduce this effect, some UV-enhanced sensors are coated with fluorescent plastics, which absorb in the blue and emit at longer wavelengths. Others have their backs etched away and are turned over to permit the illumination to reach the light-sensitive area from the back side.[6] Figure A3.4 shows the intrinsic QE of different types of CCD (not Q_{eff}, which would take into account the light lost if some of the sensor is covered by charge storage areas). The effective QE can usually only be determined by actual measurement or by very careful evaluation of the published specifications ($QE_{effective} = QE_{intrinsic} \times$ fill factor).

Edge Effects

In early CCDs, PEs were often "lost" in the crystalline imperfections that are always present at the Si/SiO_2 junction. To avoid this, ion implantation is now used to make an N-doped, sub-surface layer called the buried channel about 1 µm below this surface (Fig. A3.2). This channel attracts the free PEs, keeping them away from the edge of the Si crystal. Any serious CCDs will have a buried channel but the need for ion-implantation keeps CCD chip prices high! Figure A3.5 shows the readout noise, in root-mean squared (RMS) electrons/pixel, for surface and buried-channel CCDs having two different pixel sizes. From this you can see that small pixels (here ~5.5 × 5.5 µm) have lower read noise than larger ones (~17 × 17 µm), mostly because the larger ones have higher capacitance and capacitance is the most important parameter of read-amplifier noise. One can also see that at readout speeds higher than 1 MHz (or 1 second to read out a 1024 × 1024 CCD), the read noise increases with the square root of the read speed.

Charge Loss

The lifetime of a PE (before it drops back into the ground state) depends on the purity and crystalline perfection of the Si and on other factors such as temperature. Generally it is long enough that little charge is lost during the exposure times commonly used in fluorescence microscopy. If necessary, it can be increased by cooling the detector, something often done to reduce dark charge.

Leakage or "Dark Charge"

Dark charge is the charge that leaks into a pixel during the exposure time in the absence of light. It can be thought of as the **dark current**[7] deposited into one pixel. Many processes other than photon absorption can add PE to the charge packet. The magnitude of this dark charge depends on the length of the exposure, and is substantially reduced by cooling. The rule of thumb is that for every 8°C of cooling, the dark charge is halved. As noted above, dark charge is principally a problem because it produces Poisson noise equal to the square root of its magnitude, and if this is left unchecked, it can significantly increase the noise floor of the CCD.

Since ~1987, a process called multipinned phasing (MPP) has been available to reduce dark charge build-up by about a factor of 1000, making it immeasurable in exposures up to a minute or so. This feature should be specified if one expects to use exposures longer than a few seconds without deep-cooling.

[6] Back-illuminated CCDs have to be thinned to 7–10 µm so that conduction electrons created near what would have been the back surface can respond to the fields created by the buried channel and the CC electrodes. Thinning increases cost and also reduces QE at longer wavelength where the absorption distance of the photons becomes comparable with the actual thickness. Back-illuminated CCDs are also more expensive because it is difficult to create electrical contacts with electrodes, etc., that are now on the bottom side of the chip.

[7] A current is a flow of charge measured in charge/time. The unit of charge is the Coulomb (c). The unit of current is the Ampere (A). One Amp represents a flow of one Coulomb/s or 6.16×10^{18} electrons/s.

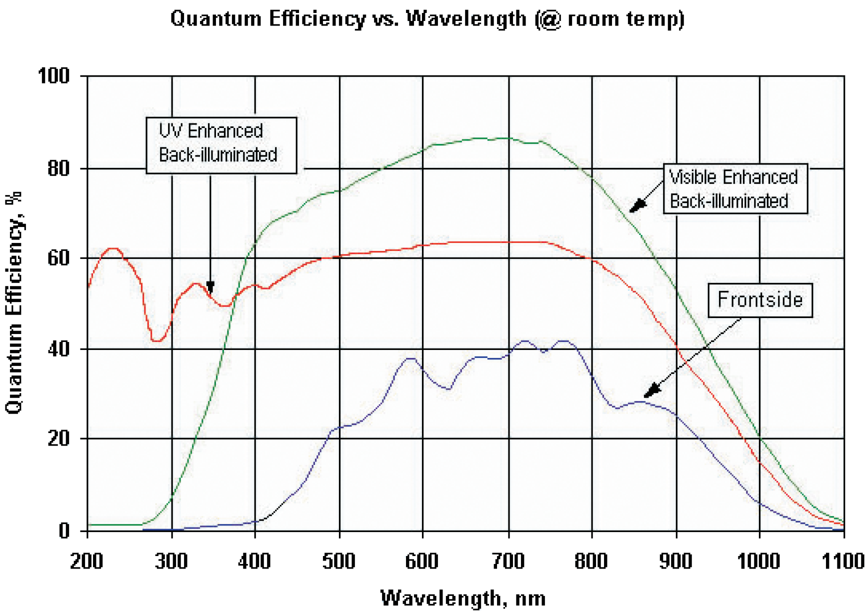

FIGURE A3.4. Intrinsic QE as a function of wavelength for a front-illuminated CCD (blue), a visible-enhanced, back-illuminated CCD (green) and a UV-enhanced CCD (red).

It should also be remembered that, while dark charge is never good, its average value can be measured and subtracted on a pixel-by-pixel basis, by subtracting a "dark image" from each recorded image as part of flat-fielding. However, because, by definition "dark" images contain very few photons/pixel, they have relatively high Poisson noise and low S/N. Therefore, a number of such images must be averaged to produce a correction mask that is statistically defined well enough that subtracting it from the data does not substantially increase the noise present in the final, corrected image.

This is not a problem when there are many counts in each pixel because the subtractive process of dark-charge normalization involves a change that is small compared with the intrinsic noise present in a large signal. It can be a problem when the black mask image is subtracted from a faint image that also contains only a few counts/pixel.

What cannot be removed by flat-fielding is the Poisson noise associated with the dark charge. This is equal to the square root of the number of electrons/pixel it represents. CCDs should always be operated such that the noise on the dark charge is less than the readout noise. On conventional CCDs this condition can usually be met quite easily by slightly cooling the sensor (0°C or about −20°C from ambient). The use of lower temperatures is complicated by the risk of condensing atmospheric water, a process that can be avoided only by enclosing the sensor in a vacuum chamber. Generally, a vacuum-hermetic enclosure, combined with good outgassing prevention, carries with it the significant benefits of more effective cooling, long-term protection of the sensor from moisture and other degrading organic condensates as well as the prevention of front-window fogging. At video rate, where exposures are short, dark charge is only a problem when the readout noise is reduced to <1 e/pixel, as it is when an "electron-multiplier" (EM) charge amplifier is used (see below and also Chapters 4 and 10). In EM-CCDs the read noise is so low that dark current becomes the main source of noise and cooling to −80°C becomes necessary.

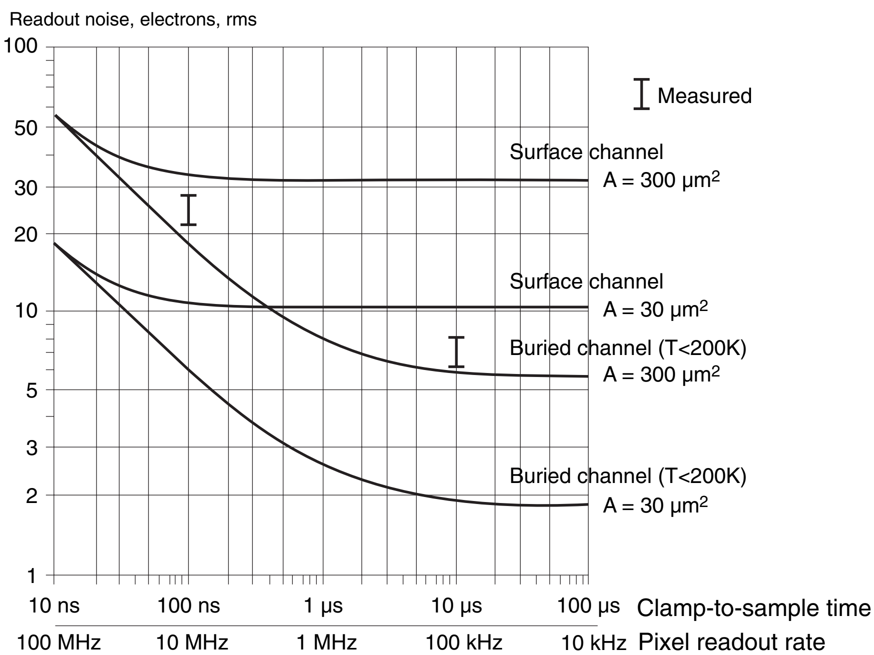

FIGURE A3.5. CCD field effect transistor (FET) noise as a function of pixel dwell time for large and small pixels and when using buried channel vs. surface channels. Smaller pixels have less read noise because they have less capacitance. Buried channels have almost 10× less read noise than surface channels.

TABLE A3.1. Typical Performance of Various Types of CCD Cameras. The "Sensitivity" Column Is a Reasonable Estimate of the Relative Suitability of the Camera for Detecting Very Faint Signals. It Spans a Very Large Range of Performance!

All CCDs are not equal						
Type	Grade	QE % (effective)	Noise (e/pix)	Sensitivity (relative)	Bit depth	Dynamic Range
Video	commercial color	10	200	1	10	1,000
	monochrome	20	200	2	10	1,000
Digital	1 Mhz, color	15	50	12	12	4,000
	1 Mhz, mono	30	50	24	12	4,000
	Back. Illum/ slow-scan	90	5	720	15	40,000
	LLL-CCD (EMCCD)	45	0.1	(18,000)*	?*	200,000

*Because the gain of the electron multiplier amplifier is unknown and large, it is not simple to measure, or even define, the sensitivity and bit depth of the EM-CCDs.

Blooming

As more photons are absorbed, the charge packet clustered around the buried channel grows and mutual repulsion between these electrons renders the field imposed by the charge-transfer electrode ever less successful in keeping the packet together. The maximum charge packet that can be stored without it overflowing into nearby pixels can be estimated by multiplying the pixel area (in square micrometers) by 600 PE/pixel (i.e., 27 k PEs for a 6.7 × 6.7 µm pixel, 540 k PEs for a 30 × 30 µm pixel). This overflow problem is referred to as "blooming" and, in CCDs for the home-video market it is limited by the presence of an *n*-layer, deeper in the Si. When the charge packet gets too big, mutual repulsion between the PEs forces some of them into this overflow layer, through which they are conducted to ground.

While this anti-blooming feature is convenient for removing the effects of the specular reflections found in images from everyday life,[8] it is not incorporated into many full-frame or frame-transfer scientific CCDs because it reduces QE for long-wavelength light. As this light penetrates farther into the Si crystal before being absorbed, much of it reaches the overflow layer where any PEs produced are lost.

Incomplete Charge Transfer

Sometimes, an imperfection in the Si will produce a pixel that "leaks" charge. Charge deposited into, or transferred through, this pixel will be lost, producing a dark vertical line above it. In addition, if one pushes the pixel clock too fast, some PE in the packet will not move fast enough and they will be left behind. In general however, on a slow-scanned, scientific CCD, fewer than 5 PEs out of a million are lost (or gained) in each, slower, vertical transfer and only 50 (0.005%) are lost during each, faster horizontal transfer. In such devices, the main noise term is Poisson noise for any signal level above ~20 PE/pixel,[9] and it seems hard to imagine doing much better than this except for signal levels <16 PE/pixel.

On the other hand, it is also true that the vast majority of CCDs made (those for camcorders, surveillance cameras and even many scientific applications), operate with much (100×?) less perfection. In microscopy today, we find CCDs that span this range of performance (Table A3.1).

All CCDs are not equal!!

CHARGE AMPLIFIERS

So far, I have described an image sensor in which up to 90% of the impinging photons make free PEs and explained how the charge packets that result from many photons hitting a given pixel can be conveyed to the **charge amplifier**, in a time-labeled manner and almost without change. Clearly the performance of the entire image detector will depend crucially on the capabilities of this amplifier.

What Is a Charge Amplifier?

Although most scientists have had some exposure to electronic circuits that amplify an input voltage or current, they may be less familiar with the operation of the type of charge amplifier found in a CCD. The following outline is presented to enable the reader to understand enough about the process to appreciate some of the important differences between the various types of CCD.

Because of the pulsatile nature of the CCD charge delivery system, the optimal way to measure charge packet size is to deposit it into a (very) small capacitor (the "read node") and then measure the voltage on this capacitor with a high impedance amplifier. As a field-effect transistor (FET) has an almost-infinite input impedance, it is ideal for this purpose and in fact, its existence makes charge-amplification possible.

There are two basic types of conventional CCD readout amplifier, non-destructive and destructive.[10] Both employ FET amplifiers.

Non-destructive ("skipper") amplifiers use an FET with a "floating gate" to sense the size of a charge packet by responding to the moving field that is produced as the packet is transferred along a charge-coupled register. Because the charge packet itself is not affected by this process, the process can be repeated hundreds or even thousands of times. If the results of all these measurements are averaged, very low readout noise levels (>±1 electron/pixel) can be obtained, but at the cost of a substantial

[8] Features such as the image of the sun reflecting off a shiny automobile can be over 1,000× brighter than the rest of the scene. Fortunately, such extremely bright features are seldom found in microscopic images unless a crystal of fluorescent dye occurs in the field of view.

[9] This calculation assumes that the read noise is 4 e/pixel, and this will be less than the Poisson Noise for any signal >16 PE. However, as many CCDs used in microscopy have >4 e/pixel of noise, this cut-off point should not be considered inflexible.

[10] The "electron-multiplier" amplifiers mentioned previously, act essentially as pre-amps to the conventional FET amps described here. They will be covered later in this Appendix.

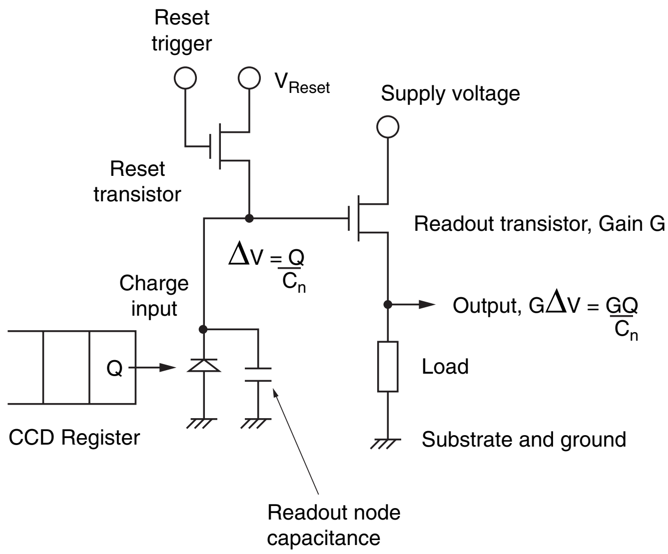

FIGURE A3.6. Destructive read-out amplifier for a CCD chip.

(×100 or ×1,000) increase in readout time and logical complexity. This approach might make sense on a Mars probe but it has not been used in microscopy to my knowledge.

Destructive readout amplifiers are more common, probably because they can operate more rapidly (Fig. A3.6). As implemented in a scientific CCD, the charge amplifier consists of the following components:

- overlying charge transfer electrodes to drag the next charge packet into the "read node"
- the read node itself: a 0.03–0.1 picofarad capacitor
- the sense FET
- the reset FET

In operation, fields from the overlying charge-coupling electrodes force a charge packet into the read-node capacitor, creating a voltage, V_c, that is proportional to the amount of charge in the packet. This voltage is sensed by the sense FET and the output is passed, via additional amplifiers, to the analog-to-digital converter (ADC) where the signal is converted into a digital number. Finally, just before the next charge packet is coupled into the read node, a reset FET discharges the capacitor, forcing V_c to zero, and allowing the read-FET to sense it again.

FET Amplifier Performance

The signal current (signal charge/s) coming from a CCD sensor is very small. Suppose that there were, on average, 400 PE in every pixel of a 512 × 512 pixel sensor.[11] Reading this out in one second would constitute a current of only 10^{-11} Amps. The current through the bulb in a home flashlight is 10^{10} times more. A very good conventional electronic amplifier designed to amplify this current would, itself, produce a random electronic noise signal larger than this, and electronic noise increases with readout speed, read-node capacitance and, to a lesser extent, temperature.

The success of the CCD in overcoming this limitation depends on two factors:

- The extremely small capacitance of the read node compared to that of any other photosensor such as a photodiode.
- Special measurement techniques such as correlated double-sampling

Clearly there are a lot of tricks to making the perfect CCD amplifier and not all CCDs employ them. Table A3.1 lists typical performance for a variety of common camera types.

NOISE SOURCES IN THE CHARGE-COUPLED DEVICE

Fixed Pattern Noise

When exposed to a uniform level of illumination, some pixels in a CCD array will collect more charge than others because of small differences in their geometry or their electrical properties. Consequently, it can be necessary to use stored measurements of the relative sensitivity of each pixel to normalize, or "**flat field**," the final dataset on a pixel-by-pixel basis. This is accomplished by first recording an image of a featureless "white" field. This is often approximated by a brightfield transmission image with no specimen, a process that will also record "inhomogeneity," or mottle, in the optical system. Differences in gain between pixels are evident as visible as nonuniformities in the digital signal stored in the memory and these are used to derive multiplicative correction coefficients.[12]

Unfortunately, one can only preserve the high precision of the CCD output if the coefficient used to normalize each pixel is equally precise. In any event, these correction coefficients vary with both the photon wavelength and the angle at which the light passes through the polysilicon electrodes on its way to the buried channel. This, in turn, depends on the details of the precise optical path in operation when an image is recorded and may even change with microscope focus! As the intrinsic noise of a pixel holding 360k, PEs is only ±600 electrons or 0.16%, pixel-to-pixel normalization for changes in sensor gain is seldom perfectly effective and consequently there is usually some level of "Fixed-pattern noise" superimposed on the final data.

In addition, the "white" image that must be used for pixel-level sensitivity normalization is itself subject to intrinsic noise (±600 electrons for a signal from a pixel with a full well charge of 360k electrons) and so multiplicative normalization may actually add some noise to the raw, uncorrected signal! Fortunately, if the white image can be defined by a multi-frame averages of several, nearly full-well "white" images, this normalization noise should only be noticeable when the image data to be corrected is similarly noise free. Without details of the signal levels present or the optical system in use, it is difficult to estimate the magnitude of normalization noise but it will be comparatively less important for images of faint objects containing few counts/pixel because these measurements are themselves less precise.

[11] Although this number may seem small, it is actually quite high compared to some uses in biological confocal microscopy. Many authors have found that in normal" use, a single-beam confocal microscope used to image a fairly faint stain will count 4–8 PE/pixel in bright areas of the image. Allowing that the effective QR of a good CCD will be ~10× higher than that of the photomultiplier tube used in the confocal microscope, this makes the expected peak CCD signal in an image from a disk-scanning confocal microscope only 40–80 PE/pixel.

[12] These correction coefficients are small and only needed when operating on images involving large numbers of photons (and consequently having relatively low Poisson noise and good S/N).

It should be also noted that the vignetting and "mottle" visible in images characteristic of video-enhanced contrast microscopy will produce small intensity errors in the data obtained by both widefield and confocal. However, this noise term will be more noticeable in widefield where more photons are used and hence the precision of the data is greater. Mottle is produced by dirt and surface imperfections on any optical components that are not located exactly at aperture planes, as well as by non-uniformities in the image sensor. Fortunately, to the extent that it is stable with time, mottle will be removed by the flat-field correction for CCD sensitivity just discussed.

What will not be removed is any change in signal caused by stray light (room light, light that goes through filters designed to remove it, etc.). The simplest test of any CCD set-up is to record an image of "nothing" (i.e., room dark, no excitation, no specimen etc.). Then do the same with 100× longer exposure time with the room lights at your normal operating level. Now adjust the display look-up tables so that you can "see the noise" in both the images on the screen. Although the only difference between the two images should be increased dark noise in the image with the longer exposure, this is seldom the case.

Noise from the Charge Amplifier

Noise is generated by both the readout and the reset FETs in the charge amplifier. Noise generated in the readout FET reaches the ADC directly. If thermal noise in the reset FET prevents it from completely discharging the read-node capacitance, it produces a random offset at each pixel (i.e., the read-node voltage is not reset exactly to zero). This is referred to as Reset Noise and has the effect that the dark charge seems to vary from pixel to pixel. Fortunately, Reset noise can be almost eliminated by employing the technique of Correlated Double-sampling (CDS) in the readout amplifier. In CDS, the circuitry of the charge-to-voltage amplifier is modified so that the output is proportional to the difference between the value of V_c just after the reset pulse and its value after the next charge packet has been inserted.

Although CDS essentially eliminates the effect of reset noise, it also distorts the noise spectrum. On the one hand, this distortion has the beneficial effect of converting the low frequency, 1/F noise from the FET into broadband noise which is more easily treated theoretically and which is less visually distracting than the short, horizontal flashes characteristic of 1/F noise.[13] On the other hand, it means that the input to the ADC must be carefully frequency-filtered. This filtering can be implemented either by employing RC circuits or by using dual-slope integration (DSI) in the ADC itself. If there are large intensity variations between neighboring pixels, the use of RC circuits will effectively compromise the large dynamic range of the CCD. Therefore, ADCs using DSI are employed on most slow-scan scientific, cooled-CCDs.

The fact that CDS and, in particular, DSI work best at low readout speeds is a final reason why most scientific CCDs operate best at relatively low readout speeds (Fig. A3.5). The other two reasons are improved charge transfer efficiency and the reduction in broadband electronic noise from the FETs (noted above.)

[13] In a CCD without CDS, noise features will seem to be smeared sideways, while in one with CDS, they will appear as one-pixel-wide stipple with no directionality.

Where Is Zero?

A final important feature of the CCD readout is that, compared to the photomultiplier tube (PMT), it is relatively difficult to determine the exact output signal level that corresponds to a zero-light signal. A properly operated PMT never records negative counts. However, as the electronic readout noise of a cooled-CCD is an RMS function with both positive and negative excursions, there will be some pixels that measure lower than the mean value of the zero-light pixel intensity distribution.

To ensure that no data is "lost," scientific CCDs are usually set up so that the zero-light signal is stored to be a few tens of digital units (ADU) above zero. A histogram of numbers stored from a "black" image will show a Gaussian-like peak centered at the offset and with a half-width equal to 2× the RMS read noise (see Fig. 4.20). This offset makes it more difficult to apply the gain and offset normalization procedures to images that record only a few detected photons in each pixel, a factor that will become more important as CCDs are increasingly used to image living cells that cannot tolerate intense illumination and which therefore produce substantially lower signal levels.

A NEW IDEA: THE GAIN REGISTER AMPLIFIER!!

Early in 2002, a new type of readout amplifier was introduced by Texas Instruments (Houston, TX) and E2V Technologies (Chelmsford, UK). As only E2V makes back-illuminated sensors, I will describe their system but both work along similar lines. E2V originally referred to their device as the "gain register" and its purpose is to amplify the size of the charge packet before it arrives at the read node. Although the term gain register has recently been replaced by the term "electron multiplier", it is important to remember that these new detectors work on a completely different principle from that employed in intensified-CCDs.

The gain register superficially resembles an additional HR, with two important differences:

- There are 4 phases rather than the usual 3 and the new phase consists of a grounded electrode located between ϕ_1 and ϕ_2.
- The charge transfer voltage on ϕ_2, is now variable, between +35 and +40 volts rather than the usual +15 volts.

As a result, when ϕ_2 is excited, there exists a high electric field between it and the grounded electrode. The high field accelerates the electrons in the charge packets more rapidly as they pass from ϕ_1 to ϕ_2 with the result that each PE has a small (but finite; usually in the range of 0.5% to 1.5%) chance of colliding with a lattice electron and knocking it into the conduction band (Fig. A3.7). Assuming the 1% gain figure, this means that for every 100 PE in the packet, on average one of these will become two electrons before it reaches the space under ϕ_2. Although this seems like a trivial improvement, after it has been repeated as part of the 400 to 590 transfers in the gain register, a total average gain of hundreds or even thousands is possible. If the voltage on ϕ_2 is reduced to normal levels, the sensor operates as a normal CCD.

As a result, a single PE can be amplified sufficiently to be safely above the noise of the FET amplifier, even when it is operating at speeds considerably higher than video rate (35–50 MHz, vs. 13 MHz for video). As the amount of gain depends exponentially on the exact voltage on ϕ_2, it is possible to "dial in" the amount of gain needed to keep the signal level well above the noise of the FET amplifier. However, it is important to remember that

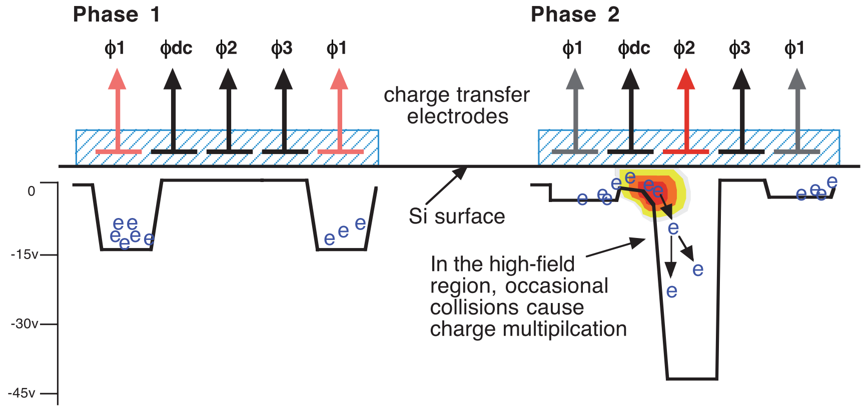

FIGURE A3.7. Energy diagram of an electron-multiplier CCD amplifier. The high field region that occurs between ϕ_2 and ϕ_{DC} when ϕ_2 goes strongly positive (right) causes about 1% of the electrons passing this region to collide with a lattice electron with sufficient energy to boost it to the conduction band. Repeated over hundreds of transfers, this process is capable of providing an average amplification of hundreds or even thousands of times.

the use of high EM gain will tend to saturate the "full-well" capacity of later pixels in the gain register, reducing intra-scene dynamic range.[14] Although this effect can be reduced to some extent by making each pixel in the gain register (and the read node) larger, this approach is limited by the fact that one triplet of electrodes can control a band of silicon only ~18 µm wide and because a larger read-node capacitance increases the read noise of the FET amplifier.

In sum, the gain-register CCD works like a normal fast-scan CCD with no read noise. The high scan speed makes focusing and sample scanning quick and easy and the device preserves the full spatial resolution of the CCD because the charge packet from one pixel is always handled as a discrete entity (unlike in an intensified-CCD). Of course, with fast readout, there is less time to accumulate much signal and the resulting image may have considerable Poisson noise. But this is not the camera's fault!

Alternatively, the output of many frames can be summed to reduce Poisson noise or, if the signal is bright, one can turn off the EM gain and have a fully functional scientific CCD.[15]

If the gain-register CCD is read out fast, there is so little time for dark charge to accumulate that cooling would seem unnecessary until one remembers that one can now "see" even one PE of dark-charge above the read noise. Because multi-pinned phasing (MPP) is less effective during the readout clockings, significant dark charge can be generated during readout. If the exposures are short, this source of dark charge becomes significant, and in an EM-CCD, even one electron is significant! In practice, the best performance is obtained when the EM-CCD is cooled to between −80° and −100°C.

[14] If a register designed with enough pixel area to hold a normal full-well charge of 30,000 electrons, is used with a gain of 100×, then the pixels near the end of the gain register will become full whenever the original charge packet has >300 electrons.

[15] Because, as noted above, the read node of the FET amplifier at the end of the gain register in an EM-CCD has a relatively large capacitance, E2V offers two separate FET readout amps. The one mounted at the end of the gain register is optimized for fast readout. The other is mounted at the end of the HR not connected to the gain-register, has low input capacitance and is optimized to read out slowly with low noise. Signal is sent to the latter by reversing the charge transfer sequence applied to the HR.

EM-CCDs have one other important form of "dark noise" called Clock Induced Charge (CIC, also known as spurious noise). CIC typically consists of the single-electron events that are present in any CCD, and are generated by the vertical clocking of charge during the sensor readout. The process involved is actually the same impact ionization that produces multiplication in the gain register; however, levels are much lower because lower voltages are involved. In conventional CCDs, CIC is rarely an issue as single-electron events are lost in the read noise. However, in EM-CCDs where the read noise is essentially zero and dark charge has been eliminated through effective cooling, CIC is the remaining source of single-electron, EM-amplified noise. If left unchecked, it can be as high as 1 event in every 7 pixels. Fortunately, it can be minimized by careful control of clocking voltages and by optimizing the readout process to cope with faster vertical clock speeds (down to 0.4 µs/shift). This leaves a detector with less than one noise pulse in every 250 pixels: a detector extremely well adapted for measuring zero!

Of Course, There Is One Snag!

The charge amplification process is not quite noise free because the exact amount by which each electron in the packet is amplified varies in a stochastic manner (i.e., some electrons are "more equal" than others.). The statistical arguments are discussed in a paper found at the URL listed below and in Chapter 4. In summary: as the multiplicative noise inherent in the charge multiplication process creates noise that has a form very similar to that produced by Poisson statistics, the easiest way to think of its effect is to assume that the amplifier has no noise at all but that the signal being fed into it is half as big as it really is. In other words, the camera will work perfectly but it will work as though it has a QE that is only half of what it really is. Back-illuminated sensors are now available with an intrinsic QE of about ~90% or ~45% when used in the gain-register mode. This is 5–10× better than the performance available from most PMTs especially in the red end of the spectrum.

It is worth noting that one can use electron multiplication and still maintain the full QE by using the detector in photon-counting mode, as is now being done by many astronomers. Photon counting is only possible when one is able to confidently see a single-photon event as different from any dark event and when the number of photons collected during an exposure is so low that there is little probability of >2 photons falling into a pixel. To implement photon-counting, one records a sequence of short exposures containing "binary"-type single-photon data, and combines them to generate an image that is free of multiplication noise.[16] To be useful for recording dynamic events in living cells, an extremely fast frame rate would be needed. This may be more possible with some future EMCCD sensors.

(More info on EM-CCDs at http://www.emccd.com and http://www.marconitech.com/ccds/lllccd/technology.html)

PART II: EVALUATING A CHARGE-COUPLED DEVICE

A. Important Charge-Coupled Device Specs for Live-Cell Stuff!

Although in Part I, much time was spent discussing cooled, slow-scan, scientific CCDs, in fact, these have not been much used in biological microscopy since Sony introduced the ICX085, 1k × 1.3k, micro-lens-coupled, interline-transfer chip in the late 1990s. Although initially developed not for the scientific market but to meet the needs of the Japanese high-definition TV standard, these chips offered a set of practical advantages that biologists found very appealing:

— As an interline transfer chip, it needed no mechanical shutter and could be run so as to produce a continuous stream of images.
— The high readout speed (up to 20 MHz) allowed real-time imaging compared with the 5–10 s/frame readout then common.
— The 6.45 × 6.45 µm pixels were small enough to sample the image produced with high-NA 40× and 60× objectives.
— The 1k × 1.3k raster size was both sufficient for most biological microscopy and significantly higher than that of the other scientific chips then available.
— Mass production allowed the development of a micro-lens array that increased the QE_{eff} to an acceptable level for a front-illuminated, interline chip and did so at a price biologists could afford.

As a result, the majority of CCDs sold for use in biological microscopy today use this chip or its higher-QE cousin, the ICX285. Although mass production made quality CCDs available to many who formerly could not have obtained one, it is important to remember that the read noise of ±8–24 electrons/pixel (depending on read speed) is substantially higher than ±2–3 electrons/pixel that characterized the best, slow-scan, scientific CCDs. Although, as noted below, the difference is only important if the dimmest pixel records fewer than ~50 electrons, and this seldom occurs in widefield fluorescence microscopy, the disk-scanning confocal fluorescence microscopes now available do provide an image in which this difference is significant (Chapter 10).

1. Quantum Efficiency (QE):

QE is the ratio of photons striking the chip to electrons kicked into the conduction band in the sensor. It should be at least 40% and on back-illuminated chips, it can reach 90% (with somewhat higher fixed-pattern noise).

The **fill-factor** is the fraction of the sensor surface actually sensitive to light. On the best frame-transfer CCDs, it can be almost 100%. On interline transfer CCDs it may be only 40%. Light not absorbed in a sensitive area is lost, reducing the QE_{eff} of the sensor proportionally.

Factors affecting QE:

Front-illuminated chips

- Light is scattered by the transparent, polysilicon charge-transfer electrodes that overlie the photosensitive silicon surface.
- This scattering is more severe at shorter wavelengths. Light that is scattered is not detected.
- As blue light is absorbed nearer to the surface than red light, and "deep electrons" may go to the wrong pixel, CCD resolution may be a bit lower than the pixel count at longer wavelengths, especially on chips with small pixels.
- Best QE: ~20% blue, ~35% red/green
- Two efforts to improve the QE of front-illuminated chips include "Virtual Phase" (one phase "open," Texas Instruments, Houston, TX) and the use of indium oxides for the transfer electrodes (Kodak, Rochester, NY). These have increased peak QE to the range of 55%.

Micro-lens array chips

- Sony has pioneered a process in which a micro-lens is mounted above each pixel of a front-illuminated, interline-transfer CCD. The lenses focus most of the impinging light onto a part of the CCD where reflection losses are least, pushing the QE to 65% in the green, less in the red (because of losses to the overflow drain) and purple.

Back-illuminated chips

- Made by thinning the silicon and then turning it over so that the light approaches the pixel from what would have been the back side. This avoids scattering in the transfer electrodes and increases the QE to about 90% in the green and >70% over the visible range.
- More expensive because of the extra fabrication.
- Slightly less resolution and more fixed-pattern noise, caused by imperfections in the thinning operation, and the presence of two sets of surface states.

Color Chips

- One-chip color sensors employ a pattern of colored filters, one over each pixel. Light stopped by any such filter cannot be detected and is therefore lost. The QE of such sensors is therefore at least 3× lower than for an otherwise comparable monochrome chip.
- 3-chip color sensors use dichroic mirrors to separate the "white" light into three color bands, each of which is directed to a separate monochrome CCD sensor. While this would seem to ensure that "all photons were counted somewhere," because such systems seldom employ microlenses, their effective QE is not much better than the 1-chip color sensors and alignment of the signal light is important.[17]

[16] There is no multiplicative noise because any spike above the FET noise floor counts as one electron, no matter how much it has been amplified.

[17] While the QE is not much better, the resolution of the 3-chip camera is the same as that of each chip, without the interpolation needed to disentangle the 3 colored images from the output of a 1-chip color sensor.

- Color can be detected by making sequential exposures of a monochrome chip through colored glass or LCD filters. This produces the same QE losses as the patterned filter but has the advantage that it can be removed when higher sensitivity is needed. This design is not suited for imaging moving objects.

2. Readout Noise:

This spec is a measure of the size of the pixels and the quality of the circuitry used for measuring the charge packet in each pixel. It is measured in "±RMS electrons of noise" (i.e., 67% of a series of "dark" readings will be ± this much).

- A good scientific CCD camera should have a noise level of <±5 electrons at a readout speed of 1 M pixels/second.
- The readout noise increases with the square root of the **readout speed** (see Table A3.2).
- **NO Free Lunch!** A chip that has ±5 e RMS of noise when readout at 100 k pixels/sec (or 10 seconds to read out a 1024 × 1024 chip), should produce ±50 e RMS of noise if read out at 10 M pixels/sec (or 0.1 sec to read the same chip).

What Is "Good Enough"?

Very low readout noise is only essential when viewing very dim specimens: luminescence, or low level fluorescence. Read noise is only a limitation when it is more than the statistical noise on the photon signal in the dimmest pixel (i.e., >sqrt of the number of detected photons = sqrt # electrons).

Consider the signal levels that you plan to use. Will the darkest important part of your image have zero signal or do you expect some background signal from diffuse staining or out of focus light? If the dimmest pixel in your image represents ~100 electrons, then the Poisson or statistical noise on this background signal will be ±10 electrons. "Adding" an additional ±10 electrons of readout noise will not make much difference to a measurement of this background signal and it will be even less significant when added to the even greater Poisson noise present in pixels where the stained parts of the image are recorded.

This is especially true because RMS noise signals add as the "sqrt of the sum of the squares" (i.e., the total noise from ±10 electrons of readout noise and ±10 electrons of Poisson noise is only sqrt (100 + 100) = ±14 electrons).

On the other hand, if you are really trying to keep those cells alive and you find that 2,000 electrons in the bright areas is enough, the dark areas may now be only 50 electrons. As the sqrt of 50 is about ±7, an additional ±10 electrons of readout noise may no longer be acceptable, but only if you have to make measurements in the dark areas on your image. In this case, the obvious choice is a slower, quieter CCD or an EM-CCD.

While in widefield fluoresecence, the background stain level is seldom so low that the sqrt of the signal recorded is lower than the read noise, the disk-scanner does provide such an image (Chapter 10). As one of the main advantages of disk-scanning is that one can scan an entire image plane very rapidly, the fact that one can read out the EM-CCD very rapidly without increasing the read noise makes it the ideal detector for this type of scanner (or, indeed for high-speed line scanning confocal microscopes).

3. Pixel Size:

- **Nyquist sampling:** The size of a pixel on the CCD is, in itself, not very important **BUT** one must satisfy the Nyquist criterion: The pixels on the chip must be **~4+** smaller than the smallest **features** focused onto it[18] (see Chapter 4): **Pixel size on the chip determines the total specimen-to-chip magnification needed!**

TABLE A3.2. Dynamic Range and Pixel Size

	12-bit camera w/small pixels	14-bit camera w/large pixels
Pixel Size	6.7 × 6.7 µm	24 × 24
Full Well	27,000	345,000
Least significant bit =	**6.5 electrons**	**21 electrons**
Implied noise level	**±13 electrons**	**±42 electrons**

Two examples:

a. 1.4 NA 100× objective and a 1× phototube.
   - The Abbe Criterion resolution @ 400 nm is about 0.22 µm. Magnified by a total magnification of 100×, this becomes 22 µm at the CCD.
   - A CCD having 8 × 8 µm pixels samples such an image adequately (~2.8 pixels/resolution element).

b. 1.3 NA 40× objective and a 1× phototube.
   - The Abbe Criterion resolution @ 400 nm is now 0.25 µm. Magnified 40× this becomes, 10 µm.
   - A CCD having 8 × 8 µm pixels is inadequate to sample this lower-mag, high-resolution image.

If you must use this objective, you need either a higher mag phototube (2.5×) or a chip with 3 × 3 µm pixels or (as CCD pixels are seldom this small), some combination.

- **Saturation signal level:** The maximum amount of signal that can be stored in a pixel is fixed by its area. The proportion is 600 electrons/square µm, so a 10 µm × 10 µm pixel can store a maximum of 60,000 electrons before they start to bleed into neighboring pixels. In practice, as fluorescent micrographs of living cells seldom produce signals this large, large pixels are usually unnecessary.

However, the saturation level also represents the top end of another spec, the **dynamic range**. This is usually quoted as 12-bit (4000 : 1) or 14-bit (16,000 : 1) etc., and represents the ratio between the full-well saturation level and the readout noise. **Therefore**, a camera with relatively high readout noise can still look good in terms of dynamic range if it has large pixels and hence a high full-well capacity. Conversely, a 12-bit camera with small pixels can have less actual noise-per-pixel intensity measurement than a 14-bit camera with large pixels.

In this case, the noise level of the 14-bit camera is >3× that of the 12-bit camera. Your signal/pixel would have to be 3× larger in order to be "seen" when using this particular 14-camera.

4. Array Size:[19] The argument for small

- Assuming 0.1 µm pixels (referred to the object plane), a 512 × 512 pixel chip will image an area of the specimen that is about 51 × 51 microns. If this is enough to cover the objects you need to see, this small chip has a lot of advantages over chips that are 1024 × 1024, or larger.
- Lower cost

[18] Of two times smaller than the "resolution," as defined by Rayleigh, or Abbe.
[19] The array size refers to the number of lines and pixels in the sensor, not to its total area.

- 4× fewer pixels to read out, meaning either:
  - 4× slower readout clock, giving 2× lower readout noise.
  - Same clock speed and noise level but 4× faster frame time. (Easier to scan specimen to find the interesting part! Time is money!)
- 4× less storage space needed to record data.

The argument for big:

- Manufacturing improvements are reducing readout noise levels at all readout speeds, and CCDs with more pixels often also have smaller pixels which can lead to lower read noise. If your labels are bright, having a larger chip allows you to see more cells in one image (as long as they are confluent!). Assuming that Nyquist is met in both cases, a large print of an image recorded from a larger sensor always looks sharper than one from a smaller array.
- **Binning:** Binning refers to the process of summing the charge from neighboring pixels before it is read out. This increases the size of each charge packet read (making it look brighter) and reduces the number of pixels. For example: 2 × 2 binning allows the owner of a 1024 × 1024 chip to obtain the speed/noise performance similar to the smaller chip (512 × 512) and to do so in a reversible manner. However, the optical magnification may need to be increased to preserve Nyquist sampling.

Before deciding that you need a larger chip, compare what you would get if the same money were spent on another scope/CCD/graduate student!

Bottom line:

- If more pixels means smaller pixels, they will each catch fewer photons unless the magnification is reduced proportionally. More pixels at the same frame[20] rate mean somewhat higher read noise because the pixel clock must go faster.

5. Readout Speed:

Although readout speed has been discussed above, we haven't mentioned that some good CCD cameras have variable speed readouts and the new EM-CCDs impose no high read speed penalty (Table A3.3).

It is convenient to be able to read out the chip faster when searching and focusing as long as one can then slow things down to obtain a lower read noise in the image that is finally recorded. However, the read speed is only one limitation on the frame rate: if the signal level is so low that 1 s/frame is required to accumulate enough signal to be worth reading out, then reducing the read time much below 0.1 s loses some of its appeal.

Faster readout speeds are particularly important for moving specimens, especially when doing widefield/deconvolution or when following rapid intracellular processes, such as vesicle tracking or ion fluxes.

TABLE A3.3. CCD Specifications

Array size	Pixel Clock Rate	Noise level*	Frame time	Frame rate/s
640 × 480	13 MHz (video rate)	200 e/pixel**	0.033	30
512 × 512	100 kHz	5 e/pixel	2.5 sec	0.4
	1 MHz	15 e/pixel	0.25 sec	4
	5 MHz	35 e/pixel	0.05 sec	20
1024 × 1024	100 kHz	5 e/pixel	10 sec	0.1
	1 MHz	15 e/pixel	1 sec	1
	5 MHz	35 e/pixel	0.2 sec	5

*Assumes conventional FET circuits. **The readout noise is relatively higher at video rate because the higher speed often precludes the use of various techniques, such as correlated double sampling, that reduce readout noise.

6. Shutter Stability:

Though not strictly a CCD spec, electronic (LCD) or mechanical shutters are often built into modern CCD cameras.[21] The latter have the disadvantages of producing vibration and having a limited lifetime but the advantage that they transmit all of the light when they are open (even an "open" LCD can absorb >50% of the light, other electronic shutters may be better).

There seems little point in having a camera capable of recording (say) 40,000 electrons/pixel with an accuracy of ±200 e (or 0.5%!) if the shutter opening time is only accurate, or even reproducible, to ±10%. If one shutters the light source instead of the camera, similar limitations apply.

7. User-friendliness:

State-of-the-art cameras often seem to have been designed to make sure that no one unwilling to become a devotee of "CCD Operation" can possibly use them efficiently! Start off by asking to see an image on the screen, updated and flat-fielded at the frame-scan rate and showing as "white" on the display screen, a recorded intensity that is only ~5% of the full-well signal. This is where you should do most living-cell work. Then ask the salesman **to help you to save time-series of this image**. Increase the display contrast until you see the noise level of the image, both before and after "flat-fielding." Put a cursor on one pixel in the top frame of the stack and plot its intensity over the series.

8. "The Clincher" (Well, at least sometimes . . .):

Ask him/her what the intensity number stored in the computer for some specific pixel means, in terms of the number of **photons** that were recorded at that location, while the shutter was open. To answer, the salesperson will have to know the QE, the fill factor and the conversion factor between the number of electrons in a pixel and the number stored in the computer memory (sometimes called the gain-setting). To help them out, any "real" scientific CCD camera has the latter number written, by hand, in the front of the certification document (usually a number between 3 and 6). If the salesman doesn't understand the importance of this fundamental number, what hope is there for you? (Hint: It is important because the Poisson noise is the sqrt of the number of electrons in the well, not the sqrt of some arbitrarily proportional number stored in your computer.)

B. Things That Are (Almost!) Irrelevant When Choosing a Charge-Coupled Device for Live-Cell Microscopy

1. Dynamic Range:

This is the ratio of the "noise level" to the "full-well" (or maximum) signal. Although 16-bit may sound a lot better than 12-bit, you need to think before you are impressed.

The noise level should not be more than 5 electrons/measurement. Period!

[20] The readout speed of a 2 × 2 binned 1024 × 1024 is a bit slower than an actual 512 × 512 because twice as many vertical clock cycles are needed, and one still needs to read out pixel by pixel in the horizontal direction.

[21] Often the same advantage can be gained by shuttering the light source. This may become more common as pulsed laser or light-emitting-diode light sources are introduced (see Chapters 5 and 6).

Twelve bits is 4,000 levels. If the first level represents 5 electrons (in fact, it should represent half the noise or 2.5 e), then the 4,000th represents 20,000 electrons or (assuming a QE of 50%), about 40,000 photons/pixel/measurement.

How often do you expect to be able to collect this much signal from an area of a living cell only 100 × 100 nm in size? You should be able to get a good, 8-bit image using only 6% of the dynamic range of a 12-bit CCD (Fig. A3.8).[22]

As the "full-well" signal is only proportional to the area of the pixel on the chip (area in sq. μm × 600), the dynamic range is only really impressive if it is high **AND** a chip has small pixels. Then it means that the readout noise is low. A test for actual dynamic range is described below.

Bottom line: For disk-scanning confocal microscopy, a large dynamic range is only important if it reflects a low readout noise level.

Easier to just check the readout noise!

2. High Maximum Signal (high, full-well number, because of large pixels):

On living cells, you will probably never have enough light to reach a full-well limit of even 20,000 electrons. Even if you do, there are better ways to use it (more lower-dose images to show time course?).

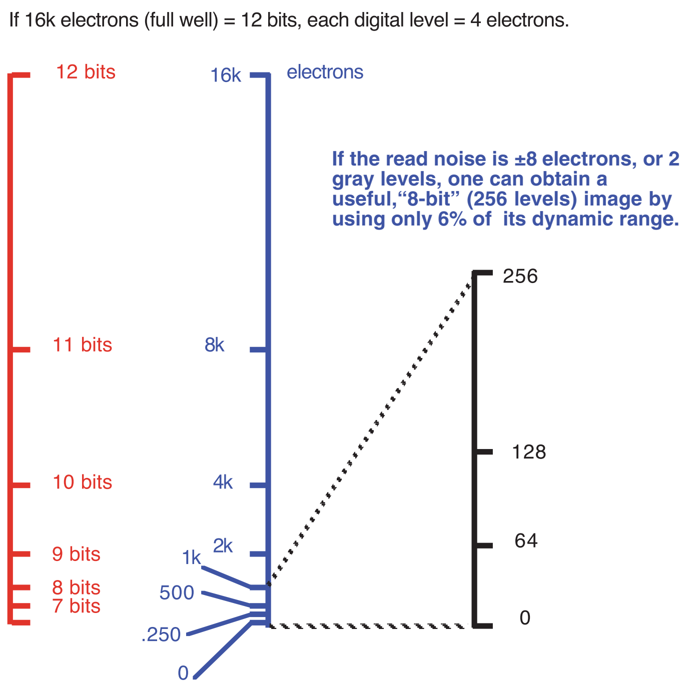

FIGURE A3.8. Not using the full dynamic range of a CCD. As most scientific CCDs have more dynamic range than one "needs" in live-cell fluorescence microscopy, the excitation dose to the specimen can be reduced if one sets up the CCD control program to display an 8-bit image using only the bottom 1,024 levels of a 12-bit image. Such an image is more than adequate for many functions in live-cell biological microscopy (particularly when other factors such as dye-loading etc., may cause larger errors) and will require only 6% as much signal as would a "full-well" image.

[22] Remember, given optical and geometrical losses, you can collect no more than about 3–10% of the photons produced, and, each fluoroscein molecule will only produce perhaps 30,000 excitations before "dying."

3. "Imaging Range" "Sensitivity" (or anything measured in LUX):

Stick to something you (and I?) understand: Photons/pixel or electrons/pixel. The other conversions are not straightforward.

4. "Neat Results":

Unless you know how well stained the specimen is, **you cannot evaluate an image of it in a quantitative manner**. (Though you may not want to admit this!) By all means, view your own specimens, but viewing "test specimens" that are not expected to fade and have a known structure (fluorescent beads in some stable mounting medium?) facilitates A/B comparison. If you do use your own test specimens to compare cameras, be sure to view them on the same scope, and with the same conditions of pixel size and readout time etc.

Better still . . .

C. A Test You Can Do Yourself!!!

Set up each camera that you want to evaluate on a tripod, add a C-mount lens, and an ND 3 or ND 4 filter. Hook up a monitor or computer and view some scene in your laboratory under ordinary illumination (avoid light from windows which may vary from day to day).

Close the lens aperture down until you can no longer discern the image (see Fig. 4.20). This is the "noise-equivalent light level": the signal level at which the electron signal (i.e., photons/pixel × QE) just equals the total noise level. Your measure is the aperture at which the image disappears.[23] Because it is sensitive to both QE and readout noise level, this is a very useful measure of what we all think of as the "sensitivity." Of course, the signal level depends not only on the light intensity but also on the exposure time and the pixel area, so make sure to keep the former constant and make allowances for the latter.

If you do not have even these meager facilities (a C-mount lens, an ND filter, a tripod and some time), take an image of nothing. Look at "no light" for one second, and for 100 seconds. Ask to see a short line profile that plots intensity vs. position along a line short enough that one can see the intensity of each individual pixel. The difference in the average intensity between the short and long exposure is a measure of the leakage.[24] With a little calibration from the published full-well specs (a spec less open to "interpretation" than "noise"), you can even get a direct measure of the read noise level from these dark images. (It should be the standard deviation of the values as long as they are counted in electrons, not "magic computer units" and as long as fixed-pattern noise is not a factor.) And just trying to work it all out will give you some idea if the salesman knows anything . . .

D. Intensified Charge-Coupled Devices

Intensified CCDs (ICCDs) are just that: the mating of an "image intensifier" to a CCD. The idea is that the photon gain of the intensifier (can be 200–2000×) will increase the signal from even a

[23] If the lens doesn't have a calibrated aperture ring, you can open the aperture all the way and reach the "threshold" exposure level by reducing the **exposure time** and adding ND filters. Remember to also correct for pixel area. Larger pixels intercept more photons.

[24] With a good EM-CCD, this measurement can be done using a short exposure and high EM gain, then counting the number of amplified dark charge/CIC spikes across a typical line of the raster.

single photoelectron above the read noise of the CCD. This occurs, and can be particularly useful where fast readout is needed such as when measuring ion transients. Finally, pulsing the voltage on the intensifier section makes it possible to shutter ("gate") the camera on the ns time scale, making the ICCD useful for making fluorescence lifetime measurements (Chapter 27, *this volume*).

However, ICCDs do not have the photometric accuracy of normal CCDs for a number of reasons:

- The relationship between number stored in memory and the number of photons detected is generally unknown and variable.
- The intensifier photocathode has low QE[25] (compared to that of a back-illuminated CCD).
- The "resolution" is generally only dimly related to CCD array size because of blooming in the intensifier. To check this, reduce light intensity until you can see the individual flashes produced by single photoelectrons. See how many lines wide they are. (They should be one line wide.)
- They have additional noise sources: phosphor noise, ions in intensifier section create flashes, high multiplicative noise in the intensifier section greatly decreases QE_{eff}, etc.
- Photocathode resistivity can produce "dose-rate" effects: non-linearities in which the recorded intensity of the brightest areas may depend on (and affect) the brightness of nearby features.

Because I expect that EM-CCDs such as those mentioned above will soon supplant ICCDs except where fast gating is needed, I have not gone into more detail here. For more info, go to: http://www.stanfordphotonics.com/

[25] And the GaAsP photocathode with better QE, have to be cooled, making the assembly very expensive.

ACKNOWLEDGEMENTS

The author would like to thank Dr. J. Janesick, formerly of the Jet Propulsion Lab (California Institute of Technology, Pasadena, CA), for many conversations about CCD operation and for the original sketches for Figures A3.1, A3.4, and A3.5 and to Colin Coates, (Andor Technologies, Belfast, UK) for his helpful comments on the manuscript and for Figure A3.6.

REFERENCES

Inoue and Spring, 1997, Video Microscopy, Second Edition, Plenum, New York, 1-741, particularly Chapters 5–9.

Pawley, J.B., 1994, The sources of noise in three-dimensional microscopical data sets, *Three Dimensional Confocal Microscopy: Volume Investigation of Biological Specimens*, (J. Stevens, ed.), Academic Press. New York, 47-94.

http://www.springer.com/978-0-387-25921-5

Handbook of Biological Confocal Microscopy
Pawley, J. (Ed.)
2006, XXVIII, 985 p. 603 illus. in color., Hardcover
ISBN: 978-0-387-25921-5