




CCD Analogy




Exposure finished, buckets now contain samples of rain.




Conveyor belt starts turning and transfers buckets. Rain collected on the vertical conveyor is tipped into buckets on the horizontal conveyor.





Vertical conveyor stops. Horizontal conveyor starts up and tips each bucket in turn into the measuring cylinder .





After each bucket has been measured, the measuring cylinder is emptied , ready for the next bucket load.



















A new set of empty buckets is set up on the horizontal conveyor and the process is repeated.




Structure of a CCD

The image area of the CCD is positioned at the focal plane of the telescope. An image then builds up that consists of a pattern of electric charge. At the end of the exposure this pattern is then transferred, pixel at a time, by way of the serial register to the on-chip amplifier. Electrical connections are made to the outside world via a series of bond pads and thin gold wires positioned around the chip periphery.




CCDs are manufactured on silicon wafers using the same photo-lithographic techniques used to manufacture computer chips. Scientific CCDs are very big, only a few can be fit onto a wafer. This is one reason that they are so costly. The photo below shows a silicon wafer with three large CCDs and assorted smaller devices.




Spectral Sensitivity of CCDs and the Atmosphere

The graph below shows the transmission of the atmosphere when looking at objects at the zenith. The atmosphere absorbs strongly below about 330nm, in the near ultraviolet part of the spectrum. An ideal CCD should have a good sensitivity from 330nm to approximately 1000nm, at which point silicon becomes transparent and therefore insensitive.



Over the last 25 years of development, the sensitivity of CCDs has improved enormously, to the point where almost all of the incident photons across the visible spectrum are detected. CCD sensitivity has been improved using two main techniques : back-side illumination and the use of anti-reflection coatings.

Thick Front-side Illuminated CCD




These are cheap to produce using conventional wafer fabrication techniques. They are used in consumer imaging applications. Even though not all the photons are detected, these devices are still more sensitive than photographic film.

They have a low Quantum Efficiency due to the reflection and absorption of light in the surface electrodes. Very poor blue response. The electrode structure prevents the use of an Anti-reflective coating that would otherwise boost performance.

Anti-Reflection Coatings

Silicon has a very high Refractive Index n . This means that photons are strongly reflected from its surface.




Fraction of photons reflected at the interface between two mediums of differing refractive indices

$$= \left[\frac{n_t - n_i}{n_t + n_i} \right]^2$$

n of air or vacuum is 1.0, glass is 1.46, water is 1.33, Silicon is 3.6. Using the above equation we can show that window glass in air reflects 3.5% and silicon in air reflects 32%. Unless we take steps to eliminate this reflected portion, then a silicon CCD will at best only detect 2 out of every 3 photons.

The solution is to deposit a thin layer of a transparent dielectric material on the surface of the CCD. The refractive index of this material should be between that of silicon and air, and it should have an optical thickness = $1/4$ wavelength of light. The question now is what wavelength should we choose, since we are interested in a wide range of colors.

With an Anti-reflective coating we now have three mediums to consider :




The reflected portion is now reduced to :

$$\left[\frac{n_t \times n_i - n_s^2}{n_t \times n_i + n_s^2} \right]^2$$

In the case where $n_s^2 = n_t$ the reflectivity actually falls to zero! For silicon we require a material with $n = 1.9$, fortunately such a material exists, it is Hafnium Dioxide. It is regularly used to coat astronomical CCDs.

Thinned Back-side Illuminated CCD




The silicon is chemically etched and polished down to a thickness of about 15-40 microns. Light enters from the rear and so the electrodes do not obstruct the photons. The QE can approach 100% .

These are very expensive to produce since the thinning is a non-standard process that reduces the chip yield. These thinned CCDs become transparent to near infra-red light and the red response is poor. Response can be boosted by the application of an anti-reflective coating on the thinned rear-side. These coatings do not work so well for thick CCDs due to the surface bumps created by the surface electrodes.


Astronomical CCDs are Thinned and Backside Illuminated.

Quantum Efficiency Comparison

The graph below compares the quantum efficiency of a thick front-side illuminated CCD and a thin back-side illuminated CCD.




Absorption length in Si




- 300K
- ⋯ 173K
- - - 77K

Indirect-gap semiconductors



Germanium




Silicon

Silicon: $h\nu > 2.5$ eV ($\lambda \sim 500$ nm): Intense absorption via direct gap

$h\nu < 2.5$ eV: Less efficient absorption requiring phonons for momentum (wave vector k) conservation ($h/\lambda \ll h/a$ where a is the lattice constant and λ is the photon wavelength)

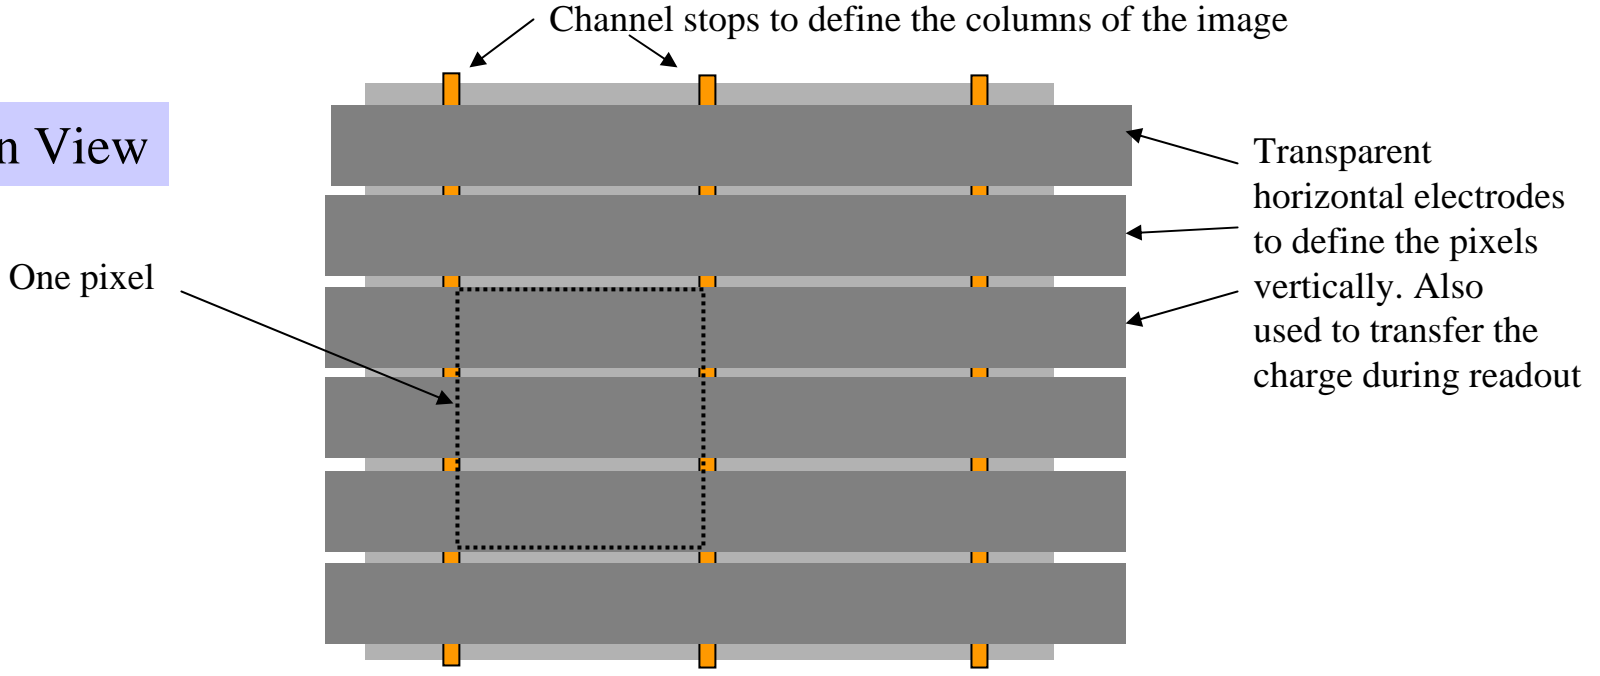
Quantum Efficiency of Back-Illuminated CCDs




Internal structure of a CCD

A small section (a few pixels) looking down onto the image area of a CCD. This pattern is repeated.

Plan View




Cross section



Every third electrode is connected together. Bus wires running down the edge of the chip make the connection. The channel stops are formed from high concentrations of Boron in the silicon.

Below the image area (the area containing the horizontal electrodes) is the 'Serial register' . This also consists of a group of small surface electrodes. There are three electrodes for every column of the image area




Cross section of serial register



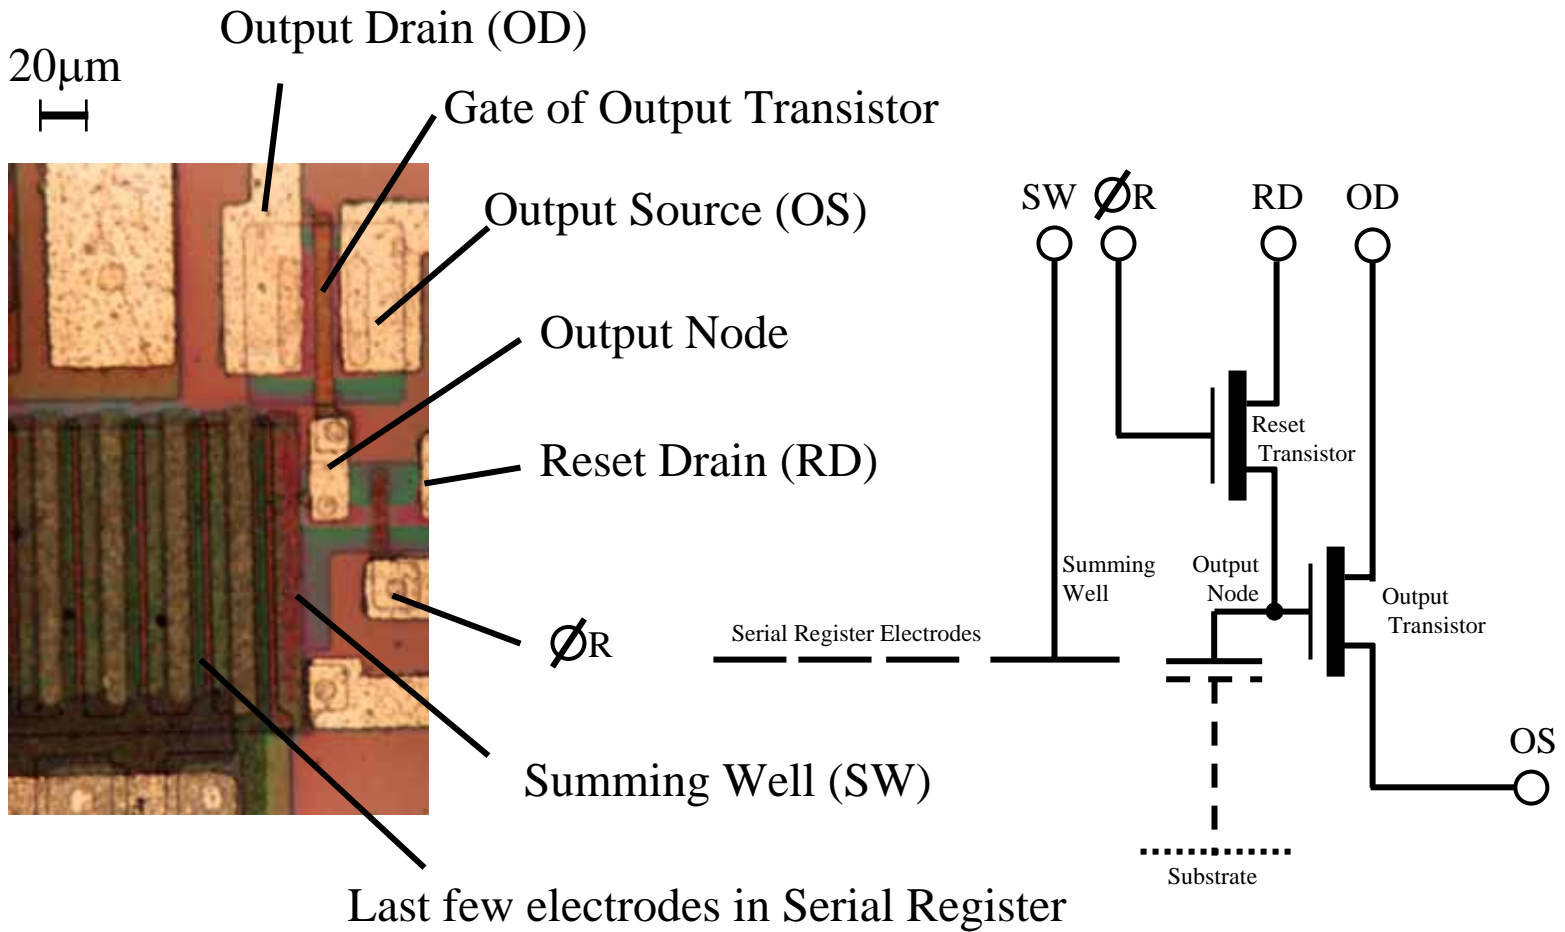
Once again every third electrode is in the serial register connected together.

Photomicrograph of a corner of a CCD




The serial register is bent double to move the output amplifier away from the edge of the chip. This is useful if the CCD is to be used as part of a mosaic. The arrows indicate how charge is transferred through the device.

Photomicrograph of the on-chip amplifier of a CCD and its circuit diagram




Electric Field in a CCD

The n-type layer contains an excess of electrons that diffuse into the p-layer. The p-layer contains an excess of holes that diffuse into the n-layer. This structure is identical to that of a diode junction. This diffusion creates a charge imbalance and induces an internal electric field. The electric potential reaches a maximum just inside the n-layer, and it is here that any photo-generated electrons will collect. All science CCDs have this junction structure, known as a 'Buried Channel'. It has the advantage of keeping the photo-electrons confined away from the surface of the CCD where they could become trapped. It also reduces the amount of thermally generated noise (dark current).




Cross section through the thickness of the CCD

During integration (exposure) of the image, one of the electrodes in each pixel is held at a positive potential. This further increases the potential in the silicon below that electrode and it is here that the photoelectrons are accumulated. The neighboring electrodes, with their lower potentials, act as potential barriers that define the vertical boundaries of the pixel. The horizontal boundaries are defined by the channel stops.



Deep Depletion CCDs

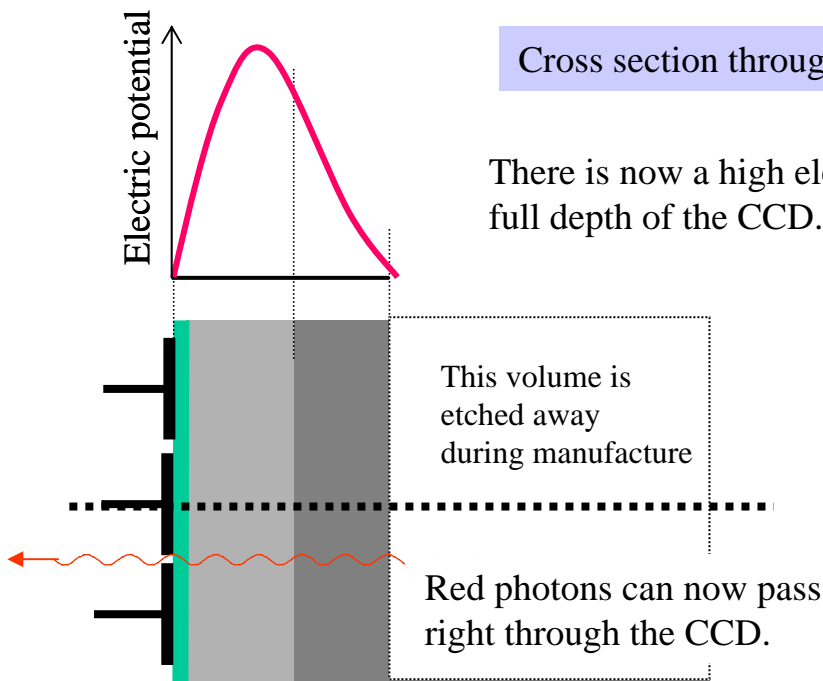
The electric field structure in a CCD defines to a large degree its Quantum Efficiency (QE). Consider first a thick front-side illuminated CCD, which has a poor QE.



Any photo-electrons created in the region of low electric field stand a much higher chance of recombination and loss. There is only a weak external field to sweep apart the photo-electron and the hole it leaves behind.

Deep Depletion CCDs

In a thinned CCD, the field free region is simply etched away.



Cross section through a thinned CCD

There is now a high electric field throughout the full depth of the CCD.

This volume is etched away during manufacture


Red photons can now pass right through the CCD.

Problem : Thinned CCDs may have good blue response but they become transparent at longer wavelengths; the red response suffers.

Photo-electrons created anywhere throughout the depth of the device will now be detected. Thinning is normally essential with backside illuminated CCDs if good blue response is required. Most blue photo-electrons are created within a few nanometers of the surface and if this region is field free, there will be no blue response.

Deep Depletion CCDs

Ideally we require all the benefits of a thinned CCD plus an improved red response. The solution is to use a CCD with an intermediate thickness of about $40\mu\text{m}$ constructed from Hi-Resistivity silicon. The increased thickness makes the device opaque to red photons. The use of Hi-Resistivity silicon means that there are no field free regions despite the greater thickness.



Cross section through a Deep Depletion CCD

Problem :

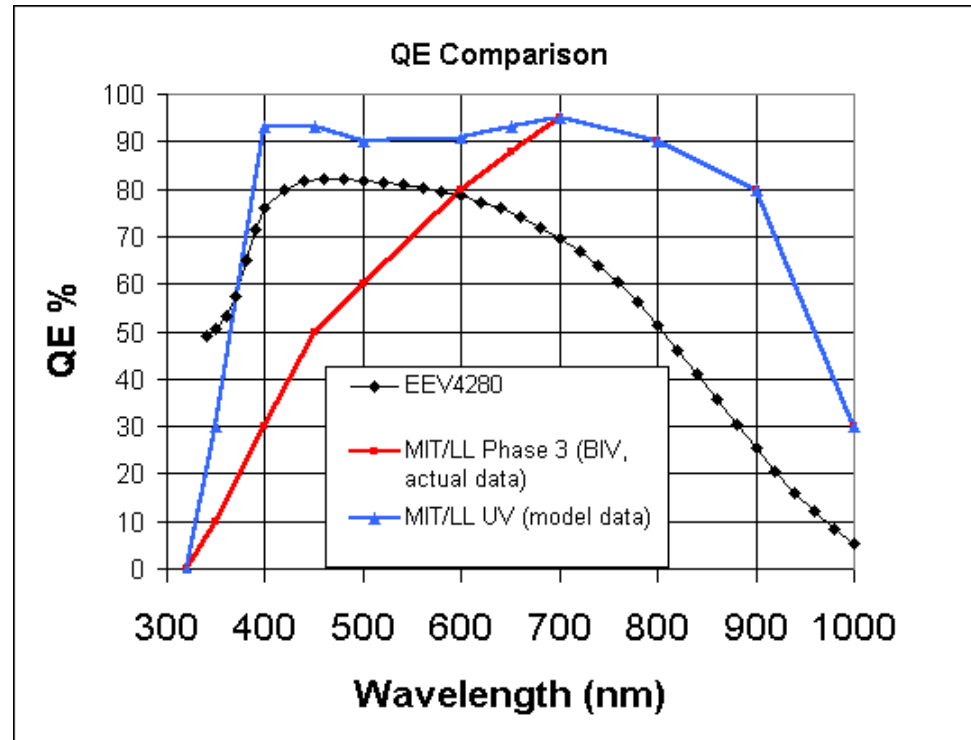
Hi resistivity silicon contains much lower impurity levels than normal. Very few wafer fabrication factories commonly use this material and deep depletion CCDs have to be designed and made to order.

Red photons are now absorbed in the thicker bulk of the device.

There is now a high electric field throughout the full depth of the CCD. CCDs manufactured in this way are known as Deep Depletion CCDs. The name implies that the region of high electric field, also known as the 'depletion zone' extends deeply into the device.

Deep Depletion CCDs


The graph below shows the improved QE response available from a deep depletion CCD.



The black curve represents a normal thinned backside illuminated CCD. The Red curve is actual data from a deep depletion chip manufactured by MIT Lincoln Labs. This latter chip is still under development. The blue curve suggests what QE improvements could eventually be realised in the blue end of the spectrum once the process has been perfected.

Another problem commonly encountered with thinned CCDs is 'fringing'. This is greatly reduced in deep depletion CCDs. Fringing is caused by multiple reflections inside the CCD. At longer wavelengths, where thinned chips start to become transparent, light can penetrate through and be reflected from the rear surface. It then interferes with light entering for the first time. This can give rise to constructive and destructive interference and a series of fringes where there are minor differences in the chip thickness.

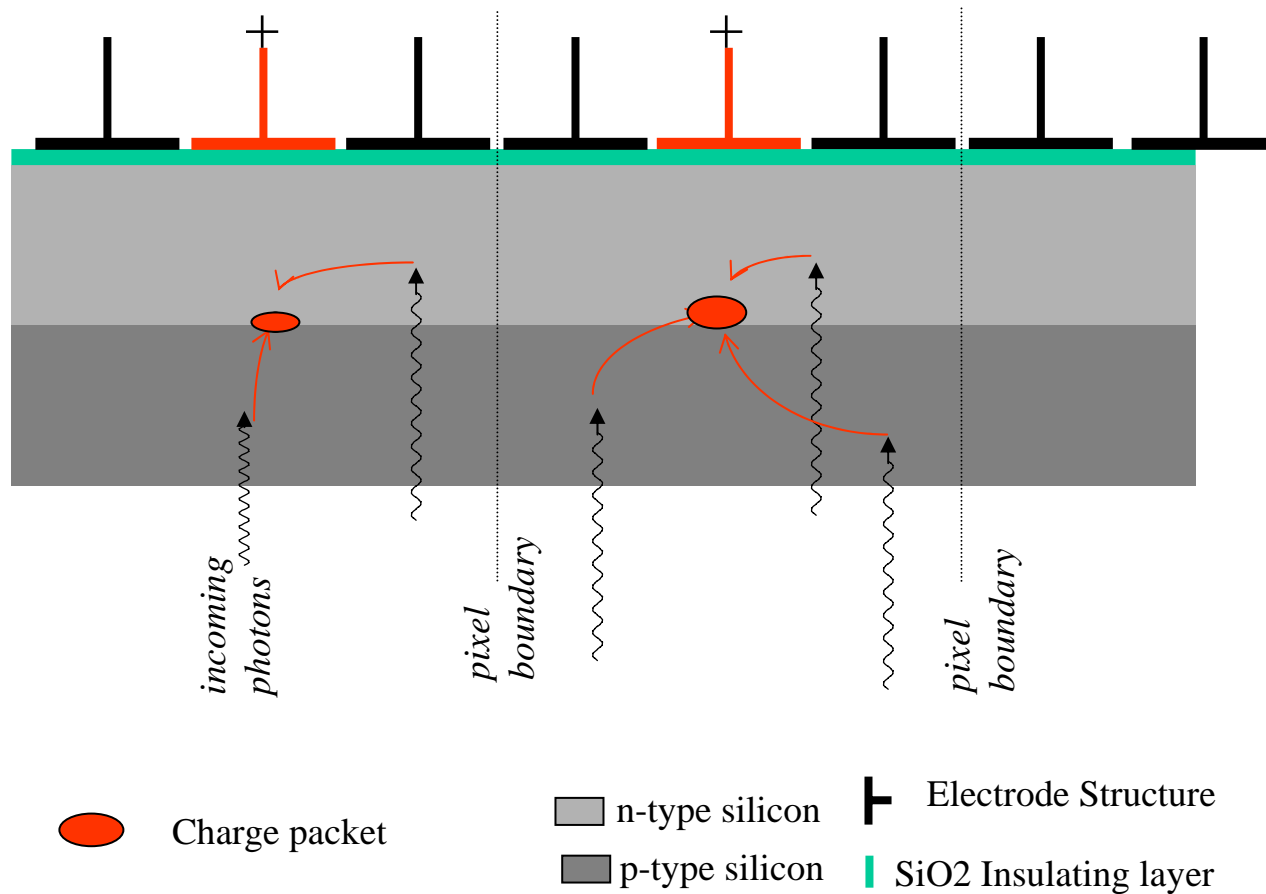
The image below shows some fringes from a thinned CCD



For spectroscopic applications, fringing can render some thinned CCDs unusable, even those that have quite respectable QEs in the red. Thicker deep depletion CCDs, which have a much lower degree of internal reflection and much lower fringing are preferred by astronomers for spectroscopy.

Charge Collection in a CCD


Photons entering the CCD create electron-hole pairs. The electrons are then attracted towards the most positive potential in the device where they create 'charge packets'. Each packet corresponds to one pixel





Charge Transfer in a CCD

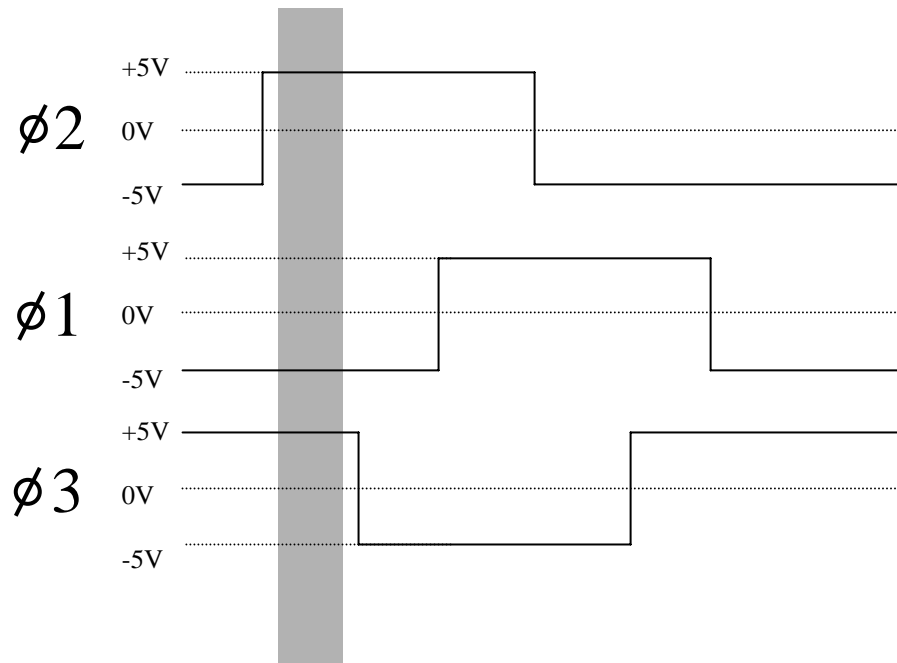

In the following few slides, the implementation of the 'conveyor belts' as actual electronic structures is explained.

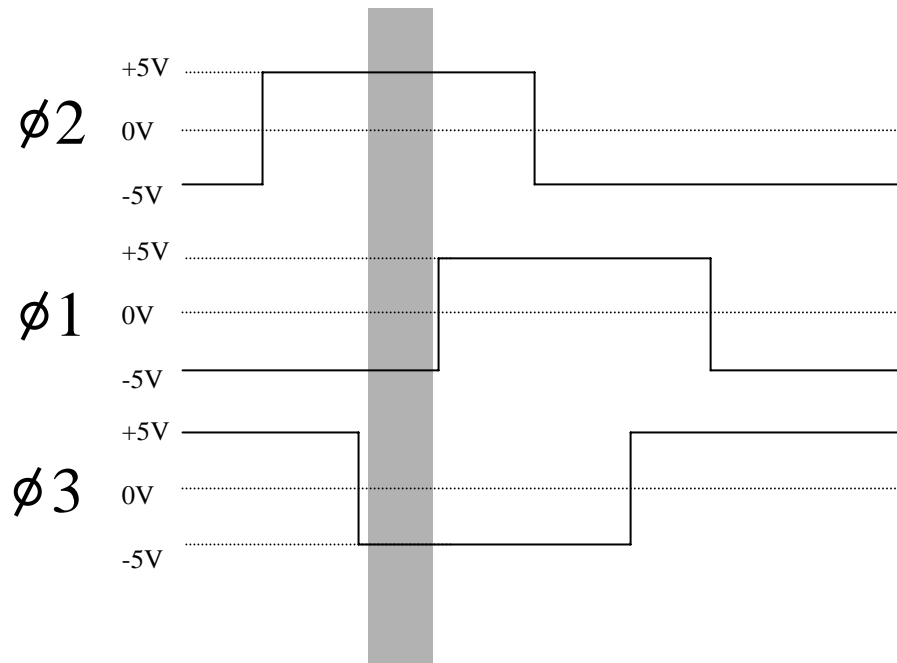

The charge is moved along these 'conveyor belts' by modulating the voltages on the electrodes positioned on the surface of the CCD. In the following illustrations, electrodes color coded red are held at a positive potential, those colored black are held at a negative potential.

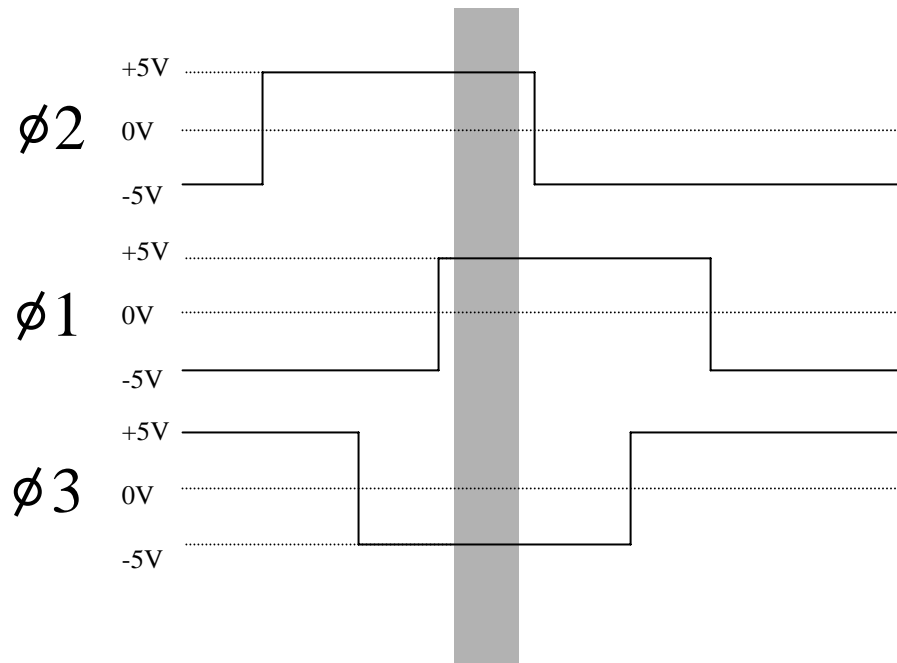







Time-slice shown in diagram













Charge packet from subsequent pixel enters from left as first pixel exits to the right.


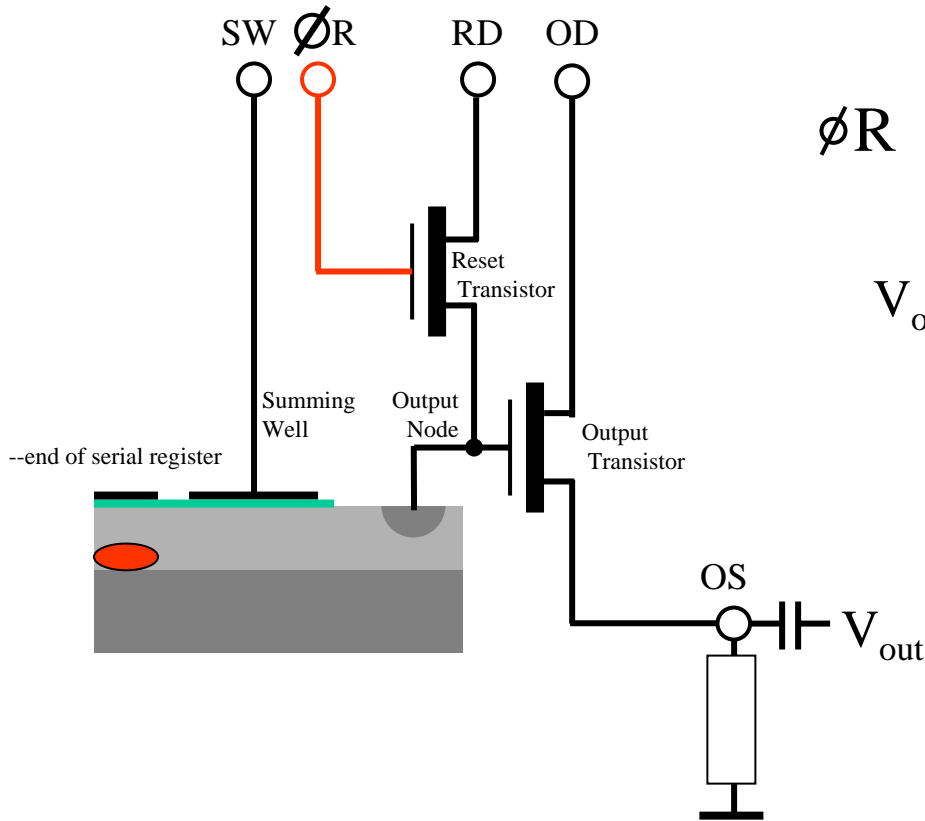




On-Chip Amplifier

The on-chip amplifier measures each charge packet as it exits the end of the serial register.



RD and OD are held at constant voltages



(The graphs above show the signal waveforms)


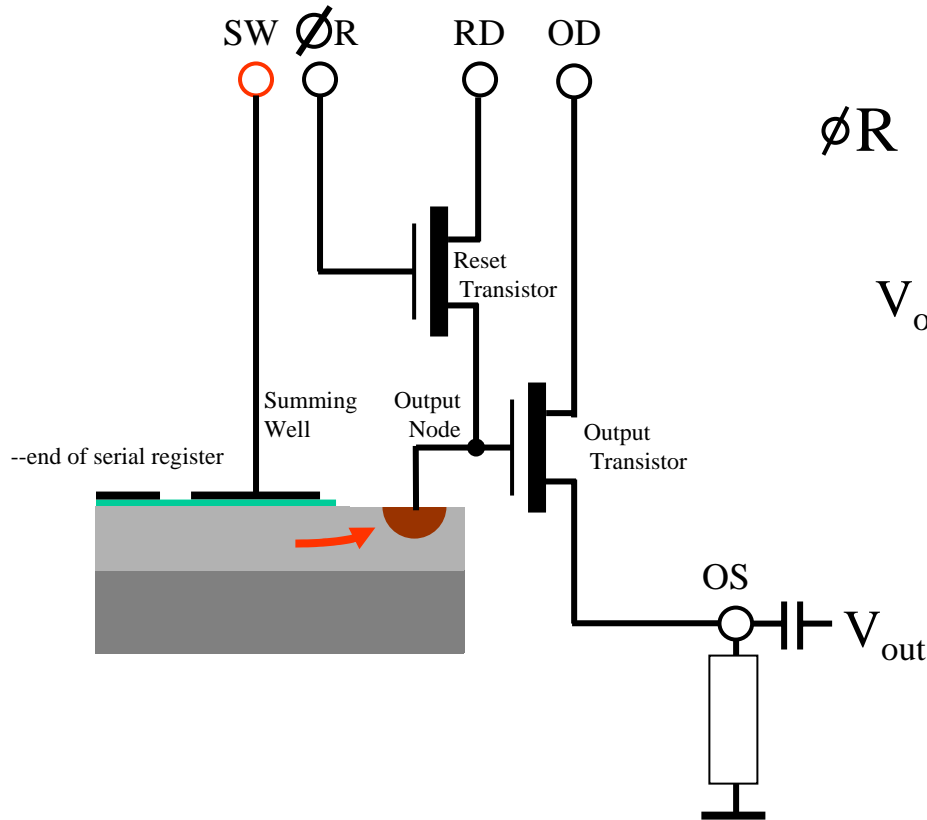
The measurement process begins with a reset of the 'reset node'. This removes the charge remaining from the previous pixel. The reset node is in fact a tiny capacitance ($< 0.1\text{pF}$)

The charge is then transferred onto the Summing Well. V_{out} is now at the 'Reference level'



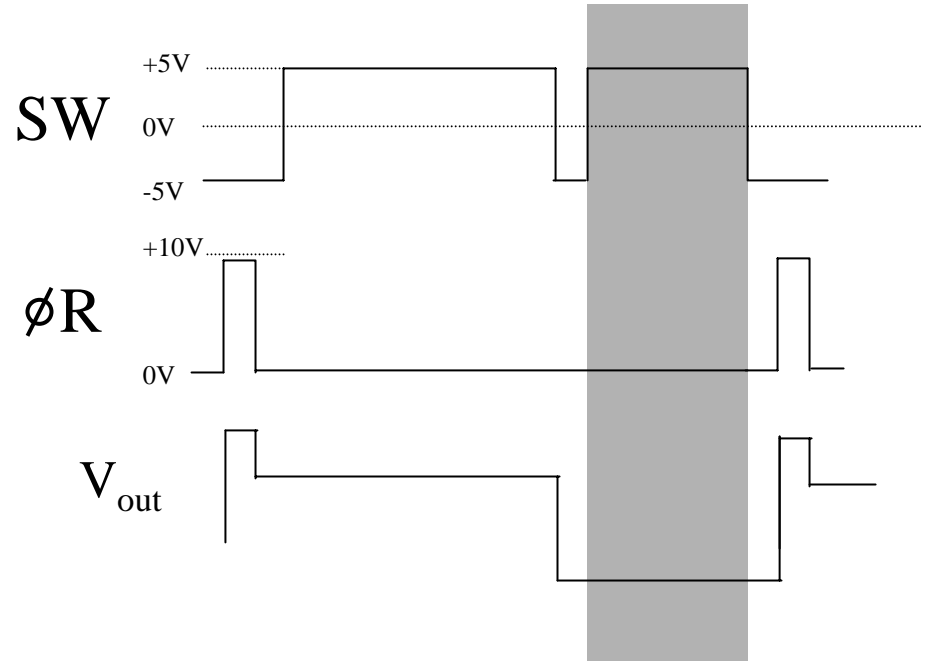

There is now a wait of up to a few tens of microseconds while external circuitry measures this 'reference' level.

The charge is then transferred onto the output node. V_{out} now steps down to the 'Signal level'



This action is known as the 'charge dump'
The voltage step in V_{out} is as much as several μV for each electron contained in the charge packet.

V_{out} is now sampled by external circuitry for up to a few tens of microseconds.

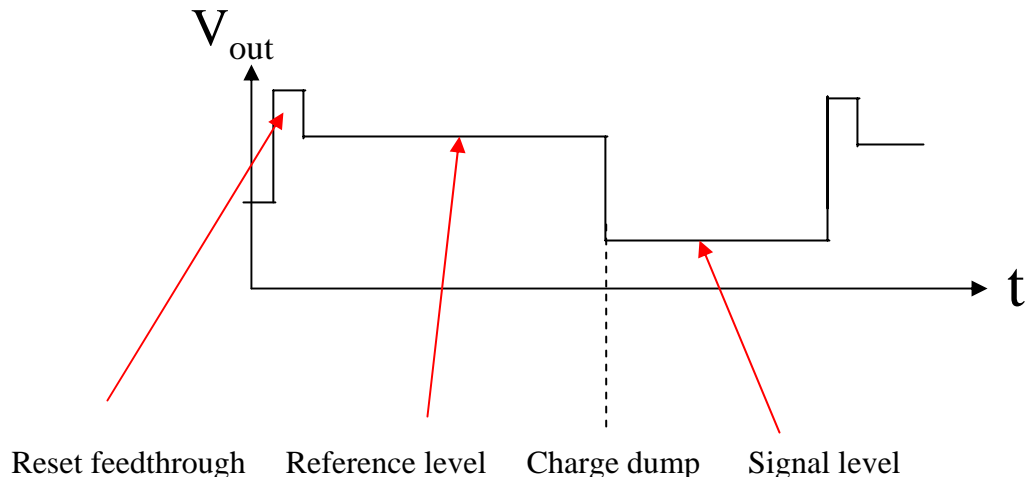


The sample level minus the reference level will be proportional to the size of the input charge packet.

Correlated Double Sampler

The video waveform output by a CCD is at a fairly low level : every photo-electron in a pixel charge packet will produce a few micro-volts of signal. Additionally, the waveform is complex and precise timing is required to make sure that the correct parts are amplified and measured.


The CCD output video waveform is shown below for the period of one pixel measurement




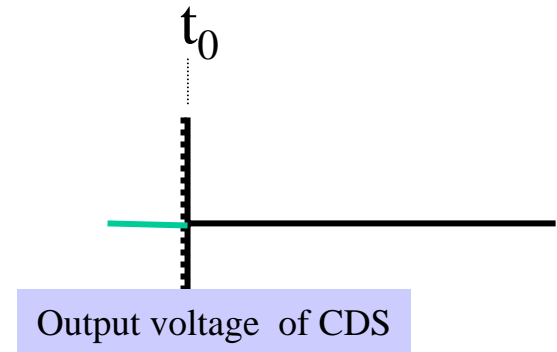

The video processor must measure, without introducing any additional noise, the Reference level and the Signal level. The first is then subtracted from the second to yield the output signal voltage proportional to the number of photo-electrons in the pixel under measurement. The best way to perform this processing is to use a 'Correlated Double Sampler'.

Correlated Double Sampler (CDS)




The CDS design is shown schematically below. The CDS processes the video waveform and outputs a digital number proportional to the size of the charge packet contained in the pixel being read. There should only be a short cable length between CCD and CDS to minimize noise. The CDS minimizes the read noise of the CCD by eliminating 'reset noise'. The CDS contains a high speed analog processor containing computer controlled switches. Its output feeds into an Analog to Digital Converter (ADC).






The CDS starts work once the pixel charge packet is in the CCD summing well and the CCD reset pulse has just finished. At point t_0 the CCD wave-form is still affected by the reset pulse and so the CDS remains disconnected from the CCD to prevent this disturbing the video processor.





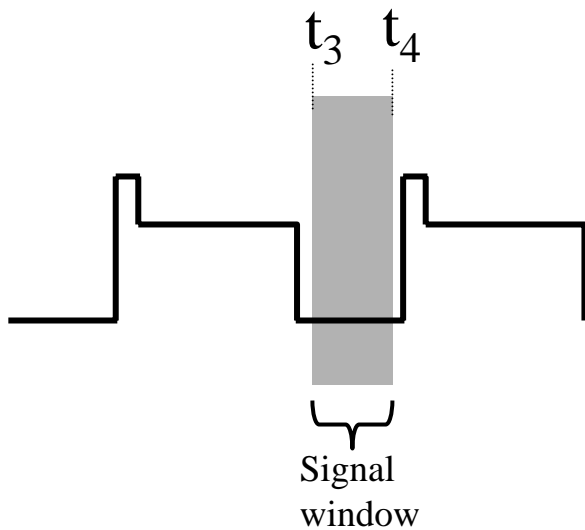
Between t_1 and t_2 the CDS is connected and the 'Reference' part of the waveform is sampled. Simultaneously the integrator reset switch is opened and the output starts to ramp down linearly.



Between t_2 and t_3 the 'charge dump' occurs in the CCD. The CCD output steps negatively by an amount proportional to the charge contained in the pixel. During this time the CDS is disconnected.



Between t_3 and t_4 the CDS is reconnected and the 'signal' part of the wave-form is sampled. The input to the integrator is also 'polarity switched' so that the CDS output starts to ramp-up linearly. The width of the signal and sample windows must be the same. For Scientific CCDs this can be anything between 1 and 20 microseconds. Longer widths generally give lower noise but of course increase the read-out time.



The CDS is then once again disconnected and its output digitized by the ADC. This number, typically a 16 bit number (with a value between 0 and 65535) is then stored in the computer memory. The CDS then starts the whole process again on the next pixel. The integrator output is first zeroed by closing the reset switch. To process each pixel can take between a fraction of a microsecond for a TV rate CCD and several tens of microseconds for a low noise scientific CCD.

This type of CDS is called a 'dual slope integrator'. A simpler type of CDS known as a 'clamp and sample' only samples the waveform once for each pixel. It works well at higher pixel rates but is noisier than the dual slope integrator at lower pixel rates.

