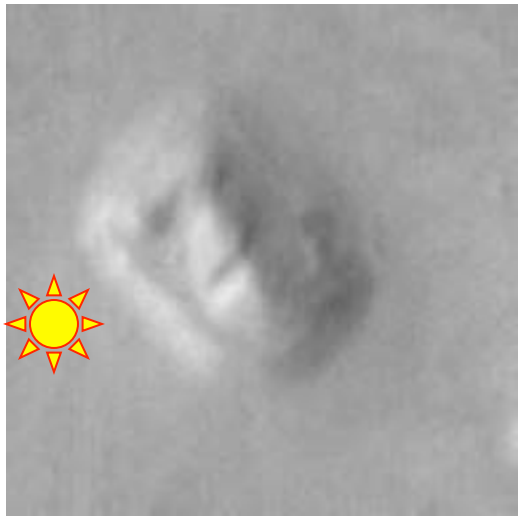

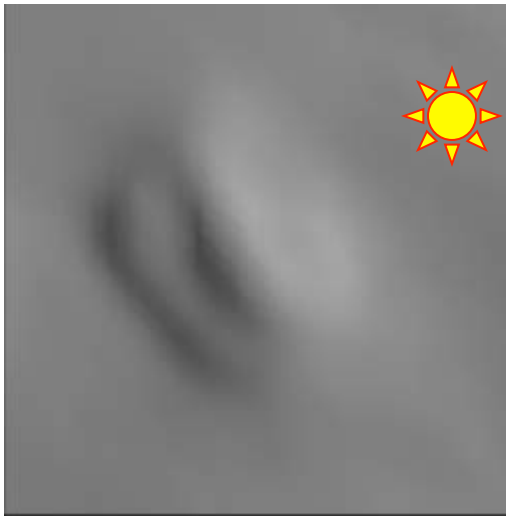

Lighting,
viewing geometry,
& resolution

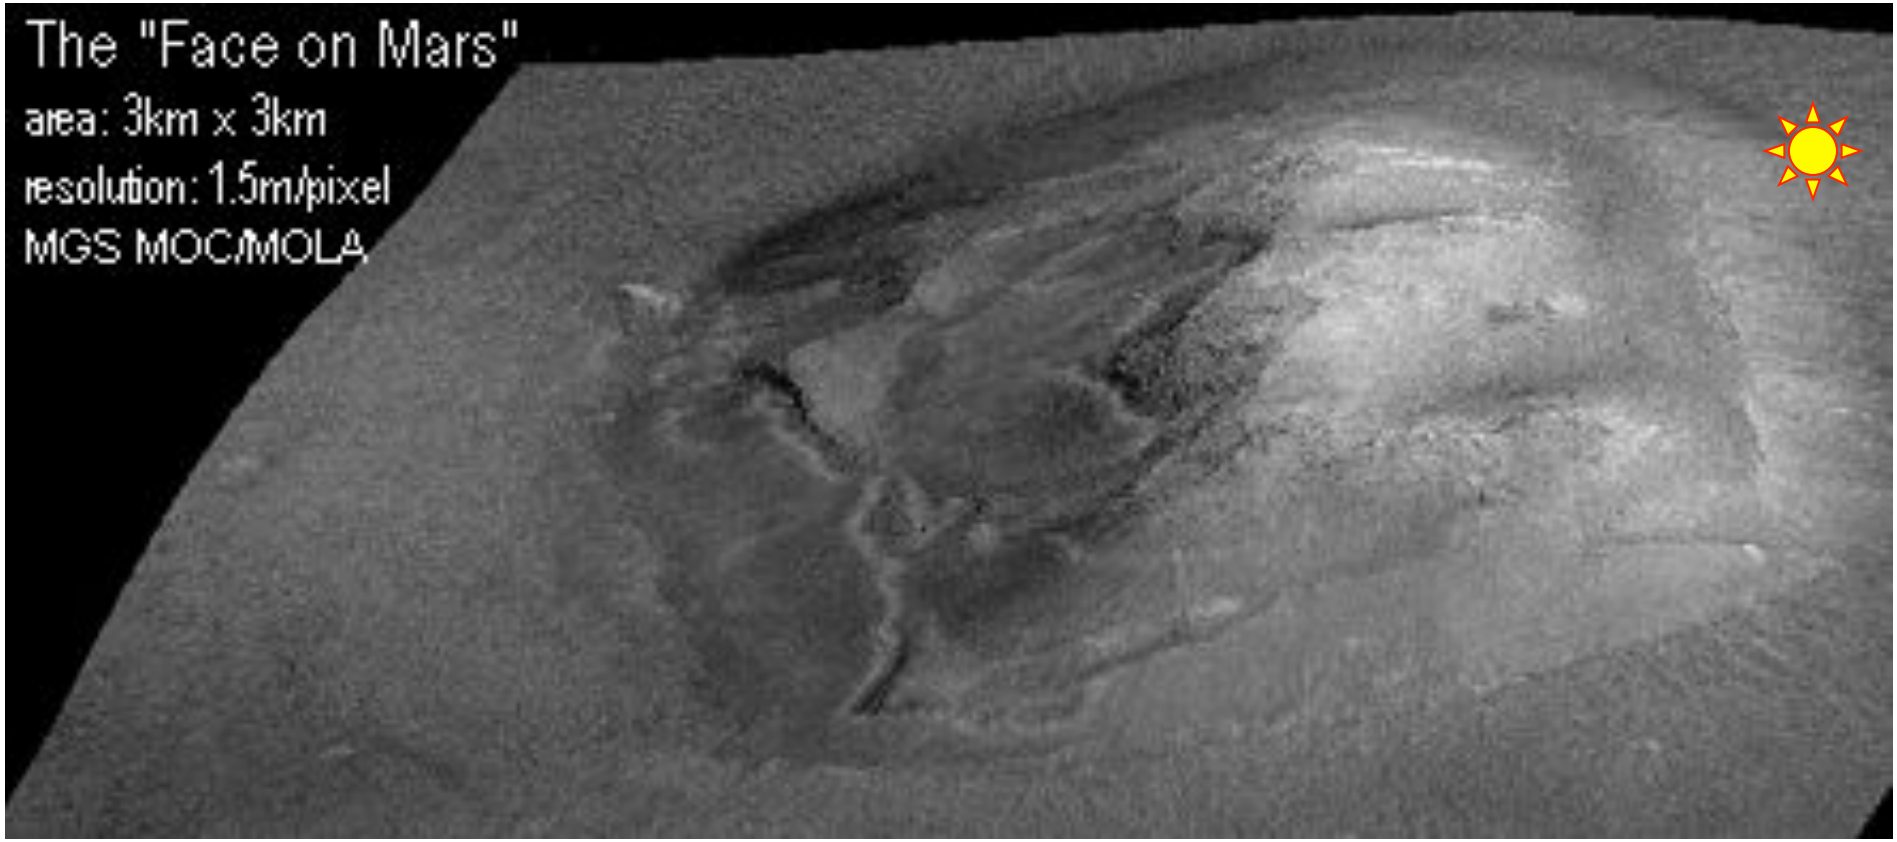

?

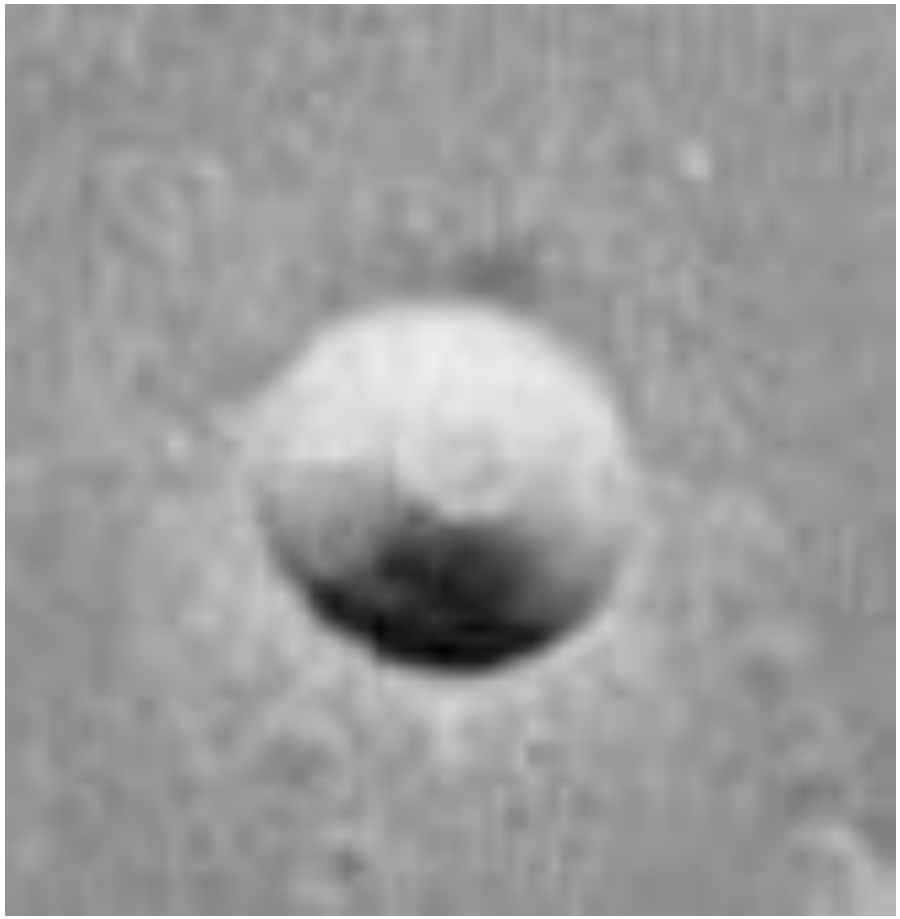

Absence of contextual clues permits ambiguity

Hill?

Or hollow?

?

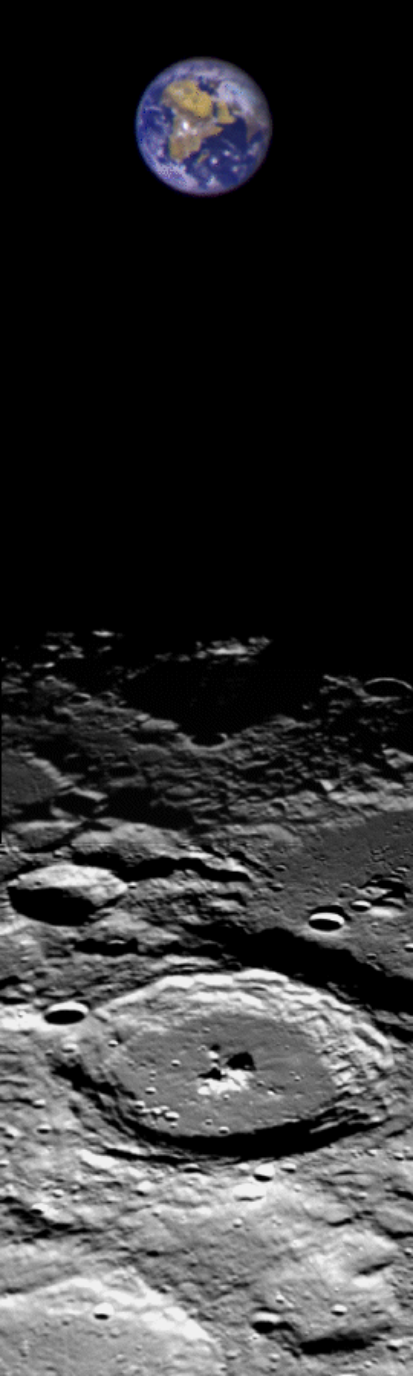

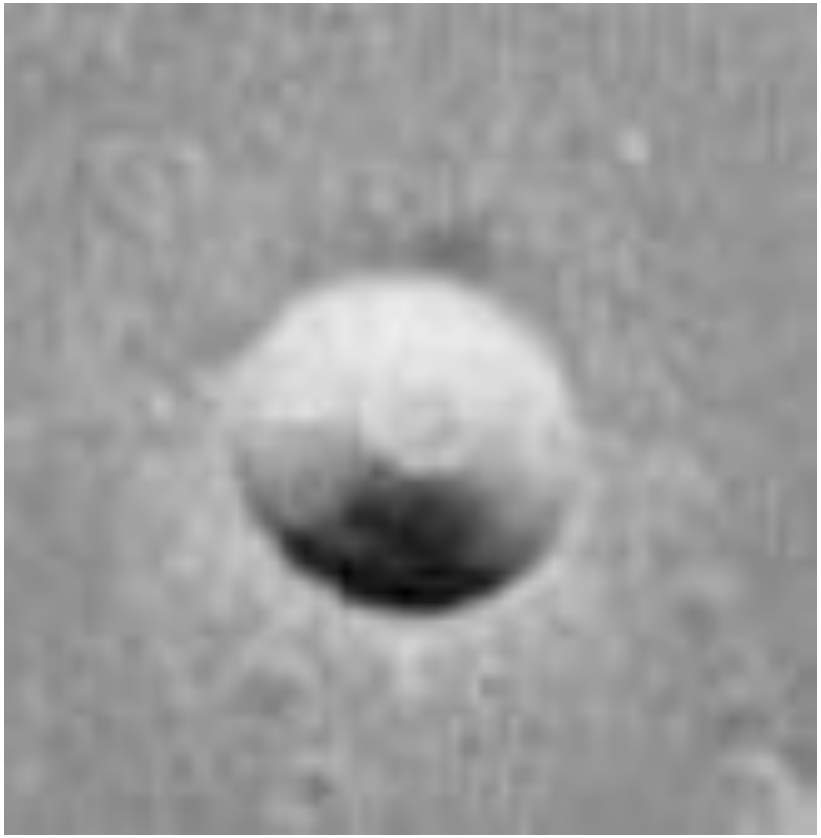

Expand the FOV for more contextual clues

Familiar scenes are interpreted more easily

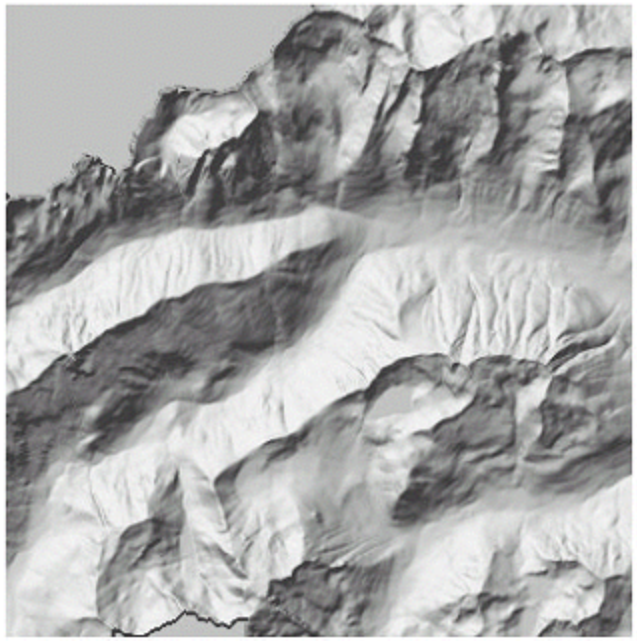

"Illuminated" DEM

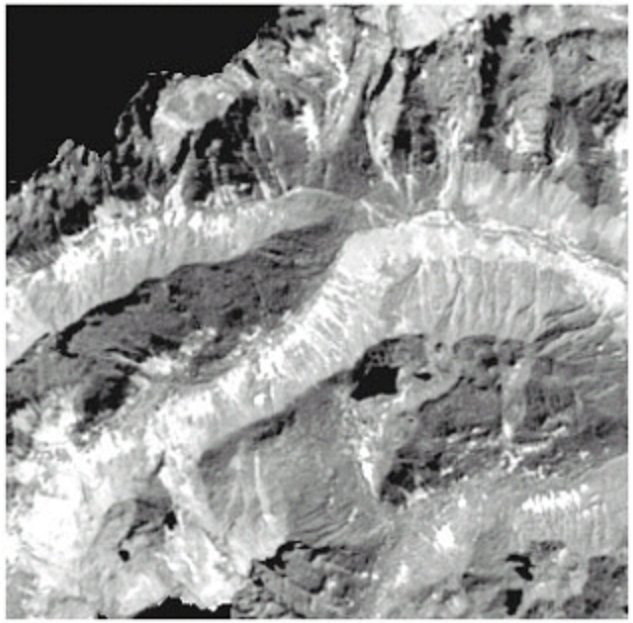

B/W image

Albedo variations, shape & lighting both contribute to B/W images

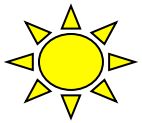

Voyager Imaging Science Subsystem

"Vidicon" video camera tube with selenium-sulfur photoconductor detector

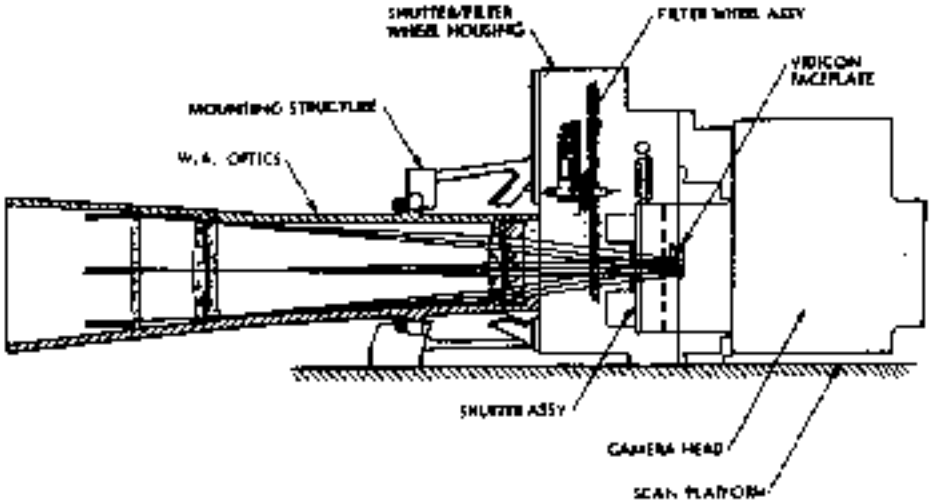

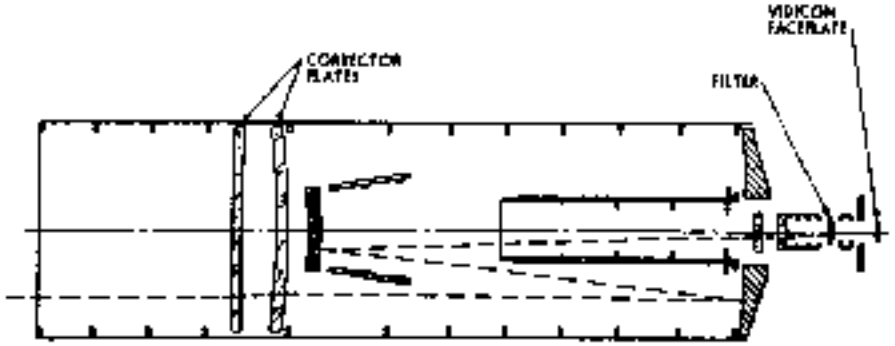

CCDs have many advantages

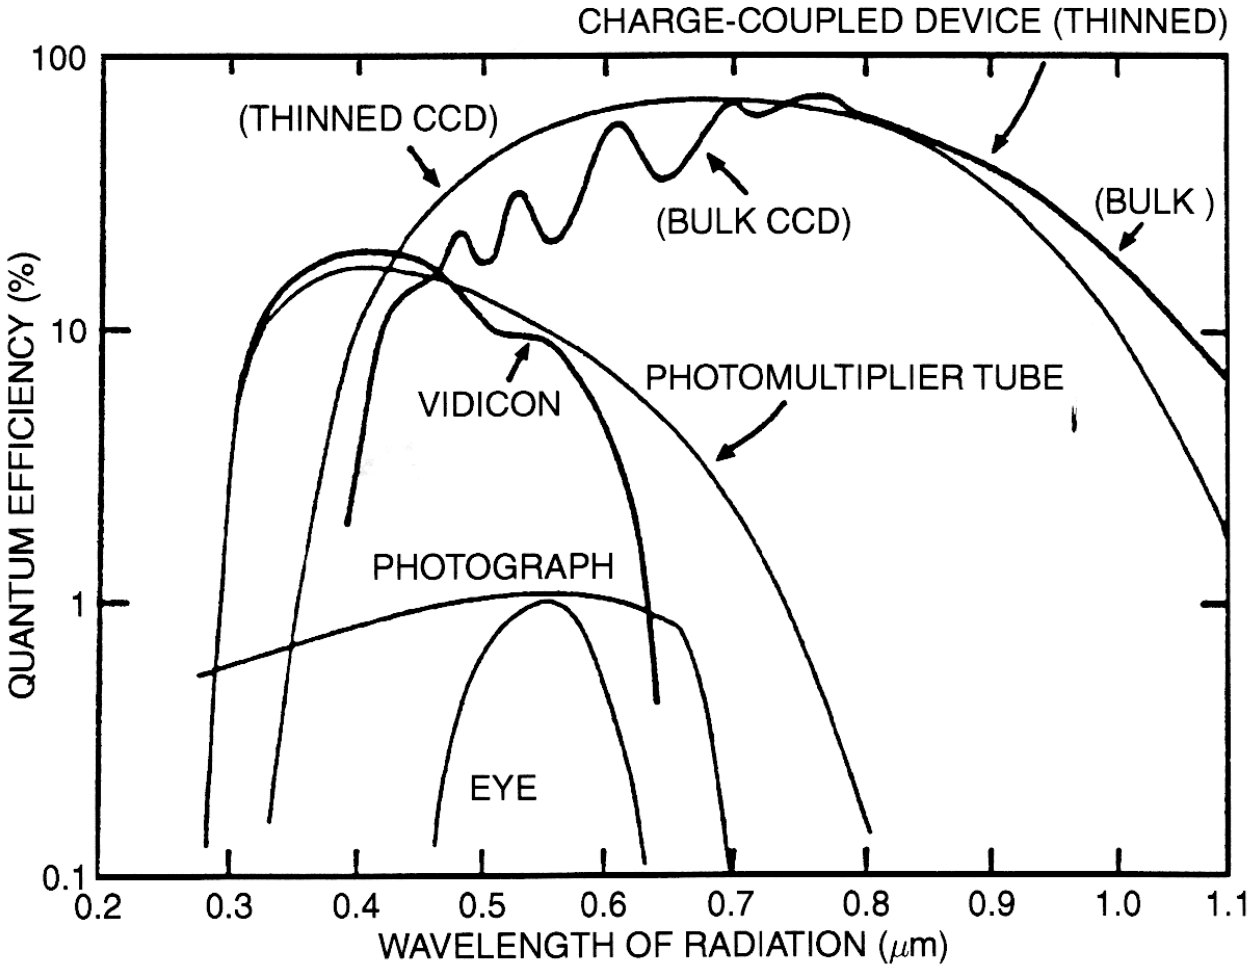

Fig. 3.2. QE curves for various devices, indicating why CCDs are a quantum leap above all previous imaging devices. The failure of CCDs at optical wavelengths shorter than about 3500 Å has been essentially eliminated via thinning or coating of the devices (see Figure 3.3).

- High QE = fraction of incoming photons that are actually detected
- Broad spectral response

CCDs: A Nobel-worthy innovation

CCDs: rows of buckets on conveyor belts

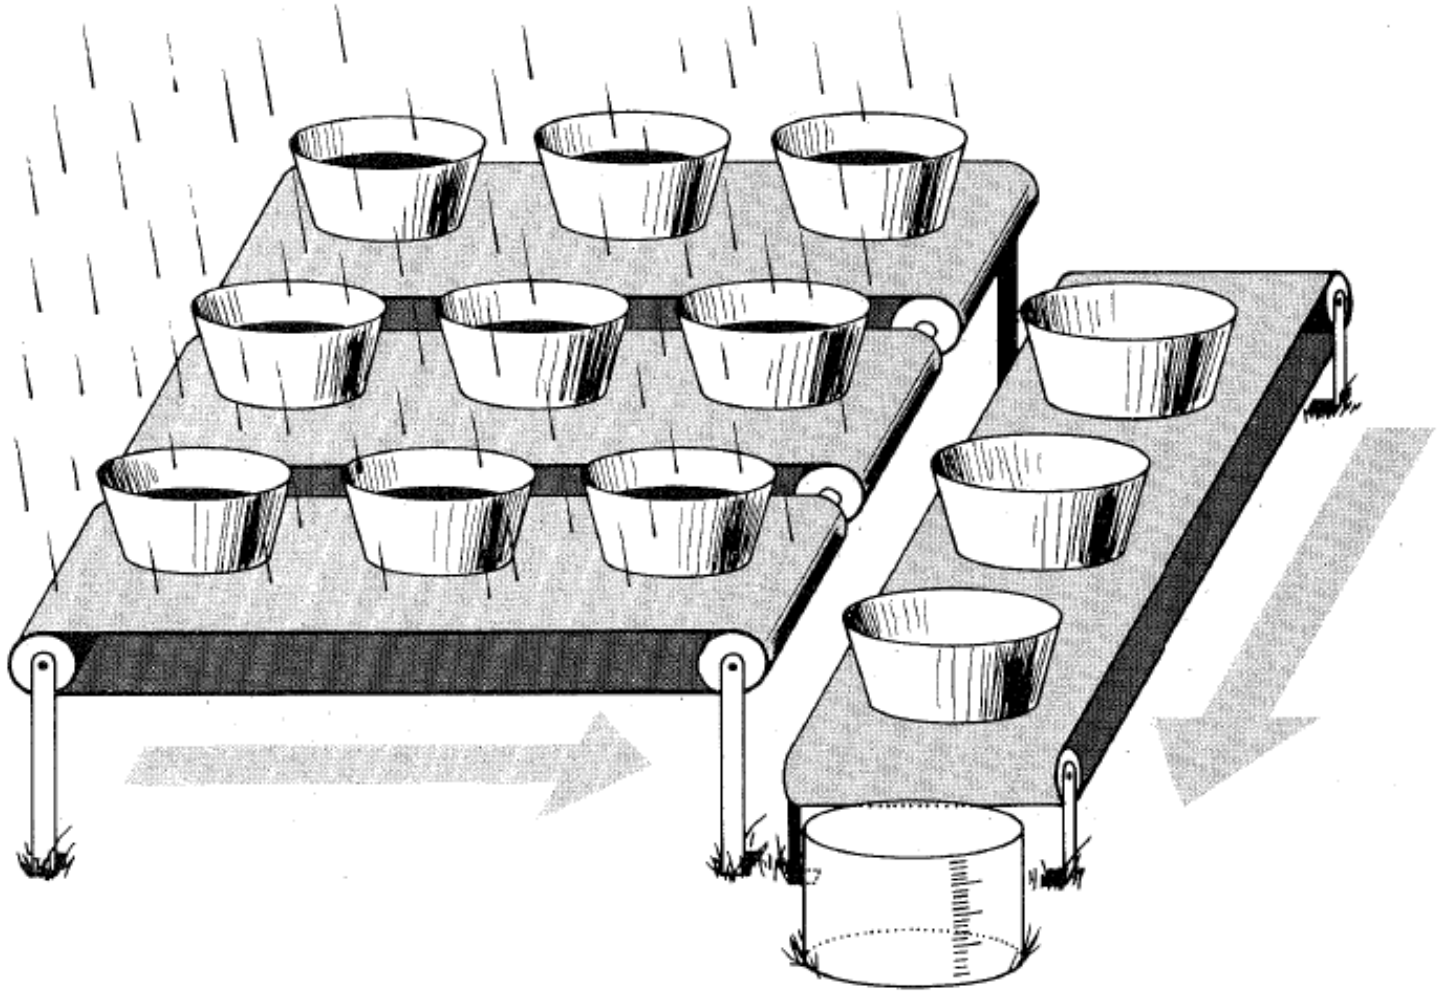

Determining the brightness distribution in a celestial object with a charge-coupled device can be likened to measuring the rainfall at different points in a field with an array of buckets. Once the rain has ceased, the buckets in each row are moved horizontally across the field on conveyor belts. As each one reaches the end of the conveyor, it is emptied into another bucket on a belt that carries it to the metering station where its contents are measured. Artwork by Steven Simpson.

CCDs: rows of buckets on conveyor belts

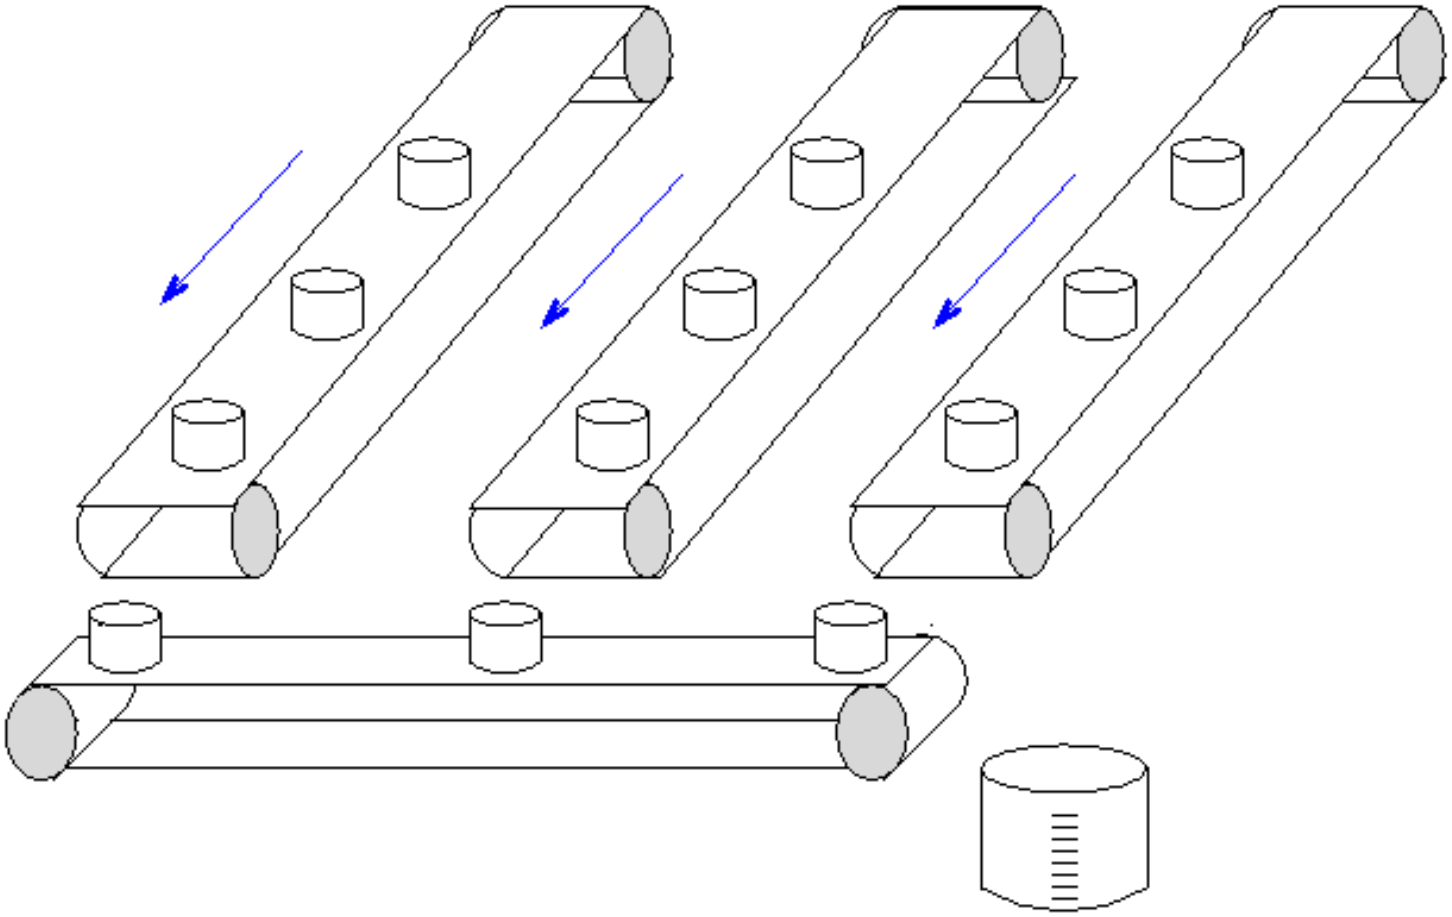

- Each “bucket” is a pixel
- They are actually (in simple CCDs) directly adjacent to each other (no space between)
- Buckets are stationary during image integration, then transfer occurs along rows
- After each pixel-width transfer, the output shift register must empty, one row at a time

CCDs: what's really going on

- Every photon ~300-1100 nm has 1 to a few times silicon's bandgap energy; absorption causes an electron to move from valence to conduction band
  ➔ *photoelectric effect*
- Freed electrons collected in "gate" structures (specifically, in 1 of 3 per pixel)
- After imaging, gate voltages are varied to shift electrons in controlled manner
- *Charge transfer efficiency* (CTE) per pixel >99.9995%!

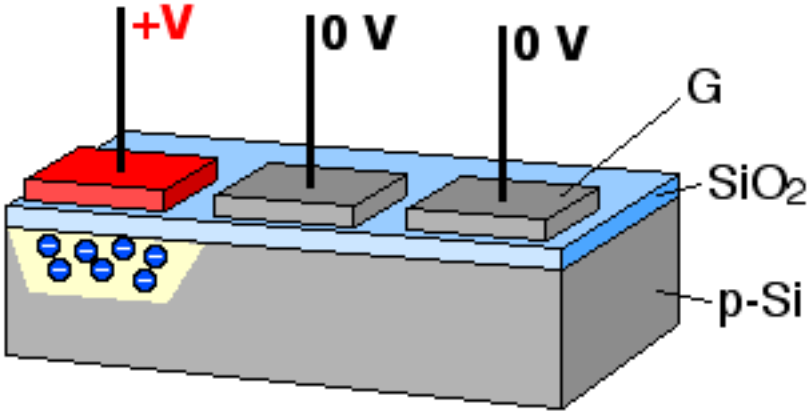

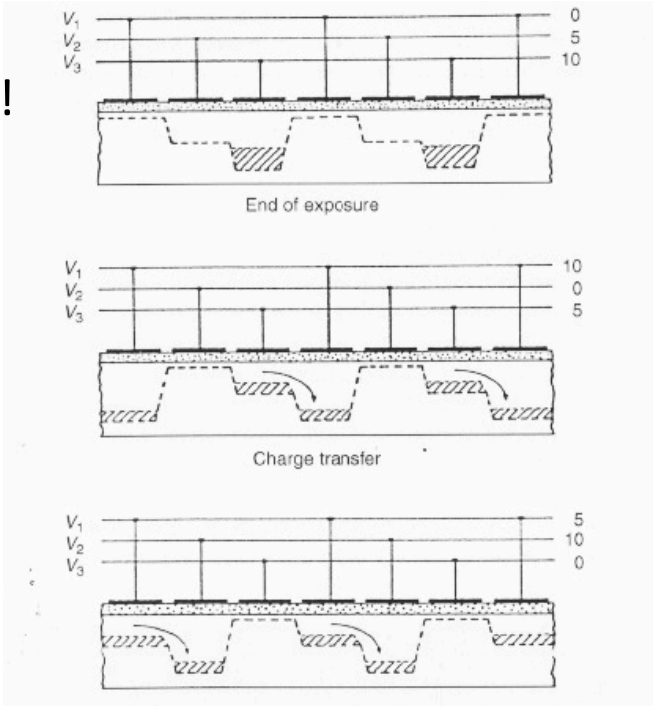

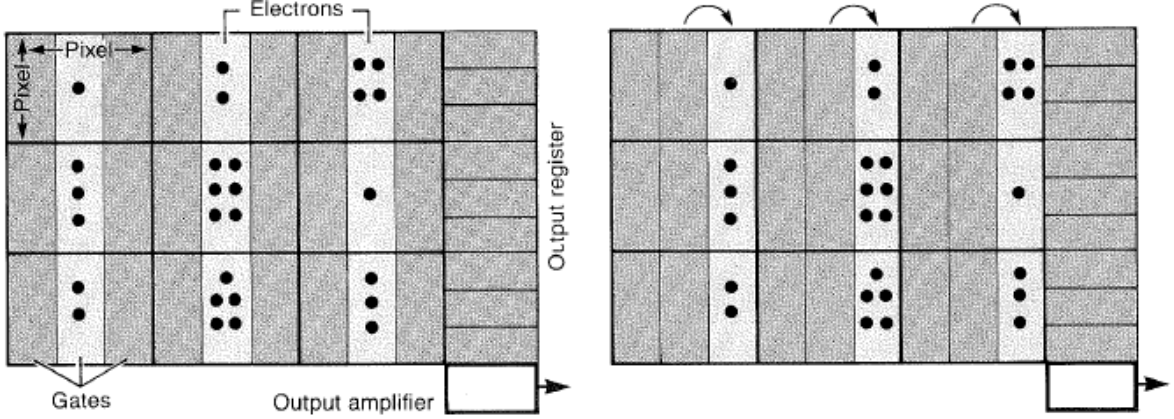

The operation of a CCD is illustrated schematically by this simplified chip consisting of nine pixels, an output register, and an amplifier. Every pixel is subdivided into three regions, or gates; each is an electrode whose voltage can be varied. *Left:* During an exposure the central gate of each pixel is "on" (yellow areas) and its neighbors are "off" (green areas). This creates "electron buckets" under the middle electrodes and barriers between adjacent pixels. *Right:* At the end of the exposure the gate voltages are changed and electrons shifted one gate to the right as new potential wells are created and old ones are destroyed.

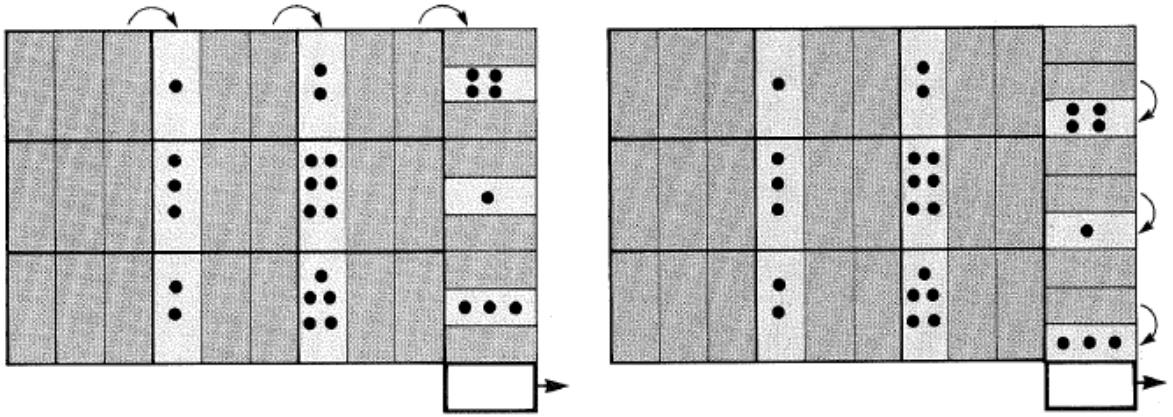

Left: As the voltages are cycled again, electrons flow from the right-most gate of one pixel to the left-most gate of its neighbor. The electrons in the right-hand pixels are transferred to the output register. *Right:* Before the pixel array can be shifted again, charge must be transferred, one pixel at a time, through the output register and amplifier. When this register has been completely emptied, another cycle of pixel-array transfers is executed. These steps continue in a systematic way until every charge packet has moved horizontally along its row, vertically down the output register, and into the amplifier, where it is measured. In a real astronomical CCD, this process can take as long as 10 seconds.

CCDs: what's really going on

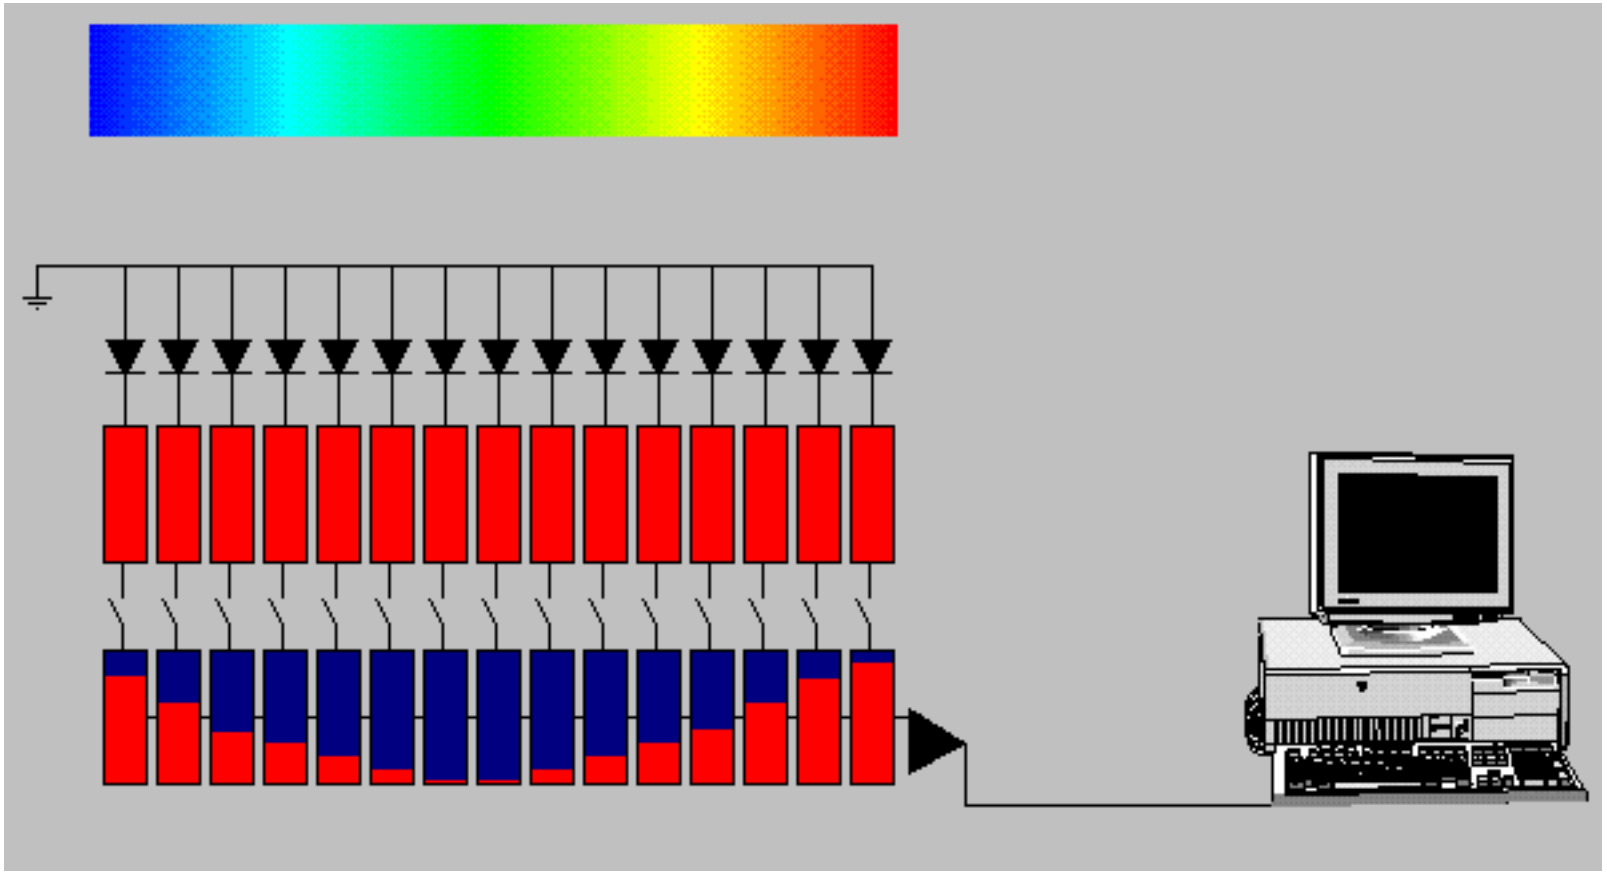

- Each pixel's accumulated charge must eventually pass through amplifier and A/D converter (converts voltages to DNs)

Image of an actual CCD

Low light level CCD (L3CCD) has extended *gain register*

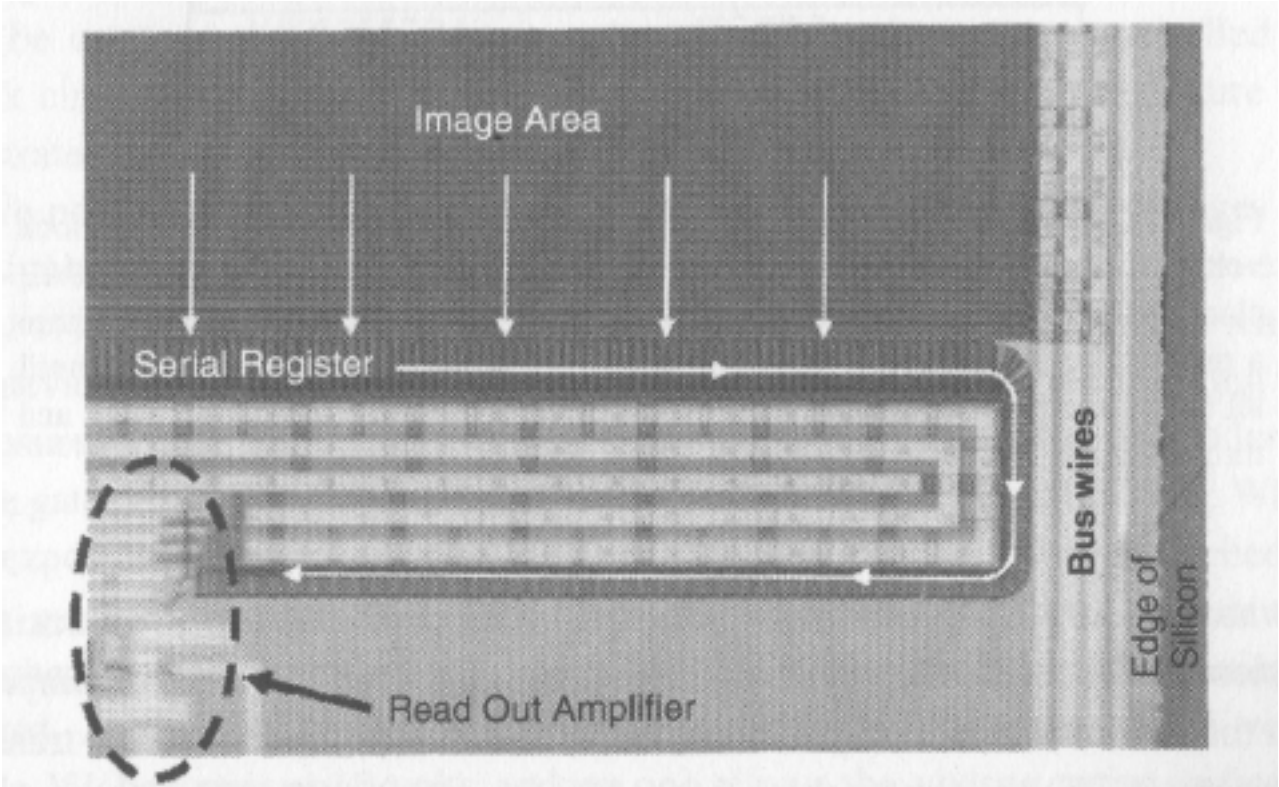

Fig. 2.3. Microphotograph of a E2V L3CCD (see Section 2.2.7) showing the image area (pixels), the serial register, and the on-chip readout amplifier. The other wiring and the bus wires are electrical connections that carry the clock signals and bias voltages to use. Added on to the normal CCD components is an extended serial register through which the readout occurs (the arrow indicates this flow) where the half after the bend is the gain register.

Case study: Mars Reconnaissance Orbiter *HiRISE*

Figure 2. HiRISE instrument inside a clean room tent at Ball Aerospace and Technology Corporation in Boulder, Colorado.

Case study: Mars Reconnaissance Orbiter *HiRISE*

Table 1. HiRISE Camera Capabilities at 300 km Altitude

Parameter	Characteristics
Ground sampling dimension	30 cm/pixel (1 μrad IFOV)
Resolution	~90 cm (3 pixels across an object)
Swath width (RED CCDs)	6 km (1.14° FOV)
Color swath width	1.2 km (0.23° FOV)
Maximum image size (pixels)	20,000 × 63,780 (14-bit data)
SNR (anywhere on Mars in the optimal season)	From 90:1 to 250:1 in RED channels with TDI 128 and full resolution
Color band passes (at half maximum of Mars- and solar-weighted spectral response; see Figure 12)	RED: 570–830 nm BG: <580 nm NIR: >790 nm
Stereo topographic precision	~25 cm vertical over ~1 m^2 areas
TDI lines	8, 32, 64, or 128
Pixel binning	none (1 × 1), 2 × 2, 3 × 3, 4 × 4, 8 × 8, 16 × 16
Bits per pixel	14, can be compressed to 8 via look-up tables (LUTs)
Compression (8-bit images only)	FELICS, compression >1.6:1

McEwen et al. (2007, JGR)

Ground sampling slightly finer than true resolution

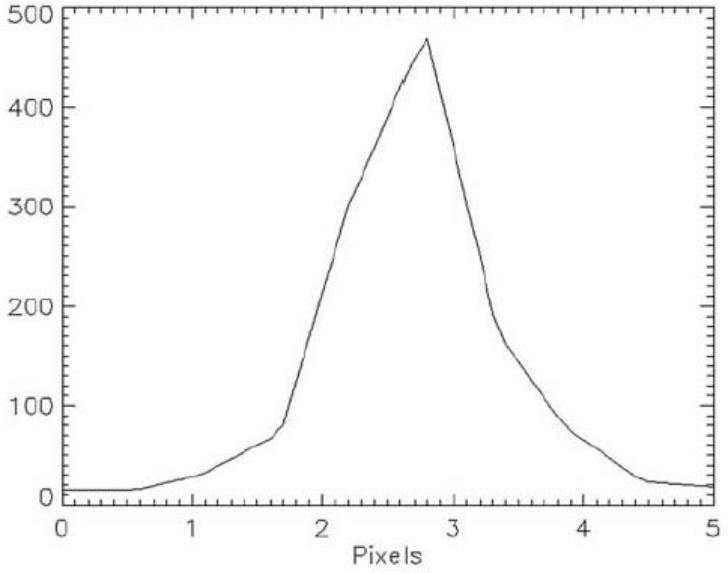

Figure 9. Typical Point-Spread Function (PSF) for a HiRISE pinhole image, acquired at Ball Aerospace with a 30-inch collimator.

Case study: Mars Reconnaissance Orbiter *HiRISE*

Table 1. HiRISE Camera Capabilities at 300 km Altitude

Parameter	Characteristics
Ground sampling dimension	30 cm/pixel (1 μrad IFOV)
Resolution	~90 cm (3 pixels across an object)
Swath width (RED CCDs)	6 km (1.14° FOV)
Color swath width	1.2 km (0.23° FOV)
Maximum image size (pixels)	20,000 × 63,780 (14-bit data)
SNR (anywhere on Mars in the optimal season)	From 90:1 to 250:1 in RED channels with TDI 128 and full resolution
Color band passes (at half maximum of Mars- and solar-weighted spectral response; see Figure 12)	RED: 570–830 nm BG: <580 nm NIR: >790 nm
Stereo topographic precision	~25 cm vertical over ~1 m^2 areas
TDI lines	8, 32, 64, or 128
Pixel binning	none (1 × 1), 2 × 2, 3 × 3, 4 × 4, 8 × 8, 16 × 16
Bits per pixel	14, can be compressed to 8 via look-up tables (LUTs)
Compression (8-bit images only)	FELICS, compression >1.6:1

Not a single 20,000-pixel detector! (for which readout would take forever);
20 separate output registers (2 per CCD) across the array

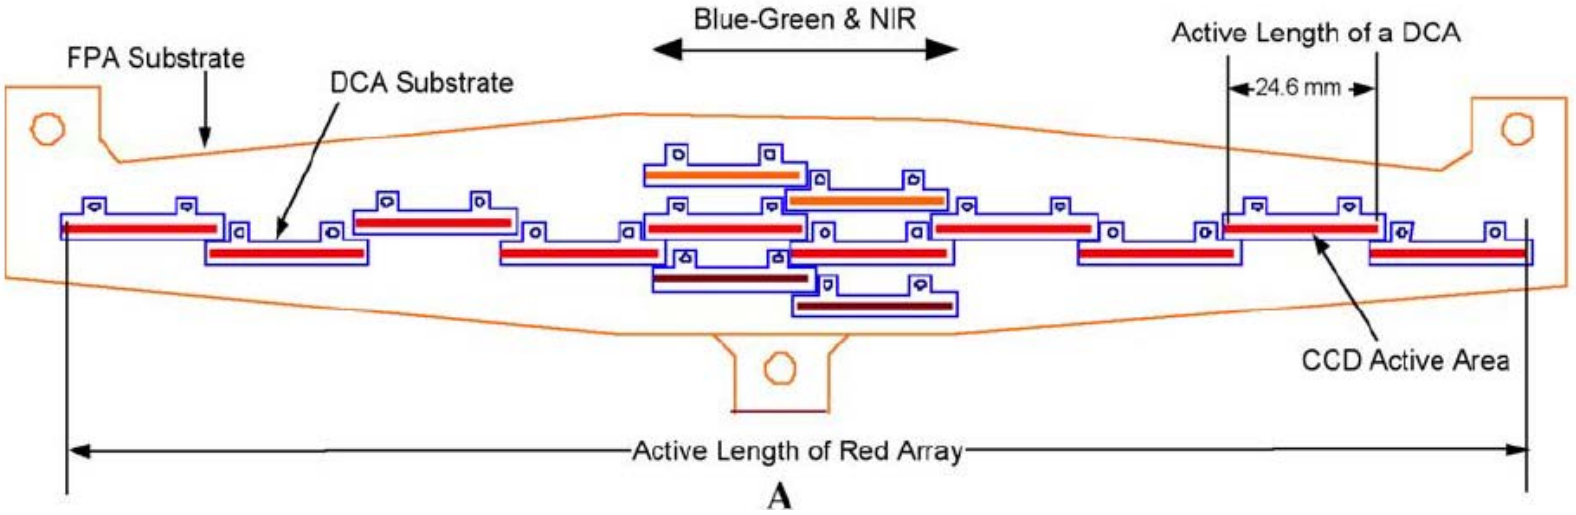

Case study: Mars Reconnaissance Orbiter *HiRISE*

Table 1. HiRISE Camera Capabilities at 300 km Altitude

Parameter	Characteristics
Ground sampling dimension	30 cm/pixel (1 μrad IFOV)
Resolution	~90 cm (3 pixels across an object)
Swath width (RED CCDs)	6 km (1.14° FOV)
Color swath width	1.2 km (0.23° FOV)
Maximum image size (pixels)	20,000 × 63,780 (14-bit data)
SNR (anywhere on Mars in the optimal season)	From 90:1 to 250:1 in RED channels with TDI 128 and full resolution
Color band passes (at half maximum of Mars- and solar-weighted spectral response; see Figure 12)	RED: 570–830 nm BG: <580 nm NIR: >790 nm
Stereo topographic precision	~25 cm vertical over ~1 m^2 areas
TDI lines	8, 32, 64, or 128
Pixel binning	none (1 × 1), 2 × 2, 3 × 3, 4 × 4, 8 × 8, 16 × 16
Bits per pixel	14, can be compressed to 8 via look-up tables (LUTs)
Compression (8-bit images only)	FELICS, compression >1.6:1

12 µm-wide pixels for HiRISE;

can be as small as 2 µm when compactness/resolution are key

BUT capacity scales with $(\text{width})^2$, so e.g. Kepler has 27 µm pixels to maximize dynamic range (full well capacity ~10^6 electrons)

Anatomy of a CCD pixel

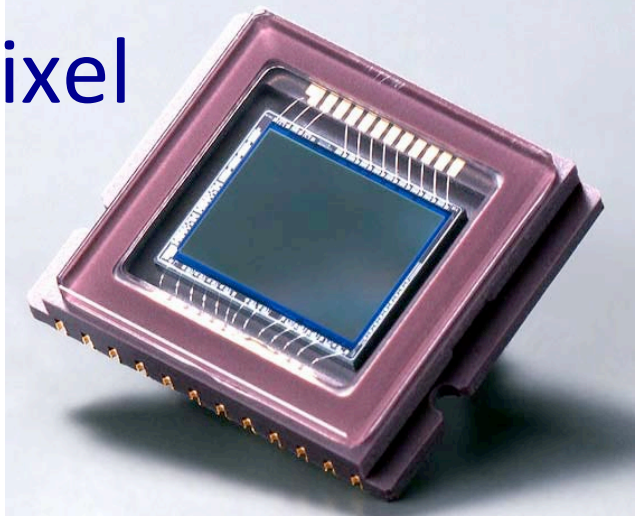

- Coatings to minimize reflection off outer surface
- Can operate upside down (*backside illumination*) if the CCD is *thinned*
  - Generally more expensive
  - Higher QE because no gate structures in the way
  - Less absorption of red photons (because thinner)
  - Potential non-uniform thickness must be calibrated

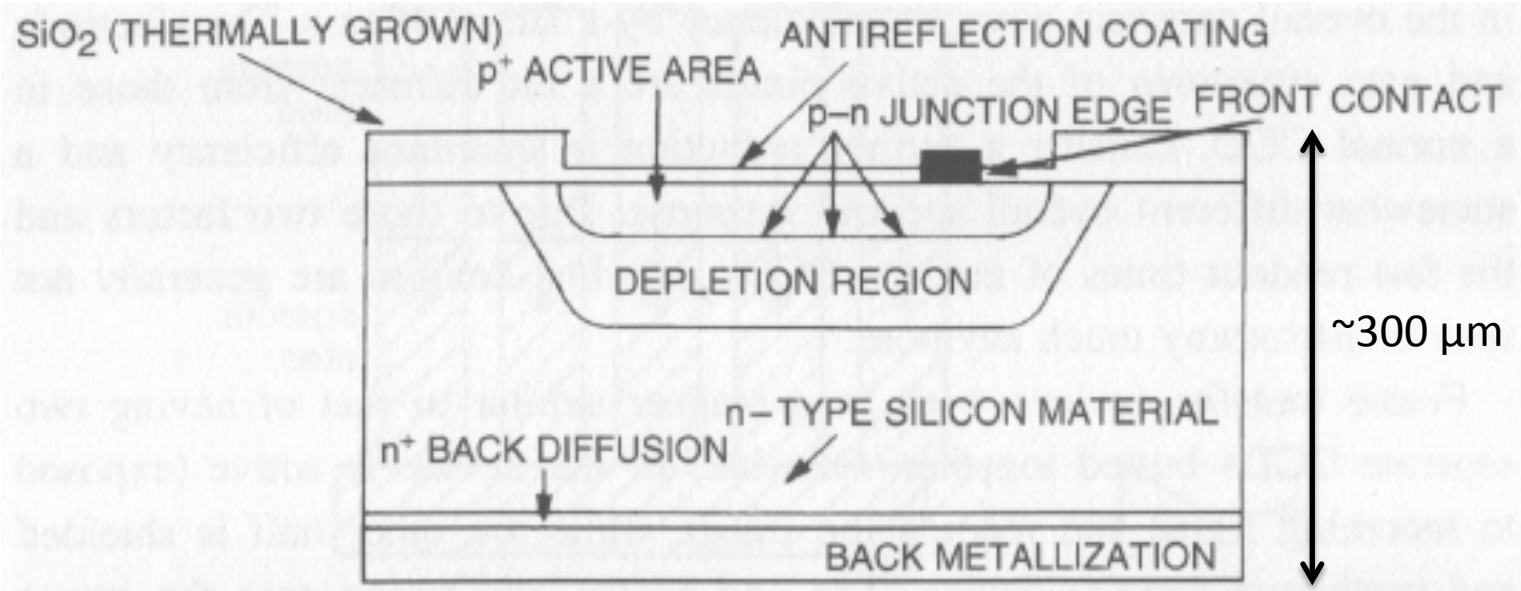

Fig. 2.4. Schematic view of a single front-side illuminated CCD pixel. The square labeled "front contact" is a representation of part of the overall gate structure. The letters "p" and "n" refer to regions within the pixel consisting of silicon doped with phosphorus and boron respectively.

Front-side vs. back-side (thinned) CCDs

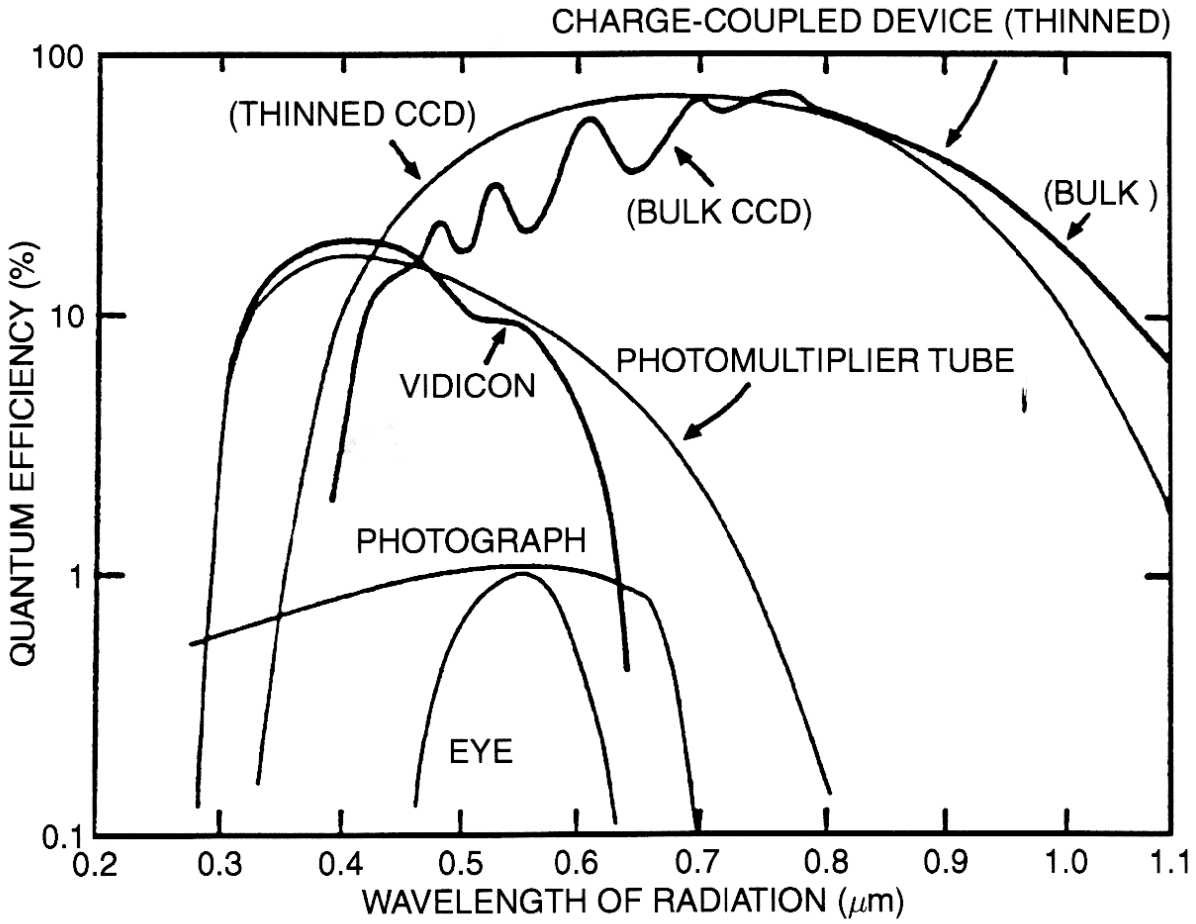

Fig. 3.2. QE curves for various devices, indicating why CCDs are a quantum leap above all previous imaging devices. The failure of CCDs at optical wavelengths shorter than about 3500 Å has been essentially eliminated via thinning or coating of the devices (see Figure 3.3).

One of the keys to HiRISE's excellent S/N: back-side illumination

Designing HiSCI

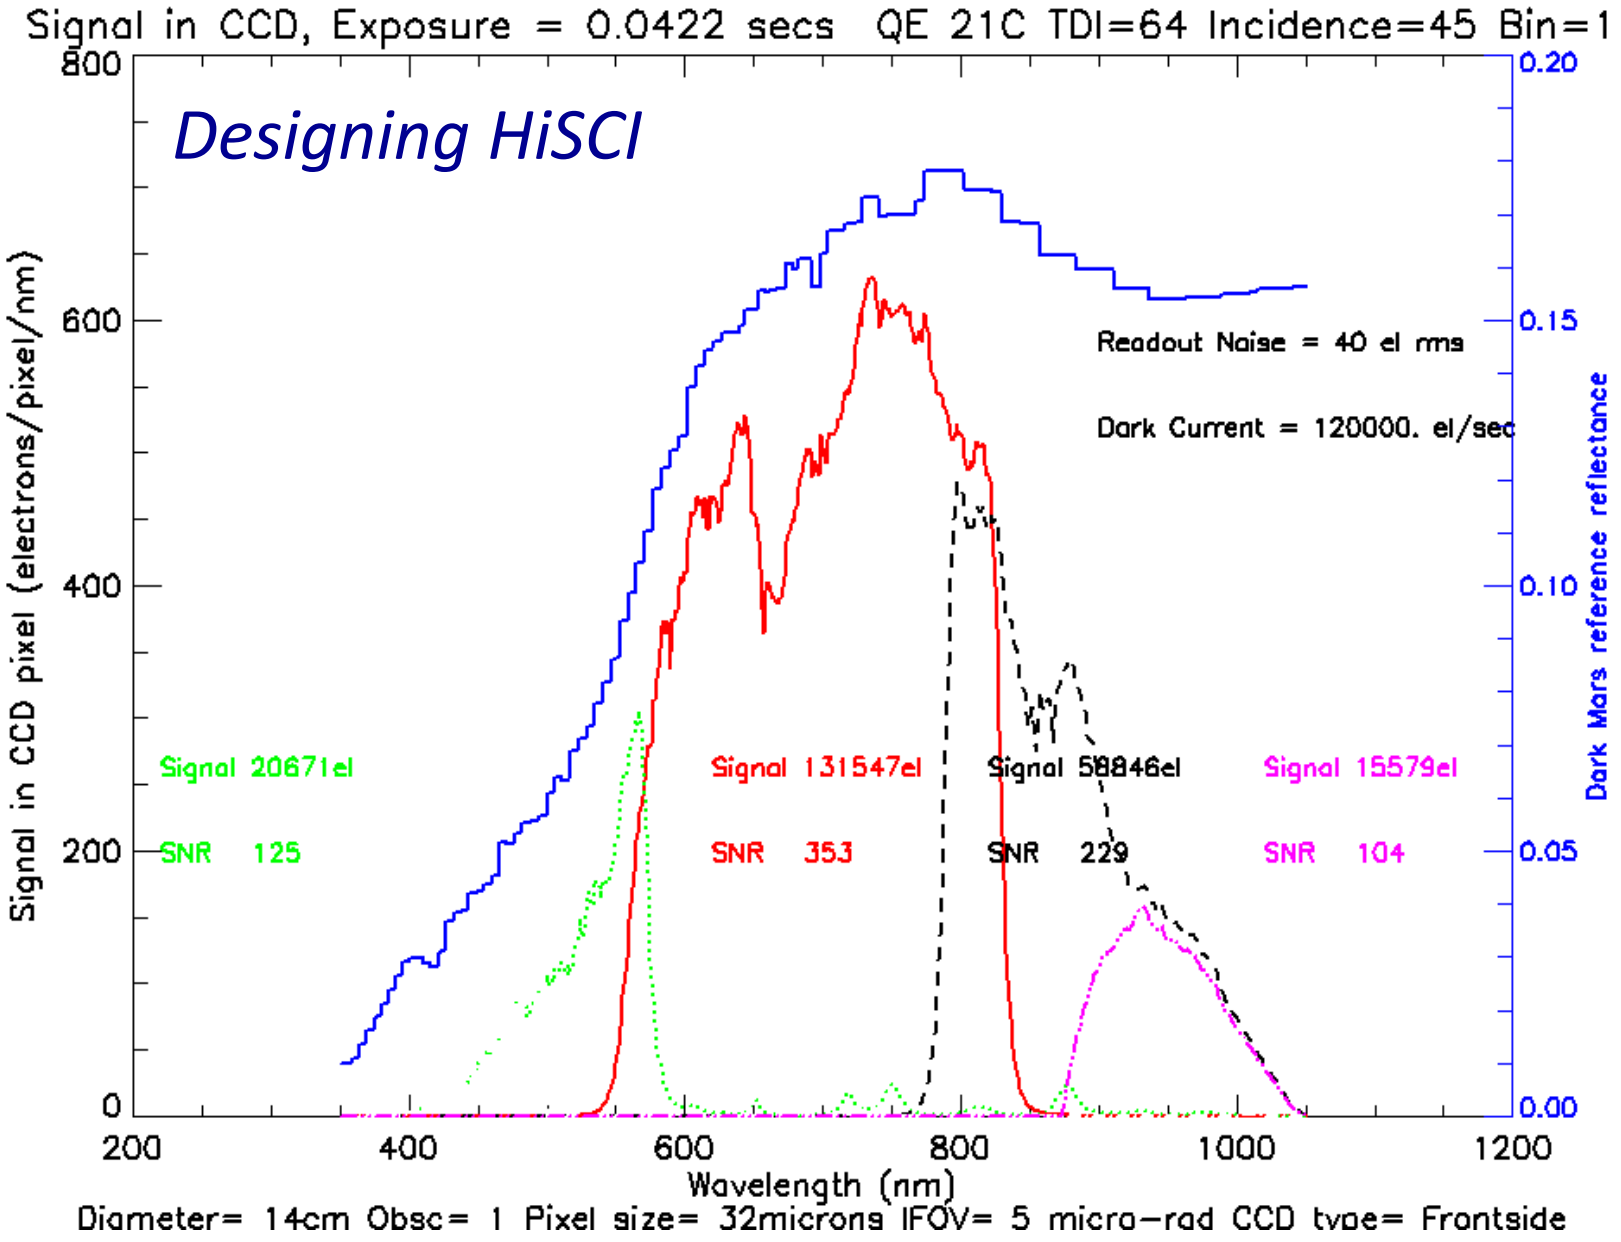

2010 Frontside QE data

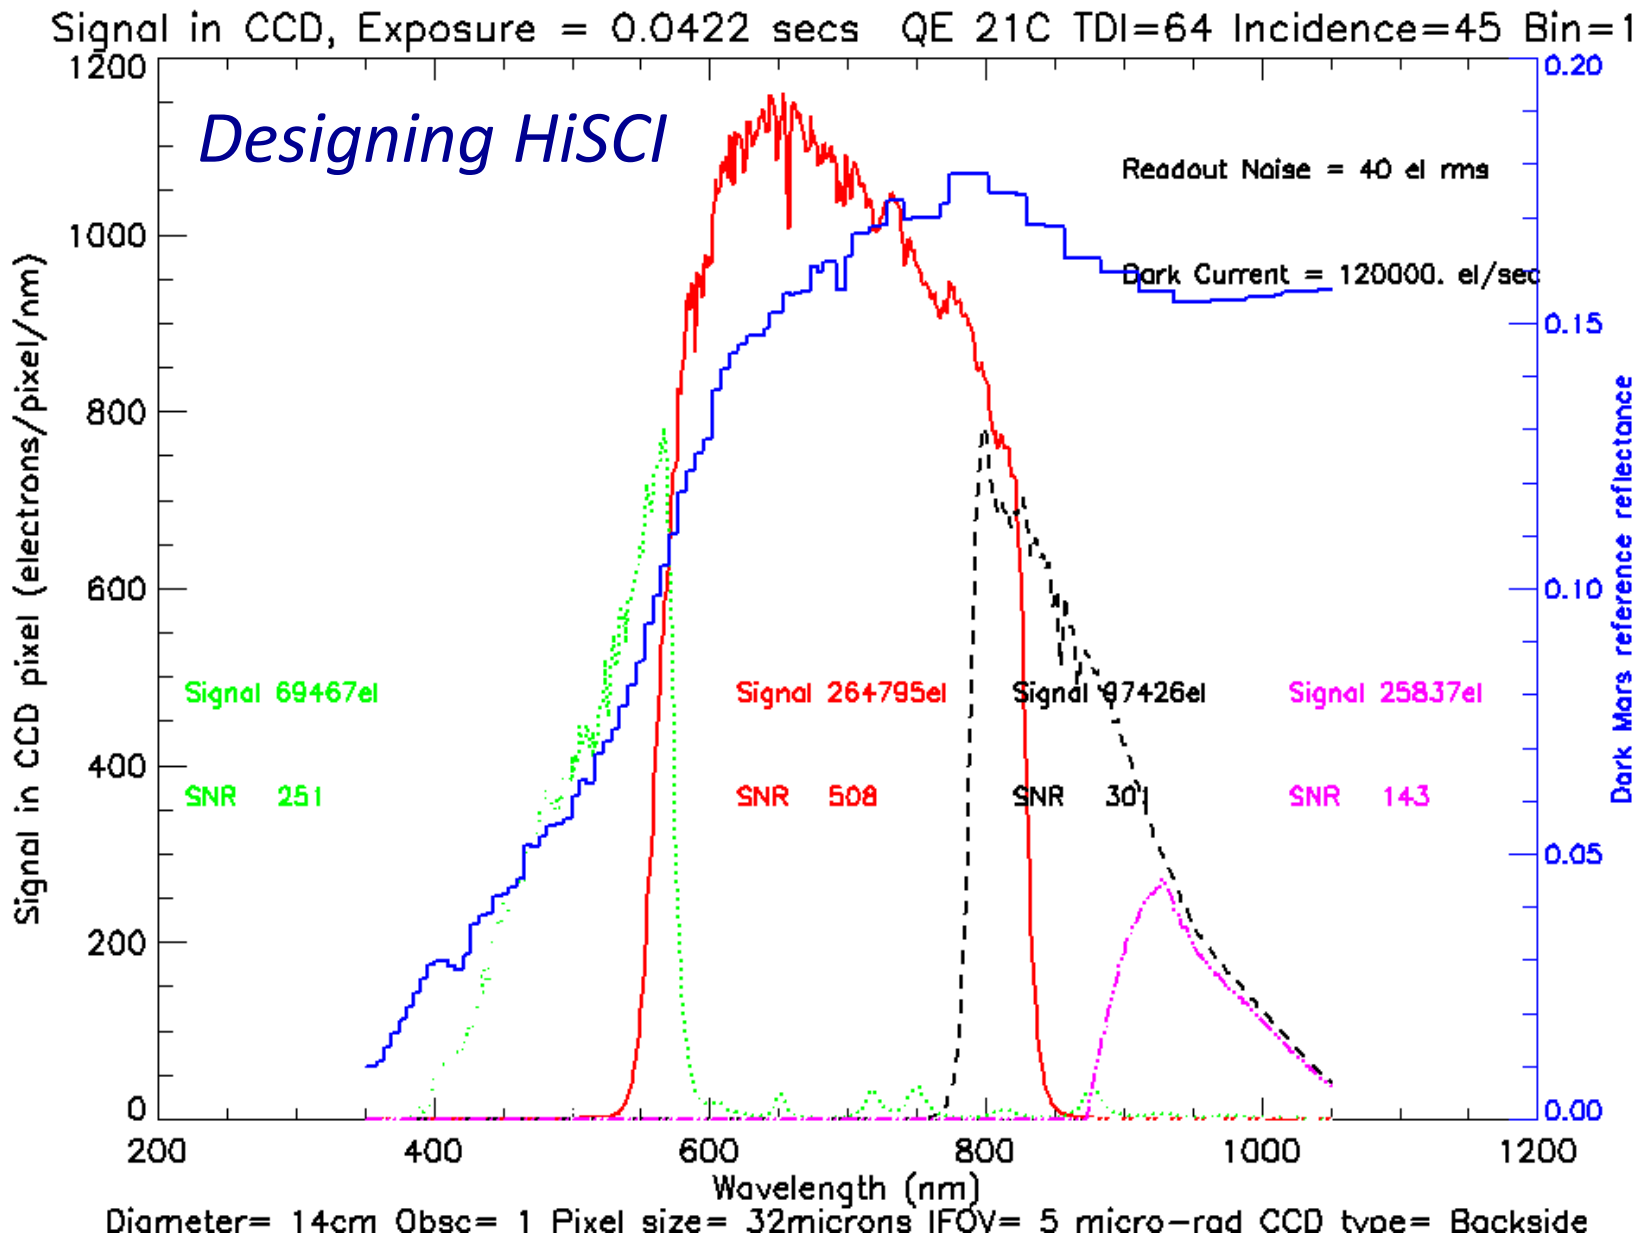

2010 Backside QE data – Morley's Model

Coatings can boost blue/UV absorptivity

Generally a phosphorescent material that “converts” UV photons into detectable longer-wavelength photons

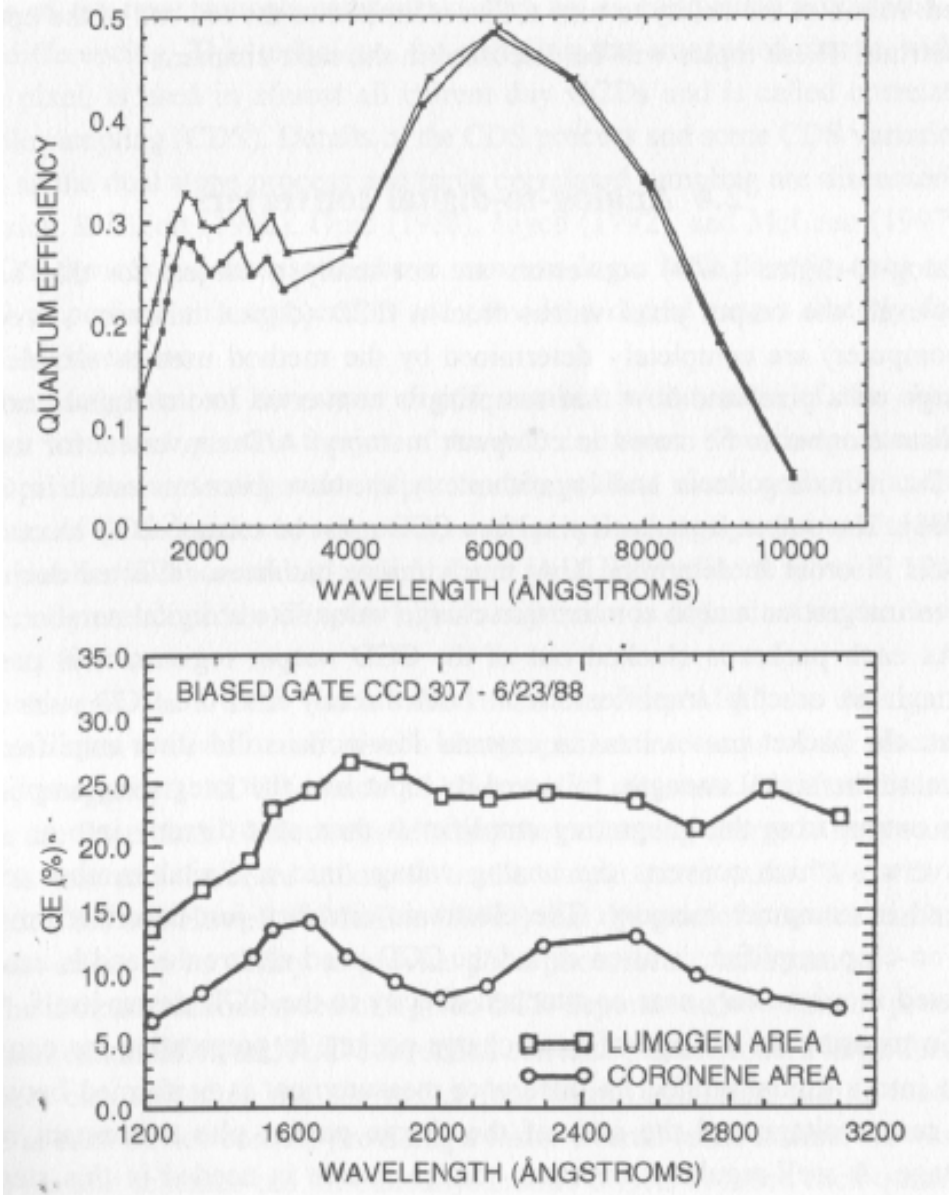

Fig. 2.9. The top plot shows QE curves for a Hubble Space Telescope WF/PC prototype CCD before and after being coated with lumogen. Note the increased UV response of the coated CCD. The bottom plot shows the QE properties of a WF/PC prototype in the far-UV spectral region. Presented are two curves, one for a coronene coated CCD and one for a lumogen coated CCD. From Trauger (1990).

Coatings can boost blue/UV absorptivity

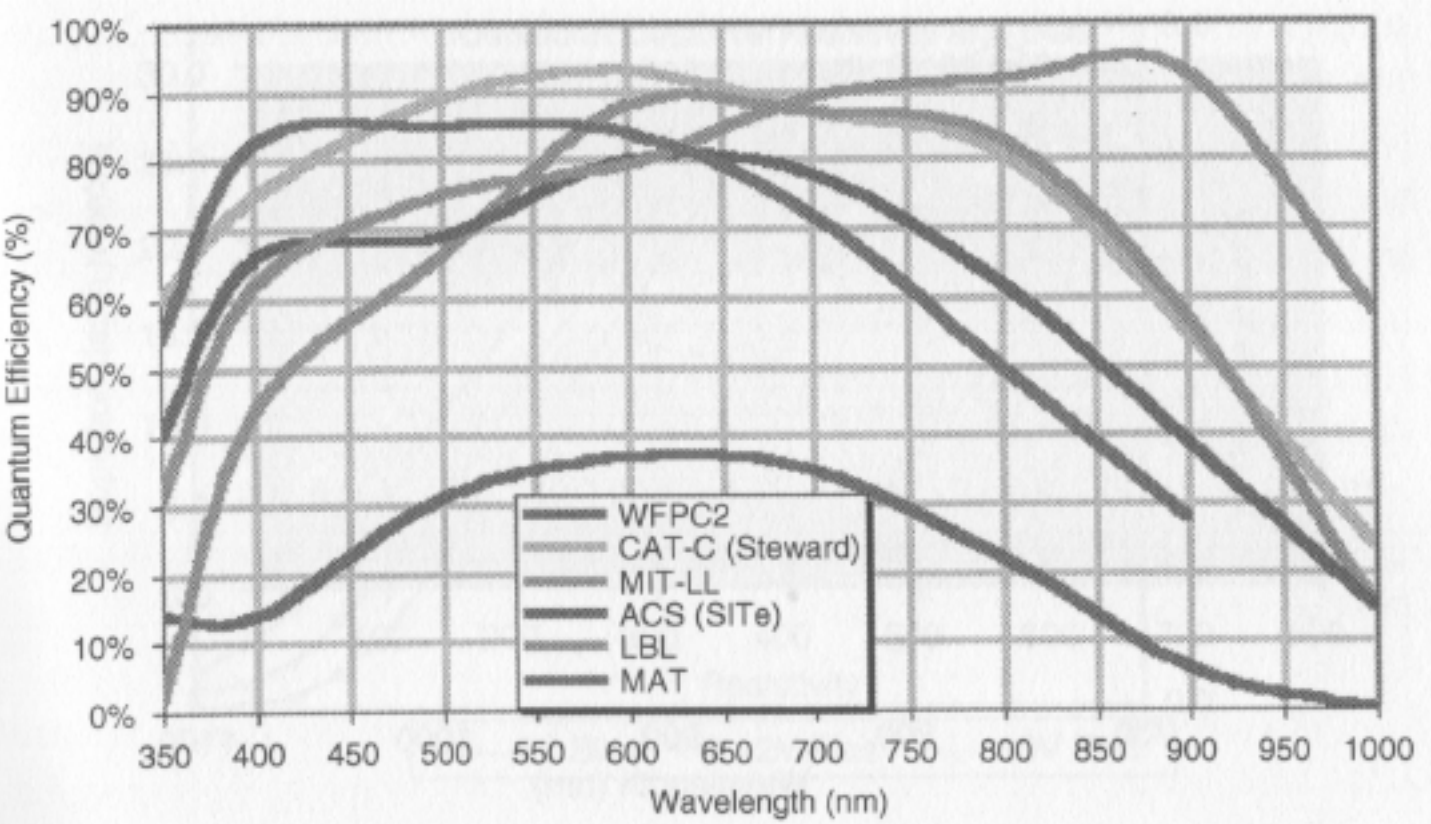

Fig. 3.3. QE curves for a variety of CCDs. WFPC2 is the second generation wide-field/planetary camera aboard HST, CAT-C is a new generation SITe CCD used in a mosaic imager at the University of Arizona's 90" telescope on Kitt Peak, MIT-LL is a CCD produced at the MIT Lincoln Laboratories and optimized for red observations, ACS is the Hubble Space Telescope Advanced Camera for Surveys SITe CCD, LBL is a Lawrence Berkeley Lab high resistivity, "deep depletion" CCD with high red QE, and MAT is a front-side, processed CCD showing high blue QE.

QE also depends on temperature

(higher T → higher QE … but also higher noise)

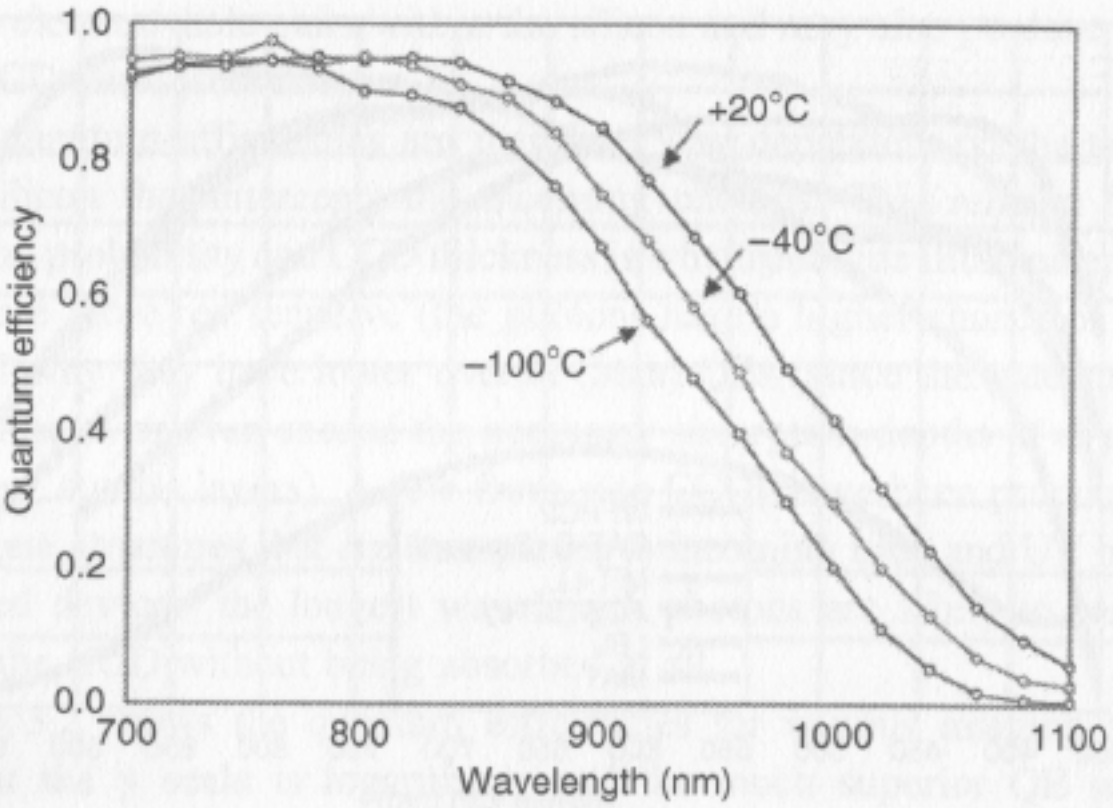

Fig. 3.4. Sensitivity of the quantum efficiency of a MIT/LL CCD for three operating temperatures. The blue sensitivity is little affected by a change in operating temperature but the red QE can change by a factor of two. The use of such devices requires a balance of higher operating temperature and keeping the dark current under control.

QE also depends on temperature

*(higher T → higher QE … but also higher noise — **dark current**)*

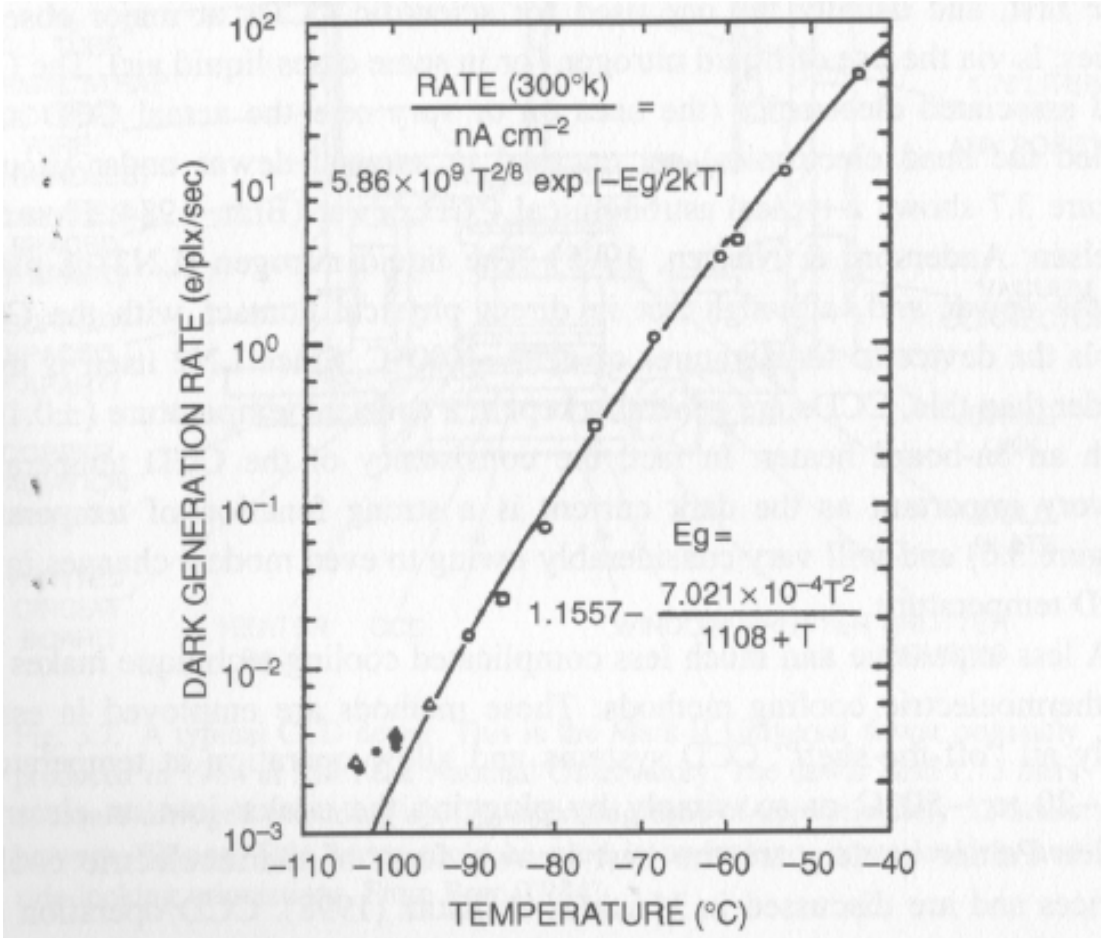

Fig. 3.6. Experimental (symbols) and theoretical (line) results for the dark current generated in a typical three-phase CCD. The rate of dark current, in electrons generated within each pixel every second, is shown as a function of the CCD operating temperature. E_g is the band gap energy for silicon. From Robinson (1988a).

CCDs have exceptionally good linearity…

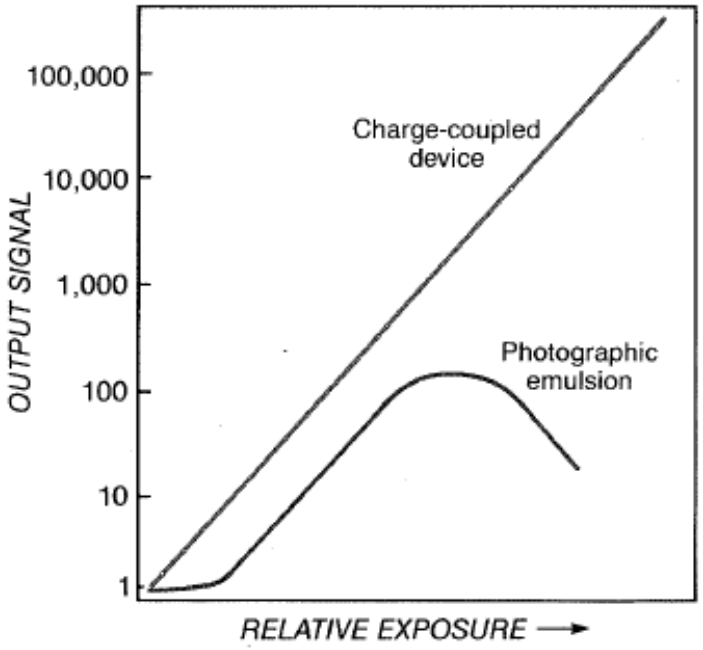

CCD's are exceptionally linear; doubling the exposure yields a signal exactly twice as strong. This holds over an enormous exposure range, enabling CCD's to detect simultaneously very bright and very faint objects in a single image. This behavior contrasts strikingly with that of photographic emulsions, which are only linear over a narrow range of exposure. Emulsions must also overcome a threshold before they start to form an image at all. Worse, beyond a certain point emulsions saturate; their ability to detect light actually decreases if the exposure is too long.

…until you approach the full well capacity

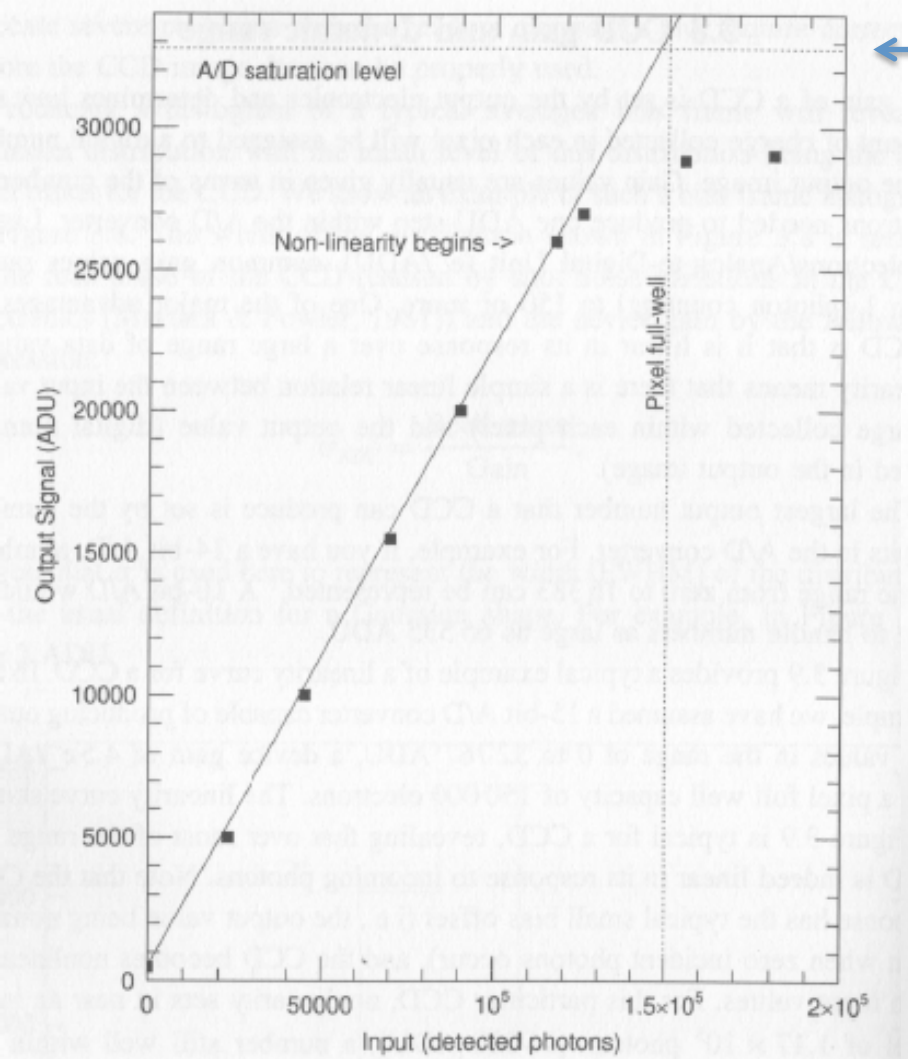

Fig. 3.9. CCD linearity curve for a typical three-phase CCD. We see that the device is linear over the output range from 500 ADU (the offset bias level of the CCD) to 26 000 ADU. The pixel full well capacity is 150 000 electrons and the A/D converter saturation is at 32 767 ADU. In this example, the CCD nonlinearity is the limiting factor of the largest usable output ADU value. The slope of the linearity curve is equal to the gain of the device.

← HiRISE is 14-bit, so A/D saturation at half this DN

…and there can be deviations even before that

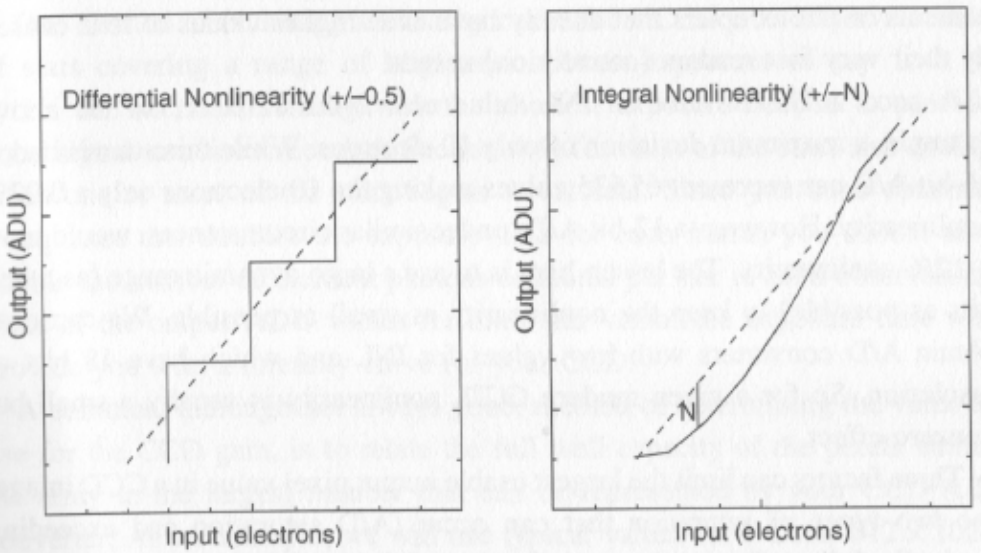

Fig. 3.10. The two types of CCD nonlinearity are shown here in cartoon form. Differential nonlinearity (left) comes about due to the finite steps in the A/D conversion process. Here we see that the linearity curve (dashed line) cuts through each step at the halfway point yielding a DNL of ±0.5 ADU. Integral nonlinearity (right) is more complex and the true linearity curve (solid line) may have a simple or complex shape compared with the measured curve (dashed line). A maximum deviation (N) is given as the INL value for an A/D and may occur anywhere along the curve and be of either sign. Both plots have exaggerated the deviation from linearity for illustration purposes.

Deviations from CCD linearity

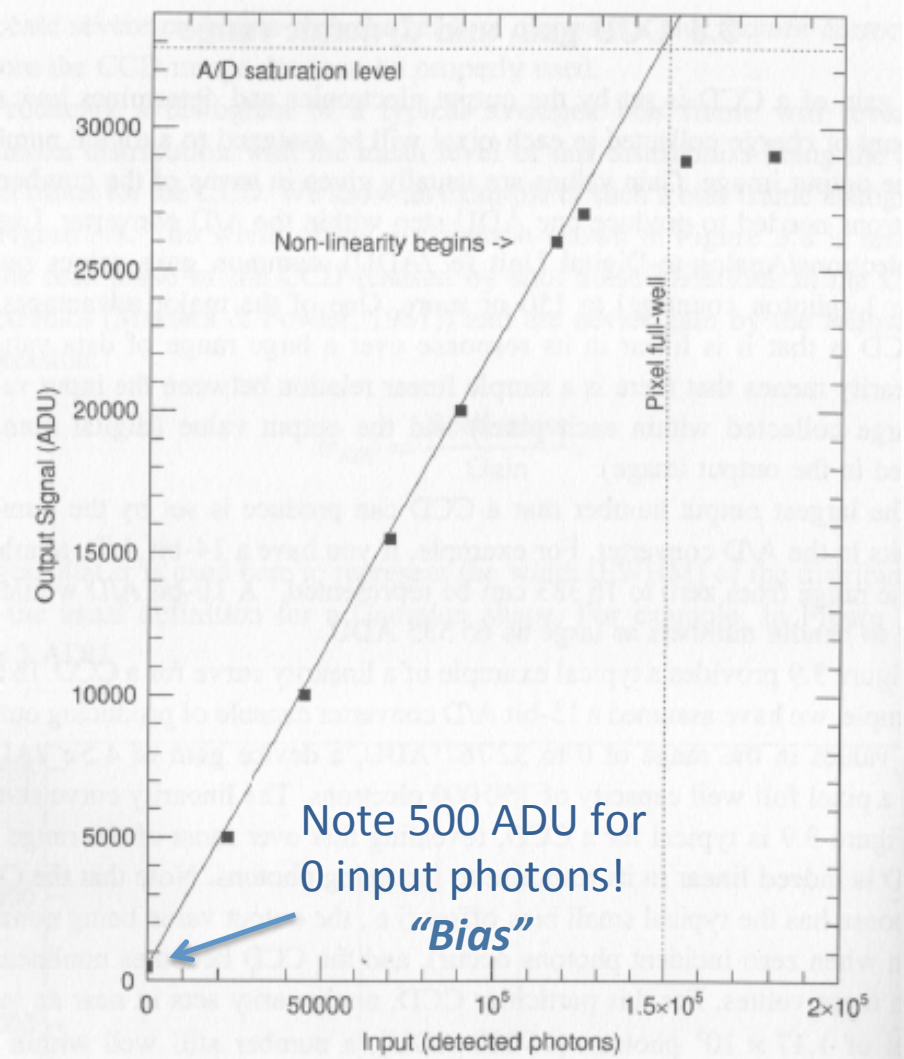

Fig. 3.9. CCD linearity curve for a typical three-phase CCD. We see that the device is linear over the output range from 500 ADU (the offset bias level of the CCD) to 26 000 ADU. The pixel full well capacity is 150 000 electrons and the A/D converter saturation is at 32 767 ADU. In this example, the CCD nonlinearity is the limiting factor of the largest usable output ADU value. The slope of the linearity curve is equal to the gain of the device.

What's a simple way to check your CCD's linearity?

Integrate for 1 second, 2 seconds, 4 seconds, …

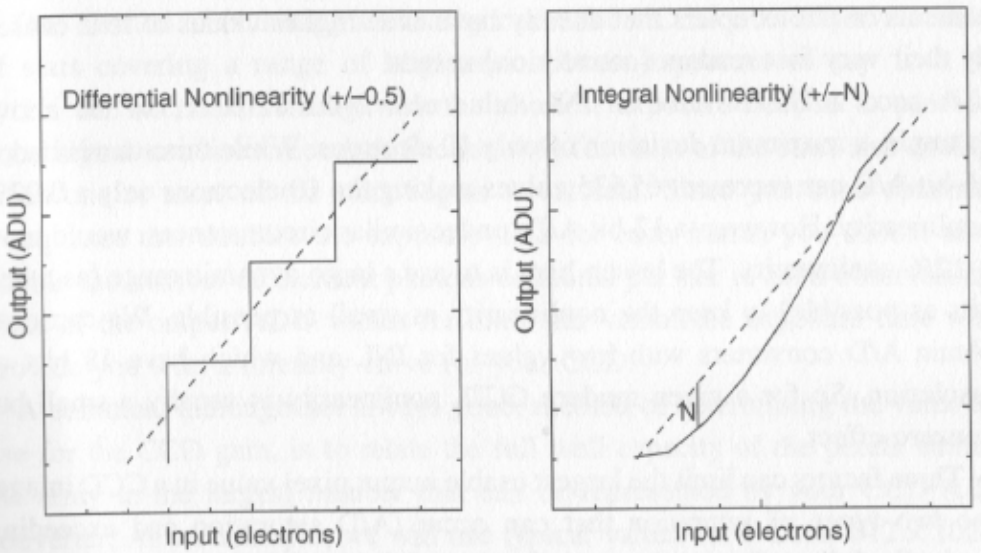

Fig. 3.10. The two types of CCD nonlinearity are shown here in cartoon form. Differential nonlinearity (left) comes about due to the finite steps in the A/D conversion process. Here we see that the linearity curve (dashed line) cuts through each step at the halfway point yielding a DNL of ±0.5 ADU. Integral nonlinearity (right) is more complex and the true linearity curve (solid line) may have a simple or complex shape compared with the measured curve (dashed line). A maximum deviation (N) is given as the INL value for an A/D and may occur anywhere along the curve and be of either sign. Both plots have exaggerated the deviation from linearity for illustration purposes.

Measuring the bias

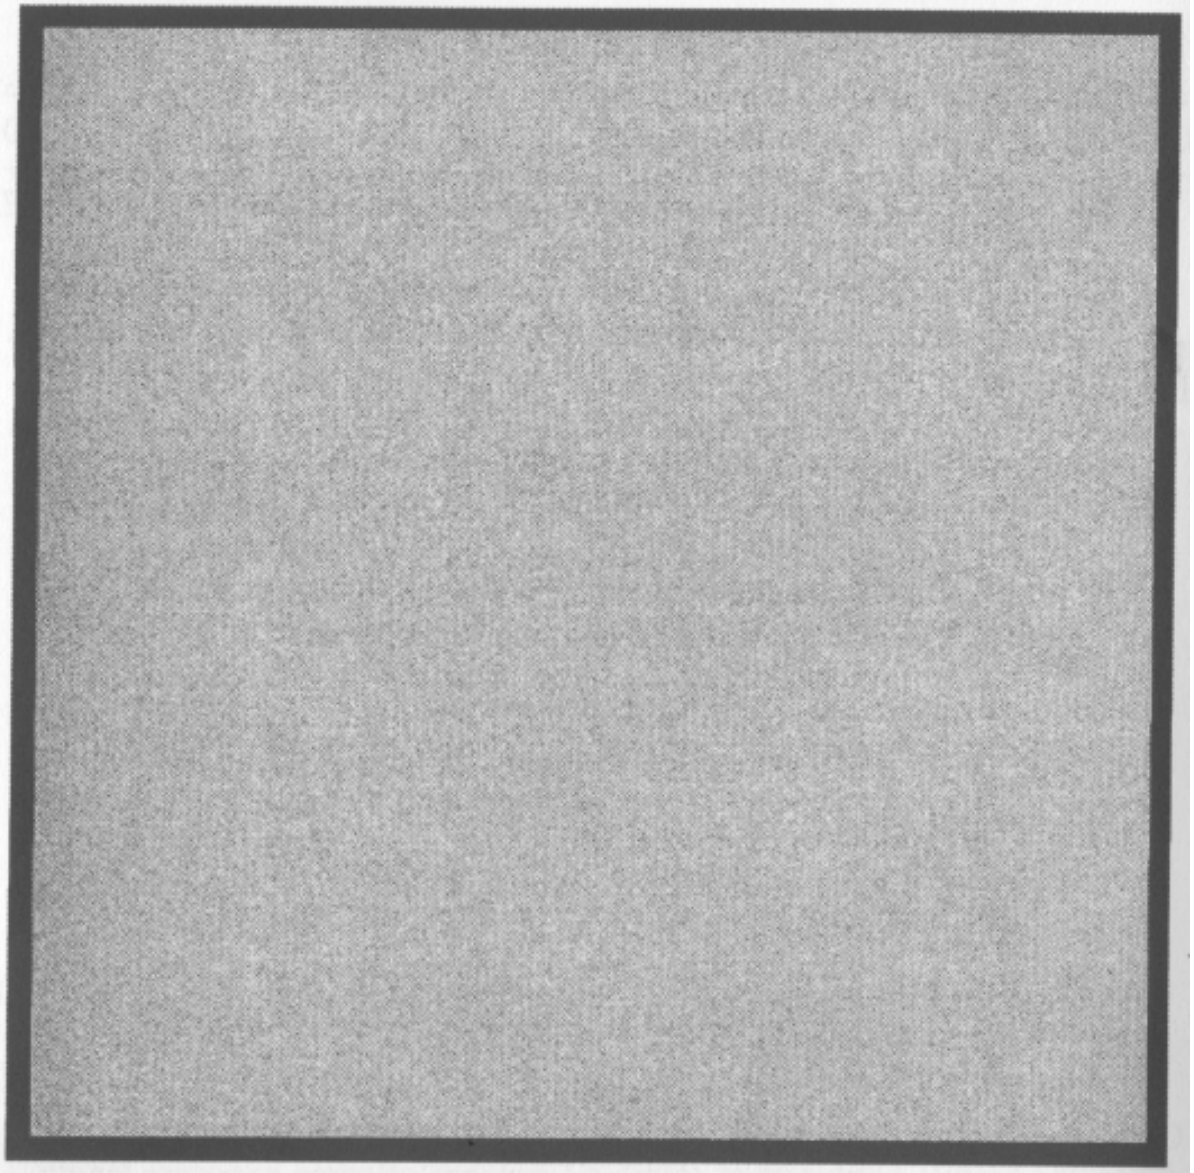

Fig. 4.2. Shown is a typical CCD bias frame. The histrogram of this image was shown in Figure 3.8. Note the overall uniform structure of the bias frame.

- A/D conversion purposely includes offset (bias) such that 0 e^- → >0 DN
- Measure the bias by reading out "image" acquired with integration time of 0 seconds

Bias is used because of ***read noise***

- Amplifier, A/D converter, electronics don't give perfectly repeatable results; i.e. the same pixel detecting the same radiance will yield distribution of DNs varying by ~ a few e^-

- Bias imposed so that read noise doesn't yield DNs < 0

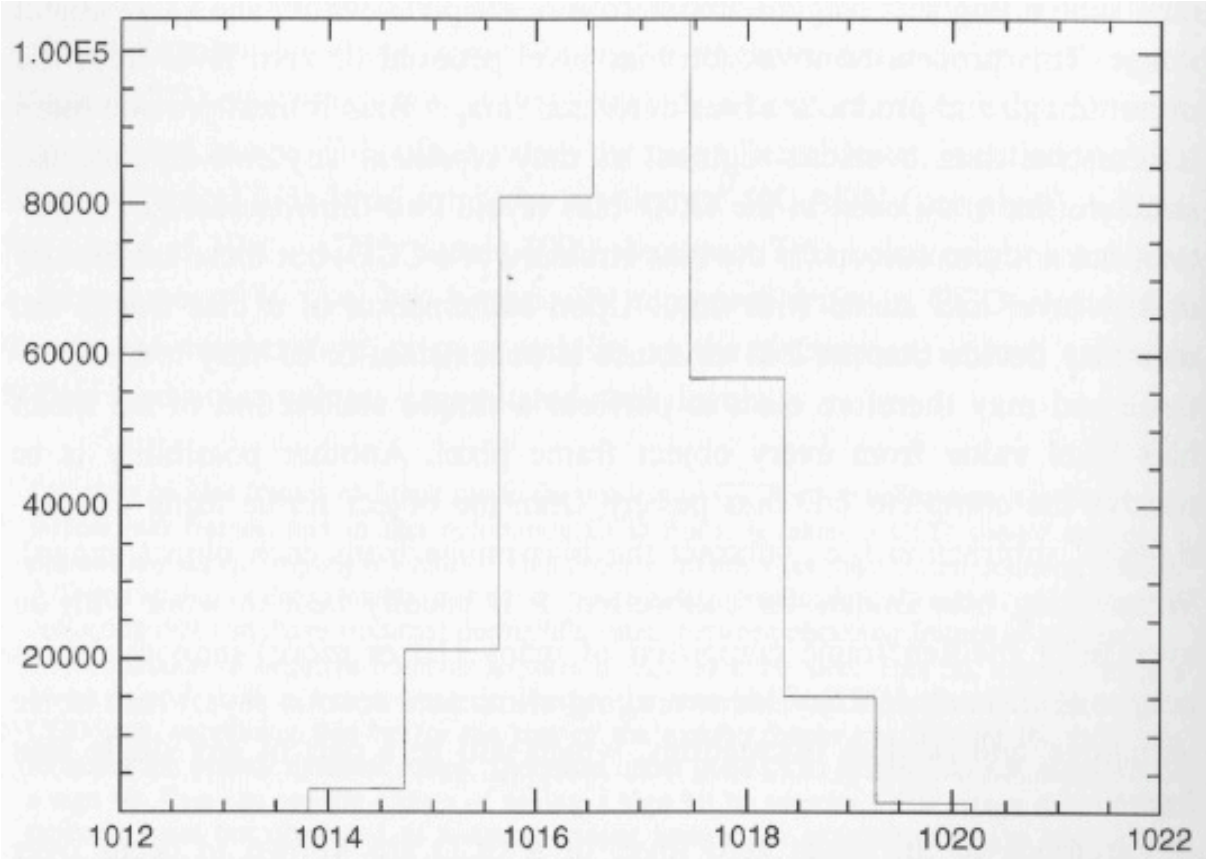

Fig. 3.8. Histogram of a typical bias frame showing the number of pixels vs. each pixel ADU value. The mean bias level offset or pedestal level in this Loral CCD is near 1017 ADU, and the distribution is very Gaussian in nature with a FWHM value of near 2 ADU. This CCD has a read noise of 10 electrons and a gain of 4.7 e^-/ADU.

Dark current

- Thermal noise

- Can vary from pixel to pixel

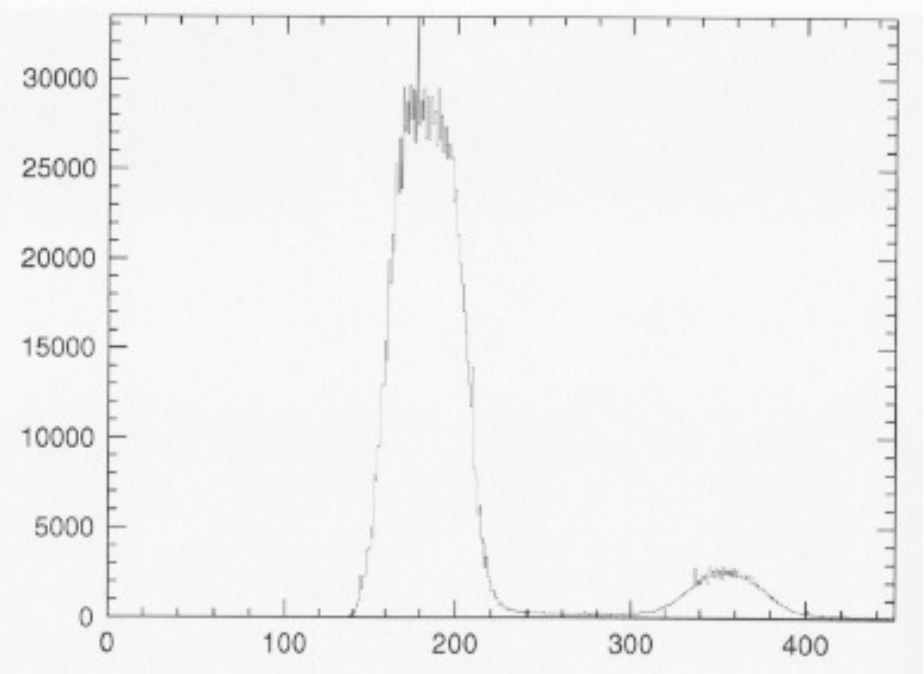

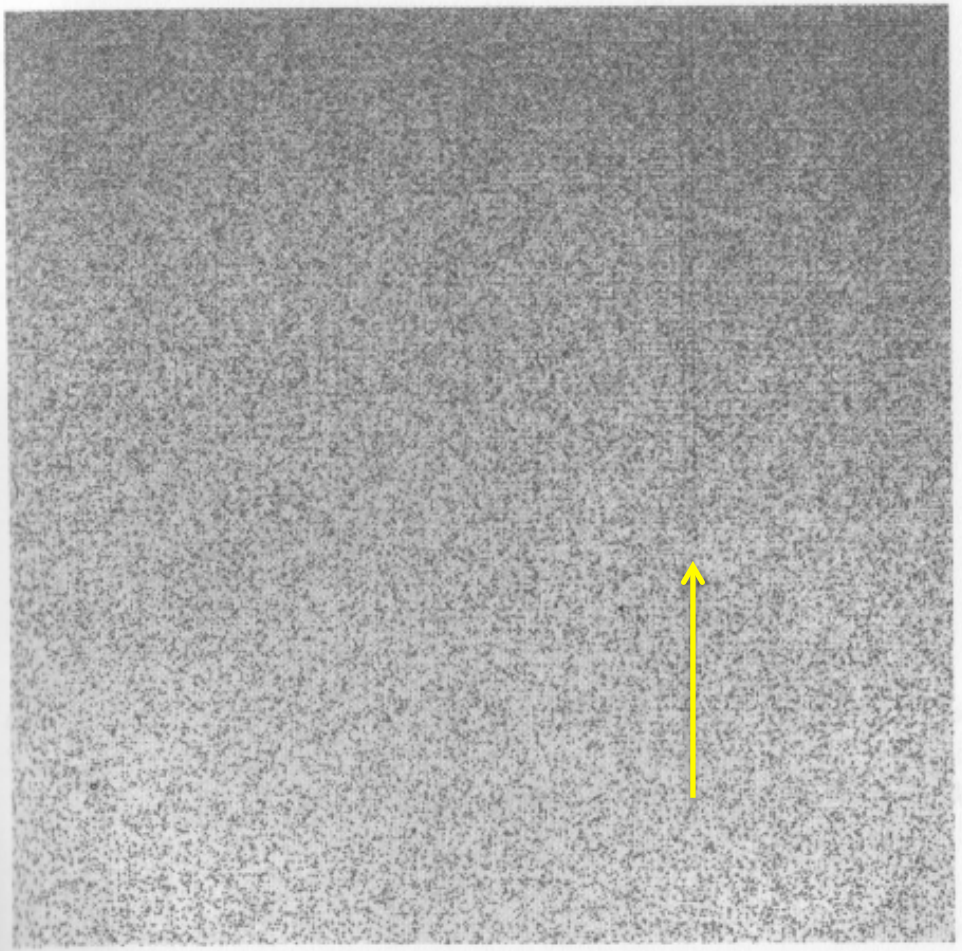

Fig. 4.3. Shown is a typical CCD dark frame. This figure shows a dark frame for a Kodak CCD operating in MPP mode and thermoelectrically cooled. Notice the nonuniform dark level across the CCD, being darker (greater ADU values) on the top. Also notice the two prominent partial columns with higher dark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

Dark current

- Negligible when cooled to near liquid N temperatures (77 K)

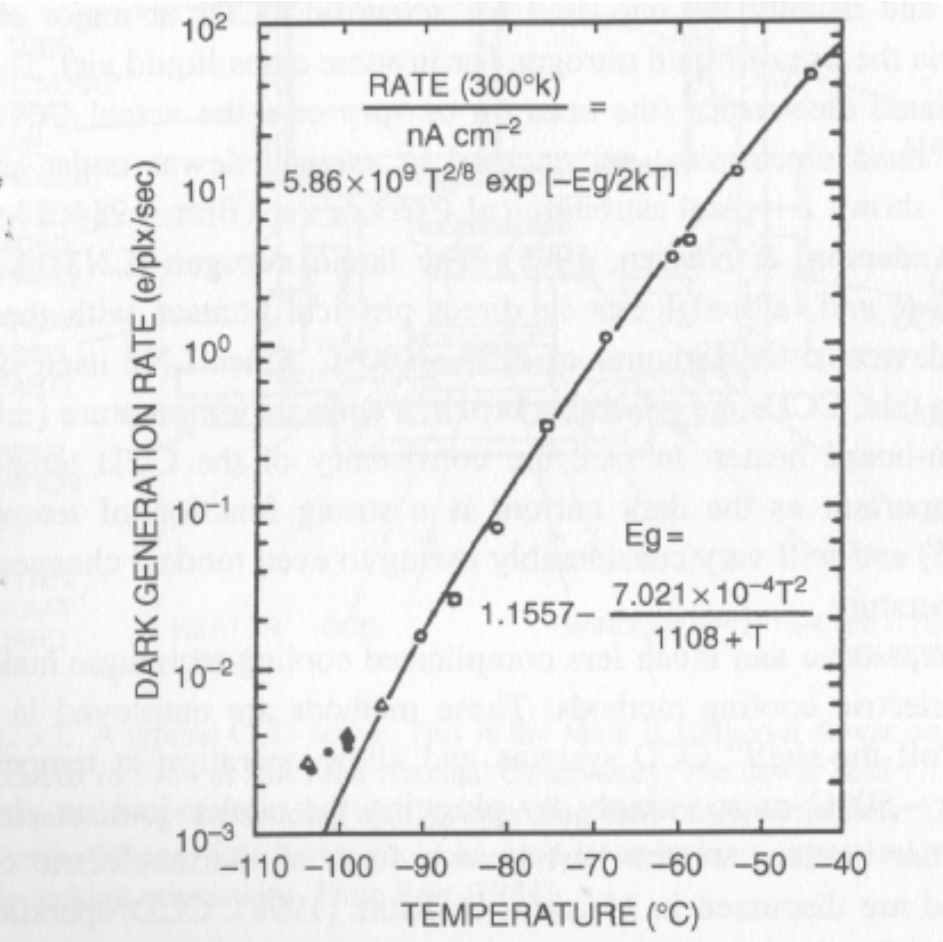

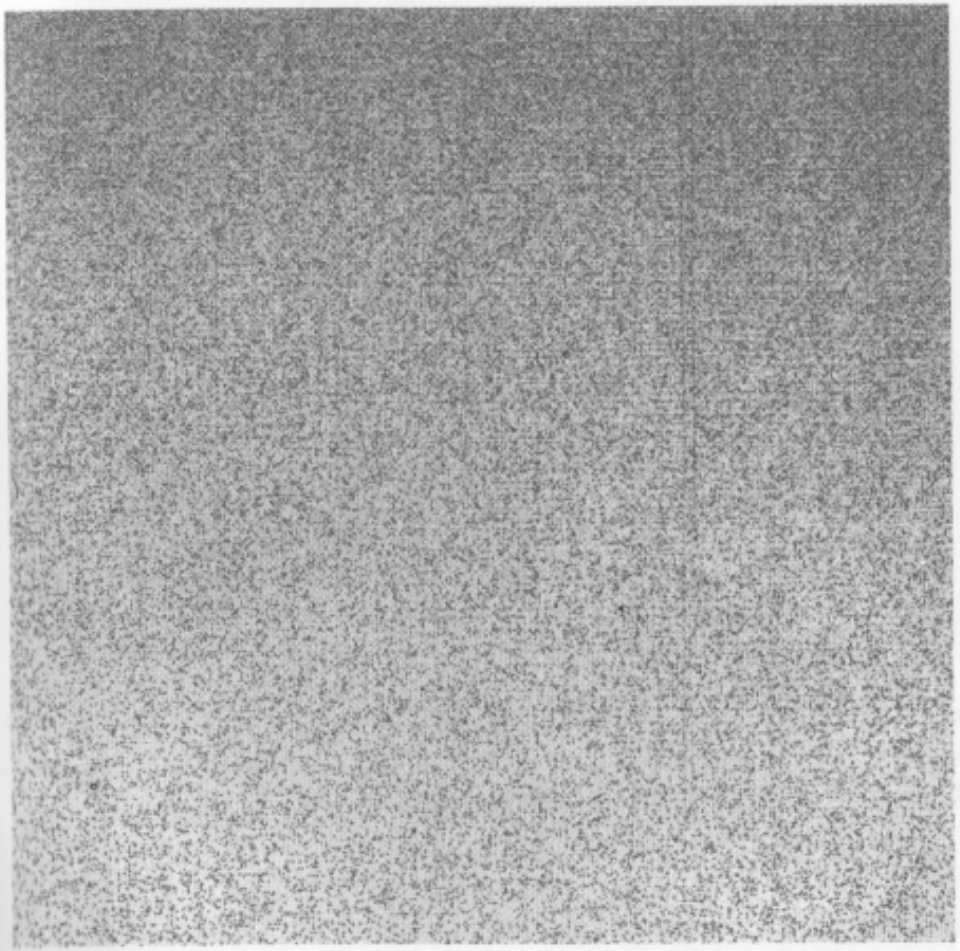

Fig. 4.3. Shown is a typical CCD dark frame. This figure shows a dark frame for a Kodak CCD operating in MPP mode and thermoelectrically cooled. Notice the nonuniform dark level across the CCD, being darker (greater ADU values) on the top. Also notice the two prominent partial columns with higher dark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

Also need to account for varying pixel sensitivity

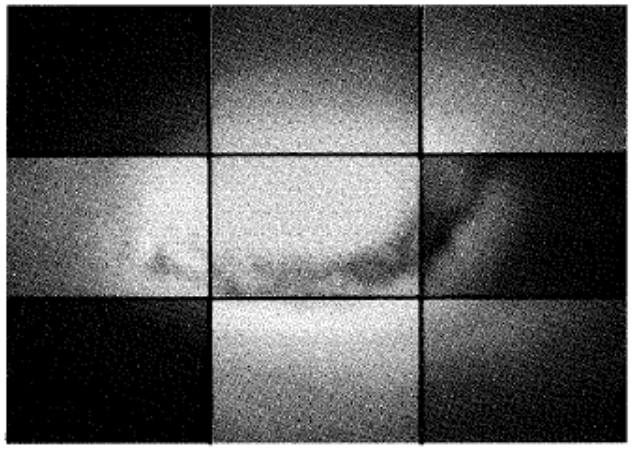

This simulated "raw" image from a charge-coupled device does not give a realistic representation of the well-known Black-eye galaxy, M64 in Coma Berenices, because individual pixels have varying sensitivities to light. This effect is shown here schematically by dividing the image into nine regions of grossly different sensitivities. Images such as this are obviously not very useful to astronomers or anybody else.

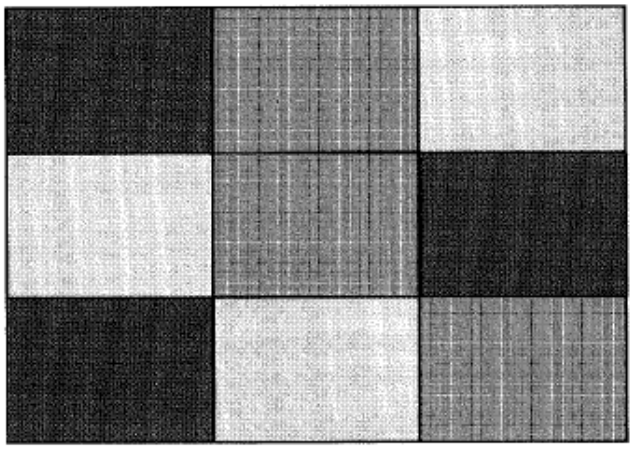

To overcome this nonuniformity problem, observers use a technique called flat-fielding. During an observing run an exposure is made of a uniform source such as the dawn sky or inside of the telescope dome to map the variations among the CCD's pixels. This image represents the flat-field frame corresponding to the simulated raw image of the Black-eye galaxy shown above.

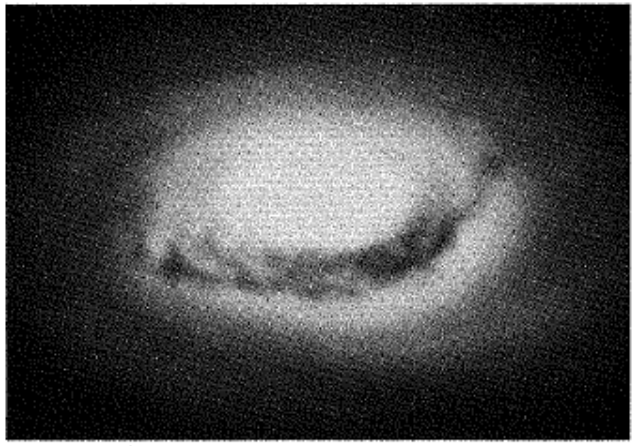

Once the response of each pixel is known, observers can scale the data by an amount determined from the flat-field exposure and thereby correct for each pixel's individual response. The result is the same as that which would be registered by a CCD with pixels of identical sensitivity; it can now be analyzed by astronomers. Image courtesy Rudolph Schild.

Flat field images

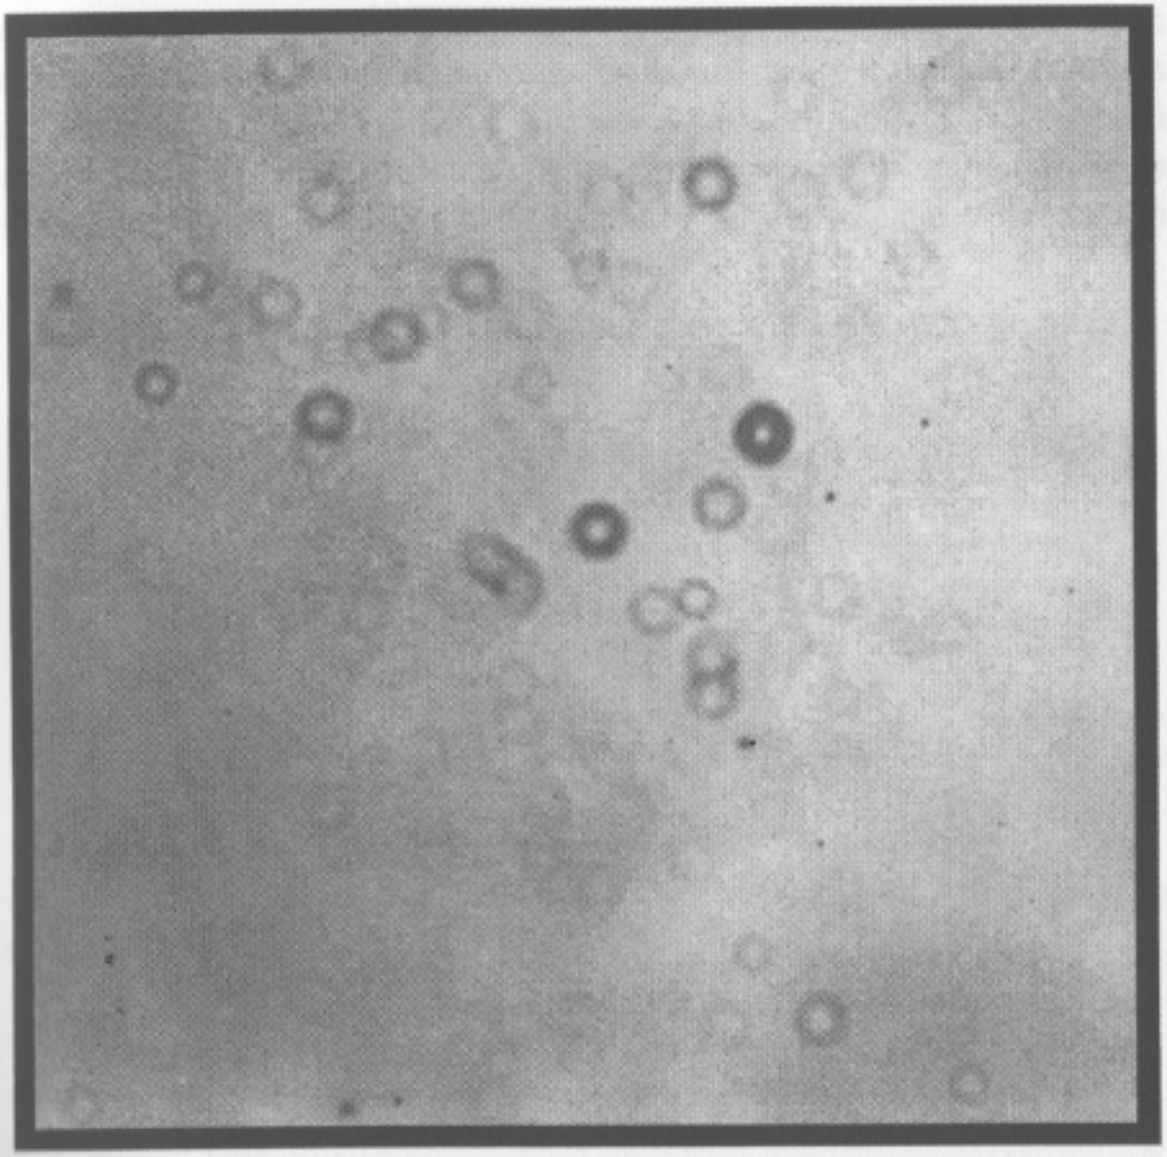

Fig. 4.4. Shown is a typical CCD flat field image. This is an R-band flat field image for a 1024 × 1024 Loral CCD. The numerous "doughnuts" are out of focus dust specks present on the dewar window and the filter. The varying brightness level and structures are common in flat field images. As seen in the histogram of this image (Figure 4.1) this flat field has a mean level near 6950 ADU, with an approximate dispersion of (FWHM) 400 ADU.

- Illuminate detector uniformly, preferably with color spectrum similar to what you actually want to image
- Note dust specs on detector surface, ***and*** "low spatial frequency" variations
- With telescopes, point at dome wall or white screen; not so easy with instruments in space…

Bottom line of image processing ("data reduction")

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

- Replace "Bias" with "Dark" if not at extremely low temperature (dark image will include the bias anyway)

- "Flat Field Frame" assumed to also be Bias/Dark-subtracted

The all-important ***signal-to-noise ratio***

- Estimate from the "CCD Equation":

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

- N_* = photons (or equivalent e^-) from target object(s)
- n_{pix} = # of pixels
- N_S = photons/e^- from "background" ("empty" sky in astronomy)
- N_D = dark current
- N_R = read noise

Note that if N_* >> the noise terms, then S/N ~ √(N_*)

Increases as *square root of* **integration time:**

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

Time Delay Integration / Drift Scanning

- Read out CCD ***while imaging***

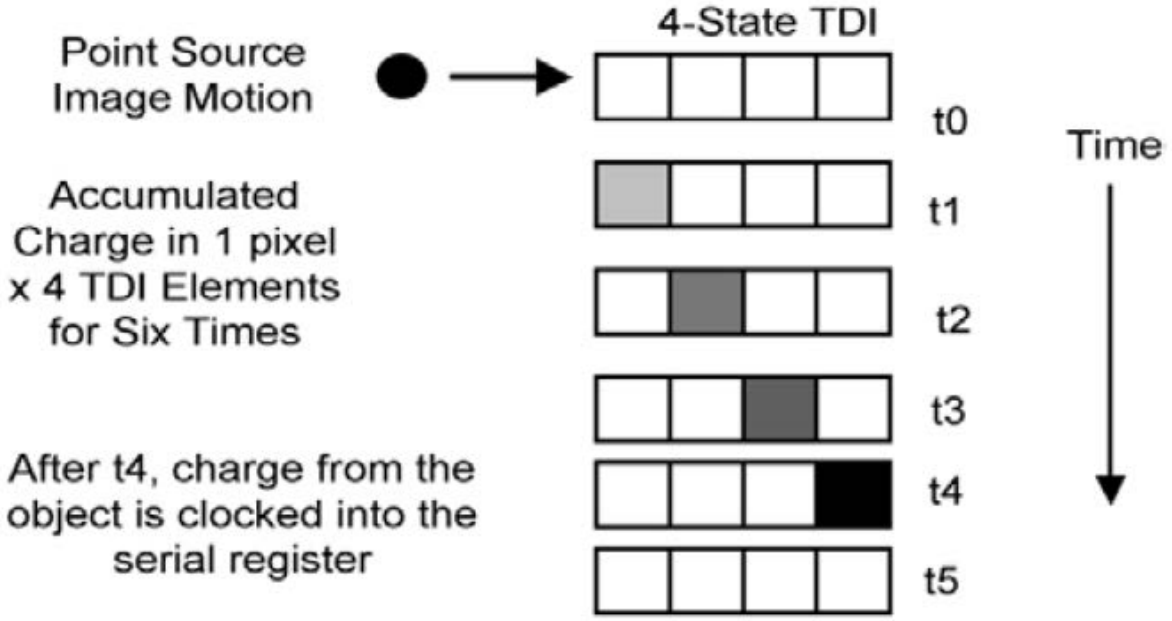

Figure 6. Illustration of 4-line Time Delay and Integration (TDI). HiRISE can use 8, 32, 64, or 128 TDI lines.

- Must transfer charge at exact rate of target motion relative to CCD
- Can greatly improve integration time (and thus S/N)

Time Delay Integration with HiRISE

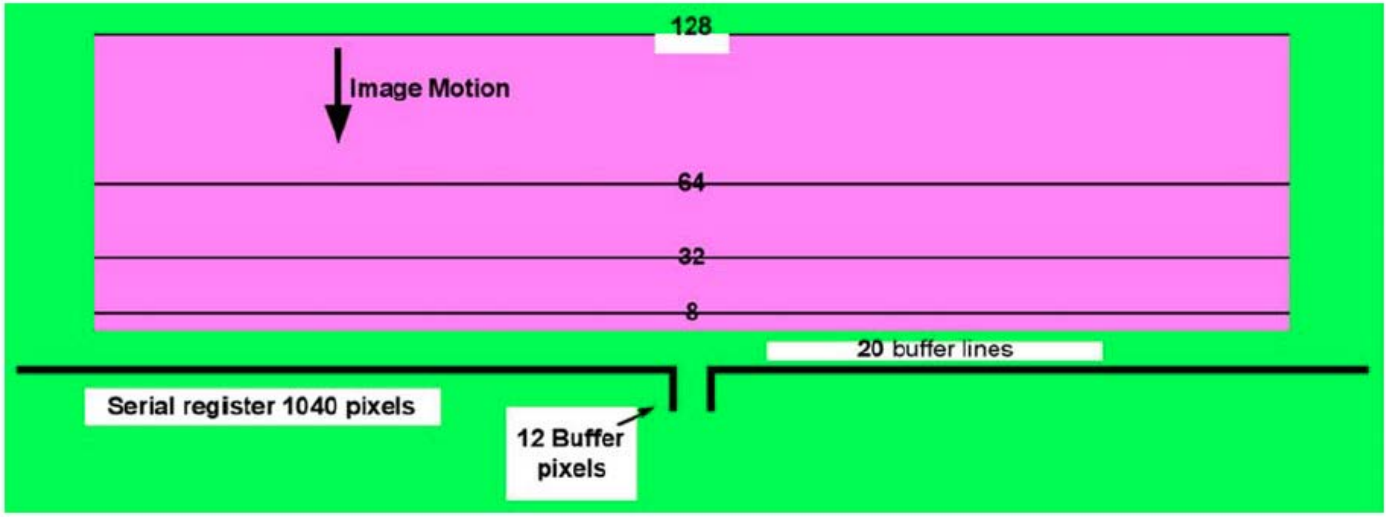

Figure 7. Layout of each HiRISE CCD. There are 4 sections of lines corresponding to the 4 TDI choices, plus 20 buffer lines (read out once at the beginning of each image). The data is read out into 2 channels from a tap in the middle of the array, beginning each line with 12 buffer pixels and ending with 16 masked pixels. Channel 0 is on the right and channel 1 on the left side, as viewed here, so the data from channel 1 must be "mirrored or flipped left to right to restore the proper image orientation. (Note that spacecraft motion is up here, whereas it is down in Figure 5.)

dark

bias

HiRISE flat fielding

- Pre-launch data, onboard LEDs, and …
  imaging Mars sideways!

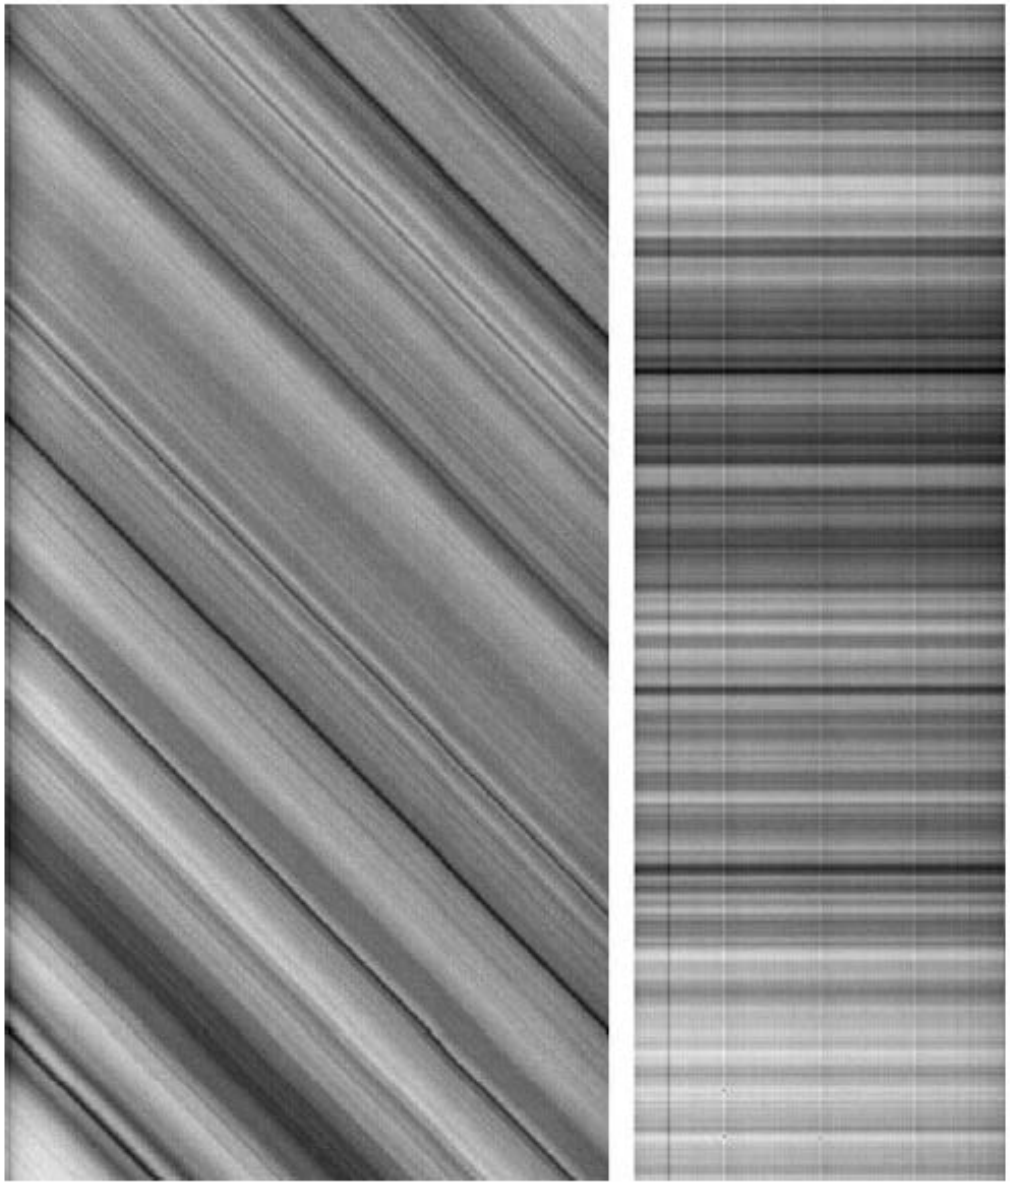

Fig. 5. HiRISE "flat-field" images produced by imaging when MRO was yawed 90° to the ground track. At left is a TDI 64 bin 2 single channel image. The slanted appearance of bright and dark bands is the result of the spacecraft motion during TDI integration, recording brightness variations on Mars that are smeared along each band. At right is a realigned TDI 128 bin 1 single-channel image useful for calibration.