Sensors & noise

CS 448A, Winter 2010

Marc Levoy
Computer Science Department
Stanford University

Camera pixel pipeline

sensor → analog to digital conversion (ADC) → processing: demosaicing, tone mapping & white balancing, denoising & sharpening, compression → storage

- every camera uses different algorithms
- the processing order may vary
- most of it is proprietary

Example pipeline

sensor → analog to digital conversion (ADC) → processing: demosaicing, tone mapping & white balancing, denoising & sharpening, compression → storage

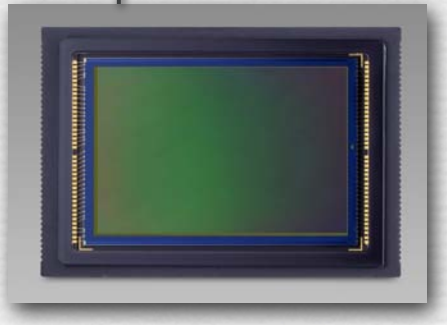

Canon 21 Mpix CMOS sensor

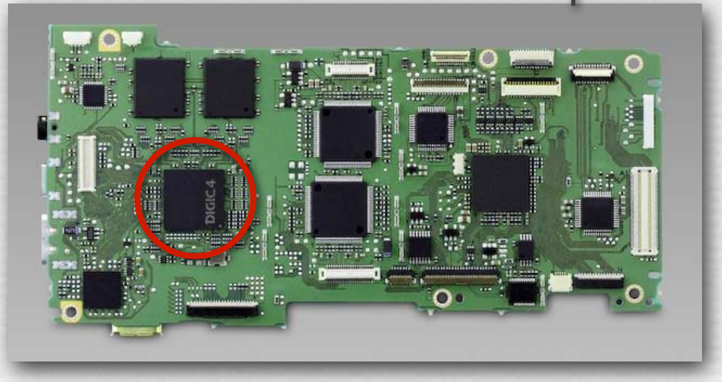

Canon DIGIC 4 processor

Compact Flash card

Example

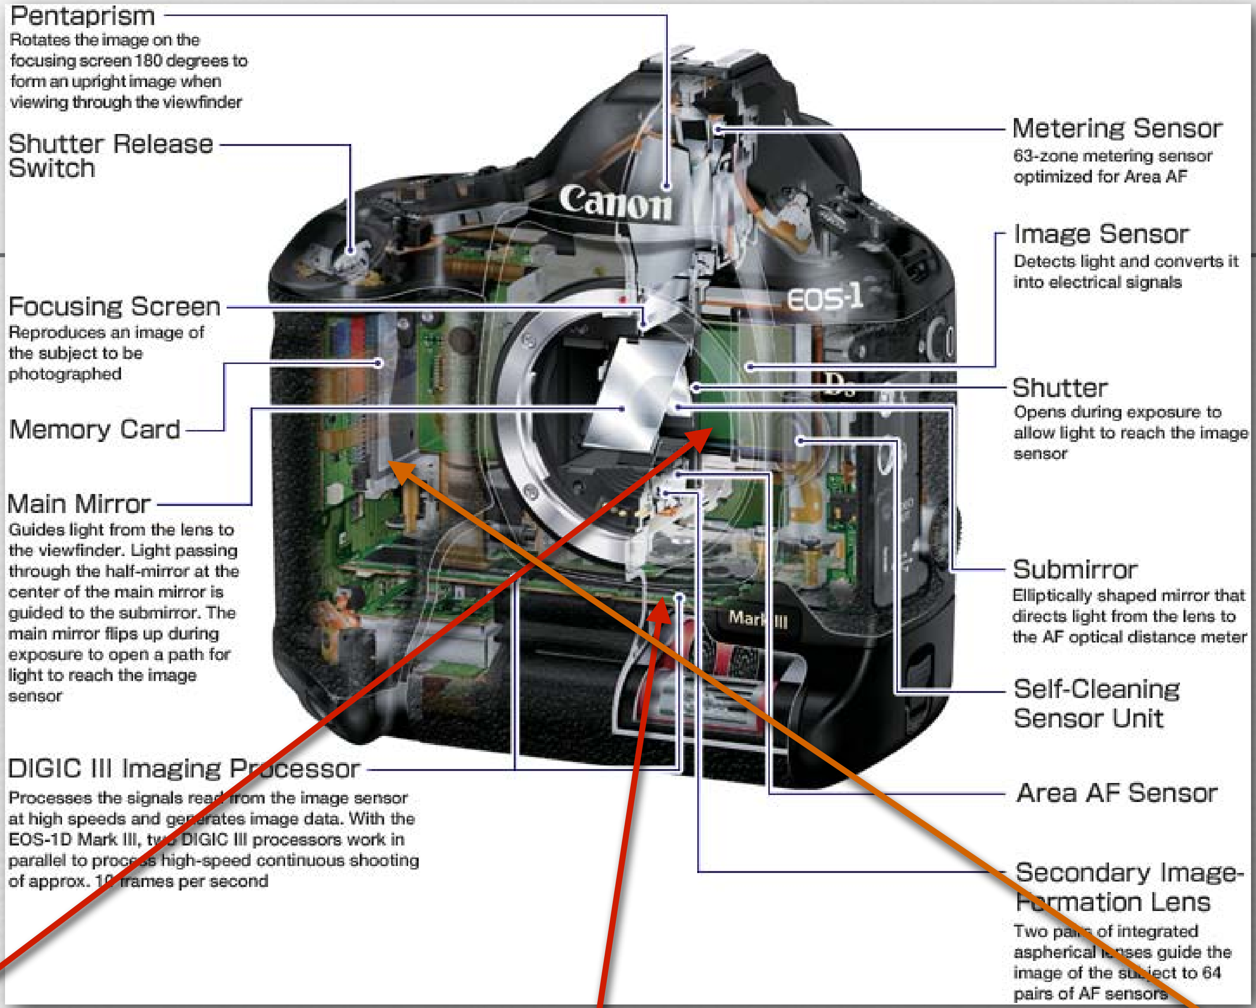

Pentaprism
Rotates the image on the focusing screen 180 degrees to form an upright image when viewing through the viewfinder

Shutter Release Switch

Focusing Screen
Reproduces an image of the subject to be photographed

Memory Card

Main Mirror
Guides light from the lens to the viewfinder. Light passing through the half-mirror at the center of the main mirror is guided to the submirror. The main mirror flips up during exposure to open a path for light to reach the image sensor

DIGIC III Imaging Processor
Processes the signals read from the image sensor at high speeds and generates image data. With the EOS-1D Mark III, two DIGIC III processors work in parallel to process high-speed continuous shooting of approx. 10 frames per second

Metering Sensor
63-zone metering sensor optimized for Area AF

Image Sensor
Detects light and converts it into electrical signals

Shutter
Opens during exposure to allow light to reach the image sensor

Submirror
Elliptically shaped mirror that directs light from the lens to the AF optical distance meter

Self-Cleaning Sensor Unit

Area AF Sensor

Secondary Image-Formation Lens
Two pairs of integrated aspherical lenses guide the image of the subject to 64 pairs of AF sensors

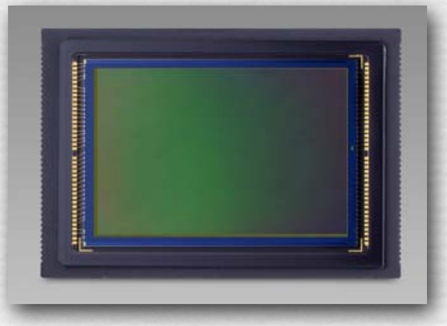

Canon 21 Mpix CMOS sensor

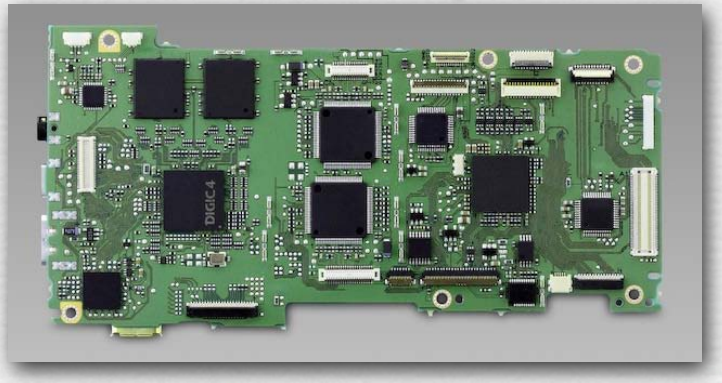

Canon DIGIC 4 processor

Compact Flash card

Outline (1st half of lecture)

- converting photons to charge
- getting the charge off the sensor
  - CCD versus CMOS
  - analog to digital conversion (ADC)
- supporting technology
  - microlenses
  - antialiasing filters

- -

- noise

The photoelectric effect

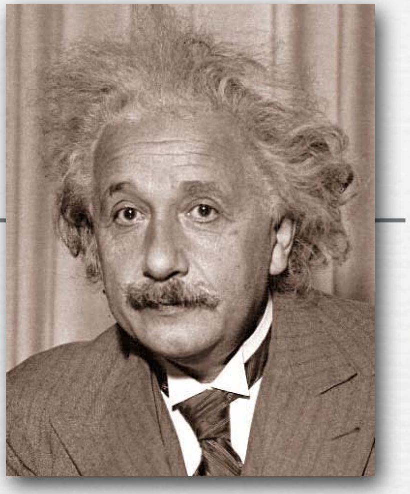

Albert Einstein

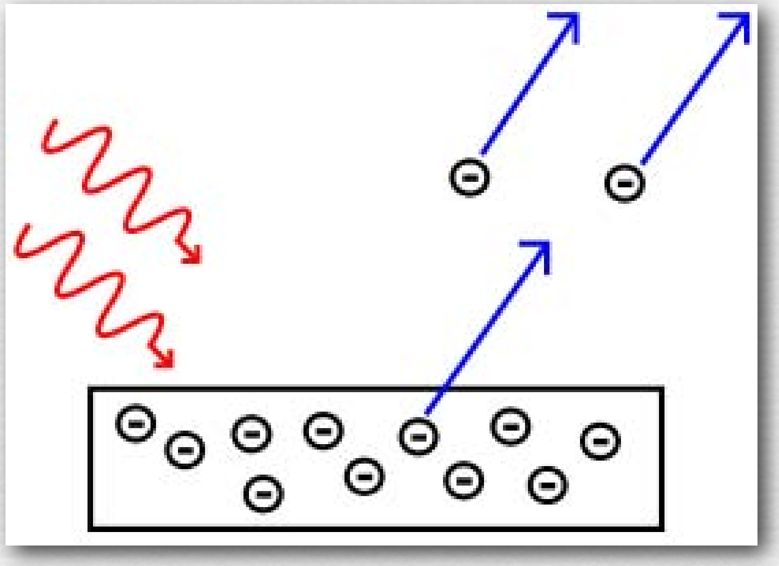

(wikipedia)

- when a photon strikes a material, an electron may be emitted
  - depends on wavelength, not intensity

$$E_{photon} = \frac{h \times c}{\lambda}$$

Quantum efficiency

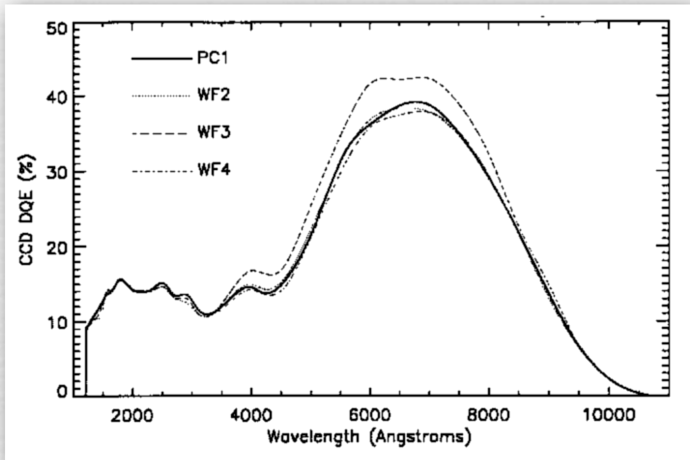

Hubble Space Telescope Camera 2

- not all photons will produce an electron
  - depends on quantum efficiency of the device

$$QE = \frac{\#\,electrons}{\#\,photons}$$

  - human vision: ~15%
  - typical digital camera: < 50%
  - best back-thinned CCD: > 90%

back-illuminated
CMOS (Sony)

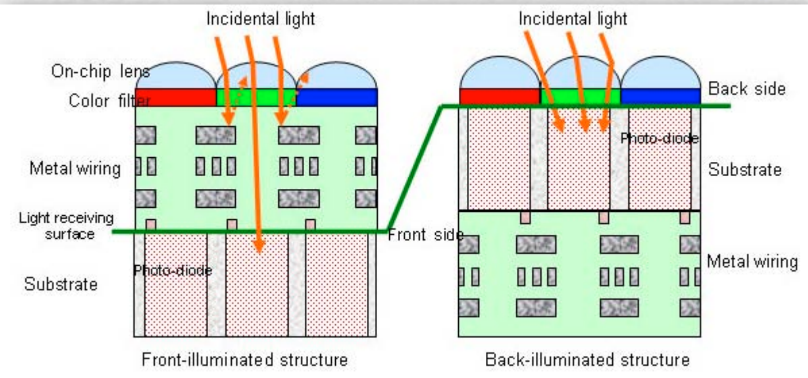

Pixel size

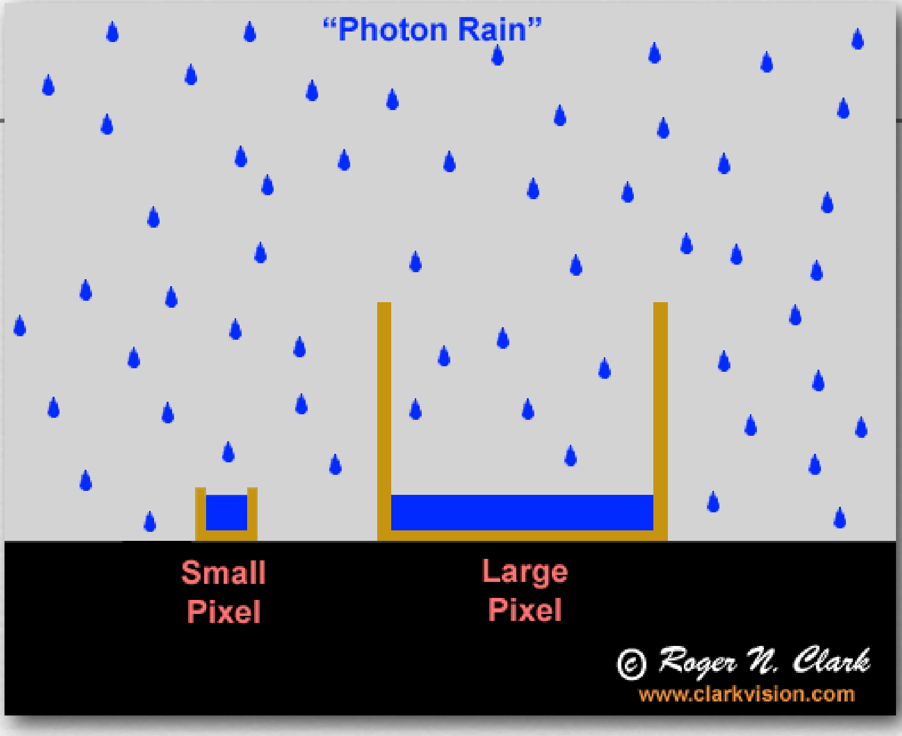

- ✦ the current from one electron is small (10-100 fA)
  - so integrate over space and time (pixel area × exposure time)
  - larger pixel × longer exposure means more accurate measure
- ✦ typical pixel sizes
  - casio EX-F1: $2.5\mu \times 2.5\mu = 6\mu^2$
  - Canon 5D II: $6.4\mu \times 6.4\mu = 41\mu^2$

Full well capacity

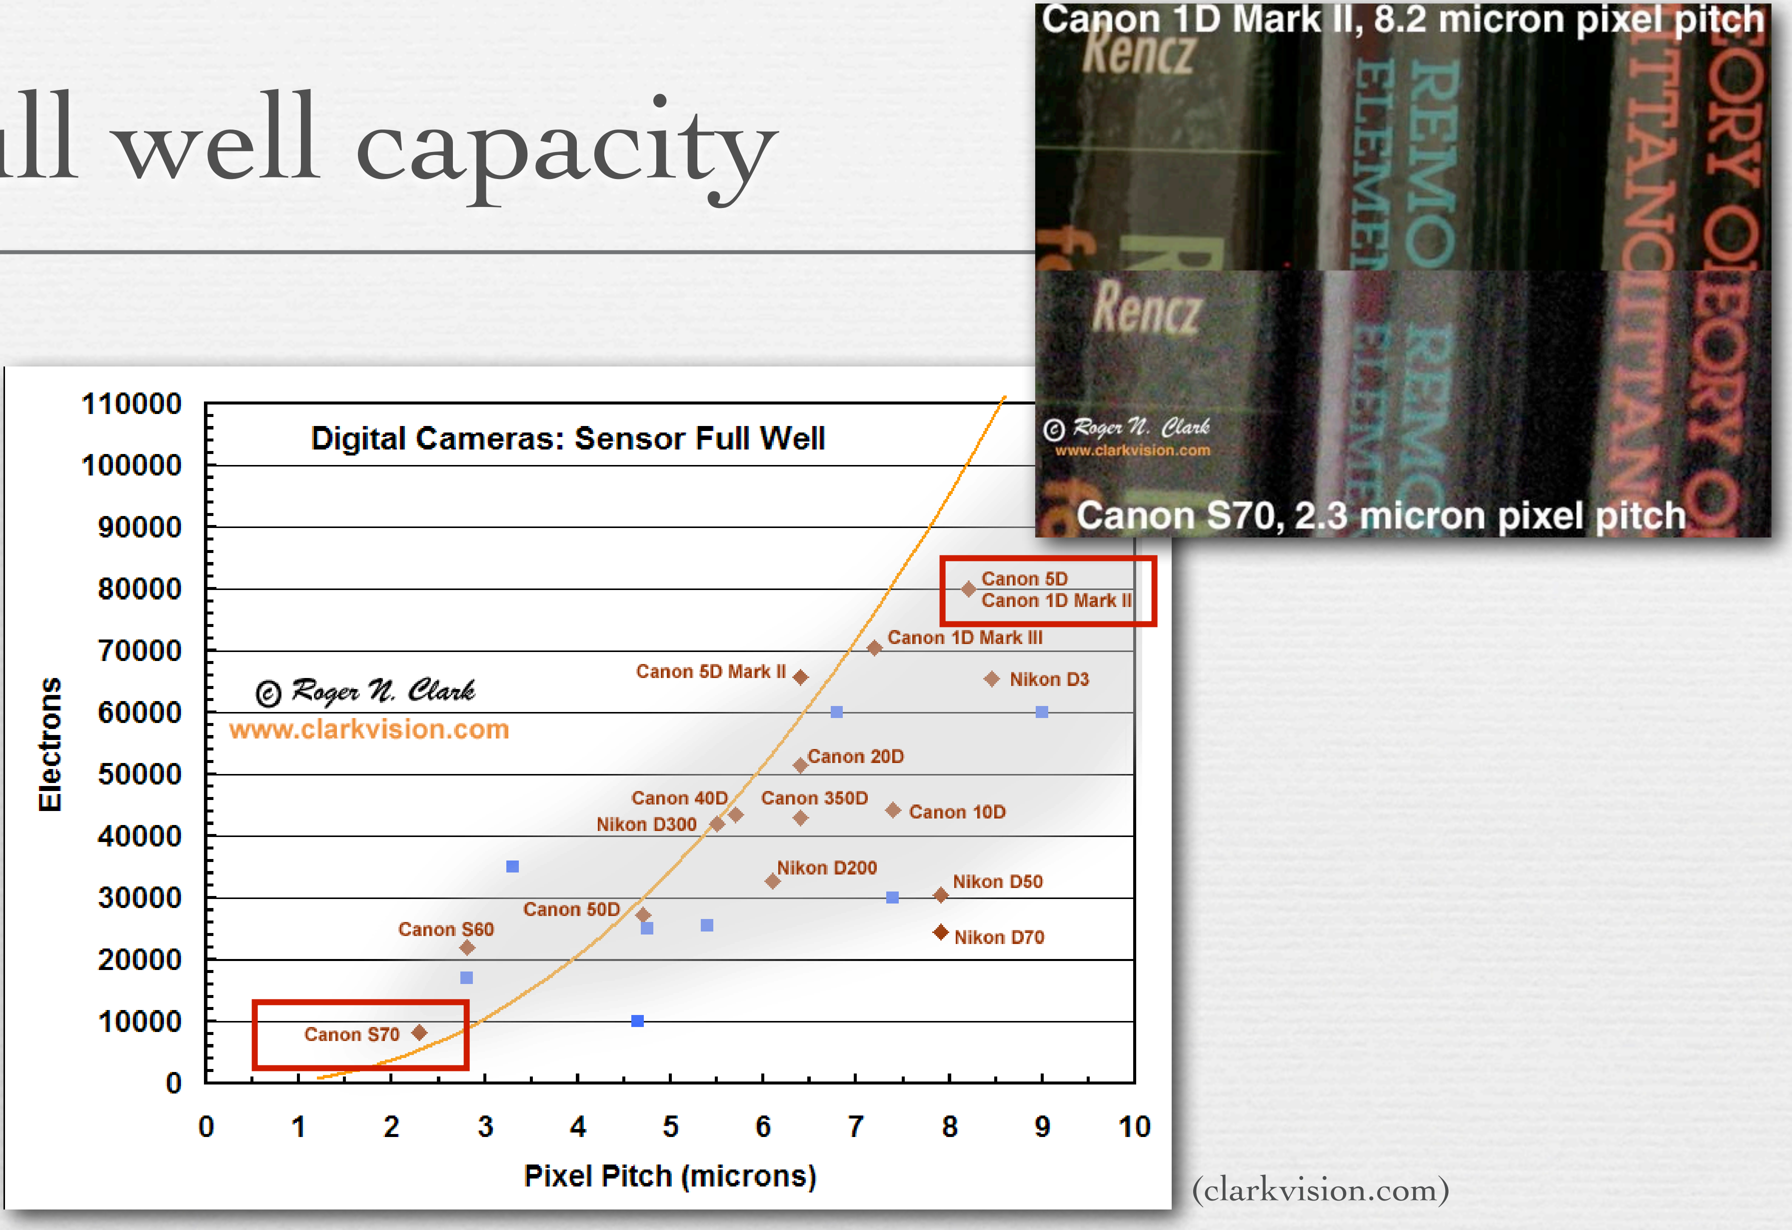

(clarkvision.com)

- ✦ too many photons causes *saturation*
  - larger capacity leads to higher *dynamic range*
  - but the *noise floor* is also a factor, as we'll see

Blooming

(ccd-sensor.de)

- charge spilling over to nearby pixels
  - can happen on CCD <u>and</u> CMOS sensors
  - don't confuse with glare or other image artifacts

Image artifacts can be hard to diagnose

(http://farm3.static.flickr.com/2102/2248725961_540be5f9af.jpg?v=0)

Q. Is this blooming?

CMOS versus CCD sensors

CMOS

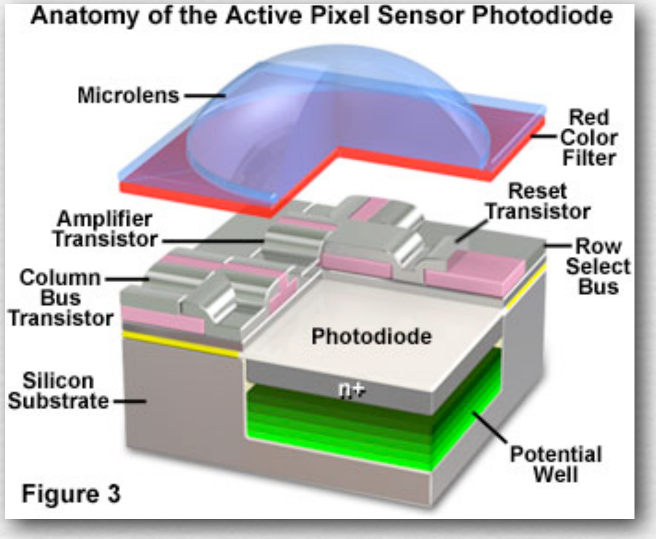

CCD

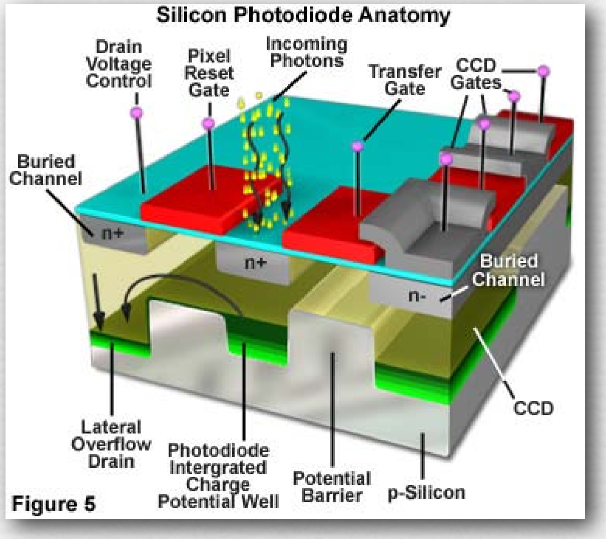

- CMOS = complementary metal-oxide semiconductor
  - an amplifier per pixel converts charge to voltage
  - low power, but noisy (but getting better)
- CCD = charge-coupled device
  - charge shifted along columns to an output amplifier
  - oldest solid-state image sensor technology
  - highest image quality, but not as flexible or cheap as CMOS

Gratuitous animation showing a CCD "bucket brigade" readout

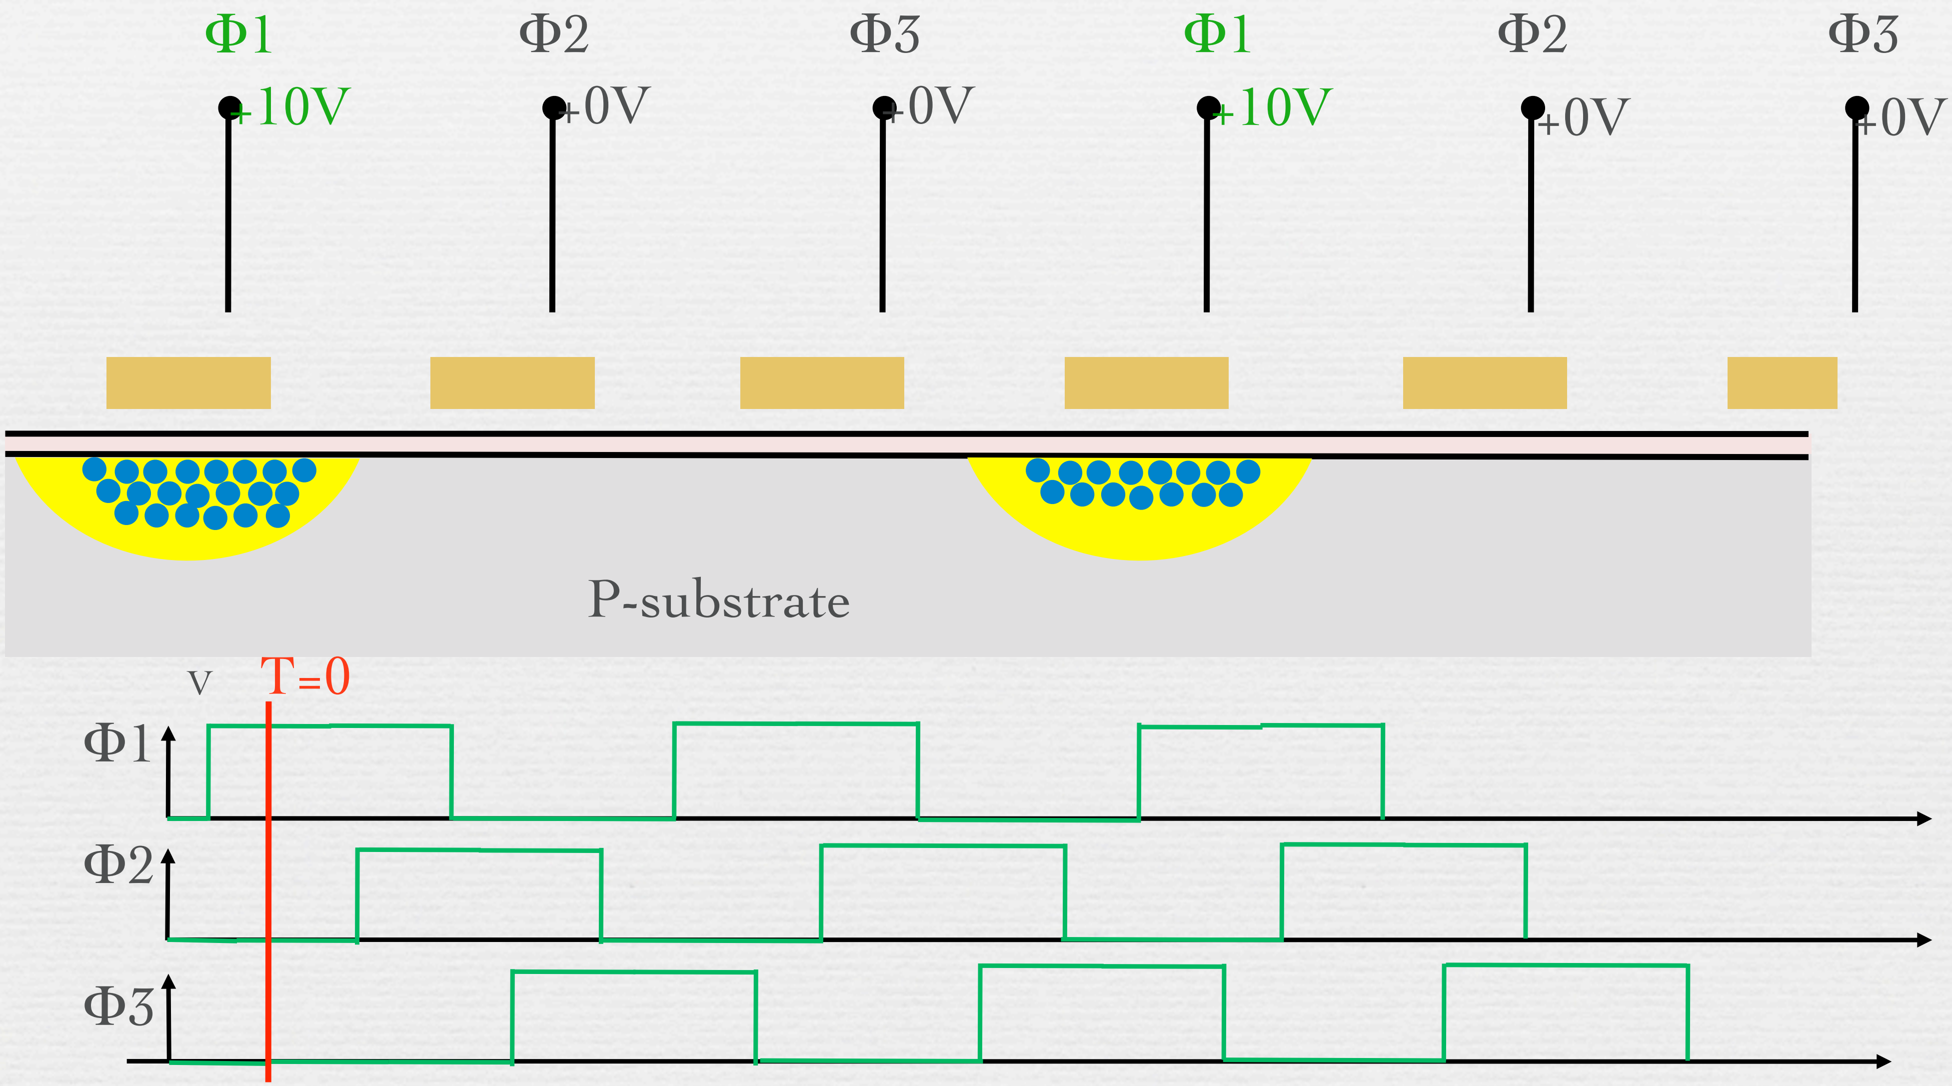

Gratuitous animation showing a CCD “bucket brigade” readout

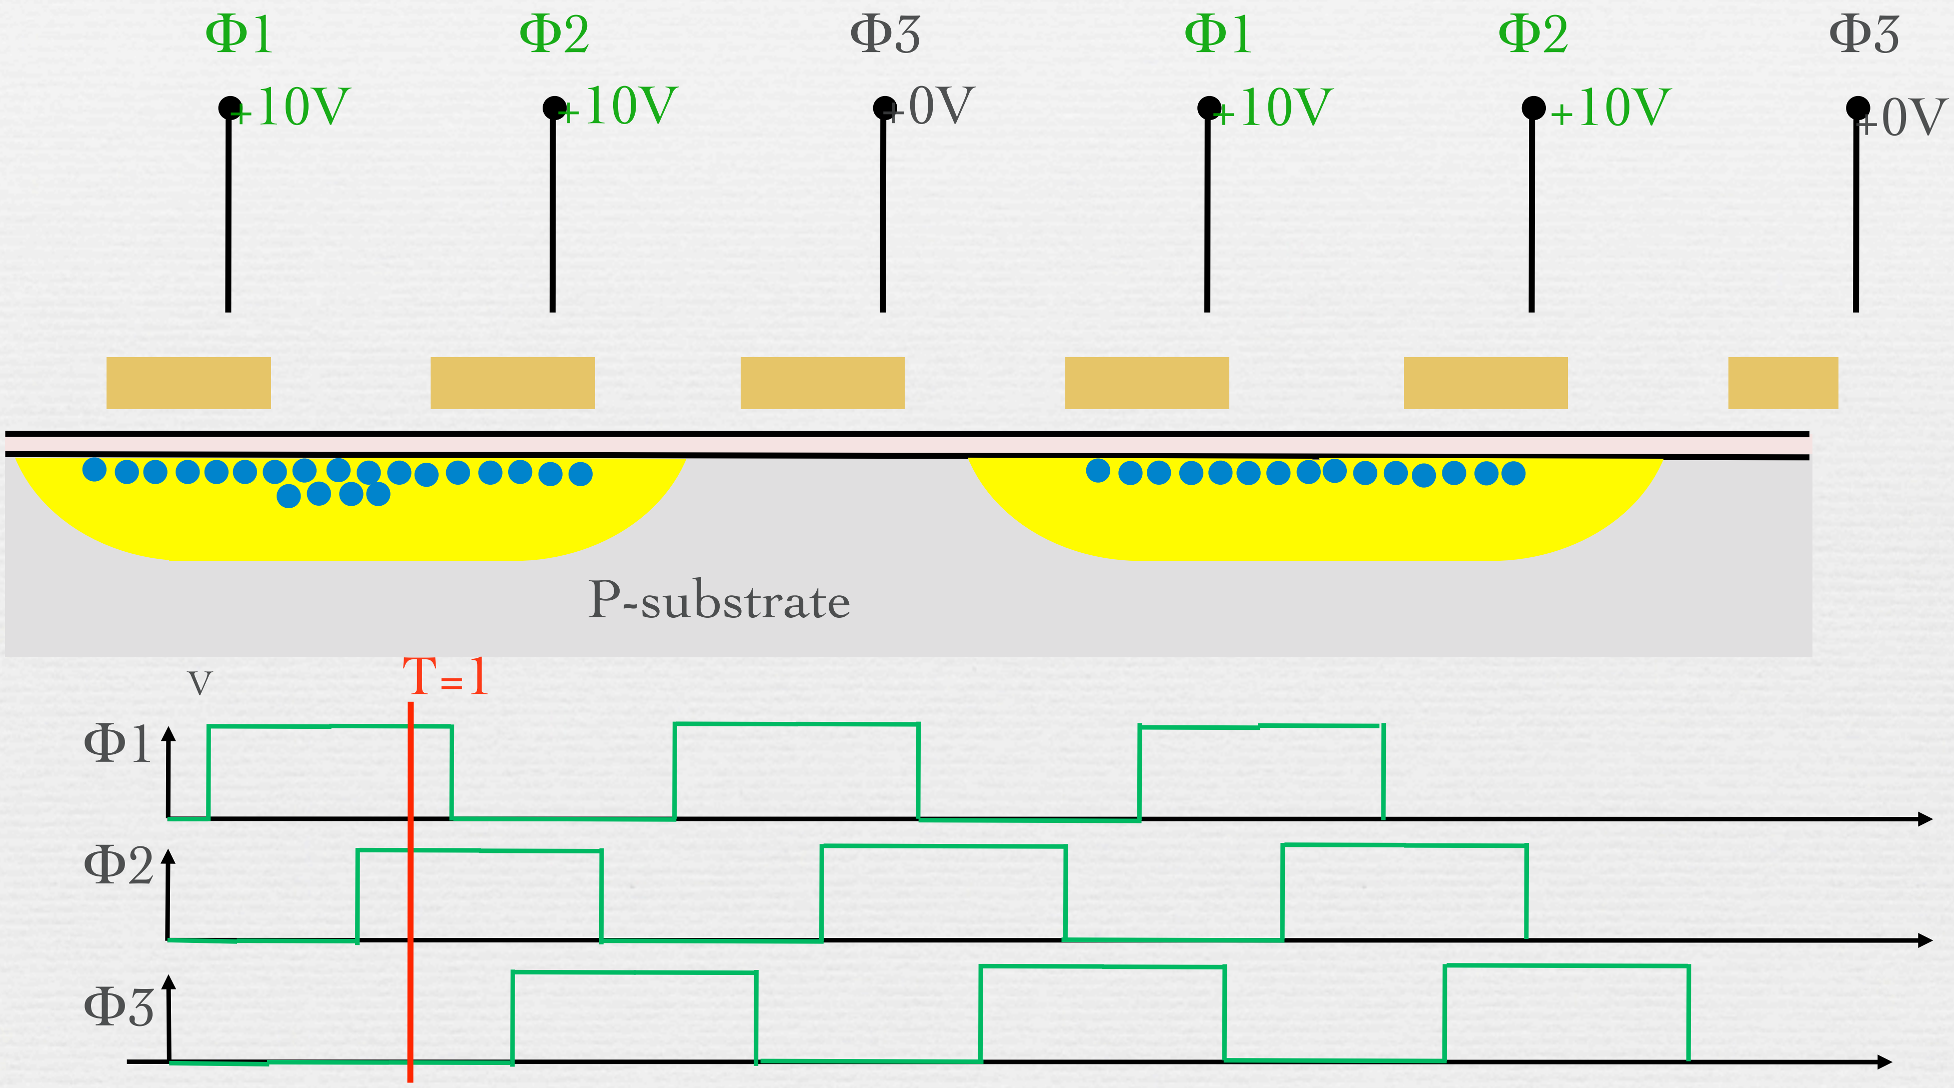

Gratuitous animation showing a CCD "bucket brigade" readout

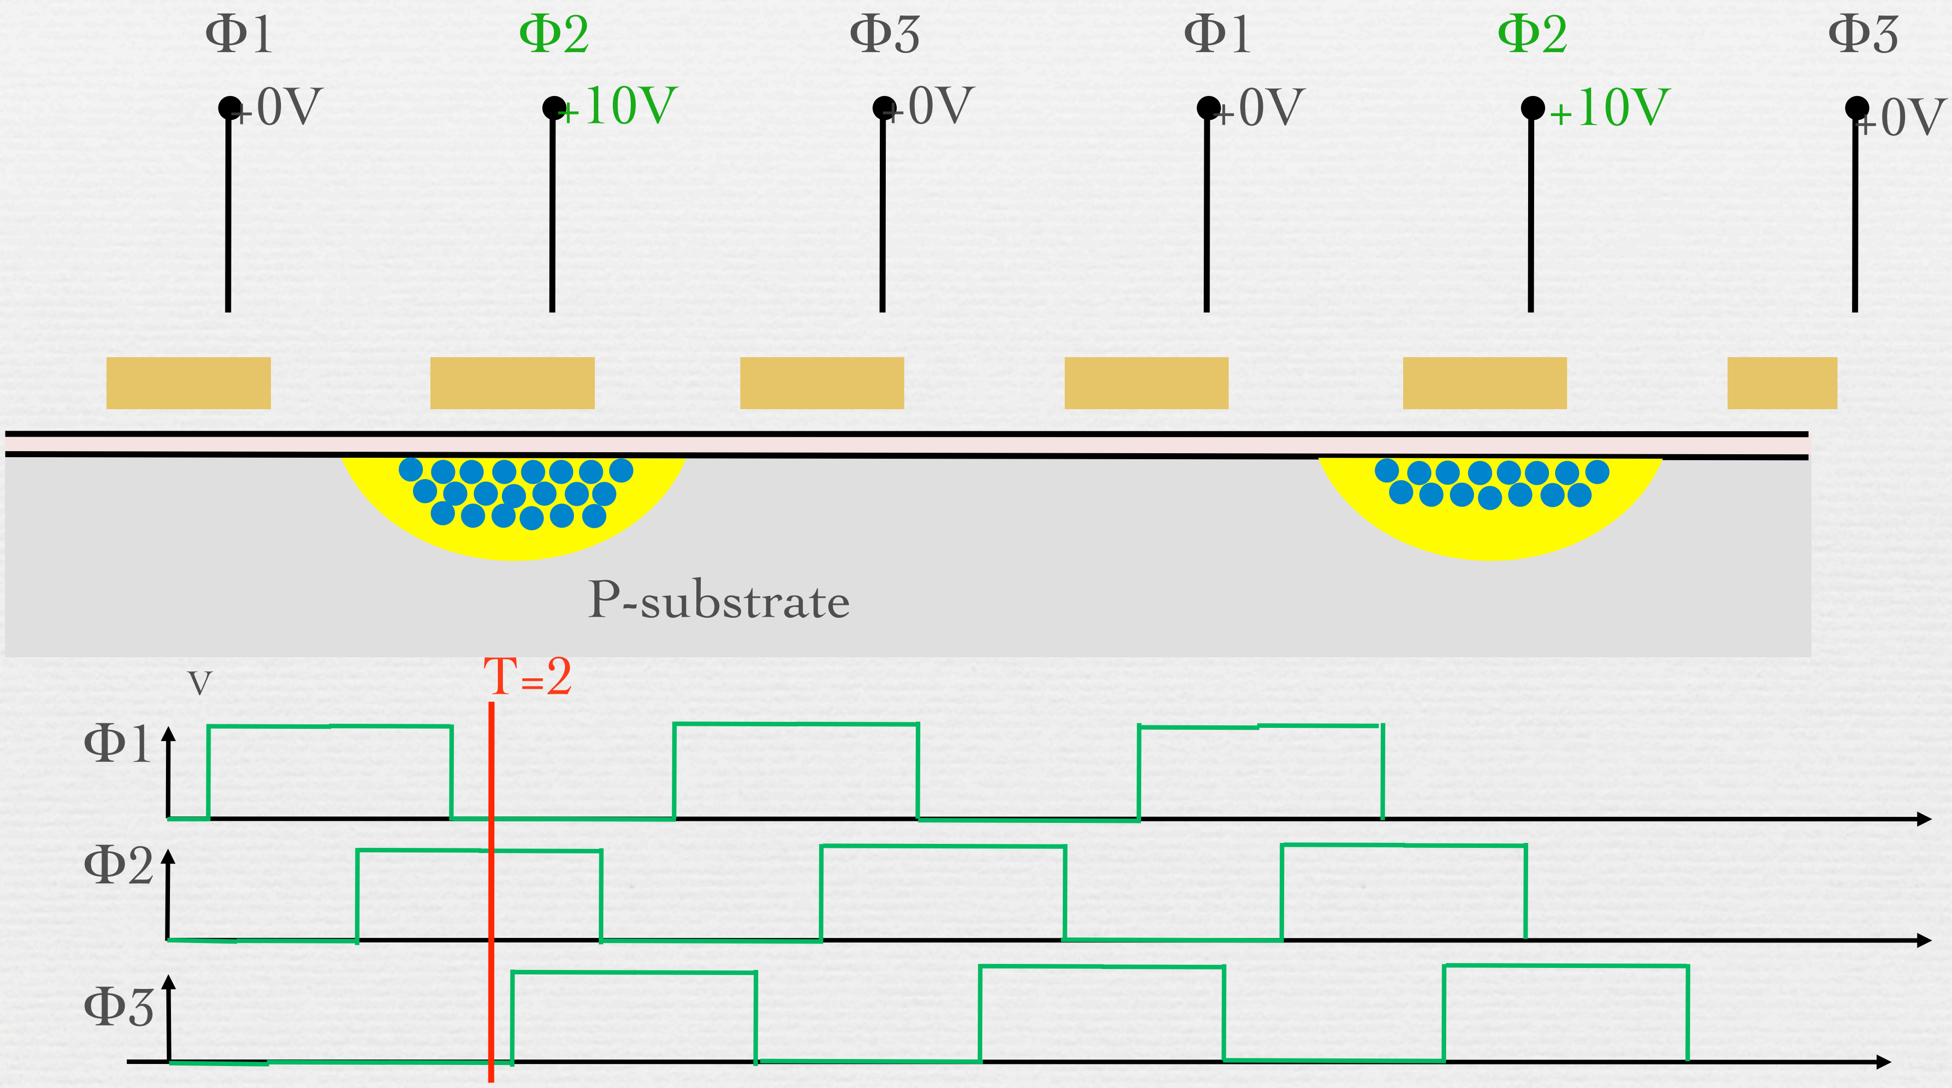

Gratuitous animation showing a CCD "bucket brigade" readout

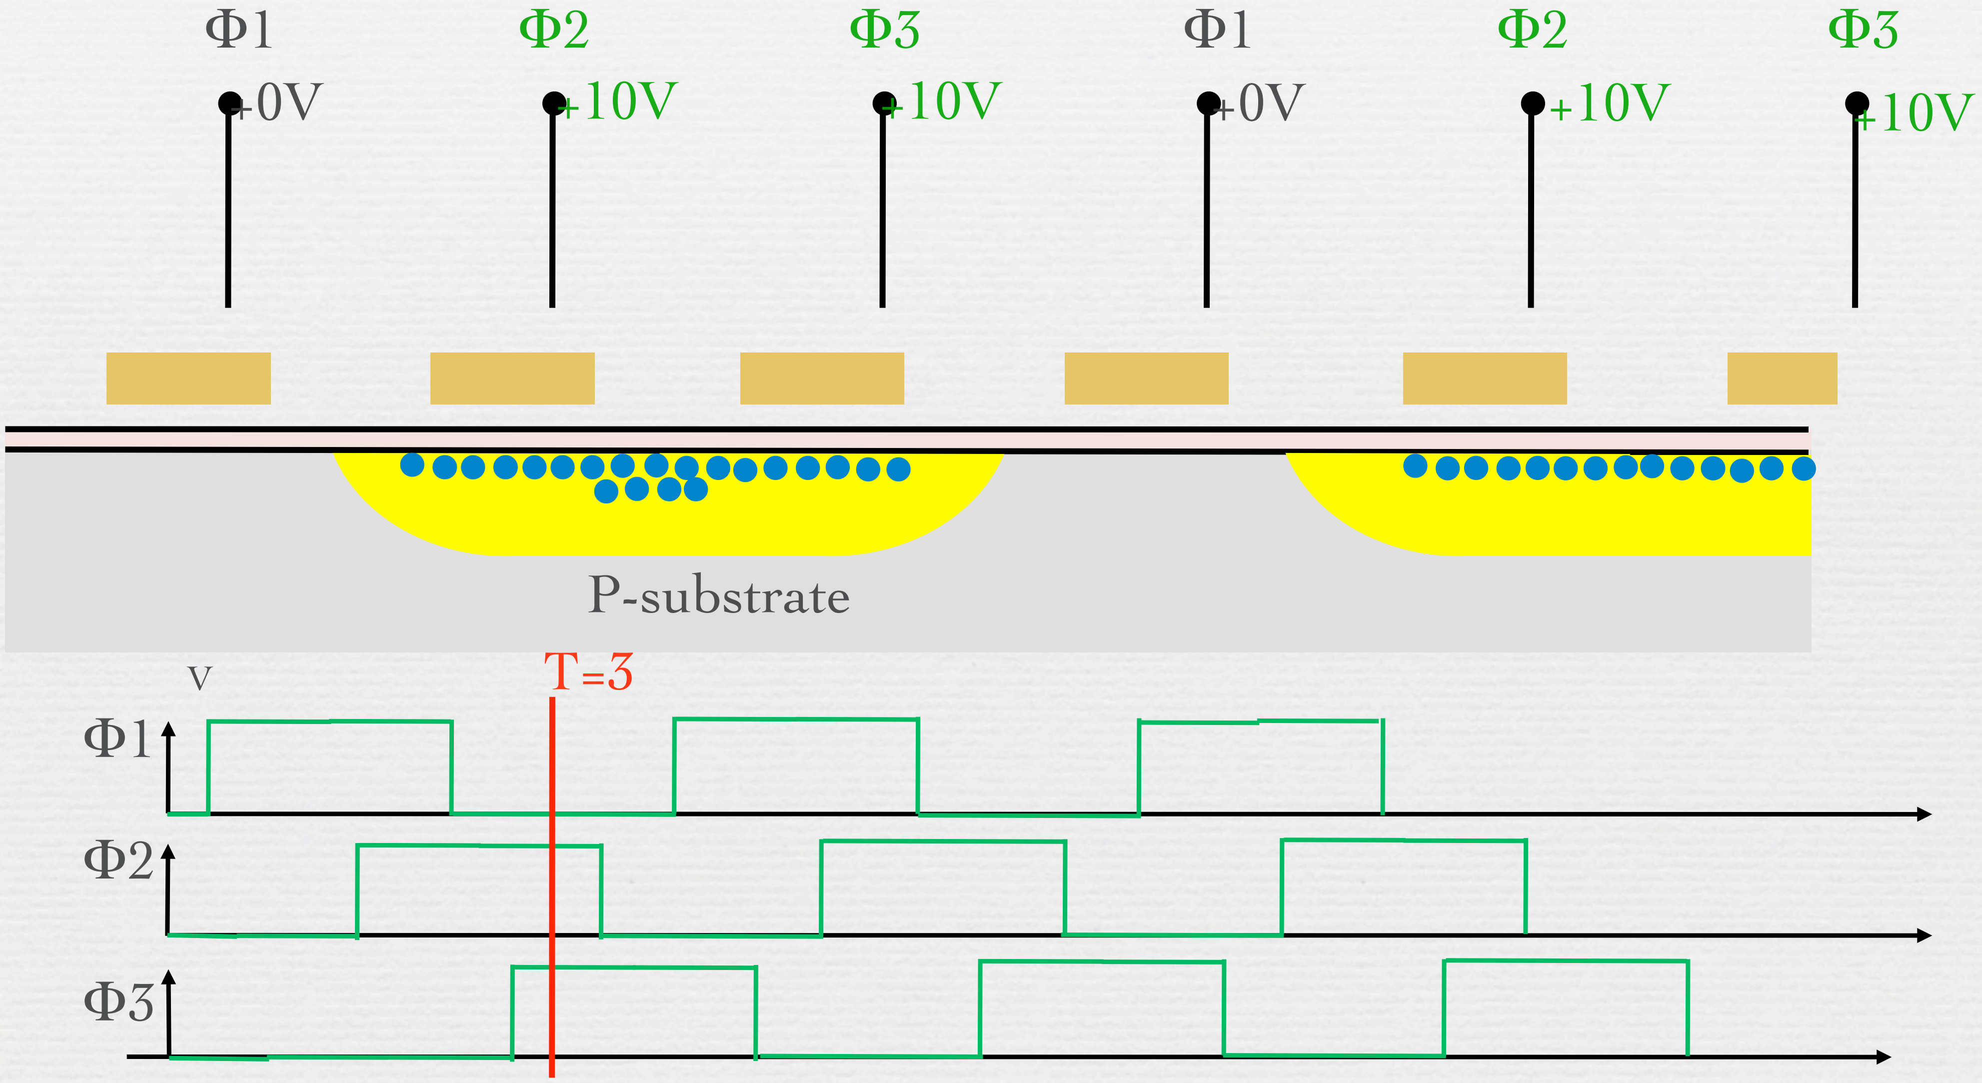

Gratuitous animation showing a CCD “bucket brigade” readout

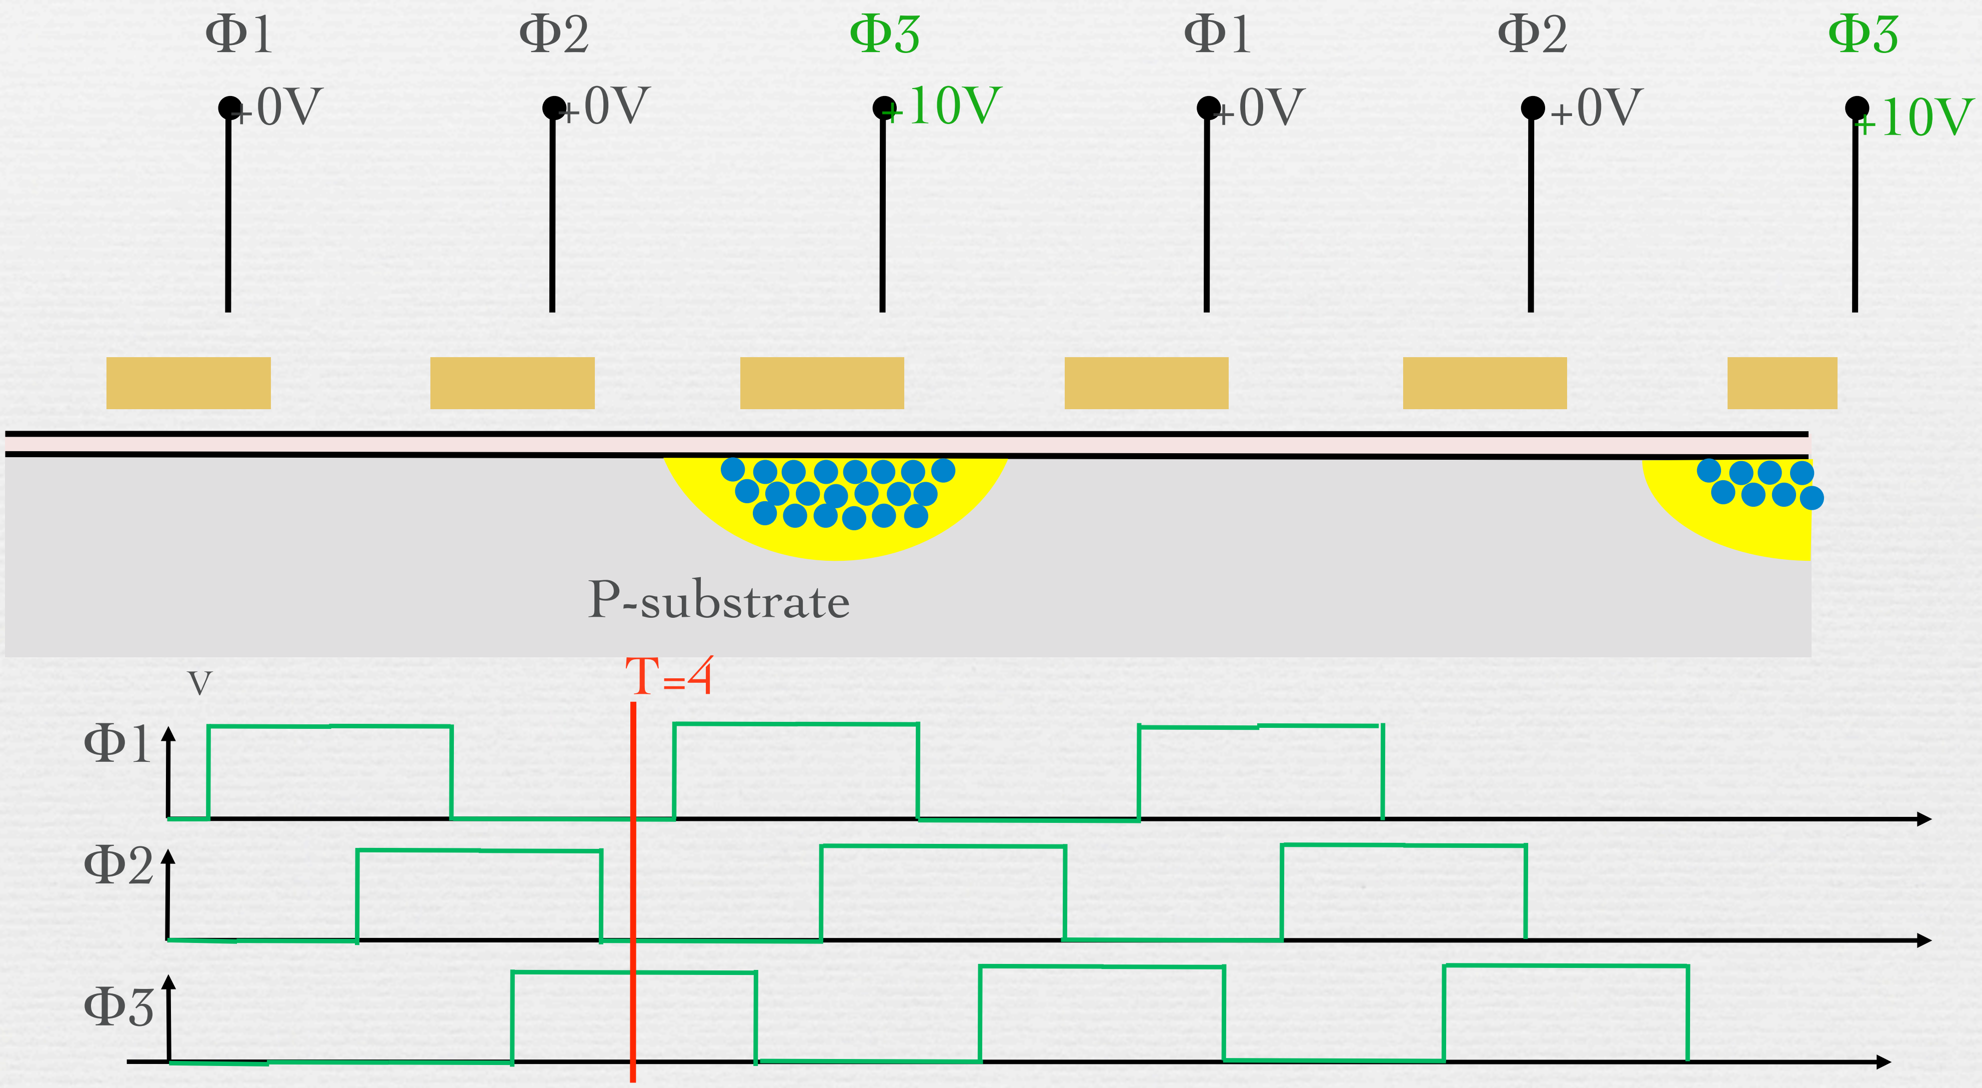

Gratuitous animation showing a CCD "bucket brigade" readout

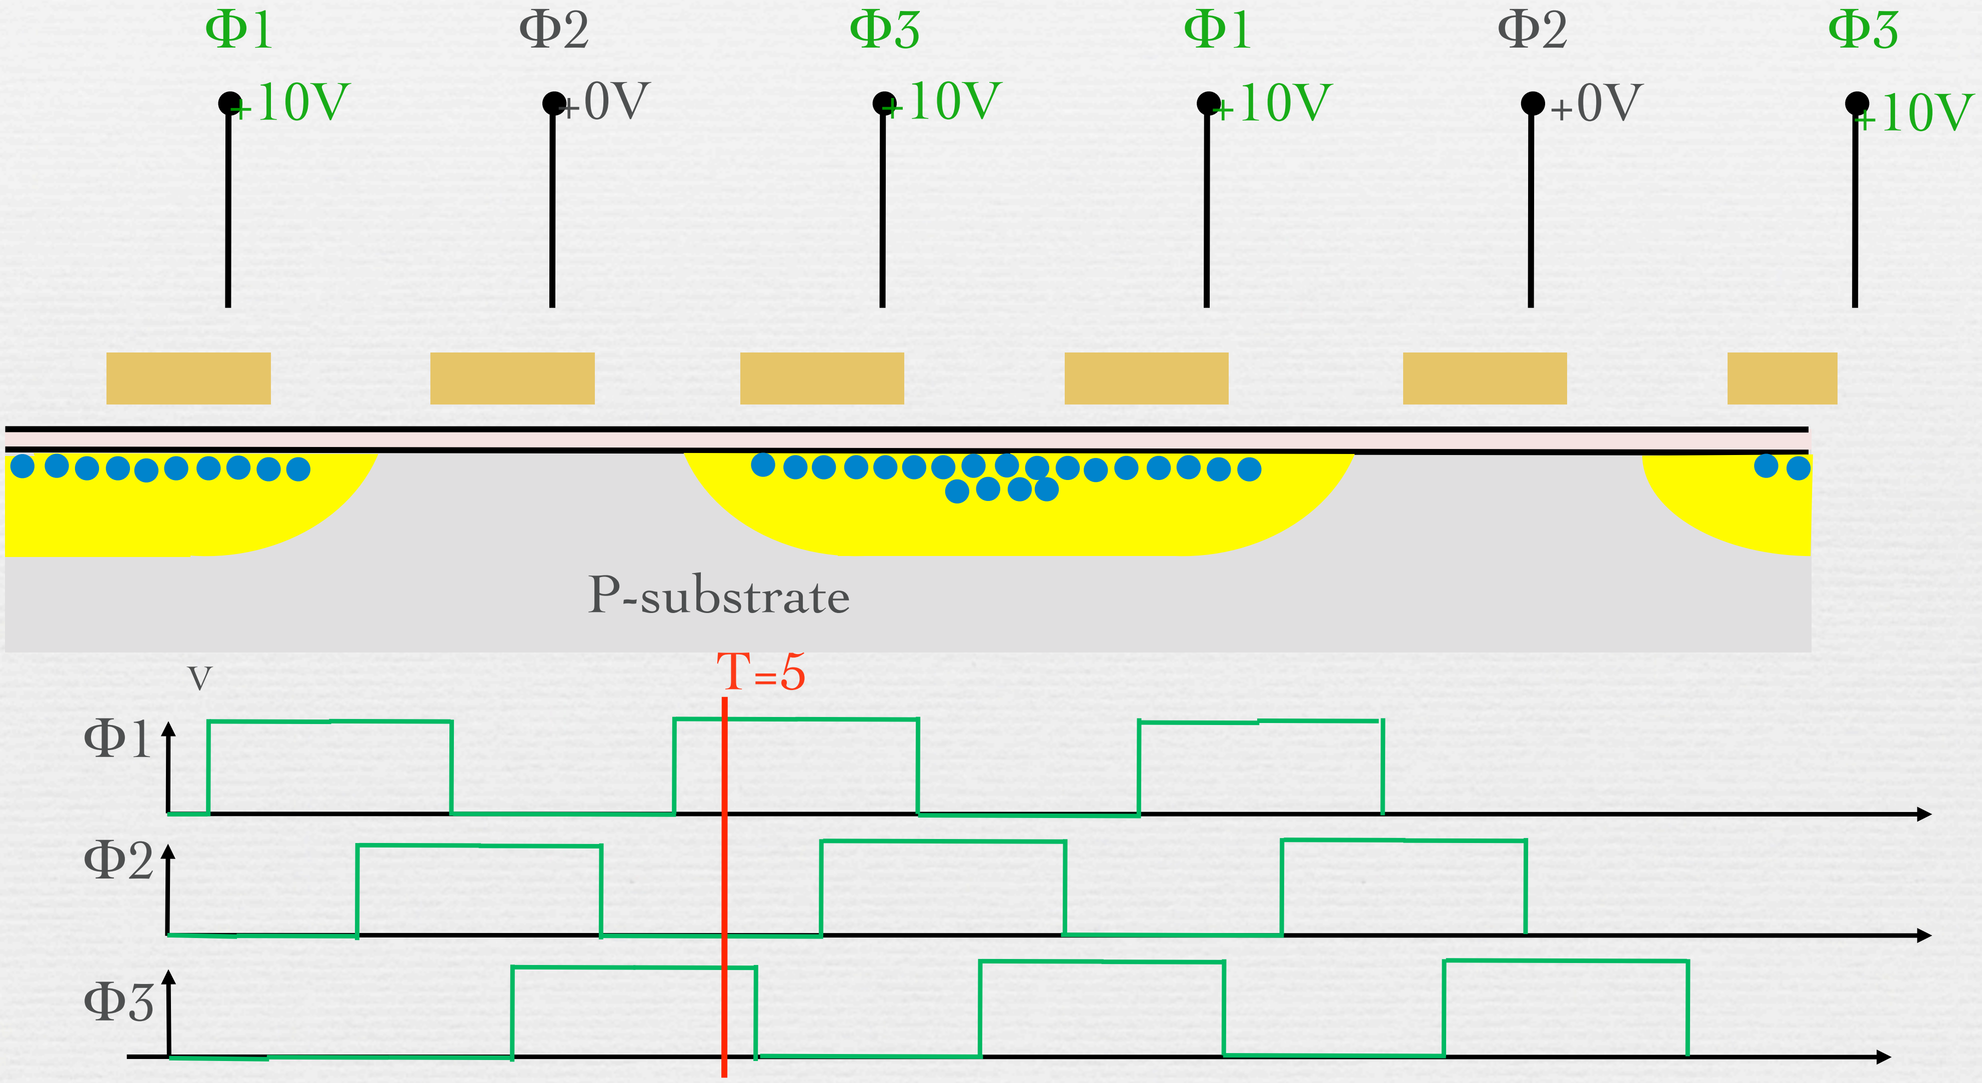

Smearing

(dvxuser.com)

CCD

CMOS

- side effect of bucket-brigade readout on CCD sensors
  - only happens if pixels saturate
  - doesn't happen on CMOS sensors

Analog to digital conversion (ADC)

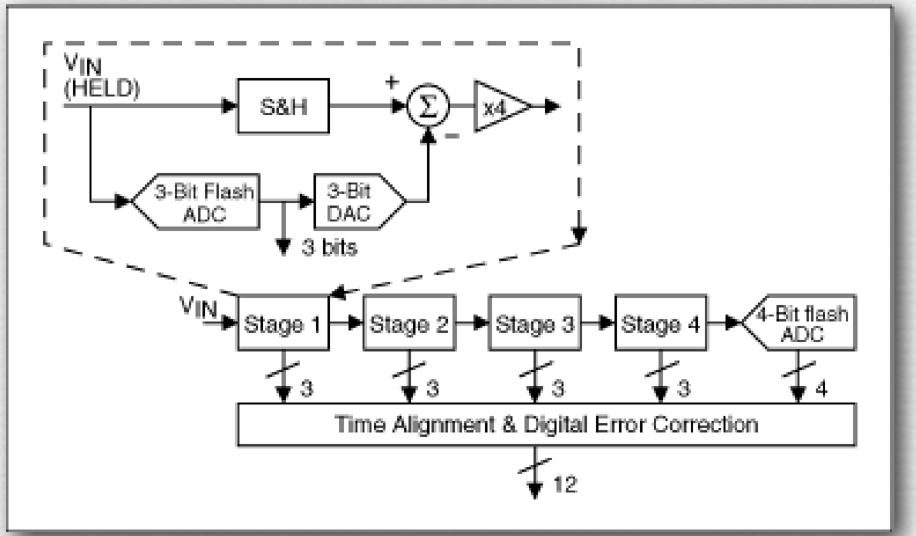

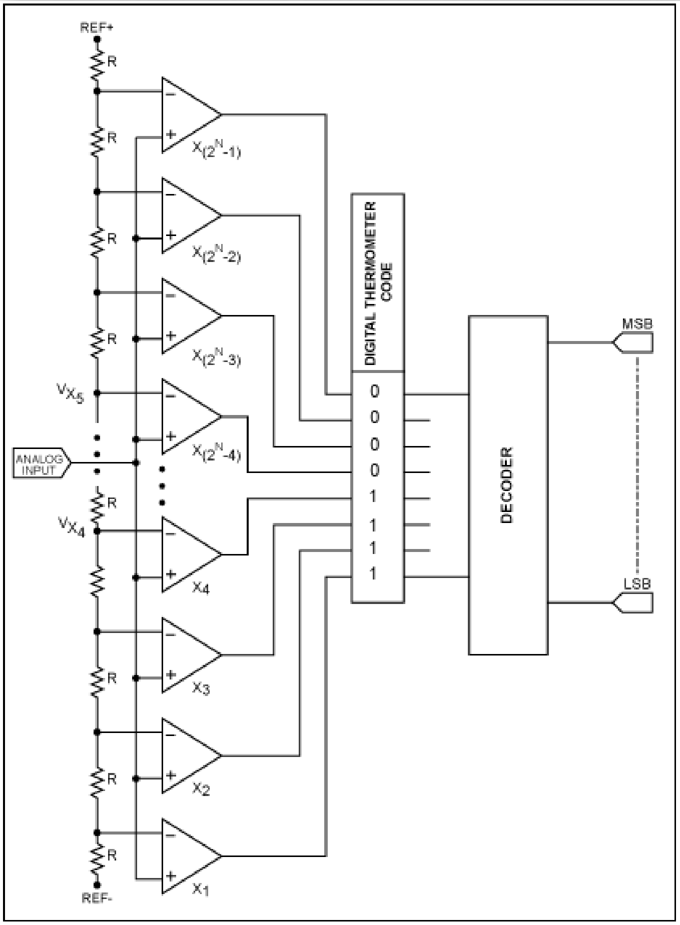

(maxim-ic.com)

- flash ADC
  - voltage divider → comparators → decoder
  - for n bits requires 2^n comparators
- pipelined ADC
  - 3-bit ADC → 3-bit DAC → compute residual → 4× → repeat
  - longer latency, but high throughput
  - some new sensors use an ADC per column

Fill factor

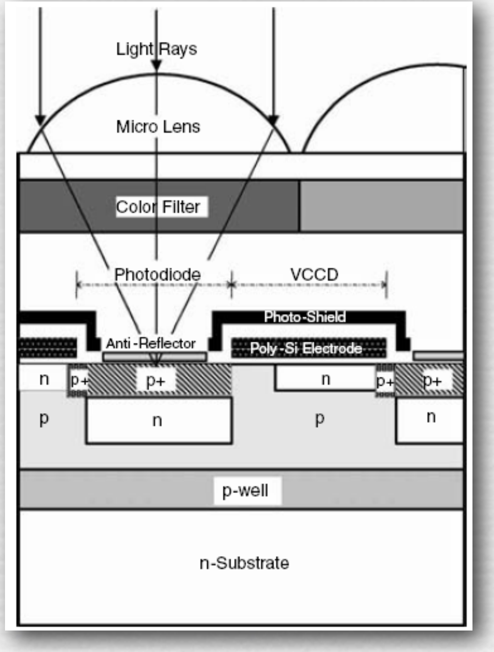

on a CCD sensor

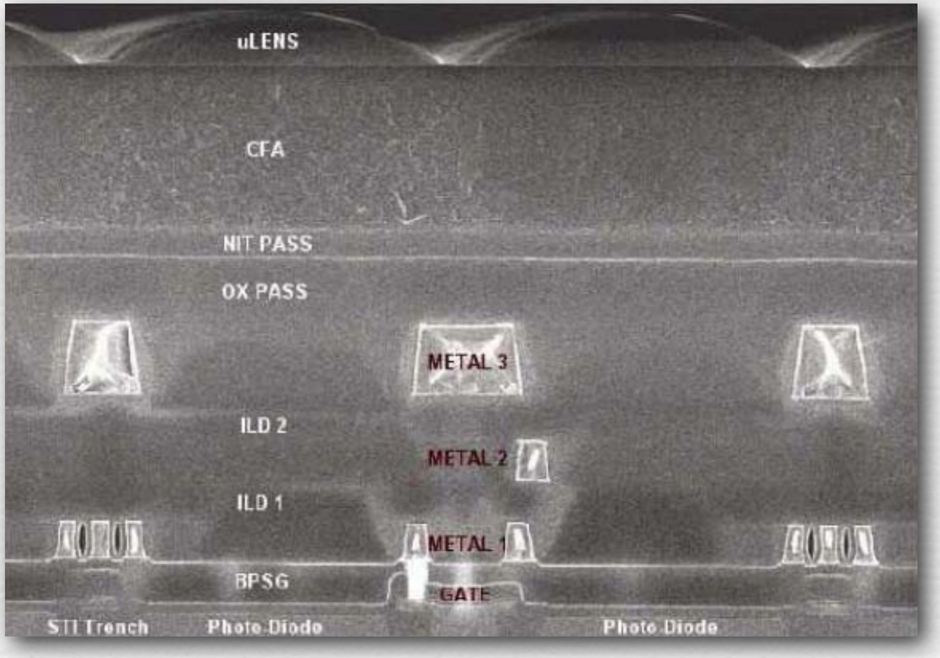

on a CMOS sensor

- ✦ fraction of sensor surface available to collect photons
  - can be improved using per-pixel microlenses

Q. An image sensor performs 2D sampling. What is the prefilter, with and without microlenses?

What per-pixel microlenses do

- integrating light over a pixel serves <u>two</u> functions
  - capturing more photons, to improve *dynamic range*
  - convolving the image with a prefilter, to avoid *aliasing*
- if the pixel is a rectangle, then this prefilter is a 2D rect

$$rect(x) \;=\; \Pi(x) \;=\; \begin{cases} 0 & if\ |x| > \frac{1}{2} \\ \frac{1}{2} & if\ |x| = \frac{1}{2} \\ 1 & if\ |x| < \frac{1}{2} \end{cases}$$

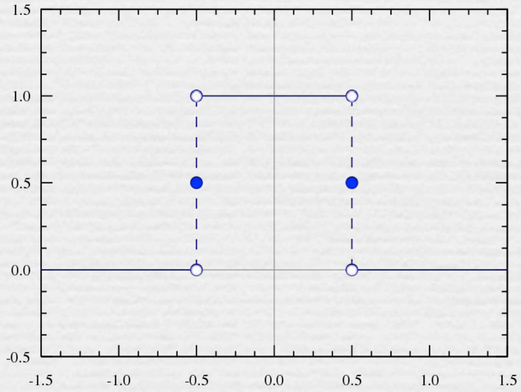

  - if only a portion of each pixel site is photo-sensitive, this rect doesn't span the spacing between pixels, so the prefilter is poor
- microlenses both gather more light and improve the prefilter
  - with microlenses, prefilter width roughly equals pixel spacing

Antialiasing filters

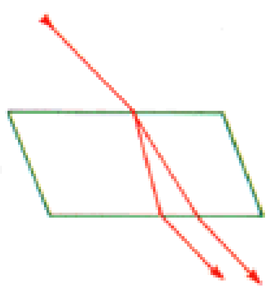

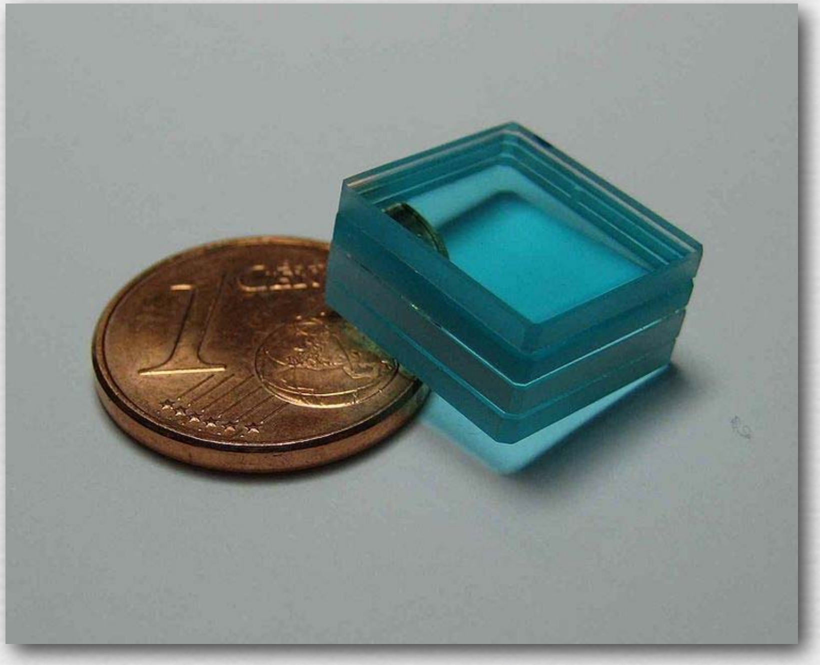

antialiasing filter

birefringence in a calcite crystal

- ✦ improves on non-ideal prefilter, even with microlenses
- ✦ typically two layers of birefringent material
  - splits 1 ray into 4 rays
  - operates like a 4-tap discrete convolution filter kernel!

Removing the antialiasing filter

- "hot rodding" your digital camera
  - $450 + shipping

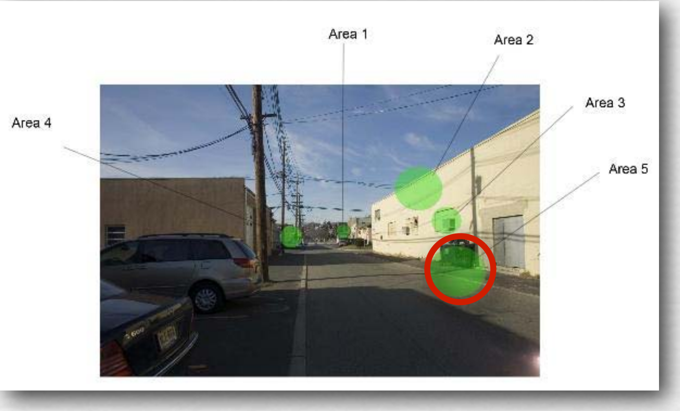

(maxmax.com)

anti-aliasing filter removed

normal

Removing the antialiasing filter

- “hot rodding” your digital camera
    - $450 + shipping

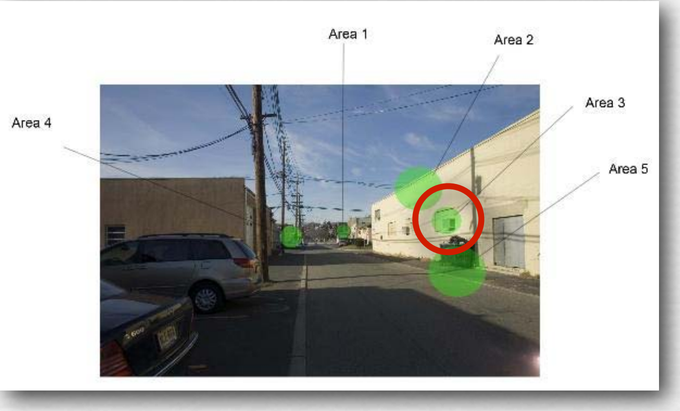

(maxmax.com)

anti-aliasing filter removed

normal

Outline (2nd half of lecture)

- examples of camera sensor noise
  - don't confuse it with JPEG compression artifacts
- probability, mean, variance, signal-to-noise ratio
- laundry list of noise sources
  - photon shot noise, dark current,
    hot pixels, fixed pattern noise, read noise
- SNR (again), quantization, dynamic range, bits per pixel
- ISO
- denoising
  - by aligning and averaging multiple shots
  - by image processing will be covered next week

Nokia N95 cell phone at dusk

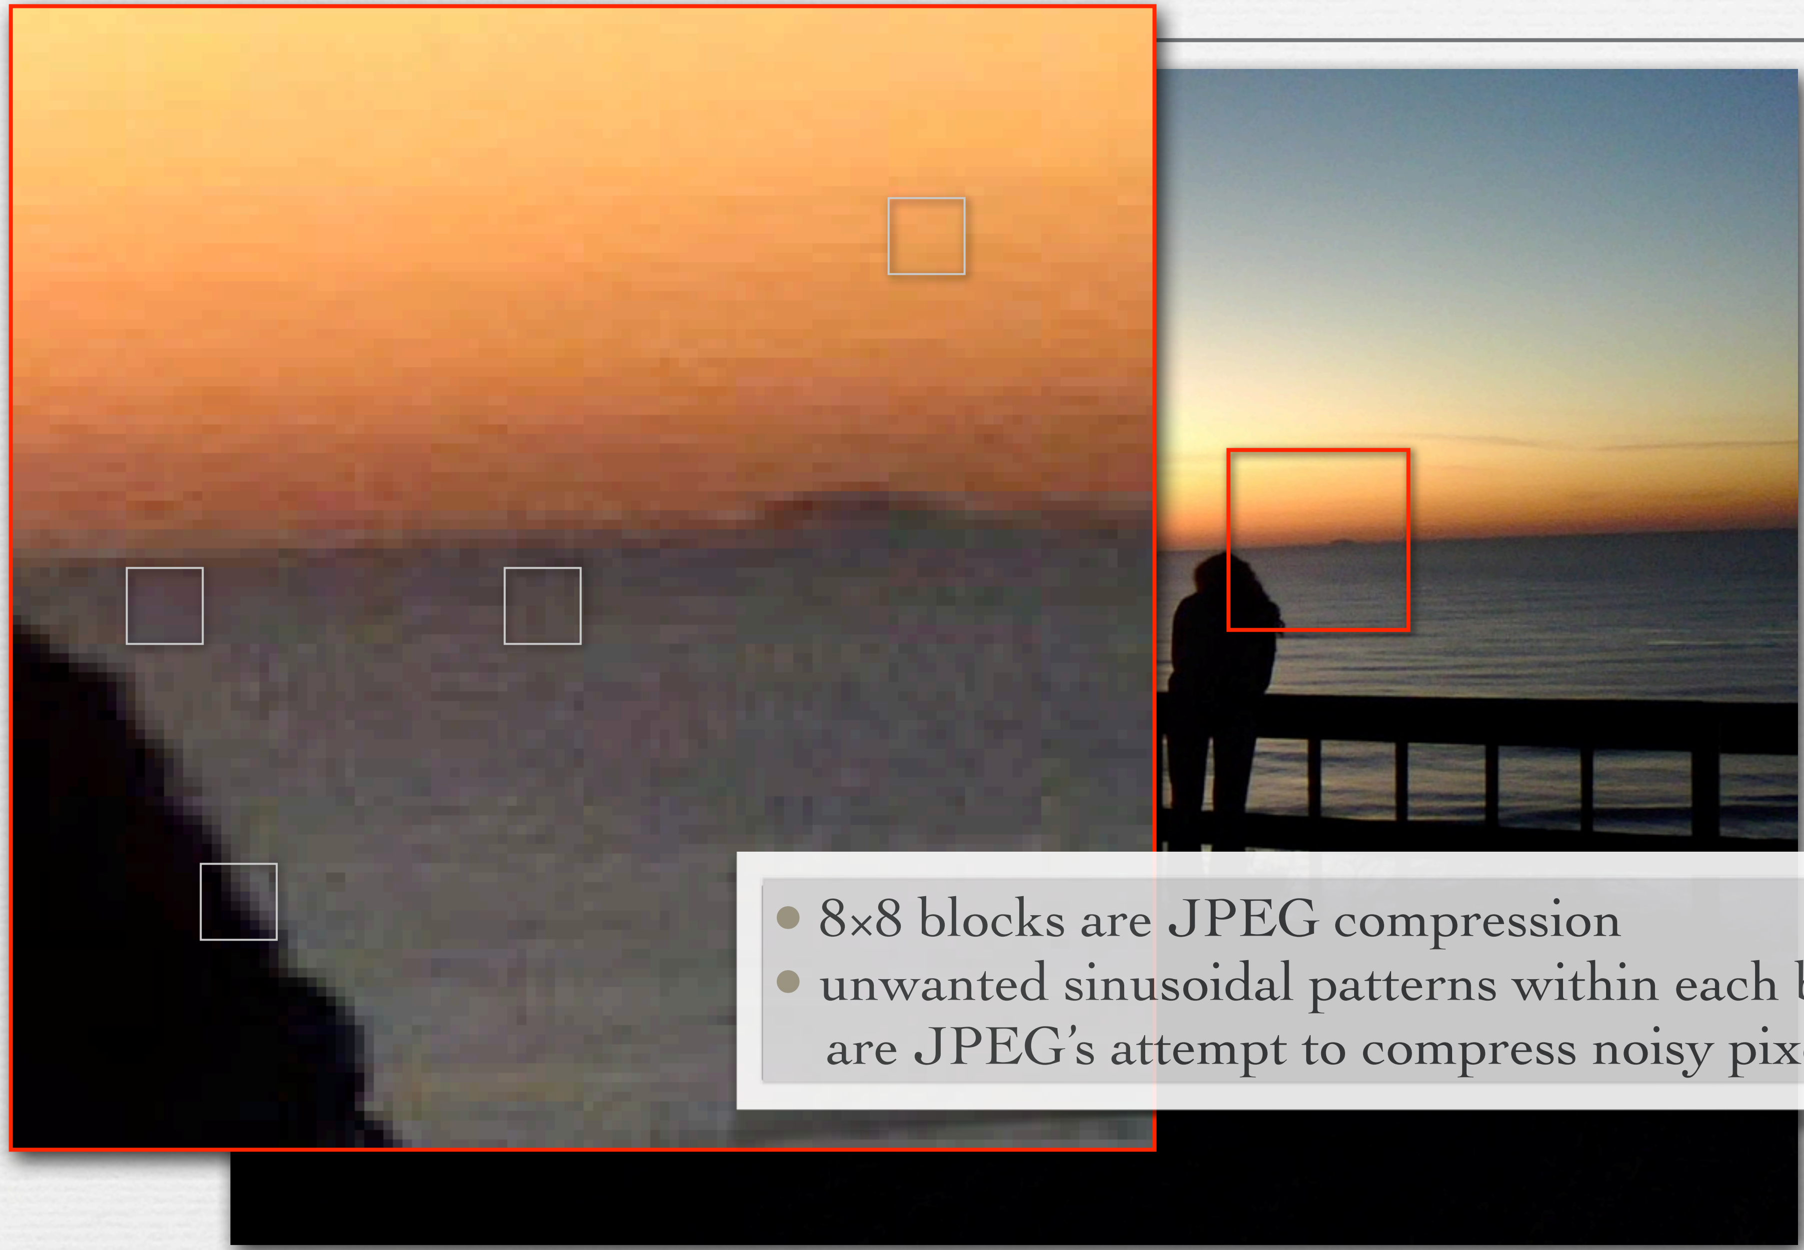

- 8×8 blocks are JPEG compression
- unwanted sinusoidal patterns within each block are JPEG's attempt to compress noisy pixels

Canon 5D II at noon

- ISO 200
- f/13.0
- 1/320 sec
- RAW w/o denoising

Canon post-processing

- ISO 200
- f/13.0
- 1/320 sec
- RAW w/o denoising

Canon 5D II at dusk

- ISO 6400
- f/4.0
- 1/13 sec
- RAW w/o denoising

Canon 5D II at dusk

Canon 5D II at dusk

Photon shot noise

- the number of photons arriving during an exposure varies from exposure to exposure and from pixel to pixel
- this number is governed by the Poisson distribution

Poisson distribution

- expresses the probability that a certain number of events will occur during an interval of time
- applicable to rare events that occur
  - with a known average rate, and
  - independently of the time since the last event
- if on average λ events occur in an interval of time, the probability *p* that *k* events occur instead is

$$p(k;\lambda) = \frac{\lambda^k e^{-\lambda}}{k!}$$

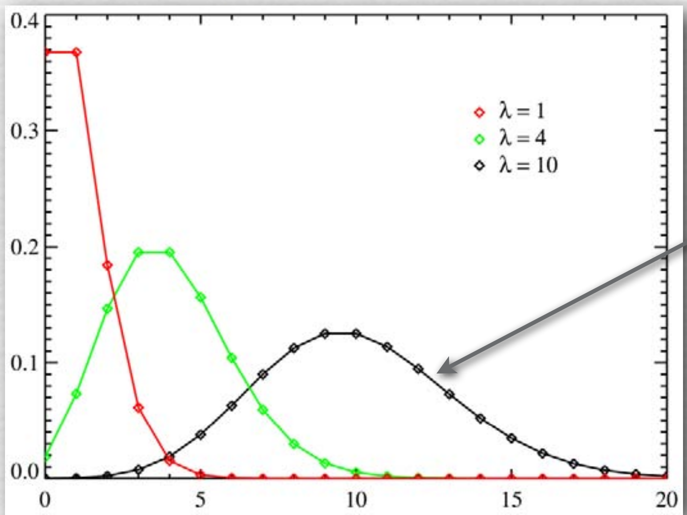

probability density function

Mean and variance

- the mean of a probability density function is

$$\mu = \int x\, p(x)\, dx$$

- the variance of a probability density function is

$$\sigma^2 = \int (x - \mu)^2\, p(x)\, dx$$

- the mean and variance of the Poisson distribution are

$$\mu = \lambda$$

$$\sigma^2 = \lambda$$

- the standard deviation is

$$\sigma = \sqrt{\lambda}$$

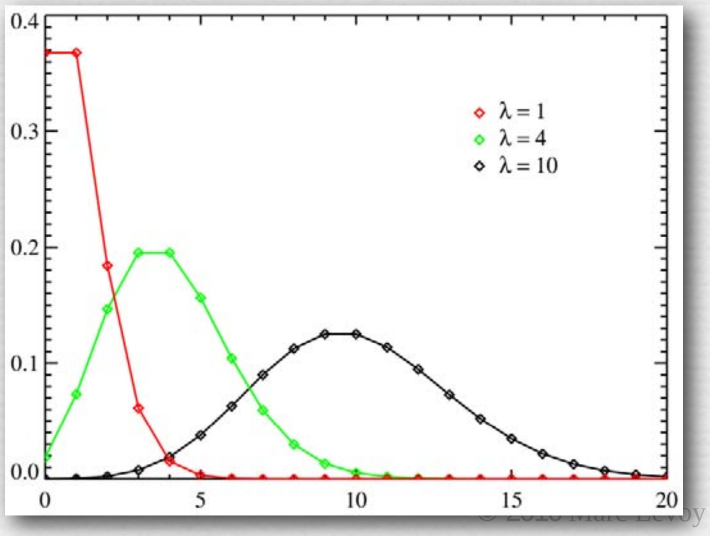

Deviation grows slower than the average.

Signal-to-noise ratio (SNR)

$$SNR = \frac{\text{mean pixel value}}{\text{standard deviation of pixel value}} = \frac{\mu}{\sigma}$$

$$SNR\ (\text{dB}) = 20 \log_{10}\left(\frac{\mu}{\sigma}\right)$$

- example
  - if SNR improves from 100:1 to 200:1,
    it improves $20 \log_{10}(200) - 20 \log_{10}(100) = +6$ dB

Photon shot noise (again)

- photons arrive in a Poisson distribution

$$\mu = \lambda$$

$$\sigma = \sqrt{\lambda}$$

- so

$$SNR = \frac{\mu}{\sigma} = \sqrt{\lambda}$$

- shot noise scales as square root of number of photons
- examples
  - doubling the width and height of a pixel increases its area by 4×, hence # of photons by 4×, hence SNR by 2× or +6 dB
  - opening the aperture by 1 f/stop increases the # of photons by 2×, hence SNR by √2 or +3 dB

Empirical example

- Kodak Q14 test chart

- Canon 10D, ISO 1600, crop from recorded image

(http://www.imatest.com/docs/noise.html)

- SNR improves with increasing signal

Dark current

- electrons dislodged by random thermal activity
- increases linearly with exposure time
- increases exponentially with temperature
- varies across sensor, and includes its own shot noise

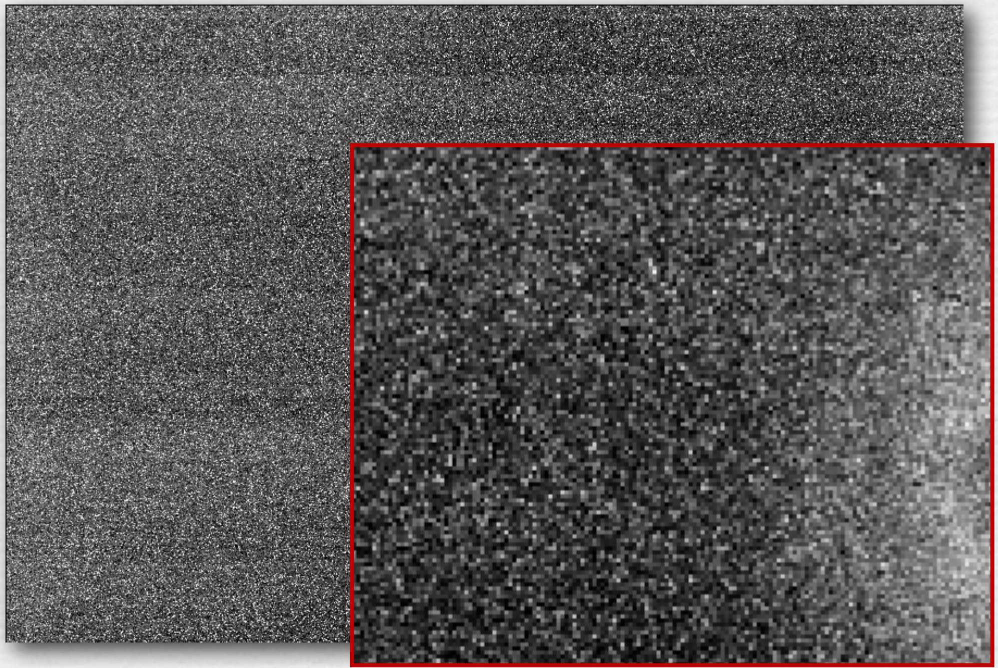

(http://theory.uchicago.edu/~ejm/pix/20d/tests/noise/)

Canon 20D, 612 sec exposure

Time exposures in astonomy

Lee Frost, star trails

(Palomar 200-inch)

- 30-minute exposure (on film)
- telescopes can rotate to avoid smearing stars
- What is the unmoving star in the middle?

Hot pixels

- electrons leaking into well due to manufacturing defects
- increases linearly with exposure time
- increases with temperature, but hard to model
- changes over time, and every camera has them

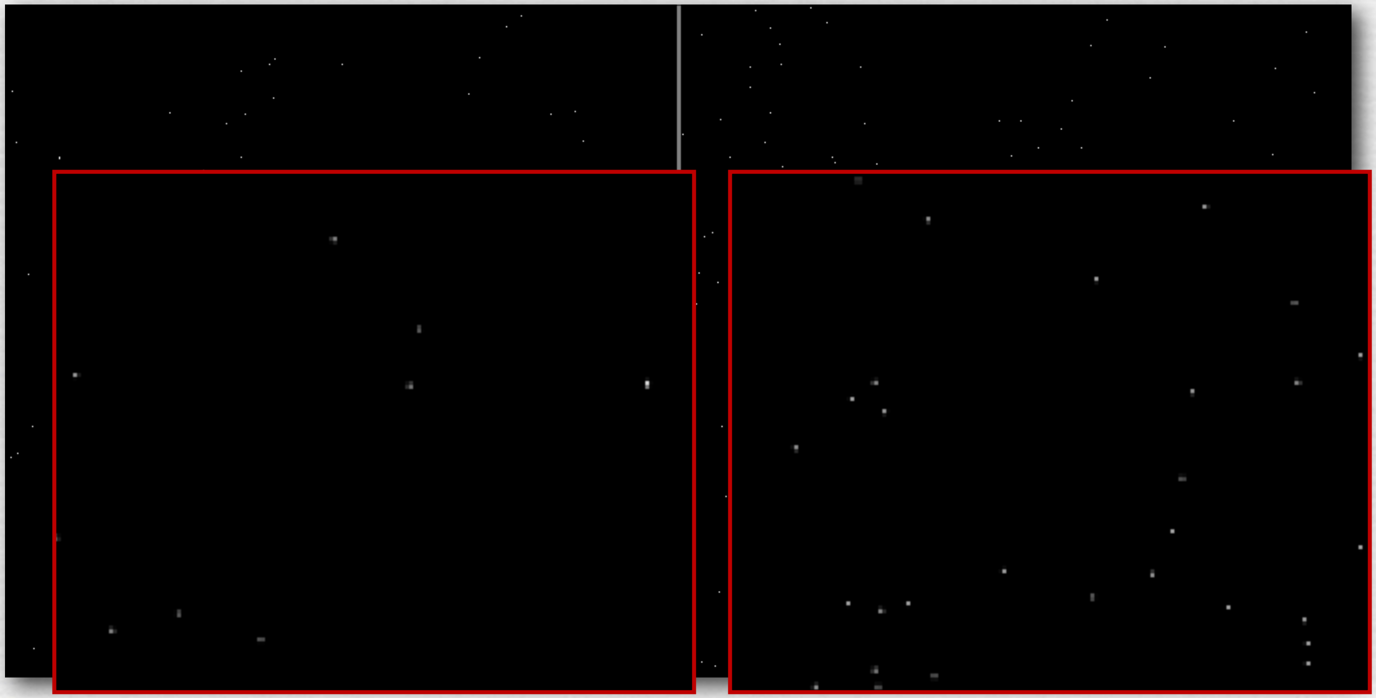

Canon 20D, 15 sec and 30 sec exposures

Fixing dark current and hot pixels

- example
  - Aptina MT9P031 (in Nokia N95 cell phone)
  - full well capacity = ~8500 electrons
  - dark current = 25 electrons/pix/sec at 55°C

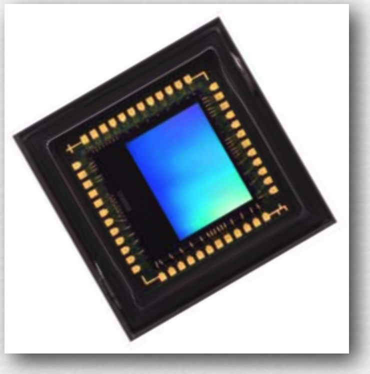

- solution #1: chill the sensor
  - Retiga 4000R bioimaging camera
  - Peltier cooled 25°C below ambient
  - full well capacity = 40,000 electrons
  - dark current = 1.64 electrons/pix/sec

- solution #2: dark frame subtraction
  - available on high-end SLRs

Fixed pattern noise (FPN)

- manufacturing variations across pixels, columns, blocks
- mainly in CMOS sensors
- doesn't change over time, so read once and subtract

Canon 20D, ISO 800, cropped

Read noise

- ✦ thermal noise in readout circuitry
- ✦ again, mainly in CMOS sensors
- ✦ not fixed pattern, so only solution is cooling

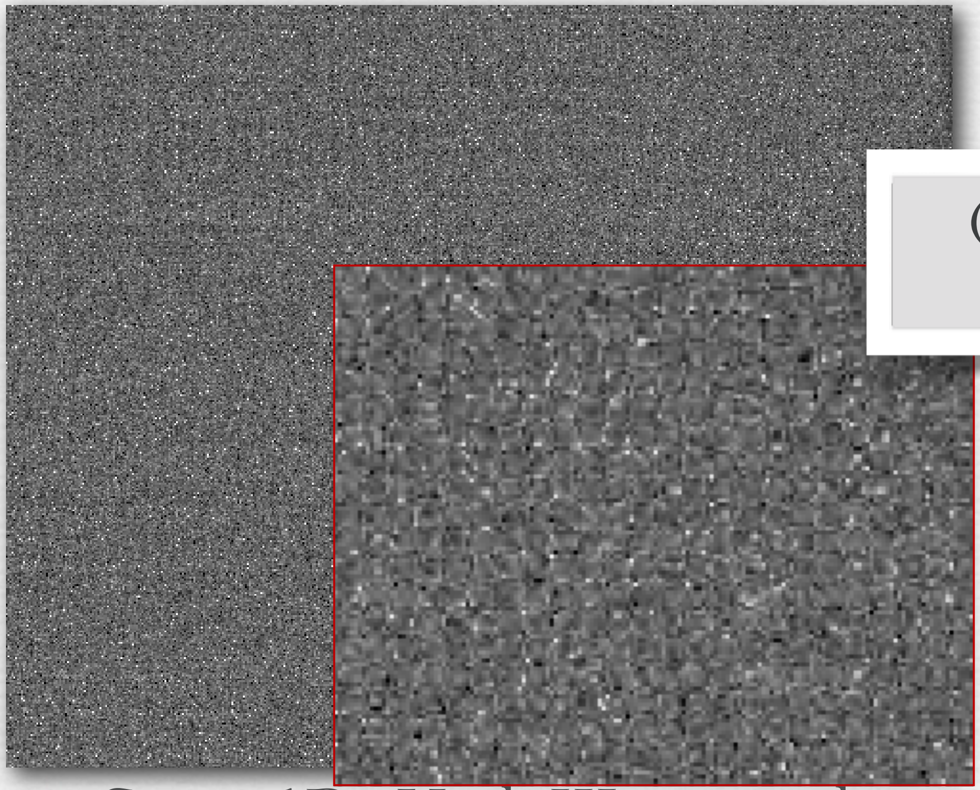

(image tainted by JPEG artifacts?)

Canon 1Ds Mark III, cropped

Signal-to-noise ratio (again)

$$SNR = \frac{\text{mean pixel value}}{\text{standard deviation of pixel value}} = \frac{\mu}{\sigma}$$

$$= \frac{P\,Q_e\,t}{\sqrt{P\,Q_e\,t + D\,t + N_r^2}}$$

- where

P = incident photon flux (photons/pixel/sec)

Q_e = quantum efficiency

t = exposure time (sec)

D = dark current (electrons/pixel/sec), including hot pixels

N_r = read noise (rms electrons/pixel), including fixed pattern noise

(formula from http://learn.hamamatsu.com/articles/ccdsnr.html)

Signal-to-noise ratio (again)

$$SNR = \frac{\text{mean pixel value}}{\text{standard deviation of pixel value}} = \frac{\mu}{\sigma}$$

$$= \frac{P\,Q_e\,t}{\sqrt{P\,Q_e\,t + D\,t + N_r^2}}$$

✦ examples

- Retiga 4000R = (1000 × 55%) / √(1000 × 55% + 1.64 + 12^2)
  = 20.8:1 assuming 1000 photons/pixel/sec for 1 second
- Aptina MT9P031 = (1000÷11×69%) / √(1000÷11×69% + 25 + 2.6^2)
  = 6.5:1 assuming pixels are 1/11 as large as Retiga's

Dynamic range

$$DR = \frac{\text{max output swing}}{\text{noise in the dark}} = \frac{\text{saturation level} - D\,t}{\sqrt{D\,t + N_r^2}}$$

- examples
    - Retiga 4000R = $(40{,}000 - 1.64) / \sqrt{(1.64 + 12^2)}$ electrons
      = 3,313:1 (11.7 bits) for a 1 second exposure, and
      = 3,333:1 (11.7 bits) for a 1/60 second exposure

    - Aptina MT9P031 = $(8500 - 25) / \sqrt{(25 + 2.6^2)}$
      = 1500:1 (10.5 bits) for a 1 second exposure, but
      = 3200:1 (11.6 bits) for a 1/60 second exposure
- determines useful ADC precision
- after gamma correction (for JPEG), you only see ~8 bits

ISO

- amplifies signal before analog-to-digital conversion
  - avoids losing low signal due to quantization and any noise introduced after quantization (yes, there is some)
  - doubling ISO doubles the signal, which is linear with light, so equivalent to doubling exposure time, or minus 1 f/stop
- maximum ISO on Canon 5D II is 6400
  - higher ISOs implemented using multiplication after ADC?
- raising ISO improves SNR relative to multiplication after ADC, or equivalently, brightening in Photoshop
- but raising exposure time improves SNR faster, so
- maximize exposure time to the limits imposed by object motion, camera shake, or sensor saturation, then maximize ISO to the limit imposed by ADC saturation

Averaging several short-exposure, high-ISO shots to avoid camera shake & reduce noise

Aligning a burst of short-exposure, high-ISO shots using the Casio EX-F1

1/3
sec

burst
at 60fps

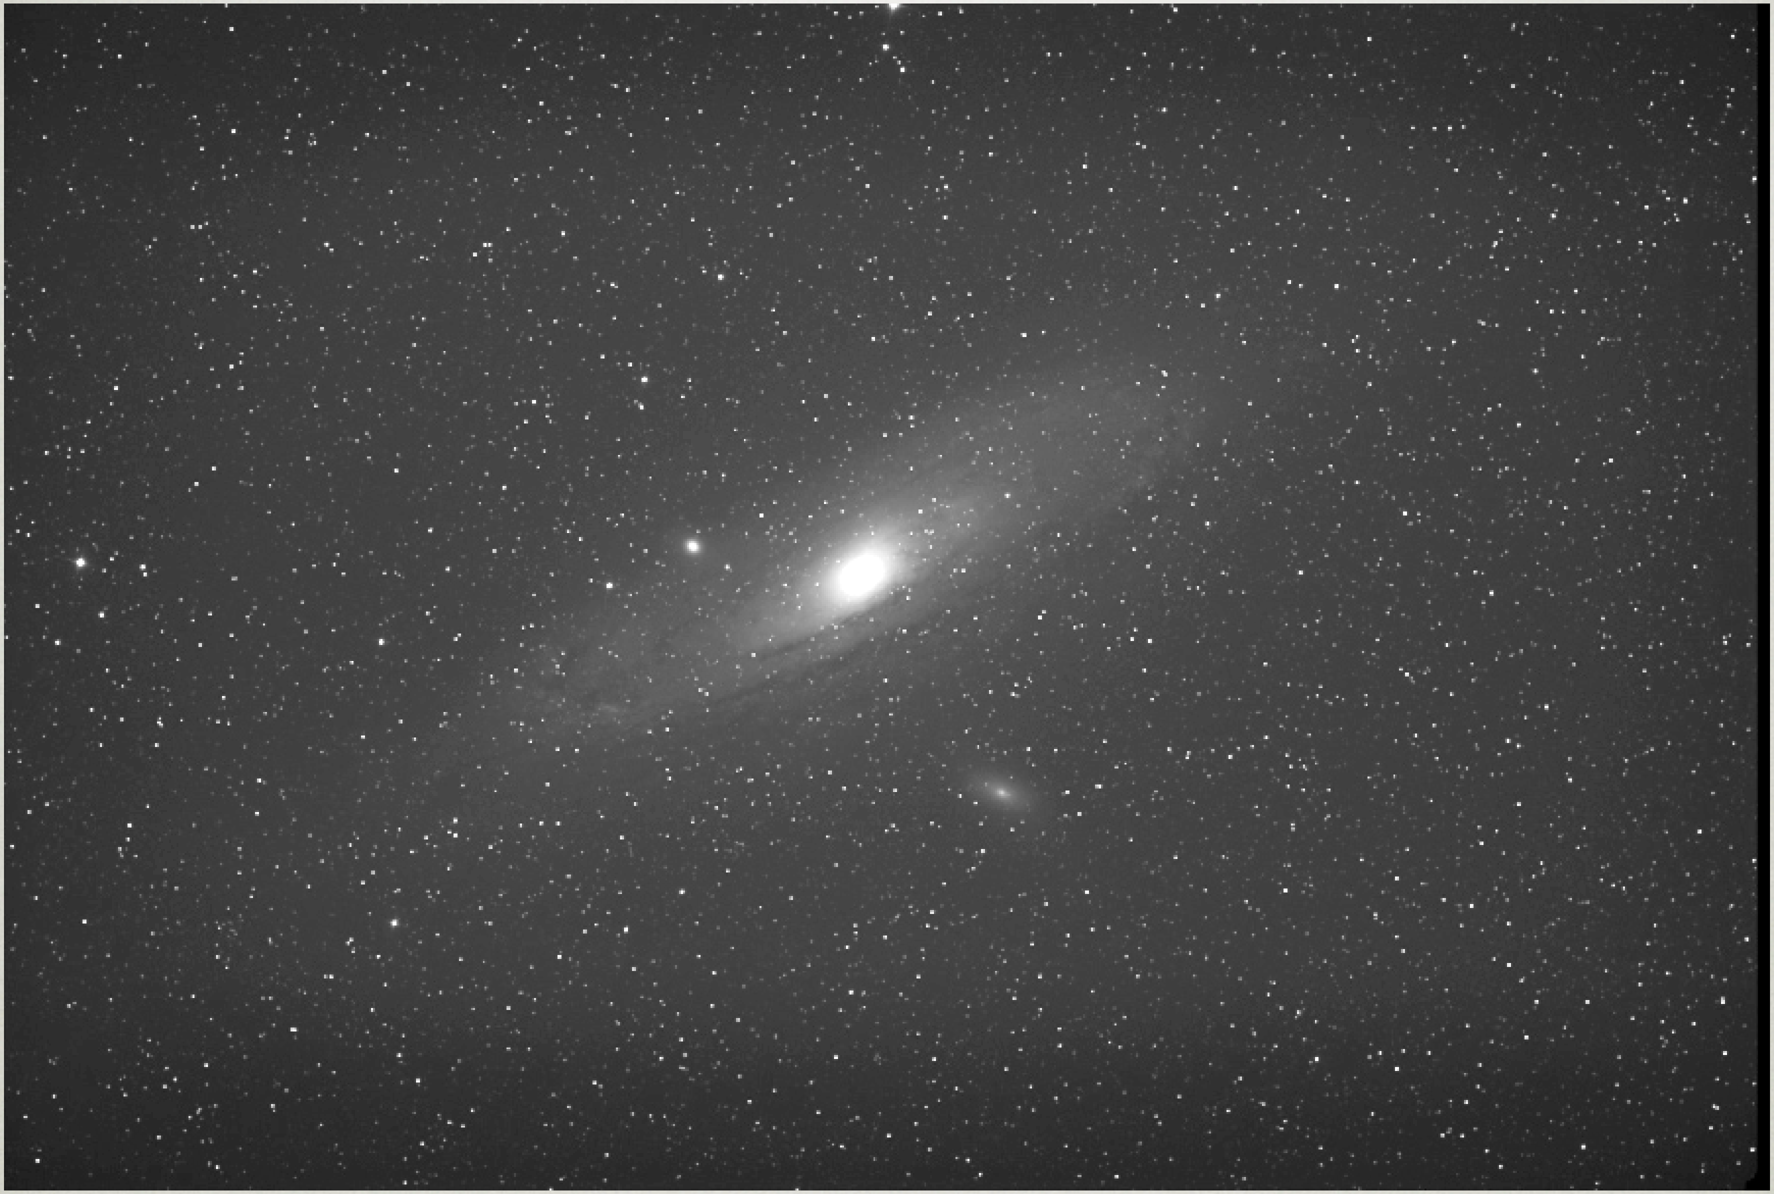

Jesse Levinson, Andromeda
(single exposure, 3 minutes)

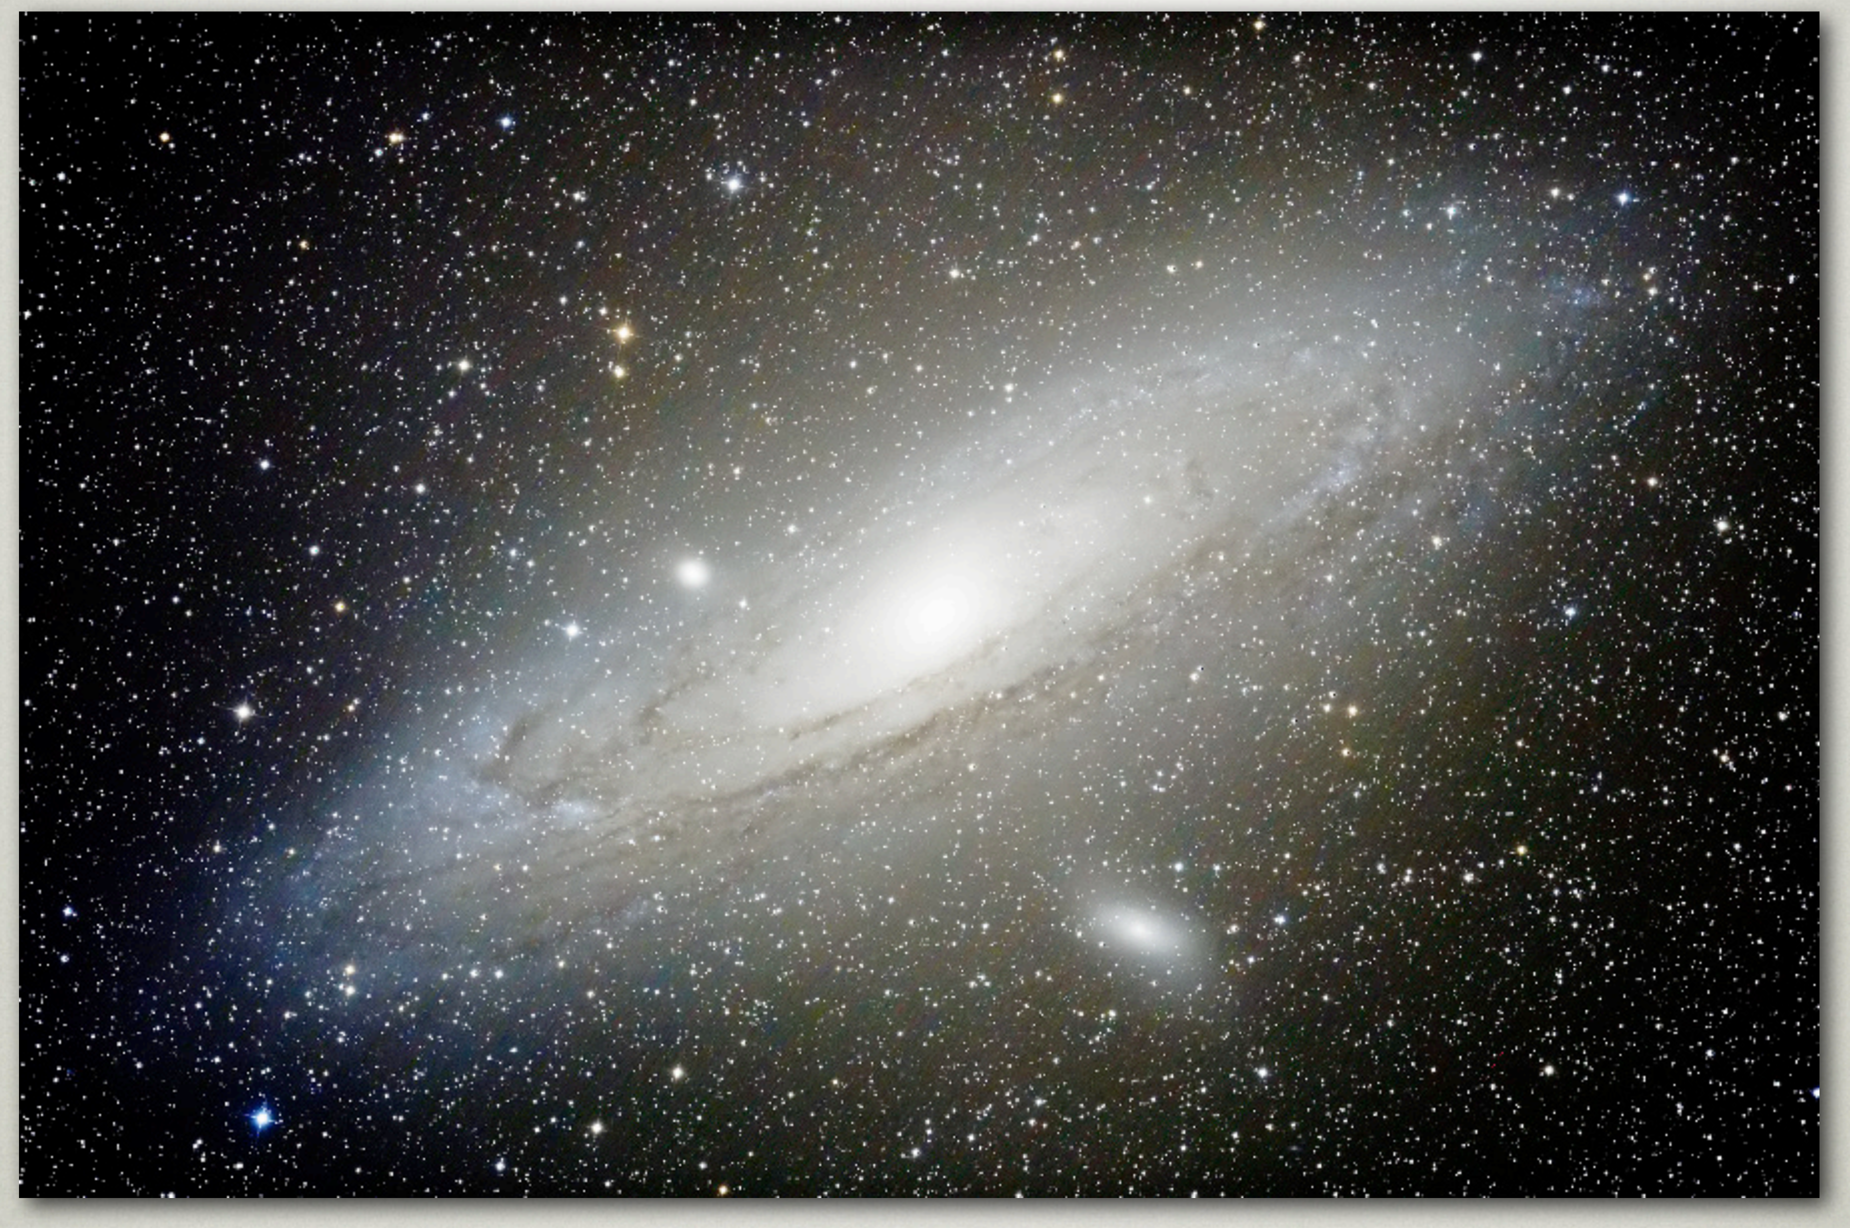

Jesse Levinson, Andromeda
(50 exposures of 3-minutes each)

Slide credits

- Brian Curless
- Eddy Talvala
- Abbas El Gamal

- Theuwissen A., *Solid-State Imaging with Charge-Coupled Devices,* Kluwer Academic Publishers, 1995.
- Filippov, A., *How many bits are really needed in the image pixels?* (sic), http://www.linuxdevices.com/articles/AT9913651997.html