

Art Competition #2 Results

Art Competition #2 – 3rd Place Winner




Shannon Shih

What's up, homie?

Made completely from scratch using Sculptris. Rigged in Blender with armature shamelessly stolen from Corrin's dragon form in Fire Emblem Fates.

Art Competition #2 – 2nd Place Winner



Cheng Cao

Extruding, subdivision, splitting, flipping, translations, simple operations! Recreation of the post-apocalyptic Las Vegas from Blade Runner 2049, using all these simple techniques with blender, starting off from a box and a circle, with a GLSL shader that adds atmospheric scattering simulation, phong shading with "Sun" type light source that creates the moody image of futuristic but totally dead Las Vegas.

Art Competition #1 – 1st Place Winner



Rachel Thomas

10,731,969 polygons
and many segfaults later
here we are, a copper dragon at last

Lecture 29:

Image Sensors

Computer Graphics and Imaging
UC Berkeley CS184/284A


Why Study Image Sensors?
Only A Sparse Random Sampling...

Imaging for Robotics




Google's "Arm Farm"

Imaging for Computer Vision



ImageNet: 15M images, 22K categories
<http://image-net.org>

Imaging in Mapping



Maps, satellite imagery, street-level imaging,...

Imaging in Mapping



Maps, satellite imagery, street-level imaging,...



Ubiquitous Imaging



Cameras everywhere

Photon Capture

The Photoelectric Effect



Albert Einstein



Einstein's Nobel Prize in 1921 "for his services to Theoretical Physics, and especially for his discovery of the law of the photoelectric effect"

Charge Coupled Devices (CCD)




Developed by Wilford Boyle (L) and George Smith (R) at Bells Labs in 1969
Nobel Prize 2009 - "for the invention of an imaging semiconductor circuit
– the CCD sensor"

Charge Coupled Devices (CCD)





CCD


Interline CCD



CMOS APS (Active Pixel) Sensor



Anatomy of the Active Pixel Sensor Photodiode



<http://www.olympusmicro.com/primer/digitalimaging/cmosimagesensors.html>

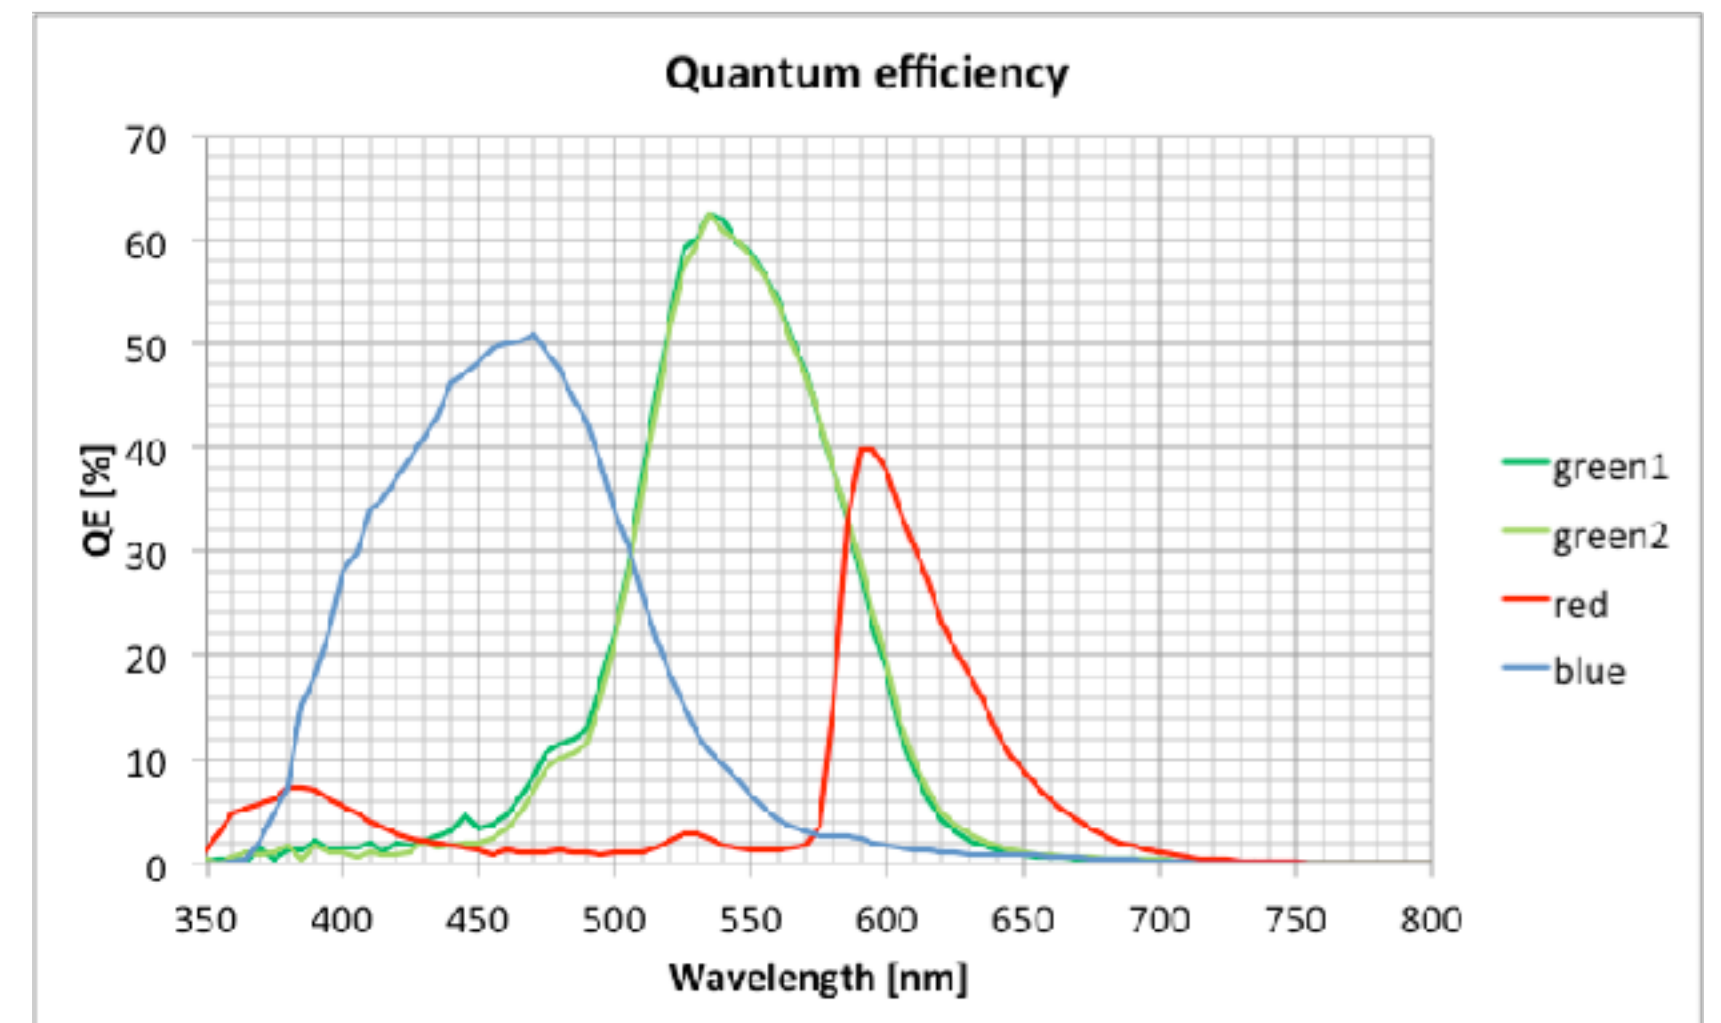
Quantum Efficiency

Not all photons will produce an electron

- Depends on quantum efficiency of the device

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

- Human vision: ~15%
- Typical digital camera: < 50%
- Mobile camera: 60%
- Best back-thinned CCD: > 90%
- Scientific CMOS (sCMOS) 95%




Meynants et al. IISW 2013


QE of a 24MP CMOS full-frame sensor

Color Architectures


Color Filter Arrays (Mosaics)



Bayer pattern
(most common)




Sony RGB+E
wider color gamut





Kodak RGB+W
higher dynamic range

Why more green pixels than red or blue?




- Because humans are most sensitive in the green portion of the visible spectrum
- Sensitivity given by the human luminous efficiency curve



Demosaicking Algorithms



Demosaicking Algorithms



Demosaicking Algorithms

Interpolate sparse color samples into RGB at every output image pixel


Simple algorithm: bilinear interpolation

- Average 4 nearest neighbors of the same color

Consumer cameras use more sophisticated techniques


- Try to avoid interpolating across edges

Due to demosaicking, 2/3 of image data is “made up”!




3-Sensor Color Architecture

- Prismatic optics
- No demosaicking
- Three (smaller) sensors and optical alignment



Philips Total Internal Reflection Dichroic Prism





Wavelengths Penetrate to Different Depths

Long-wavelength photons penetrate deeper than short in silicon

The spectral response of electrons at the surface differs from electrons deeper in the material



Foveon X3[®] Capture




Dynamic Range

CCD & CMOS Response Functions Are Linear

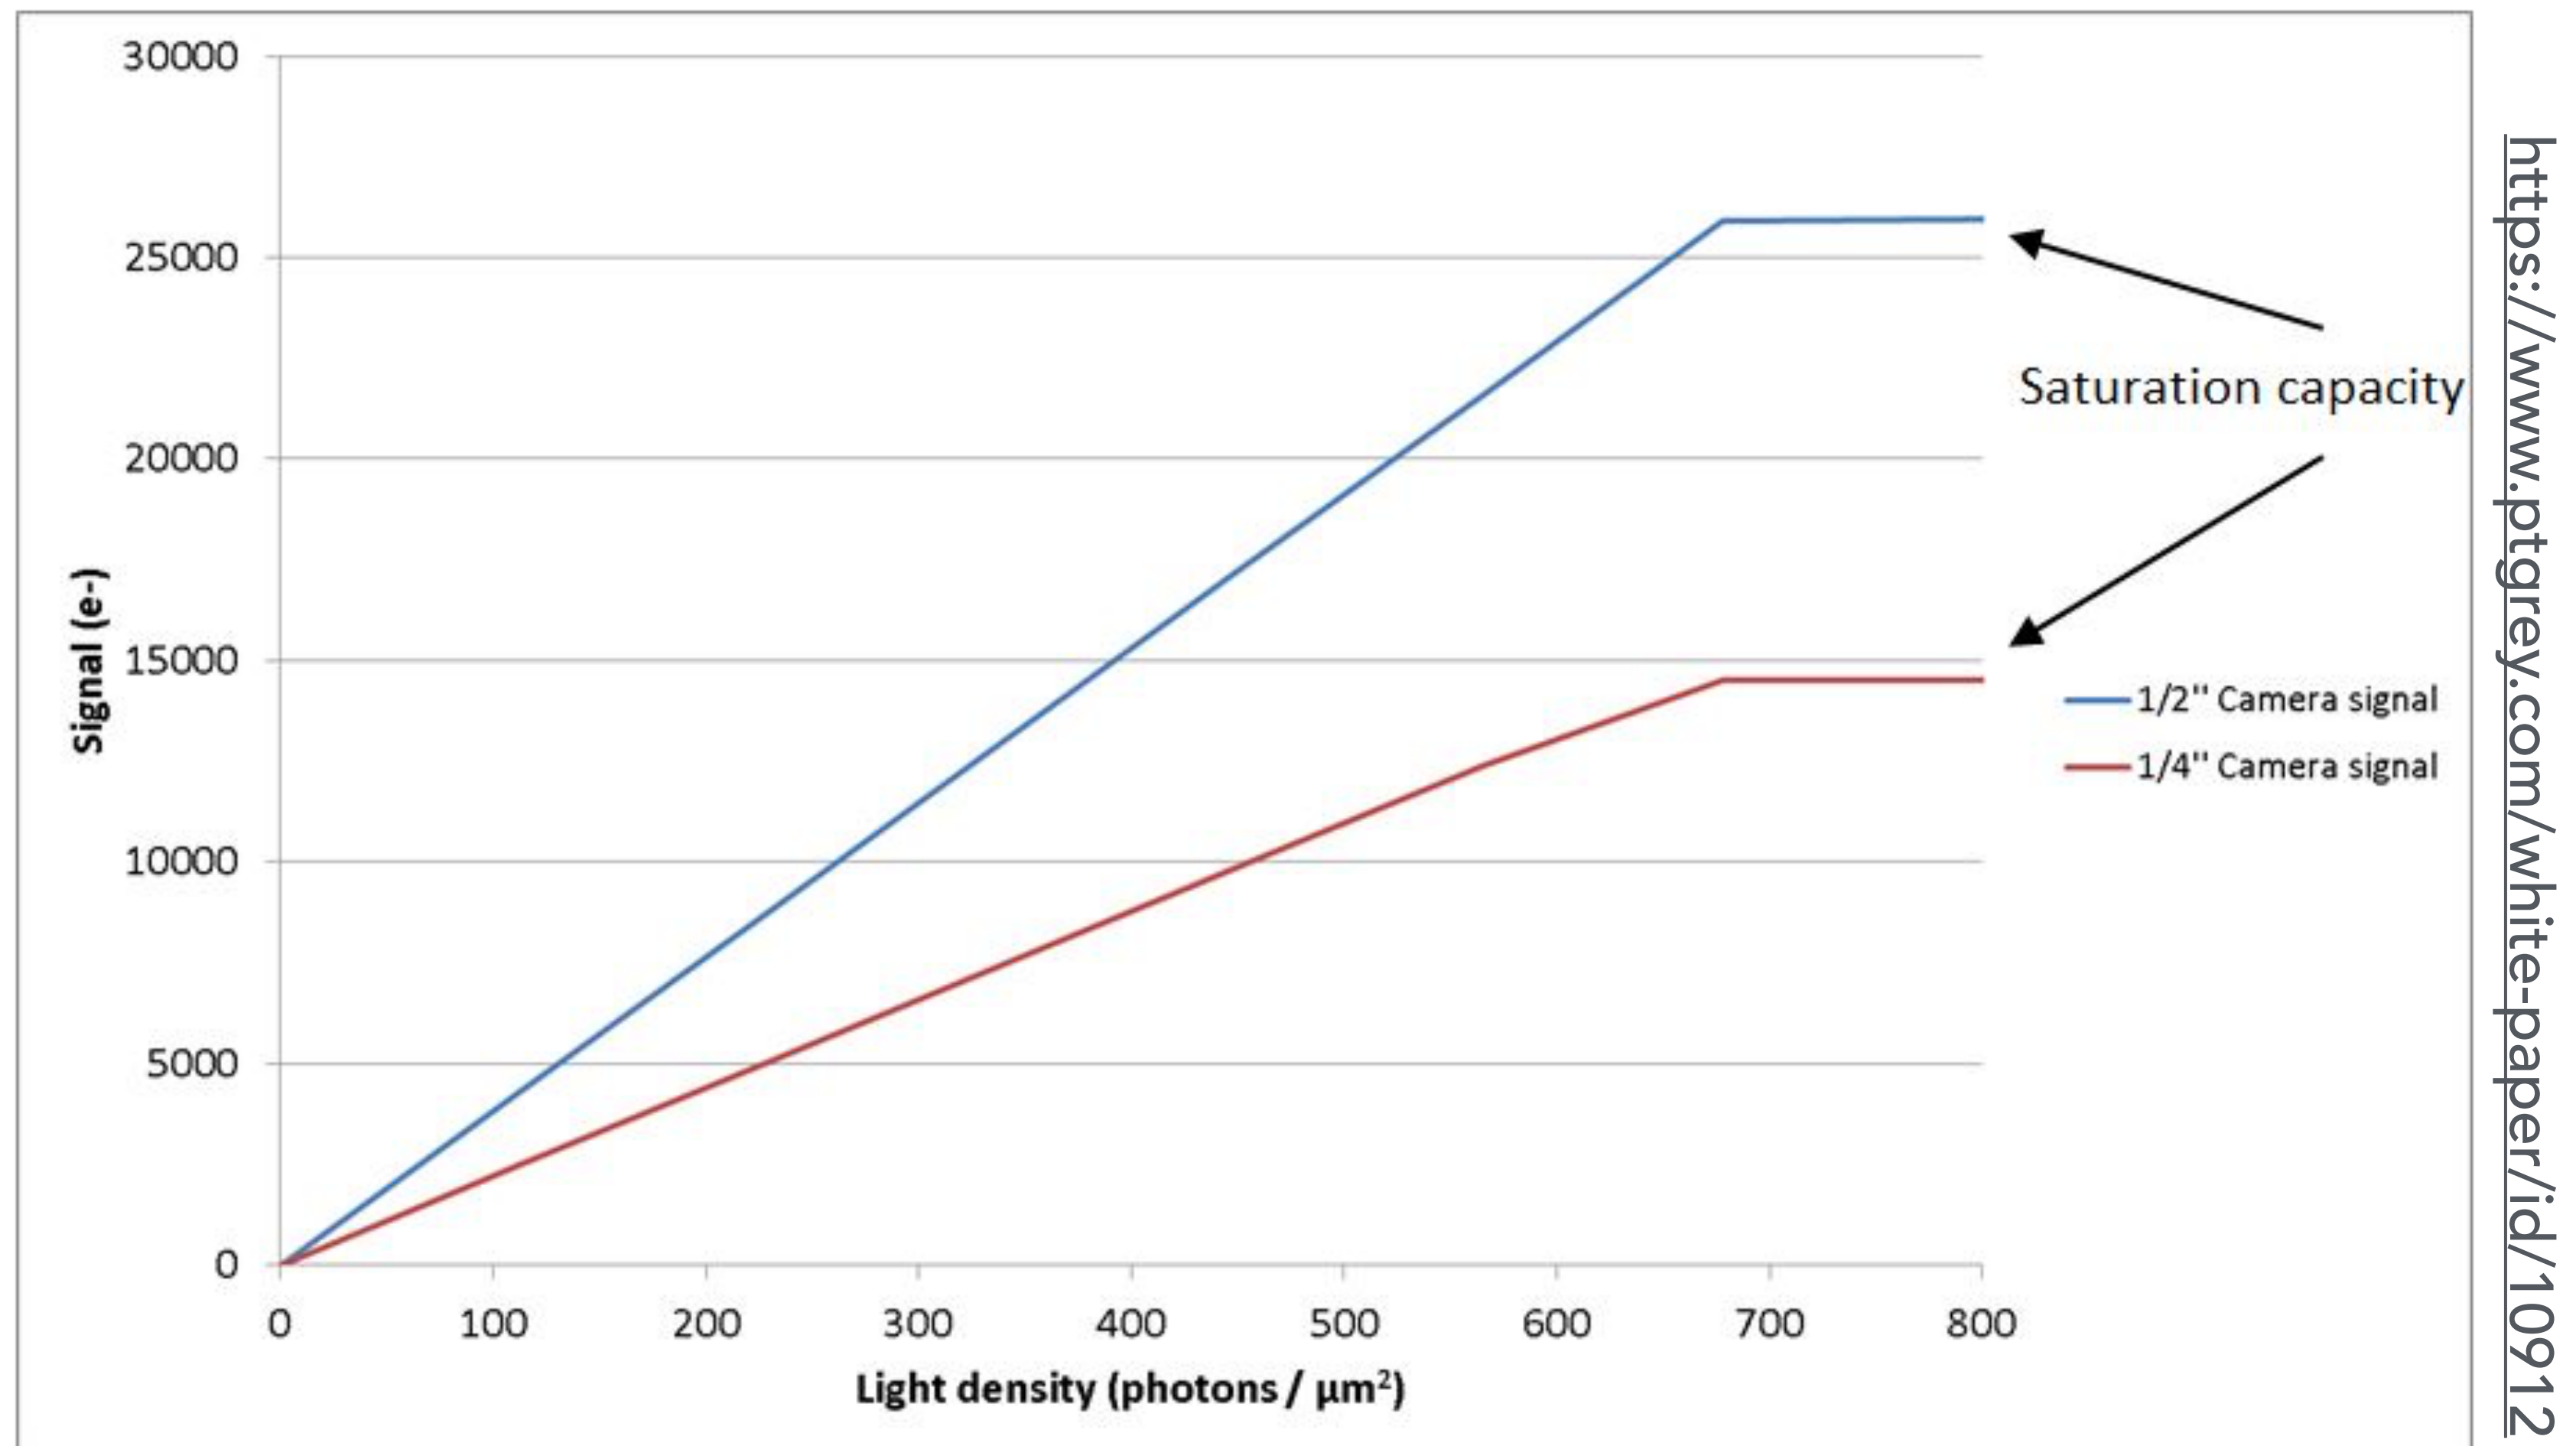
Photoelectric effect in silicon:

- Response function from photons to electrons is linear
- May have some nonlinearity close to 0 due to noise, and near pixel saturation



(Epperson, P.M. et al. Electro-optical characterization of the Tektronix TK5 ..., Opt Eng., 25, 1987)

Finite Dynamic Range: Real Sensor Pixels "Saturate"



Saturated Pixels: "Blown Out" Parts of Photo



ROMANAS NARYŠKIN, <https://photographylife.com/understanding-histograms-in-photography>

CS184/284A

Normal exposure

Ren Ng

Saturated Pixels: "Blown Out" Parts of Photo



ROMANAS NARYŠKIN, <https://photographylife.com/understanding-histograms-in-photography>

CS184/284A

Overexposure ("blown-out" or "clipped" pixels)

Ren Ng

exposure:
+0 stops



image: Paul Debevec

Slide courtesy Steve Marschner

exposure:
-8 stops



image: Paul Debevec

Slide courtesy Steve Marschner

exposure:
+6 stops



image: Paul Debevec

Slide courtesy Steve Marschner

High Dynamic Range (HDR) Through Multiple Exposures




Figure 6: Sixteen photographs of a church taken at 1-stop increments from 30 sec to $\frac{1}{1000}$ sec. The sun is directly behind the rightmost stained glass window, making it especially bright. The blue borders seen in some of the image margins are induced by the image registration process.

HDR Through Multiple Exposures




(a)



(b)



(c)



(d)



(e)




(f)


HDR Through Multiple Exposures



**Synthetic Motion Blur
Normal 8-bit image**



**Synthetic Motion Blur
HDR image**



Real Photo

DIY HDR



Slide courtesy Marc Levoy

- 2 shots
- Photoshop

DIY HDR



Slide courtesy Marc Levoy

- 2 shots
- Photoshop


DIY HDR





Slide courtesy Marc Levoy

- 2 shots
- Photoshop



HDR "Mode" On Smartphones



HDR By Pixel Mosaicking




HDR By Pixel Mosaicking




HDR By Pixel Mosaicking



(a) Exposure: T




(b) Exposure: $4T$



(c) Exposure: $16T$



(d) Exposure: $64T$




HDR in Circuit Design of Image Sensor


- **Many approaches tried: well adjusting, multiple capture, time to saturation, logarithmic, local adaptation, ...**
- **Multiple capture approach requires multiple, high-speed, non-destructive reads of the pixel's value during exposure**
- **Not common, not used in cameras for consumer photography today**

Pixel Structure & Micro Optics


Front-Side-Illuminated (FSI) CMOS

Building up the CMOS
imager layers







Polysilicon
& Via 1




Metal 1




Metal 2




Metal 3



Metal 4



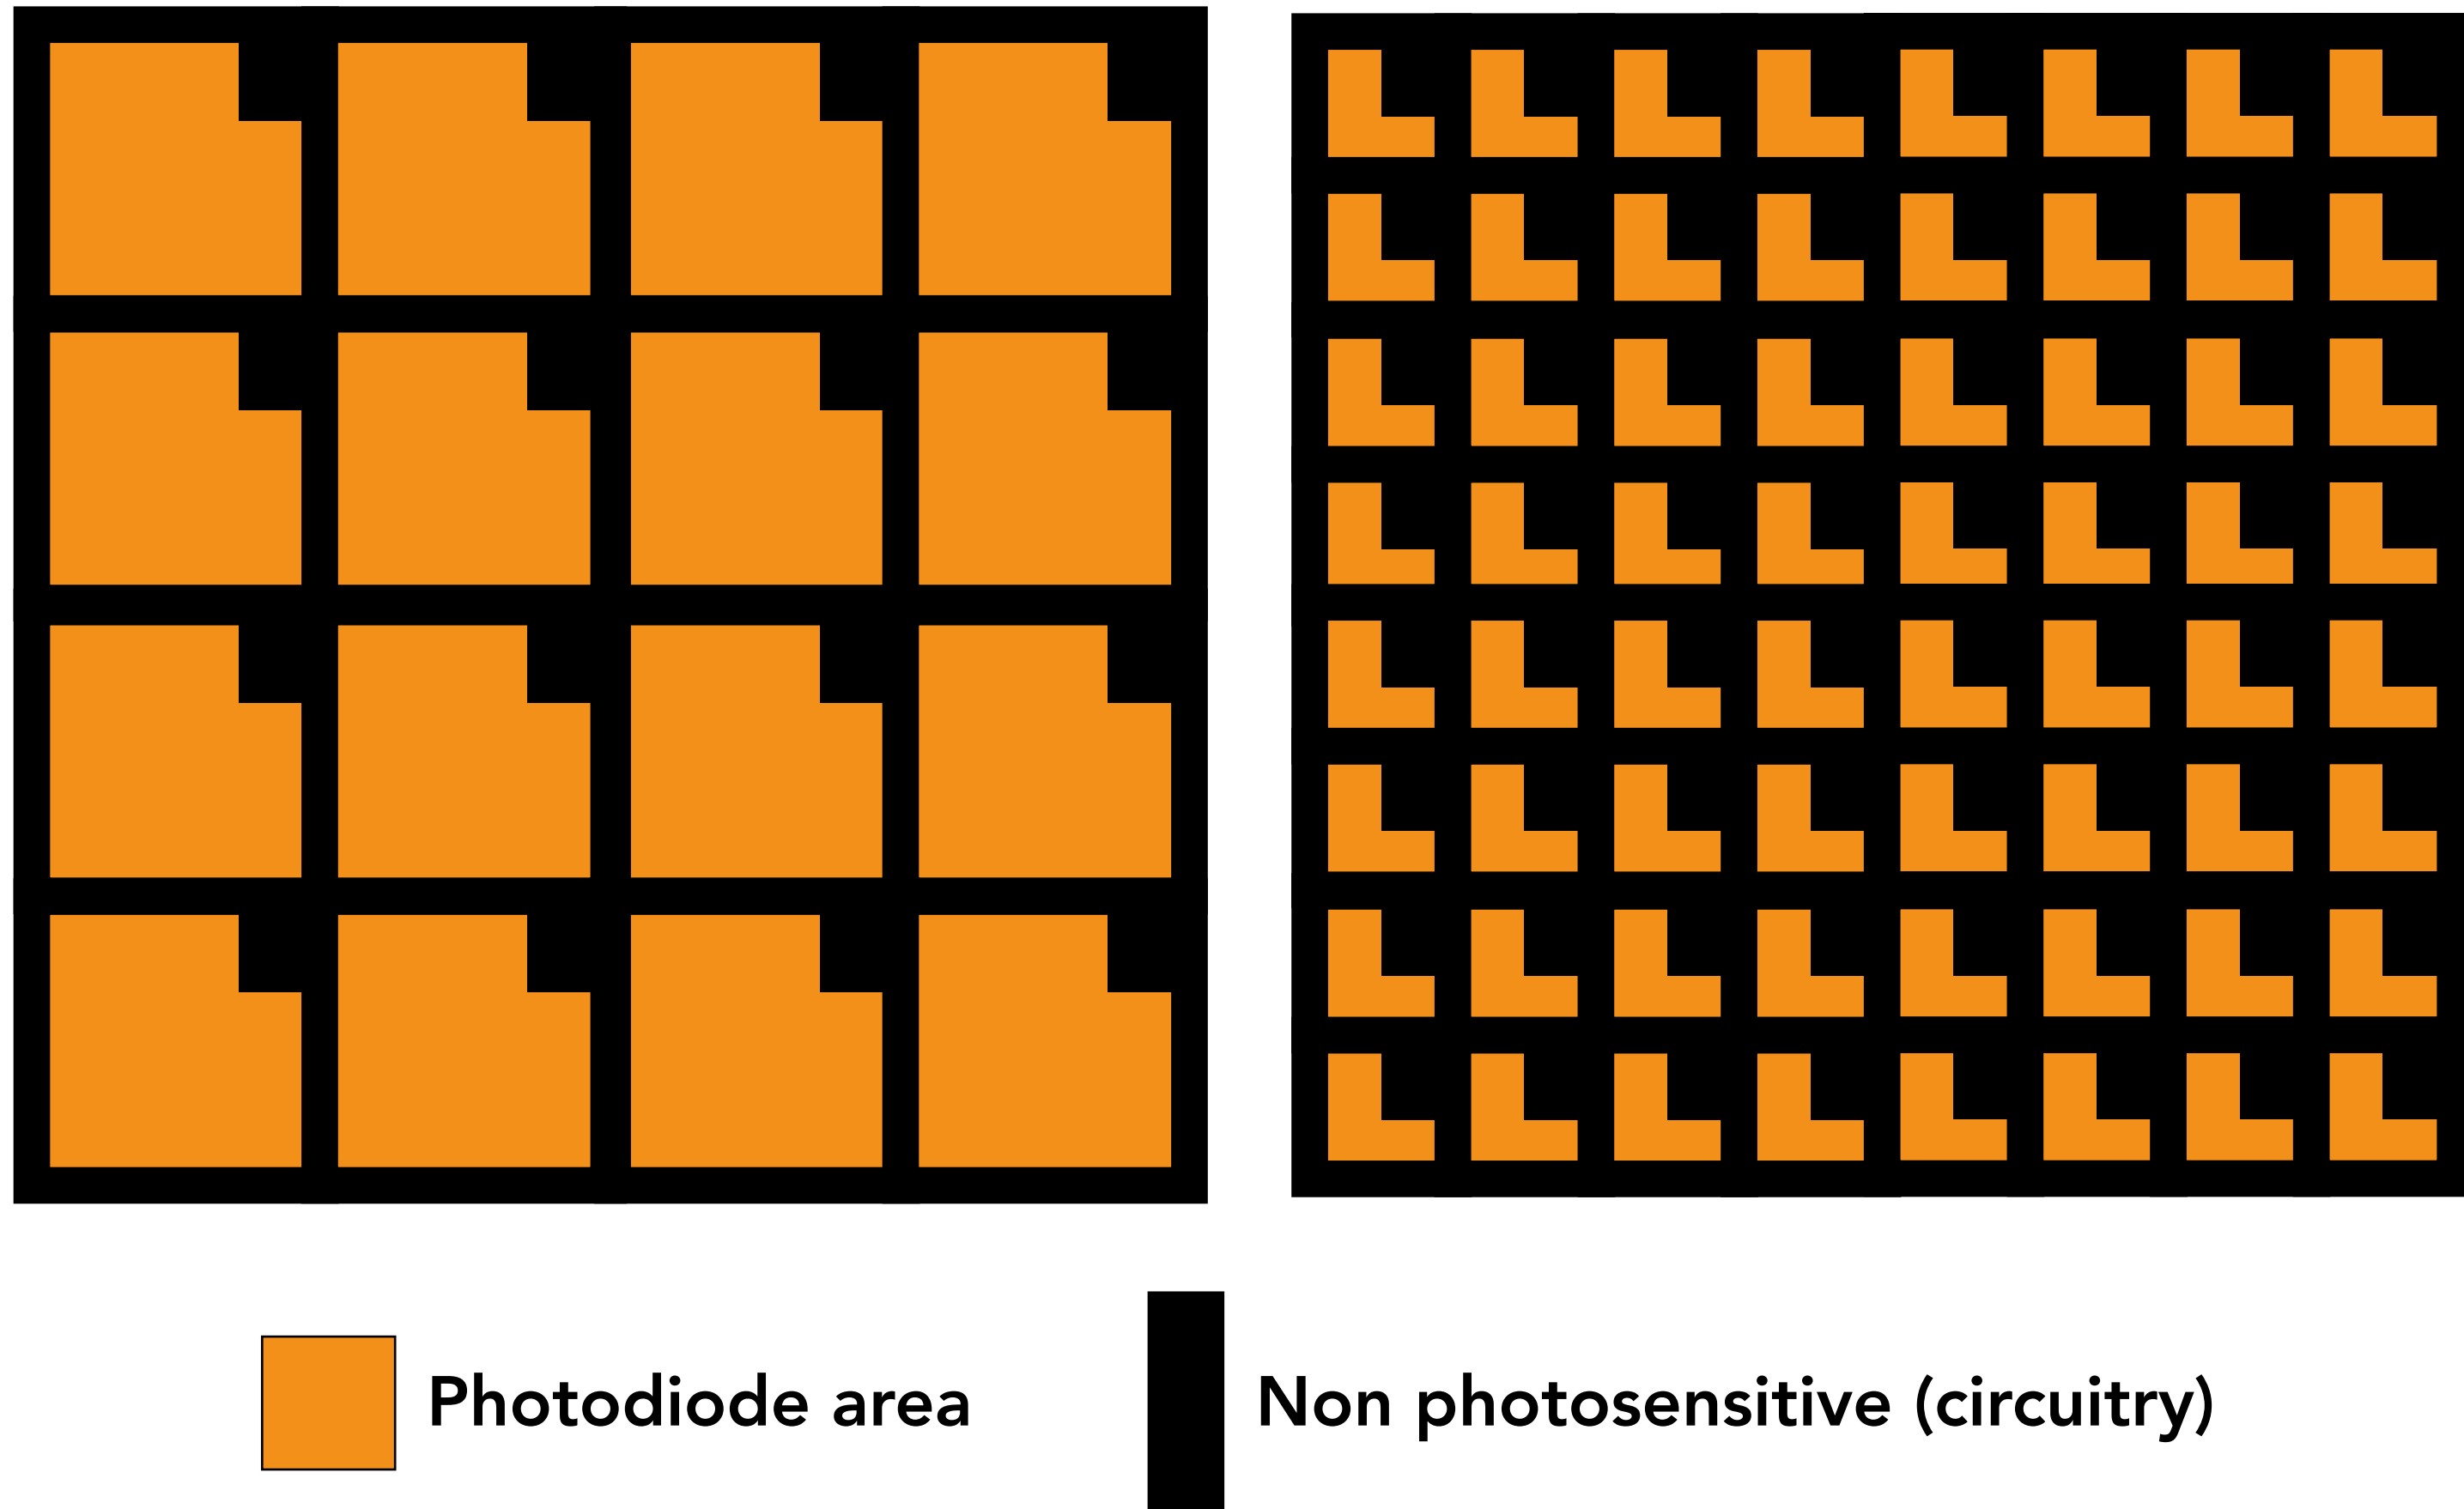
Color filter array



Courtesy R. Motta, Pixim

Pixel Fill Factor


Fraction of pixel area that integrates incoming light.




Pixel Fill Factor

Fraction of pixel area that integrates incoming light.

Optimize with per-pixel microlenses.




Microlenses on a CMOS sensor



Microlenses on CCD pixel

Pixel Fill Factor




Shifted microlenses on M9 sensor.




Leica M9

Optical Cross-Talk



<http://gmpphoto.blogspot.com/2012/09/the-new-leica-max-24mp-cmos-sensor.html>

Pixel Optics for Minimizing Cross-Talk



<http://gmpphoto.blogspot.com/2012/09/the-new-leica-max-24mp-cmos-sensor.html>


Image Example of Cross-Talk

Color desaturation due to pixel cross-talk



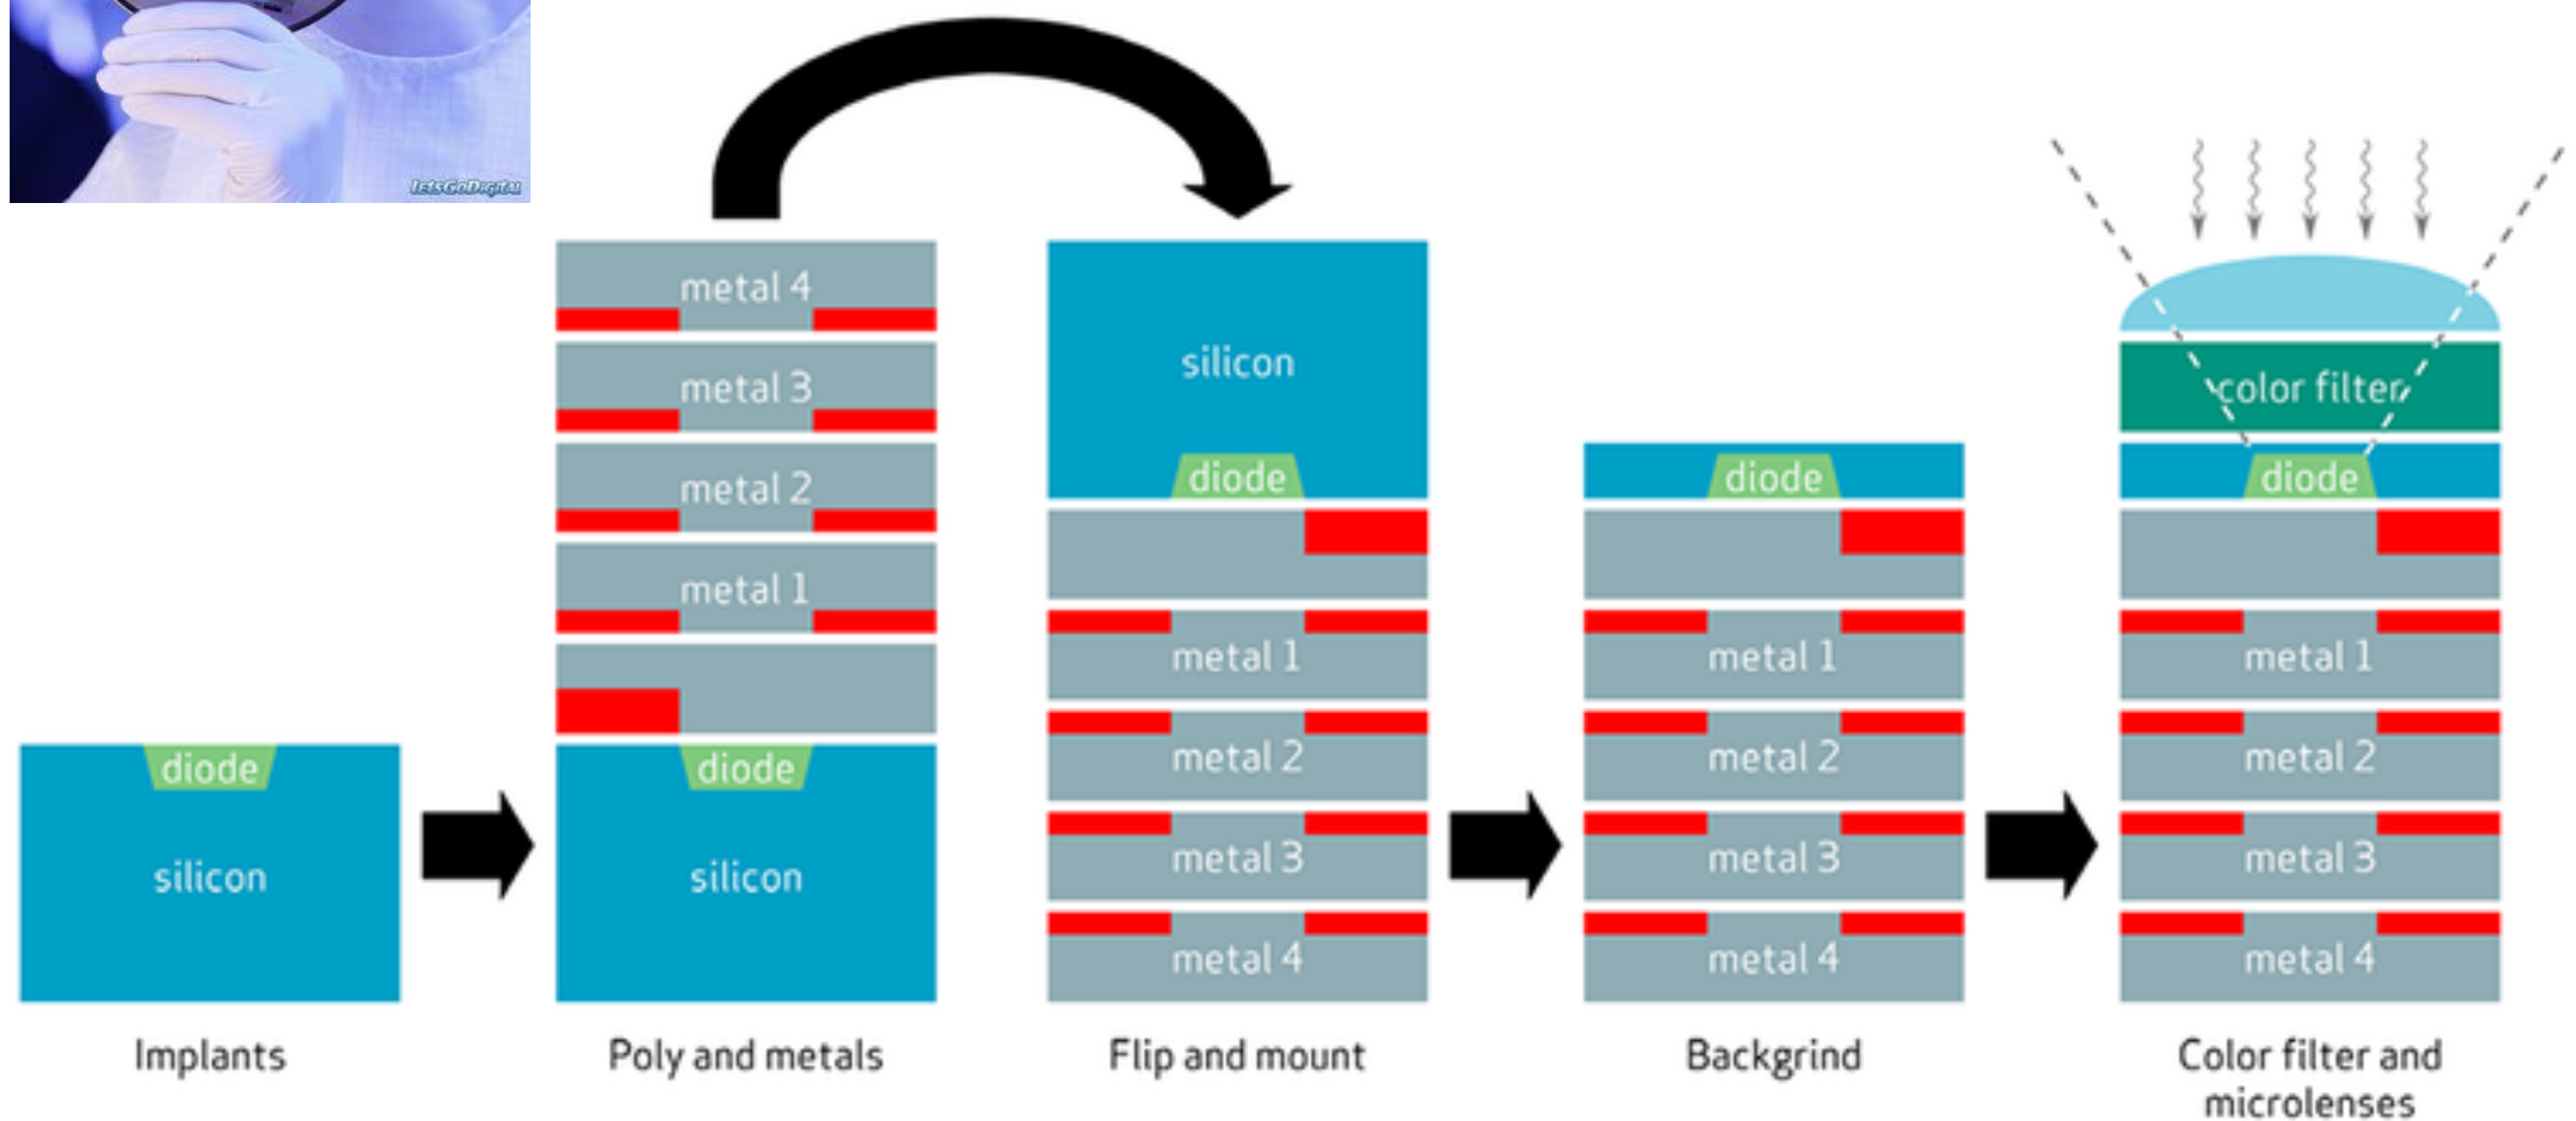
Kohyama et al. IISW 2009

Recall: FSI (Front-Side Illuminated) Pixel Structure




Humrick & Yankulin, tomshardware.com

BSI (Back-Side Illumination) Sensor Fabrication Process



Humrick & Yankulin, tomshardware.com

FSI vs BSI Pixel Structure



Humrick & Yankulin, tomshardware.com

Majority of CMOS Sensors are BSI Today



Smartphones




Some cameras

Good BSI sensors can provide higher QE and lower cross-talk.

Pixel Aliasing, Antialiasing

Pixel Fill Factor

Fraction of pixel area that integrates incoming light.





Pixel Sampling & Aliasing



What is going wrong in the image on the right?
Simulation of pixels with 25% fill factor

Pixel Sampling & Aliasing



Source of aliasing includes imperfect fill-factor, and color subsampling in color filter array.

Discussed techniques to improve fill-factor (e.g. microlenses)


Antialiasing Filter

Optical low-pass filter

- Use layer of birefringent material, splits each ray into two that overlaps each pixel
- Use two layers oriented at 90 degrees to split each ray over 2x2 pixels



Birefringence



Effect of one birefringent
OLPF layer (2D cross-section)

With and Without Antialiasing Filter @ 36 MP



D800E JPEG (default settings)



D800 JPEG (default settings)



D800E JPEG (default settings)




D800 JPEG (default settings)

With and Without Antialiasing Filter @ 36 MP



Without AA Filter (D800E)

With and Without Antialiasing Filter @ 36 MP



With AA Filter (D800)



Imaging Noise Fundamentals

(Most slides courtesy of Marc Levoy)

Image Noise



Image credit: [imaging-resources.com](http://www.imaging-resources.com)



<http://www.imaging-resource.com/PRODS/sony-a7s/YDSC00295.ARW.HTM>

Grain in image. Generally worse in low light, long exposures, shadows in images.

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, then it improves by $20 \log_{10}(200) - 20 \log_{10}(100) = +6 \text{ dB}$

Photon Shot Noise

The number of photons arriving during an exposure varies from exposure to exposure and from pixel to pixel, even if the scene is completely uniform

This number is governed by the Poisson distribution




Poisson Distribution

Probability that a certain number of random events will occur during an interval of time:

- Known mean rate
- Independent events

If on average λ events occur in an interval of time, the probability p that k events occur instead is

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



Poisson Distribution Mean and Variance


The mean and variance of the Poisson distribution are:

$$\begin{aligned}\mu &= \lambda \\ \sigma^2 &= \lambda\end{aligned}$$

The standard deviation is:

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

The error grows slower than the mean



Photon Shot Noise SNR

Photons arrive in a Poisson distribution

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

so

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

Shot noise scales as the square root of number of photons

Examples:



- A pixel that collects 10,000 photoelectrons vs. 1,000 has an SNR improvement of $\sqrt{10}$ or +10 dB
- Opening the aperture by 1 f/stop increases the number of photons by 2x, hence SNR by $\sqrt{2}$ or +3 dB

Sensor Noise Sources

(Most slides courtesy of Marc Levoy)

Pixel Noise: 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




don't confuse with
photon shot noise

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

Pixel Noise: 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

Pixel Noise: Fixing Dark Current & Hot Pixels

Example


- Aptina MT9P031 (in Nokia N95 cell phone)
- full well capacity = ~ 8500 electrons/pix
- 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/pix
- dark current = 1.64 electrons/pix/sec



Solution #2: dark frame subtraction

- available on high-end SLRs
- compensates for average dark current
- also compensates for hot pixels and FPN



Pixel Noise: 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

Pixel Noise: Read Noise

- Thermal noise in readout circuitry
- Again, mainly in CMOS sensors
- Not fixed pattern, so only solution is cooling



Canon 1Ds Mark III, cropped

Noise Recap

Photon shot noise

- Unavoidable randomness in number of photons arriving
- Grows as the square root of the number of photons, so brighter lighting and longer exposures will be less noisy

Dark current noise

- Grows with exposure time and sensor temperature
- Minimal for most exposure times used in photography
- Correct by subtraction, but only corrects for average dark current

Hot pixels, fixed pattern noise

- Caused by manufacturing defects, correct by subtraction

Read noise

- Electronic noise when reading pixels, unavoidable

Signal-to-Noise Ratio Revisited

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

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


**SNR changes with scene
brightness, aperture,
and exposure time**

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>)

Effect of Downsizing on Image Noise



averaged
down

point
sampled



Levoy

Noise Reduction by Image Averaging

Single frame
in dark room
using iPhone 4



Noise Reduction by Image Averaging

Average of
~30 frames
using SynthCam
app by Marc
Levoy

SNR increases as
 $\sqrt{\text{\# of frames}}$
(neglecting read
noise)



ISO - Signal Gain

Doubling ISO doubles the signal

- **Linear with light, so same as 2× exposure time, or brighten by one f-stop**
- **Implemented as analog or digital amplification**
 - **Analog before ADC on Canon 5D II up to ISO 6400; digital multiplication at higher ISOs?**

Ideal to amplify as early as possible during readout

- **If amplification occurs before read noise is added, and read-noise is independent of signal amplitude, then the amplified signal will have better SNR**



Nikon D3S, ISO 3200, photograph by Michael Kass



Nikon D3S, ISO 6400, photograph by Michael Kass



Nikon D3S, ISO 25,600, denoised in Lightroom 3, photograph by Fredo Durand
(Too dark to read menu)



Nikon D3S, ISO 25,600, denoised in Lightroom 3, photograph by Fredo Durand



RAW image from camera, before denoising in Lightroom



Tone mapped to show the scene as Fredo might have experienced it

Things to Remember

Photoelectric effect

Imager revolution: CCD and CMOS sensors

Sensor saturation

High dynamic range (HDR) imaging

Color architectures, Bayer filter array, demosaicking

Pixel stack, fill factor, microlenses

FSI vs BSI pixel designs

Pixel sampling, aliasing, optical low-pass filters

Noise: photon shot, pixel noise sources, SNR

Acknowledgments

Many thanks to Marc Levoy, Brian Wandell, and Pat Hanrahan, who created many of these slides.