

AY105 IR Detectors, May 2010

This week we will move into the COO detector lab under the guidance of Roger Smith and Gustavo Rahmer, who develop astronomical detector systems for various ground and space based applications.

The aim is to demonstrate infrared detector readout principals, behavioral problems and some calibration techniques, with particular emphasis on how IR detectors differ from CCDs.

Rather than work through a set of recipes to acquire the data, you will focus on the examination of data that has already been acquired for you. You will be using image display and analysis tools in the IRAF astronomical image analysis package. We will begin with an introduction by Roger to IR detectors and the test system used to acquire your data, and Gustavo and Roger will be on hand to answer questions and help you through the quite extensive activities planned....

TASKS TO BE PERFORMED

Read the introductory material (up to page 9 of this document) beforehand. We will review some of this at the start of class and take questions, but time will be limited.

Day 1:

The first set of analysis tasks will explore different readout schemes. You will see some of the signal offsets that are hidden by the signal processing but can affect dynamic range.

Astronomy requires measurement of very weak signals so we will focus on various noise sources which limit sensitivity:

- Fixed patterns (large but easily removed)
- Random noise originating in the sensor and electronics (“read noise”)
- Photon shot noise (fundamental property of the signal)

Day 2:

To make sense of the numbers obtained you will need to figure out how to convert from the numerical values (ADU) back to photons detected. To do this, we will use photon shot noise to infer the system gain and touch on some complications of these kinds of detectors such as InterPixel Crosstalk.

Finally, you will look at “dark current” which is the spurious time dependent signal generated by the detector, which has shot noise just like photogenerated current. This is important since it is the ultimate limit to detector sensitivity. You will see that it

depends on temperature, so we operate these sensors very cold, allowing us to detect fluxes as low as ~ 10 photons/pixel/hour.

THE TECHNOLOGY

Astronomical detectors must be sensitive and accurate.

To detect lower energy photons special materials with lower bandgap than silicon are required. This makes IR detectors harder to make than optical sensors based on silicon, such as CCDs. They are more costly to manufacture, more complex to operate and require more elaborate calibration to correct for non-ideal behavior.

Manufacture:

Maximum detectable wavelength is that where the photon energy ($h.c/\lambda$) equals the energy required to excite an electron from the valence to conduction band of the semiconductor, the “bandgap”. Beyond this “cutoff wavelength” the material becomes transparent.

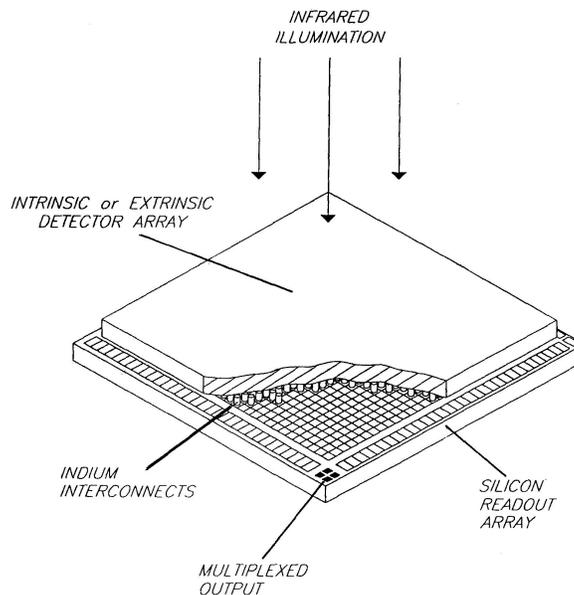


Figure 1: Infrared photodiode arrays are made from low bandgap semiconductors, such as HgCdTe, which are connected vertically by Indium columns to underlying electronics built on Silicon using conventional CMOS transistor fabrication technology. Courtesy I.McLean

In the 1980's some attempts were made to build CCDs using low bandgap (IR sensitive) material. However processing techniques are still not available, which allow fabrication of IR CCDs with adequate charge transfer efficiency. So manufacturers turned instead to using an array of photodiodes *fabricated in IR sensitive material*, connected to a readout circuit, implemented in Silicon, on an underlying layer.

The readout circuit (to be described below) requires just three CMOS transistors per pixel (and associated electrical traces), which occupy only a fraction of the area of a typical 18 μm wide pixel. This leaves space for the relatively large electrical contact pad that provides the interconnect path to the diode array, lying in the plane above it.

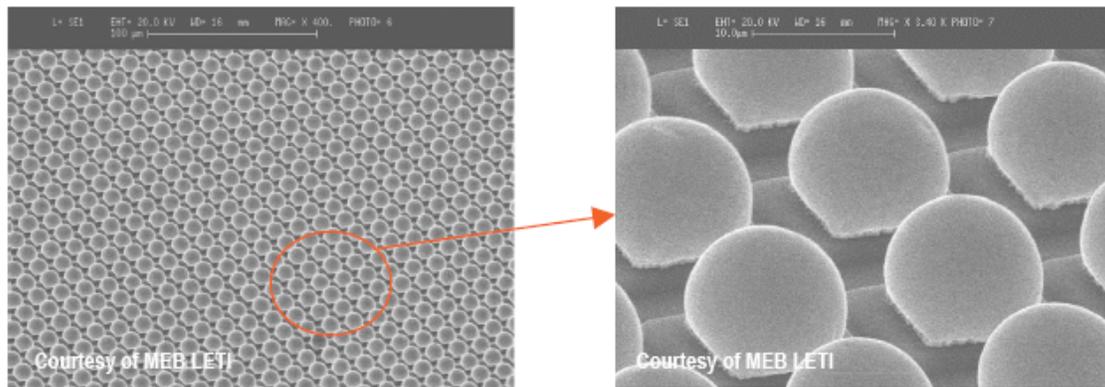


Figure 2: View of a field of indium micro-bumps with a pitch of 15 μm

Figure 2

The technology for making the vertical connection between the dissimilar materials in the light sensing and signal processing layers is the key to IR detector manufacture. The contact is constructed by depositing a thick layer of Indium on each pad, one per pixel, of the readout IC (through an etched photo-resistive mask). Matching “Indium bumps” are deposited on the underside of the photodiode array. The tops of the Indium bumps must be accurately coplanar and very clean so that when the bumps on the detector layer and the silicon layer are precisely aligned then squeezed together, a cold weld is formed making a permanent electrical and mechanical connection – one per pixel. Currently it is possible to connect 4 million pixels with only a few hundred failures. A low viscosity epoxy is then wicked into the <10 μm wide spaces between the Indium columns and the detector layer is then polished and etched until it is only ~10 μm thick.

This is a very complex and delicate process with yield problems at every step, so the top quality devices carry price tags in the \$250-500K range, making IR detectors five to ten times as expensive as CCDs.

Operation:

See Figure 3. Before the exposure, the reset switch is closed, so that the photodiodes are reverse biased by a hundreds of millivolts. The CMOS transistor, which buffers the diode voltage has essentially zero gate leakage at the low temperature required for optimal photodiode performance, so the change in voltage on the photodiode is dominated by the electron-hole pair generation by photons. During the exposure the diode junction acts as a capacitor, which stores this photogenerated charge. The electric field induced by the reverse bias has driven all mobile charges (carriers) from the P-N junction (depletion region), which then acts as an insulator (dielectric) between the capacitor “plates”, which are formed by the conductive outer regions of the diode where

mobile charges reside. The voltage across the diode capacitance is reduced by the accumulation of the photo-generated charge (and thermally generated “dark current”). Since the separation of the “capacitor plates” (width of the depletion region) is a function of accumulated charge, the voltage change is slightly non-linear, but we will ignore this for the moment.

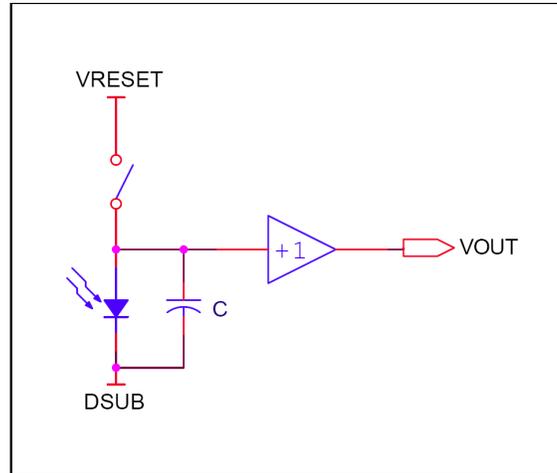


Figure 3: schematic representation of one pixel.

The pixel architecture described is able to detect very faint fluxes only because the leakage currents in the reverse biased photodiodes and the fully depleted channel of the reset MOSFET are incredibly low. This is achieved by cooling to $<150\text{K}$. All mobile (conduction band) charge has been driven from the MOSFET channel or diode junction by the application of suitable voltages, so that the only source of spurious mobile charge is from excitation of valence electrons by lattice vibrations. At low temperatures, the probability of thermal excitations exceeding the bandgap energy becomes very small. The dependence of dark current on bandgap means that the IR detector material has much larger dark current than the Silicon reset transistor.

This dark current sets the minimum detectable flux, which explains why these detectors are operated at such low temperatures. Impurities or lattice defects create localized intermediate energy levels producing higher dark current than would be expected for the bandgap, so much effort is invested in obtaining the best quality materials. With cooling, these reach a dark current floor at about 0.002 to $0.010\text{e}^-/\text{s}$ just 10^{-21}A .

To record an image, all we need to do is reverse bias (“reset”) the diode, then disconnect any current path (open the reset switch). We then measure the rate of voltage change induced by the light incident on the diode. The circuitry to sense the accumulated charge is all within the pixel footprint. Unlike a CCD the image is not shifted across the detector surface. Since the image is read *in situ*, no shutter is required.

The exposure time is the time between samples and not the time since reset. Signal arriving prior to the first sample is ignored (subtracted from the final sample), so the exposure duty cycle is not 100% but approaches it when the exposure times are long.

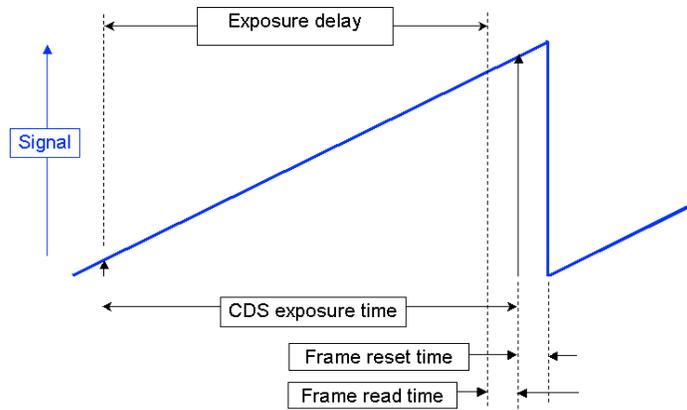


Figure 4: Sample timing for the last pixel in a CDS frame.

It is reasonable to ask why two samples are necessary. The most obvious effect of the first sample is to remove the DC offsets which are intrinsic to the readout circuit and which vary from pixel to pixel. One might consider calibrating these offsets less than once per exposure, however our inability to correct perfectly for changes in these offsets is one of the primary sources of error. There are uncertainties in the charge left on the diode after reset, both due to thermal noise in the reset switch resistance and the voltage to which it is connected. The buffer transistor offset voltage varies rapidly with temperature changes, and is sensitive to supply and detector substrate bias voltage changes. Even if external noise sources could be reduced, the initial sample would still be required to subtract offset drift due to low frequency noise (drift) in the buffer transistors.

When the diode voltage is sampled only once at the start, and again at the end of the exposure, we call this “Correlated Double Sampling” (CDS). This is similar to the CDS performed when reading a CCD, except that the signal samples span the exposure time rather than a pixel time. Clearly the 10^5 to 10^8 greater time between samples makes infrared detectors very much more susceptible to sources of zero-point drift than CCDs.

MOSFETs generate a drain-source current in response to a voltage on the gate. The gate is insulated from the FET channel so leakage at cryogenic temperatures is negligible and independent of drain-source current. Thus the accumulated charge is measured *non-destructively*. We will take advantage of this later to allow a more precise estimate of the incoming flux in the presence of electronic noise (referred to as “read noise”).

Thus far we have considered how charge is integrated on the sensing node, and how the exposure time is defined, but we have not considered how one accesses the outputs of millions of pixels packed into a small area.

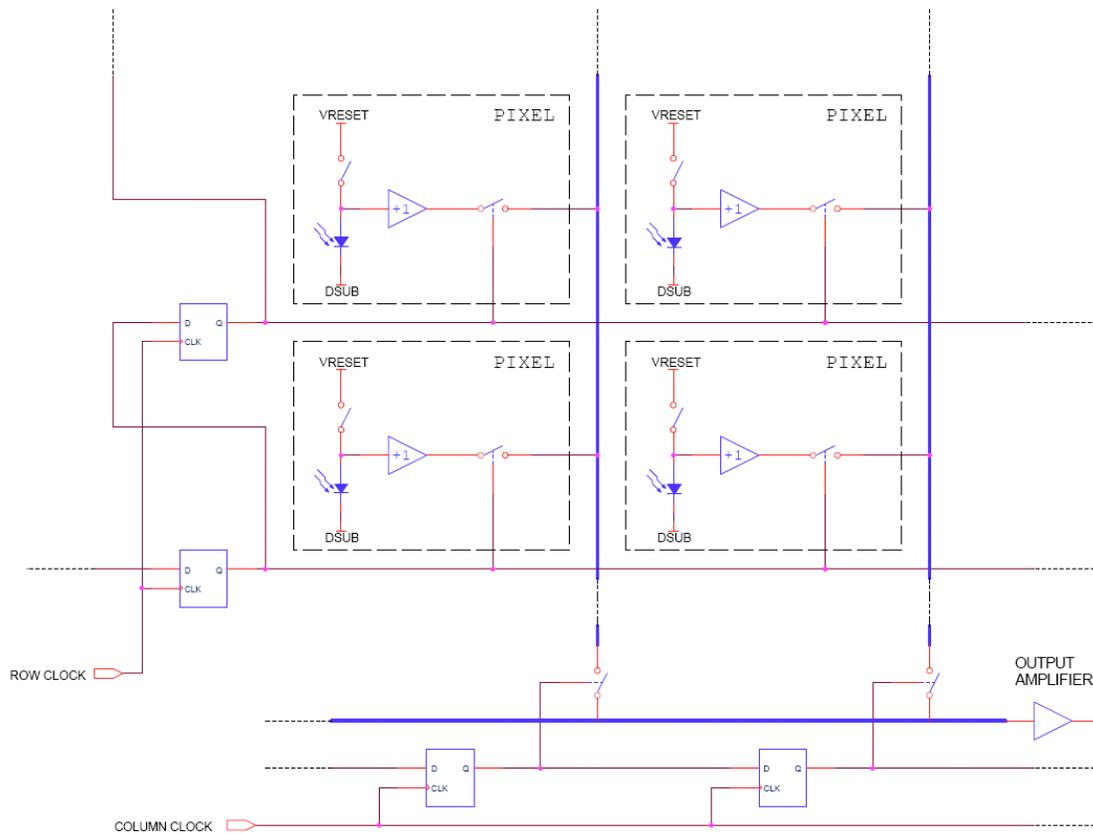


Figure 5: Schematic layout of a readout multiplexor for an infrared detector array.

The pixels are accessed via a 2D multiplexor. Each pixel output has a single MOSFET which is driven hard on or off, to act as a switch. Outputs of all switches in the same *column* are connected to a bus. The control lines for the pixel-select switches in each *row* are ganged together. Only one row-enable line is active at a time: every pixel in the selected row is connected to a different column bus. At the edge of the array, each column bus is connected via a switch to the output buffer. To raster through the pixels one enables a row then sequentially selects columns, then repeats for the next row. To operate a 2048x2048 array one does not need to generate 2048 row and column select signals. These are the outputs of serial registers running along two edges of the multiplexor. A single bit is loaded into the row/column-select shift register at frame/line-start then shifted by the row/column clock. In some devices the row clock is derived from the output of the column shift register so only a frame start pulse, one clock and a reset control line are required.

From 3 to 10 μ s are typically needed to access each pixel. To reduce the time needed to raster through all pixels the columns are subdivided into groups each served by a separate output buffer. The 2048x2049 pixel array you will be using has 32 outputs each serving 64 consecutive columns.

Note that since the exposure time is defined by the times at which the pixels are read, then the exposure for the last pixel is displaced from the first by the time it takes to scan through whole array. ($\sim 10\mu\text{s} * 2048 * 2048 / 32 = 1.3\text{s}$)

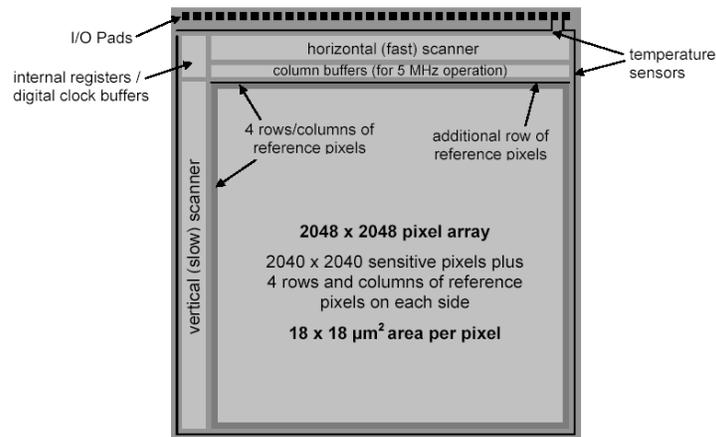


Figure 6: Teledyne HAWAII-2RG multiplexor layout.

Consequences of on-pixel integration

The good news:

- *Electronic shuttering*: Exposure time is the time between initial and final reads, not time since reset. Charge accumulated between reset and first sample has no effect on noise but can consume dynamic range and affect linearity.
- Since *readout is non-destructive*, noise reduction is possible by combining multiple samples. (Fowler Sampling, Sample Up the Ramp). However the improvement isn't as good as \sqrt{N} due to temporal correlations: there is significant noise power or systematic drift on frame-to-frame timescales.

The bad news

- *CDS occurs across the exposure time* (unlike a CCD where CDS occurs within one pixel time). Consequently IR detectors must be DC coupled and are at least 1000 times more sensitive than CCDs to electronic drifts and temperature changes.
- Because each pixel has a different signal path, *there is no "overscan"* (as in a CCD) to provide an accurate zero-point reference. The next best thing is to use dark pixels in the image area, or unconnected pixels around the edges to mimic the zero-point drift of the image pixels. These "reference pixels" don't tell us where zero is but do tell us how much it has changed.

- *Read noise increases with exposure time.* With good electronics, the read noise is reduced to the 1/f noise in the pixel buffer transistor and detector material.
- The charge to voltage conversion is *non-linear*. The signal is accumulated on the detector diode capacitance. The reverse bias applied by the reset is discharged by the photocurrent causing the width of the depletion region to be reduced, so the diode capacitance increases: the voltage change of for a given charge increment drops.
- When observing bright sources or in high background, substantial charge can be accumulated between the reset and first sample, eating into the apparent dynamic range.

Other peculiarities of IR detectors (compared to CCDs)

- Many parameters vary considerably from pixel to pixel. (e.g. dark current, QE, noise, temperature sensitivity)
- Dark current is higher and is more steeply dependent on temperature. In the best IR detectors this is just due to the lower bandgap, but it is often the case that imperfect surface passivation during manufacture degrades the dark current and causes large dark current variations from pixel to pixel.
- Dark current takes several hours to fully stabilize after a perturbation such as a temperature or bias voltage change (eg cycling power)

Experimental setup

The test system is optimized for characterizing the noise of large format astronomical IR detectors in the dark or with very low illumination. The detector is mounted inside a light-tight, cold chamber with only a narrow serpentine, black painted vent tube to allow evacuation. All wiring is potted in silver filled epoxy to thermally short it to 77K and block light leaks.

LEDs with emission wavelengths 1050, 1300 and 1550 nm are mounted within the dark chamber to provide a light sources without light leaks. Both LED current and temperature are precisely controlled to assure high flux stability.

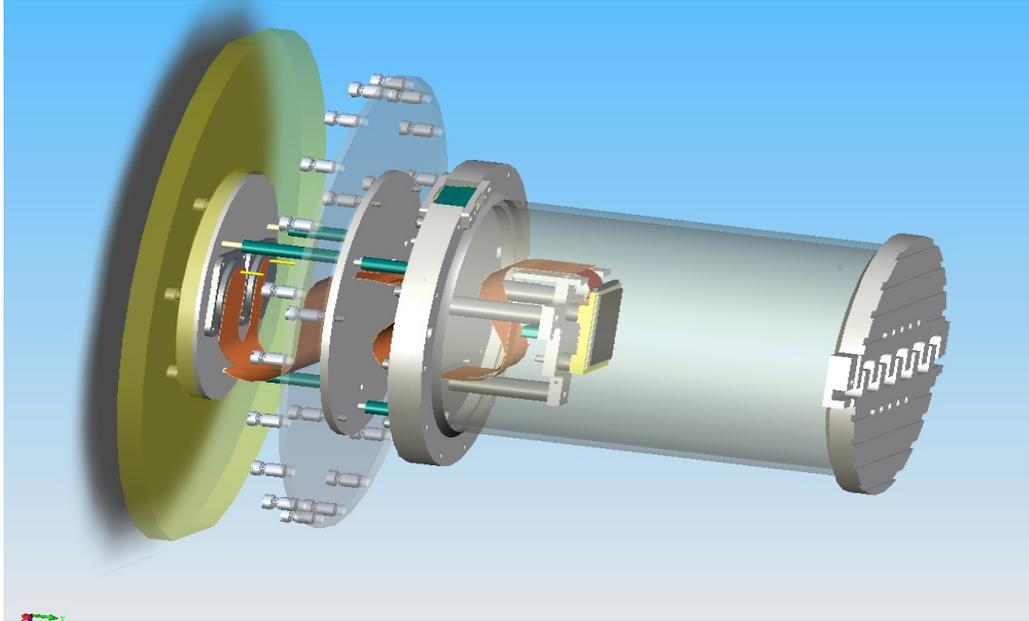


Figure 7: detector mounted on temperature-controlled plate in light-tight chamber. The surface at right is bolted to the liquid N₂ tank (not shown). For clarity, only the hatches for the vacuum enclosure and radiation shield have been shown.

A mask ~ 1mm above the detector deliberately obstructs ~20% of the image area to provide some recognizable image structure. The image area is fairly uniformly illuminated. A neutral density filter is placed across the mask aperture to provide 1000 fold flux attenuation, so that the LED can be operated at a high enough current (1mA) to minimize the intrinsic fluctuations in LED brightness.

The detector and LED temperatures are servo-controlled to millikelvin accuracy.

Electronic (inverse) gain will be set to 4.2 e-/ADU.

The detector to be tested is a Teledyne HAWAII-2RG with 1.7 μ m cutoff HgCdTe detector layer operated at 140K. The CdZnTe substrate layer has been removed so that typically only one pixel is affected by each cosmic ray.

Precautions --- not an issue for your data analysis !

The large format IR detector to be used in this experiment is being actively studied as part of the research program for the Super Nova Acceleration Probe satellite. Science grade detectors such as this have a market value approaching \$500K. They can be destroyed by

- a small mechanical shock such as striking the dewar with, or against, a hard object when the detector is cold;
- electrostatic discharge;
- heating or cooling faster than 0.5K/minute;
- heating above 60C.

The temperature control system is designed to prevent excessive slew rate, including when the power is turned off. To avoid electrostatic discharge, please avoid touching the circuit cards or cabling.

The fragility of the device when cold is due to the stress built up between layers with dissimilar coefficient of expansion. There are more than a few examples of early devices shattering.

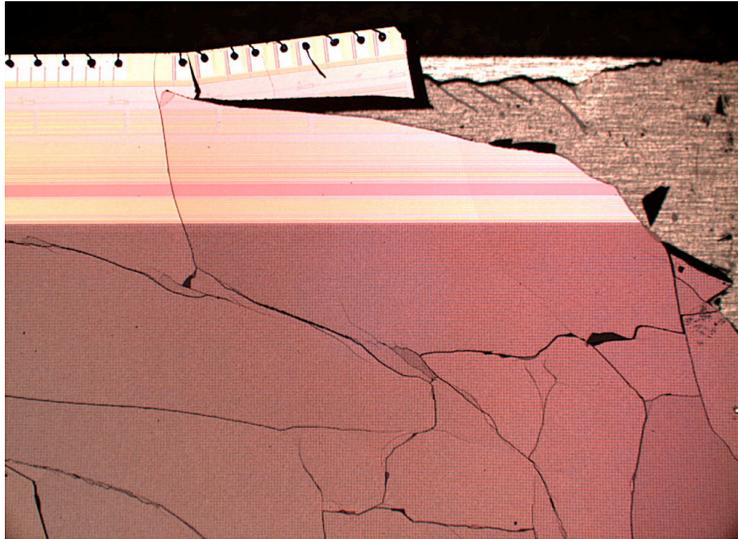


Figure 8: This bare HAWAII-2RG multiplexor shattered due to the stress induced primarily by differential thermal contraction. It was pushed over the edge when the stiff dewar was placed (gently!) on the hard lab floor.

The newer device under test is more shock resistant but you don't want to be the one who proves this to be wrong, so please move carefully in the proximity of the dewar to avoid bumping it.

Instructions...

Pairs of students will work together on the analysis tasks below. Each pair will be working with a separate user account and with their own copy of the data. VNC will be used to access the host machine.

Familiarization with basic read modes.

1. Take minimum length dark exposure.

[<.../Ay105/basic_modes/minexp/>](#)

This is best done manually using the GUI.

<Your TA or detector team member will provide a demo of the test setup.>

Record subtracted and unsubtracted frames.

Display initial, final and subtracted data. [IRAF “display” command; ximtool]

Measure stddev and mean of each frame in a (large) statistics box which is clear of major cosmetic defects. [IRAF imstat]

Observe how:

- The mean and stddev are dominated by a large fixed pattern.
- Compare the histograms of raw and subtracted frames. [IRAF imhist]
- Adjust the video input offset bias voltage. *<See your TA or detector team member to learn how.>* Why might we care about this offset, given that it has so little effect on the signal in subtracted frames? i.e. What could go wrong?

Some notes on signal polarities:

- All diode cathodes are connected in common (to Dsub). For reverse bias Dsub (cathode) is at a higher voltage than Vreset (anode). So, signal integrates positive as photocurrent discharges the reverse bias.
- The multiplexor is non-inverting.
- External electronics are non-inverting for signal and inverting for input offset voltage.
- **The ADC is peculiar in that it produces code 65535 at zero input voltage and 0 at maximum input. When looking at unsubtracted data the signal integrates from the high numbers towards zero.**
- Increasing the input offset voltage (or ADU) has the effect of producing higher output codes for raw data.
- The inversion in the ADC is removed by subtracting signal from reset level for CDS instead of the reverse: CDS frames integrate from (near) zero to positive numbers as one would expect.

There also exists a command line interface for those who prefer but this is probably an unnecessary complication. For reference only:

Script= observe

http://www.astro.caltech.edu:8080/DETECT/lab/iraf_arcview_pkg#observe)

2. **Take 2s exposure with light.**

[<.../Ay105/basic_modes/exp2s/>](#)

(LED current was manually adjusted for a signal intensity ~1200 ADU/s/pixel.)

Display initial, final and subtracted frames.

Measure stddev and mean of each frame in the same statistics box as above.

Compare to the equivalent dark frames in the section 1.

Note how the signal pattern is quite different from offset structure.

3. **Take CDS frames at various exposure times and intensities.**

[<.../Ay105/basic_modes/cds/>](#)

[Script= [DOSEQUENCE](#)]

[http://www.astro.caltech.edu:8080/DETECT/lab/iraf_lab_pkg#dosequences]

Exposure times = 2,4,8,16,24,3248s

Fluxes ~ 0, 1000 ADU/s [Manually adjust LED current each time.]

Plot our curves showing mean of CDS frame vs. exposure time for each flux. This is the only data the astronomer normally sees.

[Optional script= [GENPLOT](#), generates arbitrary plots over sequences of images]

[http://www.astro.caltech.edu:8080/DETECT/lab/iraf_lab_pkg#genplot]

Do the curves show the same saturation level?

Extrapolating back, do they all intersect at exposure time = 0 ? Let's investigate further...

4. **Sample Up the Ramp**

[<.../Ay105/basic_modes/sur/>](#)

LED brightness adjusted to ~1200 ADU/s.

Switch to SUR Readout mode at 3s per sample; expose for 80 sec.

Plot mean versus time.

[Script= DOSUR [acquires](#) and (over)plots]

[http://www.astro.caltech.edu:8080/DETECT/lab/iraf_lab_pkg#dosur]

Repeat at 0, 300, 600, 1200 ADU/s [use same LED currents as above] ; Overplot;

Why is the saturation level now independent of flux?

Fit straight line to the low intensity points on each curve. (eg in Excel)

Extrapolate back ... do the linear fits intersect at the same point?

Infer time from reset to first sample.

What is the significance of the offset left behind after Reset? How does it affect CDS frames?

Shot noise measurement

1. Photon shot noise

[<.../Ay105/shotnoise/>](#)

Take a sequence of eleven CDS frames, 5s each, with LED brightness set to 1200ADU/s.

Select an area of flat illumination, at least 50cols * 200rows, such that all columns are read by the same output amplifier. How stable is the lamp and signal gain?

[Optional script= [GENPLOT](#)]

[http://www.astro.caltech.edu:8080/DETECT/lab/iraf_lab_pkg#genplot]

Compute standard deviation for each frame. Is this a measure of the shot noise?

Compare stddev of difference frame divided by $\sqrt{2}$, to the stddev for the single frame. Which is the correct estimate of shot noise?

Why might these estimates differ.

Let's compare this spatial noise estimate with a temporal noise ie. take many images and calculate noise for each pixel independently. Combine the 11 images above to produce mean and stddev images. [IRAF imcombine]

If we had taken thousands of images the stddev image would be a temporal noise map, however with only 11 images the sample is too small to provide an accurate estimate of the noise on any given pixel. Nonetheless, the *mean* of the noises is a precise estimator of the shot noise, once one takes into account the tendency for a small sample to underestimate the standard deviation of the parent population...

Plot the histograms. [IRAF imhist]

Compare the spatial noise (calculated above using difference frames divided by $\sqrt{2}$ to the mean of the noise map. If the latter is multiplied by $\sqrt{N/(N-1)}$ with $N=11$, to convert sample variance to population variance, then the two noise estimates should agree.

Let's take this to the extreme. Make a stddev frame from only two frames [imcombine], then multiply the mean stddev by $\sqrt{N/(N-1)} = \sqrt{2}$ to see if you get the same estimate of for shot noise.

Temporal noise includes all sources of variation, while the spatial noise is insensitive to variations which are common to all pixels. At low light levels the two noise measurement methods will only agree, if the zero point drifts are negligible. At high enough light levels the lamp intensity variations will again produce correlated signal

variations, which will make temporal and spatial noise disagree. Once the equivalence of temporal and spatial noise has been verified for dark frames in a particular detector system, spatial noise is generally used since it is an easier measurement.

Read noise

So far the primary noise source examined has been the Poisson noise in the photon flux, usually called “photon shot noise”. We call the noise originating in the electronics the “read noise”. We distinguish read noise from the Poisson noise associated with thermally generated “dark current”. The word “glow” is reserved for photo-emission from transistors or photo-emitting defects (“PEDs”) within the detector itself, to distinguish this from thermally generated dark current. This device exhibits none.

The read noise is strictly the Y intercept on a plot of variance vs signal. In “science grade detectors” dark current and glow are so low that noise in a minimum length exposure is an adequate measure of “read noise”.

- Take eleven 5s exposure *dark* frames, then compute temporal and spatial noises as you did for shot noise: [<../Ay105/CDS_readnoise/dark5s/>](http://www.astro.caltech.edu/~Ay105/CDS_readnoise/dark5s/)
Note that the difference of two short exposures is used for the spatial noise estimate, in this case to suppress fixed patterns due to hot pixels (high dark current), or residual offset structure which sometimes remain after the CDS subtraction. How well do spatial and temporal noises agree once the temporal noise has been scaled by $\sqrt{11/10}$?
- {Skip this} Now compute spatial noise from a pairs of 5s, 50s, 150s, and 300s dark exposures. Does spatial noise increase with exposure time? Such an increase can be due to 1/f noise sources within a pixel, or to the shot noise associated with dark current. How can you tell the difference?
[Optional script= [GENPLOT](#)]
[http://www.astro.caltech.edu:8080/DETECT/lab/iraf_lab_pkg#genplot]

Bias voltage or temperature variations during the exposure will affect all pixels to approximately the same extent. Temporal noise will increase at longer exposure times as a result of these drifts, more than the spatial noise which is only affected by differences from pixel to pixel in the response to bias and temperature. These correlated drifts are easily detected as increased scatter in the means in a sequence of dark images. The test system you are using has been optimized to minimize such effects, which can also be mitigated to some extent by reference pixel subtraction.

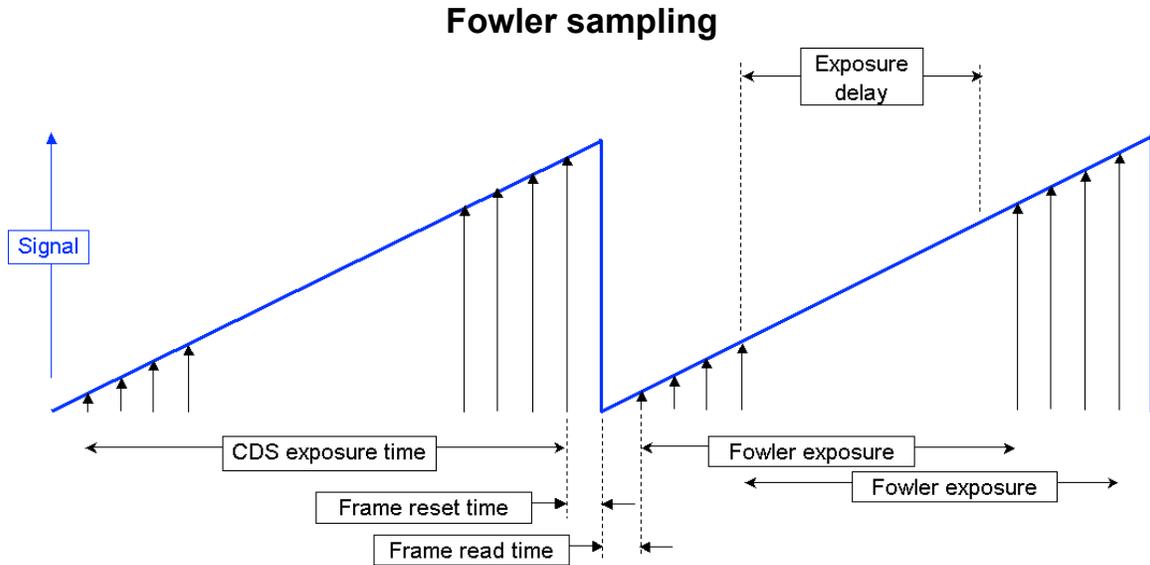


Figure 9: Comparison of CDS timing and Fowler Sampling,

Classical “Correlated Double Sampling” involves nothing more than sampling the signal after reset and subtracting this from the final value. Any signal accumulated outside these two samples is invisible, except for its effect on linearity and dynamic range.

Fowler sampling is a simple variant for improving the read noise. Since the signal is read non-destructively, multiple samples at the beginning and end of the exposure can be averaged to reduce the effective read noise. (The detector readout software will do this calculation in real time). An alternative way to think of this is as a set of partially overlapping CDS pairs, which are averaged. The effective exposure time is the difference between the Nth sample in each group and not time between first and last sample. See Figure 9.

In the limiting case where the Fowler samples are tightly packed at either end of the exposure, the exposure duty cycle is high, and the signal and photon shot noise is nearly identical in each “fowler pair”. When the fowler sampling time becomes a larger fraction of the total integration time then the shot noise in each pair is only partially correlated between pairs.

- Can you see why it is better to scan the whole array N times rather than reading each pixel N times, averaging, then moving to the next pixel?

1. Measure and plot spatial Read Noise vs #samples for *constant* exposure time.

[<.../Ay105/fowler/1_dark60s/>](http://.../Ay105/fowler/1_dark60s/)

Take darks with *constant* effective exposure time = 60s

#samples= (1,2,4,8,16,32)
[Script= DOFOWLER (acquires & plots)]
What could cause the noise reduction to be slower than \sqrt{N} ?

2. **Measure effective exposure time vs # samples.**
<.../Ay105/fowler/2_flat5s/>

Take pairs of flats (flux~300 ADU-/s) at various fowler sampling depths with constant exposure delay (5s): plot mean signal vs. number of samples:
#samples= (1,2,4,8,16,32)
[Script= DOFOWLER (acquires & plots)]
[http://www.astro.caltech.edu:8080/DETECT/lab/iraf_lab_pkg#dofowler]
[GENPLOT, generates plot of mean versus number of samples.]

Infer effective exposure time from signal vs. #samples. Estimate of frame scan time.

3. **Measure and plot Photon Shot Noise vs #samples**
<.../Ay105/fowler/3_flat60s/>

Take exposures with illumination [300 ADU-/s] for *constant* effective exposure time =60s
#samples= (1,2,4,8,16,32)
[Script= DOFOWLER (acquires & plots)]
Does photon shot noise scaling with number of samples the same way as read noise?

Variance curve: to measure e-/ADU

All of the measurements so far have been in the arbitrary units produced by the A/D converter. While we can use test instruments to measure the voltage gain (ADU/ μ V) of the electronics and multiplexor, it is the capacitance of the input node, which determines the conversion from electrons to microvolts. Last year a clever way was demonstrated by Gert Finger at ESO to directly measure this tiny capacitance (typically 45fF). However, we will take the easier route and derive the overall inverse gain (e-/ADU) of the signal path, from a basic property of Poisson statistics:

$$\text{variance} = \text{mean} \\ \Sigma^2 = X$$

This statement is only true for the fundamental unit of quantization, electrons (or photons). To rewrite this relation in derived units we define the conversion factor (1/gain):

$$\alpha = e-/ADU$$

Then by definition: $X = \alpha x$ $x = \text{mean in ADU}$
 $\Sigma = \alpha \sigma$ $\sigma = \text{standard deviation in ADU}$

So $\Sigma^2 = X$
 becomes $(\alpha \sigma)^2 = \alpha x$
 $1/\alpha = \sigma^2/x$

Thus system gain, $1/\alpha$, is the slope of a plot of variance versus mean, often referred to as a “photon transfer curve”. Astronomers usually speak of “gain” when they mean inverse gain, α . This confusion is alleviated by quoting the units, e-/ADU.

- For pairs of flat field exposures at various intensities, which may include zero (or not), plot variance of the difference vs. sum of means. Slope = ADU/e-.
 [Script= [DOVARIANCE](#)]
 [http://www.astro.caltech.edu:8080/DETECT/lab/iraf_lab_pkg#dovariance]

[<.../Ay105/variance/>](#)

- What do you think the effect of fowler sampling would be?
- Optional.... Compare e-/ADU estimates using only low flux points, only high flux points, and all data points. Why might these estimates vary slightly?

Moore Correction

[skip this section if time is short]

As usual, matters are more complicated for infrared detectors than for CCDs. In a CCD, the charge stored in neighboring pixels does not effect the measurement of charge on the pixel being read. IR detectors integrate their photogenerated charge not only on the photodiode capacitance but also the input capacitance of the buffer MOSFET and any stray capacitance to neighboring pixels. Stray capacitance can occur between bump bonds or in the multiplexor, though not within the detector material since each pixel is surrounded and separated by regions of high carrier density, which serve as a Faraday shield.

In 2003, Andrew Moore of the University of Rochester pointed out how the interpixel capacitance causes us to underestimate the shot noise when using pixel to pixel variance. Since the inter-pixel capacitance has no effect on the mean, the photon transfer curve under-estimates gain ($1/\alpha$). (An aside: this discovery led to the widespread realization that everyone’s quantum efficiency measurements had been too optimistic for decades, and providing the impetus for the manufacturer find the cause and fix the QE deficit.)

Dr Moore wrote the software you are about to use, to derive the capacitances to the surrounding pixels and the necessary correction to inverse gain derived from photon

transfer, by computing the 2D autocorrelation function of the difference of frames with identical flat illumination with sufficient intensity for photon shot noise to dominate.

- Apply Moore Correction and correct your estimate of e-/ADU.
[Script= [DOMOORE](#) (acquires images & calculates correction)]
[http://www.astro.caltech.edu:8080/DETECT/lab/iraf_lab_pkg#domoore]
<.../Ay105/moore/>

Dark current

<.../Ay105/Dark_current/>

Measure dark current, at different temperatures using long SUR exposures:
140K, 150K, 160K

[Script= [DODARK](#) acquires, changes temperature & plots results]
[http://www.astro.caltech.edu:8080/DETECT/lab/iraf_lab_pkg#dodark]

- Both DC offset and dark current vary with temperature. In addition, dark current may take longer to stabilize than the temperature. Can you see evidence for all three effects in your data? [Use IRAF hselect to get detector temperature from image headers]
- Plot equilibrium dark current in e-/s against temperature.
- Compare with theoretical dark current equation:

$$I_{\text{dark}} = C * T^{1.5} e^{-g/(2kT)}$$

where,

C = constant determined by pixel geometry and material quality

g = bangap = $h*c$ / cutoff_wavelength

k = Boltzman's constant

T = temperature (kelvin)

~~ THE END ~~

APPENDIX

The following discussion is provided to show how the parameters you have measured can affect astronomical observations. It is here for your benefit, but is not a required part of the lab....

Exposure time calculations

Telescope time is expensive, so lets examine how long it takes to reach a given signal to noise ratio and see how this is affected by detector performance.

$$S/N = \text{flux} * t / \sqrt{\{ \text{flux} * t + I_{\text{dark}}(T) * t + N_R^2(t, \text{fowler_samples}) \}}$$

The first two limiting cases are well known :

- Photon shot noise dominated:
 $S/N \cong \sqrt{(\text{flux} * t)}$
- Read noise dominated:
 $S/N \cong (\text{flux} * t) / N_R$
- Dark current dominates (long exposures on very weak sources)
 $S/N \cong \text{flux} * \sqrt{t / I_{\text{dark}}}$

The third case often becomes relevant when doing spectroscopy where the light is dispersed sufficiently for the flux per pixel (between sky emission lines) to be very low. In CCDs exposure times are limited by the need to subdivide exposure time and median filter the frames to remove the localized bright spots caused by high-energy particle detection, so dark current usually doesn't dominate ... unless pixel binning is invoked.

However dark currents tend to be higher in IR detectors. Furthermore, the particle events can be mitigated through signal processing of a series of non-destructive readouts, so that a single long exposure is possible. This in combination with noise reduction via multiple sampling tends to make the dark current more relevant:

	1.7μm cut off	2.5μm cut off
Read noise at 85% exposure duty cycle	6 e-	4 e-
Dark current floor (at very low T)	0.002 e-/s	0.01 e-/s
Exposure time where dark shot noise degrades read noise by 20%	3000 s	320 s

A spreadsheet will be provided on request which plots time required time to achieve a given S/N, for fluxes = 0.001, 0.01, 0.1 ... 1000 e-/s. You can plug in read noise and dark current numbers to compare two different scenarios as shown in Figure 10.

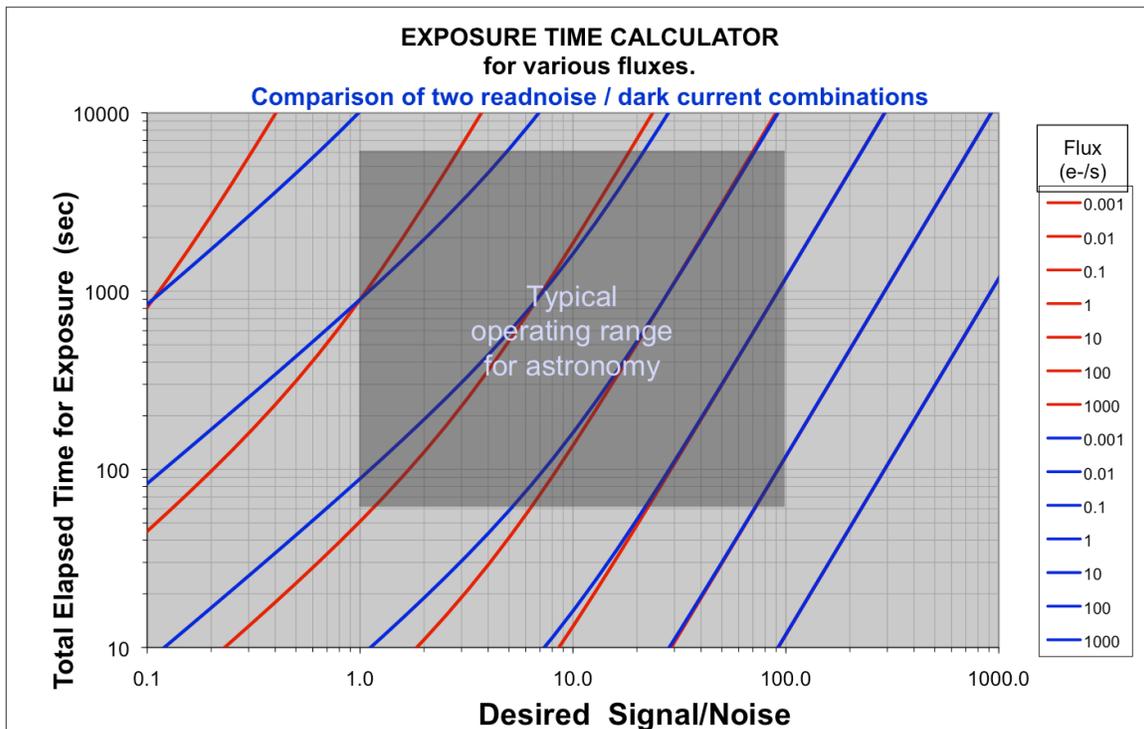


Figure 10: To save you time, a spreadsheet will be supplied which over plots two sets of curves, as shown here, allowing you to compare detector performance combinations. Here, red is for a typical of a 2.5um cutoff detector with $I_{\text{dark}}=0.05\text{e-/s}$ and $N_R=3.5\text{e-}$, while blue is for a typical 1.7um cutoff detector with $I_{\text{dark}}=0.002\text{e-/s}$ and $N_R=6\text{e-}$ (for 85% duty cycle when fowler sampling).

If Fowler Sampling is used “t” in the above formula becomes is the effective exposure time and not the elapsed time. Initially we can ignore this small adjustment.

- Make the noise the same on both detectors and play with the dark current values.
- Make the dark currents the same and vary the noise.
- Identify the read noise, dark current and photon shot noise limited regimes.
- The noise values used to make the curve sets in Figure 10 are typical of fowler sampling. Instead plug in the noise you measured for CDS. Compare with the $\sim 6\text{e-}$ possible with fowler sampling. Set $I_{\text{dark}} = 0.01 \text{ e-/s}$ in both cases.

How might one decide what is the optimum fowler sampling depth?....

The noise behavior of these devices is more complex than for CCDs and poorly understood by many users. The noise increases slightly with exposure time, while noise is reduced by multiple sampling. As the number of fowler samples increases, the observing efficiency (effective exposure time) decreases, so the attempt to reduce the noise also reduces signal. Confused?

How does one sort out these competing effects? Developing approximate analytical expressions for each effect then solving for maximum S/N, is commendable since it will provide the answer. However it is possible to do this without gaining any insight into the principal effects.

For the detectors you are testing, we have found that this complex behavior can be reduced to a simple trade between noise and observing efficiency. Figure 11 is derived from measurements and tells us that the noise increase due to longer exposures (compare different curves) can be compensated by the extra fowler sampling.

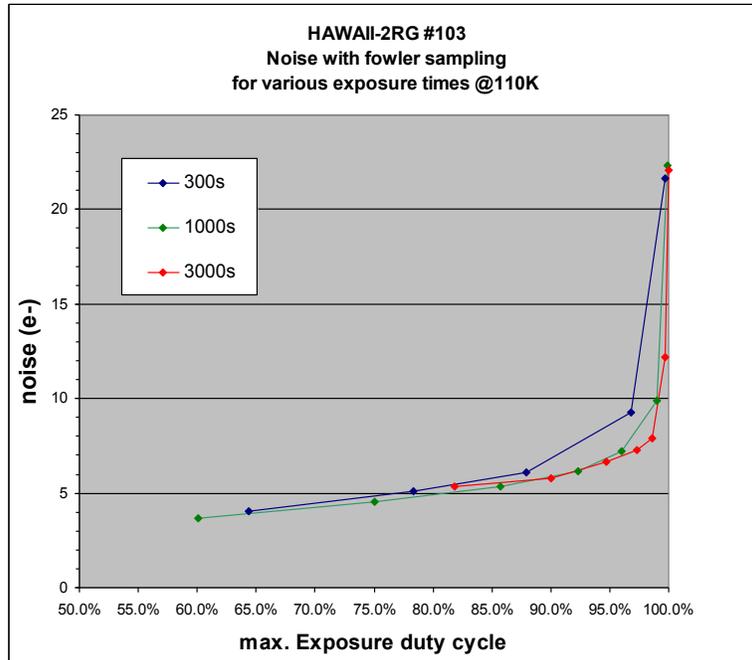


Figure 11: When noise is expressed as a function of exposure duty cycle rather than number of fowler samples, we find that degradation in noise at longer exposures is compensated by the larger number of fowler samples possible (for a given duty cycle). [110K, Dsub-Vreset=1V]

At high fluxes, we want the effective exposure time to be maximized: simple correlated double sampling is best.

When read noise dominates, some fowler sampling is desirable. How much is optimal? The existence of a sharp knee in the above curve and the fairly slow dependence of noise on exposure duty cycle below this knee, suggests that the optimum fowler sampling depth may be just below the knee.

As a first guess one could choose Fowler sampling depth to give 95% exposure duty cycle and read the noise from Figure 11. How much does this affect the exposure time contours? Is it worth optimizing further? Are there circumstances when one would choose a lower exposure duty cycle?

To find out, one would need to modify the exposure time calculator to take into account the loss of exposure duty cycle when fowler sampling. As you change noise you need to

rescale the elapsed time. You can try plugging in a few combinations of noise and exposure duty cycle read from Figure 11 into the exposure time calculator spreadsheet.