CCD Cameras

Stefan Ivanov

This note is geared towards students preparing for IOAA, but I think some bits will be interesting for a general audience too. Before going in, you should feel comfortable with stellar magnitudes and telescope optics. You'll also need to know how a p-n junction works, which in turn requires learning about doped semiconductors. Check out this video for an excellent explanation.

1. CCD Operation

First we'll touch on the photodiode, which is the main component of charge-coupled devices (CCD). Consider a p-n junction where the p-type material is connected to the negative terminal of a voltage source, and vice versa. This junction is reverse-biased: the electric field is such that the charge carriers at both ends of the depletion region are pulled away, so the depletion region will widen. There won't be any substantial current in the junction by default, but that's not our aim at the moment anyway!

In the depletion region the charge carriers on opposite sides of the junction have recombined, leaving a net charge density associated with the lattice. Thus, locally there is a large electric field which serves to block any further diffusion between the p-type and the n-type material in the steady state. This is useful. We can choose to work with a material (e.g. silicon) for which photons of a given wavelength are able to knock off the bound electrons in the depletion region, resulting in a free electron and a hole. Usually this pair would recombine and that would be the end of the story. But here the large electric field will send our hole-electron pair in opposite directions, which, if you squint your eyes, is just a net current.

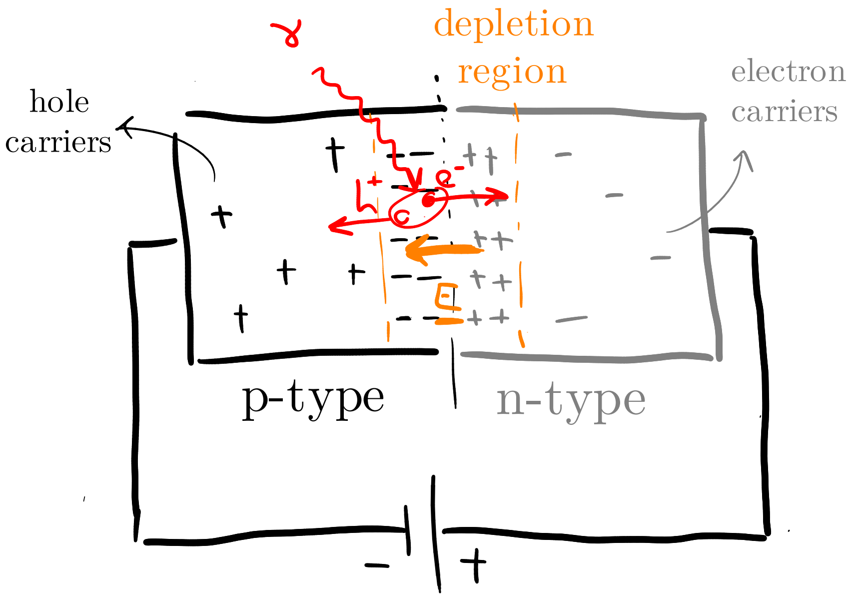

Figure 1

Notice that the number of electrons accumulated this way indicates the number of photons that have struck the depletion region. In reality, not every photon is converted to a free charge, but the conversion rate is still impressively high. This quantity is called the quantum efficiency η. It varies with the wavelength of the incident light; at some wavelengths CCD cameras can have $\eta > 90\%$.

A CCD is essentially a grid of photodiode-like detectors. Each such detector is one pixel. When you take a photo with some fixed exposure time, the pixels will absorb some photons and the resulting electrons will accumulate locally. After we're done with the exposure, we want to read out the number of electrons at each pixel. The electrons are moved along the grid by applying appropriate voltages. They finally accumulate at a capacitor, where we get a voltage signal proportional to the charge. The signals from all the pixels are then combined to give you the full image. This chain of events is similar to human vision. Photons strike the photoreceptive cells in the retina, giving rise to neural signals.

The brain then processes these signals to create a first-order approximation of the outside world.

Now, the CCD gives you just a single number per pixel – the counts. This is the number of recorded electrons times some arbitrary constant factor (called the gain) which depends on the internals of the device. The information from the counts is sufficient to recover only a black-and-white image. The way to get a colour image is to stack together shots of the object taken under a few different filters. Indeed, if any given pixel has more counts in band B (450 nm) compared to band V (550 nm) or band R (650 nm), you'd know to expect a bluish colour there.

There's another approach available if you want your colour image immediately and don't mind cutting corners on image quality – it is to use a filter mosaic. Perhaps the most popular arrangement is the Bayer filter, shown on Figure 2. Cover a quarter of the pixel array with tiny red filters, a quarter with blue, and the remaining half with green. The space under any little filter will then carry information for a single band only. However, we can still guess its expected counts in the other bands using extrapolation from those neighbouring bits that do carry legitimate information for those bands.

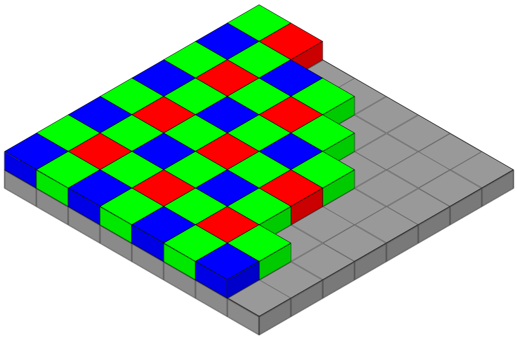

Figure 2

That's how everyday cameras work, and incidentally, this is also similar to what the eye does. The cone cells in the retina come in three types, which are sensitive to red, green, and blue light, respectively. The signals from these cells are weighted in the brain, and this gives us a sense of colour.

Example 1. Digitisation (Russia 2025). A telescope with aperture $D = 20\,\text{cm}$ and a focal ratio $A = 1/10$ was used to take a picture of a star. The table below shows the counts at each pixel of the CCD array.

(a) Find the pixel size d of the CCD array.

(b) Find the apparent magnitude of the star m.

(c) Find the night sky brightness μ (in mag/arcsec^2), assuming that 60% of the background counts are due to noise.

Assume that the detector's gain is $1\,\text{e}^-/\text{count}$. The quantum efficiency of the detector is $\eta = 86\%$ and the exposure time is $\tau = 10\,\text{s}$. The losses in the optics are $q = 10\%$. The astronomical seeing (the diameter of the disk from a point source blurred due the atmosphere) is $\phi = 1.1''$. The photon flux from a star of magnitude zero is $\Phi_0 = 10^6\,\text{cm}^{-2}\,\text{s}^{-1}$.

101	98	97	98	99	100	99
99	98	101	170	100	99	97
97	96	766	777	771	102	100
101	190	795	809	802	202	99
100	102	780	791	771	100	95
104	105	98	196	100	99	97
105	104	101	104	102	104	103

Solution. (a) There are actually two reasons why the star isn't seen as a point source. One is the seeing, and the other is diffraction at the aperture. The diameter of the Airy disk is $\theta = 2.44\lambda/D = 1.4''$. Seeing and diffraction can be treated as two independent sources of random error. Hence, the angular diameter of the star's image can be estimated as

$$\alpha = \sqrt{\phi^2 + \theta^2} = 1.8''.$$

The focal length of the telescope is $F = D/A = 2\,\mathrm{m}$. Looking at the counts, it seems that the stellar disk's diameter α corresponds to about 3 pixels. Thus $F\alpha = 3d$, and $d = \frac{D\alpha}{3A} = 5.7\,\mu\mathrm{m}$.

(b) We need to sum the counts at the pixels that contain signal from the star, subtracting the background (100 counts) from each pixel. We get $N' = 70 + 666 + 677 + 671 + 90 + 695 + 709 + 702 + 102 + 680 + 691 + 671 + 96 = 6520$ counts. The number of photons is then $N = N'/(1-q)\eta = 8424$. From this we can obtain the photon flux $\Phi = N/(\pi D^2/4)\tau$, and all that is left is to compare with the standard source:

$$m - 0 = -2.5\lg\left(\frac{\Phi}{\Phi_0}\right) \quad \Rightarrow \quad m = 13.9.$$

(c) We're told that 40% of the background is due to the night sky, meaning that there are $(0.4 \times 100)/(0.9 \times 0.86) = 52$ incident photons per pixel. One pixel corresponds to $(d/F)^2 = 8.1 \times 10^{-12}\,\mathrm{sr} = 0.35\,\mathrm{arcsec}^2$, so there are about 148 photons from $1\,\mathrm{arcsec}^2$ of the sky. We divide this by $\pi D^2/4$ and τ to get the photon flux $\varphi = 470\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}\,\mathrm{arcsec}^{-2}$. Thus

$$\mu - 0 = -2.5\lg\left(\frac{\varphi}{\Phi_0}\right) \quad \Rightarrow \quad \mu = 18.3\,\mathrm{mag/arcsec}^2.$$

2. Signal and noise

Given that the detector's output depends on the photon count per pixel, it's fairly intuitive that what we call signal correlates with the number of incident photons rather than their total energy. Denoting the photon flux per pixel by N (with units [s^{-1}]), the total signal[1] on n pixels after an exposure time t is ηnNt.

However, there are a few effects because of which the counts registered by the detector won't match this expression exactly. First and foremost, the number of incident photons is subject to deviations from nNt. When we say that the observed source has a constant photon flux (say, 4/s), we mean that this is what we expect on average if we could sample the results of a trillion observing sessions. In reality, any two photon emissions are uncorrelated, meaning that they don't know about each other, and they cannot possibly coordinate. So, the notion that you receive a photon, you wait exactly 0.25 s, then you receive another photon, then wait 0.25 s again, is completely false. On average, within a given second, you expect to receive 4 photons – that's just how the source is – but it may occur that you get 6 photons, or even 8, though this gets increasingly unlikely. Processes with independent events and a constant mean rate obey the so-called Poisson distribution. Here, the probability that you get k occurences when the value expected from the mean is λ will be given by the expression

$$P(x = k) = \frac{\lambda^k e^{-\lambda}}{k!}.$$

A plot of this is shown on Figure 3. This is all rather unintuitive and requires further explanation – for example, where did that e even come from? If you're interested in the nitty-gritty, I direct you to Chapter 30.8 of [1], but now we need to make haste.

For our purposes the important thing is that Poisson distributions have a standard deviation of $\sqrt{\lambda}$. The standard deviation is a measure of the expected variation from the mean value, and it is

[1] This is in electrons rather than counts because we don't account for the gain here. But the gain just multiplies everything by a constant factor, so it doesn't really matter for what we'll cover next.

used to indicate your error when you declare that the output of a Poisson process is indeed the mean. If there's a Poisson process with a rate of 16/s running for one second, the standard deviation is $\sqrt{16 \times 1} = 4$. Hence we claim that we expect to observe 16 ± 4 events within this one second. This is good enough – provided that λ is large, we can show that there's about 70% chance that the number of observed events falls between $\lambda - \sqrt{\lambda}$ and $\lambda + \sqrt{\lambda}$ (here, 12 and 20).

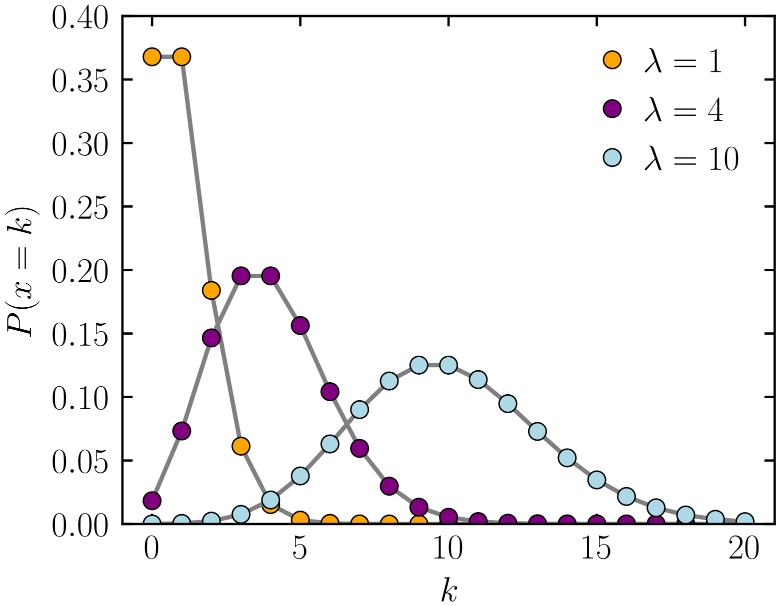

Figure 3

In photography, deviations from the mean photon flux are problematic. What you actually want to probe is the mean photon flux from each object, because that's what really tells you how objects compare in brightness (or, how the world really *looks*). But you can't avoid the inherent variations. We call these Poisson noise, or shot noise. If the expected signal is $\eta N n t$, then the Poisson noise is $\sqrt{\eta N n t}$. The relative error is thus

$$\frac{\text{noise}}{\text{signal}} = \frac{\sqrt{\eta N n t}}{\eta N n t} = \frac{1}{\sqrt{\eta N n t}}.$$

This decreases with time, so a longer exposure means a more accurate image (and, of course, brighter too). Here's a comparison of different integration times:

Figure 4

Now we'll discuss another source of noise, namely dark current. Even when a photodiode isn't pointed to some specific source, it is still influenced by the background. As the device is kept at some fixed temperature T, hole-electron pairs can form spontaneously due to thermal excitations. These still lead to counts and can't be distinguished from the electrons generated due to source photons. We call the current due to such pairs dark current. We'll denote it by D, the units being electrons per second per pixel (quantum efficiency doesn't play a part here). We see that on top of the useful signal ηNnt we get an extra signal Dnt with standard deviation $\sqrt{Dnt}$. The dark current grows exponentially with temperature, so we want to keep our sciencey CCDs cool. This is done with liquid nitrogen or a Peltier element. Conversely, in everyday CCD use exposure times are shorter and there are lower demands on quality, so we don't bother with cooling.

It is possible to remove the extra signal from the dark current in a picture by taking another picture with a closed shutter at the same CCD temperature. This is called a dark frame. By subtracting the counts of the dark frame from the counts of the original picture, we get a "cleaned-up" signal. However, the dark frame doesn't exactly match the dark currents in the original, because this is still a Poisson process. Therefore, there's an error term $\sqrt{Dnt}$ that we can't do much about even with the subtraction.

Apart from that, the CCD also has read-out noise. The electronics are not perfect, and they will miscount the number of electrons with some standard error R (a fixed number per pixel). Generally, faster read-out will bring about an increase in R. We'll conclude with a couple of examples.

Example 2. Counting Photons (IPhO 2022 2B). The absorption of a photon by a CCD camera leads to the emission of an electron within the apparatus. This occurs only if the photon has sufficient energy to excite an electron across an energy gap ΔE_{g}. Assume that every photon with sufficient energy succeeds. There is also leakage of electrons across the gap caused by the temperature of the CCD camera; this is the dark current i_{d}, measured in number of electrons per second. It depends on the temperature T according to

$$i_{\mathrm{d}} = i_0 e^{-|\Delta E_{\mathrm{g}}|/6k_{\mathrm{B}}T},$$

where i_0 is a constant. The graph below shows how dark current varies with temperature for the James Webb Space Telescope.

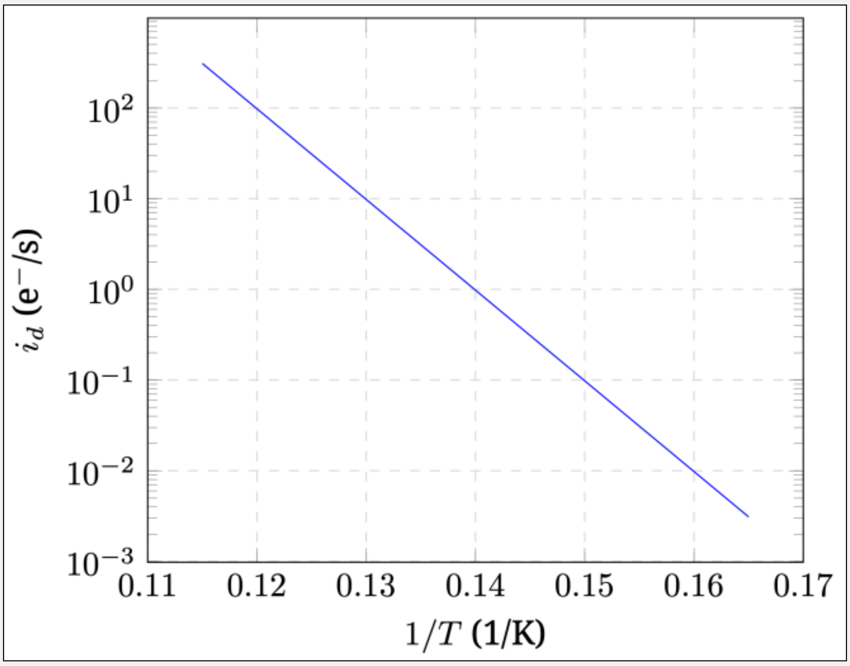

(a) Using the dark current graph, provide an order of magnitude estimate for the temperature T_{s} of a distant source of thermal photons that would just be capable of exciting an electron on the pixel.

The electrons are collected in a capacitor, and after an exposure time τ, the electrons are counted. There are three main sources of uncertainty in the process: a fixed uncertainty in the counting process called read-out noise; a Poisson distribution error associated with the dark current, and a Poisson distribution error associated with the detected incoming photons.

(b) Write an expression for the total count uncertainty σ_t if there is a read-out noise σ_r, a dark current i_d, an incoming photon rate p, and an exposure time τ.

The measured photon count is equal to the number of electrons in the capacitor, minus the number of electrons associated with the dark current. From now on, assume the exposure time is $\tau = 10^4\,\mathrm{s}$ and the read-out noise is a fixed $\sigma_r = 14$.

(c) Assume an operating temperature of $T_p = 7.5\,\mathrm{K}$. Calculate the minimum photon rate p so that the photon count is ten times the count uncertainty.

(d) Assuming all photons are barely capable of exciting an electron across the band gap, what is the intensity of the source of photons found in (c) on the primary mirror, which has an area $A = 25\,\mathrm{m}^2$? Express your answer in $\mathrm{W/m^2}$.

Solution. (a) If you look on the exponents on the y-axis, you'll notice that this is a graph of $\lg i_d$ against $1/T$. We know that

$$\lg i_d = \frac{\ln i_d}{\ln 10} = \frac{\ln i_0}{\ln 10} - \frac{|\Delta E_g|}{6k_B \ln 10}\left(\frac{1}{T}\right),$$

so the slope of the graph corresponds to $-\frac{|\Delta E_g|}{6k_B \ln 10} \equiv a = -100\,\mathrm{K}$. The photons from the distant source have a typical energy of $k_B T_s$, which should equal $|\Delta E_g|$. Thus

$$T_s = -6a \ln 10 = 1380\,\mathrm{K}.$$

(b) We're told in the problem statement that the quantum efficiency is 100%. There are three separate sources of error – the Poisson noise $\sqrt{p\tau}$, the dark noise $\sqrt{i_d\tau}$, and the read-out noise σ_r. The total error is obtained by adding them in quadrature, i.e.

$$\sigma_t = \sqrt{\sigma_r^2 + (p + i_d)\tau}.$$

(c) Using the standard jargon, we want a signal-to-noise ratio (SNR) of 10:

$$\mathrm{SNR} = \frac{p\tau}{\sqrt{\sigma_r^2 + (p + i_d)\tau}} = 10.$$

We use the graph to find that at $T_p = 7.5\,\mathrm{K}$ the dark current is $i_d = 5\,\mathrm{s}^{-1}$. From then on, finding $p\tau$ is a matter of solving the quadratic equation

$$(p\tau)^2 = 100(\sigma_r^2 + (i_d\tau) + (p\tau)).$$

We reach $p = 0.23\,\mathrm{s}^{-1}$.

(d) The energy of each photon is $|\Delta E_g| = 1.9 \times 10^{-20}\,\mathrm{J}$. The power on the mirror that gets sent to the CCD is $p|\Delta E_g|$, and the intensity is

$$I = \frac{p|\Delta E_g|}{A} = 1.8 \times 10^{-22}\,\mathrm{W/m^2}.$$

Here we assumed no losses in the optics.

Example 3. GOTO (IOAA 2017). The Gravitational-Wave Optical Transient Observer (GOTO) aims to carry out searches of optical counterparts of any gravitational wave (GW) sources within an hour of their detection by the LIGO and VIRGO experiments. The survey needs to cover a large area on the sky in a short time so as to search all possible regions constrained by the GW experiments before the optical burst signal, if any, fades away. The GOTO telescope array is composed of 4 identical reflective telescopes, each with 40-cm aperture and an f-ratio of 2.5, working together to image large regions of the sky. For simplicity, we assume that the telescopes' fields-of-view (FoV) do not overlap with one another.

(a) Calculate the projected angular size per unit length at the focal plane, i.e. the plate scale, of each telescope. Answer in arcmin/mm.

(b) If the zero-point magnitude (i.e. the magnitude at which the count rate detected at each pixel is 1 count per second) of the telescope system is 18.5 mag, calculate the minimum time needed to reach 21 mag at signal-to-noise ratio (SNR) = 5 for a point source. We'll initially assume that the noise is dominated by both the read-out noise (RON) at 10 counts/pixel and the CCD dark (thermal) noise (DN) rate of 1 count/pix/minute. The CCDs used with GOTO have a 6-micron pixel size and a gain (conversion factor between photo-electron and data count) of 1. The typical seeing at the observatory site is around 1.0 arcsec. Neglect diffraction.

(c) Normally, when the exposure time is long and the source count is high, Poisson noise from the source is also significant. Determine the relation between SNR and exposure time in the case where the noise is dominated by Poisson noise of the source.

Recalculate the minimum exposure time required to reach 21 mag with SNR = 5 from (b) if Poisson noise is also taken into consideration apart from dark noise and read-out noise. Note that the background sky brightness is also an important source of Poisson noise, but you should neglect it in in your calculations.

(d) The typical localisation uncertainty of the GW detector is about $100\,\mathrm{deg}^2$. We would like to cover all possible locations of any candidate within an hour after the GW is detected. Estimate the minimum side length of the square CCD needed for each telescope (in number of pixels). You may assume that the time taken for the CCD read-out and the aiming are negligible.

Solution. (a) The focal length of each telescope is $f = 2.5 \times 0.4 = 1\,\mathrm{m}$. One mm on the CCD corresponds to an angle of $\frac{1\,\mathrm{mm}}{1000\,\mathrm{mm}} = \frac{1}{1000}\mathrm{rad} = 3.44\,\mathrm{arcmin}$. The scale is $s = 3.44\,\mathrm{arcmin/mm}$.

(b) The seeing disk's diameter $\phi = 1.0\,\mathrm{arcsec}$ corresponds to $\phi/s = 4.8\,\mathrm{\mu m}$, which is less than the size of a pixel ($6\,\mathrm{\mu m}$), so the star can fit within one pixel. Next, assume that the standard source emits photons similar to those of our 21 mag source (so that the quantum efficiency of the CCD is the same for both). Then, the count rate p, quantum efficiency included, can be found from

$$m - m_0 = -2.5 \lg\left(\frac{p}{p_0}\right),$$

$$21 - 18.5 = -2.5 \lg\left(\frac{p}{1\,\mathrm{s}^{-1}}\right) \quad \Rightarrow \quad p = 0.1\,\mathrm{s}^{-1}.$$

Since we're working with a single pixel, the total read-out noise is $R = 10$ and the total dark current is $D = 0.01667\,\mathrm{s}^{-1}$. The equation for the SNR looks like this:

$$\mathrm{SNR} = \frac{p\tau}{\sqrt{R^2 + D\tau}}.$$

Again, a quadratic equation with respect to τ. After some legwork, we get $\tau = 8.7\,\mathrm{min}$.

(c) The signal is pt, while the Poisson noise is $\sqrt{pt}$. In the first part of this subtask there's only Poisson noise, so

$$\mathrm{SNR} = \frac{pt}{\sqrt{pt}} \propto \sqrt{t}.$$

We also need to redo part (b), this time accounting for Poisson noise as well:

$$\mathrm{SNR} = \frac{p\tau}{\sqrt{R^2 + (D+p)\tau}}.$$

The answer is $\tau = 11.1\,\mathrm{min}$.

(d) Each snapshot takes time $\tau = 11.1\,\mathrm{min}$, and we have 60 min available in total. We can fit in a maximum of 5 snapshots with each telescope. There are 4 telescopes, so we can take a total of 20 snapshots. Each of these needs to cover $5\,\mathrm{deg}^2$. This corresponds to a picture with a side of $\sqrt{5\,\mathrm{deg}^2} = 2.2\,\mathrm{deg} = 134\,\mathrm{arcmin}$. We divide this by the plate scale to find that the size of the CCD should be at least 39 mm. Dividing by the pixel size, this is 6505 pixels.

There are two more issues that we'll touch on briefly. First there's the bias, which, like the read-out noise, comes from the electronics. However, the bias is there on purpose. When we're doing read-out at the capacitor, the read-out noise might lead to a net negative voltage. The circuitry is designed to handle positive voltages only, and it can't convert such a signal to counts. To avoid this, we apply a constant voltage offset for each pixel, which of course will add a few extra counts. But because the bias is systematic rather than random, it's not too difficult to remove it entirely. One way to do this is to take a number of pictures (called bias frames) with an exposure of zero seconds. That way you don't even get a dark current, only read-out noise and the bias. After averaging[2] the bias frames, we're left with just the bias, which can be promptly subtracted from our raw CCD image. And yet, this whole procedure is redundant for most use cases. The reason is that dark frames also contain bias, so subtracting an averaged-out dark frame gets rid of the bias too.

Finally, there's the problem that a uniform signal on all pixels simply doesn't produce an uniform output, even if the error sources listed above didn't exist. There are multiple reasons for this. The quantum efficiency will vary pixel-to-pixel due to fabrication flaws; additionally, the optics may cause uneven illumination on the detector, which can create undesirable effects like vignetting. Luckily, this is all systematic and can be corrected for. The trick is to perform what is called a flat-field correction. This involves taking an image of a uniform source, for example the sky background at dusk or dawn. The output is termed a flat field. It will present as nonuniform owing to the fixed error pattern. Next, for each pixel we divide the counts in the raw image by the counts in the flat field. This essentially calibrates the raw image, because pixels that consistently overshoot are normalised by larger count values, and vice versa. We then take the resulting quotients and scale them back by a constant, usually the average count number within the flat field – after all, we want the end product to resemble the raw image. In more detail, the procedure for each pixel is as follows:

$$\text{corrected} = \frac{\text{raw} - \text{dark}}{\text{flat} - \text{dark}} \cdot \text{constant}.$$

The dark frame subtraction is there because it's better to get rid of everything except the fixed-pattern errors before applying the flat-field correction.

Even if you find the bookkeeping of errors tedious (many do!), remember that this is a science in and of itself, and it has enabled some extraordinary feats of precision measurement. Data analysis is perhaps the core skill of a professional astronomer.

Last updated: 8th August 2025

Send comments and corrections to sivanov.mail@proton.me

[2] The jargon for this is "stacking".

Problems

Problem 1. The eye as a camera (USAAAO 2025). Derive the limiting magnitude for naked-eye visibility of a star in a dark-sky environment, by considering an analogy to a CCD camera. Assume that the dominant noise source is $\sqrt{n}$ photon shot noise, and that "visibility" requires a signal-to-noise ratio (SNR) > 1. You'll need the following information:

- The dark-adapted human pupil dilates to an aperture of 6 mm. The quantum efficiency of the rod cells used for night vision is around 5% (averaged across the full visible spectrum), and the effective "exposure time" is roughly 30 ms.
- For a convenient magnitude reference, the intensity of sunlight within the visible spectrum is $5 \times 10^2\,\mathrm{W/m^2}$, and the apparent magnitude of the Sun is -26.7.

Approximate all visible photons as having wavelength $\lambda = 500\,\mathrm{nm}$ when calculating photon energy. This is near the peak sensitivity for human night vision.

Problem 2. Cluster photography (IOAA 2024). An astronomer takes pictures, in the V-band, of a faint celestial target, from a place with no light pollution. The selected target is the globular cluster Palomar 4, which has an angular diameter of $\theta = 72.0''$ and a uniform surface brightness in the V-band of $m = 20.6\,\mathrm{mag/arcsec^2}$. The observation equipment consists of one reflector telescope, with diameter $D = 305\,\mathrm{mm}$ and f-ratio $f/5$, as well as a prime focus CCD with quantum efficiency $\eta = 80\,\%$ and square pixels with size $l = 3.80\,\mathrm{\mu m}$.

(a) Calculate the plate scale (the angle of sky projected per unit length of the sensor) of the observation equipment in arcmin/mm.

(b) Estimate the number of pixels n_p covered by the cluster image on the CCD.

(c) Using an exposure time of $t = 15\,\mathrm{s}$, the astronomer obtains a signal-to-noise ratio SNR $= 225$. Compute the brightness of the sky at the observation site, knowing that the CCD has a read-out noise (standard deviation) $R = 5\,\mathrm{counts/px}$ and dark noise of $D = 6\,\mathrm{counts/px/min}$. Answer in $\mathrm{mag/arcsec^2}$. Note that the total read-out noise is $\sigma_{RON} = \sqrt{n_p} \cdot R$.

Work with the following data:

- V-band central wavelength: $\lambda_V = 550\,\mathrm{nm}$
- V-band bandwidth: $\Delta\lambda_V = 88.0\,\mathrm{nm}$
- Photon flux for a zero-magnitude object in V-band: $10\,000\,\mathrm{photons/nm/cm^2/s}$

Solutions

1. You can find the solution here.

2. You can find the solution here.

References

[1] K.F. Riley, M.P. Hobson, and S.J. Bence, *Mathematical Methods for Physics and Engineering: A Comprehensive Guide*, 3rd edition, 2006

[2] Michael Newberry, CCD Camera Gain Calibration, 1996, https://www.mirametrics.com/tech_note_ccdgain.php