

# Characterizing and Calibrating CCDs. III. Multiplicative Terms

A large, glowing nebula with a mix of teal, cyan, and magenta colors is centered in the image. The background is a dense field of stars of various colors (white, yellow, orange, blue) against a black space.

Ast 401/580  
Fall 2019

# Multiplicative Factors

Not all the pixels in a CCD are equally sensitive to light. Another way to put this is that all of the buckets aren't quite identical: some of them have slightly bigger tops than others. Differences are on the order of a percent or two to ten percent or more. Wow!

But, we want to be able to do photometry down to a percent or a fraction thereof. What to do?


# Flat-fielding

Optimally you expose the CCD to a “uniform” light source that will illuminate the chip the exact same way that the sky does at night. Well, good luck with that. Your typical choices are the following:

- 1) Dome flats (white spots)
- 2) Bright twilight skies

In general, you probably want both, and see how well they agree and which one works better (more on this in a moment...)





Dome Flat

# Dome flats

Why isn't the exposure completely "flat"?


★ Small things:

- pixel-to-pixel sensitivity variations


★ Big things:

- dust specks, vignetting


# Pixel to pixel variations



NOAO/IRAF V2.16 massey@stillwater.local Sat 10:10:35 09-Nov-2013  
Line 1572 of ccd027.fits  
dome flat wn



NOAO/IRAF V2.16 massey@stillwater.local Sat 10:11:20 09-Nov-2013  
Line 1572 of ccd027.fits  
dome flat wn





# Pixel to pixel

Let's do the math. RMS we measure is 60.8 ADU. Average is amount of light is 7450 ADUs. Gain is nominally 3e/ADUs. So we expect:

$$7450 \text{ ADUs (above bias)} \times 3 \text{ e/ADUs} = 22,350e$$

$$\text{RMS expected (e)} = \text{sqrt} (22350) = 149e$$

$$\text{RMS (ADU)} = 149e / 3e/\text{ADU} = 49.7 \text{ ADUs.}$$

Since  $60.8 > 49.7$  sensitivity variations are quite significant.

# Pixel to pixel

Since  $60.8 > 49.7$  sensitivity variations are quite significant.

How much? Remember, uncertainties add in quadrature, so estimating that the contribution due to sensitivity variations is:

$(\text{sigma\_sens})^2 + 49.7^2 = 60.8^2$ , which implies

$\text{sigma\_sens} = 35 \text{ ADU}$ . In percentage, that's then  $35 \text{ ADU} / 7450 \text{ ADU} = 0.005$  or 0.5%. Well, okay, that's not too bad. One could almost not flat-field the data if one wanted 1% photometry. Except...

# Pixel to pixel variations


Two reasons that at the pixel to pixel level things aren't completely uniform:

- ★ Photon statistics

- ★ Slight differences in the sensitivity of each pixel. Remember: the only **flat** CCD is a **dead** CCD.

Need to make sure that one has the equivalent of 10,000e for 1% photometry. 30,000–50,000e is much better. Take 5 or so at 10,000e say.


# Big things




Dust spots

Dome Flat






Dome Flat



Dome Flat


Dust spots



Dust spots


Dome Flat








NOAO/IRAF V2.16 massey@stillwater.local Sat 08:01:28 09-Nov-2013  
Line 2970 of ccd027.fits  
dome flat wn



NOAO/IRAF V2.16 massey@stillwater.local Sat 08:02:25 09-Nov-2013  
Line 2970 of ccd027.fits  
dome flat wn



# Dust spots

If there is a speck of dust on the filter, it will create a very large donut on the image.

If there's a speck of dust on the dewar window, it will create a smaller donut on the image.

# Dust spots

How can you tell?

You could CALCULATE the expected size IF you knew the distance from the CCD to the dewar window and the filter.


But the easiest way is to see if the donuts are in the same place when you change filters!

# Dust spots

Dust donuts will flat-field out perfectly as long as the dust doesn't move between your flat-field and your program frame:



- If on the dewar window, this requires that gravity or a breeze doesn't disturb them.
- If on the filter, this then requires that the filters reposition "exactly" as well as the dust doesn't move with respect to the filter.






**Cosmic rays?**  
**Bad pixels?**

**Dome Flat**



Line 2101 of ccd027.fits  
dome flat wn



NOAO/IRAF V2.16 Massey@stillwater.local Sat 08:17:30 09-Nov-2013  
Line 2101 of ccd027.fits  
dome flat wn



# Bad pixels

So, these features are clearly bad pixels rather than cosmic-rays: cosmic-rays would be positive, not negative, features.

Bad pixels might flat-field out. It all depends. If they don't, you can mask them out in your analysis. Even if they do flat-field out, the photometric error you have for any star will be vastly underestimated.



# Vignetting




Dome Flat

Vignetted region?  
Or okay?



NOAO/IRAF V2.16 massey@stillwater.local Sat 09:12:17 09-Nov-2013  
Average of columns 1264 to 1273 of ccd027.fits  
dome flat wn



# Vignetting

The region at the lower right is low on the dome flat. There are two possible reasons:

(1) The telescope/camera/CCD vignettes a bit down there, or the CCD is just not as sensitive in the lower right corner. Could be! If so, this is good news.

(2) The illumination provided by the dome flat is not quite right. If so, this is bad news. But not awful.

# Vignetting

Typically dome flats match the illumination function of the sky to 1–2%, and sometimes they do much better. To be off by 10% is possible but rare.

# Dome flats

SUMMARY: We are going to use our dome flat to remove multiplicative factors by dividing our program data by them, after additive effects have been removed (both in the program data AND in the dome flats). We are hoping on the dome flats to remove:

- ★ Pixel-to-pixel sensitivity variations (few percent)
- ★ Dust donuts (few percent)
- ★ Really low pixels (good luck with that!)
- ★ Illumination function (vignetting).

# Dome flats

We'll average a 5 dome flats (say) rather than just 1 for two reasons. What are they?

- a) We can get lots of counts (50,000e ).
- b) We can filter out cosmic rays (radiation events).
- c) We can beat down read-noise.

# Dome flats

We'll average a 5 dome flats (say) rather than just 1 for two reasons. What are they?

- a) We can get lots of counts (50,000e ).
- b) We can filter out cosmic rays (radiation events).
- c) We can beat down read-noise.



# Dome flats

Before we divide by our flat fields, we will need to normalize them by the average value in the flats. This is crucial if we are going to be able to estimate our uncertainties in the final reduced data.

# Illumination Function

We're still stuck on the vignetting issue, though. Is it real or not?

ALWAYS a good idea to check against the sky. Two options:

(1) Bright twilight flats.


- Pros: If they are bright enough you can use them as your primary flat field
- Cons: If the tube is open, scattered light can play havoc.

(2) Dark sky flats.

- Pros: Illuminates the sky exactly the same way as your object exposures
- Cons: Take away from valuable observing time.


# Sky flats

Can either be bright twilight or “blank” sky (which is never actually blank).



Sky Flat






Dome Flat



# Sky flat

Looks very similar to dome flat, except for some subtle variations in the overall shape of the illumination (“illumination function”). What happens if we divide the sky flat by the dome flat?



# Sky / Dome


Lots of good news here:

- ★ The dust donuts are gone!
- ★ The low pixels divided out!!

Some bad news though:


- ★ There's one heck of a gradient in the illumination function.






Line 1572 of test


WN Sky Flat





Column 1010 of test

WN Sky Flat







# Dome Flats vs Sky Flats

5% difference! But, we're after <1%. What do we do?

First, we must decide which is right—sky flats or dome flats!


# Dome Flats vs Sky Flats

Two ways to test:


(1) On a good, clear (photometric) night put a star in the center of the frame and take an exposure. Now move the star to where the two disagree and take an exposure. Which flat gives you consistent answers between the two places?

(2) If you have enough counts in the sky of your program frames you can try both kinds of flats and see which gives you the flattest sky. In other words, which flat actually works?










# Subtle!

Five percent is hard to discern by eye!

But you can always make plots...


NOAO/IRAF V2.16 massey@stillwater.local Sun 08:26:26 10-Nov-2013  
Average of columns 1496 to 1505 of dometest  
lmc-field144 WN



NOAO/IRAF V2.16 massey@stillwater.local Sun 08:22:05 10-Nov-2013

Average of columns 1493 to 1507 of skytest

lmc-field144 WN









# Sky flat correction

So we really do want to correct the data for the gradient in this case. The skies are better! We can either correct the dome flats by the sky flats, or we can simply get good enough counts in the sky flats to use them both for pixel-to-pixel and for illumination function.

# Side-bar: averaging frames

When we average a bunch of nearly identical frames there are a couple of ways we can do this:

**Straight average.** In the new frame, every pixel is simply the average of the values for the same pixel in all of the other frames.

**Filtered average.** First, you throw out anything anything that is several sigma high or low.

# Side-bar: averaging frames

Consider the pixel (100,100). Imagine you have 5 frames and the values are: 2000, 2100, 1900, 5000, and 2000. If you do a straight average, you would get 2600. But that has been strongly affected by the one outlier value (5000). What you could do instead is construct the median (the middle value) by sorting these numbers: 1900, 2000, 2000, 2100, 5000. The median is then 2000. You could then find the average deviation of each value from this (-100, 0, 0, +100, +3000) and quickly determine which value(s) weren't like the others. Or you could calculate the expected sigma based on read-noise and gain. (If the gain was 5 and the read-noise small, we would expect the standard deviation to be 100, wouldn't we?). So, 3000 would deviate by 30 sigma. That's rather unlikely, statistically speaking.



# Side-bar: averaging frames


For bias frames, I like to average them by throwing out the highest and lowest value at each pixel.

For flat-fields, I like to first scale each frame by the mode, and then throw out any values that are 3 sigma above the expected noise. That gets rid of cosmic rays (dome flats) and stars (in sky flats).

# Why we need to filter

Sky frames (a particularly egregious example)







5.2e+02 5.4e+02 5.7e+02 5.9e+02 5.2e+02 5.4e+02 5.7e+02 5.9e+02 7.2e+02 7.4e+02



2.6e+02 2.7e+02 2.8e+02 2.9e+02 3e+02 3.2e+02 3.3e+02 3.4e+02 3.5e+02 3.7e+02







40

44

48

52

56


60

64

71

75

80



# Combining frames

If you're trying to filter something like dome flats or sky flats, you really must scale each exposure by the average (median or modal, not mean) value in order first.

# Flat fielding

Flat fielding **MUST** be done filter-by-filter. Why?

a) Dust specks might be on the filter

b) Wavelength dependency in the sensitivity of pixels.