Characterizing and Calibrating CCDs. II. Getting Rid of the Additive Terms

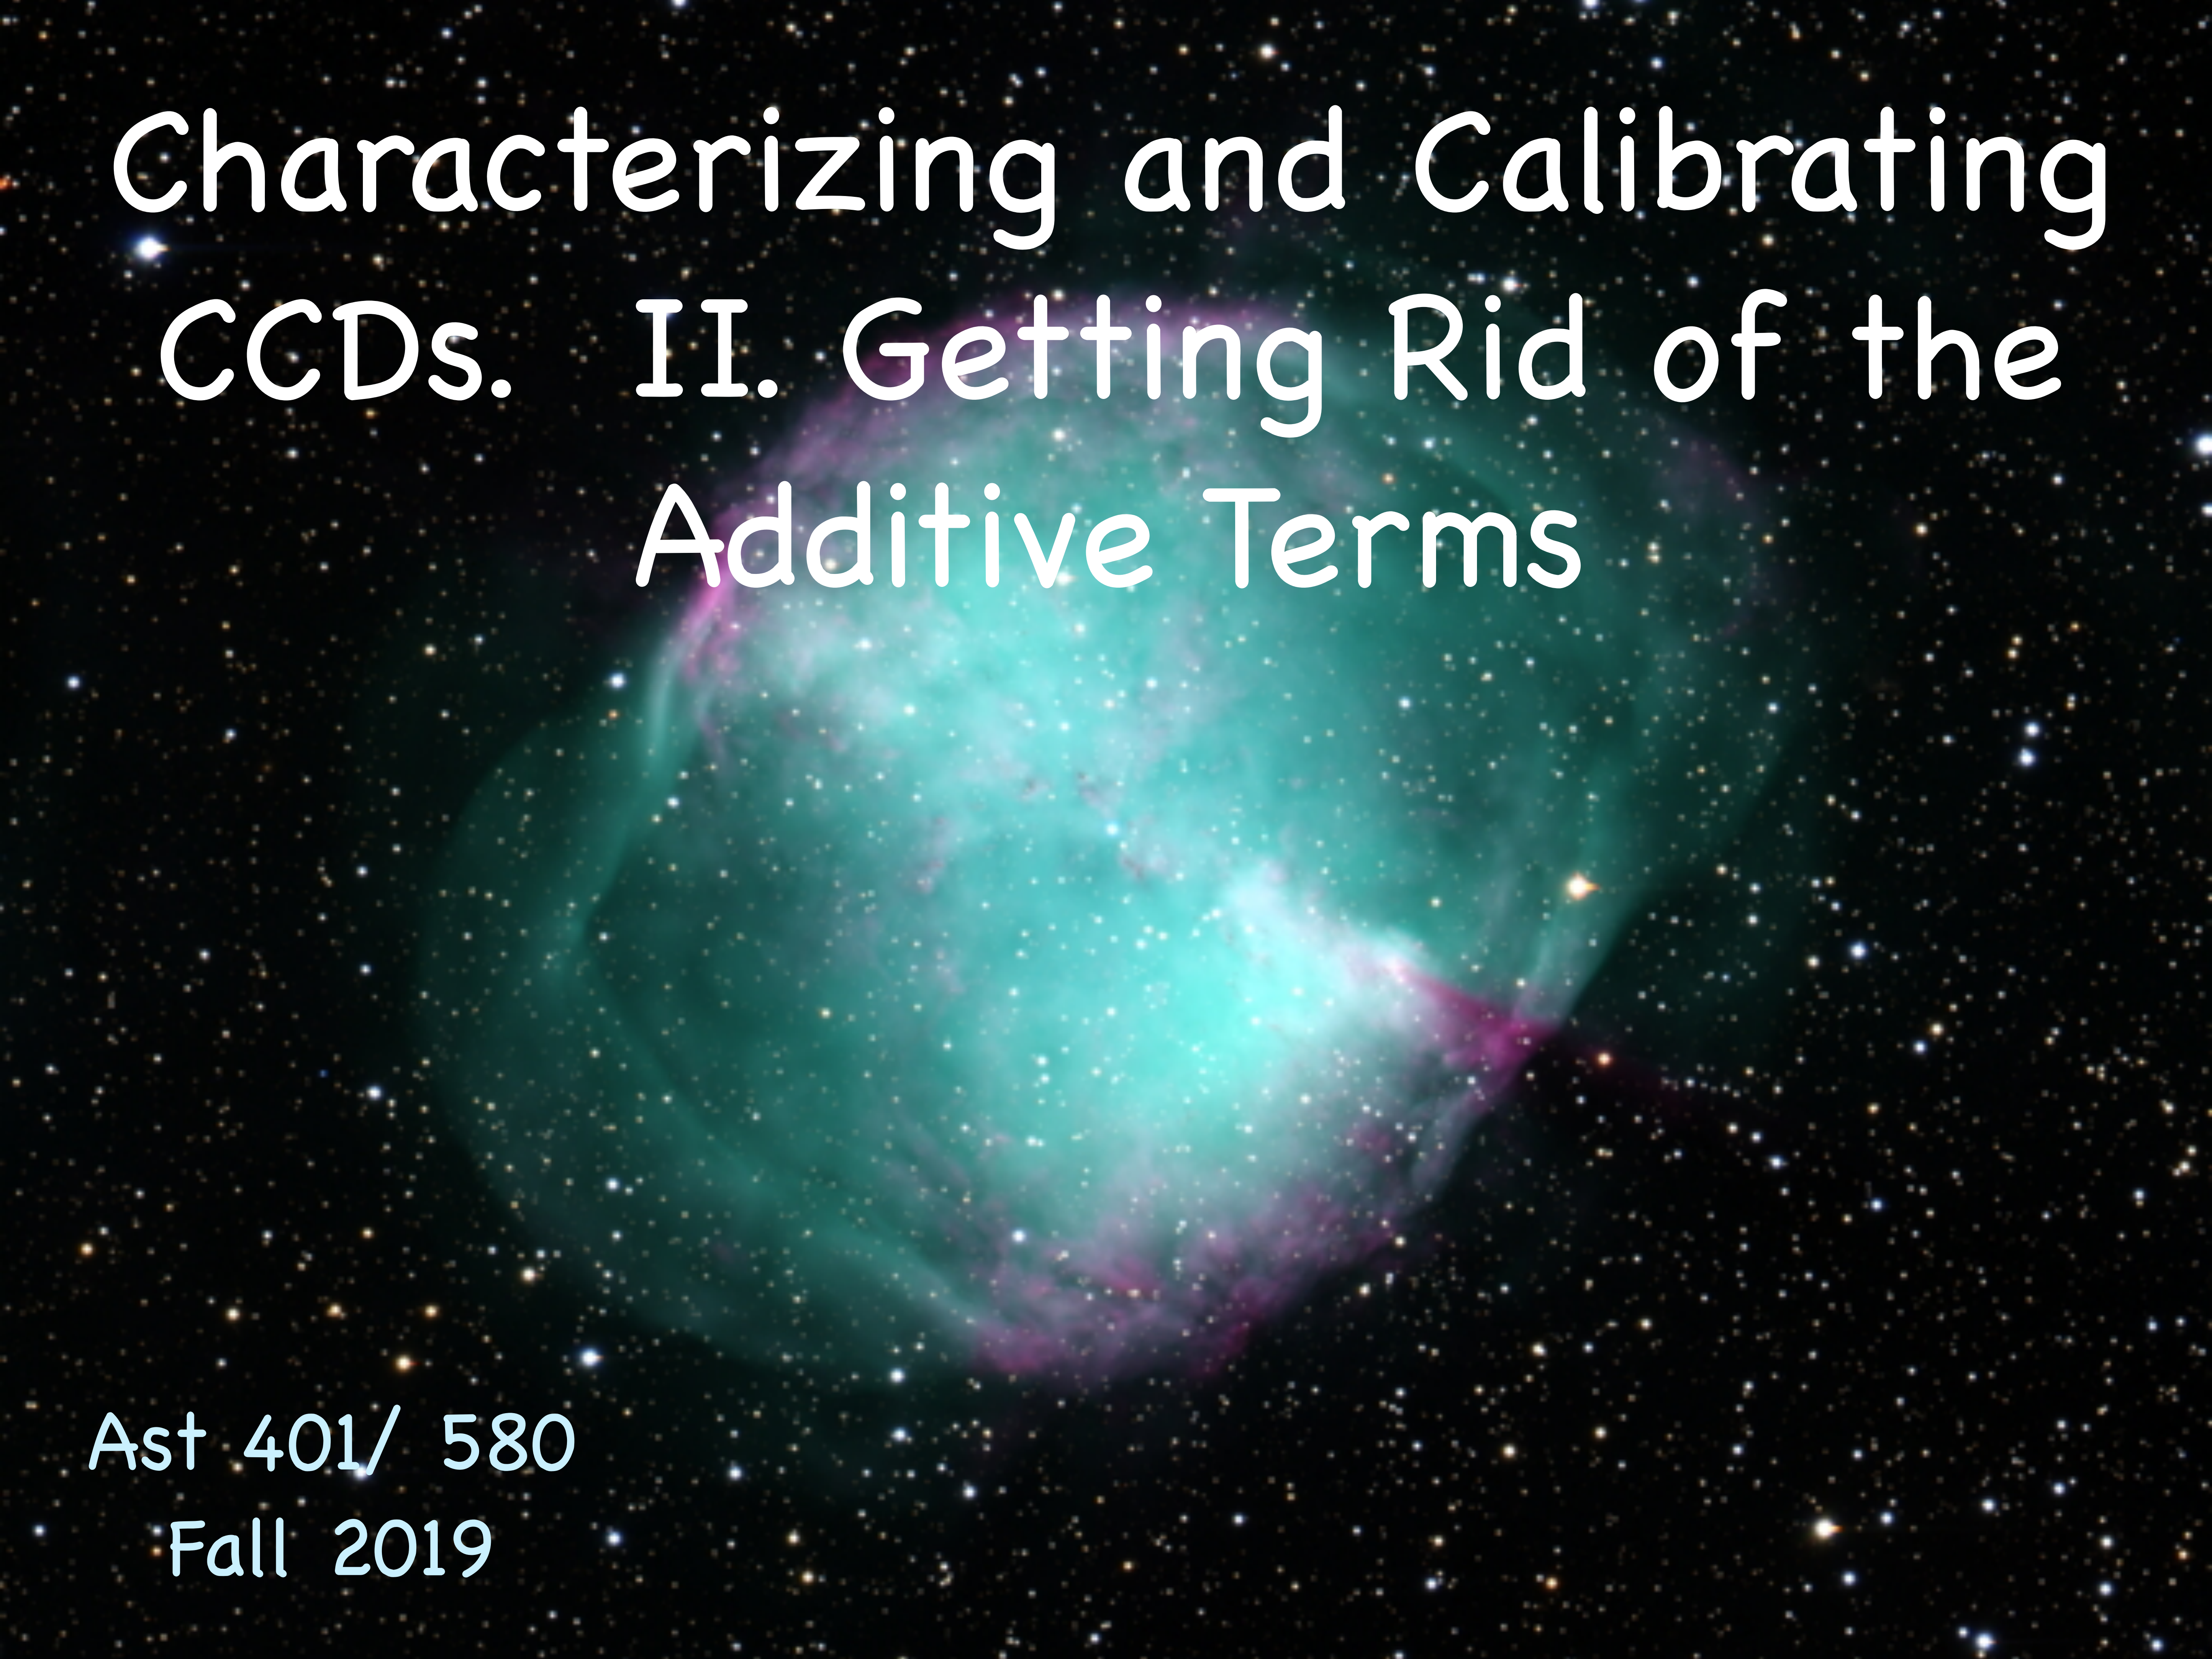

Ast 401/ 580

Fall 2019

Anatomy of a Raw CCD Image

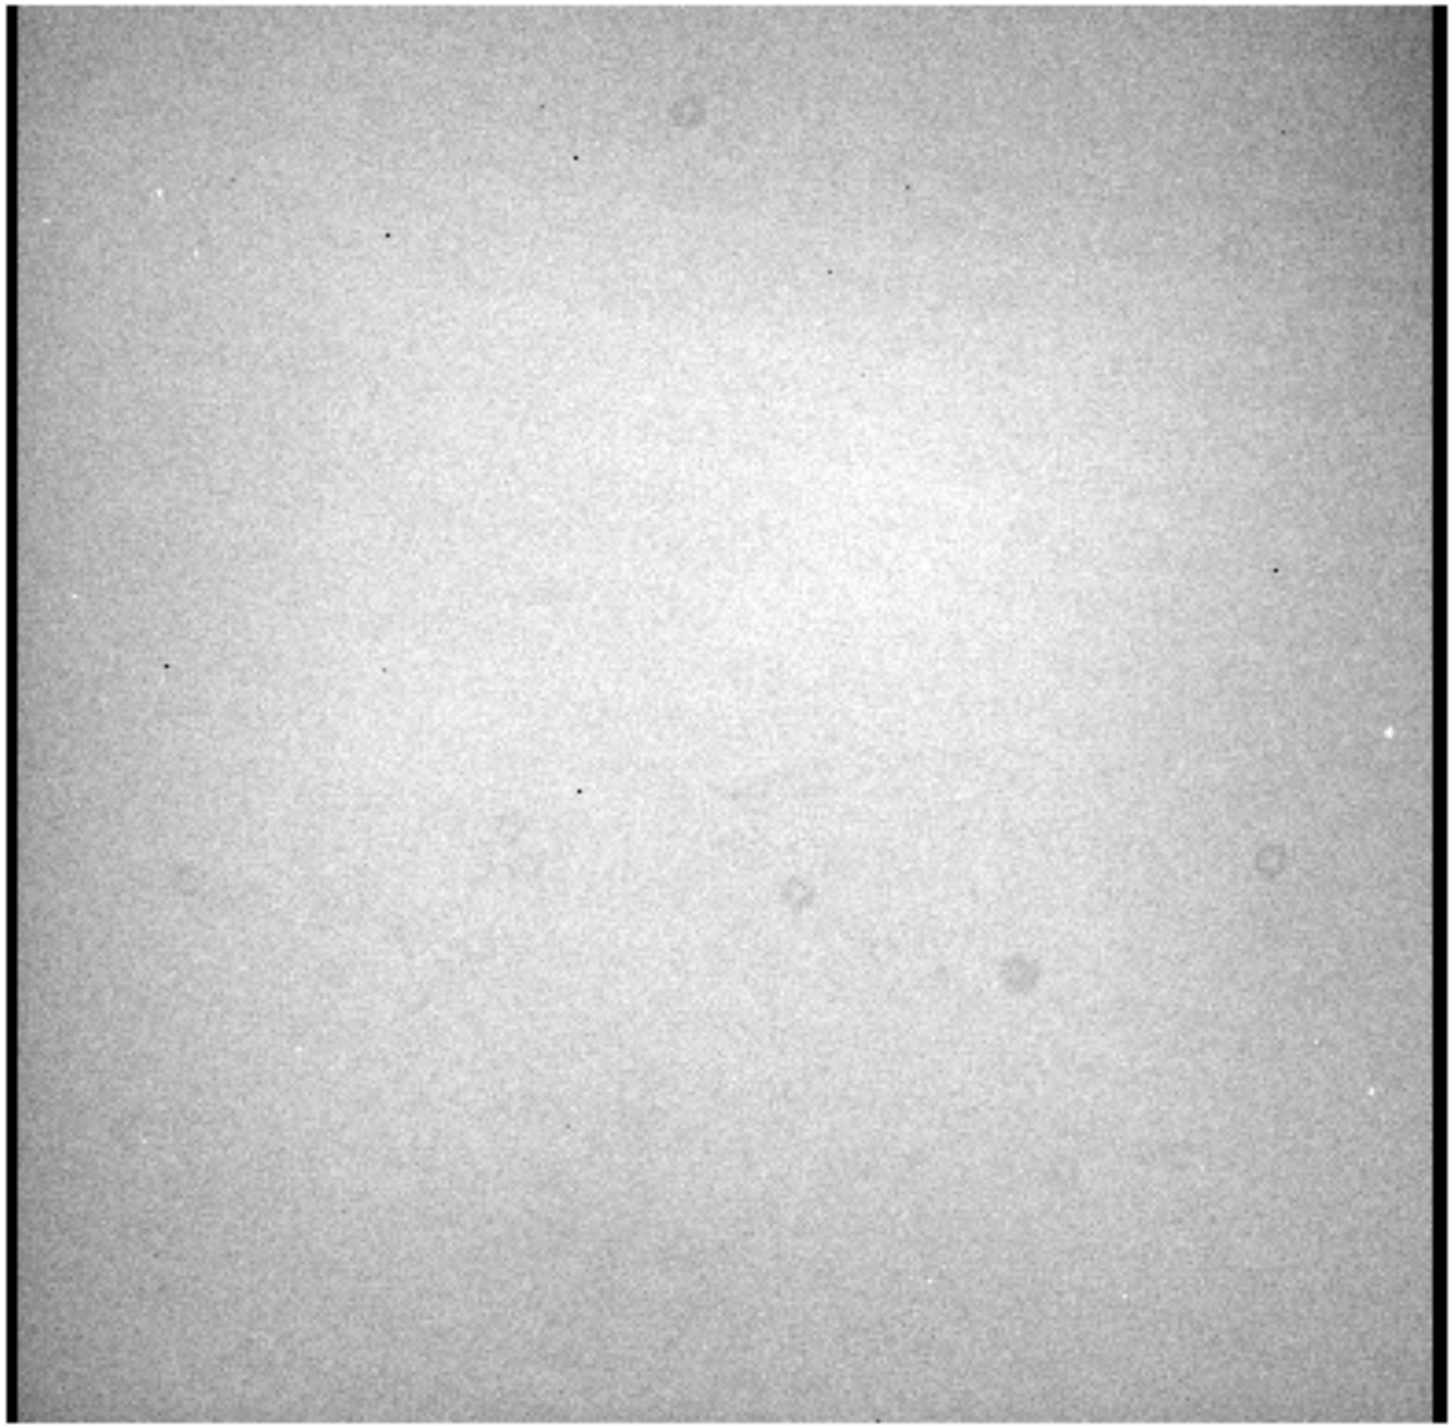

Anatomy of a Raw CCD Image

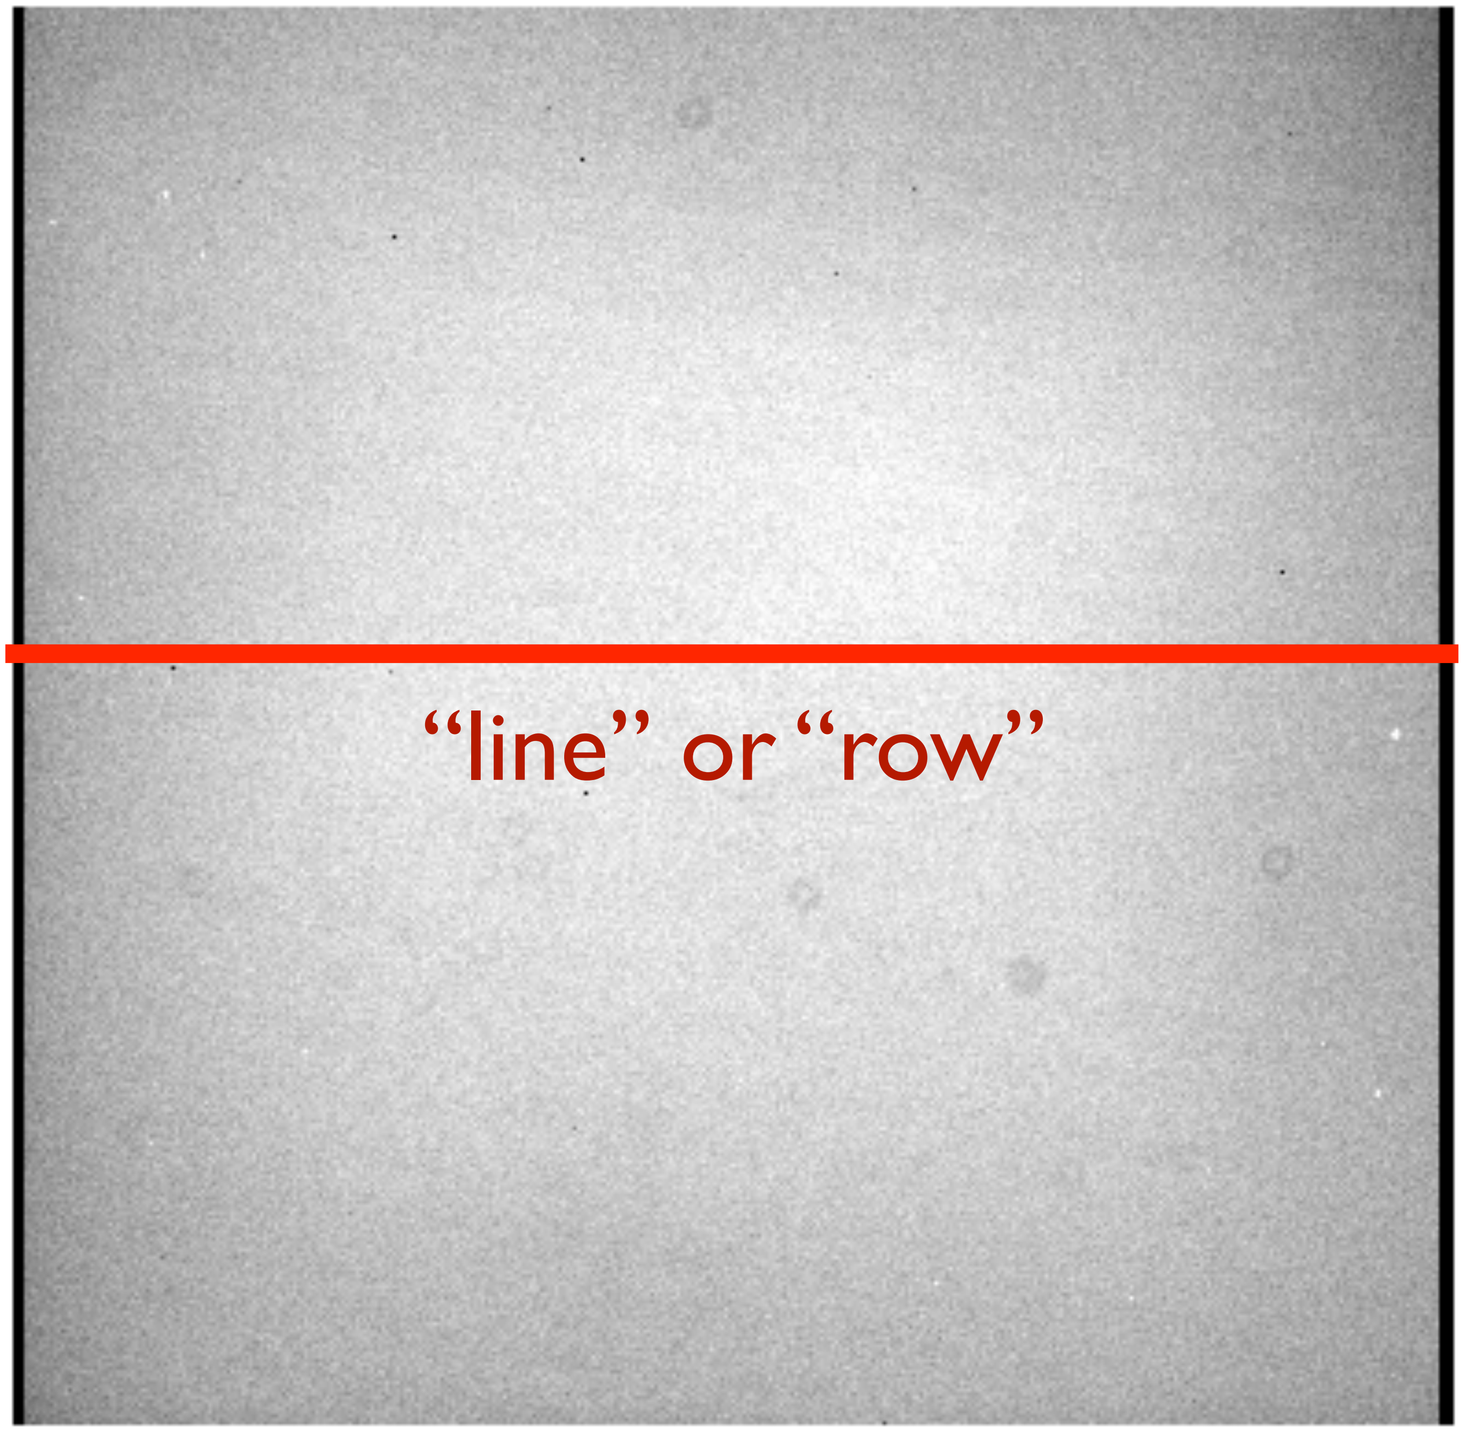

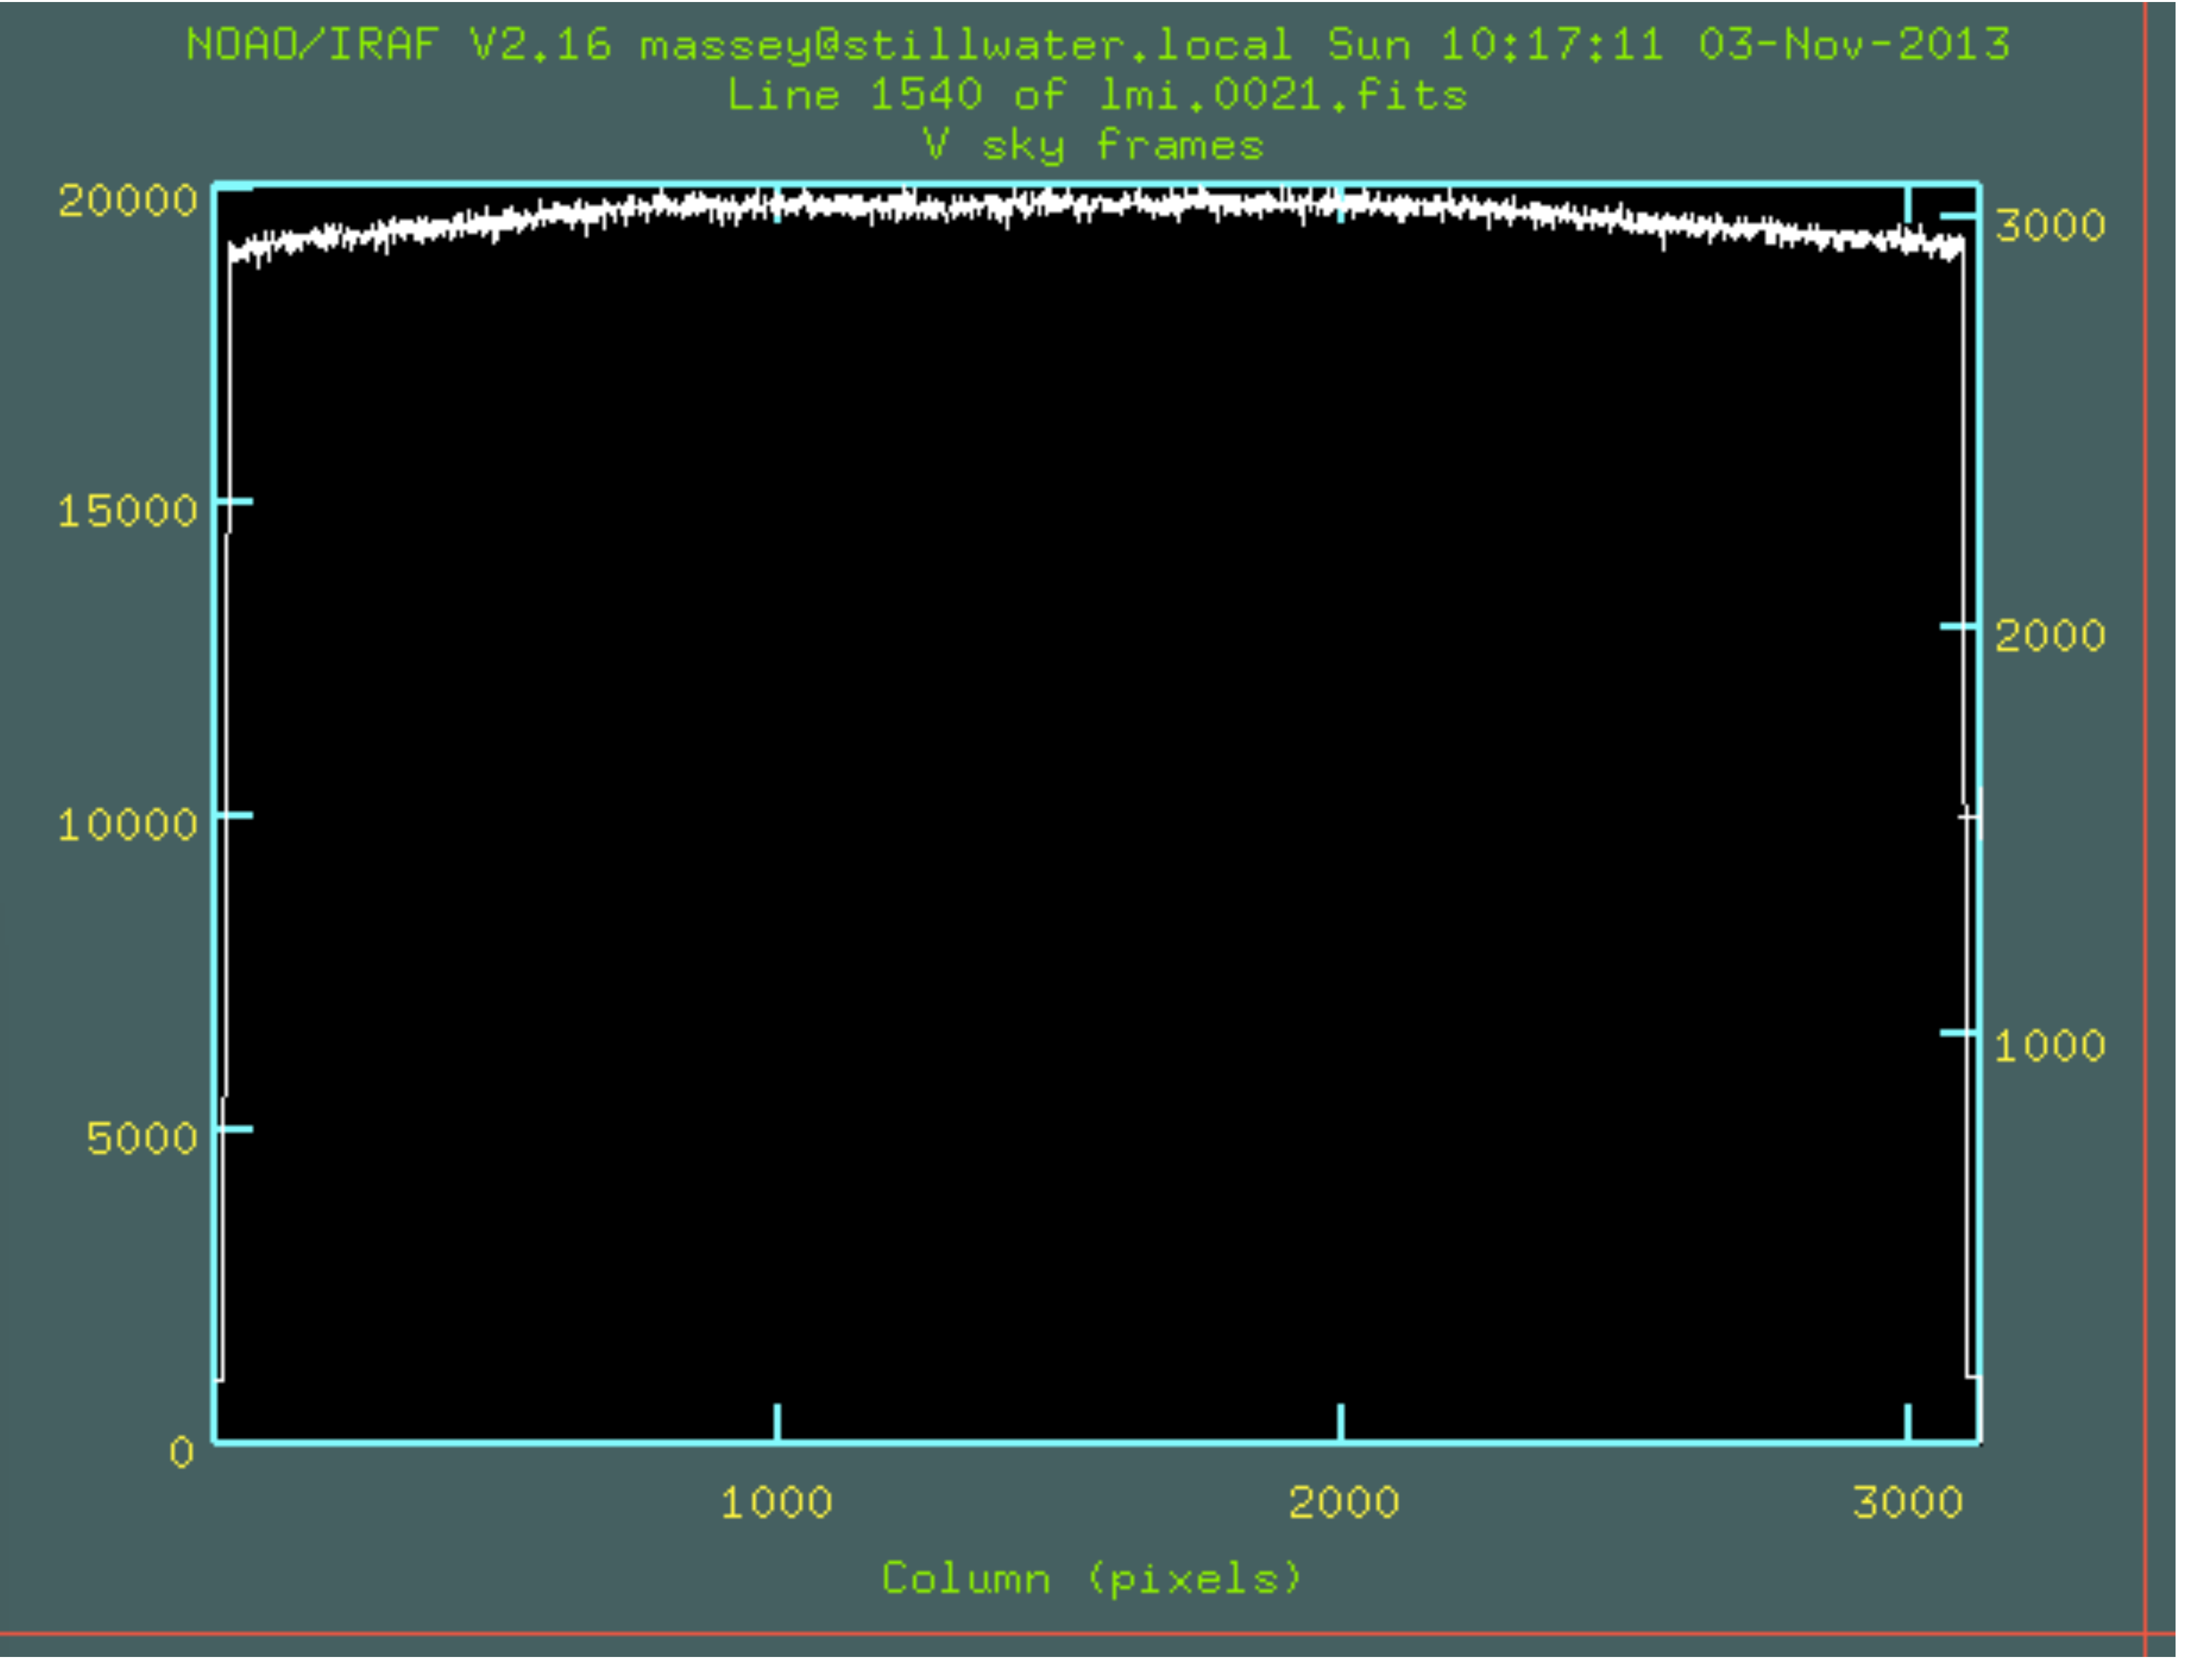
NOAO/IRAF V2.16 massey@stillwater.local Sun 10:17:11 03-Nov-2013
Line 1540 of lmi.0021.fits
V sky frames
20000
15000
10000
5000
0
3000
2000
1000
1000
2000
3000
Column (pixels)

Anatomy of a Raw CCD Image

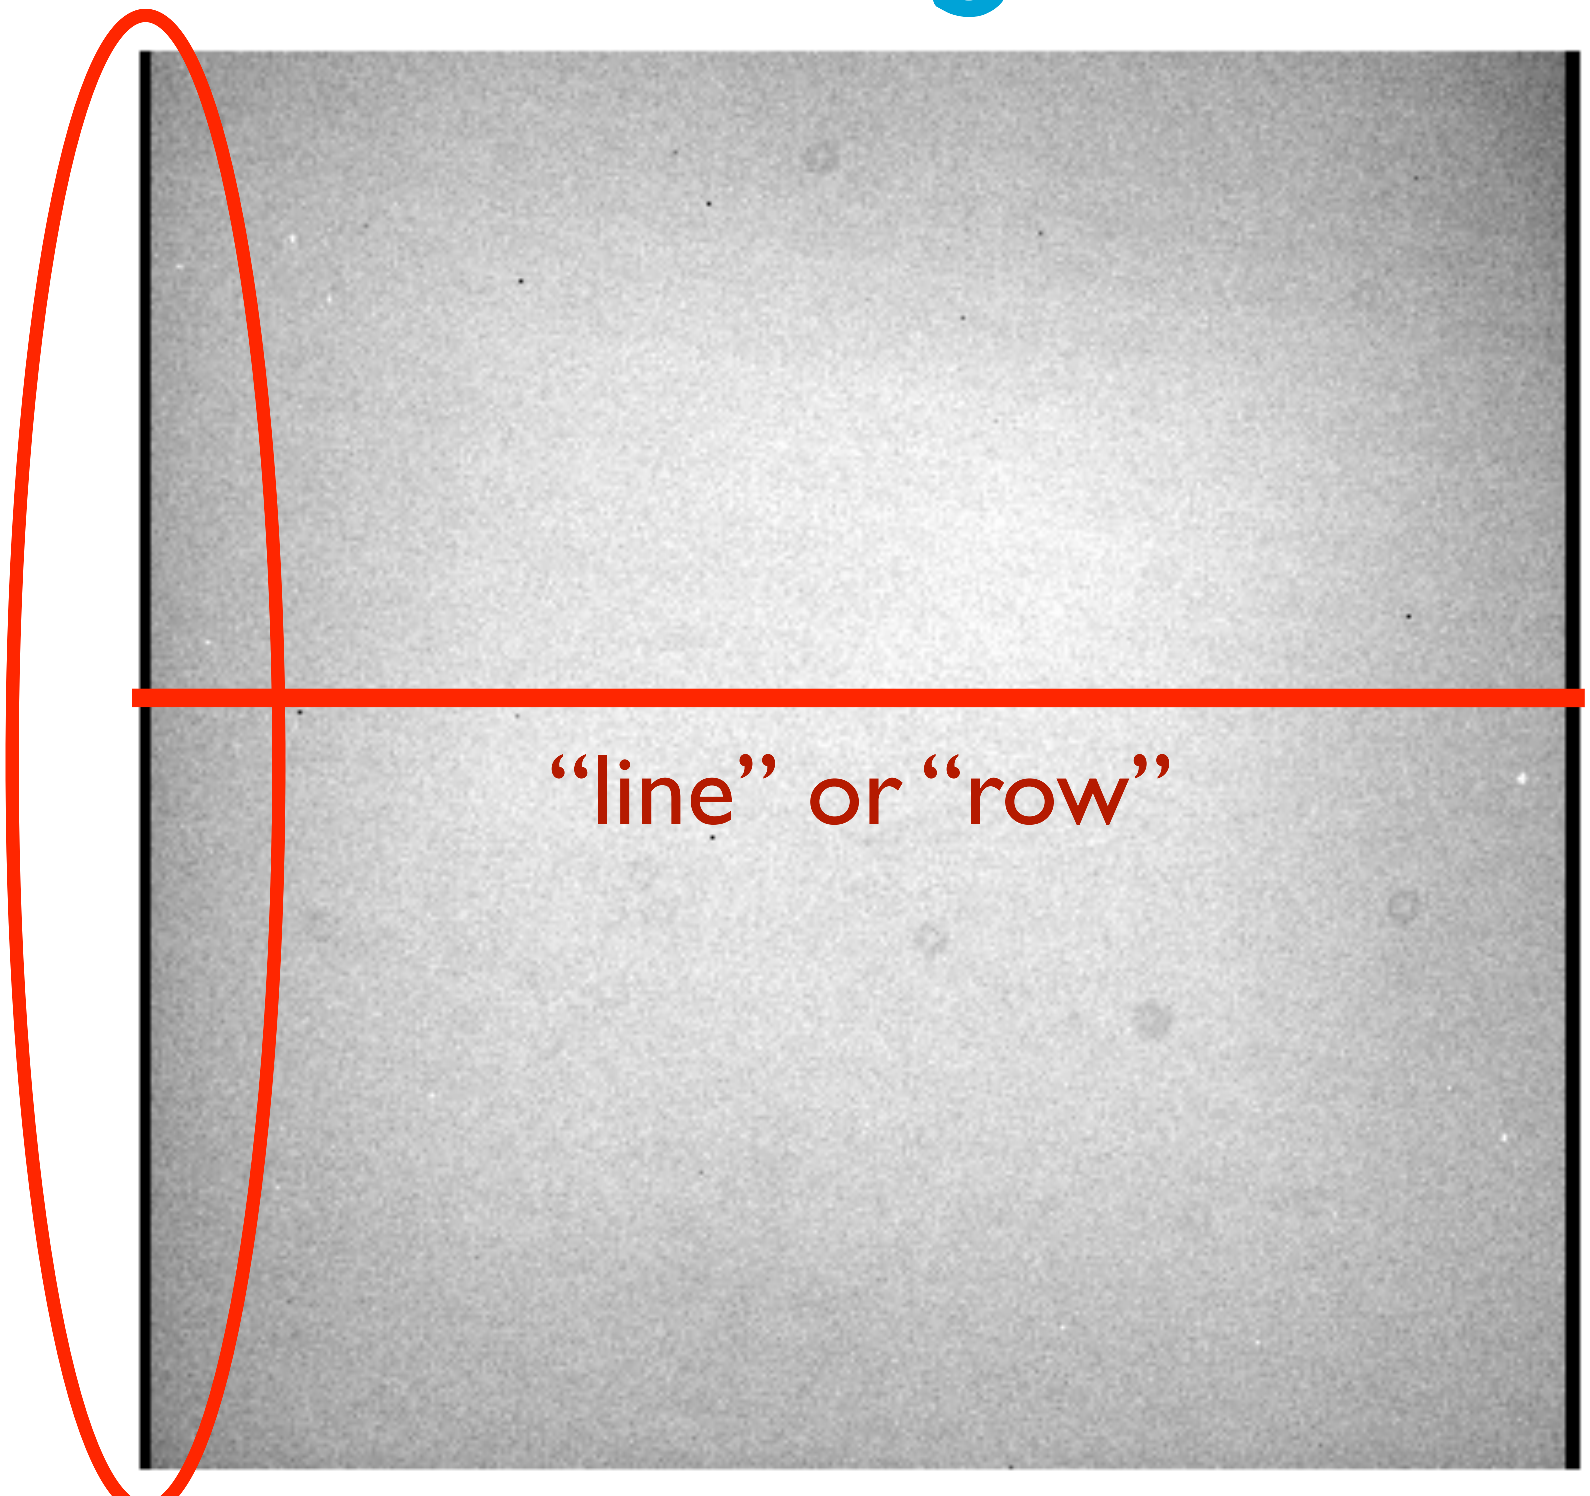

What the heck is that?

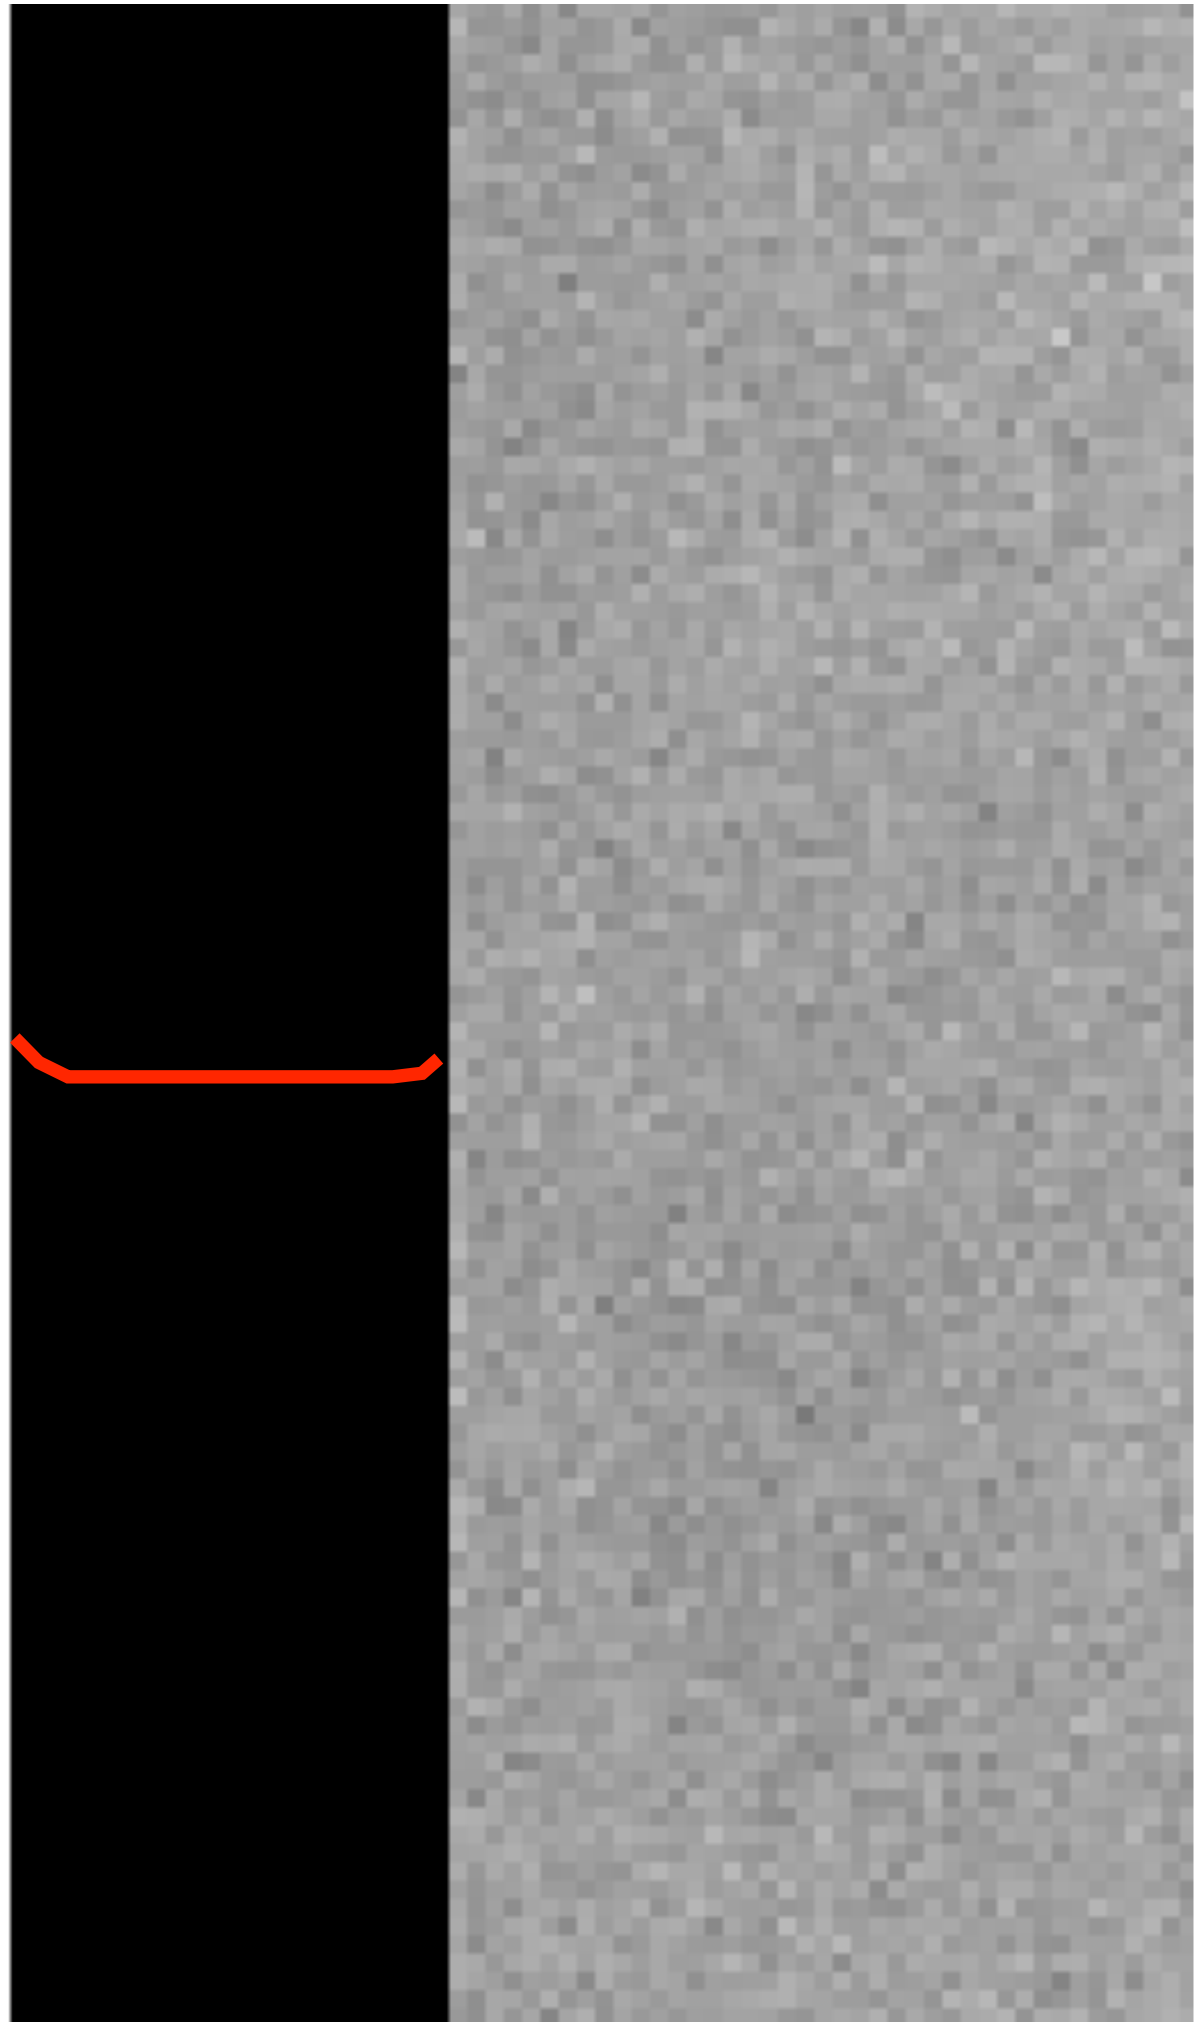

What the heck is that?

PRE-SCAN

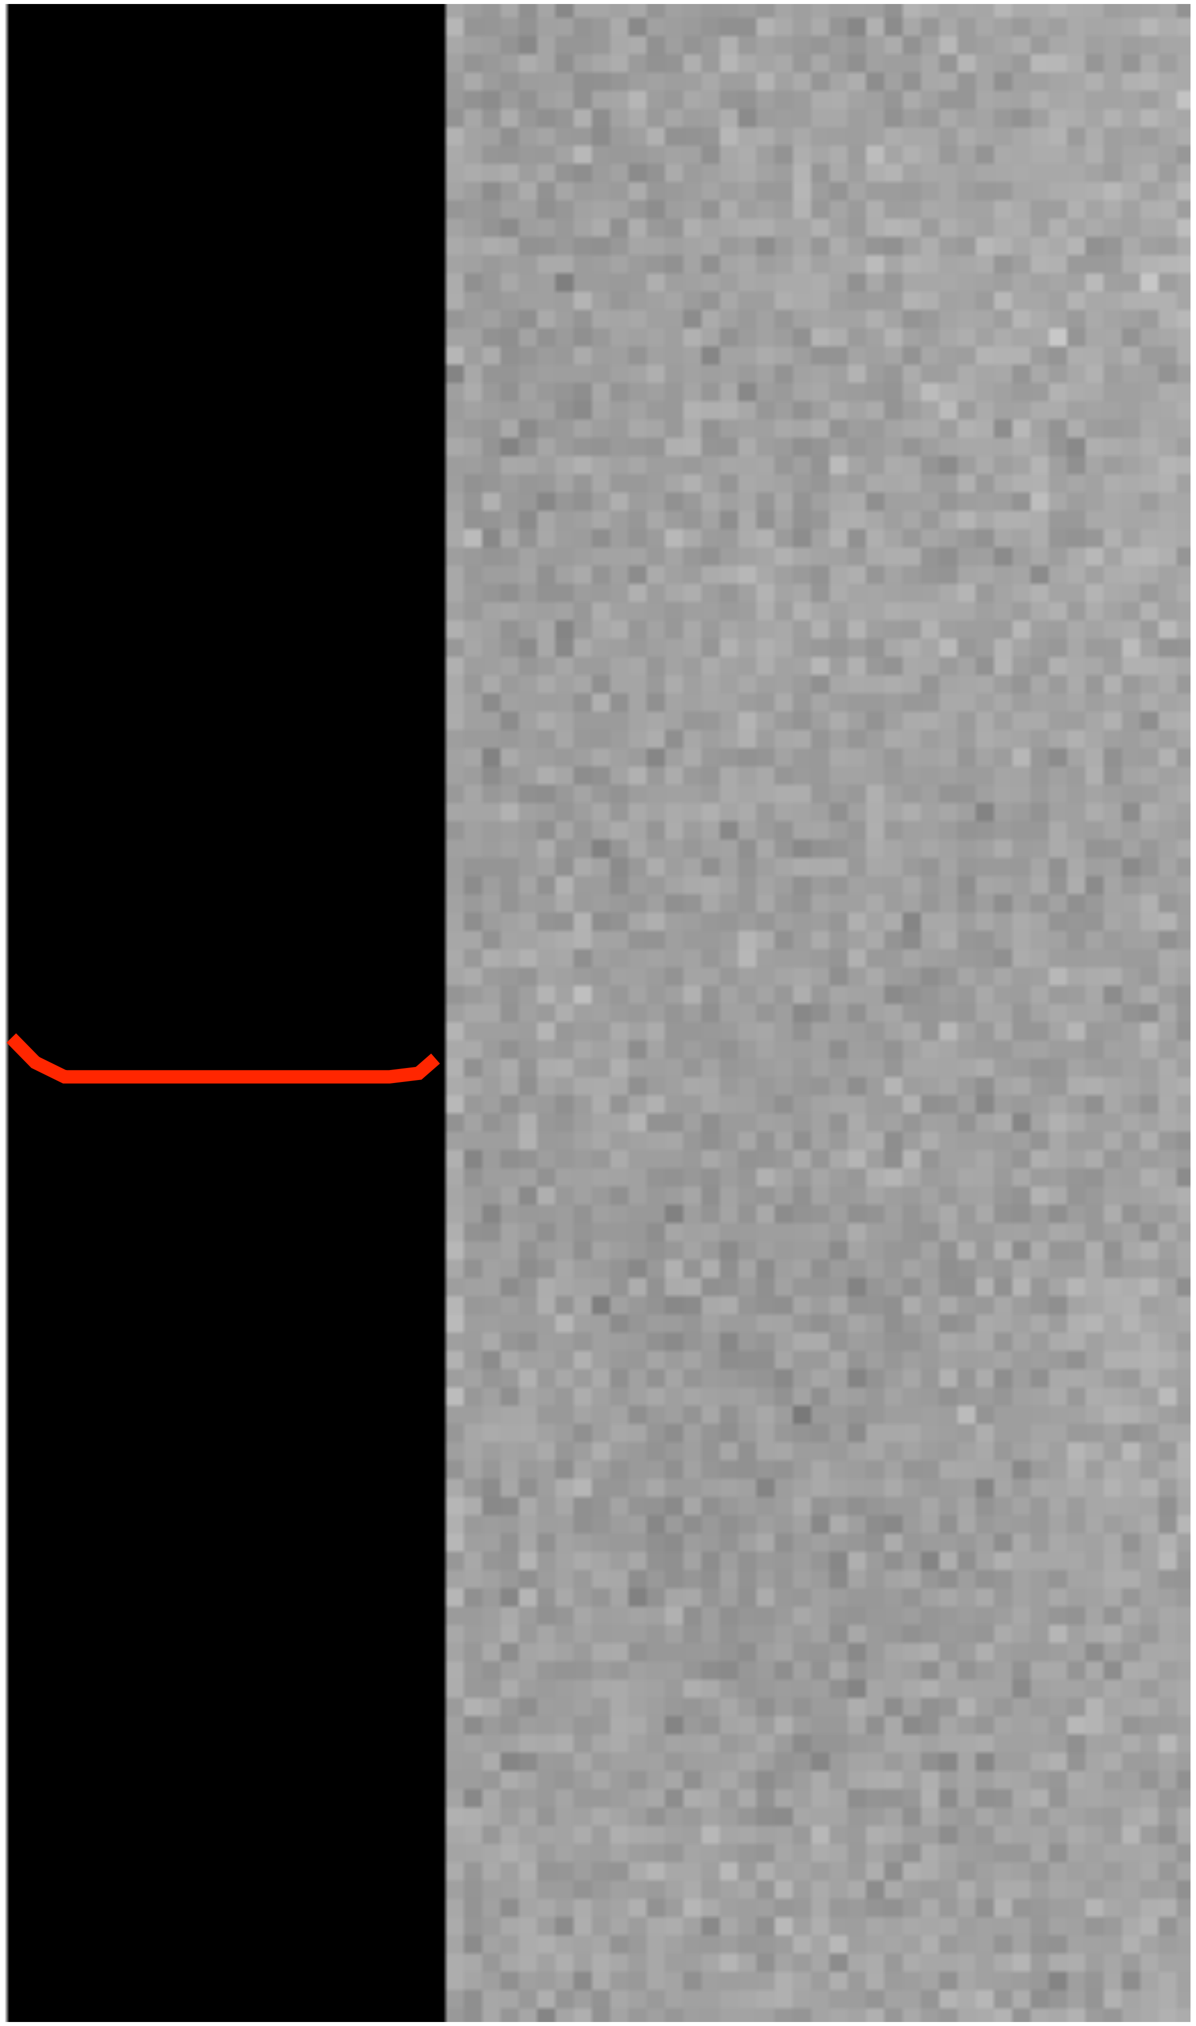

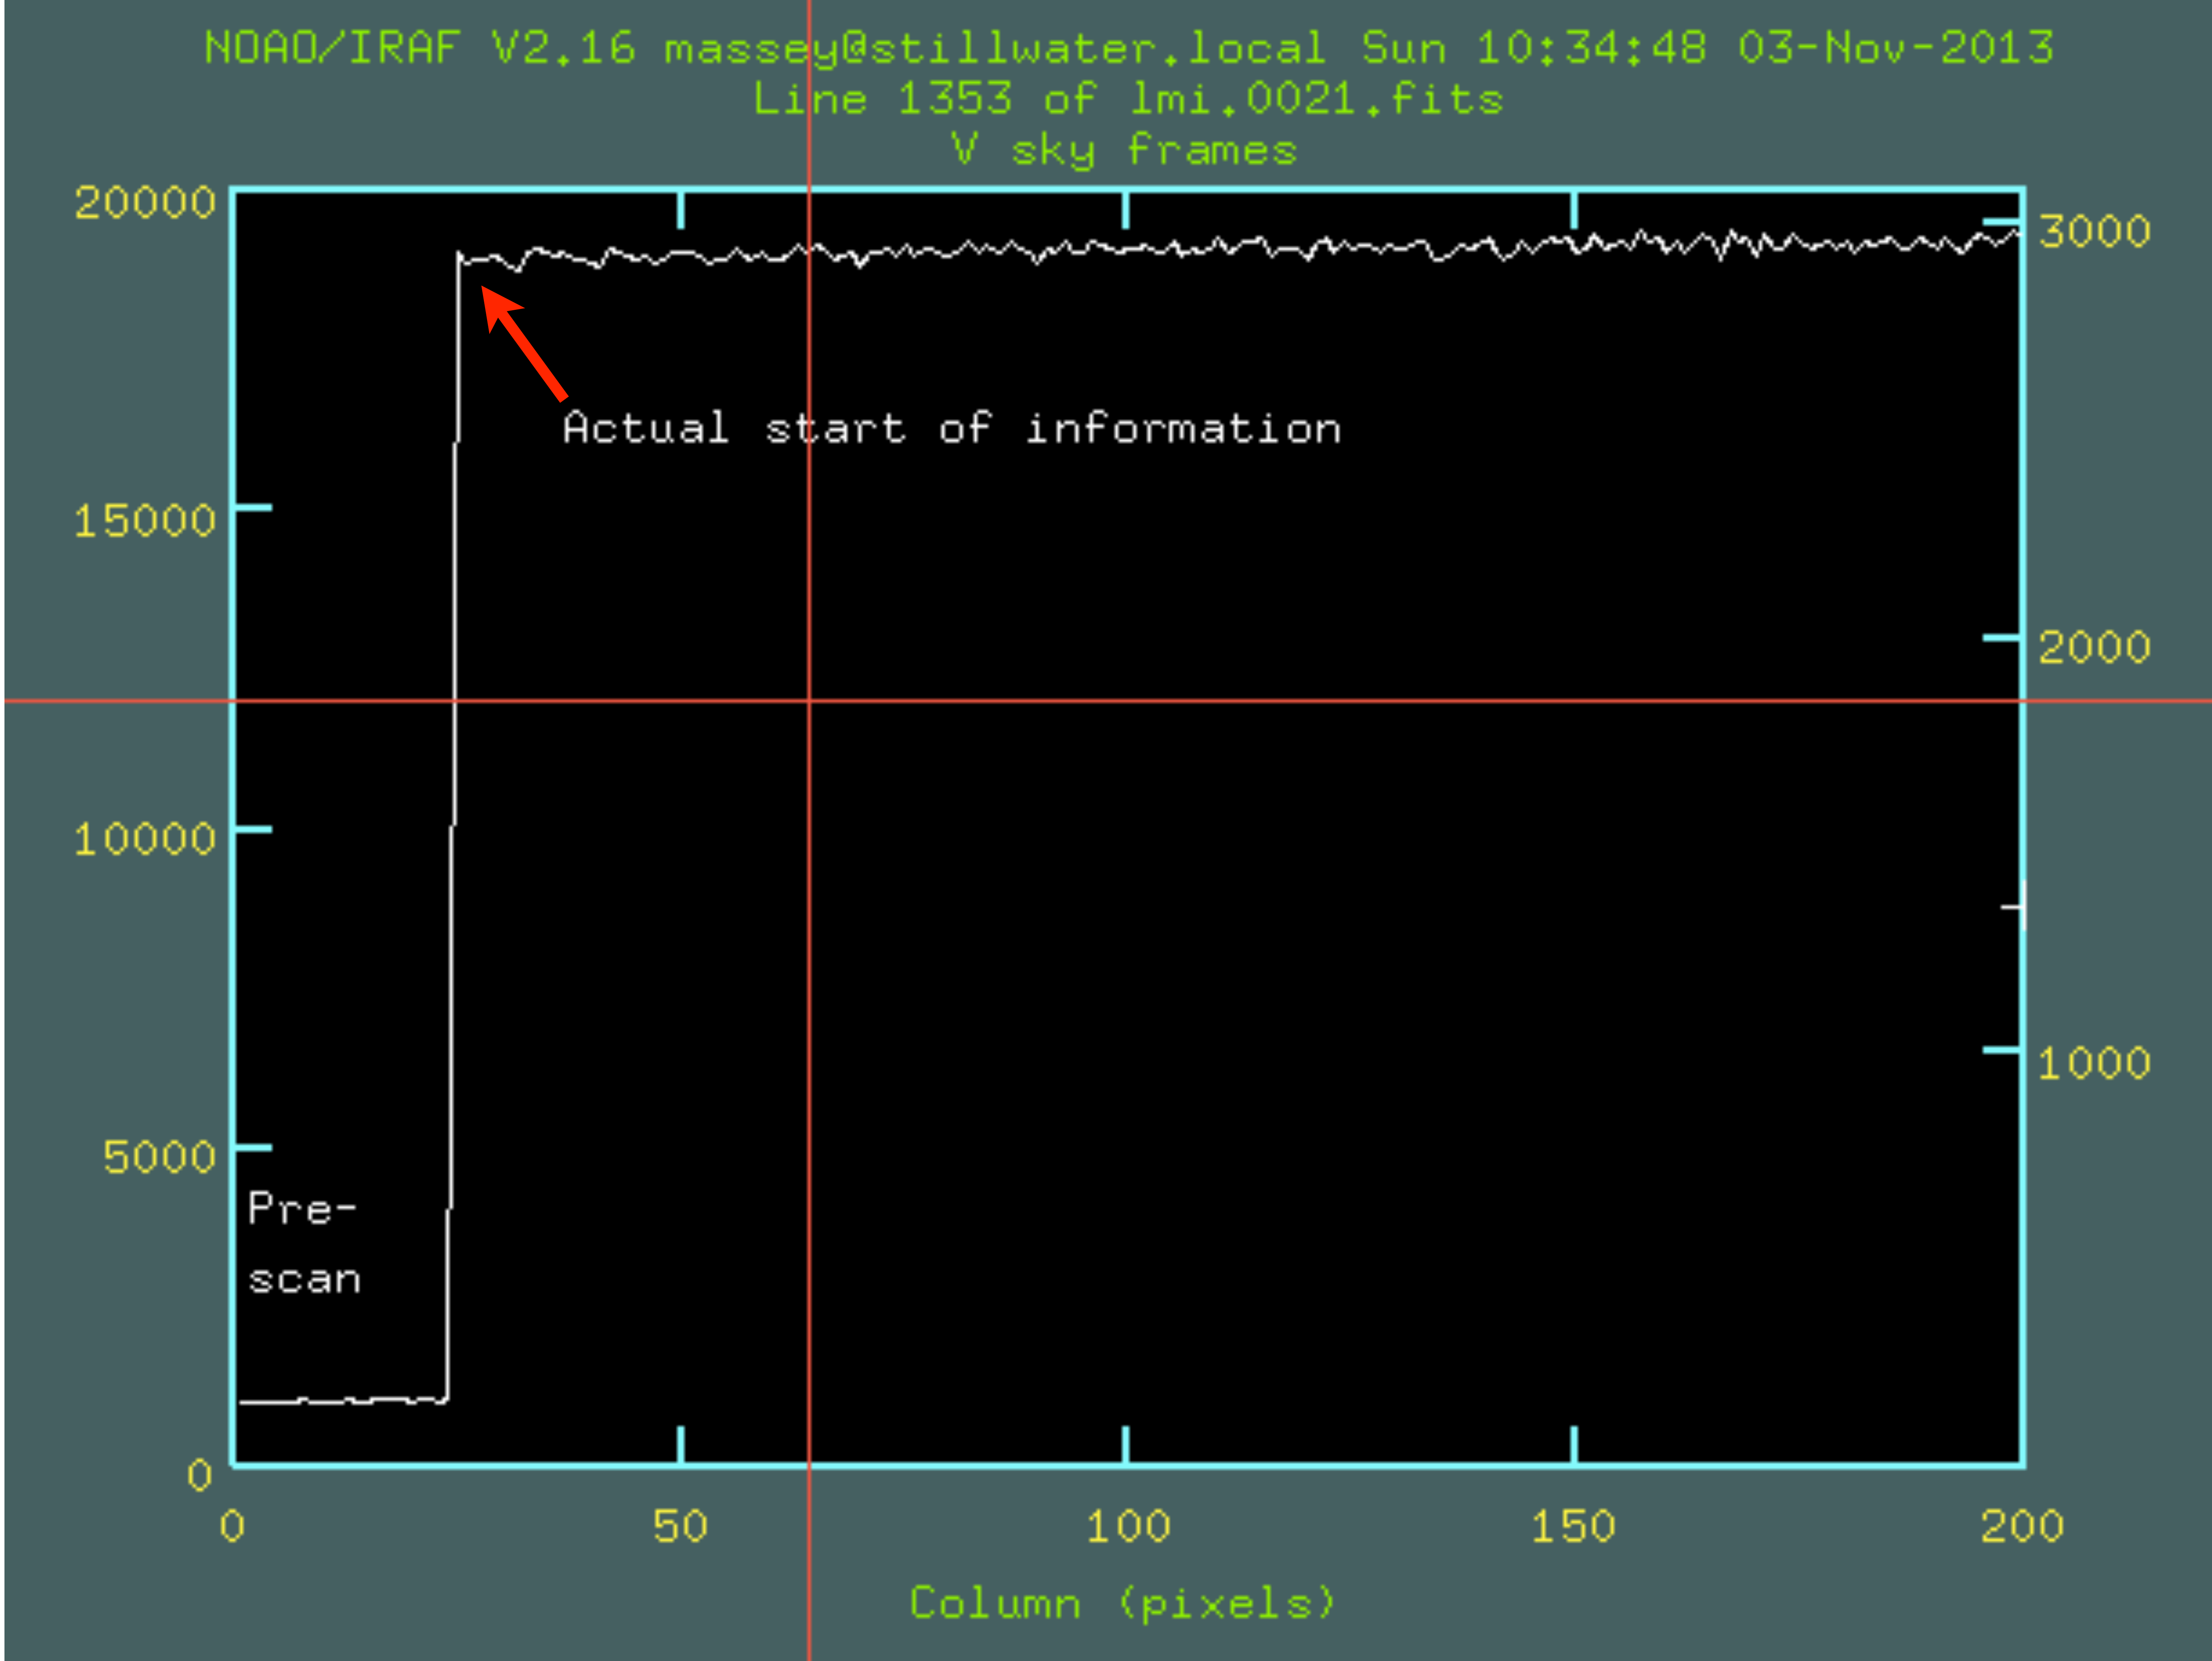

NOAO/IRAF V2.16 massey@stillwater.local Sun 10:34:48 03-Nov-2013
Line 1353 of lmi.0021.fits
V sky frames
Actual start of information
Pre-
scan
20000
15000
10000
5000
0
3000
2000
1000
0
50
100
150
200
Column (pixels)

Prescan

- The "prescan" a group of pixels generated at the start of each row by the ccd chip itself: roughly equivalent to the tuneup notes of an orchestra prior to commencing upon the oeuvre, but for each row…They do not contain information that can be considered useful for ordinary data reduction." ---Peter Collins, Lowell CCD guru
- These are physical pixels in the serial register but which aren't connect to the imaging area.
- In other words, the pre-scan should be thrown away.

Anatomy of a Raw CCD Image

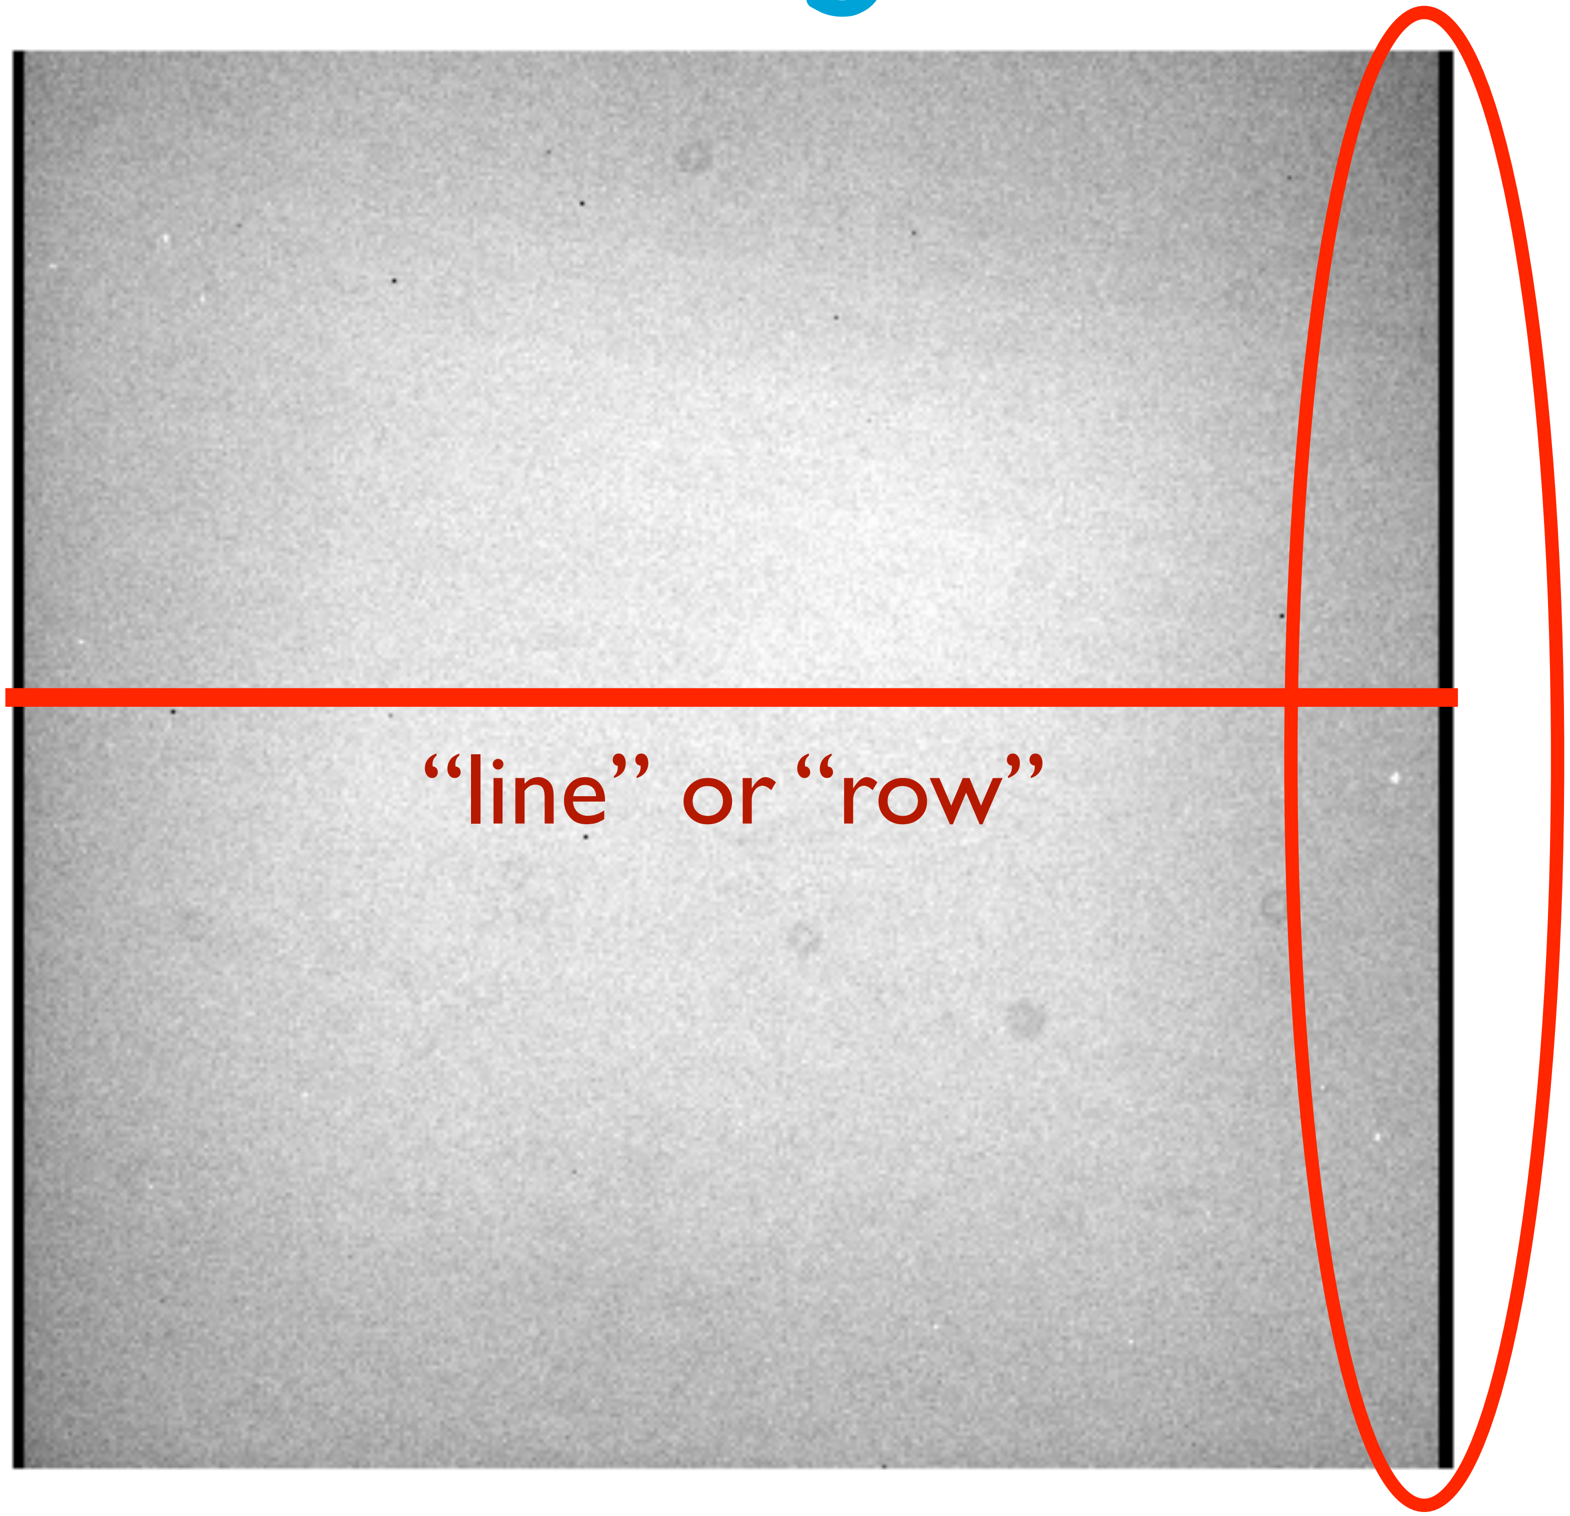

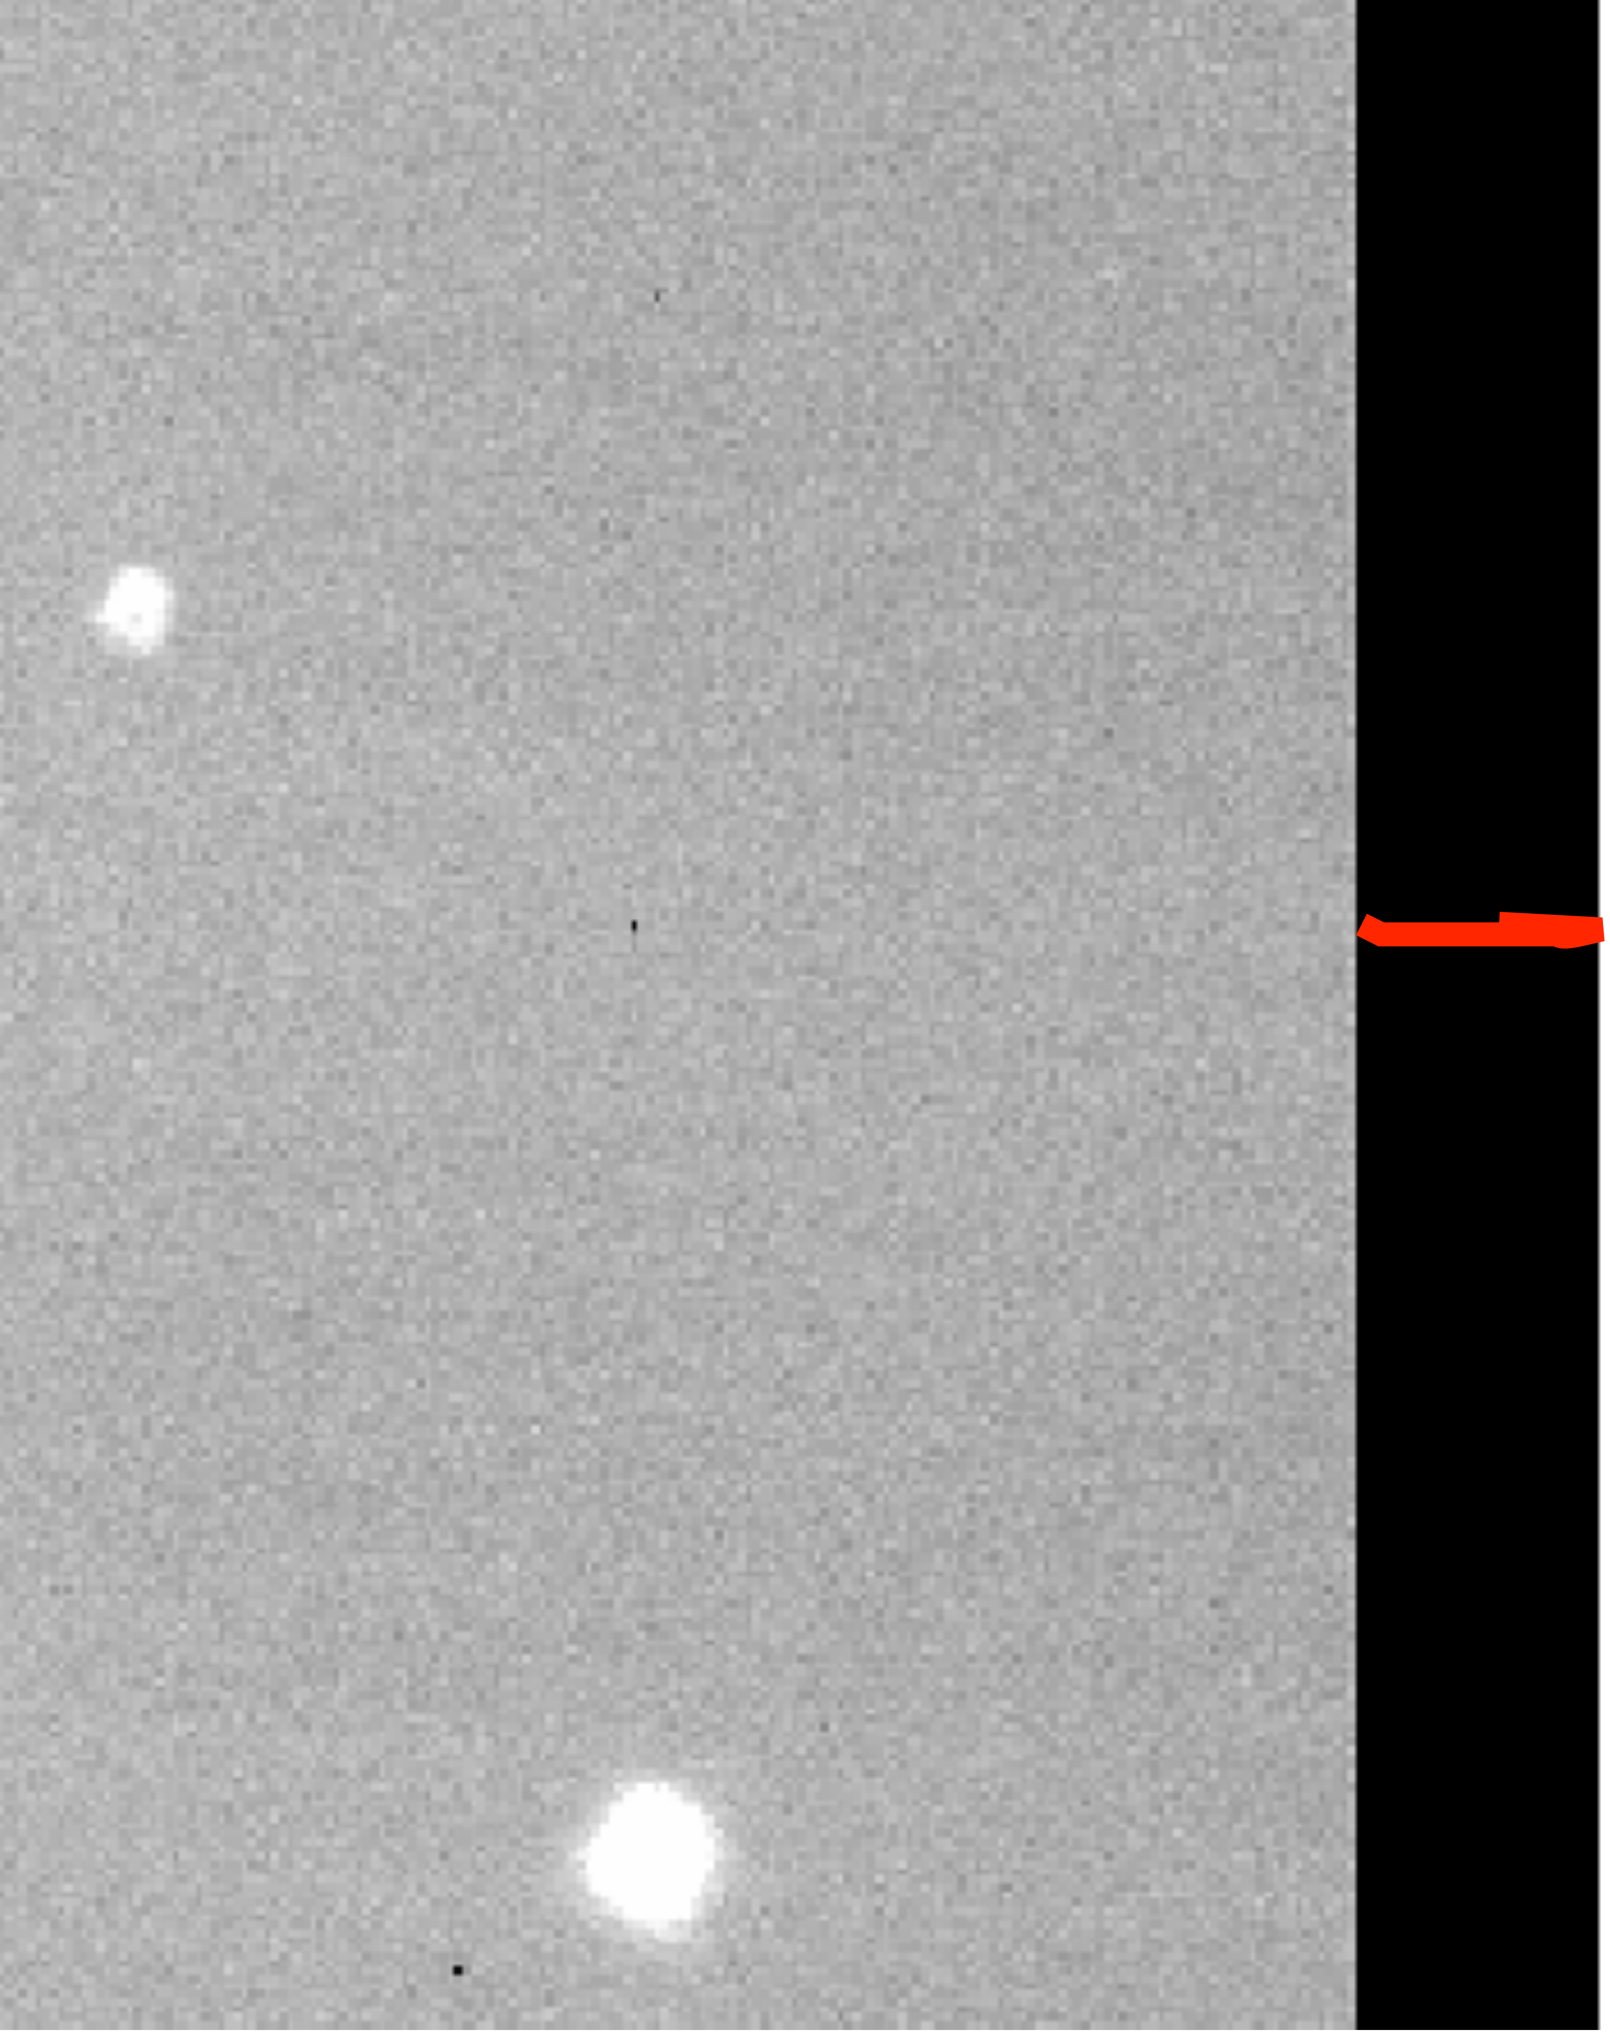

This region is the OVERSCAN
and you really need it

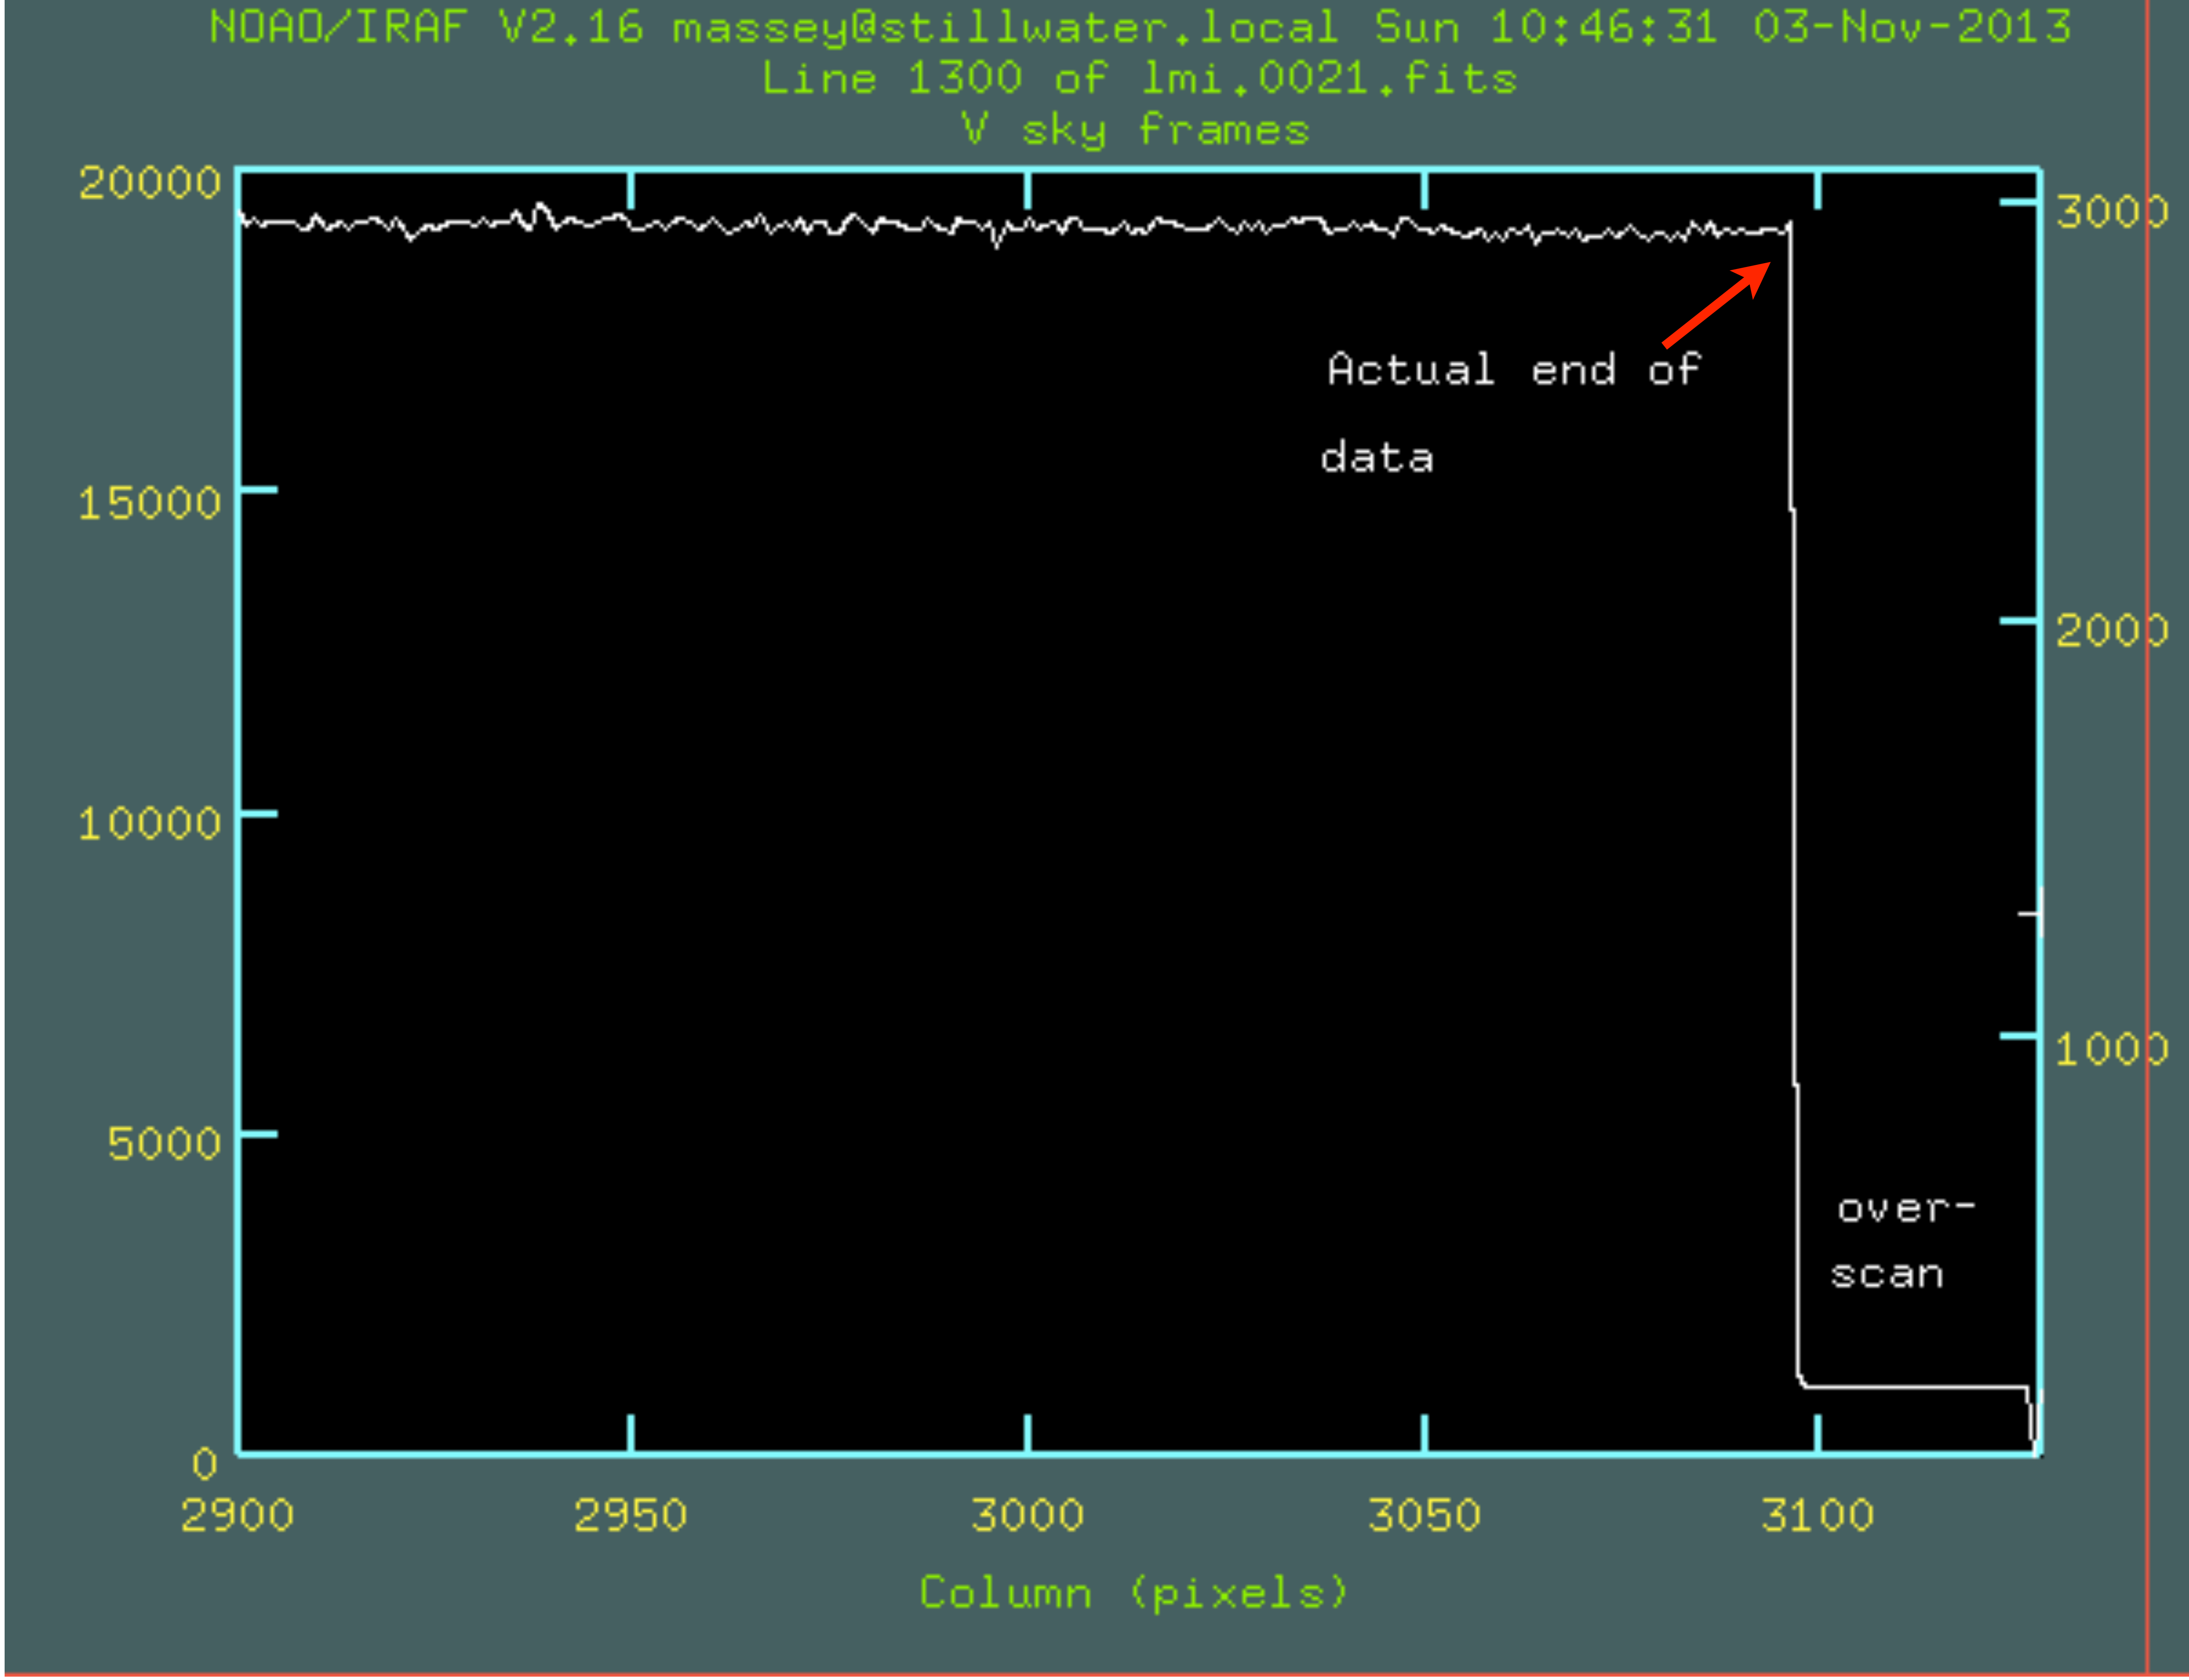
NOAO/IRAF V2.16 massey@stillwater.local Sun 10:46:31 03-Nov-2013
Line 1300 of lmi.0021.fits
V sky frames
Actual end of
data
over-
scan
20000
15000
10000
5000
0
3000
2000
1000
2900
2950
3000
3050
3100
Column (pixels)

The overscan

- This is how you keep track of the actual bias level for each frame. But to do that we want to make sure we identify the "flat" part of the overscan. They are never quite perfect.

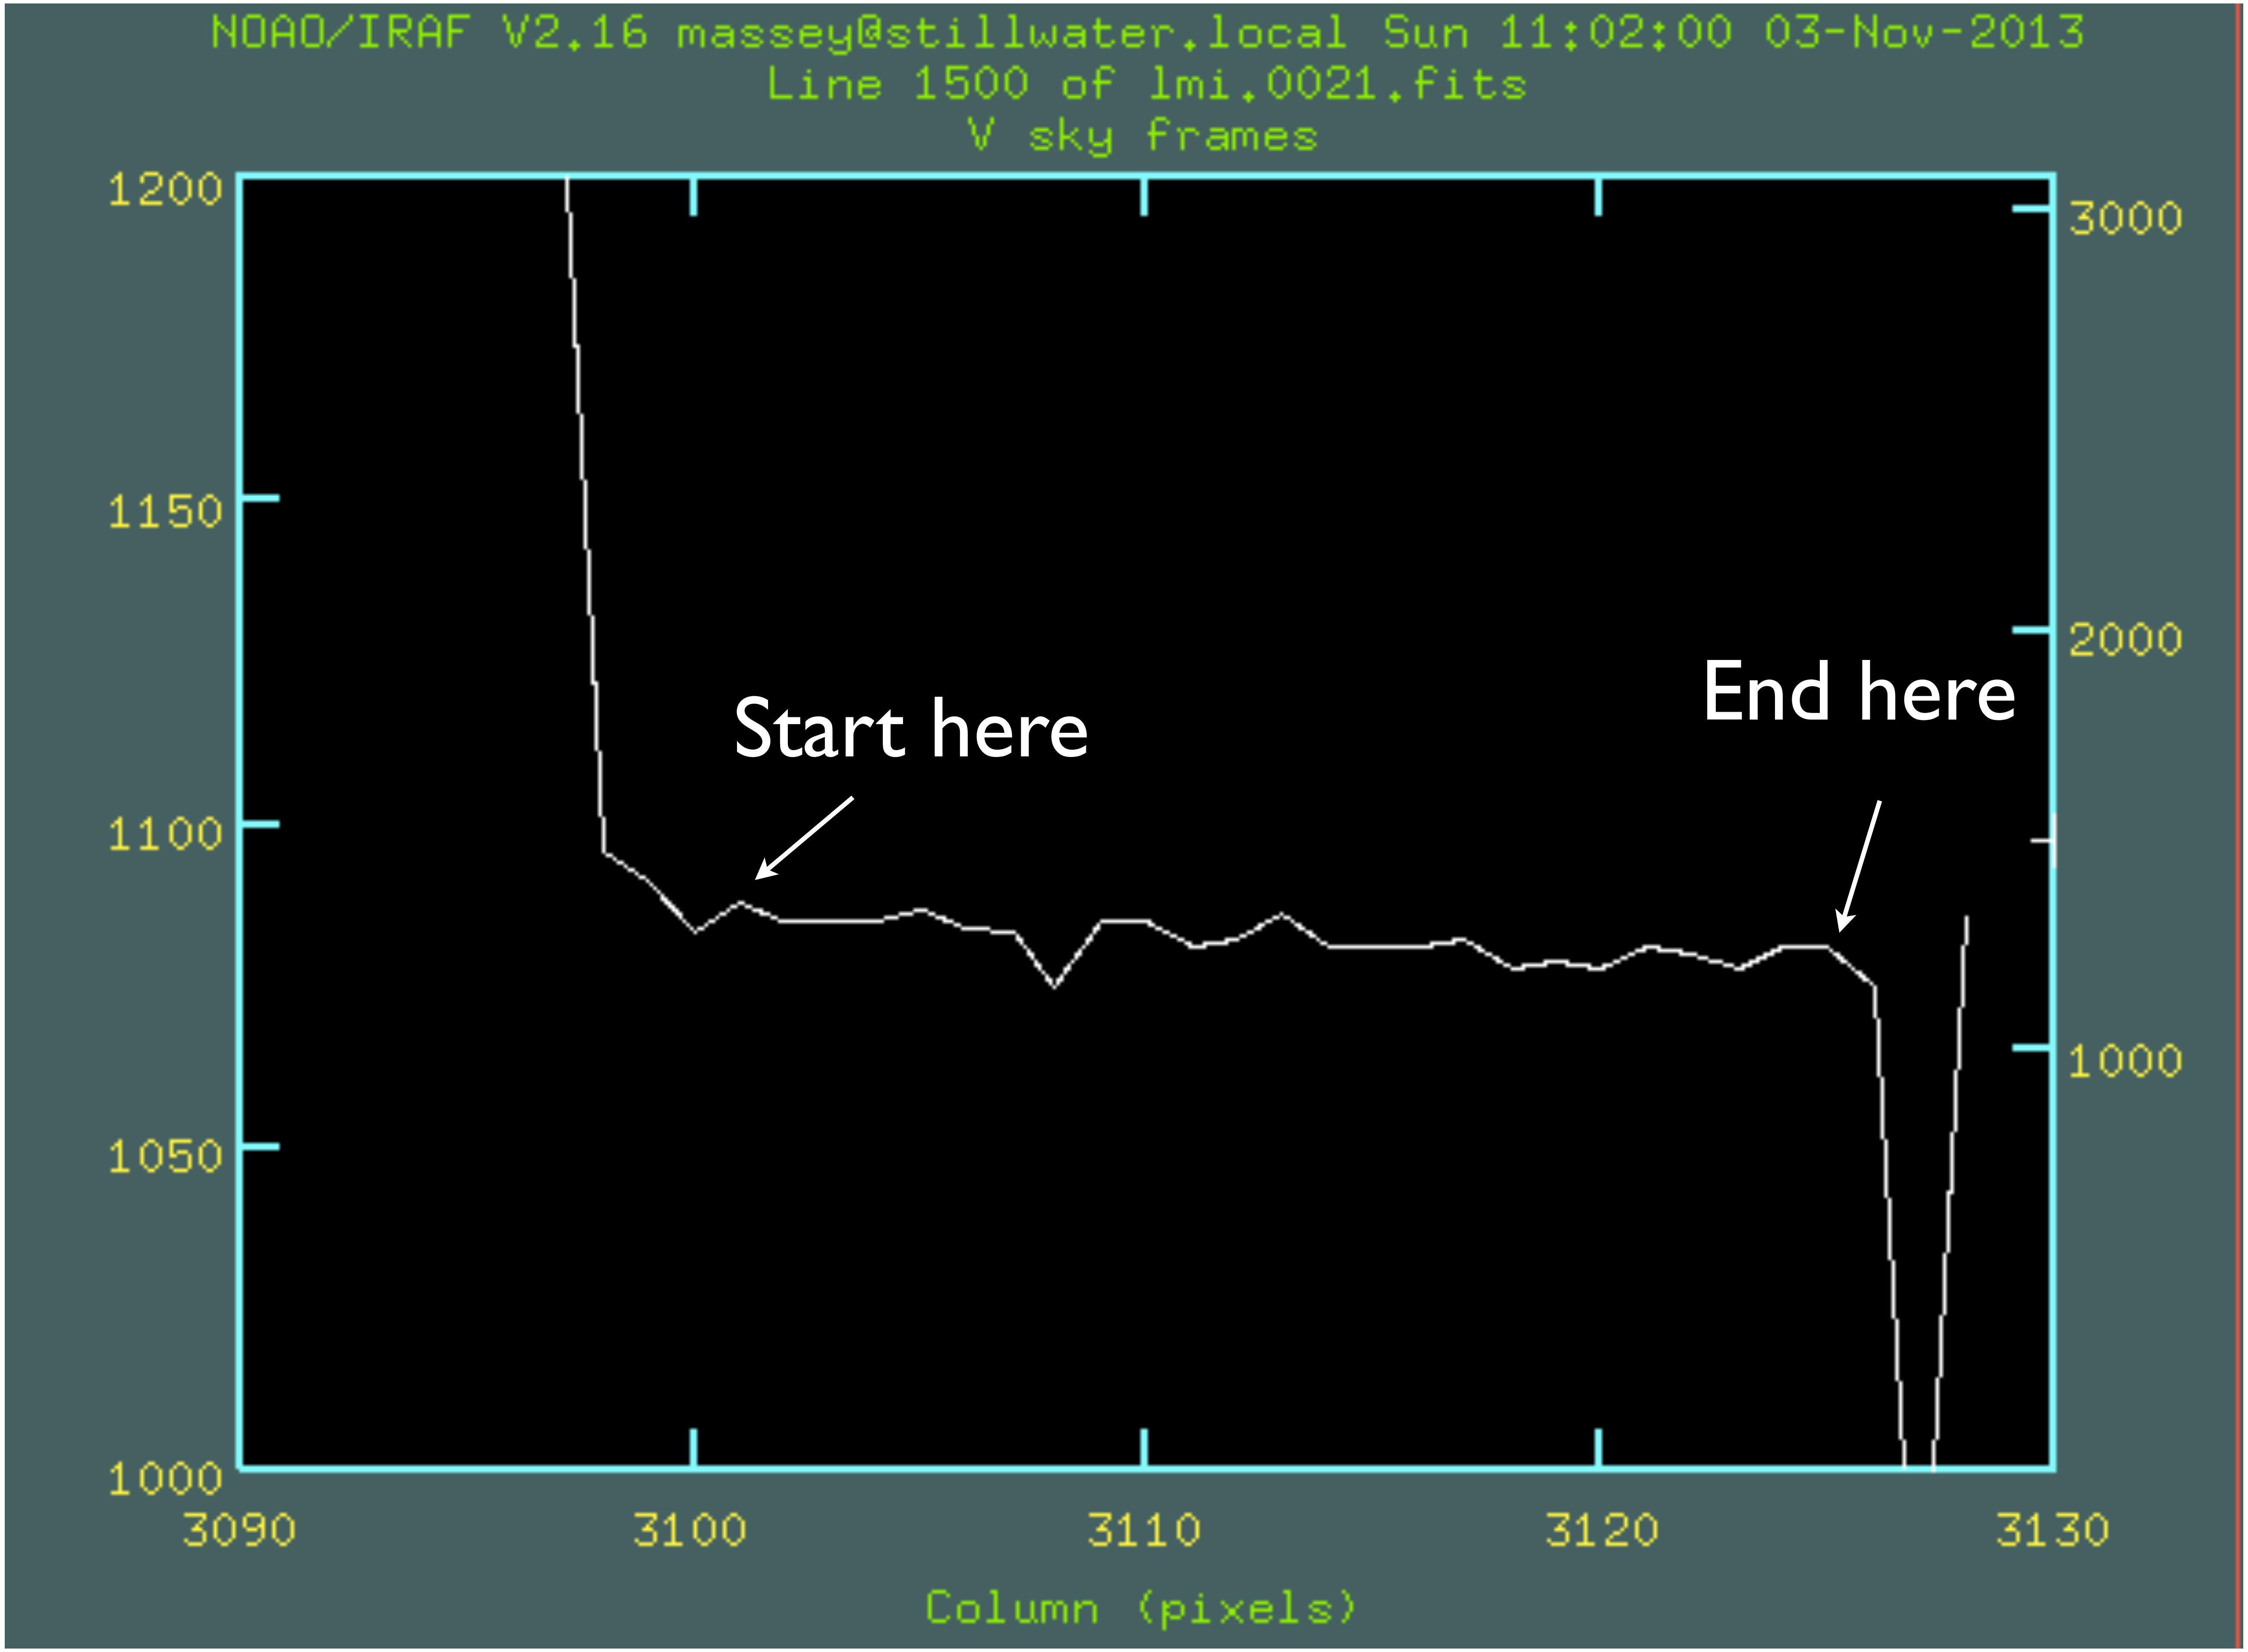

NOAO/IRAF V2.16 massey@stillwater.local Sun 11:02:00 03-Nov-2013
Line 1500 of lmi.0021.fits
V sky frames
1200
1150
1100
1050
1000
3000
2000
1000
Start here
End here
3090
3100
3110
3120
3130
Column (pixels)

Anatomy of a Raw CCD Image

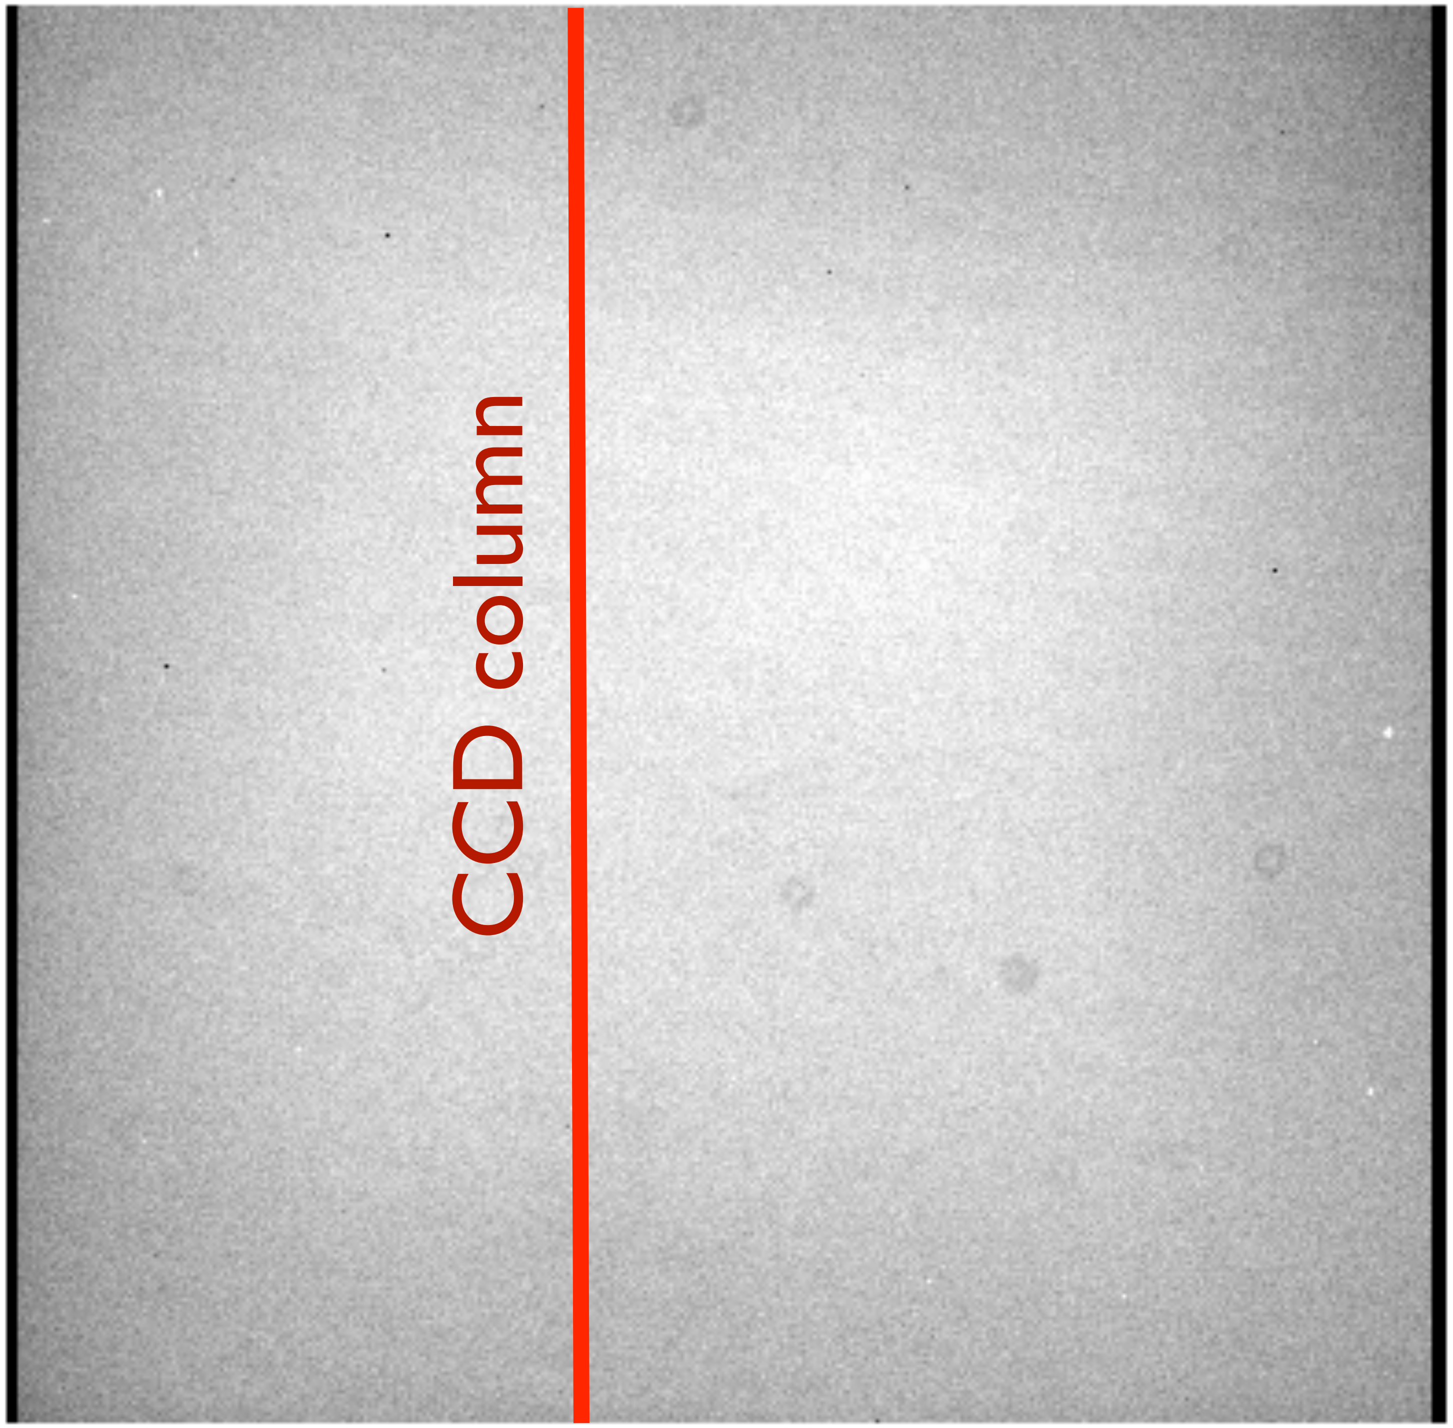

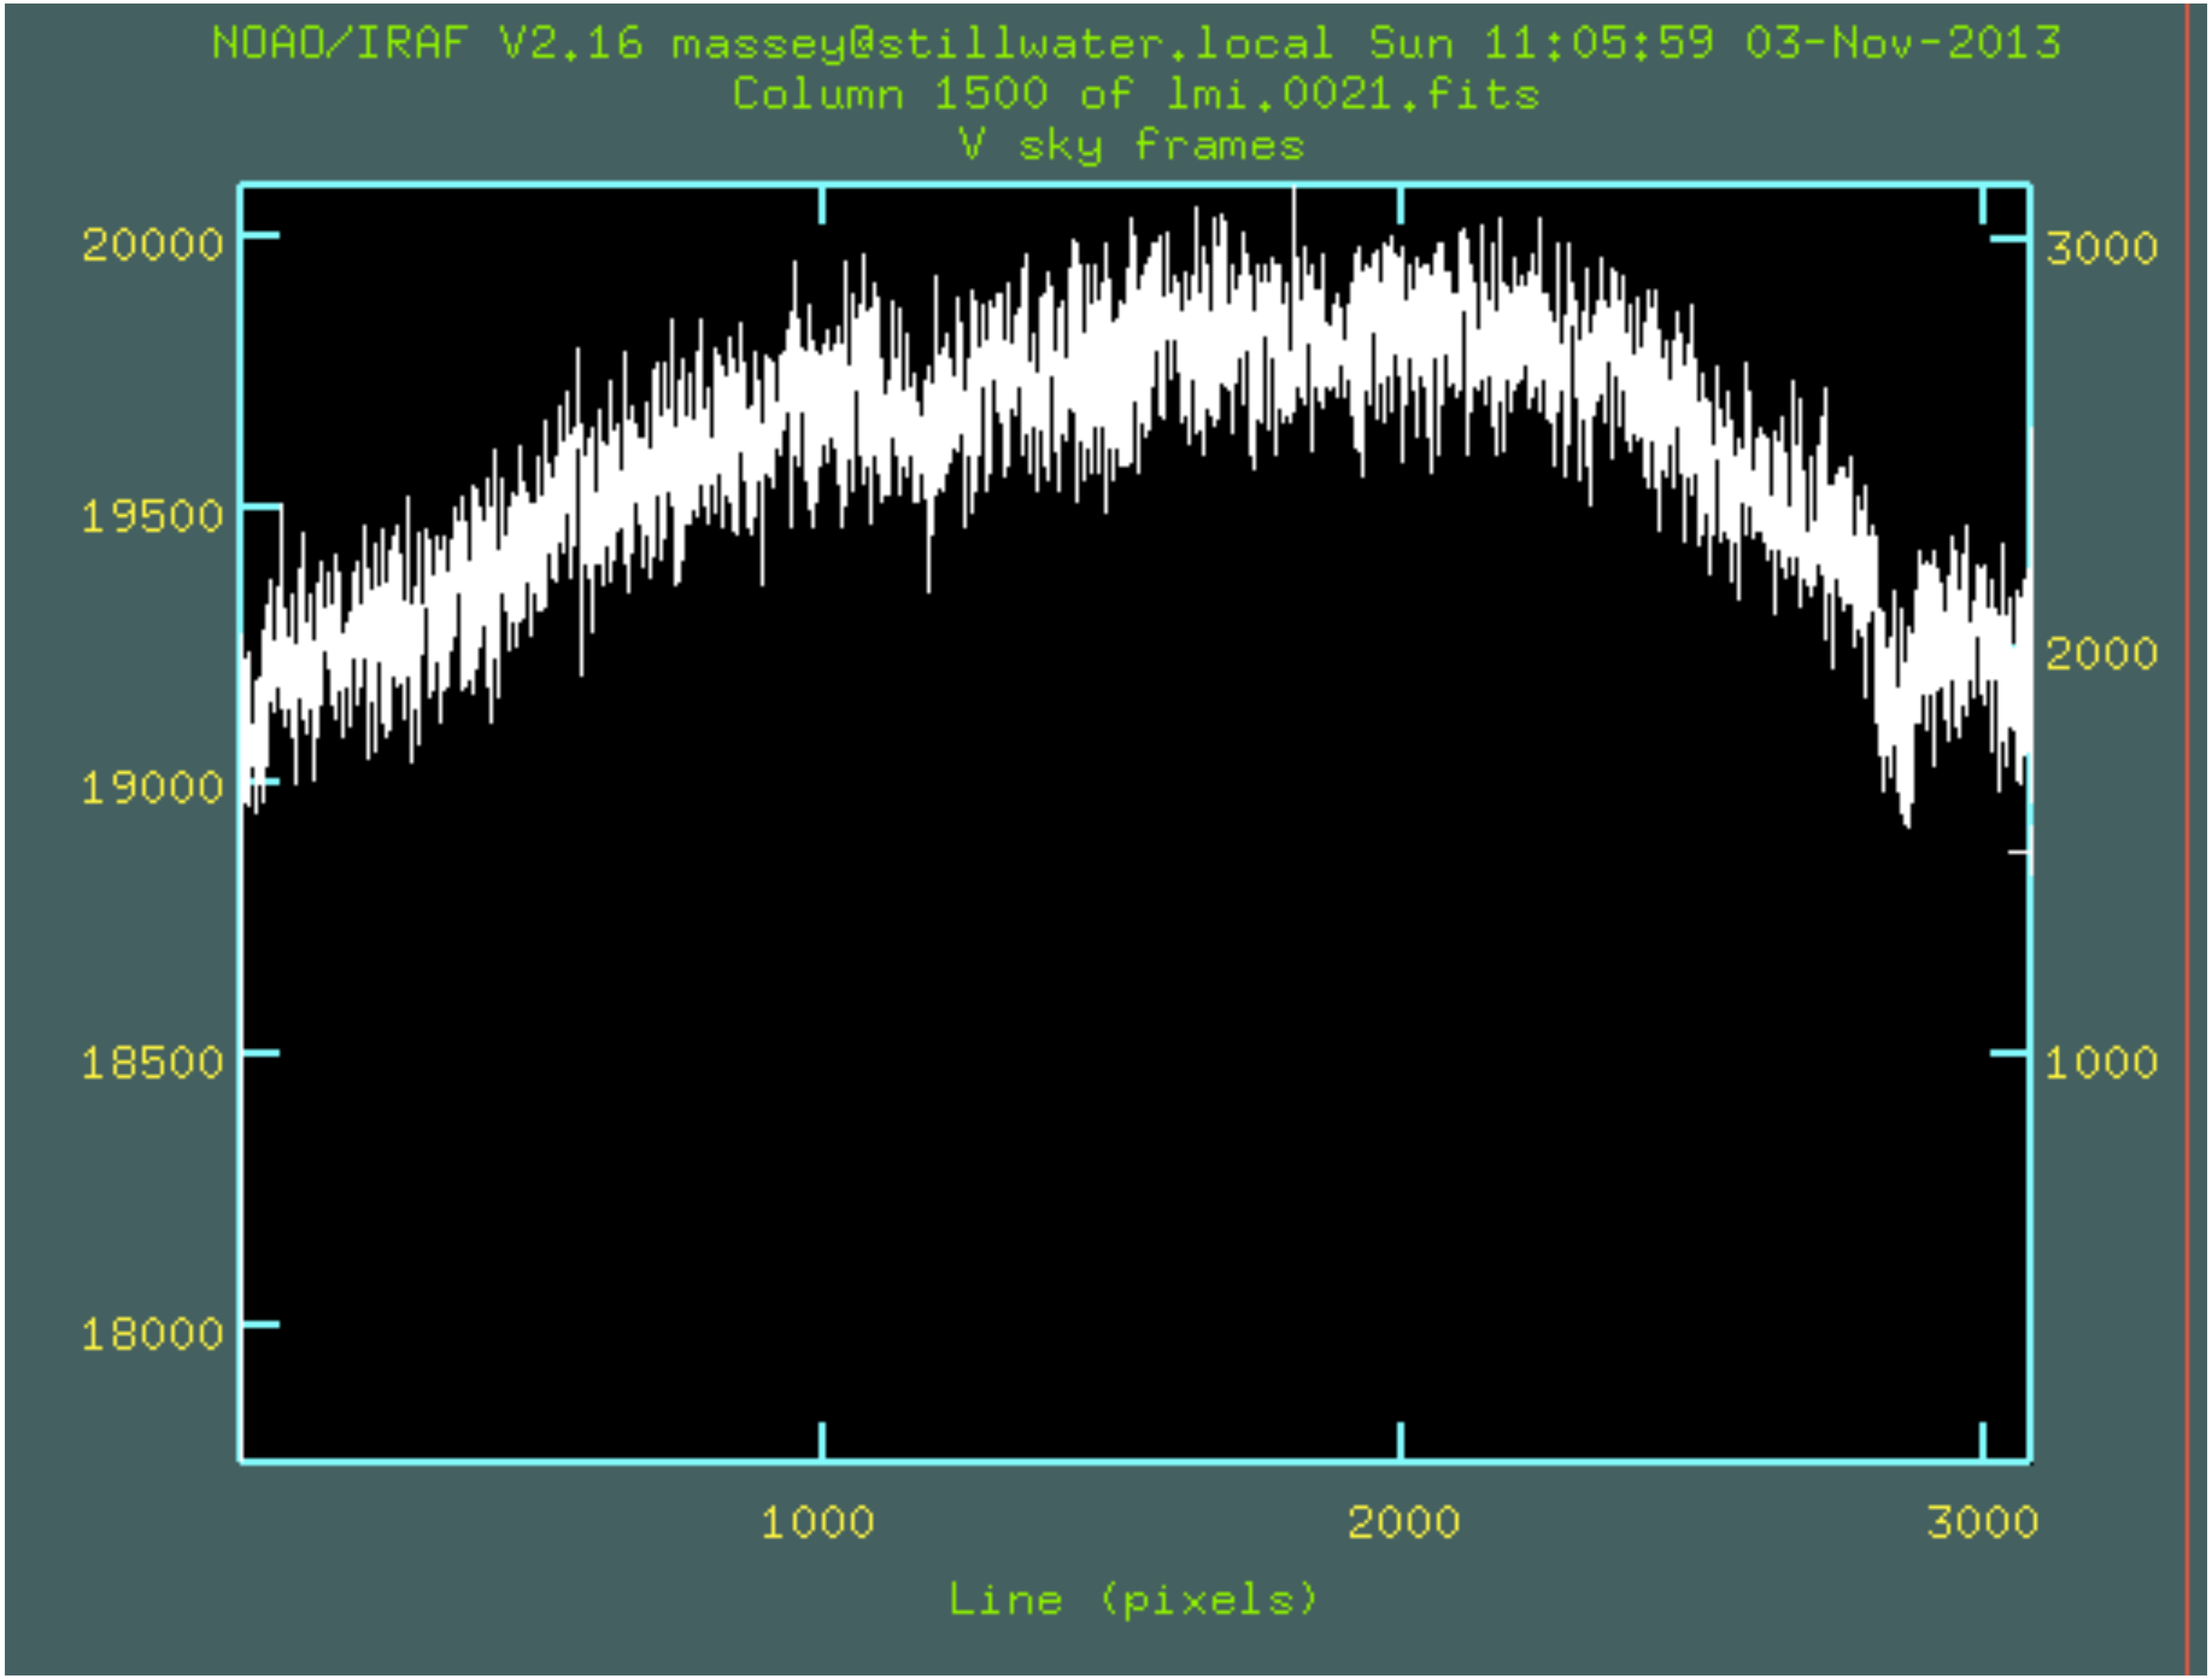
NOAO/IRAF V2.16 massey@stillwater.local Sun 11:05:59 03-Nov-2013
Column 1500 of lmi.0021.fits
V sky frames
20000
19500
19000
18500
18000
3000
2000
1000
1000
2000
3000
Line (pixels)

Pixels!

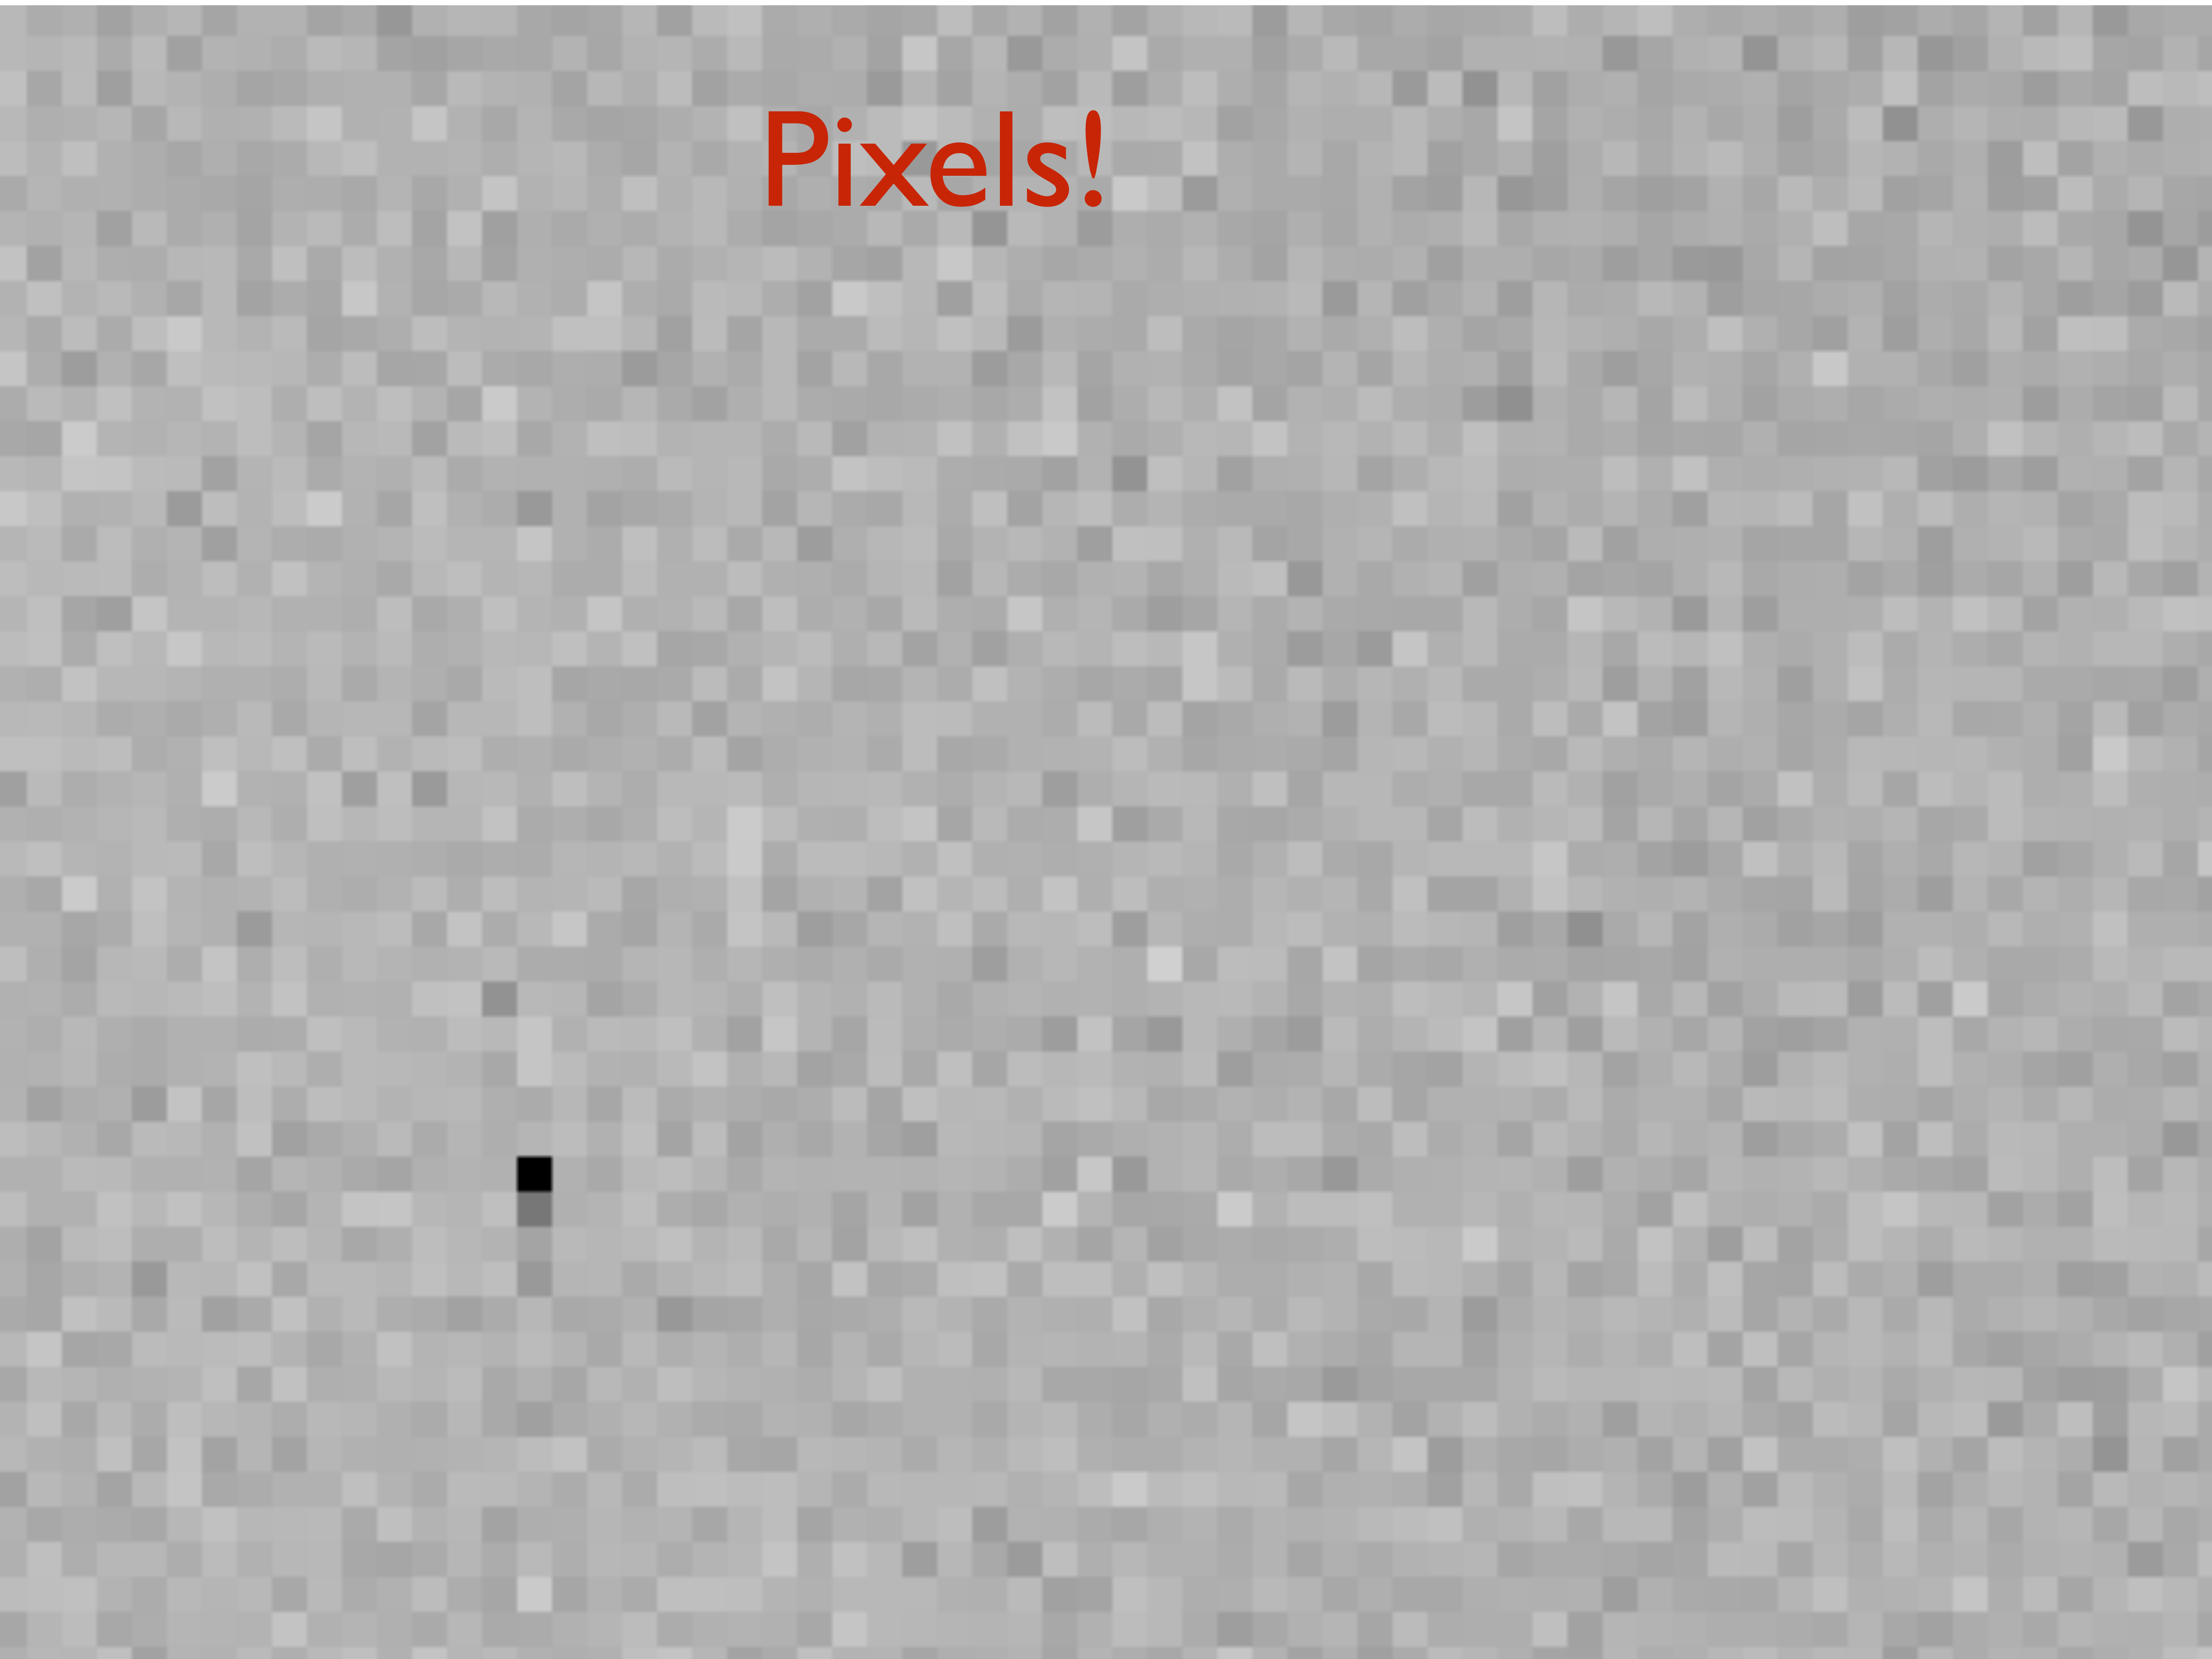

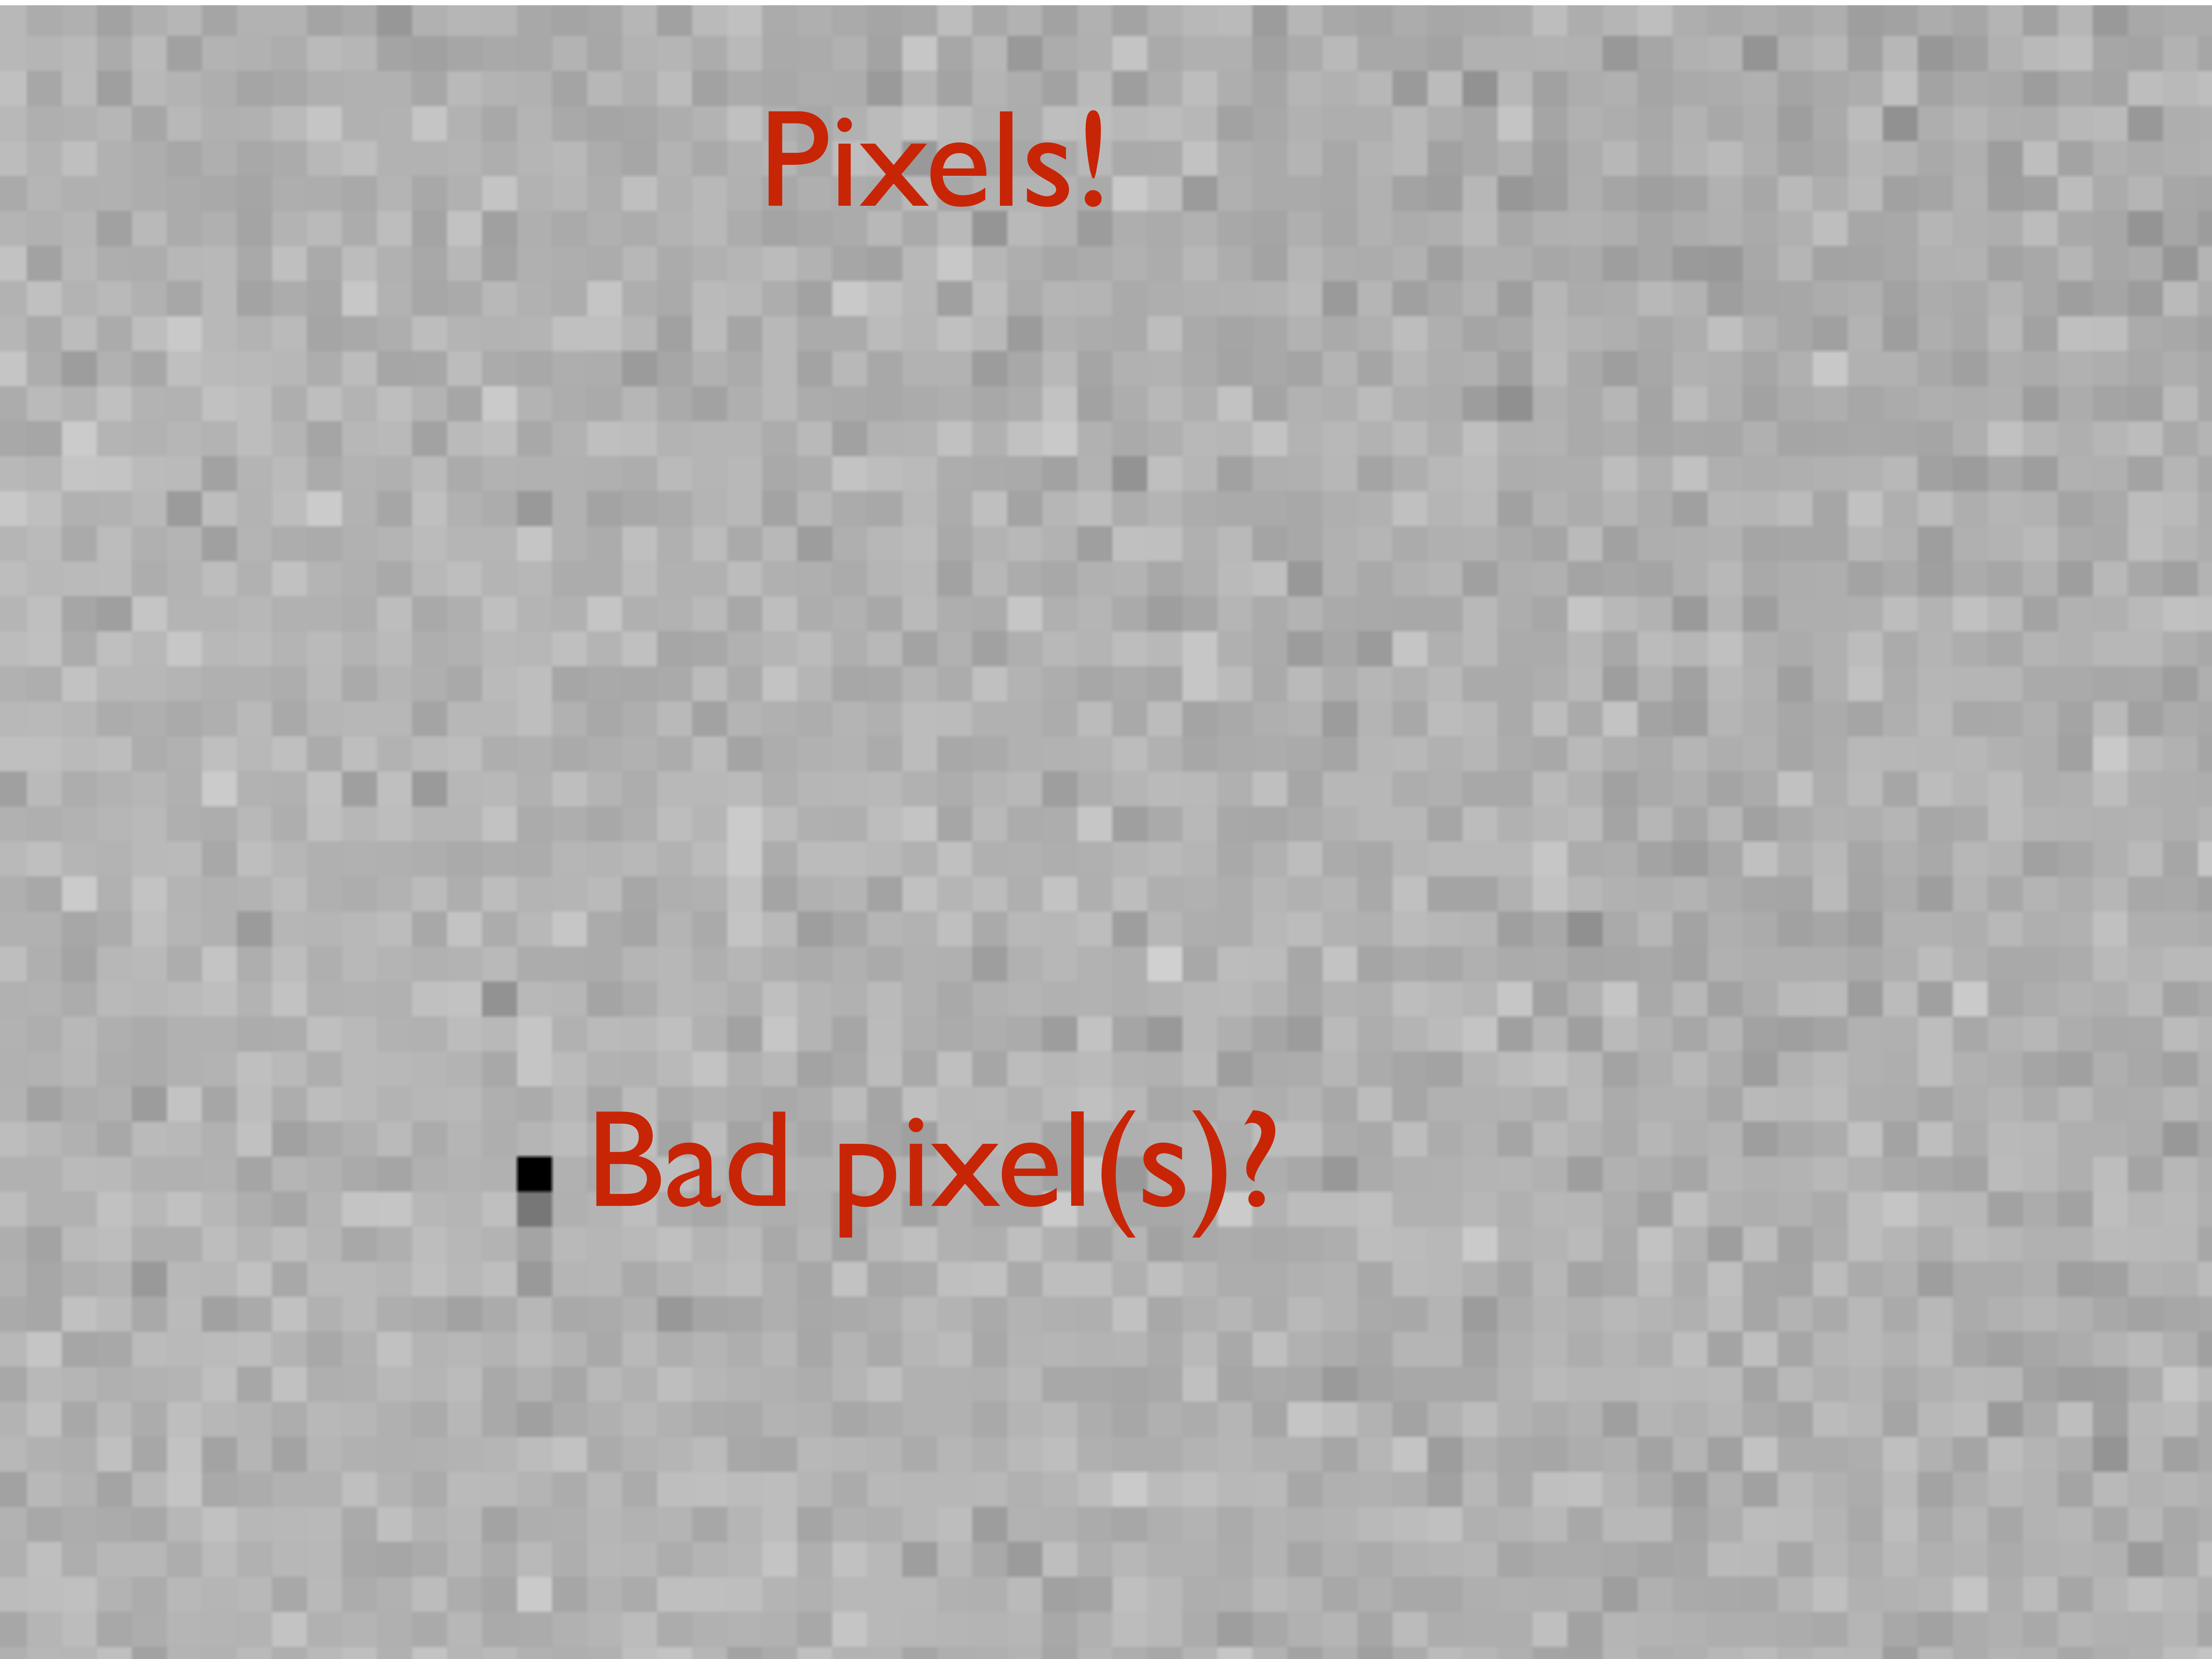
Pixels!
Bad pixel(s)?

Bad pixels

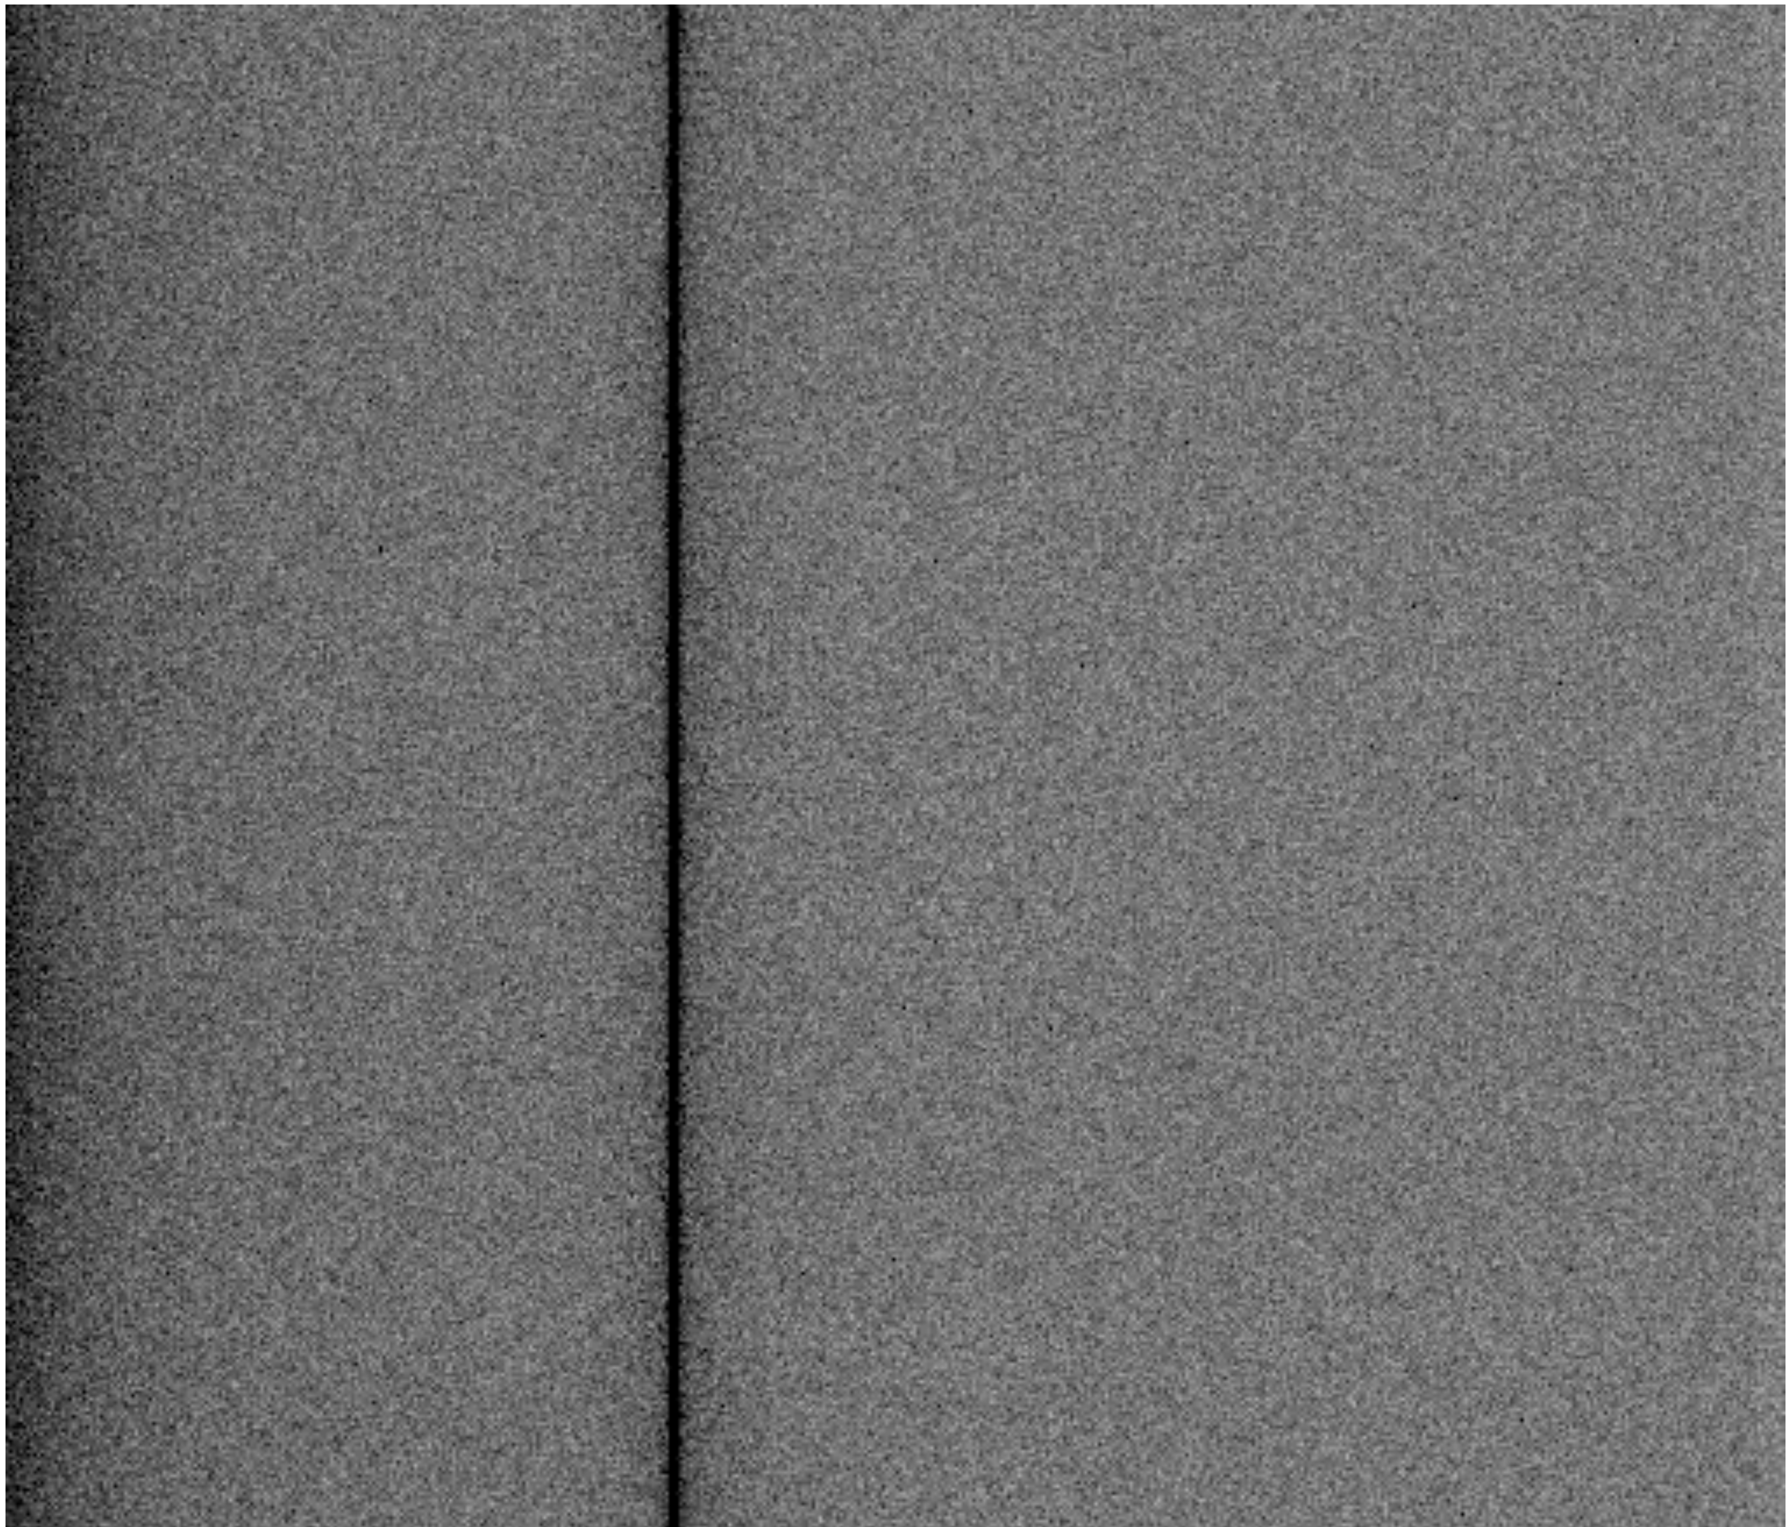

Nomenclature

- (Slightly IRAF-ese-ish but convenient):
- dataframe.fits[x1:x2,y1:y2] means the image section of the frame from column x1 to column x2 and from row y1 to row y2, inclusive. Thus the pixel at 100,200 would be called
- dataframe.fits[100:100,200:200]

SIDEBAR:
Technology vs Astronomy

Device dependent:

How a CCD works, bias, bias structure, how to reduce CCD data

Astronomy:

Errors add in quadrature, signal-to-noise

Summary of CCD

1000 photons falls on a pixel on a CCD. 80% of them result in electron/hole pairs. Those 800 electrons go through an amplifier, which adds a DC offset [bias] and amplifies the signal with some gain. This is then converted to an analog-to-digital unit by an A/D converter.

To recover the number of photons is unnecessary. But we do need to recover the number of electrons if we are to do S/N calculations. We must take the signal we see (number of counts, or ADUs), subtract bias, and multiply by the gain (e/ADUs).

Reminder about gains…

What is a typical gain for a CCD camera (i.e., e/ADUs) and WHY? Hint: there are two competing issues, one driving it to a lower number, and one driving it to a higher number.

And what IS the gain?

The gain is the proportionality constant between the counts you get out (ADUs) and the number of electrons that were detected. *You need to know this so you can calculate errors!* Sigma_photon = sqrt(N) = sqrt(gain x ADUs-above-bias)

Another Side Bar:
Propagation of Errors!

Propagation of Errors

Imagine that you're making measurements using multiple "machines." The result of each machine is a, b, c, d... The end result of combining these all (in some way) is x, so that x is dependent on a, b, c, etc. Each time you "run the machine" you get a slightly different a, b, and c, because there's some error associated with each. What is the ultimate effect on x?

$$x = f(a,b,c\ldots)$$

$$(dx)^2 = (dx/da)^2 \ (da)^2 + (dx/db)^2 \ (db)^2 + (dx/dc)^2 \ (dc)^2$$

where (dx/da) is really the partial derivative of x wrt a.

Propagation of Errors

$(dx)^2 = (dx/da)^2 \ (da)^2 + (dx/db)^2 \ (db)^2 + (dx/dc)^2 (dc)^2$

Consider the simple case then that $x = a + b + c$. Take the partial derivatives:

$dx/da = 1$, $(dx/db)=1$, and $(dx/dc)=1$. You're left with:

$(dx)^2 = (da)^2 + (db)^2 + (dc)^2$

in other words:

$\sigma^2_{tot}= \sigma^2_a + \sigma^2_b + \sigma^2_c$

Propagation of Errors

$(dx)^2 = (dx/da)^2 \ (da)^2 + (dx/db)^2 \ (db)^2 + (dx/dc)^2 (dc)^2$

What if instead $x = a\,(b/c)$? This is also straight-forward, depending upon how well you remember your Calculus!

In the end, the answer is simple enough:

$(\sigma_x/x)^2 = (\sigma_a/a)^2 + (\sigma_b/b)^2 + (\sigma_c/c)^2$

Calibration of CCD Data

Two kinds of effects we must correct for:

a) Additive Effects

- bias level (from the overscan)
- bias structure (in y direction from overscan fitting, and in x direction (and any residual y) from bias frames)
- dark current (if any)

b) Multiplicative Effects

- Pixel-to-pixel variations (from flat field)
- Large scale illumination corrections (from dark sky flats, perhaps)

Types of calibration data

1) Additive

- Data frames: (what you care about). Contain the overscan.
- Bias frames: Determines residual bias structure. Series of 0-sec exposures with the shutter closed.
- Dark frames: Seldom necessary. Take 3 each of typical exposure times if >5 minutes.

2) Multiplicative

- Dome flats (bright source, not necessarily uniform).
- Sky flats (dim source, but uniform)

Calibration of CCD Data

- Important that we do these in the right order. You MUST take care of the additive effects FIRST.
- That's because (A-B)/F != A/F - B. Note that it IS equal to A/F - B/F so if you know what you're doing you can compensate. If.

Bias levels

- As you take each CCD frame, the overall pedestal (bias) level will be a little bit different. The "extra" voltage that has been added to the signal is never quite the same.

Bias levels

```
d002.fits: Oct  3 16:50 Trim data section is [2:2047,2:3149]
d002.fits: Oct  3 16:50 Overscan section is [2050:2080,2:3149] with mean=505.1579
d002.fits: Oct  3 16:50 Trim data section is [2:2047,2:3149]
d002.fits: Oct  3 16:50 Overscan section is [2050:2080,2:3149] with mean=505.1579
d003.fits: Oct  3 16:50 Trim data section is [2:2047,2:3149]
d003.fits: Oct  3 16:50 Overscan section is [2050:2080,2:3149] with mean=505.0079
d004.fits: Oct  3 16:50 Trim data section is [2:2047,2:3149]
d004.fits: Oct  3 16:50 Overscan section is [2050:2080,2:3149] with mean=504.5082
d005.fits: Oct  3 16:50 Trim data section is [2:2047,2:3149]
d005.fits: Oct  3 16:50 Overscan section is [2050:2080,2:3149] with mean=503.9669
d006.fits: Oct  3 16:50 Trim data section is [2:2047,2:3149]
d006.fits: Oct  3 16:50 Overscan section is [2050:2080,2:3149] with mean=503.6663
d007.fits: Oct  3 16:50 Trim data section is [2:2047,2:3149]
d007.fits: Oct  3 16:50 Overscan section is [2050:2080,2:3149] with mean=503.4144
d008.fits: Oct  3 16:51 Trim data section is [2:2047,2:3149]
d008.fits: Oct  3 16:51 Overscan section is [2050:2080,2:3149] with mean=503.0336
d009.fits: Oct  3 16:51 Trim data section is [2:2047,2:3149]
d009.fits: Oct  3 16:51 Overscan section is [2050:2080,2:3149] with mean=502.8745
d010.fits: Oct  3 16:51 Trim data section is [2:2047,2:3149]
d010.fits: Oct  3 16:51 Overscan section is [2050:2080,2:3149] with mean=502.6663
d011.fits: Oct  3 16:51 Trim data section is [2:2047,2:3149]
d011.fits: Oct  3 16:51 Overscan section is [2050:2080,2:3149] with mean=502.2831
```

Fitting the overscan

- Remember that the overscan is an image section over on the right that extends from the bottom to the top of the chip. We can either choose to fit this as a constant (average of all values in the "clean" overscan region) or you can fit it in y.

Fitting the overscan

- Step 1: Identify the overscan. Might as well decide how much of the chip you're going to save (trim section) at the same time.
- Ironically, you don't want to identify the overscan (bias) region on a bias frame. That's because there's no contrast: the average value on the bias frame is very similar to that of the overscan. So, instead, make some plots of a flat-field exposure.

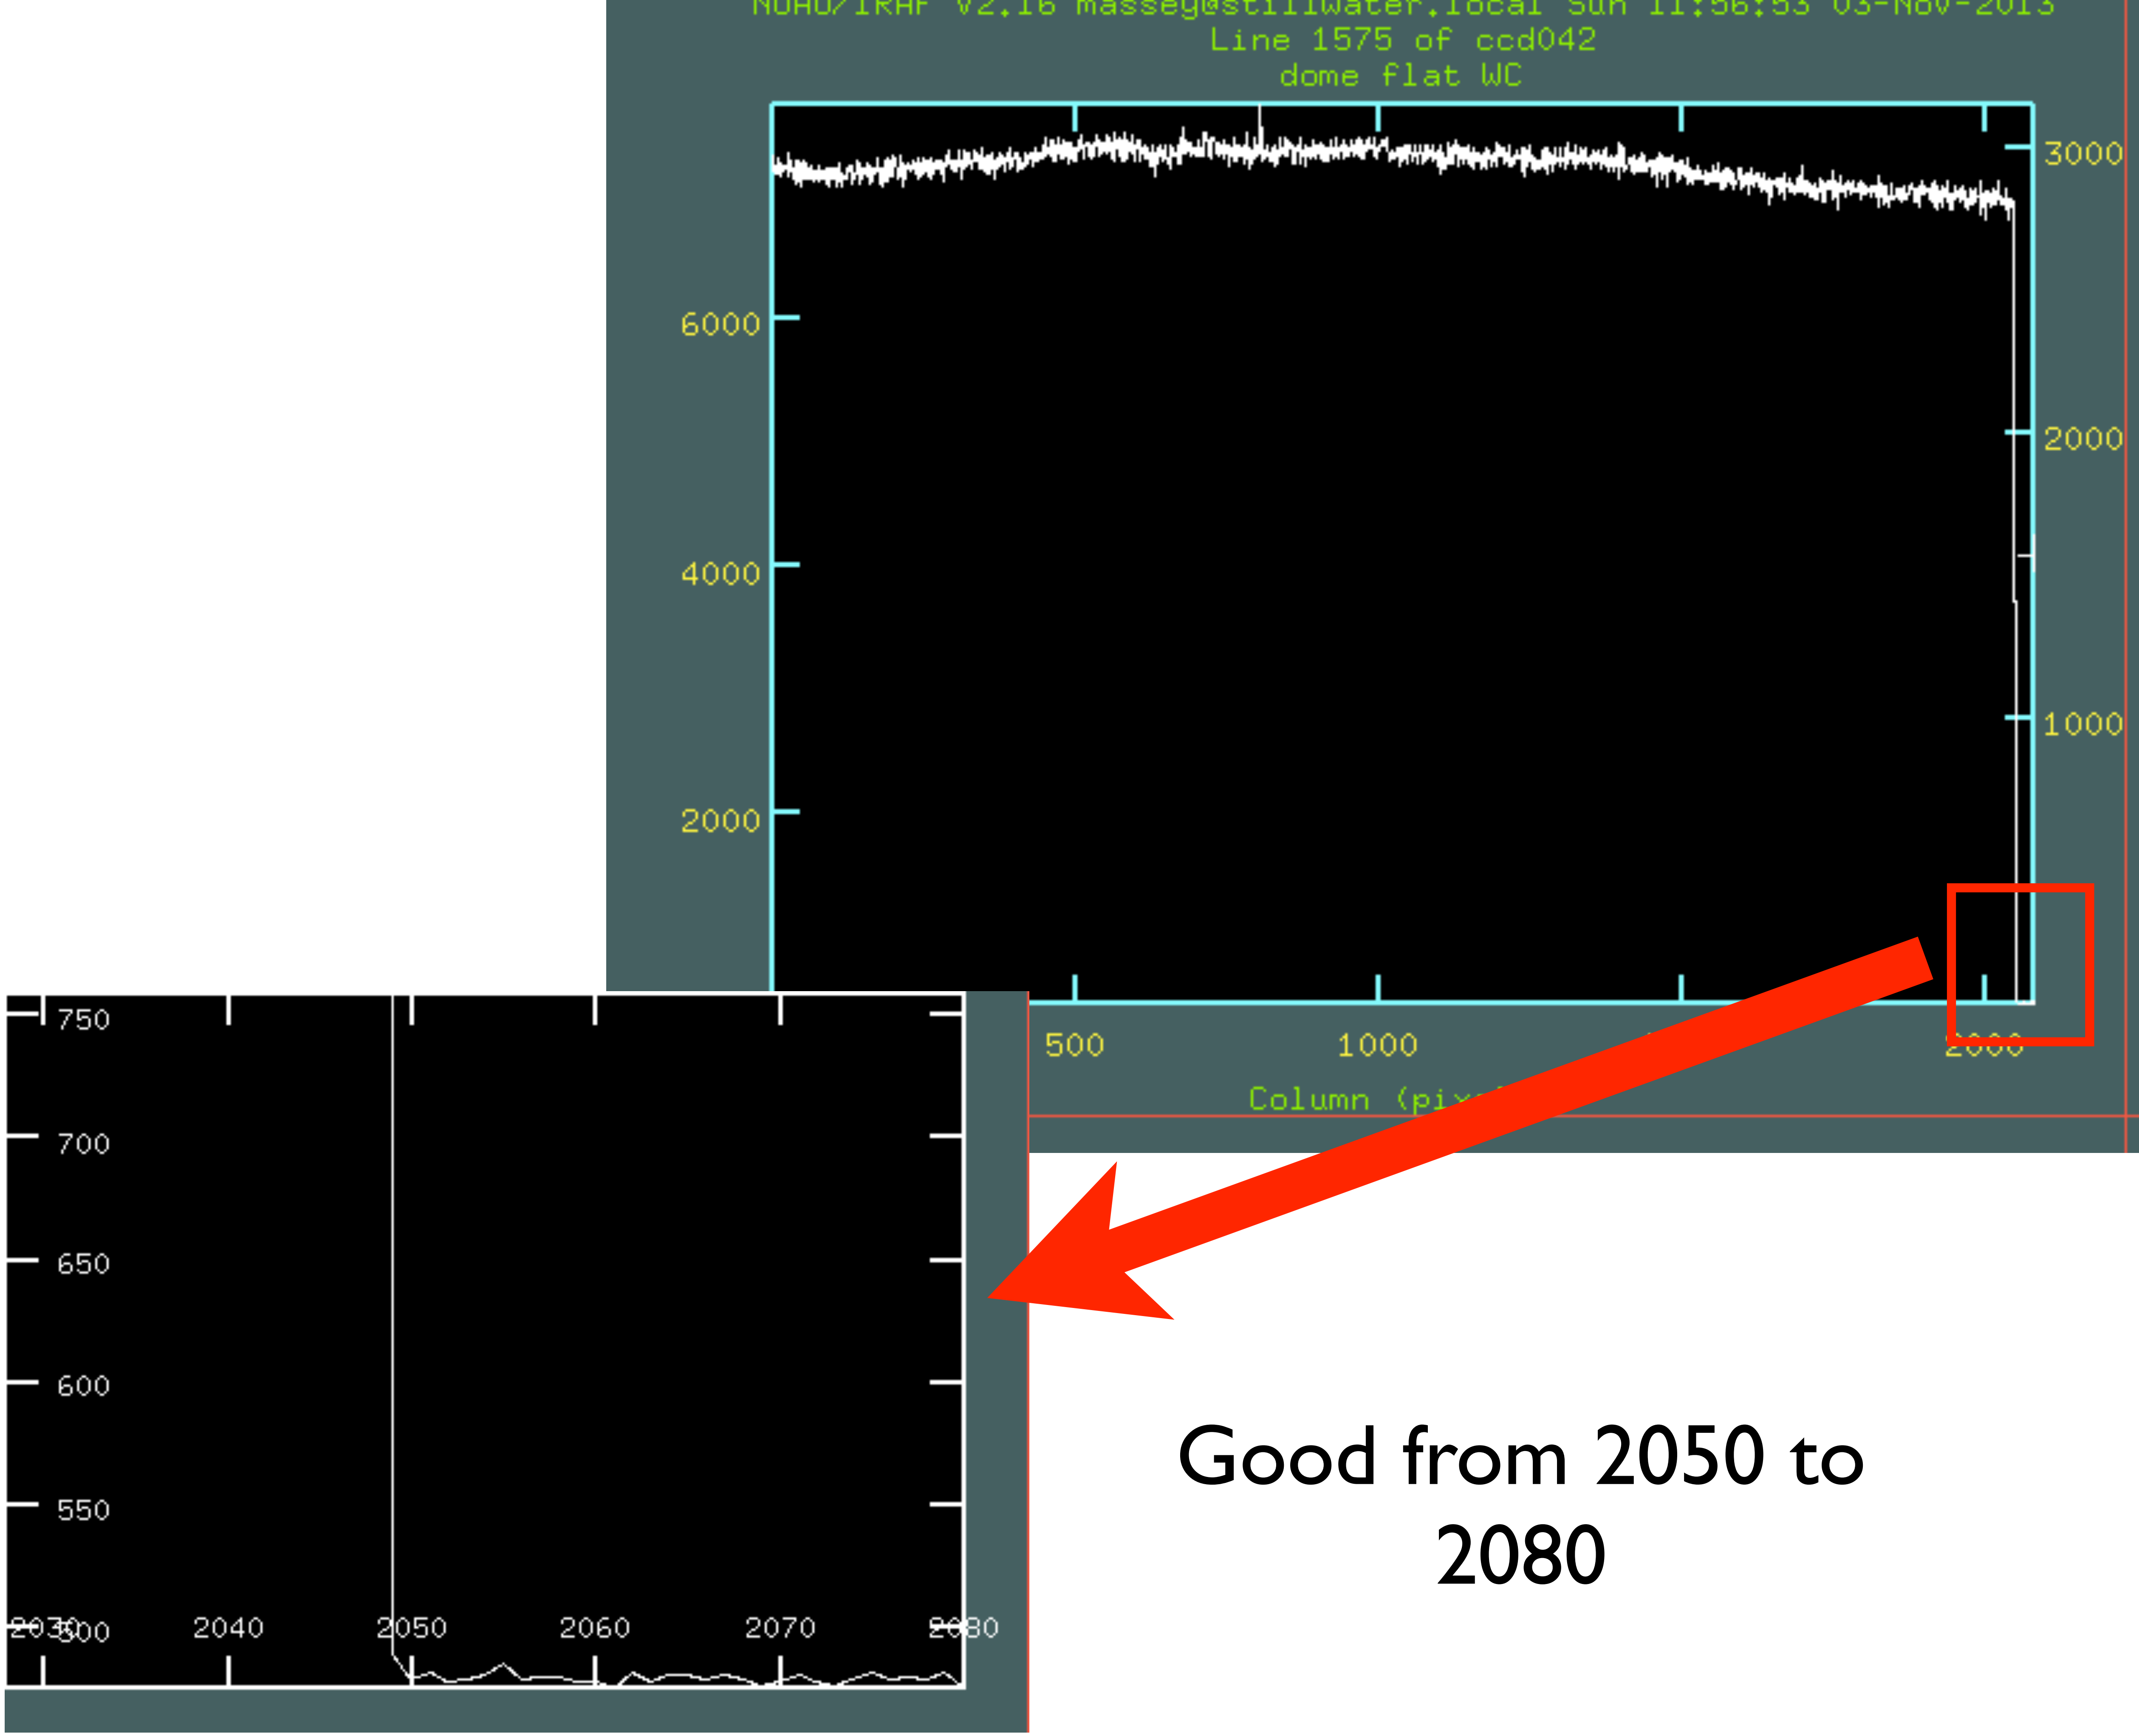

Good from 2050 to 2080

Determining the section you want to save

- Check the edges---usually 1 or 2 bad pixels even if there's no prescan

No problem at left edge!

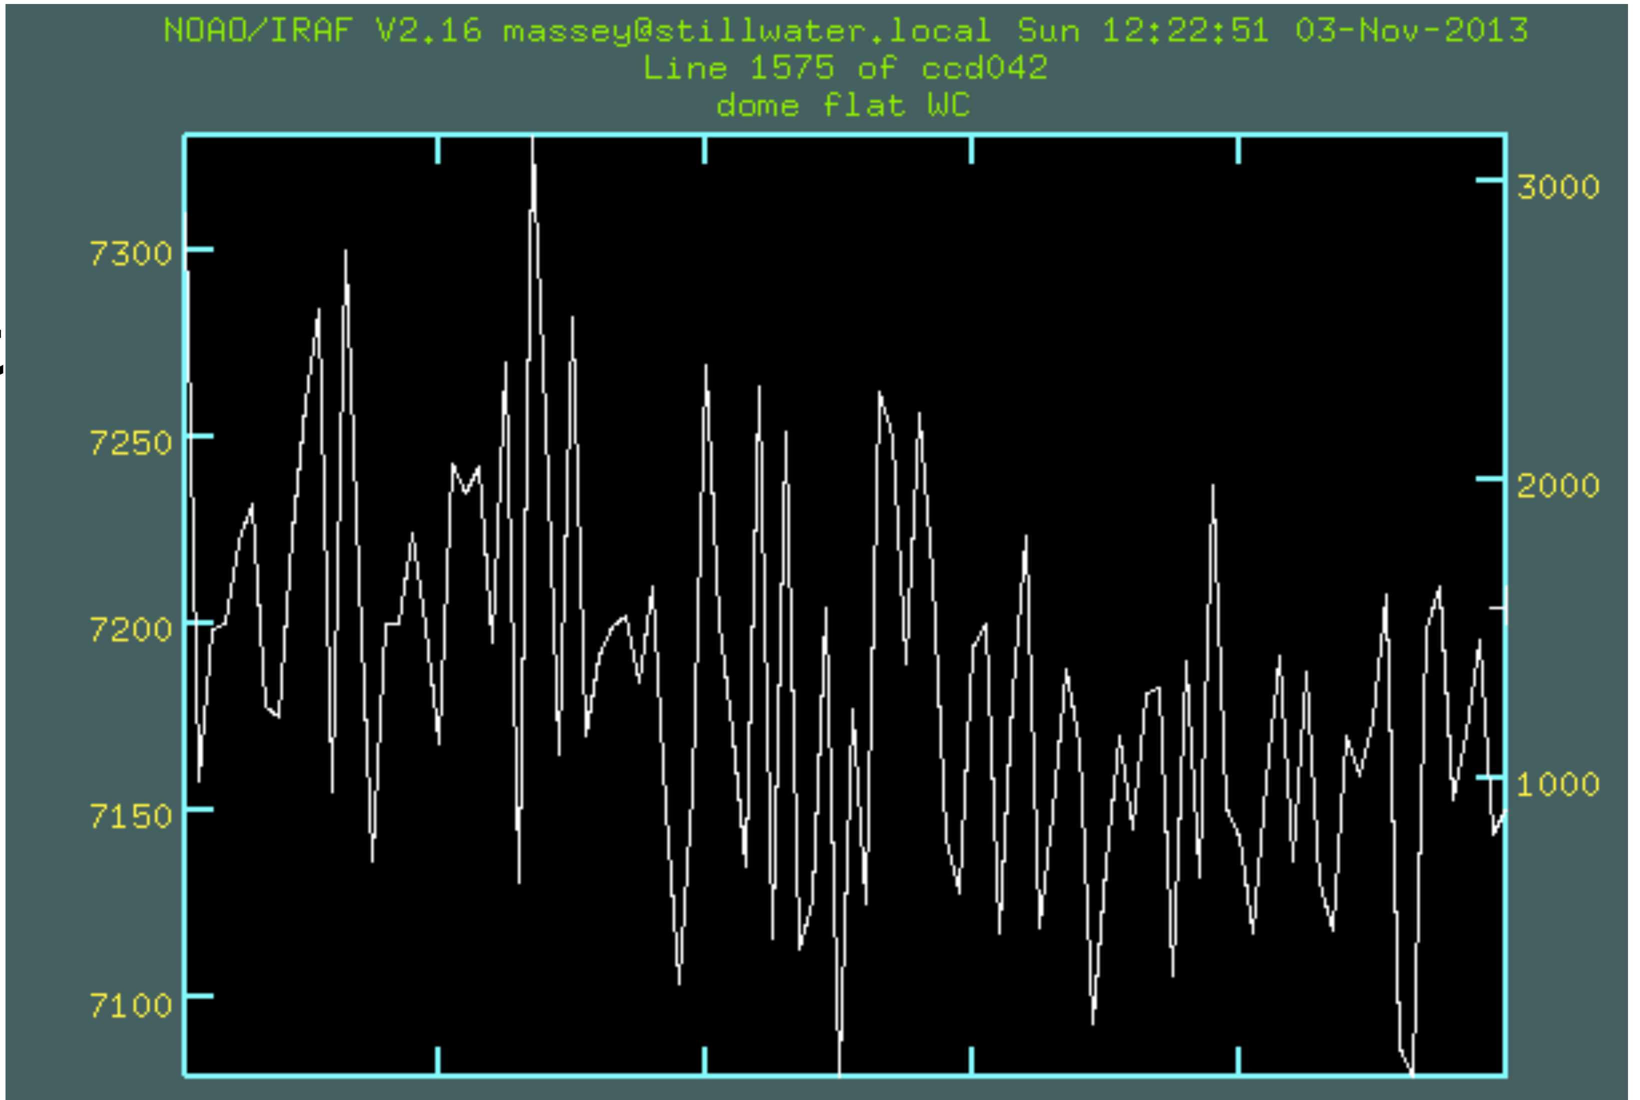

Determining the section you want to save

- Check the edges---usually 1 or 2 bad pixels even if there's no prescan

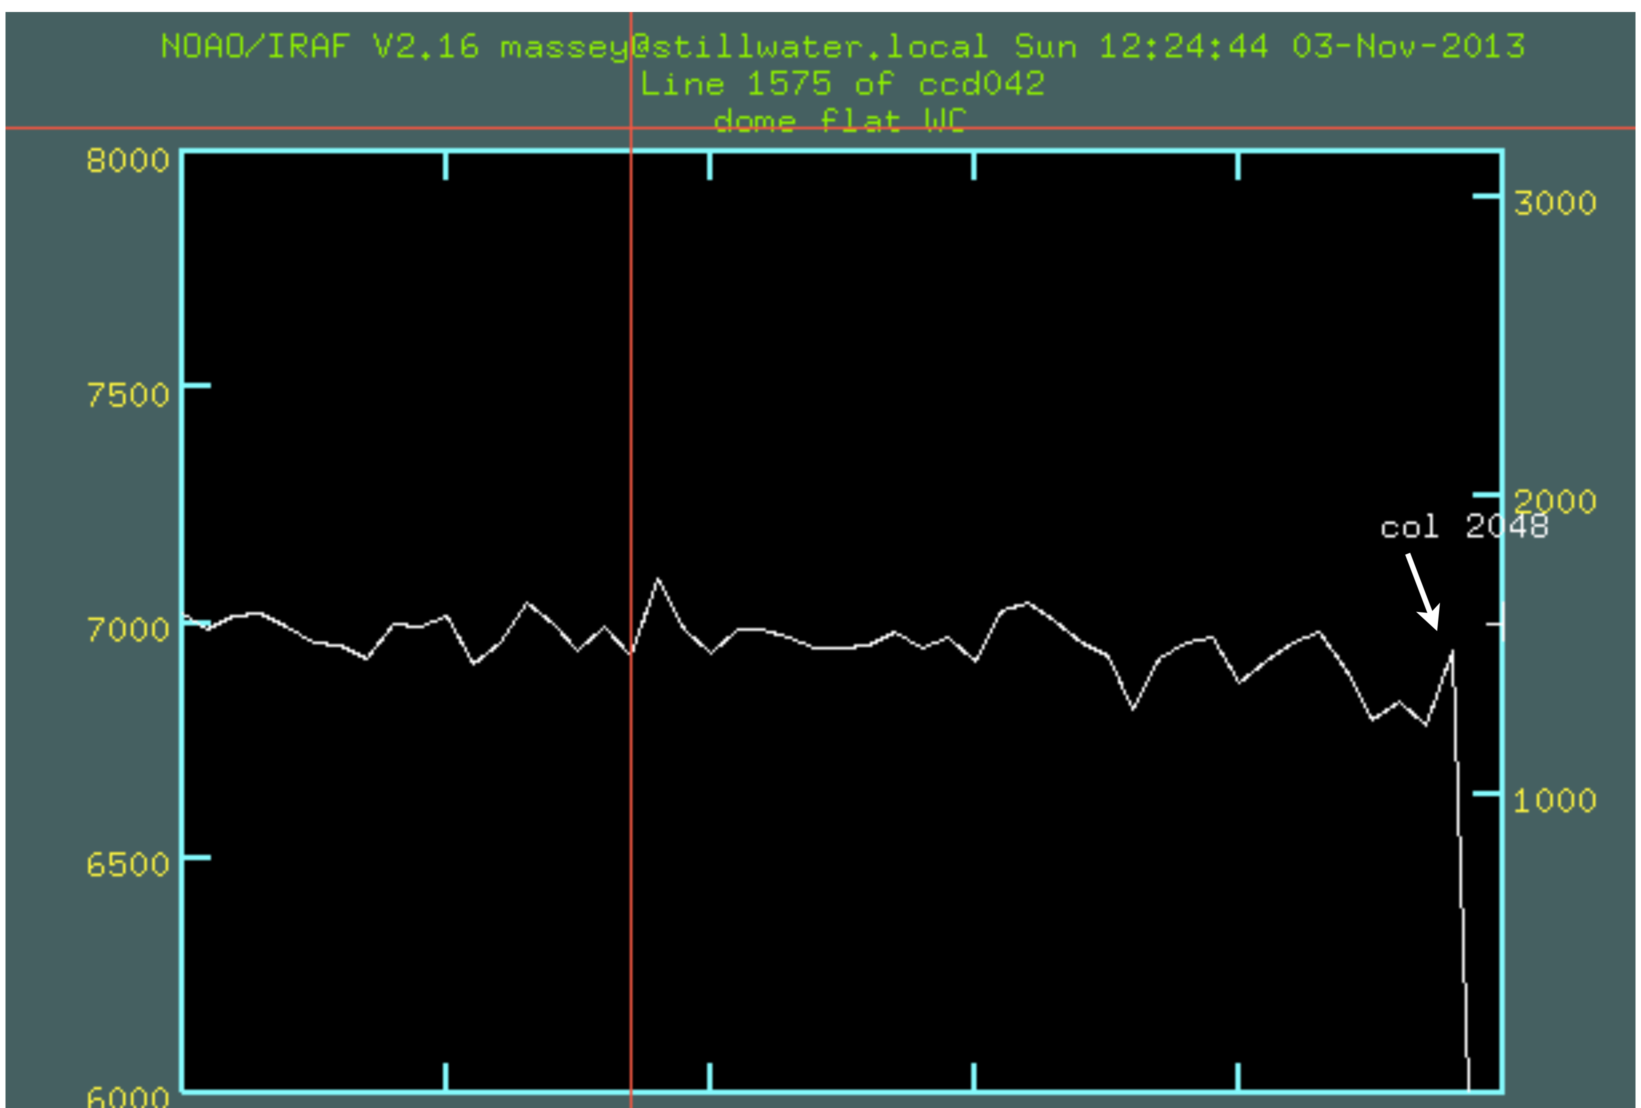

Right edge good through column 2048

Determining the section you want to save

- Better check things out in the y direction.
- Not so perfect. Start at y=6

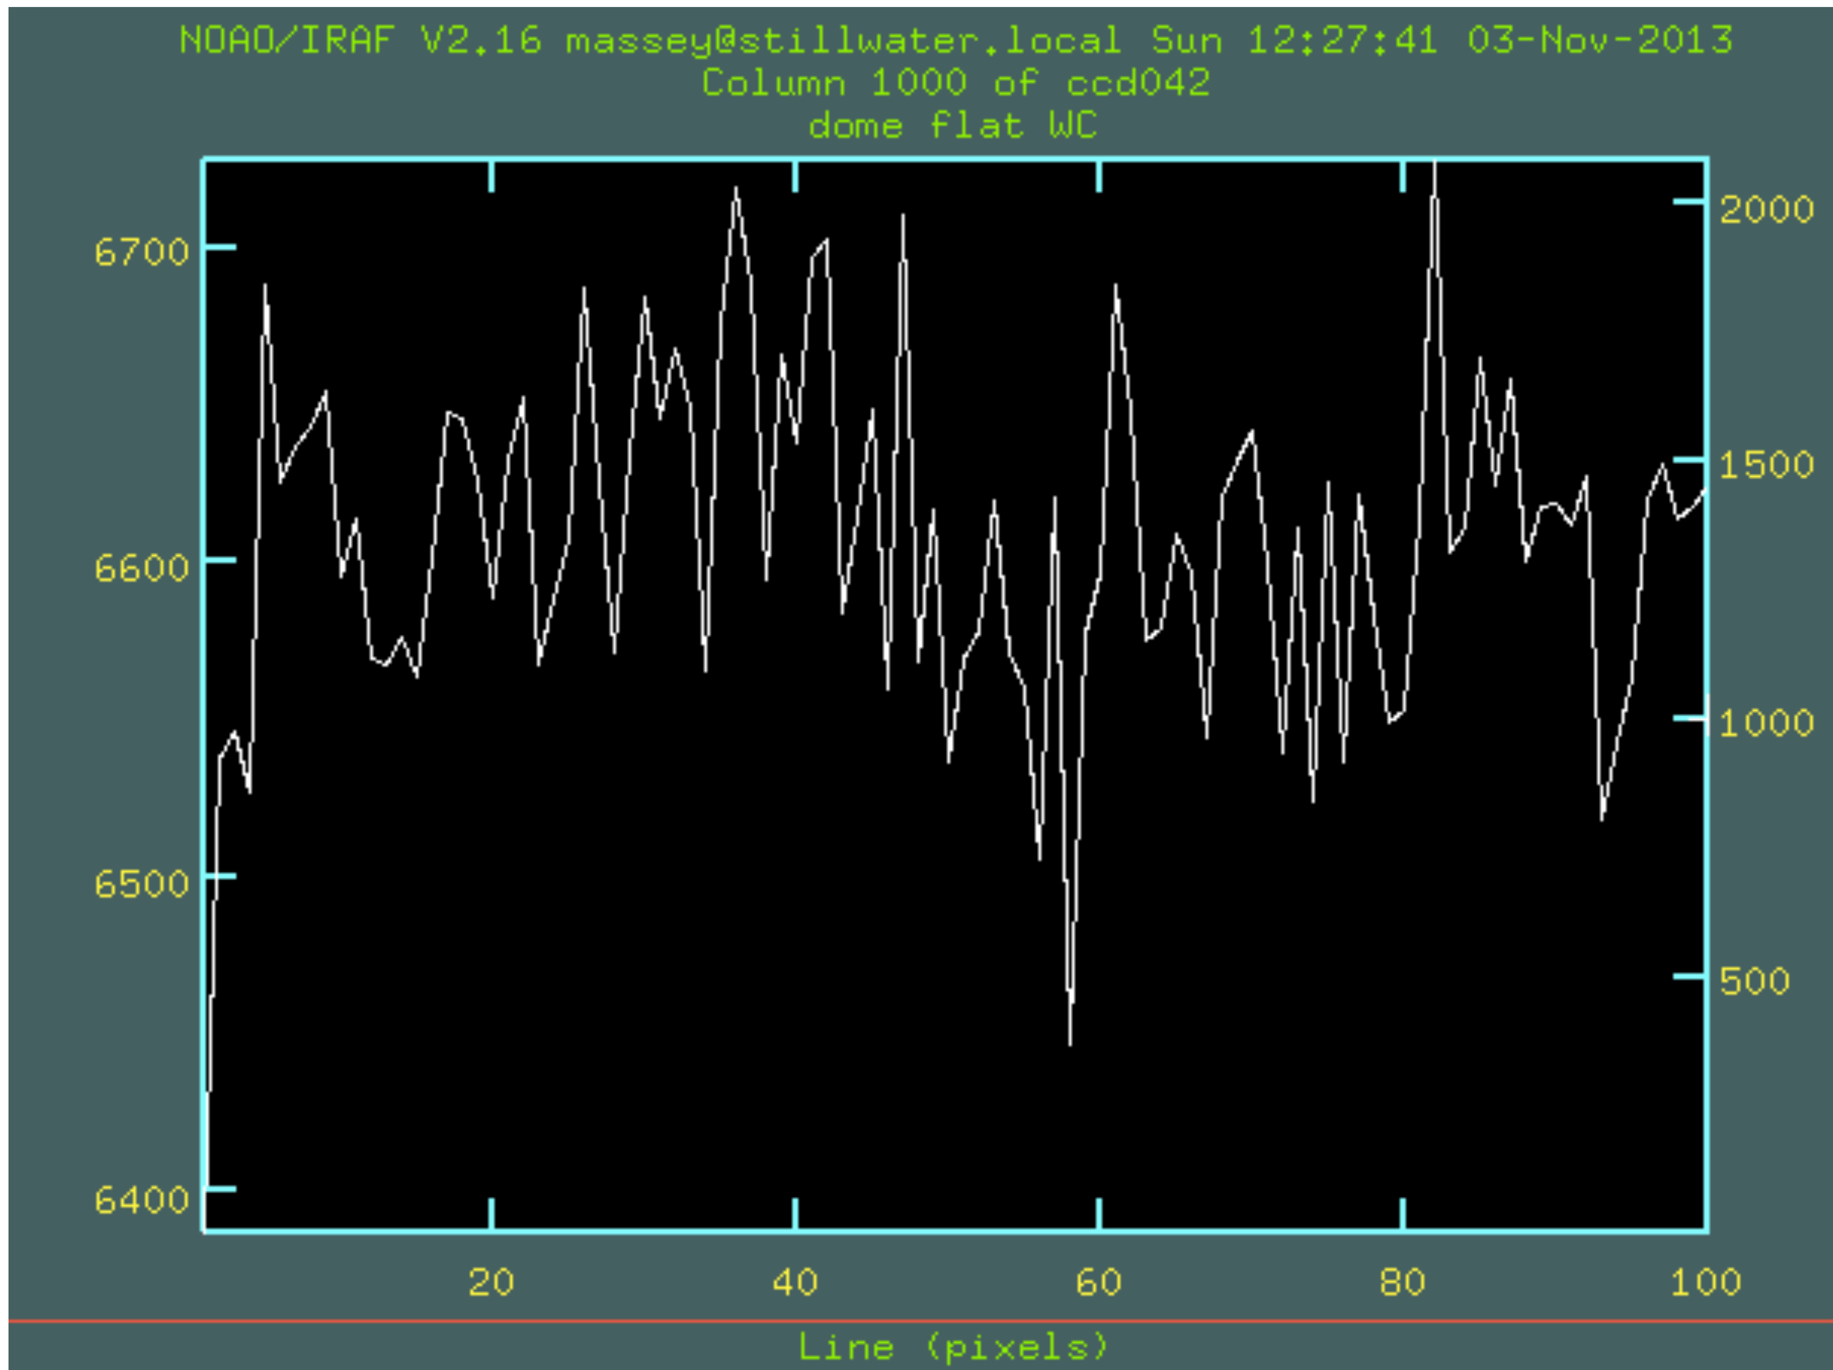

Determining the section you want to save

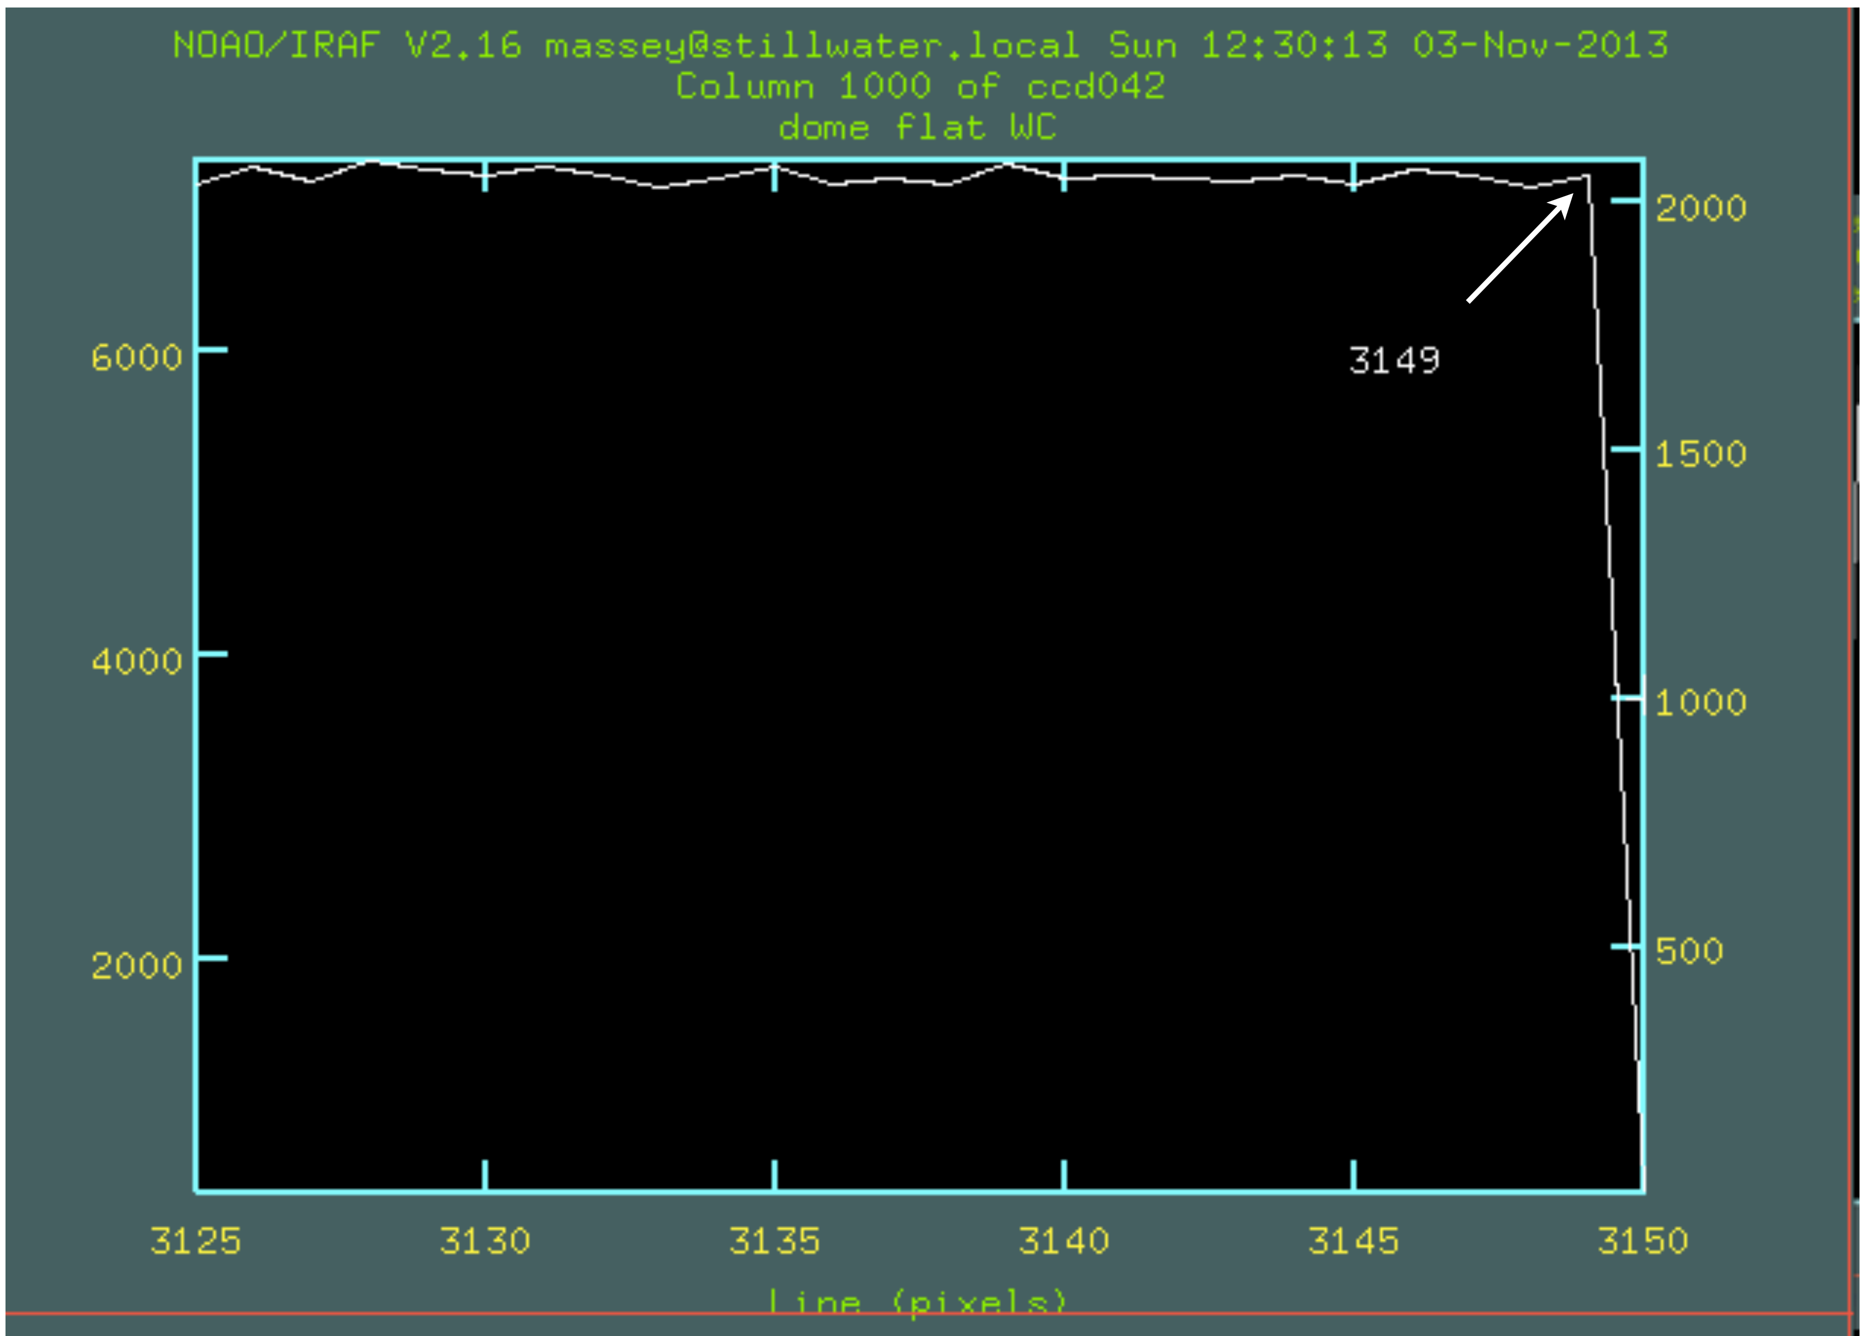

- Better check things out in the y direction.
- Good up until y=3149

Trim and bias sections

trimsec=[1:2048,6:3149]

biassec=[2050:2080,6:3149]

Always a good idea to keep the y values in the bias section the same as in the trim section

Trim and bias sections

So what does this mean?

1) Only the data from x=1 to 2048 and y=6 to 3149 will be saved.

2) The data from x=2050 to 2080 will be averaged line by line, and plotted against y, from 6 to 3149.

Fitting the overscan

Might as well fit this with a constant

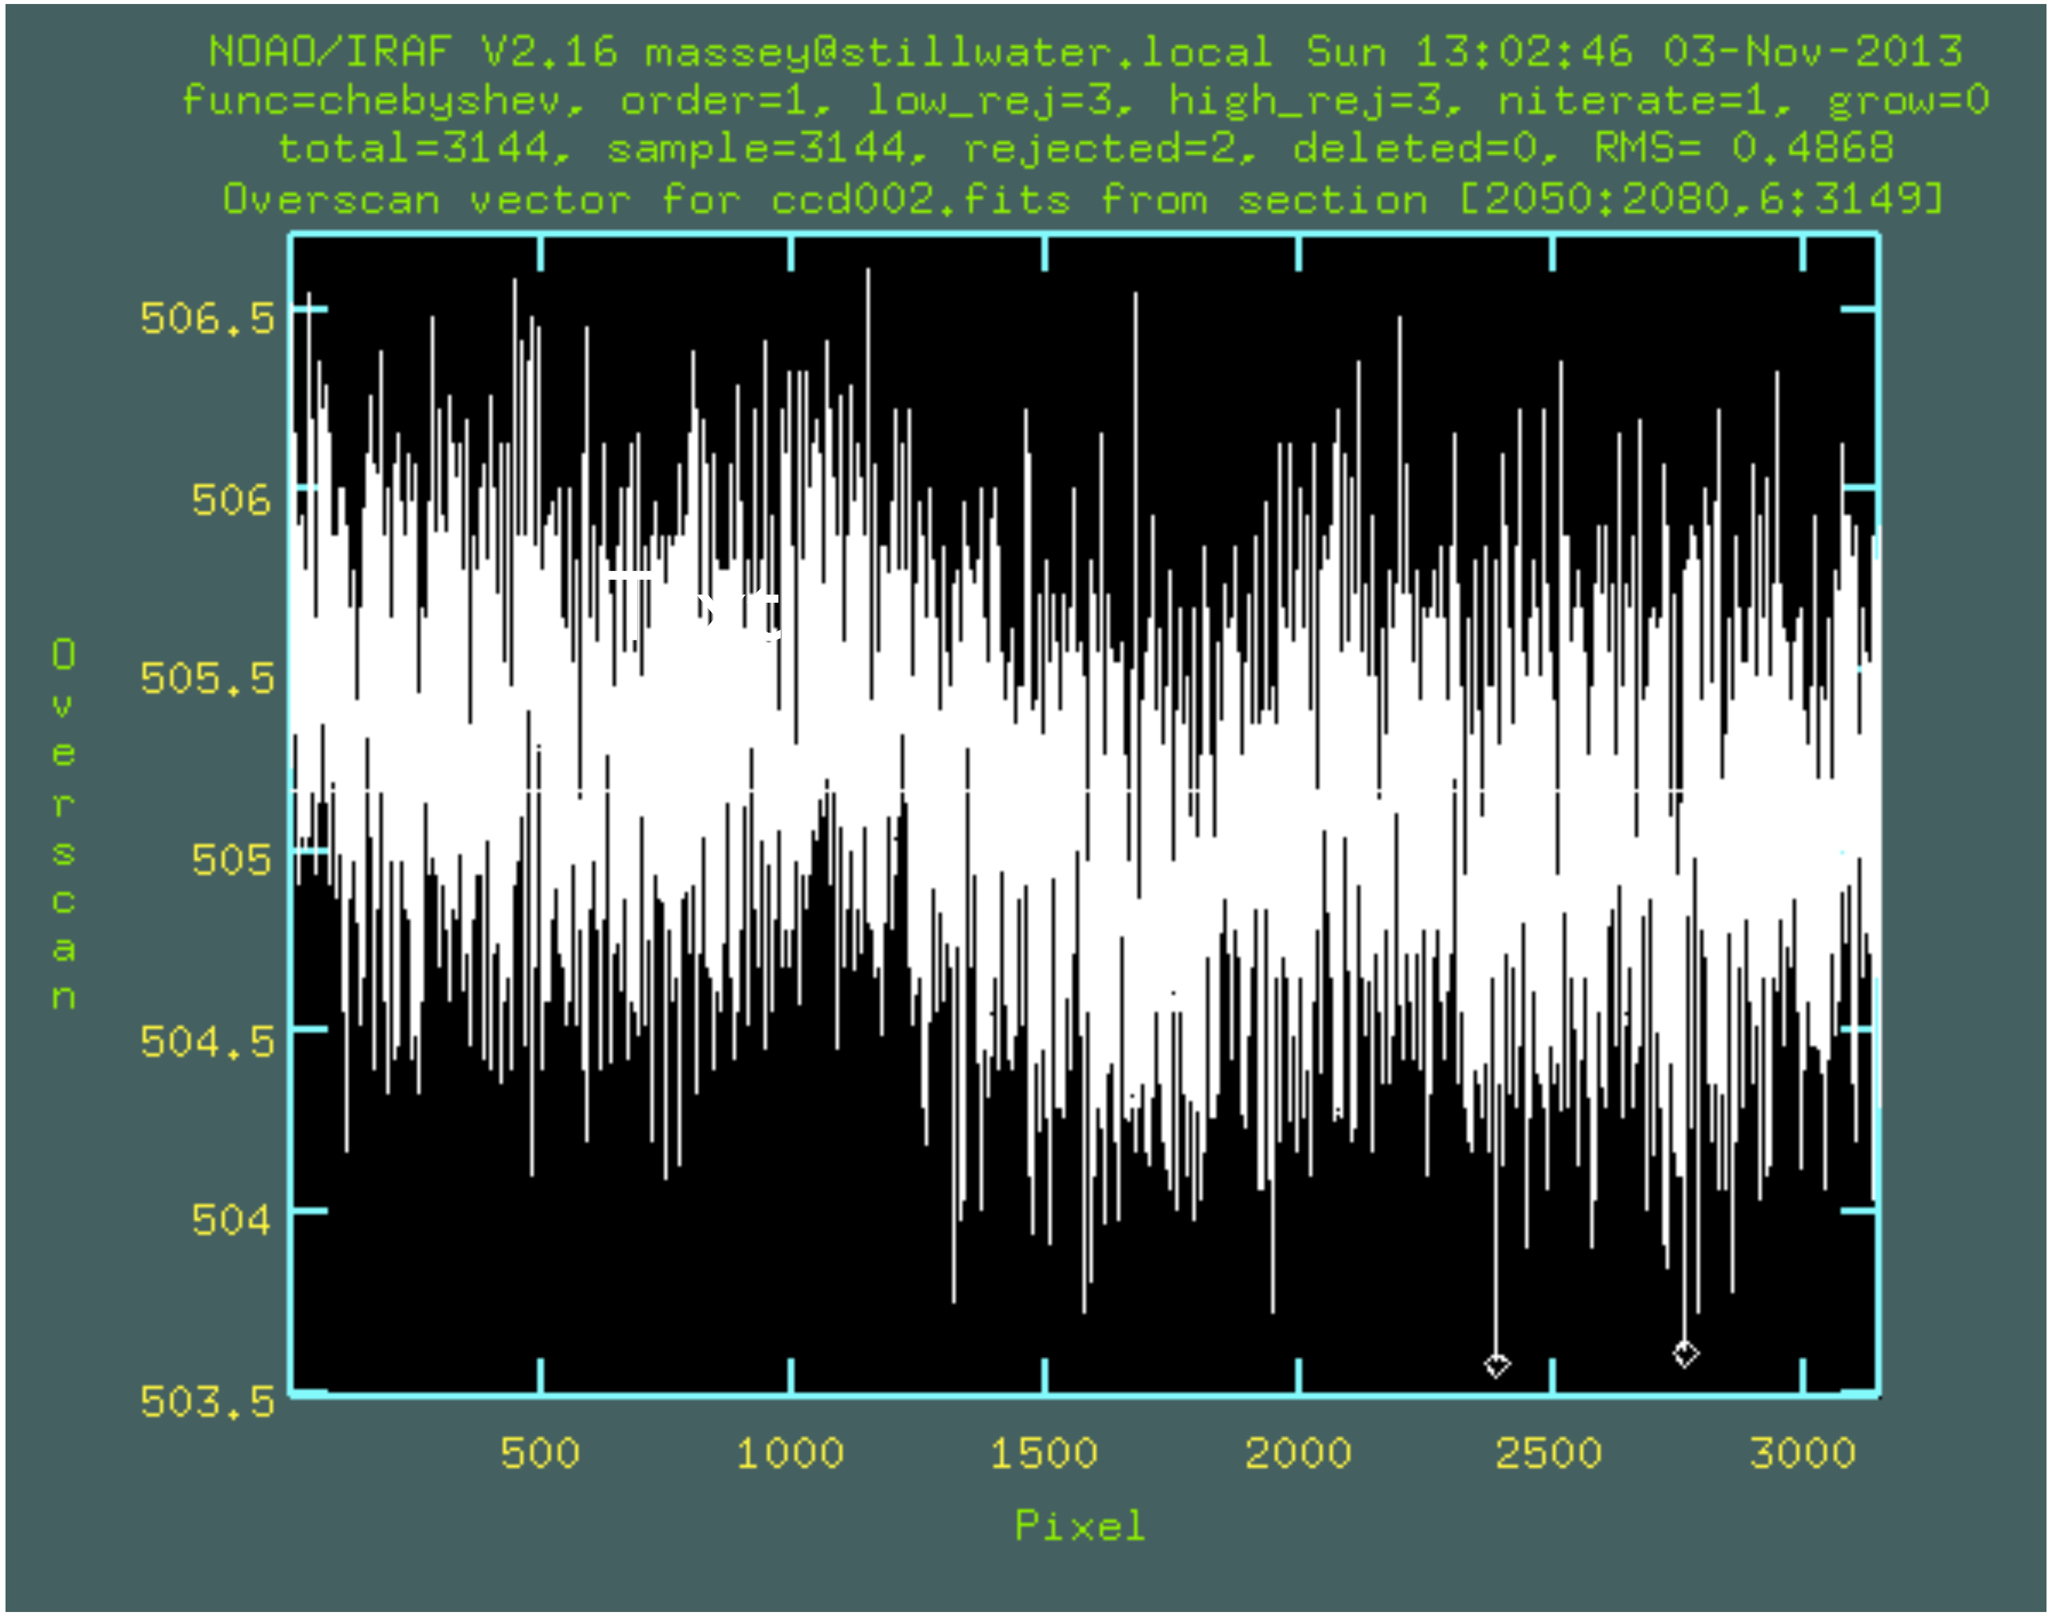

Bias structure?

- After we've removed the overscan from the bias frames (and probably everything else while we're at it) we need to combine the biases and see if there's any structure left over. Sometimes there is, and sometimes there isn't.

Bias structure?

- Question: WHY do we want to combine (average) a bunch of biases? Why not just take a single bias and subtract it?

Bias structure?

If we subtracted a single bias frame from our data, what would happen to the equivalent read-noise? It would

a) double

b) increase by $\sqrt{2}$

c) stay the same

d) go down by a factor of 2

e) Huh?

Bias structure?

If we subtracted a single bias frame from our data, what would happen to the equivalent read-noise? It would

a) double

b) increase by $\sqrt{2}$

c) stay the same

d) go down by a factor of 2

e) Huh?

Bias structure?

Regardless of how many bias frames we average, we are going to increase the read-noise. So, we should subtract an average bias only if we need to.

Bias structure?

If we average 9 bias frames, what is the noise in the combined frame?

a) 9 x greater

b) 3 x greater

c) 3 x less

d) Indeterminable

Bias structure?

If we average 9 bias frames, what is the noise in the combined frame?

a) 9 x greater

b) 3 x greater

c) 3 x less

d) Indeterminable

Bias structure?

Wait?! How can that be? We subtract a frame from another frame and the noise goes UP by the sqrt (2). But if we average two frames the noise goes down by the sqrt (2)? That doesn't seem fair!

Sidebar: Remember How errors add!

Function	Variance
$f = aA$	$\sigma_f^2 = a^2\sigma_A^2$
$f = aA \pm bB$	$\sigma_f^2 = a^2\sigma_A^2 + b^2\sigma_B^2$
$f = AB$	$\left(\frac{\sigma_f}{f}\right)^2 \approx \left(\frac{\sigma_A}{A}\right)^2 + \left(\frac{\sigma_B}{B}\right)^2$
$f = \frac{A}{B}$	$\left(\frac{\sigma_f}{f}\right)^2 \approx \left(\frac{\sigma_A}{A}\right)^2 + \left(\frac{\sigma_B}{B}\right)^2$

Sidebar: Remember How errors add!

Function	Variance
$f = aA$	$\sigma_f^2 = a^2\sigma_A^2$
$f = aA \pm bB$	$\sigma_f^2 = a^2\sigma_A^2 + b^2\sigma_B^2$
$f = AB$	$\left(\frac{\sigma_f}{f}\right)^2 \approx \left(\frac{\sigma_A}{A}\right)^2 + \left(\frac{\sigma_B}{B}\right)^2$
$f = \frac{A}{B}$	$\left(\frac{\sigma_f}{f}\right)^2 \approx \left(\frac{\sigma_A}{A}\right)^2 + \left(\frac{\sigma_B}{B}\right)^2$

Sidebar: How errors add!

We're going to AVERAGE 9 frames. That means we add all N of them together and then divide by N. In that case final_σ^2=Nσ^2/N^2 . In this case, the N is 9:

final_σ^2 = 9σ^2/ 9^2

or final_σ= σ/ 3

In other words, if you average N things, the errors will decrease by $N^{0.5}$. But if you add N things, the errors will INCREASE by $N^{0.5}$.

Bias structure?

What's a reasonable number of bias frames? Well, 9 is pretty good. What will the final noise be?

σ^2_final = $\sigma^2 + (\sigma/3)^2 = \sigma^2 + \sigma^2 / 9 = (10/9)\ \sigma^2$

σ_final = $(10/9)^{0.5}\sigma = 1.05\sigma$

So, 5% isn't so bad.

Bias structure?

- Okay, okay: we combine 10 biases (Nidia must have won that discussion) and we take a look…

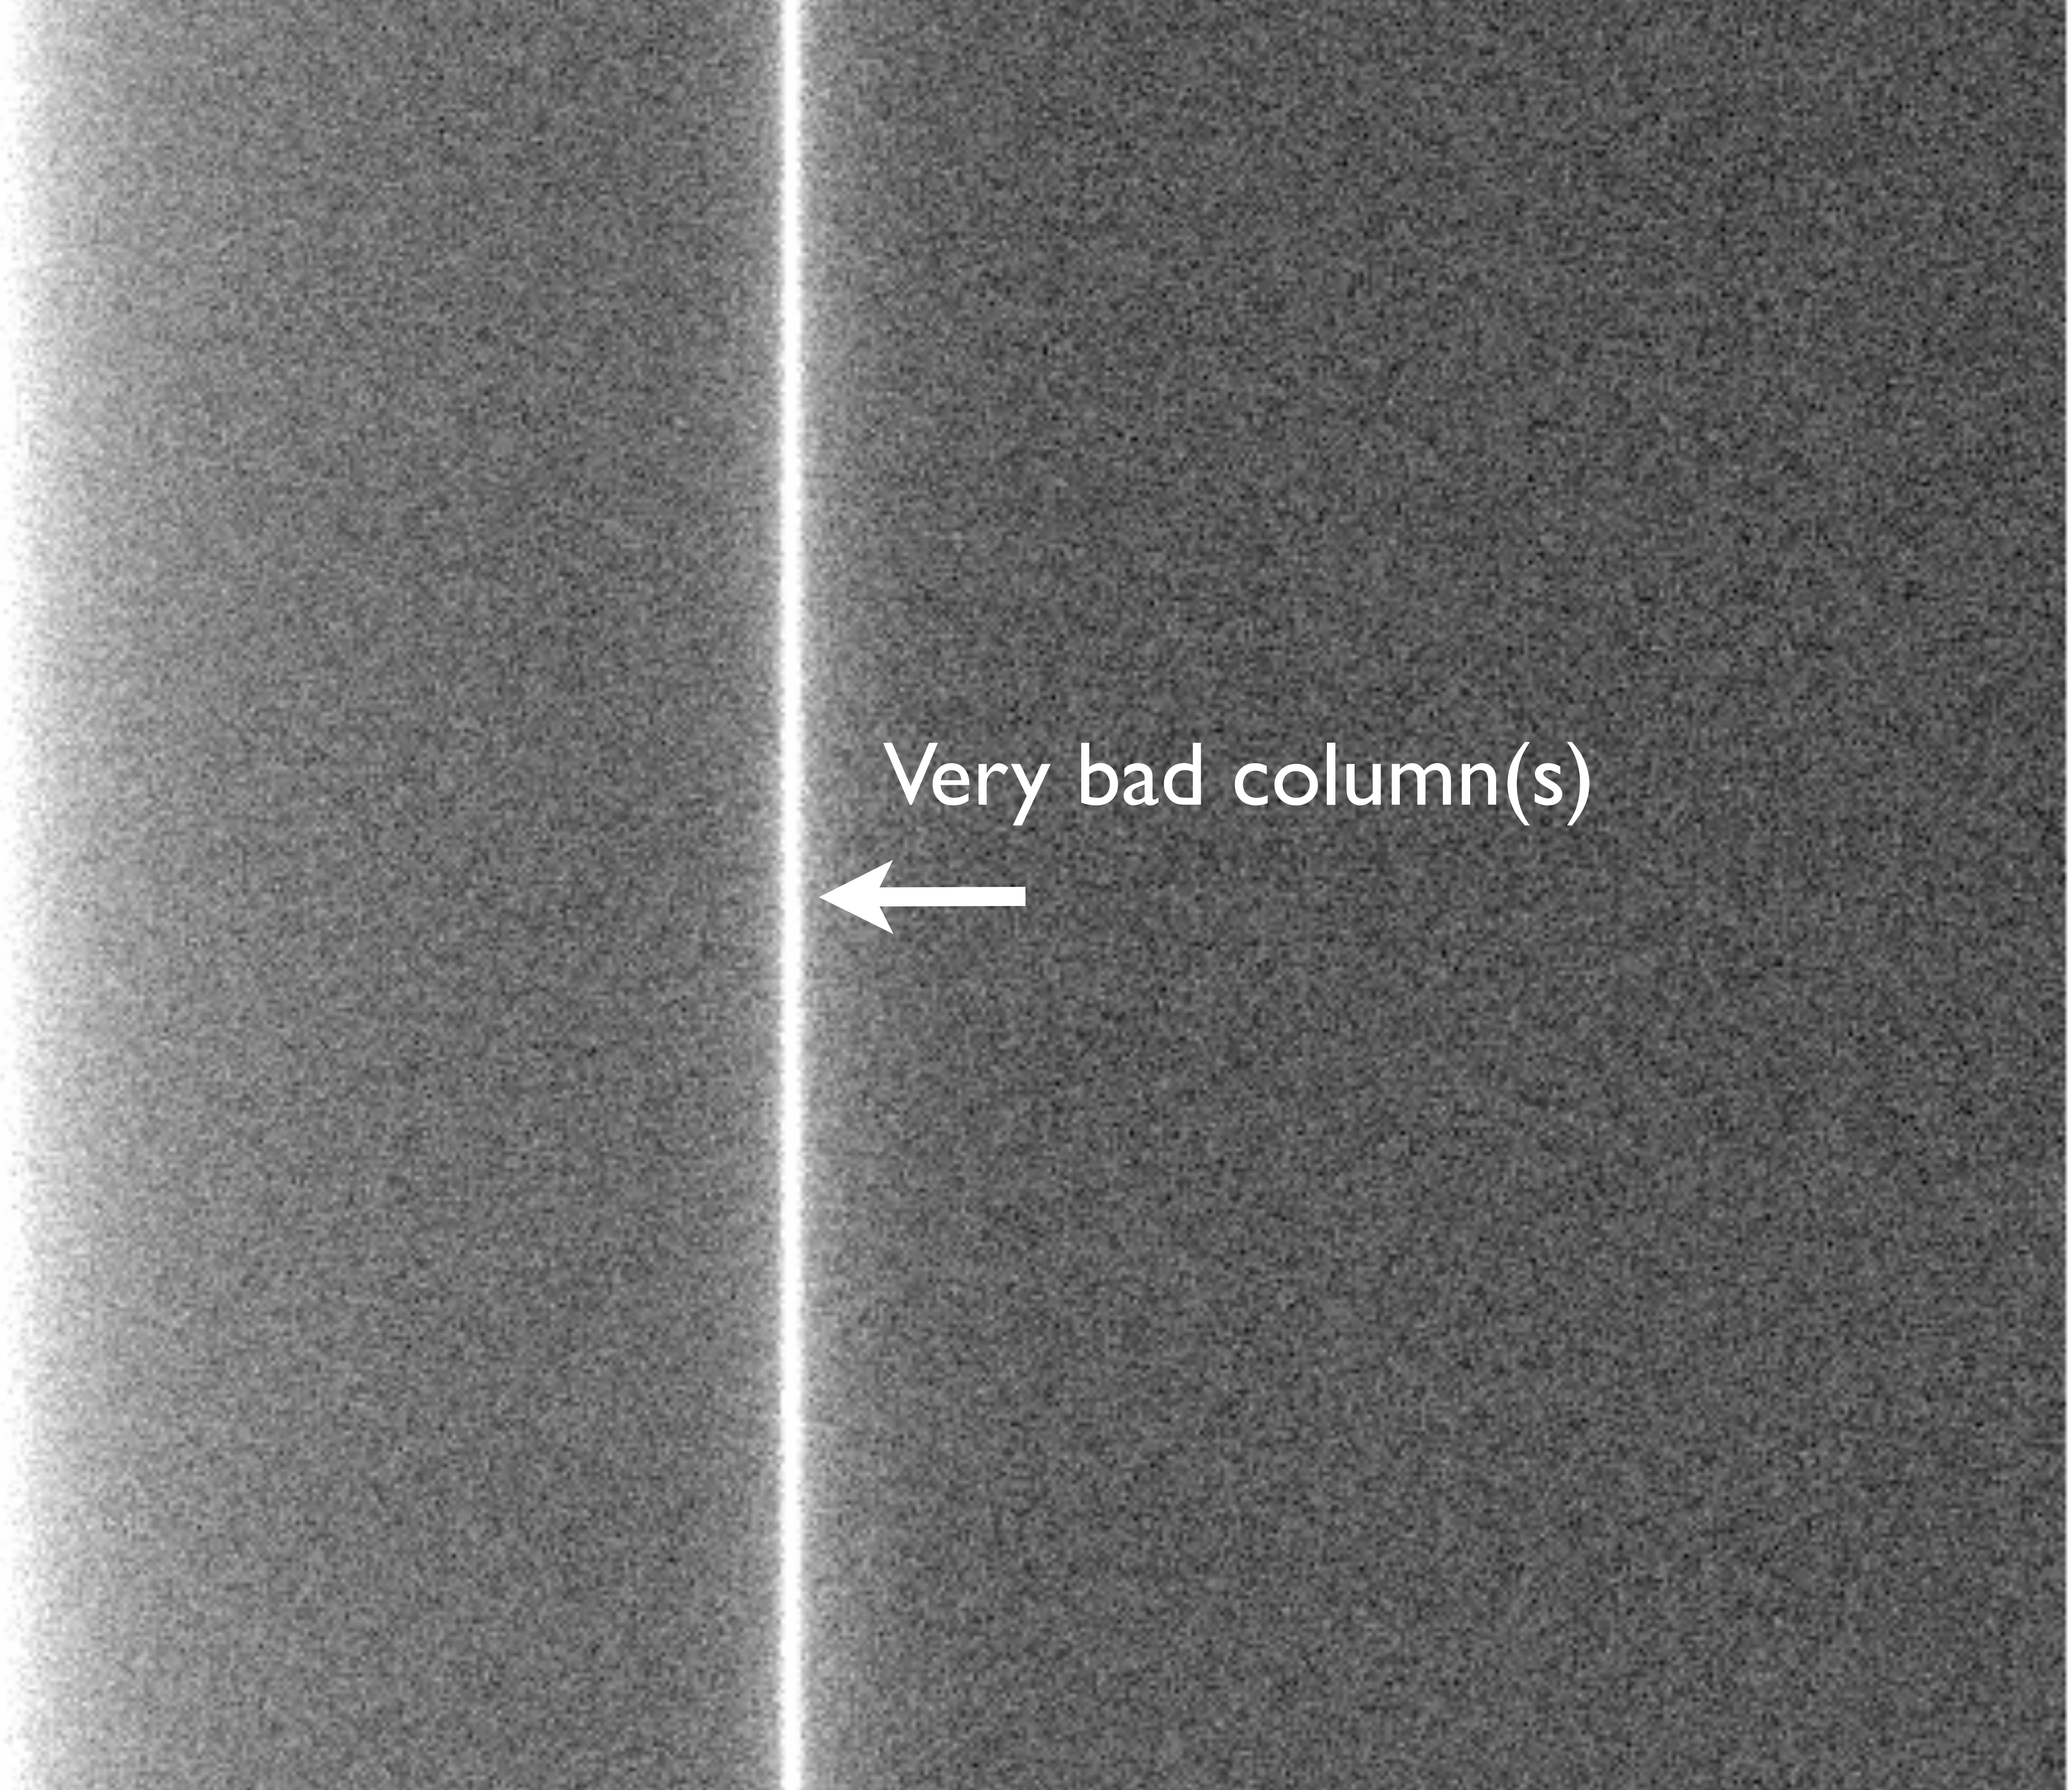
Very bad column(s)

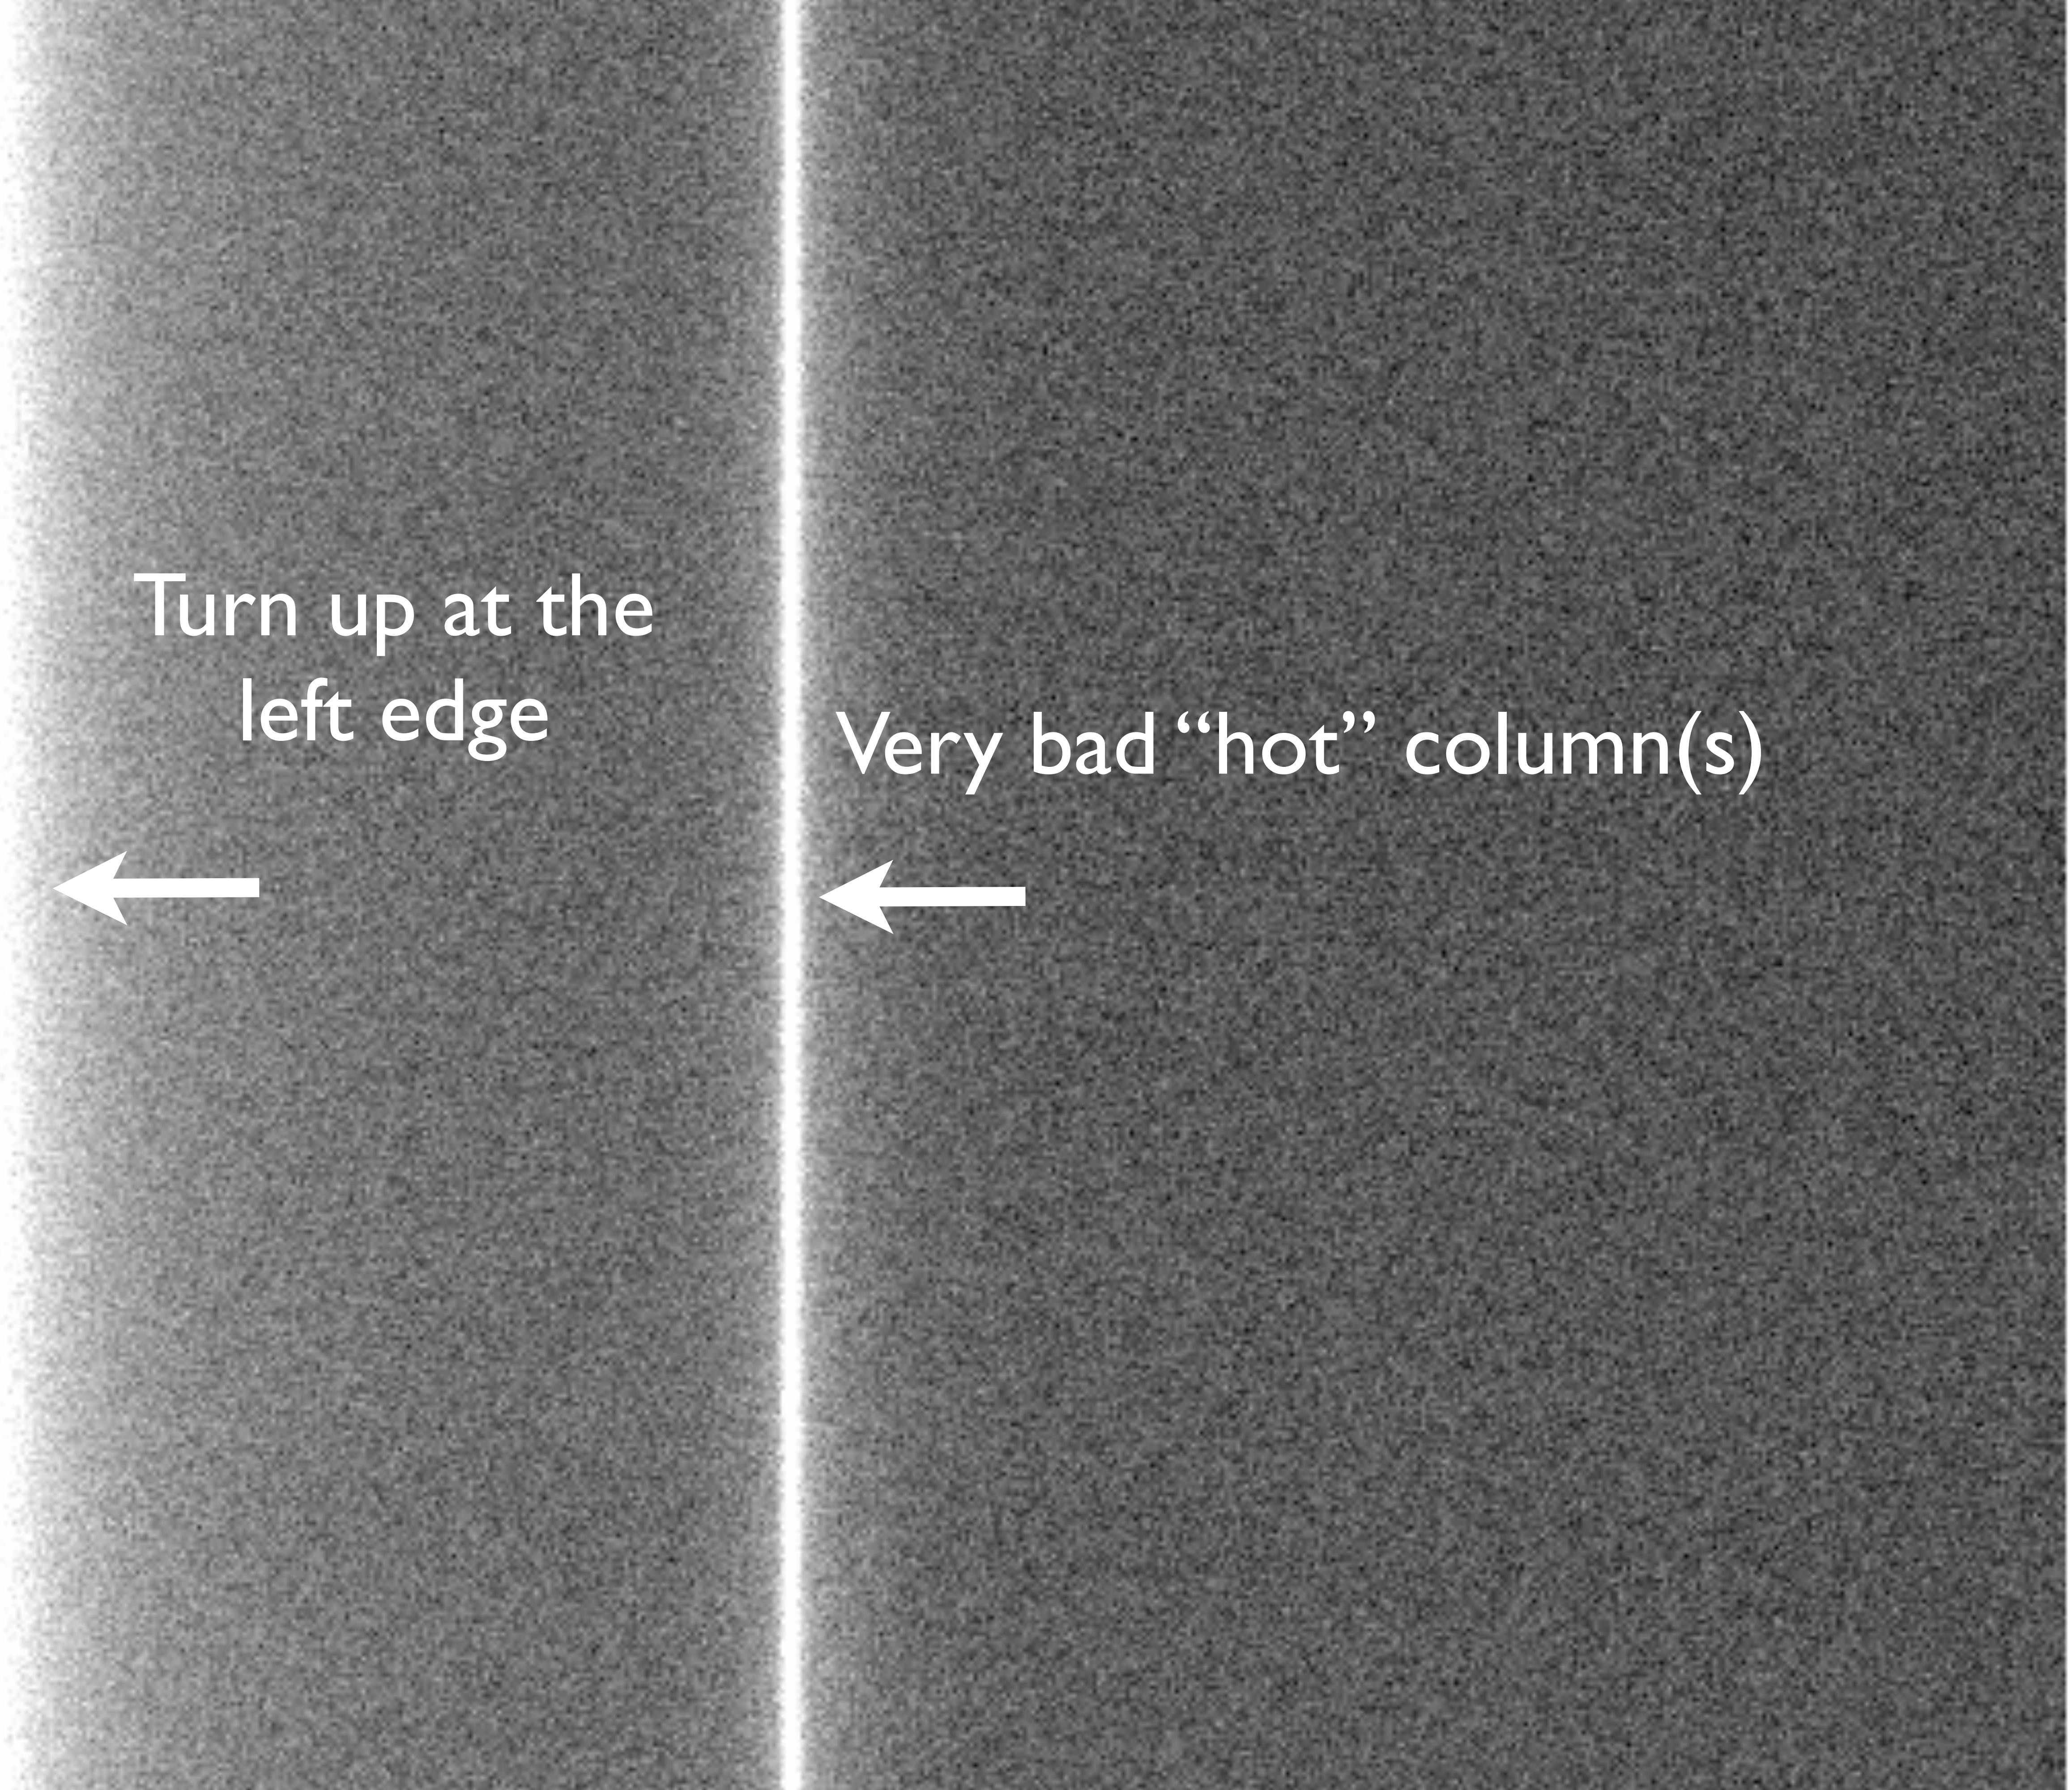
Turn up at the left edge
Very bad "hot" column(s)

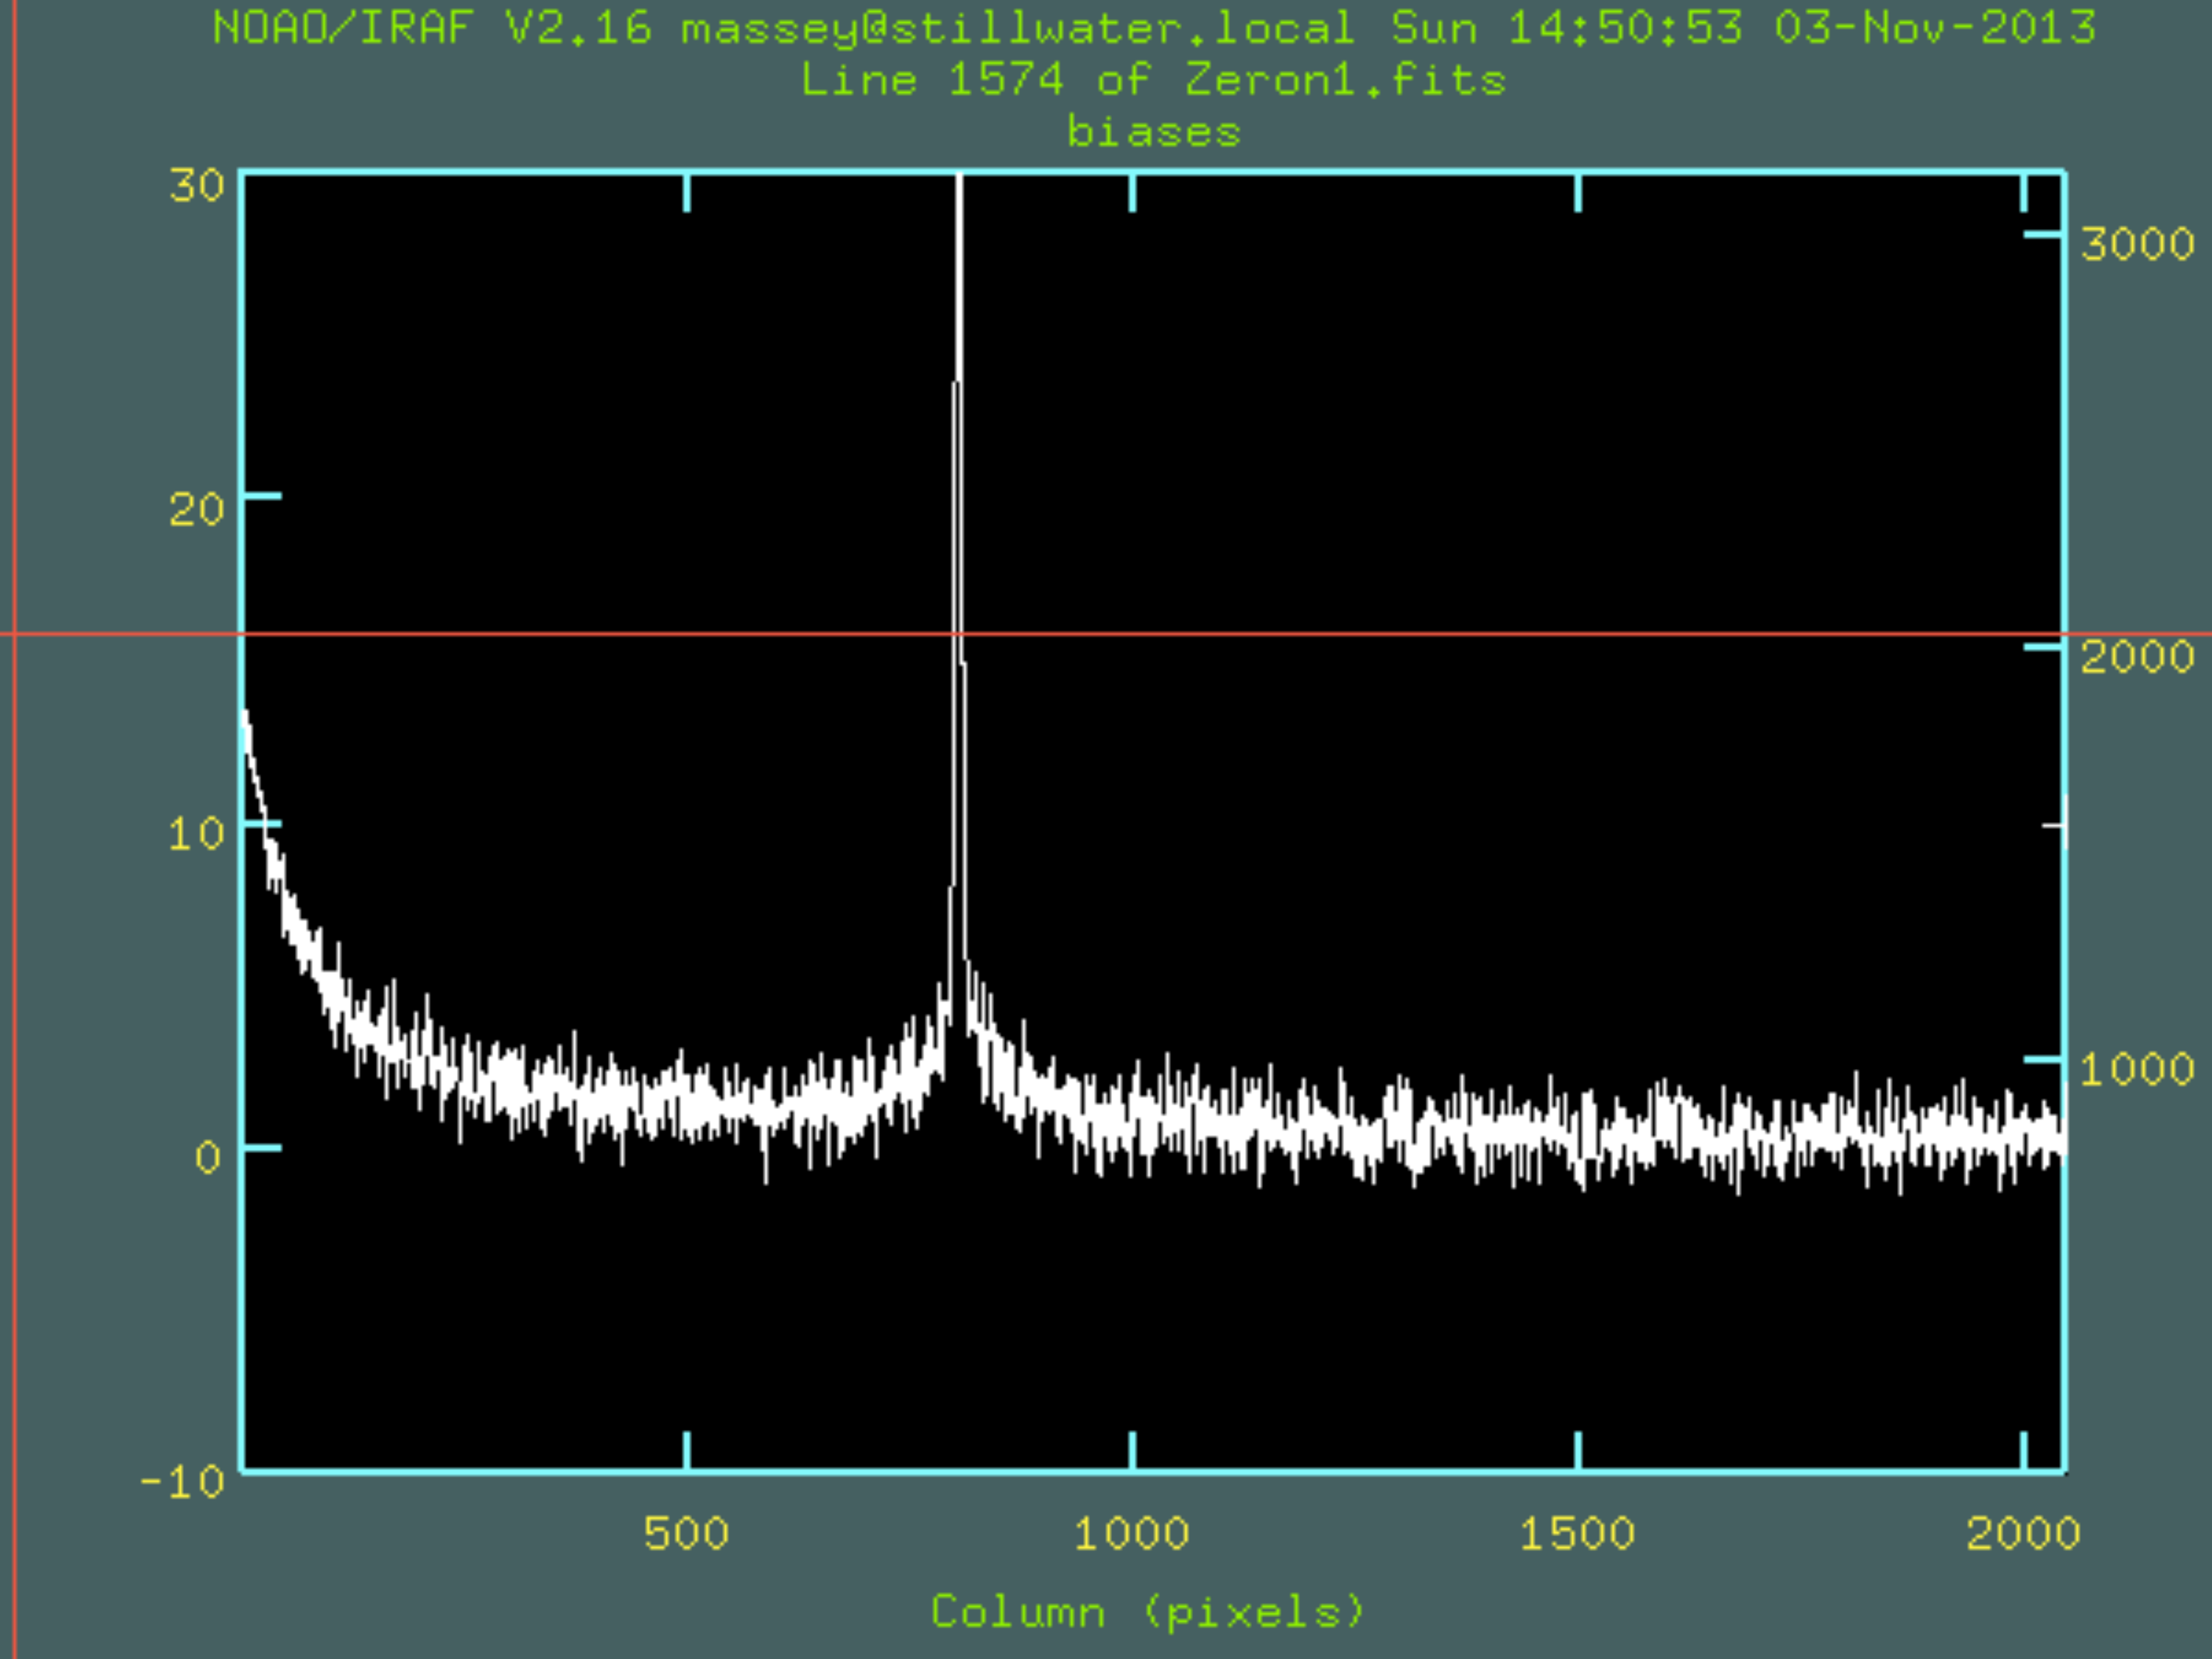
NOAO/IRAF V2.16 massey@stillwater.local Sun 14:50:53 03-Nov-2013
Line 1574 of Zeron1.fits
biases
30
20
10
0
-10
3000
2000
1000
500
1000
1500
2000
Column (pixels)

So, yes, there is bias structure!

- Subtracting it is a Very Good Thing

Subtracting bias structure

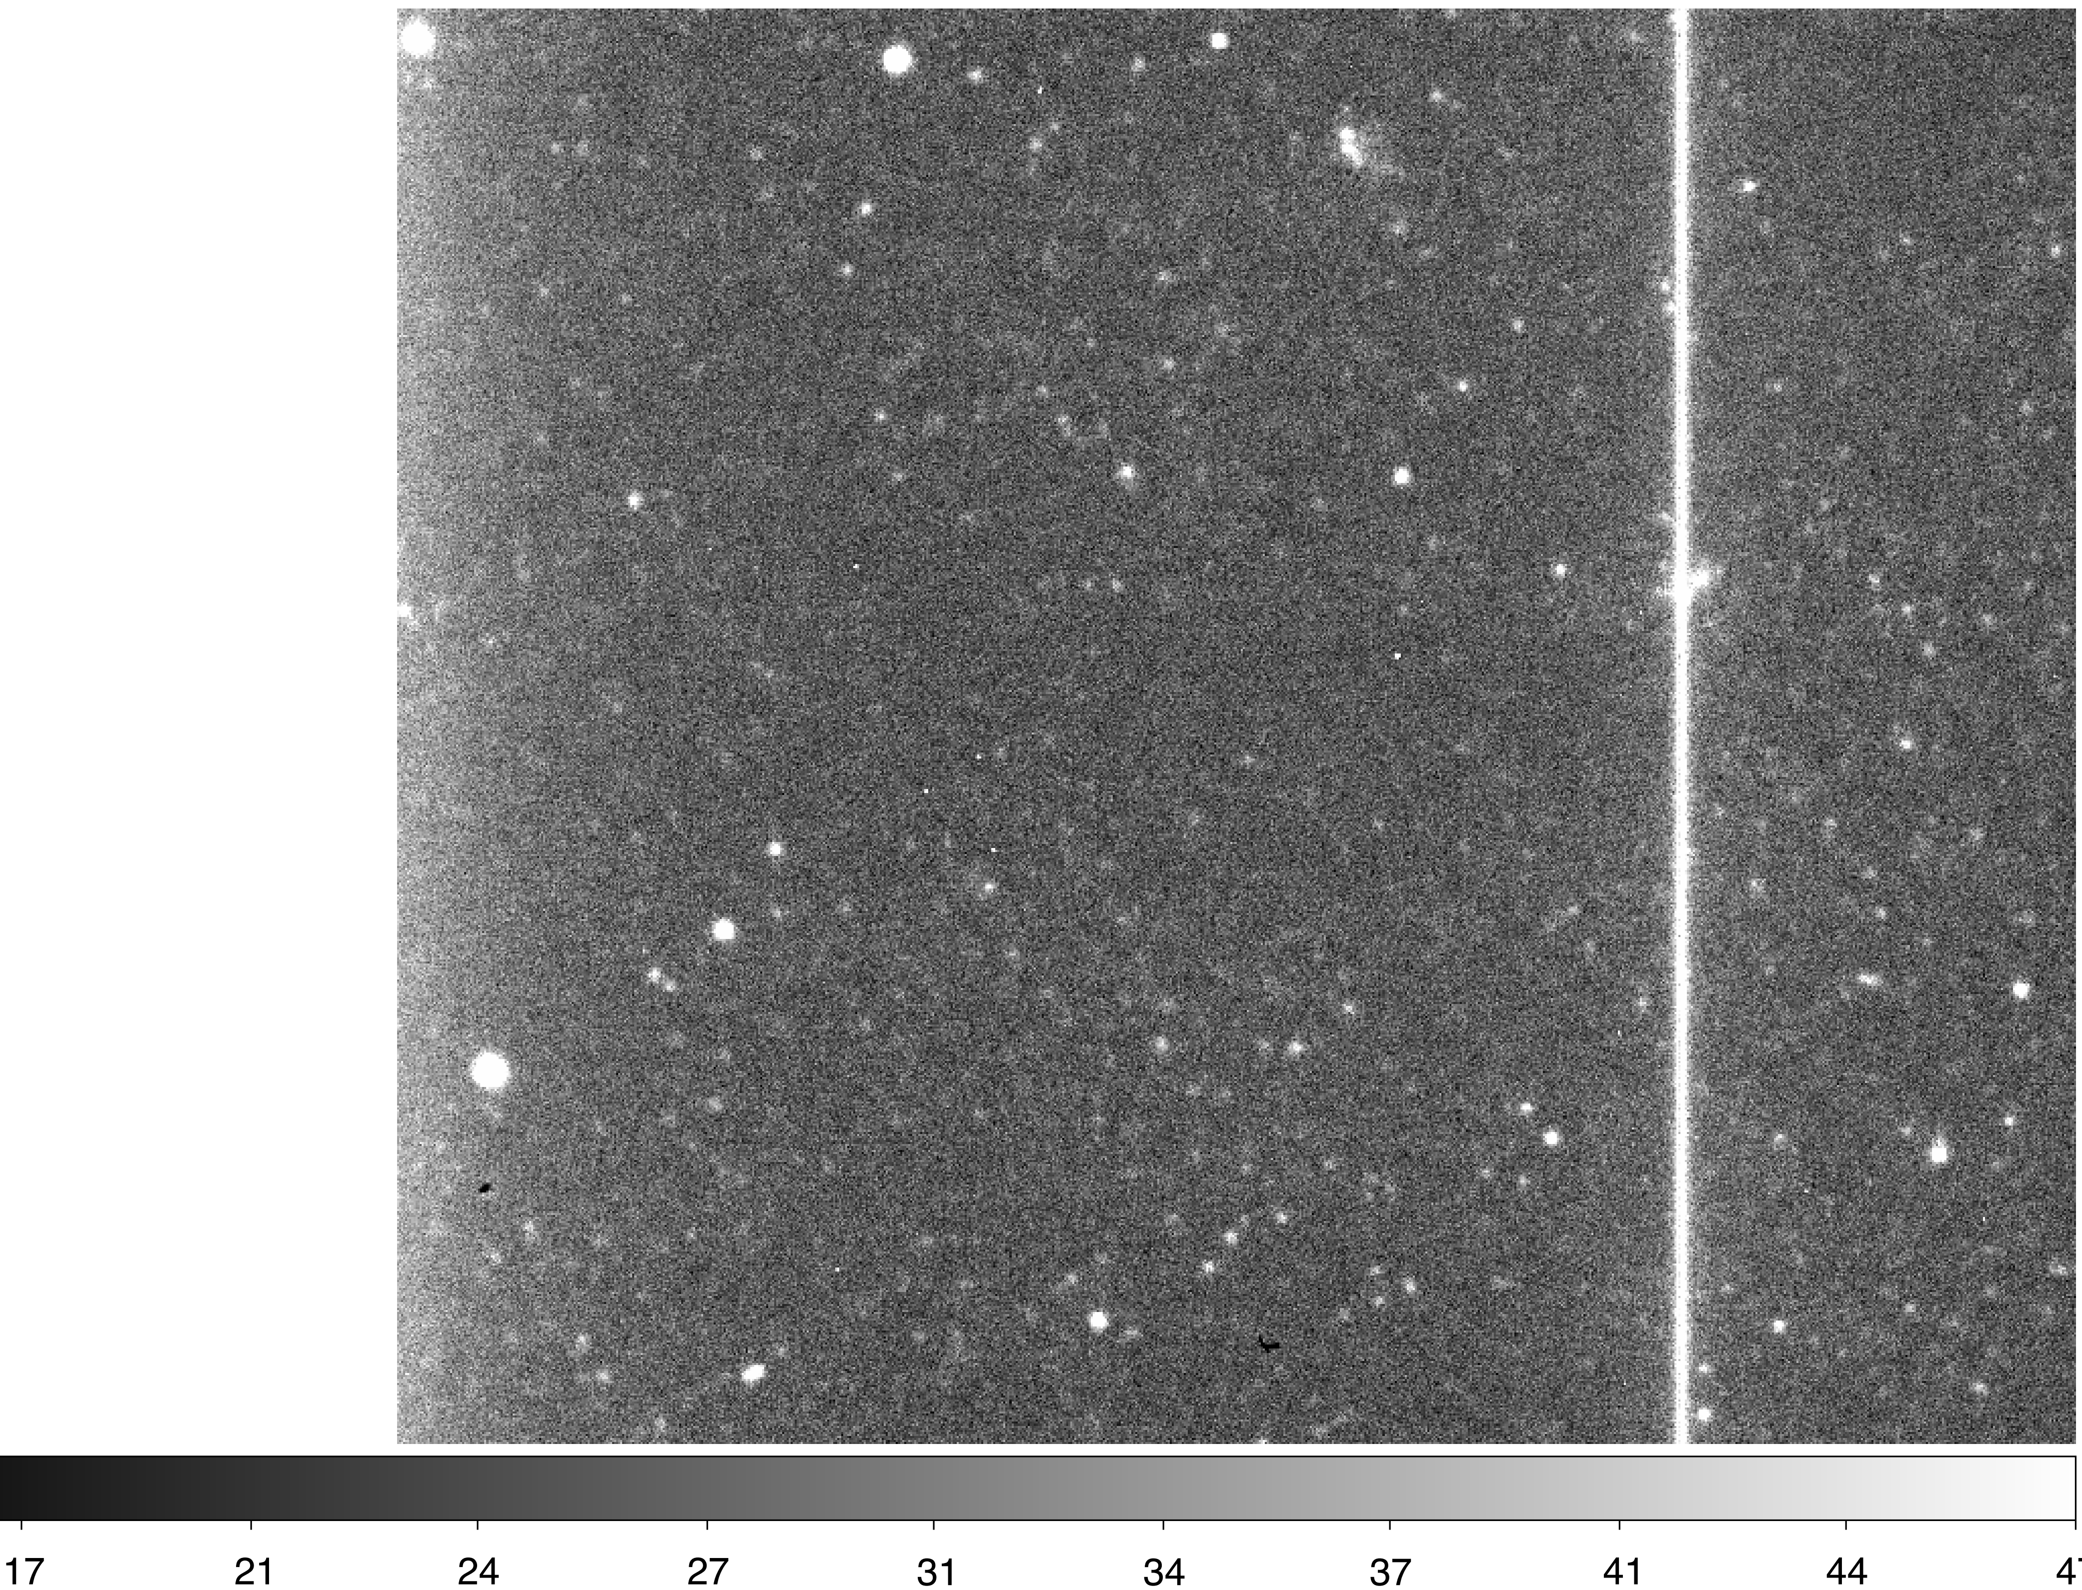

Subtracting bias structure

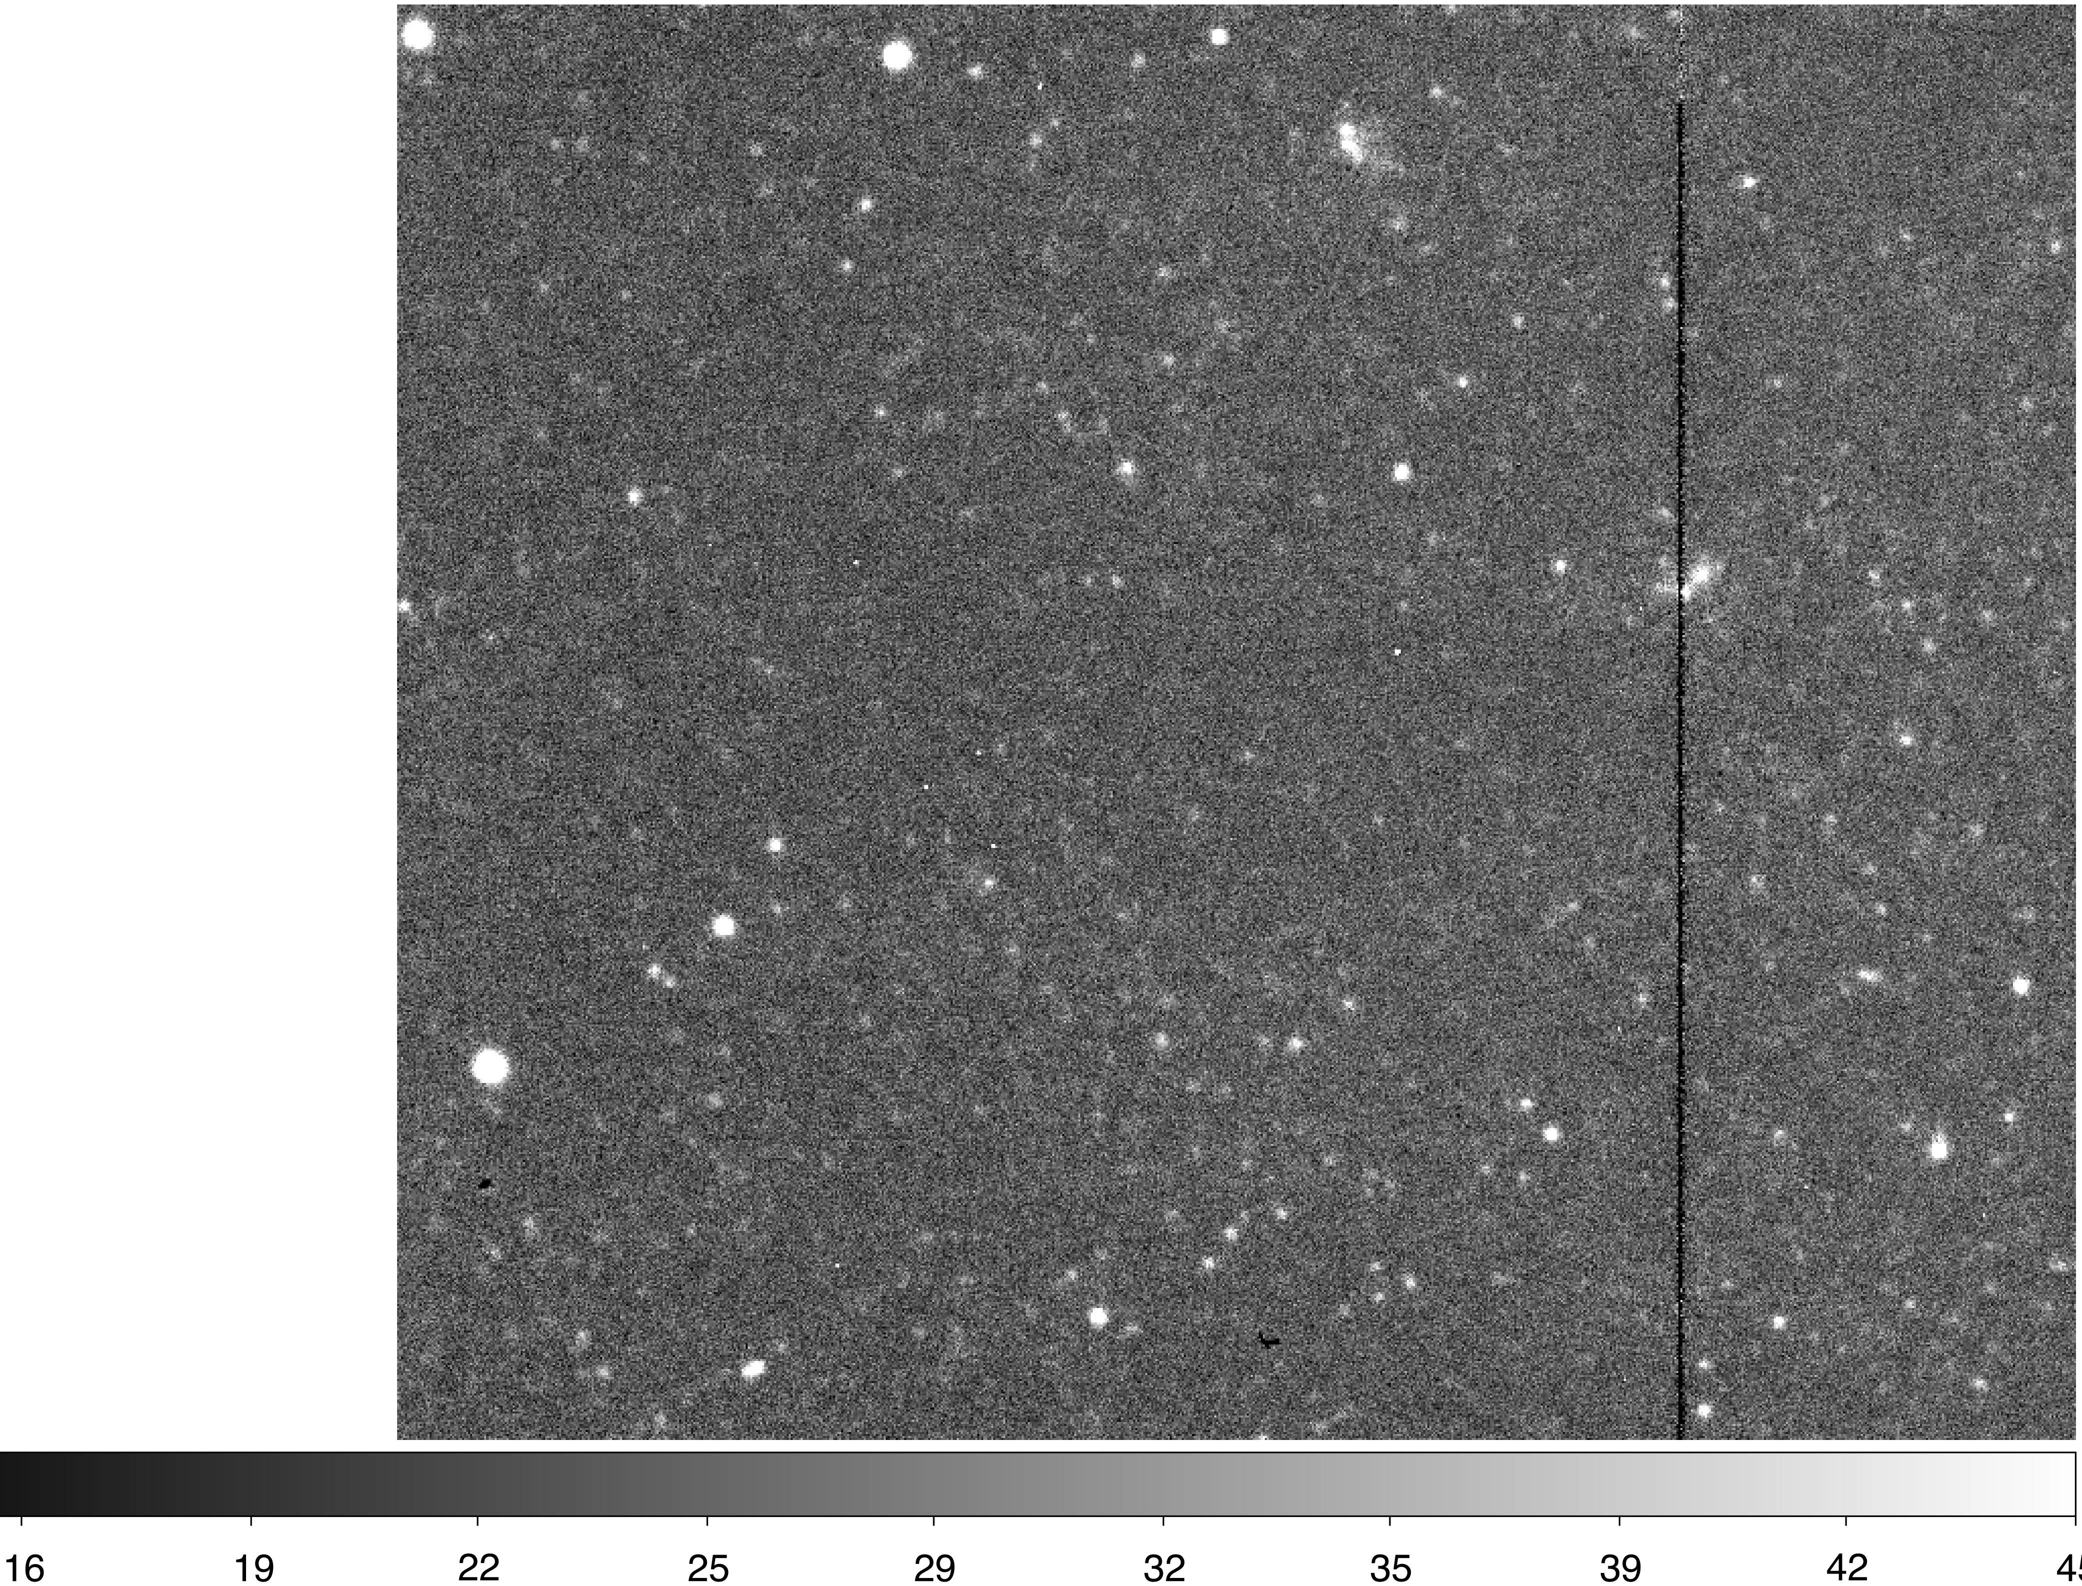

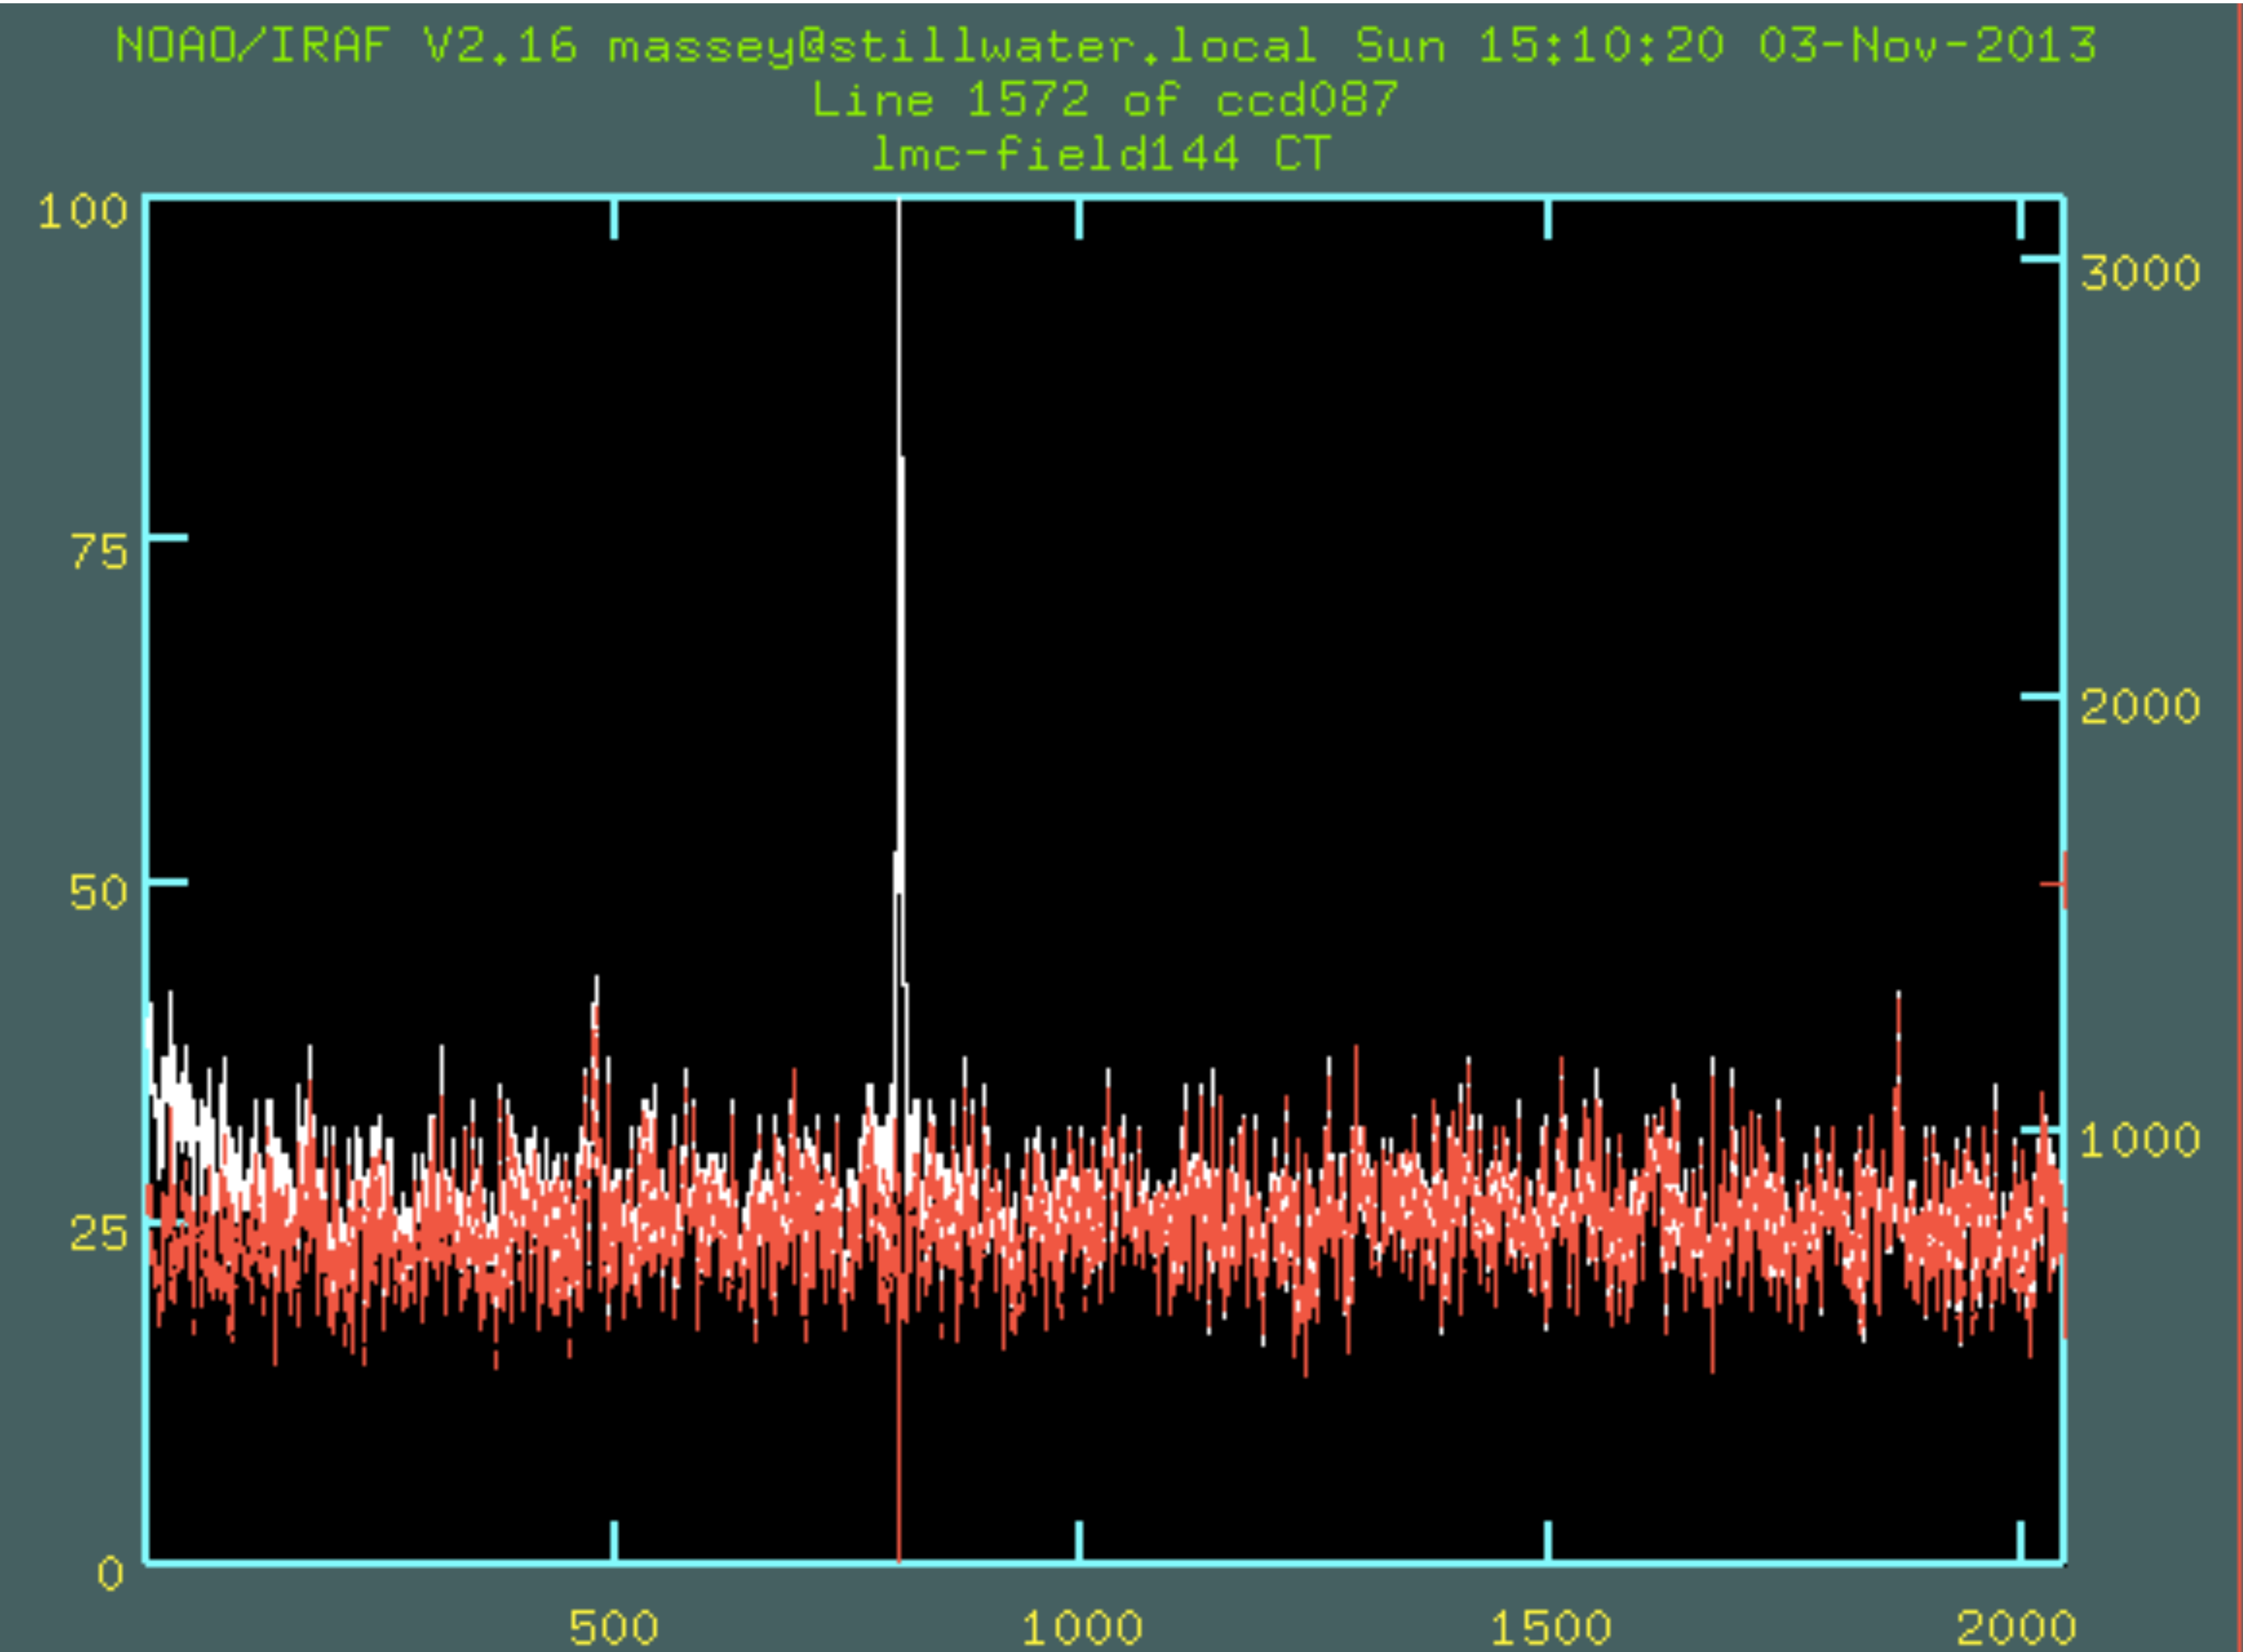

NOAO/IRAF V2.16 massey@stillwater.local Sun 15:10:20 03-Nov-2013
Line 1572 of ccd087
lmc-field144 CT
100
75
50
25
0
3000
2000
1000
500
1000
1500
2000

Let's look at a dark

- Take an exposure that's as long as your longest exposure (5 minutes in my case). Remove the overscan, remove the bias structure, and see if there's anything significant left.

Not much there.

- Next: multiplicative effects!