Calibration

January 30, 2024

University of Rochester

Instrumental Sensitivity & Calibration

Useful references

- Adam Block, *Deep sky imaging: workflow 1* in Rob Gendler (ed.) 2013, *Lessons from the masters* (New York: Springer), 159–192
- Ralph Bohlin et al. (2014), PASP 126, 7111B (HST master collection)

HII regions NGC 7635 (left) and NGC 7538 (right), LRGB (Mees Observatory image)

Calibration

To get the ideal sensitivity previously discussed requires accurate calibration:

1. Subtraction of camera-electronics offsets (bias) and average dark current
   - Requires additional data: low-noise dark and bias frames at the same CCD temperature as for the target data.
2. Correction of variation in responsive quantum efficiency (flat field)
   - Requires additional data: low-noise flat fields at the same CCD temperature and filter as for the target data.
3. Correction of permanently-troublesome (hot or cold) pixels
   - Many instrument mavens create maps of these pixels that can be used in correction.
4. Correction of atmospheric extinction from the variation of signal with zenith angle (normalization)
5. Identification and removal of transients, such as cosmic-ray hits or satellite trails
6. Scaling the signal to that of objects — usually stars at visible wavelengths — with accurately-known flux density or magnitude (flux calibration)
   - Requires additional data: observation of flux standards — at visible wavelengths, usually faintish A0V stars — over a range of zenith angle.

Steps 4 and 6 can sometimes be combined or rearranged.

Dark current and bias correction

Dark current varies from pixel to pixel, and decreases exponentially with decreasing temperature for good detectors.

- All intrinsic semiconductor photodetectors are like this: for a bandgap ΔE, carriers can be thermally produced at rates proportional to $\exp\left(-\Delta E/kT\right)$.

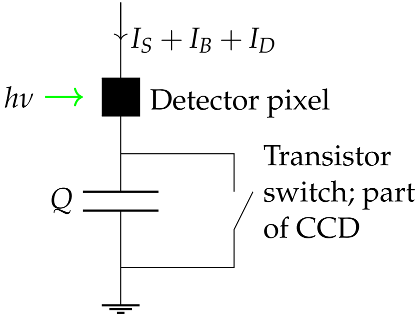

The semiconductor-transistor switches in CCD circuitry always inject charge abruptly when opened at the beginning of exposures. This is temperature dependent, too.

These effects lead to offsets in the charge, Q, that must be subtracted to leave the signal.

Dark current has shot noise, too. Subtraction gets rid of the average value, but not the shot noise.

Dark current and bias correction

Also, the difference between (or sum of) two statistically independent images has noise no smaller than either one:

$$\overline{Q} = \overline{Q_2} - \overline{Q_1} = q_e(\overline{N_2} - \overline{N_1})$$

$$\overline{(\Delta Q^2)} = \overline{(\Delta Q_2^2)} + \overline{(\Delta Q_1^2)}$$

$$= q_e^2(\overline{N_2} + \overline{N_1}) \qquad \text{if Poisson-distributed}$$

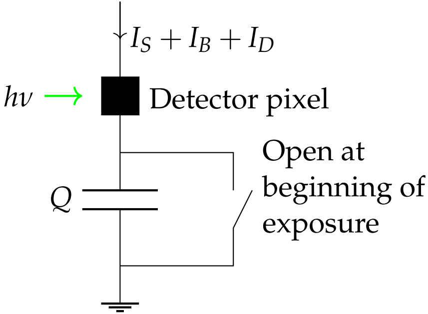

Therefore, detector arrays are operated at a low enough temperature to make both the offset and shot noise small…

And the dark current and bias correction are measured with averaging times long enough that the noise in these two corrections contribute insignificantly to the system sensitivity.

Dark current and bias correction

For example, suppose that frame 2 is signal plus background plus dark current, and frame 1 is dark current, measured with exposure times Δt_2 and Δt_1, respectively.

If the average value of the dark current in each pixel is constant, and we average n frames of dark current instead of 1 before subtracting, then the signal plus background, and the noise in this quantity, is

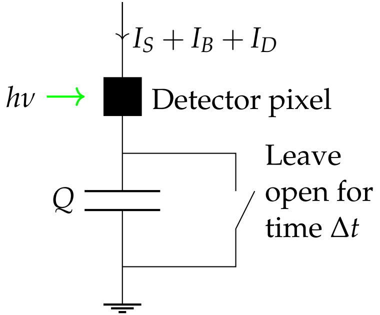

$$\overline{Q} = \overline{Q_2} - \frac{\Delta t_2}{\Delta t_1}\overline{Q_1} = q_e\left(\overline{N_2} - \frac{\Delta t_2}{\Delta t_1}\overline{N_1}\right)$$

$$\overline{(\Delta Q^2)} = \overline{(\Delta Q_2^2)} + \frac{1}{n}\overline{(\Delta Q_1^2)}$$

$$= q_e^2\left(\overline{N_2} + \frac{1}{n}\frac{\Delta t_2}{\Delta t_1}\overline{N_1}\right) \to q_e^2\overline{N_2} \quad \text{for } n \gg 1, \Delta t_1 > \Delta t_2$$

Dark current and bias correction

In practice, **for each temperature and binning in use**, we usually

- Take 44 dark frames, each with 30-minute exposure time
- Take 280 bias frames, each with the minimum possible exposure time (∼50 ms)
- Generate master dark and bias frames by averaging corresponding pixels, after discarding the smallest 4 (10) and largest 8 (14) values for each pixel in the stack of dark (bias) frames to reject cosmic ray hits on pixels.
- Then the dark current and bias correction is

$$\overline{Q_2'} = \overline{Q_2} - \frac{\Delta t_2}{\Delta t_1}\left(\overline{Q_{M,\mathrm{dark}}} - \overline{Q_{M,\mathrm{bias}}}\right) - \overline{Q_{M,\mathrm{bias}}}$$

  where M denotes the master frames.
- Dark frames can be taken during the day if the camera surroundings are truly dark.

Dark current and bias correction

The resulting corrected frame has, to good approximation, **no noise added by the master**. Consider the dark frames:

- If the CCD temperature is stable, then each dark integration would have approximately the same average charge and noise per pixel, $q_e\overline{N}$ and $q_e^2\overline{\Delta N^2}$. Since the noise in each frame is statistically independent, the variance in the master is the quadrature sum of the variances of the n dark frames:

$$\overline{Q_M} = \frac{q_e}{n}\sum_{i=1}^{n}\overline{N_i} = q_e\overline{N}$$

$$\overline{\Delta Q_M^2} = \frac{q_e^2}{n^2}\sum_{i=1}^{n}\overline{\Delta N_i^2} = \frac{q_e^2}{n^2}\sum_{i=1}^{n}\overline{N_i} = \frac{q_e^2\overline{N}}{n}$$

- So the dark master has rms noise smaller than the noise in the dark current contributed (unavoidably) to the on-target exposures by a factor of $\sqrt{n\Delta t_2/\Delta t_1} \approx 13$ for five-minute frames.
- Similarly, the bias master has rms noise lower by $\sqrt{n} = 16$.

Camera 2 master dark and bias

-20°C, linear scale

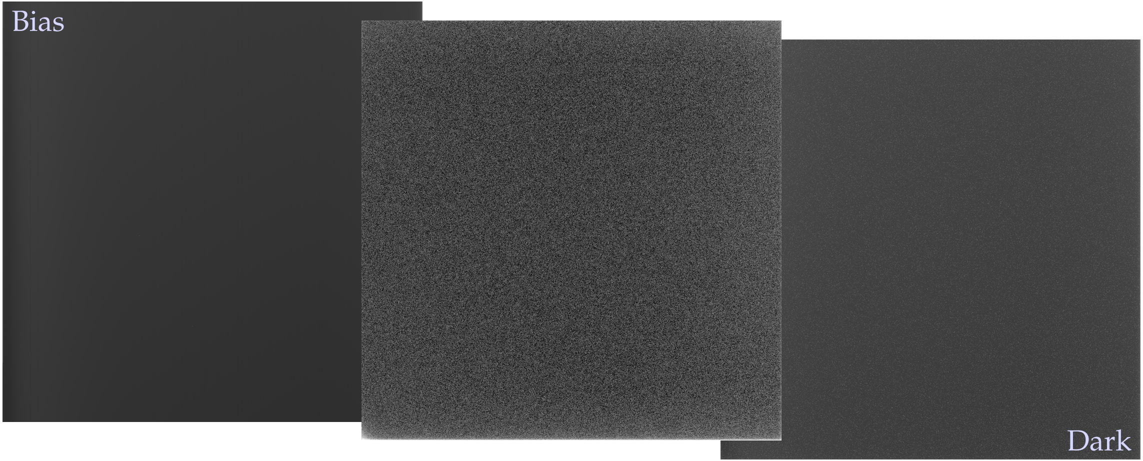

Dark − Bias, Linear scale, 0.01–0.02 e^-/s.

Flat field correction

- Good CCD detector arrays have small pixel-to-pixel variation in response, meaning that the responsive quantum efficiency is nearly uniform across the array.
- But the variation is not zero, and it would look like noise or systematic error in long enough exposures.
- Any optical surface close to the focal plane — e.g. windows and filters — can also contribute small variations in transmission among different areas of detector pixels.
- Ideally, an image of a uniformly-bright object would be perfectly uniform.
- So take a high signal-to-noise, dark and bias corrected, image of a uniform-brightness object, normalize it, and divide each target-exposure frame by that image.
- This restores uniformity to the focal-plane response, and it is called flat-fielding.

Taking flat field frames

The hard part about flat-field correction turns out to be getting a uniformly bright object, with light taking the same optical path as that from celestial objects. Many methods that you might think are no-brainers actually do not work.

- Good: twilight-sky flats (professional-observatory standard)
  - Point at zenith, take many frames with large DN per pixel in each filter after sunset and before the sky gets dark enough to see stars. (Reverse for morning.)
  - Disadvantages: Has to be done at twilight; usually not enough time to take data for all of the filters; sky emission is scattered sunlight and therefore polarized.
- Better: electroluminescent-panel flats (Mees Observatory standard)
  - Completely cover the top of the telescope with seamless EL panel, take many frames with a large DN per pixel in each filter. Can be taken any time you want, even during the day.
  - At Mees: as good as the best twilight flats, and unpolarized as well.

Flat field correction

- Currently, our flats are taken with an electroluminescent panel that lives on the telescope side of the remote-control telescope cover.
- Take 32 frames in each filter, with an exposure time set for a signal that is typically half of the maximum signal (so around 32,000 DN).
  - Recall that the maximum signal is 2^{16} DN = 65536 DN. This signal level is often called full wells.
- This requires about 0.5 s exposures for L, 0.3–1 s for RGB at 2×2 binning, and 5–60 s for the narrowband filters at 2×2 binning.
- Generate a master flat for each filter by bias-subtracting the individual frames and averaging the results.
- This takes close to two hours to flat-field all seven filters, so it is important to either start taking flats well before sunset, or only take the flats that you need.

Master L flats

October 28, 2019

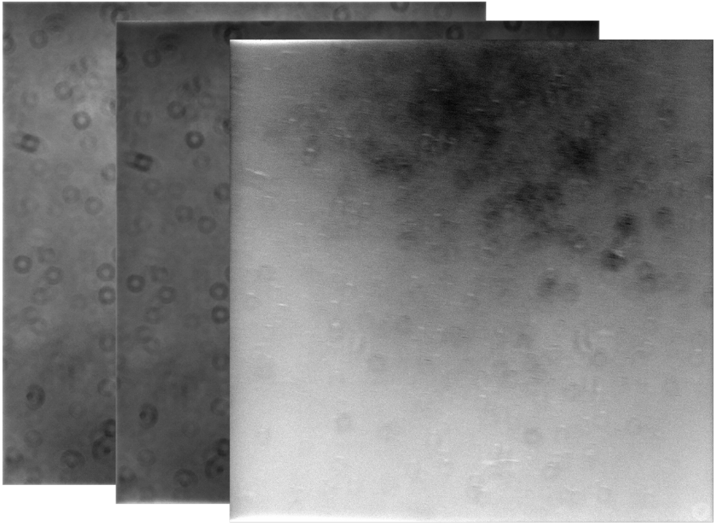

Electroluminescent planel. Linear scale

Twilight sky. Linear scale

Sky/EL. Linear scale

Most of the dust is on CCD cover window and is illuminated very slightly differently by the EL panel and the sky.

Master EL-panel flats

October 28, 2019

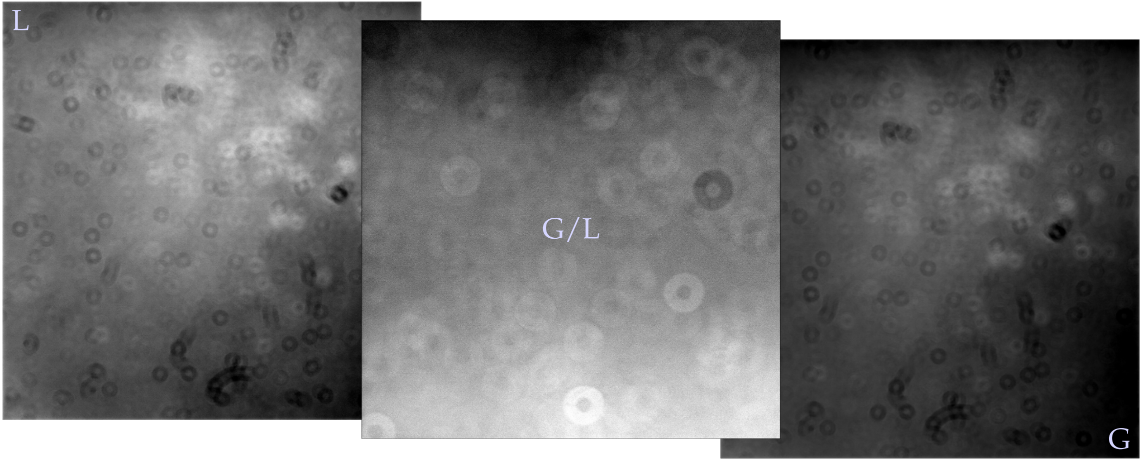

Linear scale, 1.00 ± 0.03. The CCD cover window was so dusty that the dust on the filters only shows up well in the quotients.

Atmospheric absorption

Presuming that atmospheric conditions are the same throughout the observations:

- Because of atmospheric scattering and absorption, the signal from each star is largest when observed at its highest elevation / smallest zenith angle z.
- To a good approximation, the atmosphere is plane-parallel, and the path that light takes through it varies according to

$$\ell = \frac{\ell_0}{\cos z} = \ell_0 \sec z$$

  where ℓ_0 is the vertical thickness of the atmosphere.
- If atmospheric opacity and scattering is small, light is attenuated according to

$$\begin{aligned} f &= f_0 \exp\left(-\sigma_0 \ell_0 \sec z\right) = f_0 \exp\left(-\tau_0 \sec z\right) \\ &\cong f_0 - f_0 \tau_0 \sec z \end{aligned}$$

  where f_0 is the **unattenuated** stellar flux in DN, τ_0 is the zenith optical depth of the atmosphere (not to be confused with filter transmission), and $(1 - \tau_0) f_0$ is the signal in DN that would be received at zenith.

Correcting atmospheric absorption

1. Measure f from stars in the images as a function of z at which each image is taken.
2. Fit a line to the results, determining f_0 from the $\sec z$ intercept, and $f_0\tau_0$ from the slope.
   - τ_0 should be the same for all the stars in an image, though f_0 varies from star to star. If conditions are good, τ_0 will be the same for all images in the same filter.
3. Calculate a correction factor $\overline{f_0/f}$ for the stellar signals to which lines were fit.
   - If $\overline{f_0/f}$ does not have a small variance, look for evidence of changing atmospheric conditions.
4. Multiply each image by its $\overline{f_0/f}$.

Images of the same object and filter can now be averaged, making a high S/N image of the target object.

Atmospheric-absorption correction

During observations of NGC 891 on October 28, 2019, a certain relatively isolated 13.3-magnitude star appeared in every frame. In the L frames, photometric measurements of this star, listed to the right, were made for each frame. Find the L-band zenith optical depth, τ_0, and the absorption-corrected flux density, f_0, for this star.

First, note that conditions were deteriorating as time went on, as is evident here in the seeing. So we may not want to use all of the data.

After editing the data, though, it is just plot-and-fit.

Time	Altitude [deg]	sec z	Signal [$\times 10^5$ DN]	Noise, rms [DN]	Seeing [arcsec]
19:45	34.399	1.7701	5.90	921	2.2
19:56	36.136	1.6958	5.86	909	2.4
20:01	36.933	1.6642	5.80	912	2.4
20:07	37.894	1.6281	5.99	932	2.2
20:12	38.699	1.5994	6.01	915	2.0
20:17	39.508	1.5719	5.98	911	2.2
20:22	40.321	1.5454	6.10	924	2.3
20:28	41.301	1.5151	6.17	895	2.1
20:33	42.121	1.4910	6.03	895	2.4
20:38	42.945	1.4678	6.04	937	2.5
20:43	43.772	1.4455	6.16	909	2.2
20:49	44.769	1.4199	6.22	923	2.2
20:54	45.603	1.3996	6.19	906	2.3
22:09	58.438	1.1736	6.31	901	2.4
22:15	59.486	1.1608	6.35	923	2.4
22:20	60.361	1.1505	6.39	947	2.4
22:25	61.238	1.1407	6.40	936	2.4
22:30	62.117	1.1313	6.39	940	2.4
22:36	63.174	1.1206	6.32	981	2.6
22:41	64.056	1.1121	6.28	968	2.8
22:46	64.939	1.1039	6.37	953	2.6
22:51	65.824	1.0961	6.38	950	2.6
22:57	66.887	1.0873	6.36	977	2.5
23:02	67.774	1.0803	6.34	979	2.6
23:07	68.662	1.0736	6.41	1000	2.6
00:15	80.744	1.0132	6.19	1130	3.1
00:20	81.619	1.0108	6.18	1170	3.0
00:26	82.661	1.0083	6.04	1200	3.5
00:31	83.517	1.0064	6.14	1090	3.3

Atmospheric-absorption correction

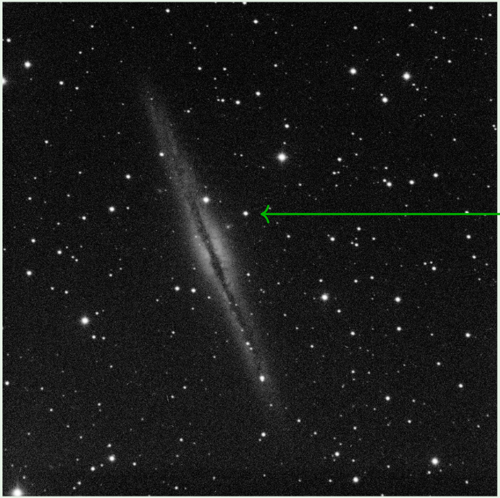

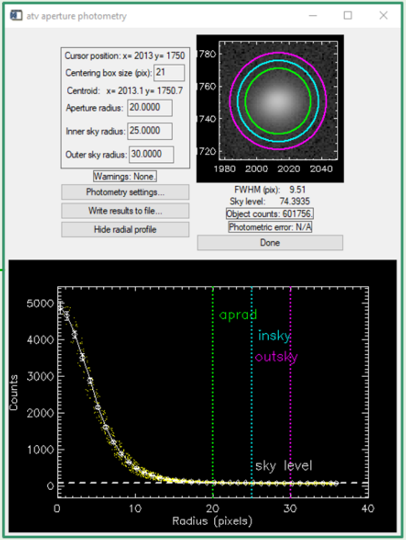

Atmospheric-absorption correction

Here is the plot and fit, omitting the last six frames (points scattered well below the line on the y-axis) from the fit.

Note that the scatter is much larger than the noise: conditions were "not photometric."

From the fit:

$$\tau_0 = \frac{7.67 \times 10^4}{7.22 \times 10^5} = 0.106$$

$$f_0 = 7.22 \times 10^5 \text{ DN}$$

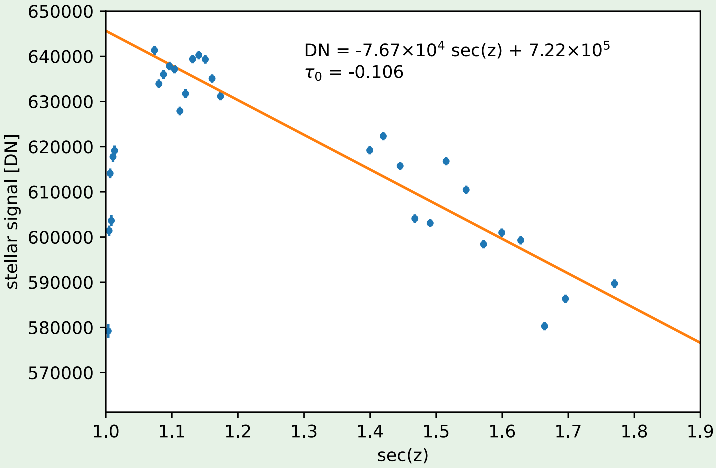

Normalization in CCDStack

Some image-reduction software packages have handy routines for normalizing the stack of images to a chosen part of one of the images. CCDStack is one of those packages.

- This automatically corrects images so that the stars in the stack have the same signal (on average) as those in the chosen part of the chosen image, thus significantly reducing the photometric scatter (though it often does not reduce the scatter to the level of the background noise).
- To correct the stack for atmospheric correction after that, scalar-divide every frame (i.e. divide every pixel) by the factor $(1 - \tau_0 \sec z_0)$, where

  τ_0 is the zenith optical depth as determined above

  z_0 is the zenith angle at which the "chosen frame" was observed

Transient removal

In one of the examples from Lesson 1, we designed a G observation that was background-limited and reached 24^{th} magnitude in six 10-minute frames.

Why not take just one 60-minute frame?

- The **average of six 10-minute background-limited frames has precisely the same sensitivity as one 60-minute background-limited frame**, as is clear from both that example and today (slide 5).
- Multiple frames allow other problems to be easily avoided, such as

  Overexposed bright stars: Sometimes, this cannot be helped, but it should be avoided whenever possible.

  Cosmic ray hits: They occur in random locations and are easy to identify and correct in a stack of images.

  Satellite trails: 1–2 per hour, especially at high declination.

Transients

Five-minute G frame on NGC 5866, with

1. an overexposed star
2. a satellite trail
3. many cosmic-ray hits, of which only a couple of oblique-incidence ones are labeled

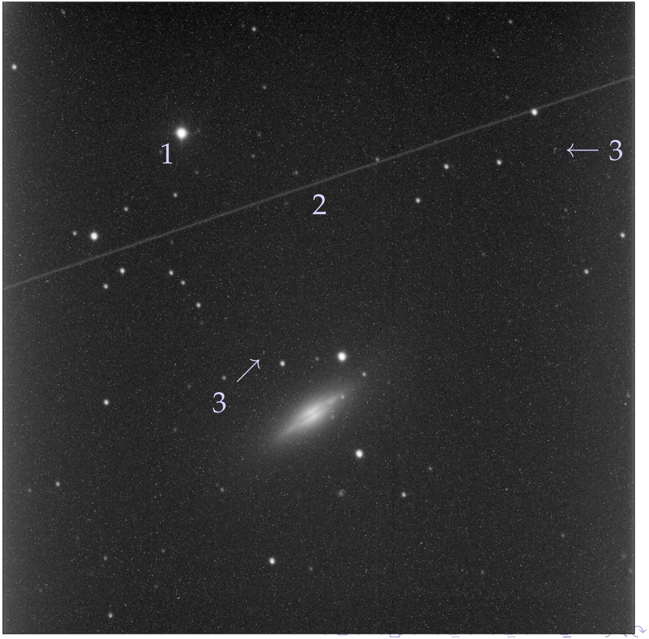

Transients

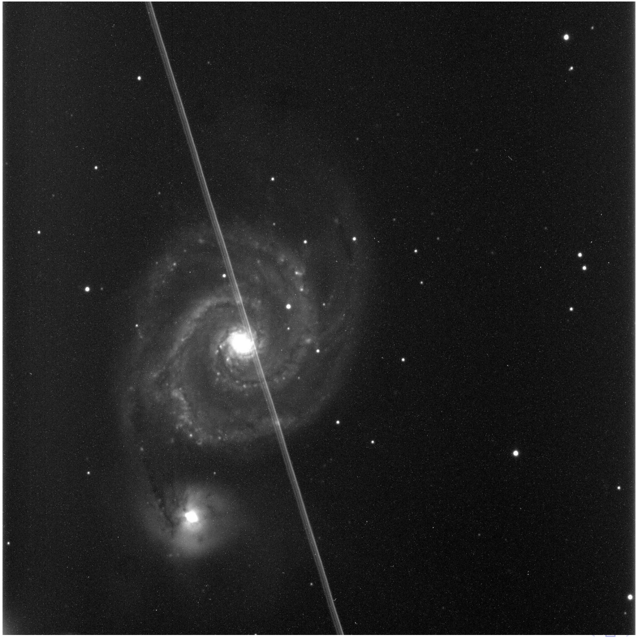

Transient removal

There is not much that you can do about overexposed stars, unless they were your main target.

To get rid of transients,

- Align the stack of frames so that every star occupies the same pixels in each frame.
- Compute the median and the standard deviation, σ, for each pixel from the stack.
- Reject each pixel that deviates by more than 2.2σ from the median, replacing its value with that interpolated from its un-rejected nearest neighbors. This procedure is one form of median filtering.
  - It can take a few iterations to remove satellite trails and occasional clusters of cosmic-ray hits, for lack of enough un-rejected neighboring pixels.