Preliminary ZTF Skyflat Results

Jason Surace

May 14, 2018

PTF was calibrated by means of skyflats, wherein the night sky glow was used to map the pixelwise relative throughput of the detectors. ZTF, however, uses a more complex flatfield derived from observations of a flatfield screen near the telescope entrance aperture, to be supplemented by means of a low order correction derived from photometric observations of stars (the latter is not currently implemented).

Here I describe the derivation of skyflats for ZTF via an offline analysis pipeline. Variations of this pipeline have been used extensively for ground-based instruments such as WIRC, as well as providing the fundamental flatfield calibration of the IRAC instrument on Spitzer. I then compare these to the existing ZTF pipeline "high frequency" flats, which are the observations of the flatfield screen. Two primary differences are expected. First, the night sky is a true uniform illuminator located at effectively infinite distance. We might expect to see differences across the focal plane, as well as small scale variations in the illumination of "dust doughnuts", e.g. out of focus dust particles on various optical surfaces intermediate between the focal and pupil planes, as a result of small angle differences between the flat illuminators. Second, the night sky is a different color than the flatfield illuminator, so some effects should be seen due to any variations in sensitivity as a function of color, or shifts in filter effective wavelength.

1. Derivation of Skyflats

Like the online pipeline, all the data for any given quad is processed independently.

a. In a given night, for each filter/ccd/quad combination, each image is unpacked along with it's associated overscan strip. The skyflats are made on a purely nightly basis.

b. For each image, the median bias level in each line in the overscan strip is subtracted from the accompanying quad, creating a bias-subtracted image.

c. Since the sky brightness varies as a function of time and telescope position, each image is multiplicatively renormalized to one. The median count level per pixel (for the image) is retained as metadata.

d. An initial flat is made by median stacking and renormalizing. This flat is discarded after processing.

e. The original bias-subtracted images are flattened with this initial flat.

f. A source detector (sextractor) is run on each bias-subtracted, flattened frame to create an object mask.

g. The bias-subtracted images are then renormalized using this object mask to further refine the "sky" level.

h. This renormalized set of images is then stacked, using temporal and spatial outlier rejection (image stats and object masks, respectively) and weighting based upon the actual input counts. The result is then normalized again to create the final flat.

For memory reasons, I am currently limiting the input data to the first 400 frames. In R-band, for the night I am currently looking at, this is over one million counts per pixel. Given the system gain, the result should be accurate to better than one tenth of a percent in terms of sheer poisson statistics. For analytic reasons, I also dump as diagnostic files a depth of coverage map (number of samples in the stack) as well as the reduced error map (observed scatter per pixel divided by root-N in the stack).

Additionally, I have made no corrections for the scattered light issues along the ccd edges.

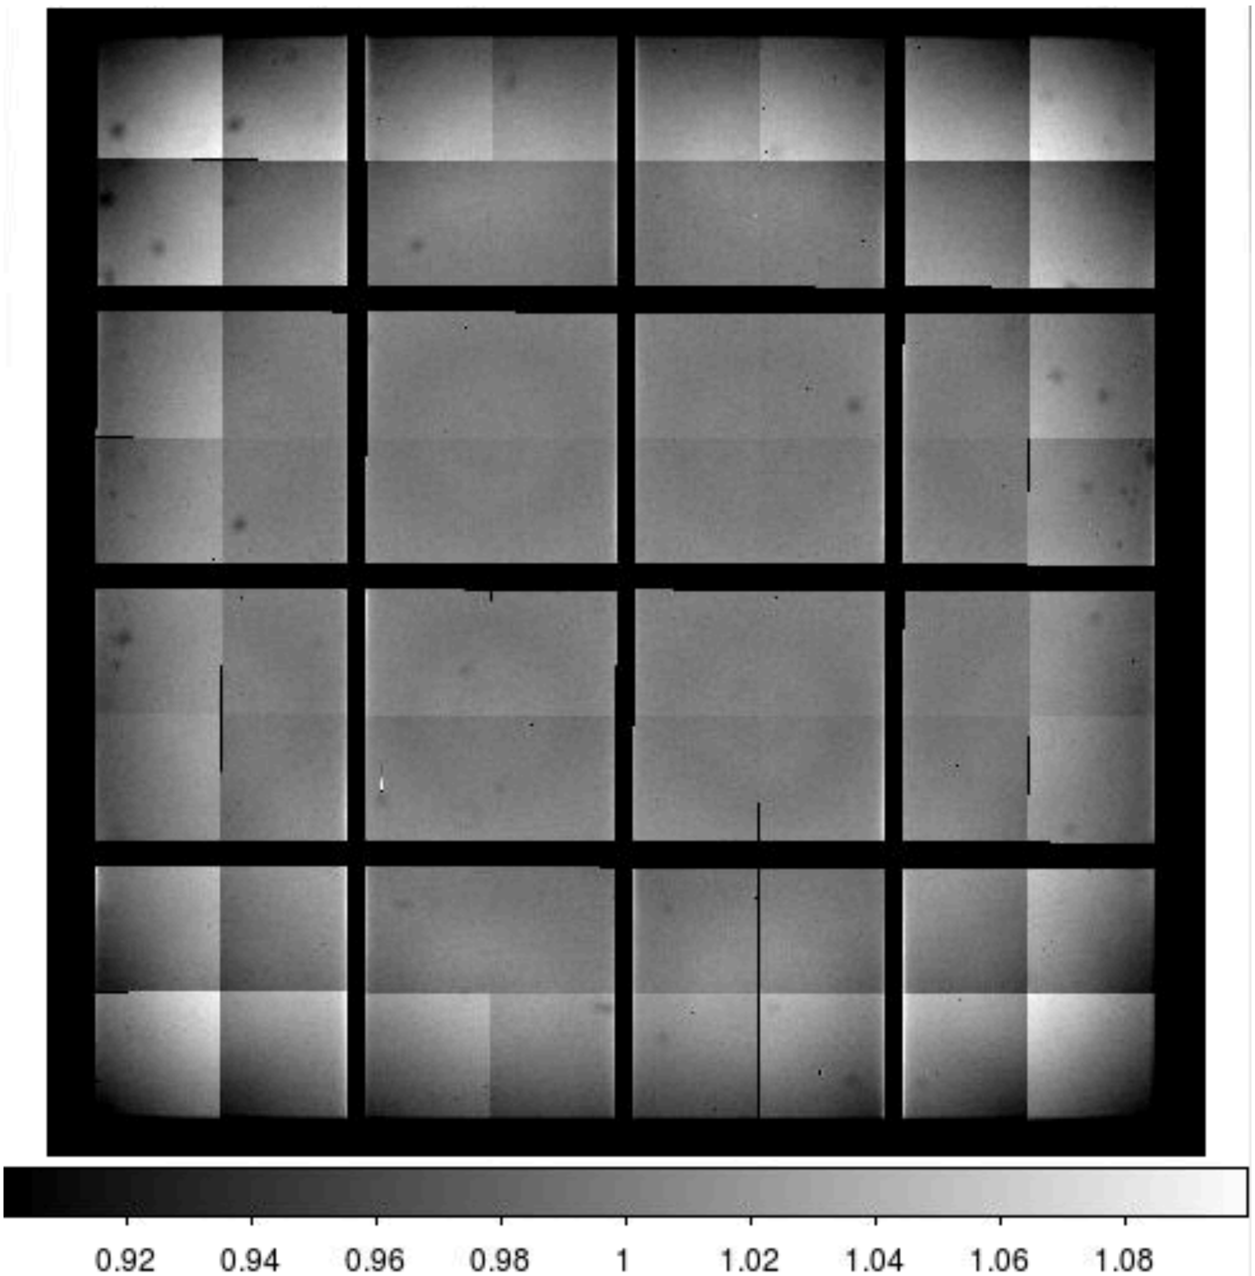

R-Band skyflats for ZTF from 03/06/2018. Like the online pipeline, in this image each quad is independently normalized. The quads have been glued together via "swarp" using the astrometry for the final input exposure, hence some small visual tilt.

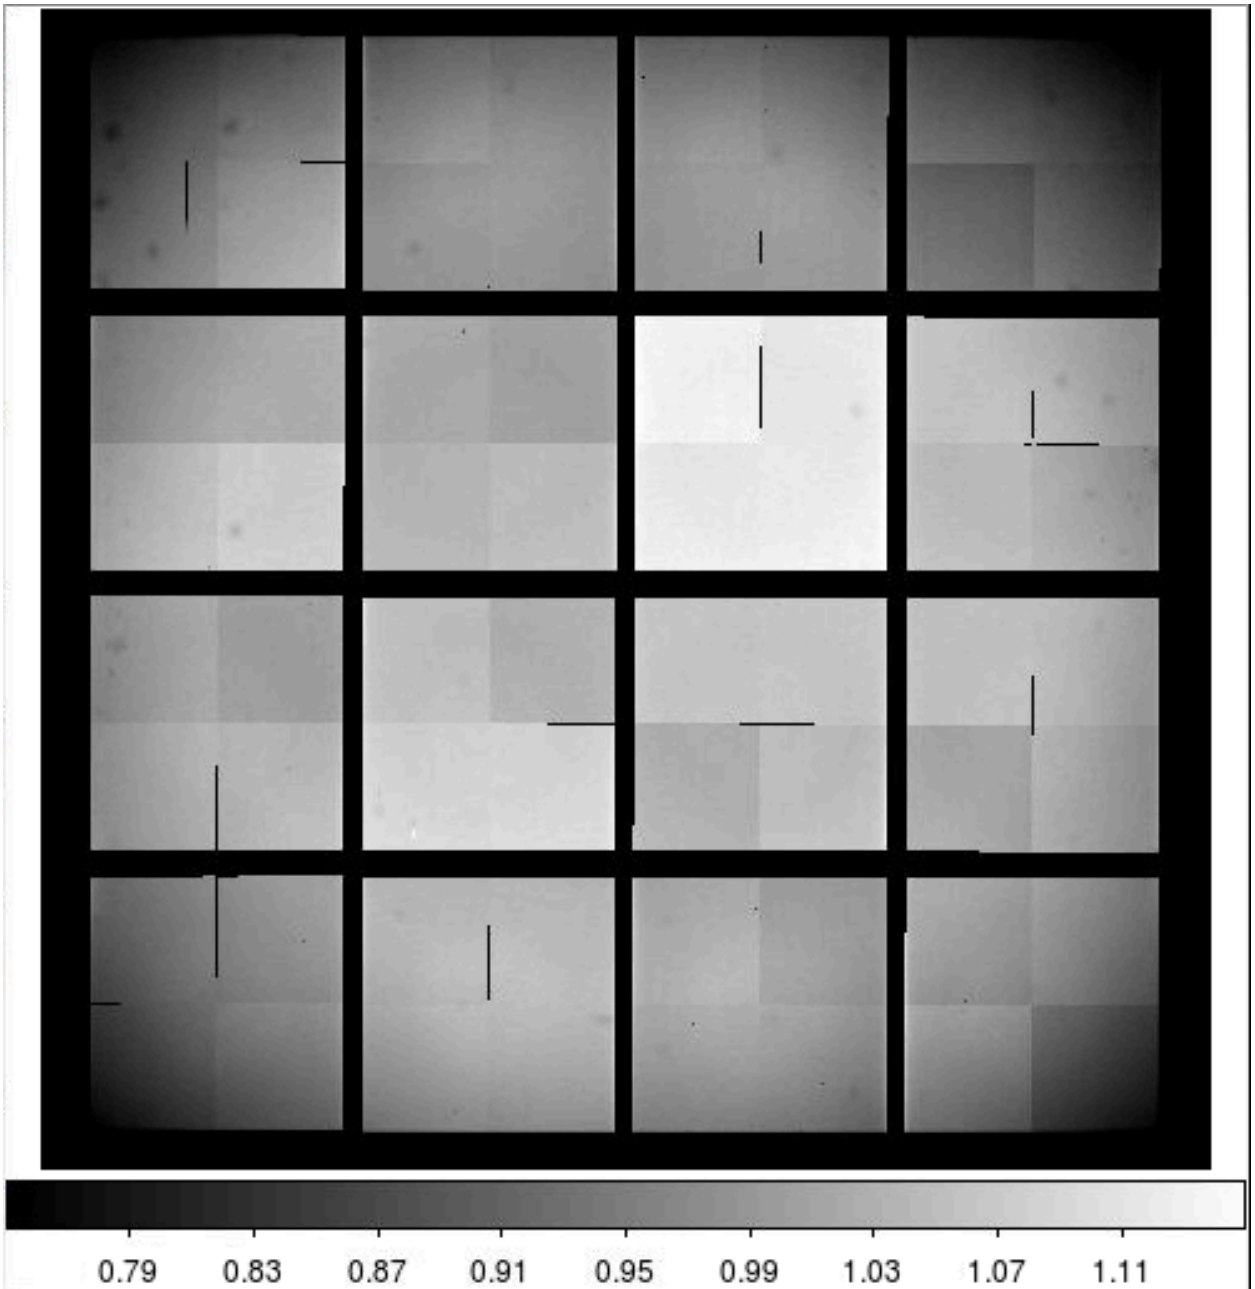

Same skyflat as above, but each quad has been rescaled according to it's median sky count level. This shows the true vignetting, as well as the sensitivity variations between quads.

2. Differences Between Sky and High Frequency Flats

To compare the two sets of flats, I have quad-by-quad taken the ratio of skyflat to high frequency (i.e. “dome”) flats created by the online pipeline in the same night. In order to ensure consistent normalization, both sets of flats were renormalized to a mean of one, excluding the outer 60 pixels in each quad. This should be sufficient to exclude most of the edge scattering. The result is show below.

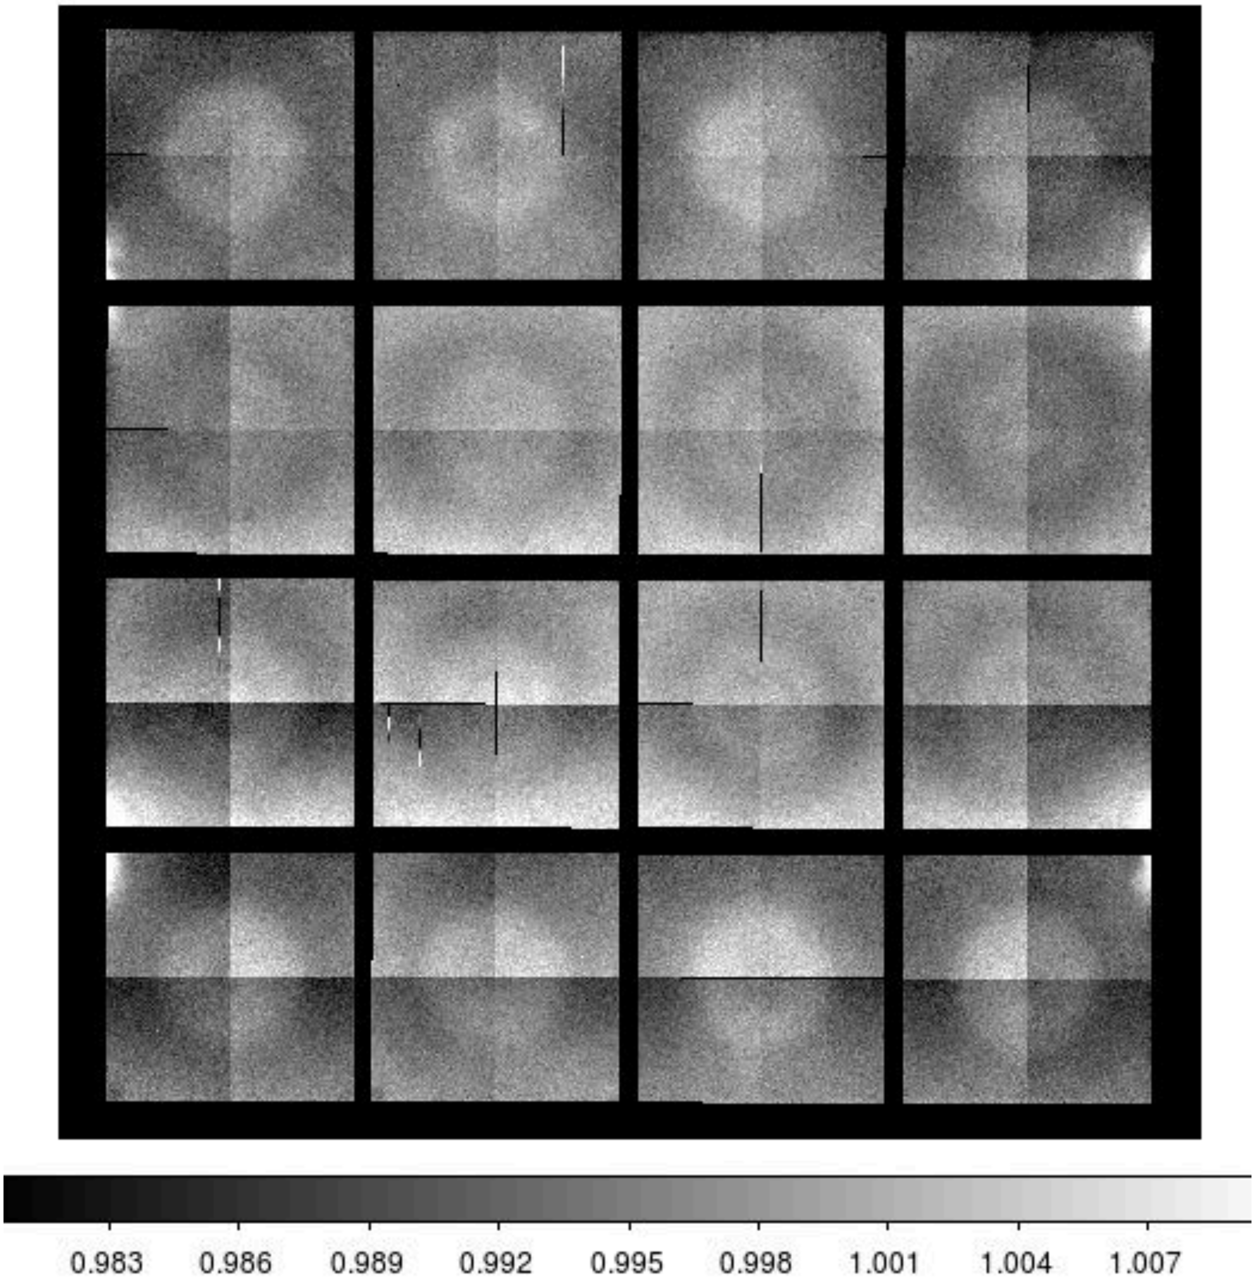

The large-scale variations are typically small, and at the level of a percent. It’s clear that the circular pattern is replicated in every ccd, so it’s something either in the ccds or the optical elements directly above each one. I believe Roger has attributed the observed rings to variations in color response. It is particularly notable that while the flats have an extremely strong and clear fringing pattern, this pattern is reproduced in both sets of flats.

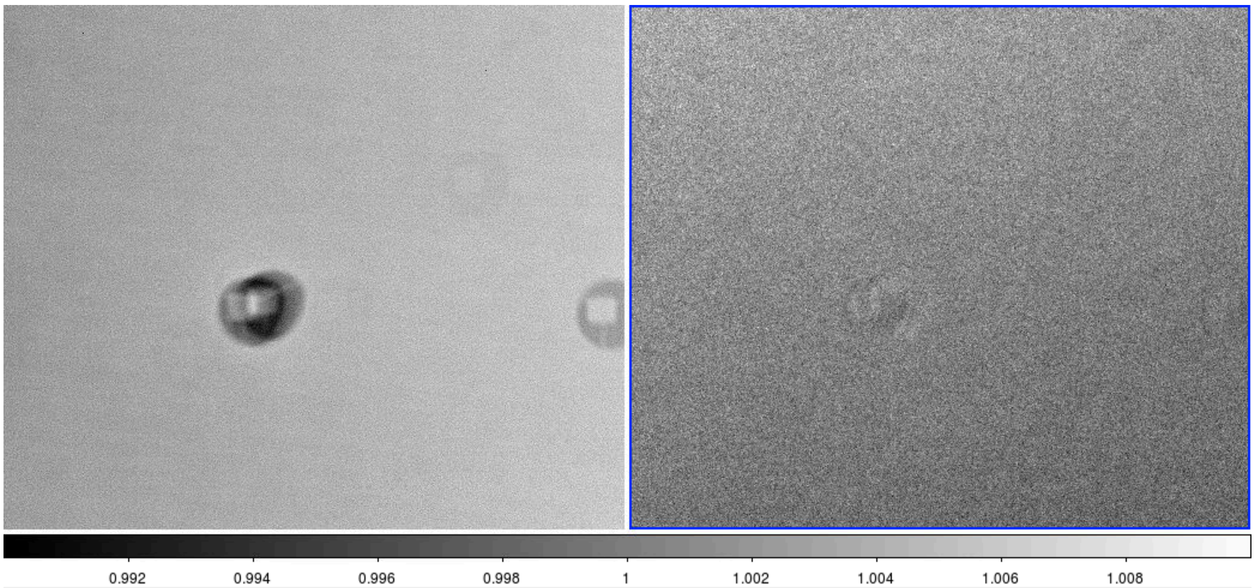

Shown above is a small variation of a dust doughnut on ccd 5 quad 4. The skyflat is at left, the ratio of skyflat to dome flat is on the right. While the spatial pattern variation is obvious, the magnitude of this effect is very small, near 0.1-0.2%. These seem to be sharpest dust doughnuts, presumably the closest surface to the focal plane.

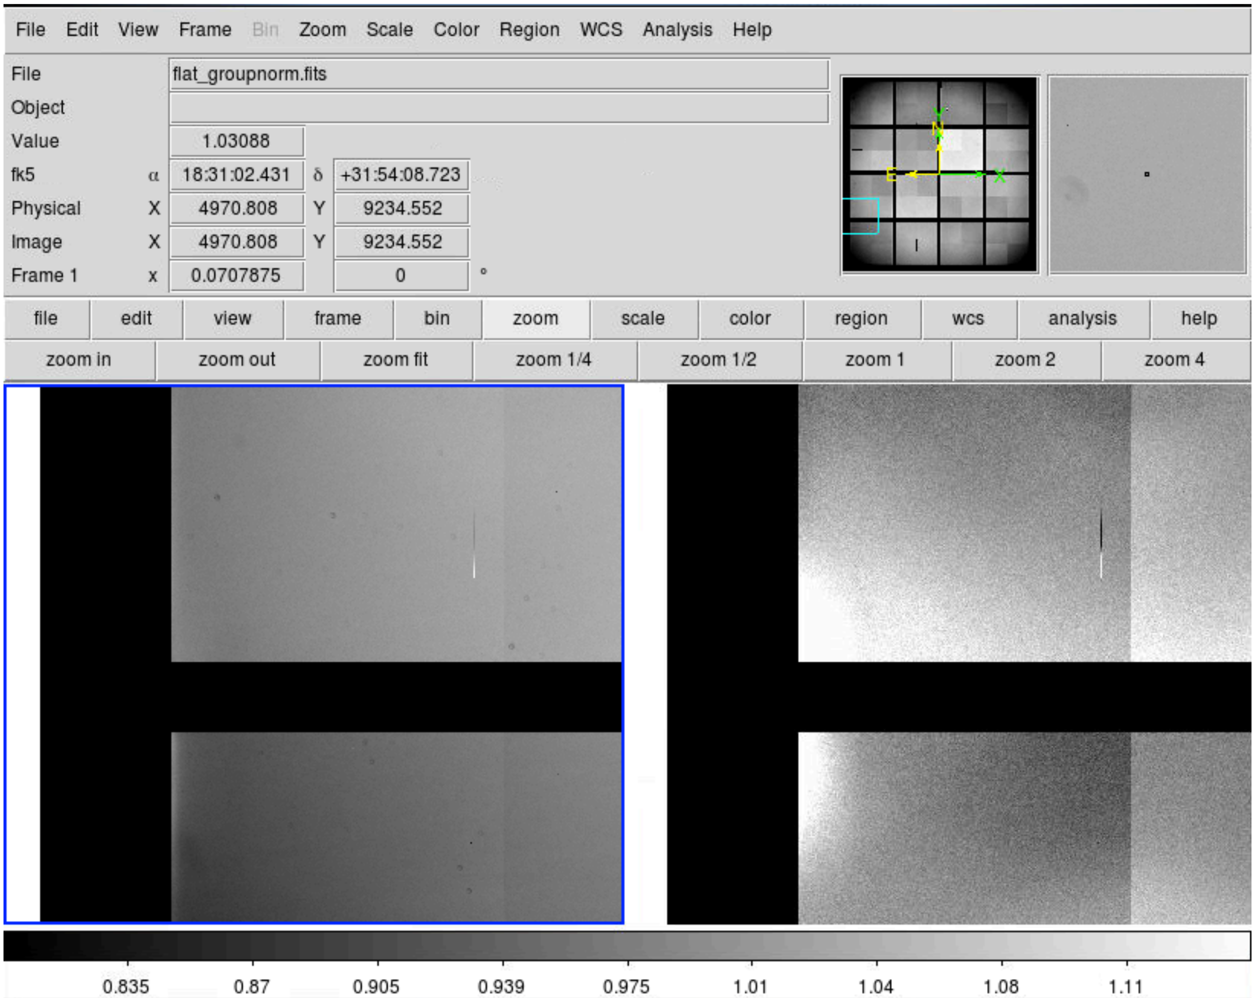

There are also some differences at the 1-2% level along the chip boundaries along the left and right sides (see the DS9 inset at upper right). Given that these are also the strongest examples of edge-scattering, this may imply some difference in scattering off the chip edges for these two illuminators (sky and flat screen).

3. Conclusions

Overall, the skyflats and the high frequency flats are substantially the same. The observed differences in illumination of dust, while detectable, seem to be a very small effect. The primary differences are the same patterns observed in the current starflat experiments. I note that the S/N of the skyflats is higher than that of the flat screen flats, however the moon phase was a waning gibbous on this night. The flat screen clearly has advantages in terms of predictability.