Introduction

Flat field correction is a key requirement for high quality data, whether for images or any other data. The science and art of flat fielding is challenging to say the least. With the introduction of CCDAutoPilot2, a tool is provided to simplify acquisition of the flat field data using different sources, including the sky itself. With this tool, you should be able to achieve field flatness better than 0.5%, not assuming any light pollution or other gradients, of course.

What is a flat field correction?

A flat field correction frame, or flat field for short, consists of evenly illuminating the imaging system, consisting of the CCD camera, filters if used and the optical tube assembly, to create an image of the non-uniformity. Since this non-uniformity represents a system transfer function, it must be divided into the data to remove the non-uniformity. This is different from the familiar dark frame subtraction since the dark signal adds to the data as a function of exposure time and temperature. Additionally, in order to not impact the magnitude of the data, the flat field correction must be normalized before dividing. This means dividing all the pixel values in the flat field by the maximum pixel value. Thus, a normalized flat field will have a maximum value of 1.0 and a minimum value as low as 0.0. Fortunately, the process of normalization and subsequent division is handled by the calibration software we use.

A flat field corrects for a number of things, some of which are familiar and some are less so.

- **Vignetting:** This is the light fall-off that occurs in any optical system. The result is the data is brighter at the center and dimmer at the edge. A typical fall-off for a well-designed F/9 system might be on the order of 3-5%; for an F/5 refractor, it can be 7-10%. Since vignetting might not be perfectly symmetrical, new flat fields are required any time the camera is rotated.

- **Dust:** Dust particles anywhere near the CCD chip itself will contribute varying size dark areas. Because they are usually not on the chip itself, their effect is one of a shadow. The shadow will have a hole in it for an OTA with a central obstruction such as a reflector or not as with a refractor. The size, and therefore the intensity of the dust shadow will be a function of the optical speed (F/#) of the system and the distance from the chip with those dust particles farthest from the chip having the lease effect.

- **Pixel variations:** Every pixel in the array has a slightly different sensitivity. While the difference is small, it is nonetheless real. A flat field will serve to adjust the lower sensitivity pixels to match the higher sensitivity ones.

As can be seen, the effects above go from very broad ones such as vignetting down to the individual pixel level. Well done flat fields can reduce the impact of all of these factors.

Why bother?

Flat fields are not strictly necessary for aesthetic imaging but if you wish to bring out all of the faint details of your target, you will need a uniform background. If you don't have a uniform background, you will end up setting the center "hot spot" to black and lose all of the data away from the center. Also, those pesky dust donuts will still be there even in the center no matter how meticulous you are at keeping your system clean. Incredibly small dust particles are very visible and by their very size are electrostatically bonded to optical surfaces. Anyone who has ever tried to clean a CCD cover slip can attest to this!

With color imaging, each filter will certainly have different dust patterns and could in fact have slightly different vignetting, depending on the optical speed of the system and how carefully the filters are mounted perpendicular to the optical path. Most serious imagers take flat fields for each filter used, as well as for a channel without a filter.

Obtaining a flat field

A flat field is obtained by uniformly illuminating the optical system and imaging the result. Sounds simple but it isn't. We have to worry about how uniform the illumination source is and whether the CCD is linear. (It isn't over its whole range.)

Uniform illumination can be obtained by carefully designed light sources such as light boxes where the individual light sources are properly placed. See http://www.hiddenloft.darkhorizons.org/notes/lb.htm for a link to such a reference paper on light source placement. There is one and only one location of the light source from the first reflector that guarantees uniform illumination. Everything else depends on diffusion and light scatter to achieve a uniform result.

An excellent source of uniform illumination is the sky itself at twilight. A paper written by, it turns out a college classmate of mine, Fred Chromey, demonstrates there is a position in the sky where the sky is most uniform over a couple of degrees. This point is defined relative to the sun's azimuth and elevation. See "The Flat Sky: Calibration and Background Uniformity in Wide-Field Astronomical Images", Chromey & Hasselbacher, PASP 108: 944-949, October 1996.

Unfortunately, taking a number of flats for a number of filters during the changing twilight illumination requires a fair amount of skill and patience. Things get especially frantic during the exercise, which is why many resort to the more stable light box. Nevertheless, many hold the sky is the best light source from which to build a flat field since its spectrum more closely approximates the night sky.

The CCD linearity enters the equation this way. On one hand, we want to illuminate the CCD with as high a signal as possible to achieve the best possible signal-to-noise ratio (SNR). The SNR is the square root of the peak value. On the other hand, at some point of illumination the CCD becomes non-linear. This non-linearity can impact the success of the flat field correction. While there are techniques for assessing the non-uniformity of a given CCD camera, a rule of thumb is to expose to approximately 1/3 of the total signal range. For modern 16-bit cameras, this means aiming for a peak level of around 20,000 ADU.

So now we need to figure out the exposure time for 20,000 ADU for each filter. And, if we are using the more desirable sky flats, this exposure changes as twilight changes. Even with a light box, we need to record and hopefully find later, the exposure times that were used for each filter.

Oh, one more thing. If we our exposure time is too short because our light source is too bright, we may have the impact of the moving shutter distorting our flat fields.

Lastly, since we are interested in dust donut removal, we need to insure we are in focus and in focus for each filter. This is necessary to clearly image the dust donuts for removal. Improper focus of the flats will leave vestiges of the dust donut.

So, we have to keep track of linear operation, maintaining focus, staying in the linear portion of the CCD response, avoiding shutter effects and get enough flat field frames to get a reasonable SNR. That is a lot to keep track of. Fortunately, there is an easier way – CCDAutoPilot2.

Flats, the CCDAutoPilot2 way

CCDAutoPilot2 provides an automatic exposure routine that dynamically adjusts the exposure time to reach your pre-determined ADU level. All you do is specify the longest and shortest exposure time you want to use, enter the number of flats, the filter and binning used and the target ADU level and it does the rest. Here is what the flats tab looks like:

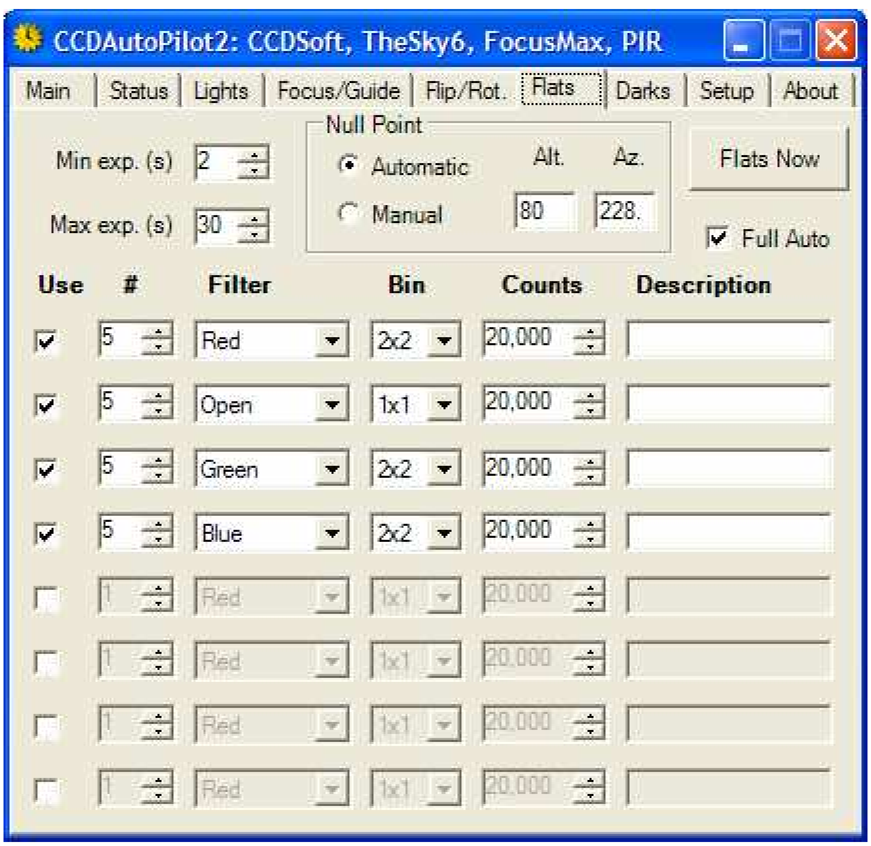

Enter your minimum and maximum exposure time at the top. Enter the number of flats, which filter and binning, your counts target and you are all set to go.

If you are using a light box, uncheck Full Auto and hit Flats Now. If you had entered focus offsets for different filters on the Focus/Guide tab, the focus point will be automatically adjusted as required. The exposure is entirely automatic. With Full Auto unchecked, the exposure is measured and adjusted one time at the beginning of each flat series. Of course this assumes your light box is stable which most are.

Now, if you want to use the sky as a light source, you can use the full power of CCDAutoPilot2. For this situation, you want to check Full Auto, so that each exposure is metered before taking a full flat. In this way, the changing light of twilight is compensated by the changing exposure time to achieve the desired ADU level. Here is a shot of the main tab:

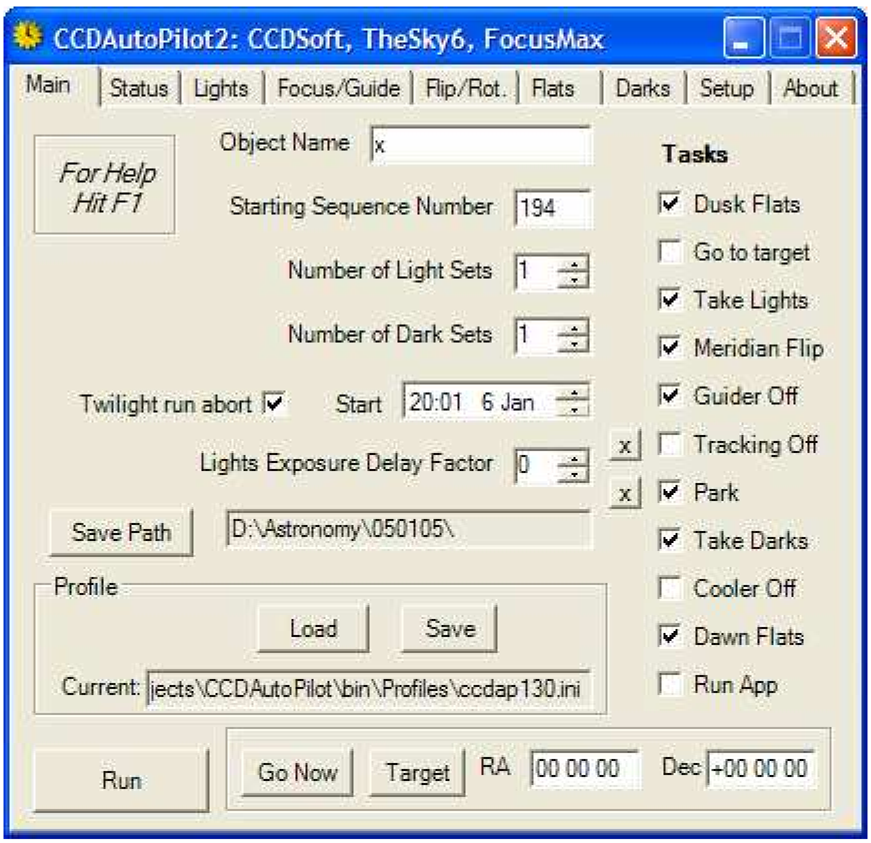

Note that there are options for Dawn and Dusk flats. Let's examine Dawn Flats. Assume the telescope is parked. The dawn flats routine will monitor the sun's altitude. When it gets to a suitable point, the telescope slews to the automatic null point. This is the point on the sky where the illumination is most uniform. Once at this point, the auto-exposure routine tries the various flat series, checking for sufficient light to achieve the desired ADU within the user specified exposure range. Once achieved, the frames for that particular filter/bin/ADU combination are taken until the series is complete. The routine then moves to another with the same exposure guidelines. Depending on the number of flats taken, the total flats run will complete before twilight ends. So the obvious question is: How many flats are enough?

How many flat frames?

Earlier, it was mentioned that the SNR is the square root of the peak value of the flats. This is measured in terms of electrons. So, for a given camera, the signal level in electrons is the selected ADU times the EGAIN in e/ADU. The EGAIN is in the FITS header for SBIG cameras. It may be measured for other cameras. See the instructions referenced at http://www.ccdware.com/resources/subexposure.cfm

A typical ST-10XME has an EGAIN of 1.4. So an ADU of 20,000 results in a signal level of 28,000e and the resultant SNR is the square root of 28,000 or 167.3. Taking 5 frames and using a median combine results in a SNR of a little over 300. A sigma combine would give an even higher result. When compared with faint detail where the SNR is on the order of 3 to 10, the flat SNR is high enough to no longer be a factor.

So, as a working number, 5 flats would seem to be adequate. Higher SNR could be achieved by a combination of dusk and dawn flats.

What about calibration?

With modern cameras properly cooled, there really isn't a need for dark frames. Review the calculator mentioned above and consider the dark signal flux in e/min. Divide by the gain and you can see the dark current ADU contribution is much below the 20,000 count level. Consider for example the ST-10XME operating at a temperature of -20°C. According to the calculator, the dark signal flux is around 3e/min or 1.5e for 30sec. A typical gain for the ST10 is 1.3 so the average "dark current" is around 1.2 ADU. Assuming a hot pixel is 10x that for 12 ADU. The noise for a 20,000 ADU signal is the square root of 20,000 or 141 ADU. The inherent noise in the light field is over 10 times the worst case dark current. Just subtract a high quality master bias frame and that is all you need to calibrate your flat exposures with exposures of up to 30 sec. Since bias frames are dominated by read noise, a master should be made up of a large number, perhaps 50 or so.

Another reason to subtract bias is to reduce the error due to the offset represented by the average bias level. The ST line of cameras mostly has an average bias level of a relatively low 100 ADU. But other cameras have higher levels. The STL-11000 has around 700 ADU. Others can be considerably larger. By subtracting the bias, you increase the accuracy of your flat field correction.

Summary

Hopefully you now have a better idea of how to obtain good flats and how CCDAutoPilot2 can help. Some experimentation may be required to achieve the optimal number of sky flats but it sure is nice to know that you are using a light source that is more closely related to the sky you are imaging at night. And even if you use a light box with the Flats Now feature, CCDAutoPilot2 takes the hassle out of any flat field technique you use.