Chapter 8

The Art and Science of Flatfielding

Beginners have more concern when dealing with the concept of flatfielding than for any other image processing step. At the same time, flatfielding is probably the most important calibration operation, and so rightfully deserves careful consideration. We have separated this step to give you details of what is actually being accomplished, guidelines for producing the best flatfield frames, and designs of practical flatfielding assemblies.

We mentioned in an earlier chapter that flatfielding is necessary in order to correct for both defects in the pixel-to-pixel response and to remove optical aberrations such as vignetting that may be present in your telescope system. However, we did not advocate any particular method to accomplish this flatfielding task.

1. Why Flatfielding is Necessary

There are several factors that modify the final response of a CCD to incident light. We consider these in order working from the front of the telescope. You may have vignetting or scattered light due to problems in the optical path. The light then can pass through a filter, which modifies the spectral response of the system. It will also modify the amount of light transmitted because even at the center of the band pass light is absorbed. In addition, the filter may be smaller than the light beam, might have non-parallel surfaces that add ghosts and reflections, and certainly there is a place where dust, insect parts etc. can accumulate. After passing through the filter, light will be either stopped or permitted to pass by a shutter. Finally, the light enters the CCD camera itself, usually through a front entrance window and optionally through a cover glass before finally reaching the actual CCD detector. Therefore, you can separate the light modification into two parts: modifications external to the CCD, and those modifications internal to the CCD. Flatfields can correct both problems, and in fact can be used in tandem in some cases to improve the flatfielding even further.

The net result is a method to remove as many of these modifications to the incident light as possible. Figure 1 shows some of the external effects that contribute to nonuniform response. This image is of a uniformly illuminated screen from an amateur telescope. Notice that the light is brightest near the center, but falls off towards all corners; this is vignetting, caused in this case by the use of a focal reducer in the telescope. The vignetting pattern is not perfectly centered on the CCD; this is due to the CCD not being at the center of the optic axis (if it is mechanically centered, then most likely the telescope collimation is slightly off). In addition, you can see two faint dust 'donuts' near the top of the image, caused by dust on the entrance window into the CCD. Note the one larger dust donut located towards the lower left. This one has much larger diameter and is correspondingly fainter; it probably arises from the filter. More discussion about dust donuts is given below.

Figure 2 is a plot along one row of a professional CCD, again illuminated by a uniform light source. In this case, the light source was strong enough that if all pixels responded equally, you would not be able to see the scatter on this scale. Instead, we see not only a low frequency wave, caused by uneven thinning of this CCD, but also pixel-to-pixel variations in the few percent level. This high-frequency variation is caused by differences in the shape and collection ability of individual pixels, and would introduce equivalent errors in your final photometry if not removed.

2. What Flatfielding Does

Most flatfielding techniques for CCDs assume that the average response of the detector is linear, that for zero input flux you have zero net signal, and that all you are doing is correcting for different gain factors of the pixels. If the detector is nonlinear, as is the case for most infrared detectors, or if there is a nonlinear portion to the light response as for antiblooming CCDs, the process can be considerably more complex. We will avoid those cases for the time being.

If a detector is linear, knowing the response to zero flux and to a uniform flux over the detector yields two points. A straight line through these two points has a slope that is equal to the gain of the detector in some arbitrary units. Each pixel of a CCD is a separate detector with its own value of gain. In the flatfielding process we are trying to balance the gain, so that the response to a uniform flux is identical for all pixels.

In addition to balancing for uniform illumination, you must also balance the wavelength response of each pixel. While through the mid-part of the spectral response of the CCD most pixels respond about the same, at the blue and red extremes you will find differences between neighboring pixels. This may be due to a thinning process, differing thickness of antireflection coatings, or impurities in the silicon. So you need to have a light source for the flatfield that is uniform in wavelength response as well. Such light sources are typically blackbodies, such as heated filaments, since they are smoothly varying without bumps, spikes or wiggles.

Let us assume a uniform illumination on a CCD, with an average response of M ADU. Then, for any given pixel, assume its response is N ADU.

$$N = g_i M$$

where g_i is the gain for pixel i. Note that, if you use a uniform illumination on the CCD, you can directly solve for the gains, since

$$g_i = \frac{N}{M} \tag{8.1}$$

This means that the gain is just the pixel response to the uniform illumination, divided by the average pixel response, or in other words, just the normalized flat field image. Therefore, to correct the pixel response to any incident flux on another exposure, you would just perform

$$M_x = \frac{N_x}{g_i}$$

$$M_x = \frac{N_x}{f}$$

where f is the normalized flatfield image. In a normalized flatfield image, the average pixel value is 1.000, with any pixel with a value below the mean will be assigned a value less than one and those greater than the mean assigned a value greater than one. So division of any image by the normalized flatfield image will "brighten" the pixels with a gain that is too weak and "dim" those with a gain that is too strong. This is an extremely simple operation, and works remarkably well as long as the pixel response is linear.

Remember that when processing your images, you are flatfielding all like images with a single master flat. This master flat should have sufficient signal/noise that it does not add significant noise to your raw data frames. Usually, you expose for flatfields to about half full-well charge capacity, so that you have lots of signal, and average several such frames together to get an even larger S/N master image. For astrometry, one can often get by with a very crude flatfield, perhaps not even uniformly illuminated, as long as you can correct the high-frequency pixel to pixel gain variations. For one percent photometry with reasonable psf sampling, a reasonable goal is to have a S/N of at least 1000 in the flatfield; for millimag photometry your flatfield has to be substantially better. Many amateur CCDs with small pixels have a very small full well, so that you need to average many flatfield frames together to yield these required S/N values.

You must be careful with flatfielding not to introduce systematic errors. For example, if the illumination used to make the flatfield image is not uniform, you may introduce a gradient across your processed images. This gradient will be constant for a single night but subsequent nights might have different illumination patterns and gradients. Even if the gradient remains constant over extended periods of time, the derived magnitudes will differ depending on whether the object is on one side or another of the frame. This is particularly bothersome if you are trying to calibrate a variable star field, or if you are using published comparison stars that are scattered over a field, because the differential magnitudes will have induced scatter.

When you take flats through a set of photometric filters, you are actually also adjusting, to first order, the spectral response of all pixels to light with the spectral distribution of the flatfield light source. If you are using projector flats, that light source can be quite 'red' in comparison to the night sky or any astronomical object. To first order, the flatfield has direct impact on the transformation coefficients you obtain for your system. If you use a dome flat on one night, and a twilight flat on another night, you will have different transformation coefficients.

3. Color Dependence of Flatfields

Shown in figure 3a and 3b are two flatfields from an early thinned CCD, the Texas Instruments 800x800. The flat on the left is the response of the CCD to blue light, and the one on the right is the response to red light. Notice that there are large differences between the two flats; flatfields are definitely color-dependent. Thinned CCDs exhibit more of this wavelength dependence than front-illuminated CCDs; an equivalent front-illuminated flat set is shown in figure 4a and 4b. For back-illuminated CCDs, the thinning process, is nonuniform since it is a chemical etchant. In some regions, the epitaxial material is overthinned and part of the light sensitive region has been removed, thereby lessening the overall QE. These regions can be seen in both the blue and red flats. In other regions, the CCD was not thinned enough, and is less sensitive to blue light. These regions show dark (lower sensitivity) on the blue flat, but often show light on the red flat because they might actually be more sensitive to red light with the thicker layer. Blue photons are absorbed closer to the front surface of the CCD and are therefore more affected by imperfections in that surface. The grid pattern seen in the red flat is due to the way the photomask was replicated on the silicon when manufacturing the CCD. Note that a front-illuminated CCD does not show nearly as much wavelength dependence, but it is also not as sensitive to blue light compared to the thinned CCD. Front-illuminated CCDs usually do not have antireflective (AR) coatings. In general, front-illuminated CCDs have cleaner large-scale flats, though again the biggest variation will be seen in the blue. Both types will have the pixel-pixel variation.

For these reasons a separate flat field frame should be used for each filter. The flatfield illumination source is usually a blackbody of a particular temperature. The mean wavelength of the light passing through the filter can shift the pixel response. So the illumination source should have the same color temperature as the object to be measured. This sounds particularly difficult. A tungsten lamp is about 3800K, yet a G-class star like the Sun is 5000K, and upper main sequence stars can be much hotter yet. How do you correct frames that contain stars with differing temperature? You can't, exactly. Luckily, this is a second-order effect, and the usual procedure is to use a light source that is close to a typical star temperature, like 6000K, and hope that the shifts are small. In fact, the usual target temperature is that of the night sky, so that the sky background can be subtracted nearly perfectly.

Holst[2] gives tables of color temperature of various sources. Table 7.1 below is a shortened version.

Because flatfield exposures have finite integration and readout times, you will detect cosmic rays on each frame; these can be removed by taking multiple flats. One rejection technique is to use a median flat, where each pixel is the median of all values for that pixel for all images in the flat set. Another technique is to use an average, but to reject the highest pixel in the set (which is presumably the cosmic ray). A third technique is to reject pixels that exceed the mean pixel value by some large factor. This retains most pixels in the averaging process while rejecting only those few pixels in all of the images in the set that have been hit by a cosmic ray. We usually use the average plus rejection of the highest pixel technique, but you should investigate all of the possible methods to decide for yourself which works best for your images.

Flatfields also correct for aberrations in your optical system. The most common aberrations are vignetting (due to improper baffling, an obstruction in the optical path, or the addition of a focal reducer in the system) and dust (on the filter, dewar window, cover plate, or on the CCD itself). These obscurations, in effect, lessen the QE response for some pixels. The pixel response can be adjusted to compensate. Note, however, that the dust pattern will change from filter to filter (another reason not to use a single flatfield for all filters) and if the filter does not reposition identically each time it is placed in the light path. Dust in reflecting telescopes results in 'dust donuts'; round halos that are out of focus, with the diameter of the donut dependent on the f-ratio of the telescope and the distance of the dust from the CCD. These donuts do not always disappear after application of the flatfield correction. This is especially true if there is some scattered light due to, say, the nearby moon. An example of the effects of nearby moonlight are shown in figure 8-5. The best solution is to clean your optical elements scrupulously. Dust donuts are similar to a prenumbral eclipse; the amount of light is decreased, but a flatfield should be able to compensate. However if the dust is close enough to the CCD that in the converging beam there is an umbral eclipse, the flatfielding is no longer effective and the pixels beneath the dust are effectively lost. This typically occurs when there is dust on the CCD itself, or on a cover plate. You should keep these items scrupulously clean.

You should take flatfields with the telescope properly focussed. Dust donuts will not be removed properly unless this condition is met. If you are taking sky flats, having the stars focussed improves the ability to reject them in the median-reject process. If you use out-of-focus images for lowering the saturation level, you need corresponding out of focus flats.

Note that if flats look different in different filters, then the pixel wavelength response differs. This means that in the strictice sense, the transformation coefficients should be determined on a per-pixel basis. This is a limitation to using CCDs for accurate photometry. Using flats in each filter does a first-order correction for pixel response and is another reason why flats are necessary.

How often do you need to take flats? They should definitely be retaken if the optical configuration has been changed (for instance, removing and replacing the CCD camera). In dusty environments, dust specs will appear on a daily basis and taking flats nightly will lessen their impact. We have on occasion seen new dust features occur during a night, so sometimes you have to use subsequent night's flats to correct the previous night's images. If you have a closed telescope tube, such as for a Schmidt-Cassegrain configuration, and if your filter wheel/camera configuration is light- and dust-tight, you may not need to take flats very often. The time interval over which you can use a single set of master flats is something each observer will have to determine. For our open-tube R/C telescopes at NOFS, we take nightly flats whenever possible. We recommend starting with nightly flats and then relaxing this requirement as you gain experience with your particular setup.

The proper exposure length for a flatfield is important. An exposure which attains half the saturation level is a good starting point. Ideally, one should use the same exposure depth as the object of interest, so that even minor linearity errors are lessened. However, the object of interest may be low signal/noise in the final image, which means individual flatfield images would also be low signal/noise and many of them would have to be combined. For this reason, most people just assume the CCD is linear and get good signal/noise in each individual flatfield image.

In general, the best results are achieved if the flatfield light source has the same spectral distribution as the object of interest. However, for faint objects near the sky background, background subtraction is the limiting factor on the photometry, and so getting good uniformity of response in the sky-illuminated pixels is even more important than balancing the response for the object of interest. So most observers try to match the flatfield light source to the sky background color. For a dark site, twilight flats do a good job. In an urban area, where the night sky is illuminated by street lights, getting a good flat that allows scientific programs that work close to the sky background limit will be a challenge. Perhaps a cloud flat will be effective under these conditions, and is a good use of an otherwise lost night; just watch out for rain!

4. Types of Flats

Producing a high quality flatfield image requires even illumination at the front of the telescope with a light of a broad wavelength response. There have been a variety of techniques developed over the years to achieve this. These include:

- twilight flats
- master sky flats
- projection flats

Each of these techniques will be discussed in detail below.

4.1. twilight flats

The easiest technique is to use twilight flats. These are images taken during twilight with the telescope pointed towards the zenith. There is a narrow window of about 20-30 minutes when the twilight sky is still bright, but does not saturate the detector, and not so faint that many stars are detected. Star images can removed by tracking the telescope during each exposure, but shifting the position by an arcmin or so between each exposure. The idea is to combine these images in such way as to exclude the stars. This is done by comparing the pixels in each image that are located at the same physical position on the CCD chip. These corresponding pixels are averaged after excluding the brightest one. Because the telescope was moved between exposures it is unlikely that any given pixel contained the light of a star in more than one image. So n rejecting the high pixels before averaging the stars will be removed because they do not fall on the same pixels in all images. You can also achieve the same effect by using a median combine of the individual flatfield images.

Twilight sky flats sound remarkably simple. They use a 6500K light source (i.e. the Sun), have even illumination, and don't require any special equipment to perform. They do a very good job of correcting images at the one percent level. However, this technique does have its faults. As the brightness of the twilight sky changes there is a narrow time window in which to collect the required images before the sky is either too bright (morning twilight) or too dim (evening twilight). If there are several filters that require master flats, there may not be enough time to take enough exposures per filter to get the required signal/noise. If the Milky Way happens to be passing overhead during twilight, you may have many stars in each flat, and even shifting between flats may not be sufficient to remove those cases where two stars fall on a given pixel. This may mean taking more flats, and rejecting more high pixels to remove such instances. The twilight sky is not flat, though it is very close to even illumination for exposures taken near the zenith. For wide field telescopes, though, there may see a gradient across your frame. Chromey and Hasselbacher (1996) measured the gradients in the surface brightness of the twilight and night sky. They found a null spot near the zenith where the gradient reaches zero. This null spot position may vary on a night to night basis, but pointing towards the zenith is a good generic location. Typical results are shown in figure 6. So for common CCD fields of view, taking twilight flats near the zenith will give the smallest gradient across the image. Note that for wide-field systems, even pointing at the null point means that there will be a radial gradient across the image. There is also significant polarization in the scattered light, and if there is some optical element that is polarization sensitive, the flatfield may be systematically wrong. There are peculiar effects that occur near sunset, such as the illumination of the atmospheric sodium layer, that may effect flats. Clouds may cause uneven illumination. You may not be able to get set up by sunset to take flats, or may not be in focus until you can see stars.

Chasing the exposure times while the sky darkens can be a challenge, and means that high-pixel removal requires correcting each exposure to some common average value (usually by determining the mode for each image and scaling appropriately), which is an imperfect process. One of the hardest things with twilight flats is determining what the next exposure length will be. With a SITe 1024x1024 CCD that has a 30-second readout time, we find that we can take 15-18 flats during the twilight window. Usually it's best to take U and B flats first, since the QE of the system is lowest in those band passes and therefore they require the brightest skies. The order for VRI depends on your personal preference and which filter is more important for your project, as all three take about the same basic exposure time.

If you leave the exposure time constant and just take a series of flats through a single filter as the sky gets dark, you end up with only a few flats that have similar signal/noise that you can combine. Flats with higher signal levels will be close to saturation. Flats with lower signal levels actually degrade the master flat signal/noise, so you need to restrict the number of such flats that you combine. We generally do not recommend the single-exposure-level kind of twilight flat.

If the exposures are made when the sky is still fairly bright the exposures will be short and the camera shutter will create noticeable vignetting. This comes about because a mechanical shutter can not act instantaneously across the whole CCD chip. This is not a real issue as long as the exposures are long compared to mechanical response time. So for an evening twilight sky you need to wait until the sky dims enough that exposure times of 3-5 seconds are acceptable. Then shoot quickly before the sky gets too dark (say, exposures of 30-60 seconds) and too many stars appear in the images. Note that the longer the read time of the CCD, the more twilight is missed and the harder it is to judge the exposure increase. Tyson and Gal (1993) gave a more analytical model for a guide to exposure lengths as a function of twilight time. They also provide on an anonomyous ftp site the Fortran source code necessary to generate an exposure guide for particular CCD readout times. Their experimental evidence shows that the rate of decline of the evening sky brightness is a constant at all wavelengths, and amounts to about a factor of two decrease every 3.3 minutes. If you take one exposure, you can then calculate how long the next exposure should be in an iterative fashion.

One consideration is that any gradient in the evening twilight sky will be reversed in the morning twilight sky. Therefore, you can form a mean twilight flat for both evening and morning, and ratio them to see if any gradient exists. If there is a gradient, the average of the evening and morning master twilight flats will reduce the effect. Avoid fields that contain any bright stars, as these will leave the outer wings of the seeing disk profile when frames are combined. Taking flats on nights when the Milky Way passes directly overhead at twilight takes care; usually the best method is to use the other twilight (e.g., if you have the Milky Way overhead in the evening, then take flats during morning twilight and vice-versa). Nights with a bright moon will cause gradients in the flats, increasing as the sunlight fades and the moonlight becomes a larger portion of the illumination. Like with the Milky Way, you can mitigate this effect by using morning twilight flats if the moon is less than full, or evening twilight flats if the moon is waning. Full moon is always a problem!

A modification of twilight flats is to place a diffusing screen in front of your telescope, such as a piece of white plexiglass. This removes stars from your image, lessens gradients since they become averaged over the entire telescope aperture, yet uses the twilight sky as the light source. A colorless diffuser is preferred as "white" to the eye may not be "white" to the CCD. Some observers suggest the use of a T-shirt as a diffuser.

One problem with twilight flats is the possibility of light leaks, especially for open-frame telescopes. The observatory dome can be used to block the western horizon for evening twilight flats to reduce the possibility of scattered light getting into the system. Having a bright moon in the sky during twilight flats may both change the gradients in the sky for flats as well as causing light leaks and reflections in telescope tubes. We generally avoid twilight flats near full moon, using either archived flats or dome flats during those nights.

In general, do not take twilight flats when there are clouds in the sky. The exposures are relatively short, and any structure in the clouds will impact the flatness of the flatfield. In addition, at sunset and sunrise the clouds will often be colored and not have the right spectral response. If you must take flats with clouds, take multiple flats to average out any effect. However, for some wide-field systems, uniform clouds offer the opportunity to get good flats since there is less twilight gradient. Park (2001) mentions that for the LOTIS system they have used fog-flats. They often get sky conditions where there is a fog layer lying a few hundred feet above the observatory. In general, this is not a recommended technique because the illumination source is often ground-based lights, that can be very non-blackbody in character (e.g. containing emission lines), but the results can be quite good if care is taken.

4.2. Master Sky Flats

Another flatfielding variant is to use the raw data frames themselves. Application of a median filter to the frames taken during the night will average out the sky gradient and polarization problems. While there is never a perfect match between color temperature of the flatfield illumination and the color of measured objects, with master sky flats you can get a perfect match between the flat and the night sky. Since the night sky is the dominant noise source for faint objects, at least getting its flatfield correct is enormously helpful. Tyson [17] gives a good description of this technique.

Master sky flats work best for deep images of relatively blank fields. This gives the maximum sky signal per image, so that readnoise and digitization noise are not the dominant noise sources. There needs to be many such frames in order to obtain sufficient signal/noise in the flatfield image so the flatfielding process does not add noise to the data frames.

Making master sky flats from object frames which have extended objects, such as galaxies or comets, will not normally be successful. No amount of median filtering can remove the extended background source.

Observatories have found sky regions that contain few stars, referring to these regions as "blank fields". A table of these (from KPNO and ESO) is shown below. You can use these fields for either master sky flats or twilight sky flats.

4.3. Projection Flats

The term "projection flats" includes any flatfield images made using an artificial light source. Such flats are extremely convenient. They can made during the night, or even during daytime (with a light-tight light box). The light source can be regulated to give you consistent exposures, and can be adjusted to be just bright enough in each filter. As with any flat-making technique, projection flats have their problems and complexities. However, a projection system has some real advantages and is almost a necessity when performing system tests. We recommend every observatory have a projection flat system, if not as their primary flat-taking device, then as a backup system and for lab tests.

The basic types of projection flats are:

- integrating spheres
- dome flats
- light boxes

These will each be discussed in turn below, after considering the source of the light used in the projector.

4.4. light sources

Light Emitting Diodes (LEDs) are extremely interesting devices. They last almost forever, are more efficient than incandescent sources, require only milliamps of current, and run cool from low voltage DC sources. They can be made of different materials to emit light from blue to near infrared. At the same time, they do have problems. First, the typical response for a given LED is only about 50 nm wide; see the chart in figure 7. While its possible to use LEDs to give reasonable light output in several band passes like B, V and R, there has not been a U-band LED manufactured to date, and the I-bandpass is not well covered either. So a general-purpose flatfield box would have to be made with a multitude of different LED types, and still not produce light in some of the important Johnson-Cousins bandpasses.

A single LED has been produced that covers the 380-750nm bandpass. Called a *white LED*, this is actually a blue LED with a phosphor coating to shift some of the blue light into redder wavelengths. It appears white to the eye, but as seen in figure 8. The output is definitely not uniform over the entire range. LEDs can be pulsed to gain higher, temporary output. With pulse duty cycles of 10 percent, the LED can run overcurrent by a factor of 4 to 5. The LED output is a strong function of current, so a constant current source is recommended. Even under constant current conditions, the light output varies for a few hundred msec after turn on as the junction temperature first rises until thermal equilibrium is reached, as can be seen in figure 9. The spectral distribution of a typical LED will also change with the ambient air temperature; typical peak wavelength changes are about 0.2nm/K. Another problem with LEDs, especially for the inexpensive ones from surplus electronic houses, is that the light output will be quite different from one device to another. Balancing an array of them to give uniform light can be problematic. The peak wavelength also will vary, especially for the blue LEDs. When using LEDs for linearity tests, it is best to leave them turned on continuously, and to do the tests either inside a lab or when the outside temperature is reasonably stable. We use LEDs for specific calibration purposes, but not as general-purpose flatfielding light sources. A great web site for gaining understanding of LEDs and finding links is http://www.pioneernet.net/optoeng/LED_FAQ.html.

Tulloch (1996a) indicates that the temperature coefficient for LED output is about 1% per degree K. He describes a temperature stabilization technique using a thermistor mounted near the LED as part of a feedback loop. Tulloch (1997) describes another novel light source for linearity measurements. A specialized LED is manufactured by Integrated Photomatrix Ltd. (IPL) that contains a green LED, a monitor photodiode and a transconductance amplifier. The output of the on-chip amplifier is proportional to the LED optical output, so feedback loop can be made to stabilize the output.

Fluorescent lamps can be purchased that come close to matching the daylight spectrum. An example is shown in figure 10. Note that the continuum is close to daylight, but that there are emission lines corresponding to the mercury light source. Emission line sources should be avoided when making flatfields. Also, fluorescent lamps have a 60Hz frequency power source which can couple into CCD systems. We do not recommend the use of fluorescent lamps as flatfield light sources.

Some observers have suggested the use of Xenon flash lamps as a flatfield light source. These have quite blue response and are of course very bright. With a pulse-type output, one could adjust the exposure by the number of pulses applied to the lamp. A typical spectrum from a Xenon flash lamp is shown in figure 11. As you can see, there is output over the entire CCD sensitivity range. However, the curve is dominated by Xenon emission lines. In general, light sources with narrow band emission should be avoided as they will not mimic the night sky brightness and will systematically affect the response through any given filter. We do not recommend Xenon flash lamps.

The simplest filament lamp is a tungsten incandescent lamp. These come in all sizes and shapes;

particularly attractive ones for amateur use are the grain-of-wheat bulbs. These are available in 1.5V and 2.5V sizes, much like Christmas tree light bulbs but smaller and with long leads. Several electronics surplus houses carry these lights. They are pretty typical 3000K light sources, with typical black-body emission over the entire spectral range. Like other black-body sources, filters can be used to mimic much hotter lights. You should run the lamp at its full rated voltage. Turning the voltage down makes the lamp redder, the wrong direction to go. If you have too much light, then either insert a neutral density filter or use a lower-wattage lamp.

Massey and Jacoby[4] describe KPNO's projection system lights. KPNO uses two banks of studio lamps. These lamps have housings with built-in slots that can be used as filter holders. The high bank consists of four Cool-Lux FOS-4 12V50W (quartz-halogen) lamps. These have an average lifetime of 4000 hours as delivered from the factory. They modify them by applying an aluminum coating to the inside surface of the reflector; otherwise light from 800-850 nm would pass through the rear of the bulb housing, causing a dip in the spectral light distribution. The low bank has four Cool-Lux FOS-5 12V25W lamps. In front of each of the low bank lamps is a stack of filters – 4mm of BG-34 blue color balancing glass and a 0.5 ND filter (to allow the voltage to be as high as possible to increase the color temperature). The bulbs, projectors and filter holders are all purchased from Cool-Lux Lighting Industries, Inc., North Hollywood, CA 91606.

A new Xenon High Intensity Discharge (HID) automobile headlight is currently on the market. They are an arc-discharge tube that contains xenon gas and mercury as well as metal halides (salts). These lamps require ballasts and AC voltage. They are designed to be ignited and then left on, rather than blinking, so for flatfield use they require a shutter. Note that, since these lamps contain mercury, they commonly have a UV cutoff glass that prevents U-band light from emerging. A headlight system using HID lamps is quite expensive at the time of this writing, and so it is unlikely that you would want to use an HID headlight as a flatfield projection light source.

Another kind of automotive light is just a typical halogen bulb with an added blue filter (such as the Sylvania CoolBlue). These are similar to a studio lamp with an outboard blue filter, and have some potential interest as a flatfielding light source. Automotive lights are leaning towards bluer response and so bear watching in the future. Their main problem is that the output light pattern tends to be beamed rather than omnidirectional.

Our recommendation is to use a filament lamp of some sort; low voltage ones offer the potential of being controlled with a voltage regulated power supply. Use a minimum of 4 bulbs to smooth out the light pattern, and consider using a blue filter when creating blue flats.

4.5. integrating spheres

Integrating spheres are a common flatfielding scheme for spectrographs, where very uniform light is needed but within a relatively small beam. The concept is very simple: a sphere has its inner surface coated with a granular reflective material (usually lambertian, where the reflection can be at any angle irrespective of the incoming light, rather than specular, where the light is reflected with an angle that depends on the angle of the incoming light). A hole is cut in the sphere to allow light to leave, and the telescope is pointed towards that exit hole. Off to one side, another hole allows light to enter the sphere. The light bounces around the inner sphere surface several times before finally exiting, thereby randomizing the direction the light came from. The result is a perfectly uniform exit beam. Properly designed, an integrating sphere will give the most uniform light pattern of any projection system. A crude diagram of an integrating sphere is shown in figure 12.

Typical integrating spheres have an output hole that is about 1/3 of the diameter of the sphere. That is, if the telescope aperture is 10cm, then a 30cm sphere is needed. These spheres get awkwardly large as the telescope diameter increases, but are probably an ideal flatfielding source for wide-field optics. Integrating spheres can be purchased commercially in arbitrarily large size, but they tend to be very expensive. An integrating sphere can be constructed that will work adequately using paper mache or fiberglass around a balloon or beach ball (these are quite inexpensive from local discount stores). Once the material sets up, remove the balloon or ball, cut the appropriate light source and exit holes, and paint the inside with a white paint, perhaps textured but at least flat, to reduce specular reflections.

Labsphere, Inc. (http://www.labsphere.com) has technical guides that give information about integrating spheres. These devices are typically used in industry for measuring the flux from a light source such as a laser; we use them in the reverse direction for astronomy.

4.6. dome flats

Dome flats are made by illuminating a white screen mounted on the inside surface of the dome. The screen is illuminated either by projection lamps on the end of the telescope or projectors mounted elsewhere in the dome. The telescope is pointed toward the screen and exposures are taken through each filter. Dome flats have a number of advantages. Remember with twilight flats that the exposures for a given filter were often of different depths, so that you had to adjust each frame by its mode before combining so that cosmic rays could be properly removed. With dome flats, the exposures use a constant illumination source so that the individual depths are very similar; you can often combine the frames without any kind of mode adjustment. Another advantage is the lack of time constraints; you can take the flats anytime that it is dark. The lamp intensity can be adjusted so that exposures are relatively short. Usually four to eight lamps are used to make the screen illumination as uniform as possible. The lamps must be shielded to prevent their light from reaching the telescope optics directly. To approximate the desired color the lamps should be run near full voltage (therefore, as hot as possible) It may be necessary to include a neutral density filter to achieve proper exposure times. The lamps should be on an adjustable, regulated power supply.

Dome flats have a number of disadvantages. For example, it is hard to get the projection lamp very hot, so the light source is always too red. Uniform illumination is difficult, and because the screen is not far from the front of the telescope, the reflected light is not coming from infinity and any dust is illuminated differently than it will be from the night sky. Proper screen material is more difficult to obtain than you might imagine. One company produces Soric, a cloth-like material similar to the home projection screen but with better UV reflectivity and more Lambertian in the reflection directivity. Many observatories including KPNO use a painted surface. The surface has to be stiff so that flexure won't flake the paint off. KPNO found that plywood and fabrics don't work well, but an aluminum sandwich material (that is, like cardboard) seemed to be suitable.

Labsphere has a Guide to Reflectance Coatings and Materials on the web (http://www.labsphere.com) that is a very useful reference when searching for paints and other materials for flatfield screens, and when deciding how to apply such materials to the supporting surface.

The trick then is to find a paint that has high UV reflectivity and which is suitably Lambertian. Most white paints are actually quite good absorbers of blue light; the human eye is just not sensitive enough to notice this 'feature' of titanium oxide.

Wu developed a barium sulfate paint formula that seems to work well. The recipe is as follows. To make one pint, heat 50 ml of distilled water to 52-66C (125-150F) and slowly add 2.25 grams elvanol polyvinyl alcohol (Grade 72-60), mixing until all the alcohol dissolves. To this mixture add 150 ml of distilled water and 200 ml of 200 proof anhydous ethyl alcohol. Then slowly add (stirring continuously) 227 gm of USP grade barium sulfate to the mixture. In practice it seems best to mix this paint just before use. We applied it to an undercoat of Krylon 1502 flat white paint using a Pace air brush with a No. 2 head.

If U-band flats are not needed, it may be possible to use normal flat white paints and little filtration of the projection system. Even at B, however, it is useful to go to the extra effort and get the colors right.

4.7. lightboxes

A variant of projection flats is the use of a transmissive system rather than a reflective system. The simplest version of this is to just cover the front of the telescope with some translucent material such as opal glass, translucent plastic, or even non-opaque paper, and let ambient light pass through the medium into your optical system. More elaborate systems provide an external source of light, much like the dome flats described later. The main difference is that you are dealing with the transmission of light rather than its reflection.

One of the nicest light boxes that we have seen was built by Gregory Pyros (http://www.pyros.com/greg/html/l Another good one can be found in Berry and Burnell (2001). The opal plastic can be obtained at a plastics supply shop; McMaster-Carr (http://www.mcmaster.com) also supplies Acetal Copolymer Sheets in Opaque White color at about US$7 per foot-square sheet. There are several other lightbox designs that can be found on the web.

Many light boxes are made out of polystyrene foamcore board. Richard Seavey mentioned in an email that one good method of cutting holes in this foam is to use a hot wire. Just put a nail through the foam into a flat piece of wood, then spin the foam piece around the nail with the hot wire held stationary at the correct radius. Hot glue is the best method of gluing foam.

Tulloch (1996b) gives the design and use of a flatfield lightbox using LEDs for simple single-wavelength calibration tests. In his case, he wanted an illumination source for a bare CCD. He obtained good results using four LEDs, one in each corner of a box, shining on a screen. This arrangement gives good illumination if the box side lengths are exactly equal to the height of the box. This neight is fairly sensitive; values of 0.8 and 1.2 times the box sides yields very uneven illumination. To improve the flatness even further, increase the basic dimension so that the CCD is seeing the center of the illumination pattern. When converting this into a standard lightbox, his table indicates that the proper ratio of height to box width for a reflective light box is about 0.64. That is, if your light box is 12 inches square, it should be about 8 inches high.

5. Scattered Light

Scattered light is present in almost every telescope. Normally it is at a low-level, but any scattered light impacts photometry since it represents light that has not properly passed through the optical system. The biggest impact of scattered light is to the flats. Shown in figure 13a is a flatfield from NOFS, using the 1.0 m telescope and a 1024x1024 CCD. Note the symmetrical donut in the center of the flat. Such a feature is almost always caused by scattered light, as most flats will be much more uniform. In this particular case, the scattered light is caused by the retaining ring holding the dewar entrance window in place. The ring is bare aluminum, and reflections from its sides are focussed towards the center of the CCD. When taking data frames with this system, you would see the normal star field plus sky background, with an enhancement of the sky background at the position of the donut. If you used the flat to process the raw frame, the result would be a flat-looking frame since the amount of scattered light in the sky is the same fraction as the scattered light in the flat. However, the stars that fell within that donut would be suppressed. This effect resulted in about a 0.06 mag error in the photometry for this particular system. We corrected this problem by placing a baffle over the window, of sufficient size so that incoming light never reaches the sides of the retaining ring. After solving this problem, a flat taken with the system is shown in figure 13b.

After this problem was solved, the flatfield that results is shown in figures 14a and 14b. The flat looks better, though there are definite large scale patterns. The question to ask is whether this is a good flat or whether there are additional scattering surfaces impacting the flat. In our case, with a thinned, backside illuminated CCD, we had one additional piece of information. The left panel shows a U-band flat; U-band performance is highly dependent on the amount of thinning of the CCD. The right panel shows an I-band sky frame, with the normal airglow fringing pattern. Note that the fringes follow the basic flatfield pattern, indicating that the reason for the large-scale variations is just thickness variations in the CCD itself.

A similar problem was seen by Balcells (1996) in the Isaac Newton Telescope (INT) prime-focus camera. They solved the light leak by placing plastic masks over each filter, plus wrapping the camera in black cloth. Flatfielding errors decreased from a few percent to much less than one percent. Even larger errors (10 percent) had been present for narrow-band filters, so much so that the images were usually considered unusable for photometry. With the changes, narrow-band errors decreased to less than 0.5 percent.

If the presence of scattered light is suspected in a camera system, a method is needed to track down its source. The simplest method is to pull the CCD camera off the telescope and just look through the optical system at sunlight or the illuminated flatfield screen. Look for any light leaking around baffles; if you can see it by eye, so can the CCD system. Look for reflections, especially from tube sides or clips in the optical path. If you cannot see anything visually, most of the major problems have been covered.

Gutierrez et al. (1996) describe a more sophisticated method for checking for scattered light. The basic idea is to look at the telescope exit pupil. If there is no scattered light, the pupil should look like the front of your telescope: a bright ring with a central obstruction from the secondary mirror. Any scattered light will show as additional components in the pupil image.

You form the pupil by creating a *camera obscura*. A high-quality pinhole is suspended above the CCD window, centered as best as possible. The pinhole will create an image of the telescope pupil on the CCD itself.

While such a technique will indicate that scattered light is present in the system, it doesn't indicate how that scattered light component is being produced. You need to understand your optical system and perhaps ray-trace the system to see where the problem resides. While solving a scattered light problem can be very time-consuming, you will be compensated by both improved flatfielding and lower background in the images, allowing you to work fainter and with more precision than before.

6. Optical Distortions

Optical distortions are present in many telescope systems. Amateur telescopes often have a focal reducer. For professional telescopes, there may be a corrector or field flattener. Understanding the impact of these optical systems on photometry is important.

Jacoby et al. (1998) describe a wide-field corrector for the KPNO Mayall 4 m telescope. A mention is made of the effect of the optical distortions on proper flatfielding. If the solid angle subtended by a pixel changes over the field, then a uniform illumination will not be properly reproduced on the CCD; pixels with smaller angular size will appear to be vignetted. This can have serious effect on photometry. Consider the case of undersampling, where all of the flux from a star falls within the smallest angular sized pixel. Flatfielding the image will boost the flux in the small pixels relative to the larger pixels and make a star that fell in the small angular sized pixel seem brighter than one that fell in a larger angular sized pixel. The same argument holds for more properly sampled stellar profiles. In general, if there is not a uniform scale over the CCD, there will be flatfielding errors. These errors are also be dependent on the type of distortion (purely radial, for instance, only affects pixels in one dimension). The proper way to remove these errors is to reproject your data into a uniform tangent-plane scale.

7. Nonuniform Illumination

Wild (1997) has an interesting way to determine flatfields. The light illumination from all of the techniques listed above is nonuniform at some level. Since the CCD response itself is invariant, and the light source can assumed to be invariant but uneven, one should be able to solve for the CCD response by itself.

If L(x,y) is the integrated signal on the CCD focal plane, and F(x,y) is the CCD response, then the total response from combining many images is

$$< H_1(x,y) >= L(x,y)F(x,y) \tag{8.2}$$

If the illumination pattern is moved on the CCD by a small amount δ by moving the telescope, the new pattern yields

$$< H_2(x-\delta,y) >= L(x,y)F(x-\delta,y) \tag{8.3}$$

Dividing these two equations gives

$$\frac{< H_2(x-\delta,y) >}{< H_1(x,y) >}F(x,y) = F(x-\delta,y) \tag{8.4}$$

where the L(x,y) cancels out. The task now becomes one of solving for F(x,y) given the measurements condensed in the H(x,y) responses.

In theory, this can be solved for the CCD response. This is especially true for a space-borne instruments where atmospheric perturbations do not occur. In practice, this is a very difficult problem since you have readnoise, psf variations, atmospheric scintillation and extinction variations, etc. However, it is a computational intensive technique and may be a method to obtain higher-quality flatfields in the future.

8. Testing Flats

In order for the flatfield correction process to work the flatfield image must accurately map deviations from "flatness" introduced by the optical system and the CCD camera. This requires that the flatfield illumination entering the telescope is itself flat. This turns out to be a difficult thing to test.

For projector flats, the CCD can be rotated 90 degrees between flatfield exposures. A ratio image, produced by dividing the first flatfield by the second rotated flatfield, will reveal any non-axisymmetric gradients. For lightboxes it is often easier to rotate the lightbox instead of the CCD. Another trick with lightboxes was mentioned by Berry(2001). The telescope is pointed horizontally, and a nearfield flatfield image is taken. The lightbox is then removed and moved far away from the telescope, so that an image of it will illuminate only a small portion of the CCD (10 percent or less). The telescope is then moved so this small illuminated spot is raster scanned across the CCD. The spot has constant illumination. If you divide the raster-scanned images by the nearfield flatfield, you can compare the spot brightness at different places on the CCD and see if any gradients remain; these gradients would then be present in the nearfield flatfield/full illumination image.

For twilight flats the light light source can not be turned and there is not enough time to rotate the CCD. One method is to compare flats taken in evening twilight to those taken in morning twilight. The image formed by the division of one set by the other will reveal any gradients in the sky illumination because any such gradient is expected to be 180 degrees different between the two sets.

For master sky flats, you also cannot rotate the CCD. A technique for testing flatness was suggested by Manfroid (1995,1996). A star field with a large number of standard stars is used, for example, a cluster like M67. Differential photometry is used to generate an error matrix across the face of the CCD. It is also possible to use smaller star fields by raster-scanning them across the CCD. The object is to find if any systematic differences in the photometry with respect to the center of the field. This technique works if the filter system is close to the standard system so transformation errors do not dominate any systematic pattern, and if the sky is photometric (especially in the case of rastering). Use plenty of stars and expose them so that Poisson error is much less than one percent. For wide field cameras, using Tycho2 stars can often highlight systematic errors.

REFERENCES

Wu, et. al. 1972. AAS Photobulletin 1, 9.

Holst, G. C. 1998. CCD Arrays Cameras and Displays (Bellingham: SPIE).

Jacoby, G. H. et al. 1998, SPIE 3355, 721.

Massey, P. and G. H. Jacoby 1992, in Astronomical CCD Observing and Reduction Techniques, ASP Conf. Ser. 23, S. B. Howell, ed., p256.

Gutierrez, L. et al., 1996, Rev. Mex. de Astron. y Astrof. 32, 63.

Chromey, F. R. and D. A. Hasselbacher 1996, PASP 108,944.

Tyson, N. D. and R. R. Gal 1993, AJ 105, 1206.

Wild, W. J. 1997, PASP 109, 1269.

Tulloch, S. 1996, RGO Technical Note 108 (http://vela.ing.iac.es/Astronomy/observing/manuals/ps/tech_notes).

Tulloch, S. 1997, RGO Technical Note 111.

Tulloch, S. 1996, RGO Technical Note 107.

Berry, R. and J. Burnell, 2001, The Handbook of Astronomical Image Processing (Richmond: Willmann-Bell).

Manfroid, J. 1995, Astr. Astrop. Sup Ser 113, 587.

Manfroid, J. 1995, Astr. Astrop. Sup Ser 118, 391.

Balcells, M. RGO Technical Note 102.

Park, Hye-Sook 2001, priv. comm.

Tyson, J. A., 1990, in CCDs in Astronomy, PASP Conf Ser 8, ed. G. H. Jacoby,1.

This plano tables was prepared with the AAS LATEX macros v4.0.

TABLE 7.1

Color Temperature of Various Sources

Source	Temperature (K)
Average Daylight	6500
Xenon(arc or flash)	6000
Studio tungsten lamps	3200
Floodlights	3000
Domestic tungsten lamps	2800-2900
Sunlight at sunset	2000
Candle flame	1800

TABLE 7.2

Blank Fields

Field	RA(1950)	DEC(1950)	Comments
SGP	00:48:30.3	-28:03:01	10 objects in 6.4'x6.4'
0102-265	01:02:48.5	-26:34:25	15 objects in 6.4'x6.4'
0248-196	02:47:46.6	-19:46:11	
KPNO-1	04:25:46	+54:09:00	
0427-36	04:27:26.1	-36:25:16	ok 5'x5'
9090-7	09:09:32.6	-07:38:26	
CT1	10:04:27.0	-02:19:00	
	10:55:49.7	-77:08:36	
1101-264	11:02:23.5	-26:43:57	
CT5	12:26:08.7	-06:38:29	
KPNO-2	13:04:33	+29:50:50	
CT2	12:54:58.7	-02:07:04	
SER1	15:13:13.9	-00:31:48	
KPNO-3	19:19:09	+12:22:05	
2335-40	23:34:34.6	-39:49:04	
2345+007	23:45:45.9	+00:40:40	
SPKS 02	11:07:06.0	-77:05:54	
SPKS 03	12:28:38.0	-62:28:18	
SPKS 04	15:39:43.0	-33:56:54	no stars 5'x5'
SPKS 05	16:02:43.0	-45:48:18	few stars 5'x5'
SPKS 06	16:25:40.0	-25:09:48	Few stars in 5' x 5'
SPKS 07	17:07:47.0	-40:34:06	
SPKS 08	17:31:06.0	-25:44:24	
SPKS 09	18:02:01.0	-04:32:36	
SPKS 10	19:00:48.0	-37:21:18	
SPKS 11	13:04:33.0	+29:50:49	
SPKS 12	16:49:42.0	-15:21:00	No stars
SPKS 13	19:19:09.0	+12:22:05	
SPKS 14	21:26:54.4	-08:51:41	

Contents