

IRIS Technical Note 16: Flat Fields

October 24, 2013

Prepared by:

Sarah Jaeggli, Montana State University

P.O. Box 173840 Bozeman, MT 59717-3840

jaeggli@solar.physics.montana.edu, (406)994-6072

with Bruce Lites, Mats Carlsson, Rock Bush, Bart de Pontieu, Charles Kankelborg, and J.P. Wuelser

Contents

1	Summary	2
2	Pre-Flight Flat Field Characterization	2
2.1	Lamp Flat Fields	2
2.1.1	Strategy	3
2.1.2	Results	4
2.2	Chae Flat Field Method	4
2.2.1	Description	4
2.2.2	Results	15
2.3	Comparison	15
3	In-Flight Flat Field Characterization	18
3.1	Slit-Jaw Imager Flats	18
3.1.1	Implementation	18
3.1.2	Results	20
3.2	Spectrograph Flats	20
3.2.1	Method	20
3.2.2	Implementation	25

3.2.3	NUV Results	25
3.2.4	FUV Results	26
3.3	Sensitivity Monitoring using Flat Fields	29

1 Summary

Raw images from IRIS may contain a variety of undesirable features that need to be removed to carry out good data analysis. These include the pattern of gain variation on the CCD, dust on various optical surfaces near focal planes in the spectrograph and imagers, and possible vignetting. Spectrograph images also contain fixed intensity variation pattern along the spatial dimension due to slight imperfections in the slit. In this document we describe the techniques used to construct flat fields from pre- and in-flight data.

Regular flat field calibration will be essential to characterizing any instrument changes, especially a loss of sensitivity for the FUV spectrograph.

Goal: The flat field should correct the images to a level equivalent to the shot and read noise of the brightest regions of typical images. Based on FO2-1 of the concept study report, typical bright count rates in active regions are a little bit less than 10,000 counts/s/spatial pixel/spectral pixel. This implies a signal to noise of 100, therefore any flat field used should be accurate to about 1%.

2 Pre-Flight Flat Field Characterization

2.1 Lamp Flat Fields

Prior to their installation in IRIS, standard UV flat field images for all of the CCDs were taken using a mercury UV lamp source filtered to 2537 Å. The flat field image tests are summarized in Table 1. A single dark and two flat images were taken for every detector. Each detector half (identified as side *h* and *g*) was read out separately to produce an individual fits file. The images for each detector half also contain an overscan region which contains no intensity signal and is useful for the calculation of the detector bias and read noise. For the flat field strategy laid out in the next section, the dark frame was subtracted from each flat frame and the two flat images were averaged together. The overscan was trimmed away and the two halves of the detector were combined to produce a single image.

Table 1: Summary of Flat Images

Array ID	serial	exposure [sec]	Date taken
FUV 1	20_02	20	2011-11-29
FUV 2	13_09	20	2011-12-01
NUV	18_04	20	2011-11-10
SJI	18_06	10	2011-10-27
Spare	20_06	20	2012-02-25

2.1.1 Strategy

The flat field images taken with the lamp include four main effects:

1. Fine-scale pattern on the CCD surface from annealing
2. Photon-counting error
3. Read noise
4. Detector gain
5. Illumination pattern of the lamp

The lamp illumination pattern must be removed from the flat images without also removing the fine-scale pattern from the CCD or the relative gain level between detectors, as any data will also suffer from a difference in gain. If we assume:

1. The anneal pattern is generally flat across the CCD.
2. The detector response (gain) is uniform across the array half and can be characterized by a single value.
3. The illumination pattern is continuous between the two halves of the CCD and can be described by a smoothly-varying polynomial function.

To this end I have written a routine to fit a 2D quadratic surface to the reduced flat image with independent intercept terms for each array half,

which has the functional form:

$$\begin{aligned}
 f(x, y) = & a \\
 & + b_1x + b_2x^2 \\
 & + c_1y + c_2y^2 \\
 & + d_1xy + d_2x^2y + d_3xy^2 + d_4x^2y^2
 \end{aligned} \tag{1}$$

where a has the value a_1 for the left half of the array and the value a_2 for the right half of the array.

2.1.2 Results

For each CCD, the illumination pattern was calculated from the fitted parameters assuming a constant value of $a = (a_1 + a_2)/2$ for the whole CCD. The resulting illumination pattern was divided from the original image to produce the master flat field. The original flat image, fitted quadratic surface, and resulting master flat field are shown in Figures 1, 3, 5, 7, and 9 for each of the CCDs. Figures 2, 4, 6, 8, and 10 show plots of the average value of the pixels in the “spatial” and “spectral” directions for each of the three images. (Note: Obviously the SJI CCD has no spectral dimension, I use “spatial” and “spectral” to refer to the directions parallel and perpendicular to the slit which are the short and long dimensions respectively for all the CCDs.)

The difference in gain between the two halves of the detector is most extreme for the FUV 1 CCD. This jump in gain is less noticeable, but still present, for the other CCDs, and can be seen more easily in the plots of the averaged counts in the horizontal direction.

2.2 Chae Flat Field Method

2.2.1 Description

The flat field method put forth by J. Chae (2004, Sol. Phys., 221, 1) can be used to build a flat field from a set of non-uniform images where the illumination pattern has been shifted relative to the detector in each image. The technique extracts the illumination pattern (or object) from the flat field pattern using a least squares method, keeping the object shifts and illumination level as free parameters. This technique has been proven more robust than other similar methods (e.g. Kuhn, Lin, & Loran 1991). We have inherited a Chae method code from Ted Tarbell used for producing flat fields for Hinode and AIA.




Figure 1: (left) Original flat image, (center) fitted illumination pattern, and (right) master flat field for the FUV 1 CCD.




Figure 2: Average behavior of the FUV 1 CCD original flat image (black), fitted surface (red), and corrected flat (blue) in the spatial (top), and spectral (bottom) directions.




Figure 3: (left) Original flat image, (center) fitted illumination pattern, and (right) master flat field for the FUV 2 CCD.




Figure 4: Average behavior of the FUV 2 CCD original flat image (black), fitted surface (red), and corrected flat (blue) in the spatial (top), and spectral (bottom) directions.




Figure 5: (left) Original flat image, (center) fitted illumination pattern, and (right) master flat field for the NUV CCD.




Figure 6: Average behavior of the NUV CCD original flat image (black), fitted surface (red), and corrected flat (blue) in the spatial (top), and spectral (bottom) directions.

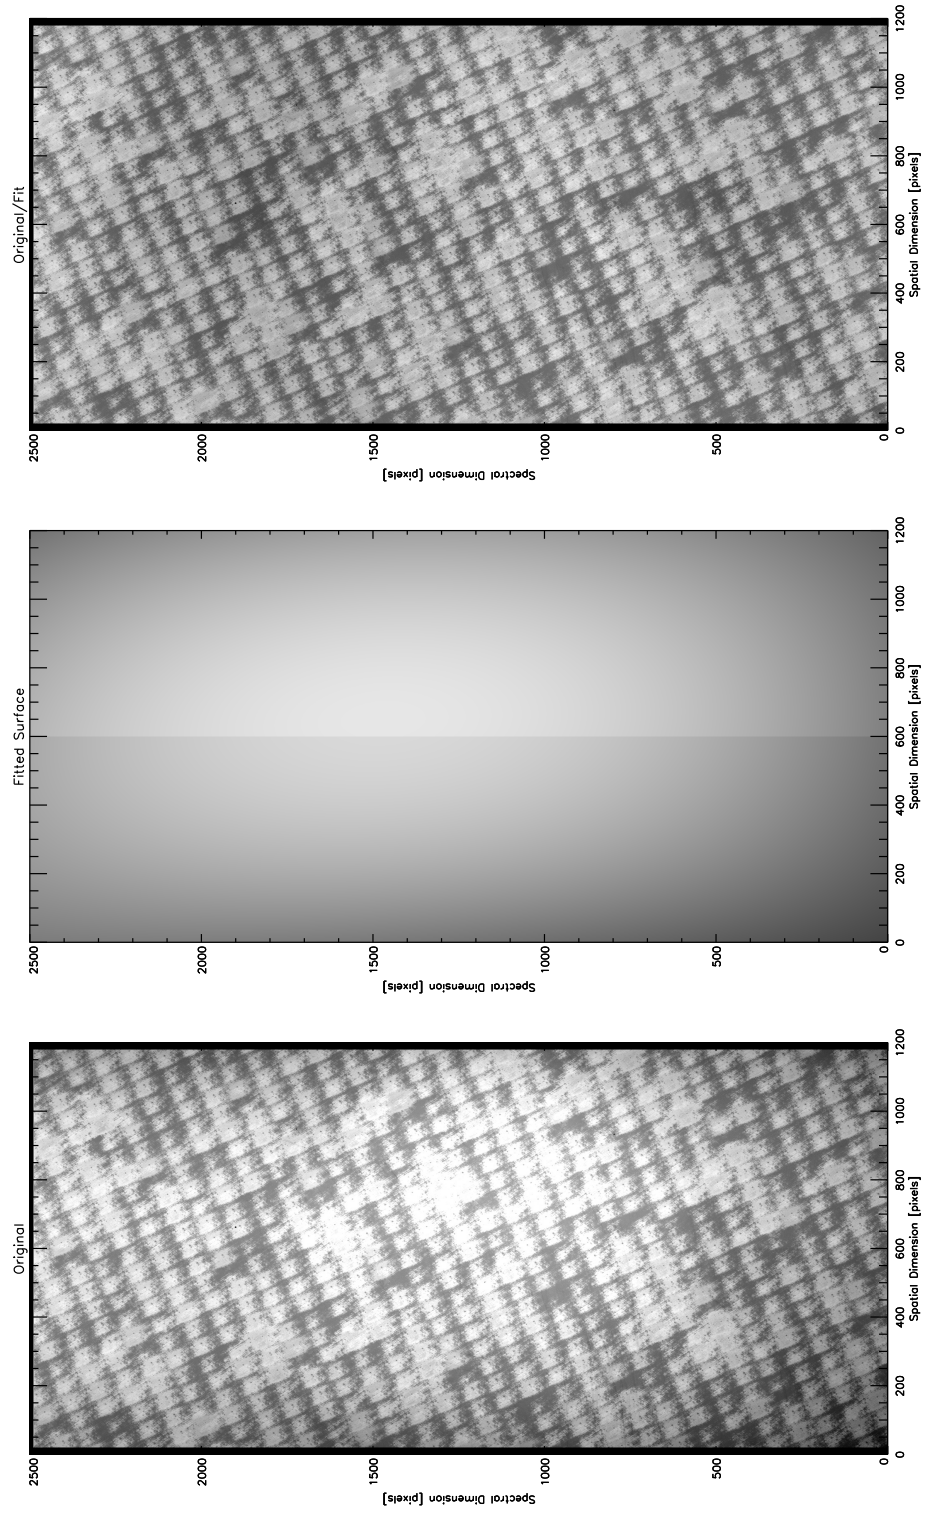


Figure 7: (left) Original flat image, (center) fitted illumination pattern, and (right) master flat field for the SJI CCD.




Figure 8: Average behavior of the SJI CCD original flat image (black), fitted surface (red), and corrected flat (blue) in the spatial (top), and spectral (bottom) directions.




Figure 9: (left) Original flat image, (center) fitted illumination pattern, and (right) master flat field for the Spare CCD.




Figure 10: Average behavior of the Spare CCD original flat image (black), fitted surface (red), and corrected flat (blue) in the spatial (top), and spectral (bottom) directions.

To test this method of flat fielding, and verify the validity of the Hg lamp flat fields at the FUV native wavelength, three sets of images were taken with the FUV SJI using the deuterium lamp source during the Thermal Vacuum tests conducted in June-July 2012. The illumination pattern of the deuterium lamp was much smaller than the detector, so it was positioned in a semi-random 7x7 grid to ensure even coverage of the detector.

2.2.2 Results

The position of the smoothly varying lamp illumination pattern was difficult for the Chae method shift determination to guess. The method was intended to be used on detector-filling high contrast scenes with relatively small shifts in the image. The shifts relative to the center image needed to be determined by hand and input into the code.

Figure 11 shows the resulting flat field and object image. The lower two frames of the figure show the flat field applied to an image in the set. Many features are apparent in the flat field: the pattern of the detector, the marks of dust particles on various optics, and of course the slit. In addition, the upper and lower right hand corners of the image appear vignetted. The flat fielded image in the lower right panel appears to be corrected very well, although some remnant of the slit is still present.

2.3 Comparison

Testing the detector flat fields produced with the Hg lamp against other flats (or any uniform image) at the native observing wavelength is the ultimate test of success for the Hg lamp flats. With the FUV SJI flat produced with the Chae technique above it is possible to see whether the 2537 Å flat applies to observations at the FUV wavelength. Figure 12 shows the Chae-method flat divided by the Hg lamp flat for the SJI. While some of the structure due to the CCD gain pattern is removed, there is still a residual pattern which is apparent. In addition the magnitude of the gain jump between the two detector halves is not reproduced in the Chae flat. In the center panel of Figure 13 we show a small section of the Chae-method flat field that does not contain significant features due to dust on the optics. Comparing this directly with the Hg lamp flat field on the left we can see that there are subtle differences in the detector flat field patterns at each wavelength. The variation in the flat field pattern for both methods is about 11%, the division of the two results in an image with about 4% variation due to the residual flat field pattern.




Figure 11: The upper left frame show a flat field for the FUV SJI constructed from a scan of the deuterium lamp across the detector using the Chae method. The upper right shows the resulting object image, and the bottom two frames show an example raw image and the flat fielded equivalent.




Figure 12: The residual of the Chae-method deuterium lamp flat field divided by the Hg lamp flat field.




Figure 13: Sub-field of the FUV SJI detector flat field for the Hg lamp flat (left), Chae-method deuterium lamp flat (center), and the division of the two.

3 In-Flight Flat Field Characterization

The Hg lamp flats only account for the detector, they do not include dust and features due to the optics in IRIS, which may be a significant source of contamination in the flat field. The wavelength-dependent effect of charge spreading has been noted during the calibration of detector gain (see ITN 25 on CCD stats), this may tend to smooth or otherwise alter the flat field pattern at shorter wavelengths. Longer wavelength photons penetrate more deeply into the CCD, changing the affect of the annealing pattern, possibly reducing the contrast. As a result the correction provided by the Hg lamp flat at FUV wavelengths may not be adequate for high signal-to-noise measurements as there will still be residual structure present. It is therefore necessary to build solar flat fields for each of the detectors and every imaging path while in orbit.

3.1 Slit-Jaw Imager Flats

Write a brief intro here.

3.1.1 Implementation

In Chae (2004), the author suggests a dither method (a Reuleaux triangle), the ideal size of the pattern, and the best number of points in the pattern. A Reuleaux “triangle” is a commonly used dither pattern in optical and radio astronomy which is formed from the intersection of three circles of equal




Figure 14: An example of a Reuleaux triangle with 15 points and a pattern size of 20 arcsec.

radius, an example of the dither pattern used for IRIS is shown in Figure 14. Based on his analysis, Chae recommends a pattern size of $L/N_x = 0.1$ to 0.25.

The area for each slit-jaw image is approximately 1000×1000 pixels, therefore the best size for the dither pattern spans 100-250 pixels. With a spatial sampling is 1/6 arcsec/pixel, the equivalent size on the sky is 17-42 arcsec. We've selected a pattern size of 20 arcsec. Having more points in the dither pattern increases the accuracy of the resulting flat field, but a longer sequence also increases the chances that the scene might change, especially in the dynamic chromosphere that IRIS observes, therefore we keep the sequences short with only 15 pointing, however each observation is repeated three times so that flat fields from different sequences can be averaged together to produce a more accurate final flat field.

Initial tests of the Chae method on FUV data did not produce good flat fields. The contrast in these channels is very high and some regions

of the CCD were receiving very little intensity, so that the CCD annealing pattern (5% rms) did not rise above the photon counting and read noise limit. To get a more even intensity pattern the telescope is taken out of focus (to a focus position=0) for the 1330, 1400, 1600, and 2793 Å filters. The sequences for the 2832 and 5000 Å filters remain in focus. Results are shown in the next section.

Details of observation: The filters and filter wheel position change the flat field slightly. Darks are taken before and after the flat field sequence for each filter. Targets for the SJI flats are quiet network plage or quiet sun, the potential for evolution in brighter targets is too high.

3.1.2 Results

Important to note that images shown here are in their level 0, unflipped, orientation. The NUV object image still shows a strong residual in the upper right corner due to vignetting of the Solc filter mask. The Chae technique assigns about half of the darkening to the object frame and the flat field only corrects for part the drop in intensity due to the mask. A larger dither pattern may help the technique to distinguish this pattern. The flat field correction away from the upper right corner is quite good, and the other corners do not seem to show vignetting at all.

The FUV flat field from the Chae method still contains some residual structure due to slight changes in the object with time. Multiple sequences processed with the Chae method, then averaged help to even out this structure in the final flat field.

3.2 Spectrograph Flats

3.2.1 Method

Flat fields for the spectrographs are not as simple to obtain and process. Quiet-Sun spectra are taken with the telescope is defocused and slewing semi-randomly. These are used to build up a spatially smooth average spectrum. The spectrum averaged over all the spatial pixels on the array is then divided to retrieve the residual flat field pattern which contains the detector pattern, the fixed pattern of the slit, and any dust or features internal to the spectrograph optics. Thermal shifts decouple flat field features on the detector from the intensity pattern due to the slit, making it necessary to determine and apply these patterns separately based on the alignment of the spectrum and array.




Figure 15: An example of the resulting flat field and object frame from the Chae method applied to the 2832 Å channel of the SJI. The two additional frames show an original image from the observed sequence and its correction by the flat field.




Figure 16: An example of the resulting flat field and object frame from the Chae method applied to the 1330 Å channel of the SJI. The two additional frames show an original image from the observed sequence and its correction by the flat field.




Figure 17: The composite Chae method flat fields for all of the SJI filters, the flats are in their Level 1 orientation. Images are scaled between the values of 0.4 and 1.3.

IRIS observes high-contrast structures, the higher the contrast, the more images are needed in a flat field series. In the FUV even the quiet-Sun is full of discrete bright features surrounded by dark regions. The ratio of the standard deviation to the mean signal tells us how uniform an image is. The ratio of the standard deviation compared to the mean signal of structures in a reduced SUMER image of the quiet-Sun is roughly 0.5. We can think about this in the way of standard image co-addition. If we want to achieve a flat field with 1% variation (this is somewhat arbitrary) we need $(0.5/0.01)^2 = 2500$ images. This brute force method (with the telescope in focus) is not very attractive. Two things help increase the uniformity in individual images are 1) changing pointing during individual exposures and 2) taking the telescope out of focus.

How much #1 helps depends on how much area is sampled during an exposure. #2 can be quantified based on pre-flight focus testing for the SJI. Figure ?? show the best (focus=-60) and worst (focus=-140) focused images from a series taken for the NUV SJI using a star-pattern and an NUV LED for illumination. From these images I've taken the mean signal and standard deviation in a sub section of the annulus encompassed by the red radii. The ratio of the standard deviations at the good and bad focus positions is given in Figure ?. Taking the telescope out of focus decreases the contrast (as we have defined it above) by a factor of about 3. This significantly decreases the number of images needed to build a 1% flat field to $(0.5/3*0.01)^2 = 278$ images.

In the case of the NUV, the signal levels throughout the spectrum are relatively high because there is still some continuum contribution and it should be possible to build a high-quality flat field rapidly. The image contrast (ratio of standard deviation to mean) is 0.2 in Bifrost simulated NUV continuum images, but still high in the Mg II line core, about 0.45. Fewer image sums are necessary to build a flat field for the NUV, but taking the telescope out of focus is still probably desirable.

The FUV contains mostly emission lines with a low-lying continuum. It is likely the continuum will be dominated by scattered light, significantly increasing the count rates in the regions between the lines. While flat field exposures for the FUV spectrograph may still take significantly longer than for the NUV, they will be a high priority. During the flat fields care should be taken to avoid dynamic activity which may be present even in the quiet Sun chromosphere to avoid saturation in the lines and large wavelength shifts due to high velocities.

3.2.2 Implementation

Images from a flat field sequence are flipped, dark subtracted, and averaged together to produce the intermediate flat field. This flat field is then up-sampled using a lossless Fourier technique (fftrebin) to twice its original size and interpolated using the geometric correction determined for the IRIS spectra which makes the spatial and spectral dimensions rectilinear to the array grid. The FFT up-sampling helps mitigate the errors inherent in resampling the image.

In spite of the rigorous averaging done throughout the observations, there is still some spatial variation in the spectrum, especially in the line cores. A running average with a width of $N_y/8$ is taken along the spatial dimension in an attempt to produce a spatially smooth spectrum, what we will call the spectral flat field. This spectrum still contains a residual of the annealing pattern of the CCD. The annealing pattern is also present in the lamp flat field produced in Section 2.1, so in an attempt to reduce the annealing pattern residual in the spectral flat we have sent it through the same process of resampling and smoothing. This pattern is then divided from the spectral flat field. The contrast of the spectral flat is slightly different, but the smoothed pattern appears to be roughly the same.

An average along the spectral dimension is taken in the image where the lines have been removed to produce the spatial flat field. Similar to the spectral flat field, the annealing pattern residual is divided from the result using the pattern from the lamp flat field. The fiducial marks are removed from the spatial flat simply by marking them as missing data and interpolating over them, so the fiducial marks remain in the data processed by the flat field.

The spectral and spatial flat field are transformed back to the original (warped) coordinate frame of the Level 1 spectra and resampled back down to size. The spectral flat field is divided from the intermediate flat field. The spatial flat field is retained separately so it can be shifted and applied to mitigate differences in the spatial pattern from thermal shifts in the spectrograph.

3.2.3 NUV Results

Figure 18 shows the results of this technique applied to the NUV data. The top left panel shows the intermediate flat field which is an average of a sequence of 150 images. The top right shows the final spectral flat field, and the bottom left shows the final spatial flat field with the fiducials removed.

The bottom right panel shows the result of removing the all the spatial and spectral features from the intermediate flat field. The master flat field is similar to the NUV lamp flat field, but it has a slightly different contrast and contains some subtle features which are not present in the lamp flat. Comparison with flat field sequences taken at other times should help, us determine the accuracy of the various aspects of the flat field.

3.2.4 FUV Results

The technique outlined above was not as successful when applied to the FUV data. The FUV has a low count level away from the lines, so we might expect a flat field generated with this technique to be almost exclusively noise away from the line center. In particular, it is difficult to characterize the spatial pattern that works well both in and away from lines. This is unusual, as the CCDs are expected to have a very linear response at low count rates, and the dark correction works with good accuracy on data binned 1×1 .

It appears that there is an additional background contribution to the FUV spectra. This contribution is most apparent on the FUV long wavelength detector, as shown in Figure 19, where it has an irregularly shaped edge with respect to the spectrum, indicated by the yellow boundary. This background does not seem to have the annealing pattern of the CCD imprinted on it, indicating that it is cannot from FUV stray light, but probably due to long wavelength visible or infrared photons, which penetrate more deeply into the surface of the CCD which has been optimized for UV light.

The FUV background is discussed in depth in a forthcoming technical note. The results of this testing indicate:

1. The background contributes approximately 5 DN/pixel/sec at disk center.
2. The background level has a smoothly varying spatial distribution, much like a smoothed limb darkening function which is offset from disk center
3. It is fixed with respect to the telescope/spectrograph orientation (we haven't actually tested this but it should be true).
4. The background exhibits photon counting statistics consistent with visible or infrared light (i.e. one electron per photon).
5. The background illumination pattern weakly dependent on pointing.





Figure 18: The NUV spectrograph flat fields.



Based on these findings, the background should not be difficult to characterize and remove from the image based on pointing, roll angle, and exposure time. The background must be removed before the flat field for the FUV channels can be successfully characterized and removed. This is especially true for long exposures and binned data where the background is more evident.

3.3 Sensitivity Monitoring using Flat Fields

Changes in sensitivity of the CCD, or charge burn-in, is mainly a concern for the bright C II and Si IV lines in the FUV but can also be measured for the Mg II lines in the NUV. For the spectrographs burn-in might occur isotropically along the spatial dimension of a line, causing an apparent deepening in the emission line core with respect to the dimmer wings, but more likely the damage will show some spatial variation, with the most burn-in occurring at the middle of the slit, where bright interesting targets tend to be centered. For the SJI, even anisotropic burn-in is not an issue, and can easily be measured and removed using the flat fielding technique.

Due to the running smooth which is taken during the spectrograph flat field processing, the final flat field may not properly account for burn-in. The intermediate flat field from a 150 image sequence produces a C II line profile which is spatially smooth to an RMS of 8% of the average intensity in the line. Smooth changes in the CCD sensitivity of less than this level cannot be distinguished from the residual spatial variation. A larger number of images in an individual sequence, or an average of multiple sequences would be necessary to distinguish burn-in at a lower level. Isotropic changes in the line profile can be easily measured from by-products of the flat field processing. It will be important to keep track of this line profile.