

# Automated Spectral Reduction Pipelines

Robert J. Smith<sup>a</sup>, Andrzej S. Piascik<sup>a</sup>, Iain A. Steele<sup>a</sup>, and Robert M. Barnsley<sup>a,b</sup>

<sup>a</sup>Astrophysics Research Institute (ARI), Liverpool John Moores University, L3 5RF, UK

<sup>b</sup>University of Oxford, Department of Physics, Denys Wilkinson Building, Keble Road, Oxford  
OX1 3RH, UK

## ABSTRACT

The Liverpool Telescope automated spectral data reduction pipelines perform both removal of instrumental signatures and provide wavelength calibrated data products promptly after observation. Unique science drivers for each of three instruments led to novel hardware solutions which required reassessment of some of the conventional CCD reduction recipes. For example, we describe the derivation of bias and dark corrections on detectors with neither overscan or shutter. In the context of spectroscopy we compare the quality of flat fielding resulting from different algorithmic combinations of dispersed and non-dispersed sky and lamp flats in the case of spectra suffering from 2D spatial distortions.

**Keywords:** data reduction, processing pipelines, CCD, bias, dark current, spectral flat field

## 1. INTRODUCTION

The Liverpool Telescope<sup>1</sup> (LT) is a 2.0-m f/10, fully automated, robotic facility owned and operated by Liverpool John Moores University at Roque de los Muchachos Observatory on the Canary Island of La Palma. There are no staff or visiting astronomers and no observations are performed interactively. Astronomers make the science-based decisions of picking their targets, signal-to-noise ratio required etc and the procedural tasks of scheduling, observation, data handling and processing are automated.<sup>2</sup> Facilities are offered to create unique observation sequences where specific science requirements demand it, but most observers rely on the software “wizards” which automatically generate sequences for them, removing most of the need to understand the opto-mechanical constraints of each individual instrument. This leads to largely homogeneous data sets which is advantageous for the time-domain research in which the observatory specialises.

The highly automated operations paradigm of this observatory lends itself well to fully automated data reduction pipelines. All imaging data are processed as they are obtained during the night. This provides immediate “quicklook” data products to those observers interested in transient sources. Just as importantly these images provide data quality metrics, such as atmospheric seeing, which are dynamically fed back into the autonomous scheduling,<sup>3</sup> allowing the Robotic Control System to adapt to changing conditions. Spectral data, the main topic of this paper, are processed asynchronously in parallel with subsequent observations and typically available in less than half an hour. Notwithstanding the unique capabilities and requirements of the Liverpool Telescope, autonomous pipelines are potentially of significant value to all observatories. Aside from the obvious benefit of freeing scientists from routine tasks to concentrate on science exploitation, pipelines provided by the observatory allow development of bespoke algorithms optimized to the unique features of each instrument. In this way, reduction recipes are defined by the people most intimately familiar with the hardware.

Frequently pipelines can be developed using widely available, generalized libraries and processing packages. Some specific requirements might require development of bespoke code. For example the LT real time pipeline prioritizes speed over almost all other requirements and though it typically performs only the very familiar CCD bias, dark and flat-field operations, has been developed in carefully speed-optimized C. Tests showed it was easy for a bespoke pipeline to decrease execution time by factors of several times compared to using more generalized packages such as IRAF and *figaro*.

---

Further author information: (Send correspondence to R.J.S.)

R.J.S.: E-mail: R.J.Smith@ljmu.ac.uk

LT Observatory: <http://telescope.livjm.ac.uk/>

In this paper we report on the investigation of various data processing algorithms for the purpose of removing instrumental signatures (i.e., bias, dark, flat-field) from the Liverpool Telescope spectrographs. These are not the time critical ‘real-time’ pipelines, but those that run the following morning to produce the final archival quality results. Section 2 describes the instruments themselves and their unique features which necessitated the development of these various algorithms. In Sec. 3 we look at alternative methods of determining the bias level of a data frame that do not rely on either an overscan strip or a zero-exposure image. Section 4 describes experiments with various flat-fielding strategies for long slit spectra. We discuss how different types of flat field are best suited to different tasks. In Sec. 4.3 we illustrate the strengths of lamp flats and in Sec. 4.4 we make use of the advantage that twilight sky flats have passed through the identical optical path as science data. In Sec. 4.5 we look at ways of combined these two different flats whilst simultaneously taking account of geometric distortions in the data.

## 2. THE INSTRUMENTS

The Liverpool Telescope is broadly instrumented to allow a diverse range of science programmes to run concurrently. All instruments are permanently mounted and available at the Cassegrain focus. A flat pick-off mirror directs the beam to the requested instrument making instrument changes possible in a few seconds which is critical to both rapid response time-domain research and to the dynamic scheduling in response to changing observing conditions. There are three primary, facility instruments. IO:O<sup>4</sup> and IO:I<sup>5,6</sup> are a pair of visible and near-infrared imagers and FRODOSpec<sup>7,8</sup> is the intermediate resolution, integral field, fibre-fed, dual arm spectrograph. There is then a suite of smaller instruments, all of which were originally conceived and designed for the specific science goals of particular research projects, but are integrated into the automated workflow of the robotic observatory and available to all telescope users. It has been our experience that by each of these instruments providing unique capabilities, many research groups find novel uses we may not have anticipated. In fact around 40% of applications to the observatory are for these specialised instruments and frequently for projects that go far beyond their original design requirements. By carefully concentrating the design process on their distinctive strengths rather than trying to provide a huge range of configuration options, these instruments provide high science impact for very low build costs. As of 2016, the installed suite consists of

- RINGO3,<sup>9</sup> a simultaneous, multi-band polarimeter, originally designed for temporal resolution ( $\sim 1$  Hz) of gamma ray burst polarisation, but now used for a wide range of time domain polarimetry and even as a high-speed, simultaneous multi-colour, non-polarimetric photometer,
- RISE,<sup>10</sup> a moderate speed ( $\sim 1$  Hz) frame-transfer broad-band imager designed for exoplanet occultation timing which also functions as a back-up for our primary CCD imager, IO:O,
- SPRAT,<sup>11</sup> a very high throughput visible light spectrograph, designed for rapid typing of newly discovered supernovae but now rivals the original facility spectrograph FRODOSpec in terms of number of observation requests,
- LOTUS,<sup>12</sup> a UV-optimized, low resolution spectrograph, designed and commissioned to provide ground-based observations in support of the *Rosetta* spacecraft.

With this suite, we offer imaging across almost the entire optical atmospheric window 3000–18000 Å, low and intermediate resolution spectroscopy from 3200–9400 Å and 1 Hz time-resolved polarimetry from 3500–10000 Å. The particular instruments of interest in the context of this paper are the two newest low-resolution spectrographs SPRAT and LOTUS.

### 2.1 SPRAT

SPRAT<sup>11</sup> is a low cost, low resolution ( $R\sim 350$ , 4000–8000 Å), high throughput ( $\sim 40\%$  total system efficiency at 5000–7000 Å including telescope, instrument and detector), long slit spectrograph designed specifically for the task of rapid, early-time spectral classification of transient objects detected by survey projects such as Pan-STARRS<sup>13</sup> PTF<sup>14</sup> and Gaia<sup>15</sup> which are highly effective at detecting transient sources, but generally then require independent spectroscopic identification of the source. It is an imaging spectrograph with deployable

slit and VPH grism allowing fully automated operation in which the imaging mode performs an iterative target acquisition before deploying the slit to obtain the spectrum. Collimator and dispersion optics are designed using all off-the-shelf optical components. The camera optics use a Nikkor SLR photographic lens and the detector is an Andor iDus 420 BV camera.

## 2.2 LOTUS

The science requirements for the LOTUS<sup>12</sup> ('LOW-cost Ultraviolet Spectrograph') design were driven by the 2015 *Rosetta* spacecraft<sup>16,17</sup> encounter with comet 67P/Churyumov-Gerasimenko. *Rosetta* has provided extraordinary images and direct in situ measurements of cometary nuclear activity, but from its unique vantage point it cannot observe the wider coma emission to associate that nuclear activity with the large scale behaviour of comets as observed from Earth. Strong near-ultra-violet emission lines (e.g., the NH(0,0) band at 3340 Å and the CN band at ~3880 Å) may be used to determine the gas content of the coma. Comparison of the emission line strengths with broad-band imaging allows mapping of the gas-to-dust ratio structure of the coma.<sup>18</sup>

To address this, the main science requirements for the instrument design were a spectral resolution  $R \sim 300$  over at least the range 3200–6000 Å. Engineering requirements limited the instrument to a small size and weight (<25 kg) and a hardware budget of <£10,000. We selected all off-the-shelf components for both low cost and rapid deployment. The detector is a Kodak Truesense KAI11002M interline CCD in a Starlight-Xpress Trius-SX35 camera, designed for the amateur astrophotography market. This device combines low cost, a physical detector area well matched to the instrument focal plane and good UV sensitivity. Its quantum efficiency of 25% at 3400 Å exceeds that of many thinned, back-illuminated CCDs more typically used in astronomical instruments.

LOTUS is not dedicated to the 67P/Churyumov-Gerasimenko observations. It is integrated into the common instrument infrastructure of the observatory and available to all users using the same flexible scheduling tools as all other instruments.

## 3. BIAS AND DARK ESTIMATION WITHOUT OVERSCAN STRIPS

As described in Sec. 2, the instruments discussed here are all extremely low cost, fast-track designs that were targeted to address specific science requirements rather than being facility class instruments. Selection of the detectors was therefore determined by unique criteria for each instrument. In all cases the detectors were required to be integrated, commercial, off-the-shelf (COTS) cameras with available control software. This minimizes development cost and allows adoption of a rapid prototyping design process that can focus available resources on the novel features of each instrument.

Use of these COTS cameras does however sometimes lead to compromises where the device was never designed for a research class instrument or indeed not designed for astronomical use at all. Common examples are non-photometric shutters which create uneven field illumination or lack of a bias overscan in the read out. Frequently these limitations are outweighed by the hardware's advantages and can be overcome in software. The cameras in SPRAT and LOTUS are very different devices, but both suffer the same disadvantage of having no overscan region to monitor bias voltage changes during read-out.

Bias images typically comprise of a zero time integration which contains two sources of signal. First, a constant pedestal signal set by the amplifier bias voltage and second there may be structure in that constant signal caused by voltage drift during the read or intrinsic electrical properties of individual pixels.

Dark images typically comprise of a long integration time obtained with the shutter closed so it contains only the bias image plus the signal that derives in the detector itself in the absence of flux from the telescope.

Imaging flat fields are typically observations of the twilight sky which thus contain the bias image, the dark image scaled to the relevant integration time and the flux from the telescope itself modulated by the beam vignetting and any pixel-to-pixel sensitivity fluctuations of the detector.

The bias image is an additive correction that adds a constant offset to every frame irrespective of the length of the integration or the incident flux. The dark current is an additive correction that scales with integration time but is independent of incident flux. The flat-field is a multiplicative correction which scales with the total incident illumination, so depends on flux and integration. We can use these differences between additive and

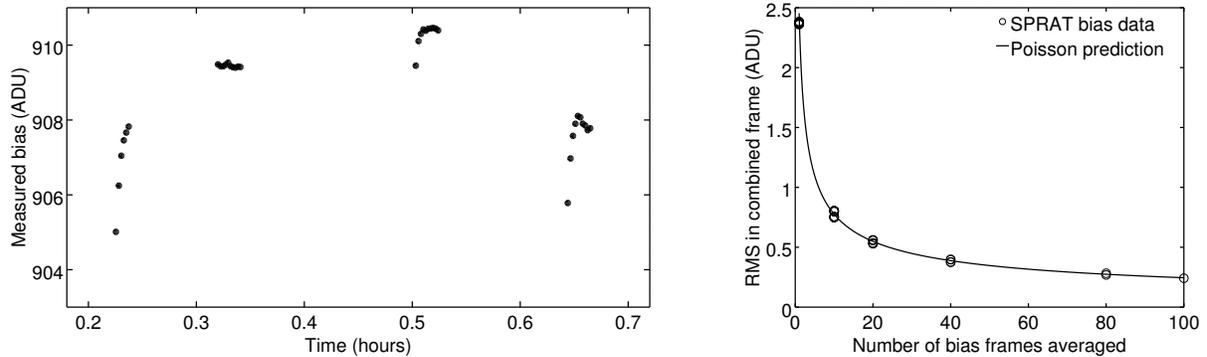


Figure 1. Left panel: Example data showing the bias does vary on short time scales. The behaviour exhibited here of the bias changing rapidly at the start of a consecutive sequence of exposures is typical of the camera and means that simply taking a bias frame immediately before the science frame does not work. Needs re-plotting. Right panel: Combining multiple bias frames reduces the pixel value RMS in excellent agreement with Poisson prediction assuming all structure in the bias image is amplifier read noise.

multiplicative terms to disentangle the multiple effects on a per frame basis without needing direct independent measurements of each.

### 3.1 SPRAT: Measuring Bias that Varies Frame-to-Frame Without an Overscan

Repeated, short exposures taken over a timescales of minutes to hours (Fig. 1) demonstrate the bias level varies sufficiently to require correction on a frame-by-frame basis. It is not sufficient to collect a bias exposure once per night or even immediately before the science observation. Consecutive frames have bias levels that differ by up to 1 ADU and the level can vary by up to 50 ADU from night to night. These data also show that the bias is extremely clean in terms having no photometrically significant 2D structure across the detector. This is demonstrated in Fig. 1 where we plot the standard deviation of pixel values in averaged stacks consisting of different numbers of bias frames. The actual observed standard deviation closely matches the prediction of Poisson statistics that when  $N$  frames are averaged, the standard deviation falls by a factor of  $\sqrt{N}$ . If the bias frame contained structure that was appearing consistently in multiple frames, the measured noise would fall more slowly than this prediction. We therefore do not need a 2D bias image and are able to use a single constant bias value for every pixel in the frame. Our task is to determine that value individually for each science frame taken.

With slit and grating deployed, there are regions of the detector that are completely unilluminated thanks to the slit geometry and 4000 Å long-pass, order-blocking filter. This dark region of the detector provides a direct measurement of the bias+dark signal in every observation and effectively provides the same functionality as would an overscan strip. However, with the slit and grating stowed, the entire detector area is used in imaging mode and the lack of overscan region becomes a problem.

It should be noted that in the absence of strong vignetting and with a “perfect” detector that shows little pixel-to-pixel quantum efficiency fluctuation then not knowing the bias voltage has little real impact on photometric accuracy. The bias in that case simply looks like extra background flux and is perfectly corrected during sky subtraction. Where the dynamic range of the flat-field is small, precision determination of the bias level is less important. In the case of SPRAT however the imaging mode is strongly vignetted with the throughput at the image edges being less than 25% of the frame centre. This vignetting does not effect the spectral mode because the slit is on axis, out of the vignetted region, but has a major impact on observational efficiency since our ability to use the imaging mode for target acquisition is constrained by field of view. The more field of view that can be accurately corrected for cross referencing to reference star catalogues, the better the effectiveness of the automated acquisition.

The strong vignetting in this case makes determination of the bias level possible. We know the geometric form the field vignetting from the twilight flat-field and we can measure its form in the image frame. The amplitude of

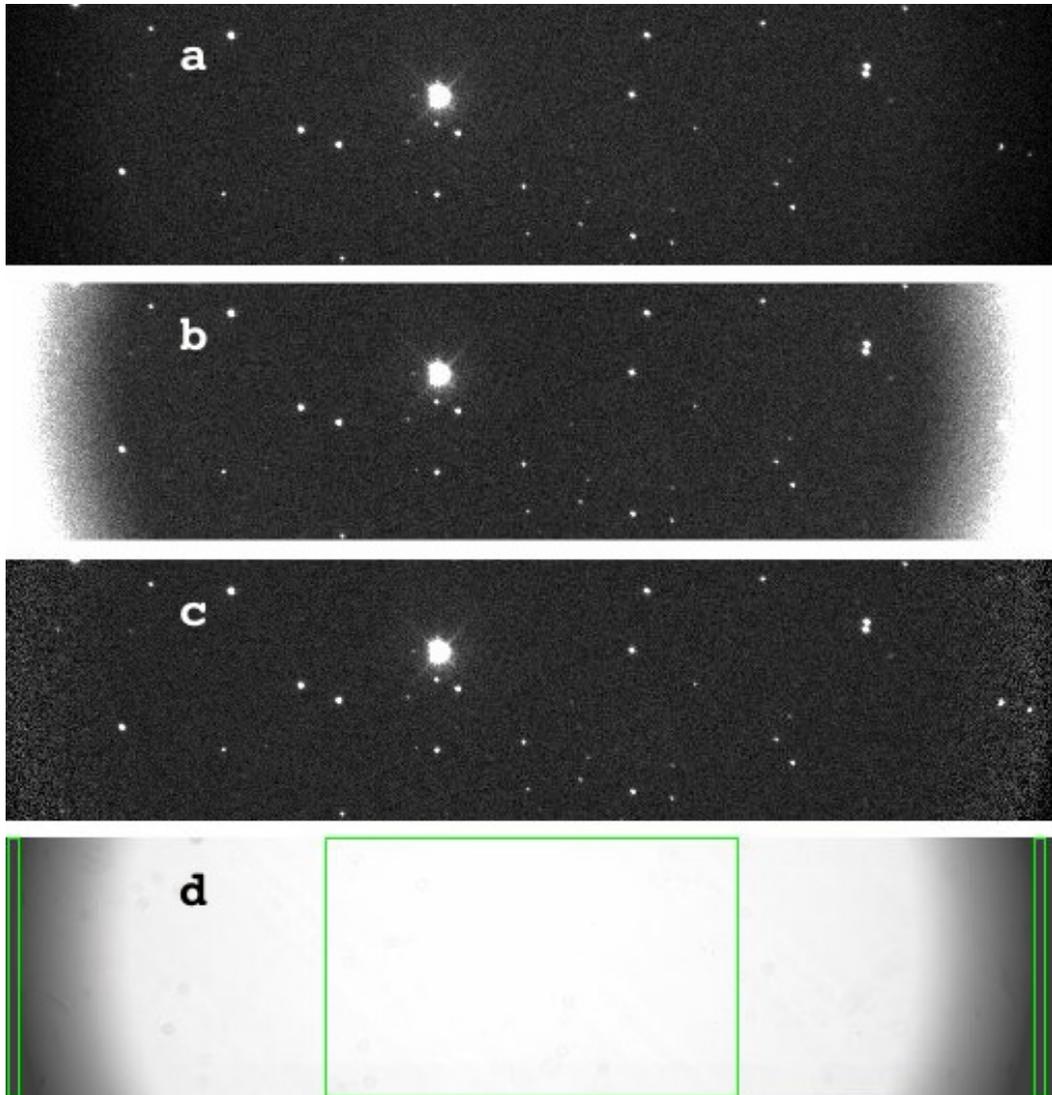


Figure 2. Example raw and reduced data from SPRAT in its imaging mode (grating and slit stowed). Raw data direct from the instrument are shown in the top panel (a). Panel (b) shows the gross over-correction of the flat-field when relying on a traditional bias frame and the bias level has drifted between exposures. Panel (c) is our final reduction, using a bias level inferred from the shape of the vignetting. The lower signal means the Poisson noise is exaggerated, but the mean sky level remains constant across the image. The bottom panel (d) is the flat-field used, to clearer illustrate the vignetting present in the raw data. The overlaid boxes are regions from which vignetting statistics are drawn as described in the text.

the vignetting function in the observation frame is diluted by the addition of uncorrected bias signal. The correct bias level may be algebraically and uniquely determined simply as the constant signal that must be subtracted off to make the vignetting match between the data frame and the flat. In other words, it is simply the bias level that once it is subtracted off allows the flat-field to most effectively create a featureless sky background.

Besides the raw data to be processed, the required inputs are a previously obtained dark frame and a flat-field. We store the dark frame on disk with the bias already subtracted. The flat field is stored on disk bias- and dark-subtracted and normalized to have mean counts of unity in the unvignetted field centre.

Given these inputs, the algorithm runs as follows.

1. A dark frame is selected from the library, choosing the available frame obtained closest in time to the image being processed.
2. Determine the mean dark current per pixel  $D$  as the average pixel value of the whole frame.
3. A twilight flat-field is selected from the library, choosing the available frame obtained closest in time to the image being processed.
4. Pixel value statistics are then taken from two regions. The mean pixel value in the unvignetted image centre and in a region near the image edge where vignetting is greatest are denoted  $F_0$  and  $F_i$  respectively. The precise regions selected are not critical. They should simply be large enough to provide a statistically fair sample of the throughput in that region and not so large that it varies greatly across the selected region. These values are used in the formula given below to represent the relative throughput of the field centre as compared to the vignetted region,  $F_0/F_i$ .
5. Repeating the same operation for the data being processed, the unvignetted sky brightness is  $S_0$  and the sky brightness measured at the frame edge is  $S_i$ . In order to fairly determine the sky background level, an iterative sigma-clipping algorithm is used to reject flux from stars in contrast to the simple mean used for the flat field frame.
6. The inferred bias value  $B_i$  is then given by

$$B_i = \frac{S_0 - S_i(F_0/F_i)}{1 - (F_0/F_i)} - \frac{T_{\text{data}}}{T_{\text{dark}}} D,$$

where  $T_{\text{data}}$  and  $T_{\text{dark}}$  are the integration times of the science exposure and the dark frame respectively.

7. Generate a dark frame that contains only the pixel-to-pixel dark current variations. I.e., constant bias and mean dark current levels have both been subtracted off.
8. Apply this dark frame to the data to remove signal from warm pixels.

This algorithm is a very simple implementation of the concept. It may be extended to use multiple sample regions at different places in the vignette pattern or even full 2-dimensional surface fits to the counts in the flat and the image data. For the specific case of SPRAT, a two region implementation was found to be computationally very fast and operationally robust. We use two vignetted regions, one on each edge of the frame and the bias level applied by the pipeline is the average of the two,  $B = (B_1 + B_2)/2$ . The image sections selected do not run quite to the very edge of the image, but are in the region where throughput is approximately in the range 25–30% of the field centre. The pixel regions used are shown in Fig 2, panel d. They are ten pixel columns wide and extend to the full height of the detector array.

### 3.2 LOTUS: Distinguishing Bias Voltage from Dark Current

Similarly to SPRAT, the CCD read-out electronics in LOTUS exhibit drifting of the bias level between frames requiring requiring its determination on a frame-by-frame basis. In this case however there is also significant 2-dimensional structure in the bias frame necessitating the creation of separate bias and dark images.

As well as having no overscan strip, the Starlight-Xpress Trius does not lend itself well to taking bias frames. The camera automatically switches between two read-out modes, one for long integrations and another for video frame-rate with short exposures. Since we did not want to use bias frames (by definition zero or very short exposures) obtained in one camera mode to process deep integrations taken in the other mode we need to be able to determine what fraction of the counts in any given pixel are bias signal and how much is dark current. We can then synthetically create separate bias and dark frames which can be applied independently by the pipeline software. We do this using their differing time dependence. Dark current scales linearly with exposure and bias is a constant.

The algorithm we use is to observe classical “dark frame” calibration images at a range of integration times. We can then derive the bias a dark current values in a manner analogous to up-the-ramp sampling used in CMOS

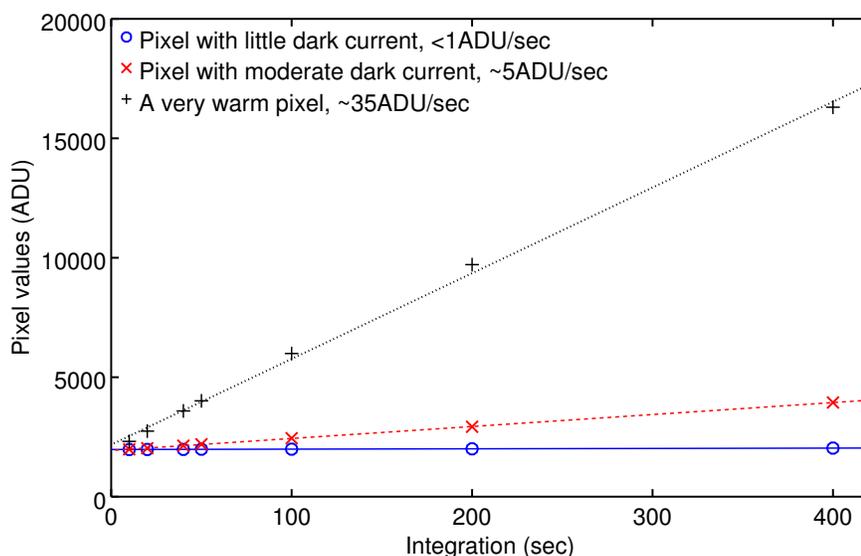


Figure 3. Fits of pixel value vs. integration time for three individual pixels on the LOTUS detector. For each pixel on the detector, the dark current (ADU/sec) is the fit gradient and bias level is the Integration=0 intercept value. The three pixels plotted were selected to illustrate a range of dark current values.

detectors. A linear fit to the observed signal as a function of integration time provides dark current (slope) and bias (zero exposure intercept) independently for each pixel. This is illustrated in Fig. 3 which shows the linear least-squares fits to the values from three different pixels. The same procedure is followed for every pixel in the image. The slope for each fit is the pixel's dark current and the intercept value is the bias.

Having created these calibrators the pipeline proceeds largely in the conventional way. The only exception is that we still need to estimate the constant pedestal level of the bias from each frame. Like SPRAT, the slit on LOTUS does not illuminate the entire frame so there is a section of the detector from which we can measure the mean bias pedestal level and set it to zero. The unusual feature here is that we must make this bias level correction after dark subtraction rather than the more typically used sequence of bias first. The reduction sequence is

1. A synthesized bias frame is selected from the library, choosing the available frame obtained closest in time to the image being processed.
2. Bias frame is subtracted from the data frame, removing the 2D bias structure.
3. A synthesized dark frame is selected from the library, choosing the available frame obtained closest in time to the image being processed.
4. The dark frame is scaled by the ratio of the dark and data exposure times ( $T_{\text{data}}/T_{\text{dark}}$ ) to generate a dark frame at the same integration as the data.
5. Dark frame is subtracted from the data frame.
6. Average signal level is measured from an un-illuminated image region which is being treated effectively as an overscan.
7. This single pedestal bias level is subtracted from or added to the entire data frame, depending on whether the bias level has drifted up or down since the reference bias frame was created.
8. Imaging mode data are flat fielded as normal. The following section addresses the flat fielding of spectral data.

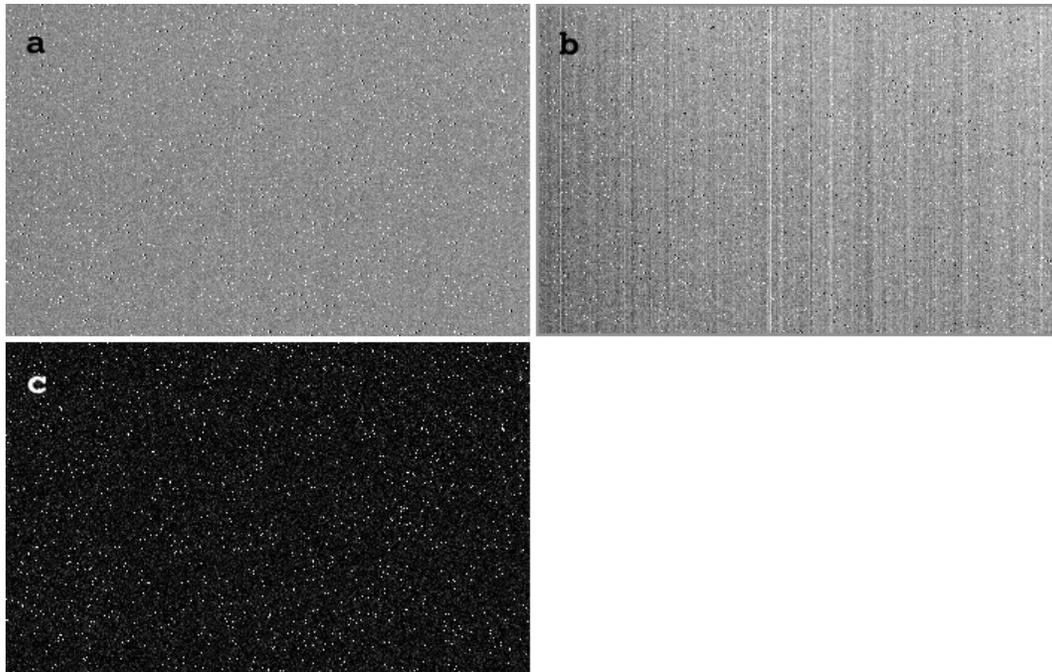


Figure 4. A demonstration of separating a long (200 sec) exposure into bias and dark components. Shown is a  $350 \times 220$  pixel section from the centre of the LOTUS CCD. Panel (a) is an unprocessed 200 sec exposure with the shutter closed. This contains contributions from both bias and dark. Exposures were collected at 10, 20, 40, 50, 100, 200, 400 sec. Panel (b) is the synthesized bias frame, showing that portion of the signal from (a) which does not change with increasing exposure time. This reveals the vertical stripes which are hidden by the hot pixels in the unprocessed frame. Panel (c) is the synthesized dark frame, containing the signal which does scale with exposure time.

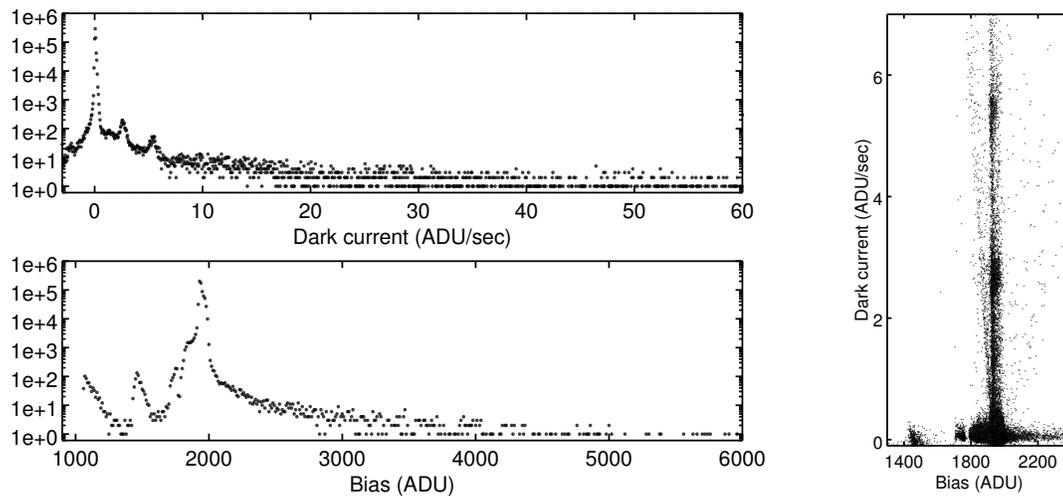


Figure 5. Frequency histograms of dark current (top-left) and bias (bottom-left) values for LOTUS. As seen in Fig. 4 panel (c), the great majority of pixels show almost no measurable dark current, but a significant minority of warm pixels generate as much as 60ADU/sec. The peaks in the bias plot (lower panel) at 1100 and 1800 ADU all arise in pixels adjacent to the detector edge so may be disregarded. The peaks at 2.5 and 5.5 ADU/sec in the dark current (top panel) are more interesting. They are warm pixels scattered apparently at random across the entire image area.

## 4. FLAT FIELDING

Flat fielding of a spectral image is a more complex problem than for a simple imaging instrument. As well as correcting the standard small spatial scale “pixel to pixel” variations, there are larger scale, wavelength dependent variations of the detector sensitivity. The flat field is also required to correct for any spatially dependent variations in the sensitivity of the instrument due to the effect of the spectrograph optics. These variations will typically be on a much larger spatial scale than pixel to pixel variations. Overall we expect the flat field may be affected by:

- Non-uniform illumination of the slit generated by vignetting and mis-alignment within the telescope as well as small scale dust and other contaminants.
- Vignetting by the optics within the instrument itself.
- De-focused images of surface contamination on the optics.
- Surface features on the CCD created during manufacture (polishing marks).
- Variation in sensitivity between individual pixels.

In the following sections we will describe methods for deriving flat-fields that will correct the pixel-to-pixel variation, larger scale errors and the effect of slit illumination, including the effect of geometric distortions on the derived corrections. A “final” flat field can be constructed by the multiplication of the various individual corrections.

### 4.1 Measuring Pixel to Pixel Variation Without the Slit and Dispersing Elements.

In some spectrographs (e.g., SPRAT, but not LOTUS) it is possible to withdraw the slit and dispersing elements from the optical beam. This allows measurement of detector flat field effects using an undispersed twilight sky image. Figure 6 shows an example SPRAT exposure taken in such a configuration and indeed is the flat-field used to correct data taken in imaging mode as described in Sec. 3 and Fig. 2. A large scale vignetting (which does not affect data taken through the slit) is apparent. The effect of the large scale vignetting can be removed by applying a 2D Gaussian smoothing function to the bias-subtracted image,  $I_{\text{sub}}$  producing a smoothed image,  $I_{\text{smo}}$ . The effect of this is to remove the pixel-to-pixel variations on spatial scales approximately less than the Gaussian full width half maximum (FWHM).  $I_{\text{smo}}$  is then divided into  $I_{\text{sub}}$  to produce a normalized pixel-scale flat field,  $F_{\text{pix}}$ . The effect of different levels of smoothing on the standard deviation (s.d.) of  $F_{\text{pix}}$  is shown in Fig. 7, with a knee at FWHM~2 pixels and flattening at FWHM~3.5 pixels. Larger scale smoothing FWHM (>5 pixels), causes a bleed effect from large features, particular evident at the image edges. For SPRAT a smoothing with FWHM = 3.5 pixels smoothing was chosen. Similar values are likely to be appropriate for most systems. The objective is simply to select a smoothing length that captures all the pixel-scale signal intrinsic to the detector whilst rejecting the larger scale features caused primarily in the optical beam.

### 4.2 The Effect of Geometric Distortions

Geometric distortions of the instrument focal plane are a common feature of long slit spectrographs. Care is usually taken to align the rotations of the slit, grating and detector such that “on average” the spatial and dispersion directions are orthogonal and aligned with the CCD detector edges. In the discussion that follows we will assume that the spectrum is oriented on the detector such that the  $x$ -axis (roughly) corresponds to the dispersion (wavelength) direction and the  $y$ -axis (roughly) to the spatial (slit) direction. We will also assume we are always working with bias subtracted images.

In this scenario the tilts and curvatures of the local spatial and dispersion vectors with respect to the overall focal plane mean that (for example) the projected traces of imperfections and dust in the slit will tend to curve slightly across the detector (Figs 8,9). To add further complication, the trace of an astronomical source across the image will have an additional curvature due to the effect of differential atmospheric refraction.<sup>19</sup> The combination of these effects can lead to subtle errors in the large scale photometric calibration of the spectrum as a function of wavelength as the spectrum “curves through” the dispersed slit profile. In a similar fashion the

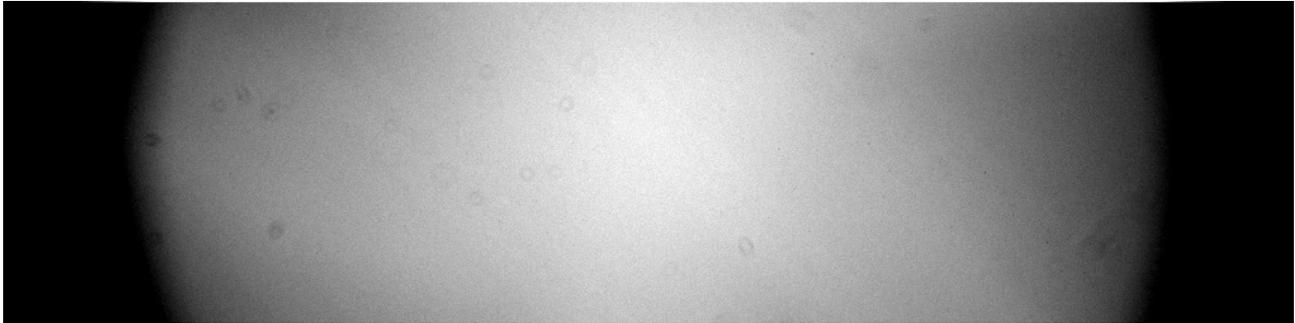


Figure 6. SPRAT sky flat in imaging mode with slit and dispersing elements removed from the beam. The image shows some of the features requiring correction. A large scale vignetting is evident across the field. This vignetting is caused by the size of the field lens, and does not affect spectral data where the slit restricts the field size to the region of the focal plane corresponding to the centre of the image. De-focused dust produces small circular features. Manufacturing polishing marks on the CCD chip surface create faint  $45^\circ$  diagonal features.

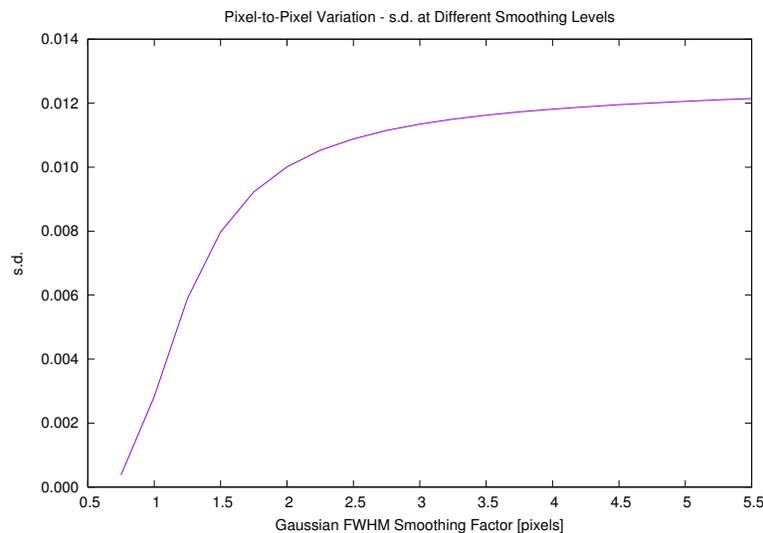


Figure 7. Flat field pixel value standard deviation vs Gaussian FWHM smoothing applied. In this graph, we aim to maximize the measured standard deviation (indicating we have captured the full pixel-to-pixel variation) whilst keeping the length scale smaller than, and thus ignoring, features which are not intrinsic to the detector. We selected FWHM=3.5 pixels.

wavelength calibration (and hence location of sky emission lines) may curve slightly as a function of the spatial direction due to optical distortion. This can lead to errors in sky subtraction.

The most common type of flat field (which is essentially available to any spectrograph, including LOTUS and SPRAT) is a twilight “sky” flat taken through the slit and dispersive elements. Such sky flats are illuminated in a similar fashion to incoming stellar light, but contain various time-dependent emission and absorption lines and bands (Fig. 9). Some spectrographs also provide the ability to illuminate the slit via an internal tungsten or similar incandescent bulb. Such “lamp” flats have the advantage of a smooth spectral energy distribution (SED) but the disadvantage that the illumination of the slit is usually not via the telescope optics, and therefore may be different to incoming stellar light (Fig. 8). Such a tungsten lamp facility is available on SPRAT but not LOTUS.

### 4.3 Use of the Lamp Flat

Nominally lacking emission lines and with a smooth SED, the lamp flat is well suited to the derivation of pixel and larger scale flat field features (e.g. CCD polishing marks). To do this we must remove the contaminating

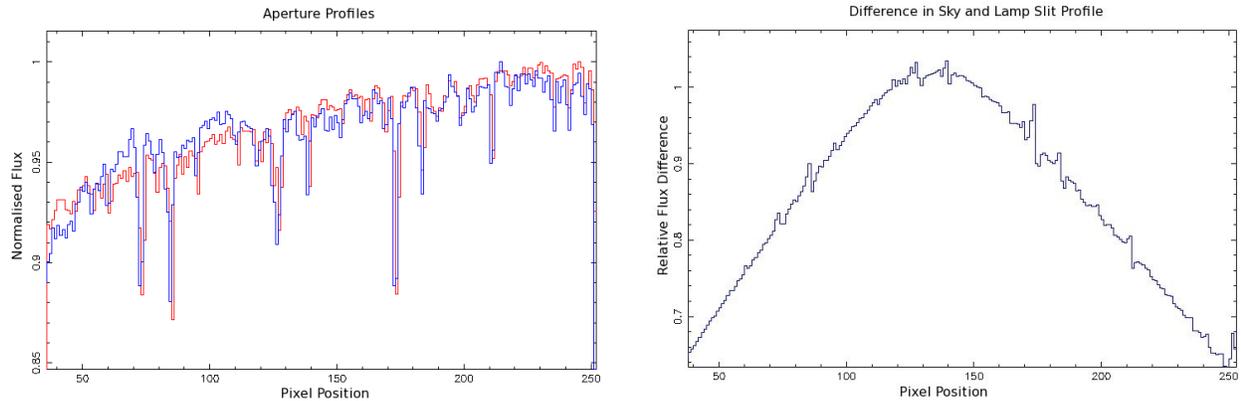


Figure 8. Left panel: Slit profiles extracted from 30 pixel wide sections at the red and blue ends of a SPRAT twilight sky spectrum. The 1-3 pixel wide decreases in flux are caused by dust particle contamination on the slit. The relative shift of these features is indicative of misalignment and distortion. The overall slope of the line showing a slight brightening of the slit from top to bottom is caused by non-uniform slit width or illumination. Right panel: Ratio of collapsed lamp and sky illumination profiles of the SPRAT slit. The overall flux distribution incident on the slit from the flat lamp does not match that from the sky via the telescope, producing the large scale variation across the image. The smaller peaks show that dust line positions are not co-incident indicating probable illumination angle differences between the two flat sources.

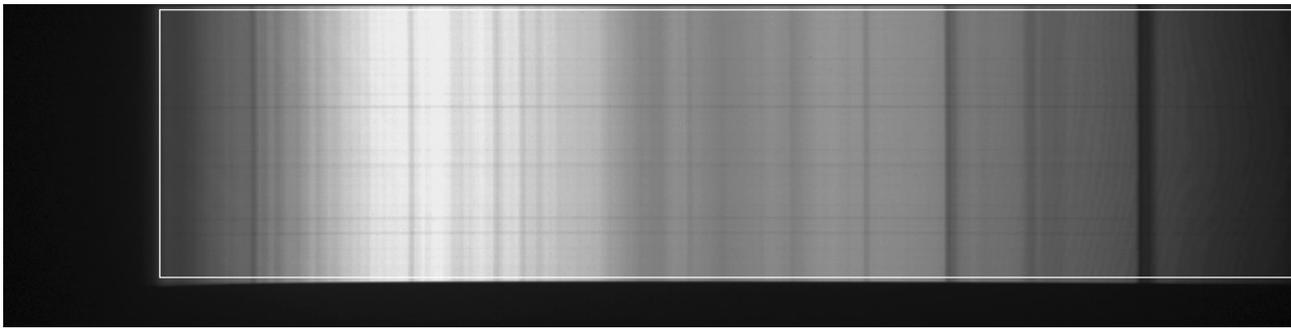


Figure 9. A twilight sky flat spectrum observed using SPRAT. The  $900 \times 216$  pixel region used in deriving a flat field is within the white rectangle. Solar Fraunhofer and telluric atmospheric lines are evident. The left black region marks the  $4000 \text{ \AA}$  high-pass filter cut-off wavelength and it is from this region that the bias level was estimated in Sec. 3.1. The lower black band is a mask on the slit. The horizontal dust lines and vertical absorption lines can be used to map geometric distortion (see Section 4.5).

effect of the slit spatial profile, which as already discussed does not accurately reflect the illumination pattern for sources observed through the telescope. We have evaluated four methods for this process, all of which begin with a bias subtracted lamp flat  $L(x, y)$  that has already been corrected for the effects of pixel to pixel variation using the flat field derived in Section 4.1. Images of the derived flats for all four methods are presented in Figure 10. The fringing on the right of the images is an interference effect between the surfaces of the CCD silicon and, being an additive effect, is not correctable by flat fielding. Our methods are:

1. **Division by  $x$ - and  $y$ - axis profiles.** A simple flat fielding method is to rescale each axis with an average illumination profile using a sky flat spectrum. A slit ( $y$ -axis) illumination profile,  $P_y$ , is created by summing  $L(x, y)$  along each row. Similarly, a wavelength ( $x$ -axis) illumination profile,  $P_x$  is created by summing  $L(x, y)$  along columns. Dividing each column of  $L(x, y)$  by  $P_x$  and each row by  $P_y$  then attempts to remove the spatial and wavelength dependent features, which following normalization to a mean pixel value of 1.0, forms the flat field  $F(x, y)$ . The outcome of this procedure is shown in Figure 10(top). As hoped for, the flat field shows the larger scale (e.g. polishing marks) features that we wished to characterize. However although the central region is well corrected for the effect of the slit profile there are

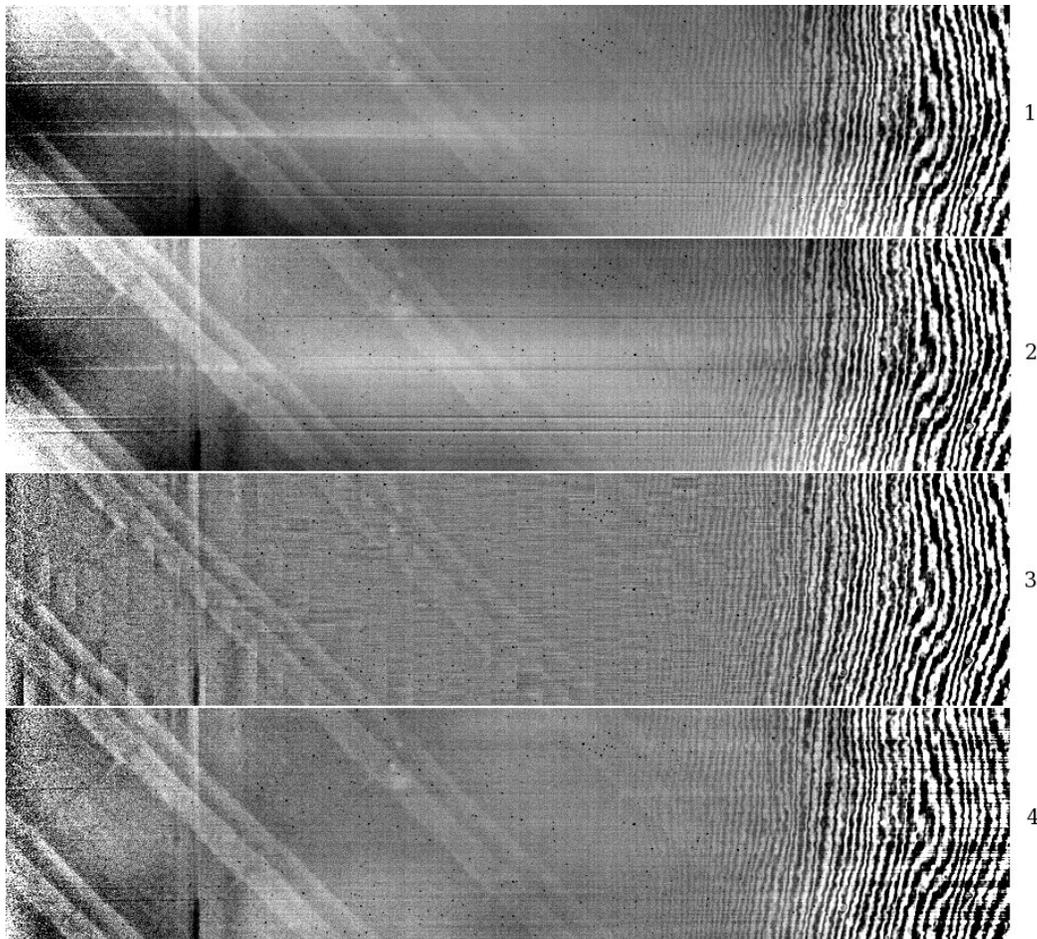


Figure 10. Flat fields derived from lamp flat observations using methods 1–4 (top–bottom). Method 1 assumes a constant mean slit profile for all positions on the  $x$ -axis. Method 2 derives the slit profile as a function of position on the  $x$ -axis via a linear interpolation of the profiles on the left and right hands sides. Method 3 derives individual slit profiles in windows of length 23 pixels along the  $x$ -axis. Method 4 fits a polynomial to each row in the image to derive how individual rows of the slit profile change as a function of  $x$ . The grey scale represents between  $-5\%$  (black) and  $+5\%$  (white) deviations from the whole-image mean.

obvious strong adjacent-row “positive-negative” slit residuals at the spectral extremes caused by geometric distortions.

2. **Application of a merged “left-right” profile.** Method 1 shows that application of single profile along an axis introduces errors if the image is distorted. This is particularly evident at the ends of dust and spectral lines. A method employing 2 merged slit profiles derived from each end of the slit was therefore investigated. This is identical to method 1 apart from in the construction of  $P_y$ . The correction applied in the spectral direction ( $P_x$ ) is the same in all four methods since we are investigating different methods of removing the slit illumination  $P_y$ . Two half image-width slit profiles were created, a left-hand side,  $P_{lhs}$ , and right-hand side,  $P_{rhs}$  profile. For an image of width  $W$ , the correction applied for a column at offset  $x$  is a linear proportion of the flux from the left and right profiles:

$$P_y(x) = \left(\frac{x}{W}\right)P_{lhs} + \left(1 - \frac{x}{W}\right)P_{rhs}. \quad (1)$$

The outcome of this procedure is shown in Figure 10(second from top). Comparison with method 1 shows a improved correction of the slit profile, although residuals associated with it are still present.

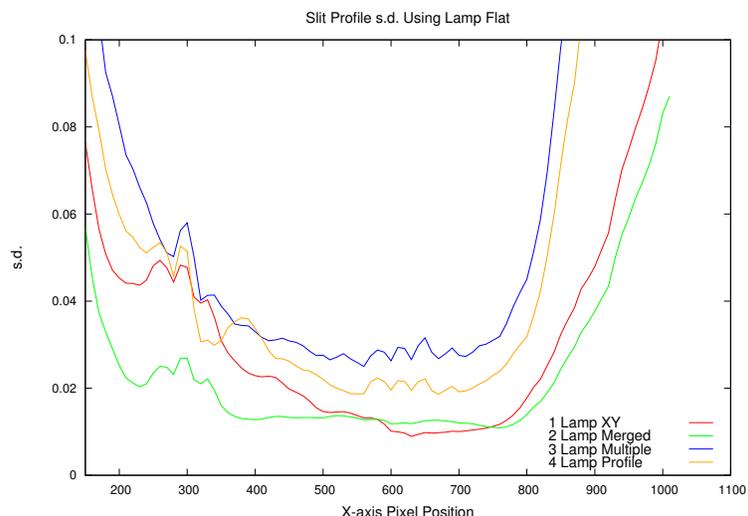


Figure 11. Standard deviation as a function of  $x$ -axis position for the four methods. Standard deviations have been measured over a sliding window of width 20 pixel and the full height of the detector.

Table 1. Standard Deviations of a  $50 \times 50$  pixel “featureless” region and (in parentheses) the whole image (apart from the image edges and fringed areas). The “Lamp” results are based on measurements of the flat-field. The “Sky” results are based on measurements of the residuals in the sky flat divided by the original input image. In both cases the results are presented for using both the normal and distortion corrected flat-fields.

Method	Normal Lamp	Normal Sky	Corrected Lamp	Corrected Sky
1. Conventional (XY)	0.007 (0.043)	0.024 (0.079)	0.008 (0.050)	0.024 (0.078)
2. LR Merge	0.008 (0.032)	0.026 (0.079)	0.012 (0.046)	0.025 (0.077)
3. Multiple Regions	0.022 (0.127)	0.029 (0.079)	0.009 (0.049)	0.034 (0.089)
4. Row Polynomial Fit	0.014 (0.103)	0.012 (0.087)	0.009 (0.022)	0.031 (0.086)

- Multiple profiles.** In this method a number of local slit profiles are generated and applied sequentially across an image. The slit profile is generated by collapsing image subsections, each of width  $L$  pixels. These are then joined together to form  $P_y(x)$  which is then applied as per methods 1 and 2. It can be seen as a potential improvement to method 2 in that it does not assume the contribution to a given pixel from a slit feature is a *linear* function of distance along the  $x$ -axis. The results of this approach are shown in Figure 10(third from top) for  $L = 23$  pixels. The residuals associated with the slit features are very well suppressed at all  $x$ -axis positions using this method. Another obvious advantage compared to the other methods is the overall very large scale “flatness” of the image. We note however that horizontal correlated noise of length  $L$  can be seen in the images, as well as some weak vertical features.
- Polynomial fitting to each row.** In this method  $L(x, y)$  is divided by  $P_x$  to remove the spectral lamp dependence. A fifth order polynomial is then fitted to each row in turn to model the slit response on a row by row basis. Row spectra generated from each polynomial are then multiplied by  $P_x$  to create an image which is then divided by  $L(x, y)$  to create the flat field. The results from this procedure are shown in Fig. 10(bottom). The procedure is more effective than methods 1 and 2 (but not as effective as 3) in removing the residuals due to the slit profile variation.

We now consider which of these methods is “best”. Overall the best flat field will be the one that accurately maps all defects affecting the image apart from any resulting from the slit profile. In Table 1 we list standard deviations for a  $50 \times 50$  pixel sub-region near the centre of the images which avoids any large scale features. Since

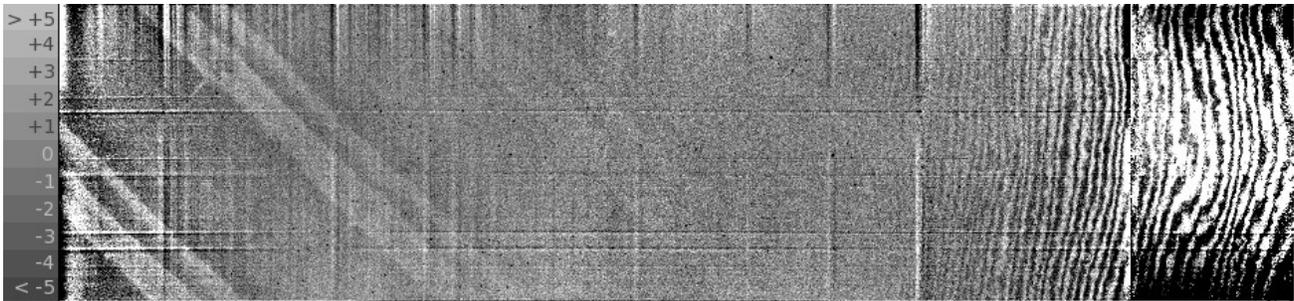


Figure 12. Division of a sky flat spectrum by X and Y profiles. Strong vertical features associated with features in the sky spectrum are apparent due to the geometric distortion.

we have already corrected our input image for pixel-to-pixel variations using the undispersed imaging flat, the *lowest* standard deviation within this small featureless region is therefore our target. This is assuming that any deviation present is caused by having imperfectly removed the slit illumination pattern. In this interpretation, the order of method preference (best to worst) is 1–2–4–3.

Looking to the large scales you can argue that we prefer a flat field that has the *largest* standard deviation, based on an interpretation that this is one is capturing the maximum amount of information about image structures. On this assumption (Tab. 1 (figures in parentheses)) the order of preference derived from “whole image” standard deviations would be 3–4–1–2. This argument is obviously flawed though if applied naïvely, since the algorithm might have introduced erroneous features. However, in the example of SPRAT data tested here, the same sequence (3–4–1–2) would be chosen by visual inspection of the images to evaluate qualitatively how well line residuals are removed.

Finally we consider the “medium” scale. In Fig. 11 we present the measured standard deviation as a function of  $x$ -axis position for the four methods. Standard deviations have been measured over a sliding window of width 20 pixels and full height of the detector. The low standard deviation for method 3 in this plot, especially at the extremes of the  $x$ -axis, demonstrates how well it removes the slit profile residuals.

Ultimately, the choice is best made through testing against simulated spectra rather than the simple, single parameter metrics discussed above, and this will be addressed in a future publication.

Overall the best choice appears to be method 3, with the caveat that this method appears to introduce additional small scale noise. The choice for any particular spectrograph will therefore depend on the relative importance of the slit profile and pixel to pixel noise to the overall flat field quality.

#### 4.4 Use of the sky flat

As described previously, the sky flat contains strong wavelength dependent features. It is however always available and does accurately reflect the slit profile on the recorded spectrum. Attempts to use our four methods on the sky flat in a similar fashion to the lamp flat do not succeed since the curvature of the vertical spectral features adds an additional component that generates strong vertical residuals. For example, the application of method 1 yields Figure 12. We therefore have focused our attention on using the sky flat to create a correction only for the effect of the slit profile. We used the same four methods as for the lamp flat, however the input image was the bias subtracted sky-flat and the resulting product was the  $P_y(x)$  images containing only the spatial contribution along the slit length. We do not include the contribution from  $P_x(y)$  along the spectral direction. These  $P_y(x)$  images are shown in Fig. 13.

To evaluate the performance of the four methods, we ratio the derived flat fields with the twilight sky flat. Such ratio images (Figure 14) are contaminated by uncorrected vertical features associated with the wavelength dependence of the sky spectrum. These vertical features should be ignored in the comparison of the images. The standard deviations for methods 1–4 in the  $50 \times 50$  pixel “featureless” and whole detector region are listed in Table 1. The whole detector values are meaningless as they are dominated by the spurious vertical residuals. However for the  $50 \times 50$  pixel regions there is an apparent preference for method 4 as that which adds the least additional pixel to pixel noise. Based on visual inspection of the images, method 3 seems to be the most

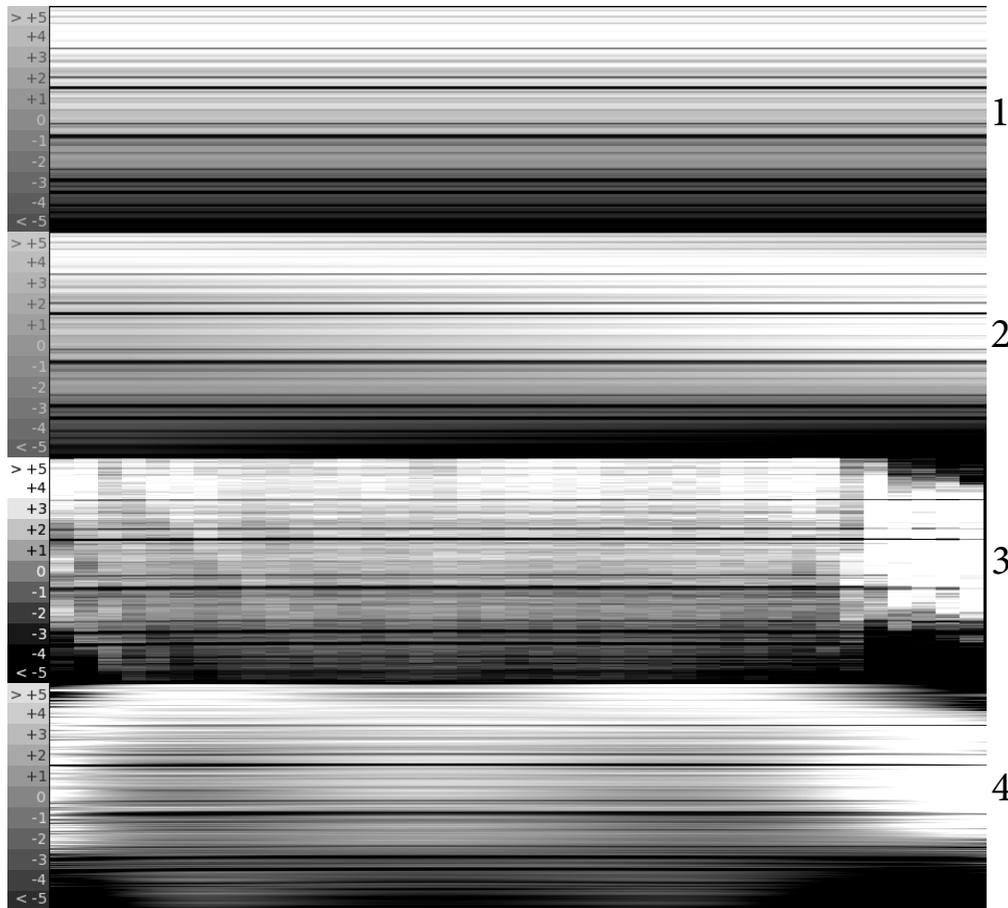


Figure 13. Flat fields generated using sky flat images and the derived  $y$ -axis profiles in methods 1-4.

effective at deriving a flat field that can remove the slit profile, again at the cost of it being the most noisy on a pixel-to-pixel level.

#### 4.5 An alternative approach? Correcting Image Distortion

Rather than trying to map the curved effect of the slit profile and remove it (for lamp flats) or use it (for sky flats), an alternative approach would be to apply a correction for the geometric distortion of the images to the data before using the conventional (method 1) approach. Either all data (science and calibration) could be corrected this way, or after deriving the flat field, it could then be “re-distorted” back to the original geometry before being applied to the science frames. This second approach would be computationally less intensive if processing a large number of science frames and prevents the need for resampling the science target pixel data which obviously risks degrading resolution. As we have already seen with a sky flat the diffraction grating disperses slit defects producing easily identifiable horizontal lines. Similarly, strong spectral lines provide a dispersed, vertical image of the aperture,  $\sim 4$ – $5$  pixels wide, (Fig. 9). These two sets of orthogonal lines extend across the image and their deviation from linear alignment with CCD pixels provides a measure of the distortion along the  $x$ - and  $y$ -axes (which also characterizes rotational skewing around the  $z$ -axis).

Using SPRAT data and various `figaro` routines we traced and corrected the horizontal and vertical lines in the sky flat sequentially. The correction used 5<sup>th</sup> order polynomials  $f(x)$ ,  $f(y)$  to characterize the distortion independently in each direction. We also wished to create the complement of the un-distort functions,  $\overline{[f(x)]}$ ,  $\overline{[f(y)]}$ . This allows the finally derived flat-field to be re-distorted for use in the existing data pipe-line. Initially  $\overline{[f(x)]}$ ,  $\overline{[f(y)]}$  were derived from a line template with simulated straight dust and spectral lines. The template

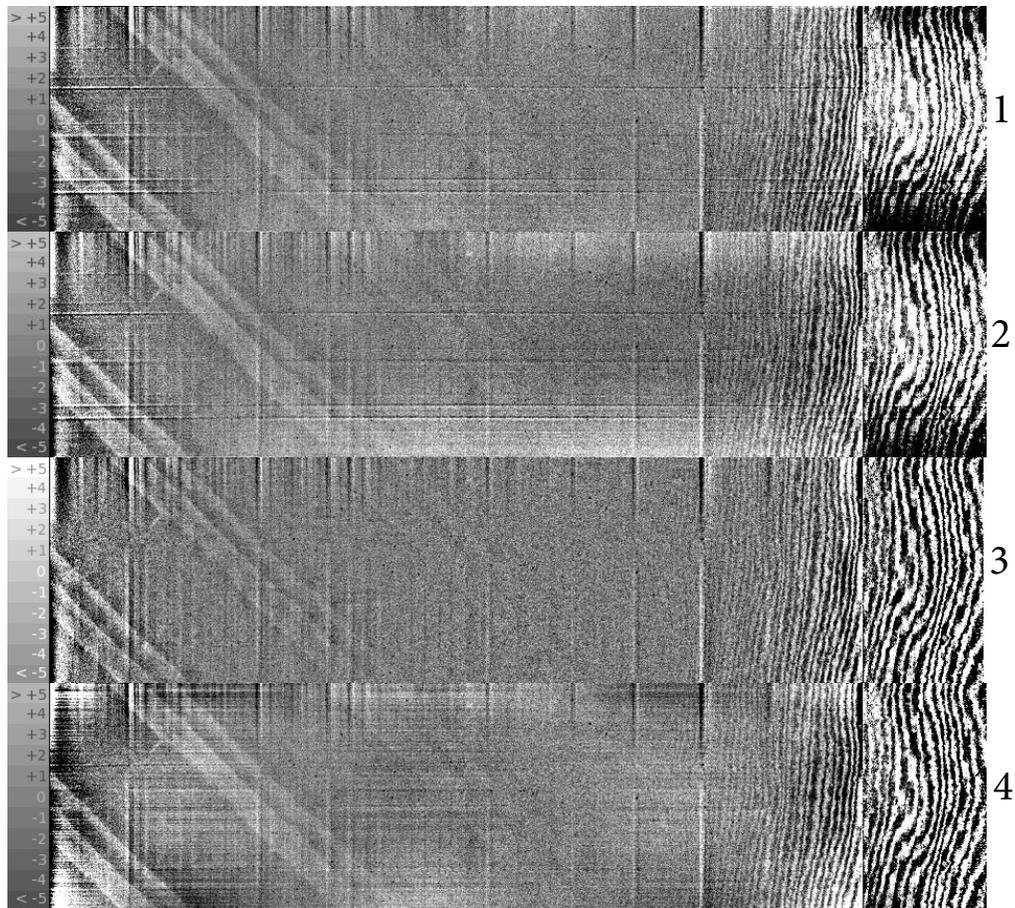


Figure 14. Ratio of flat fields from Figure 13 to raw sky flat field frame. The vertical features are expected residuals from the geometric distortion of sky features and this figure is intended only to assess the effectiveness of removing the slit illumination pattern by removing the horizontal features.

was distorted using  $[f(x)]$  &  $[f(y)]$  from a sky-flat to create a pre-distorted image. We then derived a second correction to the pre-distorted image to generate approximate inverse functions,  $[f(x)]$  &  $[f(y)]$ . Having used Legendre polynomials to characterize the distortion, creation of the inverse transformations is trivial by simply negating the polynomials coefficients.

Since the geometric correction involves a redistribution of flux across pixels, we wished to investigate its effect on small scale features in the images and whether it was fully reversible. For a uniform blank field set to pixel value 1.0 the standard deviation increased from 0.0 to 0.004 as a result of the distortion and undistortion of the image largely as a result of numerical rounding errors. For a field containing random values in the range 0.0–1.0 the standard deviation decreased from 0.280 to 0.226 illustrating there is some degree of smoothing. In a simulated one pixel wide line image it can be seen how flux redistribution results in a broadening of the line. The progressive change in cross-section for vertical line near the centre (where distortion is low) and edge (high distortion) of an image is shown in Figure 15. Overall the flux re-distributions associated with geometry corrections can be seen to distort pixel to pixel scale structures in an image. This is a natural consequence of the fact that a structure with a pixel scale of 1 pixel is by definition Nyquist under-sampled by a factor of 2. The effect on pixel to pixel scale flat fielding of geometric corrections is therefore potentially severe, in general resulting in an image smoothing. We note that it may be that an oversampling approach to flux redistribution would help alleviate this problem, and will be the subject of a future investigation.

With the above caveat in mind regarding the risks associated with flux re-distribution, we investigated the

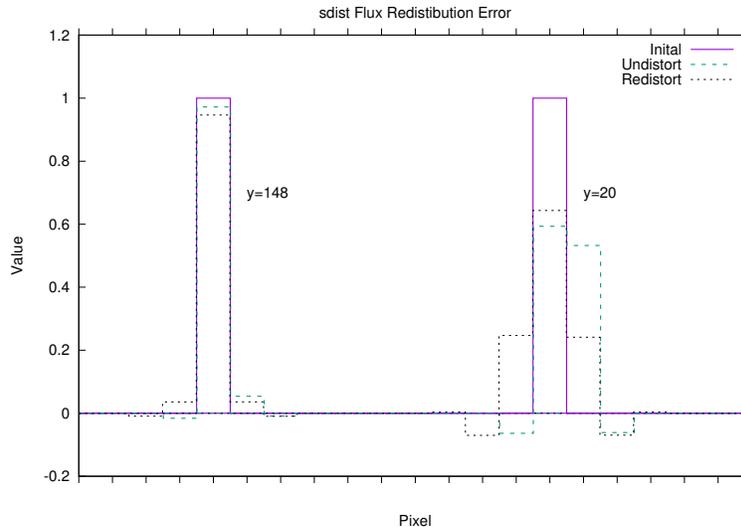


Figure 15. The progressive effect of undistorting and re-distorting for a vertical line in the simulated line field. Left shows broadening of a line near the centre at  $y=148$ , the overall amplitude is reduced by 3% and the width increase to  $\sim 3$  pixels. Towards the lower edge of the image at  $y=20$  the amplitude is reduced by 40% and broadened to 5 pixels with overcompensation indicated by the negative values.

effect of correcting our lamp flat images for geometric distortions (with the distortion polynomials derived from our sky flats). The outcome of this investigation can be seen in Figure 16. Comparison with the original analysis of the lamp flat shows that the horizontal residuals associated with the slit profile are indeed improved. The order of effectiveness of the different methods in slit profile removal is the same. The beneficial effect of the distortion on the correction of small scale variations can be seen in the increased apparent contrast of the dust features in this method. The standard deviations for the  $50 \times 50$  pixel “featureless” regions are listed in Table 1 and show lower noise than the normal methods, presumably due to the smoothing introduced by flux distribution. Standard deviation as a function of  $x$ -axis coordinate is plotted in Fig. 17. Comparison with Fig. 11 shows overall reduced standard deviations on the medium scale also.

In the specific case of SPRAT, this smoothing is not anticipated to be a problem and application of the geometrically distorted flat is potentially useful. The very first step of the procedure was removal of all the pixel-scale structure using the undispersed imaging flat. No genuine optical feature will be under-sampled at the detector due to the resolution of the grating, atmosphere and optics. Therefore any structure in the flatfield smaller than the gaussian filter applied in Sec. 4.1 must simply be noise and any slight smoothing caused by the distortion re-sampling is unlikely to be detrimental.

The same is not true of LOTUS where the undispersed imaging flat is not available. If the flat field generation procedure does not start by removing all the pixel structure and if for the detector in question the pixel to pixel and other unresolved features are significant in magnitude compared to the larger scale features such as slit illumination variations, then this procedure is probably best avoided.

## 5. SUMMARY / CONCLUSIONS

In this paper we present various software algorithms which we have tested in the data processing pipelines of the Liverpool Telescope. In particular it is valuable to note that trying out imaginative algorithms can increase the science impact of an instrument, particularly where it exhibits limitations or distortions for which the commonly used software tools do not provide facilities. Most of the time, excellent data reduction can be achieved using generic processing packages, but in some cases bespoke algorithms must be developed.

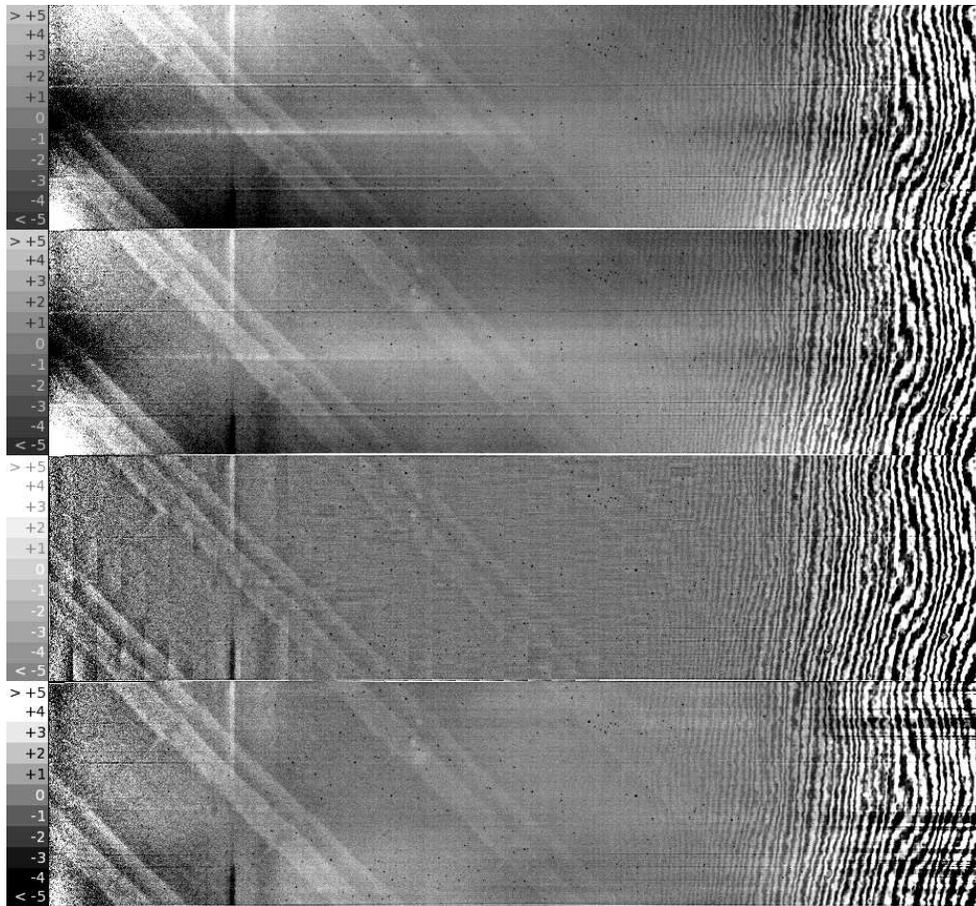


Figure 16. Flat fields generated using geometrically undistorted versions of the lamp flat images and then redistorted using methods 1–4. (Compare with Figure 10)

### 5.1 Bias and Dark Estimation

At instrument integration and commissioning the detectors in both SPRAT and LOTUS were found to exhibit large temporal variations in the bias voltage between observations and neither offered the tools common to many cameras for monitoring this fluctuation. We demonstrated two very closely related algorithms to perform this task, one optimized to each instrument. Both rely on the simple observation that bias voltage, dark current and flat fielding correction each have distinctive dependencies on integration time and incoming flux levels. For

Table 2. Taking SPRAT as one particular example, we list which of the various flat generation techniques contributes to correcting each feature in the raw data. The final flat applied is simply the product of all the desired components.

Raw data feature	Method
Non-uniform slit illumination	$P_y$ slit illumination map from the twilight sky spectrum. Section 4.4.
Optical vignetting	$P_y$ slit illumination map from the twilight sky spectrum, though it does not significantly affect SPRAT. Section 4.4.
Optics surface contamination	Spectral lamp flat after removal of the slit pattern. Section 4.3.
Pixel QE variations	High-pass filtered, undispersed imaging flat. Section 4.1.

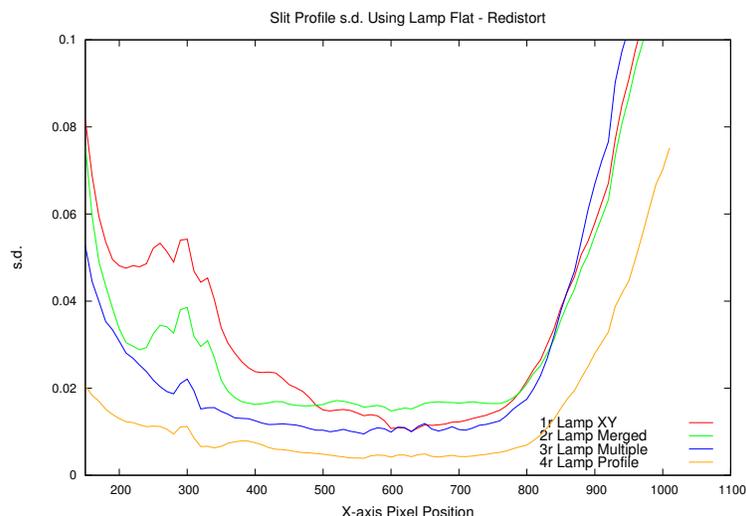


Figure 17. As Fig. 11 but for flat fields that have been undistorted and redistorted as part of their construction.

SPRAT we disentangle bias and flat-fielding terms by using the fact that bias voltage varies with time but is independent of the incident flux. For LOTUS we disentangle bias voltage from dark current by using the fact that dark current is proportional to exposure integration and bias is not.

## 5.2 Flat Fielding

The problem of flat-fielding when the spectrograph shows geometrical distortions in the focal plane was investigated. In particular we explored a number of methods of correcting for nearly-linear features associated with the slit profile. Methods based on a simple collapse down the wavelength axis, a merging of the slit profiles from either side of the detector and polynomial fitting down each row of the spectrograph gave increasingly better results. However the optimal method to remove the slit profile appeared to be dividing the spectrum into small ( $\sim 20$  pixel long) regions and deriving the slit correction locally for each region. A caveat to this conclusion was that the pixel-to-pixel standard deviation in the flat field appeared to increase using this method. The same four methods were also applied after correcting the images for the geometric distortions by line tracing and flux redistribution. In general this method showed a greater degree of correction while retaining the same sequence of effectiveness between the methods. However this method causes a loss of small spatial scale information due to under-sampling. The method chosen for a particular instrument will depend on the relative strength of the slit profile variations, the strength of the pixel scale variations and the degree of geometric distortion. We aim to investigate these trade offs further in a future paper using simulated data (where we will have knowledge of the “ground truth” of all of the contributing effects to the flat-field).

## ACKNOWLEDGMENTS

The Liverpool Telescope is operated on the island of La Palma by Liverpool John Moores University in the Spanish Observatorio del Roque de los Muchachos of the Instituto de Astrofísica de Canarias with financial support from the UK Science and Technology Facilities Council. ASP acknowledges financial support from the EU FP7.2 OPTICON project.

## REFERENCES

- [1] Steele, I. A., Smith, R. J., Rees, P. C., et al., “The Liverpool Telescope: performance and first results,” in [*Ground-based Telescopes*], Oschmann, Jr., J. M., ed., *Proc. SPIE* **5489**, 679–692 (Oct. 2004).
- [2] Smith, R. J., Clay, N. R., Fraser, S. N., Marchant, J. M., Moss, C. M., and Steele, I. A., “Switching the Liverpool Telescope from a full-service operating model to self-service,” in [*Observatory Operations: Strategies, Processes, and Systems III*], *Proc. SPIE* **7737**, 773711 (July 2010).

- [3] Fraser, S. and Steele, I. A., “Robotic telescope scheduling: the Liverpool Telescope experience,” in [*Optimizing Scientific Return for Astronomy through Information Technologies*], Quinn, P. J. and Bridger, A., eds., *Proc. SPIE* **5493**, 331–340 (Sept. 2004).
- [4] Steele, I. A., Mottram, C. J., Smith, R. J., and Barnsley, R. M., “Using dummy and pseudo-dummy amplifiers to correct for common mode CCD noise,” in [*High Energy, Optical, and Infrared Detectors for Astronomy VI*], *Proc. SPIE* **9154**, 915428 (July 2014).
- [5] Barnsley, R. M., Jermak, H. E., Steele, I. A., Smith, R. J., Bates, S. D., and Mottram, C. J., “IO:I, a near-infrared camera for the Liverpool Telescope,” *Journal of Astronomical Telescopes, Instruments, and Systems* **2**, 015002 (Jan. 2016).
- [6] Barnsley, R. M., Steele, I. A., Bates, S. D., and Mottram, C. J., “An update on the development of IO:I: a NIR imager for the Liverpool Telescope,” in [*Ground-based and Airborne Instrumentation for Astronomy V*], *Proc. SPIE* **9147**, 914737 (July 2014).
- [7] Morales-Rueda, L., Carter, D., Steele, I. A., Charles, P. A., and Worswick, S., “The Liverpool Telescope Spectrograph: FRODOSpec,” *Astronomische Nachrichten* **325**, 215–215 (Mar. 2004).
- [8] Barnsley, R. M., Smith, R. J., and Steele, I. A., “A fully automated data reduction pipeline for the FRODOSpec integral field spectrograph,” *Astronomische Nachrichten* **333**, 101–117 (Feb. 2012).
- [9] Arnold, D. M., Steele, I. A., Bates, S. D., Mottram, C. J., and Smith, R. J., “RINGO3: a multi-colour fast response polarimeter,” in [*Ground-based and Airborne Instrumentation for Astronomy IV*], *Proc. SPIE* **8446**, 84462J (Sept. 2012).
- [10] Steele, I. A., Bates, S. D., Gibson, N., Keenan, F., Meaburn, J., Mottram, C. J., Pollacco, D., and Todd, I., “RISE: a fast-readout imager for exoplanet transit timing,” in [*Ground-based and Airborne Instrumentation for Astronomy II*], *Proc. SPIE* **7014**, 70146J (July 2008).
- [11] Piascik, A. S., Steele, I. A., Bates, S. D., Mottram, C. J., Smith, R. J., Barnsley, R. M., and Bolton, B., “SPRAT: Spectrograph for the Rapid Acquisition of Transients,” in [*Ground-based and Airborne Instrumentation for Astronomy V*], *Proc. SPIE* **9147**, 91478H (July 2014).
- [12] Steele, I. A., Marchant, J. M., Jermak, J. E., Barnsley, R. M., Bates, S. D., Clay, N. R., Fitzsimmons, A., Jehin, E., Jones, G., Mottram, C. J., Smith, R. J., Snodgrass, and de Val-Borro, “LOTUS: A low cost, ultraviolet spectrograph,” *MNRAS* (2016 in press).
- [13] Jewitt, D., “Project Pan-STARRS and the Outer Solar System,” *Earth Moon and Planets* **92**, 465–476 (June 2003).
- [14] Law, N. M., Kulkarni, S. R., Dekany, R. G., et al., “The Palomar Transient Factory: System Overview, Performance, and First Results,” *Proc. Ast. Soc. Pacific* **121**, 1395–1408 (Dec. 2009).
- [15] Prusti, T., “The promises of Gaia,” *Astronomische Nachrichten* **333**, 453 (June 2012).
- [16] Glassmeier, K.-H., Boehnhardt, H., Koschny, D., Kührt, E., and Richter, I., “The Rosetta Mission: Flying Towards the Origin of the Solar System,” *SSRv* **128**, 1–21 (Feb. 2007).
- [17] Taylor, M. G. G. T., Alexander, C., Altobelli, N., Fulle, M., Fulchignoni, M., Grün, E., and Weissman, P., “Rosetta begins its Comet Tale,” *Science* **347**, 387–387 (Jan. 2015).
- [18] A’Hearn, M. F., Millis, R. C., Schleicher, D. O., Osip, D. J., and Birch, P. V., “The ensemble properties of comets: Results from narrowband photometry of 85 comets, 1976-1992.,” *Icarus* **118**, 223–270 (Dec. 1995).
- [19] Filippenko, A. V., “The importance of atmospheric differential refraction in spectrophotometry,” *Proc. Ast. Soc. Pacific* **94**, 715–721 (Aug. 1982).