

# Reduction of long slit spectra with IRAF

*Francisco Valdes*

IRAF Group, Central Computer Services, National Optical Astronomy Observatories  
P.O. Box 26732, Tucson, Arizona, 85726  
March 1986

## ABSTRACT

Tools for the reduction of long slit spectra within the Interactive Data Reduction and Analysis Facility (IRAF) at the National Optical Astronomy Observatory (NOAO) are described. The user interface (commands and special features) and the algorithms are discussed. Application of the reduction package to multi-slit images is briefly outlined. The author developed and supports the package at NOAO.

## 1. Introduction

This paper describes the tools currently available within the Interactive Data Reduction and Analysis Facility (IRAF) at the National Optical Astronomy Observatories (NOAO) for the reduction of long slit spectra. The reduction tools, called tasks, are organized as an IRAF package called **longslit**. The tasks in the package are summarized below.

|   |   |
|---|---|
| apdefine - Define apertures for 1D aperture extraction          | identify - Identify features                              |
| apextract - Extract 1D aperture spectra                         | illumination - Determine illumination calibration         |
| background - Fit and subtract a line or column background       | reidentify - Reidentify features                          |
| extinction - Apply atmospheric extinction corrections to images | response - Determine response calibration                 |
| fitcoords - Fit user coordinates to image coordinates           | setimhdr - Set longslit image header parameters           |
| fluxcalib - Apply flux calibration to images                    | transform - Transform longslit images to user coordinates |

Since there are many types of long slit spectra, detectors, and astronomical goals we do not describe a reduction procedure or path. Reduction manuals giving cookbook instructions for the reduction of certain types of data at NOAO are available from the Central Computer Services Division. Instead, each task is discussed separately. The primary emphasis is on the algorithms.

The following terminology is used in this paper. A *long slit spectrum* is a two dimensional image. The two image axes are called *axis 1* and *axis 2* and the pixel coordinates are given in terms of *columns* and *lines*. The long slit axes are called the *dispersion axis* and the *slit axis*. The reduction tasks do not require a particular orientation of the dispersion and slit axes, however, these axes should be fairly closely aligned with the image axes. **In the remainder of this paper the slit axis will correspond to image axis 1 and the dispersion axis with image axis 2.**

There are five types of operations performed by the tasks in the **longslit** package: (1) detector response calibration, (2) geometric distortion and coordinate rectification, (3) background sky subtraction, (4) flux calibration, and (5) aperture extraction of one dimensional spectra. These are listed in the order in which they are usually performed and in which they are discussed in this paper. There is also an initialization task, **setimhdr**, and a general routine, **icfit**, used in many of the long slit tasks. These are described first.

### SETIMHDR - Set long slit image header parameters

The tasks in the **longslit** package use information contained in the IRAF image header. The task **setimhdr** sets a required parameter in the image header advising the long slit tasks which image axis corresponds to the dispersion axis; the tasks work equally well with the dispersion axis aligned with the image lines or the image columns. This is generally the first task executed when reducing long slit spectra.

### ICFIT - The IRAF Interactive Curve Fitting routine

Many of the tasks in the IRAF which fit a one dimensional function utilize the same powerful interactive curve fitting routine called **icfit**. This routine allows the user to perform sophisticated function fitting interactively and graphically or to specify the function fitting parameters in advance and run the task non-interactively. That this routine is used in many tasks also has the advantage that the

user need not learn a new set of commands and features for each task requiring function fitting.

The features of the this curve fitting tool include:

- (1) A choice of four fitting functions; Chebyshev polynomial, Legendre polynomial, a linear spline, and a cubic spline.
- (2) A choice of the polynomial order or the number of spline pieces.
- (3) Deletion of individual points from the fit.
- (4) Selection of a sample or subset of points to be fit (excluding the rest).
- (5) Iterative deletion of points with large residuals from the fitted function.
- (6) Binning sets of neighboring points into averages or medians which are then fit instead of the individual points.

In addition to the above features the interactive graphics mode allows the user to:

- (1) Iterate any number of times on the fitting parameters.
- (2) Display the fit in several different ways; residuals, ratios, and the fit overplotted on the data points.
- (3) Manipulate the graphs using a large set of commands for forming and expanding any part of a graph for detailed examination.
- (4) Produce copies of the graphs with a snap-shot command.

For the applications described in this paper the most important features are the ability to adjust the function order, exclude bad points, and select subsets of points to be fit. Other useful features are taking the median or average of a set of points before fitting and iteratively rejecting deviant points. When used non-interactively the user selects the function and the order. The **longslit** tasks using the interactive curve fitting routine are **background**, **identify**, **illumination**, and **response**.

## 2. Detector Response Calibrations

The relative response of the pixels in the detector and the transmission of the spectrograph along the slit are generally not uniform. Outside of the **longslit** package are IRAF tasks for creating *flat fields* from quartz lamp calibration images which correct for small scale response variations. Flat fields, however, do not correct for spectrograph transmission variations or any large scale response patterns. The tasks **response** and **illumination** are specially designed for long slit spectra to correct both the small scale variations as well as larger scale response patterns and slit illumination and transmission effects.

These algorithms make the assumption that the wavelength and slit axis are very nearly aligned with the image lines and columns. If this is not true then the images must be aligned first or alternate response calibration methods used.

### RESPONSE - Determine response calibration

The task **response** is used with calibration images which (1) do not have any intrinsic structure along the slit dimension and (2) have a smooth spectrum without emission or absorption features. Typically the calibration images consist of quartz lamp exposures. The idea is to determine a response correction that turns an observed calibration image into one which is identical at all points along the slit.

From (1) a one dimensional spectrum is obtained by averaging along the slit; i.e. averaging the columns. Based on (2) a smoothing function is fit to the one dimensional spectrum to reduce noise and eliminate response effects which are coherent in wavelength such as fringing. The response correction for each pixel is then obtained by dividing each point along the slit (the columns) by the smoothed one dimensional spectrum.

The purpose of fitting a function to the one dimensional spectrum is to reduce noise and to remove coherent response effects which are not part of the true quartz spectrum. Examples of coherent response effects are fringing and regions of low or high response running along the slit dimension which are, therefore, not averaged out in the one dimensional spectrum. The choice of smoothing function is dictated by the behavior of the particular detector. Difficult cases are treated with the interactive graphical function fitting routine **icfit**. For the automated case the user specifies the smoothing function and order.

This calibration algorithm has the advantage of removing spatial frequencies at almost all scales; in particular, there is no modeling of the response pattern along the slit dimension. The only modeling is the fit to the **average** spectrum of the calibration source. In tests at NOAO this algorithm was able to reduce the response variations to less 0.2%, to correct for a broad diagonal region of low response in one of the CCD detectors (the CRYOCAM), and to remove strong fringing in spectra taken in the red portion of the spectrum where the detector is particularly subject to fringing.

One feature common to **response** and **illumination** is that the algorithm can be restricted to a section of the calibration image. The response corrections are then determined only within that section. If a response image does not exist initially then the response values outside the section are set to unity. If the response image does exist then the points outside the section are not changed. This feature is used with data containing several slits on one image such as produced by the multi-slit masks at Kitt Peak National Observatory.

When there are many calibration images this algorithm may be applied to each image separately or to an average of the images. If applied separately the response images may be averaged or applied to the appropriate long slit spectra; typically the one nearest the object exposure in time or telescope position. The task allows a list of calibration images from which a set of response corrections is determined.

Figure 1 shows a portion of an average quartz spectrum ratioed with the smooth fit to the spectrum. It is one of the graphs which can be produced with the **icfit** routine and, with the other figures in this paper, illustrates the formatting, zooming, and snap-shot capabilities in IRAF. The figure shows considerable structure of periodic high response lines and fringing which, because they are primarily aligned with the image lines, are still present in the average quartz spectrum. Note that this is not the response since it is the average of all the columns; an actual response column would have much larger variations including pixel-to-pixel response differences as well as large scale response patterns such as the diagonal structure mentioned previously.

### **ILLUMINATION - Determine illumination calibration**

The task **illumination** corrects for large scale variations along the slit and dispersion dimensions due to illumination or spectrograph transmission variations (often called the *slit profile*). When the detector response function is determined from quartz calibration images, using **response**, an illumination error may be introduced due to differences in the way the spectrograph is illuminated by the quartz lamp compared to that of an astronomical exposure. This violates the assumption that the calibration spectrum has no intrinsic structure along the slit. **Illumination** is also used when only the small scale response variations have been removed using a flat field correction.

The approach to determining the response correction is similar to that described for **response**. Namely, the response correction is the ratio of a calibration image to the expected calibration image. Again, the expected calibration image is that which has no structure along the slit. Calibration images may be quartz lamp exposures, assuming there is no illumination problem, and blank sky exposures. In the worst case, object exposures also may be used if the extent of the object in the slit is small.

There are several important differences between this algorithm and that of **response**:

- (1) The spectra are not required to be smooth in wavelength and may contain strong emission and absorption lines.
- (2) The response correction is a smooth, large scale function only.
- (3) Since the signal-to-noise of spectra from blank sky and object images is lower than quartz calibration images, steps must be taken to minimize noise.
- (4) Care must be taken that the spectral features do not affect the response determination.

The algorithm which satisfies these requirements is as follows. First the calibration spectrum is binned in wavelength. This addresses the signal-to-noise consideration (3) and is permitted because only large scale response variations are being determined (2). Next a smoothing function is fit along the slit dimension in each bin; i.e. each wavelength bin is smoothed to reduce noise and determine the large scale slit profile. Then each bin is normalized to the central point in the slit to remove the spectral signature of the calibration image. Finally, the binned response is interpolated back to the original image size.

The normalization to the central point in the slit is an assumption which limits the ability of the illumination algorithm to correct for all wavelength dependent response effects. There is a wavelength dependence, however, in that the slit profile is a function of the wavelength though normalized to unity at the central point of the slit.

The wavelength bins and bin widths need not be constant. The bins are chosen to sample the large scale variations in the slit profile as a function of wavelength, to obtain good signal statistics, and to avoid effects due to variations in the positions and widths of strong emission lines. This last point means that bin boundaries should not intersect strong emission lines though the bin itself may and should contain strong lines. Another way to put this criterion is that changes in the data in the wavelength bins should be small when the bin boundaries are changed slightly.

The bins may be set interactively using a graph of the average spectrum or automatically by dividing the dispersion axis into a specified number of equal width bins. When the number of bins is small (and the number of wavelength points in each bin is large) bin boundary effects are likely to be insignificant. A single bin consisting of all wavelengths, i.e. the sum of all the image lines, may be used if no wavelength dependence is expected in the response. Illumination effects introduced with **response**, however, appear as wavelength dependent variations in the slit profile.

Smoothing of each bin along the slit dimension is done with the interactive curve fitting routine. The curve fitting may be done graphically and interactively on any set of bins or automatically by specifying the function and order initially. The fitting should be done interactively (at least on the first bin) in order to exclude objects when the sky is not truly blank and contains faint objects or when object exposures must be used to determine the slit profile.

As with **response**, several blank sky images may be available (though this is less often true in practice). An illumination correction may be determined for each calibration image or one illumination correction may be computed from the average of the calibration images. Also the illumination response correction may be determined for only a section of the calibration image so as to be applicable to multi-slit data.

Figure 2 shows the fit to one of the wavelength bins; lines 1 to 150 have been summed and the sum is plotted as a function of slit position (column). The data is from a response image produced by **response**. This figure illustrates a number of things. **Illumination** may be run on a response image to remove the large scale illumination and slit transmission effects. This creates a flat field in a manner different than normal surface fitting. The figure shows that response effects occur at all scales (keeping in mind that the pixel-to-pixel response has been largely averaged out by summing 150 columns). It also illustrates how the illumination algorithm works for a typical slit profile. In this example about half the large scale variation in the slit profile is due to illumination effects and half is real slit transmission variations. For a blank sky or object image the main differences would be larger data values (hundreds to thousands) and possibly objects present in the slit to be excluded from the fit.

### 3. Distortion Corrections and Coordinate Transformations

The removal of geometric distortions and the application of coordinate transformations are closely related. Both involve applying a transformation to the observed image to form the desired final image. Generally, both steps are combined into a single image transformation producing distortion corrected images with linear wavelength coordinates (though the pixel interval may be logarithmic). This differs from other systems (for example, the Kitt Peak IPPS) which perform distortion corrections on each axis independently and then apply a dispersion correction on the distortion corrected image. While this approach is modular it requires several transformations of the images and does not couple the distortions in each dimension into a single two dimensional distortion.

To transform long slit images requires (1) identifying spectral features and measuring their positions in arc lamp or sky exposures at a number of points in the image, (2) determining the distortions in the slit positions at a number of points along the dispersion axis using either calibration images taken with special masks or narrow objects such as stars, (3) determining a transformation function between the image coordinates and the user coordinates for the measured wavelength and slit positions, (4) and interpolating the images to a uniform grid in the user coordinates according to the transformation function. The coordinate feature information and the transformation functions are stored in a database. If needed, the database may be examined and edited.

An important part of this task is the feature center determination. This algorithm is described in a separate section below.

#### IDENTIFY - Identify features

The tasks **identify** and **reidentify** are general tools used for one dimensional, multi-aperture, multi-slit, echelle, and long slit spectra. The tasks are also general in the sense that they are used to identify features in any one dimensional vector. For long slit reductions they are used to identify and trace objects in the slit and to identify, trace, and determine wavelength solutions for spectral features from arc calibration images and from sky and object exposures.

**Identify** is used to identify emission or absorption features in a one dimensional projection of an image. This projection consists of an image line or column or the average of many lines or columns. Averaging is used to increase the signal in weak features and provide better accuracy in determining the one dimensional positions of the features. The identified features are assigned user coordinates. The user coordinates will ultimately define the final coordinates of the rectified images.

For determining the distortions along the slit, the positions of object profiles or profiles obtained with multi-aperture masks in the slit are measured at a reference line. The user coordinates are then taken to be the positions at this reference line. The coordinate rectification will then correct for the distortion to bring the object positions at the other lines to the same position. (Note that it is feasible to make an actual coordinate transformation of the spatial axis to arc seconds or some other units).

For wavelength features arc calibration images are generally used, though sky and object exposures can also be used if necessary. After marking a number of spectral features and assigning them wavelength coordinates a *dispersion solution* can be computed relating the image coordinate to the wavelength;  $\lambda = f(l)$ , where  $\lambda$  is wavelength and  $l$  is the image line. The dispersion solution is determined using the **icfit** routines described earlier. This dispersion solution is used in the long slit package only as an aid in finding misidentified lines and to automatically add new features from a wavelength list. The dispersion solution actually used in transforming the images is a two dimensional function determined with the task **fitcoords**.

Figure 3 shows a graph from **identify** used on a Helium-Neon-Argon arc calibration image. Only three lines were identified interactively and the remainder were added automatically from a standard line list. Note that the abscissa is in wavelength units and the ordinate is displayed logarithmically. The latter again illustrates the flexibility the user has to modify the graph formats. Each marked feature is stored in a database and is automatically reidentified at other columns in the image with **reidentify**.

#### REIDENTIFY - Reidentify features

The task **reidentify** automatically reidentifies the spectral and object features and measures their positions at a number of other columns and lines starting from those identified interactively with **identify**. The algorithms and the feature information produced is the same as that of **identify** including averaging a number of lines or columns to enhance weak features. The automatic tracing can be set to stop or continue when a feature fails to be found in a new column or line; failure is defined by the position either becoming indeterminate or shifting by more than a specified amount (*radius* defined in the next section).

### CENTER1D - One dimensional feature centering

The one dimensional position of a feature is determined by solving the equation

$$\int (I - I_0) f(X - X_c) dX = 0 \quad (1)$$

where  $I$  is the intensity at position  $X$ ,  $I_0$  is the continuum intensity,  $X$  is the vector coordinate, and  $X_c$  is the desired feature position. The convolution function  $f(X - X_c)$  is a sawtooth as shown in figure 4. For absorption features the negative of this function is used. The figure defines the parameter  $fwidth$  which is set to be approximately the width of the feature. If it is too large the centering may be affected by neighboring features and if it is too small the accuracy is worse.

For emission features the continuum,  $I_0$ , is assumed to be zero. For absorption features the continuum is the maximum value in the region around the initial guess for  $X_c$ . The size of the region on each side of the initial guess is the sum of  $fwidth/2$ , to allow for the feature itself,  $radius$ , to allow for the uncertainty in the feature position, and  $fwidth$ , for a buffer. Admittedly this is not the best continuum but it contains the fewest assumptions and is tolerant of nearby contaminating features.

Equation (1) is solved iteratively starting with the initial position. When successive positions agree within 0.1% of a pixel the position is returned. If the position wanders further than the user defined distance  $radius$  from the initial guess or outside of the data vector then the position is considered to be indefinite.

### FITCOORDS - Fit user coordinates to image coordinates

Let us denote the image coordinates of a point in the two dimensional image as  $(c, l)$  where  $c$  is the column coordinate and  $l$  is the line coordinate. Similarly, denote the long slit coordinates as  $(s, \lambda)$  where  $s$  is the slit position and  $\lambda$  is the wavelength. The results of **identify** and **reidentify** is a set of points  $(c, l, s)$  and  $(c, l, \lambda)$  recorded in the database.

Two dimensional functions of the image coordinates are fit to the user coordinates for each set of slit and wavelength features,  $s = t_s(c, l)$  and  $\lambda = t_\lambda(c, l)$ , which are stored in the database. Note that the second function is a two dimensional dispersion solution. It is this function which is used to transform the long slit images to linear wavelength coordinates. Many images may be used to create a single transformation or each calibration images may be used separately to create a set of transformations.

This task has both an interactive and non-interactive mode. For the non-interactive mode the user specifies the transformation function, either a two dimensional Chebyshev or Legendre polynomial, and separate orders for the column and line axes. When run interactively the user can try different functions and orders, delete bad points, and examine the data and the transformation in a variety of graphical formats. The interactive option is quite useful in initially setting the transformation function parameters and deleting bad points. The two dimensional function fitting routine is similar in spirit to the **icfit** one dimensional function fitting routine. It is possible that this routine may find uses in other IRAF tasks.

Figure 5 shows a graph from **fitcoords**. The feature image coordinates of four objects in the slit (the first of which is very weak) from **identify** and **reidentify** are plotted. This information is used to measure the distortion of the spectrograph in the slit axis. This example shows particularly gross distortions; often the distortions would not be visible in such a graph, though expanding it would make the distortion visible. The transformation surface fit to this data removes this distortion almost entirely as seen in the residual plot of figure 6. Figure 7 shows the equivalent residual plot for the wavelength coordinates; a two dimensional dispersion solution.

### TRANSFORM - Transform long slit images to user coordinates

The coordinate transformations determined with the task **fitcoords** are read from the database. The transformations are evaluated on a grid of columns and lines,  $s_i = t_s(c_i, l_i)$  and  $\lambda_i = t_\lambda(c_i, l_i)$ . If no transformation is defined for a particular dimension then a unit transformation is used. If more than one transformation for a dimension is given then a set of points is computed for each transformation. The inverse transformations are obtained by fitting transformation functions of the same type and orders to the set of slit position and wavelength points. Note how this allows combining separate transformations into one inverse transformation.

The inverse transformations,  $c = t_c(s, \lambda)$  and  $l = t_l(s, \lambda)$ , are used to rectify a set of input images. The user specifies a linear grid for the transformed images by defining some subset of the starting and ending coordinates, the pixel interval, and the number of points. In addition the pixel interval can be specified to be logarithmic; used primarily on the wavelength axis for radial velocity studies. The inverse transformations define the image column and line to be interpolated in the input image. The user has the choice of several types of image interpolation; bilinear, bicubic, and biquintic polynomials and bicubic spline. In addition the interpolation can be specified to conserve flux by multiplying the interpolated value by the Jacobian of the transformation.

The wavelength of the first pixel and the pixel wavelength interval are recorded in image headers for later use in making plots and in the **onedspec** package. In addition a flag is set in the header indicating that the image has been dispersion corrected.

## 4. Background Subtraction

## BACKGROUND - Fit and subtract a line or column background

If required, the background sky at each wavelength is subtracted from the objects using regions of the slit not occupied by the object. This must be done on coordinate rectified images since the lines or columns of the image must correspond exactly to the same wavelength. A set of points along the slit dimension, which are representative of the background, are chosen interactively. Generally this will consist of two strips on either side of the object spectrum. At each wavelength a low order function is fit to the sky points and then subtracted from the entire line or column.

Ideally the response corrections and coordinate rectification will make the background sky constant at all points on the slit at each wavelength and the subtracted background is just a constant. However, if desired a higher order function may be used to correct for deficiencies in the data. A possible problem is focus variations which cause the width of the sky emission lines to vary along the slit. One may partially compensate for the focus variations by using a higher order background fitting function.

The background fitting uses the interactive curve fitting routine **icfit** described earlier. Figure 8 shows a graph from **background** illustrating how the user sets two sample regions defining the sky (indicated at the bottom of the graph).

## 5. Flux Calibration

### EXTINCTION - Apply atmospheric extinction corrections to images

A set of coordinate rectified images is corrected for atmospheric extinction with the task **extinction**. The extinction correction is given by the formula

$$\text{correction factor} = 10^{0.4 E_{\lambda} A}$$

where  $E_{\lambda}$  are tabulated extinctions values and  $A$  is the air mass of the observation (determined from information in the image header). The tabulated extinctions are interpolated to the wavelength of each pixel and the correction applied to the input pixel value to form the output pixel value. The user may supply the extinction table but generally a standard extinction table is used.

The air mass is sought in the image header under the keyword AIRMASS. If the air mass is not found then it is computed from the zenith distance, ZD, using the approximation formula from Allen's "Astrophysical Quantities", 1973, pages 125 and 133

$$A = (\cos(ZD)^2 + 2s + 1)^{1/2}$$

where  $s$ , the atmospheric scale height, is set to be 750. If the zenith distance is not found then it must be computed from the hour angle, the declination, and the observation latitude. The hour angle may be computed from the right ascension and the sidereal time. Computed quantities are recorded in the image header. Flags indicating extinction correction are also set in the image header.

### FLUXCALIB - Apply flux calibration to images

The specified images are flux calibrated using a flux calibration file derived with the **onedspec** package using standard stars. The standard stars are extracted from response corrected, coordinate rectified, and background subtracted long slit images using the tasks **apdefine** and **apextract**. The standard stars must not be extinction corrected because this is done by the **onedspec** flux calibration algorithms. The user may specify flux per unit wavelength,  $F_{\lambda}$ , or flux per unit frequency,  $F_{\nu}$ . The flux is computed using the exposure time and dispersion from the image headers and a flux calibration flag is set.

## 6. Extraction of One Dimensional Spectra

The user may wish to extract one dimensional spectra at various points along the slit. As mentioned earlier, this is necessary if observations of standard stars are to be used to calibrate the fluxes. The flux calibration values are determined from one dimensional spectra of standard stars using the **onedspec** package. The tools to extract one dimensional aperture spectra from long slit spectra are **apdefine** and **apextract**.

### APDEFINE - Define apertures for 1D aperture extraction

Extraction apertures are defined as a list consisting of an aperture number and lower and upper limits for the aperture. The aperture limits are specified as column or line positions which need not be integers. The user may create a file containing these aperture definitions with an editor or use the interactive graphics task **apdefine**.

**Apdefine** graphs the sum of a number of lines or columns (depending on the dispersion axis) and allows the user to interactively define and adjust apertures either with the cursor or using explicit commands. If an aperture definition file exists the apertures are

indicated on the graph initially. When the user is done a new aperture definition file is written.

### **APEXTRACT - Extract 1D aperture spectra**

One dimensional aperture spectra are extracted from a list of long slit images using an aperture definition file. The extraction consists of the sum of the pixels, including partial pixels, at each column or line along the dispersion axis between the aperture limits.

More sophisticated algorithms than simple strip extraction are available in IRAF and will soon be incorporated in the long slit package. The other extraction tasks trace the positions of features, i.e. the aperture is not fixed at certain columns or lines, and allow weighted extractions and detecting and removing bad pixels such as cosmic rays. The weighted extractions can be chosen to be optimal in a statistical sense.

### **Conclusion**

The IRAF long slit reduction tasks have been used at NOAO for about six months and have yielded good results. The package does not contain specific analysis tasks. Some analysis task will be added in time. The package is part of the software distributed with release of the IRAF. The author of this paper wrote and supports the tasks described here. Any comments are welcome.

### **Captions for Figures:**

Figure 1. Ratio of average quartz spectrum to fit of a 20 piece cubic spline for determination of response correction using **response**.

Figure 2. Fit of 4 piece cubic spline to the slit profile from the average of the first 150 lines in a response image using **illumination**.

Figure 3. Identification of emission lines from the central column of a Helium-Neon-Argon spectrum using task **identify**.

Figure 4. Sawtooth convolution function of width *fwidth* used in the profile centering algorithm.

Figure 5. Graph of stellar object positions identified with **identify**, traced with **reidentify**, and graphed by **fitcoords** showing the spectrograph distortions.

Figure 6. Residuals of the fit of a two dimensional 6th order Chebyshev polynomial to the data of figure 5 using **fitcoords**.

Figure 7. Residuals of the fit of a two dimensional 6th order Chebyshev polynomial to the image positions of wavelength features using **fitcoords**.

Figure 8. Constant background fit to a line of an object spectrum using **background**. The marks at the bottom of the graph indicate the set of points used in the fit.