

The Reduction of Fiber-Fed Echelle Spectrograph Data: Methods and an IDL-Based Solution Procedure

JEFFREY C. HALL¹

Lowell Observatory, Flagstaff, Arizona 86001
 Electronic mail: jch@olin.lowell.edu

ELIZA E. FULTON² AND DAVID P. HUENEMOERDER

Massachusetts Institute of Technology, Cambridge, Massachusetts 02139
 Electronic mail: eli@athena.mit.edu, dph@space.mit.edu

ALAN D. WELTY AND JAMES E. NEFF

Department of Astronomy and Astrophysics, Pennsylvania State University, University Park, Pennsylvania 16802
 Electronic mail: welty@astro.psu.edu, jneff@astro.psu.edu

Received 1993 August 13; accepted 1993 December 9

ABSTRACT. Echelle spectrograph data present several challenges in data reduction. In this paper we address the general problem of accurately extracting spectra from a night's worth of raw CCD fiber-fed echelle data frames. We first briefly review echelle spectroscopy: properties of the basic echelle spectrograph, how the orders are arranged on the CCD, and what demands and constraints this data format places on reduction algorithms. We then discuss solutions to the various problems for fiber-fed data, with particular emphasis on the removal of the scattered-light background. Finally, we discuss our implementation of these solutions. We have written a package using the Interactive Data Language (IDL) that uses the methods described in this paper to give accurate extractions of spectra from fiber-fed echelle frames with any number of spectral orders of arbitrary tilt and curvature. We describe how interested persons may obtain the package through anonymous FTP.

1. INTRODUCTION

The advent of echelle spectroscopy in the past few decades has provided a powerful tool for observing and analyzing the spectra of astronomical objects. Echelle spectrographs allow one to do high angular dispersion spectroscopy with extensive wavelength coverage. Substantial progress in the authors' areas of research on RS CVn binaries, FU Orionis stars, T Tauri stars, and long-term variability in solarlike stars has been made with the ability to observe simultaneously important spectral features formed in widely separated parts of the stellar atmosphere. These advantages, however, are accompanied by some complex problems in data reduction. Packing thirty or more spectral orders onto a single CCD presents several challenges, including the accurate removal of extraneous light and extraction of spectra from curved orders. In Sec. 2 of this paper we summarize the characteristics of the echelle spectrographs we use and briefly review the general format of echelle spectra. In Sec. 3 we discuss our reduction methods and present sample reductions for three spectra obtained with the different instruments. In Sec. 4 we describe an implementation of these routines in the Interactive Data Language (IDL).

¹Visiting observer, Kitt Peak National Observatory, National Optical Astronomy Observatories, which is operated by the association of Universities for research in Astronomy, Inc. (AURA), under cooperative agreement with the National Science Foundation.

²Lowell Observatory Summer Undergraduate Research Assistant supported by NSF Research Experience for Undergraduates funds under NSF Grant No. ATM-9114733.

2. THE FORMAT OF ECHELLE SPECTROGRAPH DATA

2.1 Instruments Discussed in This Paper

We present results in this paper based on data obtained with three different fiber-fed echelle spectrographs. We illustrate the variations on the basic echelle format and discuss the robustness of our reduction methods to these variations. These three instruments are the Penn State Fiber-Optic Echelle (FOE, located at Kitt Peak National Observatory; see Ramsey and Huenemoerder 1987), the Solar–Stellar Spectrograph (SSS, built by the High Altitude Observatory and located at Lowell Observatory's Anderson Mesa site near Flagstaff), and the Quick and Dirty Survey Spectrograph (QDSS, located at Penn State's Black Moshannon Observatory—an instrument that is better than its name implies!).

A summary of the characteristics of these instruments appears in Table 1, and the general optical layout is shown in Fig. 1 (this is the layout of the SSS echelle system). Starlight is fed from the telescope focus to the spectrograph by an optical fiber (typically of 200 to 320 μm diameter) and is collimated by lens *C* (a collimating mirror can also be used, as is the case with the FOE). It is diffracted by the echelle grating *E* and passes through a cross-dispersing prism *P* (in the case of the FOE and the SSS; the QDSS uses a reflection grating as its cross-dispersing element). The camera lens *L* focuses the resulting spectrum on a CCD detector. In Fig. 2 we show sample spectra (the frames are of different dimensions and are not displayed

TABLE 1
Properties of the Spectrographs and Raw Data Frames

	FOE	SSS	QDSS	Units
Echelle	79	79	79	ℓmm^{-1}
Cross disperser	prism	prism	grating	
Fibers	200	320	320	μm
Telescope	0.9m	1.1m	1.6m	
Raw frames	2048 ²	512 ² × 2	320 × 512	
Trim	900 × 430	512 × 400	512 × 230	
Overscan	900 × 32	none	none	
Gain	2.6	37	22.5	$e^- \text{ADU}^{-1}$
Read noise	4	42	40	e^-
# orders	33	20	22	

here with their true relative sizes). The FOE and SSS spectra show the enhanced order curvature caused by their prism cross dispersers, and the FOE and QDSS have much smaller order separation than does the SSS. General reduction techniques must account for these and other peculiarities of the instruments.

2.2 Format of the Echelle Spectrum

The format of the spectra in Fig. 2 is a result of the unique properties of the echelle grating. The diffraction of light by the echelle's surface is governed by the equation for a reflection grating

$$m\lambda = \sigma(\sin \alpha + \sin \beta), \quad (1)$$

where m is the diffraction order number, σ is the grating groove spacing, and α and β are the angles of incidence and reflection, respectively. (See inset of Fig. 1.) Rearranging this equation and differentiating with respect to λ shows that the angular dispersion of the diffracted light is

$$\frac{d\beta}{d\lambda} = \frac{m}{\sigma \cos \beta}. \quad (2)$$

Combining Eqs. (1) and (2) we can also write the angular dispersion as




FIG. 1—The basic configuration of optical elements in our echelle spectrographs. The components are the optical fiber, the collimator lens C, the echelle E, the cross disperser P, the camera lens L, and the CCD. The inset shows the echelle grating with the various angles labeled as they are referred to in the text.




FIG. 2—Raw CCD frames obtained with the three spectrographs discussed in the text. At top is a FOE spectrum of α Tau, at center is an SSS spectrum of the F8 dwarf HD 35296, and at bottom is a QDSS spectrum of the RS Canum Venaticorum binary AR Lac. Echelle order numbers m are indicated at the left of each spectrum. Important spectral features, labeled on the right, include the infrared triplet of singly ionized calcium in order 26, H α in order 34, He D₃ and Na D in order 38, and Ca II H and K in order 56. Apparent in the FOE and SSS spectra is the pronounced order curvature introduced by the prism cross disperser. The SSS has two CCD cameras; one obtains the echelle spectrum while the other obtains the single-grating Ca II H and K order which is grafted onto the main image by the observing software (this order is the single straight order at the bottom of the frame).

$$\frac{d\beta}{d\lambda} = \frac{\sin \alpha + \sin \beta}{\lambda \cos \beta}. \quad (3)$$

A useful result is obtained when the spectrograph is configured such that $\alpha = \beta$ (called *Littrow mode*)

$$\frac{d\beta}{d\lambda} = \frac{2 \tan \alpha}{\lambda}. \quad (4)$$

From Fig. 1 we can see that pure Littrow mode is of more theoretical than practical interest for the spectrograph configurations in question, since the diffracted beam must be separated from the collimated undiffracted beam. Our instruments operate in Quasi-Littrow Mode (QLM), where $\theta \equiv \alpha - \beta$ is small and the Littrow condition provides a good approximation to the true dispersion properties.

Equations (2) through (4) reveal that high angular dispersion can be achieved by working at large angles of incidence and diffraction. *Echelles* are gratings that typically have $\sigma=79$ lines mm^{-1} or 316 lines mm^{-1} and are *blazed*—ruled with tilted grooves—so that the blaze angle $\delta \equiv (\alpha + \beta)/2$ is large (typically $\delta=63^\circ 26'$; i.e., $\tan \delta = 2.00$).

Because of the large blaze angle, the maximum efficiency of an echelle lies at high orders of diffraction; for example, the FOE spectrum in Fig. 2 spans $m=25$ to $m=57$. We can see from Eq. (2) that for a given angle of incidence there are many wavelengths, corresponding to different values of m , that will be diffracted in the same direction. The wavelength difference between two such wavelengths λ_m and λ_{m+1} in successive orders is called the *free spectral range* and may be expressed simply as

$$\lambda_{m+1} - \lambda_m = \frac{\lambda_m}{m}. \quad (5)$$

Clearly for large m this is a problem; a singly dispersed echelle spectrum is a hopeless jumble of overlapping orders. The solution to the problem is the cross disperser, which reflects or diffracts the diffracted beam from the echelle in a direction perpendicular to the initial dispersion direction. This separates the previously superimposed orders so that they appear in parallel on the CCD as in Fig. 2.

In this way order overlap is removed and a large range of wavelengths is arranged in a 2D fashion ideally suited to a CCD. Other problems crop up, however: the orders will in general be both tilted and curved, and the separation between the orders will change as a function of order number. These effects depend on the type of cross disperser used (prism versus grating) and arise from the change in wavelength, and hence the change in angular dispersion of the diffracting elements, along each order. Additionally, grating cross dispersers have orders as well. So, for example, if one operates in the first order of the cross disperser and the echelle (actually the whole system up to and including the echelle) provides a broad enough range of wavelength (such that $\lambda_{\max} > 2\lambda_{\min}$), there will be contamination from the second order of the cross disperser. This is the case for the QDSS, although the effect is small. We will not discuss the detailed reasons for these effects here; the interested reader can find a thorough discussion in Schroeder (1987). Order curvature and tilt are clearly visible in Fig. 2, however, and they make plain the need for flexible reduction algorithms.

3. REDUCING THE DATA: PROBLEMS AND SOLUTIONS

With this brief background in place, we now tackle the problem of extracting the data from the frames in Fig. 2. To make the following discussion as concrete as possible, we will supplement the descriptions of the various algorithms with quantitative results of our reduction of these three frames.

Table 1 also gives basic parameters of the frames displayed in Fig. 2. The SSS frames are somewhat unusual: they are a composite obtained with two 512×512 CCDs. Columns 1 through 368 contain the “red” camera echelle spectrum and columns 369 through 400 contain the single-order “blue” camera spectrum of Ca II H and K. The FOE and QDSS data are taken with one CCD, with the FOE frames trimmed to keep only the subsection with useful signal.

3.1 Averaging the Data

Biases, flats, darks, and sometimes wavelength calibration frames should always be taken in sets and averaged to produce a reduced-noise composite. It is important to minimize the number of bad pixels in the composite because they will be propagated throughout the entire set of object frames if allowed to remain.

We average our frames pixel-by-pixel, using a bad-pixel filter to iteratively reject outlying values at each pixel. The outlier-rejection scheme removes contamination caused by cosmic ray hits. It will not remove chronically hot pixels or dead pixels, though; for cosmetically poor CCDs with bad pixels or columns a set of dark frames (see below) is usually needed to identify and mask out these pixels.

3.2 Debiasing the Data

This reduction step has two purposes: (1) quantification and subtraction of the electronic zero exposure level for a frame and (2) removal of any exposure-time-dependent dark-current effects present on the frame.

The first of these is achieved by subtracting the count level in the overscan strip (obtained by reading an extra set of columns after reading out the frame; this provides an estimate of the count level in an “empty” pixel) or a full two-dimensional bias frame (a frame produced by averaging together a number of zero-second “exposures” of the CCD with the camera shutter closed). Typically it is not necessary to remove both an overscan and a bias frame, although in some cases it may be (for example, if one is preflashing the CCD before starting the exposure). For the Kitt Peak TK2B used for the present FOE dataset, removal of the overscan from the averaged bias frame yielded a frame with count level of 0.26 ± 0.81 ADU and with normally distributed data values.

Using a two-dimensional bias frame to remove the electronic zero exposure level is slightly riskier: one must be sure the bias level does not change during the night (or one can take several sets of bias frames during a night if, for example, the CCD temperature is changing slightly). Also, one must be sure that the noise in the bias frame is below the read noise of the CCD, and this will not be the case if one is preflashing. We typically average 10 to 25 bias frames to produce a low-noise composite for full-frame subtraction.

Dark frames are used to remove the longer-term effects of dark-current accumulation of the CCD (e.g., hot pixels). This effect depends on the temperature of the CCD, and since it scales with time (perhaps nonlinearly), a truly

comprehensive dark frame set will have averaged exposures of a variety of durations (1800, 2600, 5400 s, and so forth) which allows accurate estimation of the scaling factor necessary to remove the dark-current effects from the frames. Development of such a dataset can be quite time consuming since several frames of each duration must be averaged to reduce the noise as well as to remove spurious cosmic ray events, but the calibration may be obtained during the daytime. In practice, it is common to assume linearity of the dark current with time, in which case it is removed by scaling an appropriately averaged dark current by the exposure time and subtracting it from the object frame.

3.3 Creation of an Order Map

We use our flat-field frames to locate the orders we wish to extract. We obtain the flat field by illuminating the input end of the optical fiber with a flat field and obtaining 10 to 20 strong exposures of the resultant spectrum. We average these to produce a clean composite. It is important that this composite have low noise. We also use the flat fields to make order maps; to maximize the accuracy of these order maps throughout the night we take a new set of flat fields every 2 to 3 hr during the night. We find that one or two sets of flat fields is *not* sufficient for accurate order identification over, say, a 10-hr night. Small drifts in the location of the orders are common for instruments with thermal flexure, and failing to keep track of these drifts with frequent flat-field sets during the course of the night can have disastrous effects on attempts to reduce the data.

Figure 3 shows the average of the central eleven rows of the QDSS flat fields. The figure is the classic echelle "dimetrodon plot," as we call it, and the diamonds at top mark the locations of the orders we wish to extract. These positions are found interactively or from a reference file loaded from disk. Using these positions as a starting point, our routine creates the order map in three steps, with the results of each step used as the basis for initiating the next step.

First we do a simple trace of the maxima along a given order. The routine traces to the ends of the CCD or until the absolute strength of the order or its strength relative to the background becomes too weak for accurate tracing.

We then fit Gaussians to the order profiles at a user-specified number of points along the order, using a least-squares fit that allows the peak value, centroid, and full width at half maximum (FWHM) of the profile to vary. This generates a set of points that contains accurate order centroids and FWHMs for the present order. A typical fit is displayed in Fig. 4. The rms error between these fits and the actual data points expressed as a percentage of the profile maximum is typically 2%–3% when one includes the noisier wings of the profile. If one considers only the profile core (about 5 pixels wide for the profiles in Fig. 4), the residuals are less than 0.4%: the instrumental profile of a fiber-fed echelle spectrograph is very nearly Gaussian!

Finally we perform a least-squares polynomial fit of user-specified order to the set of centroids generated in step




FIG. 3—A crosscut of a QDSS flat-field spectrum with 22 orders marked. The order-tracing routine must handle varying order widths and separation in addition to the tilt and curvature.

2 above. For all three of our datasets our fits have rms deviations from the set of centroids of typically less than 0.05 pixels for a polynomial of second order. Higher-order fits reduce this error by varying amounts depending on the data set; a 4th order fit to QDSS data yields rms's of about 0.01 pixel over the narrower orders.

This method is very rapid and consistently produces accurate order maps. Table 2 quantifies the quality of the fits found for the flat fields for each dataset. We store the map as a two-dimensional template of the actual locations, not just the fit coefficients, to ensure that it is applied with absolute consistency in every step in which it is needed.

3.4 Flat Fielding the Object Spectra

In general, the spectrum of a uniform source such as a flat field will not be a smooth continuum. Transmission variations in the fiber and the spectrograph optics, echelle fringing, and pixel-to-pixel gain variations in the CCD all




FIG. 4—The fit of a Gaussian to the H α order profiles of the flat-field spectrum obtained with the SSS. The data are shown by the solid line; the fit is shown by the dashed line. Fits to profiles from other instruments are of comparable quality. The fitted values in the profile cores are quite close to the actual data values.

TABLE 2
Accuracy of the Order Maps

	Gaussian fits			Polynomial fits to order 34 centroids			
	Order 26	Order 34	Order 46	10 pts 2nd ord	40 pts 2nd ord	10 pts 4th ord	40 pts 4th ord
FOE	1.45%	4.76%	0.91%	0.04 pix	0.03 pix	0.01 pix	0.01 pix
SSS	0.81%	1.35%	0.87%	0.03 pix	0.02 pix	0.02 pix	0.02 pix
QDSS	1.23%	3.50%	3.14%	0.02 pix	0.02 pix	0.01 pix	0.02 pix

contribute to creating unphysical variations in the data values along the order. The flat-fielding step in the reduction process removes these variations from the raw object frames before removal of scattered light and extraction of the spectrum.

We remove these variations from our object frames in a two-step process. We first extract a flat-field spectrum for each of the several averaged flat-field frames we have taken over a given night. We then divide each order in the object frames by the corresponding order of the flat-field spectrum, typically the one nearest in time to the object frame.

Creation of the flat-field spectrum: Three steps are necessary to prepare a flat-field spectrum for use.

First, the user must define the “aperture width” to be used for each order. This can be done interactively or with reference widths, determined during creation of the order map, loaded from a disk file. Assuming the order width does not change over a given order is reasonable for fiber-fed data, so long as the spectrograph is aligned and focused properly.

Second, the flat-field spectrum is extracted using a simple straight-sum extraction. For a given pixel j in the dispersion direction along order m we obtain the extracted flat-field spectrum data value F_{mj} as follows:

$$F_{mj} = \sum_{i=1}^N \epsilon_{mji} D_{mji}, \quad (6)$$

where there are N pixels across the profile (the specified aperture width) at pixel j along the dispersion direction of order m , D_{mji} is the data value of that pixel, and ϵ_{mji} is the fractional weight of the pixel. The result is a $J \times M$ flat-field spectrum, where J is the length of the order in pixels and M is the total number of orders being extracted.

We include the weighting factor to account for fractional pixels at the order boundaries. For example, if the centroid at a given location is at pixel 84.6 and an order width of 7 pixels is specified, the aperture will span pixels 81.1 to 88.1. In this case ϵ is 0.9 for pixel 81, 0.1 for pixel 88, and 1.0 for all other pixels. This method does not attempt to account for the distribution of light within the pixel, but since the flat-field spectra are well exposed, the count levels at the aperture edges are usually very small compared to the profile center and no significant errors are introduced. The high signal to noise (S/N) of these spectra also justifies the use of the straight-sum extraction of Eq. (6) (see the discussion of extraction in Sec. 3.6 for details).

The final (and optional) step in preparing the flat-field spectrum is to normalize the extracted flat-field orders to

unity. We use a cubic spline interpolation with a fairly widely spaced grid of user-selected fitting points to perform this normalization.

Application of the flat-field spectrum: Consider pixel j in the dispersion direction along order m of a given object spectrum. The extraction aperture, defined by the flat-field spectrum, is N pixels wide; the data values in the object spectrum over this aperture are S_{mji} . Then at each pixel i across the profile

$$S'_{mji} = \frac{S_{mji}}{F_{mj}}. \quad (7)$$

Dividing each pixel across the object profile by the single flat-field spectrum value at that pixel serves two purposes. First, the flux in the object spectrum is preserved since the flat-field spectrum is normalized. Second, the profile of the object spectrum is also preserved, since each pixel across the profile is divided by the same value.

Of concern here is that the various effects are indeed divided out by this method; we would expect variations across the profile to be lost in creation of the flat-field spectrum. This concern is mitigated in large part by the sameness of the fiber-fed profiles of the flat-field and object spectra. Any deviations in the profile of the flat field are echoed in the profile of the object and scale out when the flat-field division is performed.

Figure 5 shows a typical product of our flat-fielding algorithm. Three SSS spectra are displayed. At top is the normalized flat-field spectrum of the Ca II infrared triplet order. This spectrum shows a rather monstrous variation between pixels 260 and 360 and several other smaller variations. In the middle is the spectrum of HD 35296 without the flat-field correction applied, and at the bottom is the unnormalized, extracted spectrum of HD 35296 after division by the flat-field spectrum.

3.5 Removal of Background Scattered Light

3.5.1 Statement of the Problem

After removing the electronic background from an object frame, there will typically be a background remaining that arises from scattered light. Accurate estimation of the scattered light level at the center of each order to be extracted is a complex problem and constitutes the most involved step of the reduction procedure.

The problem is illustrated in Fig. 6. This figure is a subsection of the central crosscut of the α Tauri object frame. The three orders present are labeled $m-1$, m , and $m+1$. We wish to model the scattered light background




FIG. 5—An illustration of the problems in flat fielding. Extraction of the Ca II infrared triplet order flat-field spectrum in SSS data reveals a significant variation, apparent in the top plot (this flat-field spectrum is normalized to 1 to preserve counts upon division). However, there are minimal variations across the profile at a given pixel. The extracted spectrum of HD 35296 with no flat-field correction is the middle plot; Ca II $\lambda 8498$ and Ca II $\lambda 8542$ are clearly present at pixels 78 and 220, but are significantly corrupted. Finally, the lower plot is the HD 35296 spectrum with the flat-field spectrum divided out. The continuum is recovered, along with the infrared triplet lines and several weaker lines.

that is present at the base of order m . Let the peak intensity in order m be denoted I_m and let the half width at half maximum of the order be denoted W_m . The order is located at pixel x_m and the separation between it and the next order on the frame is Δx_m . The interorder minimum is located at pixel $x_{m+(1/2)}$ and the intensity at that point is $I_{m+(1/2)}$.

The background scattered light may be considered to consist of two components: a *global* background and a *local* background. The character of this background will be dependent to some extent on the configuration of the particular instrument, but its primary sources are the diffracting elements. Cardelli et al. (1990) have examined the scattered light sources for the Goddard High-Resolution Spectrograph using this two-component model. We present here a general model for optical echelle spectrographs.

In this model, we consider the global component to be a convolution of the echelle efficiency envelope across the entire frame, the flux distribution of the object, and any sources of ambient light in the spectrograph environment. As such it forms a smoothly (and usually slowly) varying surface across the frame. We consider the local component of the background at a given point between two orders to arise from scattering out of the adjacent orders only.

A simple solution to the problem of scattered light removal is to identify the counts at the interorder minima and pass a grid of these counts to a polynomial surface-fitting routine. In this case one will get a surface that matches the measured interorder minima and that when subtracted will leave the order bases aligned at zero counts.




FIG. 6—This figure illustrates the general problem of removing scattered light from an object frame. The goal is to produce model values S_m of the scattered light at the base of each order to be extracted from the measured values at the nearest interorder minima ($I_{m\pm 1/2}$). The problem is that the measured interorder minima might not reflect the actual scattered light background values for spectra where the orders are close together, because the scattered light profiles from the adjacent order overlap and artificially increase the count level at the minima.

However, this solution is applicable *only* when the echelle orders are well separated on the CCD, since it amounts to fitting only the global component of the scattered light background. If the orders are close enough so that the local scattered light profiles overlap, the measured interorder minima will not accurately reflect the true background value. In spectra with very closely spaced orders, as is the case with our FOE and QDSS data, this problem can be severe. The naive method described above will lead to serious oversubtraction of the scattered light background, which will in turn impact the accuracy of any optimal extraction that is attempted for the object (see Sec. 4) and will preclude accurate measurement of absolute line strengths.

3.5.2 Accurate Calculation of the Scattered Light Background

Gehren and Ponz (1986, hereafter GP) investigated the problem of scattered light on an echelle frame. They assumed three things: (a) the global background is constant; (b) the local background contributed by a given order is proportional to the peak intensity in that order; and (c) the local background at an interorder minimum is inversely proportional to some power of the order separation. These assumptions were motivated by the appearance of the background. The debiased and flat-fielded frame of α Tau has a scattered light background that depends at least in part on the strength of the nearby orders: the background is largest where the orders are strongest. GP found for their echelle frames that at a given crosscut on the frame the interorder minimum intensity could be fit at all orders by the relation

$$I_{m+(1/2)} = a + b \left(\frac{(I_m - a) + (I_{m+1} - a)}{[(1/2)\Delta x_m]^2} \right), \quad (8)$$

where a describes the constant global background and b describes the linear variation of the scattered light contribution. They found empirically that an inverse-square dependence on the order separation described the local scattered light satisfactorily. However, the frames from which they derived this relation covered only a small wavelength range ($\lambda 4000\text{--}\lambda 5000\text{ \AA}$), whereas our frames cover more than 5000 \AA . Although Eq. (8) leads to excellent background subtraction for the limited wavelength range case, it fails for the general case of full echelle frames such as we have with all three of our datasets. We have generalized GP's work as follows:

$$I_{m+(1/2)} = a + bX + c \left(\frac{m^4 W_m [I_m - (a + bX)]}{(\frac{1}{2}\Delta x_m)^2} + \frac{(m+1)^4 W_{m+1} [I_{m+1} - (a + bX)]}{(\frac{1}{2}\Delta x_m)^2} \right). \quad (9)$$

Equation (9) expands on GP's original relation in a number of ways. We now allow the global background to vary linearly ($a + bX$, where X is the coordinate of a given point on the crosscut). We allow for the possibility of changing order widths across the frame by incorporating the FWHM W_m into the intensity I_m available for scattering. Also, we assume the scattering follows a $\lambda^{-4} \propto m^4$ relation, a reasonable assumption for a scattering medium such as a prism, and include that in our formulation as well. We experimented with a number of relationships for the denominator of this expression, including different powers of the order separation as well as explicit profiles (e.g. Lorentzians and Gaussians). Although the Δx^{-2} dependence seems somewhat *ad hoc*, we find that it does give the best fit to the local scattered light background. (The fits obtained with Lorentzians, i.e., with a $[1 + (\Delta x_m / W_m)^2]^{-1}$ dependence, were good but definitely inferior to the Δx^{-2} fits.)

We find the constants a , b , and c with a nonlinear least-squares fit to the $I_{m+(1/2)}$ values for the given crosscut. Once these are known, the model background values S_m can be found as follows:

$$S_m = a + bX + c \left(\frac{(m-1)^4 W_{m-1} [I_{m-1} - (a + bX)]}{\Delta x_{m-1}^2} + \frac{(m+1)^4 W_{m+1} [I_{m+1} - (a + bX)]}{\Delta x_m^2} \right). \quad (10)$$

This relation describes the global component with the a and b terms and the local component by the c terms. In the last term we sum the local scattered light contribution to a given order from each of the adjacent orders only.

To calculate the full-two dimensional scattered light image, we perform fits to a number of cross-dispersion coordinates along the frame (obtaining different values of a , b , and c at each coordinate), regularize the resulting grid of coordinates and fitted data values (since it will be irregular if the orders are tilted or curved), and perform a 2D polynomial fit of user-specifiable order to the regularized grid. The resulting surface is subtracted from the frame and the

“descattered” image is written to disk. In our implementation of this equation (see Sec. 4 below) we also allow the user to vary the order of the fit to the global component; i.e., for certain frames Eq. (10) may need to be modified

$$S_m = a + bX + cX^2 + d[\text{local}], \quad (11)$$

where [local] is the quantity inside the large parentheses in Eq. (10).

If optimal extraction of the spectrum is to be performed (see Sec. 3.6) it is necessary to save the scattered light surface. The scattered light surface found for the FOE α Tau image is shown as a contour plot in Fig. 7.

3.5.3 Qualitative Assessment of the Procedure

Let us qualitatively examine the results of this algorithm applied to our datasets. Figure 8 shows the effect of Eqs. (9) and (10) applied to the FOE spectrum of α Tau. The measured interorder minima are marked with diamonds, and the fit to these minima [cf. Eq. (9)] is indicated by the boxes. The “X”s are the model order-center background, found via Eq. (10), and the dotted line running through them is the final fit (after grid regularization and 2D polynomial fitting).

The order separation in this frame decreases from right to left (i.e., from blue to red), and from pixel 300 to about pixel 100 the X's lie progressively farther below the total measured counts. This is expected since the scattered light profiles overlap more noticeably with decreasing order separation and since the order strengths themselves are increasing. The order separation continues to decrease slowly to the red end of the CCD at pixel 20, but the rapidly decreasing order strength more than compensates, and the X's and boxes reconverge. Note that between pixels 80 and 120 the actual background is only about 50% of the measured interorder minima.

In Fig. 9 we show similar plots for SSS and QDSS spectra. The symbols used in this figure are the same as in Fig. 8. The orders in the SSS data are more widely separated than the FOE orders, and as a result the model background never lies less than about 30% below the measured values.

The lower panel of Fig. 9 shows a crosscut of the QDSS AR Lacertae spectrum; this dramatically illustrates the ef-




FIG. 7—The scattered light image of the α Tau FOE spectrum produced by our algorithm. The contour reversal at bottom is a result of very slight ringing in the polynomial fit where there are no useful orders and consequently no fitting points. This does not impact the accuracy of the scattered light estimation in the useful part of the frame. The maximum contour level at top is 300 counts while the level running along pixel 300 is zero.




FIG. 8—A crosscut from the α Tau spectrum. Diamonds indicate the measured values of the interorder minima and boxes show the fit to the minima obtained with Eq. (9). The “X”s show the calculated order-center scattered light background [as per Eq. (10)] and the dotted line is the fit of the regularized, two-dimensional polynomial to this crosscut.

fect of decreasing order separation. The QDSS has a grating cross disperser, so the separation decreases rapidly toward the blue. As it does so, the model values fall away to about 40% of the measured values.

The fits (boxes) to the measured minima (diamonds) for the QDSS data are worse than those of the FOE data or the SSS data. The QDSS has noticeable ghost orders which complicate accurate determination of the data values at the interorder minima. These ghosts, which are visible between the redward orders in the lower panel of Fig. 9, become blended with the orders when the order separation becomes small. The fluctuations in the measured values for the blue orders are a result of overlapping order profiles blended with the ghost orders. Although the SSS data have more pronounced ghost orders (upper panel of Fig. 9), the primary orders are widely spaced enough to permit measurement of the interorder minima by tracing along the aperture edges and averaging the data values found. The fit to the measured minima is therefore better (see below).

3.5.4 Quantitative Assessment of the Procedure

FOE data: We constructed a scattered light image of the FOE α Tau spectrum using 40 crosscuts, a linear global background and a 4th order polynomial for the surface fit. The standard deviations of the fits to the measured interorder minima for the set of crosscuts varied from 4.5 counts to 30.3 counts, with an average over the entire frame of 16.8 counts, or about 0.3% of the maximum counts in the orders. (The crosscut in Fig. 8 shows the poorest fit, with a standard deviation of 30.3 counts.) When these uncertainties are propagated across the extraction apertures and compared with the total flux within the apertures, we find an average uncertainty of 1.3% across the entire frame.

SSS data: Despite the pronounced ghost orders in these data, we are able to estimate the interorder flux accurately at all but one point (at pixel 215). The standard deviation of the fit to the measured minima for the crosscut shown is





FIG. 9—This figure shows crosscuts of SSS data (upper panel) and QDSS data (lower panel) with the same symbols as in Fig. 8. The true background in the SSS data is closer to the measured interorder minima because the orders are more widely spaced. Confusion caused by the ghost orders and the more closely spaced main orders causes the fit to overestimate the background by about 5 counts in the first order. The fluctuations in the measured minima (indicated by the diamonds) for the QDSS data occur because the blue orders are extraordinarily close together—the actual order profiles overlap slightly in their wings—and because the ghost orders visible between the red orders at left become blended with the blue orders.

2.7 counts; the average standard deviation across the frame is 1.6 counts, which translates to average uncertainty of 0.2% over the entire frame. Again we performed the fit with a linear global background and a 4th-order surface polynomial. We did not fit the rightmost order (at pixel 383) since it is the single blue-camera order and is unrelated to the set of echelle orders on the rest of the frame.

QDSS data: The fit in the lower panel of Fig. 9 was made with a 0th-order (constant) global background and a 2nd-order surface polynomial. The uncertainty over the whole frame, calculated as for the other examples, is about 6%—clearly much poorer than for either the FOE or SSS data. The reason of this poor fit is the two minima at pixels 178 and 196, which are much higher than the surrounding minima. These fluctuations are probably not due to actual scattered light fluctuations; they more likely arise from the ghost orders and the small order separation (as discussed at the end of the previous section). There are a few such

deviations for each crosscut and their locations vary randomly among the ten blueward orders. Their effect for the displayed crosscut is to cause overestimation of the background between pixels 140 and 170. Additionally, the model finds a downturn in the scattered light in the last two orders (due to the decreasing flux at the order peaks), while the measured minima increase as the orders crowd closer together.

We conclude that the blue end of this spectrum cannot be fit accurately by our algorithm. The order profiles themselves overlap slightly; this fact and the ghost orders (visible at pixels 10 and 30) prevent estimation of the inter-order minima in this region. There are random fluctuations in the measured minima as well as a pronounced, large-scale upturn that is due to a $\approx 3\%$ overlap between the actual profiles rather than scattered light. Our routine allows the user to select an aperture range to be used for fitting the scattered light. When we select only the order between pixels 20 and 160 for fitting and use a low-order surface polynomial (1st or 2nd) the standard deviation of the fit drops from 41.3 counts to 9.6 counts, with an uncertainty over the explicitly fitted region of the frame of 0.7%. Using a low-order polynomial prevents ringing over the part of the frame where there are no fitting points.

It is possible to fit the entire background of this frame by using a 2nd-order global background—there is by inspection an overall quadratic trend to the minima. However, doing so produces unphysical results; the model background (shown by “X”s in the figure) ends up substantially higher than the measured minima. This is a clear indication of overfitting the global background and rendering the local fitting term meaningless. Ultimately it is up to the user to decide how to fit the global background in such a way that both maximizes the accuracy of the fit and preserves the accurate estimation of the local background.

General accuracy of the algorithm. We can assess the accuracy of our scattered light removal routine by considering the residual intensities in extracted heavily saturated lines. In a continuation of their initial study, Cardelli et al. (1993) have verified the accuracy of their algorithm using interstellar absorption lines. As this is not practical in the optical, we have used the atmospheric absorption bands between $\lambda 7590$ and $\lambda 7620$ to investigate the accuracy of our method. At high enough resolution, these lines are black; at our medium resolution they will be blended and the line cores will lie somewhat above the zero point.

We used the relevant portion of a high-resolution solar spectrum taken with the Fourier Transform Spectrometer.³ We placed these data on a wavelength scale and convolved them with a Gaussian at the resolution of the SSS at $\lambda 7600$. Finally we rebinned the data onto a linear grid at the SSS dispersion of $0.285 \text{ \AA pix}^{-1}$ (as expected for adequate sampling, this last step had negligible effect). The minimum residual intensity in the atmospheric lines was 0.08 at $\lambda 7602$, while the average residual intensity between $\lambda 7600$ and $\lambda 7620$ was 0.15. We examined the atmospheric bands

in a set of 10 representative SSS solar spectra and found that the line cores matched these values to within 2% of the continuum value of unity. The average residual intensity for the HD 35296 spectrum in the wavelength range under consideration is 0.17. Typically the deepest line cores in solar spectra extracted without using the scattered light removal routine were 3%–4% of the continuum intensity higher, with residual intensity of about 0.11 to 0.12 at $\lambda 7602$ and averaging 0.18 to 0.19 over the wavelength range under consideration. Comparable errors are introduced by oversubtracting the scattered light background (i.e., by simply using the measured interorder minima), and these errors become especially severe, of order 6%, for spectra with closely spaced orders such as those of the FOE.

3.6 Extracting the Spectrum

We use two methods for extracting the data from the debiased, flat-fielded, and descattered frames: *sum* and *optimal*.

Sum extraction is the simpler method: we add the total counts within the extraction aperture to obtain the extracted data value. We account for fractional pixels near the aperture boundaries, so the extraction proceeds in exactly the same way as for extraction of the flat-field spectrum [see Eq. (6)]. This is an unweighted extraction procedure; Eq. (6) may be properly written as

$$S_{mj} = \sum_i W_{mji} \epsilon_{mji} D_{mji}. \quad (12)$$

Here we write out explicitly the weight term W_{mji} . For flat-field spectrum creation and the sum extraction described here we have $W_{mji}=1$ for all i .

The concept of “optimal” extraction as a method of increasing S/N in the extracted spectra was developed in the last decade in a series of papers (Robertson 1986; Horne 1986; Marsh 1989; see also Mukai 1990 for a different approach to the same problem addressed by Marsh). Horne (1986) demonstrates succinctly how to select weighting factors that minimize the variance of the extracted spectrum; the bottom line is

$$S_{mj} = \frac{\sum_i M_{mji} P_{mji} D_{mji} / V_{mji}}{\sum_i M_{mji} P_{mji}^2 / V_{mji}}, \quad (13)$$

where P_{mji} is the order profile image, normalized to unity (i.e., $\sum_i P_{mji}=1$), D_{mji} are the data values of the scattered-light-subtracted object image, V_{mji} is the variance image, determined from the data values and the noise properties of the CCD, and the fractional pixel weights ϵ_{mji} have been included in the calculation of the data values and the model profiles. Knowledge of the variances allows one to create a bad-pixel mask M_{mji} which can be used to iteratively remove cosmic ray hits from the extracted spectrum.

The critical step in this procedure is construction of an accurate profile model, and several methods have been devised. Horne (1986) and Marsh (1989) describe methods of fitting low-order polynomials to straight and curved orders. The spectrum extraction routines developed by F.

³NSO/Kitt Peak data used here were produced by NSF/NOAO.

Valdes for the Image Reduction Analysis Facility (IRAF) use averaging of nearby profiles, while Kinney et al. (1991) fit binned, normalized profiles with high-order polynomials to generate profile models for *IUE* data.

For fiber-coupled spectrographs this problem is especially simple. The object profile is the flat-field profile; in all of our datasets the normalized flat-field profile replicates the object profile to within 1%.

It is important that scattered light be removed from the flat-field frame before this step is performed. The object frame will have been descattered as well (optimal extraction is not possible without the background image available), and it is necessary to ensure that the flat-field profiles match the object profiles as closely as possible.

The results of optimal extraction using the flat-field profile model are shown in Fig. 10. We present results for two orders of the SSS spectrum of the short-period binary UV Leonis. We show the UV Leo spectrum here as its S/N is lower than that of the HD 35296 spectrum, and the improvement produced by optimal extraction is more pronounced in the lower S/N regime. The top pair of spectra are the extracted H α order; the upper spectrum is sum extracted and the lower spectrum is optimally extracted. There is approximately a 30% gain in S/N in the continuum of the optimally extracted spectrum. The extracted spectra from order 40 (approximately $\lambda 5500 \text{ \AA}$) are shown below. Here there is a 40% gain in the S/N of the continuum in the optimally extracted spectrum: the S/N of the sum-extracted spectrum between pixels 50 and 130 is 39, while that of the optimally extracted spectrum is 55. The improvements are comparable with those expected analytically (Horne 1986) and with those obtained experimentally by Valdes (1992) using IRAF V.2.10. Note also that the cosmic ray hits in each order are removed by the bad-pixel mask.

A full comparison of sum and optimal extraction with the above assumptions for fiber-fed data appears in Table 3. In this table, we show the SSS UV Leo results and FOE results for the 10th magnitude binary SS Camelopardalis rather than α Tau for the reason given above. The greatest improvement is seen in the UV Leo spectrum; even weaker exposures obtained with the SSS show up to a factor of 2 improvement. The SSS CCDs have high read noise and the benefit of accounting for it explicitly is evident in Table 3. Conversely, our FOE data benefit less from the optimal extraction. The SS Cam spectrum was one of the weakest ones we had, and we obtained the expected result that optimal and sum extraction converge in the high S/N and low read noise limit. Even on spectra of the brightest stars, however, one can still benefit from using optimal extraction, since there will typically be weak orders containing useful lines (e.g., Ca II H and K) as well as well-exposed orders.

3.7 Wavelength Calibration of the Spectrum

A good way to calibrate the wavelengths in an echelle frame is that employed by IRAF: after marking lines in the various orders of an emission-line spectrum such as that




FIG. 10—The extracted spectrum of the H α orders and order 40 of UV Leo, taken with the SSS. (We show UV Leo here rather than HD 35296 because UV Leo is fainter and the spectrum of lower S/N; optimal extraction is most advantageous in low S/N spectra.) The x axis is scaled in pixels. The upper spectrum in each pair is shown at its actual count level; the lower spectrum has been shifted downward by 200 counts for visibility. The upper spectrum in each pair is sum extracted; the lower is optimally extracted. The S/N between pixels 60 and 120 in the H α spectra are 62 and 73, respectively. The S/N between pixels 40 and 130 in order 40 are 39 and 55, respectively. Note also the elimination of the cosmic ray hits.

from a Thorium–Argon hollow cathode, one performs a two-dimensional calibration of the entire frame.

The premise is as follows. Consider two wavelengths, λ_1 and λ_2 , which lie in different orders; from Eq. (1) we have

$$m_1 \lambda_1 = \sigma (\sin \alpha + \sin \beta_1) \quad (14)$$

and

$$m_2 \lambda_2 = \sigma (\sin \alpha + \sin \beta_2). \quad (15)$$

Now suppose that the wavelengths in question are diffracted in the same direction; then $\beta_1 = \beta_2$ and the above two equations yield

$$\lambda_1 = \frac{m_2}{m_1} \lambda_2. \quad (16)$$

If the spectrum is aligned with the columns on the CCD, so that a given CCD column corresponds to equal angles of reflection across all orders, then we can scale the wavelengths of lines in different orders by the ratio of their order numbers to produce a single “superorder” which can be calibrated as a single entity. The fit coefficients thus determined are then scaled back to the proper value for the different orders. An important advantage of this technique

TABLE 3
Signal-to-Noise in Sums and Optimally Extracted Spectra

	H α (Order 23)		Na D (Order 38)		Other	
	Optimal	Sum	Optimal	Sum	Optimal	Sum
UV Leo (SSS)	73	62	72	56	55	39
AR Lac (QDSS)	138	132	72	67	72	62
SS Cam (FOE)	83	83	75	73	31	31




FIG. 11—The initial calibration of a Thorium–Argon hollow cathode spectrum obtained with the FOE. The large-scale trend in the residuals arises from the tilt of the spectrum relative to the columns of the CCD. The smaller bumps are caused by a number of effects, for example, thermal flexure of the spectrograph, which are removed by triangulation and interpolation of the surface shown in this plot.

is that orders with only a few lines can be accurately calibrated.

Examination of the spectra in Fig. 2 shows qualitatively that the condition described above is not true for any of them: the orders are not “lined up” with the columns of the CCDs. (This is verified quantitatively by the presence of nonzero tilt coefficients in the order map fit.) This necessitates a two-part procedure to produce the full wavelength calibration.

First, we calibrate the superorder as described above. The effect of the overall tilt of the spectrum is the large-scale trend in the residuals in Fig. 11 from positive to negative.

In the second part we fit this large-scale trend as well as the smaller-scale fluctuations which are produced by a number of effects including thermal flexure of the dewar and/or the spectrograph and potential nonuniformities in the CCD surface. Rather than accounting for these effects analytically, we instead perform a second fit which proceeds in three steps. We first convert the grid of line coordinates, which is irregular and nonrectangular, into a rectangular one by fitting polynomials to the residuals order-by-order and interpolating to a uniform set of N points along the dispersion direction. The resulting grid has dimensions $N \times M$, where M is the number of orders, but is still irregular if the orders are tilted or curved. We use a Delaunay triangulation and a quintic polynomial interpolation, as implemented in the IDL built-in routines `TRIANGULATE` and `TRIGRID`, to produce a smooth fit to this irregular grid of residuals. The resulting fit for the FOE data is shown in Fig. 12. After removing a small number of the most deviant points, we obtain a final fit with an average residual across the frame of -0.001 \AA and a standard deviation of 0.040 \AA , which translates to -0.002 ± 0.073 resolution elements.

3.8 Normalization of the Spectrum by the Continuum

A final step for datasets in which absolute flux measurements are not needed is continuum normalization of the




FIG. 12—After removal of the remaining residuals, an excellent fit across the frame is obtained. After applying a “clean-up” command, which eliminates points deviant by more than a user-specified threshold, we obtain a fit with residuals of $-0.001 \pm 0.040 \text{ \AA}$ with 913 lines marked in 33 different orders.

spectra. In general some method such as cubic spline interpolation is required to remove the full continuum shape in optical data. Interpolation of a cubic spline to carefully selected points along the order is adequate for all but the most heavily blanketed orders of cool stars; in general we can define the continuum accurately for all orders containing lines of interest such as $H\alpha$, $H\beta$, and $He D_3$.

For more heavily blanketed orders, it is difficult to reliably identify a continuum, especially since the flat field was normalized and the spectra retained the strong blaze function. A “trick” we can use to address this problem is to normalize a hot-star spectrum (typically observed as a water-vapor or flux standard), output the *fit* to the continuum, and then divide the cool-star spectrum by this fit. Then a second pass is required to remove the scale and a color term, but this is easy since a linear fit (to each order) is sufficient. This technique has another advantage even for spectra or orders in which the continuum can be easily defined; it can be used for a more batch-oriented processing, since the detailed interaction need be done only once, then the set of spectra can all be divided by the continuum fit and renormalized by a linear function. We have provided the option in continuum normalization of outputting the fit, rather than the ratio. (An additional possible use of the fit is for a relative flux calibration of spectra, but this has not yet been implemented.)

The end result of our labors is shown in Fig. 13. The upper panel is the extracted, calibrated, normalized spectrum of HD 35296 obtained with the SSS. The entire echelle spectrum is shown, with lines of interest expanded in the insets. As discussed previously, the SSS also is configured to obtain simultaneous observations of the Ca II H and K through H_8 region ($\lambda 3860 \text{ \AA}$ to about $\lambda 4000 \text{ \AA}$); this spectrum is shown in the lower panel.

4. REDUCE 2.0: A PACKAGE FOR REDUCTION OF FIBER-FED ECHELLE DATA

To address the problems specific to reduction of fiber-fed data, particularly the accurate removal of background




FIG. 13—The extracted, calibrated, normalized spectrum of HD 35296 obtained with the SSS. The upper panel shows the entire echelle spectrum with lines of interest displayed in the insets, while the lower panel shows the Ca II H and K spectrum.

scattered light and the optimal extraction using the flat-field profile explicitly, we have developed a package that implements the reduction steps described above. This package is written in the Interactive Data Language (IDL)⁴ and uses the IDL “widgets” to provide interfaces between the user and the algorithms. It was developed primarily to address scattered light removal in reduction of our SSS data, but we also wish to provide the package to

⁴IDL is distributed by Research Systems, Inc., of Boulder Colorado.

the IDL community. Version 1.0 of this package was presented at the 1993 AAS meeting in Phoenix, Arizona (Hall and Neff 1993), and the many comments we received at the meeting were of great use in developing the present version of the package.

We have set up the package in a public FTP site at Lowell Observatory. Potential users are invited to contact J.C.H. [jch@olin.lowell.edu; (602)-774-3358] for information on how to obtain this package via anonymous FTP.

5. CONCLUSION

We have addressed several problems pertaining to the reduction of fiber-fed echelle spectrograph data and have implemented solutions to them in the IDL language. We were chiefly concerned with the accurate removal of background scattered light from the echelle frame and in optimal extraction of the data to improve the signal-to-noise in spectra of fainter objects or weakly exposed orders. Our implementation of these algorithms, an IDL package called REDUCE 2.0, is available through anonymous FTP to interested users.

J.C.H. acknowledges support from NSF Grant No. ATM-9114733 to Lowell Observatory and from Modification No. 2-93 to Subaward UCAR S9240 of NSF Grant No. ATM-8709659 to the High Altitude Observatory. E.E.F. acknowledges REU support available through NSF Grant No. ATM-9114733.

REFERENCES

- Cardelli, J. A., Ebbets, D. C., and Savage, B. D. 1990, *ApJ*, 365, 789
 Cardelli, J. A., Ebbets, D. C., and Savage, B. D. 1993, *ApJ*, 413, 401
 Gehren, T., and Ponz, D. 1986, *A&A*, 168, 386
 Hall, J. C., and Neff, J. E. 1993, *BAAS*, 24, 1139
 Horne, K. 1986, *PASP*, 98, 609
 Kinney, A. L., Bohlin, R. C., and Neill, J. D. 1991, *PASP*, 103, 694
 Marsh, T. R. 1989, *PASP*, 101, 1032
 Mukai, K. 1990, *PASP*, 102, 183
 Ramsey, L. W., and Huenemoerder, D. P. 1987, *SPIE*, 621, 282
 Robertson, J. G. 1986, *PASP*, 98, 1220
 Schroeder, D. J. 1987, *Astronomical Optics* (San Diego, Academic)
 Valdes, F. 1992, in *Astronomical Data Analysis Software and Systems I*, ed. D. Worrall, C. Biemesderfer, and J. Barnes (San Francisco, ASP), ASP Conf. Ser. 25, p. 398