

A New Scheme for Chiron Fiber-Mode Data Reductions

Frederick M. Walter

Stony Brook University

20 October 2015

Chiron (Tokovinin et al. 2012) is a stable high-throughput bench-mounted echelle spectrograph fed by the 1.5m telescope on Cerro Tololo, operated by the SMARTS partnership. Three years of experience with *Chiron*, observing types of objects that *Chiron* was never designed to observe, have revealed some deficiencies in the standard data processing. I discuss here a new data reduction scheme for data obtained in fiber mode. This provides different treatments of the local background and cosmic rays rejection, permits Gaussian extractions in addition to boxcar extractions, extracts 13 additional orders (down to about 4085Å), and provides more order overlap and flatter orders. This may enable new and different kinds of science from the *Chiron* spectrograph.

Table of Contents

1. Background	1
2. Calibration Data	4
3. Data Reduction	8
4. Data Organization	10
5. Cookbook	13
6. Examples	20
7. Caveats	30
8. The Software	30
9. Coming Soon ...	30
10. Summary	31
11. Acknowledgments	31
12. References	31

1 Background

The reduced Chiron spectra available from the Yale web site are in a 2-dimensional fits file which stores the wavelength solution and the extracted data¹. The reduction scheme is described by Tokovinin et al. (2013). The wavelength solution is obtained by running the IDL procedure *thid* (Valenti 1999), which does a global 2-dimensional 6th order fit (including cross-terms) to the echelle

¹These fits files have the prefix **achi**, and will be referred to herein as the **achi** spectra or **achi** data.

format data. The data are extracted using a boxcar extraction scheme. A cosmic-ray correction is applied.

The reduction scheme is optimized for the type of data for which the spectrograph was envisioned to be used: short exposures of bright stars, primarily for radial-velocity work. For this, the data reductions appear to work well. The radial velocities are certainly very stable².

The extracted spectra (orders 64-125) are divided by a flattened flat spectra, preserving the gross echelle response. Orders 126-138 (roughly 4100-4600Å) are not extracted. Furthermore, over 225 points well off the blaze are trimmed from orders. Exclusion of lower S/N data and retention of the blaze function may make sense for high precision radial velocity data analyses, but I am interested in something completely different.

I have been using Chiron in a way that was probably never envisioned by its creators: synoptic observations of galactic novae. The challenges with the novae are many-fold. While bright at discovery (typically 10th mag, but up to 3rd mag), they fade: we follow them down to V~18 in some cases (at this point essentially all the flux is in emission lines). This is below the nominal sensitivity of the acquisition camera, and so presents some challenges for the telescope operator. Other challenges arise from the fact that we use 10-30 minute exposures, so there is significant background and numerous cosmic rays. Finally, nova spectra are dominated by broad emission lines, with FWZI up to 10,000 km/s. Background determination is non-trivial when the emission line is broader than the echelle order!

In examining the extracted spectra (the **achi*.fits** files) I noticed the presence of narrow “absorption features” (figure 1) near the tops of some of the emission lines in some of the novae. Originally I thought that this might be a new phenomenon attributable to our unprecedented high-resolution synoptic coverage, but some of the physics failed to make sense. My conclusion was that at least some of these “features” might be due to incorrectly subtracted cosmic rays³, but further examination suggests a more insidious problem: in some cases the bright emission lines themselves were being treated as cosmic rays, and being removed⁴.

At that point I decided it was time to look into re-extracting the spectra. The code described here, written in IDL, does just that. It allows the user to re-extract the data, using both boxcar and (optionally) Gaussian extraction techniques. It extracts all 1028 spectral points along the order, and 75 orders, from m=64 through m=138. The broader wavelength coverage requires a new wavelength calibration. A two-step cosmic ray masking is employed, and inter-order background is subtracted.

The inputs for all the code described below are the **chi*.fits** raw data files.

²I measured the stability of the interstellar sodium D lines in σ Ori in slicer mode over the course of a year. After applying a heliocentric correction, the line velocities are stable to better than 250 m/s. The limiting accuracy is likely due to the variable telluric line contamination.

³First suggested by U. Munari in the case of an apparent narrow reversal in H α on one occasion in N Oph 2015.

⁴Sometimes being slow to publish has its benefits

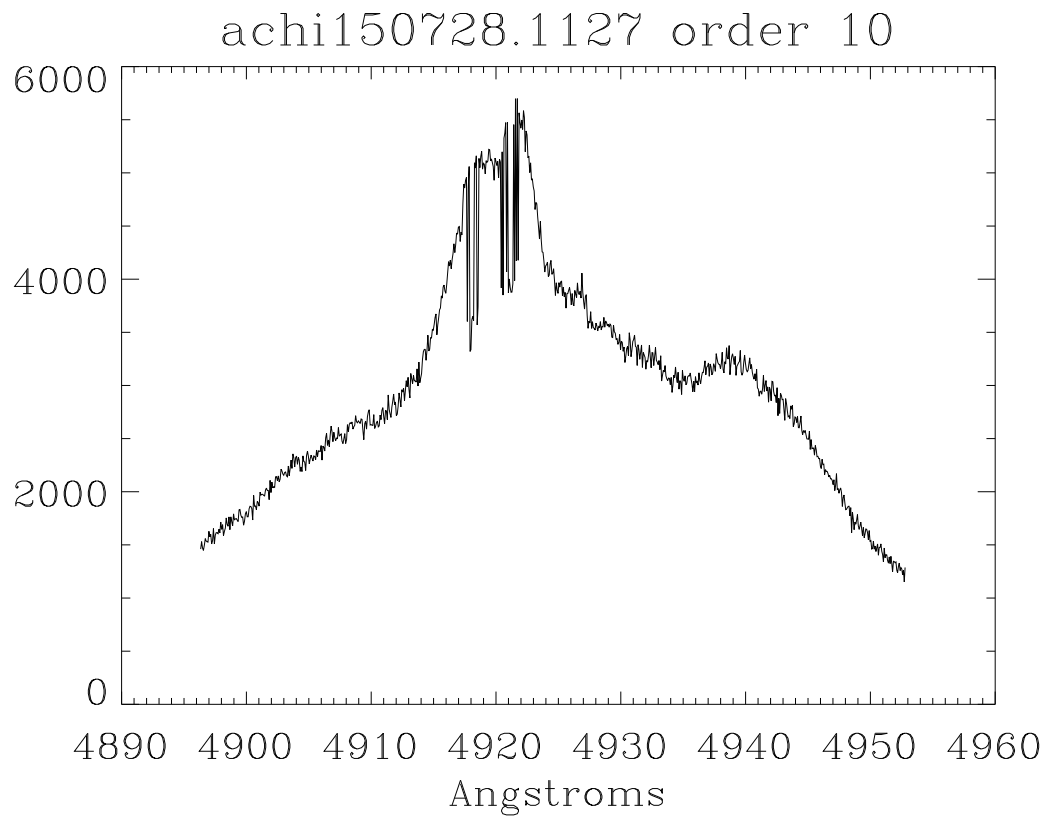


Figure 1: Order 10 of **achi150728.1127.fits**. The target is the nova N Sgr 2015b. The emission line is Fe II λ 4921. The narrow dropouts are an artifact of the standard data reduction process.

2 Calibration Data

There are 3 essential components to the calibration process:

- order trace extraction
- flat extraction, and
- wavelength calibration.

In preparation for the first two items, I co-add all the flat images taken in the appropriate mode and save that as **quartz.fits**. I trim the data and subtract the overscan, using the regions defined in the fits header, prior to co-adding the images.

2.1 Trace Extraction

I use the coadded flat image **quartz.fits** for tracing the orders. Orders are aligned roughly vertically in the images. I extract a subimage from $y=500$ to 550 (the middle of the chip is at $y=513.5$), and collapse this into a vector using the *median* function. I identify the 75 highest peaks. I do test for missing orders and discrepant peaks, and identify the central positions of the 75 orders. I then trace these orders by centroiding the sum of 21 points centered every 10^{th} point, and then fit this with a third order polynomial. The trace is stored in a 1028×75 array, and is saved into a file named **trace.sav** in the calibration data directory.

Because of the stability of the instrument, it is not necessary to extract a new trace each night. I find that using a master trace file, shifted as necessary in the cross-dispersion direction, suffices.

2.2 Flat Extraction

Flats are extracted two ways: boxcar extraction and Gaussian extraction. These techniques are outlined below.

2.2.1 Boxcar Extraction

In this scheme, the data within $\pm wid$ pixels of the location of the trace at that Y position are summed. The width of the extraction slit is ± 3 pixels (set empirically). Background is extracted on either side of the flat with a width set by the distance to the adjacent orders, less twice the extraction slit width. Following extraction, the two background spectra are filtered to remove discrepant narrow features, and averaged.

Because I keep track of the counts in the extracted spectrum and the background regions, I keep track of the \sqrt{N} statistics and propagate the errors. The errors are stored as a S/N vector (= S/error), because this is invariant when the fluxes are scaled.

2.2.2 Gaussian Extraction

In principle, extraction using the PSF can be optimal because cosmic rays and other defects will deviate from the “known” profile. In practice, here I fit cuts across the spectrum with a Gaussian profile, which is probably not the appropriate functional form. Because the closeness of the orders can make it hard to constrain the background, I simultaneously fit three adjacent orders, except for the first and last orders. where I only fit two. There are 12 free parameters (3 in the background and three in each Gaussian). The fits to the outlying orders are discarded. The background is extrapolated under the central Gaussian. The relevant parameter is the analytic integral of the Gaussian (which is presumably less affected by outlying points than is a straight summation of the data).

The fitting is done using the *mpfitfun* (Markwardt) function with constrained fits. The initial fit is made to the sum of the 5 pixels at the center of the image (Y=514). Subsequent fits going outward towards the top and bottom of the image use the previous fit as initial estimates. The position of the Gaussian is constrained to deviate from the previous fit by less than 1 pixel, and the width is constrained to deviate by less than 20%. Note that because the initial estimates depend on the previous fit, the solution can run away and settle on another order. This is not a problem for the flats or for bright sources, but can be for faint objects. Therefore I have an optional iterative fitting process, where the position and width of the initial extraction are forced to those of the trace extraction. In the second iteration these constraints are relaxed (see §3.2), with the position free to move by ± 0.5 pix, and the width free to vary by $\pm 10\%$ from the values from the fit to the trace. This reduces the likelihood that the solution will run away.

2.2.3 Comparison of Extractions

The boxcar- and Gaussian-extracted flats agree with each other (to within a normalization factor) to better than 1%, although there is significant structure (see figure 2). The structure correlates with the width of the Gaussian fit across the order.

The boxcar extraction is flatter, relative to a low order polynomial, by a factor of 2, than is the Gaussian extraction, so at first look it seems that the boxcar extraction is to be preferred. However, because of the stability of the instrument, the structure in the Gaussian flat largely cancels out (figure 4), and the two extractions are comparable, even for reasonably high S/N spectra. In principle the Gaussian extraction provides better cosmic-ray rejection, but it also takes significantly longer to run (see §5.8).

2.3 Wavelength Calibration

I was unable to get good wavelength solutions using the *thid* code. Eventually I decided to do a wavelength solution the old-fashioned way. I selected one order and, using the *thid* solution as a starting point (see §5.2), obtained a wavelength solution. One can also use the wavelength solution in the appropriate **achi** file for a starting point. Then it becomes a matter of identifying the correct




Figure 2: The normalized ratio of the boxcar-extracted to the Gaussian-extracted flat for order 108. The green trace is the scaled width of the Gaussian used to extract the flat. They track well, with the larger ratio (less flux in the Gaussian extraction) when the fit is broader. This occurs when the trace is centered between two pixels. I suspect this is attributable to the non-Gaussianity of the PSF. In any event, the effect seems to divide out (see figure 4).

Table 1: **ths Structure Tags**

Tag	Format	Description
cfits	dblarr(9)	coefficients of polynomial fit (up to 8 th order)
sdv	float	standard deviation of polynomial fit: units=pix
np	integer	Number of lines fit
m	integer	order number (138 - 64)
pixfit	ftarr(np)	X pixel of fit position
wavfit	ftarr(np)	air wavelength of fit position
diff	ftarr(np)	wavelength - fit wavelength
wid	ftarr(np)	Gaussian FWHM of line
mlam	dblarr(3)	m λ for pixels 0, 514, and 1027
xcut	intarr(2)	currently unused
mres	double	mean resolution E/ δ E

wavelengths and fitting the data positions of the lines. I fit these with a polynomial. Using the properties of the echelle, that m λ is constant as a function of the Y coordinate between the orders, I made an initial wavelength solution for each order, and then edited each order, removing blends and weak lines. The air wavelengths of the lines come from **thid_data.xdr**.

My master solution, derived for the night of 150728, contains some 1300 lines in 75 orders. The RMS scatter in the solutions for each order are good to better than 1 km/s.

This solution can then be migrated to other dates, doing a cross correlation of the “good” line list with the extracted Th-Ar spectrum to generate a gross shift, and then automatically re-fitting the individual lines. The goodness of the solution is estimated by plotting m λ as a function of X-position for all the lines fit. There is an irreducible order-dependent scatter of up to $\pm 60\text{\AA}$ in m λ after subtracting the best fit third-order polynomial. The 150728 solution maps well into all the Th-Ar spectra obtained (starting with 120324) - a strong testament to the stability of the instrument.

Verification of the wavelength solution is based on looking for significant deviations in m λ from that expected, and consequently in the length of the order in wavelength space. Deviant lengths are generally due to poorly-constrained solutions near the ends of orders with few lines near the order extrema. Refitting with a lower order polynomial usually fixes the problem.

The wavelength solution is saved in a file named *MMYYDD_cfits.sav*, where MMYDD is the date of the observations. The save file contains a structure, *ths*. The structure has one element per order fit. The tags are listed in Table 1. Note that the *pixfit*, *wavfit*, *diff*, and *wid* tags are padded to 60 element arrays.

It is important to realize that, while I need a reliable wavelength solution, I do not need one with the fidelity of the Yale solution. Consequently, I have neither examined nor tweaked the solutions

on most nights for the best possible solutions. But they look pretty good (see figure 3).

3 Data Reduction

3.1 Cosmic Rays

In most cases, we have taken single exposures in order to maximize S/N in a given time. This comes at the expense of contamination by cosmic rays. The propensity of the default reductions to flag and remove strong emission lines is a major motivation for this reduction effort.

There are two cases to consider: cosmic rays on or near the order, and those in the background.

Cosmic rays in the background are the less serious issue. Filtering, and then averaging, the background spectra on either side of the order efficiently removes most small-scale blemishes. In the Gaussian extractions we fit 4 interorder regions, ensuring that small features will have a minimal effect.

However, most cosmic ray detection algorithms look for sharp edges, very much like the dispersed spectrum, and this can cause difficulties. I originally tried a direct application of the *la_cosmic* procedure (van Dokkum 2001), and found that it tended to chew up the spectra. After some experimentation, I decided to pre-filter the data with a 1x5 median filter. This has the beneficial effect of removing sharp features in the dispersion direction, while only minimally reducing resolution for sharp features. Then an application of *la_cosmic* seems to do a better job, but it still has trouble with the bright H α emission line in many novae. So I incorporated a keyword to let one ignore any masked data in the immediate vicinity of H α .

Alternatively, the *dcr* (Pych 2004) algorithm can be used too, by specifying the *dcr* keyword in the call to *ch_reduce*. The user will need to install the *dcr* executable. A customized **dcr.par** file is written in the working directory. The *THRESH*, *XRAD*, *YRAD*, *NPASS*, and *GRAD* values can be set via keywords. I've modified the default values for *XRAD* and *YRAD* to 4 and 31, respectively. When *dcr* works, it gives about the same results as does *la_cosmic* as implemented here, though it seems to miss more cosmic rays, especially in high S/N spectra. It may be possible to fine tune the parameters to do a better job. *dcr* does, however, run much faster than *la_cosmic*. However, it seems to be unstable, in that varying the *XRAD* and *YRAD* can give wildly divergent results, including clearly unphysical continua. *dcr* seems to remove low points as well as high points, which can get it into trouble near the ends of orders.

The data quality vector (see Tables 2 and 5) encodes pixels affected by cosmic rays, as well as saturated pixels.

3.2 Data Extraction

The basic extraction techniques are described in §2.2.

The trace is cross-correlated (in X) against the image in order to find any bodily shift between the spectrum and the flat.

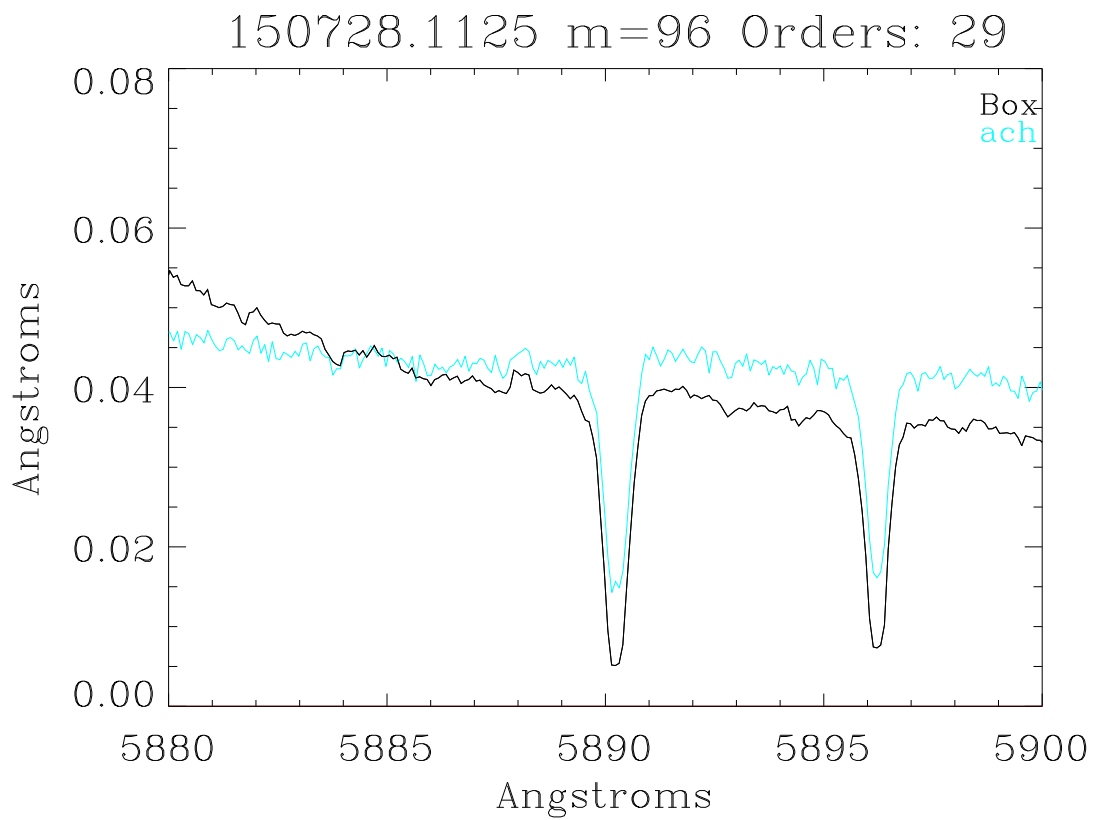


Figure 3: The Na D lines in N Oph 2015 on 150728. There are few narrow lines in novae - these are interstellar. The **achi** spectrum is in aqua. The wavelength solutions are in good agreement.

Table 2: **xchi*.fits** file contents

Index	Description
0	Wavelength solution
1	Boxcar-extracted data
2	Signal-to-noise ratio
3	data quality
–	Gaussian-extracted data
–	Flux-calibrated data (see §5.4)

Boxcar extraction uses the same slit width employed for the flat extraction.

Gaussian extraction can be specified, either with 1 or 2 passes, as described in §2.2.2.

The data are flattened by division by the appropriate flat extraction. Because the blaze function is not removed beforehand, the continua are more-or-less flat.

The extracted data are written out into an **xchi*.fits** file, with a format similar to the **achi*.fits** file from the standard extractions, except that the data are an Nx1028x75 array, rather than a 2x801x62 array. The first two images (N=0,1) are the wavelength solution and the boxcar-extracted, as in the **achi*.fits** files. The other indices are described in Table 2. An absent index means that the image is assigned dynamically, with the index recorded in the fits header.

More data, defined in Table 3 are saved in the **xchi*.sav** file. In Table 3 *nord* refers to the number of extracted spectra, nominally 75.

4 Data Organization

All data and calibration data are stored in a directory referred to as *chirondata*. The cshell definition is

```
setenv chirondata YourDataDirectory
```

Create a subdirectory *cal* under *chirondata*.

Tar files downloaded from the Yale archive should be saved and unpacked in the *chirondata* directory. Each night will be in a subdirectory named YYMMDD, the observation date.

Calibration data should be downloaded to and unpacked in *chirondata/cal*. There will be a series of directories of the form YYMMDD_cals.

Extracted and derived calibration data are saved in the *chirondata/cal* subdirectories. Extracted spectra are saved in the *chirondata* subdirectories.

All procedure file names are lower case; on occasion I have spelled them with mixed cases to improve readability.




Figure 4: The ratio of the boxcar-extracted to the Gaussian-extracted spectra, after flat division, for order 108 of chi131014.1143 (μ Col). The RMS scatter of the ratio is 0.6%. The median is 0.999; no normalization has been applied. The structure visible in the flat ratio (figure 2) is greatly reduced. The SNR of the spectrum, from counting statistics alone, peaks about 900, and exceeds 500 between pixels 250 and 850. Within this region the RMS is about 0.4%, comparable to the 0.3% expected from counting statistics alone.

Table 3: **xchi*.sav file contents**

Variable	format	description
FILELOG	structure	header information
H	strarr(*)	FITS header
MASK	bytarr(1024, 1028)	mask image from <i>la_cosmic</i>
THS	structure(nord)	defined in Table 1
XFITS	fltarr(1028,nord)	traces
<u>Boxcar Extraction Parameters</u>		
BACK	fltarr(1028, nord)	background spectra
BDATA	fltarr(1028, nord)	net extracted spectrum
BDQ	intarr(1028, nord)	data quality flags
BDX	fltarr(1028, nord)	extracted, flat-fielded spectra
SNR	fltarr(1028, nord)	SNR for BDATA, BDX
WID	INT	extraction width =2*wid+1
<u>Gaussian Extraction Parameters</u>		
GDATA	fltarr(1028, nord)	net extracted spectrum
GDQ	intarr(1028, nord)	data quality flags
GDX	fltarr(1028, nord)	extracted, flat-fielded spectra

5 The Cookbook

The code is written in IDL, and was tested under version 8.1. The process of calibrating and reducing a night's data is easily divided into three parts⁵.

5.1 Basic Calibrations

The data must be in the *chirondata/cal/YYYYMDD_cals* directory (see §4). *Ch_Cals* reads the calibration data, and

- co-adds the flat images, creating a file called `quartz.fits`
- traces the orders in the flat images, and writes `trace.sav`
- extracts and writes the flat orders to `flats_YYYYMDD.sav`
- extracts and writes the Th-Ar spectrum to `tharYYYYMDD.sav`

```
Ch_Cals: reduce Chiron calibration images
* calling sequence: Ch_Cals,day
*   DAY: 6 digit string (yymmdd); required
*
* KEYWORDS:
*   EXAMINE:  set to examine each extracted order (sets plt)
*   FORCE:    set to remake the quartz.fits file
*   MKTRACE:  set to make new trace'
*   MODE:    default=fiber
*   NOGAUSS:  set to skip Gaussian extraction of flat
*   NOWAVCAL: if set, skip wavelength solution migration
*   PLT:     set to follow the extractions by plotting the extracted orders
*   USEFLAT:  set to use pre=extracted flat for this day
```

Unless the keyword *nowavcal* is set, the default wavelength solution will be migrated to this day, using the Th-Ar spectrum for this day. See §5.2 for more details.

If multiple Th-Ar observations exist, only the first one is used. There is currently no provision for flexure or drift during the night.

Set the *plt* keyword to see the extracted flat orders. Set the *examine* keyword to pause after each order is plotted. While paused, type “z” to get an IDL prompt; type any other key to continue.

The wrapper *Do_ChCals* runs *Ch_Cals* on multiple directories. You can use *Ch_CalStatus* to list directories that have yet to be calibrated.

⁵ *Omnes data in tres partes divisa est.*

Table 4: **WCal Keystroke Commands**

Key	Description	Notes
e,E	examine fit to line	2
m,M	mark line (2 keystrokes)	1
q,Q	exit	1,2
r	refit data	1,2
R	refit data after changing polynomial order	2
s,S	restore all deleted points	2
x,X	delete point closest to cursor	1,2
z,Z	stop	1
>	jump ahead	1
<	jump back	1
space	mark line center	

Notes: 1: used in *ch_setwcal*

2: used in *plwfit*

5.2 Wavelength Calibration

For many of these routines, you will need the Th-Ar line list in **thid_data.xdr**. This file should be saved as **chirondata/cal/thid_data.xdr**. You will also need the **chirondata/cal/tharYYMMDD.sav** file generated by *ch_cals*. The wavelength solutions are stored in **chirondata/cal/YYMMDD_cals/YYMMDD_cfits.sav**.

These routines use a number of common subroutines. In *plwfit*, commands are entered by single keystrokes as defined in Table 4.

The basic routines used in wavelength calibration are:

ch_setwcal: Initial setup. This runs *thid* on the Th-Ar lamp image to determine the initial wavelength solution. *thid* does not converge as well as I think it should (or as it used to, for the BME), but typically 300-500 lines are fit. The user specifies the order to fit (using the *startord* keyword, and the order of the polynomial fit (*ford*; default=5). The number of Th-Ar lines plotted is set by the *minflux* keyword. By default, all lines stronger than 0.8 of the mean intensity are plotted. Decrease this value to see more lines. Note that the tabulated line intensities do not necessarily correlate with the observed intensities. The Ar lines are also used. No checking is done for saturated lines.

Two plots are generated. One shows the difference in pixels between the observed and predicted line location as a function of wavelength. The other (the active window) shows a portion of the spectrum with lines from the this list marked. Vertical magenta lines have

been fit; dotted yellow line have not. Use the `>` and `<` keystrokes to navigate. Going beyond the long wavelength end of the spectrum or typing `Q` exits this part of the program. Use the `M` key to add lines to the fit. First mark the yellow dotted line, then mark where it should be in the data. The Y level of the cursor sets the baseline for the Gaussian fit. The fit line center will be shown in green. Mark a line for exclusion with the `X` key.

Use the `>` key to get past the end of the spectrum to move on. Then the data will be fit using *plwfit*. Useful keystrokes (see table 1) include *e*, *x*, and *r*. Only the *r* command is case-sensitive.

ch_tweakwcals: use to tweak an existing wavelength solution for one order. Call it as **ch_tweakwcal**,**YYMMDD**,**order=x**.

ch_migratewcals: use to migrate a wavelength solution from one day to another.

```
* Ch_MigrateWcal - migrate one Chiron wavelength solution to a different day
* calling sequence: Ch_MigrateWcal,date
*   DATE: date in format YYMMDD, no default
*       Ch_Cals must have been run for this date
*
* KEYWORDS:
*   AUTOFIT: def=2; 1 to examine each fit
*   EXCLSIG: automatic line exclusion, def=2.5 sigma from initial fit
*   MINFLUX: def=0.8 of mean of Th-Ar intensities
*   ORDERS:  set for select orders
*   REFFIT:  reference fit date (YYMMDD), def="chiron_master_wcal"
*   THRESH:  threshold counts in line for autofit, def=20
```

ch_examinewcals: non-interactive examination of the solution. It plots the deviations of the fit $m\lambda$ after subtracting a third order polynomial. Outliers are flagged, and orders that are too long or too short are identified so they can be refit. Do this by running **ch_examinewcal**,/**redord**, which uses **ch_tweakwcal**,/**redord** on the flagged orders.

ch_extendwcals: use to use the solution for one order to another order on the same night.

ch_wcalsummary: generates a summary of each wavelength solution (one line per order). The last column, the mean resolution, gives a good data quality indicator. It should be about 28,000 for fiber-mode data.

Once you have a master wavelength calibration, you should be able to use **ch_migratewcal**s and **ch_examinewcal**s to calibrate all other days.

5.3 Data Extraction

The calibration files including a wavelength solution, must exist before data can be reduced. The data must be in the *chirondata/YMMMDD* directory (see §4). To reduce a full night's data, merely run *ch_reduce, YMMMDD*. To reduce a particular spectrum, type *ch_reduce, YMMMDD, file=xxxx*, where *xxxx* is the 4 digit file number (**chiYMMMDD.xxxx.fits**). The file number can be either a string or an integer.

```
Ch_Reduce reduce raw Chiron Spectra
* calling sequence: Ch_Reduce,date
*   DATE: 6 digit string (yymmdd); required
*
* KEYWORDS:
*   DOGAUSS:   1 or 2 for Gaussian extractions
*   EXAMINE:   set to examine each extracted order (sets plt)
*   FILE:      name of file; def = all
*   MODE:      default=fiber
*   NOCLEAN:   set to skip extrapolation over bad points
*   NOCR:      set to skip CR filtering
*   NOHAMASK:  set if H-alpha is in emission
*   NOGAUSS:   set to skip Gaussian extraction
*   NOPROC:    set to skip post processing
*   NOSAVE:    set to run code but to not save the data; stops when done
*   NSIG:      sigma cut for replacing bad points in Gauss fit. Not yet implemented.
*   OPTIMUM:   set for optimum extraction. Not yet implemented.
*   PREFILTER: 1xn filter, def=5
*   STARTORD:  starting extraction order, def=0
*   UPDATE:    set to update single order
```

ch_reduce options:

- The boxcar extraction is the default.
- A Gaussian extraction is specified (in addition to the boxcar) by setting the *dogauss* keyword to 1 or 2, for either single pass unconstrained fit or a constrained double pass fit (§2.2.2). Note that this option significantly slows down the execution time.
- If there is a strong H α emission line, the cosmic ray filtering may wipe it out. See figure 5 for an example. This will happen in *ch_reduce* unless the *nohamask* keyword is set. This prevents the code from identifying the top of the H α line as a defect, unless saturated.
- The clean algorithm identifies the pixels flagged as bad and interpolates over them. Use the *noclean* keyword to turn this off. All bad pixels are identified in the *bdq* vector (see Table 5).

Table 5: **Data Quality Flag**

Flag	Description
1	pixel flagged by <i>la_cosmic</i> in lower background slit
2	pixel flagged by <i>la_cosmic</i> in upper background slit
4	extracted data interpolated over single flagged pixel along spectrum
8	extracted data interpolated over single flagged pixel across spectrum
16	pixel flagged by <i>la_cosmic</i> in extraction slit
32	gaussian flat = 0
512	saturated pixel

- set *nocr* to see what happens in the absence of cosmic ray filtering.
- Use the *update* keyword to reprocess a single order. Set *update* to the order number (0-74).

The output from *Ch_Reduce* is the files **xchiYYMMDD.xxxx.fits** and **xchiYYMMDD.xxxx.sav**, as described in §3.2.

5.4 Post Processing

Ch_reduce writes out the extracted, flat-fielded counts spectrum. Post-processing consists of a.) flux-calibration, and b). creating a spliced one-dimensional spectrum.

Flux Calibration: I calibrate the response using a spectrum of μ Col obtained on 131014. In each order the counts are ratioed to the modeled spectrum⁶. The modeled spectrum is low resolution: I smooth over Balmer lines. The conversion (F_λ/count) is saved in **fluxcal.fiber.sav** in the *chdata/cal* directory.

There are issues with the calibration near the telluric A and B band heads.

The procedure *ch_calflux* is used to convert the counts spectrum to a flux spectrum. To save this to disk, use *ch_xupdate,flux=1*. The **xchi*.sav** and **xchi*.fits** files are updated. The array in the fits file is expanded. *Rd_Chiron* extracts this as d.f in the d structure.

Note that spectroscopic flux calibration is predicated on the assumptions that both the calibrator and the target are observed under photometric conditions, and with equal slit losses. This is often not the case.

Order Splicing: Splicing of the individual orders is accomplished by the *ch_splice* function, $z=\text{ch_splice}(xfile)$, where *xfile* is the **xchi*.fits** file name or number. *z* is an expanded structure described in

⁶available at <https://www.eso.org/sci/observing/tools/standards/spectra/hr1996.html>

Table 6: **Ch_Splice Output Structure**

Tag	Type	Size	Description
W	DOUBLE	Array[1028, 75]	xchi 2-D wavelength vector
S	DOUBLE	Array[1028, 75]	xchi 2-D counts vector
E	DOUBLE	Array[1028, 75]	xchi 2-D signal-to-noise vector
F	DOUBLE	Array[1028, 75]	xchi 2-D flux vector
H	STRING	Array[140]	fits header
W1	DOUBLE	Array[82887]	spliced 1-D wavelength vector
F1	DOUBLE	Array[82887]	spliced 1-D flux vector
E1	DOUBLE	Array[82887]	spliced 1-D signal-to-noise vector
DQ	DOUBLE	Array[82887]	spliced 1-D data quality vector
SPLICE	DOUBLE	Array[138]	splice points
SKY	FLOAT	0.00000	sky spectrum (not currently used)

Table 6. The output is a save file named **object_YYMMDD.xxxx.sav**, where *object* is from the fits header (with all spaces removed).

The 1-dimensional wavelength vector is linear, with a bin size equal to the smallest in z.w, so that no resolution is lost. The flux and signal-to-noise vectors are interpolated to this linear wavelength scale. The method of splicing is either to do a straight sum (keyword *sum*), to cut by the signal-to-noise vector (keyword *cut*, or to weight by the signal-to-noise vector (default).

When the data are weighed by the errors (signal/signal-to-noise), all negative points are first excluded as unphysical.

Because the splice regions are at the ends of the orders, they tend to have low S/N. The file **chdata/cal/ch_deftrims.dat** lists the empirically-determined “best” cuts, based on visual inspection of the μ Col orders. You can change these with the *trim* keyword, which is added to the defaults (subtracted from the RHS) . *trim* can be a scalar, a 2-element vector (left, right), or a full (2x75) vector.

Setting the *scale* keyword will perform a scaling of the two spectra in the overlap region. This should not be necessary, since the flux calibration serves the same purpose. Furthermore, in regions where the flux calibration may be suspect (e.g., the edges of strong lines, or the terrestrial A and B bandheads), this may introduce an unwanted shift in the flux levels.

By retaining all 1028 points in the orders, interorder gaps are eliminated through order 71 (8506Å) or through order 70 (8382Å) using the default order trims). The user can override the number of points retained for splicing by specifying the *trim* or by editing **ch_deftrims.dat**.

The flux calibration and the order splicing, as well as checking and updating the wavelength solution, are handled by the *ch_postproc* procedure.

5.5 Pairs of Spectra

Cosmic ray rejection is simpler, and more robust, when you have two or more spectra of an object obtained on the same night under the same conditions. The routine *ch_reduce2* reduces pairs of spectra. Cosmic ray rejection is handled by searching for excursions in the difference image that exceed a multiple (set by the keyword *nsig*) of the expected variance. After that, *ch_reduce2* runs much as *ch_reduce*, with the exception that the output file names are numbered 9xxx rather than xxxx, where xxxx is the number of the first of the pair.

See Figure 13 for an example of a coadded pair of deep exposures.

There are problems co-adding high S/N pairs with different exposure times, or which differ in instrumental throughput due to variable cloudiness.

5.6 Ancillary Software

In almost all cases, typing *procedure,/help* will bring up brief on-line help.

ch_calstatus: check for the existence of calibration data.

ch_sat: identify saturated pixels in **chi*.fits** image.

ch_status: check for unreduced data.

ch_xy: convert between pixel and order/wavelength coordinates.

comp_chx: compare various extraction methods.

plch: quick plot of 1-dimensional *Chiron* spectra

rd_chiron: read fits file, and convert to an IDL structure. Works for both **achi*.fits** and **xchi*.fits** files. Use the *xfiles* keyword to read a **xchi*.fits** file if both types exist.

5.7 Master Calibration Files

The following files are in **chirondata/cal**:

chiron_master flats.sav A flat spectrum. It is advisable not to use this, but rather to extract the flat on each individual night and use the **flats_YYMMDD.sav** file.

fluxcal_fiber.sav Flux calibration for the fiber mode spectra.

chiron_master_trace.sav Traces. This is used by default unless *mktrace* is specified in *ch_cals*.

chiron_master_wcal.sav The master wavelength calibration file which is migrated to and tweaked for other nights.

ch_deftrims.dat Default number of pixels to be trimmed from the orders prior to splicing.

The procedure *ch_cals* writes the following files into the **YYMMDD_wcal** directories:

YYMMDD_cfits.sav Wavelength solution.

flats_YYMMDD.sav Boxcar and Gaussian extractions.

quartz.fits Co-added fits image of all fiber-mode flat images.

tharYYMMDD.sav Boxcar extraction of the first fiber-mode Th-Ar file found.

trace.sav Optional, if the *mktrace* keyword is specified..

5.8 Run Times

All run times refer to a desktop running CentOS 6.7. The hardware is an Intel Core Duo CPU clocking at 3 Ghz, with 3.8 Gb of memory. The IDL version is 8.1.

Full calibration of a single night requires about 19 minutes, unless the Gaussian extractions are turned off.

Boxcar extraction and post-processing of an image requires less than one minute.

Extraction and calibration of a spectrum, including the Gaussian extraction, requires about 20-25 minutes for high S/N spectra, and up to 50 minutes for noisier spectra. The (lack of) speed is mostly attributable to my using *mpfitfun* to fit 12 free parameters to 1028 slices of the spectrum in each of 75 orders. This can probably be sped up.

6 Examples

Here I present a few examples comparing my reductions with the Yale reductions. In all plots my reductions are in black. Red ticks at the bottom of the plot indicate pixels flagged as bad.

Figures 5 through 9 are 1800 second integrations taken of N Sco 2015 on 28 or 29 July 2015. The *V* mag was 15.0 at this time. Figure 5 shows the effect of the cosmic ray detection algorithm on the peak of the H α emission line. Figures 6 and 7 show the region of the sodium D lines, and illustrate the consequences of subtracting the instrumental background.

Figures 10 through 12 are of N Sgr 2015b. At this time the novae was near the minimum of its dust dip, but was still reasonably bright, at *V*=12.5. This is an 1800 second integration.

Figure 13 shows the quality of spectrum one can get on a faint object (*V*745 Sco at *R* > 16). It also illustrates the quality of the order splicing.

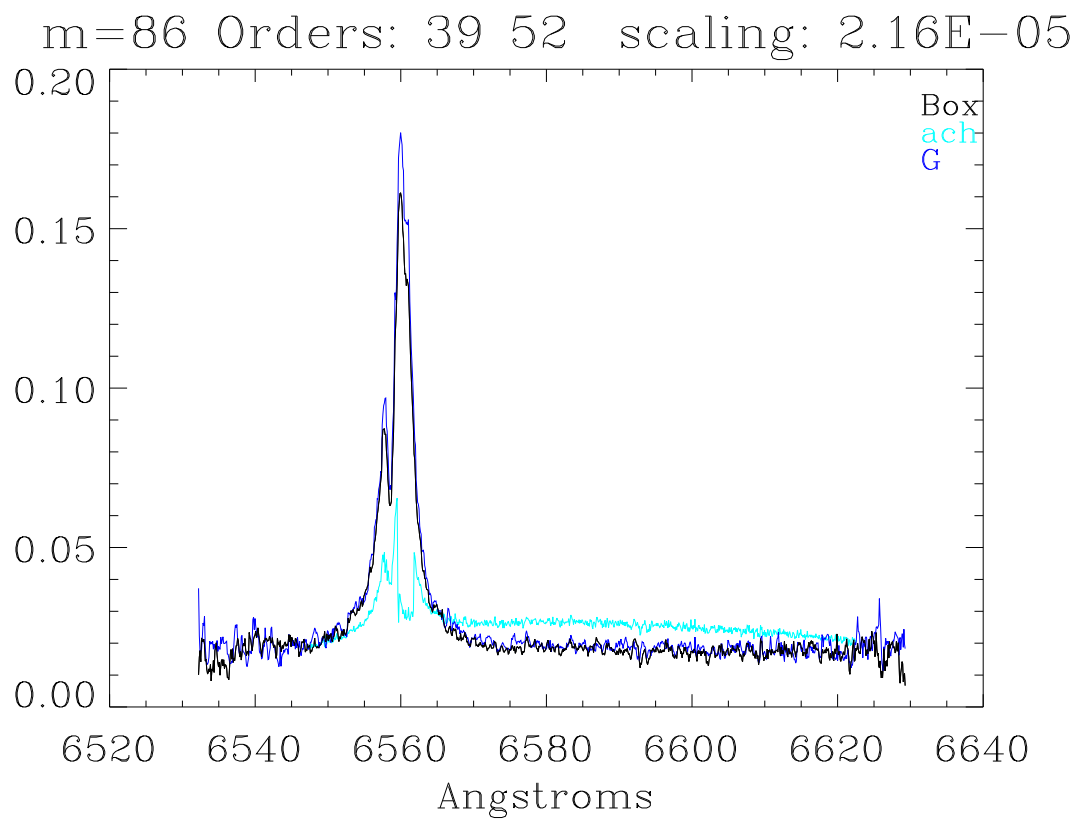


Figure 5: A comparison of the boxcar extraction (black), Gaussian extraction (dark blue) and the Yale extraction (aqua) for the $H\alpha$ of Nova Sco 2015 (image **chi150728.1124**). The Yale extraction is scaled so the medians match. Note that the central peak of the $H\alpha$ line is missing in the Yale extractions. It has apparently been flagged as a cosmic ray or some other type of defect. The top of the line is not saturated.

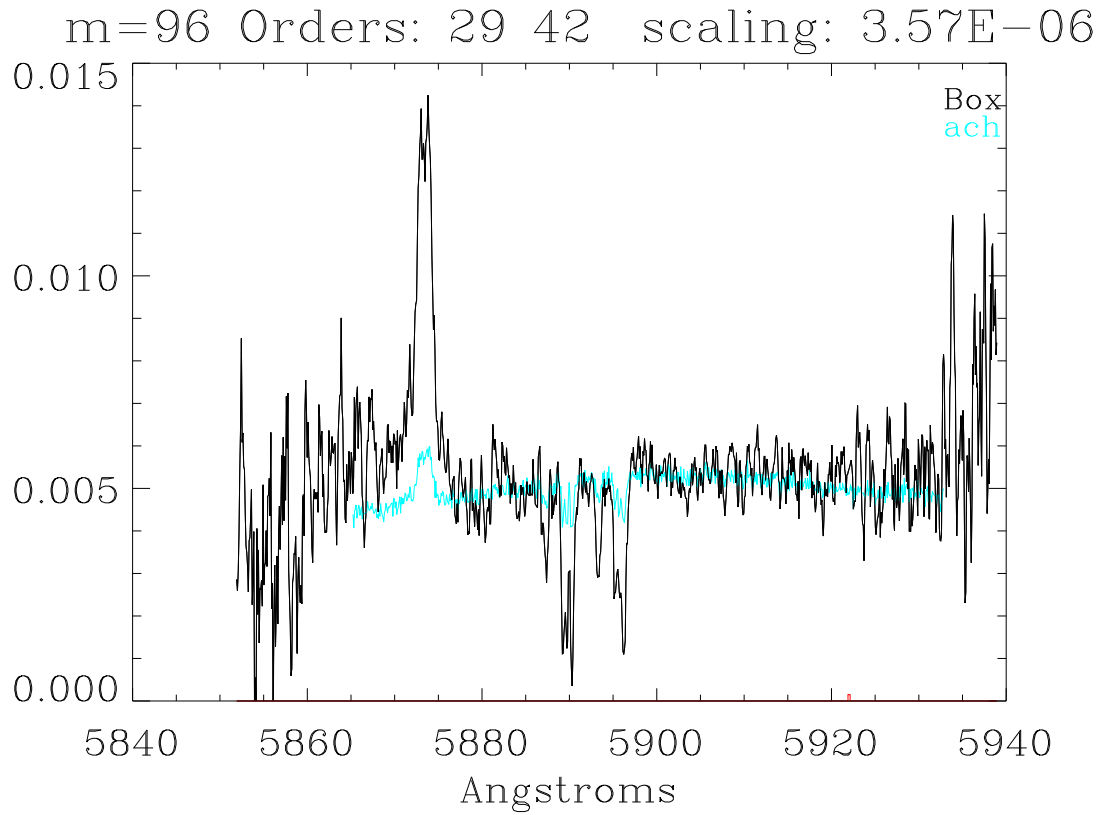


Figure 6: A comparison of the boxcar extraction (black) with the Yale extraction (aqua) for the sodium D line region of Nova Sco 2015 (image **chi150729.1124**). The Yale extraction is scaled so the medians match. Note that the line strengths are much weaker in the Yale reduction, while the S/N is higher. This suggests that the background has not been subtracted in the Yale reductions (see figure 7). The emission line is He I λ 5876. Multiple velocity components are visible in the Na D₁ and D₂ lines. The low velocity components are galactic foreground; the blue-shifted absorption lines may be ejecta from the nova.




Figure 7: This is the same data as shown in figure 6, except that the data (black) are not flattened, and the background has not been subtracted. This provides a reasonably good match to the Yale reductions (magenta).

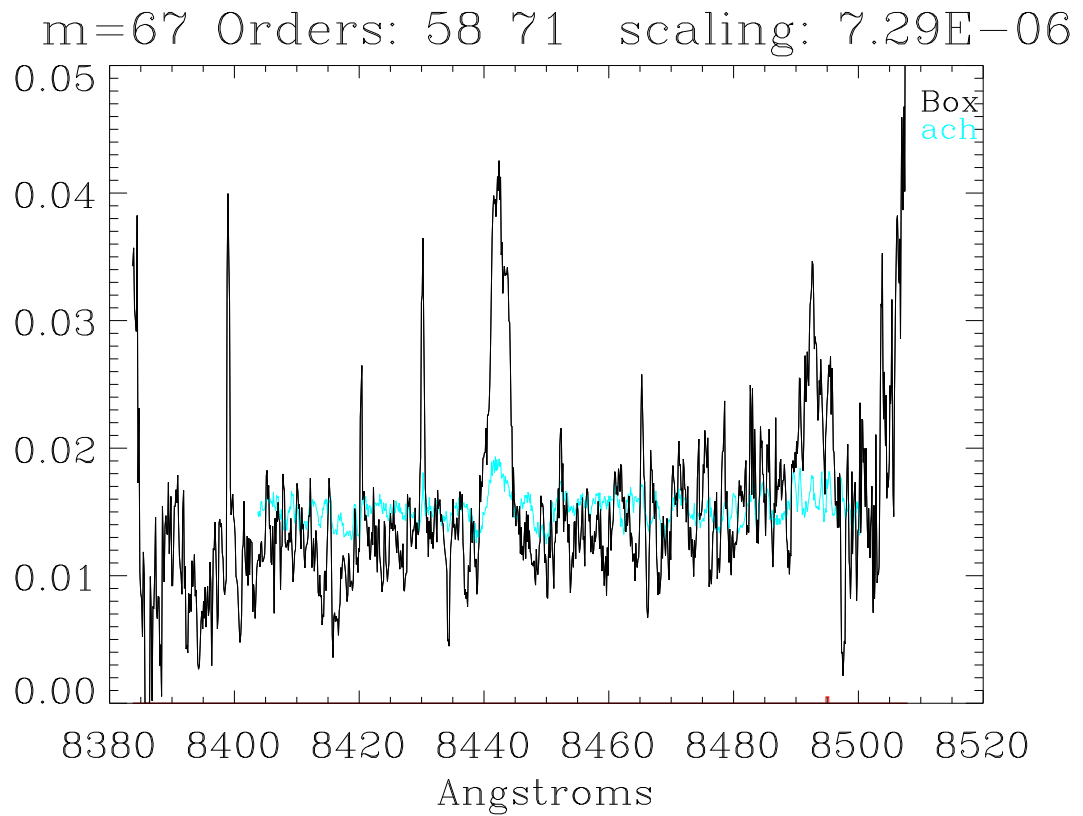


Figure 8: Order 67, showing strong emission from the $\lambda 8498$ component of the Ca II IR triplet. The line is not evident in the Yale reductions. The apparent strengths of the $\lambda 8446\text{\AA}$ O I line are also very different.




Figure 9: Order 66, showing strong emission from the $\lambda 8542$ component of the Ca II IR triplet. As in figure 8, the line is not evident in the Yale reductions.

m=90 Orders: 35 48 scaling: 2.23E-05

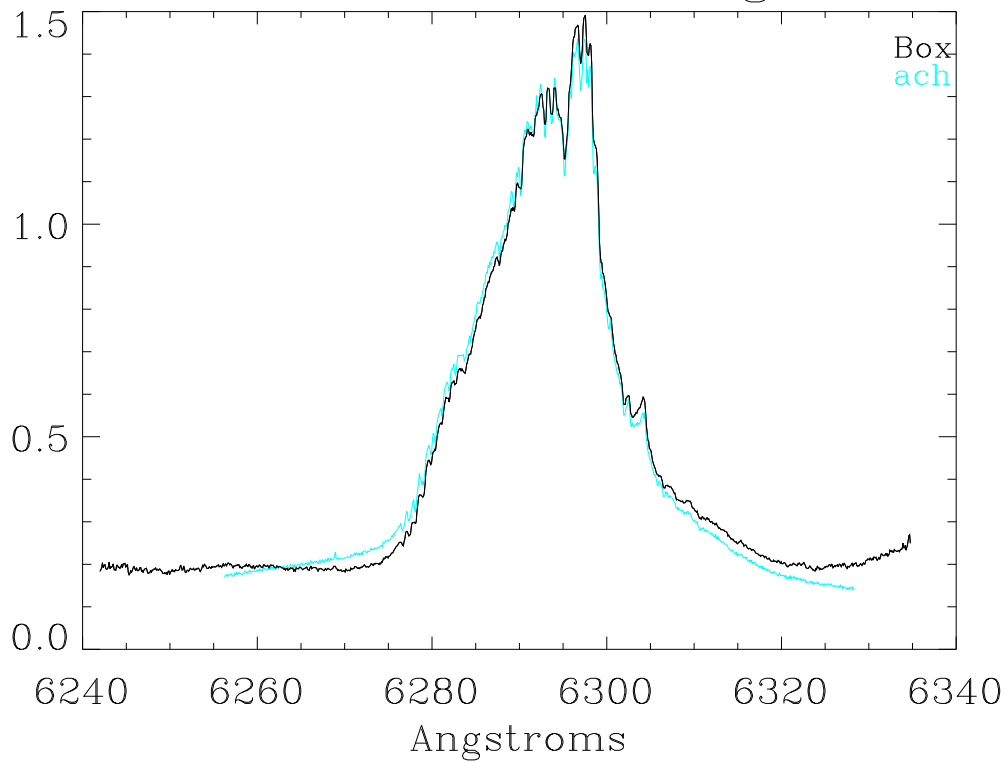


Figure 10: The bright [O I] $\lambda 6300$ line in N Sgr 2015b. The two extractions agree well when the background is negligible. Differences in the slopes are attributable to the flat-fielding.

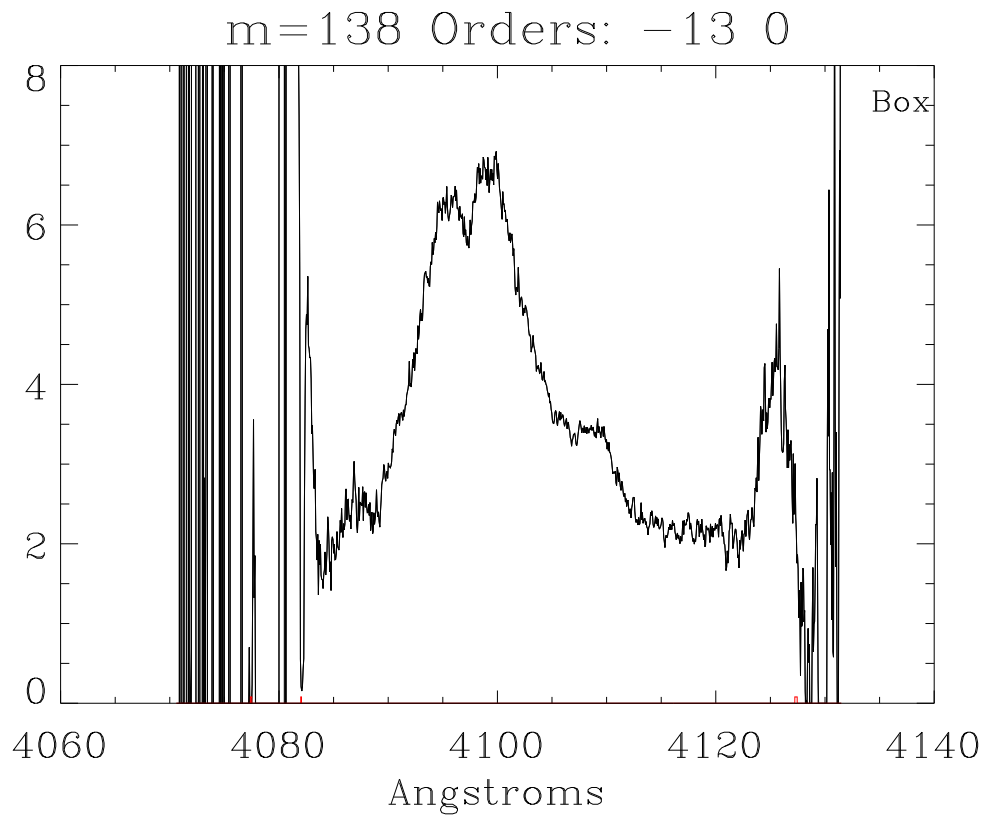


Figure 11: The $H\delta$ $\lambda 4100$ line in N Sgr 2015b. The other line is probably Fe II 4128 (multiplet 27). This is as blue as Chiron goes. All 1028 pixels are plotted; nothing has been masked. There is not much sensitivity at the ends of the order.

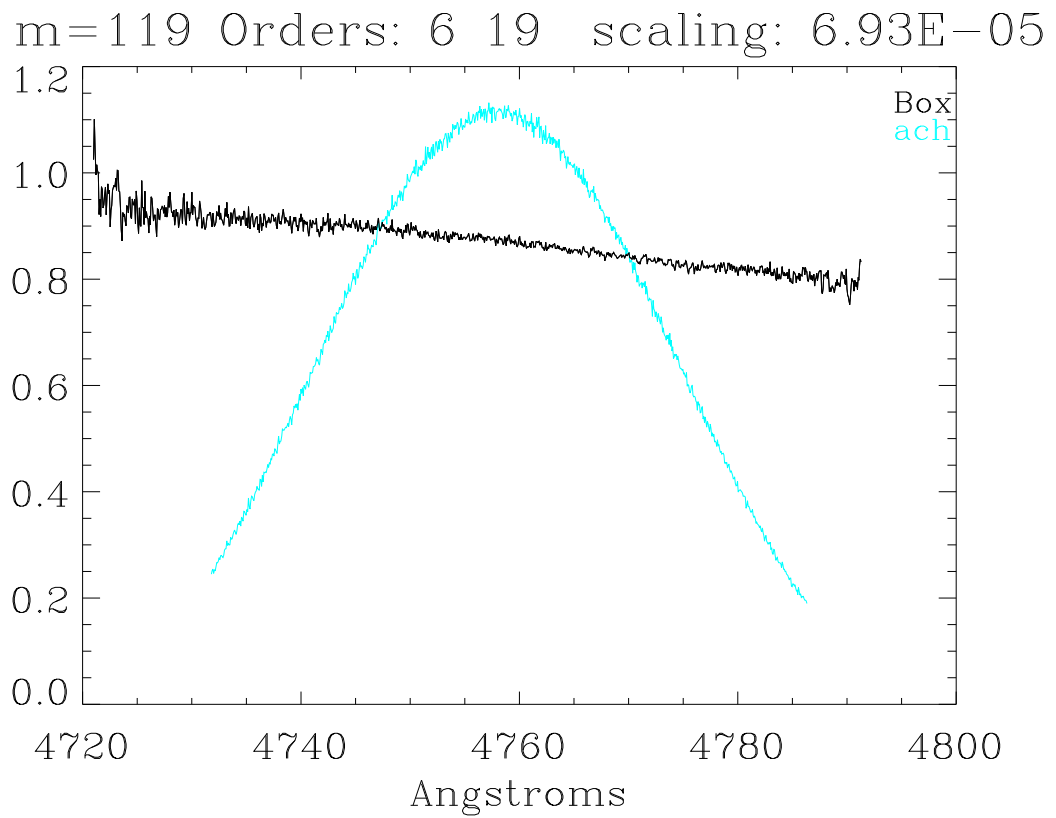


Figure 12: A particularly boring part of the spectrum of the O9.5V star μ Col, a spectrophotometric standard. Note that division by the actual flat extraction yields a very flat spectrum.

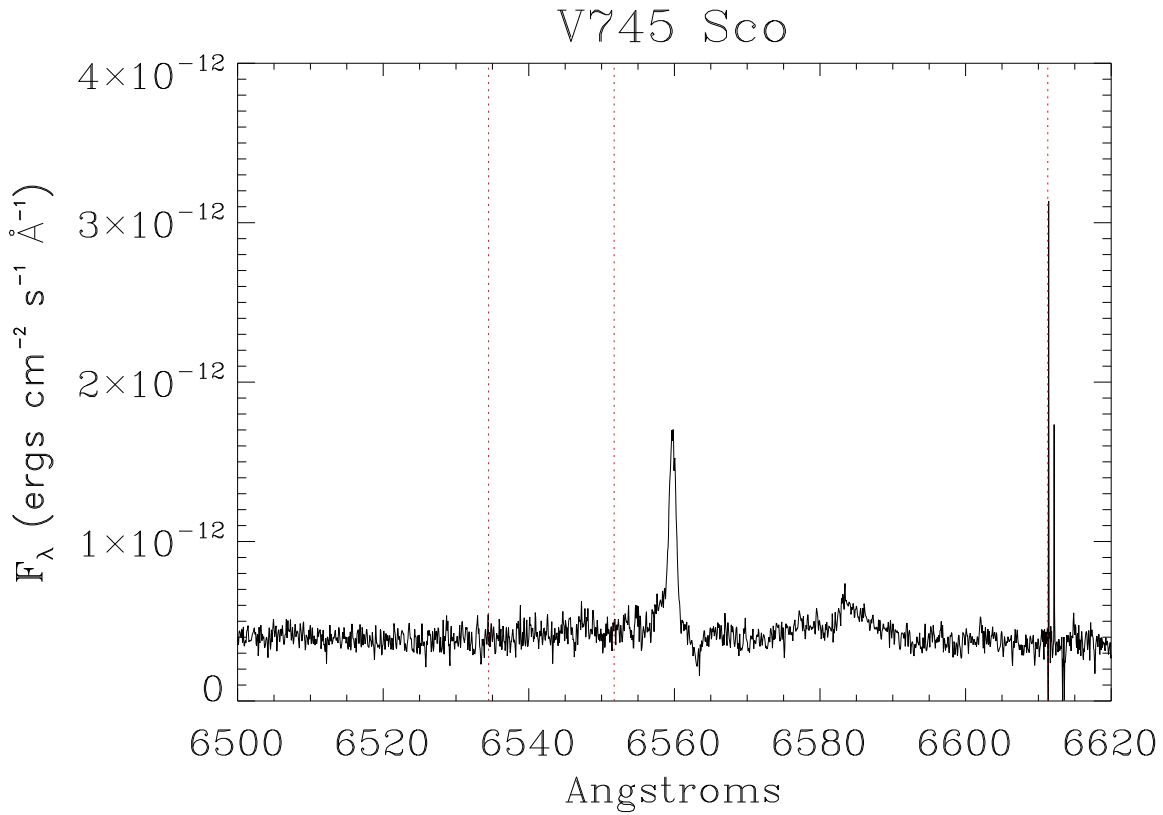


Figure 13: The H α region of V745 Sco on 150603. This is the sum of two 1800 second integrations, reduced using *ch_reduce2*. The *R* mag at the time of observation was about 16.2. The dotted red vertical lines mark the spliced regions. No background has been subtracted; the data are unsmoothed. The default trims have been expanded by 10 pixels. The high points are due to low counts at the edge of the detector rather than to cosmic rays. The narrow emission is H α ; the source of the broad emission near 6585 \AA is not currently known. No sky has been subtracted; much of the continuum and the 6563 \AA absorption may be from the night sky.

7 Caveats

This code currently runs only on fiber-mode data.

The code was developed under IDL version 8.1, and in principle will run under version 7. It does not use widgets or object graphics. It does require the IDL Astronomy User's Library⁷.

The Gaussian-extraction part of the code is slow. I'll look into speeding it up later. In the mean time Gaussian extraction is turned off by default in *ch_reduce*. It can be turned off in *ch_cals* by setting the *nogauss* keyword.

While the wavelengths match up well with those of the Yale reductions, I have not tested the fidelity of the solutions. Use this code for precise radial velocity work with caution.

Although this code has run on 1059 images from 566 nights without breaking, there is no guarantee that there are no bugs left, or that all the analysis is being done correctly⁸. The user always has the final responsibility for verifying data quality. Bugs or unexpected features may be brought to the author's attention at frederick.walter@stonybrook.edu

8 The Software

The procedures and the master calibration files are available in two tar files. Open http://www.astro.sunysb.edu/fwalter/SMARTS/NovaAtlas/ch_reduce/ch_reduce.html in your browser. You will find links to this documentation and to *ch_pros.tar* (the procedures) and *ch_calibs.tar.gz* (the master calibrations).

9 Coming Soon ...

Planned augmentations include:

- Sky subtraction capabilities, using actual spectra of the night sky.
- Testing of optimal extraction codes. At present the Gaussian fit is just that - a Gaussian fit (ignoring data previously flagged as bad). It is possible to flag significant deviations in the data against the model, and refit excluding those points. If properly tuned, this can remove cosmic rays otherwise missed.
- Extension of the code to slicer-mode data.

⁷<http://idlastro.gsfc.nasa.gov/>

⁸No process can ever be made completely foolproof, because fools are so ingenious.

10 Summary

This reduction package provides for better quality spectral extractions for *Chiron* fiber-mode observations of faint targets than those provided by default. Thirteen more orders, from 4085Å through 4510Å, are extracted, and all 1028 points along each order are available. Error bars (more precisely, the signal-to-noise ratio) are tracked for each point. The extraction does a better job subtracting the background than does the standard reductions. The extracted spectra, when divided by the flats, are flat. This simplifies the problems of flux calibration and order splicing. The products are either an array of flux-calibrated orders, or a single spliced spectrum with no gaps shortward of 8300Å. The data are stored in both fits format and as IDL save files.

Data are only as good as the extraction and calibration tools used. This code may enable more kinds of scientifically-useful observations with the *Chiron* spectrograph.

The IDL procedures and the master calibration files are available through my nova web page.

11 Acknowledgments

I thank the Provost of Stony Brook University, Dr. Dennis Assanis, for supporting my continued participation in SMARTS. I thank the SMARTS Chiron scheduler, Emily MacPherson, for service above and beyond the call of duty in scheduling Chiron observations of novae in a timely fashion. I thank Andrei Tokovinin for a careful reading of and comments on a draft of this document. I thank various experts, among them U. Munari, S. Shore, and R. Williams, for looking critically at and questioning some of the spectra I tried to feed them.

Not a penny of NSF money has been expended on this project.

12 References

- Markwardt, C. <https://www.physics.wisc.edu/~craigm/idl/fitting.html>
Pych, W. 2004, PASP, 116,148
Tokovinin, A. et al., 2013, PASP, 125, 1336
Valenti, J. 1999. available at <https://github.com/mattgiguere/idlutils/blob/master/thid.pro>
van Dokkum, P.G. 2001, PASP 113, 1420