Extremely Low-Cost Point-Source Spectrophotometry (ELCPSS)

John Beaver
University of Wisconsin – Fox Valley
1478 Midway Rd.
Menasha, WI 54952
john.beaver@uwc.edu

Charles Conger
University of Wisconsin – Plattville
a1478 Midway Rd
Menasha, WI 54952
congerc@uwplatt.edu

Abstract

We describe preliminary tests of a low-cost method for obtaining, reducing and calibrating stellar spectra. Instead of a post-focus spectrometer, we use an inexpensive, low-dispersion objective grating made by printing onto acetate with a laser printer, coupled with a low-end digital SLR camera as the detector. Although originally intended for educational use, we consider the possibility of using this technique to obtain accurately-transformed B and V magnitudes of stars without the need of an expensive photometric filter and filter wheel system. The results of two nights of observations of several bright stars are presented. Future plans are presented for more tests using a wider range of gratings, telescopes and detectors, and more advanced observing techniques that are likely to produce higher-quality data. But we show that even the crude observing techniques used for the test data can produce calibrated B and V magnitudes. General methods for reducing and calibrating the data are described, and some of the educational uses for ECLPSS are also considered.

1. Introduction

The usual method for taking spectra of stars and other astronomical objects is to use a telescope to gather and focus the light, and then allow the light to pass into a spectrograph attached to the eyepiece end of the telescope. The focused light passes through a slit, where it is dispersed by a diffraction grating into different wavelengths and then detected by a CCD array digital detector. A spectrograph such as this, even a relatively low-cost one, can be quite expensive for an individual amateur astronomer or even a high school or small-college astronomy program.

Extremely Low-Cost, Point-Source Spectrophotometry (ELCPSS) is an attempt to bring stellar spectrophotometry to those who otherwise would have little opportunity. We do this by using a simple low-dispersion diffraction grating, produced for almost nothing by printing a pattern of lines onto acetate with an ordinary laser printer. This grating is placed over the objective of a typical small amateur telescope and an ordinary low-end commercial digital SLR camera (DSLR) is used as the detector. Thus the entire telescope becomes the spectrograph, and all of the equipment is of the sort a typical small college, high school or amateur astronomer may already own (Beaver and Robert, 2011a).

These slitless spectra produced by ELCPSS are of rather low resolution, since the spectral resolution is limited by the observed size of the star image, and this is degraded by the acetate base of the objective grating itself. The point of ELCPSS is not to produce research-quality spectra in and of themselves. Rather, we hope to inexpensively obtain low-resolution spectra that are nonetheless fully wavelength and flux-calibrated, and using much of the same procedure as at any major observatory. This means ELCPSS has a fairly clear educational use. But we also show that ELCPSS may be useful as a roundabout but inexpensive method for obtaining calibrated BV (and perhaps BVR) stellar magnitudes.

Here we present the results of two nights of ELCPSS data. We demonstrate proof of concept regarding using it as a back door to stellar broadband photometry, and we consider its potential as an educational tool. We also lay out a strategy for making the rather complex data reduction and calibration process more accessible to the uninitiated. Future work to be done is outlined, and collaborations are solicited.

2. Useful Numerical Relations for Objective-Grating Spectroscopy

For objective-grating spectroscopy, the scale of the spectrum on the detector (in angstroms per pixel, for example) is set only by the line spacing of the grating, the focal length of the telescope and the physical size of the detector pixels. These data, combined with the number of pixels on the detector in the dispersion direction, then fix the possible range of wavelengths that can be imaged for a particular combination of grating, telescope and detector. We define the following quantities, with units chosen to give typical values near unity:

d = the grating spacing, in hundredths of an inch.

m = the order number of the spectrum.

p = the pixel size on the detector, in microns.

F = the focal length of the telescope, in meters.

With these definitions, a given number, *N*, of pixels in the dispersion direction, corresponds to a given range of wavelength in Angstroms, $\Delta\lambda$, given by:

$$\Delta\lambda = 2.540\,\frac{d\,p\,N}{m\,F} \qquad (1)$$

As an example, for the instrumentation used in this paper the pixel size was 7.88 micron, the grating was 200 lines/inch (*d*=0.5) and the focal length of the telescope was 2.0 m. And so for this case $\Delta\lambda/N$=5.0 Angstroms/pixel for the first order spectrum (*m*=1). As another example, to get everything from the direct image of the star (*m*=0) to 7000 Angstroms in the first order onto the detector, solving Equation 1 for *N* shows that about 1400 pixels are needed for this combination of grating, detector and telescope.

The lack of a slit ties the spectral resolution directly to the angular resolution. For a given angular resolution element, θ_R, in arcseconds, the corresponding wavelength resolution element, λ_R, in Angstroms, is given by:

$$\lambda_R = 12.3\,\frac{d\,\theta_R}{m} \qquad (2)$$

Since θ_R is set by the seeing, tracking, and especially the degrading of the image by the acetate grating, this means that one would want the smallest value of *d* possible, so as to maximize the wavelength resolution (i.e. minimize the size of the resolution element), even if this means the data is oversampled. But of course we also want *d* chosen such that the combination of telescope focal length and detector size puts the desired wavelength range onto the detector. For the data presented here, we estimate the size of the direct star image through the grating to have been about 10 arcsecond, and this corresponds to a spectral resolution element of about 60 Angstrom for our grating (d=0.5).

3. Observations

The observations presented here were carried out on two nights, June 7, 2010 (Night 1) and October 1, 2010 (Night 2). All observations were made from Appleton, Wisconsin, with moderately light-polluted skies (naked-eye limit m=4.5). Both nights were judged to be photometric. The observing log for the observations is given in Table 1. For reasons discussed below, multiple short exposures were combined to give the total exposure times listed in Table 1. The individual exposures were all between 1 and 30 seconds. Preliminary results from Night 1 were previously published (Beaver and Robert, 2011a; Beaver and Robert, 2011b), but errors in the reduction and calibration process have been corrected for the results presented here.

Star	Night	X	Exp. (s)
α Lyr	1	1.49	12
γ UMa	1	1.10	20
γ Dra	1	1.22	68
α Lyr	2	1.09	20
α Lyr	2	1.17	22
58 Aql	2	1.44	210
β Cyg	2	1.10	60
ζ Her	2	1.67	105

Table 1: The observing log for the stars observed for this paper.

For a post-focus slit spectrograph the spectrum automatically appears at the same location on the detector each time. But with an objective-grating, the spectrum could in principal appear anywhere on the detector, depending on the pointing of the telescope. And so if one is to be able to relate a given position in the observed spectrum to wavelength, either there must be a reference for wavelength in the image itself or there must be a way to exactly repeat the positioning of the star image in the telescope (and thus the location of the spectrum) from frame to frame.

Thus there are two obvious approaches. For what we call zero-order referenced (ZOR) spectra, we use the zero-order spectrum (the same as the direct image of the star) as a wavelength reference for each frame separately. Thus, for ZOR spectra, both the zero and first orders must be present on each frame, but there is no need to ensure that each exposure is taken with the same positioning of the object. Since the zero-order, direct image of the star is much more intense than the first-order spectrum, short exposures must be used so as to not overly saturate the zero-order spectrum. With proper guiding, however, one could in principal take a series of exposures for each object – long exposures to record the first-order spectrum, and short exposures to accurately determine the location of the zero-order spectrum (although one would still need to worry about scattered light and blooming from the saturated pixels).

The other approach, zero-order guided (ZOG), is to guide directly on the zero-order spectrum during the exposure, perhaps with an off-axis guider or by using the separate guide chip included with some astronomical CCD cameras. The ZOG technique would seem to have a big advantage over ZOR spectra, as the much brighter zero-order spectrum is out of the frame. But ZOG gives fewer choices of dispersion, given a particular off-axis guiding arrangement. For example, if one is using a fixed-prism off-axis guider, then the correct grating to put the desired maximum wavelength at the far edge of the detector will not necessarily put the desired minimum wavelength at the near edge. Put another way, a choice of grating that places the desired wavelength range across the detector may not at the same time put the direct image of the star in the guiding field.

For the spectra presented here, we have used the ZOR technique, in part to first test the efficacy of the crudest and lowest-cost observing techniques. A Celestron 8 (with uncoated corrector plate) was used with a 200 line/inch objective grating and a Nikon D40 SLR shooting in RAW mode (with noise reduction on) at ISO 1600. Exposure times were limited by the tracking of the (unguided) telescope and the fact that we could not too-badly saturate the zero-order spectra, as they are needed as a reference for the wavelength calibration.

We took the standard array of basic CCD reduction frames, including bias, dark and flat field. Flat fielding was accomplished by means of observations of a white screen illuminated by an old slide projector. We also used twilight sky flats to produce a final illumination correction. These frames were all taken with the objective grating in place. Since the sources of light for the flat field frames were diffuse over the entire field, so too was the spectrum produced, the zero and first order spectra mixing over the entirety of the image so as to produce images very much as if the grating were not present at all. Dark frames were taken implicitly for the longer exposures, as it is a property of the "noise reduction" feature of this particular DSLR to take a dark frame immediately after the exposure and to automatically subtract it. This feature can be turned off, but it was left on for these data. The built-in noise reduction also performs some median filtering, and this would have been a concern except that the spectral data is grossly oversampled in wavelength, and so this could not have removed "real" features of the spectra.

In addition to these basic CCD imaging calibration frames, we have taken others related directly to the calibration of the spectra. For wavelength calibration we observed a low-pressure sodium light, occluded by a screen except for a small pinhole, placed about 30 m from the telescope. A similar arrangement was used to make an artificial point source of a 300-W, 3100-K color-temperature quartz lamp. This was needed to flatten the spectra somewhat in order to better accomplish the final flux calibration.

4. Data Reduction and Calibration

The RAW Nikon images were extracted to FITS format (Wells et al., 1981) with the low-cost proprietary software package ImageTOOLSca (http://arnholm.org/astro/software/ImageTOOLSca/). This program effectively converts many consumer digital RAW formats to FITS, and writes a proper header with standard fields such as exposure time and date and time of observation. The data were extracted summing the three separate RGB mask images to single-band, 16-bit FITS images.

We used IRAF (Tody et al., 1986, Tody et al., 1993) to reduce and calibrate the data. The basic image reductions, incorporating bad-pixel repair, bias subtraction, flat-fielding (using the white-screen flats), and illumination correction (using twilight flats), were carried out in the usual way with the IRAF task CCDPROC.

The 1-d spectra were extracted from each 2-d image using the APEXTRACT package of IRAF, which fits a function to a trace of the spectrum on the 2-d image. Care was taken to use an aperture large enough to account for differences in tracking and focus, except for the extractions of the double star β Cyg, which required a smaller extraction aperture than we would have liked.

A dispersion in Angstrom/pixel was determined by measuring the separation in pixels of the zero and first order spectrum of the sole line in the observed sodium spectrum. This dispersion was then applied to the extracted 1-d spectra by directly editing the image headers. The zero-point for the wavelength calibration was determined separately for each spectrum using the interactive IRAF task SPLOT. Only a simple linear dispersion relation was applied, since for ZOR spectra nothing else would be meaningful given the fact that the individual spectra were at different locations on the detector.

A multi-order cubic spline was fit to the quartz lamp spectrum and all of the stellar spectra were divided by this spectrum in order to remove medium-scale variations before attempting flux calibration. This greatly facilitates the process of flux calibration because the instrumentation introduces medium-scale features in the spectrum due to, for example, the RGB pixel mask and interference coatings on the intervening optics (for example, the IR and UV blocking filter Nikon places over the CCD).

Flux calibration was performed with the IRAF tasks STANDARD, SENSFUNC and CALBRATE, using α Lyr as the sole spectrophotometric standard. For the second night, only the second observation of α Lyr was used for the calibration, as the first observation was inconsistent in overall scaling. This was the first observation of the night and the observing log noted that dewing may have been a problem for that observation (and this is consistent with the measured flux level being too low for that observation). No extinction curve was determined from the data; instead the mean extinction curve for KPNO was applied.

5. Results

We compare our observed, calibrated spectra to catalog spectra from the STScI HST Calibration Database Systems (CDBS) Stellar Spectral Atlases (STScI, 2012a). Figure 1 shows, plotted along with a Kurucz model-atmosphere α Lyr spectrum (Colina, 1995), the three α Lyr observations and the observations of γ Uma, also an A0V star. Figure 2 shows the observed spectrum of 58 Aql, along with a comparison spectrum of that star from the Bruzual-Persson-Gunn-Stryker (BPGS) Atlas of the CDBS. Figure 3 shows the observed spectra of β Cyg A and B, along with CDBS spectra of stars of similar spectral type. Figure 4 shows the observed spectrum of ζ Her, along with a spectrum of that star from the Wisconsin Halfwave Polarimeter (HPOL), retrieved from http://archive.stsci.edu, and Figure 5 shows the observed spectrum of γ Dra, compared to a BPGS catalog spectrum of a star of similar spectral type. The CDBS spectra were all obtained via the STScI Specview software package (STScI, 2012b). Finally, Figure 6 shows the flux-calibrated spectrum of the observed quartz lamp, compared to a 2900 K blackbody, in rough agreement with the rated color temperature of 3100 K for the bulb used.

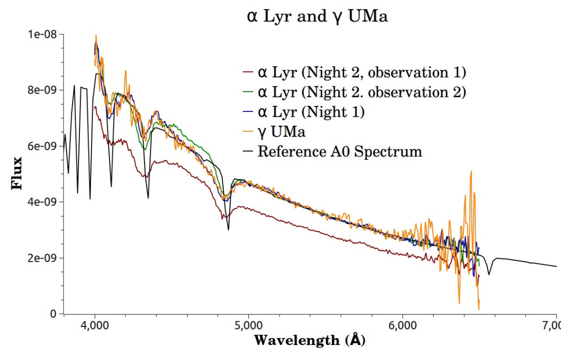

Figure 1: A Kurucz model atmosphere of α Lyr, along with our three α Lyr observations and the A0V star γ UMa.

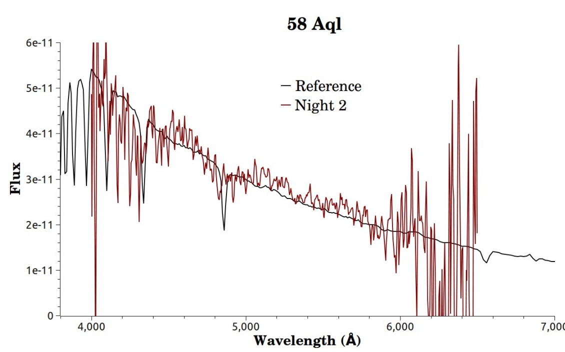

Figure 2: Our observed spectrum of 58 Aql, along with a comparison spectrum of that star from the BPGS atlas.

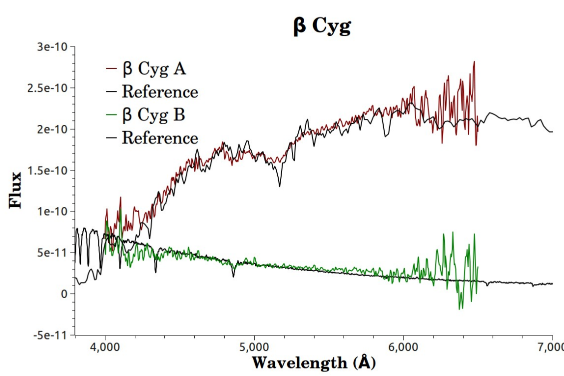

Figure 3: Our observed spectra of β Cyg A and B, along with CDBS spectra of stars of similar spectral type.

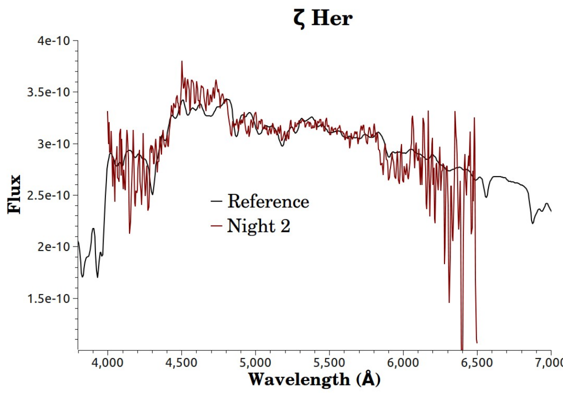

Figure 4: Our observed spectrum of ζ Her along with a spectrum of that star from the Wisconsin Halfwave Polarimeter.

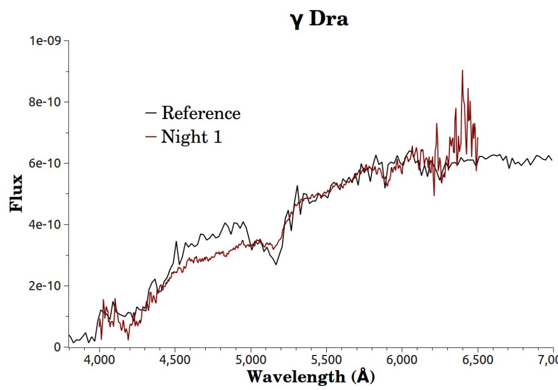

Figure 5: Our observed spectrum of γ Dra, compared to a BPGS catalog spectrum of a star of similar spectral type.

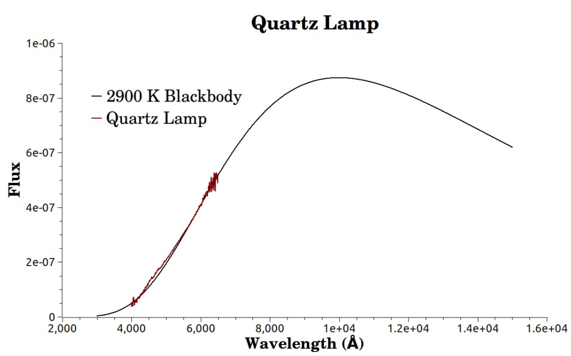

Figure 6: The flux-calibrated spectrum of the observed quartz lamp, compared to a 2900 K blackbody.

We note that α Lyr was used as the spectrophotometric standard for both nights, and so the agreement between these spectra and the catalog spectrum is unremarkable. The first α Lyr spectrum from Night 1, however, was not used for the flux calibration, and so its agreement (or lack thereof) is meaningful. This spectrum shows good relative flux calibration, but a noticeable error in absolute calibration, appearing too faint by the equivalent of a few tenths of a magnitude. This is not surprising as the observing log for that particular observation reports "dewing problem?" The rest of the observed spectra agree well with the comparison spectra, with the exception of β Cyg B, which shows a noticeable error in relative flux calibration, appearing too red. We suspect this may be due to the difficulty in extracting this spectrum without contamination from its much brighter close neighbor, β Cyg A.

With the exception of α Lyr, the comparison spectra presented in these graphs have all had a scale factor applied to best match the observed spectra. This is because the comparison spectra we found had only relative flux calibration, and in some cases we could only find comparison spectra of the same spectral class (but not the same star). We will see in the next section, however, that we have good reason to believe the absolute flux calibrations of the spectra are reasonable, with the exceptions of β Cyg B and the first Night-2 α Lyr spectrum, as previously noted.

6. ELCPSS as a Back-Door to Broad-Band Photometry

Broad-band B and V filter photometry takes complex spectral information over a range of nearly 3000 Angstrom, and compresses it into just two numbers. The success of broad-band filter photometry depends on the spectral response of the instrument-detector system, and the spectrum of the observed source, being similar enough to that which set up the standard system. If not, then a one-to-one correspondence between the instrumental and standard magnitudes may not be obtainable.

Even a low-resolution spectrum contains far more information than a set of broad-band photometric measurements. In particular, the problems of flux calibration and extinction correction for spectra are in many ways simpler than the corresponding problems of transforming instrumental broad-band magnitudes to standard magnitudes.

Thus once a spectrum is accurately wavelength and flux calibrated, a standard magnitude can be extracted simply by convolving the appropriate magnitude bandpass function with the spectrum.

We have done this for the spectra here presented using the IRAF task SBANDS, and the filter profiles for B and V from Ažusienis and Straižys (1969). These profiles are normalized, and so we used the filter extractions of our standard star (α Lyr) to set the zero points for the magnitudes. Table 2 shows the results of these extracted magnitudes, as well as the values from SIMBAD for comparison. As can be seen, even with only one observation of a single standard star and the use of a mean extinction curve for a different site, the values are all to within 0.2 magnitude in V and 0.13 in B-V, except for β Cyg B and the first α Lyr spectrum from Night 2, as already noted.

Star	X	V	V'	B-V	B'-V'
α Lyr	1.09	0.27	0.03	0.01	0.00
58 Aql	1.44	5.45	5.63	0.18	0.08
ζ Her	1.67	2.71	2.80	0.69	0.63
γ Dra	1.22	2.23	2.23	1.65	1.52
γ UMa	1.10	2.38	2.43	0.01	0.04
β Cyg A	1.10	3.20	3.09	1.19	1.09
β Cyg B	1.10	5.23	5.09	0.16	-0.06

Table 2: The extracted V and B-V magnitudes for the program stars, compared to the values from SIMBAD (the primed columns).

In principal, one should be able to do much better by adopting a hybrid technique. Instead of treating these extracted magnitudes as the final values, think of them instead as raw instrumental magnitudes for a program of broad-band photometry that includes observations of broad-band photometric standards as well as the spectrophotometric standards used for the spectroscopic flux calibrations.

Even using one spectrophotometric standard star and a generic mean extinction curve to establish the calibration, the extracted magnitudes should still be closer to the standard system than many detector-filter-instrument combinations in use. Thus one could in principal use ELCPSS to observe a regular program of photometric standards in order to determine accurate standard magnitudes from these almost-calibrated extractions. The extracted magnitudes should transform very smoothly to the standard system since they are almost there to begin with.

This would also solve the problem of a relative lack of bright spectrophotometric standards (as compared to broad-band photometric standards), thus allowing ELCPSS to be useful for differential photometry.

7. Educational Uses of ECLPSS

Both the process of carrying out ELCPSS and the calibrated spectra that result can be of educational use. We choose to illustrate this with an example. Although we present this as a hypothetical case, we are attempting to actually carry out much of this at the University of Wisconsin – Fox Valley, a small two-year campus with a primary mission of transfer to one of the four-year campuses of the University of Wisconsin system.

A beginning astronomy major starts a project to record and calibrate ECLPSS spectra, with the goal of compiling an atlas of stellar spectra. Some of the spectral types have already been recorded by a student from the previous year. The current student adds to the previous work, getting spectra of multiple stars for some spectral types. In the process, she has to plan and execute an observing run, and then reduce and calibrate the data with IRAF using a pre-configured virtual machine environment set up especially for this purpose (see Section 8). Although all of this is done with modest equipment, the basic observing and data reduction techniques and the software used are all likely to be similar to what she will eventually use in graduate school.

Although the data themselves may not contribute directly to basic research, the results still have a purpose beyond that of a mere student exercise. This is because the spectra are to be used for introducing the concept of spectral classification to the beginning Survey of Astronomy courses for non-science majors. Instead of the students learning about stellar spectra from examples in a textbook or on the Internet, obtained by a process to which the students have no direct connection, they are presented with data obtained locally by one or more of their peers.

In fact, the astronomy major makes a presentation to the class, explaining how the spectra were taken, and the actual telescope used is set up for the students. The Survey of Astronomy students gain a learning experience that has a visible and immediate connection to their experience, and the astronomy major is motivated by contributing to their learning with data she gathered herself. This is “service learning” in the true sense of the term.

8. Facilitating the Reduction and Calibration Process

Since the goal of ELCPSS is to make spectrophotometry more accessible, the complex process of reducing and calibrating the spectra is a concern. To address this, we have developed a virtual machine (VM) with an open-source operating system (Ubuntu), with all of the necessary software (also free and open-source) pre-configured and including documentation and tutorials, for reducing and calibrating ELCPSS spectra, along with other common astronomical research projects (Beaver, 2012). The goal is to provide a self-contained research-computing environment that can be shared between users, thus sidestepping for beginners the difficult process of establishing such a computing environment from scratch.

Since the virtual machines can be portable and cross-platform, individual VMs can be customized for particular combinations of projects and users. We see this, then as a key component to making ELCPSS practical and useful. For example, the coauthor of this paper had no prior training in astronomy before beginning this project, and he learned much of the process via a virtual machine specifically developed for this task.

9. Extending the Wavelength Range

The usable wavelength range for the current set of observations is from 4000-6500 Angstrom. The short wavelength cutoff is mostly from the small amount of light output at short wavelengths by the spectral flat-fielding lamp used (see above), but also from absorption by the corrector plate and the bandpass filter built into the CCD camera.

The long wavelength cutoff is mostly due to the IR blocking properties of the CCD bandpass filter, and it is possible to have this filter removed and replaced with a broad-spectrum window, extending the spectral response out to much longer wavelengths. This would possibly allow for extraction of the R-band as well as B and V.

The traditional complication of near-IR spectroscopy is that a diffraction grating produces a second-order short-wavelength spectrum that overlaps the first-order, long-wavelength spectrum. The usual method of dealing with this is to take a second spectrum through an order-separating filter that blocks the short-wavelength first-order spectrum. Since the point of ELCPSS is that it is simple and low-cost, the introduction of a filter-wheel mechanism and a specialized order-separating filter is counter to the goals of the project.

There may be another approach however. Since the spacing of the lines in the grating is relatively large compared to a traditional grating, it may be possible to execute control not only over the spacing of the lines, but also over their actual transmission profile. This raises the possibility of making a grating that produces no significant second-order spectrum in the first place, thus eliminating the need for an order-separating filter and allowing for one to record a spectrum from 3500-1050 Angstrom in one exposure. Thus one could in principal obtain BVR magnitudes with a single spectrum.

The correct transmission profile for a grating that produces only odd orders (and thus no second-order spectrum) is that of equally-sized adjacent dark and clear bands (Jenkins and White, 1950). Indeed, the 100 line/inch grating we produced already approximates this, as it was printed from a digital file with this pattern, and the 600 dpi laser printer was able to approximate this pattern. We believe that with a high-end inkjet printer, it may be possible to produce gratings that effectively suppress the second-order spectrum. Thus ELCPSS could perhaps be used as a roundabout means for extremely low-cost BVR photometry (ELCBVR?).

10. Conclusions

We find that the use of a low-cost objective grating on a small telescope, coupled with a low-end consumer DSLR camera, can yield wavelength and flux-calibrated spectra. Standard B and V magnitudes can be extracted from the spectra, and our early results suggest this may be a useful way in some circumstances to obtain accurate B and V stellar magnitudes at very little cost in equipment.

Although the process itself has clear educational uses, much work remains to be done to determine whether ELCPSS could find practical use as a tool for extending broad-band photometry to users who otherwise would be limited to visual magnitude estimates. Further tests need to be done using ZOG techniques, and the practical limits in magnitude and precision need to be determined, for both absolute spectrophotometric calibration and for “hybrid” techniques that use ELCPSS to obtain nearly-calibrated “instrumental” magnitudes with broad-band photometric standards providing the final calibration. Experiments should be carried out with other grating materials and printing methods, and a wider array of telescopes and detectors should be tested. Individuals wishing to collaborate on the ELCPSS project should contact the authors.

11. Acknowledgements

We engaged in many helpful discussion with Doug Fowler. This work was supported in part by the following grants:

- UW - Fox Valley Foundation Grant, 2010
- UW - Fox Valley Professional Development Grants, 2010, 2012
- UW Colleges Department of Computer Science, Engineering and Physics Professional Development Grants 2010, 2012

12. References

Ažusienis, A. and Straižys, V., (1969). *Soviet Astronomy* **13**, 316.

Beaver, J.E., "A Virtual Astronomical Research Machine in No Time (VARMiNT)," (2012). *2012 SAS Symposium Proceedings*, submitted.

Beaver, J.E and Robert, D.M., "A CCD Spectrograph for One Dollar ," (2011a). in *Earth and Space Science: Making Connections in Education and Public Outreach*, ASP Conference Series, **433**, 425.

Beaver, J.E., Robert, D.M., "A CCD Spectrometer for One Dollar," (2011b). in *Cosmos in the Classroom 2010: A Hands-on Symposium on Teaching Introductory Astronomy*, ed. Andrew Fraknoi, C1 (Astronomical Society of the Pacific)

Colina, L., "CDBS Kurucz Stellar Atmosphere Atlas," (1995). *STScI OSG# Instrument Science Report* SCS CAL-006.

Jenkins, F.A., White, H.E., (1950). *Fundamentals of Optics* (New York, McGraw Hill) p. 310, 327.

Space Telescope Science Institute, (2012a). Calibration Database System (CDBS), http://www.stsci.edu/hst/observatory/cdbs/astronomical_catalogs.html, Retrieved, 11 April 2012.

Space Telescope Science Institute, (2012b). Specview, *www.stsci.edu/institute/software_hardware/specview*, Retrieved, 11 April 2012.

Tody, D., "The IRAF Data Reduction and Analysis System," (1986). in *Proc. SPIE Instrumentation in Astronomy VI*, ed. D.L. Crawford, **627**, 733.

Tody, D., "IRAF in the Nineties," (1993). in *Astronomical Data Analysis Software and Systems II*, A.S.P. Conference Ser., eds. R.J. Hanisch, R.J.V. Brissenden, & J. Barnes, **52**, 173.

Wells, D.C., Griesen, E.W., and Harten, R.H., "FITS - a Flexible Image Transport System," (1981). *Astronomy and Astrophysics Supp. Series* **44,** 363–370.