CONFIRMATION OF VARIABILITY IN THE SU UMA-TYPE DWARF NOVA V1504 CYG

IAN J. CROSSFIELD
REU Program, Dept. of Physics and Astronomy, University of California, Irvine
Electronic mail: icrossfi@uci.edu

Supervisor: Prof. T. Smecker-Hane (UCI)
Postdoctoral Researcher: W. Skidmore (UCI)

ABSTRACT

The cataclysmic variable V1504 Cyg is observed at the UCI Observatory for periodic variability, showing that the observatory is capable of performing time-resolved photometry on a dwarf nova system. The software package IRAF is used to calibrate the images and to measure the flux from the target and several reference stars with aperture photometry. These same stars are then used to compute differential photometry of the target. A statistical test is employed to empirically determine the presence of variability in the target star. Finally, a theoretical light curve is fit to the data in an effort to confirm the known orbital period of the system.

1. INTRODUCTION

1.1 Dwarf Novae

Dwarf novae are binary star systems composed of a white dwarf and (typically) a lower main sequence red dwarf: the primary and secondary stars, respectively. The secondary star grows in radius until it begins to transfer mass to an accretion disk surrounding the primary star. These systems undergo recurring outbursts on the order of days to weeks, during which their brightness increases by 2 to 5 magnitudes. The source of the energy released comes from the release of gravitational energy as mass accretes from the accretion disk onto the primary star itself.

1.1.1 Equipotentials

Around any set of objects in space, there exists a set of gravitational equipotentials, along which matter experiences the same gravitational force. For two objects of masses M_1 and M_2 orbiting each other and separated by a distance r, the two-dimensional gravitational potential in co-rotating coordinates is given by

$$\Phi(x,y) = -\frac{GM_1}{(x^2+y^2)^{\frac{1}{2}}} - \frac{GM_2}{[(x-r)^2+y^2]^{\frac{1}{2}}} - \tfrac{1}{2}\Omega_{orb}^2[(x-\mu r)^2+y^2]$$

$$\mu = \frac{M_1}{M_1+M_2}$$

Figure 1 shows the equipotentials in such a system. The closed surfaces surrounding M_1 and M_2 are known as the Roche lobes; the points marked L_x are the binary system's Lagrange points, at which the gravitational potential is zero. In a

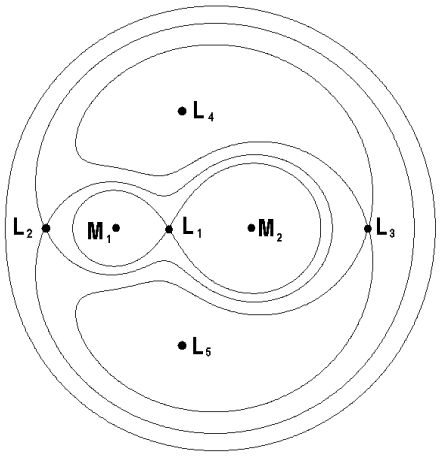

Figure 1 - Equipotentials and Lagrange points

dwarf nova system, the secondary star is less evolved than the primary. When the primary has already become a white dwarf, the secondary is still burning hydrogen and filling its core with helium. As this proceeds, the secondary's radius increases until it distorts and fills its Roche lobe completely. Matter then spills over the L_1 point, and falls down the potential well toward the primary star.

The infalling material is called the matter stream. It flows down, down until, at a fraction of its original height, it goes into orbit around the primary star, forming what is called the accretion disk. It does not impact the primary directly due to the conservation of angular momentum, since evena very small initial transverse velocity at L_1's height gives the stream significant angular momentum.

1.1.2 Accretion Disk and Outbursts

When a dwarf nova is not in outburst, it is said to be quiescent. During this time, matter is falling onto the accretion disk from the secondary star, but not from the accretion disk onto the primary; the accretion disk is not viscous enough to dissipate the energy necessary to lower the orbit of its constituent matter and deposit material onto the star.

The manner in which the outburst begins is still uncertain; the currently favored theory hypothesizes that that they result from increased viscosity in the accretion disk. As more and more matter is deposited, the surface mass density of the disk increases. At some critical point, this increased density triggers a rise in the viscosity of the disk. This increases the amount of energy lost due to friction, to the point that matter accretes from the disk onto the primary star in significant quantities. The gravitational potential energy lost by the accreting matter is emitted as radiation, which we observe as in increase in the brightness of the dwarf nova. Unlike classical novae, dwarf novae's outbursts are *not* due to nuclear burning.

Of course, angular momentum must still be conserved. So as matter is accumulating on the primary star, the accretion disk expands outwards. Though the matter stream is still delivering matter to the disk, this expansion, combined with the mass transfer from disk to star, lowers the surface mass density of the disk. Eventually this causes the viscosity to decrease to pre-outburst levels, and the star slips back into quiescence.

1.1.3 Types of Dwarf Novae

Dwarf novae are classified based on the behavior of their light curves. There are three classes of dwarf novae, each named after the star which is the prototype for that class: SS Cyg, Z Cam, and SU Uma. All have the outbursts characteristic of cataclysmic variables, but the differ in important ways.

SS Cyg systems are the least unusual. Their outbursts occur at regular intervals and are always of about the same magnitude. Z Cam systems occasionally go into 'standstills,' during which the system's brightness is nearly constant. SU Uma systems have the shortest orbital periods (1-2 hours) and have occasional 'superoutbursts' that are about a magnitude brighter than the regular outbursts.

1.1.4 What We Can See

The most obvious observable feature of a dwarf nova is its outburst, during which the star's brightness dramatically increases. Since the star is at its brightest during outburst, it follows that this is the best time to collect data on the star. To alert astronomers to currently bright stars, the American Association of Variable Star Observers (www.aavso.org) releases daily news flashes that tell what stars have recently been observed to be in outburst. I selected my target star based on one of these news flashes.

The other feature that should be visible in a dwarf nova is an ellipsoidal modulation of its brightness curve. There should be measurable changes in the brightness of the system on a timescale comparable to that of the orbital period. Since the orbital periods of these stars are only a few hours, they should be detectable with a single night's observing.

When observing a dwarf nova soon after its outburst, one should thus expect to find two elements: a long-term downward trend as the outburst fades, and a short-term periodic oscillation of the same length as the orbital period. Since it takes days for a star to settle out of outburst, one expects that over the course of one night's observations (a few hours) the downward trend would be approximately linear. I used this assumption when determining what function to fit the data to, as described below in Section 3.3.

1.2 UCI Observatory

The facility used to obtain these data was the University of California at Irvine Observatory (Figure 2). The observatory's main telescope is a 24 inch reflector, to which is attached an SBIG ST-9XE CCD. This camera is 512 x 512 pixels, has a scale of 0.809 arcsec/pix, and has a gain of 2.839 e^-/adu. Both telescope and CCD can be controlled either from the observatory or remotely, though none of my data was collected remotely.

The two most serious problems I encountered with the UCI Observatory were excessive sky brightness and coarse guiding. The former was certainly the most serious. Located in central Orange County, the Irvine night sky is lit up with the reflected glow of the greater Los Angeles area. Though this effect is most severe in the northwestern portion of the sky, it is enough to limit the observatory to viewing stars of about magnitude 16 or less – and that with exposure times of ten minutes or more.

The long exposure times – 300 seconds – that were required to observe V1504 Cyg resulted in a second difficulty: the lack of fine guiding at the observatory. The telescope has the capability to track stars by using an ST6 CCD on a smaller 5 inch telescope, which is affixed to the side of the main telescope. Though this guides the telescope accurately enough, it is only precise to within a few arcseconds. Even for short exposures, seeing at the UCI Observatory is usually limited to 3"-4", on a good night, and so the stars in my images had full widths at half maximum (FWHM) ranging from 5" to 8" (4 to 6 pix).

Actually, the guiding was more of an issue than even this FWHM increase suggests. If it simply spread out the stars light into a larger, but still-symmetrical distribution, that would not have been a problem. Instead, the stars were 'smeared' unevenly back and forth across the CCD, which resulted in highly irregular stellar profiles. This was unfortunate, because it forced me to forego PSF-fitting photometry for the less-accurate aperture photometry method.

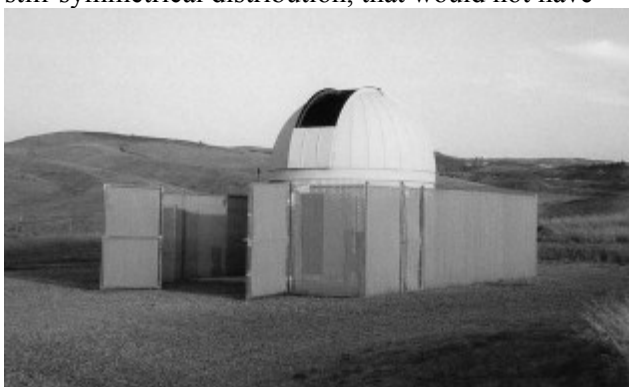

Figure 2 - UCI Observatory

2. DATA EXTRACTION

2.1 Frame Corrections

Once a series of CCD images has been taken, they must undergo a series of corrections before data can be extracted. These calibrations take into account the non-uniform quantum efficiency of each pixel in the CCD. Both bias subtraction and flat-fielding were performed with the Image Reduction and Analysis Facility (IRAF)[1], public software developed by astronomers at the NOAO and sponsored by the NSF. No dark subtraction was performed due to the low dark current (≈ 1 adu/pix/sec).

2.1.1 Bias Subtraction

The first correction to be made is bias subtraction. When read out, a CCD will always have some non-zero number of photon counts – even for a zero-second exposure, for which the entire CCD would ideally read zero counts. The bias of an image is the extra and undesired contribution to the photon count from this quality of the CCD, and it is not necessarily constant from pixel to pixel.

To correct for bias, a series of three zero-second exposures are taken and median-stacked together using the ZEROCOMBINE command. In median stacking, a new image is created for which each pixel's value is the median pixel value of the three bias frames. The median is preferred over a mean because it tends to discount anomalously high photon counts (so-called 'cosmic rays'). This new bias image is taken to be a reasonable representation of the bias of the CCD, and is then subtracted with the command CCDPROC from all other images that were taken.

[1] Public software developed by astronomers at the National Optical Astronomy Observatory and sponsored by the NSF.

2.1.2 Flat-Fielding

The second and more important correction to be performed on the images is flat-fielding. Flattening adjusts the data images by correcting for the variable response of the CCD pixels. This is necessary because not all 2^{18} pixels are created equal. Some will tend to be more sensitive to light and will record higher-than-desired counts; others will be less sensitive, and will record values lower than would be preferred. If the stars do not move appreciably across the CCD from the beginning of a series to the end, then flat-fielding may not be necessary. In my images, the star field both rotated and moved across the CCD as the night went on, so flat-fielding was definitely required.

Again, a series of three flat-field images are taken and median-stacked, this time with FLATCOMBINE. A flat-field image is an image of a region without a significant brightness gradient. The most common choices are the sky at twilight (before stars are visible) and the inside of the observatory dome. Since the CCD has such a small field of view (6.9" x 6.9") there is a negligible different in brightness across these images. The new flat frame is normalized to a sensitivity of one by dividing each pixel in an image by that image's median pixel value. The data images are then divided by this normalized flat frame, again using CCDPROC. For example, a pixel in one image that measured 5 000 counts but whose sensitivity was determined to be 1.2 would be corrected to have a new value of 5 000/1.2 = 4 167 counts.

2.2 Aperture Photometry

A favorite method of measuring the brightness of stars on CCD images is that of aperture photometry. Using some software package (in IRAF, the IMEXAM task is used), a synthetic aperture of specified radius is constructed around the star being measured. The pixel values of every pixel inside this aperture are then summed, which gives the total measured flux from that region in the sky. However, there is also a significant amount of background light in the sky that must be subtracted. An annulus of given width is constructed some distance past the edge of the aperture. The median pixel value inside this annulus is taken as the sky brightness. The sky brightness times the number of

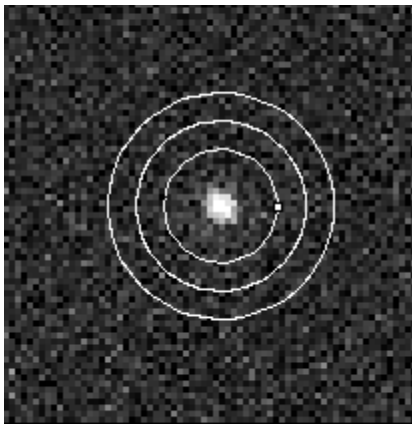

Figure 3 - Aperture photometry

pixels in the aperture is subtracted from the total pixel count inside the aperture, giving the measured flux coming from the target star. Figure 3 shows an example of aperture photometry being performed on a star: an aperture radius of 10 pixels, a buffer of 5 pixels, and a sky annulus 5 pixels wide.

Once fluxes are measured for the stars of interest, instrumental magnitudes can be computed with the familiar formula

$$m_* = m_{ref} - 2.5\log\left(\Phi_* / \Phi_{ref}\right)$$

where for IRAF $M_{ref} = 25$ and $\Phi_{ref} = 1$. Thus a star with a measured flux of 10 000 adu would have an instrumental magnitude $M_* = 15$.

2.3 Differential Photometry

Magnitudes can therefore be easily computed for a star in any number of images. However, if a thin cloud passes in front of the star during an image, less flux will necessarily be measured, and the star will appear to have decreased in magnitude. Stars can also appear to brighten as they rise overhead, and then dim again as they set. When a star is lower in the sky, its light will travel through more air, and more of its light will be deflected than when the star is directly overhead.

Fortunately, a method exists to deal with these difficulties: the method of differential photometry, first described by Howell (86). In differential photometry, several constant-brightness companion stars are used to correct the target star's brightness for non-photometric effects. The target star is traditionally called V, because it is expected to be variable. One reference star is called C, for Comparison, and another called K, for checK. I

used three companion stars, C, K1, and K2 – see Figure 4. This technique works because any detrimental atmospheric conditions (clouds, changing airmass, etc.) will be on a scale much larger than the telescope's field of view, and will therefore multiply all stars' fluxes by the same factor, and thus decrease all stars' magnitudes by the same constant.

For example, if a thin cloud that absorbed 10% of all starlight passing through it moved in front of our magnitude 15 star, one would measure only 9 000 counts, and one would calculate the magnitude to be 15.11. Using only absolute photometry, it would be impossible to know whether this magnitude increase was due to the star dimming or due to some terrestrial factor. Using differential photometry, one would see that the reference stars had all dimmed by 0.1 magnitudes, and that what was initially perceived as a change in the brightness of the target star was actually due only to atmospheric effects.

In differential photometry, then, the important data are the *differences* in magnitude. Usually several stars are used as reference, and their average magnitude subtracted from the target star's. This is done for each image in the series, creating the final data set V-(avg).

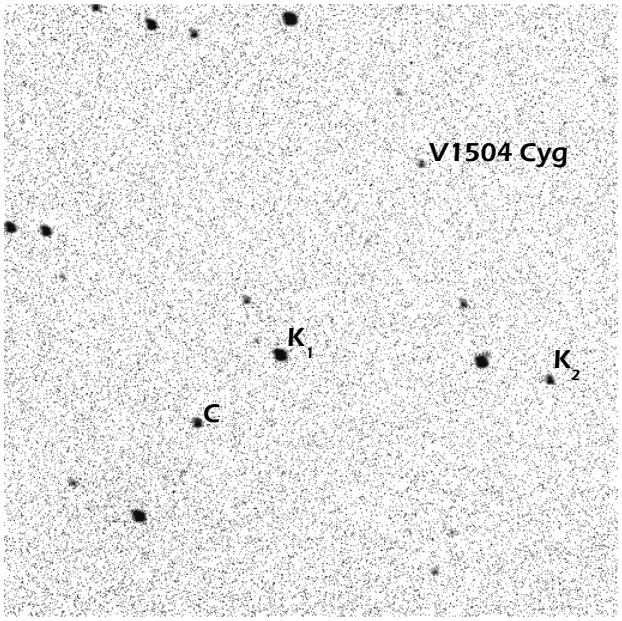

Figure 4 - V1504 Cyg and companion stars

3. DATA ANALYSIS

3.1 The Data

The star V1504 Cyg, an SU Uma type cataclysmic variable, was observed in an attempt to measure its orbital period. The series of 300-second exposures (40 in all) was collected on the night of August 6, 2003 using a V-band filter. The total time from start to finish was about six hours.

Bias corrections and flattening were performed on the images, and aperture photometry was conducted using several different aperture radii. Figure 5 shows the raw light curves for the three reference stars and the target star plotted versus the fractional Julian Date. Immediately, two things are clear from this graph. First, the effect of clouds on instrumental magnitude at time index 0.87, 0.92, and 0.95 is plain. The increasing cloud cover ultimately resulted in the termination of the image sequence. Second, although the C, K1, and K2 light curves all appear quite similar, the V curve differs substantially from the other three. This is a good initial sign that the V star may be fundamentally different from the others.

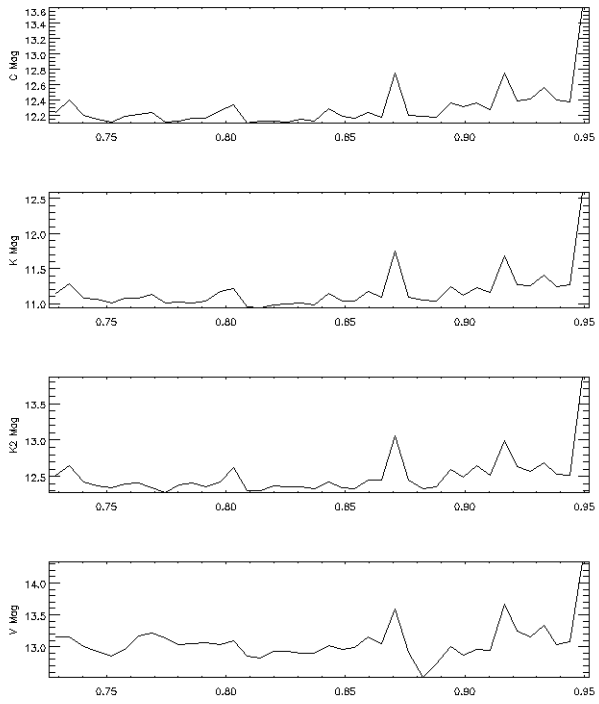

Figure 5 - Raw light curves for aperture radius r = 6 pixels

3.2 Statistics

3.2.1 Determining Uncertainties

Of prime importance in any experiment is the determination of experimental uncertainty. Statistical techniques for differential photometry are well established by, e.g., Howell (88,89). For any experimental variance σ^2, there are two components: σ^2_{stat}, statistical variance based on the expected distribution, and σ^2_{inst}, the instrumental variance, which must be estimated from collected data, and minimized if possible. These add in quadrature, so that

$$\sigma^2 = \sigma^2_{stat} + \sigma^2_{inst}$$

For the statistical uncertainty, we use the Poisson distribution (as is standard practice for counting experiments). The uncertainty in the stellar flux is then the square root of the total measured flux,

$$\sigma_{stat,\Phi_*} = \sqrt{N_* + n_P\left[N_S + N_b + \left(\frac{N_r}{g}\right)^2\right]}$$

N_* is the total sky-subtracted flux from the target star, n_P is the number of pixels in the aperture used, N_s is the median sky count, N_b is the median count from the bias frame, N_r is the read noise, calculated at 20.6 e^-/pix, and g is the gain, calculated at 2.84 e^-/adu.

Taylor (97) describes the method of error propagation, in which one variable's uncertainty is transformed into the uncertainty in a function of that variable. The uncertainty in measured magnitudes becomes

$$\sigma_{stat,m_*} = \left|\frac{\partial m}{\partial \Phi}\right|\sigma_{stat,\Phi_*} = \frac{2.5}{\Phi_* \ln(10)}\sigma_{stat,\Phi_*}$$

So for a star with a stellar flux of 80 000, an aperture 10 pixels wide, a sky count of 2 500, and a bias count of 100 (all typical values for this run), the statistical uncertainty would be ±0.013 mag. All statistical magnitude uncertainties were of this order.

An uncertainty of ±0.01 mag would be excellent, were there negligible instrumental error. However, the uncertainty stemming from the telescope itself (for reasons described above in Section 1.2) is significantly higher. A good estimate of the instrumental uncertainty, described in Howell (88), is to look at the subtraction of the check star from the comparison star (C-K). Ideally, this value would be identically zero for all the images. The more widely distributed (C-K) is, the more uncertain are the measurements. The best way to minimize the instrumental uncertainty is to choose the aperture radius that minimizes σ^2_{C-K}. In Figure 6, you can see plots of C-K for four different aperture radii – 4, 6, 8, and 10 pixels. Of the four, a radius of 6 pixels provided the lowest C-K variance (0.00145 mag^2), and thus the smallest uncertainty in C-K. I consequently used the data extracted with the 6-pixel-radius aperture for the remainder of the data analysis.

The instrumental magnitude uncertainty, then, is $\sigma_{inst,m}$ = 0.038 mag. This is increased only slightly by the statistical uncertainty, which raises

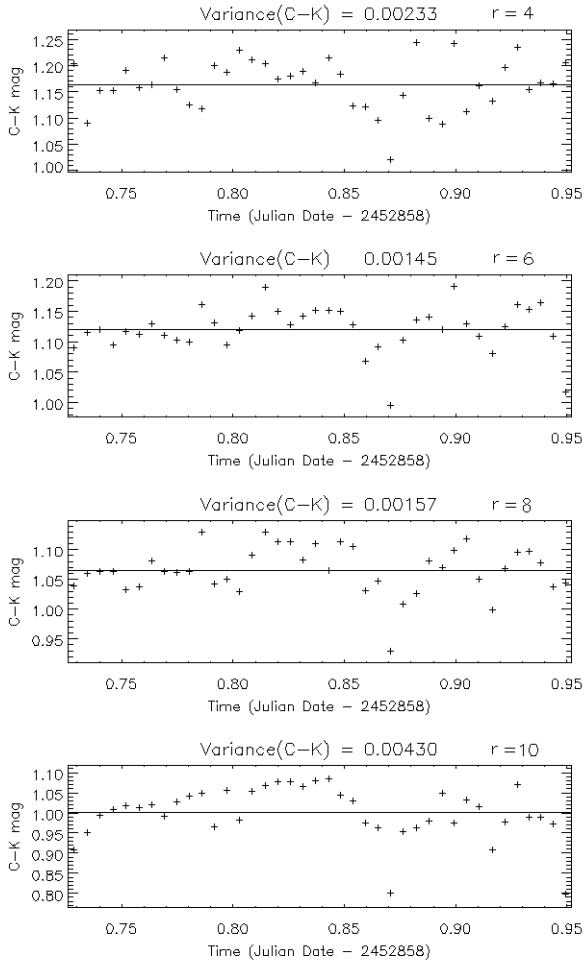

Figure 6 - C-K at four aperture radii

the final uncertainty in a single star's magnitude to around σ_m = 0.04 mag (depending on the uncertainty in the star's magnitude).

With these uncertainties, a weighted average of the three comparison stars' magnitudes was then constructed:

$$\bar{m} = \frac{\sum_{i=1}^{3} m_i w_i}{\sum_{i=1}^{3} w_i}, w_i = \frac{1}{\sigma^2_{m_i}}, \sigma_{\bar{m}} = \frac{1}{\sqrt{\sum_{i=1}^{3} w_i}}$$

The uncertainties of the final subtracted magnitude measurements are then

$$\sigma_{V-\bar{m}} = \sqrt{\sigma^2_{stat,m} + \sigma^2_{inst} + \sigma^2_{\bar{m}}}$$

3.2.2 Determining Variability

A method for statistically testing whether a star is variable – regardless of periodicity – is described by Howell (88). It compares the variances of V-C and C-K and determines whether V-C is significantly variable, given C-K's variance. The test is based upon the test statistic

$$F_{test} = \frac{\sigma^2_{V-C}}{\Gamma^2 \sigma^2_{C-K}}$$

$$\Gamma^2 = \left(\frac{N_K}{N_V}\right)^2 \left[\frac{N_C^2(N_V + P) + N_V^2(N_C + P)}{N_K^2(N_C + P) + N_C^2(N_K + P)}\right]$$

$$P = n_P\left[N_S + \left(\frac{N_R}{g}\right)^2\right]$$

The terms are: Γ^2, a factor that scales one variance based on the other; N_V, N_C, and N_K, the median over the image series of the sky-subtracted fluxes from the V, C, and K stars, respectively; N_S, the median sky brightness over the image series.

The statistical test used is an F-test, commonly used to compare two variances to determine whether one is significantly different from the other. Still following Howell (88), we formally test the hypothesis

$H_0 : F_{test} \le 1$ (V is not significantly variable)

against

$H_A : F_{test} > 1$ (V significantly varies)

H_0 is rejected at the confidence level α if

$F_{test} > F(1-\alpha, n_1 - 1, n_2 - 1)$

using the number of observations for both n_1 and n_2.

Over the series of images observed, we find that F_{test} = 2.81, and so V1504 Cyg is variable at a confidence level of 99.8%. The impression given by Figure 4 – that the V star's lightcurve is substantially different from that of the other stars – is borne out by this test.

3.3 Fitting the Data

Statistically, then, V is variable. The next question is, is the star periodic? To answer that question, I used the Interactive Data Language (IDL).[2] I used IDL's CURVEFIT routine to fit different curves to the data, and to determine which curve proved to be the best fit. I considered three curve types: an ordinary sinusoid, a sinusoid plus an exponential, and a sinusoid plus a first-order polynomial. The first of these was quickly rejected because of the measurable decrease in magnitude over the course of the night. The system was still in the process of outbursting when it was observed, and the observations clearly show that brightening trend. The second and third models returned virtually identical curves: when they were superimposed on top of each other, no significant difference could be detected. The time between outbursts for V1504 Cyg is approximately 5.8 days, and images were taken over only six hours, or 4% of that time. Over a time span so relatively short, there was essentially no difference between exponential and linear terms. So that a simpler theoretical model could be used, the sinusoid and first-order polynomial was used. It is the function

$$M(t) = A_0 \sin(A_1 t + A_2) + A_3 t + A_4$$

with five free parameters: A_0, the amplitude of variation; A_1, the angular frequency; A_2, the phase; A_3, the rate of change with time; A_4, the vertical offset. Figure 7 shows the data set with uncertainties; over it is plotted the fit function with the lowest calculated χ^2 value.

CURVEFIT uses minimization of χ^2 to 'home in' on a best-fit solution. This works well, except insofar as it means that the routine is unable to distinguish a local χ^2 minimum from the absolute χ^2 minimum over some frequency range. The resultant best-fit parameters can therefore change based on the initial input parameters – the most interesting parameter, of course, is the angular frequency, from which the orbital frequency and

[2] Commercial software sold by Research Systems, Inc.

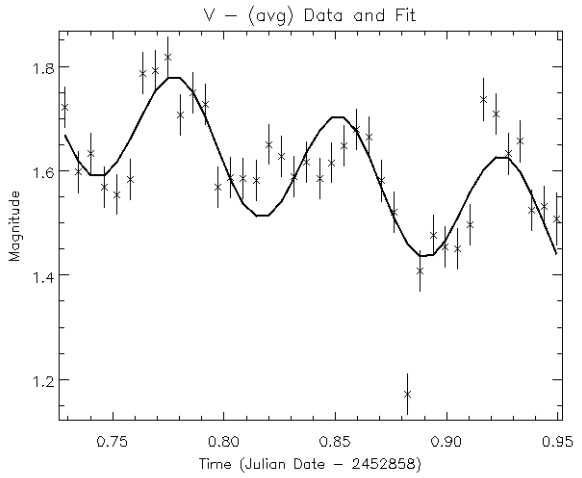

Figure 7 - Data and best-fit curve

orbital period can easily be determined. In fact, two different frequencies *were* returned – 6.5 cycles/day and 13.7 cycles/day, corresponding to 220 and 106-minute periods, respectively. The relation between these two frequencies (13.7/6.5 ≈ 2) hinted that one might well be a harmonic or subfrequency of the other. There was thus uncertainty as to which of the frequencies might be correct. Table 1 lists the final best-fit parameters with one-sigma errors.

4. CONCLUSION

V1504 Cyg was a discovered to be an SU Uma star by Nogami (1997), and Thorstensen (1997) measured its orbital period to be 100.10 ± 0.07 min. I set out to confirm the star's variability, and to detect a modulation in the star's brightness curve with a similar period.

Had V-C *not* varied significantly more than C-K, it would have meant that either V was not variable within the detection limits of the telescope, or that one of the comparison stars used was not, in fact, constant in brightness. The F-test emonstrated the variability of V1504 Cyg at a high confidence level, as expected.

Parameter	Value	Standard Error
A_0 – Amplitude	0.116 mag	± 0.009 mag
A_1 – Ang. Freq.	86 rad/day	± 1 rad/day
A_2 – Phase	4.1 rad	± 1.0 rad
A_3 – Growth Rate	-1.04 mag/day	± 0.10 mag/day
A_4 – Offset	2.47 mag	± 0.08 mag
T – Period	106 min	± 3 min
χ^2	4.1	----------

Table 1 - Best-Fit Parameters

The period measured by this experiment, 106 ± 3 min, is in reasonable agreement with the accepted orbital period. Based upon the estimated statistical and instrumental uncertainties, this is within two standard deviations of Thorstensen's value. Another method of determining how accurately the period was measured is by constructing a periodogram, as shown in Figure 7. The extremely broad peak indicates that the period is quite uncertain, probably because of the short (six hours) data set used. With such coarse resolution, it is most likely that the correct orbital period *was* detected.

This experiment could have been improved with shorter exposure times. The image series was taken eight days into the lunar cycle, a day after first quarter. If another series of images were taken nearer to the new moon, sky brightness would be much reduced, and time resolution of two or three minutes (instead of five) could be achieved. This would help in several ways. First, it would decrease the amount of noise from the already-bright Southern California sky, increasing the precision of the magnitude measurements. Second, short-exposure images could help to detect and measure any possible flickering. Third, better time resolution translates into an increased number of data points. More data would give a more accurate measure of the instrumental uncertainty, and would also lend itself to a more complete picture of V1504 Cyg's brightness curve. Lacking finer time resolution, another night or two's worth of data

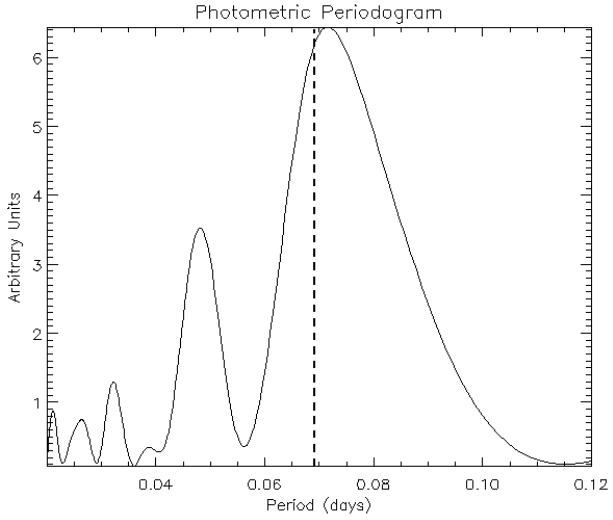

Figure 7 - Periodogram of the time series photometry; the orbital period is indicated.

would still be of assistance in charting the star's oscillations.

V1504 Cyg's variability and known orbital period were both confirmed by this experiment. This helps to pave the way for my senior thesis project, which will attempt to detect the periods of other dwarf nova systems whose characteristics are not well known.

6. ACKNOWLEDGEMENTS

I wish to thank the REU program for its support of this research, the NSF for its sponsorship of the REU program, Tammy Smecker-Hane and her research group for their assistance with research and for fielding my many questions, Tony Shoup for help with differential photometry and troubleshooting at the observatory, and Warren Skidmore for his help in choosing a target and analyzing the data I collected.

5. REFERENCES

DaCosta, G.S. 1992, "Basic Photometry Techniques," *Astronomical CCD Observing and Reduction Techniques.* Astronomical Society of the Pacific, San Francisco, 90.

Downes, R., et. al. 2003, *A Catalog and Atlas of Cataclysmic Variables*. http://icarus.stsci.edu/~downes/cvcat/

Howell, S.B., and Jacoby, G.H. 1986, *Publications of the Astronomical Society of the Pacific,* **98**, 802.

Howell, S.B., Mitchell, K.J. and Wanock, A. 1988, *Astronomical Journal*, **95**, 247.

Howell, S.B. 1989, *Publications of the Astronomical Society of the Pacific*, **101**, 616.

Massey, P. 1992, *A User's Guide to CCD Reductions With IRAF.*

Nogami, D., Masuda, S. 1997, IBVS 4532.

Skidmore, W. 1998, *A Multi-Aspect Study of WZ Sagittae and Other Dwarf Novae*. Ph.D. Thesis, University of Keele.

Sparke, L.S., and Gallagher, J.S. III. 200, *Galaxies in the Universe: An Introduction.* Cambridge University Press, Cambridge.

Taylor, J. R. 1997, *An Introduction To Error Analysis*. University Science Books, Mill Valley, CA.

Thorstensen, J. R., Taylor, C. J. 1997, *Publications of the Astronomical Society of the Pacific*, **109**, 1359.