Physical properties of IP Pegasi: an eclipsing dwarf nova with an unusually cool white dwarf

C. M. Copperwheat,[1]★ T. R. Marsh,[1] V. S. Dhillon,[2] S. P. Littlefair,[2] R. Hickman,[1] B. T. Gänsicke[1] and J. Southworth[1]

[1]*Department of Physics, University of Warwick, Coventry CV4 7AL*
[2]*Department of Physics and Astronomy, University of Sheffield, Sheffield S3 7RH*

ABSTRACT

We present high-speed photometric observations of the eclipsing dwarf nova IP Pegasi (IP Peg) taken with the triple-beam camera ULTRACAM mounted on the William Herschel Telescope. The primary eclipse in this system was observed twice in 2004, and then a further 16 times over a 3 week period in 2005. Our observations were simultaneous in the Sloan u', g' and r' bands. By phase-folding and averaging our data, we make the first significant detection of the white dwarf ingress in this system and find the phase width ϕ of the white dwarf eclipse to be 0.0935 ± 0.0003, significantly higher than the previous best value of $0.0863 < \phi < 0.0918$. The mass ratio is found to be $q = M_2/M_1 = 0.48 \pm 0.01$, consistent with previous measurements, but we find the inclination to be $83°\!.8 \pm 0°\!.5$, significantly higher than previously reported. We find the radius of the white dwarf to be $0.0063 \pm 0.0003 \, \mathrm{R}_\odot$, implying a white dwarf mass of $1.16 \pm 0.02 \, \mathrm{M}_\odot$. The donor mass is $0.55 \pm 0.02 \, \mathrm{M}_\odot$. The white dwarf temperature is more difficult to determine, since the white dwarf is seen to vary significantly in flux, even between consecutive eclipses. This is seen particularly in the u' band, and is probably the result of absorption by disc material. Our best estimate of the temperature is 10 000–15 000 K, which is much lower than would be expected for a cataclysmic variable star with this period, and implies a mean accretion rate of $<5 \times 10^{-11} \, \mathrm{M}_\odot \, \mathrm{yr}^{-1}$, more than 40 times lower than the expected rate.

Key words: stars: dwarf novae – stars: individual: IP Pegasi – white dwarfs.

1 INTRODUCTION

Cataclysmic variable stars (CVs; Warner 1995) provide examples of white dwarfs accreting from low-mass companions at rates of $\sim 10^{-11}$–$10^{-9} \, \mathrm{M}_\odot \, \mathrm{yr}$. The mass transfer in CVs and similar semidetached binary systems is driven by angular momentum loss. For long orbital period (~ 3 h or greater) systems, the angular momentum loss is thought to be driven by magnetically coupled stellar winds. For the shorter period (<3 h) systems, losses due to gravitational radiation are thought to dominate the mass transfer process. This is observationally supported to some extent: the shorter period population is dominated by lower $\dot{M}$ dwarf novae, whereas the peak of the population of high-$\dot{M}$ 'steady state' CVs is in the 3 to 4 h period range (Shafter 1992). However, there is no one-to-one relationship. The distribution of steady-state, bright CVs tails off sharply at longer periods (Ritter & Kolb 2003), and a number of dwarf novae are observed to have very long periods. It is generally thought that all non-magnetic CVs are essentially the same, with only a change in $\dot{M}$ required to move between the different subclasses. While it may be the case that the high- and low-$\dot{M}$ systems at a given orbital period are different populations with a different evolutionary history, a more likely explanation is that the instantaneous $\dot{M}$ we measure is a poor indicator of the mean $\dot{M}$, so that even high mean $\dot{M}$ systems spend some fraction of their time as slowly accreting dwarf novae.

One independent means of inferring a longer term (10^3–10^5 yr) value of $\dot{M}$ is through measurement of the white dwarf temperature. The temperature is a good tracer of the long-term $\dot{M}$ since it is determined not by the accretion heating, but by the compression of the underlying white dwarf by the accreted matter (Townsley & Bildsten 2003; Townsley & Gänsicke 2009). The existing measurements agree well with the high-$\dot{M}$/long-period, low-$\dot{M}$/short-period consensus (fig. 5 of Townsley & Gänsicke 2009). However, these data are likely to be biased. These white dwarf temperatures are spectroscopically determined, and the white dwarf must dominate

★E-mail: c.copperwheat@warwick.ac.uk

at UV wavelengths in order for the temperature to be determined. Long-period CVs are generally brighter than shorter period systems due to their higher accretion rates and more luminous donor stars, and so the white dwarfs in long-period systems must be hotter in order to be measurable. An alternative strategy which reduces this bias is to make fast, multiwavelength photometric observations of CVs where the inclination is high enough for the white dwarf to be eclipsed. With sufficient time resolution the ingress and egress of the white dwarf can be identified independently of the other components, and hence the white dwarf colours can be determined.

An ideal system for a study of this nature is IP Pegasi (IP Peg). This dwarf nova was found to be eclipsing by Goranskii et al. (1985b) and has a period of 3.8 h, in the 'long-period' tail of the dwarf nova distribution. Following Townsley & Bildsten (2003), this would imply a high mean $\dot{M}$ and a correspondingly high white dwarf temperature. However, the temperature of the white dwarf has yet to be determined with precision. Previous studies of IP Peg have shown that the white dwarf ingress feature is obscured by the ingress of the bright spot on the accretion disc, and it has been claimed that the white dwarf egress feature varies considerably in size and length (Wood & Crawford 1986; Wood et al. 1989). The difficulty in detecting the white dwarf throws the system parameters of IP Peg into doubt. In an attempt to resolve this and to measure the temperature of the white dwarf, we took a total of 18 separate observations of the IP Peg eclipse with the high-speed CCD camera ULTRACAM mounted on the William Herschel Telescope (WHT). In this paper, we present these data and our subsequent parameter determinations.

2 OBSERVATIONS

The high-speed CCD camera ULTRACAM (Dhillon et al. 2007) was mounted on the 4.2-m WHT in 2005 August. Observations were made with this instrument between August 9–15 , and August 25 to September 1. Over these periods, 16 separate observations of the IP Peg eclipse were made. In addition, two observations of the eclipse were made in 2004 August. ULTRACAM is a triple beam camera and all observations were made using the Sloan Digital Sky Survey (SDSS) u', g' and r' filters, except for the first 2004 observation, in which the i' filter was used in place of r', for scheduling reasons. Most observations were taken with an exposure time of 1.5 s, which gave a sufficient count rate in all three bands and adequately sampled the white dwarf egress, which lasts $\sim$30 s. The dead time between exposures for ULTRACAM is $\sim$25 ms. The data were unbinned and the CCD was windowed in order to achieve this exposure time. This exposure time was varied by up to 1 s in order to account for changing conditions. A complete log of the observations is given in Table 1.

All of these data were reduced with aperture photometry using the ULTRACAM pipeline software, with debiassing, flat-fielding and sky background subtraction performed in the standard way. The source flux was determined using a variable aperture (whereby the radius of the aperture is scaled according to the full width at half-maximum by a factor of 1.7). Variation in observing conditions were accounted for by dividing the source light curve by the light curve of a nearby comparison star. We use the star at $\alpha = 23^\mathrm{h} 23^\mathrm{m} 06^\mathrm{s}$, $\delta = +18°24'40''$ as the comparison. We checked the stability of this star against other stars in the field, and calculated

Table 1. Log of the observations.

Eclipse number	Cycle	Date Start	UT Start	UT End	Exposure time (s)	Phase range	Comments
1	−2200	2004 August 29	01:10	06:05	1.5	−0.60 to 0.70	Clear, seeing 0.6–0.8 arcsec i'-band filter used in place of r'
2	−2194	August 30	02:04	02:24	1.5	−0.06 to 0.02	Clear, seeing 0.6–0.8 arcsec Partial light curve: white dwarf egress only
3	−1	2005 August 12	00:41	01:17	1.5	−0.09 to 0.07	Clear, seeing 0.6–1.0 arcsec M-dwarf flare just before white dwarf egress
4	0	August 12	04:17	05:47	2.5	−0.14 to 0.24	Mostly clear, seeing 0.6 arcsec
5	6	August 13	03:06	04:00	1.5	−0.10 to 0.10	Clear, seeing 0.6 arcsec. More readout noise than usual in the r' band due to a CCD problem
6	18	August 15	00:33	01:34	1.5	−0.16 to 0.10	Clear, seeing 0.6 arcsec
7	19	August 15	04:20	05:31	1.5	−0.03 to 0.14	Clear, seeing 0.6 arcsec Partial light curve: white dwarf egress only
8	24	August 15	23:17	00:17	1.5	−0.17 to 0.08	Clear, seeing 0.6 arcsec
9	87	August 25	22:28	23:27	1.5	−0.08 to 0.08	Partial light curve due to cloud Seeing 0.6–1.0 arcsec. Dust
10	88	August 26	01:47	03:18	1.5	−0.18 to 0.09	Partial light curve due to cloud Seeing 0.6–1.0 arcsec. Dust
11	101	August 28	03:30	04:35	0.9	−0.21 to 0.07	Seeing 0.5 arcsec. Dust
12	106	August 28	22:43	23:38	1.2	−0.15 to 0.08	Seeing 1–2 arcsec. Some dust
13	107	August 29	02:24	03:21	1.2	−0.17 to 0.07	Seeing 1–2 arcsec. Some dust
14	112	August 29	21:47	22:18	2.5	−0.07 to 0.06	Seeing 1–3 arcsec. Dust Partial light curve – no ingress
15	113	August 30	01:18	02:10	1.5	−0.14 to 0.08	Seeing 1–2 arcsec. Dust
16	119	August 31	00:02	00:57	0.9	−0.16 to 0.08	Seeing $\sim$0.8 arcsec. Dust
17	120	August 31	03:50	04:43	0.9	−0.16 to 0.07	Seeing 0.5–0.6 arcsec. Much dust
18	133	September 02	04:25	06:11	1.2–2.0	−0.36 to 0.08	Seeing 1–2 arcsec. Dust

Note. Each observation is a single eclipse of IP Peg. Unless otherwise stated, the Sloan u', g' and r' filters were used. We give the phase range centred on the white dwarf egress.

apparent magnitudes of 14.45, 13.20 and 12.62 in u', g' and r', respectively.

Our short eclipse observations were interleaved with longer observations for other programmes. We used the field stars in these long data sets to determine atmospheric absorption coefficients in the u', g' and r' bands, and subsequently determined the absolute flux of our targets using observations of standard stars (from Smith et al. 2002) taken in evening twilight. We use this calibration for our determinations of the apparent magnitudes of the two sources, although we present all light curves in flux units determined using the conversion given in Smith et al. (2002). Using our absorption coefficients, we extrapolate all fluxes to an airmass of 0. For all data we convert the MJD times to the barycentric dynamical time-scale, applying light travel times to the solar system barycentre.

3 LIGHT CURVES

In Fig. 1, the 18 eclipses that were observed in 2004 and 2005 are plotted. In some cases, we only have a partial light curve, due to either poor weather or scheduling constraints. We comment on weather conditions in Table 1. In particular, note that the observations made on August 25 and later (eclipses 9 to 18) suffered from high levels of Saharan dust in the atmosphere. This dust has a grey absorption curve but the loss of flux is particularly apparent in the u'-band light curves.

These light curves show a single ingress feature, followed by separate egresses of the white dwarf, and later, the bright spot. While the egresses of these two components are distinct, it is impossible in these plots to disentangle the two ingresses. There is considerable accretion-driven flickering out of eclipse, most of which is due to the bright spot, although in some cases there is a degree of variation immediately after the white dwarf egress. There is notable variation in the strength of the white dwarf egress feature. This will be addressed in detail in Section 4.2. The 2004 observations are at a different point in the outburst cycle to the 2005 observations, and so the shape of eclipses 1 and 2 are different to the rest, due to the accretion disc flux. In addition, in eclipse 1, there is a 'hump' between the white dwarf and bright spot egresses, which may be
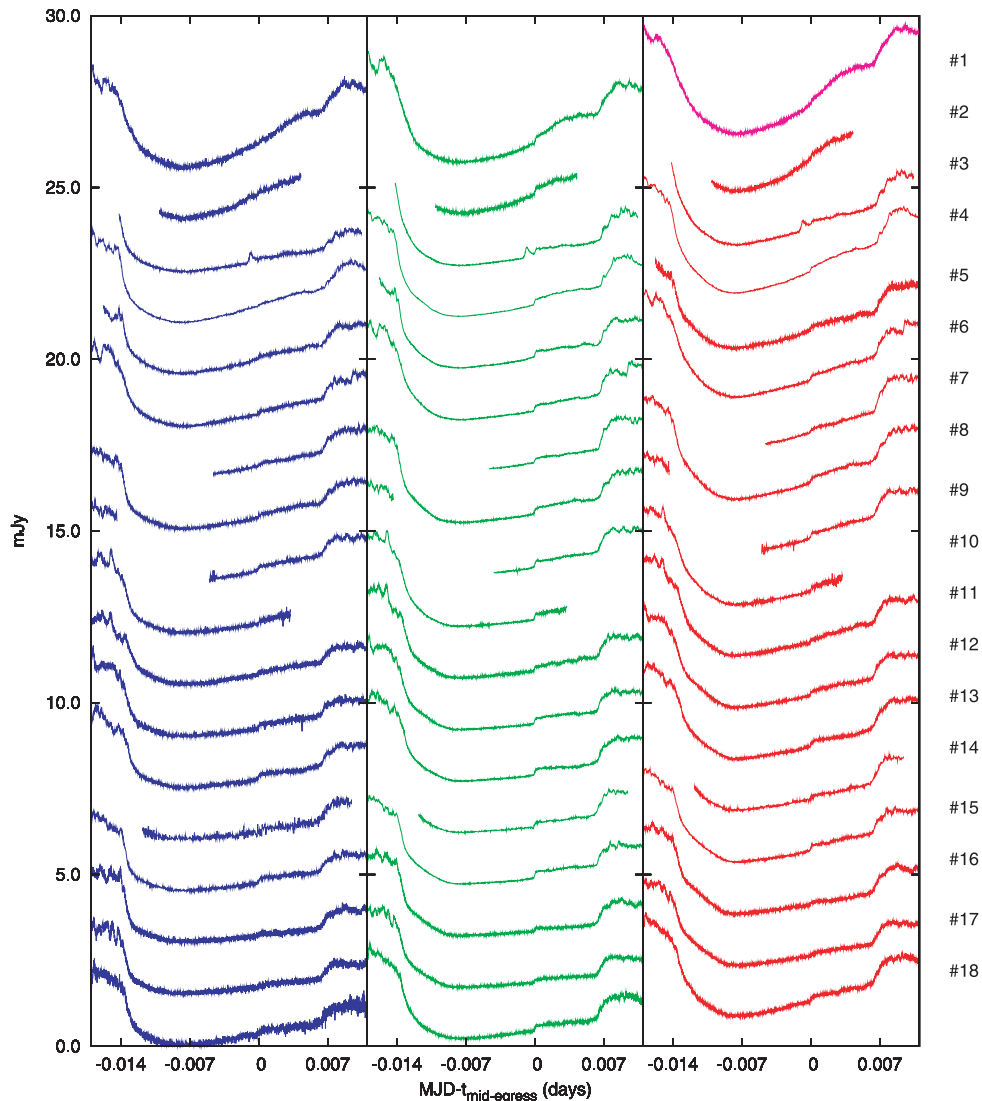
Figure 1. Light curves of IP Peg eclipses, observed in 2004 August and 2005 August. We label the eclipses according to the numbers given in the observation log (Table 1). All data were collected simultaneously in the u' g' and r' bands (blue, green and red lines, respectively) except for eclipse 1, in which an i' filter was used in place of r'. Each light curve in this plot is offset from the previous by 1.5 mJy.

due to flickering. Finally, it can be seen in eclipse number 3 that there is a brief brightening in the light curve just prior to the white dwarf egress. This brightening is most likely a flare event from the donor star, and makes it impossible to model the egress for that eclipse accurately.

4 ECLIPSE MODELLING

In this section, we report system parameters determined by fitting the data with the model described in Appendix A. This was a three stage process. We first obtained an initial fit to each individual eclipse using the simplex and Levenberg–Marquardt methods (Press 2002). These initial fits were used to obtain eclipse timings, and the resultant ephemeris was used to phase-fold all of the complete eclipse light curves. This combined light curve was used to determine the mass ratio q and inclination i of the system. We use a Markov Chain Monte Carlo (MCMC) algorithm for minimization and determination of uncertainties. These parameters were used as a basis for MCMC fits of the individual light curves, so as to determine parameters which would be expected to vary over the course of our observations, such as the accretion disc radius. Since the white dwarf egress feature appears to vary in height, and since there have been reports of variable egress duration (Wood & Crawford 1986; Wood et al. 1989), the white dwarf flux and radius were determined individually for each eclipse, as well as more refined eclipse timings which we used to determine a new ephemeris for this system.

4.1 Fitting the phase-folded data

We selected the 11 light curves in 2005 where the entire eclipse is covered (numbers 4, 5, 6, 8, 11, 12, 13, 15, 16, 17 and 18), and phase-folded them using the ephemeris given in Section 4.3. The combined and binned light curves are plotted in Fig. 2, using all 11

complete eclipses. Each bin is ∼2.7 s in length. With the flickering reduced, a white dwarf ingress feature is clearly apparent, distinct from and preceding the ingress of the bright spot. We measured the phase width of the white dwarf eclipse to be 0.0935 ± 0.0003, which is greater than the value of 0.0863 suggested as being most likely by Wood & Crawford (1986) and is also greater than the upper limit of 0.0918 given in the same paper. This implies either a higher orbital inclination or a higher mass ratio or some combination of the two.

It might be thought that this feature could be due to flickering. However, in all three bands it has a height and duration that is consistent with it being the white dwarf ingress. Additionally, by changing the number of light curves used to create the phase-folded plot, we find that the existence and size of this feature in the combined data set is not due to any one individual light curve. This is not true of other significant variations in the light curve, such as the two peaks immediately preceding the ingress at a phase range of $-0.11 < \phi < -0.10$, which are definitely due to flickering. We are therefore confident that we have detected the ingress feature.

A model consisting of a white dwarf primary, a Roche lobe filling secondary and accretion disc and bright spot components was fitted to these data using the MCMC method. A complete description of this model is given in Appendix A. In Table A1, we list the parameter determinations we make with this method, although some of these are refit on an eclipse-by-eclipse basis in Section 4.2. The uncertainties listed in this table are the formal errors, scaled so as to give a reduced χ^2 of 1. These uncertainties are underestimates, as shown by the inconsistencies between the three photometric bands. In Section 5.1, we will discuss this issue further, and calculate more realistic uncertainties which account for the systematics. The best model fits are plotted in Fig. 2. The fits are generally good, although there is some variation in the residuals out-of-eclipse due to flickering which has not been completely averaged out in the combined light curve. Additionally, our fit to
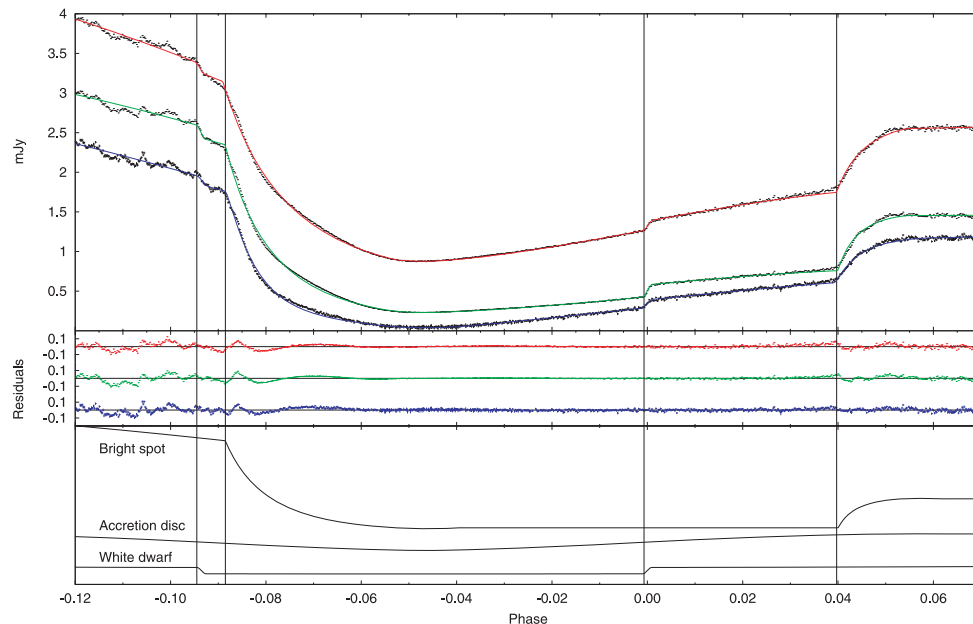
Figure 2. Top panel: phase-folded and binned light curves (top, r'; middle, g'; bottom, u'), comprised of 11 complete eclipse observations made in 2004 and 2005. We plot the average flux in mJy against the binary phase, where a phase of 0 corresponds to the mid-egress of the white dwarf. We plot the data points with uncertainties in black, and the best model fits to these data in red, green and blue (for r', g' and u', respectively). The four vertical lines indicate the start of (from left to right) the white dwarf ingress, the bright spot ingress, the white dwarf egress and the bright spot egress. Middle panel: residuals in the three bands (top, r'; middle, g'; bottom, u'). Bottom panel: the three components of the g'-band model plotted separately, showing the relative strengths of the bright spot, accretion disc and white dwarf. These three lines have been offset for clarity.

the bright spot ingress and egress is not perfect; we discuss this in Appendix A4.

One additional parameter which our results are sensitive to is the limb-darkening coefficient for the white dwarf. An initial fit to the data was made with the Levenberg–Marquardt method in order to obtain approximate white dwarf parameters, and these determinations were used with the white dwarf model atmospheres of Gänsicke, Beuermann & de Martino (1995) to determine the coefficients in the three ULTRACAM bands. We found linear coefficients of 0.352, 0.253, 0.216 in u', g' and r' to be appropriate for the white dwarf in IP Peg. These values were used in all our MCMC fits.

4.2 Fitting the individual eclipses

For some parameters, the results from the phase-folded light curves are insufficient. Some variations in the white dwarf egress feature are clearly apparent in the individual eclipses, so it was necessary to fit these separately in order to obtain a proper determination of the white dwarf radius and fluxes. Obviously, this was also necessary in order to determine the eclipse timings, and some variation in the accretion disc radius over the course of our observations would also be expected. We used our model fit to the phase-folded light curve as a basis for fitting the individual eclipses separately. It is clear from Fig. 1 that the amplitude of the flickering is much larger than the white dwarf eclipse features in these individual light curves, so we fixed most of the parameters, including the mass ratio and inclination angle, to the values determined from the phase-folded light curve. We allowed the white dwarf radius and the time of the white dwarf mid-eclipse to vary. We also wanted to track variations in the disc radius, but in these individual light curves the outer disc radius R_{disc} is poorly constrained. We therefore set $R_{disc} = R_{spot}$ and fitted R_{spot}, where R_{spot} is the distance between the white dwarf and the maximum of the surface brightness of the bright spot. This parameter is well constrained by the data, since the bright spot is so strong in flux. Note that in the previous section (the phase-folded light curves) the data are of sufficient quality that we could fit R_{disc} and R_{spot} separately, and we found R_{spot} to track the accretion disc radius closely. To allow for a degree of flexibility in modelling a variable bright spot, we also allow l (the scalelength of the bright

spot) and f_c (the fraction of bright spot flux that is taken to be constant) to vary. The remaining bright spot parameters were fixed to the values determined from the phase-folded light curve. Finally, the surface brightness of the accretion disc is modelled as a power law, and the exponent of this power law is also allowed to vary. A complete description of all parameters is given in Appendix A. Each individual light curve was fitted in the same way as the phase-folded light curve (as described in Section 4.1).

The results for the g'-band fits (for which the signal to noise is highest) are listed in Table 2. We report the white dwarf radii and timings for 17 of the 18 eclipses, omitting the eclipse with the flare near the egress feature. There is a considerable degree of scatter in our model determinations of the white dwarf radius. However, given the uncertainties, the variations are not significant. The timings of the white dwarf mid-eclipse were converted to timings of the mid-point of the egress feature so as to be consistent with previous authors, using a phase width of 0.0935. We list also our determinations of the accretion disc radius R_{disc} scaled by the binary separation a. This parameter is given for 15 of the 18 eclipses, omitting eclipses which are not complete enough for the disc radius to be determined with any precision. Also listed in Table 2 are the white dwarf fluxes in all three bands. These were calculated by measuring the height of the white dwarf egress feature in the model fits and then converted back to a magnitude scale following Smith et al. (2002). In Fig. 3, the u'-, g'- and r'-band light curves of the region around the white dwarf egress are plotted, with the accretion disc component of the emission subtracted from the data. This plot shows clearly that the height of the egress feature varies significantly from eclipse to eclipse (e.g. compare egresses 15 and 16, taken on consecutive nights). This variability is apparent in all three bands, although in r' the egress feature is very small and in u' the signal-to-noise ratio is much lower, making it harder to determine the size of the egress accurately.

4.3 Eclipse timings

We used the timings of the white dwarf mid-egress given in Table 2 in order to calculate an updated orbital ephemeris for IP Peg. These

Table 2. White dwarf parameters obtained from MCMC fits to the individual light curves, as detailed in Section 4.2.

Eclipse	Cycle	$t_{mid-egress}$	R_1/a	R_{disc}/a	u' mag	g' mag	r' mag
1	−2200	53 246.145 370 (10)	0.0037 ± 0.0003	0.262 ± 0.002		18.83 ± 0.02	
2	−2194	53 247.094 620 (12)	0.0043 ± 0.0002		19.08 ± 0.08	18.53 ± 0.02	19.57 ± 0.06
4	0	53 594.198 840 (8)	0.0031 ± 0.0008	0.279 ± 0.002	21.29 ± 0.18	18.64 ± 0.01	19.21 ± 0.02
5	6	53 595.148 038 (5)	0.0044 ± 0.0011	0.269 ± 0.001	19.50 ± 0.04	18.51 ± 0.01	19.19 ± 0.05
6	18	53 597.046 564 (5)	0.0052 ± 0.0008	0.264 ± 0.002	19.51 ± 0.05	18.38 ± 0.01	18.64 ± 0.01
7	19	53 597.204 771 (6)	0.0062 ± 0.0008	0.271 ± 0.002	19.14 ± 0.03	18.62 ± 0.01	18.89 ± 0.01
8	24	53 597.995 814 (5)	0.0060 ± 0.0017	0.270 ± 0.001	19.01 ± 0.03	18.36 ± 0.01	18.61 ± 0.01
9	87	53 607.962 749 (5)	0.0044 ± 0.0005	0.268 ± 0.003	18.73 ± 0.03	18.29 ± 0.01	18.54 ± 0.01
10	88	53 608.120 980 (8)	0.0040 ± 0.0006		19.15 ± 0.06	18.36 ± 0.01	18.84 ± 0.02
11	101	53 610.177 655 (4)	0.0039 ± 0.0007	0.253 ± 0.001	19.63 ± 0.07	18.74 ± 0.01	18.23 ± 0.02
12	106	53 610.968 692 (5)	0.0043 ± 0.0005	0.257 ± 0.001	18.91 ± 0.04	18.45 ± 0.01	18.87 ± 0.01
13	107	53 611.126 905 (5)	0.0050 ± 0.0008	0.254 ± 0.001	18.47 ± 0.03	18.20 ± 0.01	18.41 ± 0.01
14	112	53 611.917 924 (2)	0.0056 ± 0.0010	0.250 ± 0.002	19.15 ± 0.17	18.48 ± 0.02	18.94 ± 0.05
15	113	53 612.076 134 (4)	0.0044 ± 0.0005	0.258 ± 0.001	18.64 ± 0.02	18.25 ± 0.01	18.65 ± 0.01
16	119	53 613.025 377 (7)	0.0043 ± 0.0006	0.256 ± 0.001	19.08 ± 0.05	18.88 ± 0.01	19.23 ± 0.02
17	120	53 613.183 571 (7)	0.0035 ± 0.0007	0.260 ± 0.001	19.57 ± 0.09	19.07 ± 0.01	19.56 ± 0.03
18	133	53 615.240 256 (10)	0.0054 ± 0.0010	0.254 ± 0.002	18.86 ± 0.07	18.45 ± 0.01	18.65 ± 0.02

Note. We give the time of the white dwarf mid-egress $t_{mid-egress}$ and the white dwarf R_1 and accretion disc radii (R_1 and R_{disc}) scaled by the binary separation a, as obtained from the g'-band fits. $R_{disc} = R_{spot}$ in these model fits. We list also the u'-, g'- and r'-band magnitudes of the white dwarf, determined from the step height of the white dwarf egress feature.
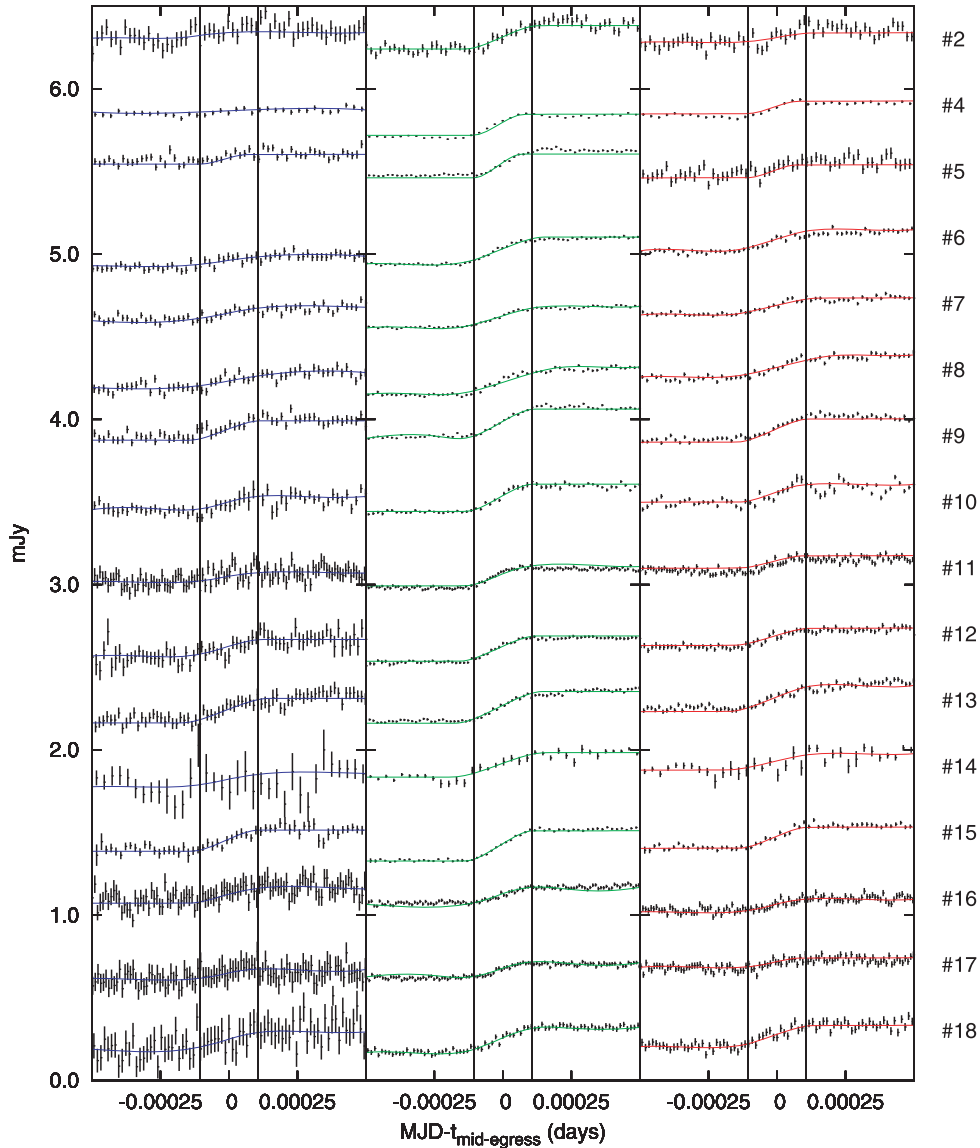
Figure 3. The region of the light curves around the white dwarf egress for 16 of the 18 eclipses. We label the eclipses according to the numbers given in the observation log (Table 1). We plot the *u'*-band data in the first column, the *g'*-band data in the second and the *r'*-band data in the third. For the *x*-axis timings, we used the ephemeris given in Section 4.3. The dashed vertical lines identify the mean phase width of the white dwarf egress.

data were complemented with eclipse timings given by previous authors (Goranskii, Lyutyi & Shugarov 1985a; Wood et al. 1989; Wolf et al. 1993). All times were converted to the barycentric dynamical time-scale, corrected for light travel to the barycentre. A least-squares fit to all of these data yielded the ephemeris

$$BMJD(TDB) = 53594.206270(1) + 0.1582061029(3)E$$

for the white dwarf mid-egress. This ephemeris fits well with our 2004/2005 data (Fig. 4), but the historical data still show the $\sim$100 s deviation from a linear ephemeris discussed by Wolf et al. (1993), which they argued was caused by a third body in the system. An alternative explanation is that some of the historical timings are unreliable: the flux variations we observe in the egress feature (Fig. 3) suggest that some observations with small telescopes may have failed to detect the egress feature at all. This is particularly likely in the instances when a very long egress was reported (Wood & Crawford 1986), for which we see no evidence.
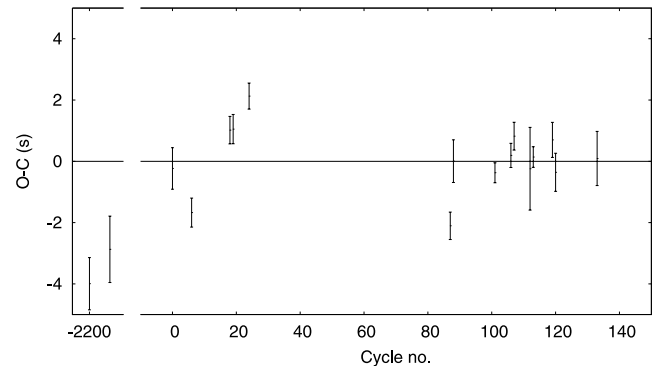
Figure 4. $(O - C)$ values plotted against cycle number, using the eclipse timings listed in Table 2 and the linear ephemeris given in Section 4.3. We plot only our 2004/2005 data points, but used all available historical data in the determination of the ephemeris. Note the break in the *x*-axis of this plot between the 2004 and 2005 data.

5 DISCUSSION

In this section, we discuss the results of our model fits in more detail. The photometric method for parameter determination in eclipsing CVs is dependent on precise measurement of the contact phases of the white dwarf and bright spot eclipses. These measurements are used to constrain the geometry of the binary system, and derive the mass ratio and inclination. Our data affords the possibility of more precise determinations than previous authors, since we have for the first time identified the white dwarf ingress, and so can accurately determine the phase width of the white dwarf eclipse. We begin in Section 5.1 by discussing our determinations of the mass ratio and binary inclination, and the uncertainties on these measurements. We compare our results to determinations made both photometrically and spectroscopically by previous authors. In Section 5.2, we go on to investigate the individual components, and give masses, radii and radial velocity semi-amplitudes for both the white dwarf and the donor star. In Section 5.3, we use these new parameter determinations to provide an updated estimate of the distance to IP Peg. In Section 5.4, we use our determinations of the white dwarf colours to determine its temperature, by comparing our measurements to theoretical cooling models. This task is complicated by the variability we observe in the white dwarf egress feature. In Section 5.5, we detail a photoelectric absorption model which could account for this variability. As discussed in Townsley & Bildsten (2003), the temperature of the white dwarf is driven by compressional heating of accreted matter, and so in Section 5.6 we discuss the implications of our temperature measurement for the accretion history of this system. Finally, in Sections 5.7 and 5.8, we examine the bright spot and accretion disc in more detail. In the case of the bright spot, we determine its temperature and flux distribution, and for the accretion disc we examine changes in disc flux and radius over the course of our observations.

5.1 Mass ratio and inclination

The phase width of the white dwarf eclipse is an observable quantity that is intrinsically linked to two physical properties: the mass ratio and the binary inclination. For a higher binary inclination the duration of the eclipse will be greater, thus to maintain the same phase width, as the inclination is increased, the size of the donor, and hence the mass ratio, must be decreased. There is therefore a unique relationship between these two properties (Bailey 1979). This degeneracy can be broken since we have an additional geometric constraint due to the ingress and egress of the bright spot. The path of the accretion stream and hence the position of the bright spot is modified by the mass ratio. With this additional information, we can determine the mass ratio and inclination in this system.

Previous studies of this system did not identify the white dwarf ingress feature, and so there has been some uncertainty in the phase width $\Delta\phi$ of the white dwarf eclipse. Wood & Crawford (1986) suggested that the white dwarf ingress was blended with the bright spot ingress, implying $0.0863 < \Delta\phi < 0.0918$. We find the phase width to be 0.0935 ± 0.0003, implying either a higher mass ratio or a higher inclination than previously suggested, or some combination of the two. We find a much better fit for the high-inclination case, due to the constraints implied by the bright spot.

Our MCMC results for the mass ratio versus inclination are plotted in Fig. 5. The coloured contours and associated points show the single parameter $1\,\sigma$ errors and 95 per cent joint confidence contours. The other points plotted are various q and i determinations taken from the literature (Wood & Crawford 1986; Marsh
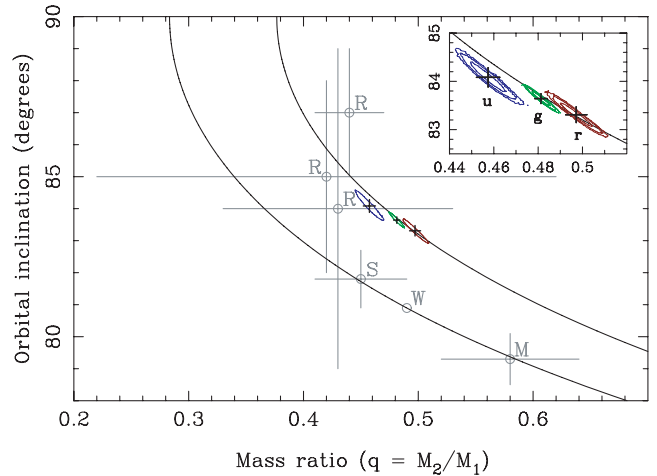

Figure 5. Mass ratio versus inclination, from our MCMC fits. We show the results of our fits to the phase-folded light curve, plotting the results in u', g' and r' separately (blue, green and red, respectively). We plot points with the single parameter $1\,\sigma$ errors and 95 per cent joint confidence contours. The two solid lines running diagonally across the plot indicate the q/i relationship for a phase width of 0.0863 (bottom) and 0.0935 (top). The inset shows an expanded view and the 68, 95 and 99 per cent joint confidence contours. The grey points are measurements of q and i from the literature, and are labelled according to source as follows: *W*, Wood & Crawford (1986); *M*, Marsh (1988); *S*, Smak (2002); *R*, Ribeiro et al. (2007).

1988; Smak 2002; Ribeiro et al. 2007). The two black lines running across the plot show the unique q/i relationship for phase widths of 0.0863 and 0.0935. Wood & Crawford (1986) suggested 0.0863 was the most likely value for the phase width: a number of authors have assumed this value and hence underestimated the inclination.

If we now look at the values from our MCMC fits as plotted in Fig. 5, it can be seen that the results from the different filters are not consistent at the 95 per cent level. There are systematic uncertainties in our parameter determinations which dominate the uncertainties implied by the individual fits. We believe this to be due to the uncertainty in the measurement of the phase widths. While the position of the white dwarf egress is well defined, the white dwarf ingress, and to some extent, the bright spot ingress and egress, are still affected by flickering, even in the phase-folded light curve. In order to quantify the uncertainty in our system parameters, it was necessary to estimate the range of variation in the phase width as caused by the flickering. Our phase-folded light curve is the average of the eleven complete light curves. If the position of the white dwarf ingress is affected by flickering, it follows that if we choose a subset of these 11 to make a new phase-folded light curve, then we will measure a different phase width. It is likely that in some light curves the flickering influences the apparent position of the ingress more than in others, and so if the phase width is measured in enough of these subsets, the uncertainty in the phase width can be determined, from which a more realistic uncertainty for the physical parameters can be calculated.

We created a series of phase-folded light curves from different combinations of the complete light curves. Each phase-folded light curve was made up of 11 light curves (like the original), but these 11 were selected using the bootstrap method (Efron 1979; Efron & Tibshirani 1993). For each of these a minimization was performed using the MCMC method, as before, and the contact phases of the ingress and egress features were computed. We additionally calculated the phase width of the white dwarf egress itself, using our
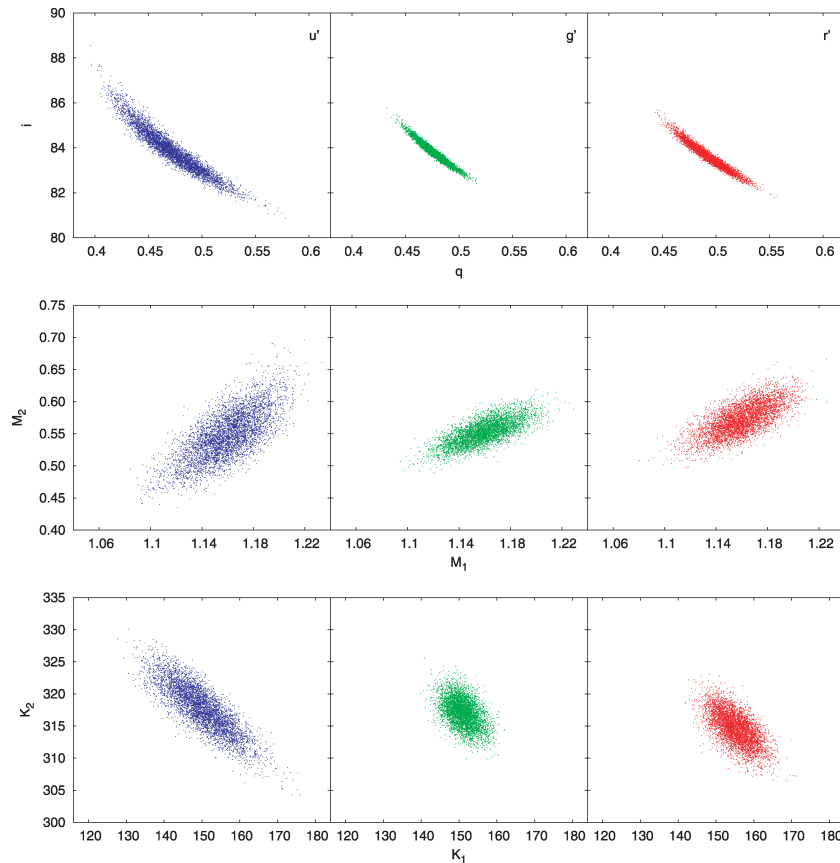
Figure 6. The results of our Monte Carlo simulation, which we use to determine the parameter uncertainties listed in Table 3. We plot the mass ratio q versus the inclination i, the white dwarf mass M_1 versus the donor mass M_2, and the white dwarf radial velocity semi-amplitude K_1 versus the donor semi-amplitude K_2. The three columns separate out the results from the individual bands (left to right, u', g' and r').

Table 3. Binary parameters for IP Peg.

		u' band	g' band	r' band
Phase width		0.0933 ± 0.0006	0.0935 ± 0.0003	0.0941 ± 0.0003
q		0.47 ± 0.03	0.48 ± 0.01	0.49 ± 0.02
i	($°$)	83.89 ± 0.96	83.81 ± 0.45	83.61 ± 0.53
a	($R_\odot$)	1.47 ± 0.02	1.472 ± 0.009	1.48 ± 0.01
M_1	($M_\odot$)	1.16 ± 0.02	1.16 ± 0.02	1.16 ± 0.02
R_1	($R_\odot$)	0.0081 ± 0.0013	0.0063 ± 0.0003	0.0064 ± 0.0004
K_1	(km s^{-1})	150 ± 7	151 ± 3	155 ± 4
M_2	($M_\odot$)	0.55 ± 0.04	0.55 ± 0.02	0.57 ± 0.02
R_2	($R_\odot$)	0.46 ± 0.02	0.47 ± 0.01	0.47 ± 0.01
K_2	(km s^{-1})	318 ± 4	317 ± 2	315 ± 3

Note. We list the phase width of the white dwarf eclipse, the mass ratio q, the inclination i and masses, radii and radial velocity semi-amplitudes of the two components. All of these parameters except the radii are determined from the Monte Carlo simulations described in Section 5.1. The white dwarf radius is derived from the weighted mean of the values determined in the individual MCMC fits, and the donor radius is calculated using the approximation of Eggleton (1983).

previously determined q and i values, along with the weighted mean and variance of R_1/a calculated from the values listed in Table 2. We subsequently calculated the covariance of the contact phases. These covariances are input into a Monte Carlo simulation through which we determined the uncertainty on the physical parameters of the system: q, i, the component masses M_1 and M_2, the binary separation a and the radial velocity semi-amplitudes K_1 and K_2. One additional input into this calculation is an estimation of the extent to which the size of the white dwarf exceeds the Eggleton zero-

temperature mass/radius relation (quoted in Verbunt & Rappaport 1988). For this the white dwarf temperature estimate from Section 5.4 is used with the white dwarf cooling models of Bergeron, Wesemael & Beauchamp (1995) and Holberg & Bergeron (2006). We find an oversize factor of $\sim$1.04 to be appropriate. The results of this simulation are plotted in Fig. 6, and all of these parameters are listed in Table 3. Unlike the earlier MCMC fits, the uncertainties on these determinations are more realistic values which account for the systematic errors due to the uncertainty in the phase widths, and

so the determinations are consistent in all three bands. The g'-band determinations are of course the most precise, due to the higher signal-to-noise ratio in these data.

We find the mass ratio in the g'-band data to be $q = 0.48 \pm 0.01$. This is consistent with, although more precise than, many of the previous determinations: the photometric studies of Wood et al. (1989) for example suggested a value of 0.49, and Smak (2002) reported $q = 0.45 \pm 0.04$. Wolf et al. (1993) suggested a value of 0.6 which is much higher than our finding, although previous photometric studies were limited by the difficulty in identifying the white dwarf ingress feature in the light curves. However, the spectroscopic study of Marsh (1988) also reported a higher value for q (0.58 ± 0.06).

The constraints on the binary inclination were reported by Wood & Crawford (1986) to be $80°.9 < i < 90°$, and previous authors have suggested values towards the low end of that range (e.g. Smak 2002, $i = 81°.8 \pm 0°.9$). Our g'-band value of $83°.81 \pm 0°.45$ is higher than most previous determinations. As we remarked earlier, a confounding factor in previous studies has been the difficulty in identifying the white dwarf ingress feature, and the subsequent underestimation of the phase width of the white dwarf's eclipse.

5.2 The white dwarf and donor star parameters

In Table 3, we list our determinations of the masses, radii and radial velocity semi-amplitudes of the two components, as determined from our Monte Carlo simulation. The white dwarf radius R_1 is calculated from the weighted mean of the R_1/a values listed in Table 2, and the donor radius R_2 is calculated using the Roche-lobe volume radius approximation of Eggleton (1983).

If the g'-band results are used, the white dwarf mass is found to be $1.16 \pm 0.02 \, M_\odot$. This is larger than some previous determinations, e.g. Smak (2002) reported a value of $0.94 \pm 0.10 \, M_\odot$, but is in very close agreement with Watson et al. (2003), who report a value of 1.16–$1.18 \, M_\odot$. We find the donor mass to be $0.55 \pm 0.02 \, M_\odot$, which is greater than previous measurements. Smak (2002) reported a value of $0.42 \pm 0.08 \, M_\odot$, and Watson et al. (2003) found M_2 to be $0.5 \, M_\odot$. Our donor radius of $0.47 \pm 0.01 \, R_\odot$ is consistent with previous authors, but it is low given our mass determination, suggesting a donor that is slightly undersized compared to an equivalent main sequence star (Baraffe et al. 1998). This is unusual for an accreting binary: donor stars tend to be oversized, due to either evolution or being out of thermal equilibrium as a result of mass loss. Our finding is consistent with a donor in thermal equilibrium. As we will show in Section 5.6, we find the accretion rate to be much lower than would be expected for this system, and this might explain why the donor is not significantly perturbed. If the donor is in thermal equilibrium, then this suggests the mass transfer rate has been low for a very long period of time, in excess of the thermal time-scale ($\sim 10^8$ yr, for a star with these parameters). Alternatively, the mass transfer might have begun very recently.

Recent work by Knigge (2006) has suggested a single, semi-empirical donor sequence for CVs with orbital periods <6 h. This sequence suggests a donor mass of $0.28 \, M_\odot$ and radius of $0.38 \, R_\odot$ for a $P = 3.8$ h CV. These values are much lower than our measurements; for our donor mass of $0.55 \, M_\odot$ to be consistent with the sequence, the binary period would need to be ~ 5.3 h. IP Peg is not the only system that fits poorly to the semi-empirical donor sequence of Knigge (2006); we see in fig. 9 of that paper that there are a number of other systems, such as RX J0944, which have a donor that is of an earlier spectral type than would be expected for their binary period. Again, the low accretion rate we will report in

Section 5.6 may imply an unusual accretion history for this system, which could account for this discrepancy.

Using our binary parameters, we calculated the radial velocity semi-amplitudes K_1 and K_2, and find them to be 151 ± 3 and 317 ± 2 km s^{-1}, respectively. Our value of K_2 is larger than the long-established value (298 ± 8 km s^{-1}; Martin, Jones & Smith 1987; Martin et al. 1989), which was more recently supported by Beekman et al. (2000). However, the much better quality data of Watson et al. (2003) give a value of 323.2 ± 3.5 km s^{-1}. The consistency of our value derived purely from photometry with the best spectroscopic measurement is a rare test of the consistency of the photometric method (see also Littlefair et al. 2008; Tulloch, Rodríguez-Gil & Dhillon 2009). The systematic uncertainty in spectroscopic determinations of K_1 is much higher than in K_2, and so previous values of K_1 in the literature as determined from the emission lines vary significantly (175 ± 15 km s^{-1}, Marsh 1988; 118 ± 10 km s^{-1}, Hessman 1989). More recently, Smak (2002) used the three-body approximation to find a value for K_1 of 134 ± 15 km s^{-1}. This is consistent with our finding of 151 ± 3 km s^{-1}.

5.3 Distance to IP Peg

Using our new parameters for the donor star, we can update the estimation of the distance to IP Peg. The surface brightness method of distance determination (Bailey 1981) uses the K-band magnitude to determine the distance to CVs. This method assumes all of the K-band flux originates in the donor star and is thus a lower limit, since the accretion disc will contribute an indeterminate amount of this flux (although this is small in the case of IP Peg). Szkody & Mateo (1986) measured the K flux and estimated the distance to be 130–142 pc. Bailey (1981) assumed a constant donor surface brightness in his relation, but subsequent authors have shown this quantity varies with the $V - K$ colour (Ramseyer 1994; Beuermann 2006). Froning et al. (1999) reported an updated value of 121 pc using this refinement of Ramseyer (1994) and the K-band measurements of Szkody & Mateo (1986). We applied the Beuermann (2006) update of Bailey's method, and found the distance to IP Peg to be 151pc. For this, we used the Szkody & Mateo (1986) K-band measurement, the donor radius listed in Table 3 and the same $V - K$ that was used by Froning et al. (1999), which is consistent with a donor spectral type of $\sim$M4.8. The uncertainty in our distance determination is likely to be dominated by the uncertainty in $V - K$. We took ± 1 to be an appropriate estimate of the uncertainty in the spectral type, and determined the subsequent range in $V - K$ using Table 3 of Beuermann (2006). We hence estimate the uncertainty in our distance determination of 151 pc to be ± 14 pc. The best trigonometric estimate of the distance is 152^{+44}_{-29} pc (Thorstensen, private communication), which is consistent with our result.

5.4 White dwarf colours and temperature

The main motivation for this work was a photometric determination of the white dwarf temperature in this system. We see in Fig. 3 that, while the length of the white dwarf egress appears constant, its flux varies significantly on time-scales of days. Using the white dwarf magnitudes listed in Table 2, we calculate the $u' - g'$ and $g' - r'$ colours of the white dwarf for each individual eclipse. We plot these data in the left-hand panel of Fig. 7. We plot also the synthetic white dwarf tracks of Holberg & Bergeron (2006) for a white dwarf mass range of 0.3 to $1.2 \, M_\odot$.
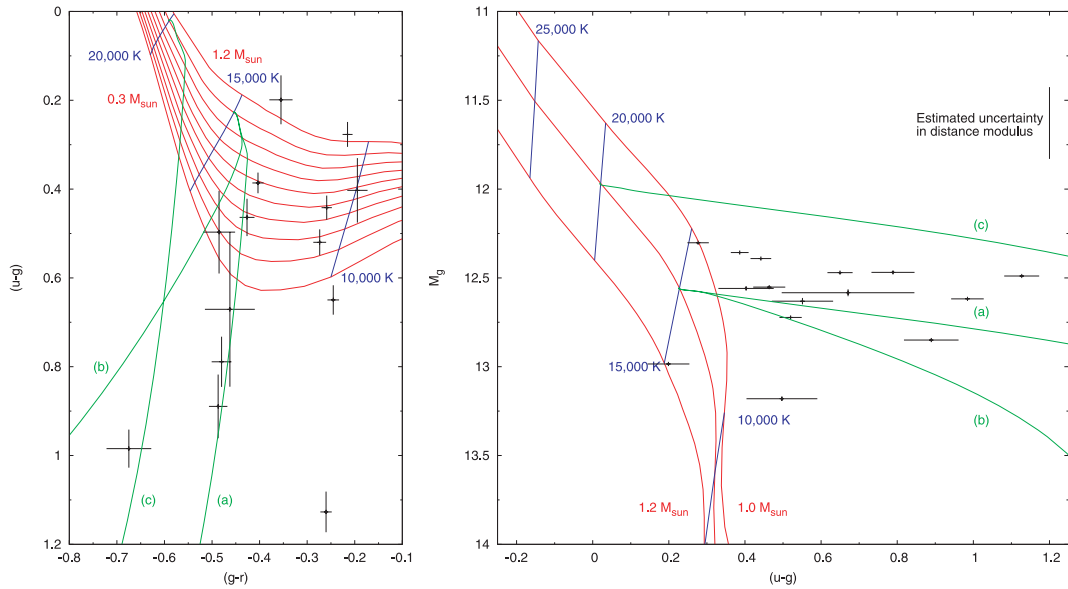
Figure 7. Using the eclipse determinations listed in Table 2, we plot data points showing the colours and magnitude of the white dwarf for each individual eclipse. In the left-hand panel, we plot the $u' - g'$ and $g' - r'$ colours for the white dwarf, and in the right-hand panel we show the $u' - g'$ colour versus the absolute magnitude M'_g. For this, we use our distance to IP Peg of 151pc (Section 5.3), but we indicate to the right of the plot our estimate of the uncertainty in the distance modulus. The red lines that run from the top left area of each panel are the synthetic white dwarf tracks of Holberg & Bergeron (2006). In the left-hand panel, we plot tracks for a white dwarf mass of 0.3 to 1.2 $M_\odot$, plotted every 0.1 $M_\odot$, but for clarity in the right-hand panel we just show the tracks for a mass of 1.0, 1.1 and 1.2 $M_\odot$. Perpendicular to these tracks we plot isotherms (blue) every 5000 K. The green lines are representative of the results of our opacity model, discussed in more detail in Section 5.5. Line (*a*) assumes a white dwarf with a temperature of 15 000 K and a mass of 1.1 $M_\odot$, and absorbing material with a temperature of 8000 K and $N_H < 10^{23}$ cm^{-2}. Line (*b*) uses the same values as line (*a*), except for the absorber temperature which is assumed to be 14 000 K. Line (*c*) uses the same values as line (*a*), except for the white dwarf temperature which is assumed to be 20 000 K.

This plot shows that as well as variation in all three bands, there are considerable variations in the white dwarf colours, particularly $u' - g'$. There is a suggestion of an inverse correlation between the $u' - g'$ and $g' - r'$ colours. This is due to large variation in the u' and r' fluxes and comparatively little change in the g' flux, so things become a little clearer when we plot $u' - g'$ against the absolute g'-band magnitude M'_g (Fig. 7, right-hand panel). We assume a distance of 151 pc (Section 5.3), but plot the error bar corresponding to our estimate of the uncertainty in this determination. Despite this uncertainty, we see in Fig. 7 a clear trend in the data. The data points with the lowest $u' - g'$ colours intersect with the synthetic tracks, and there are many other points which are much redder in colour. Not only are many of these points too red to correspond to a white dwarf of mass $\sim$1.1 $M_\odot$, they are too red to correspond to a white dwarf of any mass.

We consider two possibilities for the variation in the white dwarf colours. First, it is possible that there is some variable contamination from an additional temperature component, such as a band of hot, accreted material around the equator of the white dwarf (Piro & Bildsten 2004). If this was the case, one expects the trend in the egress data points to be approximately in line with the direction of increasing white dwarf temperature in the Holberg & Bergeron (2006) tracks. However, what we actually see is a trend that is roughly perpendicular to the direction of increasing white dwarf temperature. The inverse correlation between the $u' - g'$ and $g' - r'$ colours is due to the fact that there is comparatively little change in the g' flux compared to the (correlated) change in the u' and r' fluxes. It is difficult to reconcile this with a model that includes an additional temperature component. It is more likely that the variation in the white dwarf colours is due to photoelectric absorption, as a result of obscuration of the white dwarf by material from the accretion

disc. This model can approximately explain the colour variations, and we examine this in more detail in Section 5.5. An occulting region of disc material could explain past reports of very extended white dwarf egresses in this system (Wood & Crawford 1986). Alternatively, these might have been failures to detect the variable egress at all. Given the short period of time over which we observed this system, it is difficult to quantify the extent of this phenomenon.

We now attempt to estimate the white dwarf temperature in this system via comparison with the synthetic tracks. Fig. 7 suggests a temperature of between 10 000 to 15 000 K; it is difficult to be more precise than this due to the flux variations. Even so, this range is remarkably cool. Previous studies have suggested that a white dwarf in a CV with a period >3 h would have a typical temperature of between 25 000–35 000 K (Townsley & Bildsten 2003; Winter & Sion 2003). Such a temperature is clearly inconsistent with our IP Peg data: this would require the white dwarf to be more than a magnitude more luminous in g' and more than 1.5 mag brighter in u'. In order to plot the absolute magnitudes used in the right-hand panel of Fig. 7, we assumed a distance to IP Peg and the uncertainty in this measurement could allow a higher temperature, but judging from the parallax measurement this uncertainty is unlikely to be sufficient to reconcile a 25 000 K white dwarf with the observations. In addition, the $u' - g'/g' - r'$ plot is independent of distance and also implies a temperature range of 10 000 to 15 000 K for the white dwarf.

In making this temperature determination, we have assumed that our points with the bluest $u' - g'$ colour are indicative of a white dwarf that is largely unobscured. However, we also investigate the possibility that the accretion disc is opaque, blocking the bottom half of the white dwarf from view and so producing a situation where the obscuration of the white dwarf is at least 50 per cent at all times.

With this modification, we find the white dwarf temperature implied by our model fits is at most $\sim$21 000 K. This is higher than our previous findings, but still significantly smaller than the predicted 25 000–35 000 K. Additionally, we will show in Section 5.5 that the colour trends in the data points can be explained by an absorption model if we assume a 15 000 K white dwarf, but not if we assume a temperature of 20 000 K or higher; hence, we prefer the fits which imply a cooler white dwarf.

5.5 An absorption model

We propose that the variation in the white dwarf colours is due to occultation of the white dwarf by disc material, causing a significant and varying photoelectric absorption of the white dwarf flux. To test this, we computed an opacity model and compared it to the data. Our intention here is not to make a precise and quantitative description of the absorption in this system, but to determine if a varying opacity is an adequate description for the flux variations we observe.

We calculated the opacity for a range of absorber temperatures and neutral hydrogen column densities, allowing for hydrogen bound-free, free–free, bound–bound and H^- opacities. A white dwarf spectrum was used to determine the source flux and we assumed local thermodynamic equilibrium, an electron density of 10^{14} cm^{-3}, and we approximated the line broadening by assuming it to be Gaussian with an rms of 50 km s^{-1}. We convolved the resultant spectra with the SDSS filter profiles and hence calculated the change in the u', g' and r' magnitudes for each absorber temperature and column density. These calculations give us the change in the source flux; we chose the unabsorbed colours to be those of a 15 000 K, 1.1 M$_\odot$ white dwarf, which is appropriate given our earlier measurements. The results of this model calculation are plotted in Fig. 7 as tracks (*a*) and (*b*). These two tracks denote a constant absorber temperature of 8000 and 14 000 K, respectively. The hydrogen column density increases along these tracks. If we compare the tracks to the white dwarf measurements, we see that they lie approximately in line with the trend we observe in the data points, suggesting that absorption is an adequate explanation for the variations in white dwarf flux. There are some discrepancies; in particular, a number of data points at low temperatures are not consistent with the tracks we plot, but we attribute this to the simplicity of our model.

It was noted in Section 5.4 that it is possible that in all of our observations we see only a small fraction of the white dwarf flux, which would make our determination of the white dwarf temperature a significant underestimation. To test this, we plot a third model track in Fig. 7, labelled (*c*). This track uses the same parameters as track (*a*) with the exception of the white dwarf temperature, which is assumed to be 20 000 K. We see this track is a poor fit to our observations, and so a white dwarf temperature of 20 000 K or greater is not supported by our opacity model. In addition, if the temperature of the white dwarf is greater than $\sim$17, 000 K, we find that there is no adjustment in the temperature of the absorbing material that can lead to the redward movement in $g' - r'$ that would be required to fit the data.

Based on the neutral hydrogen column densities we use in our model, we calculate an N_H of up to $\sim$10^{23} cm^{-2} is necessary to account for the variation we observe in the white dwarf flux. For comparison, Horne et al. (1994) reported absorption of the white dwarf emission by disc material in the dwarf nova OY Carinae, but calculated an N_H value that is an order of magnitude lower than we find for IP Peg. The higher N_H value we find for IP Peg may be due to the very high inclination of this system. If this value is correct,

it will be extremely difficult to detect the white dwarf in IP Peg at ultraviolet wavelengths. The ultraviolet continuum out-of-eclipse flux has been measured in IP Peg (Hoard et al. 1997; Saito, Baptista & Horne 2005), and is consistent with a cool white dwarf and/or heavy absorption of the white dwarf flux.

5.6 The long-term accretion rate

Townsley & Bildsten (2003) noted that the temperature of accreting white dwarfs is a good tracer of the medium term (10^3–10^5 yr) accretion rate. Our value for the IP Peg white dwarf implies a mean accretion rate of $<5 \times 10^{-11}$ M$_\odot$ yr^{-1}. The expected rate for a dwarf nova of this period is $\sim$2 $\times$ 10^{-9} M$_\odot$ yr^{-1}: more than 40 times greater. The value we find for the mean accretion rate is much closer to the instantaneous rate, which we estimate to be $\sim$1.8 $\times$ 10^{-10} M$_\odot$ yr^{-1} by following the calculation of Marsh (1988) and substituting in our new distance determination.

Our determination of the mean accretion rate implies a mass transfer time-scale of $\sim$10^{10} yr, much greater than the thermal time-scale for the donor, which is $\sim$10^8 yr. We would therefore expect the donor in IP Peg to be unperturbed by the mass transfer, in contrast to the donors in many other CVs which are observed to be oversized when compared to isolated main sequence stars of equivalent mass. As we noted in Section 5.2, we do find mass and radius for the donor that are consistent with those of a main sequence star in thermal equilibrium.

The low temperature of this system suggests that there could be a selection bias in the sample of long-period systems presented in Townsley & Bildsten (2003), and further detailed studies of eclipsing systems are necessary in order to determine the extent of this bias. At this stage, for example, it is possible that mass transfer in IP Peg has begun very recently and so the temperature in this system is a poor indicator of the long-term rate. One possible mechanism is the 'hibernation scenario' proposed by Shara et al. (1986), in which, following a classical nova eruption and a white dwarf cool-down period, the binary is detached for 10^3–10^6 yr. Shara et al. (1986) argue that an old nova emerging from this period of detachment (due to angular momentum loss bringing the donor back into contact with its Roche lobe) might first appear as a dwarf nova.

Alternatively, we can take our low value for the mass transfer rate over the past 10^3–10^4 yr in this system to be correct. This is surprising, given the contention that long-period dwarf novae have a high mean $\dot{M}$ despite their observed instantaneous values. The low temperature of IP Peg would suggest that it has sustained an accretion rate well below the expected mean rate for more than 1000 yr. If all 'steady-state' CVs go through low accretion rate dwarf nova phases, then this result implies the length of these phases can be very long indeed.

The third possibility is that there are problems with the classical picture of CV evolution as driven by angular momentum loss. The orbital period distribution of CVs (Spruit & Ritter 1983) has traditionally been explained by the disrupted magnetic braking model (Robinson et al. 1981), whereby angular momentum loss above the 2–3 h 'period gap' is due to a combination of gravitational radiation and magnetic stellar wind braking, and below the gap is due to gravitational radiation only. It follows that angular momentum loss is higher at longer periods, and since angular momentum loss is proportional to $\dot{M}$, $\dot{M}$ is expected to be high at long periods. The existence of low-$\dot{M}$ dwarf novae like IP Peg at long-period conflicts with this picture. Some authors suggest alternatives, such as Willems et al. (2007), who consider the effect of circumbinary discs on CV evolution. The theoretical orbital period distribution
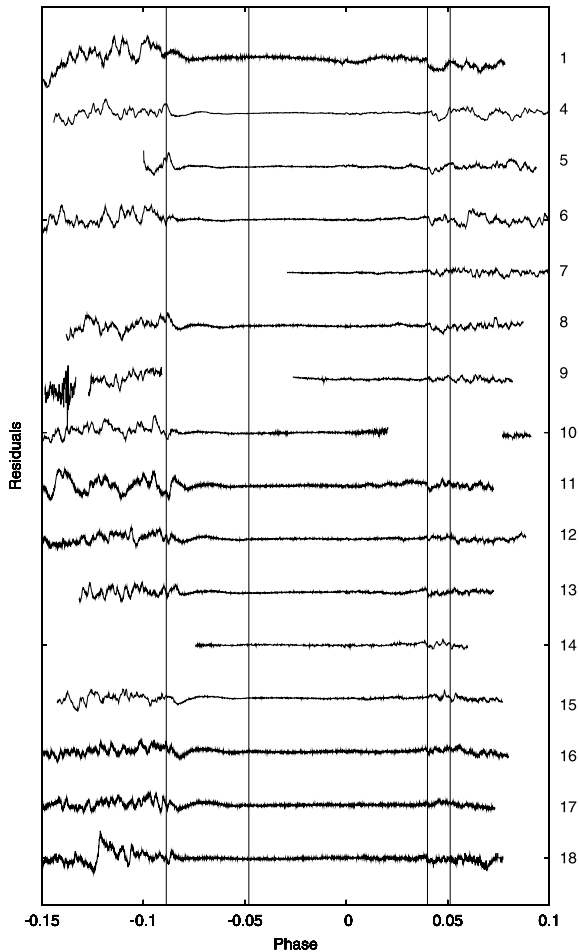
Figure 8. The residuals for the model fits to the individual g'-band light curves. The four vertical lines show (from left to right) the beginning and endpoints of the bright spot ingress, followed by the beginning and endpoints of the bright spot egress. A phase of zero corresponds to the mid-point of the white dwarf egress. We omit eclipse 2 because it includes no data for the bright spot, and eclipse 3 because the flare event near the white dwarf egress precludes a precise fit.

of Willems et al. (2007) is a reasonable reproduction of the observed distribution, and they find that systems with unevolved donors need not detach and evolve below the period gap as in the disrupted magnetic braking model. In their model, different CVs can have very different long-term $\dot{M}$ s, and some systems can exist as dwarf novae throughout their lifetime. This would explain the $\dot{M}$ we measure in IP Peg, although study of further systems is necessary to distinguish between this and the other possibilities.
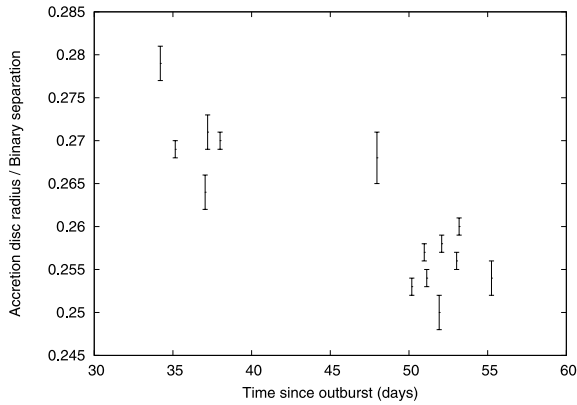
5.7 The bright spot

It is apparent from the light curves (Fig. 1) that most of the flickering originates in the bright spot. The bright spot is modelled as a linear region which approximately follows the outer circumference of the disc, with a luminosity which decreases approximately exponentially from the point of impact of the accretion stream (see Appendix A). The geometry is such that it is the impact region which is eclipsed first, and is first to emerge from eclipse. In Fig. 8, we plot the residuals from the model fits to the individual light curves, and we indicate on this figure the beginning and endpoints of the bright spot ingress and egress features. It is apparent in this figure that the majority of the variation due to flickering is localized in the impact region, since we see the residuals very quickly flatten off following the beginning of the ingress, but return immediately following the beginning of the egress.

In Fig. 9, we plot the colour variation along the bright spot. For this, we used the bright spot parameters listed in Table A1 with equation (A3) to calculate analytically the flux in each filter at each point along the length of the spot. We plot the bright spot colour from the point of impact of the gas stream on the accretion disc (labelled 'A' in Fig. 9) to a point 10 scalelengths (l; see Appendix A4) away from the impact region (labelled 'B'). We plot also a line showing the colours of blackbodies with temperatures in the $\sim$6000–17 000 K range. At the point of impact, the bright spot is close in colour to a 13 000 K blackbody, and as we move away from the impact region the colours of the bright spot decrease approximately in the direction of decreasing blackbody temperature. The spectroscopic study of Marsh (1988) also found the closest blackbody fit to the bright spot was for a temperature of 13 000 K. The shape of the bright spot line is not very important since it is very model dependent and as we note in Appendix A4, our bright spot model does not describe the
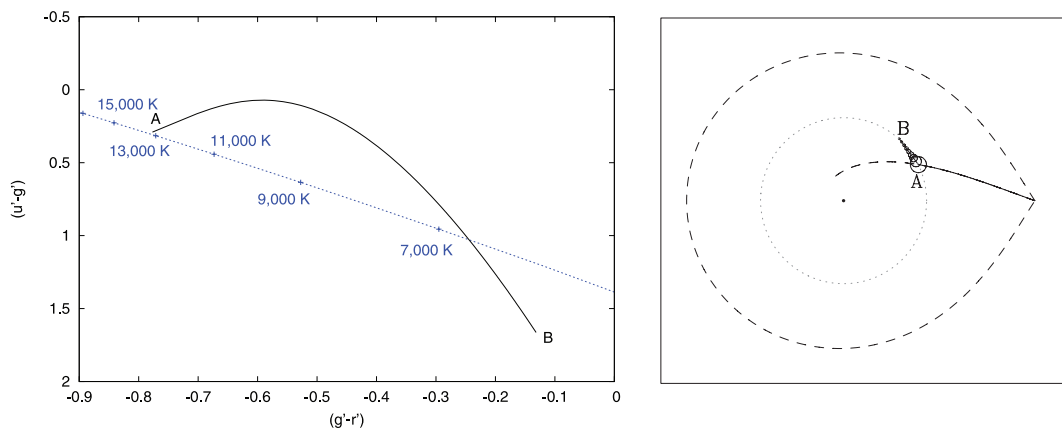
Figure 9. Left: the solid line (black) shows the colour variation along the bright spot, from the point of impact of the gas stream (labelled 'A'), to a point 10 scalelengths away from this impact region (labelled 'B'). The dashed line (blue) shows for comparison the blackbody colours for a temperature ranging from $\sim$6000–17 000 K. Right: schematic of the primary Roche lobe. The dashed line shows the extent of the Roche lobe, the dotted line shows the accretion disc radius and the central point is the white dwarf (all to scale). The line originating at the L1 point is the gas stream, and the bright spot is plotted from the impact region 'A' to 'B' as a series of circles, with the area of each circle representing the surface brightness at that point.

Figure 10. Variations in accretion disc radius over the course of the 2005 August/September observations. We plot the accretion disc radius R_{disc} scaled by the binary separation against the MJD time since the last outburst.

IP Peg bright spot perfectly. However, the fact that the general trend of the line is in the direction of decreasing blackbody temperature is clear evidence for cooling along the line of bright spot elements. The variable temperature of the bright spot in IP Peg will cause it to vary according to the wavelength of observation. For instance, near-infrared data suggest a cooler 10 000 K (Froning et al. 1999), and hotter parts than we see may dominate in the far-ultraviolet.

The points 'A' and 'B' are also labelled in the schematic of the primary Roche lobe. An additional point to note from the schematic is that the angle of the bright spot line of elements is such that it runs approximately along the circumference of the accretion disc. This behaviour is not enforced in our model, but seems physically appropriate: the incoming material from the gas stream is thermalized at point 'A' when it strikes the dense accretion disc, and as it cools, it propagates around the disc along a line of approximately constant radius and disc density.
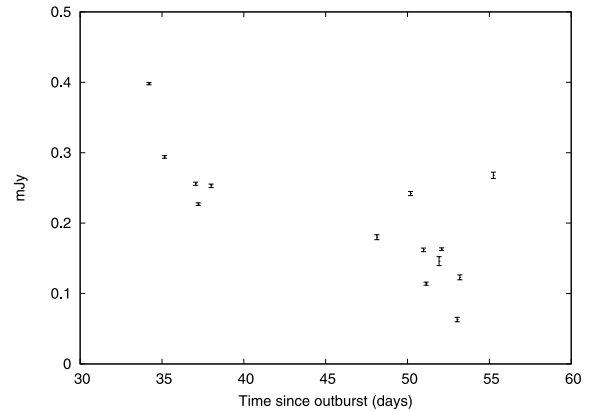
5.8 The accretion disc

5.8.1 Accretion disc radius

Data collected by the American Association of Variable Star Observers[1] shows that IP Peg was observed to go into outburst approximately 35 d prior to the beginning of our 2005 August observations. In Fig. 10, we plot the accretion disc radius scaled by the binary separation against the MJD time since the detection of the outburst, using the g'-band model fits. We see there is a decline in disc radius over the course of the ~20 d covered by our observations. During outburst the accretion disc radius increases as angular momentum is transported outwards due to the disc viscosity, and following the outburst the disc radius begins to shrink as low angular momentum material is added to the disc via the gas stream. It is this decline we observe in Fig. 10. We compared these values with previous determinations of the disc radius given by Wood et al. (1989) and Wolf et al. (1993). In Fig. 6 of Wolf et al. (1993), disc radius is plotted as a function of time since outburst, and our measurements are consistent with the rate of decline shown in their figure.

5.8.2 Accretion disc flux

The current and well-established paradigm for outbursts in dwarf novae is the thermal-viscous accretion disc instability model (Osaki

[1] http://www.aavso.org/

Figure 11. Variations in accretion disc flux over the course of the 2005 August/September observations. We plot the g'-band accretion disc flux at the point of the white dwarf egress against the MJD time since the last outburst.

1974). This model proposes that the outbursts are due to the bistable nature of accretion discs when $T_{disc} \sim 10\,000$ K, at which point hydrogen in the disc changes from a neutral to an ionized state (Meyer & Meyer-Hofmeister 1981). During quiescence, matter is stored in the disc and at a critical amount the instability sets in and a significant fraction of this mass is very rapidly accreted on to the central object. There are a number of detailed one-dimensional calculations of the disc instability model (see e.g. Hameury et al. 1998; Buat-Ménard, Hameury & Lasota 2001), and among their predictions is a significant increase in the disc luminosity during quiescence, since as matter is stored in the disc the surface density increases, prompting an increase in the effective temperature. However, this is not observed: long-term studies of dwarf novae show the visible flux is approximately constant during quiescence (Cannizzo & Mattei 1992). Much of the existing data has been collected by amateurs and is of variable quality, but the models predict an increase of ~1–3 visual magnitudes, which should be easily detectable in these data (Smak 2000). It has been proposed that these observations could be explained by a hot, high-viscosity inner region of the disc which slowly cools during quiescence, and the decrease in flux from this region counterbalances the increase in the effective temperature of the outer disc (Truss, Wynn & Wheatley 2004).

Eclipsing systems such as IP Peg allow for a more precise measurement of disc flux variations, since during the phase range over which both the white dwarf and the bright spot are eclipsed, the only contributors to the observed flux are the donor star and the accretion disc. The donor star component should be essentially constant, and so by measuring the average flux at a constant phase for every eclipse, we can determine the luminosity evolution of the accretion disc. We chose to measure the flux at the phase of the white dwarf egress, since at this point we observe a complete half of the disc and so our results should be unbiased by temporal radial variations. We used our model fits to subtract the white dwarf contribution at this point. We plot the average flux over a 10 s period at mid-egress in Fig. 11, and see an apparent decrease in disc flux over the period covered, in contrast to the increase which is predicted by the disc instability model. Our findings are consistent with the 'mirror eclipse' phenomenon reported by Littlefair et al. (2001), which they interpreted as eclipses of the donor star by the optically thin outer edge of the accretion disc. This conflicts with the disc instability model, which requires an optically thick disc for physically viable values of the disc viscosity parameter. Note also that the decrease in

brightness we observe (from ∼0.4 to 0.1 mJy) cannot be explained purely in terms of the disc radius variations we discussed in Section 5.8.1, since we see in Fig. 10 that the change in disc area over the course of our observations is only ∼25 per cent. Our results demonstrate a genuine dimming of the disc over the course of our observations. We also do not observe any evidence for the hot inner disc region proposed by Truss et al. (2004).

6 CONCLUSIONS

In 2004 and 2005, we took 18 separate observations of the eclipse in the 3.8 h period dwarf nova IP Peg. We used the triple-band camera ULTRACAM mounted on the William Herschel Telescope, and observed simultaneously in the Sloan u', g' and r' bands. By phase-folding these data, we identify for the first time the white dwarf ingress in this system. The phase width of the white dwarf eclipse is 0.0935 ± 0.0003, significantly higher than the standard value used by previous authors. This new value implies a higher orbital inclination or a higher mass ratio or some combination of the two. Our model fits strongly favour a high orbital inclination, which we find to be $83°.81 \pm 0°.45$. Our determinations of the masses of the two components are higher than those reported by previous authors.

We find the brightness of the white dwarf egress feature to be highly variable in flux and colour. We suggest that this is due to photoelectric absorption as a result of obscuration of the white dwarf by material from the accretion disc. We find the duration of the egress feature to be consistent in length, and so suggest that past reports of very extended egresses were due to failures to detect the egress at all, as a result of this varying obscuration.

It is difficult to determine an accurate temperature for the white dwarf due to the obscuration, but by comparing our data with synthetic white dwarf models we find the most likely temperature to be in the range of 10 000–15 000 K. This is very low for a dwarf nova above the period gap and, following Townsley & Bildsten (2003), implies a mean accretion rate over the past 1000–10 000 yr of $< 5 \times 10^{-11} \, M_\odot \, yr^{-1}$. This is more than 40 times lower than the expected rate. Additionally, we find a mass and radius for the donor that are consistent with those of a main sequence star in thermal equilibrium, which also suggests a very low accretion rate. Unless the mass transfer in IP Peg has begun very recently, these findings imply either that CVs can sustain accretion rates well below the expected rate for very long periods of time, or that the classical picture whereby the long-term accretion rate scales with period is not correct.

of IP Peg. We would additionally like to thank the referee, Robert C. Smith, for comments which led to a number of improvements to this paper.

REFERENCES

Bailey J., 1979, MNRAS, 187, 645
Bailey J., 1981, MNRAS, 197, 31
Baraffe I., Chabrier G., Allard F., Hauschildt P. H., 1998, A&A, 337, 403
Beekman G., Somers M., Naylor T., Hellier C., 2000, MNRAS, 318, 9
Bergeron P., Wesemael F., Beauchamp A., 1995, PASP, 107, 1047
Beuermann K., 2006, A&A, 460, 783
Buat-Ménard V., Hameury J.-M., Lasota J.-P., 2001, A&A, 366, 612
Cannizzo J. K., Mattei J. A., 1992, ApJ, 401, 642
Dhillon V. S. et al., 2007, MNRAS, 378, 825
Efron B., 1979, Ann. Stat., 7, 1
Efron B., Tibshirani R. J., 1993, An Introduction to the Bootstrap. Chapman & Hall, New York
Eggleton P. P., 1983, ApJ, 268, 368
Feline W. J., Dhillon V. S., Marsh T. R., Brinkworth C. S., 2004, MNRAS, 355, 1
Feline W. J., Dhillon V. S., Marsh T. R., Watson C. A., Littlefair S. P., 2005, MNRAS, 364, 1158
Froning C. S., Robinson E. L., Welsh W. F., Wood J. H., 1999, ApJ, 523, 399
Gänsicke B. T., Beuermann K., de Martino D., 1995, A&A, 303, 127
Goranskii V. P., Lyutyi V. M., Shugarov S. Y., 1985a, SvAL, 11, 293
Goranskii V. P., Shugarov S. Y., Orlowsky E. I., Rahimov V. Y., 1985b, Inf. Bull. Var. Stars, 2653, 1
Hameury J.-M., Menou K., Dubus G., Lasota J.-P., Hure J.-M., 1998, MNRAS, 298, 1048
Hessman F. V., 1989, AJ, 98, 675
Hoard D. W., Baptista R., Eracleous M., Horne K., Misselt K. A., Shafter A. W., Szkody P., Wood J. H., 1997, MNRAS, 288, 691
Holberg J. B., Bergeron P., 2006, AJ, 132, 1221
Horne K., Marsh T. R., Cheng F. H., Hubeny I., Lanz T., 1994, ApJ, 426, 294
Knigge C., 2006, MNRAS, 373, 484
Littlefair S. P., Dhillon V. S., Marsh T. R., Harlaftis E. T., 2001, MNRAS, 327, 475
Littlefair S. P., Dhillon V. S., Marsh T. R., Gänsicke B. T., 2006, MNRAS, 371, 1435
Littlefair S. P., Dhillon V. S., Marsh T. R., Gänsicke B. T., Southworth J., Baraffe I., Watson C. A., Copperwheat C., 2008, MNRAS, 388, 1582
Lubow S. H., Shu F. H., 1975, ApJ, 198, 383
Marsh T. R., 1988, MNRAS, 231, 1117
Martin J. S., Jones D. H. P., Smith R. C., 1987, MNRAS, 224, 1031
Martin J. S., Jones D. H. P., Friend M. T., Smith R. C., 1989, MNRAS, 240, 519
Meyer F., Meyer-Hofmeister E., 1981, A&A, 104, L10
Osaki Y., 1974, PASJ, 26, 429
Piro A. L., Bildsten L., 2004, ApJ, 610, 977
Press W. H., 2002, Numerical Recipes in C++ : The Art of Scientific Computing. Cambridge Univ. Press, Cambridge
Pyrzas S. et al., 2009, MNRAS, 394, 978
Ramseyer T. F., 1994, ApJ, 425, 243
Ribeiro T., Baptista R., Harlaftis E. T., Dhillon V. S., Rutten R. G. M., 2007, A&A, 474, 213
Ritter H., Kolb U., 2003, A&A, 404, 301
Robinson E. L., Barker E. S., Cochran A. L., Cochran W. D., Nather R. E., 1981, ApJ, 251, 611
Saito R. K., Baptista R., Horne K., 2005, A&A, 433, 1085
Shafter A. W., 1992, ApJ, 394, 268
Shara M. M., Livio M., Moffat A. F. J., Orio M., 1986, ApJ, 311, 163
Smak J., 2000, New Astron. Rev., 44, 171
Smak J. I., 2002, Acta Astron., 52, 189
Smith J. A. et al., 2002, AJ, 123, 2121

Southworth J., Hickman R. D. G., Marsh T. R., Rebassa-Mansergas A., Gänsicke B. T., Copperwheat C. M., Rodríguez-Gil P., 2009, A&A, 507, 929

Spruit H. C., Ritter H., 1983, A&A, 124, 267

Szkody P., Mateo M., 1986, AJ, 92, 483

Townsley D. M., Bildsten L., 2003, ApJ, 596, L227

Townsley D. M., Gänsicke B. T., 2009, ApJ, 693, 1007

Truss M. R., Wynn G. A., Wheatley P. J., 2004, MNRAS, 347, 569

Tulloch S. M., Rodríguez-Gil P., Dhillon V. S., 2009, MNRAS, 397, L82

Verbunt F., Rappaport S., 1988, ApJ, 332, 193

Warner B., 1995, Cataclysmic Variable Stars, Cambridge Astrophys. Ser. Vol. 28. Cambridge Univ. Press, Cambridge

Watson C. A., Dhillon V. S., Rutten R. G. M., Schwope A. D., 2003, MNRAS, 341, 129

Willems B., Taam R. E., Kolb U., Dubus G., Sandquist E. L., 2007, ApJ, 657, 465

Winter L., Sion E. M., 2003, ApJ, 582, 352

Wolf S., Mantel K. H., Horne K., Barwig H., Schoembs R., Baernbantner O., 1993, A&A, 273, 160

Wood J., Crawford C. S., 1986, MNRAS, 222, 645

Wood J. H. et al., 1989, MNRAS, 239, 809

APPENDIX A: THE LIGHT CURVE MODELLING CODE

In this section, we provide a concise description of the code, LCURVE, that we wrote in order to model the light curves of this paper; the same code has been used in Pyrzas et al. (2009) and Southworth et al. (2009). The aim of the code was to provide a flexible framework for fitting light curves characteristic of eclipsing dwarf novae and detached white dwarf/M dwarf binary stars. LCURVE was developed as a generalization of the code used in Horne et al. (1994) to model the dwarf nova OY Car and subsequently employed in the analysis of ULTRACAM photometry of several other systems (Feline et al. 2004, 2005; Littlefair et al. 2006).

The binary is defined by up to four components: white dwarf, secondary star, accretion disc and bright spot. In the case of detached systems, the latter two are disabled. Each component is described by a set of small, flat elements, each of specified area, position, orientation and brightness. The light curves are calculated by summing the contributions from all faces which are oriented towards the observer, allowing for eclipses by the two stars. Additionally, it is possible to make the accretion disc opaque or transparent. The contribution of a face depends upon how it is oriented with respect to the line of sight. Denoting the angle between the perpendicular to a face and the vector pointing towards Earth as θ, there is no contribution if $\theta > \pi/2$ radians, or equivalently if $\mu = \cos\theta < 0$. Otherwise, the contribution is typically specified using a limb-darkened Lambert's law behaviour:

$$I \propto \mu\,(1 - \epsilon + \epsilon\mu)\,, \tag{A1}$$

where ϵ is a linear limb-darkening coefficient (the code also allows for quadratic limb darkening).

The effect of finite length exposures is allowed for by calculating multiple points covering an exposure and using trapezoidally weighted averaging. Thus, 30 s long exposures of a white dwarf eclipse can in the model be split into seven points separated by 5 s for example.

In Horne et al. (1994), the overall scaling of the various components was handled using singular-value decomposition (SVD), allowable if all components contribute linearly to the light curve. The advantage of this approach is speed: during the fitting process, several potential free parameters are immediately removed at an early stage, allowing more time for optimization of the remaining

model parameters. However, it is not always a valid approach, in particular if 'reflection' (more properly reprocessing) is significant, then one cannot assume a simple linear combination, since doubling the contribution of the white dwarf, say, may also imply an altered contribution from the heated secondary star. Thus, in LCURVE one can switch off the SVD scaling, falling back on the physical model to handle the effects of reflection.

We experimented with several forms of minimization: the simplex method implemented via 'amoeba' (Press 2002), Powell's method and the Levenburg–Marquardt method. The simplex method has the advantage of robustness, while Levenburg–Marquardt can deliver uncertainty estimates once it has reached a minimum. However, we found neither to be particularly secure in reaching the minimum χ^2 for the rather complex fits that we undertook in this paper. Early on in the fitting we realized on several occasions that we were trapped in local minima, but even when we were not, the routines would often report that the fits were converged when subsequent work revealed this not to be the case. We believe that this was largely caused by degeneracies in the model which are very hard to remove entirely. There is for instance a fundamental degeneracy between mass ratio q and orbital inclination i that cannot be avoided. Ultimately, the method that helped above all in obtaining the final fits, and which, after extensive use, gave us confidence that we did achieve global minima, was the MCMC method. Our standard technique evolved to become (i) carry out an initial minimization with the simplex method and (ii) then run an MCMC chain until it reaches an essentially flat run of χ^2 with fit number, consistent with the statistical fluctuations characteristic of the MCMC method. For our main fits, we performed the same process from 10 independent starting models to check that we obtained consistent results. We typically required chains of 10 000 fits or more to reach the desired quasi-static state. We strongly recommend this technique in all similar work to guard against premature claims of 'convergence'.

Apart from the general physical parameters mentioned above and others associated with the individual components, the other parameters needed to define the model, were computational, such as the number of elements needed to define the various components. Our strategy with these was simply to make them large enough that they had no significant effect upon the best-fitting parameters. We finish by describing the model used for the individual components in more detail.

A1 The white dwarf

We modelled the white dwarf as a linearly limb-darkened sphere. Although we also enabled the possibility of Roche-distortion, we never used it in the fits presented here. The eclipse of the three other components by the white dwarf was included by default, although its effect is tiny in the case of IP Peg since the white dwarf is so small.

A2 The secondary star

The secondary star was modelled with full accounting for distortion in Roche geometry. Again linear limb darkening was assumed, and heating by flux from the white dwarf was included along with gravity darkening. None of this is especially significant for IP Peg as the secondary star contributes relatively little. The important feature is of course the eclipse, again calculated accounting for Roche geometry.

Table A1. Results from our MCMC fits to the phase-folded light curves, as detailed in Section 4.1.

	u' band	g' band	r' band
q	0.457 ± 0.005	0.481 ± 0.003	0.497 ± 0.005
i	84.08 ± 0.20	83.64 ± 0.10	83.31 ± 0.16
R_1/a	0.0058 ± 0.0004	0.0045 ± 0.0002	0.0051 ± 0.0003
R_{disc}/a	0.333 ± 0.004	0.285 ± 0.006	0.281 ± 0.004
R_{spot}/a	0.263 ± 0.001	0.261 ± 0.001	0.257 ± 0.0001
β	0.51 ± 0.06	0.19 ± 0.02	0.25 ± 0.03
γ	0.727 ± 0.022	0.532 ± 0.006	0.505 ± 0.009
δ	0.041 ± 0.069	-0.262 ± 0.065	-0.100 ± 0.045
l	0.0102 ± 0.0012	0.0097 ± 0.0001	0.0096 ± 0.0001
f_c	0.711 ± 0.008	0.723 ± 0.004	0.723 ± 0.006
ϕ	105.9 ± 1.6	126.6 ± 0.5	125.1 ± 0.6
ψ	59.1 ± 1.7	32.7 ± 0.8	39.6 ± 0.8

Note. q, the mass ratio; i, the binary inclination; R_1, the white dwarf radius; R_{disc}, the accretion disc radius; R_{spot}, the distance between the maximum of the surface brightness of the bright spot (equation A4) and the white dwarf; a, the binary separation; β and γ, the power-law exponents for the bright spot; δ, the exponent of surface brightness over the accretion disc; l, the bright spot scalelength; f_c, the fraction of the bright spot taken to be equally visible at all phases; ϕ, the angle made by the line of elements of the bright spot, measured in the direction of binary motion from the line from accretor to donor; ψ, the angle away from the perpendicular to the line of elements at which the light from the bright spot is beamed, measured in the same way as ϕ.

A3 The accretion disc

We model the accretion disc as a symmetric, flattened disc with user-defined inner and outer radii. In the case of IP Peg, we forced the inner radius to match the radius of the white dwarf. When fitting the individual light curves, the outer disc radius R_{disc} was set to match the radius of the bright spot (R_{spot}, next section) although in the case of the phase-folded light curves we found a small improvement in the fit when we allowed R_{disc} to be a free parameter. These fits show R_{spot} to be somewhat smaller than R_{disc}, particularly in the u' band. This may in part be due to the disc being non-circular, but note also R_{disc} is rather poorly constrained and the uncertainties in these values are likely much greater than the formal errors listed in Table A1.

In between the user-defined inner and outer radii, the height of the disc was defined to follow a power law in radius, i.e.

$$h(R) = h_0 R^\alpha, \tag{A2}$$

where h_0 and α are parameters of the model. The height parameter h_0 and the radius R, along with all other lengths, were scaled in terms of the binary separation, a. We fixed $h_0 = 0.02$ and $\alpha = 1.5$. Since the outer disc radius $\approx 0.3\, a$, this is equivalent to a very thin disc. The feeble emission from the disc meant that h_0 and α have little overall impact, and so they were not optimized during the fits. The surface brightness of the disc was also a power law in radius, with the overall level defined by the temperature at the outermost radius of the disc, translated into a surface brightness assuming a blackbody spectrum given the central wavelength of the filter concerned. It is important to realize that this does not mean that we are saying that the radiation from the disc is that of a blackbody; it is just one way of normalizing the overall contribution from the disc. Nevertheless, our model can certainly be criticized for the assumption of a power-law variation in surface brightness and for its symmetry. Neither is necessarily correct, on the other hand there is nothing in the data that definitely points the way to a better alternative. We did allow the power-law exponent of the surface brightness to be optimized,

since it has no clear value *ab initio* for quiescent discs, and since it has a significant effect, particularly during the egress phases of the eclipse.

A4 The bright spot

The bright spot was the hardest component to model, and indeed, as Fig. 2 shows, we never achieved an entirely satisfactory fit. The bright spot is of course a uniquely strong feature of IP Peg, and so it is to be expected that it would cause difficulties. We model the bright spot as a series of elements which lie along a straight line in the orbital plane (as we illustrate in the Fig. 9 schematic). The angle (ϕ) that the line makes to the line of centres between the two stars is a key defining feature of the bright spot since it is this, in conjunction with the distance of the bright spot from the white dwarf and the mass ratio and inclination, which determines the relative lengths of the bright spot ingress and egress features, which are both well-resolved in our data. Each element was then assumed to contribute some fraction (f_c) that was constant in phase together with some fraction of light modulated according to the angle of the element. In our model, the elements were assumed oriented perpendicular to the orbital plane (although we can allow them to have arbitrary orientation), so they give light curves that peak when we see them face-on. The angle ψ is defined as the angle away from the perpendicular to the line of bright spot elements at which the light from the bright spot is beamed. In this manner, the variation of the orbital 'hump' and the relative size of the bright spot ingress and egress features can be fitted. Note that due to the way they are defined, there is a strong correlation between the angles ϕ and ψ. This degeneracy can be broken if the observed phase range covers the entire orbital 'hump', but this is not the case for our IP Peg data. These angles are therefore poorly constrained by our model and we see in Table A1 that the fitted values for ϕ and ψ vary significantly between the different bands. Note however that $\phi + \psi$ is consistent, and so the direction of peak emission is the same for all photometric bands. Finally, we should point out that we find no correlation between these parameters and q or i, and so the fact that these bright spot angles are poorly constrained does not introduce any additional uncertainty in the physical parameters listed in Table 3.

The surface brightness of the elements in the bright spot were specified by

$$S \propto \left(\frac{x}{l}\right)^\beta \exp\left[-\left(\frac{x}{l}\right)^\gamma\right], \tag{A3}$$

where x is the distance along the line defining the bright spot, β and γ are two more power-law exponents which allow some flexibility in how the brightness varies, and l is a scalelength. The final element of the bright spot is its position within the binary. The maximum of the surface brightness occurs at

$$x = l \left(\frac{\beta}{\gamma}\right)^{1/\gamma}. \tag{A4}$$

We define this location to lie upon the ballistic gas stream, at a specified distance from the white dwarf R_{spot}, this 'bright-spot radius' being the final parameter needed to define the bright spot. The path of the gas stream is calculated by integration of the equations of Lubow & Shu (1975).

Compared to the modelling of the bright spot in Horne et al. (1994) and the related papers mentioned before, the parameter γ

allows increased flexibility in the specification of the bright spot, since in the earlier work, $\gamma = 1$ by default. This extra flexibility comes at the cost of increased degeneracy; the MCMC method was essential to cope with this. Despite this, as mentioned earlier, the strong bright spot of IP Peg is still not modelled perfectly by our code, but we felt that adding yet more complexity to the model would not be justified by any increased confidence in the fundamental parameters such as q and i.

This paper has been typeset from a TEX/LATEX file prepared by the author.