Mon. Not. R. Astron. Soc. **000**, 000–000 (0000) Printed 3 November 1998

Indirect imaging of the accretion stream in eclipsing polars I: method and tests

M.K. Harrop-Allin [1], P.J. Hakala [1,2] & Mark Cropper [1]

[1] *Mullard Space Science Laboratory, University College London, Holmbury St. Mary, Dorking, Surrey, RH5 6NT, UK*
[2] *Observatory and Astrophysics Laboratory, PO Box 14, FIN-00014 University of Helsinki, Finland*

3 November 1998

ABSTRACT

We present a method to deduce the brightness distribution of the accretion stream in eclipsing polars using photometric eclipse profiles and an indirect imaging technique based on Maximum Entropy. Our method is developed from Hakala (1995) and includes improvements to the assumed trajectory of the stream and to the numerical algorithm used to extract the information from an eclipse profile. We present tests of the method using synthetic data, and show that the method is able to retrieve details of the stream brightness distribution from an eclipse profile in the presence of moderate noise, provided the phase resolution of the light curve is good. The method can be applied to the eclipse profiles of any polar with a bright accretion stream to identify regions of enhanced emission between the two stars. The method also provides estimates of the radius at which the accretion stream becomes threaded onto the magnetic field of the white dwarf, and of the mass transfer rate through the stream.

Key words: accretion, accretion discs — methods: numerical — binaries: eclipsing — stars: magnetic fields — cataclysmic variables

1 MOTIVATION

The accretion flow in polars (AM Herculis stars) is dominated by the strong magnetic field ($\sim 10-240\,$MG) on the white dwarf. The basic characteristics of the accretion flow between the mass-donating secondary star and the white dwarf have been described by Liebert & Stockman (1985), Lamb (1985), Hameury, King & Lasota (1986), Lamb (1988) and Mukai (1988) and others, yet the detailed physics of the threading process remains poorly understood. It is thought that the stream falls initially on a ballistic trajectory from the L1 point until the radius $R_{\rm th}$ where the thermal pressure in the gas becomes comparable to the magnetic pressure in the white dwarf's magnetosphere. The stream is denser at its core than in its envelope, and the effect of the magnetic field will be significant initially in the outer parts of the stream (Mukai 1988). The threading region, where the stream becomes coupled to the field lines of the white dwarf, is bounded by $R_{\rm th}$ and the radius R_μ at which the ram pressure of the stream is equal to the magnetic pressure. In the threading region the stream plasma is pinched by the magnetic field and is shattered into a mixture of dense diamagnetic filaments ('blobs') and finely divided droplets. At first, the blobs can cross the field lines, since their dynamical time scale is much shorter than the time scale for penetration of the blobs by the field. The fine droplets follow the field lines almost immediately, since the field diffuses across them on time scales much shorter than the dynamical time scale (Hameury et al. 1986). The stream is compressed by the field and accelerated by distorted field lines (Liebert & Stockman 1985). The stream may produce a cavity in the magnetosphere of the white dwarf in which plasma accumulates and is heated by a series of hydromagnetic shocks, magnetic field reconnection and/or turbulence (Lamb 1985). Within R_μ, the magnetic field is able to withstand the ram pressure of the stream, and the stream is diverted above and/or below the orbital plane (depending on the system configuration). Any blobs that have remained intact to within R_μ are now threaded, with at most slow drifts across the field lines (Hameury et al. 1986). Finally, the stream material accretes quasi-radially near the magnetic poles of the white dwarf.

While this model remains the accepted qualitative picture, observational constraints on the characteristics of the accretion flow are hard to obtain. This is because, in many systems, emission from the stream is dominated at most wavelengths by that from the accretion region on the white dwarf. Nevertheless, progress has been made on several fronts. Observations of dips in X-ray and optical light curves of high-inclination systems are consistent with an inhomogeneous accretion stream consisting of blobs of plasma with radii $\lesssim 10^6$ cm embedded in a much rarer medium (e.g. V1432 Aql: Watson et al. 1995). A 'blobby' stream is also the accepted explanation for the enhanced soft/hard X-ray fluxes seen in many polars (Ramsay et al. 1994 and references therein). The flares seen in optical and X-ray light curves on time-scales of 5–100 s can be explained in terms of an inhomogeneous flow containing large blobs of plasma that have survived within R_μ (e.g. Warner 1995 and references therein). Analyses of trailed spectra (e.g. V834 Cen: Rosen, Mason & Cordova 1987) and Doppler tomograms (e.g. HU Aqr: Schwope, Mantel & Horne 1997) of the emission lines that originate in the accretion stream have shown that the line components can be identified with emission from the stream prior to the threading region, the threading region itself, and the flow within R_μ. Schwope et al. (1997) show that the Doppler tomograms of the He II λ4686 Å and Hγ lines of HU Aqr are consistent with a threading region measuring $\sim 6\times10^9$ cm along the stream trajectory.

In this paper we present a method that provides further observational constraints on the properties of the accretion stream. We have developed an indirect imaging technique which uses the broadband optical eclipse profiles of polars to determine the distribution of brightness along the accretion stream. For a fixed stream trajectory, we show that the stream brightness distribution can be retrieved from an eclipse profile, even in the presence of moderate noise. Using this technique, we aim to explore characteristics of the accretion flow in polars and to identify the regions of enhanced emission along the accretion stream. The technique can also be used to obtain estimates of the radius R_μ at which the ram pressure of the accretion stream equals the magnetic pressure, and hence the mass transfer rate through the stream.

2 DESCRIPTION OF THE METHOD

2.1 Previous work

Hakala (1995) (hereafter H95) presents a method to deduce the accretion stream brightness in the eclipsing polar HU Aquarii from photometric eclipse profiles. The method has close parallels with eclipse mapping of accretion discs in non-magnetic systems (e.g. Marsh & Horne 1988) in that both reconstruct the brightness of the region between the stellar components and both make use of Maximum Entropy regularization to constrain the problem (in general, there are more model parameters than there are data points). The main difference between eclipse mapping and accretion stream mapping is in the precise form of the regularization term: in eclipse mapping, the default brightness of an emission point is taken from a uniform image, while in accretion stream mapping the default brightness of an emission point is the local geometric mean of the adjacent brightnesses on the stream (see equation 1). Here we give a brief summary of the method; further details are available in H95, and we describe our improvements to the model in section 2.2. A model accretion stream is created by placing evenly-spaced emission points along a pre-set stream trajectory. The accretion regions on the white dwarf are modelled as single bright points, the brightnesses of which are fixed for each model calculation. This stream is then observed through an eclipse by the Roche lobe-filling secondary. A model light curve is generated by summing the brightnesses of the points that are visible at each orbital phase. The brightnesses of the emission points are optimized using a genetic algorithm (GA). Genetic algorithms mimic biological natural selection: the optimization proceeds by 'breeding' the 'fittest' solutions from an initial population of random solutions (for an introduction to GAs and their applications in Astronomy see Charbonneau 1995). The brightnesses are adjusted in order to minimize the quantity

$$F = \sum_{i=1}^{N} \left(\frac{\mathrm{model}_i - \mathrm{data}_i}{\sigma_i} \right)^2 - \lambda \sum_{j=1}^{M} S_j \qquad (1)$$

with $S_j = p_j - mj - p_j \ln(p_j/m_j)$, where $\lambda > 0$ is a Lagrange multiplier, N is the number of data points in the original light curve, M is the number of brightness points along the model accretion stream, model$_i$ is the brightness of the ith model light curve point, data$_i$ is the brightness of the ith point in the original light curve and σ_i its uncertainty, p_j is the normalized brightness of model stream point j, and m_j is the local geometric mean of the brightness of the emission points neighbouring point j on the model stream; this is $m_j = \sqrt{p_{j-1}p_{j+1}}$ for an emission point flanked by two neighbouring points. At the endpoints of the 1-dimensional distribution m_j

is given by $\sqrt{p_j p_{j+1}}$ for a point j at the beginning of a section of the stream and $\sqrt{p_j p_{j-1}}$ for a point j at the end of a section of the stream. The first summation on the right hand side of equation 1 is the χ^2 for the fitted model light curve, while the second sum is a Maximum Entropy (ME) regularization term (H95 uses a slightly different form of the regularization term). The χ^2 term ensures that the optimization finds a solution that is consistent with the original data, while the ME term ensures that the problem is not under-constrained. A local geometric mean is used as the default brightness in the ME term so that the stream brightness distribution obtained is locally smooth. The degree of smoothness required is adjusted by changing the value of λ (see section 2.4 for more details).

2.2 Improvements to the method of H95

We have developed the method of Hakala (1995) by including projection effects of the emission point brightnesses, by using a more realistic accretion stream trajectory, and by improving the GA optimizing procedure. In addition, we refine the solution found by the GA by using a line minimization technique.

Projection effects in the stream can be important in the high accretion state of polars (e.g. HU Aquarii: Schwope et al. 1997). The accretion stream path leads the line of centres between the two stars (Lubow & Shu 1975), and thus presents a larger projected area to the observer after the eclipse than before. The result is that the observed pre-eclipse light level is lower than the post-eclipse level. To account for this effect in our stream model, the brightness of each point along the model accretion stream is weighted by a projection factor equal to the sine of the angle between the line of sight and the tangent to the stream at that point.

The main difference between the model in H95 and that presented here is the assumed trajectory for the accretion stream. H95 approximates the accretion stream trajectory as an arc connecting the L1 point to the white dwarf; the stream is assumed to be confined to the orbital plane, and the white dwarf is treated as a dimensionless point. The trajectory used here is a combination of a ballistic trajectory from the L1 point (Lubow & Shu 1975), followed by a magnetically-channeled trajectory that follows a dipole field line out of the orbital plane. The white dwarf's magnetic field is assumed to be a centred dipole that can be inclined to the orbital plane and rotated with respect to the line of centres. The point where the ballistic trajectory meets the magnetically-channeled trajectory occurs at R_μ, and from this point the stream follows the field line above and/or below the orbital plane. The stream trajectory ends at the surface of the white dwarf at the footpoint(s) of the dipole field line. The white dwarf is not assumed to be a point source: allowance is made for self-occultation of the accretion regions by the white dwarf as well as by the Roche lobe-filling secondary. The radius of the white dwarf is calculated using the relation

$$R_9 = 1.12\,(1 - M_1/M_{\rm ch})^{3/5} \qquad (2)$$

for $0.7 < M_1 < 1.3$ (Warner 1995), where R_9 is the radius of the white dwarf in units of 10^9 cm, M_1 is the mass of the white dwarf in solar masses and $M_{\rm ch} = 1.44\,{\rm M}_\odot$ is the Chandrasekhar mass. We used ~ 200 emission points for a stream accreting onto both poles, and ~ 125 points for a stream accreting at one pole.

The brightnesses of the emission points along the stream are optimized in the first instance using a GA. We have made some improvements to the GA algorithm described by H95: these include increasing the selection pressure on the better solutions as the population of solutions evolves, and by the use of additional mutation operations that diversify the solutions and explore the parameter space more effectively.

Parent solutions in the GA are ranked according to their fitness F (evaluated using equation 1). The rank number of a selected parent solution determines the probability that the solution will be used to produce 'child' solutions by, in our case, uniform crossover (the brightness of each emission point on the stream in the child solution is taken either from one parent or the other). We have used a selection scheme that becomes increasingly biased towards better solutions as the GA proceeds, as follows. The rank numbers of selected parents are obtained using

$$S = {\rm INT}\left(-\frac{{\rm N_{pop}}}{10} s^{i-1} \log(RND)\right) + 1 \qquad (3)$$

where S is the solution's rank number after sorting, i is the generation number, ${\rm N_{pop}}$ is the population size, $0 < s < 1$ is a scale factor and RND is a random number uniformly distributed between 0 and 1. At the first generation ($i = 1$), 90 per cent of the chosen parent solutions are in the first 50 ranked solutions. The scale factor is chosen such that 90 per cent of the chosen solutions are in the first ~ 5 ranked solutions by the end of the GA run. The higher-ranked solutions are thus chosen as parents with increasing frequency as the GA proceeds.

Mutation is applied individually to each child solution. The mutation procedure alters each parameter p_j in the solution with a small probability (we used a probability of 0.005). We find that the way in which a mutated value of the brightness of an emission point is calculated has a significant effect on the final solution produced by the GA. We used a scheme where, for each individual mutation, there is an equal probability that the mutated parameter is assigned a random number from a uniform distribution, *or* a random number chosen from a Gaussian distribution centred on the old value of the parameter, *or* the mean brightness of the two neighbouring points on the stream. The mutation procedure thus has a multiple effect on the solution: depending on the selected mutation method, the solution is either diversified (so as to explore new regions of the parameter space) or smoothed.

The chief motivation for using a GA instead of a more conventional optimization routine is to maximize our chances of finding the global optimum in our multi-dimensional space. GAs cannot guarantee that the global optimum will be reached but, unlike conventional steepest descent algorithms, they are capable of climbing away from local optima. It remains possible that the GA will fail to find the global minimum before it is terminated (see section 2.3), and the algorithm will converge on a local minimum in the parameter space. Although we cannot prove that our solutions do reach the global minimum, we can demonstrate empirically the stability of the solutions obtained by using different random initial populations and evolving them in the same way. We illustrate this in section 3.3.

The best solution at the end of the GA stage is thus likely to be in the neighbourhood of the global optimum, but may not be a precise estimate of this optimum. To refine the solution for the stream brightness distribution, we use the best solution from the GA as the starting point for a more conventional line minimization routine (we used Powell's method, e.g. Press et al. 1992). In this way we exploit the complementary strengths of the two optimization methods. The genetic algorithm performs an efficient search of the multi-dimensional optimization space because of its implicit parallelism, but the amount of computer time required to obtain a precise estimate of the global optimum is prohibitive. The line minimization routine, given a starting point, will converge with high precision to the optimum nearest that starting point. We therefore use the GA to place the solution in the neighbourhood of the global minimum and then use Powell's method to refine the solution within this neighbourhood. Another example of the improvement of a GA solution using a line minimization routine can be seen in Potter, Hakala & Cropper (1998).

2.3 Stopping criteria

It is not possible to define a formal convergence criterion for the GA in terms of gradient information (e.g. where the gradients of the χ^2 and regularization terms are anti-parallel). This is because the GA does not make use of gradient information to improve its solutions from one generation to the next, but rather makes random jumps throughout the multi-dimensional space. The algorithm then breeds the solutions that give the smallest values of the fitness function to form the next generation of solutions. Certain stochastic convergence criteria have been proposed for GAs (e.g. Beasley, Bull & Martin 1993) but these are dependent on the specific algorithm used. In practice, the best one can do is to monitor the progress of the algorithm and to terminate it when the rate of improvement of the solutions from one generation to the next has decreased to a very low value. Our stopping criterion for the GA is the following: if the value of F does not improve by more than than 1 per cent over 20 successive generations, the algorithm is terminated. During the subsequent stage of the algorithm where the solution is refined using Powell's method, convergence is reached (in theory) when the local gradient of the parameter space is zero. In practice, we use the stopping criterion in Press et al. (1992) where the algorithm is terminated if the fractional decrease in F between successive iterations is less than a certain tolerance (we use 10^{-5}).

The usual criterion for a good fit is $\chi^2 = N$, the number of data points in the original light curve (this condition is used in other applications of maximum entropy image reconstruction, e.g. Cropper & Horne 1994). For synthetic data the criterion is usually satisfied after application of Powell's method. For real data, the solutions seldom reach $\chi^2 = N$. This is because the very steep changes in the eclipse profile (due to the ingress and egress of the accretion region) are not instantaneous, whereas the model assumes that the accretion regions are point sources. In the model, the ingress and egress of the accretion region are thus unresolved, irrespective of the phase resolution. A very slight misalignment between the model curve and a data point on the steep component of a real eclipse profile results in a significant contribution to χ^2. For real data, therefore, we are satisfied with $\chi^2/N \approx 3.0$. The alternative is to allocate artificially large errors to the light curve during the ingress and egress of the accretion region, in which case fits with $\chi^2 \lesssim N$ are possible.

2.4 Parameters

The fixed parameters used to model the brightness distribution of the accretion stream fall into two groups: parameters that determine the properties of the numerical algorithm, and the physical and geometric parameters that define the characteristics of the binary system and the accretion stream trajectory. The remaining set of parameters are those that are adjusted to minimize the fitness function F (equation 1) and thus produce the optimal solution: these are the brightnesses p_j of the emission points along the stream trajectory.

Parameters that determine the properties of the algorithm include the population size used in the GA, the mutation rate (described above), the number of generations (iterations) performed by the GA, and the value of the Lagrange multiplier λ used in the GA and in the line minimization routine.

The smoothness of the final solution is determined by the Lagrange multiplier λ in equation 1. If the value of λ is too small, the model light curve will fit noise features in the original data; if λ is too large, the model cannot fit the small-scale features of the original data and the resulting stream brightness distribution will lack resolution. To find the balance between these two extremes, we have found that a good choice for λ is where the ME sum (the magnitude of the second sum in equation 1) is larger than χ^2 for the initial generations, and towards the end of the GA, χ^2 is larger than the ME sum. It is possible to find such a λ because the magnitudes of the two sums in equation 1 decrease at different rates during the GA. The same value of λ is used for Powell's method. The range of values of λ allowed by the above criterion is large, and the precise value chosen is arbitrary within this range. Although there are no formal methods for choosing λ, it is fairly easy to find a suitable value by trial and error; this is demonstrated in section 3.5. We have tried other schemes for λ, such as using a value that evolves during the course of the GA (as used by Potter et al. 1998), but the results are not noticeably different to those obtained using a constant λ. We have thus opted for the simpler scheme.

The stream trajectory and the properties of the binary system are pre-set in each model calculation using various geometric and physical parameters. These include the mass ratio q of the two stars, the mass M_1 of the white dwarf (required to determine the radius of the white dwarf via equation 2) and the binary inclination i. The parameters that determine the accretion stream trajectory are the magnetic colatitude β, the magnetic longitude ζ of the dipole field line, and the radius R_μ at which the stream changes from the ballistic trajectory to follow the dipole field line. The number of emission points along the stream trajectory and the brightnesses of the points on the white dwarf that represent the accretion regions are also held constant during the optimization.

3 TESTS WITH SIMULATED DATA

To confirm that our procedure reproduces accurately the stream brightness distribution for an assumed stream trajectory, we tested the method using synthesized data. In model streams that accrete onto both poles, in particular, we have to be sure that there is sufficient information in the eclipse profile to determine uniquely the distribution of brightness along both halves of the magnetically-channeled parts of the accretion stream.

A synthetic stream brightness distribution is constructed by manually assigning brightnesses to points along a pre-existing stream trajectory. A synthetic light curve is obtained by 'observing' the synthetic stream through an eclipse by the secondary star (given values for q and i). The synthetic light curve is then used as an input for the optimizing algorithm. The stream brightness distribution produced by the algorithm can then be compared to that of the original stream to assess the reliability of the method. We discuss both two-pole and one-pole stream geometries, and discuss the effects of increased noise and decreased phase resolution.

3.1 Ideal test cases

We first establish that the method is able to retrieve information from an eclipse profile in the ideal case, i.e. when applied to a light curve with no added noise, very small uncertainties σ_i and a high phase resolution.

3.1.1 One-pole geometry

The input stream has a brightness enhancement at the junction between the ballistic trajectory and the magnetic trajectory. The motivation for chosing this brightness distribution is to check whether the optimizing algorithm is able to reproduce any brightening that might be associated with the threading region. The model parameters used are $R_\mu = 0.18a$, $i = 85.0^\circ$, $q = 0.25$, $M_1 = 0.9$, $\beta = 30^\circ$, $\zeta = 10^\circ$ and a normalized brightness of 0.23 for the accretion region. The stream has a total of 127 emission points along its length. The size of the phase bins is $\mathrm{d}\phi = 0.0001$: this corresponds to an integration time of $\sim 0.72\,\mathrm{s}$ in a binary with a 2 h orbital period (the effect of using larger phase bins is described in section 3.2.1).

The input stream is shown in the top panel of Fig. 1 projected onto the plane perpendicular to

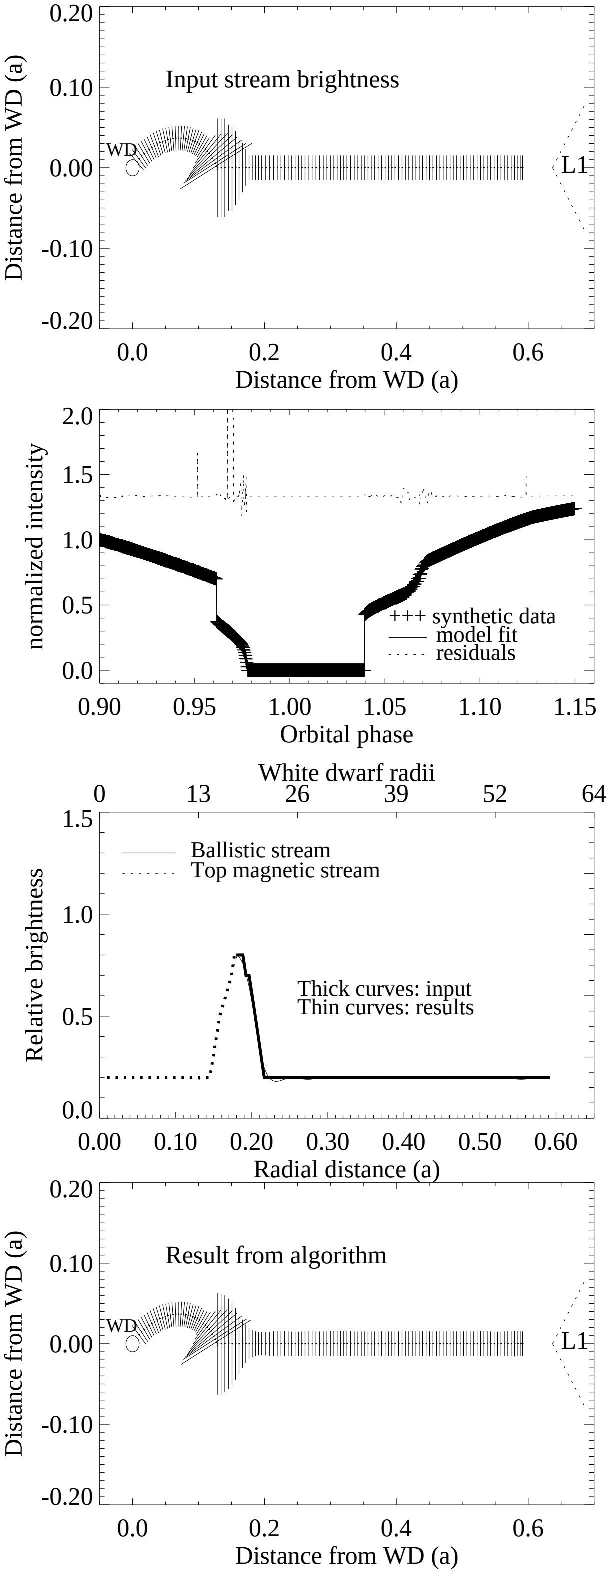

Figure 1. Simulation for a stream accreting onto one footprint of a (dipole) magnetic field line. The details of the figure are described in the text in section 3.1.1.

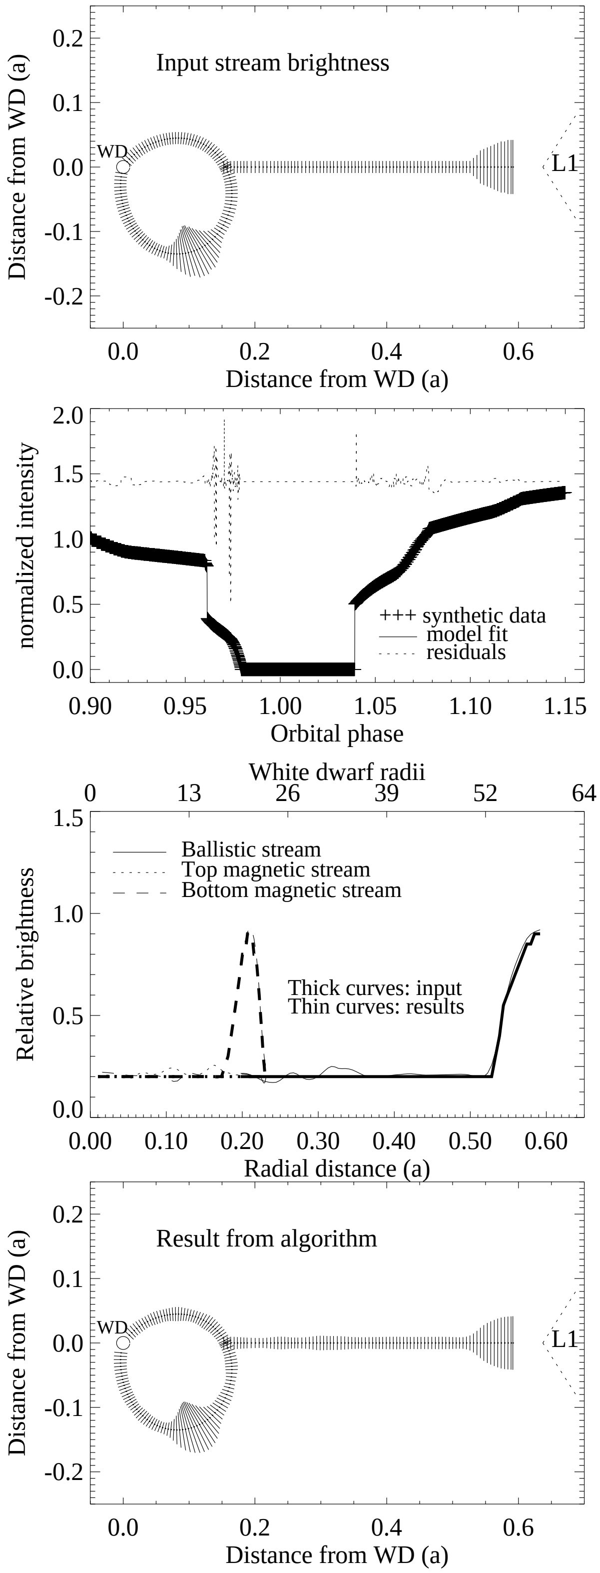

Figure 2. As for Fig. 1 but for a stream accreting onto both footprints of a dipole magnetic field line. The residuals in the second panel are plotted on a scale of -15 to 15 and are shifted vertically for clarity.

the orbital plane and passing through the centre of both stars. The brightness of each emission point is shown as a line through the point, perpendicular to the stream; the length of each line indicates the brightness of the point. The emission points along the magnetically-confined part of the stream appear more closely spaced than those on the ballistic stream due to the projection of the curved stream onto a flat plane; in reality the points are evenly spaced along the trajectory. The white dwarf is shown to scale as a circle (labelled 'WD'), and the secondary is shown (not to scale) to mark the position of the L1 point (labelled 'L1' on the figure). The second panel shows the synthetic light curve ('+' signs) and the model fit superimposed as a solid line. The uncertainties σ_i required in equation 1 are calculated using $\sigma_i = \mathrm{model}_i/SNR$, where SNR is an assigned signal-to-noise (we used $SNR = 1000$ for this ideal test case). No artificial noise is added to the synthetic light curve. The fit has $\chi^2/N = 0.13$. The dotted line shows the residuals $(\mathrm{model}_i - \mathrm{data}_i)/\sigma_i$ on a scale of -15 to 15, shifted vertically by $+7$ for clarity. The third panel of Fig. 1 shows the relative brightness of the stream as a function of radial distance from the white dwarf centre. The input values are shown as thick lines and the results are the superimposed thin lines. The model stream is shown in the bottom panel of Fig. 1 in the same projection as in the top panel. The main features of the input model are reproduced well by the algorithm. In particular, the position, width and amplitude of the feature at R_μ are reproduced. The third panel of Fig. 1 shows that the artefacts introduced by the algorithm (e.g. around $0.23a$) are at a very low level ($\lesssim 3$ per cent); it is difficult to distinguish between the input and the model stream brightness distributions.

3.1.2 Two-pole geometry

The simulated light curve is constructed using a synthetic stream with a brightening towards the L1 point and a brightness enhancement on the magnetic field line below the orbital plane (see Fig. 2). We chose this test case to check whether the method is capable of reconstructing a brightening near the L1 point which might arise as a result of irradiation of the secondary (as is observed in e.g. HU Aqr in its high state: Schwope et al. 1997). This test example also allows us to determine whether the method can reproduce bright regions on one of the two halves of the magnetically-channeled stream, and thus locate bright regions within the magnetosphere. The parameters used in the model are $R_\mu = 0.20$, $i = 85.0^\circ$, $q = 0.25$, $M_1 = 0.9$, $\beta = 25^\circ$, $\zeta = 10^\circ$, a normalized brightness of 0.23 for the upper accretion region and 0.15 for the lower. The stream has a total of 197 emission points along its length. As in the previous simulation, the data have $\mathrm{d}\phi = 0.0001$, and a signal-to-noise ratio of 1000. No noise is added to the synthetic light curve.

The model fit to the synthetic data is shown in the second panel of Fig. 2 and has $\chi^2/N = 0.51$. The model stream brightness distribution is shown in the bottom panel of Fig. 2: the position, amplitude and width of the main features of the input stream are reproduced well by the model. The artefacts introduced by the optimizing procedure are at a slightly higher level ($\lesssim 6$ per cent) than in the one-pole geometry simulation: the reason for this is explained in section 3.2.1.

3.2 The effect of noise and decreased phase resolution

For a given stream trajectory and a fixed number of emission points along the stream, there are two factors that affect the artefact level in the stream brightness distributions. These are the noise level in the light curve, and the phase resolution of the light curve. First we discuss the effect of reducing the phase resolution in the noiseless case, and then show the effect of noise in light curve with a high phase resolution. Finally, we discuss the results from a noisy synthetic light curve at a reduced phase resolution.

3.2.1 Reduced phase resolution

For a fixed number of emission points along the stream, the artefact number and amplitude is increased when the phase resolution is decreased. Fig. 3 contrasts the results for the two-pole stream example in section 3.1.2 at reduced phase resolutions $\mathrm{d}\phi$ of 0.0005 and 0.001 ($\mathrm{d}\phi = 0.0005$ corresponds to $\sim 4\,\mathrm{s}$ and $\mathrm{d}\phi = 0.001$ to $\sim 8\,\mathrm{s}$ in a binary with $P_{\mathrm{orb}} = 2\,\mathrm{h}$). These results should be compared to the results in Fig. 2 where the light curve has $\mathrm{d}\phi = 0.0001$. No noise is added to the light curves, and the data have a signal-to-noise of 1000. The fit to the data with $\mathrm{d}\phi = 0.0005$ has $\chi^2/N = 6$, while the fit in the $\mathrm{d}\phi = 0.001$ case is much worse, with $\chi^2/N = 85$. The model with $\mathrm{d}\phi = 0.0005$ has located the bright regions of the input stream and has retrieved their amplitudes and shapes, but there is an excess of brightness around $0.33a$ (labelled 'c' in the top right-hand panel of Fig. 3) and a deficit on the lower half of the magnetic trajectory just below the junction between the ballistic and the magnetic trajectories (labelled 'b' on the top right-hand panel of Fig 3). With $\mathrm{d}\phi = 0.001$ the artefact problem is worse: the model has located the two bright regions, but their shapes are distorted and their amplitudes are reduced. In addition, there is an excess of brightness on the ballistic stream near $0.33a$ (labelled 'c' on the bottom right-hand panel

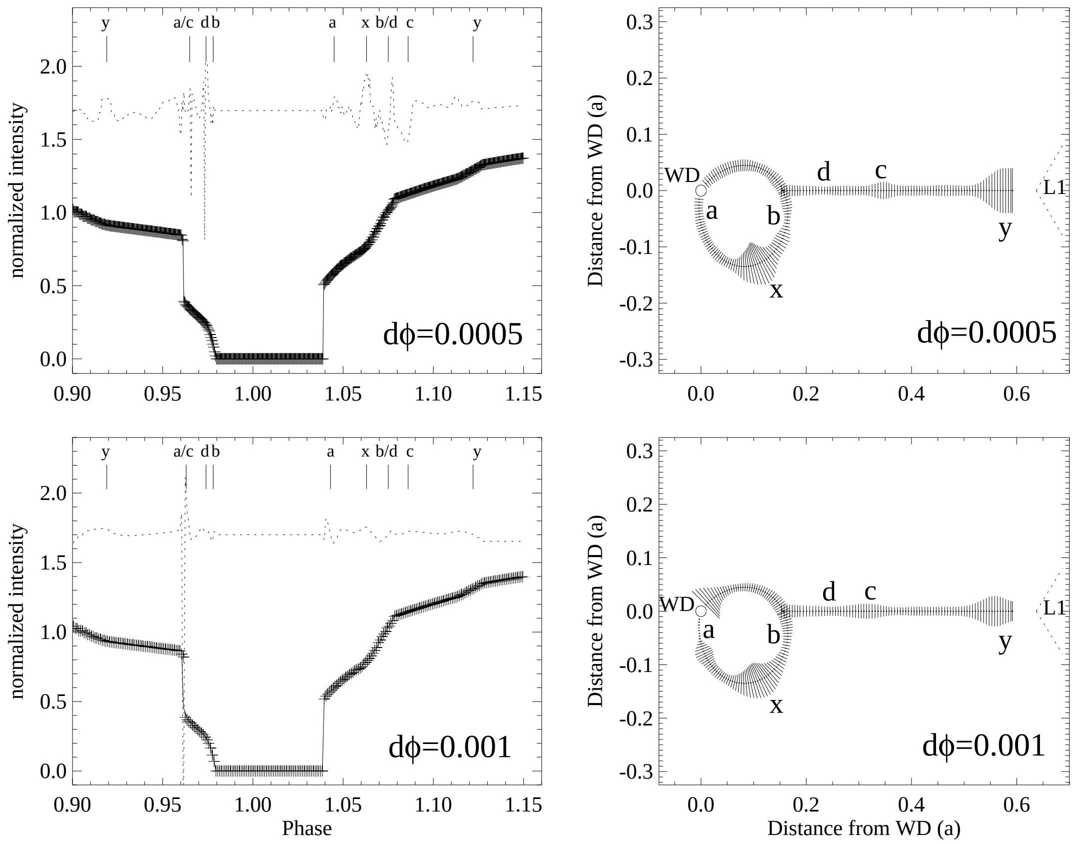

Figure 3. The effect of lowering the phase resolution of the light curve is shown for phase resolutions dϕ of 0.0005 (top) and 0.001 (bottom). As the phase resolution is decreased, the number and amplitude of the artefacts in the stream brightness distribution increases. The residuals are plotted on a scale of -30 to 30 (displaced vertically by +15 for clarity) for the dϕ = 0.0005 case, and -300 to 300 (displaced by +150) for the dϕ = 0.001 case. The regions of the stream labelled 'a', 'b', 'c', 'd', 'x' and 'y' are discussed in the text; the lines above the light curves indicate the approximate phases at which the stream points around 'a', 'b', 'c', 'd', 'x' and 'y' are eclipsed or come into view.

of Fig. 3), and the region near the white dwarf is badly misrepresented.

The artefacts in the brightness distributions, the distortion of the shapes of the bright regions on the stream, and the larger of the residuals in the fits can be explained by examining how well the model is able to resolve certain critical parts of the stream. These critical parts lie just below the orbital plane near the lower pole (labelled 'a' on the right-hand panels of Fig. 3), and just below the junction between the ballistic and the magnetic trajectories (labelled 'b'). The group of points at 'a' lies parallel to the limb of the secondary during eclipse ingress at phase $\phi \sim 0.965$, and the group of points at 'b' is parallel to the limb of the secondary during eclipse egress at $\phi \sim 1.075$. At these phases, several adjacent points at 'a' or 'b' are eclipsed or come into view in a single phase bin. As a result, the model is not able to assign brightnesses unambiguously to these points. Errors are propagated into the points on the ballistic stream that are eclipsed or that come into view in the same phase bin (in order to preserve the χ^2 of the model at that phase). For example, at $\phi \sim 0.965$ as the points at 'a' are eclipsed, the points along the ballistic trajectory at $\sim 0.32a$ (labelled 'c' on the right-hand panels of Fig. 3) are also being eclipsed. The points on the ballistic stream that come into view when the points at 'b' egress, are at $\sim 0.23a$ (labelled 'd' in Fig. 3). The brightnesses of emission points at 'c' and 'd' are thus not well-determined, and this accounts for the artefacts on the model streams at these positions (an excess of brightness at 'c' and diminished flux at 'd'). We expect the largest residuals in the fit to occur where the points 'a','b','c' and 'd' ingress and come into view: these phases are indicated by the relevant letters above the light curves in Figure 3. The artefacts on the low resolution $d\phi = 0.001$ model are particularly severe around the white dwarf, and the largest residuals occur in the light curve between $0.961 \lesssim \phi \lesssim 0.964$. This phase interval comprises just four phase bins during which 16 adjacent points just below the lower pole are eclipsed, the upper pole is eclipsed, and five points just above the upper pole are eclipsed. This large number of poorly-resolved points around the white dwarf leads to the spurious structure in the stream brightness distribution in this region, and hence to the poor fit to the light curve in the interval $0.961 \lesssim \phi \lesssim 0.964$.

The artefacts in the model stream are created as the algorithm attempts to preserve the χ^2 at a particular phase, but they tend to increase the *entropy* of the model light curve: this is because the knots of excess or diminished flux in the artefacts reduce the local smoothness of the stream brightness. Because the algorithm minimizes the entropy of the brightness distribution as well as the χ^2 of the fit, the increase in entropy caused by the artefacts must be compensated for by a *decrease* in entropy elsewhere on the stream. This occurs at the bright region on the field line below the orbital plane (labelled 'x' on the model streams). In the input stream (shown in the top panel of Fig. 2) this feature extends over 19 emission points and has a sharply peaked shape. In the low-resolution models, the feature is spread over larger numbers of emission points, and is increasingly distorted (flattened) as the phase resolution is decreased. We thus expect large residuals in the model fits where this bright region is eclipsed and comes into view: these phases are labelled 'x' in Figure 3 (the ingress of 'x' is not labelled due to lack of space: it occurs between the ingress of features 'd' and 'b'). The algorithm also smooths and spreads out the feature near the L1 point in order to reduce the entropy of the model stream: the ingress and egress of this feature are labelled 'y' in Fig. 3.

In order to test our understanding of the origin of the artefacts, we performed a fit to eclipse profiles that were generated using a stream where the emission points have a uniform brightness. We expect the artefacts to disappear in this case, irrespective of the phase resolution. This is because the algorithm should not be able to compensate for the increase in entropy (caused by artefacts) by smoothing other regions of the stream, since the stream brightness distribution is already the smoothest possible. Indeed, we find that for a light curve with the lower phase resolution ($d\phi = 0.001$), the algorithm does not produce artefacts on the model stream.

Lowering the phase resolution of the light curve thus has two main consequences: the misrepresentation of the stream brightness at ambiguous points along the stream, and the smoothing of local extrema in the stream brightness distribution to compensate for the resulting increase in entropy. Comparing the results of the $d\phi = 0.0005$ and $d\phi = 0.001$ models, we note that the quality and reliability of the model decreases dramatically if $d\phi \gtrsim 0.001$ (see also section 3.2.3).

When applying the model to real data, it is possible to distinguish between artefacts and real brightenings along the stream. Once a model stream trajectory is found that produces acceptable fits to the original data, the ambiguous points in that stream can be identified. An analysis similar to the one above can then be performed to identify the points on the stream that are likely to have poorly-determined brightnesses, and thus the regions of the stream that are likely to show artefacts.

The precise location of the ambiguous points on the stream depends on the system geometry i.e. the values of M_1, q, i, β, ζ and R_μ. For the stream trajectory assumed in section 3.1.1, the one-pole stream has no points that lie parallel to the limb of the secondary during ingress or egress. This explains why the artefact level in the one-pole case is much lower than in the two-pole example in section 3.1.2, even though the two light curves have the same phase resolution and neither have added noise.

3.2.2 Noisy data at high phase resolution

We construct a noisy light curve by multiplying the normalized intensities of the synthetic light curve by a constant to convert them to 'counts' and then adding a background 'sky' value to each point. The sky contribution is assumed to be a constant fraction of the mean intensity of the light curve. The level of noise added to each point depends on its location in the eclipse profile. One of the dominant sources of noise in a polar light curve is the flickering from the accretion region; the phase intervals most affected by flickering are thus before the accretion region ingress and after its egress. The accretion stream ingress and the points during totality are not affected (since the accretion region is obscured during these phases) and therefore have a much lower noise level. For this reason, we have assigned lower noise levels to the phases where the accretion region is obscured than when it is in view. The noise level for phases where the accretion region is in view is calculated using

$$\mathrm{noise}_i = GRND \times \sqrt{\mathrm{model}_i + \mathrm{sky}} \qquad (4)$$

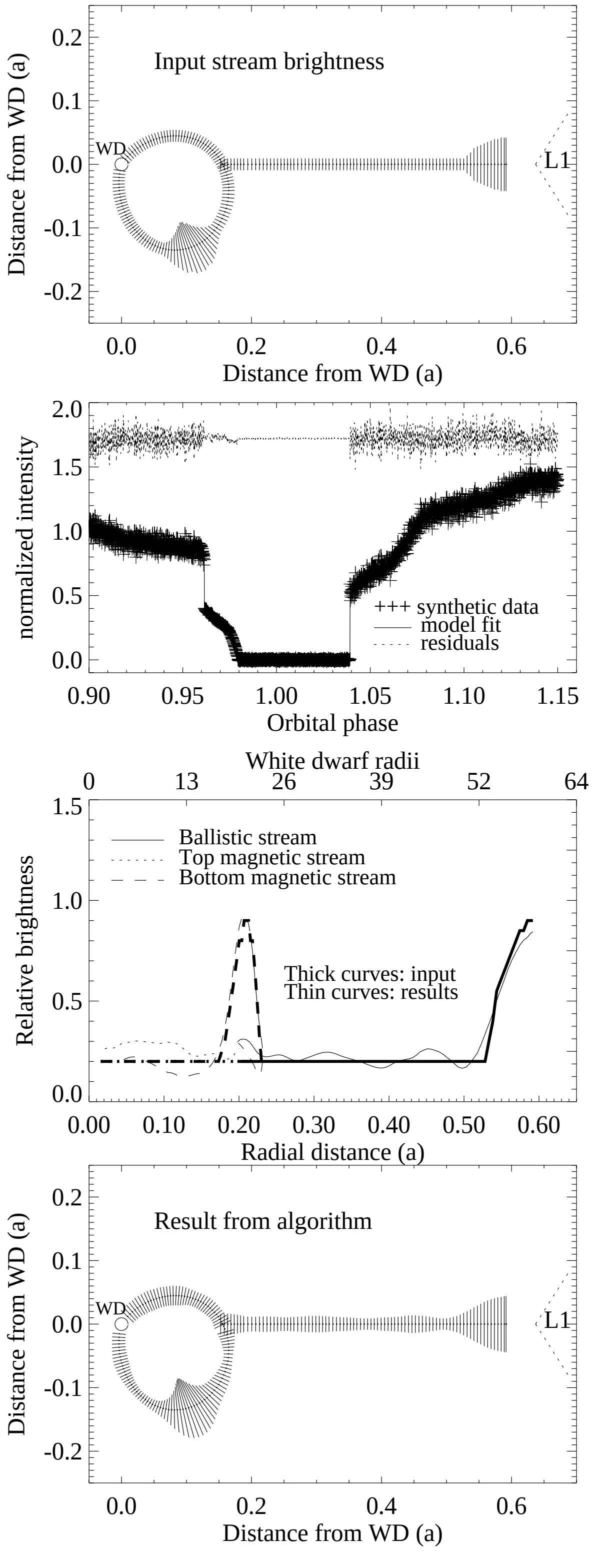

Figure 4. As for Fig. 2 but for a light curve with added noise (see text for details). The signal-to-noise ratio out of eclipse is approximately 20, and the phase resolution is $d\phi = 0.0001$. The residuals for the fit (second panel, dotted line) are plotted on a scale of -15 to 15 and are shifted vertically by $+11$ for clarity.

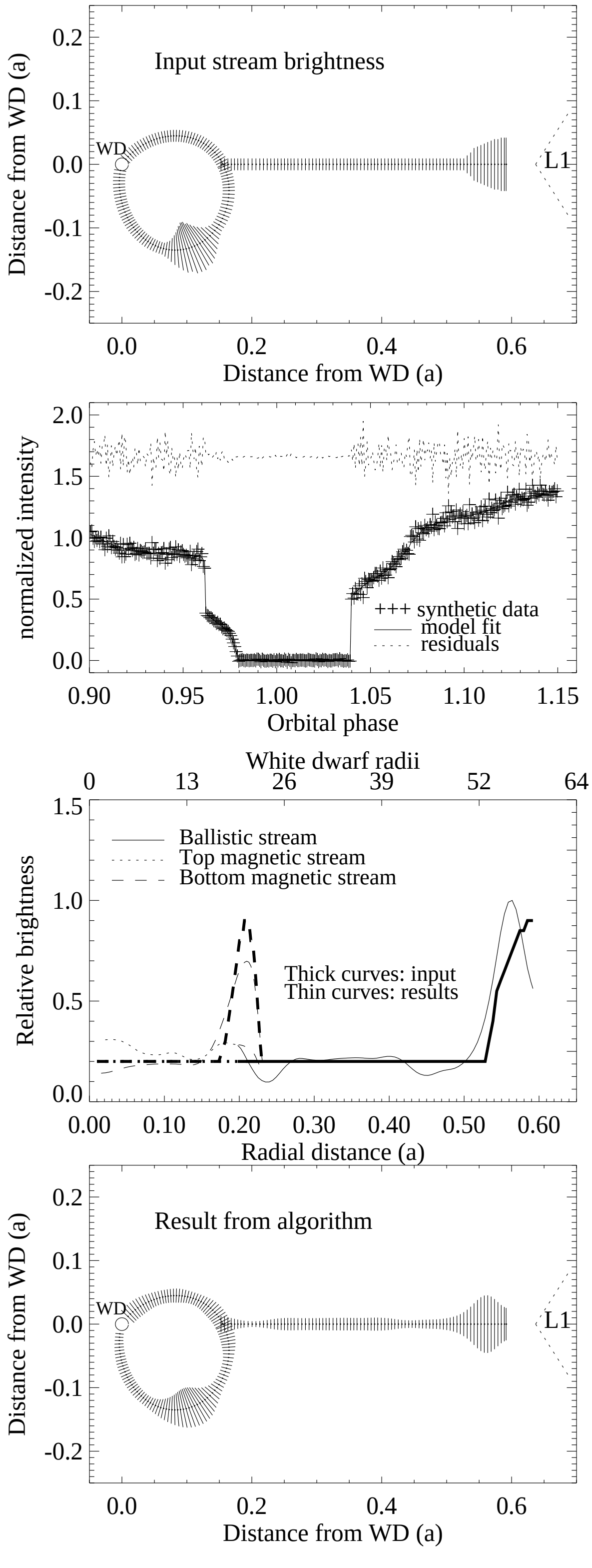

Figure 5. As for Fig. 2 but for a light curve with added noise and reduced phase resolution (see text for details). The signal-to-noise ratio out of eclipse is approximately 20, and the phase resolution is $d\phi = 0.0007$. The residuals in the second panel (dotted line) are plotted on a scale of -10 to 10 and have been shifted vertically by $+6$ for clarity.

where $GRND$ is a random number taken from a Gaussian distribution with zero mean and unit variance; the noise level for the phases where the accretion region is obscured are assumed to be one quarter this quantity. Once the appropriate noise level is added to each model point, the light curve is re-normalized so that the pre-eclipse level is 1.0 as before.

Fig. 4 shows the results derived from a light curve with a phase resolution of $\mathrm{d}\phi = 0.0001$ and with added noise such that the 'sky' background is equal to the mean intensity. The resulting signal-to-noise is ~ 20 out of eclipse and about 8 mid-way through the stream ingress. The fit to the input light curve has $\chi^2/N = 0.7$. There are some discrepancies between the input and the model streams. The model reproduces accurately the amplitude of the bright region below the orbital plane, but underestimates the brightness at L1; the two peaks differ in brightness by $\sim$ 10 per cent while in the input stream they have the same relative brightness. The bright region near the L1 point is $\sim$ 20 per cent wider than in the input stream. The number of artefacts has increased and their amplitude is at the $\sim$ 10 per cent level. However, there are no artefacts in the model with amplitudes comparable to those of the bright peaks in the input stream. In general, the model has located accurately the bright regions on the stream, and has reproduced their shapes well.

3.2.3 Noisy data at reduced phase resolution

Fig. 5 shows the stream brightness distribution predicted for a light curve with $\mathrm{d}\phi = 0.0007$ (about 5 s in a binary with $P_{\mathrm{orb}} = 2\,\mathrm{h}$) and with added noise. The signal-to-noise ratio out of eclipse is ~ 20 and mid-way through the stream ingress is about 8. Despite the noise and the reduced phase resolution, the model fit to the light curve is good, with $\chi^2/N = 0.9$. The model stream brightness distribution in the fourth panel of Fig. 5 shows that the algorithm has located accurately the main features of the input stream, although their shapes and amplitudes are not exact. The shape of the bright region near L1 in the model has its peak brightness at $0.56a$ instead of at $0.58a$. The relative amplitudes of the two main features differ by 40 per cent, whereas they are the same in the input stream. The artefact level is 15 per cent, only slightly higher than in the noisy, high resolution case (section 3.2.2, Fig. 4). The algorithm has not added any spurious structure with amplitudes comparable to those of the bright regions in the input stream.

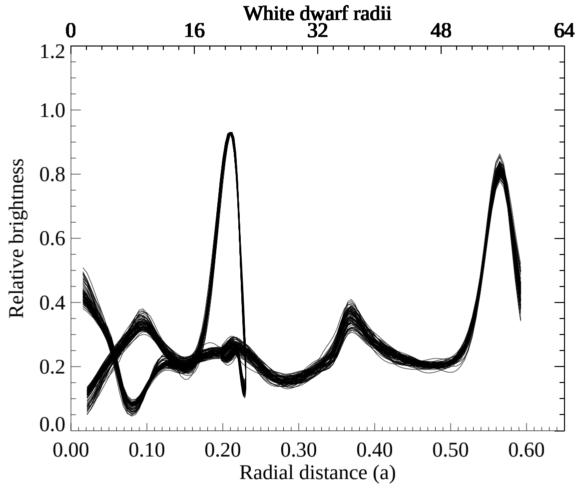

Figure 6. The stability of the solutions found by the optimizing algorithm. The plot shows brightness versus radial distance from the white dwarf for 70 solutions of the stream brightness distribution where each solution was generated using a different random initial population. This demonstrates that the algorithm finds the same minimum of equation 1, irrespective of the starting values for the solution parameters.

3.3 The stability of the solution

As mentioned in section 2.2, it is not possible to define a formal convergence criterion for our algorithm, since the GA does not improve the solutions at each stage by using gradient information. We can, however, demonstrate the stability of the solution found by the algorithm by performing a number fits to a given eclipse profile, each starting with a different random population. The results of this exercise are shown in Figure 6 for 70 fits to a set of noisy data. It can be seen that the algorithm finds essentially the same solution, irrespective of the initial parameter values. This is generally not the case with more traditional steepest descent-based optimization methods.

3.4 Ambiguity of the model geometry

When applying the method to real data, one may not know a priori whether to fit an eclipse profile with a model stream accreting onto one or onto two poles. We investigated whether, on the basis of the quality of the fit alone, it is possible to determine which of the two accretion geometries is the better approximation for a given eclipse profile. To do this, we fitted a synthetic light curve constructed from a two-pole stream using a model stream accreting only at one pole; and conversely, fitted a stream obtained from a one-pole geometry with a two-pole model.

Fig. 7 shows the results of applying a two-pole

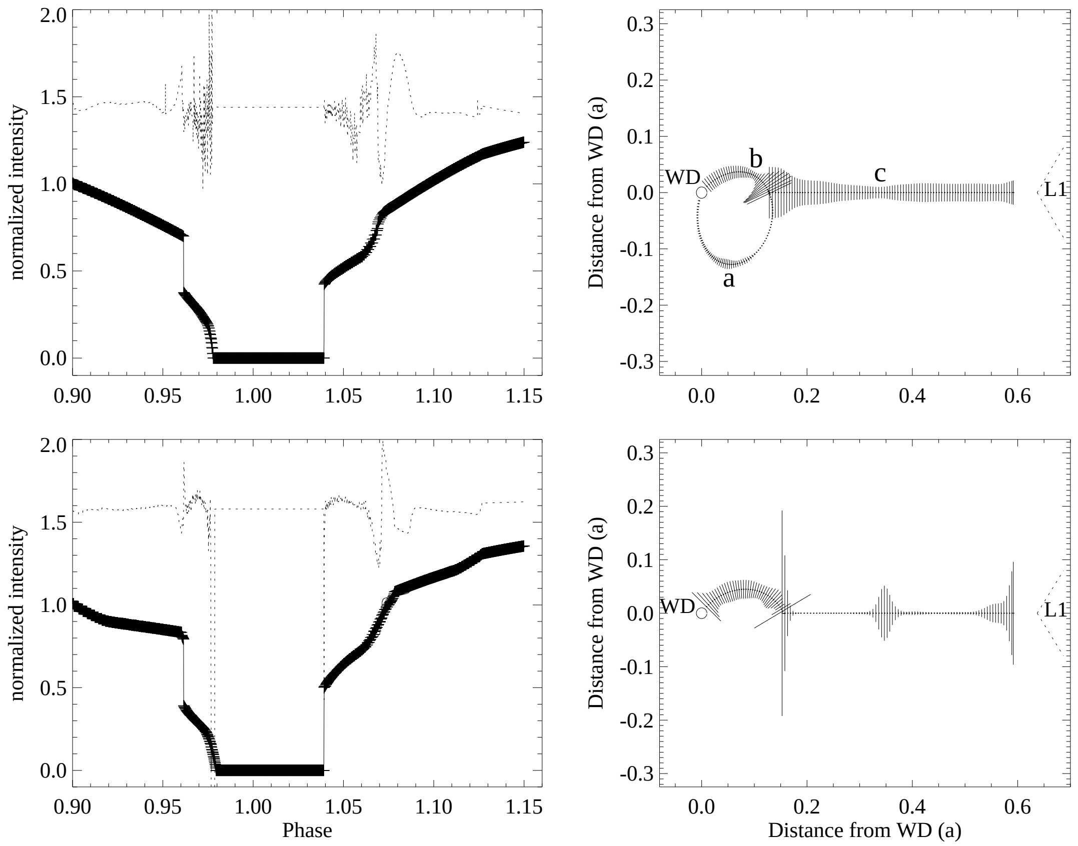

Figure 7. Top panel: fitting a synthetic light curve constructed from a stream accreting onto one pole, with a two-pole model stream. The model fit is on the left; the residuals of the fit are plotted on a scale of -30 to 30 and are displaced vertically by $+14$. The resulting stream brightness distribution is on the right. The regions of the stream labelled 'a', 'b' and 'c' are discussed in the text. Bottom panel: as for the top panel, but for a two-pole synthetic light curve fitted by a one-pole model. The residuals of the fit are plotted on a scale of -500 to 500 and are displaced by $+300$.

model to the single-pole light curve of section 3.1.1 (shown in the top panel of Fig. 1). We expect the model to produce a two-pole model stream with no brightness on the field line below the orbital plane. The two-pole model fit (top left panel of Fig. 7) has $\chi^2/N = 7.5$. The model stream (top right panel of Fig. 7) has placed the bulk of the brightness on the magnetic trajectory above the orbital plane, but there is some lower-level brightness on the field line below the orbital plane, particularly in the region marked 'a' on Fig. 7. This feature is eclipsed during the interval $0.970 \lesssim \phi \lesssim 0.975$ and comes into view again during $1.056 \lesssim \phi \lesssim 1.063$. During its ingress, the other points on the stream that are also being eclipsed are those on the upper field line just before R_μ (labelled 'b' on Fig. 7), and those on the ballistic stream around $0.32a$ (labelled 'c'). During the egress of the feature 'a', the only other points that are coming into view are the points on the upper magnetic trajectory at 'b' – the same points that are eclipsed with 'a' during $0.970 \lesssim \phi \lesssim 0.975$. There is a brightness deficit on the model stream at 'b', and also on the ballistic stream near 'c'. The reason for this is that the model cannot distinguish between the emission points at 'a' and the points at 'b', because they are eclipsed *and* come into view during the same phase intervals. The brightness below the orbital plane is taken from 'b' (and to a lesser extent from 'c' on the ballistic stream) to preserve the χ^2 during these phases. This accounts for the deficit of brightness in these regions. The artefact produced at 'a' increases the entropy of the solution, since the stream is now less locally smooth. To compensate for this, the algorithm has reduced the entropy of the solution by smoothing out the bright region at the junction between the ballistic and the magnetic trajectories.

We note that the points at 'a' are the only ones on the lower magnetic trajectory which are eclipsed *and* come into view during the same phase interval as one other region on the rest of the stream (in this case, the points at 'b'). These points can be identified in any model with a two-pole geometry, and any brightness assigned to the model stream in this region will be known to be uncertain.

The lower half of Fig. 7 shows the results obtained by applying a one-pole model to a two-pole light curve. We expect a poor fit to the light curve, since the light curve contains more information than the model we are using to fit it. The fit is shown in the bottom left-hand panel of Fig 7 and is exceptionally poor, with $\chi^2/N = 3130$. This is two orders of magnitude worse than any of the fits we have obtained using light curves with reduced phase resolution or added noise. The poor fit, in particular the large systematic errors in the fit around ϕ=0.978, 1.04 and 1.07, might prompt us to try the alternative two-pole model.

In principle therefore, on the basis of the model results alone, it is possible to determine whether to use a model with a two-pole or a one-pole geometry. However, in the examples discussed above, we have known in advance the correct values of M_1, q, i, β, ζ and R_μ to use. In practice the situation may be more complicated: there may be a combination of parameters that produce very poor fits even though the correct geometry is being used. In cases like this, other sources of information (e.g. polarimetry) can be used to substantiate the choice of the model geometry.

3.5 The choice of lambda

Fig. 8 shows the effect of using different values of λ in solutions computed using the noisy synthetic data in section 3.2.2 (the input stream is shown in the top panel of Fig. 4). These models illustrate the trade-off between minimizing χ^2 and minimizing the entropy by adjusting the value of λ. The stream image in the top panel of Fig 8 has $\lambda = 1.0 \times 10^6$, about an order of magnitude larger than necessary. With a large λ, the fitness function is dominated by the entropy term, and the algorithm attempts to reduce the entropy at the expense of χ^2. The input features of the stream are not reproduced accurately because they are local maxima in the stream image brightness distribution: the entropy in these parts of the stream is comparatively high and the algorithm smooths preferentially these parts of the stream. The solution has $\chi^2/N = 1.2$. The stream image computed using $\lambda = 1.0 \times 10^5$ is shown in the middle panel: this value of λ is appropriate for these data (this is the solution presented in Figure 4). The solution has $\chi^2/N = 0.7$, much better than the solution obtained using the large value of λ. The solution with too small a value of λ (bottom panel of Fig.8) has $\lambda = 1.0 \times 10^4$, about an order of magnitude smaller than required. The solution thus follows the original data (including the noise in the data) very closely, and has $\chi^2/N = 0.5$. The features of the input stream are discernible, but there are many artefacts in the stream distribution: these are the result of the model fitting the noise in the data as well as the original structure in the stream.

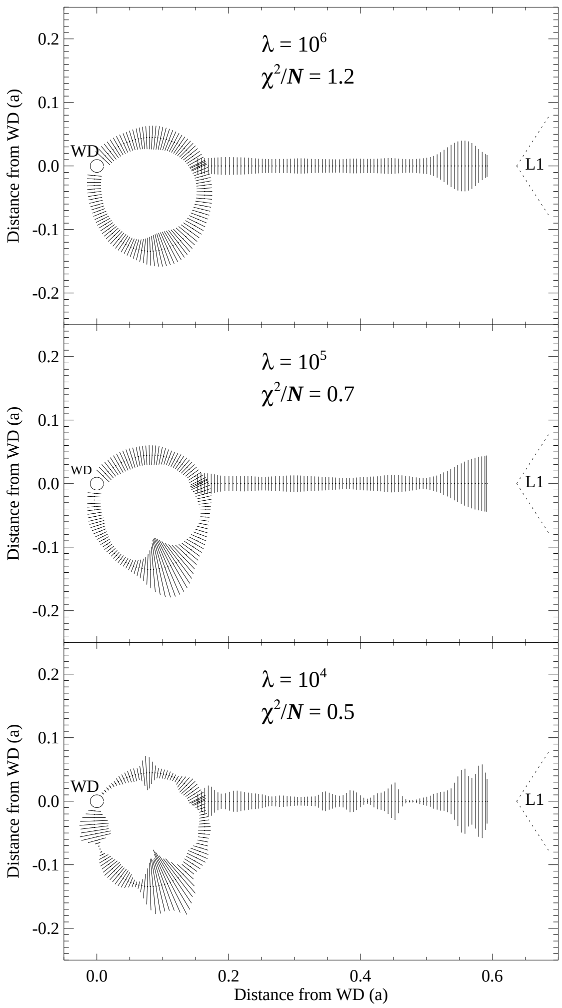

Figure 8. The effect of the Lagrange multiplier λ in solutions computed using the noisy synthetic data shown in Figure 4. The top panel shows the stream image computed using a value of λ an order of magnitude too large: the entropy of the solution is minimized at the expense of χ^2. The lower panel shows a solution with λ an order of magnitude too small: the algorithm fits the noise in the data at the expense of the smoothness of the solution. The middle panel shows the solution calculated using an appropriate value of λ.

4 DISCUSSION AND CONCLUSIONS

4.1 Information obtainable from the stream brightness distributions

The method can retrieve the features of the stream from an eclipse profile. The method performs well in the presence of noise, but requires a high phase resolution to reduce the artefact level in the model streams. The requirement for a high sampling rate may restrict the application of this method to eclipse profiles observed with photomultipliers, as opposed to CCDs. The method is able to reconstruct features such as a brightening of the stream towards the L1 point (which may be expected if the secondary is strongly irradiated by the primary in a high accretion state), a bright region where the stream becomes threaded onto the field lines (as suggested by e.g. Lamb 1985 and Hameury et al. 1986) and can identify regions of enhanced emission at specific locations within the magnetosphere of the white dwarf. The stream brightness distributions can therefore be used to to extract information on such structures from real data, and thus provide observational constraints for the properties of the accretion stream.

Given the complexities of the accretion stream–magnetosphere interaction, it is unlikely that the accretion stream would follow precisely the trajectory that we have used in our model. Although our models are unable to place constraints on the three-dimensional size of the threading region, they *can* show whether there is a bright region where we expect threading to occur, and this would allow us to place constraints on the extent of this region along the direction of our pre-set trajectory. This may then provide evidence for the existence of a magnetospheric cavity containing a pool of hot plasma, as predicted by Lamb (1985).

We have assumed a dipole magnetic field for the white dwarf. In some systems, the inclusion of higher order multipoles may be important to specify the stream trajectory (especially near the white dwarf). Higher order multipoles could easily be included in the model if warranted by the data.

Once an estimate of R_μ has been obtained from the model, an estimate of the mass transfer rate through the stream can be obtained by using

$$R_\mu = 1.7 \times 10^{10} B_{30}^{4/11} R_9^{12/11} r_9^{4/11} M_1^{-1/11} \dot{M}_{16}^{-2/11} \quad (5)$$

(Mukai 1988) where B_{30} is the polar magnetic field strength in units of 30 MG, R_9 is the white dwarf radius (as given by equation 2), r_9 is the radius of the stream in units of 10^9 cm, and $\dot{M}_{16}$ is the mass transfer rate in units of $10^{16}\,\mathrm{g\,s^{-1}}$. This equation assumes that the incoming stream does not distort the magnetic field of the white dwarf and that threading of the stream by the field is instantaneous. The radius r_9 of the stream can be obtained from Lubow & Shu (1975) or from measurements of the width of the pre-eclipse dip in systems that show this feature (as in Watson et al. 1995).

Our method in its current form is applicable to systems where the component of the eclipse due to the accretion stream can be distinguished from that due to the accretion region, and where the accretion stream contributes significantly to the overall flux. These include HU Aqr (e.g. Hakala et al. 1993, Schwope et al. 1997), V2301 Oph (e.g. Barwig, Ritter & Bärnbantner 1994) and UZ For in its high state of accretion (Bailey 1995). Our application of the method to *UBVR* eclipses of HU Aqr will be the subject of a subsequent paper.

4.2 Future improvements to the method

A straightforward way to improve the method would be to use a more realistic stream trajectory. We have assumed that the stream makes an instantaneous transition from a ballistic to a magnetic trajectory at R_μ. This is unlikely to be the case, since the dynamic timescale for dense blobs in the threading region may be shorter than the magnetic timescale on which they become coupled to the field (Hameury et al. 1986). The stream-field coupling is essentially a damping of the motion of the blob perpendicular to the stream, and can be expressed as a drag force which is proportional to the cross-field velocity and to the square of the local field strength. The trajectory of the blob can be calculated as it moves under the influence of gravity and the drag force (King 1993, Wynn & King 1995). Schwope et al. (1997) make use of a drag term to calculate their model stream trajectories, although they do not describe the precise form of the drag term used.

In principle, our method could be modified so that we remove entirely the need for a pre-set stream trajectory. A three-dimensional grid could be set up between the two stars with an emission point at each node of the grid. The number of emission points M in the model would be several orders of magnitude greater than in the current models, and the time taken to compute solutions would be substantially longer than in the quasi-one-dimensional case, since the time taken to evaluate the fitness function for each solution increases dramatically with M. Formally, the problem would not require additional constraints other than the ME regularization. In practice, however, because many points in the grid would be eclipsed at each phase step, the results may be much more artefact-prone than those from the quasi-one-dimensional model. The results would be more difficult to interpret than those presented here.

ACKNOWLEDGMENTS

We would like to thank Keith Horne, Nick Achilleos, Jon Mittaz and Kinwah Wu for useful advice. MKHA acknowledges funding support from University College London (Graduate Schools Research Scholarship), the Committee of Vice Chancellors and Principals (Overseas Research Students' Award) and the University of Cape Town (Max & Lillie Sonnenberg Scholarship). PJH wishes to thank MSSL for their hospitality and making his visit most enjoyable. PJH is supported by an Academy of Finland Research Fellowship.

References

Bailey J., 1995, in Buckley D.A.H., Warner B., eds, ASP Conf. Ser. 85, Cape Workshop on magnetic cataclysmic variables. ASP, California, p. 10

Barwig H., Ritter H., Bärnbantner O., 1994, A&A, 288, 204

Beasley D., Bull D.R., Martin R.R., 1993, Univ. Computing, 15, 58

Charbonneau P., 1995, ApJSS, 101, 309

Cropper M., Horne K., 1994, MNRAS, 267, 481

Hakala P.J., 1995, A&A, 296, 164

Hakala P.J., Watson M.G., Vilhu O., Hassall B.J.M., Kellett B.J., Mason K.O., Piirola V., 1993, MNRAS, 263, 61

Hameury J.M., King A.R., Lasota J.P., 1986, MNRAS, 218, 695

King A.R., 1993, MNRAS, 261, 144

Lamb D.Q., 1985, in Lamb D.Q., Patterson J., eds, Cataclysmic variables and low-mass X-ray binaries. Reidl, Dordrecht, p. 179

Lamb D.Q., 1988, in Coyne G.V. et al., eds, Polarized radiation of circumstellar origin. Vatican Observatory, p. 151

Liebert J., Stockman H.S., 1985, in Cataclysmic Variables and Low-Mass X-ray Binaries, eds. Lamb D.Q., Patterson J., Reidl, Dordrecht, p 151

Lubow S.H., Shu F.H., 1975, ApJ, 198, 383

Marsh T.R., Horne K., 1988, Adv. Space Res., 8, (2)127

Mukai K., 1988, MNRAS, 232, 175

Potter S., Hakala P.J., Cropper M., 1998, MNRAS, in press

Press W.H., Teukolsky S.A., Vetterling W.T., Flannery B.P., 1992, Numerical Recipes in FORTRAN. Cambridge Univ. Press, Cambridge

Ramsay G., Mason K.O., Cropper M., Watson M.G., Clayton K.L., 1994, MNRAS, 270, 692

Rosen S.R., Mason K.O., Cordova F.A., 1987, MNRAS, 224, 987

Schwope A.D., Mantel K.H., Horne K., 1997, A&A, 319, 894

Warner B., 1995, Cataclysmic variable stars. Cambridge Univ. Press, Cambridge

Watson M.G., Rosen S.R., O'Donoghue D., Buckley D.A.H., Warner B., Hellier C., Ramseyer T., Done C., Madejski G., 1995, MNRAS, 273, 681

Wynn G.A., King A.R., 1995, MNRAS, 275, 9