Mon. Not. R. astr. Soc. (1985) **213,** 129–141

Images of accretion discs – I. The eclipse mapping method

Keith Horne *Institute of Astronomy, Madingley Road, Cambridge CB3 0HA*

Accepted 1984 October 17. Received 1984 September 27; in original form 1984 July 9

Summary. A method of mapping the surface brightness distributions of accretion discs in eclipsing cataclysmic binaries is described and tested with synthetic eclipse data. Accurate synthetic light curves are computed by numerical simulation of the accretion disc eclipse, and images of the disc are reconstructed by maximum entropy methods. The conventional definition of entropy leads to a distorted image of the disc, from which only the locations of compact bright spots are reliably determined. A modified form of entropy, sensitive to the azimuthal structure of the image but not to its radial profile, suppresses azimuthal structure but correctly recovers the radial structure of the accretion disc. This eclipse mapping method permits powerful tests of accretion disc theory by deriving the spatial structure of discs from observational data with a minimum of model-dependent assumptions. Mass transfer rates for eclipsing cataclysmic variables can be measured by this method with an uncertainty of about a factor of two.

1 Introduction

Cataclysmic variables are short-period binary systems containing a white dwarf that collects matter escaping from a red-dwarf companion star that overflows its Roche-lobe. In most cases the accretion flow takes the form of a physically thin disc in the plane of the binary orbit. Rapid conversion of gravitational potential energy into heat causes the accretion disc to outshine other sources of light in the system at most wavelengths. In high-inclination systems, occultation of the accretion disc by the red dwarf produces a deep eclipse. This eclipse offers a unique opportunity to determine the spatial structure of the disc directly from observations.

The shape of the eclipse light curve is mainly determined by the brightness distribution on the face of the disc. The brightness distribution cannot be uniquely recovered from the eclipse data, however, since a one-dimensional light curve does not fully constrain a two-dimensional image. Many different images of the disc produce essentially identical eclipse light curves. This indeterminacy may be treated by confining the image to the parameter space of a specific model. The model parameters are then adjusted to obtain the best possible fit to the eclipse data. This model-fitting approach is taken by e.g. Frank & King (1981), who fit a steady-state blackbody accretion-disc model to visual and infrared light curves of the eclipsing nova-like variable RW Tri.

The model-fitting approach is limited by our present incomplete knowledge of severa important aspects of accretion disc theory. The theory treats only vertically-averaged propertie of discs, since the vertical structure depends sensitively on the as yet unidentified viscosit mechanism. While energy budget considerations suffice to predict the bolometric surfac brightness (effective temperature) as a function of disc radius, the theory is ambivalent about th spectral form of the emitted radiation. We do not know, for example, if the optically-thick region of an accretion disc emit blackbody or stellar atmosphere spectra (Wade 1984). Strong emissio lines seen in the optical spectra of cataclysmic variables indicate that emission from optically-thi regions is also important.

A second problem is the assumption of axisymmetry. While axisymmetry is probably a goo assumption for the inner parts of the accretion disc, where the strong Keplerian shear shoul rapidly eliminate azimuthal structure (Pringle 1981), the tidal stresses exerted by the red dwar and the impact of the gas stream actively maintain azimuthal structure in the outer regions of th disc. Indeed most of the observed eclipse light curves are distinctly asymmetric. Parameter found by fitting a symmetric disc model to an asymmetric light curve are of uncertain value Asymmetry can be taken into account in an *ad hoc* way (e.g. Frank *et al.* 1981), but a large numbe of nuisance parameters is required to obtain good fits to high-quality eclipse data.

The concerns voiced above lead naturally to the image reconstruction approach developed i this paper. Image reconstruction is a limiting case of model-fitting, in which the intensity at eac point on the surface of the disc is an independent parameter. This enormous flexibility ensure that a good fit to the observed eclipse data is always achieved. Maximum entropy methods cop neatly with the indeterminacy by locating the unique, positive disc image that is consistent wit the eclipse data, and that maximizes the image 'entropy'. A flexible definition of the entrop allows for optimization of different image properties, uniformity, axisymmetry, and continuit being three examples. These methods provide a framework for studying the structure of accretio discs in a nearly model-independent way.

The method used to compute the light curve of an accretion disc undergoing eclipse, and th maximum entropy methods used to reconstruct disc images from a light curve are described i Section 2. Section 3 reports the results of simulation tests in which known accretion disc image are reconstructed from synthetic eclipse data. The simulation tests anticipate several problem that might arise in the application to real eclipse data. Section 4 summarizes the eclipse mappin method, and sketches the likely impact of eclipse mapping experiments on our understanding c the accretion discs in cataclysmic binaries.

2 Methods

2.1 LIGHT CURVE SYNTHESIS

Light curves for the eclipse of an accretion disc by its binary companion star are accuratel calculable by numerical simulation with a computer. The light curve synthesis program develope here assumes that the tidally distorted surface of the occulting star is given by the critical surfac of the dimensionless Roche potential:

$$\Psi=\frac{1}{1+q}\cdot\frac{1}{R_w}+\frac{q}{1+q}\cdot\frac{1}{R_r}-\frac{P^2}{2}, \qquad (1$$

in which R_w and R_r are distances (in units of the binary separation a) from the white-dwarf an red-dwarf centres of mass, P is the distance to the axis of rotation (intersecting the orbital plane a the centre of mass of the system), and $q=m_r/m_w$ is the binary mass ratio.

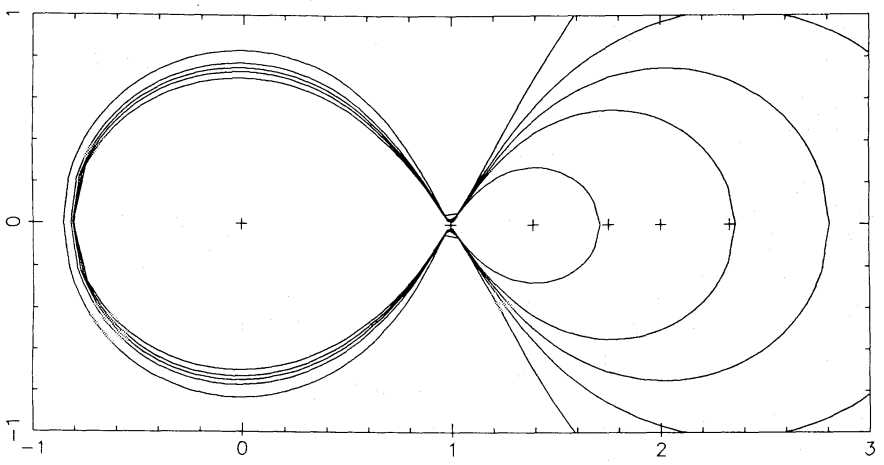

Figure 1. The Roche-lobes of a cataclysmic binary are shown in projection onto the orbital plane for mass ratios, 0.1, 0.5, 1, 2, and 10. The white dwarf is located at the origin, and the unit of distance is R_{L1}, the distance from the white dwarf to the inner Lagrangian point. This coordinate system is advantageous for eclipse mapping studies because the Roche lobe containing the accretion disc is nearly independent of the binary mass ratio.

Fig. 1 shows that the Roche-lobe occupied by the accretion disc is nearly independent of the binary mass ratio q if distances are measured in units of R_{L1}, the distance from the white dwarf to the inner Lagrangian point. As the mass ratios of cataclysmic variables are uncertain, I strongly recommend that R_{L1}, rather than the binary separation a, be adopted as the unit of distance in studies of accretion disc eclipses.

The accretion disc is assumed to be flat and to lie in the plane of the binary orbit. This flat disc approximation is not too restrictive because deep eclipses occur in cataclysmic binaries with inclinations as small as 70°, while the expected opening angles in their discs are only a few degrees. The assumption breaks down, of course, for systems with inclinations very close to 90°.

The accretion flow is assumed to be stationary in the frame that rotates synchronously with the binary. The surface of the disc is divided into a Cartesian grid of N^2 discrete elements that cover a square region centred on the white dwarf and bisected along one edge by the inner Lagrangian point. The Cartesian grid is computationally efficient, since each surface element subtends an identical solid angle, and permits the display of accretion-disc brightness distributions without interpolation.

The spatially integrated accretion disc flux f_ν at binary phase ϕ is calculated by summing contributions from all visible surface elements:

$$f_\nu(\phi)=\frac{\theta^2}{4N^2}\sum_{J=1}^{N^2} I_\nu(J)\cdot V(J,\phi). \tag{2}$$

Here $I_\nu(J)$ is the intensity along the line-of-sight to surface element J, and $V(J,\phi)$ is the fraction of element J that is visible at phase ϕ. The angular scale of the system is measured by the parameter

$$\theta^2=[R_{L1}/D]^2\cos(i), \tag{3}$$

where i is the binary inclination, and D is the distance of the system from Earth. As distances to cataclysmic variable stars are uncertain, one hopes to learn the value of θ from the eclipse analysis. Convenient units for θ are solar radii per kiloparsec (4.65×10^{-6} arcsec).

The occultation of disc elements by the red dwarf is accounted for in (2) by the visibility function $V(J,\phi)$, which normally is assigned a value of 1 or 0 depending upon whether the centre of element J is visible at phase ϕ. For elements that straddle the boundary of an occulted region, a

more accurate value is computed by refining the location of the boundary along two edges of the surface element.

To decide whether a given point on the disc is visible at a given phase, consider the ray originating at the point in question and directed toward Earth. Analytic tests first decide whether this ray intersects a cylindrical volume enclosing the companion star, and hence whether an occultation is possible. If the ray does intersect the cylinder, the Roche potential is evaluated at short intervals along the segment of the ray interior to the cylinder. An occultation is present if the Roche potential on this segment falls below the critical value defining the surface of the red dwarf. A similar algorithm was developed by Mochnacki (1971) to synthesize light curves of contact binary stars, and used by e.g. Young & Schneider (1981) to study eclipse phenomena in the emission lines of cataclysmic variables.

Three parameters fully specify the geometry of the eclipse: (1) the binary mass ratio $q=m_r/m_w$ which controls the relative sizes of the red-dwarf and white-dwarf Roche-lobes, (2) the phase of conjunction ϕ_0, and (3) the phase width of the eclipse at the centre of the disc $\Delta\phi$, which determines the inclination of the binary once the mass ratio is given. The relationship between the inclination and mass ratio for different values of $\Delta\phi$ is illustrated in Fig. 2. For deep eclipses, $\Delta\phi$ is a better parameter than the inclination because the former is well constrained by the observations while the latter is not. The inclination may of course be used for eclipse geometries in which the disc centre is not eclipsed.

Fig. 3 illustrates the eclipse geometry for the particular case $q=0.9$, $\Delta\phi=0.081$ (phase runs from 0 to 1, not 0 to 2π), which corresponds to an inclination near 75°. The red-dwarf and white-dwarf Roche-lobes are projected onto the orbital plane, and a network of ingress/egress arches is drawn on the accretion disc. Each arch outlines the disc region that is occulted at a particular binary phase. A small section of the accretion disc at the back of the Roche-lobe escapes the eclipse entirely.

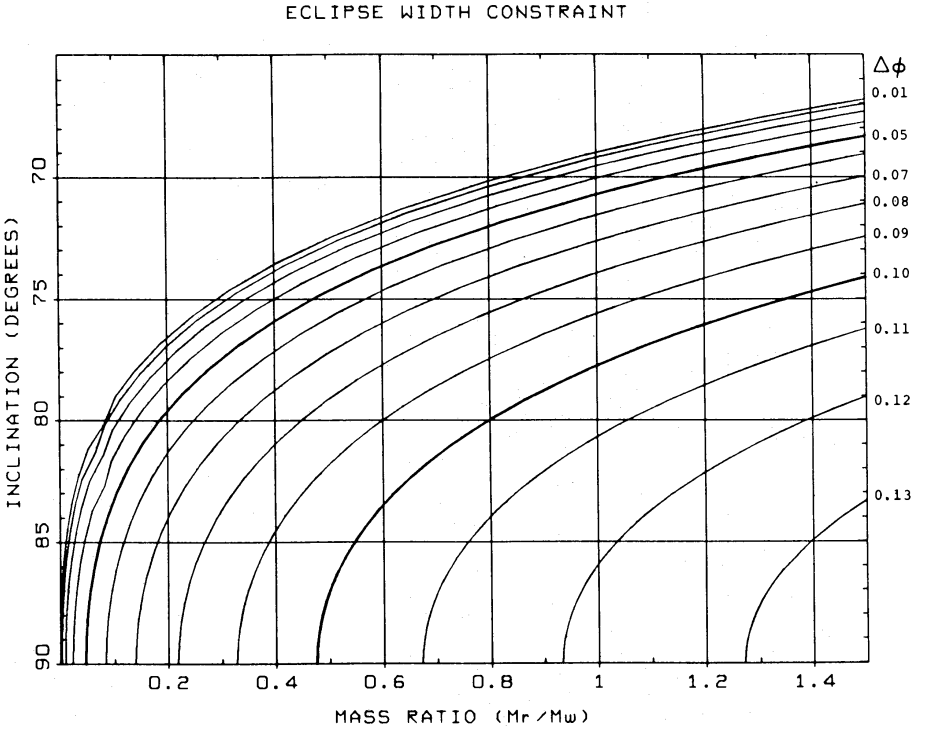

Figure 2. The relationship between the mass ratio q and the inclination i is shown for different values of the eclipse phase width $\Delta\phi$ at disc centre.

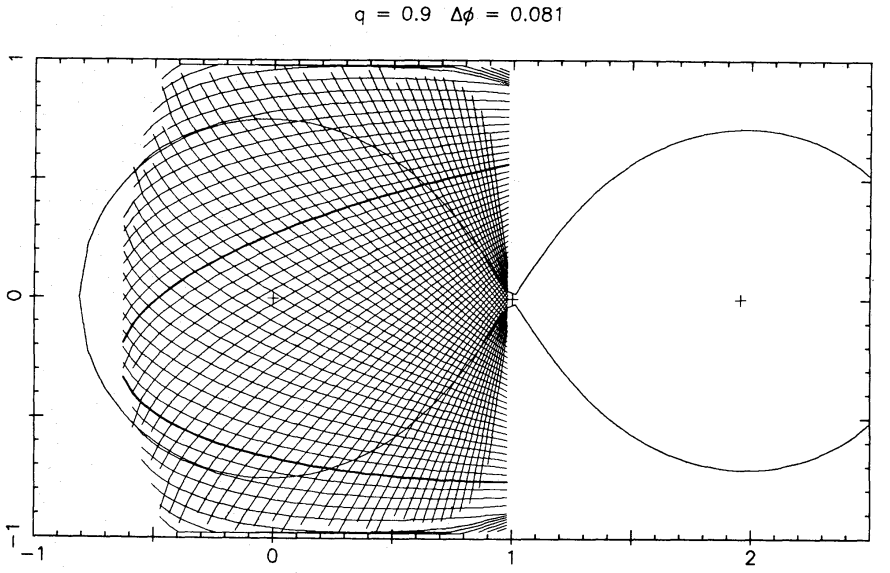

Figure 3. The geometry of an accretion disc eclipse is illustrated by a network of ingress/egress arches on the face of the disc. Each arch is the boundary of the disc region that is occulted by the red dwarf at a particular binary phase.

Examples of light curves computed by the above methods are shown in Figs 4 and 5. The eclipse geometry is the same as that in Fig. 3. Fig. 4 shows a deep but round-bottomed eclipse with broad asymmetric wings, similar to the observed eclipse curves of old novae, nova-like variables, and dwarf novae in outburst. The underlying disc intensity distribution, shown in Plate 1(a), is the

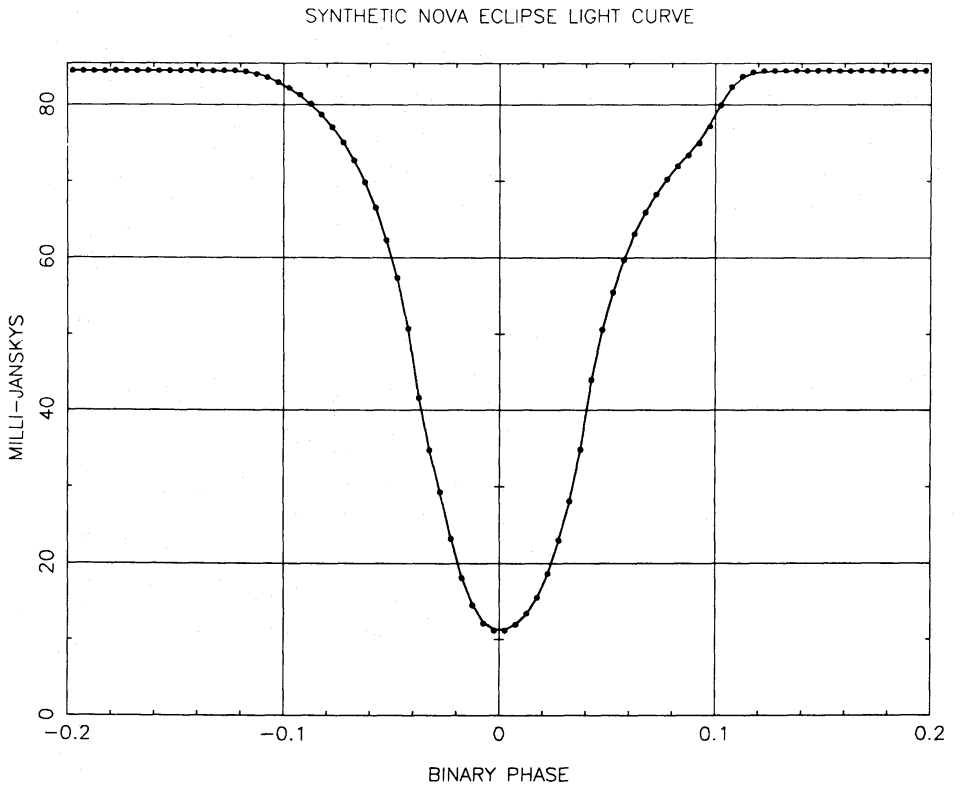

Figure 4. Points give the synthetic eclipse light curve corresponding to the accretion disc image shown in Plate 1(a), and the eclipse geometry of Fig. 3. This light curve resembles those of many old novae, nova-like variables, and dwarf novae in outburst. The continuous curve shows the fit achieved by the reconstructed image of Plate 1(b).

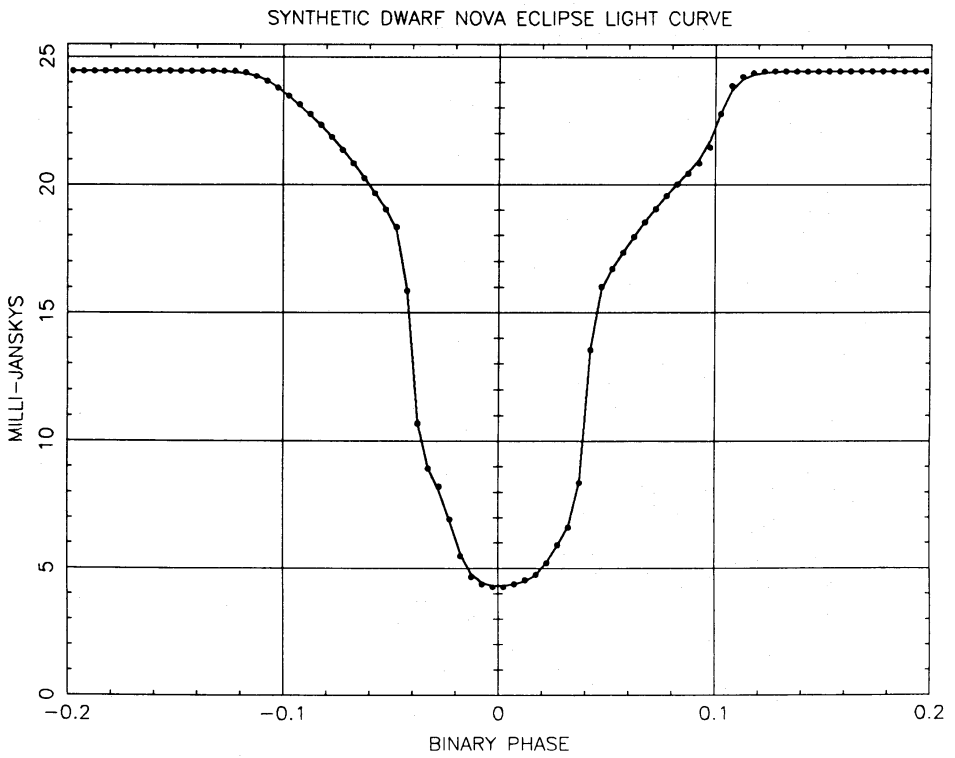

Figure 5. Points give the synthetic eclipse light curve corresponding to from the accretion disc image shown in Plate 1(a), and the eclipse geometry of Fig. 3. This light curve resembles those of quiescent dwarf novae. The continuous curve shows the fit achieved by the reconstructed disc image of Plate 2(b).

5500 Å emission from a steady-state blackbody disc, truncated at radius 0.7 and augmented by a Gaussian hot spot. Dwarf novae in quiescence exhibit eclipse light curves similar to that shown in Fig. 5, with rapid brightness changes caused by the ingress and egress of a compact bright source at disc centre, and of a bright spot forming where the gas stream impacts the outer rim of the disc. The intensity distribution, shown in Plate 2(a), consists of a uniform background within radius 0.7, and two Gaussian bright spots. These disc images and synthetic light curves serve in Section 3 as test cases for evaluating the eclipse mapping procedure.

2.2 MAXIMUM ENTROPY IMAGE RECONSTRUCTION

Maximum entropy methods have already been applied to a variety of astronomical image reconstruction problems (e.g. Gull & Daniell 1978; Bryan & Skilling 1980; Willingale 1981). Maximum entropy (ME) images are found by maximizing a scalar function on image space, the image 'entropy', subject to the constraint that the predicted data associated with the image are consistent with the observed data. The ME method thus applies to any reconstruction problem in which predicted data can be calculated from a given positive image. Its ability to reconstruct an image from incomplete data constraints is essential in the eclipse mapping problem.

Data constraints are imposed by means of the familiar consistency statistic

$$\chi^2 = \frac{1}{M}\sum_{K=1}^{M}\left[\frac{f(K)-d(K)}{\sigma(K)}\right]^2, \tag{4}$$

where $f(K)$ and $d(K)$ are the predicted and observed data, and $\sigma(K)$ is the estimated uncertainty in $d(K)$. The consistency requirement is χ^2=CAIM, where CAIM is a control parameter that strengthens or diminishes the data constraints.

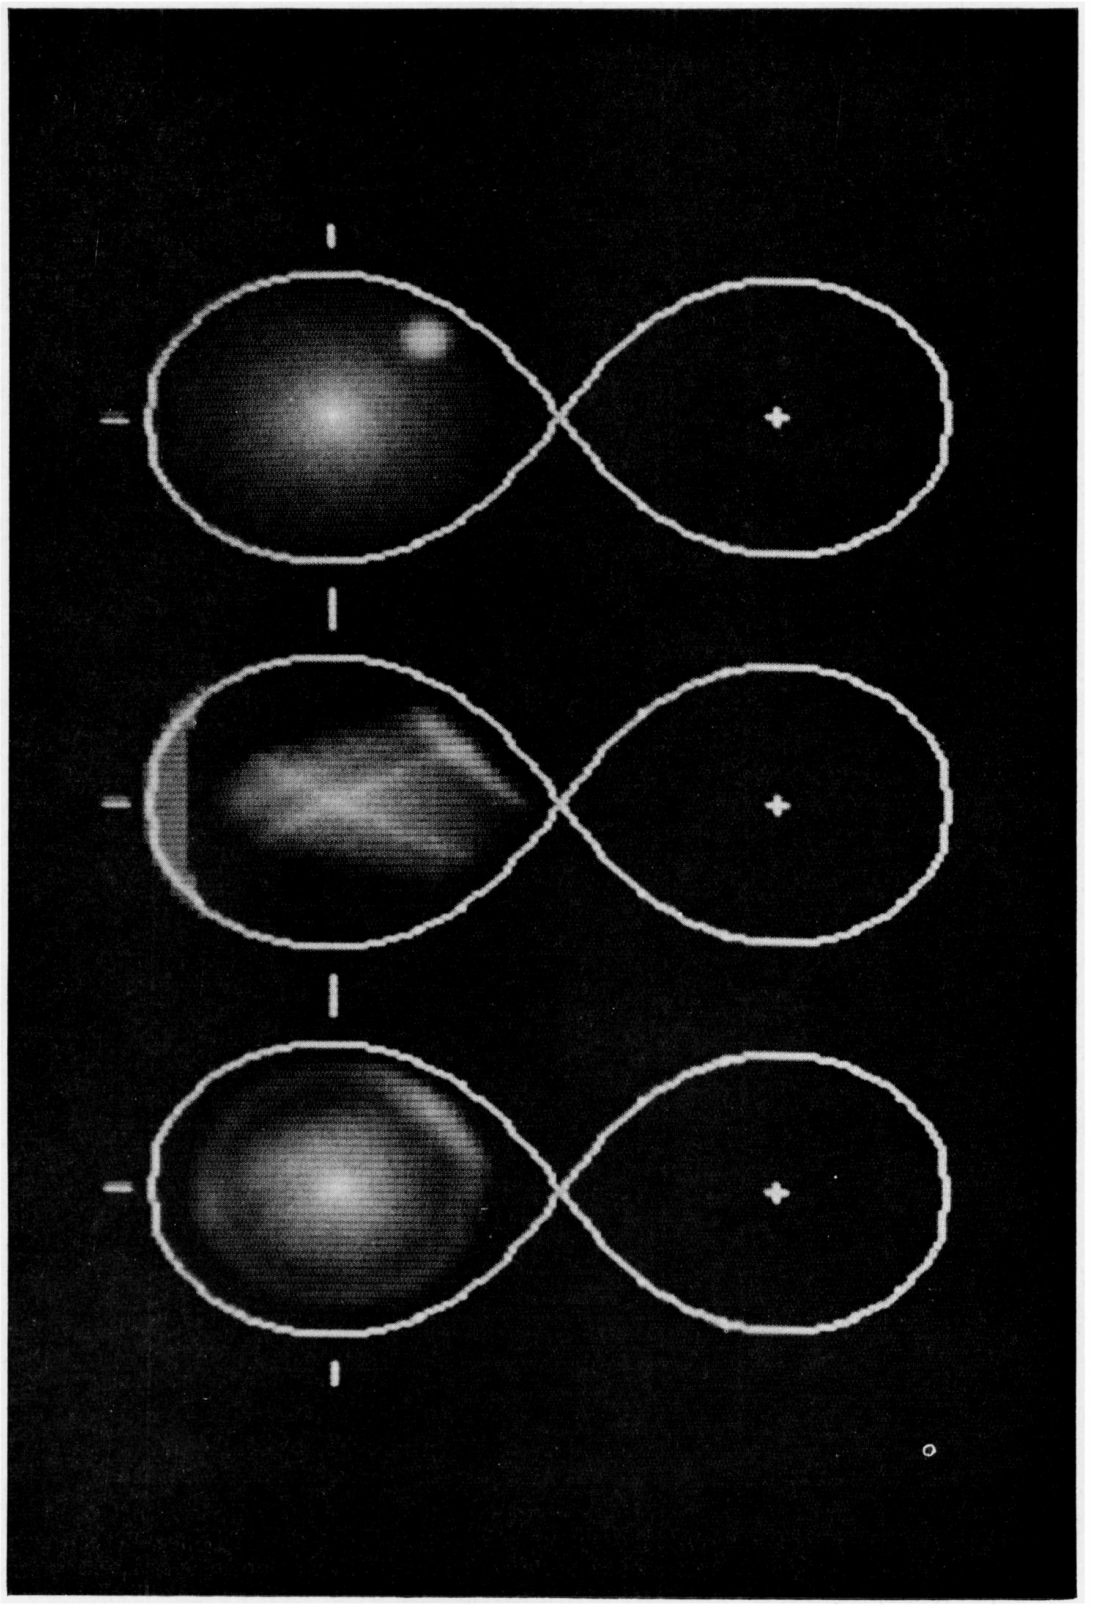

Plate 1. (a) This accretion disc image, the 5500 Å emission from a steady-state blackbody accretion disc with a Gaussian hot spot, was used to generate the light curve of Fig. 4. (b) The maximally uniform disc image reconstructed from the synthetic eclipse data of Fig. 4. (c) The maximally axisymmetric disc image reconstructed from the data of Fig 4.

[*facing page* 134]

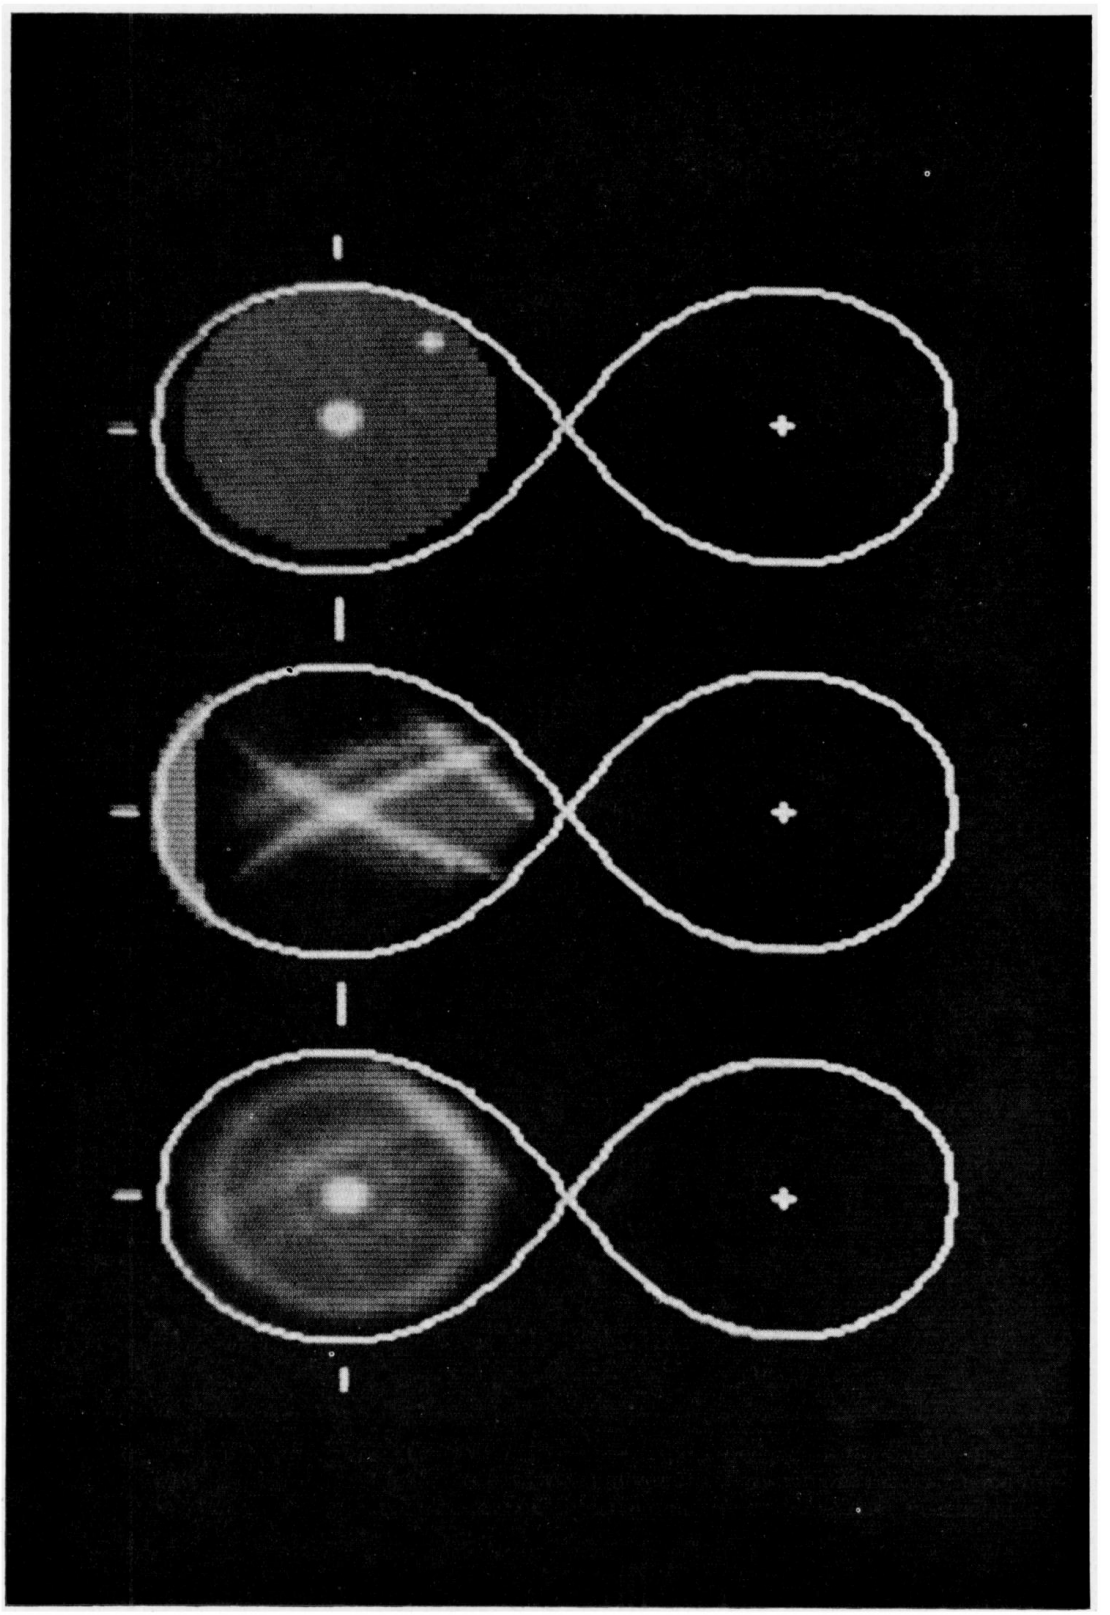

Plate 2. (a) This accretion disc image, two Gaussian bright spots on a uniform background, was used to generate the light curve of Fig. 5. (b) The maximally uniform disc image reconstructed from the synthetic eclipse data of Fig. 5. (c) The maximally axisymmetric disc image reconstructed from the data of Fig. 5.

The image entropy is defined with respect to a default image D, whose role is discussed below in Section 2.3. If $I(J)$ is the image value at pixel J, and $D(J)$ is the corresponding default value, the entropy of the image is given by

$$S=-\sum_{J=1}^{N^2} I(J)\cdot\ln\left[\frac{I(J)}{D(J)}\right]. \tag{5}$$

Shore & Johnson (1980) prove that this is the only functional form that gives a consistent result when the data are considered in trivially different forms (e.g. in a different coordinate system, or in a different order). Note that the logarithmic term in the entropy requires a strictly positive image.

By maximizing the entropy S subject to the constraint χ^2=CAIM, the ME image achieves a balance between opposing influences exerted by the entropy and by the data. When the data strongly constrain the image, the entropy has little effect. When the data was very noisy, the image is determined primarily by the entropy. With the incomplete constraints provided by eclipse data, ME images of accretion discs are influenced by both the entropy and the data.

2.3 THE DEFAULT IMAGE

The default image controls the image property that is measured by the entropy, and hence optimized by entropy maximization. Most previous applications of ME reconstruction employ a uniform default value equal to the mean value of the image. For the eclipse mapping problem, we consider a more general default image of the form

$$D(K)=\sum_{J=1}^{N^2} w(K,J)\cdot I(J)\Big/\sum_{J=1}^{N^2} w(K,J). \tag{6}$$

The default value at each pixel is thus a weighted average of the image values. Different weights may be used for different pixels. Upon substituting (6) into (5) and differentiating with respect to a specific image value $I(K)$, we find

$$-\frac{\partial S}{\partial I(K)}=\ln\left[\frac{I(K)}{D(K)}\right]+1-\sum_{J=1}^{N^2} w(J,K)\left[\frac{I(J)}{D(J)}\right]\Big/\sum_{J=1}^{N^2} w(K,J). \tag{7}$$

Since we consider weights that are symmetric in the sense $w(J, K)=w(K, J)$, this expression vanishes when the image and default image are equal. Thus the default image maximizes the entropy in the absence of data constraints.

With a uniform default image [$w(K, J)=1$], the entropy measures deviations from uniformity, and the ME image is the 'most uniform' image that is consistent with the data. The uniform default tends to suppress any structure in the image (see Bryan & Skilling 1980), and is therefore most appropriate for the detection of structure against a uniform background.

A second useful default image is specified by the weights

$$w(K,J)=\exp[-r^2(K,J)/2\Delta^2], \tag{8}$$

where $r(K, J)$ is the distance between pixels K and J. This default image is a convolution of the image with a Gaussian point-spread function of width Δ. The entropy in this case loses its sensitivity to image structure on scales broader than Δ. Broad features in the ME image are therefore freely determined by the data, while image structure on short scales is suppressed by the entropy. This form of default produces the 'smoothest' image that is consistent with the data.

The default image that proves to be most useful for reconstructing images of accretion discs is a radial profile default, specified by the weights.

$$w(K,J)=\exp\{-[R(K)-R(J)]^2/2\Delta^2\}. \quad (9)$$

Here $R(K)$ is the distance between pixel K and the centre of the disc. The radial profile default is an axisymmetric image obtained by averaging image values around annuli concentric with the centre of the disc. With this default, the entropy remains sensitive to azimuthal structure, but loses information about the radial profile of the disc. The radial profile of the ME image is therefore determined by the eclipse data alone, while azimuthal structure is suppressed by the entropy. The ME image is the 'most nearly axisymmetric' image that is consistent with the data.

2.4 MEMSYS

A robust and efficient algorithm for the constrained maximization of the entropy in a large-dimensional image space was developed by Skilling (1981). The algorithm is described by Skilling & Bryan (1984), and its general purpose FORTRAN implementation MEMSYS is described by Burch, Gull & Skilling (1983). MEMSYS adjusts an initial image by considering a quadratic approximation to the constrained entropy maximization on a 3, 4, or 6-dimensional subspace of the image space. Several transformations between image space and data space are required to compute the image space search directions. In the eclipse mapping problem, the mapping from image space to data space is given by the linear relation (2), while the transpose of (2) provides the mapping from data space to image space.

An initially uniform image evolves through iterated adjustments by MEMSYS toward the ME image. MEMSYS seeks at first to reduce the value of χ^2 until the target value CAIM is achieved. The aim subsequently is to maintain χ^2=CAIM while increasing the entropy. (In fact a smooth transition between these two stages occurs.) MEMSYS assumes a fixed default image, but the iteration converges satisfactorily when the default image for each iteration is computed from the image of the previous iteration. The iteration is terminated when the image space gradients of the entropy and of χ^2 are parallel, signalling that a correct solution to the constrained maximization problem has been found. ME images found by this procedure do not depend on the choice of initial image. The ME images are also stable to perturbations; bright spots added at different places on the image are consistently removed by about 20 additional iterations.

3 Performance tests

Maximum entropy reconstruction is a non-linear procedure that is difficult to evaluate by analytic methods. I have therefore conducted tests with synthetic eclipse data for which the underlying accretion disc images are known. Since the eclipse mapping problem is under-determined, some distortion of the reconstructed image is expected. As we shall see, the nature of the distortion can be controlled through the default image. Errors in the adopted eclipse geometry, and violation of the assumptions inherent in the light curve synthesis program also affect the reconstructed image. While it is impossible to anticipate every situation that might arise in practice, the results presented here provide a basis for interpreting ME images derived from real eclipse observations.

The main results of the simulation tests are illustrated here by two representative cases. The synthetic eclipse data are given by the points in Figs 4 and 5. The underlying accretion disc images are shown in Plates 1(a) and 2(a), with ME reconstructions given below. Plate 1(a) is the 5500 Å light from a steady-state blackbody disc model, in which the surface temperature T varies with disc radius R according to $T=10^4\,\mathrm{K}\cdot(R/R_{\mathrm{L1}})^{-3/4}$. This radial temperature profile is augmented by a circular Gaussian function to produce a hot spot in the first quadrant and near the rim of the

disc. Plate 2(a) is a uniform 10^4 K disc with two Gaussian hot spots, one near the rim of the disc and one at disc centre.

The intensity vanishes outside radius 0.7 in these models, but the reconstructed images were forced to be positive out to radius 0.85 in order to see how well the sharp outer rim of the disc is recovered. A 51×51 pixel format was used for all images. The eclipse geometry used for the synthetic light curves was that of Fig. 3, and correct values of ϕ_0, $\Delta\phi$, and q were used in the reconstructions. The effect of errors in the assumed geometry is considered in Section 3.3.

3.1 UNIFORM AND RADIAL PROFILE DEFAULT IMAGES

Plates 1 and 2 compare the original disc images to reconstructed images with uniform and radial profile defaults. The maximally uniform images are given in Plates 1(b) and 2(b), and maximally axisymmetric images appear in Plates 1(c) and 2(c). The ME images in every case are consistent with the synthetic eclipse data; the fits shown in Figs 4 and 5 are typical.

With a uniform default image, each compact bright region in the original disc image is reconstructed as a crossed pair of bright streaks. This 'crossed-arch' distortion results from an interplay between the entropy and the eclipse constraints. The entropy, acting to suppress extreme image values, is favourably increased as flux is removed from the bright region. To maintain the shape of the eclipse light curve, flux removed from the bright region is redistributed along the ingress and egress arches that pass through it. The resulting image, although severely distorted, can with care be used to ascertain the locations of the compact bright spots.

With a radial profile default image, the entropy acts to suppress any azimuthal structure in the image. Since the entropy does not sense the radial profile of the image, the latter is determined by the eclipse data alone. The maximally axisymmetric ME images (Plates 1c and 2c) recover the correct brightness distribution near the centre of the disc, but distort the distribution in the bright spot region. The bright spot flux is redistributed in azimuth to the extent that a faint annulus encircles the disc at its radius. Some remnant crossed-arch distortion is present, but its amplitude is greatly diminished because the entropy no longer tries to suppress the bright region at disc centre.

The maximally axisymmetric ME images are suitable for quantitative comparison with accretion disc models. Fig. 6 shows the blackbody brightness temperature profile derived from the ME image of Plate 1(c). The reconstructed temperature profile agrees closely with the $T \propto R^{-3/4}$ law used to derive the original image of Plate 1(a). The outer rim of the disc at radius 0.7 is also correctly reconstructed.

3.2 RESOLUTION AND NOISE

Noise in the eclipse light curve degrades the spatial resolution of the reconstructed image. This trade-off between noise and resolution is controlled through the value of CAIM, which adjusted the strength of the eclipse constraints. With a large value of CAIM, the predicted light curve need not agree very closely with the eclipse data, and the ME image relaxes toward a uniform image. With a small value of CAIM, the predicted light curve is forced to follow noise in the observed light curve, and the ME image develops a spurious 'grooved' texture oriented along the ingress/egress arches. In practice, the appearance of these features guides the choice of an appropriate value for CAIM when the estimated data errors are uncertain.

3.3 UNCERTAIN ECLIPSE GEOMETRY

It is important to consider how errors in the adopted eclipse geometry affect the ME image. Fig. 7 illustrates the distortion of the eclipse geometry produced by separate changes in ϕ_0, $\Delta\phi$, and q.

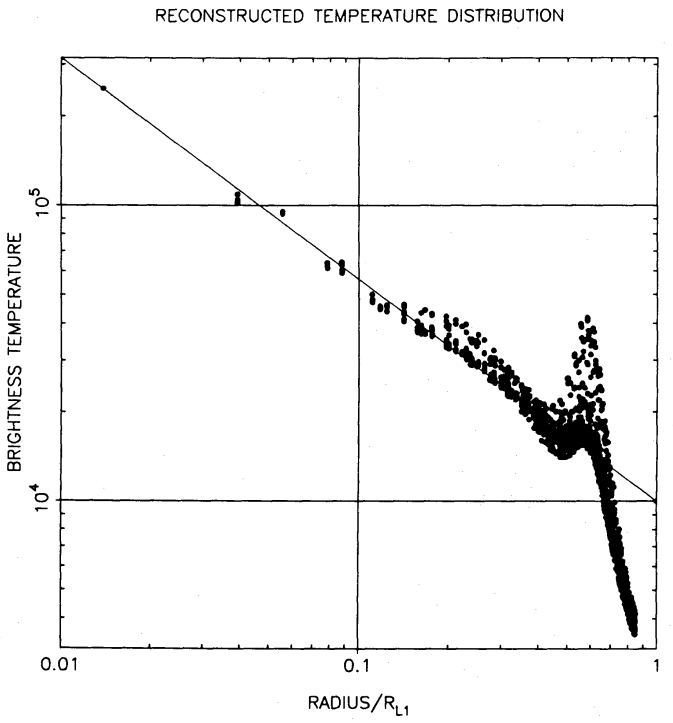

Figure 6. Each plotted point is the blackbody brightness temperature derived from a single pixel of the reconstructed disc image in Plate 1(c). This reconstructed temperature profile compares favourably to the $T \propto R^{-3/4}$ model used to generate the synthetic eclipse data. The outer edge of the disc at radius 0.7 is also correctly reconstructed.

A change in ϕ_0 causes a rotation of the occulted regions about the centre of the red dwarf (Fig. 7a), and therefore an error in ϕ_0 merely rotates the ME image. A change in $\Delta\phi$ corresponds to a change in the inclination of the system, and hence in the distance from the red dwarf to which the occulted regions extend (Fig. 7b). To first order, then, errors in ϕ_0 and $\Delta\phi$ move the position of the central bright region away from the nominal centre of the disc.

The eclipse geometry is only weakly sensitive to the mass ratio q, which controls the size of the red star and hence the width of the occulted regions. An error in q produces a second-order distortion that vanishes at the disc centre (Fig. 7c). Thus the position of a bright spot in the outer region of the disc depends on q, but the brightness distribution near disc centre, and the azimuthally-averaged properties of the disc do not.

In practice, values of ϕ_0 and $\Delta\phi$ are taken from the half-depth points of the eclipse light curve. These initial values may be adjusted until the crossed arches on an eclipse map made with a uniform default image intersect at the centre of the disc, but this fine-tuning is not usually necessary unless the eclipse is shallow. The mass ratio generally cannot be determined from the eclipse light curve (however, see Cook & Warner 1984). Spectroscopic mass ratios for cataclysmic binaries are suspect because emission lines from the disc often show unexplained phase shifts, and hence may not reveal the true orbital motion of the white dwarf. Although a range of values of q must therefore be considered, the ME image is only weakly sensitive to q.

3.4 NEGLECTED BACKGROUND LIGHT

An unknown fraction of the light at the bottom of the eclipse is contributed by the red-dwarf star. The effect of this neglected background light is determined by adding various constant fluxes to the synthetic eclipse data before reconstructing images of the disc. The added flux appears in the ME image mainly in the small section of the disc that escapes the eclipse. The reason for this is

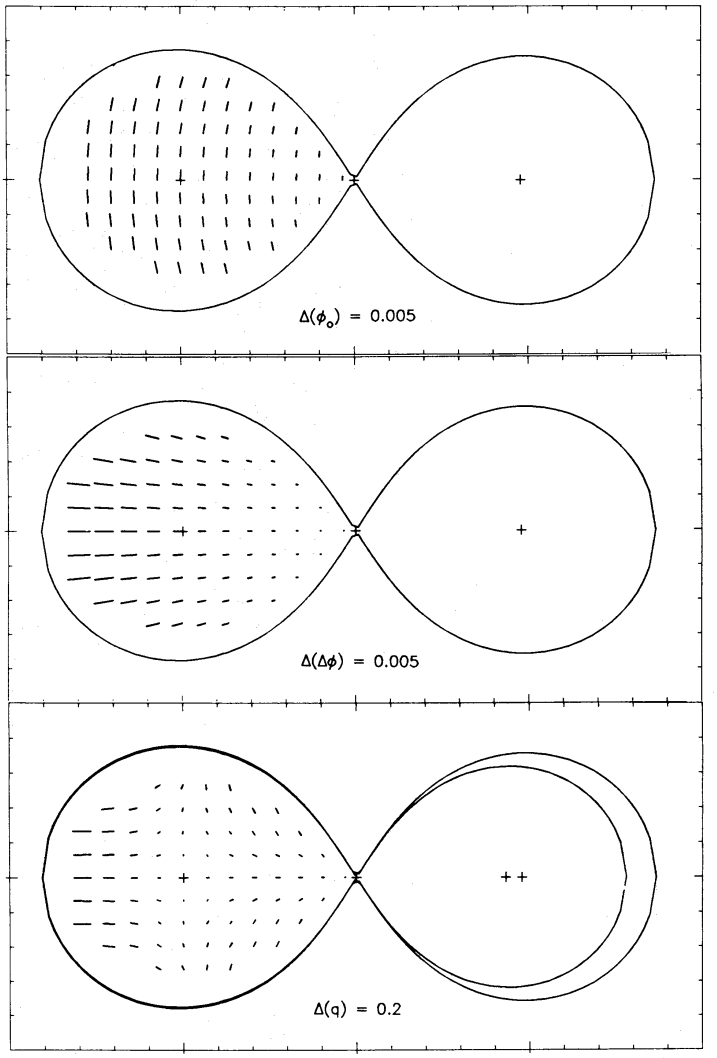

Figure 7. Vectors show the change in the eclipse geometry produced by the indiated changes in the parameters ϕ_0, $\Delta\phi$, and q.

that the eclipse data strongly constrain the total flux from the region of the disc that does participate in the eclipse. Since its effect is easily detected in the reconstructed image, the neglect of background light is not a serious problem. If the red dwarf contributes a substantial fraction of the total light (shallow eclipse), its direct contribution must be taken into account in the light-curve synthesis program.

3.5 FLAT DISC APPROXIMATION

The flat disc approximation used in the light curve synthesis program fails if the inclination differs from 90° by less than the opening angle of the disc. The eclipse mapping method formulated here therefore does not apply to cataclysmic variables with inclinations very close to 90°, or to low-mass X-ray binaries, whose accretion discs may open at an angle of 10–15° due to X-ray heating. In these systems, light from the central regions of the disc is obscured by the outer rim of the disc. The light-curve synthesis program could easily be modified to handle a disc of finite thickness, but an *ad hoc* prescription for the height of the disc photosphere at each point of the surface would be needed.

3.6 NON-STEADY EMISSION

Cataclysmic variables exhibit erratic brightness variations on short time-scales, systematic brightness variations with binary phase, and secular variations on longer time-scales. This behaviour violates the assumption of steady emission from the disc. The erratic variability, known as 'flickering', is eliminated by coherently averaging observations taken on many eclipse cycles. This treatment succeeds if the mean properties of the disc do not change significantly for many cycles while the observations are accumulated. The image reconstructed from the mean light curve then represents the time-averaged appearance of the accretion disc.

Orbital variation is a more insidious problem. The most common form is a 'hump' in the light curve during the half-cycle preceding the eclipse. The cause of the orbital hump is believed to be anisotropic radiation from the region where the gas stream from the red dwarf impacts the outer rim of the disc. Since the light-curve synthesis program models only the eclipse by the red dwarf, the observed light curve must be renormalized to eliminate brightness variations at phases away from the eclipse. This procedure is correct if every part of the disc participates equally in the orbital variation. In other cases the reconstructed image is distorted roughly in proportion to the amplitude of the orbital variation, and to the concentration of the region that varies. If the orbital variation arises from a compact spot, the reconstructed image develops a bright arch and a dark arch that cross at the location of the variable spot. In extreme cases, a good fit to the light curve cannot be achieved.

4 Conclusion

The eclipse mapping method developed in this paper produces images of an accretion disc that are fully consistent with an observed eclipse light curve. Light curves with signal-to-noise ratios as large as 200 are readily reproduced (Figs 4 and 5). Since the eclipse data do not fully constrain the image, freedom remains to optimize some property of the image. This is accomplished by maximizing the image entropy subject to the constraints of the eclipse data.

The specific image property that is optimized is controlled by the default image used in the definition of the entropy. Images formed with a uniform default image are badly distorted, but reveal the correct positions of compact bright spots on the disc (Plates 1b and 2b). Images formed with a radial profile default image suppress azimuthal structure, but recover the correct radial profile of the disc (Plates 1c and 2c). The radial profile is needed for comparison with accretion disc theory, which currently assumes axisymmetry. A preference for axisymmetry may also be motivated physically, since Keplerian shear should rapidly eliminate azimuthal structure in the disc (Pringle 1981).

By casting observations into a form that is easily visualized and directly comparable to model accretion discs, the eclipse mapping method provides unprecedented tests of accretion disc theory. Eclipse maps from light curves at different wavelengths in effect reveal the spectrum emitted at each place on the disc. This provides information about the local vertical structure of the disc, reveals the location of optically-thin regions, and maps the temperature distribution in the optically-thick regions of the disc.

The effective temperature profile of a steady-state disc in a binary system of orbital period P is predicted to be

$$T_{\rm eff}^4 \propto P^{-2}(R/R_{\rm L1})^{-3}\dot{M}f(q), \tag{10}$$

where $\dot{M}$ is the mass transfer rate, and $f(q)$ is a weak function of the mass ratio q. Temperature distributions derived from eclipse maps directly test the $T_{\rm eff} \propto R^{-3/4}$ law, which derives from basic energy balance considerations. Fig. 6 indicates that the temperature profiles of discs can be

recovered with an uncertainty of less than 20 per cent, and hence that mass transfer rates for cataclysmic variables stars can be measured to about factor of 2.

The doppler-broadened emission lines seen in the optical spectra of cataclysmic variables are known to be associated with the accretion disc, but the site of their formation is poorly known. Both photoionization of a chromospheric layer by radiation from the disc centre (Jameson, King & Sherrington 1980), and emission from an optically-thin outer disc annulus (Williams 1980) may play a role in their formation. Eclipse data with high spectral resolution can be used to map the emission regions of individual Balmer, He I, and He II emission lines.

The outbursts of dwarf novae are episodes of enhanced mass transfer through the accretion disc. The longstanding debate over the roles played by disc and red-star instabilities in triggering the outburst can be settled by eclipse maps showing the evolution of the accretion disc during the course of a dwarf-nova outburst.

The eclipse mapping method developed in this paper establishes a model-independent framework for quantitative observational studies of spatial structure in the accretion discs of cataclysmic variables. Our theoretical understanding of accretion discs has probably progressed as far as is possible in the absence of a theory for their anomalous viscosity. Further progress depends upon guidance from observations of real discs. The eclipse mapping method, so far as is possible, allows these observations to speak for themselves.

Acknowledgments

I acknowledge important discussions with A. C. S. Readhead, S. W. Mochnacki, R. P. Lindfield, S. F. Gull, and K. A. Whaler. The maximum entropy data analysis package MEMSYS, developed by J. Skilling and made available to me by S. F. Gull, was essential. Computations were performed with the VAX 11/780 computers of the Astronomy Department at Caltech, and the Cambridge node of the SERC Starlink Network. Support from a Graduate Research Assistantship at Caltech, and an SERC Postdoctoral Research Assistantship at Cambridge is gratefully acknowledged.

References

Bryan, R. K. & Skilling, J., 1980. *Mon. Not. R. astr. Soc.,* **191,** 69.
Bryan, R. K. & Skilling, J., 1984. *Mon. Not. R. astr. Soc.,* **211,** 111.
Burch, S. F., Gull, S. F. & Skilling, J., 1983. *Comp. Vis. Graph. Im. Proces.,* **23,** 113.
Cook, M. C. & Warner, B., 1984. *Mon. Not. R. astr. Soc.,* **207,** 705.
Frank, J. & King, A. R., 1981. *Mon. Not. R. astr. Soc.,* **195,** 227.
Frank, J., King, A. R., Sherrington, M. R., Jameson, R. F. & Axon, D. J., 1981. *Mon. Not. R. astr. Soc.,* **195,** 505.
Gull, S. F. & Daniell, G. J., 1978. *Nature,* **272,** 686.
Jameson, R. F., King, A. R. & Sherrington, M. R., 1980. *Mon. Not. R. astr. Soc.,* **191,** 559.
Mochnacki, S. W., 1971. *MS thesis,* University of Canterbury, New Zealand.
Pringle, J. E., 1981. *Ann. Rev. Astr. Astrophys.,* **19,** 137.
Shore, J. E. & Johnson, R. W., 1980. *IEEE Transactions on Information Theory,* **IT-26,** 26.
Skilling, J., 1981. *Algorithms and Applications,* presented at Workshop on Maximum Entropy, Laramie, Wyoming.
Wade, R., 1984. *Mon. Not. R. astr. Soc.,* **208,** 301.
Williams, R., 1980. *Astrophys. J.,* **235,** 939.
Willingale, R., 1981. *Mon. Not. R. astr. Soc.,* **194,** 359.
Young, P. & Schneider, D. P., 1981. *Astrophys. J.,* **238,** 955.