Mon. Not. R. Astron. Soc. 000, 000–000 (0000) Printed 11 March 1998

'Stokes imaging' of the accretion region in Magnetic Cataclysmic Variables – I. Conception and Realisation

S. B. Potter,[1] P. J. Hakala,[1,2] and Mark Cropper[1]
[1] *Mullard Space Science Laboratory, Holmbury St. Mary, Dorking, Surrey RH5 6NT*
[2] *Observatory and Astrophysics Laboratory, FIN-00014, University of Helsinki, Finland*

ABSTRACT

Modelling the polarized cyclotron emission from Magnetic Cataclysmic Variables (MCVs) has been a pivotal technique for determining the structure of the accretion zones on the white dwarf. To date model solutions have been obtained from trial fits to the intensity and polarization data which have been constructed from emission regions (for example arcs and spots) put in by hand. These models were all inferred indirectly from arguments based on the polarization and X–ray light curves.

We present a more analytical and objective technique using optimisation by a Genetic Algorithm, Tikhonov regularisation and the Powell's method that robustly models the details of polarized emission.

To demonstrate the success of this technique, we show the results of several simulations in which we calculated the intensity and polarization curves from arbitrarily shaped emission regions on the surface of a sphere and then applied our code to these curves to recover the original test data. We also show how adding artificial noise effects the outcome of the optimisation technique.

Key words: binaries:- close - stars : cataclysmic variables: polars : stars, polarimetry: genetic algorithms.

1 INTRODUCTION

In the past decade several authors have modelled the polarized cyclotron emission from the shock regions in magnetic cataclysmic variables (see the reviews by Cropper 1990 and Beuermann & Burwitz 1995). However, it was found very early on that it is too simplistic to assume that accretion takes place over a small circular cap at the magnetic pole: many authors (e.g. Potter et al 1997; Wickramasinghe & Ferrario 1988; Beuermann, Stella & Patterson 1987; Cropper 1986;) have found that the region must be extended into an arc or comma shape and offset from the magnetic pole in order to explain the shape of the polarized cyclotron light curves. Wickramasinghe et al. (1991) provides a good example of how high signal–to–noise ratio observations can be modelled in detail. Until now, the parameters that define the shape and location of the emission regions have been adjusted by a trial and error approach until the model gave a good fit. Clearly the number of free parameters describing the shape and location of emission regions becomes very large.

Later, Cropper & Horne (1994) used the maximum entropy method to map the emission region in ST LMi using optical photometry and X–ray light curves. They found that the resulting maps of the emission region in ST LMi showed structure within a more extended region.

In this paper we go a step further and describe the techniques of using an emission model that has the anisotropies of polarized cyclotron emission incorporated into it. This allows us to model the intensity, circular polarization, linear polarization and position angle data. However, instead of the maximum entropy method and the conjugate gradient optimisation (used in Cropper & Horne 1994), we have used Tikhonov regularisation and a genetic algorithm to obtain the best fit to the data and to produce maps of the emission regions.

2 THE CYCLOTRON MODEL

The first stage in modelling cyclotron emission is to calculate the viewing angle – the angle between the line of sight and the magnetic field line from where the emission emanates. To do this we have used a dipole magnetic field formalism as described in Cropper (1989). The magnetic dipole is offset from the spin axis of the white dwarf by an angle β and the whole system is inclined with respect to the viewer by an angle i to the spin axis. For a particular model we assume values for the polar magnetic field strength and the optical depth parameter.

We then interpolate on the cyclotron model calculations of Wickramasinghe and Meggitt (1985) to obtain the intensities and polarizations of the cyclotron flux for emission points in a magnetic field as a function of rotational phase of the white dwarf. Wickramasinghe and Meggitt (1985) calculated the cyclotron emission from hot plasmas for a wide range of values of the electron temperatures, frequency, viewing angle with respect to the magnetic field, electron number density, propagation path, and an optical depth parameter Λ.

In constructing the models we have assumed a magnetic dipole field with independent field strengths for each magnetic hemisphere. This provides a first order approximation to an offset dipole field. The electron temperature is assumed to be constant across the shock. The temperature of the cyclotron component is set principally by the mass of the white dwarf although it is affected by the strength of the magnetic field locally (Wu 1994). This temperature dependence is weak but complicated and requires detailed models for the cooling processes. Therefore, for the purposes of calculating the emission map using broad band polarizations, we have neglected the temperature dependence. This is justified in this case, but may not be appropriate when calculating cyclotron *spectra* from such a region (Potter 1998 in preparation). The height of the emission region is assumed to be a very small fraction of the white dwarf's radius. The value of the optical depth parameter changes across the emission region because of changes in the local magnetic field strength. The effects of cyclotron and free–free opacities are included in the grid of calculations. The model is a modified version of that used by Potter et al. (1997).

For the optimisation procedure the surface of the white dwarf is divided into a grid of 1740 points each representing a 6 by 6 degree square in magnetic coordinates. This provides increasing spatial resolution towards the magnetic poles where the magnetic field strength and orientation (and therefore its effect on the optical depth parameter) is changing most rapidly. Larger grids would be desirable where polarization data with high phase resolution was available. Each point on the grid can have a different number of emission points, related to the accretion rate, and thus the optical depth parameter is calculated also as a function of local density. The Stokes parameters are then calculated for each point along the emission regions and summed to produce the total polarizations (in terms of intensity, circular polarization, linear polarization and position angle) as a function of phase. Projection effects are taken into consideration (by using a cos(viewing angle) factor for each element), since most of the emission is not optically thin.

In the convention used here, the magnetic latitude of an emission point is measured from the magnetic pole in the upper magnetic hemisphere pointing towards the observer. Therefore the magnetic equator is at a magnetic latitude of 90° and the magnetic pole in the lower magnetic hemisphere pointing away from the observer is at a magnetic latitude of 180°. The magnetic longitude of an emission point is measured in an anticlockwise direction from the line joining the magnetic pole and the spin pole of the white dwarf as it rotates in a clockwise direction.

Optimisation of the model to the data then proceeds by adjusting the number and distribution of emission points across the surface of the white dwarf. During the optimisation the inclination, magnetic dipole offset (in latitude and azimuth) and the magnetic field strength at the poles are kept fixed.

3 OPTIMISATION

In this section we give a description of the way we construct the resulting cyclotron emission distribution. As noted earlier, given the cyclotron emission distribution on the white dwarf surface, the strength and orientation of a dipole field and the system inclination, we are able to compute the resulting four Stokes parameter curves (Q, U, V and I) as a function of orbital (or white dwarf rotational) phase. However, obtaining an optimal fit to the dataset is an under–constrained problem. Our typical datasets consist of about 30–50 data points over the orbital phase in each Stokes parameter, so we have all together an order of 120–200 data points. Therefore, even a very modest map (say 50×100 elements) over the white dwarf surface has about ten times more points than we have data points. It is thus necessary to employ some sort of regularisation method (for example maximum entropy) in order to find a most likely, by some acceptable criterion, solution of the set of solutions that fit the data.

3.1 The Fitness Function

Our approach is based on Tikhonov regularisation (Tikhonov 1963, Piskunov et al. 1990) for providing the additional constraint. Tikhonov regularisation is an alternative for using an entropy term. In the maximum entropy approach the function to be minimised is of the form

$$F(p) = \chi^2 + \lambda \sum_i p_i log \frac{p_i}{m_i} \tag{1}$$

where χ^2 is the χ^2 of the fit to the data, p_i is a single emission element on white dwarf surface and m_i is a local geometric mean of the pixels next to pixel i. The sum is over all the surface grid elements i. $F(p)$ is called the fitness function. Minimising this together with a suitable choice of the Lagrangian multiplier λ will produce the maximum entropy solution of the distribution of emission points. Piskunov et al. (1990) have however pointed out that, in the case of Doppler imaging, using the Tikhonov functional instead of the maximum entropy functional leads to more consistent solutions in case of noisy data. In this case the fitness function to minimise is

$$F(p) = \chi^2 + \lambda \sum_i \|\nabla p_i\|^2 \tag{2}$$

where $\|\nabla p_i\|$ is the mean gradient of the number of emission points at point i. We have selected equation 2 for our fitness function.

3.2 Genetic Algorithm

We need a robust technique to minimise $F(p)$ in order to ensure a global rather than local optimisation is achieved. We have chosen a Genetic Algorithm (GA – see for instance

Charbonneau 1995 and Hakala 1995 for astronomical applications of GA). GA is capable of climbing away from local minima. Still, we must emphasise that the method is not guaranteed to find the global minimum. There are, however, a number of applications where GAs are shown to provide better and more stable solutions than more conventional optimisation methods (Charbonneau 1995). The second advantage is that GA is a 'forward only' approach for inversion problems: we do not have to invert the problem at all, but simply compute solutions based on given sets of parameters. This makes the method very easily applicable to a range of different problems. Also, major changes to underlying models do not require very much effort in terms of 're–inverting' the problem.

The GA works by first generating a set of random solutions. The fitness of each solution is then calculated using equation 2 and the solutions are then ranked in order of their fitness. The next generation of solutions are then produced by a type of natural selection procedure – the solutions that were ranked best are more probable to breed the next generation. Formally, solutions are selected from the ranked table by using a selection algorithm of the form

$$S_i = INT(\frac{-N_{pop}}{C_s} log(RND(0,1))) + 1, \quad (3)$$

where S_i is the solution rank number after sorting and N_{pop} is the size of each population. C_s is a scale parameter used to bias the selection criteria more to the best ranked solutions. INT is the INTEGER function and RND is a random number generator. The next population is then generated by applying genetic crossover and mutation operators to the selected parent pairs.

Eventually the improvement in fitness of the GA solutions will start to level out and a more analytical approach is required to improve the fit further. For the final approach to minimum we use Powell's method, a line minimisation routine (see Press et al. 1992).

3.3 Fixed parameters

The optimisation technique makes adjustments to the number and distribution of emission points on the surface of the white dwarf to improve the model fit of the polarized light curves to the polarimetric data. However, there are also two other types of parameters that are not included in the optimisation technique and are kept constant during a particular optimisation run.

Firstly there are the physical parameters of the system in which we are interested. These include the inclination, magnetic dipole offset angles (latitude and azimuth) and the magnetic field strength at the poles. Some or all of these parameters are known for a particular system – the technique is applied for a range of values for the parameters that are not known. If none of these parameters are known then good fits can be obtained for a large range in each of these parameters. However, if one parameter, such as the inclination or the magnetic field strength, is known fairly accurately, then the range in values for the remaining parameters become significantly reduced for good fits to be obtained. For example if the inclination is known, then the circular and linear polarization light curves will be more effective in constraining the magnetic dipole offset angles.

Then there are the parameters that control the direction and/or rate at which the GA finds the global minimum. These include the size of the population, mutation probability and the Lagrangian multiplier. We consider each of these in turn.

The population size (the number of solutions that are bred) is simply a scale factor of the number of data points: the more data points to fit, the larger the population size needs to be in order for the optimisation technique to work. However its exact value depends strongly on the type of problem and can only be obtained by experiment.

A solution can be mutated by simply randomising some of its array elements. The mutation probability is a dynamical term: it has a low probability at the start of the optimisation run and gradually increases as the rate at which better solutions are found decreases. A high mutation probability at the beginning of a run serves only to re–randomise the solutions and therefore the better solutions are not bred to the next generations. However, towards the end of an optimisation run, most of the solutions are located somewhere within the global minimum and are therefore very similar. Therefore, increasing the mutation probability at this stage forces the optimisation technique to search further afield.

3.4 The Lagrangian Multiplier

The final solution is dependent on the Lagrangian multiplier λ. Ideally one needs to find the largest value for λ consistent with a good fit to the data (equivalent to the smoothest fit to the data). We ensure this by starting our search with large values for λ and gradually decrease it until $F(p)$ no longer improves. The upper curve in Fig.1 is a plot of the final values (after genetic optimisation) of the two terms in equation 2 for decreasing values of λ. The diamonds mark the points at with the GA stops and subsequently λ is halved. To begin with, for high λ, only the regularisation term is favoured resulting in a high χ^2 (top left of Fig.1). As λ is gradually decreased the χ^2 is minimised resulting in a higher final value for the regularisation term (lower right of Fig.1). Eventually reducing λ further has no effect on the final solution and the process comes to a halt (far right of fig.1).

The final solution from the GA is then used for the Powell's minimisation. λ is then chosen such that $\lambda \sum_i \|\nabla p_i\|^2$ is the same as the χ^2 term in equation 2. This ensures that the search direction is such that the smoothness of the solution and the χ^2 fit to the data are improved simultaneously. If λ is set too high the final solution is a smooth image with a poor fit to the data. If set too low, the χ^2 improves without any significant decrease in the regularisation term.

The lower curve in fig.1 shows the Powell's method improvement on the GA's final solution. Both the χ^2 and the regularisation term are minimised. The asterisks represent equal computation time intervals starting from the right where the GA finished.

4 SIMULATIONS

For the testing of our program we have generated synthetic Stokes parameter curves with 50 points over the orbital phase (typical datasets have roughly this number). This gives us 200 data points to fit. In this paper we use

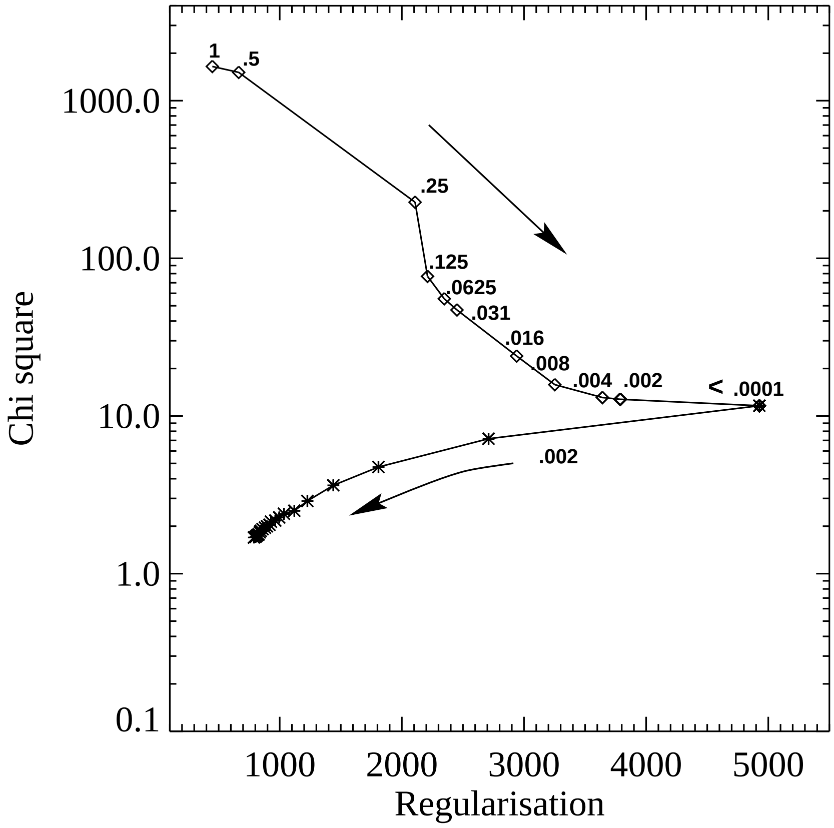

Figure 1. The upper curve traces the genetic optimisation. The numbers represent values for λ and the diamonds mark where λ was halved. The lower curve shows the Powell's optimisation. In this case λ is constant and chosen so that the regularisation term and the χ^2 are equal. The asterisks mark equal computation intervals.

a 6 degree resolution map on the white dwarf surface, which in turn yields 1740 free parameters (60×29 surface grid elements on the white dwarf surface). We have chosen an approach where our GA population size is 500. The two simulated polarized light curves presented below were produced by calculating the emission that would arise from a point source (Fig. 2a) and two extended sources (Fig. 3a) in the B–band. The resulting light curves (Fig. 2b and Fig. 3b) have infinite signal–to–noise ratio. In both cases the inclination was $80°$, the magnetic dipole offset was $10°$, the magnetic polar field strength was 60 MG and the polar optical depth parameter Λ was 1.0×10^5.

4.1 Point source

Figure 2b shows the model fit (thick smooth curves) to the input test data (histogram curves) produced by a single emission source (Fig. 2a). The technique has reproduced the polarized light curves relatively accurately. The fits are not perfect because the regularisation term smoothes the image solution and hence the sharp peaks in the test data are not exactly reproduced. It has generated the self occultation and self absorption dips in the intensity curve, the correct relative brightness of the linear polarization pulses and the double humped positive circular polarization. The variation in the position angle of linear polarization is also accurately modelled. During phases 0.5–1.0 there is little or no linear polarized flux and hence the position angle variation has no meaning.

Figure 2c. shows the model prediction for the shape and location of the emission region needed to produce the input data set. The grey scale is a measure of the number of emission points across the surface of the white dwarf. As can be seen from the figure, the optimisation code has accurately located and mapped the size of the emission region. It has also produced the correct number of emission regions and not, for example, located an emission region on the opposite magnetic hemisphere in an attempt to generate the correct amount of circular polarization.

4.2 Extended sources

Figure 3b shows the fit (thick smooth curves) to the input test data (histogram curves) produced by two extended sources (Fig. 3a).

As before, the test light curves have infinite signal–to–noise ratio. The optimisation technique has reproduced the several more complex features arising from the extended regions. These include the additional linear and intensity features such as dips and peaks. It has also correctly modelled

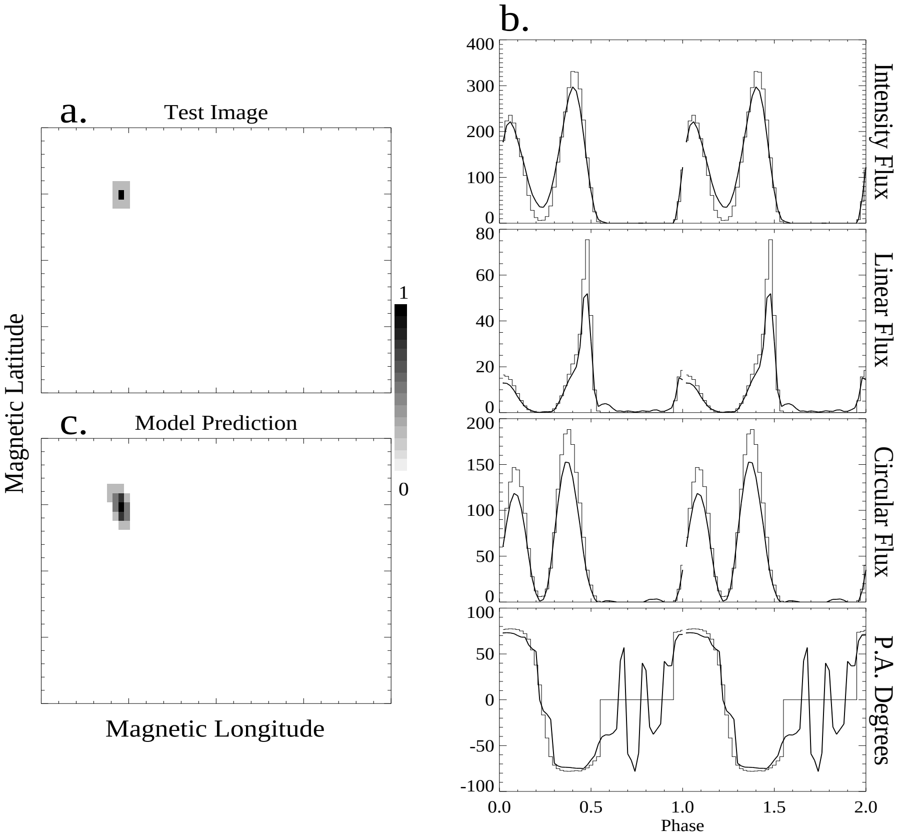

Figure 2. Point source: **a.** The emission on the surface of the white dwarf **b.** The test emission light curves (histogram) and the model fit (solid smooth curve) **c.** The predicted emission region.

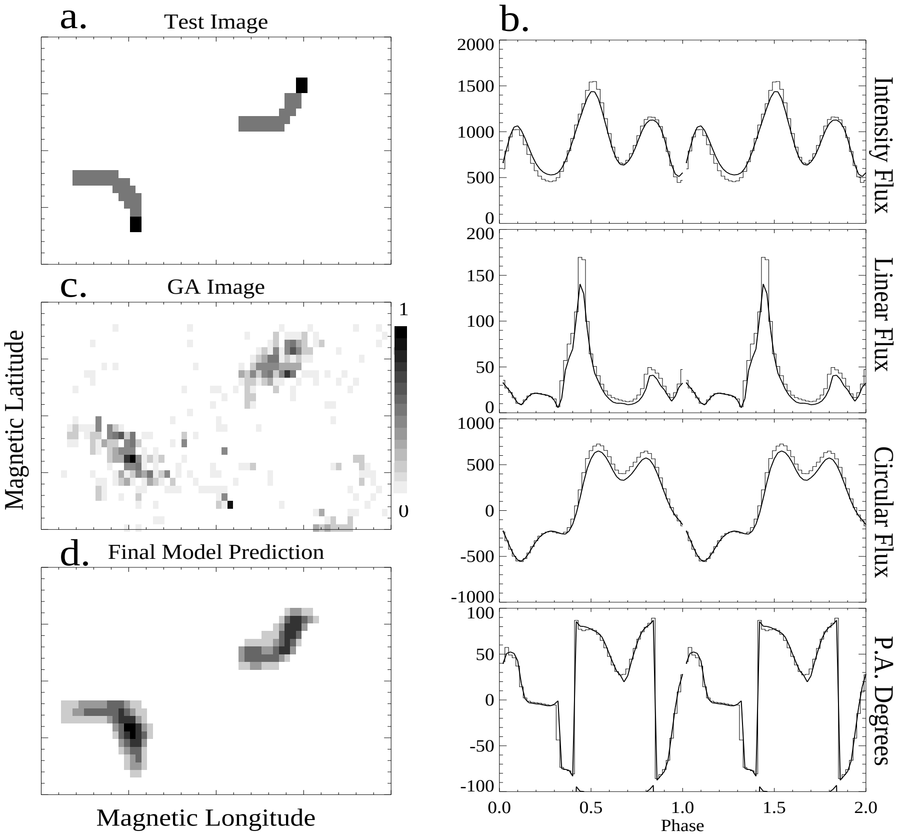

Figure 3. Extended source: **a.** The emission on the surface of the white dwarf **b.** The test emission light curves (histogram) and the model fit (solid smooth curve) **c.** GA optimised image **d.** Final model prediction after Powell's optimisation.

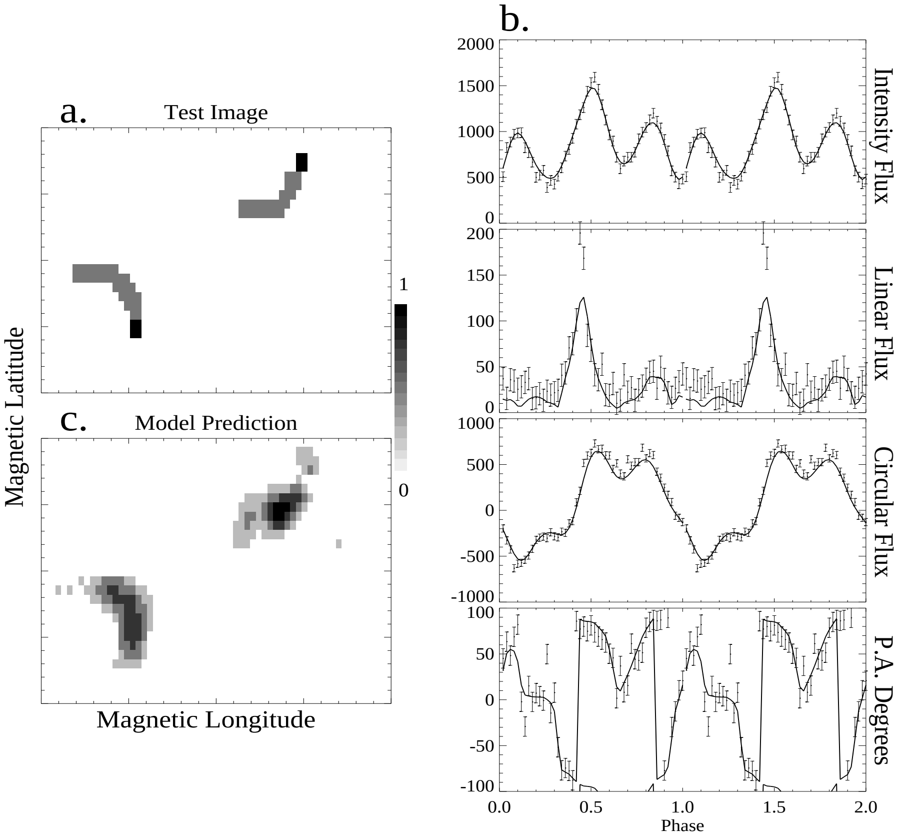

Figure 4. As for Fig.2. But for extended sources with artificial noise added.

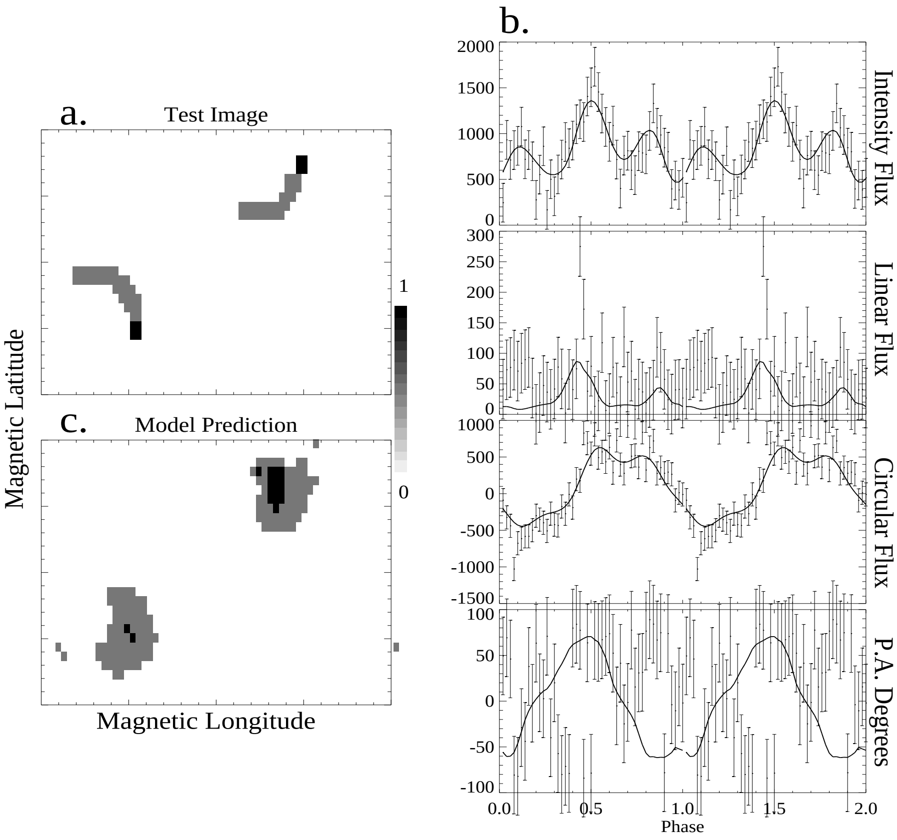

Figure 5. As for Fig.2. But for extended sources with artificial noise added.

both the positive and negative polarized light curves. The more complex linear polarization curves and the variation in position angle now extend over the whole period of rotation.

Figs. 3c and d show the optimised image solutions from the GA and the Powell's method respectively. The grey scale is as in Fig.2. As can be seen from the figures, the GA has found the location and approximate extent of the emission regions. The Powell's method has then smoothed and clearly defined these regions and also reduced the amount of other low level random emission sites.

4.3 Noisy data sets

Real data sets will always have noise to some degree superimposed on the signal. Therefore, in Figs. 4 and 5 we demonstrate how the optimisation technique handles noisy data sets. Figs. 4 and 5 use the data set shown in Fig. 3 but have different amounts of artificial noise added.

In Fig. 4b noise was introduced to each of the four Stokes parameters by using a Gaussian random number generator giving signal–to–noise ratios of 48, 3.5, 2.7 and 8.6 for the I, Q, U and V Stokes parameters respectively. This gives an average signal–to–noise ratio of $\sim$ 15.0. The error bars represent 1 sigma confidence levels. In comparison to Fig. 3, the model fit to the polarized light curves are basically the same. Upon closer inspection the sharpness of the linear polarization peaks are beginning to decrease, particularly the main peak at phase ~ 0.45. Also the map predicted by the model (Fig. 4c) is starting to show loss of spatial resolution with the sharpness of the arcs no longer so well defined. However their spatial extent and location on the white dwarf remains well determined.

Figure 5 demonstrates the effect of decreasing the average signal–to–noise ratio even further to ~ 4, corresponding to signal–to–noise ratios of $12.9, 0.94, 0.74$ and 2.3 for the I, Q, U and V Stokes parameters respectively. The prediction for the intensity light curve is still acceptable. The same is also true for the circular polarization but to a lesser extent: the dip feature at phase ~ 0.65 and the peak at phase ~ 0.1 are beginning to be lost in the noise. However, because of the lower number of counts, the noise contribution to the linear polarized light curve has a significant effect. There is no longer any detailed structure in the light curve and even the strong peak at phase $\sim$ 0.45 is no longer significantly influencing the model fit. The position angle of linear polarization is almost completely lost in the noise and the model fit is significantly different from the infinite signal–to–noise ratio data in Fig. 3.

The results of the model prediction are shown in Fig. 5c. Immediately obvious is the poor prediction for the shape and extension of the emission regions. They have become more extended and amorphous. However, they are still accurately located.

In general, for most MCVs, the linear polarization and its position angle will be the most inflicted with noise. This is true because the amount of linearly polarized flux is generally less than the intensity and circularly polarized fluxes and the intensity and circularly polarized light curves therefore have better signal–to–noise ratios. On their own, the intensity and circularly polarized light curves will put upper limits on the size of the emission regions only but can however predict the number and location of the emission regions. With linear polarimetry the shapes of the emission regions can be obtained more precisely. This is particularly true because of the sharp peak features that most MCVs produce in their linear polarized light curves. Also the variation in position angle is strongly dependent on the location of the emission sites.

5 CONCLUSIONS

We have developed a technique, using a combination of a GA and Powell's method, that produces maps of the cyclotron emission regions in AM Her systems. The technique has been applied to various simulation test cases with different source shapes and different signal–to–noise ratios. From the test cases we have shown that data with an average signal–to–noise ratio of $\sim$ 15 it is possible to determine the shape of the emission regions and their location on the surface of the white dwarf fairly accurately. Even without good linear polarization data we have shown that it is possible to locate the emission sites and put upper limits on their size.

With better signal–to–noise ratio and higher phased resolved data it is possible to obtain very detailed maps of the emission regions. For example, with 50 phase bins per rotational cycle of the white dwarf, one phase bin will correspond to just over 7 degrees of rotation by the white dwarf. To increase the spatial resolution of the optimised maps further it is necessary to obtain more phase bins per rotation cycle of the white dwarf.

Modelling of real polarimetric data will be the subject of a subsequent paper. With multi–band polarimetric data it will be possible to map the emission regions with greater accuracy. This is possible because of two reasons. Firstly there are more data points to fit which will therefore constrain the total number of possible solutions. Secondly, the shape of the emission regions are strongly wavelength dependent: solutions that are good for one wave–band may not necessarily give a good solution for another wave–band.

Having obtained maps of the accretion region and by assuming a dipole field, it is then possible to back trace the magnetic field lines that feed the emission region to the orbital plane and hence locate the threading regions where the material attaches to the magnetic field lines. Doppler mapping of emission lines can be added to obtain the velocity structure of the accretion flow and to relate this to the results obtained from the Stokes imaging.

6 ACKNOWLEDGEMENTS

SBP acknowledges the PPARC for financial support in the form of a research studentship. PJH is supported by an Academy of Finland research fellowship. We are grateful to Kinwah Wu for useful discussions during the course of this work. PJH would like to thank MSSL for hospitality during his stay.

REFERENCES

Beuermann, K., Burwitz, V., 1995, in Cape Workshop on MCVs, ASP Conf Ser, 85, 99

Beuermann, K., Stella, L. & Patterson, J., 1987 ApJ, 316, 360

Charbonneau, P., 1995, ApJS, 101, 309
Cropper, M. S., 1986, MNRAS, 222, 225
Cropper, M. S., 1989, MNRAS, 236, 935
Cropper, M. S., Sp. Sci. Rev., 1990, 54, 195
Cropper, M, & Horne, K. 1994, MNRAS, 267, 481
Hakala, P. J., 1995, A&A, 296, 164
Piskunov, N. E., Tuominen, I. & Vilhu, O. 1990, A&A, 230, 363
Potter, S. B., Cropper, M, Mason, K. O., Hough, J. H. & Bailey, J. A., 1997, MNRAS, 285, 82
Press, W. H., Teukolsky, S. A., Vetterling W.T. & Flannery, B. P., 1992, Numerical Recipes in FORTRAN, Cambridge University Press, p406.
Tikhonov, A. N., 1963 Soviet. Math. Dokl. 4, 1624
Wickramasinghe, D. T., Bailey, J. A., Meggitt, S. M. A., Ferrario L., Hough, J. H. & Tuohy I. R., 1991, MNRAS, 251, 28
Wickramasinghe, D. T., Cropper, M, Mason, K. O. & Garlick, M., 1991, MNRAS, 250, 692
Wickramasinghe, D. T., & Ferrario, L., 1988 ApJ, 334, 412
Wickramasinghe, D. T., & Meggitt, S. M. A., 1985, MNRAS, 214, 605
Wu, K., 1994, PASA, 11, 61.