Mon. Not. R. Astron. Soc. **261**, 647–656 (1993)

Microlensing light curves: a new and efficient numerical method

Geraint F. Lewis,[1] Jordi Miralda-Escudé,[1] Derek C. Richardson[1] and Joachim Wambsganss[2]

[1] *University of Cambridge, Institute of Astronomy, Madingley Road, Cambridge CB3 0HA*
[2] *Princeton University Observatory, Peyton Hall, Princeton, NJ 08544, USA*

Accepted 1992 September 28. Received 1992 September 18; in original form 1992 July 15

ABSTRACT

When a source of constant luminosity moves behind a gravitational lens composed of compact objects, its flux is observed to change due to the varying amplification of the images produced by different components of the lens. A new method is presented to calculate such microlensed light curves for any given field of point masses. The method is sufficiently efficient and fast to allow one to calculate any desired statistical property of the light curves for a large number of models in a reasonable time, with present computers. We compare the light curves generated by this method with those produced by the ray-shooting method, and find good agreement.

Key words: methods: numerical – gravitational lensing.

1 INTRODUCTION

The effect of a galaxy acting as a gravitational lens on a background source can usually be described, to a first approximation, by a smooth potential created by a continuous mass distribution which varies only on the scale of the galaxy. However, the individual stars and other compact objects introduce a graininess into the potential, and this causes each image of a point source produced by the smooth potential to be split into many images. Although these multiple images cannot be resolved with current instrumentation, since they are separated by only $\sim 10^{-6}$ arcsec, the total summed flux varies with the position of the source, so that, if the source is moving relative to the lens as seen by the observer, its flux will be observed to vary. This effect, first described by Chang & Refsdal (1979), is usually termed microlensing.

Observations of the four images of the lensed quasar 2237+0305 (Irwin et al. 1989; Corrigan et al. 1991), indicating variability on different time-scales, have proved the existence of this phenomenon. The observed variability in 2237+0305 cannot simply be intrinsic to the quasar in the absence of microlensing, since it is different and uncorrelated in each of the four images, and the time delay between the different images for this particular lens is very short, of order one day.

It then becomes of interest to find particular features or statistical properties of the observed light curves that, when compared to theoretical calculations, might enable one to derive some conclusion on either the source being lensed or the compact objects producing the microlensing. The main applications up to now have been constraint of the size of the source (Nemiroff 1988; Rauch & Blandford 1991), and the proposal that colour variations of the images can be a probe of the structure of the source (Wambsganss, Paczyński & Schneider 1990; Wambsganss & Paczyński 1991). In order to address this theoretically, one needs to compute a sufficiently large sample of light curves for many different models. This objective has proved elusive so far, due to the computational requirements to calculate the microlensing light curves.

The main difficulty is that the only way to calculate the amplification of a point source is to find all its images and add up all the amplifications. Until now, there has been no known algorithm to find all the images of a point source, or even to be sure that all of them have been found, other than a two-dimensional search with a sufficiently fine grid, which is computationally expensive (Paczyński 1986). An enormous improvement was achieved with the ray-shooting method (Kayser, Refsdal & Stabell 1986; Schneider & Weiss 1987; Wambsganss 1990), in which one does not attempt to find the images of point sources, but instead calculates rays on a grid in the lens plane, obtaining a map of the magnification in the source plane. The main disadvantage of this method to calculate light curves is that only a very small fraction of the calculated values of the magnification over the source plane are used to construct the light curves, and it is only possible to obtain a light curve with finite resolution (i.e. not for a point source, but for an extended source). What is needed is a method to find all the images of many points along a straight line in the source plane. In this paper, we present such an algorithm.

The basic idea is that it is easier to find the images of a straight line in the source plane than the images of a point. As microlensing by point masses changes only the shear, and

not the surface density, the image of a straight line consists of one curve of infinite length, plus many closed loops, all of which are connected to a star. The infinite line can easily be followed once a point is found, and the loops can simply be found by starting at each star and following them. Once the image curve going through every star has been found, we can be sure that all the images of all the points in the straight line have been found.

Thus the problem of constructing the light curve is reduced to a series of one-dimensional searches, where we follow the image curves by finding subsequent points along them. In this paper, we will discuss the method and present some examples of light curves. We will compare our results and the efficiency of our method to the ray-shooting algorithm mentioned above (Wambsganss 1990).

Witt (1993) has also thought independently of a very similar solution, based on the parametric representation of caustics (Witt 1990) to the problem of light curves of point sources. He presents alternative proofs of the basic properties of the image of a straight line. His conclusions are in basic agreement with ours.

This paper is organized as follows. In Section 2, it is proved that all the images of a straight line are in the infinite line, and loops connected to stars. In Section 3, a description of the numerical code used to implement the method is given; the inclusion of a tree code to calculate the deflection angle produced by a large number of point masses is described in the Appendix. In Section 4, we present examples of light curves, and compare our results with those obtained with the ray-shooting method. In Section 5 we discuss future applications to the microlensing observations.

2 THE IMAGE OF A STRAIGHT LINE

Let us consider a lens in the single-screen approximation, described by a smooth surface density $\Sigma(\boldsymbol{x}_l)$, where $\boldsymbol{x}_l$ is the coordinate in the image plane. The surface density varies only on scales typical of the 'macrolens' (for example, the smooth potential of a galaxy). When examined with very high resolution, there is a certain fraction of the surface density, $\Sigma_*(\boldsymbol{x}_l)$, which is not smooth and is composed of individual point masses. The total effect of the lens is then described by the superposition of a smooth lens with surface density $\Sigma(\boldsymbol{x}_l)$, and the effect of all the individual point masses, plus a negative surface density $-\Sigma_*(\boldsymbol{x}_l)$. This second component (point masses and corresponding negative smooth matter) has no effect on the macrolens properties, since the sum of the surface mass density on large scales cancels, but it introduces a granularity on small scales.

The surface density $\Sigma(\boldsymbol{x}_l)-\Sigma_*(\boldsymbol{x}_l)$ is presumed to correspond to a smooth form of dark matter. In real galaxies, the dark matter may also be clumpy on a wide range of scales. However, in this paper we will assume that the only small-scale variations of the lens potential are due to individual point masses (which change only the local shear, not the surface mass density), and that their positions are uncorrelated [this is true, to a very good approximation, for the projected positions of stars in a galaxy, except for the effects of binary and multiple systems, and star clusters (Gilmore, Reid & Hewett 1985)].

Our objective is to find the amplification at all the points along a straight line in the source plane. This gives the observed light curve, for a point source emitting a constant flux, moving with constant velocity along such a line, if the motions of the individual point masses in the lens plane are neglected (see Kundic & Wambsganss 1993 for a study of the effects of star motions).

We consider a region in the lens plane which is small compared to the scale of variation of the smooth surface density, $\Sigma(\boldsymbol{x}_l)$. If the lens consists entirely of smoothly distributed matter, the image of the straight line is simply another straight line within this small region, since the convergence and the shear are approximately constant; the curvature of the unperturbed image of the line (where 'unperturbed' means that only a smooth component is considered) is only noticeable on the scales of variation of the smooth surface density. Let us now take a circle within this small region, with the centre lying on an arbitrary point along the unperturbed image line. Consider next the effect of introducing the point masses, plus the negative surface density, $-\Sigma_*(\boldsymbol{x}_l)$, within this circle; outside the circle, the unperturbed lens is left unmodified. The image of the source line will now be a set of complicated curves. An example of the image curves of the straight line produced by a field of point masses is shown in Fig. 1. Any curve which is an image of the infinite source line cannot have an edge, so it must either have infinite length, or be a closed curve. However, far away from the edge of the circle, the influence of the point masses will become small, and the image of the line will asymptotically tend to the unperturbed image line. Therefore there is one and only one image of the source line which is an infinite curve, and this one tends asymptotically to the unperturbed image. All the other images of the source line are closed curves (Petters 1992).

Numerically, it is simple to find all the images of points along the source line which are on the image curve of infinite length: we simply find one point of this curve far away from the edge of the circle, where the curve is close to the unperturbed image line, and then we follow this curve as it wanders around inside the circle, until it comes out at the other end, and it again approaches asymptotically the unperturbed image. However, we still need to find all the closed curves as well. These closed curves could in principle be difficult to find, requiring two-dimensional searches. However, they have a property that allows them to be easily found: they are all connected to at least one star (see the example of Fig. 1). We now prove this property of the closed curves.

Let $\boldsymbol{x}_s$ be the Cartesian coordinates in the source plane, in a frame where the source line is parallel to the y-axis. Each point in the lens plane with coordinates $\boldsymbol{x}_l$ (in a frame in the lens plane having the same orientation) is mapped on to a point in the source plane, and we can consider the coordinate $\boldsymbol{x}_s$ of the mapped point as a function of $\boldsymbol{x}_l$. We call this function $f(\boldsymbol{x}_l)$. The image of the source line is then a contour of this function at a given height, and if we move the source line parallel to itself, its images will be the contours of f at different heights. The lens equation specifies that the function f is given by $f=xs=x_l-\alpha x=x_l-\partial\psi/\partial x$, where α_x is the x-component of the deflection angle, and ψ is the two-dimensional gravitational potential (cf. Schneider 1984). Taking the Laplacian of this equation, we obtain $\nabla^2 f=-2\partial k/\partial x$, where $\kappa=\Sigma/\Sigma_{cr}$ and Σ_{cr} is the critical surface density. For the type of lens that we are considering, the surface density κ only varies on scales of the smooth lens,

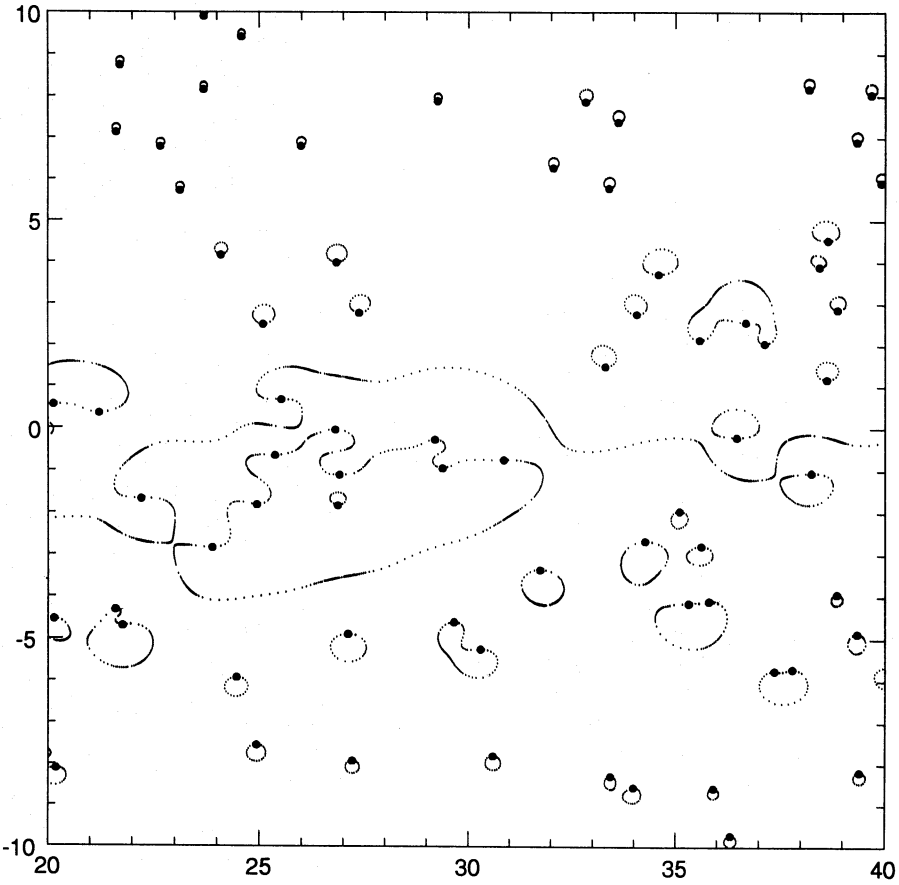

Figure 1. A sample of the image plane for a particular source line. Note the infinite curve which enters from the left and makes its way to the other side of the field, crossing a number of stars. Also note the image loops which start and end on stars. A number of these, namely those near the infinite curve and those which pass through a number of stars, are not simple in form. Those far from the infinite image curve, and far from the influence of other stars, are found to be very close to circular. The distances are in units of Einstein radii for a solar-mass star.

and we can neglect these variations. The Laplacian of f is therefore equal to zero, except on point masses, where it is equal to the derivative of a Dirac delta function.

Since a closed image curve is a contour of the function f, the gradient of f is always perpendicular to the contour, and is always directed either to the interior or to the exterior of the contour. If we integrate the gradient along the closed image curve (multiplied by the unit vector perpendicular to the curve), the integrand is always either positive or negative, so the result of the integral is not zero. However, using Gauss's theorem, this integral is equal to the integral of the Laplacian of f over the area enclosed by the contour, which is equal to zero unless one of the point masses is exactly on the contour. This shows that there is no closed image curve that does not cross a point mass.

Thus all the closed image curves can be found by searching about every star. It is simple to show that only one curve can go through any point mass. In fact, if there were two, then a critical line should lie between the two image curves arbitrarily close to the point mass, and therefore the critical line would also go across the point mass, which is impossible, since the shear is infinite on the point mass. It is also true that every point mass has an image curve going across it, since, at a sufficiently close distance to the point mass, the deflection angle can become arbitrarily large, and it is always possible to choose a position angle for a point in the lens plane with respect to the point mass, so that its mapping to the source plane will lie on the source line. Notice, however, that it is possible for a closed image curve, and for the infinite image curve, to go across more than one point mass.

3 NUMERICAL METHOD

3.1 Basic equations

For this analysis we employ the standard, dimensionless microlensing equation (Wambsganss 1990), which, when rotated into a coordinate system parallel to the external shear, can be written as

$$\boldsymbol{x}_s = \begin{pmatrix} 1-\gamma & 0 \\ 0 & 1+\gamma \end{pmatrix} \boldsymbol{x}_l - \kappa_c \boldsymbol{x}_l - \sum_{i=1}^{N_*} \frac{m_i(\boldsymbol{x}_l - \boldsymbol{x}_i)}{(\boldsymbol{x}_l - \boldsymbol{x}_i)^2}, \quad (1)$$

where $\boldsymbol{x}_s$ is the real position in the source plane and $\boldsymbol{x}_l$ is the apparent position observed in the image plane, $\boldsymbol{x}_i$ are the positions of the stars in the image plane, and m_i are their masses in solar units. All distances are expressed in units of the angular Einstein radius of a solar-mass star in the source and lens planes. The angular Einstein (or critical) radius $b_\odot$ of a solar-mass star is given by

$$\Sigma_{cr} \pi b_\odot^2 D_d^2 = M_\odot, \quad (2)$$

where D_d is the observer-lens angular diameter distance. The critical surface density, Σ_{cr}, is defined as

$$\Sigma_{cr} = \frac{c^2}{4\pi G}\frac{D_s}{D_d D_{ds}}, \tag{3}$$

where D_s and D_{ds} are the observer-source and the lens-source angular diameter distances.

The local effects of the macromodel are expressed as two parameters: the convergence, κ_c, which describes the local smoothly distributed matter in units of the critical surface density, and the shear, γ, which describes the effect of the overall matter distribution in the lens external to the microlensing region. The effect of compact bodies in the light path (e.g. stars, planets and black holes: henceforth just stars) is seen in the final term. Equation (1) maps a position in the image plane to a position in the source plane.

The calculation of the deflection angle and the amplification has to be done at many points in the lens plane in order to calculate a light curve, and this takes most of the CPU time in a calculation. We have combined our method with a tree code, which speeds up the calculation significantly (the time needed to calculate a ray with the tree code grows as $\log(N_*)$ instead of N_* for direct summation, where N_* is the number of point masses, when the tree code is used). A description of the tree code, and the changes we made for this particular application, can be found in the Appendix.

3.2 Definition of the star field

Consider a circular field cut out of the macromodel, which contains a smooth, background mass distribution, a shear term, and a distribution of stars with average surface density κ_*, expressed in units of the critical surface mass density. A point source will, in general, produce at least one image for each star in the field. However, the images at a distance R from the unperturbed image (the image that would be seen if all the mass in stars were smoothed out) have flux decreasing as R^{-4}, for large R (Paczyński 1986). Using various probabilistic arguments (Katz, Balbus & Paczyński 1986; Schneider & Weiss 1987) it is possible to determine the average number of stars, N_*, over whose images we need to sum to collect a particular fraction of the total image flux. We use

$$N_* = 300\,\frac{\langle m^2\rangle}{\langle m\rangle^2}\,\frac{\kappa_*^2}{|(1-\kappa_*-\kappa_c)^2-\gamma^2|}, \tag{4}$$

which collects over 99 per cent of the total flux. $\langle m^2\rangle$ and $\langle m\rangle^2$ are expectation values of the mass function. These stars are distributed within a circle of radius R_*, where

$$R_* = \left(\frac{N_*\langle m\rangle}{\kappa_*}\right)^{1/2}. \tag{5}$$

This radius defines the minimum region over which the stars can be distributed about the unperturbed image to achieve collection of (almost) all the flux in the final macroimage. In our case, where the source is a line, the unperturbed image is also a line. Therefore this radius defines the distance from either side of the unperturbed image line within which we need to consider image fluxes. The contribution to the light curve from images outside this region is neglected. To generate a long light curve the total star field is taken to be of radius $R_{field} = 3R_*$, and therefore $N_{Total} = 9N_*$. The light curve is effectively deduced in the region $-2R_* \rightarrow 2R_*$, ensuring that there are enough stars around the initial and final points of the light curve to collect about 99 per cent of the flux. Light curves of any length can be generated by altering these parameters.

3.3 Following the line images

3.3.1 *The infinite line*

First, the trajectory of the source is chosen. This is the source line. Then, using the mapping (1), a scalar function $f(\boldsymbol{x}_l)$, which is the separation of the associated source plane position and the source line, can be calculated over the image plane. The function $f(\boldsymbol{x}_l) = 0$ defines the set of curves which are the images of the source line. These curves are also the set of all images at all times for the passage of the source behind the star field.

To find the initial point on the infinite line, a root search is undertaken at a radius of $1.1\,R_{field}$, of the function $f(\boldsymbol{x}_l)$, effectively outside the star field, where the perturbed image line is close to the unperturbed one. The root is found by making an angular search about the central point of the star field at this radius.

From the first point, more points along the image curve are subsequently found, starting with a stepsize of 0.01 Einstein radii (for a solar-mass star). At every point, we evaluate the tangent to the image curve, which can be calculated from the amplification matrix. We use this to find a first guess for the position of the next point along the image curve. If the root is not found within an angle of $(\pi/15)$ radians of this initial guess, we divide the stepsize by a factor of two. On the other hand, if the step is completed without this division, and the root lies within $(\pi/30)$ radians, the stepsize is doubled for the next point, in order to move fast through smooth regions of the image curve. The maximum allowed stepsize is 0.2 Einstein radii for a solar mass.

The infinite curve can pass through stars. A list of the stars that are crossed is kept, and is used later when the loops have to be followed.

The infinite line travels right across the star field and exits opposite to the side that it enters. We stop following this line at a distance of $1.1\,R_{field}$, outside the region of the segment of the light curve we are calculating in the source plane. At this radius the infinite line is returning to its unperturbed form.

3.3.2 *The loops*

Once the infinite curve has been followed we start to follow the loops. Stars at a distance more than R_* away from the unperturbed image line, whose flux contribution is considered negligible, and those already crossed by the infinite line, are ignored. Each remaining star is then taken in turn, the initial point of the loop being the stellar position. The direction of the image curve very close to the star is the same as the orientation of the infinite image line far from the star field. This is used to find a first guess of position for the first point on the loop. The loops are then followed in the same way as the infinite curve.

It can be easily shown that for a particular star, far from the infinite image line, the size of the loop crossing this star is of order b_*^2/y_i, where b_* is the Einstein radius of the star and y_i is the separation from the star to the source line. As these loops have rapid directional changes, the initial stepsize is weighted with this factor to avoid many attempts of bracketing and stepsize halving to find the first point.

Again, more than one star can lie on a particular loop. When a star, other than the star used as the initial point, is crossed, it is flagged as crossed to ensure that no loop is traversed twice. The following of the loop is halted when we have returned to the position of the initial point star.

3.4 The light curves

As our numerical code follows all the image curves, with the method outlined above, we calculate the amplification $A(\boldsymbol{x}_l)$ at each point found in the lens plane. This amplification is the contribution from a particular image to the flux from the mapped source positions of the points found in the lens plane.

To generate the light curve, we divide the source line into a number of bins. For any two subsequent points along the image curve, we calculate the two corresponding source positions and the amplifications at those points, and we interpolate the amplification for all the bins between these two positions. The values of the interpolated amplifications at the centre of each bin are added to that bin (for the initial and final bin, the amplifications are multiplied by the fraction of the bin filled by the segment joining the two source positions).

The interpolation of the amplifications is done differently in three different regimes.

(1) If the amplifications of the two points are both less than 0.2, we assume that we are close to a point mass, and we

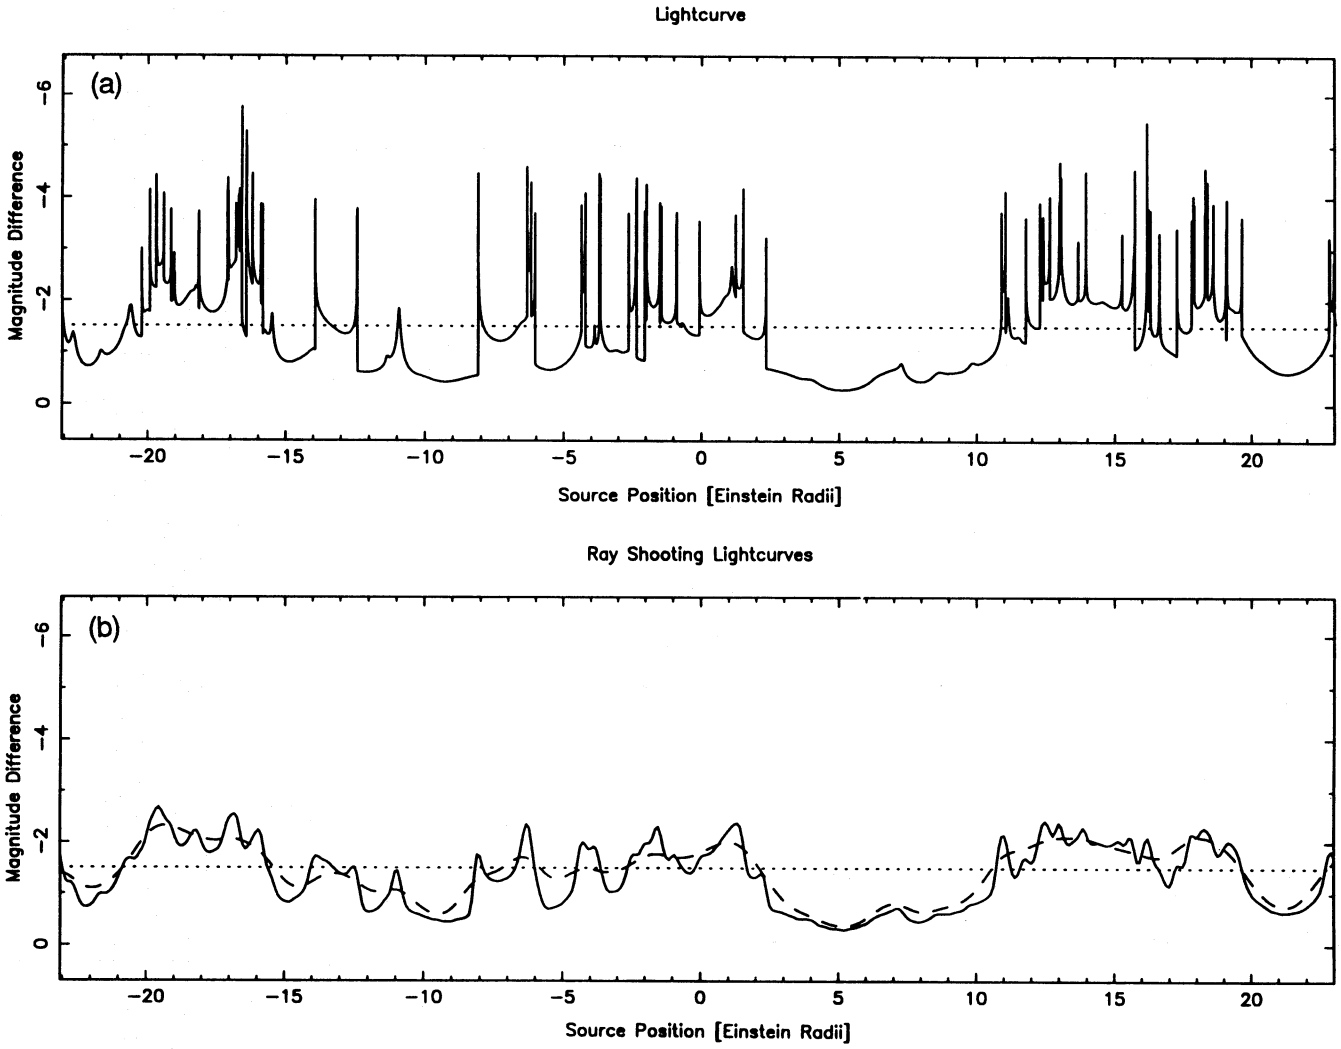

Figure 2. The light curve for $\kappa_* = 0.5$, over the range of -25 to 25 Einstein radii in the source plane, using both the method of this paper (a) and the ray-shooting method of Wambsganss (1990) (b). The x_2-axis represents the magnitude difference between an unlensed source and the images observed through the gravitational lens. The dotted, horizontal line is the astronomical magnitude difference between the unlensed source and the image produced if all the mass in compact bodies is smoothed out. For this case this is $A_s = -1.51$ mag. In Fig. 2(b), the light curves are from a magnification pattern with pixel size equal to 0.1 Einstein radii. The two light curves illustrated are for Gaussian source profiles of 1 (solid line) and 5 (dashed line) pixel half-widths.

interpolate linearly the quantity $A^{-1/4}$, which varies linearly with the source position for an image close to a point mass.

(2) If at least one of the two amplifications is larger than 5 A_s, where A_s is the amplification produced by the smooth lens, then we interpolate linearly the quantity A^{-2}, which varies linearly with source position close to a caustic.

(3) If neither of the two conditions above is true, then we interpolate the amplification directly.

In order to obtain an accurate light curve, one needs to make sure that the error made in interpolating the amplification is small. The simplest requirement is that the fractional change in the amplification is small: if $(\delta A)/A > \varepsilon$, where ε is a parameter that can be varied depending on the accuracy that is needed for the light curve, then the stepsize in the lens plane is divided by two, and the calculation is repeated until the points in the lens plane are sufficiently close that the amplification varies slowly enough. However, close to point masses and caustics, the amplification goes to zero and infinity respectively, so this requirement cannot be maintained in order to go across point masses and critical lines in a finite number of steps. To solve this problem, we have adopted the criterion that a step is sufficiently small when the amplifications A_1 and A_2 at the two points obey the following conditions in the same three regimes given above:

(1) $|A_1 - A_2| < 0.2\,\varepsilon$;

(2) $|A_1^{-1} - A_2^{-1}| < \varepsilon/(5A_s)$;

(3) $|(A_1 - A_2)/\min(A_1, A_2)| < \varepsilon$.

In the calculations presented in this paper, we have adopted $\varepsilon = 0.2$. These conditions are combined with the angular criteria in Section 3.3.

When a step crosses a caustic the sign of one of the eigenvalues of the amplification matrix changes. We then find the exact position of the caustic in the source plane. The addition of amplifications to the bins in the light curve is then done in two steps. First, the amplification is interpolated as A^{-2} between the initial point of the step and the caustic, where $A^{-2} = 0$. The amplification is then similarly interpolated between the caustic and the final point of the step.

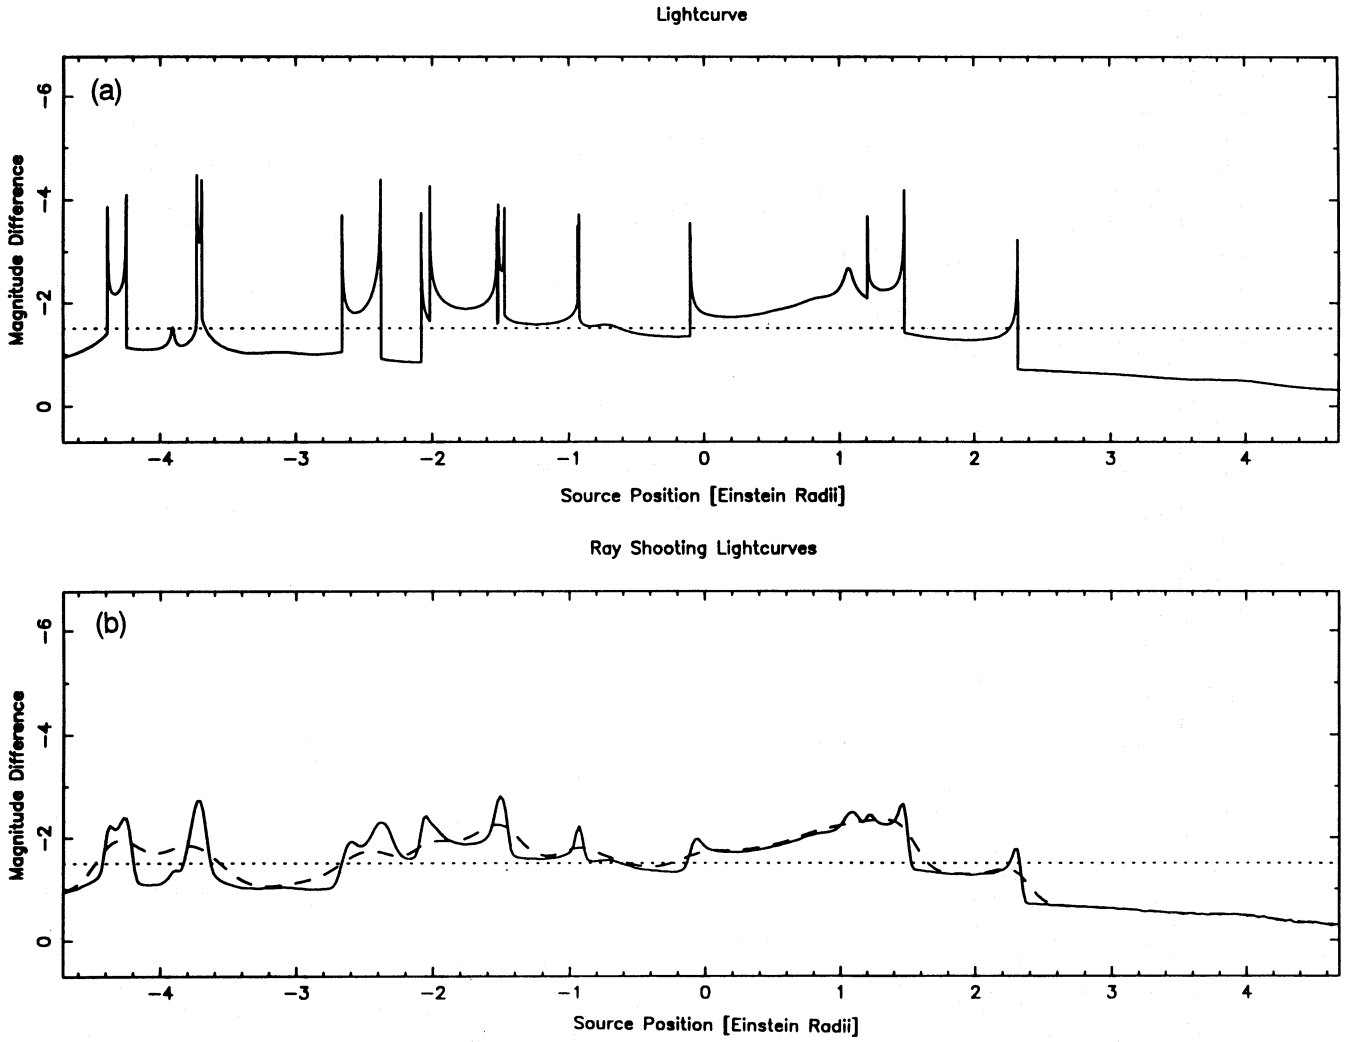

Figure 3. Enlarged central part of Fig. 2 (range −5 to 5 Einstein radii in the source plane). (a) is just an amplified version of the central region of Fig. 2(a), while (b) is a different light curve obtained from a magnification pattern with pixels of 0.02 Einstein radii. The two light curves illustrated are for Gaussian source profiles of 1 (solid line) and 5 (dashed line) pixel half-widths.

4 RESULTS

As an illustrative example of the application of this method, we present light curves for the case where all compact objects have the same mass, and their average surface density is equal to 0.5 times the critical value ($\kappa_* = 0.5$). No additional smooth surface density and no large-scale shear are added (i.e. $\kappa_c = 0$ and $\gamma = 0$ in equation 1). We generate the star field as described in Section 3.2. The total number of stars in this case was 2700.

In Fig. 1 we illustrate a small section of the star field. The large dots indicate the positions of the stars, and the small dots are the subsequent points along the images of the straight line found by our code. The infinite line, passing from the left to the right of the field, and the loops crossing the stars can be seen. Some loops, namely those crossing more than one star, or those near the infinite line, are seen to be of quite complicated form, while stars far from the infinite line have loops which are very close to circular. This result can be proved analytically. Points are not uniformly spaced along the curves, reflecting our adaptive stepsize (Sections 3.3 and 3.4).

In Figs 2, 3 and 4 we present light curves. Figs 2(a), 3(a) and 4(a) show the light curve for a point source, generated using the method outlined in this paper. The light curve required a total of 206 832 calculations of the bending angle and the amplification at different positions in the lens plane as points along the image curves were found. The light curve in Fig. 2(a) consists of 50 000 bins to provide sufficient resolution. Figs 3(a) and 4(a) are enlargements, at different scales, of the central part of Fig. 2(a).

Figs 2(b), 3(b) and 4(b) are light curves obtained as one-dimensional cuts through two-dimensional magnification patterns. In this method, a region of the source plane is divided into many pixels. The deflection angle is calculated on a grid in the lens plane, and one counts the number of rays that reach each pixel. The size of the pixel is the limiting resolution of the two-dimensional magnification pattern obtained. Thus the light curves can only be obtained for extended sources, with a size comparable to or larger than the pixel size (for full details of this method see Wambsganss 1990). Three magnification patterns, consisting of 500×500 pixels, of side length 49, 10 and 2.5 Einstein radii in the source plane, were calculated. This gives a pixel size of 0.1,

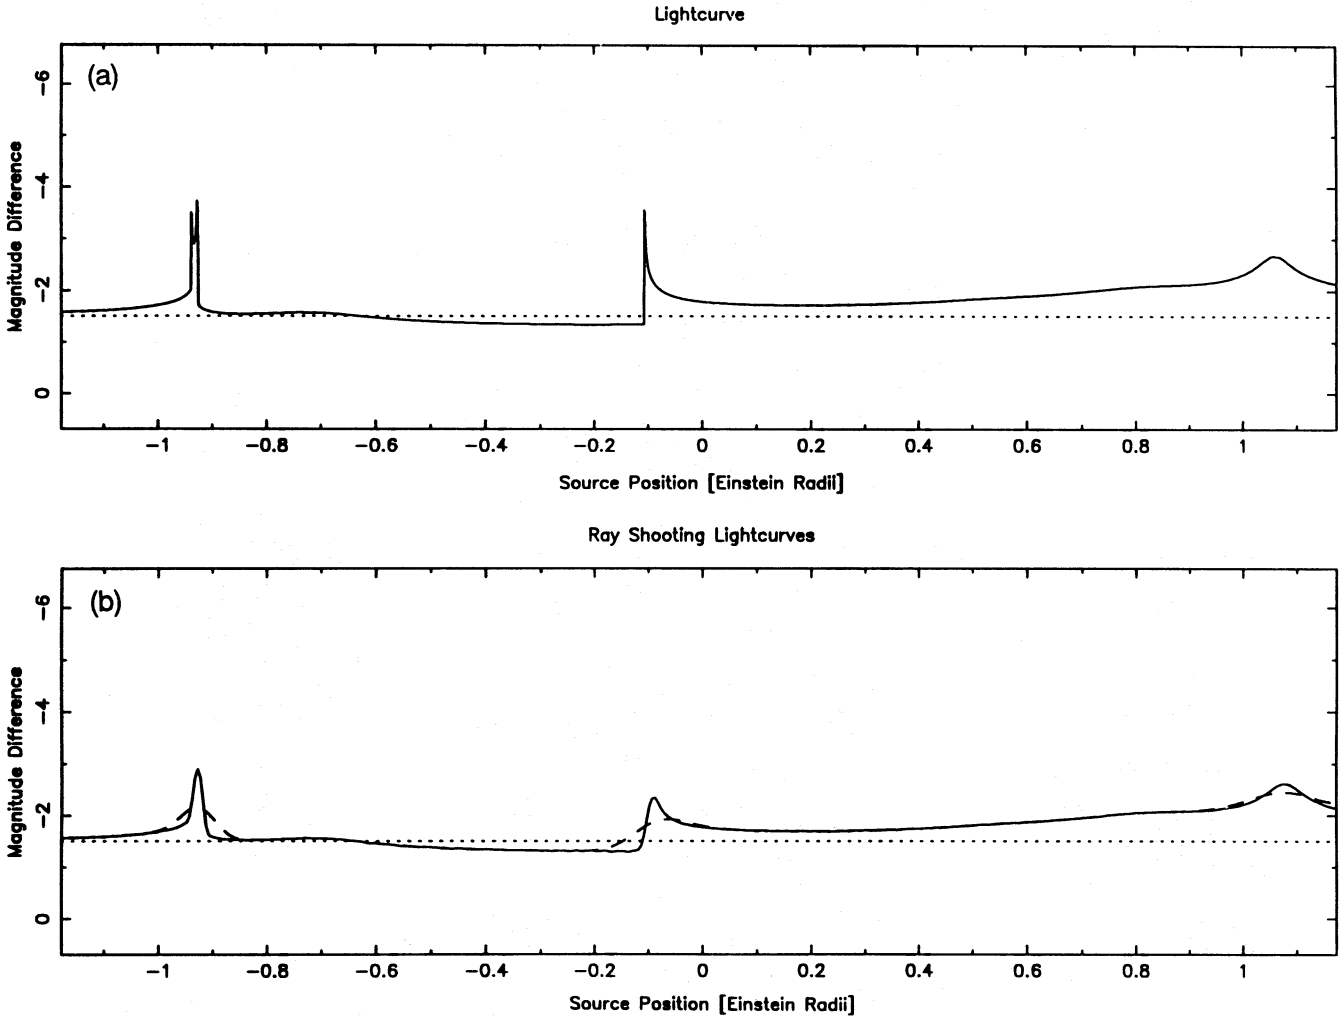

Figure 4. Enlarged central part of Fig. 2 (range -1.7 to 1.7 Einstein radii in the source plane). (a) is just an amplified version of the central region of Fig. 2(a), while (b) is a different light curve obtained from a magnification pattern with pixels of 0.005 Einstein radii. The two light curves illustrated are for Gaussian source profiles of 1 (solid line) and 5 (dashed line) pixel half-widths.

0.02 and 0.005 Einstein radii. The two light curves illustrated in each plot are for Gaussian sources of 1 (solid line) and 5 (dashed line) pixel half-widths. Each pair of light curves was generated from a different calculation of the two-dimensional magnification pattern, with increasing resolution. Note that the light curves obtained from the magnification patterns are for *extended* sources, whereas our method obtains them for *point* sources.

In Fig. 5 we present the two-dimensional magnification pattern used to generate the light curve illustrated in Fig. 3(b). The track indicates the path of the source. The events seen in the light curve in Fig. 3(a) can be identified with the features seen in the pattern, such as the sharp boundaries outlining caustics.

As can be seen, there is very good qualitative agreement between the light curves obtained using the two different methods.

5 DISCUSSION

In this paper we have presented an efficient and fast method for the generation of microlensing light curves. Our method

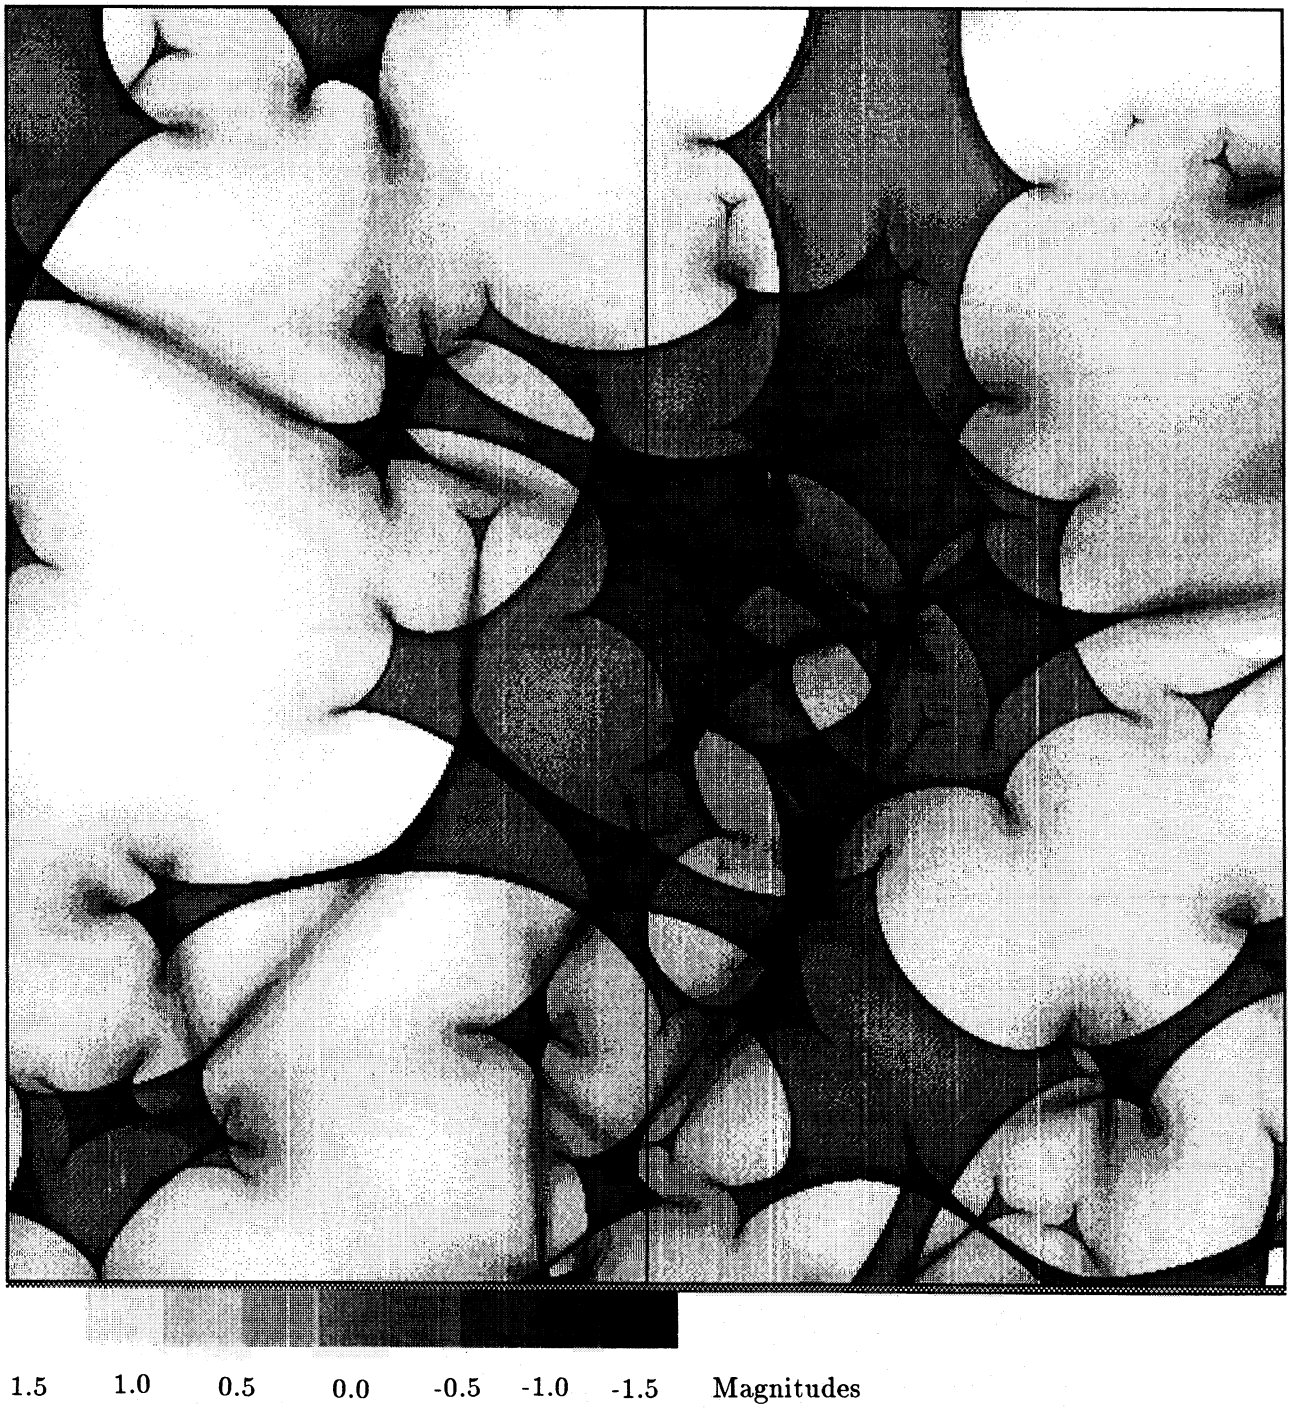

Figure 5. The magnification pattern produced with the ray-shooting method, used to obtain the light curve in Fig. 3(b).

is most efficient for point sources, but light curves for small extended sources can also be obtained by simply computing several light curves separated by a short distance and possibly calculating the derivative of the amplification in a direction perpendicular to the source trajectory. For a large enough source, some combination of the ray-shooting method and the method presented here should be most efficient.

The calculations of Wambsganss (1990) with the ray-shooting method indicate that, with a modern vector machine, a calculation of the bending angle takes of order 100 μs. This time does not grow very much with the number of point masses in the lens plane, if a tree code is being used, as in our case. Thus a calculation involving 10^9 rays can be done in about one day of CPU time, which is feasible. A calculation of 10^{10} rays would already involve enough CPU time to cause a significant slowing of the normal pace of a scientific project (usually, numerical calculations have to be repeated several times before one obtains the desired result). The magnification patterns can provide, in principle, a large number of light curves ($\approx$ few hundred) for each calculation, by taking several cuts. However, these light curves would not be statistically independent: they cross the same groups of caustics, and maxima are correlated.

With our method, we have needed 2×10^5 rays to calculate the light curve presented here. If this is considered as a typical light curve, then one can calculate 5000 light curves with the computation of 10^9 rays. Even if the number of rays needed to generate a light curve is larger for some models than in the example presented in this paper, we see that the problem of numerically evaluating the probability distribution of any observable quantity in a microlensing system from a large enough number of simulated light curves, for many different models of the lens and the source, should no longer present a computational difficulty.

In future work, we plan to examine the constraints one can place on the lens and the source in the 2237+0305 lens, where flux variations produced by microlensing have been observed, and to investigate the observational strategy in this and other lenses that would be most likely to lead to additional results.

ACKNOWLEDGMENTS

The authors would like to thank Robert Cannon, Peter Tribble and Simon White for useful discussions. We would also like to thank Hans-Jörg Witt, Peter Schneider and the referee for useful comments on the manuscript. GFL and JM are supported by SERC, DCR is supported by the British Council through a Commonwealth Scholarship and JW is supported by NSF grant AST 90-23775.

REFERENCES

Chang K., Refsdal S., 1979, Nat, 282, 561
Corrigan R. T. et al., 1991, AJ, 102, 34
Gilmore G., Reid N., Hewett P. C., 1985, MNRAS, 213, 257
Hernquist L., 1987, ApJS, 64, 715
Irwin M. J., Webster R. L., Hewett P. C., Corrigan R. T., Jedrzejewski R. I., 1989, AJ, 98, 1989
Katz N., Balbus S., Paczyński B., 1986, ApJ, 306, 2
Kayser R., Refsdal S., Stabell R., 1986, A&A, 166, 36
Kundic T., Wambsganss J., 1993, ApJ, in press
Nemiroff R. J., 1988, ApJ, 335, 593
Paczyński B., 1986, ApJ, 301, 503
Petters A. O., 1992, J. Math. Phys., 33, 1915
Rauch K. P., Blandford R. D., 1991, ApJ, 381, L39
Schneider P., 1984, A&A, 140, 119
Schneider P., Weiss A., 1987, A&A, 171, 49
Wambsganss J., 1990, PhD thesis, München
Wambsganss J., Paczyński B., 1991, AJ, 102, 864
Wambsganss J., Paczyński B., Schneider P., 1990, ApJ, 358, L33
Witt H. J., 1990, A&A, 236, 311
Witt H. J., 1993, ApJ, submitted

APPENDIX A: THE HIERARCHICAL TREE CODE

In order to use the tree code, we need an expression for the multipole expansion of a $1/r$ law. For any given box of lenses with total mass

$$m_{\rm b}=\sum_l m_l$$

and centre-of-mass position $\boldsymbol{x}_{\rm b}$, the bending angle due to any lens l in the box can be expanded about $\boldsymbol{x}_{\rm b}$ in a Taylor series given by

$$\boldsymbol{\alpha}(\boldsymbol{x}_l)=\boldsymbol{\alpha}(\boldsymbol{x}_{\rm b})+\sum_i x_{l,i}\left[\frac{\partial\boldsymbol{\alpha}(\boldsymbol{x}_l)}{\partial x_{l,i}}\right]_{\boldsymbol{x}_l=\boldsymbol{x}_{\rm b}}+\frac{1}{2}\sum_{i,j}x_{l,i}x_{l,j}\times\left[\frac{\partial^2\boldsymbol{\alpha}(\boldsymbol{x}_l)}{\partial x_{l,i}\partial x_{l,j}}\right]_{\boldsymbol{x}_l=\boldsymbol{x}_{\rm b}}+\ldots, \tag{A1}$$

where

$$\boldsymbol{\alpha}(\boldsymbol{x}_l)=m_l\left(\frac{\boldsymbol{x}_0-\boldsymbol{x}_l}{|\boldsymbol{x}_0-\boldsymbol{x}_l|^2}\right) \tag{A2}$$

is the deflection of a light ray at $\boldsymbol{x}_0$ due to lens l from equation (1). Note that all distances are in the lens plane and the subscript l here pertains to an individual lensing mass. When summed over all masses in the box, the first term becomes the monopole contribution, the second term (dipole) vanishes, and the third term becomes the quadrupole contribution. The result is most conveniently expressed in the following form:

$$\alpha_k=m_{\rm b}\frac{x_k}{|\boldsymbol{x}|^2}+\frac{1}{3}\sum_{i,j}(Q_{ij}P_k^{ij}), \tag{A3}$$

where for the two-dimensional case

$$\boldsymbol{x}=(x_1,x_2)=\boldsymbol{x}_0-\boldsymbol{x}_{\rm b}, \tag{A4}$$

$$Q_{ij}=\sum_l m_l(3x_{l,i}x_{l,j}-x_l^2\delta_{ij}), \tag{A5}$$

$$\mathbf{P}_1=\frac{1}{|\boldsymbol{x}|^6}\begin{pmatrix}x_1^3-3x_1x_2^2 & 3x_1^2x_2-x_2^3\\ 3x_1^2x_2-x_2^3 & 3x_1x_2^3-x_1^3\end{pmatrix}, \tag{A6}$$

$$\mathbf{P}_2=\frac{1}{|\boldsymbol{x}|^6}\begin{pmatrix}3x_1^2x_2-x_2^3 & 3x_1x_2^2-x_1^3\\ 3x_1x_2^2-x_1^3 & x_2^3-3x_1^2x_2\end{pmatrix}. \tag{A7}$$

We use only the quadrupole term to calculate the deflection angle from cells. Note that $\boldsymbol{Q}$, $\mathbf{P}_1$, and $\mathbf{P}_2$ all contain sym-

metries that can be exploited for improved efficiency. In fact, $\boldsymbol{Q}=(Q_{ij})$ is the quadrupole moment tensor and has the recursive property (Hernquist 1987)

$$\boldsymbol{Q}=\sum_{b}^{N_{subbox}} \boldsymbol{Q}_b+\sum_{b}^{N_{subbox}} m_b(3\boldsymbol{R}_b\boldsymbol{R}_b-R_b^2\mathbf{I}), \tag{A8}$$

that is, $\boldsymbol{Q}$ for any given box of lenses can be written in terms of the quadrupole moments of its constituent subboxes. Here $\boldsymbol{R}_b$ is the displacement vector $\boldsymbol{x}_b-\boldsymbol{x}$ between the centre of mass of subbox b and the centre of mass of the parent box, and $\mathbf{I}$ is the unit matrix. To obtain the light curves, we also need to calculate the amplification at any point in the lens plane, apart from the deflection angle. This is obtained from the derivatives of the deflection angle, which can be calculated using the tree code in the same manner as $\boldsymbol{\alpha}$, and the necessary formulae are obtained by taking the appropriate derivatives of the monopole and quadrupole terms.

A problem encountered in the implementation of the tree code is related to the discontinuities in the deflection angle, produced by changes in the tree structure as one moves around the lens plane: when a cell is substituted by smaller subcells, the error made in approximating the contribution from each star in the cells, using the quadrupole moments given above, changes. Normally, in a tree code, a cell is used to calculate the force if it subtends an angle θ smaller than a critical value θ_{crit}; if $\theta>\theta_{crit}$, smaller cells are used. This procedure introduces a discontinuity at the positions where a cell subtends an angle $\theta=\theta_{crit}$. The images of the straight line, calculated with the tree code, are therefore also discontinuous at these positions. Since our code follows the image curves by requiring a smooth change of the tangent vector along them, the presence of such discontinuities would prevent the image curves being followed. The solution we have adopted is to introduce a region $\theta_{crit}-\delta\theta<\theta<\theta_{crit}+\delta\theta$, inside which the contribution to the bending angle from a cell is calculated first as though $\theta<\theta_{crit}$ (where the multipole expansion of the cell is used), and then as though $\theta>\theta_{crit}$ (where the immediately smaller cells down the tree hierarchy are used). The two results are weighted and added together according to

$$\boldsymbol{\alpha}=\left[\frac{(\theta_{crit}+\delta\theta)-\theta}{2\delta\theta}\right]\boldsymbol{\alpha}_{(\theta<\theta_{crit})}+\left[\frac{\theta-(\theta_{crit}-\delta\theta)}{2\delta\theta}\right]\boldsymbol{\alpha}_{(\theta>\theta_{crit})}. \tag{A9}$$

We have generally used $\theta_{crit}=0.5$, and $\delta\theta=0.05$. The method has worked well with negligible CPU expense.