

VBBINARYLENSING: a public package for microlensing light-curve computation

Valerio Bozza,^{1,2★} Etienne Bachelet,³ Fran Bartolić,⁴ Tyler M. Heintz,⁵ Ava R. Hoag⁵ and Markus Hundertmark⁶

¹Dipartimento di Fisica “E.R. Caianiello”, Università di Salerno, Via Giovanni Paolo II 132, Fisciano, I-84084, Italy

²Istituto Nazionale di Fisica Nucleare, Sezione di Napoli, Via Cintia, I-80126 Napoli, Italy

³Las Cumbres Observatory, 6740 Cortona Drive, Suite 102, Goleta, CA 93117, USA

⁴Centre for Exoplanet Science, SUPA, School of Physics and Astronomy, University of St. Andrews, St. Andrews KY16 9SS, UK

⁵Department of Physics, Westminister College, New Wilmington, PA 16172, USA

⁶Astronomisches Rechen-Institut, Zentrum für Astronomie der Universität Heidelberg (ZAH), D-69120 Heidelberg, Germany

Accepted 2018 June 28. Received 2018 June 28; in original form 2018 May 7

ABSTRACT

The computation of microlensing light curves represents a bottleneck for the modelling of planetary events, making broad searches in the vast parameter space of microlensing extremely time-consuming. The release of the first version of VBBINARYLENSING (based on the advanced contour integration method presented in a previous paper) has represented a considerable advance in the field, with the birth of several analysis platforms running on this code. Here, we present the version 2.0 of VBBINARYLENSING, which contains several upgrades with respect to the first version, including new decision trees that introduce important optimizations in the calculations. The package is publicly available at <https://github.com/valboz/VBBinaryLensing>.

Key words: gravitational lensing; micro – methods: numerical.

1 INTRODUCTION

Microlensing has become one of the main methods for detecting extrasolar planets (Gould 2016). After the *Kepler* harvest (Coughlin et al. 2016), we have now a pretty large statistics on planets close to their parent stars, with periods below one year¹ (Christiansen 2018). Radial velocities also enabled the discoveries of some gas giants further away and enriched the table of habitable worlds. Direct imaging has unveiled giant planets in the outer regions of young systems (Fischer et al. 2014). At present, microlensing stands as the only way to explore the colder disc beyond the snow line, where most of the planetary embryos are believed to form through accretion of volatiles (Alibert et al. 2010).

Yet, microlensing events only probe the magnification maps generated by the lensing systems along a single line: the actual source trajectory. Different lens configurations may justify the same observed light curve, inducing severe discrete and continuous degeneracies in the parameters (Griest & Safizadeh 1997). Furthermore, a sparse sampling may complicate the problem with accidental degeneracies due to gaps in the coverage (Skowron et al. 2018). For all these reasons, the full exploration of the parameter space is excep-

tionally complicated compared to other astrophysical problems and requires the calculation of thousands of models in order to probe far apart minima in such a vast parameter space.

In addition to this, the calculation of each individual light curve takes a long computational time, because the magnification calculation has to be repeated at the epoch of each observation. If thousands of observations have been collected, we have to make thousands of magnification calculations just to evaluate an individual model.

More in detail, for any position of the source along its trajectory, in order to obtain its magnification, we need to draw the corresponding images generated by the lens and estimate their extension. For point sources, this is relatively easy, since we can recast the lens equation into a complex polynomial of order $n^2 + 1$ (n is the number of lenses), which can be solved numerically for any given source position (Witt 1990).

However, physical sources are stars with their own angular size, which cannot be neglected especially in regions where the magnification varies very quickly or is discontinuous (Griest 1991; Nemiroff & Wickramasinghe 1994; Witt & Mao 1994). The algorithms solving this problem can be roughly distinguished in two classes: inverse ray shooting and contour integration.

In inverse ray shooting (Kayser, Refsdal & Stabell 1986; Wambsganss 1990; Bennett & Rhie 1996; Bennett 2010), light rays are shot back from the observer to the lens plane. Only those rays landing within the source disc are counted for the magnification. In this

* E-mail: valboz@sa.infn.it

¹See online catalogues: <http://exoplanets.eu> (Schneider et al. 2011), <https://exoplanetarchive.ipac.caltech.edu> (Akeson et al. 2013).

case, we do not have to invert the lens map. A whole magnification map on the source plane can be constructed by shooting rays on a fine grid on the lens plane. Such maps can be re-used for different source positions.

In contour integration (Schramm & Kayser 1987; Dominik 1995; Gould & Gaucherel 1997; Dominik 1998), the lens map is inverted on a collection of points on the source boundary. In this way, one obtains a corresponding collection of points on the boundaries of the images, from which the area of the images can be recovered by use of Green's theorem. In principle, this method converts a two-dimensional problem (area of the images) to a one-dimensional problem (a summation over points on the boundary). At least for uniform-brightness sources, this solution promises to be quite effective.

In a previous work (Bozza 2010), we have presented several solutions to make contour integration much more efficient. The accuracy of the summation is improved by a parabolic correction; an error estimate is introduced on each arc of the boundary; this allows the definition of an optimal sampling, where new points are introduced on sections of the boundary in which we find the largest error; limb darkening can be included by repeating the calculation on more annuli, adaptively introduced when needed. All these optimizations have been implemented in a code that has been at the core of RTModel,² a completely automatic real-time modelling platform running on an eight-core workstation.

In the year 2016, a first version of the contour integration code has been made available on a public webpage with the name VBBINARYLENSING.³ After almost two years of extensive testing by the community, we are now ready to release version 2.0 of this code, including the implementation of some more optimizations and features to be described in this paper. The presentation of such ideas may be of interest to the whole microlensing community even from a purely theoretical perspective. The code is now freely available also on a GitHub repository⁴ under the lesser GPL 3.0 license. It is possible to compile it as a simple C++ library or to install a PYTHON package with all the functionality of the C++ code. The README.md file contains basic instructions, while several example files illustrating the use of the library from both PYTHON and C++.

In Section 2, we review contour integration and its optimizations at the base of VBBINARYLENSING. In Section 3, we discuss the details of the new root solving routine. In Section 4, we compare the effects of requiring a relative precision goal versus an absolute accuracy goal. In Section 5, we present three new tests based on the quadrupole approximation and ghost images used to decide whether the point-source approximation is sufficient or a full contour integration is necessary. Section 6 illustrates the calculation of the magnification for a single lens and a finite source. In Section 7, we describe the full light-curve calculations, explaining how parallax, orbital motion, and microlensing from satellites are implemented in the code. Section 8 contains the conclusions. An appendix lists all methods and properties of the VBBINARYLENSING library.

2 CONTOUR INTEGRATION

VBBINARYLENSING is a code based on a new advanced contour integration algorithm (Bozza 2010). In this section, we briefly summarize the core ideas and refer the reader to that paper for more details.

²<http://www.fisica.unisa.it/GravitationAstrophysics/RTModel.htm>

³<http://www.fisica.unisa.it/GravitationAstrophysics/VBBinaryLensing.htm>

⁴<https://github.com/valboz/VBBinaryLensing>

For a uniform-brightness source, the microlensing amplification is just the ratio of the total angular area of the images to the source angular area

$$\mu = \frac{A_{\text{Images}}}{A_{\text{Source}}}. \quad (1)$$

We use standard coordinates for the lens/image plane and for the source plane normalized to the Einstein radius of the total lens mass. The binary lens equation reads

$$\mathbf{y} = \mathbf{x} - m_1 \frac{\mathbf{x} - \mathbf{x}_{m_1}}{|\mathbf{x} - \mathbf{x}_{m_1}|^2} - m_2 \frac{\mathbf{x} - \mathbf{x}_{m_2}}{|\mathbf{x} - \mathbf{x}_{m_2}|^2}, \quad (2)$$

where $m_1 = 1/(1+q)$, $m_2 = q/(1+q)$, and q is the mass ratio. The lens positions can be chosen to be $\mathbf{x}_{m_1} = (-s, 0)$, $\mathbf{x}_{m_2} = (0, 0)$, so that s is the lens separation and the origin of the reference frame is in the smaller mass ($q < 1$). This choice of the origin turns out to be the most convenient for the root solving routine (see Section 3).

For a source centred at position \mathbf{y}_S in the source plane and radius ρ_* , the source boundary can be sampled as

$$\mathbf{y}_i = \mathbf{y}_S + \rho_* \begin{pmatrix} \cos \theta_i \\ \sin \theta_i \end{pmatrix}, \quad (3)$$

where $\theta_0 < \theta_1 < \dots < \theta_i < \dots < \theta_n = \theta_0$ is a collection of n angles.

By solving the lens equation, we can determine the images of \mathbf{y}_i and indicate them by $\mathbf{x}_{I,i}$. It is well known that for a binary lens there exist caustic curves in the source plane enclosing regions in which a point source generates five images, while outside the caustics only three images are generated. The index I then takes values from 1 to 3 if \mathbf{y}_i is outside any caustics or 1–5 if it is inside a caustic.

Having identified the image points belonging to the same image boundary, we can obtain the total area of the images as

$$A_{\text{Images}} = \sum_I p_I \sum_{i=0}^{n-1} A_{I,i}^{(t)} + A_{I,i}^{(p)} \quad (4)$$

$$A_{I,i}^{(t)} = \frac{1}{2} \mathbf{x}_{I,i} \wedge \mathbf{x}_{I,i+1} \quad (5)$$

$$A_{I,i}^{(p)} = \frac{1}{24} [(\mathbf{x}'_{I,i} \wedge \mathbf{x}''_{I,i}) + (\mathbf{x}'_{I,i+1} \wedge \mathbf{x}''_{I,i+1})] \Delta\theta^3, \quad (6)$$

where $p_I = \pm 1$ is the parity of image I . The trapezium approximation $A_{I,i}^{(t)}$ replaces each arc between θ_i and θ_{i+1} by a straight line. The parabolic correction $A_{I,i}^{(p)}$ adds the area of the parabolic segment between the curve and the straight line as found by checking the local derivatives at the end points of the arc. As shown by Bozza (2010), these derivatives can be easily expressed from the lens equation and used to dramatically improve the accuracy of the contour integration without increasing the sampling. Similar expressions can be given for image boundaries starting or ending at a critical point, generated by a source whose boundary is partially inside a caustic.

The magnification is then obtained by equation (1) with A_{Images} given by equation (4) and the area of the source simply being $A_{\text{Source}} = \pi\rho_*^2$ for a uniform source.

At each step, it is essential to keep track of the error committed by our approximation of the contour integral by a finite sum. This is done by defining the following error estimators for each image arc

$$E_{I,i,1} = \frac{1}{48} |(\mathbf{x}'_{I,i} \wedge \mathbf{x}''_{I,i}) - (\mathbf{x}'_{I,i+1} \wedge \mathbf{x}''_{I,i+1})| \Delta\theta^3 \quad (7)$$

$$E_{I,i,2} = \frac{3}{2} \left| A_{I,i}^{(p)} \left(\frac{|\mathbf{x}_{I,i} - \mathbf{x}_{I,i+1}|^2}{\Delta\theta^2 |\mathbf{x}'_{I,i} \cdot \mathbf{x}'_{I,i+1}|} - 1 \right) \right| \quad (8)$$

$$E_{I,i,3} = \frac{1}{10} |A_{I,i}^{(p)}| \Delta\theta^2, \quad (9)$$

where $\Delta\theta = \theta_{i+1} - \theta_i$. The first estimator checks that the parabolic approximation from the two end points of the same arc returns a consistent result; the second estimator checks that the distance between the two end points is consistent with the extrapolation from the local derivatives; and the third one just checks that the sampling is fine enough that $A_{I,i}^{(p)}$ is a truly higher order correction to $A_{I,i}^{(l)}$ and is not exploding. For each of these estimators, there is a version for image boundaries starting or ending at a critical point.

In addition to these estimators, we have also introduced a check on ghost images, i.e. the two roots of the fifth-order complex polynomial associated with the lens equation that do not correspond to physical images when the source is outside the caustic. If the distance between such ghost images has a minimum at some step i in our sampling, we add an error

$$E_{g,i} = (g_{i-1} - g_i)^2 \quad (10)$$

$$g_i = |\mathbf{x}_{g1,i} - \mathbf{x}_{g2,i}|, \quad (11)$$

and similarly for $E_{g,i-1}$. In this way, we caution that the tip of a cusp may be hidden in between our sampling.

The error estimators are essential for driving the sampling where it is actually needed. Starting from a basic sampling with two points, we gradually increase it by introducing a new θ_i partitioning the arc in which the error of the corresponding images is largest. The total error of the contour integration on the area of the images is taken as the sum of the errors of all arcs

$$E = \sum_{I,i,j} E_{I,i,j}. \quad (12)$$

Then, the absolute error on the magnification is simply

$$\delta\mu = \frac{E}{A_{\text{Source}}}. \quad (13)$$

As well known, contour integration works fine with uniform sources, but requires integration on concentric annuli to account for limb darkening. Let the source profile be given by

$$I(\rho) = \bar{I} f(\rho/\rho_*), \quad (14)$$

where \bar{I} is the average surface brightness and the profile function is normalized to 1:

$$2 \int_0^1 dr r f(r) = 1. \quad (15)$$

Let $F(r)$ be the cumulative profile at fractional radius r

$$F(r) = 2 \int_0^r dr' r' f(r'). \quad (16)$$

Then, the magnification of an annulus with inner radius r_{i-1} and outer radius r_i is approximated by

$$\tilde{M}_i = f_i [\mu_i r_i^2 - \mu_{i-1} r_{i-1}^2] \quad (17)$$

$$f_i = \frac{F(r_i) - F(r_{i-1})}{r_i^2 - r_{i-1}^2}, \quad (18)$$

where μ_i is the magnification of a uniform disc of radius $r_i \rho_*$.

The total magnification is

$$M = \sum_i \tilde{M}_i. \quad (19)$$

Also for limb darkening, we can introduce error estimators for each annulus and use optimal sampling to partition annuli only when needed. We start by the point source ($r = 0$) and the source boundary ($r = 1$). New radii are then inserted so as to divide the flux of the original annulus into two equal parts. In this way, the computational time is optimized in the best possible way.

The current version of VBBINARYLENSING uses a linear limb-darkening law

$$f(r) = \frac{1}{1-a/3} \left[1 - a \left(1 - \sqrt{1-r^2} \right) \right], \quad (20)$$

but a future extension to arbitrary profile will be straightforward.

3 ROOT SOLVING ROUTINE

The lens equation (2) can be recast into a fifth-order complex polynomial equation (Witt 1990), which can be solved by standard methods of complex algebra. In practice, all algorithms find one root at a time starting from an initial trial value. After one root is found, the polynomial is factored and the algorithm is repeated until all solutions are found.

Let $p(z)$ be our complex polynomial. If z_k is the current trial for the root to be found, we can optimize it by setting

$$z_{k+1} = z_k - \frac{1}{G}, \quad (21)$$

where

$$G = \frac{p'(z_k)}{p(z_k)}. \quad (22)$$

This is the so-called Newton algorithm, which works fine if we are close enough to the root.

A higher order algorithm was proposed by Laguerre:

$$z_{k+1} = z_k - \frac{n}{G \pm \sqrt{(n-1)(nH - G^2)}}, \quad (23)$$

where n is the degree of the polynomial, and

$$H = G^2 - \frac{p''(z_k)}{p(z_k)}. \quad (24)$$

The double sign is chosen so as to select the largest denominator and avoid singularities.

In the new version of VBBINARYLENSING, we have adopted the Skowron & Gould algorithm (Skowron & Gould 2012),⁵ which uses the Laguerre method when we are still very far from the root, so as to have a fast convergence and then switches to the Newton method, which is less demanding (it has no square roots and one derivative less), when we are close enough. The original FORTRAN code has been translated by us to C++ and incorporated in the VBBINARYLENSING library with adequate notice in the documentation. The C++ translation of the original library has also been posted on the original Skowron & Gould repository as a useful by-product of our development.

In Fig. 1, we quantify the speed-up obtained by using the Skowron & Gould algorithm with respect to a standard Laguerre method. On

⁵http://www.astrouw.edu.pl/~jskowron/cmplx_roots.sg/

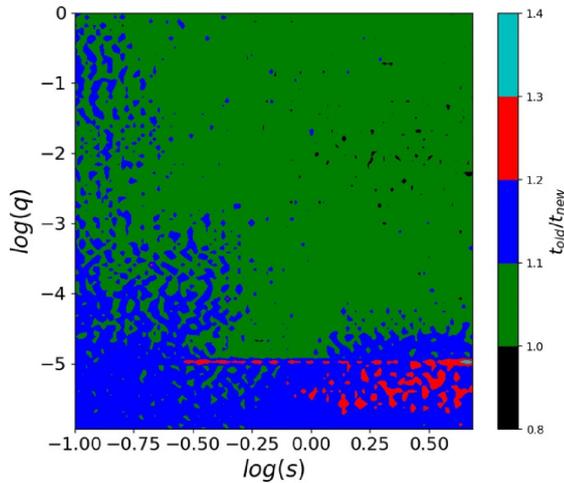


Figure 1. Speed-up obtained on a grid of possible lensing configurations (s, q) . For each point, we have averaged on source positions in the range $[-4, 4]$ for y_1 and $[0, 2]$ for y_2 .

average, it amounts to some 10 per cent. In addition, the Skowron & Gould prove to be more robust and fail-safe against testing at extreme mass ratios ($q \sim 10^{-9}$).

4 PRECISION VERSUS ACCURACY

Another upgrade of VBBINARYLENSING now implemented is the possibility for the user to specify either the absolute accuracy goal δ or a relative precision goal ϵ . Then, the calculation is stopped when either of the two thresholds is met. If $\delta\mu$ is the absolute error on the magnification as given by equation (13), the two criteria read

$$\delta\mu < \delta \quad (25)$$

$$\frac{\delta\mu}{\mu} < \epsilon. \quad (26)$$

In general, when modelling real data, we need the magnification calculation to be at least as accurate as the experimental uncertainties or better. In general, photometric observations come with statistical uncertainties of the order of 1 mmag or higher, with systematics of the same order of magnitude. In most cases, it will be sufficient to require $\epsilon = 10^{-3}$ and $\delta = 10^{-2}$ to match photometric errors.

With these numbers, the relative precision goal has no impact at low magnification ($\mu < 10$), but significantly speeds-up the computation at higher magnification ($\mu > 10$). Fig. 2 shows the speed-up obtained by the introduction of the relative precision goal as a function of magnification for a typical planetary lens configuration. Note that at a magnification $\mu \sim 100$, the speed-up is a factor of 3, while at magnification 1000, the speed-up may exceed 10. Saving computational time in the modelling of high-magnification events is very important, since these events are the most demanding in terms of computing time. A speed-up of few units may allow broader searches in the parameter space and more robust physical conclusions.

5 POINT-LIKE VERSUS FINITE SOURCE: A SIMPLE TEST TO DECIDE WHICH WAY TO TAKE

The main complication of microlensing computations indeed comes from the finiteness of physical sources. Point-like sources can be

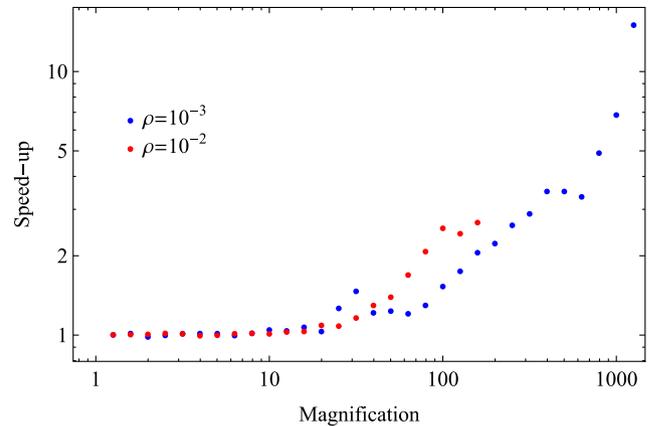


Figure 2. Speed-up obtained using a precision goal $\epsilon = 0.001$ with respect to an accuracy goal $\delta = 0.01$. The lens separation is $s = 1.2$ Einstein radii, the mass ratio is $q = 10^{-3}$. Red points are with source radius $\rho^* = 10^{-2}$, and blue points with $\rho^* = 10^{-3}$. Each point is an average over all entries in a grid over the source position falling within the same magnification bin.

easily treated by a single inversion of the lens equation. In that case, the total magnification is just given by the sum of the absolute values of the inverse Jacobian determinant evaluated at the images positions. No sampling of the boundary, no contour integration to calculate. The algorithm summarized here in Section 2 requires a minimum of eight points even when the source is very far from the caustics. With respect to the point-source calculation, the basic contour integration is therefore about 10 times slower.

While modelling a microlensing event, most of the survey observations fall close to the baseline. Even if some dense follow-up on the caustic crossings is present, the danger is that most of the computational time is spent where the event is least interesting, if we use the same contour integration algorithm everywhere.

However, choosing the right point where to switch from point to finite sources is far from straightforward and very tricky. In practice, modellers use to define an ‘anomaly window’ by visual inspection of the data, and instruct their codes to use point source outside this window and the full finite-source calculation within the anomaly window. Sometimes, some quadrupole or hexadecapole approximations are also employed in intermediate ranges around the anomaly window.

Apparently, it is impossible to transfer this approach to automatic modelling platforms or when we aim at analysing thousands of anomalous events at once. Furthermore, by a trivial extrapolation of the rate of anomalous events found in current ground-based surveys to the number of microlensing events expected by future space surveys such as *Wide-Field Infrared Survey Telescope (WFIRST)*, we must be ready to model thousands anomalous events simultaneously. The challenge posed by such data flows requires removal of any human intervention in the identification and modelling of interesting microlensing events. In this respect, the code itself must be able to decide whether to save time by the point-source approximation or to invest on the full finite-source calculation.

Let us state the terms of the problems more precisely. For any lens configuration, parametrized by the lens separation s normalized to the Einstein radius, and mass ratio q , we have to identify the regions of the source plane (y_1, y_2) in which we need to switch from the point- to the finite-source calculation. The boundary of these regions can be simply defined as the locus of points where the difference between the two calculation exactly matches the accuracy goal δ we

wish to reach. This requirement gives an operational definition of the two regimes, their respective regions of validity and their boundary. Finally, we expect that the respective extension of the two regions depend on the source size ρ_* : for smaller sources, we can safely use the point-source approximation up to small distances from the caustics, while for larger sources, we need to use the finite-source algorithm even far away.

The condition to accept the point-source result can then be formalized by the equation

$$|\mu_{\text{FS}} - \mu_{\text{PS}}| < \delta, \quad (27)$$

where μ_{PS} is the point-source magnification and μ_{FS} is the finite-source magnification for a given set of parameters (s , q , ρ_* , y_1 , y_2).

In principle, equation (27) also depends on the chosen limb-darkening profile. A limb-darkened source behaves in an intermediate way between a uniform-brightness source and a point-like source, since it downweights the limb and enhances the centre of the source disc. So, we generally have

$$|\mu_{\text{FS,ld}} - \mu_{\text{PS}}| \leq |\mu_{\text{FS,uni}} - \mu_{\text{PS}}|. \quad (28)$$

In order to keep the treatment (and the code) simple, rather than dealing with this complication, we prefer to tackle the problem referring to the worst possible situation, corresponding to the uniform disc source. As a consequence, for a limb-darkened source, we will just have some extra buffer on the safe side.

Our strategy is therefore to start with the calculation of the point-source magnification μ_{PS} and use the local information made available by this calculation to predict the discrepancy between the (uniform) finite- and point-source results. In the following subsections, we introduce three tests, respectively, based on the quadrupole expansion, the criticality of ghost images, and the distance from planetary caustics. These three tests complement each other so as to cover the whole sensible regions defined by the inequality (27) by their overlapping patches. These tests require very little additional computation to the pure point-source magnification but allow to save lots of computational time. Let us see them in detail.

5.1 Quadrupole test

It is fair to say that in any realistic calculation, we will work with an accuracy goal $\delta \ll 1$. In fact, in typical modelling problems, we would range from $\delta = 10^{-2}$ to 10^{-4} depending on the accuracy needed to minimize numerical error in the χ^2 calculation. This means that the switch from point to finite source at the boundary of the inequality (27) will occur in a perturbative regime, i.e. when the finite-source effect is just a small correction to the magnification. At this point, the first ‘failure’ of the point-source approximation will be due to higher orders in a multipole expansion in ρ_* exceeding δ . Such multipoles have been studied by several authors (Pejcha & Heyrovský 2009; Gould 2008; Cassan 2017). Following Pejcha & Heyrovský (2009), the finite-source magnification for a uniform-brightness source can be expressed as

$$\mu_{\text{FS}} = \mu_{\text{PS}} + \frac{\rho_*^2}{8} \Delta \mu_{\text{PS}} + \frac{\rho_*^4}{192} \Delta \Delta \mu_{\text{PS}} + o(\rho_*^4), \quad (29)$$

where $\Delta = \partial_{y_1} \partial_{y_1} + \partial_{y_2} \partial_{y_2}$ is the Laplacian operator. The second-order term is called the quadrupole and the fourth-order term is the hexadecapole. Gould (2008) suggests to calculate these Laplacians by sampling μ_{PS} on a collection of points on the source boundary and half-way from the boundary. Cassan (2017) proposes to calculate local derivatives at the source centre using the complex

notation, where such derivatives can be easily expressed. Since we just want to know whether the point-source approximation is sufficient, we need a test that does not involve further lens map inversions, otherwise this test would not save any time compared to the full finite-source calculation. Therefore, we follow an approach similar to Cassan (2017).

The lens equation in complex notation reads

$$\zeta = z + f(\bar{z}), \quad (30)$$

where $\zeta = y_1 + iy_2$ represents the source position in complex coordinates, $z = x_1 + ix_2$ is the generic point in the lens plane, the bar denotes the complex conjugation, and

$$f(z) \equiv -\frac{m_1}{z-s} - \frac{m_2}{z}. \quad (31)$$

The Jacobian determinant of the lens equation simply reads

$$J = 1 - f' \bar{f}'. \quad (32)$$

The point-source magnification at a source position ζ is obtained by summing the absolute values of the inverse Jacobian determinants calculated on each of the images z_I :

$$\mu_{\text{PS}} = \sum_I \frac{1}{|J(z_I)|}. \quad (33)$$

The first correction from finite source comes from the quadrupole term in equation (29). For each image position z_I , the quadrupole correction in complex coordinates reads

$$\mu_{Q_I} = -\frac{2\text{Re} [3\bar{f}'^3 f''^2 - (3 - 3J + J^2/2)|f''|^2 + J\bar{f}'^2 f''']}{J^5} \rho_*^2, \quad (34)$$

where both f and J are evaluated at each image position z_I .

In the computation of this quantity, J is already available from the point-source magnification calculation. The derivatives of $f(z)$ are easily calculated from its analytical expression, since it is just a rational function. The whole computation can be optimized in a few lines of code, which are good candidates to replace repeated calls to the root-solving routine, as required by the full finite-source calculation. So, as a first step, we shall definitely require that the full finite-source calculation is needed whenever the total quadrupole correction for all real images exceeds the accuracy goal δ . Formally, this condition reads

$$\sum_I c_1 |\mu_{Q_I}| < \delta. \quad (35)$$

The coefficient c_1 is to be chosen empirically so as to ensure enough safety margin around the dangerous regions. We will set this in the final tuning of our algorithm.

After some testing, we have soon realized that the condition (35) is not sufficient to cover all regions for which the point-source approximation fails. This could have easily expected considering that the quadrupole is only the second-order term in a power series. It may occasionally vanish while higher order corrections dominate. In particular, this happens around cusps: if we choose a specific approach angle, we may reach a cusp with zero quadrupole correction. However, we do not want to calculate arbitrary higher orders in the source size expansion, thus losing any advantage in our approach.

Cusps are defined by the vanishing of the tangential vector to the caustic. Following Petters & Witt (1996), this condition can be expressed in the complex formalism by

$$\frac{\partial \zeta}{\partial z} H_z + \frac{\partial \zeta}{\partial \bar{z}} \bar{H}_z = 0; \quad H_z = 2i \frac{\partial J}{\partial \bar{z}}, \quad (36)$$

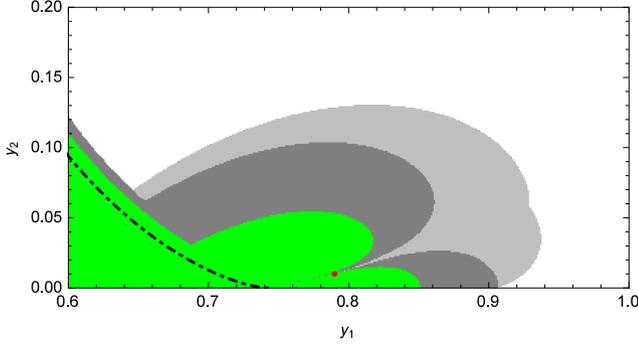


Figure 3. Example of a map around the right cusp for a configuration with $s = 1.35$, $q = 0.32$, $\rho^* = 10^{-2}$, and $\delta = 10^{-2}$. The caustic is the dotted–dashed line. In the green region, the finite-source correction is necessary, as $|\mu_{FS} - \mu_{PS}| > \delta$. This region is safely enclosed within the region covered by equation (35), shown in dark grey. However, this condition leaves one red point outside, in the thin inlet where the quadrupole correction vanishes. By adding the cusp correction according to equation (39), the safety margin is extended to the light grey region, filling the inlet, and solving the problem.

which for our binary lens without shear can be recast in the form

$$\text{Im} [\tilde{f}'^{3/2} f''] = 0. \quad (37)$$

The complex quantity appearing here is just the square root of the first term appearing in the quadrupole correction (34). Note that if equation (37) is fulfilled, then also the imaginary part of the squared quantity vanishes (the opposite is not true). Remembering that our purpose is to patch the quadrupole correction where it vanishes close to cusps, we look for the simplest function to add to equation (35) in order to achieve this goal. This function must contain the condition (37) so as to track the cusp position correctly. It turns out that a perfect job is already achieved by the minimal correction

$$\mu_C = \frac{6\text{Im} [3\tilde{f}'^3 f''^2]}{J^5} \rho_*^2. \quad (38)$$

Equation (35) is then updated to

$$\sum_l c_Q (|\mu_{Q,l}| + |\mu_{C,l}|) < \delta. \quad (39)$$

In Fig. 3, we show an example of this cusp problem solved by this correction. We see that the simple quadrupole condition (35) defines the dark grey region, which safely encloses most of the actual points in which the finite-source effect exceeds the accuracy goal (shown in green). However, one red point is left outside by this condition. Indeed, we can clearly see that the quadrupole correction vanishes along a straight line reaching the cusp by a fixed angle. The correction term we introduce in equation (39) extends the safety buffer to the light grey region, thus solving the problem.

5.2 Ghost images test

When the source approaches a cusp from the outside, one of the images is magnified well before the caustic crossing, thus triggering the quadrupole condition described before. However, when approaching a fold, none of the pre-existing images are magnified. In fact, a fold crossing shows up with a discontinuous slope change in the microlensing light curves, corresponding to the sudden creation of a new pair of images. The bump starts as soon as the limb of the source touches the fold. When this happens, there is no warning of the catastrophe from the other three images. Then, equation (39)

only intervenes for sources on the inner side (when all five images participate), but is useless on the outer side.

The fold crossing is characterized by the merger of the two ghost images, which become degenerate and then separate again as true solutions of the lens equation. Since these ghost images are made available by the root solving routine, we can try to ask how far they are from being degenerate.

The Jacobian determinant is a continuous function of z , whether we evaluate it on a real image or not. At the fold crossing, the two ghost images coincide $z_{G_1} = z_{G_2}$ and $J(z_{G_{1,2}}) = 0$. If our source centre is just outside the fold, the two ghost images of the source centre (which are available from the point-source calculation) will still be distinct but $J(z_{G_{1,2}})$ will be close to zero. A local expansion of J gives (we use z_G for either of the two ghost images to shorten notation)

$$J(z) = J(z_G) + \frac{\partial J}{\partial z} dz + \frac{\partial J}{\partial \bar{z}} d\bar{z}. \quad (40)$$

If we move from the source centre by a quantity $d\zeta = \rho^* e^{i\phi}$, we must find the corresponding shift in the ghost images. These satisfy the equation

$$\zeta = z_G + f(\hat{z}), \quad \hat{z} \equiv \bar{\zeta} - f(z_G). \quad (41)$$

Note that real images satisfy the same equation with the additional constraint that $\hat{z} = \bar{\zeta}$. For ghost images, this does not hold and we have to pair this equation with its conjugate

$$\bar{\zeta} = \bar{z}_G + f(\bar{\hat{z}}), \quad \bar{\hat{z}} \equiv \zeta - f(\bar{z}_G). \quad (42)$$

Differentiating the two equations, we obtain the following relation between $(d\zeta, d\bar{\zeta})$ and $(dz, d\bar{z})$

$$\begin{pmatrix} dz \\ d\bar{z} \end{pmatrix} = \begin{pmatrix} \frac{\hat{J}}{1-f'(\hat{z})f'(\bar{\hat{z}})} & \frac{\hat{J}}{1-f'(\hat{z})f'(\bar{\hat{z}})} f'(\hat{z}) \\ \frac{\bar{\hat{J}}}{1-f'(\bar{\hat{z}})f'(\hat{z})} & \frac{\bar{\hat{J}}}{1-f'(\bar{\hat{z}})f'(\hat{z})} \end{pmatrix} \begin{pmatrix} \rho_* e^{i\phi} \\ \rho_* e^{-i\phi} \end{pmatrix}, \quad (43)$$

where

$$\hat{J} = 1 - f'(z_G)f'(\hat{z}) \quad (44)$$

Note that this trivially collapses to the usual Jacobian matrix when $\hat{z} = \bar{z}_G$.

Now we can use this relation in equation (40) to express the change in the Jacobian function J evaluated at a ghost image when we move from the centre of the source to the limb. If $J(z_G)$ is close to zero, it is possible that $J(z)$ will vanish on some points on the limb for some values of $\phi \in [0, 2\pi]$. The maximum excursion of $J(z)$ from $J(z_G)$ will be obtained by setting the derivative with respect to ϕ to zero. This gives an equation for ϕ that can be solved to find the most dangerous point ϕ_M . After plugging this back in equation (40), we obtain an equation of the kind

$$J(z) = J(z_G) \pm \max(\Delta J|_{\rho_*}). \quad (45)$$

Both ϕ_M and $\max(\Delta J|_{\rho_*})$ can be calculated analytically, but we do not want to weigh the treatment down by lengthy expressions. In the end, if we want to test that we are far enough from a fold, we just want to make sure that $J(z)$ never vanishes on any points in the source limb, which requires

$$|J(z_G)| > |\max(\Delta J|_{\rho_*})|. \quad (46)$$

Putting the exact expressions in this equation, we get

$$\frac{1}{2} \left| J(z_G) \frac{\hat{J}^2}{\hat{J} f''(\bar{z}_G) f'(z_G) - \bar{\hat{J}} f''(z_G) f'(\bar{z}_G) f'(\hat{z})} \right| > c_G \rho_*. \quad (47)$$

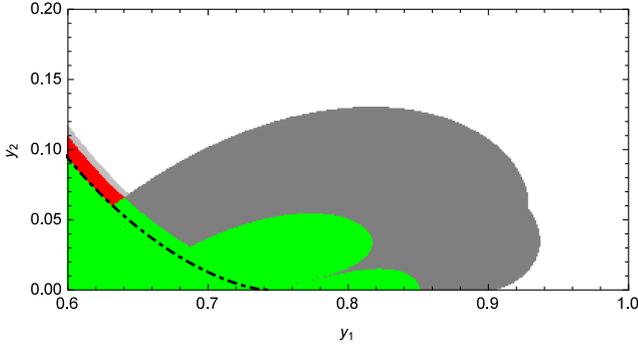


Figure 4. Same configuration as Fig. 3. Here, we have the zone covered by the quadrupole test in dark grey. This would leave out a red stripe along the fold where the limb of the source touches the fold. This stripe is safely enclosed by the ghost images test (47), reaching out to the light grey stripe..

The left-hand side represents the estimated distance on the source plane needed to send the ghost image to a critical point. If it is larger than the source size, we are sufficiently far from the fold, otherwise we have to switch to the finite-source calculation. The coefficient c_G is set empirically from test maps in order to leave enough safety buffer, as for c_Q in the case of the quadrupole test. Since we have two ghost images, both of them have to pass this test separately.

The implementation of this test requires the calculation of \hat{z} and the functions $f(\hat{z})$ and \hat{J} . With respect to the quadrupole test on the real images, we do not need f''' , so the two tests are still of comparable complexity. In Fig. 4 we can see an example of a region covered by the ghost images test: adjacent to a fold from the outside.

5.3 Planetary test

The ghost images test presented in the previous section works excellently for large binary caustics, but may fail for planetary caustics when the source is so large that its limb starts to touch the planetary caustic when its centre is still so far that the local gradient of the Jacobian is insensitive to the presence of the planet and is actually dominated by the primary component. This kind of failure, however, can be easily patched by checking that the distance from the centre of the source to the centre of the planetary caustic is greater than the source radius. In practice, the condition is

$$|\zeta - \zeta_{pc}|^2 > c_P(\rho_*^2 + \Delta_{pc}^2), \quad (48)$$

where c_P is the usual coefficient to be determined empirically by the maps. The position of the planetary caustic in our frame centred on the planet is (Griest & Safizadeh 1997)

$$\zeta_{pc} = -\frac{1}{s}, \quad (49)$$

and we have added a term Δ_{pc} in quadrature to the source radius to account for the extension of the caustic. The semi-extension of the caustic (Bozza 1999, 2000; Dominik 1999; Han 2006) is $2q^{1/2}/s^2$ for wide planets ($s > 1$) and $2q^{1/2}/s$ for close planets ($s < 1$, here we take the semidistance between the two triangular caustics). Then, we define

$$\Delta_{pc} = 3\frac{q^{1/2}}{s}, \quad (50)$$

which nicely tracks the extension of the planetary caustics in the close regime and is even too generous in the wide regime. The coefficient 3 allows enough buffer around the caustics. Fig. 5 illustrates

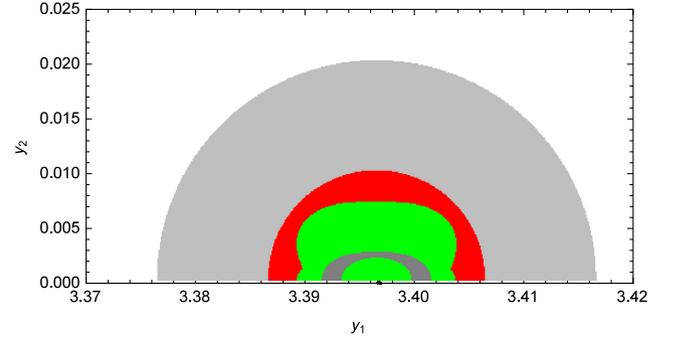


Figure 5. Map around the planetary caustic for the configuration $s = 3.67$, $q = 10^{-6}$, $\rho_* = 10^{-2}$, and $\delta = 10^{-2}$. The finite source is required in the green and red regions. The quadrupole and ghost images tests would trigger on the green region only, leaving out a considerable annulus around the caustic. Including the planetary test (48), we have enough buffer around the sensible region.

how the test succeeds for very small planetary caustics. We finally note that the planetary test is only needed for $q < 0.01$.

5.4 Broad testing on fine maps

Up to now, we have defined three test that only use local information at the centre of the source to cover all possible cases in which we need to switch to the full finite-source calculation. For each test, we can still play with the coefficients c_Q , c_G , c_P to decide how much buffer we want around the dangerous regions. Of course, it is recommendable to avoid switching at the very last time from point to finite sources, because this would generate a discontinuity in the output at the accuracy level δ . Nevertheless, we definitely want to contain the waste of time in safe areas of the source plane. The compromise proposed in our code is

$$c_Q = 6 \quad (51)$$

$$c_G = 2 \quad (52)$$

$$c_P = 2. \quad (53)$$

As can be imagined, the three conditions presented here are the output of an extensive testing with fine maps of the source plane for a wide range of parameters. We have taken $q \in [10^{-6}, 1]$, $s \in [0.1, 4]$, $\rho_* \in [10^{-4}, 10^{-1}]$, $\delta \in [10^{-4}, 10^{-2}]$, and $y_1 \in [-3, 4]$, $y_2 \in [0, 2]$ in steps of 10^{-2} or 10^{-3} for $q < 0.1$ (planetary case). These maps have allowed us to identify all possible cases and tune the coefficients so as to achieve the best performance and robustness. Since smaller sources might require an even finer grid for testing, which would be unpractical, we have replaced each ρ_* appearing in conditions (39) and (47) by $\rho_* + 10^{-3}$. In practice, we do not shrink the buffer below that of a source with radius 10^{-3} . Indeed such buffer is already so tiny that it does not make too much sense to decrease it further.

To make an example of the test maps we have built, in Fig. 6, we show a table of maps obtained for the same lens configuration with various sizes of the source and accuracy goals. We note that at lower accuracies and smaller sources, the point-source approximation works excellently down to the caustic, while larger sources and finer accuracies require an earlier switch from point to finite sources. In any case, the three conditions defined in this section perfectly enclose any dangerous regions.

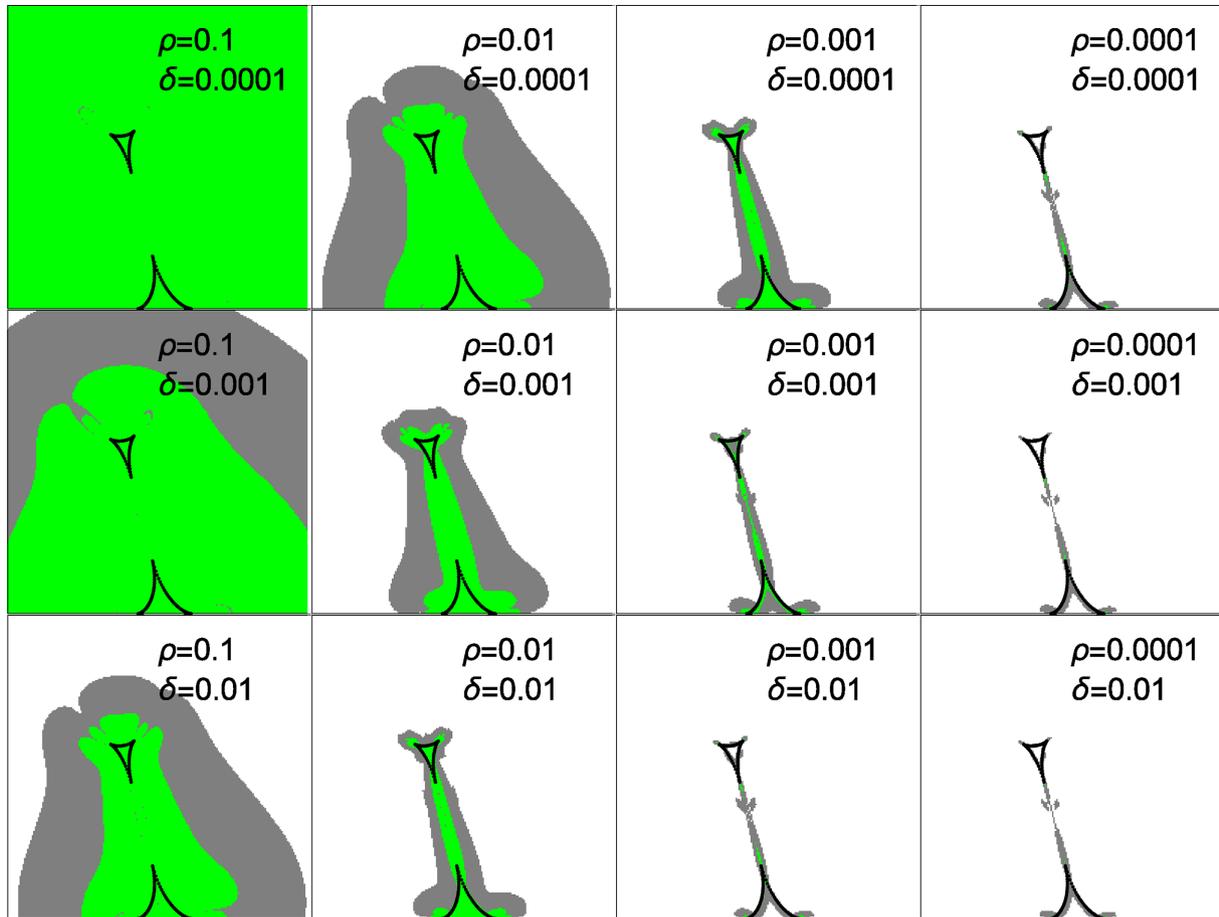


Figure 6. Maps built for $s = 0.67$ and $q = 0.56$ with the indicated values of the source radius ρ and the accuracy goal δ . Caustics are shown in black. The green regions are those in which the difference between the finite- and point-source magnification exceeds δ . These regions are well enclosed within the grey regions, defined by the three conditions described in this section.

As stated before, our investigation has been made for a uniform-brightness source, which generates the maximal discrepancy from the point source. A limb-darkened source would be intermediate between the two cases. Therefore, maps built with a limb-darkened source would have smaller green regions. Limb-darkened sources are thus automatically enclosed in our finite-source regions defined on the uniform-source case and pose no problems for our code.

Finally, an important information for the user is the parameter range in which the code is robust and can be used safely. We have already declared the ranges spanned in the detailed maps used for the definition of our finite-source condition. However, for larger source distances $|y| \gg 1$, the magnification falls to one and there are no expected problems. For the separation between the two lenses, we typically work from 0.1 to 10 without known problems. Outside this range, it is difficult to find any signatures of binarity, so we have not tested further. As for the mass ratio, we have tested down to $q \sim 10^{-9}$ finding very occasional glitches with large sources. Indeed, the source size ρ^* should be kept below 0.1 to stay on the safe side.

6 FINITE-SOURCE EFFECT ON SINGLE LENS

We have extensively discussed the upgrades on the code for the binary lensing magnification, since this represents the highlight of the `VBBINARYLENSING` library. However, the library also offers the possibility to calculate the single-lens magnification in several

flavours: point source, finite source, and finite source with limb darkening.

Traditionally, the finite-source effect for uniform-brightness discs is calculated using elliptic integrals (Witt & Mao 1994). The same approach can be extended to arbitrary limb-darkening profiles provided the source is small enough compared to the Einstein angle (Yoo et al. 2004).

After all, the calculation of elliptic integrals is not for free and requires several operations. Since the finite-source effect basically depends on the two parameters u and ρ (leaving apart the limb darkening for the moment), we have explored the possibility of pre-calculating the finite-source corrections for a wide enough range of these two parameters. These pre-calculated corrections have been stored in a binary table that is loaded at the first call of the finite-source-point-lens magnification function. We have then realized that reading a pre-calculated correction to the Paczynski magnification in a table is extremely faster than going through all elliptic integrals. Of course, the table must be fine enough to track the corrections with a simple linear interpolation between two consecutive entries in the table. We have chosen an accuracy goal of 10^{-3} . Finally, to account for limb darkening, we have adopted the same algorithm as for the binary lens magnification, i.e. we calculate the magnification on multiple annuli until we match the required accuracy.

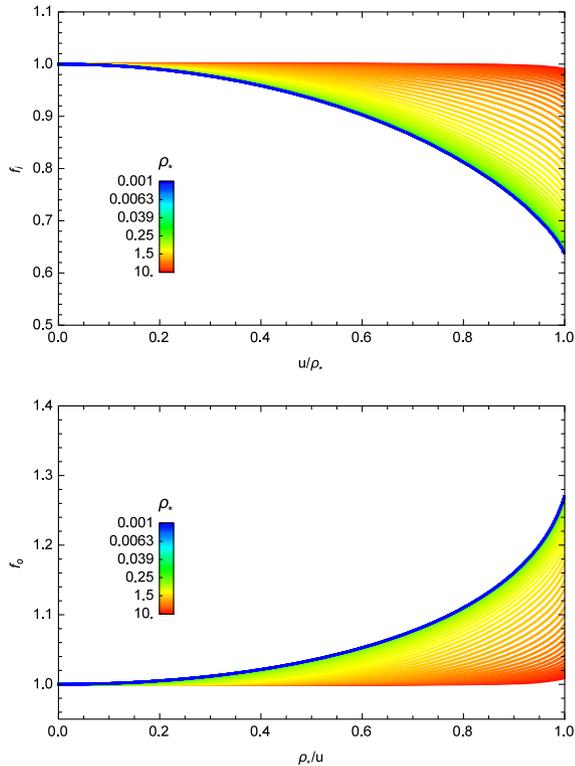


Figure 7. Finite-source correction as a function of u/ρ^* for $u < \rho^*$ (top) and $u > \rho^*$ (bottom). 101 curves are shown for source sizes in the range $[10^{-4}, 10^{+1}]$ with a logarithmic step. The larger the source, the larger f_o and the smaller f_i .

More in detail, the uniform-source magnification is

$$\mu(u, \rho_*) = \begin{cases} \frac{u^2+2}{u\sqrt{u^2+4}} f_o(\rho_*/u, \rho_*) & u > \rho_* \\ \sqrt{1 + \frac{4}{\rho_*^2} f_i(u\rho_*, \rho_*)} & u < \rho_* \end{cases} \quad (54)$$

The functions f_o and f_i can be read from Witt & Mao (1994) in terms of elliptic integrals. We show them in Fig. 7 for a variety of source radii as functions of the ratio u/ρ_* or its inverse.

Another considerable advantage in using pre-calculated tables with respect to the approach by Yoo et al. (2004) is that we are not limited to small source radii $\rho^* \ll 1$. Our approach allows to investigate sources larger than the Einstein radius without any limitations. This is particularly interesting for the study of free-floating planets, which are characterized by small Einstein radii and often fall in this case.

7 LIGHT-CURVE CALCULATIONS

The basic functions included in VBBINARYLENSING allow to calculate the binary and single-lens magnifications for specific lens and source configurations. However, it is also useful to have functions that calculate the magnification along an entire microlensing light curve as a function of typical fit parameters. For example, static binary lens models depend on the lens separation, mass ratio, and source radius, while the source position is expressed in terms of some geometric parameters (closest approach distance, time of closest approach, trajectory angle, and Einstein time).

Our code offers the possibility to calculate a single point on the light curve at a given time for any set of parameters. Alternatively, one can submit a full array of observation epochs and get back an

array of magnifications, thus calculating the whole light curve with one single call.

7.1 Parallax, satellites, and orbital motion

To a first extent, a microlensing event is short enough to approximate the relative angular motion of lens and source as rectilinear. Furthermore, the two lenses are taken as fixed. However, we know that the Earth revolves around the Sun in one year and that the two components of a binary system must orbit around their common centre of mass. In general, the time-scales of these two motions are much higher than the time-scale of the microlensing event, so that a ‘static’ model is sufficient to describe most binary events. However, the interest in such higher order effects is rapidly increasing thanks to the availability of more accurate photometry and sampling, allowing to track long-term effects. In addition, observations from space open the possibility to study the geometric parallax systematically for all events, granting access to this precious information for the near totality of events.

In our code, it is possible to include parallax in the set of parameters. The Earth orbit is calculated using the JPL approximate ephemeris formulae,⁶ which is simple and largely sufficient for microlensing parallax purposes. The user can choose between the north-east reference frame to express the components of the parallax vector and the parallel-perpendicular frame, where the reference vector in this case is the projection of the Earth-acceleration on the sky. Furthermore, it is possible to specify the reference time $t_{0, \text{par}}$, which can coincide with the closest approach time t_0 or can be separately specified. Then, all deviations from the straight motion of the Earth are calculated with respect to the chosen $t_{0, \text{par}}$ (see e.g. Skowron et al. 2011). Finally, if one wants to calculate the light curve as observed by a satellite, it is possible to use a position table of the spacecraft in the Jet Propulsion Laboratory (JPL) format.

The orbital motion of a binary lens may become important if the caustics are relatively large (as for intermediate binaries). In this case, small changes in the separation of the lenses may translate into a sensible change in the shape of the caustics, while the source is still crossing them or passing close enough for the magnification to be affected. In most studies, the linear approximation is adopted, with the addition of the two parameters ds/dt (tracking the change in the separation of the lenses) and $d\alpha/dt$ (tracking the rotation of the system) (Albrow et al. 2000). This approximation becomes unphysical if followed too far from the reference time $t_{0, \text{orb}}$ (typically equal to $t_{0, \text{par}}$). Therefore, we prefer to propose a circular orbit for the lens as a built-in function in the library, with velocity parameters $(\gamma_{\parallel}, \gamma_{\perp}, \gamma_z)$ as described in Skowron et al. (2011). Even if the radial velocity γ_z would be poorly constrained in many cases, in our opinion, it is better to stay on the safe side. The inclusion of the option for a full Keplerian orbit should come with future versions of the code without too much trouble.

8 CONCLUSIONS

With large-scale ground-based surveys detecting more and more microlensing events (Yee et al. 2018), the incoming LSST widening the microlensing window to the whole sky (Marshall et al. 2018), and the planned *WFIRST* mission promising thousands of planetary events (Bennett et al. 2018), the need for fast data processing is

⁶https://ssd.jpl.nasa.gov/txt/aprx_pos_planets.pdf

becoming more and more urgent. The distinction of genuine microlensing planets from false positives has to be made as automatic as possible, minimizing the human intervention on exceptionally puzzling cases. To this purpose, a few publicly available modelling platforms for microlensing are currently being developed: `pyLIMA`⁷ (Bachelet et al. 2017), `MuLensModel`⁸ (Poleski & Yee 2018), and `muLAn`.⁹ Although the goals of these platforms differ somewhat, they all have in common the embedding of `VBBINARYLENSING` as the basic tool for calculations of binary lens magnifications. The continuous development, maintenance, and upgrade of our software is therefore of fundamental importance of several higher level software tools that aim at upholstering microlensing with more appealing and friendly interfaces.

In the new version presented in this paper, we have introduced several novelties for speeding-up the light-curve calculation in all regimes. The Skowron & Gould algorithm for the basic root solving routine provides a modest speed-up, but more robustness. The possibility to specify a precision goal dramatically speeds-up high-magnification calculations. Three new tests now drive the software toward the more convenient branch (point or finite sources) in any configurations, thus obtaining another dramatic speed-up far from the caustics. We have adopted pre-calculated tables for the finite-source-single-lens case, which are much faster and have no practical limitations. We have finally described how full light curves for any kind of models can be obtained. For each upgrade, we have conducted extensive testing on the whole parameter space, as documented by the figures included in this paper.

Apart from the already mentioned somewhat straightforward extensions to Keplerian orbital motion and arbitrary limb-darkening profiles, the new challenges that need to be addressed in view of *WFIRST* are multiple lenses and astrometry. In this respect, we hope to come back with new versions of our code in the next future.

ACKNOWLEDGEMENTS

`VBBINARYLENSING` has benefitted from precious feedback from the community. In particular, we wish to thank Sebastiano Calchi Novati, Giuseppe D'Agò, Radek Poleski, Sedighe Sajadian, Rachel Street, and Wei Zhu for their testing and suggestions.

REFERENCES

- Akeson R. L. et al., 2013, *PASP*, 125, 989
 Albrow M. et al., 2000, *ApJ*, 534, 894
 Alibert Y. et al., 2010, *Astrobiology*, 10, 19
 Bachelet E., Norbury M., Bozza V., Street R., 2017, *AJ*, 154, 203
 Bennett D., 2010, *ApJ*, 716, 1408
 Bennett D., 2018, preprint (arXiv:1803.08564)
 Bennett D., Rhie S. H., 1996, *ApJ*, 472, 660
 Bozza V., 1999, *A&A*, 348, 311
 Bozza V., 2000, *A&A*, 355, 423
 Bozza V., 2010, *MNRAS*, 408, 2188
 Cassan A., 2017, *MNRAS*, 468, 3993
 Christiansen J. L., 2018, in Deeg H. J., Belmonte J. A., eds., *Handbook of Exoplanets*, preprint (arXiv:1803.11158)
 Coughlin J. L. et al., 2016, *ApJS*, 224, 12
 Dominik M., 1995, *A&AS*, 109, 507
 Dominik M., 1998, *A&A*, 333, L79
 Dominik M., 1999, *A&A*, 349, 108

⁷<https://github.com/ebachelet/pyLIMA>

⁸<https://github.com/rpoleski/MuLensModel>

⁹<https://github.com/muLAn-project/muLAn>

- Fischer D. A., Howard A. W., Laughlin G. P., Macintosh B., Mahadevan S., Sahlmann J., Yee J. C., 2014, in Beuther H., Klessen R. S., Dullemond C. P., Henning T., eds, *Protostars and Planets VI*, Univ. Arizona Press, Tucson, AZ, p. 715
 Gould A., 2008, *ApJ*, 681, 1593
 Gould A., 2016, in Bozza V., Mancini L., Sozzetti A., eds, *Methods of Detecting Exoplanets*, Springer, Berlin, p. 135
 Gould A., Gauchere C., 1997, *ApJ*, 477, 580
 Griest K., 1991, *ApJ*, 366, 412
 Griest K., Safizadeh N., 1997, *ApJ*, 500, 37
 Han C., 2006, *ApJ*, 638, 1080
 Kayser R., Refsdal S., Stabell R., 1986, *A&A*, 166, 36
 Marshall P., 2018, preprint (arXiv:1708.04058)
 Nemiroff R. J., Wickramasinghe W. A. D. T., 1994, *ApJ*, 424, 21
 Pejcha O., Heyrovský D., 2009, *ApJ*, 690, 1772
 Petters A. O., Witt H. J., 1996, *Jour. Math. Phys.*, 37, 2920
 Poleski R., Yee J., 2018, *Astron. Comput.*, preprint (arXiv:1803.01003)
 Schneider J., Dedieu C., Le Sidaner P., Savalle R., Zolotukhin I., 2011, *A&A*, 532, 79
 Schramm T., Kayser R., 1987, *A&A*, 174, 361
 Skowron J. et al., 2011, *ApJ*, 738, 87
 Skowron J. et al., 2018, *Acta Astron.*, 68, 43
 Skowron J., Gould A., 2012, preprint (arXiv:1203.1034)
 Wambsgans J., 1990, PhD thesis, Ludwig-Maximilians-University Munich
 Witt H. J., 1990, *A&A*, 236, 311
 Witt H. J., Mao S., 1994, *ApJ*, 430, 505
 Yee J., 2018, preprint (arXiv:1803.07921)
 Yoo J. et al., 2004, *ApJ*, 603, 139

APPENDIX A: THE VBBINARYLENSINGLIBRARY CLASS

`VBBINARYLENSING` comes with the classical source file pair `VBBinaryLens.h` containing the declarations and `VBBinaryLens.cpp` containing the corresponding implementations. A file named `instructions.cpp` contains a sample main function with documentation of all features.

In short, the `VBBINARYLENSINGLIBRARY` class contains several functions that perform the microlensing calculations and some variables that can be set by the user to control the optional settings. Here is a list with a short description for each item.

Magnification calculations:

- (i) `BinaryMag0`: magnification of a point source by a binary lens. There is also a version returning the positions of the images.
- (ii) `BinaryMag`: magnification of a uniform-brightness source by a binary lens. There is also a version returning the positions of the images.
- (iii) `BinaryMagDark`: magnification of a limb-darkened source by a binary lens.
- (iv) `BinaryMagMultiDark`: magnifications of a limb-darkened source by a binary lens in different filters with different limb-darkening coefficients. Useful for simultaneous multiband observations.
- (v) `BinaryMag2`: magnification of a generic source by a binary lens. This function (new in v2.0) implements the tests described in Section 5, thus deciding whether to go for point- or finite-source calculation.
- (vi) `ESPLMag`: magnification of a uniform-brightness source by a single lens. This uses the pre-calculated table described in Section 6.
- (vii) `LoadESPLTable`: loads the pre-calculated table for `ESPLMag` magnification. It has to be called before the first use of `ESPLMag` or related functions.

- (viii) `ESPLMagDark`: magnification of a limb-darkened source by a single lens.
- (ix) `ESPLMag2`: magnification of a generic source by a single lens. Decides whether to go for point- or finite-source calculation based on the tests in Section 5.
- (x) `ESPLMagDark`: magnification of a limb-darkened source by a single lens.

Light-curve calculations:

- (i) `PSPLLightCurve`: PSPL light curve for a given set of parameters. This and all light curve functions are available for a single epoch or for a full array of observation epochs.
- (ii) `PSPLLightCurveParallax`: PSPL light curve including parallax.
- (iii) `SetObjectCoordinates`: Sets the astronomical coordinates of the microlensing target and specifies the path where to look for the position tables of the satellites (if any).
- (iv) `ESPLLightCurve`: Extended-Source-Point-lens light curve. This uses the `ESPLMag2` function.
- (v) `ESPLLightCurveParallax`: Extended-Source-Point-lens light curve with parallax.
- (vi) `BinaryLightCurve`: static binary lens light curve for a given set of parameters. This uses the `BinaryMag2` function.
- (vii) `BinaryLightCurveW`: static binary lens light curve for a given set of parameters using the centre of the caustic of the lens on the right as a reference point for the trajectory.
- (viii) `BinaryLightCurveParallax`: binary lens light curve including parallax for a given set of parameters.

- (ix) `BinaryLightCurveOrbital`: binary lens light curve including parallax and circular orbital motion for a given set of parameters.
- (x) `BinSourceLightCurve`: light curve for a single lens and a binary source. Sources are treated as point-like.
- (xi) `BinSourceLightCurveParallax`: light curve for a single lens and a binary source including parallax.
- (xii) `BinSourceLightCurveXallarap`: light curve for a single lens and a binary source including parallax and circular orbital motion.

Settings:

- (i) `To1`: absolute accuracy goal (called δ in this paper).
- (ii) `RelTo1`: relative precision goal (called ϵ in this paper).
- (iii) `a1`: linear limb-darkening coefficient used in `BinaryMag2` and `ESPLMag2`.
- (iv) `minannuli`: minimum number of annuli to calculate for limb-darkening.
- (v) `parallaxsystem`: 0 for parallel perpendicular, and 1 for north-east.
- (vi) `t0_par_fixed`: set to 1 if you want to specify a constant $t_{0, \text{par}}$.
- (vii) `t0_par`: reference time for parallax $t_{0, \text{par}}$. Only used if `t0_par_fixed=1`.
- (viii) `satellite`: specifies the satellite number for the next calculation (0 for observations from the ground).

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.