

SPEEDING UP LOW-MASS PLANETARY MICROLENSING SIMULATIONS AND MODELING: THE CAUSTIC REGION OF INFLUENCE

MATTHEW T. PENNY

Department of Astronomy, Ohio State University, 140 West 18th Avenue, Columbus, OH 43210, USA; penny@astronomy.ohio-state.edu

Received 2013 November 5; accepted 2014 June 20; published 2014 July 16

ABSTRACT

Extensive simulations of planetary microlensing are necessary both before and after a survey is conducted: before to design and optimize the survey and after to understand its detection efficiency. The major bottleneck in such computations is the computation of light curves. However, for low-mass planets, most of these computations are wasteful, as most light curves do not contain detectable planetary signatures. In this paper, I develop a parameterization of the binary microlens that is conducive to avoiding light curve computations. I empirically find analytic expressions describing the limits of the parameter space that contain the vast majority of low-mass planet detections. Through a large-scale simulation, I measure the (in)completeness of the parameterization and the speed-up it is possible to achieve. For Earth-mass planets in a wide range of orbits, it is possible to speed up simulations by a factor of ~ 30 – 125 (depending on the survey’s annual duty-cycle) at the cost of missing $\sim 1\%$ of detections (which is actually a smaller loss than for the arbitrary parameter limits typically applied in microlensing simulations). The benefits of the parameterization probably outweigh the costs for planets below $100 M_{\oplus}$. For planets at the sensitivity limit of AFTA-WFIRST, simulation speed-ups of a factor ~ 1000 or more are possible.

Key words: gravitational lensing: micro – methods: miscellaneous – planetary systems – space vehicles

Online-only material: color figures

1. INTRODUCTION

Gravitational microlensing searches for planets are beginning to yield statistically interesting sample sizes (Gould et al. 2010; Cassan et al. 2012) that are set to increase significantly with the advent of new and proposed surveys on the ground (MOA-II, Sumi 2010; OGLE-IV, Udalski 2011; KMTNet, Kim et al. 2010), and in space (Euclid, Penny et al. 2013; WFIRST, Spergel et al. 2013). Full understanding of the results of these surveys and the planning of optimal observing strategies requires extensive computationally expensive simulations for the calculation of detection efficiencies or yield predictions.

The major bottleneck in the simulation of planetary microlensing is the computation of light curves. To compute the magnification of the binary microlensing event with a finite source for a single data point requires either the solution of multiple complex fifth-order polynomials or costly inverse ray shooting. Each light curve typically consists of thousands of data points, and each simulation typically requires thousands to millions of trial light curves to be generated in order to obtain reasonable Poisson uncertainties due to the low per-event probability of planet detection.

Significant effort has been invested in increasing the speed at which one can compute the base unit of the light curve – the binary lens magnification. Various approaches have been taken to compute finite source magnifications, either by contour integration combined with numerical solution of the lens equation (Gould & Gaucherel 1997; Dominik 1998; Bozza 2010), inverse ray shooting (Rattenbury et al. 2002; Dong et al. 2006; Bennett 2010), or a hybrid of the two (Dong et al. 2006; Dominik 2007). Others have dug deeper and improved the efficiency of the basic numerical functions these are built upon (Skowron & Gould 2012). Further gains are made by avoiding finite source calculations where approximations will suffice, e.g., the hexadecapole approximation (Pejcha & Heyrovský 2009; Gould 2008).

Perhaps more important in reducing the computational expense of microlensing calculations is the choice of an efficient parameterization. The majority of work on this front has focused on the problem of fitting an observed microlensing event with a binary lens model. In this situation, most of the parameter space is uninteresting—one would prefer to avoid a brute force exploration of parameter space and go straight to the desired answer. This is typically achieved by using a parameterization that matches the features of the event being studied in order to reduce correlations between parameters and decrease the parameter-space volume that must be searched (e.g., Albrow et al. 1999; An et al. 2002; Cassan 2008; Bennett et al. 2012). Simulations to determine detection efficiencies or predict yields have the opposite goal of exploring the entire parameter space. The specialized parameterizations developed for modeling observed events are often difficult to apply to such brute force parameter explorations because the prior distributions of specialized parameters (e.g., caustic crossing durations) are not easily mapped from the more fundamental parameters (e.g., Einstein crossing timescale, mass ratio, and projected separation; Cassan et al. 2010).

Despite the desire to search the entire parameter space, it is known that most of it contains uninteresting microlensing events that do not show signs of the planets that orbit the lens. If it is possible to identify this region of parameter space before calculating the light curve, it is possible to avoid costly light curve calculations and significantly speed-up planetary microlensing simulations. This paper introduces a new parameterization (the Caustic Region Of Influence, or CROIN) that makes it possible to easily identify uninteresting regions of parameter space, while keeping a set of geometric parameters with uniform prior distributions. This allows remarkable efficiency gains in the simulation of low-mass planetary microlensing, up to factors of a thousand in reduced computation time. The structure of the paper is as follows. Section 2 describes the standard binary microlensing parameterization, the new CROIN parameterization,

and how to transformation between them. In Section 3, we estimate the speed-up that the CROIN parameterization enables, before assessing its accuracy and any biases it introduces in Section 4. In Section 5, we discuss first the region of parameter space in which the CROIN parameterization is useful, before discussing its possible applications. We conclude in Section 6.

2. THE METHOD

Before describing the CROIN parameterization, we first review the standard parameterization and the three caustic topologies that are possible for a binary lens. We conclude the section by describing how to convert between the two parameterizations.

2.1. The Standard Parameterization

The light curve of a microlensing event is determined by the magnification pattern of the lens and the trajectory of the source passing through it. The magnification pattern of a binary lens is determined by just two parameters: s , the projected separation of the lens components in units of the Einstein ring radius r_E , and $q = M/M_*$, the mass ratio of the components, where M is the mass of the planet and M_* is the mass of the primary. The source trajectory is described by three parameters that have uniform distributions: t_0 , the time of closest approach of the source to a reference point; u_0 , the impact parameter at closest approach (normalized to the Einstein radius); and α , the angle made by the source trajectory relative to the binary axis, which points from the primary to the planet. Finally, the Einstein crossing timescale t_E , is the time taken for the source to travel one angular Einstein radius θ_E . In this work we choose the reference point as the position of the primary lens, though there are a proliferation of preferred reference points, each with their own advantages and disadvantages.

2.1.1. The Range of u_0 and t_0

When simulating microlensing events, one is forced to choose a range for the parameters u_0 and t_0 from which to draw uniformly, which is usually a compromise between efficiency and completeness. The choice for t_0 is relatively straightforward, and is usually the range of time over which data will be taken, possibly with some margins outside this range to capture planetary signals in the wings of events just outside the observing window. The choice for the maximum value of $|u_0|$, $u_{0,\max}$ is more complicated. The probability of detecting a planet falls as u_0 increases, but is finite for $u_0 \lesssim |s - 1/s|$ (where $s - 1/s$ is the approximate position of the planetary caustics; Han 2006), which, given projection effects, can become large, meaning that one is forced to compromise between completeness and computation time. A choice of $u_{0,\max} = 3$ is common for planetary microlensing simulations, but inevitably some fraction of planet detections are lost (in Section 4 we find $\sim 5\%$ for $u_{0,\max} = 3$ and $\sim 20\%$ for $u_{0,\max} = 1$).

2.2. The Caustic Region Of Influence (CROIN) Parameterization

The deviations of the binary lens magnification pattern from that of a single lens are strongest at the caustic curves (Schneider & Weiss 1986) and fall off rapidly outside the caustic (e.g., Gaudi & Petters 2002a, 2002b). The caustics are therefore a natural choice for the reference point of an improved impact-parameter-based parameterization.




Figure 1. Top panel: the CROIN impact parameter u_c plotted against projected separation s for a large number of simulated planetary microlensing events. Gray points are a sample (~ 1 of all simulated events), while colored points are events where the planet was detected, with colors coded to different ranges of mass ratio q (because of the smaller number of detections with $\log q \leq -6$, we use a larger point size to emphasize them). Note that except for the “goat-horn” features, most detections lie below the gray non-detections. The colored lines show our adopted limits for the size of the CROIN $u_{c,\max}$, with the line drawn at the maximum of the q -range corresponding to its color, i.e., the black line is $u_{c,\max}(s, q) = 10^{-3}$. Bottom panel: the same data, but now for each point u_c is normalized by $u_{c,\max}$.

(A color version of this figure is available in the online journal.)

In the CROIN parameterization, we define the time of closest approach t_c and impact parameter u_c with respect to a reference point (x_c, y_c) centered on the planetary caustic(s). With the number, position, and size of caustics differing for each lens topology, we must therefore choose a different reference point and impact parameter range $u_{c,\max}$ for each topology. We treat each topology separately below. Because we will use the center of the planet signature as our reference point, we assume the range of t_c will be the range of time over which data will be taken; this assumption relies on the planetary signature being short compared to the season length, which we show is justified in Section 4.

2.2.1. Defining $u_{c,\max}$

In order to define expressions for $u_{c,\max}$, we employed a plot of u_c against the projected separation s , as shown in Figure 1. For this we used a simulation of WFIRST-AFTA from (Spergel et al. 2013) covering a large, but non-uniform range of planet mass ratio.¹ Once the reference point is defined

¹ Specifically, the planet semimajor axis was drawn from a log-uniform distribution and the planet mass was drawn from the Cassan et al. (2012) mass function that is forced to saturate at $5.5 M_{\oplus}$ (2 dex^{-2} planets per star) as described by Penny et al. (2013). This mass function, coincidentally, roughly cancels out the microlensing detection efficiency (until it saturates) and so is useful for producing a roughly uniform number of planet detections as a function of mass. This choice of mass function here has no impact on the results.

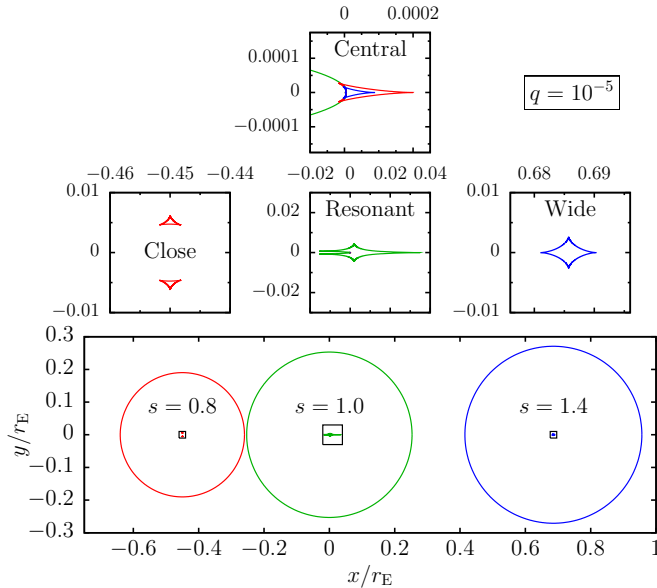


Figure 2. Plots showing the caustics for lenses with close (red), resonant (green), and wide (blue) topologies for a lens with mass ratio $q = 10^{-5}$. In each case the primary lens is located at $(x, y) = (0, 0)$ and the planet is located at $(+s, 0)$. The middle panels zoom in on the boxes shown in the lower panel. The central caustics of the close and wide topology lenses are plotted, but are too small to be resolved, in the middle panel, so the top panel zooms in on the middle panel by more than a factor of 170 to reveal their structure. The circles in the lower panel show the size of the CROIN for each example caustic topology.

(A color version of this figure is available in the online journal.)

(see the following subsections), the standard (u_0, t_0) parameters are easily transformed to (u_c, t_c) as described in Section 2.3. The plot is shown in the top panel of Figure 1, with colored points marking planet detections of differing mass ratio and light gray points marking a small fraction of the simulation’s non-detections. Here and throughout the paper we consider a planet to be detected if it causes a $\Delta\chi^2 > 160$ between the fit of a single point lens light curve model and the true light curve, which is standard for simulations of space-based microlensing surveys (see Yee et al. (2013) for a discussion).

The goal of plotting u_c against s is to identify a functional form for $u_{c, \max}$ as a function of s and q above which there are very few detections. This is very similar to the process one follows when using the rejection method for drawing random deviates from a general probability distribution function (Press et al. 1986), except that here we do not need to guarantee that the “comparison” function is always greater than the probability density function. The first attempt at defining a functional form used analytic approximations of the size of the caustics scaled by a constant factor to include planet detections via non-caustic-crossing trajectories, but we found that these did not fit the upper envelope of detections very well, regardless of the scaling factor. The final form for $u_{c, \max}$ was found by modifying these expressions with further analytic functions of s and q for each topology. Several iterations through candidate functional forms and their numerical parameters were made before arriving at those presented below. The following three subsections present the final forms for the reference points and $u_{c, \max}$ for each caustic topology in turn and justify the choices made. Figure 2 shows examples of the CROIN for a planet with $q = 10^{-5}$ relative to the caustics for each topology.

2.2.2. Close Topology

The lens has a close topology when the condition

$$\frac{q}{(1+q)^2} < s^{-8} \left(\frac{1-s^4}{3} \right)^3 \quad (1)$$

is satisfied (Erdl & Schneider 1993). The close topology lens has three caustics: a central caustic near the primary and two identical planetary caustics off the binary axis. The most prominent feature of close lenses with low-mass ratios is an elongated region of demagnification on the binary axis between the two planetary caustics. The positions of the planetary caustics can be very accurately approximated analytically (Bozza 2000b; Han 2006), therefore we choose the reference point for the close topology to be

$$(x_c^c, y_c^c) = \left(\frac{1}{1+q} \left[s - \frac{1-q}{s} \right], 0 \right), \quad (2)$$

where we have used Bozza’s slightly more accurate expression for the x position of the caustics and the superscript “c” represents the close topology. Similarly, the superscripts “w” and “r” will represent wide and resonant topologies, respectively.

The size of the CROIN that we choose for the close topology is

$$u_{c, \max}^c = \begin{cases} 0 & \text{if } s < 0.1 \\ (4 + 90s^2) \frac{\sqrt{q}}{s\sqrt{1+s^2}} & \text{otherwise} \end{cases} \quad (3)$$

For very small separations, $s < 0.1$, we have assumed that there will be no planet detections, which is reasonable even for massive planets. For larger separations, the term in square brackets is a modifier and the fractional term is an analytic estimate for the separation between the two planetary caustics (Han 2006). We chose the separation of the planetary caustics to set the CROIN size because the demagnification region lies between these two caustics. The constant term in the modifier accounts for the fact that the demagnification region is elongated along the binary axis somewhat (i.e., actually larger than the separation between the caustics). The strong s^2 term attempts to grow the CROIN to include the central caustic as the projected separation approaches resonance and the central caustic grows and begins to cause a significant number of detections in relatively high-magnification events. Detections via the central caustic form the “goat horn” feature in the scatter plot of detections in Figure 1, and it can be seen that for planets with $q \gtrsim 10^{3.5}$ the s^2 growth term will begin to significantly degrade the efficiency gains of the CROIN parameterization, but not so for lower mass ratios where the \sqrt{q} term keeps $u_{c, \max} < 1$.

2.2.3. Wide Topology

The lens has a wide topology when (Erdl & Schneider 1993)

$$s > \sqrt{\frac{(1+q^{\frac{1}{3}})^3}{1+q}}. \quad (4)$$

The wide topology has two caustics: a central caustic like the close topology lens and a single planetary caustic that lies on the binary axis. We choose the center of the planetary caustic as our reference point. Again, accurate analytical expressions for this position exist (Bozza 2000a; Han 2006)

$$(x_c^w, y_c^w) = \left(s - \frac{1}{[1+q]s}, 0 \right). \quad (5)$$




Figure 3. Plot of the CROIN impact parameter u_c against mass ratio q for resonant topology lenses from Figure 1. Gray points are non-detections while black points are detections. Lines show potential power laws that could be used to describe the upper envelope of detections. The solid line, $|u_c| = 4.5 q^{1/4}$, shows the actual choice for $u_{c,max}^r$.

We choose the CROIN size for wide topologies to be

$$u_{c,max}^w = [4 + \min(90 s^2, 160 s^2)] \sqrt{q}. \quad (6)$$

The \sqrt{q} term is simply the size of the Einstein ring of the planet at large $s \gg 1$. Similar to the close topology, contours of the deviation from the magnification of a point lens are elongated along the binary axis, and a similar constant plus growth term modifier is applied. For $s > (4/3)^2$, the $160 s^{-2}$ term is used to grow the CROIN as s shrinks closer to resonance, but below this the CROIN is large enough to cover the central caustic and the $90 s^2$ term begins to shrink the CROIN as s decreases further in a way that approximately matches the slope of the upper envelope of the right-hand goat horn.

2.2.4. Resonant Topology

The resonant topology only has a single caustic near the primary lens. While the planetward (positive x) side of the resonant caustic is larger in extent than the side closest to the primary that extends to negative x , it is significantly weaker, and negative perturbations to the magnification pattern extend away from the caustic on the negative side. We therefore choose the primary lens position as the reference point

$$(x_c^r, y_c^r) = (0, 0). \quad (7)$$

The CROIN size for resonant lenses is chosen to be

$$u_{c,max}^r = 4.5 q^{1/4}. \quad (8)$$

There is only a small range of s with a resonant configuration for small mass ratios, so there is little use in attempting to find a dependence on s . We found the simple scaling by examining the detections in Figure 3 and fitting a power law by eye to their upper envelope. A $q^{1/3}$ power law fits the distribution as well as a $q^{1/4}$, but we choose the latter in order to be lenient at low q where our statistics are poor.




Figure 4. Schematic diagram of the $(u_0, t_0) \rightarrow (u_c, t_c)$ coordinate transformation. The red line shows the source trajectory, while black lines define the geometry. For the sake of completeness, we show the general case of the caustic that lies off the reference axis. We adopt the standard convention that the impact parameter is positive if the source passes the reference point in a clockwise fashion.

(A color version of this figure is available in the online journal.)

2.3. Converting Between Parameterizations

Both the CROIN and standard parameterizations are impact-parameter-based. Converting the parameters of one to that of the other is simply a case of moving the origin. We provide the more general conversion for a CROIN center that may lie off the binary axis at (x_c, y_c) , though in all the cases we consider here the center of the CROIN lies on the binary axis. The off-axis situation may occur if you wished to extend the CROIN parameterization to smaller separations by taking the center of the secondary caustics as the CROIN center, or were considering multiplanet systems where it is no longer possible for all the planets to lie on the same axis.

Figure 4 shows the geometry of the coordinate transformation. From this geometry it is relatively straightforward to derive the following expressions

$$\frac{t_c - t_0}{t_E} = x_c \cos \alpha + y_c \sin \alpha, \quad (9)$$

$$u_c - u_0 = -x_c \sin \alpha + y_c \cos \alpha. \quad (10)$$

These expressions are general and apply to the transformation between any two impact-parameter-based parameterizations. If one's preferred reference point for the standard parameterization differs from the position of the primary lens, the above formulae can be adapted with the further transformations $x_c \rightarrow (x_c - x_0)$ and $y_c \rightarrow (y_c - y_0)$, where (x_0, y_0) is the position of the alternate reference point relative to the primary lens.

3. SPEED-UP USING CROIN

The simplest way to define and estimate the speed-up that the CROIN allows over the standard parameterization is to ask what fraction of events drawn from the standard parameter range fall within the CROIN parameter range. The speed-up is the inverse of this fraction. This can be computed trivially by Monte Carlo sampling without actually simulating events. We did this for a range of s and q and show the results in Table 1. In these computations, we drew u_0 uniformly from the range $-u_{0,max}$ to $u_{0,max}$ with $u_{0,max} = 3$, α uniformly from 0 – 2π , and assume that there is a 100% duty cycle, i.e., that the microlensing events peak

Table 1
Speed-up as a Function of s and q

$\log q$	$\log s$							
	−0.75	−0.50	−0.25	0.00	0.25	0.50	0.75	1.00
−2	1.8	1.2	1.0	2.1	1.0	2.0	8.9	26
−3	6.7	2.9	1.9	3.8	1.7	5.3	28	83
−4	20	8.0	6.0	6.6	5.5	15	88	265
−5	67	24	18	11	17	48	276	854
−6	208	76	57	20	55	155	922	2733
−7	691	248	193	36	179	462	2785	8628

Notes. Assumes 100% duty cycle and $u_{0, \max} = 3$; speed-up increases for lower duty cycles. Bold values indicate the traditional lensing zone (Wambsganss 1997).

within the period when the survey is observing (the observing season). This may not be a good assumption for surveys with short seasons, as we discuss in Section 4.1. In the traditional lensing zone ($0.6 \lesssim s \lesssim 1.6$; Wambsganss 1997) the speed-up is fairly modest until $q \lesssim 10^{-5}$. At large separations, $s \gtrsim 5$, the speed-up is always considerable, and for low mass ratios it is extreme.

4. ACCURACY AND BIAS

Our goal now is to assess the accuracy of the CROIN parameterization and make ourselves aware of any biases in the parameter distributions of detected planets that it introduces. To assess bias, we just want to measure accuracy as a function of various parameters.

The accuracy and bias is impossible to measure without reference to some observational setup, because this is what determines whether or not a planet will be detected. The simulations described in this paper are of the WFIRST-AFTA mission, which is likely to be the most sensitive microlensing survey it is cost-effective to perform, so the accuracy of the CROIN parameterization for other surveys is likely to be higher than that which we determine for WFIRST-AFTA. While one may consider the absolute number of detections per unit of time to be the best measure of survey sensitivity, for the purposes of assessing the accuracy of the CROIN parameterization this is not actually the case, because the total number of detections is primarily determined by the number of microlensing events that are monitored. Instead, it is the noise floor and the cadence that determine the most subtle features that can be detected. Of the currently imagined microlensing surveys, WFIRST-AFTA undoubtedly has the lowest noise floor due to its space-based photometry and high-resolution (minimizing the effect of blending). In the simulations here we have assumed a noise floor of 1 mmag, which would be extremely difficult to achieve from the ground. In terms of planned cadence, WFIRST at 15 minutes has one of the highest cadences for a microlensing survey,² though KMTNet has a slightly higher cadence at 10 minutes. This is not a sufficiently large difference to counter WFIRST’s space-based advantage, though we will discuss how to adjust the CROIN for different cadences in the next subsection.

The simulation data we used in Section 2.2 does not have enough low-mass planet detections to properly assess any biases that using the CROIN introduces, so we ran significantly

Table 2
Simulation Parameters

Parameter		
Survey	WFIRST-AFTA	
Cadence	15 minutes	
Season length	72 days	
Seasons	6	
Out-of-season gaps	~110 days or ~840 days	
Noise floor	1 mmag	
	Standard	CROIN
u_0 range	−3 → 3	No limit
t_0 range	0 → 2010	No limit
u_c range	No limit	− $u_{c, \max}$ → $u_{c, \max}$
t_c range	No limit	In season

Table 3
Incompleteness of Each Parameterization

Mass	Incompleteness f_{missed}	
	CROIN	Standard
1 M_{\oplus}	1.34% ± 0.22%	6.88% ± 0.10%
100 M_{\oplus}	6.87% ± 0.91%	5.22% ± 0.33%

larger simulations. These simulations were too computationally expensive to cover a large range of masses, so we only performed simulations of planets at two fixed masses: 1 Earth mass and 100 Earth masses. For each planet mass we performed two simulations. The first drew event parameters from the standard parameterization, and we used this to compute the accuracy of and biases introduced by the CROIN parameterization; we refer to this simulation as the *standard* simulation. The second simulation (referred to as the *CROIN* simulation) drew event parameters using the CROIN parameterization, and we used this to compute the accuracy of the standard parameterization, i.e., to estimate the number of planet detections that are missed by imposing a cut on u_0 .

These simulations are essentially identical to the *WFIRST* simulations in (Spergel et al. 2013), and which will be presented in detail in M. T. Penny et al. (in preparation). Full details of the simulation mechanics are given by Penny et al. (2013). The important details of the simulations are summarized in Table 2.

We measure the accuracy of each parameterization relative to the other by defining the incompleteness f_{missed} . For the CROIN parameterization, we measure it with respect to the standard parameterization, i.e., f_{missed} is the fraction of planet detections drawn from the standard parameterization limits for which the CROIN parameters are within the CROIN parameterization limits. We do the opposite to measure the standard parameterization incompleteness by running a simulation and drawing u_c and t_c uniformly and weighting events appropriately to take into account the variable $u_{c, \max}$. Our results are summarized in Table 3. We find that for 1 M_{\oplus} planets, the CROIN parameterization has 1.3% incompleteness relative to the standard parameterization. For 100 M_{\oplus} planets the CROIN parameterization’s incompleteness is 6.9%. The CROIN completeness relative to the standard parameterization is worth comparing with its opposite, the standard completeness relative to the CROIN parameterization. For 1 and 100 M_{\oplus} planets this is 6.9% and 5.2%, respectively when $u_{0, \max} = 3$. If we had chosen $u_{0, \max} = 1$, the incompleteness of the standard relative to the CROIN parameterization would have been 21% for both 1 and 100 M_{\oplus} planets. As we shall see below, the planet detections that the standard parameterization misses are primarily planets with large separations that are likely to

² Current ground-based follow-up networks achieve cadences ~1 minute when observing extremely high-magnification events (e.g., Dong et al. 2009), but the CROIN parameterization was not designed for high-magnification events, which are well served by the standard parameterization anyway.




Figure 5. Top: red and black lines plot of the number of $1 M_{\oplus}$ planet detections as a function of the time of closest approach of the source to the primary lens (t_0) for events simulated with the standard and parameterizations, respectively. The gray line shows the distribution of t_0 for all microlensing events that were detectable in the standard simulation (where detectable means the event caused a $\Delta\chi^2 > 500$ deviation relative to a flat light curve). Bottom: the same, but plotted against the time of closest approach to the center of the CROIN (t_c). Lower panels of each plot zoom in on the lowest portion of the main plot to show low-level features. Error bars show the typical uncertainties on the number of detections in the in-season bins (bin-width is 6 days). Dashed lines show the boundaries of the *WFIRST* seasons, which last 72 days.

(A color version of this figure is available in the online journal.)

be observed as “free-floating” planets with no detection of the host’s microlensing event.

4.1. Biases

Using the CROIN parameterization will necessarily bias the results of a simulation, so we must be aware of how each parameter will be affected and the magnitude of any bias. We consider how using the CROIN parameterization affects the distribution of each of the fundamental binary microlensing event parameters, beginning with t_0 .

Figure 5 shows the number of planet detections N_{det} as a function of both t_0 and t_c . For the *WFIRST* survey, 24% of detectable microlensing events (where the microlensing event causes a $\Delta\chi^2 > 500$ deviation relative to a flat baseline) peak outside the observing seasons, and a significant fraction ($\sim 10\%$) of all planet detections occur in such events. In this plot, the detections from the CROIN simulation are restricted to the same parameter range as the standard simulation ($|u_0| < 3$, $0 < t_0 < 2010$ days). The t_0 distributions from each simulation match extremely well, implying that any bias in the t_0 distribution due to the use of the CROIN parameterization is small.

The comparison of the t_0 and t_c distributions of planet detections in Figure 5 demonstrates the value of using the CROIN parameterization in low-duty cycle simulations. For $1 M_{\oplus}$ planets just 0.8 ± 0.2 of planets have t_c falling outside




Figure 6. Upper panel: the event timescale distribution of planet detections for $1 M_{\oplus}$ (red/pink) and $100 M_{\oplus}$ (black/gray) planet detections. The histogram drawn with a solid line is the result of the standard simulation. The lighter colored solid fill shows those events in the standard simulation that fall within the CROIN parameter limits. The dashed line shows the results of the CROIN simulation. Lower panel: the incompleteness f_{missed} of the CROIN parameterization relative to the standard simulation. If the incompleteness is less than 10^{-3} a point is plotted on the bottom axis (e.g., the points below $t_E = 5$ days). Color coding is the same as in the upper panel.

(A color version of this figure is available in the online journal.)

the observing season, whereas if we were to restrict ourselves to only events with t_0 in season to improve the efficiency of our simulations, we would miss $10.5\% \pm 0.7\%$ of planet detections (for $1 M_{\oplus}$ planets). For $100 M_{\oplus}$ planets, the fractions outside the season are more comparable due to the longer timescale of the planetary perturbation. In this case, $4.8\% \pm 0.8\%$ of detections have t_c fall outside the season, while $9.6\% \pm 1.0\%$ have t_0 fall outside the season.

Figure 6 plots the incompleteness of the CROIN parameterization relative to the standard parameterization as a function of the event timescale in the bottom panel, and demonstrates the bias introduced by using the CROIN parameterization in the top panel. Looking at the top panel, we can see that the bias in the timescale distribution introduced by the CROIN parameterization will be a small change in the slope of the already steeply declining large-timescale tail. However, while the apparent change in shape of the timescale distribution is difficult to notice (by comparing the solid to the dashed line), the incompleteness steadily rises as t_E increases, and approaches 100% for the extremely rare, longest-timescale events.

We can use this trend of incompleteness with the event timescale to judge the impact of changing the cadence of observations, despite running our simulations at a fixed cadence. For an microlensing event with fixed q and fixed source angular diameter relative to the Einstein radius, increasing the timescale by some factor has exactly the same effect on the χ^2 as increasing the cadence by the same factor. We can therefore estimate the necessary change in $u_{c, \text{max}}$ by finding the required $u_{c, \text{max}}$ that would make the incompleteness match that of the whole sample in bins of t_E . Figure 7 shows this for both simulations. $u_{c, \text{max}}$ scales as roughly as $t_E^{1/2}$, implying that it should also scale as $f^{1/2}$, where f is the frequency of observations. Of course, we have assumed a fixed mass and not a fixed mass ratio, so this argument may not be strictly valid.




Figure 7. Plot of the $u_{c,\max}$ that would be required to maintain the average incompleteness of 1.3% for $1 M_{\oplus}$ planets (red points) or 6.9% for $100 M_{\oplus}$ planets (black points) as a function of t_E . Arrows represent upper limits. The approximate $t_E^{-1/2}$ scaling can be used to adjust $u_{c,\max}$ to account for changes in cadence relative to the 15 minute WFIRST cadence we have simulated, because increasing the timescale will cause the $\Delta\chi^2$ of a planet detection to change in the approximately the same way as increasing the cadence for a fixed mass-ratio planet.

(A color version of this figure is available in the online journal.)

However, the mean mass ratio across the entire range of t_E only changes by a factor of ~ 2 compared to the factor of ~ 100 range of t_E , so the effect of the changing mean mass ratio is likely to be small.

Figure 8 shows the distributions of u_0 , s , α , and q for the planet detections in the standard and CROIN simulations. It is the distributions of u_0 and s where the largest differences can be seen between the standard and CROIN parameterizations.

Unsurprisingly, the CROIN parameterization performs poorly for high-magnification events, and the threshold $|u_0|$ for poor performance gets larger for larger planet masses. For the highest-magnification events, the incompleteness seems to saturate at around 10% for $1 M_{\oplus}$ planets and $\sim 20\%$ for $100 M_{\oplus}$ planets, before falling as a power law as $|u_0|$ increases with an index that seems to depend on the planet mass. The saturation at $f_{\text{missed}} < 1$ for small u_0 is not too surprising, because the CROIN parameterization deals well with high-magnification events with resonant topologies (when it is centered on the primary lens), and when s is close to 1 but not resonant the CROIN can grow large enough to encompass the primary lens. The plot of the cumulative incompleteness shows that for any reasonable choice of $u_{c,\max}$, the total incompleteness of the CROIN parameterization relative to the standard will not increase significantly, e.g., for $1 M_{\oplus}$ planets and $u_{0,\max} = 1$, $f_{\text{missed}} = 1.57\%$ compared to 1.34% for $u_{0,\max} = 3$.

Use of the CROIN does not significantly bias the distribution of s compared to using the standard parameterization. However, not using the CROIN does steepen the tail of the distribution at large s . As previously noted, however, the majority of these events that the standard parameterization misses will appear to be free-floating planets with no sign of a the host star's microlensing event. The largest incompleteness as a function of s occurs when the goat horns that are seen in Figure 1 caused by high-magnification events cross the $u_{c,\max}$ curve. This occurs

Table 4
Speed-up for Different Duty Cycles with $0.1 < a < 100$ AU^a

Duty-cycle Example Mass (M_{\oplus})	100%	80% KMTNet	25% WFIRST	Incompleteness ^b (%)
10000	1.4	1.7	5.5	18
1000	1.9	2.4	7.7	9.8
100	3.7	4.6	15	5.3
10	10	13	41	2.8
1	31	39	124	1.5
0.1	99	124	397	0.8
0.01	312	390	1247	0.4

Notes.

^a Semimajor axis is distributed logarithmically.

^b Incompleteness estimated using Equation (11) from Section 4.1 assuming an average mass ratio of $8.1 \times 10^{-6} (M/M_{\oplus})$.

at different values of s for different planet masses (with s being further away from 1 for larger planet masses).

The incompleteness as a function of the source trajectory angle α is not of much interest, but we include it in Figure 8 for completeness. Note, however, that for Earth-mass planets, the incompleteness seems to be smallest for trajectories that are either parallel or perpendicular to the binary axis. This pattern is not repeated for $100 M_{\oplus}$ planets.

The final plot in Figure 8 shows the distribution of mass ratios for our 1 and $100 M_{\oplus}$ simulations, as well as the incompleteness as a function of mass ratio. To extend the range of mass ratios that we consider, and to get a sense of the trend, we also plot the incompleteness as measured from the simulation we used to define $u_{c,\max}$, which is shown in Figure 1. The range of semimajor axis is smaller in the standard simulation ($0.3 < a < 30$ AU, red and black points) than the Figure 1 simulation ($0.1 < a < 100$ AU), but this has a negligible effect due to the strong decline in the number of detections as the separation becomes large or small. The trend is well described by a power law, with incompleteness increasing with q . A fit to the data from all the simulations yields the relation

$$f_{\text{missed}} = (1.6 \pm 0.3)\% \left(\frac{q}{10^{-5}} \right)^{0.27 \pm 0.04}. \quad (11)$$

5. DISCUSSION

5.1. Speed versus Accuracy

With estimates of the speed-up from Section 3 and the accuracy from Section 4, we can now objectively assess the usefulness of the CROIN parameterization. In Table 4, we summarize both the speed-up and accuracy it is possible to achieve for realistic simulations that will cover a range of semimajor axes and masses. We have computed the speed-up in the same way as in Section 3, but this time assumed a distribution of semimajor axes in the range $0.1 < a < 100$ AU and a range of duty cycles representative of realistic microlensing surveys. For $100 M_{\oplus}$ planets, an incompleteness of $\sim 5\%$ seems a reasonable incompleteness for most applications, and this can provide a significant speed-up of $4\text{--}15\times$. A 10% incompleteness for $1000 M_{\oplus}$ planets, however, seems too large, and the speed-up is a factor of two smaller. The CROIN parameterization therefore seems most likely to be useful for planets of $100 M_{\oplus}$ and below. For planets of Earth-mass and below the speed-ups are huge, at the minor cost of an inaccuracy of the order of 1% or less.

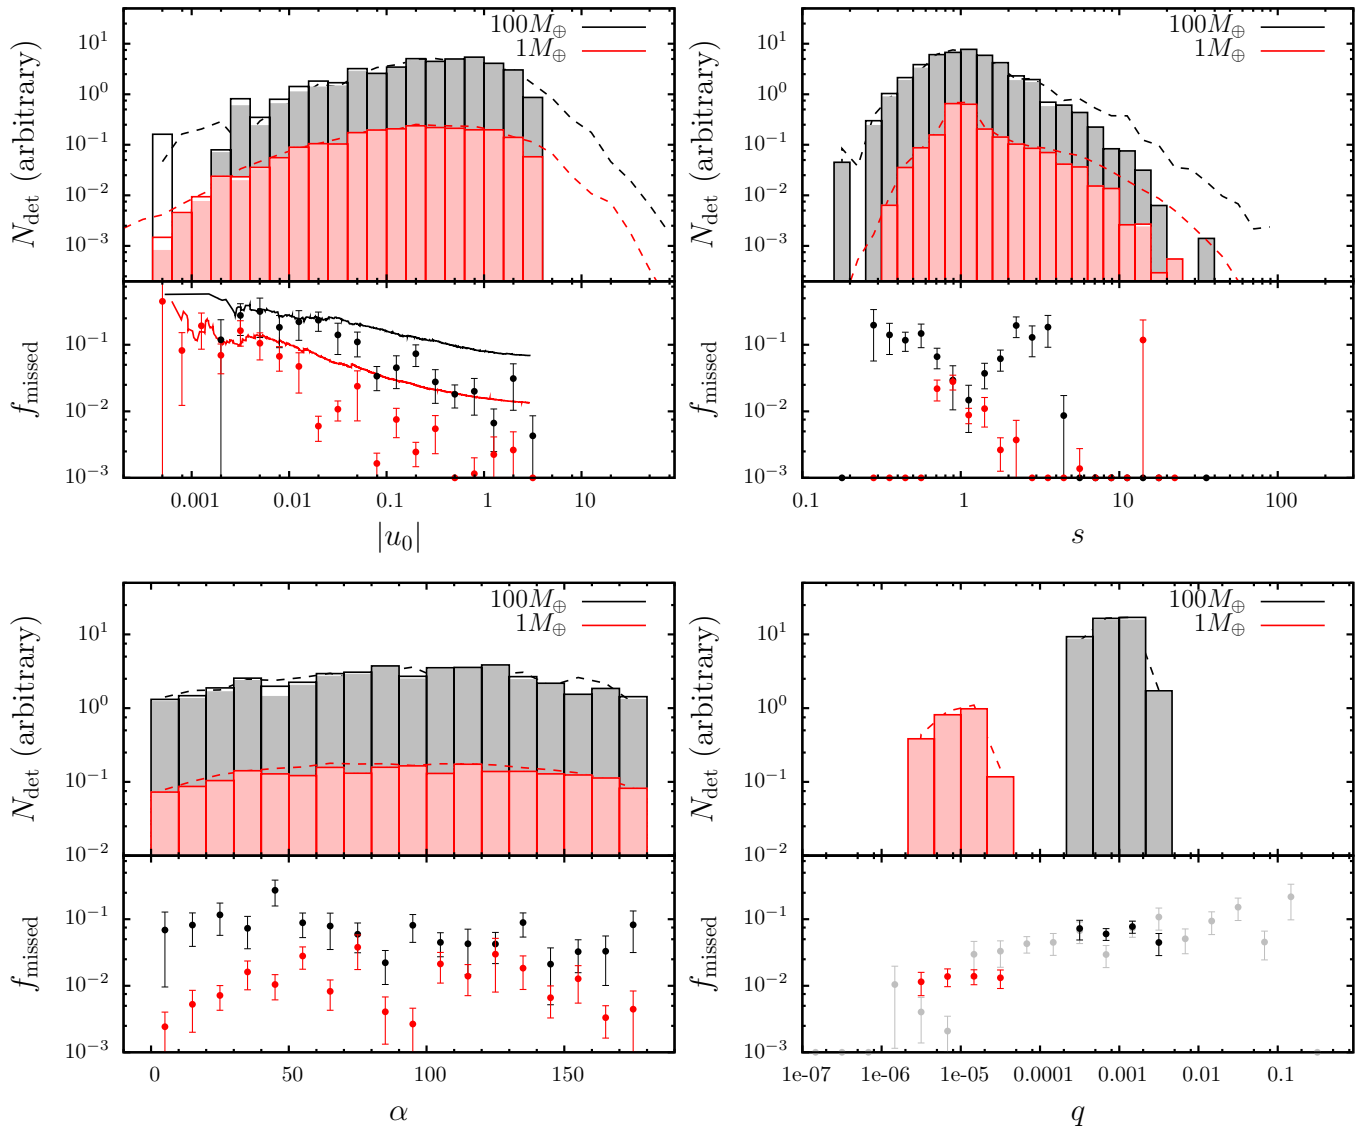


Figure 8. Same plots as Figure 6, but for the other binary microlensing parameters, u_0 , s , α , and q (arranged left to right, top to bottom, respectively). The solid lines in the incompleteness panel of the u_0 plot are the cumulative incompleteness $f_{\text{missed}}(< |u_0|)$, showing the incompleteness one would suffer if $u_{0,\text{max}}$ were chosen to be smaller than we have chosen. Gray points in the incompleteness panel of the q plot are computed from the data shown in Figure 1; note that the semimajor axis range simulated in Figure 1's data is larger ($0.1 < a < 100$ AU, gray points) than that for the test simulations shown with black and red points ($0.3 < a < 30$ AU).

(A color version of this figure is available in the online journal.)

The dominant cause of the inaccuracy of CROIN parameterization is the missing planet detections caused by the central caustic of close and wide separation planets. However, it is relatively easy to use a strategy that does not miss these detections. Rather than use the CROIN for the origin of the parameterization and drawing $|u_c| < u_{c,\text{max}}$, one would use the standard parameterization, drawing events from its usual limits (e.g., $|u_0| < 3$). Now, one would perform two tests on the parameters: the first would check if it was a high-magnification event (with some limit on $|u_0|$), and the second would convert the standard parameters to the CROIN parameters and check that the CROIN parameters were within the limits defined in this paper. If either of the tests were passed, then the light curve would be computed, otherwise the event would be assumed to be a non-detection. To do this correctly, one would need to perform a similar study to this to determine the u_0 limits for central caustic planet detections. As we were concerned primarily with speeding up low-mass planet simulations, where the number of detections due to non-

resonant central caustics is expected to be extremely small, we have not conducted this study.

Finally, it should be noted that the speed-up that will actually be achieved may be smaller than that we estimate here if light curves that contain detectable planet signatures take longer to generate than ones that do not. This will quite often be the case if one is using the hexadecapole approximation to avoid finite source calculations (Pejcha & Heyrovský 2009; Gould 2008). Even so, the actual speed-up is still worthwhile.

5.2. Potential Uses

The magnitude of speed-up it is possible to achieve with the CROIN parameterization make its application to microlensing simulations and detection efficiency calculations immediately obvious. The advent of new and proposed microlensing surveys has encouraged a flurry of new work on microlensing simulations (Shvartzvald & Maoz 2012; Green et al. 2012; Penny

et al. 2013; Spergel et al. 2013; Henderson et al. 2014; Ipatov et al. 2014). Interpretation of these new high cadence surveys will be significantly easier than for previous observations in survey plus follow-up mode (e.g., Gould et al. 2010), but will still require extensive calculations of detection efficiencies over a large parameter space (e.g., Gaudi & Sackett 2000; Gaudi et al. 2002; Tsapras et al. 2003; Snodgrass et al. 2004). In fact, the increased cadence and area of the surveys will make such analyses significantly more computationally expensive (though Moore’s law will help to a certain extent). A larger challenge will be presented by the advent of space-based microlensing planet searches. These will provide an order of magnitude more microlensing events to search for planets, but will also measure the light curves significantly more accurately, increasing the demands on the light curve computations, which are already probably close to maximum efficiency. Maximizing the scientific return of space-based surveys will require extensive simulations to optimize the various aspects that can affect the mission, from hardware to survey design. This has proved challenging so far, with only limited parameter exploration possible for low-mass planets. The CROIN parameterization represents an important way to broaden the scope of planning for these missions, increasing the size of parameter space that can be explored.

The CROIN parameterization may also find use in the modeling of individual gravitational microlensing events. As mentioned in Section 1, parameterizations centered on the caustics are already in use. The analytic limits on the impact parameters of these parameterizations could be useful in restricting a parameter search. However, as downhill fitting will quickly move a trial solution to a local minimum of the parameter space, any speed-up will be modest.

6. CONCLUSIONS

We have proposed a parameterization of binary gravitational microlensing applicable to planetary microlensing. We have empirically determined an analytic functional form for the limits of the impact parameter within which the vast majority of planetary detections can be expected. We have shown that by using this parameterization and its analytic limits it is possible to speed-up simulations of planetary microlensing by factors of 10–1000 depending on the mass of planet being investigated. This comes at a cost of excluding a small percentage of planet detections, though this is smaller than or comparable to the loss due to arbitrary truncation of the space of the standard parameterization.

I thank Scott Gaudi for helpful discussions, Radek Poleski for pointing out an oversight, Eamonn Kerins, Nick Rattenbury, Shude Mao, and Annie Robin, whose work has contributed to the development of *MaBμLS*, and finally Brett Andrews for

his zoomorphic insight. I also thank the anonymous referee for suggestions that improved the clarity of the paper. Some of the computational element of this research was achieved using the High Throughput Computing facility of the Faculty of Engineering and Physical Sciences, The University of Manchester. This work was supported by The Ohio University Department of Astronomy Jefferson Chair.

REFERENCES

- Albrow, M. D., Beaulieu, J.-P., Caldwell, J. A. R., et al. 1999, *ApJ*, **522**, 1022
 An, J. H., Albrow, M. D., Beaulieu, J.-P., et al. 2002, *ApJ*, **572**, 521
 Bennett, D. P. 2010, *ApJ*, **716**, 1408
 Bennett, D. P., Sumi, T., Bond, I. A., et al. 2012, *ApJ*, **757**, 119
 Bozza, V. 2000a, *A&A*, **355**, 423
 Bozza, V. 2000b, *A&A*, **359**, 1
 Bozza, V. 2010, *MNRAS*, **408**, 2188
 Cassan, A. 2008, *A&A*, **491**, 587
 Cassan, A., Horne, K., Kains, N., Tsapras, Y., & Browne, P. 2010, *A&A*, **515**, A52
 Cassan, A., Kubas, D., Beaulieu, J.-P., et al. 2012, *Natur*, **481**, 167
 Dominik, M. 1998, *A&A*, **333**, L79
 Dominik, M. 2007, *MNRAS*, **377**, 1679
 Dong, S., Bond, I. A., Gould, A., et al. 2009, *ApJ*, **698**, 1826
 Dong, S., DePoy, D. L., Gaudi, B. S., et al. 2006, *ApJ*, **642**, 842
 Erdl, H., & Schneider, P. 1993, *A&A*, **268**, 453
 Gaudi, B. S., Albrow, M. D., An, J., et al. 2002, *ApJ*, **566**, 463
 Gaudi, B. S., & Petters, A. O. 2002a, *ApJ*, **574**, 970
 Gaudi, B. S., & Petters, A. O. 2002b, *ApJ*, **580**, 468
 Gaudi, B. S., & Sackett, P. D. 2000, *ApJ*, **528**, 56
 Gould, A. 2008, *ApJ*, **681**, 1593
 Gould, A., & Gaucherel, C. 1997, *ApJ*, **477**, 580
 Gould, A., Dong, S., Gaudi, B. S., et al. 2010, *ApJ*, **720**, 1073
 Green, J., Schechter, P., Baltay, C., et al. 2012, arXiv:1208.4012
 Han, C. 2006, *ApJ*, **638**, 1080
 Henderson, C. B., Gaudi, B. S., Han, C., et al. 2014, *ApJ*, submitted (arXiv:1406.2316)
 Ipatov, S. I., Horne, K., Alsubai, K. A., et al. 2014, in IAU Symp. 293, Formation, Detection, and Characterization of Extrasolar Habitable Planets (Cambridge: Cambridge Univ. Press), **416**
 Kim, S.-L., Park, B.-G., Lee, C.-U., et al. 2010, *Proc. SPIE*, **7733**, 77333F
 Pejcha, O., & Heyrovský, D. 2009, *ApJ*, **690**, 1772
 Penny, M. T., Kerins, E., Rattenbury, N., et al. 2013, *MNRAS*, **434**, 2
 Press, W. H., Flannery, B. P., & Teukolsky, S. A. 1986, Numerical Recipes. The Art of Scientific Computing (Cambridge: Cambridge Univ. Press)
 Rattenbury, N. J., Bond, I. A., Skuljan, J., & Yock, P. C. M. 2002, *MNRAS*, **335**, 159
 Schneider, P., & Weiss, A. 1986, *A&A*, **164**, 237
 Shvartzvald, Y., & Maoz, D. 2012, *MNRAS*, **419**, 3631
 Skowron, J., & Gould, A. 2012, arXiv:1203.1034
 Snodgrass, C., Horne, K., & Tsapras, Y. 2004, *MNRAS*, **351**, 967
 Spergel, D., Gehrels, N., Breckinridge, J., et al. 2013, arXiv:1305.5422
 Sumi, T. 2010, in ASP Conf. Ser. 430, Pathways Towards Habitable Planets, ed. V. Coudé du Foresto, D. M. Gelino, & I. Ribas (San Francisco, CA: ASP), **225**
 Tsapras, Y., Horne, K., Kane, S., & Carson, R. 2003, *MNRAS*, **343**, 1131
 Udalski, A. 2011, in XV International Conference on Gravitational Microlensing: Conference Book, ed. V. Bozza, S. C. Novati, L. Mancini, & G. Scarpetta, arXiv:1102.0452
 Wambsganss, J. 1997, *MNRAS*, **284**, 172
 Yee, J. C., Hung, L.-W., Bond, I. A., et al. 2013, *ApJ*, **769**, 77