

Small-scale structures of dark matter and flux anomalies in quasar gravitational lenses

R. Benton Metcalf^{1,2*} and Adam Amara³

¹Max Planck Institut für Astrophysics, Karl-Schwarzschild-Str. 1, 85741 Garching, Germany

²Dipartimento di Astronomia, Alma Mater Studiorum Università di Bologna, via Ranzani 1, 40127 Bologna, Italy

³Institute of Astronomy, ETH Zürich, Wolfgang-Pauli-Strasse 27, CH-8093 Zürich, Switzerland

Accepted 2011 October 10. Received 2011 October 10; in original form 2010 July 9

ABSTRACT

We investigate the statistics of flux anomalies in gravitationally lensed quasi-stellar objects as a function of dark matter halo properties such as substructure content and halo ellipticity. We do this by creating a very large number of simulated lenses with finite source sizes to compare with the data. After analysing these simulations, we conclude the following. (1) The finite size of the source is important. The point source approximation commonly used can cause biased results. (2) The widely used R_{cusp} statistic is sensitive to halo ellipticity as well as the lens' substructure content. (3) For compact substructure, we find new upper bounds on the amount of substructure from the fact that no simple single-galaxy lenses have been observed with a single source having more than four well separated images. (4) The frequency of image flux anomalies is largely dependent on the total surface mass density in substructures and the size–mass relation for the substructures, and not on the range of substructure masses. (5) Substructure models with the same size–mass relation produce similar numbers of flux anomalies even when their internal mass profiles are different. (6) The lack of high image multiplicity lenses puts a limit on a combination of the substructures' size–mass relation, surface density and mass. (7) Substructures with shallower mass profiles and/or larger sizes produce less extra images. (8) The constraints that we are able to measure here with current data are roughly consistent with Λ cold dark matter (Λ CDM) N -body simulations.

Key words: gravitational lensing: strong – dark matter.

1 INTRODUCTION

The cold dark matter (CDM) model with a cosmological constant (Λ CDM) has become the standard model of cosmology. This model is in good agreement with a variety of observational probes of the large-scale distribution of matter and galaxies in the Universe and is in general agreement with probes of the distribution of mass in galaxy clusters and in large galaxies. In the Λ CDM model, dark matter clumps into haloes and galaxies form in the haloes. On small scales, Λ CDM predicts that dark matter haloes exist down to very small masses; the exact lower limit depending on the properties of the CDM particle and its thermal history. It has long been recognized that the number of observed dwarf galaxies in the Local Group of galaxies falls well short of the number of predicted haloes (Klypin et al. 1999; Moore et al. 1999; Diemand, Kuhlen & Madau 2007b; Springel et al. 2008). This is referred to as the *substructure problem*. Either galaxy formation is highly suppressed in small-mass haloes or Λ CDM needs to be modified in some way by, for example,

changing the properties of the dark matter particle or the initial conditions for the density fluctuation in the Universe. Warm dark matter is a popular alternative. Whether or not these small-mass haloes exist has been one of the most pressing unanswered question in cosmology for a decade.

Metcalf & Madau (2001) demonstrated that if small-scale structure exists in the distribution of dark matter it will have a strong effect on the magnifications of quasar images in strong gravitational lenses. This effect causes the flux ratio between images to disagree with any lens model with a smooth distribution of matter. These cases are called *anomalous flux ratios*. A particular case had been studied by Mao & Schneider (1998) and subsequently it was shown that anomalies are common in quasar lenses (Dalal & Kochanek 2002; Metcalf & Zhao 2002). This work and a number of subsequent studies (see Zackrisson & Riehm 2010, for a review of the subject) relied on fitting lens models to individual lens systems. It has not yet been shown clearly what can be causing these anomalies and what cannot be causing them.

In a parallel approach, we and others have tried to simulate the lenses directly from cosmological N -body simulations to determine if they are consistent with the observed frequency of flux anomalies

*E-mail: bmetcalf@mpa-garching.mpg.de

(Bradač et al. 2004; Amara et al. 2006; Macciò et al. 2006; Xu et al. 2009). The first study predicted a large number of anomalies, but it may have been strongly affected by shot noise. The two more recent and higher resolution studies found that the substructure in the N -body simulations is not sufficient to cause the observed flux anomalies [also the conclusion of Mao et al. (2004)]. This is largely because of the small number density of substructures near the radii where images form (typically around 10 kpc in projection). These studies relied on only a few projections of a small number of high-resolution haloes. It is possible that these results are a statistical fluke or that the observed anomalies are largely caused by dark matter objects along the line of sight but not inside the halo of the primary lens (Metcalf 2005a,b). Answering the question of whether the N -body simulations have enough small-scale structure in them to account for the flux ratio anomalies is one of the primary goals of this paper.

It is very difficult to realistically simulate strong quasi-stellar object (QSO) lenses from an N -body simulation. The first, and most important, problem is that shot noise from the discrete particles has a strong effect on the image magnifications. Roughly, the error in the magnification goes as $\delta\mu \sim \mu^2/\sqrt{N_s}$, where μ is the magnification and N_s is the number of particles over which the smoothing is done. Since μ can be large, 100 or larger in the best cases for detecting substructure, the amount of smoothing needed to obtain an accuracy of even 10 per cent is very large. So much smoothing can even smooth out the very substructures one wants to detect. Because of this, Xu et al. (2009) replace an N -body simulation with a simple analytic model fit to an N -body simulation. A second problem is that the highest resolution simulations do not contain baryons. Baryons have a strong effect on the profile of the lens and in some cases dominate the mass within one Einstein radius. The baryons need to be put in ‘by hand’. A third problem is that the extremely high resolution simulations required provide one, or at best a few, dark matter haloes. Variations between haloes make their lensing properties and their tendency to produce anomalies very different. It will be demonstrated in this paper that only very limited conclusions about the CDM model can be drawn from a single simulated lens.

To avoid these problems, we take a different approach in this paper. We produce a large number of analytic lens models that are meant to reproduce the population of lenses expected in the Λ CDM model. We then determine the frequency of flux ratio anomalies in these lenses and compare it to the observed frequency. We adjust the properties and abundance of the substructures to see what kind of substructure is consistent with observations. The allowed statistical properties of the substructures are compared with the properties of N -body haloes.

All previous studies, except Amara et al. (2006), have also suffered from the problem that the sources are treated as infinitely small points. The magnification of individual images is calculated by taking derivatives of the gravitational force at the position of the image. It will be shown in this paper that since the physical size of the quasar radio or mid-infrared (mid-IR) emission regions are similar to the sizes of the substructures of interest the point source magnifications are not accurate approximations. We use a new, high-speed lensing code called Gravitational Lensing with Adaptive MESH Refinement (GLAMER) (Metcalf, in preparation) that is the first one capable of producing a very large number of simulated lenses with finite sources in a reasonable amount of time. It does this through an adaptive mesh refinement algorithm that will be briefly described in Section 2.3.

In Section 2, the models and techniques used to create simulated lenses are described. In Section 3, the results of those simulations are discussed. Ways of comparing the results to the available lensing data are presented in Section 4. The results are compared with the predictions of cosmological N -body simulations in Section 5. A summary and discussion are given in Section 6.

2 LENS SIMULATIONS

Our approach in this paper is to produce a large population of realistic simulated lenses and then compare their statistical properties to the observed population of lenses. To do this, we must develop a model for the population of gravitational lens that includes the host, galaxy+dark matter halo and the substructures within the host. We will not consider the effects of companion galaxies with masses roughly equivalent to the primary lens in this paper.

2.1 Host lens model

There is significant evidence from lensing and X-ray observations that early-type galaxies have a r^{-2} mass profiles (Fukazawa et al. 2006; Gavazzi et al. 2007; Churazov et al. 2010; Humphrey & Buote 2010). In accordance with this finding, we model the host lenses as distorted singular isothermal ellipsoids (DSIEs). The surface mass density for this model is

$$\kappa(r, \theta) \equiv \frac{\Sigma(r, \theta)}{\Sigma_{\text{crit}}} \quad (1)$$

$$= \frac{r_E}{r} \left[\frac{1}{\sqrt{\cos^2(\theta) + f^2 \sin^2(\theta)}} \right. \quad (2)$$

$$\left. + \frac{1}{2} \sum_n b_n (1 - n^2) \cos[n(\theta + \phi_n)] \right], \quad (3)$$

where the Einstein radius is

$$r_E = 4\pi \left(\frac{\sigma}{c} \right)^2 \frac{D_1 D_{\text{ls}}}{D_s} \quad (4)$$

and the critical surface density is

$$\Sigma_{\text{crit}} = \frac{c^2}{4\pi G} \frac{D_s}{D_1 D_{\text{ls}}}, \quad (5)$$

where D_1 , D_s and D_{ls} are the angular size distance to the lens, to the source and between the lens and the source, respectively. The first part (2) is a singular isothermal ellipsoid (SIE) whose lensing properties have been extensively studied (see Kormann, Schneider & Bartelmann 1994, for example). The deflection angle and shear caused by the series in (3) have been worked out by Evans & Witt (2003), although with different notation.

The perturbations b_n are assumed to be of the same order as the observed perturbations in the surface brightness profile of early-type galaxies. Typical values for b_3 and b_4 are 2 or 3 per cent, but accurate statistics are not available (Bender, Doebereiner & Moellenhoff 1988; Kormendy et al. 2009). We draw random values from a Gaussian distribution with variance 0.005 for b_2 and b_3 and 0.01 for b_4 . We take $n > 4$ terms to be zero. In the observations, b_4 is usually defined with the orientation of this mode fixed to the same axis as the axis of the elliptical component to define the ‘discyness’ or ‘boxyness’ of the galaxy. Since the alignment has important effects on the lensing properties, we relax this requirement somewhat and allow $\phi_{3,4}$ to vary from the position angle of

the elliptical component. The misalignment is normally distributed with variance 3° .

We also include background shear and convergence in the model. Dalal & Watson (2005) calculated the expected distribution of γ and κ in an N -body simulation at potential lenses. They found that κ and $|\gamma|$ are both roughly lognormally distributed with a variance of $\simeq 0.03$. We assume this distribution in our model. Analytic estimates by Keeton, Kochanek & Seljak (1997) are in agreement with this result, as are observations (Koopmans et al. 2006).

The model described above is what will be called the ‘standard’ host model. To test how sensitive magnification anomalies are to the host model, we perform a series of tests where the distortions to the lens are increased. For the ‘extra distorted model’, we triple the variance in the distortion modes and decouple their orientation from the orientation of the elliptical component. For the ‘extra shear model’, we triple the variance in the background shear and convergence.

2.1.1 Distributions of host properties

Calculating the expected distribution of the lenses’ redshifts, velocity dispersions and ellipticities requires knowing not only the source luminosity and redshift distributions of lenses and sources, but also the many selection effects that might be important. The sample of lenses we wish to compare our results with was discovered in many different ways and does not have a uniform, well-defined selection criterion. Instead of trying to model these biases, we use the distributions of already known lenses when possible.

For the lens and sources redshifts, we use the observed values for the CASTLES lenses.¹ There are 60 lenses with measured source and lens redshift pairs. We draw randomly from these sets of redshifts. The lenses discussed in Section 4 are a subsample of these.

To get a sample of host velocity dispersions, σ , we use the velocity dispersions from the SLACS lenses (Koopmans et al. 2006). This sample of 61 lenses is used to make a cumulative distribution of σ . The discrete distribution is linearly interpolated to get a continuous cumulative distribution and then this is randomly sampled from. In the SLACS sample, the measured velocity dispersion of stars and the velocity dispersion of the best-fitting SIE models have statistically indistinguishable distributions. We choose to use the best-fitting SIE velocity dispersions. These values range from 160 to 396 km s⁻¹.

The axial ratios, f , are sampled independently from the SLACS lenses in the same way as the velocity dispersions. No possible correlations between the internal structure of the lenses and their redshift are reproduced in this sampling. The average of this distribution is $f = 0.75$, the standard deviation 0.14 and the range is $0.37 < f < 0.98$. The SLACS lenses are at relatively low redshift because of their selection criterion, but observations indicate that the internal structure of early-type galaxies do not evolve significantly between $z = 1$ and 0 (Thomas et al. 2005).

We consider only four image quasar lenses in this paper, while the SLACS lenses include two image lenses. The asymmetry of the lens changes the area enclosed in the tangential caustic and thus a sample of four image lenses will tend to have more asymmetric lenses than a sample that includes all multiple image cases. To correct for this bias, we calculate the ratio of the area within the tangential caustic to the area within the radial caustic (or ‘cut’ in the case of a DSIE). The number of sources used for the lens is then proportional to this ratio. More circular galaxies will have less lenses in the final

sample. This corrects for the bias in the SLACS lenses relative to the four image quasars. From 0 to ~ 100 , source positions are used for each lens model. This method of using a variable number of sources per lens is something of a compromise; ideally one would have a population of lenses that reflected the biases and one source per lens, but to do this the caustic structure of each lens would need to be calculated and then many of those with a small cross-sections for producing four images would be discarded. This would be computationally inefficient. A small number of sources per lens means that the population of high cross-section lenses will be better sampled, but if the average number of sources per lens is set too low all the lenses with small cross-sections will have zero sources. We have set the number of sources per lens so that lenses with zero sources are rare (~ 1 per cent).

For each lens model, the source centres are chosen to randomly cover a region that encloses the region within the tangential caustic. Some of these source positions give rise to less than four images (when the source intersects the caustic or is completely outside the caustic) and some give rise to more than four images (when caustics structure is more complicated). The cases with less than four images are discarded in the analysis that follows.

2.2 Substructure model

We wish to construct a substructure model that reflects the expectations we have from N -body simulation, but is relatively simple and has a small number of parameters that can be varied to measure the agreement or disagreement with Λ CDM.

Simulations show that the mass fraction in substructure within a projected radius increases roughly linearly with projected radius (Diemand, Kuhlen & Madau 2007a,b; Springel et al. 2008). With a SIE mass model, this implies that the surface mass density of substructure is constant at least near the Einstein radius and interior to it. This will be assumed in all cases.

The mass function of subhaloes in N -body simulations is found to be a power law:

$$\frac{dn}{dm} \propto m^{-\alpha}, \quad (6)$$

where n is the number of substructures in a halo. Springel et al. (2008) found that $\alpha \simeq 1.9$ up to about 1/10 of the halo mass without any resolved lower mass limit in haloes as a whole. Transforming mass function into a projected mass function in two dimensions is not straightforward because of mass segregation in the host halo. The projected substructure number density will be denoted η and the projected mass function will be $d\eta/dm$.

It is found that the substructures of different masses are distributed within host haloes in remarkably similar ways except that at each radius the mass function has an upper mass cut-off (Springel et al. 2008). If it were not for this mass cut-off, the projected mass function (surface number density) would have the same slope as the total mass function. Instead, the projected mass function will become steeper than $\alpha = 1.9$ above some mass scale. We represent this effect in our model crudely with an upper mass cut-off that is smaller than the one found for the complete mass function:

$$\frac{d\eta}{dm} \propto \begin{cases} m^{-\alpha}, & m_{\min} < m < m_{\max} \\ 0, & \text{otherwise,} \end{cases} \quad (7)$$

with $\alpha \simeq 1.9$. This is a crude model that could be improved on in the future.

The maximum mass in the mass function must be a function of host halo size. A mass scale for the host can be defined as the mass

¹ <http://www.cfa.harvard.edu/castles/>

within a fixed radius ($M \propto \sigma^2$) or the mass within a radius where the average density reaches a fixed threshold ($M \propto \sigma^3$). The latter is the one commonly used to define the mass of a halo in cosmology although the virial radius is generally larger than the radii over which one would expect the SIE model to hold. However, if the concentration of the haloes does not vary greatly within the range of host lenses, then the same scaling would be expected in the inner regions. Making the maximum substructure mass a fixed fraction of the host halo mass results in

$$m_{\max}(\sigma) = M_{\max} \left(\frac{\sigma}{\sigma_*} \right)^3. \quad (8)$$

The same scaling is assumed for the minimum mass. M_{\min} is used as an adjustable parameter to change the mass scale and test the data's consistency with a mass cut-off as would be expected in many alternative theories to CDM. The normalizing halo is fixed to $\sigma_* = 200 \text{ km s}^{-1}$.

The normalization of the mass function (7) needs to be set. To agree with N -body simulations, the fraction of mass in substructure at a fixed fraction of the virial radius should be the same in all haloes. Since $R_{\text{host}} \propto \sigma$ and (8) makes average mass scale like σ^3 , the normalization must scale like σ^{-1} . Explicitly the result is

$$\frac{d\eta}{dm} = \eta_* \left(\frac{\sigma_*}{\sigma} \right) \frac{(1-\alpha)}{[m_{\max}^{1-\alpha} - m_{\min}^{1-\alpha}]} m^{-\alpha}. \quad (9)$$

The parameter η_* is then the total surface number density of substructures, for a host with $\sigma = \sigma_*$ and is not a function of projected radius.

Although the mass fraction in substructure at a fixed fraction of the halo radius is the same for all lenses, the same is not true at the Einstein radius. Since $r_E \propto \sigma^2$, the total surface density at r_E is independent of σ for lenses and sources at the same redshift, which makes the mass fraction scale as σ^2 at this radius. As a result, we might expect substructure to be more important for larger lenses.

The internal structure of the substructures is, for simplicity, a simple power law with a cut-off radius:

$$\Sigma_{\text{sub}}(r) = \begin{cases} \frac{(2-\beta)}{2\pi} \frac{m}{R_{\text{cut}}(m,\sigma)^2} \left(\frac{R_{\text{cut}}(m,\sigma)}{r} \right)^\beta, & r < R_{\text{cut}}(m,\sigma) \\ 0, & r > R_{\text{cut}}(m,\sigma). \end{cases} \quad (10)$$

In the classical analytic treatment, the average mass density within the tidal radius is proportional to the average mass density of the host within the substructure's orbit (Binney & Tremaine 1987). This implies $R_{\text{cut}}(m,\sigma) \propto m^{1/3}$ if all the substructures are at the same distance from the centre of the host, which we assume. Since the mass density at a fixed fraction of the host halo radius is independent of the host size, it is expected that this relation is independent of the host size:

$$R_{\text{cut}}(m,\sigma) = R_{\max} \left(\frac{m}{M_{\max}} \right)^{1/3}, \quad (11)$$

where R_{\max} is a free parameter describing the size of the most massive substructures. In a more realistic model, there would be a significant scatter in the $R_{\text{cut}} - \sigma - m$ relation, but for our purposes this relation is sufficient. Using the classical tidal radius, the three-dimensional distance from the centre of the lens that this cut-off radius corresponds to is

$$R_{\text{galactic}} = 4.3 \text{ kpc} \left(\frac{\sigma}{200 \text{ km s}^{-1}} \right) \times \left(\frac{R_{\max}}{1 \text{ kpc}} \right)^{3/2} \left(\frac{10^9 M_{\odot}}{M_{\max}} \right)^{1/2}. \quad (12)$$

Our fiducial model will have $R_{\max} = 0.5 \text{ kpc}$ and $M_{\max} = 10^9 M_{\odot}$, so $R_{\max} \simeq 1.5 \text{ kpc}$ is a representative distance which is, perhaps, optimistically compact. We will vary R_{\max} from 0.25 to 4.0 kpc.

It should be noted that the appropriate R_{galactic} for lensing would be significantly smaller than the average R_{galactic} for subhaloes in general. Most subhaloes are at large radii ($\gtrsim 100 \text{ kpc}$) because there is so much volume at large radii to make up for the lower weighted number density. Projecting along the line of sight weights the inner regions of the halo more. The difference is an order of magnitude or more. This means that the substructures that are important for lensing will tend to be denser than the overall population.

In summary, the substructure model has the free parameters α , β , M_{\max} , M_{\min} , R_{\max} , η_* and the normalization host velocity dispersion σ_* which we fix at 200 km s^{-1} . However, in the simulations described in the following, α and M_{\max} are fixed and the remaining parameters are varied.

2.3 Ray-shooting

The sources that we wish to use in our simulation have sizes of $\sim 10 \text{ pc}$ and the substructures can have similar sizes. Therefore, it is essential that we be able to calculate the magnification of finite size sources. This requirement has been widely ignored in the literature because it is difficult to map the image of a finite source in a short enough amount of time to make it possible to create the large number of simulated lenses required for this problem. A new code, GLAMER, has been developed for this and other applications. This code employs a highly optimized adaptive mesh refinement scheme which allows the shapes of the images and their area to be calculated rapidly. (Because of surface brightness conservation, the area of a uniform brightness image is proportional to its magnification.) This allows us to make millions of mock lenses with a finite size source in a relatively short amount of time. Fig. 1 illustrates how the grid




Figure 1. An example of the refined grid for one particular lens and source position. The refinements continue below the resolution of this plot. The deflection angle is calculated once at the centre of each grid cell. There are four images of a 10-pc source in this case. At a higher resolution than is visible here, the lower right image breaks into two.

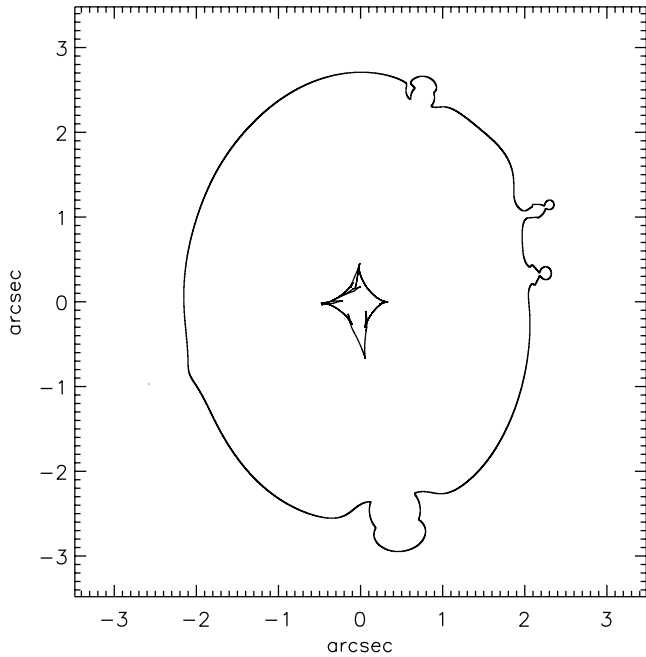


Figure 2. The critical curve (outer curve) and the caustic for the same lens as in Fig. 1. The substructure mass range is 10^7 – $10^9 M_{\odot}$ with a number density of $\eta_* = 0.5 \text{ kpc}^{-2}$ and a size scale of $R_{\text{max}} = 0.5 \text{ kpc}$. In this case, $\sigma = 214 \text{ km s}^{-1}$, $z_{\text{source}} = 1.34$, $z_{\text{lens}} = 0.41$ and $f = 0.8$.

is refined to find all the images and their areas. Fig. 2 shows the critical curve and caustic structure for one example lens. For more details on this code, see Metcalf (in preparation).

The range of positions in which a substructure will make a significant change to the magnification of an image depends on the mass of the substructure. To optimize calculations, small-mass substructures that are far away from the lens are omitted from the calculation while more massive substructures further from the lens are included. To accomplish this, a mass-dependent cut-off radius from the centre of the lens is used:

$$r_{\text{max}}(\sigma, m) = 2r_{\text{E}}(\sigma) + R_{\text{cut}}(m) + \left(\frac{2m r_{\text{E}}(\sigma)}{\pi \Sigma_{\text{crit}} \epsilon_{\text{min}}} \right)^{1/3}. \quad (13)$$

The first two terms ensure that all substructures within two Einstein radii plus the radius of the substructure are included. The third term ensures that any substructure close enough to cause a perturbation to the lens that is not well approximated as a pure shear will be included. The parameter ϵ_{min} controls how large the variation in the shear across the Einstein radius is allowed to be. We set this parameter to $\epsilon_{\text{min}} = 10^{-3}$. The contribution from substructures or companions outside this range is considered to be part of the background shear discussed in Section 2.1 as part of the host lens model.

For each lens model (host and substructure), the critical curves and caustics are found first. There are sometimes multiple, disconnected critical curves. The main tangential caustic is found by requiring its critical curve to be the one that encompasses the most area while also surrounding the centre of the lens. The area within the tangential caustic is calculated and the number of source positions that will be used for that lens is calculated as described in Section 2.1. The sources are required to have their centres inside the tangential caustic, but they are otherwise randomly distributed. Because of the finite source size, some images will be merged and this results in less than four images.

Some lenses have more than the four images that the undistorted host model alone would predict. Some of these additional images are very small and/or so close to another image that they would not be observed as separate images. We do a rough initial cut in all cases by merging together any images with centroids that are less than 0.1 arcsec apart, roughly the resolution of the *Hubble Space Telescope*. Further discussion of additional images is given in the next section.

Table 1 lists the simulation runs that were performed. They are in batches of 100 000 lenses with fixed substructure parameters. The first five sets of simulations have no substructure in them and are used to evaluate the importance of distortions to the host lens model and establish a baseline from which to measure the importance of substructure. Then the parameters for the remaining 12 simulations were chosen to explore the importance of particular substructure properties for lensing. Set 2 is taken to be a fiducial model. This is a somewhat arbitrary choice, but we do believe that it is similar to the predictions of N -body simulations except for the internal profile of the substructures which, as will be shown, has relatively little effect on the lensing properties. Relative to simulation set 2, set 1 has a higher minimum mass (and average mass), set 3 has a lower minimum mass, set 4 has a smaller source size, set 5 has more compact substructure (a smaller R_{max}), sets 6 and 8 have less compact substructure and set 7 has a shallower internal mass profile for the substructures. In sets 9 and 10, the upper mass cut-off is increased to $10^{10} M_{\odot}$, which is about 10 per cent of the host’s virial mass. Set 9 has more compact substructures than set 10. The R_{max} values are set here so that the size–mass relation is the same as in sets 7 and 8. For example, a $10^8 M_{\odot}$ substructure has the same size in sets 9 and 7. The rescaling is necessary because the size–mass relation is normalized at the maximum mass in each model which changes between these models. In sets 11 and 12, the upper mass cut-off is decreased to $10^8 M_{\odot}$. Set 11 has more compact substructures than set 12. Again, the R_{max} values are set to preserve the mass–size relation between sets 7 and 11, and between sets 8 and 12.

The range in surface number density, η_* , in the simulation sets is meant to span the credible range within a CDM-like model (Diemand et al. 2007b; Springel et al. 2008). In set 3 the number density of substructures is much higher for the same mass density, so because of computer time constraints the mass density range for this set does not go as high as in the others although the number density goes higher.

3 RESULTS

We create several million simulated lenses and save the image positions and magnifications. We also store the point source magnifications at the centroid of each image and the point source magnification for the point in the image that is closest to the centre of the source. Some of the host lens parameters are also stored. In this paper, for ease of comparison, we classify the observed and simulated lenses and reduce the position and magnification information to two parameters. The parameter $\Delta\theta$ is defined in Fig. 3. A small value of $\Delta\theta$ indicates the source is near a cusp in the caustic. Fig. 3 also describes what long- and short-axis lenses are. We have found that a good observational way of sorting the lenses into these categories is by comparing the angular distance between the centre of the lens and the singlet image to the distance between the centre of the lens and the central image of the triplet. If the former is greater, then the lens is a short-axis lens. Otherwise, it is a long-axis lens.

Table 1. Simulation runs. The top section shows models without substructure, where the source size, distortion and level of external shear are varied. The bottom rows show simulation runs with substructure. $\alpha = 1.9$ in all cases.

Set	Host model	M_{\max} (M_{\odot})	M_{\min} (M_{\odot})	R_{\max} (kpc)	β	η_* (kpc^{-2})	R_{source} (pc)	Number of simulations
	Standard	–	–	–	–	0	10	100 000
	Standard	–	–	–	–	0	1	100 000
	Extra distorted	–	–	–	–	0	10	100 000
	Extra shear	–	–	–	–	0	10	100 000
	No distortion or shear	–	–	–	–	0	10	100 000
1	Standard	10^9	10^8	0.5	1	0.013–0.13	10	10^5 per $\eta_* = 1.2 \times 10^6$
2	Standard	10^9	10^7	0.5	1	0.013–0.40	10	10^5 per $\eta_* = 3.0 \times 10^6$
3	Standard	10^9	10^6	0.5	1	0.013–0.60	10	10^5 per $\eta_* = 4.5 \times 10^6$
4	Standard	10^9	10^7	0.5	1	0.013–0.40	1	10^5 per $\eta_* = 2.0 \times 10^6$
5	Standard	10^9	10^7	0.25	1	0.013–0.40	10	10^5 per $\eta_* = 3.0 \times 10^6$
6	Standard	10^9	10^7	1.0	1	0.013–0.40	10	10^5 per $\eta_* = 3.0 \times 10^6$
7	Standard	10^9	10^7	0.5	0.5	0.013–0.40	10	10^5 per $\eta_* = 3.0 \times 10^6$
8	Standard	10^9	10^7	4.0	1	0.013–0.40	10	10^5 per $\eta_* = 3.0 \times 10^6$
9	Standard	10^{10}	10^7	1.1	1	0.013–0.41	10	10^5 per $\eta_* = 3.0 \times 10^6$
10	Standard	10^{10}	10^7	8.6	1	0.013–0.41	10	10^5 per $\eta_* = 3.0 \times 10^6$
11	Standard	10^8	10^7	0.23	1	0.013–0.49	10	10^5 per $\eta_* = 3.0 \times 10^6$
12	Standard	10^8	10^7	1.8	1	0.013–0.49	10	10^5 per $\eta_* = 3.0 \times 10^6$

The second parameter used to characterize each lens is

$$R_{\text{cusp}} \equiv \pm \frac{\mu_1 - \mu_2 + \mu_3}{\mu_1 + \mu_2 + \mu_3}, \quad (14)$$

where ‘+’ is for long-axis lenses and ‘–’ for short-axis lenses. The magnifications for the images in the triplet are μ_1 , μ_2 and μ_3 , with μ_2 being for the central image. The original motivation for this parameter was that $R_{\text{cusp}} \rightarrow 0$ asymptotically as a point source approaches a cusp in the caustic (Schneider & Weiss 1992). The R_{cusp} parameter has been widely used because of this model-independent prediction. In practice, R_{cusp} is not constrained to a very small region around zero because of finite source effects and the invalidity of the lowest order expansion of the lensing equation around the cusp. And, as will be shown, the distribution of R_{cusp} is not very model-independent.

Figs 4 and 5 show the distribution of R_{cusp} and $\Delta\theta$ for the sample of simulations listed in the captions. It can be seen that the simulated lenses occupy a well-localized regions in these diagrams when no substructure is present. Even when substructure is present at the levels investigated, the majority of lenses occupy the same regions with a smaller number of cases spread out in tails to the distribution.

Fig. 4 shows how important the ellipticity of the host lens is to the distribution of R_{cusp} values. Distortions to the SIE model and background shear do broaden the distribution, but ellipticity has a particularly strong effect. If only low-ellipticity lenses are considered, the R_{cusp} values are restricted to a much narrower band. The sample of lenses is biased towards high ellipticities relative to the general population of lenses because the cross-section for producing four images (the area within the tangential caustic) increases with increasing ellipticity. At the same time, N -body simulations might be biased towards low ellipticity since generally only dynamically well-relaxed systems are chosen for very high resolution simulations. This can explain some of the discrepancies between simulations and observations that have been reported (Macciò et al. 2006; Xu et al. 2009; Amara et al. 2006). This will be further discussed in Section 5.

Fig. 5 is similar to Fig. 4, but the effect of substructure on the $\Delta\theta$ – R_{cusp} distribution is illustrated. An additional 10 per cent error on each image’s flux is added to conservatively account for typical

observational uncertainties. Substructure has the effect of producing a population of extreme outliers in this distribution.

Fig. 6 shows the fractional error made in the magnifications when the point source magnification is used. It can be seen there that the fractional error is small for magnifications less than around five. This is a confirmation that the numerical errors made by the ray-tracing code are small. At higher magnifications, larger errors are made when the source is 10 pc. This is not a numerical effect. It can also be seen in Fig. 6 that substructure causes the errors made by using the point magnification to increase when the source size is 10 pc, but less so when the source size is 1 pc. This is in agreement with expectations because the source size of 10 pc is closer to the characteristic scale of the substructures.

Fig. 7 shows the ratio between the point source magnifications and the finite source magnifications. Again, it can be seen that numerical errors are not playing a large part. It is evident that the point source magnifications are not evenly distributed around the finite source magnifications. Centroid point source magnifications tend to overestimate the real magnification; in some cases by a large factor. This is the magnification that would be calculated when fitting a lens model to an observed lens. In the simulation, the centroid is calculated by doing a flux-weighted average over the pixels on the simulation grid. The nearest point magnification is much less biased and in the opposite direction; the magnification is underestimated. In other lensing simulations, the source position is often fixed and the images are found by an iterative minimization algorithm. This would give essentially the same result as our nearest point magnification. Both effects are much smaller for a smaller source size, as they should be.

Many images were merged because their centroids were within 0.1 arcsec. In these cases, it makes no sense to take the closest point magnification since the closest point is not unique. Unsurprisingly, the magnification at the centroid point is an even worse approximation in these cases, as can be seen in Figs 6 and 7. In exceptional cases, the centroid might not even be in one of the images that are merged. As expected, these cases only arise when substructure is present.

Figs 6 and 7 should give one pause before using the point source approximation for the magnification in any substructure lensing

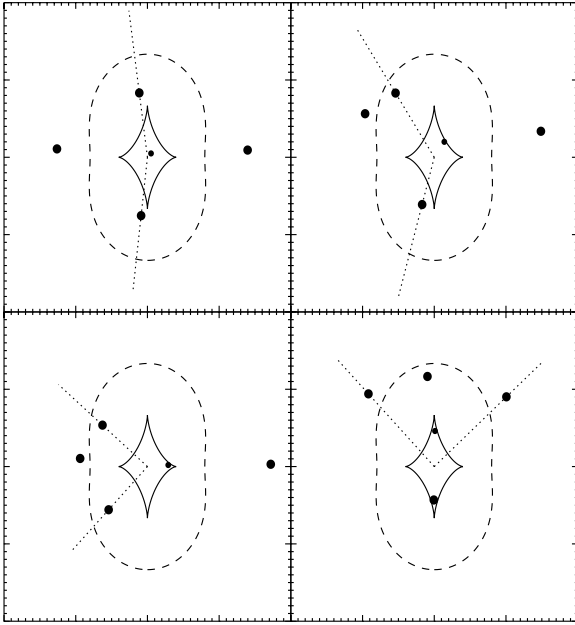


Figure 3. These diagrams represent the categorization of four image QSO lenses. The large dots represent the images while the small dot in each panel is the position of the source. The dashed curve is the critical curve (curve along which the magnification diverges) and the solid curve is the caustic curve (the curve on the source plane that bounds the region in which a source has four images). The four panels correspond to the four types of lenses. They are (clockwise from the upper left) an Einstein cross, a fold caustic, a short-axis cusp caustic and a long-axis cusp caustic. Generally, when the source is near one of the cusps in the caustic, three of the images will be close together. When the source is near the caustic but not near a cusp, two of the images will be close together. We define the *angular separation* between images as the smallest angle between the lines passing through those images and the centre of the lens. The image with the two smallest angular separations to other images is the central image of the *image triplet* which includes its neighbours. It is possible that the triplet is not well defined, but this very seldom happens in practice. The *singlet* image is the remaining image. The triplet’s opening angle, $\Delta\theta$, is the angle between the dotted lines shown in each case. When $\Delta\theta$ is small, the lens is ‘cuspy’. The categorization of observed lenses into long axis and short axis can be made by comparing the distance from the centre of the lens to the singlet image, to the distance from the centre of the lens to the central image of the triplet. If the former is larger it is a short-axis case, and if the latter is larger it is a long-axis case. In our simulations, this proves to be a very good discriminator.

study or when interpreting the results of any studies that use this approximation.

3.1 Frequency of $\Delta\theta$ – R_{cusp} outliers

To determine how often it would be expected for a lens to have $\Delta\theta$ and R_{cusp} values that are inconsistent with a smooth lens model, we define a region around the distribution in the case where no substructure is present and find how many simulated lenses lie outside this region when substructure is added. We define this region by taking bins in $\Delta\theta$ that contain 2000 simulations taking the long-axis and short-axis cases separately. Upper and lower boundaries within each bin are set such that 2.5 per cent of the simulations in the bin are greater than the upper bound and an equal number are less than the lower bound. The bins completely cover the full possible range of $\Delta\theta$. Without substructure, 5 per cent of a lens lie outside of this region. The fraction of simulated lenses outside

this region when substructure is added will be called the fraction of outliers.

Fig. 8 shows the fraction of outliers as a function of the substructure surface number density, η_* , for different substructure minimum masses (simulation sets 1, 2 and 3). A significant fraction of the lenses are found to be outliers. The top panels of Figs 9–11 show the same outlier fraction, but as a function of surface mass density.

It is surprising that in Fig. 9 the outlier fraction appears dependent only on the total surface mass density and not on the lower mass cut-off. One might think that all the lensing is being done by the most massive substructures and this is why the lower mass cut-off is not important in these cases. This does not seem to be the case; from set 1 ($M_{\text{min}} = 10^8 M_{\odot}$) to set 3 ($M_{\text{min}} = 10^6 M_{\odot}$) the mass density in the highest decade of mass (10^8 – $10^9 M_{\odot}$) drops by 60 per cent for the same total surface mass density and yet the number of outliers is unchanged.

Fig. 10 shows the importance of compactness and internal structure on the number of outliers. The substructure mass function is the same for all the models in this figure. The slope of the internal density profile, β , seems to have very little effect on the outlier fraction. On the other hand, the size of the substructures, or their compactness, does have a strong influence of the outlier fraction. Between $R_{\text{max}} = 0.5$ and 4.0 kpc, the fraction decreases significantly. Since the size–mass relation of the substructures is related to their galactocentric distance through tidal stripping, this sensitivity would provide information on where the substructures are within the lens halo or outside of it.

In Fig. 11, the upper substructure mass limit is changed to investigate further the insensitivity to mass range. It is seen again that for the same mass–size relation the fraction of outliers is dependent on the total surface mass density and relatively insensitive to the upper mass cut-off. The sensitivity to substructure compactness is again clearly present. Set 9 with $M_{\text{max}} = 10^{10} M_{\odot}$ appears to produce slightly less outliers than set 2 with $M_{\text{max}} = 10^9 M_{\odot}$. This could be because large substructures will sometimes displace the image positions and magnifications significantly while preserving a low R_{cusp} value; the cusp in the caustic is moved, but its shape remains relatively intact.

From the upper panels of Figs 9–11, it can be seen that if the size–mass relation is held fixed the outlier fraction is largely a function of the total surface mass density in substructures and not the range of substructure masses. This conclusion may depend on the function used here ($\alpha = 1.9$). Further simulations will be needed to investigate this. Changing the size–mass relation so that the substructures are less dense does reduce the fraction of anomalies (sets 6 and 8).

4 COMPARISON WITH DATA

To avoid contamination from microlensing by stars in the lens galaxy, differential extinction and variability of the source on time-scales smaller than the image delay times, we compare our simulations only to quad lenses measured in the radio and the mid-IR. Since we have not included companion galaxies to the primary lens in our simulations, we also remove lenses with nearby galaxies that appear to have similar masses to the primary. This removes 1608+656 and 1004+4112 from the list. There is a very faint dwarf galaxy within the Einstein radius of 2045+265 (McKean et al. 2007), but we will consider this to be a substructure and not a companion galaxy because it is small. Lens models show that this substructure would need to be unnaturally elongated to cause the flux anomaly




Figure 4. The distributions of $\Delta\theta$ and R_{cusp} that show the importance of distortion to the elliptical lens model and the ellipticity distribution of the lenses. The blue regions show histograms. The left column shows the long-axis lenses, and the right column shows the short-axis lenses. The top row is for simulation set ‘standard’ with $R_{\text{source}} = 10$ pc which has random distortions and the full range of ellipticities. The second row is for the simulation set ‘no distortions or shear’ which are pure elliptical models with full range of ellipticities. The third row is the same as the second, but with all the models with axial ratio $f < 0.7$ removed. The fourth row is the same as the second, but with the same axial ratio cut. The radio and IR observations are shown as red stars. It can be clearly seen that the distribution of R_{cusp} is highly dependent on the distribution of lens ellipticities and that most of the observed R_{cusp} values are not exceptionally high if the full range of ellipticities is considered. The horizontal lines show where 95 per cent of the cases are above and 95 per cent of the cases are below in bins of 3000 simulations. The observed lenses are shown in red and discussed in Section 4.

in this system, so there is probably another substructure present. The lenses must also have a detected lens galaxy which eliminates 0134–0931 and 0128+437. Table 2 lists the lenses used and their R_{cusp} and $\Delta\theta$ values are plotted in Figs 4 and 5.

The most striking thing in Figs 4 and 5 is that one of the lenses, 2045+265, has significantly higher R_{cusp} than is expected in the absence of substructure, but that all the other lenses have $\Delta\theta$ – R_{cusp} values that are not particularly anomalous. In the bottom two rows of Fig. 4, it can be seen that if only the low ellipticity lenses (axial ratio > 0.7) were considered three or four of the observed lenses would have anomalous $\Delta\theta$ – R_{cusp} values. Since the authors that have compared N -body simulations to the data using R_{cusp} values in the past have used very few simulated lenses and all with axial ratios ≥ 0.7 (Amara et al. 2006; Macciò et al. 2006; Xu et al. 2009), it is now not surprising that they concluded that the simulations did not produce enough anomalies.

It should be emphasized that just because the lenses’ $\Delta\theta$ – R_{cusp} values are not anomalous does not mean that they do not have anomalous flux ratios. Some of these cases clearly cannot be fitted by reasonable models without substructure when all the image positions and fluxes are taken into account (Metcalf & Zhao 2002; Evans & Witt 2003; Shin & Evans 2008). With so few observed lenses and only one clear anomaly in $\Delta\theta$ – R_{cusp} space, it is impossible to make any strong conclusion about the allowed properties for substructure using only the $\Delta\theta$ – R_{cusp} distribution. About one anomaly out of

the seven lenses is about what one would expect from studying the top panels of Figs 9 and 10 for a substructure surface density of $\sim 10^7 M_{\odot} \text{ kpc}^{-2}$. Other flux-based constraints are possible and will be investigated in future papers.

We introduce another constraint in the bottom panels of Figs 9–11 based on the fraction of simulations with more than four images. (This does not include the central demagnified image that forms near the centre of the lens for non-singular lens mass profiles. In our case, the mass density in the centre of the lens diverges like $\Sigma \propto r^{-1}$, and this image never appears; it is infinitely demagnified.) Even after merging images with centroids less than 0.1 arcsec apart, there are cases where the substructures cause further splitting of the images. Of the 32 QSO lenses in the CASTLES (Kochanek et al. 2000) list of lenses with more than four images and simple lenses, none has more than four images of a single source separated by more than 0.1 arcsec.² This puts a strong constraint on the allowed fraction of lenses that have more than four images, $f_{>4}$.

² 0134–0931 might have five optical images, but two of them are well within 0.1 arcsec of each other. 1933+503 has 10 radio images, but models show that the best explanation is that there are three sources with none of them imaged more than four times (Nair 1998). 1359+154 does appear to be an honest-to-goodness case of a single QSO with six images, but the lens is a group of three galaxies and thus does not pass our no companions cut.




Figure 5. The distribution of $\Delta\theta$ and R_{cusp} for four of the simulations in set 2. The blue regions show histograms. The left column shows the long-axis lenses, and the right column shows the short-axis lenses. The number densities of substructures in each row from top to bottom are $\eta_* = 0, 0.09, 0.16$ and 0.27 kpc^{-2} . Random noise of 10 per cent has been added to represent observational errors. The horizontal lines show where 95 per cent of the cases are above and 95 per cent of the cases are below in bins of 3000 simulations. These are not exactly the same bins as that are used in calculating the outliers discussed in Section 3.1, but are similar. The radio and IR data are shown as red stars.

The probability of getting zero cases of >4 images in 33, given that the probability of getting such a case is $p \simeq f_{>4}$, is a binomial distribution. There would be less than a 5 per cent chance of this happening in the observed sample if $f_{>4}$ is greater than 0.089 and less than 10 per cent chance if $f_{>4} > 0.069$. These are the dotted lines in the bottom panels of Figs 9–11. For equal surface mass density, more massive substructures cause more high image multiplicity lenses.

The multiplicity constraint does change significantly if the resolution cut-off of 0.1 arcsec is changed. There are a large number of lenses where the images are merged in some cases (up to ~ 20 per cent). With improved resolution or a more careful analysis of the data, we believe this constraint could be made significantly stronger.

Within the ranges of η_* studied here, the only models that are limited by this image multiplicity constraint are set 1 (high lower mass cut-off and compact), set 5 (supercompact) and set 9 (high upper mass cut-off and compact). The constraints are $\Sigma_* < 2.0 \times 10^7 \text{ M}_\odot \text{ kpc}^{-2}$, $\Sigma_* < 1.2 \times 10^7 \text{ M}_\odot \text{ kpc}^{-2}$ and $\Sigma_* < 1.2 \times 10^7 \text{ M}_\odot \text{ kpc}^{-2}$, respectively. The more compact and massive the substructures are the more high multiplicity cases are created. This constraint is in contrast to the R_{cusp} constraint which depends only on the mass density and compactness. With more lenses this constraint could become significantly stronger in the future.

5 EXPECTATIONS FOR SMALL-SCALE STRUCTURE WITHIN THE CDM MODEL

A good point of comparison between lens simulations and N -body simulations is the fraction of mass in substructure within a projected radius of 10 kpc. This is easily measured in the simulations and since the Einstein radius is typically around 10 kpc, it is close to what is actually constrained by the lensing data. In our model, this quantity is given by

$$f_{\text{sub}}^{10\text{kpc}} \equiv \frac{M_{\text{sub}}(R < 10 \text{ kpc})}{M_{\text{host}}(R < 10 \text{ kpc})} = \frac{G\langle m \rangle \eta_*}{\sigma_*^2} 10 \text{ kpc}, \quad (15)$$

$$= 1.08 \times 10^{-9} \text{ kpc}^2 \Sigma_*, \quad (16)$$

where the fiducial value $\sigma_* = 200 \text{ km s}^{-1}$ has been used. Note that this fraction scales with host mass in our model and in the simulations.

Diemand et al. (2007a) give $f_{\text{sub}}^{10\text{kpc}} \simeq 0.003$ for the Via Lactea simulation, and Xu et al. (2009) give $f_{\text{sub}}^{10\text{kpc}} \simeq 0.0025$ with a large scatter in the Aquarius simulations. These simulations should be resolving substructure to below 10^7 M_\odot . These translate to $\Sigma_* = 2.8 \times 10^6$ and $2.3 \times 10^6 \text{ M}_\odot \text{ kpc}^{-2}$, respectively. Accounting for the extra mass below their resolution and judging from Figs 9–11 we would expect about a 10 per cent chance of a clear outlier in the $\Delta\theta$ – R_{cusp} distribution for the high compactness cases which is




Figure 6. The range in the fractional error, $(\mu_{\text{point}} - \mu_{\text{ext}})/\mu_{\text{ext}}$, made by using the point source magnification instead of a finite size source. 90 per cent of the simulations, in running bins of 5000, fall below these curves. The top panel is for a source with a radius of 10 pc and the bottom panel is for a radius of 1 pc. The dashed curves are for no substructure ($\eta_* = 0$) and the solid curves are for $\eta_* = 0.2 \text{ kpc}^{-2}$, all from simulation set 2. All the images of all the four image systems are used. The black curves are for the point magnification calculated at the centroid of the image not including the images that were merged by the 0.1 arcsec merger requirement. The blue curves are the same but including the merged cases. The red curves are for the point source magnification calculated at the grid point in the image that is closest to the centre of the source. The errors for the small source are typically at the 1 per cent level over a wide range of magnifications, when compared to nearest point estimates, showing good convergence on the level of the numerical noise from the GLAMER ray-tracing code.

consistent with the one out of seven observed. For the larger size-mass relation (sets 8, 10 and 12), the expected fraction is increased by only a few per cent from the no substructure case, but with only one observed outlier, we do not consider this a significant contradiction.

Amara et al. (2006), Macciò et al. (2006) and Xu et al. (2009) come to the conclusion that the substructure present in the simulations is not enough to cause the observed frequency of R_{cusp} anomalies. In light of the findings in this paper we believe that these conclusions were flawed because the full range of host lens ellipticities was not represented in the simulations. Macciò & Miranda (2006) may have used too low a substructure mass range (10^5 – $10^7 M_{\odot}$) to cause enough anomalies.

There are a number of other complicating factors that make comparing observations to the true predictions of CDM difficult. For example, the baryons are not accounted for in the N -body simulations. This impacts the predictions in several ways. First, the host galaxy needs to be inserted by hand into these N -body simulations for them to be realistic lenses. The mass fraction decreases with the inclusion of baryons. Secondly, the baryons are expected to have some effect on the internal structure of the substructures, either expanding or contracting them, which will affect their tidal stripping and disruption in the host halo. The resident galaxy might also have a significant effect on the survival of substructures. As discussed in Section 2.1, the typical galactocentric distance for substructures that are important for lensing is significantly smaller than the typical distance of substructures in general. The substructure




Figure 7. The ratio of the point source magnification to the finite size magnification. The shaded regions show where 90 per cent of the simulations in running bins of 5000 are, 5 per cent above and 5 per cent below. The simulations and colour scheme are the same as in Fig. 6. As in Fig. 6, the upper panel is for a 10-pc source and the lower is for a 1-pc source. We see that, depending on the method used, magnification estimates using point sources can lead to both random errors and biases (seen as an asymmetry of the shaded region) as compared to the extended source calculation, which is closer to the observables.




Figure 8. The fraction of lenses that lie outside the region in $\Delta\theta$ – R_{cusp} space that contains 95 per cent of the lenses when there is no substructure (see text for details) as a function of substructure number density. In all cases $M_{\text{max}} = 10^9 M_{\odot}$. The curves are for $M_{\text{min}} = 10^8 M_{\odot}$ (red), $M_{\text{min}} = 10^7 M_{\odot}$ (green) and $M_{\text{min}} = 10^6 M_{\odot}$ (blue). These correspond to sets 1, 2 and 3 respectively from Table 1. There are 100 000 simulated lenses used in calculating each point.

population probed by lensing is likely to be more compact and have a steeper mass function, at least above $\sim 10^8 M_{\odot}$, than the general population. This steepening of the mass function at high masses has been only crudely accounted for in our model by the M_{max} cut-off parameter.

Because we appear to be consistent with the simulations on the frequency of R_{cusp} anomalies does not mean that some other test, such as fitting each simulated lens to a smooth lens model, would not show some inconsistency. Modelling the lens puts constraints




Figure 9. The top panel is the same as in Fig. 8, except it is now as a function of the surface mass density in substructures. The bottom panel shows the fraction of lenses with more than four images with separations of more than 0.1 arcsec and flux ratios of within a factor of 100 (excluding the cases with less than four images). The colours are the same as in Fig. 8. The dotted lines in the bottom panel show where there is only a 10 and 5 per cent chance of a sample of 32 lenses having no cases of more than four images as in the CASTLES lens sample.




Figure 10. Same as Fig. 9, but with substructure models with different internal structures. Green is set 1 (see Table 1) as in Fig. 9. Red is for denser substructures (set 5), while blue and purple are for less dense substructures (sets 6 and 8, respectively). Cyan is for substructures with less steep mass profiles, but the same sizes as green (set 7). Changing the internal mass profile has no discernible effect on the frequency of flux anomalies, but has a significant effect on the frequency of high multiplicity lenses.




Figure 11. Same as Fig. 9, but with substructure models meant to explore the importance of the upper mass limit. Green is set 1 (see Table 1) as in Figs 9 and 10. The purple and brown (sets 9 and 10, respectively) have a higher upper mass cut-off of $M_{\max} = 10^{10} M_{\odot}$ and different mass–size relations. The red and blue curves (sets 11 and 12) are for a mass cut-off of $M_{\max} = 10^8 M_{\odot}$. The R_{\max} values are set so that a substructure of the same mass will have the same size in sets 1, 9 and 11. The same is true for sets 10 and 12. The compactness clearly has a strong effect. The fraction of outliers for the green, red and purple are similar indicating that the upper mass cut-off does not have a strong effect on the number of outliers when the substructures are compact. In contrast, the number of high multiplicity lenses clearly is dependent on the mass range.

on the ellipticity. Our argument is that R_{cusp} is not a good test for the existence of substructure without further constraints.

6 CONCLUSIONS AND DISCUSSION

We have performed the largest number of lens simulations ever done with finite size sources. This was made possible by the new adaptive ray-tracing code GLAMER. We find that accounting for the finite size of the source is necessary for drawing accurate conclusions from the lensed QSO data.

We find rough consistency between the Λ CDM predictions and observations. R_{cusp} is found to be a poor discriminator between lenses with substructure and without because of its sensitivity to the ellipticity of the lens. The distribution of ellipticities used in our lens models is based on the ellipticities of observed lenses, so we do not think the ellipticities required to explain the observed R_{cusp} distribution (excepting lens 2045+265) are atypical. Other methods for comparing observations to models are likely to be more fruitful, and as the data improves more precise comparisons will be possible. In addition to the substructure within the primary lens, there should be some contribution from intergalactic small-scale structure (Metcalf 2005a,b) so one should expect the limits derived from the data to be somewhat higher than the limits derived from N -body simulations of individual dark matter haloes. The baryons also clearly play a role in shaping the lensing properties and they

Table 2. Observed lenses used in this analysis.

Name	Band	$\Delta\theta$	R_{cusp}	Reference
2045+265	Radio	35:3	0.501	Fassnacht et al. (1999)
0712+472	Radio	79:8	0.254	Jackson et al. (1998)
1555+375	Radio	108:0	0.417	Marlow et al. (1999)
1422+231	Mid-IR	74:9	0.203	Chiba et al. (2005)
0414+053	Radio	101:5	0.220	Katz, Moore & Hewitt (1997)
2237+030	Radio	146:3	0.357	Falco et al. (1996)
1115+080	Mid-IR	127:5	-0.043	Chiba et al. (2005)

are not fully taken into account in the simulations at the necessary resolution.

We have limited our study here to a substructure mass function of the form $dN/dm \propto m^{-\alpha}$ with $\alpha = 1.9$. This seems well motivated by the simulations on small mass scales, but could be steeper on larger mass scales because of tidal stripping and disruption in the central regions of the lens. With the $\alpha = 1.9$ mass function, the smaller mass substructures play a smaller part in causing flux anomalies because most of the mass resides in larger mass objects. This will make it difficult to measure any possible lower mass cut-off using monochromatic QSO lensing alone. Fortunately there are some other prospects for probing the mass function in the future such as spectroscopic gravitation (Moustakas & Metcalf 2003) and Einstein rings (Vegetti & Koopmans 2009). If the slope of the mass function is steeper than $\alpha = 1.9$, the smaller structures will play a larger role in the lensing.

It is clear that what is really required to make a more conclusive measurement of the amount of substructure in dark matter haloes is more data. With seven lenses, only limited conclusions can be made from a statistical point of view. We are also vulnerable to systematic errors. For the kind of study done here, more strong lenses measured in the radio and/or mid-IR are needed. Planned large-scale imaging surveys³ expect to increase the number of lensed QSOs in the visible by an order of magnitude, so we look forward to great improvements in this field.

ACKNOWLEDGMENTS

We would like to thank Simon White and Raul Angulo for useful discussions and comments. We would like to thank Jacqueline Chen for pointing an important error in an earlier version of this paper and the anonymous referee for very helpful suggestions and criticism.

RBM's research is part of the project GLENCO, funded under the Seventh Framework Programme, Ideas, Grant Agreement no. 259349.

REFERENCES

Amara A., Metcalf R. B., Cox T. J., Ostriker J. P., 2006, *MNRAS*, 367, 1367

³ DES (darkenergysurvey.org) (snap.lbl.gov), LSST (www.lsst.org), PanSTARS (pan-stars.ifa.hawaii.edu), VISTA (www.vista.ac.uk), EUCLID (http://www.ias.u-psud.fr/imEuclid/).

- Bender R., Doebereiner S., Moellenhoff C., 1988, *A&AS*, 74, 385
- Binney J., Tremaine S., 1987, *Galactic Dynamics*. Princeton Univ. Press, Princeton, NJ
- Bradač M., Schneider P., Lombardi M., Steinmetz M., Koopmans L. V. E., Navarro J. F., 2004, *A&A*, 423, 797
- Chiba M., Minezaki T., Kashikawa N., Kataza H., Inoue K. T., 2005, *ApJ*, 627, 53
- Churazov E. et al., 2010, *MNRAS*, 404, 1165
- Dalal N., Kochanek C. S., 2002, *ApJ*, 572, 25
- Dalal N., Watson C., 2005, in Goicoechea L. J., ed., 25 Years After the Discovery: Some Current Topics on Lensed QSOs. Preprint (astro-ph/0409483)
- Diemand J., Kuhlen M., Madau P., 2007a, *ApJ*, 657, 262
- Diemand J., Kuhlen M., Madau P., 2007b, *ApJ*, 667, 859
- Evans N. W., Witt H. J., 2003, *MNRAS*, 345, 1351
- Falco E. E., Lehar J., Perley R. A., Wambsgans J., Gorenstein M. V., 1996, *AJ*, 112, 897
- Fassnacht C. D. et al., 1999, *AJ*, 117, 658
- Fukazawa Y., Botoya-Nonesa J. G., Pu J., Ohto A., Kawano N., 2006, *ApJ*, 636, 698
- Gavazzi R., Treu T., Rhodes J. D., Koopmans L. V. E., Bolton A. S., Burles S., Massey R. J., Moustakas L. A., 2007, *ApJ*, 667, 176
- Humphrey P. J., Buote D. A., 2010, *MNRAS*, 403, 2143
- Jackson N. et al., 1998, *MNRAS*, 296, 483
- Katz C. A., Moore C. B., Hewitt J. N., 1997, *ApJ*, 475, 512
- Keeton C. R., Kochanek C. S., Seljak U., 1997, *ApJ*, 482, 604
- Klypin A., Kravtsov A. V., Valenzuela O., Prada F., 1999, *ApJ*, 522, 82
- Kochanek C., Falco E. C. I., Lehar J., McLeod B., Rix H.-W., 2000, CAS-TLES Survey (<http://cfa-www.harvard.edu/castles/>)
- Koopmans L. V. E., Treu T., Bolton A. S., Burles S., Moustakas L. A., 2006, *ApJ*, 649, 599
- Kormann R., Schneider P., Bartelmann M., 1994, *A&A*, 284, 285
- Kormendy J., Fisher D. B., Cornell M. E., Bender R., 2009, *ApJS*, 182, 216
- Macciò A. V., Miranda M., 2006, *MNRAS*, 368, 599
- Macciò A. V., Moore B., Stadel J., Diemand J., 2006, *MNRAS*, 366, 1529
- McKean J. P. et al., 2007, *MNRAS*, 378, 109
- Mao S., Schneider P., 1998, *MNRAS*, 295, 587
- Mao S., Jing Y., Ostriker J. P., Weller J., 2004, *ApJ*, 604, L5
- Marlow D. R. et al., 1999, *AJ*, 118, 654
- Metcalf R. B., 2005a, *ApJ*, 622, 72
- Metcalf R. B., 2005b, *ApJ*, 629, 673
- Metcalf R. B., Madau P., 2001, *ApJ*, 563, 9
- Metcalf R. B., Zhao H., 2002, *ApJ*, 567, L5
- Moore B., Ghigna S., Governato F., Lake G., Quinn T., Stadel J., Tozzi P., 1999, *ApJ*, 524, L19
- Moustakas L. A., Metcalf R. B., 2003, *MNRAS*, 339, 607
- Nair S., 1998, *MNRAS*, 301, 315
- Schneider P., Weiss A., 1992, *A&A*, 260, 1
- Shin E. M., Evans N. W., 2008, *MNRAS*, 385, 2107
- Springel V. et al., 2008, *MNRAS*, 391, 1685
- Thomas D., Maraston C., Bender R., Mendes de Oliveira C., 2005, *ApJ*, 621, 673
- Vegetti S., Koopmans L. V. E., 2009, *MNRAS*, 400, 1583
- Xu D. D. et al., 2009, *MNRAS*, 398, 1235
- Zackrisson E., Riehm T., 2010, *Advances Astron.*, 2010

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