BASICS OF GRAVITATIONAL LENSING

INTRODUCTION

Gravitational lensing is a consequence of one of the most famous predictions of Einstein's General Relativity—the idea that light is bent in a gravitational field. Indeed, the first calculation showing that gravitational bending of starlight could act as a lens was produced by Einstein himself, although he did somewhat pessimistically conclude that "there is no great chance of observing this phenomenon". The first gravitationally lensed quasar, Q0957+561, was discovered by Walsh *et al.* in 1979.

There are three main forms of gravitational lensing:

1. In *strong lensing*, the lens is a large mass, the geometry is favourable, and the deflection is comparatively large. The observer sees two or more separate images of the source.
2. In *weak lensing*, the lens is a large mass, but the geometry is less favourable. The image of the source is mildly distorted, with a tendency to smear into an arc centred on the lens centre: an effect known as *shear*. This means that the alignment of the background objects appears non-random, so shear can be measured statistically even if the distortions of individual objects are too small to be identified directly.
3. In *microlensing*, the lens is a small mass (usually a star), so that although the geometry is extremely favourable—source, lens and observer in a straight line—the deflection, distortion and multiple images caused by lensing cannot be resolved. Instead, the image of the source appears to brighten for the duration of the lensing event (since source, lens and observer all have relative proper motions, the alignment that creates the microlens is temporary).

Microlensing has been used to search for astrophysical dark objects—MACHOs—in the Galactic halo. Strong and weak lensing can be used to map the mass distributions of clusters of galaxies; over larger areas weak lensing can also map large-scale structure.

GEOMETRY OF GRAVITATIONAL LENSING

The basic geometry of gravitational lensing is shown in the figure below.

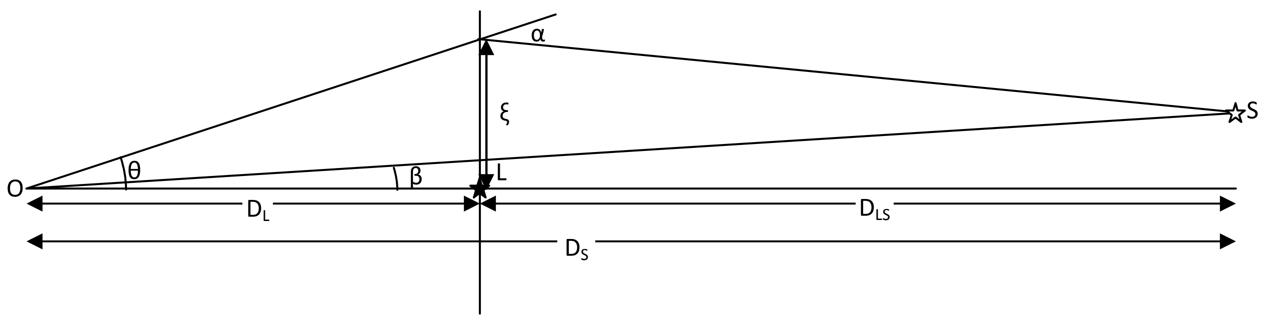

In this figure, the source S (white star) is lensed by L (whose centre is the black star) and observed at O. Note that for cosmological distances it is *not* in general true that $D_S = D_L + D_{LS}$, despite the appearance in the diagram! We assume that the lens L can be approximated by a single object at a well-defined distance D_L—the *thin lens approximation*—and that all the angles are small.

In the thin lens approximation, we assume that we can treat the paths of the light rays as straight lines, with all the deflection taking place at a single distance D_L. (In reality, the light travels along hyperbolae—we're essentially using the asymptotic straight lines of the hyperbola.)

Assuming that the lens has circular symmetry (not in general true!), we find from general relativity that the deflection angle α is

$$\alpha(\xi) = \frac{4GM(\xi)}{c^2}\frac{1}{\xi},$$

where $M(\xi)$ is the mass contained within radius ξ. In the small angle approximation

$$D_{\mathrm{LS}}\alpha = D_{\mathrm{S}}(\theta - \beta),$$

and

$$\xi = D_{\mathrm{L}}\theta.$$

In the perfectly symmetrical case where S lies directly behind L, so that $\beta = 0$, we therefore have

$$D_{\mathrm{S}}\theta_{\mathrm{E}} = D_{\mathrm{LS}}\frac{4GM}{c^2}\frac{1}{D_{\mathrm{L}}\theta_{\mathrm{E}}},$$

or

$$\theta_{\mathrm{E}} = \sqrt{\frac{4GM}{c^2}\frac{D_{\mathrm{LS}}}{D_{\mathrm{L}}D_{\mathrm{S}}}}.$$

In this situation, the light from the source galaxy is smeared into a circle of radius θ_E, known as an *Einstein ring*.

Comparing the expression for θ_E with the definition of α, we see that $D_{\mathrm{LS}}\alpha = D_{\mathrm{S}}\theta_{\mathrm{E}}^2/\theta$, and hence

$$\beta = \theta - \theta_{\mathrm{E}}^2/\theta.$$

This is a quadratic whose solution is

$$\theta_{1,2} = \frac{1}{2}\left(\beta \pm \sqrt{\beta^2 + 4\theta_{\mathrm{E}}^2}\right).$$

The implication of this is that if all lenses really were point masses, we would get an Einstein ring if $\beta = 0$, and exactly two images (one inside and one outside θ_E) if $\beta \neq 0$. In fact we often get more than two images, because lenses are not really point masses: they are frequently extended lumpy objects such as clusters of galaxies.

The quantity $D_{\mathrm{L}}D_{\mathrm{S}}/D_{\mathrm{LS}}$ has dimensions of distance, but is not simply related to any of the individual distances involved. Since it's calculated using angles, it is an angular diameter distance, and except in the low redshift regime ($z < 0.05$ or so), *angular diameter distances do not just add*. There is one "obvious" special case: if the lens is comparatively nearby, such that $D_{\mathrm{L}} \ll D_{\mathrm{S}}$, $D_{\mathrm{LS}} \simeq D_{\mathrm{S}}$ and we are sensitive mainly to D_L. Even this is not as straightforward as it looks, because angular size distance has a maximum value at $z \simeq 1$ (the exact value depends on the cosmological model; for the benchmark model it's about 1.64) and then decreases again as redshift increases—so, in fact, D_S and D_L may be very similar even though the redshifts of source and lens are very different.

In the low redshift regime, angular diameter distances are indistinguishable from any other kind of distance, and do add, so

$$\frac{D_{\mathrm{L}}D_{\mathrm{S}}}{D_{\mathrm{LS}}} \cong D_{\mathrm{L}}\left(\frac{D_{\mathrm{L}}}{D_{\mathrm{LS}}} + 1\right).$$

This would apply, for example, to microlensing of stars in the Large Magellanic Cloud by objects in the Galactic halo, or to lensing of stars in the Galactic bulge by disc stars (in which case we might further assume that on average $D_L \simeq D_{LS}$, reducing the above expression simply to $2D_L$).

MAGNIFICATION, MICROLENSING AND MACHOS

Microlensing occurs when the Einstein ring is too small to resolve: for example, if a star in the Large Magellanic Cloud, 50 kpc away, is lensed by a halo star of 0.5 solar masses 10 kpc away, the radius of the Einstein ring is 0.8 mas (milliarcseconds). In this case, the observable signature of the lensing event is not the distortion of the image, but its *magnification*.

The geometry of gravitational lensing preserves the surface brightness (flux per square arcsecond) of the source—but the distortion results in an image which covers more surface area than the original source did, so this results in an overall brightening. For a point lens, the amplification is given by

$$A(u_{1,2}) = \frac{u_{1,2}^4}{u_{1,2}^4 - 1},$$

where $u_{1,2} = \theta_{1,2}/\theta_E$, i.e. the angle between the image and the lens measured in units of the Einstein ring radius. There is always at least one image for which $A > 1$, i.e. at least one image which is brighter than the source. For small β, the amplification can be very large: gravitational lenses can act as natural 'telescopes' to study extremely distant galaxies in more detail then would be possible with unlensed systems.

The total amplification for both images combined is given by

$$A(u) = \frac{u^2 + 2}{u\sqrt{u^2 + 4}},$$

where $u = \beta/\theta_E$ (this can be derived by combining the amplifications for the two individual images with, for some reason that currently escapes me[1], a relative minus sign). For microlensing, where the individual images are not resolved, this is clearly the relevant parameter

$A(u)$ is always greater than 1, and tends to infinity as $u \to 0$ (in practice, since real lenses have finite size, the amplification is always finite, but can be very large). As $u \to \infty$, $A(u) \to 1$: a mass a large distance away has negligible effect on the image of the source, as we would expect.

If the geometry of a microlens were stable with time—as it might be for, say, a quasar lensed by a star or black hole in an intervening galaxy—it would be very difficult to identify: with no image distortion to provide a signature, how do you know that any lensing has taken place? However, *local* microlensing events are a different matter: the relative proper motion of the source and the lens means that the lensing geometry is *temporary*, and so the amplification is transient. From the diagram on the left, the value of u at time t is simply

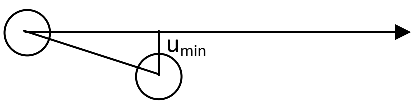

$$u(t) = \sqrt{u_{\min}^2 + \frac{(t - t_0)^2}{t_E^2}},$$

where t_0 is the time of closest approach and t_E is the time it takes the lens to move across an angular distance of θ_E relative to the source (this comes in because we're measuring u in units of θ_E). To take our earlier example of a halo MACHO and a star in the LMC, where the Einstein ring radius is 0.8 mas: if the MACHO moves at a relative velocity of say 200 km/s (reasonable for a halo object), at 10

[1] I think it's that one of the *u*-values is greater than 1, and the other less than 1, so to make *A* positive in both cases one has to have the sign reversed.

kpc distance its angular velocity is 6.5×10^{-17} rad/s = 0.012 mas/day. It will therefore take 69 days to cross 0.8 mas—the lensing event should take several months.

The duration of the lensing event depends on several factors:

- a larger lens mass increases $\theta_{\rm E}$ and therefore increases the event duration;
- a greater distance to the lens decreases $\theta_{\rm E}$, roughly as $1/\sqrt{D_{\rm L}}$, but also decreases the proper motion of the lens by a factor of $D_{\rm L}$—so the overall effect is to increase the duration;
- a smaller distance between lens and source decreases $\theta_{\rm E}$ and therefore decreases the duration;
- a faster relative velocity decreases the event duration.

Therefore, although the distribution of durations is used in MACHO analyses to estimate the typical lens mass, this result is dependent on assumptions about the dynamics and distribution of the lenses (for example, lensing of an LMC star *by another LMC star* can be mistaken for lensing by a lower mass object in the halo, because the very small value of $D_{\rm LS}$ causes the event duration to be shorter than expected for a stellar mass; likewise, if the lensing objects were to be distributed in a rotating disc instead of a non-rotating halo, the difference in expected transverse velocities would make a big difference to the inferred mass distribution). The most robust measurement is the *microlensing optical depth*, which is the probability that a random star is magnified above a certain threshold at any given time. This depends only on the density profile of the lenses, and not on their individual masses or velocities, as can be seen from the following arguments:

- The effective area of a lens is essentially $S \propto \theta_E^2 \propto M$, where M is the mean lens mass. The number of potential lenses in the line of sight to the source star is $N = nd = \rho d/M$, where n is the number density and d is the distance to the source. The probability of a lensing event must be $\propto NS$, so the mean mass M cancels.
- The event duration is inversely proportional to the mean lens velocity, but the area swept out by a lensing object in a time Δt is proportional to its velocity, so again the total number of lensing events observed in a time Δt is independent of velocity: effectively, the lens object will be in the vicinity of any given star for a shorter time, but it will potentially visit more stars.

To calculate the optical depth, we note that the linear radius of the Einstein ring is

$$\xi = D_L\theta_E = \sqrt{\frac{4GM}{c^2}\frac{D_{\rm L}D_{\rm LS}}{D_{\rm S}}} = \sqrt{\frac{4GM}{c^2}D_{\rm S}x(1-x)}$$

where $D_{\rm L} = xD_{\rm S}$. (Note that this is only valid in the low redshift regime, in which the distances behave like proper distances (so that $D_{\rm S} = D_{\rm L} + D_{\rm LS}$).)

A microlensing event will take place if the light from the source star passes through the circle of radius ξ around a halo object. The probability of this, i.e. the optical depth, is given by

$$\tau = \frac{4\pi G D_{\rm S}^2}{c^2}\int_0^1 \rho(x)x(1-x)\,{\rm d}x,$$

where $\rho(x)$ is the mass density at distance $xD_{\rm S}$. (As explained above, the mass density comes from combining the number density of lenses ($\propto \rho/M$) with the area inside the Einstein ring ($\propto M$).)

As an example, if we assume that the rotation curve of the Galaxy is flat at large *r* and that the dark halo is approximately spherical, we get

$$\rho(r) = \frac{1}{4\pi r^2}\frac{dM_r}{dr} = \frac{V^2}{4\pi r^2 G},$$

where V is the rotation speed (about 220 km/s). It is not really true that $r = xD_S$, because we are not located at the centre of the Galaxy, but if we use this as a rough approximation and assume that the integral over x has a numerical value of about 1, we find that

$$\tau \sim \frac{V^2}{c^2} = 5 \times 10^{-7}.$$

This gives an order of magnitude estimate for the microlensing optical depth through the halo.

IDENTIFYING MICROLENSING EVENTS

An observing programme capable of monitoring millions of stars on a daily basis for several years must be largely automated: typically, the same star fields are imaged repeatedly, and a computer program flags stars that have changed in brightness. Naturally, this is going to flag a lot of objects which are quite simply variable stars—the MACHO Collaboration produced a superb database of LMC Cepheids—plus other transient phenomena such as supernovae in background galaxies. The raw signal-to-noise ratio is minuscule (thousands of background events for every genuine lens), so selection criteria must be adopted to reject the spurious candidates. Fortunately, microlensing events have a number of characteristic features:

- They are one-off occurrences: since the probability of a given star's being lensed is about 10^{-7}, it's not going to happen twice. This can be used to reject periodic variables such as Cepheids, eclipsing binaries and so forth.
- They are *symmetrical*: the time distribution of $u(t)$ is identical for times before and after maximum light. Many astrophysical phenomena, such as supernovae, have asymmetric light curves, generally in the sense that they rise fast and decline slowly.
- They are *achromatic*: the lens effect has no dependence on wavelength, so the light curve will have the same form at all wavelengths. Again, this is not typical of intrinsic variability: most variable stars change their effective temperature, and therefore their colour, as they vary. For this reason, microlensing experiments take two images of each field, through different filters: any event that differs significantly between colours is rejected.
- The *shape* of the peak is well-defined: combine the expression for $A(u)$ and that for $u(t)$, and you have a formula for $A(t)$ which depends on $u_{\min}$, t_0 and t_E. A peak which cannot be fitted well using this formula is unlikely to be caused by microlensing[2].

The selection criteria for a microlens search can therefore be summarised as: no second peak; flat background outside peak region; no significant difference between colours; good fit to microlensing hypothesis for peak shape. There will also be technical cuts (e.g., are there enough points to define the peak and the baseline?), which will depend on the individual experiment.

MICROLENSING AND MACHOS

MACHOs (formally Massive Astrophysical Compact Halo Objects—but, I suspect, deliberately contrived to contrast with WIMPs!) is the generic name for dark baryonic objects of roughly stellar mass assumed to comprise some or all of the Galaxy's dark halo. Potential MACHO candidates include

- *brown dwarfs* of <0.1 solar mass ($M_\odot$);

[2] This is usually true, but in cases where the lens is not a single star, but a binary star or a star with planets, very odd-shaped light curves can be produced.

- *old white dwarfs* of perhaps 0.5 $M_\odot$;
- *neutron stars* of about 1.5 $M_\odot$;
- *black holes* of >2 $M_\odot$.

All of these would be faint enough to escape optical detection. *Red main sequence stars* of about 0.1–0.2 $M_\odot$ used to be included in this list, but would have been seen in the various deep-field images that have been taken over the last decade or so and can probably be excluded.

Given the same density profile and velocity distribution in each case, the ordering in mass is equivalent to ordering by event duration, since $\theta_E \propto \sqrt{M}$. Assuming that events are observed, their durations can therefore be used to estimate the typical lens mass, with the caveat that (as noted above) such estimates must assume a specific halo model.

RESULTS

Searches for MACHOs in the Galactic halo use the Magellanic Clouds (mostly the Large Magellanic Cloud, LMC) as a convenient collection of background stars to monitor for variation. The first results, reported by the MACHO Collaboration in 1997, suggested that a significant proportion of the halo mass was in the form of objects of order half a solar mass—presumably, since they are not visible, old white dwarfs (main-sequence stars of this mass would have absolute magnitudes of order +9 and would have been seen). This finding is difficult to understand astronomically, since we would expect that the production of stars massive enough to make *old* white dwarfs would imply lots of lower mass stars, which we should still be able to see as their lifetimes are much longer. However, the subsequent experiments EROS and OGLE have both reported much smaller numbers of candidate events, and hence a lower estimate of the optical depth: in fact, their small number of candidates is entirely consistent with *no* halo objects (they can be accounted for by LMC self-lensing, i.e. LMC stars lensed by other LMC stars, and lensing by disc stars). The results are shown in the table below.

Experiment	Stars	Elapsed time	Candidates	Optical depth	Mass range
MACHO[1]	11.9×10^6	5.7 years	13–17	$1.2^{+0.4}_{-0.3}\times10^{-7}$	$\sim0.5\ M_\odot$
EROS-2[2]	7×10^6	6.7 years	1	$<0.36\times10^{-7}$	$\sim0.4\ M_\odot$
				$<0.19\times10^{-7}$	10^{-3}–$10^{-1}\ M_\odot$
				$<0.47\times10^{-7}$	10^{-6}–$1\ M_\odot$
OGLE-III[3]	35×10^6	8 years	2–4	$(0.16\pm0.12)\times10^{-7}$	$<1\ M_\odot$

The MACHO content of the halos of other galaxies can be studied by looking for microlensing of distant quasars that are being strongly lensed by intervening galaxies: a microlensing event would amplify one of the lensed images but not the other, because of the slightly different light paths through the lens galaxy, and can hence be identified by monitoring the ratio between the fluxes from the two images. One analysis[4] studied 20 lens galaxies, and found that the fraction of the galaxy mass in compact objects was $0.05^{+0.09}_{-0.03}$. This is consistent with the expected mass fraction from stars, and precludes a significant contribution from halo MACHOs in the range of 0.1–10 $M_\odot$.

MAPPING GALAXY CLUSTERS WITH STRONG AND WEAK LENSING

Microlensing is useful for investigating the contribution to galactic dark matter of fairly massive (at least planet-sized) dark objects. Strong and weak lensing are sensitive to the overall gravitational potential of the massive lens—typically a large galaxy or cluster of galaxies—and can therefore be used to investigate the distribution of dark matter of all types.

In strong lensing (SL), multiple images of the same source object are seen. The mass distribution in the lensing cluster is reconstructed by adjusting it until the positions of the images match those observed. This is a difficult problem, because it involves inversion of a matrix (a procedure notoriously sensitive to experimental error). Most SL analyses construct model potentials with many components (typically including contributions from individual galaxies plus an overall 'cluster dark halo' contribution parameterised using a model derived from large-scale structure simulations (such as the Navarro-Frenk-White, NFW, model) and then adjust the model until the predicted positions and shapes of the multiple images match those observed. However, at least one analysis[5] uses a direct matrix inversion (and therefore does not make any assumptions about the correlation between light and mass).

Because strong lensing requires that the light from the source pass very close to the lens centre, SL analyses are typically most sensitive to the core of the lensing cluster. In contrast, weak lensing (WL) analyses study small but systematic distortions of background galaxy images (shear), producing mass profiles that extend much further out from the cluster centre. Strong and weak lensing analyses should therefore complement each other, and it is common to make combined fits where both types of data are available.

A recent example of cluster mass mapping[5] considers the cluster Abell 1689. This is a very strong lens with over 100 identified lensed images, and has been studied repeatedly with both strong and weak lensing. The image below[5] shows the result of a mass map made by direct matrix inversion, superimposed on an HST image of the cluster (the orientation, with angular and linear scale, is indicated by the arrows at bottom right). The 135 pink dots represent identified lensed images which are correctly reproduced by the mass model. The white contour represents the outer limit of the region within which the authors consider their model reliable.

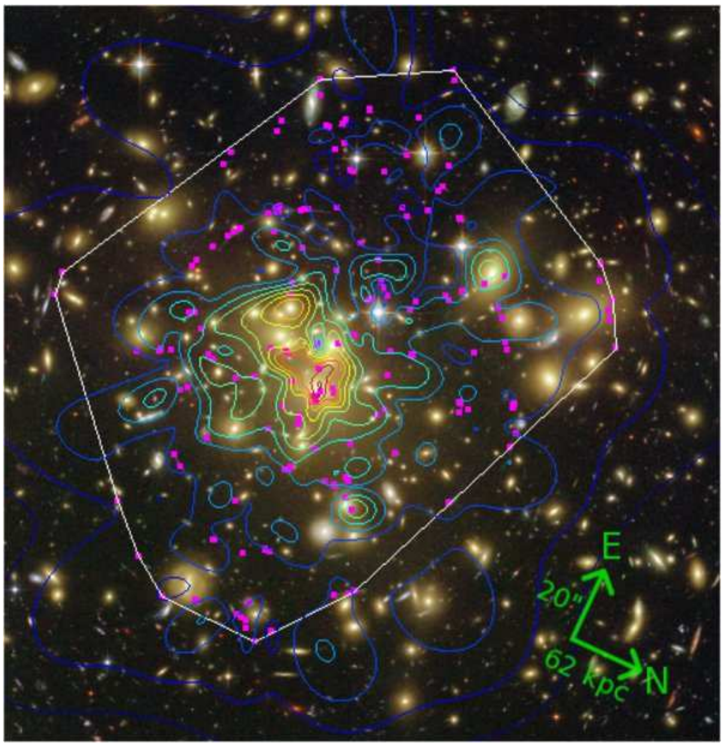

This mass map does not use the visual image as input, so it is encouraging that several of its mass concentrations do coincide with visible galaxies. The authors believe that the map resolves structures about 23 kpc across, which is approximately the size of an individual large galaxy.

Dark halos are usually fitted by the Navarro-Frenk-White profile,

$$\rho(r) = \frac{\rho_0}{\left(\frac{r}{r_s}\right)\left(1+\frac{r}{r_s}\right)^2},$$

where ρ_0 and the scale radius r_s are the fit parameters. This model has the disadvantage that it doesn't have a well-defined total mass (integrating $4\pi r^2 \rho(r)$ from 0 to ∞ produces an infinite result), but this can be dealt with by defining the 'edge' of the cluster as its virial radius (the radius within which the virial theorem holds for the cluster). The dimensionless variable *concentration*, c, is defined as $c = R_{\text{vir}}/r_S$, and the mean density within the virial radius is given by

$$\langle \rho^2 \rangle_{R_{\text{vir}}} = \frac{\rho_0}{c^3}\left(1 - \frac{1}{(1+c)^3}\right).$$

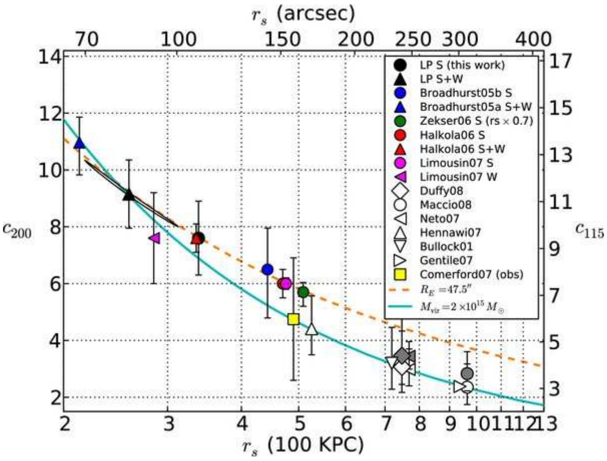

The figure on the left[5] shows fit results for Abell 1689 from a number of different observations (coloured symbols) and simulations (open symbols). Coe *et al.* point out that the variables plotted are tightly correlated, so the error bars should not be vertical but aligned along the dashed line (which corresponds to the observed Einstein ring radius of 47.5 arc seconds), as shown for one example (LP S+W, second from left). There is general agreement that the total mass of the cluster is about 2×10^{15} $M_\odot$ (the solid line), but a rather wide range of opinions as to the actual parameters of the NFW profile; Coe *et al.*[5] comment that there are systematic differences between WL and SL fits, with WL preferring a steeper mass slope than SL.

All gravitational lensing maps of clusters agree that the bulk of the cluster mass is made up of a rather smooth underlying distribution centred on the centre of the cluster—it is *not* concentrated in the individual galaxies. This is also true for the baryonic matter: most of it is actually in the form of very hot, tenuous gas pervading the entire cluster: the *intracluster medium* or ICM, which can be observed directly in X-rays (it is at a temperature of tens of million kelvin). The surface brightness and temperature profile of this gas can be used to estimate the total mass of the cluster, assuming that it is spherical and in thermal equilibrium (these assumptions are *not* guaranteed to be safe). As the gas is optically thin, the X-ray luminosity also directly measures the total mass of the ICM itself, which is consistently found to be much greater than the baryonic mass of the individual galaxies, but much less than the total mass of the cluster: for example, the Virgo cluster (our nearest large cluster) has total mass about 10^{14} $M_\odot$, of which the ICM accounts for about 14% and the galaxies about 4%. A recent study using simulations[6] concluded that the cluster masses obtained using X-ray data are systematically low by about 25%, but have little scatter, whereas in contrast WL estimates are less biased (good to 10%) but much more scattered.

The mass-to-light ratios of clusters are very large, typically several hundred, although this does depend on the wavelength used. The more massive the cluster, the greater its mass-to-light ratio.

CLUSTER COLLISIONS

A special case of mass mapping by weak lensing occurs when clusters of galaxies collide. In this case, the X-ray-emitting ICM should be much more affected by the collision than either the individual galaxies (which are spaced far enough apart that most of them don't collide) or the dark matter halos (which, being weakly interacting, should pass through each other). As the ICM accounts for most of the cluster mass, if modified-gravity models are correct, lensing maps should follow the gas, whereas if the dark matter model is correct, they should more closely follow the visible galaxies. The best-known example is the famous *Bullet Cluster*, shown in the figure on the next page. The smaller cluster on the right (the bullet cluster itself) has passed through the larger one on the left. In this image, X-ray data (pink) and a WL mass map (blue) are superimposed on an optical HST image. As

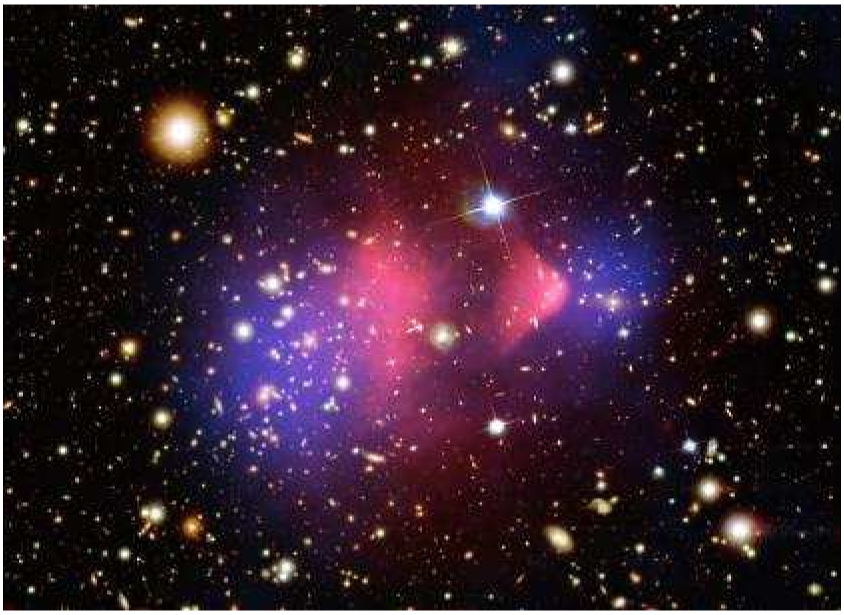

expected from dark matter models, the mass map follows the galaxies, whereas the intracluster media of the two galaxies have interacted, forming a shock front between the two clusters. This distribution is well modelled by simulations of the collision, and similar distributions have been found in other merging clusters.

[It should, of course, be noted that proponents of modified gravity models have claimed that their models *can* in fact account for the structures seen in this image. Theorists are naturally inventive!]

SUMMARY

Gravitational lensing is an important tool for studies of dark matter (also for cosmological studies in general). Of the three main forms of gravitational lensing,

- *microlensing* can be used to investigate the presence of dark objects of roughly stellar mass in the halo of our Galaxy;
- *strong lensing* can be used to map the mass distribution in the core regions of galaxy clusters (and sometimes large individual galaxies);
- *weak lensing* can be used to map the mass distribution throughout galaxy clusters, and also (using wide area surveys) the background large-scale structure.

Microlensing has shown that compact objects of planetary to stellar mass are *not* a major constituent of the Milky Way's dark halo. This is consistent with analyses of the cosmic microwave background (CMB) and Big-Bang Nucleosynthesis, which concur that the baryon density should be only about 4% of the critical density, whereas galaxy rotation curves give $\Omega_g \sim 0.1$.

Weak and strong lensing have shown that the mass of rich clusters of galaxies is not localised in the galaxies, but is mostly in the form of a fairly smooth underlying 'cluster halo' making up of order 85% of the cluster mass. This is also consistent with the ratio of baryonic to non-baryonic matter from analysis of the CMB data.

REFERENCES FOR DATA

1. C. Alcock *et al.* (MACHO Coll.), *ApJ* **542** (2000) 281.
2. P. Tisserand *et al.* (EROS Coll.), *A&A* **469** (2007) 387.
3. Ł. Wyrzykowski *et al.* (OGLE Coll.), *MNRAS* **413** (2011) 493.
4. E. Mediavilla *et al.*, *ApJ* **706** (2009) 1451.
5. D. Coe *et al.*, *ApJ* **723** (2010) 1678.
6. E. Rasia *et al.*, astro-ph/1201.1569 (2012), submitted to *New Journal of Physics*.