1

Introduction

1.1 Background

Because gravity is a long-range force, it is difficult to define precisely the concept of an 'isolated star' – and consequently also the concept of a binary or triple star. Nevertheless, many stars are found whose closest neighbouring star is a hundred, a thousand or even a million times closer than the average separation among stars in the general neighbourhood. Such pairings of stars are expected to be very long lived. There also exist occasional local clusterings of perhaps a thousand to a million stars, occupying a volume of space which would much more typically contain only a handful of stars. These clusters can also be expected to be long lived – although not as long lived as an 'isolated' binary, since the combined motion of stars in a large cluster causes a slow evaporation of the less massive members of the cluster, which gain kinetic energy on average from close gravitational encounters with the more massive members. Intermediate between binaries and clusters are to be found small multiple systems containing three to six members, and loose associations containing somewhat larger numbers. Starting from the other end, some clusters may contain sub-clusters, and perhaps sub-sub-clusters, down to the scale of binaries and triples.

Even with the naked eye, a handful of the 5000 stars visible can be seen to be double; and in the northern hemisphere two clusters of stars, the Hyades and the Pleiades, are quite recognisable. But some 2000 naked-eye stars are known to be binary (or triple, quadruple, etc.) by more detailed measurement – astrometric, spectroscopic or photometric. Observation in other wavelength ranges, such as radio, infrared, ultraviolet and X-rays, reveal further and more exotic binary companions, not so many in number, but of unusual interest. The naked-eye stars are only a tiny fraction of all the stars in our Galaxy ($\sim 10^{11}$), but are reasonably representative as far as the incidence of binarity is concerned.

Sometimes the two components are so close together as to be virtually touching; sometimes they are so far apart as to be virtually independent. Measured orbital periods range from hours (or even minutes) up to centuries. Many must have longer periods still, not yet determined but up to millions of years. The evolution of the two components of such pairs has attracted increasing interest over the last fifty years. The presence of a binary companion, if the orbital period is a few years or less, may make the evolution of a star very different from what it would have been if the star were effectively isolated. A number of these differences are now fairly well understood, but although some evolutionary problems which used to trouble astrophysicists, such as the 'Algol paradox', have been largely resolved, several still remain. New observations add new problems considerably faster than they confirm the resolution of older problems. It should be kept in mind that even single stars present many evolutionary problems, and so it is not surprising that many binary stars do.

Questions about binary stars can be divided very loosely into two categories, 'structural' and 'evolutionary'. For a particular type of binary star one can ask what physical processes are currently going on, that give this type of star its particular characteristics. In cataclysmic variables such as novae, for instance, there is little doubt that a fairly normal main sequence star of rather low mass is being slowly torn apart by the gravitational field of a very close white dwarf companion. But one can also ask how such binaries started, and subsequently evolved, so that these processes can currently take place. This evolutionary question can be harder to answer, because most evolutionary processes are very slow. An obvious further evolutionary question is: 'What will the future evolution of such systems be, up to some long-lived final state?' This book attempts to summarise progress in understanding the kind of long-term evolutionary processes involved. In the interest of brevity it will be necessary to quote, and to take for granted rather than to discuss, most of the much more substantial literature on structural problems. However, one aspect of binary stars that might be labelled 'structural', but which is certainly of vital importance for evolutionary discussions, is the determination of such fundamental parameters as masses, radii, etc.

1.2 Determination of binary parameters

If we are interested in determining the masses and radii of stars, then we have to turn almost right away to *binary* stars, since it is only by measuring orbital motion under gravity, and by measuring the shape and depth of eclipses, that we are able to determine these quantities to a good accuracy – one or two per cent in favourable cases; see Hilditch (2001). Analysis of the spectrum of an *isolated* star can determine such useful quantities as the star's surface temperature, gravity and composition. This is done by comparing the observed spectrum, preferably not just in the visible region of wavelengths but also in the ultraviolet (UV) and infrared (IR), with a grid of computed spectra for a range of temperatures, gravities and compositions. However, we do not get a mass from this process, or a radius, only the combination that gives the gravity – except in the special case of white dwarfs, where there is expected to be a tight radius–mass relation (Section 2.3.2) so that both mass and radius are functions only of gravity.

If we have an accurate parallax, as from the Hipparcos satellite, we can get closer to determining the mass of an isolated star, because the distance, the temperature (from spectral analysis), and the apparent brightness give us the radius; and hence the gravity (also from spectral fitting) gives us the mass. However, even if the parallax is good to $\sim$1%, the gravity is much less accurate, because spectra are usually nothing like so sensitive to gravity as they are to temperature. Perhaps an accuracy of $\sim$25% is achievable.

The parameters of binary systems are generally obtained from astrometric, or spectroscopic, or photometric observations, and in favourable cases by a combination of two, or even all three, of these methods. Note that terms such as 'astrometric' and 'photometric', coined originally to refer to observations in the visible portion of the electromagnetic spectrum, are now generally used to cover all parts of the spectrum, for instance radio and X-rays. If the two components of a binary are so far apart in the sky as to be resolvable from each other, which means at visual wavelengths more than $\sim 0.1''$ (0.5 μrad) apart, then the system is a 'visual binary' or 'VB', and careful astrometry, sometimes over a century or more, can reveal the orbit. Visual binaries tend to have long periods because short-period orbits are generally not resolvable. Only for systems within $\sim$5 pc of the Sun (about 50 in number) could a separation of $0.2''$ correspond to a period of $\lesssim$1 year. The upper limit of well-determined

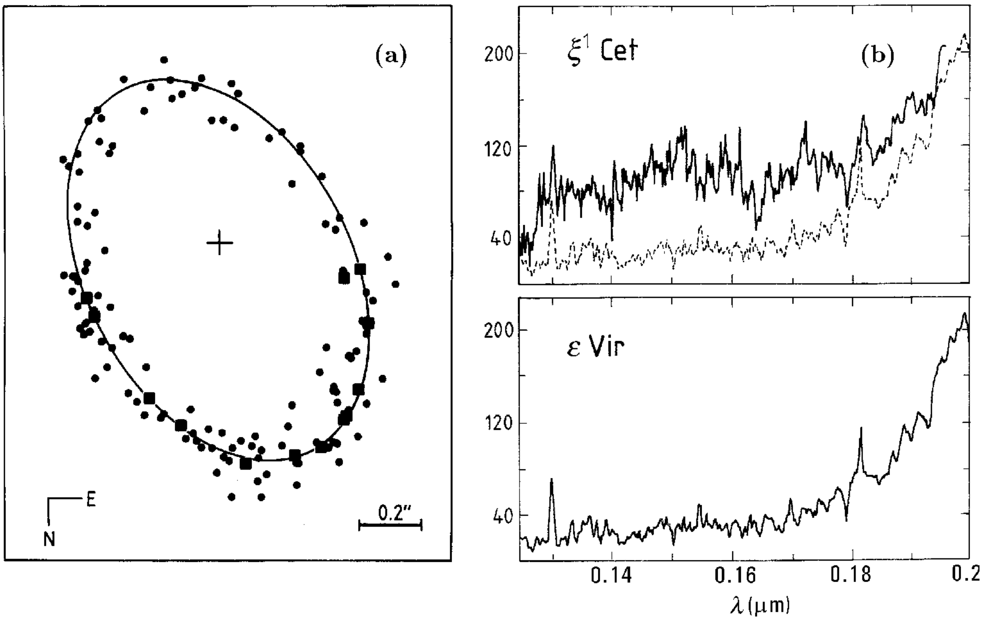

Figure 1.1 (a) The orbit of HR 3579 (F5V+G5V) from visual (dots) and speckle (square) measurements of relative position. The scatter of speckle points about the best-fit curve ($P = 21.8$ yr, $e = 0.15$, $a/D = 0.66''$, $i = 130°$) is much less than for the visual points. From Hartkopf *et al.* (1989). (b) The UV spectrum of the G8III stars ϵ Vir (bottom panel) and ξ^1 Cet (top panel, with ϵ Vir repeated). For 0.18–0.7 μm (not all shown here) the spectra are very similar. The UV excess evident in ξ^1 Cet for 0.13–0.17 μm is attributable to a white dwarf companion. From Böhm-Vitense and Johnson (1985).

visual orbital periods is about 100 years, because good accuracy is only achievable if the VB has consistently been followed for at least two full orbits. There are many orbits in the literature with periods up to 1000 years, or even longer, but these must be considered tentative – extremely tentative if the period is greater than 200 years.

Visual orbits are usually *relative* orbits, the position of one component being measured relative to the other (Fig. 1.1a). Visual orbits have been catalogued by Worley and Douglas (1984), and speckle measurements by McAlister and Hartkopf (1988). These and many other relevant catalogues can be found on the website of the Centre des Données astronomique de Strasbourg (http://cdsweb.u_strasbg.fr). From visual orbits one can determine the period (P), the eccentricity (e), the inclination (i) of the orbit to the line of sight, and the *angular* semimajor axis, i.e. the ratio of the semimajor axis a to the distance D. One can then determine M/D^3, where M is the total mass, from Kepler's law:

$$\frac{GM}{a^3} = \left(\frac{2\pi}{P}\right)^2, \quad \text{so} \quad \frac{GM}{D^3} = \left(\frac{2\pi}{P}\right)^2 \left(\frac{a}{D}\right)^3. \tag{1.1}$$

If the VB is near enough, D may be obtainable from the parallax. For Earth-based measurements, parallaxes of less than $0.1''$ are not reliable, but this has been improved by more than an order of magnitude with space-based measurements from the Hipparcos satellite. If the orbits of both the components of a visual binary can be measured *absolutely*, i.e. each orbit relative to a background of distant and approximately 'fixed' stars, then the mass ratio of the two components can further be determined. We still do not obtain the individual masses, however, unless D is separately determinable.

Even if only one component of a binary is visible at all, an astrometric orbit may in favourable cases be found by observing that the position of a star has a cyclic oscillation superimposed on the combination of its parallactic motion and its linear proper motion relative to the 'fixed' stars, i.e. faint stars most of which do not move measurably and so can be assumed to be distant. Such astrometric binaries can yield P, e and i, but information on masses is convolved with the unknown mass ratio, and also with the parallax which may or may not be measurable even if the astrometric orbit is measurable.

Some VBs can be recognised even when neither component shows measurable orbital motion. If two stars, not necessarily *very* close together on the sky, show the same substantial linear proper motion relative to the 'fixed' background, it is likely that (a) they are physically related, and (b) fairly nearby, with measurable parallaxes. Usually these parallaxes agree, confirming the reality of the pair. Such pairs are called 'common proper motion' (CPM) pairs. The two nearest stars to the Sun, V645 Cen (Proxima Cen) and α Cen, are over 2° apart, but have the same rapid proper motion and large parallax. To be pedantic, (a) they are so near the Sun, and so far apart on the sky, that actually their proper motions and parallaxes are measurably different at the 1% level, and (b) α Cen is itself a VB of two Solar-type stars, with semimajor axis 17.5″ and period 80 years, so that the proper motion of V645 Cen has to be compared with the proper motion of the centre of gravity (CG) of the α Cen pair. The period of the orbit of V645 Cen about the CG of the triple system can be expected to be about 1 megayear.

Common proper motion pairs are usually sufficiently wide that they might appear to be of little relevance to this book, which deals with pairs sufficiently close together that one component can influence the other's evolution. However the presence of a CPM companion can often reveal information on both components that would not be available if they were not paired. Several *close* pairs have a distant CPM companion; and if for example this companion has a character that suggests that it is fairly old, then one can reasonably conclude that the close binary is also fairly old. This may not be evident from the close binary alone, since the components in it may have interacted in ways that disguise the age of the system.

Modern techniques such as speckle interferometry (Labeyrie 1970, McAlister 1985), can resolve components with substantially smaller angular separations than conventional astrometry, and thus determine visual orbits of shorter period. The major limitation on resolving close components astrometrically is atmospheric 'seeing', the blurring effect of turbulence in the Earth's atmosphere. This distorts the image on a timescale of ~0.05 s. In the speckle technique the image is recorded many times a second, and so the time variation of the point-spread function can be followed and allowed for in a Fourier deconvolution. The technique of adaptive optics (Babcock 1953, Beckers 1993) is an alternative way of eliminating seeing, by continuously adapting the shape of the mirror in response to the deformation of the image of a reference point source, either a nearby single star or the back-scattered light of a laser beam pointing along the telescope. Both techniques can give resolution down to the limit of diffraction, ~0.01″ at visual wavelengths on a modern 8 m class telescope. By combining the light from two or more separate telescopes, the technique of 'aperture synthesis', long used in radio astronomy, can nowadays be applied to optical wavelengths (Burns *et al*. 1997), and should be capable of sub-milliarcsecond resolution, so that one might hope to see directly both components of nearby short-period binaries.

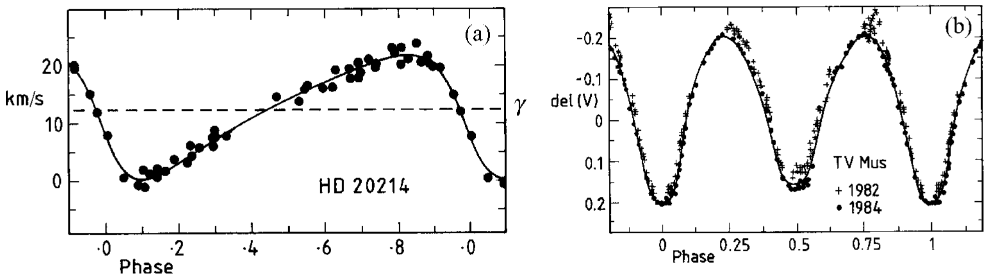

Figure 1.2 (a) The radial velocity curve of the K giant star HD20214. The rms scatter about the mean curve is only ~ 0.2 km/s. Orbital parameters are $P = 407$ days, $e = 0.41$, $f = 0.040\, M_\odot$. From Griffin (1988). (b) The light curve of a contact binary TV Mus ($P = 0.446$ days, $e = 0$, $i = 78.9^\circ$, $R_1/a = 0.59$, $R_2/a = 0.27$, $M_1/M_2 = 7.2$, $T_1/T_2 = 0.98$). A slight variation in brightness over two years, and a small distortion in the secondary eclipse, may be due to starspots. From Hilditch *et al*. (1989).

Systems may be recognisable as spectroscopic binaries (SBs) either because the spectrum is composite (Fig. 1.1b), or because it shows radial velocity variations (Fig. 1.2a), or both. In a composite spectrum, one might see for instance a combination of the relatively broad lines of H and He characteristic of a B dwarf with the narrow lines of Fe and other metals characteristic of a G or K giant. Alternatively, a star whose spectrum at visual wavelengths may seem like a K giant may be found, at UV wavelengths, to have an excess flux that can be attributed to a hot companion, sometimes even a white dwarf (Fig. 1.1b). It is not easy to disentangle composite spectra reliably, since things other than a stellar companion (for example a corona, a circumstellar disc or a dust shell) may contribute to an excess either in the UV or the IR. Even if the spectrum seems definitely a composite of two stellar spectra, we learn only that the star is a binary; we do not obtain information about the orbit unless one spectrum at least shows a variable radial velocity, consistent with Doppler shift due to motion in a Keplerian orbit.

Orbits of 1469 SBs have been catalogued in the important compilation of Batten, Fletcher and McCarthy (1989). The number of orbits is increasing rapidly, perhaps already at a rate of one or two hundred a year, and no doubt with greater rapidity in the future, partly because of cross-correlation techniques and partly because of the much-increased sensitivity of detectors. Commonly SBs are single-lined ('SB1'), but the radial velocity variation of the single spectrum seen (as in Fig. 1.2a) allows P and e to be obtained and also the amplitude K of the radial velocity variation, or equivalently (as is usual for radio pulsars) the projected semi-major axis ($a \sin i \propto K P \sqrt{1 - e^2}$). Information on masses is contained in a single function, the mass function f, convolving both of the masses with the unknown orbital inclination i:

$$f_1 = \frac{M_2^3 \sin^3 i}{(M_1 + M_2)^2} = \frac{K_1^3 P (1 - e^2)^{3/2}}{2\pi G} = 1.0385 \times 10^{-7} K_1^3 P (1 - e^2)^{3/2}$$

$$= 1.0737 \times 10^{-3} \frac{(a_1 \sin i)^3}{P^2}, \qquad (1.2)$$

where $*1$ (pronounced 'star 1') is the observed star and $*2$ the unseen component. Units are: K_1 in km/s, P in days, $a_1 \sin i$ in light-seconds and masses in Solar units. The inclination is not measurable for spectroscopic orbits because we have information on the motion in

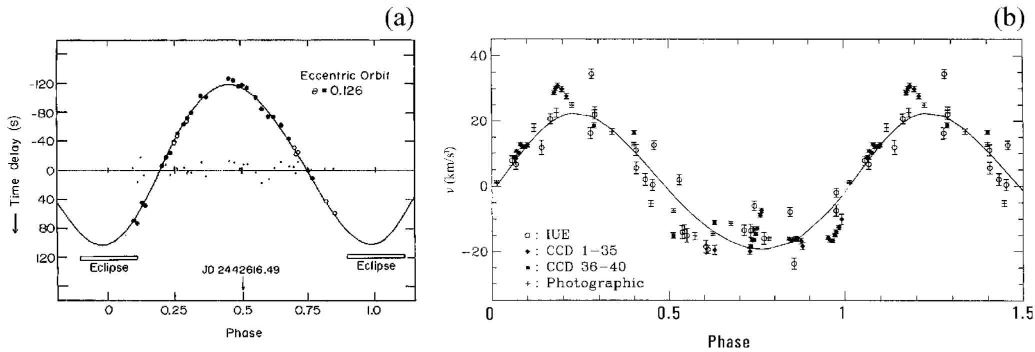

Figure 1.3 Radial velocity curves of both components of the massive X-ray binary Vela X-1 (GP Vel). (a) Doppler shift of the pulses of the X-ray pulsar: note the accurate fit to the Keplerian curve ($P = 8.964$ days, $e = 0.126$, $f_1 = 18.5\ M_\odot$). Small dots near the axis are the residuals multiplied by 2. (b) Doppler shift of absorption lines in the visible spectrum: note the larger scatter, due to irregular pulsations. From these lines $f_2 \sim 0.013$. The ratio f_2/f_1 is the cube of the mass ratio q (~ 0.09). (a) is from Rappaport *et al*. (1976), (b) from van Kerkwijk *et al*. (1995b).

only one dimension, the line of sight, whereas in visual binaries we have information in two dimensions, both perpendicular to the line of sight. In fact the red giant in ξ^1 Cet (Fig. 1.1b) does show orbital motion ($P = 1642$ days, $e = 0$, $f = 0.035\ M_\odot$, Griffin and Herbig 1981) in addition to being a composite-spectrum binary.

The mass function represents the minimum possible mass for the unseen star, which would be achieved in the somewhat improbable case $M_1 = 0$, $i = 90°$. Slightly more realistically, we might replace $\sin^3 i$ by its average value $3\pi/16 \sim 0.59$ if i is distributed uniformly over solid angle. However the value 0.59 is likely to be an underestimate, because the mere fact that a variation in radial velocity is seen implies that the lowest inclinations can be rejected. For a large ensemble of binaries we might make statistical estimates using a maximum-likelihood procedure. However, for an isolated system, with little else to guide us, we will commonly assume that a reasonable estimate of the reciprocal of $\sin^3 i$ is 1.25. We then take

$$M_1 \sim 1.25q(1+q)^2 f_1, \tag{1.3a}$$

$$M_2 \sim 1.25(1+q)^2 f_1, \tag{1.3b}$$

where $q \equiv M_1/M_2$ is the mass ratio. Sometimes we can estimate M_1 directly from the spectrum of the star, which may be similar to stars whose masses are already known from more favourable binaries (see below); then from Eq. (1.3a) q can be estimated and hence M_2. Alternatively one can often infer that $q > 1$ simply from the probability that the unseen star is less massive than the visible one. In either case both masses could be *considerably* greater than the mass function.

If the system is 'double lined' ('SB2'), and both components have measurable radial velocity variations (Fig. 1.3), we can further obtain the mass ratio, and hence the two quantities $M_1 \sin^3 i$ and $M_2 \sin^3 i$; but we still have no information on i. However, some SBs with $P \gtrsim 1$ year are also VBs, and in favourable cases all four of M_1, M_2, i and D can be separately measured, D in such cases being *independent* of parallax (which may be too small to be measurable).

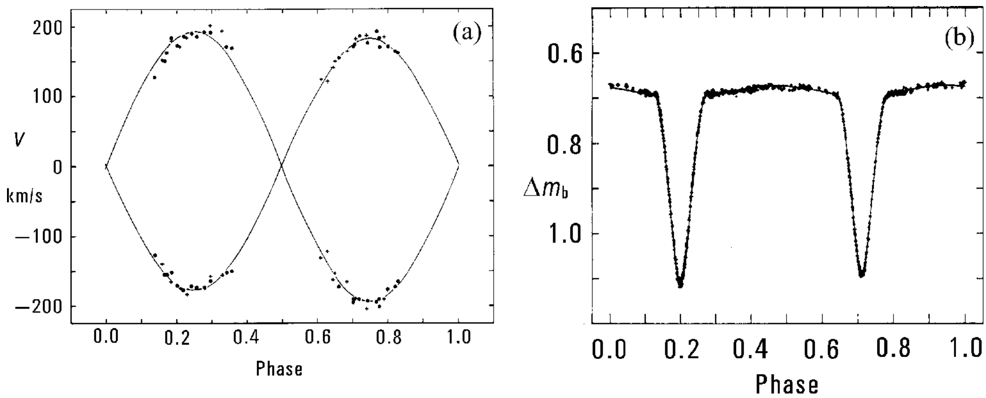

Figure 1.4 (a) The radial velocity curves and (b) the light curve of the eclipsing SB2 system V760 Sco ($P = 1.73$ days). The two components are nearly but not quite identical: in (a), $*2$ has a slightly greater velocity amplitude, and in (b) the second eclipse is slightly shallower than the first. An 'ellipsoidal variation' is seen in the nearly flat portions between eclipses. From Andersen *et al*. (1985).

Among SBs we can include both radio and X-ray pulsars, because the rapid pulsations of these objects, due to rapid rotation of an obliquely-magnetised neutron star, are often very stable and so can reveal a variable Doppler shift due to Keplerian orbital motion. Commonly, pulsar orbits are much more accurate than SB orbits based on spectral lines, so that even companions of terrestrial planetary mass can be detected (Wolszczan and Frail 1992). The much greater accuracy of radio pulsar orbits means that a number of relativistic corrections to Keplerian orbits can be measured (Taylor and Weisberg 1989, Backer and Hellings 1986). Two of these are (a) the rate Z_{GR} of advance of periastron in an eccentric orbit due to general relativity – Appendix C(a):

$$Z_{GR} = \frac{3G(M_1 + M_2)}{c^2 a(1 - e^2)} \frac{2\pi}{P}, \tag{1.4}$$

and (b) a combination γ of gravitational redshift and transverse Doppler shift:

$$\gamma = \frac{G(M_1 + 2M_2)e}{c^2(M_1 + M_2)} \frac{P}{2\pi}. \tag{1.5}$$

Along with the mass function Eq. (1.2), these two quantities allow one to determine all three of M_1, M_2 and i, even although the orbit is 'single lined'.

X-ray pulsar orbits, though commonly more accurate than radial-velocity orbits from spectral lines (Fig. 1.3), are also commonly less accurate than radio pulsar orbits, because the X-rays come from accretion of gas lost by the companion. The gas flow is normally not steady, and so the neutron star's spin rate is erratically variable by a small amount.

Photometric binaries are stars whose light output varies periodically, and in a manner consistent with orbital motion. Usually they show eclipses, but in some cases where the inclination does not permit an eclipse one may nevertheless recognise 'ellipsoidal variation' or the 'reflection effect' (see below). A light curve (Figs. 1.2b, 1.4b) can yield, in favourable circumstances, P, e and i, the ratios R_1/a, R_2/a of stellar radii to orbital semimajor axis, and the temperature T_2 provided that T_1 is known already, from a spectroscopic analysis of the brighter component. The radius ratios and i come primarily from the duration and shape of the total and partial segments of the eclipse, and the temperature from the relative depths of the deeper and shallower eclipse in each cycle. Although some light curves can be analysed crudely by assuming that both stars are spheres, the majority of eclipsers need more sophisticated modelling, usually assuming that both components fill equipotential surfaces of the combined gravitational and centrifugal field of two orbiting point masses (the Roche potential, Chapter 3). Such light curve analysis was pioneered by Lucy (1968), Rucinski (1969, 1973), and Wilson and Devinney (1971). Information on 3546 eclipsing binary stars is given in the catalogue of Wood *et al*. (1980). A catalogue by Budding (1984) gives light curve solutions for 414 eclipsers.

An eclipsing binary is also usually a spectroscopic binary, but not conversely. This is because eclipses are only probable in systems where one star's radius is $\gtrsim$10% of the separation, whereas there is no such limit on radial-velocity variations. In the best cases, where the system has eclipses and is also double lined ('ESB2', as in Fig. 1.4), we can hope to obtain all of the following fundamental data: P, e, i, a, M_1, M_2, R_1, R_2, T_1, T_2 and D (independent of parallax). The last three of these quantities depend not only on good orbital data but also on reliable modelling of stellar atmospheres, so that the effective temperature of at least one component (presumably the brighter) can be determined directly from its spectrum. This is probably reasonable for the majority of stars, but for extremes of effective temperature and luminosity (O and M stars; supergiants and subdwarfs), spectra may be affected by such difficulties as mass loss, instability, convection and metallicity, all of which are not yet well understood. A comprehensive review of data for ESB2 binary stars in the main-sequence band has been given by Andersen (1991); an earlier review by Popper (1980) also gave data for some post-main-sequence binaries. Accuracies of $\lesssim$2% for all quantities are achievable in favourable cases.

Binaries involving evolved stars (giants, supergiants, hot subdwarfs, white dwarfs, etc.) are relatively rare, especially ESB2 systems. Although the photometric and spectroscopic data may be of the same quality, or even better, it is difficult to achieve the same accuracy in the estimation of radii. This is because the two radii are of course very different in giant/dwarf binaries. The information on relative radii, as well as on inclination, is contained in the shape of the ingress/egress portions of eclipses. If one star is so much larger than the other that its occulting edge is virtually a straight line, then the inclination and hence also the ratio of radii are indeterminate. Nevertheless supplementary information from model atmospheres, and from spectrophotometry, the measurement of intensity in several wavebands that may extend from UV to IR, can reduce the indeterminacy. Recent work on such 'ζ Aur' systems (Schröder *et al*. 1997) gives parameters with sufficient accuracy that theoretical models of stellar evolution are seriously tested.

The fact that ESB2 binaries can in principle give a distance measurement that is independent of parallax implies that they could be good yardsticks for measuring distances to external galaxies. Current and developing technology means that at least OB-type binaries may be accessible in fairly nearby galaxies. Of course one does need an estimate of the metallicity in order to relate measured colours to the effective temperature of at least the hotter component.

Because stars in close binaries can be distorted from a spherical shape by the combined gravitational and centrifugal effect of an orbiting close companion, they may show a measurable light variation even when they do not eclipse. This is called 'ellipsoidal variation' – although the stars are only approximately ellipsoidal. Figure 1.4b shows this variation. The system illustrated is in fact at an inclination which also allows eclipses: the ellipsoidal

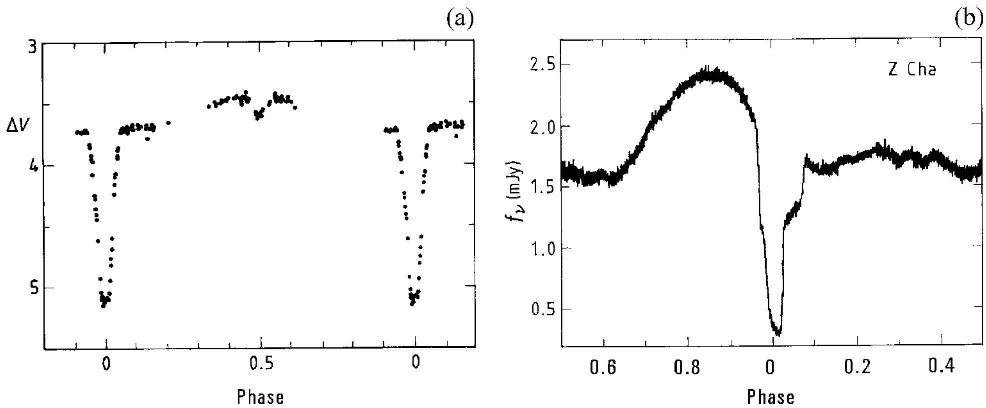

Figure 1.5 (a) The light curve of UU Sge ($P = 0.465$ days), the central star of the planetary nebula Abell 63. The hump centred on the secondary eclipse is due to a ‘reflection effect’. The fainter, cooler companion shines partly by reprocessed UV light from the very hot companion; thus it is brightest just before and after it is eclipsed, and is rather faint for half the orbit. From Bond *et al.* (1978). (b) The light curve of Z Cha, an ultra-short-period binary containing a white dwarf and a red dwarf ($P = 0.0745$ days). The hump before the eclipse, the double-stepped nature of the eclipse, and the erratic variation are all due to streams of gas flowing from the red dwarf towards, and round, the white dwarf. From Wood *et al.* (1986).

variation is the slight curvature visible between the eclipses. Such variation even in the absence of eclipses may allow at least P to be determined. Further, if $*1$ (say) is much hotter than $*2$, the hemisphere of $*2$ facing $*1$ may be substantially brighter than the other hemisphere, leading to an orbital variation (Fig. 1.5a) that also does not necessarily involve an eclipse. This is called the ‘reflection effect’ – although the light (or X-radiation, in some cases) is absorbed, thermalised and reemitted, rather than reflected.

However, not all eclipse light curves, even with high signal-to-noise ratios and with modern light-curve synthesis techniques, lend themselves to accurate measurement of fundamental data (masses, radii, etc.). Neither do all radial velocity curves, even when a non-uniform temperature distribution over the stellar surfaces due for example to the reflection effect is allowed for. This is because stars which are close enough together to have a reasonable probability of eclipse (typically, $R_1 + R_2 \gtrsim 0.2a$) are also quite likely to interact hydrodynamically and hydromagnetically, introducing the complications of gas streams, and of starspots, which are hard to model in any but an ad hoc manner. Figure 1.2b shows a light curve of a contact binary that changed appreciably over time. The changes, and slight asymmetry, can be attributed to transient starspots. Figure 1.5b shows the light curve of a dwarf nova: an eclipse of sorts is clearly recognisable, but the light variation outside the eclipse is due to gas which streams from one component into a ring or disc about the other. Modern methods of analysis such as eclipse mapping (Horne 1985, Wood *et al.* 1986) and Doppler tomography (Marsh and Horne 1988, Richards *et al.* 1995) use image-processing techniques based on maximum-entropy algorithms (Skilling and Bryan 1984). The object of eclipse mapping is to reconstruct the distribution of light intensity over (in the case of Z Cha, Fig. 1.5b) a hypothesised flat, rotating disc of gas around one star that is fed by a stream that comes from the other star. The eclipsing edge of one star as it moves across the disc and stream helps to locate the hotter and

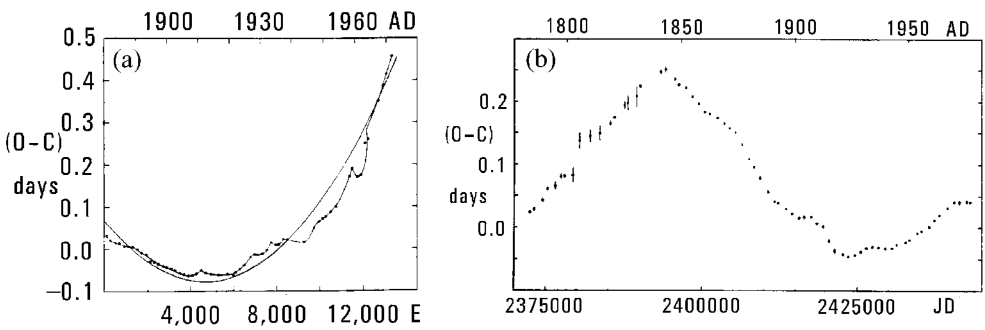

Figure 1.6 Observed times of eclipse, minus computed times obtained by assuming a constant period, plotted against cycle number (epoch) along the bottom and date along the top. (a) U Cep (G8III + B7V; 2.5 days), from Batten (1976). (b) β Per (G8III + B8V; 2.9 days), from Söderhjelm (1980). U Cep shows small erratic variations superimposed on a long term trend of increasing period; β Per also shows erratic fluctuations, but with no clear long-term trend.

cooler parts of the flow. In Doppler tomography, high wavelength resolution across a spectral line, combined with high time resolution, gives a map of intensity on a two-dimensional space of wavelength and orbital phase. This can in principle be Fourier-inverted to map intensity onto a two-dimensional velocity space, and from there one can go via some hypothesised model to a distribution in two-dimensional coordinate space. This might be either a disc-like structure, as in Z Cha, or a distribution of spots over a spherical surface, or even of spots over the joint surface of two stars that are so close as to be in contact (Bradstreet 1985). In this way one can hope to remove the distorting effect of spots and streams from the observational data, and thus be left with accurate fundamental data. But the hypothetical models of spots and streams are not in practice very strongly constrained – for example some systems may contain hot spots as well as cool spots – and so there remains considerable uncertainty in the fundamental data for many, indeed most, interacting systems.

Much information on the statistics of eclipsing binaries (and other types of variable star) comes, as a by-product, from gravitational microlensing experiments (Paczyński 1986). If a relatively nearby star happens to pass very close to the line of sight of a distant star, the apparent brightness of the distant star is temporarily increased by gravitational focusing in the field of the nearby lensing star. Such events are rare, but have been detected by several astronomical groups who monitor photometrically a large number of stars ($\sim 10^6$) in a small area of sky at frequent intervals (e.g. nightly) over several years. The light curve of a lensing event is recognisably different from the light curves of pulsators, eclipsers, novae etc.; but a large number of normal eclipsers shows up as well, and this gives a valuable database from which the statistics of orbital periods can be improved (Udalski *et al.* 1995, Alcock *et al.* 1997, Rucinski 1998). A very few lensing events also exhibit binarity directly: if the lensing object is binary it can produce a marked characteristic distortion on the light curve of a lensing event (Rhie *et al.* 1999).

Some binaries, particularly eclipsing binaries, show a measurable change of period over substantial intervals of time. Period changes are usually demonstrated by ‘O − C diagrams’ (Fig. 1.6). The difference between the observed time of eclipse, and the computed time based on the assumption of constant period, is plotted as a function of time (or of epoch, i.e. cycle