Night Lab
Eclipsing Binaries

Introduction

In this lab, we will measure the companion masses and radii, as well as the orbital period of the eclipsing binary NSVS01031772. The purpose of these notes is to provide some background on the astrophysics we will probe with these measurements, as well as the observational techniques we will use to make them.

The fundamental parameters that define the structures of stars are:

1. Mass M_*: the best way to measure stellar masses is by measuring the orbits of binary systems, as we will do in this lab! More below.

2. Composition: The "metallicity" of a star is the abundance of elements heavier than hydrogen and helium that it contains (this affects how radiation diffuses through the star), and we measure it by looking for absorption lines of atomic/molecular transitions in the star's spectrum.

3. Age $t_{\rm age}$: This is the hardest to measure!!! The most fundamental way to get stellar ages is by abundances of radioactive isotopes. As you might imagine, these are present only in trace amounts in stellar atmospheres, so measuring their abundances requires ridiculously high-resolution spectroscopy. In practice, we usually infer stellar ages by relying on models of stellar evolution and measurements of the other quantities listed here and below.

We can also measure other things about stars. Stellar evolution theory predicts the following observables given mass, age, and composition:

1. Bolometric luminosity L_*: measured by combining the flux over all frequencies with the distance to the star.

2. Temperature $T_{\rm eff}$: Stars are (mostly) blackbodies, meaning that the peak in their spectral energy distribution gives us the temperature through Wien's law. We've also seen that we can use broadband photometry and color-temperature relations.

3. Radius R_*: the most straightforward way to measure stellar radii is using the light curves of eclipsing binaries (as we will do tonight!) Otherwise, we have to use interferometry to actually resolve the stellar radius (only possible for very nearby stars) or combine other measurements (e.g., the surface temperature and luminosity via the Stefan-Boltzmann law) to infer R_*.

Hopefully you can see that mass and radius measurements move us down the path to better understanding stellar evolution, even if we won't be measuring temperatures, metallicities, or ages for now. The physics of stellar evolution is particularly uncertain for low-mass stars (when stellar astrophysicists say "low-mass," think $\lesssim 0.8\ M_\odot$) despite those stars being the most abundant in our universe! If we better understand their structure and evolution, we don't only win as stellar astronomers. Zooming "in," we learn about the irradiation and space weather around the stars, and thus the habitability of the planets they might host. Zooming "out," accurate stellar ages enable an understanding of the Milky Way's history through the field of "Galactic archaeology," and of the evolution of other galaxies through the field of "stellar population synthesis." (Happy to discuss further while we wait for our data to come in!)

Eclipsing binaries

What is an eclipsing binary? We learned about using spectroscopy to measure the radial velocities (RVs) of objects during the daytime lab. For roughly equal-mass (i.e., equal brightness) binary stars, we can simultaneously measure the RVs of both stars in the system to plot what we call the radial velocity curve. The RVs of each star oscillate around the "systemic" or center-of-mass velocity of the system. See the example below, where two stars follow elliptical orbits about their shared center of mass. We denote the *amplitudes* of the RV curves of each star by K_1 and K_2.

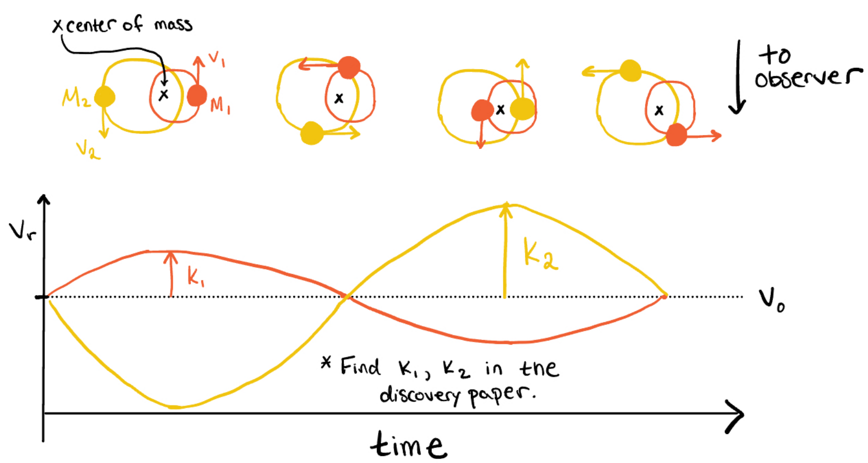

From Kepler's laws, we know that higher velocity amplitudes might indicate higher-mass stars in the binary. At this point, though, we're a little stuck, since the amplitudes of the RV curves *also* depend on the orbital inclination of the binary. For example, a nearly face-on binary would have a very small RV amplitude, no matter the masses of the stars or the intrinsic shape of the orbit, whereas a more edge-on orbit will naturally have a higher RV amplitude. It is therefore not necessarily true that $K_1 = v_1$ and $K_2 = v_2$ in the schematic figure. However, when the orbit of a binary star system is inclined so that the plane of the orbit is along our line of sight, the stars periodically eclipse each other, leading to variations in the light curve. See the figures below.

Eclipses only happen when the orbit is nearly edge-on ($i \sim 90^\circ$). Thus, when we measure an eclipsing binary, we know that the maxima in the RV curve correspond to when the stars are moving directly towards/away from us along the line of sight, giving us the true orbital speeds, or $K_1 = v_1$ and $K_2 = v_2$.

A final assumption we can make to simplify calculations going forward is that the orbit is circular, that is, the orbital speeds of both stars are constant (and so the velocity amplitudes of the RV curves give us orbital speeds at all times). For orbital periods less than a few days, circular orbits are a safe assumption since binaries undergo "tidal circularization" when tidal interactions dissipate orbital energy by deforming the stars (similar physics to what creates the tidally-locked Earth/moon system).

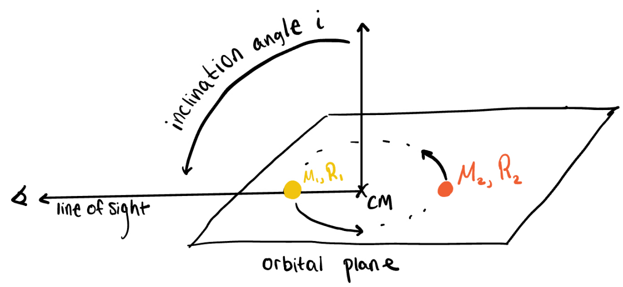

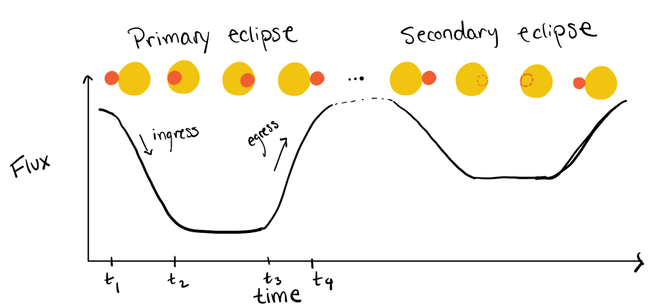

Measuring masses Given a light curve and a radial velocity curve, how can we back out the mass and radius of each star in the system? Let's start by defining the center of mass frame for a binary system. In general, the CM frame is defined such that

$$0 = \sum_i m_i \vec{r}_i, \tag{1}$$

that is, the sum of each mass in the system multiplied by its position is zero.

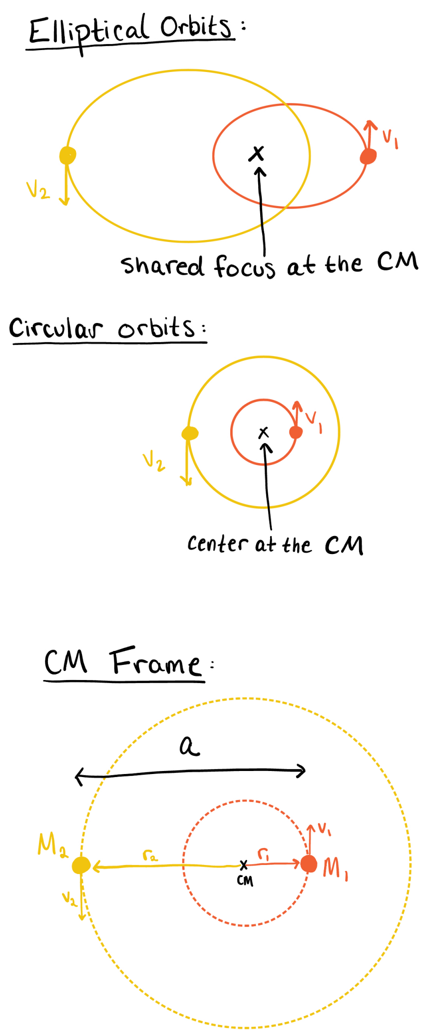

For binaries with elliptical orbits, the center of mass is located at a shared focus of the two ellipses. For binaries with circular orbits, the center of mass is located at a shared center of the two circles. See the schematic to the left.

At the bottom of the schematic figure, we show the definition of the center of mass frame in the circular orbit case:

$$0 = M_1 r_1 - M_2 r_2 \tag{2}$$

(where the negative sign captures the fact that the vector directions of the two stars' positions are opposite). We may now differentiate with respect to time and rearrange to find:

$$\begin{aligned} 0 &= M_1 v_1 - M_2 v_2 \\ \Rightarrow \frac{M_2}{M_1} &= \left| \frac{v_1}{v_2} \right|. \end{aligned} \tag{3}$$

The right hand side of the above equation is the ratio of velocity amplitudes, which we can measure from the RV curve since we've assumed an edge-on orbit. Just from the RV curve, we have recovered the binary mass ratio. Note that in general it is true that $K_1/K_2 = M_2/M_1$, regardless of the orbital inclination, since the velocity semiamplitude $K_i = v_i \sin i$, and in this ratio, the inclination terms will cancel.

To make more progress, we'll use Kepler's third law for orbital period P and orbital separation a:

$$P^2 = \frac{4\pi^2}{G(M_1 + M_2)} a^3 = \frac{4\pi^2}{G(M_1 + M_2)} (r_1 + r_2)^3 \tag{4}$$

Since we've already invoked circular orbits, we can further say that

$$\begin{aligned} v_1 &= \frac{2\pi r_1}{P} \\ v_2 &= \frac{2\pi r_2}{P} \end{aligned} \tag{5}$$

We now have the ingredients to rearrange and solve for M_1 and M_2. Note that we won't measure the RVs ourselves in this lab, but instead use published RV data to do this calculation. See the EB discovery paper in the Canvas files.

Measuring radii How about radii? This we can get by combining the light curves we measure with the known orbital speeds of the stars. The relative velocity of the stars is $v = v_1 + v_2$. During an ingress/egress, one star moves across its own diameter. Further, from the beginning of an ingress to the beginning of an egress, the star moves across the diameter of its companion. Therefore, using the time marks from the light curve sketch above (and assuming a primary eclipse),

$$\begin{aligned} R_1 &= \frac{1}{2}(t_3 - t_1)v \\ R_2 &= \frac{1}{2}(t_2 - t_1)v. \end{aligned} \tag{6}$$

Measuring orbital periods You may have noticed that in order to find the companion masses, we relied on knowing the orbital period of the binary. We can have a go at the calculation during lab using the period reported in the discovery paper if there's time, but to measure it ourselves, we'll need to measure multiple eclipses, combining data taken by all lab groups. If we were able to continuously observe the system for many days in a row, you could imagine picking out the periodic signals pretty easily. However, usually our astronomical time-domain data is irregularly sampled and has gaps (due to daytime, weather, telescope scheduling, etc). We can use a tool called the **periodogram** to find repeating signals in unevenly-spaced data.

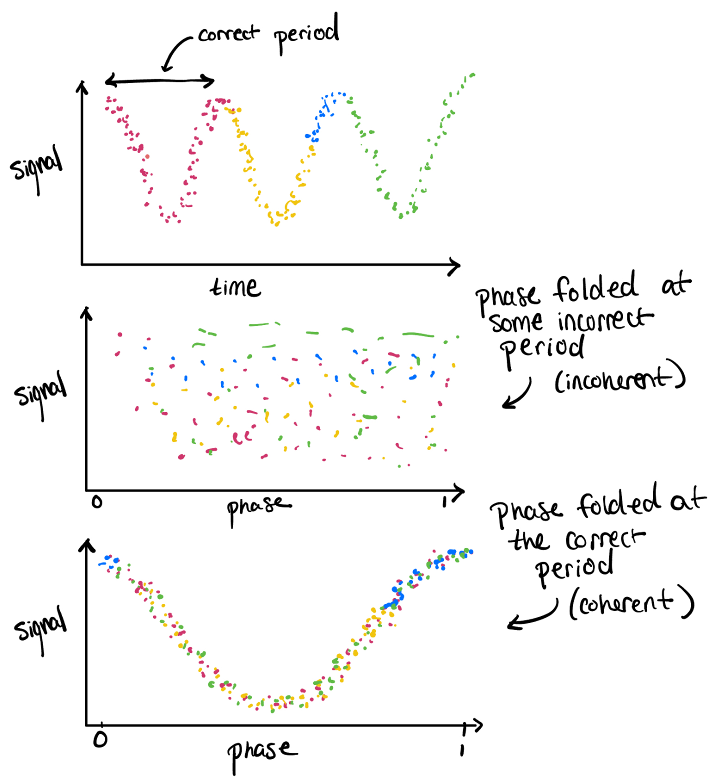

Essentially what this does is to assume a bunch of different trial periods for your data. At each trial period, we **phase fold** the data, meaning that at each point in time we ask what fraction of the way through the orbit does this time correspond to. With math, the equation for the phase is

$$\phi = \frac{(t - t_0) \bmod P}{P} \tag{7}$$

where we use the modulo operator, giving us the remainder of $(t - t_0)/P$. To find which periods yield a coherent phase-folded signal, a periodogram computes the "power" corresponding to each trial period, so that peaks in the power correspond to values likely to be the true period of the signal. The periodogram we'll use, "box least-squares" asks how well your phase-folded data matches a box function (see the example from the `astropy` BLS documentation page below) using least squares as the objective function. Astropy will do all of the phase folding under the hood for you; see the example analysis notebook.

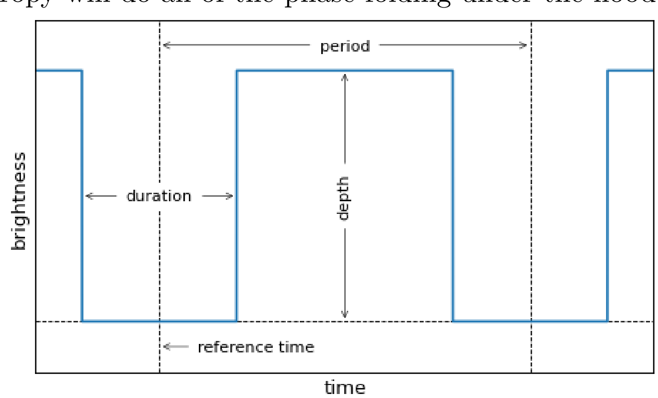

Observational considerations

There is a lot of procedural information about what we'll do during this lab in the assignment PDF itself. This is mostly meant as a supplement.

What is a CCD? Charge coupled devices (CCDs) are the detectors used by most modern astronomical instrumentation. You used one to record solar spectra during the daytime lab, but let's briefly add some detail to motivate our observing strategy and image processing. CCDs are silicon wafers that rely on the **photoelectric effect** to count the photons hitting them.

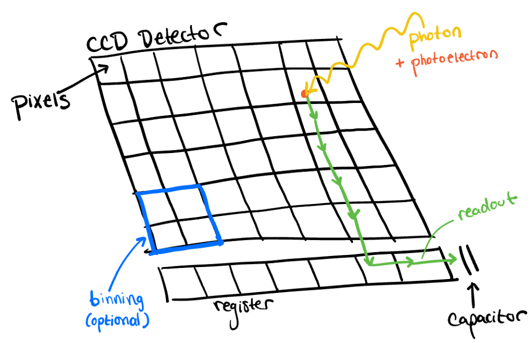

They consist of a grid of "pixels" separated by a voltage lattice such that electrons want to sit in them rather than between them (you can imagine charged wires). When a photon strikes a pixel, it liberates a photoelectron across the silicon band gap so that it may move around the device. At the end of an exposure, liberated electrons are shifted down the CCD row by row to the "serial register," then "read out" by a capacitor at the end of the register, which counts how many photoelectrons were present in each pixel, and therefore how many photons struck each pixel during the exposure.

Picking a good exposure time Our CCD measures and reports the "counts" of photoelectrons received from each pixel. However, only so many electrons are available to participate in the photoelectric effect, which means that once we reach very high counts, our detector "saturates," and adding photons will not increase the recorded counts. This is bad, since it means we won't be able to distinguish fluxes in this regime (there is slightly more nuance here—saturation can involve both overflowing pixels and limits set by the analog to digital converter that measures charges on the capacitor). Thankfully, we can avoid saturation by reducing our exposure times. We can tell that the number of photons collected equals the number of photoelectrons counted when the recorded counts are related linearly to the exposure time. Before saturation, the detector response becomes nonlinear (see the counts vs exposure time sketch), which we also want to avoid. We'll tune our exposure times so that we get detector counts in the linear regime from our target star as well as all of the reference stars.

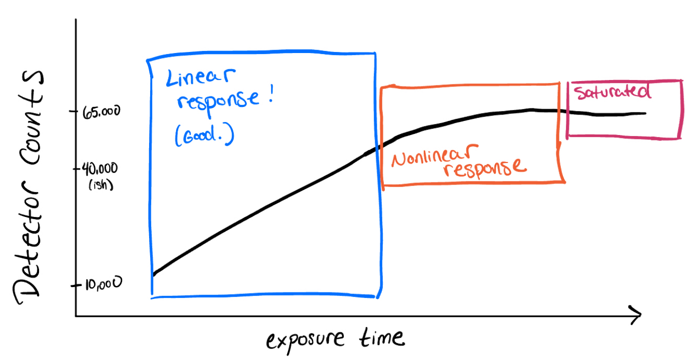

Another consideration in our exposure times will be whether to bin the CCD pixels. As an example, 2x2 on-chip binning is represented by the blue square in the CCD schematic above. Co-adding the counts from multiple pixels effectively creates a detector with fewer pixels, where each effective pixel collects a higher number of counts in a shorter time. This reduces our spatial resolution, but also lets us read out the CCD more often and with lower noise. Since atmospheric seeing will already make our stars many pixels wide, it might be a good idea to bin pixels for this lab.

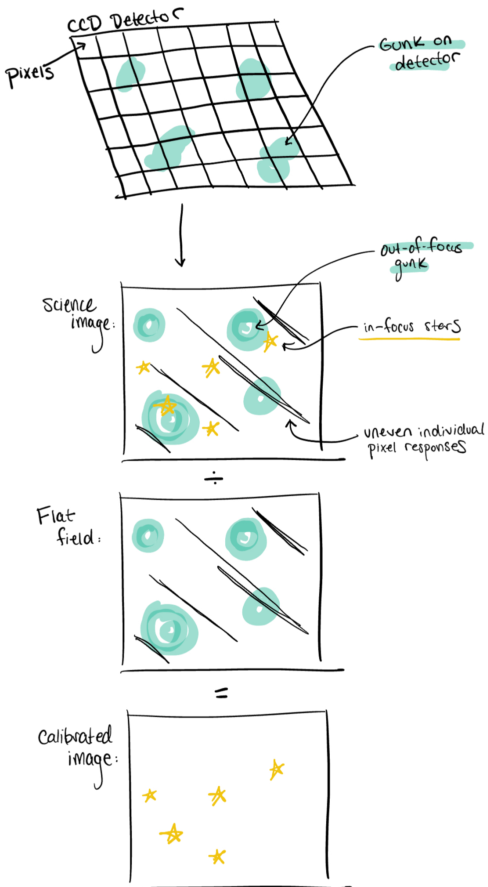

Flat-fielding When we image the sky onto a CCD chip, the pixel-to-pixel sensitivity is nonuniform. Foreground out-of-focus dust sitting on the detector also changes the sensitivity of certain pixels. Flat field images uniformly illuminate the CCD, for example, using the twilight sky, and capture each pixel's sensitivity. The flats are normalized so that the typical pixel value is 1, then the normalized flats are divided out of science images. This way, less sensitive pixels get "brightened" and more sensitive pixels get "dimmed." Note, in AstroImageJ, we'll create a "master flat" using the median pixel values from 10 flat field images (using the median rather than the mean so that we are robust to outliers like stars peaking into the flat at twilight).

Aside: other imaging data reduction Besides flat fielding, a few other steps are often taken to calibrate images from CCD detectors. We won't worry about them too much in this lab.

1. "Zeros" or "bias frames," are taken by reading out a CCD after zero exposure time. The resulting image is subtracted off of future science images to remove characteristic electronics-related noise.

2. "Dark exposures" are longer exposures with no source illuminating the CCD. This allows us to calibrate for "dark current noise," where electrons are created in the detector even in absence of the photoelectric effect. We can change settings in MaxImDL, the camera control program at the telescope, to automatically take a dark exposure at the beginning of our image series. So for example, if we want a series of 60s exposures, they will be preceded by a 60s exposure with the shutter closed.