Extrasolar Planet Detection

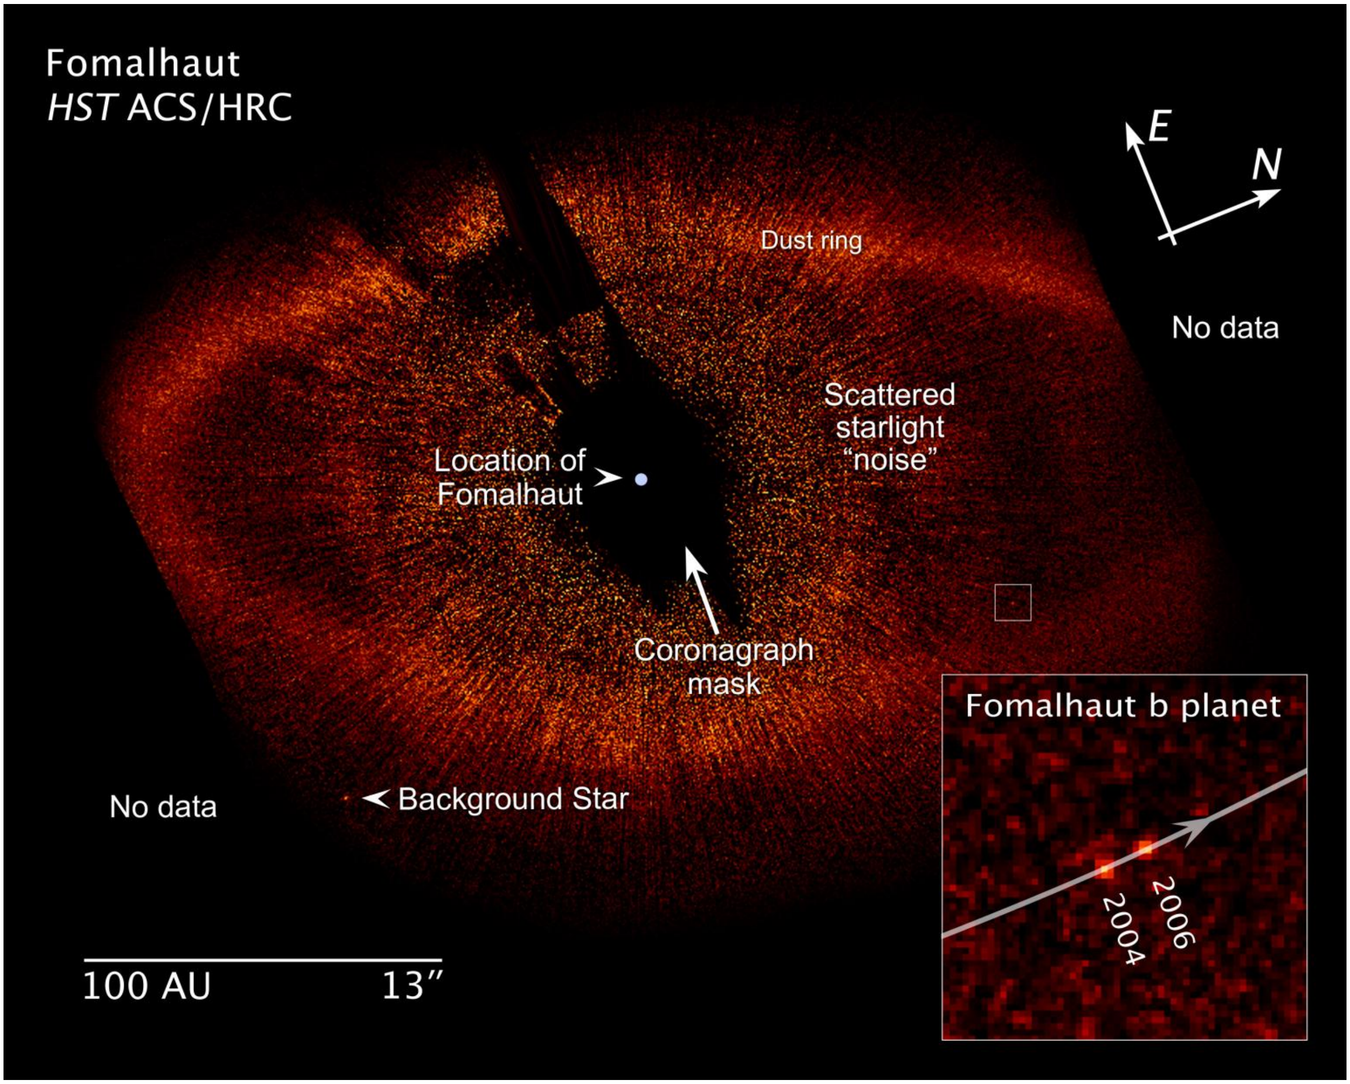

Introduction:

As of February 20, 2013, 861 exoplanets—planets that orbit stars other than our own Sun—are known to exist. Additionally, at least 128 stars are known to have more than one planet in orbit about them: more than 100 Solar Systems. Recent results from the Kepler space mission have yielded an incredible 2,000 more candidate objects that await confirmation. In fact, the Kepler team has announced that it is confident we will, in the observations to come, find a planet like the Earth, orbiting a star like the Sun, by the end of this year!

The first confirmed detection of an exoplanet about a solar-type star came in 1995[1] when the Swiss astronomers Mayor and Queloz announced the discovery of 51 Pegasi b: a planet like Jupiter, but found so close to its star that it only takes 4 days to complete one orbit. In the 15 years since the initial announcement, the discovery rate has accelerated, as well as the number of techniques used to find these planets, but the primary method for planet detection remains radial velocity measurements of their parent stars. It is this method that we will explore, using actual data collected from astronomers at Lick Observatory in the San Francisco Bay Area.

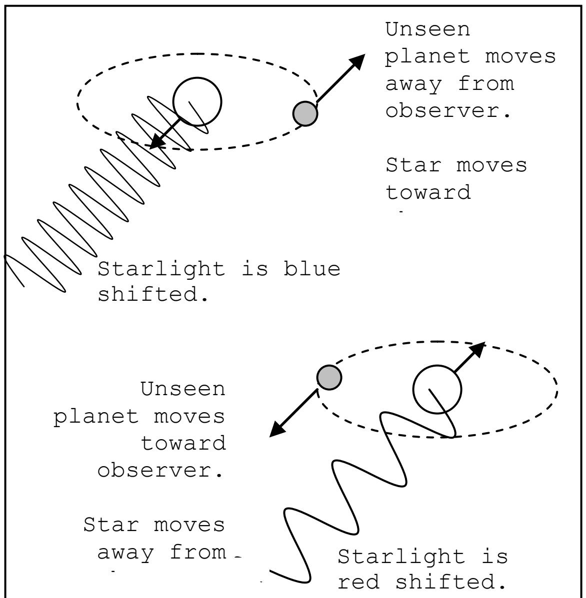

Figure 1: Doppler Shift due to motion of a star

Consider two objects: a star and its planet. While we ordinarily say that a planet orbits its star, rather than the other way around, in reality both the star and the planet orbit about a common center of mass. As a result, the star actually executes a small orbit. If the planet is far away and is much less massive than the star, the effect is negligible; the motion of the Sun caused by the Earth is slower than you can walk. However, a massive planet close to its star can cause significant acceleration. Moreover, when a star moves towards us, its light will be blue shifted slightly. On the other side of its orbit, when the star moves away from us, its light will be red shifted (Figure 1). The amount of shift seen in the light from a star is directly related to its speed, and that speed is related to the mass of what is "pulling" the star back and forth. By measuring the change in light—this stellar "wobble"—we can measure the mass of the unseen companion. To date, astronomers have been able to detect planets with masses ranging from about twice that of the Earth to several times that of Jupiter in this way.

This morning's exercise is based on data collected by the major U.S. team of astronomers using this method to find planets. The data points are velocities measured using this Doppler shift technique (for an example, see Figure 2). The software, *Systemic*, is publicly available at http://oklo.org/, and is written by astronomer Greg Laughlin at UCSC. In fact, you can download the software and analyze datasets on your own and join the hunt for planets—perhaps you will discover a new planet in the continually growing exoplanet database!

---

[1] Note that several planets were discovered in the late 1980's to be orbiting pulsars—the remnants created by Type II supernova explosions. The detection method (period perturbations in the pulsar's rotation) and the nature of the systems are quite different from the exoplanetary systems comprised of solar-like stars and (presumably) Jovian planets.

OVERVIEW:

Starting the Program:

You will use a program called *systemic.jar,* developed independently by astronomers at the University of California at Santa Cruz. The systemic console will appear with…a startling array of options. While we won't use all the possible widgets, I encourage you to explore the options.

Selecting your First System:

1. Let's start with a system that is fairly straightforward. Use the File…Choose System menu option to to select the data set called *14Her_B06K_3datasets*. The data correspond to continuous observations of a star called 14 Herculi, and display a very clear Doppler shift (see Figure 2); the data are collected from three different observatories. The vertical axis marks positive velocity (motion away from the observer) and negative velocity (motion toward the observer). The horizontal axis is marked in Julian Days, which are a unit of astronomical time using the mean solar day but ignoring calendrical.

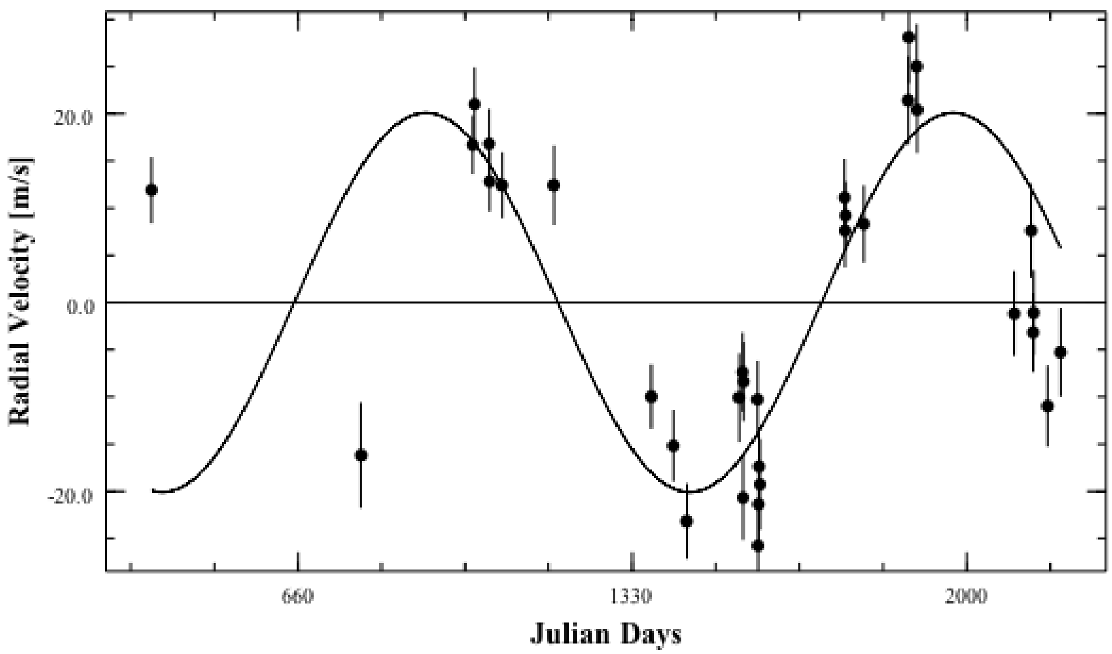

Figure 2: Sample Radial Velocity Curve. The vertical data are velocities in m/s for a star as it is pulled back and forth by an unseen (planetary companion). The horizontal axis is time measured in a special astronomical unit called a Julian Day, which are 24-hour periods measured without reference to calendar month or day.

2. There are several ways you can fit a planet to this data. The most straightforward is to click the square box under the radial velocity panel. This will open a menu with several options, starting with the orbital period of a planet (initial value is 100 days). By moving the slider bar back and forth, you can change this period; the goal is to try to match the overall shape of the data with a curve that runs through most of the points. Eyeballing the Herculi data, a period of about 1600 days can match the general trend of peaks and valleys.

3. A planet's orbit is, of course, more complicated than an arbitrary circular orbit with a given orbital period. You can also change the following:
   a. Mass: affects the maximum radial velocities
   b. Mean anomaly: This is an indication of where the planet was in its orbit when the observations take place. A mean anomaly of zero corresponds to the planet moving with maximum radial speed relative to the Earth. By convention, if the Earth is somewhere up the page, then MA = 0 means that the data collection started when the planet was at the 3 o' clock position. You can see the location of the MA by checking the "orbital view" box, and zooming out as necessary. A small circle on the orbit indicates the angle of the mean anomaly, measured counter clockwise from the 3 o'clock position.
   c. Eccentricity: This is a measure of the orbit's elongation. A circle has e = 0, with larger values of eccentricity indicating larger degrees of ellipticity (up to a maximum value of 1).
   d. Longitude of Perihelion: Once an orbit is eccentric, the position of the planet's closest approach, as seen from the Earth, can greatly affect the shape of the curve fit. The LoP is the angular location of this closest approach (technically called periastron for a star other htan the Sun).
   e. Velocity offsets (box above the lines for each planet entry): Velocity offsets are corrections to the radial velocity measurements due to each observatory and its instruments—hence in this case you'll see three different slider bars.
4. Spend a few minutes attempting to fit a single planet with an orbital period of about 1600 days to the data, varying the above orbital characteristics, and noting how the orbital view changes. As in previous exercises—and the case with much of astrophysical modeling—we can measure the quality of our fit using statistical methods. In this case we'll focus on the chi-squared value, which measures how well our line fits the data, with extra weight given to data points with relatively small error bars. A great fit will have a chi-squared value of 1; in today's practicum we'll aim for anything under about 10. You can also note your progress to a good fit by watching the residuals in the upper right. These points reflect the difference between your fit and the actual data. A perfect fit will have all points on the horizontal line.
5. When you've done as well as you can with one planet, it should be clear that it won't be enough. There is simply no way you can hit most of those points with one curve. Sharp-eyed students will also note a slight asymmetry in the sinusoidal curve, which is a tell-tale sign that another planet is lurking in the data. Rather than trying to fit all this by hand, we can move on to more sophisticated analysis tools, in particular the *periodogram*. A periodogram measures the importance of different periodicities in any data set. If you have any experience with Fourier analysis, a periodogram essentially pulls out the dominant sine curves that could be used to construct the data.

   Clear your previous work by unchecking the left-most square button and go to the View menu and select the *Periodograms* option. A window will appear that shows a logarithmic plot of the periods with their corresponding importance (measured by the "power" value—the higher the power, the more important that particular orbital period is to the data set). In this case, you should see that there is a period of 1,716.5994 days with an extremely high power. You can add a planet to your fit window by highlighting this power and clicking the '+' button.

6. And now the fun of semi-automation! With your period selected, click the circular buttons next to the Mean Anomaly, Mass, Eccentricity, LoP and Velocity Offsets boxes. Clicking the circle activates an automated optimization feature that will attempt to minimize the variance of your fit with the data. You may need to click the circle a few times to converge to a best value. Note what happens to your fit each time. When you're done, you should have a much better fit with a chi-squared value of about 15 or less.
7. To finish off the optimization, we can choose to "polish" the data, which does a global optimization. Be careful though: when polishing multiple planetary fits, the quality of fit can be quite good, but the

resulting planetary system can be a mess (see more below). Polishing also depends critically on your fit being close to the best answer—if you blindly polish your data, you can get into wildly divergent solutions. To polish your data, click the square button on the left of each box and then head up to the top of the console and hit polish. Watch what happens to your fit.

8. That's better, but we're still not hitting all the data. Go back to the periodogram window. Below the first window is a periodogram of the residuals: Basically the software takes the orbital fit you just made, subtracts it from the data, and runs another period analysis on the results. You should see a hugely important period at 3,270 days. Add the planet and follow the above procedure to get as good a fit as possible. You should get a two-planet solution with a chi-squared value of less than 2.
9. The result of the above procedure yields a fairly good fit—low chi-squared and a reasonable looking planetary system. Hitting update on the periodogram of the residuals yields no peaks above the "False Alarm Probability" lines (the horizontal dashed lines in the periodograms), and so we're done. Or are we?
10. No, we're not. Once we have a fit, we need to see if it makes sense physically. To do this, we can evolve the system using Newton's gravitation and see how the orbits change. Ideally we'd need to run the simulation for billions of years, but we have less time than that, so the standard simulation preset is for 1,000 years. Use the View…Orbital Evolution menu to bring up the simulator. About the only thing you'll want to change is the simulation time and the output time, depending on the system (having three planets can slow the calculator down, for example). You can select to watch how the semi-major axes change with time in the display box, and the window to the right will show the plot and the orbital configuration evolving in time. Finding that the system is stable in this way is telling, but not definitive, but finding instability tells you that, while your fit hits all the data points, the system is not physically realistic.

    Sadly, my fit was not particularly realistic. This is most likely due to extremely eccentric orbits created from the global optimization procedure, and you could go back and try setting eccentricities to zero and adding another planet (or, if you're really daring, try one of the more advanced options on the console).

PROCEDURE:

Now that you've run through a fairly straightforward example, you should spend the rest of the time in the lab fitting several other data sets, following the above (semi-automated) procedure. Keep an eye on the chi-squared value and the quality of fit, to be sure—but also test each system for stability. Try fits for at least four systems, including *upsand_4datasets_B06l*—the very first multiple system discovered, Upsilon Andromeda. Hiding in the data sets are the Galilean satellites of Jupiter and our own solar system, though they are particularly tricky.