

Measuring the Transit of an Exoplanet

Take Data: October 26 – Oct 31 (and possibly later)

Report Due: Thursday, November 8

Text Reference: Chapter 10

Purpose: The discovery of planets that orbit other stars is one of the most exciting recent discoveries in astronomy. One powerful technique to discover and determine the properties of exoplanets is to measure the dimming of a star as a planet transits in front of it. Such measurements can be carried out with modest-sized telescopes and we will do so in this lab.

I. Preparation for Observing:

The observing for this lab is time-critical since your chosen transit occurs at a specific time and ideally we want observations before, during, and after the transit. Thus, it is very important that you show up for your lab period on time and well prepared. This will require some collaborative effort with your lab-mates *before* the lab period. Since the transits typically last from 1 to 3.5 hours, the two groups observing that night should work together in choosing the target star and taking the data.

1. Use the Exoplanet Transit Database website (<http://var2.astro.cz/ETD/> – also linked to the class website) and use the menu near the top to go to the *Transit predictions* page. Select the date of your observing night. You will probably be asked for the east longitude (285.5 degrees) and latitude (40.5 degrees) of our observatory. The page will then display the transits happening that night for our location. In collaboration with your lab-mates, chose the star that you will observe using the following criteria. You may have to compromise on some of the criteria – life is rarely perfect! In your report you should discuss the reasons for selecting your target.
 - a. The transit occurs during your period. Note that EDT = UT – 4 hours; thus, 23:00 UT is 7:00 PM EDT.
 - b. The star is bright enough to measure an accurate brightness with exposure times of about two minutes or less. This means $V < 13.0$ or so. Brighter is better.
 - c. The transit depth is 0.01 magnitude or larger. Even when the target star is bright, various systematic errors (such as the quality of our flat-fields and darks) will probably limit the accuracy of individual measurements to about 0.005 magnitude.
 - d. The altitude during the transit is (mostly) above 45 degrees. Observations at low altitude look through more air and are more susceptible to clouds and telescope tracking problems. Unfortunately, stars that cross the meridian during the transit (which is when they are highest in the sky!) also cause problems since the telescope needs to move over the pole when this happens to avoid running into the pier. This means that the guiding has to be re-established and with a different star, since the CCD field of view has rotated by 180°. It is usually possible to follow a star for

about 40 minutes after it crosses the meridian without the telescope running into the pier. If you have a candidate transit where this would be helpful, discuss it with myself or the TA.

2. Determine how you will point at your chosen target. If the name of the star is its number in the Henry Draper (HD) catalog, then you are in luck. You can use *Find* in *The Sky* to locate and slew to your target. Most of the targets are too faint to be in the HD catalog. Then you will have to determine the Space Telescope Guide Star Catalog (GSC) number of the star using its RA and Dec and the finding chart provided by the Exoplanet Transit Database (click on the name of the star on your predictions page). Start *The Sky*, set the time and date to the beginning of your lab period, and use the *Move To* feature in the *Orientation* tab to go to the epoch 2000 right ascension and declination. This centers that star, so start zooming in. Keep zooming in until you see the “footprint” of the main CCD and guide CCD of our camera. You should then be able to match up the stars in the finder chart with those in *The Sky* and identify your target. (Turning on the orientation tool can be handy as its end often points to the location centered, i.e., your star). Click on it and record its name (probably its number in the GSC).
3. Examine the stars displayed in *The Sky* around your target. Using the camera footprint, check that there is at least one other star in the main field of view that is similar in brightness to your target (or somewhat brighter) and so can be used as a comparison star for doing relative photometry. If no such star is available, you will have to choose a different target. Then decide what star you will use to guide with by moving the field-of-view indicator around. Remember that the guider orientation is 90° if you are observing east of the meridian and 270° if you are observing west. With the guide-star in the guide CCD, your target and comparison star should be not too close to the edge of the field of the main CCD. Good guide stars are about 10th magnitude or brighter, but you may have to use a fainter one because of the small field of view of the guide CCD. If conditions are good (clear, dark skies) then stars as faint as 11th magnitude will work. If no acceptable guide star is available, you will again have to choose a different target. Finally, chose a 7th to 8th magnitude star within a degree or so of your target with which to focus the CCD.

II. Acquisition of the Data:

In broad outline, the procedure is to set up the telescope and camera, initialize the telescope pointing, focus the camera near your target star, set up the autoguider on your chosen guide star, and then take a continuous series of images for the expected duration of the transit. The lab is written to take flat-field data at the end of the observing, but it can be taken at the beginning if time permits. Most of the steps are similar to those in the previous labs, but reminders of some of the details are provided below. Do become familiar with the steps before your observing session since getting set up on your target tends to take longer than you expect and so is hectic. To obtain a good estimate of the brightness of the unobscured target star and because the predicted start and end times can be in error, try to start taking images at least 15 minutes before the predicted start of the transit and continue until at least 15 minutes after the predicted end. Exposures will use 2×2 binning (shorter readout time) and AutoDark subtraction (more convenient) unless otherwise noted. (Multiple dark exposures are not necessary because noise from the dark current is much smaller than from photon statistics.)

We will observe most of the transits through the open filter position (i.e., no filter) to maximize the amount of light. An exception is any target star brighter than about magnitude 7.5, for which exposure times are around 5 seconds or shorter. These short exposures suffer from brightness fluctuations due to atmospheric scintillation. Using the I band reduces these fluctuations (the index of refraction of the atmosphere decreases with increasing wavelength) and also slightly lengthens the exposure times due to the lower CCD efficiency.

1. Turn on the CCD system and start the *CCDSOFT* program, *IDL*, and *RUPhAst* as in previous labs. On the *Setup* tab in *CCDSOFT*, connect to the CCD and set the temperature 35 to 40 °C cooler than the ambient. Set up the *AutoSave* tab to save your images in a subdirectory named with the date and your initials in *Astrolab:/home/ph344/lab5* (e.g., *oct23tpjh*). Set the *starting number* to 1 and choose a reasonable filename prefix.
2. While the CCD is cooling, open the dome and turn on the telescope controller box. Get the telescope ready to observe. The stepper-motor controller on the observing desk is used to focus the telescope by moving the secondary mirror. Turn on the encoder readout above these controls with the green button. Use the focus motion buttons on the panel below to set the encoder to a reasonable first guess at a good focus: about 2.40 mm at an air temperature of 20 C, increasing by about 0.100 mm for every 5 °C *decrease* in temperature. (Though this temperature dependence has not been working well this year – a focus setting of about 2.35 mm seems to be working pretty well). If observing without a filter, then the initial guess for the best focus is different – about 2.55 mm.
3. Do the usual telescope initialization procedure: home the telescope, point at the zenith with the bubble levels, sync on a star near the zenith and to the right (west) of the meridian. Then slew to a bright star and refine the pointing. Center the star first in the finder with the joystick and then in the CCD field of view by taking 1-second exposures. Move the star in the CCD field using motion controls in *The Sky*. You can estimate the size of motions needed by remembering that the long axis of the CCD is oriented north/south and is about 32 arcminutes across. Don't forget to do a final *Sync* to save your pointing.
4. Slew the telescope to a 3rd magnitude or brighter star near your target star. Take a 1 second exposure and, if necessary, re-center this star and sync the coordinates.
5. Slew to your focus star and use the same procedure as in Labs 3 and 4 to focus the CCD using 1×1 binning, 5-second exposures, and a filename prefix of “focus”. Start with a full frame (*without* AutoDark subtraction to avoid the second long readout) and then click and drag a selection box in the *CCDSOFT* image display to define a subframe including the focus star and a few others. Take a set of 7 exposures reading out just the subframe (*with* AutoDark subtraction), starting with the focus 0.150 mm less than the estimated optimum value and increasing the focus value in steps of 0.050 mm between exposures. Set the focus to the best value. Remember to record the value and the temperature so that it can be included in your report.
6. Slew the telescope to your target star. Set *Subframe* off, binning back to 2×2, change the filename prefix to something other than focus, and take a 5 to 20 second exposure (depending on how faint your target and surrounding stars are). Identify the field around your target star to make sure that your pointing is correct. Use aperture photometry of your target and comparison star in this test images to decide on an exposure time that will keep

both at no more than 40,000 ADU (about 60% of the saturation value) at the center of their profiles. The cushion is to allow for higher central pixel values if the seeing improves. However, do not use an exposure longer than 2 minutes since this would result in loss of time resolution.

7. To set up the guider, select and move the red field-of-view indicator in *The Sky* to match the location of your test image – stars near the edge of the image are particularly helpful for this. Then use the telescope *Motion Controls* (in *The Sky*) to move your chosen guide star into the field of view of the guide CCD. This is a finicky process because of the small field of view of the guide CCD. If the star does not appear in the guider when you think it should be there, take another exposure with the main CCD and refine the location of the field of view indicator. Once you have acquired your guide star, do not immediately start guiding but instead do the next step.
8. Take a test image with your chosen exposure time to confirm that both your target and the comparison star are in the field of view (and preferably not in one of the vignetted corners of the field) and that the signal levels are what you expect. This also gets the autodark exposure taken, which is important for not disrupting the guiding, particularly for longer exposure times. This is a good time to set up a filename prefix for you transit data. Once you are satisfied with the field placement and exposure time, start guiding. The guider exposure times should be between 3 seconds, for brighter stars, to 6 seconds or so for fainter ones. Guider exposures shorter than 3 seconds tend to cause the telescope to fruitlessly chase rapid seeing-induced image motion. Exposures longer than 6 seconds make it harder for the guider to keep up with tracking errors and lengthens the gaps between images (see below).
9. Start taking blocks of 10 exposures. Use enough *Delay* between exposures so that the guider can take two images between the readout ending and the start of the next exposure. This will help the guiding settle down after the cessation of guiding during the readout of the main CCD. Continue taking blocks of images until at least 15 minutes after the transit is expected to end. Monitor the guiding to make sure that it is working OK (that clouds have not dimmed the star, for example). Remember to periodically move the dome so that the slit does not occult the telescope. Also check to be sure that the telescope does not run into the pier if the target star approaches the meridian.
10. Once you are done taking data for your transit, home the telescope and then manually point it at the inside of the dome. Using the procedure from Labs 3 and 4, take 7 well-exposed flat-field images of the inside of the dome with the white lights on. Use the 2×2 binning and Auto-dark subtraction.
11. Finally, cover the telescope and shut everything down. Remember to fill out the log.

III. Data Analysis (in IDL):

Data analysis can be done at any time, using the server computer (astrolab). Using VNC, you can connect to astrolab from anywhere and analyze your data.

We will use a customized version of PhAst, *RUPhAst*, to measure the brightness of the target star and one or more comparison stars in all of the images. The basic steps in the process are the following. a) Make a “flat-field” image from the your images of the inside of the dome. b) Read your science images into *RUPhAst* and use it to divide them all by the flat-field image. This removes the effect of vignetting in the camera and pixel-to-pixel sensitivity variations in the CCD. c) Read the “flattened” images into *RUPhAst* and use its aperture photometry tool to measure the brightness of the target star and one or more comparison stars in all of the images and to write the results in a file. d) Use a spreadsheet program (or some other tool of your choice) to calculate the difference in magnitude between your target and comparison star(s) and plot this *versus* time.

1. Copy all of your Lab 5 CCD images from the folder for your night in /home/ph344/lab5 to a lab5 folder in your home area and work on these copies. Also copy the file /home/ph344/lab4/phast.conf into your lab5 folder. This configuration file initializes certain quantities in *RUPhAst*, most notably the CCD gain and readnoise. Examine the contents of this file.
2. Create a subdirectory in your lab5 folder called flat and move your dome flat-field images into it. Delete any extraneous images (i.e., focus images, set-up images, ...) from your lab5 folder. Also delete any images with NOAUTODARK in the name (these are images before the automatic dark subtraction). The only images in your lab5 folder should be those of your science field. This will make automatic loading of the many images possible.
3. Start IDLDE and then change the path to the folder with your flat-field files using the IDL *cd* command:

```
cd, "/home/yourusername/lab5/flat"
```

Recall that the quotes are important, as is the comma. Use the IDL program *mkflatru* to combine your images into an average flat-field image, scaled to be close to one at its center. Call the output image something like FlatO.FIT. When combining the images, *mkflatru*, measures the average signal in the central half of each image, divides the image by this number, and then takes the median of all of these normalized images to eliminate cosmic-ray events (and stars, if these were twilight-sky images).

4. At the IDL command line, *cd* back to your lab5 folder (*cd, “..”* will do this; the *..* is a shorthand for the folder above the current one). Start *RUPhAst* and agree to the request to create an “output” directory. This directory will be created in your lab5 directory – we will make use of it shortly. Examine your final flat-field image with *RUPhAst* to check that it looks reasonable. If it does not, return to step 3 and determine what went wrong.
5. Remove all of the images loaded into *RUPhAst* (such as the final flatfield you just looked at) with *File -> Remove -> Remove all images*. Then use the menu item *File -> Read -> Read FITS directory* to load the images in your lab5 folder. You should be able to just select “.” in the navigation window that pops up (this is shorthand for the current directory). Adjust the brightness and contrast using the *Color Mouse Mode* and then examine the images. Are they all well-aligned with each other? Do they all have satisfactory image quality? Delete any bad images from your lab5 folder and then again remove and re-read the directory of images. Note the images removed and the reasons in your report.

If the position of the stars in your images shift part-way through the sequence, you will need to put the sets of images with the same alignment in different subdirectories so that you can read and photometer them easily. However, wait to do this until after dividing the images by your flat-field image in the next step. For the now, just note which ranges of image sequence numbers are aligned with each other. Changes in the positions of the stars will occur if you need to reacquire the guide star, say because of a passing cloud or because the star crossed the meridian and the telescope had to be moved to the other side of the pier.

6. Use the *Pipeline* -> *Batch calibrate* to bring up the “Batch image processing” window. Make sure that the “Current images” radio button is pushed and then push the *Flat* button and select your flat-field image created in step 3. Push the *Start* button – after some time a pop-up window will announce that batch processing is complete. Push the *Done* button to dismiss the batch processing window. The processing will have written the flattened images in the *lab5/output/images* folder. They will be named using their original sequence numbers. Remove your original images from *RUPhAst* and then load those in the *lab5/output/images* folder. The sky level in the flattened images should look reasonably uniform (at least, more uniform than the raw images). If it does not, determine what you did wrong and repeat steps 5 and 6.

If you need to divide your set of images into blocks of images where the stars are aligned, create subdirectories in your *lab5/output/images* directory and move the blocks of images into these subdirectories. Then read one block at a time into *RUPhAst* and do step 7 on each block. Then combine the individual photometry files to do steps 8 – 10.

7. Change back to the *ImExam* Mouse Mode. Identify your target and comparison star(s) and left-click on your target star to bring up the aperture photometry tool. Look at the FWHM of the star. Set the aperture radius to twice this value (rounding up to the nearest integer). Remember to hit carriage return to actually enter the new value. Set the inner sky radius to a value 5 pixels larger than the aperture radius and the outer sky radius to a value 10 pixels larger than the inner sky radius. You should actually set them in the order outer-, inner-, and then aperture-radius, since *RUPhAst* insists on having these values in increasing order. Push the button to choose the *Centroid* centering method and then bring up the *Photometry settings* window. Select output units: *Arcsecs-Magnitudes*, magnitude system: *Instrumental*, and calculate photometric errors?: *Yes*. Check that the CCD Gain is set to $1.8e^-/DN$ and the read noise to $15.5 e^-$. Push the *Apply Settings* button to apply the new values and dismiss the window.

Back in the aperture photometry tool, push the *Write results to file...* button and pick a reasonable name to store your photometry results in. Then push the *Do all* button. The radial profile plot will successively show the profile of your target star in all of the images. Watch to make sure that the photometry has been successful in every image (in particular, that the star does not change position so much that the centering algorithm fails). Then click on the next star you wish to photometer in the image. This will create an extra entry in your output file, but it is easily ignored or deleted. Push *Do all* again. Repeat this


process for every star you wish to measure. Once you have photometered your target and comparison stars, close the output file by pushing the button.

8. Examine the file with your photometry. The first column is the image in which the star was measured, with 0 being the first image in the *RUPhAst* stack, 1 the next, *etc.*. The next column is the Universal time of the middle of that image, calculated by adding half of the exposure time to the start time in the file header. There follows the name of the filter, the (x,y) of the centroid of the star, and the radii of the apertures used to do the photometry. Finally come the results of the photometry: the measured brightness of the sky in instrumental magnitudes per square arcsecond, the instrumental magnitude of the star, the uncertainty in the stellar magnitude, and the FWHM of the stellar profile. The instrumental magnitude is calculated with $m = 20.3 - 2.5 \log ((C/t))$, where C is the number of counts in the aperture and t is the exposure time. The value 20.3 was chosen to make the instrumental magnitudes be approximately correct for images taken in the V band with our telescope.

IV. Data Analysis (in Excel or another program):

Read the file with your photometry into Excel, OpenOffice Calc, or any other program of your choice. With Excel this is done with the *File -> Import* menu item. With Calc, use the *File -> Open* menu item and select file type *Text CSV* (which means that the file must end with .txt or .csv). In either case, this brings up a file import tool. Choosing a *Fixed Width* import works best for the photometry output files, though with Calc you will have to set the boundaries of the columns yourself.

9. Plot the magnitude of your target star and comparison star(s) versus time. Include the plot in your lab and comment on possible causes for any trends and features.
10. Calculate the difference between the magnitude of your target star and your comparison star(s). Plot these one or more differences versus time. If you have results from more than one comparison star, do they agree (you may need to add a constant to one, or more, of the differences to make the comparison)? This is one way to check whether a comparison star is itself a variable and also gives an indication of the level of any systematic errors present in the photometry (as a function of position in the image, for example). Produce a final best estimate of the *light curve* (magnitude difference vs. time) for the transit, by averaging your results from different comparison stars (if you have more than one).
11. Does your data show the dimming of the star caused by the transit of the planet? Estimate the depth, duration, and mid-time of the transit. Be as quantitative as you can and include your estimate of the uncertainties in these quantities. Plotting a running mean of 3 or 4 points can aid in making your estimates. Another approach is to use the simple function of time shown below and determine the parameters by least squares or by adjusting them by hand to achieve the best fit.



12. How do your values compare with those predicted by the Exoplanet Transit Database? How do they compare to results found by other observers (click on the name of the star in the long list on the left-hand side of the main ETD page to see the results reported to the database)?

13. The depth of the eclipse is a measure of the ratio of the cross-sectional area of the planet to the cross-sectional area of the star. Convert your measured depth from a difference in magnitudes into a ratio of the brightness of the star during and outside of the transit. Use this to find the ratio of the diameter of the planet to that of the star assuming that both are spherical and that the disk of the star is of uniform brightness. What is your uncertainty in this ratio? Compare your value to the ratio calculated from the stellar and planetary diameters listed for your system in the Extrasolar Planet Encyclopedia (<http://exoplanet.eu/catalog/>).