Lab #4:
Detecting a transiting exoplanet using CCD Images

10/12/2010

Your report is due on November 02, 2010 at 5:59 PM PDT.

1 Overview

One of the most exciting new topics in the field of astronomy is extrasolar planets ("exoplanets"). One of the ways we can detect and gain information about exoplanets is by observing planets that *transit,* or pass in front of their stars from our point of view. During the transit, the light we receive from the star is dimmed by up to 1-2%, providing critical insight on the properties of the planet.

In this lab we will use CCD images from the 1m Lick Observatory telescope to try and detect the transit of an extrasolar planet and determine the planetary radius. In practice, we will use a "differential aperture photometry" technique to produce a *light curve* (a plot of stellar flux as a function of time) depicting the dip induced by the planetary transit (see Figure 1). This document presents step-by-step instructions to prepare, conduct, and analyze astronomical observations with a CCD camera.

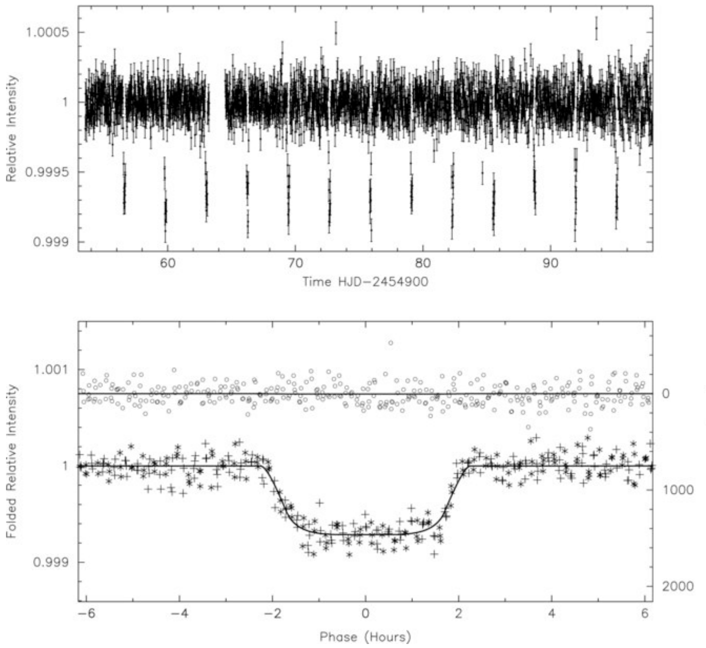

Figure 1: (*Top*) Light curve for KOI 4 obtained by the space-based Kepler mission. Notice how the star dims by less than 0.1% (!) every ~3 days. (*Bottom*) A "phase-folded" light curve, in which all eclipses have been shifted so as to be perfectly superimposed onto one another (crosses and star symbols). The transit for this system lasts about 4 hours. The open symbols represent the lightcurve of the system in opposition of phase from the transit (i.e., when the planet disappears behind the star). From Borucki et al. (2010).

1.1 Key steps

The key steps are:

1) Select appropriate targets for the Lick 1m observations (§3);
2) Gather Lick 1m observations (Oct 15, 17 and/or 18, weather permitting);
3) Read FITS files from either telescope and display the data using IDL (§4.1);
4) Apply systematic corrections including dark current and flat field corrections (§§4.2,4.3);
5) Re-align all images so that each star lies at the same pixel location in all images (§5);
6) Use an automated routine to measure the flux of the target and reference stars in all images and generate the lightcurve of the target star (§6);
7) Determine whether the transit is detected, estimate its depth and, consequently, the radius and average density of the exoplanet, placing constraints on its nature (§7).

1.2 Schedule

This is a three-week lab. Step 1 above is to be done today, in preparation for step 2 later this week (see class schedule). The remainder of the activities will start next Tuesday. You should aim at having completed step 5 above before the last week of the lab.

1.3 Recommended reading

Howell, *Handbook of CCD Astronomy*, chapter 4 & 5, especially §§5.1, 5.2 and 5.4

2 The Observatory

We will be using the 1m Nickel telescope[1] at Lick Observatory (see Figure 2), located on top of Mt Hamilton, about 12 miles East of San Jose. The Lick Observatory is administered by the University of California Observatories (UCO), along with the twin 10m Keck telescopes in Hawaii.

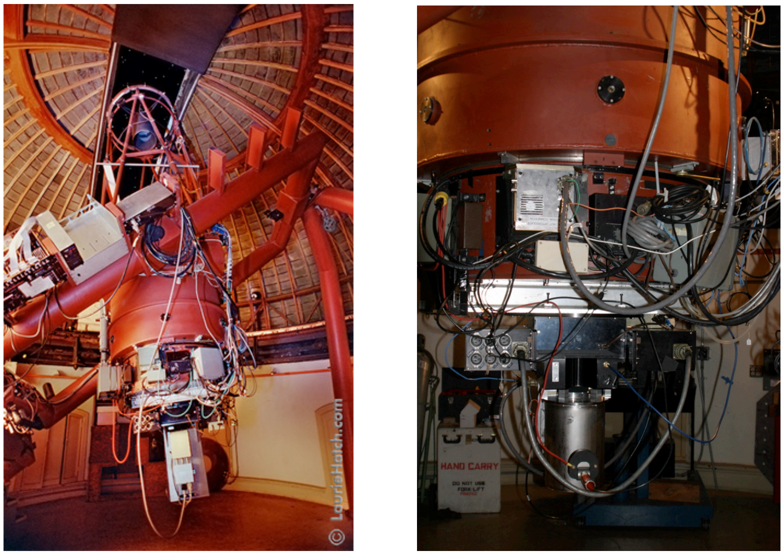

Figure 2: (*Left*) A picture from inside the dome of the 1m telescope at the Lick Observatory. (*Right*) A picture of the Direct Imaging Camera we will use during this lab.

---

[1] http:// http://mthamilton.ucolick.org/techdocs/telescopes/Nickel/intro/

The telescope is equipped with a 2048 × 2048 pixel CCD. The resultant field of view is approximately 6.3 × 6.3 arc minutes. Details of the facility are listed in Table 1.

Table 1: Nominal observatory and instrument properties. A x2 binning factor is generally used.

Site	Mount Hamilton, CA
Geographic location	37:20:24 (N) -121:38:43 (E)
Primary mirror	1 m (diameter)
Camera	Direct Imaging Camera
Detector	thinned Loral CCD
CCD format	2048 × 2048, 15 μm pixels
Pixel scale	0.184 arc seconds/pixel
Binning factors	x1, x2, x4
Field-of-view	6.3 arc minutes (square)
Optical Filters	*U, B, V, R, I*
Readout time	1’16” (x1 binning) 21” (x2 binning) 6” (x4 binning)
Gain (e^-/DN)	2 (x1 binning) 1.7 (x2 binning) 1.7 (x4 bining)

Use the Nickel telescope and Direct Imaging Camera User’s to get familiar with the properties of the instrument, the observing procedure and the software that you will use during the observations. When you are observing, you will be in charge of filling out the nightly log. Make sure to write down (in the log or a notebook) as much as you can to make sure any relevant detail is available once you analyze the data later on. Once you come back in the lab to reduce the data, it is too late to try and remember specific details of what happened during the observations!

3 Target selection

Before we observe, we have to decide which target we will point the telescope at. In this lab, our goal is to obtain a lightcurve for a known exoplanet, as opposed to searching for a new transiting exoplanet. This would require us to monitor many thousands of stars in order to detect just one transiting system! For such endeavors, dedicated telescopes (and space missions such as the Kepler mission[2]) are built and used exclusively for this purpose.

To decide which target we will look at, we will use the Exoplanet Transit Database[3] (ETD), which contains a large amount of information on already known transiting exoplanets. Click on the “Transit Predictions” link and, specifying the observatory location (see Table 1), generate a complete list of transits that are observable from Lick in the upcoming nights (see Figure 3 for an exemple of a “results” page). Pay attention to the fact that this website, as many professional astronomy catalogs and databases, uses Universal Time (UT) to avoid confusion with the local time that varies from one observatory to the next. While we are in Pacific Daylight Time, UT is

---

[2] http://kepler.nasa.gov/

[3] http://var2.astro.cz/ETD/

7h ahead of us. So sunset, which is around 6:30pm this week-end, occurs at 1:30am UT, with the date a day ahead (i.e., it is already Oct 13 in UT time!).

For this lab, we have been granted 3 nights with the Nickel telescope. Considering the time it takes to acquire a full dataset, everybody will share the same dataset, which we will obtain as soon as weather permits. To identify which transiting exoplanet(s) can be observed on each of the three nights we have access to the telescope, browse the list of predicted transit from the ETD database. You should pick the best candidate for each night based on the following criteria:

- The depth of the transit must be as large as possible to improve our chances of detecting it. Ideally, the selected targets will have depths of 0.01 mag (~1%) or larger.
- The target must remain high in the sky throughout the transit. Observing a star very low on the horizon introduces larger fluctuations in sky transparency and poorer image quality, which in turn introduces larger uncertainties in the measured fluxes. A common criterion in astronomy is to observe objects when they are above an elevation of 30°.
- The transit should occur in the first half of the night (remember that 7h UT is 12am PDT) to avoid staying up all night!
- The transit duration should be as short as possible to avoid changing weather conditions (such as fog rolling in during the night, a common occurrence on Mt Hamilton…).

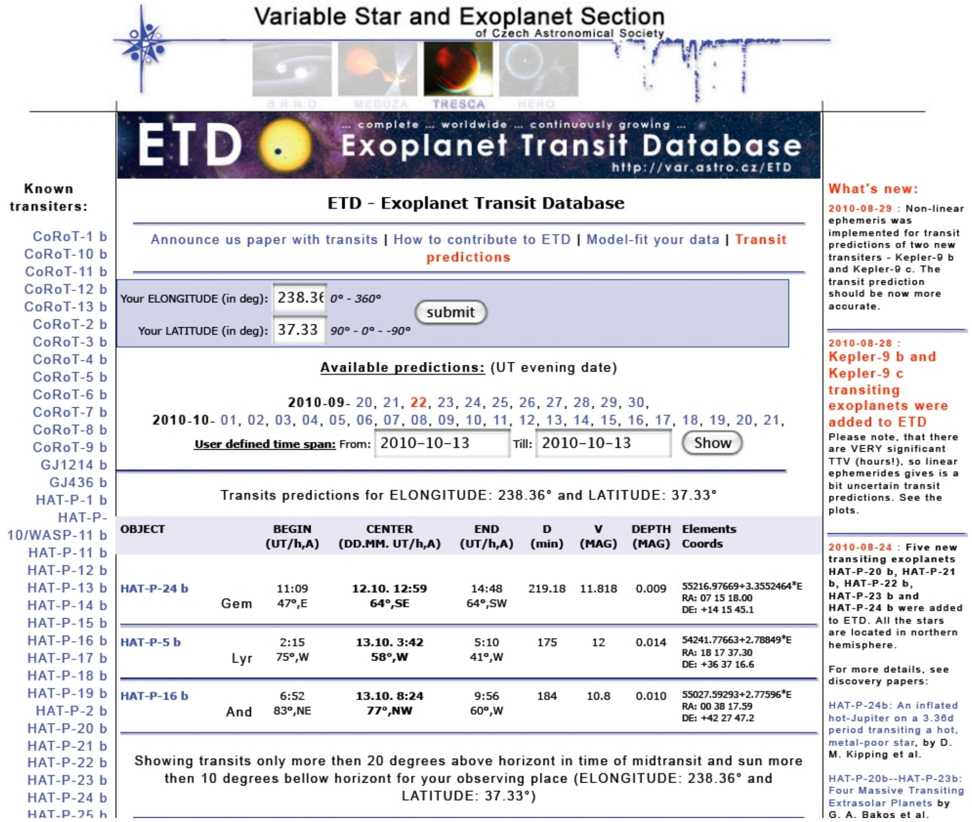

Known transiters:

CoRoT-1 b
CoRoT-10 b
CoRoT-11 b
CoRoT-12 b
CoRoT-13 b
CoRoT-2 b
CoRoT-3 b
CoRoT-4 b
CoRoT-5 b
CoRoT-6 b
CoRoT-7 b
CoRoT-8 b
CoRoT-9 b
GJ1214 b
GJ436 b
HAT-P-1 b
HAT-P-10/WASP-11 b
HAT-P-11 b
HAT-P-12 b
HAT-P-13 b
HAT-P-14 b
HAT-P-15 b
HAT-P-16 b
HAT-P-17 b
HAT-P-18 b
HAT-P-19 b
HAT-P-2 b
HAT-P-20 b
HAT-P-21 b
HAT-P-22 b
HAT-P-23 b
HAT-P-24 b
HAT-P-25 b

ETD - Exoplanet Transit Database

Announce us paper with transits | How to contribute to ETD | Model-fit your data | Transit predictions

Your ELONGITUDE (in deg): 238.36 0° - 360°

Your LATITUDE (in deg): 37.33 90° - 0° - -90°

Available predictions: (UT evening date)

2010-09- 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30,
2010-10- 01, 02, 03, 04, 05, 06, 07, 08, 09, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21,

User defined time span: From: 2010-10-13 Till: 2010-10-13

Transits predictions for ELONGITUDE: 238.36° and LATITUDE: 37.33°

OBJECT		BEGIN (UT/h,A)	CENTER (DD.MM. UT/h,A)	END (UT/h,A)	D (min)	V (MAG)	DEPTH (MAG)	Elements Coords
HAT-P-24 b	Gem	11:09 47°,E	**12.10. 12:59** **64°,SE**	14:48 64°,SW	219.18	11.818	0.009	55216.97669+3.3552464*E RA: 07 15 18.00 DE: +14 15 45.1
HAT-P-5 b	Lyr	2:15 75°,W	**13.10. 3:42** **58°,W**	5:10 41°,W	175	12	0.014	54241.77663+2.78849*E RA: 18 17 37.30 DE: +36 37 16.6
HAT-P-16 b	And	6:52 83°,NE	**13.10. 8:24** **77°,NW**	9:56 60°,W	184	10.8	0.010	55027.59293+2.77596*E RA: 00 38 17.59 DE: +42 27 47.2

Showing transits only more then 20 degrees above horizont in time of midtransit and sun more then 10 degrees bellow horizont for your observing place (ELONGITUDE: 238.36° and LATITUDE: 37.33°)

What's new:

2010-08-29 : Non-linear ephemeris was implemented for transit predictions of two new transiters - Kepler-9 b and Kepler-9 c. The transit prediction should be now more accurate.

2010-08-28 : Kepler-9 b and Kepler-9 c transiting exoplanets were added to ETD

Please note, that there are VERY significant TTV (hours!), so linear ephemerides gives is a bit uncertain transit predictions. See the plots.

2010-08-24 : Five new transiting exoplanets HAT-P-20 b, HAT-P-21 b, HAT-P-22 b, HAT-P-23 b and HAT-P-24 b were added to ETD. All the stars are located in northern hemisphere.

For more details, see discovery papers:

HAT-P-24b: An inflated hot-Jupiter on a 3.36d period transiting a hot, metal-poor star, by D. M. Kipping et al.

HAT-P-20b--HAT-P-23b: Four Massive Transiting Extrasolar Planets by G. A. Bakos et al.

Figure 3: Snapshot of the ETD website showing the list of transiting exoplanets observable from Lick tonight (Oct 12, 2010). For each entry, note in particular the beginning, central and final time of the transit (indicated in UT time, PDT is 7h behind UT) as well as the maximal "depth" of the transit (0.01 mag corresponds to 1% dimming).

In addition to the transiting planets listed in the main ETD database, explore also the "Kepler candidate" list, which includes many additional *candidate* transiting exoplanets, for which radial velocity measurements have not yet confirmed the planetary status of the obscuring body. In particular, consider KOI 217, 425, 802 and 931 for Oct 15, KOI 876 for Oct 17 and KOI 609 for Oct 18. Generally speaking, the Kepler candidates are fainter, hence more challenging, than other transiting exoplanets. Still, on some nights, they may represent the best choice for our lab.

Once each group has identified a "good" candidate for each of our three nights of observations, we will hold a lab-wide discussion to finalize our choice of targets. Since each transit lasts for 2 to 4h, we can only observe one transit per night, so it's important to choose wisely the best candidate!

4 First look at the data and calibration images

Summary of Section 4 Instructions:

- Read in all files and transform using `bscale` and `bzero` or `/fscale`
- Compute the bias and dark current rate for the Nickel CCD
- Generate master dark current images
- Create a master flat-field

4.1 FITS headers

CCD images are saved on disc in a standard astronomical format, known as FITS (Flexible Image Transport System) files. FITS files contain the image data, stored in binary format, together with some information that records the circumstances under which the data were obtained. Although the image data in FITS files are in binary format, each file also has an ASCII preamble that can be conveniently inspected at the Unix command line by typing, for example:

```
ugastro% fold image101.fits | more
SIMPLE  =                    T / NORMAL FITS IMAGE
BITPIX  =                   16 / DATA PRECISION
NAXIS   =                    2 / NUMBER OF IMAGE DIMENSIONS
NAXIS1  =                 1056 / NUMBER OF COLUMNS
NAXIS2  =                 1024 / NUMBER OF ROWS
CRVAL1  =                 2046 / COLUMN ORIGIN
CRVAL2  =                 2046 / ROW ORIGIN
CDELT1  =                   -2 / COLUMN CHANGE PER PIXEL
CDELT2  =                   -2 / ROW CHANGE PER PIXEL
OBSNUM  =                30399 / OBSERVATION NUMBER
IDNUM   =                    4 / IMAGE ID
IMTYPE  =                    0 / IMAGE READOUT GEOMETRY
AMPSROW =                    1 / AMPLIFIERS PER ROW
AMPSCOL =                    1 / AMPLIFIERS PER COLUMN
EXPTIME =            90.000000 / Exp time (not counting shutter error)
BSCALE  =             1.000000 / DATA SCALE FACTOR
BZERO   =         32768.000000 / DATA ZERO POINT
COMMENT = 'Real Value = FITS*BSCALE+BZERO'
PROGRAM = 'NEWCAM'             / New Lick Camera
VERSION = 'nickel_direct'      / Data acquisition version
TSEC    =           1241942680 / CLOCK TICK - SECONDS
TUSEC   =               753901 / CLOCK TICK - MICROSECONDS
DATE-OBS= '2009-05-10T08:04:40.75' / UTC DATE AND TIME OF OBSERVATION
DATE-STA= '2009-05-10T01:03:10.0' / START OF OBSERVATION
DATE-END= '2009-05-10T01:04:40.0' / END OF OBSERVATION
CAMERAID=                    2 / CAMERA ID NUMBER
EQUINOX =          2009.359985 / EPOCH FOR POCO POSITION IS CURRENT DATE
AIRMASS =             1.180000 / AIRMASS AT START OF OBSERVATION
```

```
INSTRUME= 'Nickel Direct Camera'
```

[…]

```
RA      = '        17:51:53.5' / RIGHT ASCENSION
DEC     = '        37:34:56.0' / DECLINATION
MPP     =                    1 / MPP STATE
HA      = '       -02:42:25.4' / HOUR ANGLE
OBJECT  = '            TrES-3'
TEMPCON =            23.400000 / CONTROLLER TEMPERATURE
NCSHIFT =                    0 / NUMBER OF CHARGE SHUFFLES
RCSHIFT =                    0 / NUMBER OF ROWS IN EACH CHARGE SHUFFLE
READ-SPD=                   80 / DCS READ SPEED
ROVER   =                    0 / NUMBER OF OVERSCAN ROWS
COVER   =                   32 / NUMBER OF OVERSCAN COLUMNS
TEMPDET =          -116.199997 / EXPOSURE START DETECTOR TEMPERATURE
TEMPDETE=             0.000000 / EXPOSURE END DETECTOR TEMPERATURE
GAIN    =                    1 / DCS GAIN INDEX
OBSTYPE = '            OBJECT' / IMAGE TYPE
RBIN    =                    2 / ROW BINNING
CBIN    =                    2 / COLUMN BINNING
CKSUMOK =                    T / CHECKSUMS MATCH
CAMCKSUM=                19345 / CAMERA-COMPUTED CHECKSUM
SFTCKSUM=                19345 / SOFTWARE-COMPUTED CHECKSUM
END
```

The header includes useful information about the size of the image (`NAXIS1` & `NAXIS2`), the date (`DATE-OBS`), the exposure time (`EXPTIME`), where the telescope was pointing (`RA` & `DEC`), the optical filter in the light path (`FILNAM`) and the binning factors used along the rows and columns directions (`RBIN` and `CBIN`, respectively). The header contains “keywords,” e.g., `FILNAM`, that have values on the right hand side of the equals sign.

Scrolling through the ASCII FITS header in a terminal is fine for browsing, but if you want to use information in the FITS header in your IDL program, you will need to read the keywords. This is accomplished with the IDL function `SXPAR`, e.g., to find the object name

```
IDL> x = readfits('image101.fits', hdr)
```

or

```
IDL> x = mrdfits('image101.fits', 0, hdr)
% READFITS: Now reading 1056 by 1024 array
```

followed by

```
IDL> print, sxpar(hdr,'OBJECT')
TrES-3
```

Two very important keywords in the header of our image are `BSCALE` and `BZERO`. Each image you read using the `MRDFITS/READFITS` function should immediately be transformed using the following equation: `x_real = x * BSCALE + BZERO`. Before this transformation, the image contains mostly negative numbers, which is obviously not normal (a detector can only count a positive number of electrons!). **NB**: this transformation is done automatically if you use the `/fscale` option of the the `mrdfits` function instead (as you did in the previous lab).

4.2 Data regions

The Lick 1m CCD FITS files contain images that have dimensions 1056 ×1024, even though there are only 1024 × 1024 “useful” pixels in the CCD array (when using a typical x2 binning factor). The extra 32 columns, those with the highest pixel numbers, are not exposed to light during the observations and therefore represent a measurement of the bias and dark counts. The actual data from each image are contained in section [`1:1024,1:1024`] whereas the dark is found in section [`1025:1056,1:1024`] (note that the `FITS` convention is that the first pixel is labeled 1 and not 0 as in IDL!). Note that there are always 32 column of "dark section", irrespective of the binning factor used in the observations. To subtract an average bias+dark count, the median count value in this region can be computed and subtracted from each exposure. Such a treatment, however, does not take into account spatial variations in the electronic properties of the detectors (e.g., a "bad" column). In this analysis, we will systematically trim the "dark section" from the images, using for instance:

```
IDL> img_trim = img(0:1023, 0:1023)
```

and use dedicated bias and dark frames instead.

4.3 Bias, darks, & flats

During the night we acquired various calibration frames, including bias, darks, and flats.

Bias Frames: A bias frame is a zero second exposure, which measures the number of counts the detector reads with no signal. These frames can be used to determine which part of the detector are affected by significant electronic defects (see left panel in Figure 4). Compute the median and pixel distribution of bias values.

Dark Exposures: A dark exposure is an image of non-zero exposure time where the shutter does not open, and is a measure of the dark current in the system. The dark signal (which contains the bias level as well) needs to be subtracted off all data images. Because the amount of dark current signal is proportional to time, the dark exposures need to be of the same exposure time as your data images, or else must be constructed by adding a bias and an appropriately scaled "dark-bias" image.

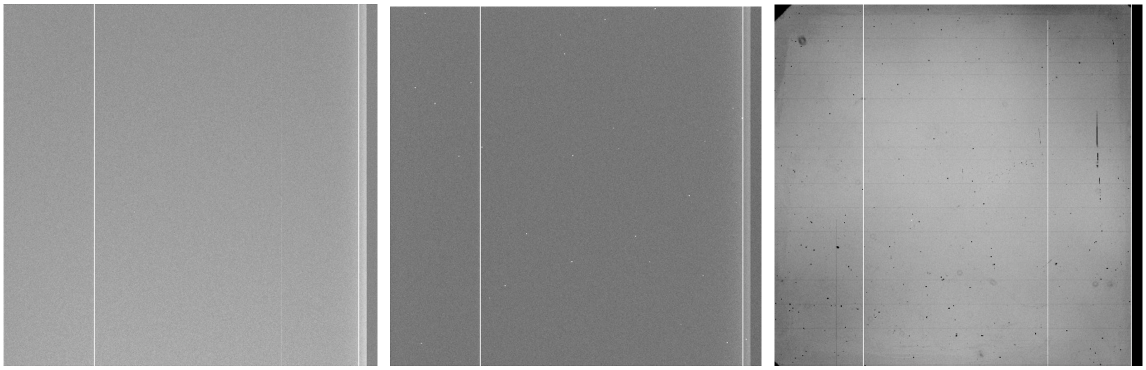

Figure 4: Raw bias (left), 90s dark (center) and flat-field (right) images for the Direct Imaging Camera on the Nickel telescope. Notice the 32 "overscan" columns to the far-right and the "bad" columns of the CCD detector. Also, the flat-field image reveals that each part of the detector has a different response to incoming light, either because of electronic defects or due to imperfection in the optical path of the instrument. For instance, most small dimples seen in the flat-field image are due to some dust motes on the entrance window of the instrument.

The Direct Imaging Camera is cooled by liquid nitrogen to a temperature of 77K to reduce as much as possible the electronic dark current. Use the bias and dark images to estimate the actual level of dark current. For the integration time used during the observations, generate a "master" dark image that will be subtracted off every image obtained during the night.

Flat Frames: Flat frames are images taken with the telescope pointing at a blank wall or at a region of the sky without any stars, often the horizon during twilight. The flat frames are a measure of the pixel-to-pixel response of the camera. You may see stars in your flat frames. Computing the median of all flat-field images will remove any star and produce a "master" flat-field (see Section 8 for an example of IDL script) that you can divide all your data images by to correct for the variation in pixel response across the CCD. Make sure the master flat-field is normalized, i.e., that its median value is unity.

5 Selecting reference stars and aligning all images

Summary of Section 5 Instructions:

- Subtract the master dark from all images and apply flat-field correction
- Correct images so that the average background count is zero, if necessary
- Find the locations of the science target and 5-10 reference stars to within ~1 pixel
- Align all images so that all stars remain at a fixed location (to within ~1 pixel)

5.1 First look at the data and selection of "reference" stars

Inspection of a data image (see Figure 5) reveals a few to a few dozen stars (depending on the location of the target in the sky). After subtraction of the master dark and division by the master flat-field, are the background counts exactly zero? Is that what you expected? If it is not null, the background must be brought down to zero. How do you perform this? Once this is done, what is the noise per pixel in the background? The resulting images are "clean", i.e., they have corrected for all known electronic "features" of the Nickel CCD. What is the apparent FWHM of stars in our images (this is what astronomers call the "seeing").

To construct the lightcurve of our target, we will compare its flux to that of other stars ("reference stars") in the image in order to correct for variations induced by other factors than the transiting exoplanet. For instance, thin clouds may come in and out of the direction toward the target. Also, as we follow the star over the course of its diurnal motion, our line of sight crosses a varying amount of atmosphere. Since the transmission of the atmosphere in the visible range is not a perfect 100%, this implies that our observations will suffer from a time-dependent transmission. By picking reference stars in the same field-of-view as our target, we can monitor *in real time* and calibrate out these variations. This method is called "differential photometry", since the fluxes we will measure are not calibrated on an absolute scale.

To select reference stars, you should pick the brightest (unsaturated) stars in the image. The more reference stars, the better, but only to the extent that their flux can be measured with good precision (i.e., avoid the bad columns!). In the image shown in Figure 5, a dozen stars are probably bright enough. Selecting 5-10 reference stars is a good number, although the exact number depends on the field that was observed and your own preferences. Record the position in the image of the science target and of the reference stars to a precision of ~1 pixel or better. Which method did you use to determine the location of the stars?

Figure 5: A raw image of a star field obtained with the Direct Imaging Camera (prior to dark subtraction and flat-field correction). The telescope is guiding to keep the target roughly in the center of the image. Many other stars are detected in the image, from which we can select "reference" star to monitor the relative brightness of the target.

5.2 Aligning all images

While we were obtaining our observations, the tracking system of the telescope compensated for the Earth's rotation and the guiding system applied finer corrections to maintain the stars in a constant location on the detector. Nonetheless, these corrections were not perfect and the location of any given star varies throughout our observations. Select a few images and check how large this effect is in your data.

In order to use a fully automated photometry routine, we must correct for this continuous "jitter" (if it is larger than half the seeing halo or so), since otherwise using an aperture at a fixed location will not provide reliable measurements. Which method can you think of to estimate the displacement of any individual image compared to a "reference" image (the one you have worked with so far, for instance)? Keep in mind that, given the large number of images we have to deal with, it is important that this process be as automated as possible. Once you have determined the necessary displacement, use the command `SHIFT()` to apply it.

Once you have selected (and tested) an alignment method, process all images through this alignment correction. We now have a collection of images that are corrected for all electronic effects and in which the stars are all in a constant location on the detector.

6 Generating the light curve

Summary of Section 6 Instructions:

- Determine the optimal aperture radius using a bright, isolated star
- Compute the flux of the science star and each of the reference stars for every image
- Compute a weighted average of the flux of the reference stars for each image
- Compute and plot the flux of the science star relative to the weighted average of the reference stars as a function of time

6.1 Aperture photometry

We must now measure the flux of the science target and of the reference stars. As you can see from your images, the light from each star doesn't fall on one pixel but rather is spread over many. In order to get the total light from each star, we need to integrate the signal from all these pixels. After we define the location of the star, we will simply add up the signal from all the pixels within a certain radius of this peak. However, defining the radius of each star on your image can be tricky; too small a radius means that we are leaving out some of the starlight, while too large a radius implies increasing the amount of noise by adding pixels that contain little stellar flux. The analysis of starlight by adding up the flux within a certain radius on an image is known as *aperture photometry*. One way (but not the only one) to compute the flux in a given aperture consist in defining a "mask" images (whose value is 1 inside of the aperture and zero outside) which can be used as follows:

```
IDL> Flux_star = total(img * mask)
```

To select the aperture radius, try using many different radii and determine which is best, as follows: Select a well-isolated and relatively bright star in the image (this can be the science target or one of the reference stars) and write a program that 1) computes the total flux in a circular aperture for a range of increasing radii, 2) estimates the total noise (this includes the background noise for each pixel in the aperture and the Poisson noise on the total number of *electrons – not DNs, remember to use the detector's gain!*) and 3) plots the resulting signal-to-noise ratio (i.e. [integrated flux]/[total noise]) as a function of aperture radius. Which aperture radius do you plan on using and why?

6.2 Generating the light curve

We are now ready to measure the flux of the science target and the reference stars in each cleaned image. Write a program that computes these fluxes for each image, as well the associated uncertainties (again, combining background and Poisson noises). In addition, to take advantage of the fact that we selected several reference stars, compute the weighted average of their fluxes (F_i) using $1/\sigma(F_i)^2$ as a weight[4]. Store the results (fluxes and uncertainties for each star, plus the weighted average of the reference stars) in a file, along with the time at which each image was taken (using the information located in the header). The function `PRINTF()` and `READF()` can be useful at this point.

Since we assume that the reference stars are intrinsically not variable (generate plots for each reference stars to test whether this assumption is correct), any variation in the ratio between these two fluxes is intrinsic to the science target. Compute and plot the ratio of these two fluxes as a function of time, along with its uncertainty. This is the lightcurve we have been trying to construct all along!

---

[4] To learn more about this choice of weight, read James Graham's handout on "the Method of Maximum Likelihood" available in the class website.

7 Analysis: Estimating the radius of the transiting exoplanet

Do you see any variation of the flux of the science target over time? Can you *quantitatively* demonstrate that we have detected the transit of the exoplanet in front of its parent star? How (hint: consider separately the "in transit" and "out of transit" data)? Is this consistent with the expected depth of the transit from the ETD database? Note that, to improve statistics, you may rebin the data (grouping the data in small packets of N individual images, averaging the fluxes and propagating uncertainties accordingly). If the transit is not detected, place an upper limit on the depth of the transit, i.e., determine the deepest transit that your data would not have been good enough to detect.

If you have detected the transit of the exoplanet, you can now use the light curve to derive some physical information about the planet. At first order, the decrease in flux during transit is a simple geometric effect whereby a small body (radius R_p) masks out a fraction of a larger, light-emitting body (radius $R_\star$). What is the relation between R_p, $R_\star$ and the depth of the transit? The best estimate of the stellar radius for your science target can be retrieved from the Exoplanet Encyclopedia[5]. If the radius of the science star has not yet been estimated, assume a radius of 1 $R_\odot$. What radius, or upper limit, do you infer for the exoplanet (in units of $R_{Jupiter}$)? How does it compare with that listed in the Exoplanet Encyclopedia? Given the mass of the planet (also available in the Encyclopedia in most cases), what is the average density of the planet (in g/cm^3)? Is this planet more like a gaseous giant (like Jupiter) or a rocky planet (like Earth)?

Finally, reflect on how the experiment could be improved to enable a better characterization of the extrasolar planet.

8 Appendix: Automated flat generation

Here's an IDL script to generate a flat field from the twilight sky frames. Frames are selected to have counts above some threshold, bias subtracted, scaled to the median value, and then median combined. Using a median rather than an average eliminates any stars in the field. Make sure you understand and adapt the script to your own data before using it!

```
; Automated sky flat generator : 10/10/2010
; first pick your Filter by uncommenting the keyword that identifies the
; frame as a V-, R-, or I-band flat

search_str = 'Flat_R  '
;search_str = 'Flat_V  '
;search_str = 'Flat_I  '

nx = 1056
ny = 1024
skythres = 3000.0 ; sky signal must be above this level to be included
; Range of filenames to consider
n0 =   1
n1 = 120
iflat = 0

; Loop over flat data and look for appropriate flats
for i = n0, n1 do begin
    fn = 'flat'+string(i,form='(i03)')+'.fits'
    print,'File ',fn
```

---

[5] http://exoplanet.eu/ For the Kepler planetary candidates (KOI xxx), the estimated stellar radius may be found at the following address: http://archive.stsci.edu/kepler/planet_candidates.html

```
    x = readfits(fn,hdr,/silent)
    if n_elements(x) gt 1 then begin
       obj = sxpar(hdr,'OBJECT')
       if obj eq search_str and median(x) gt skythres then begin
          if iflat eq 0 then begin
             fns = fn
          endif else begin
             fns=[fns,fn]
         endelse
         iflat++
     endif
   endif
endfor
nf = n_elements(fns)
print,'There are ',nf,' flats brighter than ', skythres
flats = fltarr(nx,ny,nf)
for i=0,nf-1 do begin
   x = readfits(fns[i],hdr,/silent)
   x = x(0:1023,0:1023)
   x = x - master_dark  ; Dark subtraction
   medval  = median(x) ; scale to the median value
   x /= medval
   flats[*,*,i] = x  ; stuff the frame into a 3-d array
endfor
skyflat = median(flats,dim=3) ; combine the median image
writefits,strcompress(search_str,/rem)+'-med-sky-flat.fits',skyflat,hdr

end
```

9 References

Borucki et al., 2010, The Astrophysical Journal Letters, 713, L126