Lab 3: PHOTOMETRY WITH AN INFRARED CAMERA

James R. Graham

1 Introduction

1.1 Purpose

The purpose of this lab is to make you familiar with the use of the Leuschner 30-inch robotic telescope and the infrared camera. Your goal is to show that you have all the skills to make quantitative astronomical observations with this facility. The ultimate task of this lab is to demonstrate this by making a plot which shows what exposure time is necessary to achieve a given signal-to-noise ratio for a star of known magnitude.

To accomplish this you must:

- Become familiar with celestial coordinates (right ascension and declination) and astronomical time systems (sidereal time). Learn how celestial coordinates and coordinate transformations permit you to be able to compute the elevation and azimuth angles of a star. To demonstrate this **include in your lab report a plot showing the elevation angle as a function of sidereal time** for the star that you use for your photometric calibration.

  The AY120 handout on coordinate transformations will be invaluable for this section:

  `http://www.ugastro.berkeley.edu/docs/astronomy/`

- Learn how to operate the Leuschner 30-inch telescope. Take data with the infrared (IR) camera, remove instrumental effects (dark current and flat field) and calibrate the data. Measure the gain and read noise of the IR camera. Determine the level at which the detector array saturates. **Quote these numbers in your lab report**. Construct and save as FITS files a flat-field and dark frames. **Include images of these calibration frames in your report.**

- Measure a star of known brightness and **document in your lab report the conversion factor between photoelectrons per second recorded by the camera and star brightness**. Repeat the observation several times and make a histogram of the number of photoelectrons per second recorded in your observations of this star.

- Collect all the factors that determine the signal and noise and **compute signal-to-noise ratios for stars of different magnitudes and exposure times and summarize in a plot in your lab report**.

1.2 Dates

- Field trip to Leuschner, 9/26. Bring warm clothes, water and a snack, and a **small** flashlight.

- Show and tell on 10/3: Astronomical coordinates, detector performance, saturation, gain, dark frames, and flat fields. Images of stars!

- Show and tell on 10/10: Calibrated data (dark subtracted and flat fielded), analysis of stellar photometry and signal-to-noise.

- Lab report deadline: 10/17, 6PM.

2 Some Basics of Photometry

First, we explore one of the fundamental techniques of observational astronomy—photometry, which is the measurement of the brightness of stars. Much of our understanding of stars comes from measurements of their brightness. For example, the first step in making a Hertzsprung-Russell diagram invloves the determination of the luminosity of a star, which requires a measurement of the brightness of a star combined with its distance.

The fundamental property of a star that we can measure with the infrared camera is the number of photons, N, that are detected in a given time interval, in a given frequency range (called a *frequency band*) $\nu_1 < \nu < \nu_2$. The atmosphere, telescope and detector limit the range of ν. For example, O_3 in the earth's atmosphere absorbs very strongly at $\lambda < 300$ nm. At near infrared wavelengths, H_2O absorbs at 1.1, 1.4, 1.8 and 2.9 μm.

Usually, the range of ν is further restricted by an optical filter which establishes more precisely the cut-on and cut-off frequencies. The simplest type of optical filter is a piece of colored glass.

The intrinsic property of a star that we want to measure is its spectrum L_ν, which is the electromagnetic energy radiated from the surface of the star per unit time, per unit frequency interval. The total power of the star is

$$L = \int_0^\infty L_\nu d\nu \text{ [watts] .} \tag{1}$$

If the star is at a distance, d, then the **specific flux**, F_ν, by conservation of energy is

$$F_\nu = \frac{L_\nu}{4\pi d^2} \text{ [watts m}^{-2}\text{ Hz}^{-1}\text{] .} \tag{2}$$

Stars (and even the Sun) always have $F_\nu \ll 1$ in these units, and to avoid writing large negative exponents astronomers have defined a convenient astronomical flux unit, called the jansky[1], 1 Jy $= 10^{-26}$ W m^{-2} Hz^{-1}.

It is evident that the number of photons detected by a light sensitive detector mounted on a telescope of collecting area, A, and in an exposure time, t, is

$$N = At \int_{\nu_1}^{\nu_2} \frac{F_\nu \eta_\nu}{h\nu} d\nu \tag{3}$$

where h is Planck's constant. Thus, under the two assumptions that **(1)** the spectrum, and **(2)** the efficiency for detecting photons ($0 \leq \eta_\nu \leq 1$) vary only slowly with frequency over the range $\Delta\nu = \nu_2 - \nu_1$, we can evaluate the integral approximately as

$$F_\nu = \frac{N\ h\nu}{At\ \Delta\nu\ \eta_\nu} \tag{4}$$

Thus, we conclude that $F_\nu \propto N/t$. In this lab we will directly measure N/t and hence measure a quantity that is proportional to the flux. However, we will not *directly* measure A, $\Delta\nu$, and η. We will eliminate the need to know these by using the technique of **differential** or **relative photometry**.

2.1 Magnitudes

Astronomers also like to use logarithmic units called magnitudes instead of F_ν to report the brightness of stars[2]. The magnitude of a star, m_ν, is a

[1] Jansky was a pioneer radio astronomer, who worked at Bell Labs in the 1930's.

[2] One argument is that the fluxes of stars spans many orders of magnitude, and we would have to remember their brightness in units of mJy, μJy, nJy, and pJy. For example

Table 1: Spectrum of Vega at Selected Wavelengths

Band Name	Central Wavelength	Band Width	Flux of Vega = $F_{\nu 0}$
	(μm)	(μm)	(10^{-26} W m^{-2} Hz^{-1})
U	0.37	0.066	1780
B	0.45	0.094	4000
V	0.55	0.088	3600
R	0.66	0.14	3060
I	0.81	0.15	2420
J	1.25	0.21	1570
H	1.65	0.31	1020
K	2.20	0.39	636

measure of the specific flux F_ν at some frequency ν

$$m_\nu = -2.5\ log_{10}\left(F_\nu / F_{\nu 0}\right)\ , \tag{5}$$

where the constant $F_{\nu 0}$ is a reference flux that defines the magnitude system.

The most common magnitude system is one where the star α Lyrae (Vega = HR 7001) is defined to have $m_\nu = 0.0$ at all frequencies. Astronomers use upper-case italic Roman letters to designate the magnitude in a band. Thus U is the magnitude in the ultra-violet (about 370 nm), B is the magnitude in the blue (450 nm) and V is the magnitude in the visible (550 nm). For Vega, $U = B = V = 0$. Thus the flux $F_{\nu 0}$ is just the spectrum of Vega (see Table 1).

We will use both linear (Jy) and logarithmic (magnitudes) measures of brightness.

the brightness of Vega is 3600 Jy ($V = 0$ mag.), and the brightness of the most distant quasar (SDSS J1148+5251 at $z = 6.42$ has $I = 23.3$ mag. or a flux of 1.1 μJy. Another, perhaps better reason is that flux ratios are used to estimate stellar temperatures. A flux ratio becomes a magnitude difference, which is easy to compute and to remember the conversion to spectral types.

2.2 Relative (or Differential) Photometry

For any particular telescope/filter combination it is very difficult to measure A, η_ν, and $\Delta\nu$ accurately because this would require **absolute** measurements. A few astronomers have made the painstaking efforts required to make these absolute measurements of a few stars. These are called **standard stars** and they define the primary flux scale used by astronomers. Vega is one of these standard stars.

It is much easier to compare the brightness of two stars, a technique that is called differential photometry. If we measure the brightness of a star and compare with Vega, then we can deduce the magnitude of the new star. Suppose we detect N^* photons from the star in t^* seconds, and when the same equipment is used under identical conditions to observe Vega we detect N_{Vega} photons in time t_{Vega}, then the magnitude of the new star is,

$$m_\nu^* = -2.5\ log_{10}\left(\frac{N^*\ t_{Vega}}{t^*\ N_{Vega}}\right) = -2.5\ log_{10}\left(\frac{N^*}{t^*}.\right) + 2.5\ log_{10}\left(\frac{N_{Vega}}{t_{Vega}}\right). \quad (6)$$

The challenge of this lab is to measure N/t for some stars and show that reliable photometry is possible with the IR camera on the Leuschner telescope. Of course the star you use does not have to be Vega, because many other stars have already been measured relative to Vega—it is one of these stars that you will observe. **In your lab report re-express Eq (6) for a standard star with magnitude that is not Vega**.

3 Photometric Observations

An important goal of this lab is to obtain observations of a star of known magnitude and measure the photon detection rate N/t. Repeat for several (3-5) different stars that span a factor of a few ($> \times 2$) in brightness and use a least squares fit to show that the relation between star brightness and count rate is linear and to find a good estimate of the conversion factor between Jy and photoelectrons per second. The following steps are necessary to extract numbers from the images suitable for performing photometry

- Dark/bias subtraction
- Flat field

- Sky subtract
- Extract star signal

3.1 Taking Pictures

The `qimage` program is used for acquiring data from the IR camera on `leuschner.berkeley.edu`. Log in with username `thirty`.

3.2 Saturation, Bias, Gain & Read Noise

If too many photons strike the detector, subsequent photons do not cause the output signal to increase further. The detector is said to have saturated. The purpose of this segment is to establish roughly where saturation occurs. Thus for a certain illumination level a short exposure will not saturate the detector, but a longer one will.

During the **day** there should be sufficient diffuse illumination in the dome to reach saturation in a few seconds with the K-band filter (2.2 μm). (Note: the filter in the IR camera can be selected with the `tx filter` command). With the K filter selected, the signal is mostly black-body emission within the dome, so the position of the flip-mirror does not matter. If you are not getting enough counts it is probably because the wrong filter has been selected. Note that when the IR camera is highly saturated the output signal is also zero!

Take a sequence of pictures with increasing integration times; start with the minimum exposure time, 0.66 second:

```
thirty@armin:~> qimage time=0.66 outfile=foo.fits
```

Note you have to be logged onto `leuschner.berkeley.edu = 128.32.197.194` as `thirty`, and you will have to use `scp` or `rsync` to copy your data back to `ugastro`. Choose a multiplicative factor between each exposure time so that you explore the possible range quickly. Examine the FITS images and plot the average signal as a function of the exposure time. The illumination in the dome may not be very uniform, so be sure to select a region on the images which have approximately constant number of counts. The counts in each FITS image should increase linearly with exposure time corresponding to a constant gradient. There should be a well defined knee in this curve where the slope becomes significantly shallower, beyond which point the count will no longer continue to increase.

Record the saturation value. Future IDL programs should always check to make sure that the pixel values from the camera are below the saturation threshold.

3.3 Dark Frames

You will need to subtract off the dark signal from your astronomical exposures. If your science exposure time is 5 s then you should subtract off a 5 s dark exposure. However, propagation of errors tells you that the noise in the difference image will have additional noise from the dark frame. How can you keep this noise to a minimum? How many dark frames do you need to average together to make sure that that subtraction of a dark frames increases the noise by no more than 10%?

You do not know *a priori* the length of your science exposures. So make sure that you keep a careful log and collect appropriate dark frames when you are done. You can anticipate this process by deciding to restrict your exposures to a limited menu, say 1, 3, and 10, seconds.

Save your dark frames in FITS files with unique names!

3.4 Saturation on the Night Sky

The night sky is not dark! The night sky glows due to scattering of artificial lights, scattered moon light is often very bright[3].

At infrared wavelengths ($\lambda > 700$ nm) the upper atmosphere glows due to a plethora of spectral lines from OH molecules which are excited when stratospheric water vapor and ozone react:

$$\mathrm{H_2O} + \mathrm{O_3} \rightarrow 2\mathrm{OH}^* + \mathrm{O_2}. \tag{7}$$

$$\mathrm{OH}^* \rightarrow \mathrm{OH} + \mathrm{h}\nu. \tag{8}$$

In addition, there is the scattered sunlight from interplanetary dust in the solar system and also the integrated starlight of all the stars in the Milky Way. In K-band (2.2 μm) black-body emission is also important from the

[3]From time to time the lights in the dome may get left on. If this is the case then you will not be able to observe. If you think that there is a problem check the dome web cam.

telescope's mirrors and the earth's atmosphere. They sky is relatively dark at 2.2 μm wavelength because the transparency of the atmosphere is good (Kirchoff's law) and because the mean atmospheric temperature is cool ($T \simeq$ 250 K). All of these contribute to the signal seen by the detector.

How long can you observe the sky before the chip saturates? Pointing at the night sky, repeat the exposure sequence you used for the inside of the dome. The night sky is significantly darker than the inside of the dome during the day, so exposure times can be longer. For this test we want a region of sky that is devoid of bright stars, so select coordinates that correspond to a region of high galactic latitude.

The command to point the telescope is (for example):

```
tx point ra=18:36:56.2 dec=38:47:01 epoch=2000
```

Before sending the pointing command be sure to confirm that your chosen coordinates are at a reasonable declination ($-5° > \delta > 60°$) and hour angle ($-3h30m < HA < 3h30m$).

Check the current pointing limits by asking the telescope:

```
thirty@armin:~> tx tel_status
```

It is always a good idea to confirm that the pointing is reliable by verifying that you can find a bright star near the target position

```
thirty@armin:~> bright ra=12:22 dec=33 faint=4
```

Once you have chosen your star you can pipe the output into the `dopoint` command.

```
thirty@armin:~> bright ra=12:22 dec=33 faint=4 | dopoint
```

Consult a copy of Norton's star atlas or use `xephem`. For example, if you observe in the early evening in September, avoid crowded regions of the Milky Way such as Aquila, Cygnus, Sagittarius or Scorpius. If you want to be really careful, use the digital sky survey at the Space Telescope Science Institute:

```
http://stdatu.stsci.edu/cgi-bin/dss_form
```

to inspect your coordinates. If there is a prominent star with clearly visible diffraction spikes, choose another region.

Start with a short exposure (e.g., 1 s) and work up, say in factors of two. Plot the mean signal from the array versus time (avoid stars when you calculate this statistic—try using the IDL command `defroi`). Since stars might still be present in a small number of pixels, the mean may not be the best statistic. In a star-free region most pixels will have values which reflect the sky brightness, thus the mode rather than the mean is a useful for describing the sky signal. The mode is difficult to calculate because it requires fitting a probability density function to the data. Frequently we assume that the median, which can be found by simply sorting the pixel values, is a good estimate of the mode. IDL has a convenient built-in function to compute the median. If `image` is the name of an IDL variable' containing a 2-d image, try:

```
IDL> print, median(image)
```

Plot the median pixel value of the array as a function of exposure time and establish the maximum integration time you can use without saturating the detector on the sky.

Note that the sky brightness can vary by factors of 20-50% over times of tens of minutes, so your subsequent choice of exposure times should always incorporate some head room. The apparent brightness of the sky also depends on zenith angle (angular distance from the zenith). The sky brightness at 60° from the zenith, i.e., an elevation of 30°, is about twice that at zenith because the column of air is twice as great.

3.5 Pixel gains and zero points: Using the Sky as a Flat Field

Each pixel in the CCD has a slightly different gain. To do reliable photometry we must correct for this problem. It's hard to measure the relative pixel gains because it's hard to find a uniformly illuminated area to do flat fielding. The best source is the sky itself! We can use the twilight sky as a flat-field background. The best time to take your data will be between sunset and nautical twilight (-12° twilight). Just after sunset the sky will be too bright and when it is dark you will not see enough signal to make a good flat field.

3.5.1 Observations

You need to obtain a sequence of exposures of the morning or evening twilight sky. Because the sun is either rising or setting, each exposure will have a different value of the sky brightness B. First, compare with your results on saturation to decide on a good exposure time for a typical sky exposure. Then take a sequence of pictures with that integration time, obtaining different sky brightnesses during the sequence.

During each sequence, point the telescope towards a slightly different position for every exposure so that the faint stars in the image do not occupy the same pixels in every exposure. The amplitude of the motion depends on the sharpness of the images. Typically, the full width at half maximum (FWHM) of images at Leuschner is 3-4 pixels, so it will be necessary to jog the telescope by > 10 times this width, i.e., 30-40 seconds of arc (0.01 degrees). It should not be necessary to move the telescope to an entirely new field, so keep the amplitude of your offsets to less then 300 seconds of arc (5 minutes of arc). To move the telescope a small amount (say .01 degree North or $0.'6 = 36''$) type

```
tx offset dec=.01
```

Don't forget that one second of time corresponds to $15'' cos(\delta)$. This factor is taken into account if you move in hour angle and add the `cos` parameter

```
tx offset ha=.01 cos
```

3.5.2 Analysis

The variations in B allow us to determine the pixel gains. Assume that we take N images each with the same exposure time (but different B_i), where the frame numbers run from $i = 0, 1, 2, ...(N-1)$. The signal in the i-th frame at pixel (x, y) is denoted by $z_i(x, y)$ and is given by

$$z_i(x,y) = f(x,y)B_i + c(x,y) \ , \tag{9}$$

The term $f(x, y)$ are the pixel gains that we are searching for. The constant $c(x, y)$ is dark signal plus arbitrary offset. You guessed it—we should use the technique of least squares to figure this out!

For each pixel in each image we have two unknowns: the sky brightness B_i, the pixel gain $f(x, y)$. In any given exposure we assume that the illumination is constant and B_i is the same for all pixels we have previously measured the dark contribution appropriate for a given exposure time, $c(x, y)$, so all that remains to do is to a linear fit for each (x, y) pixel in the CCD. Thus we have $2n_{pix} + N$ unknowns and $N \times n_{pix}$ data points. Unfortunately, this leads to a **non-linear** least squares problem, which cannot be solved using the techniques of linear algebra (see Appendix A).

To make the problem tractable, let us use the dark frame to subtract off the constant c term. If we define the median value of the flat-field to be unity, then we can compute the sky level in the i-th frame directly,

$$B_i = median[z_i(x, y) - c(x, y)] \tag{10}$$

The flat field then is simply the average of all the dark subtracted, median divided frames:

$$f(x, y) = \langle \frac{z_i(x, y) - c(x, y)}{B_i} \rangle. \tag{11}$$

Note for this experiment we are only interested in the *relative gain* of each pixel, i.e., the *flat field*.

Compute the errors in your flat-field using error propagation techniques and make sure that you have taken enough data so that the fractional error is $< 1\%$. Display your flat-field using `tvscl` and make sure that there are no artifacts due to stars. Read in Taylor about weighted means. How can you use a weighted mean to improve the precision of the flat field?

Save your flat-field in a fits file!

3.6 Photometric Accuracy

You are now ready to try some photometry. The first step is to show that you can perform reliable measurements of star brightness. We do this by repeated measurement of a single star.

Almost any star if fine for this experiment, so long as it is not so bright that it saturates the IRCAM detector. However, it is preferable to use the list of photometric standard stars found on the Leuschner IRCAM web page. If you use one of these stars then this measurement will also give the conversion factor between stellar magnitude and counts per second.

Note the local sidereal time and select a bright star from the list. Use the relationship between hour angle (HA), sidereal time (ST) and right ascension (RA),

$$HA = ST - RA$$

to pick one of the brighter stars ($K \simeq 10-11$ mag.) that is ≤ 2.5 hours east ($HA < 0$) of the meridian. Don't pick a star in the west ($HA > 0$)—it is already setting!

Take some test images to establish the exposure time that gives a strong signal without saturating. Our goal is to measure the brightness to better than 1%; this sets a firm lower limit to the number of photons needed of 10,000. You will probably find that $\times 2-3$ more counts are necessary, because sky noise will dominate the SNR.

Write a script to jog the star over a 3×3 raster pattern on the array so that it appears at nine distinct positions. Since the pixel scale is about $1''$ per pixel an offset of $30'' - 60''$ between pointings will be satisfactory.

Before we can extract photometry, we have to dark-subtract, flat-field and sky-subtract these data. Dark subtract and flat-field your data using the previously saved files.

There are two ways to perform sky subtraction. You can perform pair-wise subtraction, i.e., take a pair of images where the star appears in two different locations and subtract one from the other. Note, that if you do this you don't need to dark subtract—why? Another simple way to sky-subtract that avoids separate dark subtractions would be to calculate a pixel-by-pixel median value for each pixel in the stack of images. Since the median is essentially unaffected by the presence of the star (or for that matter any other faint stars), the median image provides a good measure of the sky level. To sky-subtract, simply subtract the median image from each exposure. This step applies sky subtraction and dark subtraction, but you still need to apply the flat field correction.

3.7 Photometry

When you examine your images, you will notice that the starlight is spread over several pixels. Thus, if you want accurate photometry, then you need to add up all the light in a region that contains the light. Stellar images are reasonably well described by a 2-d Gaussian, so roughly 39% of the light is contained within a circle of radius equal to the standard deviation[4]. And (86%, 99%) of the light is contained within radii of $(2\sigma, 3\sigma)$. A circular aperture that has a radius of 3.0σ (a diameter equal to $2.548\, FWHM$) includes

[4]Hey, wait a minute! I thought that 68%, not 39%, of the light would be contained within one standard deviation. Have your professors gone crazy, or what?

99% of the light.

Write an IDL program that sums all the pixels within a given radius centered on a star. (The IDL commands `cursor` and `rdpix` will be useful for exploring pixel coordinated and values). **In your lab report give an example of a plot of the pixel sum as a function of aperture radius**, and confirm that the signal reaches an asymptotic value as the radius increases. Inspect the pixels adjacent to the star using `rdpix`. Is the signal here zero? Since the sky brightness is not entirely uniform, and flat fielding may not be perfect, your dark-subtracted, flat-fielded, sky-subtracted images may not have a uniform zero background. Inspection of your images may show a gradient in uncorrected sky brightness. Therefore, you also need to establish the sky level in the vicinity of the star. Modify your routine to compute the average "sky level" in a annulus centered on your star. Subtract this local sky level from every pixel in the star before performing the sum to find the signal from the star. To minimize the noise in this process there must be several times more pixels in the sky annulus than in your star aperture. Use your knowledge of error propagation to choose the inner and outer radii of the annulus centered on the star. The more pixels you include in the annulus the more accurate your sky level determination. Also, there should be no overlap between the star aperture and the sky annulus so that the signal in the annulus is not perturbed by the light from the star. Use the IDL command `tvcircle` to show where the apertures fall on the image. **Include in your report an image of the star with circles designating the apertures drawn on it.**

Now display the results of your photometry for each of the nine exposures. How reliable is your measurement? Compare the observed scatter with what you would expect from propagating errors, including read out noise and Poisson photon noise. Record the mean value and SDOM for this experiment and repeat for few more stars with different brightness.

For extra credit: Since you observed a star of known brightness compute the photon collection efficiency of the infrared camera+telescope. Assume that the band-width of the K-filter is as listed in Table 1.

4 Observing checklist

Before you take a picture you should check:

- It is cloudy? Look out at the sky from Campbell Hall roof. If it is foggy or cloudy it may not be clear at Leuschner. Check the weather server at the Leuschner web page.

Look at the 1-km resolution satellite imagery from the Naval Research lab "West Coast & EPAC"

`http://www.nrlmry.navy.mil/sat-bin/epac_westcoast.cgi`

and click on the image for Monterey. You can see the fog spilling in through the Golden Gate and headed for Lafayette. With a little practice you will be able to tell whether or not it is clear.

- Is the dome slit open?
- Is the dome oriented so that the telescope is looking out the slit?
- Is the telescope focus correct? Adjust the secondary position so that stars images are sharp. If stars look like donuts, then the telescope is out of focus. The full width at half maximum of a star should be 2-3 pixels.
- Is the flip mirror in the right position?
- Have you selected the correct filter?
- Are you pointing at the right celestial coordinate? Always check the numbers you send to the telescope. Is the telescope pointing correctly? Ask the telescope to point at a bright star nearby—if it is clear and the bright star does not show up then you need to adjust the telescope pointing. If this test fails there is no point proceeding until you have figured out what is wrong.
- Did you choose the right exposure time? If you expose for too short a time you will not be able to see faint targets. If you expose for too long brighter stars will be saturated. Check the counts in the image with `rdpix`.

A Appendix

The least squares solution for finding the flat field starts with defining χ^2 as

$$\chi^2 = \sum_{i,x,y} [z_i(x,y) - f(x,y)B_i - c(x,y)]^2$$

and computing the partial derivatives

$$\frac{\partial \chi^2}{\partial B_i} = 2\sum f^2(x,y)B_i - 2\sum f(x,y)z_i(x,y) + 2\sum f(x,y)c(x,y) = 0,$$

$$\frac{\partial \chi^2}{\partial f(x,y)} = 2\sum f(x,y)B_i^2 - 2\sum B_i z_i(x,y) + 2\sum B_i c(x,y) = 0,$$

and

$$\frac{\partial \chi^2}{\partial c(x,y)} = 2\sum f(x,y)B_i - 2\sum z_i(x,y) + 2\sum c(x,y) = 0.$$

Solving these simultaneous equations yields the sky brightness, B_i, the flat field, $f(x,y)$ and the dark/bias, $c(x,y)$. The solution cannot be found by the techniques of linear algebra because unknown quantities such as f^2 and B^2 enter into these equations. The problem is therefore a **non-linear** least squares one.