THE UNIVERSITY OF
MELBOURNE

Observational Astrophysics

Part 3 Laboratories

March 3, 2005

Contents

1 Some words about third year labs

These notes **don't** contain all the information you need - they're here to define the terms for you, and to give you an idea of the direction you should be taking. You'll find that you need to go foraging for more detailed information in other places - often in the references, and even more often in the rubble of your demonstrators' minds. You're going to learn a lot while you're doing that, so don't waste the effort - share all the work you're doing with us in your report! Remember that your report has to be self-contained. It doesn't have to be beautiful, but it should be clear and detailed. Perhaps the most important thing to remember is that when you write your report you're trying to **teach** your reader what you did, why you did it, and what you learned. That way, we get to learn something too! Enjoy!

By the way, these notes have recenlty undergone a major revision, and may be littered with typos (and possibly more serious errors). We in part3 would love it if you could bring any mistakes to the attention of your demonstrators, or email suggestions to **part3@physics.unimelb.edu.au**

2 CCD Reduction and Analysis

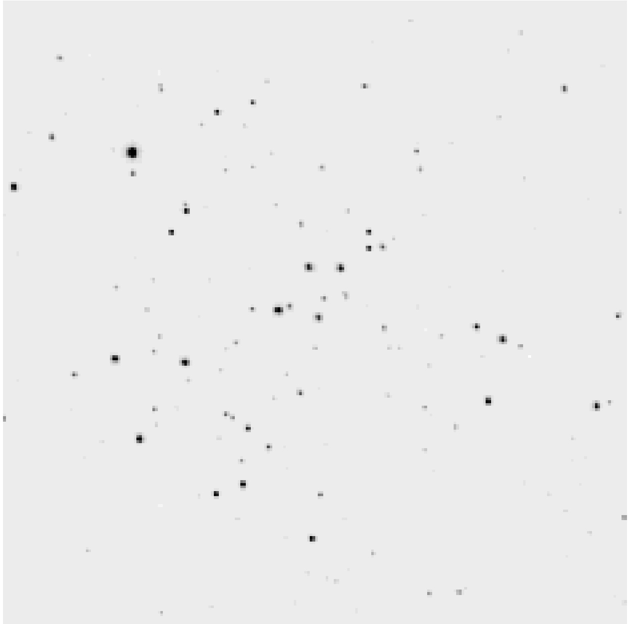

2.1 Background

In this experiment you will analyse data taken at Siding Spring Observatory, near Coonabarabran in NSW. You will measure the brightness and colours of the stars in an open cluster, and hence produce your own Hertzsprung-Russell (HR diagram). Using this information, you can work out the size, luminosity and temperature of the stars in this cluster.

In the process, you will learn the techniques of CCD data reduction in professional optical astronomy. You will know more about what's involved in observational astronomy than many professors of theoretical astronomy.

2.1.1 CCDs

Astronomers don't look through telescopes any more. In fact, it is anatomically impossible to look through most large telescopes! Starting about 1880, photographic plates replaced the human eye as the most common astronomical instrument. The reasons for this are the same as for using any camera: a permanent record is created and long exposures mean that objects too faint to see with the eye can be recorded in detail. Starting about 1975, electronic detectors began replacing the photographic plate.

The reason for the change is efficiency. A high quality astronomical photographic plate will record about 2% of the photons falling on it. The latest research-grade electronic detectors record more than 90% of the photons. They have other advantages too; since they generate digital data you can subtract off background light (a big advantage in a site like central Melbourne), and use all sorts of techniques to sharpen up your image. Most important of all, electronic detectors allow you to quantitatively measure the radiation coming from sources because they are linear devices— the number you get out of a CCD is directly proportional to the number of incident photons.[1] At their most basic, CCDs are simply photon counting devices. If you know how many photons an object has emitted, then you know how much energy it has emitted and you are on your way to understanding what it is/how it works.

CCDs are basically silicon chips, with a very special circuit pattern etched on their surfaces. When photons hit the surface they generate electron-hole pairs. The electrons are attracted to a grid of electrodes etched into the chip. Each of these electrodes forms a potential well, and it traps and stores all the electrons produced in its part of the chip. So as time goes on, a pattern of charge builds up near these electrodes. The more light falls on a given part of the chip, the more electrons will accumulate near the electrodes there. The pattern of electrons building up on the chip is thus an image of whatever the CCD is looking at.

When the observations are finished, the pattern of electrons on the grid of electrodes has to be measured and stored on computer. To do this, one end of the chip is connected to a set of wires. Under computer control, the voltages on the electrodes are set oscillating, so that a set of ripples of potential move across the chip. All the electrons are swept along by the ripples, moving to the edge of the chip where they are 'counted' (the potential they generate is measured) and the result is sent to the computer.

CCDs are now the dominant technology used in astronomy. The same technology is now widespread in devices such as video recorders and digital cameras. They are relatively cheap and reliable. It won't hurt them if you expose them to light accidentally. And they can be very efficient; most astronomy grade CCDs can collect and store an incredible 99% of the photons that fall on them. With a CCD attached, a little telescope on the roof of the School of Physics is as powerful as the biggest telescope in the world would be using naked eye observations.

2.1.2 The Hertzsprung-Russell Diagram

Stars, like people, come in many shapes and sizes. Massive blue stars burn furiously and die young (like pop singers). Boring stars like our sun chug along for billions of years emitting yellow light before they run out of fuel, swell up into cool bloated supergiants and finally splutter to an end as piddling little white dwarfs. Most stars are quite well understood (almost anything looks easy to understand when it is as distant as a typical star!), but people keep finding bizarre stars that break the rules.

In this exercise you will look at a cluster of stars, and measure their properties. You will use the available observations to figure out how bright they are, how hot, how large and more. In particular, you will plot your own Hertzsprung-Russell diagram, the famous plot of colour against brightness that allows you to classify stars into their myriad different types.

2.2 A Brief Summary of What You Should Do

1. Familiarise yourself with the observational data provided on the linux computer and process it using the IRAF (Image Reduction and Analysis Facility) package.

---

[1] Hence, the raw CCD output is usually expressed in units of 'counts'. Typically, this is just an integer between 0 and 2^{16}=65536 because the electronics on the chips utilise analog-to-digital converters with 16-bit resolution.

2. Measure the brightness of all the stars in your processed data at several different wavelengths.
3. Plot the HR diagram and work out the properties of your stars.

2.3 The Telescope

This data was taken by Dr Alicia Oshlack (a former Ph.D. student in the astro group) using the 40 inch telescope owned by Mount Stromlo and Siding Spring Observatories. This telescope is able to be used by anyone, provided they submit a proposal detailing their proposed observations and giving a scientific justification for the use. This is essentially the same process used for all major astronomical telescopes across the world.

2.4 The Data

If you look at a random bunch of stars, you get a very messy HR diagram. The reason is the stars are all at different distances. Faint red nearby stars may appear brighter than bright blue background stars, just because they are closer. For this reason, we will look at a cluster of stars. This ensures that nearly everything we see is at roughly the same distance (though we might accidentally get an occasional star in the foreground or background). All the stars in a cluster form at once, so everything we observe will have the same age. There is also a third advantage; because stars are crowded together in a cluster, we can get images of lots of stars at once.

Clusters come in two types, open clusters (relatively nearby, and containing up to a few hundred stars) and globular clusters (more distant, and containing perhaps a million stars). Unfortunately globular clusters are so distant that with our telescope all the stars will seem to blur together. We will therefore observe open clusters.

Table 1 lists several bright open clusters that can be seen from southern Australia around autumn. However, we have multi-colour images of only one of these clusters – M93 – and so this will be the one you will analyse.

2.4.1 Images to Analyse

You will need following images to do the analysis:

- Images of the cluster, in the filters B (blue), V (green), R (red) and I (infrared).
- A "flat field" image, to go with each filter used.
- A bias frame.

The last two images are described below.

Name	Right Ascension	Declination	Distance	Comments
NGC 2447	07:44:36	-23:52	1104pc	6th magnitude, also called M93
NGC 2287	06:44:54	-20:42	700pc	5th magnitude, also called M41
NGC 2422	07:34:18	-14:22	520pc	bright, also called M47
NGC 2437	07:39:18	-14:42	1530pc	6th magnitude, also called M46
NGC 2477	07:50:30	-38:25	1290pc	6th magnitude, lies in the Milky Way
NGC 2516	07:59:42	-60:44	1320pc	3.8th magnitude, very large
NGC 3293	10:31:30	-57:58	430pc	also called the Gem cluster

Table 1: Southern hemisphere open clusters visible during evenings in autumn.

2.4.2 Bias frames

Since CCDs are photon counters, one would expect that if an image was taken of nothing (i.e. just blackness) then the readout would be zero in all pixels. But this is not the case. In order to trap electrons in the pixels of the CCD, a voltage (or 'bias') must be applied across the pixels. This bias generates some additional counts in each pixel when read out. If this bias is known, then it can be subtracted from the real exposures so that only counts from real photons remain. Astronomy grade CCDs are designed so that this bias varies as little as possible between exposures so it can be accurately subtracted.

Thus, the bias frame is simply the readout from a zero-second exposure. You will notice when you examine the bias image that it is not perfectly uniform– there are random variations of a few counts between pixels. This is the "readout noise" of the CCD and is an unavoidable part of making measurements with CCDs. Astronomical CCDs are usually cooled with liquid nitrogen to minimise this readout noise.

Since the same CCD was used for all four exposures, only a single bias frame is necessary for all of the images.

2.4.3 Flat fields

Now, our telescope + CCD system, while very good, is not perfect. One fault is that the response to a uniform illumination will not be a uniform response across the CCD. Instead, there will differences in the response between different regions, and these differences will persist into observations of real sources. Typically, the throughput at the centre of the field is slightly higher than the throughput at the edges.

To account for these variations, so-called "flat field" images are taken. These images are of a uniformly illuminated source – such as a lit screen inside the telescope dome, or, as was used for these images, the sky at twilight (however, you have to be careful with twilight flats not to get too many stars appearing).

2.5 Reducing the Data

This section describes how to get your data into a suitable state to measure the magnitudes. This process is known in astro-speak as reduction (although sometimes it might create more images in the process!).

First, log in to the linux computer as `part3` (the password will be available in the lab). You will end up with a screen with an xterm in the middle. The xterm runs on the UNIX operating system, and your friendly demonstrators will be able to help you should you get stuck with using it. If you type `ls` you will get a listing of the current directory. You should see a directory called `originals` – this is where all the relevant images are kept. Please do not alter these! Instead create your own directory (I've used the name `bob`, but you can be more creative!) and copy them into it:

```
> mkdir bob
> cp originals/*.fits bob/
```

Now you want to start IRAF. This stands for *Image Reduction and Analysis Facility*, and is one of the most common tools used for reduction and analysis of astronomical images, particularly optical and infrared ones. Bring up a new window (of a slightly different type) by typing

```
> xgterm &
```

in your first window. Now type `cl` in the new window to start IRAF. You will notice that the prompt will change, and you will get a brief welcome message. You also want to start the image display tool, called `ds9`. At the IRAF prompt, type:

```
> !ds9 &
```

Images can be displayed in `ds9` by using the command `display` – you will need to put in a frame number for `ds9` to display it in, and you have a choice of four. Move into the directory where your images are (eg. `cd bob`), and display one of your images. For example:

```
> display M93_B.fits 1
```

2.5.1 A Quick and Dirty Look

`imexam` is an interactive tool for analysing astronomical data. It allows you to take a quick, dirty look at the properties of your images. To start it, just type `imexam`. If you haven't already got an image in the display, you can type the image file name after `imexam` to display automatically (eg. `imexam M93_B`).

Once you have the ring-shaped imexam cursor on your image, there are a number of keys that can be pressed to examine it:

- `l` — produces a plot of the intensity of the line the cursor is currently on.

- `c` — as for `l`, but for the current column.
- `s` — produces a 3-D surface plot of the region surrounding the cursor.
- `e` — produces a contour plot of the region surrounding the cursor.
- `a` — prints out the current position of the cursor, and results from the photometry calculations. Ask your demonstrator to translate this output for you.
- `r` — produces a radial plot of the source the cursor is currently over, as well as the output of `a`.
- `q` — quits `imexam`.

Play around with `imexam` for a bit to get a feel for what your raw data looks like. In the next section we'll use `imexam` to see how your data changes through the reduction process.

At this point it is worthwhile considering exactly what you are looking at. When an image is displayed, some stars appear to have a larger diameter than others. What is going on? Is the size of the image of each star different?[2] You can move the cursor over the display window and the count value for pixels will be displayed. Is there are difference between "large" stars and "small" stars? Make some radial profile plots of different stars. What do you notice? Is the radius over which light is spread for a "large" star actually any different than for a "small" star?

2.5.2 Processing the Data

To understand how CCD image reduction works, you need to understand that a CCD image is simply an array of values, where the value indicates the number of counts in that particular pixel. Since the images are just numbers, you can perform simple arithmetic operations on them, such as adding two images, or subtracting them and so on. This is simply what IRAF does – manipulates images arithmetically to improve the signal-to-noise, remove instrumental effects, show features not readily apparent in single images etc.

All you need to do to your images is to remove the effects of both the bias and the flat field. The total number of counts in a particular pixel (x, y) is given by

$$C(x, y) = S(x, y) \times F(x, y) + B(x, y),$$

where S is the signal (the interesting part), F is the flat field contribution and B is the bias. As you go through this procedure, use `imexam` to make a note of how the background noise changes as you first account for the bias, and then the flat field.

The way you will be removing the bias and flat field effects is by using the IRAF procedure `ccdproc`. Before we do anything, we need to set the parameters of ccdproc, so that it will do what we want. To do this, type `epar ccdproc` (epar = 'edit parameters'). You will see a screen with a lot of options (most of which we are not interested in here!) The one of relevance

---

[2] Each pixel on this CCD has an angular size of approximately 1/2 arcsecond. Assuming these stars are the same size as the Sun and 1104pc away, what is their angular size? So what is causing their images to be larger than this on the CCD?

at the moment is the zero (or bias) correction. Scroll down with the arrow keys and make the `zerocor` parameter "yes", and the others in that section "no". Then go down further and change the `zero` parameter to the name of the bias file (which should be `bias.fits`). To exit, type `:q`. To run ccdproc on an image (say bob.fits), type `ccdproc bob.fits`.

Do this for all the cluster images. The flat fields have already been bias subtracted, so it is not necessary to process them in this way.

Now you can remove the flat fields. Type `epar ccdproc` again, and make the `flatcor` parameter "yes". Scroll down and make the `flat` parameter equal to the relevant flat field image. This should be the same filter as the image you are about to run `ccdproc` on (as the flat field is different in different filters). Then run `ccdproc` on the cluster and standard images again.

To make sure which filter the image has been taken with, you can look in the "header" – this is the part of the file that contains all the technical junk about what the telescope was, what (and where) it was looking at, how long it looked for, etc. To read the header of the image bob.fits, use the command `imhead bob.fits`. (If you get just a single line, use `epar` to change the `longhead` parameter of `imhead` to yes.)

Is there any noticeable difference between image before and after bias/flatfield correction? What about the background levels (i.e. blank spaces between stars), are they zero now? Is the background the same for all filters? You will probably find the background is least for the B filter and most for the R and I filters. This is due to *sky brightness*. During the day, the Sun excites molecules in the atmosphere, which re-radiate at night. Most of the re-radiation is at specific (quantised) wavelengths in the R and I bands and the near-IR. Hence, the counts in your CCD measurements of stars also contain counts from the night sky. Luckily, the night sky uniformly illuminates the CCD just like a flat field. Your flat field corrections should make this background uniform across the CCD, hence it too can be compensated for.

2.6 Measuring the Brightness of your Stars

2.6.1 Aperture Photometry and the Point Spread Function

Once your images are fully reduced, you can go about measuring the brightness of the stars in the cluster. We will use the task `phot` in the `digiphot.apphot` package for this purpose. `phot` performs aperture photometry on an input list of coordinates. For each coordinate in the list it automatically performs the following steps:

- Determines the centre of the star
- Adds up the total flux within a circular aperture around this point
- Adds up the flux within a surrounding annulus to determine the underlying sky brightness
- Subtracts the sky brightness from the aperture sum to determine the brightness of the star

The basic idea of aperture photometry is shown in Fig. 1. We need to determine three parameters before we can run `phot`: the aperture radius, the inner radius of the sky annulus,

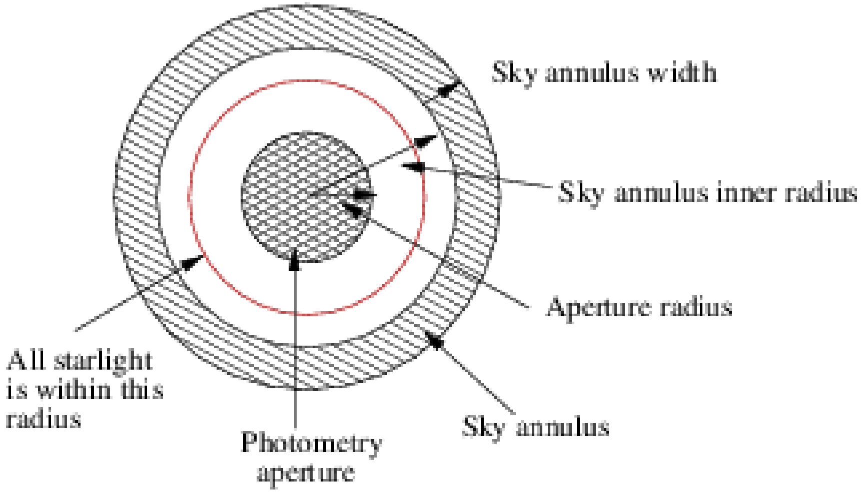

Figure 1: Parameters for aperture photometry

and the width of the sky annulus. To be able to set these, we first need to take a closer look at the shape of the light distribution in our stars.

Use `imexam` to make radial profiles and surface plots of some of your stars (the `r` key), and take note of the shapes you see. Although each star has a distinct peak, they're far from point-like (as would be expected if we had perfect resolution). Light is spread out over a wider area due to scattering by the atmosphere (mostly) and diffraction by the telescope optics (a little). The pattern produced by a point-like source on the focal plane of a telescope is called the Point Spread Function (or PSF).

For the first `phot` parameter, we need to choose an aperture radius which includes most of the light from each star, while excluding the light from surrounding stars. From the radial profile plots, find the radius at which the PSF begins to level off. Examine radial profiles for the image in each filter to see whether this radius is consistent; because the images were taken at different times of the night, the atmospheric conditions, and hence the width of the PSF, may well have changed. Record the aperture radius/radii.

For the inner radius of the sky annulus, find the point at which there's essentially no stellar light in the profile, i.e. the radius at which the profile becomes flat. Record this as the inner radius of your sky annulus.

In determining the width of this annulus you need to ensure that you include enough pixels to get a good statistical median of the local sky background, while not including too many neighbouring stars (don't worry if a couple of stars sneak in here and there — they shouldn't affect the median too much). Record this annulus width. If you can find one set of parameters which is suitable for all filters, you will make the task a little easier for yourself.

Now, load up the `phot` package by typing `digiphot` and then `apphot` and edit the `phot` parameter file (`epar phot`). Within `phot` there are sub-parameter files called "psets". Go down to the "Sky fitting parameters" pset and type `:e`. Set the parameter `annulus` to your inner radius and `dannulus` to your annulus width. Also, make sure that the parameter `skyvalue` is set to `INDEF`: this ensures that `phot` does the sky calculation for you.

Exit this pset with `:q`, and go down to the "Photometry parameters" pset. Here we want to enter a single aperture radius for the parameter `aperture`. Exit back to the main parameter

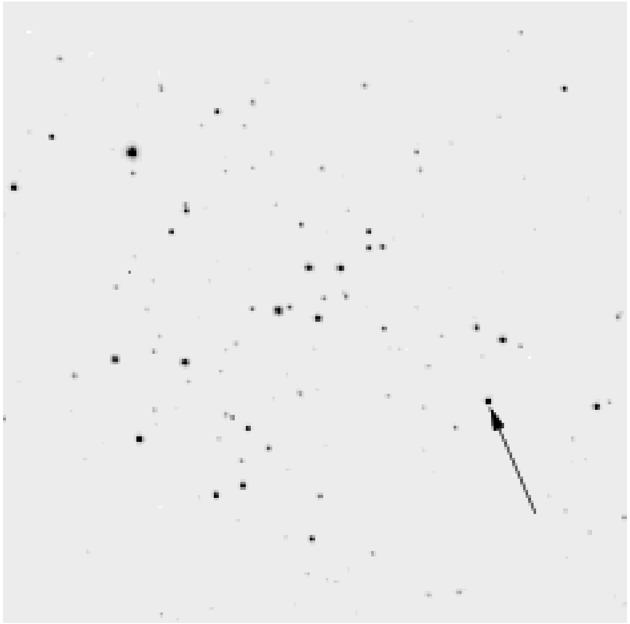

Figure 2: Open cluster M93, the arrow points to the standard star.

file again and save your entries with `:w`. Now exit the parameter file altogether with `:q`. We're not quite ready to run `phot yet` — we still need to tell `phot` where the stars are.

2.6.2 Choosing Our Stars

To find the coordinates of the stars we'll use `rimcurs`. Start the task by typing `rimcurs > coord.lst`. This will direct the output to file `coord.lst`. Place the cursor over the centre of the star and hit the space bar to register the coordinates. Start with the star indicated by the arrow in the image above (this will be your standard star — more on that later), and then try to work systematically across the field. Leave out the very faintest stars (a lot of these might be background stars), but otherwise try to get a good selection of stars at all brightnesses. Don't worry too much about finding the centres of the stars too accurately — `phot` has an algorithm which does that for you. Just get it to within a few pixels. When you're done, hit `Ctrl-z` to exit the task. The coordinate file you just generated should be good for all the cluster images, as the offsets between the different images are small.

Now it's time time actually run `phot`. Get back into the `phot` parameter file again and enter the name of the first input image, the coordinate list and the results file. Get your demonstrator to check your parameter files (this may save tears later on), and then run it by typing `:g`. Do this for each of the reduced images, (making sure you change the results file name each time.)

You should have four files full of photometric information. All we really want is the summed, sky-subtracted pixel values for each star. Ask your demonstrator how to extract the column containing the pixel values, and transfer them across to the Excel on a Windows PC.

2.7 The Nature of the Stars

2.7.1 The Magnitude System

The magnitude system is a holdover from the ancient Greek method for categorising stars. In this system the brightest star in the northern sky, Vega, was classed "magnitude 0". Fainter stars had higher magnitudes, with each increment in magnitude roughly corresponding to a factor of two reduction in brightness. With the advent of electronic detectors it was found that the ancient Greeks weren't all that good determining "a factor of two in brightness". However it turned out that they were very good at being consistent in their inaccuracies — each magnitude corresponded very closely to a factor of 2.51 difference in brightness. (More precisely, 5 magnitudes turned out to be factor of 100 in brightness, and $100^{1/5} = 10^{0.4} \approx 2.51$). In other words, the magnitude m of an object is related to its brightness f by

$$f = 2.51^{-m+c} \tag{1}$$

where c is a constant which we're going to determine. For two sources, the difference in their magnitudes is proportional to the ratio of their fluxes:

$$\Delta m = m_1 - m_2 = 2.5\log_{10}(f_2/f_1) \tag{2}$$

2.7.2 The Standard Stars

To figure out the magnitudes of our stars we need to calibrate the arbitrary flux units we've measured into something more physical. In essence, we need "c" from the above equation. This constant depends on the size of the telescope (big telescopes collect more light, so you get more counts), the efficiency of the CCD (more efficient detectors get more counts), the filter (even the best filters loose some of the light), and the atmosphere (haze or dust absorbs or scatters the starlight before it reaches the telescope). Astronomers use images of standard stars to perform this calibration. These are stars of known brightness, whose magnitudes and locations on the sky are catalogued in long lists.

In our case we don't need a separate standard star image, because we know the exact brightness of one of the stars in M93. This is the star indicated in the image in the previous section, which, hopefully, you placed at the top of your list. The apparent magnitudes for this star in our four bands are: $B = 10.23$ mag, $V = 10.14$ mag, $R = 10.15$ mag and $I = 10.18$ mag.

Use your measured fluxes for this star, along with these apparent magnitudes, to determine the calibration constant for each image. Each constant is valid for all the other stars in that image. Now calculate the apparent magnitudes of all your stars in each image. You can use the calibration constants and Equation (1) to calculate the apparent magnitudes of all of your stars, or you can just take the ratio of fluxes and use Equation (2).

From apparent magnitude we can calculate absolute magnitude (M). Absolute magnitude is defined as the magnitude a source *would* have if it were at a distance of 10 parsecs. You can find absolute magnitudes using this equation:

$$M = m + 5 - 5\log_{10} d \tag{3}$$

where d is the distance to the star in parsecs. This, of course, relies on the assumption that all the stars are at the same (known) distance. Can you figure out how this relation is derived?

Since magnitudes are logarithmic, the distance scaling factor ($(d/10pc)^2$) becomes an additive constant ($5 + 5\log_{10} d$) between apparent and absolute magnitudes. This factor is called the "distance modulus". Using the known distance (1104 pc) to the cluster, calculate the absolute magnitudes of your stars for all filters.

2.7.3 The Hertzsprung-Russell Diagram

You now have magnitudes for a group of stars, using two or more filters. You can use them to create a Hertzsprung-Russell diagram, a plot of 'colour' against brightness.

'Colour' in astronomy is not colour like the colours of the rainbow. It is defined as the *ratio of how bright a star is when viewed through two different filters*. Remember that ratio of brightnesses is the same as the difference in magnitudes (just from the usual logarithm laws). If, for example, the magnitude of a star observed through the green filter is V, and its magnitude when observed through a blue filter is B, then its colour could be measured as $B - V$. If $B - V$ is large, the object is relatively brighter in green light, compared to blue light. Astronomers would say it had a "red" colour, even though it is actually the blue and green light which is being considered. If $B - V$ was negative, it would have a "blue" colour.

For all the objects in your CCD frames, compute colours, and plot these colours against one of the absolute magnitudes (eg. M_V against $B - V$). This is an HR diagram!

Can you see a line of stars; a main sequence? If you see a main sequence, do you see stars in any other regions in the diagram? Compare your diagram to a colour-magnitude diagram in one of the text books and see if you can figure out what types of stars are present in this cluster. Note that the HR diagrams in the textbooks will probably cover a wider range of abs magnitude and colour than your plot, so be careful when identifying stars which are off the main sequence.

Extra question: Do you get the same or similar H-R diagram if you use a different set of magnitudes? (ie. R and I instead of B and V) If there is any discrepancy, can you explain why? You will need to think carefully about this: what exactly does colour measure? If stars are just black-bodies, when would you expect colours to be the same or different?

2.7.4 How Luminous are the Stars?

To work out how bright our stars really are, we must first convert from magnitudes into energies. The star Vega is magnitude zero in all filters (by definition), so we can use its known brightness at different wavelengths (Table 2) to calibrate our magnitudes to SI units. These *brightnesses* are in units of flux per unit wavelength; ie. the energy hitting a square metre every second with wavelengths in a 1nm range. The *luminosity* of a star is just in energy per second per unit wavelength; ie. it is the total energy the star radiates at this wavelength, not just the amount that hits a square metre area at the Earth.

Using Table 2 and the definition of magnitudes discussed in the preceding section, work out the flux your stars are producing at the different wavelengths. This tells you how much radiation from the stars are reaching the Earth. Using the distances to the clusters listed in Table 1 and the inverse square law, work out the total luminosity of the stars. How do they compare with

Filter	Central Wavelength	Brightness of Vega	Luminosity of the Sun
B (deep blue)	440nm	$6.3 \times 10^{-11}\mathrm{Wm^{-2}nm^{-1}}$	6×10^{23} W/nm
V (green)	550nm	$3.6 \times 10^{-11}\mathrm{Wm^{-2}nm^{-1}}$	6×10^{23} W/nm
R (dark red)	700nm	$2.2 \times 10^{-11}\mathrm{Wm^{-2}nm^{-1}}$	4.5×10^{23} W/nm
I (infrared)	880nm	$1.1 \times 10^{-11}\mathrm{Wm^{-2}nm^{-1}}$	

Table 2: The four standard filters, and the brightness per unit wavelength of Vega when observed through each.

the sun's luminosity? Are these giant stars, or are we seeing things like the Sun? Or are we seeing things fainter still, like brown dwarfs, red dwarfs or white dwarfs?

The very brightest stars known may have luminosities as much as a few thousand times that of the Sun. These are the rare supergiants; some red (very old stars bloating out in their dying moments) and some blue (very young, very hot violent stars). Do you have any of these? If so, are they red or blue?

2.7.5 Stellar Temperatures

The radiation from most stars can be approximated as black body radiation. We can use this fact, along with measurements of the stars luminosity through two different filters, to measure the temperature of the stars.

The black body radiation law (Planck function) is

$$F_\lambda = \frac{2\pi c^2 h}{\lambda^5} e^{-\frac{hc}{k\lambda T}} \tag{4}$$

where F_λ is the flux per unit wavelength emitted by a unit area, h is Planck's constant, c the speed of light, k the Boltzmann constant and T the temperature.

Use this equation to derive the ratio in flux at two different wavelengths as a function of T. Hence, with this equation, take the magnitudes measured through two different filters and use this to derive the temperatures of your stars.

How hot are they? The Sun has a surface temperature of about 6000K. How do the temperatures compare to that of the Sun? Do they all have similar temperatures or is there a range? Is there a correlation between brightness and temperature?

Use the Planck function to derive the surface area of your stars and hence their radii (be careful converting units!). For comparison, the Sun has a radius of 7×10^8 m. How do they compare?

2.8 Advanced Exercises

If you've got this far and still have time and energy left, consider a few of the following.

- If the sun were taken away and replaced by one of your stars, how hot would the Earth get? (the Sun is roughly 8 light-minutes away). How far away would a planet have to be from this sun to sustain liquid oceans? (Hint: assume planets are black bodies and use the Stefan-Boltzmann equation).
- Some of your stars are very luminous. The luminosity of a star goes as roughly their mass to the power 2.5 (ie. $L \propto M^{2.5}$) for main sequence stars like the sun and blue supergiants. How massive are these things compared to the Sun? The Sun has an expected lifetime of 10^{10} years: how long will these ones last?
- How near would one of these stars need to be, to be visible in the daylight? (Hint: the moon is only just visible in daylight).
- Precisely how bright is the night sky over Siding Spring Observatory? (Express this in magnitudes per arcsec^2)

3 Galaxy Clusters and Cosmology

Figure 3: The centre of the Virgo cluster of galaxies. Image courtesy of Jean-Charles Cuillandre (CFHT), Hawaiian Starlight, CFHT

3.1 Aims

To use observations of galaxies and clusters of galaxies to estimate Hubble's constant and the age of the universe.

3.2 Background

Cosmology is the branch of Astrophysics concerned with studying the universe as a whole. Key questions include: Where did the universe come from? Does it go on for ever? Will it end some day? Many physicists are shocked by the poor quality of the data; if you don't throw up your hands in disgust and scream "you can't assume that!" at least once during this prac, you won't have understood what you're doing. Yet the truly remarkable thing is that cosmology is possible at all; that the microscopic inhabitants of this insignificant speck of dust called the Earth can dare to try and understand the nature of the universe.

In this practical, you will use observations of two galaxy clusters to estimate several key cosmological parameters. The observations were made with the UK Schmidt Telescope at Siding Spring Observatory in NSW; see the asteroid hunting exercise for details of the telescope and the photographic plates it produces – photos of this telescope are displayed in the lab.

Most galaxies live in groups. Our own galaxy, the Milky Way, lives in a rather insignificant cluster we call the Local Group. This cluster has two big galaxies, our own and the Andromeda Galaxy (M 31), in orbit around each other. Each of these two big galaxies is a spiral galaxy, with a radius of about 10 kpc (kilo-parsecs, one parsec = 3.1×10^{16} m). We live about 8 kpc from the centre of our galaxy, so we are right out near the edge. In addition to the two biggies, there are two medium sized galaxies in our group, the Magellanic clouds, and another 10 or so dwarf galaxies, all in orbit around each other. The radius of our local group is about 1 Mpc (Mega-parsec).

When astronomers first started mapping the sky, looking for galaxies, they discovered an enormous number of them in the constellation of Virgo. There were so many galaxies that the whole area of the sky was called "The Sea of Galaxies". We have a photographic plate of the centre of this area, the plate marked 'Virgo Cluster', and/or J2137 or 12:27+13:30. Find this plate and put it on the light table. Use a magnifying lens and have a look at some of the galaxies. There are well over a thousand on this plate, though many are too faint to see without a microscope.

The Virgo cluster is a big cluster of galaxies which happens to be quite close to us, which is why it appears so large and spectacular. In fact, it has recently been discovered that our whole Local Group is probably falling into the Virgo Cluster, sucked in by its gravity.

(By the way, Virgo is by no means the biggest galaxy cluster known. In fact, it is a bit of a runt; it only looks spectacular because it is very close to us. If you look hard at the Virgo cluster plate, you will see lots of tiny galaxies in the background – try it. These are part of an absolutely colossal supercluster of galaxies, known as the 'Hydra-Centaurus Supercluster', of which Virgo is merely a tiny outlying part. And evidence coming in at the moment suggests that even this supercluster is just part of an even bigger sheet of superclusters, known as the 'Supergalactic plane', hundreds of Mpc across. Even this Supergalactic plane seems to be just one of many in the universe, forming a bubble-like network on the biggest scales we can currently observe.)

The Virgo cluster is so close to us that it doesn't all fit even on a Schmidt plate – this picture only shows the centre. So in this exercise we will concentrate on a different cluster, Hydra I (also known as Abell 1060, being the 1060th cluster in Abell's catalogue of clusters). We have a plate of this, called Abell 1060, or J7442 or 10:34-27:22. Get it out on the light table and have a look. The cluster of galaxies is obvious in the centre.

Hydra I is a cluster very similar to Virgo. It is about the same size, same mass, and has a similar number of galaxies. However, it is much further away, as is obvious from the plate. This means the whole cluster fits comfortably on one Schmidt plate.

3.3 Distance to Hydra I

In the first half of this exercise, we will estimate the distance to the Hydra I cluster. Getting distances is always the most difficult exercise in astronomy; in many ways it resembles a black art more than a science. We can't take a ruler to anything outside our Solar System (our rockets are much too slow), and triangulating only works to distances of about 20 parsecs, so we have to use guesswork. The method is this; we assume that we know some property of a distant object as it really is. For example, we might assume that we know how bright a particular distant star is. Then we measure how bright it appears from the Earth. The ratio of the two gives us the distance to the star.

The method we will use in this exercise is galaxy diameters. If we knew how big a galaxy really is, and could measure how big it appears, by simple trigonometry we could work out how far away it is. We can measure the apparent size of the galaxies in Hydra I using the magnifying eyepiece with a measuring scale. Try measuring the size of a few. The plate scale of all the plates is 1.12 arcmin per mm. So if you measure a galaxy as being 2mm long, its apparent size will be the angle 2.24 arcmin.

So we can measure the apparent size of galaxies in the cluster. But what is their real size? How big is a galaxy? We only really know the size of one galaxy at all well, our own, the Milky Way. Here's where we make our first whopping great assumption. Why don't we assume that the one galaxy we know anything about, the Milky Way, is perfectly typical? So let us assume that all the galaxies in Hydra I are just like ours; they have radii of 10 kpc.

So pick a bunch of galaxies in Hydra I. Measure their diameters. Some of them will be nearly edge-on, but that shouldn't stop you measuring them. Multiply by the plate scale to get their angular size. Now work out how far away they are! You can get an answer for every galaxy you measure, and they won't all be the same. Throw out any ludicrously discrepant values, and average the rest to get your best guess of the distance to Hydra I.

One big problem is measuring the galaxy sizes. After all, galaxies don't have sharp edges, they just fade gently out. Where then do you measure the edge? There isn't really a good answer. The most you can do is be self-consistent.

What distance do you get? Current best estimates lie in the range 35 — 80 Mpc, but don't worry if your answer is outside this, it is a very difficult measurement (I got 30 Mpc when I tried it). How might you go about improving your estimate?

By far the worst feature of this calculation is the assumption that these galaxies out in a distant cluster are just like our own galaxy. There are all sorts of reasons why this may not be true. For one, this cluster is a lot bigger than our local group. Thus just by chance, you might expect it to include a few rather large galaxies. Also, nobody understands how and why galaxies form, and it is quite possible that they form differently in big clusters. So these galaxies may be nothing like our galaxy. How can we get around this problem?

One way is to be selective in our choice of galaxies to measure. We know that our galaxy, the Milky Way, is a spiral galaxy. So try measuring a distance using only spiral galaxies in Hydra I; see the appendix for a description of different galaxy types. Has your answer changed? We also know that the Milky Way is the largest spiral in our local galaxy cluster. So perhaps it is similar to the largest spiral galaxies in Hydra I. Try getting a distance just from them.

Even this doesn't necessarily help us very much: we've no good reason to believe that even large spiral galaxies in Hydra I should be like the Milky Way. A much better assumption would be that the galaxies in Hydra I are just like the galaxies in Virgo. After all, both clusters are about the same size, mass and density. It seems very plausible that the galaxies in Virgo should at least resemble those in Hydra 1.

Here is a better method (using Virgo) to try and measure the distance to Hydra I:

1. Pick the ten biggest spirals and ten biggest ellipticals in each cluster. Throw out any that look too bizarre – you want to compare the two clusters, so if in one you find, say, some seriously warped colliding galaxies which aren't in the other cluster, you should throw them out. If you can't find ten good ones, choose eight or five – if you can find more than ten, using them all will improve the accuracy of this method.
2. Measure the sizes of all these galaxies, using the eyepiece with a scale. It doesn't actually

matter how you do this, so long as you do it the same way on each cluster. Try and think of a consistent way of measuring the sizes.

3. Find the average apparent size (in arcmin) of the ten spirals and of the ten ellipticals in each cluster, r_v^s (mean size of Virgo spirals), r_h^e (mean size of Hydra ellipticals) etc.

4. For an object of a given size r at a distance d, the angular size will be $\theta = r/d$ where θ is measured in radians (why?). This only applies if d is very much greater than r which is certainly true here. Therefore, if we assume the ten brightest of each type of galaxy are the same size in each cluster, the ratio of the distance to Hydra I (d_h) to the distance to Virgo (d_v) is given by

   $$\frac{d_h}{d_v} = \frac{\theta_v^s}{\theta_h^s}$$

   from the spirals, and

   $$\frac{d_h}{d_v} = \frac{\theta_v^e}{\theta_h^e}$$

   from the ellipticals. These two numbers should be the same, to within the errors. Are they? If you think one is more accurate, use it (justify this decision in your notes). Otherwise, average the two numbers. It is always worth measuring anything two ways to have a check.

5. The distance to Virgo is controversial, but most current estimates are around 15 Mpc. Use this, plus the ratios measured above, to find the distance to Hydra I.

How does this answer compare with the earlier one? If they are different, what is this telling us about the relative sizes of galaxies in Hydra I and our galaxy? It could be that our distance to Virgo is wrong, as this number is highly controversial. Now that the Hubble Space Telescope has been fixed, an international team (including several Australians) has been using it to spot individual stars in Virgo galaxies, which will enable them to measure the distance to Virgo accurately at last – expect results in a year or so.

3.4 Hubble's Constant

Reference: Zeilik, Gregory & Smith, Introductory Astronomy & Astrophysics, Chapter 22.

The universe is expanding; every galaxy is moving away from every other galaxy, and the further two galaxies are apart, the faster they are receding from each other. This is known as Hubble's Law, after the famous American astronomer who discovered the expansion of the Universe back in the 1920's. It can be written as:

$$v = H_0 d \tag{5}$$

where v is the velocity at which something is moving away from us, d its distance from us, and the constant H_0 is known as Hubble's constant. This constant is basically a measure of how

Galaxy	X	Y	magnitude	v (kms^{-1})
	mms from bottom left corner			
NGC 3285A	211	178	14.0	not known
NGC 3285	202	182	13.2	3049
NGC 3285B	190	171	14.0	2868
NGC 3305	171	198	14.0	4549
NGC 3307	170	178	16.0	3616
NGC 3308	169	183	13.0	3687
NGC 3309	166	179	12.7	3801
NGC 3311	165	178	13.0	3575
IC 629	169	175	16.0	2461
NGC 3312	161	176	13.1	2512
NGC 3314	159	170	14.0	2635
NGC 3315	158	181	15.0	4555
NGC 3316	154	174	15.0	3752
IC 2597	152	202	13.0	2738
NGC 3336	123	164	13.0	3689

Table 3: Galaxy Velocities for members of Hydra I

fast the Universe is expanding. If it is large, the universe is growing like crazy. We are going to use the distance we measured to the Hydra I cluster to estimate H_0.

We have already measured the distance d to Hydra I. Now we need the velocity at which it is rushing away from us v. This is measured using the Doppler effect. The gas in most Hydra I galaxies emits emission-lines of Oxygen, Hydrogen and other elements. With a spectrometer on a telescope, we can measure the Doppler shift in these lines, and hence the velocity v. These velocities are listed in Table 3, for several of the brightest galaxies in Hydra I.

Find the average velocity with which Hydra I is running away from us (why don't all the galaxies in Hydra I have the same velocity?) Use the average velocity to find Hubble's constant. What value do you get? H_0 is traditionally expressed in units of kilometres per second per Megaparsec.

The debate over the true value of Hubble's constant is without a doubt the longest and bitterest in the recent history of astronomy. For a racy account of the "Hubble Wars" read 'Lonely Hearts of the Universe' by Dennis Overbye, and several recent articles in New Scientist and Scientific American. One group of researchers, prominent amongst whom is Alan Sandage of California, regularly find low values, one of their latest being $H_0 = 42$ kms^{-1}Mpc^{-1}. Another group, including Jeremy Mould (until just recently director of Mt Stromlo Observatory near Canberra) keep finding higher values, around $H_0 = 80$ kms^{-1}Mpc^{-1}. The most recent Hubble Space Telescope result was $H_0 = 68 \pm 6$ kms^{-1}Mpc^{-1}. Who do you agree with? If your result is not near any of these, take heart, Edwin Hubble's first measurement was $H_0 > 400$ kms^{-1}Mpc^{-1}!

3.5 Age of the Universe

Now you've measured Hubble's constant, you can estimate the age of the Universe. We see everything rushing away from us, with a velocity proportional to its distance from us. This means everything was closer together in the past. Let's say a galaxy is 100 Mpc away from us today, and that $H_0 = 100$ kms^{-1}Mpc^{-1}. 100 Mpc is

$$100 \times 10^6 \times 3.1 \times 10^{16}\text{m} = 3.1 \times 10^{21}\text{km}$$

Because the galaxy is 100 Mpc away, Hubble's law tells us it is moving away from us with a velocity

$$v = H_0 \times d = 100 \times 100 = 10^4\text{km}s^{-1}$$

Now extrapolate backwards (assume that the galaxy has always moved with the same speed). How long ago was the galaxy right here? The answer t is

$$\frac{3.1 \times 10^{21}}{10^4} = 3.1 \times 10^{17} s = 10^{10}\text{years}$$

Try the calculation for galaxies at different distances – they were always here the same time ago (why?). So what was this time in the past where all the universe was piled on top of us? The Big Bang! We've found out how long ago the Big Bang was. Do the calculation for Hydra I, using your value of its distance and Hubble's constant. How old do you make the universe?

This measurement of the age of the universe is very crude, because it assumes that the universe has been expanding at exactly the same rate ever since the Big Bang. If you do a detailed calculation, solving Einstein's equations for the Universe, you find that the universe was expanding faster earlier. This means that the age you calculated is too big (why?). The correct answer turns out to be 2/3 of the age you calculated. Multiply your answer by 2/3 to get your best estimate of the age of the universe.

Compare the age of the universe with various other ages. The age of the Earth is thought to be about 4.6×10^9 years. The age of the Sun and the rest of the solar system is probably similar. The age of our galaxy is estimated as about 10^{10} years, but some star clusters orbiting our galaxy (globular clusters) are thought to be at least 1.5×10^{10} years old (their age is estimated by computer modelling the nuclear burning of their stars). Is this embarrassing for your value of the age of the universe? What is the age of the universe if $H_0 = 80$ kms^{-1}Mpc^{-1}? Anything wrong with this?

3.6 Advanced problems

Make a map of Hydra 1. Get some overhead projector sheets and washable marker pens. Put the transparent sheet on top of the Schmidt plate, and mark the positions of the galaxies. You might like to use a colour code, with different colours for spirals, ellipticals, and things just too darn faint to decide. Don't bother mapping the edges of the plate – concentrate on the centre where the cluster is.

You can't possibly map every galaxy on the plate – most are so small you would need a microscope to tell them apart from the stars. So only do the bigger ones, but try and be consistent; if you include smaller galaxies in one part of the plate than another, it will introduce spurious features into your map.

How big is the cluster (in Mpc, using your distance to the cluster)? How many galaxies does it have? What fraction are spirals and ellipticals? Do both types of galaxies live equally in all parts of the cluster? Is the cluster nice and round, or is it distorted? If a cluster has distortions and lumps, it tells you that it is a newly formed cluster, that has not had time to settle down in equilibrium.

3.7 Appendix 1: Galaxy types

Galaxies come in two main types, spirals and ellipticals. Nobody knows why. Ellipticals are fuzzy balls of stars, some spherical, others elongated like buns (oblate) or AFL footballs (prolate). They don't have any structure, and no matter what angle you look at them from, they just look like fuzzy elliptical blobs. Spirals, on the other hand, have much more structure. They are dominated by a flat disk of gas, stars and dust, looking a bit like a gramaphone record. If you look at this disk face-on, you can see all sorts of pretty patterns. Many spirals (but not all) show nice spiral patterns in this disk. Others show rings, or bars, or combinations of these. Plenty of spectacular examples of all these types can be found on the Virgo plate. In the centre of spiral galaxies, in addition to the disk you have something called the bulge, which is a lot like a small elliptical galaxy sitting at the centre of the spiral. This bulge is hard to see in a face-on spiral galaxy, but if you look at an edge-on one it can be pretty obvious – you have the flat disk, looking like a line, with a blob in the middle, the bulge. Ask your demonstrator to show you some nice examples.

The main theory for explaining these two types of galaxies these days is a violent one. The idea is that if you collide two spiral galaxies, all the stars get so muddled around, you end up with an elliptical. Can you see any galaxies that look like they are colliding on the plates? However, several very good astronomers don't believe this theory. The debate goes on…

4 Asteroid Hunting

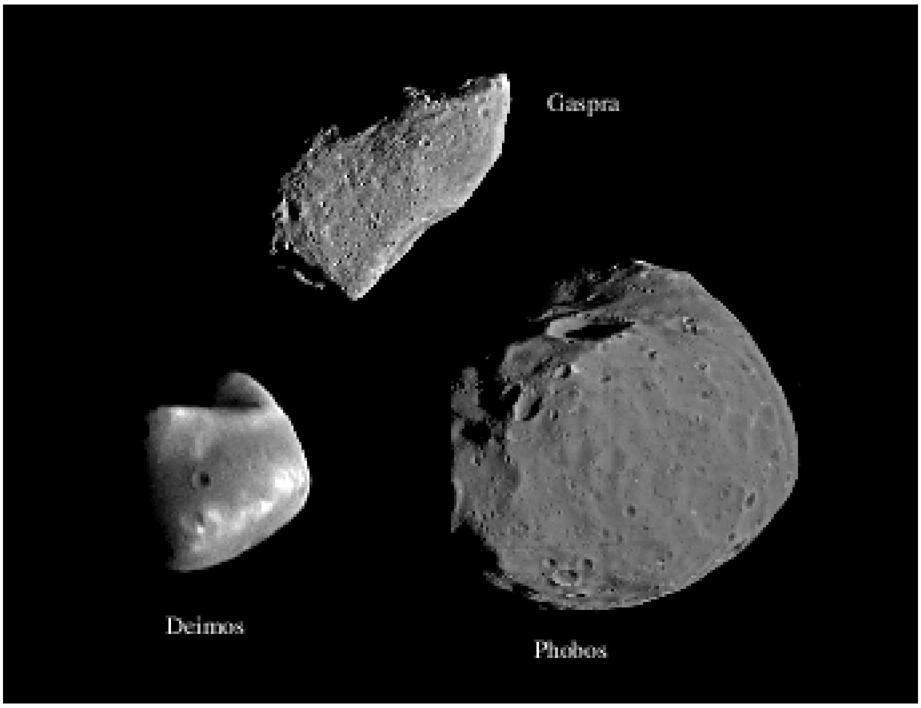

Figure 4: This image shows the asteroid Gaspra as well as the Martian moons Phobos and Deimos on the same scale. Gaspra is about 20km long. Images courtesy of solarsystem.nasa.gov

4.1 Aim

To find lots of uncatalogued asteroids. To work out where the asteroid belt is, and roughly how thick it is. To put a lower limit on the number of asteroids in the asteroid belt. You should also learn a lot about measuring photographic plates.

4.2 Background

Most large professional telescopes look at really tiny areas of sky in immense detail. This is great for studying a particular planet or galaxy, but for certain problems is rather bad. One such problem is finding asteroids. Contrary to the impressions given by various science-fiction movies, asteroids are quite spread out, so you could spend years waiting for an asteroid to blunder past the particular tiny bit of sky you are staring at.

Luckily, some clever telescope designers have come up with a solution. They devised a special type of telescope called a **Schmidt Telescope**. This has specially designed optics to take pictures of very wide areas of sky, allowing you to look for rare or spread-out objects.

You will be using data taken with the best Schmidt telescope in the world, which is jointly run by the UK and Australia, and is sited at Siding Spring Observatory, near Coonabarabran in NSW. Since it was completed in 1973, this telescope has been systematically mapping the sky. Each picture it takes covers about $6^\circ \times 6^\circ$ (for reference, the diameter of the moon is about

0.5°). In collaboration with a northern telescope (at Mt Palomar in California) the whole sky has now been photographed. The photographs are called plates, as they are thin glass sheets covered on one side with photographic emulsion.

In the lab, we have a collection of film copies of Schmidt plates. You should browse through some of them – there are some amazing pictures of galaxies, nebulae, comets etc. See how much data is on each picture; there can be as many as a million stars recorded on one picture of a bit of the Milky Way (nearly all those little dots really are stars, not dirt!). Note that these film copies are research grade material, not special teaching stuff – a lot of professional astronomers spend their lives poring over these plates trying to find quasars, galaxies, comets, the tenth planet, sputniks and so on.

Important number: the plate scale of pictures taken with this telescope is 1.12 arcmin per mm.

4.3 Measurements

You will be using plate number J2137 (also known by the coordinates of its centre on the sky, 12:27+13:30). This plate was taken for a study of a giant cluster of galaxies, the Virgo Cluster, which you can easily see in the middle. This picture was taken with a 70 minute exposure, and objects as faint as about 22nd magnitude can be seen (if you try real hard). South is at the bottom of the picture, and East is to the left. Positions on the plate should be quoted as the number of mm right and up from the bottom left-hand (SE) corner.

This plate was taken with the telescope tracking – ie following the stars as they move across the sky. This means that stars, galaxies, quasars etc appear as nice sharp images on the plate. However, if something up in space moved appreciably during the exposure, it will appear not as a dot but as a line.

This is how you hunt for asteroids; you search for things on the plate which look like short lines - about 1 mm long. Put the film copy of the plate on the light table. Use a magnifying lens and see if you can find any (hint - there are several dozen on the plate). If you have trouble, ask your demonstrator to point a few out to you. These are not known asteroids – only one of the asteroids on this plate is in any catalog (asteroid Suleika, 230mm right and 212mm up from the bottom left-hand corner). So these asteroids you are finding have no names, not even numbers, and nobody knows anything about them.

Once you've found an asteroid, make a note of where it is so you can find it again. The best way to do this is with an overhead projector transparency and a washable marker pen. Please don't write directly on the plate, or write on the overhead sheet with a pen or anything hard as you will damage it. To get your overhead projector sheets aligned right, mark a few really bright stars or galaxies on them.

As you will soon find out, there is a lot of crud on the light table and the film covers. Often things you think might be asteroids will turn out to be dirt - try moving the plate around on the light table and brushing both sides before you decide an object is an asteroid. Other things which may look like asteroids at first are edge-on spiral galaxies (there are some beauties in the Virgo cluster) and random groups of stars that happen to line up. With practice you will soon be spotting the real asteroids.

4.4 Where are the Asteroids?

The first exercise with the plate is to work out where these asteroids are in the solar system. The coordinates of the centre of the plate are 12 hours 27 min in right ascension, 13 degrees and 30 minutes in declination (northern hemisphere). See where this is on your planispheres, and on the sky globe in class (now you know why it is called the Virgo cluster). Is it near the ecliptic plane? At what angle does the ecliptic plane lie, compared to the plate boundaries?

The plate was taken on 28th March 1976. Using the Skyglobe program on the PCs, work out where the sun was at on this date (its right ascension and declination). Compare the position of the Sun and of the plate on a globe of the sky. You will see they are almost at opposite locations.

Measure the direction of motion of a few asteroids, using a protractor. Are they moving along the Ecliptic? Most of them will be; these are the asteroids in the asteroid belt, and like most of the planets they are orbiting the Sun in the Ecliptic plane. You may however find a few maverick asteroids that are moving with quite different speeds or directions. These rogue asteroids are most likely nearer the Earth on highly eccentric orbits. One much like this may well have been responsible for eradicating the dinosaurs 60 million years ago. Take careful note of where they are and tell your demonstrator – the class should build up a collection of these.

So, we now know that most of these asteroids are moving along the ecliptic, and as they are in the opposite direction from the Sun, they must be further out that the Earth. Can we work out how far away they are? The answer is yes, but only if we make a couple of key assumptions. Firstly, we must assume that they are in roughly circular orbits – not too eccentric. This is true of all known belt asteroids, though not of the maverick near-Earth ones. We will also assume that they are orbiting along the plane of the Ecliptic, in the same direction as the Earth (and all the other planets). To simplify the maths, we assume that they are exactly opposite the Sun from us and that both the asteroids and the Earth are moving along the Ecliptic in perfectly circular orbits.

The motion of the asteroid as seen from the Earth will have two causes. Firstly, the asteroids will be moving slowly around their orbits. Secondly, the Earth is speeding around its orbit, so from the Earth, the asteroids will tend to move backwards. As the Earth is moving faster than the asteroids, this latter effect is the biggest one, so they will appear to move backwards.

Measure how fast they are moving. You do this by measuring the length of the little line on the plate. As you know the plate scale (1.12 arcmin per mm) and the exposure time (70 minutes), you can work out the angular speed at which the images move. You should do this for a few asteroids – the more you measure, the more accurate your answer will be when you average them all.

Now, using Newton's law of gravity, you can work out how fast the Earth is moving round the Sun (you will need to know the mass of the Sun, 2.0×10^{30} kg, and the distance from the Earth to the Sun, 1.5×10^{11} m). If the asteroid were a distance r from the Sun, how fast would it be moving? And what would its backward angular speed be when viewed from the Earth? Using the observed backward angular speed, solve the equation and Bingo! You have r, the distance of the asteroids from the Sun (The equation in a cubic, and so not trivial to solve. However, we don't need a precise solution, s you can solve it numerically or just plot it as a graph).

So where are they? You now know how far their orbits are from the Sun, compare this with the orbits of the planets. Which planets will they pass closest to? Anything unusual about these planets which might explain why the asteroids are near them?

4.5 The Asteroid Belt

In the first exercise, you discovered that the asteroids were orbiting the Sun, and measured their distance from the Sun. Now, we will work out how thick the asteroid belt is. To do this, you should try and measure every asteroid on the plate. Search for them systematically; figure out some methodical way of working your way across the plate to make sure you don't miss any. The really faint ones will always be hard to find, so you might want to keep a separate list of possible asteroids which you are not sure of.

The first thing to check, once you've measured positions and angular speeds for all the asteroids, it whether they are concentrated near the ecliptic plane. You have already worked out where the plate lies compared to this plane. Divide the plate in half; the half nearest and furthest from the Ecliptic. Count the number of asteroids per unit area in each half. Are there more near the plane? Further from it? Is the number constant? What does this tell you about the shape of the asteroid belt?

The next thing to check is the spread in distances from the Sun. For each asteroid, measure its angular speed, and hence its distance from the Sun. Try plotting a histogram of these values. How spread out are they? Are there any asteroids well outside the normal belt?

As you will be finding out, spotting things on these plates is no slouch. You might want to think about the following biases and how to avoid them. A lot of your marks will depend on how clever you are at this. But don't feel alone and victimised by these problems; a lot of professional astronomers (including the author) have spent a lot of their lives worrying about these biases. It really is the mark of a good astrophysicist to know what is good and bad about your data, to anticipate things that could go wrong and figure out ways of solving them.

- If you just wander around the plate spotting things, you will inevitably pay more attention to some areas than others. Be systematic!
- As you get tired and bored, you will probably start missing asteroids. If you started from the top of the plate and worked down (for example), this might mean that you found a whole lot more asteroids at the top than at the bottom. This might mean you see numbers of asteroids increasing towards or away from the Ecliptic where no such effect really existed. A possible way out it to search the plate in a more random way than just starting at one end and moving through.
- The opposite of the above; as time goes on you get more experienced and spot more asteroids.
- You don't have to measure the whole plate yourself; share your results. But this introduces a new bias. Not everyone will be equally thorough in finding asteroids. You should cross check results. Perhaps you should all do a few overlapping regions, and check that you find the same number of asteroids.

- Fainter asteroids are harder to spot than bright ones, and more easily confused with dust. Make sure you aren't ignoring faint ones or bright ones; they might be in different parts of the belt.
- Short and long asteroid trails may be easier or harder to find. If you are biased in length, what effect will this have on the histogram of distances from the Sun you are measuring?

4.6 Advanced Exercises

You have measured the radial thickness of the asteroid belt. Using your counts of asteroids per square degree, work out the density (asteroids per square km) of the belt. How does this compare with asteroid belts in Sci-Fi movies? If each asteroid has a diameter of 1km and is a perfect sphere, what are the odds of a rocket hitting an asteroid as it shoots through the belt at high speed? Note that this is a lower limit; there could be lots of small dark asteroids too faint to see on this plate.

The apparent magnitude of the asteroid Suleika (230 mm right and 212 up from the bottom left-hand corner) can be measured from this plate using a microdensitometer. It is 13.3. Using this, combined with the distance of Suleika from the Sun (which you can measure from its motion), you should be able to work out how big Suleika is. You will need to assume it is a sphere, and that it reflects 6% of the light that falls on it. It will be helpful to know that the apparent magnitude of the Sun is -26.1. If the density of Suleika is typical of meteor rock (3.5×10^3) what is the mass of Suleika? If all asteroids have the same mass, what is the density of the asteroid belt? Estimate a thickness for the asteroid belt (how high it extends above and below the ecliptic), and work out the total mass of the asteroid belt. How does this compare with the mass of a planet? Is this good news for the theory that the asteroids were formed when a planet exploded?

5 Acknowledgements

Updated Feb, 2005: Randall Wayth

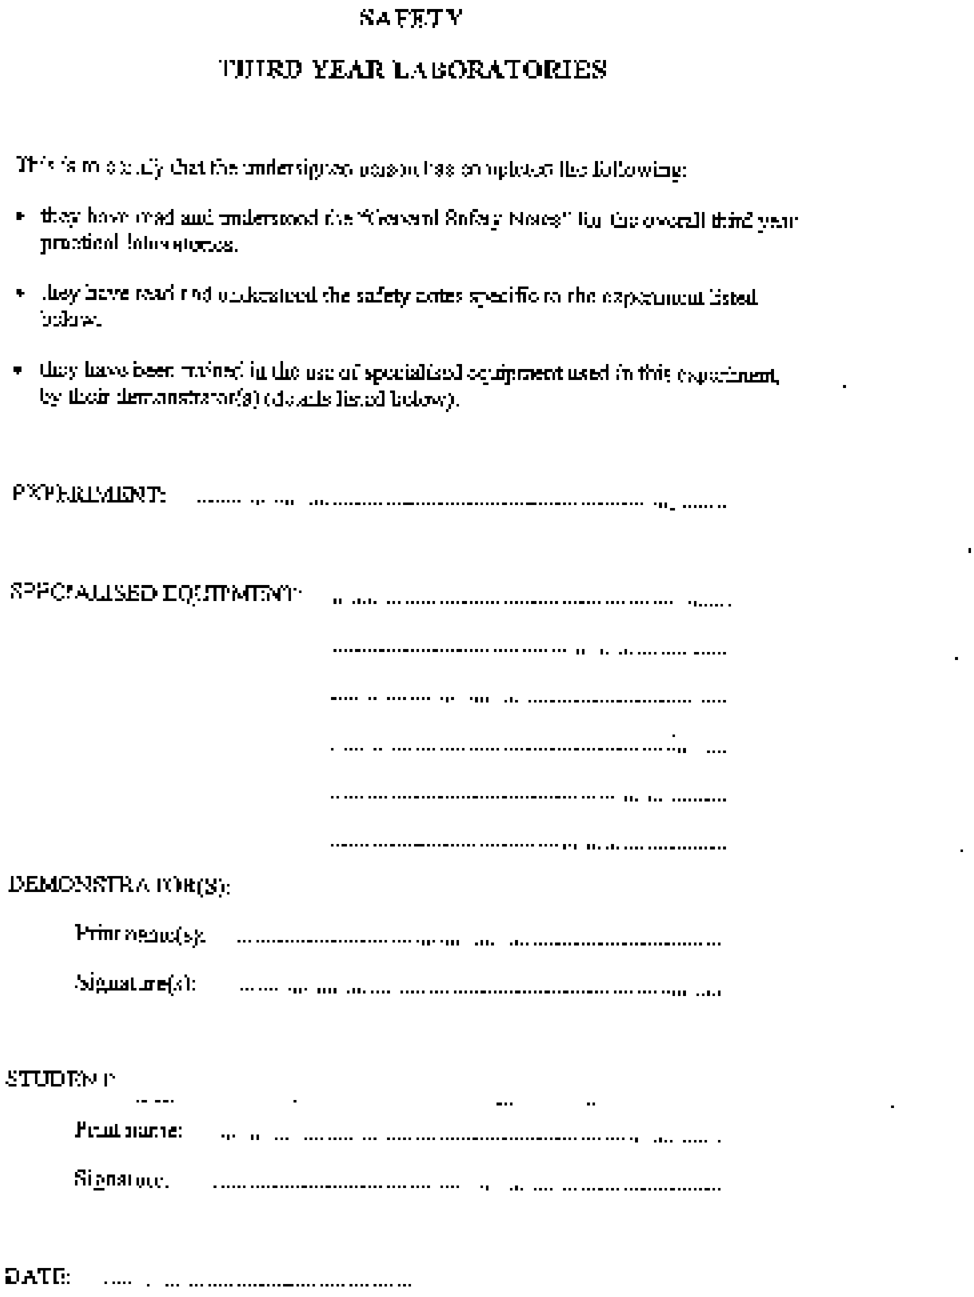

SAFETY

THIRD YEAR LABORATORIES

This is to certify that the undersigned person has completed the following:

- they have read and understood the "General Safety Notes" for the overall third year practical laboratories.
- they have read and understood the safety notes specific to the experiment listed below.
- they have been trained in the use of specialised equipment used in this experiment, by their demonstrator(s) (details listed below).

EXPERIMENT: ....

SPECIALISED EQUIPMENT: ....
....
....
....
....
....

DEMONSTRATOR(S):

Print name(s): ....

Signature(s): ....

STUDENT:

Print name: ....

Signature: ....

DATE: ....

Figure 5: Complete this safety form, detach it and hand it to your demonstrator.