

Texas Tech University Department of Physics & Astronomy
Astronomy 2401 Observational Astronomy
Lab 3:- CCD Image Processing

Objectives:

There are three principle objectives for this laboratory session.

- a) To become familiar with some of the image processing functions of CCDSoft.
- b) To understand the processes by which various types of noise are removed from a CCD image.

Introduction:

Charge-coupled devices (CCD's) are now the standard tool of astronomy for recording images of the sky. A CCD is a small light detector, measuring up to 50 mm on a side and made up of an array of up to several million micrometer-sized light collectors or photocells. Each photocell on the CCD corresponds to a pixel in the image. The photocells themselves are very small Metal Oxide Semiconductor (MOS) capacitors, which are all linked by very tiny circuits on a silicon chip. Charge is created in each MOS capacitor by the photoelectric effect whereby the light striking the silicon, removes electrons from it. These electrons are stored in the silicon substrate of each pixel until read out by the software controlling the CCD.

At the output stage the charges are converted into a voltage which is measured at the output pin of the CCD. Each row of charge is transferred across the CCD, one row at a time, and read as a " ΔV " (change in voltage) from a predetermined reference level. The ΔV is directly proportional to the charge which accumulates in each photocell and the charge is directly proportional to the number of photons collected by the photocell. The registers at the output stage are emptied in a finite period of time until all the rows of charge have been transferred across the CCD. This gives rise to a time-varying voltage, or ANALOG signal. This analog signal must be converted into a digital signal in which the ΔV s are translated into a number between 0 and 65535 (2^{16} levels). These numbers become the pixel intensity values in the image which you see. These pixel intensities are then stored, in order of their position on the chip, in a file on the computer.

The biggest advantages of using a CCD to image the sky include

- a) High sensitivity, CCD's are able to absorb up to 90% or more of the photons reaching it compared to less than 10% for film.
- b) Linearity, the charge output of each pixel is directly related to the intensity of the light to which the pixel is exposed.
- c) The output of the CCD is in digital form and so readily suitable to computer analysis.

However, despite these advantages, a CCD is still not a perfect sensor. Each image will contain a number of defects that require attention. The removal of these defects is the first goal of image processing called **calibration**.

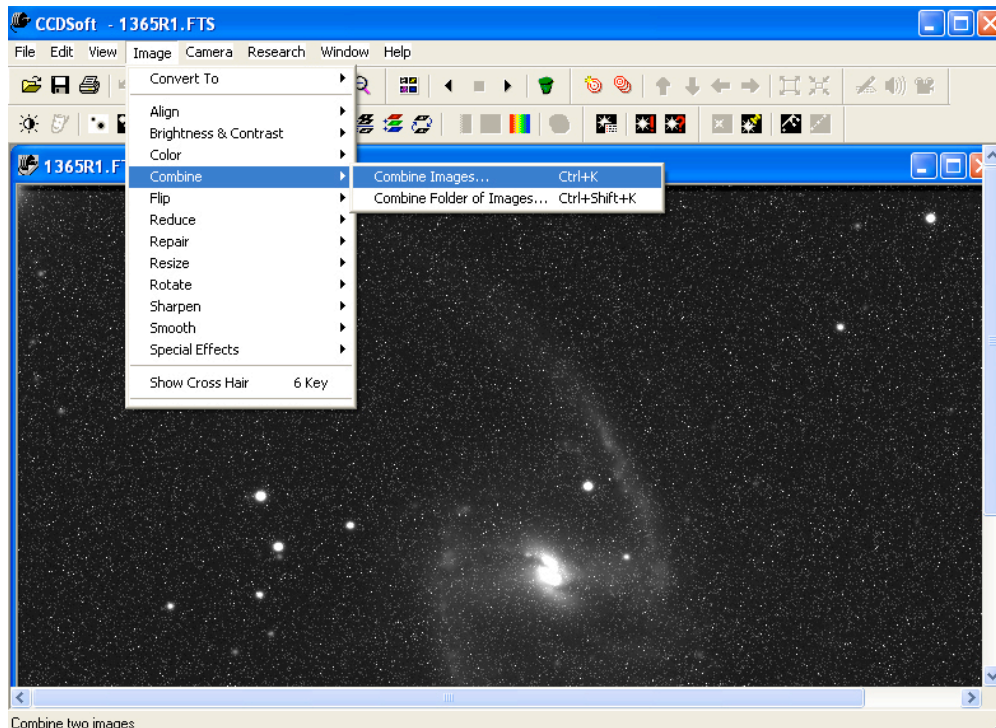
Calibration:

Calibration involves minimizing a variety of image defects that are inherent to the CCD itself. The first of these is **Dark Current**. Being a MOS device, the CCD chip generates random charge within the substrate. This charge is temperature dependant, and so can be reduced, but not eliminated, by cooling the chip. This dark current shows up in the image as noise and an elevated background. It is removed by taking what is known as a **Dark Image**, that is, an image of the same duration and at the same temperature as the image of the sky, but with the shutter CLOSED so that no light reaches the chip. In this way the image will consist only of the electronic noise generated within the chip. This dark frame is then subtracted it from the image of the sky to remove the electronic noise.

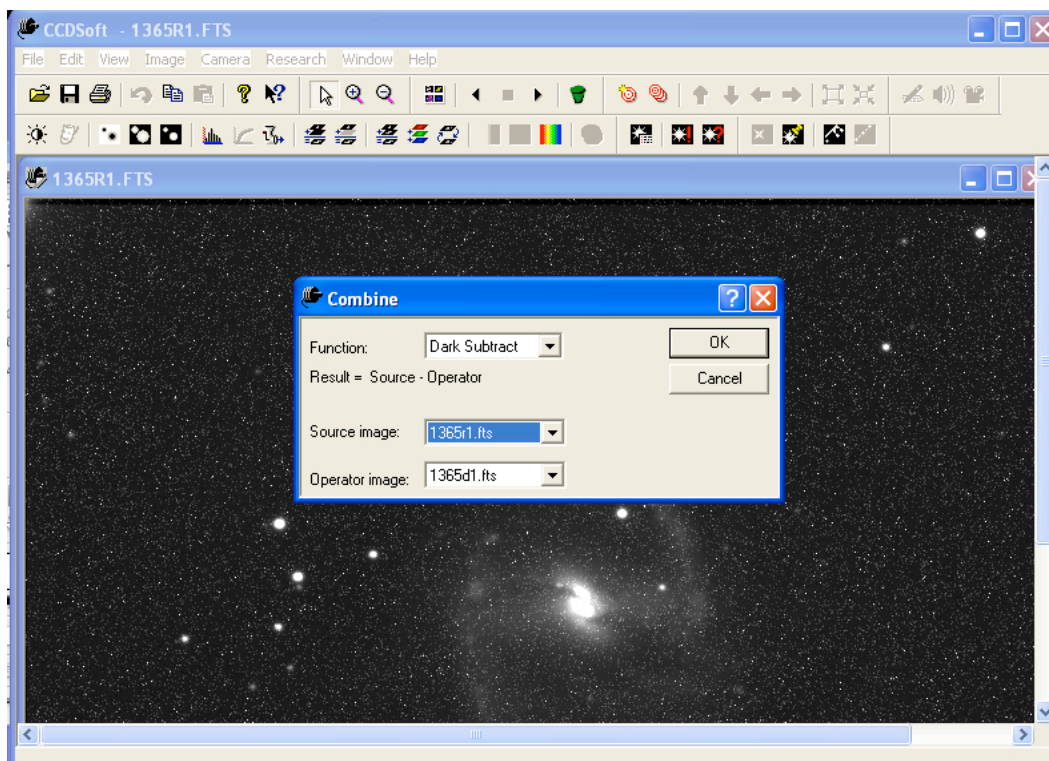
Make yourself a folder in the “**as2401**” directory. Then open the folder “**testimages**” and the subfolder “**galaxies**”. This folder contains a number of images of galaxies that Dr. Clark has taken in the past. **Copy** the images:- 1365D1.FTS; 1365R1.FTS; 1365R2.FTS; 1365R3.FTS; 1365R4.FTS: into your folder.

Open CCDSoft and then open the image 1365R1.fts in the folder you just moved. This is an image of the beautiful barred spiral galaxy NGC 1365. As you look at the image, you can see that there are a few stars visible and the galaxy is faintly visible, however there are numerous bright pixels. These are referred to as **Hot Pixels**, and are noise. Now open the image 1365D1. (In these images, the ‘r’ indicates that the image is a ‘raw’ image (ie unprocessed light image and the ‘d’ indicates that the image is a ‘dark image’.) Looking at this image, notice that it contains hot pixels in the same positions as in the other image.

To perform the subtraction, go to the “**Image**” menu and select “**Combine**” and then “**Combine Images**”.



This will bring up a box with various functions. In the box labeled “function”, select “Dark Subtract” from the options. Below this is the process that the software will perform. It now should read, “**Result = Source – Operator**”. Since you are subtracting the operator from the source, you should set the source image to be the image of the sky, in this case, 1365R1, and the operator should be the dark image, 1365D1.



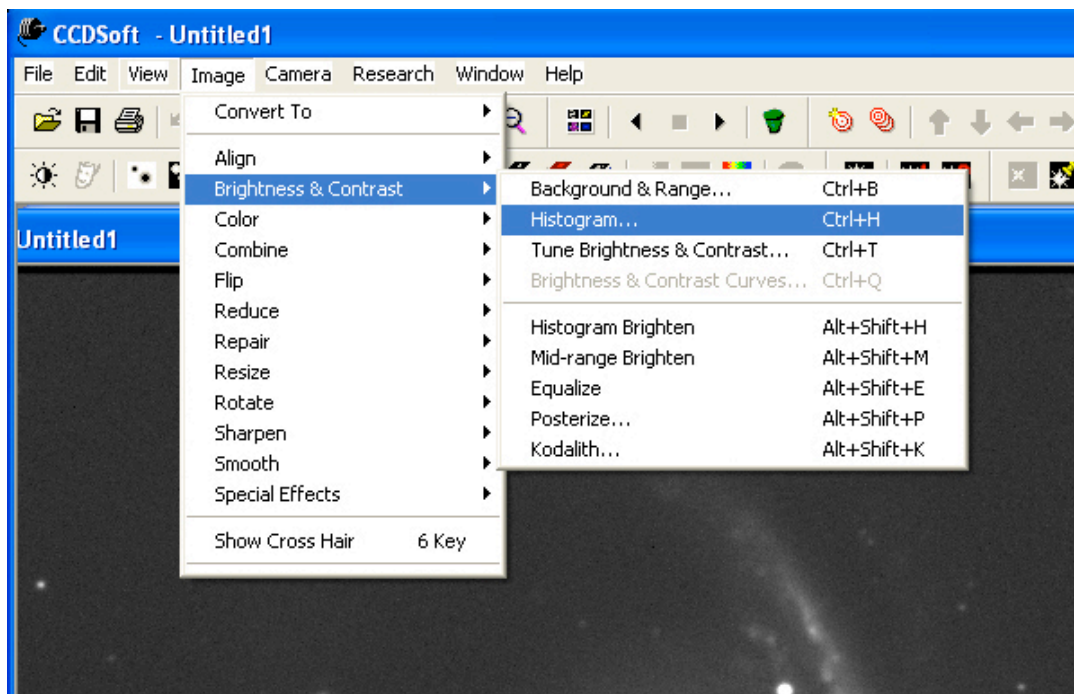
Once you have this done, click “OK” and see the result! Now the image looks much better. More stars and more of the galaxy are visible, and almost all of the noise is gone.

For many purposes, this is all that is required in calibrating the image, particularly if the image is to be just a ‘pretty picture’. However if the image is for scientific purposes, then there is a second step required. This second step is to remove the effects of non-uniform illumination of the CCD chip, and for the different sensitivity to light of each pixel. This step is known as **Flat Fielding**. This process will be covered shortly.


Image Enhancement: CCDSoft

Once the image is calibrated, it is ready for processing. Processing the image refers to enhancing the image so more detail can be seen. CCDSoft has some basic methods of doing this. They can simply adjust what the computer displays. This process does not change the actual pixel values. Or they can change the actual pixel values. We will have a look at one method called “**Histogram Stretching**”.

Histogram: Although the pixel values can have any value between 0 and 65535, in practice most times only a very small portion of the available brightness levels contain the majority of the visual information. It is possible to adjust the images so that only the important range of brightness values is displayed. A very important tool for this is the histogram. This is a graph of the various brightness values versus the number of pixels with that brightness value. To see this, go to the “**Image**” menu and select “**Brightness & Contrast**” and then “**Histogram**”.



The histogram will then appear in a separate box that can be resized if it is too small. As can be seen from the display, the majority of the pixels have very low values. Below the graph are 2 'sliders'.

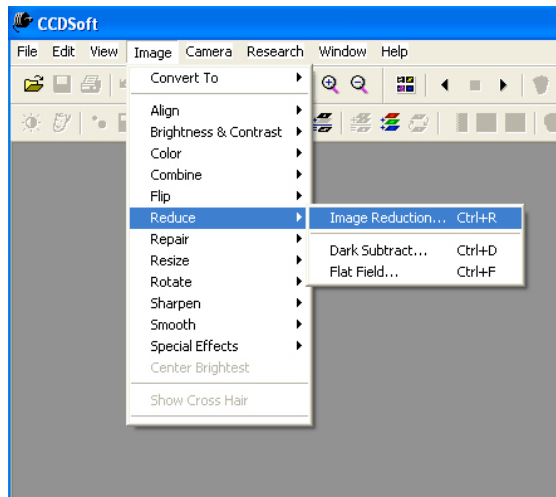


The one on the left sets the value that will be displayed as black, while the one on the right sets the value to be displayed as white. Values in between will be displayed as various shades of grey. Play around with these sliders and see what effects they have on the image. Any adjustments to the histogram do not change the actual pixels values, only the way they are displayed.

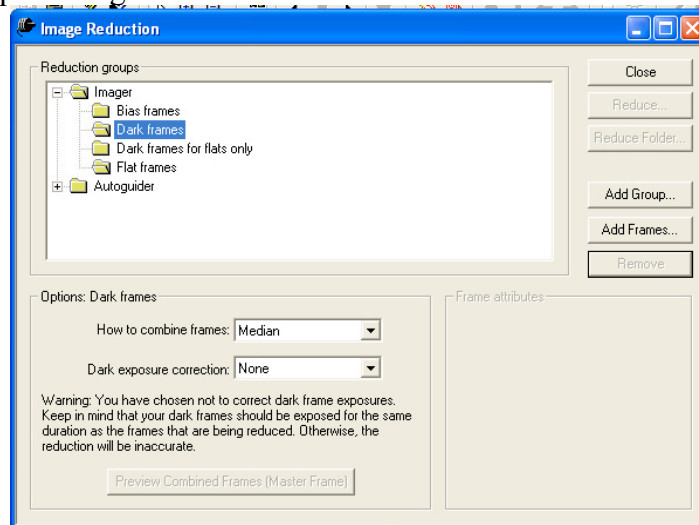
Using Multiple Images:

Just using one image really limits what can be done with that image. The main reason for this is the exposure time is limited. When imaging deep-sky objects, the longer the exposure, the more detail becomes visible, and the fainter the objects that can be imaged. However, taking one very long exposure is not very practical for many reasons, not the least of which is the gradual buildup of dark current and background glows from light pollution. Therefore, it is more usual to build up a long exposure by combining many short exposures. In this next part of the lab, you will learn how to do this in CCDSoft.

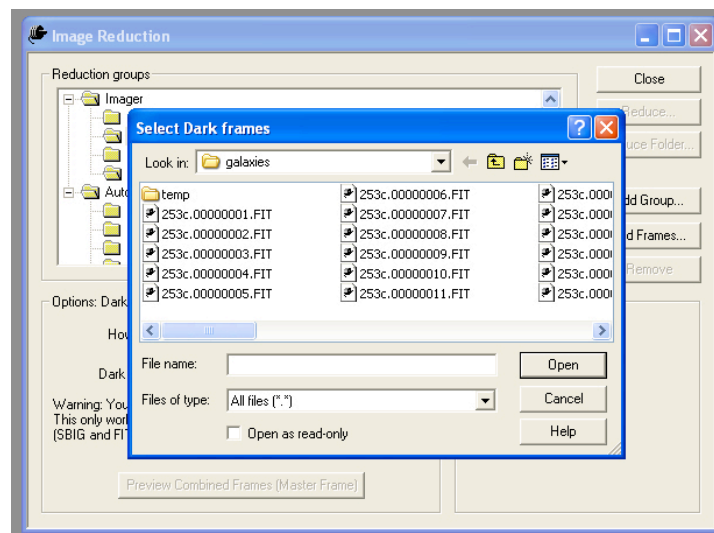
Close the images you were just using, (No need to save them). Then make yourself another folder. My convention is to label it "raw" so that I know this is where the unprocessed images are, but you can name it whatever you want, so long as you can remember which images it contains. Move the 5 files "**1365R1.fts – 1365R5.fts**" and the file "**1365D1.fts**" over to this folder. When calibrating a number of images, it can be very time-consuming to dark subtract each one individually. Fortunately, CCDSoft has a routine that will do this for all the images in one go. Go to the "**Image**" menu and the "**Reduce**" and "**Image Reduction**".



This will bring up a dialog box. Click on “Dark Frames” and then “Add Frames”.



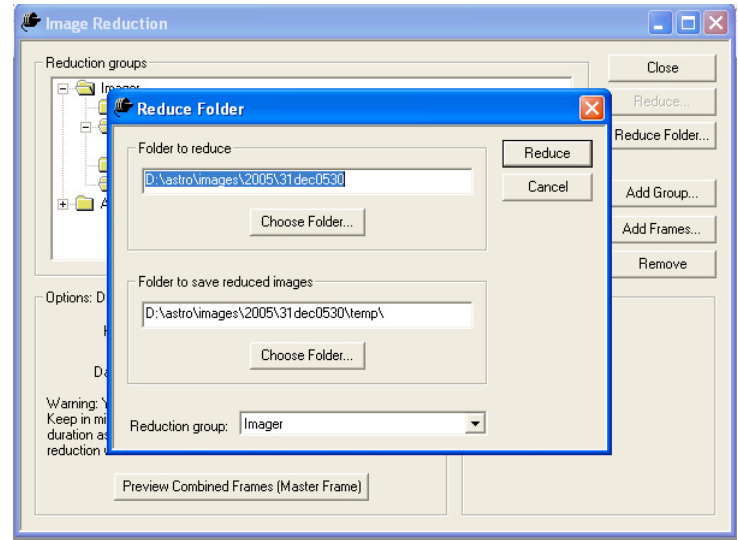
You will now have another box asking you which file to add as the dark frame.



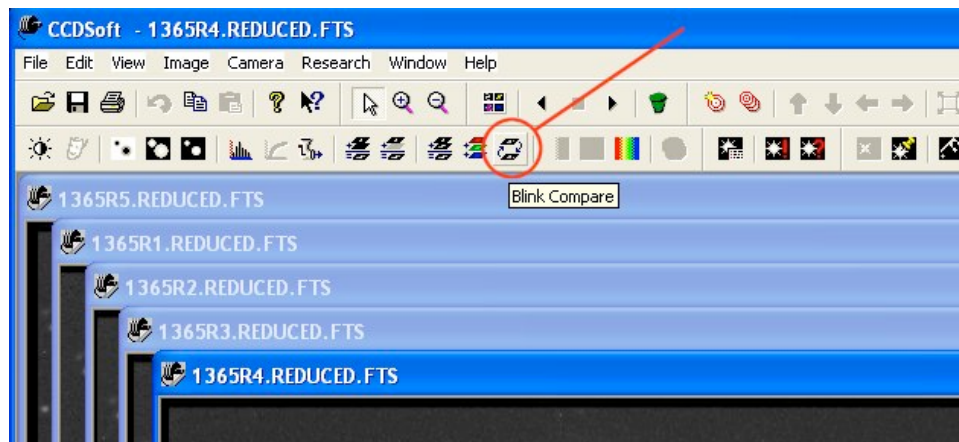
Click on the file “**1365D.fts**” and then click “**Open**”. Since we are only doing dark subtraction at this stage, this is all you need to worry about. So now click “**Reduce Folder**”. This will bring up another dialog box asking you to choose which folder contains the images to be reduced and which folder you want the reduced images to be saved into.

THESE MUST BE DIFFERENT FOLDERS.

Set the folder with the images and then create another folder to save the calibrated images in. My convention is to label this folder “calibrated” as I find this helps to know which folder is which. Once you have set the folders, click “**Reduce**”. The program will now reduce all the images in the chosen folder.



Now that we have dark subtracted all of the images, are we ready to combine them? Not quite! First, have a look at the “Calibrated” folder. In there you will see a file “**1365D1_reduced.fts**”. Since you told CCDSoft to reduce the entire folder, it also calibrated the dark frame! This file is not wanted, so delete it. Then there is another step required. To see why, open all of the reduced images in CCDSoft. You should now have 5 images of the galaxy open. Now look at the row of buttons just above the images. Almost in the middle is one that has two dark icons with arrows passing between them. This is the “Blink Compare” button. Click on this button and the program will start cycling through the images.



Watch a few cycles and you will notice that the stars in the images might seem to move. What has happened is that the telescope has moved slightly between images. This is quite common and if the images were combined, the stars would appear as trails. So before the images can be combined, they must be aligned so that the stars are all in the same positions. CCDSoft has a very simple routine for doing this.

Start by closing the images. (they should not be open when aligning) Now go to the “**Image**” menu and click on “**Align**” and then “**Align Folder of Images**”.

Now you will see a dialog box asking you which folder of images do you want aligned and which folder you want them saved into. **(They must be different)** My convention for labeling the folder in which the images will be saved is “aligned”.

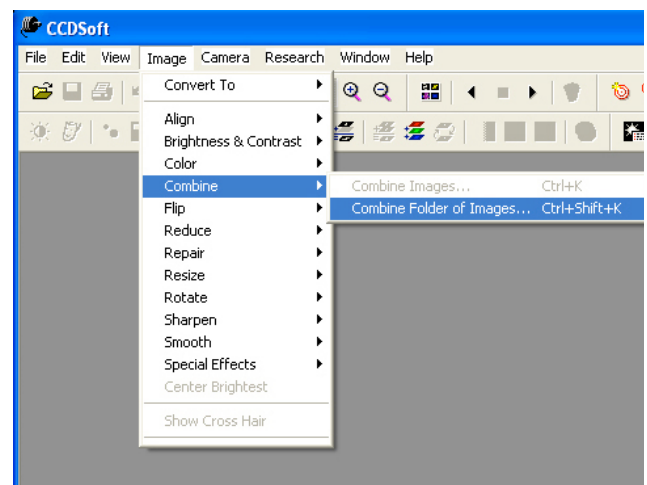

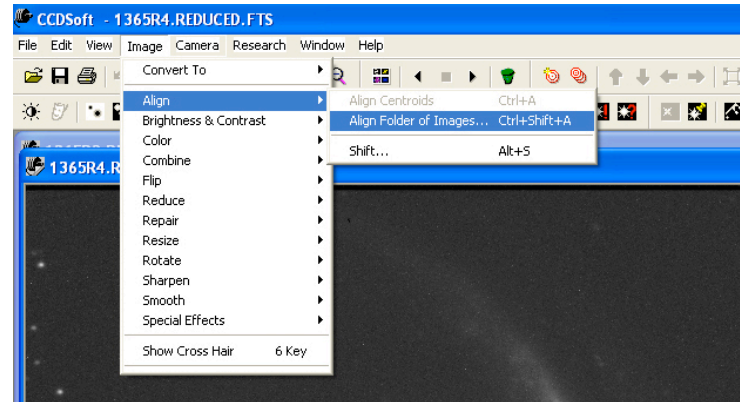
(NOTE: If you had imaged more than one object during the night, you would need to move the different calibrated images over to a temporary folder, and leave only images of the same object in the “calibrated” folder!)

Once you have set the folders, click “**Align**” and the software will align the images. Once you have aligned the images, it is a good idea to open them and blink them just to check that the alignment has been done correctly.

Once you are satisfied that they are aligned correctly, delete images in the “calibrated” folder since they are no longer needed. If you had imaged more than one object during the night you would now repeat the alignment for each different object.

With the images in a folder on their own, it is time to combine them. There are 3 ways of doing this with CCDSoft. To see these three ways, go to the “**Image**” menu and then click on “**Combine**” and “**Combine Folder of Images**”.

This will bring up a dialog box asking 2 things. Firstly, which folder of images do you want aligned? Secondly, how do you want the images combined? Set the folder and then click on the arrow in the “**How to combine**” window. The software will now display the three ways of combining the images.



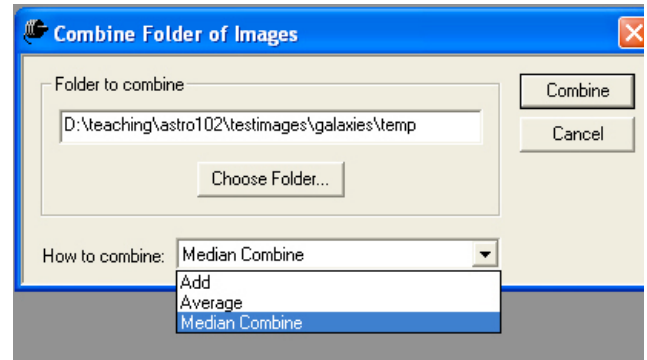
“**Add**” as the name implies, simply takes the numerical value of a particular pixel in each image and then adds them together. It does this for each pixel in the images.

“**Average**” likewise takes the numerical value of a particular pixel in each image and then adds them together, but then divides that number by the number of images to get the average.

“**Median**” finds the middle value for a particular pixel between the images.



Each of these three methods has advantages, and the one you choose will depend on the images you are trying to combine. “Add” is very useful when the object you are trying to image is very faint. The disadvantage of this method is that once a pixel reaches a value of 65535, it cannot go any higher. The technical term for this is that it is “saturated”. “Median” has the advantage of removing many remaining hot or cold pixels, while “Average” is somewhere in between. Very often I do both “Add” and “Median” and see which looks better.

Go ahead and combine the images using “Add,” “Average,” and “Median” and compare the resulting images. Save the resulting images to your portable flash drive to turn in with you lab notebook.




Flatfielding

Often after doing the dark subtraction, there are still problems with the final image. Very frequently, some parts of the image will be brighter than others; there might be dark “donuts” of varying sizes on the image, etc. All of these problems can be traced to one source:- the CCD chip is not being evenly illuminated. Below are two examples of this. The image on the left suffers from uneven illumination, while the image on the right has numerous dark “donuts”.



These effects of non-uniform illumination of the CCD chip, and different sensitivity to light of each pixel can be removed by a process known as **Flat Fielding**. The name derives from the technique of evening or “flattening” the chip illumination. What is required is a series of images of a perfectly even light source. With a perfectly even light source, any uneven illumination in the resulting images will be caused by the imaging system itself. Dividing your images of the sky by the image of the light source (called a “flat field”) will remove most, if not all of the uneven illumination defects in your images.

There are numerous methods for obtaining a flat field. The one we use in this course is to take a number of images of the twilight sky, and then combine them using a median filter. We can use images of the twilight sky, because the background sky is a uniform source of light. (at least, it is on small scales such as a CCD image!) However, it is important to make sure that there are no stars visible in the images. To ensure this, you take 10 – 20 images (the more the better) of the bright twilight sky where there are no bright stars. The exposure should be no more than 0.5 seconds (the shorter the better) but long enough to give a maximum count reading near the centre of between 15,000 and 30,000.



Final flatfield. A median combine of 8 separate images. Notice all that is visible is the background sky. All of the unevenness in the image is a result of the optical system. Dividing your images by this flatfield removes these defects as shown in the image below.

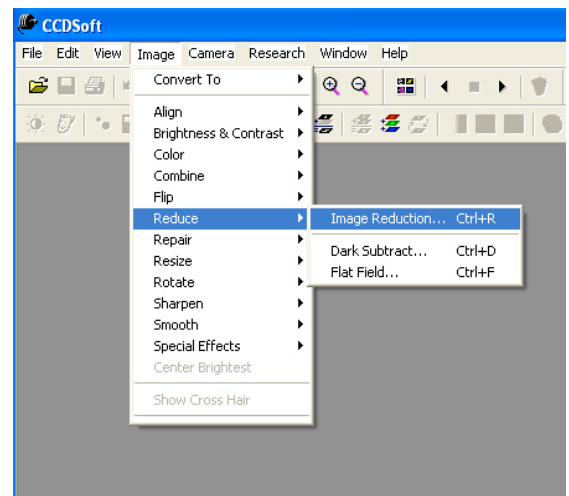


One of the above images after flatfielding. Notice how much more of the galaxy is visible and how much flatter the background is. It is still not perfect, but that is due to other problems.

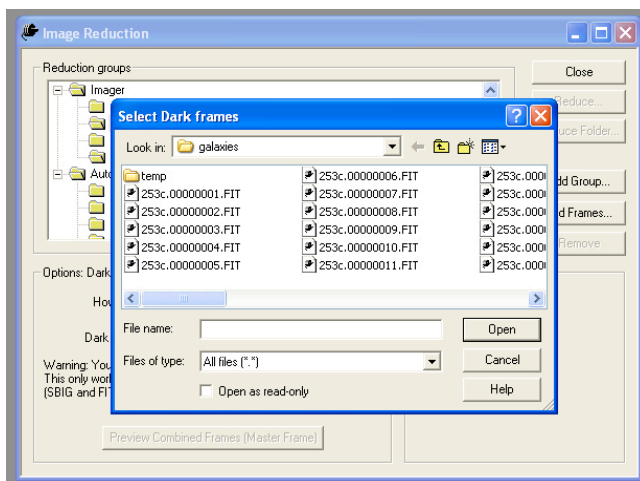
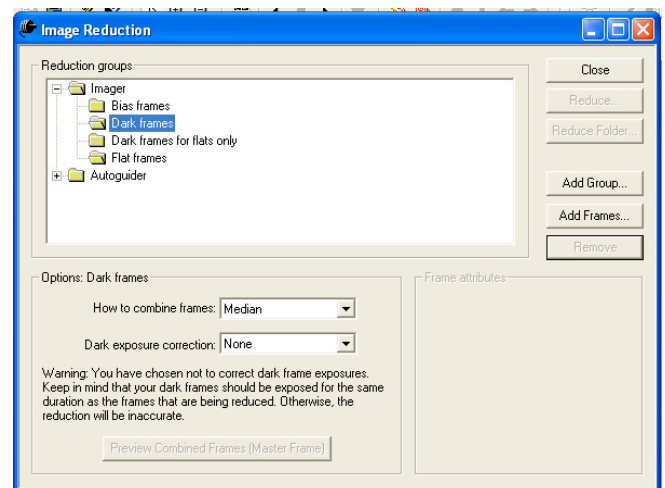
Full Image Calibration

Return to your main folder make three sub-folders. Call one folder “**calibrated**” the second “**aligned**” and the third “**combined**”. From the main “**galaxies**” folder, copy to your main folder the images starting with “**300**”. These are images of the spiral galaxy NGC 300. Also copy over the accompanying dark frames and flat field images. Notice that there are two sets of darks; one for the object images and another for the flat field images. This is because every image you take with the CCD will have dark current in the image; you must first subtract this from your images. **This includes any flatfield images you take!** Therefore you must take dark frames for your flat field images as well.

Open CCDSoft and then create your master dark by going to the “**Image**” menu and then “**Reduce**” and “**Image Reduction**”.



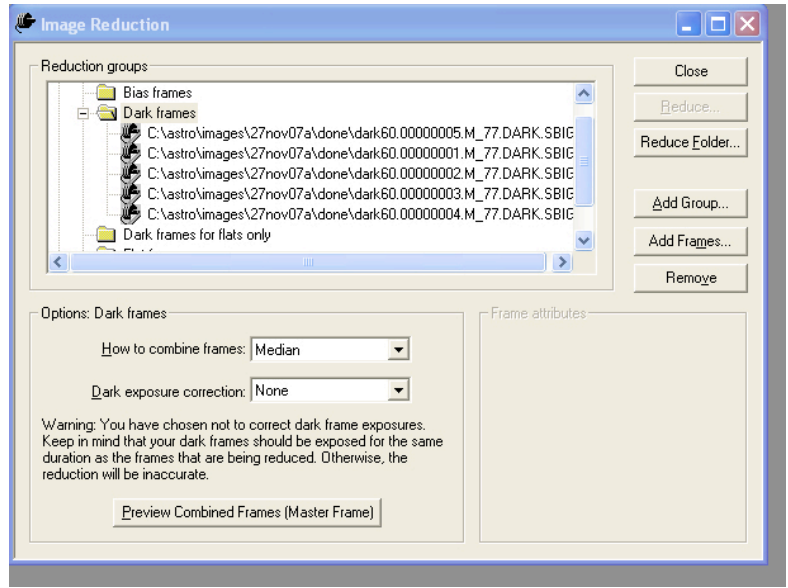
This will bring up a dialog box. Click on “**Dark Frames**” and then “**Add Frames**”.



You will now have another box asking you which files to add as the dark frame. **Select the dark frames for your object images**, and then click “**open**”.

In the windows below the main window, for the line “**How to combine frames**”, select “**Median**”, and for the line which says “**Dark exposure correction**” select “**None**”. Now click the button “**Preview combined frames**”. The program will now display a median combined image of the dark frames you selected.

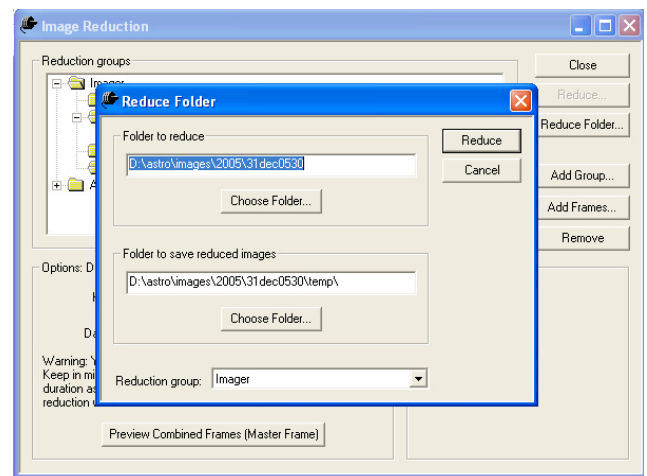
Save this image as “**MasterDark**” in your folder of images.



Now click on “**Dark frames for flats only**” and then “**Add Frames**”. This time when the box appears, **select the dark frames for your flatfield images**. Once again, for the line “**How to combine frames**”, select “**Median**”, and for the line which says “**Dark exposure correction**” select “**None**”.

Now click on “**Flat Frames**” and then “**Add Frames**”, and then select your **flatfield images**. Set the line “**How to combine frames**”, to “**Median**”, and then click the “**Preview combined frames**” button to display a median combined image of the dark flat field images you selected. Check this image carefully to make sure there are no stars left in the image. The flat frame has been normalized, so you may need to adjust contrast to see anything. If there are stars left, something went wrong when you took your flat field images and you will need to call for me to have a look. If there are no stars in the image, save it as “**MasterFlat**” in your folder of images. CCDsoft automatically applied the flat darks to the master flat, so you are good to go. Now remove all of the images in the “Dark Frames”, “Dark Frames for Flats”, and “Flat Fields”, using the “Remove” button. Then in the “Dark Frames” add the “Masterdark” that you made, and in the “Flat Fields” add the “Masterflat”.

So now click “**Reduce Folder**”. This will bring up another dialog box asking you to choose which folder contains the images to be reduced and which folder you want the reduced images to be saved into. **THESE MUST BE DIFFERENT FOLDERS**. Set the folder with the images and then save the calibrated images in the “calibrated” folder you created earlier. This is my convention as I find this helps to know which folder is which; however, you could have used any folder name. Once you have set the folders, click “**Reduce**”. The program will now



reduce all the images in the chosen folder. Pro tip: If you type in a folder name in the “Folder to save reduced images” dialog box that doesn’t exist, CCDsoft will ask to create it for you.

Once this is completed, open the calibrated images to make sure they are properly calibrated. Then align and combine the images for each object as you learned earlier. Combine with both “Add” and Median”. Then save these combined images to your flash drive to turn in with your lab manual. Be sure to include the filenames and folder locations in your lab write up. Alternatively, you may print out your images and paste them into your lab notebook.