

The Amateur Astrophotographer

Robert J. Vanderbei

PRINCETON UNIVERSITY, PRINCETON, NJ 08544
E-mail address: rvdb@princeton.edu

THE AMATEUR ASTROPHOTOGRAPHER

Copyright ©2003 by Robert J. Vanderbei. All rights reserved. Printed in the United States of America. Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a data base or retrieval system, without the prior written permission of the publisher.

ISBN 0-0000-0000-0

The text for this book was formatted in Times-Roman using $\mathcal{A}\mathcal{M}\mathcal{S}$ - $\mathcal{L}\mathcal{T}\mathcal{E}\mathcal{X}$ (which is a macro package for Leslie Lamport's $\mathcal{L}\mathcal{T}\mathcal{E}\mathcal{X}$, which itself is a macro package for Donald Knuth's $\mathcal{T}\mathcal{E}\mathcal{X}$ text formatting system) and converted to pdf format using $\mathcal{P}\mathcal{D}\mathcal{F}\mathcal{L}\mathcal{A}\mathcal{T}\mathcal{E}\mathcal{X}$. All figures were incorporated into the text with the macro package $\mathcal{G}\mathcal{R}\mathcal{A}\mathcal{P}\mathcal{H}\mathcal{I}\mathcal{C}\mathcal{X}.\mathcal{T}\mathcal{E}\mathcal{X}$.

*Many thanks to Krisadee, Marisa,
and Diana for putting up with many
weary days following long sleepless
nights.*

Contents

Preface	ix
Part 1. Equipment and Techniques	1
Chapter 1. Introduction	3
1. Location, Location, Location	3
2. Is it Worth the Effort?	4
Chapter 2. Equipment	9
1. Choosing a Telescope	9
2. Choosing a Camera	14
Chapter 3. Image Acquisition	21
1. Polar Alignment	22
2. Attachments	24
3. Focusing	25
4. Finding Faint Fuzzies	26
Chapter 4. Image Processing	29
1. Raw Image Frames	29
2. Dark Frames	32
3. Bias and Flat Frames	33
4. Image Calibration	34
5. Stacking	37
6. Analyzing the Image	37
Chapter 5. Image Enhancement and Presentation	43
1. Unsharp Mask	43
2. Gamma Stretching	47
3. Digital Development	48
4. Variants of Digital Development	53
5. Deconvolution	54
6. Correcting for Coma/Field-Curvature	56
Chapter 6. Stories	61
1. Jupiter	61

2. Starquest 2002	64
Bibliography	67

Preface

Acknowledgements. Many thanks to Al Kelly who carefully read the manuscript and made several suggestions that enhanced the final result.

Robert J. Vanderbei

January, 2003

Part 1

Equipment and Techniques

CHAPTER 1

Introduction

Over the past ten to fifteen years, there has been and still is a revolution taking place in imaging technology resulting from advances in computer hardware, image processing software, and most importantly digital imaging devices called *charge coupled devices* or *CCD's*. Everyone has seen the new breed of cameras called digital cameras. At the heart of a digital camera is a CCD chip (or, recently, a CMOS chip which is a different acronym but a similar type of device). These digital cameras have been used for astrophotography of fairly bright objects. For faint astronomical targets, however, these cameras are currently limited by the build up of thermal noise which begins to be significant after only a few seconds of exposure time. This is why these cameras are generally limited to maximum exposures of only a few seconds. But, the good news is that this thermal noise is proportional to the temperature of the chip. For this reason, there has been a parallel development of new imaging cameras for astronomical use in which the CCD chip is cooled to as much as 40° Celsius below the ambient temperature. At such cold temperatures, modern CCD chips make excellent, extremely sensitive, virtually linear, astronomical imaging devices. Furthermore, these devices are very small, smaller than conventional film, and therefore make possible a trend toward miniaturization of photography. The purpose of this book is to illustrate the amazing things that can be done with a small telescope, a personal computer, and one of these new astronomical CCD cameras.

In astronomical applications, conventional wisdom is that bigger is better. This is true, but with the advent of CCD imaging, the definition of big got a lot smaller. Today, one can take CCD pictures with a small telescope and get quality that rivals the best one could do with film technology and a much bigger telescope. I hope that this book convinces the reader that these recent developments are little less than amazing.

1. Location, Location, Location

In addition to miniaturization, CCD imaging also helps address another ever growing problem—and that is light-pollution. We'd all like



FIGURE 1.1. Light Pollution Map of the Northeast USA

to live in a great dark-sky site such as one finds in remote parts of Arizona, New Mexico, and Hawaii. But most of us aren't so lucky. We live where the jobs are—in populous areas. My home's location is shown on the light-pollution map in Figure 1.1. As New Jersey goes, my location isn't bad. But, it doesn't compare to a truly dark site. For example, the Milky Way is just barely visible when it is high in the sky late on a moonless night. But, light pollution is getting worse every year. It won't be long and I will not be able to see the Milky Way at all. With digital imaging, the images are saved on a computer and it is utterly trivial to subtract the average level of background light pollution from the image. Of course, one can scan film-based pictures into a computer and apply the same techniques but with CCD imaging the picture is already on the computer and the subtraction is trivial so one can consider it essentially automatic. With film, one needs to make a real effort to accomplish the same thing.

2. Is it Worth the Effort?

Taking pictures, whether with film or with a CCD camera, takes much more effort than one needs for simple visual observing. So, why bother? Why not just enjoy the visual observing? On this question,




FIGURE 1.2. *Left.* Approximately how M13 looks visually in a small telescope under light-polluted skies. *Right.* My first CCD image was of M13.

people hold differing views. Many prefer the visceral, dynamic feel of real-time visual observing. Others, like myself, find imaging to be fun and rewarding. The following little story gives the main reason I got into CCD imaging. I bought my first telescope, a small 3.5" Questar, a few years ago. I was, and am, very happy with the design, fit-and-finish, and optical quality of this instrument. As it was my very first telescope, I didn't know much about what to expect. Having seen lots of beautiful images on the internet and in magazines, I expected to see some of those things myself. Indeed, Jupiter, Saturn, and the Moon are certainly spectacular as are open clusters and double stars. But I was hoping also to have great views of various nebulae and even some galaxies—in this regard I was utterly disappointed. Visually, I can just barely make out that M13, the Great Globular Cluster in Hercules, is a globular cluster and M27, the Dumbbell Nebula, is a nebula. These are two very different kinds of objects but they both just looked like faint fuzzies in the eyepiece. True, M13 seemed to scintillate suggesting that it is a star cluster, but it was more of a vague impression than a clear observation. Figure 1.2 shows a rough approximation of how M13 looks to me visually in the Questar. I've purposely made the image appear rather faint to simulate the effect that light pollution has on one's eyes' dark adaptation. And, galaxies being even fainter are harder to see. From my home, I am able to see four galaxies: M31, M81, M82, and M104. That's it. I've never seen M51 or M33. And, when I say that I can see a galaxy one should be aware that what I see is just a faint glow. I do not see any dust lanes in M31, nor do I see any structure in the other three except for their overall shape. Based on the pictures I'd seen in magazines it was clear that one could expect to do a lot better by

taking pictures. So, I bought an astronomical CCD camera. The image shown on the right in Figure 1.2 is the very first image I took with it. I'd say that the difference is dramatic.

Another objection often heard is this: amateur images can't compete with those taken by Hubble, so again why bother? Of course, one can ask the same thing of the visual observer: why spend lots of time and money to look at things when instead you could download Hubble pictures of the same thing that are dramatically better than what you can hope to see. One answer is simply that it is very rewarding to do it yourself. However, there is more. Looking at a Hubble picture, one doesn't get a sense of the scale of the object, nor of how faint the object is. Also, the color in Hubble images is usually false color. These are things for which first-hand experience is the best teacher.

2.1. When to Image? Given my initial success with the first image of M13 it became tempting to take images on every clear night. The image on the left in Figure 1.3 is globular cluster M5 taken a few weeks after that first image of M13. This was taken at a time when the moon was nearly full and was not very far from M5. For many months I was quite satisfied with this image thinking that the only reason that it wasn't as good as my M13 was simply the fact that M13 is a brighter cluster. But then I noticed that others had taken pictures of M5 that were much nicer than mine. So, about 9 months later I re-imaged M5 late at night when the moon was not out and M5 was high in the sky. This new image is shown on the right in Figure 1.3. Again the improvement is obvious. There are two lessons here: (i) it is possible to image objects even when the moon is bright and nearby, but (ii) it is still better to take pictures under the best possible conditions.

Figure 1.4 shows a comparison among four different images of M5. The one on the top left is a 5 minute exposure taken with my 3.5" Ques-tar and my MX-916 camera. The one at the top right is a 7.8 second exposure from the 2 Micron All Sky Survey (2MASS). The images in the 2MASS survey come from the 1.3m telescope at Mt. Hopkins with a custom designed tri-color infrared CCD camera. The image on the bottom left is from Bruce Dale's *Learn What's Up* website (exposure details unknown) and the image at the bottom right is an 8.5 minute exposure taken by Walter MacDonald using a 10" Meade LX200 (at $f/6.3$) and an MX-916 camera. It's interesting to see how similar they are.




FIGURE 1.3. *Top Row:* Raw images of globular cluster M5. *Bottom Row:* Same images after processing. *Left Column:* With a full moon nearby. *Right Column:* Late on a dark moonless night.



FIGURE 1.4. *Top Left.* A 5 minute exposure, 3.5" Questar, MX-916 camera. *Top Right.* A 7.8 second exposure, 1.3m Mt. Hopkins, infrared CCD camera. (Credit: Image obtained as part of the Two Micron All Sky Survey (2MASS), a joint project of the Univ. of Mass. and the Infrared Processing and Analysis Center/California Institute of Technology, funded by the NASA and the NSF.) *Bottom Left.* Unknown exposure from Bruce Dale's *Learn What's Up*. *Bottom Right.* An 8.5 minute exposure, 10" LX200, MX-916 camera.

CHAPTER 2

Equipment

1. Choosing a Telescope

This book is about CCD imaging with a small telescope. There are many excellent telescopes available today. Even inexpensive ones are generally very good—at least for visual use. However, for astrophotography there are certain issues and factors that must be considered and that rule out many of the most popular and most inexpensive telescopes. Nonetheless, there are many choices that will provide both excellent visual enjoyment and superb imaging capabilities. In this section, I'll articulate the main issues to pay attention to.

Rigidity. With astrophotography, one often adds quite a bit of weight to the optical system. In addition to the CCD camera, there might be a filter-holder, a focal reducer, and various extension tubes. All together this can result in a few pounds of stuff hanging off the back of the telescope. It is important that this weight not create warpage in the optical tube assembly (OTA). Although I haven't tried astrophotography using a telescope with a plastic OTA, I strongly suspect that such a telescope will not be sufficiently solid. It is also important that the mount can handle this extra weight.

Equatorial Mount. Telescopes with so-called GoTo mounts have become extremely popular in recent years, especially with inexpensive telescopes. It is amazing how well they work for finding objects and then keeping them relatively well centered in the eyepiece over a long period of observation. The inexpensive GoTo telescopes are generally designed to be used in an Altitude/Azimuth (Alt-Az) setup. Afterall, such a setup is less expensive as one doesn't need an adjustable wedge to achieve an accurately polar aligned equatorially mounted system. And, the built-in guide computer can handle the tracking.

It turns out that, while these GoTo systems are great for visual observing, they generally aren't accurate enough to take even 10 second exposures without stars trailing. Suppose for the sake of argument that you get a system that is sufficiently accurate to take 10, maybe even 30, second exposures. You might be thinking that you will stack these

exposures in the computer to get the effect of one long exposure. This can be done, but it is important to bear in mind that, over a period of time, the field of view rotates in an Alt-Az setup. So, one is faced with either getting a field derotator, which is itself rather complicated, or living with final images that must be cropped quite severely to eliminate the ragged borders created by a stack of rotated images. In the end, it is just simpler to use an equatorial mount.

Dual Axis Drives. An equatorial mount with a motor driven right-ascension axis can be used to take short to medium length exposures, which can later be digitally stacked to achieve long total exposure images. The only problem here is that any small error in the polar alignment will eventually show itself in a drift either northward or southward of the field-of-view. Hence, after stacking the images you'll have to crop them along the northern and southern edges to get just the portion of the field of view that is present in all of the images. In the end, you will be happier with a dual axis drive system that has the capability of being controlled by an external computer. Such a computer controllable dual-axis system can be used to produce long, guided exposures that need no cropping. If you settle for single-axis drives, you will quickly wish you had a dual-axis system.

Also, it is important that the right-ascension drive be precise. The more precise it is, the longer unguided exposure one can take. Suppose you were to set up your telescope/mount with a perfect polar alignment aim at some target and let it run for 24 hours. After 24 hours your telescope would almost certainly still have the same object in the center of the field. If it didn't, your drives would have *systematic error*. This is very uncommon—essentially it is just a clock and even the cheapest clocks and watches are exceedingly accurate. However, if you observe the object in a high power eyepiece with an illuminated reticle, you will see that the tracking isn't perfect. First, it will track a little too fast for a few minutes and then too slowly. On average, the tracking is correct but there is this *periodic error* which will ruin a long exposure photograph. The better the mount, the less the periodic error. Many of the high-end mounts even can be trained to correct automatically for their own periodic error. With such mounts, it is possible to do several minute long unguided exposures (again, assuming a perfect equatorial alignment). At a minimum one would like to be able to do 20 to 30 second unguided exposures. How precise the drive needs to be to achieve this depends on the focal length of the telescope and the size of the pixels in the camera.

Quality Optics. Most flawed images are flawed because of tracking errors. But, assuming tracking is precise then quality of the optics becomes the next most critical element. The better the optics, the better the final image. It seems that mirror and lens manufacturing has progressed to the point that even inexpensive optics are really very good. Maksutov Cassegrain telescopes (MCTs) and refractors are “permanently” collimated. If manufactured correctly, they generally provide superb performance with little trouble. Classical Newtonian reflectors and Schmidt Cassegrains have adjustable collimation. When the collimation is bad, the telescope produces bad images. This is a very common problem. Usually the basic components, the primary and secondary mirrors, are of high quality but bad collimation will ruin a final image. For this reason, MCTs and refractors are less trouble. However, MCTs are usually rather slow instruments ($f/15$ is typical) and so require much longer exposures. Modern refractors, on the other hand, tend to be rather fast instruments ($f/5$ to $f/8$) and have very little chromatic aberration. They are very popular for astroimaging.

Flip Mirror System. A flip mirror system provides the capability to view the image through an eyepiece without removing the camera. With such a system, you just flip a lever to move the mirror in place and it diverts the light bundle away from the camera and toward a separately attached eyepiece. While not required, these systems certainly make some imaging tasks much easier. For example, Jupiter and Saturn are usually imaged at very high magnification (using either a Barlow lens or eyepiece projection—see Chapter ??). At these magnifications, the field-of-view is very small. It is almost impossible to find Jupiter and Saturn just using a low-power finder scope, setting circles, and images from the camera. It is a great advantage to have a high-power eyepiece at your disposal. And you need to flip back and forth between the two because first you need to find a bright star to focus on and then, without changing anything that could disrupt the focus, you need to be able to find Jupiter. It can take a very long time and be very frustrating to try to do this without a flip mirror. With one, it is a rather trivial task.

Easily Adjusted Counterweights. Once you enter the world of astroimaging, it is inevitable that you will use several different imaging configurations. You will soon find yourself doing imaging with and without a focal reducer and with and without color filters. You also might do piggyback photography where a camera rides piggyback on your telescope. You also might end up owning and using more than one camera. It is important to be able to balance each of these configurations and to be able to switch from one to another without much hassle.

Most telescope mounts make this task rather straightforward but it is important to think about it before making a purchase.

Dew Shield. On most nights and in most places, a significant amount of dew forms by morning. For example, here in New Jersey, during most of the spring, summer, and fall the late-night temperature usually falls to just above the dew point. By morning, everything is soaking wet. Dew is caused by objects cooling to a temperature below the dew point and thereby causing moisture to condense out of the air that comes in contact with the cold surface. The cooling takes place by infrared radiation. Not only do objects emit infrared but they also absorb some of the infrared emitted by others. If the net absorption matches the net emission then no overall cooling takes place. The problem is that objects that face the sky, such as the front surface of a telescope, tend to receive less infrared radiation than they emit. Hence, they cool very fast—faster than the surrounding ambient air. Other things that are shielded from direct exposure to the openness of space tend to absorb about as much radiation as they emit. Objects in such places tend remain at the ambient temperature and don't get wet (at least not until the ambient temperature falls to the dew point and which time everything gets wet).

There are two standard techniques for preventing dew from forming on the front optical surface: (a) dew shields and (b) battery operated dew zappers.

A dew shield is a tube that fits over the OTA and extends beyond the front of the telescope. It absorbs infrared radiation from surrounding objects and re-emits it in all directions. Some of that light is then absorbed by the front surface of the telescope. If the dew shield is made out of a material that absorbs infrared radiation, then it will be very effective at preventing dew from forming.

A dew zapper is a heated strip of cloth that wraps around the optical tube close to the front optical surface. It works by heating the lens just enough to keep it from forming dew. While this method works very well and is very popular, there are some drawbacks to this method. First, the heat is supplied by an electric current usually provided by a battery. If the battery goes dead, then no more zapping. Secondly, a dew zapper adds a few more wires to worry about tripping over. Finally, if too much heat is applied, then the zapper could set up air currents in front of the telescope which will quickly deteriorate an image.

In addition to the front optical surface, it is important that the rest of the equipment also works well even when wet. This may seem trivial, but just because friction based connections work reliably and smoothly

when dry and at room temperature, it doesn't mean they will work when wet and cold.

The 3.5" Questar. As already mentioned, the only telescope I own is a 3.5" Questar. I bought this telescope before I had any thoughts of doing imaging—it was to be for visual use. Like so many others starting out in amateur astronomy, I bought my first (and only) telescope with too little research. I had a friend who had, and still has, a 7" Questar. I'd looked at some of the common favorite objects through his telescope and was truly excited by what I saw. The 7" Questar was beyond the price I was willing to pay for a first telescope so I began to think about alternatives. The most obvious, and certainly the most popular, alternatives would be an 8" Schmidt-Cassegrain from Meade or Celestron. However, in my limited prepurchase investigations, the most common warning I heard was that many naive first-time buyers end up buying a telescope that is too large and too sophisticated and that they then lose interest after using it only a few times because it is just too much work. The advice is that the telescope to buy is the one you will use and the one you will use is often the one that is simplest to transport and set up. So, I considered buying a 3.5" Meade ETX-90. In fact, I almost did buy one. I made this decision during the 1998 Christmas shopping season. I decided to wait until after Christmas in the hopes of getting a sale price. In fact, I went to the local Discovery Channel store on the day after Christmas and told the sales person that I wanted to buy an ETX-90. He informed me that they didn't have any left—they'd sold them all. Bummer. Sometime a few days later I was at a party where I happened to be talking to someone about my interest in buying a telescope. I had no idea that this person was interested in astronomy but it turned out that he owns a 3.5" Questar and his advice was that I'd quickly become disappointed in the ETX and that I should buy a Questar. The idea began to make sense. Its price is similar to a 10" Schmidt-Cassegrain, which was for me a reasonable target, and yet it is small and easy to use so likely to be used. His advice and these thoughts convinced me. A few days later I ordered one and received it about a month later.

So that's the story—not a very scientific approach to the decision. Fortunately, the decision wasn't a bad one. But, as I mentioned earlier the views of deep sky objects from light-polluted New Jersey are somewhat disappointing and for this reason I began to think about CCD imaging. If I were to pick a telescope from scratch with CCD imaging in mind, I would probably pick a fast refractor designed specifically for photographic use. And my next telescope, if there is a next one, will likely be such a refractor. For now the Questar is my only telescope.



FIGURE 2.1. The 3.5" Questar telescope.

2. Choosing a Camera

There are several criteria affecting the choice of camera. When thinking about a small telescope, the most obvious issue is the size and weight of the camera. Santa Barbara Instruments Group (SBIG) makes very fine astronomical CCD cameras but they are too big and too heavy (tipping the scales at more than a kilogram) for use on a small telescope like the 3.5" Questar. The same is true for many other manufacturers. One camera maker, Starlight Express, makes very small lightweight cameras using (wouldn't you know it) Sony CCD chips. Their cameras are small partly just because that is a design objective for them but also because they have chosen to keep the workings simple. Most SBIG cameras have a built-in physical shutter, color filter wheel, off-axis smaller secondary CCD chip for guiding, and a regulated Peltier cooling system with electronics to monitor the temperature and a fan to help with the regulation. Starlight Express cameras, on the other hand, do not have shutters, built-in filter wheels or off-axis guiding chips (more on their unique approach to guiding later). Furthermore, while Starlight Express cameras do have Peltier cooling, they do not regulate the cooling and there is no cooling fan. All of this makes for

a very simple lightweight design. These cameras only weigh about 250 grams. Later we will discuss how one overcomes the apparent loss of functionality caused by these design simplifications.

Three Fundamental Facts. There are several camera models sold by Starlight Express and in the future there will likely be other manufacturers of light-weight cameras. In order to decide which model is best for a given telescope, it is important to consider three fundamental facts.

First Fact. A telescope concentrates the light entering it as an image in the *focal plane*. The actual physical scale of this image is proportional to the *focal length* of the telescope; that is, if one doubles the focal length then the image will be twice as big (in both directions and hence have 4 times the area). Now, put a CCD chip in the focal plane. The *field of view* is defined as how much of the sky fits onto the chip. It is measured as a pair of angles (one for each direction—horizontal and vertical). If the image scale is large, then the field of view is small. In fact, the angular field of view is inversely proportional to focal length. Of course, if one uses a different camera with a larger CCD chip, then the field of view increases in accordance with the increased size of the chip. We can summarize these comments by saying that the field of view is proportional to the dimension of the CCD chip and inversely proportional to the focal length of the instrument:

$$\text{Angular field of view} \propto \frac{\text{Dim of CCD chip}}{\text{Focal length}}$$

For a specific reference point, the Starlight Express MX-916 camera has a chip whose physical dimensions are: 8.5mm horizontally and 6.5mm vertically. This chip is large by current CCD standards. But, its diagonal dimension is 10.7mm, which is much smaller than standard 35mm film. Using a focal length of 1350mm for the Questar telescope, we get that the MX-916 camera has a 22×16.5 arcminute field-of-view. Many objects are too large to fit in this field. For such objects a focal reducer is required. We will discuss focal reducers in Chapter ??.

Second Fact. There are two fundamental physical dimensions that characterize a telescope: it's *aperture* and it's *focal length*. Aperture refers to the diameter of the primary mirror or lens whereas focal length refers to the distance from the primary mirror/lens to the *focal plane*, which is where the CCD chip goes. As we saw above, knowing the focal length allows one to determine the field of view. It is a common misconception that the aperture determines the brightness of faint objects at the eyepiece or, equivalently, the exposure time when imaging.

Actually, the brightness/exposure-time depend on the ratio between focal length and aperture. Because of this important fact, this ratio is give a name, *focal ratio* or *f-ratio* for short:

$$\text{Focal ratio} = \frac{\text{Focal length}}{\text{Aperture}}.$$

Exposure time is proportional to the square of the focal ratio.

To see why, it is best to break the statement into its two fundamental factors: exposure time is proportional to (a) the square of the focal length and (b) one over the square of the aperture. First let's explain (a). Consider any telescope. It's easiest to think about the telescope that one already owns. So, I'm thinking about a telescope with a 3.5" (89mm) aperture and a 1350mm focal length. Now, suppose I could find a telescope with the same aperture and twice the same focal length. According to our first fundamental fact, such a telescope would provide one-half the angular field of view. Hence, the area of the sky being imaged is one fourth of what it was before. Since the aperture hasn't changed, the same number of photons are being collected per unit area of sky. So the central one-fourth of the image in the original telescope is now being spread over the whole chip. This means each pixel receives photons at one quarter of the rate as before. Hence, the exposure time must be increased by a factor of four.

Now let's explain (b). Suppose I could find a telescope with twice the aperture and the same focal length as my base telescope. According to our first fundamental fact, such a telescope would provide the same field of view. Hence, each pixel of a CCD chip gets light from the same small patch of sky. But the aperture has four times the area and therefore collects four times as many photons per unit of time. This means that the exposure time can be one fourth as long.

These arguments can be summarized as follows. When one thinks about getting a telescope with a larger aperture, one usually thinks about getting a scaled up version of their current telescope. A telescope that has twice the aperture and twice the focal length will collect four times as much light per unit area of sky but will also collect light from a patch of sky having only one fourth the area as before. Hence, the total amount of light collected is the same. The only real difference is the larger magnification the results from the smaller field of view.

Another consideration is the pixel size of the CCD chip. If everything else is kept the same, a CCD chip with large pixels will end up accumulating photons in each pixel faster than one with small pixels.

Again, it is the area that is the relevant measure and we can say that the exposure time is inversely proportional to the area of a pixel.

These considerations can be summarized as the following fundamental fact about exposure time:

$$\text{Exposure time} \propto \frac{(\text{Focal ratio})^2}{\text{Area of pixel}}.$$

Third Fact. Because of the wave nature of light, it is impossible for a telescope to concentrate a star's light at a single point in the image plane. Instead, what one sees is a small diffuse disk of light, called the *Airy disk*, surrounded by a series of progressively fainter rings called *diffraction rings*—see Figure 2.2. If the focal length is kept constant but the aperture of the telescope is doubled, then the physical distance, in microns, from the center of the Airy disk to the first diffraction ring is cut in half. On the other hand, if the aperture is kept constant and the focal length is doubled, then this distance to the first diffraction ring is doubled. Hence, if we double both the aperture and the focal length, then these two effects cancel out and the distance to the first diffraction ring remains unchanged. This means that the size of the Airy pattern is proportional to the focal ratio:

$$\text{Airy pattern size} \propto \text{Focal ratio}.$$

At $f/15$, the first diffraction ring is 10 microns from the center of the Airy disk.

MORE TEXT HERE

3.5" Questar at $f/16$ vs. 7" Telescope at $f/8$

- Same Field-of-View.
- Quadrupled exposure times.

Conclusion: Look for a lightweight camera with large pixels.

The Starlight Express MX-916.

- Approximate equivalent ASA (ISO): 20,000
- Linear response—no reciprocity failure (prior to saturation)
- Quantum efficiency: $\approx 50\%$




FIGURE 2.2. *Left.* A computer generated image of an Airy pattern—i.e., the Airy disk and diffraction rings. *Right.* An actual image of the Airy pattern taken with a PC-164C video camera attached via eyepiece projection to a 3.5” Questar. The effective focal ratio of the imaging configuration was about $f/45$.



FIGURE 2.3. The Starlight Express MX-916 camera with its 1.25” and 2” nose pieces.



FIGURE 2.4. My observatory.

CHAPTER 3

Image Acquisition

Part of the attraction of visual observing is that it is fairly easy to do—especially with small instruments. Just set up your telescope and begin observing. For example, with my Questar it takes literally only a few minutes to set it up and begin visual observing. If the object happens to be low near the horizon or the “seeing” is for some other reason less than ideal, it’s no great loss as it only took a few minutes to check. Also, if the polar alignment is off and objects slowly drift from the center of the field, again it’s no big deal just tweak the aim as needed. If at the end of the evening you had great views of things you’ve never seen before you retire with a real sense of satisfaction. On the other hand, if things simply didn’t come together then you just put the toys away and realize you can try it all again another day.

With imaging, it’s not quite so simple. At the end of the night you have a picture to represent what you have accomplished. If it is a perfect picture, there is no limit to the amount of pride you feel in having taken that picture yourself. You can frame it and hang it up in a home gallery that you can proudly show friends and family. On the other hand, if at the end of the night the picture you’ve taken isn’t perfect, the focus wasn’t right or the polar alignment was off or whatever, then the final flawed image is a constant reminder of mistakes made. Your only recourse is to erase the picture from your computer so that you don’t have to face the continual humiliation of your mistakes. In other words, in astroimaging the emotional stakes are higher than with visual observing. Therefore, it is highly desirable to minimize failures. And, the only way to do this is to attend to every detail. The price is that it is no longer a simple two-minute process to get started. For me, it generally takes about an hour from the moment I walk out the door until I starting saving my first image. This chapter is about the various details that must be attended to. We begin with some general tips:

- If possible, image after midnight. Light pollution is less severe late at night. Also, ambient temperatures remain relatively constant late at night—the biggest and fastest changes occur in the hours just after sunset.

- Image objects high in the sky—the higher, the better. The photons from an object at 45° altitude pass through 41% more atmosphere than an object straight overhead. At 30° there is 100% more atmosphere to pass through.
- Allow enough time for the telescope to equilibrate with ambient temperature.
- Polar align carefully.
- Focus carefully.

In the next few sections, I'll elaborate a little on these last two points.

1. Polar Alignment

Polar alignment is important. It is especially critical if doing unguided exposures. There are lots of methods for aligning. The goal is to find a method that balances time and effort with the quality of the alignment. The highly accurate *drift alignment* method, in which one makes small incremental adjustments to the alignment based on observed drift of a star, yields essentially the best possible alignment. But, it can take hours to accomplish. It is appropriate for use in permanent observatory installations but is too time consuming for those of us who wish to set up and tear down every night.

Most premium German equatorial mounts have a polar alignment scope built into the polar axis. The alignment scope has a very wide field and includes an illuminated reticle that shows some of the bright stars near the north celestial pole. To use it, you rotate the reticle and adjust the mount until the star pattern shown in the reticle matches the observed star field. This method of aligning works very well and only requires a minute or two to get a very good alignment.

A popular telescope mount is the so-called *fork mount*. My Quesar has this type of mount as do most other Maksutov-Cassegrain and Schmidt-Cassegrain telescopes. To set up such a telescope for astrophotography, it is necessary to place the mount on an equatorial wedge which sits on a tripod or a pier. Polar alignment is achieved by adjusting the wedge so that it is perpendicular to the axis of rotation of the Earth. In other words, it must be adjusted so that the fork arms point parallel to the pole axis. I'll describe two quick and easy methods to achieve this.

Method 1. First, do a rough polar alignment in which we assume that Polaris is at the north celestial pole. To do this, set the declination to 90° and adjust the mount so that Polaris is centered in the eyepiece. Then aim at a star with known right ascension use it to calibrate the right ascension. Finally, point the telescope to the coordinates of Polaris

(right ascension 2h 32m, declination $89^{\circ} 16m$) and adjust the mount so that Polaris is in the center of the field. This method is easy and works well. But, how well it works depends on how accurate the setting circles are. Just because the declination circle reads 90° does not necessarily mean that the telescope is aligned with the mount's polar axis. The following method does not rely on the precision of the setting circles.

Method 2. First it is important to set the declination to 90° . As mentioned already, just because the declination setting circle reads 90° does not necessarily mean that the telescope is actually aligned with the polar axis of the mount. It is important to check this. To check it, look through the eyepiece and rotate the scope around the polar axis; i.e., rotate in right-ascension. The center of rotation is supposed to be in the center of the eyepiece. In practice it almost never is. Adjust the declination to get the center of rotation as close to the center of the eyepiece as possible. It is best to use a wide-field eyepiece for this—I use a 32mm Brandon eyepiece. You will probably find that it is impossible to get the center of rotation exactly to the center of the eyepiece. At its best, it will be a few arcminutes either to the left or right of center. This happens when the right-ascension and declination axes are not exactly perpendicular to each other. On my telescope, I can only get to within about 7 arcminutes of the center. If it is very far off, you might consider making adjustments to your mount in order to bring the axes closer to perpendicularity. Once you've got the declination set as close as possible to true 90° , lock it.

The next step is to adjust the mount so that the north celestial pole coincides with the mount's polar axis, which we identified as the center of rotation in the first step (and might not be at the center of the eyepiece). Again, a wide-field eyepiece hopefully providing a 1° field of view, or more, is desirable. Then you can star hop from Polaris to the correct location using a little finder chart such as the one shown in Figure 3.1. I recommend installing a planetarium program on your computer and using it to make a customized finder chart. With such a program, you can size the chart to any convenient size, you can flip the chart east-west to match the view in the eyepiece, and you can adjust the relative brightness of the stars and the magnitude threshold so that the picture matches very closely what you see in the eyepiece. The chart shown in Figure 3.1 was carefully prepared to match the view through the Questar using a 32mm Brandon eyepiece.



FIGURE 3.1. Finder chart for the North Celestial Pole (NCP). Polaris is the bright star at the center of the eyepiece circle. The NCP is the red crosshair near the top of the image.

2. Attachments

After getting a good polar alignment, the next step is to attach the CCD camera. Most modern astronomical equipment is designed for *T-threaded* attachments, which are 42mm in diameter and have 0.75mm pitch. Questar telescopes and Starlight Express CCD cameras do not use T-threads. They are both based instead on *P-threads*, which are 42mm in diameter but have 1.0mm pitch. These two threadings are sufficiently similar that it is easy to cross thread a P-thread accessory into a T-thread one. You get about half a turn and then the threads bind. Forcing the threads can ruin both pieces, so it's important to keep T-threaded accessories distinct from P-threaded ones.

The Starlight Express cameras, like most CCD cameras, come with *nose pieces* that allow the camera to be used like either a 1.25" or 2" eyepiece. Using the 1.25" nose piece, a CCD camera can be placed in the eyepiece port of the Questar telescope. The main disadvantage of this configuration is that the camera cannot be locked down to prevent rotation—the last thing one wants is for the camera to rotate during

a long exposure. For this reason, it is better to attach the camera to the *axial port*. To make this attachment, the Questar *swivel coupler* is attached to the axial port. The camera end of the swivel coupler is male P-threads. The Starlight Express cameras have female P-threads and hence can be attached directly to the swivel coupler. Alternatively, one might wish to add a *filter holder* and/or a *focal reducer*. We'll discuss these at more length in the gallery chapters devoted to color images and wider field imaging. The swivel coupler has two thumb screws that when loose the camera can rotate but when tightened hold the camera secure from rotation. To prevent confusion when trying to understand the orientation of an image, it is important to make sure that the CCD camera is oriented in a consistent manner. I always orient it so that rows of pixels in the camera are oriented parallel to the east-west line. One could choose to orient them parallel to the north-south line if one wants. The important thing is to do it the same way every time. Otherwise, you will certainly be confused most of the time.

In addition to attaching the camera to the telescope, you also need to attach the camera's power cable to a power supply and its computer cable to the computer. Until recently, all amateur astronomical CCD cameras attached to computers via the parallel computer port. There is a big disadvantage to the parallel port: download times are very long, 20 seconds or more is common. Recently, Starlight Express led the way in offering the option to connect via a USB port, which is much faster. Download times for the MX-916 model are about 6-seconds for a full frame image.

3. Focusing

Focusing is the most tedious step in preparing to take CCD images. The basic process is as follows. First, fire up the image acquisition software, which in my case is AstroArt, and point the telescope at a fairly bright star. Next, take a short exposure image. After downloading the first image to the computer, an out-of-focus image of the star appears. Draw a small rectangle around the out-of-focus star and then enter focus mode. In focus mode, just the pixels inside the rectangle are downloaded. Hence, downloads are much faster. Setting the focus exposure to say 0.3 seconds provides a few focus window updates per second. In addition to an image of the star, the focus window also shows the number of counts for the brightest pixel as well as the "width" of the star image expressed as a "full-width-half-max" (FWHM) value measured in pixels. Each of these pieces of information are helpful in achieving critical focus. Each one also fluctuates from one image to the



FIGURE 3.2. A Hartmann mask made from a 4" PVC endcap.

next making it a challenge to know exactly when best focus has been achieved. When changing the direction of approach to critical focus, the star's image shifts rather significantly. This *image shift* is a common problem with focus systems in which the primary mirror is moved back and forth. Changes of direction cause the mirror to tilt slightly differently and so the image moves. The amount of image shift varies from one instrument to the next. On Questars, it varies from about 20 to 60 arcseconds. This is enough to move the star out of the field of view of the focus window. The focus tool provides the capability of moving the window this way and that but it is tedious to do so. It is best to make the focus window big enough so that the star won't shift out of the field-of-view but not too big because then image download time becomes too large.

In addition to the steps described above, one can employ various focus aids such as a Hartmann mask or a diffraction-spike mask. I use a Hartmann mask as shown in Figure 3.2. It consists of a 4" PVC endcap with three holes drilled in it and painted black. Placing this mask over the dew shield, an out of focus stars appear as three smaller out of focus stars. As critical focus is approached, the three star images converge on each other to make just one in-focus star. The biggest danger of using a Hartmann mask is forgetting to remove it once critical focus has been achieved.

4. Finding Faint Fuzzies

Many of the targets of interest are completely invisible visually. An accurate method of pointing is critical. Certain accessories help a great deal:

- A *reflex sight*. These devices project a few concentric red rings directly on the sky so that you can see exactly where the telescope is pointed. Telrad makes by far the most popular unit but it is too big and heavy for use on small telescopes. I find that Rigel's QuickFinder works very well. There are some other small units that ought to work as well although I haven't tried them.
- A *flip mirror system with a parfocal eyepiece*. The Questar telescope comes with a built-in flip mirror system. However, extension tubes are needed to make an eyepiece in the eyepiece port parfocal with a CCD camera attached to the axial port. A few threaded 1"-long, 1.25"-OD extension tubes with a clamp ring are useful in making a parfocal system.
- *Digital setting circles*. These consist of small encoders that attach to the telescope mount in such a way as to measure changes in right-ascension and declination. The encoders are attached to a small, battery operated, control box which is essentially a small computer. The control box has a database of popular objects. After selecting an object in the database, the LED screen shows how many degrees to move up or down in declination and how many degrees to move east or west in right-ascension to center the object. When the screen reads 0-0, the object ought to be in the field of view—or at least close to it. Of course, for objects not in the database, one can set the LED screen to read out right-ascension and declination.

It is important to realize that DSCs are, at best, only accurate to about 10 arcminutes. This is not quite good enough to guarantee that the desired object will appear in the field of view if the field of view is say 15 by 10 arcminutes. However, if the object itself is an extended object, which is often the case, usually part of the object will be in the field. Of course, this is a best-possible scenario. If the DSCs aren't initialized carefully, or if the polar alignment is off, or if the RA and Dec axes are not orthogonal, then the accuracy of the DSCs will be accordingly less.

Not only are DSC's great for finding objects, they are also useful for taking a quick glance at the LED screen during a long exposure and verifying that the telescope is still tracking its target.

There are several manufacturers of after-market DSCs. However, currently only Jim's Mobile Inc(JMI) offers DSCs with

custom mounting hardware designed for attachment to the 3.5" Questar.

- On-computer *planetarium program*. CCD chips are small and therefore provide a small field-of-view compared to 35mm film. For example, on the 3.5" Questar, without a focal reducer, the field of view is about 20 arcminutes by 15 arcminutes. Even with everything set up as carefully as possible, it is usually the case that the target object does not appear in the first image taken. What to do? Here, a computer planetarium program can be very useful. You can display on the computer screen a field of view that covers a degree or two. If you set the image orientation in the planetarium program to match that of the images taken with the camera it is usually easy to match star patterns and figure out which way to move to center the object.

CHAPTER 4

Image Processing

At the end of a night of visual observing, one retires with pleasant memories of things observed. Memories are nice and they last a lifetime but they are hard to share with those who weren't there.

At the end of a night of CCD-imaging, you go inside and warm up to a computer that has stored on it several, possibly hundreds, of images waiting to be processed. If the excitement and anticipation of seeing the final result outweighs physical exhaustion, you can process the images immediately. With some experience it really doesn't take very long, perhaps only an extra 30 minutes. However, if exhaustion is taking over, you can always retire knowing that the images will be there in the morning waiting for you (as if you didn't have a day job!). In this chapter, I give a brief summary of what is involved.

1. Raw Image Frames

Figure 4.1 shows two raw images of NGC 3628 (the third galaxy in the Leo Trio). These are six minute guided exposures. On the night in question, I took 50 of these images—yes, that's right, 5 hours worth. As you can see, the individual images look rather noisy.

For those new to CCD imaging, there is already much to be said about these images. First of all, each image is a 752×580 array of




FIGURE 4.1. Two raw image frames of NGC 3682.

numbers which corresponds to a physical array of *pixels* on the CCD chip. On the computer, this array is stored as a 16-bit FITS file, which is the standard format for storing astronomical images. Without going into the details of this file format, it suffices to say that the light intensity associated with each pixel is a 16-bit unsigned integer. What does that mean, you ask? It means that at each pixel the light intensity is stored as a number between 0 and $2^{16} - 1$ or, in other words, between 0 and 65535. The number 0 means no light reached the pixel.

A positive number represents, roughly, the number of visible light photons that hit a given pixel. I say roughly because there are two basic steps in going from photons to numbers stored in the computer. First of all, a photon hitting a pixel in a CCD chip frees an electron which gets stored in an electron “well”. But not every photon is successful at dislodging an electron. The fraction of successful dislodgings is called the *quantum efficiency*. It depends on wavelength. A rough approximation is that about 50% of the visible light photons produce an electron. The second step is the counting of the electrons in each well at the end of an exposure. This counting process isn’t perfect and it involves applying a “gain” factor. For most CCD cameras, the gain factor is about 2. The quantum efficiency and the gain factor approximately cancel each other out and so it is approximately true that the numbers recorded in the FITS file are about the number of visible light photons that struck the CCD chip during the exposure.

In a displayed image, pixels with lots of photons are shown as white, pixels with few or none are shown as black, and pixels in a middle range are shown as gray. Neither a computer monitor nor a printed page can display 65535 different levels of gray. Even if they could, the human eye can’t discern that many gradations. Hence, one resorts to various stretchings of the image data. Look, for example, at the left-hand image in Figure 4.1. There are two sets of numbers below the image. The first one reads R=2670–25987. This means that the darkest pixel in the raw image has a count of 2670 while the brightest pixel is 25987 counts strong. This is the *range* of pixel values. The second set of numbers reads V=2803–2914. This means that in the displayed image, everything below 2803 is shown as black, everything above 2914 is shown as white, and values between 2803 and 2914 are shown with various gray levels. Note how remarkable this is. The brightest star, which is in the upper left corner, has at its core a pixel with about 23000 counts above the background level of about 2700. Yet, the galaxy, which is the smudge in the middle only shows about 100 counts above the background. It is much fainter than that bright star! And that bright star is a magnitude 10 star meaning that at the eyepiece it is a very faint

object. In fact, if this image were restretched to cover the entire range from 2760 to 25987 with a gray scale, all we would see would be the three brightest stars. And that is all one sees at the eyepiece of the 3.5" Questar.

Okay, so we've managed to detect a faint fuzzy that one could not hope to see visually through a telescope of modest aperture. But, the images are not very aesthetically pleasing. They have a variety of obvious defects. Can we fix them? For the most part, the answer is yes. To see how, it is important first to consider what is in the image. The image consists of several "signals" overlaid on top of each other. Some photons come from space—we like these. In the images in Figure 4.1, we see that there are a number of stars and the galaxy NGC 3628 is visible, barely, in the center of the frame. Another source of photons is from diffracted light in the atmosphere. These are mostly from *light pollution and the moon*—we usually don't want them. Light pollution normally gives a roughly uniform glow across the field. It is part of the reason that the smallest pixel values in the images are more than 2750 (although only a small part of the reason as we shall see). Moon glow is like sky glow but usually has a characteristic gradient to it. That is, it will be brighter in the corner nearest the moon and darker in the opposite corner. The images in Figure 4.1 do appear to have a gradient in them but these images were taken when the moon was not out. This gradient is due to another source of photons—*amplifier glow*. There is an amplifier circuit near the upper left corner of the chip in the Starlight Express cameras. This amplifier generates a small amount of heat. All warm bodies give off infra-red radiation and CCD cameras tend to be quite sensitive in the infra-red. Hence, these photons too get recorded as part of the signal that makes up the image. These infra-red photons corrupt all parts of the image—there's just more of it in the amplifier corner. Altogether the signal due to infra-red photons is called *dark current*. The colder the camera, the lower the dark current. In fact, all astronomical CCD cameras are cooled, usually to 30 or 40 degrees Celsius below ambient temperature, in an attempt to minimize the dark current. Finally, there is a certain number of counts that are generated during *readout*—that is, they are generated when the electrons on the chip are counted and sent off to the computer. These aren't due to photons, per se, but must be accounted for nonetheless. On the Starlight Express cameras, this readout signal amounts to a count of about 2635 at each pixel plus or minus about 5. If every pixel had a readout signal of 2635 all of the time, then it would be trivial to remove—just subtract it. But, unfortunately, these numbers are random and furthermore, different pixels have a slightly different average values.

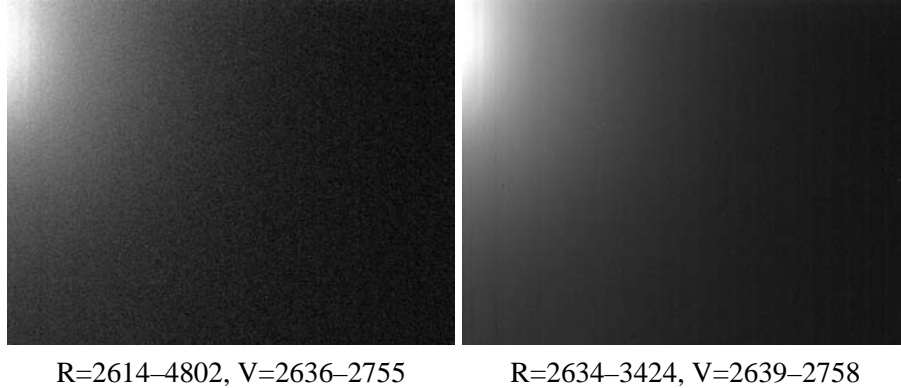


FIGURE 4.2. *Left.* A dark frames. *Right.* An average of 30 dark frames.

2. Dark Frames

So, how do we remove these unwanted signals? First, one removes the readout and dark current signals by taking another image of the same duration and at the same temperature as the original image but with the front of the telescope covered so that no space or sky photons reach the chip. This is called a *dark frame*. The image on the left in Figure 4.2 shows a dark frame that was taken to match the frames in Figure 4.1. It contains the readout signal plus the dark current—and nothing else. Hence, it can be subtracted from the original image to remove these two unwanted signals from the original. But, before doing that, we need to bear in mind one thing: if we take two grainy images and subtract them, the result is an image that is even grainier. The graininess in both the original image and the dark frame is due to the fact that the number of photons arriving at a pixel in a given time is random. Two images taken under identical conditions but at different times will be similar to each other but not identical. It is exactly the same as taking a Geiger counter, putting it near some radioactive material, and recording how many counts one gets in a fixed time interval. If one repeats the experiment several times, the results will be slightly different each time. Assuming that the half-life of the material is very long, the average number of counts will not change over time. But each individual experiment will vary around this average. If the average is about N , then the variation is about \sqrt{N} . Whether doing experiments with radioactive materials or doing CCD imaging, we are generally interested in the average value rather than the outcome of one particular experiment. The easiest way to estimate the average is to repeat the experiment many times and compute the numerical average. The image on the right in




R=94-22019, V=166-232

FIGURE 4.3. Raw image frame minus average of 30 dark frames.

Figure 4.2 shows an average of 30 dark frames. Note that the graininess has largely disappeared. We now see a nice smooth amplifier glow in the upper left corner. We also see a number of bright pixels. These are called *hot pixels*. Every CCD chip has some hot pixels. The number of dark counts in hot pixels grows with time just like all other pixels but it grows more rapidly. If we subtract this averaged dark frame from the original raw image frames, we will remove the effects of readout, amplifier glow, and other dark current including hot pixels all at once. The result is shown in Figure 4.3.

3. Bias and Flat Frames

From Figure 4.3, we learn a few more things. First, the sky glow amounts to about 150 photons per pixel, which is about the same intensity level as the galaxy itself. This galaxy has a surface brightness magnitude of 14.10. The fact that the sky glow is about the same tells us just how bad light pollution is in suburban New Jersey. Also, there are some other defects that have become evident. For example, there is some obvious dimming around the edges. This is called *vignetting*. Also, we can see a ring shaped halo around the center of the image and some dark splotches. It's hard to say exactly what caused the halo but the dark splotches are shadows of tiny dust particles on the window that covers/protects the CCD chip. These splotches are called *dust donuts*. The vignetting, halo, and dust donuts can all be removed using what is called a *flat frame*. A flat frame is an image taken of a uniformly illuminated, out-of-focus field. There are several techniques for obtaining flat frames including taking an out-of-focus image of the sky at dusk or dawn. It is important to get a very accurate flat frame. To this end, the



R=2633–2657, V=2631–2650

R=26317–33722, V=27451–36920

FIGURE 4.4. *Left.* Forty 0-second bias frames averaged.
Right. A flat frame.

most reliable method for obtaining flat frames is to build a *light box* that fits over the front of the telescope (with dew shield in place!) and provides a flat, dim, very out-of-focus, field to image. I made my light box from a battery operated, dome-shaped, closet light glued onto the end of a 4" PVC pipe coupler with a few sheets of milky plastic glued in. The whole contraption cost nothing more than pocket change and took only an hour or so to assemble. Using this light box, I obtained the flat field frame on the right in Figure 4.4. This flat field needs to be calibrated by subtracting an appropriate dark frame from it. As the exposure is typically only a fraction of a second, there is almost no dark current to subtract—the main thing to subtract is the readout signal. The readout signal can be imaged by taking a 0-second black image. Such an image is called a *bias frame*. As before it is best to average several of these. The left-hand image in Figure 4.4 shows an average of 40 bias frames. Using the averaged bias frame to calibrate the flat frame, the next step is to rescale the numbers in the calibrated flat frame so that they average about one and then divide the calibrated image frame by the normalized, calibrated flat frame. In this way, areas that are too bright get dimmed a little while areas that are too faint get brightened somewhat. The result is that the dust donuts, glow ring, and vignetting are all effectively removed as Figure 4.5 shows.

4. Image Calibration

The steps describe in the previous three sections are called *image calibration*. The process is summarized in Figure 4.6. It might seem tedious, but image processing software is carefully designed to make this task very easy. One simply selects a set of raw images, a set of dark



FIGURE 4.5. Fully calibrated image.




FIGURE 4.6. Summary of the image calibration process.

frames, a set of bias frames, and a set of flat frames and then with the click of the mouse all images will be calibrated in just a matter of seconds. Actually acquiring the calibration images may also sound rather onerous, especially if one must collect new images for every imaging session in an effort to ensure that these images are obtained under exactly the same conditions as the raw image frames. It's not as bad as it might seem.

First of all, averages of bias frames change very little over time. In fact, one can collect a large number of bias frames, average them, save the average, delete the originals, and use this one bias frame essentially indefinitely. There are some slow changes as the chip ages but they can often be ignored. If one is worried, one could remake their bias frame once a year or so. As the interexposure time is just the amount of time

it takes to download each image to the computer, a collection of 50 bias frames can be acquired in just a matter of a few minutes. And, this can be done during the daytime. Of course, most cameras operate in a few different modes and one needs bias frames for each mode. Still, a whole collection of bias frames can be created in under an hour. Do it once, and forget about it.

Conventional wisdom is that dark frames should be taken as close as possible to the same time as the raw image frames mainly so that the ambient temperature will be almost the same during the two exposures. In the early days of CCD imaging (before I was doing it), this might have been very sound advice but recent advances make this entirely unnecessary. The reason is that it is possible to scale dark frames so that the hot pixels on the dark match the corresponding pixels on the raw frames. All software packages have routines that will do this automatically. Hence, dark frames don't have to be taken at the same temperature and they don't even have to have the same exposure as the raw frames. Of course, to be safe, it is a good idea to keep things roughly similar. As an example, the 30 dark frames used in the calibration of NGC 3628 were taken during the daytime. The exposures were the same (6 minutes) but no effort was made to match the ambient temperature. The raw image frames were taken in early March on a night where the temperature dropped significantly over the 5 hours during which the images were acquired. The dark frames were taken with the camera in the refrigerator during the daytime (while no one was home so there was no danger of anyone opening the refrigerator or otherwise bumping into things). So, even though it can take several hours to collect a large set of dark frames, they only need to be collected once (or maybe twice—a summer set and a winter set) and can be collected in the daytime and the software can be programmed to collect the entire set without user intervention.

This leaves us with flat field frames. Flat fields turn out to be the most delicate frames to collect. They frames must be acquired using the telescope in exactly the same configuration as used for the raw frames. The camera should not be allowed to rotate in its coupling to the telescope. Even the dew shield must be kept to a fixed extension. And, of course, the light box must provide a perfectly flat illumination. If everything is held constant and the light box truly provides a very uniform flat field, then only one or just a few flat field frames need to be collected. It only takes a few seconds to acquire a single flat frame so collecting 5 or 10 can be done in less than a minute.

5. Stacking

The calibrated image of NGC 3628 shown in Figure 4.5 is not the work of art that we'd like it to be. But it only suffers from one deficiency—not enough photons. The sky glow is contributing on the average 150 photons per pixel and the galaxy is contributing about that many again. This means that two adjacent background pixels, which ought to have the same value of about 150 are likely to vary between $150 - \sqrt{150} = 138$ to $150 + \sqrt{150} = 162$. In fact, only two thirds of the background pixels will fall into this range. We have to make the range twice as wide, from 126 to 174 to include 95 percent of all background pixels. This is why the image looks so grainy. The technical jargon here is that the *signal* (150) to *noise* ($\sqrt{150}$) ratio is inadequate. The only solution is to collect more photons. Here again we see the beauty of CCD imaging and digital image processing. Rather than starting over and doing a (much!) longer exposure, we can collect lots of images and add them together in the computer. This is called *stacking*. Figure 4.7 shows the result of stacking 5, 15, and 50 six-minute images.

The final image, while not perfect, does begin to look quite respectable. Of course, the total image integration time is 5 hours. These images were acquired using a focal reducer which brought the Questar to about f/10. This is one of the main lessons. Imaging galaxies at f/10 requires hours of exposure to get nice results. I'm often asked what is the correct exposure time. The answer is infinity. The longer the exposure, the better the final result. The vast majority of mediocre images one sees suffer from just one deficiency—not enough photons because the total integration time was too short. Imaging faint objects at f/10 requires one important trait: patience. The final image of NGC 3628 is shown on page ??.

6. Analyzing the Image

It is instructive to analyze the final stack of 50 images shown in Figure 4.7. Here are some statistics:

- The average background level is about 248.
- The brightest pixel in the galaxy has pixel value 315.387. This value and values in general are no longer integers because we have “divided out” by a normalized flat field and we have averaged 50 frames.
- The brightest pixel in the entire image is at the center of the bright star in the upper left corner. It is the magnitude 9.85 star SAO 99572. Its value is 12953.363.



FIGURE 4.7. *Top left.* One image. *Top right.* Stack of five images. *Bottom left.* Stack of 15 images. *Bottom right.* Stack of 50 images.

It is interesting to note that the total number of “counts” associated with the bright star is 95090.586. This number is computed by subtracting the local background around this star and then adding all values associated with it. At one extreme, all 95090.586 counts could have landed on one pixel. But, achieving that would have required a smaller focal ratio (to make the star’s Airy disk and first few diffraction rings all fall onto one pixel), perfect atmospheric seeing, and precise tracking for the 5-hour total exposure time. In fact, given the difficulty of the task, it seems rather remarkable that we were able to get about 15% of the photons to land on one pixel and more than 50% of them to land in a 3×3 box centered on the brightest pixel.

The “tightness” of a star is usually reported as a number called its *full width half max* or *FWHM*. This number refers to how many pixels it is from the point on one side of the brightest pixel where the intensity has dropped by half its maximum to the analogous point on the other side. Interpolation techniques are used so that this number doesn’t have to be an integer. For the brightest star in the image, the FWHM is 2.623. A nice thing about FWHM is that it is fairly independent of which star is

used to measure it. For example, the other two bright stars in the image have FWHM values of 2.860 and 2.831. Picking two other dimmer stars at random, we find values of 2.539 and 2.703.

Measuring Image Scale and Focal Length. Using a computer planetarium program and an image processing program, it is easy to compute the image's scale. Indeed, using the planetarium program, we find that the distance from SAO 99572 to the star bright star below the galaxy (TIC 0861) is 20.308 arcminutes. Using an image processing program, we find that the same pair of stars are 476 pixel units away from each other, meaning that if the image were rotated so that these two stars were on the same row of pixels they would be 476 pixels apart. Dividing these two distances, we get the size of a pixel in arcseconds:

$$\frac{20.308 \text{ arcminutes}}{476 \text{ pixels}} \times \frac{60 \text{ arcseconds}}{\text{arcminute}} = 2.56 \text{ arcseconds/pixel.}$$

Using the additional fact that one pixel on the Starlight Express MX-916 camera is 11.2 microns across, we can compute the precise *effective focal length* of the instrument, as configured. First, we convert the pixel distance to a physical distance at the image plane by multiplying by the size of a pixel:

$$476 \text{ pixels} \times 11.2 \text{ microns/pixel} \times 0.001 \text{ mm/micron} = 5.331 \text{ mm.}$$

This physical distance corresponds to the known angular separation of 20.308 arcminutes. Hence, we can calculate the focal length using:

$$\text{focal length} = \frac{5.331 \text{ mm}}{\sin(20.308/60)} = 903 \text{ mm.}$$

This focal length is significantly shorter than the nominal Questar focal length of about 1350mm. The reason is that these images were all taken using a focal reducer. See Chapter ?? for more on using focal reducers.

Finally, we note that once the focal length is known, it is trivial to compute the *focal ratio*. In fact, one just divides by the aperture, which for the Questar is 89 mm:

$$\text{focal ratio} = \frac{903 \text{ mm}}{89 \text{ mm}} = 10.1.$$

A seven inch telescope with the same focal length would have half the focal ratio, that is, it would be about an f/5 instrument. With such a telescope, one could expect to get an image of similar quality in one fourth the exposure time or, in other words, in a little more than an hour.

Magnitudes and Sky Glow. As we mentioned earlier star SAO 99572 has peak pixel value about 12953 and total value of about 95091. In comparing stellar brightnesses one takes ratios of pixel values. In so doing, one will get about the same answer whether using the peak value or the total value. These two values appear in about the same proportion throughout an image. This is because ideally each star has the same profile, the so-called *point spread function* or *psf*, and so if one star's peak is, say, 5 times larger than another star's peak, then its total value will also be about 5 times larger. Since the total value is a larger number, it has less randomness inherent in it and hence is better for doing careful *photometry*. Also, the peak value can be significantly affected by whether the star is well centered on that pixel or straddles two pixels equally. This second problem is not too serious when the FWHM is on the large side, say 3 or more, and it becomes a critical issue as the FWHM decreases toward 1.

Given the FWHM's in the stacked image of NGC 3628, peak values should be a pretty good surrogate for stellar brightness. The advantage of peak values is that they can be compared directly with surface brightnesses of extended objects such as the galaxy or even the sky glow. The fundamental formula for the magnitude difference between two parts of the image A and B is:

$$m_A - m_B = -2.50 \log \frac{I_A}{I_B}.$$

Using the fact that the magnitude 9.85 star SAO 99572 has peak brightness over background of $12953 - 248 = 12705$ and galaxy NGC 3628 has peak brightness above background of $315.387 - 248 = 67.387$, we can compute an effective peak brightness for NGC 3628:

$$9.85 - m_{\text{galaxy}} = -2.50 \log \frac{12705}{67.387}$$

which reduces to $m_{\text{galaxy}} = 15.5$. This means that the brightest part of the galaxy is about as bright as a magnitude 15.5 star in our image. We must emphasize that this is a computation that is relative to the image at hand. If the focus, the seeing, or the tracking had been worse, then the psf would have been broader and the peak value for the star would have been smaller. The peak value for the galaxy on the other hand would not change much since it is fairly constant over several neighboring pixels. Small errors in focus, atmospheric seeing, or tracking cause faint stars to disappear into the background sky glow.

Using total values rather than peak values for stars avoids the systematic errors caused by focus errors, bad seeing, and tracking errors. As we computed earlier, one pixel corresponds to 2.56 arcseconds both horizontally and vertically. Hence a pixel covers $2.56 \times 2.56 = 6.55$ square arcseconds. Dividing the per-pixel background level of 248 by the area of a pixel and using the total value for SAO 99572, we can derive the true *surface brightness* of the sky glow:

$$m_{\text{skyglow}} = 9.85 + 2.50 \log \frac{95091}{248/6.55} = 18.3.$$

This gives a magnitude that can be compared to a star that covers exactly one square arcsecond. Sometimes surface brightnesses are computed based on smearing a star over a square arcminute. A star has to be 3600 times brighter to smear over a square arcminute with the same resulting brightness as if it had been smeared over only a square arcsecond. Hence, magnitudes based on square arcminute coverage are brighter by $2.5 \log 3600 = 8.9$ magnitudes. Hence, we see that the surface brightness of the sky glow calculated on a square arcminute basis is $18.3 - 8.9 = 9.4$.

Sky glow surface brightness is usually reported on a square arcsecond basis whereas surface brightnesses of galaxies are usually reported on a square arcminute basis. The peak surface brightness of NGC 3628 is calculated as follows:

$$m_{\text{galaxy}} = 9.85 + 2.50 \log \frac{95091}{67.387/6.55} - 8.9 = 10.9.$$

This value compares favorably with reported average surface brightnesses of from 12 to 14 depending on the source catalog. Average surface brightnesses are hard to define precisely because it is not clear how much of the galaxy to include in the calculation.

CHAPTER 5

Image Enhancement and Presentation

In some cases such as with NGC 3628, the image that results after calibrating and stacking a set of raw images cannot be improved very much by further tinkering. However, in most cases, further significant enhancements can be obtained. There are a few reasons one might wish to do more. The two most obvious are:

- To sharpen a blurred image.
- To adjust the dynamic range of the pixel values.

We will discuss these two issues in this chapter and illustrate the standard techniques for addressing them.

1. Unsharp Mask

Sometimes an image is blurred simply because of problems that occurred during image acquisition: perhaps the telescope was not focused properly or maybe there were tracking problems with the mount. In such cases, the techniques described in this section can help but really one ought to consider simply fixing the core problem and taking a new set of raw images.

However, there are certain situations when an image will be inherently blurry and it becomes desirable to sharpen it as much as possible. The two most obvious reasons for inherent blurriness are:

- Imaging at high f -number.
- Dynamic distortions for atmospheric irregularities—so called *seeing*.

For these cases, the techniques presented here work amazingly well.

Planetary imaging provides an excellent example of a situation where one typically works at a high effective f -number. Such images are normally taken using either a *Barlow lens* or a regular eyepiece for *eyepiece projection* (see Chapter ??). In either case, the f -number is often around 50 or more. At such high f -numbers, a star becomes a large blob called an *Airy disk* surrounded by a concentric sequence of successively fainter diffraction rings. An image of a planet is similarly blurred. There is nothing one can do up front about this blurring of images—it is simply

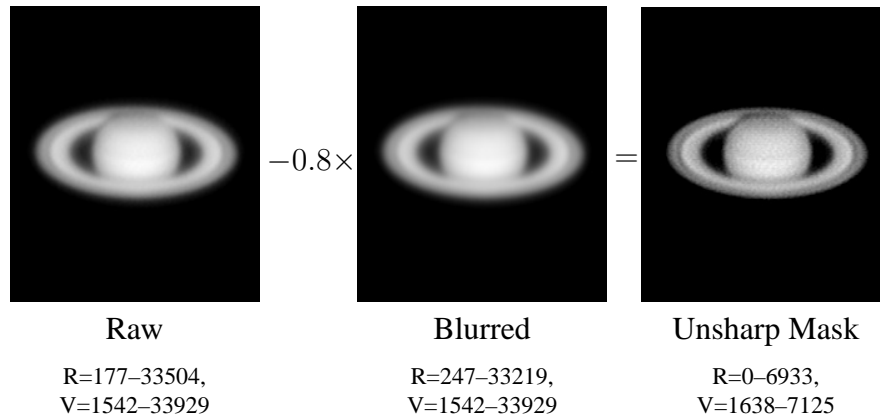


FIGURE 5.1. Summary of the unsharp masking process. The blurred image has been blurred three times with a Gaussian kernel.

a consequence of the wave nature of light. However, since this blurring is very well defined, it is possible to create specific computer procedures that can largely undo it. The simplest such technique is called *unsharp masking*, which we shall now describe.

The left-most image in Figure 5.1 is a stack of 100 calibrated 0.1 second images of Saturn taken with a CCD video camera mounted on a 3.5" Questar using eyepiece projection to raise the effective f -ratio to about 45. Note how the image looks okay but is somewhat soft. The telescope was focused very carefully prior to taking the sequence of raw images—the blurriness is due to the large f -number and perhaps the effects of atmospheric seeing (although with such a small aperture, this problem is often not the main source of blurring).

The basic idea of unsharp masking is to make a copy of the original image, further blur the copy, multiply each pixel in the blurred copy by a fixed number between 0 and 1, and then subtract the blurred copy from the original. The effect of this strange process is to remove the blurry part of the original image and leave behind the sharp part. The result shown on the right in Figure 5.1 is a rather significant improvement over the original image. Trial-and-error must be used to determine the constant by which to multiply the blurred image. For the Saturn image in the figure, 0.8 seemed to provide the best result.

Kernel Filter Smoothing. To complete the description of unsharp masking, we must explain how one obtains a blurred copy of the original. This is done using a *kernel filter*. Kernel filtering is a very general smoothing technique. We shall describe just one specific version called

the *Gaussian kernel filter*. A kernel filter is defined by a small array of numbers, called the *kernel mask*. For the Gaussian kernel, the mask is the following:

$$\begin{bmatrix} \frac{1}{16} & \frac{2}{16} & \frac{1}{16} \\ \frac{2}{16} & \frac{4}{16} & \frac{2}{16} \\ \frac{1}{16} & \frac{2}{16} & \frac{1}{16} \end{bmatrix}$$

Note that the numbers add up to one. Here's how the mask is used. For each pixel in the unsmoothed image, the current pixel value is multiplied by the value at the center of the mask ($4/16$) and to this is added the values of each of its eight neighbors weighted by the corresponding values in the mask. This total then becomes the new pixel value. This process is repeated for every pixel in the image (border pixels have to be treated slightly differently). For example, suppose that an image has the following pixel values in rows 87–90 and columns 123–126:

$$\begin{array}{cccc} & 123 & 124 & 125 & 126 \\ 87 & \left[\begin{array}{cccc} 1432 & 1510 & 1717 & 1793 \end{array} \right. \\ 88 & \left[\begin{array}{cccc} 1654 & 1723 & 1925 & 1981 \end{array} \right. \\ 89 & \left[\begin{array}{cccc} 1411 & 1522 & 1704 & 1802 \end{array} \right. \\ 90 & \left[\begin{array}{cccc} 1623 & 1730 & 1901 & 1977 \end{array} \right. \end{array} \cdot$$

Then, using the Gaussian kernel mask we can compute replacement values for the pixel in row 88 and column 124 as follows:

$$\begin{aligned} 1723 \Rightarrow & \frac{1}{16}1432 + \frac{2}{16}1510 + \frac{1}{16}1717 + \\ & \frac{2}{16}1654 + \frac{4}{16}1723 + \frac{2}{16}1925 + \\ & \frac{1}{16}1411 + \frac{2}{16}1522 + \frac{1}{16}1704 = 1648.6 \end{aligned}$$




FIGURE 5.2. Kernel filtering is an example of the convolution operation.

Other replacement values are computed in the same way:

$$1925 \Rightarrow 1786.1$$

$$1522 \Rightarrow 1645.4$$

$$1704 \Rightarrow 1782.9.$$

Applying this procedure to every pixel in an image results in a smoother image.

The technique we have just described of positioning a mask over every pixel in an image, multiplying the values in the corresponding places in the image by these mask values and adding them up is an important mathematical operation called *convolution*. It behaves in a manner similar to multiplication and hence is denoted by the $*$ symbol. Figure 5.2 illustrates how we use this symbol.

If one application of the kernel filter isn't enough, it can be applied a second, a third, or any number of times until the desired level of smoothing is achieved. The number of smoothings to apply is also a parameter in the unsharp masking algorithm. So, in the end, there are just two parameters: (a) how many times to apply the smoother and (b) what scale factor to use when subtracting the scaled smoothed image from the original. All modern image processing software packages have built-in commands to make unsharp masking an easy task. The user simply selects values for these two parameters (or accepts the defaults) and pushes a "go" button. The result appears almost instantaneously and is usually much improved relative to the original.

2. Gamma Stretching

The unsharped mask image in Figure 5.1 is a big improvement over the original but it suffers from low contrast within the rings and the disk of the planet. This is because these parts of the image all have approximately the same bright values. Yet, there is added detail in these parts of the image that we'd like to bring out. This can be done using a nonlinear stretching called *gamma stretching*.

As we mentioned at the very beginning of this chapter, all images when displayed on a computer screen or on paper are stretched in some way. The simplest sort of stretching is to specify two levels, a minimum level below which pixels should appear black and a maximum level below which pixels should appear white, and to then shade pixels with values between these two extremes using a gray level that is directly related to how far it is from the black threshold to the white threshold. This is called a *linear scaling*. Using 0 to represent black on the computer screen and 255 for white, the linear stretch used to display the unsharp mask image in Figure 5.1 is graphically illustrated in Figure 5.3.

To enhance the unsharp-masked Saturn image, we'd like to make a stretch where the input values from say about 6000 to 7200 take up a larger share of the output values. This can be done by pulling down on the sloped line in Figure 5.3 to get the curve shown in Figure 5.4. The stretching function shown in the figure corresponds to gamma (γ) value of 2. Larger values have the effect of pulling the curve down even more and hence adding contrast to the brightest parts of the image. Values of γ between zero and one actually push the curve above the original straight line and have the effect of adding contrast to dim parts of the image. Such values are often useful when trying to bring out the details in faint galaxies.

Ideally, one might like to be able to draw any nonlinear stretching curve to customize the appearance of a particular image. Most software packages give the user such a tool but using it effectively can be tedious if not downright difficult. The nice thing about gamma stretching is that there is just one parameter that needs to be chosen and then the software does the rest. Usually, one can find a good value by trial-and-error in only a few tries.

The result of a $\gamma = 2$ stretch on the unsharp-masked Saturn image is shown in Figure the center image in 5.5. The contrast is improved but now there is one more thing to worry about. The video camera used to take these images uses an "interlaced" CCD chip. This means that the odd and even lines of the chip are read off separately. As a




FIGURE 5.3. A linear stretch. Input values are shown along the horizontal axis and output values are shown on the vertical axis. Starting with an input value locate that value on the horizontal axis, move vertically up to the sloped line, and then move horizontally over to the vertical axis to read off the output value. For example, the input value of 4000 maps to the output value of 110.

consequence the precise exposure time for the two sets of lines might not match perfectly. This has now become visible in our image. The effect was always there but we have restretched the image in such a way that it is now obvious and needs attention. The simplest way to fix this is to apply the Gaussian kernel smoother one last time. The final image is shown on the right in Figure 5.5. It now seems quite acceptable, given the limitations of the telescope and the atmosphere we had to image through.

3. Digital Development

As discussed earlier, images that cover a large range of brightness levels are difficult to display effectively in a print image or on a computer screen. The problem is that white is just white when we'd like it to be dazzlingly brilliant, squint-your-eyes, white. Perhaps you've seen old art deco paintings of a night scene in which various street lights and


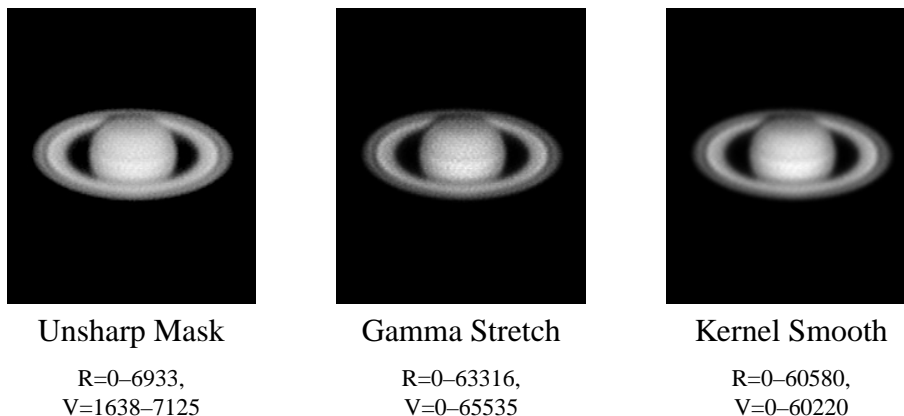


FIGURE 5.4. A gamma stretch. Note that the higher input values use up a larger share of the output range than the lower values do. Such a nonlinear stretch will bring out details that are washed out in the original image. In this stretch, we've chosen a range of 0–65535 to preserve as much of the data for further processing.



Unsharp Mask

R=0–6933,
V=1638–7125

Gamma Stretch

R=0–63316,
V=0–65535

Kernel Smooth

R=0–60580,
V=0–60220

FIGURE 5.5. Gamma stretching and Gaussian kernel smoothing.

stars have little light bulbs embedded in the painting to make these objects appropriately bright. Artists have known for a long time that white in a painting isn't bright enough. Perhaps someday computer screens

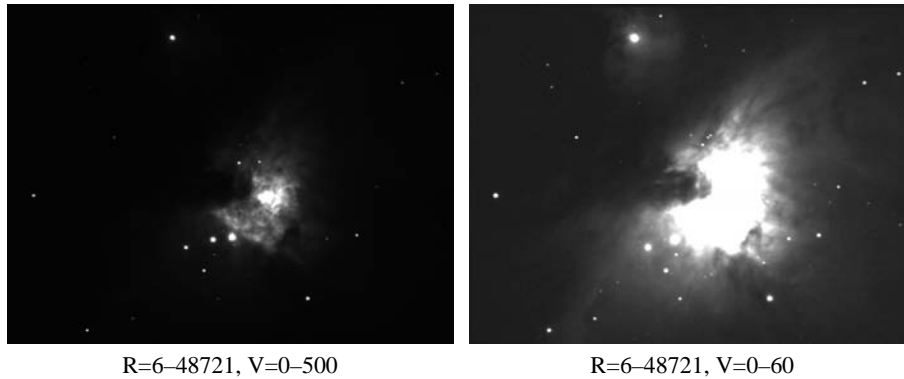


FIGURE 5.6. Calibrated image of M42 stretched two ways.

and maybe even the pages of books will be able to shine brightly at the reader. In this case, we would not need to pursue special image processing techniques to bring out the entire range of brightnesses. In the meantime, we must be clever and figure out how to display images nicely in print.

In this section and the next we confront this issue of how to display images having a huge range of brightness levels. The best example of an astronomical object with a large range is M42, the great nebula in Orion. Figure 5.6 shows a calibrated image of M42 stretched two ways. The raw image has pixels with as few counts as 6 and as many as 48721. If we were to display the image using a linear stretch over this entire range, all we would see are the four bright central stars, called the *Trapezium*, and a few other bright stars. The nebulosity would be completely lost. The two images shown in the figure are stretched 0–500 and 0–60. Even here we see each of these stretchings bring out interesting details that are lost in the other.

How can we restretch the image so that all the interesting detail can be displayed at once? The first thought is to use a gamma stretch with a gamma value less than 1. Figure 5.7 shows the result obtain using $\gamma = 0.2$. It is much improved, but the interesting details in the bright region, nearest the trapezium, have become a little flat because they must be depicted using gray-scale values over a reduced range. Essentially, each of the two images in Figure 5.6 had values 0–255 in which to display the result, but the stretched image gets about half the range, say 0–127 for the dim part of the image and 128–255 for the bright part.

The problem here is that given any two pixels in the image, say A and B, if pixel A is brighter than pixel B in the original image, then it



FIGURE 5.7. Applying a Gamma stretch ($\gamma = 0.2$).

must still be brighter in the gamma stretched image. And, this is true for any value of gamma. This preservation of relative brightness over the entire image puts a very strong constraint on what can be done with the raw image. In fact, this property must be preserved locally but it need not be preserved throughout the entire image. For example, it would be better to have a range of say 0–191 for the dim parts of the image and a range of 64–255 for the bright parts. Of course, these ranges have a significant overlap which could make the final image look “wrong” but if we make a smooth transition from the dim regions to the bright regions, then the final result will still look right.

It might sound difficult to find a technique that will have this property of preserving relative brightness only locally, but it turns out that it is quite simple. The basic method is similar to unsharp masking but instead of subtracting a blurred version of the image, we divide by it. Well, that’s almost it. We just need to make two small changes.

The first change is that we need to add some positive constant, let’s call it b , to the smoothed image before dividing. Without this constant, every part of the image, both the bright parts and the dim parts, would come out with the same general brightness level. This is not what we want—it goes too far. We still want the dim parts to appear somewhat dimmer overall than the bright parts. Adding a constant reduces the effect of the division and effectively gives us an overlap of ranges without it being complete. That is, we get something like 0–191 for the dim parts and 64–255 for the bright parts. Of course, the exact amount of the overlap depends on the constant b . The software packages that implement this method usually suggest a good value for b based on the image data but at the same time allow the user to change it to a higher or lower value as desired.




FIGURE 5.8. Summary of the Okano's digital development process. In this example, we used $a = 1000$ and $b = 40$.

The second change addresses the fact that simple division of two similar images will result in numbers that are fractions averaging about one. Some fractions will be smaller than one and others larger than one but the average value will be about one. We should multiply all the values by some positive scale factor, call it a , to make the average some reasonably large value. Some programs save images in such a way that fractional pixel values are saved as fractions. But most don't do this. Instead they round each fraction to the nearest whole number. If rounding is going to happen, it is crucial to scale up the fractions.

So, that is the method. It was invented by Kunihiko Okano. He called it *digital development processing* or DDP since it imitates the behavior of the photographic development process.

Figure 5.8 shows the result of this process on M42. As the method works best when the background pixels are close to zero, we first subtracted an appropriate background level from the raw image of M42. Then, we applied the Gaussian smoother to make a blurred copy of the image. Next, we added a constant (40) to the blurred image. Finally, we divided the original image by the blurred image and scaled the result by 1000 to make the final numbers span a reasonable range of values.




FIGURE 5.9. Digital development is very effective on globular clusters. Shown here are before and after images of M5.

The resulting image looks much better but it has unnatural dark halos around the bright stars. These halos can be minimized by a careful choice of the smoother one uses to make the blurred image—the Gaussian smoother we used was not the best choice. Implementations of DDP in commercial image processing software tend to do a better job picking an appropriate smoother but the halos are still often evident in images having very bright stars embedded in background nebulosity. In fact, DDP tends to work best on globular clusters (where there is no background nebulosity) and on galaxies (where embedded stars usually aren't too bright). Star birth regions, such as the Orion nebula pose the biggest challenge for DDP. In the next section, we'll give a minor variation on the method that works well even in the presence of embedded bright stars.

There is a second benefit to using DDP. Not only does it accommodate wide ranges of brightnesses, it also sharpens the image much in the same way that unsharp mask does. (In fact, DDP is roughly equivalent to applying a logarithmic stretch to the original image, unsharp masking that image, and then undoing the logarithmic stretch.) The sharpening effect is quite strong and is the main reason that DDP is very effective on globular clusters. As an example, Figure 5.9 shows how it does on globular cluster M5.

4. Variants of Digital Development

For those cases where DDP creates dark rings around stars, there is a simple modification to the procedure that avoids the dark rings. The idea is to replace the Gaussian smoother with a smoother that replaces each pixel value with the minimum of the values in a small rectangle




FIGURE 5.10. Two variants of DDP. On the left, the minimum value in a 7×7 rectangle around each pixel was used. On the right, the median value was used.

centered on the pixel. Figure 5.10 shows the result obtained when using a 7×7 rectangle around each point. This method loses the sharpening aspect of DDP but retains the rescaling of brightnesses so that the final image shows details in both the bright and the dark parts of the original image.

One last alternative is to use the median value in a rectangle centered on each pixel rather than the minimum. The *median* of a set of numbers is that number for which half of the remaining numbers are larger than it and the other half are smaller. Since the average pixel value in a neighborhood of a star tends to be skewed toward the maximum value, we find that the median gives a value between the minimum and the average. Hence, using the median gives a result that is somewhat in between DDP with average and DDP with minimum. It stretches brightness nicely and it sharpens the image somewhat without making dark rings around the bright stars. The result for M42 is shown on the right in Figure 5.10.

5. Deconvolution

As we saw earlier, convolving an image with a Gaussian kernel produces a blurred image. Convolution is a fundamental concept. In fact, it is so fundamental that one can view the blurry imperfections in a raw image as a convolution of a perfect image with a smoothing kernel representing the blurring caused by the diffraction properties of light combined with other sources of blur such as atmospheric turbulence, inaccurate tracking/guiding, and even bad focusing. The kernel is called the *point spread function* or *psf*. It is easy to get an estimate of the psf—simply look at the image of any star. A perfect image of a star would

be just a point of light. The fact that it is spread out in all real images is a consequence of all the blurring factors mentioned above. The star's image, after normalizing so that the total sum is one, is precisely the psf.

If we assume that the psf is known, it is natural to ask whether the convolution process can be undone to recover the underlying perfect image. Undoing convolution is called *deconvolution*. If the blur is exactly the result of convolution with a known psf, then deconvolution will recover the perfect, unblurred image. However, in practice, the psf is never known exactly and in addition to blur in the image there is also graininess resulting from the discrete nature of counting photons. The graininess gets amplified when deconvolving. Hence, it is undesirable to completely deconvolve an image. Instead, one looks for iterative procedures that slowly move from the original blurred image to the completely deconvolved image. Then, one can stop the process after just a few iterations—enough to improve the image but not so many as to amplify the graininess and other imperfections to the point of annoyance. There are three popular iterative procedures: *Richardson–Lucy deconvolution*, *VanCittert deconvolution*, and *maximum entropy deconvolution*.

We will only describe in detail the first method, Richardson–Lucy deconvolution, as this is the particular procedure that achieved wide acclaim from its success in restoring the flawed images from the Hubble Space Telescope in its early years when it had flawed optics. The first step of the process is similar to but a little bit more complicated than DDP. The first step proceeds as follows.

- (1) Make three copies of the original image; call them *numerator*, *denominator*, and *result*.
- (2) Apply the psf-based kernel smoother to *denominator*.
- (3) Divide each pixel in *numerator* by the corresponding pixel value in the blurred image *denominator*.
- (4) Apply the psf-based kernel smoother to the divided image *numerator*.
- (5) Multiply each pixel in *result* by the smoothed ratio image in *numerator*.
- (6) Now, *result* contains the result of the first iteration of the Richardson–Lucy deconvolution.

This first iteration is summarized in Figure 5.11.

Subsequent iterations are similar except that the image from the current iteration is used for the left-hand image and for the denominator image. For example, the second iteration is as shown in Figure 5.12.

As with the earlier sharpening techniques, there is a danger in overdoing it. Dark rings will start to appear around stars, areas of medium




FIGURE 5.11. Richardson–Lucy deconvolution. First iteration. Computed using the Gaussian smoother.

brightness will appear to have too much contrast, and background areas will become noisier. There are various techniques to limit these bad behaviors. The simplest is to subtract as much of the background as possible before applying the technique. If any real data was taken away with the subtraction, it is always possible, even desirable, to add the background back in at the end.

6. Correcting for Coma/Field-Curvature

Figure 5.13 shows an image of the Veil nebula. Note that the stars around the edges are not quite round. The football-shaped stars around the periphery result from an imperfect optical system. No telescope is perfect and the further one ventures from the center of the field the more the imperfections become evident. The image in the figure was




FIGURE 5.12. Richardson–Lucy deconvolution. Second iteration. Computed using the Gaussian smoother.

taken with a focal reducer in order to get the largest possible field of view. But the price one pays is to see these imperfections. In fact, a focal reducer that is not well matched to the optical system can, and often does, contribute to the problem. In this section, we show how one can address these imperfections using appropriate image-processing—it's cheaper than buying a new telescope.

Figure 5.14 shows the result of Richardson–Lucy deconvolution applied to the image from Figure 5.13. Note that the nebulosity is much improved and that the central stars have been sharpened nicely, but at the same time the stars in the corners have gotten much worse. The reason is that deconvolving an oblong star using a round psf produces not a pin-point star but rather a little line segment. This is not good. But, it



FIGURE 5.13. A calibrated image of the Veil nebula. Note the distorted star shapes around the edges and especially in the corners. This distortion results from an imperfect optical system. It becomes most evident when using a focal reducer.



FIGURE 5.14. The image in Figure 5.13 after applying Richardson–Lucy deconvolution. Note the very distorted stars in the corners.

is relatively easy to fix. What we need is to deconvolve with a psf that varies from place-to-place across the image.

None of the image processing software packages available to amateurs provide the capability to deconvolve with a psf that varies across



FIGURE 5.15. An image showing how we varied the psf over the image for the purposes of deconvolution. Note that these psf's match fairly well with the stars in the original image.

the image. So, I wrote my own. Figure 5.15 shows a model of the psf for each of 100 stars spread out over the field of view. Note that the stars shapes approximately match what we saw in the Veil nebula (Figure 5.13).

Using a psf model like the one shown in Figure 5.15 in a Richardson–Lucy deconvolution code, I was able to recover the deconvolved image shown in Figure 5.16. I think it is safe to say that the final image is much improved. As a final remark, we note that deconvolution only works well on long-exposure images (i.e., images with very little graininess).



FIGURE 5.16. The Veil nebula deconvolved using a psf that varies from place to place according to the model shown in Figure 5.15.

CHAPTER 6

Stories

1. Jupiter

I was interested in astronomy in junior high and high school. But at that time, I had no money so my interest was mostly restricted to books and visits to the local planetarium. After high-school, I devoted my energies to other activities and my interest in astronomy layed dormant for many years. Then in October of 1998, a friend of mine, Kirk Alexander, invited me to an evening of stargazing at the annual Stella-Della star party. This was my first night ever of real observing. Kirk set up his telescope and pointed it at Jupiter, which was just rising in the east. We then walked around the large field and looked at various objects through a variety of telescopes. At times during the evening we would return to his telescope for another peek at Jupiter. At one of those times, I noticed a round black spot on the face of Jupiter. I asked Kirk if maybe there might be some dirt on the eyepiece. He took a look and immediately identified the black spot as a shadow transit of one of Jupiter's moons. This was really exciting to me—to think that we were looking at a total solar eclipse taking place on another planet. Over the next hour or more we watched as the region of totality, the black spot, moved slowly across the face of the planet. I decided then and there that I would buy myself a telescope.

Just like so many before me, I made my telescope decision too hastily. Positively influenced by what I saw through Kirk's 7" Ques-tar I decided to buy the smaller 3.5" version. I also subscribed to Sky and Telescope magazine. For the next few years, every month I would check the page in Sky and Telescope that lists times of shadow transits. Those times that seemed like reasonable observing times for someone in New Jersey were then entered into my computer's calendar program so that I would be automatically reminded of the event a few hours ahead of time. I must have set up my new telescope dozens of times over a few years with the sole purpose of observing a shadow transit. I was never successful. To this day I don't know why. Was I converting from Universal Time to Eastern Standard Time incorrectly? Was my 3.5"

Questar just not up to the task? Was I using an inappropriate eyepiece? I don't know. Eventually, I gave up.

Time passed while I concentrated on other observations—double stars, the other planets, the moon, star clusters, etc. Then in July of 2001 I bought my CCD camera. This changed everything. From the day I received my camera onward I have done very little visual observing. So, in the fall of 2001 when Jupiter again became a nighttime object, I decided to try my hand at high-magnification planetary photography.

This type of imaging presents a new set of challenges. First of all, the field of view is very small and so it is hard to get the object, Jupiter, into view. Questar's built-in flip-mirror system helps a great deal but still some practice is required before one can quickly and reliably place Jupiter on the camera's CCD chip.

Furthermore, my particular camera presents certain challenges. The CCD chip is a Sony chip of an interlaced design. This means, by its fundamental design, the even lines must be read out separately from the odd lines. The download time is a few seconds and during those few seconds Jupiter can move quite a lot due to periodic tracking errors in the mount's clock drive. So, you get images of Jupiter in which the odd lines show the image shift a little bit left or right from the even lines. This is not good. Fortunately, the software that runs the camera has the option of making an image out of just one set of lines (the even or the odd, not sure which). This solves the misalignment problem but the resulting image appears squashed. Squashed images are only a minor problem since it is easy to resize them during image processing.

Another issue is that, at high magnification, even the tiniest speck of dust on the window covering the CCD chip casts a noticeable shadow onto the chip which then appears in the image. So, taking good flat-fields for calibration is very important. Also, cleaning the window prior to imaging is a good idea. The fact that the image moves from one exposure to the next, due to tracking error, also helps because the dust shadows do not move and so when the images are aligned and stacked the dust shadows get smeared out and largely disappear. Anyway, dust shadows are an issue at high f-ratio, i.e., high magnification, that must be addressed.

So, on my first night of attempting to image Jupiter, I took a series of short, 0.3 second, squashed exposures. I took some flat fields and then packed everything away and went to bed. The next day, I processed the images. I went through the usual calibration process. One frame is shown at the top in Figure 6.1. I then aligned and stacked the images and applied some deconvolution to sharpen the image a bit. All of a sudden I noticed something—a round, black spot on the face of Jupiter. My first




FIGURE 6.1. First Jupiter Images. *Top Left.* Unprocessed raw 0.3-second image. *Top Right.* Same image stretched to show moons. *Bottom.* Stack of 10 raw images deconvolved. Note the moon shadow.

thought was: is this an artifact of the deconvolution or maybe a dust shadow that is overemphasized by the deconvolution process? But, it didn't seem like a dust shadow. It was too small and too perfectly round looking. Then it dawned on me: it must be one of those long-sought but never-found shadow transits. I hadn't been thinking about shadow transits for a long time. I certainly hadn't checked *Sky and Telescope* prior to this imaging session. My only goal was to try to take a picture of Jupiter which would show its bands. But there it was—a black spot. I quickly went to check in *Sky and Telescope* and sure enough I had indeed caught a shadow transit—in my very first Jupiter image without even trying. The final image, also shown in Figure 6.1, is not one of my best Jupiter images, but the shadow transit is unmistakable.

2. Starquest 2002

Every June, the Amateur Astronomers Association of Princeton (AAAP) hosts a two-night star party at a fairly dark site in northwestern New Jersey. I first attended this weekend affair in June of 1999. It was great fun to take my recently bought Questar to its first star party. Having myself never been to Starquest, I didn't know exactly what to expect. When I got there I learned that AAAP member Bill Murray has a tradition of running an observers challenge. It is different every year but generally consists of a list of about 50 interesting things to observe. Attendees are given the list—actually there are two lists, a beginners list and an advanced list—and are instructed to record the time that they observe the various items on the list. At the Sunday morning breakfast, those observers who have logged 25 of the items are awarded a “AAAP Observers” pin. I attempted the beginner's list and had fun observing some things that I'd never before seen through my Questar—or any telescope for that matter. The most memorable discoveries were the Double Cluster and the Wild Duck Cluster. It also became apparent just how faint galaxies appear in a 3.5" f/15 instrument. On Sunday, I got my first observer's pin.

I had such a good time at my first Starquest that I decided to make it a tradition of my own to attend this party every year. June 7–9, 2002, was my first Starquest party after getting into CCD imaging. I brought my telescope and imaging equipment, i.e., camera and computer, to the party not really being sure whether I'd devote myself to visual or CCD imaging. I was pleasantly surprised to learn that Bill had selected a list of about 50 globular clusters for the observer challenge. This was perfect for me since globular clusters are objects that can be imaged nicely with fairly short exposures. The rules were that an observer must log 25 of the globulars and that no more than 15 of them can be Messier objects. I figured that it might actually be possible to image 25 globular clusters in two nights, so this is what I set out to do. My plan was to take 25 six-second unguided exposures of each globular that I could find. That's 2.5 minutes of exposure time. Since each image takes 6 seconds to download, it would take 5 minutes to do one imaging sequence. I allowed myself 5 minutes to find each target and another 5 minutes to “compose” the shot. If I could keep to this schedule, I could do four clusters per hour.

On Friday night I managed to image 16 of the *easier* clusters. On Saturday night, I imaged 11 more clusters and revisited one cluster, NGC 5053 from the previous night. Sunday morning I discovered that my image of NGC 6517 didn't actually contain the cluster. Looking




FIGURE 6.2. Images taken at Starquest'02

at star charts, I discovered that I'd missed the cluster by only 5 arcminutes. Also, due to inaccurate counting in the middle of the night, I ended up short on non-Messier objects. Even if I had correctly imaged NGC 6517, that would have only been number nine, not number ten like I had thought. My final tally was 18 Messier clusters and 8 non-Messiers. Nonetheless, Bill awarded me with a pin for the effort. Thumbnail images of the 26 clusters are shown in Figure 6.2.

Bibliography