AST 598 'Instrumentation and Data Reduction' course, ASU, Fall 2006

Astronomy with Charged Coupled Devices

R.A. Jansen

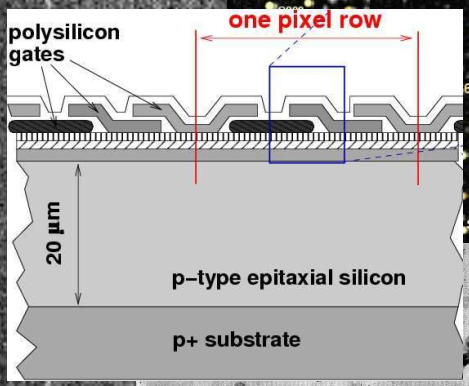

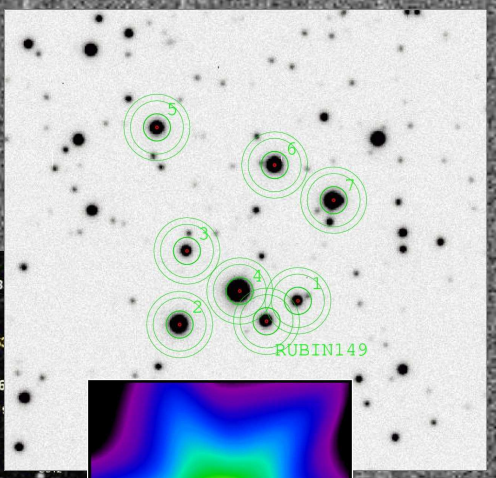

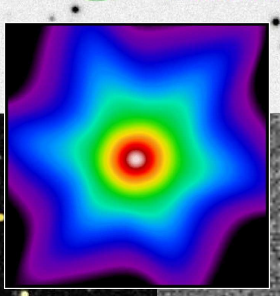

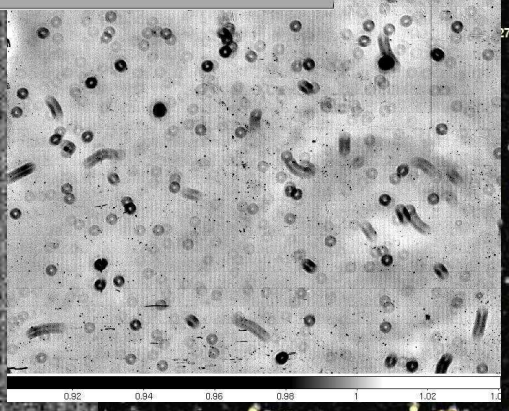

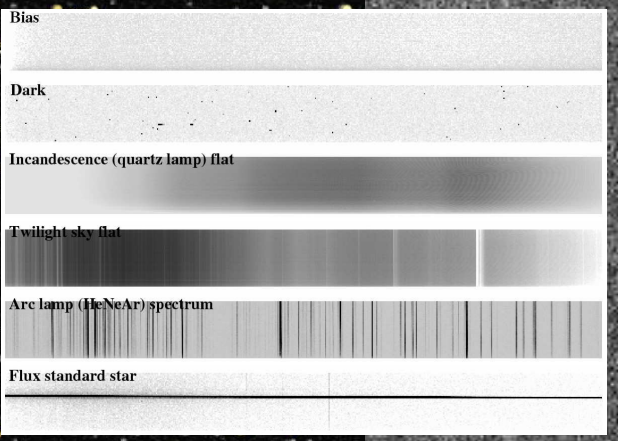

Contents

1 Introduction to CCD's

1.1 A brief history of astronomical detectors

1.1.1 The Human Eye

Throughout most of history, only one type of detector has been available to us for astronomical observation: the *human eye*. The hunan eye is a very versatile detector, but —despite its strengths— has several limitations as well. Since in modern astronomy we still encounter some concepts that are intimately tied to the earlier use of this type of detector (most notably the continued use of *magnitude*s), let us discuss some of the properties, pros and cons, of the human eye. A schematic drawing of the human eye is shown in Fig. 1.

- The amount of light that can be collected by a detector depends on the aperture through which light enters. For a well-accommodated eye, the entrance aperture (i.e., the *pupil*) diameter can be as large as $\sim$8 mm ($\sim\frac{1}{3}$ inch). This limitation can be overcome, however, by using a telescope with a larger entrance aperture to focus more light through the pupil.

- Although the most sensitive cells within the retina, the *rod cells*, have a finite probablity of responding to the absorption of even a single photon, the overall probability of absorption and response is only $\sim$1–2%.

$\triangleright$ The ratio of incoming to actually detected photons is denoted by the term *quantum efficiency (QE)*, and is usually expressed as a percentage.

- Another limitation is that the human eye has a fixed response time (taken here to be equivalent with an *integration time*) of roughly $\frac{1}{10}$ second (rod cells), followed by a recuperation period (*dead time*) of similar length, during which a given cell cannot respond to further incoming photons.

- The eye plus brain is a *logarithmic* detector: its response is not proportional to the incoming photon flux, but to the logarithm thereof.

- The range in wavelengths to which the human eye is sensitive is small. Rod cells, responsible for night vision, are only sensitive to wavelengths shorter than $\sim$570 nm. The short wavelength cut-off depends on age, but is generally longward of $\sim$370 nm (see Fig. 2).

$\triangleright$ The wavelength range over which the quantum efficiency of a detector exceeds some fiducial fraction of the peak QE value is denoted by the term *band-pass*. The fiducial fraction is commonly taken to be 10%.

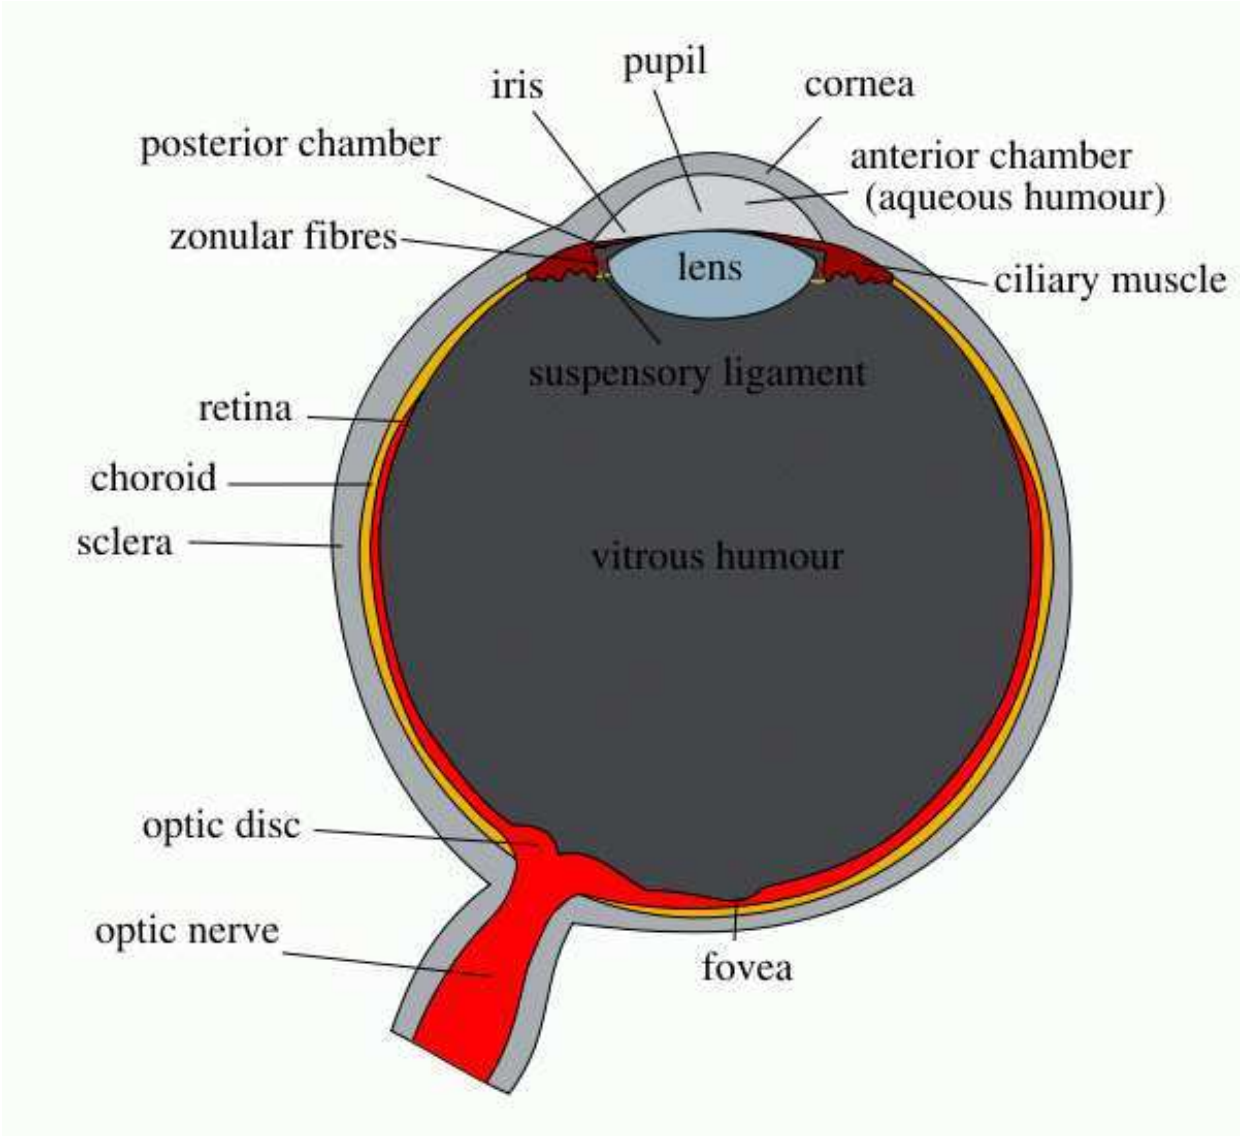

Figure 1: (a) Schematic diagram of the human eye. In a well-accomodated human eye, the diameter of the pupil (opening aperture) can reach ∼8 mm. The optic disk is better known as the *blind spot*, where the optic nerve enters the eye. Averted vision will optimize our use of the more sensitive rod cells, which are located in the retina mostly outside the fovea. The less sensitive cone cells, responsible for our perception of color, are strongly concentrated toward and densely packed within the fovea. This dense packing of photoreceptive cells also gives us our visual acuity (sharpness of vision). *Courtesy: Wikipedia*.

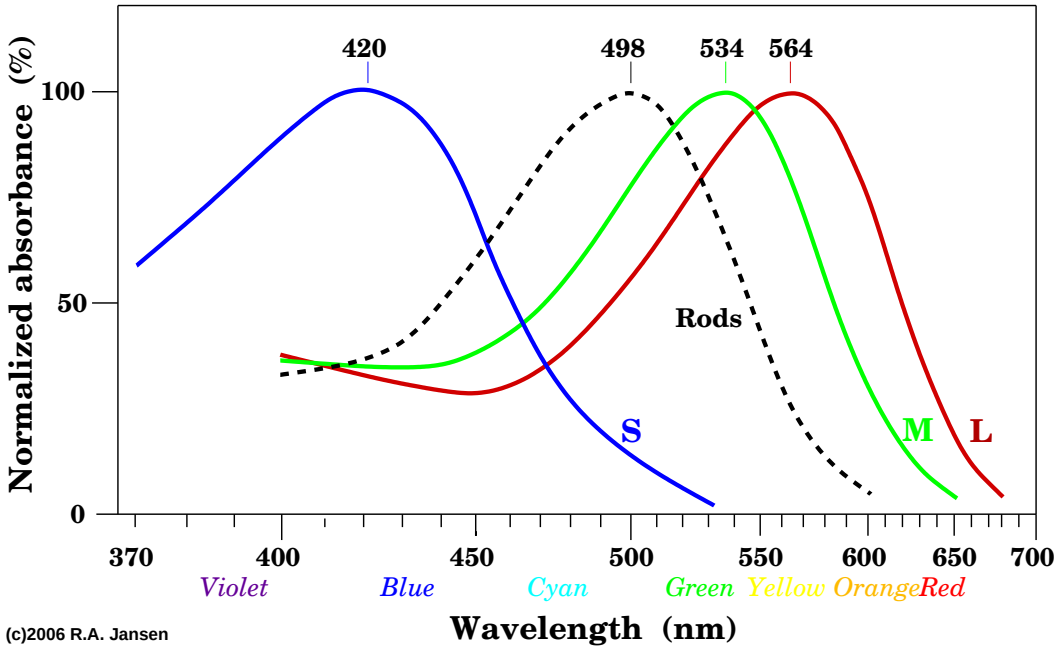

Figure 2: Normalized spectral absorption curves of the pigment in human rod cells (responsible for night vision) and of the short (S), medium (M) and long (L) wavelength pigments in cone cells. The sensitivity of the rod cells is largest in the blue-green portion of the optical spectrum. The absolute sensitivity of the rod cells is ∼100 times higher than that of the cone cells. *After Bowmaker & Dartnall 1980.*

- Thanks to the two different types of photo-receptive cells and their "inverse" physiology (an *increase* in the light level results in a *decrease* of the amount of neurotransmitters released, and high light levels cause the photo-destruction of the absorbing pigment), and thanks to an adjustable entrance aperture, the eye can function well over a very large range in brightness. Simultaneous detection of both bright and faint light sources can be a problem (e.g., think of the loss of night vision experienced due to head-lights of oncoming cars at night: it takes time for the cell to produce new absorbing pigment).

Magnitudes

The ancient Greeks (notably Hipparchus) made an inventory of stars in the sky, where the 20 brightest stars were defined to be of magnitude 1. Fainter stars were divided into 5 bins according to their perceived visual intensity, with roughly equal steps in perceived brightness between the first and subsequent bins. In more recent times, in 1856, Pogson derived his famous quantitative expression to approximate the classical magnitude scale, under the assumption that the human eye behaves as a logarithmic detector:

$$M = -2.5 \cdot \log F + C \quad ,$$

where F is the flux arriving from a star and C is an arbitrary zeropoint. The scale factor -2.5 is called the *Pogson scale*. The magnitude *difference* between two stars is then related to the *ratio* of their fluxes by:

$$M_1 - M_2 = -2.5 \cdot \log(F_1/F_2).$$

Note, that a flux ratio of 100 corresponds exactly to a difference of 5 magnitudes. To this day, astronomy students drive physics students nuts by putting seemingly arbitrary factors of -2.5 before any decimal logarithm they encounter.

1.1.2 The Photographic Plate

During the second half of the 19$^{\rm th}$ century, the then new technology of photographic emulsions found its first astronomical uses. The earliest known photograph of an astronomical object, the Moon, was made by J.W. Draper in 1840, followed five years later by Foucault and Fizeau, who photographed the Sun. By 1870, technology had improved enough to allow Jansen & Lockyer to discover the element Helium on a photograph of a solar spectrum.

With the advent of photography, it became possible to overcome the fixed exposure time of the human eye, and to faithfully record detected photons for later analysis. It also became possible to record photons at wavelengths for which the human eye is insensitive, particularly toward shorter wavelengths (near-UV; but also X-rays and ionizing particles). Photographic emulsions sensitive to longer wavelengths (near-IR) were only developed later in the 20$^{\rm th}$ century.

The glass plates used as substrates for the emulsions were resistent to bending and distortion during the observations, and have proven *extremely* stable over time (a full century and counting). Although currently available in digitized format, the two Palomar Observatory Sky Surveys (POSS in the 1950s, and POSS-II during the 1990s) and the complementary UK Schmidt surveys in the southern hemisphere were made using photographic plates.

- The quantum efficiency of normal photographic plates is low, typically only 1–2% ($\lesssim$3% for IIIaJ). Via a black magic process of *hypersensitization*, involving exposure to different gasses at various temperatures, this could be increased for some purposes to $\sim$10%, but at the cost of a reduced *dynamic range* (even in the absence of any photons, such plates may show a distinct "fog").

▷ The ratio of the maximum to the minimum measurable signal is denoted by the term *dynamic range*.

- Photographic plates are non-linear/logarithmic in their response to light. Over a limited range, between under- and overexposure, the response of photographic plates is close to logarithmic (see Fig 3). Outside this range, the response is highly non-linear. The effective dynamic range is typically only $\sim 10^2$.

  Given the very large range in brightness of astronomical sources that might be of interest within a given field of view on the sky, and even within individual extended objects such as galaxies and nebulae, often a sequence of different exposure times is required. For long exposures, *reciprocity failure*, the decrease of the sensitivity as the exposure continues, could be an issue.

- Quantitative analysis of photographic plates requires one to first scan and digitize those plates. This digitization process is far from trivial and not always perfectly repeatable.

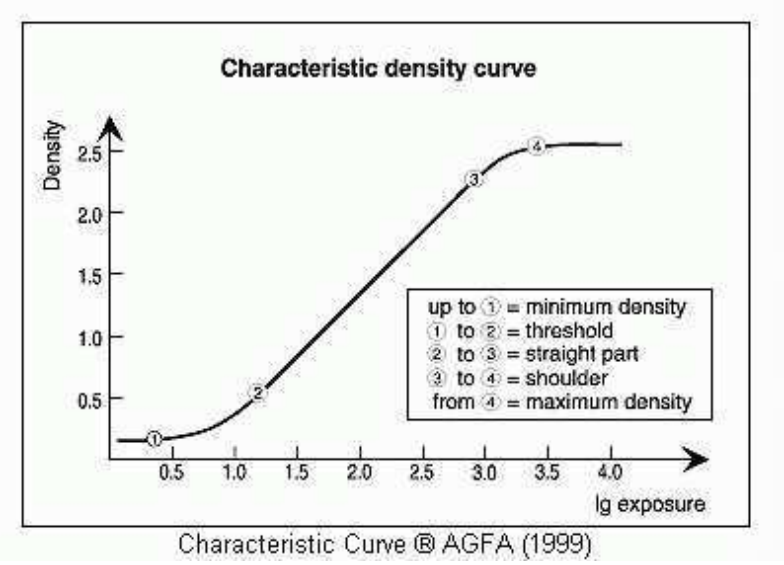

Figure 3: Schematic representation of a characteristic density curve of photographic emulsions. Between points (2) and (3), the response to light is very close to logarithmic.

- Photographic plates were the first detectors to record detected photons for later detailed study and their long-term stability is excellent. We can still compare present-day observations with plate material from 1900.
- Photographic plates have an extremely large area — unmatched to this day by any other type of 2-D detector, and well-matched to the physical size of the collimated beam of a large telescope ($\gtrsim$10 inch). In combination with Schmidt telescopes, typical areal coverage was $6^\circ \times 6^\circ$!
- The resolution of a photographic plate, defined by the grain size of the silver halide crystals (a few μm), is very high. A photographic plate has a higher resolution than a CCD chip of the same area.

The sensitivity of early photographic plates to blue light results to this day in the preponderance of blue photometric and spectroscopic standard stars.

Silverhalide photography

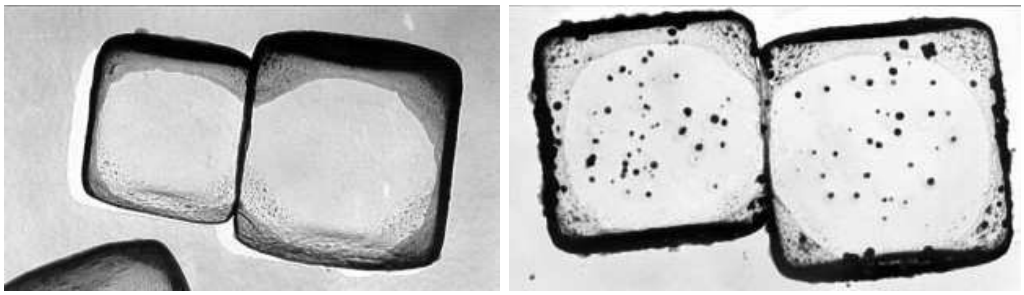

Figure 4: (a) unexposed silverbromide crystals, (b) exposed silverbromide crystals, where the metallic silver deposits are clearly seen (even after only soft development). *Courtesy: Frank Lakiere.*

Photochemical reaction equation for silver bromide:

$$\mathrm{Ag^+Br^-(crystal)} + h\nu\,\mathrm{(light)} \longrightarrow \mathrm{Br} + e^- + \mathrm{Ag^+} \longrightarrow \mathrm{Br} + \mathrm{Ag^0}$$

The free silver atom produced within the silver halide grain serves as a catalyst during subsequent development, where grains containing free silver atom(s) will form relatively large deposits of free silver (opaque), while those that do not contain free silver (grains that did not absorb photons) will remain clear.

1.1.3 The Photomultiplier Tube

A major advance for low light-level applications was the development in the first half of the 20^{th} century of vacuum-tube technology, including Photomultiplier tubes (PMTs). These detectors rely on the *photo-electric effect* and can achieve peak quantum efficiencies of upward of 20%. They are usable over a very wide range in wavelength. This type of detector *amplifies* the exceedingly faint signal of a single absorbed photon (i.e., the production of a photo-electron) by causing a cascade of secondary electrons. Amplification can be as high as a factor 10^8 (typically $\gtrsim 10^6$).

A PMT works as follows (see Fig. 5). When an incident photon strikes the thin photocathode, made of a semiconducting material such as GaAs (galliumarsenide) deposited on the inside of the entrance window, a photo-electron is dislodged. That photo-electron has some residual kinetic energy allowing it to escape from the cathode. A strong electric field (high voltage) accelerates this electron towards the first of a series of electrodes, called *dynode*s, where it arrives with sufficient energy to release a second electron. Both electrons are then accelerated toward a second dynode, each again releasing additional electrons. Each subsequent dynode in the stack is held at a more positive voltage than the previous one. The geometry of the dynode chain is such that a cascade occurs with an ever-increasing number of electrons being produced at each stage. Finally, at the *anode*, the accumulation of charge results in a sharp current pulse, the height of which is proportional to the number of original photo-electrons (and therefore photons).

- Their peak quantum efficiency is much higher than that of photographic plates and can reach $\sim$40%.
- They can be operated as virtually noise-free, highly linear and stable photon counters, allowing accurate photometry and spectrophotometry to accuracies of $\sim$1%, and hence quantitative astrophysics. They also allow high-speed photometry (few μsec).
- Although their *dynamic range* is limited by the build up of charge near the anode, which repulses arriving electrons, their dynamic range tends to be very large: typically $\sim 10^5$.
- Their opening aperture is physically large (from several cm^2 to over 50 cm diameter [the Cerenkov radiation detectors within Super-Kamiokande]) allowing industrial and medical applications (e.g., counting blood cells flowing past at high speed).
- To reduce their *dark current* due to thermal electrons being drawn away from the cathode and dynodes, they must be cooled. Typically, they are operated at dry ice (solid CO_2) temperatures ($-78\,^\circ$C).

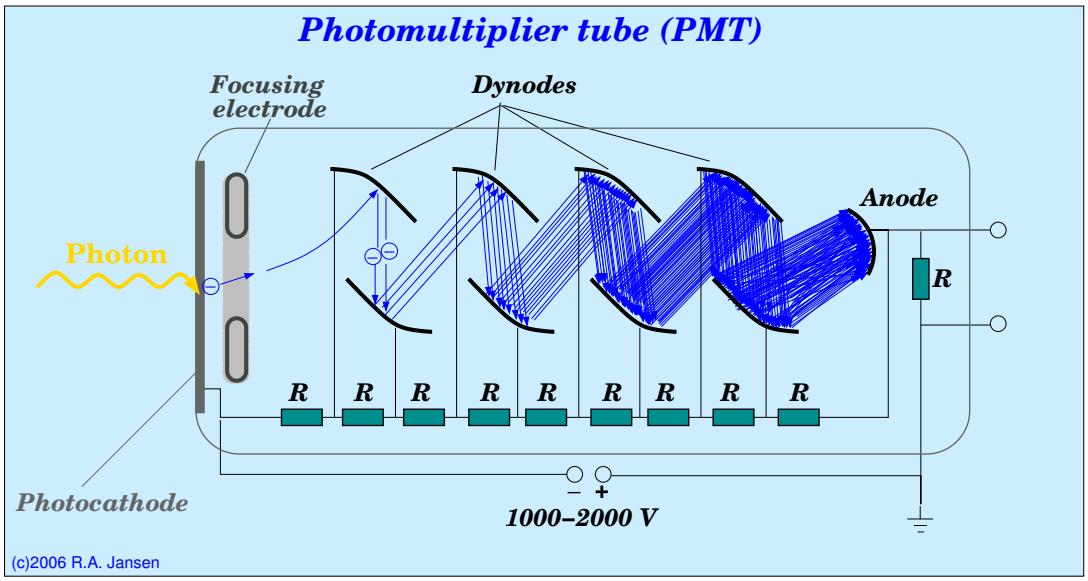

Figure 5: Schematic diagram of a photomultiplier tube. When a photon strikes the thin photocathode on the inside of the entrance window of the glass vacuum tube, a photo-electron is released and accelerated toward the first of a series of dynodes, where it arrives with sufficient energy to dislodge a second electron. The geometry of the dynode chain is such that a cascade occurs with an ever-increasing number of electrons being produced at each stage. The height of the sharp current pulse measured at the *anode* is directly proportional to the number of original photo-electrons (and therefore photons).

- Photomultiplier tubes are high-voltage devices, requiring typically 1000–2000 V for their operation. When in operation, photomultiplier tubes are prone to irreversible damage when exposed to high light levels. They also suffer from electrons dislodged by cosmic rays. Unless shielded, they are sensitive to strong ambient electric and magnetic fields.

- The main disadvantage of photomultiplier tubes, however, is that they are single-element devices and next to impossible to package into useful 1-D and 2-D arrays. They are, therefore, not very efficient for spectroscopy and not useful for astronomical imaging applications (other than in scanning and digitization of photographic plates).

To this day, astronomers and filter engineers are striving to reproduce the response curves, and in particularly the red cut-offs, of a few popular models of photomultiplier tubes, since the sensitivity of these tubes rather than the pieces of colored glass employed defined the effective "filter transmission" curves of the standard Johnson UBV and Kron-Cousins RI filters. (E.g., in the case of Johnson I, the piece of

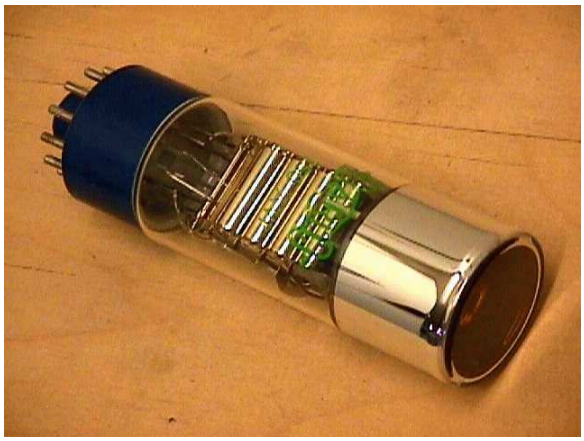

Figure 6: Example of a photomultiplier tube (PMT). *Courtesy: M. Fletcher.*

glass Johnson used was a long-pass filter, defining only the blue portion of the I transmission curve; the red portion reflected the sensitivity of his photomultiplier).

1.1.4 Image Intensified Photography

An attempt to combine the advantages of a photocathode device (photo-electric effect and signal amplification) and the 2-D imaging capability of photographic plates. Although this technology experienced a brief revival with the first baloon-borne and space-based UV imaging telescopes, it should be considered a station that was short-lived — the less said, the better.

1.1.5 Image Intensifier Vacuum Tube technology

A series of high-voltage 1-D and 2-D electronic devices were constructed starting in the 1960s based on image intensified vacuum tubes. Like in the photomultiplier tube, incident photons release photo-electrons which are accelerated and produce either a detectable current pulse or a splash of light when hitting a screen coated with a phosphor (e.g., ZnS:Cu in old-fashioned green computer screens). Unlike photomultiplier tubes, information on where exactly the incident photon arrived is preserved.
Examples: Electron Conduction Vidicon, Intensified Dissector Scanner, Intensified Reticon Scanner, Image Photon Counting System, and Multi-Anode Microchannel Array detector. The latter has experienced a revival in space-based UV astronomy (e.g., HST/STIS and HST/ACS-SBC, GALEX, and FUSE).

- Disadvantages are the high voltages required (e.g, 25 kV), image blur, and non-trivial relation between arrival location and (x, y) position.

1.1.6 The Charge Coupled Device

Charge Coupled Devices (CCDs), were originally developed at Bell labs as a solid-state memory device (Amelio, Tompsett & Smith 1970; Boyle & Smith 1970). A CCD is essentially a regular array of large numbers of coupled Metal-Oxide Semiconductor Field Effect Transistors (MOSFETs). It was immediately realised, however, that the sensitivity to light via the photo-electric effect and the ability to store charge in a regular array format was suitable for application as a 2-D detector. An excellent review paper on the workings of a CCD is MacKay (1986).

When a photon is absorbed in silicon, an electron-hole pair is created. In a CCD, two or more electrodes (called *gates* following MOSFET terminology) per isolated picture element (*pixel*) create potential wells in the silicon. The generated photo-electrons migrate to and are confined within the potential well created by the gate kept at the largest positive Voltage. The holes, on the other hand diffuse away into the bulk of the silicon, thus preventing recombination.

In a three-phase CCD (presently still the most common type in astronomical applications), three gates — each isolated from the silicon substrate and from one another — define a pixel row. Embedded channel stops of heavilly doped p-type Si define pixel columns, and prevent electrons from migrating along the length of the gates due to their negative charge. Fig 7 illustrates the structure of a three-phase CCD, and Fig 8 shows a top view of a portion of an actual device. Every third gate is electrically connected in parallel and kept at the same voltage.

To convert the 2-D array of stored photo-electrons into an image, the charge in each pixel needs to be measured (or more accurately: estimated). To this end the voltages on the electrodes of the array are cycled in such a way that the stored charges are sheparded along the columns, effectively being passed along from one pixel to the next. At the end of a column the charges are passed onto a special set of pixels with a gate structure that is orthogonal to that in the image area of the CCD. Here the charges collected in an entire pixel row are shifted one pixel at a time and passed on to on-chip electronics that amplify and estimate the stored charge. The result is subsequently converted by an Analog-to-Digital circuit into a digital number (refered to as either *Data Number* (DN), or *Analog-to-Digital Unit* (ADU), or simply *Count*).

The most common analogy for the process of reading out a CCD is that of a "Bucket Brigade", where the contents of the buckets on many parallel belts are emptied in a serial belt with one bucket per parallel belt. The buckets on the serial belt are then emptied one at a time into a measuring container. After the last bucket on the serial belt has been measured, the next row of buckets on the parallel belts is emptied into the buckets on the serial belt, and so on (see Fig. 9). A more detailed illustration of how the charges are passed along from pixel to pixel by cycling the gate voltages is given in Fig. 10.

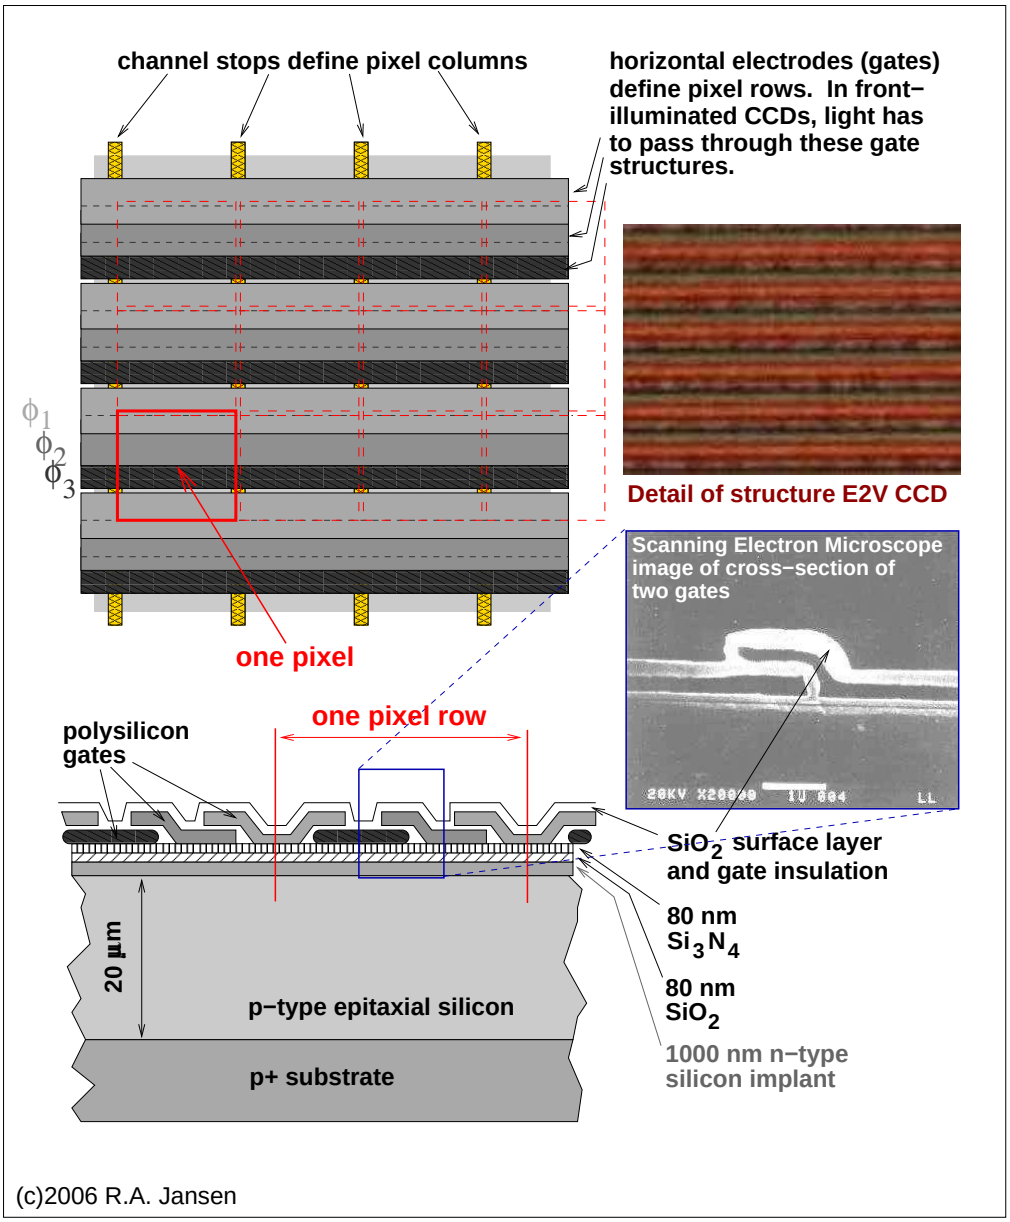

Figure 7: Schematic representation of the gate structures on a typical 3-phase CCD. Each pixel is defined by three gates and two channel stops (buried beneath and electrically isolated from the gates). When the CCD is exposed, all photo-electrons generated within the pixel volume migrate to and remain confined in the potential well of the middle gate, with the largest positive voltage (ϕ_2). The polysilicon gates partially overlap, but are isolated from one another by a thin layer of SiO_2. The relatively thick ($\sim 1\mu$m) layer of n-type Si, below the insulating SiO_2 and Si_3N_4 layers, helps reduce charge transfer losses: it keeps accumulated electrons away from the surface, where some might become trapped. At right, a detail from the image area of the CCD of Fig. 8 is shown. The repeating horizontal pattern of the gates is clear. An electron microscope scan reveals the vertical structure at the overlap of two gates.

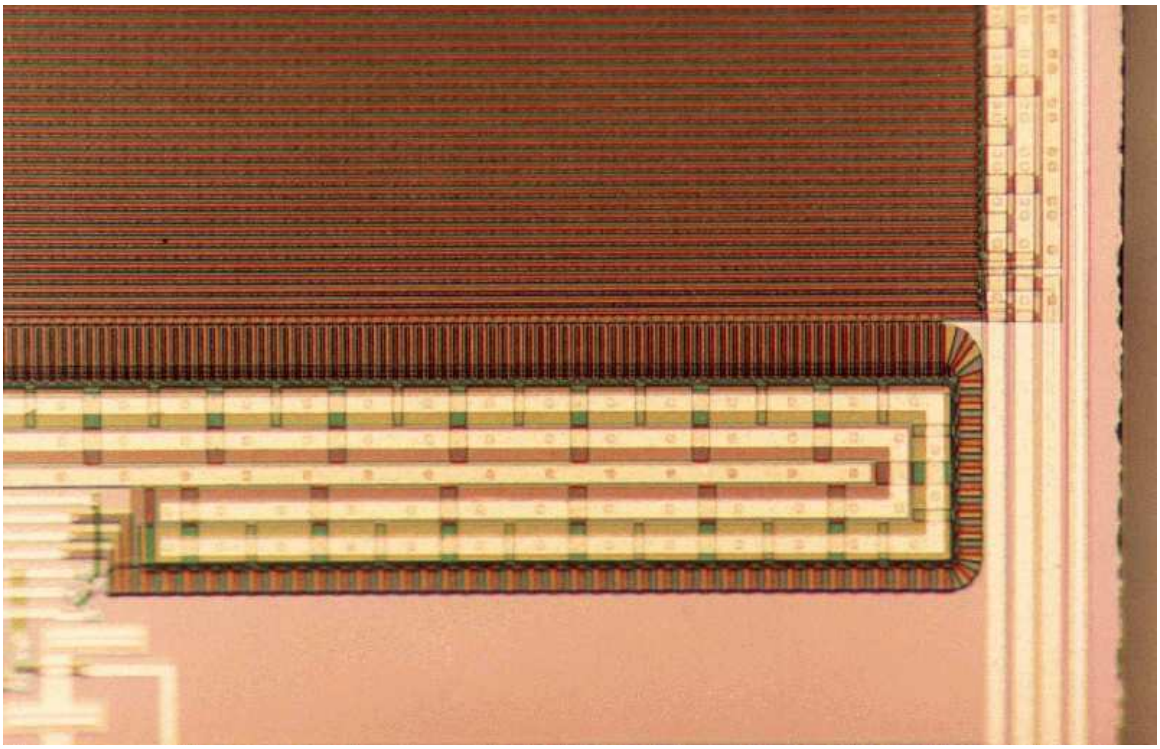

Figure 8: Photomicrograph of a corner of an E2V L3CCD, showing the image area (pixels), and the serial register (the vertically extended pixels just below the image area. On this particular CCD, the serial register bends around into an extended register, where the charge is amplified (clock voltages in the extended portion are sufficiently large to ensure a finite probability of dislodging additional electrons in each charge transfer from pixel to pixel. Part of the actual on-chip read-out electronics is visible at the lower left.

In class, an animation of the operation of a CCD, consisting of preparing the CCD for an exposure (clearing any residual accumulated charge), exposing, and read-out will be show.

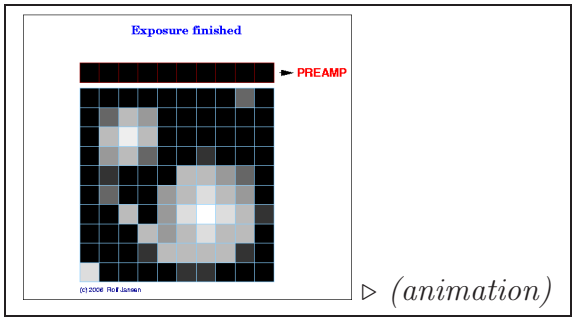

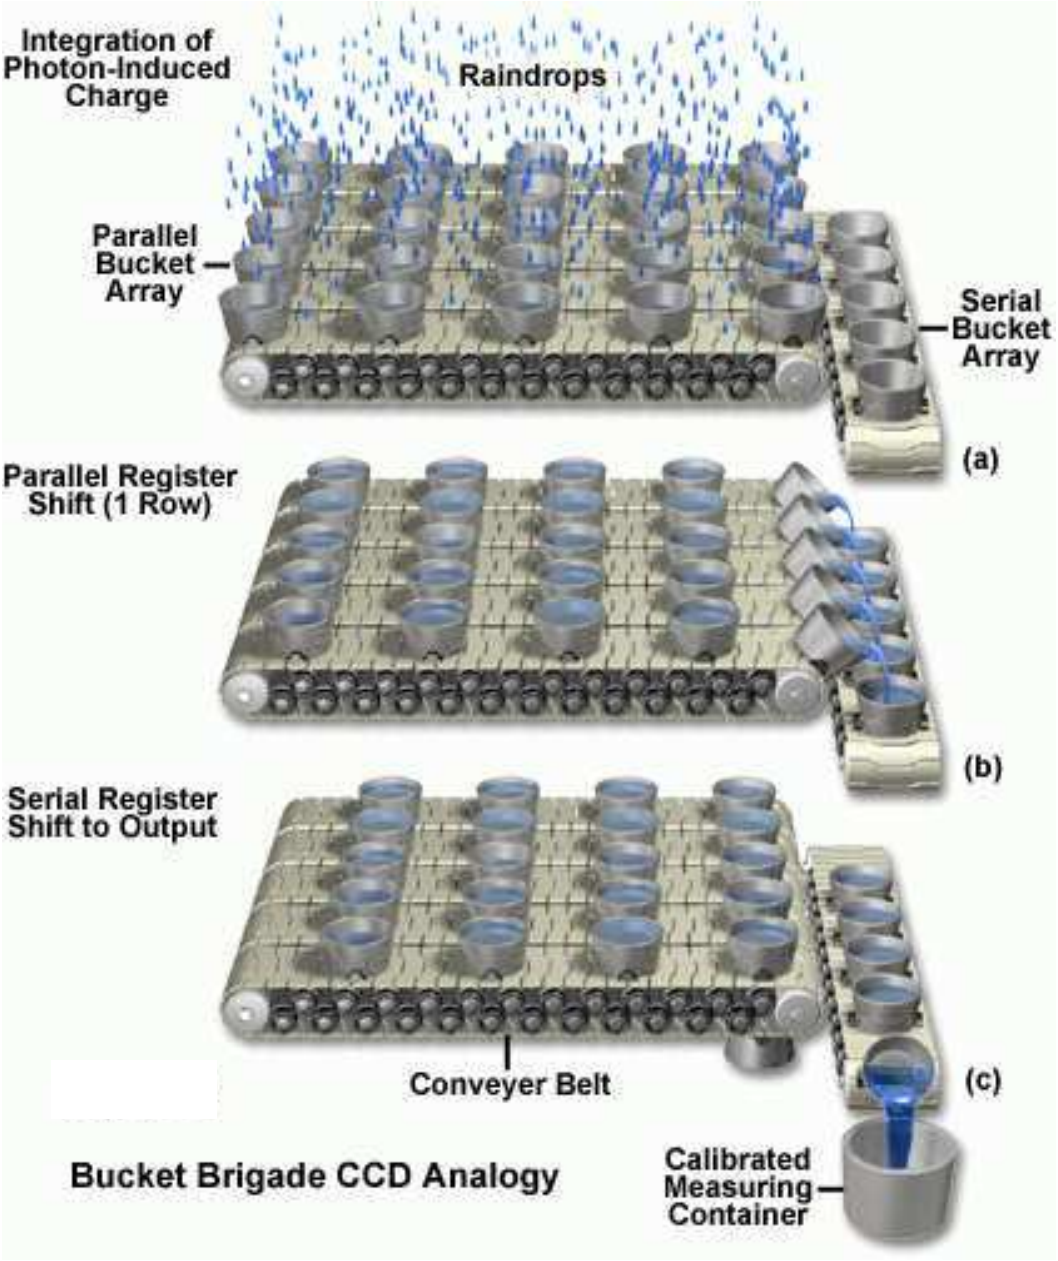

Figure 9: The "bucket brigade" analogy often used to visualize the operation of a CCD. Each bucket represents a pixel that can store charge during integration (exposure of the CCD to light). At the end of the integration, all parallel belts will simultaneously rotate their buckets toward the serial belt and dump the contents of one bucket per belt into the corresponding bucket on the serial belt. Then the serial belt rotates and presents one bucket at a time for measuring. When all buckets on the serial belt have been measured, the parallel belts will rotate and dump the next "row" of buckets into the corresponding buckets on the serial belt, and so on until all buckets have been measured. *Courtesy: Nikon/K.R. Spring, T.J. Fellers & M.W. Davidson.*

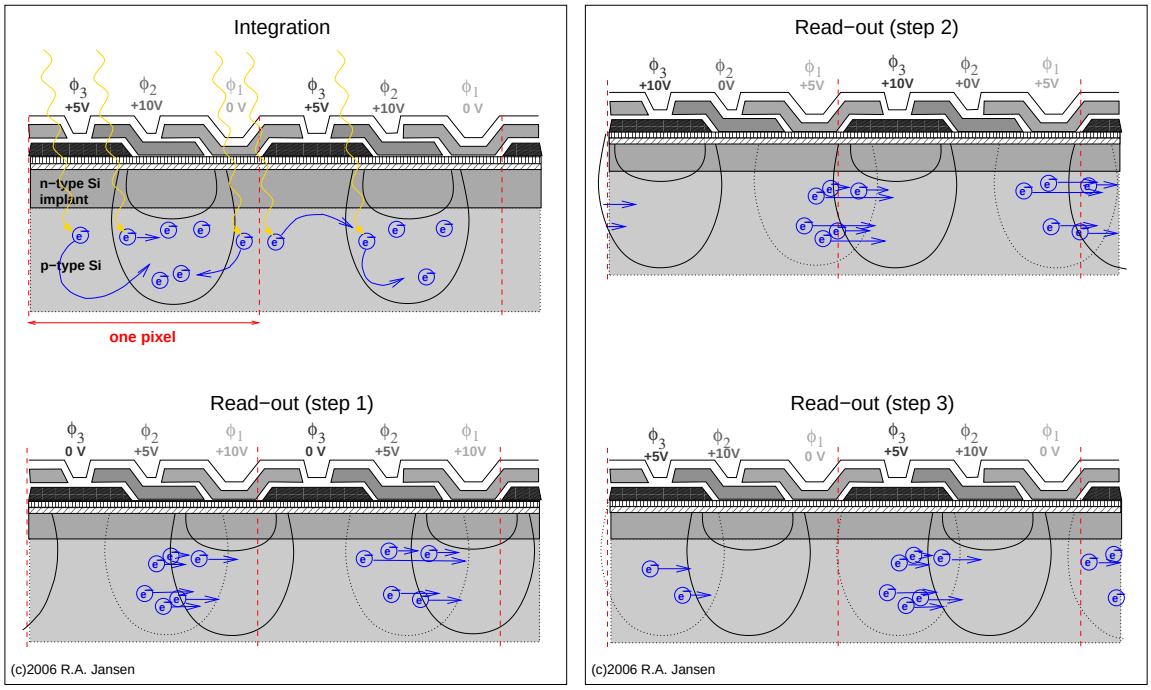

Figure 10: Schematic representation of the operation of a CCD during integration and read-out. During integration, the middle gate (ϕ_2) is kept at a large positive voltage, causing all photo-electrons generated under any of the three gates defining a pixel to migrate to the deep potential well of this gate. Note that the effect of the portion of the middle gate that overlaps adjacent gate ϕ_3 is shielded. When the CCD is read out, the voltages on the gates are cycled in succession, such that the accumulated electrons follow the deepest potential well until, at the end of a single parallel clocking cycle, they have been transferred to the adjacent pixel along the column. The illustration follows the transfer of charge within a single pixel column. Since the voltages on the gates in each pixel row along a column are connected in parallel, a single parallel clocking cycle will transfer the accumulated charge in the pixels of a row to the corresponding pixels of the serial output register (not shown). The charge is then transferred out of the serial register, one pixel at a time, to be amplified and sampled by the A/D converter. This is repeated until the entire array has been read out.

Summary of advantages and disadvantages of CCDs:

- CCDs can have peak QE in excess of 90% (when properly AR coated). Even uncoated, front-illuminated CCDs can have peak QE in excess of 70%. The QE is only a relatively slow function of wavelength within this bandpass.
- Their response to incident photons is highly linear over almost its entire operational regime.
- The natural operational band-pass of CCD's is ∼380 nm–1.1 μm, which can be extended down to ∼300 nm by optimizing antireflection coatings and by etching away most of the silicon substrate (*thinning*) and illuminating the CCD from the back. The band-pass can be extended even further toward the UV using phosphorescent coatings (like the one in your yellow marker pens), and somewhat toward the near-IR by additional doping of the silicon.
- The *full-well capacity*, the maximum number of photo-electrons that can be stored within a given pixel, is typically of the order of 10^5 e^-. The dynamic range is typically $\sim 10^4$. When more charge accumulates than can be confined by the gate potential (*full-well saturation*), the excess charge may bleed into adjacent pixels along a column.
- Saturation may also occur in another way: when the conversion of the measured charge to a digital number would require a larger number than can be stored in the available number of bits (*A/D saturation*).
- CCDs require cooling to cryogenic temperatures using liquid N_2 (-100°C–-120 °C or 170–150 K), or at least thermo-electric cooling (at ~ -80 °C, the detectors aboard the Hubble Space Telescope are at the warm end of the operational range). This cooling is required to reduce *dark current* due to thermal electrons from the bulk silicon also being dislodged and collected, and to reduce thermal noise in the read-out electronics (both on-chip and in the adjacent "head electronics").
- CCD's have small physical sizes ($\lesssim$4.5 inch) that are poorly matched to the physical size of the collimated beam of large telescopes ($\gtrsim$10 inch diameter), so that only a small portion (solid angle) of the focal plane of the telescope is actually used. CCD sizes are set by the difficulty of producing flawless large-format devices, and by the ∼12 cm diameter of the wafers used for all current silicon-based integrated circuit fabrication. Monolithic devices up to 8192×8192 pixels can now be semi-routinely produced, but the work-horse devices are currently only 2048×4096 pixels. Pixels are typically between 10μm and 24μm in size. The physical size limitation has been partly overcome by constructing arrays of multiple CCDs, that are read out simultaneously. Examples of such arrays are shown in Fig. 11.

- The temporal resolution, set by the maximum rate at which the array can be read out, is $\gtrsim$10 msec, i.e., $\sim$1,000–10,000$\times$ slower than for a PMT. In most integrating applications, the temporal resolution is *much* lower. The read-out time is largely set by the minimum sampling time per pixel to accurately estimate the accumulated charge. (Trade-off: faster read-out $\rightarrow$ less accurate; slower read-out $\rightarrow$ more accurate). Effective read-out times have been reduced in recent years by CCD designs that allow simultaneous read-out of four quadrants of the CCD through serial registers and amplifiers in each corner of the CCD, and through the development of faster and lower-noise read-out amplifiers.

- The *Charge Transfer Efficiency* (CTE) of CCDs is very close to, but less than 1. The charge losses accumulate the further a given pixel is removed from the read-out amplifier. For a 2048$\times$2048 pixel CCD read out through a single amplifier, the charge in the last pixel read experienced 2048 parallel transfers, followed by 2048 serial transfers. If we assume that parallel and serial transfers are equally efficient (not necessarily so), this results in only $(\mathrm{CTE})^{4096} \times 100\%$ of the originally generated photo-electrons to arrive at the read-out amplifier. In modern CCDs with buried channels, CTE is typically $\gtrsim$0.99998, giving $\gtrsim$92% in the above example.
  The complement of CTE, *Charge Transfer Inefficiency* ($\mathrm{CTI} = 1 - \mathrm{CTE}$) quantifies the charge losses. Hence $\mathrm{CTI} \lesssim 8\%$ in the above example. If the same CCD was read through four amplifiers (effectively forming 4 arrays of 1024$\times$1024 pixels), the losses would be $1 - (0.99998)^{2048} \lesssim 4\%$

CCDs in different flavors

CCDs come in several flavors, each optimized for specific uses, and with possible additional electronic circuitry built onto every pixel, or onto the serial output register. Examples are:

- Surface Channel CCDs. The first and simplest type of CCD, suffered from charge traps due to imperfections at or near the surface or gate structures. Traps could be reduced by a black-magic process called UV-flooding, or by uniformly illuminating the CCD at a low level prior to the exposure of the astronomical source (*pre-flashing*). Both processes effectively filled in the traps with some charge.

- Buried Channel CCDs. Since these have largely replaced the original Surface Channel CCDs, the above discussion of CCD structure already silently assumed a buried channel type. The only drawback of the buried channel design is that the storage capacity of each pixel is reduced by a factor 3–4 with respect to a surface channel design.

- Front-illuminated CCDs. In these, photons have to pass through the gates to reach the photosensitive bulk silicon. Since the gates are not entirely transparent, and their transparency depends on wavelength, the peak QE of front-illuminated CCDs is $\sim$70%, and their sensitivity steeply declines below 400 nm. The native thickness of the silicon substrate below the gate structures (typically $\sim$300 μm), does make these devices suitable for detecting far-red and near-infrared light up to $\sim$1.1 μm, but it also makes them more susceptible to charge dumped by the passage of cosmic rays.

- Back-side illuminated CCDs. The bulk silicon substrate of these CCDs is chemically etched away after manufacture. The remaining thickness is a mere $\sim$15 μm. The gate-side (front) of the device is then mounted on a rigid substrate, allowing the device to be illuminated from the back. Since the incident photons no longer have to pass through the gates, peak QEs of $\gtrsim$90% are easilly achievable, and the sensitivity to blue and near-UV light is significantly increased, while the susceptibility to cosmic ray induced signal is decreased. Drawbacks are a reduced sensitivity at the red end of optical spectrum, that the etched surface is not as flat resulting in larger variations in sensitivity, and that the remaining layer(s) of material are sufficiently thin to cause interference within the device when illuminated with (near)monochromatic light (e.g., night-sky emission lines).

- Interline and frame transfer CCDs. Light-shielded columns are either interspersed or added to one end of the CCD. During readout, the contents of one column (or the entire illuminated portion of the CCD) are rapidly shifted into the "spare" columns, without the slow step of sampling their stored charge. While the active columns are integrating on the next exposure, the spare columns are being read-out. Only half of the entire array is used to collect photons.

- Anti-blooming CCDs. An additional gate structure is added to the CCD that prevents charge from full-well saturated pixels to migrate to adjacent pixels along a column (called *bleeding*). Instead, any excess charge is drained off and lost. This allows long exposures on faint sources within fields dominated by bright sources (e.g., foreground stars). Drawbacks are, that $\sim$30% of the active pixel area is lost, reducing the peak QE of these devices, and that the total pixel well depth is reduced by a factor of $\sim$2. The tiny gaps in sensitivity between adjacent pixels can also affect the spatial resolution.

- Multipinned phase (MPP) CCDs. The much lower dark current at higher operating temperatures in these devices is achieved by a different electronic design and pixel structures, inversion of all three clock voltages and additional semiconductor doping (with Boron or Phosphorus) of the phase 3 gate. The drawback is a reduced well-depth.

- Othogonal Transfer CCDs (OTCCDs). Charge can be shifted both vertically along columns and horizontally along rows thanks to a spatially complex four-phase design, where the channel stops of regular CCDs become a fourth gate. Two of the (center) gates are triangular in shape, and the remaining two are split into pairs and rectangular in shape, surrounding the center gates (see Fig. 2.7 of Howell 2006). The additional electronic structures result in a lower QE and charge trapping can be more of an issue. By shifting the charges *while exposing*, image motions can be corrected, or PSF's can be actively shaped for high-precision photometry.

- Low Light Level CCDs (L3CCDs). The serial output register is extended and the gate voltages over the extended pixels are much higher (40–60 V), resulting in a finite probability (1–2%) of kicking loose another electrong (avalanche amplification). The total amplification factor scales as $(1+p)^N$, with $p \sim$ 1-2% and N the number of pixels in the extended register (typically $\sim$500). Since the read-noise of the output amplifier is now leveraged against the amplified signal instead of against the original signal, the S/N is sufficiently high to count individual photo-electrons (=incident photons). Fig. 8 shows this extended register for an E2V L3CCD.

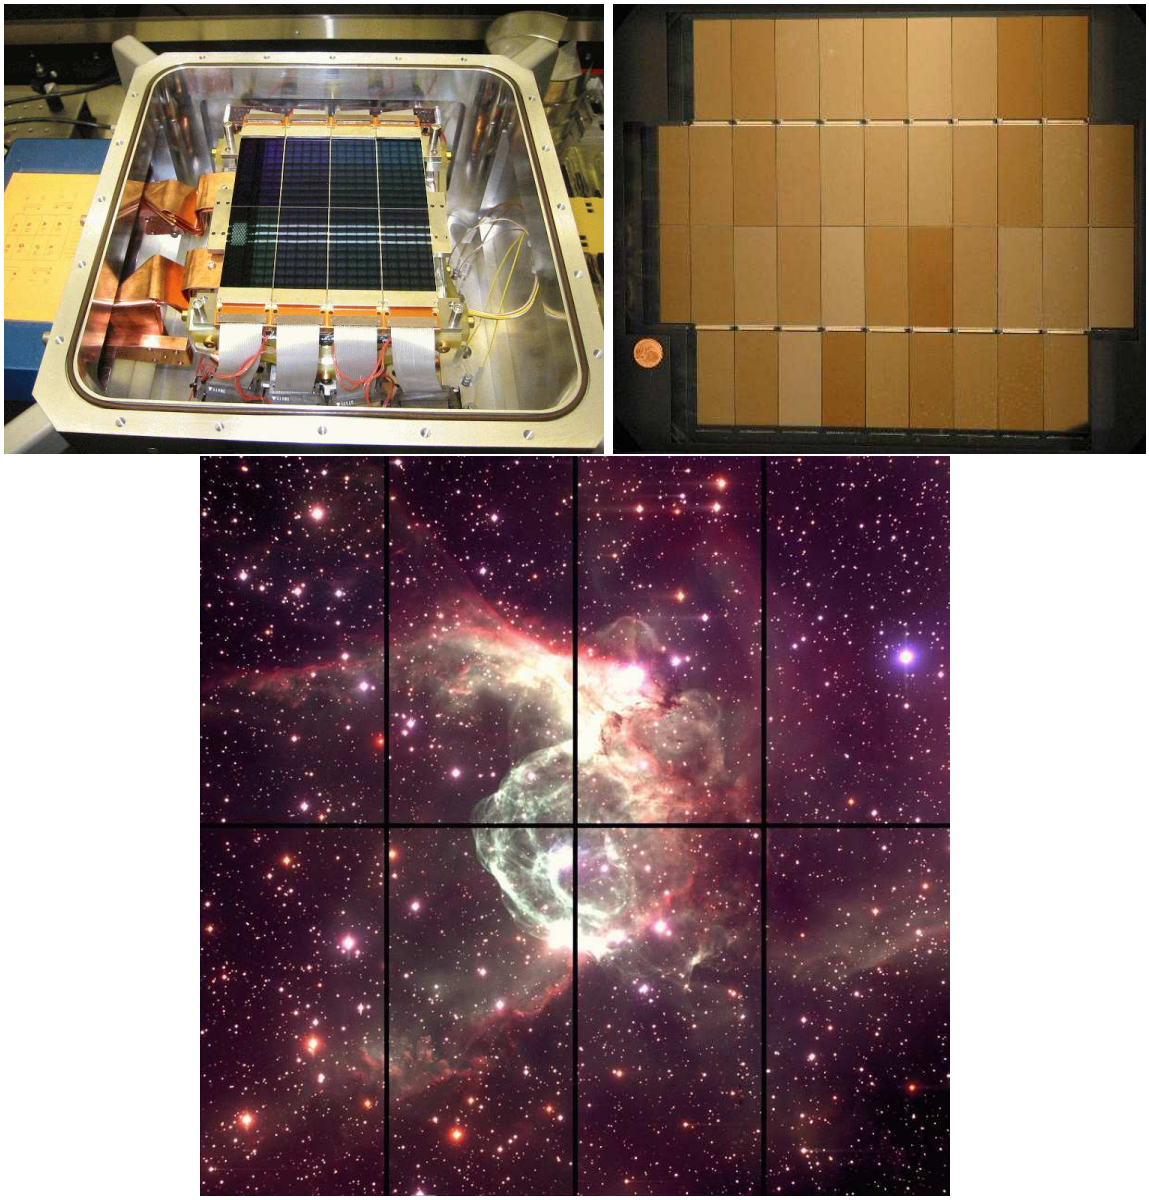

Figure 11: Example of large arrays of CCDs: (a) Magellan/IMACS, an array of 4×2 large-format CCDs. *Courtesy: OCIW*; (b) CFHT/Megacam, an array of a staggering 40! CCDs. *Courtesy: Canada-France-Hawaii Telescope/2003*. These CCD cameras allow imaging over $\sim 0.5^\circ \times 0.5^\circ$ and $\sim 0.94^\circ \times 0.94^\circ$. (c) An example of a wide-field Magellan/IMACS image. Note that the areal coverage is not complete: to fill the gaps ("butt cracks") between the CCDs, several dithered exposures need to be taken. *Courtesy: Danny Steeghs*

1.1.7 Complementary Metal Oxide Semiconductor arrays

Complementary Metal Oxide Semiconductor (CMOS) arrays are already pervasive in todays digital camera and video market, and are becoming suitable for use in astronomy as well. These devices are very similar to CCDs, but have more complex additional circuitry built into every pixel. This circuitry allows individual pixels to be addressed directly, and their charge to be read-out and converted from analog voltage to digital data number. The lower QE of CMOS devices, due to shielding of the silicon by the electronic circuitry, is not a problem for digital cameras that operate at high light levels.

- The main advantage of CMOS devices is that each pixel is individually addressable: charge in a pixel need not be transferred along a column and subsequently along a serial output register. Single-pixel read-out, A/D conversion and subsequent signal processing (such as up-the-ramp sampling to reject cosmic ray induced signal) is possible.
- The main disadvantage of current CMOS devices is that the additional circuitry leaves only a small active photosensitive region in each pixel, which results in a QE that is typically much lower ($\sim$10–20%) than that of a CCD. This handicap is being overcome in the newest state-of-the-art back-illuminated devices developed for use in future space missions.
- Because of the additional circuitry, which requires more layers to be deposited per device, their production cost is higher. (And added complexity results in a lower yield of flawless devices).

1.1.8 Superconducting Tunnel Junction arrays

Superconducting Tunnel Junction (STJ) devices may be the path of future array detectors. These detectors use a layer of supercooled polycrystalline niobium (Nb), which produces numerous pairs of free electrons per incident photon, with the total number dependent on the energy of that photon.

- Photon-counting arrays that are sensitive not only to the arrival *location* of an incident photon, but also to its approximate *energy*. Hence, where presently filters are needed, in future it may be possible to get filterless broad-band photometry, or very low resolution spectroscopy, i.e., to measure Spectral Energy Distributions (SEDs).
- Operational temperatures near absolute zero (a few **milli-K**) are required, and hence cooling with liquid Helium rather than liquid Nitrogen.

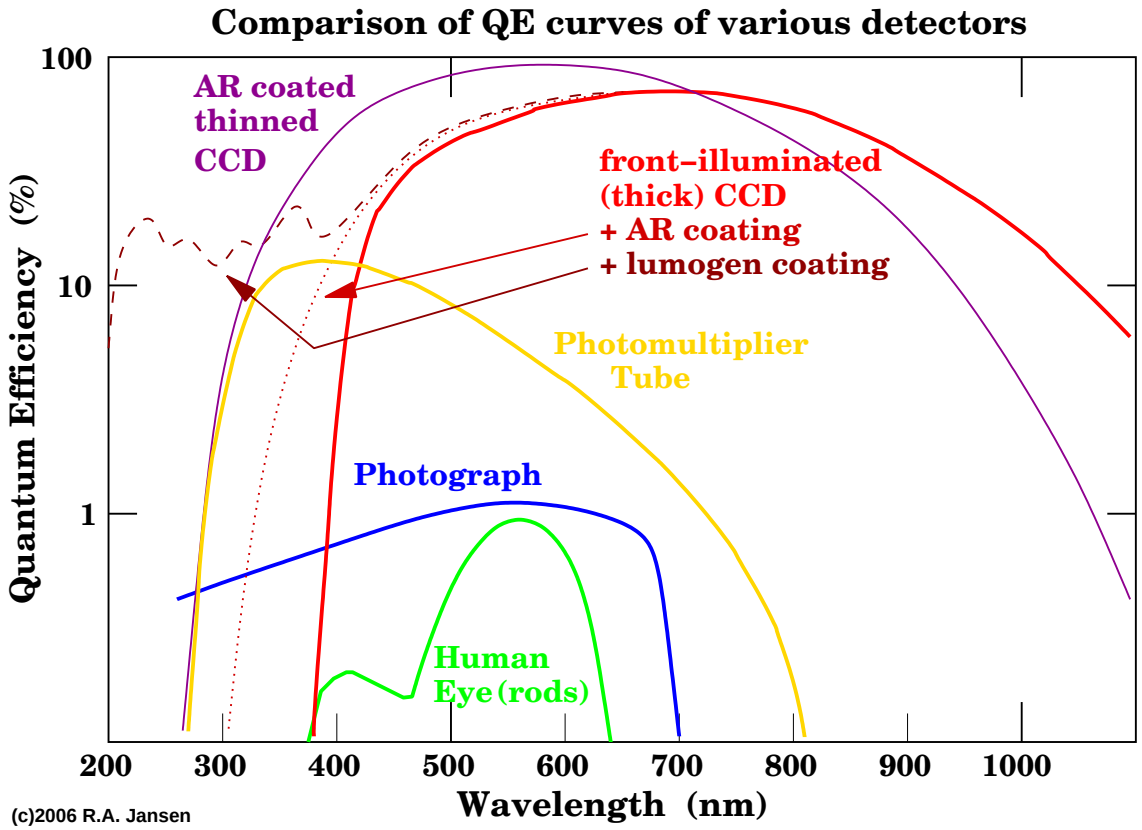

Figure 12: Comparison of typical QE curves for several of the different types of detectors discussed. The curve for CMOS devices is similar to that for CCDs, but with —at present— a lower peak QE.

- Devices are extremely sensitive to ambient electromagnetic fields, particularly their ability to resolve the photon energies, and required special shielding and possibly special measures at the telescope itself to be taken.
- Present experimental array sizes are still very small (6×6 pixels) due to their very complicated electronic structure, but will likely greatly grow in the coming decade.

1.2 Analog to Digital Conversion

When a CCD is being read out, the charge in a given pixel is eventually shifted out of the serial output register into the on-chip amplifier (*pre-amp*) circuitry, and out to the off-chip CCD head electronics. There, through a technique called Correlated Double Sampling (CDS), the pixel charge is measured as the difference between an arbitrary reset voltage and that arbitrary reset voltage plus the charge packet from the pixel. The measured charge is an analog value (assuming $N_{e^-} \gg 1$ for that pixel).

That value is subsequently represented as a digital number by the Analog-to-Digital converter (A/D). The resulting digital number is called *Analog-to-Digital Unit* (ADU), *Data Number* (DN), or simply *Count*.

- CDS sampling time is finite, typically 10–20 μsec for both $-V_0$ and $V_0 + V_{pix}$.
  faster → less accurate (higher *read-noise*)
  slower → more accurate (lower *read-noise*)

▷ The average ratio of the analog charge measurement and the assigned data number is called the *gain*. The gain is expressed in units of e^-/ADU. Because the gain value is usually ≥ 1, the assigned data number is *smaller* than the number of photo-electrons (i.e., a "reduction" rather than a "gain"). That's why in the literature one sometimes encounters the term *inverse gain* for this parameter.

- The minimum statistical error in any pixel value is 0.5 gain. Thermal noise within the on-chip (extended register +) pre-amp or in the CCD head-electronics can only increase this error.

To keep the amplifier and A/D electronics simple, and to avoid wasting an additional sign-bit for storing the digital data, the CCD electronics add a fixed offset or *bias voltage*. Hence, even when no photons were detected in a given pixel, the A/D converter will be presented with a positive measurement.

The maximum number that can be stored within an n-bit un-signed digital number is $2^n - 1$ (assuming we start counting at 0). For a typical 16-bit per pixel, the maximum number is, therefore, $2^{16} - 1 = 65535$. (Should we have used a sign-bit, then the maximum number would have been only $2^{n-1} - 1 = 32767$).

With modern, non-truncating A/D converters, the average value in ADU returned for a signal of S e^- is $\mathrm{int}(S/\mathrm{gain} + 0.5)$. For older truncating A/D converters, $\mathrm{int}(S/\mathrm{gain})$ would be returned instead.

A/D conversion is supposed to be highly linear over the entire operational range from zero electrons to full-well capacity. In practice, as electronics becomes smaller and read-out speeds become larger, while overall noise is reduced, thermally induced non-linearities may become noticeable (due to slight warming and cooling of the amplifier, depending on how much power it draws, i.e., on how much charge it saw per unit time).

2 Characterization of CCDs

2.1 Absorption length, Quantum efficiency, Charge diffusion

The active photosensitive material of CCDs is pure silicon. Silicon therefore largely determines the response of the detector to light of various wavelengths.

▷ The *absorption length* of photons in a material is that distance after which a fraction of $1/e$ remains. This is equivalent with saying: the distance after which 63% of the photons are absorbed.

Fig. 13 shows the dependence of the absorption length on wavelength from X-rays ($\sim$10Å) to the near-IR 1.1 μm cut-off (corresponding to the Si band-gap energy). Fig. 15 illustrates some of the implications of this wavelength dependence.

- Shortward of $\sim$400 nm, photons are absorbed very near the surface of the CCDs. In front-illuminated CCDs, this absorption is likely to occur already within the layers of the gate structures and/or the Si_3N_4 and SiO_2 isolation layers. In thinned, back-illuminated CCDs, this is much less of a problem.

- Pure silicon is shiny like a metal and a fairly good reflector of light, particularly at bluer colors (see Fig. 14). In addition to the small absorption length, bluer photons are more likely to be reflected off of the CCD surface. Anti-reflection coatings such as MgF_2 can greatly reduce these losses.

- Shortward of $\sim$250 nm, the penetration depth of photons increases again, but such energetic photons can generate multiple electron-hole pairs within the silicon. At very short wavelengths (X-ray and shorter) they can even cause damage to the CCD itself.

- Longward of $\sim$800 nm, thinned CCDs become transparent to incident photons, because the absorption length exceeds the typical thickness of the silicon substrate.

- At redder wavelengths, the absorption may occur deep in the bulk substrate, where the pixel gate potential well is relatively shallow, rather than in the depletion region. The resulting photo-electrons may migrate to a neighboring pixel along a column, or underneath the channel stops to a neighboring pixel along a row.

▷ Migration of photo-electrons created deep within the bulk substrate and/or well away from the center of the pixel volume to neighboring pixels along a column or row is called *charge diffusion*. Because the depth of the potential well is not uniform within the pixel volume, charge diffusion is often ascribed to *CCD sub-pixel variations*. The net effect of charge diffusion is image blurring.

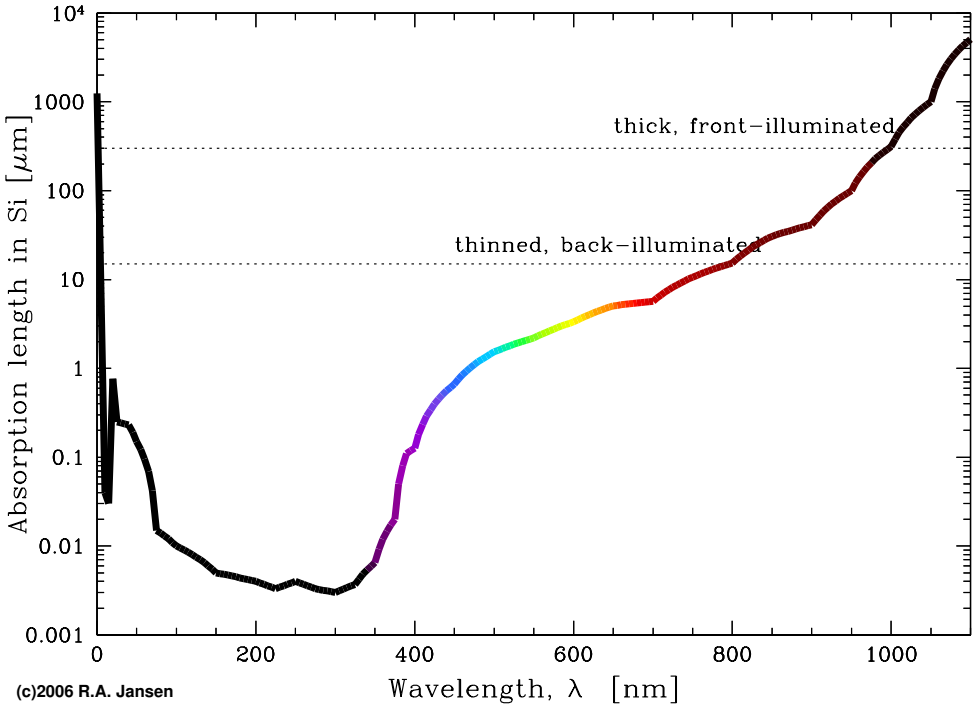

Figure 13: Photon absorption length in silicon (in μm) as a function of wavelength (in nm). Dotted lines indicate the typical thickness of the silicon substrate in front- and back-illuminated CCDs. Thinned back-illuminated CCDs are significantly less efficient at wavelengths redward of ∼800 nm. Not included in this curve is the fact that without additional measures like anti-reflection coatings, 70% or more of the photons may reflect off the front surface at short wavelengths: pure silicon looks shiny like a metal (see Fig. 14).

Figure 14: The lower of the above CCDs has had a coating of MgF_2 applied, while the upper chip is thinned but still without any coating. *Courtesy: M. Lesser/UA/ITL.*

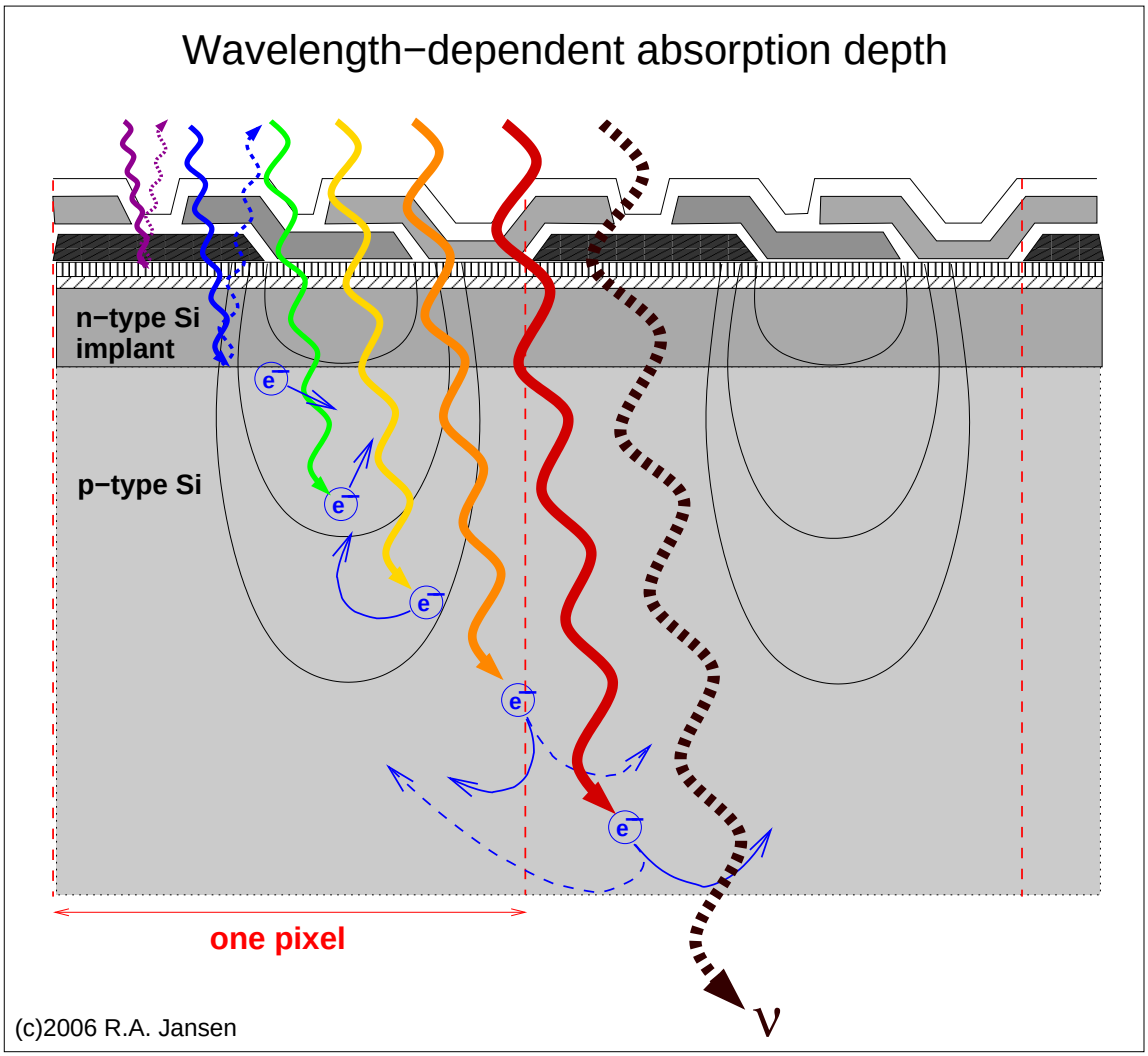

Figure 15: Implications of the wavelength-dependence of the photon absorption depth. Short wavelength photons are absorbed nearer the surface of a CCD than longer wavelength photons. Near-UV and blue photons may not even make it to the depletion region, since they are likely to be absorbed by the gate structures and isolation layers constituting the top few 100 nm. Photons in the near-IR portion of the spectrum get absorbed deep within the bulk substrate, where the pixel gate potential is relatively weak, rather than in the depletion region. The resulting photo-electrons may migrate to a neighboring pixel along a column, or underneath the channel stops to a neighboring pixel along a row. Such *charge diffusion* causes image blurring. In a front-illuminated CCD (depicted here), the silicon substrate is sufficiently thick ($\sim$300 μm) to capture all but the very reddest photons, and the thin metal (gold) layer at the rear doubles the absorption probability by reflecting photons back up through the pixel. Thinned, back-illuminated CCDs, become transparent for $\lambda \gtrsim 800$ nm (see Fig. 13).

- The QE of a CCD is better at higher operating temperatures, particularly for $\lambda \gtrsim 800$ nm. However, since raising the temperature means increasing the thermal dark current and read-noise (see also Fig. 19), there is very little net gain in doing so.

- With the purity of the Si used to fabricate a CCD, its resistivity goes up, allowing higher gate voltages, and hence larger electron storage capacity. Since higher voltages also mean *deeper* potential wells at larger distances from the gates, the probability of confining photo-electrons produced by near-IR photons increases: i.e., better QE at the reddest wavelengths. The price to pay is, indeed, the much higher price of fabrication of such a pure device.

▷ Tuning the parameters of a CCD *always* means finding a compromise between conflicting constraints.

- Measured QE curves (see, e.g., Fig. 16) are assumed to be representative *of each and every pixel* of the CCD. This is not quite so, although often a reasonable approximation. One always has to correct for pixel-to-pixel variations in the reponse.

- The QE curve delivered with a particular CCD may only be representative of *a typical device of that kind*, but not be a measurement of the wavelength dependent response of that device itself.

▷ Hence, to calibrate observations of an astronomical object of interest through a given filter, one also needs to observe *photometric standard stars* (e.g., Landolt 1992; Stetson 2000), or *spectrophotometric standard stars* (e.g., Oke 1990; Hamuy et al. 1992,1994). The apparent brightness of such standard stars — corrected for Earth's atmosphere — is accurately known either in several standard passbands (see: Bessel 1990), or as a function of wavelength in small wavelength intervals.

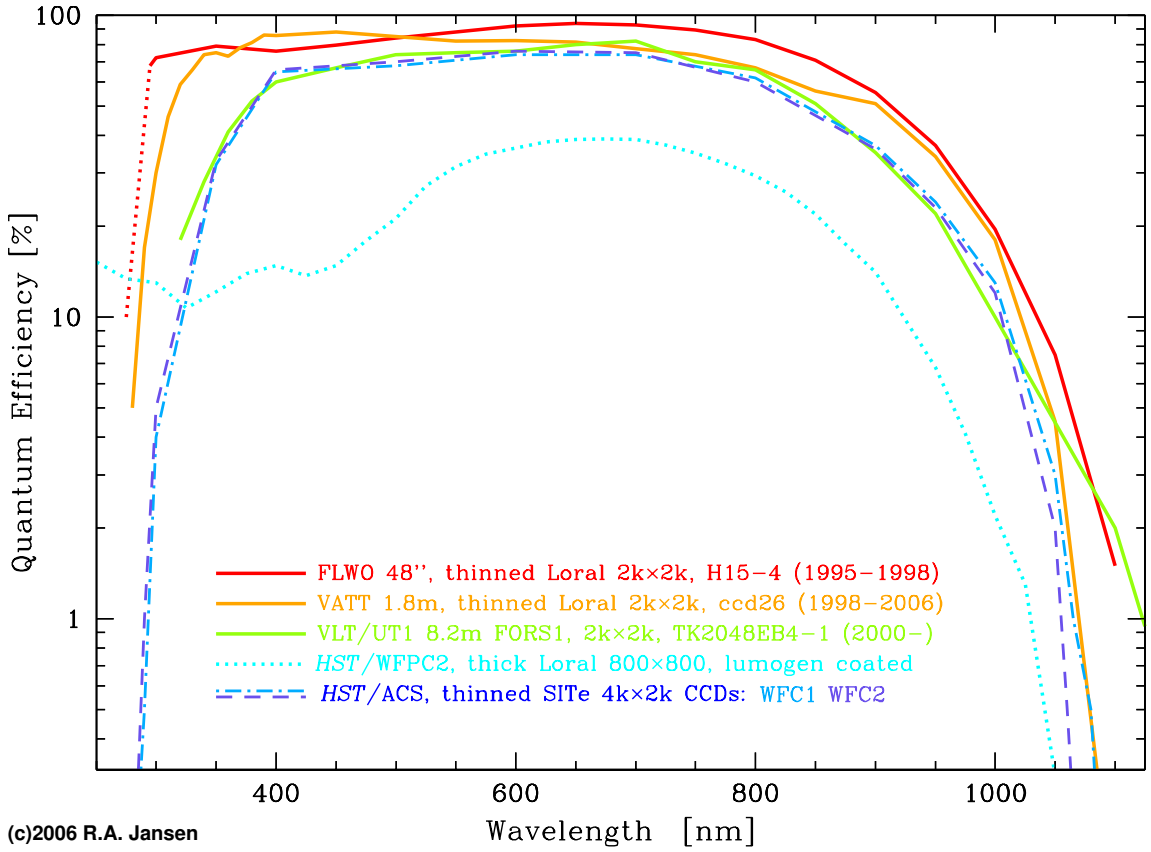

Figure 16: Comparison of QE curves for several CCDs that I happen to have used. Most of these detectors are of the same generation. The 800×800 pixel front-illuminated *HST*/WFPC2 CCDs, however, are older. Their high UV response is due entirely to a lumogen coating, which converts UV–blue photons to green ones to which the CCDs are sensitive. Current generation, red-optimized, front-illuminated CCDs perform significantly better than these WFPC2 chips.

Charge Diffusion and PSF halos

Aboard the *Hubble Space Telescope*, two of the three cameras that make up the Advanced Camera for Surveys (ACS) can observe in the optical wavelength regime: the Wide Field Channel (WFC) and the High Resolution Channel (HRC). Both cameras employ similar but differently optimized thinned back-illuminated CCDs, with the gate structures at what has become the rear. Photo-electrons from bluer photons absorbed near the illuminated surface experience a weaker electic field (farther from the gates) than those from redder photons that penetrate deeper (are absorbed nearer the gates). As a result, the effects of charge diffusion are more noticeable in observations through bluer filters.

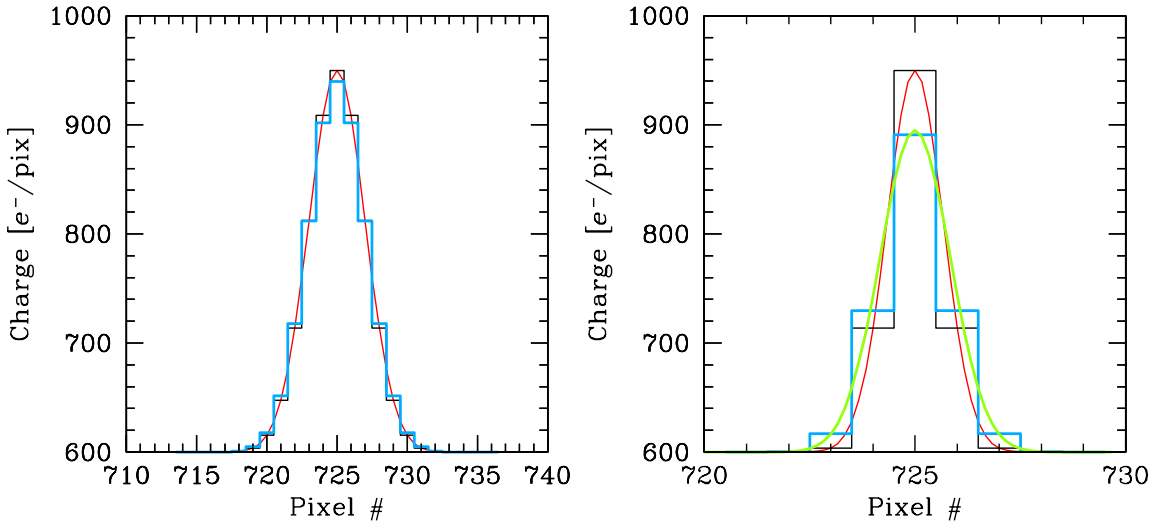

Figure 17: Charge diffusion is more detrimental when the PSF (*red*) is poorly sampled. The thin black histograms indicate the distribution of photons absorbed per pixel, while the thick blue histograms indicate the charge distribution upon read-out. The inferred green PSF overlaid on the right panel represents a ∼25% deterioration in resolution with respect to the original incident PSF.

Charge Diffusion and PSF halos (cont'd)

The thinning process, where much of the bulk substrate is etched away, is inexact, resulting in variations in thickness of the remaining silicon. Consequently, charge diffusion varies over the area of the CCD, with thicker regions suffering greater charge diffusion.

The effects of charge diffusion are more detrimental when the Point Spread Function (PSF) is not properly sampled, i.e., when the Nyquist criterion (N_{pix}/FWHM $\gtrsim 3$ for a 2-D Gaussian distribution) is not satisfied (see Fig 17).

At wavelengths $\gtrsim$800 nm, photons may pass through the entire silicon layer and either be absorbed by the substrate to which the gate side of the thinned CCD is mounted, reflected back up into the pixel (where it may be absorbed, or scattered back into adjacent pixels. The scattered photons are responsible for a large PSF halo in HRC images at $\lambda \gtrsim 700$ nm and in WFC images past 900 nm. Although the pixel scale of the WFC is coarser (0.05″/pix) than that in the HRC (∼0.027″/pix), the PSF halo will render the effective resolution for both cameras at ∼850 nm almost identical.

2.2 Charge Transfer Efficiency and Deferred Charge

The *Charge Transfer Efficiency* (CTE) of modern CCDs is typically better than 0.99999, but never exactly 1. The charge losses accumulate the further a given pixel is removed from the read-out amplifier. In a $n \times m$ pixel CCD, the charge of the last pixel to be read out, will be transferred along its column (parallel shifts) a total of n times, and along the serial output register (serial shifts) a total of m times. If n and m are large (as in most modern astronomical detectors), the CTE needs to be near perfect. The *Charge Transfer Inefficiency* (CTI), which quantifies the charge losses, is simply the complement of CTE, i.e., $1 - \text{CTE}$. The parallel CTE ($\mathcal{C}_p$) need not be identical to the serial CTE ($\mathcal{C}_s$), since the pixel sizes, gate structures, and voltages of the serial register often differ from those in the imaging area.

If N_{e^-} denotes the total number of electrons (total charge) stored in that final pixel, then the number of electrons that will actually arrive at the amplifier will be:

$$N'_{e^-} = N_{e^-} \cdot \mathcal{C}_p^n \cdot \mathcal{C}_s^m$$

Hence, the charge lost, L_{e^-}, is given by:

$$L_{e^-} = N_{e^-} - N'_{e^-} = N_{e^-} \left(1 - \mathcal{C}_p^n \, \mathcal{C}_s^m\right)$$

△ *Note that the expression given in the textbook (Howell 2006) on page 44 is* ***incorrect!*** To convince yourself of this, compute charge loss $L(e)$ with his expression and N=1000 e^-, CTE=0.99 (CTI=0.01) and a total of 2048 pixel transfers. Can you lose more charge than you had to begin with?

- Charge lost during a single transfer need not necessarilly be permanently lost. A photo-electron left behind, joins the charge that will be transferred from the next pixel in the same column, and is *indistinguishable* from that charge (electrons are all alike). Upon a subsequent clocking cycle, it again has a probability of value $\mathcal{C}_p$ to be transferred.

▷ Photo-electrons that fail to be transferred with the charge packet of the pixel in which they were generated, but that are transferred in subsequent transfers are referred to as *deferred charge*.

- Poor CTE in combination with deferred charge results in charge trails along columns behind bright sources.

- Prolonged exposure to ionizing radiation, and in particular bombardment with cosmic rays (mainly relativistic protons), results in permanent damage to the bulk silicon lattice within CCDs. Such damage results in charge traps and a general degradation of the CTE (see Fig. 18), and can also exacerbate charge diffusion. CCDs in space age on time scales of only a few years. Also, coatings on the CCD may evaporate and degrade in the high vacuum of space.

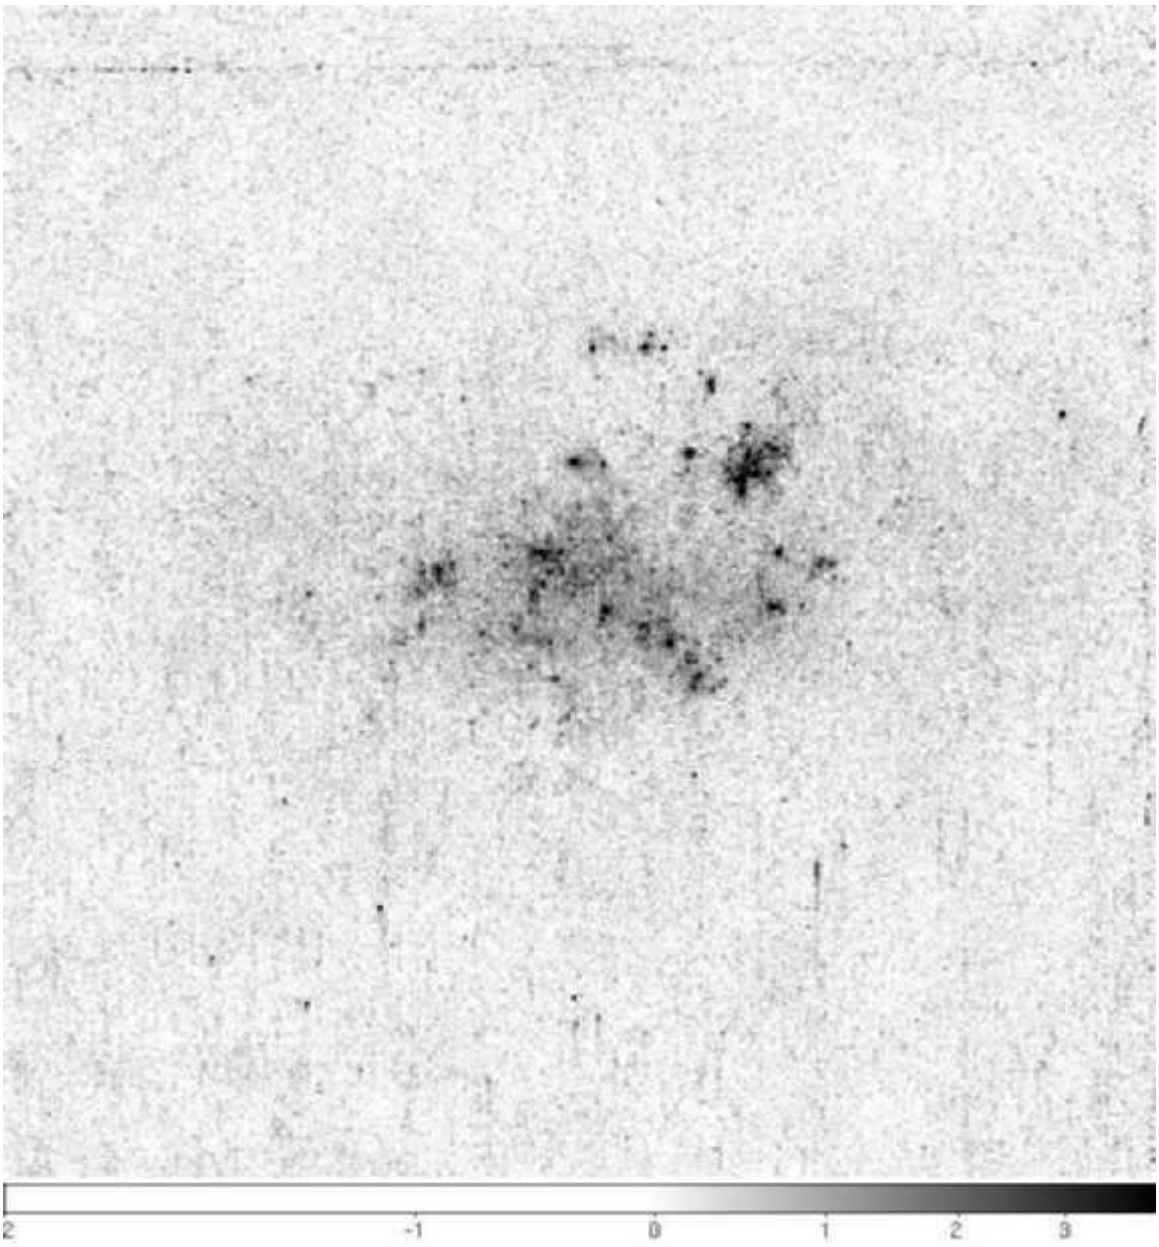

Figure 18: Example of the poor CTE and deferred charge trails behind bright sources that result from prolonged exposure to damaging particle radiation. The above near-UV (F300W) image of NGC 14 was taken with *HST*/WFPC2 during Cycle 10 after the CCDs had been in space for over 7 years and had considerably degraded with respect to their performance just after launch. The very low sky background levels ($\lesssim 5\, e^-$/pix) in the near-UV exacerbate the problem.
Note that at the tips of some trails there is no bright object! There, the bright "object" was charge left by passage of a cosmic ray particle. By combining two or more exposures, high level induced signal from such cosmic ray hits could be rejected, but the faint ($\lesssim 1\, \sigma$) trails generally remain.

2.3 Dark Current

Every material at $T \gg 0\,K$ will be subject to thermal noise within. CCD dark current is due to electrons freed from the valence band and collected within the potential well of a pixel, where they become indistinguishable from the photo-electrons.

- The dark rate depends strongly on CCD temperature (see Fig. 19). For modern CCDs, operated at $T \sim -120\,°C$, the dark rate is typically less than $10\,e^-$/pix/hr. The dark current would be lower if a CCD were to be operated at a lower temperature, but then the QE in the red and the CTE would also be lower.
  $\longrightarrow$ *compromise*.
- The dark signal also depends on the gate voltages and, for a given resistivity of the bulk Si substrate, will increase at higher voltages.
- The dark signal builts up *linearly* with exposure time. Hence, long exposures likely need to be corrected, while dark signal in short exposures may well be negligible if the dark rate is low.
- The distribution of the dark signal tends to satisfy Poissonian statistics.

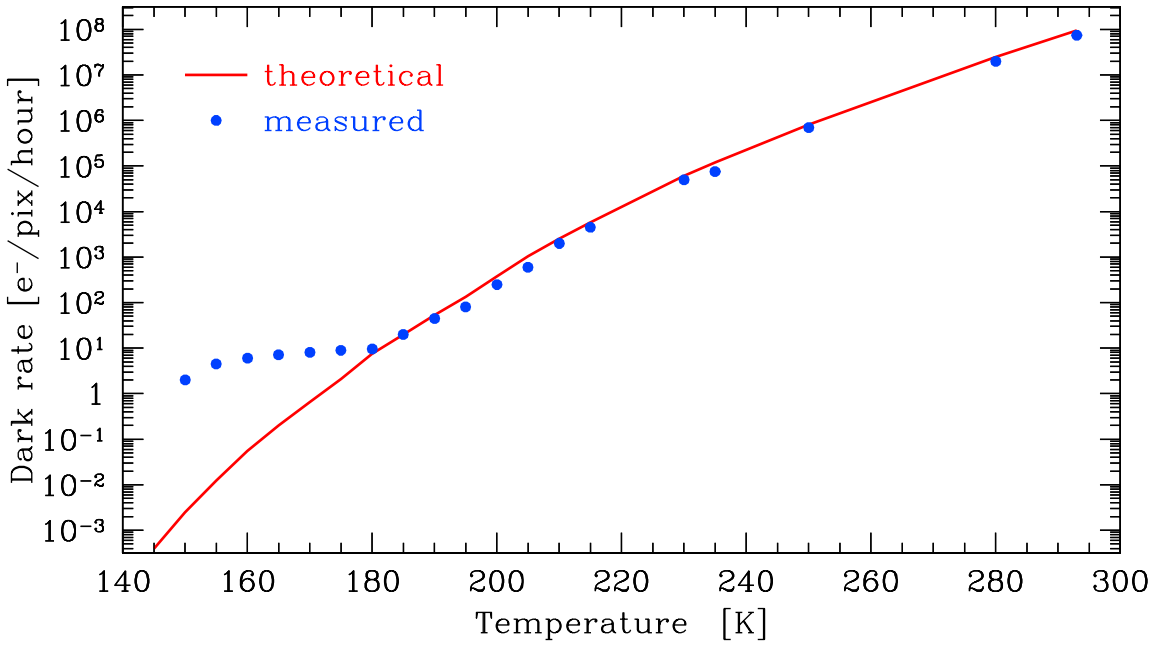

Figure 19: Measured and expected temperature dependence of the dark rate for a thinned back-illuminated 2k×4k EEV CCD delivered to ESO. Below ∼180 K, the measured dark rate does not diminish as expected due to other residual effects. Nonetheless, this particular device reaches a dark current of only $\sim 2\,e^-$/pix/hour when operated at 150 K.

2.4 Bias and Overscan

To ensure the A/D converter will never be presented with a negative number for the charge estimate — which could easily occur due to read-noise when the absorbed number of photons in a given pixel is small or zero — a sufficiently large *offset voltage* is added at the amplifier stage.

- This offset or *bias* voltage is *independent* of the length of the integration of an exposure and should theoretically be constant for all pixels. In reality, there may be some (generally low-level) structure that varies from pixel to pixel in a repeatable fashion (see Fig. 20).
- On long time scales — particularly when cryogenic cooling is not maintained when an instrument is not in use — the bias pattern may show variations. $\longrightarrow$ *show in class the animation of the changing structure in VATT "super" bias images with time: 1998 – 2006.*
- If patterns are discernable and of sufficient strength to contribute measurably to the noise, then a sufficiently large number of *bias frames* — zero-length exposures with the shutter closed — should be taken such that: $\mathcal{R}/\sqrt{N_b} \ll \Delta\mathcal{S}_b$, with $\mathcal{R}$ the read-noise, N_b the number of bias frames, and $\Delta\mathcal{S}_b$ the peak-to-peak pattern strength. $\mathcal{R}$ and $\Delta\mathcal{S}_b$ should both be expressed in either e^- or ADU.

Preceding read-out and after read-out of all pixels of the CCD has finished the read-out electronics will perform several additional (usually 8–40) clocking cycles that do not correspond to any physical pixels. These cycles will, however, exercise the amplifier and present a charge estimate corresponding to the instantaneous bias level as modulated by the read-noise to the A/D converter.

▷ These *virtual pixels* are referred to as the *overscan*, or *overscan strip(s)*. The digital numbers corresponding to these virtual pixel columns will be written to disk as if they were integral part of the CCD image.

- In the subsequent data reduction, the overscan strip(s) will be used to subtract the mean level and possibly any large-scale gradient in that level from each CCD frame.
- The bias level often displays small changes over time, e.g., as the temperature of the warm electronics changes during a night. These changes are tracked by the level measured in the overscan strip(s).

Some CCD designs also have some rows of physical pixels that are shielded from light. These are termed the *physical overscan*, and can be used to diagnose CTE and deferred charge, or other problems with the CCD electronics (e.g., reset voltages).

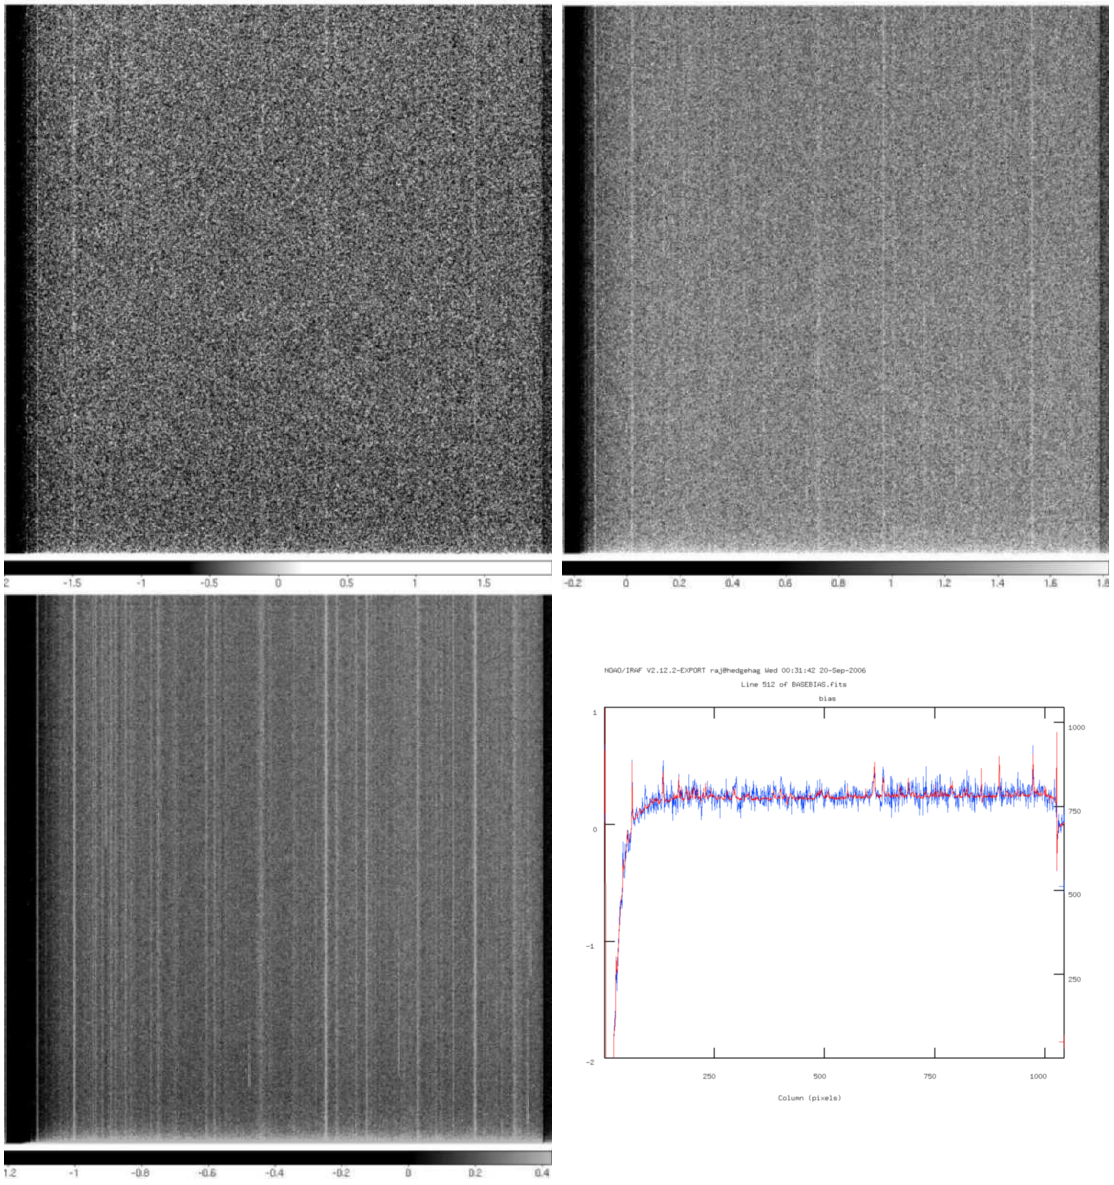

Figure 20: (a) Median average of 36 bias frames obtained at the Vatican Advanced Technology Telescope (VATT) in Nov 2001. Apart from a roll-on in the low columns, there appears to be little structure. (b) Median average of 128 bias frames obtained at the VATT in Nov 2001. Some low-level structure that remained hidden when combining only a small number of frames are becoming visible. (c) Baseline bias reference frame constructed from over 3500 frames reveals the wealth of systematic structure that is actually present. (d) A cross-cut along a single row (blue) and averaged over all rows (red) of this baseline bias reference frame, shows the strong roll-on in the low columns and the wealth of columns with bias levels that are slightly higher than average. For most studies, however, details at levels of $\lesssim 0.1\, e^-$ are probably negligible.

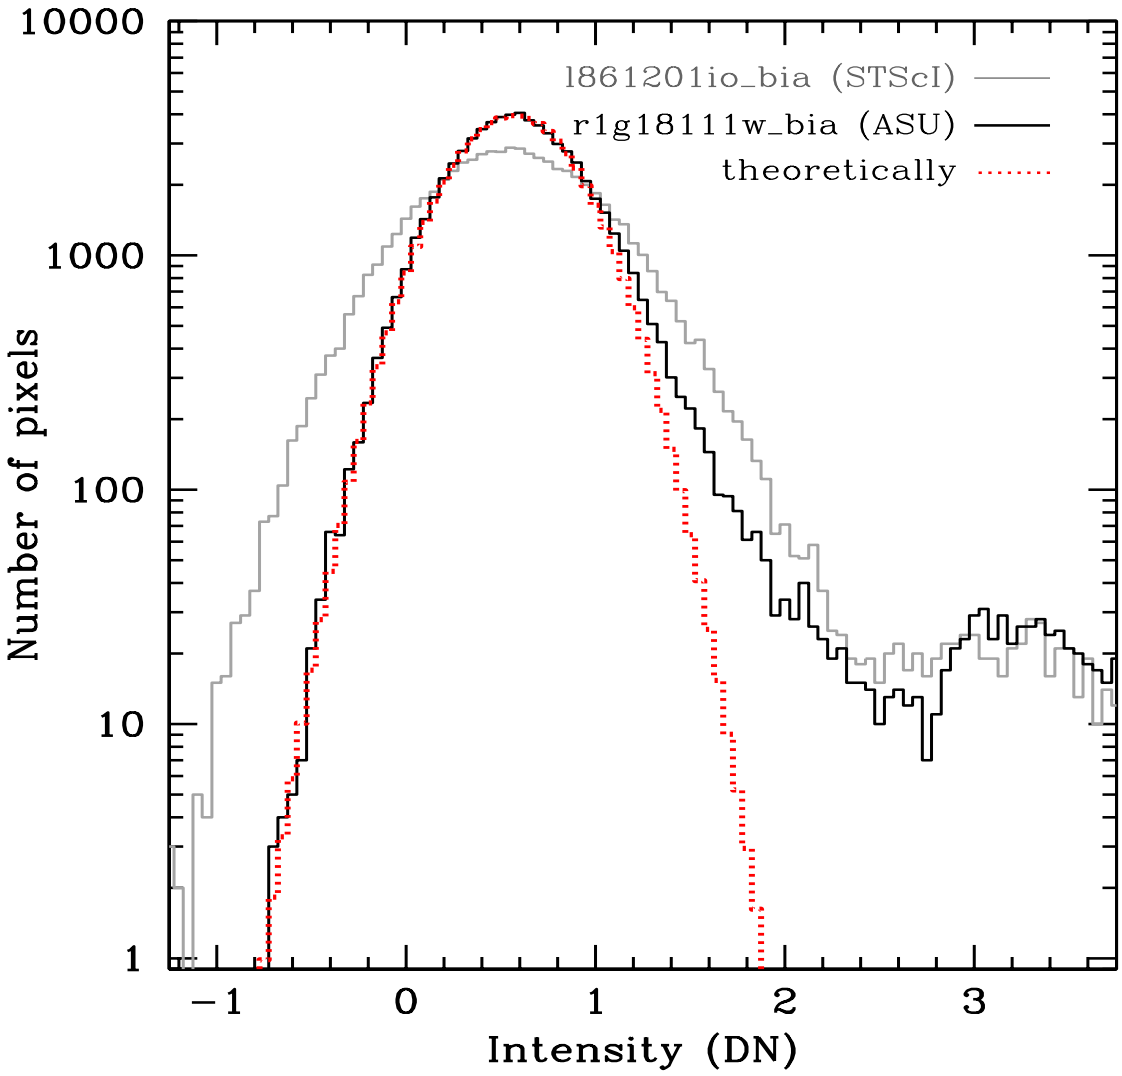

Figure 21: Examples of *HST*/STIS pixel histograms in bias reference frames. In theory, the pixel values should show a Gaussian distribution that reflects the read-noise (divided by the square root of the number of frames that were combined). Columns with higher than average bias level, and warm and hot pixels, result in a tail towards higher pixel values (*right*), but the lower half (*left*) of the histogram should be unaffected. In this particular example, a re-calibration program (labeled *ASU*; Jansen et al. 2002) resulted in a marked improvement over the bias reference frame retrieved from the *HST* Archive (*STScI*).

2.5 Gain and Read-Noise

When Silicon (i.e., a CCD) is illuminated by the 5.9 keV Kα line of a radio-active ^{55}Fe source, the charge deposited by a single energetic X-ray photon in each pixel is 1620 e^- (an extreme case of the creation of multiple electron–hole pairs). This allows one to directly measure the gain factor, $\mathcal{G}$, of the A/D conversion to a precision of better than 1% as:

$$\mathcal{G} = 1620/\mathrm{ADU} \quad [e^-/\mathrm{ADU}]$$

In practice, this is rarely done this way (usually only in the laboratory before an instrument is launched into space, or before commissioning of an expensive instrument for a major ground-based telescope).

There is another method, *Janesick's method*, where one uses the property that the noise in an image in which the signal level is high should satisfy Poissonian statistics. In that case, *photon* or *shot noise* is much larger than the read-noise, and signal, noise and gain ought to satisfy (see Fig. 22):

$$\mathcal{S}_{ADU} \equiv \frac{N_{e^-}}{\mathcal{G}} \quad (1) \qquad \text{and} \qquad \sigma_{\mathcal{S}_{ADU}} = \frac{\sqrt{N_{e^-}}}{\mathcal{G}} \quad \longrightarrow \quad \sigma^2_{\mathcal{S}_{ADU}} = \frac{N_{e^-}}{\mathcal{G}^2} \quad (2)$$

and, hence, combining (1) and (2): $\dfrac{\sigma^2_{\mathcal{S}_{ADU}}}{\mathcal{S}_{ADU}} = \dfrac{1}{\mathcal{G}}$

Of course, the assumption that any noise contributed by the detector itself or by the read-out electronics is negligible, is only an approximation. We can take such noise explicitly into account by comparing a zero-length dark exposure (a *bias* or *zero* frame), where we should only measure read-noise, with an exposure with a (fairly) uniformly high signal level, like a *dome*, *twilight* or *pupil flat* (see, e.g., Fig. 23). Since large-scale variations in signal level and pixel-to-pixel sensitivity variations would make our estimate uncertain, the trick is to use the differences of *pairs* of each, where such variations drop out.

Janesick's method will yield estimates of both gain and read-noise simultaneously, as follows. If F_1 and F_2 represent a pair of (fairly) evenly illuminated frames with plenty of signal, and if B_1 and B_2 denote a pair of bias frames, then:

$$\begin{aligned} \Delta_F &= F_1 - F_2 \\ \Delta_B &= B_1 - B_2 \end{aligned}$$

$$\mathcal{G} = \frac{(\overline{F_1} + \overline{F_2}) - (\overline{B_1} + \overline{B_2})}{(\sigma^2_{\Delta_F} - \sigma^2_{\Delta_B})} \quad [e^-/\mathrm{ADU}] \qquad \text{and} \qquad \mathcal{R} = \mathcal{G} \cdot \frac{\sigma_{\Delta_B}}{\sqrt{2}} \quad [e^-]$$

Where $\mathcal{G}$ and $\mathcal{R}$ are the gain and read-noise, averaged over all pixels of the CCD, and $\overline{X}$ denotes the mean level of frame X.

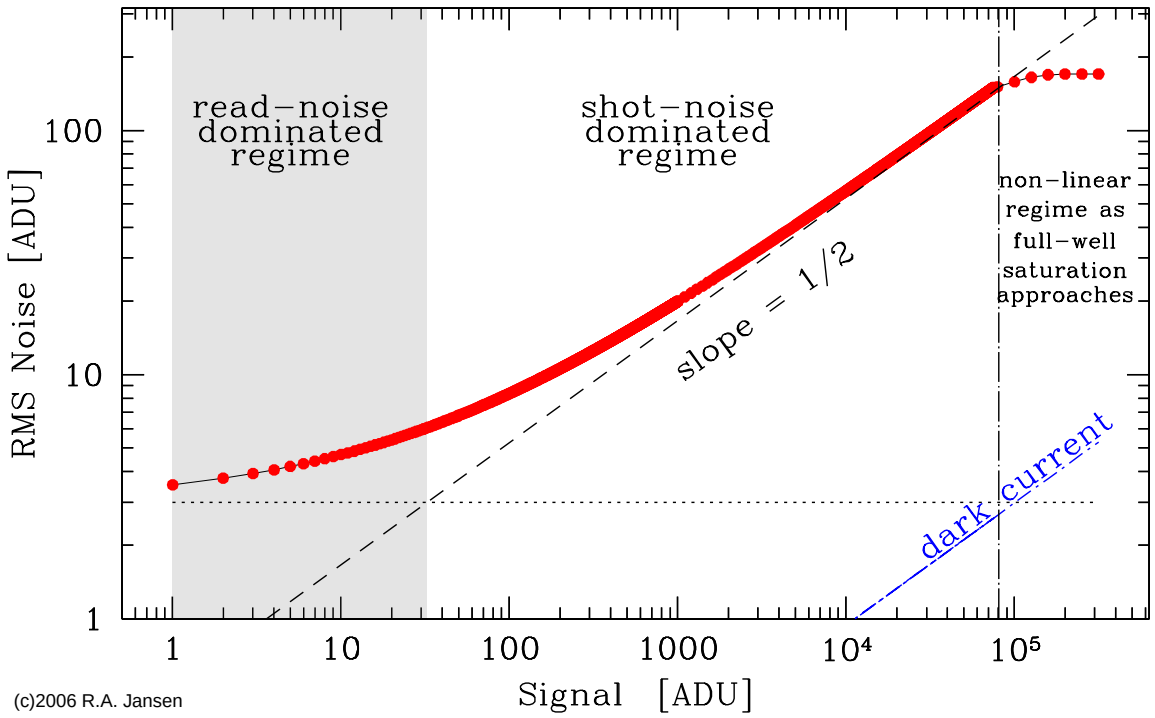

Figure 22: RMS noise as a function of signal level in a particular CCD. The read-noise sets a noise floor, to which photon noise (shot noise) increasingly adds. Once shot noise exceeds read-noise, the data are said to be shot noise limited. In the shot noise limited regime, the relation between noise and signal assymptotically approaches a slope of 1/2. If full-well capacity can be reached before A/D saturation occurs, then at very high signal levels the slope begins to deviate from 1/2. The additional noise due to dark signal is likely negligible compared to the shot noise for modern CCDs and observations through broad-band filters. It may not necessarilly be negligible compared to the read-noise, so for very faint sources in the near-UV or when observing through narrow-band filters (i.e., when the sky background is very low), the shot noise limited regime may not be reached. In the Section 3, we will discuss why the shot-noise limited regime is desirable.

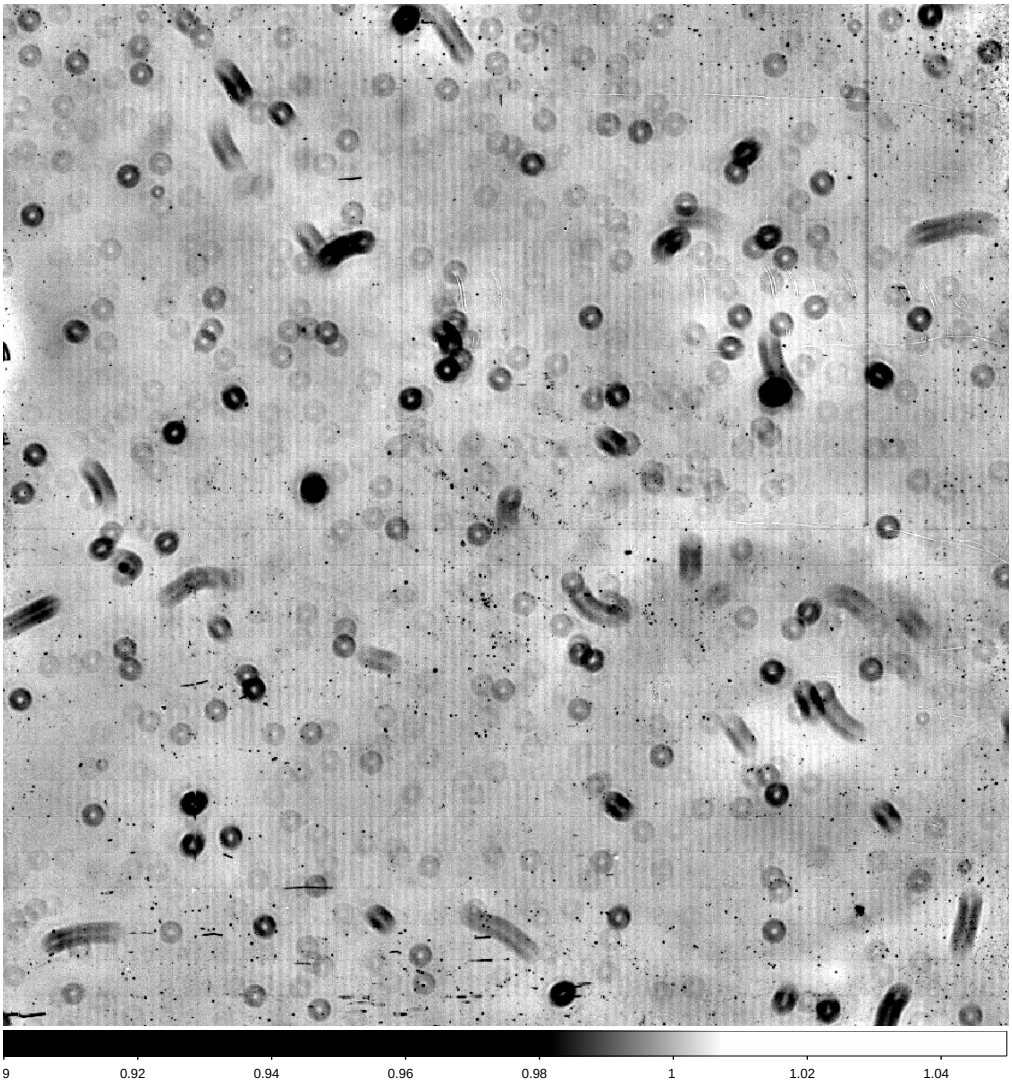

Figure 23: Example of a normalized median average of 60 bias-subtracted R-filter twilight flat frames, obtained with the CCD camera at the VATT. The range between black and white corresponds to only a $\sim$3% difference in effective response (QE). Visible are both pixel-to-pixel variations that are intrinsic to the CCD, and variations extrinsic to it. To the first category belong the pixel-to-pixel variations in the response that form the regular vertical and horizontal bands ($\pm$1%), pixels and small clusters of pixels with much lower than average response ($\lesssim$90%), and a couple of bad partial columns ($\lesssim$50%). To the second category belong the larger-scale variations due to dust particles that have settled onto the dewar window above the CCD, or on the filter (at the $\sim$50–99.9% level). Since the VATT is located within a pine forest, the vast majority of the dust particles are, in fact, semi-transparent pine pollen. If the dust particles remain stationary, division of all bias-subtracted science exposures by this normalized flat frame will correct for the variations in effective QE.

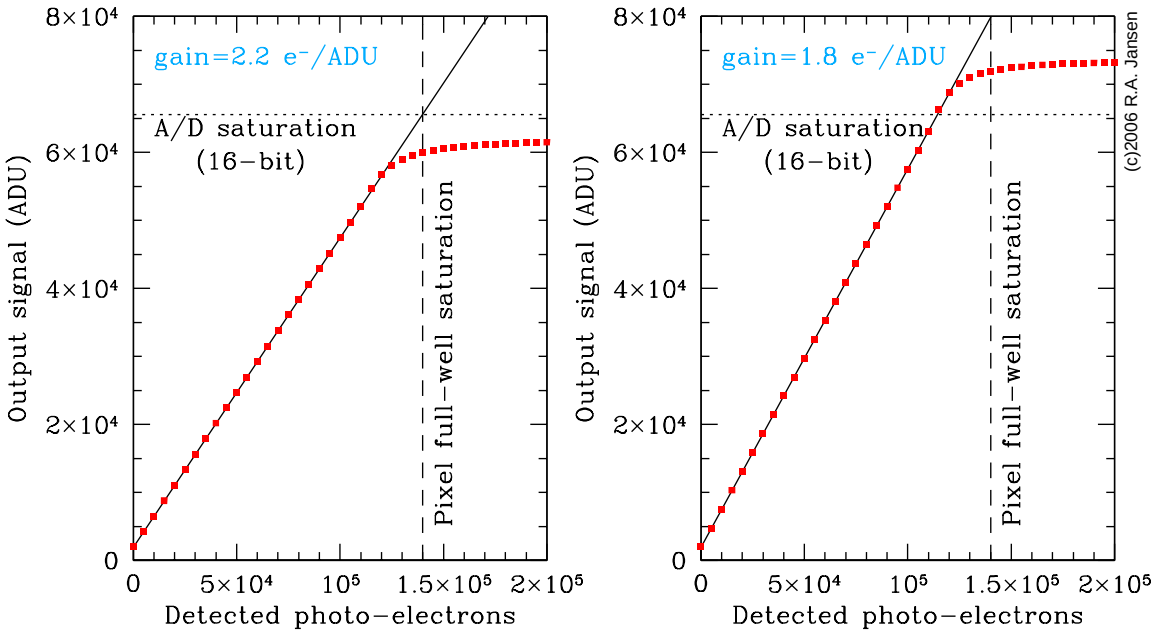

Figure 24: CCD linearity curves of a typical three-phase CCD for two different gain values. The device is linear over most of the output range from the bias level (the offset of 2000 ADU for 0 detected photo-electrons) to (*left*) ∼58,000 ADU (∼127,500 e^-), where the CCD response becomes non-linear as the pixel full-well capacity of 140,000 e^- is approached, and (*right*) 65,535 ADU, set by the saturation of the 16-bit A/D converter. In both cases, the slope of the linear part of the response curve is equal to $1/\mathcal{G}$. The gain is usually hard-wired into the CCD electronics.

2.6 Dynamic range, Saturation and Linearity

Saturation can set in two different ways:

- Pixel full-well saturation occurs when the pixel potential well can no longer confine photo-electrons to the pixel where the photon was absorbed. Shortly before full-well saturation sets in, the electrons already stored in the pixel lower the potential well to the point where additional photo-electrons may migrate deep into the substrate or even out of the pixel, resulting in a non-linear response of that particular pixel (see Fig. 24).

- A/D saturation occurs abruptly when the charge estimate presented for conversion exceeds the value that can be represented with the available number of bits per pixel. That level for an n-bit un-signed integer is $2^n - 1$ (assuming we start counting at 0). Hence, for a 12-bit A/D converter (e.g., *HST*/WFPC2) the maximum pixel value is $2^{12} - 1 = 4095$, for typical modern 15 or 16-bit A/D converters it would be 32767 or 65535, respectively.

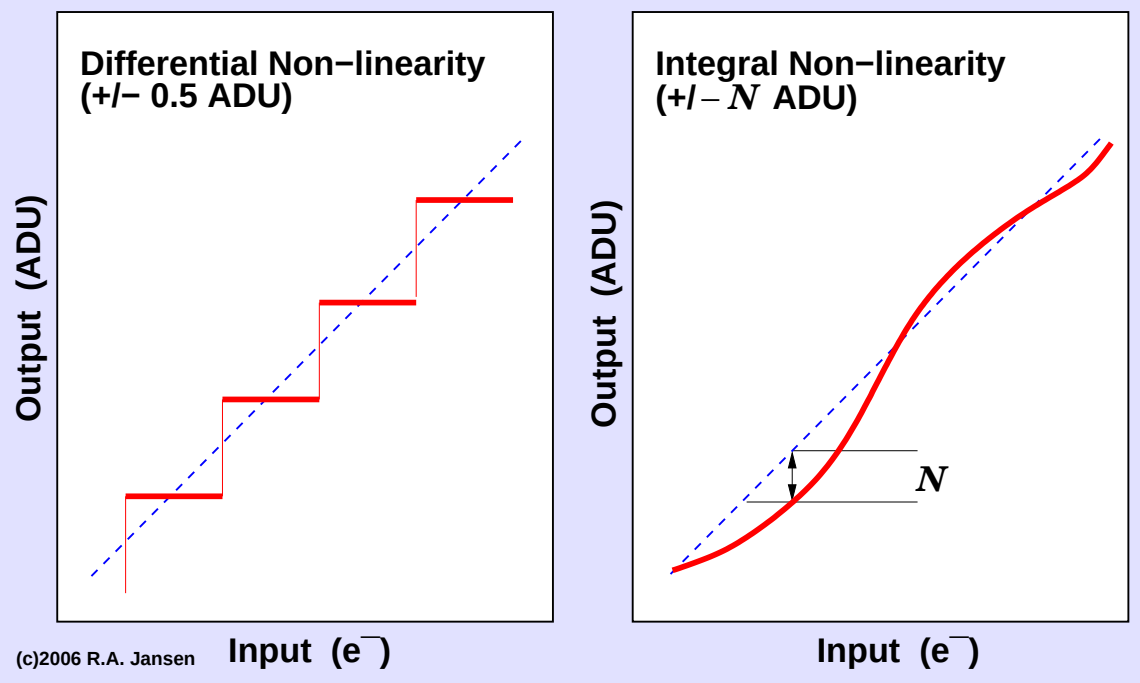

Figure 25: Illustration of the two types of CCD non-linearity. Differential non-linearity (*left*) results from the finite steps (when $\mathcal{G} > 1$) in the A/D conversion. The amplitude of this type of non-linearity is ± 0.5 ADU. Integral non-linearity (*right*) works over larger than single A/D step ranges, and can be quite complex. The magnitude of this type of non-linearity is given as the *maximum* deviation, N, with respect to a perfectly linear response.

Even when the CCD response to photons is perfectly linear at the stage of photon absorption, and storing and passing along photo-electrons, the conversion to a digital data number will result in a non-linear behavior if the gain is (significantly) larger than 1. This *differential non-linearity*, better known as *digitization noise*, is illustrated in Fig. 25*a*.

- The main gain value of *HST*/WFPC2 is $7\,e^-$/ADU, which is *larger* than the read-noise of $\sim 5.3\,e^-$ of the CCDs. Combined with the low sky background levels from space, digitization is a much more important source of noise for WFPC2 images than common in ground-based CCD images: data values below ~ 10 ADU tend to be affected, and also result in pixel distributions that satisfy neither Gaussian nor Poissonian statistics.

  △ *WFPC2 sky background values computed in UV–blue or narrow-band filter images by taking the median pixel value will be systematically in error. Median filtering to reject outlying pixels values will also be affected.*

- As read-noise of CCDs has come down over the years (from $>300\,e^-$ to typically less than $10\,e^-$ today), gain values had to come down as well, with photon-counting set-ups ($\mathcal{G} = 1$) far more common today. Digitization noise, therefore, tends to be less of a problem. This was only possible thanks to low noise A/D converters and inexpensive computer memory and disk storage, such that 16-bit/pix is now the standard rather than the exception, and even larger *bpp* values may be encountered.

A/D converters are not perfect. Some values are reported more (or less) frequently than they would be by a perfect device. This results in the second type of non-linearity, *integral non-linearity*, illustrated in Fig. 25*b*. Although the "true" data number can never be recovered, one can statistically correct for this systematic behavior. Fortunately, in most modern A/D converters, the maximum deviation, N, from the true linear response is small.

- The *HST*/WFPC2 A/D converters are known to be relatively well-behaved with $N \simeq 1.8$–2.0 ADU ($\sim$13–$14\,e^-$) for bit 12, i.e., the one set for the 2048–4095 ADU range (Baggett et al. 2002), or an integral non-linearity of $\lesssim 0.2\%$.

△ Note that the non-linearity described on page 58 and Fig. 3.9 of the textbook (Howell) is an *A/D non-linearity* encountered at high signal levels in the particular device discussed there, and *not* the intrinsic non-linearity that might occur when the charge in a pixel approaches full-well capacity that is shown here in Fig. 24! Part of the charge that is lost to a given pixel that is near full-well capacity may diffuse to neighboring pixels and be recovered for analysis when these pixels are read out. Only the fraction of photo-electrons that manage to recombine with holes deep within the substrate of a pixel are forever lost.

2.7 Transient Effects: Cosmic Rays, Cross-talk, Persistence, Hot Pixels and Bad Columns

- Cosmic rays can dump a significant amount of charge within one or more adjacent pixels, possibly saturating some, and often render the affected pixel(s) useless for subsequent scientific analysis. Thick, front-illuminated CCDs tend to be more susceptible to cosmic rays than thinned ones. Although there are several methods to detect and remove the induced signal (e.g., `la_cosmic`, by P. van Dokkum 2001), and one may sometimes decide to simply interpolate over affected pixels, cosmic rays generally necessitate splitting the desired exposure time over two or more exposures and limit any single exposure to at most ∼20 min (too avoid hits in a large fraction of pixels, and thereby reduce the chance that the same pixel gets hit in subsequent exposures).

- An example of cross-talk in the read-out electronics, when a CCD is read through more than one amplifier is visible in Fig 26, where a saturated star in one half of the CCD resulted in a ghost image at the mirrored location within the other half.

- Overexposure of bright sources can result in electrons that persist deep within the bulk substrate. Over time they may slowly migrate back up, mimicking a strongly localized increase in the dark rate, or after images of the bright source. The time constant for this *persistence* can be several hours (see Fig 26), and depends on the details of the CCD design and on how severe the overexposure was.

- As a result of defects introduced during manufacturing, or als a result of particle and radiation damage over time, some pixels will show a much larger dark current than the one characteristic for a particular CCD. Such pixels (or even entire pixel columns) are known as *warm* or *hot pixels*, depending on the severity. By warming up the device to room temperature (a process called *annealing*), a certain fraction of warm/hot pixels may return to normal.

- Unresponsive (dead) pixels and bad/dead columns typically also result at the manufacturing stage. Bad pixels can act as charge traps or shorts that drain charge away into the substrate or that prevent the transfer of charge from any pixels along a column located further away from the serial output register. As a result a (portion of a) column in the output image will appear dark. Modern science-grade CCDs show only few of these cosmetic defects.

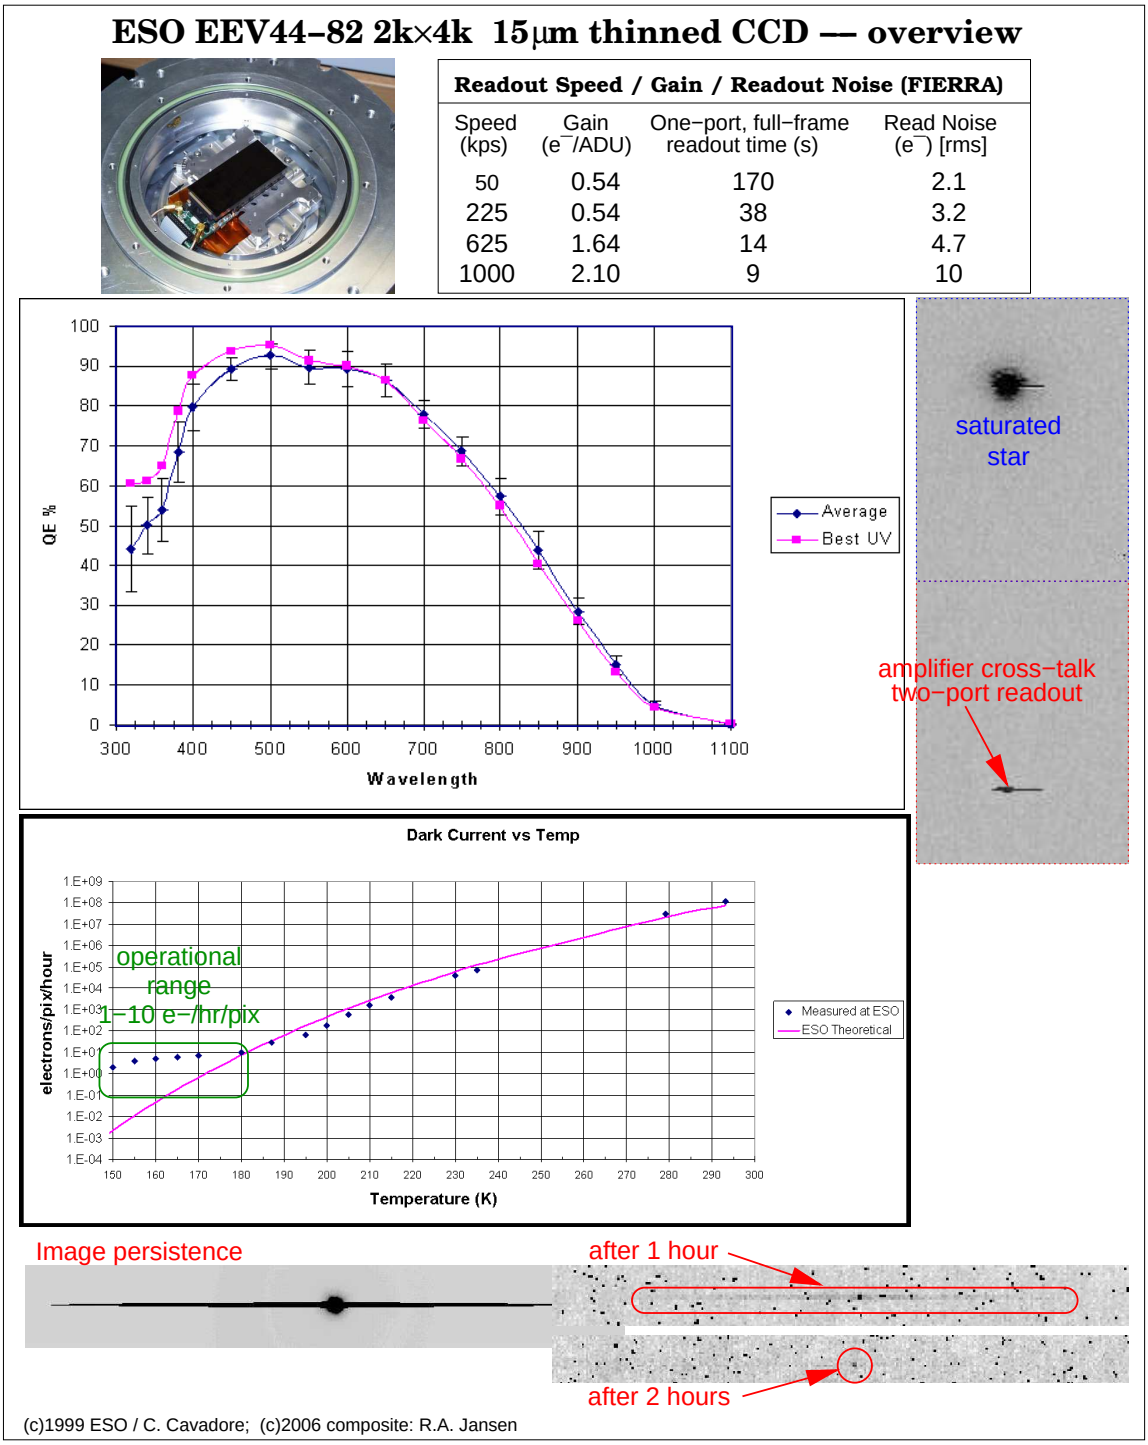

Figure 26: Performance overview of an EEV 44-82 2k×4k CCD delivered to the European Southern Observatory.

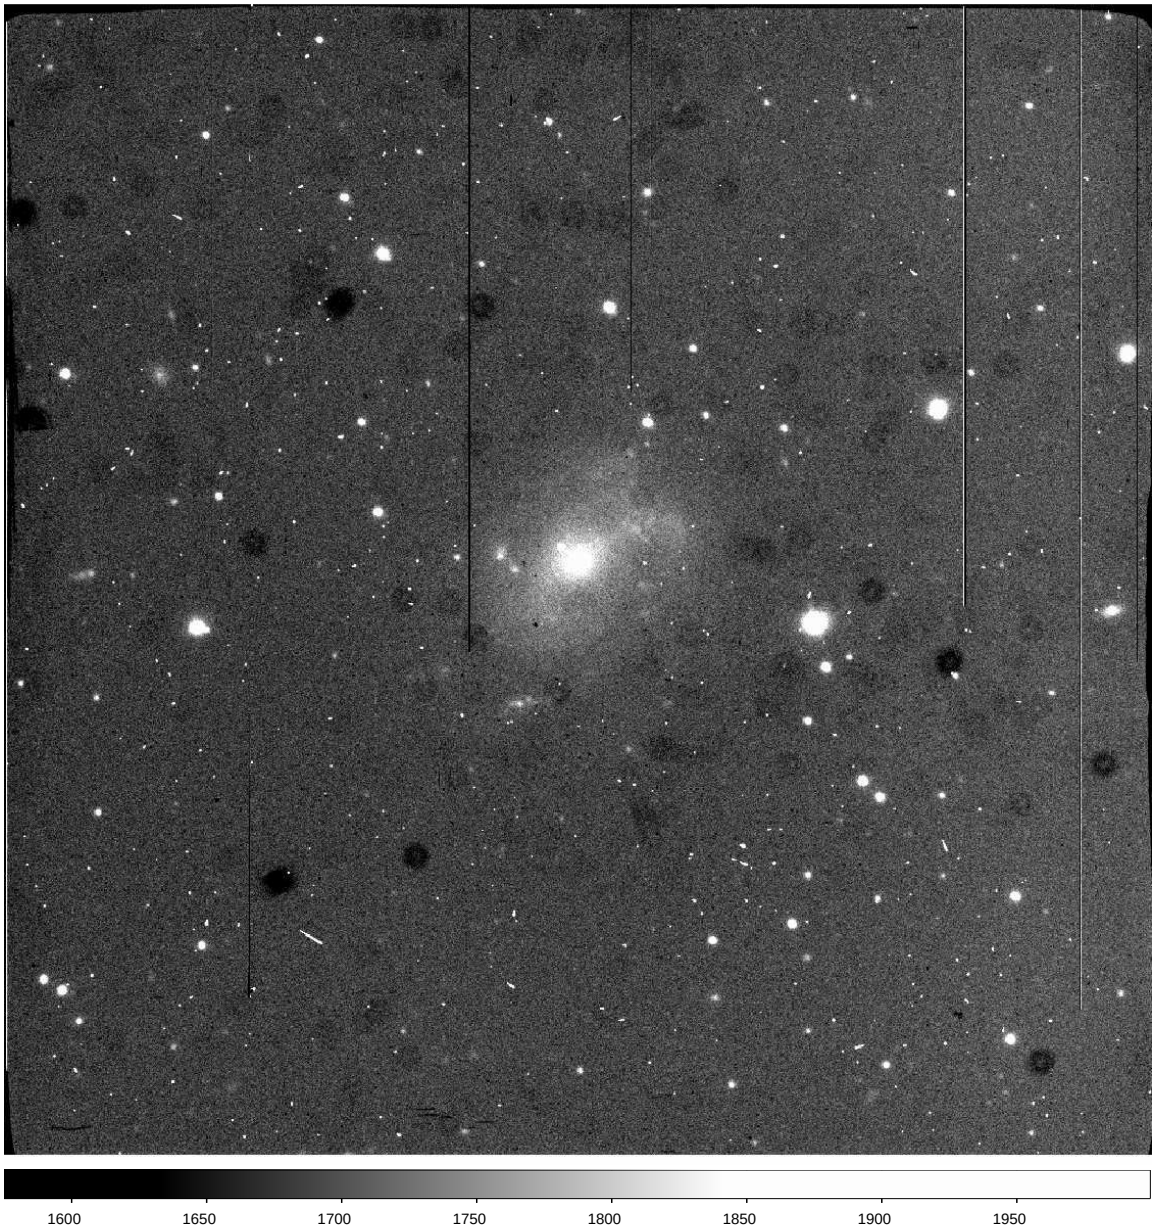

Figure 27: Example of a raw 1200 sec Hα image of UGC 10445, obtained with the CCD camera at the VATT in Sep 2003 (Jansen & Tamura, in prep.), that features several of the transient effects and defects discussed. Several bad columns are seen (some hot, some dead), and the image is riddled with cosmic ray hits. Since cosmic ray hits do not bear the imprint of the PSF, they tend to be sharper than genuine point sources (e.g., stars) whenever the PSF is properly sampled. This example also demonstrates quite effectively why it is important to flat field the data.

3 Astronomical Signals and Noise

3.1 Distributions: Poisson vs. Gauss

Read-noise closely follows the *Gaussian* or *normal* distribution associated with sampling measurements in a near-equilibrium electronic circuit. *Shot noise*, i.e., the noise associated with the arrival of discrete energy or charge packets instead follows a *Poisson* distribution.

△ *In the textbook, p. 73 and 74, the use of the term* **shot-noise** *to describe the read-noise, is confusing. On page p. 45 and 46 the read-noise is described as different from photon-noise. Here, the terms photon-noise and shot-noise have been used interchangable for the noise associated with the arrival of photons.*

An important property of the Poisson distribution is, that its *standard deviation* equals the square root of the measured number (i.e, detected photo-electrons in our case). For an average score of m, the probability of measuring n when the values of n are distributed according to a Poisson distribution is given by:

$$P_n = \frac{m^n \cdot e^{-m}}{n!} \qquad \text{, with } m = \text{constant}$$

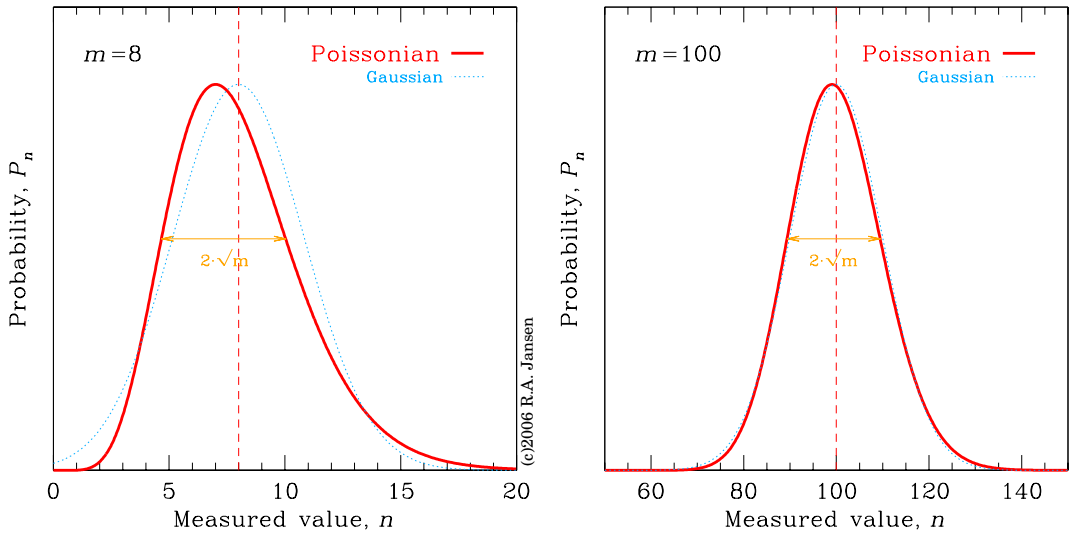

Figure 28: Examples of Poisson distributions. (*left*) For $m = 8$, the distribution is markedly skewed compared to a Gaussian distribution with the same average score m. An important property of the Poisson distribution is that the interval between $-1\,\sigma$ and $+1\,\sigma$ is exactly equal to twice the $\sqrt{m}$. (*right*) For large m, such as $m = 100$ in this example, the Poisson distribution becomes indistinguishable from the normal distribution. *Note: for visualization purposes we here plot the mathematical extension of the discrete Poisson statistics (integer n) to continuous functions!*

First, let us verify that the sum of the probabilities over all n is indeed equal to 1, and that the average score, $\bar{n}$, is indeed equal to m:

$$\sum_n P_n \;\equiv\; \sum_n \frac{m^n \cdot e^{-m}}{n!} \;=\; e^{-m} \cdot \sum_n \frac{m^n}{n!} \;=\; e^{-m} \cdot e^m \;=\; 1$$

The average of n scores is given by:

$$\bar{n} \;\equiv\; \frac{\sum_n P_n \cdot n}{\sum_n P_n} \;=\; \sum_n \frac{m^n \cdot e^{-m}}{n!} \cdot n \;=\; m \cdot \sum_n \frac{m^{n-1} \cdot e^{-m}}{(n-1)!} \;=\; m$$

Lastly, the variance can be computed as:

$$\begin{aligned}
\sigma^2 \;&\equiv\; \frac{\sum_n (P_n \cdot (n-\bar{n})^2)}{\sum_n P_n} \;=\; \sum_n P_n \cdot (n^2 - 2nm + m^2) \\
&=\; \sum_n P_n \cdot n^2 \;-\; 2m \cdot \sum_n P_n \cdot n \;+\; m^2 \;=\; -m^2 \;+\; \sum_n \frac{m^n \cdot e^{-m}}{n!} \cdot n^2 \\
&=\; -m^2 \;+\; m\, e^{-m} \cdot \sum_n \frac{m^{n-1}}{(n-1)!} \cdot n \;=\; -m^2 \;+\; m\, e^{-m} \cdot \sum_l \frac{m^l}{l!} \cdot (l+1) \\
&=\; -m^2 \;+\; m \cdot \left[\left(\sum_l \frac{m^l \cdot e^{-m}}{l!} \cdot l \right) + \left(\sum_l \frac{m^l \cdot e^{-m}}{l!} \right) \right] \\
&=\; -m^2 + m^2 + m \;=\; m
\end{aligned}$$

And, hence, we find that for the Poisson distribution the standard deviation is equal to the square root of the average score: $\sigma = \sqrt{m}$.

▷ When we talk about an $x\,\sigma$ deviation or an $y\,\sigma$ result, we mean a deviation of x times the standard deviation σ (i.e, $x\sqrt{m}$), or a result that is significant at the probability of 1 minus that of obtaining a y times σ deviation purely by chance:

$$1 - P\left(n < y\sqrt{m}\right) \qquad \text{or} \qquad 1 - P\left(|n| < y\sqrt{m}\right)$$

3.2 Signal and Noise: The CCD Equation

In science, publishing a measured value is ***meaningless*** unless the associated *measurement uncertainty* (also commonly referred to as the *error* on that measurement) is given as well. Although the specified number of significant digits may give a clue as to the precision of a measurement, it should be considered *bad practice* to only rely on this. For example, to avoid round-off errors in subsequent computations, one may want to specify one digit more than strictly merited, and sometimes the implied uncertainty of half the range of the least significant digit does not accurately represent the actual measurement error.

Furthermore, in order to obtain new observations with a telescope, one has to write an observing proposal. A proposal without a thorough discussion of (a) the minimum *signal-to-noise ratio* (S/N) required for a definite result, and (b) the S/N that a particular telescope/instrument configuration will give you in a given amount of time, is rarely awarded time.

It should, therefore, be clear that an evaluation of the measured (or expected) S/N is an important topic to address. Also, by comparing the theoretically expected S/N with that actually obtained (e.g., Fig 29), one may catch problems (i.e, user errors) in the data processing.

Earlier, we discussed how to compute the read-noise and gain values for a CCD, and we have just verified several useful properties of the Poissonian distribution which is applicable to photon-noise, so let us see which noise terms would contribute to the uncertainty in the measured number of electrons in a single pixel.

- If we denote the portion of the signal that originated from the astronomical source of interest by $S_\star$ (expressed in e^-), then the Poisson statistics give the $1\,\sigma$ uncertainty, or noise, on that signal as $\sqrt{S_\star}$.
- Usually, there is another signal present due to photons: the sky background. If we denote that portion of the signal by S_S, then the associated noise will be $\sqrt{S_S}$.

During an integration, we also collect signal that is *not* associated with arriving photons:

- The *dark signal*, S_D, due to thermal electrons from the bulk Si substrate builds up linearly with the integration time t (whether or not the CCD is actually exposed to light or not), i.e., $S_D = t \cdot dc$. The dark current (or *dark rate*) dc is usually expressed in units of e^-/pix/hour (where the "/pix" is a dummy unit, to remind us that, on average, every pixel will collect the same amount of dark signal). Here, if t is expressed in seconds, dc should be in e^-/sec. The noise associated with the discrete dark signal also satisfies Poissonian statistics and will be $\sqrt{S_D} = \sqrt{t \cdot dc}$.

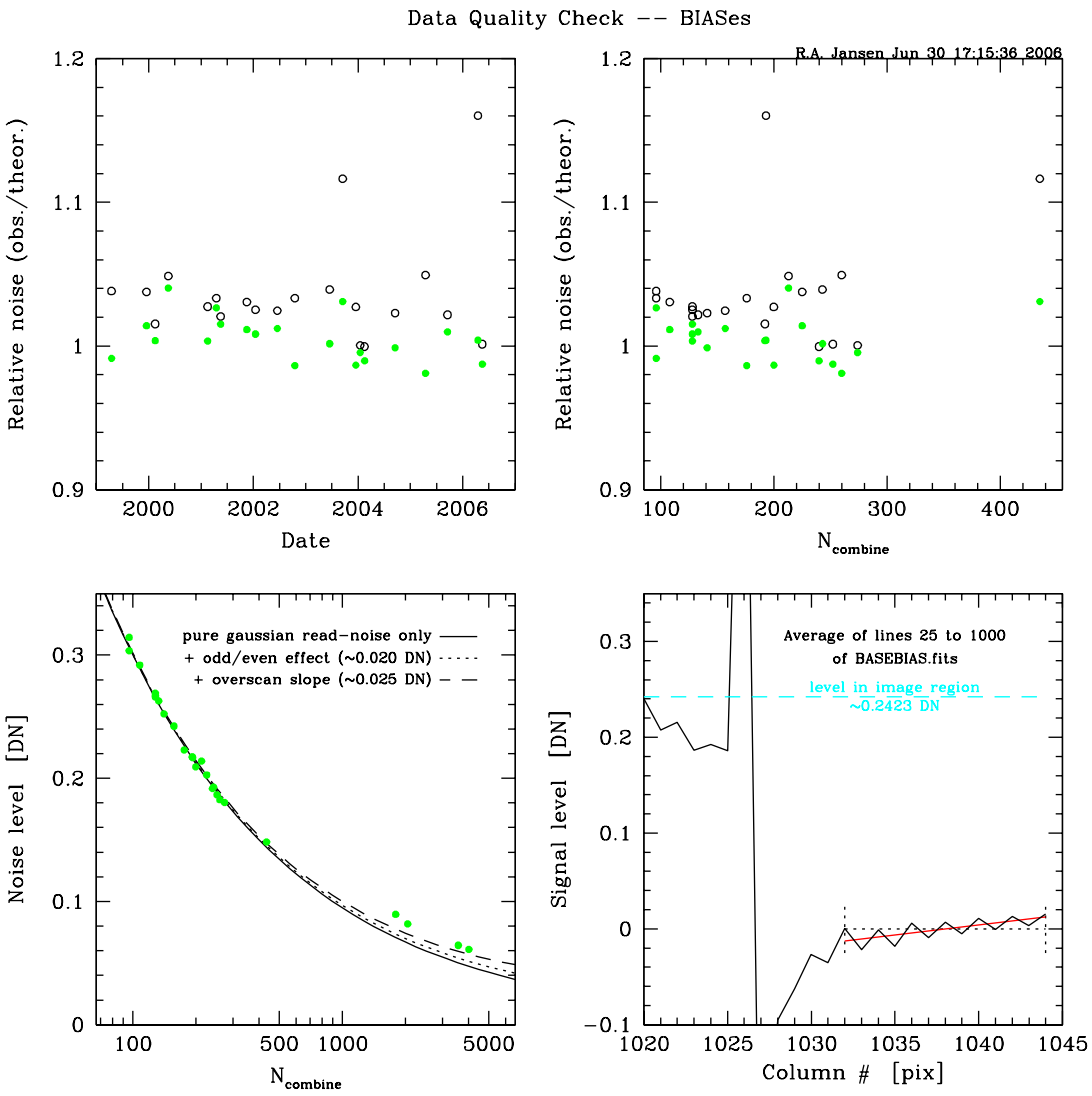

Figure 29: Example of a comparison of theoretically expected and actually observed S/N for the simplest case: the combination of many bias frames obtained at the VATT, Mt. Graham, between 1999 and 2006. Note also in the lower right panel, that the level in the overscan region (which formally starts at column #1025) is not flat: only pixels in columns 1032 through 1044 are reliable for measuring the overscan level. Furthermore, there is a small but distinct offset of ∼0.24 ADU between the bias level measured in the overscan region (virtual pixels) and that measured in the image region (actual, illuminated pixels). Such offsets are common in ageing CCD devices and need to be taken into account during the data processing.

And finally we have a noise term that has no signal associated with it at all:

- The *read-noise*, $\mathcal{R}$ (in e^-), adds directly to the total noise in the pixel under consideration.

Upon read-out of the charge in the CCD pixel, we have no way to distinguish electrons originating from the various sources. We just measure the total signal. Hence, the associated noise is also grouped together as:

$$N \;=\; \sqrt{S_\star + S_S + t \cdot dc + \mathcal{R}^2}$$

Note that the read-noise — since it should add linearly to each pixel — enters as the read-noise *squared* underneath the square root sign.

The S/N ratio for this single-pixel measurement can then be expressed as:

$$\frac{S}{N} \;=\; \frac{S_\star}{\sqrt{S_\star + S_S + t \cdot dc + \mathcal{R}^2}} \tag{1}$$

If we rather have an expression for the signal-to-noise in terms of the measured signals in ADU (but with the dark rate and read-noise still expressed in their usual units, since they are not measured from the pixel under consideration), then we need to explicitly include the CCD gain factor:

$$\frac{S}{N} \;=\; \frac{S_\star^{ADU} \cdot \mathcal{G}}{\sqrt{S_\star^{ADU} \cdot \mathcal{G} + S_S^{ADU} \cdot \mathcal{G} + t \cdot dc + \mathcal{R}^2}}$$

Of course, in real life, one often measures the astronomical signal collected within more than one pixel. If the astronomical signal is measured within $n_{\rm pix}$ pixels, then we need to modify Eq. (1) as follows:

$$\frac{S}{N} \;=\; \frac{S_\star}{\sqrt{S_\star + n_{\rm pix} \cdot (S_S + t \cdot dc + \mathcal{R}^2)}} \tag{2}$$

This equation is informally referred to as *"The CCD Equation"*. Note, that if the number of pixels containing the signal is large, the contribution of the read-noise may grow large because it enters linearly, while the other noise terms enter only as the square root of their signals.

In the above examples, one would have to gain knowledge of the amount of signal in the sky background through some magical means. If we explicitly include the uncertainty in determining the sky background level by taking the mean over $n_{\rm sky}$ background pixels, the equation is modified as follows:

$$\frac{S}{N} \;=\; \frac{S_\star}{\sqrt{S_\star + n_{\rm pix} \cdot \left(1 + \frac{n_{\rm pix}}{n_{\rm sky}}\right) \cdot (S_S + t \cdot dc + \mathcal{R}^2)}} \tag{3}$$

Note, that $n_{\rm pix}/n_{\rm sky}$ is large when $n_{\rm sky}$ is small, while for $n_{\rm sky} \gg n_{\rm pix}$ the term $(1 + n_{\rm pix}/n_{\rm sky})$ approaches 1. This properly reflects the notion that one, generally, needs to average of many background pixels to obtain an accurate estimate of the sky level.

If we, lastly, also include the effects of digitization noise as a result of the A/D conversion process, then we arrive at the complete CCD Equation:

$$\frac{S}{N} = \frac{S_\star}{\sqrt{S_\star + n_{\rm pix} \cdot \left(1 + \frac{n_{\rm pix}}{n_{\rm sky}}\right) \cdot (S_S + t \cdot dc + \mathcal{R}^2 + \mathcal{G}^2\sigma_f^2)}} \tag{4}$$

Here, σ_f is the $1\,\sigma$ error introduced by the digitization process. Although somewhat dependent on the details of the design of the A/D electronics, a likely value of σ_f may be $\sim$0.289 (Merline & Howell 1995). If the gain is large, this term may become significant. For photon-counting operation ($\mathcal{G} = 1$), it is likely very much smaller than the read-noise term.

Example: Let us compute the S/N obtained for an astronomical source in a CCD image with an exposure time $t = 300\,$sec. The CCD has a gain of $\mathcal{G} = 2\,e^-$/ADU, read-noise $\mathcal{R} = 6\,e^-$ RMS, and dark rate of $10\,e^-$/pix/hour. The sky background level is measured to be $S_S = 604\,$ADU, by averaging the pixels values in an annulus of blank sky around the object of interest with an inner radius of 20 pixels and outer radius of 30 pixels. The signal due to the source, measured in an aperture of radius 5 pixels, is found to be $S_{\star+S} - S_S = 6,487\,$ADU.

$$\frac{n_{\rm pix}}{n_{\rm sky}} = \frac{\pi \cdot 5^2}{\pi \cdot (30^2 - 20^2)} = 0.05 \quad \longrightarrow \quad n_{\rm pix} \cdot \left(1 + \frac{n_{\rm pix}}{n_{\rm sky}}\right) \simeq n_{\rm pix} = \pi \cdot 5^2$$

The gain is sufficiently close to the photon counting case, that we do not need to worry about the digitization noise $\mathcal{G}^2\sigma_f^2$.

$$dc = 10\,e^-/\text{hour} = \frac{10}{3600}\,e^-/\text{sec}$$

We then have:

$$\frac{S}{N} = \frac{6,487 \cdot 2}{\sqrt{(6,487 \cdot 2) + \pi \cdot 5^2 \,\left(604 \cdot 2 + 300 * \frac{10}{3600} + 6^2\right)}} \;\simeq\; 39$$

Note, that this is quite a bit smaller than $S_\star/\sqrt{S_\star} \simeq 114$. Had we omitted the additional noise terms in the CCD Equation, we would have grossly overestimated the S/N ratio.

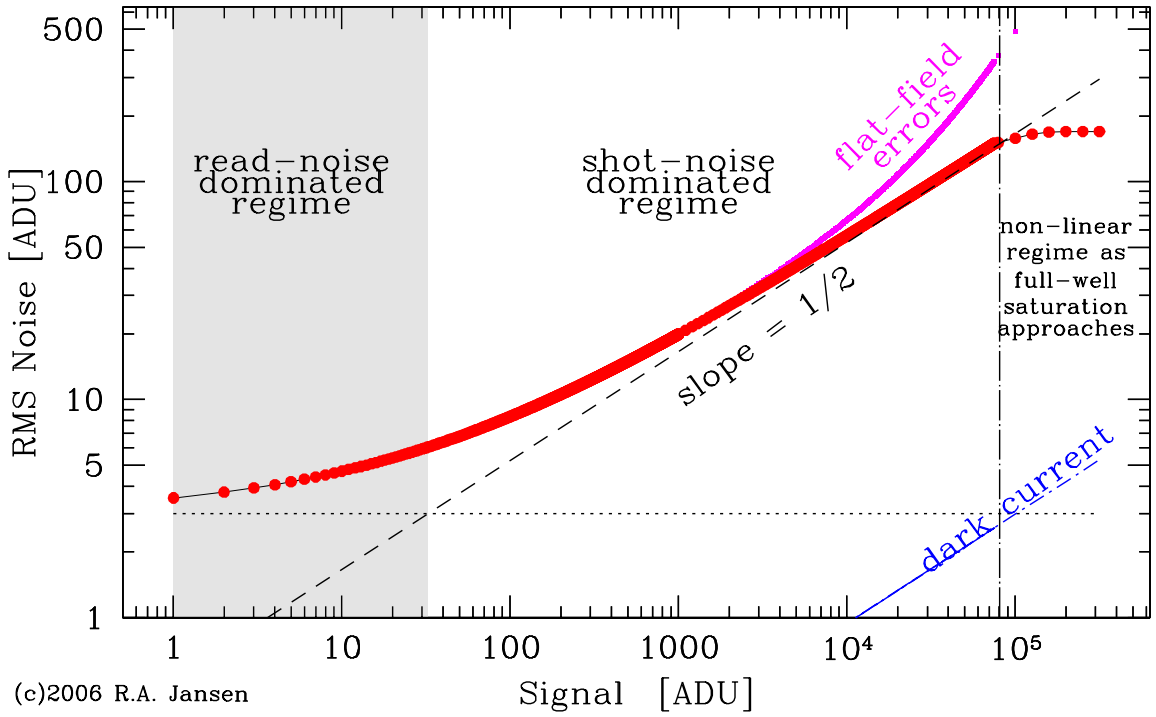

Figure 30: RMS noise as a function of signal level in a particular CCD. The read-noise sets a noise floor, to which shot noise increasingly adds. Once shot noise (photon noise) exceeds read-noise, the data are said to be shot noise limited. In the shot noise limited regime, the relation between noise and signal assymptotically approaches a slope of 1/2. If full-well capacity can be reached before A/D saturation occurs, then at very high signal levels the slope begins to deviate from 1/2. The additional noise due to dark signal is likely negligible compared to the shot noise for modern CCDs and observations through broad-band filters. It may not necessarilly be negligible compared to the read-noise, so for very faint sources in the near-UV or when observing through narrow-band filters (i.e., when the sky background is very low), the shot noise limited regime may not be reached. When the signal is shot-noise limited due to the signal from the astronomical source itself, the CCD Equation reduces to $S/N \simeq S_\star/\sqrt{S_\star} = \sqrt{S_\star}$. In this case, we speak of a "bright" object, in all other cases of a "faint" object. When the signal is shot-noise limited due to the sky, the noise scales as $\sqrt{S_\star + n_{\rm pix} \cdot S_S}$. The worst case is where the noise is dominated by the read-noise: $S/N \simeq S_\star/\sqrt{n_{\rm pix} \cdot \mathcal{R}^2} = S_\star/(\mathcal{R}\sqrt{n_{\rm pix}})$.

Computing errors on magnitudes

Using the identity $\sigma \equiv N/S$, we can now express Eq. (4) as the standard deviation on the measured magnitude:

$$\sigma_{\rm mag} \simeq \frac{1.0857\,\sqrt{S_\star + p}}{S_\star} = 1.0857\left(\frac{1}{S/N}\right) \qquad [\rm mag] \tag{5}$$

where p represents the entire term under the square root sign in Eq. (4) other than $S_\star$. The scale factor $1.0857 = 2.5/\ln 10$ comes from the definition of magnitude (factor 2.5), and from computing the error on a decimal logarithm (factor $\ln 10$).

In our example above, $S/N = 39$ and so: $\sigma_{\rm mag} \simeq 1.0857\left(\frac{1}{39}\right) = 0.028\,\rm mag$

In the above Eq. (5), we used an "approximately equal to" ($\simeq$) notation rather than an equals sign ($=$). Let's consider this for a bit.

Per definition: $m = -2.5 \cdot \log(S) + zp$

Therefore, errors of $+1\,\sigma$ and $-1\,\sigma$ in source signal S mean:

$$\begin{aligned}
\sigma_m^+ \;=\; m^+ - m \;&=\; -2.5 \cdot \log(S+N) + 2.5 \cdot \log(S) \\
&=\; -2.5 \cdot \log(\frac{S+N}{S}) \;=\; -2.5 \cdot \log(1 + \frac{1}{S/N}) \\
\sigma_m^- \;=\; m^- - m \;&=\; -2.5 \cdot \log(S-N) + 2.5 \cdot \log(S) \\
&=\; -2.5 \cdot \log(\frac{S-N}{S}) \;=\; -2.5 \cdot \log(1 - \frac{1}{S/N})
\end{aligned}$$

So we find that $|\sigma_m^+| \neq |\sigma_m^-|$ and that "upper" errors (toward *brighter* fluxes or *smaller* magnitudes) are therefore smaller than the "lower" errors (toward *fainter* fluxes or *larger* magnitudes)! This behavior is clearly illustrated in Fig. 31.

- For relatively large magnitude errors (low S/N), the upper and lower errors become progressively different in size: *magnitude errors are not symmetric.*

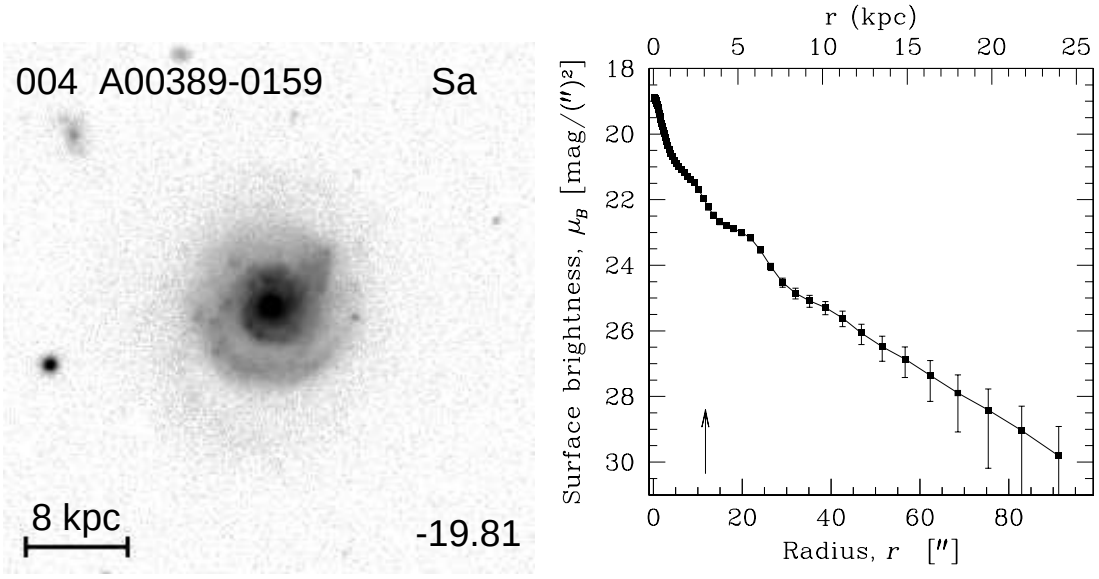

Figure 31: (*a*) *B* filter CCD image obtained at the SAO 48" telescope (Mt. Hopkins, AZ) of Sa galaxy A00389-0159 (UGC 439), observed as part of the Nearby Field Galaxy Survey (Jansen et al. 2000a,b). (*b*) Radial surface brightness profile measured from this image. Note that at low S/N (larger radii), the upper and lower magnitude errors on the profile become increasingly asymmetric.

A final important observation is that the signals $S_\star$ and S_S scale linearly with the exposure time t. The S/N ratio in the shot-noise dominated regime (see Fig. 30) will therefore scale as $t/\sqrt{t} = \sqrt{t}$.

- This implies that, to obtain a S/N that is *twice* as good, one has to integrate *4 times* as long. In the read-noise dominated regime, the S/N will increase much slower — if $n_{\rm pix} \cdot \mathcal{R}^2 \gg (S_\star + n_{\rm pix} \cdot S_S)$, the S/N remains nearly constant and < 1, no matter how long you integrate!

4 Data Storage and Transfer — FITS '101'

The standard file format for storage and exchange of astronomical data — and nowadays even direct image processing — is the Flexible Image Transport System (FITS). In the early 1980s, Wells, Greisen & Harten (1981) developed FITS to end the babylonic mess of data tapes that were unreadable on systems with different machine architectures (byte order; number of bits per byte) than the one on which they were written, and of data files that did not document their content, nor even their basic structure (i.e., the number of dimensions, the number of pixels per dimension, and the number of bits per stored pixel value). The FITS format subsequently found rapid acceptance within the astronomical community, and proved flexible enough to accommodate far more complex data structures than a simple 2-D image.

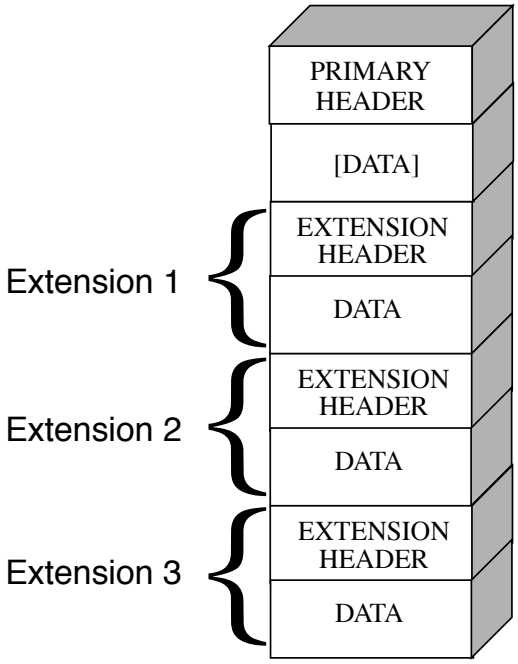

Figure 32: Schematic representation of the FITS file structure. *Courtesy: Pavlovsky et al. 2004*, *HST*/ACS Data Handbook.

A FITS file is composed of one Primary Header and Data Unit (HDU), optionally followed by an arbitrary number of extensions, each again consisting of a Header and Data Unit. A Data Unit, particularly the primary one, may in fact be empty, leaving the Primary Header Unit for storage of information that applies to all following extensions.

The simplest FITS file structure contains only the Primary Header and Data Unit, where the Header unit, now simply referred to as *the FITS header*, describes at least the essential attributes of the image stored in the Primary Data Unit that follows. Unless you observe with a mosaic CCD camera, you are likely to receive your data in this simplest flavor of FITS.

4.1 The standard FITS header

The FITS header consists of one or more plain ASCII *records* of 2880 bytes, each of which contains 36 80-character header entries (which for historical reasons are referred to as *card images*). The end of the header unit is marked by the appearance of the string `END` in the first three characters of the final 80-character card image.

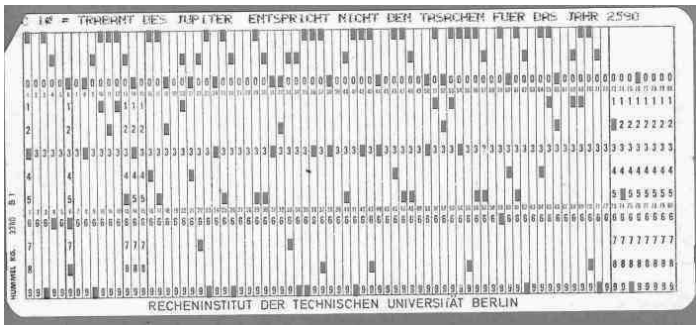

Figure 33: Standard IBM 80-column computer punch card, from which the term *card image* derives.
Courtesy: Wikipedia

This $n\times(36\times80)$-character format means that one can display and inspect the contents of the header of a FITS file using `more`, `less` or `page` in a standard 80 character wide console window (such as, e.g., `xterm`). Beware, however, that the binary data that immediately follows the header may cause your console window to get confused. If the `CFITSIO` library (by W.D. Pence) is installed on your machine you can display just the header with program `listhead`, which is shipped with this library.

Each header card image consists of a *keyword*, an optional *value*, and an optional *comment*. Blank lines (e.g., for formatting) are ignored.

KEYWORD =␣ NUMERIC_VALUE ␣/ COMMENT_STRING
KEYWORD =␣ LOGICAL_VALUE ␣/ COMMENT_STRING
KEYWORD =␣'STRING_VALUE' ␣/ COMMENT_STRING
COMMENT ␣ CHARACTER_STRING
HISTORY ␣ CHARACTER_STRING
END

A keyword is a left-justified, 8-character, blank padded, ASCII string with no embedded blanks. Only digits, hyphen, underscore, and upper-case characters are allowed. The 9-th character of the card image is either an equals sign (=), or a blank character (denoted ␣ above), in which case the card image will be interpreted as a comment line. The 10-th character must be a blank character, when the 9-th character is a =.

The value field, when used, may be string-, numeric-, or logical- (boolean-) valued. If a string, it will be left-justified and enclosed between single forward quotes. If numeric (integer, float, exponential notation in either single or double precision formats) or

boolean (T for *true* or F for *false*, without quotes), it will be right-justified to (usually) column 30 or 31.

An optional comment may follow the value and will be separated from that value by a blank space character and a forward slash (/). Although not required, the slash is usually followed by another space character. Even though recent revisions of the FITS standard do implement a provision for continuation lines, plain FITS header card images should not exceed the 80-character limit.

For a simple 2D CCD image, there are 7 mandatory header keywords to describe the main attributes of the image data that may follow in the Primary Data Unit:

- SIMPLE
- BITPIX
- NAXIS — a simple ("*flat*") 2-D CCD image has NAXIS = 2
- NAXISn, with $n = 1, \ldots,$ NAXIS — for a flat CCD image: NAXIS1 and NAXIS2
- EXTEND
- END

Very useful additional keywords to transform the actually stored digital data back to the originally recorded ADUs, are three additional keywords:

- BSCALE
- BZERO
- BUNIT — optional, to describe the physical units in which the values in the array, after application of BSCALE and BZERO are expressed.

To store within the FITS header at least some zero-th order information regarding pointing and plate scale, the following keywords are used:

- CRVALn — RA/Dec at reference pixel (in degrees)
- CRPIXn — reference pixel
- CDELTn — increment per pixel (i.e., pixel size in degrees)
- CTYPEn — coordinate system type (e.g., 'RA---TAN' and 'DEC--TAN')

Second order *world coordinate system* (WCS) keywords are:

- WCSDIM — WCS dimension
- CDi_j — diagonal element of the scale matrix
- LTVi_j — translation with respect to the original image
- LTMi_j — rotation with respect to the original image

When geometric distortions are a severe (such as for the off-axis *HST*/ACS cameras; see Figs. 36 and 37), additional terms are required for which the A_ORDER $\times$ B_ORDER distortion matrix keywords A_i_j and B_i_j were introduced. Note, however, that not every program that can display a FITS image and track the (RA,Dec) of the (x, y) pixel coordinates knows how to interpret these higher order terms — the RA and Dec values returned may not be very accurate.

```
SIMPLE  =                    T / Fits standard
BITPIX  =                   16 / Bits per pixel
NAXIS   =                    2 / Number of axes
NAXIS1  =                 1044 / Axis length
NAXIS2  =                 1044 / Axis length
EXTEND  =                    F / File may contain extensions
ORIGIN  = 'NOAO-IRAF FITS Image Kernel July 2003' / FITS file originator
DATE    = '2004-09-15T01:55:35' / Date FITS file was generated
IRAF-TLM= '16:03:21 (31/10/2004)' / Time of last modification
OBJECT  = 'SKY     '           / Name of the object observed
IMAGETYP= 'object  '           / object, dark, bias, flat, etc.
OBSERVAT= 'MTGRAHAM'           / observatory
TELESCOP= 'VATT    '           / telescope name
INSTRUME= 'CCD     '           / instrument
DETECTOR= 'ccd26   '           / detector
DATE-OBS= '2004-09-15'         / date (YYyy-mm-dd) of observation
UT      = '01:55:35'           / universal time
ST      = '18:13:27.64'        / sidereal time
RA      = '17:56:19.20'        / right ascension
DEC     = '+06:38:26.95'       / declination
EPOCH   =                2000. / epoch of ra and dec
ZD      = ' 26.1835'           / zenith distance (deg)
AIRMASS =                1.114 / airmass
EXPTIME =                  30. / actual integration time (sec)
DARKTIME=                  31. / total elapsed time (sec)
ROTANGLE= '0.000000'           / rotation angle
TELFOCUS= '-118.00 '           / telescope focus
TELFILTE= '2 1 1 0 '           / telescope filter wheel encoding
FILTER  = 'V       '           / filter name
GAIN    =                  1.9 / gain (electrons per adu)
RDNOISE =                  5.7 / read noise (electrons)
DWELL   =                   50 / sample integration time (microseconds)
PREFLASH=                    0 / preflash time (seconds)
CAMTEMP =                 -114 / camera temperature (Celcius)
DEWTEMP =                 -195 / dewar temperature (Celcius)
BIASSEC = '[1025:1044,1:1024]' / overscan portion of frame
```

↑ $1st$ 2880 bytes (=36×80)

```
TRIMSEC = '[1:1024,1:1024]'    /  region to be extracted
DATASEC = '[1:1024,1:1024]'    /  image portion of frame
CCDSEC  = '[1:1024,1:1024]'    /  orientation to full frame
ORIGSEC = '[1:1024,1:1024]'    /  original size full frame
CCDSUM  = '2 2     '           /  on chip summation
DISPAXIS= '2       '           /  dispersion axis

                        (29 blank card images)

END
```

↑ $2nd$ 2880 bytes (=36×80)

Figure 7: Example of a FITS header of a twilight sky flat exposure, before any processing, obtained at the VATT (Mt. Graham, AZ) in Sep 2004.

```
SIMPLE  =                    T / Fits standard
BITPIX  =                  -32 / Bits per pixel
NAXIS   =                    2 / Number of axes
NAXIS1  =                 1000 / Axis length
NAXIS2  =                 1000 / Axis length
EXTEND  =                    F / File may contain extensions
ORIGIN  = 'NOAO-IRAF FITS Image Kernel July 2003' / FITS file originator
DATE    = '2004-11-01T07:23:54' / Date FITS file was generated
IRAF-TLM= '00:24:02 (01/11/2004)' / Time of last modification
OBJECT  = 'SKYV    '           / Name of the object observed
IMAGETYP= 'SKY     '           / object, dark, bias, flat, etc.
OBSERVAT= 'MTGRAHAM'           / observatory
TELESCOP= 'VATT    '           / telescope name
INSTRUME= 'CCD     '           / instrument
DETECTOR= 'ccd26   '           / detector
DATE-OBS= '2004-09-15'         / date (YYyy-mm-dd) of observation
UT      = '01:55:35'           / universal time
ST      = '18:13:27.64'        / sidereal time
RA      = '17:56:19.20'        / right ascension
DEC     = '+06:38:26.95'       / declination
EPOCH   =                2000. / epoch of ra and dec
ZD      = ' 26.1835'           / zenith distance (deg)
AIRMASS =                1.114 / airmass
EXPTIME =                  30. / actual integration time (sec)
DARKTIME=                  31. / total elapsed time (sec)
ROTANGLE= '0.000000'           / rotation angle
TELFOCUS= '-118.00 '           / telescope focus
TELFILTE= '2 1 1 0 '           / telescope filter wheel encoding
FILTER  = 'V       '           / filter name
GAIN    =                 1.9 / gain (electrons per adu)
RDNOISE =                 5.7 / read noise (electrons)
DWELL   =                  50 / sample integration time (microseconds)
PREFLASH=                   0 / preflash time (seconds)
CAMTEMP =                -114 / camera temperature (Celcius)
DEWTEMP =                -195 / dewar temperature (Celcius)
CRVAL1  =           47.217250 / RA at reference pixel in degrees
```

$\uparrow 1st$ 2880 bytes (=36×80)

```
CRPIX1  =                507. / Reference pixel, axis 1
CDELT1  =  1.04055555560000E-4 / Increment per pixel on axis1, in degrees
CTYPE1  = 'RA---TAN'           / Coordinate system of axis1
CRVAL2  =            37.23057 / DEC at reference pixel in degrees
CRPIX2  =               522.0 / Reference pixel, axis 2
CDELT2  = -1.0405555556000E-4 / Increment per pixel on axis2, in degrees
CTYPE2  = 'DEC--TAN'           / Coordinate system of axis2
LTV1    =               -15.0 / Translation wrt to original image, axis 1
LTV2    =                 0.0 / Translation wrt to original image, axis 2
CD1_1   =      1.0405555556E-4 / Diagonal element of scale matrix (1,1)
CD2_2   =     -1.0405555556E-4 / Diagonal element of scale matrix (2,2)
LTM1_1  =                 1.0 / Rotation wrt to original image, M(1,1)
LTM2_2  =                 1.0 / Rotation wrt to original image, M(2,2)
```

```
WAT0_001= 'system=world'       / WCS attribute
WAT1_001= 'wtype=linear axtype=ra' / WCS attribute, axis 1
WAT2_001= 'wtype=linear axtype=dec' / WCS attribute, axis 2
HISTORY -
HISTORY ---------  CCD48PROC[V2.12.2-EXPORT]  R.A. Jansen@andromeda  ---------
HISTORY -
CC48PROC= '2004-11-01T07:17:43' / Starting primary frame processing
WCSINIT = 'COMPLETE'            / WCS initialization flag
WCSDIM  =                    2  / WCS dimension
WCSDATE = '2004-11-01T07:17:43' / WCS initialization done
FIXPIX  = 'COMPLETE'            / Bad pixel interpolation flag
FXPXFILE= 'ccd48$lib/VATT/2004/Sep/badpix.ccd26x2' / Bad pixel file
FXPXDATE= '2004-11-01T07:17:44' / Interpolation over bad pixels done
ZEROCOR = 'COMPLETE'            / Zero image correction flag
ZRCRIMAG= '../SUPERBIAS.fits'   / Name of zero image
ZRCRDATE= '2004-11-01T07:17:45' / Zero correction using image done
OVERSCAN= 'COMPLETE'            / Overscan subtraction flag
ORIGSEC = '[1:1044,1:1044]'     / Original size of full frame
BIASSEC = '[1032:1044,1:1000]'  / Overscan section
TRIMSEC = '[16:1015,1:1000]'    / Extraction section
IMAGSEC = '[1:1000,1:1000]'     / Output image section
OVSCMEAN=              1110.49  / Mean bias level (ADU)
OVSCDATE= '2004-11-01T07:17:45' / Removal bias level using overscan done
```

↑ $2nd$ 2880 bytes (=36×80)

```
DARKCOR = 'OMIT'                / Dark image correction flag
SHUTCOR = 'OMIT'                / Shutter effect correction flag
RESPFIT = 'OMIT'                / Pixel-to-pixel response fit flag
SKYCOMB = 'COMPLETE'            / Sky frame combination flag
NCOMBINE=                   14  / Number of sky frames combined
SKCBDATE= '2004-11-01T07:23:30' / Combination of sky frames @SKY_V.lis done
NORMALIZ= 'COMPLETE'            / Sky flat normalization flag
NMLZMEAN=             21575.30  / Mean level (ADU) prior to normalization
NMLZDATE= '2004-11-01T07:23:54' / Sky flat normalization done
CC48DATE= '2004-11-01T07:23:54' / Primary frame processing done
```

(25 blank card images)

```
END
```

↑ $3rd$ 2880 bytes (=36×80)

Figure 8: Example of a FITS header of a processed (bias-subtracted, trimmed, median combined and normalized) twilight sky flat image. Relevant information about the processing steps performed was recorded in the header using keywords created by the user (or his software). Any combination of up to 8 of the allowed characters that does not match a reserved FITS keyword may be used. Note also, that the values of some keywords present in the original header (Fig. 7) were updated. Other astronomers and/or their software packages may record such information, instead, in an extensive `HISTORY` or `COMMENT` section in the header. It is good practice to record sufficient processing information to later be able to trace any problems.

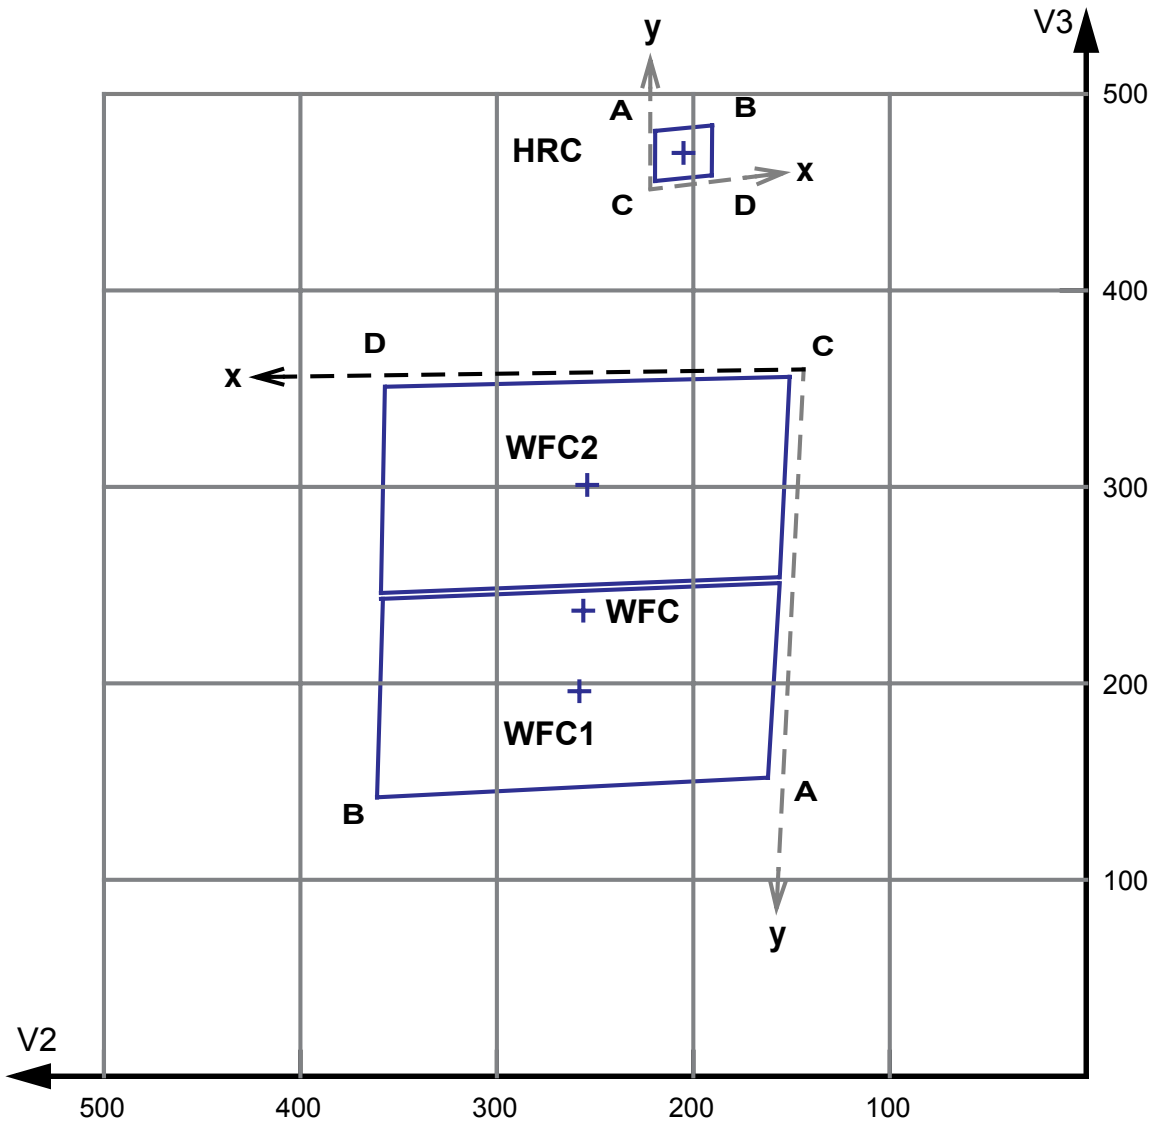

Figure 36: Geometric distortion is a significant effect for off-axis cameras like *HST*/ACS. The square grid formed by the pixels in each CCD array project onto a distorted diamond shape on the sky. (*Source: Pavlovsky et al. 2004*).

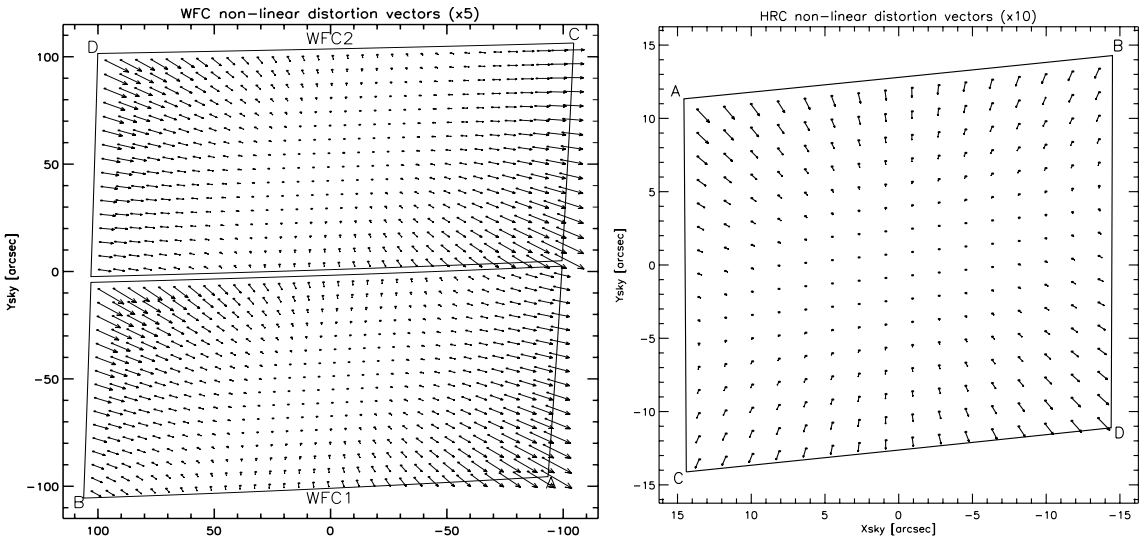

Figure 37: Non-linear component of the geometric distortion in the WFC and HRC cameras of *HST*/ACS. The arrows indicate the direction of the distortion, but are displayed at 5 times their true size. (*Source: Pavlovsky et al. 2004*).

4.2 Multi-Extension FITS (MEF)

If the FITS file does contain one or more extensions in addition to the primary header and data unit (keyword `NEXTEND`), the file said to be a *Multi-Extension FITS* (MEF) file. It is common in such files for the primary data unit to be empty, as indicated by the primary header keyword and value combination:

```
NAXIS␣␣␣=                    0␣/ Number of axes
```

The primary header unit is considered to contain information that is common to each of the following extensions, while the extension headers contain information that only applies to the data stored in the data unit of that particular extension. This avoids needless duplication of information in the extension headers and is termed *keyword inheritance*.

Data of the current generation of instruments aboard the Hubble Space Telescope (STIS, NICMOS and ACS) is stored in MEF format for two reasons:

- one wishes to store not only the data itself from a detector, but also its associated errors and data quality information, which together form a logical group or unit of data;
- the instrument may have more than one detector from which the combined data constitute a single observation (the two detectors of ACS/WFC, for example).

4.3 The ESO Hierarchical FITS header

There is another common file format, that follows the structure of the FITS format, but where most of the telescope, instrument, and detector specific information is recorded in a different way in the FITS header. This file format is known as the *ESO Hierarchical Header Format*. Here, following the obligatory FITS keywords, in the remaining header lines the first 8 characters spell `HIERARCH`, the 9-th character is a blank, and the subsequent characters form a "hierarchical tree" of space separated keywords, the first keyword of which tends to be `ESO`. After the last keyword in the tree an equals sign (`= `) and a blank space appears. The remainder of the 80-character card image is similar in format to standard FITS.

`HIERARCH␣ESO␣LEV1␣LEV2␣...␣LEV`***n*** `= ␣  NUMERIC_VALUE␣/ COMMENT_STRING`
`HIERARCH␣ESO␣LEV1␣LEV2␣...␣LEV`***n*** `= ␣  LOGICAL_VALUE␣/ COMMENT_STRING`
`HIERARCH␣ESO␣LEV1␣LEV2␣...␣LEV`***n*** `= ␣'STRING_VALUE'  ␣/ COMMENT_STRING`

The ESO Hierarchical Header Format, in fact, does *not* violate the FITS standard, but uses the provision in that standard that any line without an equals sign as the 9-th character is treated as a comment line. Standard comment keywords already present in the original FITS paper (Wells et al. 1981) are `COMMENT` and `HISTORY`.

Although almost any FITS reader/viewer will manage to read in and display the data, much of the useful information that describes those data and that may be necessary to properly process and analyse those data are ignored. In order to use standard IRAF or IDL routines to process the data, one has to translate the information in the Hierarchical keyword tree to standard 8-character FITS keywords followed by equals signs.

One way to do so, is to use a translation lookup table, where for each Hierarchical keyword, one searches for a match and writes that match to the image header (e.g., when using IRAF, task `eso2fits` in my `rjtools` package may be useful; *www.public.asu.edu/∼rjansen/iraf/"*).

```
SIMPLE  =                    T / Standard FITS format (NOST-100.0)
BITPIX  =                   16 / # of bits storing pix values
NAXIS   =                    2 / # of axes in frame
NAXIS1  =                 2080 / # pixels/axis
NAXIS2  =                 2048 / # pixels/axis
ORIGIN  = 'ESO     '           / European Southern Observatory
DATE    = '2002-10-05T23:26:56.831' / UT date when this file was written
CRVAL1  =            310.99790 / RA at ref pixel in degrees
CRPIX1  =               1029.8 / Ref. pixel of center of rotation
CDELT1  =         -0.000055479 / Increment in rows
CTYPE1  = 'RA---TAN'           / Pixel coordinate system
CRVAL2  =            -10.77578 / DEC at ref pixel in degrees
CRPIX2  =               1020.3 / Ref. pixel of center of rotation
CDELT2  =         -0.000055479 / Increment in rows
CTYPE2  = 'DEC--TAN'           / Pixel coordinate system
BSCALE  =          1.000000000 / pixel=FITS*BSCALE+BZERO
BZERO   =              32768.0 / pixel=FITS*BSCALE+BZERO
MJD-OBS =       52552.97699599 / MJD start (2002-10-05T23:26:52.454)
DATE-OBS= '2002-10-05T23:26:52.453' / Date of observation
EXPTIME =               3.0050 / Total integration time
EXTEND  =                    F / Extension may be present
INSTRUME= 'FORS1   '           / Instrument used
TELESCOP= 'ESO-VLT-U3'         / ESO Telescope Name
RA      =           310.997917 / 20:43:59.5 RA (J2000) pointing (deg)
DEC     =            -10.77578 / -10:46:32.8 DEC (J2000) pointing (deg)
EQUINOX =                2000. / Standard FK5 (years)
RADECSYS= 'FK5     '           / Coordinate reference frame
LST     =            70986.266 / 19:43:06.266 LST at start (sec)
UTC     =            84413.000 / 23:26:53.000 UTC at start (sec)
HIERARCH ESO OBS TPLNO       =          3 / Template number within OB
HIERARCH ESO OBS NAME        = 'Photom-std-MarkA-high-23nov' / OB name
HIERARCH ESO OBS PROG ID     = '70.A-0376(A)' / ESO program identification
HIERARCH ESO OBS ID          =    200121629 / Observation block ID
HIERARCH ESO OBS DID         = 'ESO-VLT-DIC.OBS-1.7' / OBS Dictionary
HIERARCH ESO OBS OBSERVER    = 'Jakobsen'   / Observer name
HIERARCH ESO OBS PI-COI ID   =          780 / ESO internal PI-COI ID
HIERARCH ESO OBS TARG NAME   = 'MarkA   '   / OB target name
HIERARCH ESO OBS GRP         = '0       '   / linked blocks
HIERARCH ESO OBS START       = '2002-10-05T23:21:20' / OB start time
HIERARCH ESO OBS EXECTIME    =         1013 / Expected execution time
HIERARCH ESO TPL ID          = 'FORS1_img_obs_crsplit' / Template signature ID
HIERARCH ESO TPL NAME        = 'jitter images' / Template name
HIERARCH ESO TPL NEXP        =            1 / Number of exposures within template
HIERARCH ESO TPL EXPNO       =            1 / Exposure number within template
HIERARCH ESO TPL START       = '2002-10-05T23:25:53' / TPL start time
HIERARCH ESO DPR CATG        = 'CALIB   '   / Observation category
HIERARCH ESO DPR TYPE        = 'STD     '   / Observation type
HIERARCH ESO DPR TECH        = 'IMAGE   '   / Observation technique
HIERARCH ESO TEL DID         = 'ESO-VLT-DIC.TCS' / Data dictionary for TEL
HIERARCH ESO TEL ID          = 'v 1.444 '   / TCS version number
HIERARCH ESO TEL DATE        = '2002-09-20T19:36:22.000' / TCS installation date
```

```
HIERARCH ESO TEL ALT         =       69.959 / Alt angle at start (deg)
HIERARCH ESO TEL AZ          =      228.951 / Az angle at start (deg)
HIERARCH ESO TEL GEOELEV     =        2648. / Elevation above sea level (m)
HIERARCH ESO TEL GEOLAT      =     -24.6251 / Tel geo latitute (+=North) (deg)
HIERARCH ESO TEL GEOLON      =     -70.4027 / Tel geo longitude (+=East) (deg)
HIERARCH ESO TEL OPER        = 'I, Condor'  / Telescope Operator
HIERARCH ESO TEL FOCU ID     = 'CA      '   / Telescope focus station ID
HIERARCH ESO TEL FOCU LEN    =      108.827 / Focal length (m)
HIERARCH ESO TEL FOCU SCALE  =        1.894 / Focal scale (arcsec/mm)
HIERARCH ESO TEL FOCU VALUE  =       -0.638 / M2 setting (mm)
HIERARCH ESO TEL PARANG START=     -135.747 / Parallactic angle at start (deg)
HIERARCH ESO TEL AIRM START  =        1.064 / Airmass at start
HIERARCH ESO TEL AMBI FWHM START=      1.06 / Observatory Seeing
HIERARCH ESO TEL AMBI PRES START=    743.50 / Observatory ambient air pressure
HIERARCH ESO TEL AMBI WINDSP =         3.42 / Observatory ambient wind speed
HIERARCH ESO TEL AMBI WINDDIR=         241. / Observatory ambient wind direction
HIERARCH ESO TEL AMBI RHUM   =           4. / Observatory ambient rel. humidity
HIERARCH ESO TEL AMBI TEMP   =        15.86 / Observatory ambient temperature
HIERARCH ESO TEL MOON RA     = 122348.524519 / ~~:~~:~~.~ RA (J2000) (deg)
HIERARCH ESO TEL MOON DEC    =  22811.42322 / ~~:~~:~~.~ DEC (J2000) (deg)
HIERARCH ESO TEL TH M1 TEMP  =        12.89 / M1 superficial temperature
HIERARCH ESO TEL TRAK STATUS = 'NORMAL  '   / Tracking status
HIERARCH ESO TEL DOME STATUS = 'FULLY-OPEN' / Dome status
HIERARCH ESO TEL CHOP ST     =           F / True when chopping is active
HIERARCH ESO TEL PARANG END  =     -135.776 / Parallactic angle at end (deg)
HIERARCH ESO TEL AIRM END    =        1.064 / Airmass at end
HIERARCH ESO TEL AMBI FWHM END=        1.06 / Observatory Seeing queried from AS
HIERARCH ESO TEL AMBI PRES END=      743.50 / Observatory ambient air pressure
HIERARCH ESO ADA ABSROT START=     43.22451 / Abs rot angle at exp start (deg)
HIERARCH ESO ADA POSANG      =      0.00000 / Position angle at start
HIERARCH ESO ADA GUID STATUS = 'ON      '   / Status of autoguider
HIERARCH ESO ADA GUID RA     =   311.012510 / 20:44:03.0 Guide star RA J2000
HIERARCH ESO ADA GUID DEC    =    -10.67274 / -10:40:21.8 Guide star DEC J2000
HIERARCH ESO ADA ABSROT END  =     43.19542 / Abs rot angle at exp end (deg)
HIERARCH ESO INS PIXSCALE    =        0.200 / pixel scale in arcsec/pixel
HIERARCH ESO INS MODE        = 'IMG     '   / Instrument mode used
HIERARCH ESO INS COLL ID     = '+1      '   / Collimator unique ID
HIERARCH ESO INS COLL NAME   = 'COLL_SR '   / Collimator name
HIERARCH ESO INS FOCU POS    =  -0.01475000 / Focus position in cm
HIERARCH ESO INS FOCU TEMP   =         15.5 / Focus temperature in C
HIERARCH ESO INS OPTI2 NAME  = 'free    '   / name for any opti
HIERARCH ESO INS OPTI2 ID    = '        '   / identification for any opti
HIERARCH ESO INS OPTI2 TYPE  = '        '   / Element type
HIERARCH ESO INS OPTI3 NAME  = 'COLL_SR '   / name for any opti
HIERARCH ESO INS OPTI3 ID    = '+1      '   / identification for any opti
HIERARCH ESO INS OPTI3 TYPE  = 'COLL    '   / Element type
HIERARCH ESO INS OPTI4 NAME  = 'free    '   / name for any opti
HIERARCH ESO INS OPTI4 ID    = '        '   / identification for any opti
HIERARCH ESO INS OPTI4 TYPE  = '        '   / Element type
HIERARCH ESO INS OPTI5 NAME  = 'free    '   / name for any opti
HIERARCH ESO INS OPTI5 ID    = '        '   / identification for any opti
```

```
HIERARCH ESO INS OPTI5 TYPE  = '        '   / Element type
HIERARCH ESO INS OPTI6 NAME  = 'free    '   / name for any opti
HIERARCH ESO INS OPTI6 ID    = '        '   / identification for any opti
HIERARCH ESO INS OPTI6 TYPE  = '        '   / Element type
HIERARCH ESO INS OPTI7 NAME  = 'B_BESS  '   / name for any opti
HIERARCH ESO INS OPTI7 ID    = '+34     '   / identification for any opti
HIERARCH ESO INS OPTI7 TYPE  = 'FILT    '   / Element type
HIERARCH ESO INS OPTI8 NAME  = 'CAMERA  '   / name for any opti
HIERARCH ESO INS OPTI8 ID    = '+3      '   / identification for any opti
HIERARCH ESO INS OPTI8 TYPE  = 'CAMERA  '   / Element type
HIERARCH ESO INS OPTI9 NAME  = 'free    '   / name for any opti
HIERARCH ESO INS OPTI9 ID    = '        '   / identification for any opti
HIERARCH ESO INS OPTI9 TYPE  = '        '   / Element type
HIERARCH ESO INS OPTI10 NAME = 'free    '   / name for any opti
HIERARCH ESO INS OPTI10 ID   = '        '   / identification for any opti
HIERARCH ESO INS OPTI10 TYPE = '        '   / Element type
HIERARCH ESO INS FILT1 ID    = '+34     '   / Filter unique ID
HIERARCH ESO INS FILT1 NAME  = 'B_BESS  '   / Filter i name
HIERARCH ESO INS IMAGE DISTOR1= 3.602000e-04 / distortion coefficient
HIERARCH ESO INS IMAGE DISTOR2= -1.228000e-06 / distortion coefficient
HIERARCH ESO INS IMAGE DISTOR3= 2.091000e-09 / distortion coefficient
HIERARCH ESO INS MOS1 POS    = 1.590765e-04 / Position of the slit in mm.
HIERARCH ESO INS MOS2 POS    = 8.707600e-05 / Position of the slit in mm.
HIERARCH ESO INS MOS3 POS    = 8.708550e-05 / Position of the slit in mm.
HIERARCH ESO INS MOS4 POS    = 3.056250e-04 / Position of the slit in mm.
HIERARCH ESO INS MOS5 POS    = -8.691450e-05 / Position of the slit in mm.
HIERARCH ESO INS MOS6 POS    = -1.069820e-04 / Position of the slit in mm.
HIERARCH ESO INS MOS7 POS    = 1.151220e-04 / Position of the slit in mm.
HIERARCH ESO INS MOS8 POS    = -5.183100e-05 / Position of the slit in mm.
HIERARCH ESO INS MOS9 POS    = -7.887350e-05 / Position of the slit in mm.
HIERARCH ESO INS MOS10 POS   = -2.739640e-04 / Position of the slit in mm.
HIERARCH ESO INS MOS11 POS   = -5.581800e-05 / Position of the slit in mm.
HIERARCH ESO INS MOS12 POS   = -1.230390e-04 / Position of the slit in mm.
HIERARCH ESO INS MOS13 POS   = 1.249610e-04 / Position of the slit in mm.
HIERARCH ESO INS MOS14 POS   = 5.529750e-05 / Position of the slit in mm.
HIERARCH ESO INS MOS15 POS   = 2.380060e-04 / Position of the slit in mm.
HIERARCH ESO INS MOS16 POS   = -3.010420e-04 / Position of the slit in mm.
HIERARCH ESO INS MOS17 POS   = -8.691450e-05 / Position of the slit in mm.
HIERARCH ESO INS MOS18 POS   = -3.960000e-06 / Position of the slit in mm.
HIERARCH ESO INS MOS19 POS   = -1.389895e-04 / Position of the slit in mm.
HIERARCH ESO INS MOS1 WIDTH  = 2.459997e+02 / Width of the MOS slit in mm.
HIERARCH ESO INS MOS1 WID    =       466.25 / MOSi slit width in arcsec
HIERARCH ESO INS MOS2 WIDTH  = 2.459988e+02 / Width of the MOS slit in mm.
HIERARCH ESO INS MOS2 WID    =       466.25 / MOSi slit width in arcsec
HIERARCH ESO INS MOS3 WIDTH  = 2.459998e+02 / Width of the MOS slit in mm.
HIERARCH ESO INS MOS3 WID    =       466.25 / MOSi slit width in arcsec
HIERARCH ESO INS MOS4 WIDTH  = 2.459994e+02 / Width of the MOS slit in mm.
HIERARCH ESO INS MOS4 WID    =       466.25 / MOSi slit width in arcsec
HIERARCH ESO INS MOS5 WIDTH  = 2.460002e+02 / Width of the MOS slit in mm.
HIERARCH ESO INS MOS5 WID    =       466.25 / MOSi slit width in arcsec
HIERARCH ESO INS MOS6 WIDTH  = 2.460002e+02 / Width of the MOS slit in mm.
```

```
HIERARCH ESO INS MOS6 WID    =       466.25 / MOSi slit width in arcsec
HIERARCH ESO INS MOS7 WIDTH  = 2.459998e+02 / Width of the MOS slit in mm.
HIERARCH ESO INS MOS7 WID    =       466.25 / MOSi slit width in arcsec
HIERARCH ESO INS MOS8 WIDTH  = 2.460001e+02 / Width of the MOS slit in mm.
HIERARCH ESO INS MOS8 WID    =       466.25 / MOSi slit width in arcsec
HIERARCH ESO INS MOS9 WIDTH  = 2.460002e+02 / Width of the MOS slit in mm.
HIERARCH ESO INS MOS9 WID    =       466.25 / MOSi slit width in arcsec
HIERARCH ESO INS MOS10 WIDTH = 2.459995e+02 / Width of the MOS slit in mm.
HIERARCH ESO INS MOS10 WID   =       466.25 / MOSi slit width in arcsec
HIERARCH ESO INS MOS11 WIDTH = 2.459991e+02 / Width of the MOS slit in mm.
HIERARCH ESO INS MOS11 WID   =       466.25 / MOSi slit width in arcsec
HIERARCH ESO INS MOS12 WIDTH = 2.460002e+02 / Width of the MOS slit in mm.
HIERARCH ESO INS MOS12 WID   =       466.25 / MOSi slit width in arcsec
HIERARCH ESO INS MOS13 WIDTH = 2.459998e+02 / Width of the MOS slit in mm.
HIERARCH ESO INS MOS13 WID   =       466.25 / MOSi slit width in arcsec
HIERARCH ESO INS MOS14 WIDTH = 2.459999e+02 / Width of the MOS slit in mm.
HIERARCH ESO INS MOS14 WID   =       466.25 / MOSi slit width in arcsec
HIERARCH ESO INS MOS15 WIDTH = 2.459995e+02 / Width of the MOS slit in mm.
HIERARCH ESO INS MOS15 WID   =       466.25 / MOSi slit width in arcsec
HIERARCH ESO INS MOS16 WIDTH = 2.459996e+02 / Width of the MOS slit in mm.
HIERARCH ESO INS MOS16 WID   =       466.25 / MOSi slit width in arcsec
HIERARCH ESO INS MOS17 WIDTH = 2.460002e+02 / Width of the MOS slit in mm.
HIERARCH ESO INS MOS17 WID   =       466.25 / MOSi slit width in arcsec
HIERARCH ESO INS MOS18 WIDTH = 2.460000e+02 / Width of the MOS slit in mm.
HIERARCH ESO INS MOS18 WID   =       466.25 / MOSi slit width in arcsec
HIERARCH ESO INS MOS19 WIDTH = 2.459993e+02 / Width of the MOS slit in mm.
HIERARCH ESO INS MOS19 WID   =       466.25 / MOSi slit width in arcsec
HIERARCH ESO INS SHUT EXPTIME= 3.004968e+00 / Real exposure time in sec.
HIERARCH ESO INS MOS CHECKSUM=            0 / Checksum of MOS slitlet configuration
HIERARCH ESO DET ID          = 'CCD FIERA - Rev 2.87' / Detector system Id
HIERARCH ESO DET NAME        = 'ccdF - fors' / Name of detector system
HIERARCH ESO DET DATE        = '05/02/2001' / Installation date
HIERARCH ESO DET DID         = 'ESO-VLT-DIC.CCDDCS,ESO-VLT-DIC.FCDDCS' / Dictionary
HIERARCH ESO DET BITS        =           16 / Bits per pixel readout
HIERARCH ESO DET RA          = -445589.27727324 / Apparent -~~:~~:~~.~ RA at start
HIERARCH ESO DET DEC         = -445589.27727324 / Apparent -~~:~~:~~.~ DEC at start
HIERARCH ESO DET SOFW MODE   = 'Normal  '   / CCD sw operational mode
HIERARCH ESO DET CHIPS       =            1 / # of chips in detector array
HIERARCH ESO DET CHIP1 ID    = 'TK2048EB4-1 160' / Detector chip identification
HIERARCH ESO DET CHIP1 NAME  = '        '   / Detector chip name
HIERARCH ESO DET CHIP1 DATE  = '31/10/1999' / Date of installation [DD-MM-YYYY]
HIERARCH ESO DET CHIP1 X     =            1 / X location in array
HIERARCH ESO DET CHIP1 Y     =            1 / Y location in array
HIERARCH ESO DET CHIP1 NX    =         2048 / # of pixels along X
HIERARCH ESO DET CHIP1 NY    =         2049 / # of pixels along Y
HIERARCH ESO DET CHIP1 PSZX  =         24.0 / Size of pixel in X
HIERARCH ESO DET CHIP1 PSZY  =         24.0 / Size of pixel in Y
HIERARCH ESO DET EXP NO      =         9391 / Unique exposure ID number
HIERARCH ESO DET EXP TYPE    = 'Normal  '   / Exposure type
HIERARCH ESO DET EXP DUMDIT  =            0 / # of dummy readouts
HIERARCH ESO DET EXP RDTTIME =       25.031 / image readout time
```

```
HIERARCH ESO DET EXP XFERTIM =       35.454 / image transfer time
HIERARCH ESO DET READ MODE   = 'normal  '   / Readout method
HIERARCH ESO DET READ SPEED  = 'normal  '   / Readout speed
HIERARCH ESO DET READ CLOCK  = 'ABCD,1x1,high' / Readout clock pattern used
HIERARCH ESO DET OUTPUTS     =            4 / # of outputs
HIERARCH ESO DET OUTREF      =            0 / reference output
HIERARCH ESO DET OUT1 ID     = 'A       '   / Output ID as from manufacturer
HIERARCH ESO DET OUT1 NAME   = 'A       '   / Description of output
HIERARCH ESO DET OUT1 CHIP   =            1 / Chip to which the output belongs
HIERARCH ESO DET OUT1 X      =            1 / X location of output
HIERARCH ESO DET OUT1 Y      =            1 / Y location of output
HIERARCH ESO DET OUT1 NX     =         1024 / valid pixels along X
HIERARCH ESO DET OUT1 NY     =         1024 / valid pixels along Y
HIERARCH ESO DET OUT1 PRSCX  =           16 / Prescan region in X
HIERARCH ESO DET OUT1 OVSCX  =            0 / Overscan region in X
HIERARCH ESO DET OUT1 CONAD  =         1.46 / Conversion from ADUs to electrons
HIERARCH ESO DET OUT1 RON    =         5.16 / Readout noise per output (e-)
HIERARCH ESO DET OUT1 GAIN   =         0.68 / Conversion from electrons to ADU
HIERARCH ESO DET OUT2 ID     = 'B       '   / Output ID as from manufacturer
HIERARCH ESO DET OUT2 NAME   = 'B       '   / Description of output
HIERARCH ESO DET OUT2 CHIP   =            1 / Chip to which the output belongs
HIERARCH ESO DET OUT2 X      =         2048 / X location of output
HIERARCH ESO DET OUT2 Y      =            1 / Y location of output
HIERARCH ESO DET OUT2 NX     =         1024 / valid pixels along X
HIERARCH ESO DET OUT2 NY     =         1024 / valid pixels along Y
HIERARCH ESO DET OUT2 PRSCX  =           16 / Prescan region in X
HIERARCH ESO DET OUT2 OVSCX  =            0 / Overscan region in X
HIERARCH ESO DET OUT2 CONAD  =         1.59 / Conversion from ADUs to electrons
HIERARCH ESO DET OUT2 RON    =         5.44 / Readout noise per output (e-)
HIERARCH ESO DET OUT2 GAIN   =         0.63 / Conversion from electrons to ADU
HIERARCH ESO DET OUT3 ID     = 'C       '   / Output ID as from manufacturer
HIERARCH ESO DET OUT3 NAME   = 'C       '   / Description of output
HIERARCH ESO DET OUT3 CHIP   =            1 / Chip to which the output belongs
HIERARCH ESO DET OUT3 X      =            1 / X location of output
HIERARCH ESO DET OUT3 Y      =         2049 / Y location of output
HIERARCH ESO DET OUT3 NX     =         1024 / valid pixels along X
HIERARCH ESO DET OUT3 NY     =         1024 / valid pixels along Y
HIERARCH ESO DET OUT3 PRSCX  =           16 / Prescan region in X
HIERARCH ESO DET OUT3 OVSCX  =            0 / Overscan region in X
HIERARCH ESO DET OUT3 CONAD  =         1.78 / Conversion from ADUs to electrons
HIERARCH ESO DET OUT3 RON    =         5.57 / Readout noise per output (e-)
HIERARCH ESO DET OUT3 GAIN   =         0.56 / Conversion from electrons to ADU
HIERARCH ESO DET OUT4 ID     = 'D       '   / Output ID as from manufacturer
HIERARCH ESO DET OUT4 NAME   = 'D       '   / Description of output
HIERARCH ESO DET OUT4 CHIP   =            1 / Chip to which the output belongs
HIERARCH ESO DET OUT4 X      =         2048 / X location of output
HIERARCH ESO DET OUT4 Y      =         2049 / Y location of output
HIERARCH ESO DET OUT4 NX     =         1024 / valid pixels along X
HIERARCH ESO DET OUT4 NY     =         1024 / valid pixels along Y
HIERARCH ESO DET OUT4 PRSCX  =           16 / Prescan region in X
HIERARCH ESO DET OUT4 OVSCX  =            0 / Overscan region in X
```

```
HIERARCH ESO DET OUT4 CONAD  =         1.65 / Conversion from ADUs to electrons
HIERARCH ESO DET OUT4 RON    =         5.38 / Readout noise per output (e-)
HIERARCH ESO DET OUT4 GAIN   =         0.61 / Conversion from electrons to ADU
HIERARCH ESO DET FRAM ID     =            1 / Image sequencial number
HIERARCH ESO DET FRAM TYPE   = 'Normal  '   / Type of frame
HIERARCH ESO DET WINDOWS     =            1 / # of windows readout
HIERARCH ESO DET WIN1 STRX   =            1 / Lower left pixel in X
HIERARCH ESO DET WIN1 STRY   =            1 / Lower left pixel in Y
HIERARCH ESO DET WIN1 NX     =         2080 / # of pixels along X
HIERARCH ESO DET WIN1 NY     =         2048 / # of pixels along Y
HIERARCH ESO DET WIN1 BINX   =            1 / Binning factor along X
HIERARCH ESO DET WIN1 BINY   =            1 / Binning factor along Y
HIERARCH ESO DET WIN1 NDIT   =            1 / # of subintegrations
HIERARCH ESO DET WIN1 UIT1   =     3.000000 / user defined subintegration time
HIERARCH ESO DET WIN1 DIT1   =     3.004814 / actual subintegration time
HIERARCH ESO DET WIN1 DKTM   =       3.1088 / Dark current time
HIERARCH ESO DET SHUT TYPE   = 'Iris    '   / type of shutter
HIERARCH ESO DET SHUT ID     = 'fors shutter' / Shutter unique identifier
HIERARCH ESO DET SHUT TMOPEN =        0.190 / Time taken to open shutter
HIERARCH ESO DET SHUT TMCLOS =        0.199 / Time taken to close shutter
HIERARCH ESO DET TELE INT    =         60.0 / Interval between two successive te
HIERARCH ESO DET TELE NO     =            3 / # of sources active
HIERARCH ESO DET TLM1 NAME   = 'CCD T1  '   / Description of telemetry param.
HIERARCH ESO DET TLM1 ID     = 'CCD T1  '   / ID of telemetry sensor
HIERARCH ESO DET TLM1 START  =       163.00 / Telemetry value at read start
HIERARCH ESO DET TLM1 END    =       163.00 / Telemetry value at read completion
HIERARCH ESO DET TLM2 NAME   = 'CCD T2  '   / Description of telemetry param.
HIERARCH ESO DET TLM2 ID     = 'CCD T2  '   / ID of telemetry sensor
HIERARCH ESO DET TLM2 START  =       164.10 / Telemetry value at read start
HIERARCH ESO DET TLM2 END    =       164.10 / Telemetry value at read completion
HIERARCH ESO DET TLM3 NAME   = 'Vacuum  '   / Description of telemetry param.
HIERARCH ESO DET TLM3 ID     = 'Vacuum  '   / ID of telemetry sensor
HIERARCH ESO DET TLM3 START  =         0.00 / Telemetry value at read start
HIERARCH ESO DET TLM3 END    =         0.00 / Telemetry value at read completion
COMMENT CRVAL1 3.109979e+02
COMMENT CRVAL2 -1.077578e+01
COMMENT CDELT1 -5.547858e-05
COMMENT CDELT2 5.547858e-05
ORIGFILE= 'FORS1_IMG_STD278.1.fits' / Original File Name
ARCFILE = 'FORS1.2002-10-05T23:26:52.454.fits' / Archive File Name
CHECKSUM= 'dDZAfAY1dAY8dAY8'   / ASCII 1's complement checksum

                    (28 blank card images)

END
```

Figure 11: Example of a FITS header written in the ESO Hierarchical format, obtained at the VLT-UT3 with the FORS1 instrument (an imager and multi-slit spectrograph, with a CCD that is read out in four quadrants through four amplifiers).

5 Looking through the Earth's Atmosphere

5.1 Atmospheric Transmission, Emission, and Scattering

All observations using ground-based instruments suffer the effects of having to look through the Earth's atmosphere. Turbulence and temperature variations in the atmosphere can blur any incident astronomical signal, and the transparency of the atmosphere is a function of wavelength — in some wavelength regimes a strong one. Below ∼310 nm, the atmosphere is essentially opaque, due to the combined effects of ozone (O_3) and Rayleigh scattering ($\propto \lambda^{-4}$). Above ∼800 nm there are only discrete windows where observations are possible (see Fig. 39). Atoms and molecules in the atmosphere are also emitting light at specific wavelengths (see Fig. 40). And lastly we have to contend with light from our Moon that is scattered by the atmosphere.

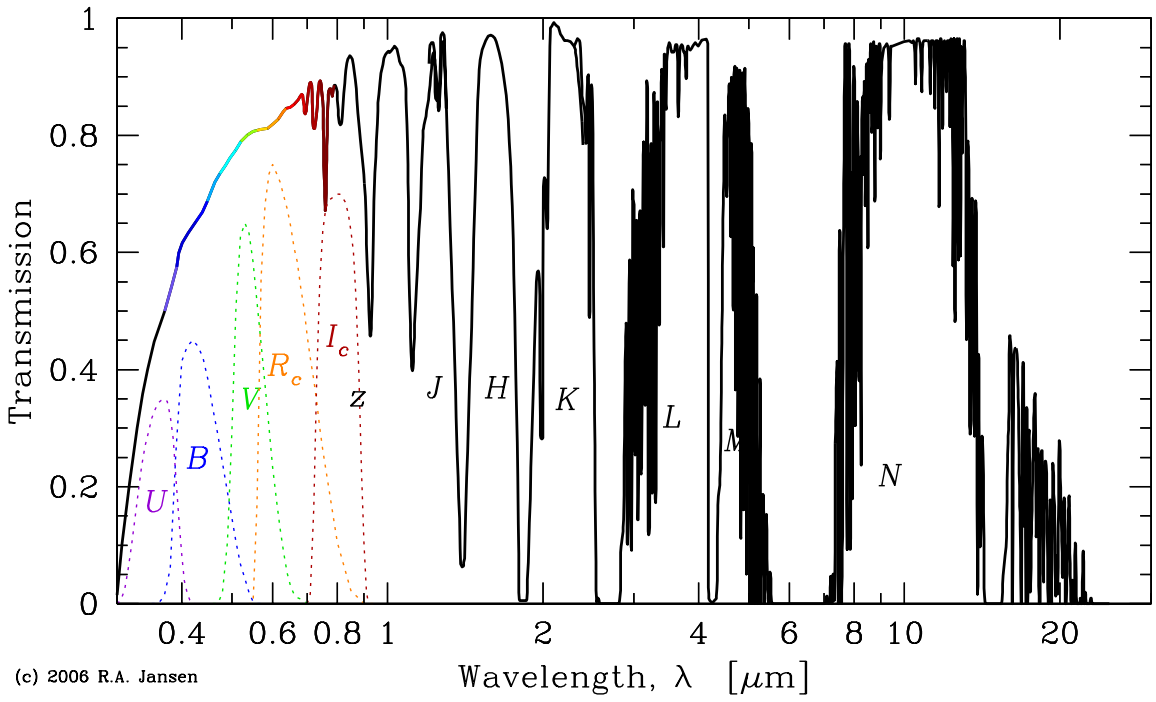

Figure 39: Atmospheric transmission from the atmospheric cut-off in the near-UV (∼310 nm) to the cut-off redward of the N window in the mid-IR (*After: RCA Electro-Optics Handbook*). The shape of several classical filter passbands ($UBVR_cI_c$) is indicated, and the atmospheric windows in the near- and mid-IR are labeled. Most of the atmospheric absorption is due to molecules: O_3 in the near-UV and visible, and O_2, H_2O, CO_2, and N_2O in the infrared. The transmission of the atmosphere is strongly dependent on the amount of precipitable water vapor in the atmosphere above the telescope. The effects of atmospheric *emission*, which limits the usable wavelength intervals even further, is not included in this plot (but see Fig. 40).

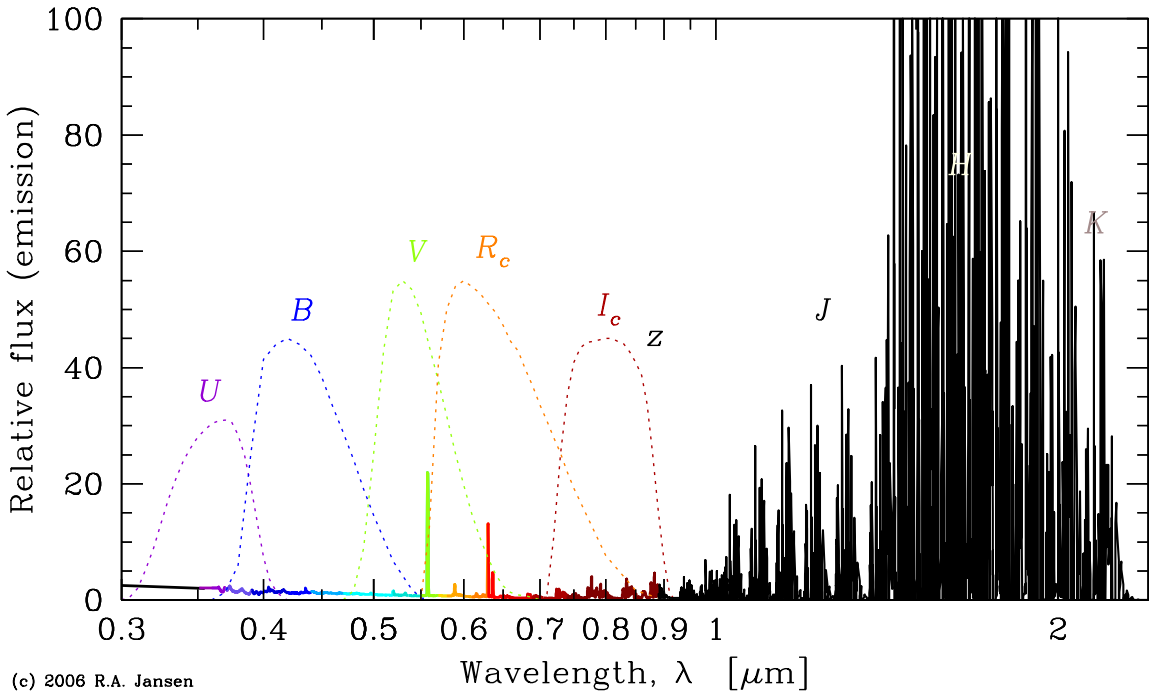

Figure 40: Atmospheric emission from the atmospheric cut-off in the near-UV through the near-IR K band (Source: *0.3–0.9 μm: based on VLT/UT3+FORS1 spectra obtained by R. Jansen & P. Jakobsen; 0.9–2.5 μm: theoretical OH night-sky lines from Rousselot et al. 2000*). Again, shapes of several classical filter passbands or their effective wavelengths are indicated. Although the optical night-sky spectrum — but for the very strong O I lines at 557.7, 630.0 and 636.3 nm — is relatively clean, the sky spectrum beyond ∼750 nm is dominated by strong series of OH lines. The strengths of night-sky emission lines tend to vary throughout a night and from night to night. Beyond ∼4 μm, the sky becomes exceedingly bright.

Table 1: Sky brightness for different lunar ages at CTIO

Lunar age (days)	Sky surface brightness (mag/arcsec^2)				
	U	B	V	R	I
0	22.0	22.7	21.8	20.9	19.9
3	21.5	22.4	21.7	20.8	19.9
7	19.9	21.6	21.4	20.6	19.7
10	18.5	20.7	20.7	20.3	19.5
14	17.0	19.5	20.0	19.9	19.2

Source: Alistair Walker, NOAO Newsletter #10

Table 1 tabulates the average surface brightness of the sky background and its dependence on lunar phase for five filters ($UBVRI$) and lunar ages of 0 (New Moon) through 14 (Full Moon). This dependence is due to scattering of moonlight in the Earth's atmosphere. The actual brightness also depends on how close to the moon you look and, for partial phases, whether the Moon is above the horizon. (Remember, that for partial phases, the Moon is up only part of the night, while at Full Moon it is up *all* night). The main trends are, that:

- The Moon's effect is strongest in the bluest bandpasses: in U, the sky is 5 mag (i.e., a factor 100!) brighter at full moon than at new Moon or when the Moon is well below the horizon. In B, the difference is $\sim$3.2 mag. In the redder I filter, the difference is reduced to a factor $\sim$2. This dependence on wavelength is partly because the moonless sky is fainter in the bluer bandpasses than it is in redder passbands — in the mid-IR the sky is so bright that there is little difference between day and night!

- Taylor, Jansen & Windhorst (2004) show that the darkest skies at the summit of Mt. Graham, AZ, where the 2$\times$8.4 m Large Binocular Telescope (LBT) is being built, have a sky surface brightness that is comparable to the darkest nights at Cerro Tollolo Inter-american Observatory (CTIO) in Chili.

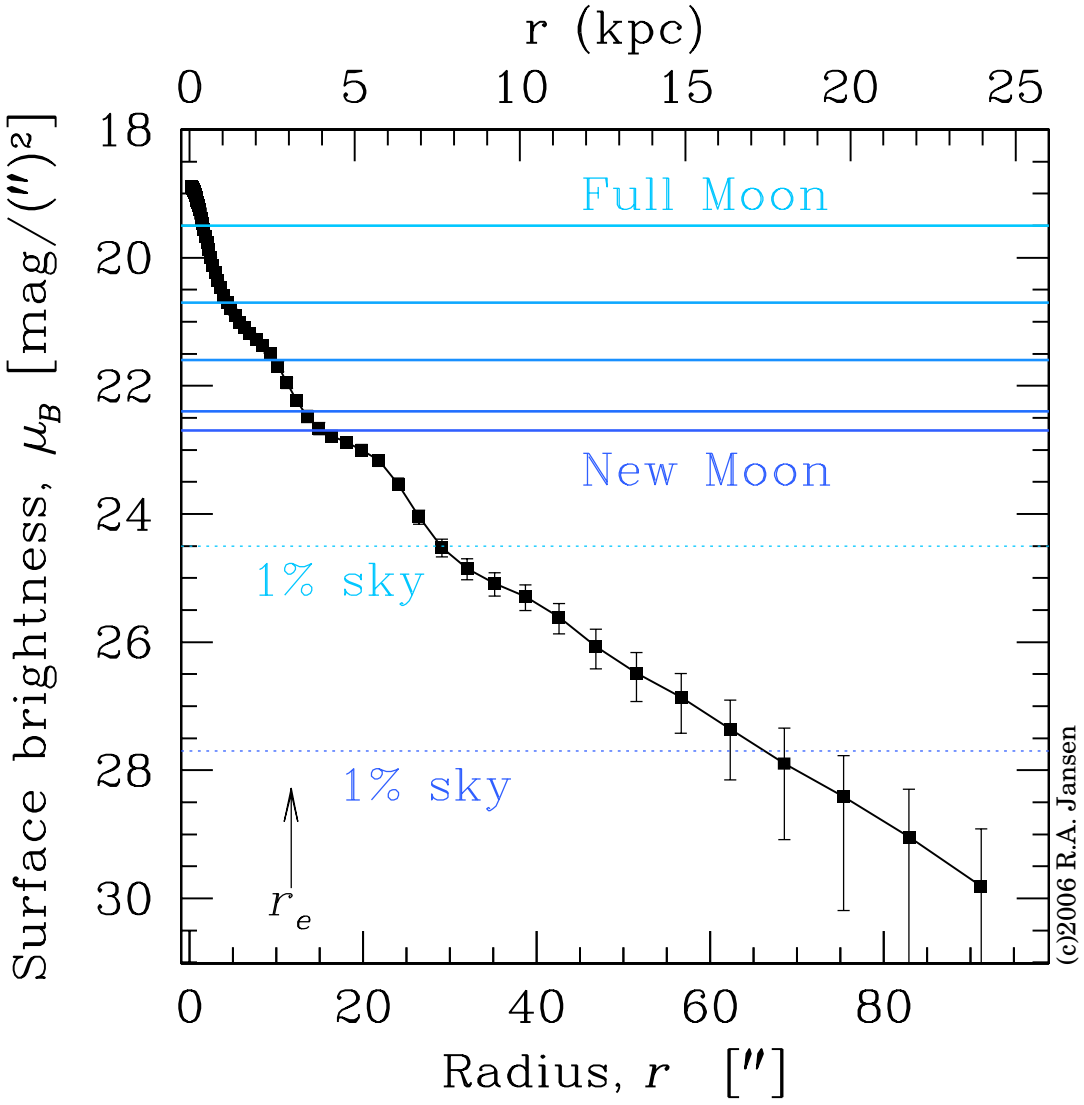

Figure 41: Comparison of the B-filter surface brightness profile of a nearby galaxy (A00389−0159; Jansen et al. 2000a) with the surface brightness of the sky, for the moon phases of Table 1. Note that, even at New Moon, most of the galaxy is fainter than the sky. With a central surface brightness, $\mu_0^B \simeq 19\,\mathrm{mag\,arcsec^{-2}}$ and a profile that declines exponentially with radius in the outer disk, this galaxy is quite typical of nearby spiral galaxies. Generally, with sufficiently long exposures, one can confidently detect signals to ∼1% of the sky background. The difference between observing during *dark* and *bright* sky conditions, corresponds to a factor of more than 2 in radius within most spiral galaxies.

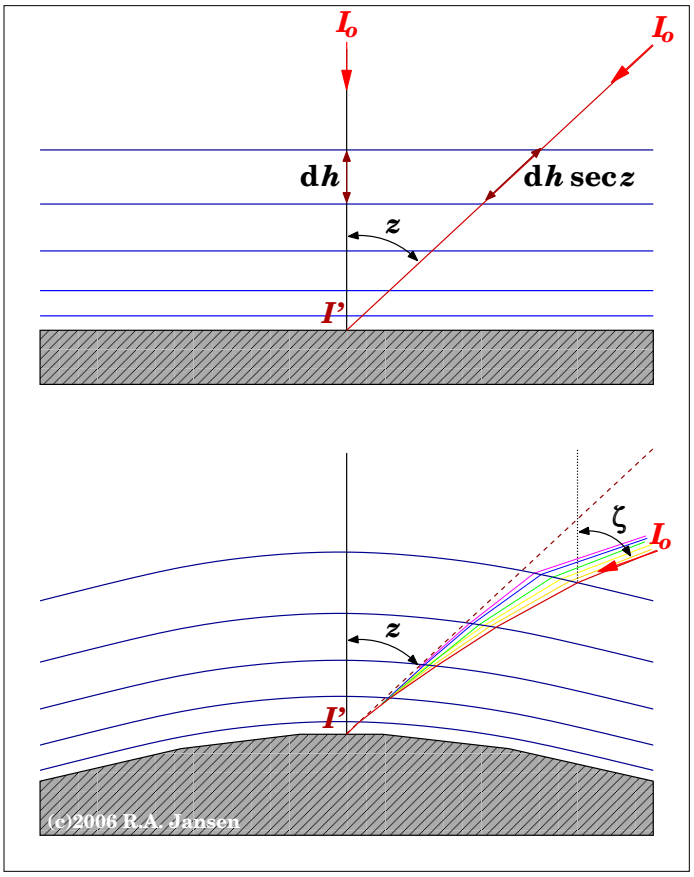

Figure 42: Schematic representation of the geometry of the Earth's atmosphere (a) in the approximation of plane parallel layers; (b) when the curvature of the atmosphere and refraction are taken into account.

Loss of intensity during the passage of light through the atmosphere results from *extinction* — the combined effects of *absorption* and *scattering*. For plane parallel layers and for monochromatic light we have at zenith (see Fig. 42):

$$\mathrm{d}I = -I \cdot \kappa \, \mathrm{d}h \qquad \text{or:} \qquad \frac{\mathrm{d}I}{I} = -\kappa \, \mathrm{d}h$$

$$\Longrightarrow \quad \ln I = -\int_0^\infty \kappa(h) \, \mathrm{d}h \; + \; \text{const.}$$

where κ is the extinction coefficient per cm — which will depend on the density of the atmosphere and, hence, on h —, and const. $= \ln(I_0)$. If we denote the extinction coefficient for one *air mass* (i.e., a full thickness of the Earth's atmosphere) by K, then the *transmission* through the atmosphere will be given by $I/I_0 = e^{-K}$.
For a *zenith distance* (angle from zenith) z this becomes:

$$I/I_0 = e^{-K \, \sec z} \quad \text{, with:} \quad \sec z \equiv \frac{1}{\cos z} \tag{6}$$

And expressed in magnitudes:

$$\begin{aligned} m - m_0 &\equiv -2.5 \cdot \log(I/I_0) = -2.5 \cdot \log(e^{-K \sec z}) \\ &= -2.5 \cdot (-K \sec z) \cdot \log e \simeq 1.0857 \cdot K \sec z \end{aligned} \tag{7}$$

5.2 Atmospheric Refraction and Dispersion

In reality, however, the Earth's atmosphere is curved, not plane parallel, and the path followed by a beam of light will be curved due to *atmospheric refraction*. This is illustrated in the bottom panel of Fig. 42. Generally, it may be necessary to use the actual air mass, $M(z)$, instead of $\sec z$.

To evaluate the importance of the difference between $\sec z$ and $M(z)$ and of atmospheric refraction, in Table 2 we compare the relative intensity losses I/I_0 for each at two different wavelengths. The correct form of Eq. (1) is:

$$I/I_0 = e^{-K \cdot M(z)} \quad \text{with:} \quad M(0) \equiv \sec 0 \equiv 1 \tag{8}$$

It is important to note that the values for $M(z)$ listed in Table 2 are valid for standard temperature and pressure *at sealevel*; $M(z)$ varies with altitude and temperature. Nonetheless, we find that the intensity losses do not differ significantly for zenith distances $\lesssim 80°$. The extinction coefficient $K(\lambda)$, however, is — as we already found from Fig. 39 — a strong function of wavelength, as evidenced here by the large difference in losses at 5500Å and 3200Å.

The zenith distance z of an astronomical object as observed at the ground is always *smaller* than the true zenith distance ζ that one would measure outside of the Earth's atmosphere. For a precise calculation of the refraction one would need to know the geometry of the atmosphere and the change in refractive index as a function of density (height) along the curved(!) light path. And the refractive index itself is also a function of wavelength. The relation between refractive index and wavelength is called *atmospheric dispersion*.

- Due to atmospheric dispersion, any object that isn't exactly at zenith will appear to the observer as a small spectra. Since the refraction will be stronger for bluer wavelengths than for redder wavelengths, the blue side of that spectrum will point to zenith and the red side away from it (i.e., the blue side points *up*).

A good approximation for the difference between true and apparent zenith distance to $\zeta \simeq 75°$ (for $\lambda = 5500$Å, 750 mbar, 0 °C) is given by the expression:

$$(\zeta - z) = 60\farcs4 \tan\zeta - 0\farcs064 \tan^3\zeta \tag{9}$$

In view of Snell's law, the wavelength dependence of this difference between true and apparent zenith distance with respect to the dispersion at 5500Å is given by:

$$(\zeta - z)_\lambda = (\zeta - z)_{5500\text{Å}} \cdot \frac{(n_\lambda - n_{\rm vac})}{(n_{5500\text{Å}} - n_{\rm vac})} \quad \text{, where} \quad n_{\rm vac} \equiv 1 \tag{10}$$

Table 3 tabulates the atmospheric dispersion coefficient with respect to that at 5500Å for five wavelengths from the atmospheric cut-off in the near-UV to the near-IR (800 nm). From this table we can then evaluate how severe the distortion of the image of that astronomical object into a little spectrum will be, for a moderate zenith distance of, say, 45°:

$$\left.\begin{array}{rcl}(\zeta - z)_{3000} &=& 1.047 \cdot (60\overset{''}{.}4 \tan 45^\circ - 0\overset{''}{.}064 \tan^3 45^\circ) \\ (\zeta - z)_{8000} &=& 0.989 \cdot (60\overset{''}{.}4 \tan 45^\circ - 0\overset{''}{.}064 \tan^3 45^\circ) \\ \tan 45^\circ &\equiv& 1 \end{array}\right\} \Longrightarrow$$

$$(\zeta - z)_{3000-8000} = (1.047 - 0.989) \cdot (60\overset{''}{.}4 - 0\overset{''}{.}064) \simeq 3\overset{''}{.}50$$

And at an air mass of 2 (zenith angle of 60°), the length of the same spectrum would have grown to $\sim 6\overset{''}{.}0$.
When trying to obtain a spectrum of an object through a typical spectrograph slit of $\sim 1''$ width, one might completely lose the blue and red parts of the spectrum when the green light is centered in the slit (Filippenko 1982).

				$(I/I_0)_{\sec z}$		$(I/I_0)_{M(z)}$	
z	$\zeta - z$	$\sec z$	$M(z)$	5500Å	3200Å	5500Å	3200Å
0°	0″	1.000	1.000	0.83	0.40	0.83	0.40
10°	10″	1.015	1.015	0.83	0.39	0.83	0.39
20°	21″	1.064	1.064	0.82	0.38	0.82	0.38
30°	34″	1.155	1.154	0.81	0.35	0.81	0.35
40°	49″	1.305	1.304	0.78	0.30	0.78	0.30
50°	1′10″	1.556	1.553	0.75	0.24	0.75	0.24
60°	1′41″	2.000	1.995	0.69	0.16	0.69	0.16
70°	2′39″	2.924	2.904	0.58	0.069	0.58	0.070
80°	5′19″	5.76	5.60	0.34	0.005	0.35	0.006
85°	9′52″	11.47	10.40	0.12	0.000	0.14	0.000
90°	35′22″	∞	38.	0	0	0.0008	0.000

Table 2: Transmission of the Earth's atmosphere for different zenith distances z and wavelengths λ. The 2nd column lists the angular deviation, $(\zeta - z)$, for $\lambda = 5500$Å (as defined in Fig. 42).

λ (Å)	3000	4000	5000	6000	8000
$\left(\frac{n_\lambda - 1}{n_{5500} - 1}\right)$	1.047	1.014	1.001	0.995	0.989

Table 3: Atmospheric dispersion with respect to the dispersion at $\lambda = 5500$Å

- To avoid this, the spectrograph slit should be aligned with the *parallactic angle*: the angle at which the spectrograph slit is normal to the horizon.

Although in imaging applications one never covers such a large range in wavelength, a stellar image observed in, say, the B filter would still be elongated by $\sim 0\farcs5$ at z=45° and $\sim 1\farcs0$ at z=60°. This is sufficiently noticeable that large telescopes at good sites have implemented *atmospheric dispersion correctors*.

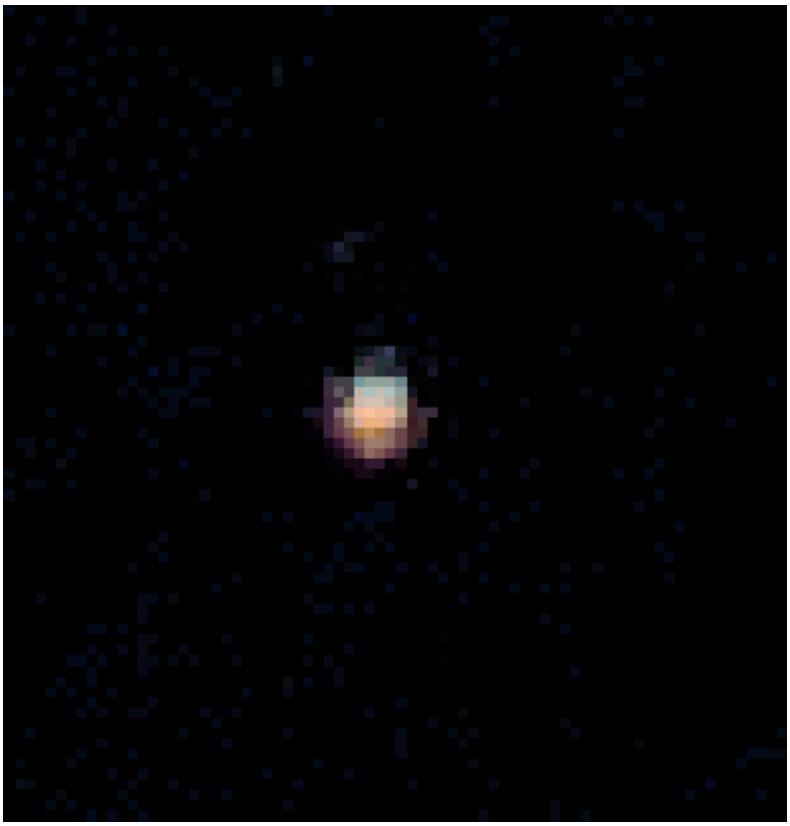

Figure 43: Example of the effects of atmospheric dispersion. The stellar image above was observed in a quick succession of BVR exposures at the 6.5 m Magellan '*Baade*' telescope with the IMACS instrument in Dec 2003 (R. Jansen & R. Windhorst) — before the atmospheric dispersion corrector was commissioned. The air mass was 1.33 (i.e., $z \simeq 41°$) and the seeing $0\farcs75$ (3.38 pixels) FWHM. Atmospheric dispersion caused a noticeable differential shift of the centroids of the B and R images on the CCD of $\sim 1\farcs0$.

5.3 Seeing and Scintillation

Along the path of a beam of light through the atmosphere, changes in intensity and direction not only occur due to extinction and refraction, but also due to *atmospheric turbulence*. Turbulence causes fluctuations in density, ρ, of the air — mostly as a result of temperature differences. Pressure fluctuations at small scales are negligible, since these would rapidly (at the sound speed!) smooth out. Temperature differences can be much longer-lived.

Let's estimate the effects due to such fluctuations. For changes in direction of a beam of light, any incremental change in refractive index, $\Delta n \equiv n - n_0$, is important. The refractive index of air is proportional to its density: $(n-1) \propto \rho$ and $n_{|\rho=0} \equiv n_{\rm vac} \equiv 1$. For an ideal gas, we have:

$$\frac{n-1}{n_0-1} \;=\; \frac{\rho}{\rho_0} \;=\; \frac{P/T}{P_0/T_0} \;=\; \frac{P\,T_0}{P_0\,T} \tag{11}$$

Because pressure differences can be ignored, $P = P_0$, from which follows (if we take $\Delta n/\Delta T \simeq {\rm d}n/{\rm d}T$):

$$\Delta n \;=\; -(n_0-1)\cdot\frac{T_0}{T^2}\cdot\Delta T \tag{12}$$

At sea level, $T \simeq T_0 \simeq 300\,{\rm K}$ and $n_0 = 1.000293$, which gives $\Delta n \simeq -10^{-6}\Delta T$. At altitude h, the pressure and, hence, the density of the air and its refractive index decrease, and they do so roughly exponentially:

$$\frac{n_h-1}{n_0-1} \;=\; e^{-h/H} \tag{13}$$

where H is the exponential scale height of the Earth's atmosphere ($H \simeq 8\,{\rm km}$). This then gives us (Eq. 8 $\rightarrow$ Eq. 6 $\Rightarrow$ Eq. 7):

$$\Delta n \;\simeq\; -10^{-6}\cdot e^{-h/H}\cdot\Delta T \tag{14}$$

As it turns out, turbulence in the atmosphere results in *two* seperate effects:

1. it causes changes in direction of the incident beam of light, which we call *seeing*;
2. it causes fluctuations in intensity, which we call *scintillation*.

In the following, we will consider a much simplified model of a single turbulence element (*turbulence cell*) of diameter L and refractive index $n_0 + \Delta n$, embedded in the ambient air (with index n_0), and reduce the problem to the two-dimensional geometry illustrated in Fig. 44.

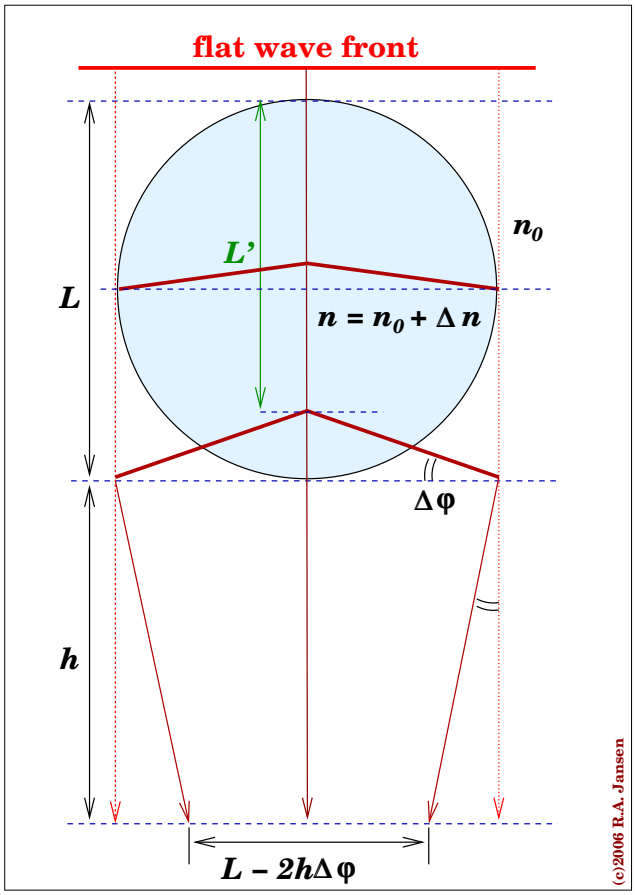

Figure 44: Deformation of a flat wave front by a turbulence cell with diameter L and refractive index $n = n_0 + \Delta n$, embedded in a medium with refractive index n_0.

a) Seeing. If a flat wave front is incident on a turbulence element, then the wave front will be distorted as a result of the difference in optical path length for refractive index $n_0 + \Delta n$ compared to that for the ambient index of n_0:

$$L' \;=\; L \cdot \frac{n_0}{n_0 + \Delta n}$$

From which follows:

$$\Delta\phi \;=\; \frac{L - L'}{\frac{1}{2}\,L} \;=\; 2\,\frac{\Delta n}{n_0 + \Delta n} \;\simeq\; 2\,\Delta n \quad (\text{since } n_0 + \Delta n \simeq n_0 \simeq 1\,)$$

Hence:

$$\Delta\phi \;\simeq\; 2\,\Delta n \;=\; -2 \times 10^{-6} \cdot e^{-h/H} \cdot \Delta T \qquad (15)$$

This explains why seeing is dominated by the turbulence within the first tens of meters above the telescope: due to the exponential decrease in density, $\Delta\phi$ rapidly decreases with increasing altitude h. Moreover, ΔT also diminishes with altitude.

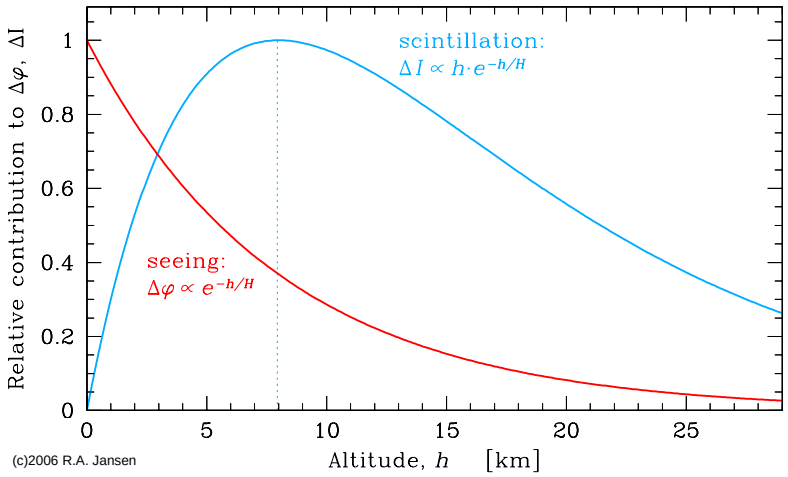

Figure 45: Relative contributions to the seeing and scintillation as a function of the altitude of the turbulence elements. The exponential scale height of the atmosphere is 8 km. The maximum contributions to seeing and scintillation are at quite different altitudes. This means that *seeing and scintillation are uncorrelated*.

b) Scintillation. The intensity fluctuations as observed at ground-level due to the passage through a turbulence element is proportional to the change in beam diameter, hence $\Delta I \propto h \cdot \Delta\phi$. So now we're dealing with $h \cdot e^{-h/H}$ instead of with $e^{-h/H}$. This means that the maximum effect of scintillation occurs higher in the atmosphere. As it turns out, at an altitude of ∼8 km.

- The difference in dominant altitude implies that *seeing and scintillation are uncorrelated*.

If D denotes the diameter of a telescope, then:

- if $D \lesssim L$: most of the time only one turbulence element will be within the beam $\Longrightarrow$ the entire (sharp) source image will be displaced but unsmeared at any one time, but that image will show strong fluctuation in intensity (i.e., scintillation);
- if $D \simeq L$: At $h \simeq 0$ and for $\Delta T = 1^\circ \quad \Longrightarrow \quad \Delta\phi \;=\; 2 \times 10^{-6}\,\mathrm{rad} \;=\; 0\farcs41$;
- if $D \gg L$: several turbulence cells will be within the beam at any one time: their combined displacements will smear the source image to a seeing-disk, but the net amplitude of scintillation is less severe.

A typical diameter of a turbulence cell of ∼10 cm at an altitude of 8–10 km subtends an angle of:

$$\tan\theta \simeq \theta = \frac{10}{(8-10)\times 10^5} \cdot \frac{180}{\pi} \cdot 3600 = 2\farcs58 - 2\farcs06$$

Because most planets subtend angles larger than $2\farcs6$ (Mars: 5–$15''$, Jupiter: 30–$40''$, Venus: 10–$60''$), planets do not scintillate. Twinkle, twinkle is a star after all...

5.4 Physical Quantities, units, photometric systems

The radiation energy measured of an astronomical source depends on several telescope/instrument parameters:

T_{atm}	transmission of the Earth's atmosphere
T_{tel}	transmission of the telescope plus instrument
D	diameter of the primary aperture of the telescope (i.e., the total light collecting surface area). In case the primary mirror has a central hole or if the telescope has a central obstruction, $D^2 = D_{out}^2 - D_{in}^2$.
$\Delta\lambda$	effective bandpass of the instrument (as modified by any filters inserted).
Δt	the effective integration time of the measurement
QE	the quantum efficiency, i.e., the probability that a photon will be absorbed and detected by the detector.

$$\log E_\lambda = -0.4 \cdot m_\lambda + zp_\lambda + \log T_{atm} + \log T_{tel} + \log \frac{\pi}{4} D^2 + \log \Delta\lambda + \log \Delta t + \log QE$$

where zp_λ denotes an absolute zeropoint relating the magnitude m_λ to the energy observed at wavelength λ of the star Vega (or that of a spectrum that is flat as a function of frequency ν, in case of the AB magnitude system). It can also be useful to express the above equation in terms of the number of detected photons (with $E = h\nu$, and with zp' containing the relevant conversion factors):

$$\log N_\lambda = -0.4 \cdot m_\lambda + zp'_\lambda + \log T_{atm} + \log T_{tel} + \log \frac{\pi}{4} D^2 + \log \Delta\lambda + \log \Delta t + \log QE$$

- For the human eye we have $T_{tel} = 0.90$, $D = 0.8\,\mathrm{cm}$, $\Delta\lambda = 1000\,\text{Å}$, $\Delta t = 0.1\,\mathrm{sec}$ and $QE = 0.01$. At zenith (*air mass* equal to 1), during conditions of excellent visibility, $T_{atm} = 0.80$ at $\lambda \sim 5500\,\text{Å}$. If the faintest stars visible to the unaided eye are of magnitude $m_V = 5.5$, and if $zp'_V = 7.02$ (assuming flux density units of $\mathrm{W\,m^{-2}\,Å^{-1}}$), what then is the number of detected photons per second?

A:
$$\log N_V = -0.4 \cdot 5.5 + 7.02 + \log 0.80 + \log 0.90 + \log \frac{\pi}{4}(0.008)^2 + \log 1000 + \log 0.1 + \log 0.01 = 0.38 \quad \longrightarrow \quad N_V \simeq 2 \text{ photons/sec}$$

- For a 1 m-telescope with a photomuliplier tube photometer behind a V filter, $T_{tel} = 0.70$, $\Delta\lambda = 500\,\text{Å}$, $\Delta t = 100\,\text{sec}$ and $QE = 0.07$ at $\lambda = 5500\,\text{Å}$. Assuming the same atmospheric conditions and the same $m_V = 5.5$ star, how many photo-electrons are detected per unit time?
  If $1\,e^-$/sec is still detectable, what then is the limiting magnitude of this telescope-instrument?

A:
$$\log N_V = -0.4 \cdot 5.5 + 7.02 + \log 0.80 + \log 0.70 + \log \frac{\pi}{4}(1.0)^2 + \log 1000$$
$$+ \log 100 + \log 0.07 = 8.0 \quad \longrightarrow \quad N_V \simeq 1.0 \times 10^8 \text{ photons/sec}$$
Hence:
$$m_{\text{lim}} = 5.5 - 2.5 \cdot \log(\frac{1}{1.0 \times 10^8}) = 25.5 \text{ mag in } V.$$

---

6 Photometric Systems & Astronomy with CCDs

6.1 Vega vs. AB magnitudes

Earlier, we have introduced magnitudes as logarithmic units that express relative brightnesses. Yet, the apparent magnitude of the Sun in the V filter is understood to be -26.76 ± 0.02 mag and the surface brightness in mag arcsec^{-2} in a galaxy was plotted as a function of radius as if they were physical units on some absolute scale. So how did we get from what is inherently an ambiguous relative unit to a unit with an absolute physical meaning?

According to the definition of *magnitude*, we have:

$$m_1 - m_2 \equiv -2.5 \log \frac{f_1}{f_2} = -2.5 \log f_1 + 2.5 \log f_2 \tag{16}$$

If m_2 were the magnitude that corresponds to a known flux density f_2 in physical units of ergs s^{-1} cm^{-2} Å^{-1} (in the case of f_λ) or of ergs s^{-1} cm^{-2} Hz^{-1} (for f_ν) — in particular if $m_2 \equiv 0$ for that known flux density —, then the term $2.5 \log f_2$ becomes equivalent to an absolute *zeropoint* (zp) for the magnitude scale:

$$m_1 = -2.5 \log f_1 + zp \tag{17}$$

6.1.1 Vega magnitudes

The ultimate standard and reference for all classical broad-band photometry is the star α Lyrae (i.e., Vega). By substituting the flux density in a given filter of Vega for f_2 in Eq. (1) and dropping the subscripts "1", we obtain:

$$m - m_{\rm Vega} = -2.5 \log f_\lambda + 2.5 \log f_{\lambda,{\rm Vega}}$$

However, this would still leave us with $m_{\rm Vega}$. The final step, then, to yield a usable physical scale is to set:

$$m_{\rm Vega} \equiv 0 \quad \text{in all filters } \textbf{\textit{per definition}} \quad \Rightarrow m - m_{\rm Vega} \equiv m$$

Hence:

$$m = -2.5 \log f_\lambda + 2.5 \log f_{\lambda,{\rm Vega}} \tag{18}$$

where the last term is the zeropoint, $zp(\lambda)$, on the Vega magnitude system. For a flux density outside the Earth's atmosphere of $(3.59 \pm 0.08) \times 10^{-9}$ ergs s^{-1} cm^{-2} Å^{-1} at 5480Å, this zeropoint $zp(5480\text{Å}) = -21.112$.

As a corollary of the requirement that $m = 0$ at all wavelengths, the *color* of Vega in any pair of filters is 0, as well.

Fig. 46(a) shows the spectrum of Vega as a function of wavelength. Note, that the flux density of Vega varies greatly over the UV–near-IR regime, yet the magnitude corresponding to that flux density is equal to 0 at every wavelength.

Since it is impractical for every observer to observe Vega, a reference set of several dozens of, generally fainter, secondary standard stars was observed. During WW II, Johnson made extensive photo-electric observations, spanning many years, through standard apertures in his UBV filter system. The UBV filter system was the first known standardized photometric system. The filter set was later extended toward the near- and mid-IR with $RIJK$ and L filters.

As progressively more precise magnitude and color measurements of an increasingly large number of stars, sampling a large range in colors, were placed onto the standard system of Johnson, slight inconsistencies in magnitude and colors became apparent. To remain as consistent as possible with earlier work, the V magnitude of Vega had to be adjusted slightly upward from exactly zero. The current best estimate is $m_V = +0.035 \pm 0.012$ mag for $f_{\lambda,V} = (3.593 \pm 0.084) \times 10^{-9}\,\mathrm{ergs\,s^{-1}\,cm^{-2}\,\AA^{-1}}$ (Colina & Bohlin 1994).

More recently, Landolt (1992) published photo-electric $UBVR_cI_c$ photometry of a large number of equatorial (Dec $\sim 0^\circ$) fields in which stars of very different colors are relatively close together on the sky (within $\sim 5'$–$15'$ on the sky). The transformations of his photometry onto the original system, again meant very slight changes to the original system. When using Landolt standard stars to photometrically calibrate you CCD images, you are actually calibrating onto Landolt's system, and not quite the original Vega system of Johnson.

Even more recently, Stetson (2000) published a very large database of BVR_cI_c standard stellar magnitudes within the Landolt equatorial fields and in many other common fields across the sky based on (largely archival) CCD observations. His photometry is reduced onto the Landolt system, but his measurement method differed. Whereas Landolt used fixed circular apertures of $14''$ with his photomultiplier tube, and hence often also measured the flux of adjacent fainters stars within that aperture, Stetson used the technique of PSF fitting, where each star is measured seperately.

- ▷ *When calibrating CCD images, care needs to be taken to reproduce as close as possible the measurement method employed to obtain the original standard star photometry.*

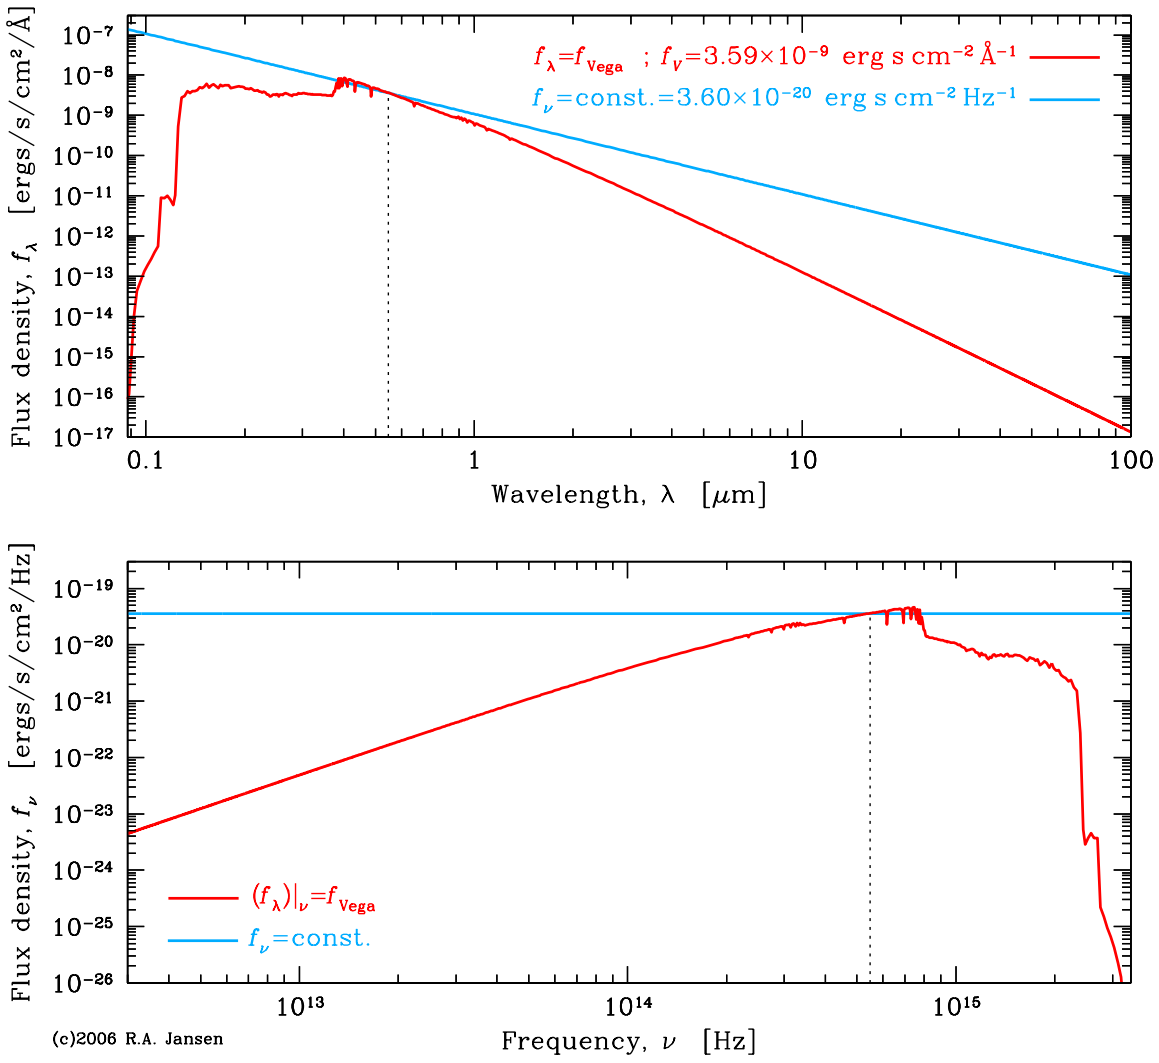

Figure 46: Comparison of the spectrum of α Lyrae (Vega) and a spectrum that is flat in f_ν. In the top panel, both are plotted as f_λ (in ergs s^{-1} cm^{-2} Å^{-1}) versus wavelength λ (in μm), while in the bottom panel they are presented as f_ν (in ergs s^{-1} cm^{-2} Hz^{-1}) versus frequency ν (in Hz). Both spectra are equal at the effective wavelength of the V filter, at 5480Å. The mid- to far-IR portion of the spectrum of Vega was replaced by the stellar atmospheric model of Kurucz (1979): an actual IR spectrum of Vega would show a significant additional, non-photospheric component due to its circum-stellar debris disk.

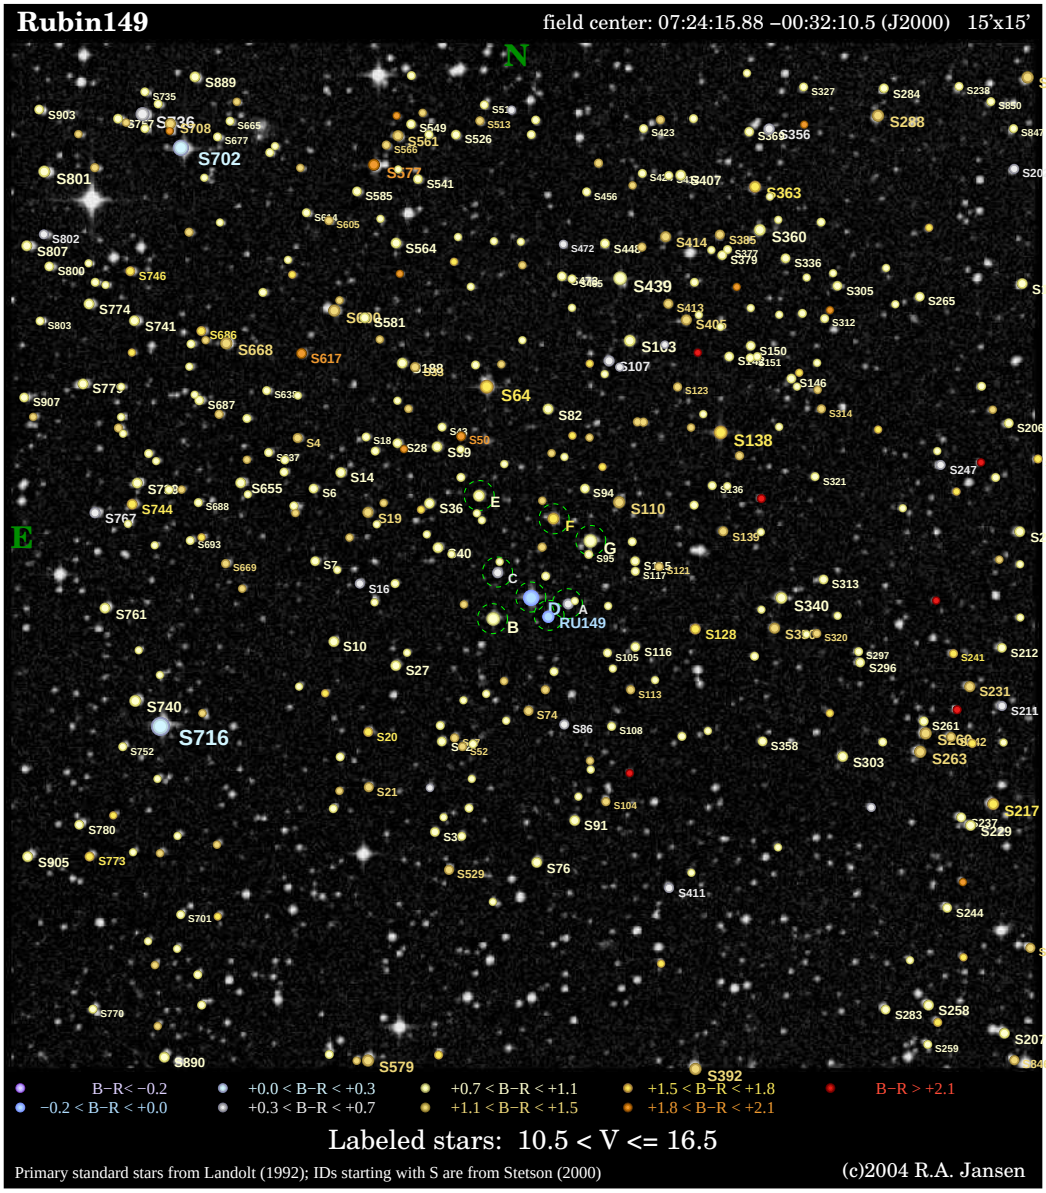

Figure 47: In standard field Rubin 149, Landolt (1992) published $UBVRI$ filter photo-electric photometry of the 8 stars labeled "Ru149" and A through G. The circular apertures used with his photomultiplier tube had a diameter of 14″. Since this field is densely populated with stars, such apertures also included neighboring stars. To use Landolt's magnitudes, one has to measure the total signal within synthetic apertures of the same diameter. Stetson (2000) published standard star photometry for all labeled stars in the larger general field of the Rubin 149 asterism. His measurements include many fainter stars that are observable even with large aperture telescopes. His measurements are reduced onto the Landolt system, but his measurement method differed. Because he fitted the stellar PSF, his published magnitudes do *not* include signal from neighboring stars. To use Stetson's photometry, one has to reproduce the same method (see also Bessell 2005).

6.1.2 AB magnitudes

To avoid the problems with Vega magnitudes — that the flux density that corresponds to $m = 0$ differs at every wavelength, and that the flux of Vega can become exceedingly small at wavelengths outside the UV–near-IR regime[1] —, another magnitude system, the *AB* or *spectroscopic* or *natural magnitude system*, was deviced (Oke & Gunn 1983). In the AB magnitude system, the reference spectrum is *a flat spectrum in* f_ν:

$$f_\nu \equiv \text{const.} \quad [\text{ergs}\,\text{s}^{-1}\,\text{cm}^{-2}\,\text{Hz}^{-1}].$$

That constant is per definition such that in the V filter: $m_V^{Vega} \equiv m_V^{AB} \equiv 0$ (or more accurately: $f_\nu\,\mathrm{d}\nu \equiv f_\lambda\,\mathrm{d}\lambda$ when averaged over the V filter, or at the effective wavelength of the V filter, $\lambda_{\rm eff} = 5480$Å.

▷ *Note that the AB magnitude system is expressed in* f_ν *rather than* f_λ!

The flux density in f_ν is related to the flux density in f_λ by:

$$f_\nu\ [\text{ergs}\,\text{s}^{-1}\,\text{cm}^{-2}\,\text{Hz}^{-1}] \;=\; \frac{\lambda^2}{c} \cdot 10^8 \cdot f_\lambda \quad [\text{ergs}\,\text{s}^{-1}\,\text{cm}^{-2}\,\text{Å}^{-1}]$$

where we used $\lambda\,\nu \equiv c$ and transformed from $f_\lambda\,\mathrm{d}\lambda$ to $f_\nu\,\mathrm{d}\nu$. The factor 10^8 is included, because the natural units of wavelength are cm, not Å. Converting the Vega magnitude zeropoint gives:

$$m \;=\; -2.5\,\log f_\nu \;-\; (48.585 \pm 0.005) \tag{19}$$

where the exact value of the zeropoint depends somewhat on the literature source (e.g., Hayes & Latham 1975; Bessell 1988,1990), based on a magnitude at 5556Å for Vega of 0.035–0.048 mag and a corresponding flux density of $(3.56\text{–}3.52)\times10^{-20}$ ergs s^{-1} cm^{-2} Å^{-1} .

The AB magnitude system is also called the spectroscopic magnitude system, because with its constant zeropoint, it is useable at any wavelength in bandpasses of any width, and hence, also for narrow-band imaging and spectroscopy. In ground-based broad-band imaging, however, the Vega magnitude system is still the most common system. Unless explicitly noted otherwise, one should assume Vega magnitudes.

[1] and not even due to the stellar atmosphere in the mid-IR–far-IR: there one actually detects the circumstellar debris ring!

6.2 Filters, filters and yet more filters

- Johnson-Morgan/"Johnson" — U, B, V, R, I broad-band filters; extended to the near-IR with J, K, and L. The main disadvantage of the U filter was that it's blue cut-off was mainly determined by the transparency of the atmosphere (rather than by the glass of the filter).

- Kron-Cousins-Glass/"Cousins" — R_c, I_c broad-band filters; better behaved in their red-tail, better positioning in wavelength with respect to UBV.

- Bessell (1979,1990), Bessell & Brett (1988)/"Bessell" — better characterization and formalization of what the $UBVR_cRI_cIJHKLM$ broad-band bandpass curves (resulting from filter plus detector) *should* look like.

- Strömgren (1966) filters — u, v, b and y; medium-band filters, specifically designed for stellar astrophysics (hot vs. cool stars; stellar composition): u and v straddle the Balmer break (actually the Ca II H+K break at $\sim$4000Å), and $u-v$ and $v-y$ colors provide the strength of that break and the continuum slope redward of the Balmer break.

- Washington filters (defined by G. Wallerstein, developed by Canterna (1976), calibrated on the Geisler (1996) system of CCD standards; see also Bessell (2001))/"Washington system" — C, M, T_1 and T_2, specifically designed for metallicity studies in old stellar populations.

- Straižys et al. (1966); Straižys & Zdanavičius (1970)/"Vilnius system" — U, P, X, Y, Z, V, and S. Mainly used for stellar classification.

- Gunn, Thuan & Gunn / "Gunn" — u, g, r (later also i, z), have filters with transmission curves with steeper cut-on and cutt-off; the precursor of the Sloan and various *HST* filters.

- 2MASS filters; Jarret et al./"2MASS system" — J, H, K_s; by reducing the effective wavelength of the K filter from 2.2 to 2.15 μm and designing a steeper red cut-off, the sky background in K_s is significantly darker than in the Johnson K filter.

- Sloan Digital Sky Survey/"Sloan" — u', g', r', i', and z'. Square filter transmission curves with minimal overlap and minimal gaps between the filters. Optimal broad-band filter set for photometric redshifts and quasar searches.

- Beijing-Arizona-Taiwan-Connecticut filter set; spectro-photometrically calibrated by Yan (2000)/"BATC" — System of 16 medium-band filters in near-UV through near-IR that avoid night-sky emission lines. Optimal medium-band filter set for photometric redshifts and (high-redshift) emission-line object searches.

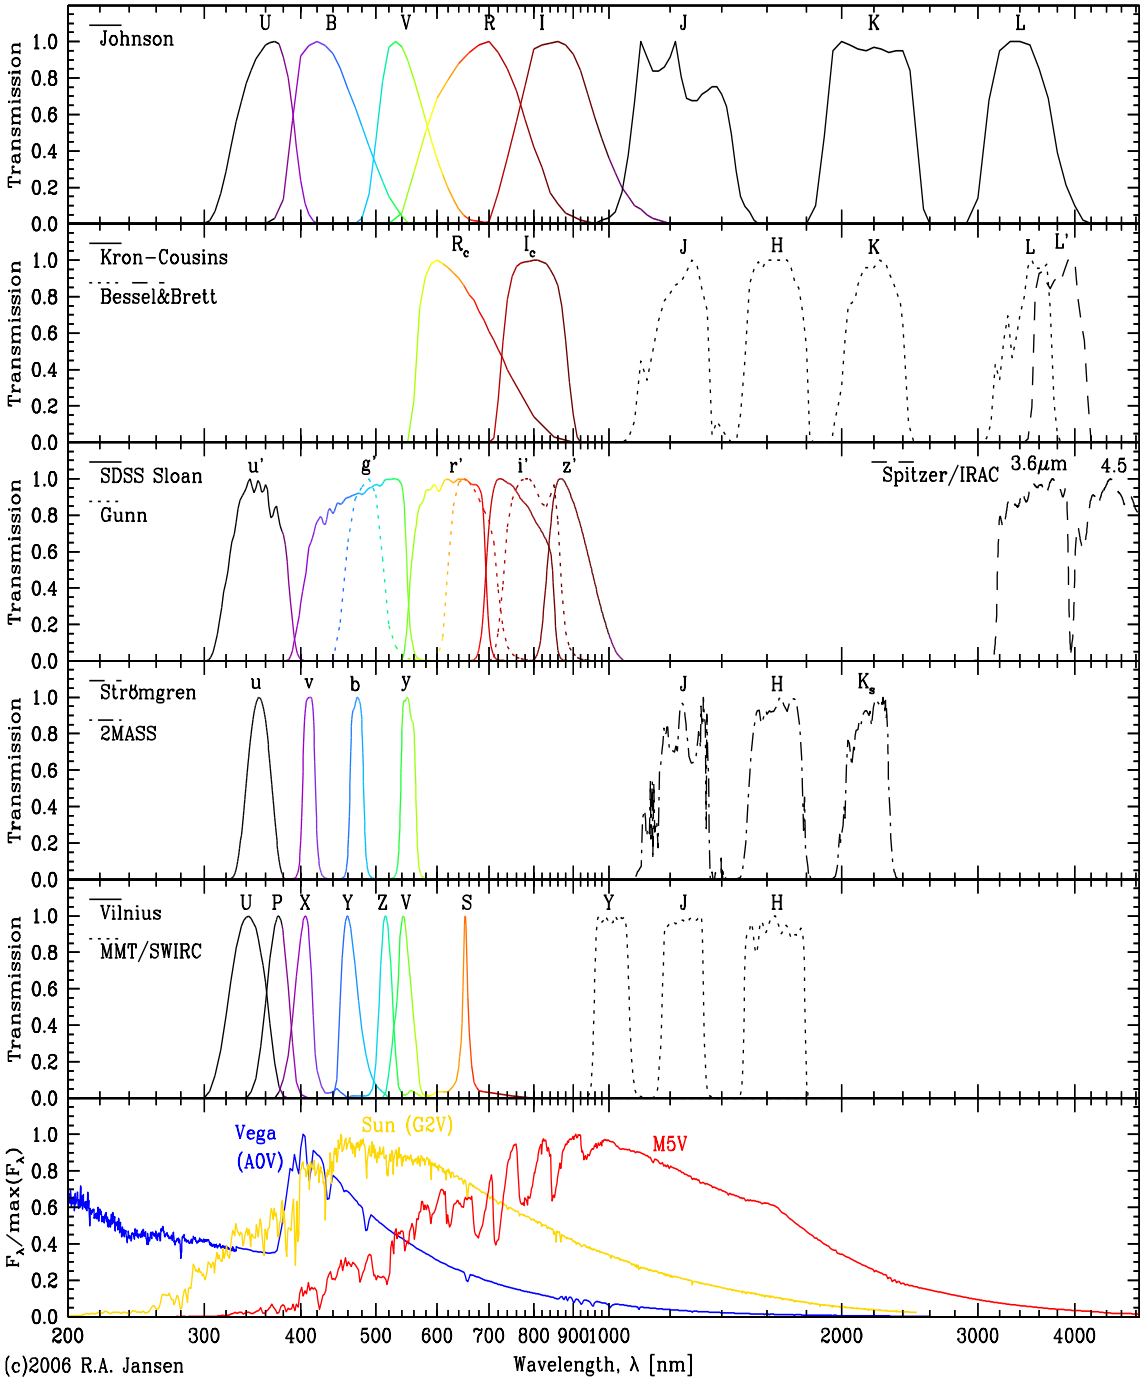

Figure 48: Overview of various filter sets (as labeled) and comparison with the spectra of an A0V (Vega), G2V (Sun) and M5V star.

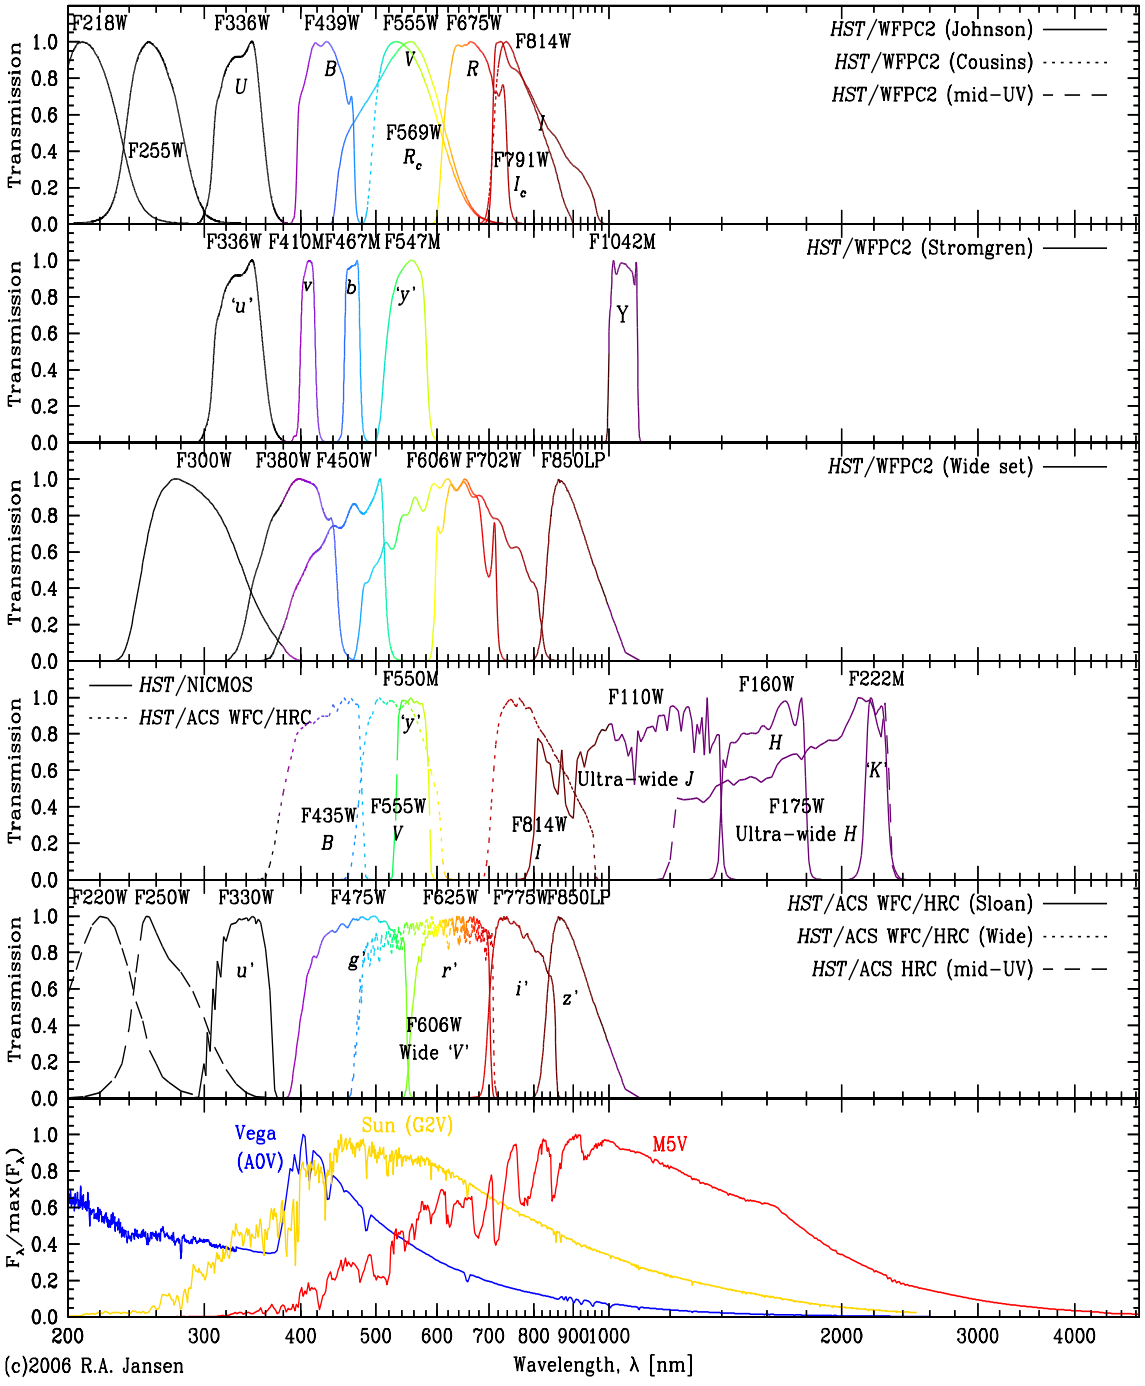

Figure 48: [*Continued*] Overview of various filter sets (as labeled).

7 Types of astronomical observations with CCDs

- *Imaging* — faithfully (in some way) record the spatial (angular) distribution of brightness on the sky
- *Astrometry* — faithfully record the relative or absolute positions of sources on the sky (regardless of brightness)
- *Photometry* — faithfully record the relative or absolute brightness of sources on the sky (regardless of position and possibly regardless of how flux is distributed on the sky)
- *Spectroscopy* — faithfully record the relative or absolute flux density as a function of wavelength or frequency
- *Kinematics* — faithfully record the relative or absolute velocities of objects or parts thereof with respect to a suitable standard of rest
- *Polarimetry* — faithfully record (relative) polarizations (degree and linear/circular)
- *Interferometry* — faithfully record (relative) phases or phase distributions of one or more sources.
- *Photon timing* — faithfully record (relative) arrival times of photons from one or more sources.
- mixtures of any of the above, e.g., *surface photometry*, *spectrophotometry*, *integral field spectroscopy*, *combining both astrometry and photometry*, etc...

For imaging, one has to consider the *plate scale* and *geometric distortions* (the latter particularly for off-axis instruments, but depending on the telescope- and instrument-design, even on-axis instruments may show significant geometric distortions (see, e.g., Fig. 36).

For photometry, one has to consider *detection* vs. *measurement* (cf. imaging), *aperture* photometry, *curve of growth* total photometry, *PSF fitting*, or *differential* photometry. Calibration onto an absolute flux system (*AB magnitudes*) or relative system (e.g., with respect to α Lyrae (Vega)).

7.1 Physical Quantities, units, photometric systems

The radiation energy measured of an astronomical source depends on several telescope/instrument parameters:

T_{atm}	transmission of the Earth's atmosphere
T_{tel}	transmission of the telescope plus instrument
D	diameter of the primary aperture of the telescope (i.e., the total light collecting surface area). In case the primary mirror has a central hole or if the telescope has a central obstruction, $D^2 = D_{out}^2 - D_{in}^2$.
$\Delta\lambda$	effective bandpass of the instrument (as modified by any filters inserted).
Δt	the effective integration time of the measurement
QE	the quantum efficiency, i.e., the probability that a photon will be absorbed and detected by the detector.

$$\log E_\lambda = -0.4 \cdot m_\lambda + zp_\lambda + \log T_{atm} + \log T_{tel} + \log \frac{\pi}{4} D^2 + \log \Delta\lambda + \log \Delta t + \log QE$$

where zp_λ denotes an absolute zeropoint relating the magnitude m_λ to the energy observed at wavelength λ of the star Vega (or that of a spectrum that is flat as a function of frequency ν, in case of the AB magnitude system). It can also be useful to express the above equation in terms of the number of detected photons (with $E = h\nu$, and with zp' containing the relevant conversion factors):

$$\log N_\lambda = -0.4 \cdot m_\lambda + zp'_\lambda + \log T_{atm} + \log T_{tel} + \log \frac{\pi}{4} D^2 + \log \Delta\lambda + \log \Delta t + \log QE$$

- For the human eye we have $T_{tel} = 0.90$, $D = 0.8\,\mathrm{cm}$, $\Delta\lambda = 1000\,\text{Å}$, $\Delta t = 0.1\,\mathrm{sec}$ and $QE = 0.01$. At zenith (*air mass* equal to 1), during conditions of excellent visibility, $T_{atm} = 0.80$ at $\lambda \sim 5500\,\text{Å}$. If the faintest stars visible to the unaided eye are of magnitude $m_V = 5.5$, and if $zp'_V = 7.02$ (assuming flux density units of $\mathrm{W\,m^{-2}\,Å^{-1}}$), what then is the number of detected photons per second?

A:
$$\begin{aligned} \log N_V &= -0.4 \cdot 5.5 + 7.02 + \log 0.80 + \log 0.90 + \log \frac{\pi}{4}(0.008)^2 + \log 1000 \\ &\quad + \log 0.1 + \log 0.01 \;=\; 0.38 \quad \longrightarrow \quad N_V \simeq 2 \text{ photons/sec} \end{aligned}$$

- For a 1 m-telescope with a photomuliplier tube photometer behind a V filter, $T_{tel} = 0.70$, $\Delta\lambda = 500\,\text{Å}$, $\Delta t = 100\,\text{sec}$ and $QE = 0.07$ at $\lambda = 5500\,\text{Å}$. Assuming the same atmospheric conditions and the same $m_V = 5.5$ star, how many photo-electrons are detected per unit time?
  If $1\,e^-$/sec is still detectable, what then is the limiting magnitude of this telescope-instrument?

A:
$$\log N_V = -0.4 \cdot 5.5 + 7.02 + \log 0.80 + \log 0.70 + \log \frac{\pi}{4}(1.0)^2 + \log 1000$$
$$+ \log 100 + \log 0.07 = 8.0 \quad \longrightarrow \quad N_V \simeq 1.0 \times 10^8 \text{ photons/sec}$$
Hence:
$$m_{\text{lim}} = 5.5 - 2.5 \cdot \log(\frac{1}{1.0 \times 10^8}) = 25.5 \text{ mag in } V.$$

---

8 CCD Image Processing & Analysis: A Practical Example from Start to Finish

The purpose of CCD image processing is to separate instrumental signals and imperfections (which includes CCD, instrument, telescope and atmosphere) from the astronomical signals. To do so requires (a large amount of) calibration observations, that observers must ensure to obtain along with the data on their science targets.

In this section, we will discuss the case of CCD imaging and point source photometry at a telescope with a camera with a single CCD detector. In the case of large CCD mosaic cameras, one has to rely to a much greater degree on calibration products provided to the observer by the observatory staff, and it may not always be possible to obtain the calibration data (in sufficient quantity or even at all) that one would ideally want. The same is true for very large aperture telescopes run in a *queue scheduling* mode. Telescopes like ESO's VLT have as their formal policy to provide calibration data sufficient to allow calibration of the science product to ∼10%. If larger accuracy is required, then *a calibration program on a smaller telescope* should be initiated in support of the program on the larger telescope.

8.1 Calibration Data to be Acquired at the Telescope

For CCD imaging applications of any type, the following calibration frames *must* be obtained at the telescope. Some will cut into the available time for observations on the scientific targets, while others may be taken in the afternoon before sunset and/or in the morning after sunrise.

- **Bias frames** (*biases*) — Bias frames are 0 sec dark exposures; the shutter stays closed. They are required to determine:
  - the *read-noise* and *gain* of the CCD (see §2.5),
  - whether the bias level shows structure across the *image region* (i.e., the illuminated pixels) of the CCD, and
  - whether the bias levels in the *overscan strip*(s) and in the image region are identical.

  Take at least 20 bias frames both before opening up in the afternoon and after closing down in the morning (this can usually be scripted, or with a command that takes as its argument the number of frames desired).

  △ Beware, that the first (few) bias frames taken in the afternoon — after the CCD has been idle for many hours —, and the first one or two frames after high light level exposures (e.g., flats) should be discarded (see also §2.7 *Transient Effects*).

- **Dark frames** (*darks*) — Dark frames are integrations with the shutter closed. Darks are required to determine:
  - the average bulk dark rate, *dc*, of the CCD [in e^-/pix/hr];
  - whether there are pixels with significantly ($\geq 4\,\sigma$) higher dark rates (i.e., warm/hot pixels or columns);
  - whether there is structure in the dark rate across the image region (e.g., a gradient in the direction of the often slightly warmer amplifier).

  The (maximum) integration time should be at least as long as your longest science exposure. Whenever possible during an observing run, take multiple sets of integrations of different lengths (e.g., 10×600 s, 7×900 s, 5×1200 s and 5×1800 s), to allow consistency checks and ensure the (*small*) dark rate is truly detected. This will likely not be possible on a single-night run. Darks can be taken in the afternoon or in the morning, or when clouded out during the night.

  △ Beware, that one should take darks only following a series of biases, to reduce the impact of any transient effects.

  △ If the instrument has light leaks, it may only be possible to obtain reliable darks at night when clouded out — and with all lights in the dome off and with the dome closed.

- **Dome/Pupil/Internal flat exposures** (*dome/pupil/quartz flats*) — Dome flats are exposures of an evenly illuminated screen on the inside of the telescope dome, or of the inside of the dome itself. Pupil flats are only possible with some telescope designs that *have* a pupil between primary and secondary mirrors (in which case they are better than dome flats, because even illumination is not a concern). Most spectrographs have a facility to take *internal incandescence flats*, using a quartz halogen lamp and integrating sphere. Flats are required to determine:
  - the pixel-to-pixel variations in effective sensitivity of the CCD (see Fig. 58), whether intrinsic to the CCD (QE) or extrinsic (e.g., dust particles)

  These exposures are usually taken in the afternoon before opening up. To be able to correct for such sensitivity variations down to 0.1% (i.e., $S/N = 1000$), one needs to accumulate a total of at least 1,000,000 e^- per filter.

  △ Aim for a level per flat frame of $\sim\frac{1}{2}$ the lesser of the full-well capacity or A/D-saturation level.

  △ Since variations in sensitivity are color-dependent and since flat field lamps rarely ever have the same spectrum as your science object or even as the sky background, the resulting corrections are only approximately correct.

- **Shutter shading exposures** (*shutter flats*) — Dome or pupil flat exposures with a sequence of exposure times starting at very short integrations (e.g, 0.1 sec) up to exposures where shutter shading should no longer be significant (e.g., ∼5–10 sec). The exposures are required to determine:

  - the effective exposure time across the CCD for a given commanded exposure time *only if the science or calibration observations require short exposures (less than a few seconds).*

  Although shutters are fast, they have to cover a physically large distance. This means that the effective exposure time of pixels that first became illuminated is slightly longer than that of the pixels that became illuminated last. Linear shutters tend to be faster than diaphragm shutters.

  Typically, one needs to take shutter flats only once per observing run in the afternoon, along with regular dome/pupil flats

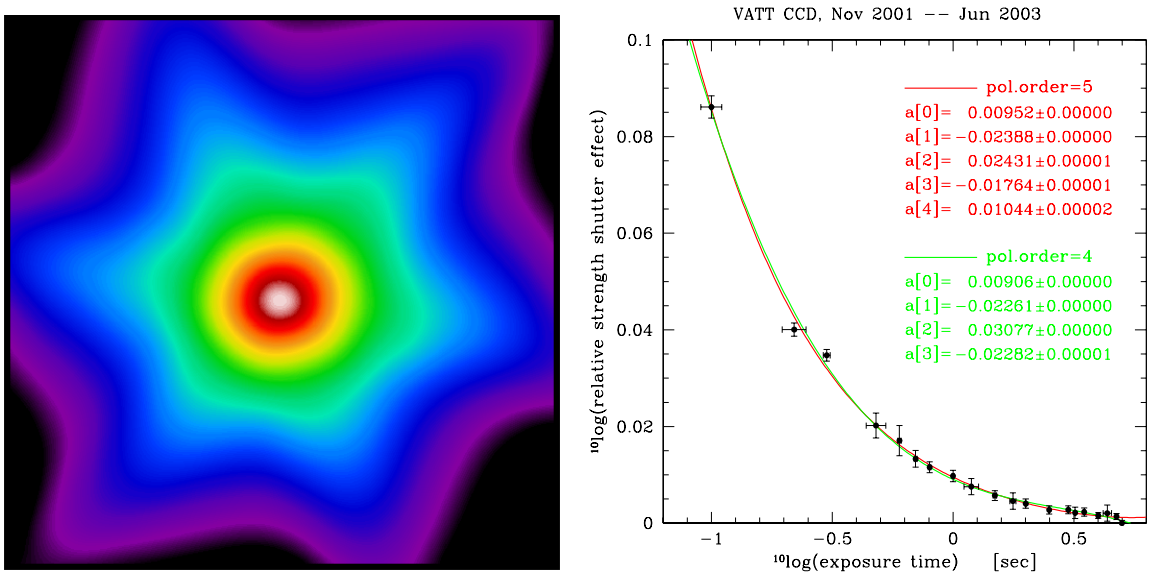

Figure 49: Example of the shading by the diaphragm shutter in the VATT CCD camera. Only the center of the CCD will be exposed during the full commanded integration time — the effective exposure time decreases toward the edges of the CCD in a hexagonal pattern that corresponds to the six blades of the shutter. Once the shape of the shading pattern has been established, a "strength" or scale factor can be fit as a function of commanded exposure time. By dividing the CCD images by the appropriately scaled pattern of the left panel, one can correct for shutter shading. For exposure times longer than $\sim 10^{0.5} \simeq 3$ sec, shutter shading becomes negligible for this particular shutter.

- **Twilight sky flat exposures** (*twilight flats*) — Exposures of the evening and morning twilight sky. They are required to determine:

  - the large-scale *illumination* of the CCD. Since the path followed by the light when observing the inside of the dome differs from that when observing the sky, and since the illumination of the dome by the flat field lamps can show gradients in brightness, twilight sky flats are needed to correct for this to first or second order. The scale of brightness gradients on the twilight sky tends to be much larger than the usable field-of-view of a telescope.

    △ Since twilight in Arizona lasts only ∼40 min, you rarely have sufficient time to obtain $\gtrsim$1,000,000 e^- per filter and measure the pixel-to-pixel variations at high S/N in the twilight flats

    △ Since you constantly have to adjust the exposure times to the changing brightness of the twilight sky, the variations in signal level between flats in the same filter are much larger than those in dome flats. This means that rejection algorithms are less efficient, and that the S/N in combining twilight flats does not scale as $\sqrt{N_{e^-}}$ (or $\sqrt{S_S}$).

    △ *Dither your exposures*, i.e, offset the telescope by ∼10″–20″ between each exposure. This will ensure that stars are rejected in the combination of the twilight flats.

    △ *Don't point your telescope at zenith*: Alt-AZ telescopes cannot track there, so stars in the darker portion of twilight will trail and affect more pixels than they should.

- **Night sky flat exposures** (*night sky flats*) — Exposures of the night sky. They are required to determine:

  - the large-scale *illumination* of the CCD *only if the very best quality flat fielding correction is required.*

  Taking night sky flats comes at a large cost, since it directly reduces the amount of time available to observe your science targets. The signal level in night sky flats tends to be a factor ∼10 or more lower than in twilight flats. For some programs it may be possible to use the exposures on the science targets to construct night sky flats.

  △ If you intend to use the science target exposures, then you must dither the exposures in a non-repeating pattern over a sufficient distance to ensure that the targets don't appear (or overlap) in the same spot on the CCD every time!

- **Orientation exposure** — A partly trailed exposure of a star. Such frame is used to determine:
  - which is E and which is W in CCD pixel coordinates.
  - how well aligned the pixel rows (or columns) are with the celestial axes, i.e., the rotation with respect to E–W.

  Unless the instrument is taken off the telescope during a run, one only needs to take one orientation exposure.

  To take such an exposure, center on a star that is bright but won't saturate in a 10 s exposure, then start an unguided exposure of approximate length

$$\sim 10\,\mathrm{s} + \frac{1}{2}\, n_{\mathrm{pix}} \cdot ps \cdot \frac{\cos(\mathrm{Dec})}{15} \quad [\mathrm{sec}] \ ,$$

  where n_{pix} denotes the size of the CCD along a row (or column), ps is the pixel scale (in $''$/pix), and $15\,''$/s is the sidereal rate (15°/hour) at the equator, which needs to be corrected for the Dec of the star. After 10 sec, turn off tracking and let the star trail across the image while the exposure continues.

  The slope of the trail gives any slight rotation of the CCD with respect to the celestial axes, while the stellar trail points West from the stellar image in the center of the CCD.

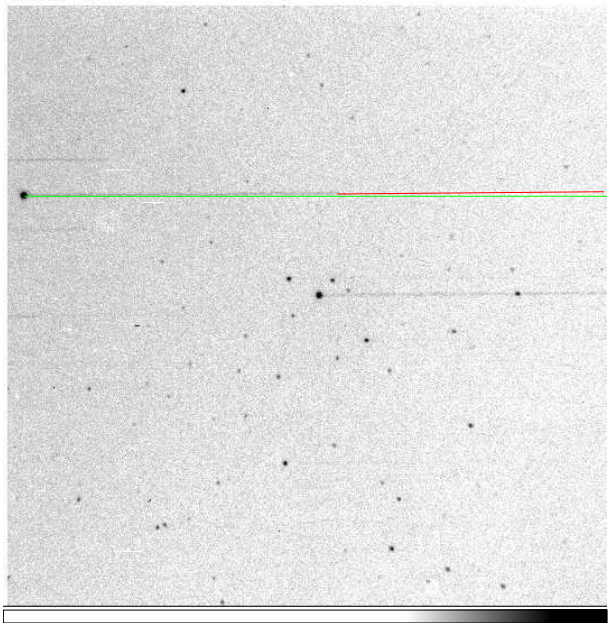

Figure 50: Example of a CCD orientation exposure, obtained at the FLWO 48" telescope. The star trail points W from the stellar image. In this case, the the CCD was rotated $\sim 0\overset{\circ}{.}4$ E of N.

- **Focus exposures** — Exposure or series of exposures at different focus settings. These frames are used to:

  - ◦ determine the best current focus (in the current filter) and adjust the focus setting accordingly

  Focus exposures need to be taken periodically throughout a night. Particularly in the first few hours of a night, the telescope may be rapidly cooling down, which causes the focus to change as well. *Autocollimation* and *autofocussing* routines may try to alleviate the problem, but human tweaking of the focus is often required. For example, at the VATT, in the first two hours after opening up, one has to adjust the focus at least once every half hour, in the next $2\frac{1}{2}$–3 hours, once per $\sim$45 min, and once per hour or so thereafter.

  △ Filters of a particular matched set are often *parfocal*, i.e., changing filters does not change the optimal focus setting. When observing through filters of different prescriptions, one needs to adjust the focus every time when changing filters. Once you know the amount of this adjustment, you do not have to take focus exposures for each change.

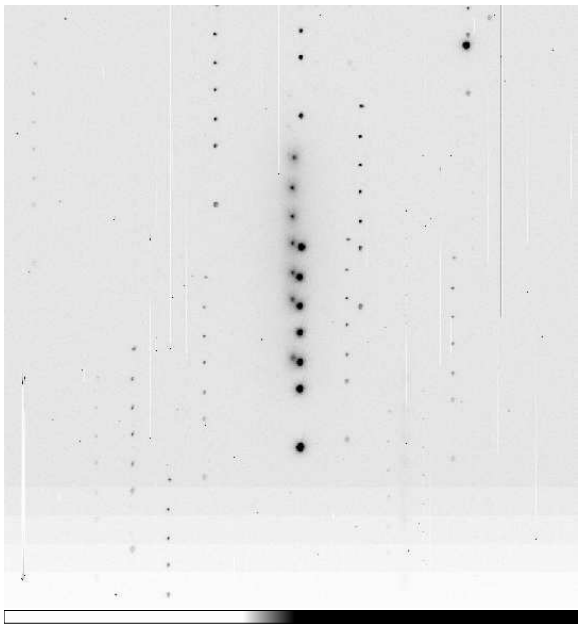

Figure 51: Example of a focus exposure. The exposure consisted of seven integrations, each at a different focus setting. In between each, the collected charge was shifted up 75 pixels along all columns. By measuring the sizes of the PSF's of a few stars at each setting, the correct focus is determined (or interpolated). Note that, in this case, the stars in the field of the science targets (appearing here as a faint and bright smudge) sufficed for taking a focus exposure.

- **(Spectro)Photometric standard star exposures** (*standards*) — Exposures of photometric or spectrophotometric standard stars or standard star fields. These frames are used to:

  - photometrically calibrate the science target exposures
  - establish the linearity of the CCD over its entire dynamic range

  The observed pixel intensities need to be calibrated onto a standard photometric system to be physically meaningfull. To do so and correct for atmospheric extinction as well, standard stars/star fields need to be observed 3 or 4 times during each night that is suspected to be (near-)photometric.

  ▷ With the term *photometric* we mean that *the transparency (per unit airmass) of the sky is* ***independent*** *of the direction in which we point the telescope.*
  (As a corollary: if there is even a single little cloud anywhere above the horizon, per definition conditions are *non-photometric*; if you hear anyone ever claiming that conditions were photometric between the clouds, *run!*).

  Exposures need to be taken in each filter in which science target exposures are taken. Even if the science targets are observed in only one filter, at least one other filter is needed for calibration.

  △ When interupting observations of science targets to observe standard stars, make sure you observe both a field at low airmass (A.M.$\lesssim$1.2) and one at intermediate ($\sim$1.5–1.6) or high airmass (A.M.$\sim$1.8–2.2).

  △ At least once, observe a standard star field in at least one filter using several exposure times that are a factor $\sim$3–5 apart, in order to establish the linearity of the CCD.

  In general, it is a good idea to observe each standard star/star field in each filter using two different exposure times, such that accurate photometry of both the brightest standards and the fainter ones is secured.

8.2 Observing Log of the Imaging Observations

In this section we reproduce the relevant portions of the observing log files of the six nights from which we will process and/or analyse data.

/data1/raj/ast598/obslogs/obslog050404.txt

```
================================================================================
CCD LOG SHEET
--------------------------------------------------------------------------------
Date:      Apr  4, 2005                Observatory: Mt.Graham Observatory, VATT
Observers: R.Jansen,K.Tamura&N.Grogin  Instrument:  ccd26 (2048x2048, binx2)
Program:   H-alpha imaging of void galaxies                   Tel.Focus:-140
Weather:   Mostly clear, but high winds up to 50mph! @ sunset  Format: FITS

                                    Filterwheel  [1]   [2]   [3]   [4]   [5]
# local time  = UT - 7 h            TOP:          -     B     U     V     I
# pixel scale = 0.3746"/pix         BOTTOM:       -     R   Ha668 Ha663   D
# gain = 1.9 e-/ADU, rdnoise = 5.7 e-
# Lat = 32:42:05 deg; Long = 109:53:31 deg W; Alt = 3191 m
================================================================================
Obs.# Object        Filt  Texp  UTC    HA   A.M.   Comments
================================================================================
 001- BIAS 15x      -        0  0:05 +0:00  1.000  Biases
 -015 BIAS          -        0  0:15 +0:00  1.000
# Filled dewar at 18:00; temperature in dome is 2 C; humidity in dome is 44%.
# >> Started up telescope
# Collimation guess: tipx=105, tipy=140, focus_guess=-140
# Turned dome tracking OFF; Pointed telescope at inside dome (AZ=97,Alt=37);
# Set RA drive bias rate to +15.0"/sec to counter tracking...
 016- FLAT 7x       R       12  1:46 -3:39  1.658  Dome flats R (25W spotlight)
 -022 FLAT          R       12  1:51 -3:39  1.660
 ...
# >> Stowed telescope and reset RA bias rate to 0 (-> siderial rate); closed
#    mirror shutters.
 083- BIAS 7x       -        0  3:24 +0:00  1.000  Biases
 -089 BIAS          -        0  3:28 +0:00  1.000
 090- DARK 3x       -      600  3:29 +0:00  1.000  Night darks
 -092 DARK          -      600  3:51 +0:00  1.000
 093- DARK 3x       -      900  4:06 +0:00  1.000  Night dark
 -095 DARK          -      900  5:03 +0:00  1.000
 096- DARK          -     1200  5:21 +0:00  1.000  Night dark
 -098 DARK          -     1200  6:04 +0:00  1.000
# OPENED UP AT 11:40 PM: humidity in dome is 60%; temperature is -5 C.
# Wind gusts are less frequent -- mean windspeed 12--17 mph from S-SW.
# Collimated telescope: tipx=105, tipy=155, focus_guess=-160

 ...

# Probably photometric, but horrendous seeing due to high winds and unstable
# atmosphere. Not much useful science data...
 0175 SKY           R       60 12:29 -0:06  1.018  ~29000 ADU
 0176 SKY           B       30 12:32 -0:03  1.018  ~46250 ADU
 0177 SKY           B       15 12:34 -0:02  1.018  ~31750 ADU
 0178 SKY           B       11 12:35 -0:01  1.018  ~30000 ADU
 0179 SKY           B       10 12:36 -0:00  1.018  ~34750 ADU
 0180 SKY           B        8 12:37 +0:00  1.018  ~36500 ADU
 ...
# >> Stowed telescope and closed down @5:45 AM; temperature -3 C; humidity 11%.
# Photometric at sunrise.
 188- BIAS 23x      -        0 13:05 +0:00  1.000  Biases
 -210 BIAS          -        0 13:20 +0:00  1.000
================================================================================
Probably photometric night, but horendous seeing (mostly >~3") due to high winds
================================================================================
```

/data1/raj/ast598/obslogs/obslog050405.txt

```
================================================================================
CCD LOG SHEET
--------------------------------------------------------------------------------
Date:      Apr  5, 2005                Observatory: Mt.Graham Observatory, VATT
Observers: R.Jansen,N.Grogin&K.Tamura  Instrument:  ccd26 (2048x2048, binx2)
Program:   H-alpha imaging of void galaxies                  Tel.Focus: -140
Weather:   Photometric, but windy @ sunset                   Format: FITS

                                     Filterwheel  [1]   [2]   [3]   [4]   [5]
# local time  = UT - 7 h             TOP:          -     B     U     V     I
# pixel scale = 0.3746"/pix          BOTTOM:       -     R   Ha668 Ha663   D
# gain = 1.9 e-/ADU, rdnoise = 5.7 e-
# Lat = 32:42:05 deg; Long = 109:53:31 deg W; Alt = 3191 m
# exposures marked * were not stored on CDROM.
================================================================================
Obs.# Object        Filt  Texp  UTC    HA   A.M.   Comments
================================================================================
 001- BIAS 15x      D        0  0:08 +0:00  1.000  Biases
 -015 BIAS          D        0  0:18 +0:00  1.000
# Filled dewar at 18:05; temperature in dome 7 C; humidity in dome 36%.
# >> Started up telescope.
# >> Opened up at 18:45PM; Conditions are photometric (but windy) at sunset.
# Collimation guess: tipx=105, tipy=140, focus_guess=-140
# Evening twilight flats...
 0029 SKY           B      2.0  2:03 -0:46  1.016  Twilight flats B; ~24500 ADU
 0030 SKY           B      2.5  2:04 -0:45  1.016  ~24750 ADU
 0031 SKY           B      3.2  2:05 -0:44  1.015  ~25500 ADU
 ...
 0036 SKY           R       12  2:09 -0:39  1.012  Twilight flats R; ~24250 ADU
 0037 SKY           R       15  2:10 -0:38  1.012  ~23250 ADU
 0038 SKY           R       20  2:11 -0:37  1.011  ~23750 ADU
 0039 SKY           R       25  2:12 -0:36  1.011  ~22250 ADU
 ...
 0044 optaxis_test  R       60  2:25 +2:59  1.270  DUST MOTE BACK AGAIN!!!
# Collimated telescope: tipx=105, tipy=130, focus_guess=-140
 0045 focus_run     R      7x5  2:44 +1:02  1.243  best focus near -150
 0046 focus_run     R      7x5  2:50 +1:07  1.253  focus set to -150; fwhm~2.0"
 047- Rubin149      R    10,30  2:57 +1:11  1.254  (using 'standards')
 049- Rubin149      B    20,60  3:00 +1:14  1.259
 0057 Rubin149      R       10  3:13 +1:27  1.287  photom.diff.~0.4%
 058- PG1047+003    R    10,30  3:20 -1:51  1.343  fwhm~1.8"
 060- PG1047+003    B    20,60  3:22 -1:49  1.332
 0068 PG1047+003    R       10  3:37 -1:34  1.294  photom.diff.~1.3%; fwhm~3"

 ...

 122- PG1047+003    R    10,30  8:11 +3:00  1.681
 124- PG1047+003    B    20,60  8:14 +3:03  1.702
 0132 PG1047+003    R       10  8:28 +3:17  1.821  photom.diff. ~0.9%
 0133 focus_run     R     7x10  8:32 -1:50  1.094  focus set to -169; fwhm~1.2"
 0134 A160006+4302  R      180  8:40 -1:42  1.084
 0135 A160006+4302  B      360  8:44 -1:38  1.079  fwhm~1.3"
 0136 A160006+4302  V      240  8:51 -1:31  1.070

 ...

 0175 PG1528+062    R       15 12:11 +2:20  1.349
 0176 PG1528+062    B       40 12:13 +2:22  1.354  jog 5" W
 0180 PG1528+062    R       15 12:20 +2:29  1.385  sky getting rapidly brighter
 0181 PG1633+099    R       15 12:22 +1:26  1.158
 0182 PG1633+099    B       40 12:23 +1:28  1.161
 0186 PG1633+099    R       15 12:31 +1:35  1.175
# Twilight sky flats at "BD+28d4211"...
 0187 SKY           R      2.0 12:33 -3:37  1.450  Twilight flat R; ~13k ADU
 0188 SKY           R      2.0 12:34 -3:36  1.449  ~16k ADU
 ...
# Stowed telescope and closed down @5:45 AM; temperature +6 C; humidity 1.8%.
# Probably photometric, but thin band of cirrus low on SW horizon @ sunrise...
 196- BIAS 15x      D        0 13:08 +0:00  1.000  Biases
 -210 BIAS          D        0 13:23 +0:00  1.000
================================================================================
(Probably) photometric night, reasonable seeing after the first couple of hours
================================================================================
```

/data1/raj/ast598/obslogs/obslog050406.txt

```
================================================================================
CCD LOG SHEET
--------------------------------------------------------------------------------
Date:      Apr  6, 2005                Observatory: Mt.Graham Observatory, VATT
Observers: R.Jansen & K.Tamura         Instrument:  ccd26 (2048x2048, binx2)
Program:   H-alpha imaging of galaxies observed with HST        Tel.Focus:-140
Weather:   Overcast with cirrus clouds @ sunset                 Format: FITS

                                    Filterwheel  [1]   [2]   [3]   [4]   [5]
# local time  = UT - 7 h            TOP:          -     B     U     V     I
# pixel scale = 0.3746"/pix         BOTTOM:       -     R   Ha668 Ha663   D
# gain = 1.9 e-/ADU, rdnoise = 5.7 e-
# Lat = 32:42:05 deg; Long = 109:53:31 deg W; Alt = 3191 m
# exposures marked * were not stored on CDROM.
================================================================================
Obs.# Object        Filt  Texp  UTC    HA   A.M.   Comments
================================================================================
 001- BIAS 15x      D        0  0:32 +0:00  1.000  Biases
 -015 BIAS          D        0  0:42 +0:00  1.000
# Filled dewar at 18:15; temperature in dome 12 C; humidity in dome 28%.
# >> Started up telescope
# Collimation guess tipx=105, tipy=130, focus_guess=-140
# Turned dome tracking OFF; Pointed telescope at inside dome (AZ=97,Alt=37);
# Set RA drive bias rate to +15.0"/sec to counter tracking...
 016- FLAT 5x       R       14  1:52 -3:35  1.638  Dome flats R (25W lamp)
 -020 FLAT          R       14  1:57 -3:35  1.640
 ...
# >> Stowed telescope and reset RA bias rate to 0 (-> siderial rate); closed
#    mirror shutters.
# >> Opened up @ 19:45 PM. Collimated telescope: tipx=105, tipy=130, focus=-140
 0052 focus_run     R     7x10  3:09 +1:21  1.326  focus set to -118; fwhm~1.2"

 ...

 0104 A16006+4302   R      180 12:00 +1:42  1.084  sky getting rapidly brighter
 0105 A16006+4302   B      360 12:04 +1:46  1.089  ~7000 ADU
# Morning twilight sky flats...
 0107 SKY           R      150 12:23 +0:08  1.035  ~40500 ADU
 0108 SKY           R       80 12:26 +0:11  1.035  ~42500 ADU
 0109 SKY           R       50 12:28 +0:13  1.035  ~43500 ADU
 ...
 0114 SKY           B      3.0 12:37 +0:23  1.037  ~44500 ADU
 0115 SKY           B      2.0 12:38 +0:23  1.038  ~30000 ADU
 0116 SKY           B      1.8 12:39 +0:24  1.038  ~35500 ADU
 0117 SKY           B      1.5 12:40 +0:25  1.038  ~30500 ADU
 ...
# >> Stowed telescope and closed down @6:00 AM.
# Filled dewar; temperature in dome is +5 C; humidity is 11.5% at shutdown.
# Bands of thin cirrus are visible low on the SW horizon at sunrise...
 131- BIAS 15x      D        0 13:15 +0:00  1.000  Biases
 -145 BIAS          D        0 13:25 +0:00  1.000
================================================================================
Non-photometric night; fair seeing at start of night, later variable up to ~2"
================================================================================
```

/data1/raj/ast598/obslogs/obslog050407.txt

```
================================================================================
CCD LOG SHEET
--------------------------------------------------------------------------------
Date:      Apr  7, 2005                Observatory: Mt.Graham Observatory, VATT
Observers: R.Jansen & K.Tamura         Instrument:  ccd26 (2048x2048, binx2)
Program:   H-alpha imaging of galaxies observed with HST        Tel.Focus: -120
Weather:   Windfeathers (thin cirrus) all over @ sunset         Format: FITS

                                    Filterwheel  [1]   [2]   [3]   [4]   [5]
# local time  = UT - 7 h            TOP:          -     B     U     V     I
# pixel scale = 0.3746"/pix         BOTTOM:       -     R   Ha668 Ha663   D
# gain = 1.9 e-/ADU, rdnoise = 5.7 e-
# Lat = 32:42:05 deg; Long = 109:53:31 deg W; Alt = 3191 m
# exposures marked * were not stored on CDROM.
================================================================================
Obs.# Object        Filt  Texp  UTC    HA   A.M.   Comments
================================================================================
 001- BIAS 15x      D        0  0:25 +0:00  1.000  Biases
 -015 BIAS          D        0  0:38 +0:00  1.000
# Filled dewar at 18:15; temperature in dome 12 C; humidity in dome 35%.
# >> Started up telescope.
# >> Opened up at 18:40 PM. Conditions are non-photometric (thin cirrus all over
#    at sunset).
# Collimation guess tipx=105, tipy=130; focus_guess=-120
# Evening twilight flats...
 ...
 0030 SKY           B      0.6  1:58 -0:43  1.014  ~36500 ADU
 0031 SKY           B      0.7  1:59 -0:42  1.014  ~34250 ADU
 0032 SKY           B      0.8  2:00 -0:41  1.013  ~32500 ADU
 ...
 0036 SKY           R      2.1  2:03 -0:37  1.011  ~20000 ADU
 0037 SKY           R      2.5  2:04 -0:37  1.011  ~19000 ADU
 0038 SKY           R      3.0  2:05 -0:36  1.010  ~19000 ADU
 0039 SKY           R      4.0  2:06 -0:35  1.010  ~21500 ADU
# Collimated telescope: tipx=100, tipy=140
 0052 focus_run     R     7x10  3:23 +0:42  1.014  focus set to -127; fwhm~1.2"

 ...

# Major system of clouds moving in.... waiting for it to clear up a bit.
# >> Stowed telescope and closed cover and dome at 3:30 AM. Wind has picked up.
 083- BIAS 5x       D        0 10:31 +0:00  1.000  Biases
 -087 BIAS          D        0 10:33 +0:00  1.000
 088- DARK 3x       D     1200 10:40 +0:00  1.000  Night darks
 -090 DARK          D     1200 11:22 +0:00  1.000
 091- BIAS 5x       D        0 11:46 +0:00  1.000  Biases
 -095 BIAS          D        0 11:49 +0:00  1.000
# Turned OFF dome tracking; pointed telescope at clear spot on inside of dome;
# set RA drive bias rate to +15.0"/sec to counter tracking...
 ...
 103- FLAT 7x       R       15 12:11 -3:37  1.658  Dome flats (25W lamp)
 -109 FLAT          R       15 12:16 -3:37  1.659
 ...
# >> Stowed telescope and reset RA drive bias rate to 0 (-> siderial rate),
# >> Shut down telescope at 6:15 AM
# Filled dewar; temperature in dome is +1.5 C, humidity is 32% at shutdown.
# Thick cloud covering at sunrise (since 1:30 AM).
================================================================================
Non-photometric night; fair seeing at start of night, later variable up to ~2"
================================================================================
```

/data1/raj/ast598/obslogs/obslog050408.txt

```
================================================================================
CCD LOG SHEET
--------------------------------------------------------------------------------
Date:      Apr  8, 2005                Observatory: Mt.Graham Observatory, VATT
Observers: R.Jansen & K.Tamura         Instrument:  ccd26 (2048x2048, binx2)
Program:   H-alpha imaging of galaxies observed with HST          Tel.Focus:-120
Weather:   Near-photometric @ sunset (small cloud low on N horiz) Format: FITS

                                       Filterwheel  [1]   [2]   [3]   [4]   [5]
# local time  = UT - 7 h               TOP:          -     B     U     V     I
# pixel scale = 0.3746"/pix            BOTTOM:       -     R   Ha668 Ha663   D
# gain = 1.9 e-/ADU, rdnoise = 5.7 e-
# Lat = 32:42:05 deg; Long = 109:53:31 deg W; Alt = 3191 m
================================================================================
Obs.# Object        Filt  Texp  UTC    HA   A.M.   Comments
================================================================================
# Filled dewar at 18:15; temperature in dome 7 C; humidity in dome 37%.
# Considerably colder than yesterday...
# >> Started up telescope.
# >> Opened up at 18:45 PM. Conditions are (near-)photometric at sunset.
# Collimation guess tipx=105, tipy=130; focus_guess=-120
# Evening twilight flats...
 ...
 0014 SKY           B        5  2:08 -0:29  1.007  Twilight flats B; ~27750 ADU
 0015 SKY           B        7  2:09 -0:28  1.007  ~30000 ADU
 0016 SKY           B        8  2:10 -0:27  1.006  ~26500 ADU
 ...
 0020 SKY           R       30  2:15 -0:22  1.005  Twilight flats R; ~19250 ADU
 0021 SKY           R       40  2:16 -0:21  1.004  ~18000 ADU
 0022 SKY           R       55  2:17 -0:19  1.004  ~16500 ADU
 ...
# Collimated telescope: tipx=120, tipy=145
 0027 focus_run     R     7x10  2:55 -2:05  1.389  focus set to -97; fwhm~1.4"
 0028 PG1047+003    R       15  3:04 -1:56  1.357  forgot to set focus
 0029 PG1047+003    B       25  3:05 -1:55  1.354  ""
 0033 PG1047+003    R       15  3:12 -1:48  1.331  photom.diff. ~1.8%
 034- Rubin149      R    10,30  3:17 +1:43  1.326
 036- Rubin149      B    20,60  3:20 +1:46  1.335
 0044 Rubin149      R       10  3:33 +1:59  1.377  photom.diff. ~0.7%
 0045 PG0918-029    R       15  3:40 +0:08  1.153
 0046 PG0918-029    B       25  3:41 +0:10  1.153
 0050 PG0918-029    R       15  3:47 +0:16  1.155  photom.diff. ~0.8%

 ...

 0105 PG1047+003    R       15  8:21 +3:22  1.866
 0106 PG1047+003    B       30  8:22 +3:23  1.882
 0110 PG1047+003    R       15  8:29 +3:30  1.947  photom.diff. ~1.4%
 0111 PG1323-086    R       15  8:35 +1:01  1.387
 0112 PG1323-086    B       30  8:37 +1:02  1.389
 0116 PG1323-086    R       15  8:43 +1:08  1.402  photom.diff. ~0.4%
 0117 PG1528+062    R       15  8:51 -0:48  1.142
 0118 PG1528+062    B       30  8:52 -0:47  1.141
 0122 PG1528+062    R       15  8:58 -0:41  1.135  photom.diff. ~1.1%

 ...

 0149 A16006+4302   R      180 11:49 +1:38  1.079
 0150 A16006+4302   B      360 11:53 +1:42  1.084
 0152 MarkA         R       15 12:09 -2:42  1.896  sky getting brighter!
 0153 MarkA         B       30 12:10 -2:41  1.886
 0157 PG1633+099    R       20 12:18 +1:34  1.172
 0158 PG1633+099    B       40 12:19 +1:35  1.175  ~13000 ADU
# Stowed telescope; tracking OFF!...
# Morning twilight flats...
 ...
 0169 SKY           B      1.5 12:40 +0:00  1.000  Twilight flats B; ~50250 ADU
 0170 SKY           B      0.8 12:40 +0:00  1.000  ~33250 ADU
 0171 SKY           B     0.65 12:41 +0:00  1.000  ~33000 ADU
 ...
# Closed down @ 5:50 AM; temperature -5 C; humidity 57% ; strongs winds
 181- BIAS          D        0 13:54 +0:00  1.000  Biases
 -200 BIAS          D        0 14:0  +0:00  1.000
================================================================================
Photometric night; but wind afffecting seeing after ~1:00 UT
================================================================================
```

/data1/raj/ast598/obslogs/obslog050929.txt

```
CCD LOG SHEET
-------------------------------------------------------------------------------
Date:      Sep 29, 2005                Observatory: Mt.Graham Observatory, VATT
Observers: R.Jansen, N.Grogin & A.Mott  Instrument:  ccd26 (2048x2048, binx2)
Program:   H-alpha imaging of void galaxies                    Tel.Focus:
Weather:   Clear @ sunset                                      Format: FITS

                                    Filterwheel  [1]   [2]   [3]   [4]   [5]
# local time  = UT - 7 h            TOP:          -     R   Ha668 Ha663   D
# pixel scale = 0.3746"/pix         BOTTOM:       -     U     B     V     I
# gain = 1.9 e-/ADU, rdnoise = 5.7 e-
# Lat = 32:42:05 deg; Long = 109:53:31 deg W; Alt = 3191 m
===============================================================================
Obs.# Object        Filt  Texp  UTC    HA    A.M.  Comments
===============================================================================
# Filled dewar at 17:20PM; humidity in dome 52%; temperature in dome 14.5 C
 001- BIAS          D        0  0:50 +0:00  1.000
 -004 BIAS          D        0  0:52 +0:00  1.000
# CCD acquisition computer 'vattccd' hung -- rebooted -> OK
 005- BIAS          D        0  1:02 +0:00  1.000
 -010 BIAS          D        0  1:   +0:00  1.000
# >> Started up telescope.
# >> Opened up at 17:45PM. Conditions are clear (photometric) at sunset.
# Collimation guess: tipx=115, tipy=120, focus=-100.
# Evening twilight sky flats...
 ...
 0025 SKY           B        8  1:31 +0:07  1.055  Twilight flats B; ~40000 ADU
 0026 SKY           B        8  1:32 +0:08  1.055  ~30500 ADU
 0027 SKY           B       10  1:33 +0:09  1.055  ~29500 ADU
 0028 SKY           B       12  1:34 +0:10  1.056  ~27250 ADU
 0029 SKY           B       16  1:35 +0:11  1.056  ~27750 ADU
 ...
 0034 SKY           R       45  1:41 +0:17  1.057  Twilight flats R; ~18000 ADU
 0035 SKY           R       60  1:43 +0:19  1.058  ~17250 ADU
 0036 SKY           R       90  1:45 +0:21  1.059  ~17500 ADU
 ...
# Collimated telescope: tipx=120 , tipy=120
 0041 focus_run     R      7x5  2:35 +2:52  1.465  focus set to -121; fwhm=1.1"
 0042 A16006+4302   R      450  2:56 +4:10  1.538  SN2005bk follow-up; fwhm~1.3"
 0043 A16006+4302   R      450  3:07 +4:21  1.604  SN2005bk follow-up
 0044 A16006+4302   B      600  3:17 +4:31  1.670  SN2005bk follow-up; fwhm~1.4"
 ...
 0050 PG2213-006    R       15  4:24 -0:35  1.204
 0052 PG2213-006    B       40  4:28 -0:30  1.201
 0055 PG2213-006    R       15  4:32 -0:26  1.198  photom.diff. in R < 0.65%
 0056 MarkA         R       15  4:37 +1:10  1.450    = 100%*abs(x-y)/(0.5*(x+y))
 0058 MarkA         B       30  4:40 +1:13  1.460
 0061 MarkA         R       15  4:44 +1:17  1.470  photom.diff. in R < 0.26%

 ...

 0081 PG0231+051    R       15  6:49 -2:27  1.385
 0083 PG0231+051    B       30  6:52 -2:23  1.372
 0086 PG0231+051    R       15  6:57 -2:19  1.355
 0087 PG2336+004    R       15  7:03 +0:41  1.196
 0089 PG2336+004    B       40  7:07 +0:45  1.200  trailed -- secondary jumped!
 0092 PG2336+004    R       15  7:12 +0:50  1.205  photom.diff. < 1.1%
 0093 PG2336+004    B       40  7:15 +0:53  1.209  (redo of 0089) OK

 ...

 0111 PG2336+004    R       15  9:42 +3:21  1.833  fwhm~1.2"
 0113 PG2336+004    B       35  9:46 +3:24  1.865
 0116 PG2336+004    R       15  9:50 +3:29  1.914  photom.diff. < 0.77%
 0117 PG0231+051    R       15  9:54 +0:38  1.139
 0119 PG0231+051    B       35  9:58 +0:41  1.142
 0122 PG0231+051    R       15 10:02 +0:46  1.146  photom.diff. < 0.19%
```

/data1/raj/ast598/obslogs/obslog050929.txt

```
 ...

 0142 Rubin149      R       10 12:40 -1:26  1.284
 0144 Rubin149      B       30 12:42 -1:    1.2    OVEREXPOSED!!
 0147 Rubin149      R       10 12:47 -1:19  1.269  too bright ~60000 ADU
# Morning twilight sky flats at "Rubin149"...
 0148 SKY           R      3.0 12:48 -1:17  1.266  Twilight flats R; ~23500 ADU
 0149 SKY           R      2.6 12:49 -1:16  1.264  ~25500 ADU
 0150 SKY           R      2.2 12:50 -1:16  1.262  ~28500 ADU
 0151 SKY           B      4.0 12:51 -1:15  1.261  Twilight flats B; OVEREXPOSED
 0152 SKY           B      2.0 12:52 -1:14  1.259  ~41500 ADU
 0153 SKY           B      1.2 12:53 -1:13  1.257  ~30000 ADU
 0154 SKY           B      0.9 12:53 -1:12  1.256  ~28500 ADU
 ...
# Stowed telescope @ 6:05 AM; closed mirror cover, then dome slit.
# Conditions are photometric at sunrise.  Most likely photometric all night!
# Dome tracking is turned OFF; Pointed telescope at clear spot on inside of
# dome. Set the RA drive bias rate to +15.0 "/sec to counter tracking...
 168- FLAT          R      2.0 13:23 -3:30  1.693  Dome flats (40W lamp)
 -172 FLAT          R      2.0 13:26 -3:30  1.694
 173- FLAT          R      0.4 13:43 -3:30  1.696  Dome flats (2x60W+25W lamps)
 -177 FLAT          R      0.4 13:
 178- FLAT          R      0.8 13:47 -3:30  1.696  Dome flats (2x60W+25W lamps)
 -182 FLAT          R      0.8 13:
 ...
 243- FLAT          R      0.1 15:03 -3:31  1.702  Dome flats (2x60W+300W)
 -247 FLAT          R      0.1 15:
 248- FLAT          R      1.2 15:12 -3:31  1.702  Dome flats (60W+25W lamps)
 -252 FLAT          R      1.2 15:
 253- FLAT          R      5.0 15:18 -3:31  1.703  Dome flats (25W lamp)
 -260 FLAT          R      5.0 15:
 261- FLAT          B       30 15:39 -3:31  1.705  Dome flats (300W + daylight
 -270 FLAT          B       30 15:                  from open door to outside).
# >> Stowed telescope and reset RA bias rate to 0 (-> siderial rate).
# >> Shutdown telescope @8:55AM.
# Filled dewar at 9:05AM; humidity in dome 47%; temperature in dome 11 C.
 271- BIAS          D       0  16:17 +0:00  1.000  Biases
 -290 BIAS          D       0  16:
================================================================================
Most likely photometric all night!
================================================================================
```

8.3 Starting IRAF and DS9, and retrieving the *raw* data

Start an `xgterm` terminal window (a modified `xterm`, with extra graphical capabilities built in specifically for use with IRAF). In the following, we will denote the command prompts of `xterm`, `xgterm` and `IRAF cl` as `$>`, `%>` and `cl>`:

```
$>  xgterm
```

In the following, a standard IRAF (version 2.12.2) installation is assumed in /iraf/, with compatible versions of the `stsdas`, `tables` and `ctio` packages already installed in /iraf/extern/. IRAF can be obtained from *http://iraf.noao.edu/* or *http://iraf.net/downloads/*.

If this is the first time you ever ran IRAF, then you first have to create a subdirectory in your home directory and name it "iraf". Then, in that subdirectory, run `mkiraf`, which will prepare a user login macro named "login.cl" and a user parameters directory named "uparm":

```
%>  mkdir iraf ; cd iraf
%>  mkiraf          (where prompted for the terminal type, enter xgterm)
```

In the "login.cl" file just created, change `imdir` to an empty string and add a `username` entry:

```
set     imdir           = ""
set     username        = "A.Student"       ⟵ insert below the userid entry
```

This needs to be done only once (unless IRAF gets upgraded to a newer version).

If not installed on your machine already (check in /iraf/extern/), download and install IRAF external layered package `rjtools`, available as a gzip-compressed tar archive from *http://www.public.asu.edu/∼rjansen/*. In your "iraf" home directory:

```
%>  (g)tar -xvpzf rjtools_2.12_16Jun2005.tgz
```

The tar file can be deleted once `rjtools` is installed.

Create a file named "loginuser.cl" (or edit the existing one) and insert the following lines (before the `keep` statement in the case of an existing file):

```
reset rjtools         = home$rjtools/
task rjtools.pkg      = rjtools$rjtools.cl
printf ("reset helpdb=%s,rjtools$lib/helpdb.mip\nkeep\n",envget("helpdb")) | cl
flpr

keep
```

A final note: in order to find the tasks within the `rjtools` package using task `apropos` within the `stsdas` package, its database needs to be updated. Since `apropos` is a simple ASCII text file, all that needs to be done is to append file "rjtools$lib/apropos.db" to "stsdas$lib/apropos.db". Unfortunately, this may require administrator privileges, so you may have to ask the administrator to type in your iraf home directory:

```
%>  cat rjtools/lib/apropos.db >> /iraf/extern/stsdas_version/lib/apropos.db
```

In the following, we'll also use a unix/linux utility called `narg`, which is a C-shell script around executable `arg`. These can also be downloaded from *http://www.public.asu.edu/∼rjansen/* as a gzip-compressed tar file, and either installed in /usr/local/src/ (as 'root') or in your home directory (as 'user'). You may need a Fortran77 compiler to recompile the source code.

Start up the IRAF command language interface in that `xgterm` terminal window:

```
%>  cd ~/iraf
%>  cl
```

Enter the working directory containing the data:

```
cl>  cd /data1/raj/ast598/
```

Load any required packages that are not loaded by default upon startup (loaded by default are: `noao`, `language`, `system`, `lists`, `dataio`, `images`, `imutil`, `tv`, `plot`, `utilities`):

```
cl>  stsdas
cl>  imred
cl>  bias
cl>  rjtools
```

(Note that in reality, the `cl>` prompt changes after loading each subsequent package). Start up an image display program that knows how to communicate with IRAF, e.g., `ds9`, and do so from within IRAF (i.e., *not* from a unix command prompt!):

```
cl>  !ds9 &
```

Since `ds9` is a unix program and not an IRAF task, we need to precede its name with an exclamation mark (`!`) to temporarily escape to the unix shell. The ampersand (`&`) causes `ds9` to run in the background and gives us back the `cl>` prompt.

Find an IRAF display buffer that is sufficiently large to include the entire 1044×1044 pixel CCD image (need not be an exact match to the actual image size — just find the next larger one):

```
cl>  gdevices
  imt1  imt512     512  512
  imt2  imt800     800  800
  imt3  imt1024   1024 1024   ⟵  cuts off the overscan strips
    ...
 imt44  imt1100   1100 1100   ⟵  choose this one
    ...
```

Set the default image display buffer (twice for good measure):

```
cl>  set stdimage=imt44
cl>  set stdimage=imt44
```

Assume that we have read from tape or copied from CDROM/DVD all "raw" FITS data we wish to process as it came from the telescope. To reduce their volume, all these raw, 16-bit integer data were stored in `gzip`-compressed format. They currently reside in a subdirectory named "raw/".

Copy the contents of the "raw/" subdirectory to the present working directory, and uncompress all data:

```
cl>  dir raw
050404    050405    050406    050407    050408    050929
cl>  !du -sh raw
388M    raw
cl>  !\cp -pr raw/* .        ⟵ don't forget the period as I did in class!
cl>  !gunzip 050???/*.fits.gz
cl>  !du -sh
1.2G    .
```

During data processing, one often makes a mistake that renders the partially processed data useless. For that reason we *always* want to keep an original copy of the raw data around. Since we will copy and also convert the data from their original 16-bit encoding to 32-bit floating point encoding, and since we may keep some partially processed data and calibration images around, the total free diskspace required for the present image processing amounts to several ($\sim$3) *Gigabytes*!

8.4 Preparations before CCD image processing

• Create a list of all data in the subdirectories per night:

```
cl>  !\ls -1 050???/*.fits > all.lis
```

Print an inventory of the unique "object" names (FITS header keyword `OBJECT`) as they occur in the image headers:

```
cl>  hselect @all.lis object yes | sort | unique
A16006+4302  <--  our science target
BIAS
DARK
FLAT         <--  dome flats
MarkA        <--  Landolt photometric standard field
PG0231+051        ""
PG0918-029        ""
PG1047+003        ""
PG1323-086        ""
PG1528+062        ""
PG1633+099        ""
PG2213-006        ""
PG2336+004        ""
Rubin149          ""
SKY          <--  twilight sky flats
focus_run
optaxis_test <--  orientation exposure
```

Now we know what they are called, create separate lists for biases, darks, dome flats, twilight flats, and science target frames:

```
cl>  hselect @all.lis $I 'object=="BIAS"' > bias.lis
cl>  hselect @all.lis $I 'object=="DARK"' > dark.lis
cl>  hselect @all.lis $I 'object=="FLAT"' > flat.lis
cl>  hselect @all.lis $I 'object=="SKY"' > sky.lis
cl>  hselect @all.lis $I 'object!="BIAS"&object!="DARK"&object!="FLAT"'>tmp.lis
cl>  hselect @tmp.lis $I 'object!="SKY"&object!="focus"' > sci.lis
```

Note, that the science exposures were defined by what they are *not*. Since that wouldn't all fit on my command line, I used a temporary list ("tmp.lis") from which I selected further. In the following, any files with names starting with "tmp" or "_" are temporary files, to be deleted when no longer needed, whose names may be re-used.

```
cl>  delete tmp.lis yes verify-
```

• Construct a bad pixel/bad column list using the median of a few twilight sky flat exposures (a convenient task to do so is `getregion`, which is part of the `rjtools` package). Make sure `ds9` is already open:

```
cl>  head sky.lis nlines=7 > tmp.lis
cl>  unlearn imcombine
cl>  imcombine ("@tmp.lis", "tmp.fits", combine="average", reject="avsigclip",
>>>       scale="mean", zero="none", weight="mean", statsec="[16:1015,1:1000]")
cl>  unlearn getregion
cl>  getregion ("tmp.fits", "tmp.reg", format="basic", chkbound+, append+,
>>>       verbose+)
# GETREGION: NOAO/IRAF V2.12.2-EXPORT raj@andromeda Oct 4 09:50:10 2006
    image = tmp.fits [1044 x 1044]
    output= tmp.reg (format="basic")
  Displaying image "tmp.fits":  z1=11624.57 z2=53420.01
    Wait for the cursor cross to appear in the image display area,
    then mark all regions by hitting "b" once at each of two dia-
    gonally oposite corners to mark an arbitrary rectangular region,
    or "c" once to mark a single pixel.  Hit "q" to quit.

    ...

# GETREGION Finished.
```

Overlay the defined regions onto the image to verify using task `tvmarkall` (which can also be found in the `rjtools` package):

```
cl>  unlearn tvmarkall
cl>  tvmarkall ("tmp.fits", gal-, star-, cosmic-, del-, region+,
>>>       regfil="tmp.reg", regclr=204)
```

If we're happy with it (no missed or incorrectly defined bad pixels/columns), save the file under a more descriptive name with a descriptive header line prepended:

```
cl>  printf("# Loral 2048 x 2048 back-illum. 15mu-pixel ", > "badpix.ccd26")
cl>  printf("thinned CCD, binned 2x2, untrimmed\n", >> "badpix.ccd26")
cl>  !cat tmp.reg >> badpix.ccd26
cl>  delete tmp.lis,tmp.reg yes verify-
```

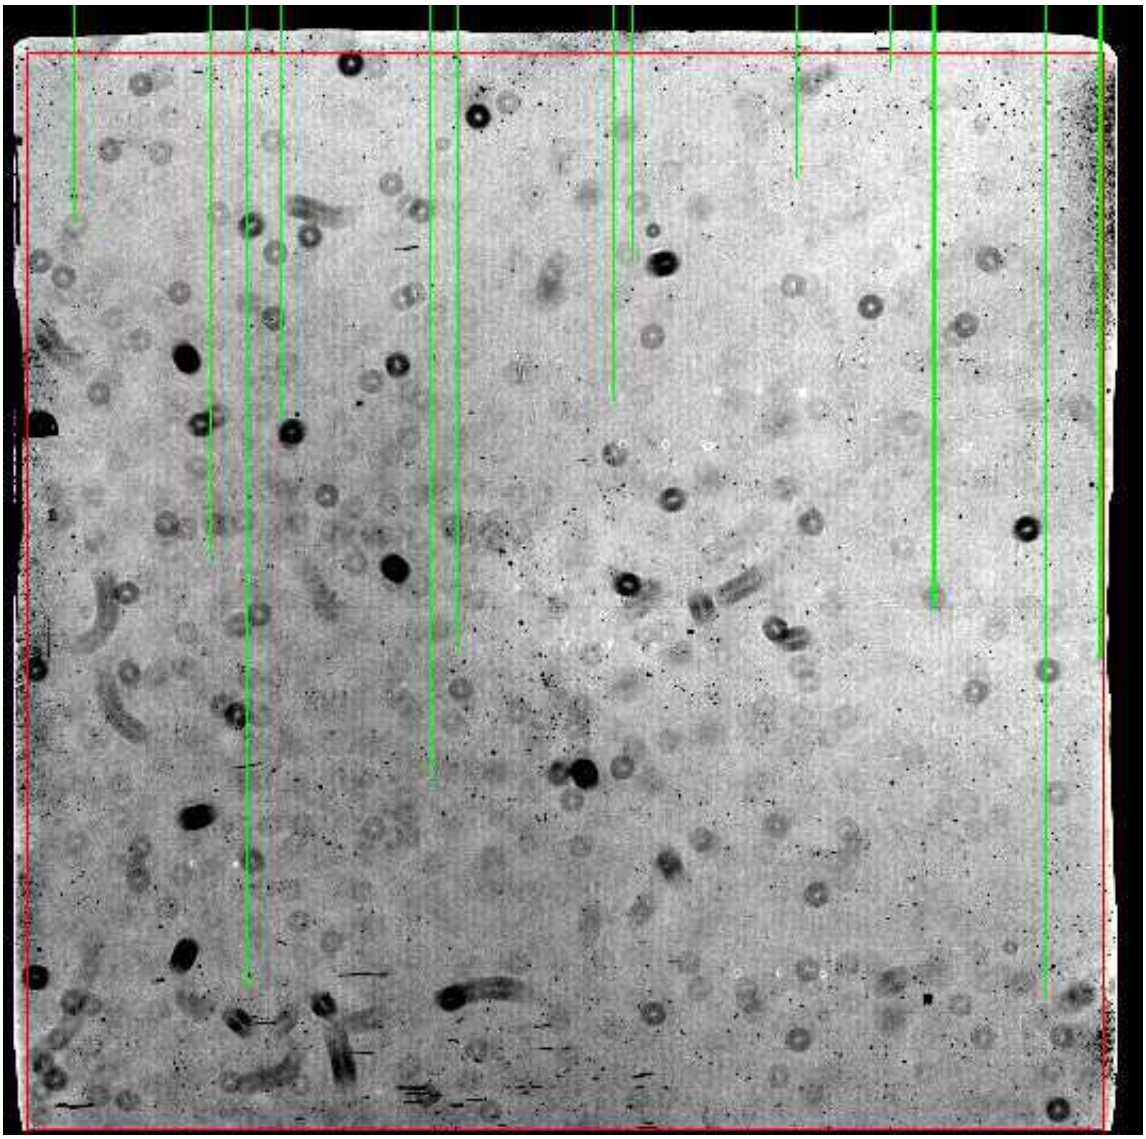

Figure 52: Bad pixel columns (*green*), interactively marked on the image display using task `getregion`. Also indicated (*red*) is the 1000×1000 pixel section that we will excise ("trim") in later processing steps, to avoid problems due to poor response near the left, right and upper edge of the illuminated portion of this CCD.

• Determine the location of the virtual overscan strip, which will be used to track the (somewhat varying) bias level. Note that the noise within the virtual overscan strip should be *only* read-noise, regardless of the signal level in the illuminated region (image region) of the CCD, while the physical overscan strip may suffer from deferred charge and leaked light, that will result in additional signal and noise:

```
cl>  implot tmp.fits
  :c 1 1044     ⟶   [16:1015,1032:1044]
  :l 1 1044     ⟶   [1032:1044,1:1000]
```

So, this particular CCD has two possible overscan strips (see Fig 53). Let's see which one is the virtual overscan strip.

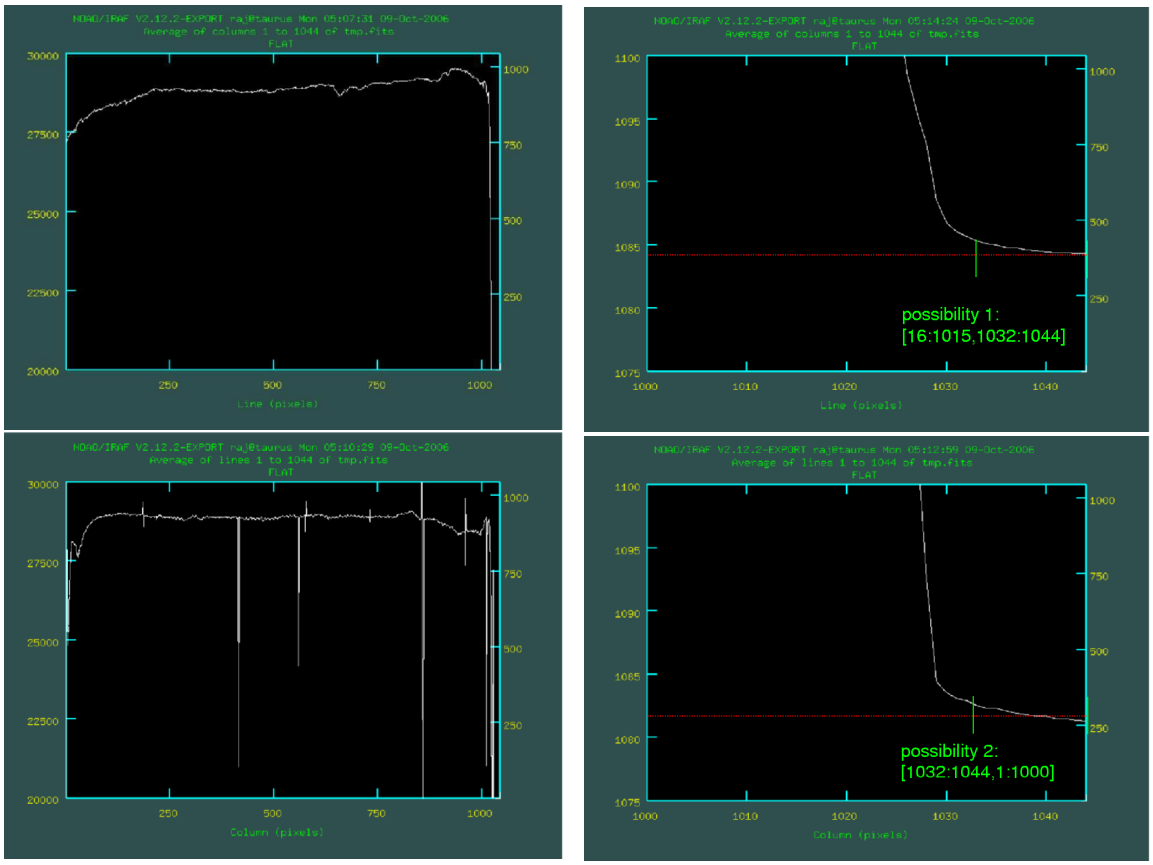

Figure 53: Screen shots of `implot` graphs of the average of columns 1–1044 (*top*) and of lines 1–1044 (*bottom*) in a temporary image created by averaging 7 twilight sky flats. Two possible overscan strips are identified, only one of which is the *virtual* overscan strip.

Compute the pixel statistics in both strips identified in Fig. 53, both in biases and in twilight sky flats:

```
cl>  !head -20 bias.lis | sed 's/$/[16:1015,1032:1044]/g' > tmp.lis
cl>  imstat @tmp.lis fields="image,mean,stddev,min,max"
050404/e0001.fits[16:1015,1032:1044]   1060.   2.952   1047.   1071.
050404/e0002.fits[16:1015,1032:1044]   1060.   2.962   1046.   1071.
050404/e0003.fits[16:1015,1032:1044]   1060.   2.921   1047.   1071.
050404/e0004.fits[16:1015,1032:1044]   1060.   2.934   1048.   1073.
 ...
cl>  !head -20 sky.lis | sed 's/$/[16:1015,1032:1044]/g' > tmp.lis
cl>  imstat @tmp.lis fields="image,mean,stddev,min,max"
050404/e0177.fits[16:1015,1032:1044]   1087.   485.9   1058.  16497.
050404/e0178.fits[16:1015,1032:1044]   1072.   15.67   1060.   1612.
050404/e0179.fits[16:1015,1032:1044]   1087.   488.1   1059.  16528.
050405/f0029.fits[16:1015,1032:1044]   1066.   14.93   1051.   1574.
 ...
```

From which we find that the mean ∼1060–1120 ADU and the rms ("stddev") ∼14–500 ADU. This is not consistent with pure read-noise which, from the statistics computed in the bias frames, we know should be ∼2.95 ADU. Hence, this is *not* the virtual overscan strip.

Let's try the other option:

```
cl>  !head -20 bias.lis | sed 's/$/[1032:1044,1:1000]/g' > tmp.lis
cl>  imstat @tmp.lis fields="image,mean,stddev,min,max"
050404/e0001.fits[1032:1044,1:1000]   1059.   2.850   1050.   1071.
050404/e0002.fits[1032:1044,1:1000]   1060.   2.829   1049.   1070.
050404/e0003.fits[1032:1044,1:1000]   1060.   2.820   1050.   1071.
050404/e0004.fits[1032:1044,1:1000]   1060.   2.859   1048.   1070.
 ...
cl>  !head -20 sky.lis | sed 's/$/[1032:1044,1:1000]/g' > tmp.lis
cl>  imstat @tmp.lis fields="image,mean,stddev,min,max"
050404/e0176.fits[1032:1044,1:1000]   1078.   3.006   1066.   1090.
050404/e0177.fits[1032:1044,1:1000]   1076.   2.934   1066.   1088.
050404/e0178.fits[1032:1044,1:1000]   1076.   2.912   1065.   1087.
050404/e0179.fits[1032:1044,1:1000]   1077.   2.928   1065.   1088.
 ...
```

Here we find that the mean ∼1060–1080 ADU and the rms ∼2.9–3.1 ADU. This *does* look like the virtual overscan strip. Note also, that the minimum and maximum values reported are all consistent with the mean $\pm 4\,\sigma$.

So, in the following, we will adopt (see also Fig. 54):

```
TRIMSEC = STATSEC = '[16:1015,1:1000]', BIASSEC = '[1032:1044,1:1000]'
```

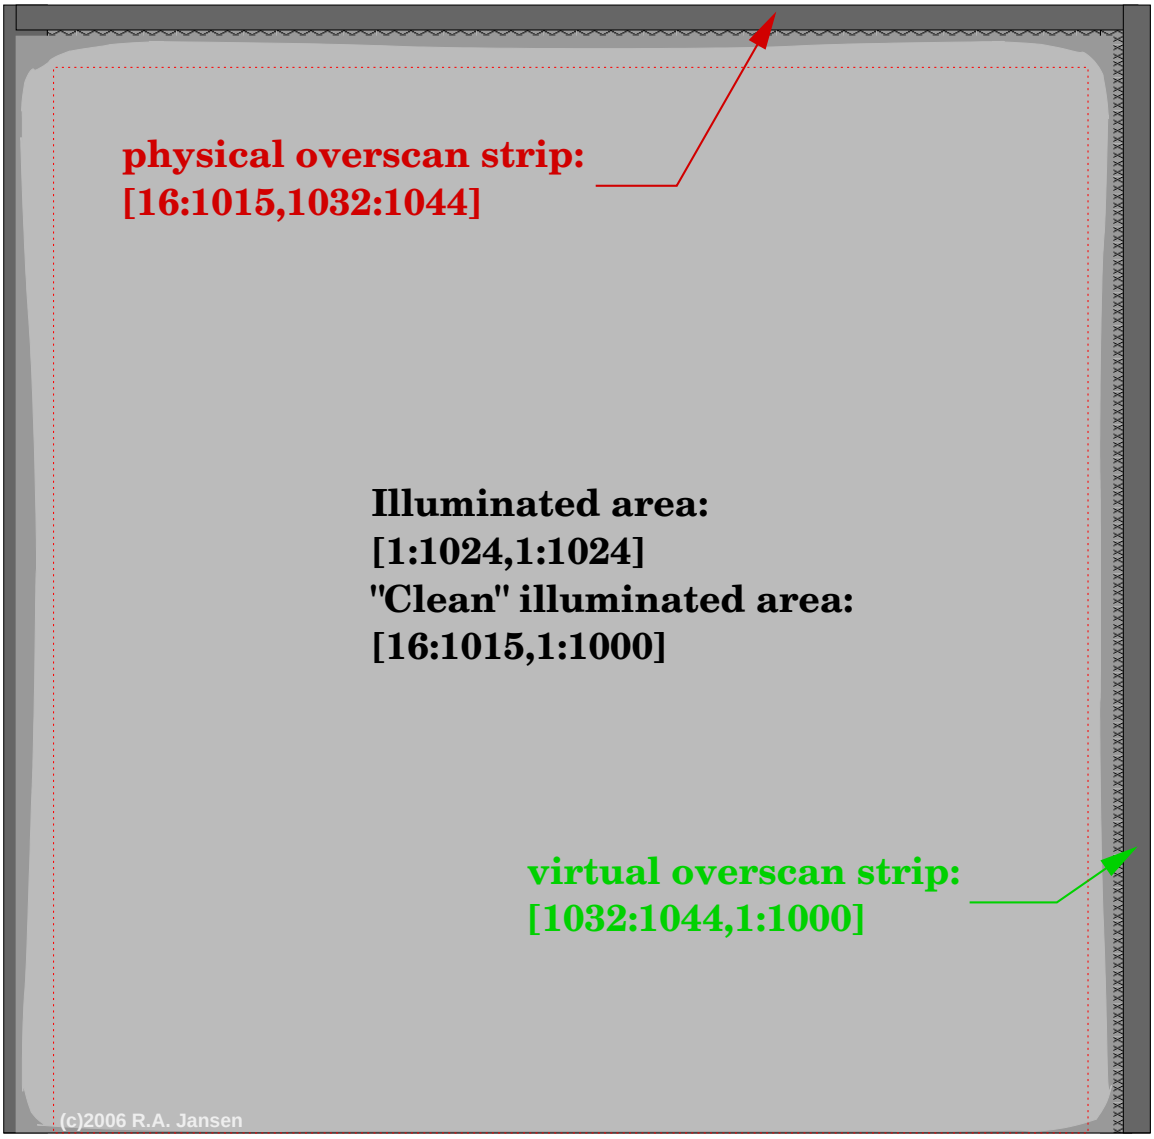

Figure 54: Layout of the CCD area in pixel coordinates, with overscan strips indicated. We will trim the outer edges of the illuminated area because their response is poor (see also Fig. 52).

• Next we need to measure the CCD gain and read-noise using Janesick's method using pairs of flats and biases. A convenient task to do so is `findgain`, which is found within the `noao.obsutil` package). First load that package (`noao` is already loaded):

```
cl>  obsutil
```

For usage information, view the parameter file of this task, or its help file:

```
cl>  lpar findgain   and/or:    help findgain
```

Prepare some matched lists of biases and (dome) flats — we almost always have fewer flats than biases:

```
cl>  !grep '050404' flat.lis | wc -l
  7
cl>  !grep '50404' flat.lis | head -6 | sed 's/.fits//g' > _f1
cl>  !grep '50404' flat.lis | tail -6 | sed 's/.fits//g' > _f2
cl>  !grep '50404' bias.lis | head -6 | sed 's/.fits//g' > _b1
cl>  !grep '50404' bias.lis | head -7 | tail -6 | sed 's/.fits//g' > _b2
```

To check these lists, type:

```
cl>  !paste _f1 _f2 _b1 _b2
050404/e0016    050404/e0017    050404/e0001    050404/e0002
050404/e0017    050404/e0018    050404/e0002    050404/e0003
050404/e0018    050404/e0019    050404/e0003    050404/e0004
050404/e0019    050404/e0020    050404/e0004    050404/e0005
050404/e0020    050404/e0021    050404/e0005    050404/e0006
050404/e0021    050404/e0022    050404/e0006    050404/e0007
```

Since we want to avoid a lot of typing, we create a little template script, where the variable entries are symbolically denoted `$1`, `$2`, etc... We then use a unix utility called `narg` to replicate this template script, while substituting the variables with entries taken from the above matched lists, i.e., `$1` from "_f1", `$2` from "_f2", etc.:

```
cl>  !echo 'findgain ("$1", "$2", "$3", "$4",' > getccdpar.tem
cl>  !echo '      section="[16:1015,1:1000]", verbose-)' >> getccdpar.tem
cl>  type getccdpar.tem
findgain ("$1", "$2", "$3", "$4",
      section="[16:1015,1:1000]", verbose-)

cl>  !narg getccdpar.tem _f1 '$1' _f2 '$2' _b1 '$3' _b2 '$4' > getccdpar.cl
```

Verify that the replication and entry substitution into the template script was successful:

```
cl>  type getccdpar.cl
findgain ("050404/e0016", "050404/e0017", "050404/e0001", "050404/e0002",
       section="[16:1015,1:1000]", verbose-)
findgain ("050404/e0017", "050404/e0018", "050404/e0002", "050404/e0003",
       section="[16:1015,1:1000]", verbose-)
findgain ("050404/e0018", "050404/e0019", "050404/e0003", "050404/e0004",
       section="[16:1015,1:1000]", verbose-)
findgain ("050404/e0019", "050404/e0020", "050404/e0004", "050404/e0005",
       section="[16:1015,1:1000]", verbose-)
findgain ("050404/e0020", "050404/e0021", "050404/e0005", "050404/e0006",
       section="[16:1015,1:1000]", verbose-)
findgain ("050404/e0021", "050404/e0022", "050404/e0006", "050404/e0007",
       section="[16:1015,1:1000]", verbose-)
```

Execute this IRAF CL script and compute the gain, $\mathcal{G}$, and read-noise, $\mathcal{R}$; the syntax to do so may look odd (we call the IRAF CL itself from within the currently running CL, but reading its command input from script "getccdpar.cl"):

```
cl>  cl < getccdpar.cl
 1.95    5.60
 1.95    5.60
 1.95    5.60
 1.95    5.60
  ...    ...
```

This is close to, but slightly different from, the advertised values that were recorded in the observing log files (§ 8.2): $\mathcal{G} = 1.9\,e^-/\mathrm{ADU}$ and $\mathcal{R} = 5.7\,e^-$.

Let's repeat this for the nights of Apr 7 2005 and Sep 29 2005 to see if the differences found are significant:

```
cl>  !grep '050407' flat.lis | wc -l
  7
cl>  !grep '50407' flat.lis | head -6 | sed 's/.fits//g' > _f1
cl>  !grep '50407' flat.lis | tail -6 | sed 's/.fits//g' > _f2
cl>  !grep '50407' bias.lis | head -6 | sed 's/.fits//g' > _b1
cl>  !grep '50407' bias.lis | head -7 | tail -6 | sed 's/.fits//g' > _b2
cl>  !narg getccdpar.tem _f1 '$1' _f2 '$2' _b1 '$3' _b2 '$4' > getccdpar.cl
cl>  cl < getccdpar.cl
 1.96    5.62
 1.96    5.62
 1.96    5.63
 1.96    5.63
  ...    ...
```

(The `range` command below is used to specify only the exposure numbers of the five 2 sec and eight 5 sec R-filter flats of Sep 29 2005):

```
cl>  !range -pre "050929/b" -d 4 168 - 172 253 - 259 > _f1
cl>  !range -pre "050929/b" -d 4 169 - 172 253 - 260 > _f2
cl>  !grep '50929' bias.lis | head -12 | sed 's/.fits//g' > _b1
cl>  !grep '50929' bias.lis | head -13 | tail -12 | sed 's/.fits//g' > _b2
cl>  !narg getccdpar.tem _f1 '$1' _f2 '$2' _b1 '$3' _b2 '$4' > getccdpar.cl

cl>  cl < getccdpar.cl
 1.96    5.72
 1.95    5.69
 1.95    5.69
 1.95    5.80
  ...     ...
```

In the mean, we find $\overline{\mathcal{G}} = 1.953 \pm 0.005$ e^-/ADU and $\overline{\mathcal{R}} = 5.66 \pm 0.07$ e^-. The read-noise does, therefore, not differ significantly from the one advertised, but the gain is slightly ($\sim$3%) higher.

Clean up the temporary lists we created:

```
cl>  delete _f1,_f2,_b1,_b2 yes verify-
```

Since we won't need the `obsutil` package further, unload it:

```
cl>  bye
```

(Note, that `bye` unloads the last package loaded. Packages can be unloaded up to the `cl` level, but one can unload (i.e., exit) the IRAF CL itself only by typing `logout`).

8.5 Processing of the calibration frames

8.5.1 Bias frames

Our goal is to construct a "master" bias reference frame to check for structure in the bias level across the CCD and for offsets with respect to the level in the virtual overscan strip. That frame should have a high enough S/N, that it will add only negligible noise when subtracted from any of the other types of images. We will:

(0) visually inspect all bias frames and convert to 32-bit floating-point format
(1) interpolate over the bad pixels/columns we defined earlier
(2) subtract the level measured in the overscan strip
(3) average the individual biases together into a single, high S/N frame
(4) verify that the level in the overscan region of the combined frame is exactly 0; if not, subtract the residual level.

• Visually inspect all bias frames (remove any biases with problems):

```
cl>  imexamine @bias.lis 1 allframes- nframes=1
```

(the `n`-key will proceed to the next image in the list; the `p`-key will go to the previous one. Hit `q` to quit.

Convert all biases from 16-bit integer ("ushort") to 32-bit floating-point ("real") format (remember that an A/D converter only returns integer ADUs). The smallest number of bits that can store the floating-point equivalent of a 16-bit integer is 32:

```
cl>  chpixtype @bias.lis @bias.lis "real" oldpixtype="all" verbose+
```

Add a line to the image headers indicating the start of the processing log entries, and add an entry logging the pixel data type conversion. Since we'll do this for every other type of data as well, it pays to put this in a little script. In the following, whenever the argument of a `type` command is a template script ("*.tem"), that template script should be created — with `!echo` commands or with your favorite ASCII text editor. The displayed output of `type` just shows what the template script should look like.

```
cl>  type hdrlog_init.tem
gdate() ; sysinfo() ; unlearn hedit ; hedit.add=yes ; hedit.verify=no
s1="CCD image processing -- "//sysinfo.username//"@"//sysinfo.host
s1="-------- "//s1//" --------"
hedit ("@%.lis", "HISTORY", s1)
hedit ("@%.lis", "CHPXTYPE", "COMPLETE")
hedit ("@%.lis", "CHPXDATE", gdate.fdate)

cl>  !sed 's:%:bias:g' hdrlog_init.tem > hdrlog_init.cl
cl>  cl < hdrlog_init.cl
```

This added the following log entry to the FITS headers of all bias frames:

```
HISTORY = '------- CCD image processing -- R.A. Jansen@andromeda -------'
CHPXTYPE= 'COMPLETE'
CHPXDATE= '2006-10-11T17:09:22'
```

• Interpolate over the bad pixels/columns identified earlier, and log relevant information to the image headers:

```
cl>  fixpix("@bias.lis","badpix.ccd26",verbose+, >> "ccdproc.log")

cl>  type hdrlog_fxpx.tem
gdate()
hedit ("@%.lis", "FIXPIX", "COMPLETE")
hedit ("@%.lis", "FXPXFILE", "/data1/raj/ast598/badpix.ccd26")
hedit ("@%.lis", "FXPXDATE", gdate.fdate)

cl>  !sed 's/%/bias/g' hdrlog_fxpx.tem > hdrlog_fxpx.cl
cl>  cl < hdrlog_fxpx.cl
```

This added the following log entry to the FITS headers of all bias frames:

```
FIXPIX = 'COMPLETE'
FXPXFILE= '/data1/raj/ast598/badpix.ccd26'
FXPXDATE= '2006-10-11T17:10:41'
```

• Subtract the bias level as measured in the virtual overscan strip. First, we'll measure that level and record it in a log file:

```
cl>  !sed 's/$/[1032:1044,1:1000]/g' bias.lis > _bias.lis
cl>  imstat ("@_bias.lis", fields="image,midpt", lower=INDEF, upper=INDEF,
>>>       nclip=3, lsigma=3., usigma=3., format-, cache-, > "overscan.dat")
cl>  delete _bias.lis yes verify-
```

Now fit and remove the overscan level from all biases using task `colbias`:

```
cl>  unlearn colbias
cl>  colbias ("@bias.lis", "@bias.lis", bias="[1032:1044,1:1000]", trim="",
>>>       median+, interactive-, function="legendre", order=1, low_reject=3.,
>>>       high_reject=3., niterate=3, logfiles="ccdproc.log")
```

And update the headers (first create the following two template scripts):

```
cl>  type hdrlog_ovsc.tem
gdate()
hedit ("@%.lis", "BIASSEC,TRIMSEC,DATASEC,CCDSEC,ORIGSEC",add-,del+)
hedit ("@%.lis", "OVERSCAN", "COMPLETE")
hedit ("@%.lis", "ORIGSEC", "[1:1044,1:1044]")      <-- original size of full frame
hedit ("@%.lis", "BIASSEC", "[1032:1044,1:1000]")   <-- overscan section
hedit ("@%.lis", "TRIMSEC", "[16:1015,1:1000]")     <-- excised from original frame
hedit ("@%.lis", "IMAGSEC", "[1:1000,1:1000]")      <-- output image section
!narg addovscmean.tem %.lis '$1' > addovscmean.cl
cl < addovscmean.cl
hedit ("@%.lis", "OVSCDATE", gdate.fdate)
```

```
cl>  type addovscmean.tem
match "$1" "overscan.dat" | fields "-" 2 | scanf("%8f", x)
hedit ("$1", "OVSCMEAN", x)
```

```
cl>  !sed 's:%:bias:g' hdrlog_ovsc.tem > hdrlog_ovsc.cl
cl>  cl < hdrlog_ovsc.cl
```

△ But we didn't actually trim the bias frames (we will trim the frames for all other data types — hence the above script), so:

```
cl>  hedit ("@bias.lis", "TRIMSEC", "[1:1044,1:1044]", add-)
cl>  hedit ("@bias.lis", "IMAGSEC", "[1:1044,1:1044]", add-)
```

• Average the overscan-subtracted biases using task `imcombine`. Note, that no scaling — and in particular no *multiplicative* scaling — should be allowed, and that each bias frame should be given equal weight.

```
cl>  unlearn imcombine
cl>  imcombine ("@bias.lis", "BIAS.fits", logfile="ccdproc.log",
>>>    combine="average", reject="avsigclip", outtype="real", scale="none",
>>>    zero="none", weight="none", statsec="[16:1015,1:1000]", lthreshol=INDEF,
>>>    hthreshold=64000., nkeep=1, mclip+, lsigma=3., hsigma=3., sigscale=0.1)
```

And update the headers:

```
cl>  gdate()
cl>  hselect BIAS.fits ncombine yes | scanf ("%d", i)
cl>  print ("Combined ",i,"individual bias frames.")
Combined 179 individual bias frames.
cl>  hedit ("BIAS.fits", "NCOMBINE", add-, del+)
cl>  hedit ("BIAS.fits", "BIASCOMB", "COMPLETE")
cl>  hedit ("BIAS.fits", "NCOMBINE", i)
cl>  hedit ("BIAS.fits", "BSCBDATE", gdate.fdate)
```

• Verify that the level in the overscan region of the combined frame is exactly 0. If not, subtract the residual level.

```
cl>  imstat BIAS.fits[1032:1044,1:1000] fields="mean" format-
-4.836262E-4

cl>  unlearn imarith
cl>  imarith ("BIAS.fits", "-", "-4.836262E-4", "MASTERBIAS.fits",
>>>       title="MASTERBIAS", divzero=0., hparams="", pixtype="real",
>>>       calctype="real", verbose+, noact-, >> "ccdproc.log")
```

Inspect the resulting "MASTERBIAS.fits" frame; does the noise agree with our theoretical expectations?

```
cl>  display MASTERBIAS.fits 1 zscale- zrange- z1=-2. z2=4.
cl>  implot MASTERBIAS.fits      ⟵ see Fig. 55

cl>  s1="MASTERBIAS[1:1044,1:1044],MASTERBIAS[16:1015,1:1000]"
cl>  s1=s1//",MASTERBIAS[1032:1044,1:1000]"
cl>  imstat (s1, fields="image,mean,midpt,stddev,min,max", lower=INDEF,
>>>       upper=INDEF, nclip=3, lsigma=3., usigma=3.)
#              IMAGE                MEAN      MIDPT    STDDEV     MIN    MAX
MASTERBIAS[1:1044,1:1044]      0.2325    0.2236    0.2349 -0.5484  0.9989
MASTERBIAS[16:1015,1:1000]     0.2368    0.2289    0.2279 -0.4581  0.9302
MASTERBIAS[1032:1044,1:1000] -8.871e-4 -0.009762 0.2128 -0.6350  0.6331
```

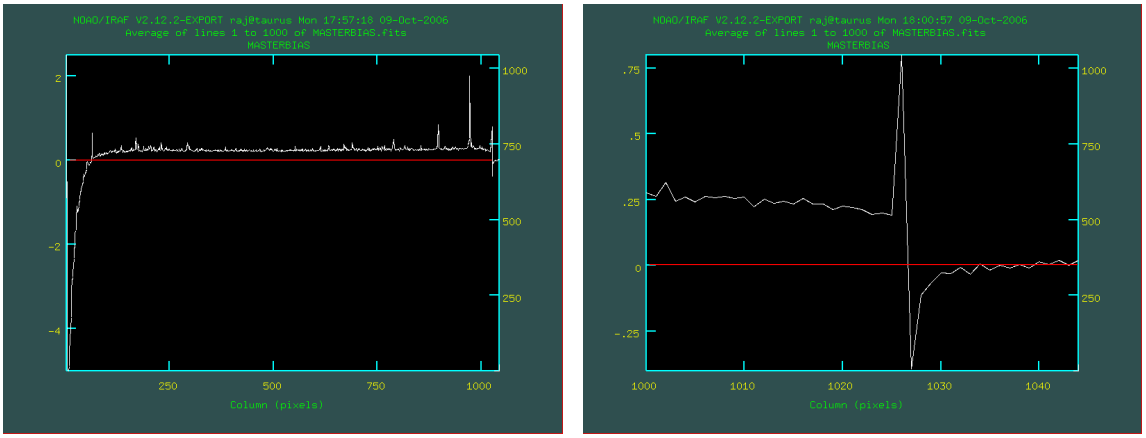

Figure 55: Screen shot of `implot` graphs of the average of lines 1–1000 of the MASTERBIAS frame and a detail thereof, showing (*a*) a strong *roll-on* of the bias level at low column numbers, and (*b*) a small but systematic offset of ∼0.24 ADU between the bias level in the overscan strip and in the image (illuminated) region of the CCD.

Assuming Gaussian read-noise, the expected noise in the absence of genuine variations in the column-to-column bias level would be:

$$\frac{1}{\sqrt{179}} \cdot \frac{5.66\,[e^-]}{1.953\,[e^-/\mathrm{ADU}]} \;=\; 0.2166\;[\mathrm{ADU}]$$

This is indeed what we measure in the overscan region ($\sigma = 0.2128$). In the image region (see Fig. 55*a*), several columns with higher levels contribute to the noise, as does the "roll-on" at low column numbers. The excess with respect to the theoretically expected noise is only:

$$100\% \cdot \frac{0.2279 - 0.2166}{0.2166} \;\sim\; 5\%$$

and can, therefore, be ignored.

From the computed image statistics, it is clear that the bias level in the image region shows a small but systematic offset with respect to that in the virtual overscan strip. The value for the mean in the "clean" portion of the illuminated region is 0.2368 ADU (see Fig. 55*b*). This offset, the few columns with elevated bias level, and the strong roll-on at low column numbers, require that we correct for the structure in the bias by subtracting the MASTERBIAS frame from all other types of CCD data in the subsequent image processing steps.

The noise level in the MASTERBIAS frame is much smaller ($\sim$25$\times$) than the read-noise, and its subtraction in subsequent image processing steps will therefore not add any appreciable noise.

If you're sure that everything is as it should be, then it is OK to delete the individual bias frames at this point:

```
cl>  imdelete @bias.lis yes verify- default+
```

8.5.2 Dark frames

Our goal is to construct a dark reference frame to try and measure the bulk dark rate of the CCD, to find whether there are pixels with significantly higher dark rates, and to check whether there is structure in the dark level.

(0) visually inspect all dark frames and convert to 32-bit format
(1) interpolate over the bad pixels/columns we defined earlier
(2) subtract the MASTERBIAS frame to correct for structure in the bias
(3) subtract the level measured in the overscan strip and trim the image (cut off the overscan and other bad/non-illuminated pixels)
(4) average the darks together, scaling by integration time
(5) measure the bulk dark rate; identify hot pixels and add them to the bad pixel file
(6) normalize the combined dark frame to a dark time of 1 sec

---

• Visually inspect all dark frames (remove any darks with problems):

```
cl>  imexamine @dark.lis 1 allframes- nframes=1
```

Convert all darks from 16-bit integer to 32-bit floating-point format:

```
cl>  chpixtype @dark.lis @dark.lis "real" oldpixtype="all" verbose+
```

Add a line to the image headers, indicating the start of the processing log entries, and add an entry logging the pixel data type conversion:

```
cl>  !sed 's:%:dark:g' hdrlog_init.tem > hdrlog_init.cl
cl>  cl < hdrlog_init.cl
```

• Interpolate over the bad pixels/columns identified earlier; log relevant information to the image headers:

```
cl>  fixpix ("@dark.lis", "badpix.ccd26", verbose+, >> "ccdproc.log")

cl>  !sed 's:%:dark:g' hdrlog_fxpx.tem > hdrlog_fxpx.cl
cl>  cl < hdrlog_fxpx.cl
```

• Correct for the structure in the bias level by subtracting the MASTERBIAS frame:

```
cl>  imarith ("@dark.lis", "-", "MASTERBIAS.fits", "@dark.lis", title="",
>>>       divzero=0., hparams="", pixtype="real", calctype="real", verbose+,
>>>       noact-, >> "ccdproc.log")
```

And update the headers:

```
cl>  type hdrlog_zrcr.tem
gdate()
hedit ("@%.lis", "ZEROCOR", "COMPLETE")
hedit ("@%.lis", "ZRCRIMAG", "/data1/raj/ast598/MASTERBIAS.fits")
hedit ("@%.lis", "ZRCRDATE", gdate.fdate)

cl>  !sed 's:%:dark:g' hdrlog_zrcr.tem > hdrlog_zrcr.cl
cl>  cl < hdrlog_zrcr.cl
```

• Subtract the bias level as measured in the virtual overscan strip. First, we'll measure that level and record it in a temporary log file:

```
cl>  !sed 's/$/[1032:1044,1:1000]/g' dark.lis > _dark.lis
cl>  imstat ("@_dark.lis", fields="image,midpt", lower=INDEF, upper=INDEF,
>>>       nclip=3, lsigma=3., usigma=3., format-, cache-, >> "overscan.dat")
cl>  delete _dark.lis yes verify-
```

Now fit and remove the overscan level from all dark frames using task `colbias`:

```
cl>  colbias ("@dark.lis", "@dark.lis", bias="[1032:1044,1:1000]", median+,
>>>       trim="[16:1015,1:1000]", interactive-, function="legendre", order=1,
>>>       low_reject=3., high_reject=3., niterate=3, logfiles="ccdproc.log")

cl>  !sed 's:%:dark:g' hdrlog_ovsc.tem > hdrlog_ovsc.cl
cl>  cl < hdrlog_ovsc.cl
```

• Average the darks using task `imcombine`. Note, that when combining darks with different integration times, multiplicative scaling according to the elapsed *dark time* should be allowed. We will also combine the darks separately for each dark time, to verify that we indeed detect the bulk dark rate.

```
cl>  hselect @dark.lis $I 'abs(darktime-600.)<=1.'  > dark600.lis
cl>  hselect @dark.lis $I 'abs(darktime-900.)<=1.'  > dark900.lis
cl>  hselect @dark.lis $I 'abs(darktime-1200.)<=1.'  > dark1200.lis

cl>  type dkcomb.tem
imcombine ("@$1.lis", output="$2.fits", logfile="ccdproc.log",
       combine="average", reject="avsigclip", outtype="real", scale="none",
       zero="none", weight="none", statsec="[1:1000,101:1000]", lthresho=INDEF,
       hthreshold=64000., nkeep=1, mclip+, lsigma=3., hsigma=3., sigscale=0.1,
       grow=1.0)
```

```
cl>  unlearn imcombine
cl>  !sed 's:$1:dark:g' dkcomb.tem | sed 's:$2:DARK:g' > dkcomb.cl
cl>  cl < dkcomb.cl

cl>  !sed 's:$1:dark600:g' dkcomb.tem | sed 's:$2:DARK600:g' > dkcomb.cl
cl>  cl < dkcomb.cl

cl>  !sed 's:$1:dark900:g' dkcomb.tem | sed 's:$2:DARK900:g' > dkcomb.cl
cl>  cl < dkcomb.cl

cl>  !sed 's:$1:dark1200:g' dkcomb.tem | sed 's:$2:DARK1200:g' > dkcomb.cl
cl>  cl < dkcomb.cl
```

Note that we use `grow=1.0`, to reject the four pixels nearest a pixel with a significantly deviant value. To reject pixels affected by cosmic rays, it is often necessary to also reject the neighbours of the pixel that received most of the charge. Note also, that we explicitly specified the header keyword to use for determining the integration time, `DARKTIME`, since the standard `EXPTIME` keyword will be 0. sec for dark frames.

Update the headers:

```
cl>  gdate()
cl>  hselect DARK.fits ncombine yes | scanf ("%d", i)
cl>  print ("Combined ",i,"individual dark frames.")
Combined 12 individual dark frames.
cl>  hedit ("DARK.fits", "NCOMBINE", add-, del+)
cl>  hedit ("DARK.fits", "IMCMB*", add-, del+)
cl>  hedit ("DARK.fits", "DARKCOMB", "COMPLETE")
cl>  hedit ("DARK.fits", "NCOMBINE", i)
cl>  hedit ("DARK.fits", "DKCBDATE", gdate.fdate)
```

Note that, because the first frame in the list of darks was a 601 second integration, all other frames were scaled to 601 sec as well, and the effective dark time of the DARK calibration image will be 601 sec:

```
cl>  hselect DARK*fits $I,darktime yes
DARK.fits       601.
DARK600.fits    601.
DARK900.fits    901.
DARK1200.fits  1200.
```

• Inspect the resulting frames. Do we see structure? Do we indeed detect the bulk dark rate?

```
cl>  imexamine DARK*.fits 1 allframes- nframes=1
cl>  display DARK.fits 1 zscale- zrange- z1=-2. z2=4.

cl>  s1="[101:900,101:900]"
cl>  s2=DARK.fits"//s1//",DARK600.fits"//s1//",DARK900.fits"//s1
cl>  s2=s2//",DARK1200.fits"//s1
cl>  imstat (s2, fields="image,mean,midpt,stddev,min,max", lower=INDEF,
>>>    upper=INDEF, nclip=3, lsigma=3., usigma=3.)
#           IMAGE              MEAN      MIDPT    STDDEV    MIN    MAX
DARK600.fits[101:900,101:900]  0.2145  0.2032    1.672   -4.808   5.237
DARK900.fits[101:900,101:900]  0.1549  0.1397    1.660   -4.835   5.144
DARK1200.fits[101:900,101:900] 0.1277  0.1198    1.192   -3.454   3.708
DARK.fits[101:900,101:900]     0.1157  0.1119    0.6041  -1.698   1.930
```

Our best estimate for the bulk dark rate from the combination of all individual darks is $dc = 3600 \cdot ((0.1119 \pm 0.6041)/601) = 0.67 \pm 3.62 e^-$/pix/hr, which is not particularly significant (only $\sim 0.18\,\sigma$). From the frames averaged for each individual exposure time we get:

600 sec: $dc = (1.22 \pm 10.0)\ e^-$/pix/hr
900 sec: $dc = (0.56 \pm 6.63)\ e^-$/pix/hr
1200 sec: $dc = (0.36 \pm 3.58)\ e^-$/pix/hr

Although, at face value, all dark rate estimates agree within a factor ∼3, and with a measurement derived over multiple observing runs from many more dark frames ($dc = 1.08 \pm 0.50\, e^-$/pix/hr; see Fig. 56), we must conclude that we did not actually detect the bulk dark rate.

Now, warm and hot pixels will have much larger ($> 4\,\sigma$) dark rates than the bulk of the pixels. So even though the S/N in the combined dark is rather low, we may still use this frame to identify such pixels. Do we find pixels that are consistently "hot" in each of the four combined frames?

The easiest way to find out is by masking the pixels that deviate by more than $+4\,\sigma$ in each of DARK600, DARK900, DARK1200 and DARK (although the latter frame is not really independent of the former three).

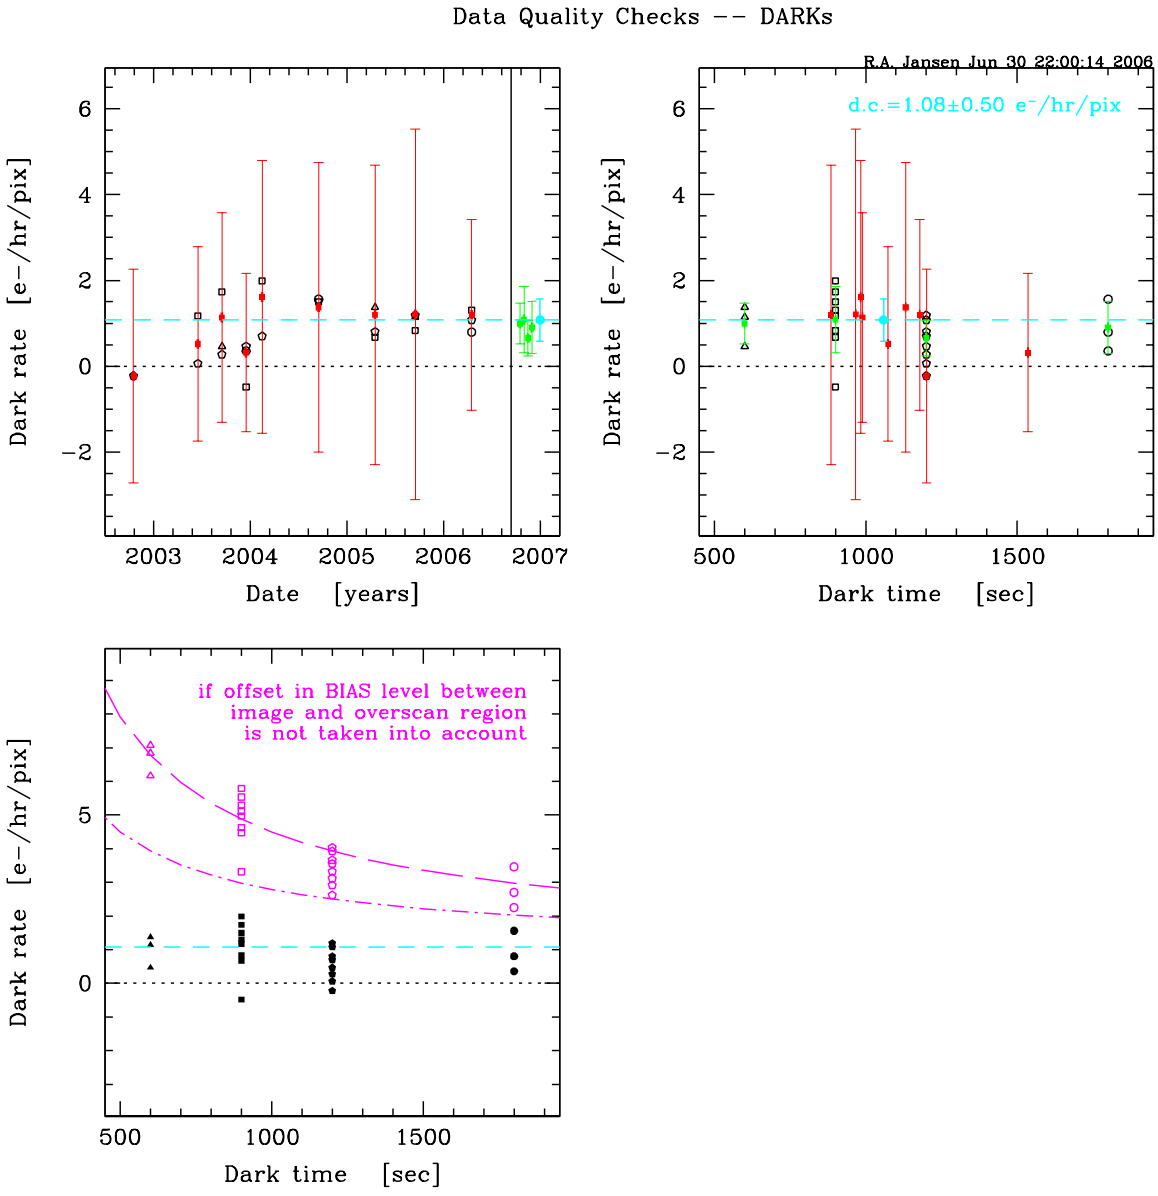

Figure 56: Example of the measurement of the bulk dark rate using data from multiple observing runs at the VATT, between 2002 and 2006. Since the dark rate of modern CCD's can be very low ($1.08 \pm 0.50\, e^-$/pix/hr here), it may not be trivial to obtain a sufficient number of high-quality dark frames for a reliable detection. In the lower left panel, the magenta data points and curve show how omitting subtracting the MASTERBIAS frame (and therefore of the offset in bias level of the image regions compared to that in the overscan strip) could lead one to an erroneous dark rate.

First, we need to fit out any gradients (the lower several tens of rows in the combined dark images appear somewhat brighter):

```
cl>  unlearn fit1d ; unlearn imreplace
cl>  fit1d.interactive=no ; fit1d.function="legendre"
cl>  fit1d.naverage=-5 ; fit1d.order=3
cl>  fit1d.low_reject=3. ; fit1d.high_reject=3. ; fit1d.niterate=3

cl>  fit1d ("DARK.fits", "tmp1.fits", "difference")
cl>  x = 4.*0.6041
cl>  imreplace ("tmp1.fits", 0., lower=INDEF, upper=x)
cl>  imreplace ("tmp1.fits", 1., lower=x, upper=INDEF)
cl>  imstat tmp1.fits fields="image,min,max"

cl>  fit1d ("DARK600.fits", "tmp2.fits", "difference")
cl>  x = 4.*1.672
cl>  imreplace ("tmp2.fits", 0., lower=INDEF, upper=x)
cl>  imreplace ("tmp2.fits", 1., lower=x, upper=INDEF)

cl>  fit1d ("DARK900.fits", "tmp3.fits", "difference")
cl>  x = 4.*1.660
cl>  imreplace ("tmp3.fits", 0., lower=INDEF, upper=x)
cl>  imreplace ("tmp3.fits", 1., lower=x, upper=INDEF)

cl>  fit1d ("DARK1200.fits", "tmp4.fits", "difference")
cl>  x = 4.*1.192
cl>  imreplace ("tmp4.fits", 0., lower=INDEF, upper=x)
cl>  imreplace ("tmp4.fits", 1., lower=x, upper=INDEF)

cl>  unlearn imsum
cl>  imsum ("tmp1,tmp2,tmp3,tmp4", "tmp.fits")
cl>  display tmp.fits 1 zs- zr- z1=-1. z2=5.
```

We want a warm/hot pixel to show up in at least 3 of the four combined images (assuming we might lose one detection due to image noise), so the lower limit for a genuine warm pixel in the combined mask image should be 3:

```
cl>  ctio  (if not already loaded)
cl>  pixselect ("tmp.fits", lower=3.0, upper=INDEF, verbose-) | count
   26
cl>  pixselect ("tmp.fits", lower=3.0, upper=INDEF, verbose-)
  698   10  3.
  843   23  4.
  184   27  3.
  982   39  4.
  ...   ..  ..
```

Since none of the identified warm/hot pixels is extremely hot (i.e., hot enough to come close to saturating that pixel in a long exposure), if we had a very high S/N DARK frame we would opt to scale and subtract that DARK image from all science exposures. That way, we would subtract the excess dark signal without distroying the collected astronomical source signal in that pixel. However, since our DARK image has relatively poor S/N, we will add the pixels identified to the bad pixel list, and interpolate over them. We will have to be careful when interpreting the signals in these pixels.

We need to reformat the x and y coordinates from the output of `pixselect` into the "`x1 x2 y1 y2`" format of the bad pixel file:

```
cl>  pixselect ("tmp.fits", lower=3.0, upper=INDEF, verbose-, >> "hotpix.lis")
cl>  !awk '{print $1,$1,$2,$2}' hotpix.lis >> badpix.ccd26
```

Clean-up of temporary files that are no longer needed:

```
cl>  delete tmp* yes verify-
```

• Normalize the combined dark image to a dark time of 1 sec to produce a "master" dark image.

If we *had* detected a bulk dark rate and, particularly, if the dark rate showed any systematic structure across the CCD, we would normalize our combined dark frame to a dark time of 1 sec (i.e, divide the frame by the integration time in seconds):

```
cl>  imarith ("DARK.fits", "/", "601.", "MASTERDARK.fits", title="MASTERDARK",
>>>       divzero=0., hparams="darktime", pixtype="real", calctype="real",
>>>       verbose+, noact-, >> "ccdproc.log")
```

Note, that by specifying "darktime" for the '`hparams`' task parameter, the value of FITS header keyword `DARKTIME` will be updated to reflect the executed arithmetic operation, i.e., will be divided by 601. $\longrightarrow$ 1.0 sec.

8.5.3 Dome flat frames

Our goal is to construct a reference frame for each filter, that records the wavelength-dependent pixel-to-pixel (small-scale) variations in the effective sensitivity of the CCD, and do so at high S/N ($\gtrsim$1000).

(0) visually inspect all dome flat frames, add filter name, and convert to 32-bit format
(1) interpolate over the bad pixels/columns we defined earlier
(2) subtract the MASTERBIAS frame to correct for structure in the bias (*optional*)
(3) subtract the level measured in the overscan strip and trim the image (cut off the overscan and other bad/non-illuminated pixels)
(4) model shutter shading correction using frames with exposures shorter than ∼3 sec
(5) median average the dome flats together, scaling by the average signal level
(6) fit a low-order surface to the combined flat, and divide by the fit to produce a normalized reponse frame, containing only the high-order structure (i.e., the pixel-to-pixel sensitivity variations)

• Visually inspect all dome flat frames and remove any flats with problems (e.g., overexposed/saturated frames, anomalous gradients, filter wheel problems, etc...):

```
cl>  imexamine @flat.lis 1 allframes- nframes=1
```

Add a proper (human readable) filter name to the image headers — at the VATT, only the *position* of each filter wheel is recorded. The observing log files (see § 8.2) give us the following translations:

Apr 2005: R = `"1 2 1 0"` (and B = `"2 1 1 0"`, but no B flats were taken)
Sep 2005: R = `"2 1 1 0"` and B = `"1 3 1 0"`

```
cl>  hselect @flat.lis $I 'telfilte="1 2 1 0"' > _flatR.lis
cl>  hselect @flat.lis $I 'telfilte="2 1 1 0"' >> _flatR.lis
cl>  hselect @flat.lis $I 'telfilte="1 3 1 0"' > _flatB.lis
cl>  hedit.add=yes ; hedit.verify=no ; hedit.show=no
cl>  hedit @_flatR.lis filter "R"
cl>  hedit @_flatB.lis filter "B"
cl>  delete _flat?.lis yes verify-
```

Convert all dome flats from 16-bit integer to 32-bit floating-point format:

```
cl>  chpixtype @flat.lis @flat.lis "real" oldpixtype="all" verbose+

cl>  !sed 's:%:flat:g' hdrlog_init.tem > hdrlog_init.cl
cl>  cl < hdrlog_init.cl
```

• Interpolate over the bad pixels/columns + warm pixels identified earlier:

```
cl>  fixpix ("@flat.lis", "badpix.ccd26", verbose+, >> "ccdproc.log")

cl>  !sed 's:%:flat:g' hdrlog_fxpx.tem > hdrlog_fxpx.cl
cl>  cl < hdrlog_fxpx.cl
```

• It is not necessary to correct (dome- and twilight-) flat frames for either structure in the bias level, or for dark current. Flat field frames (should) have high signal levels and *are normalized* to 1. Hence the maximum error one would make would be:

$$100\% \times \text{few ADU} \,/\, \text{several tens of thousands of ADU} \ll 0.1\%$$

Furthermore, flat field exposures tend to be *short* exposures (generally, less than a minute), and the dark rate is very much less than the photon rate.

• Subtract the bias level as measured in the virtual overscan strip. First, we'll measure that level and record it in a temporary log file:

```
cl>  !sed 's/$/[1032:1044,1:1000]/g' flat.lis > _flat.lis
cl>  imstat ("@_flat.lis", fields="image,midpt", lower=INDEF, upper=INDEF,
>>>       nclip=3, lsigma=3., usigma=3., format-, cache-, >> "overscan.dat")
cl>  delete _flat.lis yes verify-
```

Now, fit and remove the overscan level from all dome flats using `colbias`:

```
cl>  colbias ("@flat.lis", "@flat.lis", bias="[1032:1044,1:1000]", median+,
>>>       trim="[16:1015,1:1000]", interactive-, function="legendre", order=1,
>>>       low_reject=3., high_reject=3., niterate=3, logfiles="ccdproc.log")

cl>  !sed 's:%:flat:g' hdrlog_ovsc.tem > hdrlog_ovsc.cl
cl>  cl < hdrlog_ovsc.cl
```

• Model any shutter shading features and correct the short dome flat exposures. For the VATT CCD camera, shutter shading effects become progressively serious for exposure times shorter than ∼3 sec.

The basic approach to construct a pixel-to-pixel response frame is to combine the dome flat frames and fit out any low-order gradients, which are presumed to reflect the illumination pattern of the CCD.
Constructing the pixel-to-pixel response frame using a surface fit can create artifacts when there is a relatively high-order feature present in the flats (such as the shutter

shading feature). To fit that feature out may require a surface fit of high enough order, that genuine features (e.g., dust particles) will be (partially) fit out.
Hence, a two-step approach is required: (1) divide the combined flats for the shorter exposure times by a response frame constructed from flats with much longer exposure times to produce an image of the shutter shading pattern, and (2) fit a surface to that image (which may be a high-order fit, because all genuine variations dropped out already). Lastly, normalize the combined flat by dividing by that fit.

First, let's see what exposure times we have per filter:

```
cl>  hselect @flat.lis exptime,filter yes | sort numeric+ | unique
0.1      R       short; Sep 29 2005
0.4      R       short;     ""
0.8      R       short;     ""
1.2      R       short;     ""
2.0      R       short;     ""
5.0      R       long ;     ""
12.      R       long ; Apr 4–8 2005
14.      R       long ;     ""
15.      R       long ;     ""
30.      B       long ; Sep 29 2005
```

Create lists per exposure time and observing run:
Apr 4–8 2005:

```
cl>  hselect @flat.lis $I 'filter=="R"&&exptime>5.0' > flatR_long.lis
```

Sep 29 2005:

```
cl>  hselect @flat.lis $I 'filter=="R"&&exptime==0.1' > flatR_0.1s.lis
cl>  hselect @flat.lis $I 'filter=="R"&&exptime==0.4' > flatR_0.4s.lis
cl>  hselect @flat.lis $I 'filter=="R"&&exptime==0.8' > flatR_0.8s.lis
cl>  hselect @flat.lis $I 'filter=="R"&&exptime==1.2' > flatR_1.2s.lis
cl>  hselect @flat.lis $I 'filter=="R"&&exptime==2.0' > flatR_2.0s.lis
cl>  hselect @flat.lis $I 'filter=="R"&&exptime==5.0' > flatR_5.0s.lis
cl>  hselect @flat.lis $I 'filter=="B"' > flatB_long.lis
```

Combine flats per observing run or night, per filter, and per exposure time. Scale by the mean signal level:
Apr 4–8 2005:

```
cl>  imcombine ("@flatR_long.lis", output="FLATapr_R.fits", combine="average",
>>>       logfile="ccdproc.log", reject="avsigclip", outtype="real",
>>>       scale="mean", zero="none", weight="none", statsec="[1:1000,101:1000]",
>>>       lthreshold=INDEF, hthreshold=64000., nkeep=1, mclip=yes, lsigma=3.,
>>>       hsigma=3., sigscale=0.1, grow=1.0)
```

Normalize the combined flat. Note, that for the normalization of the flats with "long" exposures, we can use a low-order 2-D fit. In the fitting, all higher-order features are rejected iteratively. The ratio of flat and fit is then a *response frame*, containing only the pixel-to-pixel (small-scale) sensitivity variations. Task `imsurfit` performs both fitting and computing this ratio when `type_output="response"` is selected:

```
cl> unlearn imsurfit
cl> imsurfit ("FLATapr_R.fits", "RESPapr_R.fits", xorder=9, yorder=9,
>>>     type_output="response", function="spline3", cross_terms=yes,
>>>     lower=3., upper=3., ngrow=1, niter=5, div_min=0.2)
cl> imstat RESPapr_R.fits fields="image,mean,midpt,stddev,min,max" nclip=X
#           IMAGE             MEAN      MIDPT    STDDEV    MIN     MAX
nclip=0:  RESPapr_R.fits     0.9966    0.9994   0.01568   0.2119   1.131
nclip=5:  RESPapr_R.fits     0.9992    0.9999  0.005335   0.9824   1.016
```

Sep 29 2005:

```
cl> imcombine ("@flatB_long.lis", output="FLATsep_B.fits", combine="average",
>>>     logfile="ccdproc.log", reject="avsigclip", outtype="real",
>>>     scale="mean", zero="none", weight="none", statsec="[1:1000,101:1000]",
>>>     lthreshold=INDEF, hthreshold=64000., nkeep=1, mclip=yes, lsigma=3.,
>>>     hsigma=3., sigscale=0.1, grow=1.0)
```

Normalize the combined flat:

```
cl> unlearn imsurfit
cl> imsurfit ("FLATsep_B.fits", "RESPsep_B.fits", xorder=9, yorder=9,
>>>     type_output="response", function="spline3", cross_terms=yes,
>>>     lower=3., upper=3., ngrow=1, niter=5, div_min=0.2)
cl> imstat RESPsep_B.fits fields="image,mean,midpt,stddev,min,max" nclip=X
#           IMAGE             MEAN      MIDPT    STDDEV    MIN     MAX
nclip=0:  RESPsep_B.fits     0.9989    0.9995  0.008119   0.2493   1.058
nclip=5:  RESPsep_B.fits     0.9998    0.9999  0.003079   0.9905   1.009
```

Next, combine the various short and 5.0 s long R-filter flats (all from Apr 4–8 2005) per exposure time:

```
cl> unlearn imcombine
cl> imcombine.logfile="ccdproc.log" ; imcombine.combine="average"
cl> imcombine.reject="avsigclip" ; imcombine.outtype="real"
cl> imcombine.scale="mean" ; imcombine.zero="none" ; imcombine.weight="none"
cl> imcombine.statsec="[1:1000,101:1000]"; imcombine.lthreshold=INDEF
cl> imcombine.hthreshold=64000. ; imcombine.nkeep=1 ; imcombine.mclip=yes
```

```
cl>  imcombine.sigscale=0.1 ; imcombine.lsigma=3. ; imcombine.hsigma=3.
cl>  imcombine.grow=1.0
cl>  imcombine ("@flatR_0.1s.lis", output="FLATR_0.1s.fits")
cl>  imcombine ("@flatR_0.4s.lis", output="FLATR_0.4s.fits")
cl>  imcombine ("@flatR_0.8s.lis", output="FLATR_0.8s.fits")
cl>  imcombine ("@flatR_1.2s.lis", output="FLATR_1.2s.fits")
cl>  imcombine ("@flatR_2.0s.lis", output="FLATR_2.0s.fits")
cl>  imcombine ("@flatR_5.0s.lis", output="FLATR_5.0s.fits")
```

Normalize the combined flat for the long 5.0 sec exposure (where shutter shading should be negligible) into a response frame:

```
cl>  unlearn imsurfit
cl>  imsurfit ("FLATR_5.0s.fits", "RESPR_5.0s.fits", xorder=5, yorder=5,
>>>       type_output="response", function="spline3", cross_terms=yes,
>>>       lower=3., upper=3., ngrow=1, niter=5, div_min=0.2)
cl>  imstat RESPR_5.0s.fits fields="image,mean,midpt,stddev,min,max" nclip=X
#           IMAGE             MEAN      MIDPT    STDDEV    MIN      MAX
nclip=0:  RESPR_5.0s.fits     0.9991    0.9993  0.006432   0.2157   1.056
nclip=5:  RESPR_5.0s.fits     0.9999    1.0000  0.002864   0.9912   1.008
```

Construct the response frames for the shorter exposures using the 2-step approach. Since we will have many repeated commands, it's best to write a little template script:

```
cl>  type shadnorm.tem
imarith ("FLATR_%s.fits", "/", "RESPR_5.0s.fits", "SHUTR_%s.fits",
       divzero=0., pixtype="real", calctype="real", noact-)
imsurfit("SHUTR_%s.fits", "_SHUTR_%s.fits", xorder=15, yorder=15,
       type_output="fit", function="spline3", cross_terms=yes, lower=3.,
       upper=3., ngrow=1, niter=5)
imarith ("FLATR_%s.fits", "/", "_SHUTR_%s.fits", "RESPR_%s.fits",
       divzero=0., pixtype="real", calctype="real", noact-)
imstat RESPR_%s.fits fields="image,mean,midpt,stddev,min,max" nclip=5

cl>  !sed 's:%:0.1:g' shadnorm.tem > shadnorm.cl
cl>  cl < shadnorm.cl
#  IMAGE             MEAN      MIDPT    STDDEV    MIN      MAX
RESPR_0.1s.fits     0.9999    1.0000  0.003033   0.9907   1.009

cl>  !sed 's:%:0.4:g' shadnorm.tem > shadnorm.cl
cl>  cl < shadnorm.cl
#  IMAGE             MEAN      MIDPT    STDDEV    MIN      MAX
RESPR_0.4s.fits     0.9998    0.9999  0.003294   0.9899   1.010
```

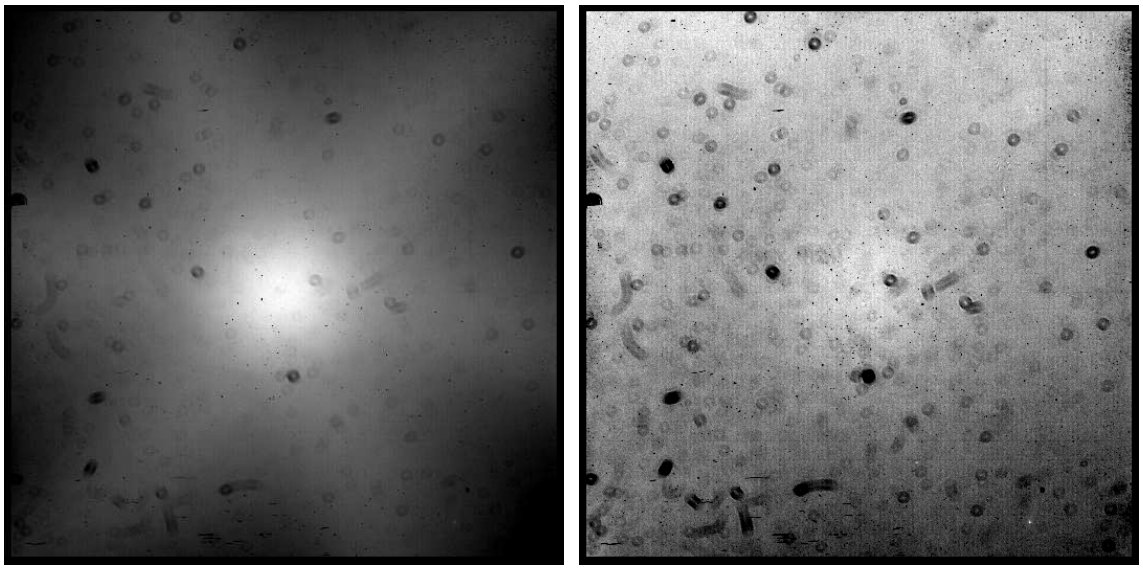

Figure 57: (*a*) Combined frame constructed from the 0.1 s dome flats in *R*. The hexagonal shutter shading pattern (see Fig. 49) is obvious. (*b*) Combined frame constructed from the 1.2 s dome flats. A central "bright spot" due to shutter shading is still apparent, but other low-order structures (e.g., a gradient from lower right to upper left) become important too. Gradients seen in dome flats need not be identical to those (if any) observed in the night-sky background.

```
cl>  !sed 's:%:0.8:g' shadnorm.tem > shadnorm.cl
cl>  cl < shadnorm.cl
#  IMAGE              MEAN      MIDPT    STDDEV    MIN     MAX
RESPR_0.8s.fits     0.9999    1.0000  0.002773   0.9915   1.008

cl>  !sed 's:%:1.2:g' shadnorm.tem > shadnorm.cl
cl>  cl < shadnorm.cl
#  IMAGE              MEAN      MIDPT    STDDEV    MIN     MAX
RESPR_1.2s.fits     0.9998    0.9999  0.003050   0.9906   1.009

cl>  !sed 's:%:2.0:g' shadnorm.tem > shadnorm.cl
cl>  cl < shadnorm.cl
#  IMAGE              MEAN      MIDPT    STDDEV    MIN     MAX
RESPR_2.0s.fits     0.9999    1.0000  0.002921   0.9910   1.009
```

Lastly, average the pixel-to-pixel response frames thus constructed together into a final high-S/N response frame:

```
cl>  !\ls -1 RESPR_*s.fits > resp_sep.lis
cl>  unlearn imcombine
```

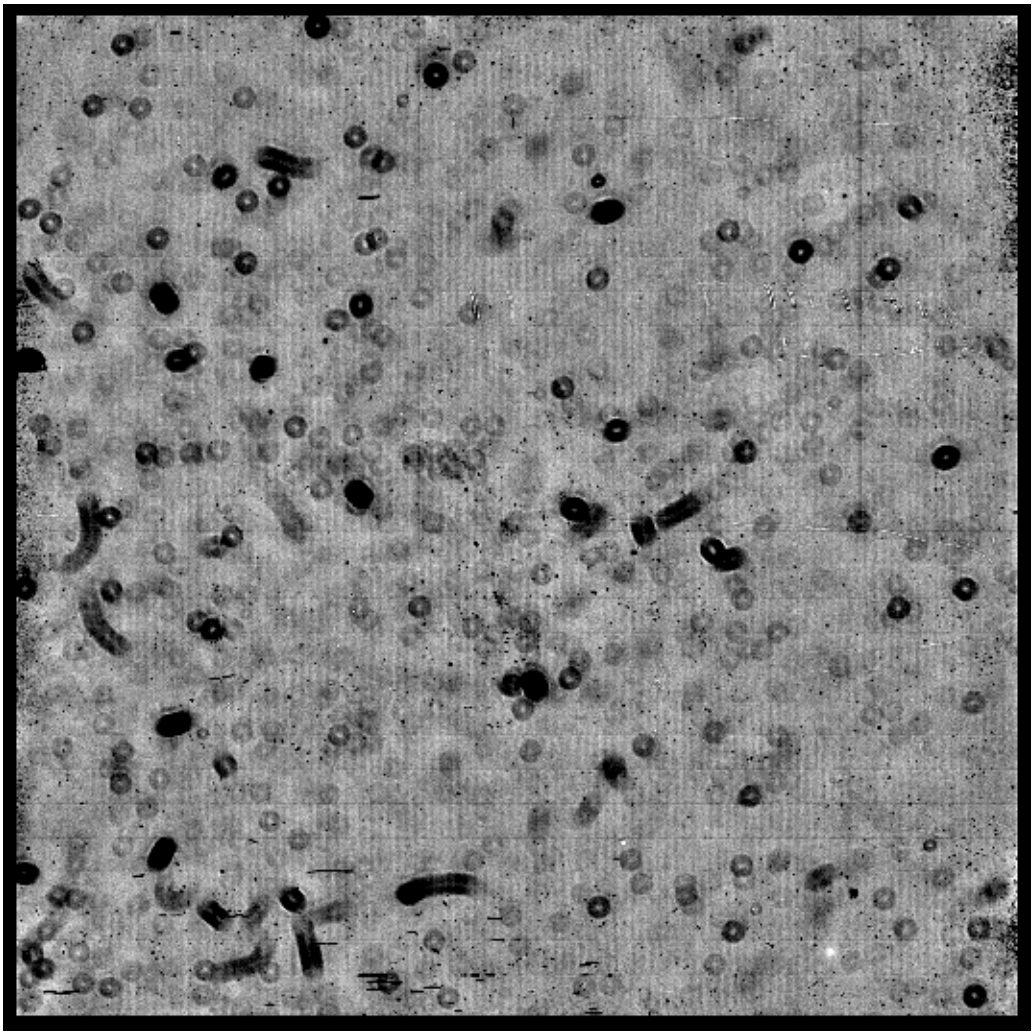

Figure 58: The final high-S/N R-filter response frame for the 2005 Sep 29 data, constructed from all dome flats taken during that night. The response frame shows only the pixel-to-pixel (small-scale) variations in the effective sensitivity of CCD images. Such variations can be *intrinsic* (pixel-to-pixel Quantum Efficiency variations) or *extrinsic* (e.g., the many semi-transparent dust and pollen particles) to the CCD. Note, that since we opted to interpolate over bad pixel columns as one of our first processing steps, they are absent from the above response frame (only the very wide column at upper right shows through).

```
cl>  imcombine ("@resp_sep.lis.lis", output="RESPsep_R.fits", combine="average",
>>>       logfile="ccdproc.log", reject="avsigclip", outtype="real",
>>>       scale="mean", zero="none", weight="none", statsec="[1:1000,101:1000]",
>>>       lthreshold=INDEF, hthreshold=64000., nkeep=1, mclip=yes, lsigma=3.,
>>>       hsigma=3., sigscale=0.1, grow=1.0)
```

Fig. 57 illustrates why we couldn't just average all individual dome flat frames together to create a high-S/N response frame. Fig. 58 shows what our resulting response frame, "RESPsep_R.fits", looks like.

Now, did we indeed reach our goal of 1,000,000 e^- per pixel?

Sep 29 2005 (see § 8.2):

0.1 s	R	$\longleftarrow$	5× 28,886.
0.4 s	R	$\longleftarrow$	5× 19,355.
0.8 s	R	$\longleftarrow$	5× 37,313.
1.2 s	R	$\longleftarrow$	5× 25,280.
2.0 s	R	$\longleftarrow$	5× 30,059.
5.0 s	R	$\longleftarrow$	8× 19,516.
			+
			860,593. ADU = 1,680,738 e^- $\longrightarrow$ OK!

30. s	B	$\longleftarrow$	10× 15,895.
			+
			158,950. ADU = 310,429 e^- $\longrightarrow \sim {}^1/_3$

Apr 4-8 2005:

12,14,15	R	$\longleftarrow$	19× 28,356.
			+
			538,764. ADU = 1,052,206 e^- $\longrightarrow$ OK

8.5.4 Twilight flat frames

Our goal is to construct a reference frame for each filter, that records the wavelength-dependent illumination pattern (large-scale variations) of the CCD.

(0) visually inspect all twilight flat frames, add filter name, and convert to 32-bpp
(1) interpolate over the bad pixels/columns we defined earlier
(2) subtract the MASTERBIAS frame to correct for structure in the bias (*optional*)
(3) subtract the level measured in the overscan strip and trim the image (cut off the overscan and other bad/non-illuminated pixels)
(4) correct frames with very short exposure times for shutter shading
(5) median average the twilight flats together, scaling by the average signal level
(6) if the $S/N \ll 1000$, then fit a low-order 2-D surface (illumination pattern) to the combined flat; if $S/N \gtrsim 1000$, use the normalized combined twilight flat directly for correction of both pixel-to-pixel sensitivity variations and illumination. If the S/N is intermediate and if the S/N of the response frame constructed from the dome flats is <1000, then the illumination pattern should be fit, but illumination-corrected dome flats and twilight flats may be combined to increase the S/N of the final response frame.

• Visually inspect all twilight sky flat frames and remove any flats with problems (e.g., overexposed/saturated frames, anomalous gradients, filter wheel problems, etc...):

```
cl>  imexamine @sky.lis 1 allframes- nframes=1
```

△ Removed frame "050929/b0151.fits" from file "sky.lis": it is saturated!

Add a proper (human readable) filter name to the FITS headers. The translations from wheel positions to filter names are as for the dome flats. Note, that wheel-position string `"2 1 1 0"` corresponds to the B filter in Apr 2005, while it denotes R in Sep 2005:

Apr 2005: $R =$ `"1 2 1 0"` and $B =$ `"2 1 1 0"`
Sep 2005: $R =$ `"2 1 1 0"` and $B =$ `"1 3 1 0"`

```
cl>  match "0504" "sky.lis" > sky_apr.lis
cl>  match "0509" "sky.lis" > sky_sep.lis
cl>  hselect @sky_apr.lis $I 'telfilte="1 2 1 0"' > _sky_aprR.lis
cl>  hselect @sky_apr.lis $I 'telfilte="2 1 1 0"' > _sky_aprB.lis
cl>  hselect @sky_sep.lis $I 'telfilte="2 1 1 0"' > _sky_sepR.lis
cl>  hselect @sky_sep.lis $I 'telfilte="1 3 1 0"' > _sky_sepB.lis
```

```
cl>  count _sky*.lis
     21      21     378 _sky_aprB.lis
     16      16     288 _sky_aprR.lis
      8       8     144 _sky_sepB.lis
      6       6     108 _sky_sepR.lis
     51      51     918 Total
cl>  hedit.add=yes ; hedit.verify=no ; hedit.show=no
cl>  hedit @_sky_aprR.lis filter "R"
cl>  hedit @_sky_aprB.lis filter "B"
cl>  hedit @_sky_sepR.lis filter "R"
cl>  hedit @_sky_sepB.lis filter "B"
cl>  delete _sky*.lis yes verify-
```

Convert all twilight flats from 16-bit integer to 32-bit floating-point format:

```
cl>  chpixtype @sky.lis @sky.lis "real" oldpixtype="all" verbose+

cl>  !sed 's:%:sky:g' hdrlog_init.tem > hdrlog_init.cl
cl>  cl < hdrlog_init.cl
```

• Interpolate over the bad pixels/columns + warm pixels identified earlier:

```
cl>  fixpix ("@sky.lis", "badpix.ccd26", verbose+, >> "ccdproc.log")

cl>  !sed 's:%:sky:g' hdrlog_fxpx.tem > hdrlog_fxpx.cl
cl>  cl < hdrlog_fxpx.cl
```

• It is not necessary to correct the sky flat frames for either structure in the bias level, or for dark current.

• Subtract the bias level as measured in the virtual overscan strip. First, we'll measure that level and record it in a temporary log file:

```
cl>  !sed 's/$/[1032:1044,1:1000]/g' sky.lis > _sky.lis
cl>  imstat ("@_sky.lis", fields="image,midpt", lower=INDEF, upper=INDEF,
>>>       nclip=3, lsigma=3., usigma=3., format-, cache-, >> "overscan.dat")
cl>  delete _sky.lis yes verify-
```

Now, fit and remove the overscan level from all twilight flats using `colbias`:

```
cl>  colbias ("@sky.lis", "@sky.lis", bias="[1032:1044,1:1000]", median+,
>>>       trim="[16:1015,1:1000]", interactive-, function="legendre", order=1,
>>>       low_reject=3., high_reject=3., niterate=3, logfiles="ccdproc.log")
```

```
cl>  !sed 's:%:sky:g' hdrlog_ovsc.tem > hdrlog_ovsc.cl
cl>  cl < hdrlog_ovsc.cl
```

• Correct the short twilight sky flat exposures for shutter shading effects.

We already produced images of the shutter shading pattern, when we processed the short dome flat exposures. The one corresponding to the very shortest exposure, "_SHUTR_0.1s.fits", has the highest fidelity (in the exposures longer than ∼1 sec, the non-uniform illumination starts to contribute, as well).

First, let's find an appropriate normalization of frame "_SHUTR_0.1s.fits", such that the center of the shutter feature, measured to be near pixel $(496, 485)$, will have a value of exactly 1, and the edges of the CCD a value of I_{border}/I_{center} :

```
cl>  unlearn imstat
cl>  imstat _SHUTR_0.1s.fits[489:503,478:492]
#               IMAGE                  MEAN  MIDPT   STDDEV   MIN     MAX
 _SHUTR_0.1s.fits[489:503,478:492]  33521. 33522.   5.119 33507.  33530.
cl>  unlearn borderstat
cl>  borderstat _SHUTR_0.1s.fits border=10 verbose+
      npix       mean     midpt       mode    stddev        min       max
     39600   27646.04  27656.68   27738.60 555.2507   26216.65  28741.52
```

So, in a 0.1 sec exposure, the exposure time for pixels on the optical axis of the shutter is 33522./27656.68 = 1.212076 times larger than the exposure time for pixels along the outer edge of the CCD.

```
cl>  unlearn imarith
cl>  imarith ("_SHUTR_0.1s.fits", "/", "33522.", "SHUTSHAD_0.1s.fits",
>>>       title="SHUTFLAT", divzero=0., pixtype="real", calctype="real",
>>>       verbose+, noact-, >> "ccdproc.log")
```

Repeating the measurements of the relative exposure times of center and border for the other exposure times, we find the following:

Exposure time, t	I_{center}	I_{border}	Relative strength, $\mathcal{S}(t)$
0.1 s	33522. ADU	27656.68 ADU	1.212076
0.4 s	20281. ADU	19134.28 ADU	1.059930
0.8 s	38218. ADU	37088.87 ADU	1.030444
1.2 s	25627. ADU	25212.62 ADU	1.016435
2.0 s	30313. ADU	30028.83 ADU	1.009463
5.0 s	19456. ADU	19462.09 ADU	0.999687

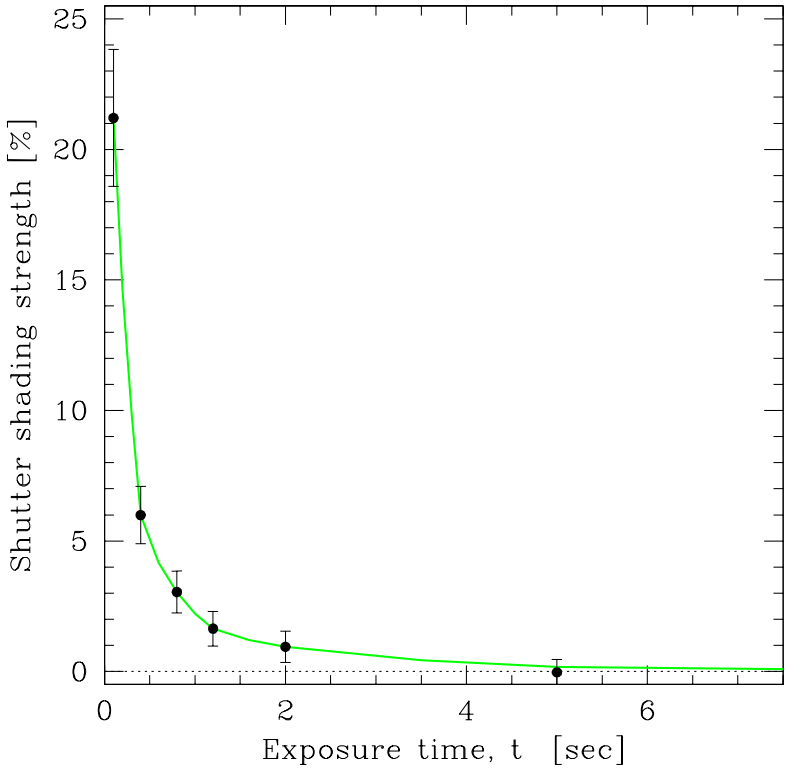

Figure 59: Relative shutter shading strength versus exposure time for the Sep 2005 run. Adopting strength $\mathcal{S}(t)$ from the fit in this graph (*green*), we can scale frame "SHUTSHAD_0.1s.fits" and correct the twilight flats with any $0.1\,\mathrm{s} < t < 5.0\,\mathrm{s}$.

Having determined the shape of the shutter shading pattern and the normalization as a function of exposure time, we can correct the twilight flats with short exposures, should the S/N in the combined longer exposures not suffice.

• Average the twilight flats together per run (or night) and per filter, scaling by the mean signal level, and fit the illumination pattern.

Create lists per observing run and filter for the longer twilight flat exposures:

```
# Apr 4–8 2005:
cl>  hselect @sky_apr.lis $I 'filter="R"&exptime>=3.0' > skyR_apr_long.lis
cl>  hselect @sky_apr.lis $I 'filter="B"&exptime>=3.0' > skyB_apr_long.lis

# Sep 29 2005:
cl>  hselect @sky_sep.lis $I 'filter="R"&exptime>=3.0' > skyR_sep_long.lis
cl>  hselect @sky_sep.lis $I 'filter="B"&exptime>=3.0' > skyB_sep_long.lis
```

Combine the sky flats per run and per filter:

```
cl>  type ftcomb.tem
imcombine ("@sky%_$_long.lis", output="SKY$_%.fits", combine="average",
       logfile="ccdproc.log", reject="avsigclip", outtype="real",
       scale="mean", zero="none", weight="none", statsec="[1:1000,101:1000]",
       lthreshold=INDEF, hthreshold=64000., nkeep=1, mclip=yes, lsigma=3.,
       hsigma=3., sigscale=0.1, grow=1.0)

cl>  unlearn imcombine
cl>  !sed 's:%:R:g' ftcomb.tem | sed 's:\$:apr:g' > ftcomb.cl
cl>  cl < ftcomb.cl

cl>  !sed 's:%:B:g' ftcomb.tem | sed 's:\$:apr:g' > ftcomb.cl
cl>  cl < ftcomb.cl

cl>  !sed 's:%:R:g' ftcomb.tem | sed 's:\$:sep:g' > ftcomb.cl
cl>  cl < ftcomb.cl

cl>  !sed 's:%:B:g' ftcomb.tem | sed 's:\$:sep:g' > ftcomb.cl
cl>  cl < ftcomb.cl
```

Normalize the combined twilight flats. Since there are no B-filter dome flats, the Apr 2005 twilight flats in B will serve both to correct for pixel-to-pixel (small-scale) sensitivity variations *and* to correct for the illumination pattern. That frame will therefore merely be normalized to a mean level of 1. The R-filter twilight flats of both Apr and Sep 2005 will be used only to fit the illumination pattern. And the Sep 2005 B-filter twilight flat will be used both to correct for the illumination pattern and to increase the S/N in the response frame constructed from the dome flats.

Apr 2005, *B*-filter twilight flat:

```
cl>  imstat SKYapr_B.fits nclip=0
#           IMAGE                MEAN    MIDPT     STDDEV     MIN       MAX
              SKYapr_B.fits     46359.   46465.     1599.    5780.    53316.

cl>  imarith ("SKYapr_B.fits", "/", "46359.", "RESPILLUMapr_B.fits",
>>>       divzero=0., hparams="", pixtype="real", calctype="real",
>>>       title="RESP+ILLUM_apr_B", verbose+, noact-, >> "ccdproc.log")

cl>  imstat RESPILLUMapr_B.fits nclip=0
#           IMAGE                MEAN    MIDPT     STDDEV     MIN       MAX
         RESPILLUMapr_B.fits        1.    1.002    0.0345   0.1247      1.15
```

Apr and Sep 2005, *R*-filter; and Sep 2005, *B*-filter, twilight flats:

```
cl>  type mkillum.tem
imarith ("SKY$_%.fits", "/", "RESP$_%.fits", "_ILLUM%_$.fits",
       pixtype="real", calctype="real")
imsurfit ("_ILLUM%_$.fits", "ILLUM$_%.fits", xorder=9, yorder=9,
       type_output="fit", function="spline3", cross_terms=yes,
       lower=3., upper=3., ngrow=1, niter=5)
imstat ("ILLUM$_%.fits", fields="mean", format-) | scan(s1)
imarith ("ILLUM$_%.fits", "/", s1, "ILLUM$_%.fits",
       title="ILLUM_$_%", pixtype="real", calctype="real")

cl>  unlearn imarith
cl>  !sed 's:%:R:g' mkillum.tem | sed 's:\$:apr:g' > mkillum.cl
cl>  cl < mkillum.cl
cl>  imstat ILLUMapr_R.fits nclip=0

cl>  !sed 's:%:R:g' mkillum.tem | sed 's:\$:sep:g' > mkillum.cl
cl>  cl < mkillum.cl
cl>  imstat ILLUMsep_R.fits nclip=0

cl>  !sed 's:%:B:g' mkillum.tem | sed 's:\$:sep:g' > mkillum.cl
cl>  cl < mkillum.cl
cl>  imstat ILLUMsep_B.fits nclip=0
```

8.6 Processing of the science frames

---

(0) visually inspect all science frames, add filter name, and convert to 32-bpp
(1) interpolate over the bad pixels/columns we defined earlier
(2) subtract the MASTERBIAS frame to correct for structure in the bias
(3) subtract the level measured in the overscan strip and trim the image (cut off the overscan and other bad/non-illuminated pixels)
(4) subtract any bulk dark signal (if significant; possibly optional)
(5) correct for shutter shading (if significant; it likely won't be)
(6) correct for the pixel-to-pixel (small-scale) variations in effective sensitivity
(7) correct for the illumination pattern (large-scale variations)

- Visually inspect all science frames and remove any that have uncorrectable problems (e.g., overexposed/saturated frames, filter wheel problems, etc...):

```
cl>  imexamine @sci.lis 1 allframes- nframes=1
```

△ Removed frame "050929/b0147.fits" from file "sci.lis": it is near saturation (∼60,000 ADU). Also removed "050929/b0089.fits", because the secondary jumped.

Add a proper (human readable) filter name to the FITS headers. The translations from wheel positions to filter names are as for the dome and twilight flats:

Apr 2005: $R =$ `"1 2 1 0"` and $B =$ `"2 1 1 0"`
Sep 2005: $R =$ `"2 1 1 0"` and $B =$ `"1 3 1 0"`

```
cl>  match "0504" "sci.lis" > sci_apr.lis
cl>  match "0509" "sci.lis" > sci_sep.lis
cl>  hselect @sci_apr.lis $I 'telfilte="1 2 1 0"' > _sci_R.lis
cl>  hselect @sci_apr.lis $I 'telfilte="2 1 1 0"' > _sci_B.lis
cl>  hselect @sci_sep.lis $I 'telfilte="2 1 1 0"' >> _sci_R.lis
cl>  hselect @sci_sep.lis $I 'telfilte="1 3 1 0"' >> _sci_B.lis
cl>  count sci.lis
     73      73    1314 sci.lis
cl>  count _sci*.lis
     27      27     486 _sci_B.lis
     46      46     828 _sky_R.lis
     73      73    1314 Total
cl>  hedit.add=yes ; hedit.verify=no ; hedit.show=no
cl>  hedit @_sci_R.lis filter "R"
cl>  hedit @_sci_B.lis filter "B"
cl>  delete _sci*.lis yes verify-
```

Convert all science frames from 16-bit integer to 32-bit floating-point format:

```
cl>  chpixtype @sci.lis @sci.lis "real" oldpixtype="all" verbose+
```

Add a line to the image headers, indicating the start of the processing log entries, and add an entry logging the pixel data type conversion:

```
cl>  !sed 's:%:sci:g' hdrlog_init.tem > hdrlog_init.cl
cl>  cl < hdrlog_init.cl
```

- Interpolate over the bad pixels/columns + warm pixels identified earlier:

```
cl>  fixpix ("@sci.lis", "badpix.ccd26", verbose+, >>"ccdproc.log")
cl>  !sed 's:%:sci:g' hdrlog_fxpx.tem > hdrlog_fxpx.cl
cl>  cl < hdrlog_fxpx.cl
```

- Correct for the structure in the bias level, and for the small offset of that level between the image region and the overscan strip, by subtracting the MASTERBIAS frame:

```
cl>  imarith ("@sci.lis", "-", "MASTERBIAS.fits", "@sci.lis", divzero=0.,
>>>    hparams="", pixtype="real", calctype="real", title="", verbose+,
>>>    noact-, >> "ccdproc.log")
cl>  !sed 's:%:sci:g' hdrlog_zrcr.tem > hdrlog_zrcr.cl
cl>  cl < hdrlog_zrcr.cl
```

- Subtract the bias level as measured in the virtual overscan strip. First, we'll measure that level and record it in a temporary log file, then we'll remove it and trim the images using task `colbias`:

```
cl>  !sed 's/$/[1032:1044,1:1000]/g' sci.lis > _sci.lis
cl>  imstat ("@_sci.lis", fields="image,midpt", lower=INDEF, upper=INDEF,
>>>    nclip=3, lsigma=3., usigma=3., format-, cache-, >> "overscan.dat")
cl>  delete _sci.lis yes verify-
cl>  colbias ("@sci.lis", "@sci.lis", bias="[1032:1044,1:1000]",
>>>    trim="[16:1015,1:1000]", median+, interactive-, function="legendre",
>>>    order=1, low_reject=3, high_reject=3, niterate=3, logfiles="ccdproc.log")
cl>  !sed 's:%:sci:g' hdrlog_ovsc.tem > hdrlog_ovsc.cl
cl>  cl < hdrlog_ovsc.cl
```

• Subtract any bulk dark signal, if significant:

We found that the bulk dark rate is very low, $\sim$1 e^-/pix/hr. Our longest science exposure is 600 s (image "050929/b0044.fits"; see § 8.2), corresponding to a bulk dark signal of $(600/3600) * 1 \simeq 0.17\,e^-$ or $\sim$0.1 ADU. Most science exposures are much shorter. We will, therefore, omit any correction for dark signal.

• Correct short science exposures for shutter shading, if significant:

The shortest science exposures are 10 s and, hence, no correction for shutter shading is required.

• Correct for pixel-to-pixel (small-scale) variations in sensitivity:

We need to process the science frames per observing run and per filter — the small-scale variations, particularly those due to dust and pollen particles, vary per run and possibly even from night to night.

```
cl>  hselect @sci_apr.lis $I 'filter=="R"' > _sci_apr_R.lis
cl>  hselect @sci_apr.lis $I 'filter=="B"' > _sci_apr_B.lis
cl>  hselect @sci_sep.lis $I 'filter=="R"' > _sci_sep_R.lis
cl>  hselect @sci_sep.lis $I 'filter=="B"' > _sci_sep_B.lis

cl>  type hdrlog_ftcr.tem
gdate()
hedit ("@%.lis", "FLATCOR", "COMPLETE")
hedit ("@%.lis", "FTCRIMAG", "/data1/raj/ast598/$.fits")
hedit ("@%.lis", "FTCRDATE", gdate.fdate)
cl>  type hdrlog_ilcr.tem
gdate()
hedit ("@%.lis", "ILLUMCOR", "COMPLETE")
hedit ("@%.lis", "ILCRIMAG", "/data1/raj/ast598/$.fits")
hedit ("@%.lis", "ILCRDATE", gdate.fdate)
```

Apr 2005, *B*-filter science frames:

```
cl>  imarith ("@_sci_apr_B.lis", "/", "RESPILLUMapr_B.fits", "@_sci_apr_B.lis",
>>>    title="", divzero=0., hparams="", pixtype="real", calctype="real",
>>>    verbose+, noact-, >> "ccdproc.log")

cl>  !sed 's:%:_sci_apr_B:g' hdrlog_ftcr.tem | sed 's:\$:RESPILLUMapr_B:g' \
>>>    > hdrlog_ftcr.cl
cl>  cl < hdrlog_ftcr.cl
cl>  !sed 's:%:_sci_apr_B:g' hdrlog_ilcr.tem | sed 's:\$:RESPILLUMapr_B:g' \
>>>    > hdrlog_ilcr.cl
cl>  cl < hdrlog_ilcr.cl
```

Apr 2005, *R*-filter science frames:

```
cl>  imarith ("@_sci_apr_R.lis", "/", "RESPapr_R.fits", "@_sci_apr_R.lis",
>>>    title="", divzero=0., hparams="", pixtype="real", calctype="real",
>>>    verbose+, noact-, >> "ccdproc.log")
cl>  imarith ("@_sci_apr_R.lis", "/", "ILLUMapr_R.fits", "@_sci_apr_R.lis",
>>>    title="", divzero=0., hparams="", pixtype="real", calctype="real",
>>>    verbose+, noact-, >> "ccdproc.log")
cl>  !sed 's:%:_sci_apr_R:g' hdrlog_ftcr.tem | sed 's:\$:RESPapr_R:g' \
>>>    > hdrlog_ftcr.cl
cl>  cl < hdrlog_ftcr.cl
cl>  !sed 's:%:_sci_apr_R:g' hdrlog_ilcr.tem | sed 's:\$:ILLUMapr_R:g' \
>>>    > hdrlog_ilcr.cl
cl>  cl < hdrlog_ilcr.cl
```

Sep 2005, *B*-filter science frames:

```
cl>  imarith ("@_sci_sep_B.lis", "/", "RESPsep_B.fits", "@_sci_sep_B.lis",
>>>    title="", divzero=0., hparams="", pixtype="real", calctype="real",
>>>    verbose+, noact-, >> "ccdproc.log")
cl>  imarith ("@_sci_sep_B.lis", "/", "ILLUMsep_B.fits", "@_sci_sep_B.lis",
>>>    title="", divzero=0., hparams="", pixtype="real", calctype="real",
>>>    verbose+, noact-, >> "ccdproc.log")
cl>  !sed 's:%:_sci_sep_B:g' hdrlog_ftcr.tem | sed 's:\$:RESPsep_B:g' \
>>>    > hdrlog_ftcr.cl
cl>  cl < hdrlog_ftcr.cl
cl>  !sed 's:%:_sci_sep_B:g' hdrlog_ilcr.tem | sed 's:\$:ILLUMsep_B:g' \
>>>    > hdrlog_ilcr.cl
cl>  cl < hdrlog_ilcr.cl
```

Sep 2005, *R*-filter science frames:

```
cl>  imarith ("@_sci_sep_R.lis", "/", "RESPsep_R.fits", "@_sci_sep_R.lis",
>>>    title="", divzero=0., hparams="", pixtype="real", calctype="real",
>>>    verbose+, noact-, >> "ccdproc.log")
cl>  imarith ("@_sci_sep_R.lis", "/", "ILLUMsep_R.fits", "@_sci_sep_R.lis",
>>>    title="", divzero=0., hparams="", pixtype="real", calctype="real",
>>>    verbose+, noact-, >> "ccdproc.log")
cl>  !sed 's:%:_sci_sep_R:g' hdrlog_ftcr.tem | sed 's:\$:RESPsep_R:g' \
>>>    > hdrlog_ftcr.cl
cl>  cl < hdrlog_ftcr.cl
cl>  !sed 's:%:_sci_sep_R:g' hdrlog_ilcr.tem | sed 's:\$:ILLUMsep_R:g' \
>>>    > hdrlog_ilcr.cl
cl>  cl < hdrlog_ilcr.cl
```

To test whether the quality of the flat fielding meets our goals (i.e., residual large-scale variations in the background level of no more than a few tenths of a percent), we can examine one or more representative frames with a non-negligible background level:

```
cl>  implot 050408/i0153.fits
:l 1 100          ⟵ plot the average of the lower 100 lines
:l 250 500        ⟵ plot the average of lines 250–500
:l 1 1000         ⟵ plot the average of all 1000 lines
:y 4095 4120      ⟵ change the vertical plot limits to 4095–4120 ADU
:c 1 1000         ⟵ plot the average of the 1000 columns
```

Some of the resulting graphs are shown in Fig. 60. The residual variations in the background level on scales of ∼100 pixels appear to be less than ∼25 ADU (peak-to-valley) on an average of ∼4102 ADU, corresponding to $\lesssim 0.61\%$. The residual variations on larger scales are less (∼0.25%).

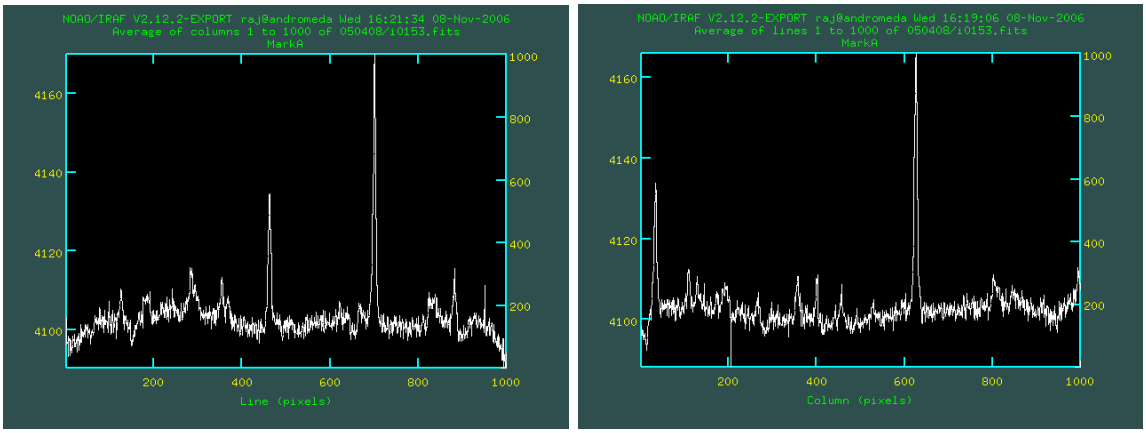

Figure 60: Graph of the average of all columns (*top*) and rows (*bottom*) in a flatfielded image (`050408/i0153.fits`) produced with task `implot`. The spikes correspond to genuine astronomical objects (mostly stars) in the image. The residual variations in the image background on scales of ∼100 pixels are less than ∼25 ADU (peak-to-valley) on an average level of ∼4102 ADU, i.e., $\lesssim 0.61\%$. When averaging 1000 lines or columns, the variations are less than ∼0.25%.

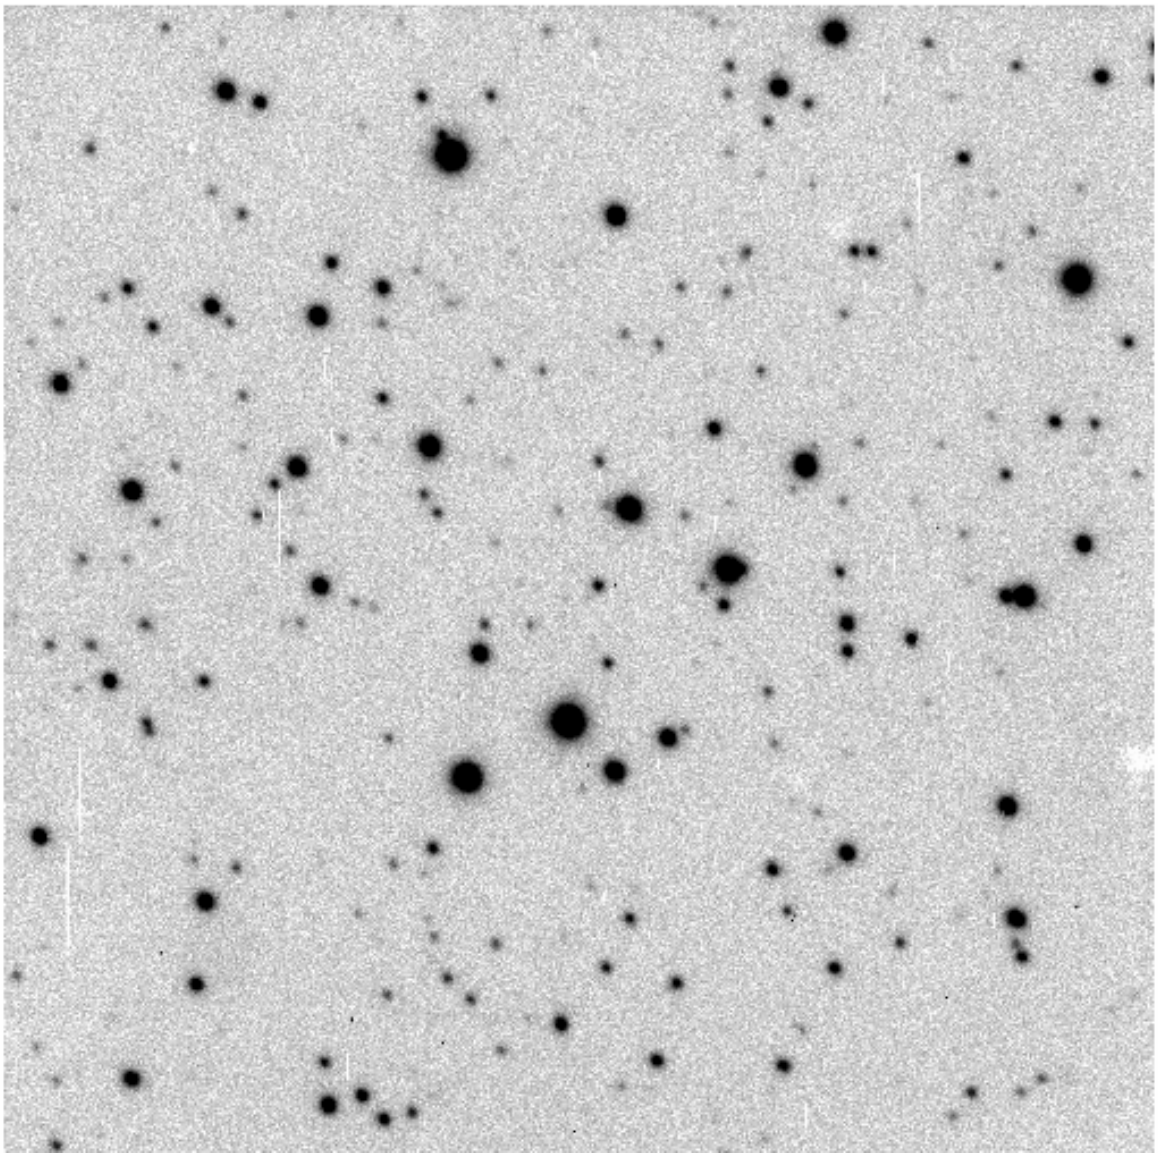

Figure 61: Processed (bias- and overscan-subtracted, flatfielded) R-filter science image of photometric standard star field Rubin 149 ("050405/f0057.fits"). Some additional bad column-sections are visible that were not visible in the flat images (whether or not a charge trap rears its ugly head can depend on the signal level!). If many frames display these columns, they should be added to the bad pixel list and all frames reprocessed. There are also a few patches visible, where the flat fielding is imperfect: dust particles moved, disappeared or were newly deposited at these locations. "Delta" flats may be constructed from dome or twilight flats taken before and after such a particle (dis)appeared/moved to correct for them.

8.7 Photometric calibration of the science frames

• Load digital photometry packages:

```
cl>  digiphot
cl>  apphot
cl>  photcal
```

• Create lists of standard star images for both April and September 2005 runs:

```
cl>  hselect @sci.lis $I,object yes | match "MarkA" | fields "-" 1 > tmp.lis
cl>  hselect @sci.lis $I,object yes | match "PG" | fields "-" 1 >> tmp.lis
cl>  hselect @sci.lis $I,object yes | match "Rubin" | fields "-" 1 >> tmp.lis
cl>  sort tmp.lis | match "0504" | sed 's:.fits::g' > standards_apr.lis
cl>  sort tmp.lis | match "0509" | sed 's:.fits::g' > standards_sep.lis
cl>  delete tmp.lis yes verify-

cl>  hselect @standards_apr.lis object yes > _std_apr.lis
cl>  hselect @standards_sep.lis object yes > _std_sep.lis
```

• Using task `markstds`, create files "*.stdcoo" listing the pixel coordinates of each photometric standard star within the observed Landolt fields. First, we need to know what number of pixels corresponds to a radius of 7″ (the aperture Landolt used). The pixel size of the VATT camera was found (through astrometric matching to 2MASS fields) to be 0.″3746, so 7″ ∼ 18.7 pixels. The choice of inner radius of the sky annulus is fairly arbitrary, but should be well away from the outer edge of the measurement aperture. The width of the annulus should be chosen such that at least a thousand pixels or so fall within the annulus.

```
cl>  type getstdcoo.tem
print ("\nLandolt field $2 ...")
markstds ("$1.fits", radstr="18.7,40,50", refstar="$2")
ctrcoo ("$1.stdcoo", "$1.stdcoo", "$1.fits",
        cntrbox=7, redispl-, verbose-)
```

Apr 2005 standards:

```
cl>  !narg getstdcoo.tem standards_apr.lis '$1' _std_apr.lis '$2' > getstdcoo.cl
cl>  type getstdcoo.cl
print ("\nLandolt field Rubin149 ...")
markstds ("050405/f0047.fits", radstr="18.7,40,50", refstar="Rubin149")
ctrcoo ("050405/f0047.stdcoo", "050405/f0047.stdcoo", "050405/f0047.fits",
        cntrbox=7, redispl-, verbose-)
print ("\nLandolt field Rubin149 ...")

    ...
```

Have the finding charts of each observed standard star field ready (e.g., Fig. 63). Then, in each image, mark the individual photometric standard stars. In each field, start with the naming (reference) star of that field, then star "A", "B", "C", etc. Note, that you may have to invert the x- and/or y-axis in `ds9` (using the `Zoom` menu) to get the same image orientation as your finding chart (N up, E to the left):

```
cl>  unlearn markstds
cl>  cl < getstdcoo.cl
Landolt field Rubin149 ...
MARKSTDS:   NOAO/IRAF2.12.2  raj@andromeda   Nov 16 01:06:17 2006
   image = 050405/f0047.fits [1000 x 1000]
Wait for the cursor cross to appear in the image display area, then mark
all (standard) stars in the image using the 'c' key.  Hit 'q' to quit.
Note:  start with the main (i.e., naming or reference) star in the field.

MARKSTDS: Finished.

Landolt field Rubin149 ...
MARKSTDS:   NOAO/IRAF2.12.2  raj@andromeda   Nov 16 01:07:03 2006

    ...
```

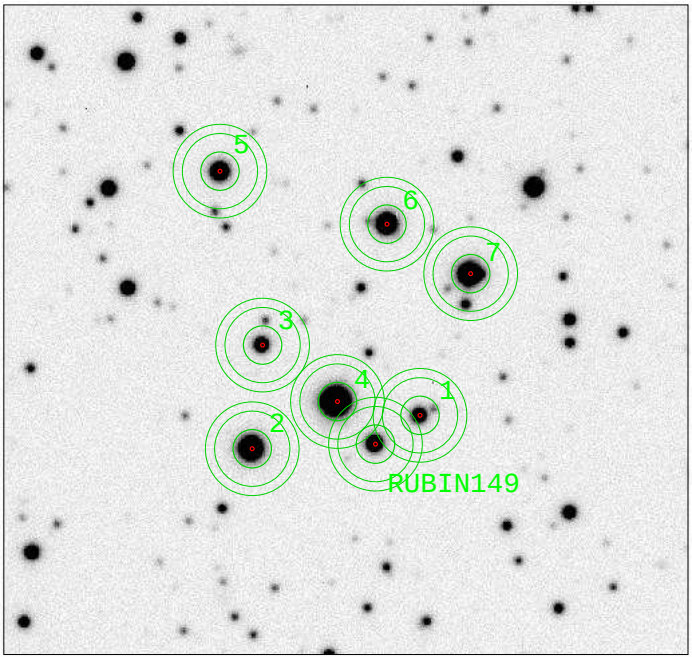

Figure 62: Display with graphical overlay after marking each of the 8 individual photometric standard stars in the field of Rubin 149 using task `markstds` on frame "050405/f0049.fits".

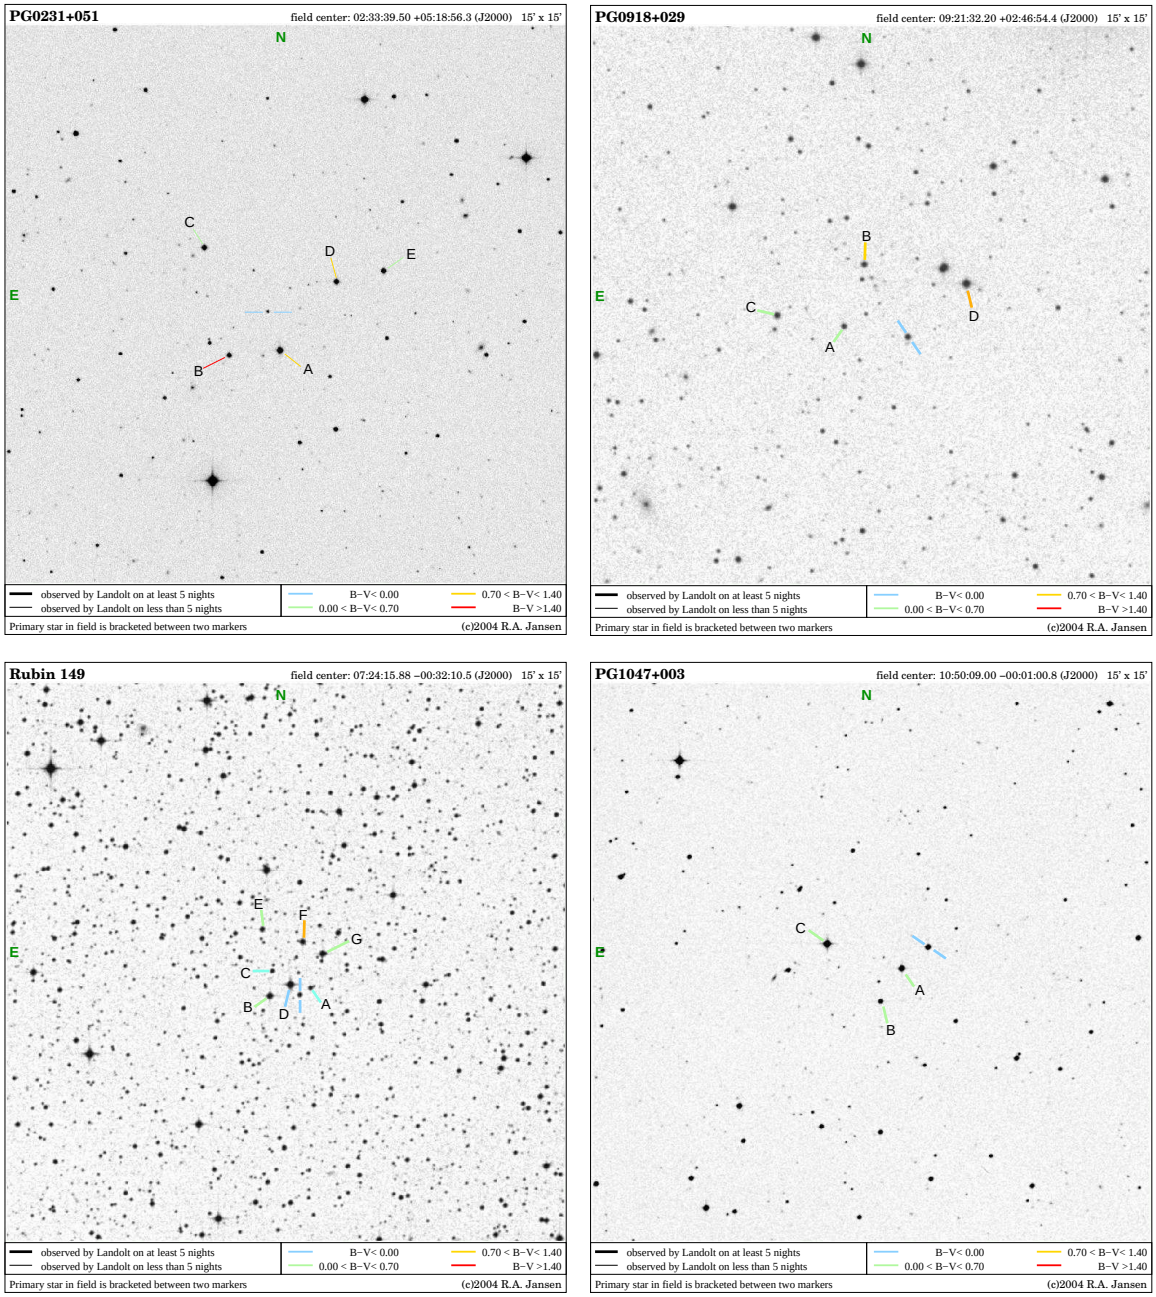

Figure 63: Finding charts for several Landolt (1992) fields with R.A. $< 12^h$. Only fields that contain multiple photometric standard stars were observed (see also § 8.2). Each field is named after the primary or reference star in the field. A convient web-accessible list of Landolt standards and finding charts can be found at *www.noao.edu/wiyn/obsprog/images/atlasinfo.html*.

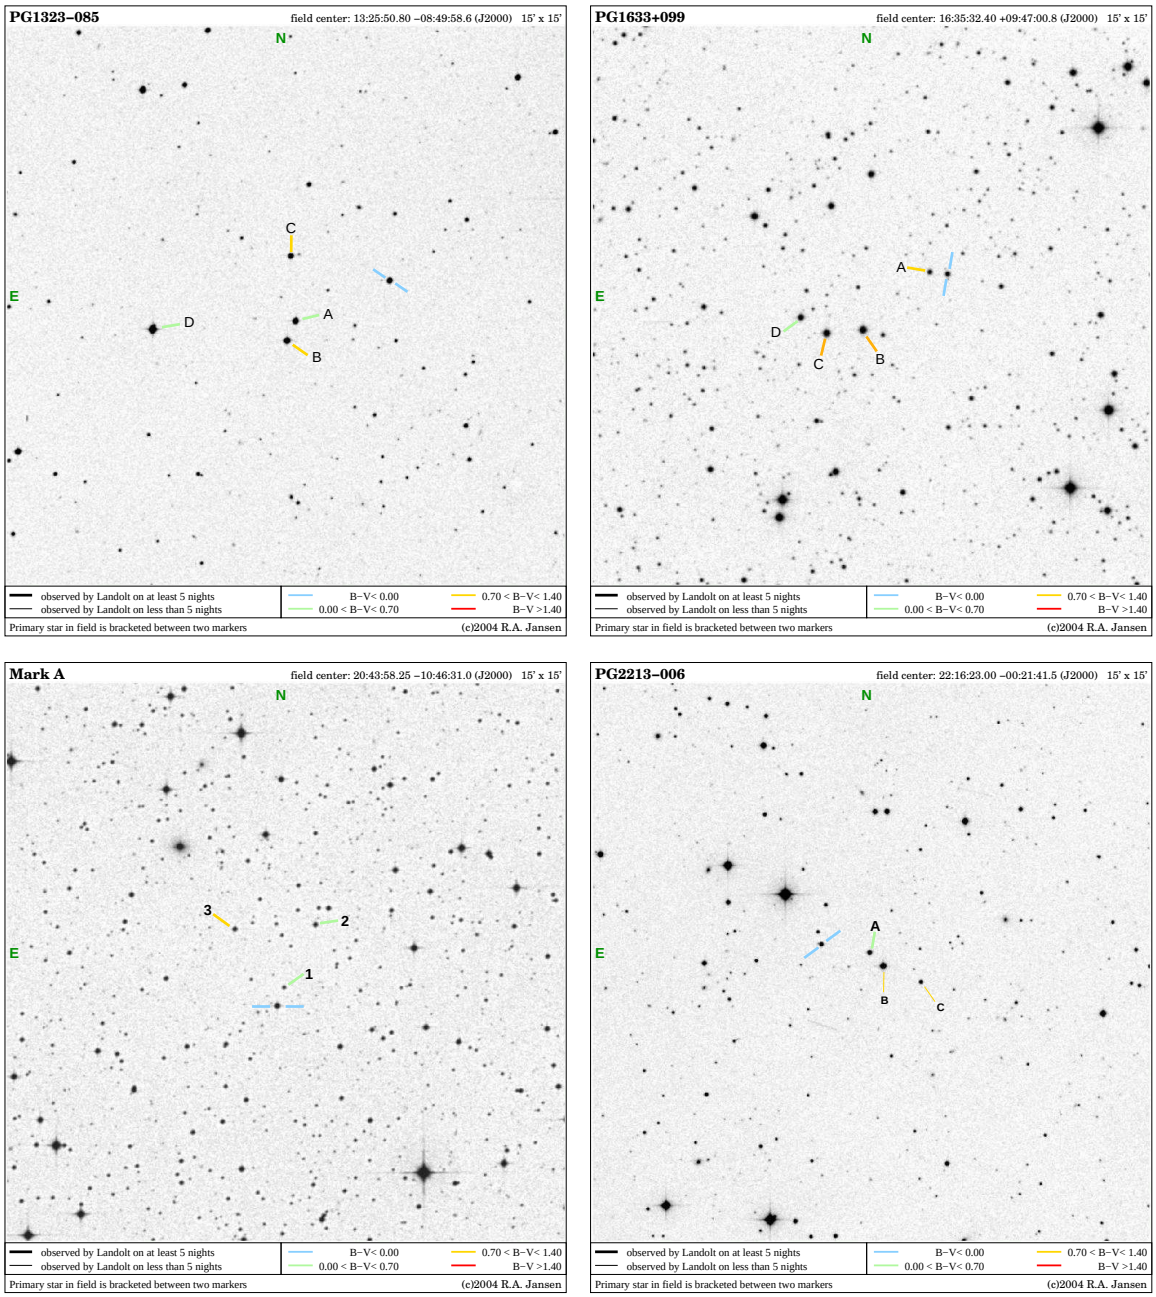

Figure 63: (cont'd) Finding charts for several Landolt standard star fields with R.A. > 12^h.

Sep 2005 standards:

```
cl>  !narg getstdcoo.tem standards_sep.lis '$1' _std_sep.lis '$2' > getstdcoo.cl
cl>  cl < getstdcoo.cl
```

• Using task `qphot`, perform aperture photometry to measure instrument magnitudes for all standard stars in each field. Note, that an arbitrary zero-point magnitude needs to be chosen. For no particular reason, I will use 26 mag throughout (any value in the 25–30 mag range would be fine, *as long as you stick with just one value throughout*). In order to correctly compute the uncertainties on the aperture magnitudes, `qphot` also needs to know the CCD gain (which we computed in § 8.4):

```
cl>  type getstdphot.tem
qphot ("%.fits", cbox=5, annulus=40., dannulus=11., apertures="18.7",
        coords="%.stdcoo", output="%.stdmag", zmag=26.,
        exposure="exptime", airmass="airmass", filter="filter", obstime="ut",
        epadu=1.953, interactive-, verbose+)

cl>  unlearn apphot ; unlearn qphot
cl>  !narg getstdphot.tem standards_apr.lis '%' > getstdphot.cl
cl>  type getstdphot.cl
qphot ("050405/f0047.fits", cbox=5, annulus=40., dannulus=11., aperture="18.7",
        coords="050405/f0047.stdcoo", output="050405/f0047.stdmag", zmag=26.,
        exposure="exptime", airmass="airmass", filter="filter", obstime="ut",
        epadu=1.953, interactive-, verbose+)

    ...

cl>  cl < getstdphot.cl
f0047.fits    447.55   674.77  40.86908   16.024  ok     ⟵ Rubin 149 (RU_149)
f0047.fits    401.74   645.45  40.06779   16.492  ok     ⟵ star RU_149A
f0047.fits    575.56   680.00  40.06578   14.409  ok     ⟵ star RU_149B
f0047.fits    564.94   572.19  39.77685   16.432  ok     ⟵ star RU_149C
f0047.fits    487.72   631.16  40.70958   13.563  ok        [etc.]
f0047.fits    608.18   392.49  40.38199   15.529  ok
f0047.fits    434.98   446.69  40.06924   15.037  ok
f0047.fits    349.10   499.03  40.25653   14.631  ok
f0048.fits    445.50   674.39  119.3431   16.031  ok     ⟵ star RU_149
f0048.fits    398.76   645.15  118.1907   16.476  ok     ⟵ star RU_149A
f0048.fits    573.51   679.53  118.1061   14.423  ok        [etc.]
f0048.fits    562.87   571.56  117.7270   16.445  ok

     ...

cl>  !narg getstdphot.tem standards_sep.lis '%' > getstdphot.cl
cl>  cl < getstdphot.cl
```

△ In what follows, we will use the tasks in the `photcal` package to simultaneously and interactively solve for the absolute photometric zeropoints, atmospheric extinction and color terms, in order to reduce our instrument magnitudes onto the Landolt photometric system. Since the syntax is rather cryptic, and since the steps required to prepare for these tasks leave little room for error, you might also want to read the package help file (`help pcintro`), the help files on the individual tasks used in the following, and Phil Massey's tutorial ("Stellar CCD Photometry using IRAF", p.22-35) for additional examples.

• Prepare matched input and configuration files. First, get an inventory of the standard star images per run, per field, and per filter:

```
# Apr 2005 standards:
cl>  hselect @standards_apr.lis $I 'object=="Rubin149"&filter=="R"'
050405/f0047 050405/f0048 050405/f0057 050408/i0034 050408/i0035 050408/i0044
cl>  hselect @standards_apr.lis $I 'object=="Rubin149"&filter=="B"'
050405/f0049 050408/i0036 050408/i0037
cl>  hselect @standards_apr.lis $I 'object=="PG0918-029"&filter=="R"'
050408/i0045 050408/i0050
cl>  hselect @standards_apr.lis $I 'object=="PG0918-029"&filter=="B"'
050408/i0046
cl>  hselect @standards_apr.lis $I 'object=="PG1047+003"&filter=="R"'
050405/f0058 050405/f0068 050405/f0122 050405/f0123 050405/f0132 050408/i0028
050408/i0033 050408/i0105 050408/i0110
cl>  hselect @standards_apr.lis $I 'object=="PG1047+003"&filter=="B"
050405/f0060 050405/f0061 050405/f0124 050405/f0125 050408/i0029 050408/i0106
cl>  hselect @standards_apr.lis $I 'object=="PG1323-086"&filter=="R"'
050408/i0111 050408/i0116
cl>  hselect @standards_apr.lis $I 'object=="PG1323-086"&filter=="B"'
050408/i0112
cl>  hselect @standards_apr.lis $I 'object=="PG1528+062"&filter=="R"'
050405/f0175 050405/f0180 050408/i0117 050408/i0122
cl>  hselect @standards_apr.lis $I 'object=="PG1528+062"&filter=="B"'
050405/f0176 050408/i0118
cl>  hselect @standards_apr.lis $I 'object=="PG1633+099"&filter=="R"'
050405/f0181 050405/f0186 050408/i0157
cl>  hselect @standards_apr.lis $I 'object=="PG1633+099"&filter=="B"'
050405/f0182 050408/i0158
cl>  hselect @standards_apr.lis $I 'object=="MarkA"&filter=="R"'
050408/i0152
cl>  hselect @standards_apr.lis $I 'object=="MarkA"&filter=="B"'
050408/i0153
```

And similarly for the Sep 2005 standards (use the field names from "_std_sep.lis").

Next, create (using your favorite ASCII text editor) a file that lists per night the matched sets of images of the same field, one R and one B filter image per line, preceded by the field identifier. "Matched" means taken during the same sequence of exposures at the same pointing (give or take a few pixels) and at very similar UT and Airmass. If no matched observation is available in B, then `INDEF` should be specified. The spaces before and after the colon following the standard star field identifier are mandatory. The files should be formatted as shown below. If you specified the image extension ".fits" as part of the image names to `qphot`, then they also need to be specified here, or `mkobsfile` will complain.

```
# Apr 2005 standards:
cl>  type 050405/standards.sets
Rubin149 : f0048.fits f0049.fits
Rubin149 : f0047.fits INDEF
Rubin149 : f0057.fits INDEF
PG1047+003 : f0058.fits f0060.fits
PG1047+003 : f0068.fits f0061.fits
PG1047+003 : f0122.fits f0124.fits
PG1047+003 : f0123.fits f0125.fits
PG1047+003 : f0132.fits INDEF
PG1528+062 : f0175.fits f0176.fits
PG1528+062 : f0180.fits INDEF
PG1633+099 : f0181.fits f0182.fits
PG1633+099 : f0186.fits INDEF
```

Note, that if we had made observations in more filters, the above file would have listed matched images in each of the observed filters, not just R and B.

```
cl>  type 050408/standards.sets
Rubin149 : i0034.fits i0036.fits
Rubin149 : i0035.fits i0037.fits
Rubin149 : i0044.fits INDEF
PG0918-029 : i0045.fits i0046.fits
PG0918-029 : i0050.fits INDEF
PG1047+003 : i0028.fits i0029.fits
PG1047+003 : i0033.fits INDEF
PG1047+003 : i0105.fits i0106.fits
PG1047+003 : i0110.fits INDEF
PG1323-086 : i0111.fits i0112.fits
PG1323-086 : i0116.fits INDEF
PG1528+062 : i0117.fits i0118.fits
PG1528+062 : i0122.fits INDEF
PG1633+099 : i0157.fits i0158.fits
MarkA : i0152.fits i0153.fits
```

Sep 2005 standards:

```
cl>  type 050929/standards.sets
MarkA : b0056.fits b0058.fits
MarkA : b0061.fits INDEF
PG2213-006 : b0050.fits b0052.fits
PG2213-006 : b0055.fits INDEF
PG2336+004 : b0092.fits b0093.fits
PG2336+004 : b0087.fits INDEF
PG2336+004 : b0111.fits b0113.fits
PG2336+004 : b0116.fits INDEF
PG0231+051 : b0081.fits b0083.fits
PG0231+051 : b0086.fits INDEF
PG0231+051 : b0117.fits b0119.fits
PG0231+051 : b0122.fits INDEF
Rubin149 : b0142.fits INDEF
```

Now, create an observations file using task `mkobsfile`, which collects and tabulates the measured instrument magnitudes for each star in each of the photometric standard star images as well as associated information extracted from the image headers (filter, UT, Airmass at the middle of the exposure). If this information is not (or incorrectly) recorded in the image headers, then the headers need to be fixed before running `mkobsfile` (*before running* `qphot`*, actually*). The default tolerance for shifts between matched images is 5 pixels. If the distances between stars in the images are relatively large, we can get away with larger shifts by selecting a larger value of the `tolerance` parameter (set here to 20 pixels). If larger shifts occur, or if the fields are relatively crowded, one would first have to construct a shifts file, in which the pixel shifts between each set of matched images is given.

```
cl>  cd 050405/
cl>  mkobsfile ("*.stdmag", "R,B", "standards.obs", imsets="standards.sets",
>>>    minmagerr=0.001, shifts="", apercors="", tolerance=20., verbose+)
Observations file:  standards.obs
        Image set:  Rubin149  8 stars written to the observations file
        Image set:  Rubin149  8 stars written to the observations file
        Image set:  Rubin149  8 stars written to the observations file
        Image set:  PG1047+003  4 stars written to the observations file
        Image set:  PG1047+003  4 stars written to the observations file
        Image set:  PG1047+003  4 stars written to the observations file
        Image set:  PG1047+003  4 stars written to the observations file
        Image set:  PG1047+003  4 stars written to the observations file
        Image set:  PG1528+062  2 stars written to the observations file
    ...
```

Check that the reported number of stars matches the number of standard stars within the field of view. If a larger number is reported, the shifts between matched images is larger than the specified tolerance.

```
cl>  cd ../050408/
cl>  mkobsfile ("*.stdmag", "R,B", "standards.obs", imsets="standards.sets",
>>>    minmagerr=0.001, shifts="", apercors="", tolerance=20., verbose+)
```

△ There are three matched sets, where there is a large shift between the exposures: "i0106" is shifted $(-105.4, -53.9)$ pixels with respect to "i0105", "i0153" $(0.4, 344.8)$ pixels with respect to "i0152", and "i0158" $(-106.3, 0.6)$ pixels with respect to "i0157". We have to take these sets out of the "standards.sets" file and handle them seperately:

```
cl>  !egrep -ve '(i0106|i0153|i0158)' standards.sets > _standards.sets
cl>  !egrep -e '(i0106|i0153|i0158)' standards.sets > standards1.sets
cl>  !\mv _standards.sets standards.sets
cl>  delete standards.obs,fstandards.obs.dat yes
cl>  mkobsfile ("*.stdmag", "R,B", "standards.obs", imsets="standards.sets",
>>>    minmagerr=0.001, shifts="", apercors="", tolerance=20., verbose+)
cl>  type standards1.shifts
i0105.fits    0.0    0.0
i0106.fits  105.4   53.9
i0152.fits    0.0    0.0
i0153.fits   -0.4 -344.8
i0157.fits    0.0    0.0
i0158.fits  106.3   -0.6
cl>  mkobsfile ("*.stdmag", "R,B", "standards1.obs", imsets="standards1.sets",
>>>    minmagerr=0.001, shifts="standards1.shifts", apercors="", tolerance=5.)
Observations file:  standards1.obs
        Image set:  PG1047+003  4 stars written to the observations file
        Image set:  PG1633+099  5 stars written to the observations file
        Image set:  MarkA  4 stars written to the observations file
```

And merge the two observations tables (omit the 3 header lines):

```
cl>  !tail +4 standards1.obs >> standards.obs
```

And for the Sep 2005 standards:

```
cl>  cd ../050929/
cl>  mkobsfile ("*.stdmag", "R,B", "standards.obs", imsets="standards.sets",
>>>    minmagerr=0.001, shifts="", apercors="", tolerance=20., verbose+)
```

Along with the "standards.obs" files, `mkobsfile` also outputs files "fstandards.obs.dat" that describe the format of the former. The files "standards.obs" for each night need to be edited, such that they list the proper identification for each individual star as published in Landolt (1992; his Table 2), or as included within IRAF's `photcal` catalogs directory in file: `noao$digiphot/photcal/catalogs/nlandolt.dat`.
If you intend to use standard stars from other sources/fields or for near-IR filters, be sure to read: `noao$digiphot/photcal/catalogs/README`

```
cl>  cd ../050405/
cl>  copy standards.obs standards.obs.orig
```

Before editing:

```
cl>  type standards.obs.orig
# FIELD         FILTER           OTIME AIRMASS  XCENTER   YCENTER     MAG   MERR
Rubin149-1      R            2:58:21.0  1.256   445.502   674.385  16.031  0.002
*               B            2:59:47.0  1.258   456.746   672.236  15.770  0.002
Rubin149-2      R            2:58:21.0  1.256   398.765   645.148  16.476  0.003
*               B            2:59:47.0  1.258   410.493   643.013  16.808  0.003
Rubin149-3      R            2:58:21.0  1.256   573.511   679.531  14.423  0.001
*               B            2:59:47.0  1.258   584.623   677.585  15.292  0.001
Rubin149-4      R            2:58:21.0  1.256   562.869   571.558  16.445  0.002
*               B            2:59:47.0  1.258   574.052   569.640  16.640  0.003
  ...           ..             ...       ...       ...       ...      ...    ...
```

After editing:

```
cl>  type standards.obs
# FIELD         FILTER           OTIME AIRMASS  XCENTER   YCENTER     MAG   MERR
RU_149          R            2:58:21.0  1.256   445.502   674.385  16.031  0.002
*               B            2:59:47.0  1.258   456.746   672.236  15.770  0.002
RU_149A         R            2:58:21.0  1.256   398.765   645.148  16.476  0.003
*               B            2:59:47.0  1.258   410.493   643.013  16.808  0.003
RU_149B         R            2:58:21.0  1.256   573.511   679.531  14.423  0.001
*               B            2:59:47.0  1.258   584.623   677.585  15.292  0.001
RU_149C         R            2:58:21.0  1.256   562.869   571.558  16.445  0.002
*               B            2:59:47.0  1.258   574.052   569.640  16.640  0.003
  ...           ..             ...       ...       ...       ...      ...    ...
```

Remove any standards stars of which the images (perhaps only the stellar cores!) are saturated from file "standards.obs". Also weed out stars with cosmic ray hits within the measurement aperture (if you have relatively long standard star exposures), which

may require visual inspection:

```
cl>  type standards.stat.tem
imexamine ("%.fits", 1, "", logfile="standards.stat", keeplog+,
        defkey="a", allframes-, nframes=1, ncstat=27, nlstat=27,
        imagecur="%.stdcoo")
cl>  !\ls -1 *.stdcoo | sed 's:.stdcoo::g' > standards.stat.lis
cl>  !narg standards.stat.tem standards.stat.lis '%' > standards.stat.cl
cl>  cl < standards.stat.cl

cl>  type standards.stat
# [1] f0047.fits - Rubin149
#  COL    LINE   COORDINATES      R    MAG    FLUX      SKY    PEAK    E    PA
 447.68  674.81 447.68 674.81  16.92  12.53  97314.    40.90   1995.  0.02  -65
 401.36  645.43 401.36 645.43  17.95  13.00  63112.    40.12   1150.  0.22   22
 575.69  679.98 575.69 679.98  16.69  10.93 425754.    44.12   8743.  0.04  -63
 564.97  572.08 564.97 572.08  16.90  12.95  65834.    40.81   1357.  0.04  -58
 487.78  631.07 487.78 631.07  17.09  10.08 929397.    49.58  18752.  0.04  -62
 608.10  392.42 608.10 392.42  16.22  12.05 151762.    41.66   3236.  0.04  -57
 435.00  446.71 435.00 446.71  16.74  11.56 237520.    43.13   4981.  0.06  -53
 349.08  499.01 349.08 499.01  16.40  11.16 344955.    44.96   7090.  0.08   -5
# [1] f0048.fits - Rubin149
#  COL    LINE   COORDINATES      R    MAG    FLUX      SKY    PEAK    E    PA
 445.51  674.22 445.51 674.22  20.70  11.35 289461.    122.6   3822.  0.01   56
 398.96  644.99 398.96 644.99  20.93  11.77 195043.    119.8   2420.  0.21   27
  ...     ...    ...    ...     ...    ...    ...      ...     ...    ...   ...
```

In output table "standards.stat", check the column containing the peak intensities (`PEAK`) for values above ∼62,000 ADU or so, and remove any found from the observations file.

• Create, edit and verify a configurations file using task `mkconfig`. The resulting configuration file should list the proper transformation equations, which (in our case) should be in terms of $B-R$ colors (instead of $B-V$ and $V-R$). Note, that `mkconfig` dumps you into a `vi` text editor. Delete the transformation equations for U, V and I, and edit the default ones for B and R to look like shown below. To exit, hit `ESQ` (if you're in 'insert' mode) and `:wq`:

```
cl>  unlearn mkconfig
cl>  mkconfig ("standards.cfg", "nlandolt", "standards.obs", "nlandolt",
>>>    verbose+)
```

```
#  Declare the new Landolt UBVRI standards catalog variables

catalog

V       4               # the V magnitude
BV      5               # the (B-V) color
UB      6               # the (U-B) color
VR      7               # the (V-R) color
RI      8               # the (R-I) color
VI      9               # the (V-I) color

error(V)  12            # the V magnitude error
error(BV) 13            # the (B-V) color error
error(UB) 14            # the (U-B) color error
error(VR) 15            # the (V-R) color error
error(RI) 16            # the (R-I) color error
error(VI) 17            # the (V-I) color error
# Declare the observations file variables

observations

TR              3               # time of observation in filter R
XR              4               # airmass in filter R
xR              5               # x coordinate in filter R
yR              6               # y coordinate in filter R
mR              7               # instrumental magnitude in filter R
error(mR)       8               # magnitude error in filter R

TB              10              # time of observation in filter B
XB              11              # airmass in filter B
xB              12              # x coordinate in filter B
yB              13              # y coordinate in filter B
mB              14              # instrumental magnitude in filter B
error(mB)       15              # magnitude error in filter B

# Customized transformation section to reduce our mags onto the Landolt system

transformation

fit   b1=0.0, b2=0.35, b3=0.000
const b4=0.0
BFIT : mB = (BV + V) + b1 + b2 * XB + b3 * (BV + VR) + b4 * (BV + VR) * XB

fit   r1=0.0, r2=0.08, r3=0.000
const r4=0.0
RFIT : mR = (V - VR) + r1 + r2 * XR + r3 * (BV + VR) + r4 * (BV + VR) * XR
```

```
   <ESQ>
   :wq
```

```
   "standards.cfg" 46L, 1554C written


** Beginning of compilation **

** End of compilation **

CATALOG VARIABLES, COLUMNS, AND ERROR COLUMNS:

 1  V   4       12
 2  BV  5       13
 3  UB  6       14
 4  VR  7       15
 5  RI  8       16
 6  VI  9       17

OBSERVATIONAL VARIABLES, COLUMNS, AND ERROR COLUMNS:

 1  TR  3       INDEF
 2  XR  4       INDEF
 3  xR  5       INDEF
 4  yR  6       INDEF
 5  mR  7       8
 7  TB  10      INDEF
 8  XB  11      INDEF
 9  xB  12      INDEF
10  yB  13      INDEF
11  mB  14      15

FIT AND CONSTANT PARAMETER VALUES:

 1  b1  0.
 2  b2  0.35
 3  b3  0.
 4  b4  0.      (constant)
 5  r1  0.
 6  r2  0.08
 7  r3  0.
 8  r4  0.      (constant)

TRANSFORMATION EQUATIONS:

 1 BFIT:  mB = (BV+V)+b1+b2*XB+b3*(BV+VR)+b4*(BV+VR)*XB
   delta(BFIT, b1) = 0.1
   delta(BFIT, b2) = 0.1
   delta(BFIT, b3) = 0.1
   delta(BFIT, b4) = 0.1
   error = ,  min = ,  max =
   weight = ,  min = ,  max =
   plot  x = ,  y =
```

```
 2 RFIT:  mR = (V-VR)+r1+r2*XR+r3*(BV+VR)+r4*(BV+VR)*XR
   delta(RFIT, r1) = 0.1
   delta(RFIT, r2) = 0.1
   delta(RFIT, r3) = 0.1
   delta(RFIT, r4) = 0.1
   error = ,  min = ,  max =
   weight = ,  min = ,  max =
   plot  x = ,  y =

Catalog input variables         = 12
First catalog column            = 4
Last catalog column             = 17

Observational input variables   = 12
First observational column      = 3
Last observational column       = 15

Fitting parameters              = 6
Constant parameters             = 2

Auxiliary (set) equations       = 0
Transformation equations        = 2

Warnings                        = 0
Errors                          = 0
```

⟶ `standards.cfg`

• Run task `fitparams` to interactively solve (fit) the transformation equations listed in `standards.cfg`, one at a time. Check that solutions indeed converge. On photometric nights, the RMS scatter should be no more than ~ 0.03 mag after rejection of outliers, and the reduced χ^2 should be close to 1.

```
cl>  unlearn fitparams
cl>  fitparams ("standards.obs", "nlandolt", "standards.cfg", "standards.sol",
>>>    weighting="photometric", nreject=5, logfile="photcal.log",
>>>    log_fit+, log_results+)

:tolerance 1e-2     ⟵ make fit less restrictive; allow poor fit to converge
h     ⟵ plot function (observed B mag + mean offset) against fit (B mag)
t     ⟵ toggle overlay with indication of fit (open boxes)
l     ⟵ plot residuals with respect to the fit versus function
=     ⟵ hardcopy of screen to default print device (optional)

d     ⟵ delete some of the most pronounced outliers
f     ⟵ refit
:tolerance 1e-3     ⟵ reduce tolerance somewhat to see if fit still converges
```

```
f     ⟵ yes, it does
k     ⟵ plot the residuals versus time of observation, UT (denoted by 'TB')
=
g j     ⟵ redefine the 'j' key to plot:
 XB,residuals     ⟵ residuals versus Airmass (symbolically denoted by 'XB')
j
=
g j     ⟵ redefine the 'j' key to plot:
 BV,residuals     ⟵ residuals versus B−V color
j
d     ⟵ delete some additional outliers
f     ⟵ refit; RMS is down to 0.0???  mag
=
g j     ⟵ redefine the 'j' key to plot:
 er(mB),residuals     ⟵ residuals vs. error in instrumental B magnitude
j
:errors     ⟵ print an overview of the errors and fitted parameters
#Wed 01:32:41 06-Dec-2006
#mB = (BV+V)+b1+b2*XB+b3*(BV+VR)+b4*(BV+VR)*XB
#Solution converged

low_reject    3.
high_reject   3.
nreject      5
grow          0.
tol        0.001
maxiter      15

niterations    10
total_points     ??
rejected      ??
deleted       ??
standard deviation 0.0??????????
reduced chi       1.?????????
average error     0.0????????
average scatter     0.0???????
RMS          0.0????????

 parameter      value       error
        b1 ?.???????  0.0??????     (fit)
        b2 0.???????  0.00?????     (fit)
        b3 0.0??????  0.00?????     (fit)
        b4 0.0000000  0.0000000(constant)
```

Since the reduced $\chi^2 \sim 1$ and the RMS < 0.03 mag, and since the fitted values appear to make sense (errors on fit <10% of the fitted values), let's accept this fit as a valid solution and move on to the next filter:

```
q
 yes
 next
```

And similar key-command sequences for the R filters, but with `TR` instead of `TB`, and `XR` instead of `XB`.

$\longrightarrow$ `standards.sol` and `photcal.log`

Check in file 'photcal.log' that nothing is listed under the header "`#UNMATCHED OBJECT`" (i.e., all stellar indentifications were indeed found in the '`nlandolt`' standard table. Check under the header "`#RESULTS: BFIT`" and "`#RESULTS: RFIT`" whether there are any stars that are consistently bad and, hence, rejected from the fit (values of `INDEF` in each column).

```
cl>  type standards.sol
```

The fitted photometric transformation coefficients, p_z, p_e and p_c, (and errors thereon) are (1) the photometric zeropoint offsets with respect to the assumed zeropoint of the instrument magnitudes of 26.0 mag (see `qphot` above), (2) the atmospheric extinction coefficient in magnitudes per airmass, and (3) the first-order color term, $\Delta(B-R)$ per 1 magnitude of $B-R$, which reflects the difference between the natural system of your telescope, instrument and detector, and that of the Landolt system.
If the system throughput in the B and R filters are reasonably similar to those of Landolt, then higher-order color terms are probably not required (besides, we here consider only two filters, so a fit of such higher-order terms will likely be uncertain or not converge at all).
Note, that the absolute zeropoint magnitude that places our instrument magnitudes on the Landolt photometric system is (26 − offset) mag (i.e., we subtract, not add, the fitted zeropoint offset from our assumed instrument zeropoint).
To photometrically calibrate any instrument magnitude measured during the given night, simply compute:

$$\begin{aligned} m &= (m_{\rm instr} - 26) + (26 - p_z) + p_e \cdot \text{airmass} + p_c \cdot \text{color} \\ &= m_{\rm instr} - p_z + p_e \cdot \text{airmass} + p_c \cdot \text{color} \end{aligned}$$

where color is $B - R$ in our case, and p_z the zeropoint offset. This equation must be applied *iteratively*, since the correct color is not known a priori (only the difference between the two instrument magnitudes, i.e., the color on the instrument's natural system, is known). Alternatively, to avoid loss of precision or introduction of additional correlated errors, one would follow the procedure given in Appendix B of Jansen et al. 2000, ApJS 126, 271.

8.8 Image analysis: differential photometry of SN 2005bk

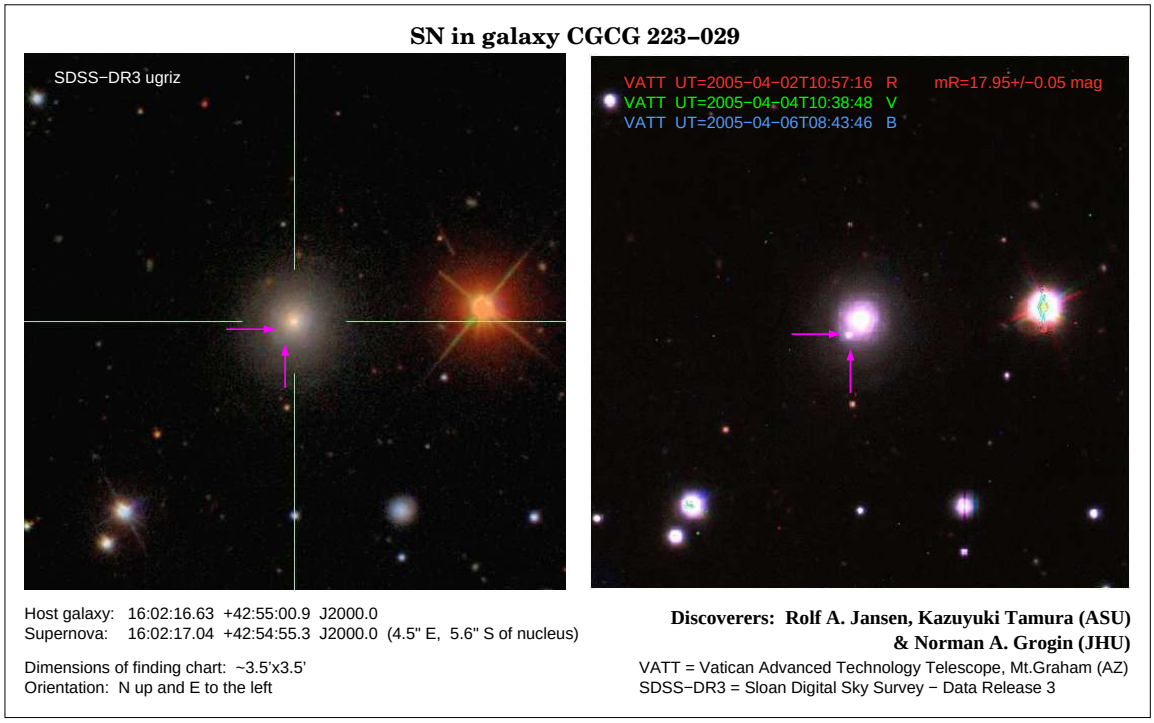

Figure 64: Finding chart with the discovery images of SN 2005bk in CGCG 223−029 (A16006+4302).

9 A practical example of CCD Spectroscopic data reduction

The basic processing steps for the processing of CCD spectroscopic data follow those discussed for CCD image processing. Here, also, our goal is to separate instrumental signals and imperfections (as result from CCD, instrument, telescope and atmosphere) from the astronomical signals. The main difference is that along one direction of each CCD frame (usually more or less aligned with either the lines or the columns of the CCD), wavelength rather than spatial information is recorded. We will discuss here the case of longslit spectroscopy, i.e., spectroscopic observations through a single slit that samples both the light from the science target and the sky background on either side. Multislit-, fiber- and slitless spectroscopy are merely special cases of long-slit spectroscopy.
Useful reading includes Chapter 6 from Howell (2006), and "A User's Guide to Reducing Slit Spectra with IRAF" by Massey, Valdez & Barnes (1992), pp. XX–YY.

9.1 (Additional) Calibration Data to be Acquired at the Telescope

For CCD spectroscopy, in addition to the types of calibration observations that were discussed in § 8 for CCD imaging, one needs to obtain:

- *wavelength reference* or *comparison* exposures (*comps*) — these are usually spectra of internal arc lamps that produce HeNeAr(Fe/Cu) (see Fig. 66 for an example) or HgCd(He/Ne/Ar) emission-line spectra. Since flexure in a spectrograph causes pointing-dependent shifts in wavelength, each science image needs to be preceded and/or followed by such a comparison lamp spectrum while the telescope is pointed at that science target. For long exposures, it is recommended to take such exposures both before *and* after each science exposure.

- If accurate redshifts are the goal, then one would also observe *velocity standards* (these can be Galactic stars, but also some well-observed galaxies), in order to take out the motion of the Earth about its axis and around the Sun, and — for extragalactic work —, the motions of the Sun around the Milky Way center and of the Milky Way with respect to a suitable standard of rest (e.g., the cosmic microwave background).

- If flux-calibrated spectra are the goal, then one would also observe *spectrophotometric standard stars* to correct for the wavelength-dependent response of the system (atmosphere, telescope, instrument and detector).

Figure 65: Example of the different types of calibration frames required or common for CCD spectroscopy of a science target.

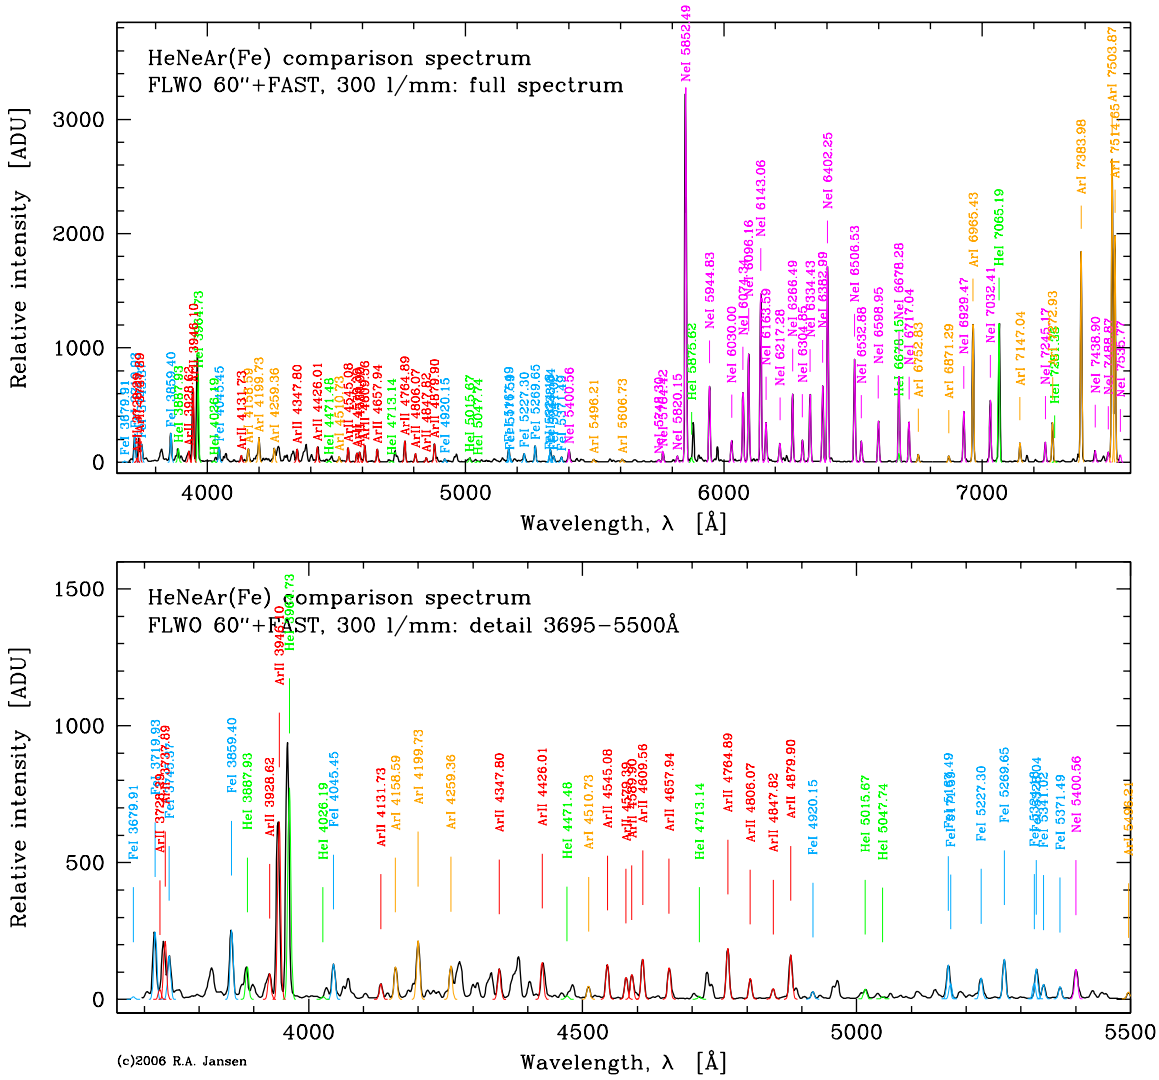

Figure 66: Example of a HeNeAr(Fe) comparison lamp spectrum as observed with the FAST spectrograph (Fabricant et al. 1998) at the FLWO 60″ telescope. Lines from Fe I, He I (generally faint and blended with other lines), Ne I, Ar I and Ar II are overlaid and labeled with their wavelengths. The top panel shows the full spectrum, while in the bottom panel the region between 3695Å and 5500Å is displayed in greater detail. Convenient tables of wavelengths and line identifications are included within IRAF (see `noao$lib/linelists/README`), but the relative line intensities may vary between lamps and spectrographs.

Another good source for wavelengths and plots of their relative strengths in comparison spectra for several types of gas mixtures is: *www.noao.edu/kpno/specatlas/*.

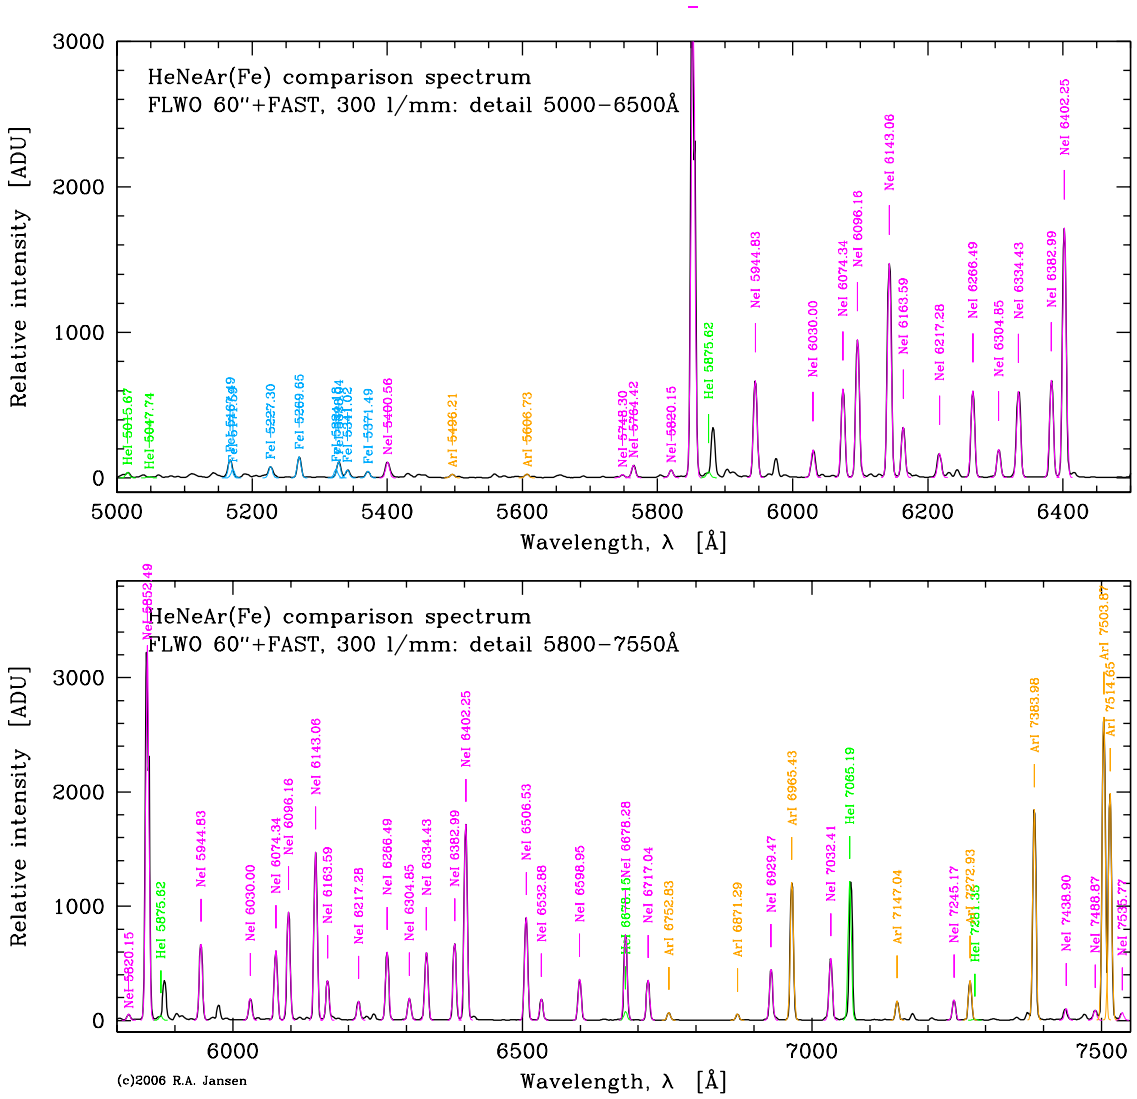

Figure 66: (cont'd) The region between 5000Å and 6500Å (*top*) and that between 5800Å and 7550Å (*bottom*).

9.2 Observing Log of the Spectroscopic Observations

MMTblueApr03/obslogs/obslog030404.txt

```
                                CCD LOG SHEET
-------------------------------------------------------------------------------
Date:      Apr  4, 2003                 Observatory: Mt.Hopkins (MMTO), MMT
Observers: Hathi, Echevarria & Jansen   Instrument:  Blue Channel Spectrograph
Program:   BBPAR Field Ellipticals      Slit:        1.0"x180"
Weather:   Clear at sunset              Tape no: X              Format: FITS

# ccd35 (2688x512, binned 1x2, spatial pixel size 2*0.28")
# 300 l/mm grism:  lamb_cen=6000 (3045--8955) Angstrom
# 600 l/mm grism:  lamb_cen=8600 (7104--10096) Angstrom

#local time  = UT - 7 h

===============================================================================
Obs.# Object     Filt Grism Texp  UTC     HA   A.M.  Comments
===============================================================================
0001 BIAS        X      300    0  23:20  0:00 1.000
0002 HeNeAr      X      300   60  23:30  0:00 1.000 spec.foc = 3.600
0003 HeNeAr      X      300   60         0:00 1.000 spec.foc = 3.800
0004 HeNeAr      X      300   60         0:00 1.000 spec.foc = 4.000
0005 HeNeAr      X      300   60         0:00 1.000 spec.foc = 3.400
0006 HeNeAr      X      300   60         0:00 1.000 spec.foc = 3.200
0007 HeNeAr      X      300   60   0:00  0:00 1.000 spec.foc = 3.500
008- FLAT        X      300    7   0:03  0:00 1.000 spec.foc @300 set to 3.550
-027 FLAT        X      300    7   0:17  0:00 1.000
0028 HeNeAr      UV-36  600   60   0:29  0:00 1.000 spec.foc = 3.550
0029 HeNeAr      UV-36  600   60   0:32  0:00 1.000 spec.foc = 3.300
0030 HeNeAr      UV-36  600   60   0:34  0:00 1.000 spec.foc = 3.700
0031 HeNeAr      UV-36  600   60   0:35  0:00 1.000 spec.foc = 3.900
0032 HeNeAr      UV-36  600   60   0:37  0:00 1.000 spec.foc = 3.800
033- FLAT        UV-36  600   15   1:03  0:00 1.000 spec.foc @600 set to 3.700
-052 FLAT        UV-36  600   15   1:17  0:00 1.000
0053 HeNeAr      X      300   60   1:20  0:00 1.000 spec.foc @300-->3.55  NB:
054- BIAS        X      300    0   1:27  0:00 1.000     there is a ~60 AA shift
-083 BIAS        X      300    0   1:40  0:00 1.000     between supposedly
084- FLAT        X      300    7   1:43  0:00 1.000     identical settings!
-088 FLAT        X      300    7   1:
# Opened up @18:50 --> Clear at sunset!
0089 SKY         X      300  1.0   1:55  0:00 1.000     ~14000 ADU
0090 SKY         X      300  1.5   1:56  0:00 1.000     ~17000 ADU
0091 SKY         X      300  1.8   1:58  0:00 1.000     ~17500 ADU
0092 SKY         X      300  2.1   1:59  0:00 1.000     ""
0093 SKY         X      300  2.4   1:59  0:00 1.000     ""
0094 SKY         X      300  2.9   2:00  0:00 1.000     ""
0095 SKY         X      300  3.4   2:00  0:00 1.000     ~18000 ADU
0096 SKY         X      300  3.7   2:01  0:00 1.000     ""
0097 SKY         X      300  4.5   2:02  0:00 1.000
0098 SKY         X      300  5.5   2:03  0:00 1.000
0099 SKY         X      300  6.7   2:03  0:00 1.000
0100 SKY         X      300    8   2:04  0:00 1.000
0101 SKY         X      300   10   2:05  0:00 1.000     ~20000 ADU
0102 SKY         X      300   14   2:05  0:00 1.000     ~22000 ADU
0103 SKY         X      300   17   2:06  0:00 1.000     ~21000 ADU
0104 SKY         X      300   21   2:07  0:00 1.000     ~20000 ADU
0105 SKY         X      300   26   2:08  0:00 1.000     ~19000 ADU
0106 SKY         X      300   32   2:09  0:00 1.000     ~18000 ADU
0107 SKY         X      300   39   2:10  0:00 1.000     ~15500 ADU
0108 SKY         X      300   50   2:12  0:00 1.000     ~14000 ADU
0109 SKY         X      300   90   2:13  0:00 1.000     ~15500 ADU
0110 SKY         X      300  180   2:15  0:00 1.000     ~18000 ADU
0111 SKY         X      300  300   2:19  0:00 1.000     ~
0112 HeNeAr      X      300   30   2:38 -0:04 1.373 for BD+75d325
0113 BD+75d325   X      300    5   2:40 -0:02 1.373
0114 BD+75d325   X      300    5   2:42 -0:00 1.373
0115 BD+75d325   X      300    5   2:43 +0:00 1.373
0116 HeNeAr      X      300   30   2:46 -2:24 1.162 for Feige34
0117 Feige34     X      300    9   2:49 -2:21 1.156 Seeing appears to be < 1"
0118 Feige34     X      300    9   2:52 -2:19 1.151
0119 Feige34     X      300    9   2:54 -2:17 1.148
0120 HeNeAr      X      300   60   3:01 -2:15 1.175 for GRB030329
0121 GRB030329   X      300  900   3:06 -2:09 1.160
```

MMTblueApr03/obslogs/obslog030404.txt

```
0122 GRB030329   X      300  900   3:24 -1:52 1.120
0123 HeNeAr      X      300   60   3:40 -1:35 1.091 for GRB030329
0124 HeNeAr      X      300   60   3:51 -3:30 1.397 for BBP1_01
0125 BBP1_01     X      300 1200   3:55 -3:25 1.377
0126 BBP1_01     X      300 1200   4:19 -3:01 1.276
0127 HeNeAr      X      300   60   4:40 -2:40 1.207 for BBP1_01
# There was a problem with the first 1200 sec exposure on BBP1_02_03: after
# exposure ended and after readout no frame was written and an error message
# returned claiming 'No write permission on File_String or something like that!
0128 HeNeAr      X      300   60   5:20 -2:01 1.110 PA=-116! BBP1_02 is bottom
0129 BBP1_02_03  X      300 1200   5:23 -1:57 1.101 one, BBP1_03 is top one.
0130 BBP1_02_03  X      300 1200   5:44 -1:37 1.069
0131 HeNeAr      X      300   60   6:06 -1:14 1.040 for BBP1_02_03
0132 No1         X      300 1200   6:17 -0:51 1.169 @#*&&! wrong grism!!!
# Changed to 600l grism @ 8600 AA + UV-36 filter, spec.focus set to 3.70
0133 HeNeAr      UV-36  600   60   6:42 -0:27 1.159 HaoJing's z=6 QSO no.1 blind
0134 No1         UV-36  600 1200   6:49       1.157 offset from FK5 star...
#                                                  Not a single object photon
#                                                  detected. Findingchart sucks
# Reset spectrograph to 300l grism @ 6000 AA , spec.focus set to 3.55
0135 BBP19_01    X      300 1200   7:30 -0:21 1.005 Exceedingly faint!!
0136 BBP19_01    X      300 1800   7:55 +0:03 1.001 ""; center may be off slit
*137 HeNeAr      X      300   60   8:29 +0:37 1.011 for BBP19_01 @#@& No header!
0138 HeNeAr      X      300   60   8:34 +0:42 1.014 for BBP19_01
0139 BBP28_01    X      300 1800   8:53 +0:05 1.069
0140 HeNeAr      X      300   60   9:28 +0:40 1.078 for BBP28_01
0141 BBP28_01    X      300 1800   9:32 +0:44 1.080
0142 HeNeAr      X      300   60  10:04 +1:16 1.103 for BBP28_01
0143 BBP28_01    X      300 1800  10:09 +1:21 1.108
0144 HeNeAr      X      300   60  10:51 +2:59 1.280 for BBP19_01; offset from
0145 BBP19_01    X      300 1800  10:53 +3:01 1.289               nearby galaxy
0146 BBP19_01    X      300 1800  11:29 +3:37 1.455
0147 HeNeAr      X      300   60  12:00 +4:08 1.662 for BBP19_01
0148 HeNeAr      X      300   30  12:07 -4:13 1.702 for BD+28d4211
0149 BD+28d4211  X      300    7  12:09 -4:11 1.682
0150 BD+28d4211  X      300    7  12:12 -4:09 1.664
0151 BD+28d4211  X      300    7  12:14 -4:07 1.650
0152 HeNeAr      X      300   30  12:17 -2:45 1.210 for CygOB2#9
0153 CygOB2no9   X      300    7  12:19 -2:43 1.204
0154 CygOB2no9   X      300    7  12:21 -2:41 1.200
0155 CygOB2no9   X      300    7  12:23 -2:40 1.196
0156 SKY         X      300  180  12:28 -3:18 1.414 ~18000 ADU; tel. pointed E
0157 SKY         X      300  120  12:32 -3:18 1.414 ~16000 ADU
0158 SKY         X      300  100  12:34 -3:18 1.414 ~25500 ADU
0159 SKY         X      300   60  12:37 -3:18 1.414 ~24000 ADU
0160 SKY         X      300   30  12:39 -3:18 1.414 ~22000 ADU
0161 SKY         X      300   25  12:40 -3:18 1.414 ~21000 ADU
0162 SKY         X      300   20  12:41 -3:18 1.414 ~21500 ADU
0163 SKY         X      300   15  12:42 -3:18 1.414 ~21000 ADU
0164 SKY         X      300   12  12:43 -3:18 1.414 ~19500 ADU
0165 SKY         X      300   10  12:44 -3:18 1.414 ~19500 ADU
0166 SKY         X      300    8  12:45 -3:18 1.414 ~19000 ADU
0167 SKY         X      300    6  12:45 -3:18 1.414 ~19250 ADU
0168 SKY         X      300    4  12:46 -3:18 1.414 ~17500 ADU
0169 SKY         X      300    2  12:47 -3:18 1.414 ~ 8000 ADU
0170 SKY         X      300    2  12:47 -3:18 1.414 ~ 9250 ADU
0171 SKY         X      300    2  12:48 -3:18 1.414 ~11000 ADU
0172 SKY         X      300    2  12:49 -3:18 1.414 ~12750 ADU
0173 SKY         X      300    2  12:50 -3:18 1.414 ~14500 ADU
0174 SKY         X      300    2  12:50 -3:18 1.414 ~17000 ADU
175- FLAT        X      300    7  13:03  0:00 1.000
-184 FLAT        X      300    7  13:09  0:00 1.000
185- BIAS        X      300    0  13:26  0:00 1.000
-214 BIAS        X      300    0  13:    0:00 1.000
215- DARK        X      300 1800  13:    0:00 1.000
-224 DARK        X      300 1800  18:47  0:00 1.000

================================================================================
Photometric night? Fair seeing ~0.8--1.5" throughout the night.
Some bands of cirrus low in the sky are visible at sun rise...
--------------------------------------------------------------------------------
```

MMTblueApr03/obslogs/obslog030405.txt

```
                                CCD LOG SHEET
------------------------------------------------------------------------------
Date:      Apr  5, 2003                 Observatory: Mt.Hopkins (MMTO), MMT
Observers: Hathi, Echevarria & Jansen   Instrument:  Blue Channel Spectrograph
Program:   BBPAR Field Ellipticals      Slit:        1.0"x180"
Weather:   Scattered clouds afternoon   Tape no: X              Format: FITS

# ccd35 (2688x512, binned 1x2, spatial pixel size 2*0.28")
# 300 l/mm grism:  lamb_cen=6000 (3045--8955) Angstrom
# 600 l/mm grism:  lamb_cen=8600 (7104--10096) Angstrom

#local time  = UT - 7 h

==============================================================================
Obs.# Object     Filt Grism Texp  UTC     HA   A.M.  Comments
==============================================================================
001- BIAS        X      300    0  23:50  0:00 1.000
-030 BIAS        X      300    0   0:06  0:00 1.000
031- FLAT        X      300    7   0:08  0:00 1.000 spec.foc @300 set to 3.55
-050 FLAT        X      300    7   0:09  0:00 1.000
0051 HeNeAr      X      300   60   0:23  0:00 1.000 spec.foc = 2.900
0052 HeNeAr      X      300   60   0:25  0:00 1.000 spec.foc = 3.100
0053 HeNeAr      X      300   60   0:27  0:00 1.000 spec.foc = 3.300
0054 HeNeAr      X      300   60   0:29  0:00 1.000 spec.foc = 3.500
0055 HeNeAr      X      300   60   0:31  0:00 1.000 spec.foc = 3.700
0056 HeNeAr      X      300   60   0:33  0:00 1.000 spec.foc = 3.900
0057 HeNeAr      X      300   60   0:35  0:00 1.000 spec.foc = 4.100
0058 HeNeAr      X      300   60   0:37  0:00 1.000 spec.foc = 4.100
0059 HeNeAr      X      300   60   0:43  0:00 1.000 spec.foc = 3.550
0060 HeNeAr      X      300   60   0:51  0:00 1.000 spec.foc = 3.400 <-- best
0061 HeNeAr      UV-36  600   60   1:03  0:00 1.000 spec.foc = 3.000
0062 HeNeAr      UV-36  600   60   1:06  0:00 1.000 spec.foc = 3.200
0063 HeNeAr      UV-36  600   60   1:08  0:00 1.000 spec.foc = 3.300
0064 HeNeAr      UV-36  600   60   1:09  0:00 1.000 spec.foc = 3.400 <-- best
0065 HeNeAr      UV-36  600   60   1:11  0:00 1.000 spec.foc = 3.500
0066 HeNeAr      UV-36  600   60   1:13  0:00 1.000 spec.foc = 3.600
0067 HeNeAr      UV-36  600   60   1:15  0:00 1.000 spec.foc = 3.700
0068 HeNeAr      UV-36  600   60   1:17  0:00 1.000 spec.foc = 3.800
0069 HeNeAr      UV-36  600   60   1:18  0:00 1.000 spec.foc = 4.000
070- FLAT        UV-36  600   12   1:31  0:00 1.000 spec.foc @600 set to 3.40
-089 FLAT        UV-36  600   12   1:46  0:00 1.000
0090 HeNeAr      X      300   60   1:49  0:00 1.000 @300 l/mm
091- BIAS        X      300    0   1:51  0:00 1.000
-093 BIAS        X      300    0   1:52  0:00 1.000
# Opened up @18:50 --> Clear at sunset
094- SKY         X      300  1.0   1:53  0:00 1.000 ~29000 some of the first few
-121 SKY         X      300  300   2:25  0:00 1.000        may be overexposed...
0122 HeNeAr      X      300   30   2:43 +0:04 1.374 for BD+75d325
0123 BD+75d325   X      300    5   2:45 +0:06 1.374 seeing seems OK for 1" slit
0124 BD+75d325   X      300    5   2:47 +0:08 1.374
0125 BD+75d325   X      300    5   2:48 +0:10 1.374
0126 HeNeAr      X      300   30   2:51 -2:15 1.143 for Feige34
0127 Feige34     X      300    9   2:54 -2:12 1.138
0128 Feige34     X      300    9   2:56 -2:11 1.135
0129 Feige34     X      300    9   2:57 -2:09 1.132
0130 GRB030329   X      300 1200   3:07 -2:05 1.150
0131 HeNeAr      X      300   60   3:28 -1:44 1.106 for GRB030329
0132 GRB030329   X      300 1200   3:33 -1:38 1.097
0133 BBP22_02    X      300 1800   4:10 -0:51 1.027 @#@#*&! nothing in the slit
0134 BBP10_01    X      300 1800   5:04 -0:40 1.042 switched to 1.5" slit.
0135 HeNeAr      X      300   60   5:35 -0:08 1.029 for BBP10_01
0136 BBP10_01    X      300 1800   5:37 -0:06 1.029
0137 BBP18_01    X      300 1800   6:29 -0:22 1.061
0138 HeNeAr      X      300   60   7:02 +0:09 1.057 for BBP18_01
0139 BBP18_01    X      300 1800   7:08 +0:15 1.058
0140 BBP06_01    X      300 1800   8:02 +2:18 1.152
0141 HeNeAr      X      300   60   8:35 +2:51 1.226 for BBP06_01
0142 BBP06_01    X      300 1800   8:40 +2:56 1.244
0143 HeNeAr      X      300   60   9:11 +3:27 1.346 for BBP06_01
# Changed to 600l grism @ 8600 AA + UV-36 filter for HaoJing's z=6 QSO no.1
0144 No1         UV-36  600 1800   9:21 +2:16 1.257 blind move after point.check
```

MMTblueApr03/obslogs/obslog030405.txt

```
0145 HeNeAr      UV-36  600   60   9:59 +2:54 1.326 @!&*@ Not a single photon.
# Reset spectrograph to 300l/mm grism @ 6000 AA
0146 BBP06_01    X      300 1800  10:12 +4:28 1.654 NB: wavelen. offset possible
0147 HeNeAr      X      300   60  10:43 +4:59 1.901 for BBP06_01
0148 BBP28_01    X      300 1800  10:49 +2:05 1.164
0149 HeNeAr      X      300   60  11:21 +2:37 1.224 for BBP28_01
0150 BBP28_01    X      300 1800  11:25 +2:41 1.232
0151 HeNeAr      X      300   60  11:56 +3:12 1.311 for BBP28_01
0152 HeNeAr      X      300   30  12:02 -4:14 1.712 for BD+28d4211; 1.5" slit
0153 BD+28d4211  X      300    7  12:05 -4:11 1.686
0154 BD+28d4211  X      300    7  12:06 -4:10 1.678
0155 BD+28d4211  X      300    7  12:07 -4:09 1.672
0156 BD+28d4211  X      300    7  12:08 -4:08 1.660 changed to 1.0" slit
0157 BD+28d4211  X      300    7  12:09 -4:07 1.653
0158 BD+28d4211  X      300    7  12:10 -4:06 1.646
0159 HeNeAr      X      300   30  12:11 -4:05 1.637 for BD+28d4211; 1.0" slit
0160 HeNeAr      X      300   30  12:14 -2:45 1.208 for CygOB2#9; 1.0" slit
0161 CygOB2no9   X      300    7  12:16 -2:42 1.203
0162 CygOB2no9   X      300    7  12:17 -2:41 1.199
0163 CygOB2no9   X      300    7  12:18 -2:40 1.197
0164 CygOB2no9   X      300    7  12:19 -2:39 1.194 changed back to 1.5" slit
0165 CygOB2no9   X      300    7  12:20 -2:38 1.192
0166 CygOB2no9   X      300    7  12:21 -2:37 1.190
0167 HeNeAr      X      300   30  12:22 -2:36 1.187 for CygOB2#9; 1.5" slit
168- SKY         X      300  180  12:25 -3:18 1.414 tel. pointed E; 1.5" slit
-193 SKY         X      300  1.0  12:52 -3:18 1.414
# Stowed and closed down telescope @ 5:54 A.M.
194- FLAT        X      300    7  13:01  0:00 1.000
-213 FLAT        X      300    7  13:11  0:00 1.000
214- BIAS        X      300    0  13:15  0:00 1.000
-243 BIAS        X      300    0  13:25  0:00 1.000
244- DARK        X      300 1800  13:29  0:00 1.000
-246 DARK        X      300 1800  16:00  0:00 1.000

================================================================================
Clear night. Fair seeing ~0.8--1.5" throughout the night again.
Also clear at sun rise.  Photometric all night?
--------------------------------------------------------------------------------
```

9.3 Processing of the calibration frames

9.3.1 CCD layout, bad pixels, read-noise and gain

• Determine the CCD gain and readnoise, and the location of the overscan strip. Also determine what portion of the full CCD frame is illuminated. Verify `DISPAXIS` is defined in the FITS headers.

```
cl>  set stdimage=imt27    ⟵ twice; 3104×512 pixel display buffer
cl>  stsdas
cl>  imred
cl>  bias
cl>  rjtools
```

△ Note: data taken at the MMT with the Blue and Red spectrographs is written to FITS in *signed integer* format, where values > 32768 are represented by values between -32767 and -1. To get floating point pixel values between 0. and 65535., we need to use a two-step conversion. First, we convert from signed integer to unsigned integer, then from unsigned integer to real:

```
cl>  !\ls -1 03040?/*.fits > all.lis
```

Removed corrupted frame "a0137.fits" from "all.lis", then:

```
cl>  chpixtype ("@all.lis", "@all.lis", "ushort", oldpixtype="all")
cl>  chpixtype ("@all.lis", "@all.lis", "real", oldpixtype="all")
```

Fix some header keywords; add two keywords identifying the grisms used (see § 9.2):

```
cl>  unlearn hedit ; hedit.verify=no ; hedit.update=yes
cl>  hedit ("@all.lis", "observat", "mmto", add-)
cl>  hedit ("@all.lis", "telescop", "mmt", add-)
cl>  hedit ("@all.lis", "instrume", "Blue Channel Spectrograph", add-)

cl>  !range -pre "030404/a" -suf ".fits" -d 4 1 - 27 53 - 132 135 - 224 \
>>>        > all300.lis
cl>  !range -pre "030405/b" -suf ".fits" -d 4 1 - 60 90 - 143 146 - 246 \
>>>        >> all300.lis
cl>  !range -pre "030404/a" -suf ".fits" -d 4 28 - 52 133 134 > all600.lis
cl>  !range -pre "030405/b" -suf ".fits" -d 4 61 - 89 144 145 >> all600.lis

cl>  hedit ("@all300.lis", "grism", "300l/mm", add+)
cl>  hedit ("@all300.lis", "lamcen", 6000., add+)
cl>  hedit ("@all600.lis", "grism", "600l/mm", add+)
cl>  hedit ("@all600.lis", "lamcen", 8600., add+)
```

Next, print an inventory of the unique 'object' name+grism combinations:

```
cl>  hselect @all.lis object,grism yes | sort | unique
# 300 l/mm grism observations:
BBP06_01          "300l/mm"       <-- Science targets: field elliptical galaxies,
BBP10_01          "300l/mm"           selected from the "B-Band Parallel" survey
BBP18_01          "300l/mm"           ""
BBP19_01          "300l/mm"           ""
BBP1_01           "300l/mm"           ""
BBP1_0203         "300l/mm"           ""
BBP22_02          "300l/mm"           ""
BBP28_01          "300l/mm"           ""
BD+284211         "300l/mm"       <-- Blue and red spectrophotometric standard stars
BD+75D325         "300l/mm"           (Massey et al. 1988; Oke 1990 [BD+75d325])
BD+75d325         "300l/mm"           ""
CygOB2No9         "300l/mm"           ""
CygOB2no9         "300l/mm"           ""
Feige34           "300l/mm"           ""
GRB030329         "300l/mm"       <-- Target of Opportunity: bright GRB afterglow!
HeNeAr            "300l/mm"       <-- Wavelength calibration comparison-lamp spectra
dark              "300l/mm"
flat              "300l/mm"
sky               "300l/mm"
zero              "300l/mm"
# 600 l/mm grism observations:
No1               "300l/mm"       <-- HaoJing's z~6 candidate (wrong grism)
No1               "600l/mm"       <-- HaoJing's z~6 candidate (correct grism)
HeNeAr            "600l/mm"       <-- Wavelength calibration comparison-lamp spectra
flat              "600l/mm"
```

And create lists of all frames per exposure type:

```
cl>  hselect @all.lis $I 'object=="zero"' > bias.lis
cl>  hselect @all.lis $I 'object=="dark"' > dark.lis
cl>  hselect @all.lis $I 'object=="flat"' > flat.lis
cl>  !fgrep -f flat.lis all600.lis > flat600.lis
cl>  !fgrep -f flat.lis all300.lis > flat300.lis
cl>  hselect @all.lis $I 'object=="sky"' > sky.lis
cl>  !\cp sky.lis sky300.lis      <-- there are no twilight exposures for 600l/mm
cl>  hselect @all.lis $I 'object!="zero"&&object!="dark"' > tmp.lis
cl>  hselect @tmp.lis $I 'object!="flat"&&object!="sky"' > sci.lis
cl>  !fgrep -f sci.lis all300.lis > sci300.lis
```

```
cl>  !fgrep -f sci.lis all600.lis > sci600.lis
cl>  delete tmp.lis yes
```

Construct a bad pixel/bad column list using the median average of a few twilight sky exposures (because of the difficulty of displaying a frame with a large dynamic range, we will fit out the large-scale signal):

```
cl>  head sky.lis nlines=7 > tmp.lis
cl>  unlearn imcombine ; unlearn fit1d
cl>  imcombine ("@tmp.lis", "tmp.fits", combine="average", reject="avsigclip",
>>>       scale="mean", zero="none", weight="mean", statsec="[5:2688,2:255]")
cl>  fit1d ("tmp.fits", "_tmp.fits", "difference", axis=1, interactive-,
>>>       sample="*", naverage=-7, function="spline3", order=23,
>>>       low_reject=3., high_reject=3., niterate=3, grow=1.)
cl>  fit1d ("_tmp.fits", "_tmp.fits", "difference", axis=2, interactive-,
>>>       sample="1:255", naverage=-31, function="legendre", order=3,
>>>       low_reject=3., high_reject=3., niterate=3, grow=0.)

cl>  unlearn getregion
cl>  getregion ("_tmp.fits", "_tmp.reg", format="basic", chkbound+, verbose+)
GETREGION:  IRAF2.12.2  raj@andromeda  Nov 24 17:28:29 2006
   image = _tmp.fits [2708 x 276]
   output= _tmp.reg  (format="basic")
Displaying image "_tmp.fits":   z1=-525.65 z2=264.06
   Wait for the cursor cross to appear in the image display area,
   then mark all regions by hitting "b" once at each of two dia-
   gonally oposite corners to mark an arbitrary rectangular region,
   or "c" once to mark a single pixel.  Hit "q" to quit.
Type b again to draw box.
Type b again to draw box.
   ...
# GETREGION Finished
```

Overlay the defined regions onto the image to verify:

```
cl>  unlearn tvmarkall
cl>  tvmarkall ("_tmp.fits", gal-, star-, cosmic-, del-, region+,
>>>       regfil="_tmp.reg", regclr=205)

cl>  printf("# MMT Blue Channel Spectrograph\n", > "badpix.ccd35")
cl>  printf("# STA0520A, 2688[+20]x512 15mu-pixel back-", >> "badpix.ccd35")
cl>  printf("illum.  CCD, binned 1x2, untrimmed\n", >> "badpix.ccd35")
cl>  !cat _tmp.reg >> badpix.ccd35
```

```
cl>  delete tmp.lis,_tmp.fits,_tmp.reg yes verify-
```

Determine the location of the virtual overscan strip (actually, for this particular CCD it's quite obvious that only [2687:2708,1:256] would qualify, since the overscan strip along the top clear shows charge trails):

```
cl>  implot tmp.fits
:l 1 255
:x 2668 2708
:y 350 1775        ⟶ [2693:2708,1:255]
:x 1 30
:y 0 14500
q
```

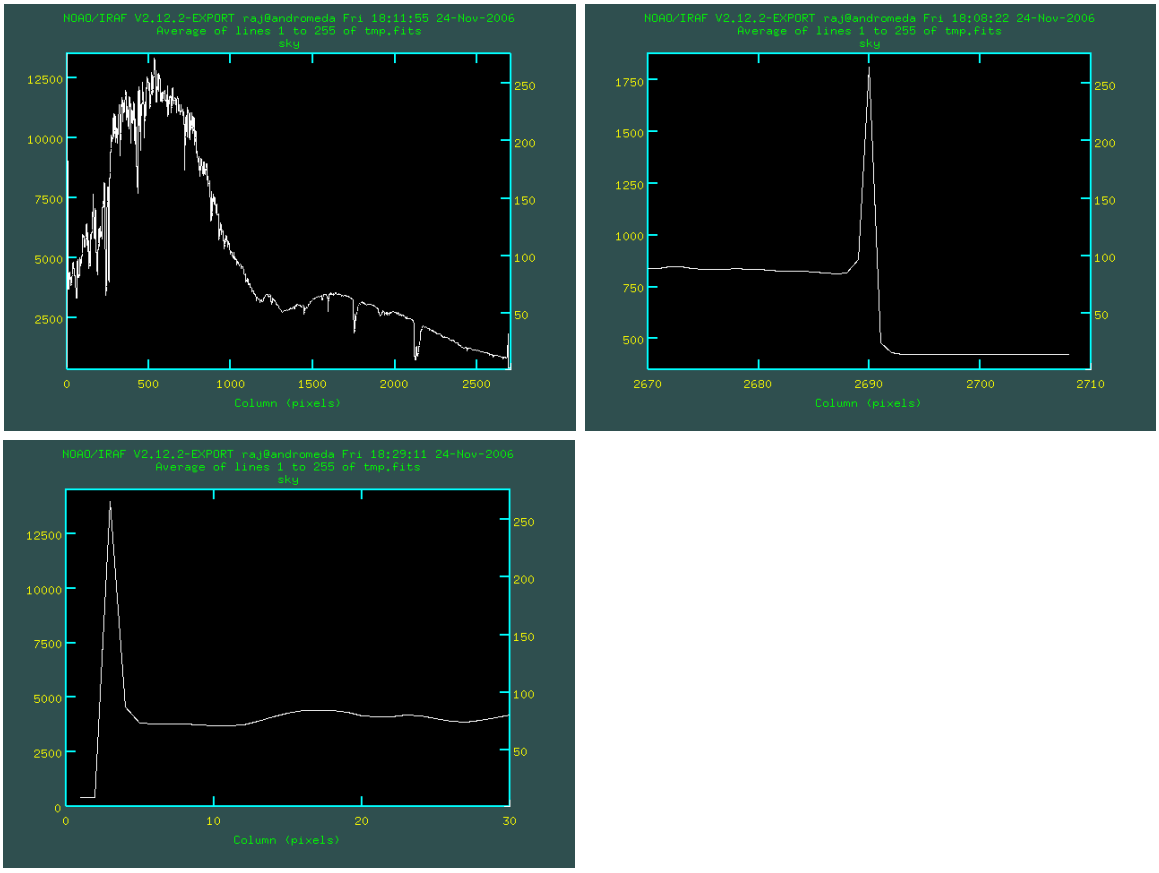

Figure 67: Screens shot of `implot` graphs of a temporary image created by averaging 7 twilight sky flats. (*a*) average of lines 1–255, spanning the full 2708 pixel column range; (*b*) detail of the final ∼40 columns, including the overscan strip; (*c*) detail of the first ∼30 columns. Note the amplifier turn-on/off spikes.

Note, that there are strong amplifier turn-on/off spikes in both columns 1–4 and 2689–2692! Also, the level in the overscan strip is not exactly flat, but falls off somewhat toward the final column. The slope is shallowest in the final 8 pixels, i.e., [2701:2708,1:255].

```
cl>  !head -20 bias.lis | sed 's/$/[2701:2708,1:255]/g' > tmp.lis
cl>  unlearn imstat
cl>  imstat @tmp.lis fields="image,mean,stddev,min,max"
#              IMAGE                      MEAN   STDDEV      MIN      MAX
 030404/a0001.fits[2701:2708,1:255]    482.6    2.264      476.     494.
 030404/a0054.fits[2701:2708,1:255]    484.1    2.268      477.     492.
 030404/a0055.fits[2701:2708,1:255]    482.7    2.197      474.     492.
 030404/a0056.fits[2701:2708,1:255]    482.4    2.125      475.     489.
 030404/a0057.fits[2701:2708,1:255]    482.3    2.215      476.     491.
             ...                         ...      ...       ...      ...
```

From which a mean of ~ 482.56 ADU and an rms ~ 2.17 ADU is found, and min/max values that are consistent with the mean $\pm 4\,\sigma$. This must, therefore, indeed be the virtual overscan strip. So, in the following, we shall adopt:

`TRIMSEC = STATSEC = '[5:2688,1:255]'` and `BIASSEC = '[2701:2708,1:255]'`

Measure the CCD gain and read-noise using pairs of flats and biases (Janesick's method):

```
cl>  obsutil
cl>  !grep '030405' flat300.lis | wc -l
cl>        40
cl>  !grep '030405' flat300.lis | head -39 | sed 's/.fits//g' > _f1
cl>  !grep '030405' flat300.lis | tail -39 | sed 's/.fits//g' > _f2
cl>  !grep '030405' bias.lis | head -39 | sed 's/.fits//g' > _b1
cl>  !grep '030405' bias.lis | head -40 | tail -39 | sed 's/.fits//g' > _b2
```

(To check:)

```
cl>  !paste _f1 _f2 _b1 _b2
```

Create a little template script:

```
cl>  !echo 'findgain ("$1", "$2", "$3",' > getccdpar.tem
cl>  !echo ' "$4", section="[5:2688,1:255]", verbose-)' >> getccdpar.tem
cl>  type getccdpar.tem
findgain ("$1", "$2", "$3",
        "$4", section="[5:2688,1:255]", verbose-)
cl>  !narg getccdpar.tem _f1 '$1' _f2 '$2' _b1 '$3' _b2 '$4' > getccdpar.cl
```

(Check:)

```
cl>  type getccdpar.cl
cl>  cl < getccdpar.cl
  1.16    2.73
  1.16    2.72
  1.16    2.71
  1.16    2.70
  1.16    2.69
   ...     ...
```

For night '03/04/05, we find $\overline{\mathcal{G}} = 1.160 \pm 0.005\, e^-$/ADU and $\overline{\mathcal{R}} = 2.69 \pm 0.05\, e^-$. Similarly, for night '03/04/04, $\overline{\mathcal{G}} = 1.160 \pm 0.005\, e^-$/ADU and $\overline{\mathcal{R}} = 2.71 \pm 0.01\, e^-$. For the two nights combined: $\mathcal{G} = 1.160 \pm 0.005\, e^-$/ADU and $\mathcal{R} = 2.710 \pm 0.035\, e^-$. The CCD gain is, therefore, $\sim$3.54% lower than the advertised $1.2\, e^-$/ADU, while the readnoise is $\sim$4.1% higher than the advertised $2.6\, e^-$.

```
cl>  delete _f1,_f2,_b1,_b2 yes verify-
cl>  bye        ⟵   (to unload the obsutil package)
```

9.3.2 Bias frames

• Process all BIAS frames (interpolate over bad pixels and subtract the overscan leven) and combine them into a "MASTERBIAS" frame.

Inspect all bias frames and run `imstat` to check for deviant frames:

```
cl>  imstat @bias.lis fields="image,midpt,stddev,min,max" > biasstat.lis
cl>  imexamine @bias.lis
```

There seems to have been a probem with frames "a0185" through "a0196" (most likely because of bright lights in the dome). Deleted these affected frames from files "bias.lis", "all.lis" and "all300.lis".

Add a line to the image headers indicating the start of the processing log entries:

```
cl>  type hdrlog_init.tem
gdate() ; sysinfo() ; unlearn hedit
hedit.add=yes ; hedit.verify=no ; hedit.show=no
s1=" CCD image processing -- "//sysinfo.username//"@"//sysinfo.host//" "
i=66-int(0.5*(66-strlen(s1)))
cpadd (s1, pchar="-", nchar=i, paddleft+, verbose-)
cpadd (cpadd.cpadd, pchar="-", nchar=66, paddleft-, verbose-)
hedit ("@%.lis", "HISTORY", cpadd.cpadd)

cl>  !sed 's:%:bias:g' hdrlog_init.tem > hdrlog_init.cl
cl>  cl < hdrlog_init.cl
```

Interpolate over the bad pixels identified:

```
cl>  unlearn fixpix
cl>  fixpix ("@bias.lis", "badpix.ccd35", verbose+, >> "ccdproc.log")
```

Update the headers:

```
cl>  type hdrlog_fxpx.tem
gdate()
hedit ("@%.lis", "FIXPIX", "COMPLETE")
hedit ("@%.lis", "FXPXFILE", "/data1/raj/MMTblueApr03/badpix.ccd35")
hedit ("@%.lis", "FXPXDATE", gdate.fdate)
cl>  !sed 's:%:bias:g' hdrlog_fxpx.tem > hdrlog_fxpx.cl
cl>  cl < hdrlog_fxpx.cl
```

Subtract the bias level as measured in the virtual overscan strip. First, measure that level and record it in a log file for later use:

```
cl>  !sed 's/$/[2701:2708,1:255]/g' bias.lis > _bias.lis
cl>  imstat ("@_bias.lis", fields="image,midpt", lower=INDEF, upper=INDEF,
>>>      nclip=3, lsigma=3., usigma=3., format-, cache-, >> "overscan.dat")
cl>  delete _bias.lis yes verify-
```

Now, fit and remove that overscan level using task `colbias`:

```
cl>  unlearn colbias
cl>  colbias ("@bias.lis", "@bias.lis", bias="[2701:2708,1:255]", trim="",
>>>      median+, interactive-, function="legendre", order=1, low_reject=3.,
>>>      high_reject=3., niterate=3, logfiles="ccdproc.log")
```

Update the headers (write the following two template scripts):

```
cl>  type hdrlog_ovsc.tem
gdate()
hedit ("@%.lis", "BIASSEC,TRIMSEC,DATASEC,CCDSEC,ORIGSEC", add-, del+)
hedit ("@%.lis", "OVERSCAN", "COMPLETE")
hedit ("@%.lis", "ORIGSEC", "[1:2708,1:276]")
hedit ("@%.lis", "BIASSEC", "[2701:2708,1:255]")
hedit ("@%.lis", "TRIMSEC", "[5:2688,1:255]")
hedit ("@%.lis", "IMAGSEC", "[1:2684,1:255]")
!narg addovscmean.tem %.lis '$1' > addovscmean.cl
cl < addovscmean.cl
hedit ("@%.lis", "OVSCDATE", gdate.fdate)
```

```
cl>  type addovscmean.tem
match "$1" "overscan.dat" | fields "-" 2 | scanf("%8f", x)
hedit ("$1", "OVSCMEAN", x)
cl>  !sed 's:%:bias:g' hdrlog_ovsc.tem > hdrlog_ovsc.cl
cl>  cl < hdrlog_ovsc.cl
```

But we didn't actually trim the bias frames (we will trim the frames for all other data types — hence the design of the above template script), so:

```
cl>  hedit ("@bias.lis", "TRIMSEC", "[1:2708,1:276]", add-)
cl>  hedit ("@bias.lis", "IMAGSEC", "[1:2708,1:276]", add-)
```

Average the overscan-subtracted biases:

```
cl>  unlearn imcombine
cl>  imcombine ("@bias.lis", "BIAS.fits", logfile="ccdproc.log",
>>>       combine="average", reject="avsigclip", scale="none", zero="none",
>>>       weight="none", statsec="[5:2688,1:255]", lthreshold=INDEF,
>>>       hthreshold=65535.)
```

Update the header:

```
cl>  gdate() ; hselect BIAS.fits ncombine yes | scanf("%d", i)
cl>  print ("Combined ",i,"individual bias frames.")
Combined 112 individual bias frames.
cl>  hedit ("BIAS.fits", "NCOMBINE", add-, del+)
cl>  hedit ("BIAS.fits", "BIASCOMB", "COMPLETE")
cl>  hedit ("BIAS.fits", "NCOMBINE", i)
cl>  hedit ("BIAS.fits", "BSCBDATE", gdate.fdate)
```

Verify, that the level in the overscan strip of the combined frame is 0. If not, subtract the residual level:

```
cl>  imstat BIAS.fits[2701:2708,1:255] fields="mean" format- | scan (s1)
cl>  print ("Mean level of overscan strip in BIAS.fits is "//s1)
Mean level of overscan strip in BIAS.fits is 0.3636922
cl>  unlearn imarith
cl>  imarith ("BIAS.fits", "-", s1, "MASTERBIAS.fits", title="MASTERBIAS",
>>>       pixtype="real", calctype="real", verbose+, >> "ccdproc.log")
```

Inspect the resulting "MASTERBIAS" frame and evaluate the observed noise:

```
cl>  display MASTERBIAS.fits 1 zscale- zrange- z1=-2.  z2=4.
```

```
cl>  s1="MASTERBIAS.fits[5:2688,75:255]"
cl>  imstat (s1, fields="image,mean,midpt,stddev,min,max", lower=INDEF,
>>>       upper=INDEF, nclip=3, lsigma=3., usigma=3.)
#              IMAGE                 MEAN    MIDPT    STDDEV     MIN       MAX
 MASTERBIAS.fits[5:2688,75:255]    0.5989   0.5998    0.2421  -0.1312    1.3260
cl>  implot MASTERBIAS.fits
:c 5 2688
q
```

The level in the lower rows of the MASTERBIAS frame is systematically higher than that in the higher rows: the level falls off rapidly from ∼1.6 ADU in the first few lines to ∼0.60 ADU at lines 75 and higher (see Fig. 68). This signal is likely not a true bias signal, but may be either a transient effect (after strong illumination), or due to light leaking onto the CCD during read-out.
Assuming the signal in the first ∼75 lines is spurious in nature, correct the MASTERBIAS frame. Since, its magnitude appears to have no dependence on column number, we can abuse task `illumination` to fit a single function to all pixels as a function of line number, regardless of column number. We have to beware, though, that `illumination` normalizes its resulting frame:

```
cl>  !\mv MASTERBIAS.fits _MASTERBIAS.fits
cl>  twodspec
cl>  longslit
cl>  illumination ("_MASTERBIAS.fits", "tmp.fits", interactive-, bins="",
>>>       nbins=1, sample="*", naverage=-5, function="spline3", order=6,
>>>       low_reject=3., high_reject=3., niterate=3, grow=0.)
```

In the original MASTERBIAS frame we found a mean value in lines >75 of 0.5998, while now it is ∼1, so:

```
cl>  imcalc ("_MASTERBIAS.fits,tmp.fits", "MASTERBIAS.fits",
>>>       "im1 - im2*0.5998 + 0.5998", pixtype="real")
cl>  imstat MASTERBIAS.fits[2701:2708,1:255] fields="mean" format- | scan (s1)
cl>  print ("Mean level of overscan strip in MASTERBIAS.fits is "//s1)
Mean level of overscan strip in MASTERBIAS.fits is -0.08864927
cl>  imarith ("MASTERBIAS", "-", s1, "MASTERBIAS.fits", title="MASTERBIAS",
>>>       pixtype="real", calctype="real", verbose+, >> "ccdproc.log")
```

Inspect the resulting corrected MASTERBIAS frame and evaluate the observed noise:

```
cl>  display MASTERBIAS.fits 1 zscale- zrange- z1=-2.  z2=4.
```

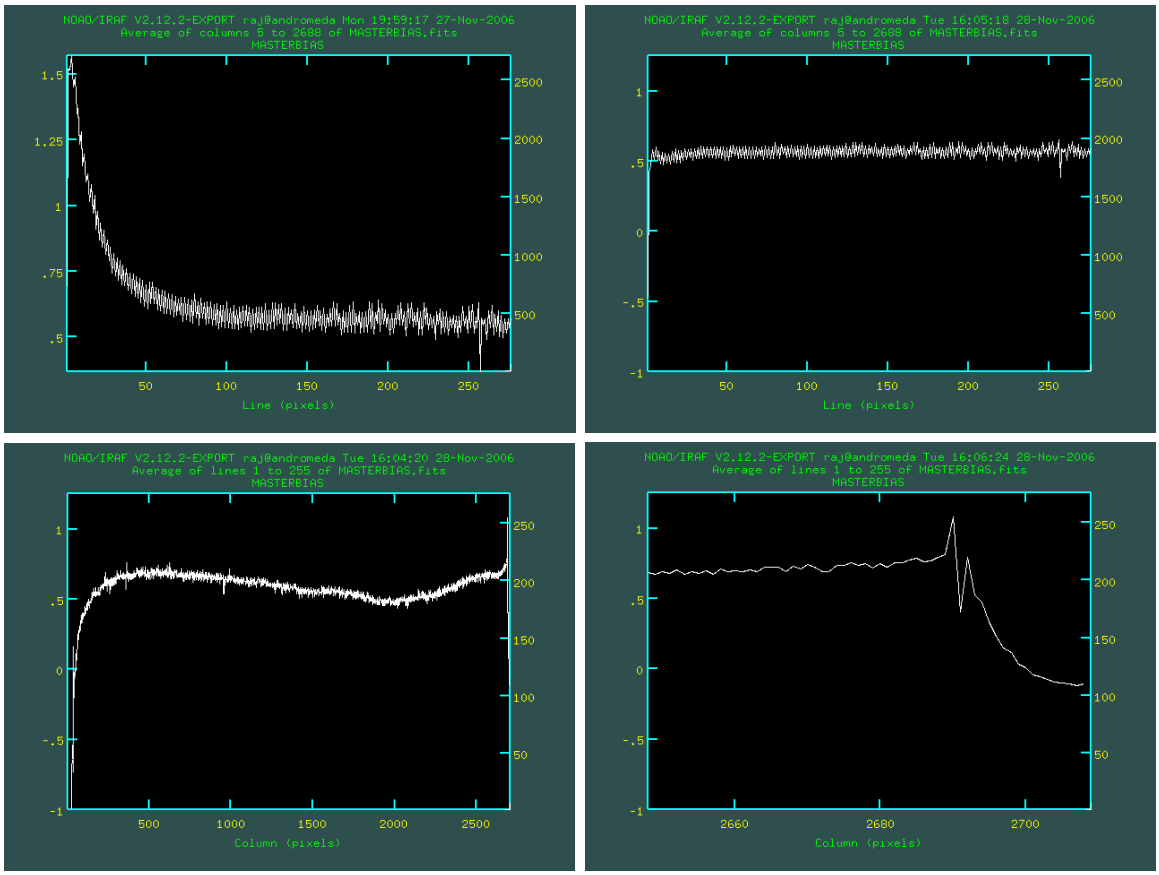

Figure 68: (a) Average of columns 5–2688 of the original MASTERBIAS frame. The "bias" level shows a strong dependence on line number in the first ∼75 lines, or so. (b) Same as (a) but for the corrected MASTERBIAS frame. Only the genuine odd-even pattern remains. (c) Average of lines 1–255 in the corrected MASTERBIAS frame, showing that the bias level varies with column number. (d) Detail of the final 60 columns of (c), illustrating the roll-off in the virtual overscan strip.

```
cl>  implot MASTERBIAS.fits
:c 5 2688
:y -1.00 1.25
:l 1 255
:x 2648 2708
q
```

The level in the image region of the MASTERBIAS frame is not constant as a function of column number: there is a strong roll-on in the first few hundred columns and a slow variation with column number thereafter. The level is reaches a local minimum of ∼0.50 near column 2000. There is also a roll-off within the nominal

bias strip, [2689:2708,*], some of which is still noticable within our adopted overscan, [2701:2708,1:255]. There is an offset between the level in the overscan strip and that in the image region of the CCD, demonstrating the need for a 2-D correction for the structure in the bias level.

```
cl>  s1="MASTERBIAS[5:2688,1:255],MASTERBIAS[1900:2100,1:255],"
cl>  s1=s1//",MASTERBIAS[2701:2708,1:255]"
cl>  imstat (s1, fields="image,mean,midpt,stddev,min,max", lower=INDEF,
>>>       upper=INDEF, nclip=3, lsigma=3., usigma=3.)
#               IMAGE                  MEAN    MIDPT    STDDEV     MIN       MAX
 MASTERBIAS[5:2688,1:255]            0.6788    0.6761    0.2435   -0.0560    1.4126
 MASTERBIAS[1900:2100,1:255]         0.5701    0.5595    0.2224   -0.0967    1.2366
 MASTERBIAS[2701:2708,1:255]      -6.69E-4   -0.0095    0.1994   -0.5967    0.5887
```

Assuming pure Gaussian read-noise, the expected noise in the absence of genuine variations in the column-to-column bias level would be: $2.710\,[e^-]/1.160\,[e^-/\mathrm{ADU}]/\sqrt{112}$ $= 0.220751\,[\mathrm{ADU}]$. This is indeed close (within $\sim$10%) of what we measure. The noise level is $\sim 12\times$ smaller than the read-noise, so the subtraction of the MASTERBIAS frame in subsequent processing steps will not add appreciable noise.

Since we no longer need the individual bias frames at this point:

```
cl>  imdelete @bias.lis yes verify- default+
```

9.3.3 Dark frames

• Process all DARK frames (interpolate over bad pixels, subtract the MASTERBIAS frame, subtract the overscan level, and trim the frames) and combine them into a "MASTERDARK" frame (if S/N is sufficient; if not, try and measure the bulk dark rate); determine whether there are any warm/hot pixels.

Inspect all dark frames and run `imstat` to check for deviant frame:

```
cl>  imstat @dark.lis fields="image,midpt,stddev,min,max" > darkstat.lis
cl>  imexamine @dark.lis allframes- nframes=1
```

There seems to have been a problem with frames "a0215" and "a0217" (likely because of lights in the dome). Deleted these affected frames from files "dark.lis", "all.lis", and "all300.lis".

Add a line to the image headers indicating the start of the processing log entries:

```
cl>  !sed 's:%:dark:g' hdrlog_init.tem > hdrlog_init.cl
cl>  cl < hdrlog_init.cl
```

Interpolate over the bad pixels and log relevant info to the headers:

```
cl>  fixpix ("@dark.lis", "badpix.ccd35", verbose+, >> "ccdproc.log")

cl>  !sed 's:%:dark:g' hdrlog_fxpx.tem > hdrlog_fxpx.cl
cl>  cl < hdrlog_fxpx.cl
```

Subtract the MASTERBIAS frame to correct for structure in the bias level:

```
cl>  imarith ("@dark.lis", "-", "MASTERBIAS.fits", "@dark.lis", title="",
>>>       pixtype="real", calctype="real", verbose+, >> "ccdproc.log")
```

Log relevant information to the headers:

```
cl>  type hdrlog_zrcr.tem
gdate()
hedit ("@%.lis", "ZEROCOR", "COMPLETE")
hedit ("@%.lis", "ZRCRIMAG", "/data1/raj/MMTblueApr03/MASTERBIAS.fits")
hedit ("@%.lis", "ZRCRDATE", gdate.fdate)

cl>  !sed 's:%:dark:g' hdrlog_zrcr.tem > hdrlog_zrcr.cl
cl>  cl < hdrlog_zrcr.cl
```

Subtract the bias level as measured in the virtual overscan strip. First, measure and record that level so we can update the headers:

```
cl>  !sed 's/$/[2701:2708,1:255]/g' dark.lis > _dark.lis
cl>  imstat ("@_dark.lis", fields="image,midpt", lower=INDEF, upper=INDEF,
>>>       nclip=3, lsigma=3., usigma=3., format-, >> "overscan.dat")
cl>  delete _dark.lis yes verify-
```

Now, fit and remove the overscan level from all dark frames:

```
cl>  colbias ("@dark.lis", "@dark.lis", bias="[2701:2708,1:255]", median+,
>>>       trim="[5:2688,1:255]", interactive-, function="legendre", order=1,
>>>       low_reject=3., high_reject=3., niterate=3, logfiles="ccdproc.log")

cl>  !sed 's:%:dark:g' hdrlog_ovsc.tem > hdrlog_ovsc.cl
cl>  cl < hdrlog_ovsc.cl
```

Combine the overscan-subtracted dark frames into a single frame "DARK.fits". Note, that since the elapsed dark time is 1800 sec for all darks, we will disallow multiplicative scaling (it *should* be allowed if one has darks with different integration times):

```
cl>  unlearn imcombine
cl>  imcombine ("@dark.lis", "DARK.fits", logfile="ccdproc.log",
>>>       combine="average", reject="avsigclip", scale="none", zero="none",
>>>       weight="none", statsec="[5:2688,1:255]", lthreshold=INDEF,
>>>       hthreshold=65535., grow=1.5)
```

Update the header:

```
cl>  gdate() ; hselect DARK.fits ncombine yes | scanf("%d", i)
cl>  print ("Combined ",i,"individual dark frames.")
Combined 13 individual dark frames.
cl>  hedit ("DARK.fits", "NCOMBINE,IMCMB*", add-, del+)
cl>  hedit ("DARK.fits", "DARKCOMB", "COMPLETE")
cl>  hedit ("DARK.fits", "NCOMBINE", i)
cl>  hedit ("DARK.fits", "DKCBDATE", gdate.fdate)
```

Inspect the combined dark frame:

```
cl>  display DARK.fits 1 zs- zr- z1=-2.  z2=4.
cl>  implot DARK.fits
:c 1 2684
:l 1 255
q
```

```
cl>  s1="DARK.fits"
cl>  s2=s1//","//s1//"[1:2684,1:85],"//s1//"[1:2684,86:170],
cl>  s2=s2//s1//"[1:2684,171:255]"
cl>  imstat (s2, fields="image,mean,midpt,stddev,min,max", lower=INDEF,
>>>       upper=INDEF, nclip=3., lsigma=3., usigma=3.)
#           IMAGE                    MEAN    MIDPT    STDDEV      MIN       MAX
 DARK.fits                        1.11130  1.12447  0.92729  -1.67070   3.89349
 DARK.fits[1:2684,1:85]           1.56229  1.59168  0.86617  -1.04690   4.16626
 DARK.fits[1:2684,86:170]         1.29481  1.29060  0.76310  -0.99457   3.58433
 DARK.fits[1:2684,171:255]        0.49659  0.49198  0.76923  -1.81107   2.80438
```

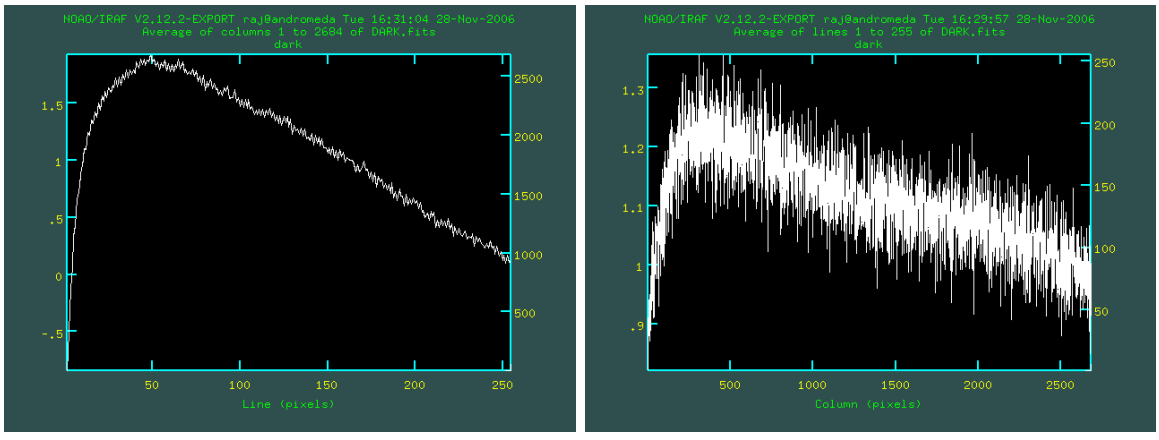

Figure 69: Structure along columns and along lines in the combined DARK frame.

The largest possible error we could make by not correcting for dark current at all would be $\sim 1.9\,[\mathrm{ADU}]\cdot 1.16\,[e^-/\mathrm{ADU}] = 2.20\,e^-$ in an 1800 sec exposure on a science target. This is smaller than the read-noise, but not by much. Subtracting this dark frame from a 1800 s object spectrum would result in a minimum probable error per pixel (from just the read-noise) of $2.71/1.16/\sqrt{13} = 0.65$ ADU. This is rather more than we would like.

Apart from a few pixels that appear "hotter" than average, there is not much small-scale structure in the dark fraem. If we assume that the observed large-scale structure represent a genuine variation in the bulk dark rate, then we can improve the S/N by fitting a surface to the combined dark frame and use that instead to correct for the CCD dark current:

```
cl>  unlearn imsurfit
cl>  imsurfit ("DARK.fits", "MASTERDARK.fits", 6, 9, type_output="fit",
>>>       function="legendre", cross_terms=yes, lower=3., upper=3.,
>>>       ngrow=1, niter=3, regions="all")
```

Locate the warm/hot pixels with dark rates $> 4\,\sigma$ above the bulk dark rate:

```
cl>  imarith ("DARK.fits", "-", "MASTERDARK.fits", "_DARK.fits", title="",
>>>       hparams="", pixtype="real", calctype="real")
cl>  imstat ("_DARK.fits", fields="image,mean,midpt,stddev,min,max",
>>>       lower=INDEF, upper=INDEF, nclip=3., lsigma=3., usigma=3.)
#           IMAGE                  MEAN     MIDPT     STDDEV      MIN       MAX
 _DARK.fits                -9.2374E-5 -0.003310   0.73448  -2.20383   2.20359

cl>  x = 4.*0.73448
cl>  imreplace ("_DARK.fits", 0., lower=INDEF, upper=x)
cl>  imreplace ("_DARK.fits", 1., lower=x, upper=INDEF)
cl>  imstat _DARK.fits fields="image,min,max"
#       IMAGE         MIN      MAX
   _DARK.fits          0.       1.

cl>  ctio
cl>  pixselect ("_DARK.fits", lower=1.0, upper=INDEF, verbose-) | count
     63     189     945 STDIN
cl>  pixselect ("_DARK.fits", lower=1.0, upper=INDEF, verbose-)
    176    1  1.
    171    3  1.
    172    3  1.
   1019    8  1.
    382    9  1.
    ...   ..  ..
```

Add these pixels to the bad pixel list so that we may interpolate over them. First, we need to reformat the x and y coordinates from the output of `pixselect` into the "`x1 x2 y1 y2`" format of the bad pixel file:

```
cl>  pixselect ("_DARK.fits", lower=1.0, upper=INDEF, verbose-, > "hotpix.lis")
cl>  !awk 'print $1,$1,$2,$2' hotpix.lis >> badpix.ccd35
```

Clean-up of temporary files that are no longer needed:

```
cl>  delete tmp.fits,_DARK.fits yes verify-
```

Our best estimate for the bulk dark rate is $dc = 3600 \cdot ((1.124 \pm 0.734) \cdot 1.16/1800) = 2.61 \pm 1.70\, e^-$/pix/hr (but at a significance of only $\sim 1.5\,\sigma$, and without the possibility to cross-check the results for difference integration times, and with the caveat that we observe a dependence on both column and line number).

Normalize the level in the MASTERDARK frame to an integration time of 1 sec (and automatically update header keyword `darktime` accordingly):

```
cl>  imarith ("MASTERDARK.fits", "/", "1800.", "MASTERDARK.fits",
>>>       title="MASTERDARK", hparams="darktime", pixtype="real",
>>>       calctype="real", verbose+, >> "ccdproc.log")
```

9.3.4 Internal flats

• Process the incandescent (internal quartz lamp) flats (interpolate over bad pixels and subtract the overscan level; also subtract the MASTERBIAS frame if the signal level at some of the covered wavelengths is low; trim the frames and combine them into a high-S/N flat frame).

Since the average signal level in the flats will be a strong function of wavelength (due both to the CCD QE and to the lamp spectrum), that dependence needs to be fit out in order to create the pixel-to-pixel response frame using task `longslit.response`. This task not only performs a fit to the average of all lines, but divides the original frame by that fit:

```
cl>  unlearn response
cl>  response ("FLAT_grating.fits", "FLAT_grating", "RESPONS_grating.fits",
>>>    threshold=0.001, sample="*", naverage=-3, function="spline3", order=31,
>>>    low_reject=3., high_reject=3., niterate=2, grow=0.)
```

There is also a lower-order dependence of the average response along the spatial direction, which reflects the illumination pattern by the internal lamp, integrating sphere and auxilliary optics, and slit, that needs to be fit out using task `illumination`:

```
cl> unlearn illumination
cl> illumination ("RESPONS_grating.fits", "_RILLUM_grating.fits", bins="",
>>>   nbins=1, sample="*", naverage=-3, function="legendre", order=5,
>>>   low_reject=3., high_reject=3., niterate=2)
```

Create the "MASTERRESP" frame(s), one for each grating and/or grating tilt angle, by dividing the "RESPONS" frame(s) by the corresponding "_RILLUM" frame(s):

```
cl> unlearn imarith
cl> imarith ("RESPONS_grating.fits", "/", "_RILLUM_grating.fits",
>>>   "MASTERRESP_grating.fits", title="MASTERRESP", divzero=0., hparams="",
>>>   pixtype="real", calctype="real", verbose+, noact-)
```

(Should there be any non-illuminated portions of the CCD, then set all pixel values in those regions to 1.000 using task `imreplace` or STSDAS task `imcalc`.)

9.3.5 Twilight sky flats

• Process the twilight sky flats as were the incandenscent ones, and combine them. Fit the average response along the spatial direction, which represents the illumination pattern for light from the sky by the spectrograph and slit, and create the "MASTERILLUM" frame.

```
cl> response ("SKY_grating.fits", "SKY_grating", "_SRESP_grating.fits",
>>>   threshold=0.001, sample="*", naverage=-3, function="spline3", order=23,
>>>   low_reject=3., high_reject=3., niterate=2, grow=0.)

cl> illumination ("_SRESP_grating.fits", "MASTERILLUM_grating.fits", bins="",
>>>   nbins=1, sample="*", naverage=-3, function="legendre", order=5,
>>>   low_reject=3., high_reject=3., niterate=2)
```

9.4 Processing of the science target and comparison lamp frames

• Visually inspect all frames, fix any problems with the FITS headers.

→ tasks `imexamine` and `hedit`.

• Process all science target and comparison lamp frames (interpolate over bad pixels, subtract the MASTERBIAS frame and subtract the overscan level, trim the images, correct for dark signal (if significant), and divide by the appropriate MASTERRESP and MASTERILLUM frames — per grating and wavelength setting). Correct the signal in, or interpolate over pixels affected by cosmic ray hits.

→ tasks `fixpix`, `imarith`, and `colbias`.

• Determine a wavelength solution by identifying emission lines in one comparison spectrum, sampling lines in a few lines near the middle of the spatial extent, and fitting a polynomial to the identified wavelengths and pixel positions along the dispersion axis. Subsequently trace the same emission lines along the full extent of the spatial axis and fit a fully 2-D wavelength solution. Given this solution, and assuming the flexure of the spectrograph causes shifts that are small compared to the separation of emission lines, transfer it to all other comparison lamp spectra by fitting the relative shifts.

→ package `twodspec.longslit`
→ tasks `identify`, `reidentify`, and `fitcoords`

• Physically apply the 2-D wavelength solutions to the science target frames, taking the solution fitted to the associated comparison lamp spectra. In this step, the signal in the science frames will be resampled onto a regular (linear or logarithmic) wavelength grid. Check whether the wavelengths reported for the four corners of the frame make sense (some versions of the associated tasks have had a history of bugs!), and check that the sky lines in the resulting frames are no longer slanted (they show at least a minimal slant in the raw frames from most spectrographs).

→ package `twodspec.longslit`
→ task `transform`

• In each science frame, determine regions that are free of object signal for fitting the sky background. Subtract the sky background level and record the average sky spectrum in line #1 of each science frame. It may be necessary to refit the brightest sky lines with a fit of a different order.

→ package `rjtools` (and `images`)
→ tasks `getregion`, `fit1d`, and `imcopy`

• Trace the spatial location of the spectra as a function of wavelength and extract all signal within some fiducial aperture into a 1-D spectrum, optionally weighting that signal by its S/N.

→ package `twodspec.apextract`
→ tasks `aptrace` and `apsum` (alternatively, one can use tasks `apall`, but beware of all the hidden parameter sets that need to be properly populated).

• If flux calibration is possible and spectrophotometric standard stars were observed, determine the sensitivity function and correct the spectra.
→ package `onedspec`
→ tasks `standard`, `sensfunc`, `calibrate`, `dispcor` and `splot` (other possible tasks, depending on the goals of the spectroscopy, are `continuum` and `deredden`).

Sources, references, and additional reading

Baggett, S., et al. 2002, *The HST/WFPC2 Data Handbook, v. 4.0*, ed. B. Mobasher (Baltimore: STScI) (*http://www.stsci.edu/hst/wfpc2/*)

Bessell, M.S. 1990, PASP 102, 1181, *"UBVRI Passbands"*

Bessell, M.S. 2005, ARA&A 43, 293, *"Standard Photometric Systems"*

Bowmaker, J.K., & Dartnall, H.J.A. 1980, J. Physiol. 298, 501–511, *"Visual pigments of rods and cones in a human retina"*

van Dokkum, P.G. 2001, PASP 113, 1420, *"Cosmic-Ray Rejection by Laplacian Edge Detection"*

Gonzaga, S., et al. 2005, *ACS Instrument Handbook, v. 6.0*, (Baltimore: STScI) (*http://www.stsci.edu/hst/acs/documents/handbooks/cycle15/cover.html*)

Hamuy, M., Walker, A.R., Suntzeff, N.B., Gigoux, P., Heathcote, S.R., & Phillips, M.M. 1992, PASP 104, 533, *"Southern Spectrophotometric Standards"*

Howell, S.B. 2006, *Handbook of CCD Astronomy*, 2nd edition (Cambridge University Press, Cambridge UK)

Jakobsen, P., Jansen, R.A., Wagner, S., & Reimers, D. 2003, A&A 397, 891, *"Caught in the Act: A helium-reionizing quasar near the line-of-sight to Q 0302−003"*

Jansen, R.A., Franx, M., Fabricant, D.G., & Caldwell, N. 2000*a*, ApJS 126, 271, *"Surface Photometry of Nearby Field Galaxies: the data"*

Jansen, R.A., Fabricant, D.G., Franx, M., & Caldwell, N. 2000*b*, ApJS 126, 331, *"Spectrophotometry of Nearby Field Galaxies: the data"*

Liller, W. 1992, *The Cambridge guide to astronomical discovery* (Cambridge University Press, Cambridge USA)

Keel, W.C., *Astronomical Techniques* course notes (*www.astr.ua.edu/keel/techniques/*)

Krist, J. 2003, Instrument Science Report ACS 2003-06, *"ACS WFC & HRC field-dependent PSF variations due to optical and charge diffusion effects"* (*www.stsci.edu/hst/acs/-documents/isrs/isr0306.pdf*)

MacKay, C.D. 1986, ARA&A 24, 255, *"Charge-coupled devices in astronomy"*

Mellier, Y., Cailloux, M., Dupin, J.P., Fort, B. & Lours, C. 1986, A&A 157, 96, *"Evaluation of the performance of the 576×384 Thomson CCD for astronomical use"*

Newberry, M.V. 1991, PASP 103, 122, *"Signal-to-Noise Considerations for Sky-subtracted CCD Data"*

O'Connell, R.W., lecture notes (*www.astro.virginia.edu/class/oconnell/astr511/lec11-f03.html*)

Oke, J.B. 1990, AJ 99, 1621, *"Faint Spectrophotometric Standard Stars"*

Pavlovsky, C., et al. 2004, *The HST/ACS Data Handbook, v. 3.0* (Baltimore: STScI) (*http://www.stsci.edu/hst/acs/*)

Pence, W.D., 2002, *CFITSIO User's Reference Guide*, v. 2.4 (GSFC, Greenbelt MD) (*http://heasarc.gsfc.nasa.gov/fitsio*)

Perryman, M.A.C., et al. 1994 in: *Frontiers of Space and Ground-Based Astronomy*, eds. W. Wamsteker et al. (Kluwer Academic Publishers, Dordrecht), p. 537

Rousselot, P., Lidman, C., Cuby, J.-G., Moreels, G., & Monnet, G. 2000, A&A 354, 1134, *"Night-sky spectral atlas of OH emission lines in the near-infrared"*

Sterken, Ch. & Manfroid, J. 1992, *Astronomical photometry – A guide* (Kluwer Academic Publishers, Dordrecht)

Stetson, P.B. 2000, PASP 112, 925, *"Homogeneous Photometry for Star Clusters and Resolved Galaxies. II. Photometric Standard Stars"*

Taylor, V.A., Jansen, R.A., & Windhorst, R.A. 2004, PASP 116, 762, *"Observing Conditions at Mt. Graham: VATT* UBVR *Sky Surface Brightness and Seeing Measurements from 1999 through 2003"*

Wells, D.C., Greisen E.W., & Harten R.H. 1981, A&AS 44, 363, *"FITS – a Flexible Image Transport System"*

Hanisch, R.J., et al. 1993, "NOST Definition of the Flexible Image Transport System (FITS)", NOST 100-1.0

Vanouplines, P., *"A note on magnitudes"* (*http://www.vub.ac.be/STER/www.astro/magnitud.htm*)

Apogee Instruments Inc., *CCD University* (*http://www.ccd.com/ccdu.html*)

ESO's CCD Performance and Results web-page (*http://www.eso.org/projects/odt/Publications/-CCDpub_99/public.html*)

Molecular Expressions' Optical Microscopy Primer, *Digital Imaging in Optical Microscopy* (*http://micro.magnet.fsu.edu/primer/digitalimaging/*)

Nikon's Microscopy U, *"Introduction to Charge-Coupled Devices"*, by: K.R. Spring, T.J. Fellers & M.W. Davidson (*http://www.microscopyu.com/articles/digitalimaging/ccdintro.html*)

SITe 2048×4096 Scientific-Grade CCD (ST-002A CCD data sheet) (*http://www.ociw.edu/-instrumentation/ccd/parts/ST-002A.pdf*)

Outreach and Education site of the Australia Telescope (*http://outreach.atnf.csiro.au/education/-senior/astrophysics/*)

Frank Lakiere's web-site on photography (*webhost.ua.ac.be/elmc/website_FL/index-eng.htm*)

Various *Wikipedia* pages (beware: information in these may neither be complete nor correct)

'Astronomy with Charged Coupled Devices' is a companion e–book to the AST598 course on astronomical instrumentation and data reduction, in addition to 'Handbook of CCD Astronomy' by S.B. Howell (2nd ed.). The aim of this section of the course is to introduce students to the concepts and techniques that are important for near–UV through near–IR observational astronomical research that uses CCD detectors.