The SALT HRS Spectrograph

LUKE MARTIN GRAHAM TYAS

TYAS, LUKE MARTIN GRAHAM (2012) The SALT HRS Spectrograph. Doctoral thesis, Durham University.

The SALT HRS Spectrograph

Luke M. G. Tyas

A Thesis presented for the degree of
Doctor of Philosophy

Centre for Advanced Instrumentation
Department of Physics
Durham University
United Kingdom

September 2011

The SALT HRS Spectrograph

Luke M. G. Tyas

Submitted for the degree of Doctor of Philosophy
September 2011

Abstract

SALT HRS (Southern African Large Telescope High Resolution Échelle Spectrograph) is a high-resolution, high-efficiency spectrograph for the 11m SALT telescope in Sutherland, South Africa. The initial optical design work was performed at the University of Canterbury, New Zealand. Revisions to the concept, the mechanical design, manufacture, assembly and testing have been handled by the Centre for Advanced Instrumentation, at Durham University in the United Kingdom.

SALT HRS is a fibre-fed échelle grating spectrograph with four operational modes: low-, medium- and high-resolution and high-stability modes, having spectral resolutions of $R \approx$16000, 37000, 67000 and 67000 respectively over a wavelength range of 370-890nm. The instrument is of a dual channel, 'white pupil' design, in which the primary mirror acts to collimate light onto a single R4 échelle grating, and also to focus dispersed light to an intermediate focus. A dichroic beam-splitter separates the dispersed light into two separate spectral channels. Spherical pupil mirrors transfer the separated beams via a fold mirror to two wavelength-specific volume-phase holographic gratings (VPHGs) used as cross-dispersers. Cross-dispersed spectra are then imaged by two fully dioptric camera systems onto optimized CCD detectors.

This thesis presents the results of the laboratory testing and specification of several critical sub-systems of SALT HRS, as well as the development of key software tools for the design verification and operation at the telescope. In Chapter 1 we first review the technical development of high-resolution spectroscopy and its specific implementation in SALT HRS. In Chapter 2 we develop a comprehensive throughput model of the entire system based on a combination of as-built performance and specific throughput measurements in the laboratory. This is used to make some specific predictions for the on-sky performance of SALT HRS and the magnitude limits for science targets. We also present a graphical exposure time calculator based on these measurements which can be used by an astronomer to plan their observations with SALT HRS.

Chapter 3 contains a detailed treatise on the optical fibre system of SALT HRS. Considerations for the use of optical fibres in astronomy are provided, as are details of an optional double scrambler, and the various instrument fibre modes. Extensive measurements of focal ratio degradation (FRD) are also presented, with testing of input beam speed; wavelength; fibre bending; variable pupil mirror illumination; and vacuum tank pressure dependency. The systems for fibre management are reviewed, as is the fibre bundle assembly process.

Testing of two further sub-systems is described in Chapter 4. Firstly the long-term stability of the mirror mounting mechanisms is determined. The advantages of cross-dispersion of échelle spectra using volume-phase holographic gratings are then discussed, and the results of diffraction efficiency measurements are given for both red and blue channel gratings.

Modern CCD technologies are examined in Chapter 5, and the blue detector is experimentally characterized using photon transfer and quantum efficiency curves. It is also used for an investigation into cosmic ray events in CCDs. Results from shielding the detector using lead are described, as is an attempt to distinguish the source of the events based on their morphology.

Finally, Chapter 6 deals with the handling of data produced by SALT HRS. Methods of wavelength calibration of the spectra are discussed, including the use of Thorium-Argon lamps and an iodine absorption cell. The implementation of a Python based quick-look data reduction pipeline is reviewed, with a description of the processes performed.

A summary of the thesis is given in Chapter 7.

Declaration

The work in this thesis is based on research carried out at the Centre for Advanced Instrumentation, the Department of Physics, University of Durham, United Kingdom. No part of this thesis has been submitted elsewhere for any other degree or qualification and it is the sole work of the author unless referenced to the contrary in the text.

Some of the work presented in this thesis has been published in the following publications:

Publications

D. G. Bramall, R. M. Sharples, L. M. G. Tyas, J. Schmoll, P. Clark, P. Luke, N. Looker, N. A. Dipper, S. Ryan, D. A. H. Buckley, J. Brink, S. I. Barnes: "The SALT HRS spectrograph: final design, instrument capabilities and operational modes", Proc. SPIE **7735**, 77354F (2010)

S. I. Barnes, P. L. Cottrell, M. D. Albrow, N. Frost, G. Graham, G. Kershaw, R. Ritchie, D. Jones, R. M. Sharples, D. G. Bramall, J. Schmoll, P. Luke, P. Clark, L. M. G. Tyas, D. A. H. Buckley, J. Brink: "The optical design of the Southern African Large Telescope High Resolution Spectrograph: SALT HRS", Proc. SPIE **7014**, 70140K (2008)

Acknowledgements

First and foremost I would like to thank my PhD supervisor, Prof. Ray Sharples, for tolerating any ineptitudes in the past four years. His time, patience, encyclopædic knowledge and clear guidance are greatly appreciated, as are the efforts he made to ensure my rushed but nevertheless timely completion of this thesis.

Thanks too must go to Stuart Barnes, the primary designer of SALT HRS, whose tome-like thesis has been my 'go to reference' for much of the background to this project. I would also like to offer my sincerest thanks to the SALT HRS project engineers David Bramall, Jürgen Schmoll and Paul Clark. All three have been helpful, understanding, quick to reply to questions, and a great support in my various experiments. David has offered his engineering expertise in working with the mirror mounts and fibre feeds, and always seems to know what's going on; Jürgen has helped align optics on numerous occasions, answered so many of my queries and shown great enthusiasm for our work; and Paul has been of tremendous help in using the detectors, as well as always seeming to know where something I'm looking for can be found. I would also like to thank the various inhabitants of offices 125b and 127 over the years, who have been great colleagues and also friends: Jonny Taylor, James Osborn and Claire Poppett in particular. I have greatly enjoyed my time among the 'beautiful freaks' who make life so interesting at the Centre for Advanced Instrumentation, and will no doubt look back on these years as some of the best of my life.

On the pastoral side, thanks go to my 'Big Bang Theory' housemates - Richard 'Leonard Hofstadter' Abel and Paul 'Sheldon Cooper' Mellor - who have made my time in Durham all the easier. I hope we will get chance to waste time watching CSI repeats and competitively cook a 'family dinner' for each other again in the future. Margarita Gelepithis, Liz Ellis and Dima Smayra also deserve a big thank you for their continued friendship over the years. Liz also earns a special mention for her surprise delivery of chocolates to keep me going through numerous 'all-nighters'. An enormous thank you must also go to Maryam Al-Azeezy, whose love and support this year has been a greatly stabilizing influence and the source of much happiness.

Finally I would like to thank my parents, Martin and Janice, and grandparents William, Stella, Winifred and Eric, for their continuous encouragement and support in life and throughout my education. I dedicate this doctoral thesis to all of them.

Contents

List of Figures

Chapter 1

Introduction

1.1 Astronomical Spectroscopy

1.1.1 A Brief History of Spectroscopic Development

Spectroscopy is one of the most fundamental tools at the modern astronomer's disposal. Whilst astrometric measurements represent some of the earliest examples of the science of astronomy, having been made by humankind since ancient times, discerning spectral information from celestial objects is a relatively modern undertaking. In 1835, French philosopher August Comte was writing of the hope to 'study by any means' anything beyond 'geometric or mechanical phenomena' of stars - unbeknown to him, Fraunhofer and many others had already taken the first steps towards modern spectroscopy (Comte, 1835).

The first spectroscopic study of sunlight was an experiment of Isaac Newton's, described in his 1672 book 'Treatise on Optics'. He used a prism to separate white light into what he thought to be an infinite array of colours, also noting that a thin slit resulted in greater intensity of colour compared to experiments performed with larger circular apertures. Although he is acknowledged as having observed the spectrum of Venus, in addition to that of the Sun, the poor quality of glass available in his prisms denied him the opportunity of discerning any line features (Loewen & Popov, 1997).

It was not until 1800 that William Herschel continued Newton's line of research, discovering that the spectrum continues beyond the red end of the visible, into the infra-red. Using thermometers he also determined that such radiation is perceived as heat, that the maximum brightness of sunlight occurs in the yellow and that the maximum energy output from the Sun is just redward of the visible spectrum (Kayser, 1900). Not long after, Johann Ritter found that the spectrum also extended in the blue direction, noticing that ultra-violet radiation blackens silver chloride on exposure (Ritter, 1803). Concurrently, Thomas Young began using a transmission grating of 500 lines per inch, noting that red light could be observed in four different directions and that the sine of their respective angles varied according to the ratio 1:2:3:4 (Young, 1803). Young was the first to distinguish colours by their wavelength. Presumably independent research on the topic was also carried out in the United States. In [illegible] Francis Hopkinson noticed that multiple images of a street lamp could be seen through a fine silk handkerchief (Loewen & Popov, 1997). His friend, astronomer David Rittenhouse determined that this was a diffraction effect and went on to build gratings (using wire wrapped around screws of a known thread) and constrain the wavelength of visible light (Rittenhouse, 1786).

Joseph von Fraunhofer built the first ruling engine for diffraction gratings at his leading Münich optics company. His measurements were remarkably detailed for the time and he understood a variety of nuances relating to grating technology, including the need for high accuracy (1% groove spacing), polarization effects, and the effect on efficiency of groove shape variations. Fraunhofer also began using collimated light (using telescopes to do so) and cross-dispersers (in the form of a prism at the entrance slit) to observe higher orders, normally overlapped with other orders. A further accomplishment was his identification of the solar spectral lines (see Figure 1.1), which he used as wavelength markers when measuring the refractive indices of glass blanks for use in achromatic lens systems. Fraunhofer also examined flames (noting that two lines occur at nearly the same wavelength as the 'double D' lines in the solar spectrum) and the continuum spectrum given off by hot glass; he also used a 100mm telescope with an objective prism to observe numerous stellar and planetary spectra (Fraunhofer, 1823).

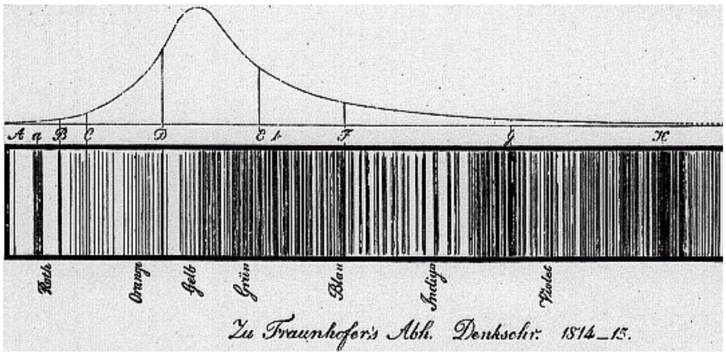

Figure 1.1: The solar spectrum as recorded by Joseph von Fraunhofer (reproduced from Tennyson, 2005).

Over the following decades, numerous incremental discoveries occurred in the areas of both absorption and emission spectra. William Herschel's son, John Herschel investigated the spectra of many substances and measured the absorption lines in coloured glass, believing the lines to be related to temperature rather than the elements of composition (J. Herschel, 1823). [illegible] Fox Talbot performed a body of tests showing various emission spectra from flames and Charles Wheatstone measured spectra of a variety of metals excited by electric sparks. Wheatstone also made the connection between the spectra being intrinsic to the elements and suggested spectroscopy as a means of determining material composition (Wheatstone, 1861). Anders Ångstrom produced a great deal of work on spectra beginning in 1855 (with his fame recognized with the non-SI unit, the Ångstrom, which represents 1×10^{-10}m) and was the first to describe the difference between the spectra of gaseous and solid bodies (Ångstrom, 1855). Julius Plücker used glass tubes blown by Heinrich Geissler (whose name is still given to the tubes) to measure gas emission lines stimulated by electric discharge, in doing so confirming that the three hydrogen lines H_α, H_β and H_γ were the same as Fraunhofer's F, C and G absorption lines in the Sun (Plücker, 1858).

The physicist Gustav Kirchhoff and chemist Rudolph Bunsen cemented the significance of spectroscopy in modern science. With his prescient publication, Kirchhoff (1859) outlined the connection between the emission and absorption of light, indicating the fundamental connection between the unique spectra emitted by different elements. Kirchhoff's law states that the ratio between the degree of emission and the degree of absorption for rays of the same wavelength is a constant for all bodies at the same temperature, or

$$\frac{\epsilon_\lambda(T)}{\kappa_\lambda(T)} = \text{constant} \tag{1.1}$$

where ϵ and κ are the coefficients of emission and absorption at wavelength λ and temperature T (Kirchhoff, 1860). Therefore a gas which radiates a given line spectrum must also absorb the wavelengths of the line at which it radiates, showing that the Fraunhofer D lines were the yellow lines of Sodium and that a Sodium flame will absorb the same yellow light from a stronger source (see Figure 1.2 for an example of absorption and emission line spectra). Kirchhoff described how the dark Fraunhofer lines were due to absorption by corresponding elements in the cooler solar atmosphere whilst the continuum spectrum came from the hotter interior. Bunsen and Kirchhoff continued to work analysing spectra, performing measurements on many pure elements, in doing so discovering hitherto unknown elements such as Caesium and Rubidium (Kirchhoff & Bunsen, 1860). The experimental equipment made use of a CS_2 filled glass prism for dispersion and a platinum coil covered in a salt of the studied compounds, heated using a gas burner which produced a hot and almost colourless flame, still bearing the name of its inventor today. Their work drew research from the scientific community and also ignited popular interest. Kirchhoff's graduate student Georg Quincke continued research in the 1870's performing experiments with both immersed gratings and reflection gratings (by coating Nobert glass ruled gratings) as well as producing the first grating replicas (Quincke, 1872).

1.1.2 The Advance of Astronomical Spectroscopy

Since Newton's first experiments with a prism, spectroscopy has been intimately tied with astronomy. Norman Lockyer continued the tradition in 1868 with his solar spectrum studies which ultimately lead to the discovery of a new element, Helium (after 'helios', the Greek word for 'Sun') and also pioneered the use of photography and Doppler measurements in spectral studies (Lockyer, 1874). Zöllner improved on this work with his measurements of the rotation speed of the Sun using shifts in solar prominences as a base for his calculations (Zöllner, 1871). Huggins ex[illegible]ed Doppler measurements to other stars, a field which is still widely explored in modern

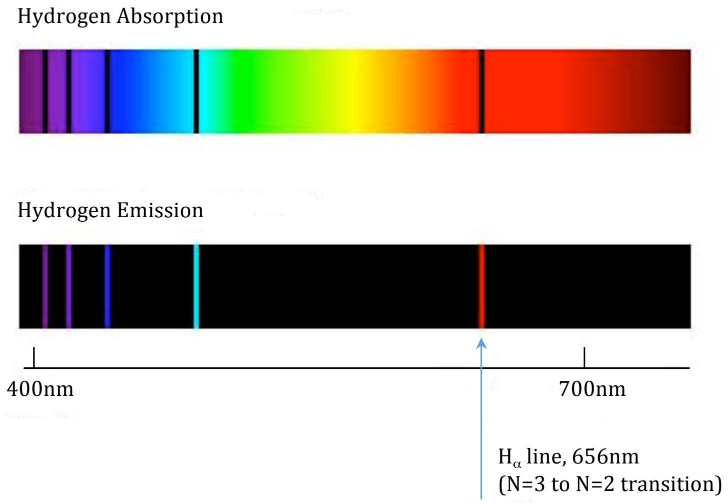

Figure 1.2: An example of both absorption and emission spectra for Hydrogen, with the H_α line indicated. Note that Hydrogen absorbs and emits at the same wavelengths.

astronomy for velocity measurements (Huggins, 1872). Others took the field into the ultra-violet regime, using fluorescent eyepieces to make measurements, whilst Italian physicist and theologian Angelo Secchi, through the accumulation of over 4000 stellar spectra, devised the first stellar classification system based on five classes (which would later be replaced by the familiar Harvard spectral classification) (Secchi, 1866).

In contemporary astronomy, a large proportion of our knowledge about the universe has its root in spectroscopy. Spectral information, coupled with greater knowledge of atomic physics has allowed astronomers to determine not only an object's composition and classify it accordingly, but also to calculate age, rotational or radial velocities, peculiar motion, distance, luminosity, temperature, seismology and rate of mass loss. Spectra derived velocity measurements also allow for indirect identification of otherwise 'hidden' companion objects such as black holes and extra-solar planets (Tennyson, 2005). Spectroscopy is also valuable for solar system studies, since planets and their satellites reflect sunlight and also add their own absorption lines characteristic of their surface composition or atmospheres. The spectra of galaxies are similar to those of stars (since the light is produced by billions of stars within the galaxy) but are nonetheless vital for studies of star formation processes and mapping the large scale structure of the Universe via their Doppler recession velocities. Edwin Hubble's spectral measurements of galaxies in the 1920's showed the fundamental relationship between a galaxy's distance and recession speed, indicating that the further away the galaxy, the faster it was receding from us. These observations provided the first evidence for the Big Bang theory of the origin of the Universe. In the following decade Fritz Zwicky also used measurements of galaxy spectra to measure the random motions of galaxies in clusters. The measurements suggested the presence of large amounts of non-luminous mass (dark matter). Spectroscopy is also of great significance in the study of gaseous nebulae and quasars. Nebulae show no strong continuous spectrum like a star or galaxy, having instead a spectrum dominated by isolated strong emission lines. Initially some of these lines were thought to correspond to a new element, but further investigation revealed the emission came from doubly ionized oxygen at extremely low density. The rarefied conditions allowed for 'forbidden lines' which are suppressed at normal densities. Quasars (strong radio sources found to be associated with dim, blue, point-like objects) were revealed to be characterized by extremely red-shifted spectra, meaning that they are extremely distant objects formed in the early Universe, now thought to be powered by accretion onto super-massive black holes.

1.1.3 Grating Spectrographs

In the following decades, into the early 20th century, spectrographs using a prism as the dispersion element had become commonplace in many observatories. The first diffraction grating spectrogram of the solar spectrum was obtained in 1873 by Henry Draper, using a grating prepared by Lewis Rutherford, who ruled reflection gratings in speculum metal (a soft copper-tin alloy) (Draper, 1873). Spectral photography became an in important tool for a variety of reasons. Firstly, fainter objects could be observed due to increased integration times. Furthermore the measurement could be repeated to gain increased precision and the time spent recording observations by hand at the telescope eyepiece could be reduced.

Lord Rayleigh had showed in 1874 that the ideal diffraction grating would be a better means for achieving high spectral resolution than the popular prismatic dispersion, writing that:

> *To obtain a diffraction spectrum in the ordinary sense, containing all the light, it would be necessary that the retardation should gradually alter by a wavelength in*

> *passing over each element of the grating and then fall back on its previous value, thus springing suddenly over a wavelength.*

However, the manufacture of gratings at that time was far from ideal and the efficiency was low since the light is dispersed into several orders. By 1882 Henry Rowland of Johns Hopkins University had perfected his ruling engine which produced gratings with much higher tolerances and in much larger dimensions. With his development of concave gratings the need for a collimating lens (and with it the associated transmission loss and wavelength limitations) disappeared. Grating spectroscopy was able to deliver a resolution greater than could be achieved with a prism 1m in base length. Rowland's successor, John Anderson manufactured a so-called 'blazed' grating which could diffract a large fraction of light into the first order (Hearnshaw, 1986). Later came the technology to vacuum deposit aluminium on ruled glass, which proved a much superior medium to the previously used speculum metal.

By 1929 the Mount Wilson observatory had a Cassegrain spectrograph which incorporated flexure control to improve stability. As Coudé telescopes became popular it was possible to remove the issue of flexure altogether whilst also making use of larger gratings (since there is essentially no space limitation at this focus) and a long focal length camera to achieve high dispersion (Hearnshaw, 1986). See Section 1.2.2 for an explanation of the benefits of a long focal length camera.

The next major step came with the move from blazed to échelle grating. The resolving power of a grating is proportional to the product of N, the number of rulings, and m, the order of diffraction. It is therefore possible to achieve high resolution by either increasing the number of rulings (either the size of the beam and grating or the ruling density) or by increasing the order of diffraction. Early work progressed by increasing the density and by blazing the grooves with improved tool shapes into the second or third order. Yet a grating diffracting into a high order can have the same resolving power with a fraction of the grooves. At the turn of the century Michelson began producing échelon gratings made from a number of parallel plates of glass and acting in transmission. By 1933 the technique was refined to produce a grating with a resolving power of 10^6, although the dimensions of such a grating was highly constrained by the

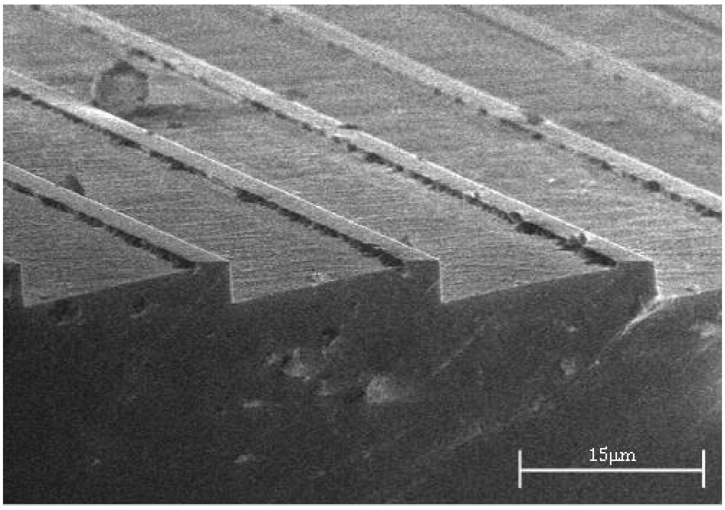

Figure 1.3: Scanning electron microscope image of an échelle grating.

manufacturing technique. In 1910 Robert Wood had developed a reflecting échelette grating for infra-red use which diffracted into only a handful of orders (Wood, 1947). The échelle grating, which acts in reflection and is more coarsely ruled than an échelette was first described in 1949 by Harrison. The technical difficulty of the échelle is that the reflecting facets need to have a groove profile accurate to $\lambda/10$ and that each groove must be positioned to the same accuracy (Harrison, 1949). By the 1970's échelle gratings with dimensions of 300 × 400mm were possible. Figure 1.3 shows a scanning electron microscope image of an échelle grating, demonstrating the precision required of the ruling process.

1.2 Échelle Spectrographs

1.2.1 Introduction to Échelle Gratings

Light incident on a grating surface is diffracted from the grooves, with each groove becoming a small source of reflected light. This is represented in the left diagram of Figure 1.4, where light incident at angle α to the normal N is diffracted through an angle β by reflective grooves with spacing σ and facet width σ_S. From the condition for constructive interference at a given

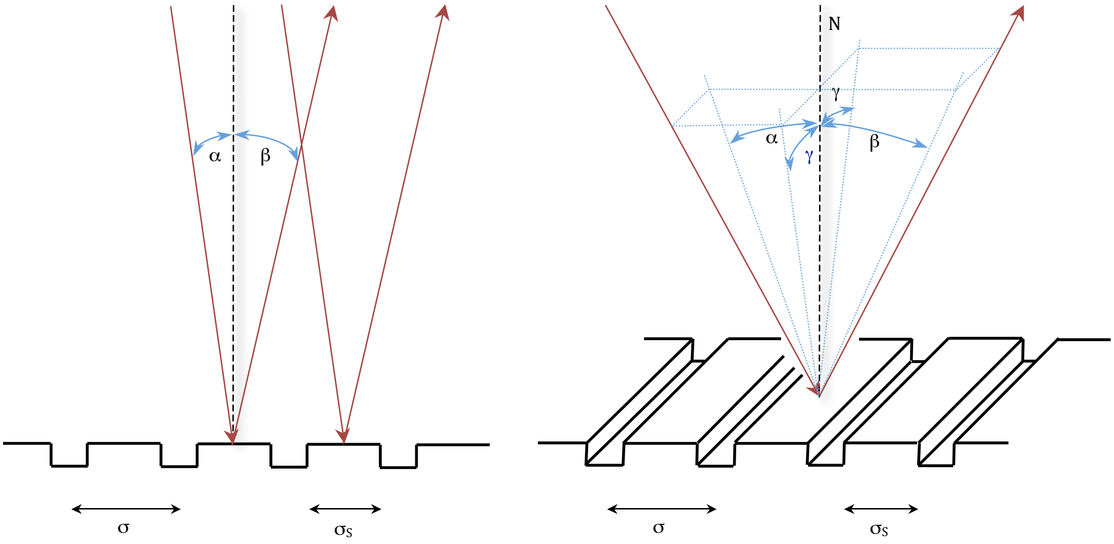

(a) Plane view (b) Profile view with $\gamma \neq 0$ illumination

Figure 1.4: Schemata of a reflection grating with illumination at angle α and diffracted rays at angle β from a grating with groove width σ.

wavelength, λ,

$$m\lambda = \sigma(\sin\alpha + \sin\beta) \tag{1.2}$$

where m is an integer multiple of the wavelength and labelled such that for the specular reflection order, $m = 0$. This is known as 'the grating equation'. m represents the number of wavelengths path difference between light reflected from successive grooves. For $m = 0$ the grating acts as a mirror with all wavelengths superimposed. For the non-specular orders of $m \neq 0$ the diffraction angle β is dependent on λ so that wavelengths are separated angularly. This is assuming that incident and diffracted rays are perpendicular to the grooves. Given the slightly more complex situation where the grating is illuminated at an angle γ with respect to the facet normal the grating equation gains an extra term to become:

$$m\lambda = \sigma(\sin\alpha + \sin\beta)\cos\gamma \tag{1.3}$$

The diagram on the right side of Figure 1.4 shows this more complicated case. As mentioned earlier in Section 1.1.3, it is enormously beneficial to diffraction efficiency to put a large fraction

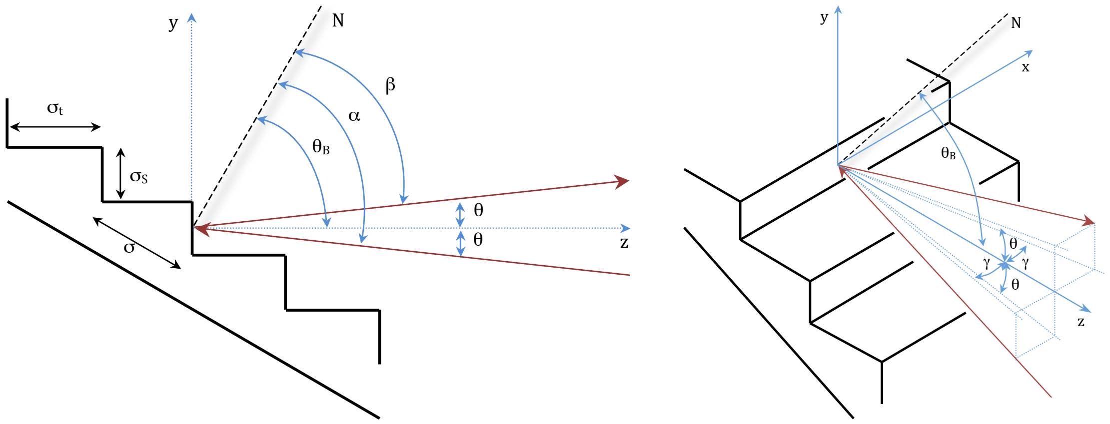

(a) Profile view of a blazed grating showing the optical step σ_t, and facet width σ_s

(b) Isometric view with $\gamma \neq 0$ illumination

Figure 1.5: Diagrams of a blazed grating with grating normal N, blaze angle θ_B, angle of illumination α and angle of diffraction β.

of the light into a single order. This can occur when the facets are angled so that for a certain wavelength in a particular order the incident (α) and diffracted (β) rays are almost at the same angle. This is made possible by ‘blazing’ the grating, which is achieved in practice by orienting the groove facets such that the diffraction angle is very close to the angle of specular reflection. The grating facet angle with respect to the grating normal is known as the blaze angle, θ_B. Échelle gratings are blazed gratings with a high θ_B, and are often described in terms of an ‘R-number’, which is simply $\tan\theta_B$. Common R-numbers include R2, R2.8, R3 and R4. Figure 1.5 shows both a profile and isometric view of an échelle grating, from which it is evident that the angles of incidence and diffraction (α and β) may be defined in terms of the blaze angle (θ_B) according to

$$
\begin{aligned}
\alpha &= \theta_B + \theta \text{ and} \\
\bar{\beta} &= \theta_B - \theta \qquad (1.4)
\end{aligned}
$$

where θ is the facet illumination angle with respect to the facet normal.

$\bar{\beta}$ is the diffraction angle for a blaze wavelength λ_B in the centre of order m. For an échelle illuminated out of the normal plane ($\gamma \neq 0$) the blaze wavelength can be defined in terms of the grating equation as

$$\begin{aligned} m\lambda_B &= \sigma(\sin\alpha + \sin\bar{\beta})\cos\gamma \\ &= 2\sigma\sin\theta_B\cos\theta\cos\gamma \end{aligned} \tag{1.5}$$

The condition whereby $\theta = 0$ (so that $\alpha = \bar{\beta}$) is known as the Littrow condition. Figure 1.5a shows the optical step and facet width (σ_t and σ_s respectively) which can be defined by

$$\begin{aligned} \sigma_t &= \sigma\sin\theta_B \text{ and} \\ \sigma_s &= \sigma\cos\theta_B \end{aligned} \tag{1.6}$$

The optical step determines the order of interference for diffracted light according to

$$m = \frac{2\sigma_t}{\lambda} \tag{1.7}$$

Quasi-Littrow illumination again sets $\theta = 0$ but allows $\gamma \neq 0$.

1.2.2 Échelle Dispersion

One of the key properties of a grating is the angular dispersion, which is calculated by differentiating equation 1.3 with respect to λ for a given α so that:

$$\frac{d\beta}{d\lambda} = \frac{m}{\sigma\cos\beta\cos\gamma} = \frac{\sin\alpha + \sin\beta}{\lambda\cos\beta} \tag{1.8}$$

which in the centre of the order at the blaze wavelength becomes:

$$\frac{d\beta}{d\lambda} = \frac{2\sin\theta_B\cos\theta}{\lambda_B\cos\beta} \tag{1.9}$$

In order to give high dispersion it is therefore necessary to either increase the angle of incidence, α, or increase the groove density, $1/\sigma$. An échelle grating is able to offer high dispersion whilst still being coarsely ruled by having a large blaze angle, θ_B and operating in high orders (typically $m = 10$ to over 100). The angular dispersion is of course a property of the échelle grating and the conditions of its illumination, being independent of the other components of the spectrograph in this sense. To understand the physical extent over which a spectral range is dispersed requires calculation of the linear dispersion of the instrument. This is dependent on the optical system since the extent δl over which a wavelength range $\delta\lambda$ is spread on the detector is given by

$$\delta l = f_{\mathrm{cam}} \frac{d\beta}{d\lambda} \delta\lambda \tag{1.10}$$

where f_{cam} is the focal length of the camera which images the dispersed spectrum onto the detector.

1.2.3 Free Spectral Range

Free spectral range, $\Delta\lambda_{\mathrm{FSR}}$ is the change in wavelength from an order m to the next order $m \pm 1$:

$$\Delta\lambda_{\mathrm{FSR}} = \frac{\lambda}{m} \tag{1.11}$$

Any wavelength in order m is also present in the $m-1$ and $m+1$ orders (although the intensity and diffraction angle will vary between the orders). In terms of the blaze wavelength, the free spectral range is given by

$$\Delta\lambda_{\mathrm{FSR}} = \frac{\lambda_B^2 \cos\gamma}{2\sigma \sin\theta_B \cos\theta} \approx \frac{\lambda_B^2}{2\sigma_t} \tag{1.12}$$

Multiplying the free spectral range by the angular dispersion gives the angular extent of one free spectral range such that

$$\Delta\beta_{\mathrm{FSR}} = \frac{d\beta}{d\lambda} \Delta\lambda_{\mathrm{FSR}} = \frac{\lambda_B}{\sigma \cos\beta \cos\gamma} \tag{1.13}$$

Under Littrow illumination this may be reduced to $\Delta\beta_{\mathrm{FSR}} = \lambda_B/(\sigma_s \cos\gamma)$. For a given diffraction angle β and order m the angular extent of a spectrum varies with the groove density. With

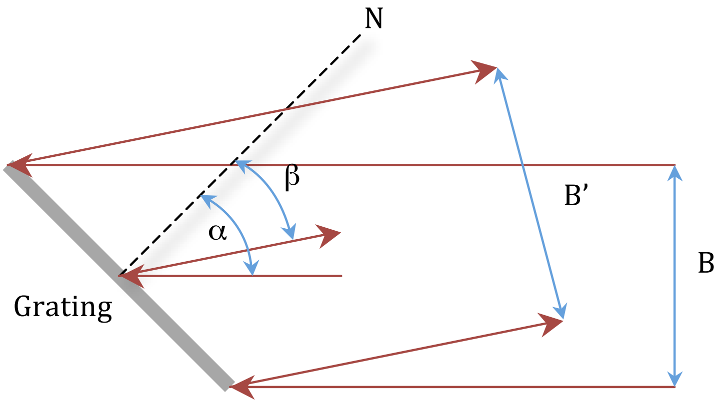

Figure 1.6: The effect of anamorphic magnification r on on beam size B.

large σ (i.e. a coarse grating) the spectrum has a smaller angular extent per free spectral range than a spectrum from a finely ruled grating.

1.2.4 Anamorphic Magnification

Anamorphic magnification, r, is the expansion of a beam with diameter B to a larger diameter B' following illumination at angle α and diffraction through angle β:

$$B' = \frac{B}{r} \tag{1.14}$$

The quantity r is given by

$$r = \left|\frac{d\beta}{d\lambda}\right| = \frac{\cos\alpha}{\cos\beta} \tag{1.15}$$

Since the magnified beam diameter $B' \propto \beta$, anamorphic magnification effects vary across a free spectral range, particularly for SALT HRS which uses a high R-number grating large free spectral ranges. Anamorphic magnification is easily visualised in a diagram such as Figure 1.6 which shows increased beam diameter B' compared to the (monochromatic) incident beam B for a given diffraction angle β.

1.2.5 Cross Dispersion

Since échelle gratings work at high orders, there are many combinations of λ and m which satisfy the grating equation (see equation 1.3). A second dispersive element is therefore required to disperse the spectra in a direction orthogonal to the main échelle dispersion direction. An in depth introduction to cross-disperser properties and theory can be found in Section 4.2.3, so they shall not be discussed further here.

1.3 The SALT Telescope

1.3.1 Overview of SALT

SALT (the Southern African Large Telescope) is a modern optical telescope situated in Sutherland, South Africa at an altitude of 1798m above sea level (Buckley et al., 2008). The primary mirror is 11m in diameter and of a segmented design, comprising 91 separate but identical 1m hexagonal segments laid out in a hexagonal mosaic (arranged so that an 11m circle is transcribed by the vertices). In many ways, SALT may be considered a sister telescope to the earlier Hobby Eberly Telescope (HET) at the McDonald Observatory in Texas, USA. Both telescopes share a common design basis (although SALT shows some evolutionary design improvements) with a mission of achieving high light collecting power for a low cost. The concept was first proposed by Ramsey & Weedman (1984) and is inspired by the Arecibo radio telescope: the SALT primary mirror mosaic is spherical and uses a tracker assembly to follow the curved focal surface. Whilst Arecibo faces directly upwards towards the sky, and must wait for an object to traverse its line of sight, the SALT primary mirror is mounted at a fixed altitude of 37° from the vertical and can move in azimuth to acquire a target. A spherical aberration corrector (SAC) is mounted concentric to and 13.08m (half the radius of curvature) away from the primary mirror, to compensate for the effects of a non-parabolic primary. The tracker assembly also features an atmospheric dispersion corrector (ADC) which compensates for the fact that different wavelengths of light are refracted by different amounts in the atmosphere. Figure 1.7 shows the general optical layout

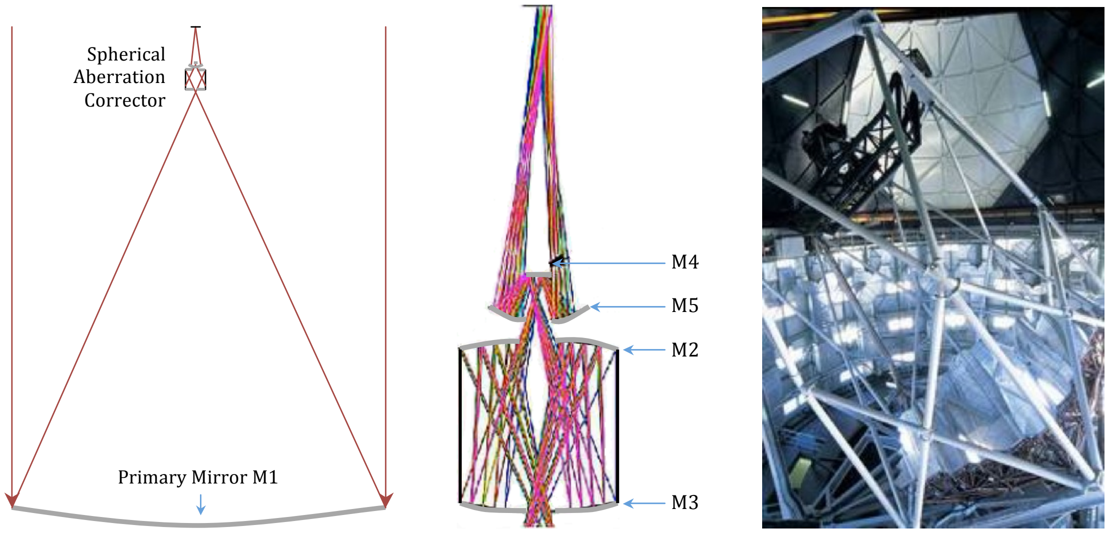

(a) SALT Optics (b) The SAC (c) SALT photograph

Figure 1.7: A pictorial overview of the SALT telescope showing the overall telescope optics (*a*), the spherical aberration corrector (SAC) optics (*b*) and a photograph for comparison (*c*).

of the SALT telescope and an expanded view of the SAC mirrors alongside a photograph of the telescope.

1.3.2 The SALT Sky Viewing Window

The SAC provides a circular 8 arcmin field of view over a declination range of $-75°$ to $+10°$ (with a 1 arcmin wide annulus 'border' for guide stars) and feeds the prime focus with an f/4.2 beam. The SAC moves during the course of an observation to track an object on sky, with the maximum object tracking time dependent on altitude and varying between 45 minutes at the northern declination range to 3 hours at the southern extreme. Although zenith angle is fixed, in between observations the telescope is capable of rotating in azimuth allowing multiple telescope re-acquisitions. The SALT mount is therefore quite different from the typical equatorial mounts of early telescopes and the alt-azimuth mounts found on other ∼10m class telescopes such as the ESO VLTs in Chile or the Keck telescopes in Hawaii. The most obvious disadvantage of the HET/SALT mount design is that it gives the telescope a restricted annulus-shaped observing patch on the sky as can be seen in Figure 1.8, which is not as flexible for target acquisition as an

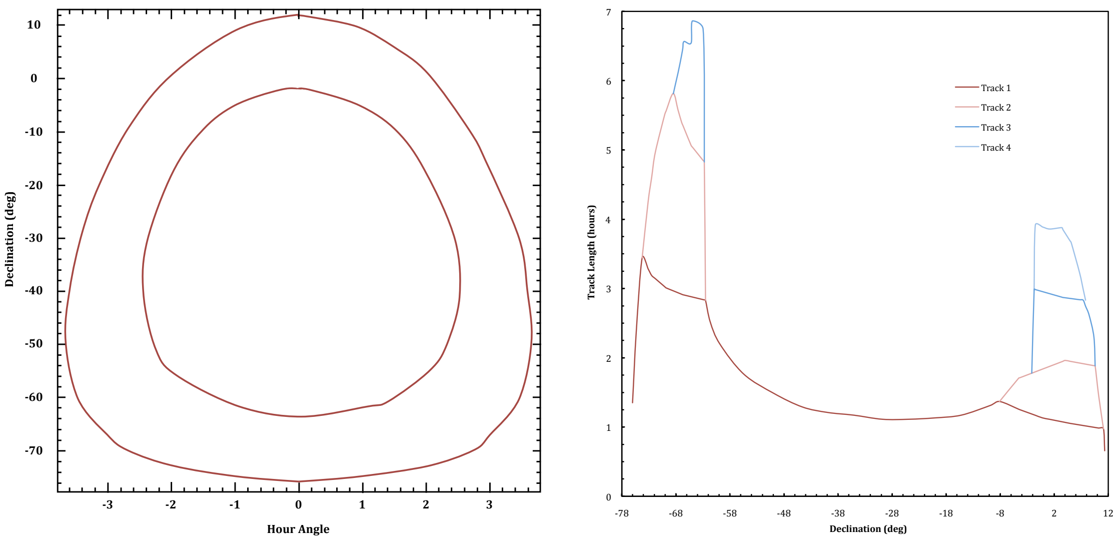

(a) SALT observing patch (b) Tracking time

Figure 1.8: The SALT sky viewing annulus as function of declination and hour leads to variance in tracking time depending on the declination of the target. In some cases, multiple telescope re-acquisitions may increase the observing time over the course of a night (from http://www.salt.ac.za/telescope/performance-characteristics/).

alt-azimuth system. However, using optimized telescope queuing systems it is possible to access ∼70% of the sky over the course of a year. Figure 1.8 also shows the tracking time as a function of both declination and the hour angle at which the track began. A target with intermediate declination may be tracked for the full 2 hour viewing zone but at the declination extremes the track time is shortened. However the source may continue to be tracked by moving the telescope in azimuth and re-acquiring. Therefore quasi-continuous viewing for 7 hours is achievable in the south and up to 4 hours in the north (see http://www.salt.ac.za/proposing/observing-constraints/). Each mirror segment is mounted in a truss arrangement at nine points with three actuators to move it in tip/tilt and piston. The mirror segments are aligned during twilight using a Shack-Hartmann wavefront camera located in a tower alongside the dome building and are kept in position during observation with a series of capacitive edge sensors in a closed-loop control system which updates every 20s. The tracker and payload also move with great precision with movements in tip, tilt, x, y, z and field rotation to an accuracy of 5µm or 1[illegible]. The tracking of an object requires x and y motion with concurrent tip/tilt changes, whilst telescope focus is maintained with z translation (O'Donoghue et al., 2006). The practical aspects of the

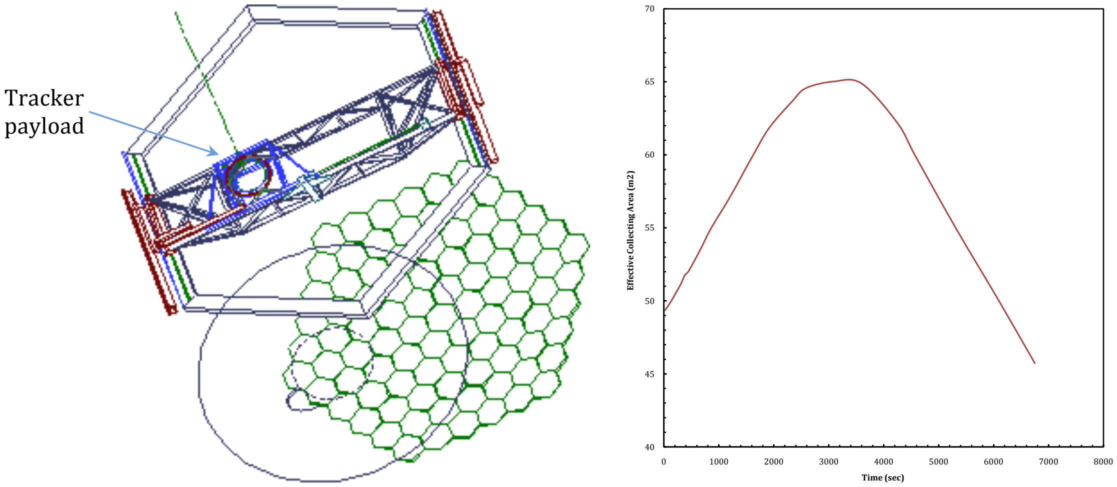

(a) Object Tracking

(b) Telescope Collecting Area

Figure 1.9: The SALT primary mirror collecting area varies as a function of time across an observation due to the variable pupil illumination effects of the unique tracking system (from http://www.salt.ac.za/telescope/observing-constraints).

novel HET/SALT tracking system for operation management and sky access is discussed above, but it should also be noted that the unique tracking method causes a variably illuminated pupil on the primary mirror. This has implications for the SALT HRS instrument performance and is discussed further in Section 3.3. Figure 1.9 shows how the telescope entrance pupil migrates over the primary mirror array and gives an example of how the effective collecting area of the telescope varies with time over the course of tracking, peaking in the middle of the track at 65m^2 when the pupil and mirror are co-aligned (note that in this case, the 'time' ordinate is arbitrary).

1.3.3 Breaking the Cost versus Diameter Curve

Although the object tracking system becomes more complex as a result of the SALT design, the limited motion of the primary mirror provides considerable cost savings and also reduces flexure and precision control issues. In contrast, alt-azimuth mount, although reducing space and cost for the mount and dome when compared with an equatorial mount, requires sophisticated control in three degrees of freedom (altitude, azimuth and field de-rotation). The use of a segmented

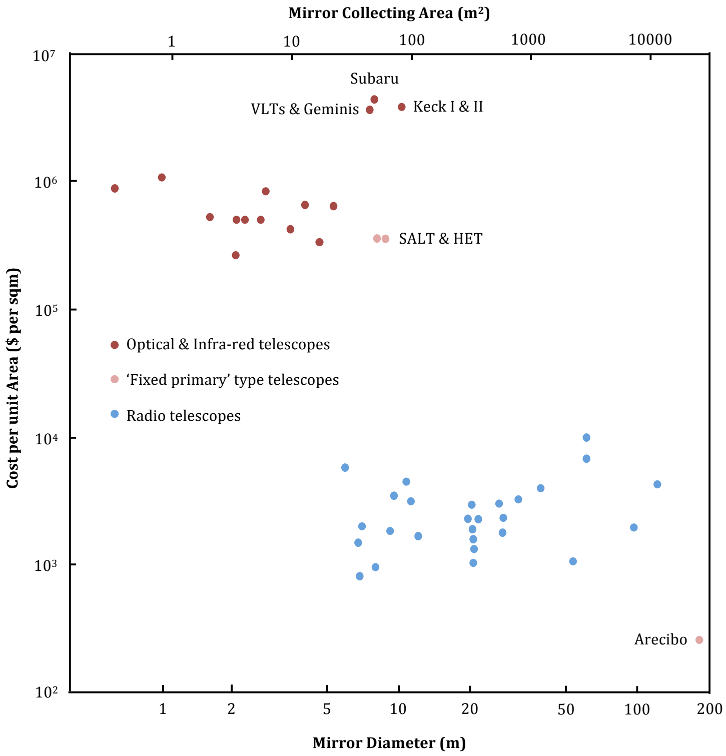

Figure 1.10: Telescope costs versus primary mirror diameter (from Buckley et al., 2000).

primary mirror in SALT also reduces costs, since a monolithic mirror becomes highly complex at 5m and unfeasible beyond 8m in diameter due to difficulty in manufacturing and properly cooling the glass. Monolithic mirrors are vastly heavier requiring expensive control and mounting systems and causing them to sag under their own weight. The lighter mosaic structure in SALT also carries the benefits of reduced transport, installation and maintenance costs. SALT was built at a cost of ~\$15M, roughly a fifth of the cost of a similarly sized telescope such as a Gemini, Keck or VLT (see Figure 1.10).

1.3.4 SALT Optical Error Budget

Buckley (1995) gives the median seeing conditions at Sutherland as FWHM = 0.9". When added in quadrature with the expected image quality of the SALT optics, the total SALT image quality error budget predicts that the encircled energies $EE(50) = 1.29''$ and $EE(80) = 2.15''$. This can

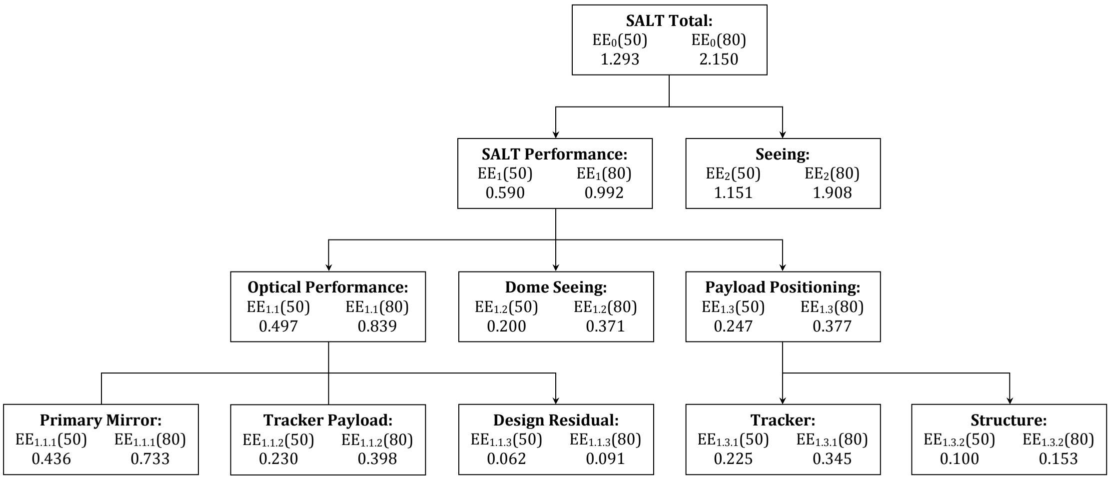

Figure 1.11: SALT optical error budget using site median seeing values converted to encircled energies for a zenith angle of 37° and wavelength of 633nm (from Swart, 2001).

be seen in the error budget shown in Figure 1.11 (Swart, 2001). Although early observations with SALT rarely achieved these specifications, this was traced to a poorly engineered mounting system for the SAC. Extensive re-work on the SAC has recently been completed (Summer 2011) and the telescope is close to achieving the design performance (O'Donoghue et al., 2010).

1.4 SALT High Resolution Spectrograph

1.4.1 Spectrograph Overview

At the time of writing, SALT HRS is one of three first generation instruments for SALT, the others being already in the commissioning phase and include an imaging camera (SALTICAM) and a Fabry-Perot imaging spectrograph (Robert Stobie Spectrograph). There is also a visitor instrument in the form of the Berkeley Visible Imaging Tube (a photocathode based instrument for microsecond time resolution photometry with various filters). The Robert Stobie spectrograph has an imaging spectrograph, spectropolarimetry potential and multi-object capability with resolutions between $R \sim 800$ to 13000.

SALT HRS is designed to complement these capabilities by providing high resolution, seeing-

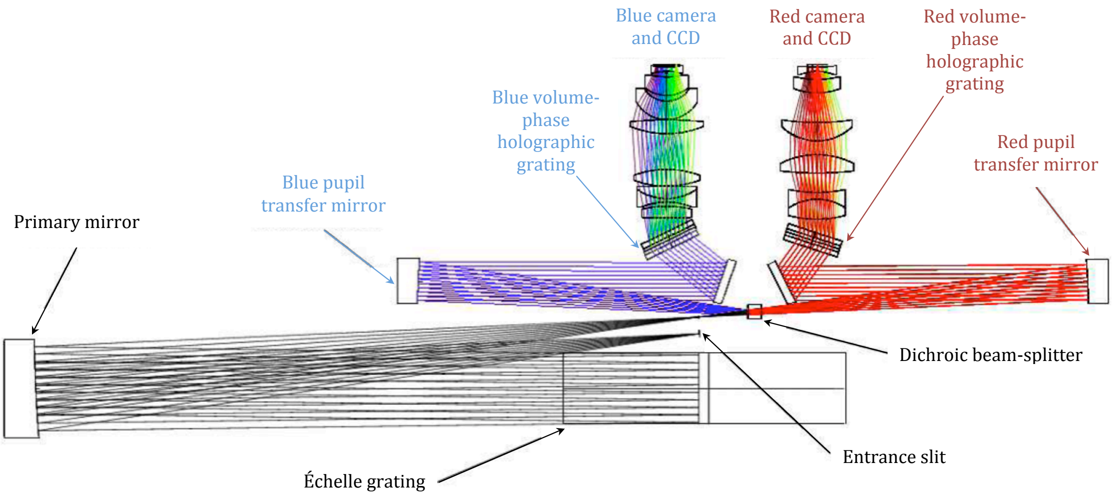

Figure 1.12: A ray trace diagram of the main instrument optics for SALT HRS - pre-slit optics not included (from Barnes et al., 2008).

limited spectroscopy. The design is of high relative efficiency and is fed with 50m long optical fibres from the fibre instrument feed (FIF) on the SALT tracker payload at the prime focus. SALT HRS uses a single Newport R4 échelle grating (blaze angle $\theta_B = 76°$) with 41.59 grooves/mm as the main dispersion element, which is fabricated as a replication mosaic of two identical 214mm × 400mm gratings on a single substrate. It is illuminated with a 200mm diameter beam at the nominal blaze angle of 76°. The pre-dispersion beam from the fore optics and entrance slit is collimated by the primary mirror onto the grating. This mirror is part of a so-called 'white pupil' design and also acts to bring the dispersed beam to an intermediate focus, de-magnifying the pupil to 111mm thereafter, which allows for more feasible cross-disperser and camera dimensions. A dichroic is used as a beam-splitter to separate two distinct 'blue' and 'red' channels operating at $370 \leq \lambda \leq 555$nm and $555 \leq \lambda \leq 890$nm respectively. Each channel has its own optimized pupil transfer and fold mirrors as well as a wavelength-tuned volume phase holographic (VPH) grating for cross-dispersion of the échelle orders. The dispersed spectra in each channel are imaged onto wavelength-specific charge coupled device (CCD) detectors by two camera systems. Figure 1.12 shows a ray-trace diagram of the main SALT HRS optical elements as described above. The pre-slit optics (input optics, image slicers and transfer optics) are not included.

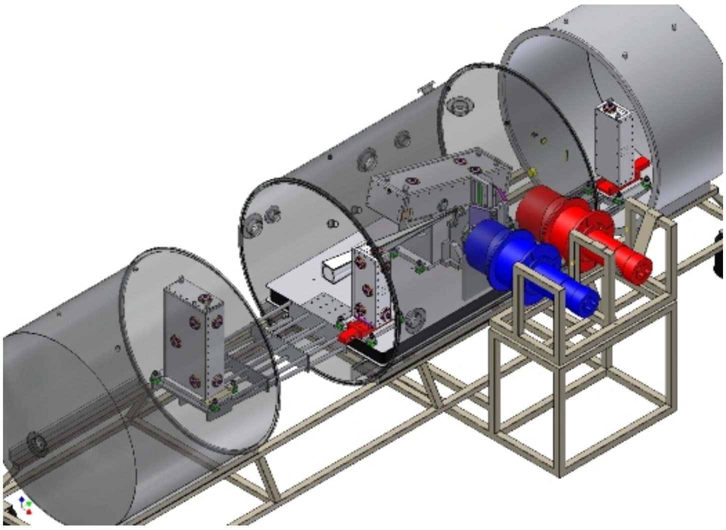

Figure 1.13: A CAD diagram of SALT HRS mounted inside the vacuum tank, with the colour-coded red and blue detectors 'hanging' from the side of the vessel.

The spectrograph optics are contained within a vacuum tank at ∼5mbar to minimize the destabilizing effects of temperature and pressure changes. The tank in turn is housed in a temperature-stabilized environment in the instrument room beneath the telescope. A closed loop control system with temperature sensors internal to the tank and an external heater element system maintain a consistent temperature slightly above the enclosure ambient temperature. Figure 1.13 shows the instrument mounted to an optical bench and inside the vacuum tank, with the colour-coded red and blue detectors mounted from the side of the tank. The two cameras straddle the tank wall and couple directly to their respective CCDs using the final camera lens element as a window to the detector dewars.

Much of the optical design of SALT HRS was undertaken at the University of Canterbury's Department of Physics and Astronomy in New Zeala[illegible] Dr. Stuart Barnes. However, some optical design modifications, the mechanical design, testing, assembly and integration [illegible] performed at the Centre for Advanced Instrumentation in the Physics Department of Durham University in the United Kingdom. Professor Ray Sharples is the principal investigator (PI).

Table 1.1: Summary of the four key SALT HRS operational modes

	Low-Resolution	Medium-Resolution	High-Resolution	High-Stability
Fibre core (µm)	500	500	350	350
Resolving Power	16000	37000	67000	67000
Image Slicing	None	Bowen-Walraven	Bowen-Walraven	Bowen-Walraven
Mode Scrambling	None	None	None	Double scrambler
Nod & Shuffle	Optional	None	None	None
Iodine Cell	None	None	None	Optional
Simultaneous Th-Ar	None	None	None	Optional[1]

1.4.2 Spectrograph Modes

SALT HRS has four key operational modes: low-, medium- and high-resolution modes plus a high-stability mode. All modes are selectable at any time depending on the observers needs. The trade-off to gain increased spectral resolution in the higher resolution modes is a decrease in instrument throughput. The high-stability mode offers the same resolution as the high-resolution mode but with increased wavelength stability for high precision Doppler studies, for example in indirect Doppler spectroscopy exo-planet detection. It adds a fibre double scrambler system to increase long-term spectral stability by ensuring complete scrambling of the light (particularly important given the variable pupil illumination of SALT as discussed in Section 1.3.2) as well as the user-selectable option for a retractable iodine cell to increase wavelength calibration precision (Marcy & Butler, 1992). Different resolutions are achieved using different combinations of fibre core sizes (350µm for the high-resolution and high-stability modes and 500µm for the low- and medium-resolution modes) in combination with image slicing (see Section 2.2.5). The four key modes are summarized in Table 1.1.

1.4.3 Example Spectra & Spectral Formats

Figure 1.14 shows a simulated blue spectrum created by Stuart Barnes, using his 'EchMod' software (Barnes, 2005, PhD University of Canterbury, New Zealand). Figures 1.15 and 1.16 show the spectral coverage of the red and blue SALT HRS detectors respectively. Note that

---

[1] The iodine cell and simultaneous Th-Ar mode cannot both be used at the same time.

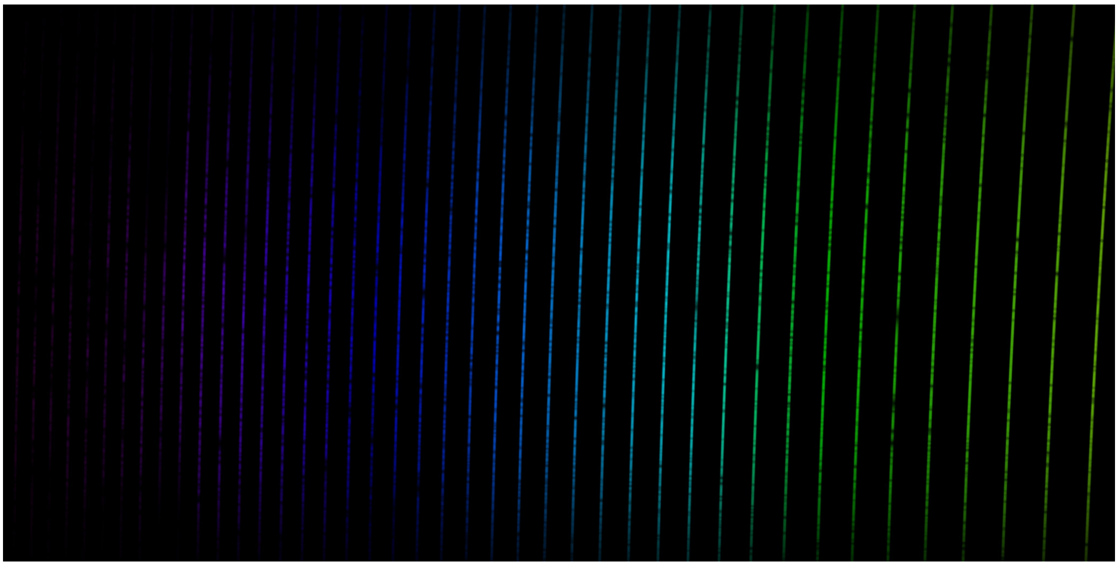

Figure 1.14: A simulated blue channel cross-dispersed échelle spectrum created using EchMod modeling code (Stuart Barnes, private communication), with the representative colours added to aid understanding.

since both detectors feature the same 15µm pixel size, the 2048 × 4096 pixel blue detector can be only 50% of the area of the larger 4096 × 4096 pixel red detector, whilst still pr[illegible]ng full wavelength coverage. This is a result of careful instrument design choices and the fact that the angular extent of a free spectral range range is smaller for the bluer wavelengths (see equation 1.13).

1.4.4 SALT HRS Science Requirements

SALT HRS is an ideal instrument to compliment RSS in that it operates over a similar wavelength range (RSS covers $320 \leq \lambda \leq 900$nm whilst HRS operates between $370 \leq \lambda \leq 890$nm) but for single sources offers a higher spectral resolution range of $R \approx 16000$ to 67000. Based on a list of 50 science drivers for SALT HRS selected from a survey of potential users, the following are the main science categories to be addressed by the instrument:

- element abundance studies in local group galaxies;
- extra-solar planet detection;

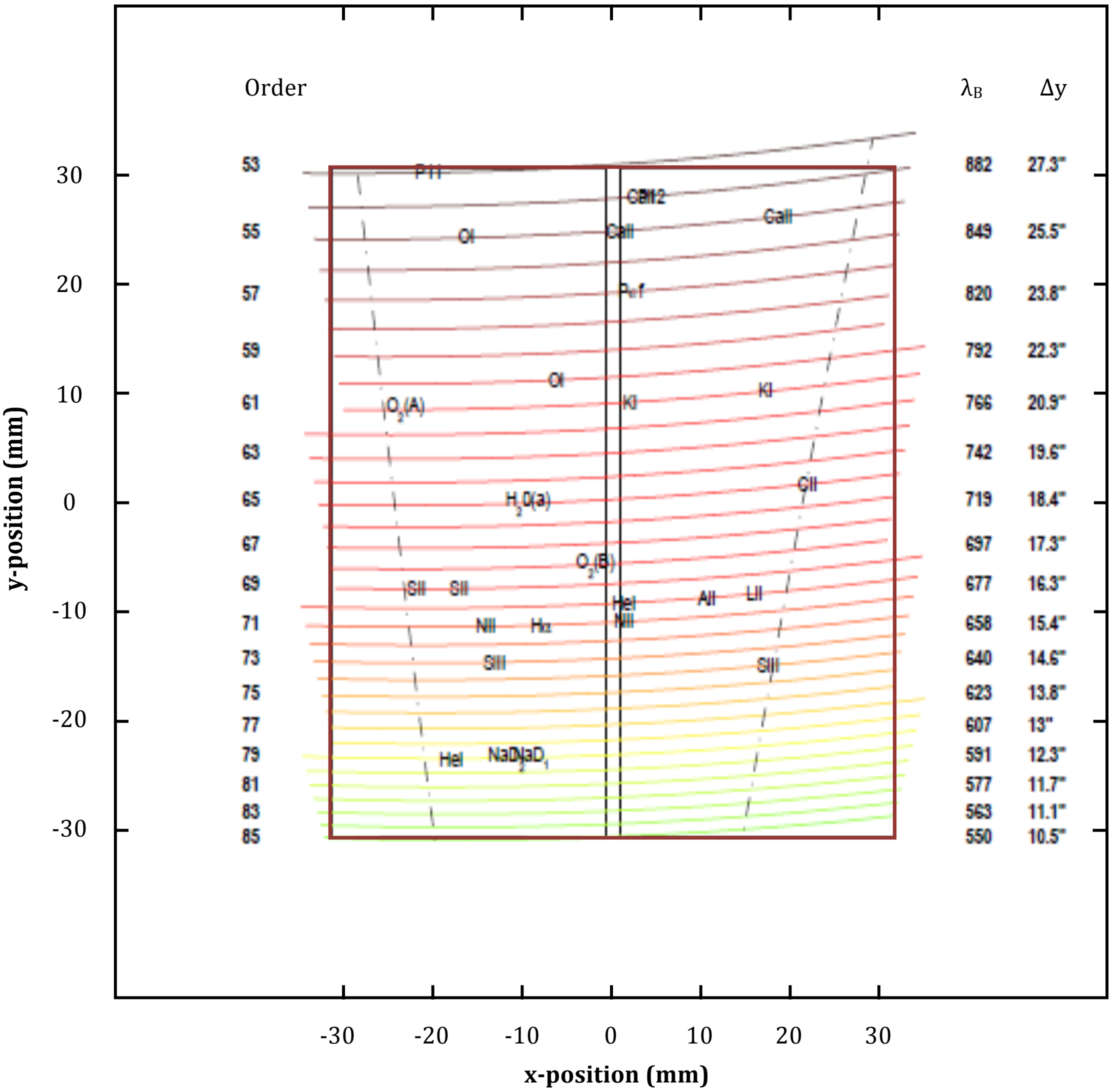

Figure 1.15: The simulated spectral format of the SALT HRS red arm with its larger square detector, showing the order number m, the blaze wavelength λ_B of each order, the inter-order spacing Δy and the position of some commonly used spectral lines. The dashed lines show the limits corresponding to one free spectral range in each order.

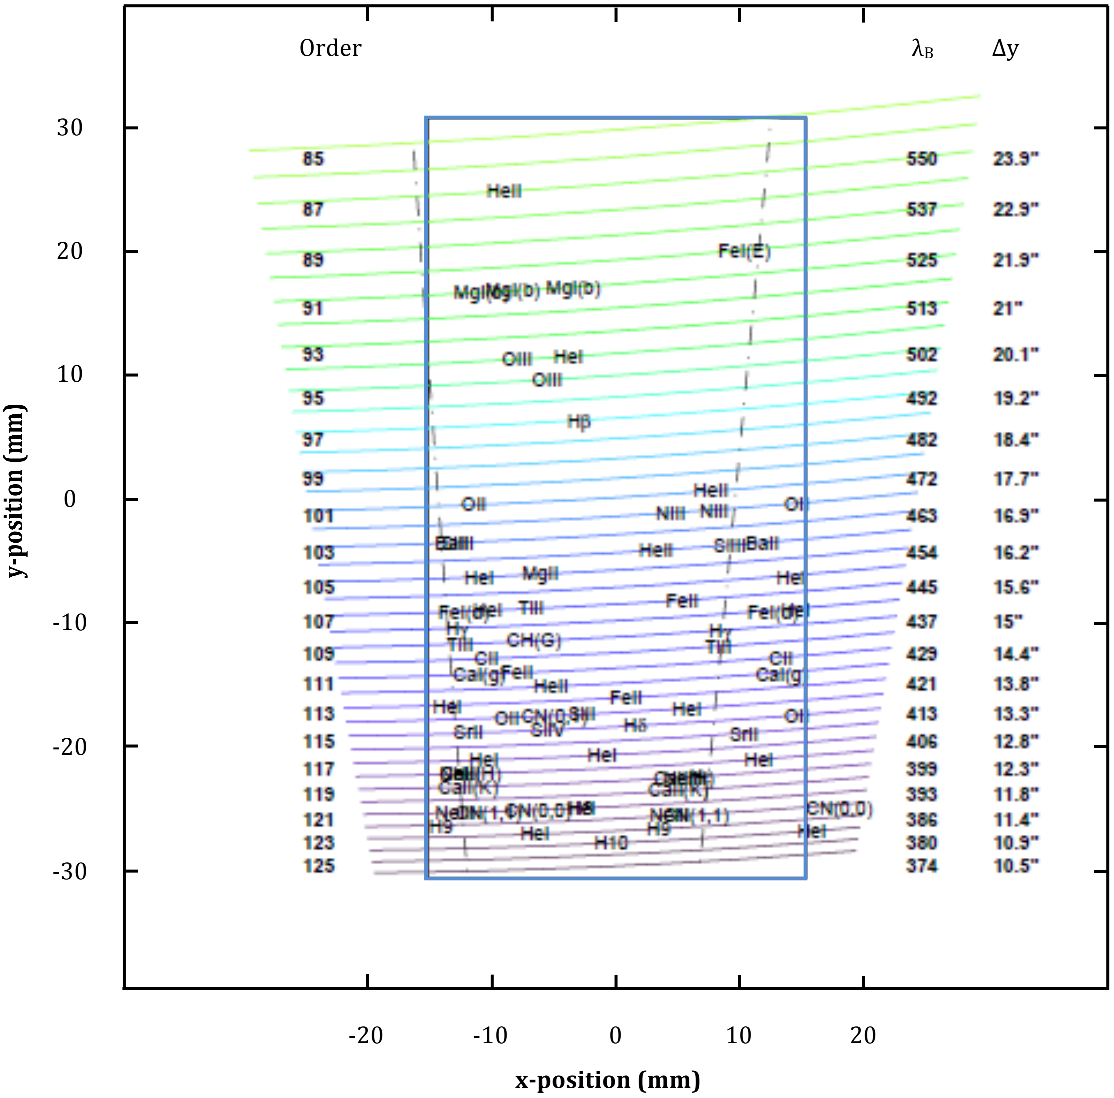

Figure 1.16: The simulated spectral format of the SALT HRS blue arm, with its narrower CCD area, again showing the order number m, the blaze wavelength λ_B of each order, the inter-order spacing Δy and the position of some commonly used spectral lines. The dashed lines show the limits corresponding to one free spectral range in each order.

- stellar internal structure and dynamics;
- star cluster and galaxy dynamics;
- outflow and accretion studies;
- high- and moderate-redshift galaxy spectroscopy.

This leads to a list of desired instrument capabilities which are:

- resolving power in the range $R = 16000$ to 65000;
- wavelength coverage $\lambda = 370$ to 890nm;
- high mechanical and thermal stability;
- limited multi-object capability.

SALT HRS is designed to be competitive with similar instruments on telescopes of the same class as SALT. See Tables 1.2 and 1.3 for examples. Since some instruments have multiple spectral channels, data relating to the blue channel appears in blue text whilst figures for the red channels appear in red. Where multiple configurations are possible, data [illegible] displayed side-by-side.

1.5 Summary

This introductory chapter has discussed the development of spectroscopy and its connection with, and importance to modern astronomy. The fundamentals of diffraction gratings and the theoretical performance and design issues of échelle gratings have been outlined. A description of the SALT telescope with emphasis on its unique design and tracking was presented, and the SALT HRS spectrograph introduced with key details of its design, operational modes and science drivers.

In the rest of this thesis, we explore in more detail the as-built performance of SALT HRS, including performance evaluation and design trade-offs of critical systems such as the detectors, fibre system and optical components. We also present solutions for key elements of the

Table 1.2: Comparison of SALT HRS with other spectrographs on large telescopes

Instrument	SALT HRS	UVES	HET HRS	HIRES
Telescope	SALT, Sutherland	Kueyen VLT, Paranal	HET, Texas	Keck, Mauna Kea
Size (m)	9.2 effective aperture	8.2	9.2 effective aperture	10
Illumination	Quasi-Littrow	Quasi-Littrow	-	-
Configuration	Dual beam, white pupil	White pupil	White pupil	-
Wavelength Range (nm)	370-555, 555-890	300-500, 420-1100	420-1100	320-1100
Resolving Power	16000, 37000, 67000, 67000	80000, 110000	30000, 120000	67000
Nod & Shuffle	Yes (LR only)	-	-	-
Radial Velocity Precision (ms^{-1})	3-4	2 (with Iodine cell)	≤ 10	-
Fibre Length (m)	50	-	-	-
Scrambler	Yes (HS only)	-	-	
Image Slicer	2 pairs Bowen-Walraven	Optional 3, 4, or 5 slice Bowen-Walraven	Yes	-
Échelle Type	R4 mosaic	R4	R4	R2.8
Grating Master	MR166	-	-	-
Groove Density (gr/mm)	41.6	41.6, 31.6	31.6	52.6
Grating Size (mm)	214 × 840 × 125	214 × 840 × 125, 214 × 840 × 125	210 × 836	-
Grating Mount	Face-down	Face-down	Face-down	-
Beam Size (mm)	200	200	-	305
Camera Type	Dioptric, dioptric	Dioptric, dioptric	Refractive	Catadioptric
Camera f/#	f/1.5, f/1.8	f/1.8, f/2.5	f/1.9	f/1.0
Camera Focal Length (mm)	166.8, 208.5	-, -	500	-
Cross-dispersers	VPH grating	Grating (×2), Grating	Grating	Grating
CCD Format	2k × 4k, 4k × 4k	2k × 4k, 4k × 4k	4k ×	2k × 2k
CCD Type	E2V44-82, E2V231-84	E2V44-82, E2V44-82/MIT ID-20 mosaic	Orbit mosaic	SITe
CCD Grade	1, 1	1, 1	-	1
CCD Pixel Size (µm)	15, 15	15, 15	-	24
CCD Coating	Broadband astronomy, extra red plus	-, -	-	-
Calibration	Continuum lamp, Th-Ar, simultaneous Th-Ar (HS only), Iodine cell (HS only)	Continuum lamp, Th-Ar, optional Iodine cell	Iodine cell	-
Environment	Vacuum chamber, temperature stabilized room	Nasmyth mount, passive enclosure	-	- -
Reference	Bramall et al., 2010	Dekker et al., 2002	Tull, 1998	Vogt et al., 1994

Table 1.3: Comparison of SALT HRS with other spectrographs on large telescopes

Instrument	SALT HRS	PEPSI	HARPS	HDS
Telescope	SALT, Sutherland	LBT, Arizona	La Silla	Subaru, Mauna Kea
Size (m)	9.2 effective aperture	11.8 effective (2×8.4)	3.6	8.2
Illumination	Quasi-Littrow	-	Quasi-Littrow	-
Configuration	Dual beam, white pupil	-	White pupil	-
Wavelength Range (nm)	370-555, 555-890	390-580, 580-1050	378-530, 533-691	320-1100
Resolving Power	16000, 37000, 67000	120000, 120000	120000	165000
Nod & Shuffle	Yes (LR only)	-	-	-
Radial Velocity Precision (ms^{-1})	3-4	-	≤1	-
Fibre Length (m)	50		38	-
Scrambler	Yes (HS only)	-	Yes	-
Image	2 pairs	-	No	-
Slicer	Bowen-Walraven			
Échelle Type	R4 mosaic	R4	R4	R2.9
Grating Master	MR166	-	MR160	-
Groove Density (gr/mm)	41.6	31.6	31.6	31.6
Grating Size (mm)	214 × 840 × 125	-	840 × 214 × 125	-
Grating Mount	Face-down	-	Side-on	-
Beam Size (mm)	200	200	-	305
Camera Type	Dioptric, dioptric	Dioptric, dioptric	Dioptric	Catadioptric
Camera f/#	f/1.5, f/1.8	f/2.3, f/2.3	f/3.3	f/0.96
Camera Focal Length (mm)	166.8, 208.5	-, -	728	-
Cross-dispersers	VPH grating	VPH grisms	Grism	Gratings
CCD Format	2k × 4k, 4k × 4k	4k × 4k, 4k × 4k	2 × 2k × 4k mosaic	2 × 2 × 4k
CCD Type	E2V44-82, E2V231-84	-, -	E2V	E2V
CCD Grade	1, 1	-, -	1	1
CCD Pixel Size (µm)	15, 15	15, 15	15	13.5
CCD Coating	Broadband astronomy, extra red plus	-,-	Single layer AR	-
Calibration	Continuum lamp, Th-Ar, simultaneous Th-Ar (HS only), Iodine cell (HS only)	-	Continuum Simultaneous Th-Ar Decommissioned Iodine cell	-
Environment	Vacuum chamber, temperature stabilized room	-	10^{-2}mbar vacuum chamber Insulated room	-
Reference	Bramall et al., 2010	Pallavicini et al., 2003	Pepe et al., 2000	Noguchi et al., 2002

observing system such as the exposure time calculator (ETC) and quick-look data reduction software.

Chapter 2 gives a more in-depth breakdown of the individual optical components of SALT HRS, and their theoretical and measured throughput performances. This is compiled into an overall instrument throughput model as a function of wavelength, which is in turn used as the basis for an exposure time calculator written by the author.

Specifics of the SALT HRS fibre-feed, as well as the procedure for protecting and bundling the fibres is presented in Chapter 3. A background to optical fibres is given as well as the motivation for and results for a range of focal ratio degradation (FRD) measurements.

Chapter 4 presents the results of the specific sub-system tests. Details of the pupil mirrors are presented along with the results of a mechanical stability test. Additionally, a theoretical background to volume-phase holographic gratings is given and compared to results of diffraction efficiency tests for both red and blue cross-dispersers.

An overview of astronomical CCDs and their differences from consumer products may be found in Chapter 5. Details of the SALT HRS detectors are presented with a description of the blue detector cleaning process and results. The blue detector is characterised with both photon transfer functions and quantum efficiency curves. The second portion of this chapter deals with cosmic ray events which are a significant noise source for long exposures of faint targets. I present the results of a variety of lead shielding tests in an effort to characterize the properties of 'cosmic ray' events on the detectors.

Chapter 6 addresses the data which will be produced by SALT HRS. Firstly, the means of properly calibrating the data with an accurate wavelength scale is discussed. A software pipeline for the reduction of échelle data is then introduced.

Finally, Chapter 7 gives the conclusions of this thesis and an overall assessment of the SALT HRS project.

Chapter 2

Throughput Model & Exposure Time Calculator

2.1 Model Rationale

It is crucial to have an accurate picture of the expected instrument throughput (as a function of wavelength) both during the design and development stages, and also post-assembly. Since it is often necessary during design to conform to budget, dimension or efficiency constraints, a throughput model can be used to predict what effect any changes will have on overall instrument performance. It is also vital post-assembly for confirming the actual instrument performance against that predicted. Moreover, knowledge of throughput is advantageous as a general rule of thumb during design to ensure performance is comparable or better than similar instruments and suitable for the science it is planned for (see for example Dekker et al., 2000 & Kaufer et al., 1997). The physical layout of the major optical components (with the exception of the pre-slit optics and CCD detectors) on the SALT HRS optical bench is shown in Figure 2.1.

In this chapter a throughput model is presented based on the measured or theoretical efficiencies

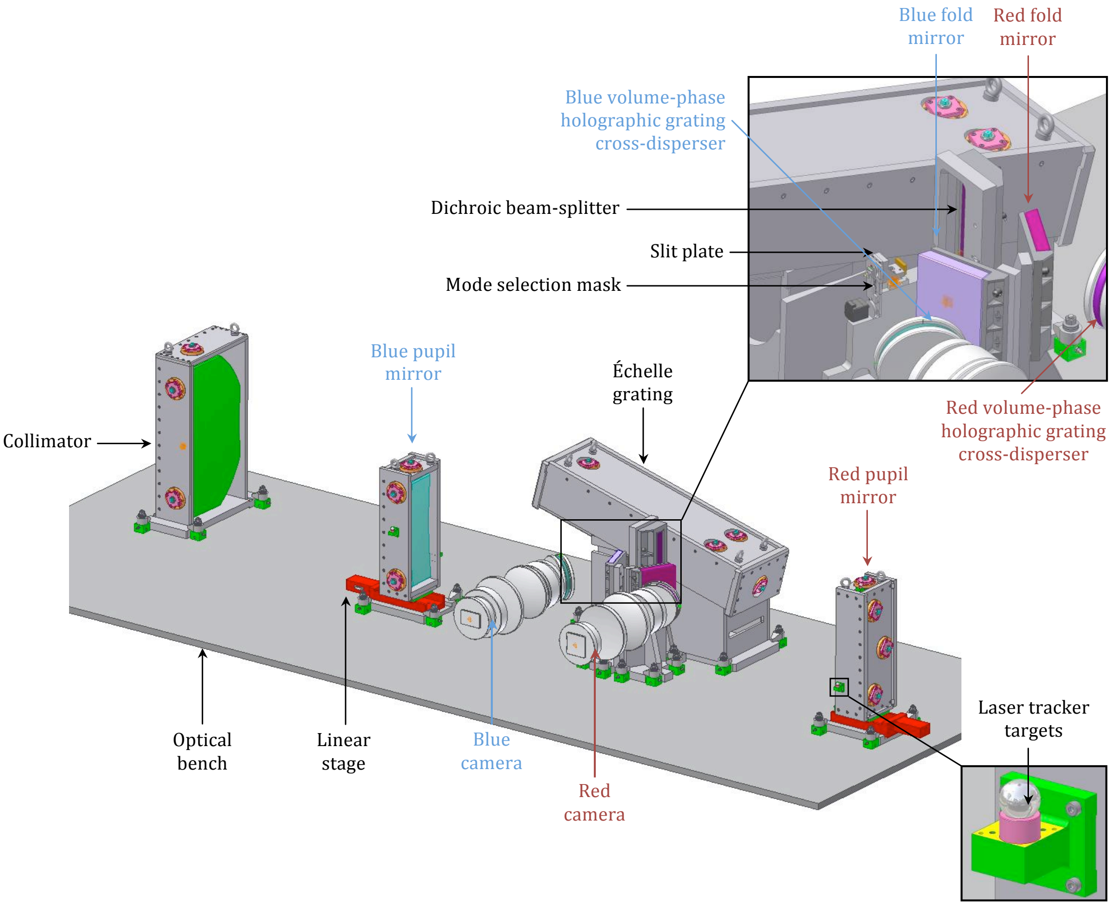

Figure 2.1: A CAD diagram of the major optical component groups (apart from pre-slit optics and detectors) as located on the optical bench. The close-up views show the compact layout around the dichroic, and laser tracker targets used for focusing the pupil mirrors.

of the SALT HRS components. The model uses wavelength bins of 50nm between 400nm and 890nm, and 10nm bins blueward of 400nm, where losses in many components quickly compound to produce a more rapid fall-off in performance. The image slicer and camera components are discussed in somewhat more detail than the other component groups. Although tangential to the main purpose of this chapter, unlike the other components described, they are not discussed in great detail anywhere else in this thesis.

2.2 Source Data

2.2.1 Optical Path and Sources of Loss

Although the optical path of the SALT HRS instrument is briefly summarized in Chapter 1, a 'block diagram' of the major component groups in their input-to-output order may be seen in Figure 2.2. The colour-coding indicates in which modes the respective components are used, with a white background denoting a component group found in all four instrument modes.

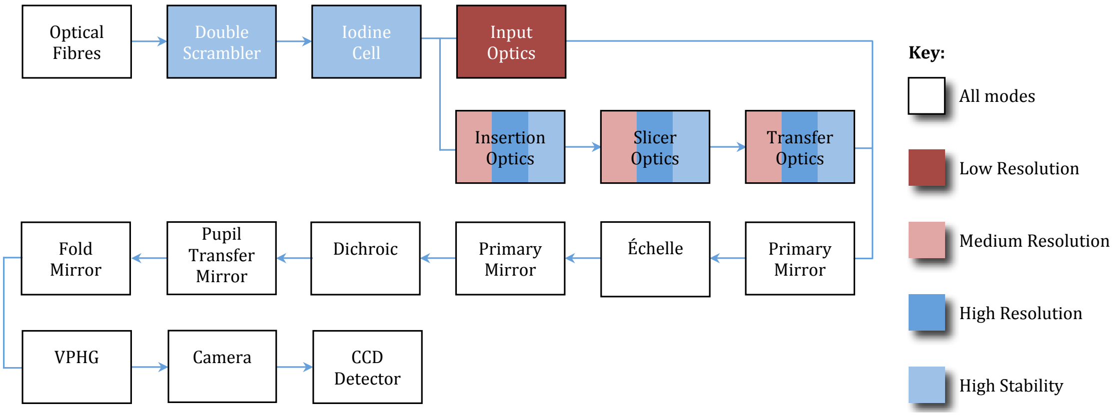

Figure 2.2: A block diagram of the SALT HRS optical path, showing the major component groups to be discussed in this chapter. Colour-coding indicates in which mode(s) the respective component groups are used.

The telescope efficiency is not included in this throughput model, which is specific to the instrument components alone. However, a plot of SALT efficiency as a function of wavelength is included in Figure 2.3 for the sake of completeness. The object light enters the telescope, and is fed into the spectrograph via optical fibre. These fibres have two sources of loss - transmission losses and losses due to the effect of focal ratio degradation (Ramsey, 1988). Specific to the high-stability mode (which is designed to maximize wavelength calibration precision at the expense of efficiency), there will be an additional loss due to the presence of the fibre double scrambler optics. Losses in the high-stability mode may increase further still should the observer add the optional iodine cell into the double scrambler air gap.

In low resolution mode, where throughput is maximized as far as possible, light from the fibres is

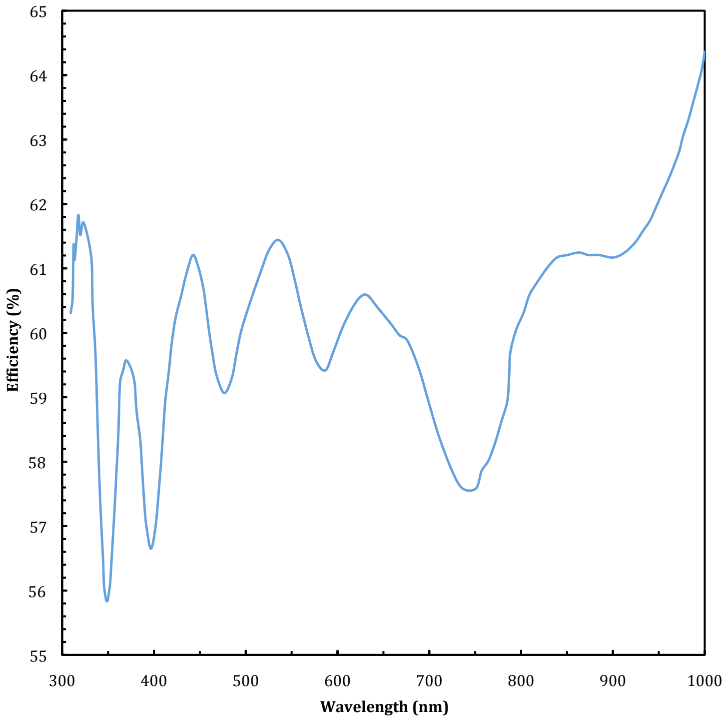

Figure 2.3: A plot of SALT telescope efficiency as a function of wavelength. The rapid fluctuations in efficiency at the ±2% level are due to the anti-reflective coatings on the mirrors and optics. The efficiency is calculated using witness sample measurements of the mirrors, including the four SAC mirrors, primary mirror and a fifth fold mirror. This curve does not however include the atmospheric dispersion corrector (ADC), which is expected to achieve greater than 95% efficiency (from Buckley, 2004).

transferred to the slit using a simple set of mirror and lenses. The medium- and high-resolution modes and high-stability modes must undergo image slicing. Light is transferred to the image slicer via a set of input optics, with the sliced fibre image conveyed to the slit using further transfer optics. The image slicers are responsible for both transmission losses and geometric losses. Light emerging from the slit propagates to the primary mirror, which acts to collimate the light onto the échelle grating, which itself causes significant dispersion losses as well as a smaller reflection loss. Spectra reflected from the grating return to the primary mirror for a second pass, which acts to focus the light onto the dichroic. Light which is r[illegible]d of 555nm is transmitted, whilst bluer light is reflected. At this point the instrument is separated into a blue arm and a red arm. In both cases, light from the dichroic is reflected from a pupil transfer mirror (which acts to re-collimate the beams), through a fold mirror and into a volume phase holographic grating (VPHG) which acts to cross-disperse the orders of spectra produced by the échelle. The cross-dispersed spectra are then imaged onto the appropriate CCD detector using a fully transmissive camera system. The VPHGs, cameras and detectors are optimized for their respective arms, whilst all other optical components are tuned to function across the full SALT HRS operational wavelength range as best as possible.

2.2.2 Optical Fibres

The PolyMicro FBP type fibres used in SALT HRS were selected for their good performance across the full wavelength coverage of the instrument, in addition to the option to procure custom drawn core diameters and for their focal ratio degradation (FRD) performance. The manufacturer provides transmission data for the FBP fibres across a range of wavelengths[1]. The attenuation of the fibre in dBkm^{-1} is tabulated as a function of wavelength, and may be converted to the transmission through 50m of fibre using equation 2.1. Figure 2.4 shows the transmission through a 50m length of fibre as a function of wavelength.

$$\text{Throughput} = \frac{1}{10^{\frac{(\text{Attenuation} \times 0.050)}{10}}} \tag{2.1}$$

In addition to transmission losses, a decrease in throughput may be expected due to focal ratio degradation (Ramsey, 1988). Measurements of Polymicro FBP fibres (400µm and 600µm core sizes) suggest that throughput of 95.7% or greater may be expected at any wavelength with an input beam speed of f/3.85. A more detailed treatise on these experimental tests and results can be found in Chapter 3. There is also a Fresnel loss associated with light reflecting back from the fibre input faces, which is not included here. Such a loss is expected to be minimized by the use of refractive index matching gel coupling at a fibre face-glass interface or using a suitable optical cement when bonding to a lens.

---

[1] http://www.polymicro.com/products/opticalfibers/products_opticalfibers_fbp.htm

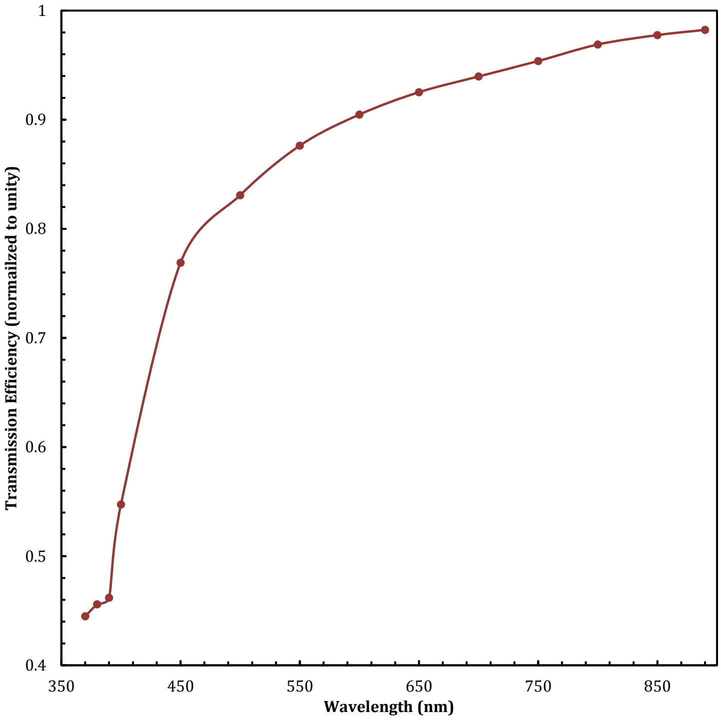

Figure 2.4: The transmission losses which can be expected through a 50m length of Polymicro FBP-type optical fibre. This curve does not include FRD or end coupling losses.

2.2.3 Double Scrambler & Iodine Cell

The double scrambler system ensures nearly complete scrambling in the optical fibres, enabling higher velocity precision for the high-stability mode. The merits, operation and design of such a system are discussed in greater detail in Section 3.3. The two lenses of the scrambler system are separated such that an iodine cell may also be inserted in the air gap, for optional use in the high-stability mode (see Section 6.2.2 for further details). The efficiency of these systems is influenced by four factors:

- Optical aberrations: the input fibre face is imaged onto an output fibre face. Therefore any rays outside of the target fibre core causes losses.
- Alignment error: misalignment of the two lenses will cause loss due to the same mechanism as for optical aberrations.
- Fresnel and absorption losses: at the scrambler lenses and cell windows.

- Numerical aperture losses: due to the extended input fibre core and the spacing between the lenses, the light cone hitting the target fibre is faster than the cone emitted by the input fibre. The input fibre feeds the system at the numerical aperture of the fibre, meaning that the output fibre is partly overfilled, leading to lossy modes.

The simulated effect of these factors can be seen in Figure 2.5 at 370, 550, 700 and 890nm, as well as the total throughput caused by these contributions, with and without the iodine cell.

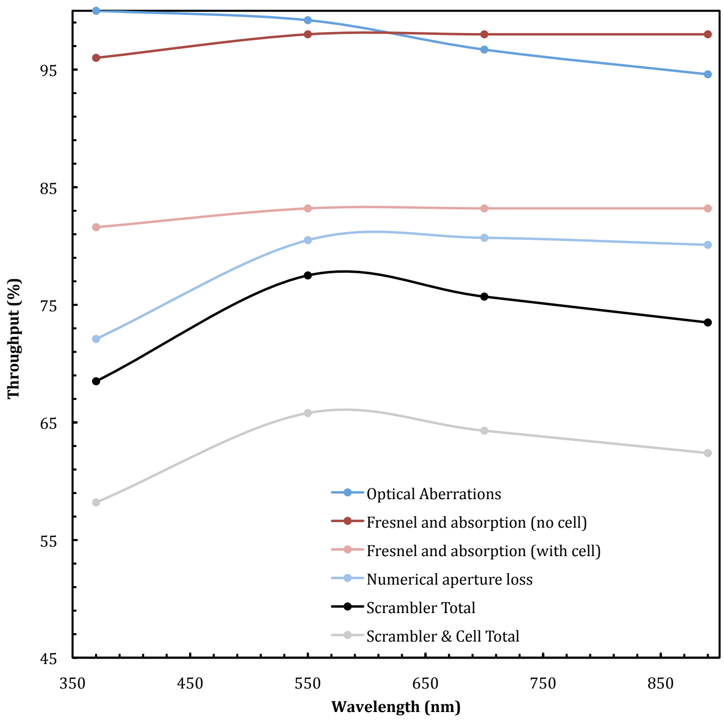

Figure 2.5: A graph of the double scrambler system efficiency, with and without the iodine cell inserted. Also shown on the same axes are the various contributions to these losses.

2.2.4 Input Optics

The low-resolution mode is optimized for maximum throughput. This has been achieved by transferring unsliced light as efficiently as possible from the fibre output to the slit plate. To achieve this, transfer optics have been designed with the minimum number of components and

Table 2.1: Low-Resolution Mode Input Optics Properties

Component	Glass Type	Central Thickness (mm)	Optical Surfaces
Lens 1.1	N-FK51	35	0
Lens 1.2	S-LAH66	2	1
Lens 2	N-FK51	18	2

surfaces as possible (see Table 2.1). A doublet lens is cemented directly to the fibre output face, and a third lens and fold mirror transfer the light to the slit entrance. The transmission losses are calculated for the central thickness of all three lenses using their appropriate glass type; anti-reflective coating properties are also included for the three air-glass surfaces; as is the reflectivity of the injection fold mirror (see Figure 2.6).

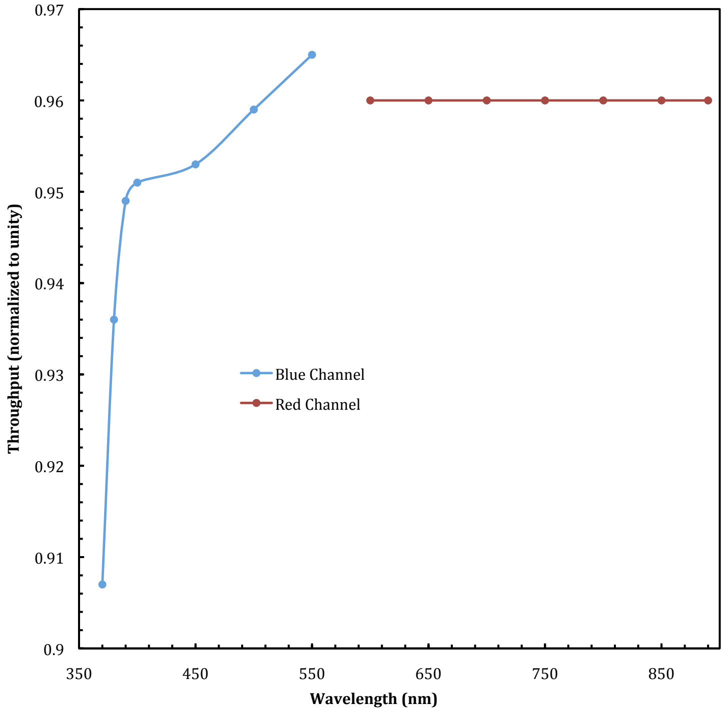

Figure 2.6: A graph of the low-resolution mode input optics total throughput as a function of wavelength, with the wavelengths corresponding to the blue and red spectral channels plotted accordingly.

2.2.5 Slicer Optics

Figure 2.7 shows the entrance aperture coupling of an optical fibre as a function of the fibre diameter, for a variety of stellar point spread functions (PSF) between 0.5arcsec and 2.5[illegible] (assuming the object PSF is described by a Moffat function and is added in quadrature to the expected SALT telescope image quality of $EE(80)$=2.15”[2]). When using a single fibre, particularly when of smaller diameter, the object coupling is low. For this reason, image slicers are employed to enable higher spectral resolutions in the three higher resolution modes. In SALT HRS, the image slicers are of the modified-Bowen-Walraven type (Bowen, 1938 & Walraven, 1972), similar to those used in FEROS (Kaufer, 1998) and UVES (Dekker et al., 2003) spectrographs. The modified design allows for slicing of two fibres simultaneously (for example sky and object fibres), both with identical optical path lengths through the slicing elements (see Figure 2.8a). The

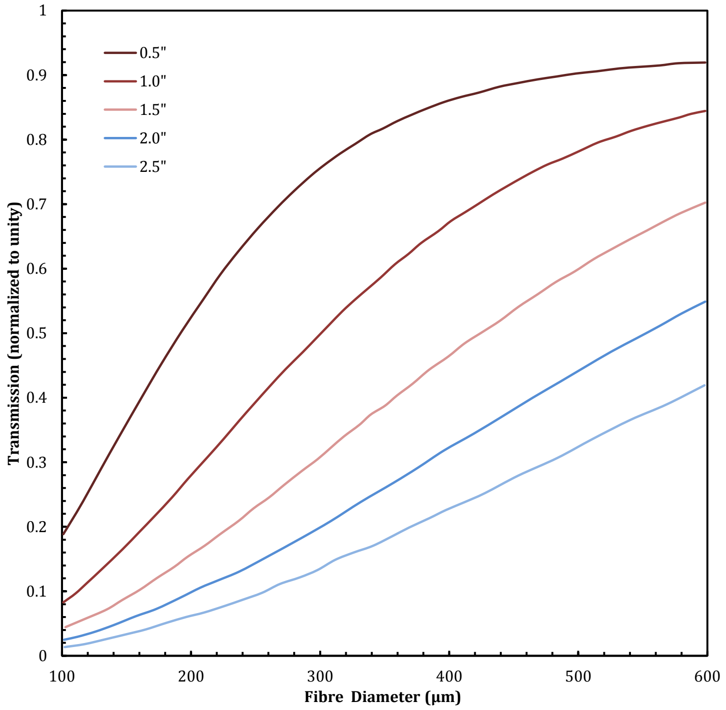

Figure 2.7: Fibre entrance aperture efficiencies for SALT HRS (assuming a Moffat function object PSF added in quadrature to the nominal SALT image quality) for a variety of seeing conditions between 0.5” and 2.5”.

---

[2] This is a nominal design value from the SALT image quality error budget. During early observations SALT has failed to achieve this value (by some margin). However recent work on the SALT spherical aberration corrector appears to have remedied this problem (see http://www.salt.ac.za/public-info/recent-highlights-observations/).

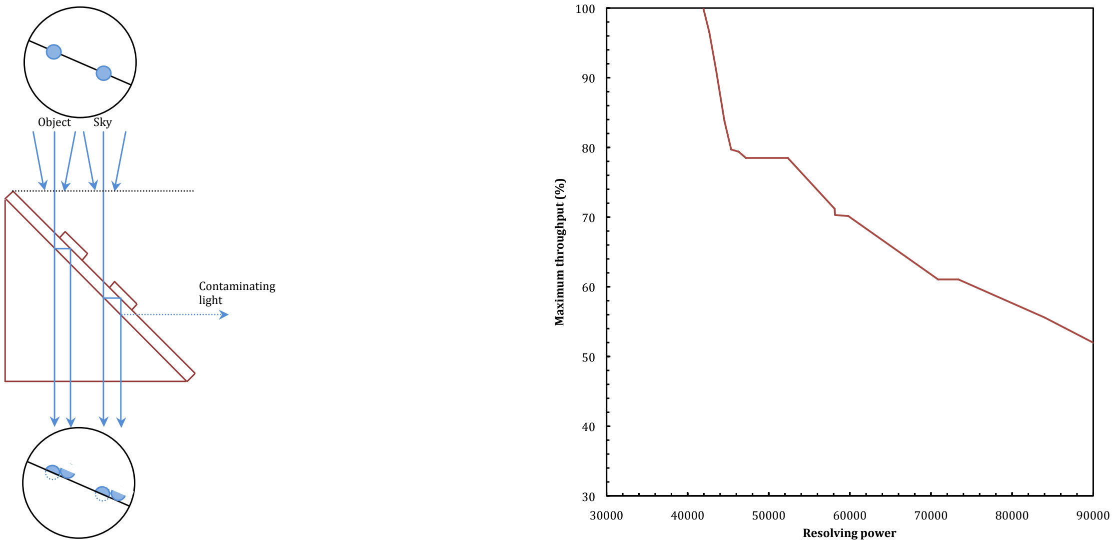

(a) Optical path in a slicer

(b) Geometric slicer throughput

Figure 2.8: The optical path followed in a modified-Bowen-Walraven image slicer (from Kaufer, 1998) and the geometrical throughput of such a system as a function of resolving power. The steps in this function arise due to the finite number of fibre core sizes and the integer number of image slices (courtesy M. Albrow).

absence of optical path differences between the sliced fibre images reduces the defocus introduced by the image-slicer. This defocus effect is minimized further by using the slicer fore-optics to convert the fast f/3.8 fibre output to a much slower focal ratio of ∼f/25 in the slicer optics. Figure 2.8*b* shows the geometrical throughput of a Bowen-Walraven type image slicer system as a function of resolving power. The throughput for a fixed object plus sky observation is shown.

Four lenses form the slicer insertion optics and precede the image slicer, which consists of the slicer optics and a base plate (see Figure 2.9). Following this, a further four lenses are used to transfer the sliced beam to the slit. In total there are ten transmissive components including the slicer and base plate, featuring thirteen air-glass surfaces. In addition to the extra transmission and reflection losses of the more complex slicer feed, geometric losses occur within the slicer itself which contribute losses of ∼20% to the medium- and high-resolution and high-stability modes (see Table 2.3). Transmission data for the lenses is taken from the O'Hara and Schott[3] glass

---

[3] http://www.oharacorp.com/ & http://www.schott.com/advanced_optics/

manufacturer data sheets, which is used to calculate transmission of individual lenses at their central thickness. The transmission qualities of the anti-reflection coatings are also included for each air-glass surface and are taken from measured witness sample data provided by the lens manufacturer (Spanoptic[4]). The performance of the coating varies slightly with the refractive index of the glass to which it is applied, and therefore there may be a slight deviation in performance between lenses. However, such differences are negligible given the very high performance of the coatings. The overall throughput of the medium- and high-resolution and high-stability mode input optics can be seen in Figure 2.10.

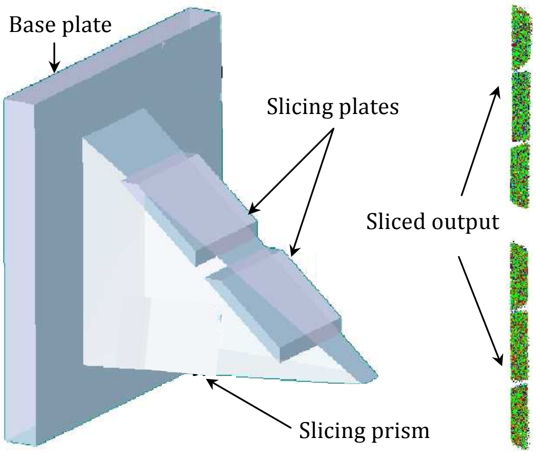

Figure 2.9: A diagram of the image slicers and the output pattern produced by such a slicer (from Barnes, 2008). The image slicer is of the modified-Bowen-Walraven type, similar to those used in FEROS and UVES (Kaufer, 1998).

2.2.6 Mirrors

The instrument design comprises four mirror surfaces, with an additional fifth surface (a small fold mirror) in the case of the low-resolution mode. The primary mirror is first used as a collimator from slit to grating, but also acts in a second pass - known as a ‘white pupil design’ spectrograph (Dekker et al., 2000) - to re-image the diffracted spectra to a point where a dichroic can be used to separate the light into red and blue spectral channels. Transfer mirrors re-collimate the beams following the dichroic, and a fold mirror is used to inject the beam into

---

[4] http://www.spanoptic.com/

Table 2.2: Medium- & High-Resolution and High-Stability Mode Insertion, Slicer & Transfer Optics Properties

Component	Glass Type	Central Thickness (mm)	Optical Surfaces
Lens 1.1	Silica	3.2	1
Lens [illegible]	N-FK51	4	1
Lens 2.2	S-LAL7	4	1
Lens 3.1	N-BK7	4	1
Bowen-Wolraven Slicer	N-BK7	8	0
Base Plate	N-BK7	3	1
Lens 4.1	N-FK51	19	2
Lens 4.2	S-LAL7	2.4	2
Lens 5.1	S-LAL7	2	2
Lens 5.2	N-FK51	20	2

Table 2.3: Throughput due to geometric slicer losses

Medium Resolution	High Resolution	High Stability
0.779	0.821	0.808

the VPHG cross-dispersers. For the low-resolution mode, an additional fold mirror is featured in the input optics design due to space constraints near the slit plate. The primary mirror coating is a broadband ‘Enhanced Silver’ coating provided under sub-contract from the mirror manufacturer KiwiStar by Advanced Telecommunications & Optical Coating[5], who also provide a ‘Protected Silver’ coating for the low-resolution injection mirror. The transfer and fold mirrors have coatings optimized to their specific wavelength range of operation. The reflective efficiency of the broadband primary mirror coating as a function of wavelength can be seen in Figure 2.11. This coating provides a high and reasonably flat performance redward of 400nm but displays a significant fall-off in performance blueward of this. At the time of writing the primary mirror coating is failing to meet specifications in the blue wavelength regime. The reflectivity of the blue and red specific coatings for the respective pupil and fold mirrors can be seen in Figure 2.12.

---

[5] http://www.kiwistar.com/ & http://www.opticalsputtering.com/

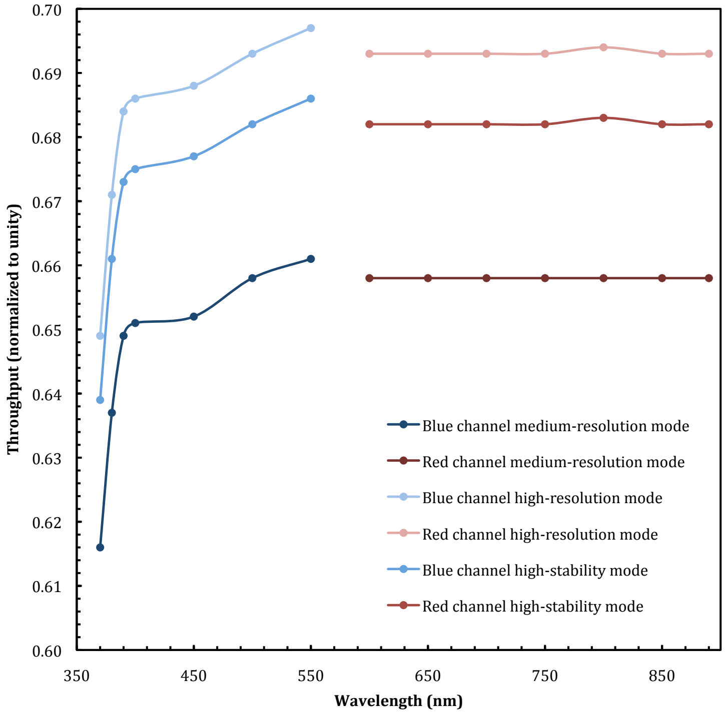

Figure 2.10: The modelled throughput of the medium- and high-resolution and high-stability mode input optics, including insertion, image-slicing and transfer components.

2.2.7 Échelle Grating

The échelle grating performance is interpolated from four data points provided as test results in the grating manufacturer's report from Newport Corporation[6]. The points indicate a minimum efficiency of 67% at 360nm; 65% at 400nm; 63% at 550nm and 62% at 900nm. The interpolated data is shown by the solid line in Figure 2.13, and was produced using a simple low-order polynomial fit to the data points provided.

2.2.8 Dichroic

The dichroic acts to separate light into the red and blue arms in order that the following optics and detectors may be better optimized for performance over a narrower range of wavelength operation. It is specified to act in reflectance blueward of 555nm and in transmission redward of

---

[6] http://gratings.newport.com/home

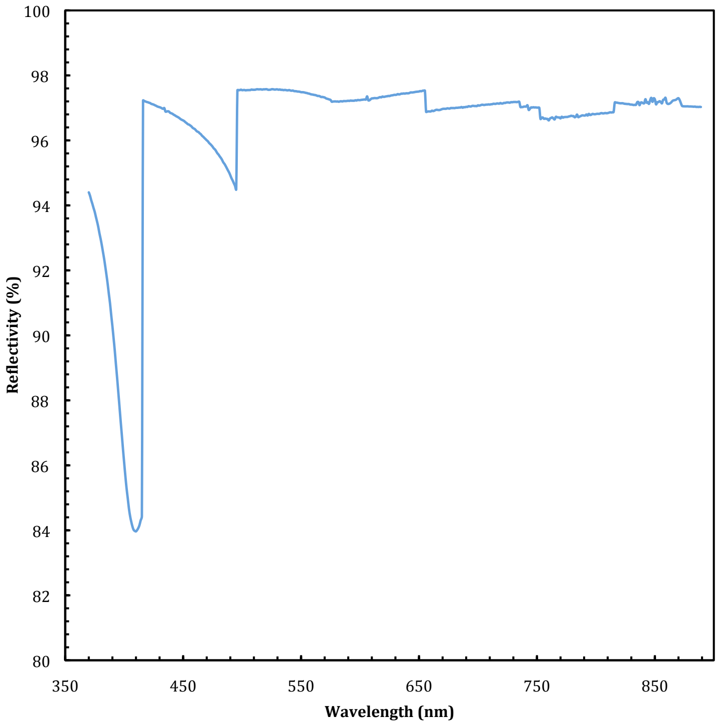

Figure 2.11: The broadband 'Enhanced Silver' mirror reflectivity of the primary mirror as a function of wavelength, as measured by the coating sub-contractors, Advanced Telecommunications & Optical Coating.

this. Transmission data is provided by the manufacturers, Asahi Spectra[7], for a collimated beam with a normal angle of incidence. Since their experimental setup would not allow for accurate measurements at 5.5° angle of illumination, the measured 0° values were used to calculate the theoretical transmission at 5.5° angle of incidence (which is the angle at which the dichroic will be used in SALT HRS). On first glance it might be expected that the blue performance is simply 100 minus the red performance (e.g. if the blue transmittance is 3% for a given wavelength of 542nm it may be expected that $100 - 3 = 97\%$ of 542 nm light is reflected). In reality this value is reduced by approximately 4% since some light will pass through the substrate and backscatter from the uncoated rear surface. A plot of the theoretical dichroic transmittance as a function of wavelength at 5.5° angle of illumination may be seen in Figure 2.14.

---

[7] http://www.asahi-spectra.com/

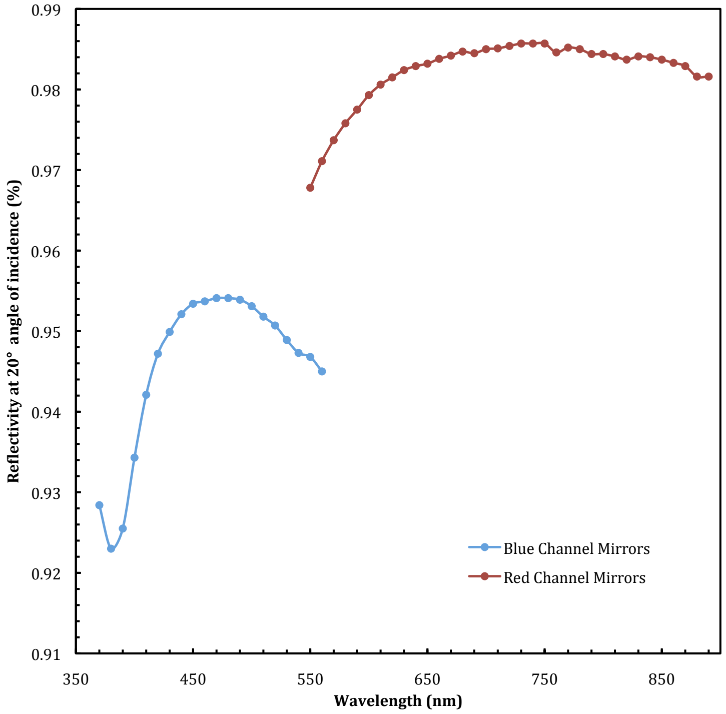

Figure 2.12: The witness sample measured reflectivity of the blue and red mirror coatings (as used in the pupil and fold mirrors) as a function of wavelength at 20° angle of illumination.

2.2.9 Volume Phase Holographic Gratings (VPHGs)

The VPHGs are used to cross-disperse the overlapping orders produced by the dispersion of the échelle grating. They disperse the spectra orthogonal to the dispersion direction of the échelle and separate each order for imaging on the detectors (see Section 4.2.3 for further explanation of cross-dispersion requirements). The VPHGs are sensitive not only to the input wavelength but also to the angle of incidence, gelatine layer depth and the contrast modulation index (the difference between the maximum and minimum refractive indices in the gelatine layer) (Baldry et al., 2004). See Section 4.2.2 for further details on the theory of VPHG diffraction.

Many of the components in the optical train have a fairly flat performance across the majority of the wavelength range, falling off in the extreme blue wavelengths. The two largest contributors to both throughput loss, and the profile of these losses with wavelength, are the dispersing elements. The ability to tune the throughput response of the VPHGs by varying the angle and contrast modulation has been the topic of much internal discussion. By modifying the grating for a peak

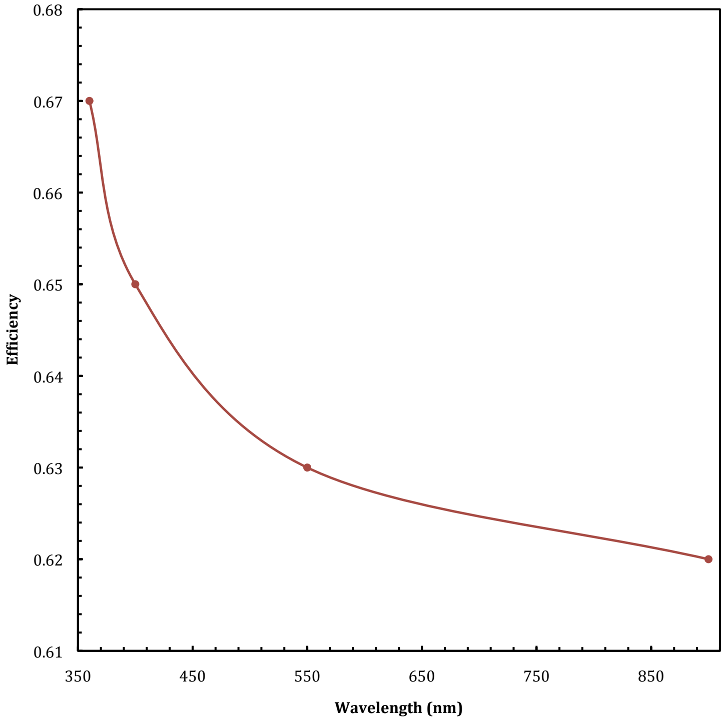

Figure 2.13: Interpolated efficiency from Newport échelle test data points.

efficiency at different wavelengths it would be possible to change the overall response of the instrument, for example to partly compensate for the rapid drop-off in the fibre transmission at short wavelengths (see Figure 2.4). Ultimately these tradeoffs are based on the science cases. For example it would not make sense to sacrifice efficiency across the blue arm solely to increase the performance between 370 and 400nm, if there are no spectral lines of importance in that range. In these cases the end-to-end throughput model becomes a highly useful tool for decision making. The final design optimizes both blaze angle and contrast modulation at 450nm for the blue VPHG and 705nm for the red VPHG, which are approximately the central wavelengths of each channel. The technical specifications of the VPHGs can be found in Section 4.2.4. Figure 4.15 shows the efficiencies of the blue and red channel VPHGs. Manufacturer predicted and measured performance are shown on the graph, as is data measured in-house (details of the in-house VPHG efficiency measurements can be found in Section 4.2.5).

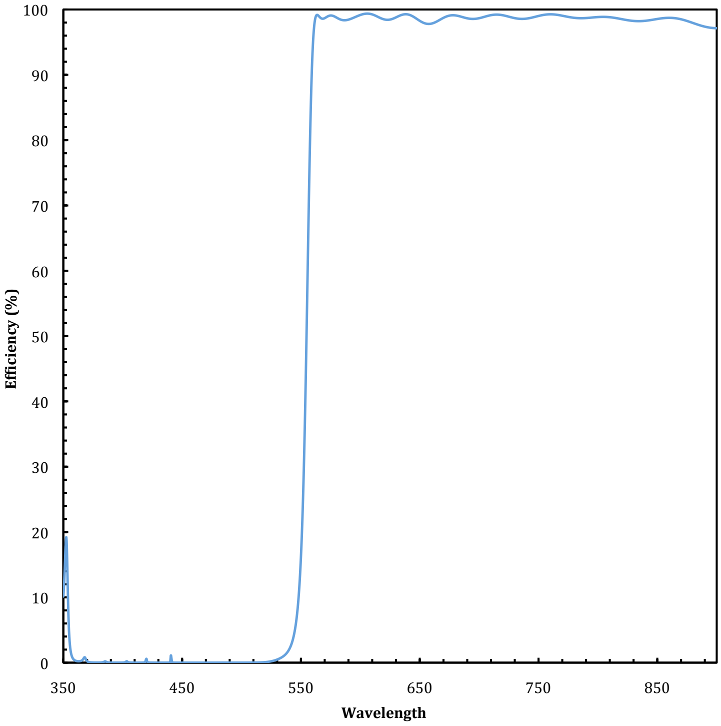

Figure 2.14: Theoretical transmittance as a function of wavelength for the Asahi Spectra manufactured dichroic illuminated at 5.5° angle of illumination, based on actual measured results using collimated light at normal illumination.

2.2.10 Cameras

The blue and red spectral channels have individual camera assemblies, optimized for their particular wavelengths of operation. Both cameras are fully dioptric and all lenses are coaxial with no tip or tilt. ZEMAX ray-trace diagrams of the red and blue cameras may be seen in Figures 2.16 and 2.17 respectively. The glass type for each lens is also shown in the diagrams. The blue camera consists of eight elements, two of which form a doublet made of glasses with similar coefficients of thermal expansion (since they are optically cemented together), with the final fused silica lens acting as both a field-flattener and as the entrance window to the detector cryostat. The red camera features six lenses, two of which form a doublet, again with similarly matched coefficients of thermal expansion and also using the final circular field-flattening lens as the detector cryostat window. The two largest lenses in the red camera (having diameters ∼220mm), L2 and L3, are made from S-BAH11 and BK7 respectively, due to their relatively low

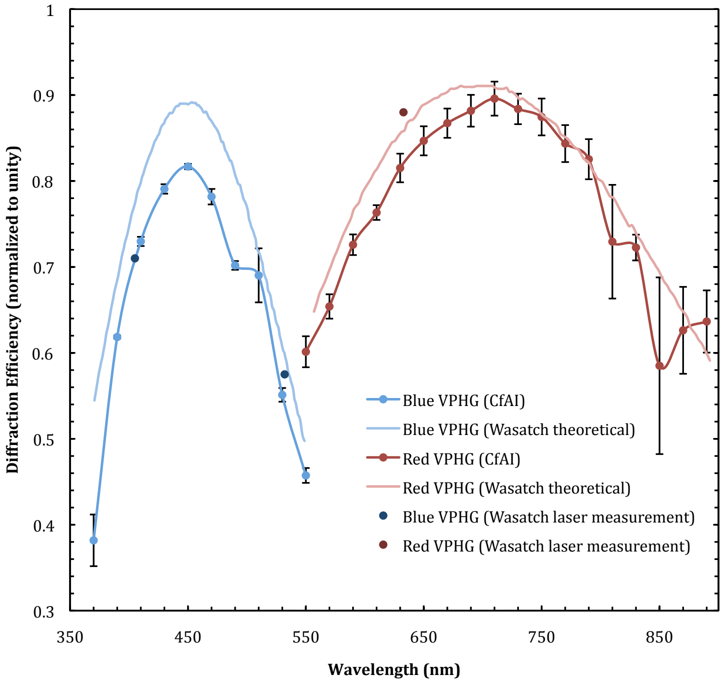

Figure 2.15: Diffraction efficiency as a function of wavelength for the Wasatch manufactured blue and red VPHGs, illuminated at 25.6° and 18.2° angles of illumination respectively. The red and blue lines represent the red and blue gratings respectively. Three data sets are shown: the lighter coloured lines represent the manufacturer's theoretical performance prediction based on rigorous-coupled-wave-analysis results; the individual data points represent manufacturer measured performance with laser light at given wavelengths. Finally, the darker coloured lines with error bars represent CfAI measured data, with the error bars representing the standard error from five repeated measurements.

costs as optical glasses.

The red camera has an effective focal length of 208.5mm and a path length of 470mm with a field of view of ±11.7°. With this focal length and a detector with 15µm pixel size, Nyquist sampling (two pixels per resolution element) up to a resolution of $R_{\text{max}} = 100000$ is achievable. The blue camera has an effective focal length of 166.8mm and a path length of 480mm with a similar field of view of ±11.6°. This enables $R_{\text{max}} = 80000$ with Nyquist sampling.

The focal length of the camera, f_{cam} is determined by noting that in order for maximum resolving power R_{max} to be achieved, the CCD must sample at least two resolution elements:

$$f_{\text{cam}} = \frac{n_{\text{samp}}}{2} R_{\text{max}} s_{\text{pix}} \cot\theta_B (1 + \tan\theta_B \tan\theta) \tag{2.2}$$

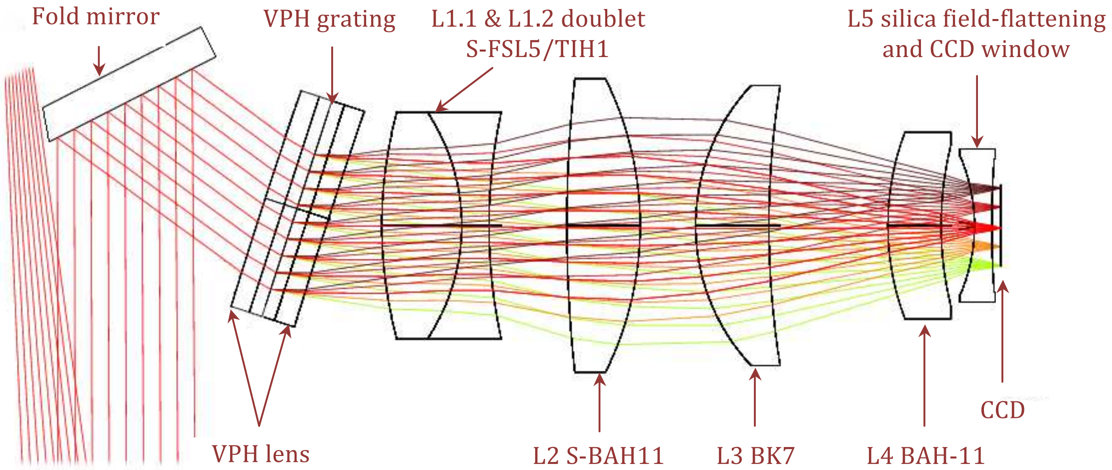

Figure 2.16: A ray-trace diagram of the red camera, also showing the fold mirror and VPH grating. The final lens acts as a field-flattening lens but also as a CCD window. Glass types for each lens are also labelled (Barnes, 2008).

where n_{samp} is the number of CCD pixels per resolution element and s_{pix} is the pixel size. Typically $n_{\text{samp}} = 2$ for the critical Nyquist sampling condition giving

$$f_{\text{cam}} \simeq R_{\text{max}} s_{\text{pix}} \cot \theta_B \tag{2.3}$$

for small θ. Therefore for an instrument such as SALT HRS, with a high blaze angle (R4 = $\tan 76°$), the cameras must have short focal lengths. Since R = Constant $\times$ $B \tan \theta_B$, (from equation 3.11 in Section 3.2), where B is the collimated beam diameter incident on the échelle, the focal ratio of the camera (with monochromatic conditions) is

$$\frac{f_{\text{cam}}}{B} = \text{Constant} \times s_{\text{pix}} \tag{2.4}$$

Therefore for a given R_{max}, f_{cam} depends only on pixel size.

As with the throughput modelling of the input optics, the glass transmissions and anti-reflection coating transmissions for the camera elements are combined, taking into account the number of optical surfaces and the central thickness of each lens in the cameras (see Tables 2.4 and 2.5). The calculated throughput of both cameras is shown in Figure 2.18.

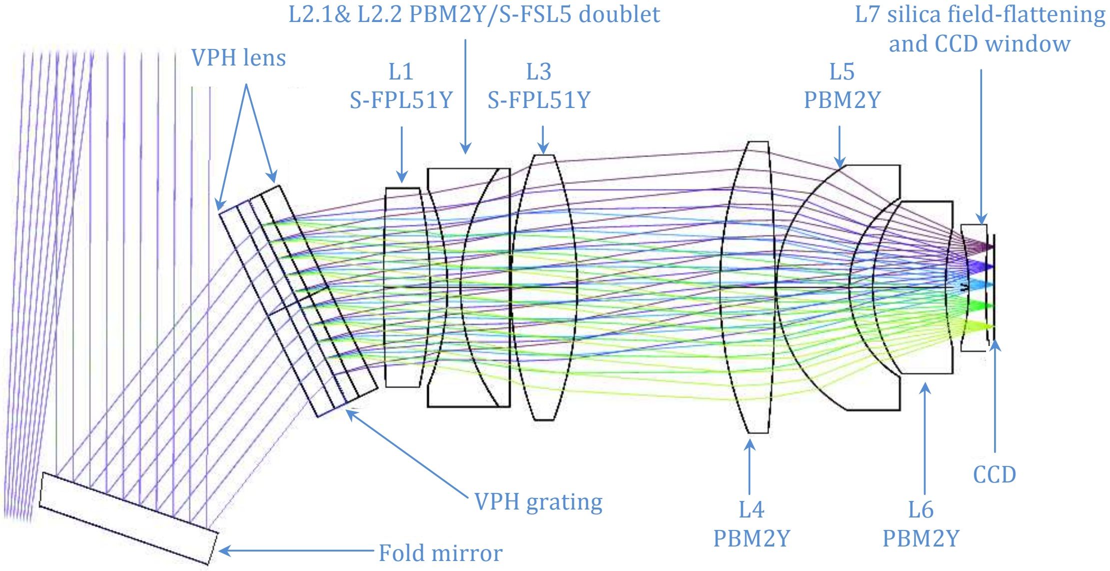

Figure 2.17: A ray-trace diagram of the blue camera, also showing the fold mirror and VPH grating. The final lens acts as a field-flattening lens but also as a CCD window. Glass types for each lens are also labelled (Barnes, 2008).

2.2.11 Charge Coupled Device Detectors

The charge coupled device (CCD) detectors for SALT HRS are highly efficient, featuring many technological advantages over typical consumer CCDs, such as backside illumination, deep depletion silicon, fringe suppression technologies and cryo-cooling. Their technical specifications and details of performance can be found in Chapter 5. Figure 2.19 shows the quantum efficiency curves for both the red and blue detectors as provided by the manufacturer, e2V[8]. Also

Table 2.4: Blue Camera Properties

Component	Glass Type	Central Thickness (mm)	Optical Surfaces
Lens 1.1	S-FPL51Y	35	2
Lens 2.1	PBM2Y	10	1
Lens 2.2	S-FSL5	37	1
Lens 3.1	S-FPL51Y	50	2
Lens 4.1	PBM2Y	41.5	2
Lens 5.1	PBM2Y	55	2
Lens 6.1	PBM2Y	55	2
Lens 7.1	Fused Silica	3	2

---

[8] http://www.e2v.com/

Table 2.5: Red Camera Properties

Component	Glass Type	Central Thickness (mm)	Optical Surfaces
Lens 1.1	S-FSL5	60	1
Lens 1.2	S-TIH1	20.5	1
Lens 2.1	S-BAH11	55	2
Lens 3.1	BK7	55	2
Lens 4.1	S-BAH11	40	2
Lens 5.1	Fused Silica	13	2

over-plotted is the in-house measured efficiency of the blue detector (see Chapter 5 for further details of the experimental procedure and results). At the time of writing, the red detector has not yet shipped to CfAI for testing.

2.3 Total Efficiency Curves

2.3.1 Total Instrument Throughput

The total instrument throughput is an amalgamation of all the optical components mentioned above. Although individual losses can appear negligible on first inspection, the sheer number of components and optical surfaces can quickly deteriorate the overall throughput, particularly in the extreme blue wavelengths. Figure 2.20 shows the throughput of each set of components discussed in the previous sections, as a function of wavelength. Also included is a plot of the overall instrument efficiency for the high resolution mode. It is clear that the response curve shape in each arm is dominated by the échelle grating, the VPHGs and the detectors (although almost all components show significant tail-off in performance below 400nm).

Figure 2.21 shows the expected throughput of SALT HRS at all four modes of operation (with the iodine cell losses included with the high-stability mode).

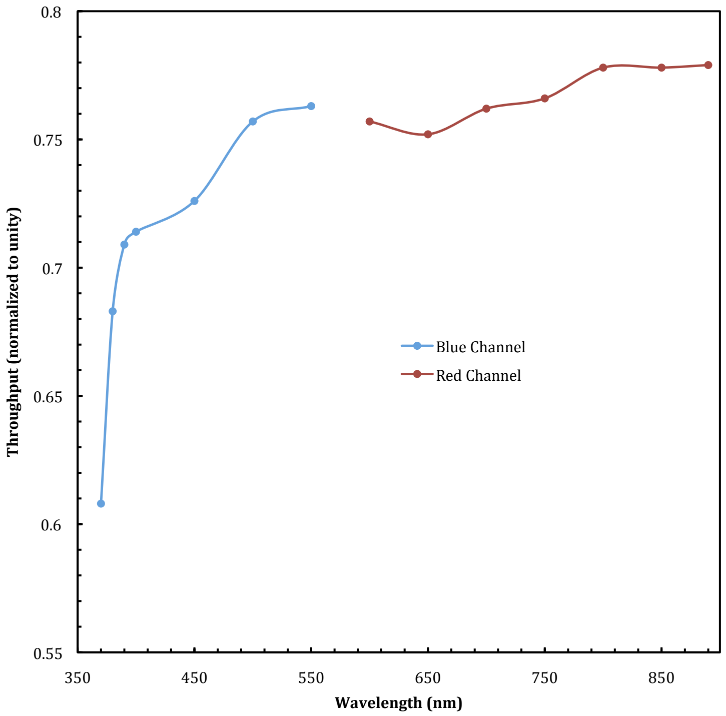

Figure 2.18: The calculated throughput of both blue and red cameras, plotted in corresponding colours. Glass transmission losses contribute to more significant losses shortward of 400nm where performance decreases rapidly with wavelength.

2.3.2 Signal-to-Noise Estimates & Sensitivity Limits

The throughput measurements may be used to estimate the global performance of SALT HRS as a function of wavelength and magnitude. Figures 2.22*a* and 2.22*b* show predicted signal-to-noise (SNR) ratios for the blue and red channels respectively. The predictions use an F-type star as the science object, assuming: no reddening; 1 arcsec atmospheric seeing conditions; 1 airmass atmospheric depth; a 1000s exposure time; new moon observation conditions and the low resolution instrument mode. The S/N is expressed for a 13, 15, 17 and 19 magnitude object in the V-band. Additional S/N calculations are presented in the sections below, which describe the exposure time calculator (used to calculate the performances shown).

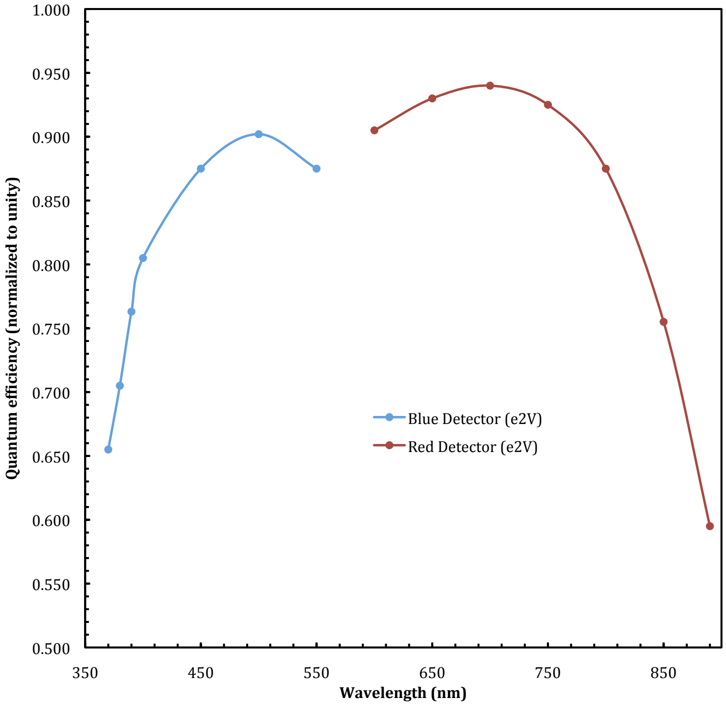

Figure 2.19: The quantum efficiency of the SALT HRS detectors as measured by the manufacturer, e2V. Also over-plotted is the in-house measurement of the blue detector efficiency.

2.4 Exposure Time Calculator

The SALT HRS exposure time calculator (ETC) takes user-defined observation parameters, and uses them to calculate the dark count, read noise, sky background flux and object flux in units of electrons per wavelength bin. The flux is then used to calculate the SNR per pixel at a variety of wavelengths. Using this information, astronomers may use the software to properly plan their observations. Additional user notices are also created. For example, atmospheric refraction in the absence of atmospheric dispersion correction is calculated. Should the exposure time lead to a large number of cosmic ray hits or detector blooming (due to saturation of some pixels), this is also mentioned as a warning. The software was written in the Java programming language at the request of SALT, and features a graphical user interface to simplify use by astronomers. Since it is compiled as a .jar file, the software is platform independent and can run on any computer with any operating system, using Java plugins (which are packaged with all modern web browsers) directly from a webpage. No developer tools or Java compilers are required, making the software

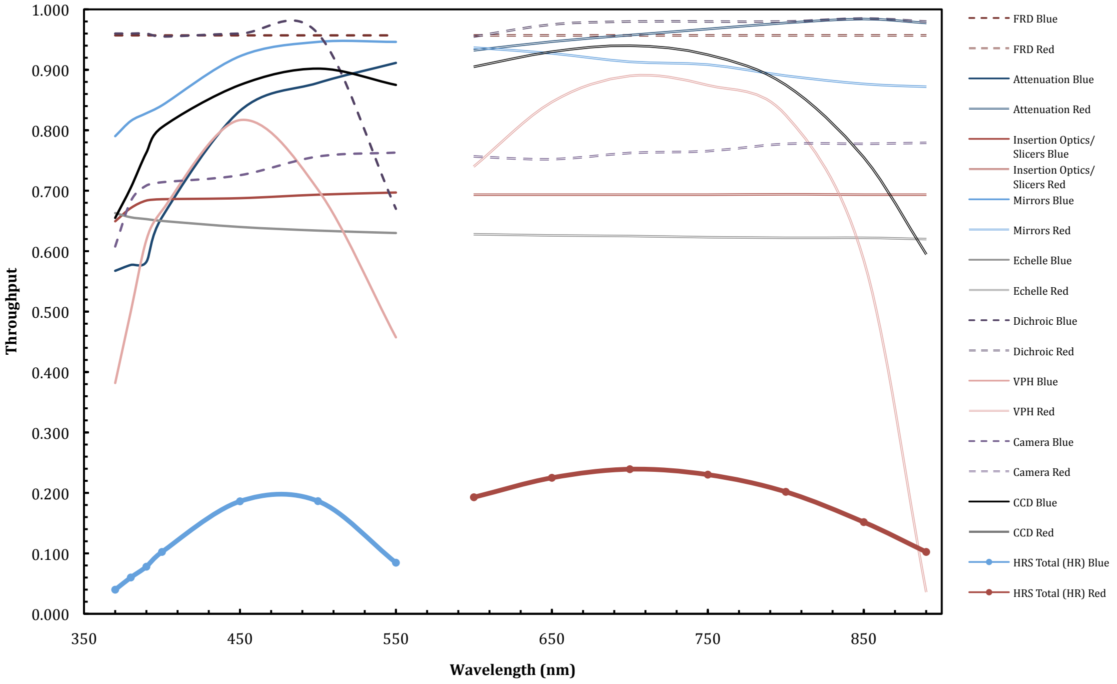

Figure 2.20: A breakdown of throughput as a function of wavelength for each major component group, also showing the combined result of performance for the high resolution mode.

readily accessible for potential instrument users. The GUI is shown running in Mac OS X 10.6 in Figure 2.23.

2.4.1 ETC User Parameters

The following parameters are required by the ETC software and must be defined by the user (default values which the software launches with are shown in parentheses):

- Science object magnitude in the V-band (12)
- Object spectral type-OBAFGKM (F)
- Reddening E(B-V) (0 mag)
- Exposure time (1000s)
- Spectrograph mode (Medium-resolution)

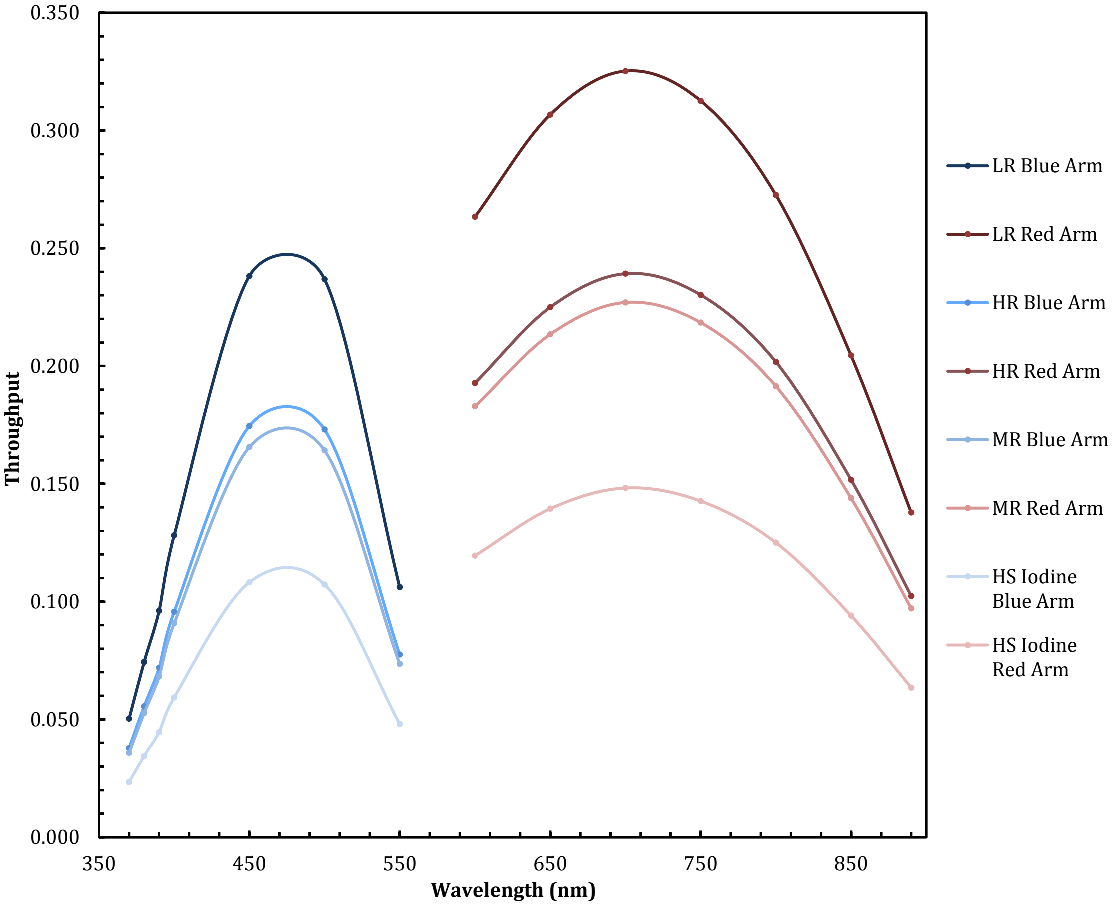

Figure 2.21: Total throughput of SALT HRS in both spectral channels and all four major operational modes (the iodine cell is also used in the high-stability mode).

- Effective mirror diameter (Average-46m^2)
- Presence of iodine cell (No)
- Moon brightness (Dark)
- Seeing (1arcsec)
- Airmass (1.5)
- Temperature (25°C)
- Atmospheric pressure (773hPa)
- Relative humidity (25%)

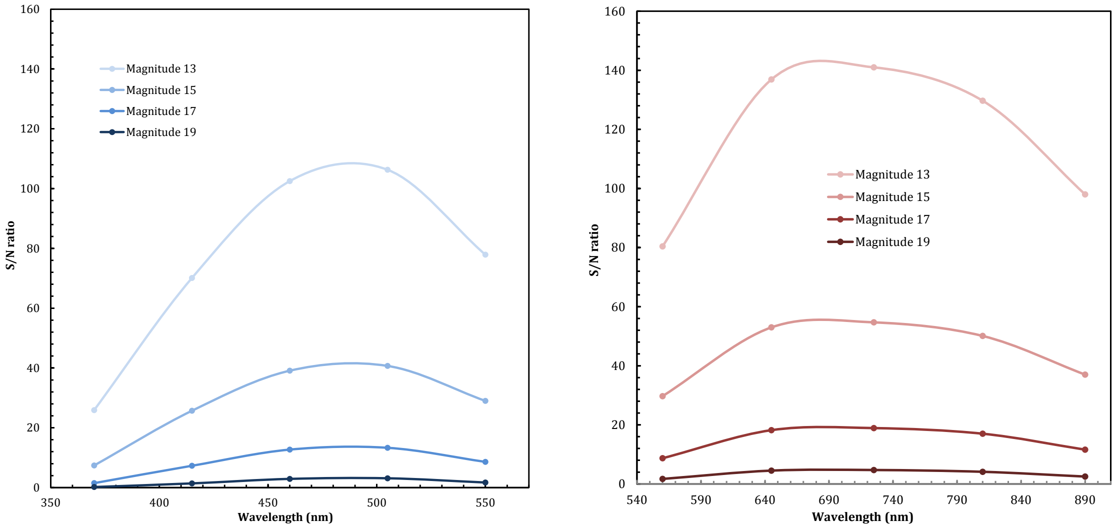

(a) Blue spectral channel

(b) Red spectral channel

Figure 2.22: SNR as a function of wavelength in a 1000s exposure for four different object magnitudes using the low resolution mode.

These parameters can be seen as a column of inputs on the left side of the GUI screen in Figure 2.23. Values can be entered by the user as text or using the drop-down menus provided which feature lists of pre-defined entries. Each parameter is set with a realistic range or pre-defined list of options to prevent erroneous data entry. The temperature, atmospheric pressure and humidity fields may be left unchanged in most cases, as these values are used only to calculate the atmospheric refraction of the object light, which is efficiently corrected for by the atmospheric dispersion corrector in the SAC optics. This information is therefore included more as a ‘for your interest’ footnote.

2.4.2 Pre-defined Input Parameters

The ‘hard-wired’ input parameters defined within the code are usually stored as an array of five elements. Two arrays generally exist for each parameter being described. For example, the instrument efficiency for a particular mode is stored in two arrays, with one array for each spectral arm. Each array contains five elements corresponding to the five wavelengths in each spectral channel for which calculations are performed (370, 415, 460, 505 and 550nm for the blue

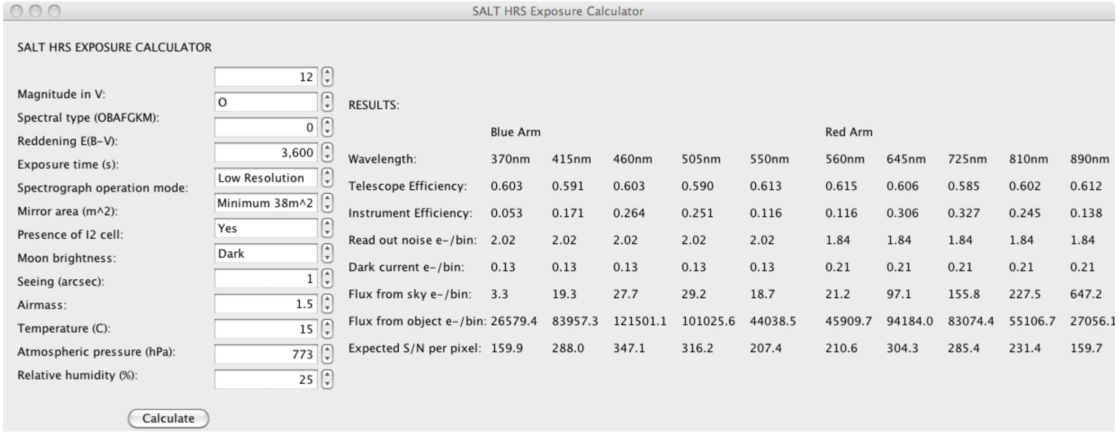

SALT HRS EXPOSURE CALCULATOR

Magnitude in V: 12
Spectral type (OBAFGKM): O
Reddening E(B-V): 0
Exposure time (s): 3,600
Spectrograph operation mode: Low Resolution
Mirror area (m^2): Minimum 38m^2
Presence of I2 cell: Yes
Moon brightness: Dark
Seeing (arcsec): 1
Airmass: 1.5
Temperature (C): 15
Atmospheric pressure (hPa): 773
Relative humidity (%): 25

RESULTS:

	Blue Arm					Red Arm				
Wavelength:	370nm	415nm	460nm	505nm	550nm	560nm	645nm	725nm	810nm	890nm
Telescope Efficiency:	0.603	0.591	0.603	0.590	0.613	0.615	0.606	0.585	0.602	0.612
Instrument Efficiency:	0.053	0.171	0.264	0.251	0.116	0.116	0.306	0.327	0.245	0.138
Read out noise e-/bin:	2.02	2.02	2.02	2.02	2.02	1.84	1.84	1.84	1.84	1.84
Dark current e-/bin:	0.13	0.13	0.13	0.13	0.13	0.21	0.21	0.21	0.21	0.21
Flux from sky e-/bin:	3.3	19.3	27.7	29.2	18.7	21.2	97.1	155.8	227.5	647.2
Flux from object e-/bin:	26579.4	83957.3	121501.1	101025.6	44038.5	45909.7	94184.0	83074.4	55106.7	27056.1
Expected S/N per pixel:	159.9	288.0	347.1	316.2	207.4	210.6	304.3	285.4	231.4	159.7

Fig[illegible].23: A screenshot of the SALT HRS exposure time calculator graphical user interface.

channel; and 560, 645, 725, 810 and 890nm for the red channel). The pre-defined parameters may be adjusted in the code (but not by a general GUI-level user), and are as follows:

- The CCD operating temperatures and readout noise. Detector dark current is calculated as a function of the operational temperature of the detectors, based on a temperature relationship described by the manufacturers[9]. This temperature had not been finalized at the time of software development and may be subject to change.
- The SALT telescope efficiency at the ten wavelength bands used (see Figure 2.3). This is used for converting the object photon flux incident on the telescope into the flux entering the instrument via the optical fibres.
- The optical fibre attenuation at each wavelength, and the length of the fibres in metres (from which the fibre losses may be calculated, since the actual length was unknown at the time of software writing). This is used to calculate (using equation 2.1 from this chapter) the photon flux delivered to the spectrograph input optics by the fibre connection from the telescope.
- The SALT HRS instrument efficiencies at the ten wavelength bands used, taking into account the different modes and the optional iodine cell (see Figure 2.21). This is used to convert the photon flux delivered from the telescope to the photon flux recorded by the

---

[9] www.e2v.com

instrument CCDs.

- Tables describing the sky background light caused by moonlight, at all ten wavelengths and for the conditions of: new moon; half moon; and full moon. This is used to estimate the sky background flux which will [illegible] the object SNR values.
- Colour correction tables for each spectral type. These are used to calculate the magnitude of an object at the ten calculation wavelengths based on its magnitude in the V-band. Since different stellar type objects have different spectra, their magnitude is not constant at all wavelengths. The object magnitude at each wavelength is used to calculate the photon flux from the object.
- A flux calibration table consisting of the number of photons/cm^2/s/Å at each wavelength produced by an A0V star at 0.0 magnitude. This allows for conversion of the magnitude of the science object at each wavelength to the photon arrival rate per unit area on the telescope. These values are taken from the Hubble Space Telescope absolute spectrophotometry of Vega (Bohlin & Gilliland, 2004), and corroborated with the Pickles stellar flux libraries (Pickles, 1998).
- Fibre radii of the 350µm and 500µm fibres are used for calculation of the fibre angular radii (used for calculating the seeing efficiency) and the fibre spatial angle (used for sky and object flux calculations).
- The bin size at each wavelength in Å/pixel, calculated for the order in which said wavelength is closest to the blaze wavelength. The bin size values were calculated using the linear dispersion of SALT HRS (see equation 1.10 in Chapter 1) and the 15µm pixel size of pixels in both detectors.

2.4.3 Signal-to-Noise Ratio (SNR)

The signal-to-noise ratio is calculated as follows:

$$S/N = \frac{\text{Signal}}{\text{Noise}} = \frac{Nt}{\sqrt{(N + N_{\mathrm{D}} + \text{[illegible]}} + \sqrt{n_{\mathrm{slt}}} R_{\mathrm{e}}} \tag{2.5}$$

Table 2.6: Atmospheric Extinction Co-efficients used in the ETC

$k(\lambda)$		$k(\lambda)$	
k(V=550nm)	0.125	k(560)	0.118
k(365nm)	0.443	k(645nm)	0.085
k(415nm)	0.318	k(725nm)	0.065
k(460nm)	0.227	k(810nm)	0.016
k(505nm)	0.165	k(890nm)	0.001

where N is the number of photons per second per resolution element from the object; N_{D} is the number of photons per second per resolution element from dark current in the CCD; N_{B} is the number of photons per second per resolution element from the sky; R_{e} is the RMS readout noise per pixel; n_{slt} is the number of pixels per resolution element and t is exposure time. The photon flux from the object (N) is dependent not only on the efficiency of the telescope, fibre feed and instrument, but also on the interstellar and atmospheric extinction.

2.4.4 Atmospheric Extinction

The transmission losses associated with the atmosphere depend on the airmass, X, of the object being observed, and also on the extinction coefficient k for a given wavelength. Atmospheric transmission, T_{atm} can be calculated using

$$T_{\mathrm{atm}} = e^{\frac{-k}{1.086}X} \tag{2.6}$$

The values of k are interpolated using a high order polynomial from those defined at the standard U, B, V, R and I bands (365, 445, 551, 658 and 806nm respectively). Atmospheric extinction is as follows in Table 2.6:

2.4.5 ETC Calculation Processes

Figure 2.24 shows a flow-chart describing the processes and calculations performed by the exposure time calculator code, in order of their execution. For further details the reader may refer to the actual ETC code, complete with comprehensive comments, which can be seen in the

Table 2.7: Parameters held constant (unless individually altered under test) in ETC results testing

Parameter	Value	Parameter	Value
Object Magnitude	13	Mirror Area	Average
Object Spectral Type	O	Moon Brightness	Dark
Reddening E(B-V)	0.0	Seeing	1.0arcsec
Exposure Time	3600s	Airmass	1.5
Instrument Mode	Low-resolution		

Appendix of this thesis. The example of the code presented there has been 'stripped down' and is specific only to the calculations performed by the ETC, since including the code to control and format the GUI would increase the volume of code by an order of magnitude. This version nonetheless could be opened using a command line and Java compiler and will print the results to screen if the user parameters are written to the correct fields.

2.4.6 Example ETC Results

The output of the ETC obviously varies according to the input parameters. In order to check for proper functioning of the ETC, comparisons are made between the results when all parameters are held constant but one. The parameters to be individually varied are shown in Table 2.7, along with the default value used when testing other parameters. These parameters are in most cases tuned to give a high flux or SNR value, so that when compared to less favourable parameters the dynamic range of the effect is large enough to clearly show the effect. Each parameter is varied in turn and the results are as follows:

- Figure 2.25 shows the change in SNR when the object magnitude is reduced. As expected, the SNR decreases with increasing magnitude, with the difference between magnitudes becoming increasingly smaller, since magnitude is a logarithmic scale of brightness. The sky flux remains constant but magnitude change alters the object flux.
- Varying spectral type shows a relative change in the SNR as a function of wavelength (see Figure 2.26). The O type object is hot and blue and therefore has higher SNR at 370nm than at 890nm due to increased object flux in the blue. The much cooler M type object has a different spectrum, and so has poor blue SNR and very high red SNR. This result meets

expectations. Again the sky flux is unchanged by this change whilst object flux varies at different wavelengths.

- As with changes in spectral type, changes in the interstellar reddening of an object causes SNR to vary with wavelength. Figure 2.27 shows how the highly reddened objects (higher E(B-V) value) show relatively lower SNR in the blue and higher SNR in the red. Sky flux values are unaffected by a reddening value change.
- Figure 2.28 shows the effect of reducing exposure time. With shorter exposures the SNR decreases proportionally across all wavelengths, as is to be expected. The result is not linear since SNR has a t component on both the numerator and denominator. Although changing exposure time alters both sky and object flux by the same amount, in the denominator the t component is square rooted (see equation 2.5).
- The expected trend is seen in Figure 2.29, in that the higher resolving power modes have a lower SNR across the wavelength range. The addition of the iodine cell also reduces performance as predicted. There is a slight wavelength dependence introduced by the use of different optical components with different performances between the modes. The curves roughly follow the shapes and relative differences of the instrument throughput curves shown in Figure 2.21 (although here the blue performance is relatively higher than the efficiency alone would suggest, since the object is a hot blue O type star). Changing instrument efficiency varies both sky and object fluxes.
- The effect of reducing/increasing the mirror area shows that SNR is reduced/increased across the wavelength range. This can be seen in Figure 2.30. The change is smooth because changes in mirror area alter the object and sky flux collected at all wavelengths by the same factor.
- Figure 2.31 shows the effect on sky background flux when the moon brightness is changed. Significant increases can be seen in some wavelengths when the moon is brighter. The object flux remains unchanged by the moon conditions, but the SNR is decreased as sky background increases.

- An increase in seeing (which is a decrease in atmospheric quality for observing as the object PSF is increasingly smeared) produces a lower SNR as would be expected. This effect can be seen in Figure 2.32. A change in seeing affects only the object flux.
- Finally, the effect of increasing airmass is seen in Figure 2.33. As more atmosphere stands between telescope and object, the SNR decreases as a result of increased extinction.

To summarize, the effects outlined above seem to confirm the correct functioning of the SALT HRS exposure time calculator. The results of parameter changes can all be explained properly when considering the change in light the instrument receives at a given wavelength, the way the instrument responds to this light, and the way the SNR equation (2.5) dictates the response according to different changes in sky [illegible] object fluxes.

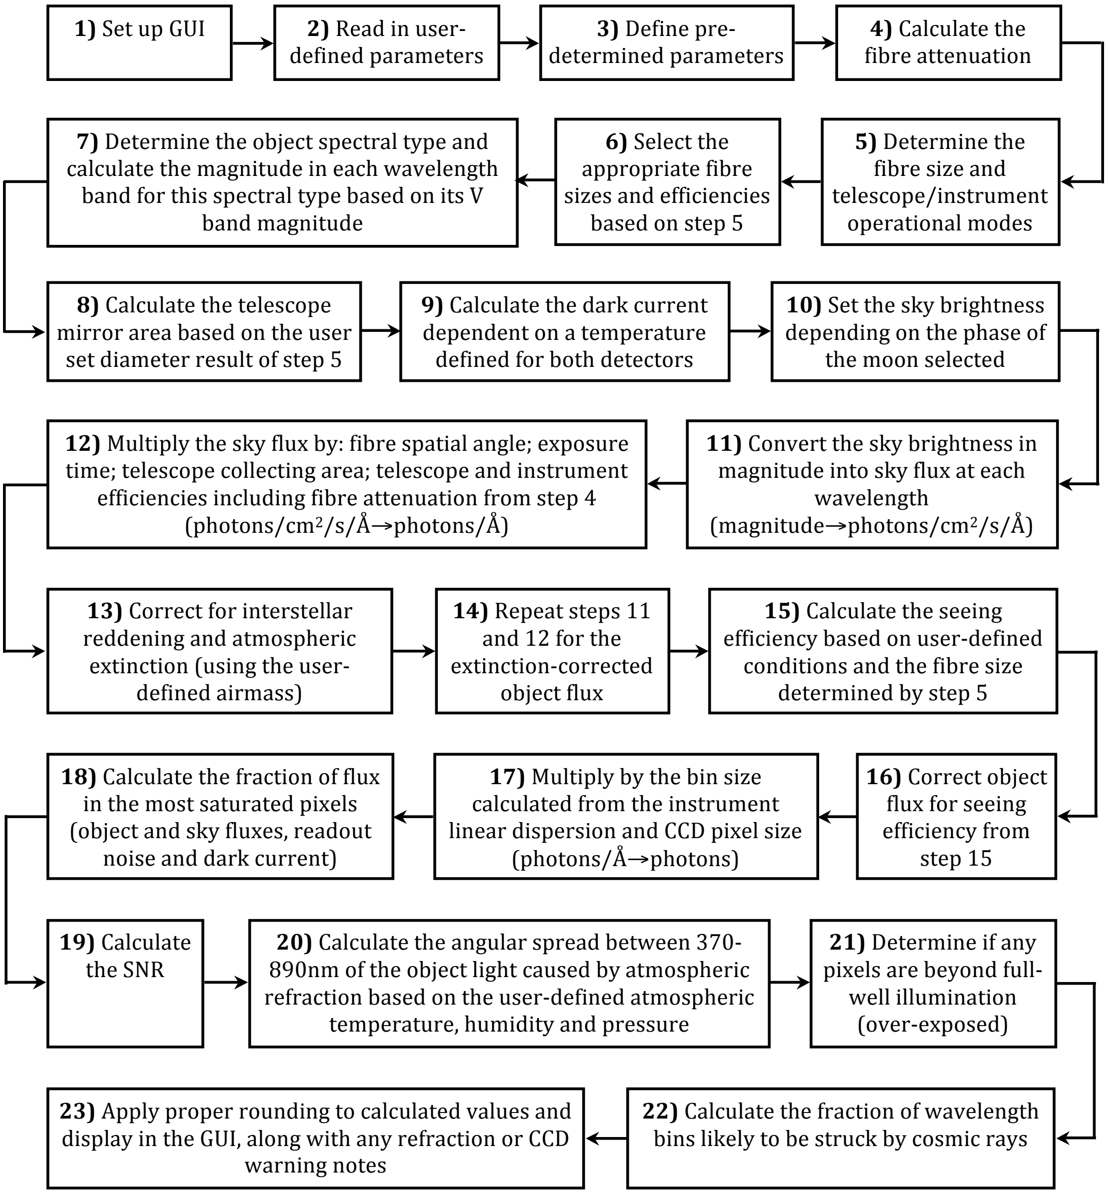

Figure 2.24: A flow-chart of the processes and calculations performed by the exposure time calculator.

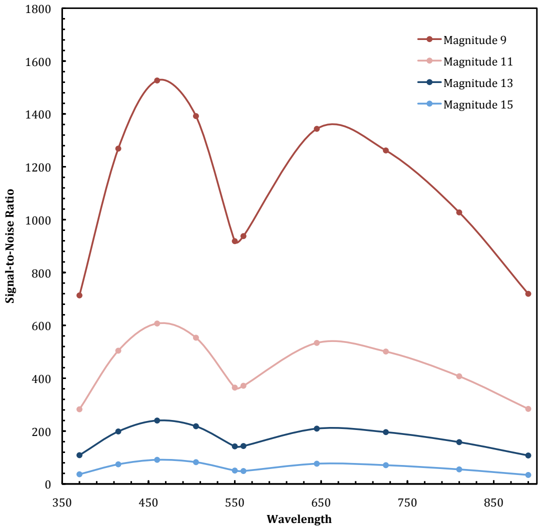

Figure 2.25: Change in object magnitude.

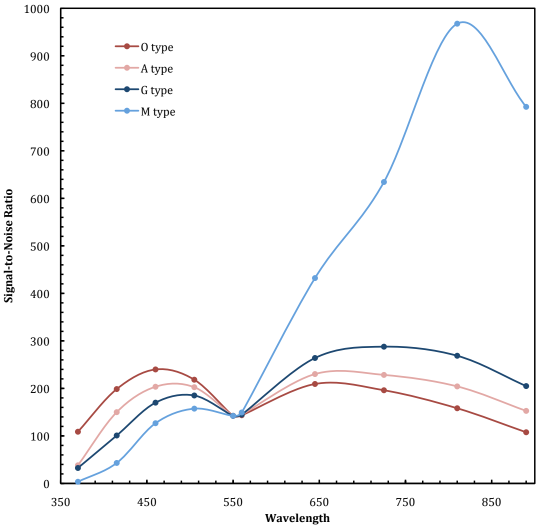

Figure 2.26: Change in object spectral type.

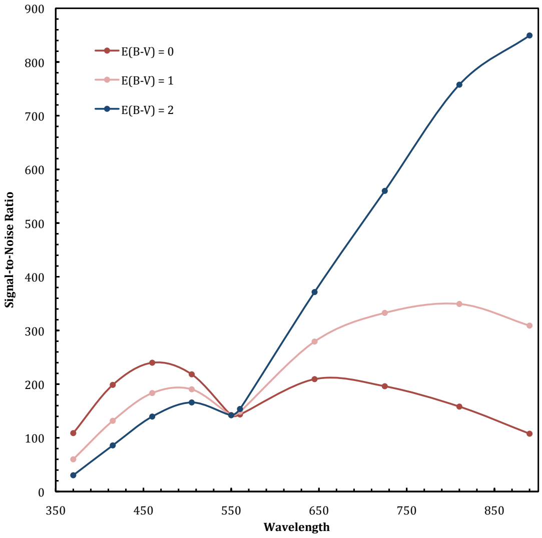

Figure 2.27: Change in object reddening.

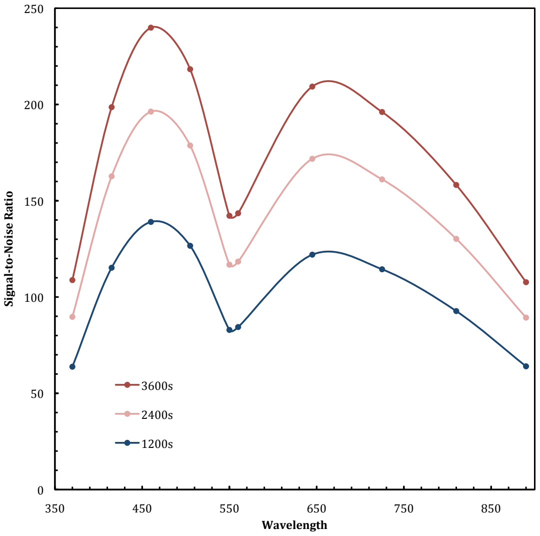

Figure 2.28: Change in observation time.

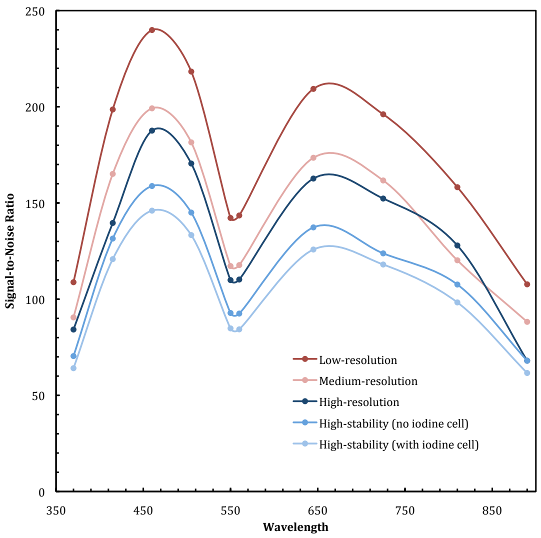

Figure 2.29: Change in instrument mode.

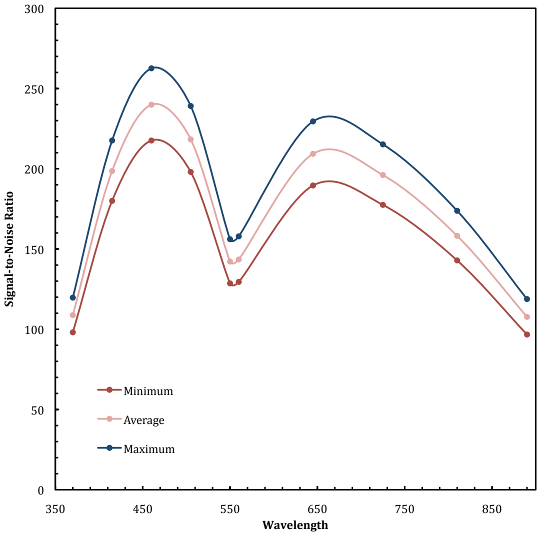

Figure 2.30: Change in mirror diameter used.

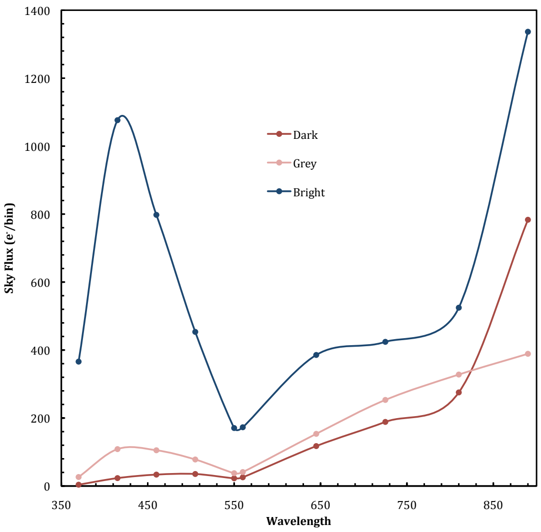

Figure 2.31: Change in moon brightness.

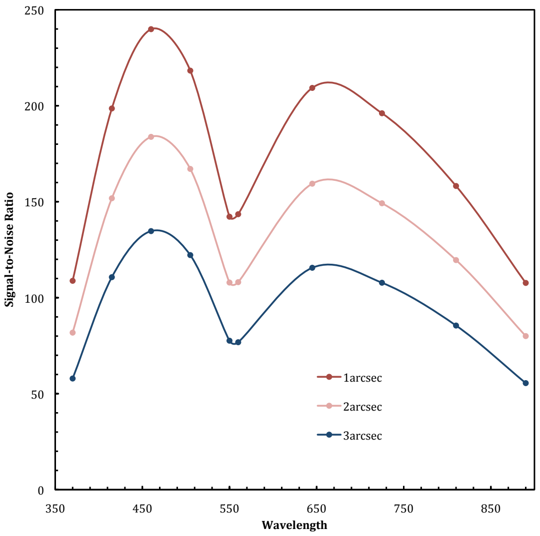

Figure 2.32: Change in atmospheric seeing.

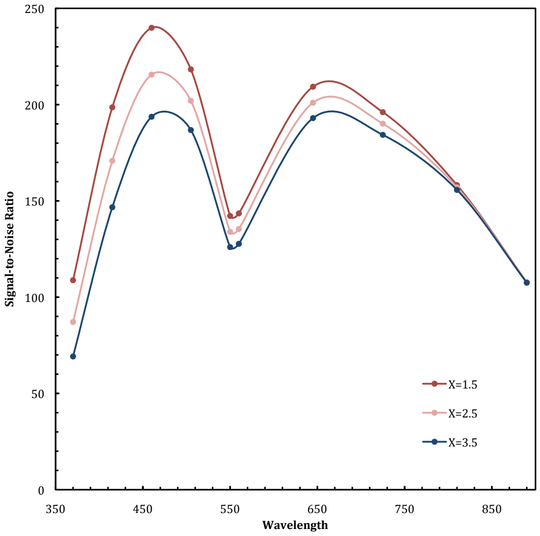

Figure 2.33: Chang irmass.

Chapter 3

Fibre System

3.1 Background to Optical Fibres

Optical fibres can be thought of as a 'light pipe' and saw development in the 1960's, predominantly driven by the telecommunications industry. A core (usually of fused silica) operates as a light guide, propagating light along its length by total internal reflection (TIR). This core is then wrapped in a cladding of lower refractive index than the core material, which helps to prevent loss at the core surface. Ideally the cladding is at least ∼10 times the thickness of the longest wavelength of light to be transmitted, to minimize loss (Lu et al., 1998). Buffer and jacket layers (often made of polyamide) then cover the cladding for protection.

Each mode can be considered as a plane wave, with different modes incident at different angles to the fibre axis. Low order modes are close to the axis and higher order modes are closer to the critical angle for TIR. A mode is 'lossy' if it exceeds the critical angle and is therefore not propagated along the fibre. Mode bumping occurs when a mode is bumped to a different order by material defects or imperfections at the core/cladding boundary. An alternative (and more accurate) model for an optical fibre is to consider it as an optical cavity within which guided waves (modes) propagate.

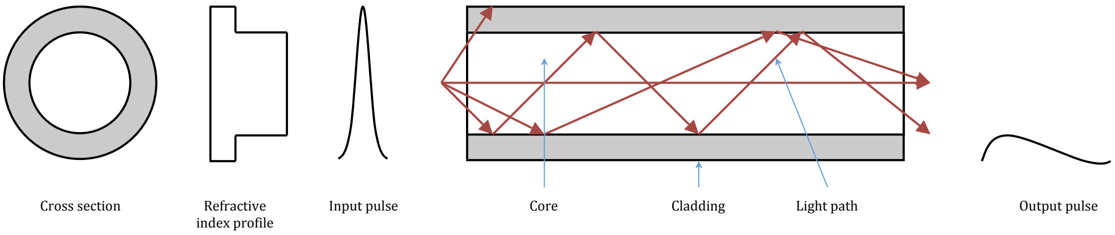

Figure 3.1: Fibre modes in a multimode stepped index optical fibre (Nagel, 1993).

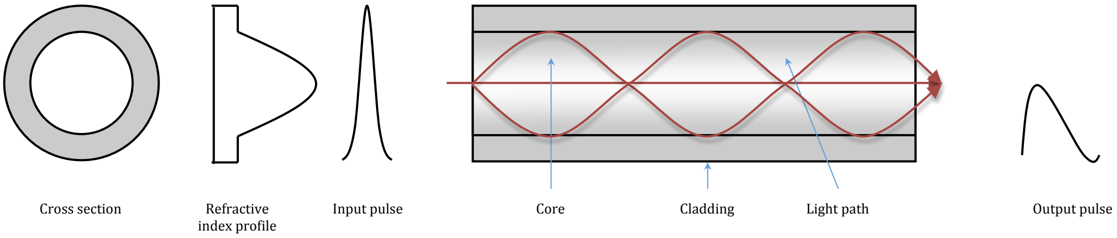

Figure 3.2: Fibre modes in a multimode graded index optical fibre (Nagel, 1993).

Aside from specialist photonic fibres or those with unusual core shapes, most optical fibres fall into one of three categories: multi-mode stepped-index (see Figure 3.1), multi-mode graded-index (see Figure 3.2) and single-mode stepped-index (see Figure 3.3). With self-explanatory nomenclature, multi-mode fibres carry multiple modes, as opposed to single-mode fibres which carry only one (fundamental) mode. Single-mode fibres are useful in interferometry - since the path length is controlled by the fibre length (Parry, 1998) - but are of little use for a spectrograph feed if the incident beam has many modes. A stepped-index fibre has a uniform refractive index across the core, whereas a graded-index fibre uses many layers of gradually decreasing refractive index. This forms a parabolic function in a step-wise manner. In a graded-index fibre, as the different order modes propagate, higher order modes travel a greater distance than the low

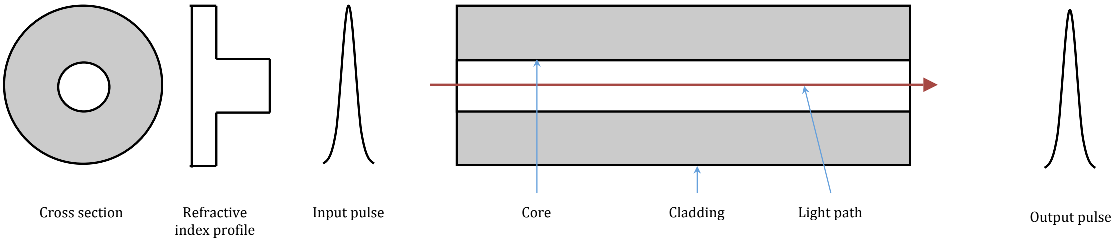

Figure 3.3: Fibre modes in a single mode stepped index optical fibre (Nagel, 1993).

order modes. The graded-index fibre therefore removes some of the temporal smearing, as the additional path length of the high order modes is somewhat mitigated by their increased speed through the lower index materials. This is useful when pulses are transmitted, such as in the telecommunications sector. However, increased spectral coverage, lower price and greater range of core sizes has made the multi-mode step-index variety, the fibre of choice for astronomical purposes (Nelson, 1988).

3.2 Considerations for Optical Fibre use in Astronomy

An obvious consideration when designing a fibre-fed instrument, is the spectral attenuation of the fibre, which will vary as a function of wavelength and cause losses to increase with fibre length. Hydroxyl (OH-) ions in the silica cause absorption in the red whilst metal impurities cause absorption in the blue. Further loss is attributed to Rayleigh scattering due to density fluctuations, which dominates losses in silica cores. Fibres with a high OH- content are considered 'wet' and feature absorption bands in the red. 'Dry' fibres with low OH- content are available but the process of manufacture breaks the silicon oxide (SiO) bonds needed for good blue transmission (Barden, 1995). See Figure 3.4 for a comparison of the transmission of 'wet' and 'dry' fibres. The attenuation of SALT HRS fibres is discussed in greater detail in Section 3.4 and also in Section 2.2.2 of the previous chapter.

In addition to attenuation, there are three other noteworthy issues for fibre use in astronomy, which are mode dependent rather than material or instrument dependent: azimuthal scrambling; radial scrambling; and focal ratio degradation (FRD). All three effects occur as a result of mode bumping, which has both beneficial and detrimental effects.

Azimuthal scrambling causes an input laser beam injected into the fibre at angle θ to emerge from the output with an annulus shaped distribution with radius θ, demonstrating symmetry in intensity about the fibre axis. The rays entering the fibre off-axis are propagated in a helical fashion. As each skew ray travels along a sightly different path length (due to small scale imperfections in the fibre core), the result is a ring structure. The azimuthal scrambling effect

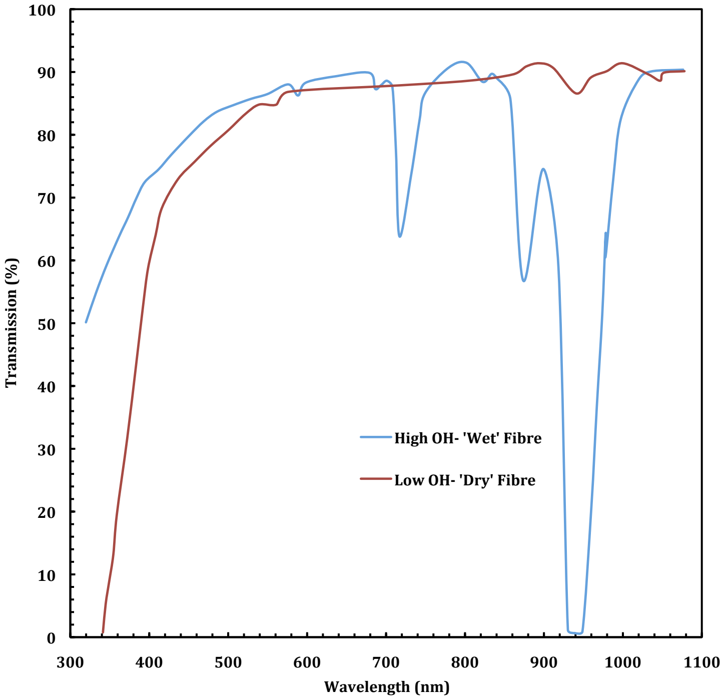

Figure 3.4: A comparison of high OH^- ‘wet’ and low OH^- ‘dry’ optical fibre transmissions from the ultra-violet to near-infra-red, showing the absorption bands in the red for ‘wet’ fibres and the poor blue performance of ‘dry’ fibres (Barden, 1995).

is complete (meaning that the input beam information is ‘smeared’ with complete symmetry), even in short lengths of fibre, according to both theoretical (Heacox, 1986) and experimental (Angel et al., 1977) results. This differs from the radial scrambling effect, which is incomplete, with longer length fibres providing more opportunities for modes to be ‘bumped’. The radial scrambling essentially determines the width $\Delta\theta$ of the output annulus. It can be thought of as a measure of the constancy of the output beam as the input is moved about the fibre input face. Figure 3.5 demonstrates the effect as a cartoon. Both forms of scrambling are of use when feeding a spectrograph, as they act to remove ‘memory’ of small scale fluctuations in the input illumination (for example due to tracking errors or telescope wind shake). In the case of SALT HRS, the scrambling effect is even more important as a means of reducing the effect of variable pupil illumination caused by the unique object tracking system of the SALT telescope. The ability to reduce the ‘memory’ of the input beam is of particular significance for the high stability mode which is optimized for high velocity precision (see also Section 3.3).

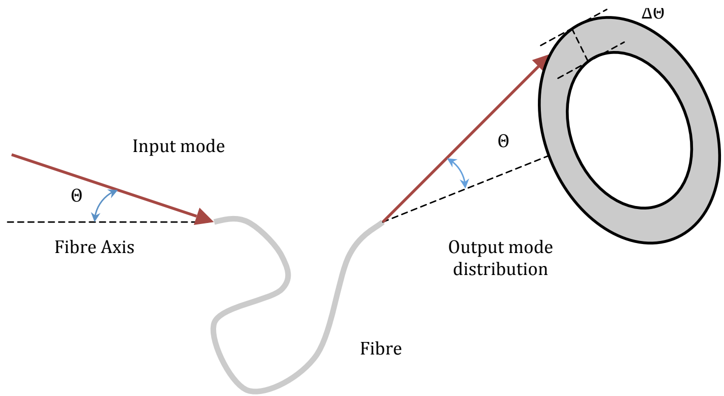

Figure 3.5: The radial and azimuthal scrambling when inserting a single mode beam of vanishingly small diameter and divergence-such as a laser (adapted from Ramsey, 1988).

Focal ratio degradation (FRD) is a phenomenon exhibited by all optical fibres in which the f-ratio of a beam of light is decreased when propagating down the fibre. The effect of FRD must be understood as a degradation in spectrograph performance since it leads to a reduction in resolving power. Figures 3.6 and 3.7 contrast the case of conventional slit and fibre-fed spectrographs. In Figure 3.6, the telescope with diameter D and focal length f_{tel} feeds a grating element through a slit with width w. Figure 3.7 shows a similarly designed spectrograph, this time illuminated via an optical fibre of diameter d rather than a slit. The angle subtended on

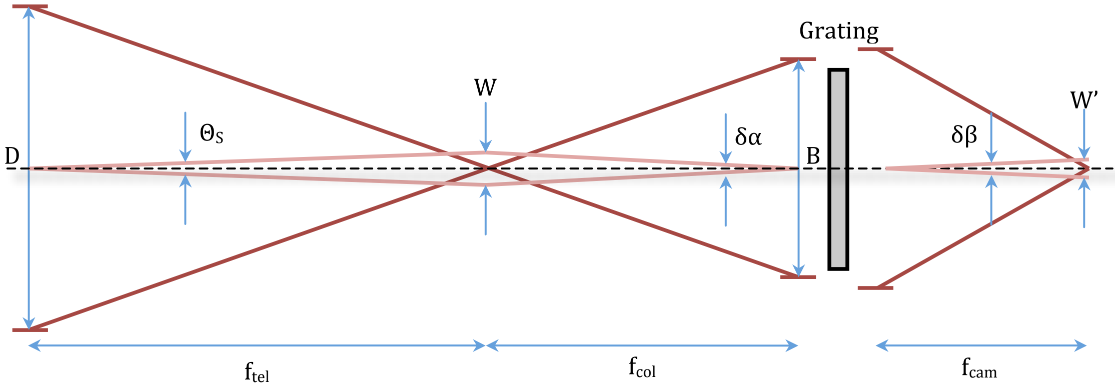

Figure 3.6: A schematic of a directly fed slit spectrograph (adapted from Schroeder, 1989).

the sky by the slit or fibre is given by

$$\theta_S = \frac{w}{f_{\mathrm{tel}}} \quad \text{or} \quad \theta\,\text{[illegible]}\,\frac{d}{f_{\mathrm{tel}}} \tag{3.1}$$

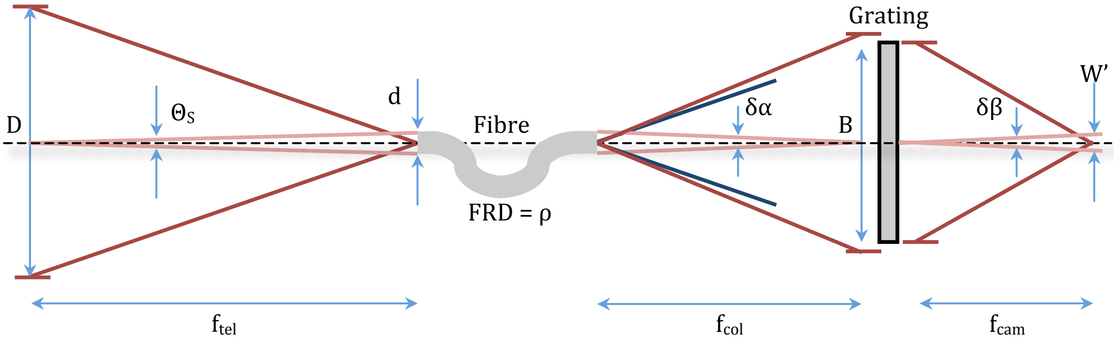

Figure 3.7: A schematic of a fibre-fed spectrograph, showing the faster beam emerging from the fibre (adapted from Schroeder, 1989).

For the directly coupled design the following situation applies:

$$\frac{f_{\rm tel}}{D} = \frac{f_{\rm col}}{B} \tag{3.2}$$

where $f_{\rm col}$ is the collimator focal length. However, since FRD causes the output beam to emerge faster than the input beam the equality above is broken and

$$F_{\rm in} = \frac{f_{\rm tel}}{D} \quad \text{and} \quad F_{\rm out} = \frac{f_{\rm col}}{B} \tag{3.3}$$

The FRD parameter, ρ describes this as

$$\rho = \frac{F_{\rm out}}{F_{\rm in}} \tag{3.4}$$

in which case a decrease in focal ratio of 10% would give an FRD parameter value of 1.1. The ρ term acts to modify the equality such that for fibre-fed instruments

$$\frac{f_{\rm tel}}{D} = \rho\frac{f_{\rm col}}{B} \tag{3.5}$$

Therefore, to maintain beam size equivalence between a direct- versus fibre-fed spectrograph, the collimator focal length must decrease, which would reduce resolution (if throughput is held constant by keeping beam size the same). Spectral resolution is defined as the ability to distinguish two wavelengths λ_1 and $\lambda_2 = \lambda_1 + \delta\lambda$ (see Figure 3.8):

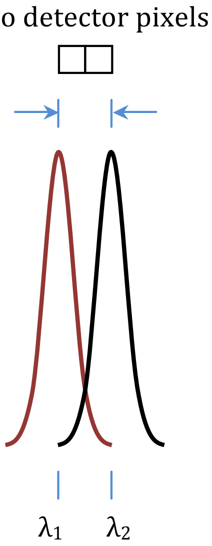

Figure 3.8: A diagram explaining the definition of spectral resolution for a spectrograph with a CCD detector.

$$R = \frac{\lambda}{\delta\lambda} \tag{3.6}$$

As the grating disperses λ_1 and λ_2 differently, the angular width between the two is $\delta\beta$. Resolution may then be written in terms of angular dispersion, $d\beta/d\lambda$, (see Section 1.2.2) as

$$R = \frac{\lambda}{\delta\beta}\frac{d\beta}{d\lambda} \tag{3.7}$$

or, noting that $\delta\beta = \delta\alpha(d\beta/d\alpha)$

$$R = \frac{\lambda}{\delta\alpha}\frac{d\alpha}{d\lambda} \tag{3.8}$$

Since

$$\frac{d\alpha}{d\lambda} = \frac{\sin\alpha + \sin\beta}{\lambda\cos\alpha} \tag{3.9}$$

resolution may be re-written in the more useful form

$$R = \frac{1}{\delta\alpha}\frac{\sin\alpha + \sin\beta}{\cos\alpha} = \frac{1}{\delta\alpha}\frac{2\tan\theta_B}{(1 - \tan\theta_B\tan\theta)} \tag{3.10}$$

where the $\cos\gamma$ term can be omitted because γ is always small so that $\cos\gamma \approx 1$. For a directly fed spectrograph equation 3.10 becomes

$$R = \frac{f_{\mathrm{col}}}{w}\frac{2\tan\theta_B}{(1 - \tan\theta_B\tan\theta)} = \frac{2B\tan\theta_B}{\theta_S D(1 - \tan\theta_B\tan\theta)} \tag{3.11}$$

This is useful for determining resolving power using telescope parameters and the collimated beam size in conjunction with the slit width (expressed here in the θ_S term which represents the angle the slit subtends on the sky). Should the collimated beam be matched to the projected échelle length, $B = L\cos\alpha$ then equation 3.11 becomes expressible in terms of L as

$$R = \frac{2L\sin\theta_B\cos\theta}{\theta_S D} \tag{3.12}$$

as first noted by Bingham (1979). In the case of the fibre-fed SALT HRS design, equations 3.11 and 3.12 gain the additional FRD factor of ρ, such that:

$$R = \frac{2B\tan\theta_B}{\rho\theta_S D(1 - \tan\theta_B\tan\theta)} = \frac{2L\sin\theta_B\cos\theta}{\rho\theta_S D} \tag{3.13}$$

Equation 3.13 shows that FRD acts to degrade the spectral resolution if the collimated beam size or length of the échelle is not increased.

In reality, the resolving power of a slit limited spectrograph is given by equation 3.12 only in the condition that the seeing disk is large compared to the slit width. This condition provides a uniformly illuminated slit. Should the seeing disk only partly fill the entrance slit (or if the slit is absent), the equations describing resolving power must be modified. Since the geometry of a fibre is circular and it is approximately uniform in surface brightness, the effective resolving power, R' of a fibre with diameter d may be calculated by subdividing the fibre into many slits of width w_i equal to the chord parallel to the dispersion direction, as shown in Figure 3.9. Each

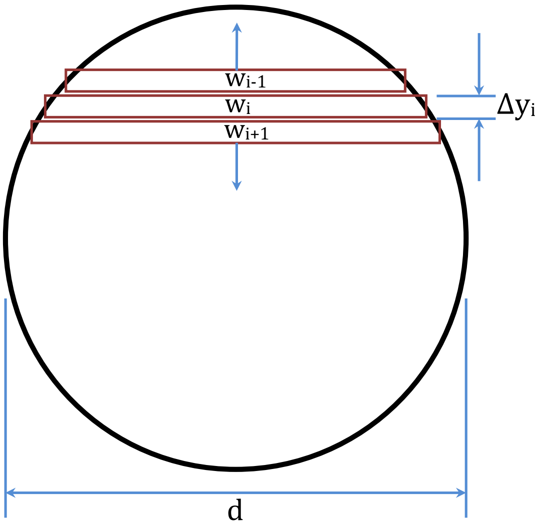

Figure 3.9: Calculation of the flux-weighted fibre width requires division of a fibre into ‘sub-slits’.

‘sub-slit’ is weighted by the fraction of total flux which it encloses. This fraction is dependent on the area of the slit where the normalized area is given by

$$A_i = \frac{w_i \Delta y_i}{\pi \frac{d^2}{4}} \tag{3.14}$$

Summing all such ‘sub-slit’ contributions gives the flux-weighted fibre width, $\bar{w}$:

$$\bar{w} = \sum_{i=1}^{n} w_i A_i \tag{3.15}$$

With infinitely many sub-slits of infinitesimal width ($n \to \infty$, $\Delta y_i \to dy$ and $A_i \to wdy$), this sum may be solved (Vaughnn, 1994) to give

$$\bar{w} = \left(\frac{8}{3\pi} \right) d \tag{3.16}$$

For a slit width w_s the flux-weighted slit width becomes

$$\bar{w} = \frac{8}{3\pi} d \left(1 - \left(1 - \left(\frac{w_s}{d} \right)^2 \right)^{\frac{3}{2}} \right) \frac{1}{T} \tag{3.17}$$

where the T term is the relative transmission of the fibre given by

$$T = 1 + \frac{2}{\pi} \left(\frac{w_s}{d} \left(1 - \left(\frac{w_s}{d} \right)^2 \right)^{\frac{1}{2}} - \cos \left(\frac{w_s}{d} \right) \right) \tag{3.18}$$

For a fully illuminated fibre the slit width $w_s = d$, transmission $T = 1$ and the flux-weighted slit width is $\bar{w} = 8/3\pi d$, so that

$$R' = \frac{d}{\bar{w}} R = \frac{3\pi}{8} R \approx \frac{R}{0.849} \tag{3.19}$$

Since the denominator in equation 3.19 is <1, a fibre can deliver a resolving power somewhat greater than the resolving power achieved using a uniformly illuminated slit of the same width as the fibre diameter.

3.3 Double Scrambler

The scrambling properties of optical fibres are of particular importance for precision radial velocity measurements. Literature suggests that the scrambling (i.e. removal of the memory of the input object position and incident angle) of multi-mode step-index fibres is incomplete (Hunter & Ramsey, 1992). Both azimuthal and radial scrambling properties have been described previously in Section 3.2. Angel et al. (1977) first demonstrated the use of optical fibres for linking a telescope to a spectrograph (first implemented on an instrument in 1979 by Serkowski et al.), and also identified the high level of azimuthal scrambling offered by fibre coupling. However Heacox (1987) demonstrated the ineffective nature of the radial scrambling component. The so-called 'double scrambler' (Hunter & Ramsey, 1992) is designed to overcome this limitation and ensure complete fibre scrambling. For this reason, a double scrambler system is fitted to the high-stability mode of SALT HRS, and it is of particular importance given the variable pupil illumination properties of SALT (as discussed in Section 1.3.2 and shown in Figure 3.10).

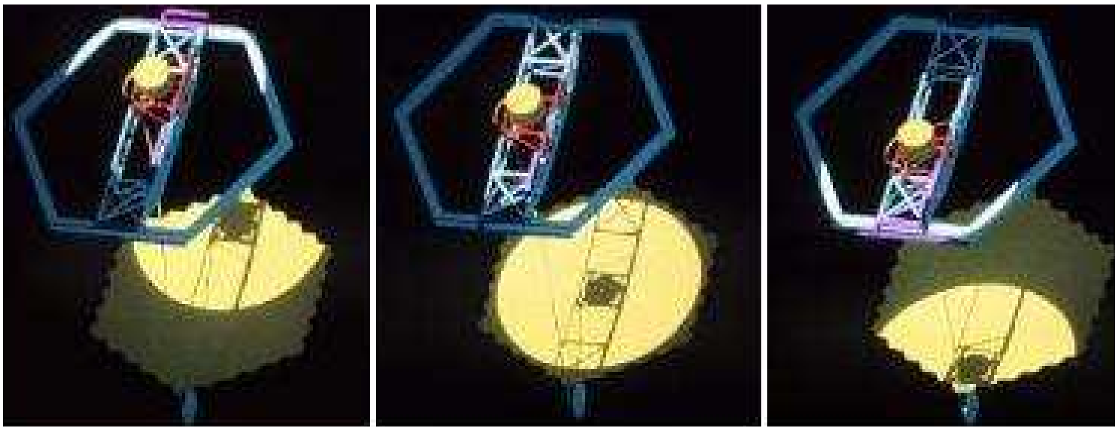

Figure 3.10: The pupil migration across the primary mirror (actually for the HET telescope in this image).

A double scrambler consists of a break in the fibre feeding the instrument, in to which a hemispherical lens is cemented at each side of the junction. The lenses cause all rays exiting the first fibre at a specific angle to be directed to a specific position on the input face of the second fibre. Conversely, rays exiting the first fibre at a given position are directed into various angles at the second fibre entrance. In this way, angles and positions of rays are interchanged. Since the azimuthal scrambling is complete even in short lengths of fibre, this allows a partially scrambled radial component to be converted in to an azimuthal component for which the scrambling is then also good. Figure 3.11 shows a schematic of the double scrambler optics and process, and also illustrates the slight over-filling caused by a double scrambler. Since a fraction of the light is lost by over-filling, the double scrambler is only used on the high-stability mode where efficiency is of lower priority than wavelength stability.

The SALT HRS double scrambler (see Figure 3.12) is designed for use across the entire 370-890nm wavelength range and has an air gap between the two lenses of 112mm. This is to accommodate a 100mm long iodine cell (see Section 6.2.2) which features a 2mm thick fused silica window at each end. The optical design[1] was created using rays at the centre and edges of the

---

[1] Done by Jürgen Schmoll

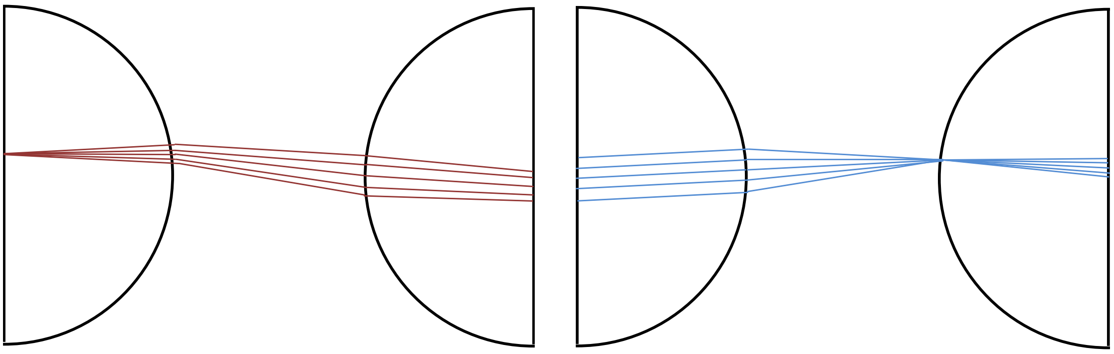

(a) Light exiting at different angles is imaged to different positions

(b) Light exiting at different positions is imaged to different angles

Figure 3.11: A double scrambler with the 'telescope end' on the left and the 'spectrograph end' on the right. Both hemispherical lenses are of focal length f and are separated by distance f (from Hunter & Ramsey, 1992).

350µm fibre (numerical aperture = 0.13) at wavelengths of 370, 550, 700 and 890nm. Both spherical rod lenses have the same diameter (8mm), central thickness (17mm) and radius of curvature (5mm) to simplify manufacture and preserve the paraxial numerical aperture. The efficiency of the double scrambler is influenced by four factors: optical aberrations; numerical aperture losses due to over-filling the fibre input; Fresnel and absorption losses in the lenses/windows; and misalignment error during assembly. The losses are broken down in greater detail in Section 2.2.3 but also summarized in Table 3.1.

Figure 3.12: A ZEMAX ray-tracing diagram of the SALT HRS double scrambler, with identical spherical rod lenses L1 and L2 on the left and right respectively, separated by 112mm to accommodate an optional iodine cell (2mm cell windows shown 100mm apart) or fold mirror for a simultaneous Thorium-Argon source feed.

Table 3.1: Double scrambler efficiency

Wavelength (nm)	Total throughput	Total throughput (with Iodine cell)
370	68.5%	58.2%
550	77.5%	65.8%
700	75.7%	64.3%
890	73.5%	62.4%

3.4 The SALT HRS Fibres

The optical fibres feeding SALT HRS connect to the telescope at the fibre instrument feed (FIF), which is located at the telescope prime focus payload and can be seen in Figure 3.13. The FIF can accommodate 12 optical fibres in two rows of six fibres each. The rows can be moved apart on a set of rails, with a second set of rails allowing for movement in the orthogonal direction, which means the fibres can be positioned anywhere between 15” to 3.7’ apart. The FIF has a re-positioning accuracy of 27µm (corresponding to 0.12”) RMS, meaning a target object can stay accurately centred throughout an exposure (Buckley & Sessions, 2004).

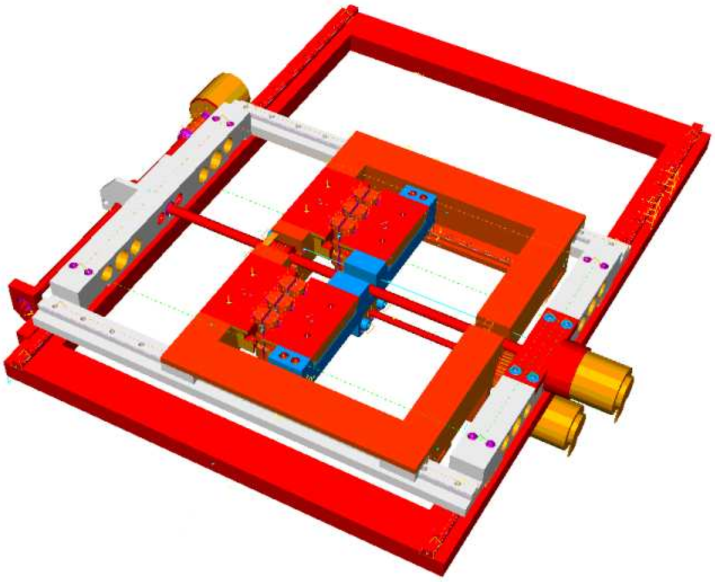

Figure 3.13: A CAD model of the fibre instrument feed (FIF), which can accommodate six pairs of optical fibres. Fibres may be positioned anywhere between 15” to 3.7’ apart with [illegible]” accuracy (Buckley & Sessions, 2004).

Due to its low FRD properties, relatively flat spectral response, low attenuation in the blue and availability in large custom core sizes of 350µm and 500µm, the fibre of choice for SALT HRS is the PolyMicro Technologies FBP broad spectrum step-index fibre. Figure 3.14 shows the transmission at various wavelengths along a 50m length of Polymicro FIP, STU, FVP and FBP fibre types. The FBP fibres clearly represent the best throughput option across the 370-890nm operational range of SALT HRS.

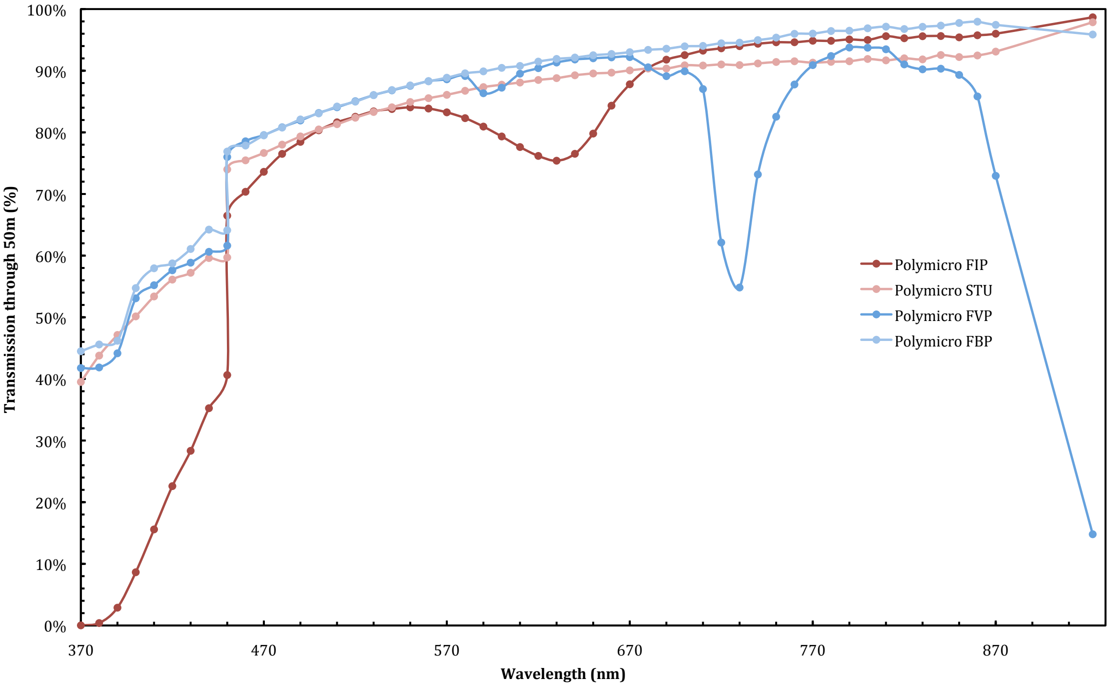

Figure 3.14: The transmission through 50m lengths of a variety of Polymicro fibre types across the SALT HRS operational wavelength range, showing the superior characteristics of the FBP fibre type selected.

The SALT HRS fibre feed carries sixteen fibres in total, eight of which are functional and eight being spares. The fibre properties are shown in Table 3.2. Following the reasoning in Section 3.7, the fibre ends are potted in glass ferrules. For the FRD testing some off-the-shelf ferrules are used, although the final instrument build uses custom sized ferrules with tight tolerances to ensure control over concentricity within the mounting.

To ensure proper protection and strain relief for the delicate fibres on their route from the telescope prime focus to the spectrograph, the fibres are carried in furcated sub-bundles. A bundle of four fibres is housed inside 5mm diameter Miniflex Optical Fibre Protection Tubing

Table 3.2: Details of science fibres used in SALT HRS

Mode	Fibre Type	Number	Nominal Core Diameter (μm)	Buffer Diameter (μm)	Minimum Bend Radius (mm)
Low resolution	FBP500	4	500	590 ± 10	80
Medium resolution	FBP500	4	500	590 ± 10	80
High resolution	FBP350 (custom draw)	4	350	425 ± 8	80
High stability	FBP350 (custom draw)	4	350	425 ± 8	80

(OFPT), which is a segmented polymer tube which prevents tight bending and crushing (both of which are known to increase FRD - Heacox 1987). Four OFPT bundles (plus a spare) are radially arranged around a central strength member made of aramid strands (to act as a strain relief under the weight of the OFPT bundles) and wound in a spiral pattern to avoid differential lengthening under bending. The aramid member is flexible yet more resistant to sagging than equivalent diameter steel wire. 60mm wide Scapa WSL2250 laminated swell tape is wrapped around the cables to form a hygroscopic barrier preventing moisture ingress. An outer PVC-covered steel tube of Adaptaflex SPL20 (with 15.9mm internal diameter) conduit protects the bundle and prevents bending tighter than an 80mm radius bend. Figure 3.15 shows a cross-section through the fibre bundle (not to scale). The estimated mass per unit length of the bundle including all components is 646 gm^{-1}.

Figure 3.16 shows the breakout box which connects to the Adaptaflex outer conduit and manages the individual fibres (protective covers are not shown in the diagram) at the FIF on the tracker. The fibres initially break-out of the Adaptaflex conduit inside the OFPT tubing before separating into individual fibres. The central aramid strength member is attached to a solid beam in the break-out box. PTFE pads hold the delicate fibres in place and grooves are cut in the housing to accommodate hypodermic syringes which protect the fibres from c[illegible]g. A heat shrink wrap also provides additional strain relief.

A strain relief box is included in the fibre run to accommodate changes in fibre length when the cable is moved around during installation or by the cable wrap on the tracker. It consists of a series of loops in the bare fibre which are free to expand and contract as required. The strain

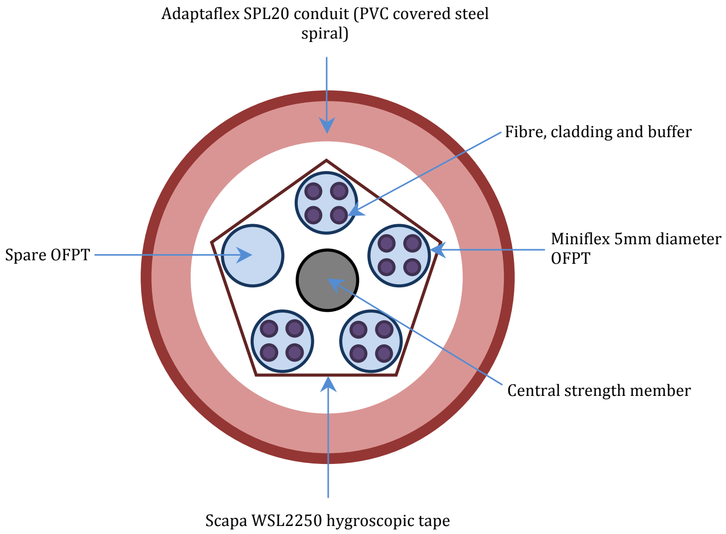

Figure 3.15: A cross section through the fibre bundle showing the optical fibres, OFPT, strength member, hygroscopic tape and outer conduit (not to scale).

relief box is mounted on the optical bench beside the vacuum tank at the instrument end, close to the Iodine cell. The strain relief box can be seen in Figure 3.17.

3.5 Fibre Bundle Assembly

The fibre bundle assembly for SALT HRS optical fibres is simple in concept but non-trivial in implementation. Bare fibres are quite delicate and very unwieldy even at short lengths. Assembling the sixteen fibres, each of length 50m, inside four Miniflex cables wound helically about an aramid cable within a larger Adaptaflex conduit, has provided an assembly challenge with some ingenious solutions required. The first stage is to insert the four fibres inside the Miniflex tubing. This is achieved by attaching a string, placed inside the conduit during manufacturing, to a bundle of four fibres using epoxy adhesive. The bare fibres are laid down a long corridor on a 'carpet' of protective foam, adjacent to the Miniflex conduit, which is held under tension using a clamp at each end. The fibres can then be gently pulled through the conduit (see Figure 3.18*a*). Winding the Miniflex sub-bundles about the aramid strength member is more complex and requires the assistance of a custom-built 'winding machine', consisting of a wooden disk on

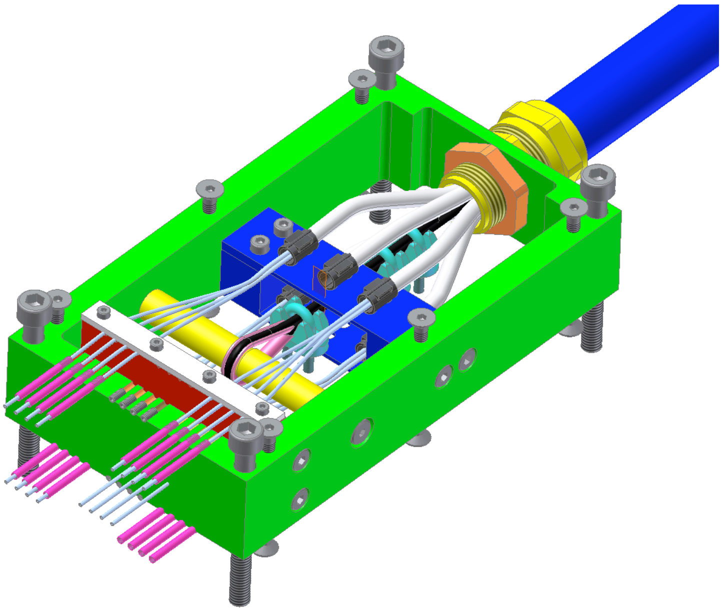

Figure 3.16: A diagram of the fibre break-out box, showing the separation of the protected fibre bundle into Miniflex sub-bundles and finally into individual fibres for connection to the FIF at the telescope. The magenta cylinders towards the bottom left of the diagram are the hypodermic syringes used to protect the fibres.

which are mounted several spools containing the Miniflex conduit (see Figure 3.18*b*). A steel cable is held under tension at each end of the fibre length, as is the aramid strength member. The steel cable passes underneath a trolley on which the winding machine is mounted, and links to a pulley system. As the trolley is pulled along, the steel cable grips the pulley and rotates the wooden disk perpendicular to the direction of trolley motion. The five Miniflex conduits are mounted on spools at equal spacing around the wooden disk circumference. The aramid cable passes through the centre axis of the rotating disk so that the Miniflex is wrapped helically about it as it unwinds from the spools. Each spool is free to rotate on its own bearings, so that it maintains a constant orientation as the disk rotates, thereby preventing torsion in the fibres. A thin adhesive tape is attached at ∼30cm intervals to bind the spiral. The pitch of the spiral is constant since it depends on the length of steel cable across which the trolley is moved, and the gearing of the pulley. The bundled Miniflex/aramid strength member can then be pulled through the outer Adaptaflex conduit using a steel cable attached to excess aramid cabling.

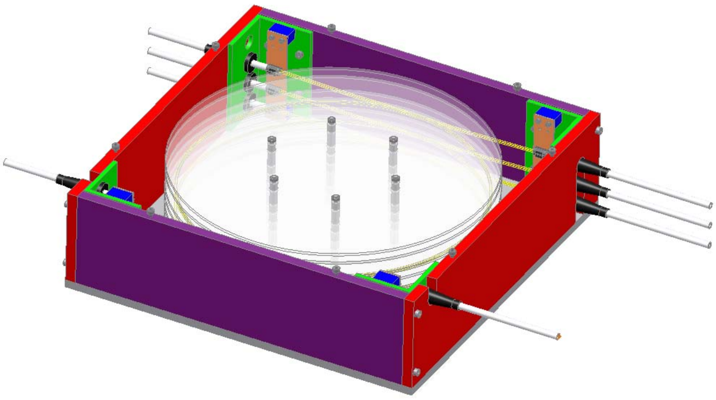

Figure 3.17: A diagram of the fibre strain relief box. Fibre loops inside the box are separated by Perspex disks to avoid tangling and are free to expand and contract with cable movement.

3.6 Operational Modes

The SALT HRS fibre design allows for multiple operational modes: fixed object, nod and shuffle, and dual-object. The first two modes are used for sky background subtraction, which is especially important given the very large aperture of the SALT telescope. The red end of the SALT HRS spectrum (>600nm) suffers particularly given the abundance of bright atmospheric hydroxyl (OH) bands.

The 'fixed object and sky' mode is self-explanatory, using a pair of fibres (A and B) at the telescope focal plane to observe an object in conjunction with a nearby patch of 'empty' sky. The sky fibre spectrum can be simply subtracted from the object fibre, or more commonly, consecutive exposures are taken with the object switched between the two fibres (ABBA model).

'Nod and shuffle' techniques are more involved, but have been shown to substantially improve sky subtraction on faint objects (Glazebrook & Bland-Hawthorn, 2001). A similar method known as beam-switching has seen widespread use in infra-red spectroscopy. In this case a secondary mirror (the primary remains fixed on the object) acts to 'optically chop' between object and sky fields whilst the detector array is continuously read-out. This rapid beam-switching yields essentially contemporaneous sampling at the expense of increased read-out noise from the detector. In the nod and shuffle system, a CCD is used to store two quasi-simultaneous images, with charge shuffled from an illuminated region to a storage region on the chip. The movement of charge

(a) Drawing four fibres through each Miniflex conduit

(b) Bundling the Miniflex conduits around a strength member

Figure 3.18: Fibre bundle assembly (sixteen optical fibres inside four Miniflex tubes, each wrapped around a strength member in a helix).

within the detector does not involve actual read-out (with the associated readout noise penalty) and can be performed several orders of magnitude faster than the read-out process. The charge shuffle is coordinated with a telescope nod so object and sky exposures are imaged side by side on the detector. Figure 3.19 describes the nod and shuffle process. Initially, Fibre A observes the sky with fibre B set to observe the science object. A telescope nod moves fibre A to capture the object and fibre B to observe sky. To ensure the spectra are not muddled, the charge on the CCD shuffles upward so object A is in the position of object B. The telescope is nodded a second time to return the sky to fibre A and the object to fibre B. Charge is shuffled back down the CCD so both object spectra overlap. In this case, the spectrum occupies three adjacent orders, the higher and lower orders show a single sky spectrum each with the central order representing a spectrum from two superimposed object exposures. When using nod and shuffle mode, wavelength calibration and flat-fielding spectra must also be obtained using the same method.

An advantage of the method is that the CCD is read out only once, meaning limited read-noise. Furthermore, since the images are taken using identical optical paths, irregularities in the slit and optics, fringing and flat-field structure (due to pixel response) cancel when removing the sky spectrum from that of the object. For SALT HRS, constraints on the detector area required for charge shuffling and inter-order spacing mean the nod and shuffle mode is only possible using the low-resolution mode (see Section 4.2.4 for more details about inter-order spacing).

Finally, because the separation of the two input fibres can be set by the FIF within the range 15 arcsec to 3.7 arcmin, and the orientation on-sky by the tracker mount, a limited dual-object mode of SALT HRS is possible where a different object is observed down each fibre. This is most likely to be of use when one of the targets is a reference object and when sky subtraction is not critical (the two objects must also be closer than 3.7 arcmin).

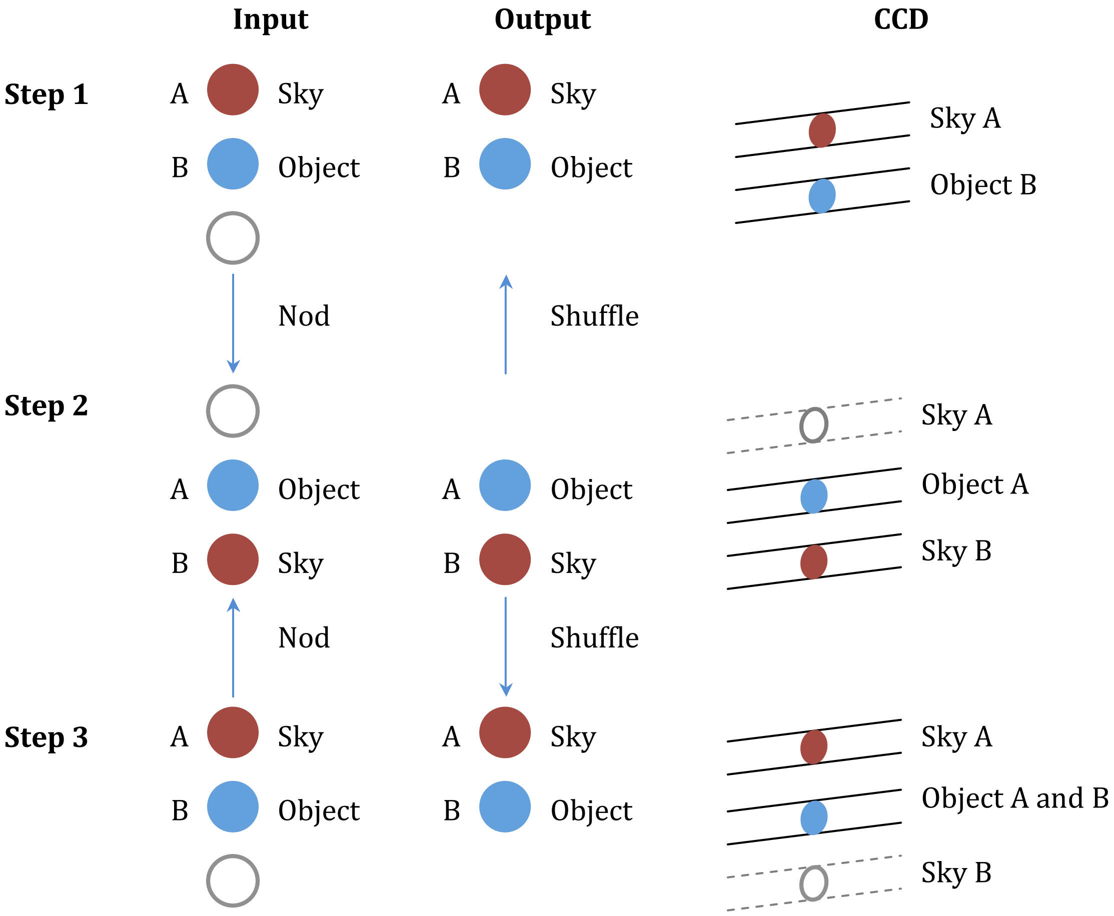

Figure 3.19: The three steps of the nod and shuffle process for obtaining accurate sky background subtraction spectra.

3.7 Focal Ratio Degradation (FRD) Testing

Establishing the FRD losses in the PolyMicro fibres is vital to ensure proper fibre selection and performance, and for creation of throughput models for SALT HRS (see Chapter 2). An extensive series of FRD tests was therefore carried out in the laboratory using the SALT HRS fibres in a variety of setups.

3.7.1 Fibre Polishing

Prior to FRD testing it is essential that the fibre ends are polished flat and smooth to ensure a proper measurement of the output beam pattern (Barden, 1998). In this instance the polishing was performed by hand, having first cleaved the fibre to length using a ThorLabs fibre cutting blade. The fibre end is then polished on progressively smoother 3M carborundum lapping paper with sizes 30µm, 9µm, 3µm and finally 1µm. The polishing is done in a figure-of-eight star shape to ensure that the polished surface is flat rather than concave.

Initially the fibres were seated in syringes, held like a drill bit inside a custom made polishing chuck. A small weight was attached to the fibre above the chuck to provide continuous downforce against the polishing paper. Despite several attempts to polish the fibres in this manner, the edges of the fibres always chipped, presumably as the fibre end changed direction on the paper and gripped along the leading edge.

A different chuck was devised, which held the fibre in place using a weighted top-piece. The top-piece is a metal disk which clips around the fibre using a slot cut into its centre, which is drilled out, and provides consistent downward force as it rests on top of the glass ferrule. To avoid fracturing to the sides of the fibres, they were set in glass ferrules (which could also be polished down at a similar rate to the fibre) using a UV curing epoxy glue (Norland 88). This epoxy expands very little upon setting, causing negligible stress to the fibre. It also has a low co-efficient of thermal expansion and low out-gassing qualities, making it suitable for use in the SALT HRS vacuum environment. The layout of the fibre polishing plate has sufficient room to accommodate the 2mm ferrules making this a practical polishing technique for the production fibre ends (although a fibre polishing machine will likely be used, to ensure consistency and allow for 0.1µm scale finishes). Care was taken at every stage of polishing to clean both sides of the polishing papers, the fibre end and chuck, and the polishing surface, with isopropyl alcohol and optical tissues. This ensures no debris can damage the surface at any stage of the procedure. A star-shaped figure-of-eight pattern is followed across the paper for roughly 3 minutes of polishing at each grade. Illumination of the fibre with an LED allowed inspection through a microscope and photography using a DSLR camera held at the eyepiece. Figure 3.20 shows the illuminated fibre ends after various levels of cleaning and polishing.

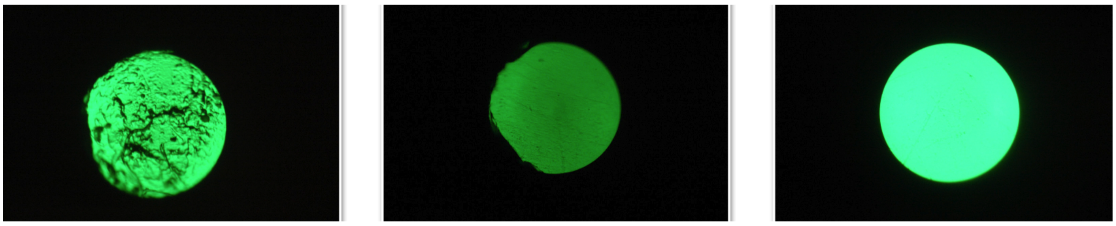

Figure 3.20: The leftmost image shows a dirty, unpolished fibre core. In the centre is a fibre polished to 9µm but with fractured edges. The image on the right shows a clean fibre polished to a 1µm finish.

3.7.2 FRD Measurement Setup

The test rig for FRD measurement is adapted from that used by Schmoll et al. (2003). A white halogen light source feeds a 75µm pinhole via a fibre bundle through a neutral density filter, which acts to remove any spatial effects of the fibre bundle. The light from the pinhole then passes through a Single Lens Reflex (SLR) camera lens set to focus at infinity, which acts to collimate the beam. The beam then passes through a filter wheel, which is used to select the wavelength of light to be used in the test (440nm, 532nm, 730nm or 830nm). This wavelength-filtered beam is focused on to the fibre input using a second SLR camera lens, also focused at infinity. Adjusting the aperture setting of the second camera lens allows the fibre input to be fed at different beam speeds (f/2, f/2.8, f/4, f/5.6, f/8, f/11 and f/16). The input fibre end is glued into a cylindrical glass ferrule (during the polishing process), which is held in place using a custom-made brass ferrule and chuck. This is mounted against a glass window in a five-axis stage atop two orthogonally mounted linear stages. The fibre output end is also mounted in a glass ferrule housed within a larger brass chuck. This ferrule sits inside a brass cylinder which plugs the CCD entrance window. The fibre output face is placed against the Retiga 1300 CCD window. Cargille refractive index matching gel[2] is used at both fibre interfaces to ensure better coupling. Images are taken in a darkened environment, with the stray light sources (such as the light source fan exhaust) covered in opaque fabric. Since the beam is stopped down as f/number

---

[2] http://www.cargille.com/refractivestandards.shtml

increases, the exposure time is also increased accordingly (exposure times were selected so that the peak pixel value is ∼60000 counts, which allows for a large dynamic range without saturating the detector which would occur above 65535 counts). Inconsistent exposure time is acceptable even when comparing spots at different f/number because the data is normalized to unity during processing. Dark frames are taken with the same exposure time as the image frame to which they correspond. The QCapture Suite software[3] was used for .tiff image capture, controlled on a PC via FireWire connection. Data reduction is performed on the image files using an IDL code written by the author.

To ensure the beam is properly collimated, a mirror is inserted in the beam path at 45 °, between the two camera lenses and a small finder-scope is used to inspect the beam. With the scope focused at infinity (both cross-hairs and an extremely distant background in focus), a collimated beam is confirmed when the pinhole is also in focus. Tip/tilt adjustment of the fibre input is verified using a pentaprism, inserted between the two camera lenses. A laser is directed through a pinhole made in a screen and onto the pentaprism. This laser spot is then reflected through 90 ° and through the second camera. The unique design of a pentaprism ensures that the laser is always reflected through the same angle which makes certain that the beam enters the camera lens on-axis. The laser spot is centred on the lens, and is reflected back from the surface of the glass window coupled to the fibre input. The pentaprism directs the reflected spot back on to the laser screen. The tip and tilt of the five-axis stage is adjusted until the reflected spot is in alignment with the pinhole from which the laser beam initially emerges. As long as the beam passes through the centre of the lens, arriving on the optical axis and is not skew, this process ensures the fibre input is completely normal to the input beam. Positioning the fibre input at the correct focus is achieved with an iterative process. The stage is moved in the x plane until spot brightness is at a peak. This method is then repeated in the y and z planes. Whilst the spot is still out of focus, it overfills the fibre input and therefore the peak brightness can be maintained during significant movement of the stage. As the spot approaches focus after further iterations, the input beam size is approximately equal to the fibre core size, so that the peak brightness is not only higher, but also falls off more rapidly with only small movement of the stage (as the

---

[3] http://www.qimaging.com/products/software/

fibre core sized spot quickly moves out of alignment).

3.7.3 Growth Curve Calculations

The IDL code produces a normalized growth curve of the enclosed energy of the spot image. The code first reads in the image and dark frame .tiff files and stores them as arrays. The dark frame is subtracted from the image frame, and centre of gravity calculations determine the coordinates of the centre of the spot. A maximum working radius is calculated by taking the minimum distance from the spot centre to the edge of the frame. Using the equation of a circle and scanning across each pixel in the image, the code works outwards in ever-increasing radii from the centre of the spot, summing the counts in all pixels within an annulus and adding this to any previous counts from previous (smaller) radii circles. This process can be visualized using Figure 3.21. The summed count across all pixels within a given size circle are normalized to unity by dividing the count by the total count in the entire image. This value is then written to a text file alongside the spot diameter in pixels, which can be converted to an f/number using calibration data.

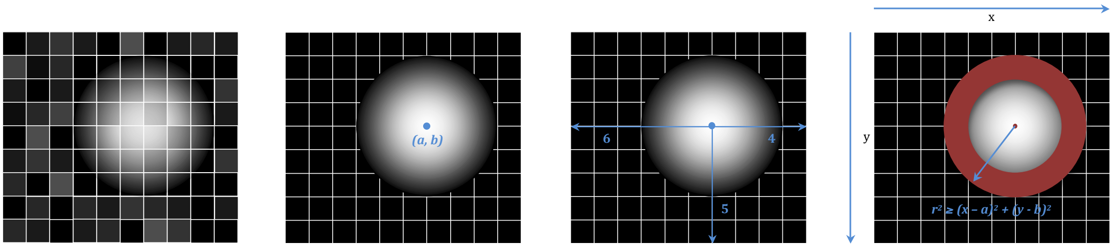

Subtract dark frame

Identify spot centre using centre of gravity calculations (6, 4)

Determine minimum working radius (4 units)

Loop across all pixels in *x* and *y*, using the equation of a circle to determine if a pixel falls within an annulus of given radius, then sum those pixels and create a cumulative total of previous annuli. Write the sum against radius size

Figure 3.21: A graphical representation of the data reduction code processes, as performed on an imported .tiff array.

3.7.4 Parameters of the Tests Conducted

In addition to FRD testing at a variety of input beam speeds and wavelengths, tests were also performed with both various degrees of input beam occultation and with the fibres stressed under curvature. The setup is mounted to a 4m long optical bench so that in the wavelength and input beam speed tests, the 2m extent of the fibre may lie in a straight line from input to output. This is to ensure that any measured throughput losses occur as a result of the measured input beam speed/wavelength and the properties of the fibre, rather than being attributable to modal noise caused by stresses along the fibre length (which may be induced by bending of the fibre). Once integrated with SALT, the instrument will be located in a spectrograph room beneath the telescope, coupled by 50m long optical fibres. The movement of the SALT tracker assembly is discussed at length in Section 1.3.2. The tracking and fibre positioning processes will naturally induce motion in the fibres, and therefore it is crucial that the effect of fibre bending on FRD losses is also characterized. In order to test the effect of fibre bending on FRD, a fibre guide has been designed which allows the bending of the fibre in a controlled way through various radii. A groove in the shape of an upside down letter ‘T’ is cut in to the surface of a thick plastic plate. The design is such that the fibre can easily be positioned into the slot and will naturally spring out to the edge of the undercut as it resists flexure (see Figure 3.22). This ensures that the fibre is uniformly curved and stressed. Clamping the fibre would likely be less precise and would introduce additional stresses at the clamp point. The groove is cut so the remaining ‘uncurled’ fibre is able to lay in a straight line between the fibre input stage and the output at the CCD. The larger diameter curves were tested to begin with, moving inwards to increasingly tight radii, thus ensuring that the fibre is not snapped and that any potential micro-fractures caused by tighter bending are not affecting measurements made at larger radii bends.

Simulation of the SALT pupil migration (discussed in Section 3.3) is performed using two masks. The first of these masks is hexagonal to simulate the primary mirror envelope, with the second mask being slightly smaller and circular to represent the entrance pupil of the SALT spherical aberration corrector (SAC). Since the lenses have a 58mm focal length and an f/4 beam speed is used for the tests, the hexagonal mask - simulating the 11m mirror array - is 14.5mm in diameter (58/4), whilst the circular mask - simulating the 9.2m entrance pupil - is 12.1mm in diameter

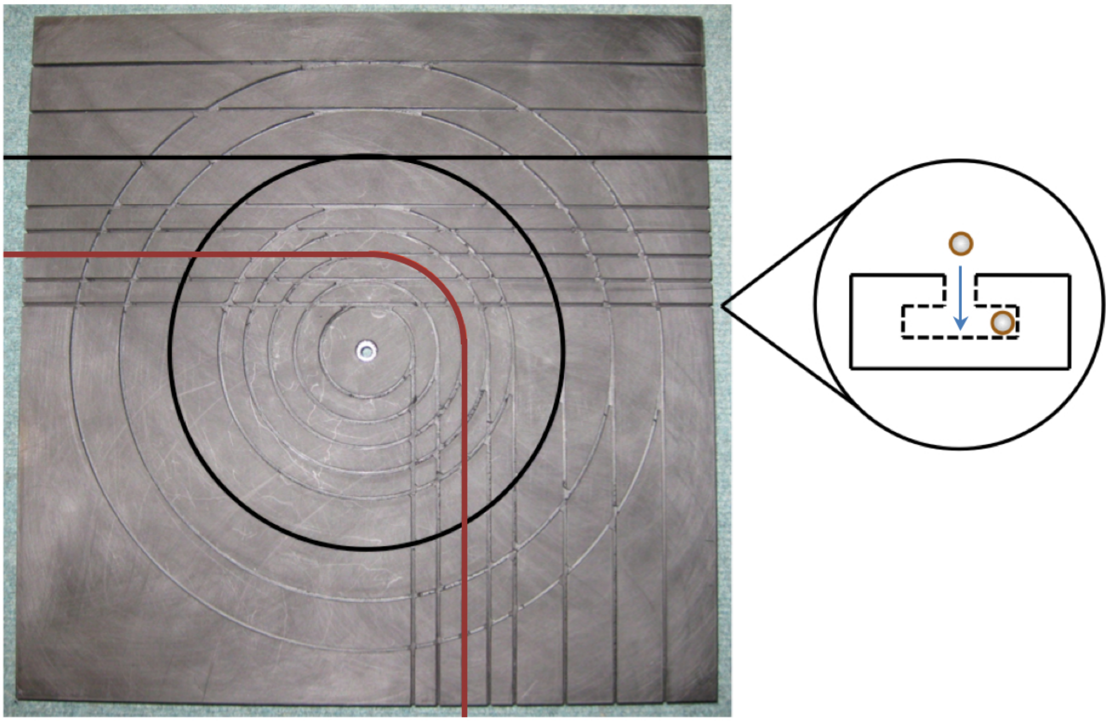

Figure 3.22: Photograph of the grooved platform for bending the fibre to a known radius.

(14.5 × 9.2/11). These masks are cut from sheet metal and clamped inside 2.5" lens holders mounted on linear stages, which are inserted between the two SLR camera lenses where the beam is collimated. The hexagonal mask is aligned so as to be completely illuminated by the beam, and then remains stationary relative to the beam. The circular mask is mounted directly in front of the hexagonal mask, and is translated relative to the beam to simulate the pupil migration of SALT tracking (see Figure 3.23).

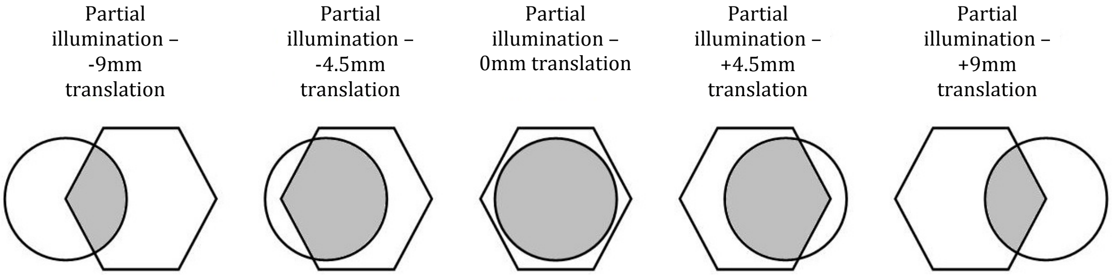

Figure 3.23: Translation of the circular mask in front of the hexagonal mask is used to simulate the pupil migration effect during SALT telescope tracking (see Figure 3.10).

3.7.5 Calibration Procedure

The calibration of the data enables a conversion between the spot size in pixels and the f/number of the output beam. The CCD, with a glass window attached a few millimetres infront, is placed at the re-imaged focus of the pinhole (where the fibre input was located). An image of the resulting defocused spot is then taken at each nominal aperture size (f/2 through f/16). A centre of gravity calculation determines the coordinates of the spot centre, and a slice through the image centre in both x and y directions is taken. An intensity plot of this image slice reveals a top-hat function (see Figure 3.24), the edges of which mark the edges of the spot at its widest point. An average of the spot width in both x and y directions is then plotted against f/number, with a straight line fit being used to calibrate the FRD data (see Figure 3.25). This produces an empirical calibration of the distance from the CCD camera to the detector

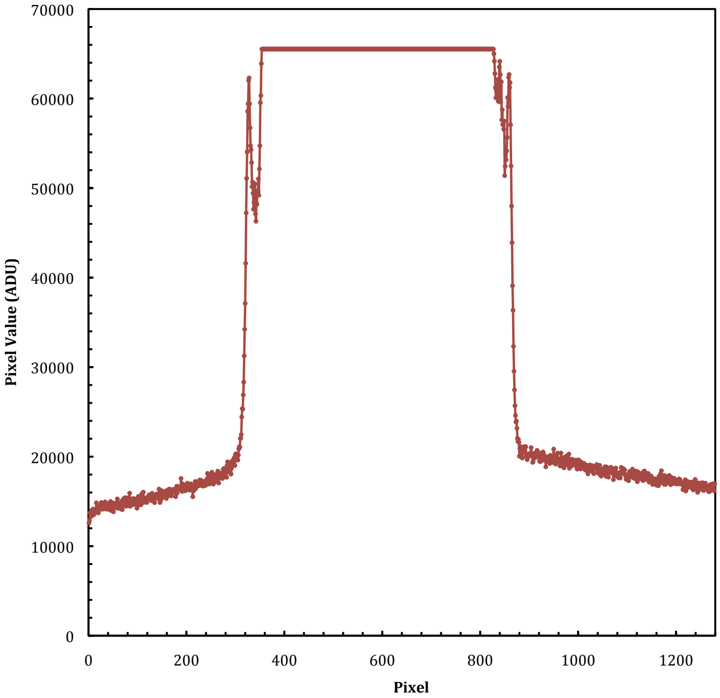

Figure 3.24: An intensity plot through a slice at the defocused spot image centre in the x direction, used for determining the spot diameters at different beam speeds for pixel-to-f/number calibration of the growth curves.

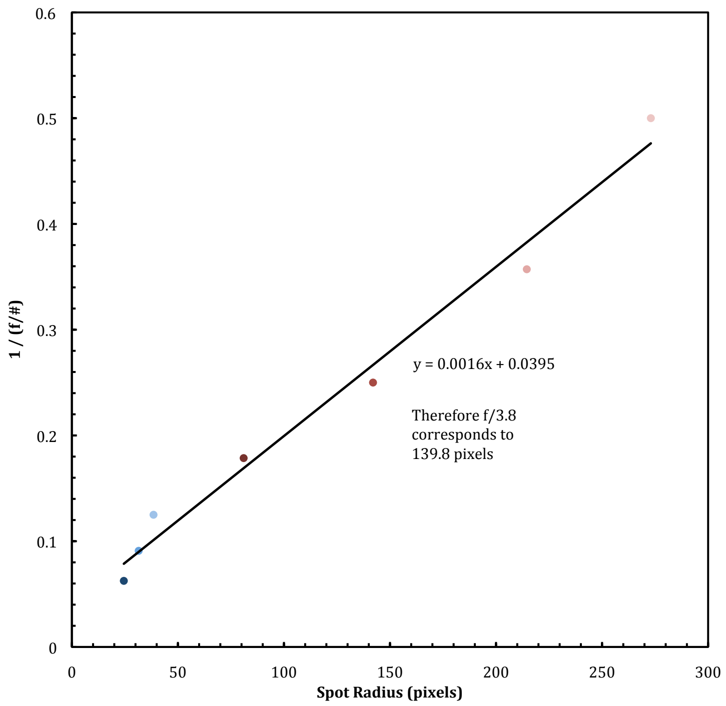

Figure 3.25: Plot of spot size in pixels against 1/(f/number), used to calibrate the FRD growth curves.

3.7.6 Initial Results

Figure 3.26 shows an example growth curve derived from the IDL software of various spot images. The data taken with the smaller diameter FBP400 fibre illuminated at 532nm. Each coloured curve depicts a different input beam speed between f/2 and f/16 and clearly shows the expected trend of increased FRD losses with slower beam inputs. The dashed lines represent a repeat measurement made at each input beam speed and show a reasonable level of repeatability (∼5%), with both data sets at least following the same trend, if not always the same values. Similar growth curves were produced at 440nm, 730nm and 830nm using the same beam inputs and using both FBP400 and the larger FBP600 fibres at 2m in length.

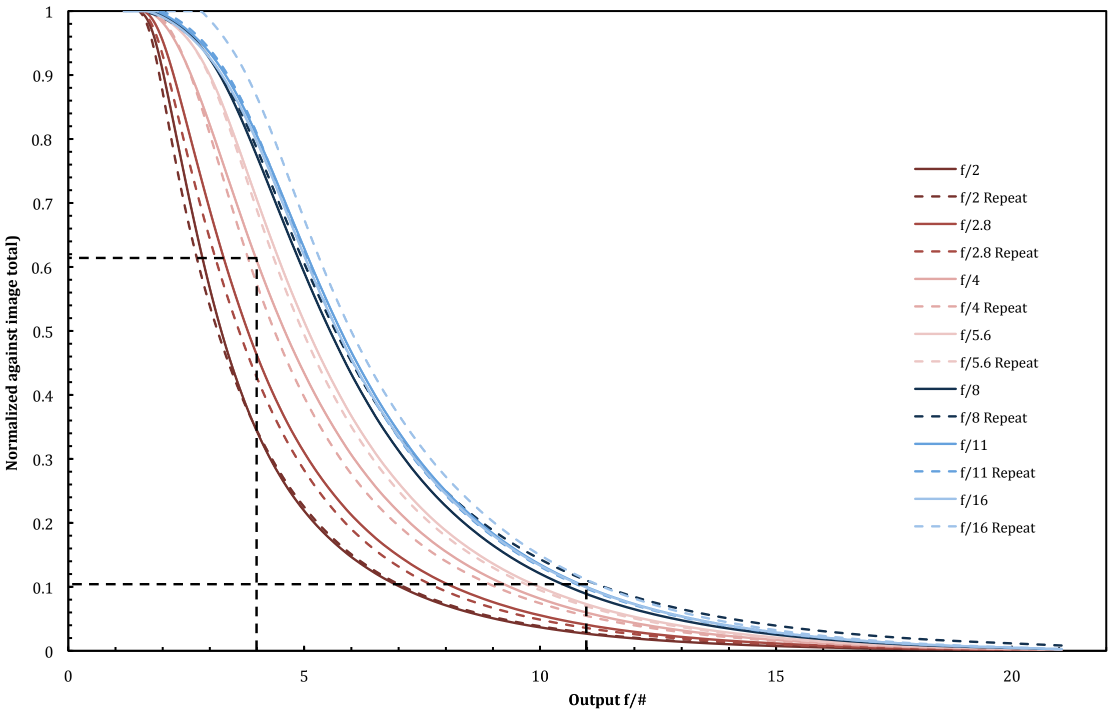

Figure 3.26: Growth curves of encircled energy within fibre image spots taken at a variety of input beam speeds (as shown by the legend) and with 532nm illumination using the FBP400 fibre. A fall-off in throughput into a matched output f-ratio with increasing input f/number is clear. The dashed lines represent throughput when input and output beam speed are equal at f/4 and f/11. The dashed line for f/4 input and output shows 61.5% efficiency, compared to the f/11 case which has only 10.5% efficiency.

Initial Wavelength Results

Our initial results suggest that there is no statistically significant wavelength dependence of FRD with either FBP400 or FBP600 fibres, irrespective of beam input speed. For both fibres, the difference in measured FRD at f/3.8 output between different wavelengths is comparable to the difference in FRD measured between repeat measurements at the same wavelength. For example, for an f/4 input and f/3.8 output with the FBP600 fibre, the difference between a measurement at 532nm and a repeat measurement at 532nm is 3.4%, whereas the difference between a measurement at 532nm and a measurement at 440nm is actually smaller at 0.5% (see Figure 3.27*b*). The lack of dependence of FRD on wavelength agrees with both previous experimental observation (Schmoll et al., 2003) and theory (Lund & Enard, 1984). Figures

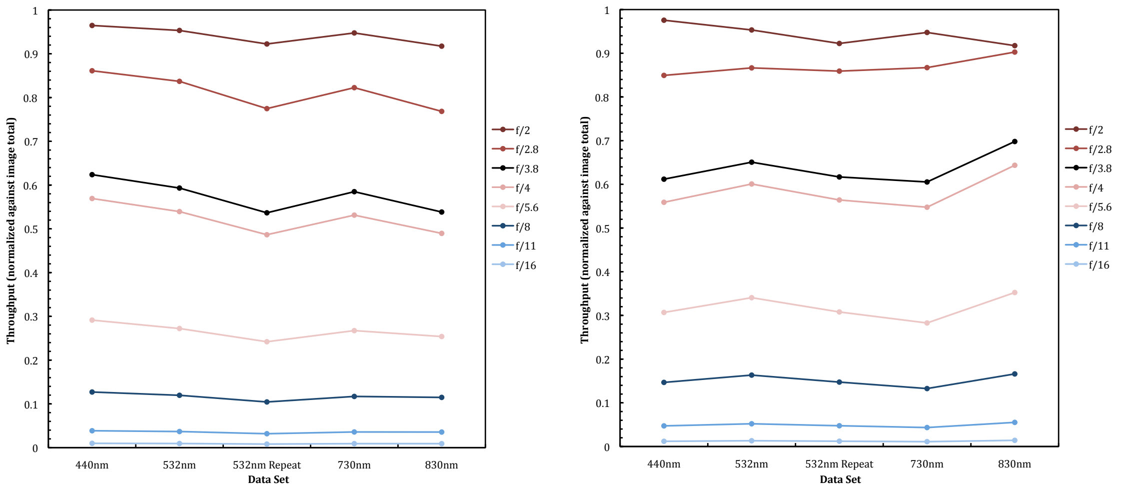

(a) FBP400 fibre

(b) FBP600 fibre

Figure 3.27: Throughput as a function of wavelength for FBP400 and FBP600 fibres with f/4 input and various output beams (shown in the legend). Although losses increase with slower beams there appears to be no discernible difference in performance either between the two fibres or in their performance when illuminated at different wavelengths.

3.27*a* and 3.27*b* show throughput as a function of wavelength for the FBP400 and FBP600 fibres respectively, with f/4 input and a variety of output beam speeds. FRD losses increase with input beam f/number as one would expect.

Initial Bending Results

Figure 3.28 shows the throughput of an FBP600 fibre illuminated at f/4 input and 532nm wavelength light, when bent through a various radii circles. The throughput has a range of 13% between the maximum and minimum value, however it is difficult to determine whether there is a trend of increasing loss at tighter radii, or if the results are merely noisey. The minimum bend radius does indicate peak throughput, but quickly falls off, before rebounding again at even smaller bend radii.

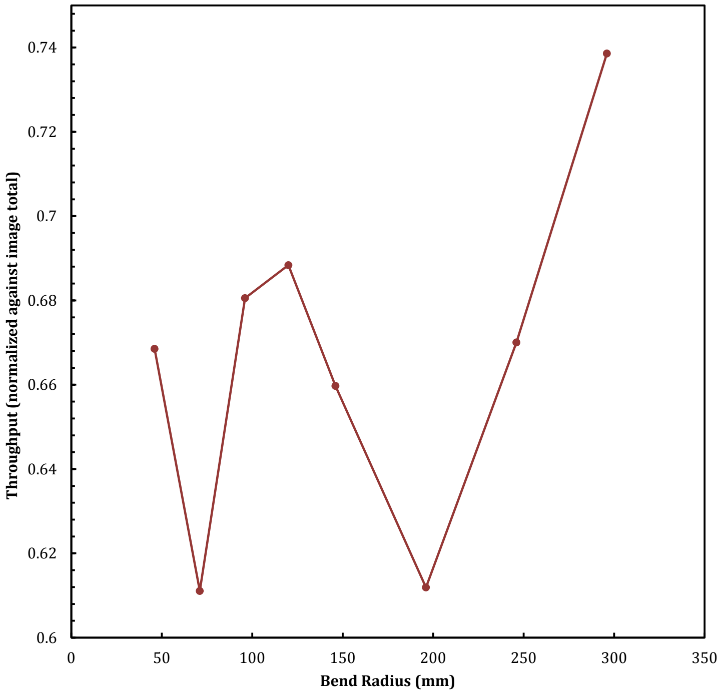

Figure 3.28: The throughput of an FBP600 fibre with f/4 532nm illumination. Although the general trend towards smaller radii bends is a fall-off in throughput, the results are noisey and difficult to interpret.

Initial Variable Pupil Illumination Results

Figure 3.29 shows the encircled energy growth curves for an FBP400 fibre illuminated at 532nm and f/4 input, with various amounts of beam occultation. The dashed line indicates the f/4 input and output condition. With the mask translated by 0.0mm (i.e. no beam occultation, simulating a fully illuminated SALT primary mirror) the FRD efficiency is ~62%. Positive and negative movement of the mask show symmetry in the results when translated by the same amount, and losses <3% can be seen up to ±6.0mm translation. More significant losses of ~10% and ~20% may be seen at ±7.5mm and ±9.0mm respectively. Under maximal illumination the SALT primary mirror has a 55m^2 collecting area. The telescope is not operated with <38m^2 collecting area, which is a 31% reduction from maximum to minimum (David Buckley, private communication). Referring to Figure 3.23 (which is to scale), it is clear that the significant losses occurring at ±7.5mm and ±9.0mm are conditions of occultation far more extreme than any effect will be on-sky at SALT. Therefore it is expected that the variable pupil illumination effect associated with the SALT tracking system is likely to result in only small losses of <3% at the spectrograph due to an increase in FRD. Results with the larger FBP600 fibre are very similar, also showing a <3% loss within the SALT operational range.

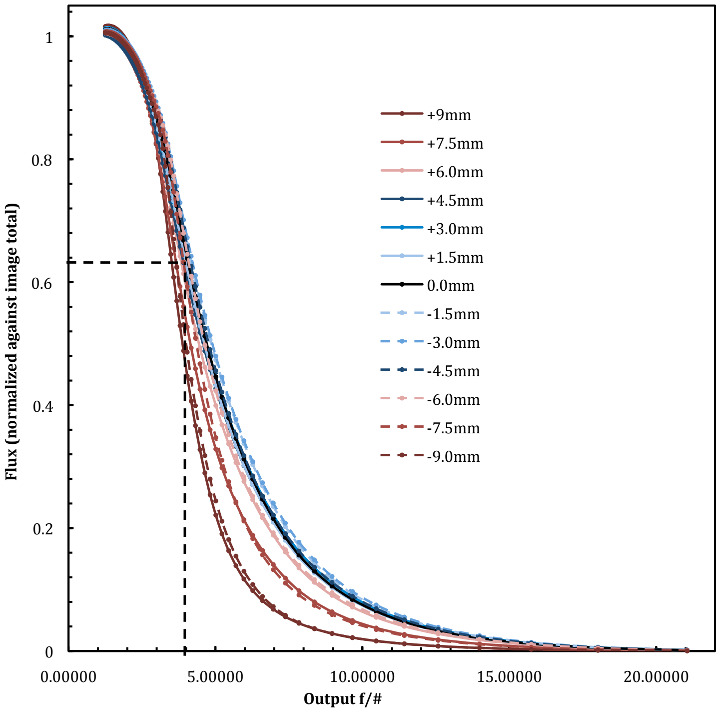

Figure 3.29: A set of encircled energy growth curves for an FBP400 optical fibre illuminated at f/4 input and 532nm, using the hexagonal mask as a variable pupil simulation device at various translations. The dashed line indicates the ∼62% efficiency of the fibre at f/4 output in the absence of beam occultation. Loss can be seen as ∼3% in cases applicable to SALT use, with more major losses outside the SALT operational range.

3.7.7 Discussion of Initial Results

Whilst it is encouraging that there seems to be little dependence on FRD on wavelength, and that pupil occultation should have a negligible effect (apart from in extreme cases, likely beyond those found in SALT HRS usage), it is quite likely that the absolute values of throughput are incorrect. The throughput values produced by the initial FRD tests are alarmingly low, typically in the order of 60% at f/4 in and f/3.85 out, irrespective of core size, wavelength of illumination or variation in input beam profile. This falls well short of the expected losses which were anticipated to be no greater than 5-10%, based on previous tests in literature. Similar tests performed at Potsdam (albeit on smaller core 120µm Polymicro FVP fibres), which provided the inspiration for these measurements, suggest that a fibre fed at f/4 input, polished and immersion gel coupled, with f/3 output cone will have >97% throughput, based on tests with 20 fibres (Schmoll et al., 2003). More recent measurements performed on optical fibres for the prototype VIRUS integral field spectrograph, used the same Polymicro FBP type fibres as measured here, although at the smaller core size of 200µm. The VIRUS measurements (Murphy et al., 2010) suggest an 86% throughput with f/3.65 input and f/3.85 output at 600nm illumination, again using immersion gel coupled fibre ends. Measurements of FBP400 and FBP600 fibres have also been made at the University of Cape Town. With gel coupled ends and input of f/4.58 and output at f/4, the FBP400 shows efficiency of 97.8% through a 9m length of fibre, and 98.5% through 30m of FBP600 fibre (Sessions, 2001). The 'worst case scenario' of throughput found in the literature suggests a roughly 78% efficiency for f/3.8 in and f/4 out, using a now discontinued PolyMicro HHP320 fibre (Ramsey, 1988). It is not mentioned that the fibre ends are coupled using index matched immersion gel, which may account for the lower value for the HHP320. Equally, it may be safe to assume that modern fibres have better FRD properties than a discontinued fibre from more than 20 years ago.

Various sources of error have been considered which may have contributed to the unusually high losses observed in the initial experiments. Since the fibres are inspected after polishing to ensure a uniform and smooth finish, well protected in a ferrule, stored in optical tissue and coupled to the input/output with index matching gel, poor fibre face finish is unlikely to be the problem, although this is known to be a cause of FRD (Carrasco & Parry, 1994). Furthermore, multiple measurements with the fibres removed and re-coupled between exposures has illustrated no large scale shift in results, meaning that any error is likely to be systematic. The iterative focusing method, pentaprism tip/tilt alignment technique and beam collimation checks should also have ensured little error in the optical setup, unless components have shifted between measurements (which would likely create more of a random error). Since the setup is mounted on a vibration isolated bench in a temperature controlled environment any such shifts are likely to be small.

The systematic nature of the error suggests that it may arise in the calibration stage of the measurements, when converting the pupil size in pixels to a corresponding f/number. Our initial calibration relies on the assumption that the camera lens manufacturer has properly calibrated the lens stops used in the tests. Since the aperture of the lens is stopped down using eight rotating blades as a mask, the spot becomes more octagonal in shape as it is stopped down further. This introduces ambiguity in the aperture size measured by the manufacturer - for example minimum diameter, mean diameter or maximum diameter - as well as a change in the beam profile.

Evidence for inaccuracy in the calibration is found using simple trigonometry. Using the notional f/ratios and 58mm focal length of the camera lens, the physical diameter of the aperture can be calculated. The physical diameter of the calibration spot may also be calculated based on the measured spot diameter and physical pixel size. This enables calculation of the distance between the focus (on the front face of the glass window) and the 'far-field' pattern recorded on the CCD chip, which is obviously constant irrespective of beam speed (in reality some small variance may exist due to the passage of the beam through the glass window at different angles, and error in measuring the pixel width of the spot image). Upon performing this calculation for each aperture size, the window-detector distance was found to vary from a minimum of 4.1mm to a maximum of 8.1mm, clearly indicating manufacturer error in the definition of the aperture sizes.

3.8 Revised Setup

To alleviate these issues, the setup was revised to enable a more accurate calibration. The improved setup does not rely on the camera lens aperture stop, instead using a physical mask to define each aperture, with the lens set at its fastest f-stop. Each mask is machined with a circular stop which can be accurately produced and measured to ensure it corresponds to the assumed f/number, removing any ambiguity and ensuring a circular beam at all input speeds, since the camera lenses may be used fully open. Switching from 58mm lenses to longer 135mm focal length lenses also enables the use of aperture stops with larger physical diameters, thereby reducing the effect of any machining tolerances.

The relative size of the detector to the fibre core was also of concern. Since the Retiga camera detector is approximately 7mm across, the pupil spot size is around 3.5mm diameter on the chip. Since the fibre core itself is 0.6mm (when using the FBP600) this is a sizeable fraction of the spot size and the measurement plane may not be far enough away to obtain the true far-field pattern. Moving to a larger 16mm Apogee detector to allow a larger pupil spot size will reduce this effect. Finally there was some concern that the filters used may have caused a slight vignetting of the collimated beam. For this reason, a system of two collimators is employed, where the filter section is decoupled from the camera lenses. Small lenses are first used to collimate a narrow beam which underfills the filter; light is then passed through the camera lens collimator to be masked with the appropriate aperture stops. Figure 3.30 shows the revised layout of the FRD test setup, including the laser tip/tilt alignment testing apparatus.

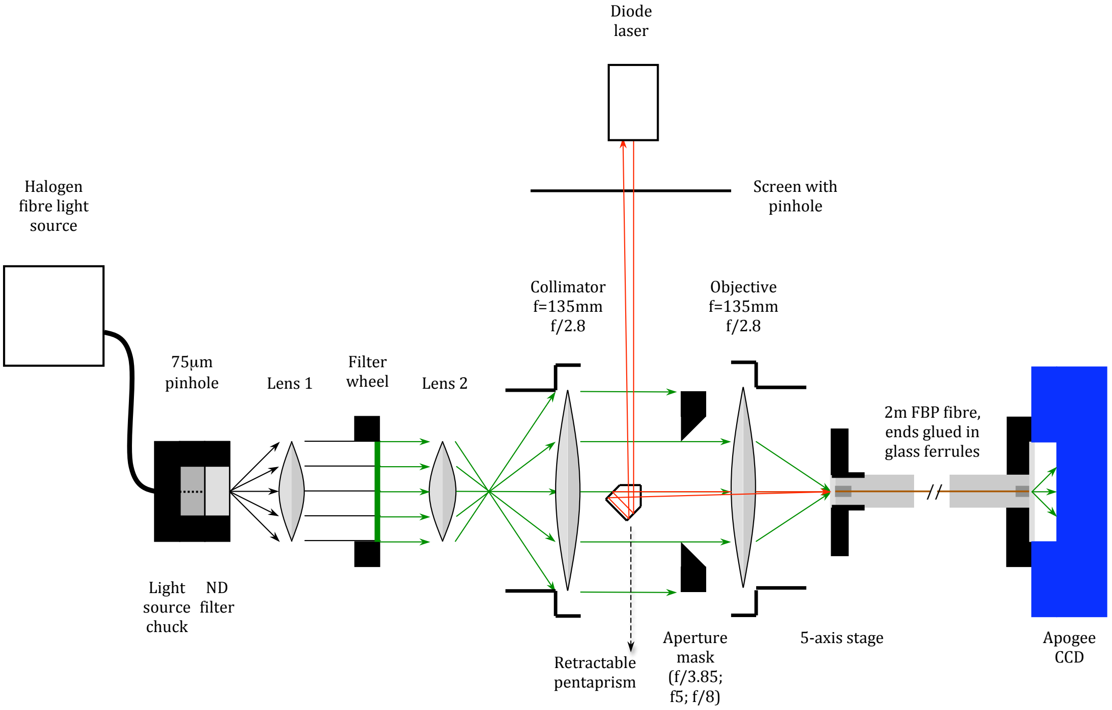

Figure 3.30: Revised FRD test setup featuring a larger area CCD, decoupled wavelength filter optics and separate precision-machined aperture stops.

3.8.1 Revised Results

Following the improvements to the experimental setup detailed above, a considerable increase in throughput was measured. The losses recorded indicated results in much better agreement with values achieved in existing and experimentally comparable literature. Figure [illegible] shows the throughput losses for an FBP400 fibre at 430, 532, 730 and 830nm for f/2.8, f/3.85, f/5 and f/8 input beam speeds, into an f/3.85 output. All four beam speeds show good agreement at all wavelengths and the minimum throughput measured is high at 95.7%. Slightly surprising is the fact that the slower beam speeds achieve similar throughput, although throughput should remain high at all input f-ratios due to the fast f/3.85 output cone.

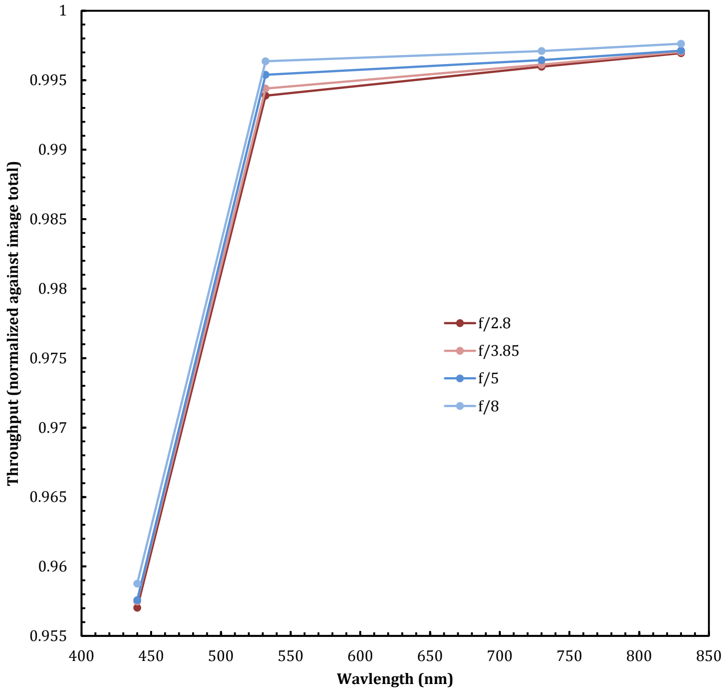

Figure 3.31: The throughput of an F[illegible]0 optical fibre at a variety of wavelengths and input beam speeds, with f/3.85 output. Minimal FRD loss is observed at all wavelengths.

Revised Bending Results

Figure 3.32 shows the effect of fibre bending on FRD losses. An effect similar to that seen in Figure 3.28 is shown, where the throughput shows a marked decrease at 246mm bend radius, before rebounding to higher throughput at still smaller radii. The protective fibre conduit (see Section 3.5) prevents bending beyond a radius of 80mm. Figure 3.28 therefore demonstrates a maximum ∼5% loss over the 'straight fibre' case for the SALT HRS fibres on-sky.

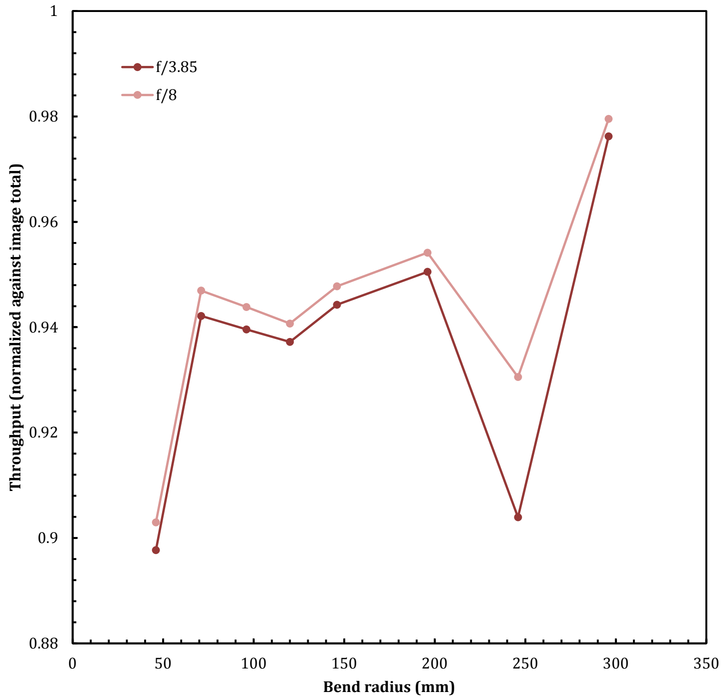

Figure 3.32: The throughput of an FBP400 optical fibre when bent at increasingly smaller radii. A similar effect is seen in Figure 3.28, where the throughput shows a marked decrease at a bend radius of 246mm, before rebounding back to higher efficiency.

Revised Variable Pupil Illumination Results

The effect of variable illumination on an FBP600 fibre with 532nm illumination is shown in Figure 3.33. The input and output of the fibre is f/5, chosen to more readily show any loss effects than the efficient f/3.85 input and output beam case. The peak-to-valley change in throughput is $\leq$ 2%, even at levels of beam occultation beyond that found in use on sky. The asymmetry in the throughput curve is unexpected and may be due to a misalignment of the hexagonal mask. If the bottom side of the hexagonal mask is non-parallel with the optical bench the beam occultation is asymmetric when the mask is translated.

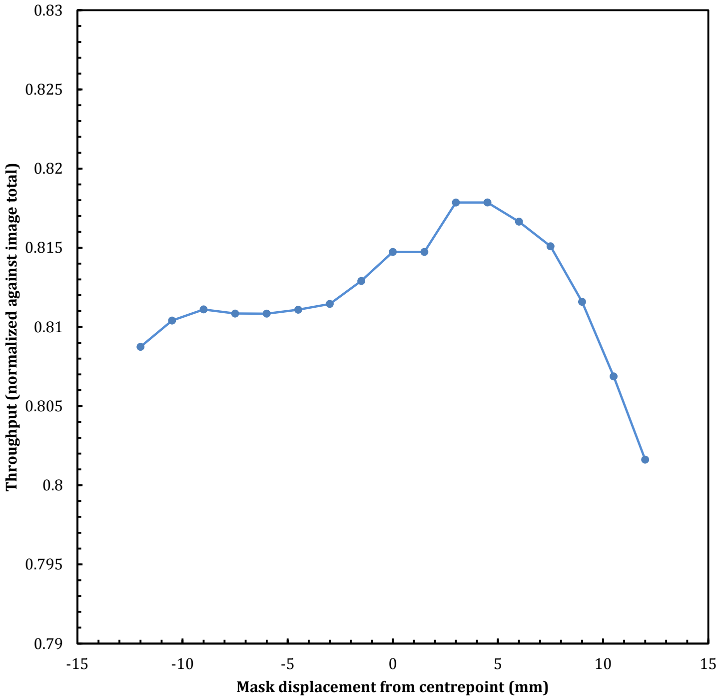

Figure 3.33: The throughput of an FBP600 fibre with 532nm f/5 input illumination and f/5 output when the input beam is masked to simulate SALT variable pupil illumination. Throughput loss of $\leq 2\%$ is seen at the most extreme beam occultation.

3.9 Vacuum Stress Results

In addition to macro-bending of the fibres, micro-bending effects may be caused by the fibre feed-through mechanism which allows the individual fibres to enter the instrument vacuum tank (see Section 3.4). A tight seal is obviously required in order to maintain a reasonable hold time on the tank and this may produce FRD increases. The vacuum feed-through was designed with this issue in mind and employs various strategies to minimize stresses (Figure 3.34). Firstly, the fibre, stripped of its buffer and cladding layers, is glued in to a custom drawn cylindrical glass ferrule. The glue, mentioned previously as Norland 88, is low out-gassing, UV cured and has a low coefficient of thermal expansion. Therefore it is perfectly suited to the vacuum conditions and should expand/contract very little with changes in temperature. The glass ferrule can then be held inside a snug fitting rubber O-ring to ensure a good seal. The O-ring slots inside a tapered hole through a fitting mounted on top of the tank, and a compressing screw (drilled through its centre along the axis of its thread to accommodate the fibre) can be screwed down on top of the ferrule to hold it in place. As the pressure inside the tank decreases the O-ring is sucked deeper into the tapered recess, effectively sealing more tightly as the vacuum increases. This obviously causes a uniform inwards radial pressure on the glass ferrule, some of which may be transmitted to the fibre. However, ideally the protection from the ferrule and the radial nature of the stresses should induce little deviation from a cylindrical core shape.

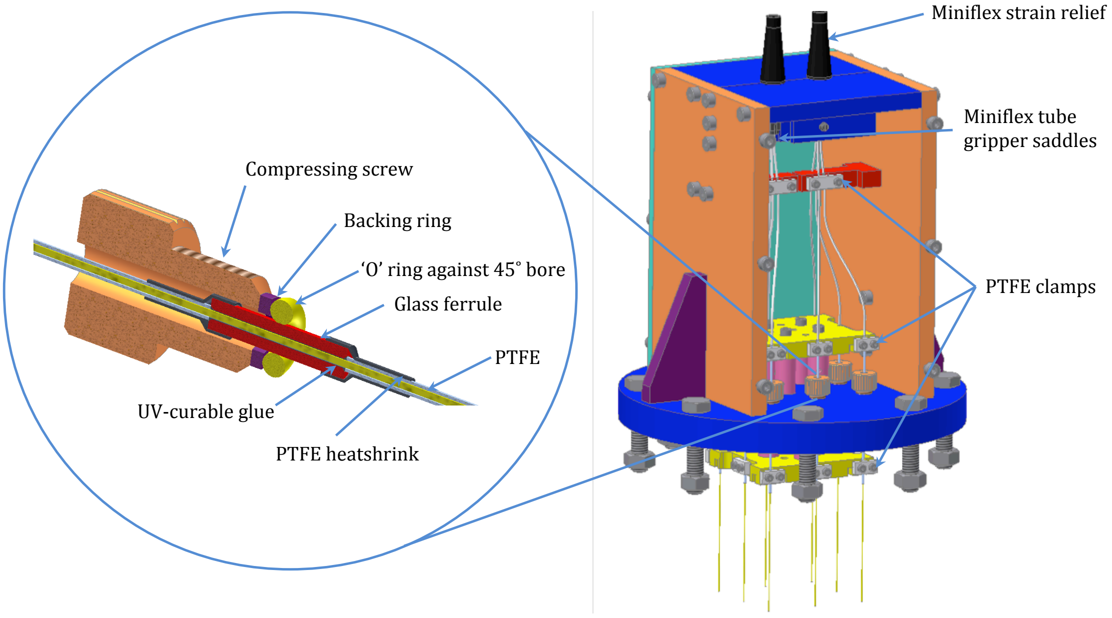

Figure 3.34: A diagram of the custom designed fibre feed-through system.

To test this, the revised FRD rig was transported to an optical bench (eventually intended for the double scrambler optics) beside the vacuum tank in the SALT HRS assembly clean tent. Image spots are recorded from a 532nm illumination at f/2.8, f/3.85, f/5 and f/8 input beam speeds. A 25m length of fibre, this time from the actual draw of FBP500 500µm core fibre to be used in SALT HRS, is used in the tests. For practicality the majority of the fibre remains spooled. A glass ferrule is glued at both input and output faces, and two additional ferrules are glued along the length of the fibre and fitted with O-ring seals. The sealed ferrules are inserted in the fibre feed and looped back out of the tank to the CCD detector output. This means the fibre will experience any vacuum induced stresses twice, unlike in the final construction of the instrument, in which the fibres terminate inside the tank. Whether this two-pass stressing of the fibre induces a cumulative effect on losses is unknown, but it should give an upper limit to the effect. The tank is sealed and pumped down to a pressure of 5.8×10^{-2}mbar. Images of the far-field illumination are recorded at this pressure. Figure 3.35 shows the layout for feeding the fibres through the SALT HRS vacuum tank. The vacuum valve is closed and the hose detached, before gradually opening the valve to allow air in. In this manner a variety of 15 different pressures, roughly doubling each iteration, between 5.8×10^{-2}mbar and 1000mbar (atmospheric pressure) are achieved with images recorded at each pressure.

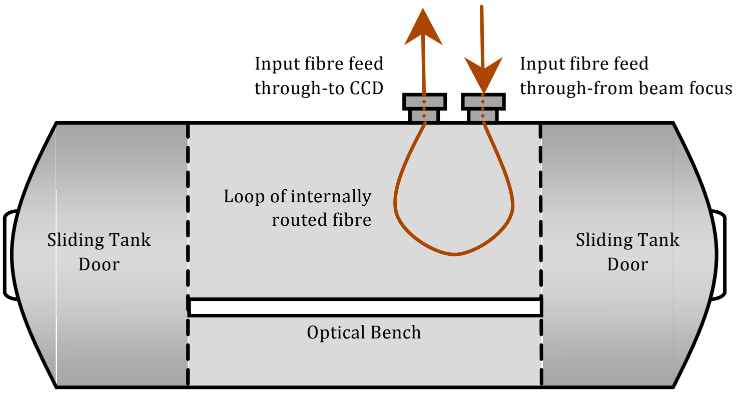

Figure 3.35: A diagram of the FBP500 fibre looped through the vacuum feedthroughs in a double pass, for stress testing.

Figure 3.36 shows the variance of throughput with tank pressure on a logarithmic scale. Although the essential results compare 5.8×10^{-2}mbar performance with FRD measurements taken under atmospheric conditions, different pressures are measured to investigate any systematic relationship between pressure (and hence potential radial stress on the fibre) and the fibre losses. Throughput is measured at f/3.85 output with various input speeds. Losses are less than 1% across all tank pressures and beam input speeds, with very little scatter and no discernible dependency on tank pressure, suggestive that the fibre feed-through is not causing additional FRD.

Although the outcome is positive for the SALT HRS design which utilizes an f/3.85 output cone, stress-induced FRD may still be occurring, but masked by the efficiency of the fibre operating here near its numerical aperture (the f/2.8 feed is also effectively a f/3.85 feed since the fibre cannot propagate any faster incident rays). Since FRD is caused by modes being ‘bumped’ to

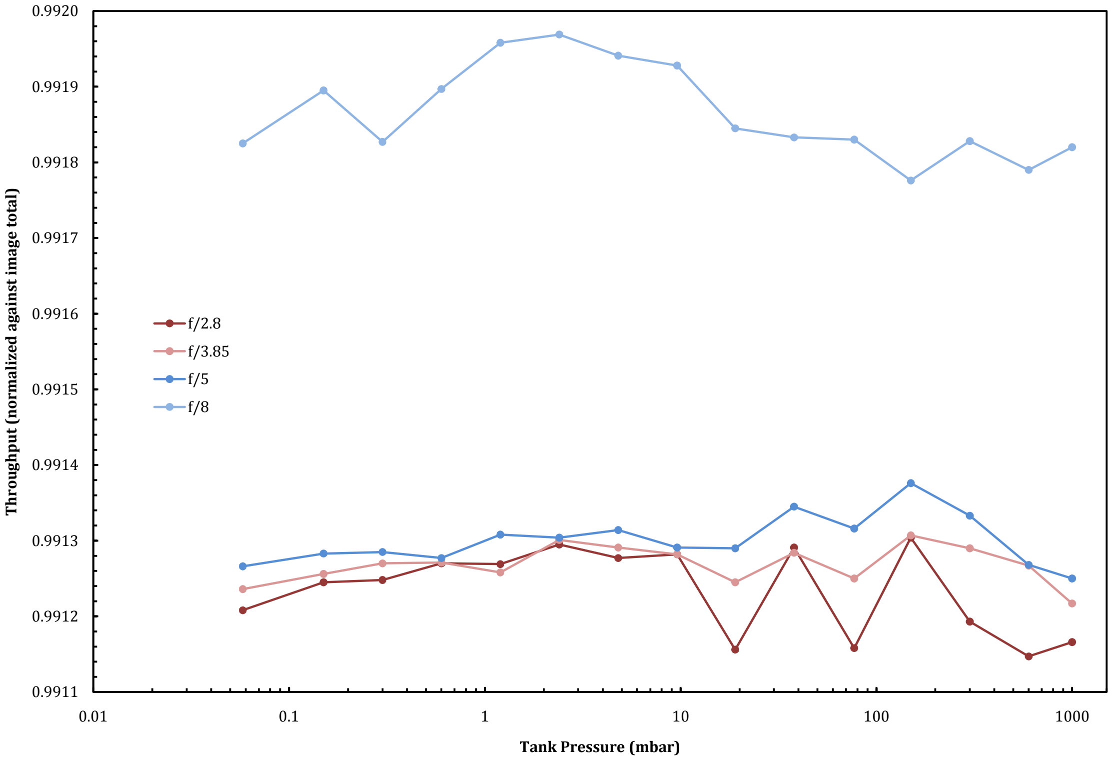

Figure 3.36: A plot of normalized throughput in an f/3.85 output cone as a function of pressure.

higher modes (some of which may be lossy, see Section 3.1), many of the modes 'bumped' up may still exit the fibre within the fast f/3.85 cone. Ergo the effect of slower modes being bumped to faster modes should appear more strongly when measured with a slower output cone which some modes will exit faster than. It is also possible that lossy modes may not be registered at the image plane at all, meaning that the encircled flux at an f/3.85 aperture equivalent is normalized against a lower total flux, thereby obscuring any effect. Figure 3.37 shows the normalized throughput of an f/3.85 input with an f/10 equivalent output. The dark red plotted data labelled 'Frame Normalized' is normalized on an independent frame by frame basis, where the encircled flux is scaled to the total flux in its own frame. However the lighter red plotted data labelled '1000 mbar Frame Normalized' is normalized against the 1000mbar pressure frame total, which is assumed to represent a reference frame of lowest stress and therefore lowest potential loss. Even in this more extreme case there is no discernible trend in FRD loss with pressure. In fact the 1000mbar frame normalized data shows even less scatter than the individually normalized frame data.

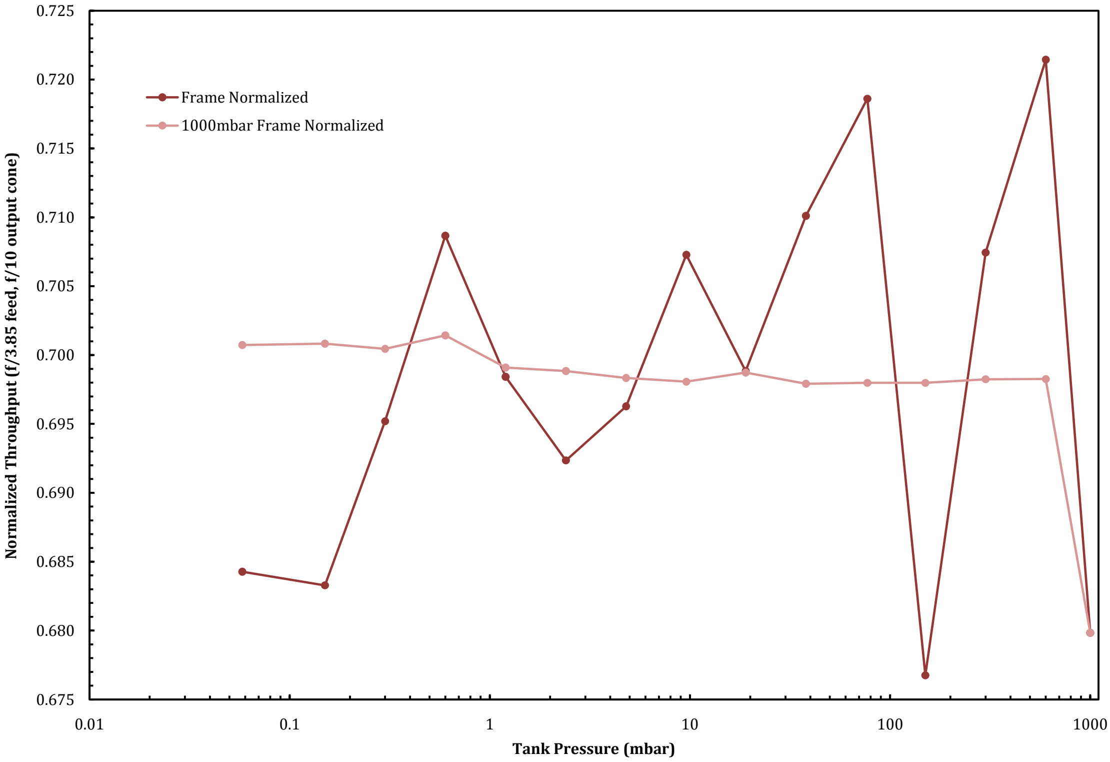

Figure 3.37: A plot of normalized throughput with an f/3.85 input and f/10 equivalent output as a function of pressure.

Figures 3.38*a* and 3.38*b* show the radial symmetry of the far-field patterns across all tank pressures. Figure 3.38*a* plots the intensity in each pixel horizontally across the spot centre (determined using centre of gravity calculations) and shows no significant deviation with pressure. Figure 3.38*b* shows the accompanying image spots using a logarithmic intensity scale to highlight any differences. Again there is a striking symmetry between all images.

3.10 FRD Measurement Repeatability

To ensure reliability of the measurements, consistency checks using repeat measurements performed. Figure 3.39*a* shows the variance in throughput (measured with 25m of spooled FBP500 fibre fed in to the instrument tank under atmospheric pressure with 532nm wavelength illumination and f/3.85 input beam speed in an equivalent output aperture of f/5) with data

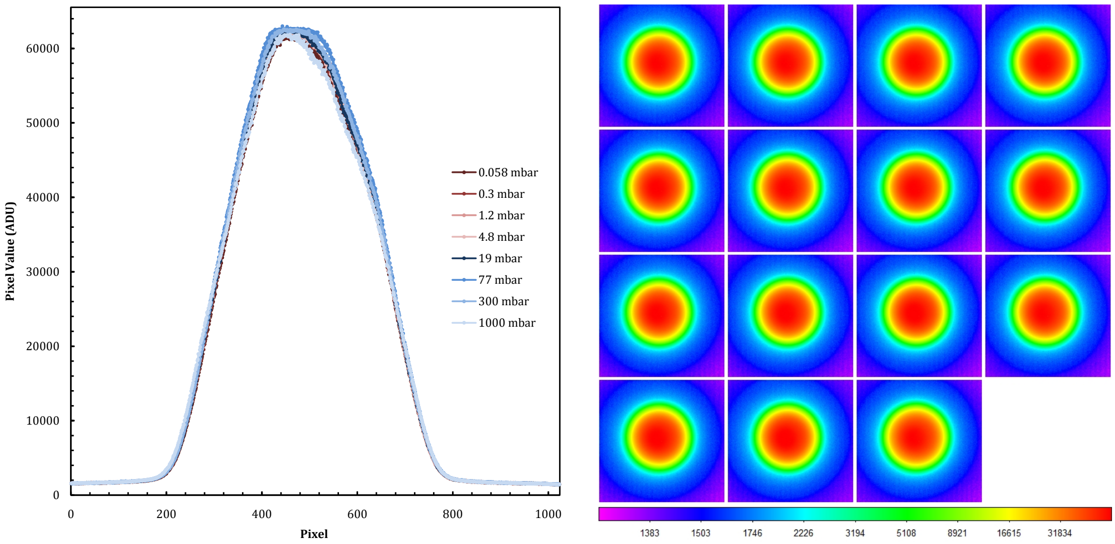

(a) Radial intensity plots horizontally cut through the spot centre

(b) Logarithmic scale spot images (increasing tank pressure from left to right and top to bottom)

Figure 3.38: Both an intensity profile through the centre of each spot and the far-field spot images shown on a logarithmic scale show symmetry at all tank pressures.

taken across a 1 hour time frame. A repeat measurement is made every 6 minutes during which no aspect of the experimental set-up is altered, with the halogen light source remaining on at full intensity and the CCD being held at a constant −25°C. As expected the variance is very small. Although the lamp intensity and detector dark current may fluctuate slightly this is clearly being successfully accounted for in the data reduction. This may be contrasted with the adjacent Figure 3.39*b*, in which the experimental set-up is altered between measurements. Here the fibre is decoupled from the input and output windows, and cleaned of index matching gel. Fresh gel is re-applied and the contacts re-made. The fibre input is also moved away from focus in x, y, z, tip and tilt, and re-focused according to the technique outlined in Section 3.7. Figure 3.39*b* shows greater variance in results and suggests that the quality of the immersion coupling between the polished fibre face and the input/output windows combined with the accuracy of the focus can cause considerable losses of up to ∼10%.

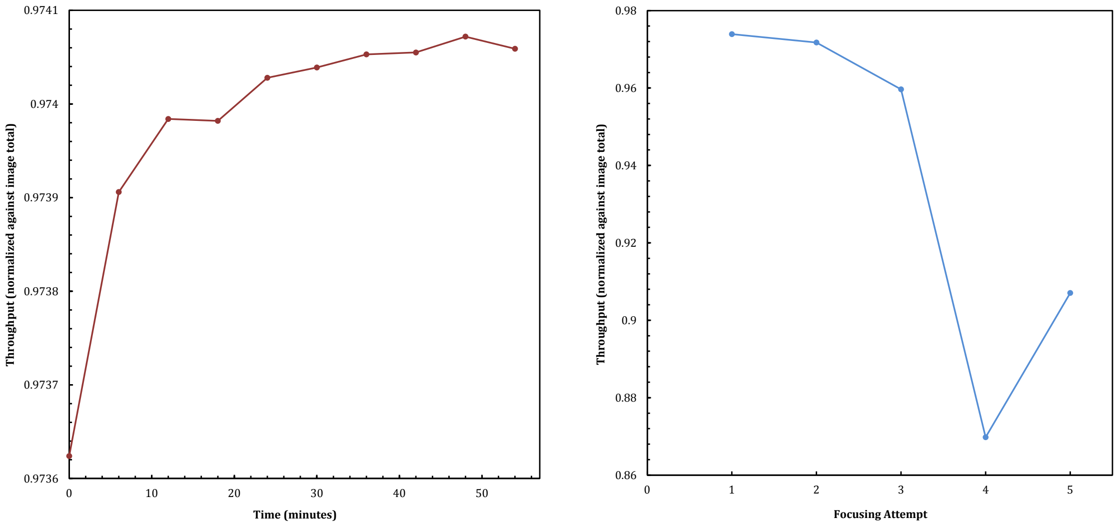

(a) Variability with time

(b) Repeated results with re-coupling and re-focusing

Figure 3.39: The repeatability of results seems excellent when the fibre coupling is not interfered with (3.39*a*, Much larger variability is seen when re-coupling the fibres and re-focusing the optics (3.39*b*).

Chapter 4

Sub-system Pre-Assembly Tests

4.1 Mirror Mount Testing

4.1.1 Introduction

This section discusses the tests which have been conducted on the mechanical mounts as used for the primary collimator mirror, and the red and blue pupil mirrors in SALT HRS.

4.1.2 The Collimator Mirror

The collimator mirror is a concave, off-axis parabolic mirror made from Zerodur, which is sourced from a larger on-axis parabolic parent and has a focal length of 2000 ± 10mm (radius of 4000 ± 20mm). The specifications and performance requirements for the collimator can be seen in Table 4.1. Since SALT HRS is of the so-called 'white pupil' design, the collimator serves two functions.

Table 4.1: Specifications and performance requirements for the collimator mirror

Conic Constant	-1.00 ± 0.05
Off-axis Value	190mm±1mm
Optical Surface Quality	60/40 S/D
Surface Roughness	$\leq$ 2nm RMS
Wavefront Error	$\leq \lambda/5$ peak-to-valley (equivalent to $\leq$ 25nm RMS at 633nm)
Minimum Reflectivity	$R \geq 85\%$ at 373nm
	$R \geq 92\%$ at 400nm
	$R \geq 94\%$ at 450nm
	$R \geq 95\%$ at 500nm
	$R \geq 96\%$ at 550nm
	$R \geq 97\%$ at 600nm
	$R \geq 97.5\%$ above 650nm

Firstly it acts to collimate the incident beam emerging from the slit on to the échelle grating. In addition to this role, the mirror also serves to re-focus the collimated light reflected from the grating to the intermediate focus before the dichroic, in doing so reducing the beam to a more manageable size for splitting into two channels and reducing the size of optics in the cross dispersers, cameras and detectors downstream.

4.1.3 The Pupil Transfer Mirrors

The red and blue arms of SALT HRS both feature a pupil mirror. Each mirror is identical in all aspects (dimensions, material, shape) with the exception of the anti-reflection (AR) coating applied, which is specific to each wavelength range. The mirrors are manufactured from Zerodur (chosen for homogeneity, very low thermal expansion properties and for ease of processing and over-coating) and are concave with a spherical radius of 2222±10mm. The performance requirements are outlined in Table 4.2. The pupil mirrors have been designed as an integrated part of the imaging system and also serve as the active focusing elements of SALT HRS. However the temperature and pressure stability of the instrument should limit the need to refocus, and a ±1mm shift in pupil mirror focus can compensate for the thermal expansion resulting from a 15 to 25°C temperature variation. The pupil mirror focus is also needed to compensate for the change in focus position when the tank changes from atmospheric pressure to a vacuum. This is possible with a 6.2mm shift of the blue mirror and a 5.7mm shift of the red (Barnes et al.,

Table 4.2: Optical performance specifications for the red and blue pupil transfer mirrors

Characteristic	Specification
Optical Surface Quality	60/40 S/D
Surface Roughness	$\leq$ 2nm RMS
Wavefront Error	$\leq \lambda/5$ peak-to-valley (equivalent to $\leq$ 25nm RMS at 633nm)
Minimum Reflectivity	$R \geq 95\%$ at 373nm (up to 560nm) $R \geq 96\%$ at 550nm $R \geq 97\%$ at 600nm $R \geq 97.5\%$ at 650nm (up to 880nm)

2008).

4.1.4 The Mirror Mounts

The pupil mirrors and the collimator share similarly designed mounts. They are mounted on linear stages for re-focusing, and are also adjustable in $x/y/z/\phi/\theta$, via a system of six spring-loaded connections. The six axial supports can be seen in Figures 4.1*a* and 4.1*b* which show the pupil mirror mount, without the mirror. Figure 4.1*c* shows the mount with the mirror installed. There is some parasitic movement in the mirror, so in some cases it is not possible to adjust the mirror in one of the five axes without inducing a smaller motion in another. Tolerancing dictates that the mirror must be adjustable reproducibly to within 0.5mrad (corresponding to 1mm shift at a distance of 2m), and show mechanical stability over extended periods at this position.

The mirror axial supports and adhesives are a proven choice having been used successfully (for échelle mounting) on FEROS and UVES (Dekker, 2000 & Kaufer, 1997). Circular pads machined from Invar (a specialist steel alloy with extremely low coefficient of thermal expansion) are bonded to the mirror using 3M ScotchWeld 2216 B/A epoxy adhesive. A polytetrafluoroethylene (PTFE) jig houses the pads (with sufficient gap for a thin adhesive layer) and allows accurate placement at the appropriate bonding sites on the mirror, also acting as a retainer until the adhesive has cured (which takes several days for each pad). The pad is then screwed onto the spring-loaded axial support assembly at one of six locations on the mirror mount. Figure 4.2 shows a horizontal slice through the support assembly, with the Invar pad (1) shown in red. Once assembled, the

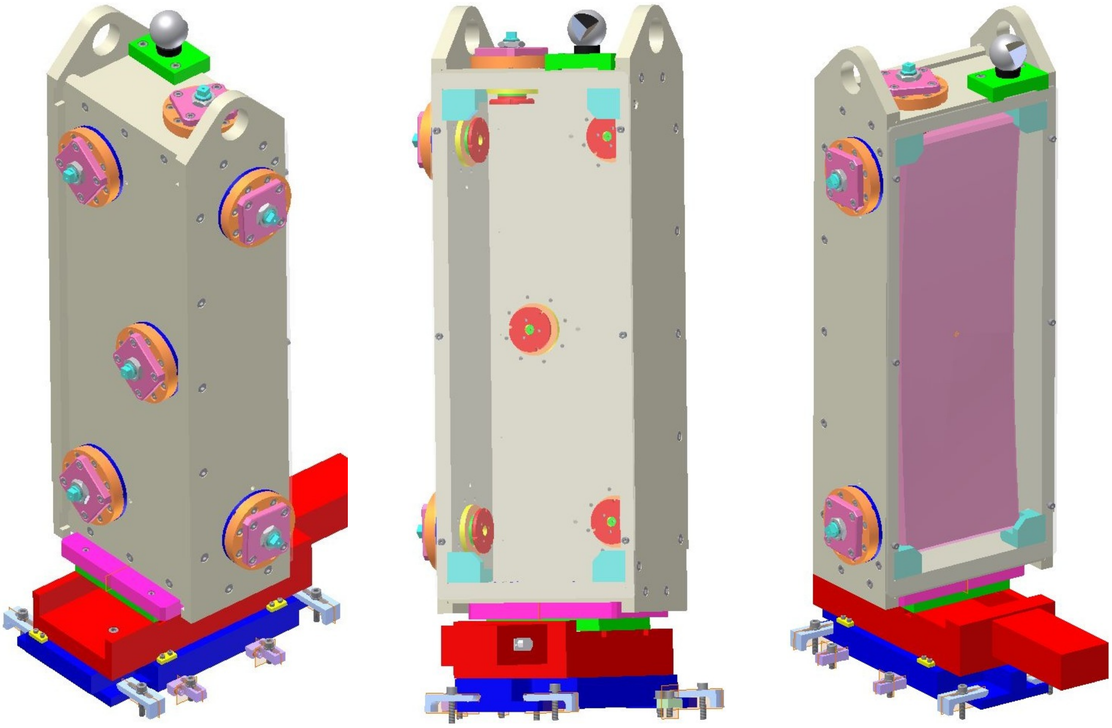

(a) Rear view (b) Front view (c) Mirror installed

Figure 4.1: Design drawings for the pupil mirror mounts, without (a and b) and with (c) the mirror installed.

spring (12) is under compression and held down with the spring retainer (3) which is screwed onto the spigot (4). The end cap (5) covers the assembly and rotation of the adjustment screw (6) applies more or less pressure to the pin (7) which acts against the spring to adjust the mirror positioning.

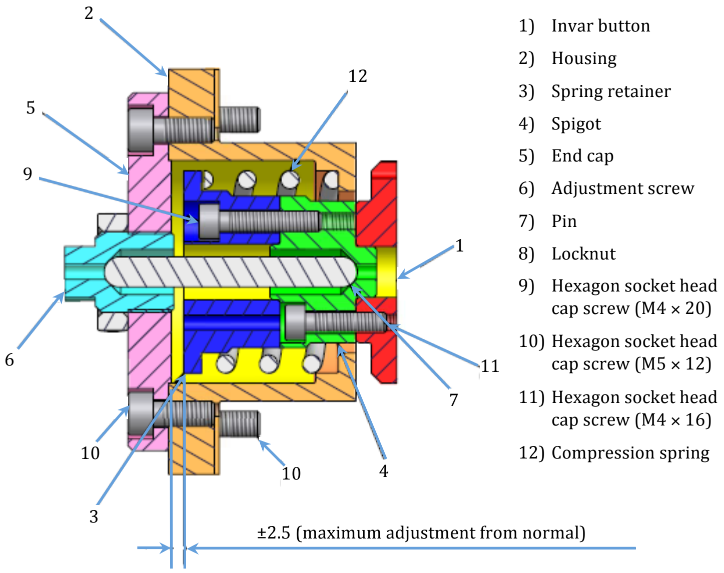

Figure 4.2: A slice-through diagram of the pupil mirror axial support assembly.

4.1.5 Procedure for Verifying the Mirror Mounts

In order to determine if the mounts meet their specifications, their adjustment and stability in $x/y/z/\phi/\theta$ must be verified. To test this, a diode laser was set up to illuminate a variable iris, stopped down to be as small as possible (approximately 0.5mm diameter). The laser beam is reduced in intensity using AstroSolar film attached to a pellicle, which acts as an ND6 neutral density filter (reducing illumination by a factor of $\times 10^{-6}$). Without the filters in place the beam is too bright and saturates the detector even at the minimum possible 0.03s exposure time. Use of the filters allows the exposure time to be set to 0.26s. This is a trade off between a long exposure (in which the optics may potentially move, smearing the spot, and background light becomes more apparent) and a short exposure (in which the effect of the mechanical shutter opening and closing can introduce noise). The stopped down, filtered beam is incident at an arbitrary angle, θ, on a flat mirror fixed to a 'dummy' pupil mirror (a rectangular block of solid aluminium) using double-sided adhesive tape. The pinhole image is reflected back through angle θ on to a large Apogee CCD camera. The CCD is cooled with a Peltier cooler to $-25°$C and attached to a computer running the MaxIM CCD control software. A second beam, also illuminating a maximally stopped down variable iris and ND6 filtered, is incident directly on the detector from a diode laser source 2m away, and acts as a reference to determine whether any potential image spot motion is caused by creep in the mirror mount assembly or by motion or expansion/contraction of the CCD mount. The experimental setup can be seen as a schematic in Figure 4.3. The CCD is separated from the mirror by 2m on a vibration-damped optical bench, which is housed in a clean tent. The laser incident on the mirror is situated as close as possible to the mirror to reduce the effect of diffraction through the pinhole on the spot image size at the detector. By tracking movement in the reflected image spot relative to the reference spot over time, and knowing the pixel size of 15µm, it is possible to measure the angular stability in the mounting. MaxIM software was setup with a script to control the detector exposures. The first pair of exposures were delayed by 200s to provide sufficient time to vacate the room and turn

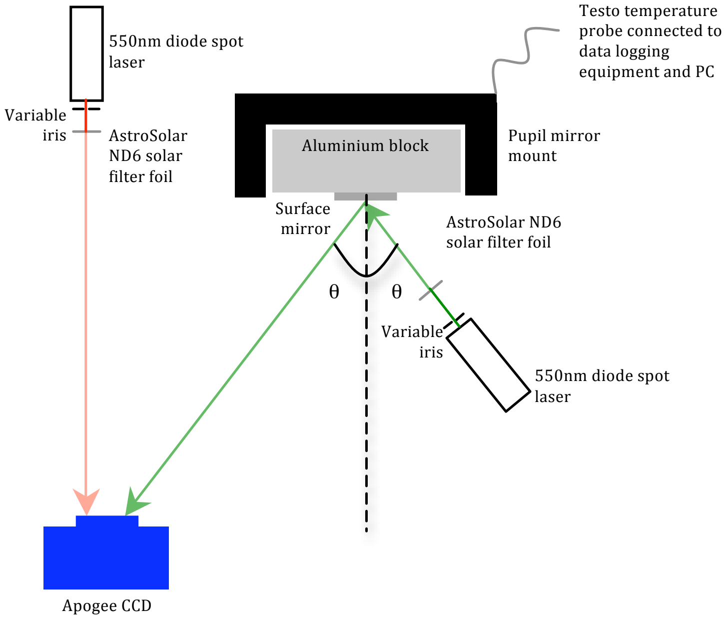

Figure 4.3: A schematic of the experimental setup used to determine pupil mirror mount positional stability, using a reflected and 'reference' laser spot incident on a CCD.

off any light sources aside from the lasers. After this delay, a 0.26s exposure is taken, repeated 354 times at an interval of 500s (plus the CCD read-out time), totaling a measurement period of 180000s (or 50 hours).

Although it would be trivial to use a centre-of-gravity type calculation to locate the barycentre of a single spot in an image, the addition of a second, potentially moving spot adds an additional level of complexity in software reduction. For this reason, the SExtractor[1] package was used to quickly locate the barycentre of each spot. The minimum detection threshold of 5σ and a minimum area of 10000 connected pixels above this threshold ensure stray light and ghost spot elimination, whilst successfully identifying both spots and their Airy disks. Figure 4.4 shows the image spots from the first exposure. The setup is not perfect in that the laser spot is relatively large on the detector, which degrades spatial resolution. However, the use of smaller pinholes resulted in larger diffraction effects at the CCD, with any pinhole below 250µm leading to a spot size greater than the detector area. In principle it should be possible to focus the pinhole image to a smaller spot size at the detector using lenses, but given the objective of the test, the introduction of additional elements and sources of motion were deemed inappropriate. Whilst

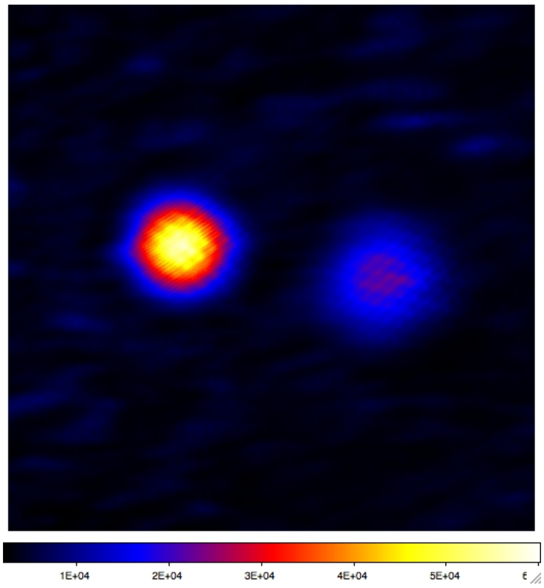

Figure 4.4: An image of both laser spots on the CCD (the brighter spot is that of the reflected beam).

---

[1] www.astromatic.net/pubsvn/software/sextractor/trunk/doc/sextractor.pdf

motion in the CCD mount may be accounted for using the reference laser (both spots would move by the same amount in the same direction), it is difficult to disentangle motion of either of the laser mounts from that of the mirror mount. Motion in the mirror laser spot may be due to a movement of either the laser mount or the mirror mount. However, it is extremely unlikely that motion of the laser mounts or detector will counteract motion in the mirror mount in every exposure, acting sometimes to reinforce and sometimes to reduce the inferred motion, so the data still provide a useful upper limit on the mount movement.

4.1.6 Mirror Mount Tracking Results

Figure shows the position of both reference (red) and mirror reflected (blue) image spots on the CCD as calculated from the SExtractor barycentre parameter. Figure shows the change in distance between spot centre co-ordinates relative to its position in the previous exposure, with the dark red data representing the reflected beam and the light red data representing the reference beam. This may be compared with ambient air temperature fluctuations recorded over the same period (shown in blue). The reference beam shows greater scatter and occasional large motions of over several tenths of a millimetre. In contrast, the reflected mirror beam shows much greater stability, with all movements less than 0.03mm. Of course since this motion is relative, it is entirely possible that the movement of the reference beam is random about a point whereas the motion of the more stable beam could be cumulative drift in one direction. There is no obvious correlation between the small scale temperature variations measured and beam motion.

An alternative means of establishing the stability may be seen in Figures 4.6*a* 4.6*b*, in which the spot separation and change in spot separation relative to the previous frame are plotted respectively. This has the advantage of removing any dependency on CCD mount motion (since both spots see the same change in position) but removes the possibility of decoupling reference and reflected beam motions. Again it is difficult to determine any temperature dependency.

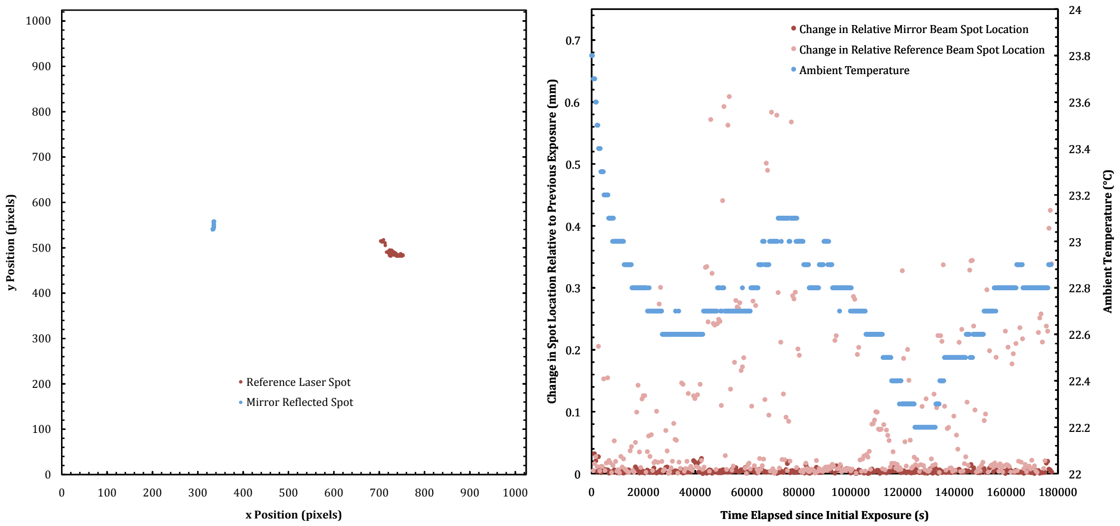

(a) Spot position on CCD

(b) Physical change in spot position

Figure 4.5: Plot *a* shows the locations on the CCD of both reference and reflected image spots for each exposure. Plot *b* shows the positional change in spot locations with respect to their location in the previous exposure, with an overplot of the mount temperature. The reference beam is seen to be considerably more unstable than the reflected beam.

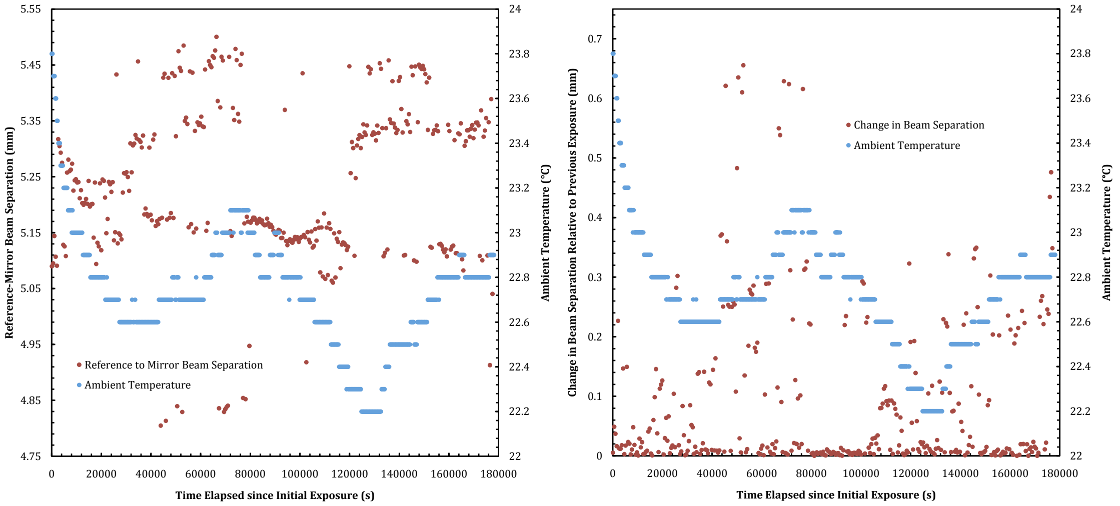

(a) Spot separation and temperature with respect to time

(b) Change in spot position relative to the previous frame

Figure 4.6: Relative change in spot separations and separations relative to the previous frame also reveal no major change in the mirror mount stability and again show no temperature fluctuation correlation.

4.1.7 Mirror Mount Stability Conclusion

Although the means for determining mirror stability are imperfect, as described previously, all three measurements indicate that the mirror mount is stable to within <1mm. The range in minimum to maximum calculated beam separation is seen to be $\leq$ 0.7mm, which falls within the required positional tolerance. It is also encouraging that the majority of this motion shown by any of the measurement techniques is more likely due to motion in the reference beam than due to mirror creep, as seen in Figure [illegible] Despite the fact that mirror stability cannot be decoupled from laser mount and/or CCD mount stability, it is highly unlikely that motion in the latter two mounts always acts to cancel out the motion of the mirror. Based on these facts the mirror mounts appear suitably stable for use in the final instrument. Although no trend of mirror stability with temperature is obvious, the mount should perform in the instrument with at least a similar level of, if not greater stability than these tests indicate, given the more controlled temperature environment of the SALT HRS vacuum tank.

4.2 Volume Phase Holographic Gratings (VPHGs)

4.2.1 Advantages of VPHGs

Volume phase holographic gratings (or VPHGs), are used in SALT HRS for cross dispersion purposes. In contrast to the gratings commonly used in astronomical instrumentation, the dispersion is caused by Bragg diffraction as light passes through a volume of modulated refractive index, rather than by a surface relief structure (as in the main échelle grating). According to Barden et al. (1999), VPHGs offer several advantages such as:

- The efficiency envelope is governed by Bragg diffraction and can be tuned by adjustment of the grating angle for different wavelengths or diffraction orders.
- High line density, high dispersion transmission gratings can still have very high efficiency.
- Large gratings can be produced using the manufacturing technology (similar to that used

for head-up displays).

- The grating is robust, environmentally stable and can easily be cleaned, since the grating itself is sandwiched between two glass substrates (unlike a surface relief grating which cannot easily be cleaned). The substrates can also be anti-reflection (AR) coated.
- Customized grating parameters are viable since each grating is an original rather than a replica from a master.
- It is possible to create both reflection and transmission grating geometries.

4.2.2 VPHG Manufacture and Theory

As mentioned in Section 4.2.1, VPHGs may be designed for either reflective or transmissive diffraction. The light is diffracted as it passes through a thin (typically 3-30μm) layer of material, which is usually dichromated gelatin (DCG). To produce diffraction, the refractive index of the DCG layer is modulated sinusoidally along the grating by exposing it to an interferogram produced by two large collimated laser beams, before being chemically processed (Baldry et al., 2004). In the simplest case, the modulated fringes are perpendicular ($\gamma = 90°$) or parallel ($\gamma = 0°$) to the grating surface for a transmission or reflection grating respectively. However it is also possible to design gratings with fringes inclined with respect to the grating surface in both cases. The SALT HRS gratings are both transmissive, and operate according to the principles demonstrated in Figure 4.7. The spacing of the fringes (given by $1/\nu$) define the dispersion properties according to the usual (transmission) grating equation

$$m\nu\lambda = \sin\alpha - \sin\beta \tag{4.1}$$

where m is the diffraction order, ν is the grating frequency, λ is the wavelength, α is the angle of incidence and β the angle of diffraction. The fringe pattern in the bulk of the DCG layer diffracts the light and therefore its properties determine how much light goes into a particular order depending on grating volume and fringe contrast. The angular and spectral relationship of the incident light relative to the Bragg condition also determine the grating performance. For a

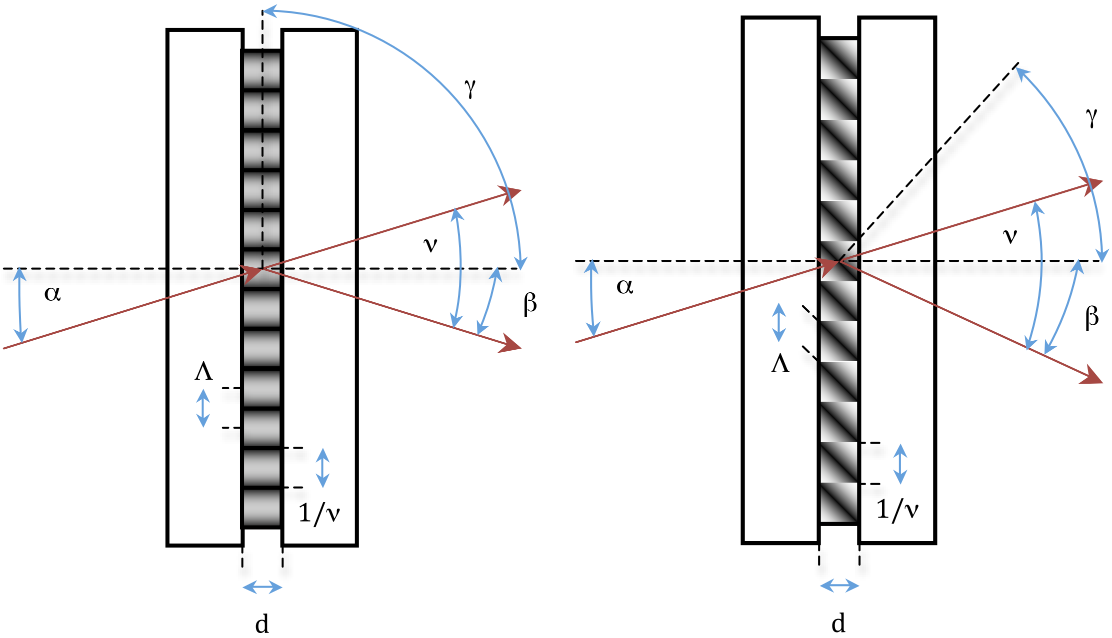

(a) Fringes perpendicular to grating surface ($\gamma = 90°$)

(b) Fringes tilted with respect to grating surface

Figure 4.7: The operation of transmission volume phase holographic gratings with fringes normal (*a*) and inclined (*b*) to the grating surface (adapted from Barden et al. 1998).

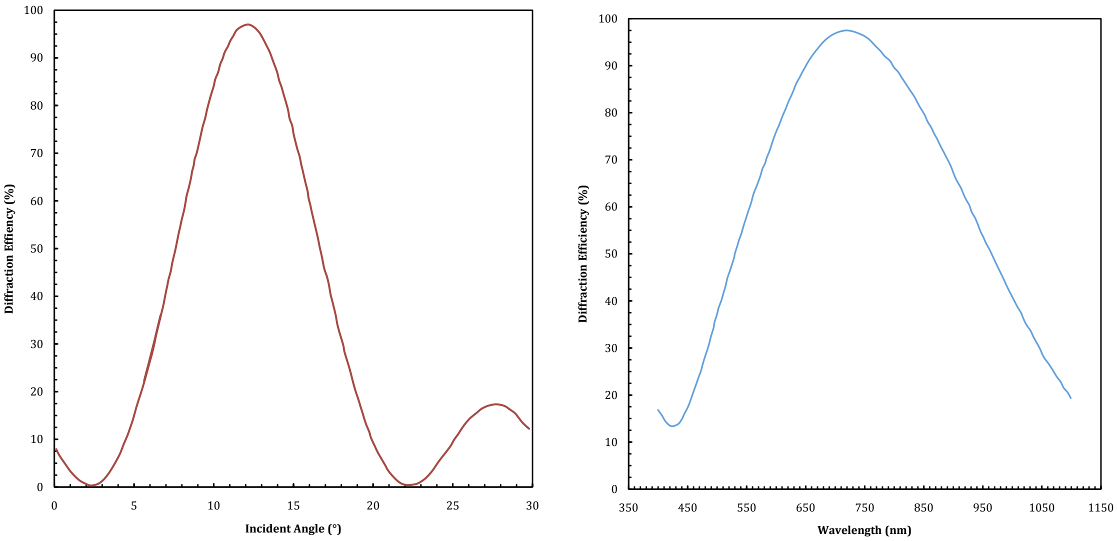

(a) Angular envelope (b) Wavelength envelope

Figure 4.8: Simulated VPH grating Bragg envelopes for a grating with 600 lines mm^{-1} and first order operation at 700nm (from Barden et al. 1998).

grating like the ones used in the SALT HRS design, with parallel fringes normal to the grating surface, the Bragg condition is

$$m\lambda = \Lambda 2 n_2 \sin\left(\alpha_{2B}\right) \tag{4.2}$$

where Λ is the fringe spacing of the grating equal to $1/\nu$ for fringe planes orthogonal to the grating surface, α_{2B} is the Bragg angle in the grating medium and n_2 is the refractive index of the DCG layer. The Bragg condition is met when m is an integer and λ and Λ are such that the angles of incidence and diffraction are equal and opposite with respect to the surface normal. Light illuminating the grating at angles significantly outside the Bragg condition pass through without being diffracted, but near the Bragg condition angle there exists a range of angles which still produce efficient diffraction. The range of angles and wavelengths for which diffraction occurs fall within these ‘Bragg envelopes’. This allows a VPH grating to be tuned towards performance in certain orders or for different wavelength ranges by tilting the grating with respect to the beam. Figures 4.8*a* and 4.8*b* show simulated angular and wavelength Bragg envelopes for a hypothetical grating which has 600 lines mm^{-1} and first order operation at 700nm.

Figure 4.7 shows the grating parameters which, along with the amplitude of the refractive index modulation, Δn_2 (with n_0 the refractive index of air and n_1 the index of the covering subtrate material), determine performance. The refractive index for the DCG layer is expressed as an average value approximated by

$$n_2(x,z) = n_2 + \Delta n_2 \cos\left[\left(\frac{2\pi}{\Lambda}\right)(x\sin(\gamma) + z\cos(\gamma))\right] \tag{4.3}$$

where γ is the fringe tilt angle (see Figure 4.7 for further explanation). Equation 4.3 holds true for the case of slanted or unslanted fringe patterns. The substrate index of refraction is given by n_1 and the refractive index for air, n_0, is assumed to be equal to 1. Kogelnik (1969) used a coupled-wave analysis using zeroth and first order diffraction only to develop a set of equations for grating efficiency which are sufficiently accurate for many cases and are commonly used in the VPH grating industry. Diffraction efficiency, spectral bandwidth and angular bandwidth are functions of the intensity of the grating modulations (Δn_2) and the thickness (d) of the grating volume. As a general rule of thumb, gratings optimized for wide bandpass have lower peak diffraction efficiency than those with a narrower bandpass. In the case of transmission gratings, with orthogonal (non-slanted) volume fringe structure, the efficiency at the first order Bragg condition is given by

$$\eta = \sin^2\left[\frac{\pi \Delta n_2 d}{\lambda \cos(\alpha_{2B})}\right] \tag{4.4}$$

where α_{2B} is the angle within the grating material (where $n = n_2$) given by the relation

$$\sin(\alpha_{2B}) = \left(\frac{n_0}{n_2}\right)\sin(\alpha) \tag{4.5}$$

The first order spectral bandwidth and angular bandwidth may be approximated by

$$\left(\frac{\Delta\lambda_{\mathrm{FWHM}}}{\lambda}\right) \sim \left(\frac{\Lambda}{d}\right)\cot(\alpha_{2B}) \tag{4.6}$$

and

$$\Delta\alpha_{\mathrm{FWHM}} \sim \left(\frac{\Lambda}{d}\right) \tag{4.7}$$

where α_{FWHM} is expressed in radians.

4.2.3 Cross-dispersion Methods

Various dispersion elements (such as prisms, gratings or a combination of the two in the form of a grism) may be used to provide cross-dispersion, as long as the dispersion direction is orthogonal to that of the échelle (so that the various orders are separated). The amount of inter-order spacing and the spectral format is dependent on the cross-disperser used. A cross-disperser typically has considerably lower angular dispersion than the main échelle grating. If the focal length of the camera used to image the dispersed orders onto the detector is f_{cam} then the physical inter-order separation may be given by

$$\Delta y = f_{\mathrm{cam}} \frac{d\beta}{d\lambda_{\mathrm{XD}}} \Delta\lambda_{\mathrm{FSR}} = f_{\mathrm{cam}} \frac{d\beta}{d\lambda_{\mathrm{XD}}} \frac{\lambda_B^2}{2\sigma \sin\theta_B \cos\theta} \tag{4.8}$$

where $d\beta/d\lambda_{\mathrm{XD}}$ is the angular dispersion of the cross-disperser, $\Delta\lambda_{\mathrm{FSR}}$ is the free spectral range of the order, σ is the groove spacing of the échelle, λ_B is the blaze wavelength and θ_B the blaze angle. Figure 4.9 shows a schematic of échelle cross-dispersion in which orders are separated by Δy in the λ_{XD} direction. Also shown is the order tilt, Ψ which is given by

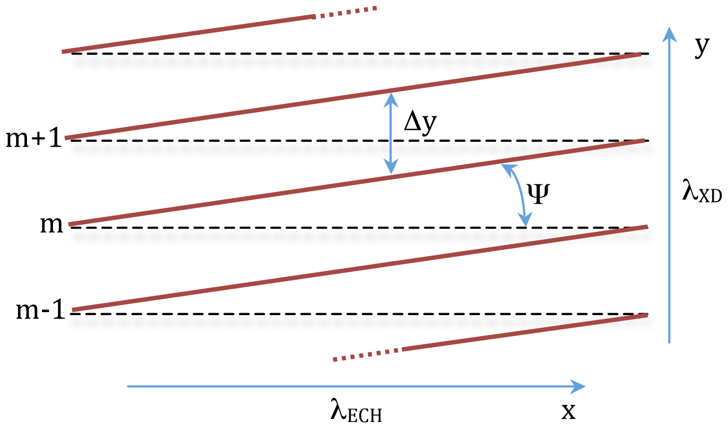

Figure 4.9: A schematic demonstrating the order separation Δy between orders m and $m \pm 1$ caused by cross-dispersion. Also shown is the order tilt Ψ as explained by equation 4.9.

$$\tan\Psi = \frac{d\beta/d\lambda_{\mathrm{XD}}}{d\beta/d\lambda_{\mathrm{ECH}}} \tag{4.9}$$

where $d\beta/d\lambda_{\mathrm{ECH}}$ is the échelle angular dispersion. Since $d\beta/d\lambda_{\mathrm{ECH}}$ is not uniform throughout the wavelength range of an order, the line tilt is also not constant throughout the order, and so each order is not only globally tilted with respect to the échelle dispersion direction, but also locally slightly curved due to varying Ψ. For this reason it is convenient to specify the order tilt at the centre of each order, which requires that equation 4.9 be re-written in terms of the blaze angle:

$$\tan\Psi_B = \frac{d\beta}{d\lambda_{\mathrm{XD}}}\frac{\lambda\cos\bar{\beta}}{2\sin\theta_B\cos\theta} \tag{4.10}$$

According to equation 4.8, the inter-order separation varies according to the angular dispersion of the cross-dispersing element. If a grating is used, the order number m_{g} is low and the grating is blazed at a shallow angle β_{g} with a high groove density σ_{g}, in order to achieve suitably low dispersion. If $d\beta/d\lambda_{\mathrm{XD}} = m_{\mathrm{g}}/\sigma_{\mathrm{g}}\cos\beta_{\mathrm{g}}$, substituting this into equation 4.8 reveals that

$$\Delta y = \mathrm{Constant} \times \lambda_B^2 \tag{4.11}$$

The relation for prism cross-dispersers is quite different and in this case

$$\Delta y = \mathrm{Constant} \times \frac{1}{\lambda_B} \tag{4.12}$$

The dispersion of a prism depends upon the angle of incidence. ‘Minimum deviation’ illumination is when the incident angle produces a dispersed beam almost parallel with the prism base. Although this illumination has least dispersion it has the advantages of minimizing reflection losses at each surface, minimizing the transmission loss by minimizing the path length in the prism, and minimizing the size (and hence cost) of the prism required (Schroeder, 1989).

From equations 4.11 and 4.12, it is clear that order separation increases rapidly with wavelength for a grating and decreases more slowly with wavelength for a prism. Given these conditions, visualized in Figure 4.10, a prism enables a larger wavelength range to be captured on a single

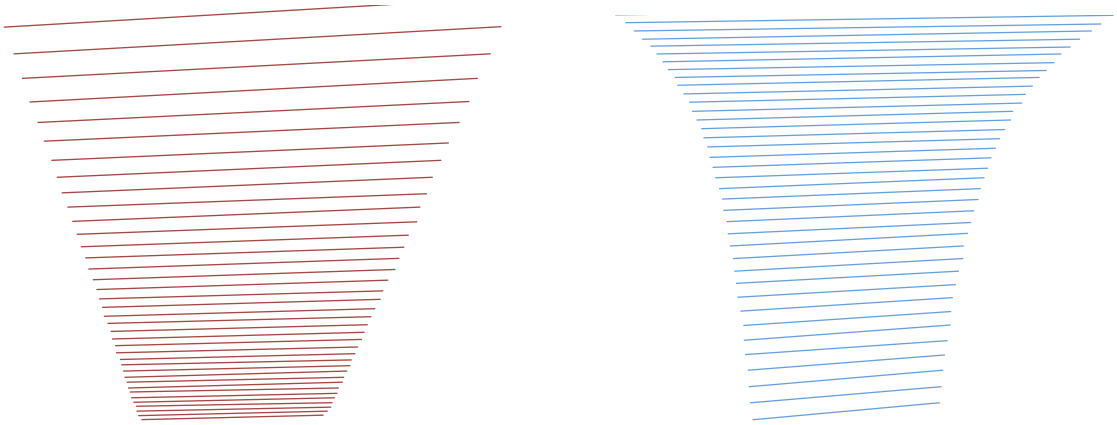

(a) Grating

(b) Prism

Figure 4.10: Grating vs. prism cross dispersion (adapted from Schroeder, 1989).

detector area and for a given cross-dispersion the prism shows considerably more uniform order spacing than a grating does. Some échelle spectrographs with grating cross-dispersers may therefore use multiple gratings to optimize the inter-order spacing for a particular wavelength range. For example UVES makes use of two pairs of different cross-dispersion gratings, with two available in each of its red and blue channels (Dekker et al., 2000). Due to the effects of the blaze function, a grating also shows a greater efficiency variation with wavelength compared with that of a prism, as shown in Figure 4.11.

It is also worth noting that the échelle grating has an effect on the inter-order spacing, as well as the dispersion of the cross-dispersing elements. From equation 4.8 is follows that the inter-order spacing is sensitive not only to the angular dispersion of the cross-dispersers, but als[illegible] the free spectral range of the échelle grating. If the wavelength extent from one order to the next is increased whilst cross-dispersion remains constant, the inter-order spacing will increase. Free spectral range is dependent on several influences, but is most sensitive by far to changes in the groove spacing σ. A more coarsely ruled grating (larger σ) has a larger free spectral range, λ_{FSR}, and therefore the number of orders over a given wavelength range will decrease. Figure 4.12 reveals how tuning the échelle ruling alters the inter-order spacing, but also how the angular width of the orders is increased as a result.

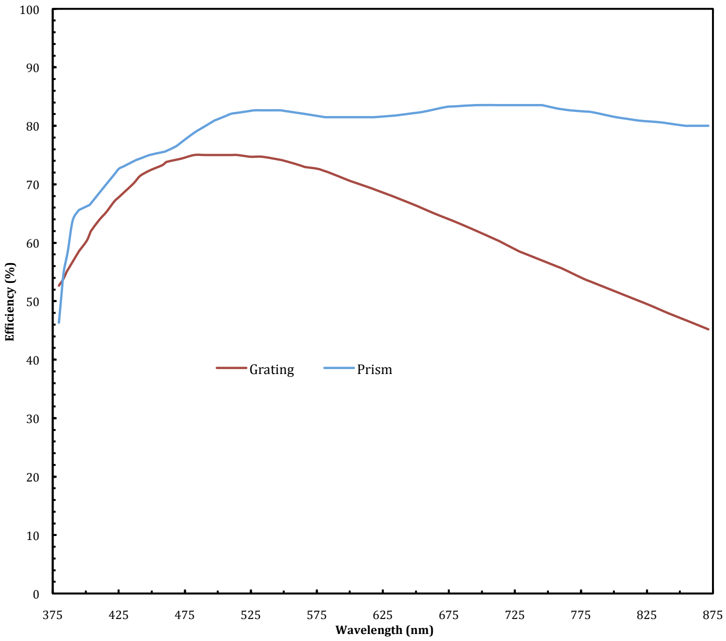

Figure 4.11: The efficiency of two BK7 prisms with AR coatings in double pass compared to the efficiency of a surface relief grating blazed at $\lambda = 550$nm. The prism shows similar losses to the grating in the extreme blue due to transmission and reflection losses, but performs 10-30% better across the rest of the wavelength range.

4.2.4 The VPHGs for SALT HRS

The SALT HRS design features two separate VPHGs for cross-dispersion of the échelle orders, one operating in each arm downstream of the respective transfer mirrors and before the cameras and detectors. Table 4.3 shows the technical specifications of the gratings. Upon inspection of the gratings manufactured under sub-contract by Wasatch, a distorted fringe pattern was clearly visible (see Figure 4.13).

As mentioned in Section 4.2.1, VPHGs are individually made rather than being reproduced replicas of a master grating. This provides opportunity to modify the design parameters of the grating to customize their performance for SALT HRS cross-dispersion purposes. In the original design, the blue VPHG was optimized for peak performance in the centre of the blue channel wavelength range (450nm). Since there are considerable losses in the instrument shortward of 400nm (primarily due to attenuation in the fibre feed, but also due to reflection losses on the mirrors and reflection and absorption losses in the input optics and camera lenses), shifting the

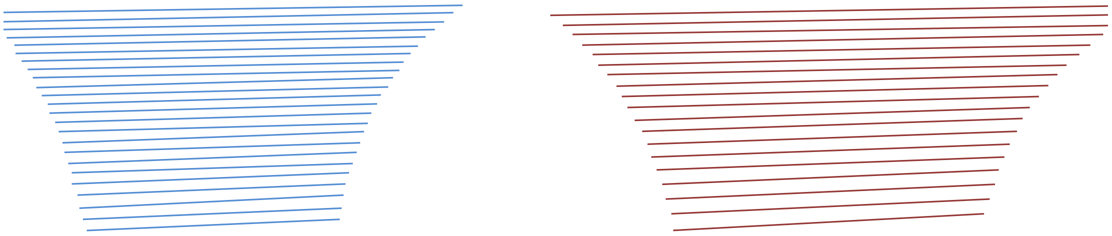

(a) Échelle grating with 100l/mm (b) Échelle grating with 75l/mm

Figure 4.12: The effect of increasing the groove spacing of the échelle grating whilst maintaining the cross-dispersion and camera focal length. With a more coarsely ruled grating the inter-order spacing increases, as does the angular width of each order.

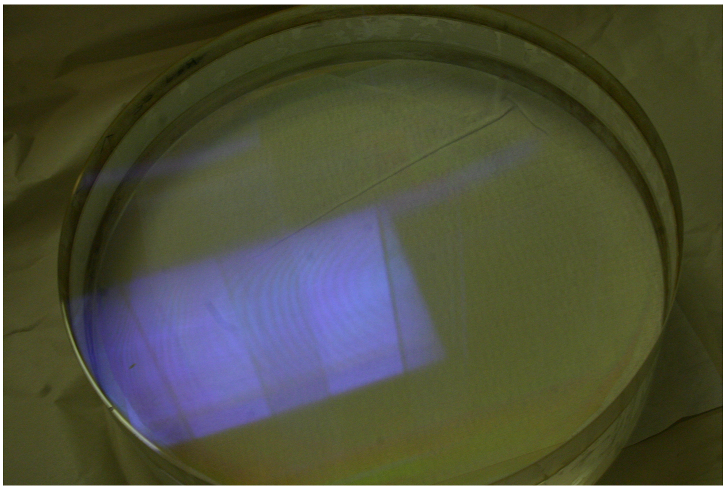

Figure 4.13: Visible distorted fringe patterns on the Wasatch VPHGs.

peak of the blue VPHG to 390nm would ostensibly provide compensation for these losses and give a flatter overall instrument efficiency. First order efficiency predictions made by Wasatch (using rigorous coupled-wave analysis techniques accounting for scatter, absorption, loss to higher orders and Fresnel reflections) show that shifting the peak from 450nm to 390nm (which increases performance by ∼10% at this wavelength) results in a reduction in maximum efficiency from ∼91% to ∼88% and reduces throughput by ∼7% at 530-540nm. The changes to the specification would also require a change in the angle of illumination of the VPHG which may have proven difficult given its proximity to the camera. In consultation with the SALT Science Working Group it was therefore decided that given the peak throughput penalty and losses elsewhere, the blue VPHG should remain optimized for 450nm performance.

Another issue for cross-disperser tuning relates to the crossover performance of the dichroic, which has a coating with a 20nm wide transition region either side of the nominal 50% transmission wavelength (located at 555nm) where some fraction of the light is reflected and the remainder transmitted. Light at these transition wavelengths is potentially transmitted to the focal plane of both detectors but if these transition orders are to be fully captured it results in a reduction of the total spectral bandwidth. The SALT HRS design permits changing the spectral range on each detector by tilting the VPHGs. It was therefore considered whether the instrument should be set up to have complete wavelength coverage (370-890nm assuming a 'perfect' dichroic) with reduced efficiency at 555nm, or sacrifice the extreme red orders for uniform efficiency. Simulation of the spectral formats for the two cases suggested that the science case was more strongly in favour of accepting a dip in throughput at the crossover region in exchange for capturing red wavelengths out to 890nm. This arrangement ensures the continuum red of the near-infra-red Ca II triplet is well measured (Bramall et al., 2010).

It would appear from Section 4.2.3 as though a prism would be the better choice for the SALT HRS cross-dispersing element, with its broader wavelength coverage and more slowly varying inter-order spacing. However, the prisms for SALT HRS would be unfeasible in size. Indeed, the original R2 échelle design for the instrument (which was rejected at its preliminary design review stage in part for this reason) would have required the largest prisms ever manufactured for astronomical purposes. Given the risk and cost for such components, as well as the potential for long lead times, a prism is therefore impractical. The fact that prisms have good transmission is also somewhat lost when considering long path lengths and the risk of inhomogeneities inherent to producing large pieces of glass. Since SALT HRS is divided into both a red and blue channel, the provision of a cross-dispersing element for each channel means that VPHGs, even with their lower wavelength range of operation compared to prisms, are perfectly acceptable. Finally, the $d\beta/d\lambda_{\mathrm{XD}} \propto \lambda^2$ relation for gratings is actually of benefit for the SALT HRS spectral format, since a large inter-order separation (minimum of 11 arcsec for full wavelength coverage) is required to allow for the nod and shuffle functionality of the low resolution mode and for image-sliced higher resolution formats.

Table 4.3: Specifications and performance requirements for the cross-dispersers

	Blue Grating	Red Grating
Type	1850l/mm VPH grating	855l/mm VPH grating
Clear Diameter	154mm	150mm
Angle of Incidence	$24.6 \pm 0.3°$	$17.5 \pm 0.1°$
Central Wavelength	450.0 ± 20nm	705.0 ± 20nm
Wavelength Range	$370 \leq \lambda \leq 555$nm	$555 \leq \lambda \leq 890$nm
Material	Fused silica	Fused silica
Size	165mm physical diameter	165mm physical diameter
Coating	None	None[2]

At the time of the SALT HRS design phase, VPHGs had not yet been used for cross-dispersion elements in high resolution échelle spectrographs. However, both SOAR (Castilho et al., 2004) and PEPSI for the Large Binocular Telescope (Strassmeier et al., 2007) are designed with VPH gratings in mind. Due to their high peak efficiency, broad wavelength coverage (compared to similar surface relief gratings) and tuneable nature, it is expected that they will become a standard for next generation instruments. Indeed ESPRESSO, a larger 'HARPS-like' high-stability instrument for the VLT (and also a pathfinder instrument for CODEX on the proposed E-ELT) will be pioneering the use of VPHGs with slanted grating profiles (Pasquini et al., 2009). This is not strictly necessary on ESPRESSO but will certainly be required on CODEX and is therefore useful to test.

As described in Section 1.4.1, the white pupil transfer optics act to de-magnify the beam. This means that the required line density for the VPH gratings is increased to a higher density (the 1.8× demagnification increases the required line density by a factor of 1.8). This higher line density is more easily and precisely manufactured using the laser interferogram technique and therefore improves the efficiency of the gratings. Furthermore, the smaller diameter beam reduces the size and complexity of the cameras. However, since the pupil transfer mirrors are spherical (see Section 4.1.3) and located 20-30mm further from the intermediate focus than their focal length, the beams incident on the VPH gratings are not perfectly collimated. An additional meniscus lens on each side of both gratings is therefore included to counterbalance this effect and provide an additional degree of freedom in the optimization of the instrument optical train

---

[2] Although the VPHG substrates are specified with no AR coating, it is worth noting that the grating is optically cemented on both sides to additional meniscus lenses which do feature an AR coating optimized for the blue and red channels.

(which is optimized as a complete system rather than in sub-sections). The disadvantage is that the cameras are therefore 'pre-aberrated' and cannot be simply tested using collimated light (Barnes et al., 2008).

4.2.5 VPHG Efficiency Testing

The VPHG manufacturer, Wasatch Photonics, provided theoretical diffraction efficiency curves for both gratings, simulated using rigorous coupled wave analysis GSolver[2] software. In addition to this, the actual measured diffraction efficiency for each grating also provided. However, the measurement contained only two data points per grating at specific laser wavelengths: 405nm and 532nm (blue grating); and 532nm and 632.8nm (red grating). A more definitive test of the grating efficiency, especially in the blue, was therefore required. Additionally, given the concerns about the visible distorted fringe patterns on the delivered VPHGs mentioned in Section 4.2.4, in-house tests were performed to confirm the quoted diffraction effiencies and to ensure that the fringing defects cosmetic (probably due to small variations in the DCG thickness).

The experimental setup used was adapted from that developed by Tamura et al. (2006) to measure throughput variation in cryogenic large-format VPH gratings. Figure 4.14 shows a diagram of the setup in which the test grating is illuminated using monochromatic light. A ThorLabs OSL1-EC fibre coupled halogen light source feeds high intensity white light into a motorized Oriel Cornerstone 130 1/8m monochromator. Newport TracQ software is used to select the appropriate grating and desired wavelength and a Fiberguide Industries fibre bundle transfers the monochromated light to a collimating lens. The collimated beam is then incident on the VPHG which is mounted on a rotation stage. The rotation stage can be adjusted to an accuracy of ±0.5° and is used to orient the grating at the appropriate angle of incidence (24.6° and 17.5° for the blue and red gratings respectively). A large rotating ring-shaped stage surrounds the grating, mounted to which is a 50mm SLR camera lens coupled to an ATIK 314e charge coupled device (CCD) detector. The camera lens and CCD are aligned and held on the same mount which is free to rotate on the stage. This setup enables the diffracted beam to be

---

[2] www.gsolver.com

'tracked' as the change in wavelength alters the expected diffraction angle from the grating. The SLR camera lens is used deliberately out of focus, so that the diffracted beam spot is spread over a larger fraction of the CCD area. This increases the number of pixels (as well exposure time), without saturating the detector to improve the statistics of the measurements.

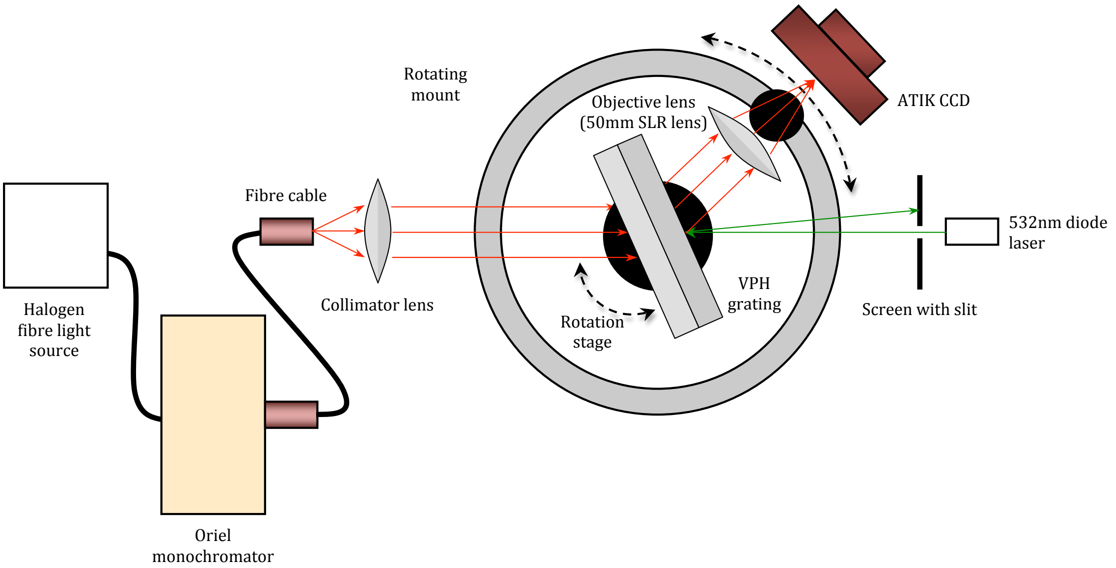

Figure 4.14: The experimental optical setup for VPHG diffraction efficiency testing as adapted from Tamura et al. (2006).

Initially, a series of reference frames and dark frames are exposed at 20nm wavelength intervals across the range of operation for each grating (i.e. ten measurements between 370-550nm for the blue grating and eighteen measurements between 550-890nm for the red grating). With the grating removed, the camera and CCD detector are aligned with the collimated beam on the rotating ring stage. An image spot is exposed at each wavelength, immediately followed by a dark frame of the same duration. The exposure time for each frame is optimized so that the brightest pixels are well illuminated to around 60000ADU, also ensuring that no pixels are saturated (which may lead to a loss of measured light in the reference frame). The exposures therefore cover a wide range between 95ms and 15000ms and depend on the output of the white light source at a particular wavelength, the diffraction efficiency of the monochromator and the efficiency of the CCD.

The grating is then added to the optical setup and is aligned on its rotation stage using a 532nm diode laser. The laser projects a vertically aligned line through a paper screen with a slit cut into it. The rotation stage is first set to 0° and clamped in position. The VPHG is rotated on top of the stage until the reflected laser line is coincident with the incident beam. This calibrates the rotation angle of the VPHG. The laser can then be switched off and the grating may be rotated on the stage to the test angle.

Measurements of the diffracted beams are also made at each wavelength, with an exposure time identical to the corresponding reference frame. Again a dark frame follows each exposure. For both the calibration and grating measurement dark frames, the camera shutter is open. The monochromated beam is blocked at the exit of the fibre guide so that the 'dark' frame records not only the dark current in the detector but also the stray light conditions at the time. The CCD is Peltier cooled to 0°C to reduce dark current and ensure consistency between measurements. Diffraction efficiency may then be calculated using equation 4.13, as a simple ratio of the total light recorded in the measurement frame to that in the calibration frame. A simple IDL program is used to read the .fits files, subtract the appropriate dark frames and calculate the efficiency ratio by summing all the counts in the measurement and calibration frames and then dividing these values.

$$\eta(\lambda) = \frac{\text{Measurement frame}(\lambda, t_{\text{exp}}) \text{ - Measurement dark frame}(\lambda, t_{\text{exp}})}{\text{Calibration frame}(\lambda, t_{\text{exp}}) \text{ - Calibration dark frame}(\lambda, t_{\text{exp}})} \tag{4.13}$$

4.2.6 VPHG Efficiency Results

The experimental procedure outlined in the previous section was performed five times, with a re-alignment of optics and rotation stages and a new set of reference and dark frames taken for each set of repeats. Figure 4.15 shows the manufacturer efficiency test results on the same axes as the averaged results measured at CfAI. The blue lines correspond to the blue grating results with red lines relating to the red grating. The error bars are calculated as the standard error, taking the standard deviation of all five measured values at a given wavelength and dividing by $\sqrt{5}$. The lighter coloured curves without data points represent the Wasatch predicted efficiency based on their GSolver results, whilst the standalone data points represent the average of an S and P polarization state laser measurement performed by the manufacturer prior to shipping. The red grating performs within ~4% of the theoretical efficiency at the majority of wavelengths. The blue grating performs more poorly compared to expectations, achieving 6−8% lower performance at most wavelengths and a considerably worse efficiency at 370nm, some ~14% below predictions. This gives more cause for concern since most of the components in the optical train suffer non-trivial performance losses in the extreme blue. However, the limited wavelength measurements from Wasatch appear to also fall on the 'real-world' performance curve, suggesting that the GSolver result may overstate the diffraction efficiency somewhat. The three 'kinks' in the curves at 510nm, 810nm and 850nm may potentially be explained by the increased scatter in results as indicated by the noticeably larger error bars at these points. Slight losses of the order of 1-2%

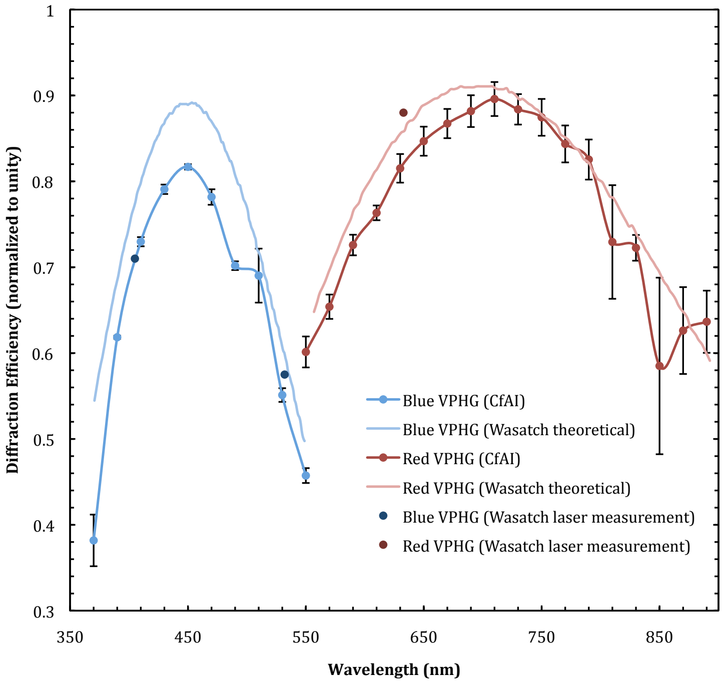

Figure 4.15: Diffraction efficiency results for the SALT HRS volume phase holographic gratings as a function of wavelength. The theoretical efficiency is higher than the actual performance as measured at CfAI and by the manufacturers.

in the measured values compared to theory may be attributable to illumination of the gratings at a slightly different angle during testing. The blue grating is tuned for peak performance at 25.6° and the the red grating at 18.2° angles of incidence. Since the rotation stage on which the grating is mounted is marked in intervals of 2°, the angle of incidence can be set to the nearest 0.5° at best, and more likely only to the nearest 1°. Since the diffraction efficiency falls off reasonably sharply on either side of the peak Bragg angle (see Figure 4.8*a*) this could cause some small losses. However, given the repeated setups it would be expected that the grating would align reasonably closely about the ideal angle of incidence, and the small scale losses this issue may account for are by no means large enough to explain the more substantial losses of the blue grating below 400nm.

There is anecdotal evidence within the astronomical instrumentation community to suggest that real-world performance of VPHGs is not quite at the same level as theoretical models would predict. For example in the case of the HERMES spectrograph (Claire Poppett, private communication), VPHGs were found to exhibit a significant variation in performance across the aperture (although these gratings are considerably larger than those used in SALT HRS and may have been manufactured as a mosaic). Another example can be found in Barden et al. (2000) where the diffraction efficiency of three KOSI[3] manufactured gratings were measured in a variety of orders at different angles and wavelengths. In the case of the 1200l/mm VPHG (the one most similar to those used in SALT HRS) the measured performance is 2-8% lower than that predicted at most wavelengths and again particularly in the blue (see Figure 4.16). Adams et al. (2008) also report that their VPHG efficiencies 'are falling below RCWA predictions' although do not specify by how much. The HETDEX instrument for which these VPHGs are being prepared uses three gratings of the same dimensions and modulation profile and line density. Two were provided by Wasatch (one as a master and one as a copied grating) with the third manufactured by KOSI. There appear to be discrepancies between the three gratings which exceed what might be expected from alignment errors. The Wastach master grating achieves the same peak performance as the KOSI grating, but performance falls off more rapidly from the blaze peak, with the most extreme loss in the blue at 340nm. The Wastach copy also performs more poorly than the master grating, with a 10% discrepancy at the peak 450nm wavelength. Given the experience of others it is therefore possible that the SALT HRS VPH gratings are performing as well as can be expected. Based on the variability of the three 'identical' gratings for HETDEX, it is clear that error in the manufacture of gratings clearly exists beyond that which is accounted for in RCWA

---

[3] www.kosi.com

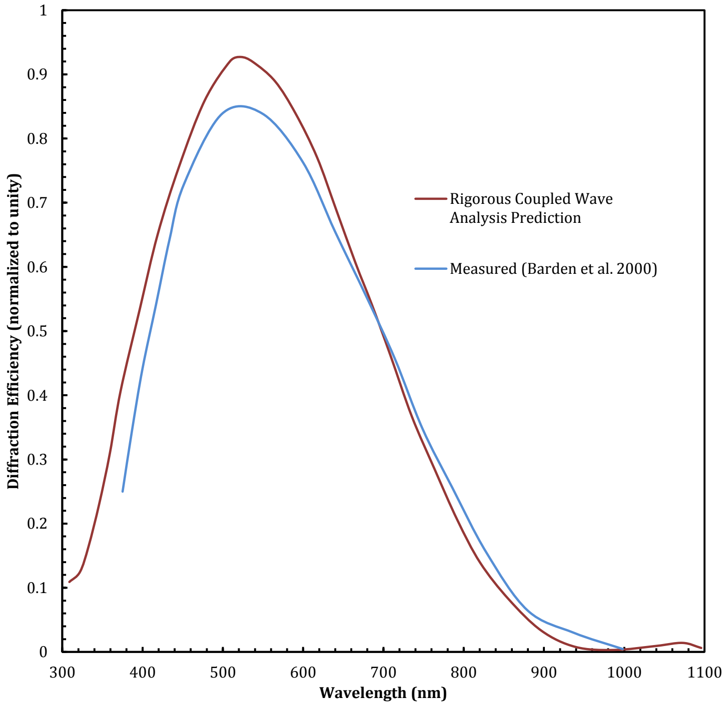

Figure 4.16: A noticeable difference in the measured versus rigorous coupled-wave analysis (RCWA) predicted diffraction efficiency for a 1200l/mm VPH grating as measured in Barden et al. (2000).

predictions. Certainly the SALT HRS red grating should be considered within specification and the blue grating tends towards achieving specification with the exception of wavelengths blue of 400nm. The visible fringe patterns do not seem to be causing any bizarre effects and are likely a result of the grating manufacturing process rather than a problem with the cementing of the meniscus lenses.

Chapter 5

Detectors

5.1 The SALT HRS Detectors

SALT HRS splits the light into two separate arms at a crossover wavelength of 555nm, to increase the efficiency of cross-dispersion and photon detection and also to avoid the need for CCD detector mosaics. For this reason, each arm has its own individual charge coupled device (CCD) detector, optimized for performance over the respective blue and red wavelength ranges. This also has a secondary cost saving benefit in that the blue arm detector can be half the physical size of the red arm detector since the free spectral range, $\Delta\lambda_{\rm FSR}$, of each order is shorter (see equation 1.12 and Figures 1.15 and 1.16 for the spectral formats on the detectors). Both detectors are manufactured by e2V Ltd. and configured in a cryostat assembly with accompanying signal processing electronics by Astronomical Research Cameras Inc. The red detector is a 4k × 4k pixel E2V231-84 deep depletion device with fringe suppression, whereas the blue detector is a 4k × 2k pixel E2V44-82 device. Table 5.1 outlines the key characteristics of the red and blue detectors as provided by the manufacturer.

Table 5.1: e2V Red and Blue CCD Detector System Properties

Property	Red Detector	Blue Detector
Detector	E2V231-84	E2V44-82
Number of Pixels	409[illegible]112	4096(+6[illegible]048
Pixel Size	1[illegible]µm	15µm
Image Area	61.4m[illegible]1.4mm	61.4m[illegible]0.7mm
Flatness (peak-valley)	$<$ [illegible]0µm	$<$ [illegible]0µm
Amplifier Sensitivity	$7\mu V/e^-$	$6\mu V/e^-$
Readout Noise	$< 5e^-$ at 1MHz, $2e^-$ at 50kHz	$< 2.5e^-$ at 20kHz
Maximum Data Rate	3MHz	1MHz
Full Well	$350000e^-$	$200000e^-$
Dark Signal	$< 3e^-$/pixel/hour (at $-100°$ C)	$< 1e^-$/pixel/hour (at $-123°$ C)

5.2 CCD Structure and Design

Both the red and blue detectors for SALT HRS are of the thinned, backside illuminated wafer variety, quite the opposite of the thicker, front illuminated wafers found in most consumer electronics products (Gregory et al., 1999). Figure 5.1 demonstrates the difference between typical consumer front illuminated wafers, and the astronomical standard backside thinned wafers.

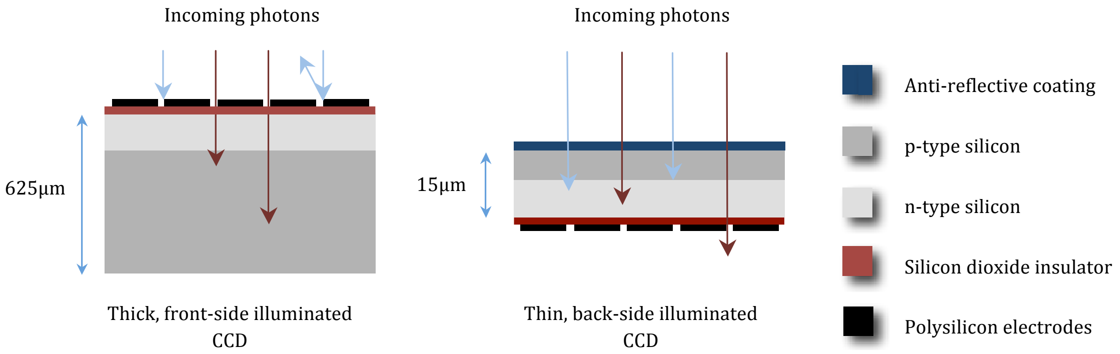

Figure 5.1: The different structures of front and backside illuminated CCDs.

The thicker front-side illuminated CCDs have lower quantum efficiency, since the electrodes cause reflection and absorption of incoming photons and also prevent the addition of an anti-reflective (AR) coating which could improve performance. In contrast, backside illumination leaves a structure compatible with AR coating and removes the issue of electrode transmission losses. However, the specialist procedure to manufacture such CCDs is non-standard and requires chemical etching and polishing. This reduces the yield and raises cost significantly. The thick p-type section in the traditional CCDs also has a low electric potential gradient which means the field preventing electron-hole recombination is weak and may lead to loss of charge. Thinned devices remove this region of low field giving a better quantum yield. Unfortunately, whilst the thin structure improves the blue response, it also makes the device more transparent to near infra-red photons.

Whilst AR coating can help to optimize red sensitivity to some extent, the red E2V231-84 CCD used in SALT HRS also features deep depletion. This technology uses a high-resistivity silicon to mimic the blue performance enhancements of a thin wafer, but means the wafer can be made at an intermediate thickness of ∼40µm, giving the opacity required for efficient detection of red photons (Peckerar et al., 1979). The high purity level required for high-resistivity silicon also drives the CCD price upwards.

Figure 5.2 shows the electric potential in the three wafer types. The wafer on the left is a typical front-side illuminated device, where a region of low electric potential gradient is sufficiently weak as to allow photo-electrons to recombine, leading to losses. The central wafer illustrates the thinned device structure, in which the low field area is etched away and a high field throughout the wafer allows for efficient detection of blue photons, but also means the optical depth is small so red photons may pass through undetected. The wafer on the right shows the 'best of both worlds' nature of a deep depletion thinned device. This allows sufficient path length for red photon opacity and removes the low electric field region associated with recombination losses.

The fraction of photons reflected at the interface of two media at different refractive indices (n), is given by:

$$\left[\frac{n_t - n_i}{n_t + n_i}\right]^2 \tag{5.1}$$

where n_t is the refractive index of silicon, and n_i is the refractive index of air. For the case of silicon in air, reflection occurs 32% of the time. Without the AR coating, two in three photons can be counted at best. The AR coating is a transparent dielectric chosen to be in the middle of the refractive indices of silicon and air, and must have optical thickness of $\lambda/4$. λ is often chosen

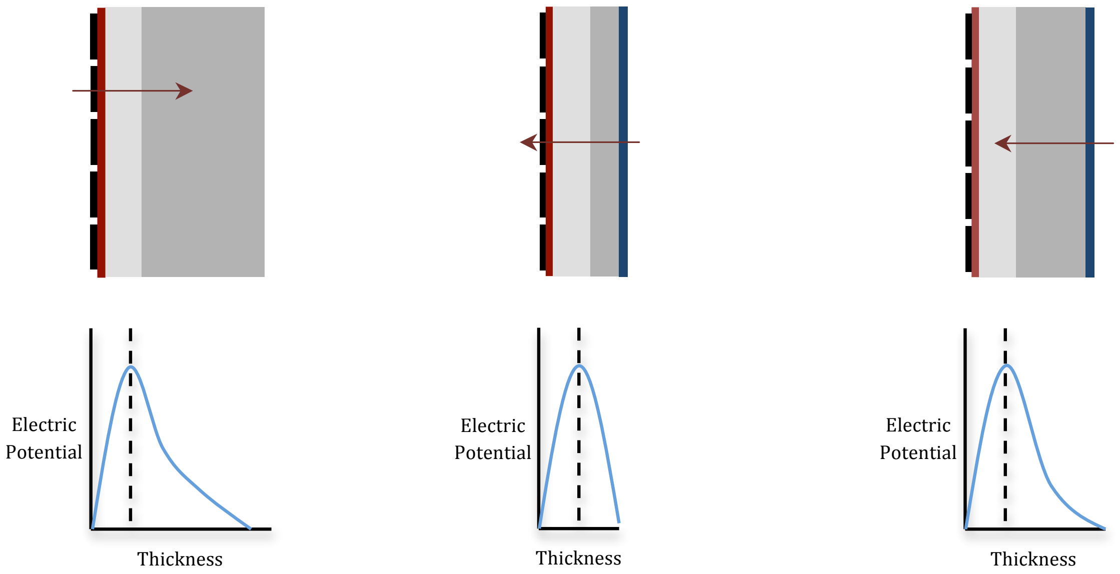

Figure 5.2: The electric field as a function of wafer depth in three CCDs. On the left is a typical thick detector, the central diagram shows a thinned detector and the diagram on the right details a deep depletion device structure (see Figure 5.1 for the CCD structure key).

to be at the centre of the visible wavelength range for a broadband coating (∼550nm). There are now three refractive indices to consider:

$$\left[\frac{n_t \times n_i - n_s^2}{n_t \times n_i + n_s^2}\right]^2 \tag{5.2}$$

where n_s is the refractive index of the coating. It is clear from equation 5.2 that should $n_s^2 = n_t$, reflectivity falls to zero. In the case of silicon and air, this means the AR coating needs to have $n = 1.9$. Hafnium dioxide fills this criteria and is used in the AR coating of most astronomical CCDs. Given the dependence of optimal optical thickness on wavelength for the coating, tuning this thickness allows for changes in the peak coating efficiency, in addition to the use of different dielectric materials.

5.3 Blue Detector Cleaning

5.3.1 Motivation

In order to reduce dark current in the CCD (thermally generated electrons in the silicon wafer), both detectors in SALT HRS are cooled to $T{\sim}170$K. To achieve this, each detector is housed inside a vacuum-pumped cryostat. A window on the front of the cryostat allows light in to the detector, which is mounted a few centimetres behind it. In the final detector system the cryostat window will be replaced by the final optical element of the cameras, but the cryostat was fitted with a flat window for test purposes. To the rear of the CCD is the coldhead, which is maintained at low temperature with refrigerant gas pumped in from an external compressor. The coldhead is the lowest temperature body inside the sealed cryostat, which is kept under vacuum conditions both to reduce contaminants inside but also to eliminate conduction to facilitate the low temperatures required. However, the cryostat is not a perfect vacuum, and as the coldest region, the coldhead becomes the site where any remaining gases condense. The coldhead is capable of reaching temperatures around 110K, whilst the detector achieves peak performance at roughly 173K. For this reason, the detector is warmed relative to the surrounding cryostat with heat provided by a small resistor attached to the metal CCD mounting plate. This also allows for a more accurate control of the detector temperature, which must be stabilized to $\pm0.1^{\circ}$.

During contractor testing for the blue detector and optimization of the electronics, the CCD surface became contaminated. This occured because when warming the detector back to room temperature, the compressor was turned off allowing the coldhead to warm up. As the coldhead warmed more quickly than the detector, condensed gases boiled off its surface and re-condensed onto the detector, which was now cooler than the coldhead. Vacuum grease and other contaminants therefore bonded to the CCD surface. In order to prevent this occurring again, a warm-up procedure was developed in which the detector is set to a temperature slightly above ambient and is heated using the resistor, and the cryostat is vacuum pumped throughout the warm-up cycle to remove potential contaminants. This ensures the detector is always warmer than the coldhead and any contaminants are immediately removed from the cryostat.

5.3.2 Flat Illumination & Cleaning Results

The blue detector was cleaned by the contractor using a solvent although concerns remained that the detector still showed some residual contamination. Once the blue detector was delivered, a flat-field setup was used to uniformly illuminate the detector to quantify the losses associated with the soiled regions. Figure 5.3 shows the experimental setup for illumination of the detector. This same setup was used for determination of detector characteristics using the photon transfer curve described later in Section 5.4, and an adaptation of the setup is also used in fringing tests mentioned in Section 5.6. A ThorLabs model OSL1-EC halogen fibre light source illuminates a Linos 75µm diameter pinhole. Any structure from the fibre is smeared using a neutral density filter, followed by a filter wheel fitted with a ThorLabs FB440-10 filter with 440nm central wavelength and 10nm full width half maximum (FWHM) bandpass. The beam is expanded and collimated by a Vixen ED102SS 10cm diameter apochromatic refractor telescope and overfills the 6cm × 3cm area of the detector. Use of a narrow bandwidth and an apochromatic telescope ensures that chromatic aberration in the beam projected onto the detector is minimal. Overfilling of the detector and the wide telescope field of view should also help to reduce any vignetting and provide a suitably uniform illumination. The setup is mounted on a vibration damped Newport optical bench in a dark room with stray light sources reduced where possible. Cladding on the optical bench sides and opaque fabric shrouding the telescope from the dew cup to the detector entrance window also reduces stray light.

The same set of exposures were made twice, with the detector rotated by 90° between two otherwise identical frames. The central 2k × 2k region (since the 4k × 2k detector will be landscape in one frame and portrait in another) of the detector should therefore be illuminated in both frames with the same beam profile and can then be compared for uniformity across two orientations. Figure 5.4 shows the uniformity of the beam across the detector. A 6 second exposure with single channel read out at 400kHz (pixels per second) and −110°C detector temperature is made at both orientations, with an accompanying 6 second dark frame.

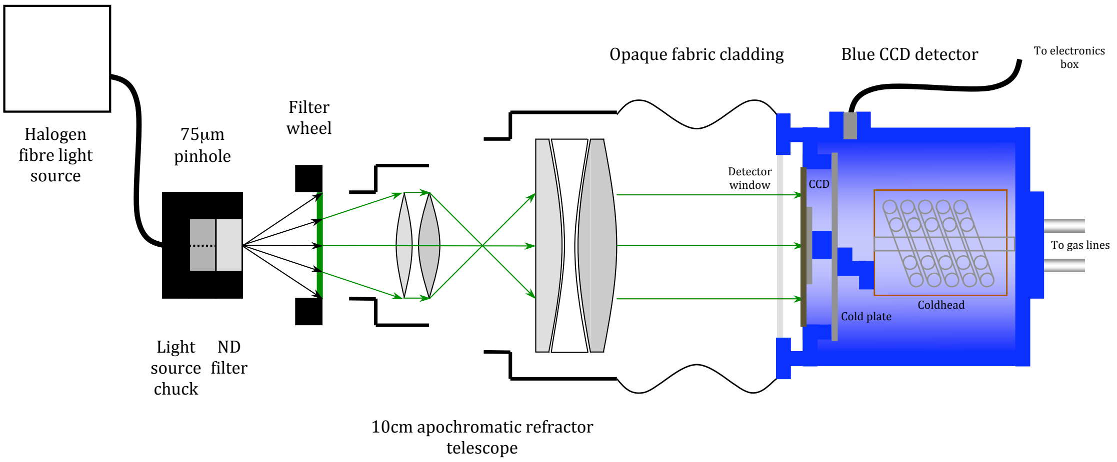

Figure 5.3: The optical setup for 440nm flat-field illumination of the blue detector.

Figure 5.5 shows a corner of the detector as illuminated by the flat-field beam. Figure 5.5*a* shows a frame taken prior to cleaning, whilst Figure 5.5*b* shows the same region after two cleans. The losses prior to cleaning were quite marked: areas showing hair lines typically exhibited 2000-3000DN less signal than surrounding pixels ($4-6\%$ of total signal); areas showing extended contaminant deposits typically exhibited 6000-7500DN less signal than surrounding pixels ($12-15\%$ of total signal). The slight ‘mottled’ pattern running diagonally across both images appears to be a small additive clock noise which disappears at other read-out speeds and in other exposures.

Figure 5.6 shows photographs of the detector corner taken under a 50× magnification microscope. The image on the left (Figure 5.6*a*) shows a deposition of vacuum grease and other contaminants in the corner of the detector, left over as a residual from cleaning undertaken by the contractor. Also visible are the edges of evaporated solvent pools left on the chip when a film of impurities is left behind. The image on the right (Figure 5.6*b*) shows the result of two in-house cleans using high-purity alcohol and lens tissue. The only non-uniformities remaining are cosmetic imperfections on the CCD surface itself.

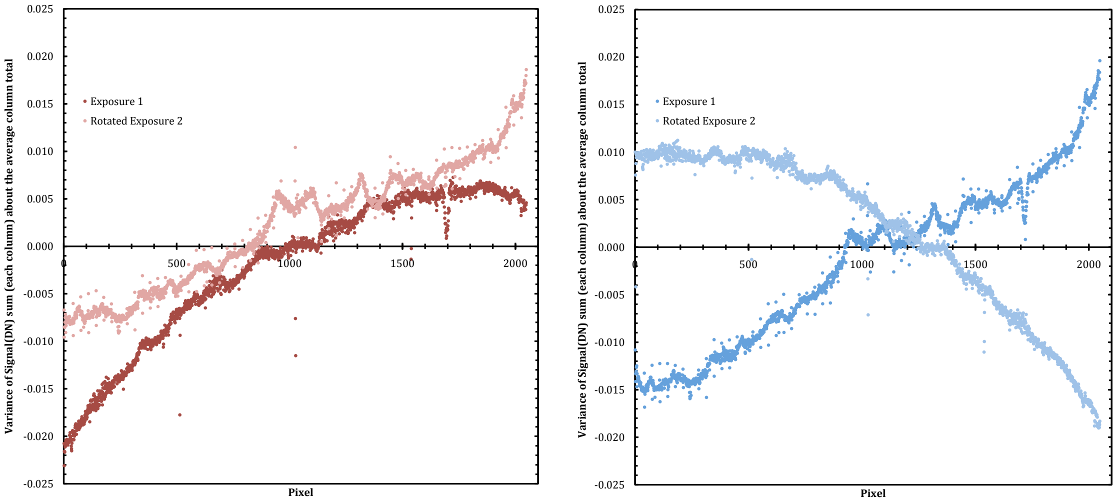

(a) x-direction

(b) y-direction

Figure 5.4: Plots showing good beam uniformity in both x and y directions indicating a suitably ‘flat’ illumination was obtained using the telescope as a collimated beam expander.

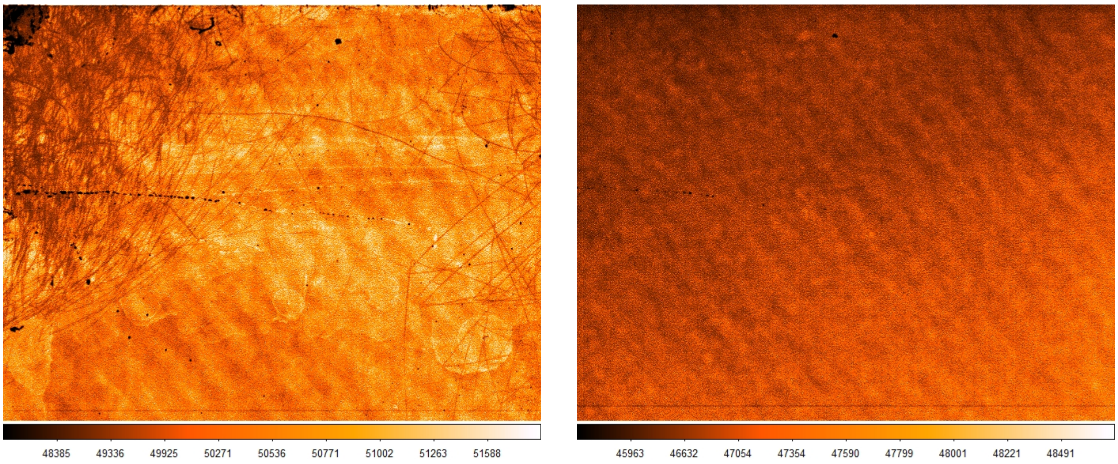

(a) Soiled detector

(b) Cleaned detector

Figure 5.5: Flat-field illuminated images from the soiled blue detector region before and after cleaning.

(a) Soiled detector (b) Cleaned detector

Figure 5.6: The result of successive cleaning of the contaminated blue CCD using high purity alcohol, from left to right, as seen under a micoscope.

5.4 Blue Detector Photon Transfer Curve

5.4.1 Motivation & Camera Gain Constant

Various techniques exist for determining the performance of CCDs, and they are often expressed in a graphical form as a transfer curve. Three of the most common ones include the quantum efficiency transfer (efficiency of the detector in producing photo-electrons as a function of incident light wavelength), x-ray transfer and photon transfer. The latter is highly useful for characterising numerous CCD parameters such as read noise, dark current, quantum yield, full well capacity, pixel linearity and non-uniformity, sensitivity, signal-to-noise, offset and dynamic range. The photon transfer curve is also useful because it includes all components of the camera system including not only the CCD detector itself, but also the clock and bias levels, signal processing and analogue-to-digital (ADC) conversion as well as acquisition software (Janesick, 2001).

The input to the camera is obviously incident photons, with the output given as an encoded digital number, $S(\text{DN})$ representing the accumulated signal of each pixel. Output signal is a function of several factors as illustrated in Figure 5.7, which shows the input to output signal path in a CCD camera system as a block diagram.

If $S(\text{DN})$ is the average signal for a group of pixels, P is the average number of incident photons

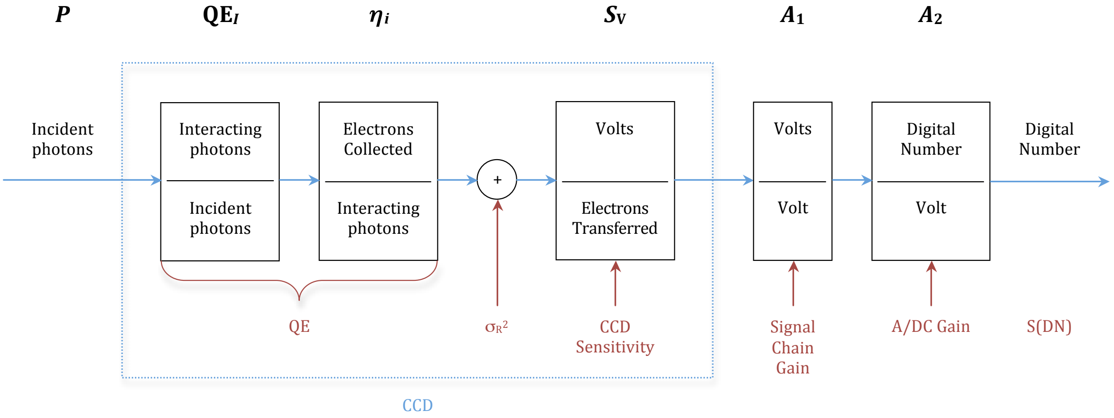

Figure 5.7: CCD Camera Block Diagram (from Janesick, 2001).

per pixel, QE_I is the interacting quantum efficiency (interacting photons per incident photons), η_i is the quantum yield (number of electrons generated, collected and transferred per interacting photon), S_V is the sensitivity of the readout amplifier (Volts/electron), A_{CCD} is the amplifier gain (Volts/Volt), A_1 is the signal processor gain (Volts/Volt) and A_2 is the ADC gain (DN/Volt), then

$$S(\mathrm{DN}) = PQE_I\eta_i S_V A_{CCD}A_1A_2 \tag{5.3}$$

The constants defining the conversion of DN into interacting photons or generated signal electrons are:

$$K = \frac{1}{S_V A_{CCD}A_1A_2} \tag{5.4}$$

and

$$J = \frac{1}{\eta_i S_V A_{CCD}A_1A_2} \tag{5.5}$$

where K and J are gain constants in units of electrons/DN and interacting photons/DN respectively. They can be related via quantum yield, η_i:

$$\eta_i = \frac{K}{J} \tag{5.6}$$

where i is the pixel identifier. The photon transfer function thus becomes a useful tool to determine J and K without independent measurement of QE_I, η_i and S_V. Although high energy photons can create multiple electron-hole pairs, at λ>400nm it can be assumed that only single pairs are generated and therefore quantum yield is unity ($\eta_i = 1$).

$$S(\mathrm{DN}) = \frac{P_I}{K} \tag{5.7}$$

where P_I is the number of interacting photons per pixel ($P_I = (QE_I P)$). Using the propagation of errors formula and assuming K has negligible variance:

$$K = \frac{S(\mathrm{DN})}{\sigma_S^2(\mathrm{DN}) - \sigma_R^2(\mathrm{DN})} \tag{5.8}$$

This relates signal and signal noise in D[illegible]nits of electrons. Another consequence of quantum yield being unity is that J and K are equal for illumination at wavelengths above 400nm.

A photon transfer curve can be generated by illuminating a detector uniformly at increasing levels, and plotting the noise, $\sigma_S(\mathrm{DN})$, as a function of average signal $S(\mathrm{DN})$. Due to the large dynamic range offered by modern CCDs, a logarithmic scale is used. Figure 5.8 shows an example curve, which typically displays three regimes: the read noise floor, σ_R (random noise under dark conditions due to dark current and amplifier readout noise); signal noise governed by shot noise of photon arrival (displaying a gradient of ½ on a logarithmic scale due to the $\sqrt{(N)}$ Poisson statistics nature of the noise); and pixel non-uniformity noise (a result of differences in sensitivity across pixels, which is proportional to signal and therefore carries a slope of unity). Also obvious is the full capacity, after which the noise level rapidly decreases. 'Blooming' of over-saturated pixels into surrounding pixels lowers the noise but signal remains high.

5.4.2 Blue Detector Characterization

The SALT HRS blue detector was vacuum pumped and cooled to its optimal operating temperature of −100 °C (a 'happy medium' temperature - dark current is minimized without reaching a temperature so low that the signal amplifiers lose efficiency or which would require exotic refrigerants). Using the method of flatfield illumination from a collimated beam expanded by a 10cm

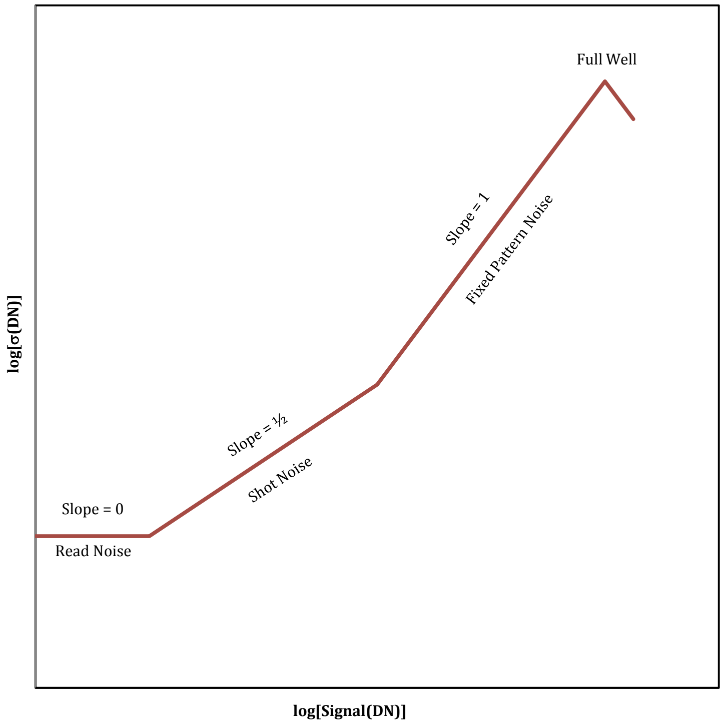

Figure 5.8: Schematic photon transfer curve showing the three distinct noise regimes (read noise, shot noise and fixed pattern noise), plus full well capacity.

telescope - previously discussed in Section 5.3 - the detector was uniformly illuminated with 440nm light (white light passed through a 440nm, narrowband filter with 10nm FWHM). The idling of the detector electronics (a constant read-out mode so that charge does not accumulate) was switched off, and three back-to-back exposures were made at a variety of exposure times. Total integration time is held constant across each pair of exposures at 11 seconds, so each frame acquires a similar dark current. However, the exposure time within this 11 second window is varied logarithmically between 0.015 seconds (the minimum period allowed by the shutter system employed) and 10.985 seconds. Also recorded was a full 11.000 second dark frame. The first frame at each exposure is discarded, and is included only as a means of reading out the full detector before commencing the second and third exposures with the idling disabled. A simple IDL routine reads in each pair of images at each exposure duration as .fits files and subtracts one from the other. The IRAF imstat function is used to calculate the mean pixel value in the exposure frames, and the standard deviation of the pixels in the subtracted frame. Figure 5.9

shows a plot of noise σ_S(DN), versus mean pixel value S(DN), i.e. the photon transfer curve. In Figure 5.9, a photon transfer curve is plotted for not only the whole detector, but also for comparison as a 200 × 200 pixel sub-array in each of the four corners of the detector.

Fitting a line to the 'slope ½' region of the photon transfer curve in Figure 5.9 and tracing it back to a σ_S value of 1DN gives a graphical determination of the camera gain constant, K, as 2.6e$^-$/DN.

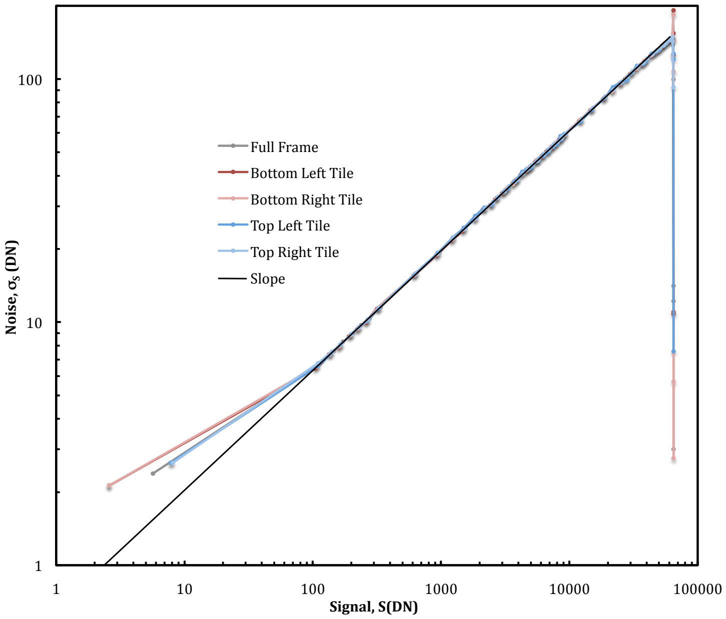

Figure 5.9: A photon transfer function curve for the SALT HRS blue detector, based on 440nm flat-field illumination and 65 different exposure period frames across the full dynamic range of the detector.

However, a more accurate means of estimating K is to make multiple S(DN) and σ_S^2 measurements in many sub-arrays and produce a histogram of camera gain. 10000 separate 50 × 50 pixel sub-arrays were therefore sampled randomly across the CCD, from a frame with a 9.5 second exposure. This frame has a mean signal of approximately 62000DN, which is sufficiently close to full well that read noise becomes a negligible contribution, yet far enough from full well to ensure that few if any pixels are saturated and at risk of 'blooming', thereby reducing the noise. σ_S^2 is determined in the usual way, by subtracting two identically exposed flat-field images, and calculating the variance across each sub-array in the subtracted frame using IRAF scripts. The mean of each sub-array is also calculated, and averaged together, from which S_{OFF} is subtracted to give an accurate estimate of $S(\mathrm{DN})$. S_{OFF} is an electrical offset representing the output in DN in the absence of any signal. In this case a 0.015 second dark frame is used to estimate S_{OFF}. Although such a frame will also contain a dark current signal, this will be a negligible contribution over the short duration of the exposure and S_{OFF} essentially represents the detector bias. S_{OFF} is also calculated on a sub-array by sub-array basis, although it is almost constant. Finally, σ_R is calculated, in the same way as σ_S, using the standard deviation of two subtracted 0.015 second dark frames. K is then calculated for each sub-array according to equation 5.4 with the results shown as a histogram in Figure 5.10. The histogram has a Gaussian distribution with mean $K = 2.90 \pm 0.08\mathrm{e}^-/\mathrm{DN}$.

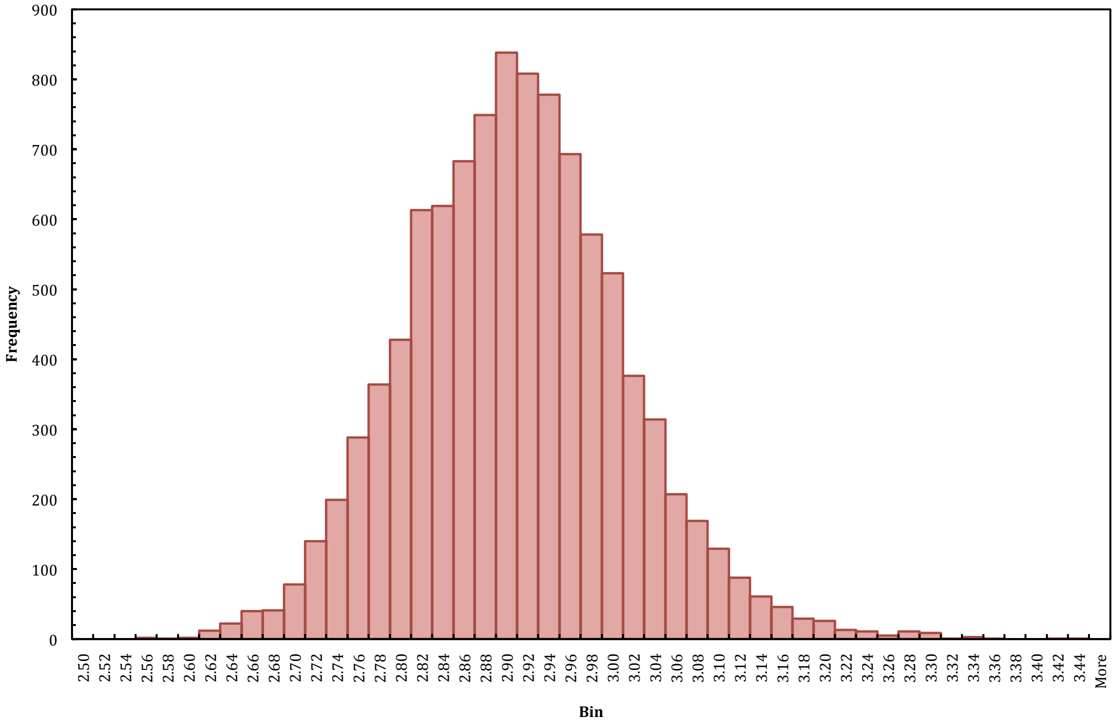

Figure 5.10: Histogram of K (camera gain constant) values, as calculated from 10000 randomly sampled 50×50 pixel sub-arrays in a pair of subtracted, close to full well illumination flat field frames.

The uncertainty in camera gain can be estimated using equation 5.9, assuming that the read noise σ_R is negligible, that photon statistics apply ($K\sigma_S(\mathrm{DN}) = [KS(\mathrm{DN})]^{1/2}$) and that $\sigma_S \ll$

S(DN):

$$\sigma_K = \left[\frac{2}{N_{\text{pixels}}}\right]^{1/2} K \tag{5.9}$$

Substituting the appropriate values of K and $N_{\text{pixels}} = 2500$ (from the 50 × 50 pixel sub-array) into equation 5.9, we find $\sigma_K = 0.08\text{e}^-/\text{DN}$. Full well capacity and read noise define the dynamic range (in decibels) of the CCD using equation 5.10 below:

$$\text{DR} = 20\log\left[\frac{S_{FW}(\text{e}^-)}{\sigma_R(\text{e}^-)}\right] \tag{5.10}$$

Based on the 16-bit architecture of the CCD giving $2^{16} = 65535$ count range, or $S_{FW} = 190054\text{e}^-$, the dynamic range is 98dB, for a read noise of $\sigma_R = 2.40\text{e}^-$. The photon transfer function also allows a measurement of the detector linearity. Negligible non-linearity between change in gain and change in signal is preferable in an ideal detector. Plotting S(DN) as a function of exposure time, t_E reveals a power law of the form

$$S(\text{DN}) = C(t_E)^{\gamma} \tag{5.11}$$

where C is a proportionality constant, t_E is the exposure time in seconds and γ becomes a measure of device linearity, with a value of unity signifying proportionality between output signal and exposure time. Figure 5.11*a* shows that for the blue detector, $\gamma = 0.994$. A linearity residual transfer curve reveals smaller detail according to equation 5.12, where LR is the linearity residual at a specific signal level (%) and S_M(DN) is the signal level at a mid-scale in the plot, made at an associated exposure time of t_{EM} (seconds).

$$\text{LR} = 100\left(1 - \frac{\frac{S_M(\text{DN})}{t_{EM}}}{\frac{S(\text{DN})}{t_E}}\right) \tag{5.12}$$

Plotting LR at each S(DN) and t_E level from the photon transfer curve exposures using S_M(DN) at 5 seconds and t_{EM} as 31086, the graph in Figure 5.11*b* shows that linearity is within ±3%.

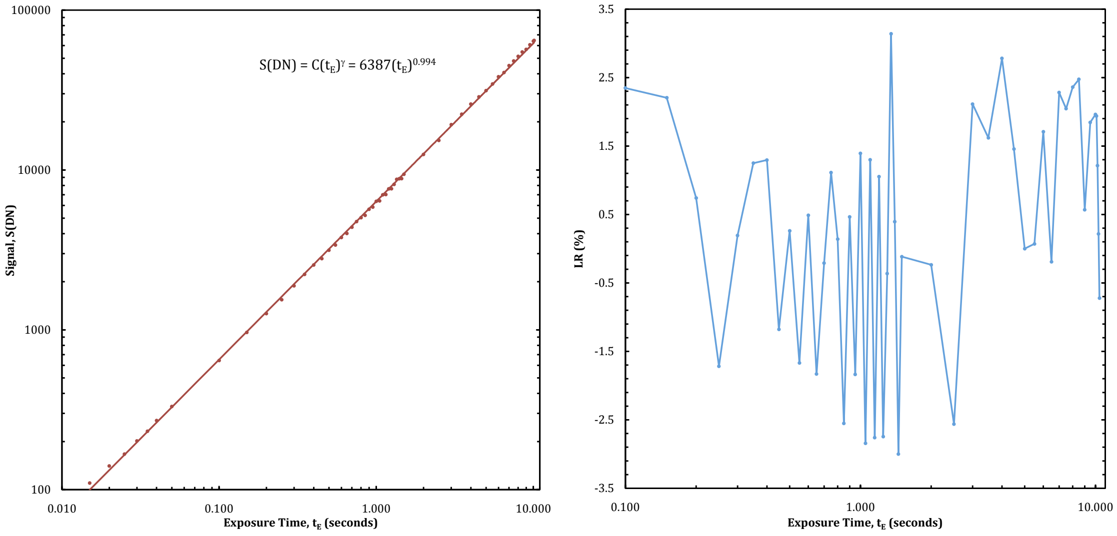

(a) Linearity transfer curve

(b) Residuals curve

Figure 5.11: Linearity transfer curves showing a detector response linearity to better than ±3%.

5.5 Blue Detector Quantum Efficiency

5.5.1 Introduction

The quantum efficiency of a detector is simply the fraction of incident photons which produce an electron-hole pair and are converted into a signal. The closer to 100% this value is, the more sensitive the detector to light. Equation 5.3 shows the dependencies of CCD signal on: P (the average number of incident photons per pixel); QE_I (the interacting quantum efficiency); η_i (the quantum yield - number of electrons generated, collected and transferred per interacting photon); S_V (the sensitivity of the readout amplifier in Volts/electron); A_{CCD} (the amplifier gain); A_1 (the signal processor gain in Volts/Volt); and A_2 (the ADC gain in DN/Volt).

Over the wavelength range of blue detector operation (370-555nm) it is reasonable to assume that quantum yield, η_i is unity, although at much shorter (UV and X-ray) wavelengths multiple electron-hole pairs may be produced. Knowledge of the camera gain constant, K from the photon transfer curve therefore allows calculation of quantum efficiency, QE_I, since it combines all of the above terms (see equation 5.4) given a measurement of the number of incident photons P

and the CCD signal recorded S(DN).

5.5.2 Blue Detector Efficiency Testing

The experimental setup for measuring the SALT HRS blue detector quantum efficiency is shown in Figure 5.12. An Oriel Apex light source feeds an Oriel Cornerstone 1/8m monochromator.

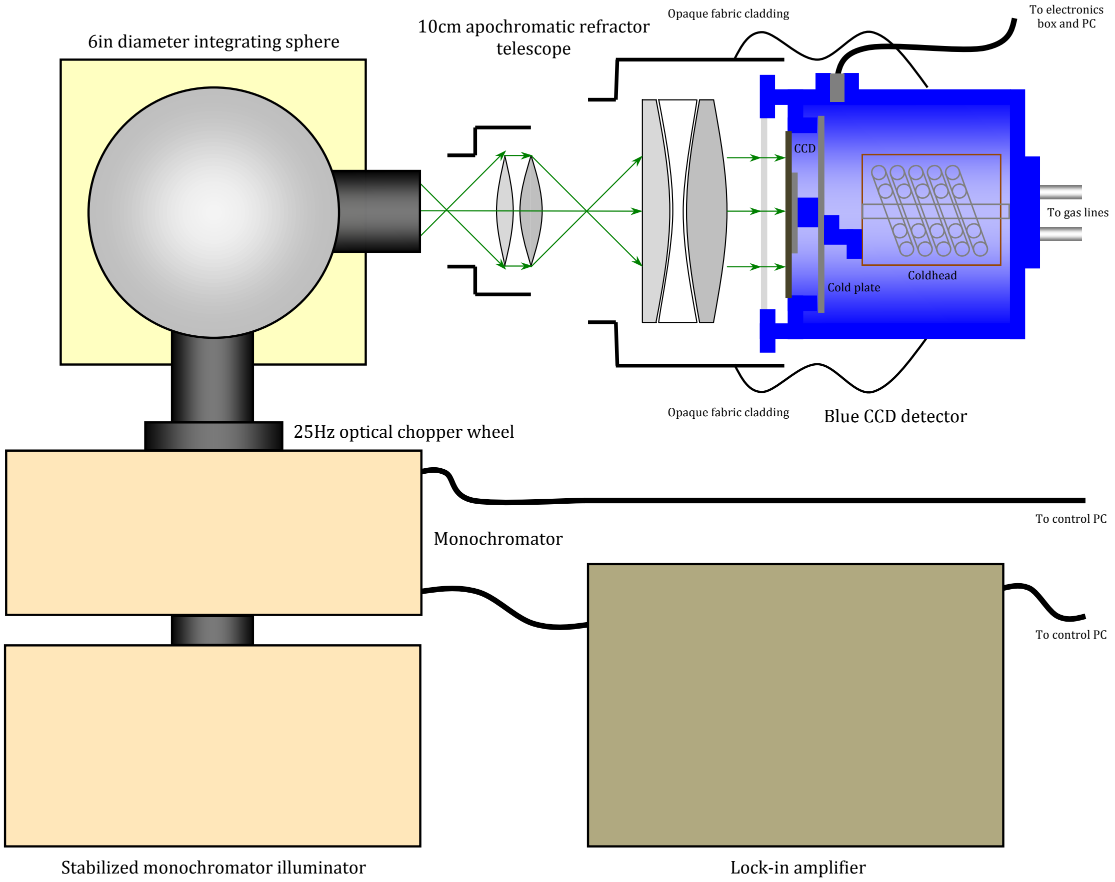

Figure 5.12: The experimental set-up for CCD quantum efficiency measurement.

The light source has a highly stable power supply to ensure consistent light intensity. The quartz tungsten halogen lamp emits white light over the full wavelength range of the blue detector, and is switched on more than 30min before use to ensure proper stabilization of the lamp. The monochromator is software controlled and motorized for wavelength switching and illuminates a Newport 70675 6in diameter PTFE-lined integrating sphere with light at the various test wavelengths. An optical chopper (see Figure 5.13) operating at 25Hz sits between the monochromator and integrating sphere. The chopper consists of a rotating mask with equally spaced slits of equal size, which pass in front of the beam. It is connected to a Merlin lock-in amplifier which identifies the peak signal and allows for automatic background light subtraction (the signal measured when the beam is blocked is due to other sources of light or noise) in the measured signal.

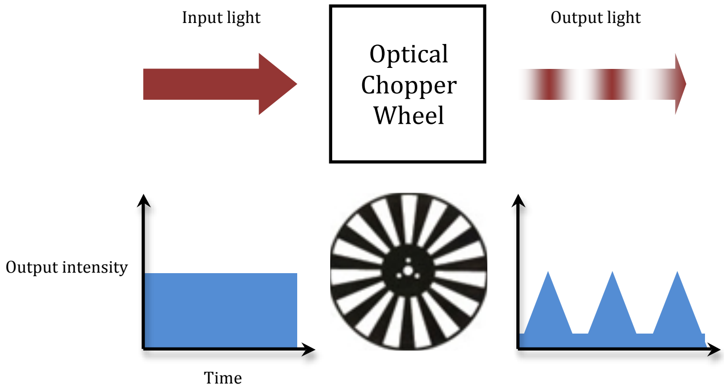

Figure 5.13: The operation of an optical chopper and the resulting signal. By phase-locking the detector to the chopper signal a measurement insensitive to extraneous background light can be made.

The integrating sphere has an interior cavity with high diffuse reflectivity, and acts to remove directional memory of the input beam. Light exiting the sphere is expanded and collimated by a 10cm diameter apochromatic refractor telescope to over-fill the detectors. The port size of an integrating sphere scales with its diameter. The 6in sphere has a 1.5in port diameter, but to test the full area of the detector without multiple tiled measurements would require a significantly larger and unfeasibly expensive integrating sphere, hence the use of the telescope. A NIST calibrated photodiode is placed on the top port of the integrating sphere to measure the actual light output at each test wavelength (370-555nm in 5nm intervals, with 25nm intervals thereafter up to 900nm). The photodiode is then moved to the telescope output and placed in the centre of the beam, inside the dew cup of the telescope and covered in opaque fabric to block stray light. Figure 5.14 shows the beam intensity as measured by the photodiode at the integrating sphere and after the telescope beam expansion (scaled by a factor of 1250). Aside from light loss, there is no appreciable consequence of using the telescope as a beam expanding device, since both intensity plots have the same shape. The 'flatness' of the telescope expanded beam has already been discussed in Section 5.3 and is presented in Figure 5.4.

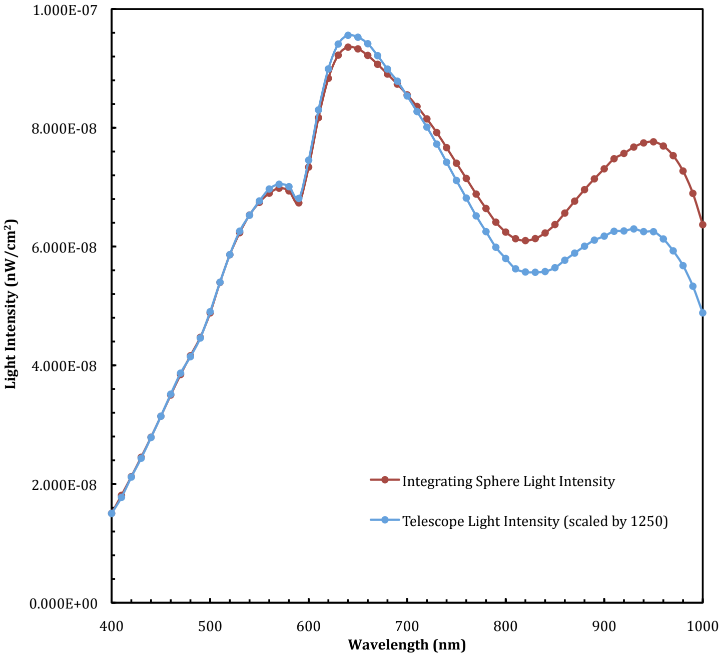

Figure 5.14: A comparison of the light intensity exiting the integrating sphere compared with that at the telescope output (scaled up by a factor of 1250) shows generally good agreement in performance, albeit with significant light loss using the telescope. Agreement is very good over the important blue detector test wavelength range up to 555nm, with the difference in the red and near-IR wavelengths likely due to absorption caused by the telescope optics.

5.5.3 Blue Detector Quantum Efficiency Results

Given the calibrated intensity of light exiting the telescope, measured in nWcm^{-2}, it is simple to calculate the total energy incident on the detectors by multiplying by the detector area ($\sim$18.874cm^2 for the blue CCD) and integrating time. Since the energy of a photon, E is given by

$$E = \frac{hc}{\lambda} \tag{5.13}$$

the energy of an individual photon at each measurement wavelength may be calculated. This leads to a value for the number of photons incident on the CCD. Using equations 5.3 and 5.4, inserting a camera gain constant, K of 2.9e$^-$/DN (from Section 5.4), the quantum efficiency can be determined at each wavelength based on the number of incident photons and the total signal recorded in each frame. The measured quantum efficiency curve for the blue detector can be seen in Figure 5.15 [illegible] blue curve). For comparison the over-plotted [illegible] blue data curve is

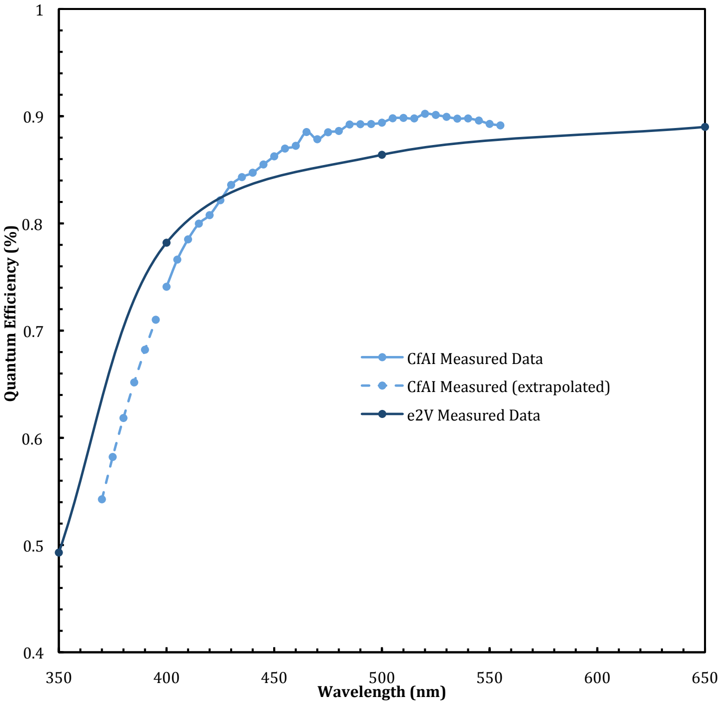

Figure [illegible] The SALT HRS blue detector quantum efficiency curve.

the data from the manufacturer for the actual CCD shipped. It would appear from Figure 5.15 that the blue detector is performing within or above expectations, apart from some slight loss of 2-4% blueward of 400nm. This may not be a property of the detector but rather the fact that the photodiode is not calibrated below this wavelength. Data in the 370-395nm range (indicated with a dashed line) is therefore based on a linear e[illegible]olation of the photodiode responsivity fall-off which may be over-stating the effect. From this result we can also be confident that the blue detector cleaning process (see Section 5.3) did not damage the detector or remove anti-reflective coating.

5.6 Red Detector Fringing

5.6.1 The Causes of Fringing

As a thin material with reflective parallel surfaces, the wafer in thinned CCD detectors can act as a resonant optical étalon. As incident light becomes redder, the wafer becomes increasingly transparent, and light bounces back and forth between the front and back surface (even with an AR coated back surface, since the refractive index of the silicon layer at $n{\sim}4$ is a large change from the silicon dioxide layer at $n{\sim}1.5$). Although this effect acts to increase the length of the light path (increasing the chance of photo-electron generation), it also allows standing waves to form in the waf[illegible]oducing constructive or destructive interference patterns. The result is an interference fringe pattern which depends on small scale variations in the wafer detector thickness (Jorden et al., 2004). The fringe signal can be very difficult to remove in software post-processing as it is not only a function of detector thickness but also incident wavelength, degree of collimation and angle of incidence of illumination.

5.6.2 Fringing in the Red Detector

The red detector in SALT HRS is expected to show fringing effects at far red wavelengths in the red channel as it features a deep depletion detector. As mentioned previously the deep depletion device is thicker than a standard thinned detector, meaning it is more opaque to red photons. However the device features a proprietary fringe suppression option. Several science grade CCD manufacturers offer such a technology although because it is a proprietary process it is difficult to be sure how it is achieved. It is believed that the process involves etching a physical step into the surface of each pixel such that the fringing destructively cancels at a particular wavelength (likely around 900nm) using a quarter-wavelength step (Paul Clark, private communication).

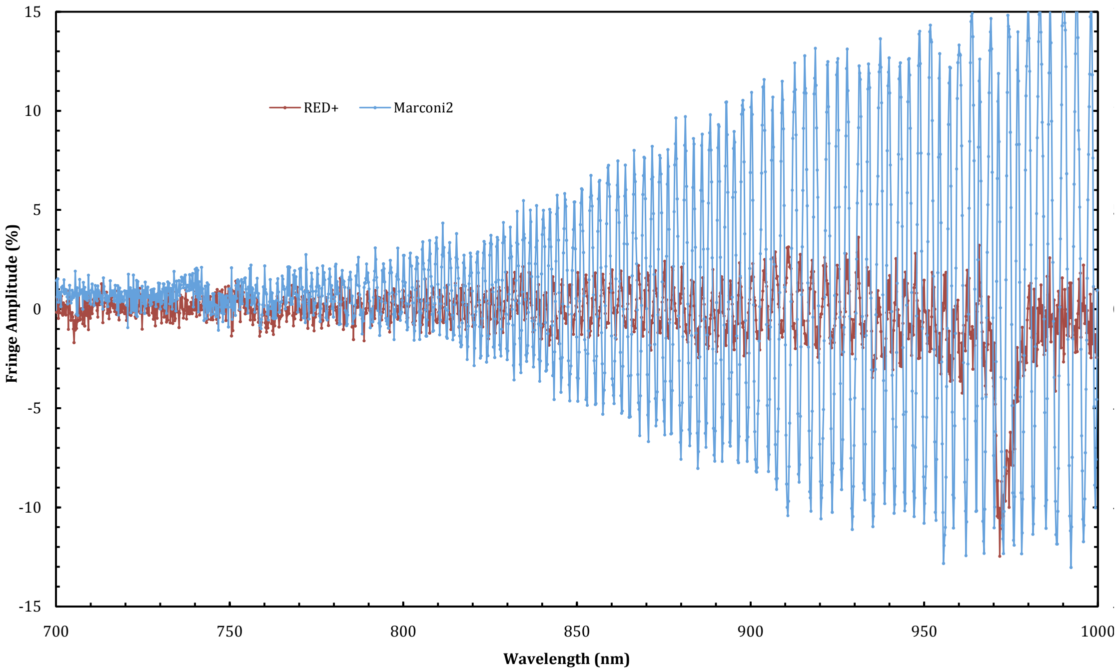

Figure 5.16: A comparison between two E2V CCDs at the Isaac Newton Group of telescopes, with and without fringe suppression. The red data points indicate the improved performance of the E2V-CCD4290 detector with fringe suppression as compared with the blue data points showing the performance of a non-fringe suppression E2V-CCD4011 device (Tulloch, 2005).

5.6.3 Red Detector Fringing Measurements

Fringing in the SALT HRS red E2V231-84 detector will be tested with an experimental setup closely resembling that used for the blue detector quantum efficiency measurements (see Section 5.5). The detector will be cooled to its optimum operating temperature of -127°C and illuminated at a variety of wavelengths using a flat-field beam expanded by a telescope and fed via a 6in diameter Newport integrating sphere from an Oriel monochromator illuminated by a stable quartz tungsten halogen light source.

Given the non-availability of the red detector at the time of writing, a photon transfer curve, quantum efficiency curve and fringing measurement have not yet been obtained.

5.7 CCD Cosmic Ray Events

5.7.1 Motivation

Cosmic ray events are a common feature of CCD imaging in astronomy. They manifest themselves by depositing significant quantities of charge in one or more pixe[illegible]heir presence in images wastes pixels, causes confusion with real signals (e.g. emission lines) and requires complex identification and cleaning in data reduction software (van Dokkum, 2001). Imaging data can be reasonably easily cleaned using median filtering since cosmic ray events may be distinguished from stars on the basis of their PSF. Échelle data is more difficult to reduce, since exposures are typically very long - as exposure time increases, so does the number of cosmic ray hits. This difficulty can be mitigated to some extent by co-adding multiple shorter exposures using median pixel values, but this technique carries the penalty of increased read-out noise from each exposure. There is also a degradation in the data should observing parameters such as seeing, sky transparency or sky emission line intensities change with time. The comparison of multiple exposures is also inaccurate if there is a non-integer pixel shift between exposures.

5.7.2 Composition and Distributions of Cosmic Rays

Primary cosmic rays are predominantly protons (79%), and occasionally helium (15%) or heavier element nuclei, which originate both in high energy processes in the sun (solar flares) and from galactic and extra-galactic sources (most likely supernova events). Protons and the nuclei of elements synthesized in stars are the truly primary sources although non-end-products of stellar nucleosynthesis such as lithium or beryllium also occur as a product of interaction of the primary sources with interstellar gas. The cosmic rays striking the detector in a typical ground-based astronomical spectrograph, are not the high-energy protons impinging on the upper atmosphere. Interaction with atmospheric particles produces a particle shower cascade, resulting initially in mesons which decay mostly into muons lower in the atmosphere.

Most muons are produced at an altitude of approximately 15km and lose about 2GeV to ionization before arriving at the surface (Gaisser & Stanev, 2009). The energy spectrum and angular distribution depends not only on the production spectrum but also on energy loss to the atmosphere and decay. For example, a 2.4GeV muon has a decay length of 15km, which is sufficient to reach the ground, but energy losses reduce the decay length to 8.7km. The mean ground level muon energy is therefore slightly positively skewed towards a peak at ∼4GeV. The energy spectrum below 1GeV is flat with gradual steepening reflecting the primary spectrum in the 10-100GeV range and showing further steepening at higher energies due to a pion contribution.

Figure 5.17*a* shows the vertical flux of cosmic rays as a function of depth in the atmosphere, broken down into constituent particles of protons/neutrons, muons, neutrinos, electrons/positrons and pions. Dashed lines on the graph represent the muon flux at the detector altitude in the optics laboratory where testing was performed (the NetPark Research Institute is at a height of 112m above sea level) and at the final installation site of the SALT observatory (which is located 1798m above sea level). Clearly the neutrino flux is inconsequential (due to the famously weak interaction of neutrinos) and the pion and electron/positron flux are very low, relatively speaking, at the ground. The dashed lines indicate a potential excess in muon flux at SALT compared with Durham of $\sim 20\text{m}^{-2}\text{s}^{-1}\text{sr}^{-1}$. This increase would represent a noticable but not troublesome ∼140 extra hits over the 18.9cm^2 area of the blue detector during a 3600s exposure.

Figure 5.17*b* shows the muon energy spectrum at sea level for two angles; $\theta = 0°$ and $\theta = 75°$. Red data points indicate $\theta = 0°$ whilst blue data points represent $\theta = 75°$. The shift in the spectrum between the two angles is a reflection of the increased energy loss of particles traveling through a greater depth of atmosphere at increasing angles. At large angles low energy muons decay before reaching the surface so that the average muon energy increases. The distribution of muons at the ground is $\propto \cos^2 \theta$ (Groom, 2002) which gives a difference of factor $\pi/2$ in flux between horizontally and vertically aligned detectors. Although a vertically oriented detector therefore receives fewer hits, it should also be noted that the track length of the hits will be longer since muon energy deposition is strongly dependent on the thickness of the depleted region.

The cosmic ray energy spectrum is also dependent on location and time. Solar wind variations

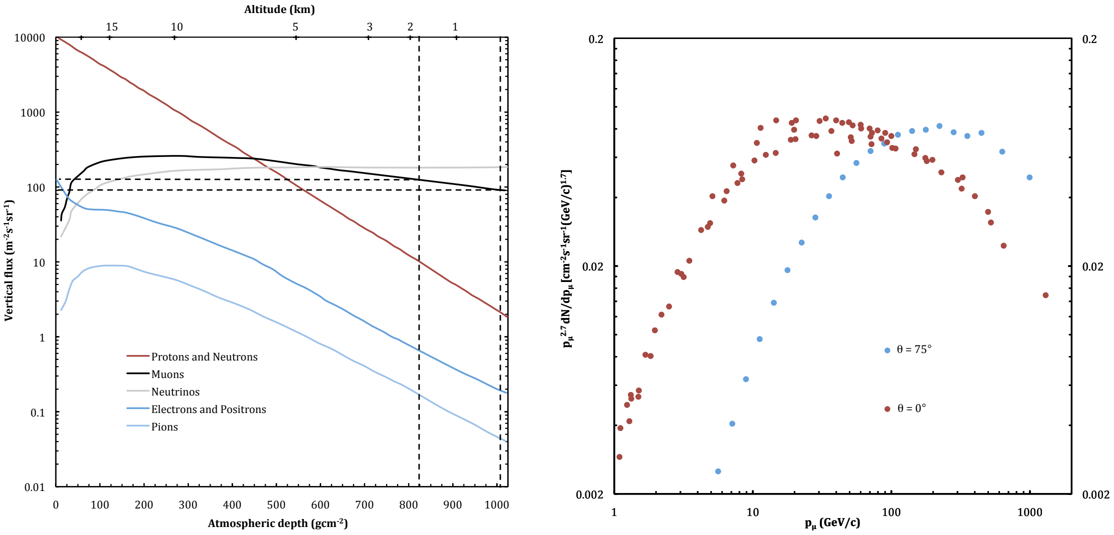

(a) Distribution of 1GeV cosmic rays as a function of atmospheric depth

(b) Muon energy spectrum at sea level for angles, $\theta = 0°$ and $\theta = 75°$

Figure 5.17: Cosmic ray distributions (both adapted from Gaisser & Stanev, 2009). Vertical dashed lines in Figure *a* indicate the altitudes of the optics laboratory and of SALT. Horizontal dashed lines indicate the muon flux at these al[illegible]es.

are known to cause variance in the lower energy region of the energy spectrum, since there is a significant anti-correlation with solar activity. The magnetized plasma in the solar wind acts to decelerate and therefore modulate cosmic rays in the inner solar system which partially excludes lower energy primaries. Lower energy ($\leq$1GeV) cosmic rays are also affected by the Earth's geomagnetic field which must be penetrated by incoming particles before reaching the upper atmosphere.

Figures 5.18*a* and 5.18*b* show a typical cosmic ray event from a selected 3600 second exposure taken with the SALT HRS blue detector, cooled to -110°C. The background values are around 720DN whilst the cosmic ray event shows values between 2000 and 3000DN in the central pixels.

5.7.3 Ground-based Sources

Other 'cosmic ray' events seen in the detector may not be strictly cosmic in origin at all, but rather the result of ground-based radiation sources. γ-rays from terrestrial radio-isotope decays are quite

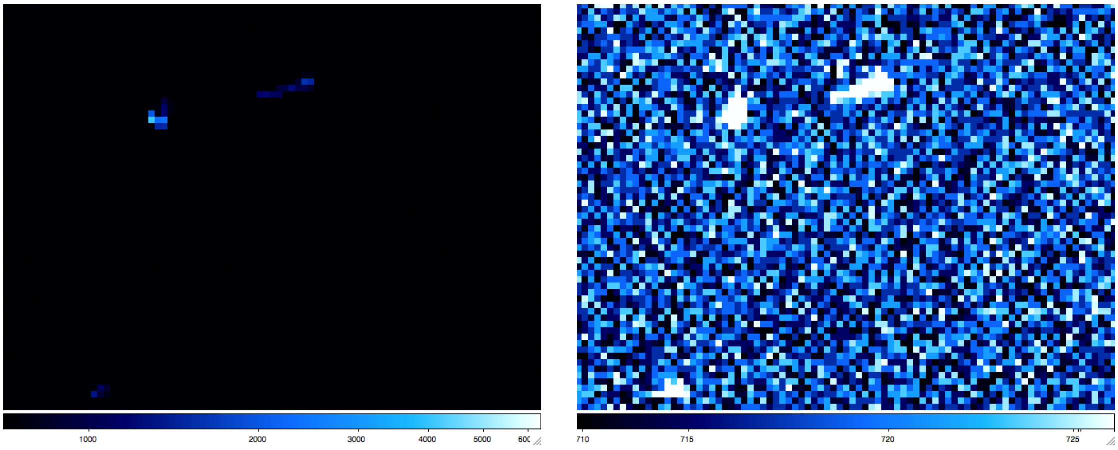

(a) Logarithmic scale

(b) Histogram equalization

Figure 5.18: A cosmic ray event viewed under different scales.

penetrating and could easily reach a detector, although α- and β-particle decay products are likely absorbed before travelling far (Groom, 2002), unless emitted by materials in the cryostat itself. Such ground-based sources (commonly produced by Potassium-40, Thorium-232 and Uranium-238) are inevitably present in the environment and in building materials such as concrete. There is also evidence for the presence of Potassium-40 in UBK7 glasses and for Cobalt-60 in a variety of metal alloys (Florentin-Nielsen et al., 1995). Cobalt-60 is introduced into many alloys during their manufacture. The furnace linings include the isotope as a radioactive marker for indicating new lining replacements, but the cobalt inevitably mixes with the alloy being produced.

The β-ray activity of a variety of cryostat materials can be seen in Table 5.2, as measured in 25mm diameter cylindrical samples with a low-level Geiger Müller counter at the Ris¢ National Laboratory, Denmark. Other samples were measured in a γ-ray spectrometer to determine the radio-isotope content, with the results available in Table 5.3 confirming the presence of large amounts of Potassium-40 in UBK7 glass and Cobalt-60 in Covar alloy. For this reason, BK7 glass is avoided for the final lens elements in the red and blue cameras which will also act as a CCD window in the final instrument configuration. Fused silica is used for the CCD flat window for detector tests.

Table 5.2: β-ray activity for materials commonly used in CCD cryostats (Florentin-Nielsen et al., 1995)

Material	Concentration (Bq kg^{-1})	Detection Limit (Bq kg^{-1})
Copper	0.00	0.11
Covar	0.03	0.07
Brass	0.06	0.07
qq Bronze	0.23	0.09
Titanium	0.28	0.15
AlMgSi Alloy	1.3	0.30

Table 5.3: Radio-isotopes identified by γ-ray spectrometer and their activities (Florentin-Nielsen et al., 1995)

Sample	K (g)	Ra-226 (Bq)	Th-232 (Bq)	Co-60 (Bq)
Steel 37	<0.02	<0.05	<0.1	<0.04
AlCuPb Alloy	<0.02	<0.05	<0.1	<0.04
Loral 2k CCD	<0.02	<0.07	<0.1	>0.7
UBK-7 lens	4.2	>0.1	<0.2	<0.07
Fused silica lens	<0.02	<0.06	<0.1	<0.04

5.7.4 Cosmic Rays in SALT HRS Detectors

Due to the high energies of true cosmic rays, protecting a CCD from their effects is likely to be unrealistic. However tests have suggested that ground-based events may be ameliorated with proper shielding. A pioneering set of experiments at the Lawrence Berkeley National Laboratory (LBNL) (Groom, 2002) indicated that cosmic ray muons produced straight tracks on the detector, whilst lower-energy Compton-scattered electrons generated by incident γ-rays from ground sources produced a 'worm' (a wandering track on the detector), or a spot. Figures 5.19*a* and 5.19*b* show a straight cosmic ray from a cosmic ray muon, and a worm from γ-ray electron scattering. Results were taken in four environments:

- an optics laboratory in which both straight muon tracks as well as worms and spots were visible, in roughly a 1:3 ratio;
- a low-background facility in which low-radioisotope concrete was used, showing muon events but a vast reduction in worms and spots;
- under 480m-water-equivalent depth at a Californian dam, where the muon count was vastly reduced;

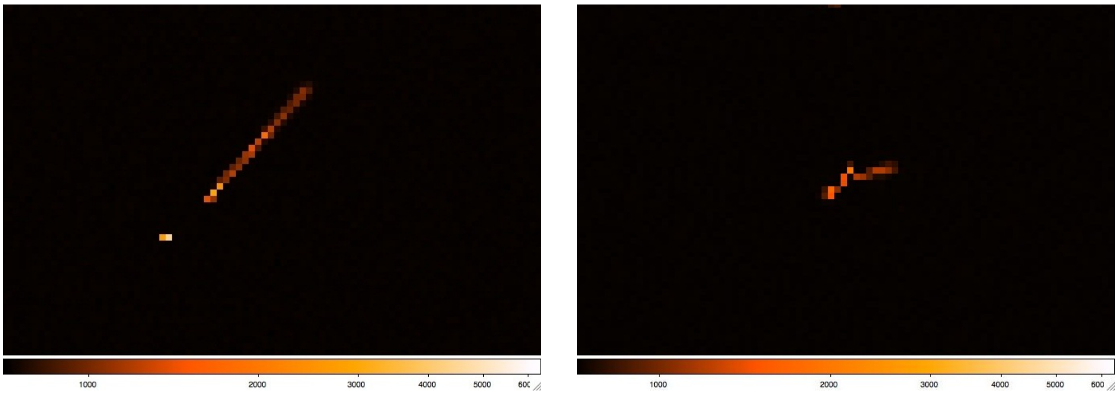

(a) A true cosmic ray

(b) A ground-based ‘worm’

Figure 5.19: Images showing the different form of muons and ‘worms’ as seen on the blue detector (both viewed on a logarithmic scale). The cosmic ray on the left has a straight track whilst the ground-based ’worm’ has the morphology of scattered electrons (as described by Groom, 2002).

- in a lead-lined clean room in the dam, where suppression of both muons, worms and spots was observed (with remaining counts thought to be inherent to the dewar assembly rather than the environment).

The work at LBNL suggests that lead shielding 1cm in thickness is sufficient to significantly lower the contamination from ground based ‘cosmic rays’, with the total count reduced by ∼80%. An experiment to confirm this using one of the SALT HRS detectors is presented below.

The blue detector was first cooled to its optimal operating temperature to suppress dark counts. Multiple 3600 second dark frame exposures were then taken at a variety of times of day, in an optics laboratory with concrete foundations, floor and walls. Background radiation in the room was measured using a Geiger counter model 900/44A from Thermo Electron Corporation and gave consistently 7-8 counts per second across a 180 second measurement period in a variety of locations within the room. The Geiger counter is sensitive to 15-500keV radiation. The astronomical standard SExtractor v2.3.1 package was used to identify cosmic ray events in the resulting blue detector images. Inspection of the resulting event co-ordinates by eye in a histogram equalized image suggest reliable identification of such events using this package.

5.7.5 Detection Thresholds

SExtractor first determines the background leve[illegible] frame before determining if a pixel belongs to the background or to an object. Those pixels which are deemed not to be background pixels are grouped into separate objects and written to a catalogue. An object is defined as having a minimum number of connected pixels above the background level (with a saddle point in the object indicating multiple objects ‘smeared’ together in the image). The ability of the software to identify cosmic rays is therefore sensitive to the input parameters selected. The parameter ‘MINAREA’ allows the user to set the number of contiguous pixels which must all be above the threshold in order for an event to be counted. For example, with a threshold of 4.5σ and MINAREA of 5 pixels, a group of three pixels 3σ above the background would not be counted as an event, since they are neither deviant enough from the background nor a large enough cluster of pixels to fulfill the criteria. Equally a group of ten 3σ or a group of four 10σ pixels would also not be counted. Before using SExtractor as a means of counting and performing geometrical measurements on cosmic ray events, the dependence on input parameters is investigated. Figure 5.20 shows the change in the number of events identified as a function of the σ threshold limit applied in the filter, for different ‘MINAREA’, using the same input image each time. As is to be expected, a greater number of events are identified when either σ or ‘MINAREA’ is reduced. In the move from σ=10.5 and MINAREA=6 to σ=1.5 and MINAREA=1, the number of detections increases by more than an order of magnitude, from ∼250 to ∼3400. The MINAREA=1 curve shows a marked increase in the number of detections for any given σ value, demonstrating a similar increase over the MINAREA=2 curve as that curve has over the MINAREA=5 curve. Given the surplus of MINAREA=1 events, it was however possible that these were hot pixel identifications rather than cosmic ray events (although it may not necessarily be expected to see a surplus of hot pixels at the σ=10.5 level).

A means of determining if a single identified pixel is a radiation event or a hot pixel was therefore required. SExtractor was run on each of the ten frames for each of the five shielding type data sets (see Section 5.7.6 for further details on the data sets) with σ=1.5 and MINAREA=1 identification parameters. SExtractor was set to create a so-called ‘segmentation’ image for each processed frame. This check frame is a .fits image in which the background is set to 0 and

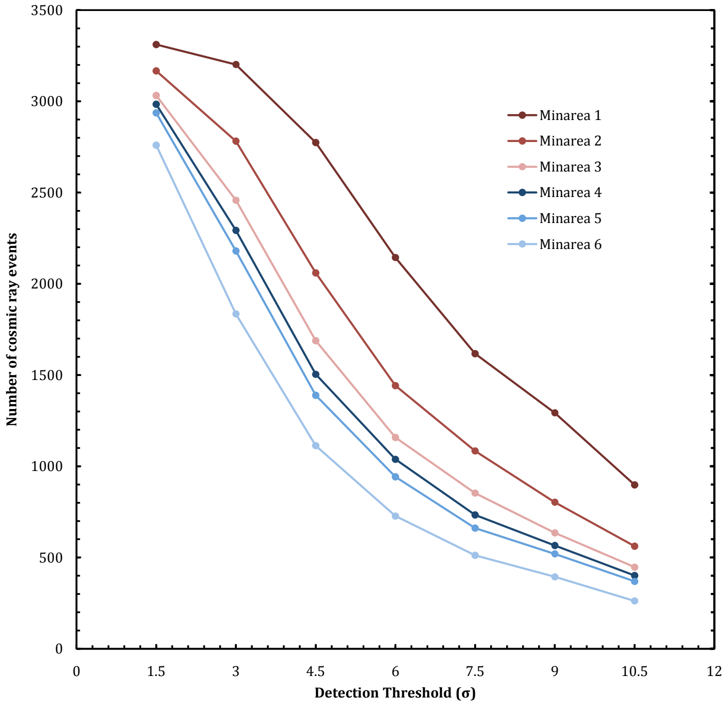

Figure 5.20: A plot of SExtractor identified objects as a function of detection threshold in σ above the background, at a variety of MINAREA values, for a 3600s full-frame exposure with the blue SALT HRS detector.

the value of any identified pixels is set to an integer depending on its order of detection by the software. For each data set, the ten 'segmentation' frames were compared using an IDL routine whose operation is outlined in Figure 5.21. The program first reads the segmentation .fits file, loops over x and y co-ordinates and finds any non-zero pixels. These non-zero pixels are then reset to the background value of 0 (should one or more of its neighbouring pixels also have a non-zero value), else set to a value of 1. This process is repeated for each of the ten frames, which are then overlaid into a $2048 \times 4096 \times 10$ pixel data cube. The code loops over x and y searching for potential hot pixels (value 1) and then compares the (x,y) position of the hot pixel to the same location in the other nine frames in z. The identification of a hot pixel in multiple frames then increases a counter accordingly. For all five data sets the results are similar, indicating that very few or none of the single pixel events identified by SExtractor are hot pixels identified in all of the frames. Several hundred pixels appeared in only one of the frames, with a few tens of pixels in two or three frames, with no single pixels above 1.5σ being identified in more than

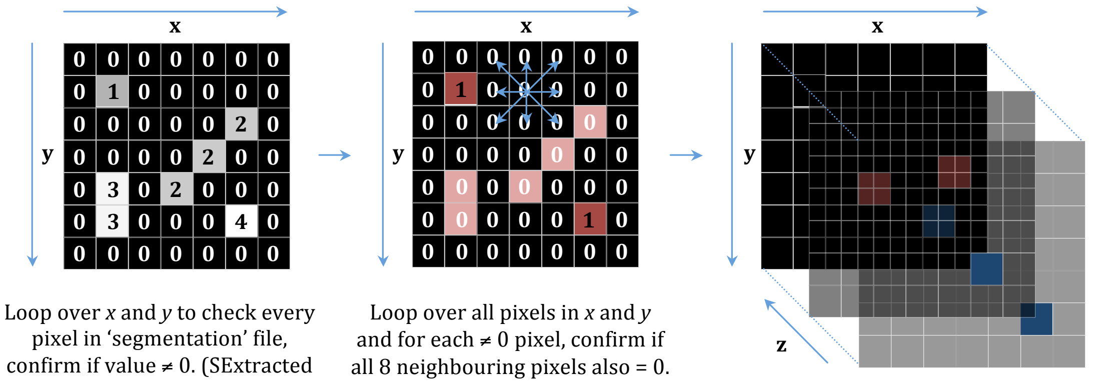

Loop over *x* and *y* to check every pixel in 'segmentation' file, confirm if value ≠ 0. (SExtracted pixel check)

Loop over all pixels in *x* and *y* and for each ≠ 0 pixel, confirm if all 8 neighbouring pixels also = 0. If so, change pixel value to 1. If not, change value to 0. (Now only potential hot pixels have non-zero value)

Add each frame as a layer of a data-cube. Loop over *x*, *y* and *z* in the cube to count in how many of the frames a pixel is 'hot' (=1) and increase a counter accordingly

Figure 5.21: The process of hot pixel identification in IDL using stacks of SExtractor segmentation image files.

three of the ten frames. This encouraging result confirms that using SExtractor it is possible to identify a considerable fraction of the potential cosmic ray events without having the number of events distorted by the presence of fixed detector artefacts.

Figure 5.22 shows the frequency distribution of the total signal (DN) in each of the dark image frames. Strong agreement across all ten frames is demonstrated, with a distinct peak in frequency for events characterized by a signal in the 1000-1500(DN) bin.

5.7.6 Blue Detector Shielding Tests

A housing for the blue detector was constructed from lead bricks, each 76mm in width, 51mm tall and 152mm in length. The housing features shielding on up to six faces of a cube (up to ∼180 bricks), with the top face supported by a particle-board plate ∼35mm in thickness. This plate was also checked for radioactivity using a Geiger-counter and gave a reading of 7 counts per second, which is identical to the background count rate in the laboratory. Sets of ten 7200s dark exposures were taken with a variety of shielding options: full 6-side shielding with 76mm thick lead walls; full 6-side shielding with 51mm thick lead walls; floor and wall only (5-side) shielding with 76mm thick lead (no 'roof' shielding); floor only (1-side) shielding with 76mm thick lead

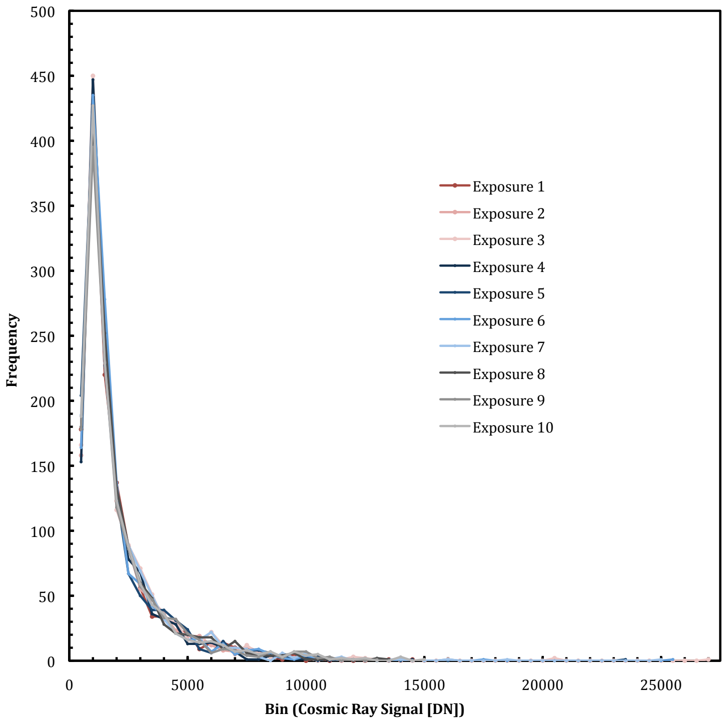

Figure A histogram displaying cosmic ray flux distribution across multiple 3600 second dark exposures made using the cooled SALT HRS blue detector in laboratory conditions.

(no ‘walls’ or ‘roof’); and finally, no shielding at all. Using identical SExtractor thresholds of 1.5σ and a MINAREA of 1 in each case (for reasons outlined previously), and averaging the data from all ten directly comparable exposures, the result plotted in Figure 5.23 was obtained. The histogram of the frequency of events identified within a given signal count bin shows a clear skew towards a large number of events at lower signal rates, exemplified by the strong peak in the 1000-1500DN signal bin. The plot also shows the expected decrease in the number of events as the volume of lead shielding increases. Given the expected high-energy nature of events caused by true cosmic ray muons, it is perhaps not surprising that the full 76mm thick shielding proves only slightly more effective than 51mm shielding at reducing the number of detected events in a given signal bin, at least in the low signal regime. This is because the 51mm shielding is of sufficient depth to attenuate the lower energy ground-based sources, yet the additional depth the 76mm shielding offers is still insufficient to block the very high energy cosmic rays. Due to the relatively low occurrence of events with a large signal (either a small number of pixels with

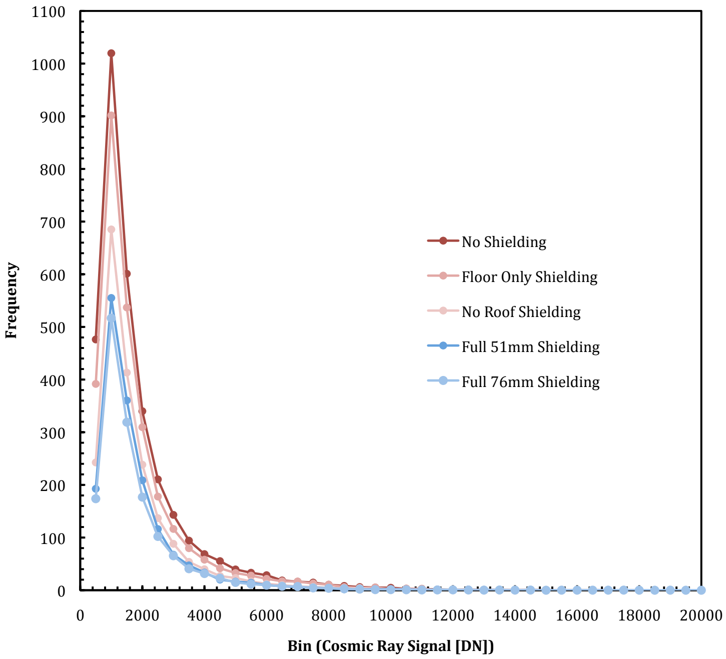

Figure 5.23: A histogram of cosmic ray flux distribution for various types and depths of lead shielding, showing a trend of reduced flux with increased shielding.

a high signal or an extended object containing a contribution from many pixels) it is difficult to accurately characterize the effect of shielding in the high signal level tail. Figure 5.24 shows the same data as Figure 5.23 but re-binned on a logarithmic S(DN) axis to more clearly show the frequency bins at high signal level. The effect of shielding is also seen to extend to the higher signal detections, although the benefit of thicker 76mm lead versus 51mm lead still seems negligible in this regime region too.

Figure 5.25 presents these results as a ratio plot of the ratio of frequency of event occurrence (in a given signal bin) with a specific level of shielding against its frequency in the case of no shielding. Again a logarithmic S(DN) axis is used for binning. Here we can see the effect of shielding more clearly. The floor only shielding offers little benefit over using no shielding at all in terms of event reduction (presumably because very few local events are coming 'up through the floor'). The effect of shielding all four sides and beneath the detector are much more pronounced, with a 40-50% reduction in the number of events recorded across the event signal level range. This suggests a large fraction of shieldable incident radiation is incident from the sides of the

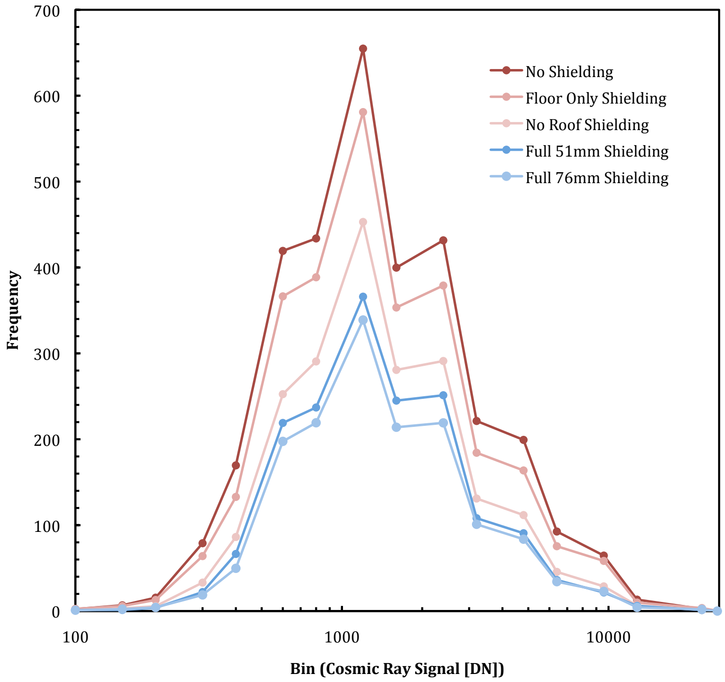

Figure 5.24: A histogram of cosmic ray flux distribution for various types and depths of lead shielding, with a logarithmic signal scale.

detector rather than 'straight down' from the sky. This could be due to the diffuse nature of ground-based radiation sources, or from a contribution of true cosmic rays at large angles from the zenith (particularly since the detector is vertically oriented). A 50-70% decrease across the range can be seen with full 51mm and 76mm thick shielding, with only a marginal difference between the two. Given the low number of events in the final 10000-20000 and 20000-30000DN bins, the random noise makes any conclusion about the effect of shielding impossible for these events.

5.7.7 Cosmic Ray Event Characterization

A suitable method for differentiating events caused by ground-based radiation from true cosmic rays is required before any conclusions about the reduction effects of shielding can be drawn. For example, there is no way to tell if the 'no roof' shielding case reduces ground-based events from a variety of angles or if it is absorbing cosmic rays from close to the horizon. Using the signal

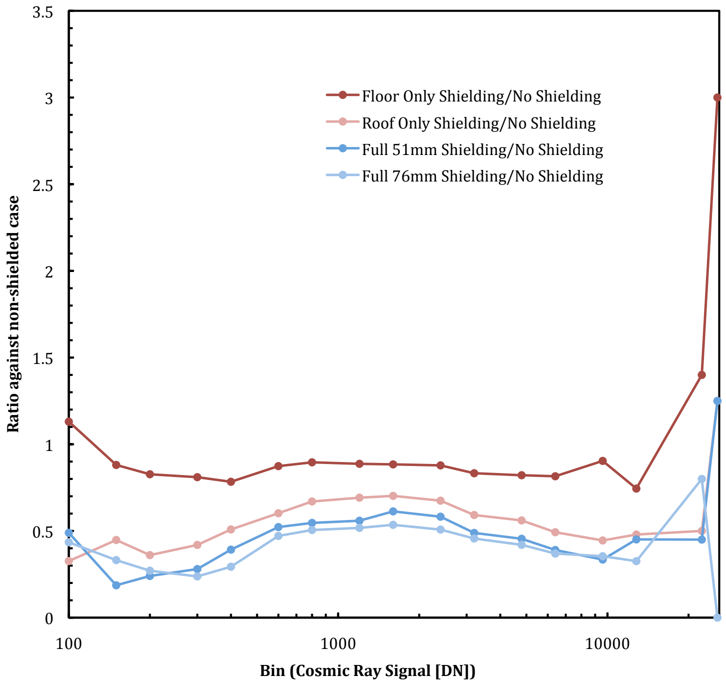

Figure 5.25: The ratio of cosmic ray flux plotted as a function of event signal st h (DN), under various shielding conditions to flux in the case of no detector shielding.

count as a means of identifying the event cause is also dubious. A high count rate could equally be due to a large energy deposition from a normally incident cosmic ray (where a large amount of energy is deposited in one pixel) or from a highly scattered ground-based source of much lower energy but at a grazed angle of incidence which spreads the event across many pixels. For this reason the geometry of each event must be investigated to separate the two sources. Here the 'worm' versus 'spot' versus 'muon streak' identification introduced by Smith et al., (2002) is useful.

Smith et al. (2002) use a system of 'perp. counts' and second moments information to catalogue the cosmic ray event zoo on a detector. The 'perp. count' quantity refers to 'perpendicular counts', a name given to the re-normalized signal in each event, as demonstrated in Figure 5.26*a*. The signal in each event is scaled by the depletion region thickness in the wafer divided by the track length, which is calculated from the observed projected length of the event. 'Perp. count' is therefore correcting each event to have the signal which it would have shown if it had been incident on the detector at 90° to the surface. The blue E2V44-82 detector has a depletion region

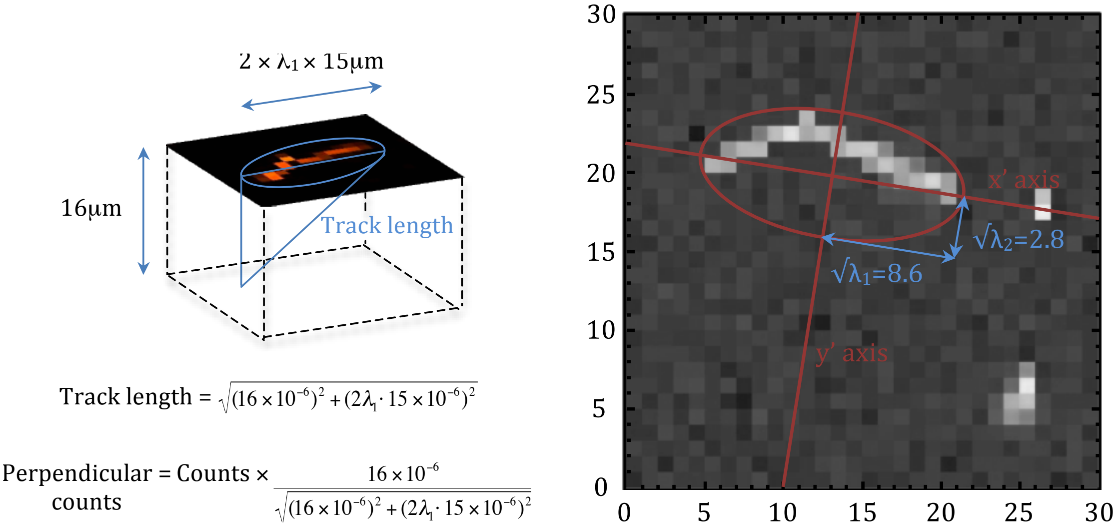

(a) Perpendicular counts calculation (b) Eigenvalue definitions

Figure 5.26: The definitions of 'perp. counts' and the second-moments matrix eigenvalues λ_1 and λ_2 as used for event charaterization in Smith et al. (2002).

of 16μm depth (E2V, private communication). Figure 5.26b shows the definitions of geometric parameters derived for each event, that is λ_1 and λ_2, the maximal and minimal eigenvalues of the second moment matrix respectively (where the principle moments identify the barycentre of the event). SExtractor is able to calculate A and B values for each identified event, where A and B are the semi-major and semi-minor axes of a fitted ellipse. This information may be used to empirically determine a 'pe[illegible]ounts' and λ_2 cut for event characterization.

Results from the LBNL data is shown in Figure 5.27. The cut on perp. counts can be seen in Figure 5.27a, which is a frequency histogram of the perp. counts parameter. The distribution follows a typical Landau curve in which a strong Gaussian-like peak is followed by a high energy 'tail'. The cut is defined conservatively to encapsulate the majority of muons. When applied to the plot in Figure 5.27b, an obvio[illegible]ıon box' exists at $0.25 \leq \lambda \leq 0.55$. They confirm this cut-off by comparison with a frame with little muon contamination (taken deep underground).

Unfortunately when applying the same principles to the SALT HRS blue detector cosmic ray data, a 'muon box' cut-off is much less obvious. A cut on perp. counts may be defined easily

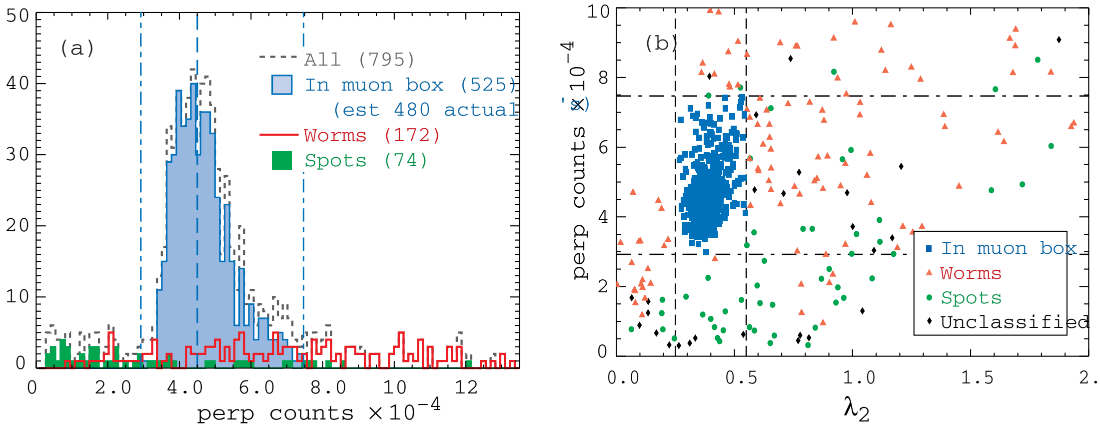

(a) Perpendicular counts distribution

(b) Perp. counts as a function of λ_2

Figure 5.27: Cosmic ray distributions for a 4800s exposure with a 300µm thick detector at LBNL (from Smith et al., 2002).

from Figure 5.28*a*, which shows the characteristic Landau distribution. However a plot of perp. counts as a function of λ_2 is less revealing (see Figure 5.28*b*).

In a reversal of the checks applied at LBNL, Figure 5.29 shows the perp. counts versus λ_2 plot with the events extracted at various levels of shielding over-plotted. Since the addition of lead shielding is expected to reduce the number of ground-based events ('worms') by a larger factor than muons, it was hoped that a region of data would 'disappear' with shielding, thereby indicating a potential remaining muon component. This result was not demonstrated in Figure 5.29 which shows a general reduction in the number of particles, but not at a specific geometry or perp. count level. Randomly sampling a smaller number of events (in case any signal is saturated by the large number of data points) also shows no clear distinctions.

For this reason, plots of λ_1 versus λ_2 and perp. counts versus ellipticity (ellipticity, e, is also calculated by SExtractor using $e = 1 - {}^{\lambda_2}/_{\lambda_2}$) were also produced for diagnostic purposes. These plots can be seen in Figure 5.30. The dashed line in Figure 5.30*a* represents the line of $\lambda_1 = \lambda_2$ equality, which explains the absence of data points to the left of this line. The solid lines indicate a tentative estimate of the λ_2 cut-offs. Comparing the distribution of shielded and non-shielded data indicates a potential 'muon band' which remains at $0.68 \leq \lambda_2 \leq 0.74$.

It may be expected that the events showing higher perp. counts are caused by muons and are

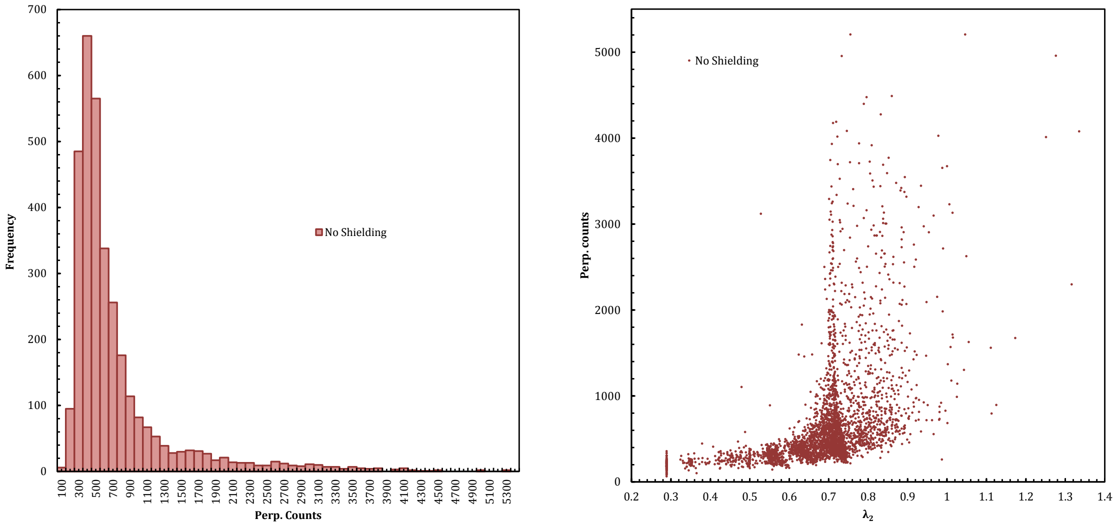

(a) Blue detector perp. counts distribution (b) Blue detector perp. counts as a function of λ_2

Figure 5.28: Cosmic ray distributions for a 7200s exposure with the SALT HRS blue detector. Whilst the perp. counts cut-off in (*a*) remains fairly clear, the 'muon box' in (*b*) is difficult to identify.

therefore 'straight'. This would present as a large perp. count contribution at large ellipticity. However, the plot of perp. count versus ellipticity (Figure 5.30*b*) shows no such contribution.

The absence of single pixel (low λ_2) events with a perp. count value above 1000 is shown in Figure 5.28*b*. This may indicate that once a deposited signal has reached a certain threshold, the charge within this pixel begins to spread into neighbouring pixels. Figure 5.31 attempts to determine whether such an effect is caused by charge diffusion or as a consequence of charge transfer efficiency (CTE) effects in the detector. If high signal events lose charge to the surrounding pixels by diffusion, the geometry of such events would be randomized amongst neighbouring pixels, therefore showing no directional trend. However, should high signal events show a preferential alignment, the effect is likely due to a charge transfer effect as charge is drawn in the direction of the readout port (a single port, left channel readout was used exclusively in all cosmic ray exposures). Figure 5.31*a* shows a plot of θ as a function of event ellipticity, where θ is the event semi-major axis angle with respect to the detector column direction as determined by SExtractor. Events with a high 'straightness', are aligned predominantly with the $\theta = 0°$ direction, and to a

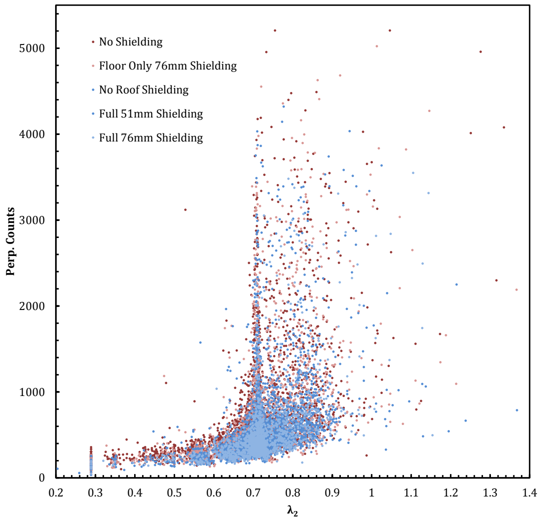

Figure 5.29: The 'muon box' distribution for the blue detector remains e[illegible]e when the events recorded under shielded conditions are over-plotted. Only a reduction in the number events is seen, with no obvious reduction for a particular perp. count level or λ_2 range.

lesser extent at $\theta = 90° = -90°$ directions. Figure 5.31*b* shows a plot of θ versus perp. counts, in which a large fraction of events (particularly those with high perp. count) cluster along the $\theta = 0°$ direction. This result suggests a propensity for high signal events to smear in the direction of the readout port as a result of CTE effects intrinsic to the detector. Given this result, it is entirely possible that some muon 'streaks' are actually caused by detector effects under high signal deposition. The smaller clustering effect at $\theta = 90° = -90°$ may be due to a smaller charge diffusion effect, since charge is unlikely to diffuse to diagonally located pixels.

5.7.8 Summary of Cosmic Ray Detection Measurements

Integrating the curves in Figure 5.23, the cosmic ray arrival rate is shown to be $\sim$1.4cm^{-2}min^{-1} without any shielding. This compares favourably with the results shown in Table 5.4[1]. The addition of lead shielding on all sides of the detector reduce this rate by more than half to $\sim$0.6cm^{-2}min^{-1}. An increase in shielding depth from 51mm to 76mm shows little additional benefit. Whether a significantly thinner, and hence more practical depth of shielding is capable of similar event rate reductions is unproven here, however literature suggests that 15mm is

---

[1] "On the Rates of Radiation Events in CCD's", from an ESO report in Messenger **47**

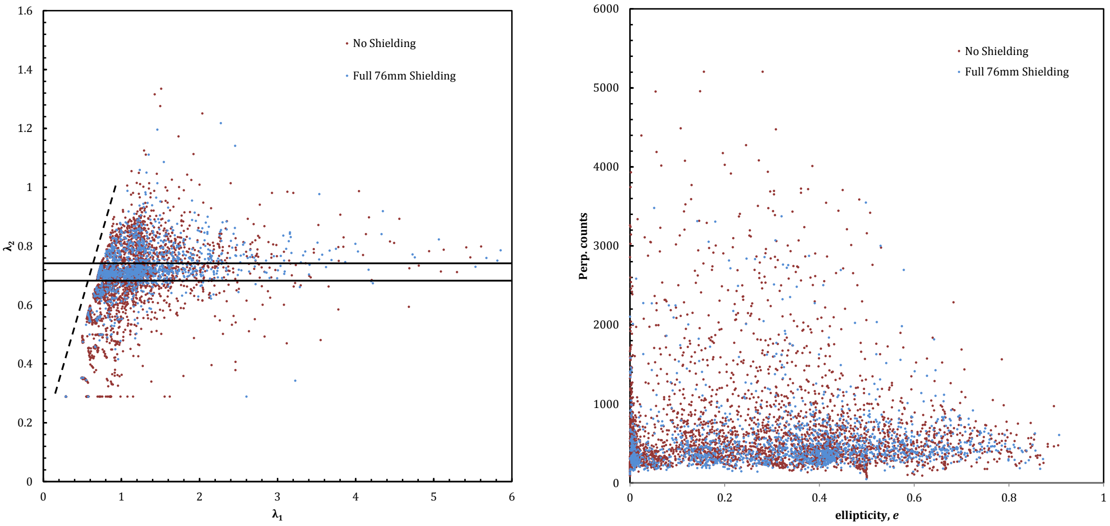

(a) Blue detector λ_1 versus λ_2 distribution

(b) Blue detector perp. count versus ellipticity distribution

Figure 5.30: Cosmic ray distributions for a 7200s exposure with the SALT HRS blue detector. (a) may demonstrate a potential 'muon band' at $0.68 \leq \lambda_2 \leq 0.74$ which seems less diminished by shielding effects than events in the rest of the distribution. (b) does not show the expected high perp. count contribution at high ellipticity caused by muons.

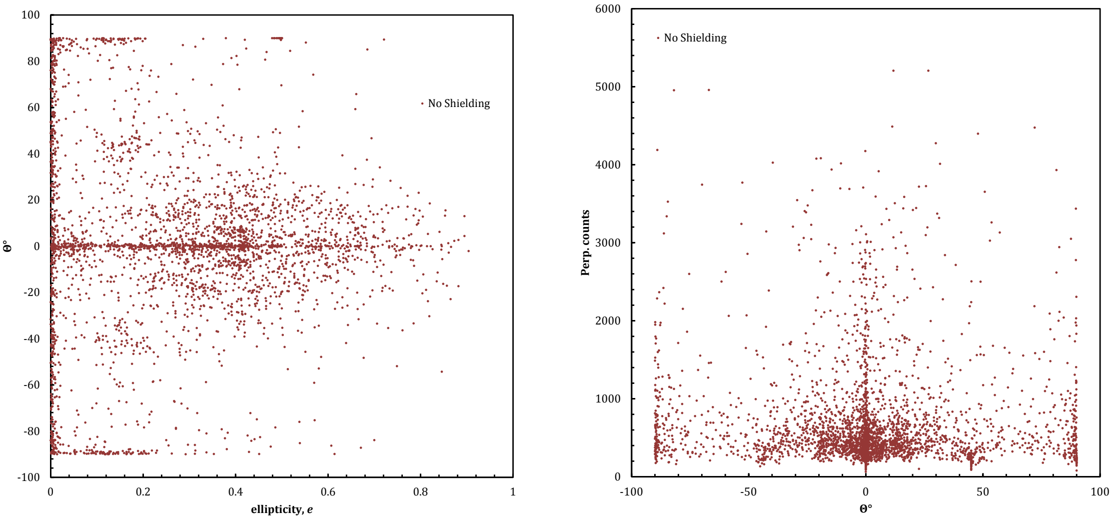

(a) Blue detector θ versus e distribution

(b) Blue detector θ versus perp. counts distribution

Figure 5.31: Cosmic ray distributions for a 7200s exposure with the SALT HRS blue detector. (a) demonstrates the preference of 'straight' events, perhaps caused by muons, to align in the direction of the readout port. b shows a similar effect for events with high perp. count signal.

Table 5.4: The frequency of radiation events in a variety of CCDs

CCD Type	Number Events cm^{-2}min^{-1}
RCA SID 501EX (thinned)	5.8
RCA SID 501EX (thinned)	5.2
RCA SID 501EX (thinned)	5.7
RCA SID 501EX (thinned)	6.6
RCA SID 501EX (thinned)	6.1
GEC 8603	1.2
GEC 8603	2.1
GEC 8603	2.3
RCA SID 006 ES	5.2
TEK 512 M-11	1.4
E2V44-82	1.27

sufficient to achieve a similar reduction factor (Smith et al., 2002). Results using shielding in limited directions indicate that there is no prevailing direction of cosmic ray incidence. Shielding on all six sides of the detector is therefore required.

Given the difficulties in identifying an events source (cosmic muons or local radioactive decay products) based on geometric or signal properties of the associated event, further investigation is required. The more clearly defined results described in Smith et al. (2002) likely stem from the use of a CCD with considerably thicker depletion region (300μm versus 16μm). Using a detector with a considerably thicker depletion region would allow for longer event tracks, thereby increasing the dynamic range of geometries observed and emphasizing the difference between 'streaks' and 'worms'. It may also be interesting to repeat measurements with the detector mounted horizontally, rather than vertically, to quantify both the change in event detection rates and morphologies.

5.8 CCD Window Deformation

As mentioned in Sections 5.3 and 5.4, many of the experiments performed on the detectors require a flat window whilst the final implementation of the camera optics and detectors calls for a concave optical element to form the final lens of the camera system. As in Tables 2.4 and 2.5 of Section 2.2.10, the final elements in the red and blue cameras are RCM5.1 and

BCM7.1 respectively. Both lenses are concave and manufactured from fused silica. Since the large instrument tank and the smaller CCD cryostats are under different vacuum pressures, the link between the final two optical elements in each camera (RCM4.1-RCM5.1 and BM6.1-BCM7.1 respectively) must be under ambient air pressure. RCM4.1 and BCM6.1 on the camera barrel side bear the light vacuum of the instrument tank with thicknesses of 40mm and 55mm respectively over their 150mm diameters. Under greater risk of pressure induced flexure are the 13mm thick RCM5.1 and BCM7.1 lenses, also 150mm in diameter. These lenses are not only much thinner but also bear a much greater vacuum load at the CCD cryostat interface. There was a concern that these thinner lenses may deflect sufficiently under pressure to impair the image quality of the camera. In order to check this, the deflection, d, of a circular plate resting on its edge may be calculated according to the equation (Miller, 2003).

$$d = \frac{Pa^4}{64D} \tag{5.14}$$

where

$$D = \frac{Et^3}{12(1-\nu^2)} \tag{5.15}$$

P is the pressure acting on the plate (10^5 Pa), a is the plate radius (0.075m), E is Young's modulus of elasticity (73 GPa for fused silica), t is the plate thickness (0.013m) and ν the Poisson coefficient (0.17 for fused silica). Substituting the aforementioned values into equations 5.15 and 5.14 in turn gives a central deflection of 3.6µm. The sag of the surfaces is related to their curvature radius, r, by the formula

$$s = r - \sqrt{(r^2 - a^2)} \tag{5.16}$$

The sag can therefore be changed by the deflection amount in order to recalculate the new radius of curvature, which can then be checked against permitted tolerances to ensure image quality is not unacceptably degraded. In reality the deflection should be less than the quoted value since the 13mm central thickness of the lenses is its minimum thickness, with the concave lenses becoming considerably thicker towards their edges. Table 5.5 shows the worst case scenario

Table 5.5: Calculated radius changes for a 13mm fused silica plate as compared with lens tolerance specifications

Channel	Symmetrical Surface				Cylindrical Surface			
	Sag s(mm)	Radius r(mm)	Δr (mm)	Allowed Δr(mm)	Sag s(mm)	Radius r(mm)	Δr (mm)	Allowed Δr(mm)
Red	18.292	162.9	0.030	±1	4.455	633.48	0.51	±5
Blue	23.425	131.775	0.017	±1	4.005	704.2	0.63	±10

case of the red channel lens which demonstrates that the deflection effect falls within ∼10% of the tolerance allowed for in the image quality error budget. Other values show a smaller effect.

Chapter 6

Calibration and Quicklook Data Reduction

6.1 Introduction

This section deals with the wavelength calibration and quick-look reduction of data from the SALT HRS instrument. The spectra recorded on the CCD are only a list of intensities at x and y positions on the detector and have no intrinsic 'colour' in the way a colour photograph does (or indeed the way the simulated spectra in Figure 1.14 of Section 1.4.3 does). The échelle spectra are also not a simple one-dimensional spectrum, but instead consist of a number of orders m, with wavelengths λ appearing in multiple orders (and at different intensities and locations within different orders). The orders are not uniformly spaced across the cross-dispersion λ_{XD} (or y) direction and are also tilted (see Section 4.2.3). Furthermore, each order has a profile of intensity in the y direction and the area in between orders may contain scattered light and ghost images. All of these effects need to be accounted for in the quick-look data reduction pipeline.

6.2 Wavelength Calibration

A means of mapping a location on the chip to an actual wavelength is required for all modes of SALT HRS operation. The precision and accuracy in the dispersed spectrum to which wavelength calibration can be performed ultimately determines the radial velocity precision of the instrument.

6.2.1 Thorium Argon Lamps

A Thorium-Argon lamp (often shortened to 'Th-Ar' or called an 'arc spectrum') is a standard means of wavelength calibration. Its rest-frame spectrum is well known and quantified from laboratory measurements and atomic physics calculations. Illuminating the instrument with a Th-Ar lamp allows for a comparison of the spectrum recorded with an atlas of known lines and their intensities. In theory this is a 'once only' calibration which can be applied to astronomical spectra. However, in reality instrumental shifts (changes in path lengths due to small movements or expansion/contraction of optical components) over short time periods can make a 'one off' measurement inaccurate. In the case of many astronomical spectrographs, the Th-Ar exposure is made immediately before or after a science object exposure to improve the accuracy of calibration. Even with these efforts the precision may still be limited to the order of several hundred ms^{-1} since the object and Th-Ar exposures are not contemporaneous. The ELODIE and HARPS spectrographs (which are designed specifically for high wavelength precision for Doppler measurements) tackle this issue by allowing a simultaneous Th-Ar lamp exposure, interlacing the spectral orders on the detector (Pepe et al., 2002) as can be seen in Figure 6.1. This system is effective but is not without limitations. Since the simultaneous spectra follow slightly different optical paths through the instrument, and are imaged onto different regions of the detector, the calibration is still imperfect. There are also other issues with the use of Th-Ar lamps besides their implementation: the lamp is active (requires an input voltage which will be unstable on small time scales) and also deteriorates over time, subtly altering the output spectra and also requiring eventual replacement. The SAC system on the tracker assembly of the SALT telescope features an illumination system capable of mimicking the illumination of the SALT pupil at any

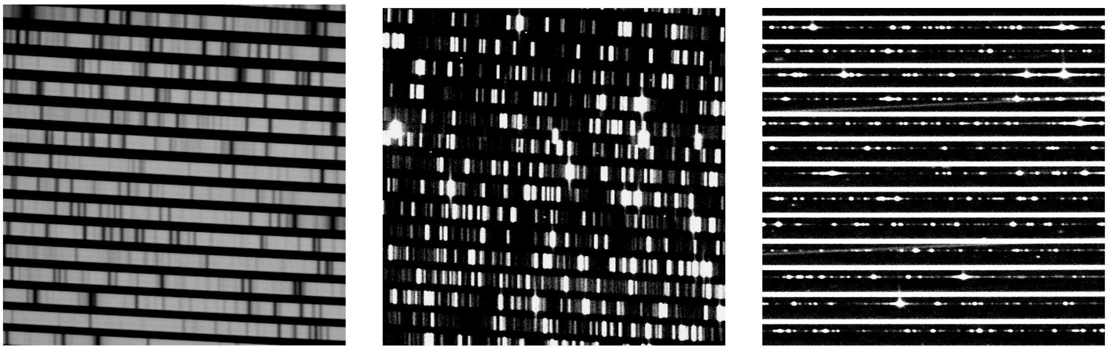

(a) Object spectrum (b) Th-Ar spectrum (c) Object and arc

Figure 6.1: An object spectrum (*a*); a non-simultaneous complimentary Th-Ar wavelength calibration exposure (*b*); and a simultaneous object/Th-Ar exposure on the same detector (*c*) (from, Pepe et al., 2002).

particular position in the telescope track. This ensures the calibration light is injected [illegible]o the instrument in a similar manner to the object source. SALT HRS also features simultaneous Th-Ar calibration within the high-stability mode only. In this case, the iodine cell (see Section 6.2.2) is removed from its position in between the two fibre double scrambler lenses (see Section 3.3), and a fold mirror slides into this gap and injects Th-Ar light via an optical fibre linked to the lamp. Although this method removes the pupil illumination simulation effects of the SAC calibration optics, the use of the double scrambler for the science light in this mode make it a moot point. Figure 6.2 shows the SAC calibration light exiting an 8mm diameter fibre bundle.

6.2.2 Iodine Absorption Cell Techniques

Due to the issues mentioned above in Section 6.2.1, a different solution for calibration of extremely high precision Doppler spectra was introduced in the 1970's in the form of the absorption cell. Light from the object is passed through a gaseous medium so that absorption lines with known wavelengths are imprinted on the object spectrum, whilst ensuring that the object and calibration light follow the same optical path at the same time. Originally this method used telluric absorption lines, whereby oxygen lines caused by the Earths atmosphere were used as the calibration lines. At best, this technique provided a precision of 30ms^{-1} however, due to

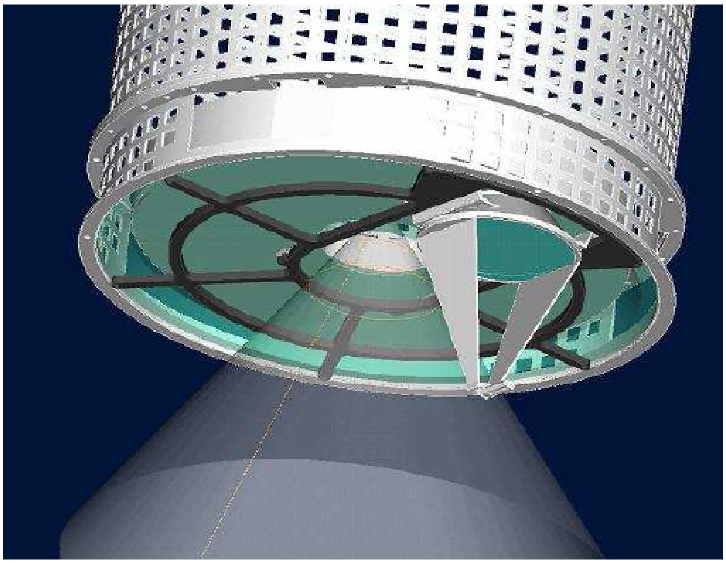

Figure 6.2: Th-Ar light exiting the SAC calibration optics system (Meiring & Buckley, 2004).

several issues:

- Atmospheric turbulence and changing winds have an effect on the pressure and temperature of the telluric gases.
- Line depths vary with airmass as an object moves across the sky.
- The lines are limited to a wavelength range of a few tens of Ångstroms.
- It is not possible to measure the object spectrum without viewing telluric lines, making a comparison impossible.

For these reasons a better method is to use an absorption cell which operates under controlled conditions and which can be removed from the beam when required. Campbell and Walker (1980) used the first of such cells to achieve an unprecedented radial velocity precision of 13ms^{-1}, using a cell containing Hydrogen Fluoride (HF). Figure 6.3 shows a plot of solar spectra taken with and without the ‘imprinted’ absorption cell lines. There were however still some drawbacks with this technique: HF is toxic; it must be re-filled after each observing run; it requires a long 1m path length; and it needs to be stabilized at the relatively high temperature of 100°C.

Following on from the HF cell, iodine cells offered several advantages:

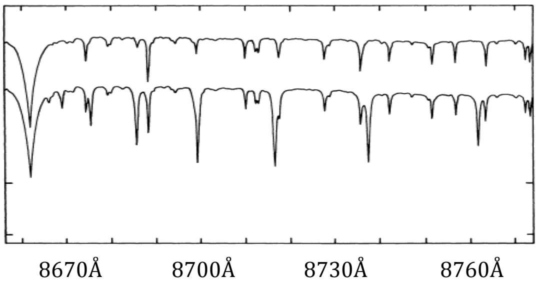

Figure 6.3: Solar spectra taken with ([illegible]) and without ([illegible]) imposed HF lines (http://www.tls-tautenburg.de/research/vorles.html).

- Large wavelength coverage of 1000Å.
- A lower stabilization temperature of ~70°C.
- The sealed cell design should provide uninterrupted service throughout the life-cycle of the instrument.
- Iodine is non-toxic.

In the SALT HRS high-stability mode, radial velocity precision is paramount, and an iodine cell was therefore chosen as the main wavelength calibration mode. The use of an iodine cell is only as precise as its own calibration however. The instrument profile (IP) of a real spectrograph, which has imperfect optics and alignments, smears a δ-function input (monochromatic beam) into a roughly Gaussian profile. The IP itself is not a major issue, but rather changes in IP, which cause shifts in the centroid of a line, which will then manifest themselves as velocity shifts. The iodine spectrum observed using the instrument can be compared with an iodine spectrum made at much higher spectral resolution using a Fourier transform spectrometer (FTS) (Butler & Marcy, 1995). The instrument spectrum is co[illegible]ed with the FTS spectrum to calculate the IP. Removal of IP asymmetries in this way allows for much greater radial velocity precision.

6.3 Exposure Meter

SALT HRS is to be fitted with a Hammamatsu H7421 photon counting module for use as an exposure meter. To avoid loss of light in the instrument, the exposure meter operates using

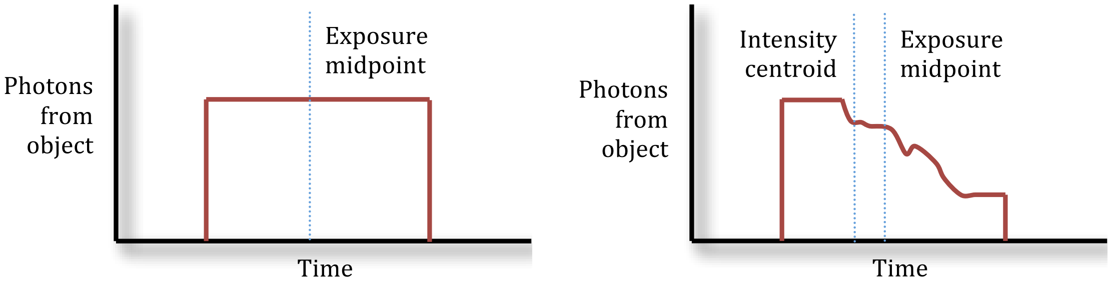

(a) ‘Perfect’ conditions (b) Varied cloud cover

Figure 6.4: With varying imaging or seeing conditions the exposure midpoint is not an appropriate measure of when the observation was taken. Here an exposure meter allows calculation of the flux-weighted midpoint.

light that would otherwise be wasted. Light incident on the gap between the two échelle grating mosaics is reflected back to feed the exposure meter, representing approximately 2.3% of the total light input. This light is fed into an optical system which illuminates the photomultiplier tube (PMT) outside of the vacuum tank. The PMT is located outside the tank to ensure heat from the electronics does not perturb the temperature stability of the instrument. Performance over the full wavelength range of the spectrograph is achieved with the PMT, but it does not have constant quantum efficiency over the whole range. It must therefore be calibrated accordingly. Since it is not possible to directly view the fibre input to HRS, the exposure meter may also be used in conjunction with the SALT imaging camera (SALTICAM) to acquire targets, using bright stars as a means of peaking the HRS signal. This will require a moving average count to be provided in real-time. A time-indexed record of the photon count during each exposure will also be stored in the FITS headers of the image so the time of the flux-weighted exposure mid-point can be calculated (see Figure 6.4). Under perfect atmospheric conditions, the photon flux from a star is roughly constant, but should conditions change, for example if clouds come overhead, the flux will decrease. This changes the effective mid-point of the exposure, which is important in radial velocity studies where a precise correction for the Earth’s orbital motion must be applied.

6.4 Data Reduction

This section describes a quick-look data reduction package which will be delivered with SALT HRS for use by SALT astronomers at the telescope. A more refined échelle data reduction package may be more appropriate for final 'science grade' reduction.

6.4.1 FIESTool

FIEStool is a Python based software package developed by Eric Stempels for the rapid reduction of échelle spectra from the FIES (FIber-fed Échelle Spectrograph) instrument at the Nordic Optical Telescope. It features a graphical user interface (GUI) for user interaction and performs various tasks using pre-existing IRAF packages, also calling on the NumArray module of Python to perform image manipulation and arithmetic. There is the option for some user control of the basic reduction properties, or the reduction may be performed automatically or on multiple files. Although the reduction relies on pre-defined parameters to simplify the process for the user and increase speed, it is relatively instrument independent and therefore suitable for adaptation to SALT HRS use.

6.4.2 Software Requirements

FIEStool requires the following software packages:

- python v2.3 or better
- python-dev
- python-tk
- python-Numeric
- f2c
- libplot-devel
- X and X-devel

A full IRAF distribution is also needed, and FIEStool was developed using version 2.12.2. Additionally, the following Python modules are required:

- Numarray v0.9
- PyFITS v0.9
- PyRAF v1.1.1
- Biggles 11.6.4

6.4.3 Input Images Required

In order to calculate a new set of reference data, the following input data is required:

- Several bias frames ($\geq$10 recommended)
- Multiple well exposed flat frames, obtained under similar conditions with a similar optical setup as the object frames ($\geq$ 10 recommended)
- One well exposed frame of a hot source, preferably an early type star, to be used for finding the order locations
- One well exposed wavelength calibration frame, preferably a ThAr spectrum, to be used for determining the wavelength solution

Additional sets of input files are required if the binning or detector read-out speeds are changed.

6.4.4 Calibration Frame Creation

Before processing science data the FIEStool calibration frames must be in place. The reference frames required for data reduction can be made within the FIEStool package itself in the 'CalibCalc' mode using the frames mentioned above in Section 6.4.3. Figure 6.5 shows the GUI interface for this mode along with instructions for its use. The tasks must be completed in the order outlined within the GUI as some tasks rely on previously generated data. Figure 6.6 is a flow chart for the 'CalibCalc' processes, explaining the actions performed and also specifying the

Python and/or IRAF tasks used to achieve this in *italic* text. The numbers labelled on buttons in Figure 6.5 correspond to the tasks described in Figure 6.6.

6.4.5 Data Reduction Processes

In processing the data, a variety of tasks are performed using the reference frames. A user is able to select between 'Quicklook', 'Advanced' and 'DoubleSpec' (interlaced spectra) modes dependent on his/her requirements (see Figure 6.7 for a screenshot of the main FIEStool user interface). The quicklook mode is optimized for speed of reduction: the science frame is pre-processed (header checks and heliocentric velocity corrections); bias levels are subtracted (with cosmic rays removed); the frame is flat-fielded; each order is extracted into a one-dimensional spectrum (with blaze shape corrected for); a wavelength definition is assigned; and the spectra are plotted automatically. In this simple mode the reduced data is not of science quality and is used only for an observer to make real-time decisions regarding the observations. The advanced mode is optimized for high data quality and adds additional tasks to the reduction process. Several of these tasks are computationally intensive and therefore the reduction process takes considerably longer than in quicklook mode. Before the extraction of spectra, the advanced mode also subtracts scattered light from a two-dimensional model in the inter-order spacing. At the extraction stage deviant pixels are also cleaned from the orders. And finally, after extracting orders, the spectra for each order are merged into one long spectrum. This requires re-binning on the wavelength scale and takes considerable time. The 'DoubleSpec' mode is designed for interlaced spectra of sky plus object or object plus simultaneous ThAr, and also differs from the quicklook reduction. In this case, scattered light is removed (as with the advanced mode) but the spectra of individual orders are not re-binned into one larger merged spectrum, as the re-binning process may introduce uncertainty in the wavelength calibration. Figure 6.8 shows a flow-chart for the data reduction tasks in the order in which they are performed. As is the case in Figure 6.6 in Section 6.4.4 for the calibration processes, the associated Python and/or IRAF tasks for each stage are again indicated in *italic* text. The mode(s) in which these tasks are performed is indicated by a 'Q' (Quicklook), 'A' (Advanced) or 'D' (DoubleSpec).

6.4.6 SALT HRS Trial Reduction

At the time of writing, FIEStool has not been fully adapted to reduce SALT HRS data, although trial reductions have been performed using data from the FIES instrument. Artificial SALT HRS Th-Ar and tungsten white light data have been created using the IRAF mkechelle task. The challenge is now to adapt the software to accept SALT HRS .fits file headers and to tune the IRAF task parameters to process SALT HRS data.

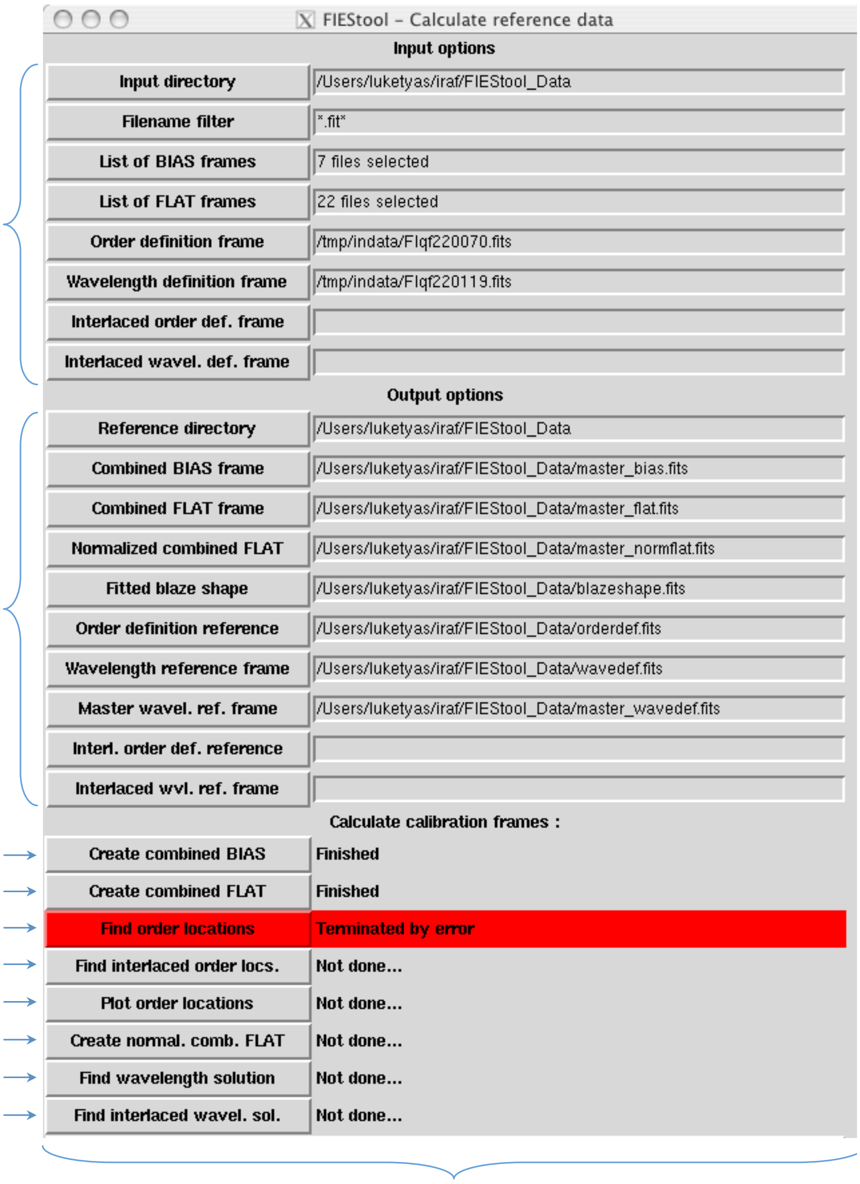

FIEStool – Calculate reference data

Click these buttons to set the path and/or filename for the input file(s). Multiple bias/flat frames are recommended for averaging into a combined bias/flat. The order definition frame can be a well-exposed frame of a hot source, such as an early-type star. The wavelength definition frame can be a well-exposed ThAr spectrum

Input options

Input directory	/Users/luketyas/iraf/FIEStool_Data
Filename filter	*.fit*
List of BIAS frames	7 files selected
List of FLAT frames	22 files selected
Order definition frame	/tmp/indata/FIqf220070.fits
Wavelength definition frame	/tmp/indata/FIqf220119.fits
Interlaced order def. frame	
Interlaced wavel. def. frame	

Click these buttons to select the path and filenames for the calibration/reference frames created in the processes below

Output options

Reference directory	/Users/luketyas/iraf/FIEStool_Data
Combined BIAS frame	/Users/luketyas/iraf/FIEStool_Data/master_bias.fits
Combined FLAT frame	/Users/luketyas/iraf/FIEStool_Data/master_flat.fits
Normalized combined FLAT	/Users/luketyas/iraf/FIEStool_Data/master_normflat.fits
Fitted blaze shape	/Users/luketyas/iraf/FIEStool_Data/blazeshape.fits
Order definition reference	/Users/luketyas/iraf/FIEStool_Data/orderdef.fits
Wavelength reference frame	/Users/luketyas/iraf/FIEStool_Data/wavedef.fits
Master wavel. ref. frame	/Users/luketyas/iraf/FIEStool_Data/master_wavedef.fits
Interl. order def. reference	
Interlaced wvl. ref. frame	

Calculate calibration frames :

1) Create Combined Bias →	Create combined BIAS	Finished
2) Create Combined Flat →	Create combined FLAT	Finished
3) Find Order Locations →	Find order locations	Terminated by error
4) Find Interlaced Orders →	Find interlaced order locs.	Not done...
5) Plot Order Locations →	Plot order locations	Not done...
6) Create Normalized Flat →	Create normal. comb. FLAT	Not done...
7) Find Wavelength Solution →	Find wavelength solution	Not done...
8) Find Interlaced Wavelength Solution →	Find interlaced wavel. sol.	Not done...

Frame calculation tasks are labelled 'Not done...' when incomplete, illuminated green whilst in progress, labelled 'Finished' when complete and turn red should they fail to execute correctly

Figure 6.5: A screenshot of the 'CalibCalc' window, used for generating reference frame data for following reduction processes. The numbered tasks correspond to the same tasks in the flow-chart of Figure 6.6.

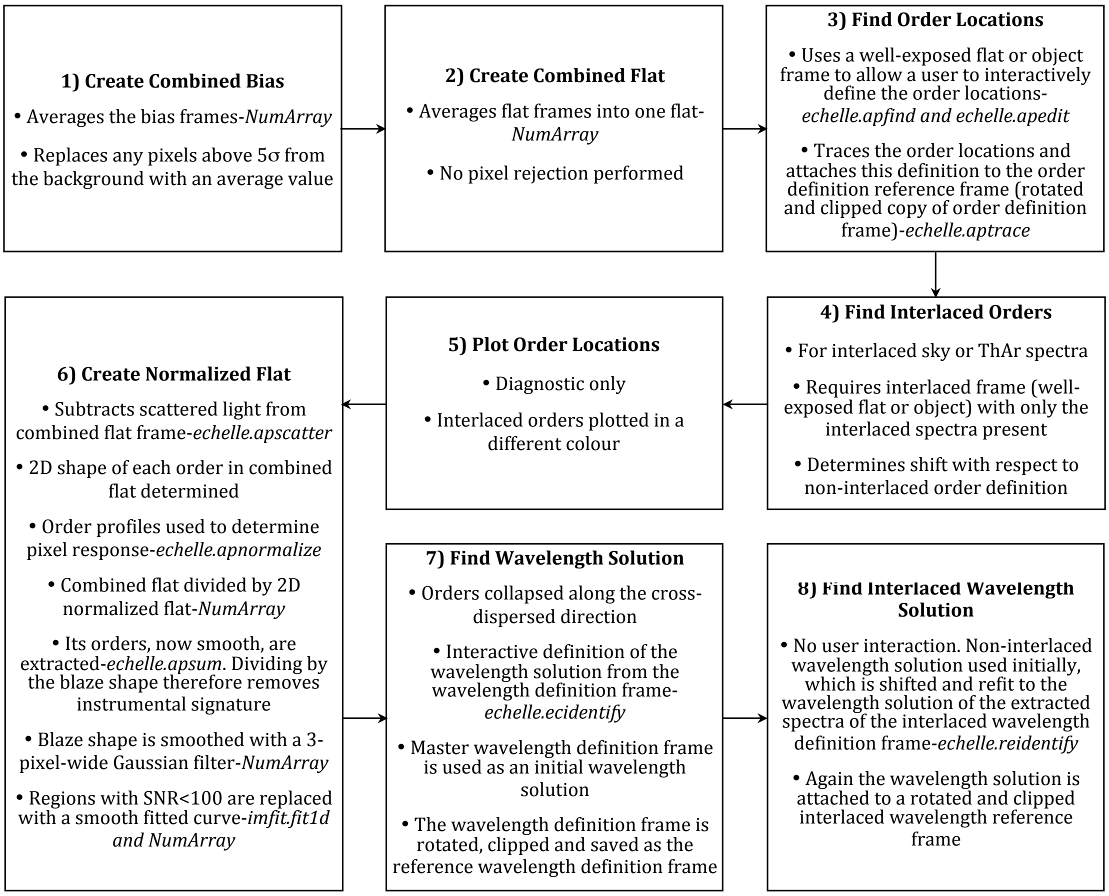

1) Create Combined Bias

- Averages the bias frames-*NumArray*
- Replaces any pixels above 5σ from the background with an average value

2) Create Combined Flat

- Averages flat frames into one flat-*NumArray*
- No pixel rejection performed

3) Find Order Locations

- Uses a well-exposed flat or object frame to allow a user to interactively define the order locations-*echelle.apfind and echelle.apedit*
- Traces the order locations and attaches this definition to the order definition reference frame (rotated and clipped copy of order definition frame)-*echelle.aptrace*

4) Find Interlaced Orders

- For interlaced sky or ThAr spectra
- Requires interlaced frame (well-exposed flat or object) with only the interlaced spectra present
- Determines shift with respect to non-interlaced order definition

5) Plot Order Locations

- Diagnostic only
- Interlaced orders plotted in a different colour

6) Create Normalized Flat

- Subtracts scattered light from combined flat frame-*echelle.apscatter*
- 2D shape of each order in combined flat determined
- Order profiles used to determine pixel response-*echelle.apnormalize*
- Combined flat divided by 2D normalized flat-*NumArray*
- Its orders, now smooth, are extracted-*echelle.apsum*. Dividing by the blaze shape therefore removes instrumental signature
- Blaze shape is smoothed with a 3-pixel-wide Gaussian filter-*NumArray*
- Regions with SNR<100 are replaced with a smooth fitted curve-*imfit.fit1d and NumArray*

7) Find Wavelength Solution

- Orders collapsed along the cross-dispersed direction
- Interactive definition of the wavelength solution from the wavelength definition frame-*echelle.ecidentify*
- Master wavelength definition frame is used as an initial wavelength solution
- The wavelength definition frame is rotated, clipped and saved as the reference wavelength definition frame

8) Find Interlaced Wavelength Solution

- No user interaction. Non-interlaced wavelength solution used initially, which is shifted and refit to the wavelength solution of the extracted spectra of the interlaced wavelength definition frame-*echelle.reidentify*
- Again the wavelength solution is attached to a rotated and clipped interlaced wavelength reference frame

Figure 6.6: A flow chart of the processes performed in creating new reference frames for FIEStool data reduction, with the Python and/or IRAF tasks required to perform these tasks in *italic* text.

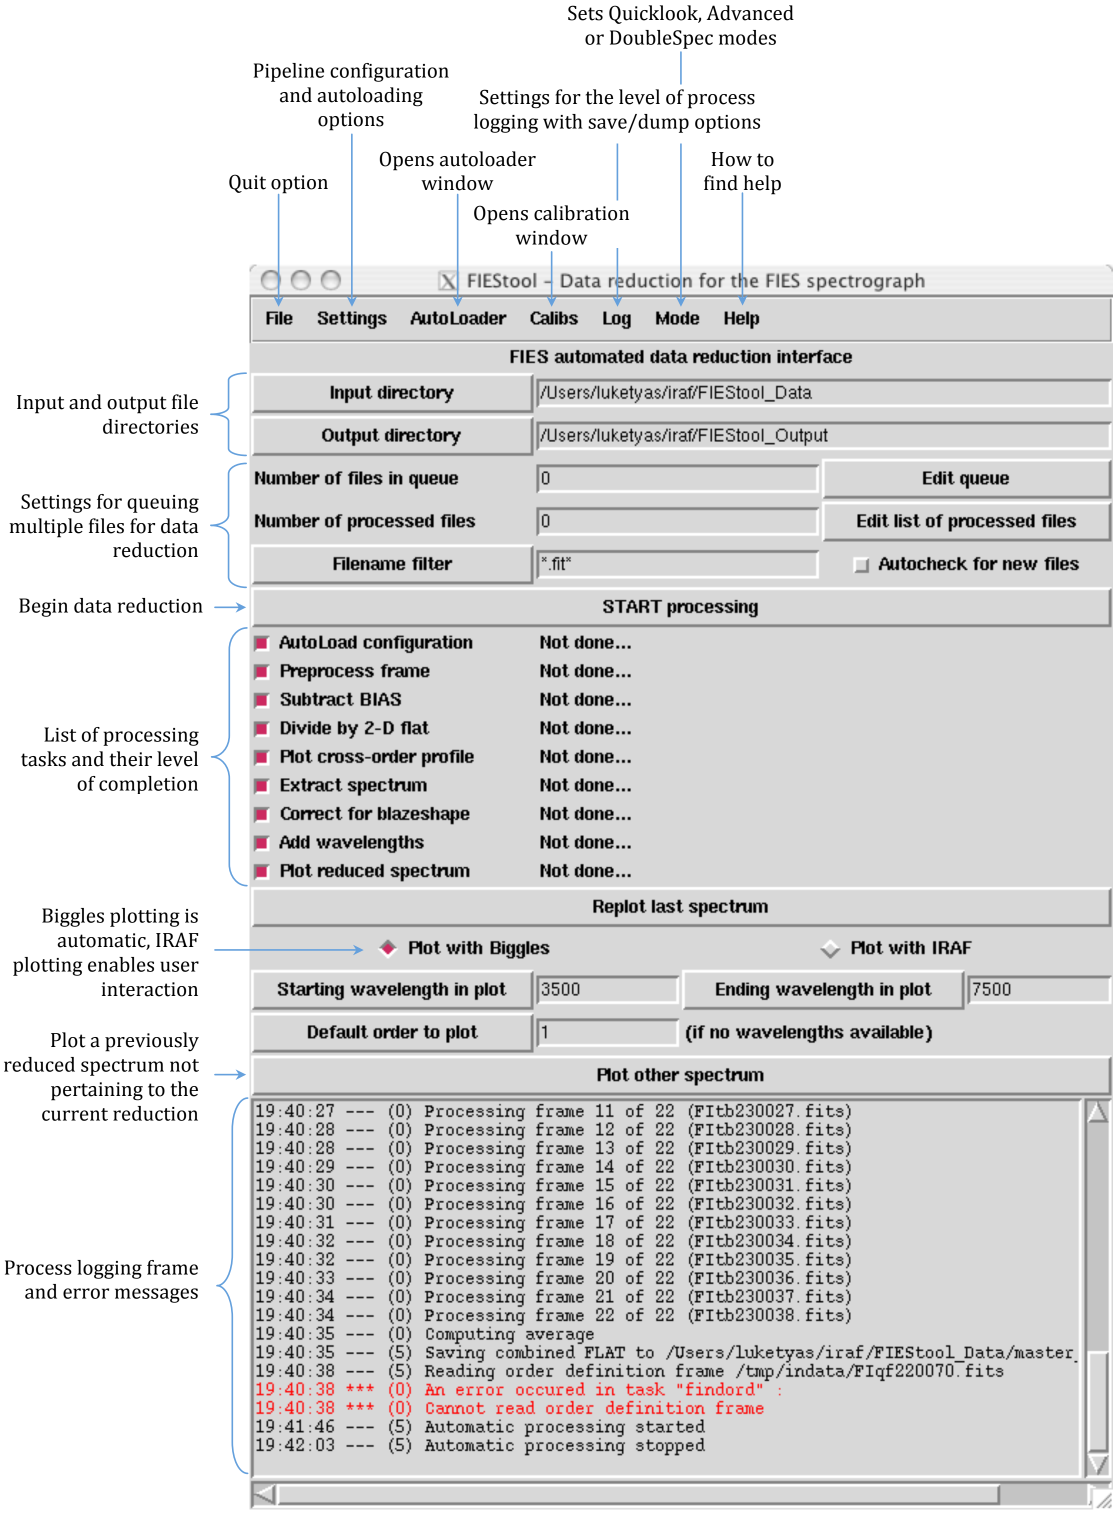

Sets Quicklook, Advanced or DoubleSpec modes

Pipeline configuration and autoloading options

Settings for the level of process logging with save/dump options

Quit option

Opens autoloader window

How to find help

Opens calibration window

Input and output file directories

Settings for queuing multiple files for data reduction

Begin data reduction

List of processing tasks and their level of completion

Biggles plotting is automatic, IRAF plotting enables user interaction

Plot a previously reduced spectrum not pertaining to the current reduction

Process logging frame and error messages

Input directory /Users/luketyas/iraf/FIEStool_Data
Output directory /Users/luketyas/iraf/FIEStool_Output
Number of files in queue 0
Number of processed files 0
Filename filter *.fit*

AutoLoad configuration Not done...
Preprocess frame Not done...
Subtract BIAS Not done...
Divide by 2-D flat Not done...
Plot cross-order profile Not done...
Extract spectrum Not done...
Correct for blazeshape Not done...
Add wavelengths Not done...
Plot reduced spectrum Not done...

Starting wavelength in plot 3500
Ending wavelength in plot 7500
Default order to plot 1 (if no wavelengths available)

```
19:40:27 --- (0) Processing frame 11 of 22 (FItb230027.fits)
19:40:28 --- (0) Processing frame 12 of 22 (FItb230028.fits)
19:40:28 --- (0) Processing frame 13 of 22 (FItb230029.fits)
19:40:29 --- (0) Processing frame 14 of 22 (FItb230030.fits)
19:40:30 --- (0) Processing frame 15 of 22 (FItb230031.fits)
19:40:30 --- (0) Processing frame 16 of 22 (FItb230032.fits)
19:40:31 --- (0) Processing frame 17 of 22 (FItb230033.fits)
19:40:32 --- (0) Processing frame 18 of 22 (FItb230034.fits)
19:40:32 --- (0) Processing frame 19 of 22 (FItb230035.fits)
19:40:33 --- (0) Processing frame 20 of 22 (FItb230036.fits)
19:40:34 --- (0) Processing frame 21 of 22 (FItb230037.fits)
19:40:34 --- (0) Processing frame 22 of 22 (FItb230038.fits)
19:40:35 --- (0) Computing average
19:40:35 --- (5) Saving combined FLAT to /Users/luketyas/iraf/FIEStool_Data/master_
19:40:38 --- (5) Reading order definition frame /tmp/indata/FIqf220070.fits
19:40:38 *** (0) An error occured in task "findord" :
19:40:38 *** (0) Cannot read order definition frame
19:41:46 --- (5) Automatic processing started
19:42:03 --- (5) Automatic processing stopped
```

Figure 6.7: A screenshot of the main FIEStool GUI for reducing échelle spectra, including reduction modes, queuing options, plotting options and process logging functionality.

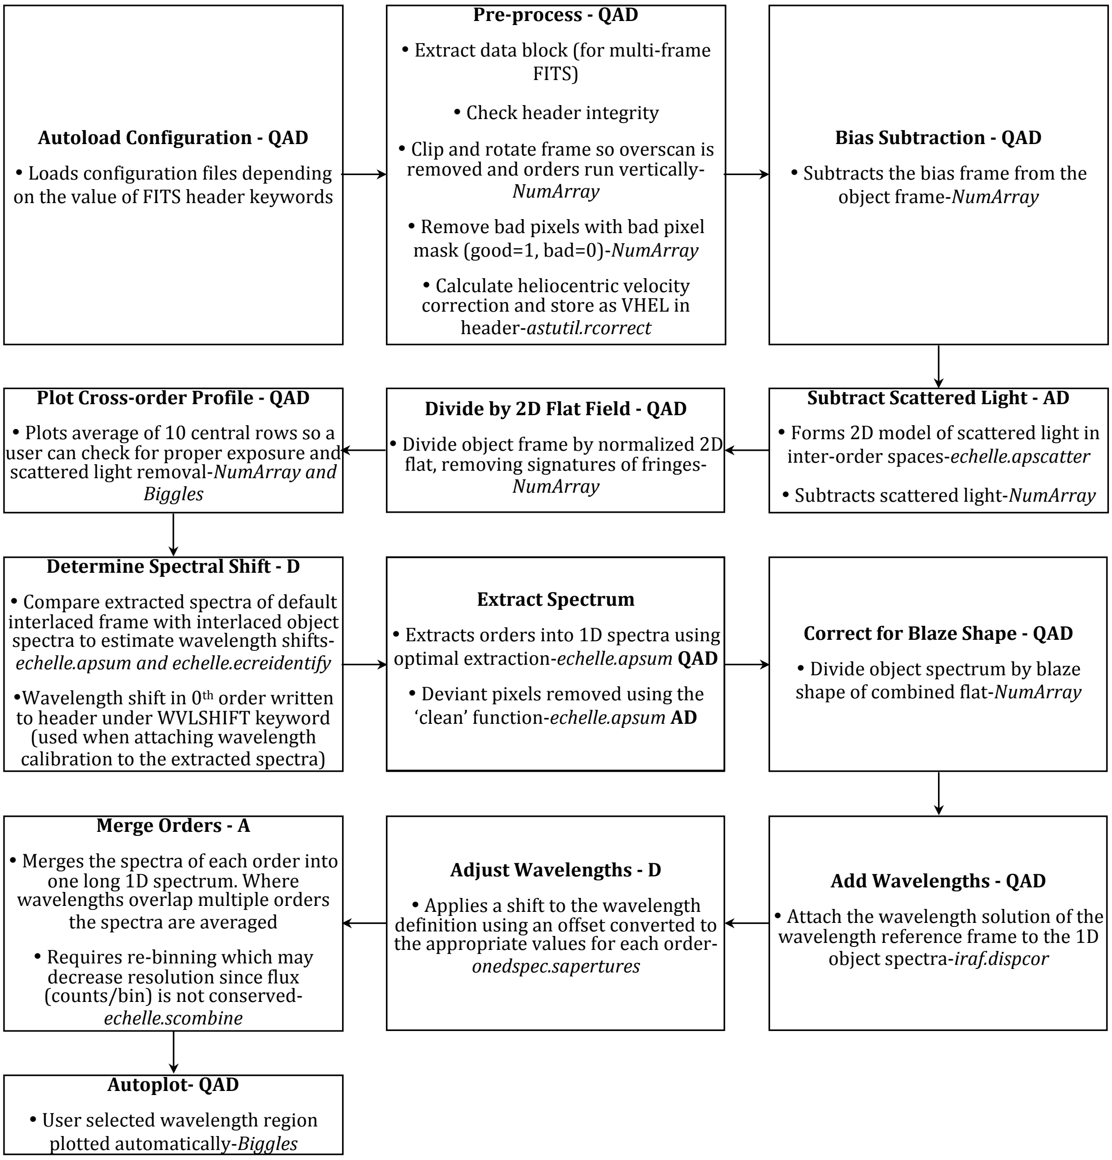

Figure 6.8: A flow chart of the processes performed in FIEStool data reduction, with the Python and/or IRAF tasks required to perform these tasks in *italic* text. The mode(s) in which these tasks are performed is indicated by a 'Q' (Quicklook), 'A' (Advanced) or 'D' (DoubleSpec).

Chapter 7

Summary & Conclusions

The material in this thesis describes the theoretical background to cross-dispersed échelle spectrographs, and introduces the SALT telescope and the specific SALT HRS design. Measurements of the efficiency and/or stability of key SALT HRS components are presented and used within a complete instrument throughput model to predict final performance. The blue channel SALT HRS detector has been used to quantify the effect of cosmic ray events on long exposures. Additionally, a potential quick-look data reduction pipeline is described and the instrument throughput information has been used to develop a GUI-based exposure time calculator for observer use.

The SALT HRS high-resolution échelle spectrograph was described in Chapter 1. The instrument was seen to compare favourably with similar instruments in terms of wavelength coverage, observational capabilities offered, spectral resolution and stability. The dual spectral channel design allows for a large number of components to be optimized for a narrower wavelength range of operation, limits the size of the expensive cross-diserpser and camera optics required and removes the need for CCD detector mosaics. With its four key modes and other optional features, the SALT HRS design is shown to be a powerful and highly flexible tool for the astronomical community. A user is able to select a mode depending on his/her specific need for efficiency or spectral resolution. The high-stability mode, added to the baseline design by CfAI, also adds the potential for high velocity precision measurements, useful particularly in the emerging and exciting field of extra-solar planet detection. High stability is facilitated by the double scrambler in the fibre feed of this mode (see Chapter 3), which acts to ensure nearly complete scrambling of the variable pupil illumination in SALT, thereby removing memory of the input beam and increasing spectral stability at the expense of throughput loss. The optional simultaneous Th-Ar mode and iodine absorption cell (see Chapter 6) in this mode allow for increased precision in wavelength calibration of the spectra. Time-indexed flux data during exposures also allow for precise heliocentric velocity correction when using the exposure meter. A unique nod-and-shuffle capability available in the low-resolution mode also allows for highly accurate sky background subtraction for work on very faint targets (see Chapter 3).

In Chapter 3, the focal ratio degradation (FRD) losses were comprehensively measured for Polymicro FBP fibres used in SALT HRS. The results as a function of input beam speed and wavelength were as expected: there was no discernible FRD wavelength dependence and the need to feed the fibre input with a fast f-ratio input beam to minimize FRD loss was confirmed. Results suggested a loss of $\leq$5% in the conditions applicable to SALT HRS. Encouragingly, the increase in FRD losses as a result of fibre bending induced stress appeared limited to a radius of curvature well below that allowed by the protective fibre bundle on the SALT HRS fibres. Furthermore, the effect of the change in illumination of the SALT primary mirror as an object is tracked seemed to have increased FRD losses by no more than 3%. Radial stress at the fibre feedthrough caused by the change to vacuum conditions also appeared to have no measurable effect on FRD losses. Since the fibres will eventually be bundled in the protective fibre conduit, further measurements may be required to ensure any strain caused by the increased weight on the fibres is not causing FRD to increase. Other details concerning the operation of optical fibres, double scrambler design, SALT HRS fibre modes, the fibre instrument feed (FIF), and the fibre bundle assembly were also presented within Chapter 3.

The use of volume-phase holographic gratings as cross-dispersion elements is quite pioneering in that they are a relatively new concept in astronomical instrumentation and feature in only a few current instrument designs. However, literature would suggest their use in astronomical spectroscopy will increase with next-generation instruments. Chapter 4 gives details of the cross-dispersion requirements and describes the advantages of using VPHGs in SALT HRS. The results of performance measurements presented in Chapter 4 indicate that the SALT HRS VPHGs may be operating some 4-15% below their theoretical efficiency, particularly in the extreme blue, but still meet specification at most wavelengths.

Chapter 5 characterizes the blue CCD detector system (at the time of writing the red detector system has still not been delivered from the supplier, ARC Inc.). The effect of contaminant material left on a few small regions of the detector was measured before and after cleaning using flat-field exposures. Following this the measurement of a photon transfer curve using the same flat-field illumination method allowed for calculation of the camera gain constant as $K = 2.90\mathrm{e}^-$ using 440nm illumination and a 400kHz readout speed. At [illegible] same readout speed, and using the value of K from the photon transfer curve, a blue detector quantum efficiency curve with 5nm wavelength resolution was produced and calibrated with a photodiode of known response. The detector was shown to exhibit excellent quantum efficiency performance in line with its specification. The non-availability of the red detector prevented similar characterization.

The blue detector was also used to investigate cosmic ray events during long exposures. The event rate has been shown to be consistent over time and can be reduced by ~50% using lead shielding on all sides of the detector. The use of 76mm thick shielding over 51mm shielding is shown to be minimal, so long as shielding from all directions is provided. The different morphology of ground-based events and true cosmic ray muons was described in an attempt to characterize the events observed according to their geometry. The availability of the red CCD in the future may allow for comparison of the event geometry in a thicker detector, which is expected to better highlight the different geometry exhibited by muon 'streaks' from primary cosmic rays and ground-based 'worms' from local radioactive decay products.

A feature of the SALT HRS design is the limited number of moving parts and its operation inside a vacuum vessel within thermally stabilized conditions, both of which allow for a stable optical path essential for high wavelength stability. The effect of motion, pressure change and thermal drift are therefore mitigated against. However the optical components must still be mounted in a stable manner. This can be difficult to achieve when considering the weight of some components, as well as the need for precise adjustability in up to six degrees of freedom. The mount must also not over-constrain the components in case of expansion/contraction under temperature change. Therefore the mirror mounts in SALT HRS have been tested to ensure they achieve the required level of stability. In the case of the pupil mirror mounts, the stability required is 0.5mrad, which corresponds to beam movement of 1.0mm at a distance of 2m. Tests described in Chapter 4 confirm beam motions of $\leq$0.7mm over a 50 hour period in the laboratory. Given the limitations of the measurement method and the larger temperature changes experienced during the test compared with those inside the vacuum tank, the mounts are expected to achieve much better stability than this in operation in the final instrument configuration.

The results of the subsystem tests presented in this thesis were included in a complete instrument performance model in Chapter 2. For components which have not been tested in this thesis, predicted throughput values were calculated using manufacturer witness sample or test data. The throughput model proved invaluable when considering the tuning of the volume-phase holographic gratings to enable complete wavelength coverage, the order coverage at the dichroic cross-over, and the tuning of the blue grating to the centre of its wavelength range, as opposed to increasing performance in the extreme blue wavelengths.

The work presented in this thesis has been concerned not only with the performance and construction of SALT HRS but also with its operation following delivery and integration. Chapters 2 and 6 both describe software tools useful for the operation of the instrument at the telescope. The accessible GUI-based exposure time calculator is useful to astronomers for planning observations given the expected mode of operation and conditions, and can be launched within a web browser. The calculator relies on data from the throughput model for telescope and instrument efficiency, as well as other parameters dependent on observing conditions. Data reduction should be simplified and rapid using the quick-look pipeline presented in Chapter 6, which will prove invaluable to a SALT astronomer who requires near-real-time feedback on spectra measured by the instrument.

Given measurements obtained thus far, SALT HRS is expected to meet its on-sky performance requirements and should prove an important addition to southern hemisphere capability for high-resolution spectrographic observations. Its wide applicability and flexibility, coupled with the high light collecting power of the SALT telescope, promise[illegible]llow it to meet a large and diverse set of science objectives.

Appendix

This appendix contains a copy of the exposure time calculator Java source code, as described in Section 2.4.5 of Chapter 2.

```
class App_Without_GUI {
public static void main(String[] args) {

// ##################################################################################
// Define user set parameters in variables and arrays:
// ##################################################################################

// Observation and atmosphere variables
//float vmag, ebv, obstime, seeing, airmass;
//double temp, atmpres, humid;

// Select values from user data in GUI text fields
float vmag = 12.0f; // Magnitude in V
float ebv = 0.0f; // Reddening E(B-V)
float obstime = 3600f;         // Exposure time (s)
float seeing = 1.0f; // Seeing (arcsec)
float airmass = 1.5f; // Airmass
double temp = 15.0d; // Temperature (C)
double atmpres = 773.0d; // Atmospheric pressure (hPa)
double humid = 25.0d; // Relative humidity (%)
String spec = "O";            // Star type (OBAFGKM)
String mode = "Low Resolution"; // Spectrograph operation mode
String mirrorArea = "Minimum 38m^2"; // Mirror diameter (cm)
String iodineCell = "No"; // Presence of I2 cell
String moon = "Dark";       // Moon brightness

// Detector values
double RedRON = 3.4f; // Read out noise (e-)
RedRON = Math.sqrt(RedRON);
double BlueRON = 4.1f; // Read out noise (e-)
BlueRON = Math.sqrt(BlueRON);
float FWCRed = 350000.0f; // Full well capacity (e-)
float FWCBlue = 600000.0f; // Full well capacity (e-)
double ccdTemp = 153d; // CCD detector temperature (K)

// Telescope efficiency
float[] teleff = {0.603f, 0.591f, 0.603f, 0.590f, 0.613f, 0.615f, 0.606f, 0.585f, 0.602f, 0.612f};

// UBVRI is turned in to Blue: 370-415-460-505-550 and Red: 560-645-725-810-890nm
```

```
// ##################################################################################
// Define pre-set parameters:
// ##################################################################################

// Intrinsic magnitude
float[] imag = {0.0f, 0.0f, 0.0f, 0.0f, 0.0f, 0.0f, 0.0f, 0.0f, 0.0f, 0.0f};
// Interstellar reddened magnitude
float[] rmag = {0.0f, 0.0f, 0.0f, 0.0f, 0.0f, 0.0f, 0.0f, 0.0f, 0.0f, 0.0f};
// Atmosphere reddened magnitude
float[] emag = {0.0f, 0.0f, 0.0f, 0.0f, 0.0f, 0.0f, 0.0f, 0.0f, 0.0f, 0.0f};
double[] fluxBlue = {0.0d, 0.0d, 0.0d, 0.0d, 0.0d}; // At Earth
double[] fluxRed = {0.0d, 0.0d, 0.0d, 0.0d, 0.0d}; // At Earth

// Interstellar extinction
float[] intext = {4.66f, 4.37f, 4.08f, 3.80f, 3.51f, 3.44f, 2.90f, 2.39f, 1.84f, 1.33f};
// (mag/E(B-V)
float[] atmext = {0.443f, 0.318f, 0.227f, 0.165f, 0.125f, 0.118f, 0.085f, 0.065f, 0.016f, 0.001f};


// ##################################################################################
// Intrinsic colours of stars in bands:
// ##################################################################################

// Defined in the OBAFGKM order within the array. Used to calculate the magnitude of an
// OBAFGKM object at the ten different wavelength bins

// e.g. to call an A class object, use element 2

float[] mag_505_550 = {-0.09f, -0.04f, 0.07f, 0.17f, 0.26f, 0.69f, 0.60f};
float[] mag_460_505 = {-0.11f, -0.04f, 0.08f, 0.20f, 0.27f, 0.12f, 0.54f};
float[] mag_415_460 = {-0.20f, -0.08f, 0.04f, 0.25f, 0.47f, 1.14f, 1.44f};
float[] mag_370_415 = {-1.45f, -0.60f, -0.11f, -0.65f, -0.66f, -0.10f, 0.66f};

float[] mag_550_560 = {0.01f, -0.03f, -0.01f, -0.02f, -0.01f, 0.00f, -0.07f};
float[] mag_560_645 = {0.16f, 0.08f, -0.02f, -0.21f, -0.31f, -0.67f, -1.26f};
float[] mag_645_725 = {0.09f, 0.01f, -0.03f, -0.12f, -0.23f, -0.36f, -0.92f};
float[] mag_725_810 = {0.09f, -0.04f, -0.12f, -0.20f, -0.21f, -0.34f, -1.27f};
float[] mag_810_890 = {0.11f, 0.02f, -0.06f, -0.07f, -0.09f, -0.17f, -0.23f};


// ##################################################################################
// Determine fibre attenuation:
// ##################################################################################

double[] fibreLength = {0.035d, 0.035d, 0.035d, 0.035d, 0.035d, 0.035d, 0.035d, 0.035d, 0.035d, 0.035d};
double[] attenuation = {70.3d, 46.8d, 21.7d, 15.5d, 11.5d, 10.8d, 6.95f, 4.95d, 2.5d, 2.83d};
double[] fibreEfficiency = {0.0d, 0.0d, 0.0d, 0.0d, 0.0d, 0.0d, 0.0d, 0.0d, 0.0d, 0.0d};

for (int i=0; i<10; i++) {
fibreEfficiency[i] = 1.0d / Math.pow(10.0d, ((attenuation[i] * fibreLength[i]) / 10.0d));
}


// ##################################################################################
// Instrument efficiencies:
// ##################################################################################

// LR mode blue instrument efficiency
float insteffLowResBlue[] = {0.094f, 0.249f, 0.315f, 0.284f, 0.127f};
// LR mode red instrument efficiency
float insteffLowResRed[] = {0.127f, 0.324f, 0.340f, 0.250f, 0.141f};
```

```
// MR mode blue instrument efficiency
float insteffMedResBlue[] = {0.067f, 0.174f, 0.219f, 0.198f, 0.088f};
// MR mode red instrument efficiency
float insteffMedResRed[] = {0.088f, 0.226f, 0.235f, 0.150f, 0.099f};
// HR mode blue instrument efficiency
float insteffHighResBlue[] = {0.070f, 0.150f, 0.232f, 0.209f, 0.093f};
// HR mode red instrument efficiency
float insteffHighResRed[] = {0.093f, 0.238f, 0.250f, 0.200f, 0.076f};
// HS mode blue instrument efficiency
float insteffHighStabBlue[] = {0.051f, 0.134f, 0.168f, 0.153f, 0.068f};
// HS mode red instrument efficiency
float insteffHighStabRed[] = {0.068f, 0.173f, 0.183f, 0.146f, 0.076f};


// ###################################################################################
// Determine spectrograph opertation mode and set-up:
// ###################################################################################

float FibreRadius = 0.0f; // Half fibre angle (arcsec)
float nfibre = 2.0f; // Number of fibres in use
float npixel = 13.0f; // Number of pixels in the cross order direction
// Blue arm overall instrument efficiency (fibres, scrambler, I2 cell, instrument, CCDs)
float insteffBlue[] = {0.0f, 0.0f, 0.0f, 0.0f, 0.0f};
// Red arm overall instrument efficiency (fibres, scrambler, I2 cell, instrument, CCDs
float insteffRed[] = {0.0f, 0.0f, 0.0f, 0.0f, 0.0f};
float iodineCellEff[] = {0.851f, 0.851f, 0.851f, 0.851f, 0.851f, 0.851f, 0.851f, 0.851f, 0.851f, 0.851f};

for (int i=0; i<5; i++) {
if (mode.equals("Low Resolution")) {
nfibre = 2.0f;
npixel = 13.0f;
FibreRadius = 1.115f;
insteffBlue[i] = (float)fibreEfficiency[i] * insteffLowResBlue[i];
insteffRed[i] = (float)fibreEfficiency[i+5] * insteffLowResRed[i];
} else if (mode.equals("Medium Resolution")) {
nfibre = 2.0f;
npixel = 13.0f;
FibreRadius = 1.115f;
insteffBlue[i] = (float)fibreEfficiency[i] * insteffMedResBlue[i];
insteffRed[i] = (float)fibreEfficiency[i+5] * insteffMedResRed[i];
} else if (mode.equals("High Resolution")) {
nfibre = 2.0f;
npixel = 13.0f;
FibreRadius = 0.780f;
insteffBlue[i] = (float)fibreEfficiency[i] * insteffHighResBlue[i];
insteffRed[i] = (float)fibreEfficiency[i+5] * insteffHighResRed[i];
} else if (mode.equals("High Stability") && iodineCell.equals("No")) {
nfibre = 2.0f;
npixel = 13.0f;
FibreRadius = 0.780f;
insteffBlue[i] = (float)fibreEfficiency[i] * insteffHighStabBlue[i];
insteffRed[i] = (float)fibreEfficiency[i+5] * insteffHighStabRed[i];
} else if (mode.equals("High Stability") && iodineCell.equals("Yes")) {
nfibre = 2.0f;
npixel = 13.0f;
FibreRadius = 0.780f;
insteffBlue[i] = (float)fibreEfficiency[i] * insteffHighStabBlue[i] * iodineCellEff[i];
insteffRed[i] = (float)fibreEfficiency[i+5] * insteffHighStabRed[i] * iodineCellEff[i+5];
}
}
```

```
// ####################################################################################
// Correct the intrinsic colours:
// ####################################################################################

// If statement finds corresponding element to spectral type, e.g. if spectral type
// is A, element for bv etc. becomes 2

int x = 0;

if (spec.equals("O")) {
x = 0;
} else if (spec.equals("B")) {
x = 1;
} else if (spec.equals("A")) {
x = 2;
} else if (spec.equals("F")) {
x = 3;
} else if (spec.equals("G")) {
x = 4;
} else if (spec.equals("K")) {
x = 5;
} else if (spec.equals("M")) {
x = 6;
}

// Now correct colours, using correct filters and spectral types:

imag[4] = vmag;
imag[3] = vmag + mag_505_550[x];
imag[2] = imag[3] + mag_460_505[x];
imag[1] = imag[2] + mag_415_460[x];
imag[0] = imag[1] + mag_370_415[x];

imag[5] = vmag + mag_550_560[x];
imag[6] = imag[5] + mag_560_645[x];
imag[7] = imag[6] + mag_645_725[x];
imag[8] = imag[7] + mag_725_810[x];
imag[9] = imag[8] + mag_810_890[x];


// ####################################################################################
// Correct observed V magnitude to intrinsic magnitude:
// ####################################################################################

float off = 0.0f;

off = intext[4] * ebv;
// intext[4] is essentially 3.1
// extinction R(V) over the average line of sight in the Milky Way

// ####################################################################################
// Calculate mirror area:
// ####################################################################################

float telArea = 0.0f;

if (mirrorArea.equals("Minimum 38m^2")) {
telArea = 380000.0f;
} else if (mirrorArea.equals("Average 46m^2")) {
telArea = 460000.0f;
} else if (mirrorArea.equals("Maximum 55m^2")) {
telArea = 550000.0f;
```

```
}


// ##################################################################################
// CCD Dark Current/Full Well Capacity/Read-out Noise Calculations:
// ##################################################################################

// ##################################################################################
// Red detector:

double QdRed = 0.0d;

// Dark current in e- per pixel per hour
QdRed = 55336103d * 122d * ccdTemp * ccdTemp * ccdTemp * Math.exp(-6400 / ccdTemp);

double DarkRed = 0.0d;
DarkRed = (QdRed / 3600.0d) * npixel * obstime; // Recorded dark current


// ##################################################################################
// Blue detector:

double QdBlue = 0.0d;

// Dark current in e- per pixel per hour
QdBlue = 33584001d * 122d * ccdTemp * ccdTemp * ccdTemp * Math.exp(-6400 / ccdTemp);

double DarkBlue = 0.0d;
DarkBlue = (QdBlue / 3600.0d) * npixel * obstime; // Recorded dark current


// ##################################################################################
// Correct for moon brightness:
// ##################################################################################

float[] skymag = {0.0f, 0.0f, 0.0f, 0.0f, 0.0f, 0.0f, 0.0f, 0.0f, 0.0f, 0.0f};

float[] MP_NewMoon = {22.11f, 22.69f, 22.76f, 22.52f, 22.13f, 22.04f, 21.28f, 20.69f, 19.85f, 18.03f};
// Sky at new moon
float[] MP_HalfMoon = {20.04f, 21.02f, 21.52f, 21.66f, 21.58f, 21.53f, 20.99f, 20.37f, 19.66f, 18.79f};
// Sky at half moon
float[] MP_FullMoon = {17.19f, 18.53f, 19.32f, 19.75f, 19.94f, 19.97f, 19.99f, 19.81f, 19.15f, 17.45f};
// Sky at full moon

for (int i=0; i<10; i++) {
if (moon.equals("Dark")) {
skymag[i] = MP_NewMoon[i];
} else if (moon.equals("Grey")) {
skymag[i] = MP_HalfMoon[i];
} else if (moon.equals("Bright")) {
skymag[i] = MP_FullMoon[i];
}
}


// ##################################################################################
// Calibrate flux from sky:
// ##################################################################################

// For a star of magnitude 0.0 in photons: mag -> photons/cm^2/s/A

float[] fluxcal = {551.0f, 1638.0f, 1426.0f, 1177.0f, 1005.0f, 956.0f, 740.0f, 588.0f, 449.0f, 390.0f};
```

```
// From HST absolute spectrophotometry of Vega
// (corroborated with Pickles stellar flux library for A0V star)

double[] skyBlue = {0.0d, 0.0d, 0.0d, 0.0d, 0.0d};
double[] skyRed = {0.0d, 0.0d, 0.0d, 0.0d, 0.0d};

for (int i=0; i<5; i++) {
skyBlue[i] = Math.exp(-0.92104d * skymag[i]) * fluxcal[i];
skyRed[i] = Math.exp(-0.92104d * skymag[i+5]) * fluxcal[i+5];
}


// ##################################################################################
// Multiply by fibre spatial angle (arcsec^2):

float FibreOmega = 0.0f;

FibreOmega = 3.141592654f * FibreRadius * FibreRadius;

for (int i=0; i<5; i++) {
skyBlue[i] = skyBlue[i] * FibreOmega;
skyRed[i] = skyRed[i] * FibreOmega;
}


// ##################################################################################
// Multiply by telescope area (and efficiency) and observation time:

// Units reduced from photons/cm^2/s/A to photons flux

for (int i=0; i<5; i++) {
  skyBlue[i] = skyBlue[i] * obstime * telArea * teleff[i];
skyRed[i] = skyRed[i] * obstime * telArea * teleff[i+5];
}


// ##################################################################################
// Multiply by bin size (A/pixel):

float binsize[] = {0.024f, 0.026f, 0.029f, 0.032f, 0.035f, 0.038f, 0.043f, 0.049f, 0.056f, 0.060f};

for (int i=0; i<5; i++) {
skyBlue[i] = skyBlue[i] * binsize[i];
skyRed[i] = skyRed[i] * binsize[i+5];
}


// ##################################################################################
// Multiply by instrument efficiency (which includes fibre and detector efficiencies):

for (int i=0; i<5; i++) {
    skyBlue[i] = skyBlue[i] * insteffBlue[i];
    skyRed[i] = skyRed[i] * insteffRed[i];
}


// ##################################################################################
// Correct for atmospheric and interstellar extinction:
// ##################################################################################

for (int i=0; i<10; i++) {
rmag[i] = imag[i] + intext[i] * ebv - off;
```

```
}

for (int i=0; i<10; i++) {
emag[i] = rmag[i] + atmext[i] * airmass;
}


// ################################################################################
// Calibrate flux from object:
// ################################################################################

for (int i=0; i<5; i++) {
fluxBlue[i] = Math.exp(-0.92104d * emag[i]) * fluxcal[i];
fluxRed[i] = Math.exp(-0.92104d * emag[i+5]) * fluxcal[i+5];
}


// ################################################################################
// Multiply by telescope area (and efficiency) and observation time:

for (int i=0; i<5; i++) {
    fluxBlue[i] = fluxBlue[i] * obstime * telArea * teleff[i];
fluxRed[i] = fluxRed[i] * obstime * telArea * teleff[i+5];
}


// ################################################################################
// Multiply by bin size:

for (int i=0; i<5; i++) {
    fluxBlue[i] = fluxBlue[i] * binsize[i];
fluxRed[i] = fluxRed[i] * binsize[i+5];
}


// ################################################################################
// Multiply by instrument efficiency (which includes fibre and detector efficiencies):

for (int i=0; i<5; i++) {
fluxBlue[i] = fluxBlue[i] * insteffBlue[i];
fluxRed[i] = fluxRed[i] * insteffRed[i];
}


// ################################################################################
// Multiply by seeing efficiency:

// First calculate seeing efficiency:

double seeingeff = 0.0d;

seeingeff = 1.0d - Math.exp((-FibreRadius * FibreRadius)
/ (2.0d * (seeing / 2.345d) * (seeing / 2.345d)));

for (int i=0; i<5; i++) {
fluxBlue[i] = fluxBlue[i] * seeingeff;
fluxRed[i] = fluxRed[i] * seeingeff;
}


// ################################################################################
// Signal to Noise Calculations:
```

```
// ##################################################################################

// Arrays to be used defined:

double PeakPixelBlue[] = {0.0d, 0.0d, 0.0d, 0.0d, 0.0d};
double PeakPixelRed[] = {0.0d, 0.0d, 0.0d, 0.0d, 0.0d};

double noisesqrBlue[] = {0.0d, 0.0d, 0.0d, 0.0d, 0.0d};
double noisesqrRed[] = {0.0d, 0.0d, 0.0d, 0.0d, 0.0d};

double snrBlue[] = {0.0d, 0.0d, 0.0d, 0.0d, 0.0d};
double snrRed[] = {0.0d, 0.0d, 0.0d, 0.0d, 0.0d};

// Fraction of flux in most saturated pixel:

for (int i=0; i<5; i++) {
PeakPixelBlue[i] = (fluxBlue[i] + skyBlue[i]) * 0.23f + BlueRON + (DarkBlue/npixel) * obstime;
PeakPixelRed[i] = (fluxRed[i] + skyRed[i]) * 0.23f + RedRON + (DarkRed/npixel) * obstime;
}

// Noise squared:

for (int i=0; i<5; i++) {
noisesqrBlue[i] = fluxBlue[i] + (2.0f-nfibre) * skyBlue[i] + nfibre * npixel *
(BlueRON * BlueRON + (DarkBlue/npixel) * obstime);
noisesqrRed[i] = fluxRed[i] + (2.0f-nfibre) * skyRed[i] + nfibre * npixel *
(RedRON * RedRON + (DarkRed/npixel) * obstime);
}

// Signal to Noise Ratio:

for (int i=0; i<5; i++) {
snrBlue[i] = fluxBlue[i] / Math.sqrt(noisesqrBlue[i]);
snrRed[i] = fluxRed[i] / Math.sqrt(noisesqrRed[i]);
}


// ##################################################################################
// Atmospheric Refraction Calculations:
// ##################################################################################

double lambda[] = {3700.0d, 8900.0d}; // Wavelength range
double R[] = {0.0d, 0.0d};

for (int i=0; i<2; i++) {
double z0 = Math.atan(Math.sqrt(airmass*airmass-1));
double sig = 10000.0d / lambda[i];
double T = 273.15d + temp; // Temperat ure in Kelvin from Centigrade user input
double x1 = Math.log((humid + 0.01d)/100.0d);  // Log base 10 of humidity (as a fraction of 1)
double td = 238.3d * ((temp+238.3d) * x1 + 17.2694 * temp)
/ ((temp+238.3d) * (17.2694d-x1) - 17.2694d * temp);
double pw = 4.50874d + 0.341724d * td + 0.0106778d * (td*td) +
0.000184889d * (td*td*td) + 0.00000238294 * (td*td*td*td) + 0.0000000203447 * (td*td*td*td*td);
double Pw = 1.333224d * pw;
double Ps = atmpres - Pw;
double beta = 0.001254 * T / 273.15d;
double kappa = 1.0d; // For a spherical Earth

double Ds = (1.0d + Ps * (0.000000579d - 0.00093250d/T + 0.25844d/(T*T))) * Ps/T;
double Dw = (1.0d + Pw * (1.0d + 0.00037*Pw) * (-0.00237321 + 2.23366d/T -
710.792d/(T*T) + 77514.1d/(T*T*T))) * Pw/T;
double n0 = 1.0d + 0.00000001d *((2371.34d + 683939.7d / (130.0d - (sig*sig)) +
```

```
4547.3d / (38.9d-(sig*sig))) * Ds + (6487.31d + 58.058d *(sig*sig) - 0.71150d * (sig*sig*sig*sig) +
0.08851d * (sig*sig*sig*sig*sig*sig)) * Dw);
double gamma = n0 - 1.0d;

R[i] = kappa * gamma * (1.0d-beta) * Math.tan(z0) - kappa * gamma * (beta-gamma/2.0d) *
((Math.tan(z0))*(Math.tan(z0))*(Math.tan(z0)));
R[i] = 180.0d / 3.14159d * 3600.0d * R[i];
}

double deltaR = R[0] - R[1];

System.out.println("Without atmospheric dispersion correction, the spectral length between
3700A and 8900A would be " + deltaR + " arcseconds");


// ####################################################################################
// Over-exposure Warning:
// ####################################################################################

int IsOverExposed = 0;

for (int i=0; i<5; i++) {
if (PeakPixelBlue[i] > FWCBlue) {
IsOverExposed = 1;
}
else if (PeakPixelRed[i] > FWCRed) {
IsOverExposed = 1;
}
}

if (IsOverExposed == 1) {
System.out.println("Warning: The CCD is overexposed!");
}

if (IsOverExposed != 1) {
System.out.println("The CCD is not overexposed");
}


// ####################################################################################
// Cosmic Ray Warning:
// ####################################################################################

float hitPixelFraction = 0.0f; // Fraction of pixels hit by cosmic rays
float hitBinFraction = 0.0f; // Fraction of bins hit by cosmic rays
float CosmicHits = 1.0f; // Cosmic radiation hits (cm^-2s^-1)
float pixelSize = 0.0015f; // Pixel size (cm)
float CCDArea = (pixelSize * 4096.0f) * (pixelSize * 4112.0f); // Detector area (cm)
float pixelNumber = 4096.0f * 4112.0f; // Number of pixels

hitPixelFraction = ((CosmicHits * CCDArea * obstime) / pixelNumber) * 100.0f;

if (hitPixelFraction > 100.0f) {
hitPixelFraction = 100.0f;
}

hitBinFraction = hitPixelFraction * npixel;

if (hitBinFraction > 100.0f) {
hitBinFraction = 100.0f;
}
```

```
if (hitBinFraction >= 10.0f) {
System.out.println("Warning: "+ hitPixelFraction + "% of CCD pixels are struck by cosmic rays.
This represents" + hitBinFraction + "% of your wavelength bins, and means optimum extraction
algorithms may experience difficulty masking these hit pixels. If possible, use a shorter exposure time");
}
else if (hitBinFraction < 10.0f) {
System.out.println("Using this exposure time, there are no issues with cosmic ray hits");
}

// ##################################################################################
// Round values for display:
// ##################################################################################


for (int i=0; i<5; i++) {
insteffBlue[i] = (float)Math.round(insteffBlue[i] * 1000) / 1000;
insteffRed[i] = (float)Math.round(insteffRed[i] * 1000) / 1000;
}

BlueRON = (double)Math.round(BlueRON * 100) / 100;
RedRON = (double)Math.round(RedRON * 100) / 100;

DarkBlue = (double)Math.round(DarkBlue * 100) / 100;
DarkRed = (double)Math.round(DarkRed * 100) / 100;

for (int i=0; i<5; i++) {
skyBlue[i] = (double)Math.round(skyBlue[i] * 10) / 10;
skyRed[i] = (double)Math.round(skyRed[i] * 10) / 10;

fluxBlue[i] = (double)Math.round(fluxBlue[i] * 10) / 10;
fluxRed[i] = (double)Math.round(fluxRed[i] * 10) / 10;

snrBlue[i] = (double)Math.round(snrBlue[i] * 10) / 10;
snrRed[i] = (double)Math.round(snrRed[i] * 10) / 10;
}


// ##################################################################################
// End of calculations:
// ##################################################################################


}
}
```

Bibliography

Adams, J. J., Hill, G. J., MacQueen, P. J.: “Volume phase holographic grating performance on the VIRUS-P instrument”, Proc. SPIE **7014**, 1471 (2008)

Angel, J. R. P. Adams, M. T. Boroson, T. A., Moore, R. L: “Relative contributions of scattering, diffraction and modal diffusion to focal ratio degradation in optical fibres”, ApJ **218**, 776 (1977)

Ångstrom, A. J.: “Optical investigations’, Phil. Mag. **9**, 327-342 (1855)

Baldry, I. K., Bland-Hawthorn, J., Robertson, J. G.: “Volume Phase Holographic Gratings: Polarization Properties and Diffraction Efficiency”, PASP **116**, 403-414 (2004)

Barden, S. C.: “Review of Fiber-Optic Properties for Astronomical Spectroscopy”, Proc. SPIE, **2476**, 2-9 (1995)

Barden, S. C., Arns, J. A., Colburn, W. S.: “Volume-phase holographic gratings and their potential for astronomical applications”, Proc. SPIE **3355**, 866-876 (1998)

Barden, S. C.: “Review of Fiber-Optic Properties for Astronomical Spectroscopy”, Fiber Optics in Astronomy III. ASP Conf. Series **152** (1998)

Barden, S. C., Arns, J. A., Colburn, W. S.: “Volume-Phase Holographic Gratings and the Efficiency of Three Simple VPH Gratings”, PASP (2000)

Barnes, S. I., Cottrell, P. L., Albrow, M. D., Frost, N., Graham, G., Kershaw, G., Ritchie, R., Jones, D., Sharples, R. M., Bramall, D. G., Schmoll, J., P. Luke, P., Clark, P.,Tyas, L. M. G., Buckley, D. A. H., Brink, J.: “The optical design of the Southern African Large Telescope High Resolution Spectrograph: SALT HRS”, Proc. SPIE **7014** (2008)

Bingham, R. G.: “Grating spectrometers and spectrographs re-examined”, QJRAS, **20**, 395-421 (1979)

Bramall, D. G., Sharples, R. M., Tyas, L. M. G., Schmoll, J., Clark, P., Luke, P., Looker, N., Dipper, N. A., Ryan, S., Buckley, D. A. H., Brink, J., Barnes, S. I.: “The SALT HRS spectrograph: final design, instrument capabilities and operational modes”, Proc. SPIE **7735** (2010)

Bohlin, R. C., Gilliland, R. L.: “HST Absolute Spectrophotometry of Vega from the Far-UV to the IR”, 4th June Astron. J., 3712 (2004)

Bowen, I. S.: “The Image-Slicer: a Device for Reducing Loss of Light at Slit of Stellar Spectrographs”, *ApJ*, **88**, 113 (1938)

Buckley, D. A. H.: “Annual Report of the SAAO” 35-36 (1995)

Buckley, D. A. H.: “SALT Science Requirements”, SALT *Document DB000531* (2000)

Buckley, D. A. H., Sessions, N.: “FIF pre-CDR”, *SALT document SALT-3400-AE-0024* (2004)

Buckley, D. A. H., Cottrell, P. L., Nordsieck, K. H., O'Donoghue, D. E., Williams, T. B.: "The first-generation instruments for the Southern African Large Telescope (SALT)", Proc. SPIE, **5492**, 60-74 (2004)

Buckley, D. A. H., Swart, G. P., Meiring, J. G.: "Completion and commissioning of the Southern African Large Telescope", Proc. SPIE **6267**, 62670Z (2006)

Butler, R. P., Marcy, G. W.: "Determining Spectrometer Instrumental Profiles Using FTS Spectra", PASP **107**, 966-976 (1995)

Campbell, B., Walker, G. A. H.: *IAU Colloquium 88, Stellar Radial Velocities* PASP, **91**, 550 (1979)

Carrasco, E., Parry, I. R.: "A method for determining the focal ratio degradation of optical fibres for astronomy". MNRAS, **271**, 1-12 (1994)

Castilho, B. V. Delabre, B., Gneiding, C. D.: "A new concept for échelle spectrographs: the SOAR Telescope Échelle Spectrograph", Proc. SPIE **5492**, 433-444 (2004)

Comte, A.: *Cours de Philosophie Positive*, **chap. II**, 19th lesson (1835)

Dekker, H., D'Odorico, S., Kaufer, A., Delabre, B., Kotzlowski, H.: "Design, construction and performance of UVES, the échelle spectrograph for the UT2 Kueyen Telescope at the ESO Paranal Observatory", Proc. SPIE **4008**, 534-545 (2000)

Dekker, H., Nissen, P. E., Kaufer, A., Primas, F., D'Odorico, S., Hanuschik, R. W.: "High S/N, high resolution Image Slicer observations with UVES", Proc. SPIE, **4842**, 139-150 (2003)

Draper, H.: "On diffraction spectrum photography", Phil. Mag. **46**, 417-425 (1873)

Florentin-Nielsen, R., Andersen, M. I., Nielsen, S. P.: "Cosmic Ray Events and Natural Radioactivity in CCD Cryostats", IAU Symposium, **167**, 207 (1995)

Fraunhofer, J.: "Kurtzer Bericht von der Resultaten neuerer Versuche über die Gesetze des Lichtes, und die Theorie derselbem", Gilberts Ann. Phys. **74**, 337-378 (1823)

Gaisser T. K. & Stanev, T.: "Cosmic Rays", Review of Particle Physics, Particle Data Group (2009)

Gregory, J. A., Burke, B. E., Kosicki, B. B., Reich, R. K.: "Developments in X-ray and astronomical CCD imagers", Nuclear Instruments and Methods in Physics Research A, **436**, 1-8 (1999)

Glazebrook, K., Bland-Hawthorn, J.: "Microslit Nod-Shuffle Spectroscopy: A Technique for Achieving Very High Densities of Spectra", PASP, **113**, 197-214 (2001)

Groom, D.: "Cosmic Rays and Other Nonsense in Astronomical CCD Imagers", Lawrence Berkeley National Laboratory (2002)

Harrison, G. R.: "The production of diffraction gratings II: The design of échelle gratings and spectrographs", J. Opt. Soc. Am. **39**, 522-528 (1949)

Heacox, W. D.: "On the Application of Optical-Fibre Image Scramblers to Astronomical Spectroscopy", Astron. J., **92**, 219-229 (1986)

Heacox, W. D.: "Radial image transfer by cylindrical, step-index optical waveguides", J. Opt. Soc. Am. A **4**, 488 (1987)

Hearnshaw, J.: *The analysis of starlight: one hundred and fifty years of astronomical spectroscopy*, Cambridge University Press (1986)

Herschel, J. F. W.: "On the absorption of light by coloured media, and on the colour of certain flames", Edinb. Trans. **9 II**, 445-460 (1823)

Huggins, W.: "On the spectrum of the great nebula in Orion and on the motions of some stars towards and away from the Sun", Proc. Royal Soc. **20**, 379-394 (1872)

Hunter, T. R. & Ramsey, L. W.: "Scrambling Properties of Optical Fibers and the Performance of a Double Scrambler", PASP, **104**, 1244-1251 (1992)

Janesick, J. R.: *Scientific Charge-Coupled Devices*, SPIE Press, Washington, USA (2001)

Jorden, P. R., Pool, P., Tulloch, S. M.: "Secrets of E2V Technologies CCDs ex Marconi CCDs", Astrophysics and Space Science Library, **300I**, 115-122 (2004)

Kaufer, A., Wolf, B., Essternwarte, L., Pasquini, L.: "FEROS, the Fiber-fed Extended Range Optical Spectrograph for the ESO 1.52-m telescope", The ESO Messenger **89** (1997)

Kaufer, A.: "A two-beam two-slice image slicer for fiber-linked spectrographs", Fiber Optics in Astronomy III, ASP Conf. Series (1998)

Kayser, H.: *Handbuch der Spectroscopie*, Vol. **1**, Hirtzel, Leipzig, (1900)

Kirchhoff, G.: "Über den Zusammenhang zwischen Emission und der Absorption von Licht und Wärme", Monatsber. d. Berlin. Akad., 783-787 (1859)

Kirchhoff, G. & Bunsen, R.: "Chemische Analyse durch Spectralbeobachtungen", Pogg. Ann. **110**, 160-189 (1860)

Kogelnik, H.: "Coupled wave theory for thick hologram gratings", *The Bell System Technical Journal*, **48**, 2909-2947 (1969)

Lockyer, J. N.: "Spectroscopic Notes, I to III", Proc. Royal Soc. **22**, 371-380 (1874)

Loewen, E. G & Popov, E.: *Diffraction Gratings and Applications*, Taylor and Francis (1997)

Lu, G.: "Optical Fibre for UV-IR Broadband Spectroscopy", Proc. SPIE, **3355**, 884-891 (1998)

Lund, G. & Enard, D.: Proc. SPIE, **V 445**, 65 (1984)

Marcy, G. W., Butler, R. P.: "Precision radial velocities with an iodine absorption cell", PASP, **104**, 270 (1992)

Meiring, J. G. & Buckley, D. A. H.: "South African Large Telescope (SALT) project progress and status after four years", Proc. SPIE 5489, **592** (2004)

Miller, J. L. & E. Friedman, E.: *Photonics Rules of Thumb*, McGraw Hill, New York, USA (2003)

Murphy, J. D., MacQueen, P. J., Hill, G. J., Grupp, F., Kelz, A., Palunas, P., Roth, M., Fry, A.: "Focal Ratio Degradation and Transmission in VIRUS-P Optical Fibers", Proc. SPIE, **7018**, 7018T-12 (2008)

Nagel, S. R.: "Lightwave Communications Using Optical Fibers", E12, printed in D. Halliday, R. Resnick, & J. Walker, *Fundamentals of Physics, Extended (4th Ed.)*, John Wiley & Sons, USA (1993)

Nelson, G. W.: "Introduction to Fiber Optics", PASPC, **3**, 2-22 (1988)

Noguchi, K., Aoki, W., Kawanomoto, S., Ando, H., Honda, S., Izumiura, H., Kambe, E., Okita, K., Sadakane, K., Sato, B., Tajitsu, A., Takada-Hidai, T., Tanaka, W., Watanabe, E., Yoshida, M.: "High Dispersion Spectrograph (HDS) for the Subaru Telescope", Pub. Ast. Soc. Japan, **54**, 855-864 (2002)

O'Donoghue, D., Buckley, D. A. H., Balona, L. A., Bester, D., Botha, L., Brink, J., Carter, D. B., Charles, P. A., Christians, A., Ebrahim, F., Emmerich, R., Esterhuyse, W., Evans, G. P., Fourie, C., Fourie, P., Gajjar, H., Gordon, M., Gumede, C., de Kock, M., Koeslag, A., Koorts, W.P., Kriel, H., Marang, F., Meiring, J. G., Menzies, J. W., Menzies, P., Metcalfe, D., Meyer, B., Nel, L., O'Connor, J., Osman, F., du Plessis, C., Rall, H., Riddick, A., Romero-Colmenero, E., Potter, S.B., Sass, C., Schalekamp, H., Sessions, N., Siyengo, S., Sopela, V., Steyn, H., Stoffels, J., Stolz, J., Swart, G., Swat, A., Swiegers, A., Tiheli, T., Vaisanen, P., Whittaker, W., van Wyk, F.: "First science with SALT:

peering at the accreting polar caps of the eclipsing polar SDSS J015543.40+002807.2", Mon. Not. Roy. Astron. Soc. **372**, 151-162 (2006)

O'Donoghue, D. E., O'Connor, J., Crause, L. A., Strumpfer, A., Strydom, O. J., Brink, J. D., Sass, C., Wiid, E., Atad-Ettedgui, E "Saving SALT: repairs to the spherical aberration corrector of the Southern African Large Telescope (SALT)", Proc. SPIE **7793**, 7739Q (2010)

Pallavicini, R., Zerbi, F. M., Spano, P., Conconi, E., Molinari, R., Mazzoleni, Strassmeier, K. G.: "A new concept and a preliminary design for a high resolution (HR) and very-high resolution (VHR) spectrograph for the LBT", Astronomische Nachrichten **323** 488-493 (2003)

Parry, I. R.: "The Astronomical Uses of Optical Fibers", ASP Conf. Series, **152**, 3-13 (1998)

Pasquini, L., Manescau, A., Avila, G., Delabre, B., Dekker, H., Liske, J., D'Odorico, S., Pepe, F., Dessauges, M., Lovis, C., Megevand, D., Queloz, D., Udry, S., Cristiani, S., Bonifacio, P., Dimarcantonio, P.: "ESPRESSO: A High Resolution Spectrograph for the Combined Coudé Focus of the VLT", *ESO Messenger Review* (2009)

Peckerar, M. C., McCann, D., Blaha, F., Mend, W., Fulton, R.: "Deep Depletion Charge-Coupled Devices for X-Ray and IR Sensing Applications", *Westinghouse Electric Corporation 21203* (1979)

Pepe, F., Mayor, M., Queloz, D., Udry, S., Delabre, B., Kohler, D., Lacroix, D., Sivan, J. P., Benz, W., Bertaux, J. L.: "HARPS: a new high-resolution spectrograph for the search of extrasolar planets", Proc. SPIE **4008**, 582 (2000)

Pepe, F., Mayor, M., Rupprecht, G. Avila, G.; Ballester, P.; Beckers, J. L.; Benz, W.; Bertaux, J. L.; Bouchy, F.; Buzzoni, B.; Cavadore, C.; Deiries, S.; Dekker, H.; Delabre, B.; D'Odorico, S.; Eckert, W.; Fischer, J.; Fleury, M.; George, M.; Gilliotte, A.; Gojak, D.; Guzman, J.-C.; Koch, F.; Kohler, D.; Kotzlowski, H.; Lacroix, D.; Le Merrer, J.; Lizon, J.-L.; Lo Curto, G.; Longinotti, A.; Megevand, D.; Pasquini, L.; Petitpas, P.; Pichard, M.; Queloz, D.; Reyes, J.; Richaud, P.; Sivan, J.-P.; Sosnowska, D.; Soto, R.; Udry, S.; Ureta, E.; van Kesteren, A.; Weber, L.; Weilenmann, U.; Wicenec, A.; Wieland, G.; Christensen-Dalsgaard, J.; Dravins, D.; Hatzes, A.; Krster, M.; Paresce, F.; Penny, A: "HARPS: ESO's coming planet searcher. Chasing exoplanets with the La Silla 3.6-m telescope", The Messenger, **110**, 9-14 (2002)

Pickles, A. J.: "A Stellar Flux Library: 1150-25000Å", PASP, **110**, 749 July (1998)

Plücker, J.: "Forgesetze Beobachtungen über die electrishe Entladung", Pogg. Ann. **104**, 113-128 (1858)

Quincke, G.: "Optische Experimentaluntersuchunger: XV. On diffraction gratings", Ann. der Physik, **146**, 1-65 (1872)

Ramsey, L. W. & Weedman, D. W.: "The Penn State spectroscopic survey telescope", Proc. ESO Colloquium, Garching, April 9-12, 851-860 (1984)

Ramsey, L. W.: "Focal Ratio Degradation in Optical Fibers of Astronomical Interest", ASP Conf. Series, **3**, 26-40 (1988)

Rittenhouse, D.: "An optical problem proposed by F. Hopkinson and solved", J. Am. Phil. Soc. **201**, 202-206 (1786)

Ritter, J.: "Versuche über das Sonnenlicht", Gilberts Ann. **12**, 409-415 (1803)

Schmoll, J., Roth, M. M., Laux, U.: "Statistical Test of Optical Fibers for use in PMAS, the Potsdam Multi-Aperture Spectrophotometer", PASP, 02-129 (2003)

Schroeder, D. J.: *Astronomical Optics*, Academic Press, San Diego, USA (1989)

Secchi, P.: "Analyse spectrale de la lumire de quelques étoiles, et nouvelles observations sur les taches solaires", Comptes Rendus des Sances de l'Acadmie des Sciences, **63**, 364-368 (1866)

Serkowski, K., Frecker, J. D., Heacox, W. D., Roland, E. H.: Proc. SPIE, **172**, 130 (1979)

Sessions, N. J. A.: “An Investigation into the Fibre Instrument Feed for the Southern African Large Telescope (SALT)”, *University of Cape Town MEng thesis* (2001)

Smith, A. R., MacDonald, R. J., Hurley, D. L., Holland, S. E., Groom, D. E., Brown, W. E., Gilmore, D. K., Stover, R. J., Wei, M.: “Radiation events in astronomical CCD images”, Proc. SPIE **4669**, 172-183 (2002)

Stempels, E.: “FIEStool v1.0”, NOT Observatory Document (2005)

Strassmeier, K. G., Woche, M., Andersen, M., Ilyin, I.: “PEPSI: The Potsdam Échelle Polarimetric and Spectroscopic Instrument for the Large Binocular Telescope (LBT)”, *Astronomische Nachrichten* **328**, 627 (2007)

Swart, G.: “SALT Specification and Error Budget Status”, *a presentation to* SSWG (2001)

Tamura, N., Murray, G. J., Luke, P., Blackburn, C., Robertson, D. J., Dipper, N. A., Sharples, R. M., Allington-Smith, J. R.: “Cryogenic tets of volume-phase holographic gratings: results at 100K”, Appl. Opt. **45**, 5923-5928 (2006)

Tennyson, J.: *Astronomical Spectroscopy: An Introduction to the Atomic and Molecular Physics of Astronomical Spectra*, Imperial College Press (2005)

Tull, R.: “High-resolution fiber-coupled spectrograph of the Hobby-Eberly Telescope”, Proc. SPIE Kona Meeting, **387** (1998)

Tulloch, S.: “Investigation of Low Fringing Detectors on the ISIS Spectrograph”, *ING La Palma Technical Note*, **No. 130** (2005)

van Dokkum, P.: “Cosmic-Ray Rejection by Laplacian Edge Detection”, PASP **113**, 1420-1427 (2001)

Vaughnn, D.: “What’s wrong with the throughput-resolution product? A fiber-fed spectrograph forces a reevaluation of instrument design parameters”, Proc. SPIE **2198**, 31-43 (1994)

Vogt, S. S., Allen, S. L., Bigelow, B. C., Bresee, L., Cantrall, T., Conrad, A., Couture, M., Delaney, C., Epps, H. W., Hilyard, D., Hilyard, D. F., Horn, E., Jern, N., Kanto, D., Keane, M. J., Kibrick, R. I., Lewis, J. W., Pardeilhan, G. H., Pfister, T., Ricketts, T., Robinson, L. B., Stover, R. J., Tucker, D., Ward, J. M., Wei, M.: “HIRES: the high-resolution échelle spectrometer on the Keck 10-m Telescope”, Proc. SPIE, **362**, 2198 (1994)

Walraven, T.: in S. Lausten and A. Riez (eds.), “Aux. Instrumentation for Large Telescopes”, 175-180 (1972)

Wheatstone, C.: “On the prismatic decomposition of the electric, voltaic and electro-magnetic sparks”, Chem. News, **3**, 198-201 (1861)

Wood, R. W.: “The use of échelette gratings in high orders”, J. Opt. Soc. Am. **37**, 733-737 (1947)

Young, T.: “On the theory of light and colours”, Phil. Trans. **II**, 399-408 (1803)

Zöllner, F.: “Über die spectroscopische Beobachtung der Rotation der Sonne und ein neues Reversionsspectroscop”, Pogg. Ann. **144**, 449-456 (1871)