Birmingham Solar Oscillations Network: The Next Generation

by

Steven James Hale

A thesis submitted to the
University of Birmingham
for the degree of
Doctor of Philosophy

School of Physics and Astronomy
College of Engineering and Physical Sciences
University of Birmingham
December 2018

Abstract

The aim of this thesis is to ensure the continued operation and possible expansion of the Birmingham Solar Oscillations Network (BiSON) through development of a new instrumentation paradigm known as *BiSON-Mini*.

The first half of the thesis discusses the emergence of the field of helioseismology, and the history of the BiSON network. A comparison baseline is described by assessing the historic network performance in terms of duty cycle and noise levels, and in terms of atmospheric extinction. The baseline informs performance to be matched, or ideally exceeded, while developing prototype miniaturised instrumentation to complement the existing network.

With a performance baseline described, in the second half of the thesis we develop prototype instrumentation that meets the requirements of simplicity and reduced physical footprint. The overall noise characteristics of the prototype system were assessed during an observing campaign over summer 2018 at Izaña, Tenerife, and found to be comparable with the large-scale existing BiSON infrastructure allowing seamless integration. Following the results from this project, BiSON is certainly well placed to observe the next solar cycle and beyond.

Acknowledgements

First and foremost thanks go to my supervisors Prof. Bill Chaplin and Prof. Yvonne Elsworth, without whom this work would not exist.

I would like to thank all those who are, or have been, associated with BiSON. In no order: George Isaak, Bill Brookes, Bob van der Raay, Clive McLeod, Roger New, Sarah Wheeler, Clive Speake, Rachel Howe, Guy Davies, Anne-Marie Broomhall, Graham Verner, Bernard McCarty, Richard Lines, Phil Pavelin, Barry Jackson, Hugh Williams, Joe Litherland, Tony Taylor, Ian Barnes, Richard Bryan, John Allison, and Eddie Ross. Particular thanks must go to Brek Miller for the endless pro-bono consulting.

Technical help was received from David Hoyland regarding PCB design techniques, and John Bryant on surface-mount reflow soldering – thanks to you both. In Physics Mechanical Workshop – Stephen Brookes and the team for help with the more technical machining that I was unable to manage myself. In Physics stores – Ian Barnett, Mohammed "Chaw" Mughal, and Paul Liddington. Technical and building managers – Tendai Makuwatsine and Mark Wicks.

None of this work would be possible without the extensive help of the on-site staff at the remote observatories:

- At Mount Wilson: Prof. E. J. Rhodes, Jr., Stephen Pinkerton, the team of USC undergraduate observing assistants, former University of Southern California staff members Maynard Clark, Perry Rose, Natasha Johnson, Steve Padilla, and Shawn Irish, and former University of California Los Angeles staff members Larry Webster and John Boyden.

- At Las Campanas: Patricio Pinto, Andrés Rivera, Emilio Cerda, Frank Perez, Marc Hellebaut, Patricio Jones, Gastón Gutierrez, Juan Navarro, Francesco Di Mille, Roberto Bermudez, and all the staff of Las Campanas Observatory.

- At Izaña: Dr. Pere Pallé, Dr. Teo Roca Cortés, Santiago López, Antonio Pimienta, and all the staff at the Instituto de Astrofísica de Canarias who have contributed to running the Mark-I instrument over many years.

- At Sutherland: Pieter Fourie, Willie Koorts, Jaci Cloete, Reginald Klein, John Stoffels, and the staff of the South African Astronomical Observatory.

- At Carnarvon: Les Bateman, Les Schultz, Sabrina Dowling-Giudici, Inge Lauw of Williams and Hughes Lawyers, and NBN Co. Ltd.

- At Narrabri: Mike Hill and the staff of the Paul Wild Observatory and the Commonwealth Scientific and Industrial Research Organisation.

Almost all of the work on this project was completed using open source free software, and so I would like to thank the community as a whole for the effort put into open source projects. Specific thanks to Jasem Mutlaq for help with the INDI Library, and Jean-Luc Geehalel for advice on the EQMOD API. All data analysis was completed using Python, with figures produced using Matplotlib. Vector drawings were produced with Inkscape. Raytracing was completed using the free OSLO-EDU package. The thesis was typeset using LaTeX.

Finally, I would like to acknowledge the funding assistance of the School of Physics and Astronomy at the University of Birmingham. BiSON is funded by the Science and Technology Facilities Council (STFC) in the UK. Additional funding was secured via the STFC Impact Accelerator account, with the assistance of Alan Tibbatts from Alta Innovations.

Sincere apologies to anyone that has been inadvertently forgotten.

S. J. Hale

Birmingham, UK.
December 2018

Contents

List of Figures

List of Tables

Nos esse quasi nanos gigantum humeris insidientes.

We are as dwarfs sitting on the shoulders of giants.

Said of Bernard of Chartres in *The Metalogicon* (1159)
John of Salisbury, c. 1120 – 1180

If I have seen further it is by standing on the shoulders of Giants.

Letter to Robert Hooke (5 February 1675)
Isaac Newton, 1643 – 1727

Chapter 1

Introduction

The aim of this project is to ensure the continued operation and possible expansion of the Birmingham Solar Oscillations Network (BiSON) through development of a new instrumentation paradigm known as *BiSON-Mini*. The network, observing oscillations of the Sun, has been in operation since the early 1990s. Now approaching nearly 30 years of age the hardware and electronics, some of which are still original, are showing increasing failure rates. Simultaneously, the funding for in-house technical support has reduced. With reduced funding and limited technical support the instrumentation needs to become simpler to maintain and more cost effective.

The current BiSON network consists of six sites, of which up to four may have the Sun within view simultaneously. The locations of the sites are listed in Table 1.1, and shown on a map in Figure 1.1. There are two stations in each 120° longitude band, and all of the sites lie at moderate latitudes, around ±30°. Four of the sites are housed inside classical observatory domes with equatorial mounts, and two make use of cœlostats mounted on top of specialist buildings. All sites measure oscillations that are used as a probe of the internal structure of the Sun, shown in Figure 1.2.

The core of the Sun has a temperature of around 15×10^6 K, and occupies about 0.2 to 0.25 solar radii. Energy is released via nuclear fusion in the core, through the p–p (proton–proton) chain and to a smaller extent through the CNO (carbon–nitrogen–oxygen)

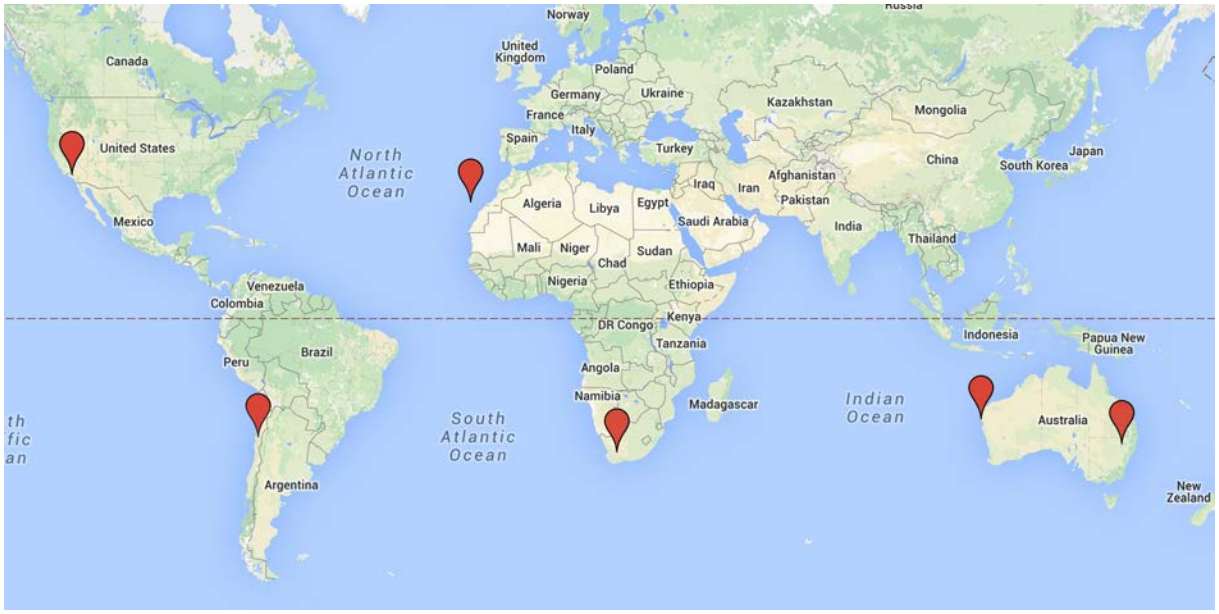

Figure 1.1: The six station Birmingham Solar Oscillations Network (BiSON). Image credit: Google Maps

Location	Longitude (deg E)	Latitude (deg N)	Altitude (metres)	Commissioned (year)
Mount Wilson, California, USA	-118.08	+34.13	1742	1992
Las Campanas, Chile	-70.70	-29.02	2282	1991
Izaña, Canary Islands	-16.50	+28.30	2368	1975
Sutherland, South Africa	+20.82	-32.38	1771	1990
Carnarvon, Western Australia	+113.75	-24.85	10	1985
Narrabri, NSW, Australia	+149.57	-30.32	217	1992

Table 1.1: The six station Birmingham Solar Oscillations Network (BiSON).

cycle. Energy is transported from the core primarily via radiation, since to 0.7 solar radii the layers are stable against convection. Throughout the radiative zone the temperature drops from about 7×10^6 K to 2×10^6 K. In the convection zone, between about 0.7 solar radii and the visible surface, the radiative gradient is steeper than the adiabatic gradient and so the layer becomes unstable to convection, and energy here is transported primarily by movement of thermal cells. The boundary at the transition between the radiatively stable layer and the convection zone is known as the tachocline, since this is the region where a large velocity shear occurs between uniform rotation and differential rotation in the two layers, respectively. The shear at the tachocline is thought to be responsible for

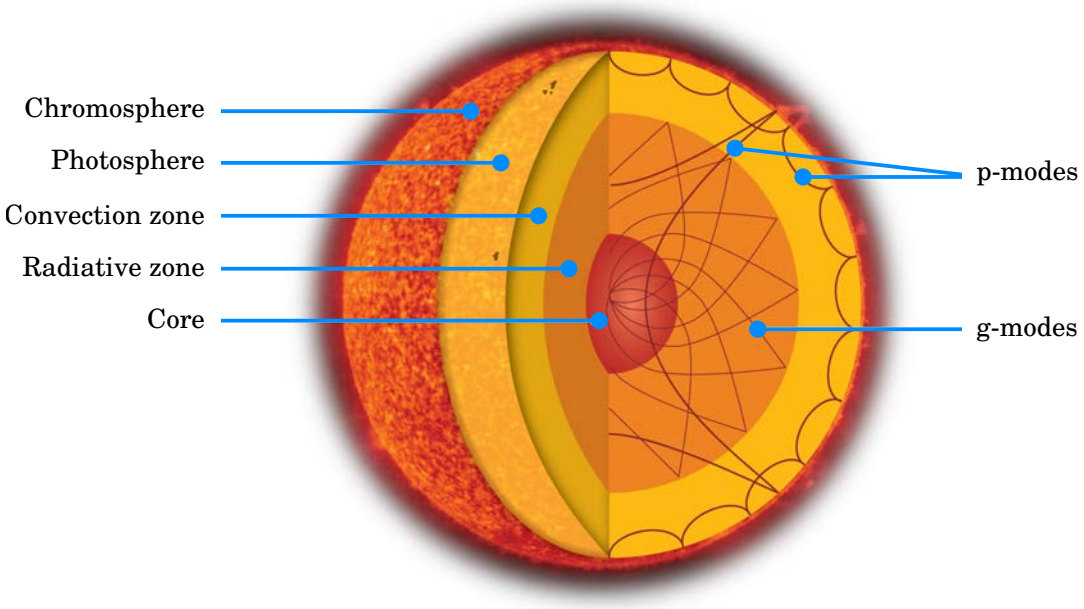

Figure 1.2: Cutaway diagram showing key regions of the solar interior, from the core, the inner radiative zone, the turbulent outer convection zone, followed by the photosphere and chromosphere. Image credit: ESA; Sun's chromosphere image based on SOHO image; credit: SOHO (ESA & NASA).

the dynamo that drives the solar magnetic field. The visible surface of the Sun is known as the photosphere, where the temperature has dropped to around 5777 K.

BiSON measures many so-called p-modes (pressure modes), which are the result of interference of acoustic waves generated by turbulence in the outermost layers of the convection zone. The sound waves are stochastically excited and intrinsically damped, and there is a broad spectrum of waves both radial and non-radial on a range of spatial and temporal scales. The Sun acts as a natural cavity to trap these waves, since a sound wave generated in the convection zone heading towards the solar interior is refracted as the temperature increases, and a sound wave heading towards the solar surface is returned inwards by reflection due to the decreasing scale height at the top of the convection zone. A wave injected perpendicular to the surface passes all the way through. Waves having the correct characteristics constructively interfere and set up standing wave patterns that can be detected from their effect on the photosphere.

The modes of oscillation are characterised by the number of nodes n in the radial direction, known as the radial order, and two numbers that describe the spatial pattern of

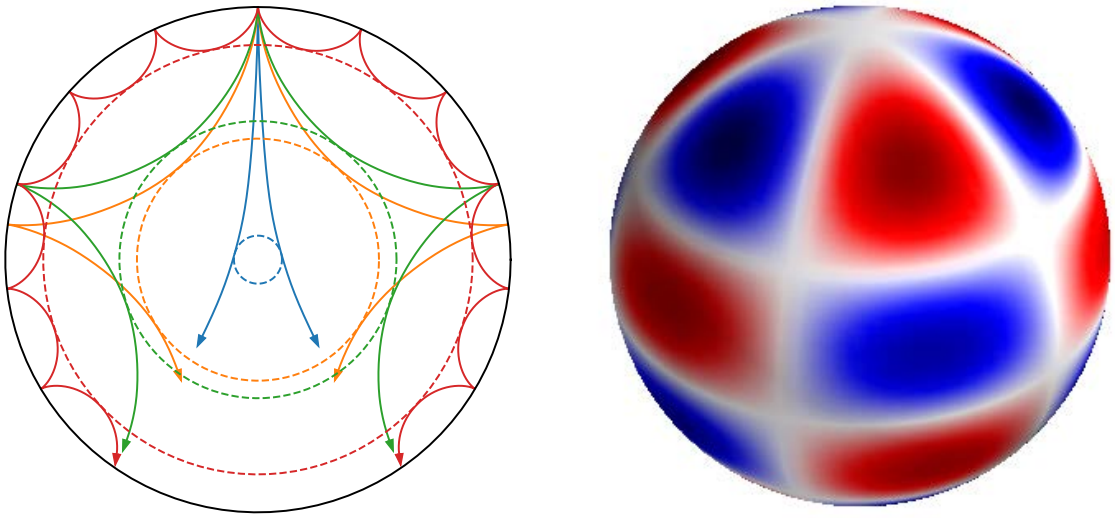

(a) (b)

Figure 1.3: (a) shows mode formation rays for a range of ℓ, computed from the standard reference helioseismic "Model S" (Christensen-Dalsgaard et al., 1996). (b) shows the spherical harmonic for $\ell = 6$, $m = 3$, where blue areas are approaching the observer and red areas are receding. Image credit: Warrick H. Ball

the oscillations across the surface, which may be described as spherical harmonics. The number of nodal lines on the surface is the angular degree ℓ, and the number of nodal lines crossing the equator is the azimuthal order m. Some examples are shown in Figure 1.3. The largest amplitude p-modes have frequencies of $\sim$3150 µHz and a period of around five minutes. The lowest frequency modes are of particular interest, since these are more sensitive to the conditions inside the deep solar interior. Few modes with frequencies less than 1000 µHz have been detected – see, e.g., García et al. (2001); Chaplin et al. (2002); Broomhall et al. (2007); Davies et al. (2014). Observing such modes would allow better constraints to be put on the structure of the solar core – see, e.g., Houdek et al. (1998, 1999, 2001); Chaplin et al. (2004a).

Additionally, so-called g-modes (gravity modes) are expected to have even lower frequencies ($<$470 µHz), and so offer greater sensitivity to conditions within the core of the Sun. Since their restoring force is buoyancy, g-modes can only exist in regions that are stable against convection meaning they are trapped within the radiative zone. Such g-

modes become evanescent in the convection zone and so are expected to have very low amplitude effects on the photosphere. Observations of these g-modes would be a far more sensitive probe of the conditions inside the radiative interior (García et al., 2007; Mathur et al., 2007, 2008; Appourchaux et al., 2010). Unfortunately noise, of both solar and instrumental origin, is highest below 1000 µHz, and so observing these modes is a significant observational challenge that has not yet been incontrovertibly achieved. A recent claim of detection by Fossat et al. (2017) has been convincingly refuted (Schunker et al., 2018; Scherrer and Gough, 2019).

Refreshing the existing ageing hardware will ensure the ongoing programme of solar activity cycle observations is successful. Long term continuous and comprehensive (so-called synoptic) observations of the Sun remain important (Elsworth et al., 2015). The Sun is a variable star, and the solar output changes on time scales that vary from short, in terms of seconds and minutes to days, and long, in terms of years to geological time. Current datasets extend to almost three 11 year solar cycles (one complete 22 year cycle of magnetic polarity reversal) and yet our understanding of the solar dynamo, the mechanism believed to be driving the solar cycle, is poor. More than 50 models give a wide range of predictions regarding the timing and amplitude of each cycle (Pesnell, 2008). It is essential that observations continue as we go into Cycle 25 and beyond. As a consequence of the reduced cost and physical footprint made available by BiSON-Mini, it may even enable the legacy BiSON network to be extended with additional nodes with the aim of detecting these very low-frequency modes of oscillation below 1000 µHz. Davies et al. (2014) show that by making more efficient use of contemporaneous data between BiSON sites it is possible to reduce the noise level in the resulting power spectrum. Similarly if an instrument makes multi-wavelength observations, for example multiple measurement points across a single spectral line, Lund et al. (2017) show that it is again possible to reduce the noise level. The improvement in noise level comes from leveraging both the incoherence of instrumental noise between contemporaneous sites, and that different points on the spectral line are formed over a range of heights in the solar atmosphere and

so intrinsically experience different realisations of noise (Espagnet et al., 1995; Lefebvre et al., 2008). This means by observing at different points on the line, measurements are made that are spatially independent and so involve coherent signal but incoherent noise. These effects combine to provide several independent realisations of the noise and so when merged improve the signal to noise ratio (García et al., 2004). The commissioning of additional BiSON-mini nodes would increase the number of incoherent noise realisations and so allow a reduction in overall noise level and an improvement in precision.

The progress of technology has been significant since the BiSON nodes were first commissioned. Electronic miniaturisation has produced micro-controllers and single-board computers, allowing old single-purpose hardware to be replaced with software-defined control systems that are both simpler to update and have considerably lower deployment cost. Additionally, taking advantage of modern fibre optics allows the instrumentation to be split with only the light collection optics on board a telescope mount rather than the entire instrument, and this results in a much smaller physical footprint. To borrow a term from the information technology community, we need to make use of inexpensive commercial off-the-shelf (COTS) components where possible. This includes stock optics and mechanics as opposed to expensive custom components. Since the observed signal from the oscillations is small with long timescales, it is essential the instrumentation has high stability over very long time periods. Any changes to the instrumentation must be moderate in order to achieve consistency in the data, and not involve a fundamental change in the observing methods. Whilst it is of course important to take advantage of improvements in technology when developing newer instrumentation, such changes must always be undertaken slowly and in parallel with existing instrumentation to allow comparison studies.

Before undertaking any major instrumentation changes, we should first consider the historical data archive and assess the performance that has been achieved. In Chapter 2, we give a brief overview of the history of observational helioseismology, and the contribution made by the University of Birmingham as the field developed in the 1970s. In

Chapter 3, we provide a qualitative discussion on the design and operation of a BiSON spectrophotometer. The level of discussion in Chapter 3 is sufficient to enable understanding of the performance analysis in the following two chapters. A more in-depth quantitative discussion on the existing instrumentation is in Chapter 7.

In Chapter 4 we discuss the history of each BiSON site, including personal detail from the author's own over 15 years experience working in the field. Each site is assessed in terms of operational fill (i.e., number of successful observing hours achieved), signal to noise ratio, and overall mean noise level at high-frequency. In Chapter 5, we make a similar performance assessment in terms of site atmospheric extinction.

With a performance baseline described, in the second half of this thesis we go on to develop prototype instrumentation that meets the required simplicity and reduced physical footprint. Although it is possible to use signal averaging techniques to recover a signal hidden in noise, it is always best to start with a system designed around the lowest practical noise floor. The effect of noise in resonant scattering spectrophotometers has been previously investigated in some detail, but the existing literature considers attempts to discriminate different noise sources via the final data product (e.g., Hoyng, 1989). This is fundamentally approaching the problem in the wrong direction. Instead, we now take an electrical engineering approach to new instrumentation design, measuring the performance and noise levels of each individual component before incorporating it into the overall system. We develop an end-to-end model of a spectrophotometer starting from the solar constant, progressing through the expected signal and noise level at the detectors, all the way to quantisation noise in the final data acquisition system. The resulting performance specification has previously been undocumented for any BiSON instrumentation.

We begin in Chapter 6 by modelling a vapour reference cell from first principles, and validating the accuracy of the model using a tunable diode-laser probing a standard BiSON reference cell. Using the model we show the effects of variation in vapour temperature and aperture size within the scattering photodetectors, and later this allows us to make an

estimate of the contribution of fluctuations in vapour temperature control to the overall noise level.

In Chapter 7, we embed the vapour cell model within a wider system model. This enables us to estimate the optical throughput of light incident on the spectrophotometer collection optics, and so estimate the power incident on the vapour cell scattering photodetectors informing later detector design. The impact of polarisation switching rates is considered, and the magnetic field strength required for optimum Zeeman splitting is determined via further modelling of the solar potassium absorption line.

The results from Chapter 7 are used in Chapter 8, where the estimated power incident on the scattering detectors informs the amplification required within the data acquisition system, and measurements of polarisation switching rates informs the bandwidth requirements. All noise sources are assessed in isolation, both through circuit simulation and empirical testing, to determine their contribution to the final overall noise level. The new instrumentation was commissioned at Izaña, Tenerife, during a summer observing campaign in 2017. The results from the new instrumentation and other contemporaneous BiSON sites over the entire 2018 summer observing campaign are compared with the established performance baseline.

With the prototype instrumentation complete, in Chapter 9 we move on to the problem of tracking the Sun and maintaining light into the spectrophotometer. Commercial products are available for collecting and using energy from the Sun. So-called *daylighting systems* aim to concentrate solar energy from a single collection point and distribute this around a building using optical fibres (Ullah and Shin, 2012; Lingfors and Volotinen, 2013). It would be convenient if we could make use of one of these off-the-shelf systems since they are already designed to track the Sun and concentrate light into a fibre. Unfortunately there is little information available on the actual design, construction, and operation of these tracking concentrators – for obvious commercial reasons. A key requirement in our system is that light from the entire solar disc is collected without vignetting. Unfortunately there is no way to ensure this would always be the case from a

commercial unit and so such systems were not considered further.

Since the original BiSON sites were commissioned there has been considerable technological advancement in amateur telescope mounts. Many inexpensive consumer-grade amateur telescope mounts now include at least partial automation. So-called "GOTO" telescope mounts typically include a handset that allows the mount to be pointed at any target automatically, following a manual alignment and calibration procedure. It is these mounts that we will consider in Chapter 9 for our tracking and guiding requirements. There are several difficulties to overcome when attempting to fully automate a partially automated consumer mount. Many COTS mounts either do not have facility for computer control at all, or the communications protocol is proprietary to the manufacturer's own control handsets. The communication and control problem can be solved by considering only manufacturers that open-source and publish their communications interface. Additionally, the mounts typically do not make use of absolute position encoders and so their control systems are unable to keep track of where the mount is pointing when power is lost or the mount is moved manually by hand. This is acceptable for amateur use only because such mounts are intended to be transported and manually aligned each time they are used and to be packed away at the end of each observing session. It is possible to obtain absolute position information by making use of micro-electro-mechanical systems (MEMS) to track the position of a telescope, such as the accelerometers used to control screen orientation in devices such as smartphones and tablets. These can easily be mounted on the telescope itself and avoid costly modifications to the mount. Throughout Chapter 9 we will assess the viability of obtaining reliable pointing information for use when tracking the Sun, and the precision achievable when guiding using commands sent over the proprietary communications motor control interface.

Finally, in Chapter 10 we take an overview of the work completed, consider the extent and limitations of the results, and discuss prospects for further work.

We will begin now with a brief history of the field of helioseismology.

Chapter 2

Observational Helioseismology: A Brief History

Contents

2.1 Introduction

> This chapter includes text from the publication "Performance of the *Birmingham Solar-Oscillations Network* (BiSON)" (Hale et al., 2016) in the journal *Solar Physics*. The published article was first released on 2015 December 7, and is included in Appendix D for reference. Whilst this publication is a result of collaborative working, the content is entirely the author's own work. Co-authors provided guidance on historical accuracy and manuscript review.

The Sun rings like a bell (Wolff, 1972; Chaplin, 2006). The consideration of sound being generated near the Sun's surface dates back to the mid 1940s. Scientists at the time had suggested it as a possible source of odd "temperature reversal" problems observed in the solar atmosphere (Biermann, 1946; Schwarzschild, 1948), where the temperature of the corona had been seen to exceed temperatures found nearer the visible surface. Both Biermann and Schwarzschild suggested that such sound waves could transport energy into

the upper parts of the atmosphere and heat the corona.

It was not until some years later that more serious consideration was made to the solar atmosphere being filled with acoustic energy, when Leighton et al. (1962) discovered oscillations in the velocity field across the Sun. Robert Leighton used a novel technique to measure the movement of the gas in the solar atmosphere, derived from the work of George Ellery Hale at the Mount Wilson (Hale) Observatory (Hale, 1908a) in Los Angeles, USA. George Ellery Hale, seen in Figure 2.1, developed a "spectroheliograph" to make observations of magnetic fields on the surface of the Sun (Hale, 1908b). Robert Leighton modified the spectroheliograph in such a way that it could measure the Doppler shift of selected Fraunhofer lines, rather than the effects of magnetic fields. The so-called Fraunhofer absorption lines are formed when a cooler gas is placed in between a hotter broad-spectrum source and the observer. The cooler gas absorbs light at its characteristic frequency, and then re-emits it isotropically. Since some of the light is now scattered in directions away from the observer, a reduction in intensity is seen – an "absorption line".

Figure 2.1: George Ellery Hale operating the spectroheliograph at Mount Wilson. Image credit: Courtesy, Huntington Library.

Light from the Sun is similar to that of a black body at a temperature of 6000 K. The emission spectrum of a black body in thermal equilibrium is described by Planck's law,

$$B(\lambda, T)\,\mathrm{d}\lambda = \frac{2hc^2}{\lambda^5}\frac{1}{e^{\frac{hc}{\lambda k_\mathrm{B} T}} - 1}\,\mathrm{d}\lambda\ , \tag{2.1}$$

where λ is the wavelength, T the temperature, k_B is the Boltzmann constant, h the Planck

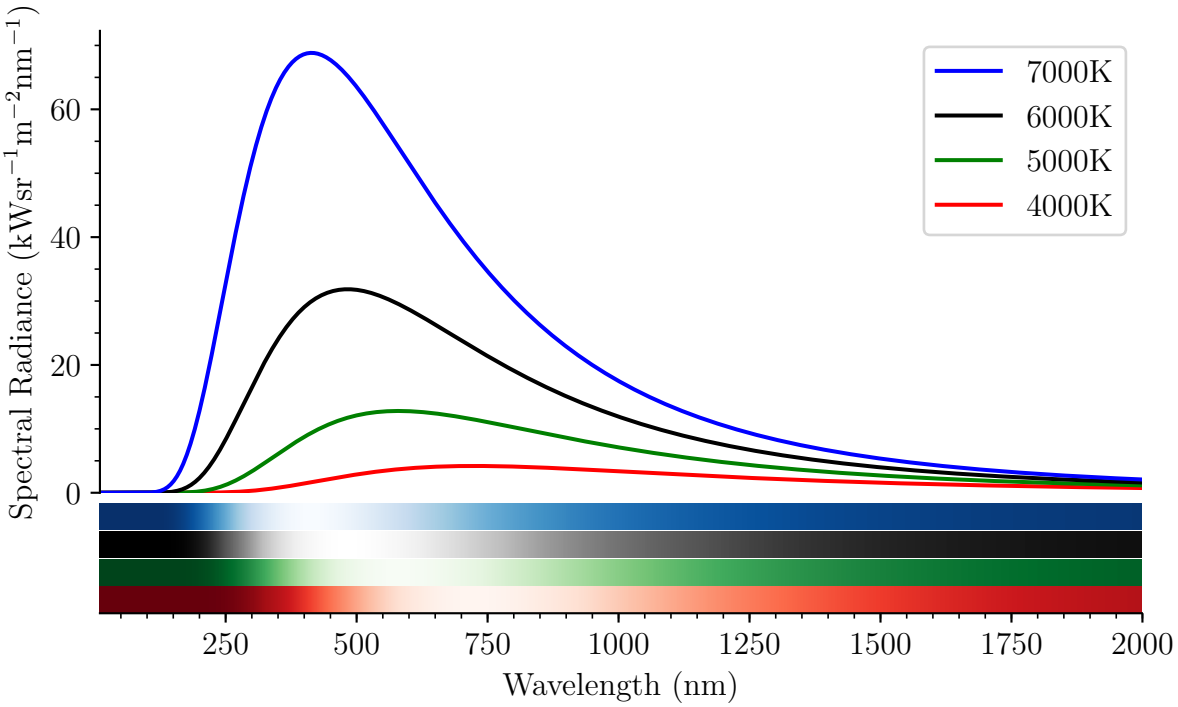

Figure 2.2: Planck's law describes the emission spectrum of a black body in thermal equilibrium at a specific temperature. Light from the Sun is similar to that of a black body at a temperature of 6000 K. The strips at the bottom of the plot show how such intensity profiles are observed through a spectrophotometer. The colours are for identification only, and are not representative of wavelength.

constant, and c the speed of light. Examples for a selection of different temperatures are shown in Figure 2.2. In the Sun, the cooler outer regions scatter light from the hotter inner regions resulting in such absorption lines appearing in the solar spectrum. These lines were observed by German physicist Joseph von Fraunhofer in the early 1800s, and have since become known as Fraunhofer lines. An example of the formation of an absorption line is shown in Figure 2.3.

Leighton measured Doppler shifts in the wavelength of such solar absorption lines. As the line shifts in wavelength due to the motion of the solar atmosphere, the light intensity passed by the spectroheliograph changed in relation to the velocity. The technique required exposure of two photographic plates, with each plate capturing light passed in a very narrow wavelength band centred on the red- and blue-shifted wings of a selected Fraunhofer line. In order to enhance the contrast of the shift, one of the photographic

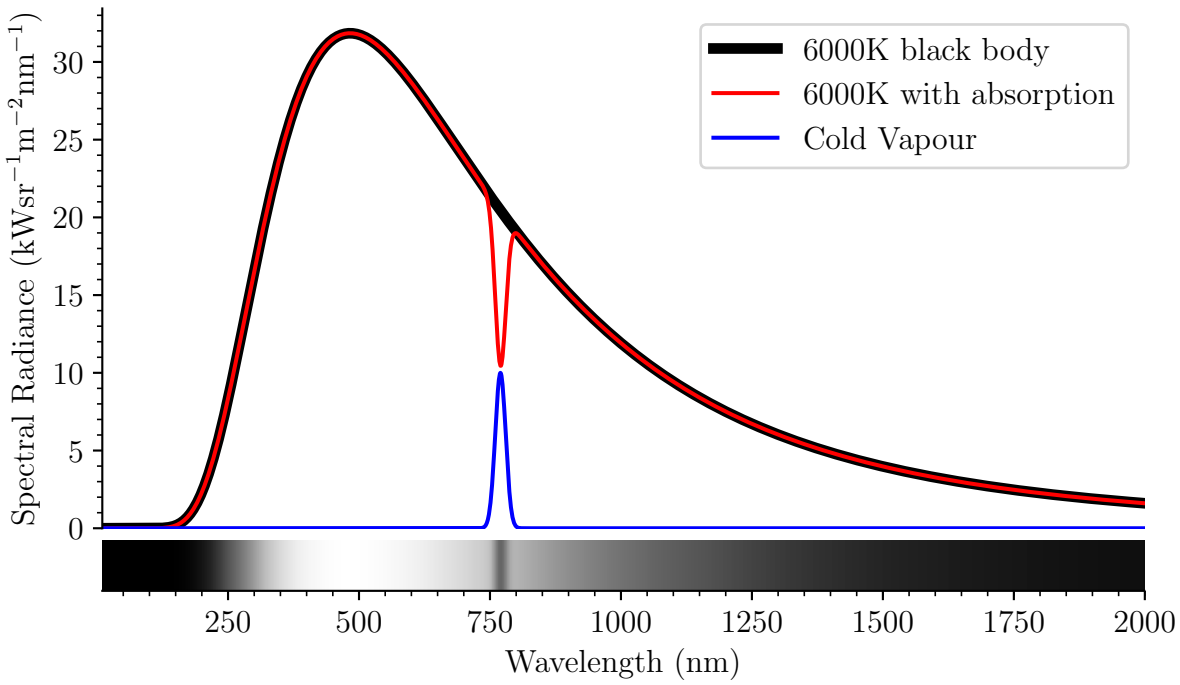

Figure 2.3: When a cooler gas is placed in between a hotter broad-spectrum source and an observer, the gas absorbs light at its characteristic frequency. It then re-emits light isotropically, and since some of the light is now scattered in directions away from the observer a reduction in intensity is seen – an "absorption line" is formed. In this plot, the black line represents a typical emission spectrum from a black body source at a temperature of 6000 K. The blue line is a typical absorption profile of potassium at 770 nm. The red line is the difference between the two, showing how light has been removed from the black body spectrum. The grey-scale at the bottom shows how such a line is observed through a spectrophotometer. Both the amplitude and width of the absorption line have been exaggerated for visibility.

plates was developed as a negative. This meant that when the two plates were superimposed the "differential" was formed. Any intensity fluctuations that were common to both sides were removed, leaving only the strong Doppler signal. The principle of this method is still common in instruments used today, as we shall see later in Chapter 3 on BiSON instrumentation.

Over the course of a few minutes, Leighton's technique would scan across a strip of the Sun, and then move back in the other direction. This meant that plates corresponding to the same section of the Sun could be compared over time. Leighton expected that any motion would have a random appearance, since the origin of the process ought to be stochastic. The lifetime of solar granulation, the columns of convection currents within the

solar plasma, was accepted as a few minutes and so for any location on the solar surface the velocity signal when compared over this time scale should be completely unrelated.

The analysis of the plates did not show this. When looking at any given point on the Sun, the area was oscillating with a characteristic timescale of around five minutes, and velocity amplitudes of order $1000\,\mathrm{m\,s^{-1}}$. The phenomenon soon received independent confirmation involving the entire Sun (Evans and Michard, 1962a,b; Howard, 1962). Solar oscillations had been discovered, marking a new field in solar physics known as helioseismology.

Over the next few years, several theories were suggested for mechanisms that could explain the observations. Many suggested the possibility of pressure waves (so-called p-modes), whilst others considered the possibility of buoyancy waves in the atmosphere due to gravity (so-called g-modes). The now accepted explanation arrived in the early 1970s. Roger Ulrich and John Leibacher independently suggested that the excitation mechanism of the oscillations were beneath the Sun's surface, in the convection zone (Ulrich, 1970; Leibacher and Stein, 1971).

Ulrich and Leibacher suggested that a sound wave generated in the convection zone heading towards the solar interior would be refracted as the temperature increases. Likewise, a sound wave heading towards the solar surface would be returned towards the core due to the decreasing scale height at the top of the convection zone. The combination of refraction and reflection would form a cavity and result in these acoustic pressure waves being trapped in the interior. Certain frequencies of oscillation would constructively interfere with other waves and set up standing wave patterns. These standing wave patterns would be observable in the photosphere as the compressions and rarefactions cause movements of the solar surface. Ulrich predicted that in a diagram of wavenumber against angular frequency a set of "ridges" should be seen, since at each wavenumber there is a set of oscillation modes corresponding to the overtones of the wavenumber.

Several years passed before sufficient quality data had been captured in order to test the theory. Roger Ulrich developed his own observing programme in 1973, working with

Ed Rhodes and George Simon. They used the Diode Array Magnetograph at Sacramento Peak observatory to look at a small square patch of the Sun where each side was approximately one-seventh of the solar diameter – less than one percent of the surface area. Simultaneously and independently, Franz Deubner had also been working on a high-resolution dataset that scanned the Doppler shift along approximately one-fifth of the solar diameter.

Just before Rhodes et al. could complete the analysis of their data, in April 1975 Deubner published his results (Deubner, 1975). Deubner presented a plot showing ten ridges, just as predicted by Ulrich. Deubner had conclusive proof of standing acoustic waves trapped beneath the surface of the Sun, and subsequent confirmation was also published by Rhodes et al. (1977). This brings us neatly back around to the opening statement – the Sun rings like a bell. In the same way that turbulence in a wind instrument creates standing sound waves and thus music, turbulence in the convention zone creates the music of the Sun.

2.2 The Birmingham Contribution

What would become known as the *High-Resolution Optical Spectroscopy* (HiROS) group at the *University of Birmingham* was founded by George Isaak. Isaak held an academic position at Birmingham from 1969 onwards. His initial research interests involved using high-power lasers, tuned to match the centre of an absorption line, as a means of interstellar communication. In developing an atomic beam spectrophotometer suitable for looking for such unusual features in stellar absorption lines, he realised it could also be used to look for exoplanets – planets orbiting different stars. His instrument could detect the "wobble" of the star caused by a sufficiently large orbiting planet.

Such an instrument would need to be very precise in order to be able to measure Doppler shifts in absorption lines of the order of just one metre per second. It would also need excellent long term stability given they would be looking for planets with or-

bital periods extending to several Earth-years. Isaak, and new PhD student John "Bill" Brookes, decided to test their design on the Sun by looking for the pull of large planets such as Jupiter. By observing absorption lines in light from the Sun, they could avoid worrying about the large collection optics required for observing a distant star and instead concentrate on the operation of their spectrophotometer.

Since the instrument was designed for stellar observations, Isaak decided not to resolve an image of the Sun and look at small areas as did Leighton, Ulrich, and Deubner. Instead he observed the "Sun-as-a-star", looking at the average signal from across the whole of the visible solar surface, just as he would have no choice but to do when making stellar observations.

The instrument was tested from the summer of 1973, on the roof of the Poynting Physics building in the School of Physics at Birmingham. We'll look in more detail at the operation of the spectrophotometer later in Chapter 3, but for now it is sufficient to say that the basic measurement was that of the shift in wavelength of an absorption line formed by potassium atoms in the Sun, and the same line formed by potassium atoms in the spectrophotometer. It is basically the same technique as used by Leighton in the early 1960s, except without the need for expensive and time consuming photographic plates.

In early 1974, just over a year before Deubner (1975) and Rhodes et al. (1977) would announce discovery of the high-wavenumber "ridges", Isaak and his team observed globally coherent five-minute oscillations over the whole Sun. These were the low-wavenumber modes of oscillation, formed by sound waves that penetrate the deep interior and core. The very lowest of these, the *radial* or "breathing" modes, involve oscillation of the whole volume of the Sun. This was confirmation of earlier predictions (Wolff, 1972, 1973) that the entire solar volume could oscillate like a variable star.

The key difference between the low-wavenumber and high-wavenumber results is due to the differences in observing technique. By resolving an image of the Sun, and looking at small local areas of the solar surface, an instrument sees only the high-wavenumber modes of oscillation. In contrast, by observing the Sun-as-a-star, the high-wavenumber

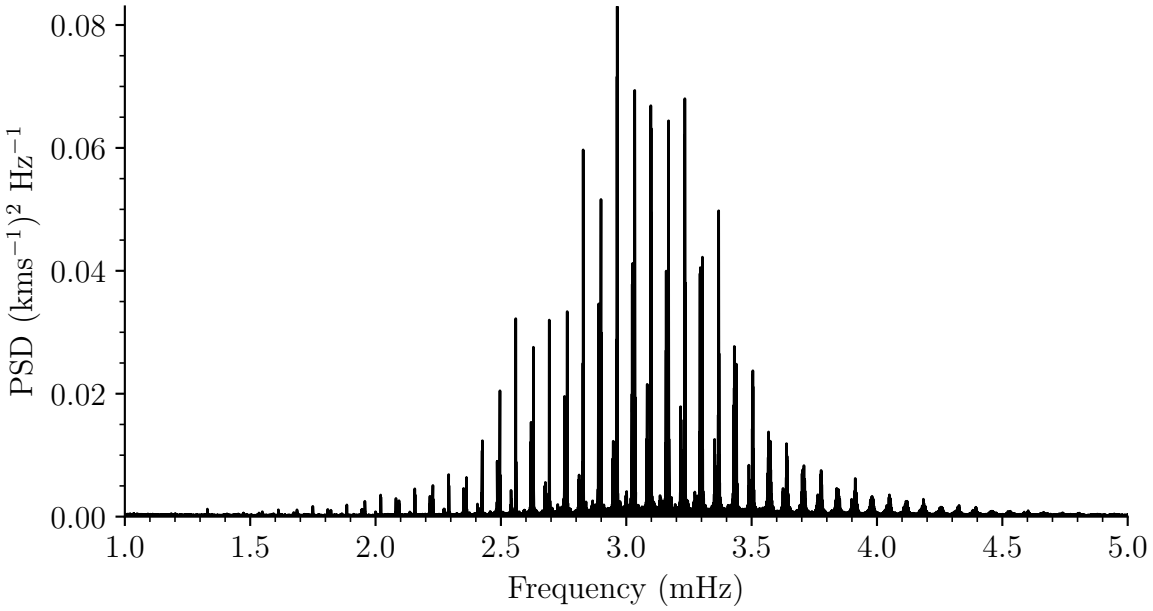

Figure 2.4: Power spectrum produced from 20 years of BiSON data (Hale, 2015a) showing the "picket fence" signature of global modes of oscillation. The dominant spacing is that between pairs of even and odd degree modes, and is around 68 µHz.

modes are "averaged out" over the visible surface, and leave only a negligible signal. The resulting signal is then dominated by the low-wavenumber modes.

With the enthusiasm of their success, Isaak was keen to leave the poor weather of Birmingham and moved the instrument to the Pic-du-Midi observatory in the French Pyrenees. PhD student Brookes observed for several months collecting data over the summer of 1974. With the longer dataset the five-minute envelope in the power spectrum started to become resolved into a collection of narrow peaks. It quickly became essential to collect even more data in order to improve the resolution of the power spectrum. By 1979, PhD student Teo Roca-Cortés was able to make power spectra from data collected over a total of three years. With the improved resolution it was undeniable that the picket-fence-like pattern was genuine, shown in Figure 2.4. The Birmingham team had found proof that the peaks were the signatures of the global modes of oscillation of the whole Sun (Claverie et al., 1979). Observations of what was believed to be the 160 minute fundamental radial "breathing mode" of oscillation (Severnyi et al., 1976; Brookes et al., 1976) unfortunately was determined to be an artefact of the gap-structure introduced by

the daily rotation of the Earth (Elsworth et al., 1989). The 160 minute oscillation is the ninth harmonic of a day, which is to say 24 hours divided by 9. From the early 1980s, the need for continuous observations became obvious. Longer datasets allow mode frequencies to be measured to a much higher precision, since the resolution of a Fourier transform goes as the inverse of the length of the time series. With greater precision comes tighter constraints on the science obtained from the data, such as conditions in the interior of the Sun. It also opens the door for new science, as previously hidden structure becomes resolved.

Observing for longer and longer periods solves only part of the problem. So far, all of the data that had ever been collected was from one site only. This means that even when many days are processed together as one time series, there are still gaps in the dataset. Primarily these are regular breaks due to the diurnal cycle repeating every 24 hours, but also random breaks due to poor weather conditions. These gaps cause problems when processing the Fourier transform of the data. Power gets "spread out" into sidebands around the real mode peaks at spacings determined by the frequency of the gaps, the so-called *window function* of the dataset. If data are collected in four-hour blocks each morning and afternoon, then that is equivalent to a frequency of approximately 69 µHz, and even if data are collected continuously each day, the 24 hour diurnal cycle creates sidebands at approximately 11.6 µHz, both of which are unfortunately close to the mean spacing between mode peaks. This makes the already busy spectra even more complicated and difficult to analyse and interpret. Random gaps due to poor weather have the effect of raising the overall background noise of the spectra, reducing the signal to noise ratio and making modes more difficult to isolate out from the noise.

Clearly, adding together more and more days of data was not going to be sufficient. Whilst it would still provide the necessary improvement in resolution of the power spectrum, the confusion caused by the side bands would make separation and identification of the mode-pairs almost impossible. As well as needing long year-round observations, the data needed to be as near continuous as possible – complete 100 % coverage, with no

sidebands and no extra noise. A different technique was going to be required.

2.3 Observational Facilities

2.3.1 Terrestrial Poles

Whilst the Birmingham group were working at Izaña, Eric Fossat had been working on the solar telescope at the Nice Observatory. Fossat and his team had been making similar Sun-as-a-star observations and had discovered the same problems caused by gaps in the observations. Having moved on in 1977 from Nice to a postdoctoral position in Boulder, Colorado, Fossat spoke with Bernard Jackson who had considerable experience working in Antarctica. At the South Pole, during the summer months the Sun does not set and so it would be possible to make observations over several days with no gaps at all in the data, possibly even weeks if the weather was cooperative.

Through Jackson, Fossat made contact with Martin Pomerantz in early 1978. Pomerantz had visited Antarctica several times and had access to the Bartol Research Foundation which is very near the geographical South Pole. Plans were soon in place to begin an observing programme, and Fossat and his new colleagues arrived in Antarctica in November 1979. The Antarctic conditions make it difficult and costly to maintain observations. Whilst setting up the instrument, temperatures plummeted to $-35\,^{\circ}\mathrm{C}$. Vehicle breakdowns meant they had to start skiing out to the observation site, which was seven kilometres away from the accommodation to avoid any atmospheric pollution from the main site.

Despite the difficulties, an observing run covering the first six days of 1980 was completed successfully, with only two hours missing due to cloud coverage. The results were published in December of 1980 (Grec et al., 1980), and at the time was the longest and most detailed set of continuous observations available. Martin Pomerantz collected a second dataset from the South Pole between 1981 and 1982, this time with Jack Harvey and Tom Duvall who had already had observational successes at the Kitt Peak National

Observatory. Their instrument was the first of its type to provide a "complete" view by detecting modes over a wide range of angular degree.

It was now possible to explain the whole range of five-minute oscillations from low- to high-degree. But observing from the Antarctic is not a permanent solution. The harsh conditions make it a less than ideal working environment, and the long continuous daylight hours are only available for a few months each year. Safe transportation to and from the site limits an observing season to at most three months, from late October to early February. Still another different technique would be required to collect continuous datasets lasting many years.

2.3.2 Space

Observing from space would seem to be the ideal solution. Free from the diurnal cycle and the Earth's atmosphere, a 100 % duty cycle of low-noise data is possible. Several space-based missions have been successful. In the early 1980s the *Active Cavity Radiometer Irradiance Monitor* (ACRIM: Willson, 1979) onboard the NASA *Solar Maximum Mission* spacecraft carried an instrument intended to monitor the total radiation output of the Sun. Its purpose was not specifically designed to monitor solar oscillations, but it was sufficiently precise to detect the small changes in intensity caused by the oscillations. The results showed variations of a few parts per million of the mean intensity, with peaks in the power spectrum corresponding to modes with periods of oscillation of around five minutes (Woodard and Hudson, 1983).

Probably the most successful space-based mission is the *Solar and Heliospheric Observatory* (SOHO), a joint project between ESA and NASA which launched in December 1995 and began normal operations in May 1996 (Domingo et al., 1995). Among the suite of instruments onboard the spacecraft are three helioseismic instruments: *Global Oscillations at Low Frequencies* (GOLF: Gabriel et al., 1995) and *Variability of solar IRradiance and Gravity Oscillations* (VIRGO: Fröhlich et al., 1995) for low-degree oscillations and *Michelson Doppler Imager* (MDI: Scherrer et al., 1995) for observations at medium and

high degree. The spacecraft orbits the "L1" Lagrangian point between the Earth and the Sun, which is the point at which the combined gravitational pull of the two masses provides precisely the centripetal force required to orbit the Sun with the same period as the Earth. This allows the spacecraft to remain stationary with respect to the Earth and provide the continuous view of the Sun required for uninterrupted data capture. The SOHO instruments have achieved duty cycles around 99 %, excluding the periods in 1998 and 1999 when there were serious issues with the spacecraft following the failure of the on-board gyroscope stabilisation system. The mission was originally planned to last for just two years, but is still in operation and currently has a mission extension lasting until December 2020 and has now been in operation for over twenty years.

MDI ceased observations in 2011 when it was superseded by the *Helioseismic and Magnetic Imager* (HMI: Schou et al., 2012) onboard the *Solar Dynamics Observatory* (SDO), but GOLF and VIRGO are still in use. SDO's geosynchronous orbit allows it to downlink a much higher volume of data, including Doppler and intensity images from HMI at high resolution, but its duty cycle is slightly reduced due to its twice-yearly eclipse seasons when the Sun is blocked by the Earth for up to an hour each day for a few weeks.

As with any space-based project, there is not only a huge cost involved, but also years of development time before a mission can be launched. The harsh environment of space means equipment is often prone to failure, and there are very few prospects of repairs being made. This means even the simplest problem can result in a mission being compromised. Mission lifetimes are also limited to how much fuel is carried by a spacecraft, with most typically carrying enough to operate for only a few years.

2.3.3 A Global Network

Constructing a network of mid- to equatorial-latitude observatories is a much more economical method for achieving the goal of a 24-hour view of the Sun. It is also much easier to visit the sites and perform repairs or upgrades. In ideal conditions only three stations would be required, equally separated in longitude, to obtain 100 % duty cycle. This as-

sumes perfect weather and no equipment failure. In the real world a level of redundancy is required, and more sites are needed to compensate for variable weather conditions.

The Birmingham group were the first to begin construction of a network of ground-based observatories, beginning with Izaña in Tenerife in 1975, and culminating in six operational sites in 1992 (Isaak, 1992). We'll look at the Birmingham Solar Oscillations Network (BiSON) in more detail Chapter 4.

Other groups have also had success with ground based networks. Eric Fossat and colleagues went on to deploy the *International Research of Interior of the Sun* (IRIS: Fossat, 1991) network. A prototype was tested at the La Silla Observatory in Chile from the end of 1986 to the middle of 1987, and roll-out of the network began later that same year. The operational strategy of IRIS was different from BiSON, requiring local observers for each instrument as opposed to automation. The instrumentation operated in a similar manner to BiSON, observing spatially unresolved global low angular-degree oscillations using an absorption line of sodium, rather than the potassium line used by BiSON. We discuss the advantages and disadvantages in the use of sodium or potassium in Section 3.2. The IRIS network spanned six sites in total (up to nine sites when considering additional collaboration (Salabert et al., 2002a)), and was operational until 2000 (Salabert et al., 2002b; Fossat and IRIS Group, 2002).

John Leibacher and colleagues later developed the *Global Oscillation Network Group* (GONG: Harvey et al., 1996). Encouraged by the success of the Birmingham group with robotic automation, their aim was to set up a six-site automated network (The GONG Project, 1985) using resolved imaging to observe modes of oscillation at medium degree (up to $l \approx 150$), complementary to the existing low-degree observations made by BiSON. The sites were chosen in early 1991, and deployment began in 1994 with instruments coming online throughout 1995. The GONG network was upgraded in 2001 – 2002 to observe up to around $l = 1000$, and is still in operation today. With full automation and a strategy of regular preventative maintenance visits to the sites, the network has achieved a mean duty cycle of 86 %.

For completeness, there have also been observations made by the *LOWL* project (Tomczyk et al., 1995), which observed at medium degree from one or two sites between 1994 and 2004, and the *Taiwanese Oscillations Network* (TON: Chou et al., 1995) for high-degree observations, which was deployed between 1993 and 1996 and operated for only a few years.

2.4 Notable Results

Helioseismology has had many achievements over the years, both within solar physics and the wider fields of physics. In the mid-1960s, observations of the solar electron neutrino flux returned values that were consistently lower than that predicted by nuclear physics. The fine structure of the solar spectrum measured by BiSON, and the precision with which results were determined, placed strong constraints on the solar models. This suggested that the so-called solar neutrino problem needed to be resolved within neutrino physics, not solar physics (Elsworth et al., 1990), and led to the discovery of neutrino oscillation.

The structure of the Sun, and the rotation rate of the core, has been debated widely. Gough (1986) suggested that the depth of the convection zone could be determined directly from the internal sound speed. Later, Christensen-Dalsgaard et al. (1991) used helioseismic data collected by Duvall et al. (1988) and Libbrecht and Kaufman (1988) to show that the depth of the convection zone is 0.713 ± 0.003 solar radii. Basu and Antia (2004) used data from GONG and MDI to further constrain this value to 0.7133 ± 0.0005 solar radii. More recently, Basu et al. (2009) used BiSON data to obtain sound speed and density profile differences between the Sun and reference models, and found that by having more low-degree mode frequencies available from the long baseline BiSON data the solar models tended to show higher correlation and with lower uncertainties.

It is possible to measure the internal rotation rate of the Sun through helioseismology since, as the Sun rotates, the frequency degeneracy of each oscillation mode is lifted causing the frequency to be rotationally split. Since the modes can be considered to be

probes of different layers within the Sun, this allows the internal rotation profile to be mapped by measuring the rotational splitting of many modes. One of the first results was to show that differential rotation observed on the surface continues throughout the convection zone, and that solid body rotation dominates within the radiative interior and the core (Rhodes and Ulrich, 1977; Gough, 1981; Thompson, 1990). Initial indications were found by Claverie et al. (1981), Duvall and Harvey (1984), and Duvall et al. (1984) that the core was rotating faster than the surface. Later results by Lazrek et al. (1996) using data from the IRIS network, and by Corbard et al. (1998) using data from GOLF and MDI, suggested that the core and surface rotation rates were similar. The long baseline BiSON data again are well placed to allow very precise measurement of the mode splitting down to very low frequencies, providing a more robust estimate of rotation in the deep radiative interior and outer parts of the core (Chaplin et al., 1999), and also allowed consideration of changes in core rotation rate over a solar cycle (Chaplin et al., 1996b).

Possibly the most important recent result from BiSON is continuous monitoring of the solar cycle, particularly in terms of rotation rates since the layer of rotational shear at the boundary between the convective zone and the radiative zone, known as the tachocline, is thought to create the dynamo that drives the solar cycle (Antia et al., 2008). Analysis of helioseismic data from the extended minimum at the end of Cycle 23 going into Cycle 24 showed that there were significant differences between the rotation rates and zonal-flow pattern during the two minima (Howe et al., 2009; Antia and Basu, 2010; Basu and Antia, 2010). Changes in the solar cycle were later validated by an updated analysis of p-mode parameter shifts from BiSON, and it was shown that previously published results on the solar cycle variation of mode parameters were reliable (Howe et al., 2015). Recent results examining the frequency shifts in low-degree modes from BiSON data made comparisons with a number of global activity proxies showed that there has been a change in the frequency response to activity during Cycle 23 suggesting that the Sun may be undergoing a fundamental change in the nature of the solar dynamo (Basu et al., 2012; Howe et al., 2017, 2018).

2.5 Conclusion

In this chapter we covered the discovery of solar oscillations, and the part played by the University of Birmingham as the field emerged in the late 1960s and early 1970s. We discussed several observational facilities and techniques, and showed how the work by the HiROS group fits within the wider field of research.

In the next chapter, we will go on to look in more detail at the instrumentation developed by the HiROS group, and how the solar oscillation data are calibrated.

Chapter 3

BiSON Instrumentation

Contents

3.1 Introduction

The aim of any instrument designed for detecting solar oscillations is to be able convert some measurement of light from the Sun into some measurement of oscillation frequency and power. The oscillations can only be seen by the effect they cause on the very outer layers of the Sun. Any instrument needs to be both very sensitive, and very stable. This is because the signal produced by the oscillations is very small, and the lifetime of modes can be very long. Since any observing programme is potentially going to be running for several years in order to obtain the required data resolution, it is important that the response of an instrument does not change from day to day, or month to month. Here we will discuss instrumentation purely at a qualitative level, giving an overview of operation sufficient to enable understanding of the performance analysis in the next two chapters. A full quantitative discussion begins in Chapter 7.

There are several ways in which the oscillations can affect the solar surface, and so there are a number of possible detection techniques. For oscillation modes that cause

the entire Sun to gently expand and contract, it is possible to observe their influence on the solar radius. There are several problems with this method. The Sun is a gas body with no well defined surface. Observations of the solar radius made from Earth are seriously affected by the atmosphere, making the already "fuzzy" photosphere even less well defined. It is almost impossible to detect the variations of only a few tens of metres in the solar diameter – a displacement of a few parts per million. For this technique to be successful it is necessary to escape Earth's atmosphere by going into space. But once in space, a better technique can be used – variations in luminosity.

As the oscillations travel through the photosphere, a transfer of energy between the wave and the gas is possible. This causes a change in luminosity as the energy transfer causes the gas to heat and cool. As we saw in Chapter 2, this principle was first demonstrated in the early 1980s by the ACRIM instrument on board the NASA Solar Maximum spacecraft (Willson, 1979; Woodard and Hudson, 1983). Observations of the intensity are, as with observations of the solar diameter, seriously affected by the Earth's atmosphere. In order to detect the oscillations using a ground-based instrument, it needs to use a technique that is insensitive to variations in atmospheric transmission. Rather than observing the change in luminosity or size of the Sun, it's possible to measure the oscillations by observing the motion of the solar atmosphere. We saw in Chapter 2 that this was the method used by Robert Leighton in 1962 (Leighton et al., 1962) using the spectroheliograph at Mount Wilson. The photographic technique used by Leighton would be unmanageable for long-term observations, and so there needed to be a way of making similar measurements without the need for manually exposing expensive photographic plates. George Isaak developed a technique by expanding on an instrument he had designed earlier in his career, an Atomic Beam Spectrophotometer (Isaak, 1961). The spectrophotometer worked by using an atomic vapour as a lab-frame reference. This meets the second requirement of long-term stability. The Doppler shift of a light source can be measured by comparing the position of a Fraunhofer line in the emission spectrum of the source with that of the same line in a laboratory atomic reference cell. So long as the atomic vapour is kept under

constant conditions the reference point will be very stable. So-called Resonant Scattering Spectroscopy is the technique that has come to be used in all BiSON instrumentation.

In this chapter we will discuss the basic operation at a qualitative level of a typical BiSON spectrophotometer. In Section 3.2 we consider what it means to produce a vapour. Section 3.3 describes the typical mechanical design and optical layout, and finally Section 3.4 goes on to consider data acquisition and calibration in terms of line-of-sight velocity.

3.2 What is a vapour?

What is a vapour? How is it different from a gas? A vapour is produced when the environmental conditions (temperature, pressure) are such that an element cannot exist in the solid or liquid state, indicated by the line ATC in Figure 3.1. This means that a vapour can be returned to solid or liquid state by increasing the pressure. In comparison, when the temperature is above the critical point, then the material remains in the gas phase no matter how high the pressure is raised. Figure 3.1 shows a typical phase-diagram for a generic pure substance in a closed system, showing the relationship between temperature and pressure and how the different phases interact.

In Figure 3.1, the line between points B and T shows how variations in pressure affect the melting point of a solid. Crossing the line corresponds to melting and freezing, and along the line the solid and liquid phases are in equilibrium. The line between point A and C is the *vapour pressure* curve, where A-T is the vapour pressure of the solid phase, and T-C is the vapour pressure of the liquid phase. Along the vapour pressure curve, the solid and liquid phases are in equilibrium with the vapour phase. The point T is special, since this is the only combination of pressure and temperature where all three phases can exist simultaneously since they are all in equilibrium. This is known as the *triple point*. Point C is also special, since this is the *critical point* above which a material exists as a *super-critical fluid*.

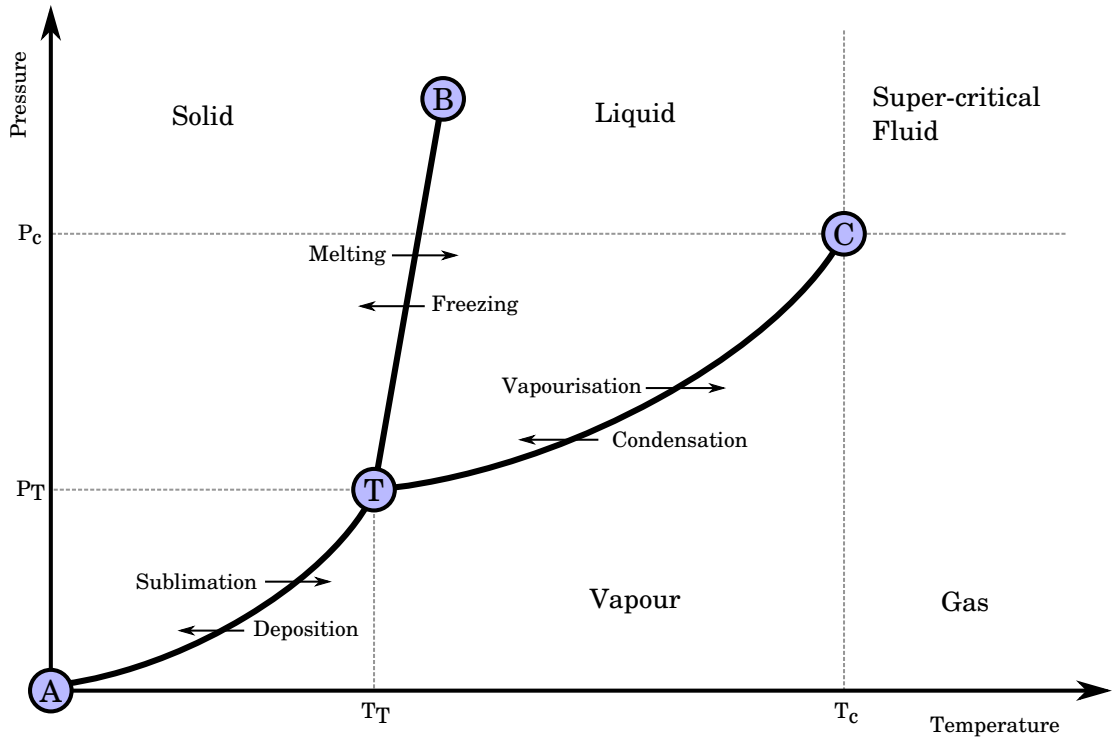

Figure 3.1: A typical phase-diagram for a closed system. Point C is the *critical point*, and T is the *triple point* where the matter can exist in any of the three states simultaneously.

The vapour pressure exists due to the constant motion of the atoms that form a material. When a material is heated, atoms at the surface of the material can acquire enough kinetic energy to escape the material and enter the gaseous phase – *vaporisation* or *sublimation*. A fraction of these atoms will collide with and re-enter the material undergoing *condensation* or *deposition*. As the number of atoms in the vapour phase increases, eventually a state of equilibrium will be reached. A steady state will exist where exactly the same number of atoms are leaving the surface as colliding with the surface. This *equilibrium vapour pressure* is characteristic of a given material at a given temperature, just like its molecular mass or melting and boiling points. It does not depend on the physical amount of material so long as at least a small amount of material is present in equilibrium with the vapour.

In a vapour reference cell, the reference material inside the cell is under vacuum or as close as realistically possible. This means that the temperature does not need to be raised very high before the material sublimates from a solid to a vapour. The *normal*

boiling point of a material is the temperature at which the vapour pressure is equal to standard atmospheric pressure. Since our potassium cell is under vacuum we do not need anywhere near the 759 °C normal boiling point of potassium. BiSON spectrophotometers typically operate around 115 °C.

There are several factors that influence the choice of which Fraunhofer line to use as a reference. Ideally the line will be a good strong feature, and not blended with transitions related to other elements. This also includes contamination from absorption features of Earth origin, so-called *telluric contamination*. For example, the Fe I solar line at 525.0 nm is known to be weakly blended with a telluric water vapour line (Livingston and Wallace, 1985), and this has caused problems for historic solar magnetograph observations taken with Fe I at the Mount Wilson 150 foot tower (Carter et al., 1992). Lines in the visible region were generally preferred, historically, because they were within the sensitivity range of photomultiplier tubes. With newer photodiodes this is no longer an issue, but visible light is still easier and safer to work with in general. Most importantly, it must be possible to actually generate a reference vapour of the same atoms within a spectrophotometer.

With most alkaline metals, a compromise has to be made between ease of producing the vapour, and the abundance of the metal in the solar atmosphere. Historically, most solar spectrophotometers have used either sodium or potassium. The sodium lines are deeper and are visible, but the lines of the D1-D2 doublet are very close together at 0.6 nm making them difficult to separate, and the D2 transition is contaminated with a telluric line. The potassium D1-D2 doublet lines are weaker and in the near infra-red, and similarly the D2 transition is contaminated, but the lines are separated by 3.4 nm making it easier to isolate the clean D1 line. Importantly, the potassium D1 line also has steeper sides making an instrument intrinsically more sensitive since a given shift gives greater change in measured intensity. Sodium has a distinct disadvantage that the metal attacks the cell glass at temperatures required to form the vapour, limiting the life of the cells compared with potassium. A comparison between the two metals in 1989 showed that a potassium vapour cell based spectrophotometer has approximately 2.2 times the sensitivity of a

sodium based system (Isaak et al., 1989). The velocity amplitudes from the sodium instrument were 1.4 times that of potassium as a result of their relative abundances – the sodium atoms being observed at a higher altitude in the photosphere where the density is reduced. Frequencies of individual modes from the two systems agreed to 0.1 µHz (Isaak et al., 1989).

All BiSON instruments use the D1 transition of potassium as the solar reference line, with a central wavelength of 769.898 nm. In the next section, we will discuss the mechanical design and optical layout typically used to build a vapour reference cell into an operational resonant scattering spectrophotometer.

3.3 A Resonance Scattering Spectrophotometer

3.3.1 Mark-I Optical Schematic

The original Mark-I BiSON spectrophotometer is described in detail by Brookes (1974) and Brookes et al. (1978a), and is shown schematically in Figure 3.2. In this section we will look only at the general operation, and we will discuss specific components and specifications later in Chapter 7.

Light from the Sun is first passed through a 15 mm diameter aperture which forms the field stop of the system. A coarse red filter is used to remove light with wavelength shorter than 500 nm and also relieve the thermal load on the subsequent optics by blocking infra-red. An interference filter with a bandwidth of 1.5 nm is used to isolate the potassium D1 Fraunhofer line at 769.898 nm. The light then enters an enclosure where it passes through some polarisation optics which allow wavelength selection, before entering the cell through a 5 mm diameter aperture stop placed immediately in front of the entrance to the vapour cell at a distance of 850 mm from the field stop. Initially lens L1 shown in Figure 3.2 was not part of the design. If we first consider the schematic with lens L1 removed, then these two apertures produce an unvignetted field of view of $40'$ and this is wide enough to pass the full solar disc at $32'$ plus the maximum expected guider error.

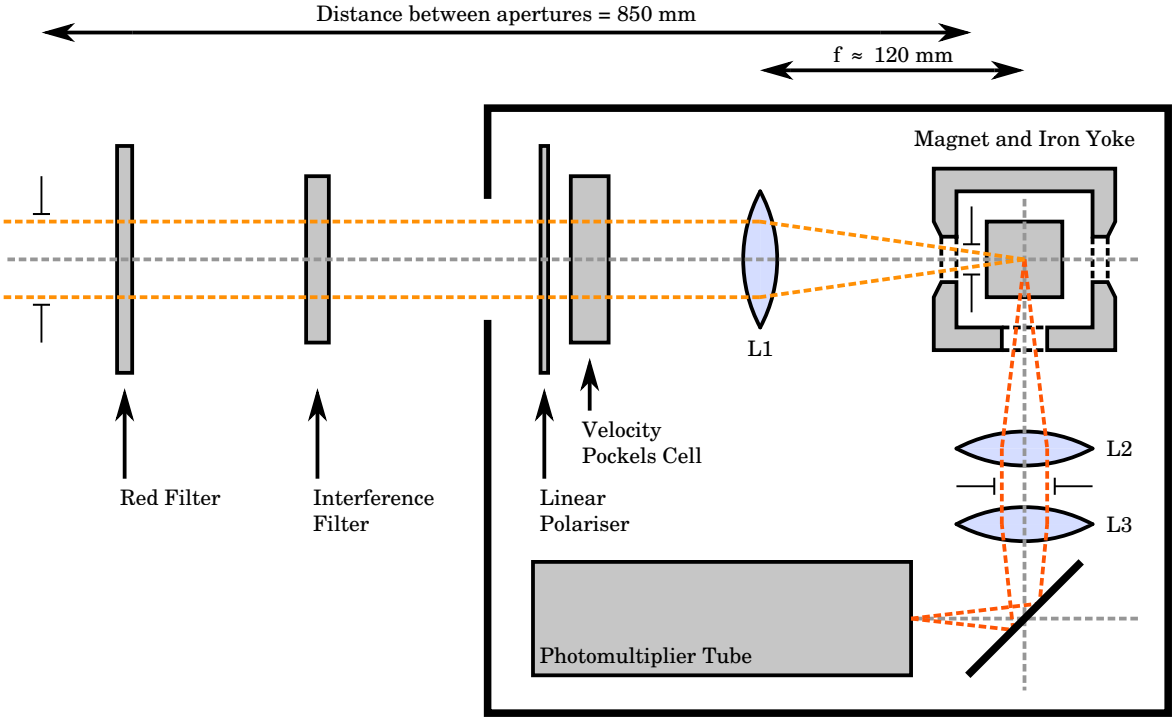

Figure 3.2: Schematic of the Mark-I resonant scattering spectrophotometer.

The addition of lens L1 increases the objective aperture and so allows more light into the cell improving the signal to noise ratio, however it also has the undesirable effect of focusing an image of the Sun into the vapour cell. Imaging the Sun in this way increases sensitivity to guiding errors which cause the image to move inside the cell. Providing the guiding is suitably precise, including the lens produces a net gain in performance and so all current BiSON spectrophotometers make use of imaging optics.

The potassium is heated to around 115 °C to form a vapour. The cell is placed in a longitudinal magnetic field which causes the line to be Zeeman-split into two components, where the separation is dependent on the magnetic field strength. Splitting the lab reference-frame in this way allows measurements to be taken at two working points on the solar absorption line. Measuring intensity changes at the steepest part of the wings of the line provides an improvement in sensitivity over measuring at the line centre. Figure 3.3 shows the splitting produced by a 0.18 T longitudinal magnetic field where the two components are separated by approximately 0.013 nm, or 5.2 km s^{-1} in terms of Doppler velocity. The solar absorption line is modelled as a thermally and rotationally broadened

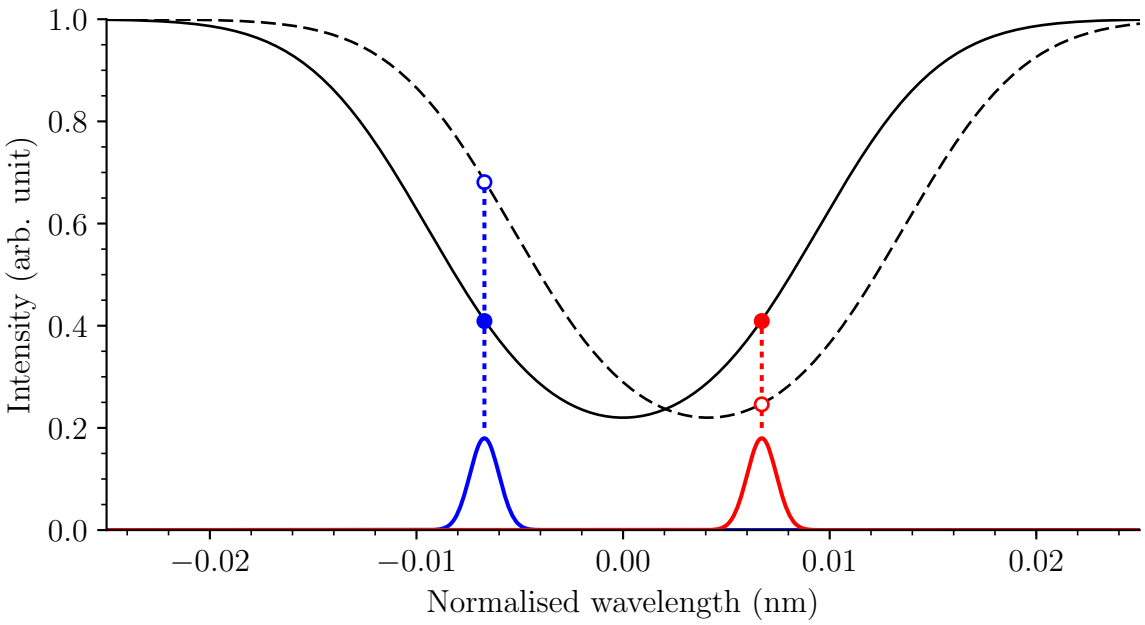

Figure 3.3: The potassium solar absorption line is modelled according to Underhill (1993) and shown with both zero line-of-sight velocity (solid line), and also red-shifted by $1600\,\mathrm{m\,s^{-1}}$ (dashed line) to demonstrate the effect on the scattering intensity. The two absorption components are split by a 0.18 T longitudinal magnetic field producing a separation of approximately 0.013 nm, or $5.2\,\mathrm{km\,s^{-1}}$ in terms of Doppler velocity.

Gaussian, where Underhill (1993) determined from BiSON data that the solar line is well described by,

$$I_\lambda = I_0 e^{-\beta_\lambda} \ , \tag{3.1}$$

where

$$\beta_\lambda = 1.514 e^{\left[-\left(\frac{\lambda-\lambda_0}{92}\right)^2\right]} \ , \tag{3.2}$$

and λ and λ_0 are the wavelength and central wavelength respectively in units of milli-Angstroms. The solar line is shown in Figure 3.3 with both zero line-of-sight velocity (solid line), and also red-shifted by $1600\,\mathrm{m\,s^{-1}}$ (dashed line) to demonstrate the effect on the scattering intensity.

The two components of the Zeeman split absorption line interact only with circularly polarised light. This means it is possible to select which component should be observed at any given time. The polarisation is selected using a linear polariser and a electro-optic Pockels-effect quarter-wave retarder. The polarisation is switched at 0.5 Hz, producing

a 1-s integration period for each if the two polarisation states. By making use of this technique, we can measure the difference between the two wings and form a ratio that reduces the variations in atmospheric transmission and maximises the sensitivity to Doppler velocity signals, just like the two-exposure technique used by Leighton that was seen in Chapter 2.

Light in the selected passband is focused into the centre of the potassium cell and resonantly scattered by the vapour. The vapour radiates isotropically, and some of the light is collected by a lens mounted perpendicular to the optic axis. An aperture in the image plane of this lens removes non-resonantly scattered light from the edges and corners of the glass cell, and then a second lens passes the light onto the photocathode of a photomultiplier tube (PMT) via a short perspex light guide. Pulses from the PMT are counted and logged by the data acquisition system. Light that passes directly through the cell unscattered is collected and logged via a second PMT (not shown in the schematic), and this acts as a measure of the total solar intensity.

3.3.2 Modern RSS Optical Schematic

A modern BiSON resonant scattering spectrophotometer is shown in Figure 3.4. The design is much the same as the Mark-I spectrophotometer, except for three important differences.

Firstly, the aperture stop diameter is increased from 5 mm to 30 mm, greatly increasing the light throughput and thus the signal to noise ratio. A pair of lenses in a Keplerian telescope arrangement is used to reduce the beam diameter, allowing the subsequent optics to be smaller and cheaper. Secondly, the polarisation switching rate of the Pockels cell is increased from 0.5 Hz to 100 Hz, and this again has a large impact on the signal to noise ratio. The characteristic amplitude spectral density of atmospheric scintillation noise has typically decayed to near zero at 100 Hz, shown later in Figure 8.24 on page 217, and so switching at this rate or better significantly reduces the atmospheric effects that are so detrimental to ground-based photometry. Finally, the single scattering PMT is replaced

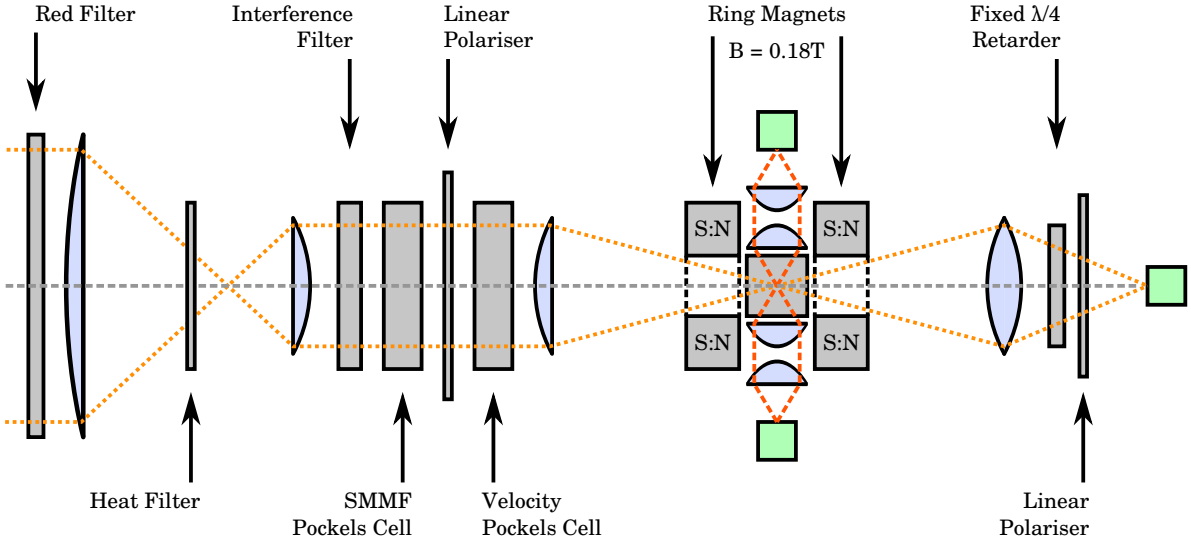

Figure 3.4: Schematic of a typical modern BiSON resonant scattering spectrophotometer.

by two photodiode-based detectors. This not only removes the high-voltage and cooling requirements of a PMT, but doubles the level of detected scattered light enabling yet another signal to noise ratio improvement.

3.3.3 Optional Components

In addition to the optical components required for the primary data product, some BiSON spectrophotometers include modifications to the basic instrumentation that allows additional data products to be captured with little additional effort.

Solar Mean Magnetic Field

The solar potassium absorption line is itself Zeeman split by the solar magnetic field. If a second Pockels cell is installed in front of the linear polariser, providing a total of four polarisation states, then it is possible to isolate the solar left and right circular components and so estimate the value of the solar mean magnetic field (SMMF) (Chaplin et al., 2003). Prior work on the BiSON SMMF data has included searching for correlation between changes in the solar mean magnetic field and coronal mass ejections (Chaplin et al., 2004b, 2005a).

Multi-height Absorption Profiles

Since only a very small bandwidth of light is resonantly scattered by the vapour cell, it is possible to include a second vapour cell and oven with a larger magnetic field strength working at a different point on the solar absorption line. Since different working points on the absorption line correspond to absorption at different heights in the solar atmosphere, this enables a single instrument to log multiple data channels with uncorrelated noise levels (Lund et al., 2017). The first spectrophotometer of this so-called "2B" type was designed by Lewis (1996) as a PhD project.

In the next section we will look at the data acquisition system, and how the raw data are calibrated in terms of line-of-sight velocity.

3.4 Data Acquisition and Calibration

3.4.1 Data Acquisition

The scattered light falls on photodiode detectors, with one on each side of the vapour cell to increase the number of captured photons and so improve the noise level. The detector output voltage is passed through a voltage-to-frequency converter to produce a pulse train. This digital signal is more robust than the delicate voltage signal, and is more suitable for sending down a long cable from the dome to the main electronics in the room below. The pulses are counted by a set of scalars that are controlled by a central gating system which ensures that the pulses are separated based on the different states of operation.

The pulses from the photodetectors are typically integrated for 3.2-s, a total of 1.6-s per wing, and then acquisition is paused for 0.8-s while the data are read out by the computer. The combined total making a 4-s readout cycle, which is subsequently written to disk. When using old dialup modems, it could take several hours to transfer a single day of data back to Birmingham. For this reason, at the end of each day the data were compacted down to a 40-s cadence, and these smaller files transferred and stored in

Birmingham.

Newer systems count for longer and have reduced dead-time since they are capable of faster readout, and this provides another improvement in noise level. Now that all sites have broadband internet connections and storage capacity is inexpensive, the raw counts are stored at the full 4-s cadence. The data are still typically compacted to 40-s cadence for subsequent analysis.

3.4.2 Data Calibration

A typical day of Doppler velocity data from a BiSON spectrophotometer is shown in Figure 3.5. The calibration of Doppler shift to velocity is discussed by Elsworth et al. (1995), with an improved data pipeline that includes correction for atmospheric differential extinction by Davies et al. (2014). The subsequent concatenation of data from the sites is discussed by Chaplin et al. (1997) and by the author (Hale, 2003). The concatenation algorithm makes use of Allen's interval algebra (Allen, 1983) for the required temporal reasoning when handling overlapping time periods within many different datasets. Here, we will discuss the initial calibration from Doppler shift to velocity.

The line-of-sight velocity between a source and an observer, v, is given by

$$v = c\,\frac{\delta\lambda}{\lambda}\ , \tag{3.3}$$

where λ is the rest wavelength, $\delta\lambda$ is the shift in wavelength, and c is the speed of light. Receding velocities are defined as positive, by convention. We first form a ratio R from the two instrumental passbands,

$$R = \frac{I_\mathrm{b} - I_\mathrm{r}}{I_\mathrm{b} + I_\mathrm{r}}\ , \tag{3.4}$$

where I_b and I_r are the intensities measured at the blue and red wings of the solar absorption line, respectively. Assuming the line is symmetric then the observed line-of-

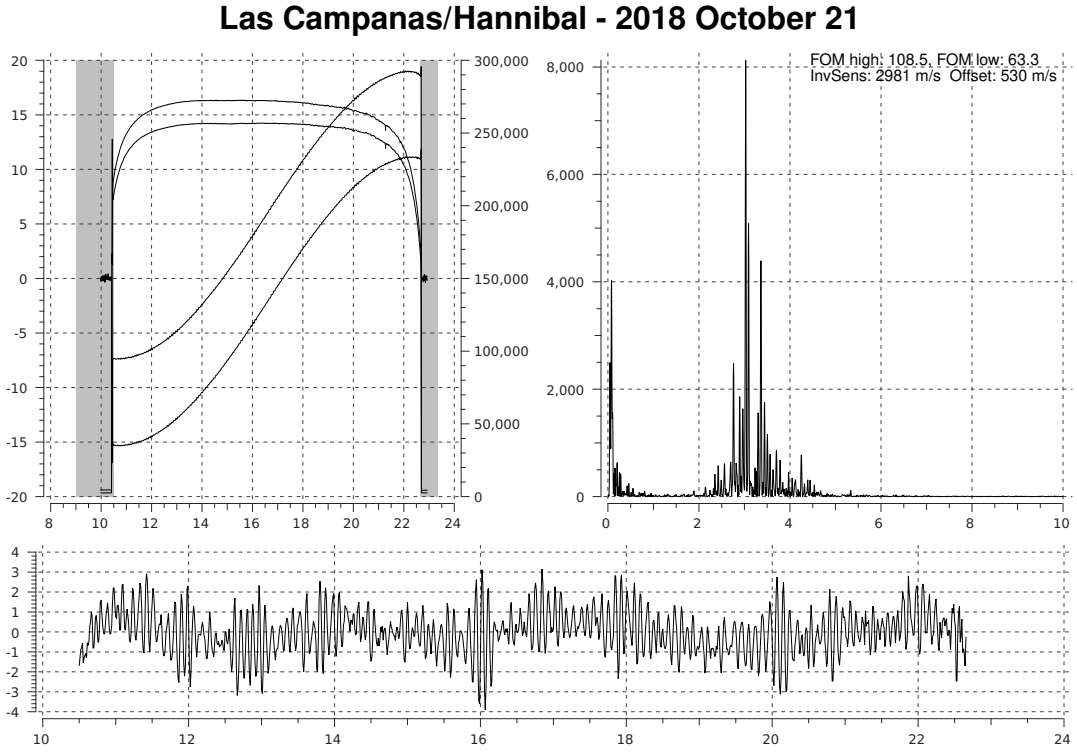

Figure 3.5: Typical daily data from a BiSON spectrophotometer after processing through the standard data pipeline. The top-left panel shows the raw data. The two S-shaped curves are the measured line-of-sight Doppler ratio output from the two scattering detectors, where the almost sinusoidal shape is dominated by the rotation of Earth. The two upper traces show the corresponding sum of the two spectroscopic measurements from each detector, and these can be considered to be a proxy for the total solar intensity throughout the day. The lower panel shows the calibrated velocity data after compensating for Earth's rotation and orbital velocity, and other offsets. The top-right panel shows the power spectral density of the daily velocity residuals.

sight velocity, $v_{\rm obs}$, is proportional to the ratio,

$$v_{\rm obs} \propto R \ , \tag{3.5}$$

and so,

$$v_{\rm obs} \propto \frac{I_{\rm b} - I_{\rm r}}{I_{\rm b} + I_{\rm r}} \ , \tag{3.6}$$

and the constant of proportionality can be found by making use of the known velocity between the Sun and the Earth to produce a daily calibration. Using the ratio in this way rather than raw intensity allows variations due to atmospheric effects to be significantly reduced, and this greatly improves the long term stability of the instrumentation.

The constant of proportionality is known as the *velocity sensitivity* and is more typically referred to as its inverse which is of the order of $3000\,\mathrm{m\,s^{-1}}$ per unit ratio.

The line-of-sight velocity v_{obs} is made up of several individual components (Brookes et al., 1976), each contributing to the total observed velocity,

$$v_{\mathrm{obs}} = v_{\mathrm{orb}} + v_{\mathrm{spin}} + v_{\mathrm{grs}} + v_{\mathrm{osc}} + v_{\mathrm{other}} \ , \tag{3.7}$$

where v_{orb} is the line-of-sight velocity due to Earth's orbit, v_{spin} the line-of-sight velocity due to Earth's spin, v_{grs} the gravitational redshift, and v_{osc} the velocity caused by the solar oscillations. The v_{other} term is used to include any line-of-sight velocity effects caused by instrumental noise and other factors. We will now look at each of these terms individually in the following subsections.

3.4.3 Earth's Orbit – v_{orb}

If the orbit of the Earth around the Sun were perfectly circular, then there would be no line-of-sight velocity component due to the orbit. Since there is a slight eccentricity in the orbit, a line-of-sight velocity effect is produced. The magnitude of this value is also affected by the Moon and the larger planets in the solar system, notably Jupiter. The value can be looked up in a standard Almanac, and BiSON uses the JPL ephemeris. The maximum and minimum line-of-sight orbital velocities of $\pm 500\,\mathrm{m\,s^{-1}}$ occur in April and October, respectively.

3.4.4 Earth's Spin – v_{spin}

The rotation of the Earth produces a signal that dominates the daily data. It creates a familiar sinusoidal shape, where the observer on Earth is approaching the Sun in the morning making v_{spin} negative, zero at local noon, and positive in the afternoon when the observer is moving away from the Sun. The line-of-sight spin velocity is determined by both the observer's altitude and latitude, as well as the declination of the Sun. It is

described by an equation of the form (Brookes et al., 1976),

$$v_{\text{spin}} = \omega_E R_E(\lambda) \cos\lambda \cos\delta \sin\left[\frac{\pi}{12}\,(t - t_0)\right] \ , \tag{3.8}$$

where ω_E is the angular velocity of the Earth, $R_E(\lambda)$ the observer's distance from the centre of the Earth, λ the observer's latitude, t_0 the time of local solar noon in hours, t the time of observation in hours, and δ the declination of the Sun.

The maximum effect of v_{spin} occurs for an observer on the equator, where the range in value at mean sea level is approximately $\pm 460\,\text{m}\,\text{s}^{-1}$.

3.4.5 Gravitational Redshift – v_{grs}

The gravitational redshift is an effect caused by a photon doing work to escape the gravitational field of the Sun. The Planck-Einstein relation describes the energy of a photon as,

$$E = \frac{hc}{\lambda} \ , \tag{3.9}$$

where h is Planck's constant, c the speed of light, and λ the wavelength. Since both h and c are constant, a reduction in energy by doing work can only manifest as an increase in wavelength.

The effect can also be considered to be a result of gravitational time dilation. As the wave moves away from the source of the gravitational field, the rate at which time passes increases. Since frequency is the inverse of time, the frequency of the wave is reduced in a lower strength gravitational field.

The gravitational redshift is given by (Brookes et al., 1976),

$$v_{\text{grs}} = \frac{GM_\odot}{c}\left(\frac{1}{R_\odot} - \frac{1}{d}\right) - \frac{GM_E}{cR_E} \ , \tag{3.10}$$

where $M_\odot$, M_E. $R_\odot$, and R_E are the Sun and Earth masses and radii respectively, and d is the distance from the Earth to the Sun. The gravitational constant is G, and c the speed of light. For a body of the size and mass of the Sun, the gravitational redshift corresponds to an equivalent velocity of $632\,\text{m}\,\text{s}^{-1}$. This value is always present and is

essentially constant.

3.4.6 Solar Oscillations – $v_{\rm osc}$

Measuring $v_{\rm osc}$ is the mission objective. The value of $v_{\rm osc}$ includes a large number of modes of oscillation. When averaged over the whole Sun, the high-order modes will be averaged out through superposition and so this means only the low-order, most penetrating, modes contribute to the signal. $v_{\rm osc}$ is typically just a few metres per second, and this is a small fraction of the total observed signal – around 0.1 %. The oscillations can be regarded as the residuals to a fit of the known parameters to the ratio data, and so they are known colloquially as simply *the residuals*.

In addition to the oscillations this value also contains components due to granulation and other solar activity, such as the convective blueshift caused by hotter and brighter material rising and cooler, dimmer, material settling away into the solar interior.

3.4.7 Other – $v_{\rm other}$

The $v_{\rm other}$ component comprises of solar and terrestrial effects, but is predominantly instrumental noise. A detailed study of noise levels in resonant scattering spectrophotometers was undertaken by Hoyng (1989), and later by Chaplin et al. (2005b). We will look at contributions to noise in resonant scattering spectrophotometers in Chapter 8 starting on page 179.

3.4.8 Ratio-Velocity Calibration

The standard calibration is performed as ratio against velocity, where the computed observer velocity is expanded as a polynomial in velocity (Elsworth et al., 1995). A linear

least-squares fit is performed to the ratio such that,

$$\begin{aligned} R(v_{\rm obs}) &= \frac{I_{\rm b} - I_{\rm r}}{I_{\rm b} + I_{\rm r}} \\ &= \sum_{i=0}^{3} a_i (v_{\rm orb} + v_{\rm spin} - v_{\rm grs})^i \end{aligned} \tag{3.11}$$

where a_i and $v_{\rm grs}$ are the fit coefficients, and $v_{\rm orb}$ and $v_{\rm spin}$ are determined as described previously. Historically a second-order function was used for calibration, but this was later updated to a third-order function in order to help reduce low-frequency offsets (Miller, 2003b). Once fitted the constant of proportionality converting ratio to velocity, known as the *velocity sensitivity*, is then $\mathrm{d}R/\mathrm{d}v_{\rm obs}$ which is defined as

$$\frac{\mathrm{d}R}{\mathrm{d}v_{\rm obs}} = \sum_{i=1}^{3} i \; a_i (v_{\rm orb} + v_{\rm spin} - v_{\rm grs})^{i-1} \; , \tag{3.12}$$

and this is more commonly discussed in terms of the *inverse sensitivity*, $\mathrm{d}v_{\rm obs}/\mathrm{d}R$, which is typically of the order of $3000\,\mathrm{m\,s^{-1}}$ per unit ratio.

The final two terms from equation 3.7, $v_{\rm osc}$ and $v_{\rm other}$, are determined by finding the residuals to the fit and multiplying by the inverse sensitivity, thus

$$v_{\rm osc} + v_{\rm other} = \frac{R(v_{\rm obs}) - \sum_{i=0}^{3} a_i (v_{\rm orb} + v_{\rm spin} - v_{\rm grs})^i}{\sum_{i=1}^{3} i \; a_i (v_{\rm orb} + v_{\rm spin} - v_{\rm grs})^{i-1}} \; . \tag{3.13}$$

Fitting a polynomial in this way can be considered to be a type of high-pass filter. It is possible to determine a basic on-the-fly calibration by simply taking a running-mean of the ratio and multiplying by an estimate of the inverse sensitivity.

More recent updates to the calibration pipeline include additional fit parameters used to compensate for atmospheric differential extinction (Davies et al., 2014), and this has allowed data to be used from much larger zenith angles.

3.5 Conclusion

In this chapter we have presented a qualitative description of the operation of a BiSON resonant scattering spectrophotometer, and of how the data are calibrated in terms of

line-of-sight velocity.

In the next chapter we will look in more detail at the commissioning and operational history of each BiSON site, and also produce quantitative performance statistics for each site in terms of operational duty cycle and instrumentation noise levels. These quality metrics will become essential when comparing the performance of instrumentation developed during this project with the existing network infrastructure.

Chapter 4

Network Performance

Contents

4.1 Introduction

This chapter includes text and figures from the publication "Performance of the *Birmingham Solar-Oscillations Network* (BiSON)" (Hale et al., 2016) in the journal *Solar Physics*. The published article was first released on 2015 December 7, and is included in Appendix D for reference. Figures have been refreshed to include additional data from 2015 up to 2018, and so no longer match the published versions. The results were also presented at the Royal Astronomical Society *National Astronomy Meeting* on 2015 July 9 during the session on *The Variable Sun: Cycles, Waves, Oscillations, and Instabilities* (Hale, 2015f). Data from this publication can be downloaded from the University of Birmingham eData archive (Hale, 2015a). Whilst this publication is a result of collaborative working, the content is entirely the author's own work. Co-authors provided guidance on historical accuracy, data processing, and manuscript review.

We saw in Chapter 2 that the High-Resolution Optical Spectroscopy group at the University of Birmingham were the first to begin construction of a network of ground-based observatories dedicated to helioseismology. The early years of the network have been outlined by Chaplin et al. (1996a). In this chapter we will look back over the history of the network, the challenges that have been encountered, and reflect on what we have learned

from almost three decades of experience in operating a semi-automated ground-based observing network.

The meta-data from the sites have been analysed to assess performance in terms of site insolation, and from this we will present updated results on the temporal coverage and noise performance of the individual sites and the network as a whole. In brief, early campaign-style observations from two sites (Haleakala and Izaña) were followed by the addition of an automated observing station in Carnarvon, Western Australia, and later by the deployment of more standardised stations in Sutherland, South Africa, Las Campanas, Chile, and Narrabri, Australia, over the period 1990 – 1992. The instrument from Haleakala was moved to California and installed in the 60 foot tower at the Mount Wilson (Hale) Observatory in 1992, forming the final network configuration in use today. Since then, with occasional instrument upgrades, the network has been operating continuously, providing unresolved-Sun helioseismic observations with an average annual duty cycle of about 82 %.

Before we can discuss site performance and data quality, we must first define some robust quality metrics in order to allow valid comparisons to be made. In Section 4.2 we will look at the two standard metrics used by BiSON, and then in Section 4.3 go on to apply the metrics to data from each site from 1995 onwards. Finally in Section 4.4 we will consider the performance of the network as a whole, making use of concatenated data from multiple sites and instruments.

4.2 Data Quality Metrics

There are two standard metrics used by BiSON. These are the five-minute Figure of Merit known simply as the FOM, and the mean high-frequency noise level. Data from BiSON are collected on a 40-s cadence, giving an upper limit in the frequency domain (Nyquist frequency) of 12.5 mHz. The FOM is a signal-to-noise ratio, and is defined as the mean power in the main "five-minute" signal band (2 mHz to 5 mHz) divided by the mean power

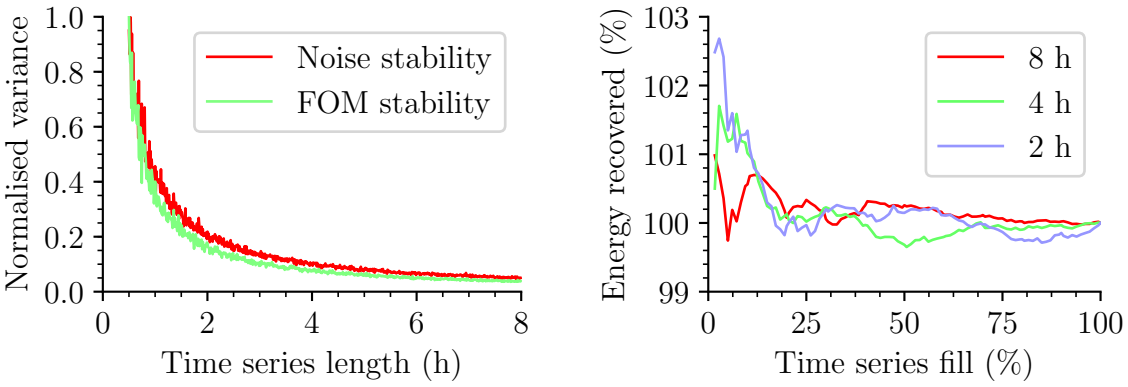

(a) FOM stability

(b) Fill stability

Figure 4.1: Estimated FFT stability over (a) varying dataset length and (b) fill. For a valid comparison between datasets from different days and different sites, a time series needs to be at least 3-h in duration and have a fill of at least 25 %.

in the high-frequency noise (5.5 mHz to 12.5 mHz). When considering only the mean noise level, and not the signal-to-noise ratio, we take the mean power in the upper 10.0 mHz to 12.5 mHz region and this is assumed to be white. The conversion from the time domain to the frequency domain is calculated using the standard Cooley–Turkey discrete Fourier transform (Cooley and Turkey, 1965) as implemented in the Python NumPy library. We will refer to these definitions of FOM and mean noise throughout this chapter.

Both of these metrics require a certain volume of data in order to ensure that the quality estimates are reliable, and that meaningful comparisons can be made between data from different days and different instruments. As the length of a dataset decreases, different realisations of random noise begin to have a larger impact on the estimate of quality. Eventually the quality estimate becomes so uncertain as to be meaningless. In order to determine the minimum dataset length required to make reliable comparisons, a simple artificial dataset was produced with one thousand realisations of random noise. The FOM and mean noise were calculated for each realisation, and the variance in the FOM and mean noise recorded. This was done for dataset lengths varying from thirty minutes up to eight hours. The normalised results are shown in Figure 4.1(a). Both the quality metrics appear to stabilise at dataset lengths of a minimum of three hours,

and so this was selected as the minimum length that could be reliably used for quality comparison.

An additional problem with data from BiSON is that due to variable weather conditions it is usually not continuous. In order to be able to compare absolute noise level we need to rescale the power spectrum produced by the FFT to compensate for any missing data. This is done by simply dividing by the percentage fill (i.e., the percentage of the total observing time where data are available). To determine how low the fill can be whilst still providing a meaningful estimate of the total energy in an equivalent gap-free dataset, a similar test was used with simple artificial data and one thousand realisations of random noise. A variable size gap was created in three datasets of two, four, and eight hours in total length. The results for percentage of total energy recovered compared with the original time series are shown in Figure 4.1(b). Even with very low fills it is possible to recover an estimate of the original total energy to within a few percent. A fill of at least 25 % was selected as being the minimum requirement. In the next section we will apply these metrics and limits to daily data from each site from 1995 onwards. We begin with Izaña, Tenerife.

4.3 Site History and Performance

4.3.1 Izaña, Tenerife

The HiROS group have built and deployed many instruments, including a total of 13 resonance-scattering solar spectrophotometers. The instrument that was to become the first "node" of BiSON was called the Mark-I (Miller, 1998a) and was installed in 1975 at the Observatorio del Teide, Tenerife, shown in Figure 4.2.

Mark-I is housed in the "Solar Pyramid" run by Pere Pallé and his team of observers (Roca Cortés and Pallé, 2014). Light is collected via two flat mirrors, known as a cœlostat. The beam is projected through an open window into the apex of the pyramid, where Mark-I sits on an optical bench. At the beginning of each observing session,

Figure 4.2: Observatorio del Teide, Izaña, Tenerife, Canary Islands
Instituto de Astrofísica de Canarias
Location: -16.50E +28.30N Elevation: 2368 m Established: 1975
Image credit: Google Maps

the on-site observer uncovers and aligns the mirrors and starts the guiding system. Operator presence is also required throughout the day to close the mirrors in the event of bad weather, and at the end of the observing session. The pyramid is shown in Figure 4.3, and the Mark-I spectrophotometer is shown in Figure 4.4.

The original control system was based around a 40-channel scalar module used for counting pulses from a photomultiplier tube (McLeod, 2002). Timing was controlled by a quartz clock producing pulses at precisely 1-s intervals, with the absolute time related to GMT once per day via a time check on a local radio station. At the end of each interval two relays would change state and reverse the voltage across an electro-optic modulator, and the data gate would be incremented to the next channel of the scalar. Once all 40 channels of the scalar had been filled the contents would be destructively written out to magnetic tape requiring a further 2-s. The process would then repeat, with each block of data separated by 42-s. The cadence was changed to 40-s at the beginning of 1990. This allows for simpler concatenation of data from different sites since there are an integer multiple of 40-s in a day, and so it provides a network time-standard.

(a) The Solar Pyramid

(b) Cœlostat Optics

Figure 4.3: The Solar Pyramid at Izaña. (a) shows the secondary mirror mounted in the summer position. In (b) the secondary mirror is fixed to the winter mount.

The system was computerised in 1984 using the BBC Microcomputer. The BBC Micro was originally commissioned on behalf of the British Broadcasting Corporation as part of their Computer Literacy Project, and it was designed and built by the Acorn Computer company. Despite the BBC Micro being discontinued in 1994, the spectrophotometer continued operating in this configuration until 2003 when the computer finally failed and was replaced by a modern PC. A dedicated microcontroller-based interface was designed to enable the PC to communicate with the original scalar system (Barnes et al., 2003, 2004), and this remains the operating configuration today.

Despite the higher photon shot noise level compared to our more modern sites, Izaña has been and still is a work-horse of the BiSON network. The duty cycle and data window-functions are shown in Figure 4.5 and 4.6. The duty cycle is plotted in terms of both number of observational hours, and also percentage insolation. The insolation compares only the potential daylight hours against actual observational hours, and does not differentiate between poor weather conditions and any periods of instrumental failure. Tenerife provides exceptional weather conditions over the summer months, but is rather poor throughout the winter where the conditions become much more variable. The regular "holes" in the window function at midday during the spring and autumn months are caused by cœlostat shadowing. The secondary mirror has two mounting positions, from

(a) Mark-I front

(b) Mark-I rear

Figure 4.4: The Mark-I spectrophotometer in Izaña. (a) shows the spectrophotometer entrance aperture and auxiliary guiding telescopes. (b) shows the rear of the spectrophotometer with the cover removed.

above or from below, which allows an unobstructed daily run to be achieved throughout the winter and summer months. However, during the change-over period between the two phases the secondary mirror unavoidably shadows the primary. The FOM and mean noise levels for Izaña are shown in Figure 4.7. The noise levels from Izaña show remarkable stability and offer an unprecedented temporal baseline of over 40 years.

Throughout 1976, George Isaak, the then head of the HiROS research group, visited several institutions around the United States looking for somewhere suitable to establish a second site. Isaak was already considering the development of a permanent network to expand beyond Tenerife long before his seminal article in 1979 on global studies of the five-minute oscillation (Claverie et al., 1979). The group began operating on the island of Haleakala in Hawaii, at the Mees Observatory, in 1981. The instrument at Haleakala, the Mark-III, operated in much the same way as Mark-I in Tenerife. From the two sites together, data were collected for up to 22 hours per day for 88 days, producing the longest time series, and most highly resolved power spectrum, that had been achieved up to that time.

When planning the expansion to a global network the group realised that it would not be practical to operate such a network manually. The key to a reliable network would be

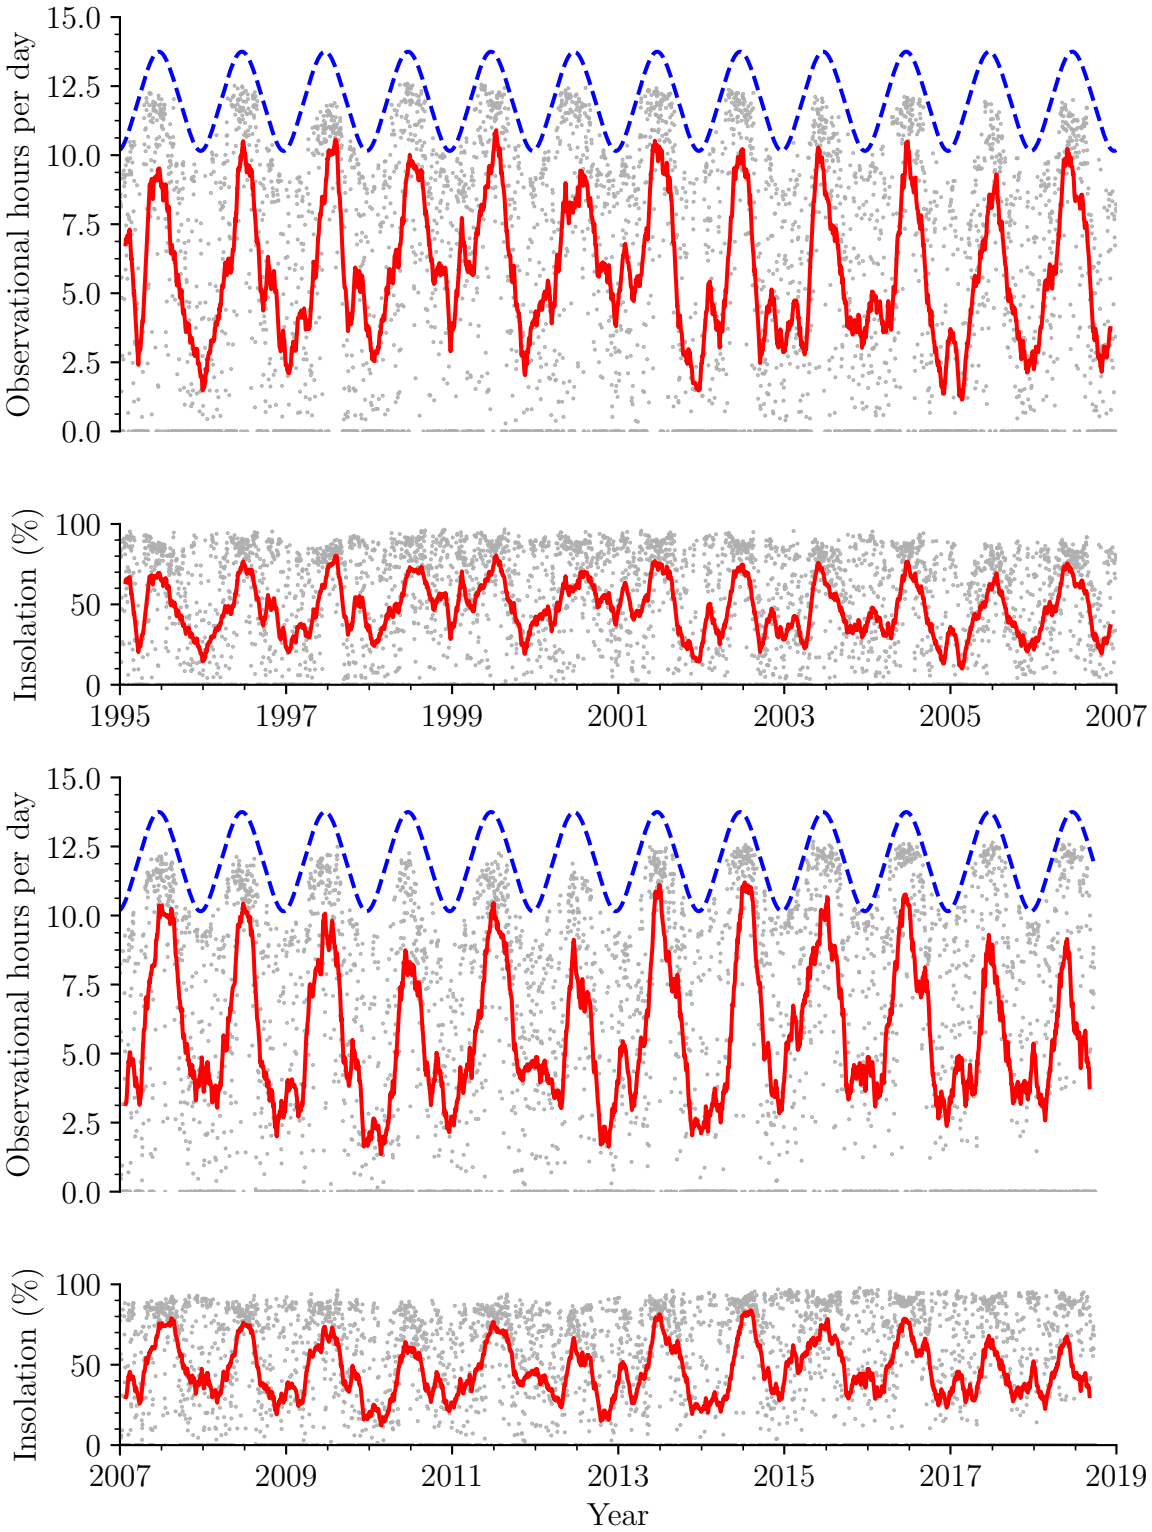

Figure 4.5: Izaña duty cycle as a function of date, plotted in hours per day, and as a percentage of potential daylight hours. There is one grey dot per day, and the solid red curve represents a 50-day moving mean. The dashed blue curve shows potential daylight hours.

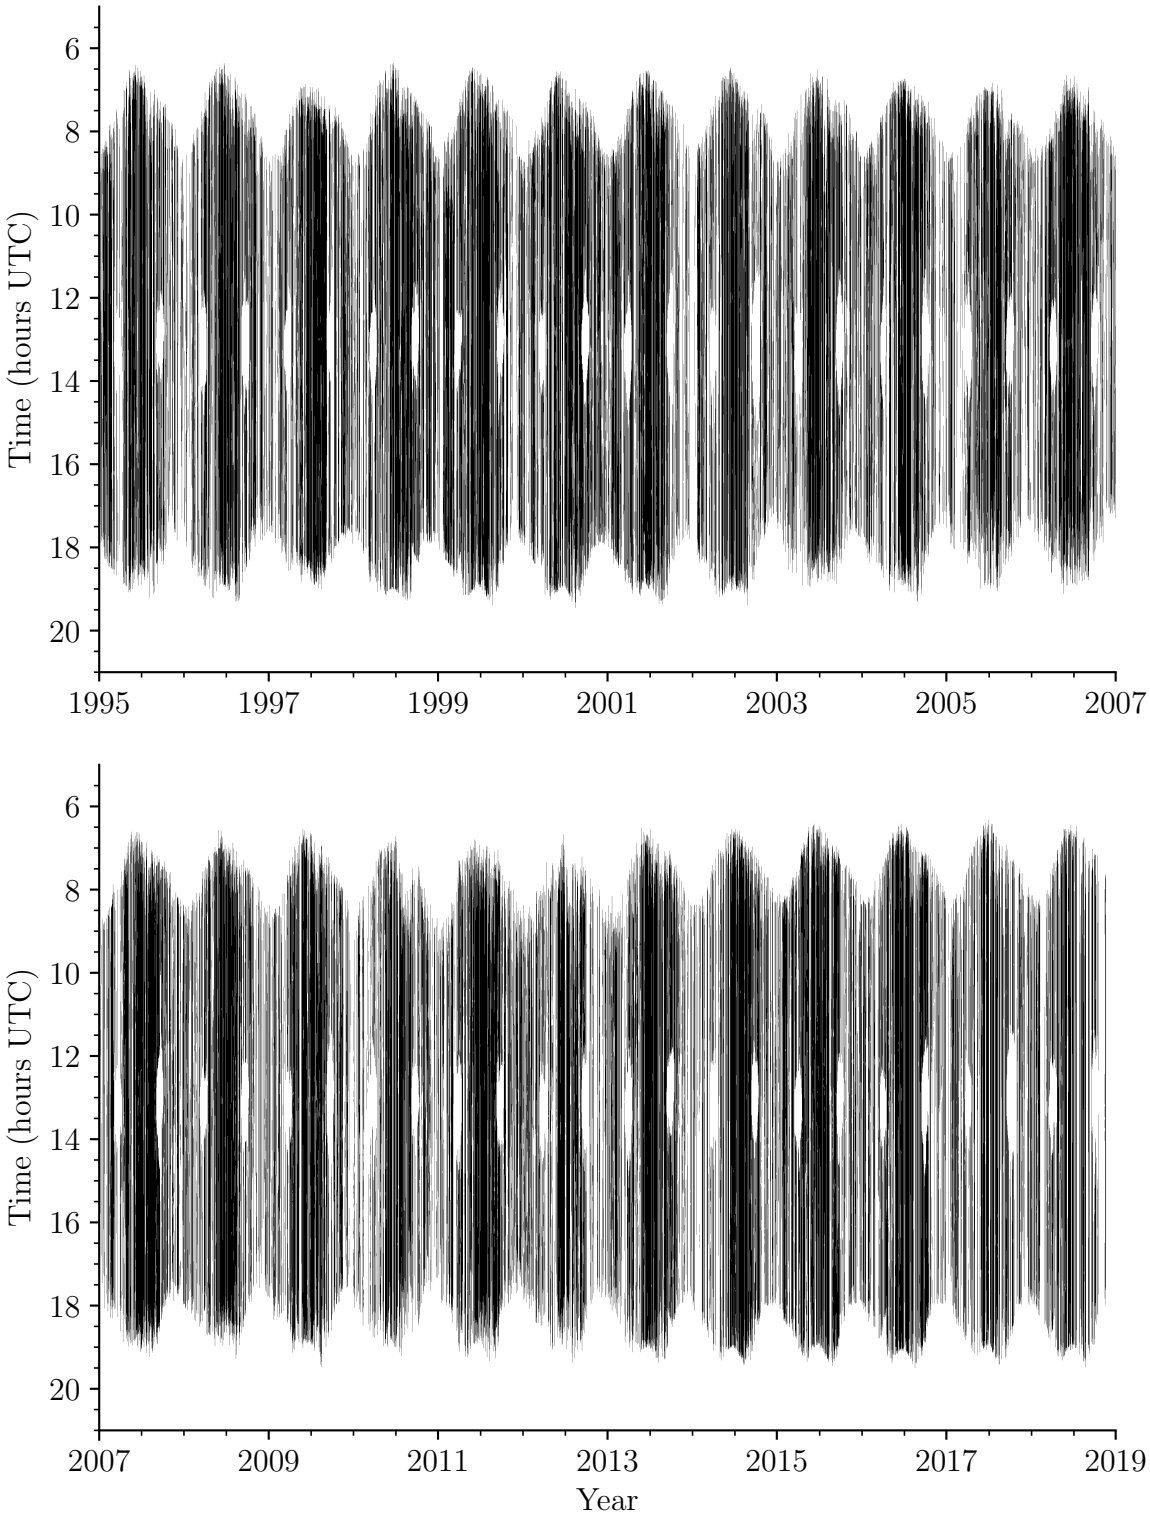

Figure 4.6: Izaña data window-function. The regular "holes" in the data around midday each spring and autumn are due to the cœlostat secondary mirror shadowing the primary mirror.

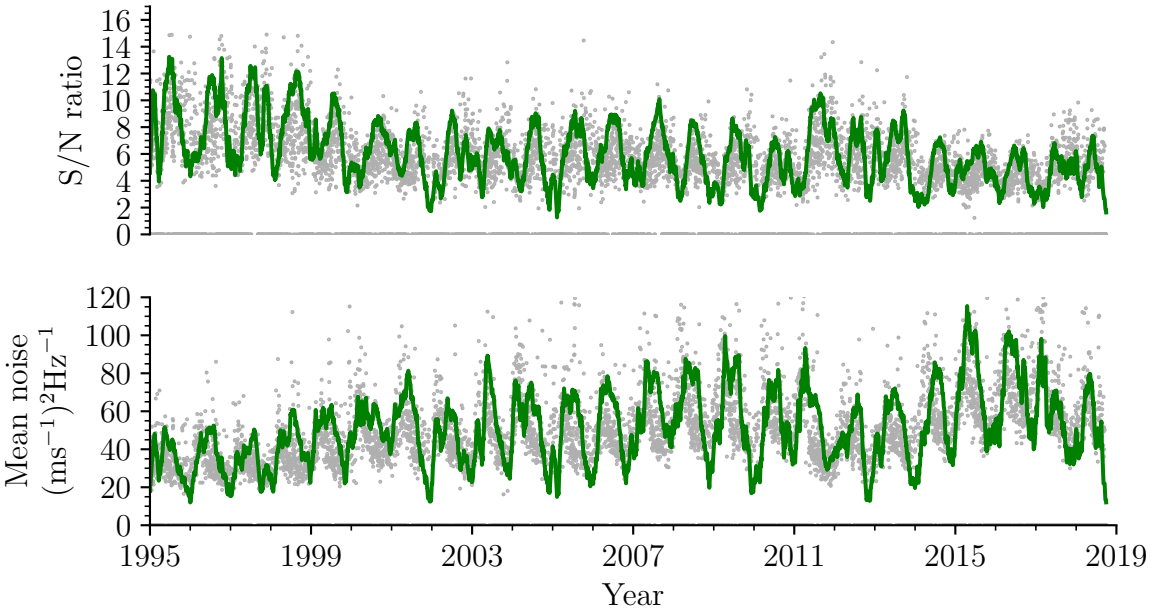

Figure 4.7: Izaña data quality as a function of date. Top: Signal to noise ratio, higher is better. Bottom: Mean noise level, lower is better. There is one grey dot per day, and the solid green curve represents a 50-day moving mean.

automation. From 1981 the group worked on an innovative new robotic design. Testing in Haleakala during 1983 proved that the system worked, and the group began looking for a permanent installation site. Western Australia was selected as a good longitude to complement coverage from the existing campaigns.

4.3.2 Carnarvon, Western Australia

Four sites in Western Australia were investigated: Woomera Rocket Testing Range, Learmonth, Exmouth, and Carnarvon. Isaak decided that Woomera was too dusty, but the other three towns were all suitable. The group settled on Carnarvon, some 900 km north of Perth, since it had a slightly better history of cloud-free days, and typically experienced lower temperatures than the others making for a less hostile working environment. Learmonth would later become home to one of the GONG nodes. Carnarvon also had the advantage of being home to the Overseas Telecommunications Commission (OTC), the Satellite Earth Station on Brown's Range (Dench and Gregg, 2010). Having technical help nearby would prove to be very useful in the event that a problem occurred with the

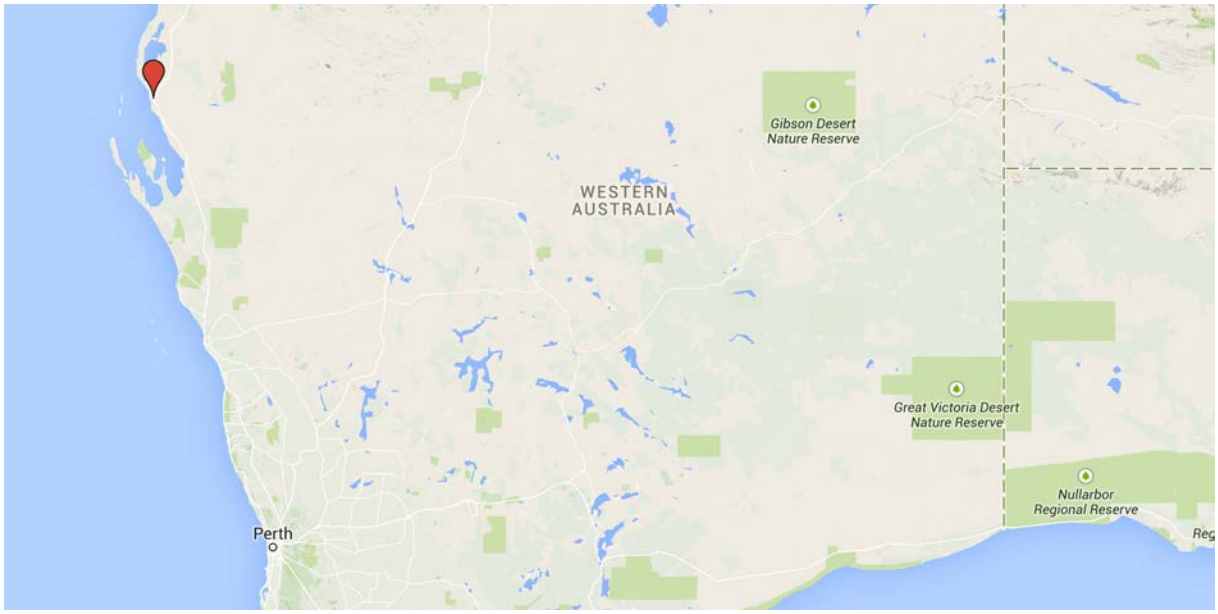

Figure 4.8: Carnarvon, Western Australia
Overseas Telecommunications Commission – Satellite Earth Station Carnarvon
Location: +113.75E -24.85N Elevation: 10 m Established: 1985
Image credit: Google Maps

automation. The location of Carnarvon is shown in Figure 4.8.

The group arrived in Carnarvon in 1984. Similar testing was carried out to that which had been successful in Haleakala. The system was set up on a flat section of the OTC site, and at night it was simply covered up with a tent, shown in Figure 4.9. Carnarvon was sufficiently remote that there were few concerns about vandals, and certainly no fear of rain. By the end of the year it was clear that full automation was practical, and the decision was made to push ahead.

Isaak secured funding for a prototype, and work started in Carnarvon in earnest. A concrete foundation was poured during the middle of 1985, and a central pier was built to support the mount and instrument. The structure was enclosed in a cylinder of marine plywood and corrugated iron, topped with a commercial 12 foot (3.7 m) dome from Ash Manufacturing in Illinois, USA. The observatory was split into two floors: the instrument and mount in the upper section, and the control electronics in the air-conditioned room below. The construction process is shown in Figure 4.10.

The data-acquisition and control system was initially based around a Hewlett-Packard

Figure 4.9: Carnarvon initial outdoor testing. Image credit: Clive P. McLeod

3421A Data Acquisition device, known colloquially as a "data logger". These were quickly retired from service and replaced with a standard desktop personal computer, and a Keithley System 570 digital input/output interface. The computer ran Microsoft DOS, and data would be retrieved each day via dial-up modem and a long-distance telephone call from Birmingham. This was long before the days of multi-tasking operating systems, and so data could only be retrieved during a scheduled "window" when the computer would switch from running the data-acquisition program to running a data-transfer program. Initially the data collection systems operated on a 42-s cadence like the existing systems at Izaña and Haleakala. Carnarvon was migrated to the newer 40-s cadence in April 1992. Data obtained previous to this date are interpolated onto the newer standard cadence.

The mount and all the associated dome electronics were controlled via the Keithley. The mount had position encoders on both the RA and declination axes so that the computer could know where it was pointing, and command it to move to specific coordinates. Likewise, the dome had position encoders so that it could be pointed in azimuth and elevation. System modules would receive encoder positions from the computer, and then decide how to drive the motors to move the mount and dome to the requested position.

Whilst the computer could, of course, calculate the position of the Sun for any given date and time to a very high precision, the mount encoders were not precise enough on their own to control the mount. In addition to the encoders, a two-stage "autoguider" was used. Each stage has a progressively narrower field-of-view. The encoders can point the mount accurately enough to bring the Sun into the view of the *coarse* guider, which then takes control and "pulls in" until the Sun is then within the very narrow field-of-view of the *fine* guider. Once the fine guider has acquired the Sun and achieved *finelock* then the guider takes over completely and controls the mount with no further input from the computer. If *finelock* is lost, for example if the sky becomes cloudy, then the mount drops back to computer control and the encoders are again used to keep the mount pointing in roughly the correct direction. When the weather clears, the mount is ready to go and the autoguider takes over again.

Since the whole system is designed to run completely unattended, careful weather monitoring is also required. Both rain and wind sensors are used to close the dome in the event of precipitation or excessive wind. The dome motors and computer are both powered through uninterruptible power-supplies, to ensure than the dome can close safely and become weather-proof in the event of a power failure.

Over the months following installation, the system worked extremely well. It produced high-quality data and proved to be reliable. Most glitches were due to the Australian wildlife: rodents eating through cables, or cockatoos making nests in the upper parts of the dome. The Sun itself caused some problems, with certain types of connector and cable insulation quickly decaying under constant exposure to solar UV. A simple and cheap solution was found for this, which was to wrap exposed cables and components in aluminium cooking foil.

Naturally over the years repairs have been made and upgrades installed. Probably the most significant is the upgrade to the control software. The original Keithley System 570 data-acquisition system in Carnarvon ran for more than a decade of continuous use. The group also maintained several spares, and so were confident in keeping them all running

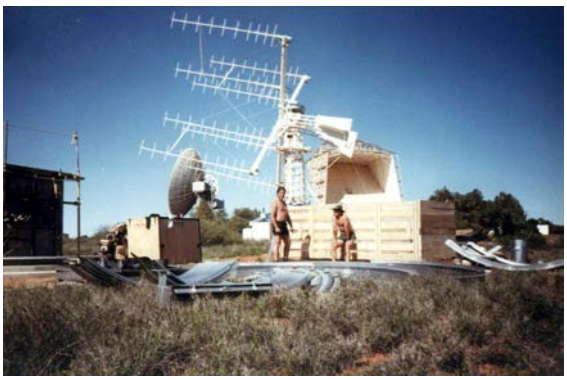

(a) Dome kit

(b) Base and pier

(c) Walls

(d) Almost complete

Figure 4.10: Construction of the Carnarvon observatory. (a) The dome arrives as a flat-pack kit. (b) A concrete base is poured and a central pier built. (c) The walls are fabricated from marine plywood and corrugated iron. (d) The dome is in place and almost complete, with only the shutter remaining to be installed. Image credit: Clive P. McLeod

for many more years. However, the PCs to which they were connected had a much shorter lifespan. The Keithley required an ISA peripheral interface slot on the motherboard. This was the standard interface on all *IBM compatible* PCs. But as the industry moved on from ISA, to EISA, VESA, VESA Local Bus, and PCI, the old ISA slots were eventually phased out. Whilst the Keithleys were still perfectly acceptable and functional interfaces, they could no longer be used with the modern PCs that replaced the old ones as they failed. The PC in Carnarvon failed in August 2002. The interface used to replace the Keithley was an Advantech PCI-1753 96-channel digital input/output card connected via a modern PCI-slot. Whilst it would have been possible to update the old DOS-based code to work with new hardware, by this point DOS itself was already obsolete. The open-source GNU/Linux operating system (Stallman, 1983; Torvalds, 1991) had become the de facto standard for network servers, and was an ideal choice for an industrial control system. The new hardware was first rolled out in August 2002 (Miller, 2002, 2003a).

Since GNU/Linux is a modern multi-tasking operating system, careful consideration has to be given to "locking" hardware resources to ensure that only one software process attempts to use a given piece of hardware at any one time. This was never required with DOS, since it could only ever run one program at once. If a program was running it knew it had exclusive access to the entire hardware and would run everything in one big loop, known as the *cyclic executive* approach. For simplicity, the hardware is split up into separate subsystems – dome, mount, temperatures, and so on. Each of these has an associated resource lock. If a program wants to access the dome hardware, it must first acquire the resource lock on the dome. To make this easier to remember, a "zoo" analogy is used. Each program, library, or device driver is an animal. The "Zoo" has several large paddocks for the animals to go in. When an animal has acquired a resource lock it is said to be in one of the paddocks. Only one animal can ever be in a paddock. The animal names are chosen to reflect their purpose. For example, the dome is a big heavy structure, and so a big heavy animal was chosen to control the dome – the elephant. The program that handles temperature control is called the iguana, because iguanas are cold-blooded

creatures and have to carefully control their own body temperature.

Another advantage to using GNU/Linux was that it had a complete communications and network stack built-in. The standard operating system tools could be used for dial-up networking, it was no longer required to write dedicated code to enable data to be sent home to Birmingham through a telephone link, and the multitasking nature also meant that this could be done at any time. There was no longer a special time at which it was possible to dial-up the computer. System time-keeping was upgraded to make use of the GPS time-signal. As always-on broadband connections became commonplace, GPS receivers were replaced by the online Network Time Protocol (NTP). The Zoo worked well and proved reliable, and was soon rolled out to all sites.

Carnarvon has not been without its problems. The failure of the PC in August 2002 caused the first drop-out in the Carnarvon duty cycle and window-function plots, Figures 4.11 and 4.12. In May 2005 a freak rainstorm produced over three inches of rain in just two hours. This is almost the same amount of rain that Carnarvon expects in a whole year. It was not a good time to find out that the rain detector had failed, and so the dome did not close (Barnes and Hale, 2005; New and Hale, 2006). At the time, Carnarvon had two instruments in operation. The whole dome was thoroughly flooded, completely destroying the control electronics for one instrument and severely damaging the second instrument. Water was literally poured out of some of the electronics. The primary instrument, Mark-V, was able to be repaired on-site after designing new detectors and control electronics in Birmingham. However, the secondary instrument, spectrophotometer-J (Jabba), was written off and had to be returned to Birmingham (Barnes et al., 2007) for a considerable programme of repairs and upgrades. The new instrument was installed in 2009 (Barnes and Miller, 2009) and provided such an improvement in noise level and data quality (Figure 4.13) that it became the primary instrument, seeing the retirement of the original spectrophotometer. The dome was shut down for just over a month in October 2013 during nearby construction work (Hale, 2014c).

Much of the old analogue electronics have been gradually replaced with new digi-

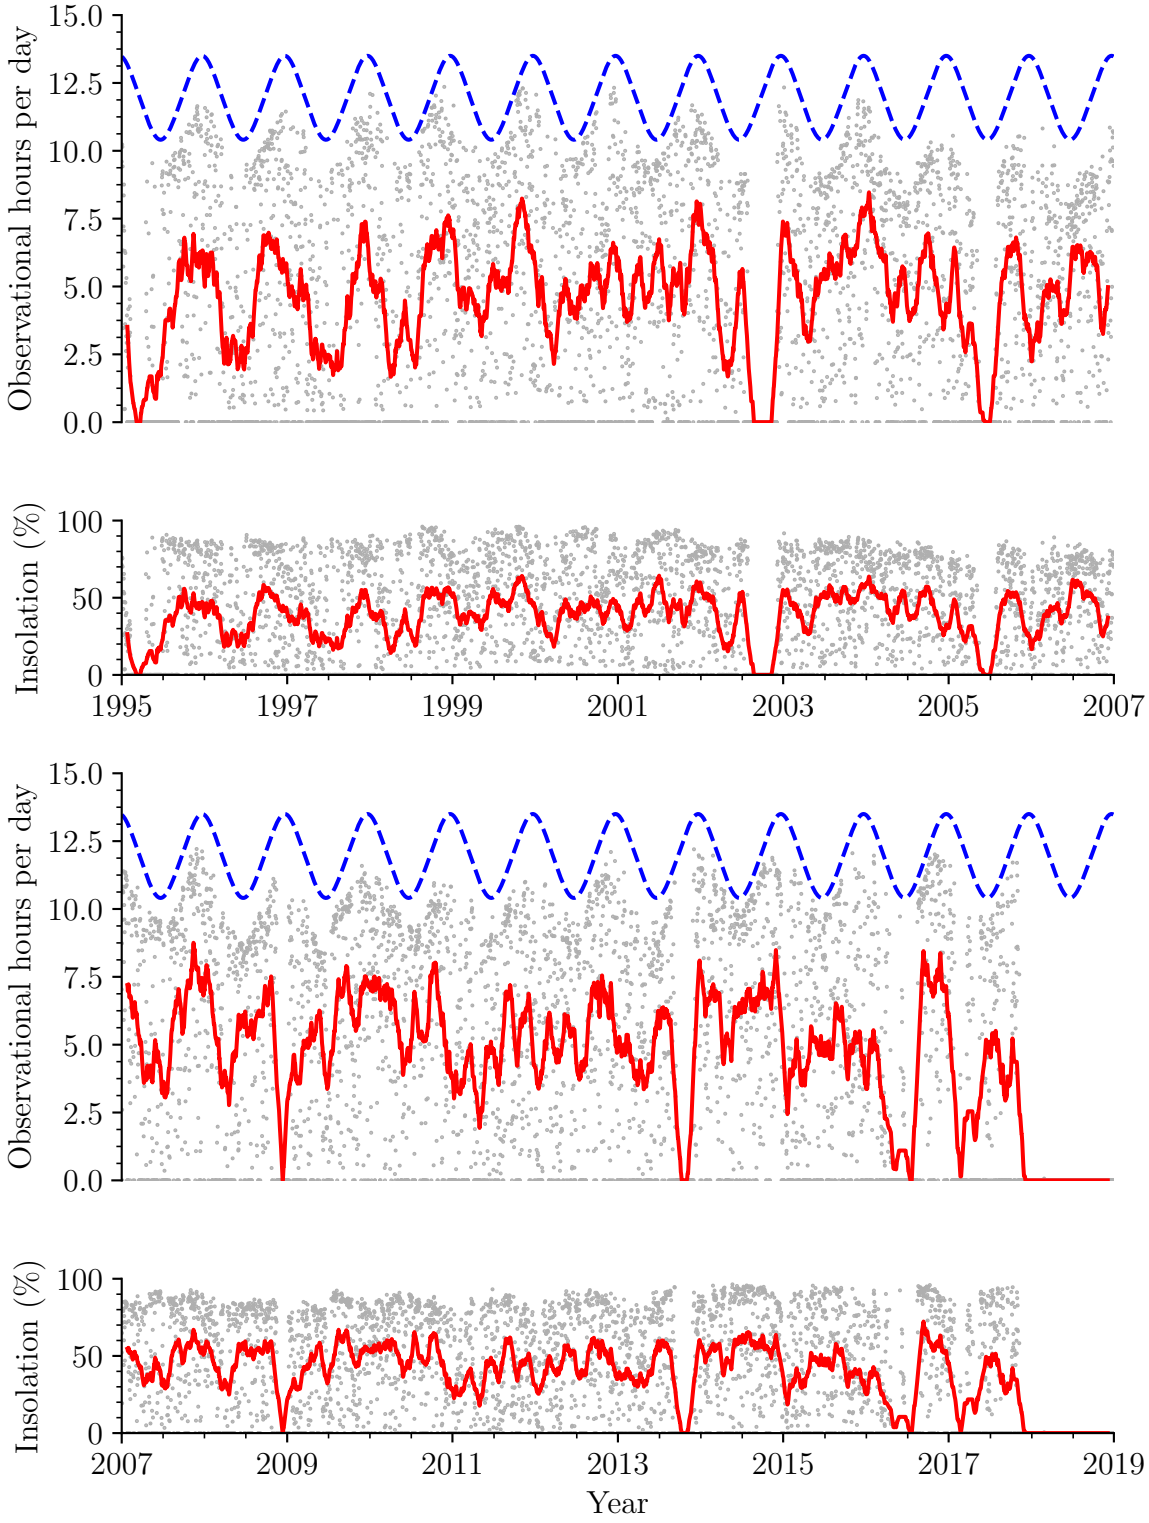

Figure 4.11: Carnarvon duty cycle as a function of date, plotted in hours per day, and as a percentage of potential daylight hours. There is one grey dot per day, and the solid red curve represents a 50-day moving mean. The dashed blue curve shows potential daylight hours.

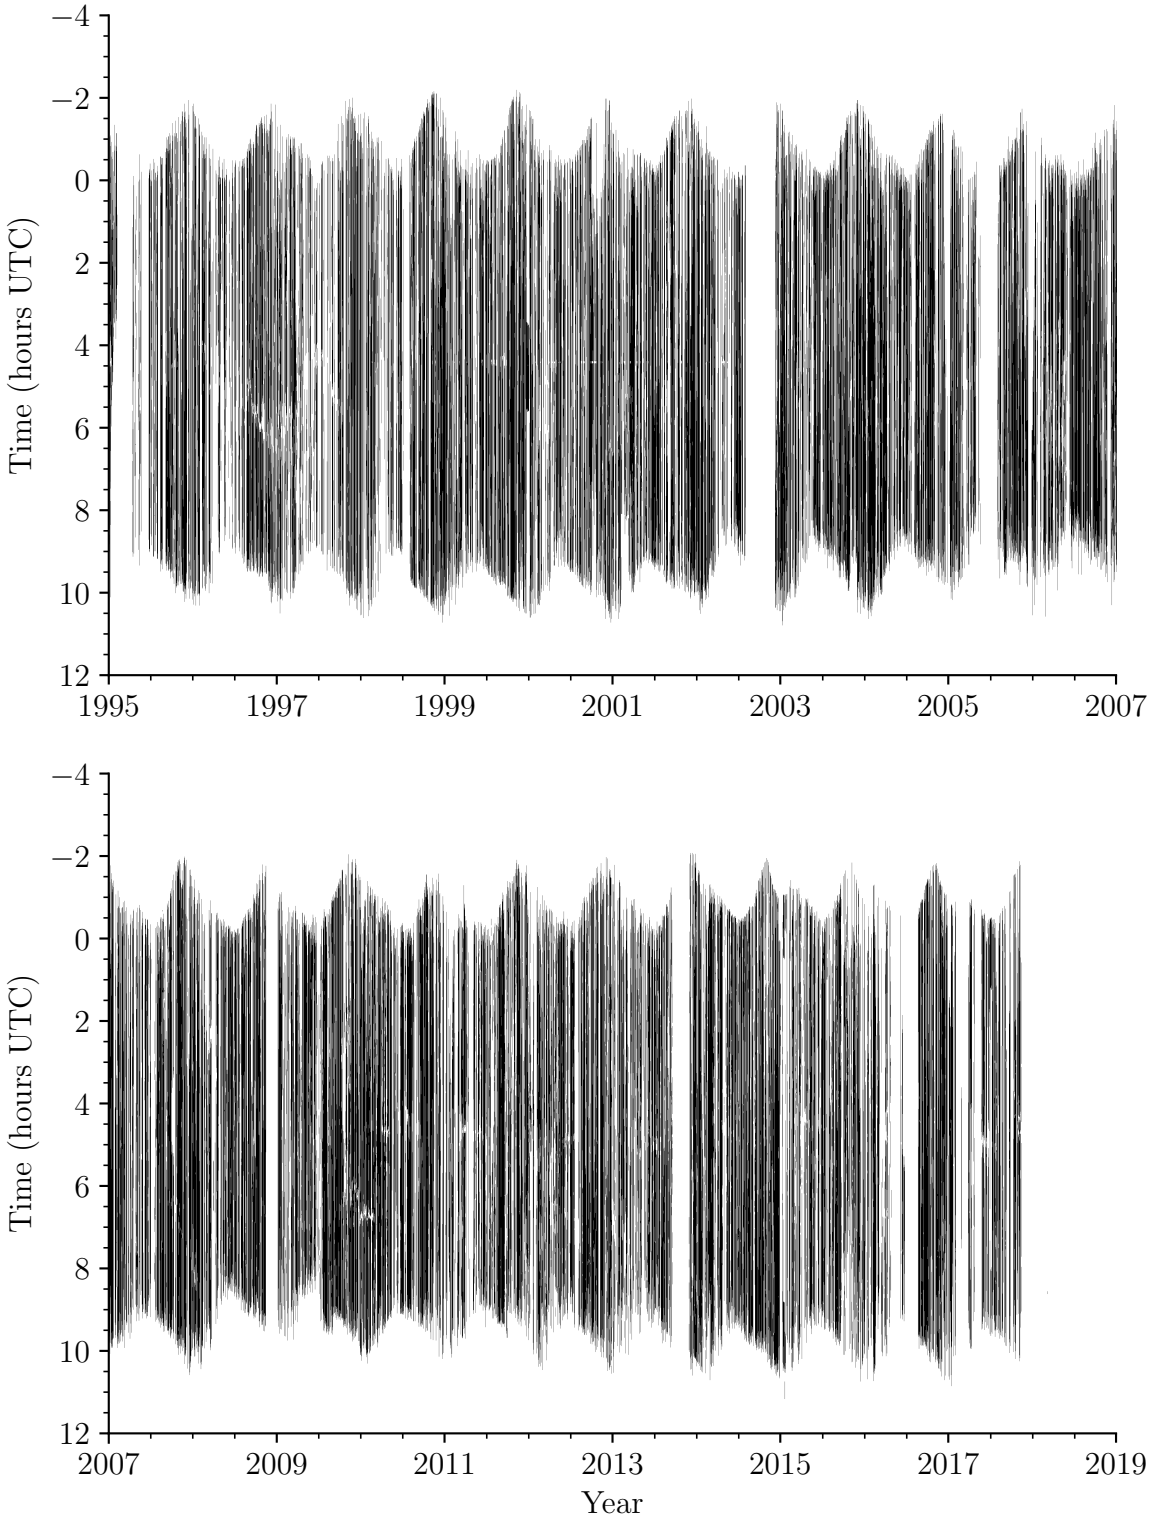

Figure 4.12: Carnarvon data window-function. The missing data in 2002 are due to a computer failure, in 2005 due to severe flooding, in 2009 due to spectrophotometer replacement, and in 2013 due to nearby construction work.

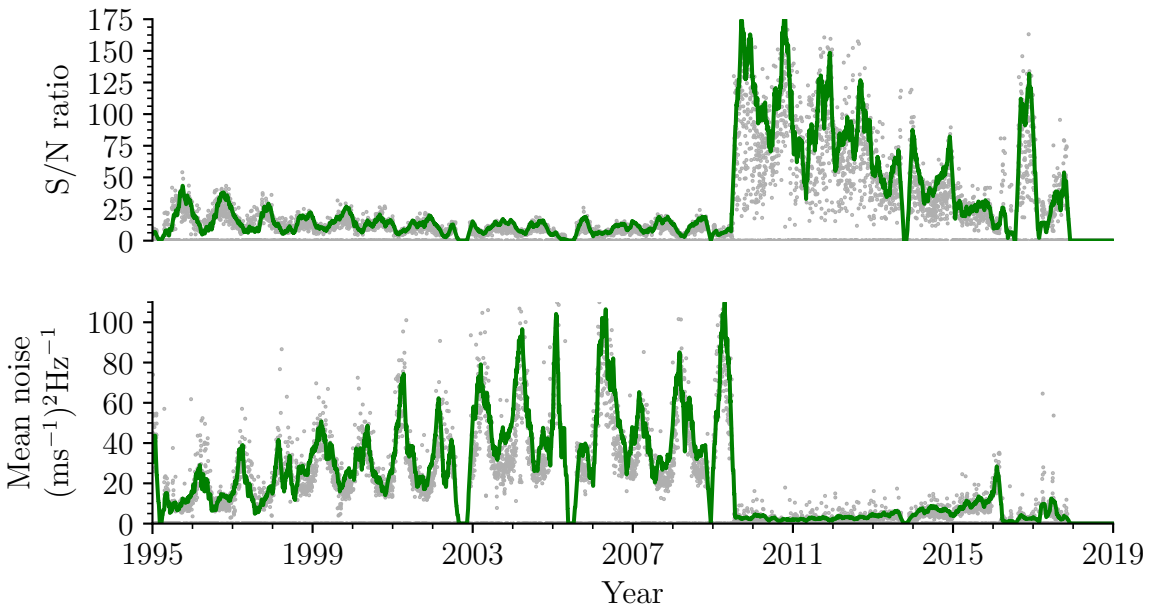

Figure 4.13: Carnarvon data quality as a function of date. Top: Signal to noise ratio, higher is better. Bottom: Mean noise level, lower is better. There is one grey dot per day, and the solid green curve represents a 50-day moving mean. The large step in 2009 is due to installation of a new upgraded spectrophotometer.

tal variants. Rather than the whole system consisting of modules in a central rack and communicating through a single interface, the new designs are independent with internal microcontrollers and communicate with the PC through dedicated RS-232 serial ports. This makes subsequent repairs and upgrades considerably easier since units can be inspected in isolation. So far Carnarvon has been given upgraded digital temperature controllers (Barnes, 2009b), a new digital autoguider (Hale and Miller, 2012; Barnes and Miller, 2012), and new counters (Hale, 2013a; Taylor et al., 2009). Problems with old optical filters were found in March 2016 (Hale, 2016a).

A internal leak on the water coolant system caused the instrument to be almost destroyed for a second time in August 2016 (Hale, 2016c), with severe damage to the power and signal interface PCBs. During development work on this project the site was converted to fibre-feed in April 2018 (Hale and Ross, 2018), and this has allowed the instrument to be moved to the ground floor of the dome inside the air-conditioned area removing the need for water cooling. Unfortunately, shortly following the conversion a failure of the

Figure 4.14: The Carnarvon dome, still standing proudly on the OTC site in 2015. The antennas on the left are the new satellite internet base station operated by NBN, the Australian National Broadband Network, designed to bring fast internet access to the Australian outback. The historic casshorn and dish antennas are visible in the background.

dome shutter motor, and subsequently the PC, resulted in extended downtime.

The Carnarvon BiSON dome continues to collect data year-round. A panorama of the current site condition, including the two historic OTC antennas, is shown in Figure 4.14. The goal of demonstrating that an automated system could work had been successfully achieved. Following the success of the Carnarvon station in the late 1980s, three more sites were commissioned.

4.3.3 Birmingham, UK

During the commissioning of the new sites, some minor changes were made to the original design. The main building became rectangular and made of brick, rather than cylindrical and made from plywood clad with corrugated iron. This substantially increased the floor-area, allowing for more storage space and easier access to the control electronics. The mount was also made considerably larger, allowing for bigger and heavier instrumentation.

The first of this new design was built in Birmingham in 1988, and is located on the roof of the Poynting Physics building. Given the generally poor weather conditions in Birmingham, this dome has unsurprisingly not been used for a regular observing programme. The Birmingham dome was primarily for testing new equipment, and training new staff before they visit a real site. It is still used for this purpose today. With the new observatory design finalised and tested, the first robotic site to be commissioned was in Sutherland, South Africa.

4.3.4 Sutherland, South Africa

The robotic BiSON station at Sutherland began construction in 1989 (Davidson and Williams, 2004) and was completed in 1990. The South African Astronomical Observatory (SAAO) itself was established in 1972, and is run by the National Research Foundation of South Africa. It is located in the Karoo Desert, in the Northern Cape of South Africa approximately 350 km northeast of Cape Town, shown in Figure 4.15.

Like Carnarvon, Sutherland started out with a Keithley System 570 data-acquisition system, and later got its Zoo upgrade in 2006 (Barnes and Miller, 2006). New digital temperature controllers were installed in 2007 (Barnes, 2008, 2009a), new counters in 2012 (Hale, 2012), and a new digital autoguider in 2013 (Hale, 2013b). The spectrophotometer is the Mark-F, codenamed "Fred".

The Sutherland duty cycle and data window-function are shown in Figures 4.16 and 4.17. The horizontal line that runs through midday in the window-function plot for 1995 – 2007 is caused by a beam-chopper. Since the detectors are unipolar, an offset voltage is injected to ensure the output never goes negative. The beam-chopper was used to check the dark counts from the detectors during the day. In 2006 it was decided that this was no longer necessary, and so the regular gap is no longer present.

The drop-out in late 1997 through early 1998 was caused by a range of faults. The tracking motor on the mount failed, and problems with the declination limit switches caused the mount to be unresponsive. Also, the dome moved out of alignment and began shadowing the instrument just before sunset. Later, the electronic polarisation modulator failed, and this was not replaced until late January 1998.

In early 1999, a fault developed on the scalar system that counts the pulses produced by the detectors. Following extensive investigation, simply re-seating all the chips on the scalar cards fixed the problem, but not until almost a month of data had been lost. The problem reoccurred several times over the years, until the scalar system was finally replaced completely in 2012. The gap in 2007 was due to some downtime whilst the temperature control systems were upgraded, and in 2008 due to a heavy snow storm.

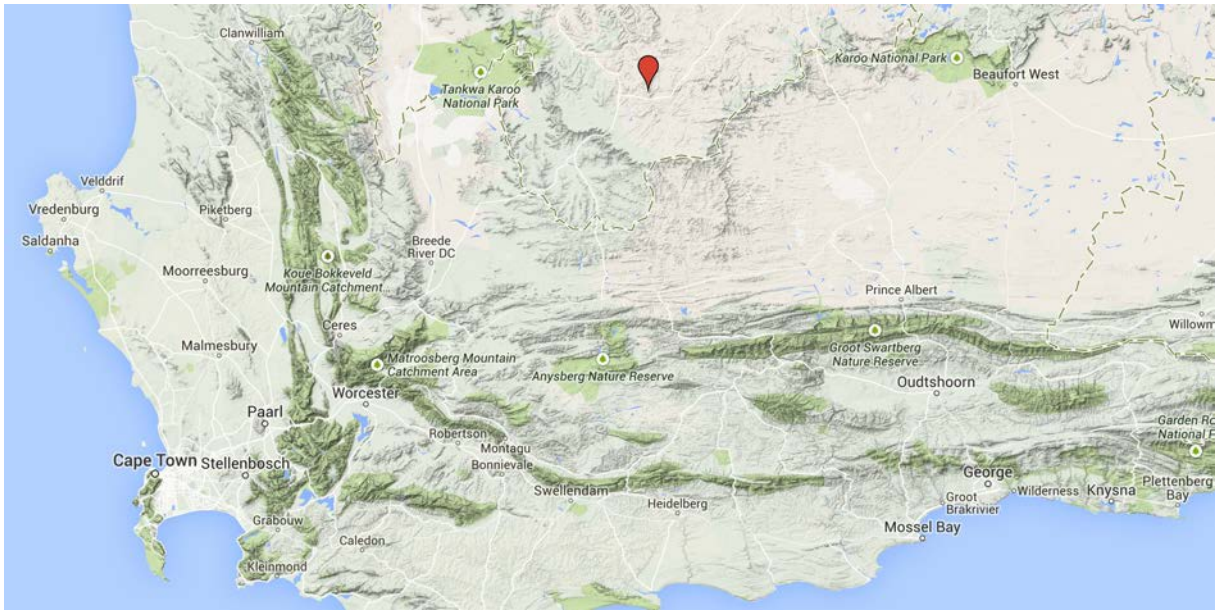

Figure 4.15: Sutherland, South Africa
South African Astronomical Observatory
Location: +20.82E -32.38N Elevation: 1771 m Established: 1990
Image credit: Google Maps

Looking at the data quality and noise levels (Figure 4.18) we see that the site has shown consistent performance with gradual improvement as systems were upgraded. The higher noise level in 2000 was due to a sticky declination gearbox. Weather conditions in Sutherland show solid year-round performance.

4.3.5 Las Campanas, Chile

After the completion of Sutherland, a third automated site was opened in Chile in 1991, at the Las Campanas Observatory operated by the Carnegie Institution for Science. Las Campanas Observatory is located in the southern Atacama Desert of Chile, around 100 km northeast of La Serena shown in Figure 4.19. The observatory was first established in 1969 (Di Mille et al., 2018), and was a replacement for the Mount Wilson Observatory near Pasadena in the USA which had started to experience too much light pollution from the growing city of Los Angeles. The main office is in La Serena, whilst the headquarters remain in Pasadena.

Las Campanas operates two spectrophotometers, Hannibal and Ivan. The site received

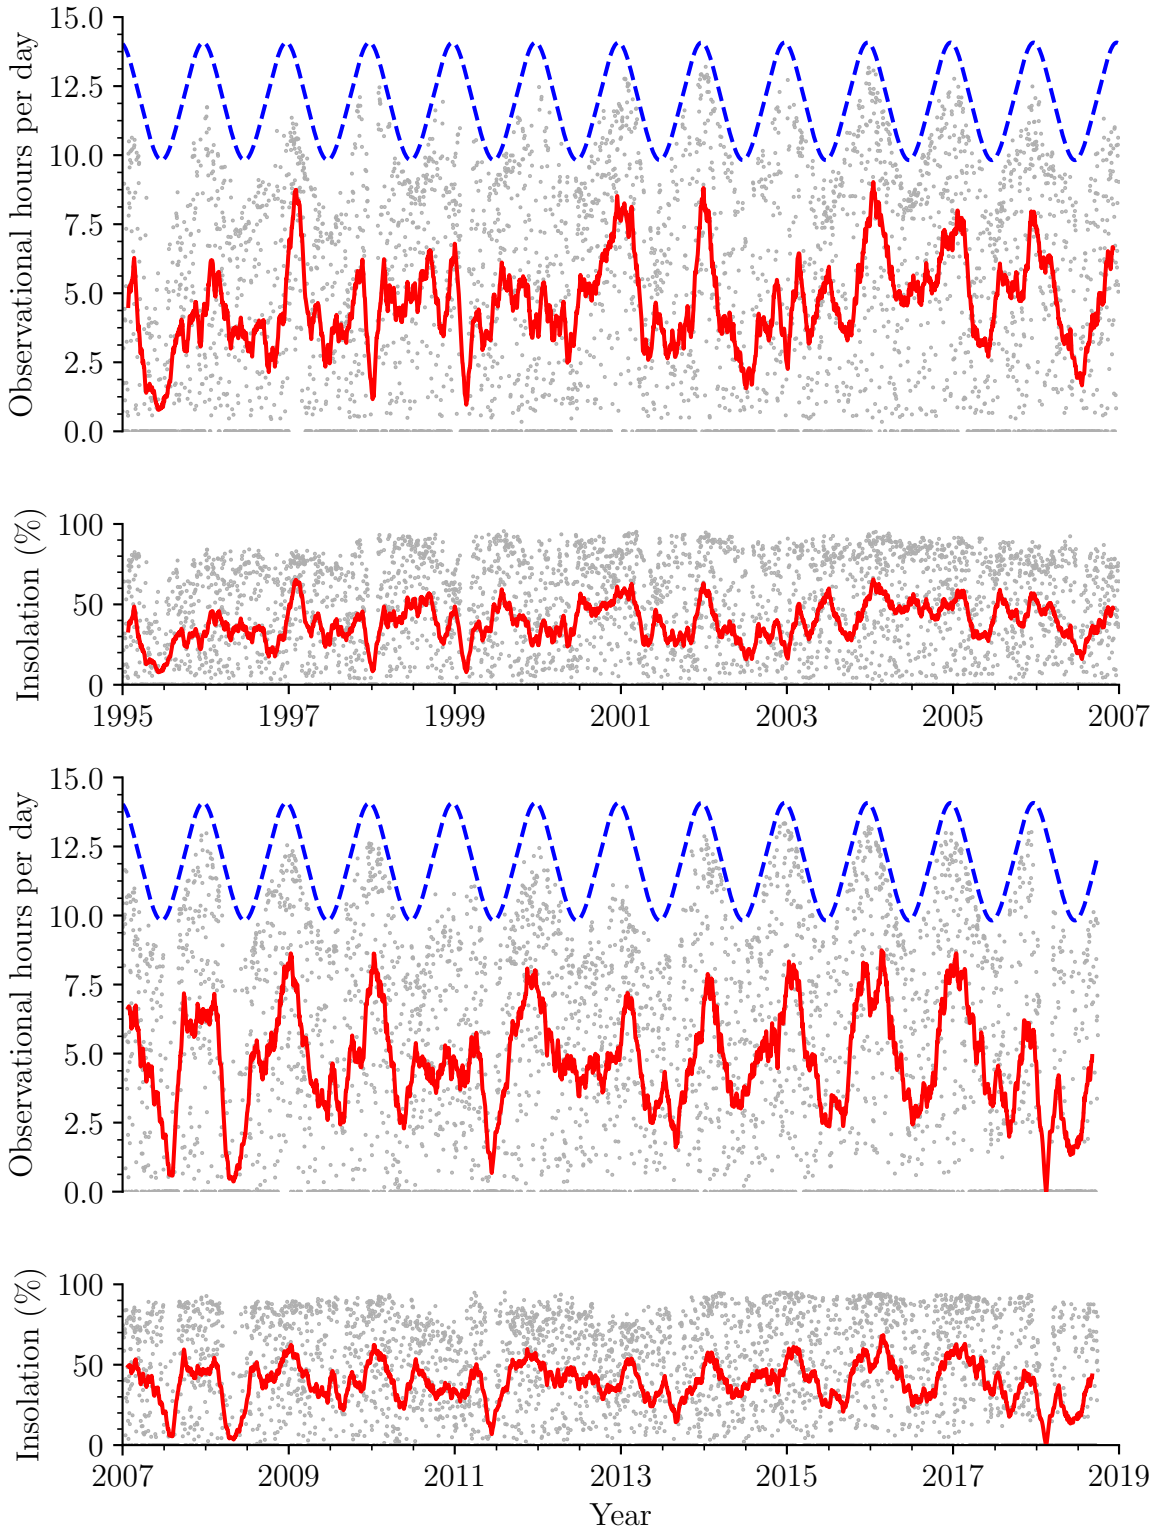

Figure 4.16: Sutherland duty cycle as a function of date, plotted in hours per day, and as a percentage of potential daylight hours. There is one grey dot per day, and the solid red curve represents a 50-day moving mean. The dashed blue curve shows potential daylight hours.

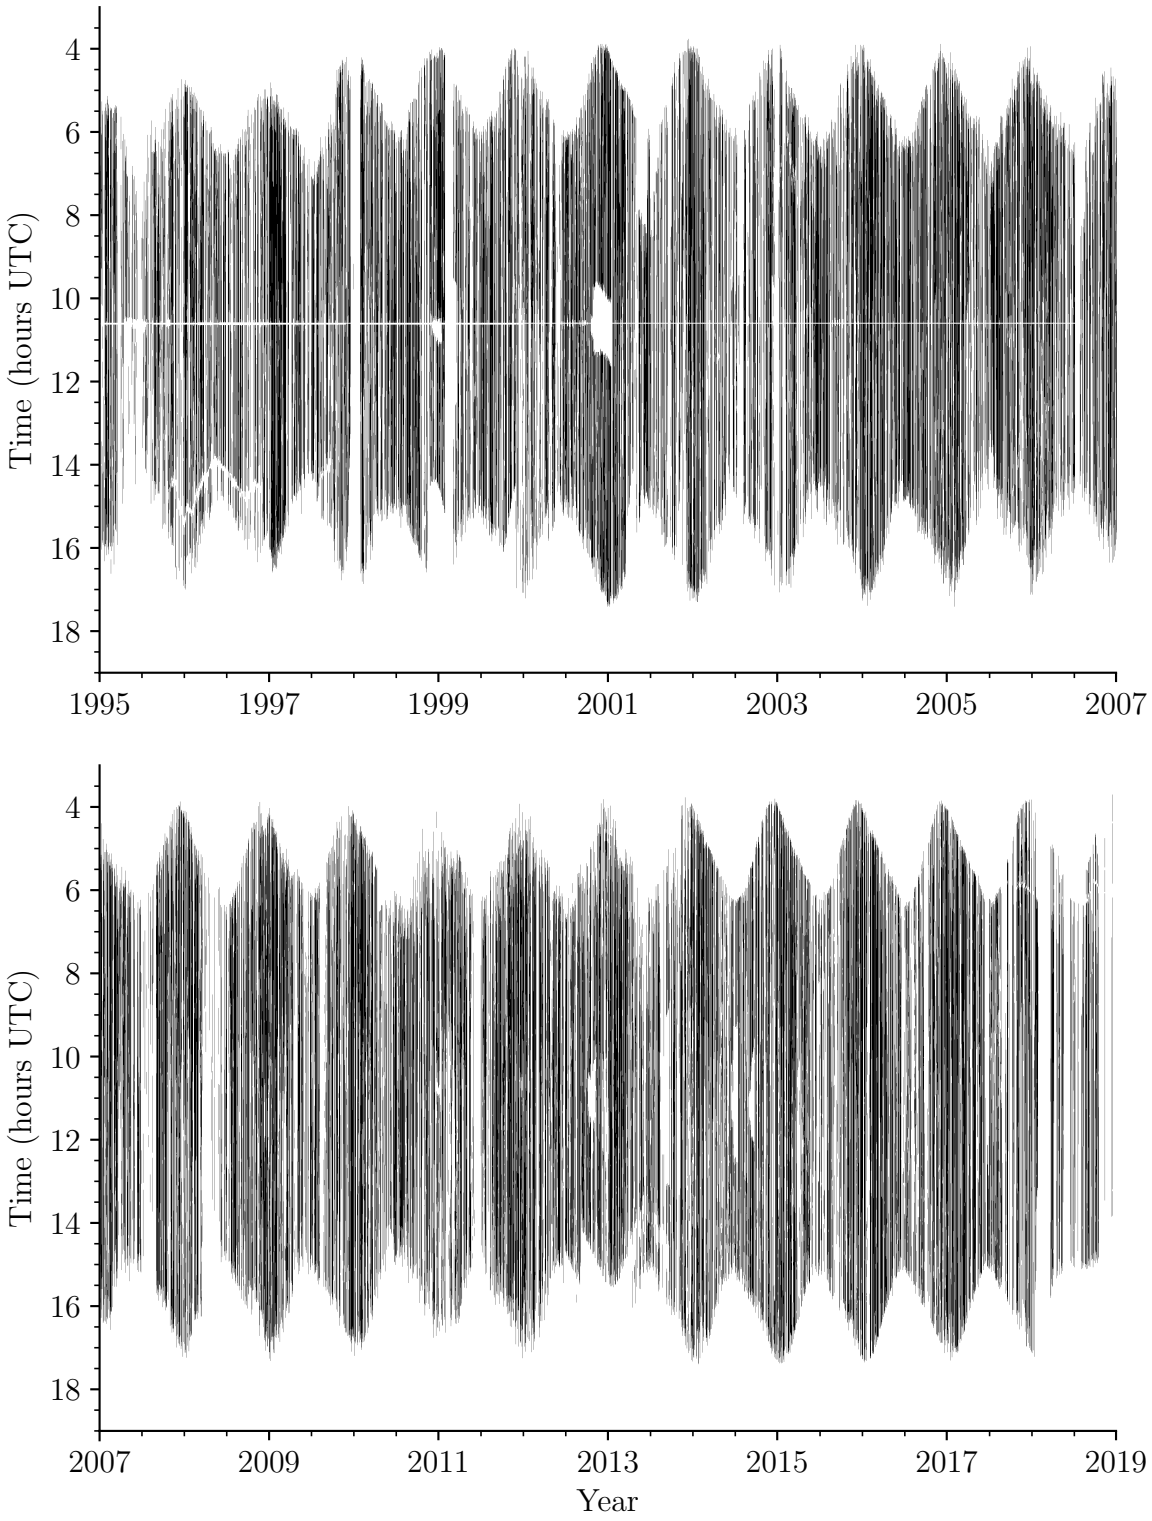

Figure 4.17: Sutherland data window-function.

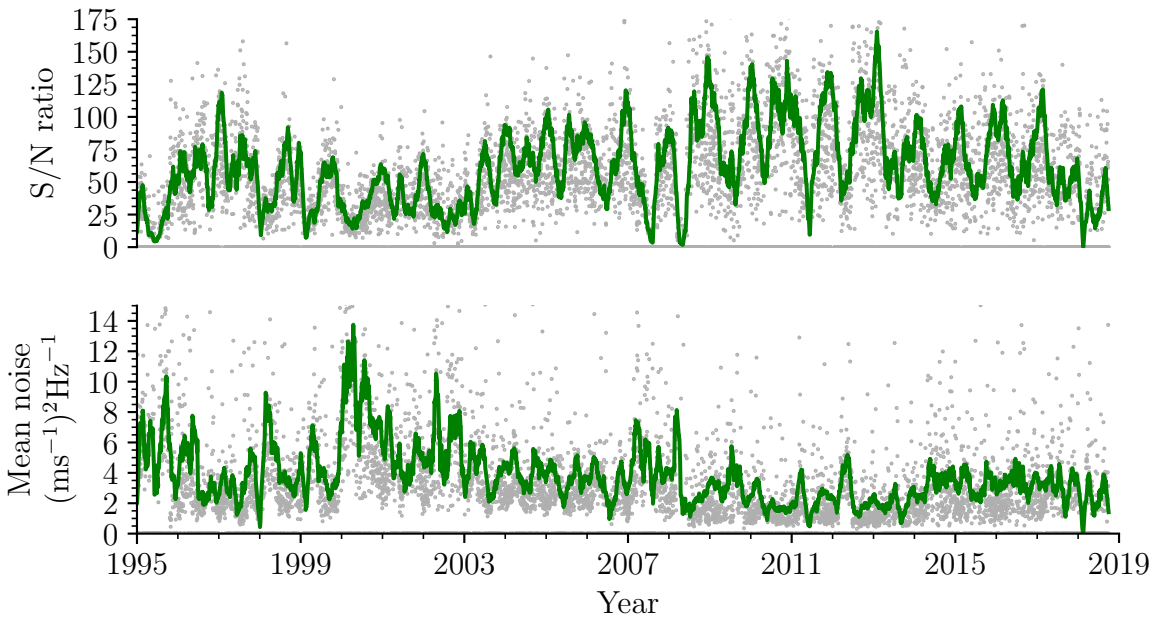

Figure 4.18: Sutherland data quality as a function of date. Top: Signal to noise ratio, higher is better. Bottom: Mean noise level, lower is better. There is one grey dot per day, and the solid green curve represents a 50-day moving mean.

its Zoo upgrade in December 2005 (Hale and Miller, 2006), a new digital autoguider in February 2011 (Miller, 2011; Barnes and Miller, 2011b), and a new temperature controller in 2015 (Hale, 2015b,c).

The Las Campanas duty cycle and data window-function are shown in Figures 4.20 and 4.21. The horizontal line that runs through midday in the window-function plot for 1995 – 2007 is again caused by a beam-chopper. As with Sutherland, in 2006 it was decided that this was no longer necessary, and so the regular gap is no longer present.

Las Campanas has had a range of problems. One notable experience was in July 1997 when a direct lightning strike to the dome destroyed several pieces of sensitive electronics. An emergency trip was organised with personnel arriving on site in August (Miller, 1997). Whilst repairs were completed, weather conditions deteriorated. Wind speeds reached a sustained 80 mph making it difficult to even walk to the dome. Eventually the wind abated, only to be replaced with a heavy snow fall that left the observatory cut off by road, without internet or phone connections, and dependent on old diesel generators for electrical power. As the snow melted, the 4.2 kV transformer junction-box supplying the

Figure 4.19: Las Campanas, Chile
Carnegie Institution for Science
Location: -70.70E -29.02N Elevation: 2282 m Established: 1991
Image credit: Google Maps

BiSON dome was first flooded – requiring drainage holes to be drilled in the box as a temporary fix – and subsequently destroyed by another lightning strike. The dome was without power for more than a month, until early November 1997.

Two further visits to finalise repairs were needed in November 1997 (Lines, 1998) and January 1998 (Miller, 1998b), which were related to problems with the computer, the dome azimuth motor, and the water-loop system that stabilises the instrumentation temperatures. Noise problems in both detectors persisted until early 1999, clearly visible in the plot of data quality, Figure 4.22. A large part of 2000 was lost due to a broken declination gearbox on the mount, and also to the failure of the dome azimuth motor. More electrical problems occurred in May 2014 when faults with both the shutter and blind limit switches caused the circuit breakers to trip repeatedly. A site visit was required to replace the limit switches, and also work on additional problems with the water-loop pump and the uninterruptible power supplies (Hale, 2014b). The new mount controller had to be replaced in November 2017 due to issues reading the mount position encoders (Hale, 2017d) caused by corrosion around some of the PCB header connectors.

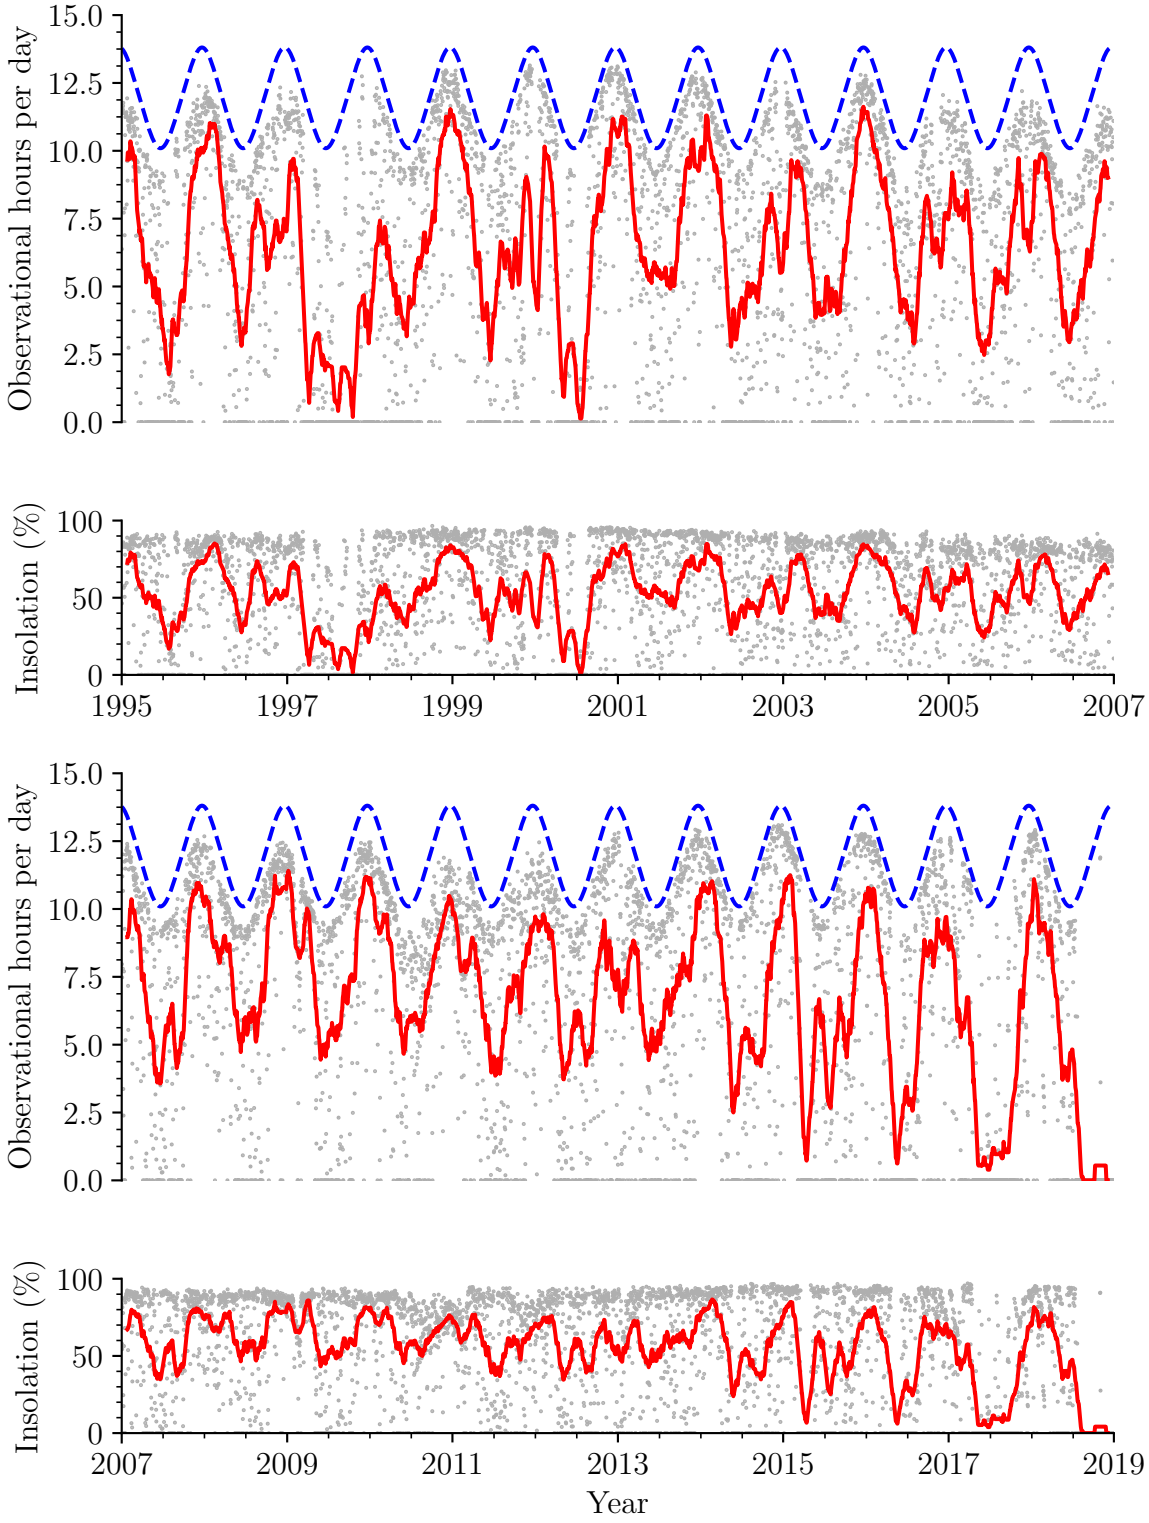

Figure 4.20: Las Campanas duty cycle as a function of date, plotted in hours per day, and as a percentage of potential daylight hours. There is one grey dot per day, and the solid red curve represents a 50-day moving mean. The dashed blue curve shows potential daylight hours.

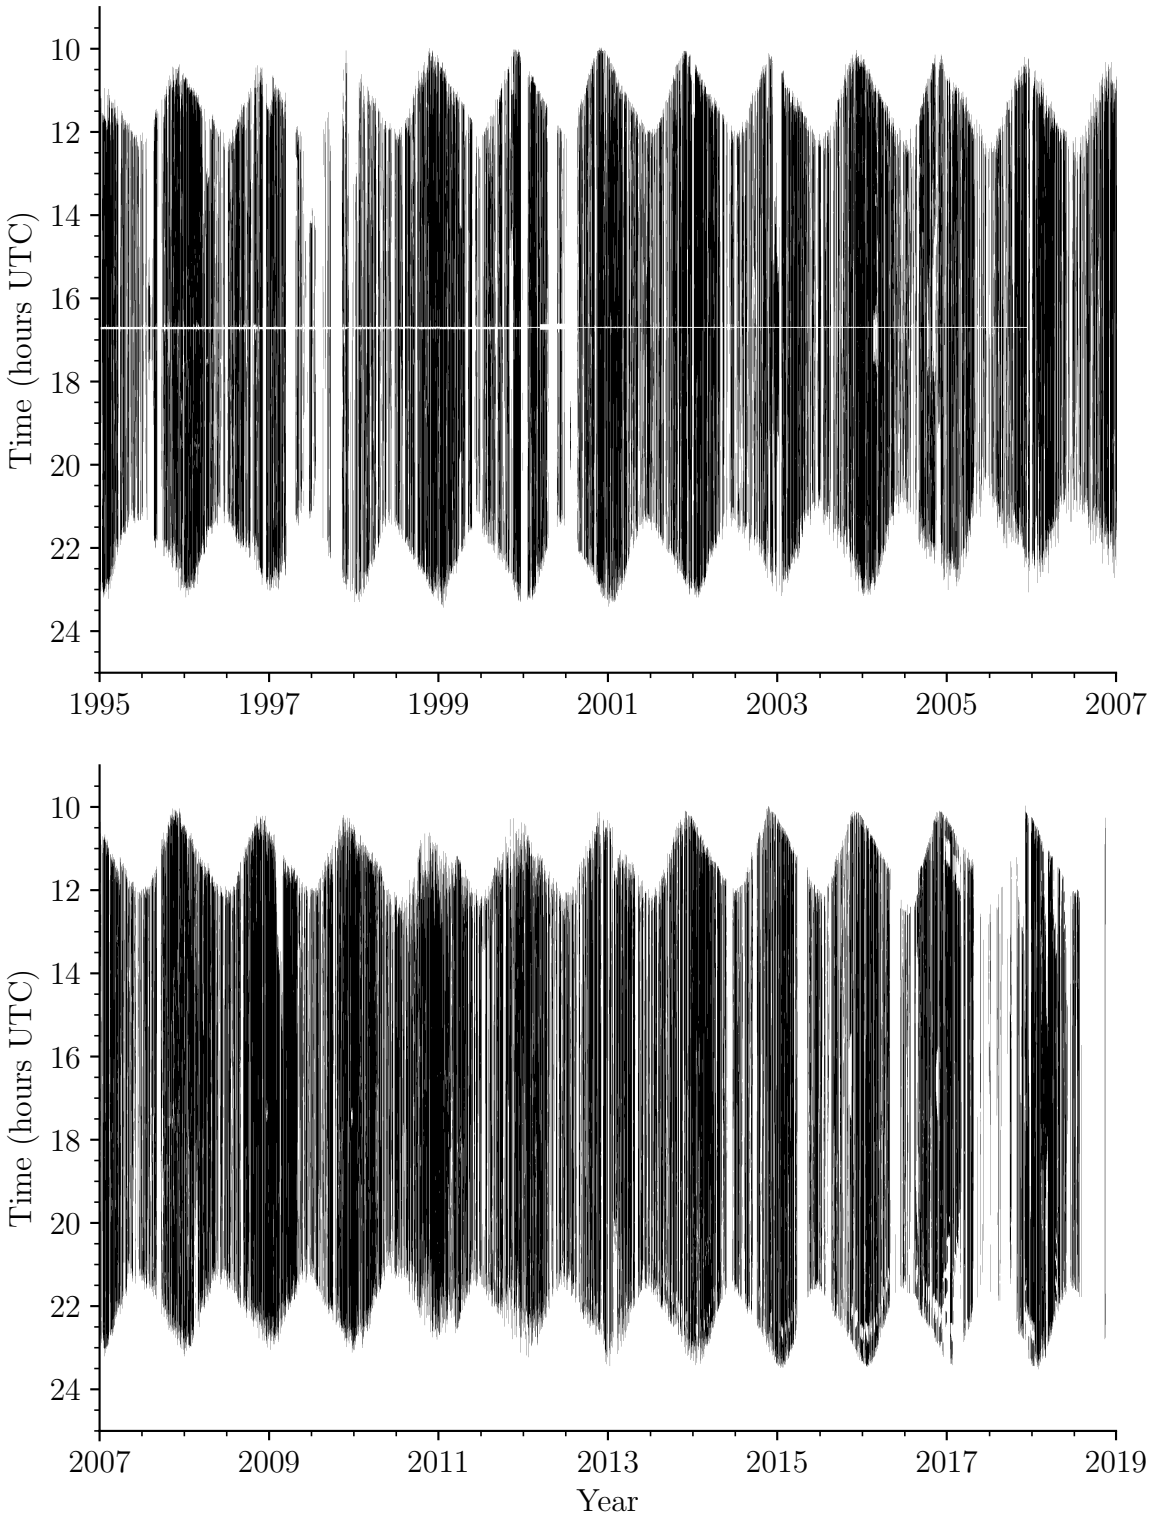

Figure 4.21: Las Campanas data window-function.

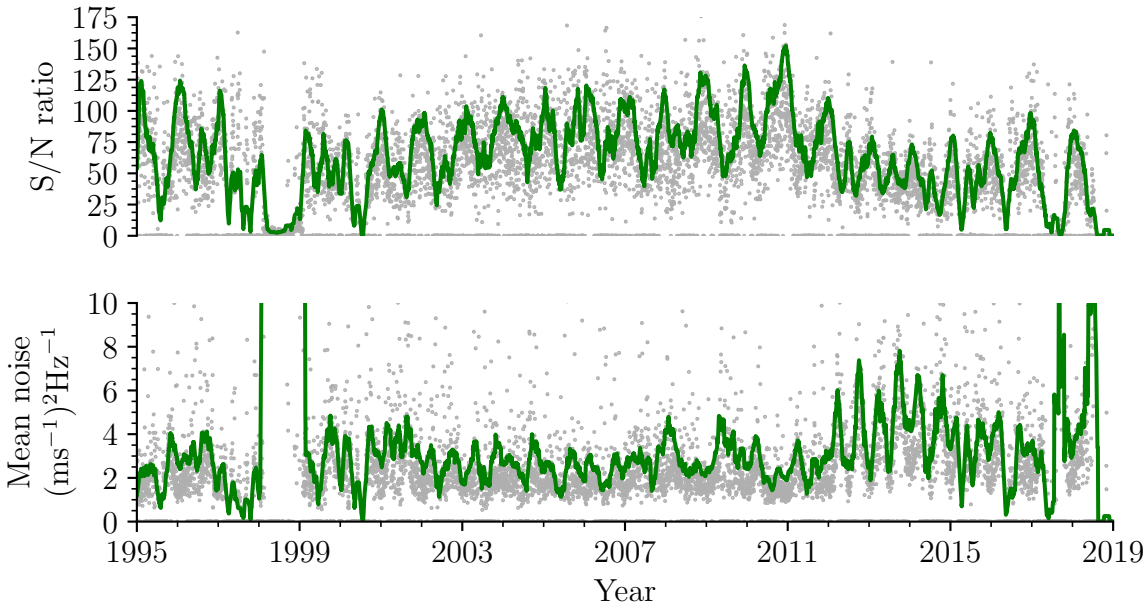

Figure 4.22: Las Campanas data quality as a function of date. Top: Signal to noise ratio, higher is better. Bottom: Mean noise level, lower is better. There is one grey dot per day, and the solid green curve represents a 50-day moving mean.

Las Campanas is the best performing station in the network, consistently supplying duty cycles above 80 % in the summer and regularly above 40 % even during the winter months. From 2012 the noise performance has deteriorated slightly due to reduced performance of the potassium scattering vapour cell. In October 2018 an unrelated fault with the weather module caused the vapour cell to overheat and be damaged (Hale, 2018c). The cell was replaced producing a noticeable improvement in performance.

4.3.6 Narrabri, NSW, Australia

The final fully-automated site was installed in Narrabri, Australia, in 1992 (Williams, 2004). It is on the site of the Australia Telescope Compact Array (ATCA) at the Paul Wild Observatory operated by the Commonwealth Scientific and Industrial Research Organisation (CSIRO). The observatory is around 550 km northwest of Sydney, shown in Figure 4.23.

The spectrophotometer at Narrabri is the Mark-G, Ghengis. Narrabri received its Zoo upgrade in June 2004 (Jackson and Miller, 2004), new temperature controllers in

Figure 4.23: Narrabri, New South Wales, Australia
Paul Wild Observatory, Australia Telescope National Facility, CSIRO
Location: +149.57E -30.32N Elevation: 217 m Established: 1992
Image credit: Google Maps

January 2010 (Barnes and Hale, 2010; Barnes, 2010), and a new digital autoguider and counters in April 2013 (Hale and Davies, 2013).

The Narrabri duty cycle and data window-function are shown in Figures 4.24 and 4.25. The horizontal line that runs through midday in the window-function plot for 1995 – 2007 is again caused by a beam-chopper. As with Sutherland and Las Campanas, in 2006 it was decided that this was no longer necessary, and so the regular gap is no longer present.

The gap in 1997 is due to problems with the scalars and correct termination of the signal cables from the voltage-to-frequency output of the detectors. The missing data in 2003 and 2004 is due to blind motor problems (Miller, 2000; New and Isaak, 2003). In 2013 a fault on the anemometer in June caused problems with the weather module, and this kept the dome closed unnecessarily. In October 2015 the water coolant pump was replaced, and some other temperature issues identified as being caused by a broken drive cable (Hale, 2015d). Further problems were experienced in February 2018 related to the detector temperature stabilisation, which required replacement of some thermoelectric Peltier modules (Hale, 2018a).

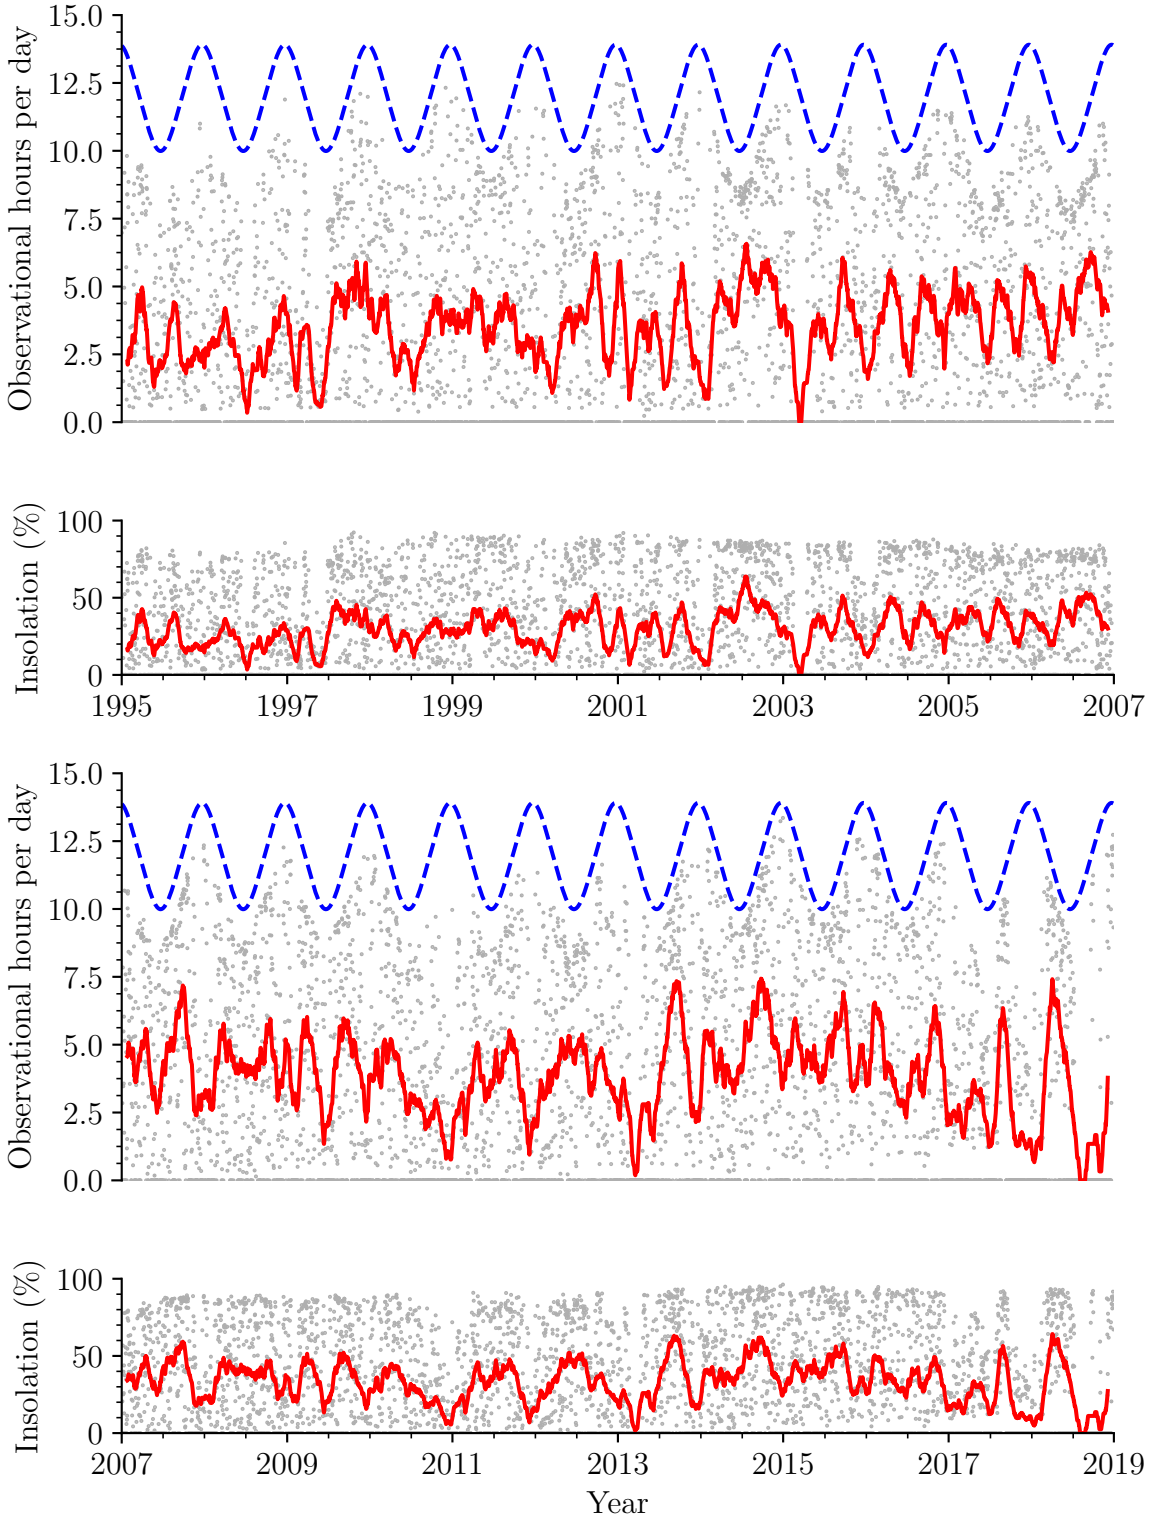

Figure 4.24: Narrabri duty cycle as a function of date, plotted in hours per day, and as a percentage of potential daylight hours. There is one grey dot per day, and the solid red curve represents a 50-day moving mean. The dashed blue curve shows potential daylight hours.

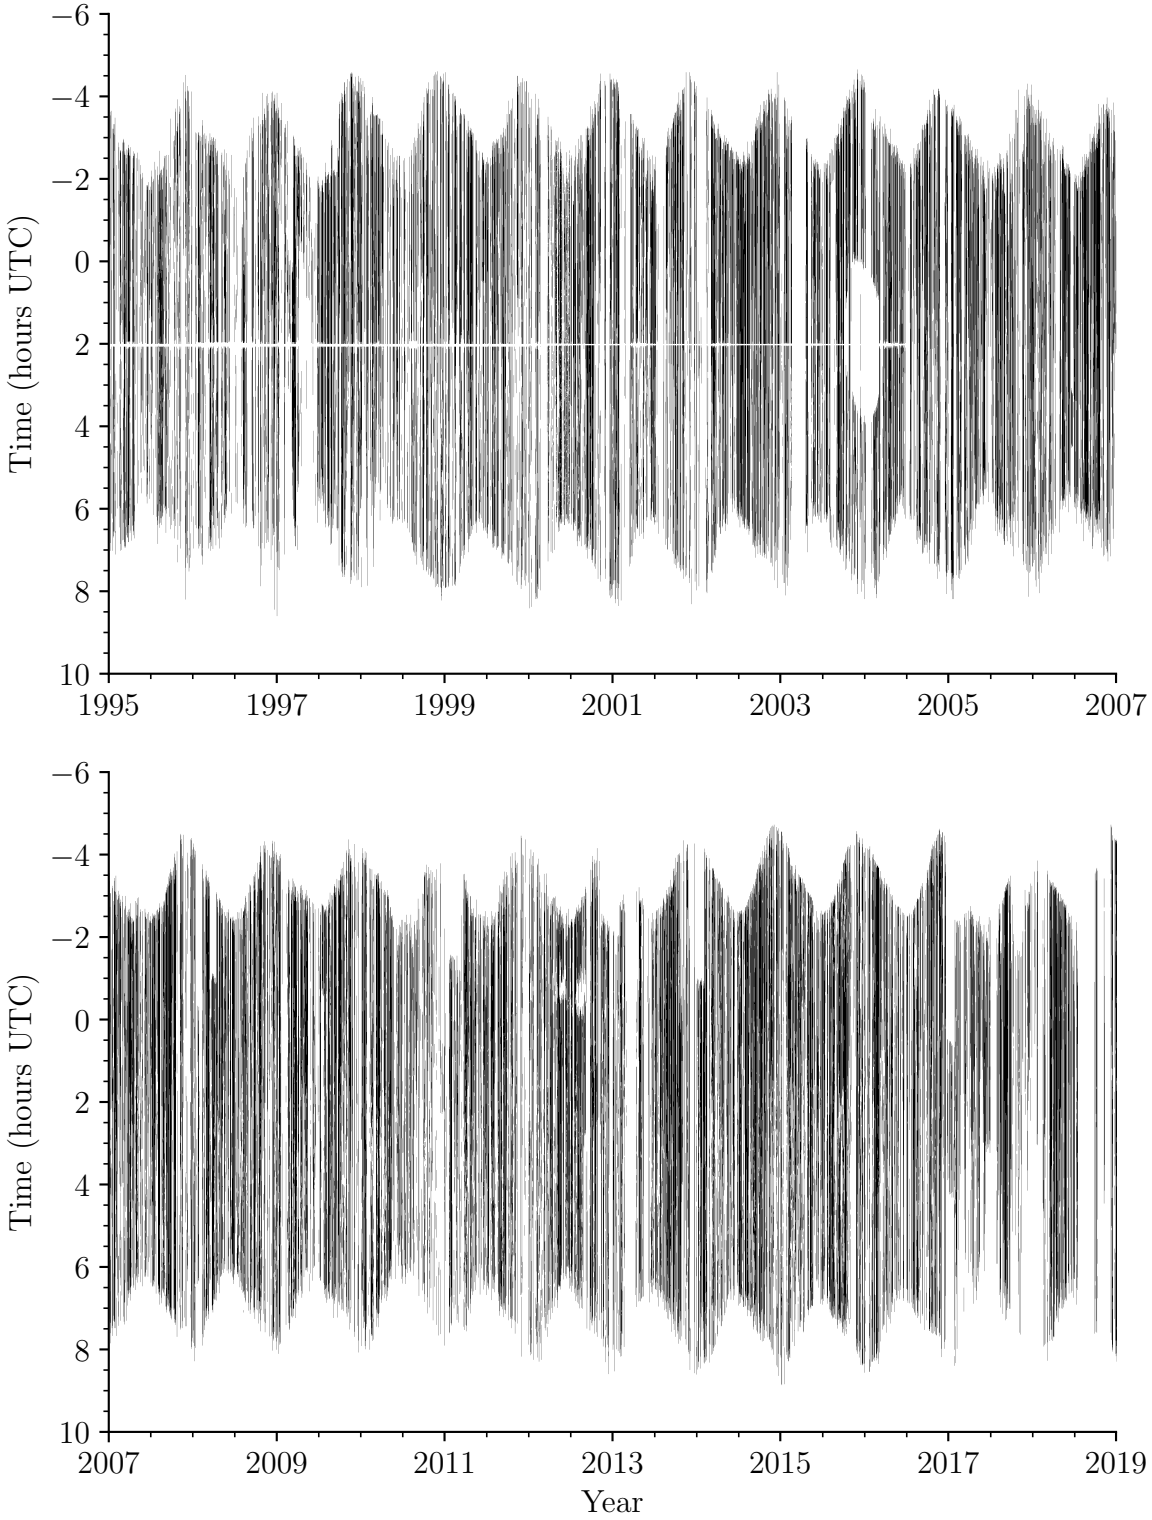

Figure 4.25: Narrabri data window-function.

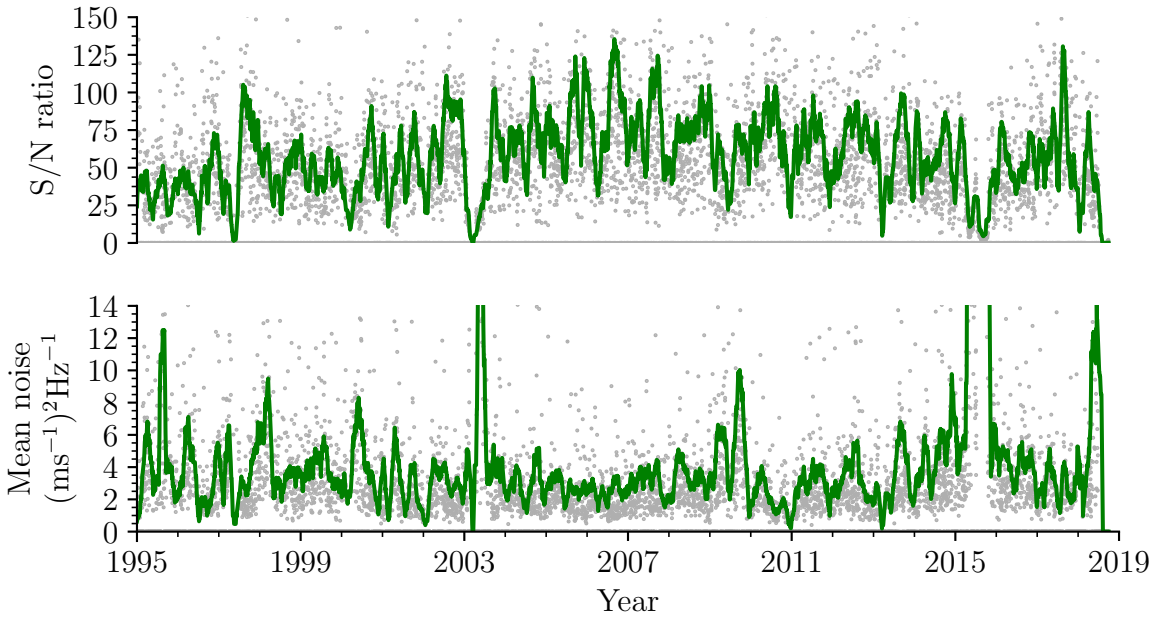

Figure 4.26: Narrabri data quality as a function of date. Top: Signal to noise ratio, higher is better. Bottom: Mean noise level, lower is better. There is one grey dot per day, and the solid green curve represents a 50-day moving mean.

In the plot of data quality, Figure 4.26, the increased noise level in 2003 is due to a failure of an interference filter temperature controller, and that in 2009 is due to a faulty power supply producing under-voltage power rails. Other than these faults, the site shows solid performance.

4.3.7 Mount Wilson, California, USA

The final change to arrive at the existing network configuration was also made in 1992. The Mark-III instrument was moved to the 60 foot tower at the Mount Wilson (Hale) Observatory in California (Elsworth, 1992), the same observatory where Leighton et al. originally discovered solar oscillations. The observatory is located in the San Gabriel Mountains near Pasadena, around 80 km northeast of Los Angeles, shown in Figure 4.27.

The tower is operated by Prof. E. J. Rhodes, Jr. and his team of undergraduate volunteer observers. Like Izaña the tower uses a cœlostat to direct light down the tower into the observing room below, and as such requires someone to open the dome and align the mirrors each morning. It also requires presence throughout the day to close the dome in

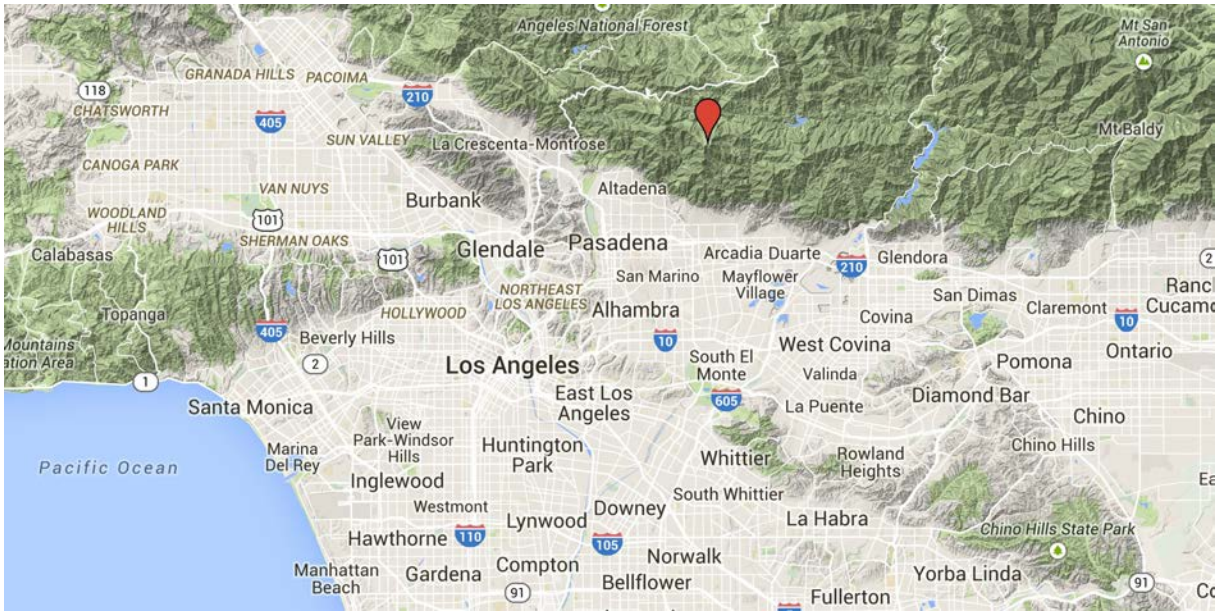

Figure 4.27: Mount Wilson, Los Angeles, California, USA
Mount Wilson (Hale) Observatory
Location: -118.08E +34.13N Elevation: 1742 m Established: 1992
Image credit: Google Maps

the event of bad weather, and at the end of the observing session.

The Mount Wilson duty cycle and data window-function are shown in Figures 4.28 and 4.29. The horizontal line that runs through midday in the window-function plot for 1995 – 2007 is again caused by a beam-chopper. As with Sutherland and Las Campanas, in 2006 it was decided that this was no longer necessary, and so the regular gap is no longer present. There are also additional gaps at midday during the winter months, when the primary mirror has to be moved from the east side of the tower in the morning to the west in the afternoon, in order to avoid shadowing from the secondary mirror. Unlike Izaña, the secondary at Mount Wilson has only one mounting configuration.

In July 1996 the Mark-III instrument from Haleakala was retired and replaced with the eleventh spectrophotometer designed by the group, code-named Klaus (Miller and New, 1999). The new design offered much improved data quality in line with the other newer instruments in the network, and the reduction in noise level can be seen clearly in the plot of data quality, Figure 4.30. The large step in noise level during 2004 corresponds to an autoguider problem, where it failed to drive in right ascension. A transformer

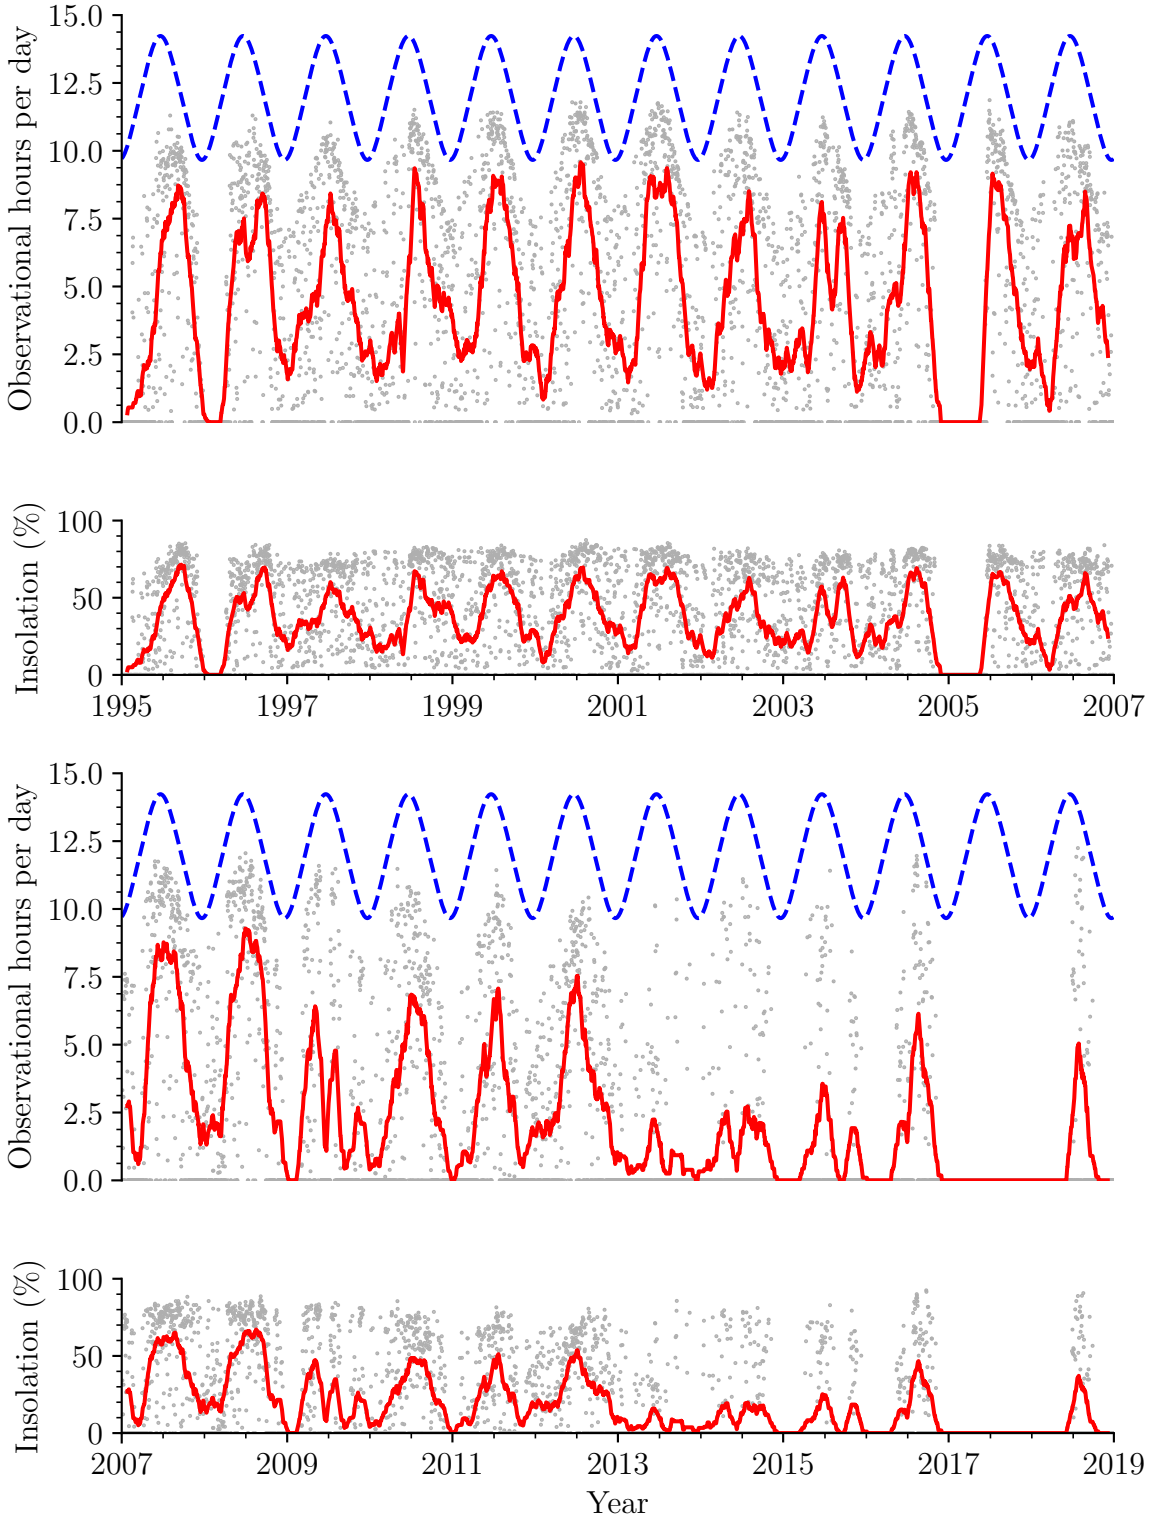

Figure 4.28: Mount Wilson duty cycle as a function of date, plotted in hours per day, and as a percentage of potential daylight hours. There is one grey dot per day, and the solid red curve represents a 50-day moving mean. The dashed blue curve shows potential daylight hours.

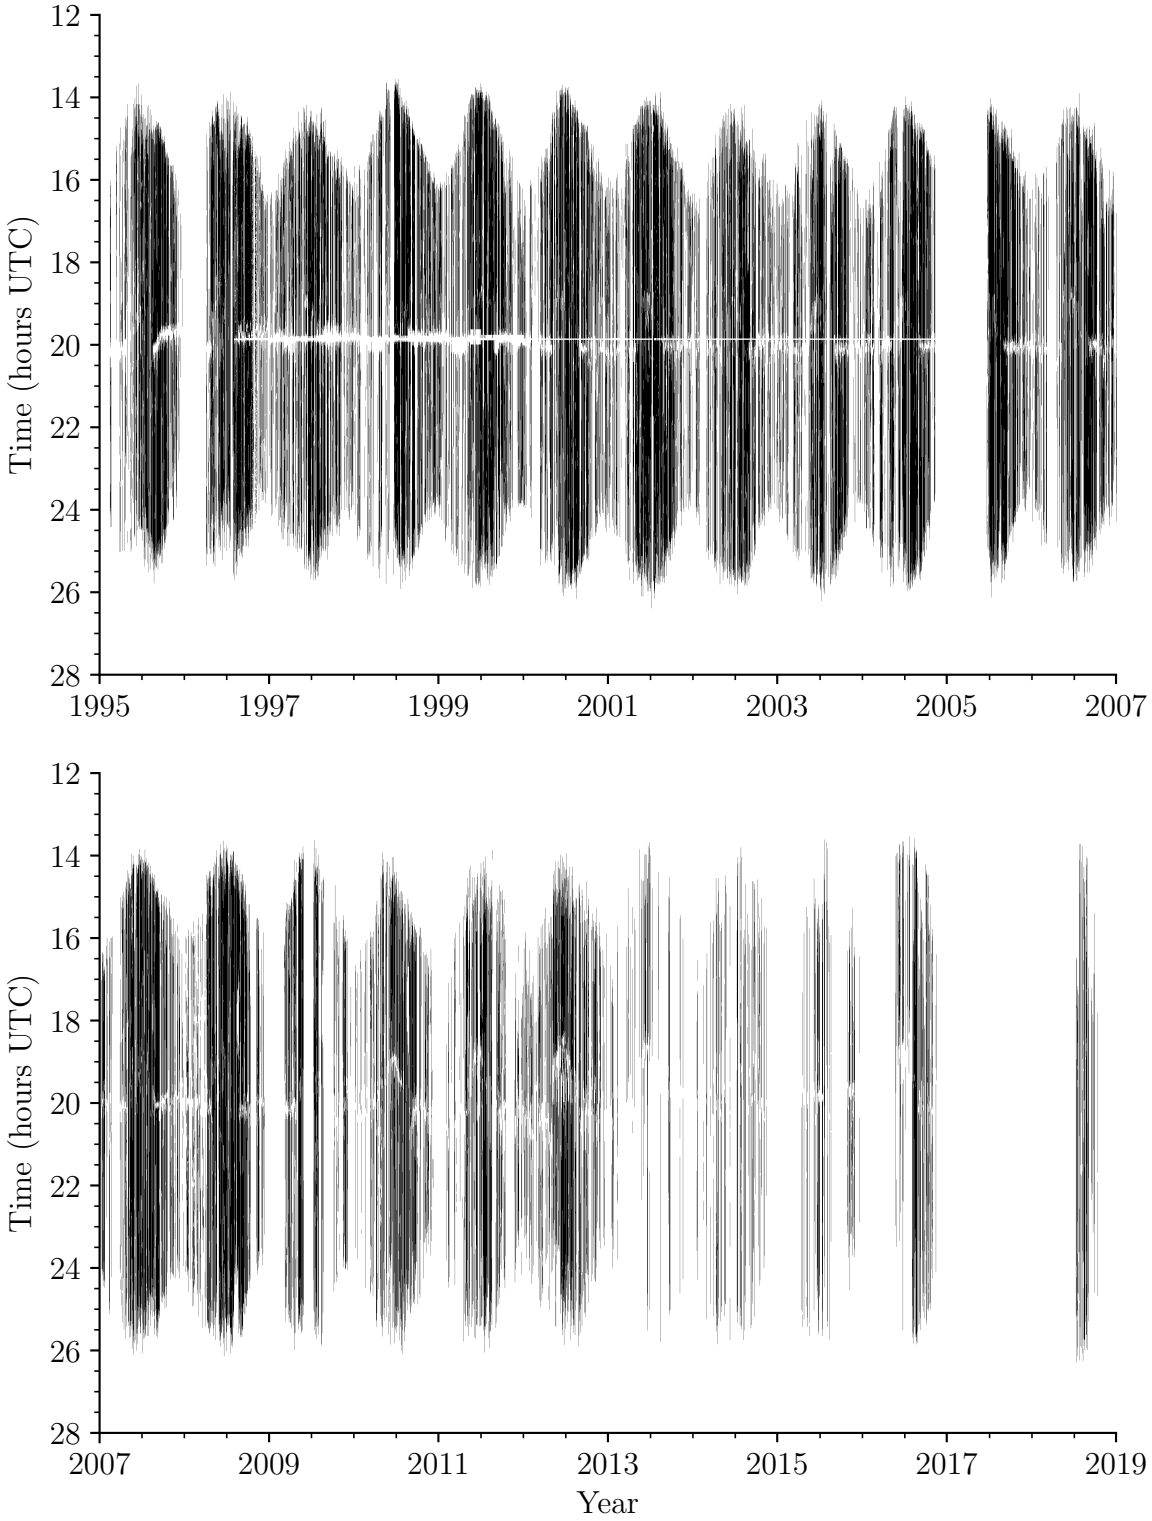

Figure 4.29: Mount Wilson data window-function.

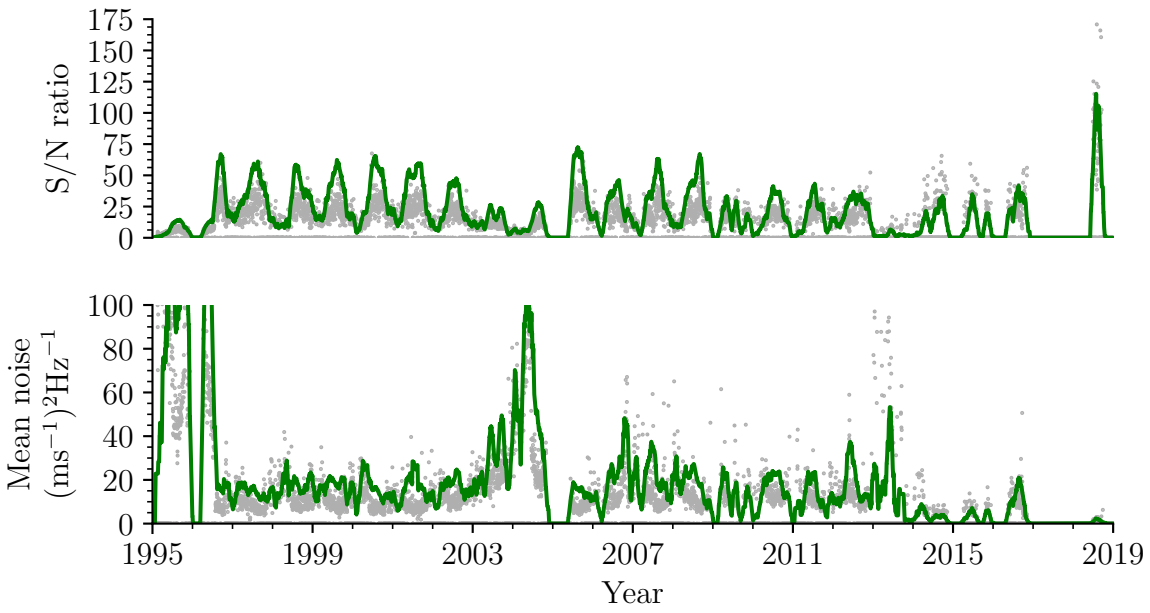

Figure 4.30: Mount Wilson data quality as a function of date. Top: Signal to noise ratio, higher is better. Bottom: Mean noise level, lower is better. There is one grey dot per day, and the solid green curve represents a 50-day moving mean. The large step in 1996 is due to installation of a new upgraded spectrophotometer, and in 2004 due to autoguider problems.

was replaced, and a PCB repaired. Mount Wilson received its Zoo upgrade in September 2005 (Miller, 2005), which included alignment of all the optics returning the site to its original performance.

In 2009 a variety of faults caused loss of data. The computer failed at the beginning of the year, and a site visit to replace it could not be arranged until March. Shortly after installing the new computer, the UPS failed in May and caused problems with the temperature controller. At the end of July the mirrors were removed for two weeks for re-aluminisation. In August the entire observatory was evacuated due to the nearby "Station Fire". There was concern that the entire observatory could be lost. Fortunately that was not the case, thanks to efforts from the fire service and the US Forestry Service.

Further guider problems were experienced in 2011, where one of the motors failed again and needed to be replaced. Through 2013 and 2014, significant problems were experienced with the autoguider electronics in the tower. On a site visit from Birmingham, the guider system built into the second flat of the cœlostat was completely rewired and refurbished,

removing some very old and corroded cables and producing a significant decrease in noise level (Hale, 2014a). Problems with the guiding electronics have been ongoing, with further attempts at repair in December 2015 (Hale, 2015e) and April 2016 (Hale, 2016b). During development work on this project, Klaus was trialled with a fibre-feed conversion in September 2016 (Hale, 2016d), and later converted permanently in June 2018 making use of a dedicated mount for the fibre collection optics (Hale, 2018b). This conversion bypasses the damaged mirrors and faulty guider, resulting in a substantial decrease in noise level and a typical FOM that has more than doubled when compared with previous years.

In the next section we will consider the performance of the network as a whole, where all the data are concatenated into one long time series.

4.4 Whole Network Performance

The whole-network duty cycle and data window-function are shown in Figures 4.31 and 4.32. In a total of 7305 days, 45.5 % of the available time was covered by one site, 30.3 % by two sites, 6 % by three sites, 0.06 % by four simultaneous sites. Just 9 days had no coverage, and 633 days have a fill greater than 99 % of which 312 days achieved 100 %. The average fill is 82 % for the whole dataset.

The original data pipeline for calibration of raw data from the BiSON spectrophotometers through to velocity residuals is described by Elsworth et al. (1995). The next stage of analysis involves combining the residuals into an extended time series, and transforming into the frequency domain where the mode characteristics can be analysed. This is described by Chaplin et al. (1997) and Hale (2003). An updated pipeline that includes fitting of differential atmospheric extinction was produced by Davies et al. (2014), and this corrects for the gradient of extinction across the extended source of the Sun at high airmasses. Uncorrected differential extinction causes the line-of-sight velocity due to solar rotation to become unequally weighted, and so introduces offsets at the beginning and end

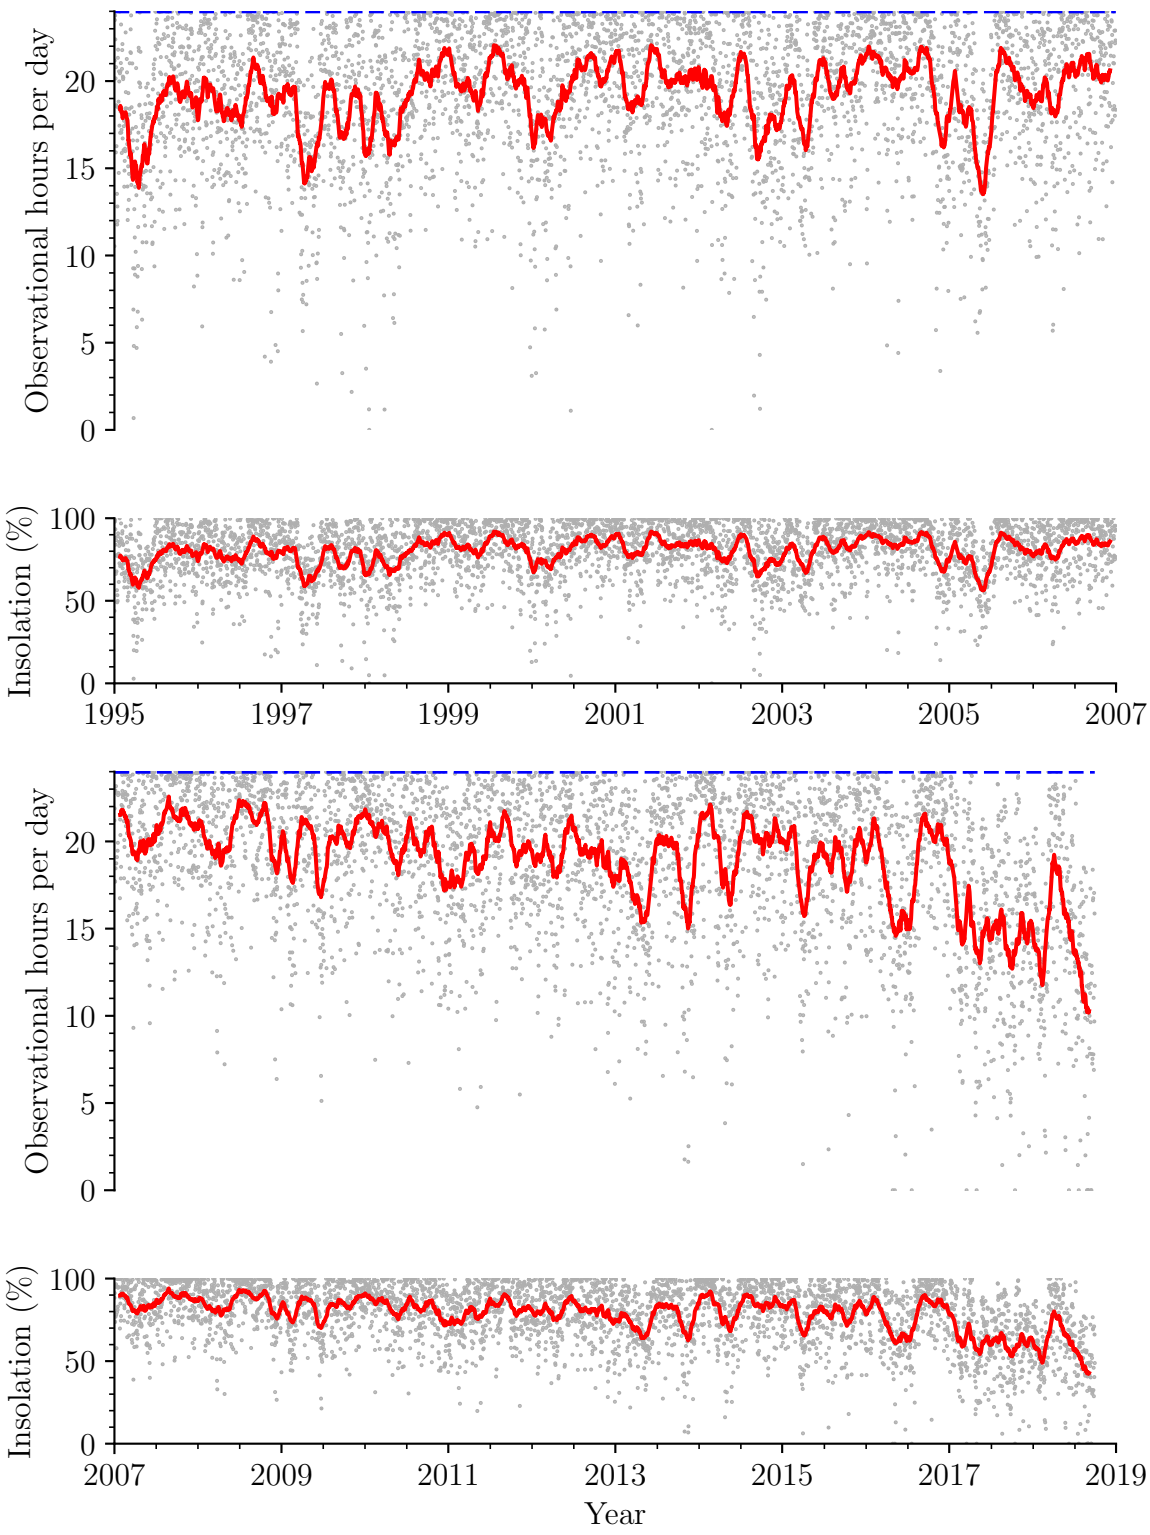

Figure 4.31: All-station duty cycle as a function of date, plotted in hours per day, and as a percentage of potential daylight hours. There is one grey dot per day, and the red curve represents a 50-day moving mean.

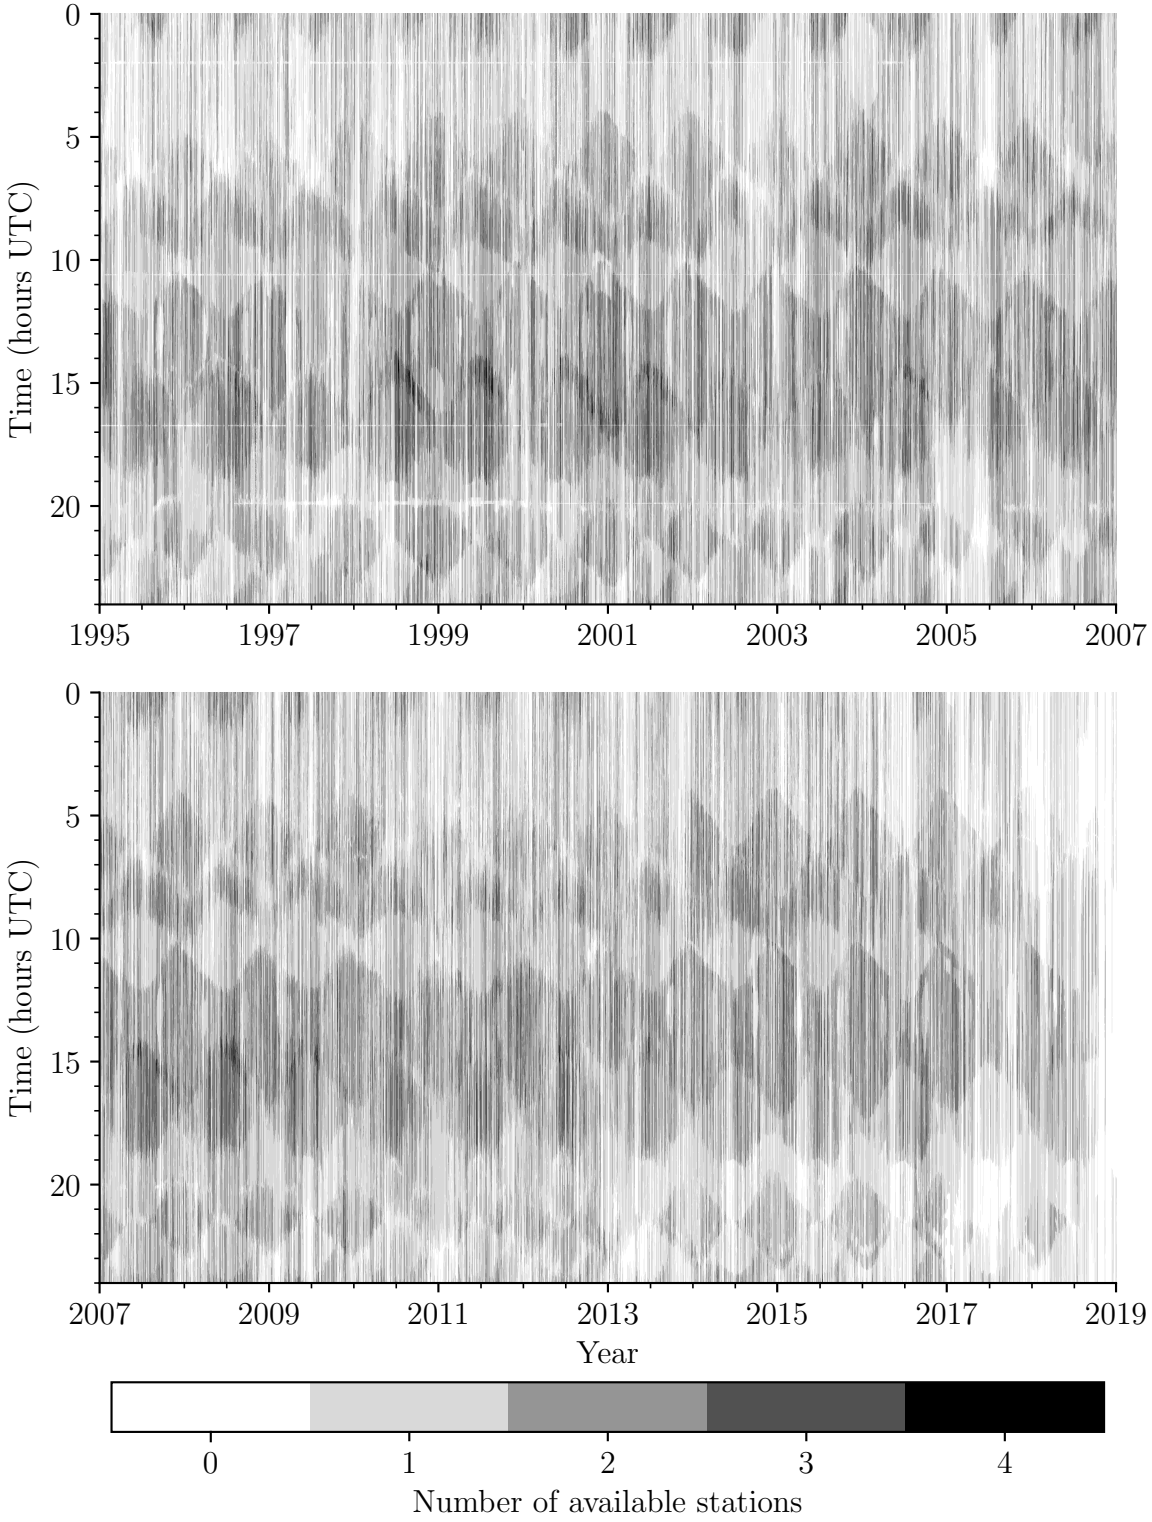

Figure 4.32: All-station data window-function.

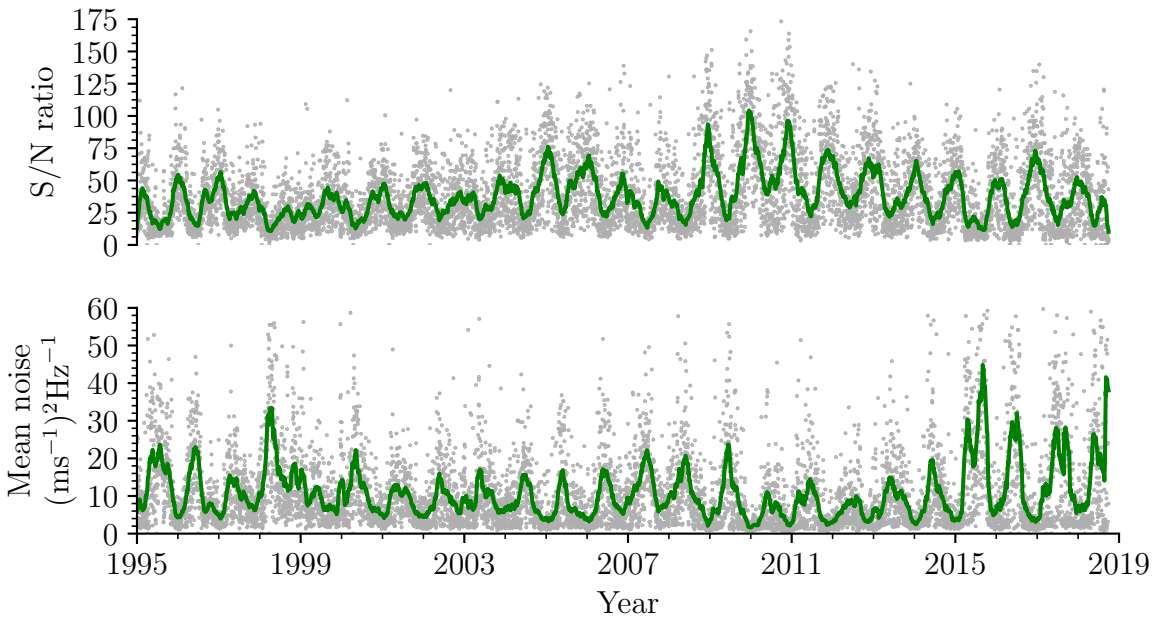

Figure 4.33: All-station data quality as a function of date. Top: Signal to noise ratio, higher is better. Bottom: Mean noise level, lower is better. There is one grey dot per day, and the solid green curve represents a 50-day moving mean.

of each day. By applying differential extinction correction we have removed most of the low-frequency drifts in the dataset that were previously filtered using a 25-sample moving mean, and this allows further investigation of the very low-frequency modes of oscillation that were previously lost in the noise background. We are now also able to make better use of weighted averaging of overlapping sites to produce a further improvement in signal to noise ratio.

The entire network archive of velocity residuals has been regenerated using the latest data pipeline, and a new concatenated time series has been produced for the period from 1995 to the end of 2014 using the latest weighted-averaging techniques. A noise ceiling of 100 (ms^{-1})^{2}Hz^{-1} was selected to reject data above this level, and this has reduced the overall fill from 82 % to 78 %. The plot of data quality and noise levels (Figure 4.33) shows excellent stability through the entire period. The seasonal variation in noise level, and subsequent change in FOM, is due to an effect described by Chaplin et al. (2004c, 2005b). As we saw in Chapter 3, the basic measurement of a BiSON RSS is that of intensity change due to the shift of a solar Fraunhofer line. As the Fraunhofer line shifts

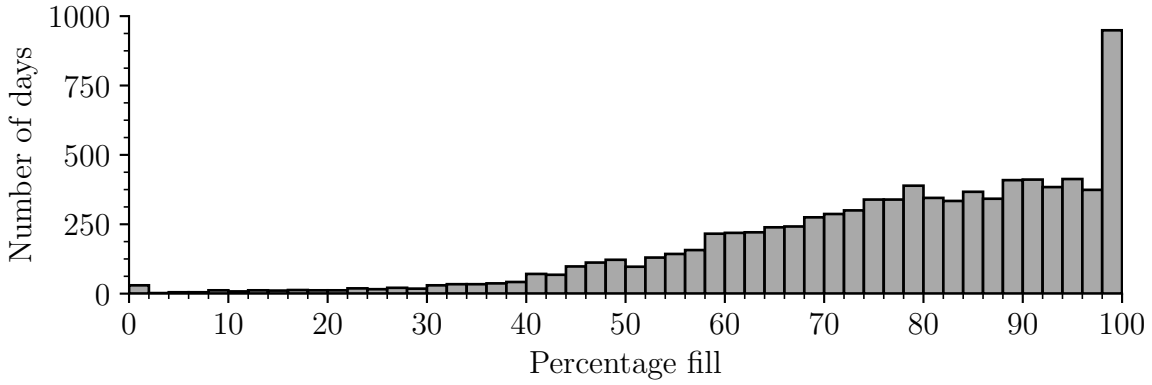

Figure 4.34: The distribution of fill per day for the overall time series.

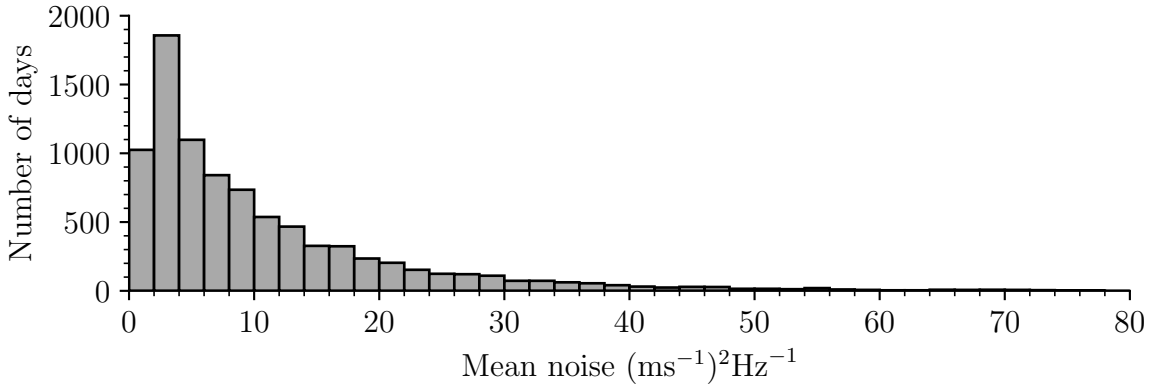

Figure 4.35: The distribution of mean noise per day for the overall time series. Lower is better.

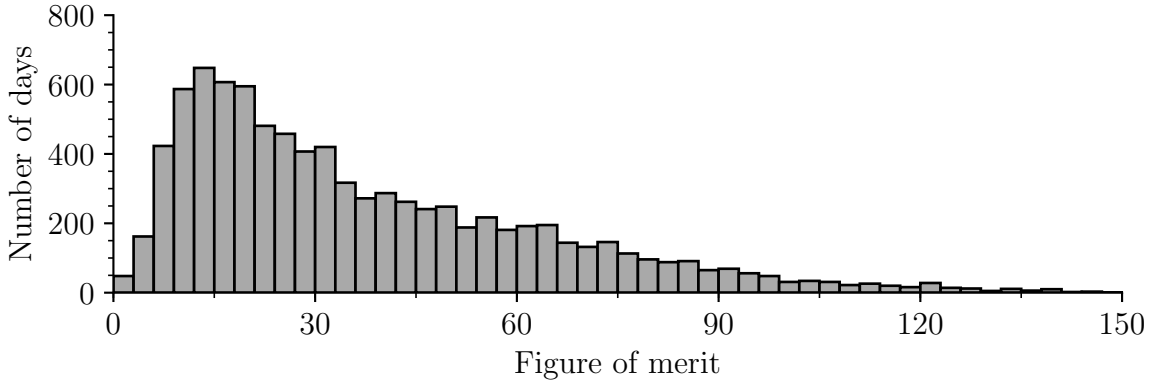

Figure 4.36: The distribution of FOM per day for the overall time series. Higher is better.

due to the line-of-sight relative motion of the solar surface with respect to the laboratory, the measured intensity changes. The data pipeline calibrates intensity measurements into velocity, taking into account the non-linearity of the Fraunhofer line shape. However the propagation of noise on the original intensity measurement through the calibration process is non-linear (Hoyng, 1989) due to the varying types of noise. The system encounters white noise, for example due to photon statistics and analogue-to-digital conversion, multiplicative noise due to gain fluctuations, and additive noise due to offset fluctuations, in addition to more random effects due to pointing errors and temperature fluctuations. We will look at different types of noise in resonance scattering spectrophotometers in Chapter 8. As the Fraunhofer line shifts daily due to Earth's rotation, and seasonally due to the eccentricity of the Earth's orbit, the line gradient at the operating point of the instrument changes. This causes the noise level in the derived quantities to change even if the noise level in the basic intensity measurement remains constant. Figure 4.34 shows a histogram of the daily fill, and Figures 4.35 and 4.36 show histograms of noise level and FOM, respectively. The new data pipeline provides an improvement across the entire historic archive, not just new data, and as such provides an exciting opportunity for new science.

4.5 Conclusion

Throughout this chapter we have assessed the data quality from each BiSON site. In Section 4.2 we introduced two robust quality metrics, and in Section 4.3 applied the metrics to data from each site for the period 1995 to 2018. Finally in Section 4.4 we considered the performance of the network as a whole, making use of concatenated data from multiple sites and instruments, and also discussed improvements to the data through a more sophisticated data processing pipeline.

In the next Chapter we will assess site performance in terms of atmospheric extinction, complementing the work here on weather conditions and instrumentation noise levels. We

will make use of these statistics from Chapter 7 onwards when comparing data quality from new instrumentation with that from the existing network.

Chapter 5

Atmospheric Extinction

Contents

5.1 Introduction

> This chapter includes text and figures from the publication "Atmospheric Extinction Coefficients in the I_c Band for Several Major International Observatories: Results from the BiSON Telescopes, 1984–2016" (Hale et al., 2017) in the journal *Astronomical Journal*. The published article was first released on 2017 August 11, and is included in Appendix E for reference. Data from this publication can be downloaded from the University of Birmingham eData archive (Hale, 2017a). Whilst this publication is a result of collaborative working, the content is entirely the author's own work. Co-authors provided guidance on data processing and manuscript review.

In the previous chapter we considered the performance of the network in terms of duty cycle and overall noise levels. In this chapter we will consider site performance in terms of total atmospheric column-aerosol extinction coefficient.

Since BiSON observations are disc-integrated Sun-as-a-star, one might assume that atmospheric extinction has little effect on the data quality, aside from the usual difficulties associated with ground-based photometry. However, only the final data product is a

single-pixel average. The initial observations involve imaging by virtue of the Sun being an extended source. At high airmasses the Sun suffers differential extinction across the disc, and this causes the line-of-sight velocity due to solar rotation to become unequally weighted. Without a uniform Sun-as-a-star average, the line-of-sight velocity measurement becomes biased by the solar rotation (Davies et al., 2014). Variations in atmospheric scintillation across the disc also contribute to noise within the line-of-sight velocity measurement due to non-uniform disc averaging. We see this variation in noise level as a seasonal effect.

We begin in Section 5.2 by describing how the total atmospheric column-aerosol can be derived from BiSON data in terms of aerosol optical depth (AOD). In Section 5.3 we consider how line-of-sight velocity offsets can introduce a small systematic offset on the AOD estimate. In Section 5.4 we present the historic extinction coefficients of each site, and then finally in Section 5.5 consider the overall performance of the whole network.

5.2 Deriving extinction coefficients

Atmospheric extinction has three main components: Rayleigh scattering, scattering due to aerosols, and molecular absorption. The strongest absorption effects are due to molecular oxygen and ozone which both absorb in the ultraviolet, and water vapour which absorbs in the infrared. At the BiSON observational wavelength of 769.9 nm, Rayleigh scattering is at a level of a few percent, and there is no molecular absorption: the observed atmospheric extinction is dominated by the contribution of aerosols.

The Beer-Lambert law states that the transmittance, T, of a material is related to its optical depth, τ_λ, by

$$T = \frac{I}{I_0} = e^{-\tau_\lambda A} \,, \tag{5.1}$$

where I_0 is the solar extraterrestrial radiance (i.e., at zero airmass), and I is the direct-Sun radiance. In this case τ_λ is the column atmospheric aerosol optical depth (AOD) per unit airmass, and A is the relative optical airmass as a function of solar zenith angle.

The aerosol optical depth is typically quoted as unitless when considering only the unit airmass at the zenith. By taking the natural logarithm of both sides we obtain,

$$\ln(I/I_0) = -\tau_\lambda A \ , \tag{5.2}$$

which gives a convenient linear relationship where the gradient of the relationship is a measure of τ_λ, and I_0 is now simply a normalisation factor taken as the maximum intensity measured on a given day. For astronomical use, we rescale AOD in terms of magnitude,

$$\begin{aligned} \kappa_\lambda &= -2.5 \log_{10}(e)(-\tau_\lambda) \ , \\ &= 1.086\tau_\lambda \ , \end{aligned} \tag{5.3}$$

where κ_λ is the atmospheric extinction coefficient, with units of magnitudes per airmass. More accurately, this is the column extinction coefficient since we do not include any knowledge on the vertical structure of the atmosphere. Readers in the climate modelling and aerosol communities should divide the values presented here by 1.086 in order to recover the total column-aerosol in terms of aerosol optical depth (AOD).

In this analysis, the known zenith angle was used to calculate the airmass based on Kasten and Young (1989) who define the airmass as,

$$A = \frac{1}{\cos z + 0.50572(6.07995 + 90 - z)^{-1.6364}} \ , \tag{5.4}$$

where the zenith angle z is in degrees. This model gives an airmass of approximately 38 at the horizon, producing good results for the whole range of zenith angles.

The extinction coefficients were determined using two methods: Firstly by making a standard linear least-squares fit of the magnitude-like value from equation 5.2 against airmass, and secondly by calculating the non-overlapping independent first-differences and then obtaining statistical estimations from the histogram of a time series of,

$$\frac{\mathrm{d}\, m}{\mathrm{d}\, A} = \frac{m_i - m_{i-1}}{A_i - A_{i-1}} \ , \tag{5.5}$$

where m is the magnitude-like value from equation 5.2, A is airmass, and i is the sample

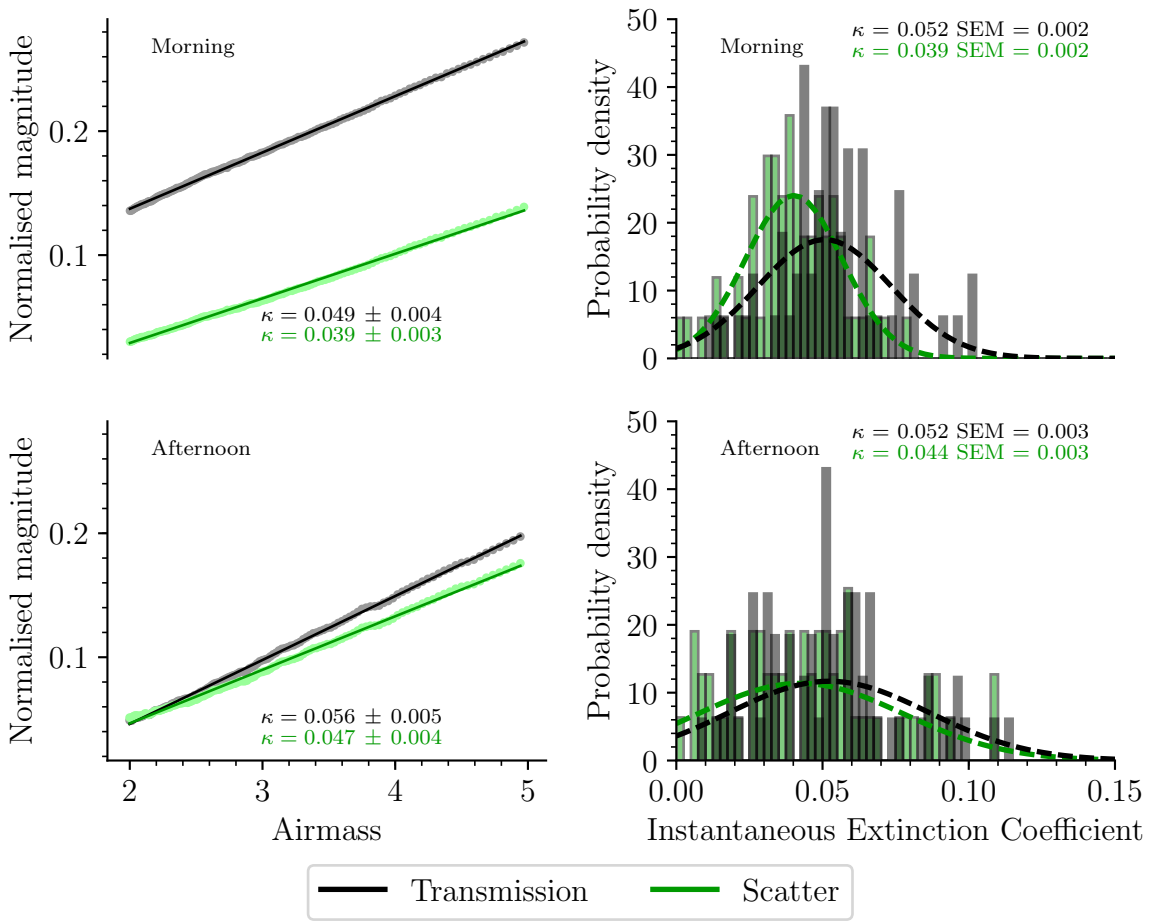

Figure 5.1: Morning and afternoon extinction coefficients for 2014 December 1 at Sutherland, South Africa. Left column: The formal uncertainty on the fit is calculated from the covariance matrix returned by the fitting function. Right column: The uncertainty on the histogram is the standard-error of the median (SEM). The dashed lines on the histograms represent for comparison the equivalent Gaussian profile for the measured mean and standard deviation.

index incremented in steps of two. In our fits we consider only air masses in the range of 2 to 5, corresponding to zenith angles between approximately 60° to 80°, since this is the region where airmass is changing most linearly. Below two airmasses the change does not follow a strictly linear relationship and is not well described by a straight line fit. The rate-of-change is also too low to allow good fitting. Additionally, we need to ensure we remove the seasonal variations in minimum airmass due to the changing maximum altitude of the Sun throughout the year, since this could introduce an artificial seasonal effect in the derived extinction values. Above five airmasses the rate of change is too high, producing

differential extinction across the extended source of the Sun (Davies et al., 2014). The pre-meridian and post-meridian values (hereafter referred to as 'morning' and 'afternoon') are fitted separately, since it is expected that these will differ due to local environmental considerations. The results from both techniques for a typical day are shown in Figure 5.1. Extinction coefficients are derived from both the vapour cell scattering detectors (green trace), and from the directly transmitted light (black trace) where transmission detectors are available. The estimates produced from the scattering detectors are typically a slight underestimate in comparison to that measured from the directly transmitted light, and this is due to line-of-sight velocity offsets that we will consider further in the next section.

The coefficient estimation technique and the selection of airmass range-limits affects the value of the determined extinction. In order to investigate the robustness of our parameters, a randomisation trial was performed where rather than fixing the lower and upper airmass limits at 2 and 5 respectively, they were randomly selected each day between 2–3 and 4–5 airmasses. Five realisations of fitted extinction gradients from each site were generated, where each day in each realisation used a randomised upper and lower airmass limit between the stated bounds. These realisations were then compared with the extinction gradients measured when the airmass limits were fixed. A similar comparison was made between the time series of fitted gradients, and a time series of median first-differences. In all cases, the mean difference was less than 4 % of the mean extinction. The standard deviation of the difference was 3 to 6 times lower than the measured extinction standard deviation. Any systematic offsets or increase in scattering due to the processing techniques are at a level significantly lower than that due to real physical effects. These two techniques are considered to be equivalent, and the results presented here are produced solely from the first method using linear least-squares fitting.

In the next section we will look at how line-of-sight velocity offsets can introduce a small systematic offset on the AOD estimate, depending on whether the solar intensity data is taken from the scattering detectors or the transmission monitor in each instrument.

5.3 Line-of-sight Velocity Effects

We saw in Chapter 3 that BiSON solar spectrophotometers use a potassium vapour reference cell to precisely measure the intensity at two points within the wings of the potassium Fraunhofer line at 770 nm. From this, the line-of-sight velocity between Earth and the solar surface can be determined. It is possible to use the sum of these two intensities as a proxy for a measurement of the intensity of the continuum outside the absorption line that would be measured from the light transmitted directly through the instrument. However, the scattered sum does not correctly recover the true total intensity, but rather suffers a systematic underestimation that varies with Doppler velocity.

If the solar absorption line is modelled as a thermally and rotationally broadened Gaussian, as we saw in Chapter 3, then we can assess the effect of this underestimate of intensity. Figure 5.2 shows the expected effect of this dependency on line-of-sight velocity in the absence of atmospheric extinction over the range of line-of-sight velocities experienced during a typical day. The amplitude of the daily change in effective intensity varies throughout the year with changes in Earth's orbital velocity, and this creates a very small seasonal effect in the scattering extinction coefficients. The increased effective intensity is most pronounced at high red-shifts experienced during the local afternoon of each site. The gradual increase in effective intensity will cause any extinction coefficients derived from the scattered sum to have a slight overestimate in the morning and a larger underestimate in the afternoon.

If the precise shape and depth of the solar absorption line were known, then it would be possible to correct for this effect and recover the equivalent transmitted intensity. Unfortunately the solar disc-integrated line profile changes with magnetic activity and so one cannot correct based on a single line profile. Barreto et al. (2014) have resolved this problem for our data from Izaña by making use of a calibration technique where they have achieved a *mean bias* (defined as the mean difference between the transmission and scattering extinction coefficients) of ≤ 0.01.

In practice, extinction data from the scattering detectors offers a slight underestimate

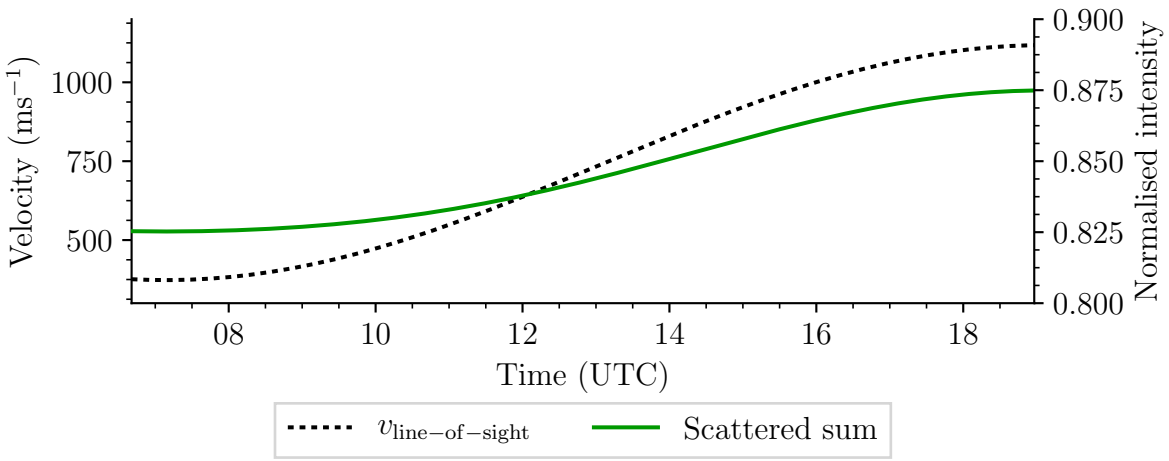

Figure 5.2: The expected variation in scattered sum over typical line-of-sight Doppler velocities, in the absence of any atmospheric extinction. Velocity is measured against the left axis, and normalised sum intensity measured against the right axis. The amplitude of the daily change in effective intensity varies throughout the year with changes in Earth's orbital velocity, and so creates a very small seasonal effect in extinction coefficients derived from the scattered sum.

of extinction in comparison to transmitted measurements for both morning and afternoon periods. However, this is simply a small systematic offset on the absolute value, and the extinction coefficients observed from real data are otherwise identical to those derived from the transmitted light and show the same trends.

In the next section we present fitted extinction coefficients for all the historic data from each BiSON site. For clarity the units of extinction will no longer be stated on each value in the text. All extinction values are specified in magnitudes per airmass. During the discussion for each site, we quote either mean, mode, or median values depending on the values required for comparisons with other studies. For consistency, a summary of the coefficients is given in Table 5.2. Since there are significant seasonal differences, we also present the values measured over two months for each mid-summer (July-August in the northern hemisphere, January-February in the southern hemisphere) in Table 5.3, and during mid-winter (January-February in the northern hemisphere, July-August in the southern hemisphere) in Table 5.4.

5.4 Site Extinction

5.4.1 Izaña, Tenerife

The BiSON node at Tenerife (Roca Cortés and Pallé, 2014) is based at the Observatorio del Teide, which is operated by the IAC (Instituto de Astrofisica de Canarias). The Canary Islands are located about 100 km to the west of the North African coast. The islands are close to the Western Sahara and so during the summer months they frequently experience high concentrations of mineral dust in the atmosphere. Events where the Saharan Air Layer passes over the Canary Islands are known as *Calima*, and these cause a fog-like reduction in visibility.

There are a number of telescopes providing weather information at the IAC. The STELLA Robotic Observatory provides several environmental parameters, including measurements of ground-level dust (STELLA Robotic Observatory, 2017) from a VisGuard 2 In-situ Visibility Monitor manufactured by SIGRIST Photometer (SIGRIST Photometer, 2017). The VisGuard is a device intended for use in vehicle tunnels to monitor visibility and carbon monoxide concentration, and used to control the ventilation system to minimise energy use or close the tunnel if limits are exceeded. It is also used as an early warning system for fires. The VisGuard is essentially a photometer that measures the intensity of scattered light from an air sample drawn into the instrument by a fan. The operation is much the same as our own BiSON spectrophotometers. The photometer illuminates the sample at a wavelength of 880 nm and compares the light scattered at a 30° angle with the directly transmitted light. This two-beam system eliminates effects due to temperature, ageing of the electronics, or fluctuation of the light source. The output is a measure of Polystyrol-Latex-Aerosols (PLA) which is then multiplied by some factor to obtain the extinction coefficient. For example, for use in tunnels the PLA is multiplied by 30 to produce extinction per metre (E/m). The STELLA web page provides values in terms of PLA per cubic-metre. However, in terms of dust measurement the absolute value is not particularly important so long as we have some qualitative idea of what is

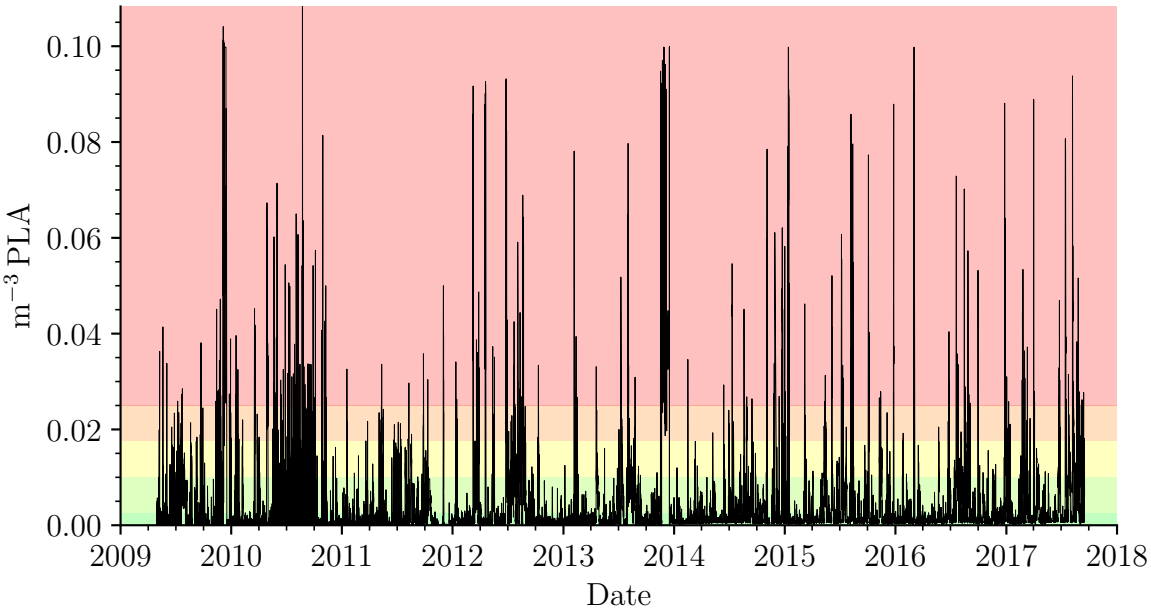

Figure 5.3: Dust data from the STELLA telescope. The coloured bands indicate the qualitative limits described in Table 5.1.

considered low and what is considered to be high. The IAC have defined some threshold levels, and these are detailed in Table 5.1. The data from the sensor is shown in Figure 5.3 with coloured banding from green to red to indicate the threshold levels.

The VisGuard is a useful real-time estimate of dust levels, but since it is measuring dust at ground level rather than a column through the whole atmosphere the results do not necessarily correlate with those determined from a whole day of BiSON data. The dust and aerosol concentration as measured by Mark-I is easily seen in Figure 5.4 as the strong seasonal variation in extinction. The modal values of the extinction distributions

Table 5.1: Qualitative dust limits used at the IAC.

Qualitative Dust Measurement		Range (m^{-3} PLA)	
Spanish	English	Min	Max
Poca	Little		<0.0025
Media	Half	0.0025	<0.0100
Bastante	Quite	0.0100	<0.0175
Mucha	A lot	0.0175	<0.0250
Fuera limites	Outside limit	0.0250	

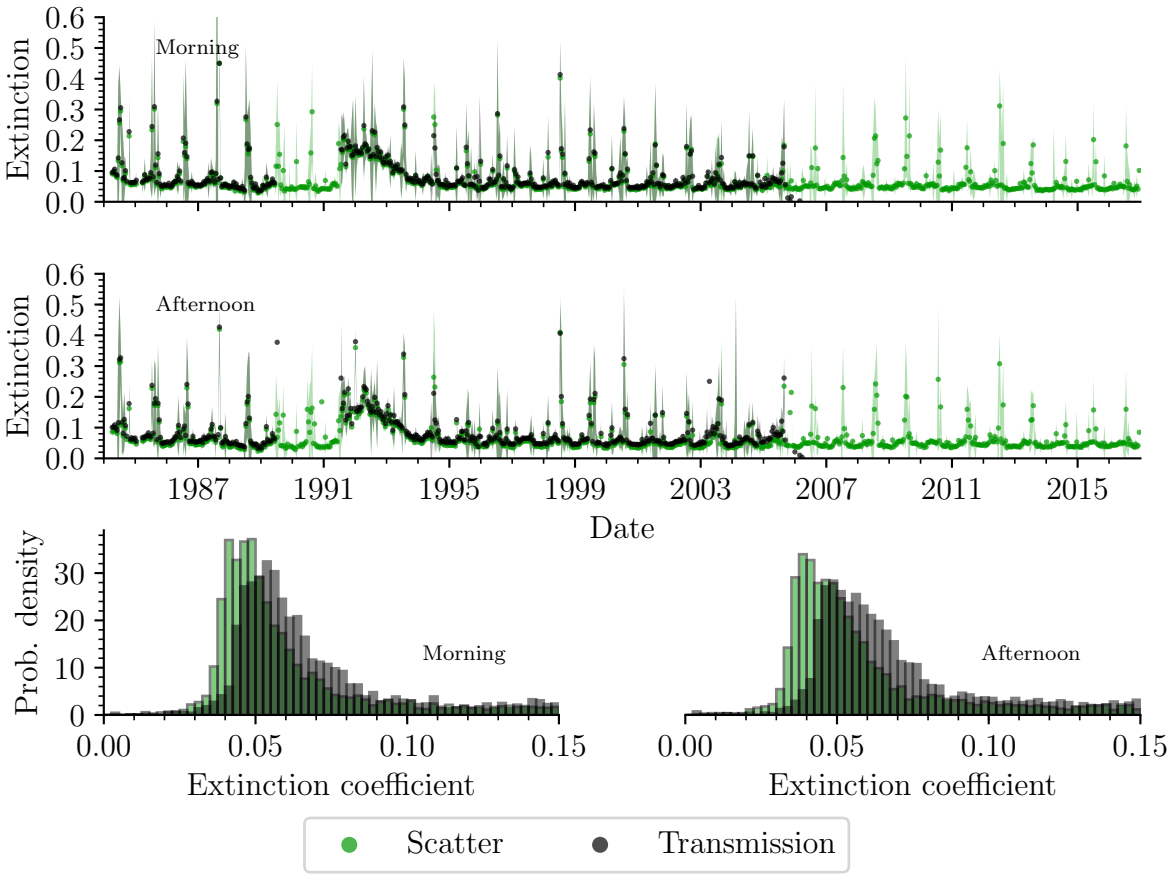

Figure 5.4: Extinction coefficients and statistical distribution from Izaña, Tenerife. Each dot represents the median value over 14 days. The coloured banding represents ±3 times the standard-error on each median. Dates indicate January 1 for each year.

are 0.054 in the morning and 0.046 in the afternoon for transmission. The standard deviation on these values is 0.09, the high value being indicative of the large scatter in values between summer and winter periods. The difference in the values pre-meridian and post-meridian is likely due to the surrounding geography, where morning extinction effects are through atmosphere over North Africa and afternoon is over the potentially clearer North Atlantic ocean. If only the winter months are considered, then the modal value for both morning and afternoon drops to 0.045 with correspondingly reduced scatter. Mineral dust events are typical between June and October, where extinction values anywhere from 0.1 to 0.8 may be experienced. The extinction values derived from the scattered-light are a slight underestimate as expected, but otherwise show the same trends. The distribution of extinction coefficients appears to show two combined trends, the first a

set of normally-distributed values centred on approximately 0.05, and a long positive tail that corresponds to the periods of mineral dust events. If the summer and winter periods are analysed separately, then the winter does indeed show a mean of approximately 0.05 to 0.06, and the summer a much higher mean of 0.14 to 0.15 with correspondingly higher standard deviation. A similar atmospheric analysis on data from telescopes on the Canary Archipelago, which included our BiSON data from this instrument, was made by Laken et al. (2016) and the modal values are consistent given the uncertainties. Laken et al. (2016) provides a thorough investigation into the occurrence of dust events, and how they change over both short seasonal periods and longer time scales. Several authors have investigated this in more detail – see, e.g., Guerrero et al. (1998); Jimenez et al. (1998); Siher et al. (2004); García-Gil et al. (2010); Laken et al. (2014).

Siher et al. (2002) present extinction values for the IRIS (*International Research of the Interior of the Sun*) site based at Izaña. IRIS was a similar network to BiSON and used a similar observational technique but made use of the shorter wavelength sodium absorption line at 589.6 nm. They quote an average extinction value of 0.111, which is slightly higher than the mean values found here. The higher value is expected due to their use of the shorter wavelength, since shorter wavelengths tend to suffer greater extinction. Jimenez et al. (1998) found the extinction to vary during 1984 to 1989 between 0.04 to 0.07 at 680 nm, which is in agreement with the values found here for Izaña. During dust storms Jimenez et al. (1998) reported values up to 0.8, which again is in agreement with our findings. The GONG (*Global Oscillation Network Group*) site-survey (Hill et al., 1994) at Izaña measured an average extinction value of 0.1169 from 1985 September to 1993 July. GONG is another network similar to BiSON. Their initial site-survey used a normal incidence pyrheliometer (NIP) manufactured by Eppley Laboratories (Fischer et al., 1986) which has a wide spectral sensitivity range of 250 nm to 3000 nm. Light from the Sun is broadly like that of a black body at a temperature of 6000 K, meaning that the intensity peaks at a wavelength of approximately 500 nm and decays quickly towards the infrared. The value measured by the pyrheliometer will be strongly weighted towards the

peak wavelength of the solar spectrum. King (1985) discusses the wavelength dependence of typical expected atmospheric extinction from 300 nm to 1100 nm at the Roque de los Muchachos Observatory, on the adjacent island of La Palma. At 500 nm an extinction coefficient of 0.1244 can be expected, very close to the value determined during the GONG site-survey. At 300 nm, near the short end of the pyrheliometer sensitivity range, values well over 3 can be expected. The average extinction measured by the pyrheliometer will be much higher than the monochromatic values measured by BiSON at 769.9 nm, and unfortunately this means that no comparison can be made between the results from BiSON and the GONG site-survey.

Probably the most striking feature in Figure 5.4 is the increase in extinction following the eruption of Mount Pinatubo in the Philippines on 1991 June 15. There is also a hint of the tail-end of effects from the El Chichón eruption in Mexico in 1982 April where the start of the data in 1984 show extinction values around 0.1, double the typical value expected outside of a dust intrusion. Both Guerrero et al. (1998) and García-Gil et al. (2010) have previously observed these features from telescopes based at the Canary Islands.

5.4.2 Carnarvon, Western Australia

The BiSON node at Carnarvon is based at the historic *Overseas Telecommunications Commission Satellite Earth Station*, around 900 km north of Perth in Western Australia. The measured extinction coefficients over the operational lifetime of the site are shown in Figure 5.5.

The median morning extinction is 0.072, and the afternoon is 0.078. This increases to approximately 0.1 in the summer and decreases to around 0.06 in the winter. The standard deviation shows a similar increase in summer compared to winter. In the summer the standard deviation is noticeably lower in the afternoon compared to the morning, at 0.05 and 0.07 respectively. Morning data are collected over the plains of Western Australia, and afternoon data are through air over the Indian ocean, and so it is expected that the sandy environment of Carnarvon would have greater impact on extinction during morning

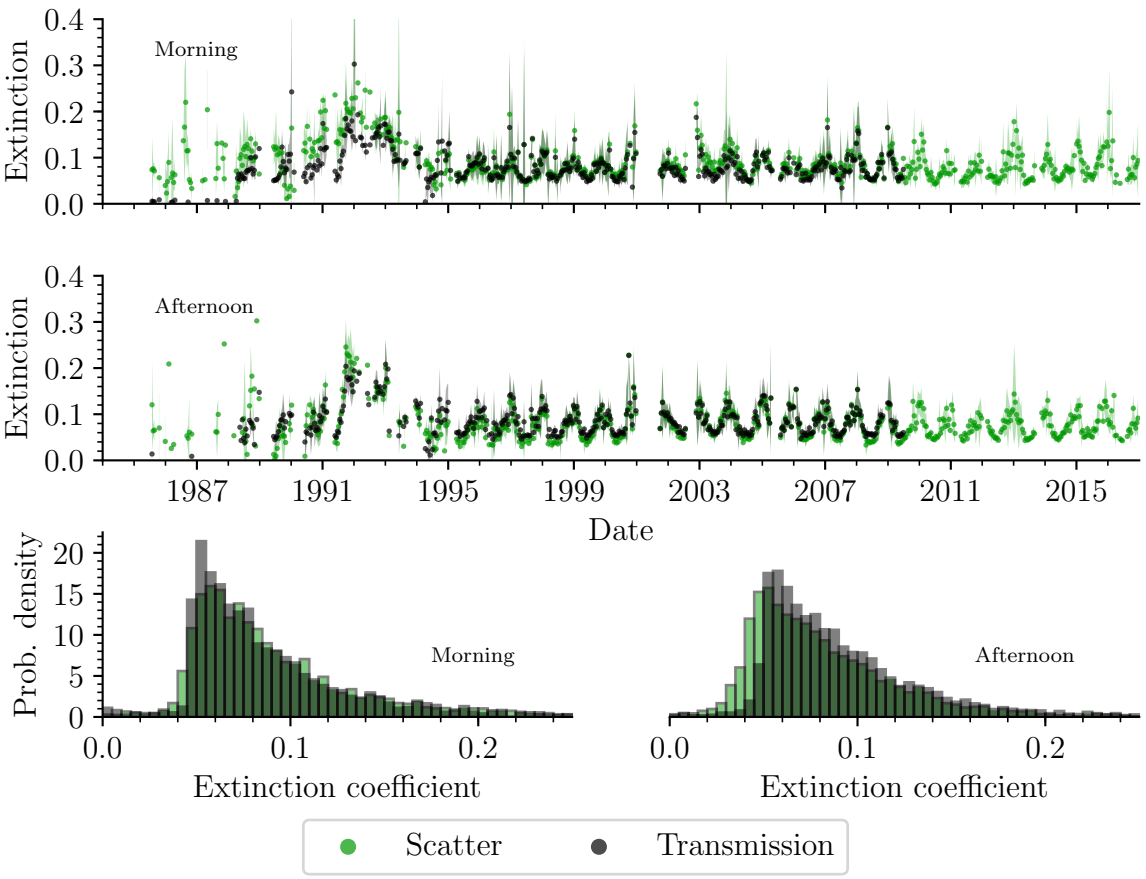

Figure 5.5: Extinction coefficients and statistical distribution from Carnarvon, Western Australia. Each dot represents the median value over 14 days. The coloured banding represents ±3 times the standard-error on each median. Dates indicate January 1 for each year.

observations. There are no comparison sites available near Carnarvon.

5.4.3 Sutherland, South Africa

The BiSON node at Sutherland is situated 360 km north-east of Cape Town, at the *South African Astronomical Observatory* (SAAO). The measured extinction coefficients over the operational lifetime of the site are shown in Figure 5.6.

The median morning extinction is 0.046, and afternoon is 0.045. There are no significant differences in environment around Sutherland, and this is reflected by the stability of the extinction coefficients between morning and afternoon. The site shows stable performance with little change in the median values throughout a year. The atmosphere

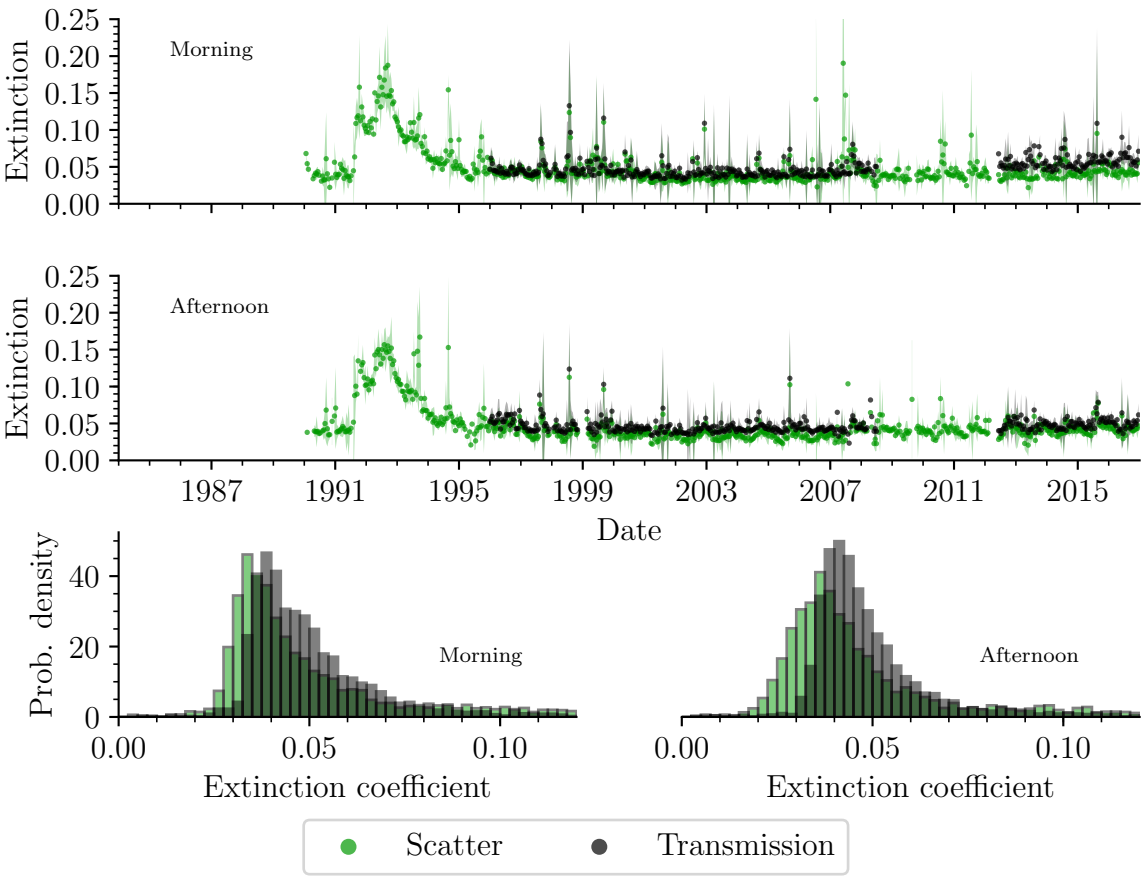

Figure 5.6: Extinction coefficients and statistical distribution from Sutherland, South Africa. Each dot represents the median value over 14 days. The coloured banding represents ±3 times the standard-error on each median. Dates indicate January 1 for each year.

is particularly stable, with standard deviations of between 0.04 and 0.06 in the winter, and approximately 0.02 during the summer. Kilkenny (1995) state a mean extinction coefficient of 0.07 which is slightly higher than found here, but is within our measured standard deviation.

5.4.4 Las Campanas, Chile

The BiSON node at Las Campanas is situated 630 km north of Santiago, at the *Las Campanas Observatory*. The measured extinction coefficients over the operational lifetime of the site are shown in Figure 5.7.

Transmission data from this site are unreliable due to instrumentation issues, and so

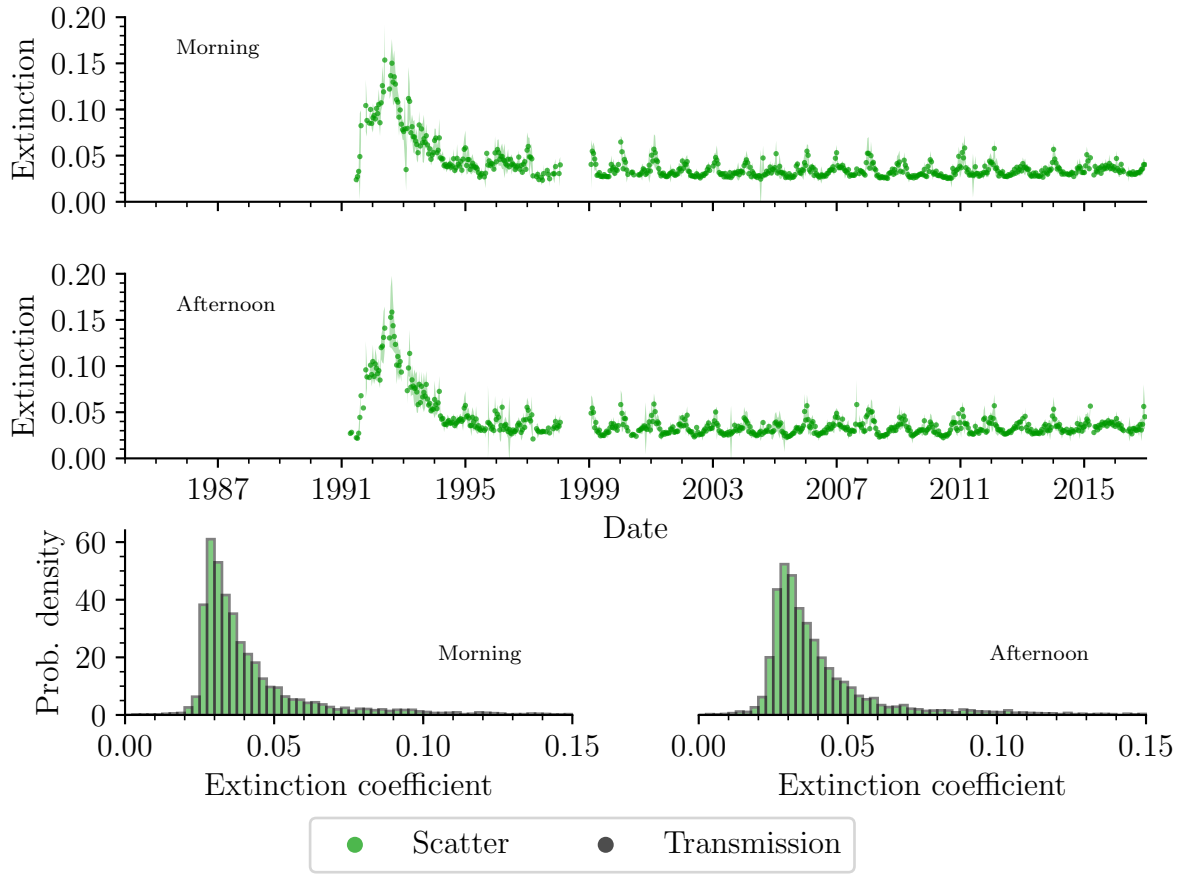

Figure 5.7: Extinction coefficients and statistical distribution from Las Campanas, Chile. Each dot represents the median value over 14 days. The coloured banding represents ±3 times the standard-error on each median. Dates indicate January 1 for each year. There is no transmitted-light measured at Las Campanas, and so only the scattering values are presented.

only the scattered-light data have been used to derive the extinction coefficients. The median extinction is 0.033 for both morning and afternoon, with standard deviation of less than 0.02. There is also little variation between winter and summer periods with mean extinction at 0.028 and 0.042 respectively. Morning data are collected through airmasses over South America, and afternoon data are through air over the South Pacific ocean. There appears to be no significant difference between the two zones, and Las Campanas has the most stable atmosphere of all BiSON sites, showing the lowest standard deviation.

Siher et al. (2002) performed a similar atmospheric analysis for the IRIS site based at La Silla, the adjacent mountain ridge to Las Campanas. They quote a value of 0.097, which is higher than found here, but we again have to consider the shorter wavelength used

by IRIS, and also that we expect the extinction from the scattered light to be a slight underestimate of the equivalent extinction from the transmitted light. The measured standard deviation is similar at 0.028.

5.4.5 Narrabri, NSW, Australia

The BiSON node at Narrabri is situated 525 km north-west of Sydney, at the *Paul Wild Observatory*. The measured extinction coefficients over the operational lifetime of the site are shown in Figure 5.8.

The median morning extinction is 0.074, and afternoon is 0.077. The standard deviation is approximately 0.04. There are no significant differences in environment around Narrabri and again this is reflected in the stability of the extinction coefficients between morning and afternoon. There is, however, a strong seasonal variation with median values dropping to 0.06 in the winter and rising to almost 0.08 in the summer. Narrabri tends to have lower extinction and lower standard deviation than the coastal town of Carnarvon, although it does suffer from a higher percentage of cloudy days overall. No datasets could be found for comparison with Narrabri.

5.4.6 Mount Wilson, California, USA

The BiSON node at Mount Wilson is situated 52 km north-east of Los Angeles, at the *Mount Wilson (Hale) Observatory*. The measured extinction coefficients over the operational lifetime of the site are shown in Figure 5.9.

The median morning extinction is 0.085, and afternoon is 0.081. Mount Wilson is an atmospherically-interesting site due to its location close to the major city of Los Angeles. Morning data are collected over the San Gabriel mountains, while afternoon data are observed through air over the city, and so it may be expected that afternoon extinction values would be worse. The median-values found here show that morning and afternoon are generally similar, however the modal values do show an increase from 0.048 in the morning to 0.070 in the afternoon. There is an upward trend in extinction levels from

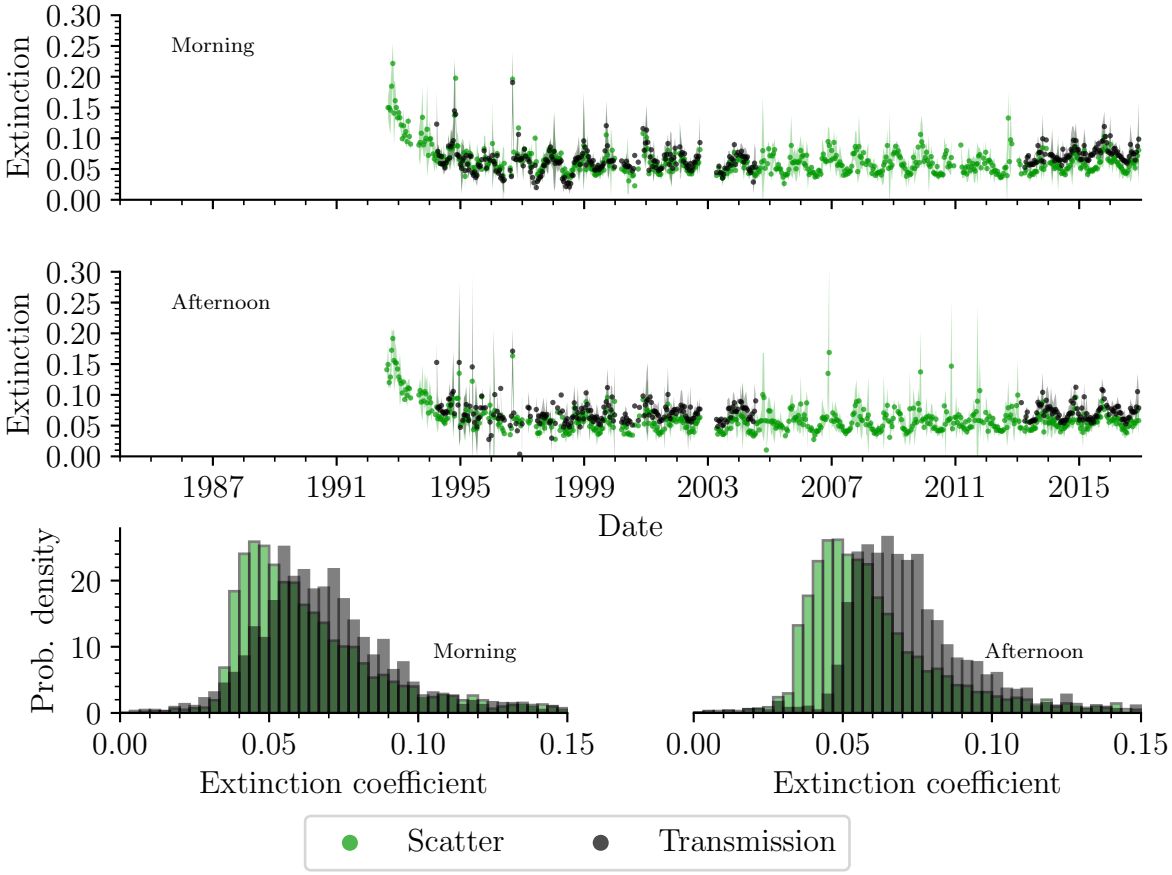

Figure 5.8: Extinction coefficients and statistical distribution from Narrabri, eastern Australia. Each dot represents the median value over 14 days. The coloured banding represents ±3 times the standard-error on each median. Dates indicate January 1 for each year.

around 1999, becoming from 2007 part of the large scatter in values that could be considered to be a result of the highly variable atmosphere near a large city. In clear sky conditions during our observations the inversion boundary layer traps any pollution at a relatively low altitude – the well-known Los Angeles smog. Since the observatory elevation is well above the inversion layer, the pollution would not be seen except at very high zenith angles that have been specifically excluded from this analysis. At this site, light is collected via two mirrors, known as a cœlostat, and it is possible that the increase in scatter is due to gradual reduction in performance due to deterioration of the mirrors; however, since the extinction is determined from a daily calibration, any variation in long term performance is removed completely and so this is just speculation and the cause of

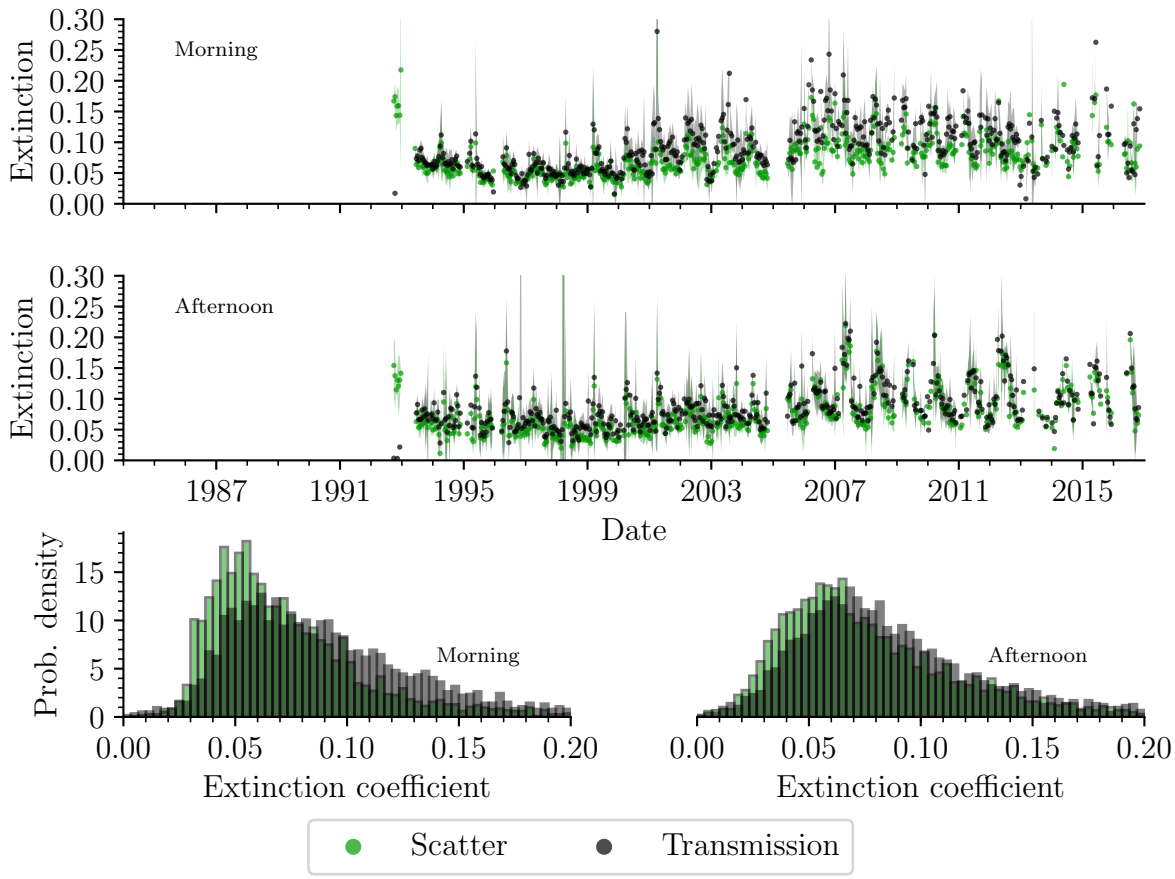

Figure 5.9: Extinction coefficients and statistical distribution from Mount Wilson, California, USA. Each dot represents the median value over 14 days. The coloured banding represents ±3 times the standard-error on each median. Dates indicate January 1 for each year.

the increased scatter is not clear. During the winter months, Mount Wilson does suffer a high proportion of cloudy days, but when the sky is clear there appears to be little variation between seasons, both in terms of absolute extinction and standard deviation.

5.5 Whole Network Summary

A summary of the extinction coefficients from 1995 over all seasons is given in Table 5.2. In order to compare seasonal differences, the values measured over two months for each mid-summer (July-August in the northern hemisphere, January-February in the southern hemisphere) are presented in Table 5.3, and during mid-winter (January-February in the northern hemisphere, July-August in the southern hemisphere) in Table 5.4.

Table 5.2: Extinction coefficients from all sites. Only 1995 onwards has been considered in order to remove any exceptional atmospheric events. The data-sets are sufficiently large that the standard error of the mean (i.e., $\sigma/\sqrt{N}$) is very small and therefore not presented.

Location	Detector	Mode	Median	Mean	Sigma
Mount Wilson	Transmission (Morning)	0.048	0.085	0.096	0.055
Mount Wilson	Transmission (Afternoon)	0.070	0.081	0.093	0.057
Mount Wilson	Scatter (Morning)	0.056	0.064	0.074	0.043
Mount Wilson	Scatter (Afternoon)	0.061	0.070	0.082	0.053
Las Campanas	Transmission (Morning)				
Las Campanas	Transmission (Afternoon)				
Las Campanas	Scatter (Morning)	0.029	0.033	0.037	0.018
Las Campanas	Scatter (Afternoon)	0.028	0.033	0.036	0.014
Izaña	Transmission (Morning)	0.054	0.058	0.092	0.090
Izaña	Transmission (Afternoon)	0.046	0.060	0.092	0.091
Izaña	Scatter (Morning)	0.048	0.051	0.082	0.084
Izaña	Scatter (Afternoon)	0.039	0.049	0.079	0.087
Sutherland	Transmission (Morning)	0.038	0.046	0.054	0.040
Sutherland	Transmission (Afternoon)	0.039	0.045	0.051	0.026
Sutherland	Scatter (Morning)	0.034	0.039	0.048	0.039
Sutherland	Scatter (Afternoon)	0.037	0.038	0.044	0.027
Carnarvon	Transmission (Morning)	0.052	0.072	0.085	0.050
Carnarvon	Transmission (Afternoon)	0.052	0.078	0.088	0.043
Carnarvon	Scatter (Morning)	0.057	0.074	0.087	0.049
Carnarvon	Scatter (Afternoon)	0.052	0.071	0.081	0.043
Narrabri	Transmission (Morning)	0.053	0.066	0.070	0.029
Narrabri	Transmission (Afternoon)	0.061	0.068	0.073	0.025
Narrabri	Scatter (Morning)	0.048	0.055	0.062	0.028
Narrabri	Scatter (Afternoon)	0.050	0.053	0.059	0.027

Table 5.3: Summer extinction coefficients from all sites. In the northern hemisphere the measured summer months each year are the beginning of July to the end of August, and in the southern hemisphere the beginning of January until the end of the February. Only 1995 onwards has been considered in order to remove any exceptional atmospheric events. The data-sets are sufficiently large that the standard error of the mean (i.e., $\sigma/\sqrt{N}$) is very small and therefore not presented.

Location	Detector	Mode	Median	Mean	Sigma
Mount Wilson	Transmission (Morning)	0.053	0.077	0.084	0.041
Mount Wilson	Transmission (Afternoon)	0.050	0.082	0.095	0.062
Mount Wilson	Scatter (Morning)	0.053	0.057	0.065	0.031
Mount Wilson	Scatter (Afternoon)	0.037	0.076	0.088	0.056
Las Campanas	Transmission (Morning)				
Las Campanas	Transmission (Afternoon)				
Las Campanas	Scatter (Morning)	0.032	0.042	0.046	0.017
Las Campanas	Scatter (Afternoon)	0.040	0.041	0.044	0.016
Izaña	Transmission (Morning)	0.057	0.073	0.151	0.133
Izaña	Transmission (Afternoon)	0.057	0.075	0.147	0.136
Izaña	Scatter (Morning)	0.048	0.070	0.142	0.130
Izaña	Scatter (Afternoon)	0.047	0.070	0.141	0.135
Sutherland	Transmission (Morning)	0.043	0.046	0.050	0.022
Sutherland	Transmission (Afternoon)	0.044	0.046	0.048	0.013
Sutherland	Scatter (Morning)	0.036	0.040	0.044	0.021
Sutherland	Scatter (Afternoon)	0.037	0.040	0.042	0.014
Carnarvon	Transmission (Morning)	0.085	0.100	0.115	0.078
Carnarvon	Transmission (Afternoon)	0.095	0.107	0.114	0.045
Carnarvon	Scatter (Morning)	0.119	0.106	0.119	0.068
Carnarvon	Scatter (Afternoon)	0.081	0.103	0.110	0.046
Narrabri	Transmission (Morning)	0.076	0.074	0.081	0.030
Narrabri	Transmission (Afternoon)	0.073	0.077	0.086	0.041
Narrabri	Scatter (Morning)	0.060	0.068	0.074	0.026
Narrabri	Scatter (Afternoon)	0.060	0.061	0.068	0.035

Table 5.4: Winter extinction coefficients from all sites. In the northern hemisphere the measured winter months each year are the beginning of January until the end of the February, and in the southern hemisphere the beginning of July to the end of August. Only 1995 onwards has been considered in order to remove any exceptional atmospheric events. The data-sets are sufficiently large that the standard error of the mean (i.e., $\sigma/\sqrt{N}$) is very small and therefore not presented.

Location	Detector	Mode	Median	Mean	Sigma
Mount Wilson	Transmission (Morning)	0.048	0.077	0.086	0.051
Mount Wilson	Transmission (Afternoon)	0.075	0.070	0.074	0.032
Mount Wilson	Scatter (Morning)	0.046	0.060	0.070	0.035
Mount Wilson	Scatter (Afternoon)	0.053	0.056	0.059	0.027
Las Campanas	Transmission (Morning)				
Las Campanas	Transmission (Afternoon)				
Las Campanas	Scatter (Morning)	0.028	0.028	0.031	0.015
Las Campanas	Scatter (Afternoon)	0.027	0.028	0.030	0.010
Izaña	Transmission (Morning)	0.045	0.048	0.059	0.035
Izaña	Transmission (Afternoon)	0.045	0.048	0.062	0.062
Izaña	Scatter (Morning)	0.042	0.043	0.051	0.030
Izaña	Scatter (Afternoon)	0.038	0.038	0.047	0.045
Sutherland	Transmission (Morning)	0.038	0.049	0.061	0.042
Sutherland	Transmission (Afternoon)	0.039	0.044	0.056	0.046
Sutherland	Scatter (Morning)	0.033	0.041	0.057	0.056
Sutherland	Scatter (Afternoon)	0.029	0.037	0.049	0.045
Carnarvon	Transmission (Morning)	0.051	0.056	0.065	0.028
Carnarvon	Transmission (Afternoon)	0.056	0.060	0.065	0.026
Carnarvon	Scatter (Morning)	0.051	0.060	0.068	0.032
Carnarvon	Scatter (Afternoon)	0.052	0.053	0.059	0.026
Narrabri	Transmission (Morning)	0.057	0.058	0.061	0.024
Narrabri	Transmission (Afternoon)	0.064	0.061	0.064	0.016
Narrabri	Scatter (Morning)	0.040	0.046	0.050	0.017
Narrabri	Scatter (Afternoon)	0.044	0.046	0.049	0.016

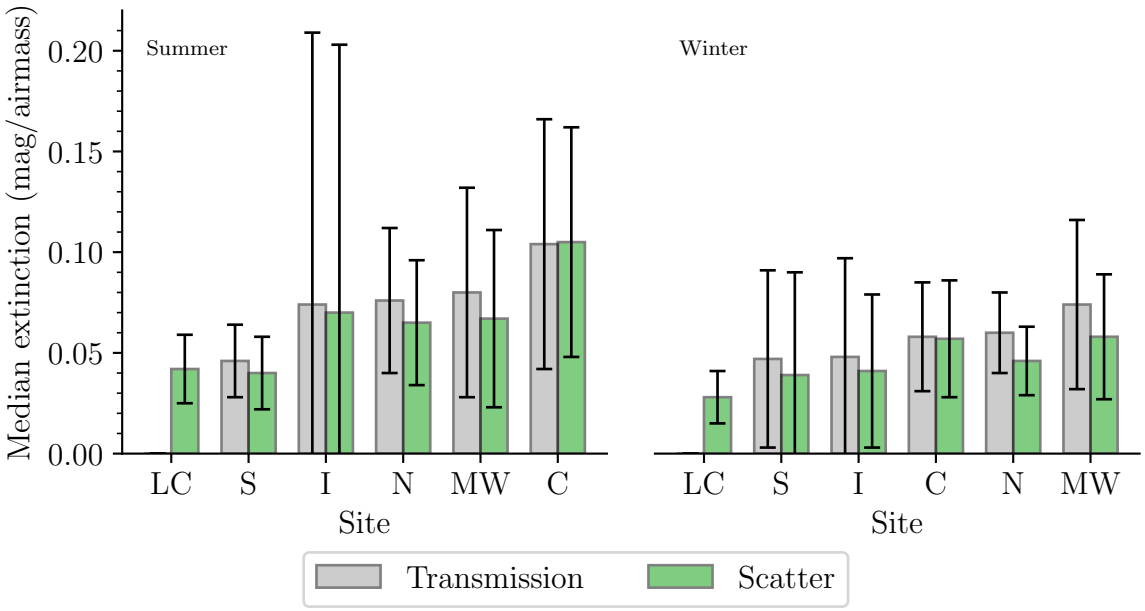

Figure 5.10: Summary of median extinction coefficients and distribution standard deviations. Only 1995 onwards has been considered in order to remove any exceptional atmospheric events. Left panel: Summer periods from Table 5.3. Right panel: Winter periods from Table 5.4. The mean of the morning and afternoon median coefficients have been taken for each site, and the extinction distribution standard deviation shown as error bars. The x-axis labels are the initials of each site name.

The median and standard deviation values for mid-summer periods from Table 5.3, and for mid-winter periods from Table 5.4, are shown graphically in Figure 5.10. The mean of the morning and afternoon values were taken for each site, and the extinction distribution standard deviation shown as error bars. We find the best results are from the Las Campanas and Sutherland observatories with consistent year-round performance. Izaña offers comparable high-performance during the winter months, but becomes the worst site during mid-summer due to high concentrations of mineral dust in the atmosphere from the Western Sahara on the North African coast. Carnarvon similarly suffers degradation during the summer months, which again is most likely due to wind-borne sand in the dry environment of northern Western Australia. Narrabri and Mount Wilson both offer consistent performance.

5.6 Conclusion

In this chapter we have made innovative use of BiSON data to derive measurements of atmospheric opacity, and we have presented the column atmospheric extinction coefficients from each site over the years 1984 to 2016. This is an important contribution to the literature since there are limited data on aerosol optical depth from other sources prior to the mid-1990s.

This completes our assessment of the historic overall performance of the BiSON network. In the next chapter, and those beyond, we will develop the new miniature instrumentation required when extending the existing network. We begin with modelling the potassium vapour cell at the heart of all BiSON instrumentation.

Chapter 6

Vapour Resonance Modelling

Contents

6.1 Introduction

We have already looked at the basic operation and calibration of data from the BiSON resonance scattering spectrophotometers in Chapter 3. In this chapter, we will look in more detail at the fundamental physical principles of operation by developing and applying a complete model of a typical vapour cell used in a BiSON spectrophotometer.

The effect of noise on velocity measurements from full-disc helioseismic observations made by resonant scattering spectrophotometers has been investigated in some detail by Hoyng (1989), and later specifically for BiSON by Chaplin et al. (2005b). Hoyng presented a theoretical analysis making use of qualitative arguments to estimate various contributions to instrumental sensitivity limits, such as photon noise, gain and offset fluctuations, and variations in vapour temperature. As Hoyng explains, this technique is only possible due to power spectra of noise sources obeying specific rules and while it is not always possible to predict the frequency dependence of noise sources using this technique,

it is possible to estimate parameters by comparing with a reference noise level such as that from solar granulation. Chaplin et al. extended the theoretical approach by analysis of simulations and application to real data from BiSON instrumentation. By making use of redundant data channels within a single instrument it is possible to separate the effects of some noise sources into two components – sources correlated between detectors which may be of either solar origin or instrumental origin such as vapour temperature fluctuations and atmospheric scintillation, and also sources uncorrelated between detectors such as amplifier noise. While an improvement on theoretical estimation this still does not allow the discrimination between all the different noise sources that contribute to the overall instrumental noise performance, which would potentially allow identification of sources that could be reduced either in hardware or through subsequent data processing. Rather than attempting to determine the effect of noise sources from data *a posteriori*, it would be better to assess noise characteristics during instrumental design and commissioning when individual components can be isolated and profiled. Whilst this is no longer possible for existing BiSON instrumentation, here we will begin an electrical engineering approach to new instrumentation design that will produce a performance specification that has previously been undocumented for any BiSON instrumentation.

In Sections 6.2 and 6.3 we will develop the theory of spectral line formation from first principles, and produce a model of a potassium vapour cell. In Section 6.4 we will use a tunable laser to probe the performance of a typical BiSON vapour cell, and in Section 6.5 validate the model against the measured performance in order to show that the model can be used to compare expected performance between existing BiSON vapour cells and commercially available alternatives.

Some work has been done previously on the optical depth of potassium vapour in BiSON instrumentation – see, e.g., Brookes et al. (1978b); Broomhall (2007); Davies (2011). These works concentrated on estimating the effect of instrumental weighting of solar rotation between different scattering detectors on opposite sides of the vapour cell based on an unknown value for the optical depth – an effect known as Doppler Imaging.

The new work here instead determines the precise quasi-polychromatic optical depth from the vapour temperature and from this calculates the expected scattering intensity. We begin with a brief review of resonance radiation.

6.2 Resonance Radiation

It is well known that elements of matter have atomic transitions that correspond in energy to certain wavelengths of light (Mitchell and Zemansky, 1961). If an inbound photon has an energy that matches the energy of an atomic transition for a given element, then it will be absorbed and an electron will be raised to an excited state. After some lifetime, the electron will decay and a new photon will be emitted. Since an excited atom has no preferential alignment, the scattering is isotropic. This has the effect of creating absorption lines in broadband light that is passed through an atomic vapour. Light that interacts with the vapour is scattered isotropically, and so at that wavelength appears darker than comparable light that would have otherwise reached the observer. It is this isotropic nature that is leveraged by the BiSON resonance scattering spectrophotometers to move a photon away from the instrumental optic axis and into a detector.

In principle, any given quantum atomic resonance transition should produce a well defined absorption line that corresponds exactly to the energy of the transition. In reality, there are several natural broadening processes and we will look at these in turn now.

6.2.1 Natural Line Width

In a gas that forms an absorption line, the absorption coefficient κ_ν of the gas is defined in terms of,

$$I_\nu = I_0 e^{-\kappa_\nu x} \ , \tag{6.1}$$

where I_0 is the intensity of incident light, I_ν is the intensity of the transmitted light, and x is the thickness of the absorbing material. The units of the absorption coefficient are always expressed in terms of the reciprocal of the units of x such that the exponent

becomes a unitless ratio.

Mitchell and Zemansky (1961, p. 95) show that the integral of κ_ν over all frequencies is equal to,

$$\int \kappa_\nu \mathrm{d}\nu = \frac{\lambda_0^2}{8\pi}\frac{g_2}{g_1}\frac{N}{\tau}\left(1 - \frac{g_1}{g_2}\frac{N'}{N}\right) , \tag{6.2}$$

where λ_0 is the wavelength at the centre of the absorption line, τ the lifetime of the excited state, and g_1 and g_2 are the statistical weights of the normal and excited states respectively. N is the number density of atoms of which $[N - N']$ are capable of absorbing in the frequency range ν to $\nu + \mathrm{d}\nu$, and similarly N' is the number of excited atoms that are capable of emitting in this range. In the case where the ratio of excited atoms to normal atoms is very small ($\leqslant 10^{-4}$) then N' can be neglected and equation 6.2 reduces to,

$$\int \kappa_\nu \mathrm{d}\nu = \frac{\lambda_0^2}{8\pi}\frac{g_2}{g_1}\frac{N}{\tau} , \tag{6.3}$$

and this result shows that the the integral of the absorption coefficient is constant if N remains constant. This assumption can normally be made when the formation of excited atoms is due to absorption of a beam of light. Where a gas is electrically excited at high current densities, for example in a gas discharge lamp, the number of excited atoms may become a high fraction of the total number and so the simplification in equation 6.3 cannot be made. In the case where all atoms capable of absorbing at a certain frequency are already excited, the gas is said to be saturated and no further absorption can take place at that frequency.

Although the energy of an atomic transition is fixed, and therefore similarly the frequency of light with which it can interact, the absorption profile has a natural line width which arises from the uncertainty principle. When considered in the form,

$$\Delta E \Delta t > \frac{\hbar}{2} , \tag{6.4}$$

where ΔE is the energy, Δt the lifetime, and $\hbar$ the reduced Planck constant, the result suggests that for extremely short excited state lifetimes there will be a significant uncertainty in the energy of the photon emitted. If the energy of many photons are measured

the distribution formed is Lorentzian in shape, where Γ is the width parameter for a Lorentzian profile. If we assume the lifetime in the excited state is the uncertainty in time (i.e., $\tau = \Delta t$) then,

$$\Delta E = \frac{\Gamma}{2} = \frac{\hbar}{2\tau} \ , \tag{6.5}$$

and so,

$$\Gamma = \frac{\hbar}{\tau} \ , \tag{6.6}$$

which shows that the natural line width is inversely proportional to the lifetime of the atom in the excited state. Mishra (1950) observed the natural lifetime of the potassium $4^2P_{\frac{1}{2}} - 4^2S_{\frac{1}{2}}$ transition to be 27 ns. Using equation 6.6 this corresponds to a Lorentzian width of 2.44×10^{-8} eV. We can convert this to frequency using the relation $E = h\nu$, and so equation 6.5 becomes,

$$\begin{aligned} \Delta E &= \frac{h}{4\pi\tau} \ , \\ h\Delta\nu &= \frac{h}{4\pi\tau} \ , \\ \Delta\nu &= \frac{1}{4\pi\tau} \ , \end{aligned} \tag{6.7}$$

producing a width of 5.9 MHz. For the line centre of the $4^2P_{\frac{1}{2}} - 4^2S_{\frac{1}{2}}$ atomic transition at 769.898 nm this is a full width at half maximum (FWHM) of just 1.166×10^{-5} nm, which is equivalent to a Doppler velocity of $4.54\,\mathrm{m\,s^{-1}}$.

6.2.2 Pressure Broadening

Pressure acts to modify the natural line width. When an excited atom collides with another particle, it can be stimulated to decay earlier than would be the case for its natural lifetime. Continuing the argument from the uncertainty principle as before, if the lifetime is reduced than the uncertainty on the energy emitted must increase and hence broaden the line width. As the pressure increases, the likelihood of collisions increases.

There are three types of pressure broadening. If the excited atom collides with another of the same type, it is known as Holtsmark broadening. If the atom is of a different

"foreign" type in a mixed gas then it is known as Lorentz broadening. Lorentz broadening can also cause other effects such as a shift of the line centre and asymmetry of the line profile. Finally, collisions with electrons and ions is known as the Stark effect. Holtsmark broadening depends on the pressure of the absorbing gas, and Lorentz broadening on the pressure of the foreign gas. In general, these can be considered to be identical under the common name pressure broadening.

Experimental conditions are usually selected to ensure that pressure broadening is minimised, and this is the case for vapour reference cells which are typically placed under vacuum before being filled with their reference element. We will neglect this effect.

6.2.3 Doppler Broadening

Since the atoms in a gas are moving, the wavelength of light absorbed or emitted by the gas will be shifted according to the standard non-relativistic Doppler formula,

$$\frac{\Delta\lambda}{\lambda_0} = -\frac{\Delta\nu}{\nu_0} = \frac{v}{c}\,, \tag{6.8}$$

where $\Delta\lambda$ is the shifted wavelength, λ_0 is the rest wavelength, v the speed, and c the speed of light. Similarly, $\Delta\nu$ is the shift in frequency, and ν_0 is the frequency at rest. The atoms in a gas are moving at speeds defined by the Maxwell-Boltzmann distribution, and so a spectral line will be broadened into a range of possible wavelengths with a Gaussian distribution. The full width at half maximum (FWHM) due to Doppler broadening is given by,

$$\Delta\nu = \sqrt{\frac{8k_{\mathrm{B}}T\ln 2}{mc^2}}\nu_0 \tag{6.9}$$

which can be used to calculate the expected Doppler width of the absorption line. Potassium has an atomic weight of 39.0983 u. For a reference vapour at 100 °C the broadening is 0.0017 nm. This is equivalent to a Doppler velocity of $663\,\mathrm{m\,s^{-1}}$, almost 150 times larger than the width produced by natural broadening. On the Sun at a temperature of 5777 K, the D1 line width due to Doppler broadening is 0.0067 nm.

Mitchell and Zemansky (1961, p. 99) show that when considering only Doppler broadening, the absorption coefficient of a gas is given by,

$$\kappa_\nu = \kappa_0 e^{-\left[\frac{2(\nu-\nu_0)}{\Delta\nu}\sqrt{\ln 2}\right]^2} , \tag{6.10}$$

where κ_0 is the ideal maximum absorption for Doppler broadening alone, and $\Delta\nu$ is the Doppler breadth defined in equation 6.9. The integral of κ_ν over all frequencies is equal to,

$$\int \kappa_\nu \mathrm{d}\nu = \frac{1}{2}\sqrt{\frac{\pi}{\ln 2}}\kappa_0 \Delta\nu , \tag{6.11}$$

and this depends only on the temperature and the atomic mass of the atoms forming the vapour, which for a given material means temperature is the only variable.

6.2.4 Other Broadening

There are other effects that result in the broadening of spectral lines. A significant effect on the Sun is caused by solar rotation, meaning that light forming a line on one side of the disc will be red-shifted, and on the opposite side blue-shifted. The two combine to produce an overall broadening of the spectral line, an effect equivalent to a Doppler velocity of several thousand metres per second.

Another solar contribution is from velocity fields due to convection in the photosphere, known as microturbulence and macroturbulence (de Jager, 1954). The effect is similar to Doppler broadening, in that the distribution in line-of-sight velocity due to the movement of gas in convection cells causes a broadening effect. Rather than being derived from the movement of atoms, macroturbulence represents the velocity dispersion of granulation.

In the next section, we will use equation 6.3 and 6.11 to develop a model of the optical depth of potassium vapour, and then go on to model the response of a typical BiSON spectrophotometer.

6.3 Vapour Optical Depth

To be able to interpret the results of any experiments involving a gas reference cell, we need a mathematical expression for the absorption of the gas under the combined broadening conditions. We saw in equation 6.1 that for a gas that forms an absorption line the absorption coefficient κ_ν of the gas is defined by,

$$I_\nu = I_0 e^{-\kappa_\nu x} \ , \qquad \text{(6.1 revisited)}$$

and so it is necessary to determine the value of κ_ν. We know from equation 6.3 that the integral of κ_ν due to natural broadening over all frequencies is defined by,

$$\int \kappa_\nu \mathrm{d}\nu = \frac{\lambda_0^2}{8\pi} \frac{g_2}{g_1} \frac{N}{\tau} \ , \qquad \text{(6.3 revisited)}$$

and from equation 6.11 that the integral of κ_ν due to Doppler broadening over all frequencies is defined by,

$$\int \kappa_\nu \mathrm{d}\nu = \frac{1}{2}\sqrt{\frac{\pi}{\ln 2}} \kappa_0 \Delta\nu \ , \qquad \text{(6.11 revisited)}$$

where all terms are as defined previously. We know that no matter what physical processes are responsible for the formation of a spectral line, the integral of the absorption coefficient over all frequencies must remain constant if the number of atoms remains constant. This can be explained more simply by considering that any given atom can have the appropriate physical properties (i.e., speed, direction) to interact with only one particular frequency. If a spectral line is broadened, then the absorption depth must simultaneously reduce due to there being fewer atoms available to interact at the line centre. From this we can equate equations 6.3 and 6.11 thus,

$$\frac{\lambda_0^2}{8\pi} \frac{g_2}{g_1} \frac{N}{\tau} = \frac{1}{2}\sqrt{\frac{\pi}{\ln 2}} \kappa_0 \Delta\nu \ , \qquad (6.12)$$

and solving for κ_0 produces (Mitchell and Zemansky, 1961, p. 100),

$$k_0 = \frac{2}{\Delta\nu}\sqrt{\frac{\ln 2}{\pi}} \frac{\lambda_0^2}{8\pi} \frac{g_2}{g_1} \frac{N}{\tau} \ , \qquad (6.13)$$

which is the maximum absorption possible. This can be combined with the expected line profile to produce the final value for κ_ν. We have already seen such an example in equation 6.10 where for a line profile dominated by Doppler broadening κ_ν is defined by the distribution,

$$\kappa_\nu = \kappa_0 e^{-\left[\frac{2(\nu-\nu_0)}{\Delta\nu}\sqrt{\ln 2}\right]^2} , \qquad \text{(6.10 revisited)}$$

which, as stated, has a Gaussian profile. Where a spectral line is dominated by natural or pressure broadening a Lorentzian distribution may be a more appropriate line profile. If several absorption processes are balanced then a convolution of the two distributions may be required, known as a Voigt profile after Woldemar Voigt.

The absorption coefficient and thickness in equation 6.1 can be combined into one term τ_ν, known as the optical depth,

$$I_\nu = I_0 e^{-\tau_\nu} , \qquad (6.14)$$

where as before I_0 is the intensity of incident light, and I_ν is the intensity of the transmitted light. Unit optical depth is therefore defined as the optical thickness that reduces the photon flux by a factor of $1/e$. It should be noted that τ_ν is specifically a different quantity from τ the lifetime of the excited state.

We saw in sections 6.2.1 and 6.2.3 that the 769.898 nm line profile of a potassium vapour is easily dominated by Doppler broadening, and so equation 6.10 is valid to describe the distribution. In order to evaluate κ_ν we need to know the value of all the terms in equation 6.13 – temperature T, the atomic mass m, the frequency and wavelength of the line centre ν_0 and λ_0, the number density of absorbing atoms N, the lifetime of the excited state τ, and the degeneracy of the $4^2P_{\frac{1}{2}} - 4^2S_{\frac{1}{2}}$ energy levels g_1 and g_2, respectively.

We already know the lifetime τ to be 27 ns, the atomic mass to be 39.0983 u, the central wavelength to be 769.898 nm, and Mitchell and Zemansky (1961, p. 97) show the ratio g_2/g_1 to be unity. We can easily determine and control the temperature of the reference cell, and so the only remaining term is the number density. Unfortunately, it is almost impossible to know the exact amount of potassium in any given cell. The process

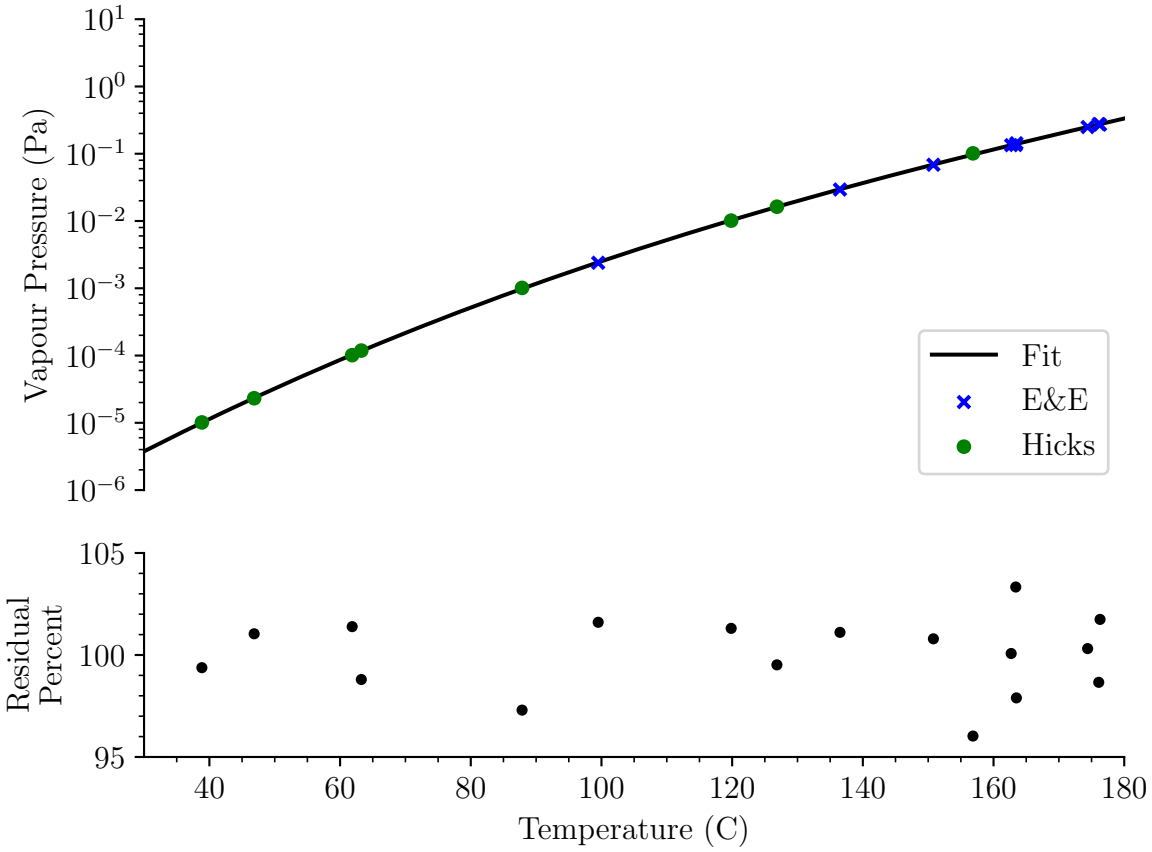

Figure 6.1: Upper: The blue crosses and green dots show measurements of the vapour pressure of potassium made independently by Edmondson and Egerton (1927) and Hicks (1963) at various temperatures. The black line shows a cubic fit to the data. Lower: The residuals to the fit in terms of percentage (i.e., $100 \cdot (\text{model}/\text{data})$). The uncertainty on the source data are not available, and so no error bars are plotted.

of moving potassium from an ampoule, through a manifold, and into a cell is very skilled but not precise (Jackson, 2003). However, we can determine the number density based on the vapour pressure, since as we saw in Section 3.2 it depends only on the temperature and not on the amount of material.

The vapour in a gas reference cell can be treated as an ideal gas since the pressure is by design very low. The cells are evacuated and baked out before being filled with a small amount of potassium. Since the pressure is low, there will be few interactions between particles. The ideal gas law is defined as,

$$pV = Nk_{\mathrm{B}}T \ , \tag{6.15}$$

where p is the pressure, V the volume, N the number of atoms or molecules, k_B the Boltzmann constant, and T the temperature. If we combine the number of atoms and the volume into one parameter and define the pressure in terms of temperature then equation 6.15 simplifies to,

$$n(T) = \frac{p(T)}{k_\mathrm{B} T} \ , \tag{6.16}$$

where $n(T)$ is now the number density which depends only on the vapour pressure and the temperature.

Both Edmondson and Egerton (1927) and later Hicks (1963) have independently made measurements of the vapour pressure of potassium at various temperatures. These data have been used to fit a polynomial in log space,

$$p_\mathrm{vapour} = e^{a\,T^3 + b\,T^2 + c\,T + d} \ , \tag{6.17}$$

where p_vapour is the vapour pressure and T is the temperature in Kelvin. The values of the four fitted coefficients are shown in Table 6.1. A plot of results from this model against the original data are shown in Figure 6.1. The model residuals are all within $\pm 5\,\%$ of the original value, with an upper limit to the fit of 449 K (175 °C). For a typical vapour temperature of 100 °C the number density required is $4.9 \times 10^{17}\,\mathrm{m}^{-3}$. This is equivalent to $3.2 \times 10^{-8}\,\mathrm{kg}$ of potassium – if you can see solid potassium in the cell then there is enough to form the required vapour.

We now have everything in place to be able to calculate the optical depth of a potassium vapour from equation 6.14. We know all of the parameters required to calculate the absorption coefficient κ_0 in equation 6.13 by determining the vapour pressure from

Table 6.1: Potassium vapour pressure model coefficients.

Coefficient	Value	Uncertainty
a	6.005×10^{-7}	8.747×10^{-8}
b	-8.912×10^{-4}	1.007×10^{-4}
c	4.890×10^{-1}	3.840×10^{-2}
d	-9.555×10^{1}	4.844

equation 6.17, and we can calculate the wavelength-dependent absorption coefficient from equation 6.10. In the next section we will verify the model by probing a potassium vapour cell using a Toptica DLCpro tunable diode laser.

6.4 Laser Calibration

The Toptica DLCpro tunable diode laser is a Littrow-style external cavity diode laser (ECDL, Toptica, 2017). In order to understand the operation and calibration of the laser we must first briefly discuss the design of an ECDL, starting with the difference between a laser diode and a diode laser.

A laser diode is the small, relatively inexpensive component that emits light. They are used in a variety of products such as optical disc-drives and barcode scanners. The emission spectrum of a laser diode is relatively broad and the lasing wavelength is not well defined. The laser cavity is a Fabry-Pérot resonator which is formed from two facets of a semiconductor crystal by cleaving two parallel edges. Light is emitted by the semiconductor material in the same manner as a typical light-emitting diode (LED), except in the case of the laser diode photons are reflected several times from each end face before they exit the diode. The light is amplified by stimulated emission as it passes through the cavity. Some light is lost due to absorption, but as the electrical drive current increases the amplification eventually exceeds the loss and the diode is said to "lase". A laser diode generally has a wide gain profile that supports many lasing modes simultaneously and so the emission spectrum is broadened by many different wavelengths. By comparison, a diode laser is the term used to define not just the light emitting element, but the control system as a whole – i.e., a diode laser system contains a laser diode.

The laser diode used here has an emission spectrum, and so a coarse tuning range, of 765 nm to 805 nm. Narrow line width is achieved by introducing frequency-selective feedback into the laser cavity using a grating mounted in front of the laser diode. A second laser cavity is formed externally between the grating and the semiconductor rear

facet, creating a grating-stabilised external cavity diode laser. The grating is in the Littrow configuration, which means that it is manufactured such that the diffraction angle and incidence angle are identical. Ensuring the back reflected beam is parallel to the incident beam allows the external cavity to be formed. The primary tuning mechanism is by altering the angle of incidence on the grating, which for coarse adjustment is done using a micrometer screw, and for fine adjustment by varying the applied voltage to a piezo actuator. Tuning can also be achieved by changing the diode drive current and temperature. The diode suffers mode-hopping as the tuning parameters are adjusted, but this can be mitigated by varying all parameters simultaneously. The typical mode-hop-free tuning-range (MHFTR) is quoted at 20 GHz to 50 GHz, which at a wavelength of 770 nm is equivalent to 40 pm to 98 pm. Achieving this range can be tricky due to the interaction of all the various tuning methods. The typical line width achieved by a Littrow-style ECDL over a duration of 5 µs is 10 kHz to 300 kHz, which again at a wavelength of 770 nm is less than 0.6 fm. The main contributions to the line width are electronic noise, acoustic noise, and other vibrations that affect the cavity length.

The final component of the diode laser system is an optical isolator intended to prevent light re-entering the laser and causing stability issues inside the cavity. The isolator used here produces a nonreciprocal propagation direction by making use of a Faraday-rotator between two linear polarisers aligned at 45° with respect to each other. Isolation in the reverse direction is due to the polarisers becoming crossed, and provides better than −38 dB rejection. The optical components are also angled at 1° to eliminate collinear back reflections.

The output of the ECDL has no intrinsic wavelength knowledge, except that the wavelength must be somewhere within the broadband output of the model of installed laser diode. In order to coarse-tune the laser to approximately 769.9 nm the output wavelength was measured using an Agilent 86142B optical spectrum analyser (Agilent, 2005, 2012), shown in Figure 6.2. The 86142B has a resolution bandwidth (RBW) of 0.06 nm and so the measured laser line width is considerably broadened. This resolution is sufficient to

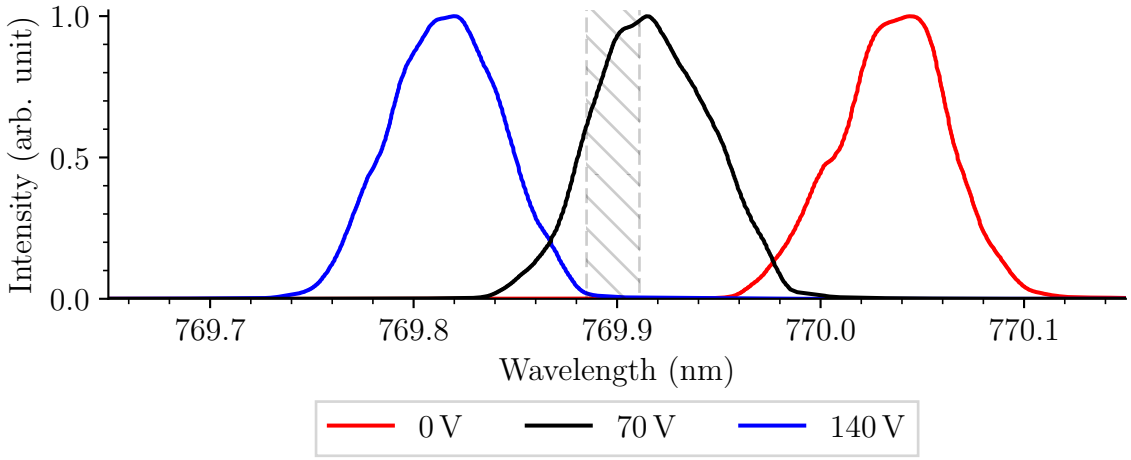

Figure 6.2: Toptica DLCpro tunable diode laser profile tuning range, measured using an Agilent 86142B optical spectrum analyser. The line width is much narrower than indicated in the figure. The measured FWHM is broadened to the minimum 0.06 nm resolution bandwidth of the spectrum analyser. The red and blue lines indicate the extent of the voltage-controlled piezo fine-tuning range after prior coarse tuning to approximately 769.9 nm. The hatched area indicates the extent of our wavelength region of interest.

allow use of the coarse manual micrometer screw adjustment to ensure the required wavelength is positioned within the narrow scan range of the piezo actuator. The resolution is insufficient to tune the wavelength precisely to the desired absorption line, but once coarse calibration is achieved simply scanning between the piezo limits allows the absorption line to be found. In order to calibrate the width and depth of the absorption line at varying temperature, it is necessary to calibrate the piezo voltage in terms of wavelength and to normalise the beam intensity. This calibration is non-trivial and we will now go on to discuss the steps required to achieve the final overall calibration.

We saw in Chapter 3 that BiSON spectrophotometers use a magnetic field to Zeeman-split the reference absorption line and produce two separate instrumental passbands on either side of the solar absorption line. This technique can be used to produce a known wavelength separation and so calibrate the piezo. Figure 6.3 shows the two components of the Zeeman-split absorption line with the cell heated to 90 °C in a 0.3 T longitudinal magnetic field. The beam intensity is not uniform across the piezo scan due to a drive-current feed-forward mechanism employed to increase the MHFTR of the diode laser. It

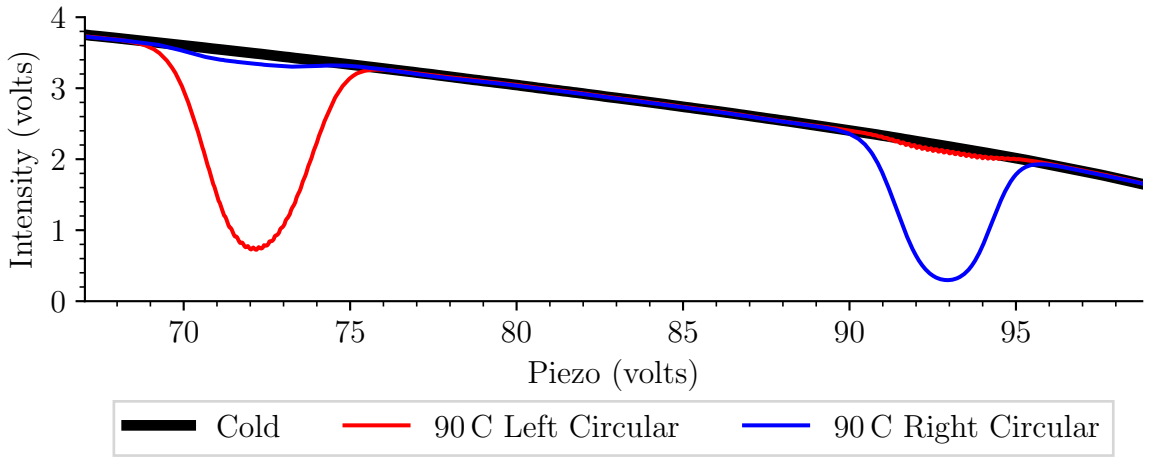

Figure 6.3: Laser intensity flat-field calibration.

is imperative that the diode does not mode-hop during the scan since this would result in different wavelength calibrations either side of the mode-hop, destroying the overall calibration. The feed-forward adjusts the drive-current proportionally with the piezo position-voltage by a factor of $-0.32\,\mathrm{mA\,V^{-1}}$, which helps prevent mode-hopping but has the undesired side-effect of changing the beam intensity. If a mode-hop continues to appear at an inconvenient wavelength, then the temperature of the diode can be adjusted to move the mode-hop elsewhere. The mean drive-current and detector-gain has to be selected carefully to ensure that the current always stays above the lasing threshold of 109 mA, but also that the beam intensity does not exceed the dynamic range of the detector throughout the whole piezo scan range. Once a mode-hop-free scan has been achieved with a beam intensity that is comfortably within the dynamic range of the detector, a flat-field can be captured and used to normalise the intensity of the scan, shown in Figure 6.4.

Zeeman splitting is an effect that arises from the interaction between an applied magnetic field $\vec{B}$ and the intrinsic magnetic moment of the atom $\vec{\mu}$. If the applied magnetic field is in the z-direction then magnetic potential energy ΔE_z is defined by,

$$\Delta E_z = g_j m_j \mu_B |\vec{B}| \ , \tag{6.18}$$

where m_j is the quantum number that describes the total angular momentum that is

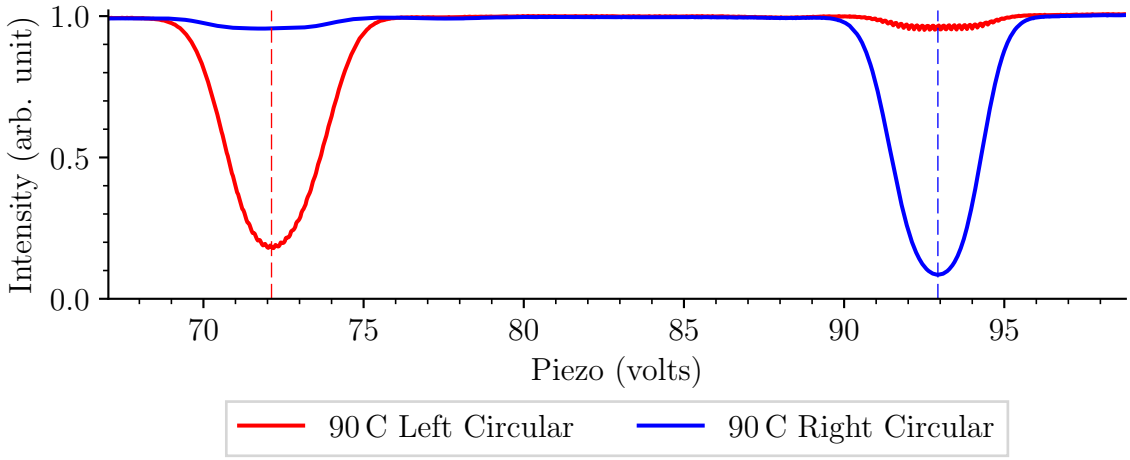

Figure 6.4: Piezo voltage calibration. The line centres, indicated by the dashed vertical lines, occur at 72.2 V and 92.9 V.

parallel to the axis of the applied field, μ_B is the *Bohr magneton* which is the magnetic dipole moment associated with unit orbital angular momentum, and g_j is the *Landé-g factor* which is dependent on the relative contributions of electron spin and orbital angular momentum.

For potassium the values of g_j are 2 for the ${}^2S_{\frac{1}{2}}$ state and $\frac{2}{3}$ for the ${}^2P_{\frac{1}{2}}$ state. There are two possible values for m_j of $\pm\frac{1}{2}$, and in the absence of an external magnetic field these states are degenerate. When a field is applied, both states are split into two further levels, giving a total of four possible transitions. Two of these transitions are forbidden if the photon is travelling parallel to the direction of the applied field. Using these values we can determine that,

$$\Delta E_z = \pm\frac{4\mu_B B}{3} \ , \tag{6.19}$$

and converted to frequency this energy becomes,

$$h\Delta\nu = \pm\frac{4\mu_B B}{3} \ , \tag{6.20}$$

which in terms of wavelength is,

$$\Delta\lambda = \pm\frac{4\mu_B B\lambda_0^2}{3hc} \ , \tag{6.21}$$

where the higher frequency, shorter wavelength transition is referred to as the *blue* transi-

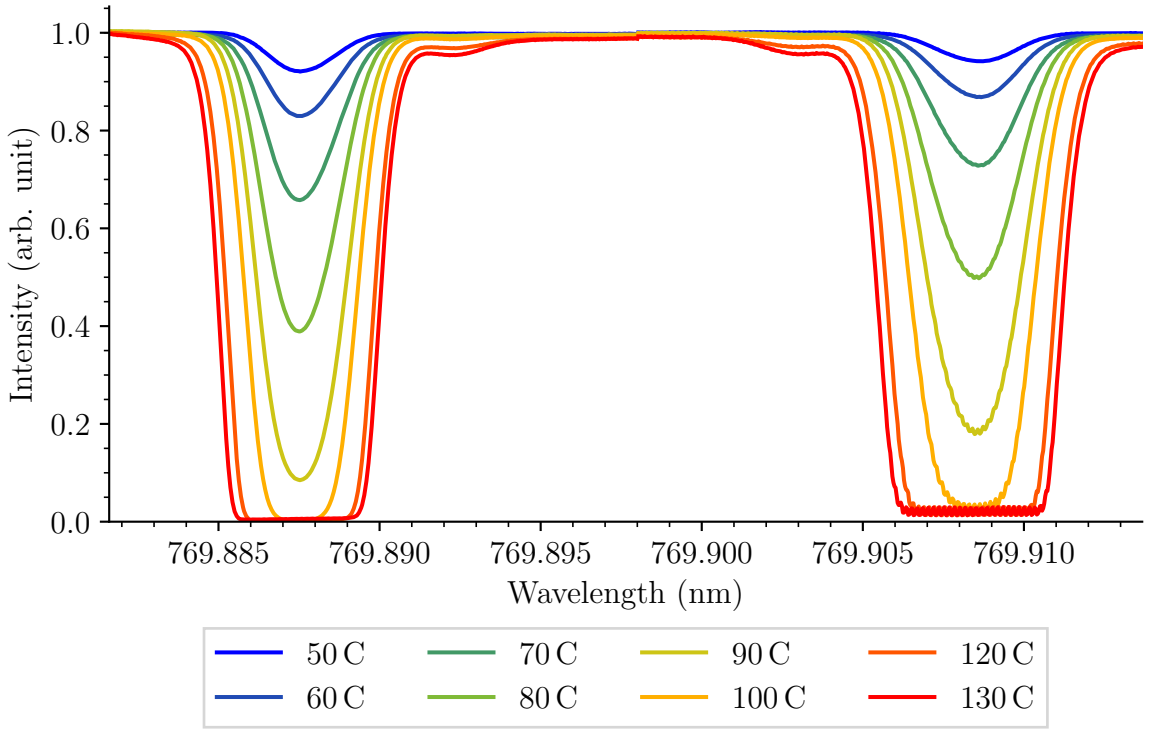

Figure 6.5: Vapour cell absorption intensity profiles at a range of temperatures.

tion, and the lower frequency longer wavelength as the *red* transition. In each transition, the change in atomic state Δm_j is unity and this means that the atom can interact only with circularly polarised light, with each being sensitive to one hand of polarisation. The sensitivity to opposite directions of circular polarisation is what allows the spectrophotometer to differentiate between the two passbands.

By making use of the result from equation 6.21 and applying the splitting to the piezo voltages at the measured line centres we can complete the calibration. The calibrated absorption profiles for eight different cell temperatures are shown in Figure 6.5. There is some evidence of slight misalignment of the polarising filters indicated by differences in absorption depth between the two polarisation states, and also some non-linearity in the wavelength scan indicated by a slight difference in width between the two polarisation states. These errors are small and will be mitigated against by fitting the model to both absorption components simultaneously, allowing all equal-and-opposite errors to cancel. In the next section, we will validate the model developed throughout Sections 6.2 and 6.3

against the absorption profiles measured here.

6.5 Model Validation

In Section 6.2, equation 6.1, we defined I_ν the frequency-dependent intensity of light transmitted through an atomic vapour,

$$I_\nu = I_0 e^{-\kappa_\nu x} \ , \qquad \text{(6.1 revisited)}$$

where I_0 is the intensity of incident light, x is the thickness of the absorbing material, and κ_ν is the frequency-dependent absorption coefficient. In equation 6.10 we defined κ_ν for a line profile dominated by Doppler broadening as,

$$\kappa_\nu = \kappa_0 e^{-\left[\frac{2(\nu-\nu_0)}{\Delta\nu}\sqrt{\ln 2}\right]^2} \ , \qquad \text{(6.10 revisited)}$$

and then in Section 6.3 equation 6.13 determined κ_0 as,

$$k_0 = \frac{2}{\Delta\nu}\sqrt{\frac{\ln 2}{\pi}}\frac{\lambda_0^2}{8\pi}\frac{g_2}{g_1}\frac{N}{\tau} \ , \qquad \text{(6.13 revisited)}$$

where N is the number density of absorbing atoms determined from the vapour pressure in equation 6.16 and all other terms are as defined previously.

The laser calibration developed in Section 6.4 is dependent on the accuracy and uniformity of the vapour cell longitudinal magnetic field. The field strength at several points inside the cell oven was measured using a Hirst GM04 Gaussmeter and the field profile is shown in Figure 6.6. It is clear that the target field strength of 0.3 T is only achieved at the cell centre and drops to approximately 0.26 T at the front and rear of the cell oven. This means that the absorption profile will be broadened by the non-uniform magnetic field, and the effect can be seen in the left panel of Figure 6.7. If the laser piezo is calibrated based on a mean field strength of 0.283 T then it is clear that the measured absorption profile is broadened by the modelled components at the lower and higher field strengths. If the FWHM determined from equation 6.9 is modified to allow an additional degree

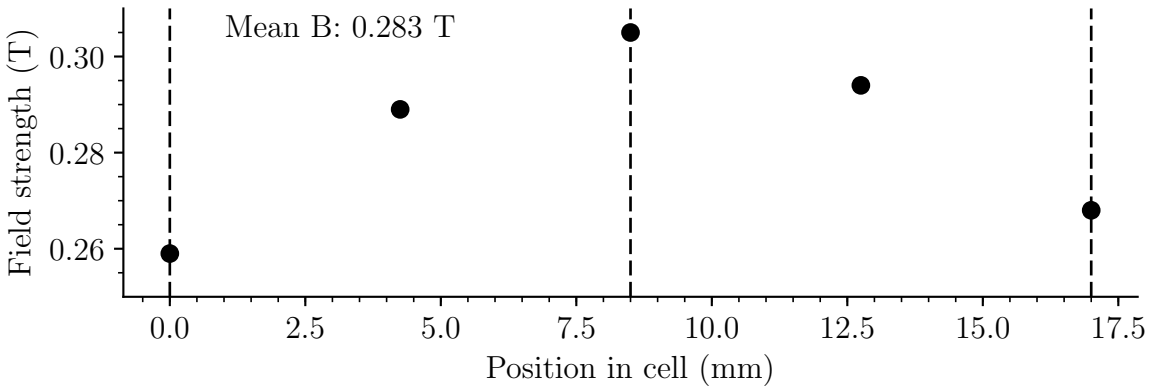

Figure 6.6: On-axis magnetic variation inside the cell oven. The field inhomogeneity adds an additional source of spectral broadening that must be included in the model.

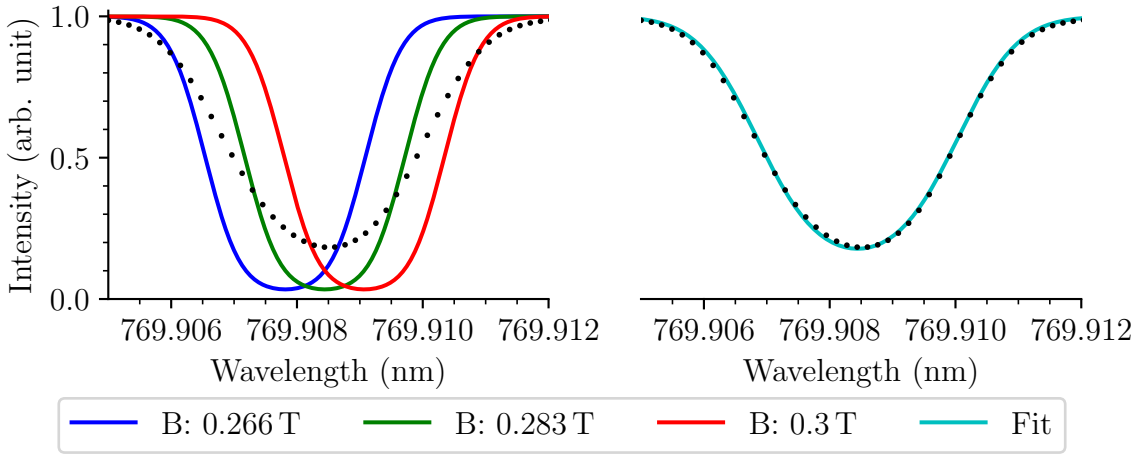

Figure 6.7: Left: Spectral broadening due to magnetic inhomogeneity. Right: Improved fitting with broadening scaling factor. The vapour is modelled at a typical working temperature of 90 °C.

of freedom to the fit by adding a broadening scaling factor to compensate for the inhomogeneity then the fit is significantly improved, shown in the right panel of Figure 6.7. The model has been fitted against absorption profiles for all eight vapour temperatures in Figure 6.5. Both absorption components are fitted simultaneously allowing systematic equal-and-opposite errors to cancel. The results are shown in Figure 6.8, and the fitted parameters and uncertainties listed in Table 6.2.

All fitted temperatures indicate that the vapour pressure is controlled by the temperature of the solid potassium in the stem of the cell. The temperature of the upper cube of the cell simply has to be higher than the stem and high enough to ensure that deposition

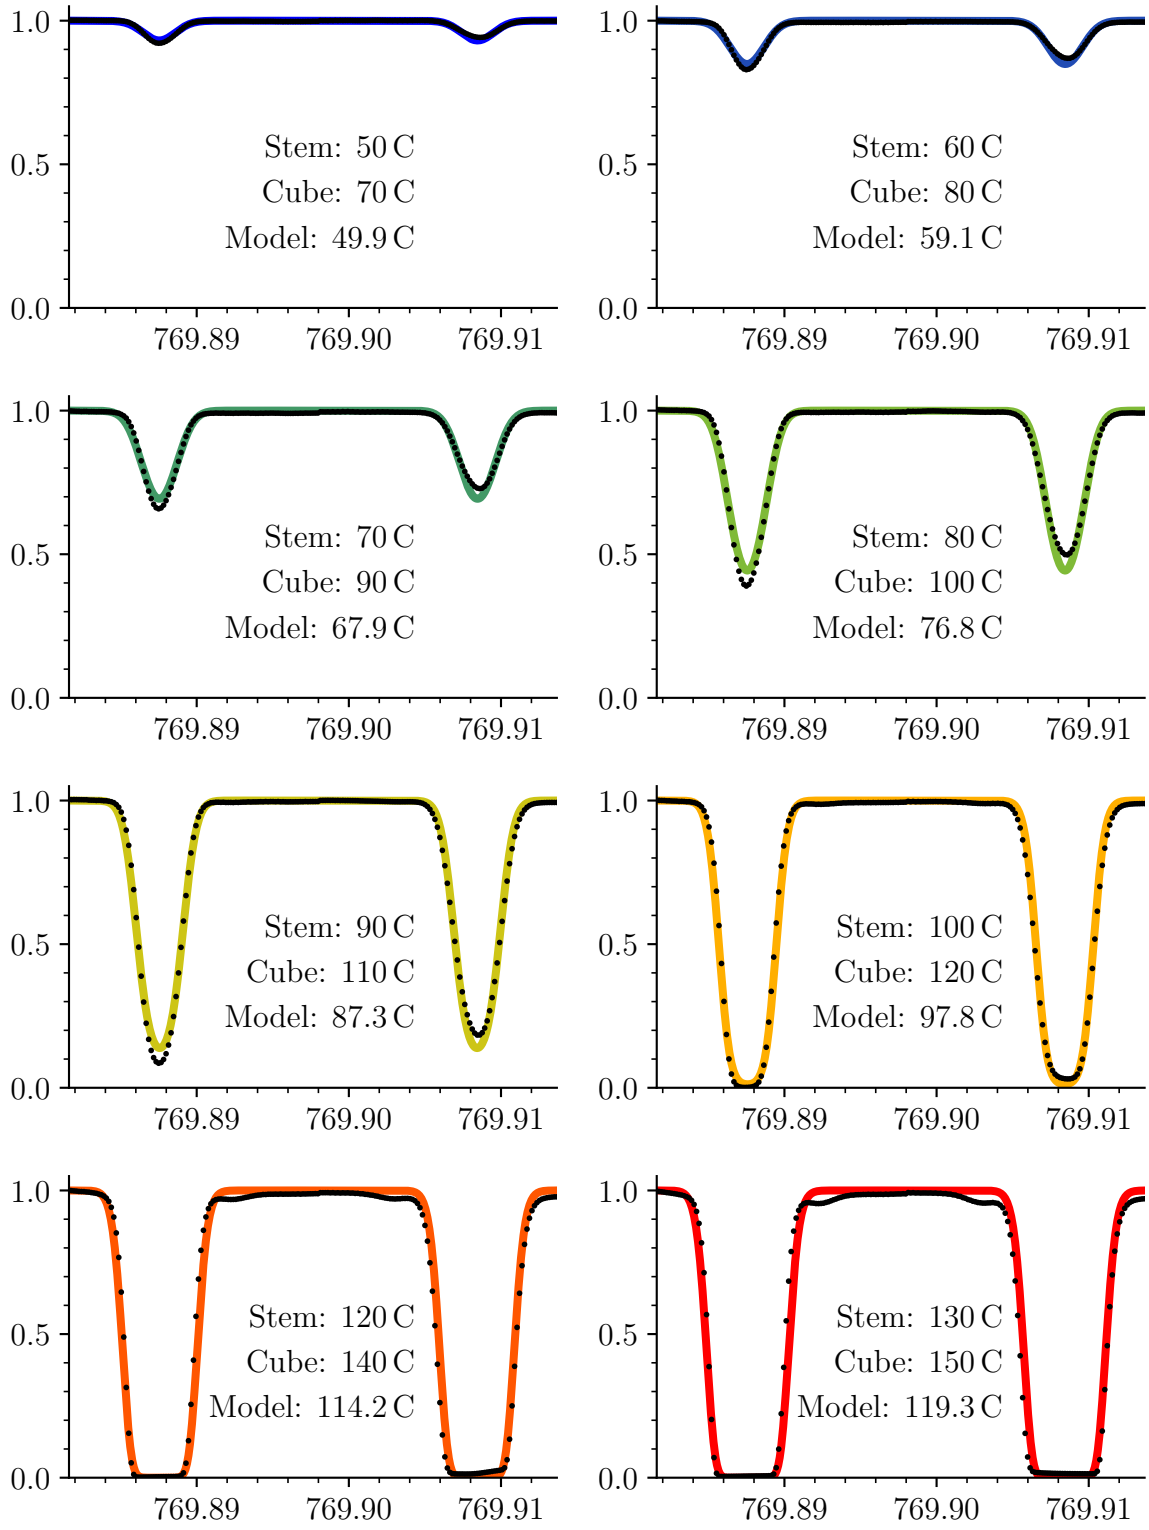

Figure 6.8: Modelled absorption profiles at a range of vapour temperatures. Both absorption components are fitted simultaneously allowing systematic equal-and-opposite errors to cancel.

Table 6.2: Fitted model parameters at a range of vapour temperatures.

Oven Temp		**Model Temp**			**Broadening Factor**	
Stem	Cube	Stem	Uncertainty	ΔT	Scale	Uncertainty
50	70	49.9	0.1	−0.1	1.55	0.01
60	80	59.1	0.1	−0.9	1.59	0.01
70	90	67.9	0.1	−2.1	1.55	0.01
80	100	76.8	<0.1	−3.2	1.48	0.01
90	110	87.3	0.1	−2.7	1.40	<0.01
100	120	97.8	0.1	−2.2	1.36	0.01
120	140	114.2	0.6	−5.8	1.41	0.01
130	150	119.3	0.8	−10.7	1.47	0.02

of solid potassium onto the windows does not occur. The increasing difference between the model temperature and the stem temperature at higher values is an indication of the difficulties of coupling heat efficiently into the cell. Heating of the stem is achieved by a foil heating element wrapped around the glass, and heat does of course leak into the surrounding chassis as well as heating the cell. At higher temperatures the losses become more significant and it becomes progressively more difficult to efficiently heat the cell, and this is indicated by the real temperature of the vapour as determined from the model being a few degrees below measured temperature of the stem heating element.

The model can also be used to estimate the scattering performance of the cell and the effect of different detector aperture sizes and positions. Figure 6.9 shows the expected intensity change as light passes through the cell from left to right determined from equation 6.1. The absorption coefficient is modelled as monochromatic from equation 6.13 but with the initial intensity scaled by the expected FWHM to compensate for the differences in spectral bandwidth. We make an assumption that any photon is scattered only once. The white circle in the centre indicates the typical 6 mm aperture of a scattering detector, with scattered light perpendicular to the page coming out towards the reader as the beam passes through from left to right. It is clear that at low temperatures there is very little scattered light due to the small absorption coefficient. At high temperatures almost all the light is scattered at the front of the cell and so not captured by the detector. This

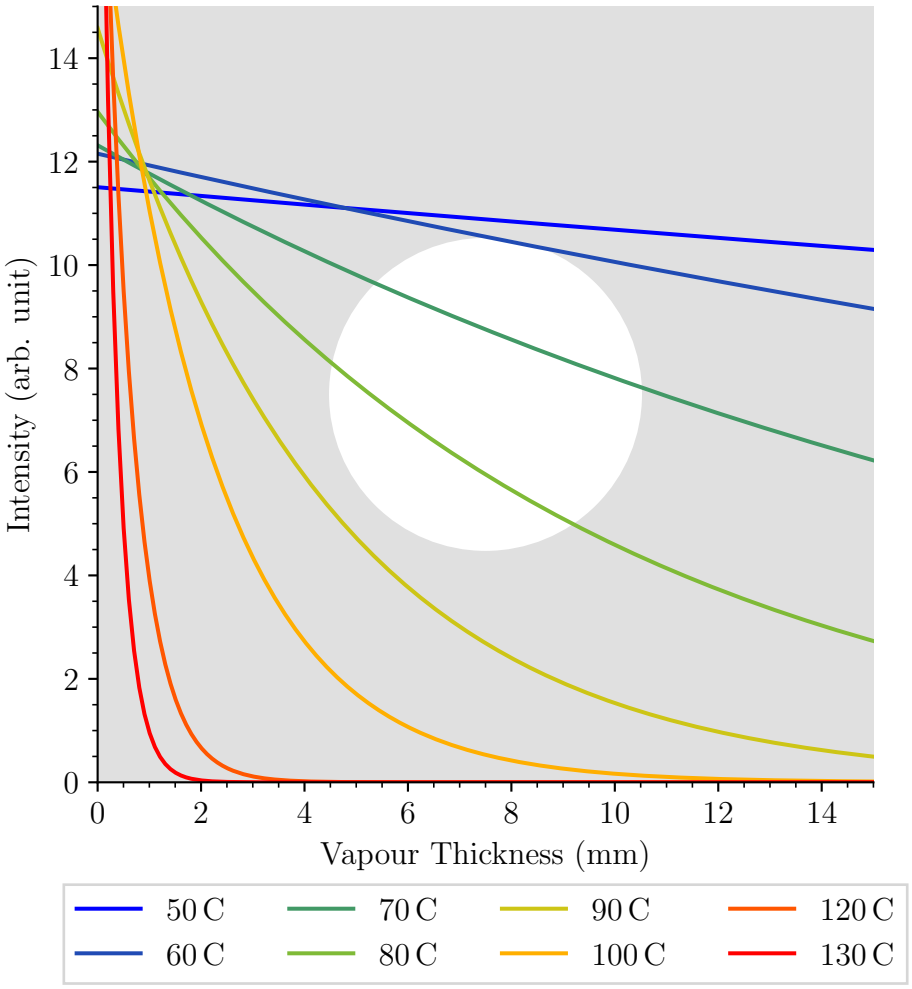

Figure 6.9: Absorption through a typical 15 mm BiSON vapour cell at various temperatures. The absorption coefficient is modelled as monochromatic but with the initial intensity scaled by the expected FWHM to compensate for the differences in spectral bandwidth. The white circle in the centre indicates the typical 6 mm aperture of a scattering detector.

implies that there will be some optimised temperature at which the measured scattering intensity is greatest. Figure 6.10 shows the estimated scattering intensity for several aperture sizes and vapour temperatures. If the aperture is equal to the size of the vapour cell then all the light scattered towards the detector is captured, and the intensity continually increases with temperature. As the aperture size is reduced a temperature plateau forms. One might expect that the ideal aperture is as large as possible in order to maximise the intensity of scattered light and so maximise the signal-to-noise ratio. However, there are several reasons why this is not the case. Applying a smaller aperture and causing the

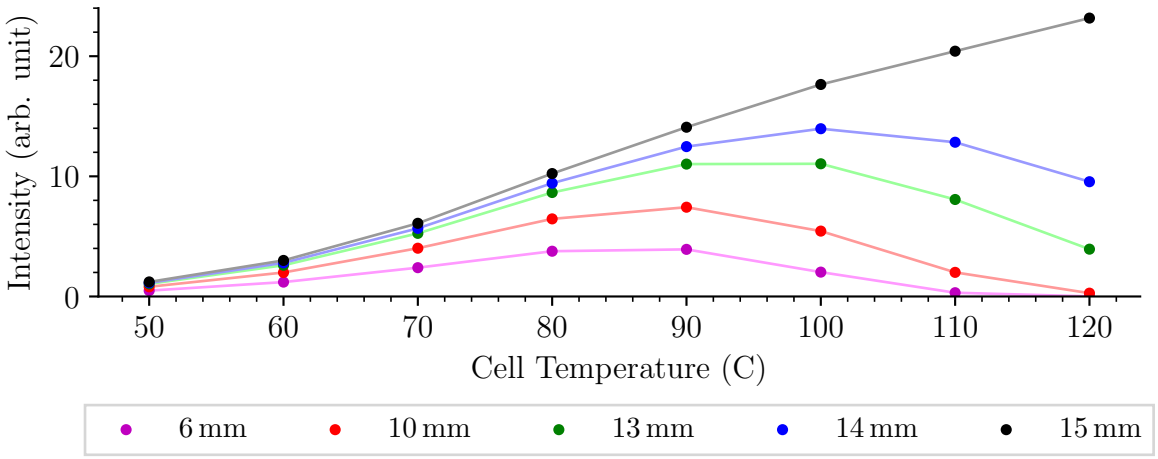

Figure 6.10: Optimum vapour temperature for maximum scattering intensity at a range of aperture sizes.

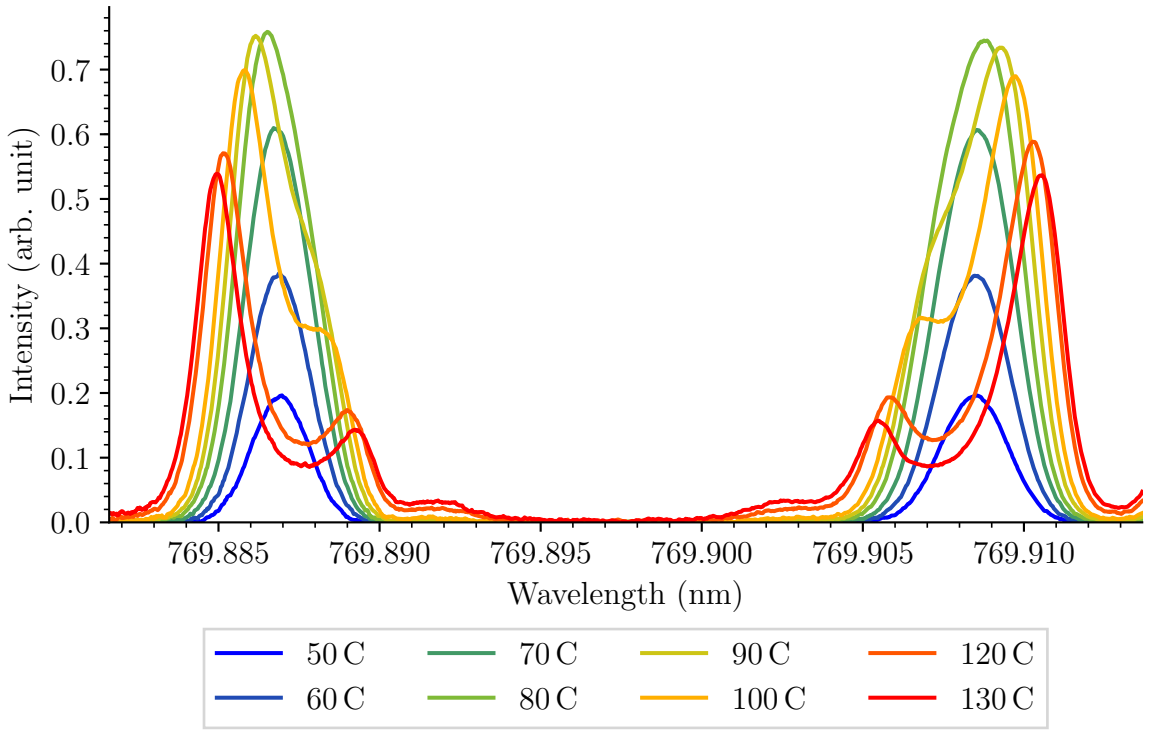

Figure 6.11: Vapour cell scattering intensity profiles at a range of temperatures. The profile asymmetry seen at temperatures above the ideal temperature is a result of the inhomogeneity of the magnetic field.

temperature plateau to form minimises the effect of fluctuations in vapour temperature, and the temperature-width of the plateau is greatest at smaller apertures. Also, it is imperative to eliminate non-resonantly scattered light from reaching the detector, such as specular reflections from the cell glass especially at the edges and corners of the cell, and again this is best achieved via small apertures effectively acting as a baffle.

The scattering detectors used here employ a 6 mm circular aperture which suggests that the maximum intensity will be found at approximately 80 °C to 90 °C. Figure 6.11 shows the scattering profiles taken simultaneously with the absorption profiles shown in Figure 6.5 over a range of vapour cell temperatures, and the intensity peaks at approximately 80 °C to 90 °C as expected. The profile asymmetry seen at temperatures above the ideal temperature is a result of the inhomogeneity of the magnetic field. At high temperatures the majority of the light is scattered at the front of the cell within the lower-magnitude magnetic field and not captured by the detector. This results in the light scattered within the range of the aperture being artificially weighted towards the higher-magnitude field at the centre of the cell, and so the scattering intensity is weighted towards the larger-separation of the stronger field.

6.6 Conclusion

In this chapter we developed a model of a BiSON potassium vapour cell and validated the output of the model using a tunable diode laser. The model was found to effectively reproduce the performance of a vapour cell and can be used to compare expected performance between existing BiSON vapour cells and commercially available alternatives.

There are a number of conflicting parameters that combine to produce the final cell characteristics. The cell temperature needs to be as high as possible to maximise the scattering cross-section, but this results in a high optical depth causing all the light to be scattered very close to the front of the cell and not reach the detector. Also, a high optical depth increases the probability of multiple scattering, meaning that light that

is initially scattered towards a detector may moved away from the detector axis again. However, if the temperature is too low then the small absorption coefficient results in too little light being scattered towards a detector. Likewise the size and position of the aperture is critical to both eliminating non-resonantly scattered light, and also causing a temperature plateau to form minimising the effect of temperature fluctuations.

At the ideal temperature of approximately 88 °C for a scattering aperture of 6 mm diameter, the optical depth is approximately 1 if measured from the front of the cell to the centre, or approximately 2 from the front of the cell to the rear. If the magnetic field were uniform, rather than the magnetic inhomogeneity causing a mean magnetic broadening of 1.5 times the Doppler-dominated FWHM, then the ideal temperature would reduce by approximately 5 °C to 83 °C. Many BiSON instruments run at a stem temperature of 90 °C to 100 °C and a cube temperature of 110 °C to 120 °C and so these results agree with expectation based on experience.

In the next chapter we will go on to extend this model into a full end-to-end simulation from the solar constant through to the output of the detectors, taking into account changes in étendue and spectral bandwidth at critical points, and assessing all noise sources in isolation.

Chapter 7

A Miniature Spectrophotometer

Contents

7.1 Introduction

The first task of any telescope, be it solar or stellar, is the collection and concentration of light. A solar telescope has the distinct advantage of having a much higher photon flux available for collection, and so it can be made much smaller than its stellar counterpart. The optical design principles required are the same for either type of telescope. One of the key principles to miniaturising a BiSON node is splitting the collection optics from the spectrophotometer. Once this is done by making use of fibre optics, a much smaller optical mount can be used and with it a much smaller housing for the whole system.

In this chapter we continue the model from Chapter 6 and build around it the framework of components required for it to become an operational solar spectrophotometer. Our design specification requires the spectrophotometer to be small and inexpensive. To

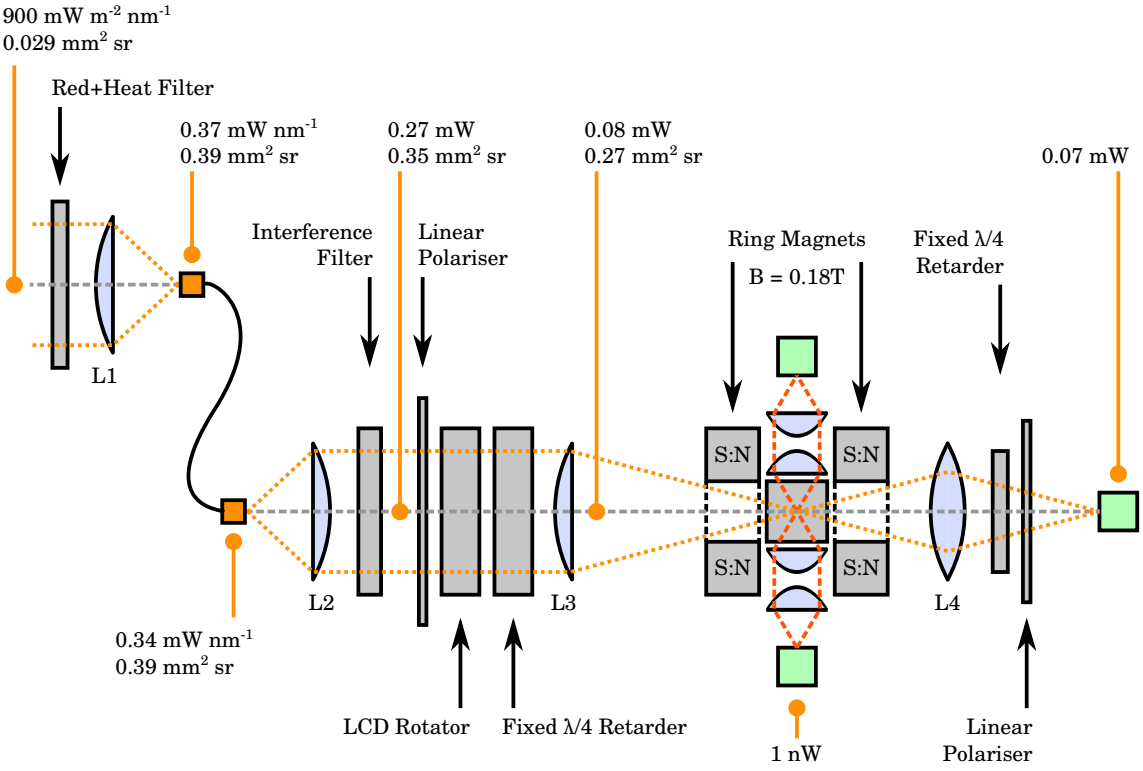

Figure 7.1: Schematic of a fibre-based BiSON resonant scattering spectrophotometer. The étendue and optical power (or power spectral density, as appropriate) is indicated at key points in the system by the orange markers. The green boxes indicate detector locations.

achieve this we need to eliminate as much hardware as possible and make the design as simple as it can be. To achieve the system shown schematically in Figure 7.1 it will be necessary to use some optically-fast lenses and we will not be able to make use of first-order approximations in our optical calculations. We must instead consider the effects of all aberrations. Some background theory supporting the work in this chapter is available in Appendix A.

We begin the end-to-end system model in Section 7.2 by considering the total output power of the Sun and use the results from Chapter 5 to determine the power available to the instrument. In Section 7.3 we develop the optics required to couple sunlight into a fibre and in Section 7.4 consider the filters required for safe operation. In Section 7.5 we measure the polarisation switching rate of a Pockels-effect variable retarder and compare with a liquid-crystal variable retarder. In Section 7.6 we make use of the results from Chapter 6 to assess the expected performance of an off-the-shelf potassium vapour cell and

oven with comparison to a standard BiSON vapour cell. In Section 7.7 we model the solar potassium absorption line and determine the optimum magnetic field strength that maximises instrumental sensitivity to changes in line-of-sight velocity. Finally in Section 7.8 we consider control of the light that is transmitted directly through the instrument and not scattered by the vapour cell.

7.2 Power Budget – The Solar Constant

The Sun radiates an enormous amount of energy. As we saw in Chapter 2, light from the Sun is broadly like that of a black body at a temperature of 6000 K, or more precisely 5777 K. The emission spectrum of a black body in thermal equilibrium is described by Planck's law,

$$B(\lambda, T)\,\mathrm{d}\lambda = \frac{2hc^2}{\lambda^5}\frac{1}{e^{\frac{hc}{\lambda k_\mathrm{B} T}} - 1}\,\mathrm{d}\lambda\ , \qquad \text{(2.1 revisited)}$$

where λ is the wavelength, T the temperature, k_B is the Boltzmann constant, h the Planck constant, and c the speed of light. The SI unit of spectral radiance, B, is Watts per steradian per square metre per nanometre ($\mathrm{W\,sr^{-1}\,m^{-2}\,nm^{-1}}$). We require the radiant exitance, which is the total radiated power per unit area. The radiant exitance can be determined by integrating equation 2.1 over all wavelengths, and then again with respect to the solid angle over the hemisphere into which the surface is radiating. After integration, the total radiant exitance for a black body such as the Sun is approximately $63\,\mathrm{MW\,m^{-2}}$. At Earth this reduces by inverse square law to approximately $1365\,\mathrm{W\,m^{-2}}$ which is known as the solar constant, measured at the top of Earth's atmosphere, and assumes radiation incident on a surface perpendicular to the rays.

To calculate the power reaching the ground we must take into account attenuation by Earth's atmosphere, an effect we saw in Chapter 5 known as atmospheric extinction. Figure 7.2 shows the American Society for Testing and Materials (ASTM) G173-03 Reference Solar Spectral Irradiances (ASTM, 2012; NREL, 2012). The orange curve shows the Extraterrestrial Radiation (ETR) solar spectrum at the mean Sun-Earth distance, at

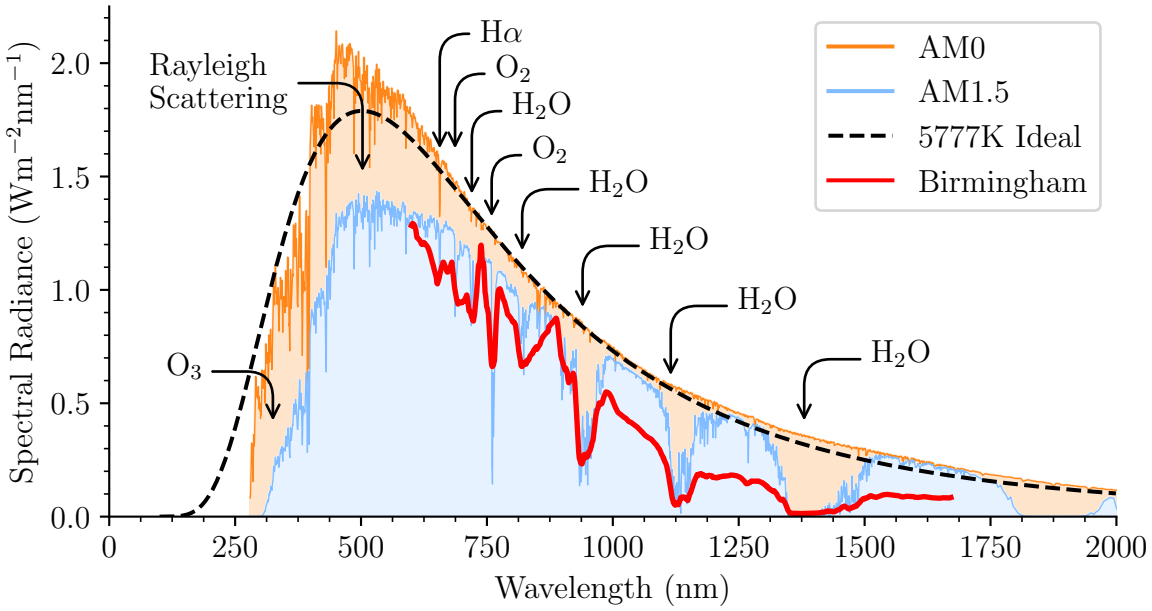

Figure 7.2: Reference Solar Spectral Radiance. AM0 means zero airmass, i.e., at the top of the atmosphere. AM1.5 is 1.5 air masses. The dashed black line represents a black body spectrum at 5777 K, and the solid red line the profile measured from Birmingham after scaling to the expected value. The arrows indicate regions of significant molecular absorption.

the top of the atmosphere (airmass zero, AM0). The blue curve shows Direct Normal Irradiance (DNI) through 1.5 airmasses (AM1.5), on a surface with the surface normal tracking the Sun, and where the atmospheric conditions are modelled according to the 1976 United States Standard Atmosphere (NOAA, 1976). The temperature, pressure, aerosol density, air density, and molecular species density are specified in 33 layers. The defaults for several parameters as specified by NOAA are an aerosol optical depth (AOD) of 0.084 measured at 500 nm, with a total column water vapour equivalent of 1.42 cm, and a total column ozone equivalent of 0.34 cm. For comparison, the black dashed line in Figure 7.2 shows a 5777 K black body spectrum, and the red line shows the solar profile between 600 nm to 1650 nm measured using an Agilent 86142B optical spectrum analyser at Birmingham through unspecified atmospheric conditions, and scaled to the expected value. After scattering and absorption within Earth's atmosphere, the solar constant at mean sea level is reduced to approximately 1000 W m^{-2} and this varies greatly with atmospheric conditions. The solar constant corresponds to an equivalent radiant exitance of a

black body at 364.4 K, just below the boiling point of water. This is a convenient value for domestic solar hot water systems, where no concentration is required and a simple flat panel absorber can be used to provide hot water.

Near our target wavelength of 770 nm, the power spectral density is 56.7 kW m^{-2} nm^{-1}. At the top of the Earth's atmosphere, the power spectral density is reduced to just 1.22 W m^{-2} nm^{-1}. We saw in Chapter 5 that atmospheric extinction can vary between 0.03 at Las Campanas, and anything up to 0.5 during dusty periods at Izaña. At one airmass with an aerosol optical depth of 0.03, the power spectral density at sea level near 770 nm is 1.18 W m^{-2} nm^{-1}. If we estimate six airmasses with an aerosol optical depth of 0.1, this reduces the power spectral density at sea level to 0.67 W m^{-2} nm^{-1}. We will need to consider both these limits when specifying the dynamic range of the photodetectors. For the remainder of this chapter, we will assume a typical optical power spectral density budget of 0.9 W m^{-2} nm^{-1} near 770 nm, and a minimum acceptance angle of 0.27° based on the maximum angular diameter of the Sun, a field angle of 0.54°. These values will be required when determining the necessary aperture size, numerical aperture, and bandwidth of the system.

The first stage of instrumentation is to collect sufficient light to meet the constraints of the detection and acquisition electronics. In the next section we will specify a custom solar concentrator that meets the design requirements we have determined here, and make use of commercial off-the-shelf (COTS) optical components and fibres.

7.3 Fibre Coupling

7.3.1 Étendue and Acceptance Angle

An optical fibre has an acceptance angle governed by the critical angle of total internal reflection between the core and cladding of the fibre, shown in Figure 7.3. The angles of incidence and refraction when a ray of light passes through a boundary between two different materials are described by Snell's Law. A ray of light will propagate along the

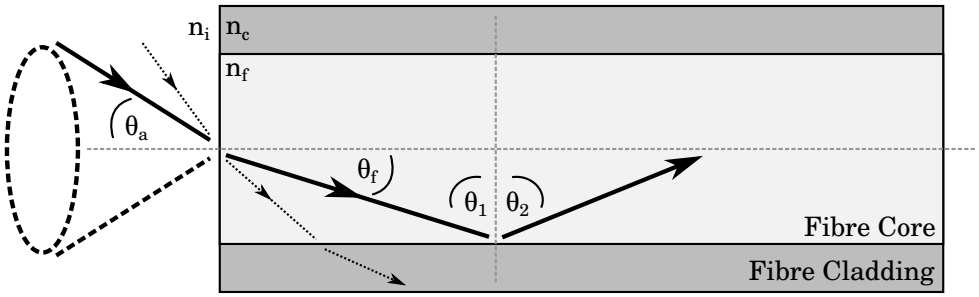

Figure 7.3: An optical fibre has an acceptance angle governed by the critical angle of total internal reflection between the core and cladding of the fibre. The large solid arrows show a bound ray entering at the acceptance angle and propagating along the fibre. The smaller dashed arrows show an unbound ray entering at greater than the acceptance angle and being lost into the cladding.

fibre only if it experiences total internal reflection at the boundary between the fibre and the cladding. This can only occur where θ_2 has an absolute maximum value of 90°, and the ray runs along parallel with the fibre-cladding boundary. We can calculate the maximum value of θ_1, and from there determine the acceptance angle θ_a. Solving Snell's Law,

$$\frac{\sin\theta_1}{\sin\theta_2} = \frac{n_2}{n_1} \ , \tag{7.1}$$

for θ_1 and setting $\theta_2 = 90°$ we get,

$$\theta_\mathrm{c} = \theta_1 = \sin^{-1}\left(\frac{n_\mathrm{c}}{n_\mathrm{f}}\right) \ , \tag{7.2}$$

where n_f is the refractive index of the fibre core, and n_c is the refractive index of the cladding. We can now use equation 7.1 again to find the acceptance angle, θ_a,

$$\begin{aligned} n_\mathrm{i}\sin\theta_\mathrm{a} &= n_\mathrm{f}\sin\left(90-\theta_\mathrm{c}\right) \\ &= n_\mathrm{f}\cos\theta_\mathrm{c} \\ &= n_\mathrm{f}\cos\left(\sin^{-1}\left(\frac{n_\mathrm{c}}{n_\mathrm{f}}\right)\right) \ . \end{aligned} \tag{7.3}$$

Using the trigonometric identity,

$$\cos(\sin^{-1}x) = \sqrt{1-x^2} \ , \tag{7.4}$$

we get,

$$\begin{aligned} \mathrm{NA} = n_\mathrm{i} \sin\theta_\mathrm{a} &= n_\mathrm{f}\sqrt{1 - \frac{n_\mathrm{c}^2}{n_\mathrm{f}^2}} \ , \\ &= \sqrt{n_\mathrm{f}^2 - n_\mathrm{c}^2} \ , \end{aligned} \tag{7.5}$$

which is the numerical aperture (NA) of the fibre. For a given fibre core diameter, the NA is considered to be the light gathering capacity of an optical fibre. The higher the value of NA, the more light is collected for propagation by the fibre. An FT1000EMT multimode fibre from Thorlabs was selected for use here as a standard off-the-shelf fibre with a silica core diameter of 1000 µm, a low hydroxyl content for low attenuation at 770 nm of 7 dB km^{-1}, and an NA of 0.39 producing an acceptance angle of approximately 22.9°. The large core diameter was selected in order to both increase the light throughput, and maximise the focal ratio of the collection optics by allowing for longer focal lengths which naturally produce larger images. A larger, 1500 µm diameter, fibre is available off-the-shelf, however the minimum bend radius becomes very large at these diameters. The minimum short term bend radius of the 1000 µm is 50 mm, the long term limit is 100 mm, and these were considered to already be at the limits of what might be required during the operation of the system. Additionally, a smaller core diameter is beneficial at the output of the fibre in order to reduce the divergence of the resulting collimated beam, and so restrict the off-axis angles through subsequent optics. The core size is a compromise between these requirements. Ideally we would like a large fibre with high NA at the input, and a small fibre with low NA at the output, but clearly this is not possible. The core diameter and NA of the fibre applies a constraint on the maximum aperture of the objective lens we can use to couple sunlight into the fibre. A simple way of assessing the maximum aperture is to consider conservation of étendue.

Étendue, or throughput, is a measure of the size and angular spread of a beam as it passes through a system. Chaves (2008, p. 65) introduces the concepts of "spatial room" and "angular room", and defines étendue as a geometrical quantity formed from the product of area and solid angle. If the spatial room is decreased (smaller aperture)

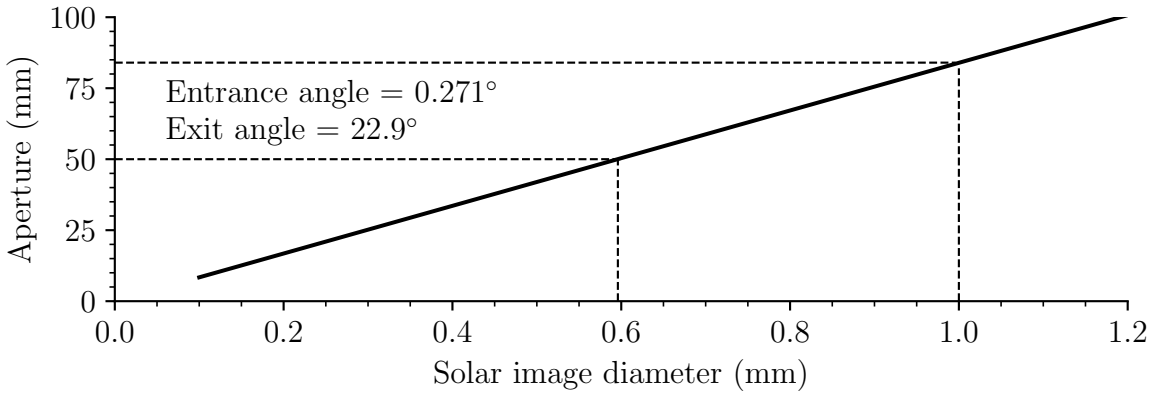

Figure 7.4: Maximum objective aperture based on conservation of étendue. If we require a 1 mm solar image diameter then the maximum objective aperture diameter is 84 mm. Similarly, for a 50 mm objective aperture diameter the minimum solar image diameter we can expect is just under 0.6 mm.

then the angular room (divergence) must increase. Étendue is the optical equivalent of entropy defined in the second law of thermodynamics and similarly étendue can never decrease, it can only be conserved or increased by, for example, light passing through a diffuser. Once diffused, the étendue cannot be reduced again since an "undiffuser" does not exist. Étendue can, however, be reduced by losing light through issues such a vignetting causing a loss in efficiency. If we compare the diameter and acceptance angle of the fibre with the required solar field angle, we can determine the maximum objective aperture that allows étendue to be conserved and no light lost.

The solid angle of a cone whose cross-section subtends the angle 2θ, shown schematically at the entrance to a fibre in Figure 7.3 on page 144, is given by,

$$\Omega = 2\,\pi\,[1 - \cos(\theta)]\ , \tag{7.6}$$

where Ω is the solid angle in steradians. For a 1 mm diameter fibre with an NA of 0.39 the étendue is equal to $3.9 \times 10^{-7}\,\mathrm{m^2\,sr}$. The solar field angle is 0.27°, and by applying conservation of étendue the maximum objective aperture diameter is 84 mm to match the fibre. The maximum aperture diameter for a range of fibre diameters is shown in Figure 7.4. We can calculate a first-order estimate of the required focal length using the

field angle equation,

$$\theta_\alpha = \tan^{-1} \frac{h_{\text{im}}}{f} \ , \tag{7.7}$$

where θ_α is the solar field angle and h_{im} is the radius of the fibre. Solving for the focal length f gives a value of 105 mm. BiSON instruments typically use a front aperture of around 30 mm, and so using the full étendue of the fibre would require a focal ratio of $f/3.5$. As with all telescopes, more light is better and so it would be useful to be able to increase this aperture. It will also be necessary to increase the acceptance angle in order to minimise the effect of system imperfections, such as optical aberrations and guiding errors, meaning the final system will likely be much faster than this.

7.3.2 A Solar Concentrator

Since we require fast optics, we cannot make use of first-order approximations and instead need to consider all aberrations. Whilst the required calculations are not particularly difficult they are laborious and time consuming, and so several commercial optical design packages have been developed for computational ray tracing. Here, we make use of OSLO-EDU, a version of OSLO TracePro which is free for educational use.

Several options have been designed and trialled for the fibre-feed collection system, shown schematically as lens L1 in Figure 7.1 on page 140. All make use of cemented achromatic doublets, which are designed to reduce as far as possible chromatic aberration. Since our system is quasi-monochromatic we do not need good broadband performance, but achromat lenses are also well optimised to correct for spherical aberration which we do require with fast focal ratios. A doublet lens corrects for aberrations by making use of two spherical lenses manufactured from different glass types. The curvature of the cemented interface between the two lenses is optimised to allow the aberrations of the first spherical lens to be corrected by equal-and-opposite aberrations in the different glass type of the second spherical lens. Several off-the-shelf lenses were modelled in OSLO-EDU, and are presented below as options to be trialled during an observing campaign.

Collection Optics – Option A

The Thorlabs AC508-080-B-ML lens was found to predict a 0.96 mm image diameter from an 80 mm focal length when autofocussed by OSLO-EDU for minimum on-axis spot size. The focal ratio corresponds to $f/1.6$. This is the longest focal length that would almost exactly fill a 1 mm diameter fibre core. The lens is manufactured from Schott N-BAF10 and N-SF6HT glass and has a clear aperture equal to 90 % of the lens diameter, 45.7 mm. This is the maximum focal length we can consider. It is very likely that completely filling the core diameter of the fibre will cause issues with tracking errors as the solar image vignettes at the edge of the fibre core, and so it is sensible to maximise the acceptance angle. The acceptance angle can be considered to be a measure of how precisely the optics must track the apparent movement of the Sun. The smaller the acceptance angle, the more precise the tracking needs to be, and so it is a tolerance budget to be used up on all the system imperfections such as tracking inaccuracy due to gearing, flexure of the mount, movement due to wind, and optical aberrations etc. If the Sun moves outside the acceptance angle then the image will move off the fibre and vignette the solar disc. The image will need to be smaller than the fibre to allow for such tracking inaccuracy, meaning a shorter focal length is desirable.

The shortest focal length lens Thorlabs offer in 50.8 mm diameter is the AC508-075-B-ML, which is 75 mm. However this is manufactured from different glass (N-LAK22/N-SF6HT) and the aberration correction is significantly worse. The minimum on-axis spot size is almost three times larger for the same configuration, despite having a notionally 5 mm shorter focal length. Since we are already at $f/1.6$ we will need to reduce the aperture diameter to allow us to move to shorter focal lengths.

Collection Optics – Option B

The Thorlabs AC254-030-B-ML lens is manufactured from N-BAF10/N-SF6HT glass, and has a focal length of 30 mm with 25.4 mm diameter. The focal ratio corresponds to $f/1.2$. The predicted image diameter when autofocussed for minimise on-axis spot size

is 0.6 mm, and so this will be much easier to focus and align with little risk of vignetting of the solar disc. Unfortunately the clear aperture of 22.8 mm is a 42 % reduction by area compared with a typical BiSON aperture of 30 mm, which would reduce the collected flux from 0.64 mW nm^{-1} to 0.37 mW nm^{-1} based on our assumed power budget determined in Section 7.2. We will consider the effect of this reduction in aperture later in the next section, along with differences in filter transmission coefficients.

Collection Optics – Option C

Both of the above options suffer from potential problems caused by dirt or damage to the end of the fibre, since they rely on an unobscured image of the Sun being accepted and transmitted along the fibre. Some alternative, more esoteric, options have been considered that would allow all spatial information to be scrambled before coupling the resulting light into the fibre, and this would potentially allow for more uniform transmission. A standard instrument in photometry is the *integrating sphere*, a spherical cavity with a diffuse internal coating that allows the total power emitted by a source to be measured with one detector independent of any directional characteristics of the source. This component would be ideal for coupling sunlight into an optical fibre without relying on any imaging onto the fibre, but unfortunately while the measured flux inside the sphere is indeed directly proportional to the total output of the source the absolute transmission is of the order of only a fraction of a percent. Likewise diffusing the light immediately before the fibre entrance causes too much light to be lost.

One potential option is to use a second lens in a typical Keplerian telescope configuration in order to couple quasi-parallel light into the fibre. In doing this we convert spatial-position into angular-position, and so the fibre receives uniform illumination reducing the effect of guiding errors and inhomogeneities of the polished end of the fibre. Unfortunately, again, this configuration lacks transmission efficiency. As a simple estimate, if we used a AC060-010-B-ML 6 mm diameter, 10 mm focal length lens as the "eyepiece" of the Keplerian telescope, then the transmission into a 1 mm fibre would be

of the order of 2.7 % by ratio of aperture area. If we couple this with a AC508-075-B-ML lens from before with a clear aperture of 45.7 mm we produce an effective aperture of just 7.6 mm. This will likely be too small, but could be optimised if it becomes obvious that focusing an image of the Sun onto the end of the fibre suffers insurmountable problems, and so it is worth testing. Another option to maximise both the aperture and acceptance angle is to switch to using non-imaging optics, such as a compound parabolic concentrator – see, e.g., Winston (1974); Rabl (1975, 1976); Rabl and Winston (1976); Winston et al. (2005); Chaves (2008); O'Gallagher (2008). Due to the desire to avoid custom-made components this is considered outside the scope of this project.

All optics were purchased with broadband anti-reflective coatings. Without AR coating, typically 4 % of the light is lost at each optical surface due to reflections. The Thorlabs "B" coating is designed for a wavelength range of 650 nm to 1050 nm and achieves 0.19 % reflectance at 770 nm. An alternative "V" coating is available which achieves 0.065 % reflectance at 770 nm, however the "V" coating is less tolerant to off-axis angles and performance becomes worse than the broadband coating at greater than approximately 25° angle-of-incidence (AOI). The broadband coating provides better uniformity over all AOI and so was considered to provide the best overall performance.

7.3.3 Fibre Output Collimation

There are a number of factors that determine the optimisation of the fibre output collimation (L2 in Figure 7.1 on page 140). Ideally we would like a narrow beam to limit the required diameter of subsequent optics, with a low divergence to limit the angle-of-incidence through subsequent polarisation optics. Furthermore, since an image of the fibre is going to be formed by lens L3 at the centre of the vapour cell, we need to consider the resulting magnification and ensure we do not produce an image that is too large for the cell. These requirements cannot all be met simultaneously, and so some compromise is required.

A 30 mm focal length collimating lens produces a beam ideal for a 25.4 mm optic, with

a divergence of 0.95°. The Thorlabs AC254-030-B-ML lens as used on the input side of the fibre is ideal for also collimating the output. If we compare these parameters with the Jabba spectrophotometer in Carnarvon, Jabba uses a Keplerian telescope to collect sunlight and compress it into a narrow beam inside the instrument. The objective lens has a focal length of 150 mm and the collimating lens a focal length of 50.8 mm, producing a total of almost three times magnification. We know the solar field angle to be 0.271° and therefore the divergence of the beam inside the spectrophotometer is 0.8°, comparable to the divergence achieved with the fibre.

7.3.4 Fibre Output Homogeneity

An optical fibre is known to scramble image information due to the varying entrance angle causing changes in the optical path length through the fibre. Rays launched into the fibre at larger angles of incidence undergo more internal reflections than those launched at angles closer to the optical axis, and this increases the effective path length. Since BiSON spectrophotometers observe in the Sun-as-a-star configuration it would be ideal if all spatial information from the solar disc was lost, and so the additional of an optical fibre to the design is a major enhancement.

To assess the uniformity of the fibre output, a short 2 m section of FT1000EMT was illuminated with a diffuse source and the output butt-coupled to a CCD. Half of the fibre input was then masked off across the diameter, creating an image discontinuity with the same diffuse source. The output from both configurations is shown in Figure 7.5. The light output remains almost completely uniform, even with this worst-case discontinuity in the source image, and so we can be sure that light from all areas of the solar disc is completely scrambled as it leaves the fibre and enters the spectrophotometer. This is important since it allows us to disregard vignetting inside the spectrophotometer caused by varying clear aperture of different components.

In the next section we will look at the filters required to reduce the spectral bandwidth of the collected light, both in terms of coarse filtering to relieve the thermal and UV load

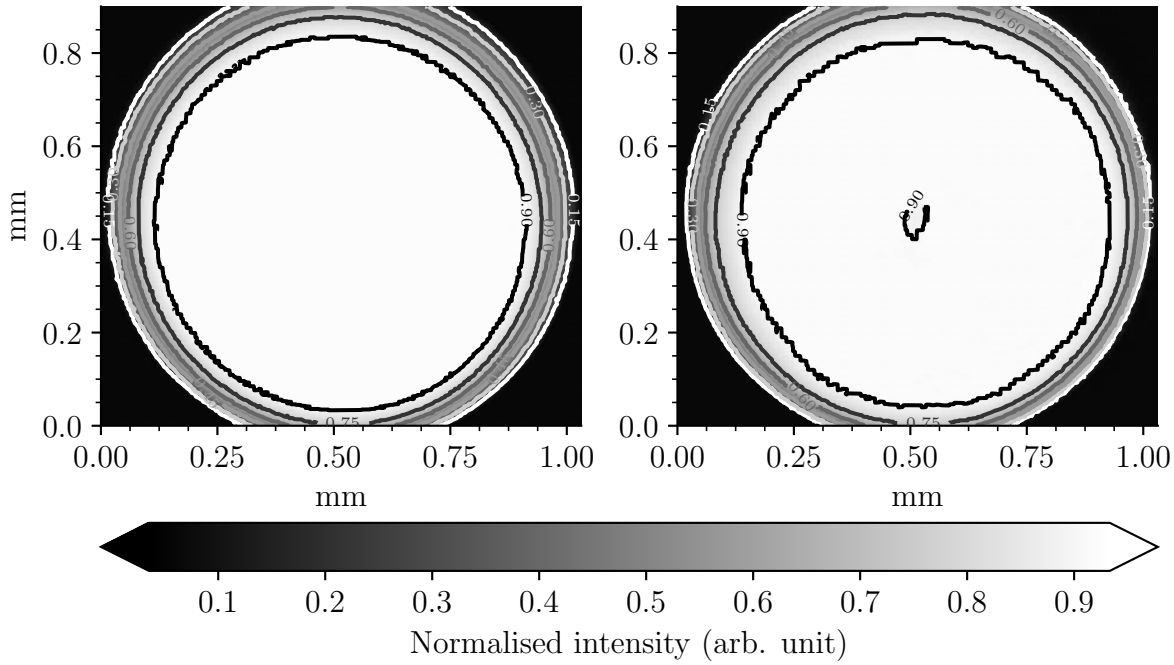

Figure 7.5: Fibre output homogeneity with variable illumination. Left: the output of the fibre with full illumination of the entrance aperture. Right: the output of the fibre when one half of the entrance aperture has been masked across the diameter. The light source was natural diffuse daylight from a window, and the intensity is normalised due to automatic gain control on the imaging sensor.

on the instrument, and fine filtering for wavelength selection.

7.4 Optical Filters

BiSON spectrophotometers typically use three spectral bandwidth filters. The first is a 50 mm diameter, 3 mm thick, Schott RG9 coarse red filter which is placed before the telescope objective lens. It is the first optical element in the system, and is used to limit the spectral range between approximately 750 nm to 1000 nm, relieving much of the infra-red thermal load on the instrument. The second filter is a Schott KG4 short-pass filter that further attenuates the infra-red. Schott KG4 has only 71 % transmission at the standard 2 mm thickness, and so the thickness is specified at only 1 mm to increase the transmission to 81 %. The reduced thickness filter is very fragile, and so the filter diameter is also reduced to 25 mm and placed after the objective lens in the converging

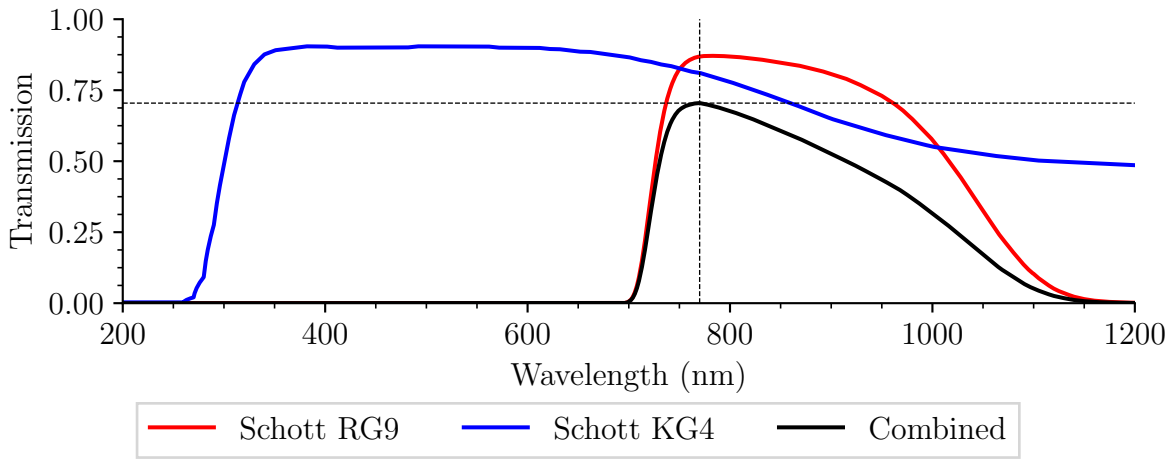

Figure 7.6: Schott RG9 and KG4 filter profiles as used in many BiSON spectrophotometers. The KG4 filter has been discontinued.

part of the beam. Finally, an interference filter with a narrow 1.5 nm bandwidth is the third filter in the system, and this is used to isolate the potassium D_1 absorption line at 770 nm from the nearby D_2 absorption line at 766 nm. We will now look in more detail at the coarse filters, and then later consider the narrow interference filter.

The filter profiles of Schott RG9 and KG4 are shown in Figure 7.6, along with the combined profile of both filters. At 770 nm the transmission is 70 %, which is not optimal. Furthermore, Schott have discontinued the KG4 filter and none of the other KG filters are suitable since they suffer higher attenuation at 770 nm. While some suppliers have remaining stocks of KG4 available, it would not be wise to continue to design new instrumentation around it. We are not the only group to have problems with the loss of Schott KG4. The Johnson-Cousins UVBRI filters have been the standard for photometric measurements for many decades. They were designed in the 1950s by H. Johnson and later evolved by A.W.J. Cousins in the 1970s for use predominately, although not exclusively, with photomultiplier tubes (PMT). The infra-red "I" filter of the UVBRI set was manufactured using Schott RG9 in conjunction with Schott KG4, the same combination as used by BiSON. Additionally, the original PMTs of the time were sensitive to only around 900 nm whereas modern CCDs are sensitive to around 1100 nm, and other detector types extend further into the infra-red. This means that not only can the original filter

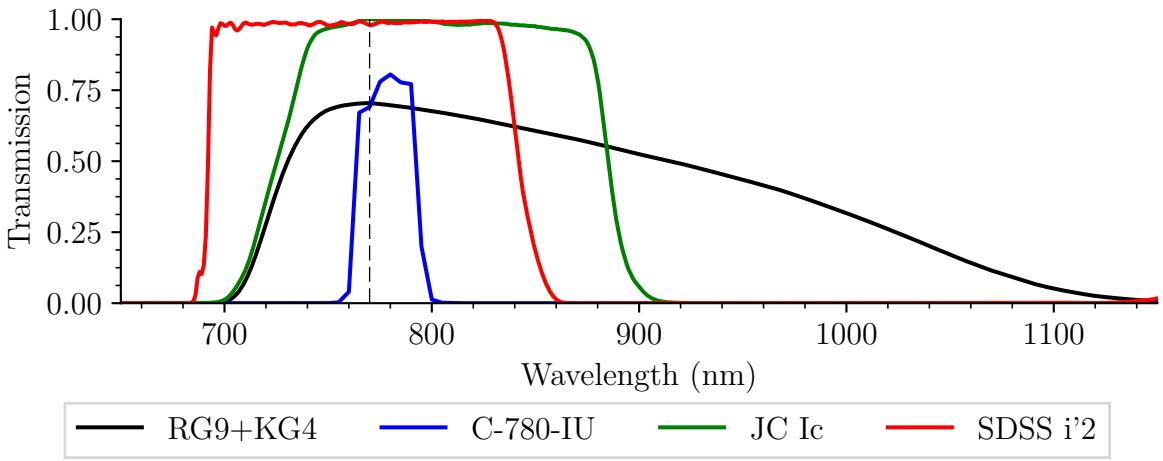

Figure 7.7: Alternative coarse filters in replacement of Schott RG9 and KG4.

specification no longer be purchased, but the filter needs to be updated with a 900 nm short-pass cut-off if the filter is to produce a comparable response with other detector types as the original PMT passband. The most common filter set in use today is the photometric system designed by Fukugita et al. (1996) for the *Sloan Digital Sky Survey* (SDSS), in response to the requirement for five mostly non-overlapping filters covering the 300 nm to 1100 nm range of silicon CCD cameras. The *Hubble Space Telescope* has an SDSS filter set, and much of the photometry with metre-class telescopes is done in this system.

We considered several potential replacement filters, and produced a shortlist of the alternatives. Astrodon in the United States have developed new filters for both the UVBRI and SDSS sets. The filters are manufactured using hard-oxide sputtered dielectric coatings, rather than coloured glass, and so achieve very high transmission with close to ideal cut-on and cut-off passbands. Figure 7.7 shows the profiles of the Astrodon Johnson/Cousins I_c and SDSS i'2 filters. A third option from Comar Optics, model C-780-IU, and also the RG9/KG4 profiles are included for comparison. Either of the Astrodon filters would be ideal for our use, but the SDSS i'2 filter was selected as the most appropriate due to its shorter cut-off for better infra-red rejection. The i'2 has a bandwidth of 695 nm to 844 nm and at 770 nm the transmission is 98 %, an increase of 36 % in comparison to the original Schott KG4/RG9 filters. We will now look at the

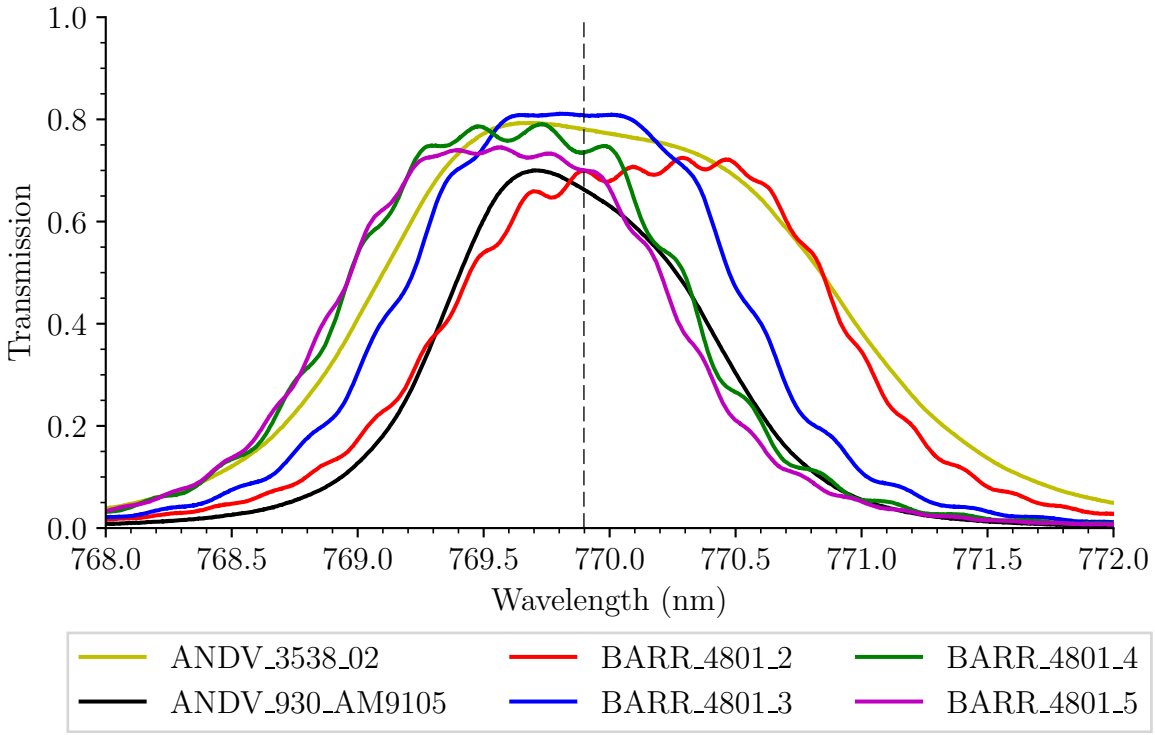

Figure 7.8: Measured profiles of several 1.5 nm interference filters from group stock.

final filter stage, the narrow wavelength selection.

The final filter at the entrance to the spectrophotometer is a 1.5 nm interference filter, which isolates the 770 nm absorption line from other nearby absorption lines. Several filter types were already in stock having been purchased by the group previously. These were all assessed using an Agilent 86142B optical spectrum analyser, and the measured profiles are shown in Figure 7.8. The four filters manufactured by Barr Associates, numbered 2–5, all show widely varying performance despite being from the same batch. The two filters manufactured by Andover Corporation are also quite different. The best filter was selected as being the Barr Associates filter labelled as number three, since it has the highest transmission at 81 %, is well centred at 770 nm, and shows a wide plateau over the central region which will provide better stability. The bandwidth of interference filters is sensitive to changes in temperature and off-axis angle. The off-axis angles through BiSON spectrophotometers are typically small and have negligible effect since the 1.5 nm bandwidth is considerably larger than the bandwidth of the potassium absorption line,

and so small changes in filter bandwidth are not seen by the instrument. Where large changes in ambient temperature can be experienced, such as in our BiSON domes, the interference filter is thermally controlled. In the case here where the instrument is held in a thermally controlled room separated from the external observation area, small fluctuations in ambient temperature are insignificant and temperature control is not required.

Now that we have fully defined the apertures and filters used in the input stage of the spectrophotometer, we can estimate the total power through the system. In Section 7.2 we determined an optical power budget of $0.9\,\mathrm{W\,m^{-2}\,nm^{-1}}$ near 770 nm. The clear aperture of the new instrument is 22.8 mm collecting $0.37\,\mathrm{mW\,nm^{-1}}$. After passing through the filters, and also accounting for losses in the fibre of $7\,\mathrm{dB\,km^{-1}}$ and 4 % at each surface, the total power into the spectrophotometer is 0.27 mW. By comparison, the Jabba spectrophotometer has a 30 mm aperture collecting $0.64\,\mathrm{mW\,nm^{-1}}$, and after passing through the coarse RG9, KG4 and fine filters a total of 0.36 mW enters the spectrophotometer. By making use of filters with higher transmission we have somewhat mitigated the losses caused by smaller objective aperture, and can expect a similar signal to noise ratio.

In the next section we will look at polarisation control before passing light into the potassium vapour cell and estimating the power reaching the detector photodiodes.

7.5 Polarisation Control

A variety of linear polarisers and quarter-wave retarders have been used to control and monitor the polarisation state of the beam inside BiSON spectrophotometers. These vary from fixed filters and retarders, photo-elastic modulators, to most recently Pockels-effect cells. Different types of retarder offer different rates of switching between polarisation states. A motorised rotating quarter-wave retarder, such as that used by the GOLF instrument, takes several seconds to switch between states. An electro-optic Pockels-effect cell is capable of switching rates of at least 100 Hz, while a photo-elastic modulator typically operates in excess of 50 kHz.

The Pockels-effect cells currently in use by the group are custom designed field-widened devices, and require the use of custom in-house designed high-voltage power supplies. Throughout this project we would like to simplify construction by making use of COTS components. The most common electrical retarders are LCD based, and these switch states at much lower rates than can be achieved with Pockels-effect cells. In order to assess the potential impact of reduced switching rates we must first profile the performance of a Pockels-effect cell in terms of slew rate. We will do this next, and then go on to consider the best slew rate achieved with an LCD retarder.

7.5.1 Pockels-effect Cells

The Pockels-effect cells used in BiSON instrumentation are constructed using potassium dideuterium phosphate KD_2PO_4 (KD*P), which belongs to the tetragonal crystal system (Phillips, 2011). This crystal exhibits birefringence and is normally optically uniaxial, however when placed in an electric field it becomes biaxial. The KD*P crystal is z-cut, which means that the two faces of the crystal are normal to the crystallographic z-axis. The on-axis rays through the spectrophotometer are parallel to the crystallographic z-axis, and the beam propagates in the positive z-direction. The cell acts as a linear retarder to a beam passing through the Pockels cell parallel to the optical axis. An electric field can be applied via transparent electrodes applied to the input and output faces of the crystal. The retardance varies linearly with the voltage applied across the electrodes, and at approximately 2300 V the retardance reaches 45° and the Pockels cell behaves like a quarter-wave plate. When combined with a linear polariser this causes one of the polarisation components to be retarded by 90° in phase (a quarter of a wavelength) in relation to the other component, producing circularly polarised light. Circular polarisation is considered to be left- or right-handed depending on direction of rotation of the electric field vector of the wave. If the electric field across the cell is reversed then the retardance becomes −45° and the opposite handedness of circular polarisation is produced. This is true only for on-axis rays, and as the off-axis angle increases the degree of circular polarisation

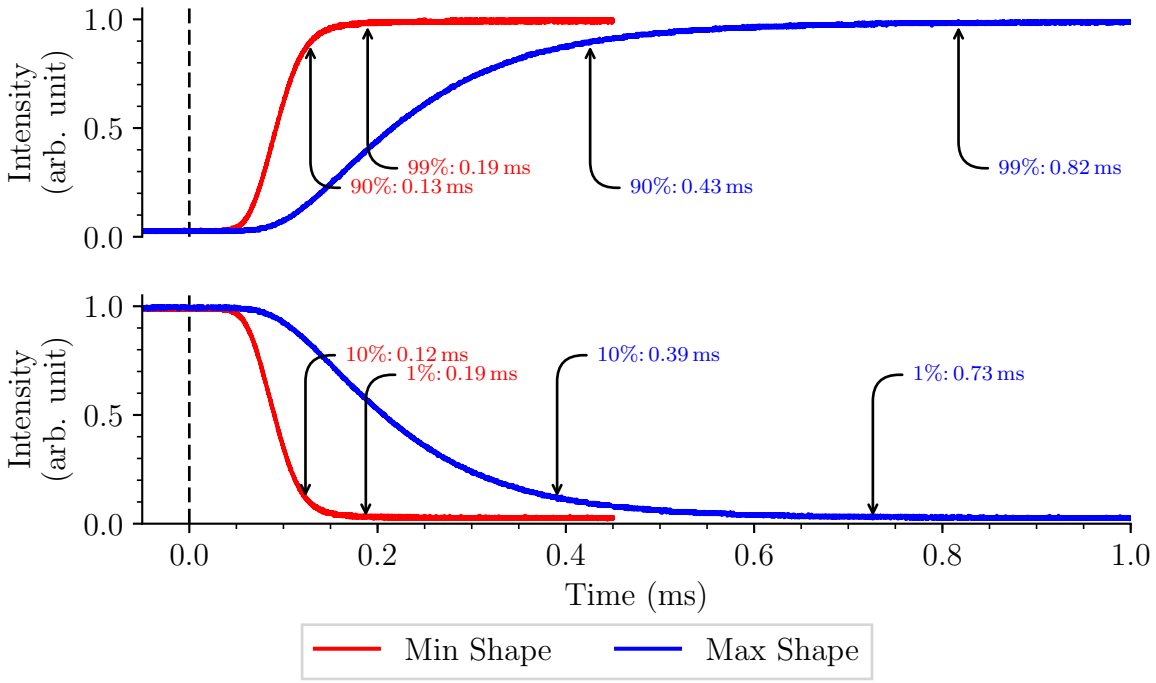

Figure 7.9: Pockels-effect cell switching rate. The driver has a "shape" control that smooths the transition and helps to relieve mechanical stress on the crystal at the expense of a longer transition period. The arrows indicate the time at which the stated percentage of the final value is reached. The dashed vertical lines indicate the switching command trigger point.

changes. We saw in Section 7.3 that the divergence of the beam in the quasi-parallel part of the optical path is almost 1°, and so it is important to understand how the Pockels cell behaves at these angles.

The optical properties of KD*P crystals have been modelled in significant detail by Miller (2012a,b,c,d) in terms of Stokes parameters (Stokes, 1851) and Mueller matrices. Miller (2012e) shows that a Pockels cell manufactured from a single 5 mm z-cut KD*P crystal has a useful operating field of view of approximately ±1°. More recent instruments use a field-widened Pockels cell developed within the group by Prof. Y. P. Elsworth and Prof. C. C. Speake, which consists of two 5 mm z-cut KD*P crystals separated by a 4.57 mm z-cut quartz crystal acting as a 90° rotator. The transparent electrodes on each crystal are wired such that the electric fields are in opposite directions and the operating voltage is reduced to approximately 1150 V, meaning that each crystal contributes an eighth-wave retardance. Miller (2012f) shows that the field of view of this three-element

Pockels cell is approximately ±2°, allowing greater tolerance to off-axis rays.

It is not intended to reproduce any of the modelling work completed by Miller here since, for our purposes, it is unnecessary to have a rigorous mathematical understanding of the fundamental principles. It is necessary only to understand the Pockels cell acts as a variable retarder, where the retardance is proportional to the applied voltage. What is of much greater importance is the rate at which the device can switch between ±45° retardance, since this determines how fast we can switch between sampling the red- or blue-wing of the solar potassium absorption line.

A field-widened dual-crystal Pockels cell was taken from group stock (cell labelled "L16") and the switching rate was tested using a photodetector with 100 kHz bandwidth. In industry it is typical to quote the slew-rate of a device as the time taken to transition from 10 % to 90 % of the final value. Here we quote time taken from the instant the change is commanded to when the output reaches 90 % and 99 % of the final value, and so performance appears lower than datasheet advertised specifications. The switching rate for transitions in both polarisation directions is shown in Figure 7.9. The driver has a "shape" control that smooths the transition and helps to relieve mechanical stress on the crystal at the expense of a longer transition period. At minimum shape the switching time to 99 % of the final value is approximately 0.2 ms. At maximum shape the switching time is approximately 0.8 ms, and this is broadly symmetric in both directions. We'll see later in Chapter 8 on the data acquisition sub-systems that the high-gain scattering detectors do not have sufficient bandwidth to see this transition rate. The detectors require at least 3 ms to fully pass a transient signal, and so 3 ms is our target rate for comparison. We'll look next at the switching rate of a liquid crystal retarder.

7.5.2 Liquid Crystal Retarder

The Thorlabs LCC1223T-B is a nematic liquid crystal cell that can be used as a variable wave plate, and is anti-reflection coated over the range 650 nm to 1050 nm. Many nematic crystal organic molecules are uniaxial with one long axis and two others of equal length,

similar to a cylinder. In the nematic phase the cylinders have no preferential order, but they self-align with their long axes approximately parallel and this creates an optical anisotropy. When an electric field is applied the molecules tilt and align to the field, and so changes in the field strength cause changes in the effective retardance. The driving signal is typically between 0 V to ±25 V, considerably lower, more convenient, and safer than the 2300 V required by a Pockels-effect cell. The bipolar drive voltage is switched at 2 kHz and is required to minimise the DC bias, and so minimise damage caused by the build up of ions in the liquid crystal material. The LCC25 controller can be computer-controlled to allow switching between two predetermined control voltages, while automatically ensuring zero DC bias across the device.

LCD retarders are typically slower to switch between states than a Pockels-effect cell, and so we need to test if we will be able to achieve the standard 100 Hz switching rate used by existing BiSON spectrophotometers. Switching rate is affected by both the applied voltage, and the crystal temperature. When switching from a low voltage to a high voltage, the liquid crystal becomes "tense" and the change occurs rapidly. In the reverse direction the rate is limited by how long it takes the cell to "relax" back to its natural state and this is determined by the liquid viscosity, and so by the temperature. The low-to-high voltage transition is much faster than the high-to-low voltage transition. Higher temperatures and larger voltage differences allow for faster switching. The overall retardance changes slightly with temperature, and so it is important to calibrate the control voltages at the intended operating temperature.

Figure 7.10 shows the time for a half-wave transition between two different states at room temperature, with times to 90 % and 99 % indicated. The calibrated control voltages are listed in Table 7.1. When considering $\pm\lambda/4$, switching in the fast direction (low-to-high voltage) takes between 15 ms to 26 ms. Switching in the slow direction (high-to-low voltage) takes between 67 ms to 120 ms. This rate is unacceptably slow due to the use of low, similar, control voltages. The transition between $3\lambda/4$ and $\lambda/4$ requires control voltages of 1.78 V and 4.03 V respectively. The switching time can be reduced by

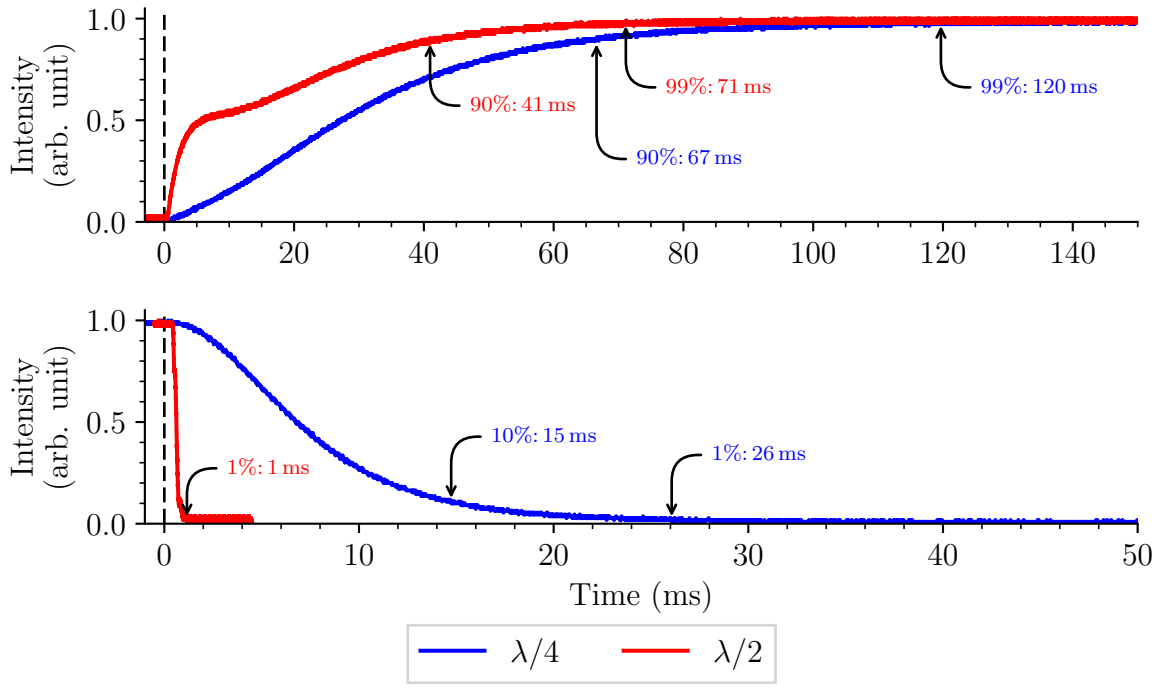

Figure 7.10: Liquid-crystal switching rate. Top: high-to-low voltage transitions. Bottom: low-to-high voltage transitions. The arrows indicate the time at which the stated percentage of the final value is reached. The dashed vertical lines indicate the switching command trigger point.

Table 7.1: Calibrated LCD control voltages at room temperature.

	LCD Control Voltage				
Retardance	λ	$\frac{3\lambda}{4}$	$\frac{\lambda}{2}$	$\frac{\lambda}{4}$	zero-wave
Voltage	1.34	1.78	2.32	4.03	25.00

adding a fixed quarter-wave plate and instead operating the LCD between zero- and half-wave retardance. The transition between zero-wave retardance and $\lambda/2$ requires control voltages of 25 V and 2.32 V respectively, reducing the fast transition to just 1 ms and the slow transition to between 41 ms to 71 ms.

The fast switching direction is now better than the scattering detector bandwidth, but the slow direction is still unacceptable. In order to reduce the transition period further we need to warm the liquid crystal to higher temperatures, and we will look at that next.

7.5.3 LCD Temperature Control

The LCC1223T-B LCD is capable of being temperature-controlled in order to stabilise the retardance and reduce the switching-time between polarisation states. The device contains a heating element and a TH10K thermistor temperature sensor. In order to control the temperature, we need to readout and calibrate the output from the temperature sensor.

A negative temperature coefficient (NTC) thermistor such as the TH10K relies on the property of intrinsic (non-doped) semiconductors that electrical resistance decreases exponentially with increasing temperature. The effect is described by the *Arrhenius equation* which states that,

$$R_{\mathrm{T}} = A\, e^{\beta/T} \ , \tag{7.8}$$

where R_{T} is the resistance at a given temperature T, and A and β are constants. The constant β is a parameter that defines the energy required to move electrons from the valence band to the conduction band, and so higher values of β mean the material is more insulating. A thermistor can be calibrated by introducing a reference temperature producing a standard resistance,

$$R_{\mathrm{T}_0} = A\, e^{\beta/T_0} \ , \tag{7.9}$$

and then solving equation 7.8 and 7.9 simultaneously, such that,

$$\ln R_{\mathrm{T}} - \ln R_{\mathrm{T}_0} = \ln \frac{R_{\mathrm{T}}}{R_{\mathrm{T}_0}} = \beta \left(\frac{1}{T} - \frac{1}{T_0} \right) \ , \tag{7.10}$$

and so,

$$\beta = \frac{\ln \frac{R_{\mathrm{T}}}{R_{\mathrm{T}_0}}}{\frac{1}{T} - \frac{1}{T_0}} \ . \tag{7.11}$$

We can then solve for either R or T,

$$R_{\mathrm{T}} = R_{\mathrm{T}_0} e^{\beta(\frac{1}{T} - \frac{1}{T_0})} \ , \tag{7.12}$$

and,

$$T = \frac{1}{\frac{1}{\beta} \ln \left(\frac{R_{\mathrm{T}}}{R_{\mathrm{T}_0}} \right) + \frac{1}{T_0}} \ , \tag{7.13}$$

and so calibration can be completed using a minimum of two measurements. Generally the value of β is provided by the manufacturer in the component datasheet, and so this can be used directly with equation 7.13 to immediately begin converting resistance values to temperature with no additional calibration requirement. Unfortunately, β is temperature dependent and so such calibration is valid over only a small range of temperatures – typically 20 °C at most. Steinhart and Hart (1968) investigated about 100 models to better fit the observed variation in resistance from two datasets produced using sensors immersed in a laminar-flow constant-temperature bath with milli-Kelvin precision over a temperature range of 0 °C to 35 °C, and a third dataset over a range of −68 °C to 134 °C. The one consistently best model was,

$$T^{-1} = A + B \ln R_{\mathrm{T}} + D(\ln R_{\mathrm{T}})^3 \ , \tag{7.14}$$

where A, B, and D are the fitted constants, R_{T} is the measured resistance, and T^{-1} the inverse temperature in Kelvin. The mean of the residuals from fits to the model in equation 7.14 were approximately the same size as the experimental precision. The work of Steinhart and Hart (1968) was later independently confirmed by Mangum (1983) and Koehler and von Herzen (1986), and equation 7.14 has become known as the Steinhart-Hart Equation.

The datasheet for the TH10K thermistor suggests a calibration equation similar to the

Table 7.2: TH10K thermistor calibration coefficients reproduced from the datasheet. Resistance values are in Ohms, with higher values corresponding to lower temperatures. These values can be used with equation 7.15 to convert resistance to temperature.

Resistance		A	B	C	D
Min	Max				
32 770.0	692 600.0	3.357×10^{-3}	2.521×10^{-4}	3.374×10^{-6}	-6.496×10^{-8}
3599.0	32 770.0	3.354×10^{-3}	2.562×10^{-4}	2.140×10^{-6}	-7.241×10^{-8}
681.6	3599.0	3.353×10^{-3}	2.542×10^{-4}	1.143×10^{-6}	-6.938×10^{-8}
187.0	681.6	3.354×10^{-3}	2.538×10^{-4}	8.543×10^{-7}	-8.791×10^{-8}

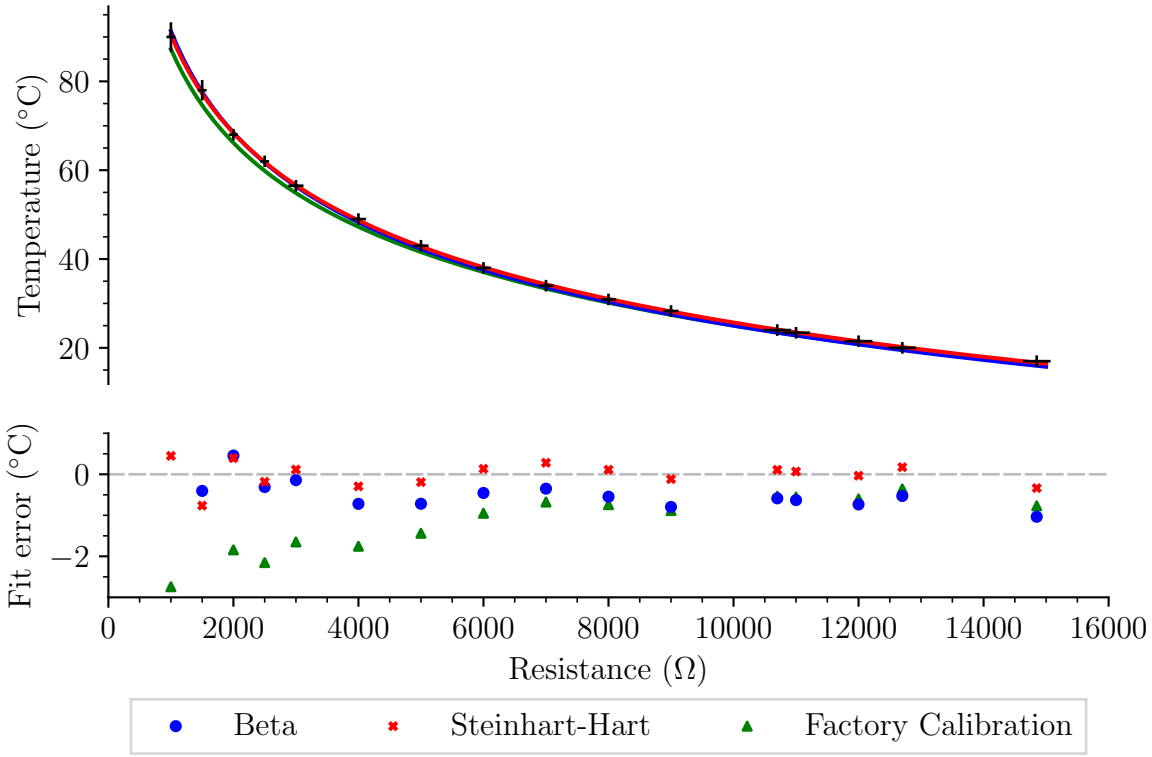

Figure 7.11: Fits of the Beta model and the Steinhart-Hart model against measured resistance and temperature data. The Steinhart-Hart model produces better results, with lower residual variance and an almost zero residual mean. The factory calibration using the datasheet values is shown for comparison.

Steinhart-Hart Equation but additionally including the squared-term in the polynomial,

$$T^{-1} = A + B \ln\left(\frac{R_\mathrm{T}}{R_{25\,^\circ\mathrm{C}}}\right) + C \ln\left(\frac{R_\mathrm{T}}{R_{25\,^\circ\mathrm{C}}}\right)^2 + D \ln\left(\frac{R_\mathrm{T}}{R_{25\,^\circ\mathrm{C}}}\right)^3 , \tag{7.15}$$

where $R_{25\,^\circ\mathrm{C}}$ is the reference resistance at 25 °C which is equal to 10 000 Ω. Four sets of the four coefficients, A, B, C, and D are provided for use within different resistance ranges making a total of 16 parameters. The factory-calibrated coefficients are listed in Table 7.2.

To check the temperature calibration, the LCD was heated and the sensor output stabilised based on a range of different resistance values. The temperature at each resistance was then measured using a *Fluke 62 MAX+* non-contact infra-red thermometer, which has a stated accuracy of ±1 °C and a precision of ±0.1 °C. The actual experimental error is somewhat larger than these values due to the LCD having only one heating ele-

Table 7.3: Fitted coefficients for the Beta model and the Steinhart-Hart model.

	Model	
Coefficient	Beta	Steinhart-Hart
β	3772 ± 23	
A		$103.256 \pm 6.389 \times 10^{-5}$
B		$24.540 \pm 1.149 \times 10^{-5}$
D		$6.850 \pm 5.406 \times 10^{-8}$

Table 7.4: Statistical performance of the residuals from each thermistor calibration technique.

Model	Variance	Sigma	Mean
Beta	0.302	0.549	−0.386
Steinhart-Hart	0.087	0.295	−0.005
Factory calibration	0.721	0.849	−1.313

ment and a finite size, resulting in slightly different temperature values depending on the measurement point on the device. The estimated uncertainty in measured temperature is approximately ±1 °C at 20 °C, increasing to ±3 °C at 90 °C as the thermal gradient across the device increases. Additionally, the measured resistance has uncertainty of approximately ±200 Ω at 20 °C where the sensitivity to small temperature fluctuations is greatest, reducing to approximately ±50 Ω at 90 °C. Since the uncertainty has been defined on both the dependent and independent variables, an orthogonal distance regression (ODR) technique was used to fit both the β model (equation 7.13) and the Steinhart-Hart model (equation 7.14) to the acquired data. The results are shown in Figure 7.11, and the fitted coefficients listed in Table 7.3. The statistical performance of the model residuals is listed in Table 7.4. While both models produce results better than the factory calibration, the Steinhart-Hart model does produce better results with lower residual variance and an almost zero residual mean. The difference between the measured values and the factory calibration is likely due to the temperature gradient across the device resulting in different measurements depending on the installation location of the sensor. If the sensor is mounted on the opposite side of the device to the heating element, the points nearer the heater will be at a higher temperature than indicated by the sensor. This tempera-

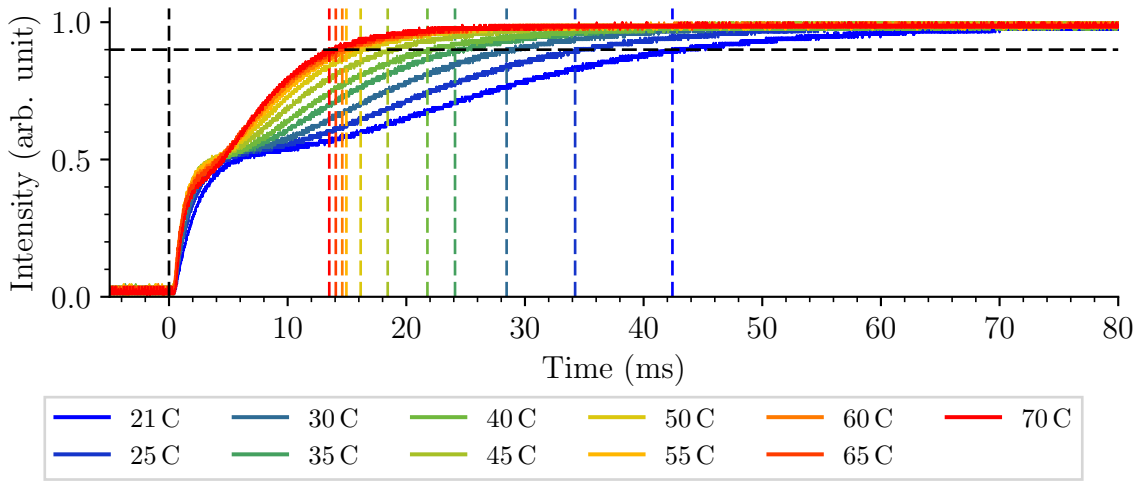

Figure 7.12: Temperature dependence of the LCD switching rate. The black dashed vertical line indicates the switching command trigger point. The dashed horizontal line indicates the 90 % complete point.

Table 7.5: LCD switching rates at a range of temperatures.

	Switching rate (ms) at Temperature (°C)										
Temp	21	25	30	35	40	45	50	55	60	65	70
90 %	42	34	28	24	22	18	16	15	15	14	14
99 %	74	57	48	42	43	31	27	26	27	27	25

ture gradient effect should reduce given sufficient time for the temperature to stabilise. The Steinhart-Hart model was selected as the calibration technique for the temperature control requirements of the device using equation 7.14 and the coefficients in Table 7.3 based on the actual measured temperatures of the device.

Since the fast (low-to-high voltage) switching is already acceptable at under 1 ms, we only need to consider improvements to the slow (high-to-low voltage) switching rate. Figure 7.12 shows the high-to-low voltage switching time from zero-wave to half-wave for a range of calibrated temperatures. Table 7.5 shows the switching time for both the 90 % and 99 % transition completion points. There is little improvement above 50 °C and so this was selected as the ideal operating temperature, producing a switching rate of between 16 ms to 27 ms or approximately 4.4 times faster than the original test at 120 ms shown in Figure 7.10 on page 161. However, even if we accept a 90 % transition as suf-

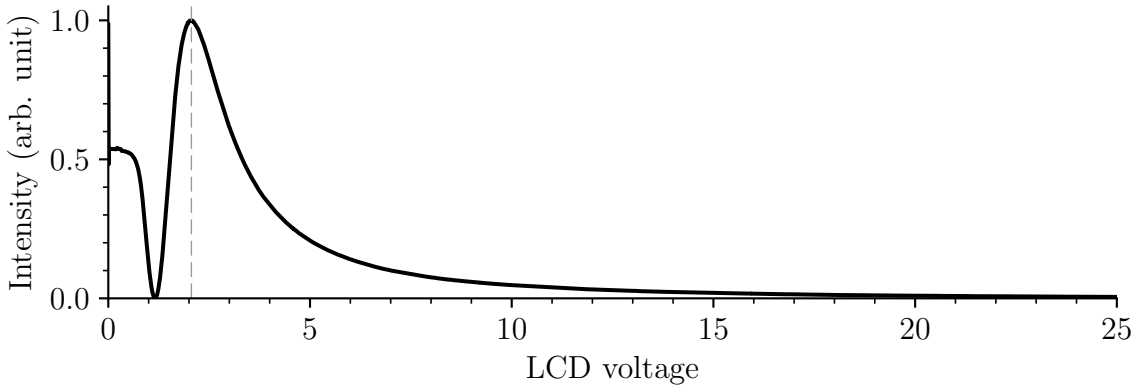

Figure 7.13: Liquid-crystal retardance calibration. Izaña 2017 May 13. Half-wave is achieved at 2.06 V at 50 °C.

ficiently stabilised to begin data acquisition, the 16 ms dead period is still over 5 times longer than the 3 ms maximum we would like to achieve. We will not be able to switch polarisation states at 100 Hz since this would require a stabilisation time no longer than 10 ms without even considering a period of data acquisition after stabilisation is complete. A reduced switching rate of 5 Hz limits the dead time per 4-s acquisition period to 25 %, and this is increased from 15 % when switching at 100 Hz with a Pockels cell. Figure 8.24 on page 217 shows that atmospheric scintillation noise has reached −3 dB by about 5 Hz, and so switching at a rate of 5 Hz will provide an improvement in noise level compared with the 0.5 Hz rate used at Izaña. Ideally we would need to reach about 20 Hz to reduce scintillation noise to −9 dB and achieve performance similar to 100 Hz switching. We will consider switching rates and dead time in more detail in Chapter 8 on data acquisition.

Since the aim is to remove custom Pockels-effect cells and associated drivers, we will continue with the COTS LCD as the polarisation control component for this design despite the expected reduced performance. Figure 7.13 shows a retardance calibration voltage-scan after the LCD was commissioned at Izaña and stabilised at 50 °C. Half-wave retardance is achieved at 2.06 V, and zero-wave retardance at the maximum 25 V, and these are the final operating settings.

The LCC1223T-B has only a 20 mm clear aperture, despite being marketed as a standard 25.4 mm optic. This results in some slight vignetting causing the étendue to reduce

to 0.27 mm^2 sr. The initial linear polariser provides 42 % transmission of incident unpolarised light, and so these two issues reduce the optical power to 0.08 mW at this point in the system as it enters the vapour cell.

In the next section we will go on to look at a commercial off-the-shelf potassium vapour cell and oven.

7.6 Vapour Cells

In Chapter 6 we determined that for a typical BiSON vapour cell and oven the ideal cell stem temperature is approximately 88 °C, and for the main cell-cube anything sufficiently higher than this ensures the glass walls remain clear of solid potassium – typically a differential of 20 °C is maintained. Many BiSON cell ovens, such as in Figure 7.14, run at a stem temperature of 90 °C to 100 °C and a cube temperature of 110 °C to 120 °C. The ideal cell stem temperature varies, as we saw in Figure 6.10 on page 135, due to differences in the size and position of the scattering detector aperture.

By making use of the model developed in Chapter 6 we can assess the expected performance of an off-the-shelf vapour cell and oven, such as the Thorlabs GC25075-K cell and GCH25-75 oven shown in Figure 7.15. This cell is quite different from the type currently used by BiSON. It is cylindrical, 72 mm long, and 25 mm in diameter. The filling stem is cut off very short to leave only a small nub. The oven is formed from two separate heating elements that sit at each end of the cell. Figure 7.16 shows the expected absorption profile through the cell for a range of temperatures, and the shaded areas represent the extent of the two ovens. At high temperatures, above approximately 90 °C, all of the light is scattered within the first oven and so it is impossible to detect any of the scattered light. If we assume that we could capture scattered light from the entire central section of the cell, between the ovens, then the optimum scattering temperature determined from the model developed in Chapter 6 is approximately 70 °C. This is the temperature at which the greatest reduction in transmitted intensity is seen between the ovens, i.e., maximised

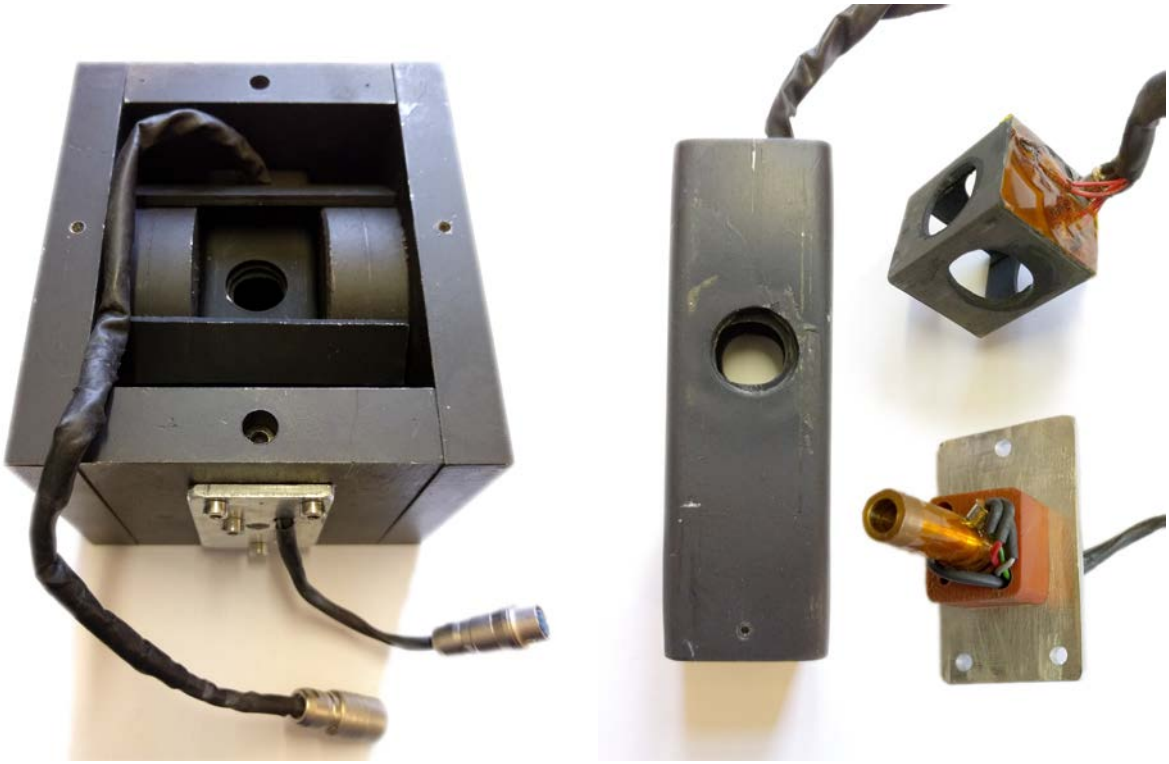

(a) BiSON magnet housing (b) BiSON vapour cell oven

Figure 7.14: A typical BiSON magnet and vapour cell oven. (a) shows the magnet and yoke, with the cell oven installed inside. (b) shows the cell oven, including the separate cube and stem heaters.

scattering in the detector area. However, there are significant engineering problems when attempting to apply this technique.

The typical use case for these cells is in Saturated Absorption Spectroscopy (SAS) providing a stable locking reference for a tunable diode laser at the centre of an atomic transition, and as such only the entrance and exit windows are designed to remain clear. Without a long stem, the coldest part of the glass becomes the central point between the heaters and so this is where deposition of solid potassium occurs, which makes the walls opaque precisely where we intend to observe resonant scattering. We know from Chapter 6 that the most critical temperature is that of the coldest point of the cell, i.e., the temperature of the solid potassium, and this heater configuration allows little stabilisation of the temperature of the coldest point. Additionally it is difficult to design a photodetector that has a sufficiently large acceptance angle to cover the whole central

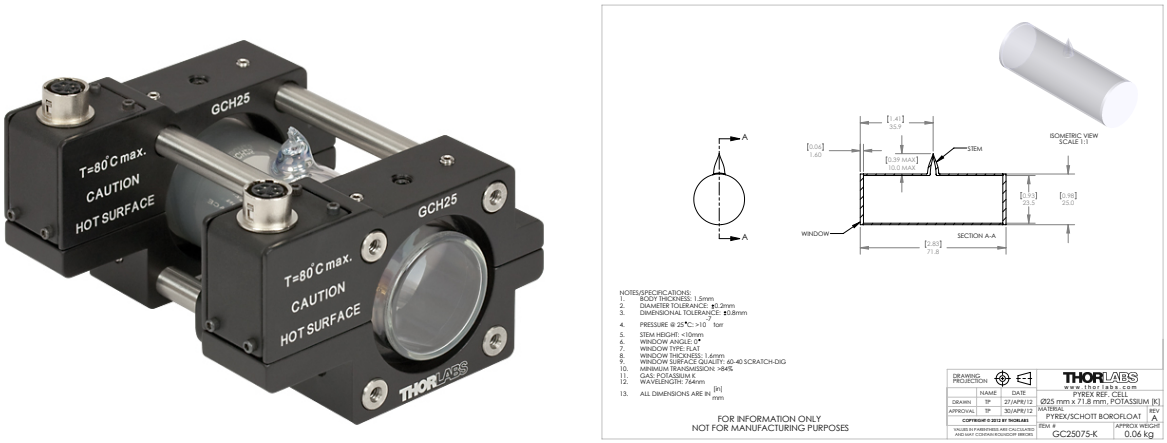

(a) Thorlabs GCH25-75 oven (b) Thorlabs GC25075-K cell

Figure 7.15: A typical commercial off-the-shelf vapour cell and oven. (a) shows the Thorlabs GCH25-75 oven with installed cell. (b) shows a schematic of the Thorlabs GC25075-K vapour cell. Image credit: Thorlabs, Inc.

section of the cell, whilst also rejecting all non-resonantly scattered light.

A long cylindrical cell similar to the GC25075-K has been developed for the GOLF-NG mission (Turck-Chièze et al., 2008), and is shown in Figure 7.17. This custom cell has a long stem which is required to keep the main body of the cell clear of solid potassium. Foil heaters are used to stabilise the temperature of the whole length of the upper cylindrical section, and a separate second foil heater is used to control the stem at a lower temperature, essentially making a large version of a standard BiSON cell. The custom cell housing consists of many exit ports, allowing many detectors to be used rather than a single detector with large acceptance angle.

Whilst the manufacturer can easily produce a cell with the stem left longer than normal, we are then moving away from an off-the-shelf component and back towards a custom part. We would also need to develop an oven and housing similar to that designed for GOLF-NG, and at this point we have completely missed our off-the-shelf specification requirement. We have to conclude that it is unlikely to be possible to develop a resonant scattering spectrophotometer using solely off-the-shelf components. In order to continue developing and testing the new system, we make use of the existing vapour cell and oven shown in Figure 7.14.

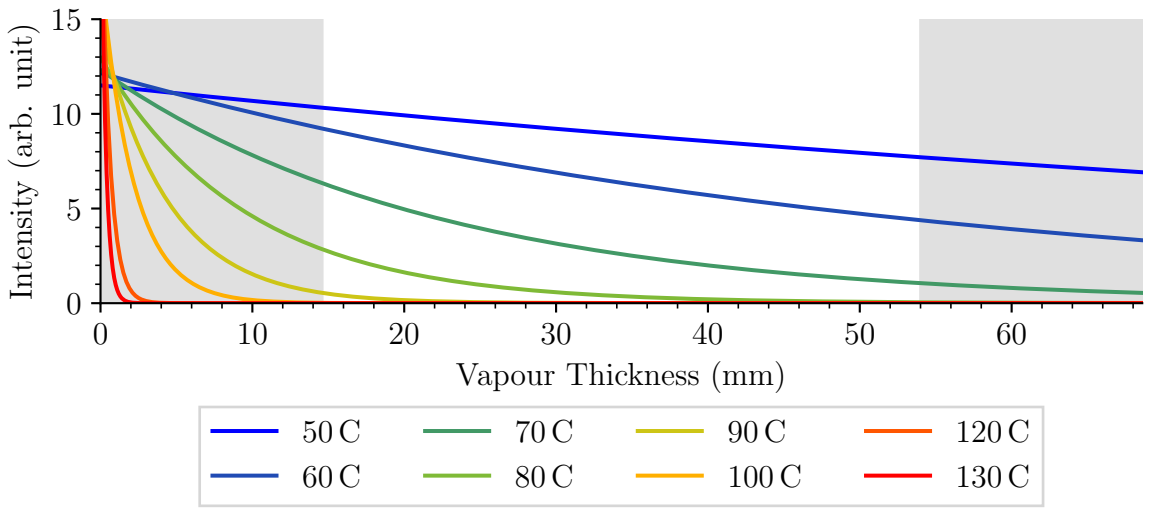

Figure 7.16: Absorption through a Thorlabs GC25075-K vapour cell at various temperatures. The absorption coefficient is modelled as monochromatic but with the initial intensity scaled by the expected FWHM to compensate for the differences in spectral bandwidth. The shaded areas at each end represent the extent of the oven heating elements, and scattering cannot be observed from these regions.

When developing the data acquisition system in the next chapter, it will be necessary to know the power incident on both the scattering and transmission photodiodes. In the previous section we determined that the optical power into the vapour cell is 0.08 mW, with a bandwidth of approximately 1 nm through an aperture 20 mm in diameter. It is possible to estimate the fraction that reaches the scattering photodiodes by considering the geometry of the cell, the scattering cross section, and the collection optics. From Chapter 6 we know that at the optimum vapour temperature the optical depth is unity at the centre of the cell and so we first scale by,

$$I = I_0\, e^{-1} \ , \tag{7.16}$$

making an assumption of single-scattering. This is acceptable since we require only a simple power estimate and not high accuracy. We then scale by the scattering cross-section,

$$I = I_0\, \frac{\lambda_{cs}}{\lambda_{beam}} \ , \tag{7.17}$$

where λ_{cs} is the scattering cross-section of approximately 4 pm and λ_{beam} is the incoming

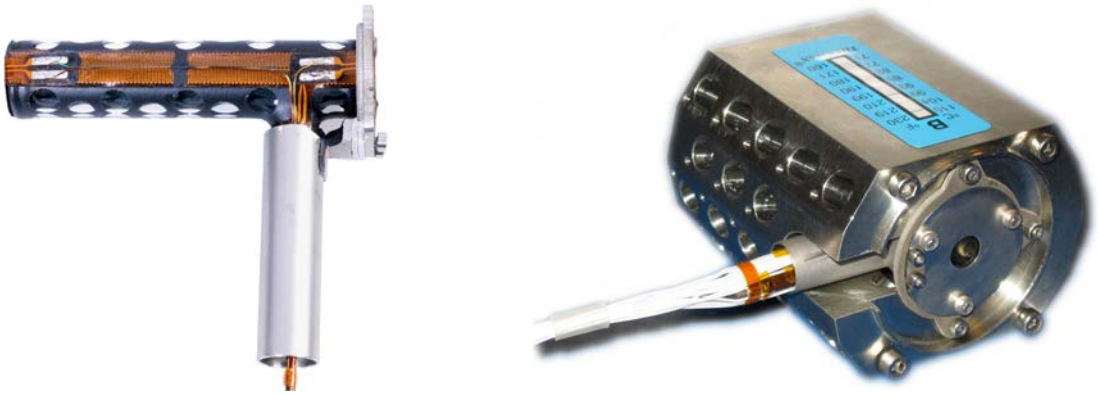

(a) GOLF-NG vapour cell (b) GOLF-NG oven housing

Figure 7.17: The GOLF-NG vapour cell and oven. (a) shows the long cylindrical vapour cell with foil heaters along the entire length, and an extended stem for controlling the position and temperature of the coldest point. (b) shows the cell oven housing with multiple scattering exit ports. Image credit: Reproduced from Turck-Chièze et al. (2008).

beam bandwidth of approximately 1 nm. The cell scattering is isotropic, and so if we assume a spherical cell we can make use of the inverse square law to calculate the power incident at the collection optics,

$$I = I_0 \frac{d^2}{D^2} \ , \tag{7.18}$$

where d is the image diameter at the centre of the cell of approximately 3.3 mm and D is the distance to the collection optics of approximately 9 mm. Finally we need to scale by the aperture area,

$$I = I_0 \frac{A_{det}^2}{A_{beam}^2} \ , \tag{7.19}$$

where A_{beam} is the beam aperture of 20 mm and A_{det} is the detector aperture of 6 mm. Combining these terms we obtain a final value of approximately 1.5 nW of light incident at the scattering detector when measuring the solar continuum. The solar potassium absorption line has an absorption depth of around 77 %, and so the final value will be between 0.35 nW and 1.4 nW depending on the line-of-sight velocity offset. We will assume an approximate value of 1 nW to be typical when measuring near the mid-point of the absorption line wings. Figure 7.1 on page 140 shows the instrument characteristics at each key point in the design.

Lens L3 in Figure 7.1 is positioned to focus quasi-parallel light, forming an image of the

fibre at the centre of the cell. The requirements are that the focal length is long enough to ensure the beam has converged sufficiently before entering the aperture in the side of the magnet, but short enough so as not to cause unnecessarily high magnification of the image. A focal length of 100 mm is ideal, and so the Thorlabs AC254-100-B-ML lens is used here. We saw earlier than the fibre collimating lens has a focal length of 30 mm, and so the system magnification of the 1 mm fibre is 3.3 times producing an image in the cell 3.3 mm in diameter. Typical BiSON cell image sizes vary by design style, but this image size is of the order expected and is acceptable.

In the next section, we will look at the magnetic field strength required to optimise the sensitivity of the instrument and maximise the line-of-sight velocity ratio.

7.7 Magnetic Field Strength

In order to determine the optimum position of the instrumental passbands, and thus calculate the ideal magnetic field strength for Zeeman splitting, it is necessary to know the shape of the solar potassium absorption line. We saw in Chapter 3 that Underhill (1993) showed BiSON data are well fitted if the solar line is described by,

$$I_\lambda = I_0 e^{-\beta_\lambda} \ , \qquad \text{(3.1 revisited)}$$

where

$$\beta_\lambda = 1.514 e^{\left[-\left(\frac{\lambda-\lambda_0}{92}\right)^2\right]} \ , \qquad \text{(3.2 revisited)}$$

and λ and λ_0 are the wavelength and central wavelength respectively in units of milli-Angstroms. The line profile is shown in Figure 3.3 on page 34. From the calculations in Chapter 3 we know that the solar line-of-sight velocity will shift between approximately $-328\,\mathrm{m\,s^{-1}}$ to $1592\,\mathrm{m\,s^{-1}}$. The optimum magnetic field strength has lower and upper bounds that can be determined by considering how the passbands change with field strength in Figure 3.3. If the field strength is too low then the two passbands will not be completely separated, and also the red component becomes situated on the wrong side of

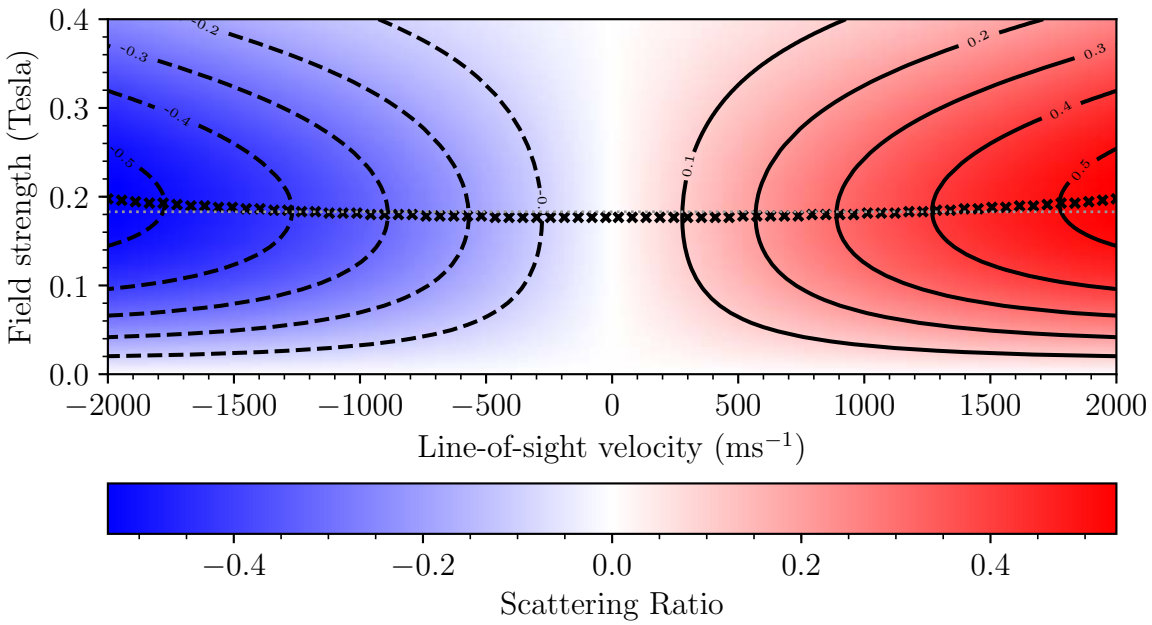

Figure 7.18: Magnetic field strength optimisation. The crosses indicate the point of maximum ratio for each velocity and therefore the ideal magnetic field strength. The horizontal grey dotted line indicates the mean ideal field strength of 0.183 T.

the solar line at the maximum velocity offset. If the field strength is too high then the blue component moves completely off the absorption line and into the continuum at maximum offset. Thus, the suitable range of magnetic fields is therefore approximately 0.11 T to 0.3 T.

Figure 7.18 shows the predicted scattering ratio over this range of velocity offsets and magnetic field strengths, drawn symmetrically for completeness. As the velocity offset increases, the ideal magnetic field strength for a given velocity offset, indicated by the crosses on Figure 7.18, naturally increases in order to maintain the passbands on the steepest part of the absorption line and so providing the maximum sensitivity. The ideal field strength varies from 0.176 T to 0.198 T, respectively from zero velocity offset to the maximum offset expected. Current BiSON spectrophotometers operate at 0.18 T. It should be noted that the sensitivity is not largely dependent on the field strength in this region, and that even at high velocity offsets a field anywhere in the range 0.15 T to 0.25 T would be acceptable. A 0.18 T longitudinal magnetic field splits the two components by approximately 0.0133 nm, or 5.2 km s^{-1} in terms of Doppler velocity.

Some BiSON so-called "2B" instruments use two vapour cells in order to probe the absorption line at four working points, and these are designed with a second magnetic field at 0.3 T. Whilst this suffers reduced sensitivity, it is the minimum field strength that completely separates the two stronger-field components from the two weaker-field components. The weaker field is produced using Swift Levick Magnets designed using Remko soft-iron pole pieces attached to an AlNiCo-V horseshoe yoke. The stronger field is produced using Samarium-cobalt-alloy (SmCo) pole pieces and a EN1A steel yoke. The SmCo is preferable since the magnets are smaller and lighter for a given field strength.

7.8 Transmission Monitor

It is useful to monitor the light that is not scattered by the cell and is transmitted directly through the spectrophotometer. This measurement allows the quality of beam polarisation to be monitored, and also to estimate the daily atmospheric extinction coefficient as we saw in Chapter 5.

Lens L4 shown schematically in Figure 7.1 on page 140 is positioned to focus light diverging from the cell into an image onto the photodiode of the transmission detector. In Section 7.5 we determined that the optical power entering the vapour cell is 0.08 mW, and this can be also be considered to be the total light exiting the cell and incident on the transmission photodiode since very little power is removed from the beam by the scattering cell. The optional quarter-wave plate and linear polariser shown before the transmission detector allows for the quality of polarisation to be monitored. When these are in place the power incident on the photodiode is reduced to 0.07 mW when the polarisers are uncrossed, and practically zero when crossed. The filters and the photodiode are all mounted with a small wedge angle in order to prevent detrimental back-reflections into the cell.

7.9 Conclusion

In this chapter we developed an end-to-end model of the entire system, from the output of the Sun, through Earth's atmosphere, through each component of the spectrophotometer, and finally onto the photodiode detectors. We determined that the scattering detectors will receive between 0.35 nW and 1.4 nW of light depending on atmospheric conditions and the line-of-sight velocity offset. The transmission monitor will receive of the order of 0.07 mW. These values will be used in the next chapter to inform the specification and design of the photodiode amplifiers.

There are two critical components in the system. The first is the variable retarder that allows electronic polarisation switching, and so wavelength selection. In this chapter we determined that a Pockels-effect variable retarder can switch polarisation states in under 1 ms, and can easily achieve switching rates of the order of 100 Hz, but in operation they require several thousand volts and custom high-voltage driver units. A liquid-crystal variable retarder operates at a maximum of ± 25 V and off-the-shelf driver units are readily available. Unfortunately they are much slower to respond, taking 16 ms to reach 90 % transition, and 27 ms to reach 99 % transition, when heated to a temperature of 50 °C. At these low slew rates we will not be able to switch polarisation states at 100 Hz, and the results found here will be used in the next chapter to assess the impact of reduced switching rates.

The second critical component is the potassium vapour cell. We used the vapour cell model from Chapter 6 to assess the expected performance of a COTS Thorlabs GC25075-K cell and GCH25-75 oven. The typical use case for these cells is in Saturated Absorption Spectroscopy (SAS), where absorption within the cell is used to provide a stable locking reference for a tunable diode laser at the centre of an atomic transition. Our use case requires the ability to detect the light scattered from the vapour cell, and unfortunately the mechanical construction of these types of cells does not permit this to be done easily. At high temperatures with a large scattering cross-section, all of the light is scattered within the first oven making it impossible to reach a detector. At lower temperatures

it is necessary to collect scattered light along the centre of the length of the cell, and this is difficult to achieve while simultaneously ensuring that detection of non-resonantly scattered light is blocked. Additionally, the coldest point on the cell occurs in the middle, between the two ovens, and so this is where the solid potassium forms exactly where we need the glass to be clear to observe scattering. Unfortunately, we have to conclude that it is unlikely to be possible to develop a resonant scattering spectrophotometer using solely off-the-shelf components. Any new instrumentation will still require a custom vapour cell and oven manufactured to the required specifications.

We will now go on to develop a new data acquisition system based around a "Raspberry Pi" single-board computer, and consider the effects of the design choices and compromises we have made by comparing results with contemporaneous output from other BiSON sites. In order to continue developing and testing the new system, we make use of an existing BiSON vapour cell and oven in the next chapter.

Chapter 8

Data Acquisition and Control

Contents

8.1 Introduction

In the previous chapter we developed the framework of components required for a solar spectrophotometer, and produced an end-to-end model of the system from the solar constant through to the expected intensity incident on the photodetectors. In this chapter we will now go on to consider data acquisition using a Supervisory Control and Data Acquisition (SCADA) control system architecture.

Existing BiSON solar observatories make use of a single computer as a centralised controller, handling all systems from dome and mount pointing, to temperature stabilisation and data readout. Communication with peripherals is handled either through direct manipulation of digital input/output lines made available through a PCI interface card, or RS232 serial ports. In order to reduce the physical footprint of the control hardware we

now move to a Distributed Control System (DCS), where the necessary control loops are handled by individual autonomous controllers distributed throughout the system. The spectrophotometer data acquisition is handled by a Raspberry Pi single-board computer, with a second Pi controlling mount pointing (discussed later in Chapter 9 starting on page 229). Supervisory control, telemetry transport, and event logging is over Ethernet, or WiFi, using SSL (Secure Sockets Layer) encryption and authentication.

We begin in Section 8.2 by considering existing photodiode amplifier design, and in Section 8.3 assess through circuit simulation the performance in terms of noise, bandwidth, and stability. In Section 8.4 we go on to build a new photodetector based on the optical constraints determined in Chapter 7, paying careful attention to best practice circuit design, PCB layout, and mechanical shielding. Section 8.5 discusses input filtering and analogue-to-digital conversion. Section 8.6 discusses control of the instrumental polarisation state and rate of switching. Section 8.7 considers temperature control through proportional/integral/derivative (PID) control loop feedback techniques. Finally in Section 8.8 we consider initial results from the system following installation and commissioning at Izaña, Tenerife, and in Section 8.9 the results from a full summer observing campaign throughout 2018. Some background theory supporting the work on transimpedance amplifier design in this chapter is available in Appendix B.

8.2 The Spectrophotometer-J (Jabba) Detector

The typical amplifier circuit used in a BiSON resonantly scattered light detector was designed in 1996 (Lines and Miller, 1996) and installed in Spectrophotometer-K, codenamed Klaus, at Mount Wilson Observatory. The same circuit was subsequently used, with some revisions, in the new detectors for Spectrophotometer-J, codenamed Jabba, during the 2006–2009 refit (Barnes, 2011).

The circuit uses a transimpedance amplifier for initial photocurrent-to-voltage conversion, followed by a second voltage amplifier for additional gain. These two stages use a

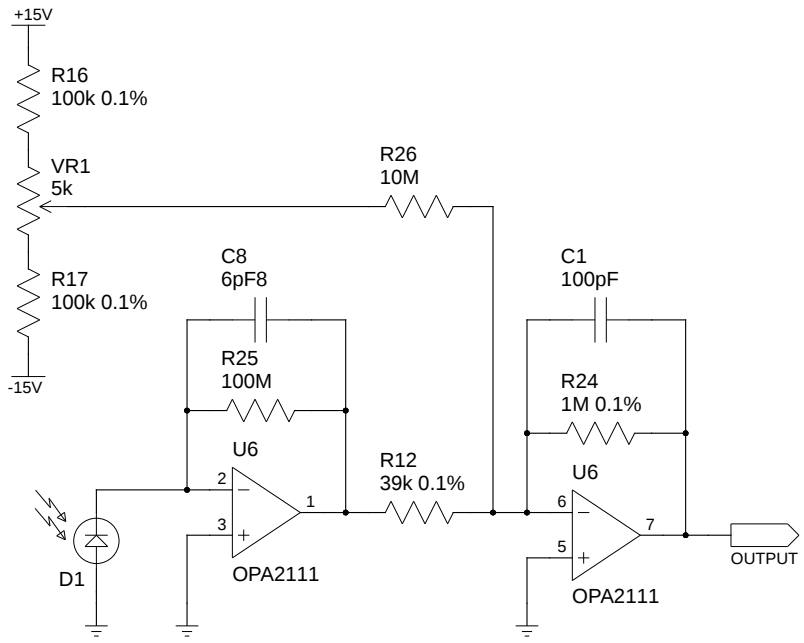

Figure 8.1: The JDet transimpedance amplifier.

high-precision low-noise Burr-Brown OPA2111KP op-amp (Burr-Brown, 1984). The amplifier circuit is shown in Figure 8.1. The signal is then passed through a 4-pole low-pass Bessel filter, which is intended to remove high-frequency noise and also provide a little more gain to boost the output dynamic range to between 0 V to 10 V. The filter stages use a lower quality, less expensive, TL082 op-amp since at this stage the signal has already been significantly amplified and is more robust. The Bessel filter circuit is shown in Figure 8.2. The amplified signal is subsequently fed into a voltage-to-frequency converter, and the output pulses sent to a set of counters for acquisition and post-processing.

Since the original design was produced, no quantitative performance assessment has ever been made. The circuit has been copied from detector to detector with only minor modifications. In the following sections we will develop a photodiode model, and use it to analyse the noise levels, frequency response, slew rate, and stability of the photodiode amplifier circuit we saw in Figures 8.1 and 8.2.

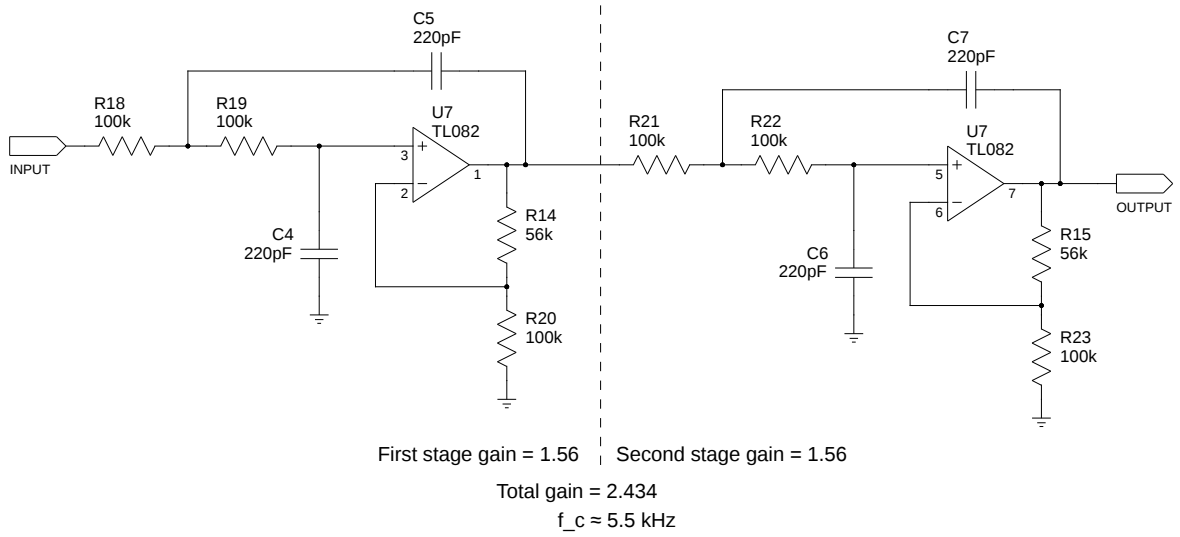

Figure 8.2: The JDet low-pass Bessel filter.

8.3 A Photodiode Model

A photodiode can be broken down into an equivalent circuit of several discrete components (Graeme, 1996). In Figure 8.3, I_S is the current generated due to the incident flux on the diode. The leakage current, I_L, is more generally known as the *dark current* when discussing photosensitive devices, and this current flows even when there is no incident flux on the device so creating a temperature dependent offset. In addition to dark current, a third current source, I_N, models the random noise intrinsic to the diode. This noise is typically quoted in terms of the *noise equivalent power* (NEP) in units of $\mathrm{W}\sqrt{\mathrm{Hz}}^{-1}$, i.e., the incident flux required to produce a signal-to-noise ratio of one with a bandwidth of 1 Hz. The NEP is calculated from the noise spectral density in $\mathrm{A}\sqrt{\mathrm{Hz}}^{-1}$ divided by the responsivity of the diode in $\mathrm{A\,W^{-1}}$. Since the responsivity is wavelength dependent, the NEP is always quoted at a particular wavelength. The junction capacitance, C_D, is determined by the depletion layer depth and any bias voltage applied to the diode. It is generally high at zero bias and reduces as the reverse bias voltage is increased. The two resistance values, R_D and R_S, are properties determined by the manufacturing process and can vary significantly between batches of any given photodiode. The diode shunt resistance, R_D, is usually tens or hundreds of MΩ, and is an inverse exponential function

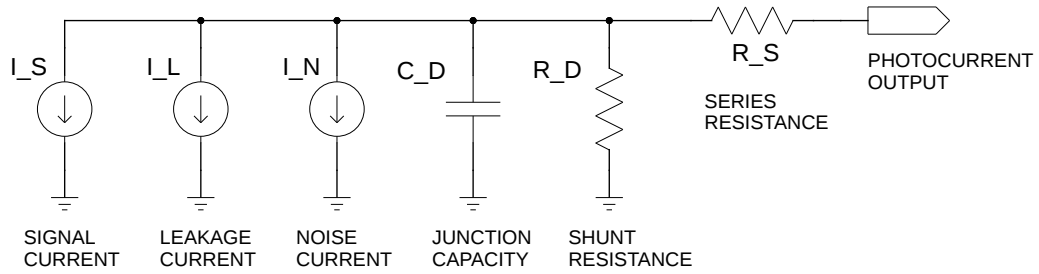

Figure 8.3: A photodiode can be modelled from a set of discrete components.

of the forward voltage across the diode. The series resistance, R_S, is usually low, a few Ω, and is determined by such things as the type of case terminals. A selection of common photodiode electrical and optical specifications is shown in Table 8.1.

If a load is placed across the output of the photodiode and ground, then the output voltage, V_{out}, can be determined from Ohm's law to be,

$$V_{out} = (I_S + I_L + I_N)\left(\frac{R_L R_D}{R_L + R_D + R_S}\right) \ , \tag{8.1}$$

where R_L is the load resistance. However in the case where R_L is much larger than the diode shunt resistance the photocurrent that is generated will flow through R_D. This creates a forward bias across the diode and so the value of R_D drops exponentially as the incident flux increases, causing problems not only because the voltage output is now a logarithmic function of incident flux rather than a linear response, but mainly due to the wide variation in R_D between production batches. Better results are obtained by ensuring R_L is very much less than R_D, causing the photocurrent to flow through the fixed R_L. The simplest way to achieve a low load resistance is to connect the photocurrent to an operational amplifier, forming a *transimpedance amplifier* (TIA), shown in Figure 8.4(a). The feedback resistor R_F sets the gain of the amplifier, and the capacitor C_F is added for stability. This configuration has a very linear response and also very low noise, since the dark current is almost completely eliminated. However, depending on the properties of the diode chosen and the gain required this circuit can be slow to respond and may not meet the bandwidth requirements of the system. One way to improve the speed of response

Table 8.1: Common photodiode electrical and optical specifications. Reproduced from Centronic (2017). Characteristics are measured at $22 \pm 2\,°C$ ambient. Shunt resistance is measured at 10 mV.

Type	Active Area		Responsivity A W^{-1} $\lambda = 436$ nm		Dark Current nA		NEP W $\sqrt{\text{Hz}}^{-1}$ $\lambda = 436$ nm	Capacitance pF		Shunt Resistance MΩ		Risetime ns $\lambda = 820$ nm
								$V_r = 0$	$V_r = 12$			$R_L = 50\,\Omega$
	mm^2	mm	Min	Typ	Max	Typ	Typ	Max	Max	Min	Typ	Typ
OSD1-5T	1	1.13 ⊘	0.18	0.21	1	0.2	2.5e-14	35	7	250	1000	7
OSD5-5T	5	2.52 ⊘	0.18	0.21	2	0.5	3.3e-14	130	35	100	600	9
OSD7.5-5T	7.5	2.75	0.18	0.21	3	1	4.6e-14	180	40	60	300	10
OSD15-5T	15	3.8	0.18	0.21	5	1	5.5e-14	390	80	50	200	12
OSD35-5T	35	5.9	0.18	0.21	10	2	7.5e-14	950	200	20	100	20
OSD50-5T	50	7.98 ⊘	0.18	0.21	15	5	1.6e-13	1300	270	5	25	26
OSD60-5T	62	7.9	0.18	0.21	25	6	2.3e-13	1800	310	3	12	30
OSD100-5T	100	11.3 ⊘	0.18	0.21	30	8	2.1e-13	2500	520	2	15	45
OSD300-5T	300	19.54 ⊘	0.18	0.21	200	30	3.5e-13	7500	1500	1	5	125

is to reverse bias the diode, i.e., the cathode is made positive and the anode negative, shown in Figure 8.4(b). This configuration reduces the capacitance of the photodiode due to an increase in the width of the depletion region and so improves the speed of response. Unfortunately the dark current is no longer eliminated resulting in poorer noise performance than the photovoltaic mode of operation.

From the information above it is clear that best results are obtained when using a small photodiode, a bright source, and a low bandwidth requirement. Since the number of component variables and performance metrics is high, the best way to optimise a potential circuit for a specific application is to run simulations of the proposed specification. Over the next few subsections, we will model and simulate the existing detector circuit performance, and then go on to develop a new detector optimised for both high-gain resonance scattering sensitivity, and low-gain transmitted light detection.

When modelling the photodiode response using the equivalent circuit shown in Figure 8.3, the model has been simplified by ignoring the leakage current and noise current. This is acceptable since we make use of the low-noise photovoltaic mode of amplification where we have shown the leakage current to be almost completely eliminated. We can also ignore the photodiode intrinsic noise current since in a high-gain environment this noise is insignificant in comparison to that from the rest of the system. The series re-

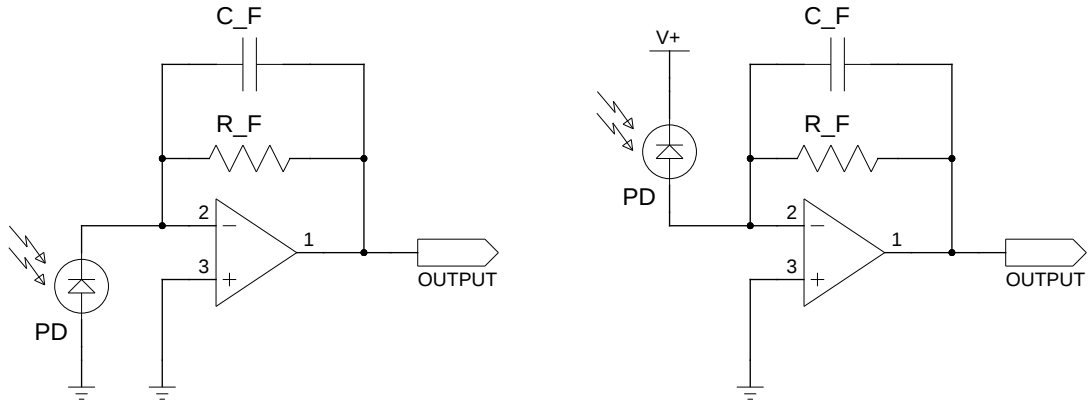

(a) Photovoltaic mode amplifier. (b) Photoconductive mode amplifier.

Figure 8.4: (a) shows a transimpedance amplifier providing current-to-voltage conversion in photovoltaic mode. (b) shows a similar configuration with the photodiode reversed biased in photoconductive mode. The feedback resistor R_F sets the gain of the amplifier, and the capacitor C_F is added for stability.

sistance is generally not stated by any manufacturers, and so this has been assumed to be low at 10 Ω. The final photodiode model is that of a signal current source, a parallel capacitance, a parallel shunt resistance, and a series resistance. We can use this model along with SPICE (Simulation Program with Integrated Circuit Emphasis) simulation to analyse the performance of the Jabba detectors shown in Figures 8.1 and 8.2. SPICE is the industry standard electrical simulation tool, an open-source package developed by Laurence Nagel and Donald Pederson at the Electronics Research Laboratory of the University of California, Berkeley, in 1973 (Nagel and Pederson, 1973; Nagel, 1975). In the following subsections we will use SPICE to determine noise levels, frequency response, slew rate, and finally amplifier stability.

8.3.1 Noise Performance

We first look at the most simple type of noise on passive components. In the same way that an atomic absorption line in a vapour is broadened by thermal motion of the atoms, electronic noise is produced in a conductor due to thermal motion and agitation of electrons and occurs at equilibrium regardless of any applied voltage. This type of noise

was first measured by John Johnson, later explained by Harry Nyquist, and has become known as Johnson-Nyquist noise.

The Johnson-Nyquist thermal noise in an ideal resistor is approximately white, meaning it is spectrally flat with a uniform power density, and has an RMS amplitude given by (Johnson, 1928; Nyquist, 1928),

$$\text{Noise} = \sqrt{4k_{\text{B}}TRB}\ , \tag{8.2}$$

where T is the temperature of the resistor in Kelvin, k_{B} the Boltzmann constant, R the resistance, and B the bandwidth in Hz. Since the bandwidth is generally unspecified, values are usually stated in terms of the amplitude spectral density in $\text{V}\sqrt{\text{Hz}}^{-1}$. Thus for a $1\,\text{k}\Omega$ resistor at room temperature, the value for Johnson noise is given by,

$$\text{Noise} = \sqrt{4\ .\ 1.38 \times 10^{-23}\ .\ 293\ .\ 1000}\ \sqrt{B}\ , \tag{8.3}$$

which is equal to $4\,\text{nV}\sqrt{\text{Hz}}^{-1}$. Similarly a $100\,\text{M}\Omega$ resistor produces Johnson noise of $1.27\,\mu\text{V}\sqrt{\text{Hz}}^{-1}$. It is important to note that even when cooled with liquid nitrogen to $77\,\text{K}$ this value drops only slightly to $0.65\,\mu\text{V}\sqrt{\text{Hz}}^{-1}$. When using high value feedback resistors, Johnson noise remains significant even at very low temperatures and is usually the dominating noise in high-gain systems.

Over a $5\,\text{kHz}$ bandwidth at room temperature for a $100\,\text{M}\Omega$ gain resistor this produces $0.09\,\text{mV}_{\text{RMS}}$ of thermal noise. The circuit in Figures 8.1 and 8.2 have a combined gain of 62.4 after the initial current-to-voltage conversion, and so this means we can expect to obtain around $5.6\,\text{mV}_{\text{RMS}}$ of noise just from Johnson noise alone. In a transimpedance amplifier the gain increases linearly with R, and so even though the output noise voltage increases proportionally to $\sqrt{R}$ better noise performance is achieved through higher gain feedback resistance. The other important source of noise in the circuit is produced by the operational amplifiers. We need to consider input bias current, input noise voltage density, and also input noise current density.

An ideal op-amp has infinite voltage gain, infinite input impedance, zero output

impedance, and infinite bandwidth. This leads to the well-known *golden rules* for op-amps (Horowitz and Hill, 2015), which state that for an op-amp with external feedback the output will attempt to do whatever is necessary to produce zero voltage difference between the inputs, and that the inputs draw zero current. Real life op-amps do not meet this ideal, although some get very close. Actual values can be found per-device by reading the relevant product datasheet. The input bias current is the current which actually flows at the inputs, and for an OPA2111KP has a typical value listed as 3 pA (Burr-Brown, 1984). This may sound small but it has to be processed through the gain of the system and so the resulting voltage offset can become quite significant. The input bias current is also temperature dependent, and so it is important to select a device with a low temperature coefficient where long term stability is crucial. Is it very easy for a sensitive circuit to be turned into a thermometer due to a poor temperature coefficient. The input voltage noise is most significant at low frequency, below around 10 kHz, and so has become known as $1/f$ noise, *flicker noise*, or *pink noise*. Input current noise is a measure of the white shot noise. Although independent of temperature, this too passes through the gain of the system and so can become significant on the output. In a system that is amplifying pA of photocurrent, an amplifier with input current noise of the order of fA is required. The OPA2111KP has a typical current noise of $0.8\,\mathrm{fA}\,\sqrt{\mathrm{Hz}}^{-1}$. The 3 dB point at which the $1/f$ noise starts to rise out of the white noise with decreasing frequency is known as the *corner frequency* and for systems operating at low frequency it is important to select a device with as low a corner frequency as possible.

All of these uncorrelated noise sources have to be processed through the system gains, added in quadrature at the output, while also taking into account the bandwidth of the system. For any even mildly complex circuit this is very difficult to do manually, and so it is appropriate to take advantage of computational simulation. Having calculated a ball-park noise figure, we now go on to simulate the response of the more complicated non-linear components, add in noise from all the other large resistors in the circuit, and also include the modelled photodiode response.

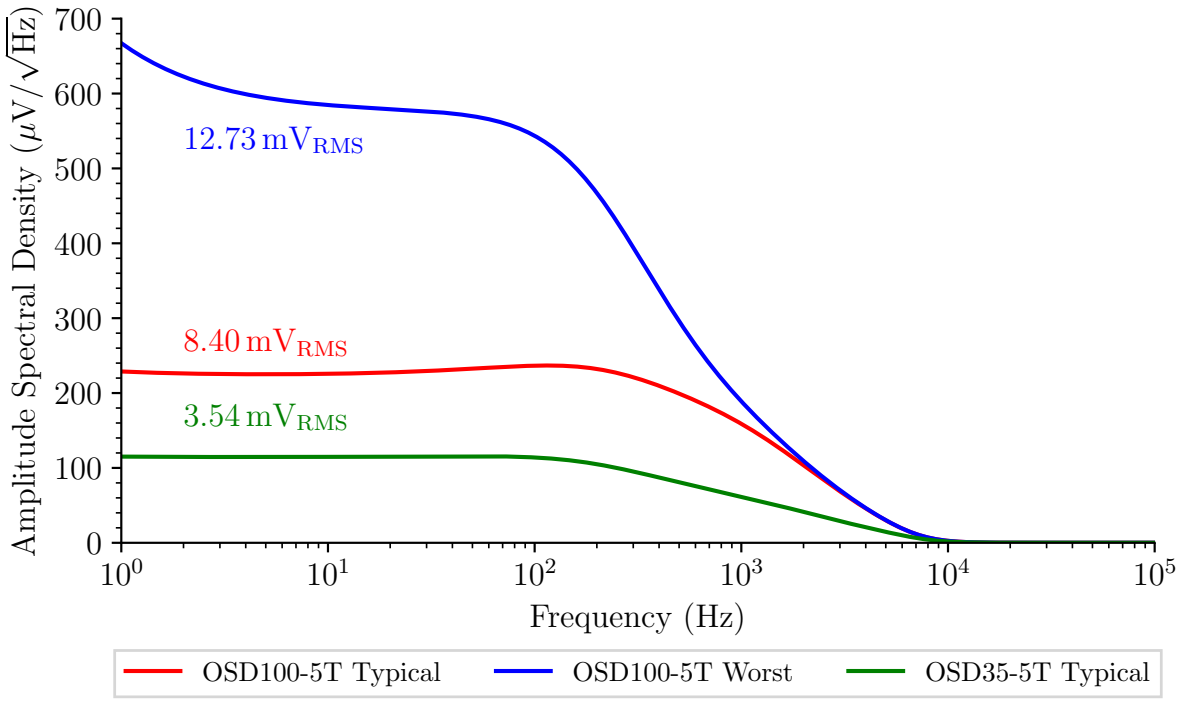

Figure 8.5: Jabba scattering detector noise analysis. Typical and worst-case performance is shown for the OSD100-5T photodiode based on the datasheet specification, and also for comparison the typical performance of a diode with a sensor of half the diameter.

SPICE models for the OPA2111KP and TL082 op-amps were downloaded from their respective manufacturers websites, and the circuit in Figures 8.1 and 8.2 set up for simulation. Figure 8.5 shows the results of the noise analysis for three different photodiodes. These are the OSD100-5T, the standard photodiode used in BiSON spectrophotometers, with parameters taken from the datasheet for typical and worst-case performance (Centronic, 2017). For comparison a simulation was also performed for a smaller diode, the OSD35-5T, to show what improvement could be expected if a smaller sensitive area were acceptable. The three results are all very close to our simple calculation for a single gain resistor, and shows that the system is indeed Johnson noise dominated, with variation in performance due to the specification of the diode. Predominantly the difference is due to the change in diode shunt resistance. As the parallel shunt resistance increases the total resistance decreases in terms of noise, and so the noise level also reduces.

8.3.2 Frequency Response

Since the solar oscillation signal has a frequency envelope of the order of a few mHz one might consider that our bandwidth requirements are insignificant. However the instrumentation multiplexes different polarisation states through one photodetector and so it is necessary to pass transient changes in incident flux with the minimum amount of dead time between transitions. This significantly increases the bandwidth requirements of the system, and means we have to consider both analogue and high-speed digital design techniques.

A transient can be considered as an ideal square wave, which can be represented as a sum of the odd-integer harmonics of a fundamental frequency using Fourier expansion,

$$\begin{aligned} A(t) &= \frac{4}{\pi} \sum_{k=1}^{\infty} \frac{\sin\left(2\pi(2k-1)f_0 t\right)}{2k-1} \ , \\ &= \frac{4}{\pi} \left[\sin\left(2\pi f_0 t\right) + \frac{1}{3}\sin\left(6\pi f_0 t\right) + \frac{1}{5}\sin\left(10\pi f_0 t\right) + \ldots \right] \ , \end{aligned} \tag{8.4}$$

where f_0 is the fundamental frequency, k is an odd-integer harmonic, and t is time. Figure 8.6 shows a 100 Hz fundamental sine wave, followed by summing a total of 2, 5, and 10 waveforms. At 10 waveforms (the fundamental and 9 odd-integer harmonics) the maximum frequency $f_{\max}$ is 1.9 kHz. If we continued to 25 waveforms (the fundamental and 24 odd-integer harmonics) the maximum frequency would be 4.9 kHz. Ideally a square wave changes between low and high states instantaneously. In real physical systems such as electronic circuits this is impossible since it requires equation 8.4 to be summed to infinity, which would require infinite bandwidth. All physical systems act as a low-pass filter, with the cutoff frequency determining the bandwidth and so how much smoothing is applied to the ideal square wave. The response of a circuit at high frequencies affects its processing of short-time events. Johnson and Graham (1993, p. 5) describe a rule of thumb for the maximum practical frequency component as,

$$f_{\text{knee}} = \frac{0.5}{t_r} \ , \tag{8.5}$$

where t_r is the desired rise time of a 10 % to 90 % transition, and f_{knee} the inflection point

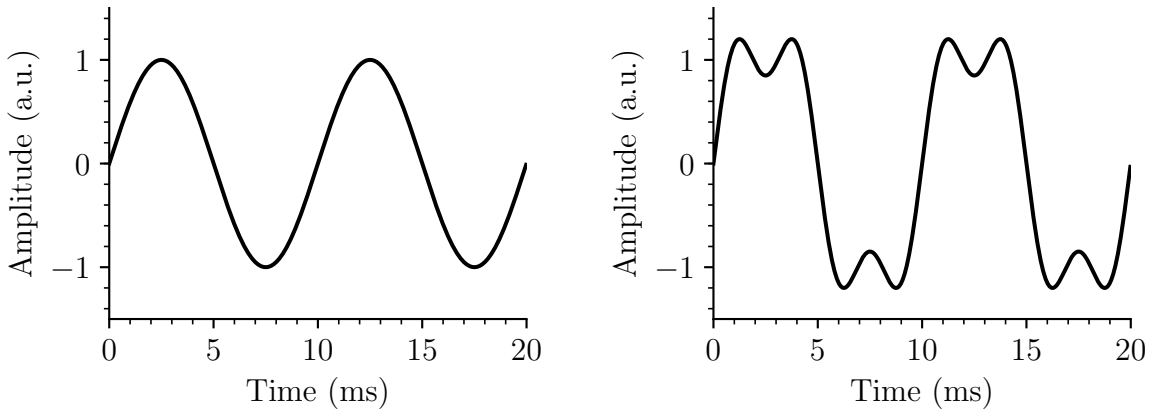

(a) nWaveforms = 1 $f_{\max} = 100\,\text{Hz}$

(b) nWaveforms = 2 $f_{\max} = 300\,\text{Hz}$

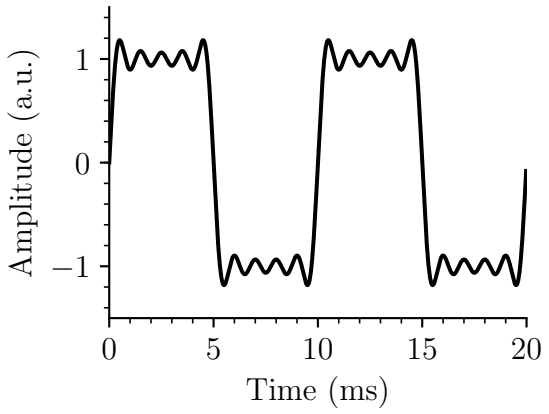

(c) nWaveforms = 5 $f_{\max} = 900\,\text{Hz}$

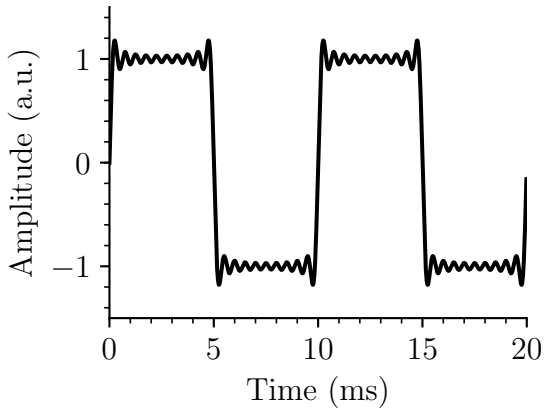

(d) nWaveforms = 10 $f_{\max} = 1.9\,\text{kHz}$

Figure 8.6: A square wave can be considered to be built up from a collection of sine wave harmonics.

or "knee" in the frequency spectrum above which frequency components are insignificant in the determining the shape of the signal. If we require a rise time of 0.5 ms to match the slew rate of a Pockels cell, then this suggests a bandwidth requirement of 1 kHz.

The oscillations around the transitions in Figure 8.6 as the waveform becomes square are known as *ringing* and are due to the Gibbs phenomenon, which describes how a Fourier series behaves at a discontinuity. It is important to note that the overshoot does not decay as the number of harmonics increases but only achieves a limit. The time domain representation of a perfect low-pass filter is a sinc function, and the ringing is a manifestation of the impulse response of this bandwidth limitation. In terms of

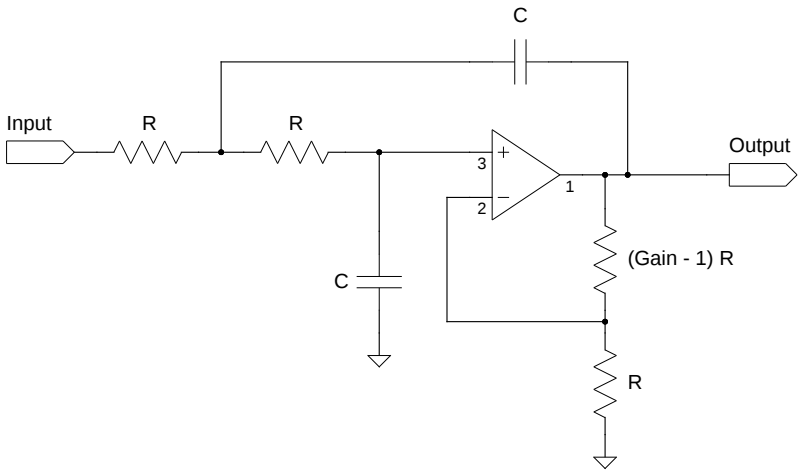

Figure 8.7: Generic Sallen-Key topology voltage-controlled voltage-source (VCVS) low-pass filter schematic.

signal processing, while the time domain oscillations are known as ringing, frequency domain oscillations are known as *ripple*. These artefacts are undesirable, but they can be reduced or eliminated by selecting the most appropriate low-pass filter configuration for the circuit requirements. There are several types of low-pass filter, but they all use the same basic filter design, known as the Sallen-Key topology (Sallen and Key, 1955), shown in Figure 8.7. Appropriate values for R and C can be calculated using (Horowitz and Hill, 2015, p. 408),

$$\mathrm{RC} = \frac{1}{2\pi c_\mathrm{n} f_\mathrm{c}} \ , \tag{8.6}$$

where c_n is a normalising factor for each cascaded filter stage, and f_c is the $-3\,\mathrm{dB}$ cutoff frequency. Example values for c_n and K are shown in Table 8.2 reproduced from Horowitz and Hill (2015, p. 408) for three filter types and four different pole configurations. Typically components would be specified using filter-design software rather than a look-up table. The Butterworth filter has a "maximally flat" passband response, but compromises with a poor rate of decay from the passband to the stopband. If the system specification can allow some ripple in the passband then the Chebyshev filter can be used to increase the sharpness of f_knee. The Bessel filter optimises for linear phase-shift rather than amplitude or frequency response, and so produces equal time delays for all signals within

Table 8.2: Voltage-controlled voltage-source (VCVS) lowpass filters coefficients. Reproduced from Horowitz and Hill (2015, p. 408).

Poles	**Butterworth**	**Bessel**		**Chebyshev** (0.5 dB)		**Chebyshev** (2 dB)	
	K	c_n	K	c_n	K	c_n	K
2	1.586	1.272	1.268	1.231	1.842	0.907	2.114
4	1.152	1.432	1.084	0.597	1.582	0.471	1.924
	2.235	1.606	1.759	1.031	2.660	0.964	2.782
6	1.068	1.607	1.040	0.396	1.537	0.316	1.891
	1.586	1.692	1.364	0.768	2.448	0.730	2.648
	2.483	1.908	2.023	1.011	2.846	0.983	2.904
8	1.038	1.781	1.024	0.297	1.522	0.238	1.879
	1.337	1.835	1.213	0.599	2.379	0.572	2.605
	1.889	1.956	1.593	0.861	2.711	0.842	2.821
	2.610	2.192	2.184	1.006	2.913	0.990	2.946

the passband and minimises distortion of the waveform. The Bessel filter can be said to have maximally flat time delay, and this is the reason the Jabba detector makes use of a Bessel-type low-pass filter.

The SPICE simulation of frequency response is shown in Figure 8.8 for both the amplifier-stage shown in Figure 8.1, and the filter-stage shown in Figure 8.2 on page 182. We can see that due to the very high gain of the initial transimpedance amplifier the bandwidth is just 239 Hz, and so from equation 8.5 we can expect the rise time to be longer than 2 ms for a 10 % to 90 % transition. The bandwidth of the Bessel low-pass filter is actually higher at 5.6 kHz than the bandwidth of the first stage amplifier, and so the filter is doing nothing except providing additional gain. It is possible that the filter was designed before the required gain and bandwidth were known, and then never updated. It is important to specify the output bandwidth when amplifying a signal with a narrowband spectrum in order to maximise the signal-to-noise ratio, otherwise the amplifier will keep adding noise power while the signal power remains constant. Since we know the total noise is dominated by the bandwidth-limited first stage transimpedance amplifier, we can

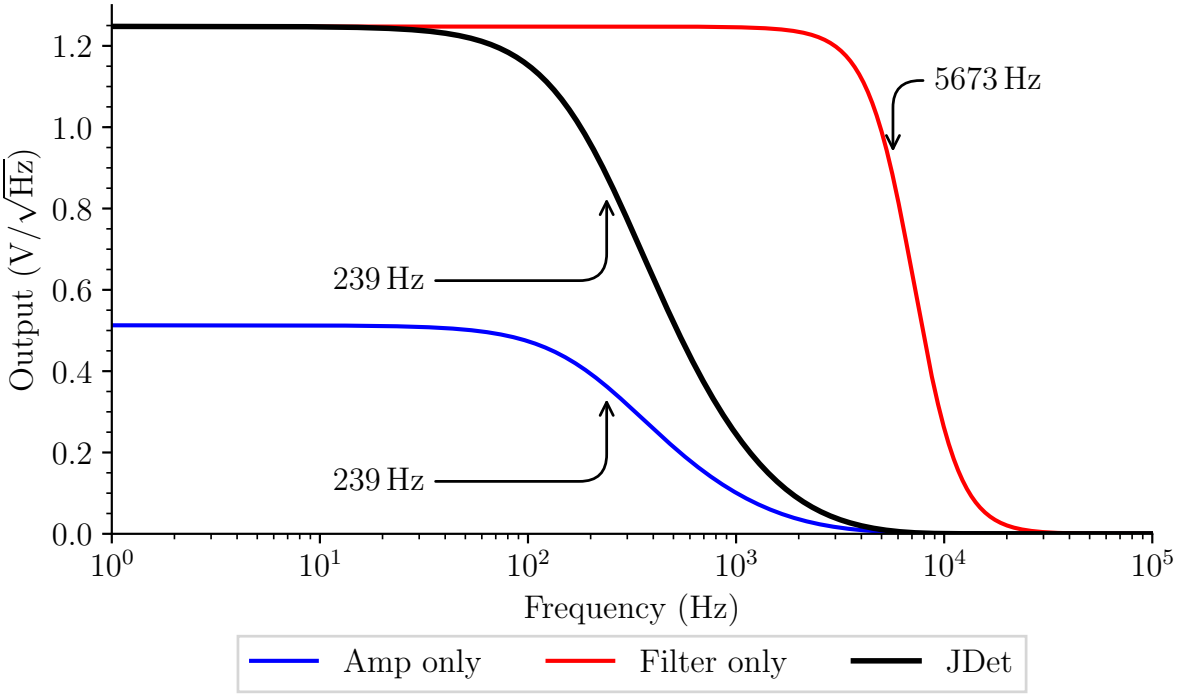

Figure 8.8: Jabba scattering detector frequency response analysis. The response of the initial twin-amplifier stage, and the low-pass filter stage, are shown separately. The heavier weight black line shows the response of the whole detector. The input signal is 200 pA of alternating photocurrent. The arrows indicate the −3 dB point for each response curve, and equation 8.5 can be used to estimate the time of a 10 % to 90 % transition from these values.

simplify the detector design by removing the filter completely and increasing the gain of the second stage amplifier, eliminating more than half of the components.

8.3.3 Slew Rate and Transient Response

Figure 8.9 shows the SPICE transient response analysis. The simulated photocurrent is stepped from 25 pA to 500 pA, resulting in rise time of 2.5 ms to 4 ms depending on the acceptable percent of completion. Current BiSON data acquisition systems are configured for a settling time delay of just 0.5 ms since detector bandwidth has never been adequately considered, and so this means that some polarisation state mixing is reducing the instrumental sensitivity.

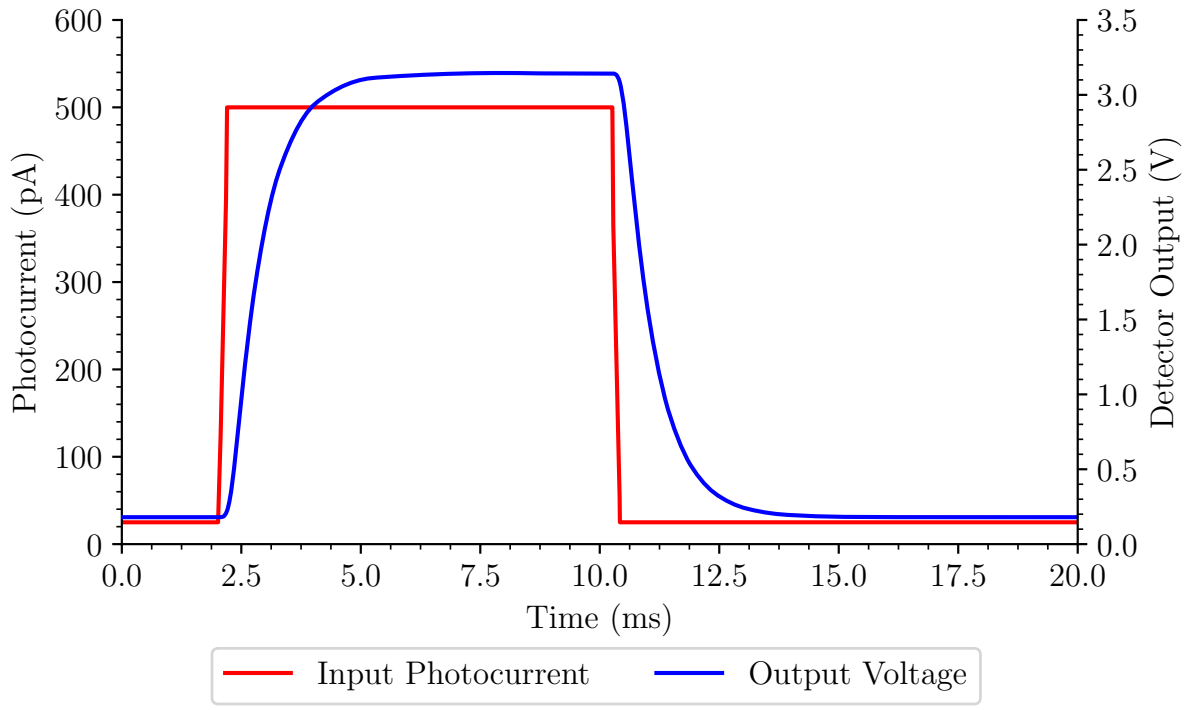

Figure 8.9: Jabba scattering detector transient response analysis. The transient response of the system is shown when stepping between 25 pA of photocurrent and 500 pA of photocurrent.

8.3.4 Stability

Figure 8.10 shows the SPICE Bode plot of the first and second stage amplifier. The phase margin of both amplifiers is greater than 45° and so can be considered stable and reliable.

In the next section we will develop a new miniature transimpedance amplifier with a minimal number of components, and compare the performance with the Jabba detector.

8.4 A New Photodetector

8.4.1 Amplifier Gain and Stability

In Chapter 7 we developed an end-to-end model of the entire system, and we determined that the scattering detectors will receive between 0.35 nW and 1.4 nW of light depending on the line-of-sight velocity offset. The transmission monitor will receive approximately 0.07 mW. We also know that there is a high bandwidth requirement due to the

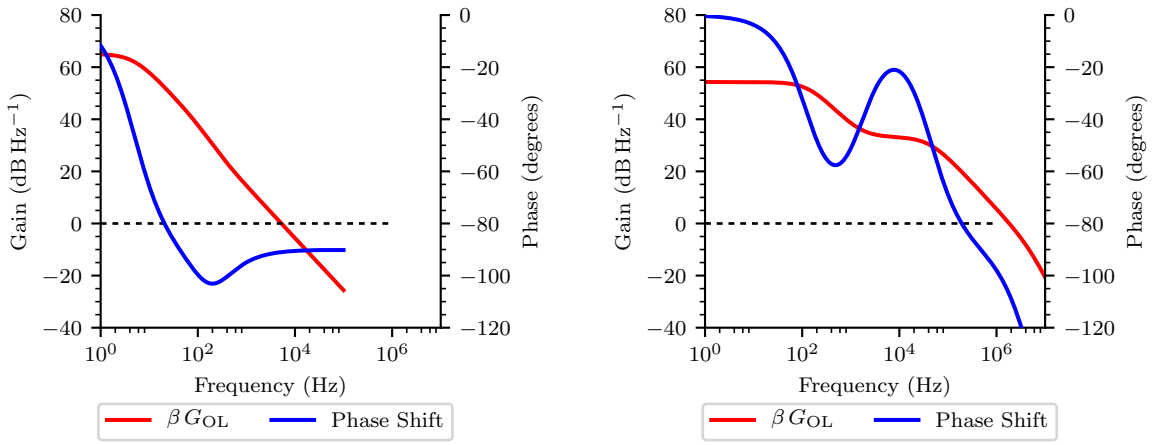

(a) Transimpedance amplification.

(b) Voltage amplification.

Figure 8.10: Jabba scattering detector stability analysis.

need to process intensity transients during changes of polarisation state. Unfortunately we cannot purchase an off-the-shelf photodiode amplifier suitable for these requirements. For example, the Thorlabs PDA200C bench-top photodiode amplifier has a minimum measurement range suitable for up to 100 nA, two orders of magnitude larger than our signal, and also costs £800 for only one channel of amplification. By designing a custom amplifier we can optimise our signal-to-noise ratio and reduce the component cost to a tenth that of a commercial unit.

The Centronic OSD100-5T photodiode used in most BiSON instrumentation has responsivity at 770 nm of approximately 0.45 A W^{-1} (Centronic, 2017). We can expect the photodiodes to be producing between 160 pA to 630 pA. The existing scattering amplifier used in BiSON instrumentation, as we saw in Figure 8.1 on page 181, makes use of 100 MΩ transimpedance gain resistance (R8), followed by 62.4 times voltage amplification providing an output between 0 V to 10 V into a VFC32 voltage-to-frequency converter. At this gain a photocurrent of 400 pA gives an output of approximately 2.5 V. Similar gains are used in the new amplifier, shown in Figure 8.11, with a 100 MΩ transimpedance gain resistance, followed by 30.9 times voltage amplification. The lower second-stage amplification is due to the reduced 0 V to 5 V dynamic-range input of the ADS1210 analogue-to-digital converter. The transimpedance compensating capacitor (C7) is set at 6.8 pF. This value

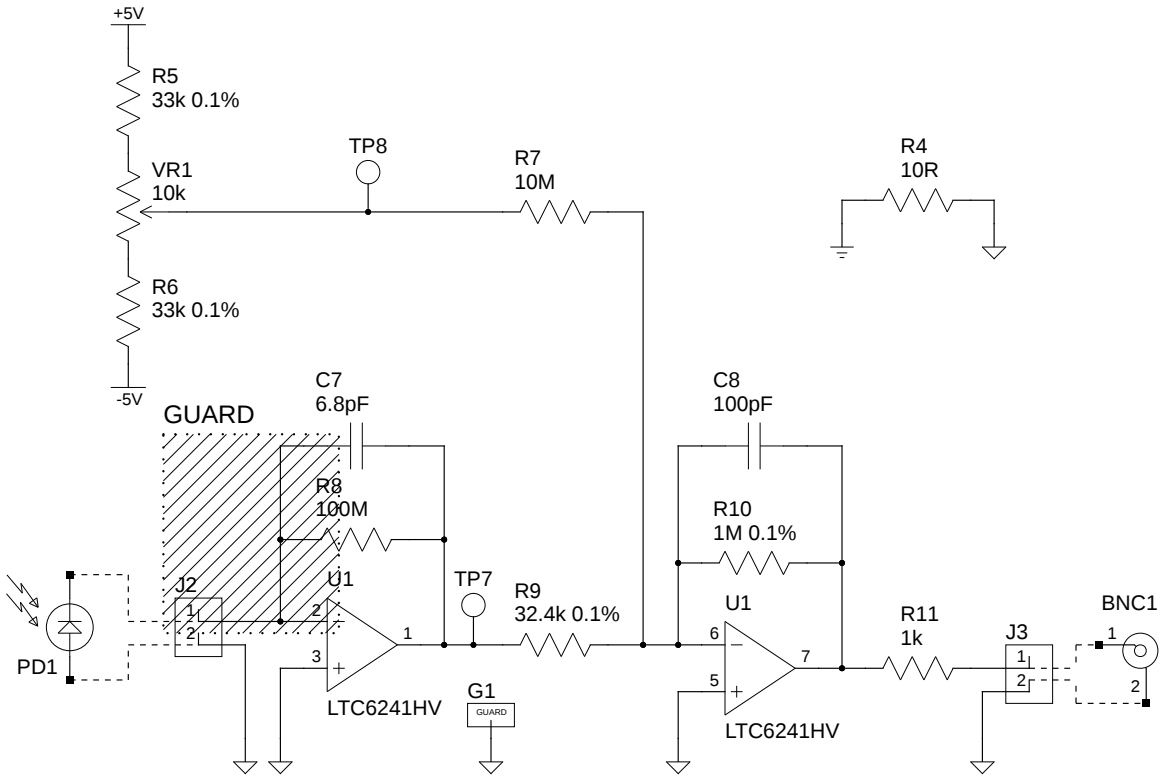

Figure 8.11: Photodiode transimpedance amplifier schematic.

could be made slightly lower to give a small bandwidth improvement at the expense of some gain peaking. Bandwidth could also be improved by reducing the transimpedance gain and compensating by increasing the second-stage voltage gain, at the expense of an increased noise level. The transmission detector requires a much lower gain and uses just 30 kΩ transimpedance resistance followed by a unity inverting amplifier producing a typical 1 V output, and can be considered to be essentially noiseless in comparison to the scattering detectors.

The Linear Technology LTC6241HV (Linear, 2005) was selected for both amplification stages due to its low noise characteristics. The input bias current typical value is just 0.2 pA, compared with the OPA2111KP (Burr-Brown, 1984) typical value of 3 pA. The input noise is 0.56 fA $\sqrt{\text{Hz}}^{-1}$, down from 0.8 fA $\sqrt{\text{Hz}}^{-1}$. Both devices have similar low drift temperature coefficients, 2.5 µV °C^{-1} for the LTC6241HV down from 2.8 µV °C^{-1} on the OPA2111KP. We use the "HV" version of the LTC6241 since this allows bipolar ±5 V

operation. The LTC6241 is a rail-to-rail device, meaning that the output stage can swing within 30 mV of either supply rail to maximise the signal dynamic range.

Both the scattering and transmission detectors are designed with a phase margin greater than 45° and so can be considered to be stable and reliable under all conditions.

8.4.2 Grounding, Decoupling, Bypassing, and Guarding

Although the amplifier schematic itself is relatively simple, great care must be taken in the design and construction to ensure the specified performance levels are reached. Unlike the thermal noise and amplifier noise we discussed in Section 8.3, many sources of noise can be reduced to an insignificant level with proper circuit design, PCB layout, and mechanical construction. Likewise they can also destroy the performance of an expensive instrumentation amplifier without careful layout and grounding, plus appropriate electrostatic and magnetic shielding.

The supporting components for the system in Figure 8.11 are shown in Figure 8.12. The incoming supply is an unregulated ±12 V, and these power rails are first passed through ±5 V regulators to ensure a smooth clean supply to the amplifiers. The regulated power rails are bypassed with 1 µF capacitors as specified by their respective datasheets in order to short high-frequency noise direct to ground. In addition, the supply voltages to the low-level amplifiers are decoupled from the rest of the circuit using a small RC network (Horowitz and Hill, 2015, p. 583).

A total of three different grounds are defined in the circuit. These are power, signal, and chassis. A common problem with grounding is forgetting about currents that are flowing through a ground trace can appear as a signal in another part of the circuit (Horowitz and Hill, 2015, p. 582). The important consideration is to look at where large currents are flowing and ensure their voltage drops do not connect to an input, such as the supply or reference point of an op-amp. Whilst this can be done at the PCB design stage, it is useful to force good grounding techniques by defining different ground symbols that are all tied to one common ground at a single point. By working in this way we ensure,

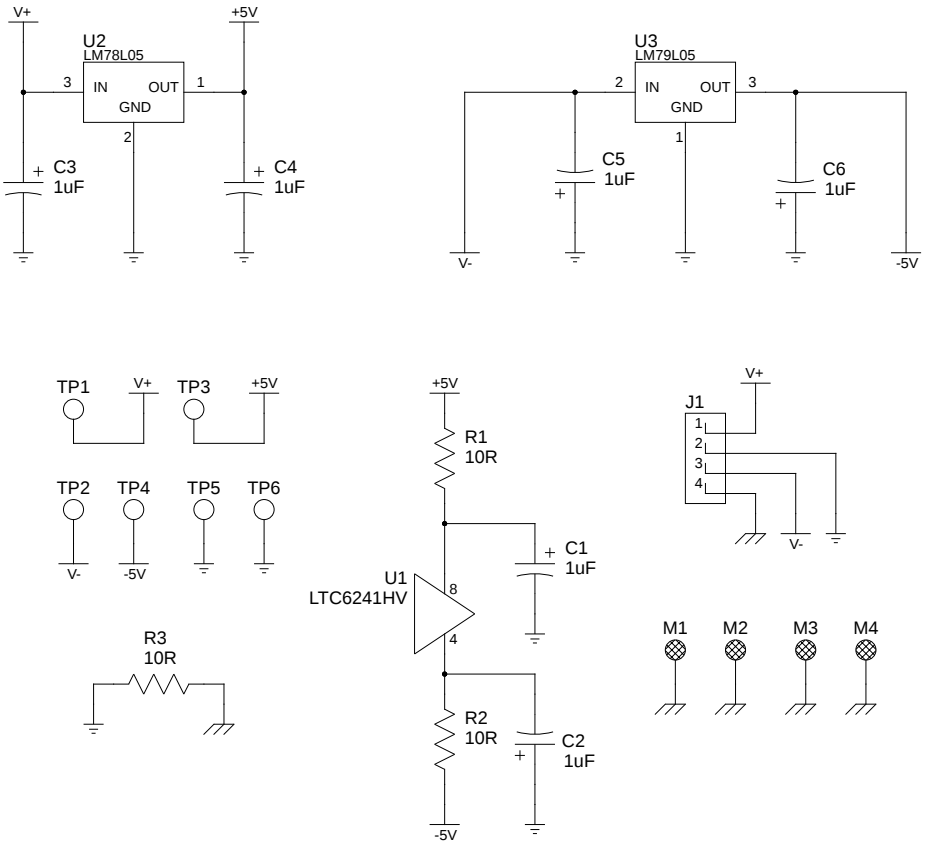

Figure 8.12: Photodiode amplifier power supply schematic.

for example, that noise filtered by the power regulator bypass capacitors is returned to power-ground and not signal-ground where it would likely be connected directly to the grounded input of an op-amp, which is precisely where we do not want the noise to go. A flood-filled copper ground plane is applied to both sides of the PCB in order to reduce capacitive coupling between traces and components – the ground plane "swallows" the electrostatic fringing fields (Horowitz and Hill, 2015, p. 581). The ground plane is cutout underneath the op-amp on both sides, since capacitance here between the surfaces of the PCB can cause the op-amp to oscillate at high frequencies. This would not be a problem in our relatively low-frequency use case, but we follow best practice.

One particularly important area of concern is the high-impedance connection between the photodiode output and the inverting input of the transimpedance amplifier. If we consider 10 GΩ of resistance across the PCB between a 5 V supply trace and a high-

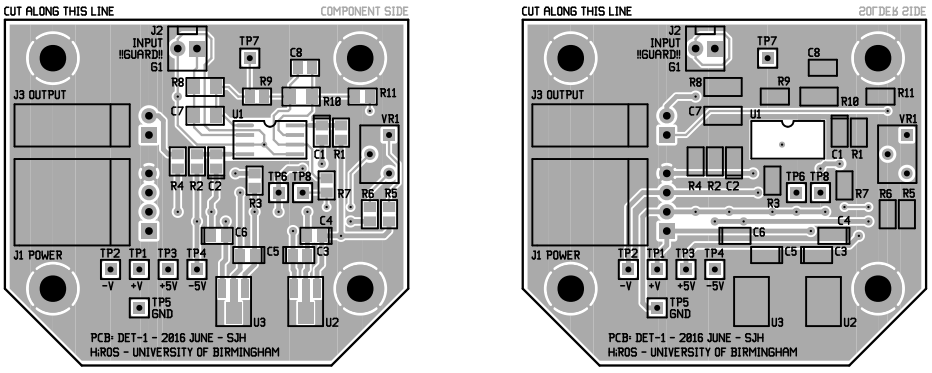

(a) Top-side (b) Bottom-side

Figure 8.13: Amplifier PCB rendering. (a) shows the top-side of the PCB. (b) shows the bottom-side of the PCB.

impedance input this adds 500 pA of leakage current – equal to the signal we are are trying to amplify and three orders of magnitude higher than the input bias current of the operational amplifier. A technique to reduce the effects of input capacitance and leakage currents for small signals at high impedance levels is a signal guard electrode (Horowitz and Hill, 2015, p. 587). A low-impedance ring connected to signal-ground is used to surround the high-impedance input and this protects everything inside the ring from external leakage currents. The soldermask is removed from this section of the board to avoid creating a leakage bridge over the guard ring.

The top and bottom PCB layout is shown in Figure 8.13, along with a photo-realistic rendering in Figure 8.14. Note the visible guard ring around the connector J2, the gain resistor R8, and the compensating capacitor C7, and the inverting input to op-amp U1. The board is finished with soldermask and silk screen, and the pads are finished with electroless nickel immersion gold (ENIG) surface plating. The surface-mount components were soldered in a vapour-phase reflow oven, and the through-hole components added later by hand.

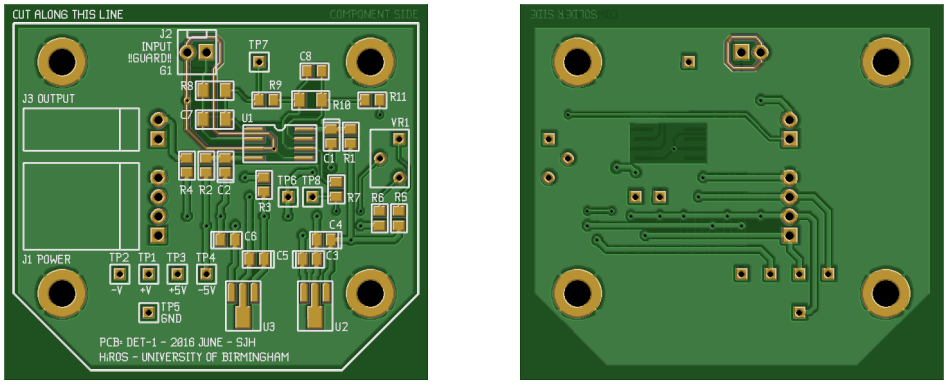

(a) Top-side (b) Bottom-side

Figure 8.14: Amplifier PCB photorealistic rendering. (a) shows the top-side of the PCB. (b) shows the bottom-side of the PCB. Note the guard-ring around the high-impedance amplifier input with solder-mask removed.

8.4.3 Housing

The housing of the detectors is a key component of the design and contributes significantly towards the low noise performance of the system. After taking great care to acquire and amplify our signals, we now need to consider how to export them from the detector. We have discussed the idea of one controlled ground point within the detector, but now we need to transmit signals between two instruments which both have their own idea of "ground" despite how well we might try to tie them together. We use shielded coaxial cable connected to the amplifier output and signal ground at the driving end, but keep the shield isolated from ground at the receiving end. We'll see in the next section that since the ADS1210 analogue-to-digital converter has a differential amplifier as an input buffer, the small amount of "ground signal" is ignored (Horowitz and Hill, 2015, p. 583). The coaxial cable is connected via insulated BNC connectors, ensuring that the shield is connected directly to signal-ground and not to the noisy chassis-ground via the case. The power, temperature sensor, and Peltier drive connections are made via a 9-way D-connector which requires no significant consideration, except to ensure the LM35 temperature sensor cables are routed well away from the photodiode output to ensure the diode does not pickup noise from the temperature controller. The connectivity pin-out is shown in Figure 8.15.

The housing itself is a diecast metal box connected to chassis ground, and painted

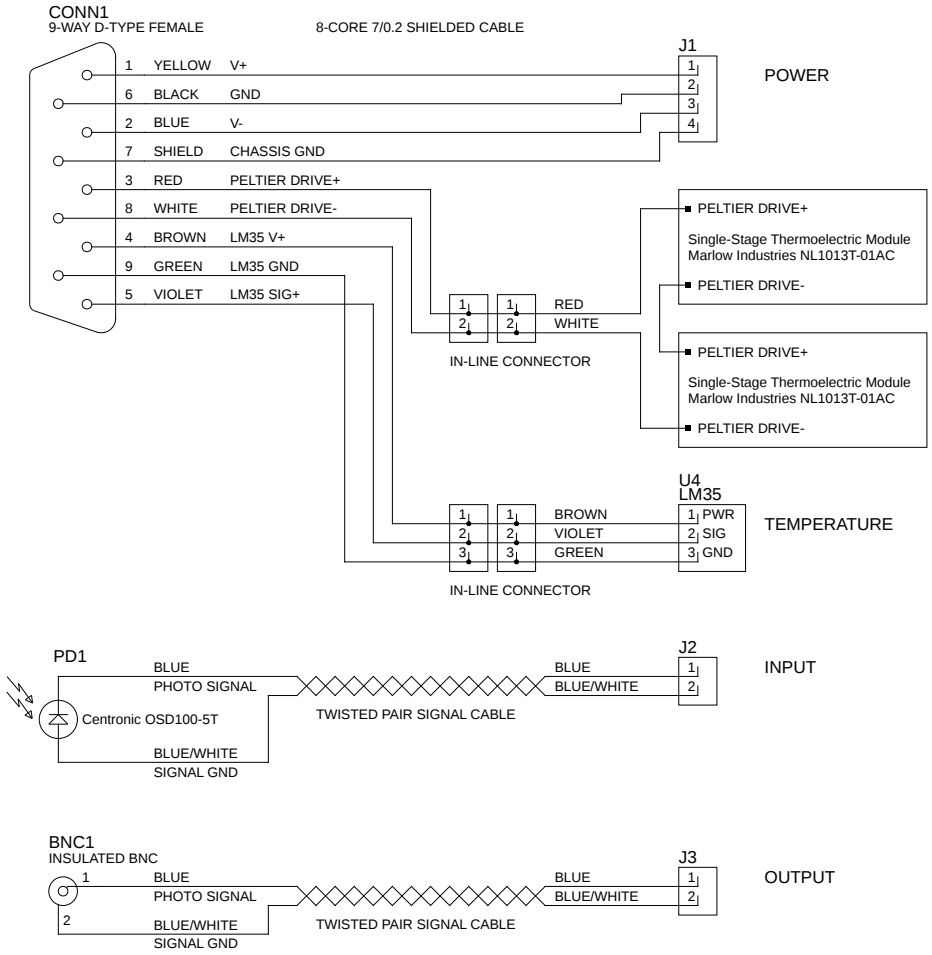

Figure 8.15: Amplifier external connections and pin-out.

matt black to reduce internal reflections. Ensuring the box is grounded helps to shield the internal amplifier from most interference, but magnetic coupling can still be a problem particularly in terms of 50 Hz mains pickup. Pickup has been mitigated by avoiding as much as possible any traces that close around in a loop and become an antenna. The connections from the photodiode, and from the board to the case BNC connector, are made using twisted pairs of wires since this allows the signals induced in successive twists cancel (Horowitz and Hill, 2015, p. 581). The photodiode is mounted in a metal frame on top of two Peltier thermoelectric modules for temperature stabilisation, along with the LM35 temperature sensor. One of the assembled detectors is shown in Figure 8.16.

(a) Assembled components (b) Detector housing and faceplate

Figure 8.16: Detector housing assembly. (a) shows the internal components with the photodiode and PCB mounted. (b) shows the completed assembly with mounting faceplate.

8.4.4 Electronic Noise Performance

The noise performance of the new scattering detectors has been simulated using SPICE, shown as the red line in Figure 8.17. To compare real performance with the simulation, noise from an isolated detector was captured with a Keysight InfiniiVision MSO-X 2012a oscilloscope. The oscilloscope runs at several mega-samples per second and was configured to decimate to a total of 62 500 samples over a period of 0.5-s. This produces a final sampling rate of 125 kHz with 8 µs per sample, providing a 62.5 kHz Nyquist frequency limit. A total of 30 realisations of noise were captured in this configuration. For each realisation the Fourier transform was calculated and the amplitude spectral density determined, removing all phase information. Finally, the mean of all amplitude spectra was found in order to improve the signal to noise ratio in the absence of phase differences. The result is shown as the black line in Figure 8.17.

Despite our best efforts the physical detector suffers from some slight mains pickup at 50 Hz, and the peak around 30 kHz is due to powering the detector from a switched-mode power-supply (SMPS). The noise is minor and the performance matches expectation. We can conclude that the detector is performing at the physical limit, and the only

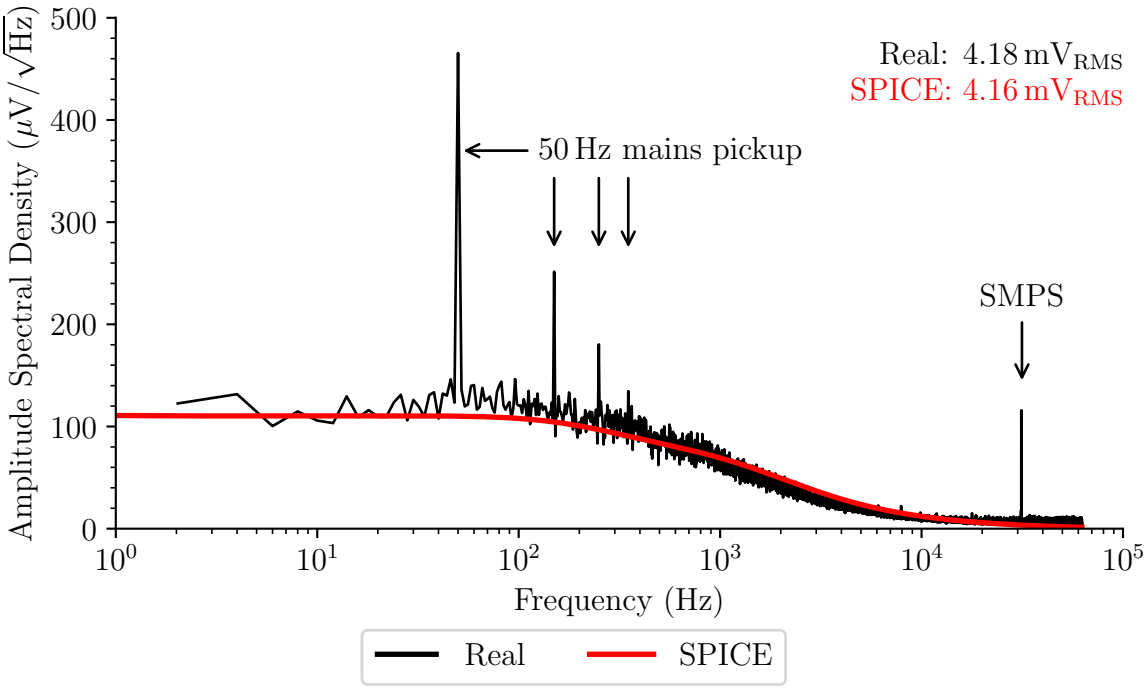

Figure 8.17: Measured scattering detector noise, along with the expected noise profile from SPICE simulation. The physical detector suffers from some minor mains pickup at 50 Hz, and a peak around 30 kHz due to powering the detector from a switched-mode power-supply (SMPS), but broadly matches expectation.

way to reduce the noise level is to either reduce the gain (brighter source) or reduce the size of the photodiode (higher shunt resistance and lower capacitance). The RMS noise amplitude is approximately 4.2 mV$_{RMS}$, which is comparable to the Jabba detector estimate at 8.4 mV$_{RMS}$ when taking into account the gain difference due to requiring half the dynamic-range. The typical noise value of 100 µV $\sqrt{\text{Hz}}^{-1}$ can be rescaled by the amplifier gain enabling the noise equivalent power (NEP) to be calculated. This is the equivalent signal that would be required to produce the same output from an ideal noiseless detector. The electronic NEP is 0.07 pW $\sqrt{\text{Hz}}^{-1}$.

In the next section we will consider the process of analogue-to-digital conversion and data logging using a Raspberry Pi single-board computer.

8.5 Analogue-to-Digital Conversion

The data acquisition system is based around a number of Texas Instruments ADS1210 24-bit delta-sigma ($\Delta\Sigma$) analogue to digital converters which are capable of data output rates of several kHz (Texas Instruments, 1996), all read out via a Raspberry Pi single-board computer over the on-board Serial Peripheral Interface (SPI) bus. In this section we'll look first at the operation of the ADC, and then go on to consider the data acquisition system as a whole.

The ADS1210 is an integrated package containing a programmable gain amplifier (PGA), a second-order delta-sigma modulator, a programmable digital low-pass filter, and a micro-controller with several control registers. The $\Delta\Sigma$ modulator is essentially a combined voltage-to-frequency converter and a pulse counter (Baker, 2011a). The modulator samples the incoming signal at a very high rate, clocked by an external crystal, and converts it into a 1-bit stream of pulses by comparing the difference between the current input signal (the Δ) and passing it through an integrator (the Σ) and a comparator. A digital filter then reduces and converts the pulses into a high-resolution measurement of the input signal. This means that there are two sampling rates – the input sampling rate and the output data rate. The input sampling rate, f_{SAMP}, is defined by,

$$f_{\mathrm{SAMP}} = \frac{\mathrm{TMR} \,.\, \mathrm{PGA} \,.\, f_{\mathrm{XIN}}}{512} \,, \tag{8.7}$$

where f_{XIN} is the frequency of the crystal oscillator controlling the main clock of the ADS1210, PGA is the multiplier of the programmable gain amplifier, and TMR is the turbo mode rate multiplier which can boost the sampling rate for higher precision. The ADC is installed with a 5 MHz crystal, the TMR is set to the maximum 16 times multiplier, and the PGA set to the minimum unity-gain. This produces an input sampling rate of 156.25 kHz. By oversampling at a rate much higher than the band of interest, the quantisation noise is pushed up to higher frequencies and then removed by the low-pass decimation filter producing an effect known as noise shaping. The effectiveness of noise shaping can be improved by using multiple integrators. The ADS1210 uses a second-order

modulator with two integrators meaning that the noise term depends on the previous two errors rather than just one.

The output data rate, f_{DATA}, and the bandwidth of the digital low-pass filter are both determined by the decimation ratio (Baker, 2011b). The output data rate can be calculated from,

$$f_{\mathrm{DATA}} = \frac{\mathrm{TMR} \,.\, f_{\mathrm{XIN}}}{512 \,.\, (N+1)} \ , \tag{8.8}$$

where N is the decimation ratio. For the ADS1210, valid decimation ratios range from 19 to 8000. If values are defined outside of this range then the digital filter will have either not enough, or too much, data resulting in incorrect results. By rearranging equation 8.8 we can determine the required decimation ratio for the desired data rate,

$$N = \frac{\mathrm{TMR} \,.\, f_{\mathrm{XIN}}}{512 \,.\, f_{\mathrm{DATA}}} - 1 \ , \tag{8.9}$$

where terms are as defined before. Lower decimation ratios produce higher data rates with higher bandwidth, and as such include more quantisation noise. Lower data rates provide better signal-to-noise ratio and so increase the effective number of bits (ENOB) of the ADC. The digital filter used by the ADS1210 has a $(\sin(x)/x)^3$ response known as a sinc^3 filter, and it is described by the following transfer function,

$$|H(f)| = \left| \frac{\sin\left(\frac{\pi f N}{f_{\mathrm{MOD}}}\right)}{N \sin\left(\frac{\pi f}{f_{\mathrm{MOD}}}\right)} \right|^3 \ , \tag{8.10}$$

where N is again the decimation ratio, f_{MOD} the frequency of the $\Delta\Sigma$ modulator, and f the input signal frequency. Figure 8.18 shows the frequency response in terms of multiples of the data rate. The $-3\,\mathrm{dB}$ point of the filter occurs at 0.262 times the data rate. Signals in the stopband, defined as frequencies higher than the first notch frequency at the data rate, have at least $-40\,\mathrm{dB}$ rejection. The low-pass filter is essential not only for noise shaping but also to prevent aliasing from frequencies above the Nyquist limit, and so the design specification has to be chosen carefully. The ADCs are configured with a decimation ratio of 155 producing a 1 kHz data rate with a low-pass $-3\,\mathrm{dB}$ point at 262 Hz. Conversion cycles run continuously and set a data-ready flag once each conversion cycle is complete.

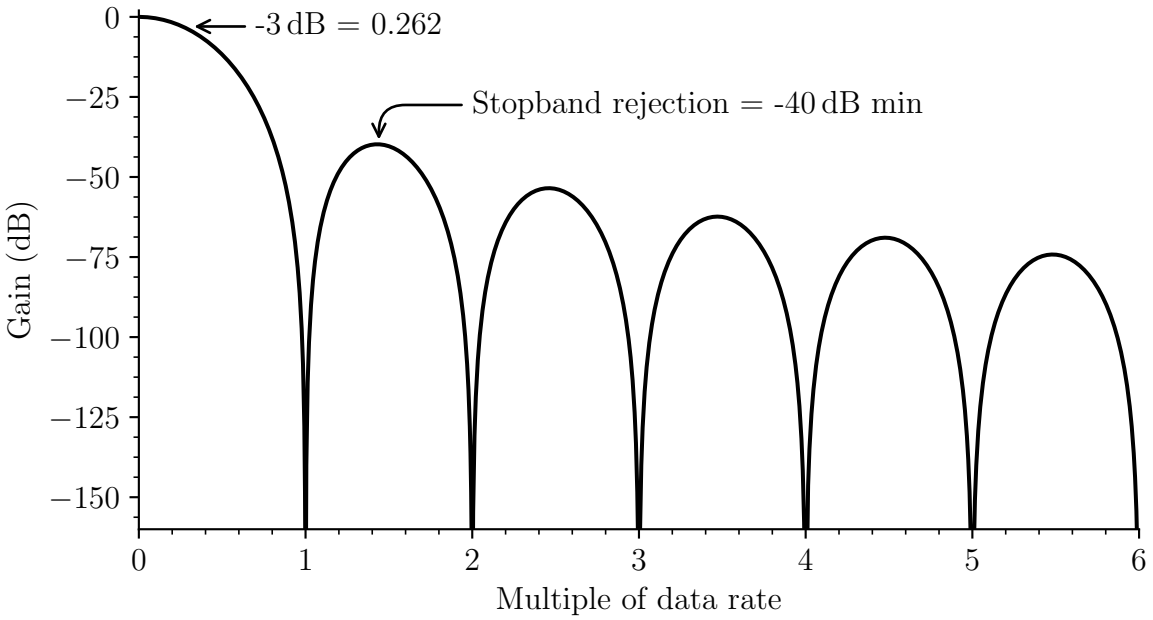

Figure 8.18: ADS1210 low-pass filter response. The −3 dB point of the filter occurs at 0.262 times the data rate. Signals in the stopband, defined as frequencies higher than the first notch frequency at the data rate, have at least −40 dB rejection.

The data-ready line is used to trigger an interrupt on the Raspberry Pi causing it to read out the data from the IC and clear the data-ready flag ready for the next conversion cycle. We'll now look at the control side of the hardware and how everything is integrated into a complete data acquisition system.

The Raspberry Pi is a small single-board computer based around a Broadcom System-on-a-Chip (SoC) with an integrated ARM compatible CPU, and runs a version of the GNU/Linux operating system (Stallman, 1983; Torvalds, 1991). We use the current latest Raspberry Pi 3 which is powered by the Broadcom BCM2837 SoC with a 1.2 GHz 64-bit quad-core ARM Cortex-A53 processor, with 1 GB of RAM, and 512 kB shared L2 cache. The Pi makes available a 40 pin General Purpose Input-Output (GPIO) connector which includes Serial Peripheral Interface (SPI), Inter-Integrated Circuit (I^2C), and Universal Asynchronous Receiver-Transmitter (UART) communications interfaces in addition to basic digital IO. We saw earlier that the ADS1210 ADCs connect via the SPI bus, which is a synchronous serial communication interface used primarily in embedded systems. SPI devices use a master-slave architecture, with a single master and multiple slave devices

which are selected through asserting individual slave select (SS) lines. Communication is full-duplex with the master controlling clock cycles on the SCLK line. During each clock cycle, the master sends a bit on the master-out-slave-in (MOSI) line, while the slave sends a bit on the master-in-slave-out (MISO) line. Transmissions typically consists of 8-bit words, meaning that three transmissions are required for each 24-bit ADC value. Several Texas Instruments CD40109BE voltage level shifters are used to convert between the 5 V logic level of the ADS1210 and the lower 3.3 V logic level of the Pi. The integration period of each conversion cycle is handled in hardware by the ADC system clock, and global network synchronisation is handled in software by the Pi. The Raspberry Pi does not have a hardware real-time clock, and so timekeeping is maintained by the Network Time Protocol (NTP) which is capable of synchronising all participating computers to within one millisecond of Coordinated Universal Time (UTC) under ideal conditions. Even over the public internet with variable network latency, accuracy to better than tens of milliseconds is typically achieved. The final data product is compacted to a cadence of 40-s, and so even a 100 ms clock error is only 0.25 %. Clock errors of less than 10-s are not visible in archive data when comparing timing between contemporaneous observations. The acquisition code is written in Python and runs as a daemon (Disk And Execution Monitor), meaning that it runs in the background with no controlling terminal (Stevens and Rago, 2013, p. 463) and is fully autonomous. The code runs continuously, handling date changes and closing/opening appropriate output files at sunrise and sunset with no operator intervention. In the next section we will look at polarisation control and switching rates between observational states.

8.6 Polarisation Control

Instrumental polarisation state is controlled via a TTL logic level signal from the Pi to the LCC25 LCD controller. A key consideration in reducing noise is the rate of polarisation switching. In the previous section we saw that for Johnson noise, the noise voltage

measured depends on the measurement bandwidth, i.e., how much noise you see depends on how fast you look. The same is true for atmospheric scintillation, where measuring different polarisation states as close to contemporaneously as possible, i.e., at high switching rates, minimises the atmospheric noise contribution. In between each polarisation phase is a delay known as the "front porch" to allow mechanical and electrical stabilisation of the polarisation state. Existing BiSON instrumentation applies a front porch delay of 0.5 ms followed by an integration period of 5 ms. This is repeated for 170 cycles over four polarisation states giving a total acquisition period of 3740 ms, after which the data are read out and the next 4-s integration period is started. The switching rate is 90.9 Hz for the line-of-sight velocity Pockels cell, and 45.45 Hz for the solar mean magnetic field (SMMF) Pockels cell. There is 15 % dead time due to stabilisation delays and data readout within each 4-s integration period. However we saw in Section 8.3 that the scattering detector amplifier rise time is 2.5 ms to 4 ms depending on the acceptable percent of transition completion, and this means that polarisation state mixing between data channels is reducing the instrumental sensitivity. If the front porch is increased to 2.5 ms at the same data rate then the dead time increases to 49 %, and at 4 ms the dead time rises to 74.5 %. If we wish to maintain the existing 15 % dead time with a 2.5 ms front porch then the velocity switching rate must reduce from 90.9 Hz to just 22 Hz.

In Section 7.5 we found the the LCD retarder used in replacement of a Pockels-effect cell has a rise time of 16 ms to 27 ms depending on the acceptable percent of transition completion. When applying a 16 ms front porch with only two polarisation phases (no SMMF measurement) we can achieve a switching rate of 5 Hz if we allow the dead time to rise to 20 %. We will assess the impact of this reduced switching rate later in Section 8.8 while commissioning the new instrument at Izaña. In the next section we will look at temperature control of various components.

8.7 Temperature Control

There are five components of the spectrophotometer that require active temperature control. These are the cell top and stem, the LCD variable retarder, and the two photodetectors. The first three of these are unipolar channels – they require only heating above ambient temperature. The two detectors require bipolar control, both cooling and heating, since they require stabilisation at all ambient temperatures. Temperature control is achieved via a five-channel computer-controlled current-source that can dynamically change the power to the heating and cooling elements as required based on readout from several temperature sensors (Barnes and Miller, 2011a).

Figure 8.19 shows a generic temperature sensor amplifier circuit that is used for all five channels. Each channel can be configured for either LM35, AD590, thermistor, or pure voltage sensor input. The photodetectors use LM35 sensors, the two oven channels use AD590 sensors since they require larger dynamic range, and the LCD uses a thermistor. The input stage is configured appropriately for each sensor, and all signals are amplified using an INA114 instrumentation amplifier producing an output that varies between 0 V to 5 V. The values are sampled using the same ADS1210 24-bit analogue to digital converter as the photodetector outputs. The control voltages for the current sources are produced by a MAX536 12-bit digital to analogue converter (DAC) and an LM675 power op-amp as the current source, shown in Figure 8.20. Components with value SOT (set on test) are configured as required by the use case. The current-limit and unipolar/bipolar operation are selected by the configurable components. If we set $R_a = R_b$ and $R_x = R_y$ then the op-amp will supply a constant current I,

$$I = (V_+ - V_-)\frac{R_x}{R_a R_z} \ , \tag{8.11}$$

where V_+ is the DAC output varying between 0 V to 5 V and V_- is the voltage reference level, either 0 V for a unipolar channel or 2.5 V for a bipolar channel.

The five ADCs and the DAC are all controlled via a PIC micro-controller which communicates over RS232. The temperature control calculations are made using a propor-

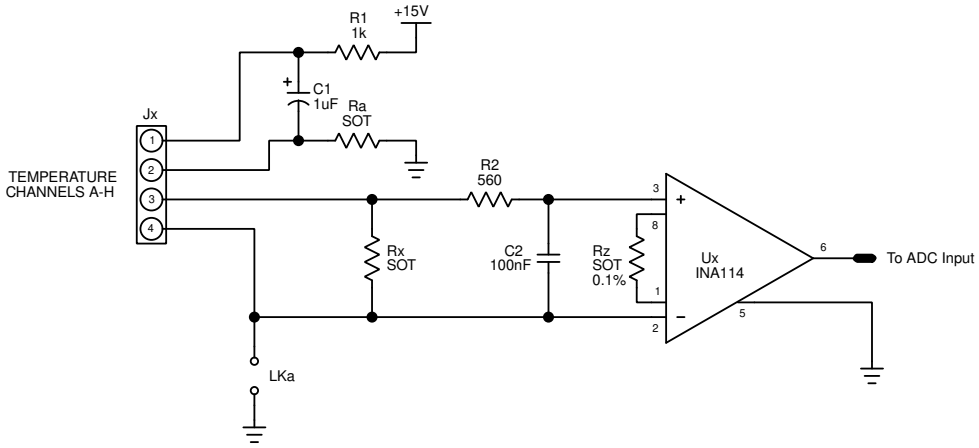

Figure 8.19: Generic temperature sensor input-stage amplifier. The circuit can be configured for LM35, AD590, thermistor, or pure voltage inputs. Components with value SOT (set on test) are configured as required by the use case. Image credit: Reproduced from Barnes and Miller (2011a).

tional/integral/derivative (PID) control loop feedback mechanism running in Python on a Raspberry Pi single-board computer. Simple PID control can be defined as,

$$\text{Output} = K_P\, e(t) + K_I \int e(t)\, dt + K_D \frac{\mathrm{d}e(t)}{\mathrm{d}t} \,, \tag{8.12}$$

where e is the control error defined as the desired setpoint minus the current value, and K_P, K_I, and K_D are the proportional, integral, and derivative coefficients. The algorithm implemented on the Pi is more sophisticated with improvements to prevent problems such as integral wind-up and derivative-kick, but equation 8.12 shows the most basic implementation. The proportional coefficient causes the output to change in a manner that is directly proportional to the error. This term alone is sufficient for many systems that do not require a constant drive signal in order to maintain their position, such as a position actuator. In the case of temperature control, when the temperature setpoint is reached the proportional term goes to zero, and so the output goes to zero. This causes a consistent setpoint error since the actual setpoint cannot be reached. The integral term corrects for setpoint error by acting in relation to the amount of time for which the error has persisted. In this way, the integral term can apply a constant power offset to maintain the setpoint while the proportional term acts to stabilise fluctuations. The derivative term acts as a damping force to reduce the rate of change of error, with the aim of reducing

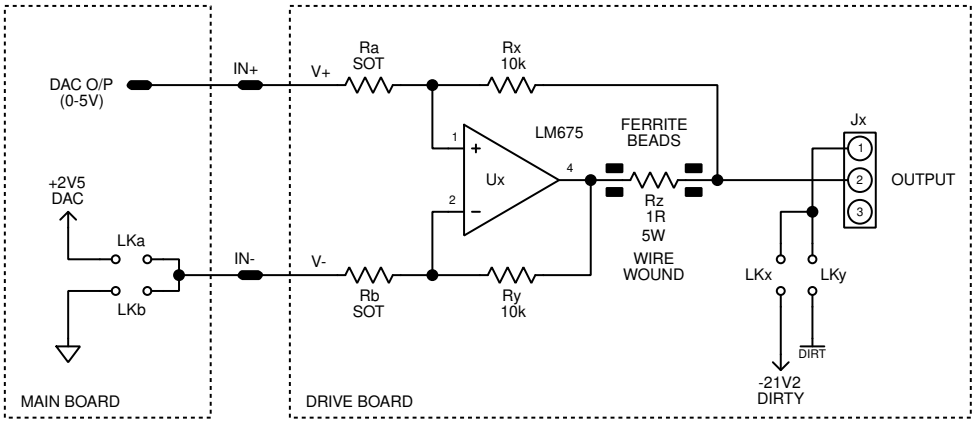

Figure 8.20: Computer-controlled digital current source for temperature drive outputs. The circuit can be configured for unipolar and bipolar temperature control elements via the LK links, and for current limits up to 2 A. Components with value SOT (set on test) are configured as required by the use case. Image credit: Reproduced from Barnes and Miller (2011a).

overshoot.

The temperature control implemented here generally makes use of only the proportional and integral terms, with the derivative being unnecessary. The tuning technique is to increase the proportional coefficient until the temperature begins to oscillate about the setpoint, and then back it off until the temperature stabilises. This is the maximum proportional gain allowed and results in a constant setpoint error slightly below the commanded setpoint. The integral coefficient is then increased until the setpoint error is minimised, again ensuring that the system does not oscillate. Once the PID coefficients are optimised it is typical to achieve temperature stabilisation to better than $1 \times 10^{-3}\,^{\circ}\mathrm{C}$ through the implemented algorithm.

In the next section we will discuss commissioning the new instrument, the challenges encountered, and compare the data quality with other instruments in the BiSON network.

8.8 Commissioning

The new prototype system was commissioned at Izaña, Tenerife, over two site visits. A summer observing campaign began with initial setup during 2017 May (Hale, 2017b), with secondary updates during 2017 September (Hale, 2017c). Light from the pyramid

primary cœlostat was shared between Mark-I and the new system, such that direct comparison could be made between the two instruments which would be more difficult if using independent light-feeds. We will look at the development of the dedicated light-feed later in Chapter 9.

Over the next subsections we will continue the assessment of noise contributions to the data acquisition system, compare results from the three types of collection optics developed in Chapter 7, and finally compare the data quality with other instruments in the BiSON network.

8.8.1 Electronic Noise

We saw in Section 8.4 that the RMS noise amplitude from the scattering detectors is approximately $4.2\,\mathrm{mV_{RMS}}$, and this is exactly as predicted for the circuit and components. While care was taken in the detector design to ensure good grounding techniques and clean signal pathways, the attached data acquisition system does unfortunately induce additional noise in the scattering detector outputs. The majority of the additional noise is above 1 kHz and so does not impact seriously on the final data due to the ADC low-pass filter. The noise RMS increases from $4.2\,\mathrm{mV_{RMS}}$ to between $7.2\,\mathrm{mV_{RMS}}$ (Figure 8.21) and $27.4\,\mathrm{mV_{RMS}}$ (Figure 8.22) depending on the Raspberry Pi CPU load. It is not clear precisely what is causing the noise, but it is most likely due simply to the Pi being completely unshielded. Installing the ADCs within the same enclosure, in close proximity to the Pi, is allowing digital noise to be picked up at the ADC inputs.

When the low-pass filter is applied to the output from the scattering detectors, shown as the green line in Figure 8.22, the noise RMS reduces from 27.4 mV to 4.8 mV which is close to the original clean output of the isolated detector. After rescaling the detector output voltage by the amplifier gain the electronic NEP is of the order of $3\,\mathrm{pW_{RMS}}$. Whilst the low-pass filter removes most of the noise induced by the Raspberry Pi, it would be better if the source of the noise could be eliminated. When the filter is applied to the clean detector output the noise RMS reduces to 2 mV, and so there is potential for a

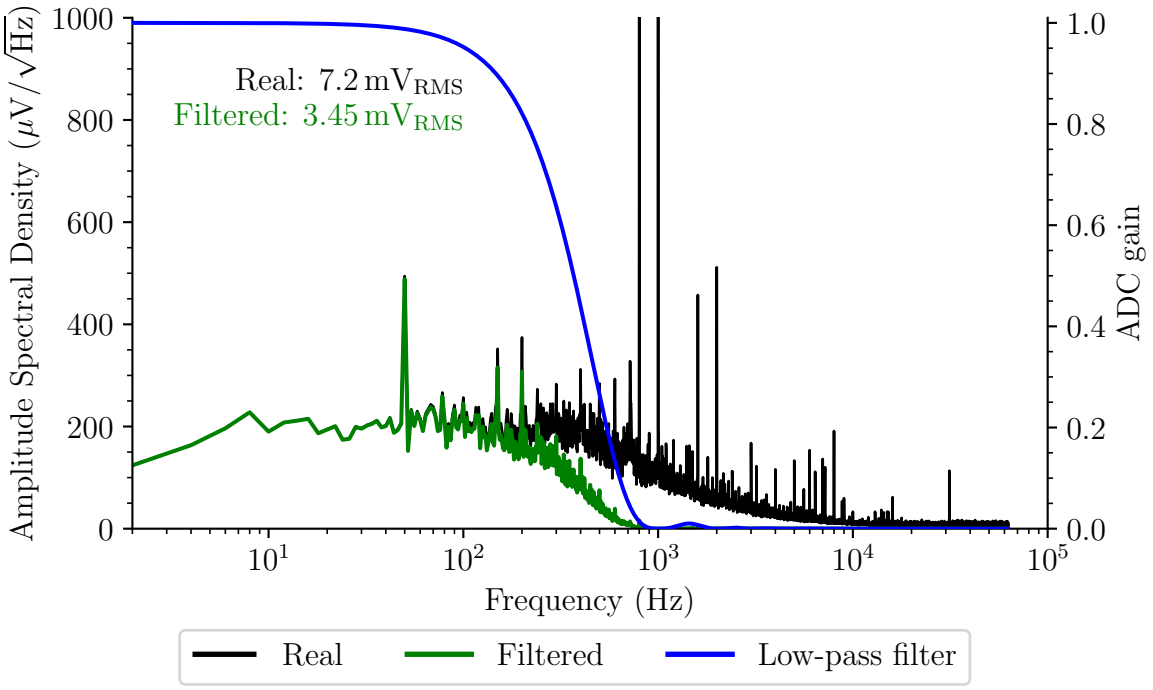

Figure 8.21: The black line shows the noise output from the detector when the temperature controller and Raspberry Pi are operating with low CPU load. The additional noise is predominantly caused by the Raspberry Pi and is above 1 kHz. The green line shows the detector output after passing through the low-pass filter stage of the ADC. The blue line shows the gain of the low-pass filter, measured against the right-hand axis.

factor of 2.4 reduction in the electronic noise level with further work in addition to that achievable by changes in gain and photodiode model.

Better shielding of the ADCs from the Pi within the enclosure, or installing the ADCs directly at the detectors and designing for a longer cable run to the Pi within realistic limits of the SPI interface, will likely eliminate the pickup. However, the electronic noise is already low compared with other noise sources, and so it is unlikely that further reduction in noise level would produce significant improvement to the final data product.

8.8.2 Photon Shot Noise

We know from the work in Chapter 7 that the typical incident flux from the vapour cell on the scattering photodiode is of the order of 1 nW. The energy of a photon in Joules is

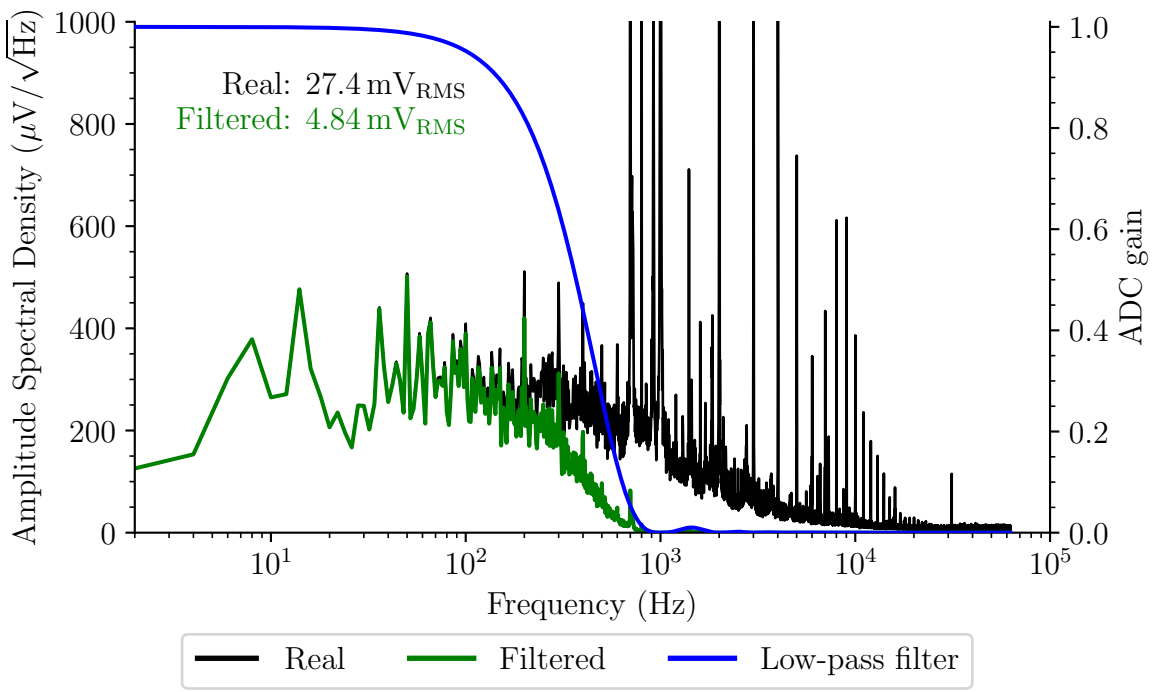

Figure 8.22: The black line shows the noise output from the detector when the temperature controller and Raspberry Pi are operating with high CPU load. The additional noise is predominantly caused by the Raspberry Pi and is above 1 kHz. The green line shows the detector output after passing through the low-pass filter stage of the ADC. The blue line shows the gain of the low-pass filter, measured against the right-hand axis.

given by,

$$E = \frac{h\,c}{\lambda}\ , \tag{8.13}$$

where λ is the wavelength in metres, h the Planck constant in Joule-second, and c the speed of light in metres per second. For 1 nW of light at 770 nm this is equivalent to approximately 4×10^9 photons per second. The standard deviation of shot noise is equal to the square-root of the average number of events, and converting this back to power gives a value of 0.02 pW$_{RMS}$, two orders of magnitude smaller than the electronic noise.

8.8.3 Guider Noise

In BiSON instrumentation where the Sun is imaged inside the vapour cell, an additional source of noise is caused by guider errors. Guider fluctuations cause the image of the Sun to move around inside the vapour cell which changes the vapour optical depth seen

by the scattering detectors, and this in turn causes changes in the measured velocity ratio. The pointing sensitivity of Jabba in Carnarvon was estimated during a site visit to be $0.018\,\mathrm{m\,s^{-1}}$ per arcsecond (Barnes and Miller, 2009), and this confirmed a requirement for sub-arcsecond guiding accuracy achieved with current BiSON autoguiders.

When an image of the Sun is instead focused onto the end of a fibre, the output of the fibre is uniform regardless of the absolute position of the image provided that the image falls entirely within the core of the fibre. The constraint on guiding accuracy is relaxed and guider noise is considered to be eliminated completely.

8.8.4 Thermal Noise

The most important component temperature is the vapour cell stem, since as we saw in Chapter 6 this controls the vapour temperature and thus the measured scattering intensity. The upper plot of Figure 8.23 shows the cell stem temperature controlled over several hours. The mean stem temperature is 90 °C and the standard deviation is 0.7×10^{-3} °C. Using the vapour cell model from Chapter 6 we can convert the temperature noise into scattering intensity noise, assuming a detector aperture of 6 mm and scaling for a typical 1 nW intensity.

The lower plot of Figure 8.23 shows the amplitude spectral density of the effect of thermal noise on the measured scattering intensity, indicating a typical $0.2\,\mathrm{pW}\,\sqrt{\mathrm{Hz}}^{-1}$ which is of a similar order to the detector electronic noise. Regarding the practical aspects of reducing noise, it is essential that inside the detector the cabling for the LM35 sensor is routed well away from the photodiode output and the sensitive amplifier inputs, otherwise noise from the temperature controller is picked up in the photosignal output.

8.8.5 Atmospheric Scintillation

The new spectrophotometer prototype uses an LCD retarder to switch between circular polarisation states, rather than the Pockels-effect cell typically installed in our existing spectrophotometers. As we saw in Section 7.5 an LCD retarder is much slower to switch

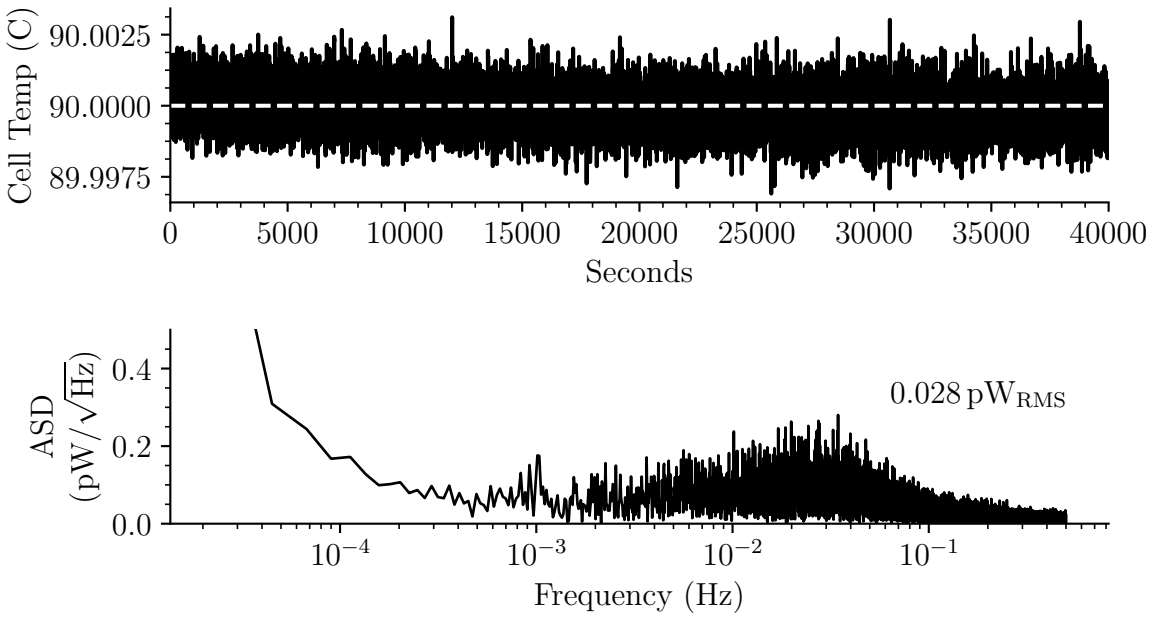

Figure 8.23: Upper: Vapour cell stem controlled at a mean of 90 °C. The standard deviation is 0.7×10^{-3} °C. Lower: The amplitude spectral density of the effect of thermal noise on the measured scattering intensity, indicating a typical $0.2\,\mathrm{pW}\,\sqrt{\mathrm{Hz}}^{-1}$ which is of a similar order to the detector electronic noise.

between polarisation states and this means that the switching rate needs to be reduced from the typical 100 Hz down to the order of 5 Hz while still maintaining a reasonable percentage of dead-time per acquisition period. Since the polarisation state switching is slower, atmospheric turbulence will have a larger impact on the data noise level.

In order to determine the effects of the reduced switching rate, the atmospheric scintillation noise was measured using the output from the transmission monitor at the rear of the instrument. This was done in much the same was as measuring the detector electronic noise. With the polarisation switching disabled, data from the transmission monitor was logged for 10-s at 6.25 kHz. A total of 24 realisations of noise were captured in this configuration, making two minutes of data in 10-s chunks. From each 10-s chunk the amplitude spectral density was calculated and then stacked to improve the final signal to noise ratio. This was done at approximately the same time each day during periods of varying atmospheric stability, and the results are shown in Figure 8.24. The qualitative dust levels are described in Table 5.1 of Section 5.4.1 on page 96. After rescaling the

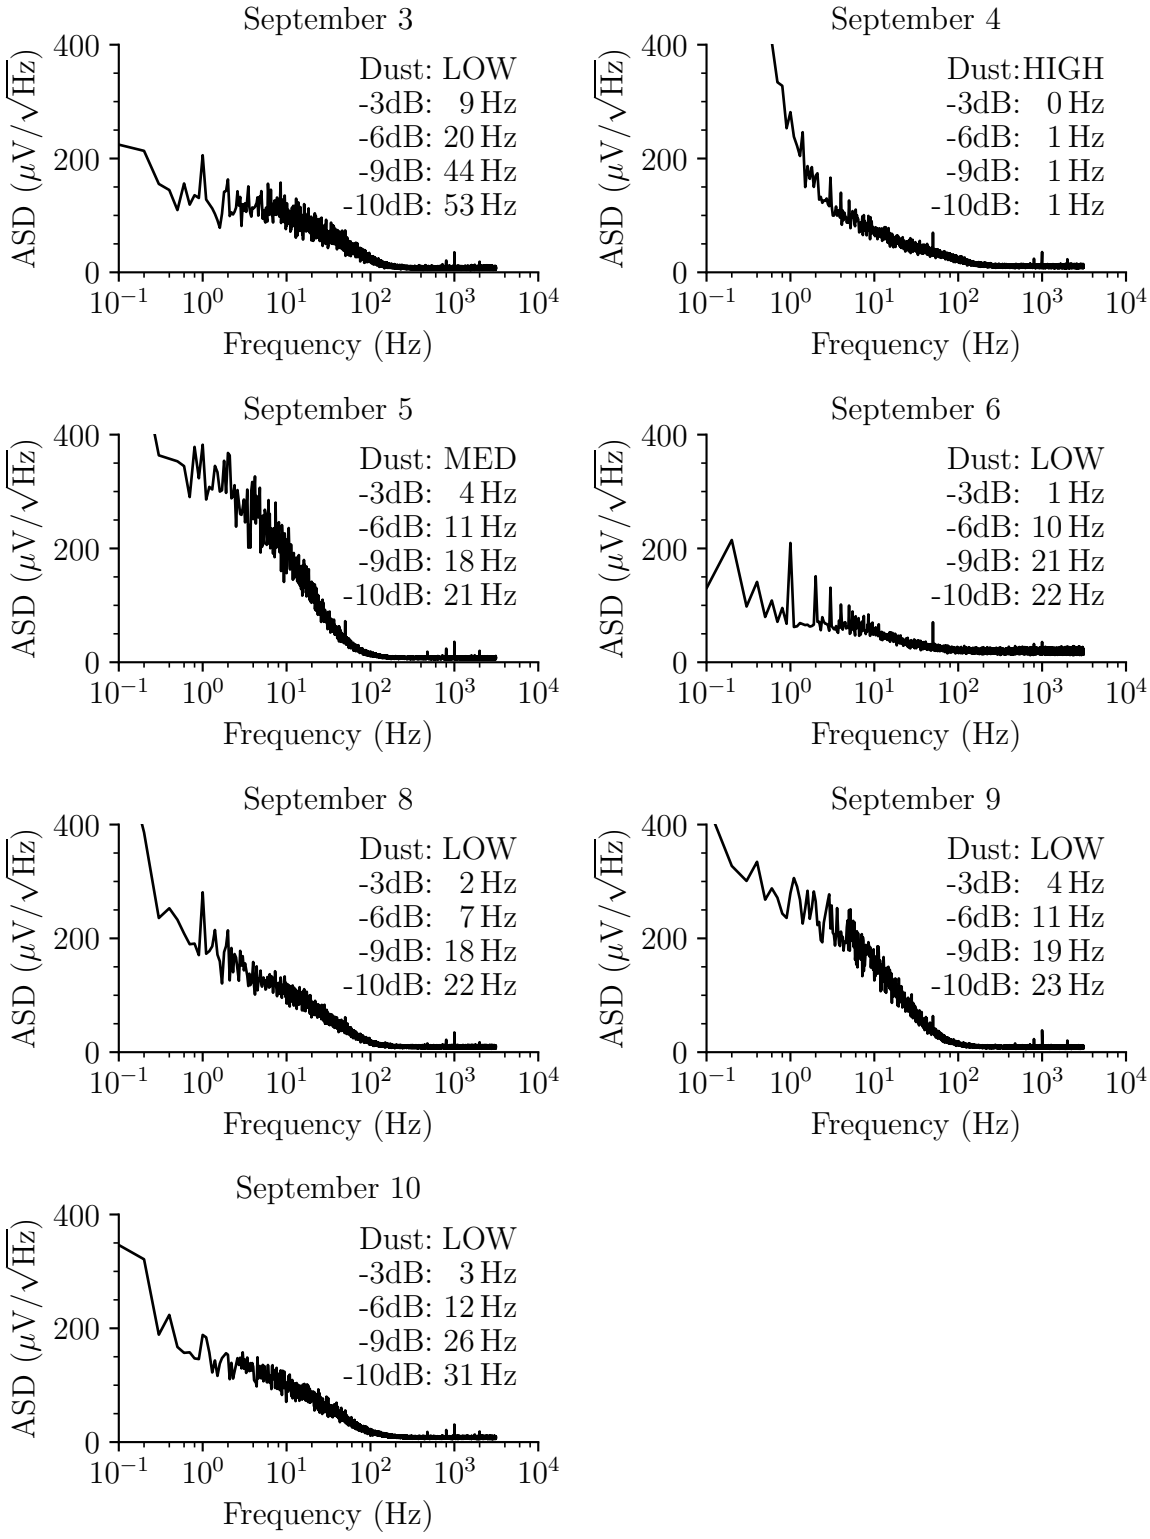

Figure 8.24: Daily scintillation measurements collected during 2017 September. The qualitative dust levels are described in Table 5.1 of Section 5.4.1 on page 96.

detector output voltage by the amplifier gain the scintillation NEP on a typical clear day is $14.8\,\mathrm{nW}\,\sqrt{\mathrm{Hz}}^{-1}$. This is over 10^5 times greater than the electronic noise, and so the system is dominated by atmospheric noise. It is unlikely that we will be able to achieve the same performance as our other instruments switching at 100 Hz, as these operate in the frequency region where the scintillation noise has almost completely decayed. The consequence of this is a higher background noise level within the resulting FFT, and so greater difficulty in resolving low amplitude modes of oscillation. However, the −3 dB point (noise power reduced by 50 %) typically occurs at around 5 Hz, and the −6 dB point (noise power reduced to 25 %) typically occurs at between 5 Hz to 15 Hz, and so even small improvements in switching rate give a potentially large improvement in noise level. The −10 dB point (noise power reduced to 10 %) typically occurs above 20 Hz.

8.8.6 Data Quality

In Section 7.3.2 on page 147 we defined three candidate solar collection optics intended to concentrate sunlight into an optical fibre. We will now examine the data quality from the new instrument, and compare the performance of the three options.

The figures presented here are produced by the BiSON data calibration pipeline "Red". The five-minute FOM calculated by Red is slightly different from the values presented in Chapter 4 since Red calculates a Lomb normalised periodogram as described by Press et al. (2002), elsewhere we use the standard Cooley-Turkey discrete Fourier transform (Cooley and Turkey, 1965) as implemented in the Python NumPy library.

Collection Optics – Option A

Option A was a Thorlabs AC508-080-B-ML doublet lens, with focal length 80 mm and objective aperture 50.8 mm, producing a focal ratio of $f/1.6$. Optical simulation predicted that this lens would form an image of the Sun 0.96 mm in diameter, almost completely filling the 1 mm diameter fibre core. These collection optics were installed and commissioned at Izaña, along with the main instrumentation, during the site visit in 2017 May (Hale,

2017b). The results showed issues with low-frequency performance due to the solar image overfilling the fibre core causing vignetting, and so unequal weighting of light from the solar disc. This was caused partly by the control of aberrations being slightly poorer than suggested by the lens datasheet, and partly by the difficulty in correctly focusing and aligning the image on the end of the fibre with such a small error margin. Two further options were trialled during the Izaña site visit in 2017 September (Hale, 2017c).

Collection Optics – Option B

Option B was a Thorlabs AC254-030-B-ML doublet lens, with focal length 30 mm and objective aperture 25.4 mm, producing a focal ratio of $f/1.2$. The predicted solar image diameter is 0.6 mm, which comfortably falls entirely within the fibre core. Figures 8.25, 8.26, 8.27, 8.28, and 8.29 show the daily data collected with this configuration during the site visit. The regular gap at midday is caused by cœlostat shadowing which is unavoidable at this time of year when moving from summer into winter, and means that it is difficult to assess the true data quality that would be obtained from a completely clear day. Nevertheless the low-frequency performance appears excellent, confirming the previous issues were due to the solar image size formed by the collection optics slightly overfilling the fibre, and the high-frequency noise is of the order of that expected.

Collection Optics – Option C

Both options A and B suffer from potential problems caused by dirt or damage to the end of the fibre, since they rely on an unobscured image of the Sun being accepted and transmitted along the fibre. Any dirt on the end of the fibre would vignette part of the solar disc causing an unequal weighting to the solar rotation,and so give rise to guider noise and loss of sensitivity. Optical arrangement option C proposed to eliminate this possibility by using a second lens in a typical Keplerian telescope configuration in order to couple quasi-parallel light into the fibre rather than forming an image of the Sun. In doing this we convert spatial-position into angular-position, and so the fibre

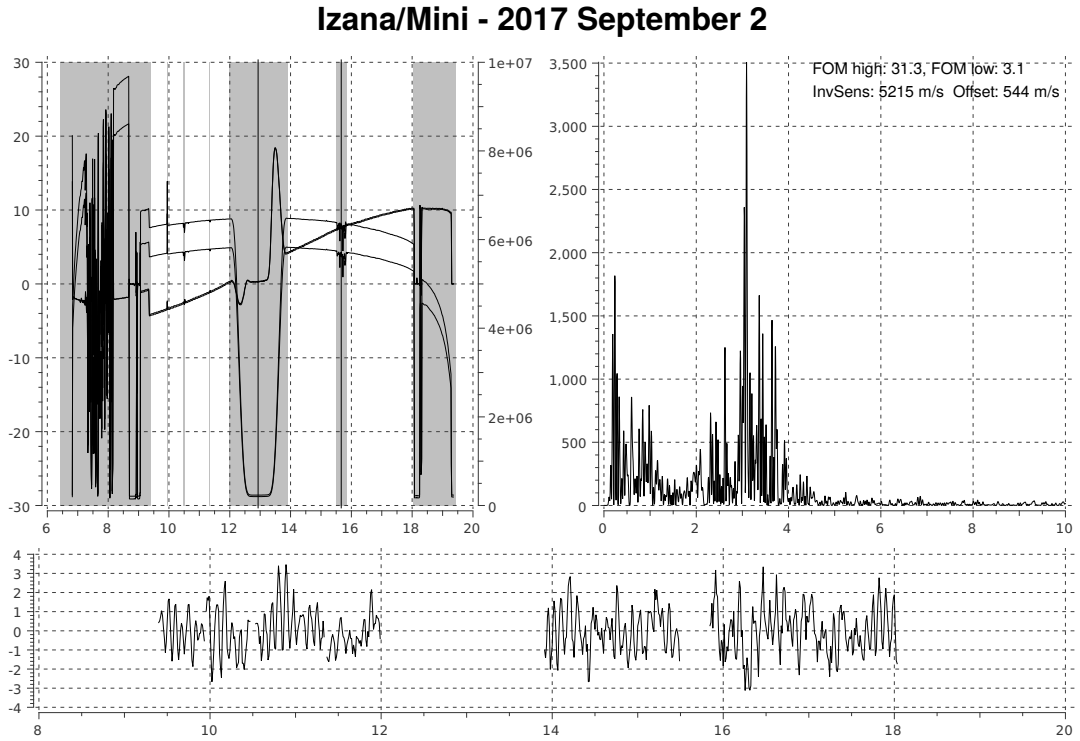

Figure 8.25: 2017 September 2 using optics option B.

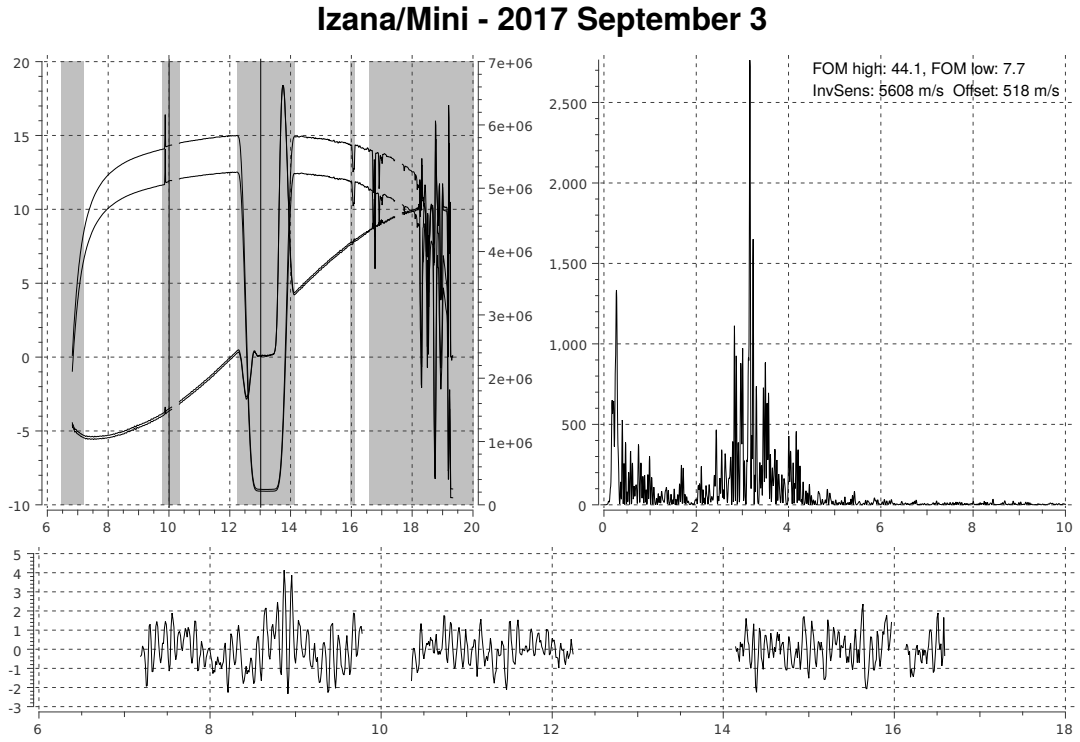

Figure 8.26: 2017 September 3 using optics option B.

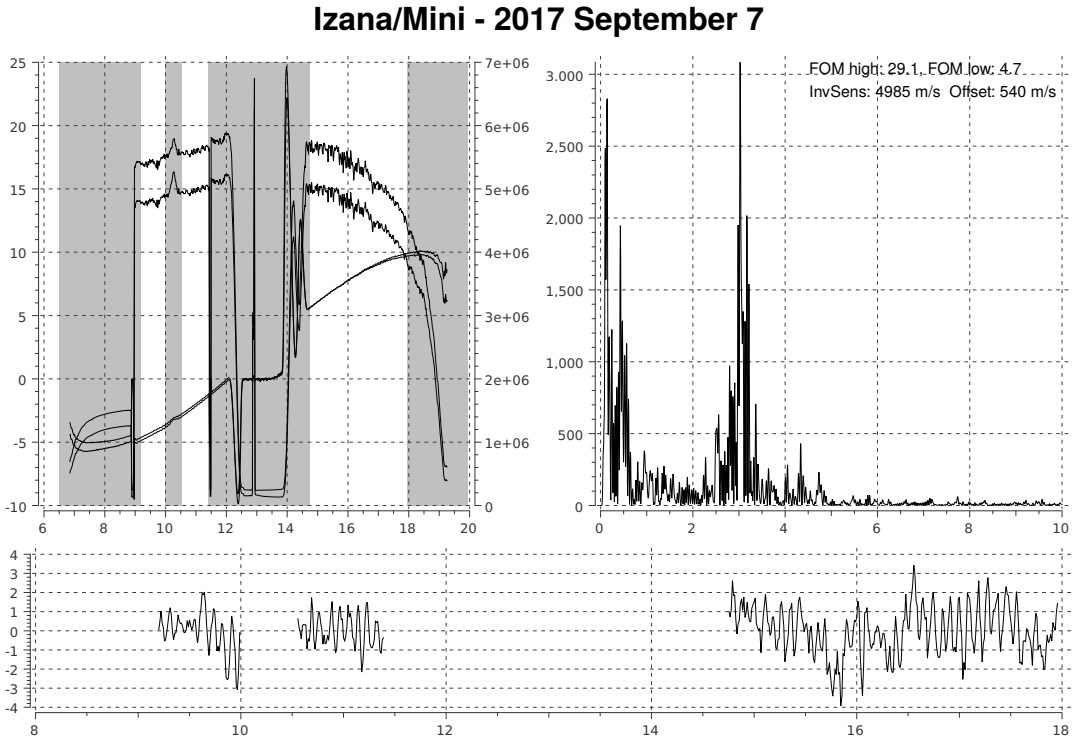

Figure 8.27: 2017 September 7 using optics option B.

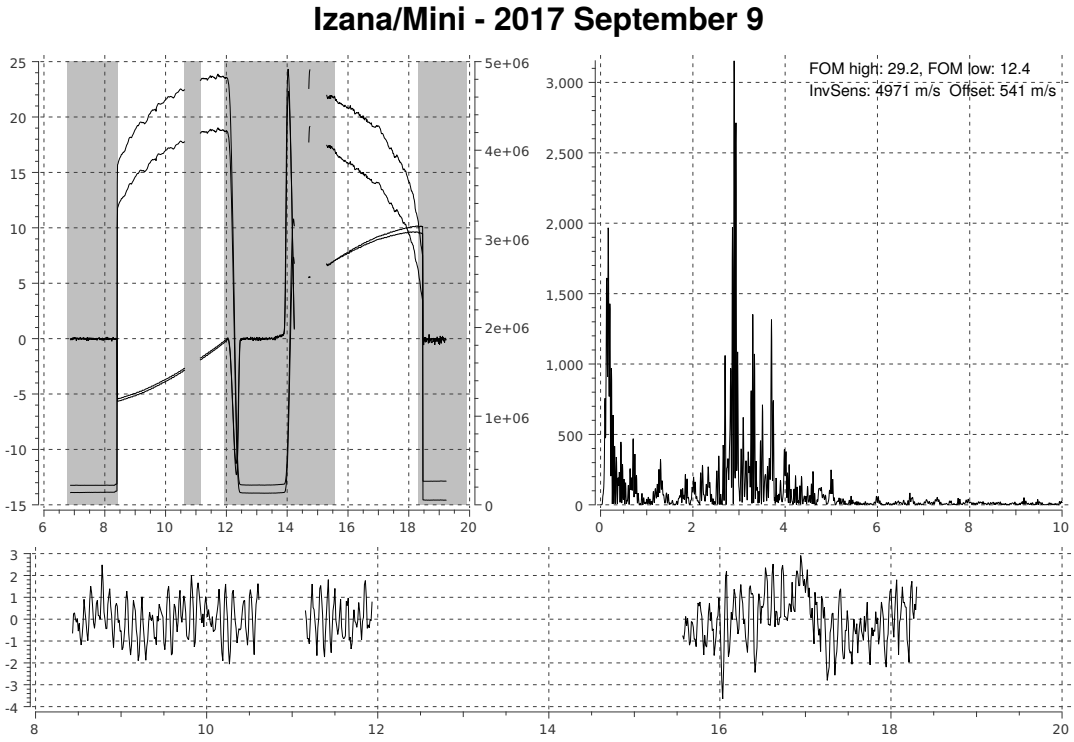

Figure 8.28: 2017 September 9 using optics option B.

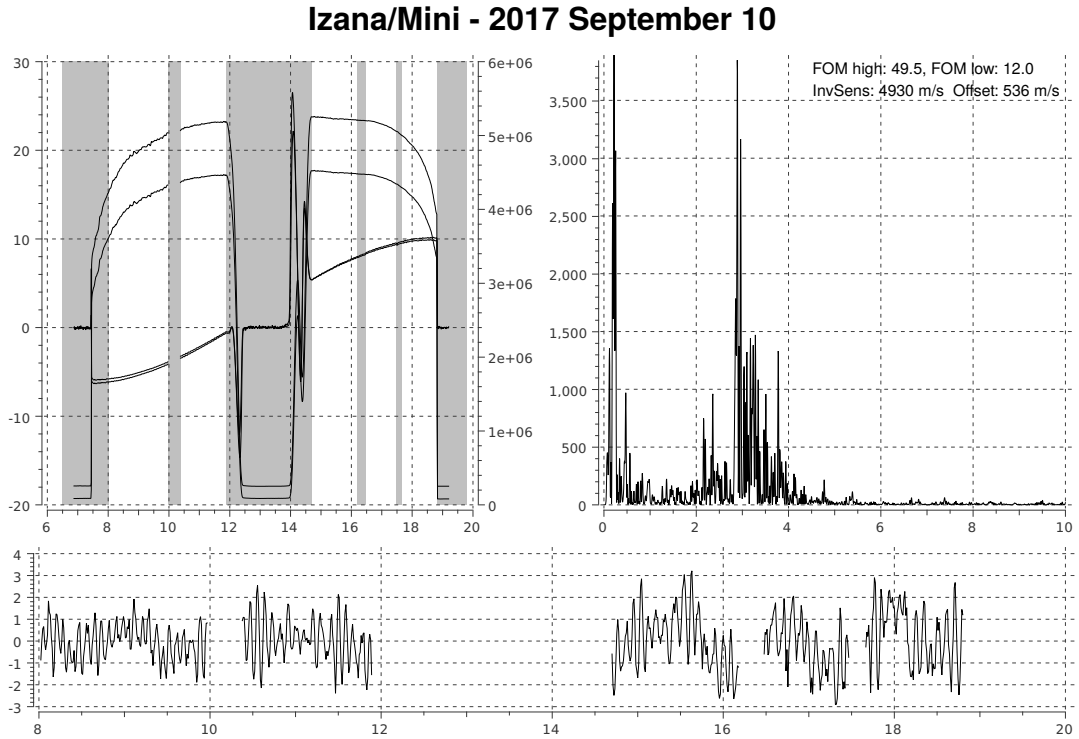

Figure 8.29: 2017 September 10 using optics option B.

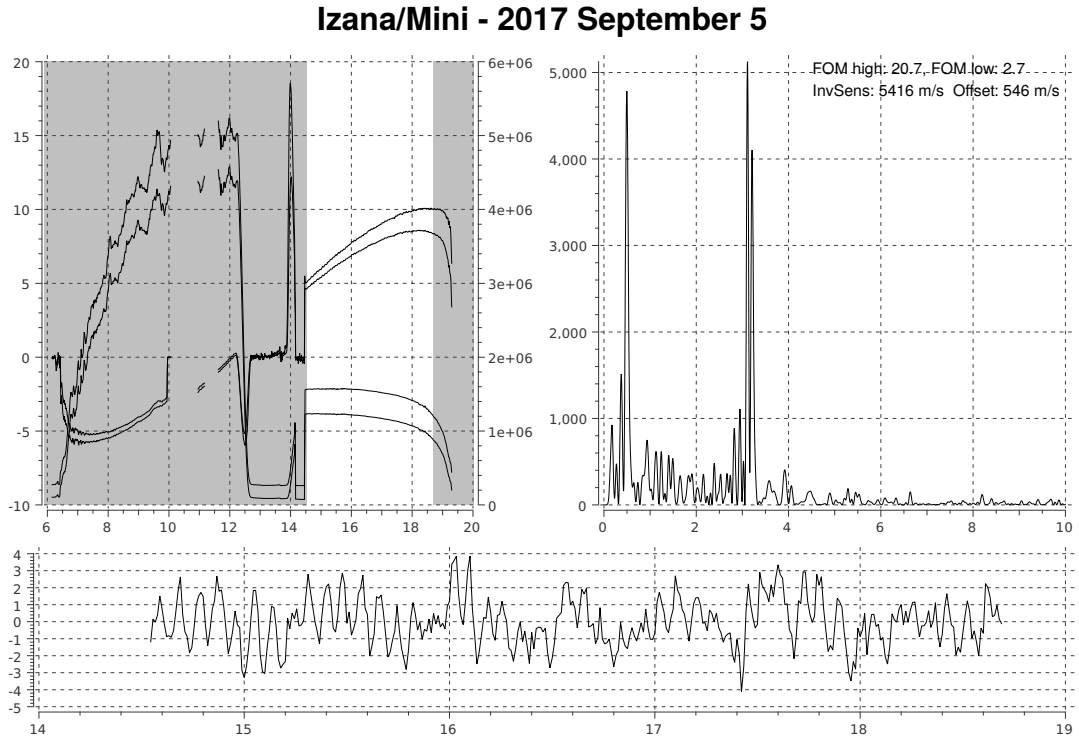

Figure 8.30: 2017 September 5. The second half of this day was collected using the quasi-parallel telescope arrangement, optics option C.

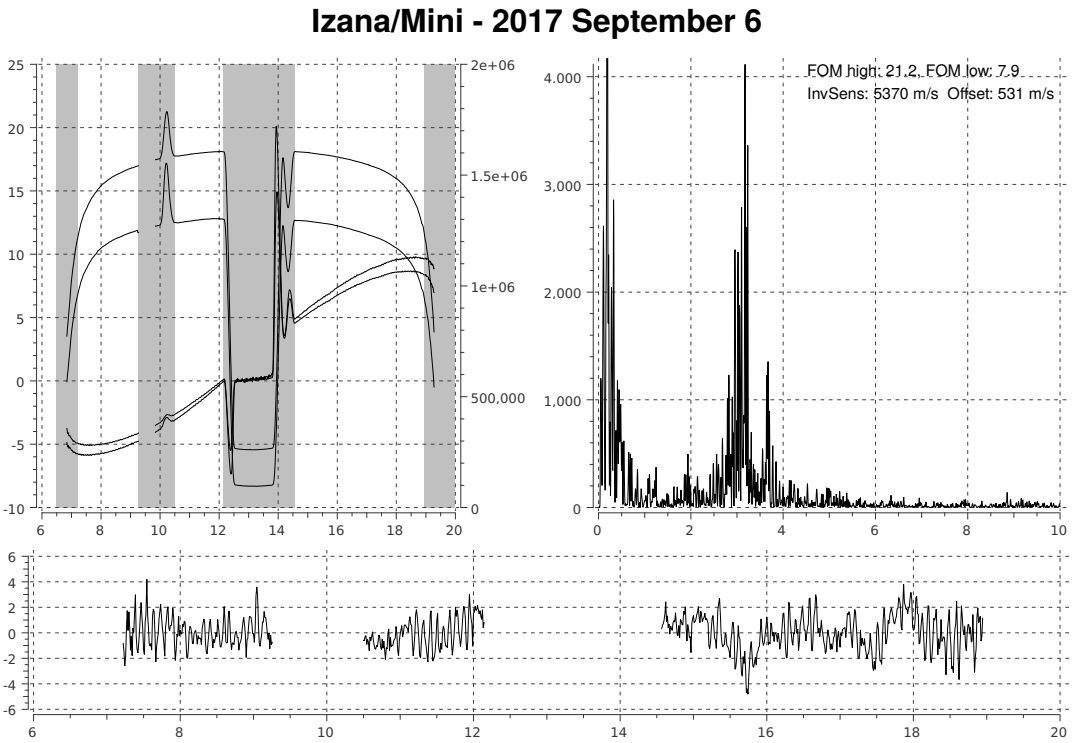

Figure 8.31: 2017 September 6. This day was collected using the quasi-parallel telescope arrangement, optics option C.

receives uniform illumination reducing the effect of guiding errors and inhomogeneities of the polished end of the fibre. Figures 8.30 and 8.31 show the daily data collected with this configuration. These two days show a noticeably higher noise level due to the lower light throughput and efficiency of coupling to the fibre. There also appears to be some unexplained Doppler-imaging issues apparent from the large divergence in velocity ratio from each detector at the beginning and end of each day. With some further work and development this configuration clearly has potential, should the simple approach in optics option B prove to have unresolvable issues. After testing was completed, optics option B was selected for final commissioning.

In the next section we will compare data from the prototype with contemporaneous results from other BiSON spectrophotometers over an entire 2018 summer observing campaign.

8.9 Observing Campaign: Summer 2018

With commissioning of the prototype BiSON Mini instrumentation completed at Izaña in 2017 September (Hale, 2017c), the instrumentation continued to run alongside Mark-I capturing contemporaneous data. The Pyramid cœlostat was used to feed light to both instruments simultaneously. Figure 8.32 shows the power spectral density of concatenated time series from each BiSON site for the summer 2018 observing campaign. Specifically, data from 2018 April 14 to 2018 August 26 were selected from all sites where available since this is the complete period between cœlostat shadowing at Izaña. A 50 µHz moving mean has been applied to smooth the results and aid clarity. Table 8.3 provides a summary of the performance statistics from each site. The FOM and mean noise are defined as previously in Section 4.2 on page 46. If the mean noise is assumed to be white, the noise equivalent velocity can be determined and this is shown in the last column of Table 8.3. The new prototype instrument achieves a mean noise level of $9.9\,(\mathrm{ms}^{-1})^2\,\mathrm{Hz}^{-1}$, comparable to a noise equivalent velocity RMS of $35.2\,\mathrm{cm\,s}^{-1}$, and this is similar performance to both Las Campanas and Narrabri. There is clearly an issue with low-frequency stability apparent from the increased noise level around 1 mHz, but this is also visible within the data from Mark-I and is considered to be a result of using the light feed from the Pyramid cœlostat. Whilst making use of a fibre feed does eliminate many guiding related problems, the cœlostat does still introduce problems with day to day alignment and drifts. Alignment issues can cause vignetting of the image of the solar disc at the edge of the fibre core, and this can give rise to the low-frequency effects.

During parallel development work on this project, both Jabba at Carnarvon and Klaus at Mount Wilson were converted to use a fibre feed based on the same optical schematic as the new prototype instrumentation (Hale and Ross, 2018; Hale, 2018b). At Carnarvon, the existing equatorial mount was used for the optical feed. Unfortunately, shortly following the conversion a failure of the dome shutter motor, and subsequently the PC, resulted in extended downtime and so no data are available for comparison. At Mount Wilson, a commercial semi-automatic solar-tracker was installed and this bypassed the damaged

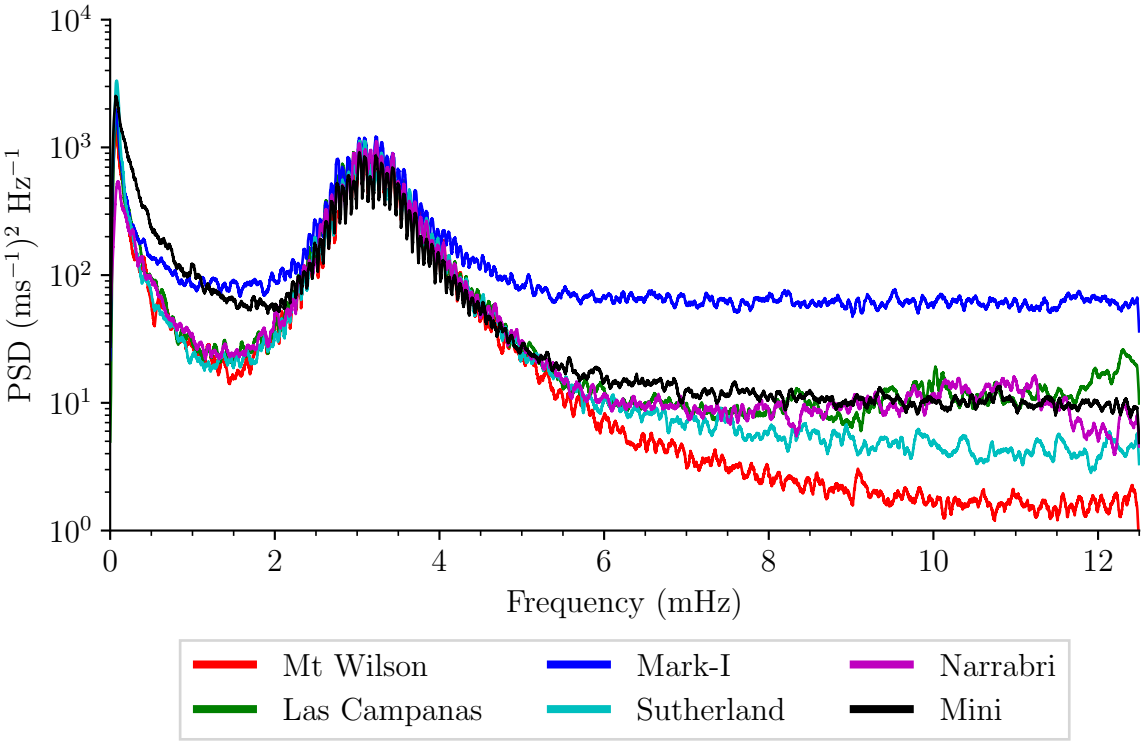

Figure 8.32: Comparison of data from an observing campaign over 2018. The power spectral density has been smoothed with a 50 µHz moving mean to aid clarity.

Table 8.3: Summer 2018 observing campaign performance metrics.

Site	FOM	Mean Noise $(\mathrm{ms}^{-1})^2\,\mathrm{Hz}^{-1}$	Noise Equivalent Velocity $\mathrm{cm\,s}^{-1}$ RMS
Mount Wilson	82.6	1.6	14.2
Sutherland	42.9	4.3	23.1
BiSON Mini	19.6	9.9	35.2
Narrabri	27.9	10.4	36.1
Las Campanas	25.0	13.5	41.1
Mark-I	5.5	61.6	87.7

mirrors and faulty guider at the site. From Table 8.3 we can see that Mount Wilson is now the best performing site, with a mean noise level of $1.6\,(\mathrm{ms}^{-1})^2\,\mathrm{Hz}^{-1}$ comparable to a noise equivalent velocity RMS of just $14.2\,\mathrm{cm\,s}^{-1}$. The prototype requires a factor of 6 reduction in noise power to equal the performance of Mount Wilson, and this is of the order expected to be obtained by reducing scintillation noise via faster polarisation switching. Mount Wilson can be considered to demonstrate the performance achievable when making use of both low-noise fibre optics and fast 100 Hz polarisation switching.

8.10 Conclusion

In this chapter we developed the necessary data acquisition and control system architecture required for automated instrumentation. Through Sections 8.2 to 8.7 we assessed the performance of all instrument components in isolation, from photodiode amplifier performance in terms of noise, bandwidth, and stability, through to considering the impact of changes in polarisation switching rate and vapour temperature fluctuations. Finally in Sections 8.8 and 8.9 we presented results following installation and commissioning at Izaña, Tenerife, and from a full summer observing campaign throughout 2018.

A summary of per-component instrumental noise levels determined throughout this chapter is given in Table 8.4. We can see that the overall noise level is easily dominated by atmospheric scintillation, as can be expected from ground-based photometric measurements. Whilst the reduced frequency polarisation switching of the new instru-

Table 8.4: Instrumentation noise levels.

	Detector Output		Noise Equivalent Power	
Noise Type	ASD	RMS	ASD	RMS
Electronic	$100\,\mu\mathrm{V}\,\sqrt{\mathrm{Hz}}^{-1}$	$4.2\,\mathrm{mV_{RMS}}$	$0.07\,\mathrm{pW}\,\sqrt{\mathrm{Hz}}^{-1}$	$3\,\mathrm{pW_{RMS}}$
Photon shot noise		$0.03\,\mathrm{mV_{RMS}}$		$0.02\,\mathrm{pW_{RMS}}$
Temperature	$285\,\mu\mathrm{V}\,\sqrt{\mathrm{Hz}}^{-1}$		$0.2\,\mathrm{pW}\,\sqrt{\mathrm{Hz}}^{-1}$	
Scintillation†	$200\,\mu\mathrm{V}\,\sqrt{\mathrm{Hz}}^{-1}$		$14.8\,\mathrm{nW}\,\sqrt{\mathrm{Hz}}^{-1}$	

† Voltage from low-gain transmission monitor.

ment does result in reduced performance on a daily cadence, over longer periods the new instrument demonstrates equivalent performance to the large-scale existing BiSON infrastructure allowing seamless integration. These results suggest we can move forward with a design that delivers low high-frequency noise, which will give excellent signal-to-noise in the core modes for solar-cycle studies, and good stability at low-frequency for detecting low-frequency modes to help further constrain the structure of the core. The new instrumentation achieves significant miniaturisation in comparison to the existing control systems. What was once a full-height rack of electronics is reduced to just three boxes approximately 4U (four rack units) in height, which allows the system to be easily deployed within existing infrastructure.

In the next chapter we will discuss development and automation of an inexpensive consumer-grade telescope mount, and how this can be used to add the final piece to a fully autonomous solar spectrophotometer.

Chapter 9

A Consumer-Grade Telescope Mount

Contents

9.1 Introduction

In this chapter we will consider the problem of acquiring and tracking the Sun. As we saw in Chapter 4, existing BiSON sites use either a large mount to point the instrument at the Sun directly, or observe via cœlostat mirrors. It will not be financially viable to deploy more sites in this manner due to the large housings required to weather-proof these systems. Given that the fibre-optic system developed throughout the project is very small and lightweight, a small inexpensive commercial off-the-shelf (COTS) mount will be mechanically more than sufficient. Unfortunately, in addition to their reduced load-bearing ability, small commercial telescope mounts tend to also suffer from missing features and reduced performance in comparison to their larger siblings.

Two key issues when considering attempting to automate a COTS telescope mount is access to a communications protocol with a published application programming interface (API), and the ability to keep track of where the mount is pointing when power is lost or the mount is moved manually by hand. Many COTS mounts either do not

have facility for computer control at all. Where an interface is provided the communications protocol is often proprietary to the manufacturer's own control handsets, and while the handsets themselves are typically ASCOM compatible (Denny, 1997) it is not obvious without actual testing what level of control is available in terms of relatively simple pointing commands or more advanced absolute motor speeds. Additionally, such mounts are intended to be manually aligned each time they are used, since they are expected to be packed away at the end of each observing session, and so they offer no facility to store a permanent calibration. Typically movement is tracked by counting stepper-motor steps once initial alignment has been completed, and so the calibration is lost as soon as the mount is moved by hand or the power is interrupted. Larger, more expensive, mounts employ absolute position encoders to avoid this problem.

The communication and control problem can be solved by simply choosing only manufacturers that open-source and publish their control API. The issue of alignment and recovery from power failure without user intervention is rather more tricky. Potentially a mount could be retro-fitted with absolute position encoders, or limit switches to indicate a home position, but this is known to be very challenging for precise absolute encoders and also does begin to move away from the COTS ideal. There are only so many modifications that can be made before it would be simpler and better to custom build a mount to our own requirements. It may be possible to use microelectromechanical systems (MEMS), such as the accelerometers used to control screen orientation in devices such as smartphones and tablets, to track the position of a telescope. These could easily be mounted on the telescope itself and avoid costly modifications to the mount.

Some background theory supporting the work in this chapter is available in Appendix C, where we review the standard astronomical algorithms for determining the position of the Sun at any time and location, and develop the rotation matrices required for a full three-dimensional pointing model of a telescope mount. In Section 9.2 we will assess the pointing precision achieved with MEMS sensors, and in Section 9.3 we will use a *Sky-Watcher HEQ5-Pro* COTS mount and the published control API to determine the

guiding precision that can be achieved. Finally in Section 9.4 we will consider options for a weatherproof enclosure that will be required for full automation.

9.2 MEMS Sensors

9.2.1 Determining attitude using a three-axis accelerometer

A microelectromechanical system (MEMS) accelerometer measures linear acceleration, and this includes the acceleration due to the Earth's gravitational field vector. Inside the device, a micro proof mass is suspended by restoring springs, and deflection of the proof mass due to acceleration is detected by measuring changes in capacitance between the proof mass and sensing plates. An embedded micro-controller digitises the signal and allows the data to be output over a digital serial bus. A total of three measurement axes are required in order to sense acceleration in all directions. The typical use case of an accelerometer is in smartphones and tablets, in order to detect the orientation of the device and allow the display to be rotated. Rotating a display to the correct orientation requires only a very coarse estimate of which way is "up". Perhaps such a device can provide an output with sufficient resolution to enable coarsely pointing a telescope? A thorough discussion of calibration of a three-axis accelerometer and extracting pitch and roll angles is given by Pedley (2013, 2015). We will now summarise the calibration process described by Pedley and go on to apply the techniques to an Analog Devices ADXL345 accelerometer fixed to the optical tube of an equatorially mounted telescope.

The Analog Devices ADXL345 accelerometer is factory calibrated to output values in units of g, with a precision of 4 milli-g per least-significant bit and a dynamic range of $\pm 16\,g$. The accelerometer measures both the gravitational field vector, and linear acceleration due to motion. Measurement accuracy of the gravitational field is reduced when the device is subject to additional external accelerations, however on a quasi-static design such as a telescope mount this will not be a concern. The factory calibration is a 6-term model providing gain and offset parameters for each of the three channels, and

this generally provides adequate results. Re-calibrating after the device is installed allows for greater precision to be achieved by correcting for thermal stresses introduced during soldering, and the convenience of allowing installation at a arbitrary angle – the device axes do not need to align with the system axes. A general 12-parameter calibrated output $\mathbf{G}_{12}$ can be defined in terms of the factory calibration $\mathbf{G}_f$ by (Pedley, 2015, eq. 35),

$$\mathbf{G}_{12} = \begin{pmatrix} G_{12x} \\ G_{12y} \\ G_{12z} \end{pmatrix} = \mathbf{W}\mathbf{G}_f + \mathbf{V} = \begin{pmatrix} W_{xx} & W_{xy} & W_{xz} \\ W_{yx} & W_{yy} & W_{yz} \\ W_{zx} & W_{zy} & W_{zz} \end{pmatrix} \begin{pmatrix} G_{fx} \\ G_{fy} \\ G_{fz} \end{pmatrix} + \begin{pmatrix} V_x \\ V_y \\ V_z \end{pmatrix} , \tag{9.1}$$

where the gain matrix $\mathbf{W}$ includes any rotation of the integrated circuit package and also corrects for all possible cross-talk interactions, and $\mathbf{V}$ are the channel offsets.

If M measurements are used for the 12-parameter calibration, then the i-th measurement at pitch angle $\theta[i]$ and roll angle $\phi[i]$ becomes (Pedley, 2015, eq. 36),

$$\begin{pmatrix} G_{12x}[i] \\ G_{12y}[i] \\ G_{12z}[i] \end{pmatrix} = \begin{pmatrix} W_{xx} & W_{xy} & W_{xz} \\ W_{yx} & W_{yy} & W_{yz} \\ W_{zx} & W_{zy} & W_{zz} \end{pmatrix} \begin{pmatrix} G_{fx}[i] \\ G_{fy}[i] \\ G_{fz}[i] \end{pmatrix} + \begin{pmatrix} V_x \\ V_y \\ V_z \end{pmatrix} = \begin{pmatrix} -\sin(\theta[i]) \\ \cos(\theta[i])\sin(\phi[i]) \\ \cos(\theta[i])\cos(\phi[i]) \end{pmatrix} , \tag{9.2}$$

which can be decomposed into three equations,

$$\begin{aligned} W_{xx}G_{fx}[i] + W_{xy}G_{fy}[i] + W_{xz}G_{fz}[i] + V_x &= -\sin(\theta[i]) \\ W_{yx}G_{fx}[i] + W_{yy}G_{fy}[i] + W_{yz}G_{fz}[i] + V_y &= \cos(\theta[i])\sin(\phi[i]) \\ W_{zx}G_{fx}[i] + W_{zy}G_{fy}[i] + W_{zz}G_{fz}[i] + V_z &= \cos(\theta[i])\cos(\phi[i]) \ , \end{aligned} \tag{9.3}$$

with residuals,

$$\begin{aligned} r_x[i] &= -\sin(\theta[i]) - W_{xx}G_{fx}[i] - W_{xy}G_{fy}[i] - W_{xz}G_{fz}[i] - V_x \\ r_y[i] &= \cos(\theta[i])\sin(\phi[i]) - W_{yx}G_{fx}[i] - W_{yy}G_{fy}[i] - W_{yz}G_{fz}[i] - V_y \\ r_z[i] &= \cos(\theta[i])\cos(\phi[i]) - W_{zx}G_{fx}[i] - W_{zy}G_{fy}[i] - W_{zz}G_{fz}[i] - V_z \ . \end{aligned} \tag{9.4}$$

If we now consider only the x-component, equation 9.4 can be simplified to (Pedley,

2015, eq. 41),

$$r_x = Y_x - \mathbf{X}\beta_x \ , \tag{9.5}$$

where r_x is the M-length array of residuals to the calibration fit, Y_x is the array of x-components of the gravitation field for the true measured angle, X is the matrix of accelerometer measurements, and finally β_x is the solution vector for four of the calibration parameters. The optimum least squares solution for β_x can be found by making use of the Normal Equations for least squares optimisation,

$$\beta = (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\mathbf{Y} \ , \tag{9.6}$$

such that (Pedley, 2015, eq. 46),

$$\beta_x = \begin{pmatrix} W_{xx} \\ W_{xy} \\ W_{xz} \\ V_x \end{pmatrix} = (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T Y_x = (\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T \begin{pmatrix} -\sin\theta[0] \\ -\sin\theta[1] \\ \cdots \\ -\sin\theta[M-1] \end{pmatrix} \ , \tag{9.7}$$

and similarly for the remaining two axes. In order to solve for 12-parameters the minimum number of measurements M is four, and these should be well distributed in the measurement space. More measurements give a more robust calibration. The calibrated accelerations are subsequently determined using equation 9.1.

Figure 9.1 shows a range of measurements made by an Analog Devices ADXL345 accelerometer fixed to the optical tube of an equatorially mounted telescope, whilst the mount was slewed slowly through its full range of motion. Both factory calibration and improved 12-parameter calibration are shown. The roll and pitch angles can be computed from the calibrated measurement matrix $\mathbf{G}_{12}$ using (Pedley, 2013, eq. 25–26),

$$\phi_{xyz} = \tan^{-1}\left(\frac{G_{12y}}{G_{12z}}\right) \ , \tag{9.8}$$

and,

$$\theta_{xyz} = \tan^{-1}\left(\frac{-G_{12x}}{\sqrt{{G_{12y}}^2 + {G_{12z}}^2}}\right) \ , \tag{9.9}$$

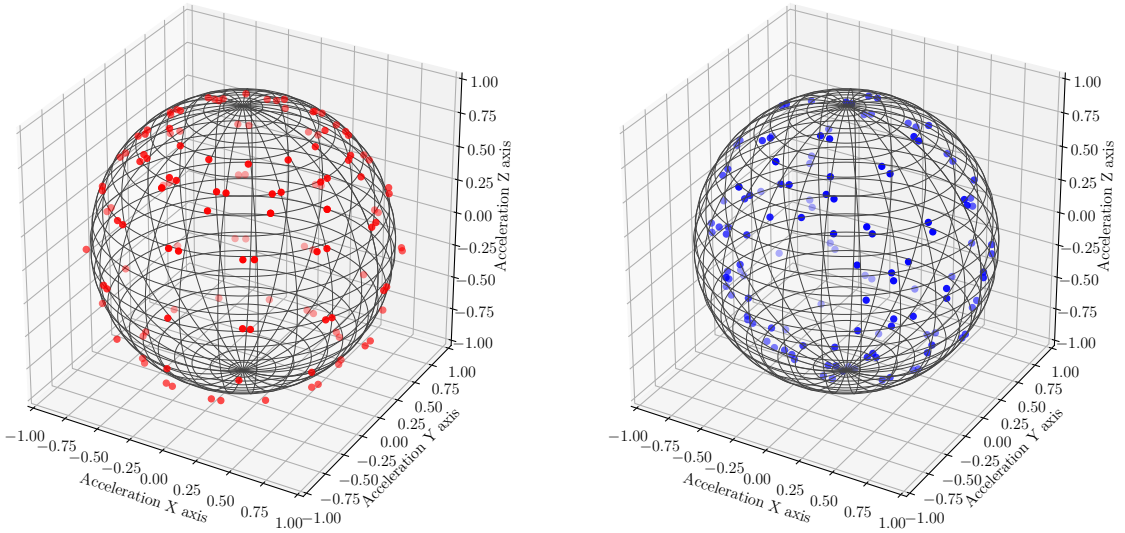

(a) Raw accelerometer data

(b) Calibrated accelerometer data

Figure 9.1: Data from an Analog Devices ADXL345 accelerometer fixed to the optical tube of an equatorially mounted telescope, slewed through its full range of motion. Panel (a) shows the raw data with standard factory calibration. Panel (b) shows the data after 12-parameter calibration. All points now sit on the surface of a sphere of radius 1 g.

where ϕ and θ are the measured roll and pitch angles respectively, and the subscript xyz notes that the angles are according to the aerospace rotation sequence $\mathbf{R}_{xyz}$ (see Appendix C.3 for more details).

Equation 9.8 becomes unstable when the telescope is pointing near the zenith, since the x-axis becomes aligned with the gravitational field vector and enters a condition known as Gimbal Lock. Any rotation in roll can no longer be detected and roll becomes undefined. A common work-around for this problem when using aerospace rotation sequence $\mathbf{R}_{xyz}$ is to modify equation 9.8 to include in the denominator a fraction μ of the x-axis measurement whilst remembering to maintain the sign of G_{12z} after taking the square root (Pedley, 2013, eq. 38),

$$\phi_{xyz} = \tan^{-1}\left(\frac{G_{12y}}{\pm\sqrt{{G_{12z}}^2 + \mu{G_{12x}}^2}}\right) , \tag{9.10}$$

such that ϕ is slowly driven to zero as the telescope approaches a vertical orientation. Figure 9.2 shows the resulting accuracy of the estimated angles from the calibration scan in Figure 9.1, with $\mu = 0.01$. The accelerometer achieves an accuracy of approximately $\pm 4°$

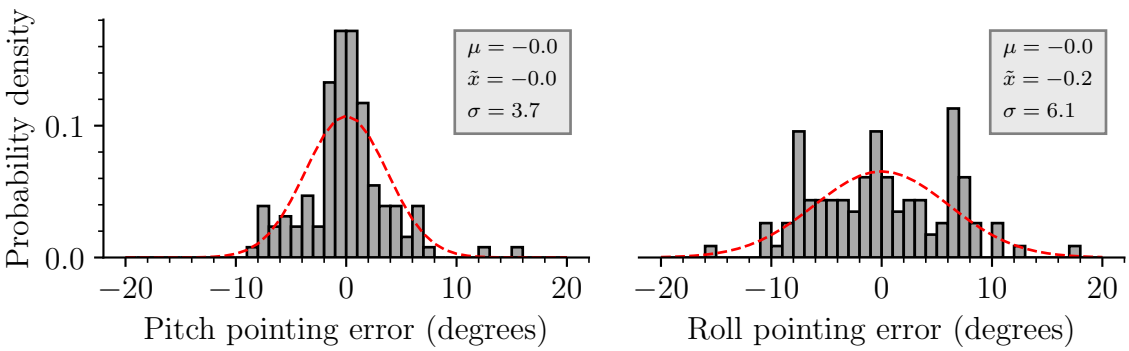

Figure 9.2: Accuracy of the Analog Devices ADXL345 accelerometer after 12-parameter calibration, comparing known reference orientations with the measured estimates. The dashed red line indicates the equivalent Gaussian for the measured mean and standard deviation.

in pitch and ±6° in roll. The reduced performance in roll is due to the above work-around at high pitch angles. This performance is adequate for a solar autoguider with a capture angle of a few degrees.

An additional ambiguity is caused while the telescope is vertical. It becomes impossible to determine if the telescope is on the east or west side of the pier. This ambiguity can be resolved by the addition of a second accelerometer, or even a simple tilt switch, mounted directly on the polar axis of the mount. This second sensor provides the necessary up/down signal to completely determine the telescope attitude. Mészáros et al. (2014) trialled this two-accelerometer technique by mounting two Freescale MMA8453Q accelerometers on their telescope at Konkoly Observatory located at the Piszkéstetö Mountain station. Mészáros et al. achieved fit residuals of 0.12° by making use of a custom jig for sensor calibration, with full temperature compensation and taking good care of power supply stability via a custom printed circuit board. Clearly with further work the results achieved here can be improved, and telescope attitude sensing with accelerometers is certainly viable.

The technique presented here works only with an equatorial telescope mount. With an altitude-azimuth mount, only changes in altitude (pitch) can be measured. Rotations in azimuth (yaw) are aligned with the gravitational field vector and cannot be detected, and a different sensor is required. Next, we will look at using a three-axis magnetometer

to detect the orientation of Earth's magnetic field and directly measure the telescope "heading" angle.

9.2.2 Determining heading using a three-axis magnetometer

In order to detect rotation about a vector parallel to Earth's gravitational field vector (i.e., yaw, heading, azimuth) we need to be able to detect rotation within Earth's magnetic field, otherwise known as an electronic compass. To do this we will make use of a Honeywell HMC5883L three-axis magnetometer, which uses magneto-resistive sensors to measure both the direction and the magnitude of Earth's magnetic field. The Honeywell HMC5883L magnetometer is factory calibrated to output values in units of µT, with a dynamic range of 800 µT, and an embedded 12-bit ADC providing a claimed 1° to 2° heading accuracy. As with the accelerometer, a magnetometer needs a final in-situ calibration in order to achieve the best precision. A thorough discussion of calibration of a three-axis magnetometer is given by Ozyagcilar (2015a,b,c,d). We will now summarise the calibration process described by Ozyagcilar and go on to apply the techniques to the Honeywell HMC5883L, including tilt-compensation based on earlier results from the Analog Devices ADXL345 accelerometer.

A magnetometer vector measured at the device $\mathbf{B}_d$ can be defined after arbitrary device rotation in terms of the local geomagnetic field vector $\mathbf{B}$ by (Ozyagcilar, 2015a, eq. 5),

$$\mathbf{B}_d = \mathbf{W}\mathbf{R}_x(\phi)\mathbf{R}_y(\theta)\mathbf{R}_z(\psi)\|\mathbf{B}\| \begin{pmatrix} \cos\delta \\ 0 \\ \sin\delta \end{pmatrix} + \mathbf{V} \ , \tag{9.11}$$

where ϕ, θ, and ψ are roll, pitch, and yaw angles as previously, δ is the magnetic inclination at the measurement location, $\mathbf{V}$ is the "hard-iron" offset vector, and $\mathbf{W}$ is the "soft-iron" gain matrix. So-called hard-iron offsets are magnetic fields generated by nearby permanent magnets, such as other components on the PCB, and motors in the telescope mount. These components are all in fixed positions and rotate with the device, and

so they appear as an additive magnetic field vector within the reference frame of the magnetometer. So-called soft-iron interference is due to temporary induction of magnetic fields in otherwise normally unmagnetised components, such as the sheet steel of the telescope mount and housing, caused by the geomagnetic field itself. Soft-iron effects are much more complicated to model since they depend on the orientation of the device within the geomagnetic field, and typically suffer magnetic hysteresis effects as the device rotates. The soft-iron matrix $\mathbf{W}$ is a 9-element matrix similar to that used during calibration of the accelerometer in equation 9.1, and in addition to calibrating soft-iron effects also calibrates rotation of the integrated circuit package and corrects for all possible cross-talk interactions and gain variations between channels. As with a typical analogue compass which works only when held level, the measurements from a magnetometer need to be derotated back to the flat plane where $\phi = \theta = 0$, in addition to removing offsets. Rearranging equation 9.11 for ψ we get (Ozyagcilar, 2015a, eq. 6),

$$\begin{aligned}\mathbf{R}_z(\psi)\|\mathbf{B}\| \begin{pmatrix}\cos\delta\\0\\\sin\delta\end{pmatrix} &= \begin{pmatrix}\cos\psi & \sin\psi & 0\\-\sin\psi & \cos(\psi) & 0\\0 & 0 & 1\end{pmatrix}\|\mathbf{B}\| \begin{pmatrix}\cos\delta\\0\\\sin\delta\end{pmatrix}\\ &= \mathbf{R}_y(-\theta)\mathbf{R}_x(-\phi)\mathbf{W}^{-1}(\mathbf{B}_d - \mathbf{V}) \ ,\end{aligned} \tag{9.12}$$

and so (Ozyagcilar, 2015a, eq. 9),

$$\begin{pmatrix}B_{fx}\\B_{fy}\\B_{fz}\end{pmatrix} = \begin{pmatrix}\cos\psi\cos\delta\|\mathbf{B}\|\\-\sin\psi\cos\delta\|\mathbf{B}\|\\\sin\delta\|\mathbf{B}\|\end{pmatrix} = \mathbf{R}_y(-\theta)\mathbf{R}_x(-\phi)\mathbf{W}^{-1}(\mathbf{B}_d - \mathbf{V}) \ , \tag{9.13}$$

where $\mathbf{B}_f$ is equal to the magnetometer measurements with both soft-iron and hard-iron effects removed, and derotated to a flat plane where the z-component B_{fz} is equal to $\|\mathbf{B}\|\sin\delta$. The yaw, or compass heading, or azimuthal angle is found from (Ozyagcilar, 2015a, eq. 10),

$$\psi = \tan^{-1}\left(\frac{-B_{fy}}{B_{fx}}\right) \ , \tag{9.14}$$

and this simply requires the addition of the known local magnetic declination to convert

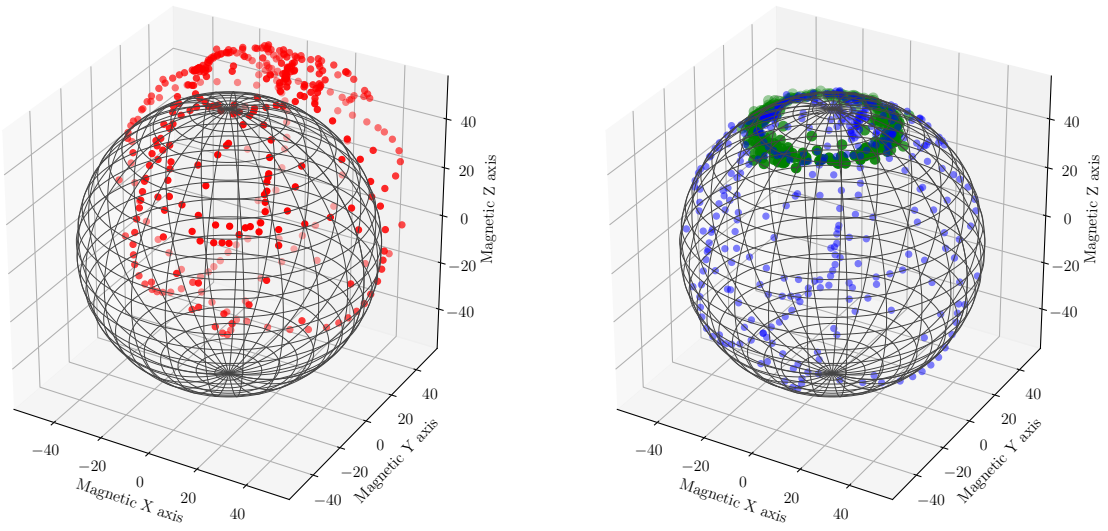

(a) Raw magnetometer data

(b) Calibrated magnetometer data

Figure 9.3: Data from a Honeywell HMC5883L magnetometer slewed randomly through the full range of motion. Panel (a) shows the raw data with a hard-iron offset. In panel (b), the blue dots show the data after 4-parameter calibration removing the hard-iron offset. All points now sit on the surface of a sphere centred at the origin, with radius equal to the local geomagnetic field strength in Birmingham calibrated at 54.5 µT. The green dots show the same data after derotation to the flat plane where $\phi = \theta = 0$ and the z-component B_{fz} is equal to $\|\mathbf{B}\| \sin \delta$.

from magnetic north to true north.

In many cases the soft-iron effects are insignificant, and only the hard-iron offsets dominate. This allows a simplification during calibration since only four parameters need to be determined, and these are the magnitude of the geomagnetic field strength $\|\mathbf{B}\|$, and the three components of the hard-iron vector $\mathbf{V}$. The soft-iron matrix $\mathbf{W}$ becomes the identity matrix. With these assumptions we can follow the same calibration procedure as for the accelerometer, by developing a performance function to be minimised by optimising the calibration fit and again using equation 9.6 to solve the fit through matrix algebra.

Figure 9.3 shows a range of measurements made by a Honeywell HMC5883L magnetometer slewed randomly through the full range of motion. It is clear from Figure 9.3(a) the data do indeed suffer from a hard-iron offset, which moves the data away from the origin. After applying the 4-parameter calibration, shown by the blue dots in Figure 9.3(b),

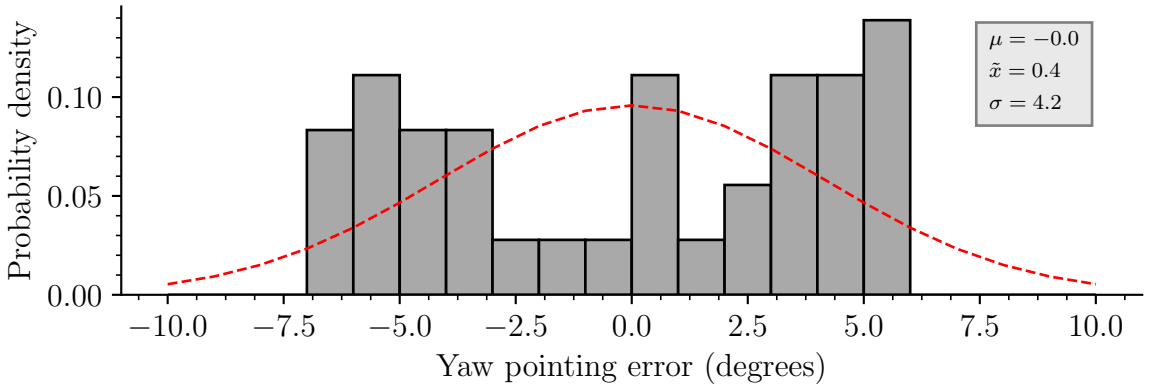

Figure 9.4: Accuracy of the Honeywell HMC5883L magnetometer after 4-parameter calibration, comparing known reference orientations with the measured estimates. The dashed red line indicates the equivalent Gaussian for the measured mean and standard deviation.

the data now sit on the surface of a sphere centred at the origin, with radius equal to the local geomagnetic field strength in Birmingham calibrated at 54.5 µT. Our assumption that the soft-iron effects are insignificant is shown to be true, since these would have the effect of distorting the data away from a perfect sphere and into an ellipse. The green dots show the same data after derotation, by making use of the pitch and roll angles determined by the accelerometer, to the flat plane where $\phi = \theta = 0$ and the z-component B_{fz} is equal to $\|\mathbf{B}\| \sin \delta$. These data are now calibrated and ready for the heading to be extracted using equation 9.14.

Figure 9.4 shows the resulting accuracy of the estimated angles, with the magnetometer achieving an accuracy of approximately ±5°. The distribution is bimodal due to slight differences in gain between the magnetometer channels, or possibly the influence of small soft-iron effects. With further work, there should be no reason why the manufacturers stated performance of 1° to 2° resolution should not achieved. Even with this simple calibration, and as with the accelerometer, this performance is adequate for a solar autoguider with a capture angle of a few degrees.

Having shown that it is certainly possible to obtain sufficiently accurate absolute pointing information from an inexpensive telescope mount, we will now go on to consider fine guiding of a mount once the target has been acquired by the coarse pointing.

9.3 Solar auto-guiding using computer vision

A range of consumer-grade amateur telescope mounts were considered for testing, and the *Sky-Watcher HEQ5-Pro* was selected and trialled at the Mount Wilson (Hale) Observatory in 2016 September (Hale, 2016d). The HEQ5 Pro is a mid-range consumer mount costing less than £1000, but offering good stability and tracking performance with a stated positioning accuracy up to $1'$. Like most mounts in this price range, it is a so-called GOTO mount and is supplied with a computerised handset to enable ease of alignment and use. The SynScan handset allows one-star, two-star, or three-star alignment procedures to be completed by the user, after which the SynScan handset can automatically position any known target within the field-of-view of the attached telescope eyepiece.

It is possible to control the mount from a PC by connecting to the SynScan handset via an RS-232 serial port cable. Of particular interest with the Sky-Watcher mount is that it is also possible to bypass the handset and connect directly to the micro-controller allowing direct querying of the position encoders and control of the motors. Sky-Watcher publish the control API and make writing a custom control library rather easy. With many other manufacturers this could only be possible, if at all, by reverse engineering the motor controller communications protocol. This is important since the SynScan handset is intended for typical night-time astronomy and does not include features for aligning, tracking, or guiding on the Sun. Regarding the position encoders, as expected these are not absolute encoders and are simply incremental counters based on stepper-motor steps. Due to this, the mount has no absolute position knowledge when it is first powered, and any subsequent alignment is lost if the mount is moved manually without making use of the motors. Throughout the rest of this section we will assume that the mount has already been aligned by making use of the techniques developed earlier in this chapter.

The Sky-Watcher HEQ5-Pro was configured with a binocular arrangement consisting of a *Starlight Express Superstar* CCD camera for guiding, and a fibre-pickup for data acquisition. It was temporarily installed for testing at the Mount Wilson 60 foot tower, shown in Figure 9.5. Discussion of data from the temporarily fibre-converted Klaus spec-

Figure 9.5: Sunrise at the Mount Wilson (Hale) Observatory. The Sky-Watcher HEQ5-Pro was configured with a binocular arrangement of a CCD camera for guiding, and a fibre-pickup for data acquisition. This image won the Gatsby Foundation *#TechniciansMakeitHappen* competition, a campaign to raise the public profile of STEM technicians such as the author, and encourage young people to choose a technical career (Hale, 2017e).

trophotometer is beyond the scope of this chapter – for more on fibres see Chapter 7 starting on page 139. The mount was guided throughout the day using a custom algorithm developed in Python and run on a Raspberry Pi single-board computer. The Starlight Express Superstar autoguider camera has a 6.4 mm by 4.75 mm CCD with 4.65 µm square pixels in a 1392 by 1040 array. When coupled with an 80 mm focal length objective lens this produces an approximate field of view of 4.6° by 3.5°, where each pixel has about 12″ field of view. The Sun has an extent of about 32′ and so produces an image about 160 pixels in diameter on the camera. Two filters were used to bring the image within the dynamic range of the CCD chip. These were a neutral density filter with optical density of 5 (i.e, transmission of approximately 10^{-5}), and a bandpass filter with 10 nm width centred at 760 nm. Images were read from the camera as frequently as possible, approximately once per second. The image exposure time was 79 ms, but the cadence is

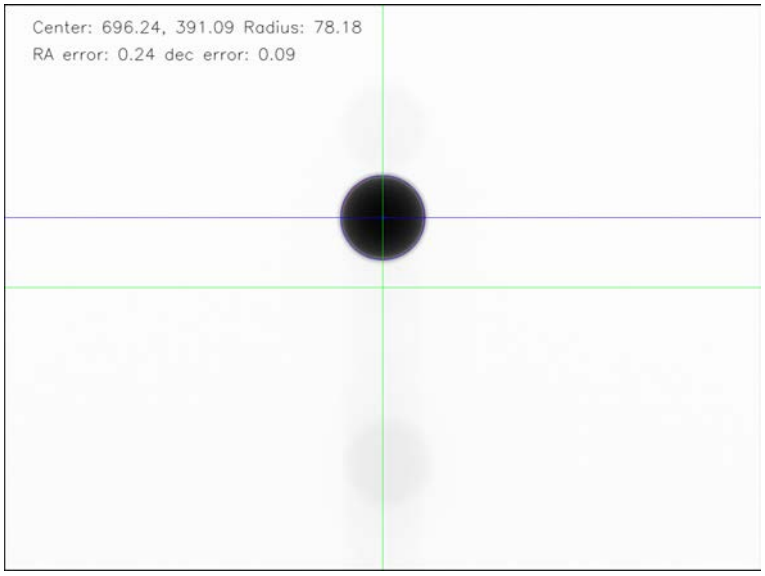

Figure 9.6: Guider image processing with OpenCV. The image has been inverted for printing. Values shown are in pixel units.

restricted by the CCD read-out time and USB transfer rate. An example of the image captured by the camera is shown in Figure 9.6.

Images from the camera were processed to determine the position of the Sun to sub-pixel accuracy using the OpenCV (Open-source Computer-Vision) library. The centroid position of the solar disc was found using OpenCV by applying a black and white threshold, finding the contours in the image, and subsequently reading out the contour centroid position. The guiding error was determined by comparing the current solar centroid position with a desired target value. The position error for each axis was then passed through a proportional–integral–derivative (PID) control loop feedback mechanism in order to determine the correct mount drive rate, and the mount motors updated with the new drive rate using the published Sky-Watcher API. The PID control algorithm uses the same code as developed earlier for temperature control in Section 8.7 on page 210, where we saw that a servo feedback system can be defined simply by,

$$\text{Output} = K_P\, e(t) + K_I \int e(t)\, dt + K_D \frac{\mathrm{d}e(t)}{\mathrm{d}t} \ , \qquad \text{(8.12 revisited)}$$

where e is the control error defined as the desired setpoint minus the current value, and K_P, K_I, and K_D are the proportional, integral, and derivative coefficients. The solar

guiding algorithm makes use of only the proportional and integral parameters.

Guider telemetry was logged in order to assess the performance and stability. Results from 2016 September 25 are shown in Figures 9.7 for the RA axis and 9.8 for the declination axis. The green traces log the position error in pixels, while the red, blue, and black traces log the proportional, integral, and total motor drive rate respectively in microsteps per second. The zero point of the integral RA drive rate is offset based on a standard solar tracking rate of 105 microsteps per second. The jumps in RA position error shortly after sunrise and shortly before sunset are due to backlash in the RA drive-train as the telescope weight shifts on the axis. Both axes have a minimum drive rate of 6 microsteps per second, and this is visible in the declination drive rate, where the PID control becomes essentially an on/off switch at minimum-rate around zero-error. The non-zero mean error in declination is possibly due to mechanical issues with the axis. At the end of the day the movement of the declination axis was found to be very notchy. The movement cleared as the axis was manually exercised several times through the full range of motion, but is still a concern regarding long term reliability.

The distribution of position-error for both axes, calibrated in arcseconds, is shown in Figure 9.9. The developed control algorithm achieves performance on both axes to better than $\pm 8''$ when logged at an approximate 1 Hz cadence over an entire day. This is more than sufficient for our requirements when considering the target object is over $1800''$ in diameter. The performance also compares favourably with data captured using the standard BiSON full-size equatorial mount in Carnarvon on 2016 August 16, where the RA guiding accuracy was measured to be $\pm 9''$ and declination $\pm 4''$ over a whole day. Historically, some BiSON telescope mounts have been believed to achieve $1''$ guiding but there is no documentation showing this performance. When an image of the Sun of formed within the vapour cell, guider movement results in Doppler velocity offsets due to variations in effective optical depth, and so precise guiding is necessary. The pointing sensitivity of Spectrophotometer-J in Carnarvon was estimated during a site visit to be $0.018\,\mathrm{m\,s^{-1}}$ per arcsecond (Barnes and Miller, 2009). When an image of the

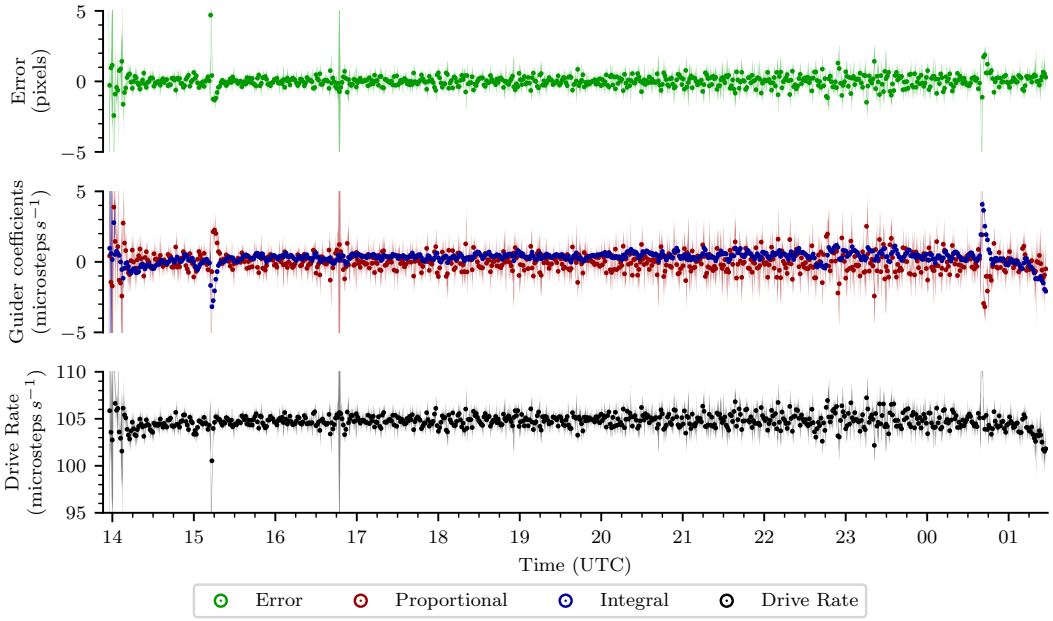

Figure 9.7: Guider RA performance on 2016 September 25. Each dot represents the median of 50 points, and the banding represents ±1 sigma on each median value.

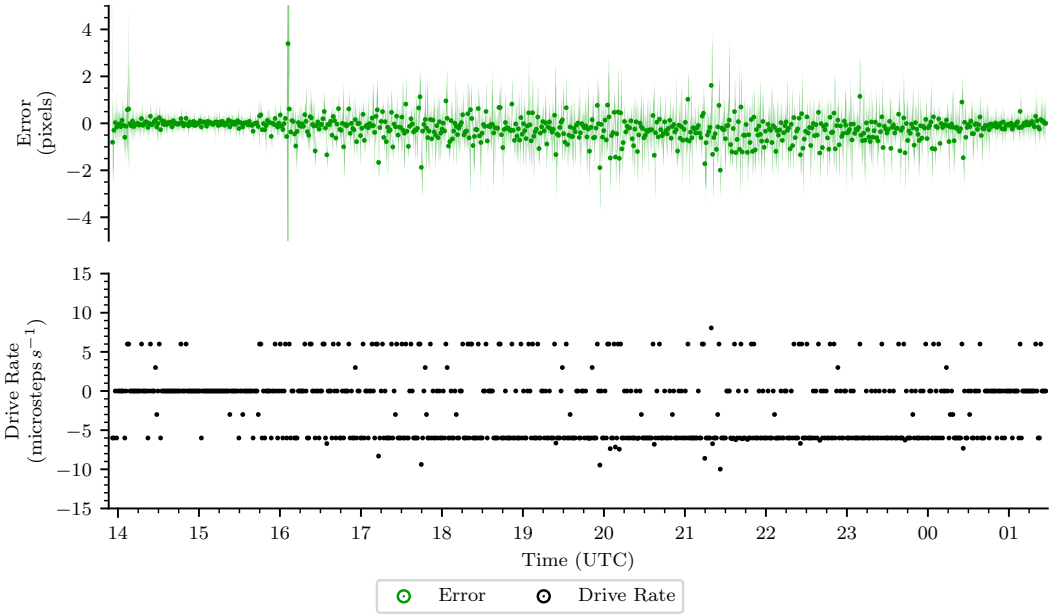

Figure 9.8: Guider declination performance on 2016 September 25. Each dot represents the median of 50 points, and the banding represents ±1 sigma on each median value.

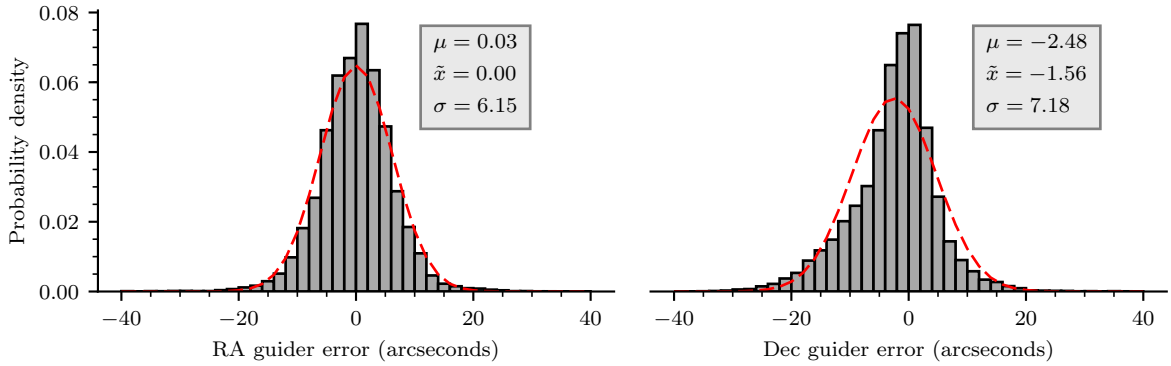

Figure 9.9: Guider performance histogram on 2016 September 25. The dashed red line indicates the equivalent Gaussian for the measured median and standard deviation.

Sun is instead focused onto the end of a fibre, the output of the fibre is uniform regardless of the absolute position of the image provided that the image falls entirely within the core of the fibre. The constraint on guiding accuracy is relaxed and guider noise is considered to be eliminated completely.

9.4 Enclosure and Weatherproofing

The telescope mount will require a weatherproof housing to keep it safe and secure when not in use. Whilst there are many options for amateur observatory enclosures, these are all typically large enough to hold at least one or two people. They are intended to be installed in a back garden, where the amateur astronomer sits outside when using their telescope. This will be far too large for our purposes. We need to keep the physical footprint as small as possible to maximise the ease of installation. Our only requirement is that the enclosure completely contains the mount, and so it can be little bigger than the mount itself. For this we have no choice but to make use of a bespoke design.

Figure 9.10 shows one simple option for an enclosure. This design uses a perspex dome rather than a mechanical opening, enabling observations to be made directly through the enclosure itself. This is similar to the design in use for the solar light feed to the HARPS-N spectrograph at La Palma (Dumusque et al., 2015). The advantage of this design is that it is simple, requiring no motors or weather sensors. There is of course concern

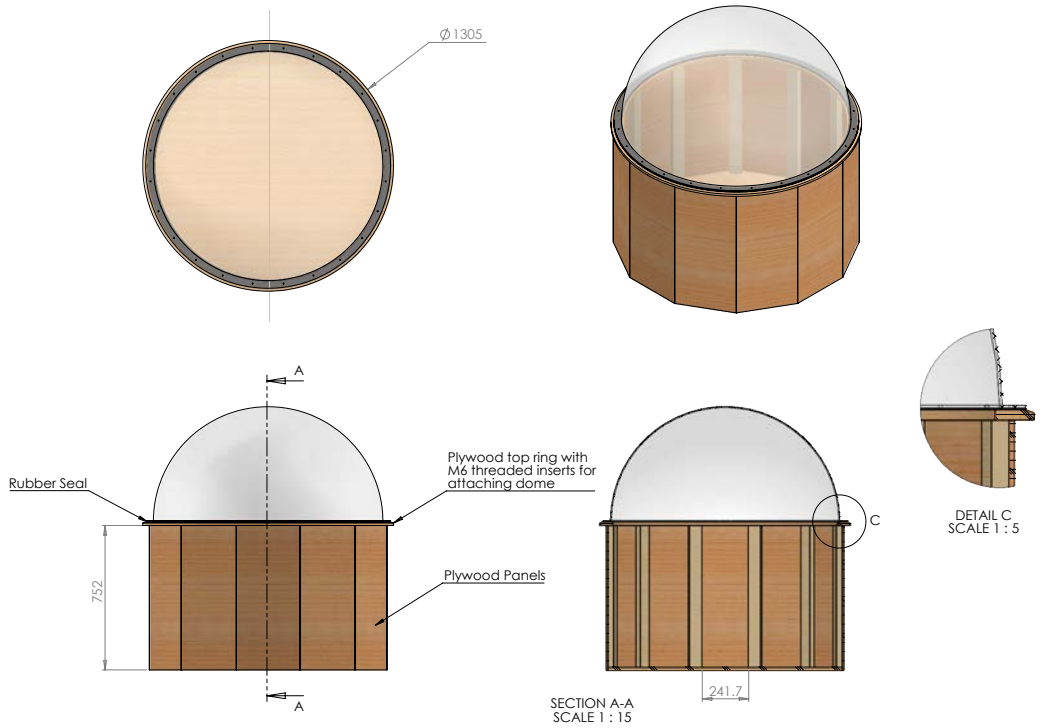

Figure 9.10: Perspex dome enclosure option. Drawing credit: Stephen Brookes

with keeping the perspex clean and free from scratches when used without a regular maintenance schedule and on-site observers.

Figure 9.11 shows two alternative designs that are mechanically simple, and allow the telescope to observe through an opening in the typical fashion. This style of enclosure is more complicated due to the requirement of mechanical drive systems and weather sensors. However, the automation requirements are trivial and the only real concern is the additional cost over a non-opening style enclosure. The final design and construction of a suitable enclosure is subject to further work.

9.5 Conclusion

In this chapter we developed all of the necessary algorithms and software required to fully automate a COTS consumer-grade telescope mount. In Section 9.2 we showed

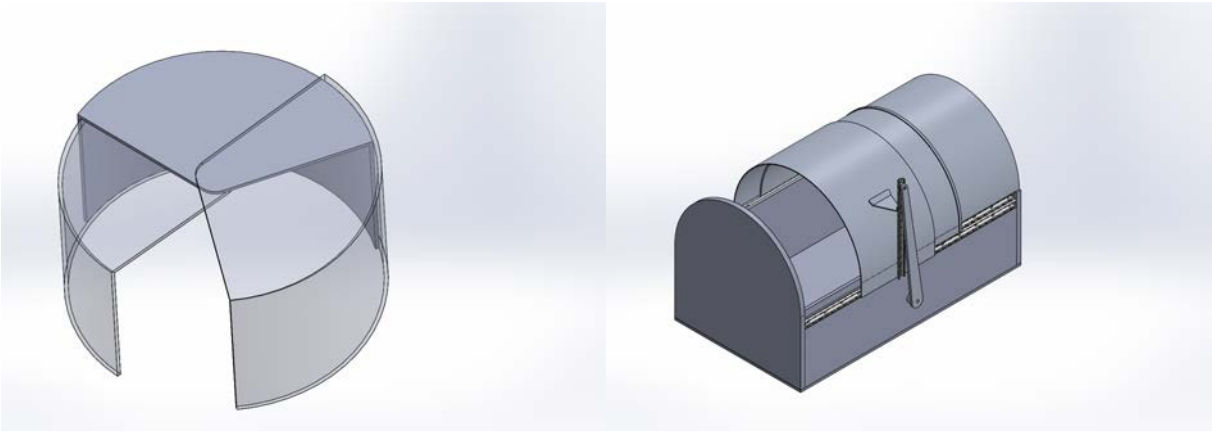

Figure 9.11: Mechanical enclosure options. Drawing credit: Stephen Brookes

that it is possible to point a low cost mount to within a few degrees by making use of absolute attitude information from low-cost MEMS sensors. with an expectation of further improving the performance with better calibration and more development. In Section 9.3 we demonstrated solar guiding capability to better than $\pm 8''$, which is more than sufficient for our requirements when considering the target object is over $1800''$ in diameter. Finally, in Section 9.4 we considered some initial options for a weatherproof enclosure that is required for full automation.

In addition to the control system and electronics miniaturisation developed in Chapter 8, in this chapter we have achieved another significant miniaturisation success. With the use of a small mount and enclosure, it is possible to easily and inexpensively commission new instrumentation on the roof of an existing building. This avoids the need for expensive and time consuming construction of dedicated observatories.

Since this work was completed, Sky-Watcher have released their own inexpensive solar mount making use of MEMS accelerometers. The Sky-Watcher Solarquest AltAz mount includes *Heliofind Solar Alignment Technology* to provide semi-automation and make it particularly easy to use. The only user controls on the mount are an 8-way slide switch and a power button. When the mount is turned on, it acquires a GPS fix to provide location and time, from which it can calculate the expected position of the Sun. The mount drives to the expected altitude determined using an accelerometer, and then slews clockwise until

the Sun is acquired in the built-in quadrant-photodiode based autoguider. The 8-way slide switch can then be used to adjust the final guiding position offset. In order to prevent the attached cables and fibre becoming wrapped around the pier, it is necessary to ensure the mount is pointing to the left of the Sun at startup and so limit the amount of clockwise rotation required to acquire the Sun. The mount has no azimuth alignment ability, such as a magnetometer, and so it cannot be started in cloudy conditions where the Sun is not visible in the autoguider, although it does continue tracking through clouds once the initial acquisition has been completed. The Solarquest mount would be perfect meeting all our requirements with its built-in guiding ability if were not for one missing feature – there is no facility for computer control. This means that there is no way to handle automatic startup and shutdown, and so an observer is always required to be present. The Solarquest mount is perfect for a site with daily observers such as Mount Wilson. A Solarquest mount was installed at Mount Wilson during the fibre conversion in 2018 June in order to bypass the damaged mirrors and faulty guider at the site (Hale, 2018b), and is showing excellent performance.

In the next chapter we will review the work completed throughout this project, and summarise the findings.

Chapter 10

Conclusions and Further Work

The aim of this project was to ensure the continued operation and possible expansion of the Birmingham Solar Oscillations Network (BiSON) through development of a new instrumentation paradigm known as BiSON-Mini. In Chapter 1 we considered the structure of the Sun and the role of BiSON in the application of helioseismology. We then went on to discuss the processes and the formation of the oscillation modes and the importance of long term synoptic observations of the Sun, both in terms of monitoring changes in the 11 year solar cycle, and the difficulties involved in observing the low-frequency oscillations below 1000 µHz that are most sensitive to conditions within the solar core. Observational projects, such as developed throughout this thesis, are essential if solar g-modes are to be detected and confirmed.

In Chapter 2 we gave a brief overview of the history of observational helioseismology, and the contribution made by the University of Birmingham as the field developed in the 1970s. We then went on to describe qualitatively the operation of the existing BiSON spectrophotometers in Chapter 3. We assessed the historic network performance in terms of duty cycle and noise levels in Chapter 4, and in terms of atmospheric extinction in Chapter 5. This initial work enabled us to establish a performance baseline to be matched, or ideally exceeded, when developing prototype miniaturised instrumentation to complement the existing network.

With a performance baseline described, in the second half of this thesis we then went on to develop prototype instrumentation that meets the requirements for simplicity and reduced physical footprint. We began in Chapter 6 by modelling a vapour reference cell from first principles, and we validated the accuracy of the model using a tunable diode-laser probing a standard BiSON reference cell. We determined that for a typical BiSON 15 mm cube vapour cell, the ideal stem temperature is approximately 88 °C when the scattering detectors have a 6 mm diameter aperture adjusted to observe the centre of the cell. This result agrees with expectation based on experience with existing BiSON infrastructure. Many BiSON instruments run at a stem temperature of 90 °C to 100 °C and a cube temperature of 110 °C to 120 °C. The variation is dependent on the point at which the temperature is measured, and the exact optical configuration of the scattering detectors in each spectrophotometer.

In Chapter 7 we developed an end-to-end model of the entire system, from the output of the Sun, through Earth's atmosphere, through each component of the spectrophotometer, and finally onto the photodiode detectors. We embedded the vapour cell model from Chapter 6 at the heart of the wider system model. We determined that the scattering detectors will receive between about 0.35 nW and about 1.4 nW of light depending on the line-of-sight velocity offset, and the transmission monitor about 0.07 mW under typical site conditions determined from Chapter 4 and Chapter 5. We also assessed the performance of a typical Pockels-effect cell as used in existing BiSON instrumentation, where it was found that the variable retarder can switch polarisation states in under 1 ms, and can easily achieve switching rates of the order of 100 Hz. A Pockels-effect cell requires a high-voltage high-bandwidth power supply, and these are currently fabricated in-house from a bespoke design. While off-the-shelf Pockels cells are available, they are intended for use with lasers. This means that they do not have sufficiently wide aperture or field of view to replace our existing small stock of custom Pockels cells. Additionally, the necessary high-voltage driver units have very limited off-the-shelf availability. An off-the-shelf liquid-crystal variable retarder and driver that operates at low voltage was trialled.

Unfortunately such LCDs are much slower to respond, taking up to 27 ms to transition, when heated to a temperature of 50 °C. At these low slew rates the new instrumentation is not able to switch polarisation states at 100 Hz and the reduced switching rate is detrimental to the daily data quality due to an increased level of atmospheric scintillation noise. The consequence of this is a higher background noise level within the resulting power spectrum, and so greater difficulty in resolving low amplitude modes of oscillation. We will discuss data quality further in a moment. These parameters were used to inform the subsequent design of the detectors, amplifiers, and the data acquisition system.

In Chapter 8 we discussed the data acquisition system, and analysed through circuit simulation the noise levels and slew rates of the existing BiSON photodiode transimpedance amplifiers. We then went on to design a new data acquisition system based around a Raspberry Pi single board computer and several 24-bit ADCs. It is here that a significant miniaturisation success is found, where a full-height rack of electronics is reduced to just three boxes approximately 4U (four rack units) in height. We used the modelled results from Chapter 7 to inform the amplification and bandwidth requirements of the new detectors. All noise sources were assessed in isolation either through simulation or direct empirical measurement. The resulting performance specification has previously been undocumented for any BiSON instrumentation. By far the dominant noise source is that from atmospheric scintillation. This is not an unexpected result – ground-based photometry is hard. However, the −3 dB point of scintillation noise decay typically occurs at around 5 Hz, and signal to noise ratios of 100:1 are achieved when considering the five-minute modes of oscillation. The −6 dB point typically occurs at between 5 Hz to 15 Hz, and so even small improvements in switching rate give a potentially large improvement in noise level. The −10 dB point typically occurs above 20 Hz.

Finally in Chapter 9 we considered the task of tracking the Sun and collecting light into the spectrophotometer. Improvements introduced here produce the single largest reduction in footprint of the whole project. Traditional BiSON instrumentation is housed in a 12 foot diameter dome and held on a large equatorial mount. By making use of

modern and inexpensive consumer-grade amateur "GOTO" telescope mounts, along with again a Raspberry Pi single board computer running bespoke automation software, it is possible to reduce the size of the total housing to something little larger than the mount itself. We demonstrated that inexpensive MEMS inertial sensors can be used to coarsely point the telescope and acquire the Sun within the moderate field of view of an autoguider CCD camera. Using a Sky-Watcher HEQ5-Pro off-the-shelf mount and the published control API, we demonstrated solar guiding capability to better than $\pm 8''$ when logged at an approximate 1 Hz cadence over an entire day. While this may be poor by astronomical standards, it is similar to that achieved with existing BiSON mounts, and is more than sufficient for our requirements. The weatherproof housing for such a mount is subject to further work, with several options currently under consideration. The use of a small mount and enclosure avoids the need for expensive and time consuming construction of dedicated observatories. It will be possible to easily and inexpensively commission new instrumentation on the roof of existing buildings, simplifying potential expansion of the BiSON network.

The new instrumentation was commissioned during 2017 at Izaña, Tenerife, and shared the light feed from the Mark-I cœlostat. Comparison of data from the entire 2018 summer observing campaign with other contemporaneous BiSON sites shows similar performance to existing BiSON infrastructure. On a daily cadence the reduced polarisation switching rate does result in a higher noise level due to the increased effect of atmospheric scintillation noise. The difference is equal to the expected increase in scintillation noise when moving from 100 Hz to 5 Hz polarisation switching. When concatenated into longer time series the effect of scintillation noise is reduced and performance becomes equivalent to other sites. When techniques from BiSON-Mini are applied to a spectrophotometer with 100 Hz polarisation switching, as shown with parallel testing at Mount Wilson, there is potential for almost an order of magnitude reduction in noise level over the best historic BiSON performance.

Throughout this thesis we have demonstrated that it is certainly possible to achieve the

stated aims by assembling many commercial off-the-shelf components. The vapour cell, oven, and magnets are still bespoke components due to issues discovered with commercial alternatives. In Chapter 7 we used the vapour cell model from Chapter 6 to assess the expected performance of an off-the-shelf Thorlabs GC25075-K cell and GCH25-75 oven. The typical use case for these long cylindrical cells is in Saturated Absorption Spectroscopy, where absorption within the cell is used to provide a stable locking reference for a tunable diode laser at the centre of an atomic transition. Our use case requires the ability to detect the light scattered from the vapour cell, and unfortunately the mechanical construction of these types of cells does not permit this to be done easily. At high temperatures with a large scattering cross-section, all of the light is scattered within the first oven making it impossible to reach a detector. At lower temperatures with a reduced scattering cross-section and so reduced scattering intensity, it is necessary to collect scattered light along the entire central length of the cell. This is difficult to achieve while simultaneously ensuring that detection of non-resonantly scattered light is blocked. In order to make successful use of a typical off-the-shelf vapour cell it would require a bespoke oven and magnet housing such as that designed for the GOLF-NG mission. Even with such a housing, it is not clear that an off-the-shelf cell would work since the filling-stem is cut short, preventing a cold section from being maintained well away from the scattering windows. This means that the coldest point on the cell occurs in the middle, between the two ovens, and so this is where the solid potassium forms right where we need the glass to be clear to observe scattering. The group no longer have access to sufficient expertise to fill our own vapour cells or design custom magnets, and so with the desire to move away from custom components this remains a concern and is the subject of further work.

It was a design criterion to avoid bespoke components, and whilst we do in fact still have some custom designs this does not necessarily result in increased risk to the mission. When relying on custom components made either in-house or in partnership with industry, there is significant risk that the skills or manufacturing ability will be lost. However, when

relying on commercial products there is instead a risk that the manufacturer will simply discontinue the product with no guarantee that any replacement will be able to drop in without forcing significant and unexpected changes. When swapping in-house risk for commercial risk, it is not clear that there has actually been a reduction in the level of risk to the mission. The most favourable situation is what we have achieved here, where relatively simple components are sourced inexpensively off-the-shelf, and mission-critical components remain a bespoke design that can be outsourced at greater cost to one of many industrial partners.

Projects such as BiSON, and the extended observational network potential of BiSON-Mini, provide the necessary low-frequency data that probes the structure and rotation of the solar core, and continues to improve our knowledge of the dynamo that drives the solar cycle. Refreshing the existing ageing hardware will ensure the ongoing programme of solar activity cycle observations is successful. Long term synoptic observations of the Sun are important if we are to detect p- and g-modes with frequencies below 1000 µHz, and so improve constraints on the structure of the deep interior and core.

Early predictions for the amplitude and timing of Cycle 24 produced many disagreements with little consensus (Pesnell, 2008), and this has led to substantial progress over the last decade in terms of understanding and predicting the solar cycle – see, e.g., Yeates et al. (2008); Karak and Nandy (2012); Muñoz-Jaramillo et al. (2013). With the suggestion that we may be leaving a grand maximum and entering a period of reduced solar activity, it is essential that observations continue as we go into Cycle 25 in order to obtain the necessary data required to test these updated models. Determining the accuracy of models and predictions is important for assessing the long-term impact on climate change and space weather.

The next stage of BiSON-Mini is to develop the prototype at Izaña into a fully automated instrument, replacing the original Mark-I spectrophotometer (Hale, 2017b,c). Once fully operational the design will be ready for commissioning of new nodes to extend the current BiSON network. During work on the fibre prototype, instruments at both

Carnarvon and Mount Wilson were converted to fibre using techniques learned throughout this project (Hale and Ross, 2018; Hale, 2018b), and these conversions have resolved many ongoing issues that were threatening continued operation at the sites. Following the work and results from this project, BiSON is certainly well placed to observe the next solar cycle and beyond.

Appendix A

Optical Theory

Contents

A.1 First-order Optics

All optical systems are based on the refraction or reflection of light. When a ray of light propagates from one material to another, its direction of travel changes. This is because the speed of light is different for different materials. The measure of the change, the *index of refraction*, is inversely proportional from the speed of light in the material to the speed of light in a vacuum. Therefore, the index of refraction in a vacuum is exactly unity and is dimensionless. The index of refraction of air is 1.00029, usually approximated to unity. For water it is 1.33, and for common glass 1.51. The index also depends on the wavelength of light.

Snell's Law states that the ratio of the sines of the angles of incidence and refraction is equal to the reciprocal of the ratio of the refractive indices,

$$\frac{\sin\theta_1}{\sin\theta_2} = \frac{n_2}{n_1} \ , \qquad \text{(7.1 revisited)}$$

where each angle θ is measured from the normal to the boundary, and n the refractive index of each medium.

Carl Friedrich Gauss showed that an optical system can be characterised by a set of cardinal points from which it is possible to calculate optical properties. In so-called Gaussian optics the assumption is made that angles of deviation are small and so the approximation $\sin\theta = \theta$ can be made. Equation 7.1 then simplifies to

$$\frac{\theta_1}{\theta_2} = \frac{n_2}{n_1} \ . \tag{A.1}$$

A.1.1 Thin Lenses

For a thin lens having radii of curvature R_1 and R_2, and manufactured from a material with refractive index n, the approximate focal length is given by,

$$\frac{1}{f} = (n-1)\left(\frac{1}{R_1} - \frac{1}{R_2}\right) \ . \tag{A.2}$$

For example a biconvex lens where R_1 is 50 mm and R_2 is -100 mm, with refractive index of 1.5, the focal length would be approximately 66.7 mm. Note that R_2 is negative due to the sign convention of optics – a negative radius means the surface appears concave to the incoming light.

A convex, or *positive*, lens acts to converge a collimated beam of light. This means it will focus a beam of parallel rays to a common focal point. A concave, or *negative*, lens causes a parallel bundle of rays to diverge.

When the object distance is finite, the distance from which the image is formed from the lens is given by,

$$\frac{1}{f} = \frac{1}{d_{\text{ob}}} + \frac{1}{d_{\text{im}}} \ , \tag{A.3}$$

where f is the focal length, d_{ob} is the object distance, and d_{im} the image distance. The magnification M of the lens is,

$$M = \frac{d_{\text{im}}}{d_{\text{ob}}} = \frac{h_{\text{im}}}{h_{\text{ob}}} \ , \tag{A.4}$$

where h_{ob} is the object height, and h_{im} the image height.

Often, equation A.3 is written in a more convenient form where names are given to

the inverse quantities. The inverse of the focal length is the *refracting power*, and the inverse of the distances are known as *vergences*. The power of a lens is the sum of the object and image vergences, thus

$$P = V_{\text{ob}} + V_{\text{im}} \ , \tag{A.5}$$

where the units are m^{-1} and known as the *diopter*. This is the optometrist form of the thin lens formula, and the power of eyeglasses and contact lenses are usually expressed in this unit. The human eye has a power of 50 D when focusing on distant objects.

The focal ratio, F, of a lens is defined as,

$$F = \frac{f}{a} \ , \tag{A.6}$$

where f is the focal length and a is the diameter of the limiting aperture. If the focal length of a system is 200 mm and the aperture is 20 mm, then the focal ratio is 10. This is usually written as being an $f/10$ system. It is easy to be ambiguous when using words such as "larger" or "smaller" to compare systems, and for this reason the terms "faster" and "slower" are defined. For example, an $f/5$ system is faster than an $f/10$ system since it has either a larger aperture, or a smaller focal length, and so allows more light through the system meaning a shorter (i.e., faster) detector shutter speed can be used.

Assuming the focal length is known, then constructing the image distance and height for a thin lens can be done by drawing three rays. Firstly, the *principal ray* from the object through the centre of the lens. Then a ray from the object, parallel to the optical axis and passing through the focal point f_2 behind the lens. Finally, a ray from the object, passing through the focal point f_1 in front of the lens, emerging parallel to the optical axis. The object and image planes lie at the intersections of the three rays. A typical optical configuration is shown in Figure A.1.

If an optical system is focusing a parallel entering bundle of light into a point, then it is called a focal system or more commonly an *objective*. If an objective is used with a detector such as a CCD or photographic film, then it is called a camera.

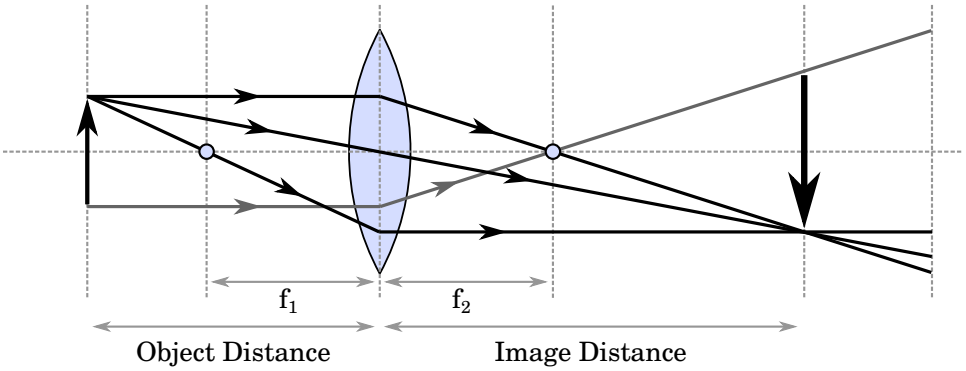

Figure A.1: Image formation for a thin lens with finite object distance.

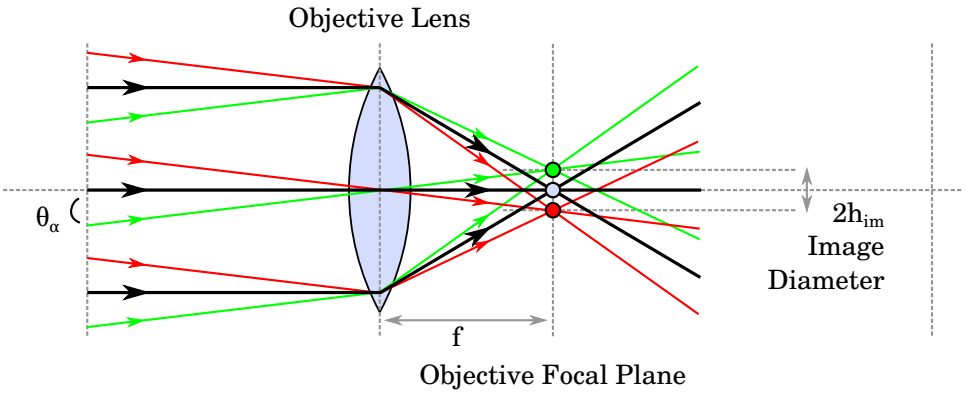

Figure A.2: Image formation for a thin lens with infinite object distance.

The size of the image formed when the object is at infinity can be calculated with simple trigonometry, as shown in Figure A.2. The image is formed at the focal point of the lens, on the objective focal plane. The field angle is given by,

$$\theta_\alpha = \tan^{-1} \frac{h_{\text{im}}}{f} \ , \qquad \text{(7.7 revisited)}$$

and so

$$h_{\text{im}} = -f \tan \theta_\alpha \ , \qquad \text{(A.7)}$$

where the negative sign indicates that the image is inverted.

A.1.2 Thick Lenses

In real life, lenses are not infinitesimally thin. They must have some thickness. Gaussian optics still applies for a thick lens provided the distance between the surfaces is taken into consideration. A thick lens has four *cardinal points*. Two we have already seen, the

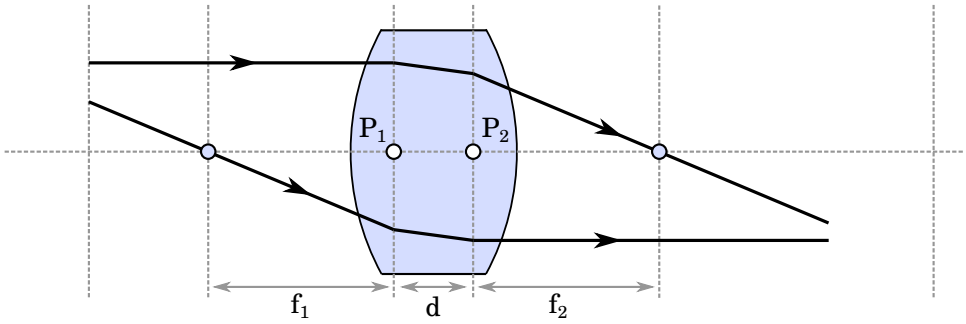

Figure A.3: Image formation for a thick lens with finite object distance.

focal points f_1 and f_2, and two additional *principal points* P_1 and P_2. These represent the intersection with the optical axis of two *principal planes*, which are defined as being the plane where the converging or diverging beam has the same diameter as the entering bundle. The positions of the principal planes vary for each design of lens.

For a thick lens having radii of curvature R_1 and R_2, and manufactured from a material with refractive index n, and where the distance between the principal planes is d, then equation A.2 becomes,

$$\frac{1}{f_{\text{eff}}} = (n-1)\left(\frac{1}{R_1} - \frac{1}{R_2} + \frac{(n-1)d}{nR_1R_2}\right) , \tag{A.8}$$

where f_{eff} is the effective focal length. If the position of all the cardinal points are known then it is possible to construct the image location in a similar way to a thin lens, as shown in Figure A.3.

Expressing f_{eff} purely in terms of the focal lengths of the two surfaces, equation A.8 becomes,

$$f_{\text{eff}} = \frac{f_1 f_2}{f_1 + f_2 - d} , \tag{A.9}$$

and equation A.9 also applies when calculating the overall effective focal length for multiple lenses in a system. The effective focal length may be longer or shorter than the physical length of the system, which is a useful property for telescopes where long focal lengths but short physical dimensions are desired.

A.2 Third-order Optics

Under first-order Gaussian optics, all rays forming a bundle converge to one point. When the $\sin\theta = \theta$ Gaussian approximation is dropped and replaced with the Maclaurin expansion (the Taylor expansion about zero) of sine,

$$\begin{aligned}\sin\theta &= \sum_{i=0}^{\infty} \frac{(-1)^i}{(2i+1)!}\theta^{2i+1} \ , \\ &= \theta - \frac{\theta^3}{3!} + \frac{\theta^5}{5!} - \frac{\theta^7}{7!} + \dots \ ,\end{aligned} \tag{A.10}$$

defects begin to appear in the system and not all rays converge to one point. Generally it is necessary to take only the third-order of the sine expansion,

$$\sin\theta = \theta - \frac{\theta^3}{6} \ , \tag{A.11}$$

which is known as "third-order theory" or sometimes "Seidel theory" after the 1857 publication by Ludwig von Seidel. Seidel took the first-order monochromatic defects, known as geometric *aberrations*, and decomposed them into five constituent components (Seidel, 1857). These components are now typically referred to as the five *Seidel Aberrations*.

These aberrations depend on several factors. The most important are the size of the aperture of the system, and the size of the field. Aberrations due to aperture go to zero as the aperture becomes infinitesimal, whilst those due to field angle go to zero as a ray bundle becomes closer to the optical axis. In the following subsections we will discuss the common forms of aberration.

A.2.1 Spherical Aberration

Spherical aberration causes the paraxial rays to come to a focus at a different point from the marginal rays. If the focal length of the marginal rays is less than that of the paraxial rays, then the aberration is said to be positive. If the focal length of the marginal rays is longer than that of the paraxial rays, then the aberration is said to be negative. There is no point on the optical axis where a perfect image is formed. There is a point,

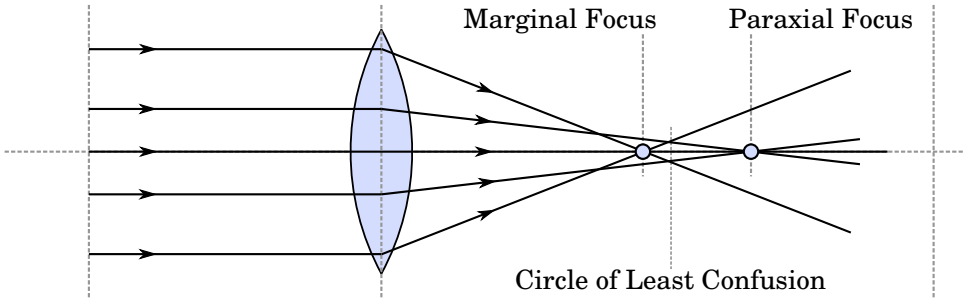

Figure A.4: Example of positive spherical aberration. The rays forming the paraxial focus are too far from the optical axis to be considered paraxial, but this has been done to more clearly demonstrate the effect.

somewhere between the marginal focus and the paraxial focus, where there is a "circle of least confusion" and the best image will be formed. An example of positive spherical aberration is shown in Figure A.4.

The diameter of the circle of least confusion for a spherical mirror is given by (James and Sternberg, 1969),

$$\Delta = \frac{f}{64F^3} \ , \tag{A.12}$$

where f is the focal length and F the focal ratio. For a single element lens with spherical surfaces the diameter of the circle is given by (James and Sternberg, 1969),

$$\Delta = -\frac{D^3}{16}\left[\left(\frac{n}{n-1}\right)^2 \frac{1}{f^2} - \left(\frac{2n+1}{n-1}\right)\frac{1}{R_1 f} + \left(\frac{n+2}{n}\right)\frac{1}{R_1^2}\right] \ , \tag{A.13}$$

where D is the lens diameter, R_1 the radius of curvature of the first surface, and n the refractive index of the medium. For an object at infinity the spherical aberration is at a minimum when the lens is plano-convex with the curved side towards the infinite conjugate, since this divides the deviation of the rays equally between both surfaces. When $n = 1.5$ equation A.13 reduces to,

$$\Delta = \frac{7}{48}\frac{f}{F^3} \ , \tag{A.14}$$

and if the value of $\Delta \leqslant |8\lambda/F^2|$ then the image of a point source is indistinguishable from that produced by a system with zero spherical aberration. The system can be said to be *diffraction limited* and not limited by spherical defects.

Spherical aberration can be avoided by using non-spherical optics, such as aspheric lenses and parabolic mirrors, or by reducing the size of the aperture stop to remove the more extreme marginal rays. Alternatively it can be corrected by introducing additional optical elements into the system. The Schmidt-Cassegrain telescope is a so-called catadioptric system, meaning that it uses both mirrors and lenses. An inexpensive spherical mirror is used, and the spherical aberration corrected using a lens designed with opposite spherical aberration to cancel that introduced by the rest of the system, but without added any extra power. Similarly, a Maksutov-Cassegrain telescope uses a meniscus lens designed with zero power, but with negative spherical aberration equal and opposite to the positive aberration of the main spherical mirror. For a lens, two spherical lenses can be combined together to form a compound *doublet* where aberrations are reduced by using two different types of glass and the surfaces designed to have equal and opposite aberrations.

A.2.2 Comatic Aberration

Comatic aberration, or *coma*, degrades images of off-axis point sources. When an incoming ray is both meridional and off-axis, it is known as a *skew ray*. Coma is particularly a problem when using parabolic mirrors. As we saw earlier, a parabolic mirror is free of spherical aberration. A parallel bundle of rays that are also parallel to the optical axis will come to a single focus. If the bundle of rays are not parallel to the axis of the parabola, then they will be incident on the mirror at an angle which results in off-axis images appearing wedge-shaped. Coma means *comet* in Latin, and is so-called because an image of a point source showing coma aberration takes the shape of a comet. For a telescope, this results in star images towards the edge of the field having comet-like tails that spread out either away from the optical axis (external or negative coma) or inwards towards the optical axis (internal or positive coma). A similar effect occurs with a lens, where skew rays cause an imbalance of refractive power on opposite sides of the lens. An example of negative coma in a lens is shown in Figure A.5.

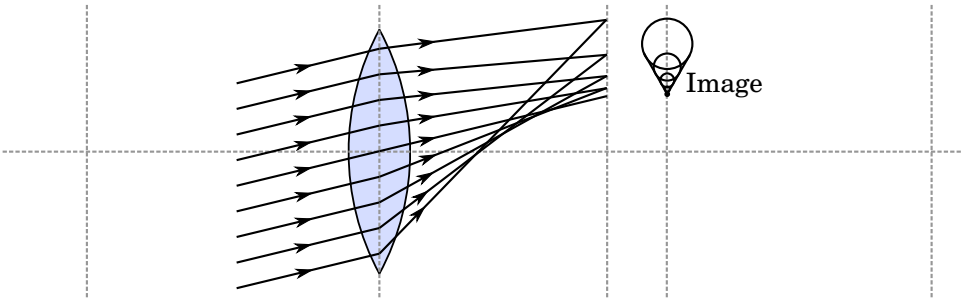

Figure A.5: Example of negative comatic aberration.

A.2.3 Astigmatism

Astigmatism occurs when there is a difference between the power of a lens in tangential (vertical) plane and in the sagittal (horizontal) plane. This causes there to be two focal surfaces formed in the two planes, which means it is impossible to create a sharp off-axis image. When a star is imaged by an astigmatic system it will appear as a vertical line in the sagittal focal plane and as a horizontal line in the tangential focal plane. At some point in between there will be an approximately circular area of blur that reaches its smallest dimensions. When astigmatism is fully corrected the system is said to be anastigmatic and the tangential and sagittal focal surfaces coincide.

A.2.4 Curvature of Field

Curvature of field means that the sharpest image is formed on a curved focal plane. When a flat sensor is used such as a photographic plate or a CCD, the off-axis image will be blurred due to being out-of-focus. Most systems will suffer from curvature of field unless measures have been deliberately taken to prevent it.

A.2.5 Distortion

Image distortion occurs when the image scale varies across the field, with distance from the axis. If the image scale is too large then it is known as pincushion, or positive, distortion. When too small, it produces barrel or negative distortion where the edges of the image become more rounded.

A.2.6 Chromatic Aberration

Chromatic aberration is not a geometrical aberration, but it causes similar effects and so it is appropriate to be included in a discussion on aberrations. Chromatic aberration is the failure of an optical system to focus or magnify all wavelengths equally. Longer wavelengths are refracted less than shorter wavelengths, which visibly spreads out different colours in an effect known as *dispersion*.

There are two forms of chromatic aberration. Longitudinal chromatic aberration is an error in the focus, where different wavelengths have different focal planes. Lateral chromatic aberration is an error in magnification, where different wavelengths are dispersed at different powers from opposite sides of the aperture, resulting in an effect similar to comatic aberration for off-axis sources.

Non-diffracting reflective optics do not suffer from either type of chromatic aberration. The typical method of minimising chromatic aberration in diffracting systems is to use two spherical lenses made from different materials, known as a doublet, in a similar manner to correcting spherical aberration. Such lenses are known as *achromatic*, and designs typically correct both spherical and chromatic aberration simultaneously. For this reason, even in monochromatic systems it is useful to use a doublet since they also correct for the spherical aberration. Sometimes a third lens is used to further improve the achromatic doublet, and such systems are called *apochromatic*.

A.2.7 Vignetting

Vignetting is not an aberration as such, but it is another defect worth considering in a discussion on optical systems. Vignetting occurs when not all the intended rays make it through the system to form part of the image. Figure A.6 shows a single lens with a diaphragm aperture mounted in front of it. In this case, the diaphragm forms the limiting aperture stop and not the lens itself. The diaphragm controls both the light-gathering power of the system, and the position of the marginal rays as they pass through the lens. Every objective system also contains a field stop, and this limits the size of the image

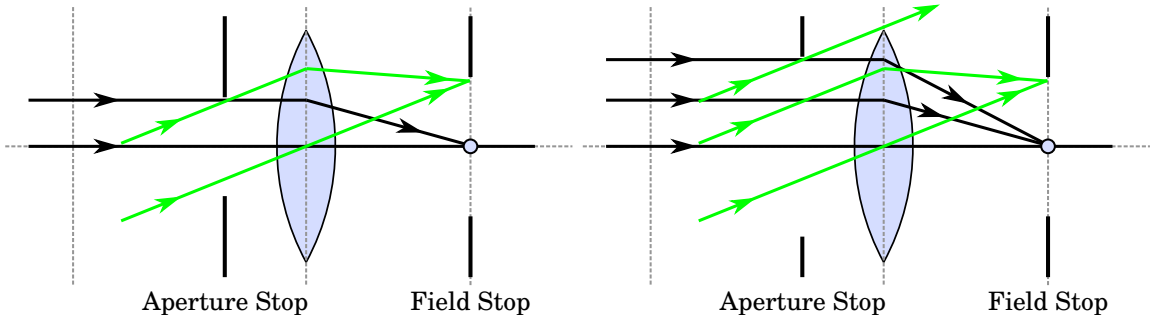

(a) Narrow Aperture Stop (b) Wide Aperture Stop

Figure A.6: Example of vignetting. In (a) the aperture stop is narrow and a similar amount of light is reaching the lens from both on-axis rays and off-axis rays. As the aperture stop is made wider, as in (b), more on-axis rays reach the image but some of the off-axis rays are lost resulting in light fall-off at the edge of the field compared with the centre.

field and thus determines the most oblique beam that can reach the focal plane. The field stop is typically determined by the size of the sensor being used to capture the image.

In Figure A.6(a) the aperture stop and the field stop entirely determine the maximum position and field angle of the marginal rays. As the aperture stop is made wider, as in A.6(b), part of the oblique beam is cut off by the edge of the lens. This results in more on-axis rays reaching the image, but some light is lost at the edge of the field resulting in a gradual light fall-off across the field from the centre.

A.2.8 The Abbe Sine Condition

We saw earlier that the focal ratio, F, of a lens is defined as,

$$F = \frac{f}{a} \ , \qquad \text{(A.6 revisited)}$$

where f is the focal length and a is the diameter of the limiting aperture. When considering the resolving power of a lens, a useful performance metric is the the *numerical aperture* or NA introduced by Ernst Abbe. For an object at infinity the numerical aperture of a lens is defined as,

$$\text{NA} = n\frac{a}{2f} \ , \qquad \text{(A.15)}$$

where n is the refractive index of the medium in the image space, and f and a are as defined before. Abbe went on to show what has become known as the *Abbe sine condition* which states that both on-axis and off-axis rays will be sharply imaged only if,

$$\frac{\sin u}{\sin U} = \frac{\sin u'}{\sin U'} \ , \tag{A.16}$$

where u and U are angles relative to the optical axis of any two rays launched from the object, and u' and U' are the angles of the same rays as they reach the image. Another way of stating this is,

$$\frac{h}{n \sin \alpha} = C \ , \tag{A.17}$$

where h is the distance of the incoming ray from the axis, α is the angle between the ray and the axis as it travels towards the focus in image space, and C is a constant. This implies that there is an imaginary spherical surface connecting the entrance and exit rays where the centre is at the focal point. This constant is the effective focal length, and the imaginary spherical surface is the second principal plane. It would be logical to believe that the angle defined in equation A.15 is equal to,

$$\mathrm{NA} = n \tan \alpha_{\max} \ , \tag{A.18}$$

where $\alpha_{\max}$ is the half-angle of the cone formed by the marginal rays. But this would be correct only if the principal planes were actually plane, as in the thick-lens definition of first-order optics. For a lens that is well corrected for coma and spherical aberration, the Abbe sine condition applies and equation A.18 becomes,

$$\mathrm{NA} = n \sin \alpha_{\max} \ . \tag{A.19}$$

This, of course, also has implications when calculating the size of an image formed by an objective lens of an object at infinity. Equation A.17 can be written as,

$$\frac{a}{2n \sin \alpha} = f_{\mathrm{eff}} \ , \tag{A.20}$$

since h is now the input aperture and the constant C is the effective focal length as

described above. The focal ratio, equation A.6, becomes,

$$F \approx \frac{1}{2\mathrm{NA}} \ , \tag{A.21}$$

and similarly the calculation of image size, equation A.7, becomes,

$$h_{\mathrm{im}} = -f \sin\theta_{\alpha} \ , \tag{A.22}$$

where the negative sign as before indicates that the image is inverted.

Appendix B

Operational Amplifier Theory

Contents

B.1 Introduction

In this chapter we will look at some of the properties of amplifiers that will help us when designing a low-noise photodiode amplifier. The generic term amplifier used here could also mean a unity-gain voltage-follower, or even a filter that attenuates certain frequencies.

In order to more clearly demonstrate the concepts discussed, we make use of circuit simulation. The industry standard electrical simulation tool is SPICE (Simulation Program with Integrated Circuit Emphasis), an open-source package developed by Laurence Nagel and Donald Pederson at the Electronics Research Laboratory of the University of California, Berkeley, in 1973 (Nagel and Pederson, 1973; Nagel, 1975). Originally a text-based application, subsequent versions of SPICE have added graphical interfaces for schematic capture and presentation of results. Many of these derivatives are now commercial applications, but some remain available as freeware. Two notable free applications are LTspice from Linear Technology and TINA from Texas Instruments. The following circuit simulations have all been performed with LTspice.

We begin by looking at amplifier gain in Section B.2, and then go on to consider bandwidth and stability in Section B.3 and B.4. Finally we look at the effect of frequency dependence in the amplifier feedback network in Section B.5.

B.2 Negative Feedback Gain

A simple non-inverting amplifier circuit is shown in Figure B.1. Here, A is the gain of the amplifier with no feedback, the so-called *open-loop* gain, G_{OL}. The feedback can be described as,

$$V_{\mathrm{FB}} = \beta\, V_{\mathrm{OUT}} \ , \tag{B.1}$$

where the feedback factor β is the percentage of the output that is returned to the subtraction node of the amplifier and is set simply by the use of a resistor network. Thus,

$$\beta = \frac{\mathrm{R}_1}{\mathrm{R}_1 + \mathrm{R}_2} \ , \tag{B.2}$$

where R_1 and R_2 are the two resistors forming the feedback resistor divider. If we define the *closed-loop* gain, G_{CL}, as being simply the output divided by the input, then,

$$\begin{aligned} G_{\mathrm{CL}} &= \frac{V_{\mathrm{IN}}\, G_{\mathrm{OL}}}{V_{\mathrm{IN}} + (\beta\, V_{\mathrm{OUT}})} \ , \\ &= \frac{V_{\mathrm{IN}}\, G_{\mathrm{OL}}}{V_{\mathrm{IN}} \left(1 + \beta\, \frac{V_{\mathrm{OUT}}}{V_{\mathrm{IN}}}\right)}, \\ &= \frac{G_{\mathrm{OL}}}{1 + \beta\, G_{\mathrm{OL}}} \ , \end{aligned} \tag{B.3}$$

where the quantity $\beta\, G_{\mathrm{OL}}$ is known as the *loop-gain*. In typical applications the loop-gain is very much larger than 1, and so the expression in equation B.3 can be simplified to,

$$G_{\mathrm{CL}} \approx \frac{1}{\beta} \ , \tag{B.4}$$

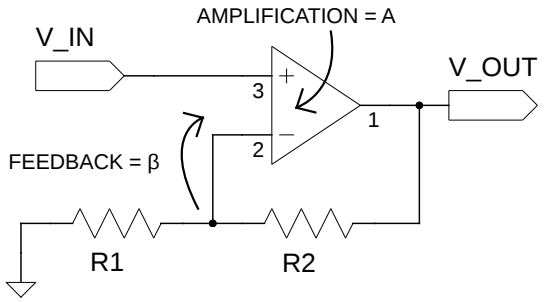

Figure B.1: A basic non-inverting amplifier circuit.

from which we find the standard expression for the gain of a non-inverting amplifier,

$$\begin{aligned} G_{\mathrm{NI}} &= \frac{1}{\beta} \,, \\ &= \frac{\mathrm{R}_1 + \mathrm{R}_2}{\mathrm{R}_1} \,, \\ &= 1 + \frac{\mathrm{R}_2}{\mathrm{R}_1} \,, \end{aligned} \tag{B.5}$$

where G_{NI} is the non-inverting gain. Thus the closed-loop gain of the system depends entirely on the external components forming the feedback network.

B.3 Gain Sensitivity and Bandwidth

It is difficult to design a general purpose amplifier with precisely controlled gain that is uniform across the whole of the required bandwidth. It is much easier to design an amplifier with very high gain. Typical values of open-loop gain G_{OL} can be millions of times the input, but this is very unstable depending on the amplifier supply voltage, temperature, and frequency. Often the maximum gain starts to roll-off at frequencies well under 1 Hz. Most amplifiers are *internally compensated* making them more stable, resulting in an open-loop gain that rolls off at 20 dB/decade and so sacrificing high-frequency response for stability. At very high frequencies the gain eventually becomes zero due to roll-off caused by parasitic capacitance and inductance, which is the reason real life amplifiers never have infinite bandwidth. We saw in equation B.4 that the closed-loop gain depends only on the external feedback network, and so negative feed back is

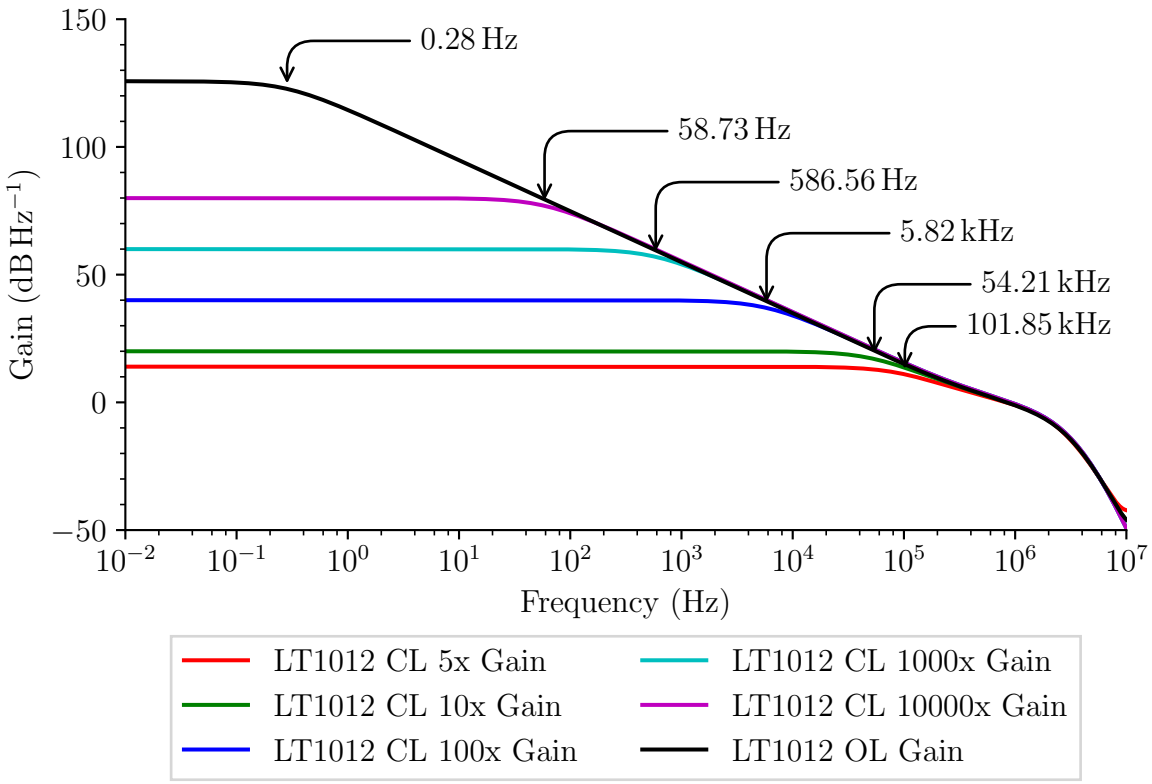

Figure B.2: Comparison between open-loop gain and several closed-loop gains. The arrows indicate the −3 dB cut-off point for each curve.

an ideal solution to the problem of open-loop gain instability. By capping the very high open-loop gain down to our desired level we stabilise the gain of the amplifier.

Figure B.2 shows the simulated open-loop gain for the LT1012 general-purpose op-amp made by Linear Technology. Also plotted are the closed-loop gain responses for five different gains. The arrows indicate the −3 dB cut-off point for each curve. It is clear that the open-loop gain is not constant across the bandwidth. The −3 dB cut-off is just 0.28 Hz, followed by 20 dB/decade roll-off as expected. The roll-off increases to almost 40 dB/decade as the gain passes through unity (0 dB). This frequency dependence means we must modify equation B.3 to include the frequency dependent open-loop gain,

$$G_{\mathrm{CL,low}\nu} = \frac{G_{\mathrm{OL,low}\nu}}{1 + \beta\, G_{\mathrm{OL,low}\nu}} \,, \tag{B.6}$$

where $G_{\mathrm{OL,low}\nu}$ and $G_{\mathrm{CL,low}\nu}$ are now the frequency dependent open-loop gain and frequency dependent closed-loop gain respectively. The closed-loop cut-off frequency $f_{\mathrm{c,CL}}$

is related to the open-loop cut-off frequency by,

$$f_{\mathrm{c,CL}} = f_{\mathrm{c,OL}}(1 + \beta\ G_{\mathrm{OL,low}\nu})\ , \tag{B.7}$$

where $f_{\mathrm{c,OL}}$ is the closed-loop cut-off frequency. This assumes that the relationship between gain and frequency is represented by an ideal one-pole response. We already know from Figure B.2 that this op-amp has two poles which can be seen from the two turning points, the first at low-frequency and the second at around 1 MHz where the roll-off increases to almost 40 dB/decade.

As β tends to zero (increasingly higher closed-loop gain) the closed-loop cut-off frequency approaches the open-loop cut-off frequency, and this means that lower gains have increased bandwidth before the gain begins to roll-off. Comparing equations B.6 and B.7 we can see that both terms are changed by the same factor. The gain is decreased by a factor of $(1+\beta\ G_{\mathrm{OL,low}\nu})$ and the bandwidth is increased by a factor of $(1+\beta\ G_{\mathrm{OL,low}\nu})$, which can be seen from the closed-loop gain examples in Figure B.2. Looking at open-loop gain in Figure B.2 we can see that at low-frequency $G_{\mathrm{OL,low}\nu}$ is 126 dB (approximately 2×10^6 times). For the the 100-times gain response β is 0.01. From equation B.6 we expect the gain to be reduced by $(1 + 2 \times 10^6\ .\ 0.01)$ which is equal to 20k-times or 86 dB, and we can clearly see the gain is indeed reduced to 40 dB. From equation B.7 we expect the cut-off frequency to increase to 5.6 kHz, and the simulation predicts a cut-off of 5.82 kHz. The difference is due to the single-pole assumption discussed earlier.

These two specification values are usually combined by manufacturers into one value known as the *gain-bandwidth product*. This is defined by,

$$f_{\mathrm{GBP}} = G_{\mathrm{OL,low}\nu}\, f_{\mathrm{c,OL}} = G_{\mathrm{CL,low}\nu}\, f_{\mathrm{c,CL}}\ , \tag{B.8}$$

where f_{GBP} is the gain-bandwidth product. In the design process, this is used to select the amplifier performance required. If a gain of 100 is required over a bandwidth of 100 kHz then the f_{GBP} required is at least 10 MHz.

B.4 Amplifier Stability

The stability of an amplifier refers to its ability to reliably amplify as designed despite variations in component tolerance, or changes in environmental and operating conditions. A marginally unstable amplifier may exhibit effects such as ringing in the time domain, or gain-peaking in the frequency domain. A completely unstable amplifier will exhibit oscillation that will likely be catastrophic to the signal. In this section we will look at what causes an amplifier to oscillate and what can be done to ensure the system will be stable.

An amplifier will start to oscillate if the feedback becomes positive rather than negative. This will occur if an induced phase-shift on the output is greater than 180°, causing the feedback to approach anti-phase with the input. All amplifiers can be considered to be low-pass filters with a response that rolls-off with increasing frequency at 20 dB/decade for a single pole, as seen earlier. The filter response is caused by one or more poles, and it is these poles that cause phase-shift in the output with respect to the input. Each pole in the circuit transfer function provides 90° of phase shift.

The amplifier stability is defined by the loop gain $\beta\, G_{\mathrm{OL}}$. Recall from equation B.6 that the closed-loop gain is,

$$G_{\mathrm{CL,low}\nu} = \frac{G_{\mathrm{OL,low}\nu}}{1 + \beta\ G_{\mathrm{OL,low}\nu}}\ , \qquad \text{(B.6 revisited)}$$

which assumes that $\beta\, G_{\mathrm{OL,low}\nu}$ must be positive. In the case where $\beta\, G_{\mathrm{OL,low}\nu} = -1$ then $G_{\mathrm{CL,low}\nu}$ becomes infinite. At this point the amplifier will become an oscillator. Even with zero input, there will be sufficient noise in the circuit for the positive feedback to take hold and dominate the output. From this we find the *Barkhausen stability criterion* which states that the loop-gain magnitude must be less than unity at the point where the phase-shift reaches 180°, thus,

$$|\beta\, G_{\mathrm{OL,low}\nu}|_{180} < 1\ . \qquad \text{(B.9)}$$

Ideally equation B.9 would be significantly less than unity (i.e., less than 0 dB) in order

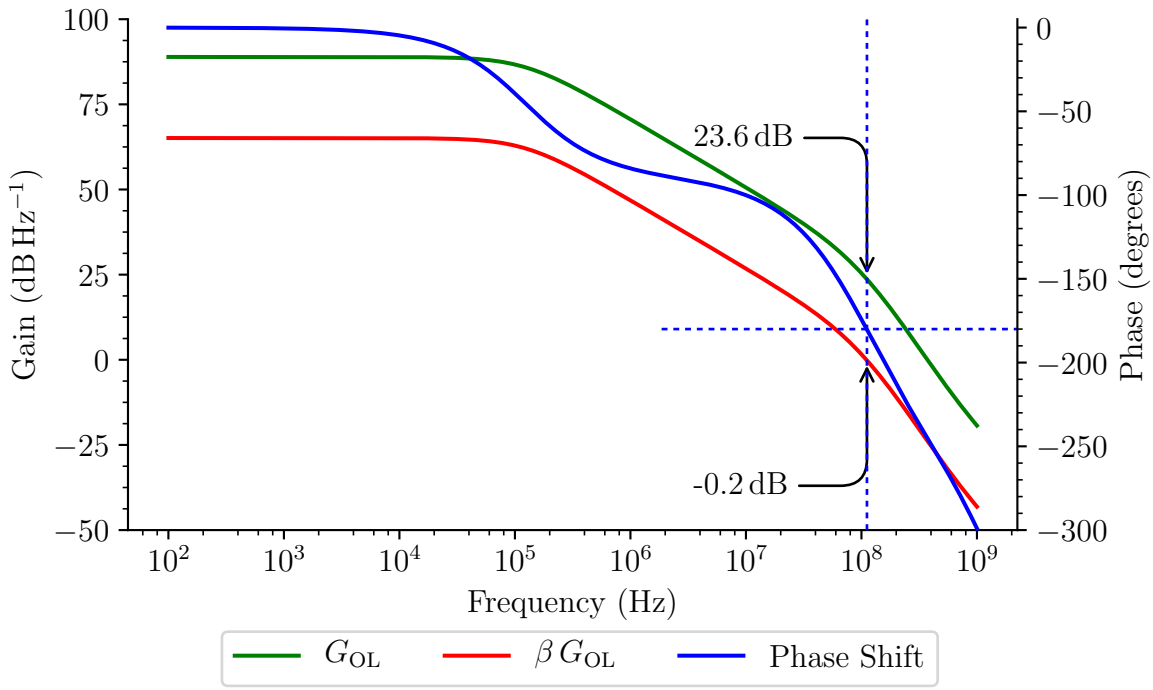

Figure B.3: Bode plot of an unstable amplifier. The green trace is the open-loop gain, and the red trace is the reduced loop-gain after adding 15.5 times closed-loop gain ($\beta = 0.065$). The blue trace shows the phase shift (measured against the right-hand axis), with a crosshair indicating the frequency at which the phase reaches $-180°$.

to avoid cases of marginal stability, where the system may work in one location during testing, and then fail when deployed due to changing environmental conditions. The parameters used to assess stability are the *gain margin* and *phase margin*, which will be discussed next.

For this simulation a generic non-inverting voltage amplifier model of the sort we saw in Figure B.1 was constructed from discrete transistors. This is necessary since real op-amps are internally compensated specifically to help avoid the stability issues we wish to discuss, and our simple model is very unstable by comparison. Figure B.3 shows a combined *Bode magnitude plot* and *Bode phase plot* developed by Hendrik Wade Bode. The blue crosshair indicates the frequency at which the phase shift reaches $-180°$. The open-loop gain at this frequency is 23.6 dB and so this amplifier will definitely not be stable at unity closed-loop gain ($\beta = 1$). To bring the amplifier close to stability we need to reduce the entire loop-gain $\beta\, G_{\mathrm{OL}}$ curve by a factor of 23.6 dB. We can do this by

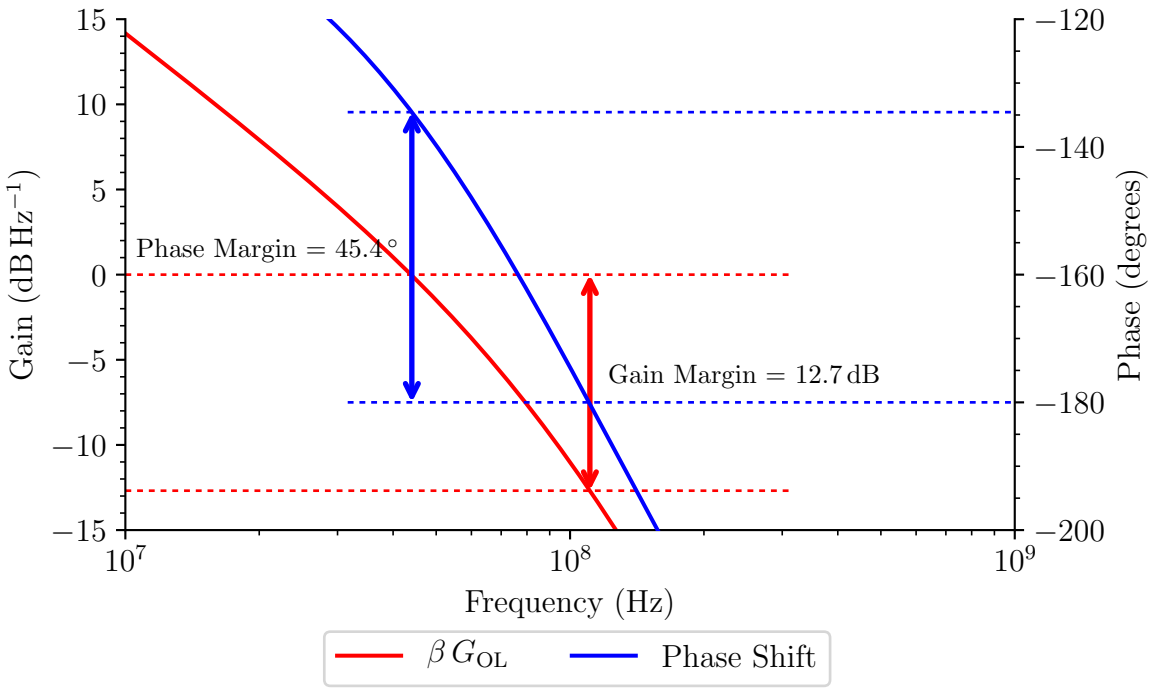

Figure B.4: Bode plot of a stable amplifier. The red trace is the reduced loop-gain after adding 62 times closed-loop gain ($\beta = 0.016$). The blue trace shows the phase shift measured against the right-hand axis. An amplifier can be considered stable if the phase margin is at least 45°.

increasing the closed-loop gain to 15.5 times, making $\beta = 0.065$, using resistors R1 = 1 kΩ and R2 = 14.5 kΩ. This is only possible when β is frequency independent, which can be assumed to be the case when the feedback network is made only from resistors. The red trace in Figure B.3 shows that with this closed-loop amplification we now have unity gain at the frequency where the phase shift reaches −180°. However, this circuit now has only marginal stability. It may work right now, but when deployed somewhere with different environmental conditions and so changed electrical parameters it could easily become unstable again. A standard rule-of-thumb states that to be considered stable an amplifier must be designed with a phase margin of at least 45°.

Figure B.4 shows how the phase margin and gain margin are determined. The closed-loop gain of this amplifier has now been increased to 62 times, making $\beta = 0.016$, using resistors R1 = 1 kΩ and R2 = 61 kΩ. The gain margin is calculated at the frequency where the phase shift reaches −180°. We can see that at this frequency the gain is now −12.7 dB,

meaning we have a 12.7 dB margin between the 0 dB point. The gain margin measures how far the loop-gain is from unity when the phase shift is −180°. If an amplifier has a gain margin of 0 dB then it will oscillate. The phase margin is calculated at the frequency where the gain is 0 dB. Again, we can see that at this frequency the phase shift is −134.6°, meaning we have a margin of 45.4° from the −180° point. The phase margin measures how close the loop-gain is to achieving a phase shift of −180° when the loop-gain is unity. If an amplifier has a phase margin of 0° then it will oscillate. This amplifier can now be considered to be stable.

The datasheet for the LT1012 general-purpose op-amp we examined earlier shows that at 0 dB this amplifier has a phase shift of around 110° meaning it has a phase margin of 70°. The design is internally compensated with a 30 pF capacitor specifically intended to reduce the open-loop gain at higher frequencies, and so ensuring that the gain will be less than unity where the phase shift becomes 180°. This op-amp is *unity-gain stable*.

We have now developed two of the key metrics of amplifier design. Increasing the closed-loop gain increases the stability of the amplifier. Decreasing the closed-loop gain increases the bandwidth of the amplifier. Up to now we have only considered feedback networks that are frequency independent. Now, we will go on to look at the effects on frequency response and phase shift when both the gain and the feedback are frequency dependent.

B.5 Feedback Poles and Zeros

We have already mentioned poles earlier when discussing the low-pass nature of an amplifier. We now need to look at these in more detail, particularly with respect to the frequency dependence of a negative feedback network. Poles are defined as the roots of a polynomial on the denominator of a transfer function, and zeros are defined as the roots of a polynomial on the numerator of a transfer function. For the generalised transfer

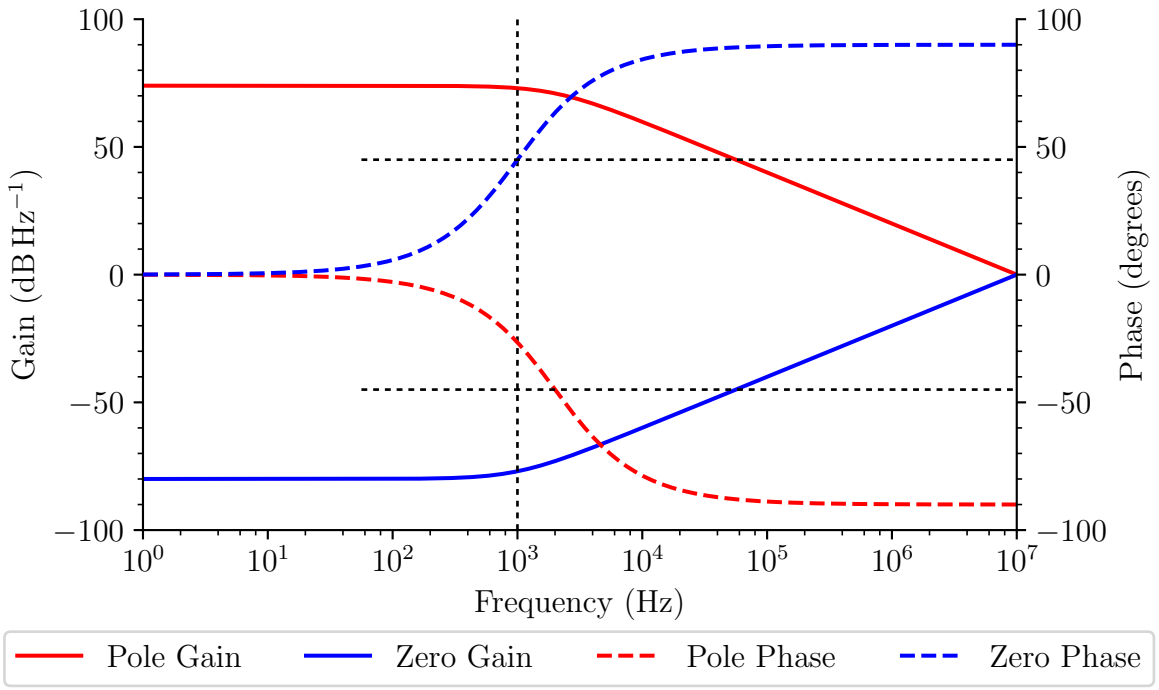

Figure B.5: Bode plot of a 1 kHz pole and zero. A pole introduces a −90° phase shift and a 20 dB/decade decay in gain. A zero introduces a 90° phase shift and a 20 dB/decade increase in gain. The crosshairs indicate the ±3 dB and ±45° phase shift points.

function,

$$G(s) = \frac{(s - z_1)(s - z_2)\dots(s - z_n)}{(s - p_1)(s - p_2)\dots(s - p_n)}\,, \tag{B.10}$$

the poles are $p_1, p_2, \dots, p_n$ and the zeros are $z_1, z_2, \dots, z_n$. A pole in the transfer function contributes terms that create an exponential decay. A zero in the transfer function contributes terms that create an exponential growth. Both induce phase shifts in the signal, shown in Figure B.5. A pole introduces a −90° phase shift and a 20 dB/decade decay in gain. A zero introduces a 90° phase shift and a 20 dB/decade increase in gain. Things become slightly more complicated when we apply this to a feedback network.

Figure B.6 shows two feedback networks, one configured with a pole and another configured with a zero. The RC network of both is defined for −3 dB at $f_c = 1$ kHz. In B.6(c) we can see the response we expect, a −3 dB corner at 1 kHz followed by 20 dB/decade decay. The phase shift begins at around a decade before the corner frequency, reaches −45° at the corner frequency, and ends at −90° around a decade later. The response for the zero shown in B.6(d) is not what we expected. This is because the circuit in B.6b actually now

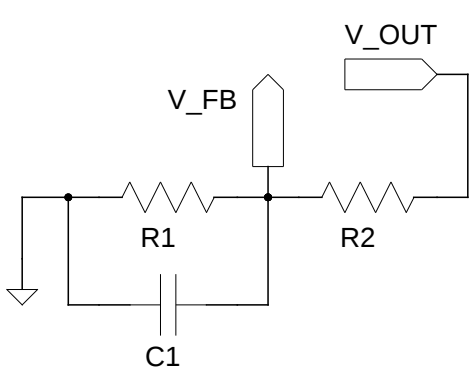

(a) Feedback with Pole

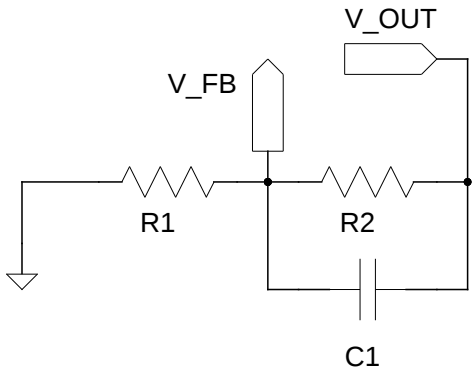

(b) Feedback with Zero

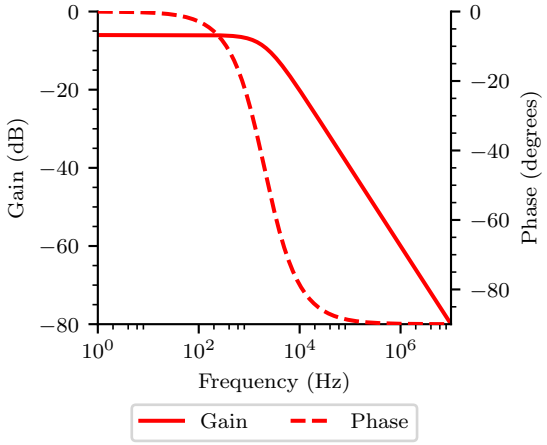

(c) Response with Pole

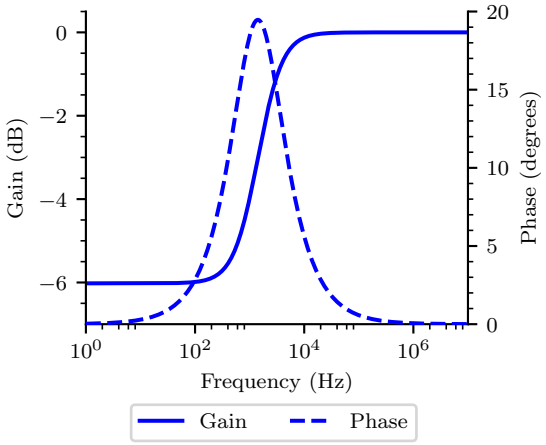

(d) Response with Zero

Figure B.6: A negative feedback network with poles and zeros.

contains both a zero *and* a pole. Rather than seeing the continued 20 dB/decade increase in loop-gain, the response combines with the pole to produce a step increase. Likewise the phase shift from the zero and the pole combine to form zero net phase shift, with only a small value phase shift during the transition. Capacitance in the feedback network can cause complicated frequency response profiles in amplifiers.

In the previous section we developed a model of an unstable amplifier and we determined that for it to be stable a closed-loop gain of 62 times ($\beta = 0.016$) was required to ensure a phase margin of at least 45°. This meant that the feedback network was made from resistors R1 = 1 kΩ and R2 = 61 kΩ. If we take that same stable amplifier circuit and add capacitance to the feedback network we can explore what effect this has on the

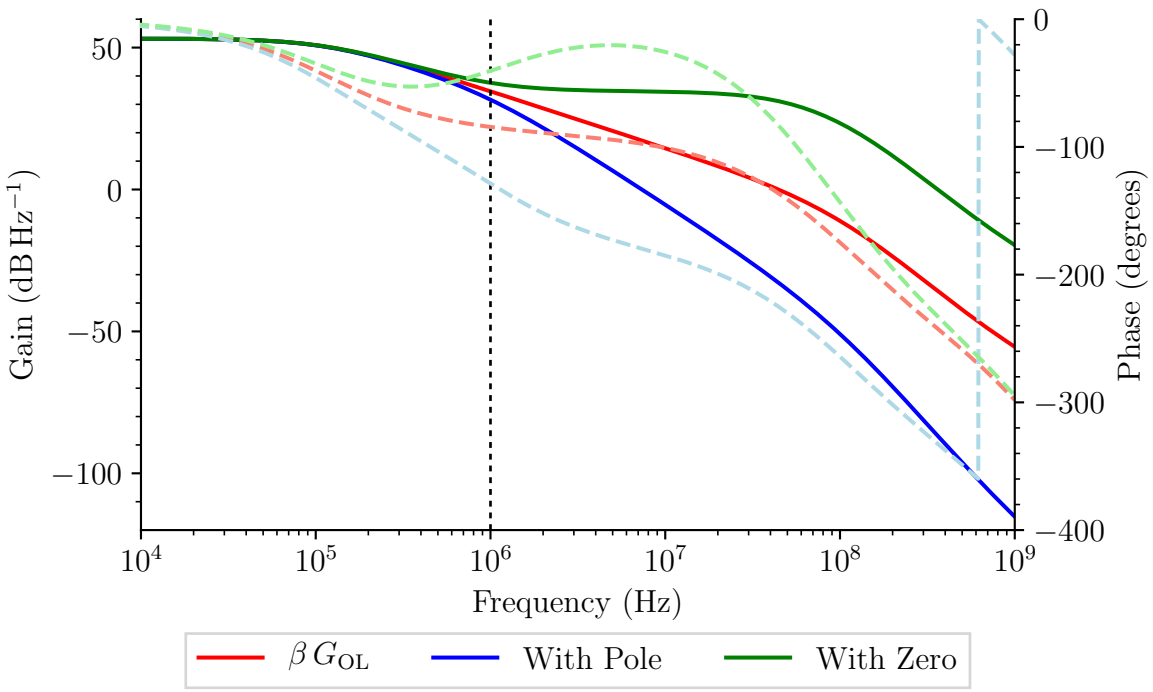

Figure B.7: Frequency dependent feedback networks

response and the stability. Figure B.7 shows the stable amplifier loop-gain response we saw in Figures B.3 and B.4, and includes the response when an additional zero or pole is added at 1 MHz to the feedback network, indicated by the dotted vertical line. The 3 dB diversion at the cut-off frequency is clearly visible, with the pole decreasing the loop-gain (decreasing β, increasing the closed-loop gain) and the zero increasing the loop-gain (increasing β, decreasing the closed-loop gain).

Looking more carefully at Figure B.7, we see the blue line indicating the feedback network with the additional pole begins to decay at 20 dB/decade, increases to 40 dB/decade at 1 MHz, before increasing again at 100 MHz to 60 dB/decade. We also see an increase in the phase shift as expected. These effects combine to reduce the gain-margin to 0 dB and the phase-margin to 0°. The amplifier is now only marginally stable. The additional pole has reduced the stability.

Looking at the red line in Figure B.7 which indicates the feedback network with the additional zero and pole, we see the zero initially acts against the decay of the first pole, reducing the decay to almost 0 dB/decade. The additional pole then rolls off again to

give the same profile as before, except with a step increased gain. There is an initial reduction in phase shift before as expected settling back to almost the same phase shift as the original red trace. The increase in loop-gain without a reduction in phase shift means that the 0 dB point is now placed well past $-180°$ of phase shift. At the point of $-180°$ there is still 15.7 dB of gain, meaning the circuit will now oscillate.

Clearly any capacitance in the feedback network has been detrimental to the stability of the amplifier. This is not always the case. Even with internally compensated amplifiers, sometimes additional capacitance is needed in the system in order to compensate for unavoidable capacitance elsewhere. Small amounts of parasitic capacitance can still have a significant affect on the system stability when present at the op-amp inputs or output. In the context of a photodiode transimpedance amplifier this can be particularly problematic since the diode junction capacitance appears exactly where we do not want it, right at the inverting input. Photodiode amplifiers can be prone to oscillation if not designed carefully.

Appendix C

Astronomical Algorithms

Contents

C.1 Introduction

In order to point a telescope at any celestial object, we must know its coordinates at any given time and date, and be able to convert between different coordinate systems. The two key texts in the field of practical astronomy are *Astronomical Algorithms* (Meeus, 1991) and *Practical Astronomy with Your Calculator* (Duffett-Smith, 1988). In this chapter we will review in Section C.2 the standard astronomical algorithms for determining the position of the Sun at any time and observer location, and then in Section C.3 develop the rotation matrices required for a full three-dimensional pointing model of a telescope mount including conversion between horizon, equatorial, ecliptic, coordinate systems.

C.2 Locating the Sun

The four most common astronomical coordinate systems are horizon, equatorial, ecliptic, and galactic. Each system uses a different frame of reference, and as such the ideal use case varies depending on the purpose of the position coordinate. The horizon coordinates,

azimuth, A, and altitude, a, are referenced from the plane of the observer's horizon. This system is ideal for pointing a telescope, but since position of all astronomical objects are continually moving with time as the Earth rotates it is poor at defining the position of stars. The equatorial coordinates, right ascension, α, and declination, δ, are referenced from Earth's equator. Right ascension is measured from the *vernal equinox*, Υ, and so does not change as the Earth rotates making the system perfect describing the position of astronomical objects. The ecliptic coordinates, ecliptic longitude, λ, and ecliptic latitude, β, are similar to the equatorial system except referenced from the plane of Earth's orbit around the Sun. Ecliptic coordinates are ideal for defining the position of the Sun, and also for describing the position of other planets in the Solar System. Galactic coordinates use the plane of our own Galaxy for reference, and so are ideal for describing the position of stars within the Milky Way.

To calculate the position of the Sun, we need to find its geocentric ecliptic longitude. The simplest way to do this is to define an epoch for which the mean ecliptic longitude is known, and use that as a starting point from which we can calculate how far the Sun has moved on a given date. If the epoch is J1990 (1990 January 0.0, Julian date 2 447 891.5) then the mean ecliptic longitude, ϵ_g, at the epoch is 279.403 303°. To find the position of the Sun at any other time, we have to add the number of degrees moved by the Sun since the epoch and also take account of the slightly elliptical orbit of the Earth. If we first consider a circular orbit, then the number of degrees moved since the epoch, the *mean anomaly* $M_\odot$, is given by,

$$M_\odot = \frac{360}{365.242191}\, D + \epsilon_g - \overline{\omega}_g \ , \tag{C.1}$$

where D is the number of days since the epoch, ϵ_g the mean ecliptic longitude at epoch, and $\overline{\omega}_g$ the mean geocentric longitude at solar perigee. To correct for the non-circular orbit and find the *true anomaly* we need to determine the *eccentric anomaly*, E, given by,

$$E - e \sin E = M_\odot \ , \tag{C.2}$$

Table C.1: Solar orbit parameters at epoch 1990.0.

Parameter	Description	Value
ϵ_g	Ecliptic longitude	279.403 303°
$\overline{\omega}_g$	Ecliptic longitude at perigee	282.768 422°
e	Eccentricity of orbit	0.016 713
r_0	Length of semi-major axis	$1.495\,985 \times 10^8$ km
θ_0	Solar angular diameter at $r = r_0$	0.533 128°

where e is the eccentricity of the orbit equal to 0.016 713. This equation does not have a closed-form solution for E when given M, and is usually solved iteratively using numerical techniques such as the Newton-Raphson method. Since we will not need to be particularly accurate when positioning the Sun into the field-of-view of the autoguider, and the eccentricity is very low, we can use a simpler approximation to the true anomaly, v, given by,

$$v = M_\odot + \frac{360}{\pi}\, e \sin M_\odot \ , \tag{C.3}$$

where both v and $M_\odot$ are expressed in degrees. The geocentric ecliptic longitude of the Sun is then given by,

$$\lambda_\odot = v + \overline{\omega}_g \ , \tag{C.4}$$

or expressed in full,

$$\lambda_\odot = \frac{360}{365.242191}\ D + \frac{360}{\pi}\ e \sin\left\{\frac{360}{365.242191}\ D + \epsilon_g - \overline{\omega}_g\right\} + \epsilon_g \ . \tag{C.5}$$

The values of ϵ_g, $\overline{\omega}_g$, and e at other epochs can be determined using the following equations,

$$\begin{aligned} \epsilon_g &= 279.696\,677\,8 + (36\,000.768\,92\,T) + (0.000\,302\,5\,T^2) \\ \overline{\omega}_g &= 281.220\,844\,4 + (1.719\,175\,T) + (0.000\,452\,778\,T^2) \\ e &= 0.016\,751\,04 - (0.000\,041\,8\,T) - (0.000\,000\,126\,T^2) \end{aligned} \tag{C.6}$$

where T is the number of Julian centuries since 1900 January 0.5,

$$T = \mathrm{JD} - 2415020.0/36525 \tag{C.7}$$

and JD is the Julian date of the desired epoch. Some useful orbit parameters at epoch J1990 are given in Table C.1.

Having determined the Sun's geocentric ecliptic longitude, $\lambda_{\odot}$, and considering that the ecliptic latitude, $\beta_{\odot}$, is zero by definition, we have completely defined the position of the Sun at any date and time. We now need to convert from ecliptic coordinates to something more convenient when pointing a telescope, and so we will now look at rotating angles between different coordinate systems.

C.3 Rotation Matrices and Coordinate Conversion

The standard equations for converting between coordinate systems typically consider only two dimensions, for example azimuth and altitude to right ascension and declination. Since our aim is to take absolute pointing information from inertial sensors such as accelerometers, we need to consider rotations in all three dimensions.

The most flexible method for converting between coordinate systems makes use of matrices. Rotation matrices can be complicated and confusing. Not mathematically, since they involve simple matrix manipulation, but in terms of definition of axes and the order of rotations. We also need to consider rotations that are *extrinsic*, i.e., rotations about the original coordinate system which remains fixed, and rotations that are *intrinsic*, i.e., rotations about the axes of the rotating body which changes with each rotation. We first need to define basic rotation matrices for the three Euler angles, for which we use the Tait-Bryan formalism typically used in aerospace since it represents the attitude of aircraft with respect to Earth's frame of reference, and this is analogous to the orientation of a telescope. The following three matrices apply counter-clockwise elemental rotations of vectors around a right-handed coordinate system. To continue the aerospace analogy,

we define the z-axis as a rotation in *yaw* which changes the heading by ψ,

$$\mathbf{R}_z(\psi) = \begin{pmatrix} \cos\psi & \sin\psi & 0 \\ -\sin\psi & \cos(\psi) & 0 \\ 0 & 0 & 1 \end{pmatrix} , \tag{C.8}$$

the y-axis a rotation in *pitch* by θ,

$$\mathbf{R}_y(\theta) = \begin{pmatrix} \cos\theta & 0 & -\sin\theta \\ 0 & 1 & 0 \\ \sin\theta & 0 & \cos(\theta) \end{pmatrix} , \tag{C.9}$$

and the x-axis a rotation in *roll* by ϕ,

$$\mathbf{R}_x(\phi) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\phi & \sin\phi \\ 0 & -\sin\phi & \cos(\phi) \end{pmatrix} . \tag{C.10}$$

In terms of astronomy, these angles are equivalent to azimuth, altitude, and field rotation, respectively. Any final orientation can be achieved by combining these three elemental rotations. It is important to consider the order of rotations. The rotation sequence $\mathbf{R}_{xyz}$,

$$\begin{aligned} \mathbf{R}_{xyz} &= \mathbf{R}_x(\phi)\mathbf{R}_y(\theta)\mathbf{R}_z(\psi) , \\ &= \begin{pmatrix} \cos\theta\cos\psi & -\cos\theta\sin\psi & \sin\theta \\ \cos\phi\sin\psi + \cos\psi\sin\phi\sin\theta & \cos\phi\cos\psi - \sin\phi\sin\theta\sin\psi & -\cos\theta\sin\phi \\ \sin\phi\sin\psi - \cos\phi\cos\psi\sin\theta & \cos\psi\sin\phi + \cos\phi\sin\theta\sin\psi & \cos\phi\cos\theta \end{pmatrix} , \end{aligned} \tag{C.11}$$

is most common and known as the aerospace rotation sequence, where rotation is first in yaw, then pitch, then roll. However it can be considered to be either combined intrinsic rotations about axes x, then y', then z'', or a combination of extrinsic rotations about the

original axes z, then y, then x. Likewise, the rotation sequence $\mathbf{R}_{zyx}$,

$$\begin{aligned}\mathbf{R}_{zyx} &= \mathbf{R}_z(\psi)\mathbf{R}_y(\theta)\mathbf{R}_x(\phi)\ , \\ &= \begin{pmatrix} \cos\psi\cos\theta & \cos\psi\sin\theta\sin\phi - \cos\phi\sin\psi & \sin\psi\sin\phi + \cos\psi\cos\phi\sin\theta \\ \cos\theta\sin\psi & \cos\psi\cos\phi + \sin\psi\sin\theta\sin\phi & \cos\phi\sin\psi\sin\theta - \cos\psi\sin\phi \\ -\sin\theta & \cos\theta\sin\phi & \cos\theta\cos\phi \end{pmatrix}\ ,\end{aligned} \tag{C.12}$$

defines the final attitude from rotations in reverse order. After rotation the new Euler angles are extracted from the elements of the rotation matrix using,

$$\begin{aligned}\phi &= \tan^{-1}\left(\frac{\mathbf{R}_{21}}{\mathbf{R}_{22}}\right), \\ \theta &= \sin^{-1}\left(-\mathbf{R}_{20}\right), \\ \psi &= \tan^{-1}\left(\frac{\mathbf{R}_{10}}{\mathbf{R}_{00}}\right),\end{aligned} \tag{C.13}$$

where ϕ, θ, and ψ are roll, pitch, and yaw respectively as before, and the subscripts to rotation matrix $\mathbf{R}$ represent the row and column elements of the matrix. The ATAN2 function is used when computing roll and yaw, in order to restrict these rotations between $\pm 180°$. Pitch angles are defined as being restricted between $\pm 90°$. We will now use these rotation matrices to convert between astronomical coordinate systems.

C.3.1 Horizon $\leftrightarrow$ Equatorial

Converting from Horizon to Equatorial coordinates is a two-stage process. First we correct for the observer's geographical location by rotating azimuth and altitude by the observing latitude (Duffett-Smith, 1988, p. 50, Table 4-A),

$$\begin{aligned}\mathbf{R}_{A,a\rightarrow H,\delta} &= \mathbf{R}_z(\pi)\mathbf{R}_y\left(\frac{\pi}{2} - \theta\right)\ , \\ &= \begin{pmatrix} -\sin(\theta) & 0 & \cos(\theta) \\ 0 & -1 & 0 \\ \cos(\theta) & 0 & \sin(\theta) \end{pmatrix}\ ,\end{aligned} \tag{C.14}$$

where θ is the observer latitude, H the hour-angle, and δ the declination. Hour-angle is a measure of the time since the object crossed the meridian, and so it increases throughout the day. When H is zero, the object is on the meridian and is said to *transit*, at which point its altitude is at a maximum. This rotation matrix is symmetric (i.e., it is a square matrix whose transpose is equal to itself) and so it can also be used in the reverse direction to convert from hour-angle and declination to azimuth and altitude.

The second stage is a rotation to correct for the observer's longitude and local-time. Hour-angle and declination are rotated by the local sidereal time, and the y-axis is flipped to convert from westward notation to eastward notation (Duffett-Smith, 1988, p. 50, Table 4-B),

$$\begin{aligned} \mathbf{R}_{H,\delta\rightarrow\alpha,\delta} &= \mathbf{R}_z(\mathrm{LST}) \ , \\ &= \begin{pmatrix} \cos(\mathrm{LST}) & \sin(\mathrm{LST}) & 0 \\ \sin(\mathrm{LST}) & -\cos(\mathrm{LST}) & 0 \\ 0 & 0 & 1 \end{pmatrix} \ , \end{aligned} \tag{C.15}$$

where α is right ascension, and LST is the local sidereal time of the observer. Again the rotation matrix is symmetric and so it can also be used in the reverse direction to convert from right ascension and declination to hour-angle and declination. Note that since we are calculating three dimensional rotation, we can also determine the degree of field rotation between the coordinate systems.

Sidereal time is the hour angle of the vernal equinox. A sidereal clock runs at a rate such that any star returns to the same point in the sky after exactly 24 hours as measured by the clock. This rate is slightly faster than the equivalent rate measuring transits of the Sun through the meridian of a particular observer due to the Earth moving nearly one degree in its orbit around the Sun, meaning the Earth has to turn slightly more than one full rotation on its axis between each solar transit. A sidereal day is slightly shorter than a solar day, with 24 hours of sidereal time taking just 23 hours 56 minutes of solar time. This means there are approximately 365.25 solar days in one year, but the Earth makes

approximately 366.25 complete revolutions about its axis. Local sidereal time corrects for the observer's longitude by adding the longitude (measured in hours) to Greenwich mean sidereal time (GST).

These two matrices can be combined to form a single-step conversion from azimuth and altitude (A, a) to right ascension and declination (α, δ),

$$\begin{aligned}\mathbf{R}_{A,a\to\alpha,\delta} &= \begin{pmatrix} -\sin(\theta) & 0.0 & \cos(\theta) \\ 0.0 & -1.0 & 0.0 \\ \cos(\theta) & 0.0 & \sin(\theta) \end{pmatrix} \cdot \begin{pmatrix} \cos(\mathrm{LST}) & \sin(\mathrm{LST}) & 0.0 \\ \sin(\mathrm{LST}) & -\cos(\mathrm{LST}) & 0.0 \\ 0.0 & 0.0 & 1.0 \end{pmatrix} \\ &= \begin{pmatrix} -\cos(\mathrm{LST})\sin(\theta) & -\sin(\mathrm{LST}) & \cos(\mathrm{LST})\cos(\theta) \\ -\sin(\mathrm{LST})\sin(\theta) & \cos(\mathrm{LST}) & \sin(\mathrm{LST})\cos(\theta) \\ \cos(\theta) & 0.0 & \sin(\theta) \end{pmatrix} ,\end{aligned} \tag{C.16}$$

where as before θ is the observer's latitude and LST is the local sidereal time. This matrix is no longer symmetric and so the conversion works only in the forward direction. However, calculating the transposed matrix allows azimuth and altitude to be obtained from right ascension and declination.

C.3.2 Ecliptic $\leftrightarrow$ Equatorial

Converting between geocentric ecliptic longitude λ and geocentric ecliptic latitude β, and right ascension and declination, requires a clockwise rotation correcting for the obliquity of the ecliptic plane (Duffett-Smith, 1988, p. 50, Table 4-C),

$$\begin{aligned}\mathbf{R}_{\lambda,\beta\to\alpha,\delta} &= \mathbf{R}_x(\epsilon) , \\ &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos(\epsilon) & \sin(\epsilon) \\ 0 & -\sin(\epsilon) & \cos(\epsilon) \end{pmatrix} ,\end{aligned} \tag{C.17}$$

where ϵ is the obliquity of the ecliptic, and as before α is right ascension, and δ the declination. Again, the transposed matrix allows the conversion in the opposite direction.

The obliquity is the axial tilt of the Earth between the equator and the orbital plane – the cause of the seasons throughout a year. The value changes between 22.1° to 24.5° on a timescale of approximately 41 000 years. The value of ϵ on J2000 epoch (2000 January 1.5, Julian date 2 451 545.0) was 23.439 292° (23° 26′ 21.448″), and it is currently decreasing.

Although the value changes only very slowly, when making a conversion between ecliptic and equatorial coordinates we should calculate the correct value on the date of the conversion. Laskar (1986) developed an expression to calculate the obliquity,

$$\begin{aligned}\epsilon = {} & 23^\circ\,26'\,21.448'' - 4680.93''\,T - 1.55''\,T^2 + 1999.25''\,T^3 - 51.38''\,T^4 \\ & - 249.67''\,T^5 - 39.05''\,T^6 + 7.12''\,T^7 + 27.87''\,T^8 + 5.79''\,T^9 + 2.45''\,T^{10}\ ,\end{aligned} \tag{C.18}$$

where T is multiples of 10 000 Julian years from J2000. This produces values accurate to better than 0.02″ over 1000 years, and within a few arcseconds over 10 000 years. Since we do not need such precision we can make use of a simpler expression developed by Lieske et al. (1977) and described by Duffett-Smith (1988, p. 41),

$$\epsilon = 23^\circ\,26'\,21.45'' - 46.815''\,T - 0.0006''\,T^2 + 0.001\,81''\,T^3\ , \tag{C.19}$$

where T is now the number of Julian centuries from J2000.

We now have all the information required to determine the position of the Sun in both horizon and equatorial coordinates on any date and time.

Appendix D

Refereed Publication

Chapters 2 and 4 include text and figures from the publication "Performance of the *Birmingham Solar-Oscillations Network* (BiSON)" (Hale et al., 2016) in the journal *Solar Physics*. The published article was first released on 2015 December 7, and is included here for reference. Data from the article are available from the University of Birmingham eData archive (Hale, 2015a). Whilst this publication is a result of collaborative working, the content is entirely the author's own work. Co-authors provided guidance on historical accuracy, data processing, and manuscript review.

Solar Phys (2016) 291:1–28
DOI 10.1007/s11207-015-0810-0

Performance of the *Birmingham Solar-Oscillations Network* (BiSON)

S.J. Hale[1] · R. Howe[1] · W.J. Chaplin[1] · G.R. Davies[1] · Y.P. Elsworth[1]

Received: 25 June 2015 / Accepted: 20 October 2015 / Published online: 7 December 2015

Abstract The *Birmingham Solar-Oscillations Network* (BiSON) has been operating with a full complement of six stations since 1992. Over 20 years later, we look back on the network history. The meta-data from the sites have been analysed to assess performance in terms of site insolation, with a brief look at the challenges that have been encountered over the years. We explain how the international community can gain easy access to the ever-growing dataset produced by the network, and finally look to the future of the network and the potential impact of nearly 25 years of technology miniaturisation.

Keywords Helioseismology, observations · Oscillations, solar

1. Introduction

The *Birmingham Solar Oscillations Network* (BiSON) has now been operating continuously as a six-station network for well over twenty years, recording high-quality spatially unresolved, or "Sun-as-a-star" helioseismic data. It therefore now seems timely to update our previous article on BiSON performance (Chaplin *et al.*, 1996). We present updated results on the temporal coverage and noise performance of the individual sites and the network as a whole and reflect on what we have learned from more than two decades of experience in operating a semi-automated ground-based observing network. These data are available to the wider community through the BiSON Open Data Portal.

The early history of the network has been outlined by Chaplin *et al.* (1996). In brief, early campaign-style observations from two sites (Haleakala and Izaña) were followed by the addition of an automated observing station in Carnarvon, Western Australia, and later by the deployment of more standardised stations in Sutherland, South Africa, Las Campanas, Chile, and Narrabri, Australia, over the period 1990 – 1992. The instrument from Haleakala was moved to California and installed in the 60-foot tower at the Mount Wilson Hale Observatory

S.J. Hale
s.j.hale@bham.ac.uk

[1] School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, United Kingdom

in 1992. With occasional instrument upgrades, the network has been operating continuously since then, providing unresolved-Sun helioseismic observations with an average annual duty cycle of about 82 %.

2. Observational Helioseismology: A Brief History

Oscillations in the velocity field across the Sun were first discovered by Robert Leighton (Leighton, Noyes, and Simon, 1962) using the spectroheliograph developed by George Ellery Hale at the Mount Wilson Hale Observatory in California. The now accepted explanation for these oscillations was developed by Roger Ulrich and John Leibacher, both independently suggesting that sound waves could be generated in the convection zone of the solar interior (Ulrich, 1970; Leibacher and Stein, 1971). Interest in the new field of helioseismology grew quickly, culminating in a dataset six days in length collected from the South Pole (Grec, Fossat, and Pomerantz, 1980), which at the time was the longest and most detailed set of continuous observations available.

It soon became clear that long-term continuous observations were required. The Birmingham group was the first to begin construction of a network of ground-based observatories, beginning with Izaña in Tenerife in 1975 and culminating in six operational sites in 1992.

Other groups have also had success with ground-based networks. Fossat and colleagues went on to deploy the *International Research of Interior of the Sun* (IRIS; Fossat, 1991) network. The operational strategy of IRIS was different from BiSON, requiring full participation of local scientists responsible for each instrument as opposed to automation. The instrumentation operated in a similar manner to BiSON, observing spatially unresolved global low angular-degree oscillations using an absorption line of sodium, rather than the potassium line used by BiSON. The IRIS network spanned six sites in total (up to nine sites when considering additional collaboration (Salabert *et al.*, 2002a)) and was operational until 2000 (Salabert *et al.*, 2002b; Fossat and IRIS Group, 2002).

Leibacher and colleagues later developed the *Global Oscillation Network Group* (GONG; Harvey *et al.*, 1996), a six-site automated network using resolved imaging to observe modes of oscillation at medium degree (up to $l \approx 150$), complementary to the existing low-degree networks. The sites were chosen in early 1991, and deployment began in 1994 with instruments coming online throughout 1995. The GONG network was upgraded in 2001 – 2002 to observe up to around $l = 1000$, and is still in operation today.

For completeness, we should also note the LOWL project (Tomczyk, Schou, and Thompson, 1995), which observed at medium degree from one or two sites between 1994 and 2004, and the *Taiwanese Oscillations Network* (Chou *et al.*, 1995) for high-degree observations, which was deployed between 1993 and 1996, but operated for only a few years.

Several space-based missions have also been successful. In the early 1980s the *Active Cavity Radiometer Irradiance Monitor* (ACRIM; Willson, 1979) onboard the NASA *Solar Maximum Mission* spacecraft was sufficiently precise to detect the small changes in intensity caused by the oscillations. Later, the *Solar and Heliospheric Observatory* (SOHO), a joint project between ESA and NASA, was launched in December 1995 and began normal operations in May 1996 (Domingo, Fleck, and Poland, 1995). Among the suite of instruments onboard the spacecraft are three helioseismic instruments: *Global Oscillations at Low Frequencies* (GOLF; Gabriel *et al.*, 1995) and *Variability of solar Irradiance and Gravity Oscillations* (VIRGO; Fröhlich *et al.*, 1995) for low-degree oscillations and *Michelson Doppler Imager* (MDI; Scherrer *et al.*, 1995) for observations at medium and high degree. The mission was originally planned for just two years, but it is still in operation and currently has a

mission extension lasting until December 2016, at which point it will have been in service for over twenty years. MDI ceased observations in 2011 when it was superceded by the *Helioseismic and Magnetic Imager* (HMI; Schou *et al.*, 2012) onboard the *Solar Dynamics Observatory* (SDO), but GOLF and VIRGO are still in use.

3. Designing an Automated Robotic Network

The principle of using resonant scattering spectroscopy to achieve stable and precise measurements of the line-of-sight velocity of the solar atmosphere was first proposed by Isaak (1961). The optical design of the instrument first used at Pic-du-Midi in 1974 to detect long-period solar oscillations is described by Brookes, Isaak, and van der Raay (1976). The basic observational parameter is the Doppler shift of the solar potassium Fraunhofer line at 770 nm. This is achieved through measurement of intensity over a very narrow range in wavelength that sits in the wings of the potassium line so that intensity changes with Doppler shift, and through comparison with the same transition in a potassium vapour in the laboratory, the change in intensity is calibrated to become a measure of velocity. Following initial tests at Pic-du-Midi, the spectrometer was then relocated to the Observatorio del Teide, Tenerife, during 1975. The updated apparatus is described by Brookes, Isaak, and van der Raay (1978a).

The six-station network of today was completed in 1992. There are two stations in each 120° longitude band, and all of the sites lie at moderate latitudes, around ±30°. The oldest site in the Birmingham Network, Izaña, has now been collecting data for nearly forty years.

The original control system was based around a 40-channel scaler module used for counting pulses from a photomultiplier tube (McLeod, 2002). Timing was controlled by a quartz clock, related to GMT at least once per day, and producing pulses at exactly 1 s intervals. At the end of each interval two relays would change state and reverse the voltage across an electro-optic modulator, and the data gate would be incremented to the next channel of the scaler. Once all 40 channels of the scaler had been filled, the contents would be destructively written out to magnetic tape, a process that required a further 2 s. The process would then repeat, with each block of data separated by 42 s. The cadence was changed to 40 s at the beginning of 1990. This allows for simpler concatenation of data from different sites since there are an integer multiple of 40 s in a day, and so it provides a network time-standard.

The system was computerised in 1984 using the BBC Microcomputer. The BBC Micro was originally commissioned on behalf of the British Broadcasting Corporation as part of their Computer Literacy Project, and it was designed and built by the Acorn Computer company. Despite the BBC Micro being discontinued in 1994, the spectrometer continued operating in this configuration until 2003, when the computer finally failed and was replaced by a modern PC. A dedicated PIC-based interface was designed to enable the PC to communicate with the original scaler system (Barnes, Jackson, and Miller, 2003, 2004), and this remains the operating configuration today.

For the subsequent fully automated solar observatories, the data-acquisition and control system was initially based around a Hewlett-Packard 3421A data acquisition unit, known colloquially as a "data logger". By the early 1990s these were retired from service and replaced with a standard desktop personal computer and a Keithley System 570 digital input/output interface. The computer ran Microsoft DOS. The limitations of the operating system meant that data could only be retrieved by Birmingham during a scheduled time

window when the computer would switch from running the data-acquisition program to running a data-transfer program. Over the years a variety of means have been used to return the data from the sites to Birmingham, ranging from tape cartridges or floppy disks sent by post, through direct dial-up modem connections over international phone lines, to the modern internet.

Initially, the instrument at the first automated dome in Carnarvon observed through a glass window. Although the concern regarding site security was low, the window was considered necessary for safe operation. Despite being cleaned regularly by the on-site support, marks on the glass were detrimental to the data, and the decision was eventually made to remove the window and operate in the conventional style of a simple shutter opening. This meant that since the whole system is designed to run completely unattended, careful weather monitoring was required. Both rain and wind sensors are used to close the dome in the event of precipitation or excessive wind. All of the subsequent observatories operated without window glass.

Naturally, repairs have been made over the years and upgrades installed. Probably the most significant is the upgrade to the control software. The original Keithley System 570 data-acquisition system in Carnarvon ran for more than a decade of continuous use, but the units eventually became obsolete as PCs failed and the replacements lacked the required ISA interface. At the same time, the MS-DOS operating system on which the original dome control software relied itself became obsolete, and the Windows software that superseded it for home and office uses was unsuitable for system automation. A new dome control system known as the "Zoo" was developed in the late 1990s (Miller, 2002, 2003) to run under the recently released GNU/Linux operating system (Stallman, 1983; Torvalds, 1991), and this continues in use to the present day.

Much of the old analogue electronics have been gradually replaced with new digital variants. Rather than the whole system consisting of modules in a central rack and communicating through a single interface, the new designs are independent with embedded microcontrollers and communicate with the PC through dedicated RS-232 serial ports. This makes subsequent repairs and upgrades considerably easier since units can be inspected in isolation.

In the following sections we look back over the performance of the network, including temporal coverage and noise levels, and also discuss some of the significant events during the life of the stations.

4. Data Quality Metrics

Before we can discuss data quality, we must first define some quality metrics. There are two standard metrics used by BiSON. These are the *five-minute figure of merit* known simply as the FOM, and the mean high-frequency noise level.

Data from BiSON are collected on a cadence of 40 s, giving an upper limit in the frequency domain (Nyquist frequency) of 12.5 mHz. The FOM is a signal-to-noise ratio and is defined as the total power in the main "five-minute" signal band (2 mHz – 5 mHz) divided by the noise (5.5 mHz – 12.5 mHz). When considering the mean noise level in isolation, we look at the mean power in the high-frequency noise (10.0 mHz – 12.5 mHz). We refer to these definitions of *FOM* and *mean noise* throughout this article.

Both of these metrics require a certain volume of data in order to ensure that the quality estimates are reliable and that meaningful comparisons can be made between data from

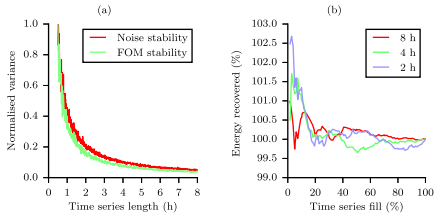

Figure 1 Measured FFT stability over varying dataset length (a) and fill (b). For a valid comparison between datasets from different days and different sites, a time series needs to be at least 3 h in duration and have a fill of at least 25 %.

different days and different instruments. As the length of a dataset decreases, different realisations of random noise begin to have a stronger effect on the estimate of quality. Eventually, the quality estimate becomes so variable as to be meaningless. To determine the minimum dataset length required to make reliable comparisons, a simple artificial dataset was produced with one thousand realisations of random noise. The FOM and mean noise were calculated for each realisation, and the variance in the FOM and mean noise recorded. This was done for dataset lengths varying from thirty minutes up to eight hours. The normalised results are shown in Figure 1(a). Both the quality metrics appear to stabilise at dataset lengths of a minimum of 3 h, and so this was selected as the minimum length that could be reliably used for quality comparison in this article.

An additional problem with data from BiSON is that due to variable weather conditions it is usually not continuous. To be able to compare absolute noise levels, we need to rescale the power spectrum produced by the FFT to compensate for any missing data. This is done by simply dividing by the percentage fill (*i.e.* the percentage of the total observing time where data are available). To determine how low the fill can be whilst still providing a meaningful estimate of the total energy in an equivalent gap-free dataset, a similar test was used with simple artificial data and one thousand realisations of random noise. A variable size gap was created in three datasets of two, four, and eight hours in total length. The results for the percentage of total energy recovered compared with the original time series are shown in Figure 1(b). Even with very low fills it is possible to recover an estimate of the original total energy to within a few percent. For this article, a fill of at least 25 % was selected as the minimum requirement.

5. Site Performance

5.1. Izaña, Tenerife

As we saw at the beginning of Section 3, the instrument that was to become the first node of the BiSON was installed at the Observatorio del Teide, Tenerife, in 1975. The Birmingham

group were the first to establish a global network of ground-based observatories dedicated to helioseismology.

George Isaak, the then head of the *High-Resolution Optical Spectroscopy* (HiROS) research group, was already considering the development of a permanent network to expand beyond Tenerife long before his seminal article in 1979 on global studies of the five-minute oscillation (Claverie *et al.*, 1979). Campaign-style operations continued throughout 1978 and 1979 at Pic-du-Midi in the French Pyrenees, and a short run at Calar Alto in Spain in 1980. In 1981 the group secured funding to operate a second site on the island of Haleakala in Hawaii, at the Mees Observatory. The instrument at Haleakala operated in much the same way as Mark I in Tenerife. From the two sites together, data were collected for 88 days, producing the longest time series and most highly resolved power spectrum that had been achieved up to that time. However, this was still a long way from year-round complete coverage.

Mark I is housed in the Pyramid Building at the Observatorio del Teide and is run by Pere Pallé and his team of observers (Roca Cortés and Pallé, 2014). Light is collected via two mirrors, known as a cœlostat. The beam is projected through an open window into the apex of the pyramid, where Mark I sits on an optical bench. At the beginning of each observing session, the on-site observer needs to uncover and align the mirrors and start the system. Operator presence is also required throughout the day to close the mirrors in the event of bad weather, and at the end of the observing session.

Despite the higher photon shot noise level compared to our more modern sites, Izaña has been and still is a work-horse of the BiSON network. The site duty-cycle is shown in Figure 2. The duty-cycle is plotted in terms of both number of observational hours and percentage insolation. The insolation compares only the potential daylight hours against actual observational hours and does not differentiate between poor weather conditions and any periods of instrumental failure. Tenerife provides exceptional weather conditions over the summer months, but is rather poor throughout the winter where the conditions become much more variable. There are regular "holes" in the site window function at midday during the spring and autumn months that are caused by cœlostat shadowing. The secondary mirror has two mounting positions, from above or from below, which allows an unobstructed daily run to be achieved throughout the winter and summer months. However, during the change-over period between the two phases, the secondary mirror unavoidably shadows the primary.

The figure of merit (FOM) and mean noise levels for Izaña are shown in Figure 3. The seasonal variation in noise level, and subsequent change in FOM, is due to an effect described by Chaplin *et al.* (2004, 2005). The basic measurement of a BiSON RSS is that of intensity change due to the shift of a solar Fraunhofer line. As the Fraunhofer line shifts due to the line-of-sight relative motion of the solar surface with respect to the laboratory, the measured intensity changes. A data pipeline calibrates the intensity measurements into velocity, taking into account the non-linearity of the Fraunhofer line shape. However, the propagation of noise on the original intensity measurement through the calibration process is non-linear (Hoyng, 1989) due to the varying types of noise. The system encounters white noise, for example due to photon statistics and analogue-to-digital conversion, multiplicative noise due to gain fluctuations, and additive noise due to offset fluctuations, in addition to more random effects due to pointing errors and temperature fluctuations. As the Fraunhofer line shifts daily due to Earth's rotation and seasonally due to the eccentricity of Earth's orbit, the line gradient at the operating point of the instrument changes. This causes the noise level in the derived quantities to change even if the noise level in the basic intensity measurement remains constant. Aside from the seasonal variation, the noise levels from

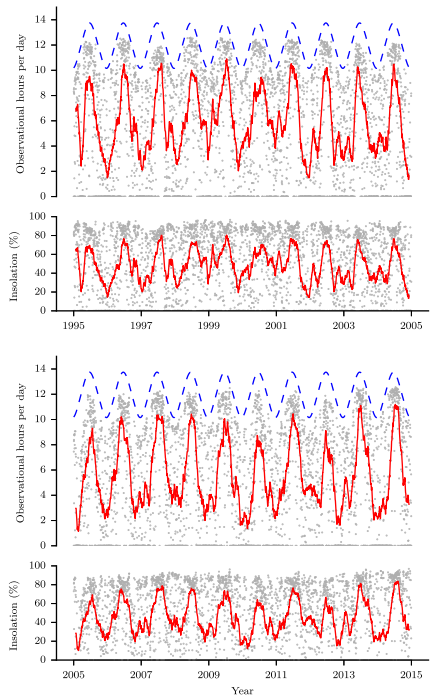

Figure 2 Izaña duty cycle as a function of date, plotted in hours per day, and as a percentage of potential daylight hours. There is one grey dot per day, and the solid red curve represents a 50-day moving mean. The dashed blue curve shows potential daylight hours.

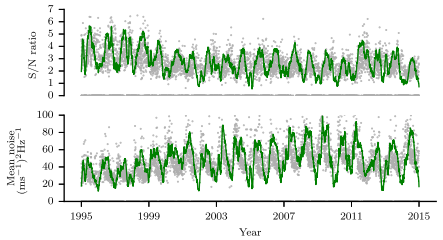

Figure 3 Izaña data quality as a function of date. Top: Signal-to-noise ratio, higher is better. Bottom: Mean noise level, lower is better. There is one grey dot per day, and the solid green curve represents a 50-day moving mean.

Izaña show remarkable stability and offer an unprecedented temporal baseline of almost 40 years.

5.2. Carnarvon, Western Australia

When considering expanding to a global network, the group realised that it would not be practical to operate such a network manually. Requiring an observer to be present on site, all day every day for 365 days per year, would be very expensive and potentially unreliable. As the number of network nodes increases, the number of observers required to operate them all year round would become untenable. The key to a reliable network would be automation. For the early 1980s, this was an ambitious endeavour. If successful, it would be one of the first automated astronomical telescopes anywhere in the world, in any field.

From 1981 the group worked on this innovative new design, and by the summer of 1983, a prototype was ready to test. The new instrument would point directly at the Sun and move on an equatorial mount like a classical astronomical telescope. Moving a mount under computer control would be significantly easier than aligning and pointing mirrors. Testing in Haleakala proved that the system worked, and the group began looking for an installation site. Western Australia was selected as a good longitude to complement coverage from Pic-du-Midi and Haleakala. Four sites were investigated: Woomera Rocket Testing Range, Learmonth, Exmouth, and Carnarvon. Isaak decided that Woomera was too dusty, but the other three towns were all suitable. The group settled on Carnarvon, some 900 km north of Perth. The group arrived in Carnarvon in 1984. Similar testing was carried out to that which had been successful in Haleakala. By the end of the year it was clear that full automation was practical, and the decision was made to push ahead.

Initially, the data collection systems operated on a 42 s cadence like the existing systems in Izaña and Haleakala. Carnarvon was migrated to the newer 40 s cadence in April 1992. Data obtained previous to this date are interpolated onto the newer standard cadence. Over

the months following installation, the system worked extremely well. It produced high-quality data and proved to be reliable. Most glitches were due to the Australian wildlife – rodents eating through cables, or cockatoos making nests in the upper parts of the dome. The Sun itself caused some problems, with certain types of connector and cable insulation quickly decaying under constant exposure to solar UV. A simple and cheap solution was found for this, which was to wrap exposed cables and components in aluminium cooking foil. The goal of demonstrating that an automated system could work, and work well, had been successfully achieved. Funding to roll out further stations was quickly secured.

Carnarvon has not been without its problems. The PC failed in August 2002, and this is the first drop-out in the Carnarvon duty cycle (Figure 4). In May 2005 a freak rainstorm emptied over three inches of rain in just two hours. This is almost the same amount of rain that Carnarvon expects in a whole year. It was not a good time to find out that the rain detector had failed, and so the dome did not close (Barnes and Hale, 2005; New and Hale, 2006). The whole dome was thoroughly flooded, completely destroying the control electronics for one instrument and severely damaging a second instrument. Water was poured out of some of the electronics. At the time, Carnarvon had two instruments in operation. The primary instrument was able to be repaired on-site after designing new detectors and control electronics in Birmingham. However, the secondary instrument was written off and had to be returned to Birmingham (Barnes, Miller, and Jackson, 2007) for a considerable programme of repairs and upgrades. The new instrument was installed in 2009 (Barnes and Miller, 2009) and provided a substantial improvement in noise level and data quality (Figure 5).

In 2013 the land was bought by NBN Co. Ltd, the National Broadband Network of Australia, as a base for a new satellite internet station. A lease was approved between the group and NBN, and the dome was shut down for just over a month in October 2013 during the NBN construction in an attempt to limit the amount of dust and dirt entering the dome. The NBN compound has two large antennas that are due east of the dome and cause some shadowing in the early morning around the summer solstice. Despite this, the Carnarvon BiSON dome continues to collect data year-round.

5.3. Sutherland, South Africa

Following the success of the Carnarvon station, three more sites were commissioned. Some minor changes were made to the original design. The main building became rectangular and made of brick, rather than cylindrical and made from plywood clad with corrugated iron. This substantially increased the floor-area, allowing for more storage space and easier access to the control electronics. The mount was also made considerably larger, allowing for larger and heavier instrumentation.

The first station of this new design was built in Birmingham in 1988. It is located on the roof of the Poynting Physics building. Once the new-style dome had been finalised, the first site to be commissioned was in Sutherland, South Africa, at the South African Astronomical Observatory (SAAO) in 1989 (Davidson and Williams, 2004) and completed in 1990. The observatory itself was established in 1972 and is run by the National Research Foundation of South Africa. It is located in the Karoo Desert, in the Northern Cape of South Africa, approximately 350 km north-east of Cape Town.

Like Carnarvon, Sutherland started out with a Keithley System 570 data-acquisition system and later got its Zoo upgrade in 2006 (Barnes and Miller, 2006). New digital temperature controllers were installed in 2007 (Barnes, 2008, 2009), new counters in 2012 (Hale, 2012), and a new digital autoguider in 2013 (Hale, 2013).

The Sutherland duty cycle is shown in Figure 6. The drop-out in late 1997 through early 1998 was caused by a range of faults. The tracking motor on the mount failed, and

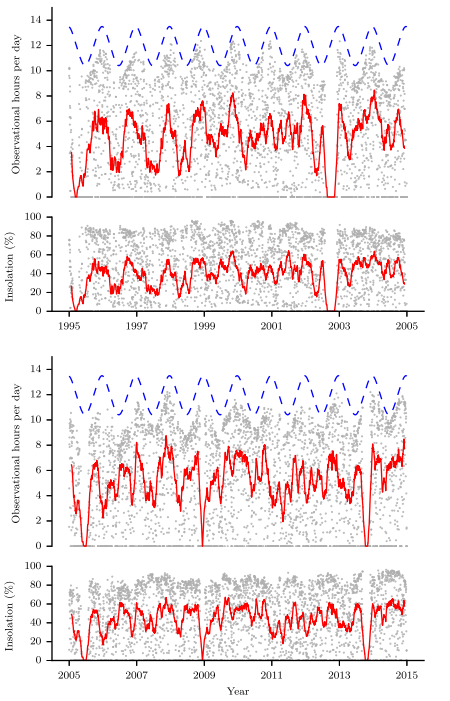

Figure 4 Carnarvon duty cycle as a function of date, plotted in hours per day, and as a percentage of potential daylight hours. There is one grey dot per day, and the solid red curve represents a 50-day moving mean. The dashed blue curve shows potential daylight hours.

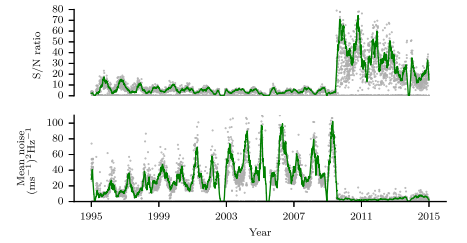

Figure 5 Carnarvon data quality as a function of date. Top: Signal-to-noise ratio, higher is better. Bottom: Mean noise level, lower is better. There is one grey dot per day, and the solid green curve represents a 50-day moving mean. The large step in 2009 is due to installation of a new upgraded spectrometer.

problems with the declination limit switches caused the mount to be unresponsive. Also, the dome moved out of alignment and began shadowing the instrument just before sunset. Later, the electronic polarisation modulator failed, and this was not replaced until late January 1998. In early 1999, a fault developed on the scaler system that counts the pulses produced by the detectors. Following extensive investigation, simply re-seating all the chips on the scaler cards fixed the problem, but not until almost a month of data had been lost. The problem reoccurred several times over the years, until the scaler system was finally replaced completely in 2012. The gap in 2007 was due to some downtime whilst the temperature control systems were upgraded, and in 2008 due to a heavy snow storm.

Looking at the data quality and noise levels (Figure 7), we see that the site has shown consistent performance with gradual improvement as systems were upgraded. The higher noise level in 2000 was due to a sticky declination gearbox. Weather conditions in Sutherland show solid year-round performance.

5.4. Las Campanas, Chile

After the completion of Sutherland, a third automated site was opened in Chile in 1991, at the Las Campanas Observatory operated by the Carnegie Institution for Science. Las Campanas Observatory is located in the southern Atacama Desert of Chile, around 100 km north-east of La Serena. The observatory was established in 1969 and was a replacement for the Mount Wilson Hale Observatory near Pasadena, which had started to experience too much light pollution from the growing city of Los Angeles. The main office is in Las Serena, whilst the headquarters remain in Pasadena.

Las Campanas received its Zoo upgrade in December 2005 (Hale and Miller, 2006), a new digital autoguider in February 2011 (Miller, 2011; Barnes and Miller, 2011), and a new temperature controller in 2015 (Hale, 2015b,c).

The Las Campanas duty cycle is shown in Figure 8. Las Campanas has had a range of problems. One notable experience was in July 1997, when a direct lightning strike to the dome destroyed several pieces of sensitive electronics. Whilst repairs were completed,

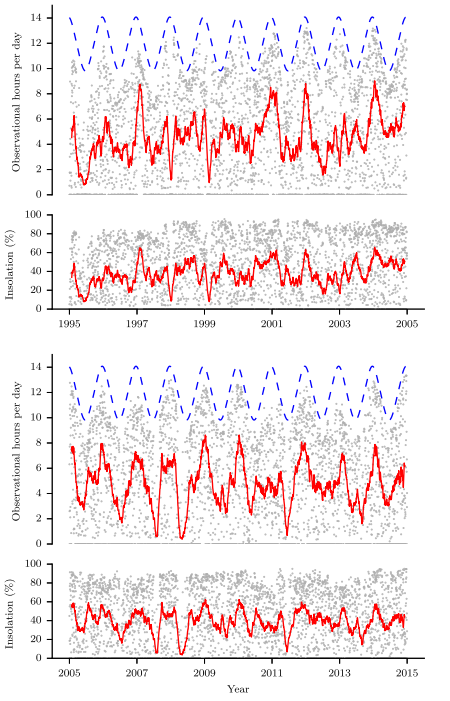

Figure 6 Sutherland duty cycle as a function of date, plotted in hours per day, and as a percentage of potential daylight hours. There is one grey dot per day, and the solid red curve represents a 50-day moving mean. The dashed blue curve shows potential daylight hours.

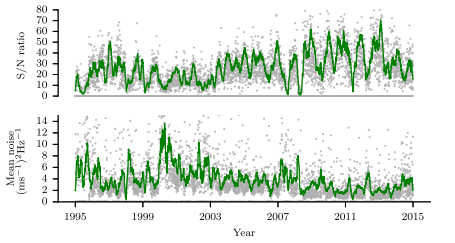

Figure 7 Sutherland data quality as a function of date. Top: Signal-to-noise ratio, higher is better. Bottom: Mean noise level, lower is better. There is one grey dot per day, and the solid green curve represents a 50-day moving mean.

weather conditions deteriorated with a heavy snow fall that left the observatory cut off by road, without internet or phone connections, and dependent on old diesel generators for electrical power (Miller, 1997). As the snow melted, the 4.2 kV step-down transformer for the BiSON dome was first flooded and subsequently destroyed by another lightning strike. The dome was without power for more than a month, until early November 1997.

Two further visits to finalise repairs were needed in November 1997 (Lines, 1998) and January 1998 (Miller, 1998), which were related to problems with the computer, the dome azimuth motor, and the water-loop system that stabilises the instrumentation temperatures. Noise problems in both detectors persisted until early 1999, clearly visible in the plot of data quality (Figure 9). A large part of 2000 was lost due to a broken declination gearbox on the mount and also to the failure of the dome azimuth motor. More electrical problems occurred in May 2014, when faults with both the shutter and blind limit switches caused the circuit breakers to trip repeatedly. A site visit was required to replace the limit switches and also work on additional problems with the water-loop pump and the uninterruptable power supplies (Hale, 2014b).

Las Campanas is the best-performing station in the network, consistently supplying duty cycles above 80 % in the summer and regularly above 40 % even during the winter months. From 2012 the noise performance has deteriorated slightly. A recent site visit indicated reduced performance from the potassium vapour cell, and this may need to be replaced soon.

5.5. Narrabri, NSW, Australia

The final fully automated site was installed in Narrabri, Australia, in 1992 (Williams, 2004). It is on the site of the Australia Telescope Compact Array (ATCA) at the Paul Wild Observatory operated by The Commonwealth Scientific and Industrial Research Organisation (CSIRO). The observatory is around 550 km north-west of Sydney.

Narrabri received its Zoo upgrade in June 2004 (Jackson and Miller, 2004), new temperature controllers in January 2010 (Barnes and Hale, 2010; Barnes, 2010), and a new

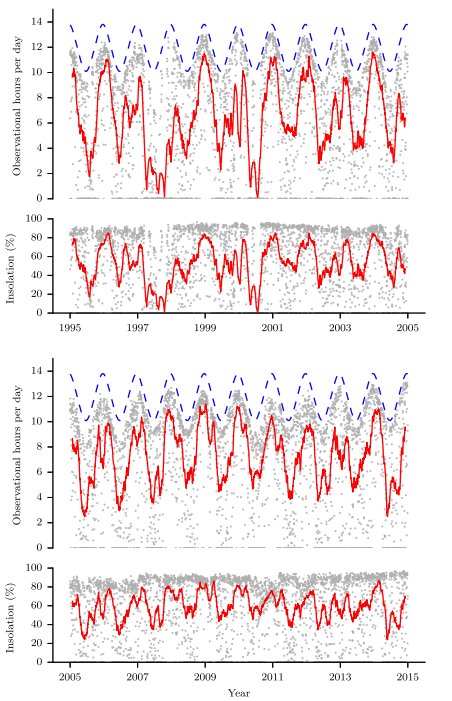

Figure 8 Las Campanas duty cycle as a function of date, plotted in hours per day, and as a percentage of potential daylight hours. There is one grey dot per day, and the solid red curve represents a 50-day moving mean. The dashed blue curve shows potential daylight hours.

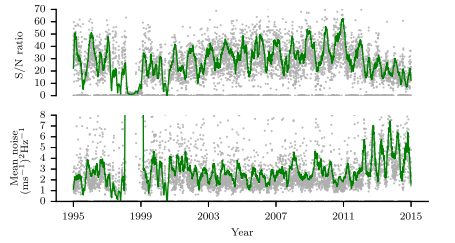

Figure 9 Las Campanas data quality as a function of date. Top: Signal-to-noise ratio, higher is better. Bottom: Mean noise level, lower is better. There is one grey dot per day, and the solid green curve represents a 50-day moving mean.

digital autoguider and counters in April 2013 (Hale and Davies, 2013). In February 2000, the dome blind motor failed (Miller, 2000) and was replaced with an identical spare from Birmingham, but failed again in March 2003, by which time the part had been discontinued. The dome manufacturer specified an alternative part for the blind mechanism, and this was installed in July 2003 (New and Isaak, 2003). Following several further motor failures, a change to the control system was identified as being required due to differences in the motor type (Jackson and Miller, 2004).

The Narrabri duty cycle is shown in Figure 10. The gap in 1997 is due to problems with the scalers and correct termination of the signal cables from the voltage-to-frequency output of the detectors. The missing data in 2003 and 2004 are due to the blind motor problems discussed earlier. In 2013 a fault on the anemometer in June caused problems with the weather module, and this kept the dome closed unnecessarily.

In the plot of data quality, Figure 11, the increased noise level in 2003 is due to a failure of an interference filter temperature controller, and in 2009 is due to a faulty power supply producing under-voltage power rails. Other than these faults, the site shows solid performance.

5.6. Mount Wilson, California, USA

The final change to arrive at the existing network configuration was made in 1992. The Mark III instrument was moved from Hawaii to the 60-foot tower at the Mount Wilson Hale Observatory in California (Elsworth, 1992). This is the very same observatory where Robert Leighton originally discovered the solar five-minute oscillations (Leighton, Noyes, and Simon, 1962). The observatory is located in the San Gabriel Mountains near Pasadena, around 80 km north-east of Los Angeles. The tower is operated by Ed Rhodes and his team of undergraduate volunteer observers. Like Izaña, the tower uses a cœlostat to direct light down the tower into the observing room below, and as such requires someone to open the dome and align the mirrors each morning. It also requires presence throughout the day to close the dome in the event of bad weather and at the end of the observing session.

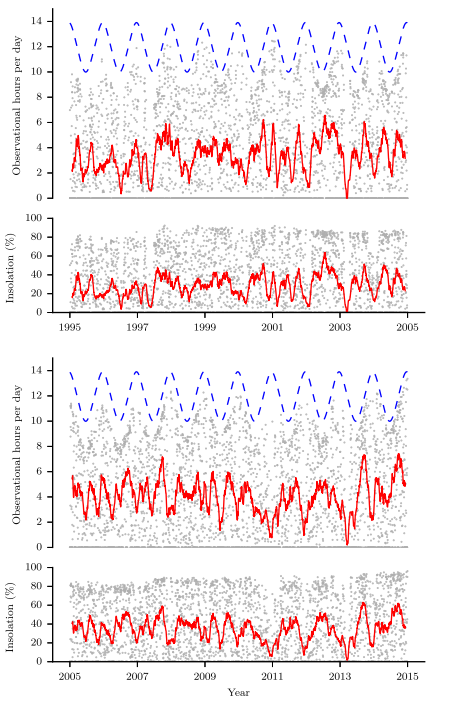

Figure 10 Narrbri duty cycle as a function of date, plotted in hours per day, and as a percentage of potential daylight hours. There is one grey dot per day, and the solid red curve represents a 50-day moving mean. The dashed blue curve shows potential daylight hours.

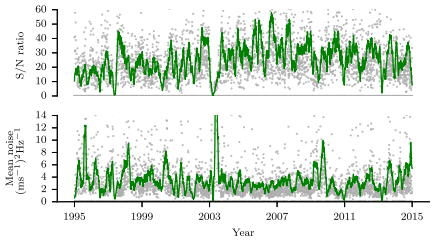

Figure 11 Narrabri data quality as a function of date. Top: Signal-to-noise ratio, higher is better. Bottom: Mean noise level, lower is better. There is one grey dot per day, and the solid green curve represents a 50-day moving mean.

The original optical configuration used by Leighton involved the cœlostat firing light 60 feet down the tower, a further 30 feet down into a pit where it would reflect from the spectroheliograph, and finally another 30 feet back up the pit to the observing room above, the total optical path length being 120 feet. The spectroheliograph is no longer in use, and so the current optical configuration is somewhat different. The primary instrument operated by Rhodes used a small objective lens near the base of the tower. Since it required only half the original optical path length, the cœlostat mirrors are now effectively oversized for their current usage. We can take advantage of this by picking off a small section of the beam and directing it to another instrument without affecting the operation of the main instrument. This is done using a periscope arrangement of mirrors mounted in the tower shaft taking a part of the beam and shifting it slightly to the south where it is directed back down the tower. Another mirror at the bottom reflects the light horizontally into the BiSON spectrometer. In total, five mirrors are used to direct light into the spectrometer.

There are regular gaps in the data at midday during the winter months, when the primary mirror has to be moved from the east side of the tower in the morning to the west in the afternoon, in order to avoid shadowing from the secondary mirror. Unlike Izaña, the secondary at Mount Wilson has only one mounting configuration. The Mount Wilson duty cycle is shown in Figure 12.

In July 1996 the Mark III instrument from Haleakala was retired and replaced with the tenth spectrometer designed by the group, code-named Klaus (Miller and New, 1999). The new design offered much-improved data quality in line with the other newer instruments in the network, and the reduction in noise level can be seen clearly in the plot of data quality (Figure 13).

The optical configuration with the pick-off mirrors is marginal, and very precise alignment is required for optimum performance of the spectrometer. Aligning five different mirrors over such large distances and ensuring no vignetting is difficult. Vibrations from the tower can also cause the mirrors or the spectrometer itself to move slightly, meaning the optical alignment has to be checked frequently. The large step in noise level during 2004 corresponds to an autoguider problem, where it failed to drive in right ascension. Mount

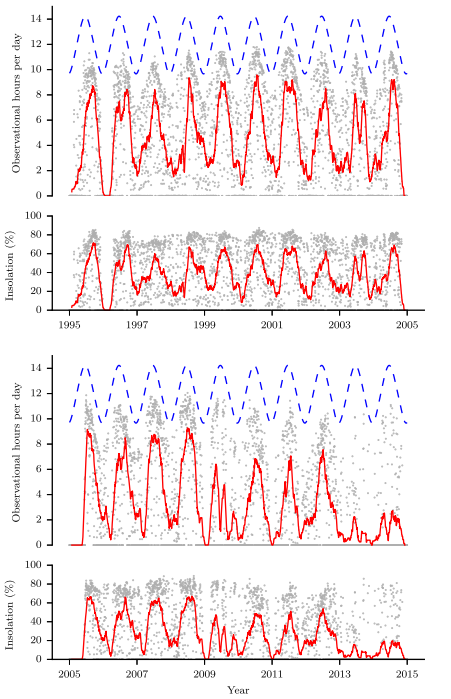

Figure 12 Mount Wilson duty cycle as a function of date, plotted in hours per day, and as a percentage of potential daylight hours. There is one grey dot per day, and the solid red curve represents a 50-day moving mean. The dashed blue curve shows potential daylight hours.

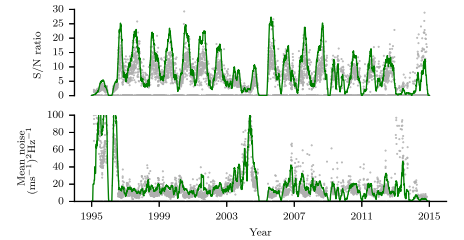

Figure 13 Mount Wilson data quality as a function of date. Top: Signal-to-noise ratio, higher is better. Bottom: Mean noise level, lower is better. There is one grey dot per day, and the solid green curve represents a 50-day moving mean. The large step in 1996 is due to installation of a new upgraded spectrometer, and in 2004 due to autoguider problems.

Wilson received its Zoo upgrade in September 2005 (Miller, 2005), which included alignment of all the optics returning the site to its original performance. In 2009 a variety of faults caused loss of data. The computer failed at the beginning of the year, and a site visit to replace it could not be arranged until March. Shortly after installing the new computer, the UPS failed in May and caused problems with the temperature controller. The on-site help investigated the problem, but could not locate the fault. Another site visit was arranged for July, where a blown fuse was found and the problem solved. At the end of July, the mirrors were removed for two weeks for re-aluminisation. Just at the point where everything was working, in August the entire observatory was evacuated due to the nearby Station Fire. There was concern that the entire observatory could be lost. Fortunately, that was not the case, thanks to efforts from the fire service and the US Forest Service.

Further guider problems were experienced in 2011, where one of the motors failed again and needed to be replaced. Through 2013 and 2014, significant problems were experienced with the autoguider electronics in the tower. On a site visit from Birmingham, the guider system built into the second flat of the cœlostat was completely rewired and refurbished, removing some very old and corroded cables and producing a significant decrease in noise level (Hale, 2014a). Significant damage to the primary mirror was also discovered during the visit. This is a serious problem since it is implicit in our analysis that we see the whole of the solar disc and that no part is vignetted. Since the Sun is rotating, any vignetting of the disc causes uneven weighting and produces an offset in the computed residuals through a process known as Doppler Imaging (Brookes, Isaak, and van der Raay, 1978b). For the Sun, when viewing a typical Fraunhofer line, the ratio in weighting of opposite sides of the disc is about 4:1. The instrument used at Mount Wilson is the same design as used at our other sites on equatorial mounts. It is optically designed to directly observe the Sun. At Mount Wilson, the cœlostat is approximately 60 feet away from the spectrometer, at the top of the tower. Optically this is near infinity, and so the spectrometer forms an image of the surface of the cœlostat mirrors, with the result that it is essential the mirrors be clean and in good condition. The damaged area of the mirror was causing part of the solar disc to be missing

from the light entering the spectrometer. The problem was compounded due to the unusual design of the tower cœlostat. Usually, the primary mirror performs both tracking and guiding to follow the Sun throughout the day. The Mount Wilson cœlostat separates these functions between the two mirrors – the primary mirror tracks whilst the secondary mirror guides. This means that any dirt or damage on the surface of the primary mirror will give rise to a signal that will be seen to oscillate at the frequency of the tracking error, as the autoguider on the secondary mirror compensates for tracking errors of the primary mirror. The worm-drive of the primary mirror uses 432 teeth, which puts the expected gear frequency at 2.5 mHz. This causes the vignetted part of the solar image to oscillate at 2.5 mHz, periodically varying the weighting of the solar disc, and corrupting the power spectrum at this frequency with sufficient amplitude to make the data unusable. Unfortunately, it is almost impossible to filter out this fault since it is within the main solar five-minute-oscillation band of interest, falling very close to one of the mode peaks and is not sufficiently coherent to be removed by subtracting a single sine wave. The fault was able to be mitigated by rotating the mirror in its housing to move most of the damage to an area not used by our pick-off mirrors. A similar fault occurs if the two mirrors are not aligned correctly at the start of an observing period. If the light from the primary mirror falls off the edge of the secondary mirror, then again the solar disc is vignetted and the 2.5 mHz tracking error becomes visible.

6. Whole Network Performance

Plots of the whole-network duty cycle and data window-function are shown in Figures 14 and 15. In a total of 7305 days, 45.5 % of the available time was covered by one site, 30.3 % by two sites, 6 % by three sites, 0.06 % by four simultaneous sites. The horizontal lines that run from 1995 – 2005 in Figure 15 are caused by a midday beam-chopper used to check the dark counts from the detectors during the day. From 2006 it was decided that this was no longer necessary, and so the regular gap is not present. Other regular gaps are those caused by cœlostat shadowing at Izaña and Mount Wilson, as discussed in Section 5. Figure 16 shows a histogram of the daily fill. Just 9 days had no coverage, and 633 days have a fill greater than 99 % of which 312 days achieved 100 %. The average fill is 82 % for the whole dataset.

The original data pipeline for calibration of raw data from the BiSON spectrometers through to velocity residuals is described by Elsworth *et al.* (1995). The next stage of analysis involves combining the residuals into an extended time series and transforming into the frequency domain where the mode characteristics can be analysed. This is described by Chaplin *et al.* (1997) and Hale (2003). An updated pipeline that includes correction for differential atmospheric extinction was produced by Davies *et al.* (2014a). By applying differential extinction correction, we have removed most of the low-frequency drifts in the dataset that were previously filtered using a 25-sample moving mean, and this allows further investigation of the very low frequency modes of oscillation that were previously lost in the noise background. We are now also able to make better use of weighted averaging of overlapping sites to produce a further improvement in the signal-to-noise ratio.

The entire network archive of velocity residuals has been regenerated using the latest data pipeline, and a new concatenated time series has been produced for the period from 1995 to the end of 2014 using the latest weighted-averaging techniques. A noise ceiling of 100 $(\mathrm{m\,s^{-1}})^2\,\mathrm{Hz^{-1}}$ was selected to reject data above this level, and this has reduced the overall fill from 82 % to 78 %. The plot of data quality and noise levels (Figure 17) shows

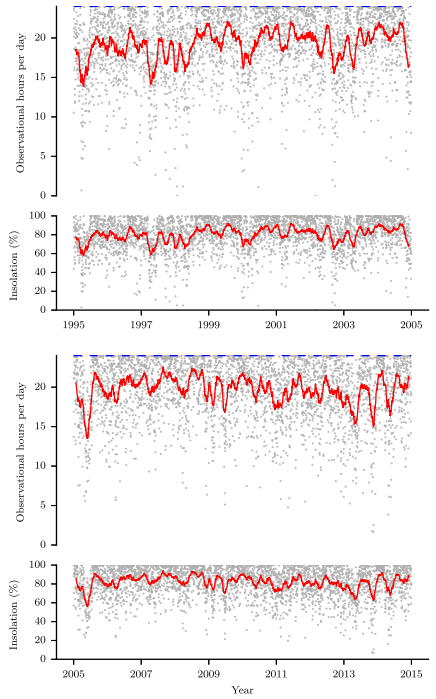

Figure 14 All-station duty cycle as a function of date, plotted in hours per day, and as a percentage of potential daylight hours. There is one grey dot per day, and the red curve represents a 50-day moving mean.

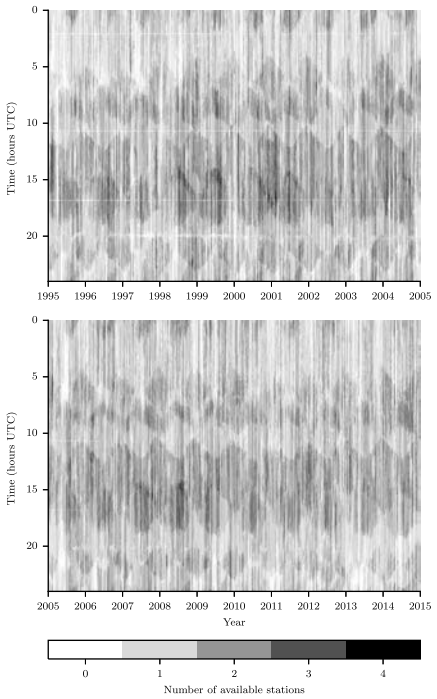

Figure 15 All-station data window-function.

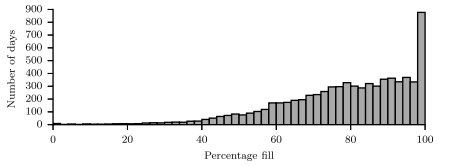

Figure 16 The distribution of fill per day for the overall time series.

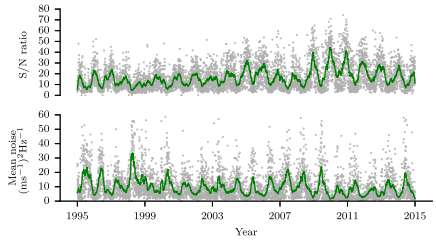

Figure 17 All station data quality as a function of date. Top: Signal-to-noise ratio, higher is better. Bottom: Mean noise level, lower is better. There is one grey dot per day, and the solid green curve represents a 50-day moving mean.

excellent stability through the entire period. Histograms of noise level and FOM (Figures 18 and 19) show the daily distribution of data quality.

The new data pipeline provides an improvement across the entire historic archive, not just new data, and as such provides an exciting opportunity for new science.

7. BiSON Open Data Portal

All data produced by the BiSON are freely available from the *BiSON Open Data Portal* – http://bison.ph.bham.ac.uk/opendata – and are also in the process of being deposited in the University of Birmingham Long Term Storage Archive (LTSA). The LTSA ensures the data will be available via a persistent URL for a minimum of ten years. Data will be available from the archive using the "FindIt@Bham" service – http://findit.bham.ac.uk – sand we also hope to provide all datasets with a Digital Object Identifier (DOI) as soon as this facility becomes available via the archive.

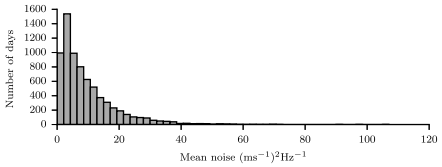

Figure 18 The distribution of mean noise per day for the overall time series. Lower is better.

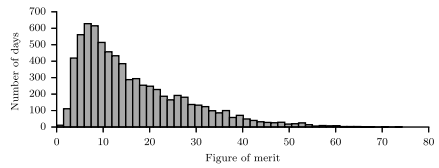

Figure 19 The distribution of FOM per day for the overall time series. Higher is better.

Data products are in the form of calibrated velocity residuals, concatenated into a single time series from all BiSON sites. Individual days of data, and also bespoke products produced from requested time periods and sites, are available by contacting the authors. We hope to also provide all raw data products via the LTSA as the archive is populated. Oscillation mode frequencies and amplitudes are available from Broomhall *et al.* (2009) and Davies *et al.* (2014b).

All data created specifically during the research for this article are openly available from the University of Birmingham ePapers data archive (Hale, 2015a) and are also listed on the BiSON Open Data Portal.

8. To the Future

It would be scientifically advantageous to increase the network duty cycle from 78 % to 100 % or better (*i.e.* to have every 40 s interval covered by one or more sites), and also to make better use of weighted averaging of multi-site data. We saw in Section 5 a seasonal variation in noise level caused by the changing Doppler offset throughout the year. A similar offset effect is seen on a daily period due to Earth's rotation, and this causes instruments at different longitudes to sample line formation at different heights in the solar atmosphere, meaning that the noise between sites is not completely coherent. Simultaneous observing at as many sites as possible allows the incoherent components of the noise to be beaten

down and potentially gives access to solar g-modes, which are expected to have very low frequencies and low amplitudes (Appourchaux *et al.*, 2010).

Technology has moved on significantly since the BiSON nodes were designed in the late 1980s. Whilst a considerable programme of upgrades has taken place over the years, the overall design is still limited by the original specification. Deploying more nodes in the classic style would be prohibitively expensive.

By taking advantage of modern fiber optics and electronic miniaturisation such as microcontrollers and single-board computers, it is possible to design a solar spectrometer with a much smaller physical footprint and considerably lower deployment cost, thus making it feasible to observe from many more sites. Work is underway on a second-generation network that will operate as a complement to the existing nodes, and this aims to guarantee better than 100 % duty cycle and much lower noise levels.

Disclosure of Potential Conflicts of Interest The authors declare that they have no conflicts of interest.

Acknowledgements We would like to thank all those who are, or have been, associated with BiSON. In Birmingham: George Isaak, Bill Brookes, Bob van der Raay, Clive McLeod, Roger New, Sarah Wheeler, Clive Speake, Brek Miller, Richard Lines, Phil Pavelin, Barry Jackson, Hugh Williams, Joe Litherland, Ian Barnes, Richard Bryan, and John Allison. In Mount Wilson: Ed Rhodes, Stephen Pinkerton, the team of USC undergraduate observing assistants, former USC staff members Maynard Clark, Perry Rose, Natasha Johnson, Steve Padilla, and Shawn Irish, and former UCLA staff members Larry Webster and John Boyden. In Las Campanas: Patricio Pinto, Andres Fuentevilla, Emilio Cerda, Frank Perez, Marc Hellebaut, Patricio Jones, Gastón Gutierrez, Juan Navarro, Francesco Di Mille, Roberto Bermudez, and the staff of LCO. In Izaña: We would like to give particular thanks to Pere Pallé and Teo Roca Cortés, and all staff at the IAC who have contributed to running the Mark I instrument over many years (see also the acknowledgements in Roca Cortés and Pallé, 2014). In Sutherland: Pieter Fourie, Willie Koorts, Jaci Cloete, Reginald Klein, John Stoffels, and the staff of SAAO. In Carnarvon: Les Bateman, Les Schultz, Sabrina Dowling-Giudici, Inge Lauw of Williams and Hughes Lawyers, and NBN Co. Ltd. In Narrabri: Mike Hill and the staff of CSIRO. BiSON is funded by the Science and Technology Facilities Council (STFC).

References

Appourchaux, T., Belkacem, K., Broomhall, A.-M., Chaplin, W.J., Gough, D.O., Houdek, G., Provost, J., Baudin, F., Boumier, P., Elsworth, Y., García, R.A., Andersen, B.N., Finsterle, W., Fröhlich, C., Gabriel, A., Grec, G., Jiménez, A., Kosovichev, A., Sekii, T., Toutain, T., Turck-Chièze, S.: 2010, The quest for the solar g modes. *Astron. Astrophys. Rev.* **18**, 197. DOI. ADS.

Barnes, I.: 2008, The installation of the Sutherland temperature controllers in 2007 August. *BiSON Technical Report Series 294*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2036/.

Barnes, I.: 2009, The Sutherland temperature controllers. *BiSON Technical Report Series 316*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2040/.

Barnes, I.: 2010, The Narrabri temperature controller. *BiSON Technical Report Series 333*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2047/.

Barnes, I., Hale, S.J.: 2005, Camarvon trip report – May 2005. *BiSON Technical Report Series 253*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2026/.

Barnes, I., Hale, S.J.: 2010, The installation of new temperature controllers at Narrabri in 2010 February. *BiSON Technical Report Series 332*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2046/.

Barnes, I., Jackson, B., Miller, B.A.: 2003, BeebSub #56A and BeebSub #56B – BBC computer substitutes. *BiSON Technical Report Series 218*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2014/.

Barnes, I., Jackson, B., Miller, B.A.: 2004, BeebSub #56C – An improved BBC computer substitute. *BiSON Technical Report Series 243*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2023/.

Barnes, I., Miller, B.A.: 2006, The grand opening of the Sutherland Zoo. *BiSON Technical Report Series 276*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2032/.

Barnes, I., Miller, B.A.: 2009, Jabba is returned to Carnarvon in 2009 July. *BiSON Technical Report Series 323*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2043/.

Barnes, I., Miller, B.A.: 2011, The mount controller: A digital autoguider for Las Campanas. *BiSON Technical Report Series 344*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2049/.

Barnes, I., Miller, B.A., Jackson, B.: 2007, The removal of Jabba from Carnarvon in 2006 November. *BiSON Technical Report Series 282*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2034/.

Brookes, J.R., Isaak, G.R., van der Raay, H.B.: 1976, Observation of free oscillations of the sun. *Nature* **259**, 92. DOI. ADS.

Brookes, J.R., Isaak, G.R., van der Raay, H.B.: 1978a, A resonant-scattering solar spectrometer. *Mon. Not. Roy. Astron. Soc.* **185**, 1. ADS.

Brookes, J.R., Isaak, G.R., van der Raay, H.B.: 1978b, The observation of a rotating body using high-resolution spectroscopy. *Mon. Not. Roy. Astron. Soc.* **185**, 19. ADS.

Broomhall, A.-M., Chaplin, W.J., Davies, G.R., Elsworth, Y., Fletcher, S.T., Hale, S.J., Miller, B., New, R.: 2009, Definitive Sun-as-a-star p-mode frequencies: 23 years of BiSON observations. *Mon. Not. Roy. Astron. Soc.* **396**, L100. DOI. ADS.

Chaplin, W.J., Elsworth, Y., Howe, R., Isaak, G.R., McLeod, C.P., Miller, B.A., van der Raay, H.B., Wheeler, S.J., New, R.: 1996, BiSON performance. *Solar Phys.* **168**, 1. DOI. ADS.

Chaplin, W.J., Elsworth, Y., Howe, R., Isaak, G.R., McLeod, C.P., Miller, B.A., New, R.: 1997, Techniques used in the analysis of data collected by the Birmingham Solar-Oscillations Network (BiSON). II. Frequency domain analysis & data merging. *Astron. Astrophys. Suppl.* **125**, 195. DOI. ADS.

Chaplin, W.J., Elsworth, Y., Isaak, G.R., Miller, B.A., New, R., Pintér, B.: 2004, Novel techniques for the identification of noise contributions to full-disc helioseismic power spectra. In: Danesy, D. (ed.) *SOHO 14 Helio- and Asteroseismology: Towards a Golden Future*, *ESA SP 559*, 360. ADS.

Chaplin, W.J., Elsworth, Y., Isaak, G.R., Miller, B.A., New, R., Pintér, B.: 2005, Noise characteristics of full-disc helioseismic observations made by resonant scattering spectrometers. *Mon. Not. Roy. Astron. Soc.* **359**, 607. DOI. ADS.

Chou, D.-Y., Sun, M.-T., Huang, T.-Y., Lai, S.-P., Chi, P.-J., Ou, K.-T., Wang, C.-C., Lu, J.-Y., Chu, A.-L., Niu, C.-S., Mu, T.-M., Chen, K.-R., Chou, Y.-P., Jimenez, A., Rabello-Soares, M.C., Chao, H., Ai, G., Wang, G.-P., Zirin, H., Marquette, W., Nenow, J.: 1995, Taiwan Oscillation Network. *Solar Phys.* **160**, 237. DOI. ADS.

Claverie, A., Isaak, G.R., McLeod, C.P., van der Raay, H.B., Cortes, T.R.: 1979, Solar structure from global studies of the 5-minute oscillation. *Nature* **282**, 591. DOI. ADS.

Davidson, C., Williams, H.K.: 2004, Progress report on the construction of the South African dome – 1989 November/December. *BiSON Technical Report Series 226*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2017/.

Davies, G.R., Chaplin, W.J., Elsworth, Y., Hale, S.J.: 2014a, BiSON data preparation: A correction for differential extinction and the weighted averaging of contemporaneous data. *Mon. Not. Roy. Astron. Soc.* **441**, 3009. DOI. ADS.

Davies, G.R., Broomhall, A.M., Chaplin, W.J., Elsworth, Y., Hale, S.J.: 2014b, Low-frequency, low-degree solar p-mode properties from 22 years of Birmingham Solar Oscillations Network data. *Mon. Not. Roy. Astron. Soc.* **439**, 2025. DOI. ADS.

Domingo, V., Fleck, B., Poland, A.I.: 1995, The SOHO mission: An overview. *Solar Phys.* **162**, 1. DOI. ADS.

Elsworth, Y.: 1992, Update on my report on the trip to Mount Wilson (Hale) Observatory. *BiSON Technical Report Series 5*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2003/.

Elsworth, Y., Howe, R., Isaak, G.R., McLeod, C.P., Miller, B.A., New, R., Wheeler, S.J.: 1995, Techniques used in the analysis of solar oscillations data from the BiSON (University of Birmingham) network. I. Daily calibration. *Astron. Astrophys. Suppl.* **113**, 379. ADS.

Fossat, E.: 1991, The IRIS network for full disk helioseismology – Present status of the programme. *Solar Phys.* **133**, 1. DOI. ADS.

Fossat, E., IRIS Group: 2002, 11 years of IRIS network exploitation. In: Combes, F., Barret, D. (eds.) *SF2A-2002: Semaine de l'Astrophysique Francaise*, EDP Sciences, Les Ulis, 521. ADS.

Fröhlich, C., Romero, J., Roth, H., Wehrli, C., Andersen, B.N., Appourchaux, T., Domingo, V., Telljohann, U., Berthomieu, G., Delache, P., Provost, J., Toutain, T., Crommelynck, D.A., Chevalier, A., Fichot, A., Däppen, W., Gough, D., Hoeksema, T., Jiménez, A., Gómez, M.F., Herreros, J.M., Cortés, T.R., Jones, A.R., Pap, J.M., Willson, R.C.: 1995, VIRGO: Experiment for helioseismology and solar irradiance monitoring. *Solar Phys.* **162**, 101. DOI. ADS.

Gabriel, A.H., Grec, G., Charra, J., Robillot, J.-M., Roca Cortés, T., Turck-Chièze, S., Bocchia, R., Boumier, P., Cantin, M., Cespédes, E., Cougrand, B., Crétolle, J., Damé, L., Decaudin, M., Delache, P., Denis, N., Duc, R., Dzitko, H., Fossat, E., Fourmond, J.-J., García, R.A., Gough, D., Grivel, C., Herreros, J.M., Lagardère, H., Moalic, J.-P., Pallé, P.L., Pétrou, N., Sanchez, M., Ulrich, R., van der Raay, H.B.: 1995, Global Oscillations at Low Frequency from the SOHO mission (GOLF). *Solar Phys.* **162**, 61. DOI. ADS.

Grec, G., Fossat, E., Pomerantz, M.: 1980, Solar oscillations – Full disk observations from the geographic South Pole. *Nature* **288**, 541. DOI. ADS.

Hale, S.J.: 2003, Scientific advancements in analysis of solar oscillation data. Master's thesis, School of Physics and Astronomy, University of Birmingham, UK. http://etheses.bham.ac.uk/5952/.

Hale, S.J.: 2012, Sutherland 2012 June installation of tiger counters. *BiSON Technical Report Series 358*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2053/.

Hale, S.J.: 2013, The installation of a digital autoguider in Sutherland in 2013 November. *BiSON Technical Report Series 362*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2057/.

Hale, S.J.: 2014a, Autoguider repairs at Mount Wilson in 2014 April. *BiSON Technical Report Series 365*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2060/.

Hale, S.J.: 2014b, Blind and mount controller repairs in Las Campanas in 2014 June. *BiSON Technical Report Series 366*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2061/.

Hale, S.J.: 2015a, BiSON – All sites – 1995 to 2014 – Performance check. Birmingham Solar Oscillations Network, University of Birmingham, UK. http://epapers.bham.ac.uk/1977/.

Hale, S.J.: 2015b, The Hannibal temperature controllers. *BiSON Technical Report Series 372*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2067/.

Hale, S.J.: 2015c, New temperature controller for Hannibal in Las Campanas in 2015 April. *BiSON Technical Report Series 373*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2068/.

Hale, S.J., Davies, G.R.: 2013, The installation of a digital autoguider and tiger counters in Narrabri in 2013 April. *BiSON Technical Report Series 360*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2055/.

Hale, S.J., Miller, B.A.: 2006, The grand opening of the Las Campanas Zoo. *BiSON Technical Report Series 261*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2029/.

Harvey, J.W., Hill, F., Hubbard, R.P., Kennedy, J.R., Leibacher, J.W., Pintar, J.A., Gilman, P.A., Noyes, R.W., Title, A.M., Toomre, J., Ulrich, R.K., Bhatnagar, A., Kennewell, J.A., Marquette, W., Patron, J., Saa, O., Yasukawa, E.: 1996, The Global Oscillation Network Group (GONG) project. *Science* **272**, 1284. DOI. ADS.

Hoyng, P.: 1989, On the sensitivity of resonant scattering spectrometers for whole-disk solar velocity oscillation measurements. *Astrophys. J.* **345**, 1088. DOI. ADS.

Isaak, G.R.: 1961, An atomic beam spectrophotometer. *Nature* **189**, 373. DOI. ADS.

Jackson, B., Miller, B.A.: 2004, The grand opening of the Narrabri Zoo in 2004 July. *BiSON Technical Report Series 241*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2022/.

Leibacher, J.W., Stein, R.F.: 1971, A new description of the solar five-minute oscillation. *Astrophys. Lett.* **7**, 191. ADS.

Leighton, R.B., Noyes, R.W., Simon, G.W.: 1962, Velocity fields in the solar atmosphere. I. Preliminary report. *Astrophys. J.* **135**, 474. DOI. ADS.

Lines, R.: 1998, A visit to Las Campanas after the lightning strike that destroyed our step-down transformer. *BiSON Technical Report Series 71*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2005/.

McLeod, C.P.: 2002, Mark I scaler system. *BiSON Technical Report Series 184*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2009/.

Miller, B.A.: 1997, The trip to Las Campanas during the big snowstorm of 1997 August. *BiSON Technical Report Series 62*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2004/.

Miller, B.A.: 1998, A visit to Las Campanas in 1998 January. *BiSON Technical Report Series 74*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2006/.

Miller, B.A.: 2000, The replacement of the blind motor and the connection of the station to the Internet in Narrabri in 2000 March. *BiSON Technical Report Series 138*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2008/.

Miller, B.A.: 2002, How do you get to the zoo? *BiSON Technical Report Series 187*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2010/.

Miller, B.A.: 2003, The grand opening of the Camarvon Zoo in 2002 November. *BiSON Technical Report Series 193*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK.

Miller, B.A.: 2005, The grand opening of the Mount Wilson Zoo. *BiSON Technical Report Series 255*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2027/.

Miller, B.A.: 2011, The installation of a digital autoguider in Las Campanas in 2011 March. *BiSON Technical Report Series 343*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2048/.

Miller, B.A., New, R.: 1999, The installation of Klaus at Mount Wilson in 1996 June. *BiSON Technical Report Series 106*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2007/.

New, R., Hale, S.J.: 2006, Camarvon trip report – July/August 2005. *BiSON Technical Report Series 260*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2028/.

New, R., Isaak, G.R.: 2003, Work carried out at Narrabri in 2003 July. *BiSON Technical Report Series 213*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2013/.

Roca Cortés, T., Pallé, P.L.: 2014, The Mark-I helioseismic experiment—I. Measurements of the solar gravitational redshift (1976–2013). *Mon. Not. Roy. Astron. Soc.* **443**, 1837. DOI. ADS.

Salabert, D., Fossat, E., Gelly, B., Tomczyk, S., Pallé, P., Jiménez-Reyes, S.J., Cacciani, A., Corbard, T., Ehgamberdiev, S., Grec, G., Hoeksema, J.T., Kholikov, S., Lazrek, M., Schmider, F.X.: 2002a, IRIS^{++} database: Merging of IRIS + Mark-1 + LOWL. *Astron. Astrophys.* **390**, 717. DOI. ADS.

Salabert, D., Jiménez-Reyes, S.J., Fossat, E., Gelly, B., Schmider, F.X.: 2002b, Variability of p-mode parameters in 11 years of IRIS^{++} data. In: Wilson, A. (ed.) *From Solar Min to Max: Half a Solar Cycle with SOHO*, *ESA SP 508*, 95. ADS.

Scherrer, P.H., Bogart, R.S., Bush, R.I., Hoeksema, J.T., Kosovichev, A.G., Schou, J., Rosenberg, W., Springer, L., Tarbell, T.D., Title, A., Wolfson, C.J., Zayer, I., MDI Engineering Team: 1995, The solar oscillations investigation – Michelson Doppler Imager. *Solar Phys.* **162**, 129. DOI. ADS.

Schou, J., Scherrer, P.H., Bush, R.I., Wachter, R., Couvidat, S., Rabello-Soares, M.C., Bogart, R.S., Hoeksema, J.T., Liu, Y., Duvall, T.L., Akin, D.J., Allard, B.A., Miles, J.W., Rairden, R., Shine, R.A., Tarbell, T.D., Title, A.M., Wolfson, C.J., Elmore, D.F., Norton, A.A., Tomczyk, S.: 2012, Design and ground calibration of the Helioseismic and Magnetic Imager (HMI) instrument on the Solar Dynamics Observatory (SDO). *Solar Phys.* **275**, 229. DOI. ADS.

Stallman, R.M.: 1983, New Unix implementation. https://www.gnu.org/gnu/initial-announcement.html.

Tomczyk, S., Schou, J., Thompson, M.J.: 1995, Measurement of the rotation rate in the deep solar interior. *Astrophys. J. Lett.* **448**, L57. DOI. ADS.

Torvalds, L.B.: 1991, Free Minix-like kernel sources for 386-AT. https://groups.google.com/d/msg/comp.os.minix/4995SivOI9o/GwqLJlPSICEJ.

Ulrich, R.K.: 1970, The five-minute oscillations on the solar surface. *Astrophys. J.* **162**, 993. DOI. ADS.

Williams, H.K.: 2004, The construction of the Narrabri dome – 1992 August. *BiSON Technical Report Series 227*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. http://epapers.bham.ac.uk/2018/.

Willson, R.C.: 1979, Active cavity radiometer type IV. *Appl. Opt.* **18**, 179. DOI. ADS.

Appendix E

Refereed Publication

Chapter 5 includes text and figures from the publication "Atmospheric Extinction Coefficients in the I_c Band for Several Major International Observatories: Results from the BiSON Telescopes, 1984-2016" (Hale et al., 2017) in the journal *Astronomical Journal*. The published article was first released on 2017 July 20, and is included here for reference. Data from the article are available from the University of Birmingham eData archive (Hale, 2017a). Whilst this publication is a result of collaborative working, the content is entirely the author's own work. Co-authors provided guidance on data processing and manuscript review.

The Astronomical Journal, 154:89 (11pp), 2017 September https://doi.org/10.3847/1538-3881/aa81d0

Atmospheric Extinction Coefficients in the I_c Band for Several Major International Observatories: Results from the BiSON Telescopes, 1984–2016

S. J. Hale[1,2], W. J. Chaplin[1,2], G. R. Davies[1,2], Y. P. Elsworth[1,2], R. Howe[1,2], M. N. Lund[1,2], E. Z. Moxon[1,2], A. Thomas[1,2], P. L. Pallé[3], and E. J. Rhodes, Jr.[4,5]
[1] School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT, UK; s.j.hale@bham.ac.uk
[2] Stellar Astrophysics Centre, Department of Physics and Astronomy, Aarhus University, Ny Munkegade 120, DK-8000 Aarhus C, Denmark
[3] Instituto de Astrofísica de Canarias, and Department of Astrophysics, Universidad de La Laguna, San Cristóbal de La Laguna, Tenerife, Spain
[4] Department of Physics and Astronomy, University of Southern California, Los Angeles, CA 90089, USA
[5] Astrophysics and Space Sciences Section, Jet Propulsion Laboratory California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109-8099, USA
Received 2017 May 9; revised 2017 July 3; accepted 2017 July 20; published 2017 August 11

Abstract

Over 30 years of solar data have been acquired by the Birmingham Solar Oscillations Network (BiSON), an international network of telescopes used to study oscillations of the Sun. Five of the six BiSON telescopes are located at major observatories. The observational sites are, in order of increasing longitude: Mount Wilson (Hale) Observatory (MWO), California, USA; Las Campanas Observatory, Chile; Observatorio del Teide, Izaña, Tenerife, Canary Islands; the South African Astronomical Observatory, Sutherland, South Africa; Carnarvon, Western Australia; and the Paul Wild Observatory, Narrabri, New South Wales, Australia. The BiSON data may be used to measure atmospheric extinction coefficients in the I_c band (approximately 700–900 nm), and presented here are the derived atmospheric extinction coefficients from each site over the years 1984–2016.

Key words: atmospheric effects – Sun: helioseismology – Sun: oscillations

1. Introduction

The Birmingham Solar Oscillations Network (BiSON) is a six-site ground-based network of solar observatories. The primary science output of the network is detection of solar oscillations. Here, we take an alternative window into these data and assess the historic atmospheric column extinction coefficients at each of our international network sites, over the life of the network. The history and performance of the network is detailed in Hale et al. (2016). In summary, the first instrument was commissioned at Observatorio del Teide in Izaña, Tenerife, in 1975, with the additional five nodes coming online throughout the mid-80s and early-90s. The observational sites are, in order of increasing longitude: Mount Wilson (Hale) Observatory (MWO), California, USA; Las Campanas Observatory (LCO), Chile; Observatorio del Teide, Izaña, Tenerife, Canary Islands; South African Astronomical Observatory (SAAO), Sutherland, South Africa; Carnarvon, Western Australia; and Paul Wild Observatory, Narrabri, New South Wales, Australia. The network operates continuously and provides an annual data duty cycle averaging around 82%. The locations of the network nodes are summarized in Table 1.

In the next section, we will take a brief look at the network instrumentation. In Section 3, we describe how the atmospheric-extinction coefficients are determined. We will then go on to present the historic extinction coefficients of each site in Section 4.

2. Instrumentation

The BiSON solar spectrometers provide very precise measures of the disc-averaged line-of-sight velocity of the solar surface. This is done by comparing the wavelength of the potassium absorption line at 769.898 nm formed within the Sun, with the same line in a vapor reference cell on Earth. The spectrometers typically have three photo-detectors. Two of the detectors measure the intensity of the light scattered from the vapor cell, and the third measures the intensity of light transmitted directly through the instrument. The light is pre-filtered using an I_c band filter (approximately 700–900 nm) formed from Schott RG9 and KG4 glass, and the bandwidth is then reduced again to 15 Å using an interference filter centered on 769.9 nm. The width of the potassium absorption line is significantly narrower than 15 Å, and so the measurement of the transmitted light can be considered to be a measurement of the direct-Sun radiance near the center of the I_c band—essentially the instrument becomes an automated solar photometer. Where a measurement of the transmitted light is not available, the light scattered from the vapor cell can also be used as a proxy for the transmitted intensity—see the Appendix for further details. Figure 1 shows a typical day of data captured from the site at Sutherland, along with the variation in airmass during the day. The data have been pre-processed to remove any periods of instrumental failure and cloudy conditions. Even thin cirrus produces an easily identifiable reduction in data quality, and so the data analyzed in this paper are from clear sky observations only.

3. Deriving Extinction Coefficients

Atmospheric extinction has three main components: Rayleigh scattering, scattering due to aerosols, and molecular absorption. The strongest absorption effects are due to molecular oxygen and ozone, which both absorb in the ultraviolet, and water vapor which absorbs in the infrared. At the BiSON observational wavelength of 769.9 nm, Rayleigh scattering is at a level of a few percent, and there is no molecular absorption: the observed atmospheric extinction is dominated by the contribution of aerosols.

Table 1
Coordinates of the Six-station Birmingham Solar Oscillations Network (BiSON)

Location	Longitude (deg E)	Latitude (deg N)	Altitude (m)	Commissioned (year)
Mount Wilson, California, USA	−118.08	+34.13	1742	1992
Las Campanas, Chile	−70.70	−29.02	2282	1991
Izaña, Tenerife, Canary Islands	−16.50	+28.30	2368	1975
Sutherland, South Africa	+20.82	−32.38	1771	1990
Carnarvon, Western Australia	+113.75	−24.85	10	1985
Narrabri, New South Wales, Australia	+149.57	−30.32	217	1992

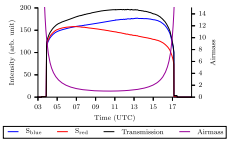

Figure 1. Data from 2014 December 1 at Sutherland, South Africa. The black line shows the transmitted solar intensity. The blue and red lines show the measured intensity at points on the shorter-wavelength (S_{blue}) and longer-wavelength (S_{red}) "wings" of the potassium absorption line at 769.898 nm, respectively. The variation in airmass is shown by the magenta line, measured against the right-axis.

The Beer–Lambert law states that the transmittance, T, of a material is related to its optical depth, τ_λ, by

$$T = \frac{I}{I_0} = e^{-\tau_\lambda A}, \quad (1)$$

where I_0 is the solar extraterrestrial radiance (i.e., at zero airmass), and I is the direct-Sun radiance. In this case, τ_λ is the column atmospheric aerosol optical depth (AOD) per unit airmass, and A is the relative optical airmass as a function of solar zenith angle. The AOD is typically quoted as unitless when considering only the unit airmass at the zenith. By taking the natural logarithm of both sides we obtain,

$$\ln(I/I_0) = -\tau_\lambda A, \quad (2)$$

which gives a convenient linear relationship where the gradient of the relationship is a measure of τ_λ, and I_0 is now simply a normalization factor taken as the maximum intensity measured on a given day. For astronomical use, we rescale AOD in terms of magnitude,

$$\begin{aligned} \kappa_\lambda &= -2.5\log_{10}(e)(-\tau_\lambda), \\ &= 1.086\tau_\lambda, \end{aligned} \quad (3)$$

where κ_λ is the atmospheric extinction coefficient, with units of magnitudes per airmass. More accurately, this is the column extinction coefficient, as we do not include any knowledge on the vertical structure of the atmosphere. Readers in the climate modeling and aerosol communities should divide the values presented here by 1.086 in order to recover the total column-aerosol in terms of AOD.

In this analysis, the known zenith angle was used to calculate the airmass based on Kasten & Young (1989) who define the airmass as,

$$A = \frac{1}{\cos z + 0.50572(6.07995 + 90 - z)^{-1.6364}}, \quad (4)$$

where the zenith angle z is in degrees. This model gives an airmass of approximately 38 at the horizon, producing good results for the whole range of zenith angles.

The extinction coefficients were determined using two methods: first, by making a standard linear least-squares fit of the magnitude-like value from Equation (2) against airmass, and second, by calculating the non-overlapping independent first-differences and then obtaining statistical estimations from the histogram of a timeseries of,

$$\frac{d\,m}{d\,A} = \frac{m_i - m_{i-1}}{A_i - A_{i-1}}, \quad (5)$$

where m is the magnitude-like value from Equation (2), A is airmass, and i is the sample index incremented in steps of two. In our fits, we consider only airmasses in the range of 2–5, corresponding to zenith angles between approximately 60°–80°, as this is the region where airmass is changing most linearly. Below two airmasses, the change does not follow a strictly linear relationship and is not well described by a straight line fit. The rate-of-change is also too low to allow for good fitting. Additionally, we need to ensure we remove the seasonal variations in minimum airmass due to the changing maximum altitude of the Sun throughout the year, as this could introduce an artificial seasonal effect in the derived extinction values. Above five airmasses, the rate-of change is too high, producing differential extinction across the extended source of the Sun (Davies et al. 2014). The pre-meridian and post-meridian values (hereafter referred to as "morning" and "afternoon") are fitted separately, as it is expected that these will differ due to local environmental considerations. The results from both techniques for the same day as in Figure 1 are shown in Figure 2.

The coefficient estimation technique and the selection of airmass range-limits affects the value of the determined extinction. In order to investigate the robustness of our parameters, a randomization trial was performed where, rather than fixing the lower and upper airmass limits at 2 and 5, respectively, they were randomly selected each day between 2–3 and 4–5 airmasses. Five realizations were then generated for the full timeseries of fitted extinction gradients from each site, and the absolute difference was compared between each realization and the gradients measured when the airmass limits were fixed. A similar comparison was made between the

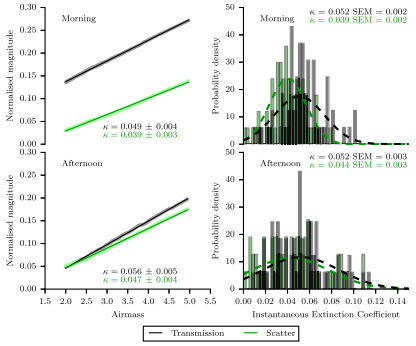

Figure 2. Morning and afternoon extinction coefficients for 2014 December 1 at Sutherland, South Africa. The values determined are shown in the plot. Left column: the formal uncertainty on the fit is calculated from the covariance matrix returned by the fitting function. Right column: the uncertainty on the histogram is the standard error of the median (SEM), i.e., 0.741 times the interquartile range of the distribution divided by the square-root of the number of points. The dashed lines on the histograms represent, for comparison, the equivalent Gaussian profile for the measured mean and standard deviation.

timeseries of fitted gradients and a timeseries of median first-differences. In all cases, the mean difference was less than 4% of the mean extinction. The standard deviation of the difference was 3–6 times lower than the measured extinction standard deviation. Any systematic offsets or increase in scattering due to the processing techniques are at a level significantly lower than that due to real physical effects. These two techniques are considered to be equivalent, and the results presented here are produced solely from the first method using linear least-squares fitting.

In the next section, we present fitted extinction coefficients for all of the historic data from each BiSON site. For clarity, the units of extinction will no longer be stated on each value in the text. All extinction values are specified in magnitudes per airmass. During the discussion for each site, we quote either mean, mode, or median values depending on the values required for comparisons with other studies. For consistency, a summary of the coefficients is given in Table 2. Because there are significant seasonal differences, we also present the values measured over two months for each mid-summer (July–August in the northern hemisphere, January–February in the Southern hemisphere) in Table 3, and during mid-winter (January–February in the Northern hemisphere, July–August in the southern hemisphere) in Table 4.

4. Sites

4.1. Izaña, Tenerife

The BiSON node at Tenerife (Roca Cortés & Pallé 2014) is based at the Observatorio del Teide, which is operated by the IAC (*Instituto de Astrofísica de Canarias*). The Canary Islands are located about 100 km to the west of the North African coast. The islands are close to the Western Sahara and so, during the summer months, they frequently experience high concentrations of mineral dust in the atmosphere. This aerosol concentration is easily seen in Figure 3 as the strong seasonal variation in extinction.

The modal values of the extinction distributions are 0.054 in the morning and 0.046 in the afternoon for transmission. The standard deviation on these values is 0.09, the high value being indicative of the large scatter in values between summer and winter periods. The difference in the values pre-meridian and post-meridian is likely due to the surrounding geography, where morning extinction effects are through the atmosphere over North Africa and afternoon is over the potentially clearer North Atlantic ocean. If only the winter months are considered, then the modal value for both morning and afternoon drops to 0.045 with correspondingly reduced scatter. Mineral dust events are typical between June and October, where extinction values anywhere from 0.1 to 0.8 may be experienced. The

Table 2
Extinction Coefficients from All Sites

Location	Detector	Mode	Median	Mean	Sigma
Mount Wilson	Transmission (Morning)	0.048	0.085	0.096	0.055
Mount Wilson	Transmission (Afternoon)	0.070	0.081	0.093	0.057
Mount Wilson	Scatter (Morning)	0.056	0.064	0.074	0.043
Mount Wilson	Scatter (Afternoon)	0.061	0.070	0.082	0.053
Las Campanas	Transmission (Morning)	…	…	…	…
Las Campanas	Transmission (Afternoon)	…	…	…	…
Las Campanas	Scatter (Morning)	0.029	0.033	0.037	0.018
Las Campanas	Scatter (Afternoon)	0.028	0.033	0.036	0.014
Izaña	Transmission (Morning)	0.054	0.058	0.092	0.090
Izaña	Transmission (Afternoon)	0.046	0.060	0.092	0.091
Izaña	Scatter (Morning)	0.048	0.051	0.082	0.084
Izaña	Scatter (Afternoon)	0.039	0.049	0.079	0.087
Sutherland	Transmission (Morning)	0.038	0.046	0.054	0.040
Sutherland	Transmission (Afternoon)	0.039	0.045	0.051	0.026
Sutherland	Scatter (Morning)	0.034	0.039	0.048	0.039
Sutherland	Scatter (Afternoon)	0.037	0.038	0.044	0.027
Carnarvon	Transmission (Morning)	0.052	0.072	0.085	0.050
Carnarvon	Transmission (Afternoon)	0.052	0.078	0.088	0.043
Carnarvon	Scatter (Morning)	0.057	0.074	0.087	0.049
Carnarvon	Scatter (Afternoon)	0.052	0.071	0.081	0.043
Narrabri	Transmission (Morning)	0.053	0.066	0.070	0.029
Narrabri	Transmission (Afternoon)	0.061	0.068	0.073	0.025
Narrabri	Scatter (Morning)	0.048	0.055	0.062	0.028
Narrabri	Scatter (Afternoon)	0.050	0.053	0.059	0.027

Note. Only 1995 onwards has been considered in order to remove any exceptional atmospheric events. The data sets are sufficiently large that the standard error of the mean (i.e., $\sigma/\sqrt{N}$) is very small and therefore not presented.

Table 3
Summer Extinction Coefficients from All Sites

Location	Detector	Mode	Median	Mean	Sigma
Mount Wilson	Transmission (Morning)	0.053	0.077	0.084	0.041
Mount Wilson	Transmission (Afternoon)	0.050	0.082	0.095	0.062
Mount Wilson	Scatter (Morning)	0.053	0.057	0.065	0.031
Mount Wilson	Scatter (Afternoon)	0.037	0.076	0.088	0.056
Las Campanas	Transmission (Morning)	…	…	…	…
Las Campanas	Transmission (Afternoon)	…	…	…	…
Las Campanas	Scatter (Morning)	0.032	0.042	0.046	0.017
Las Campanas	Scatter (Afternoon)	0.040	0.041	0.044	0.016
Izaña	Transmission (Morning)	0.057	0.073	0.151	0.133
Izaña	Transmission (Afternoon)	0.057	0.075	0.147	0.136
Izaña	Scatter (Morning)	0.048	0.070	0.142	0.130
Izaña	Scatter (Afternoon)	0.047	0.070	0.141	0.135
Sutherland	Transmission (Morning)	0.043	0.046	0.050	0.022
Sutherland	Transmission (Afternoon)	0.044	0.046	0.048	0.013
Sutherland	Scatter (Morning)	0.036	0.040	0.044	0.021
Sutherland	Scatter (Afternoon)	0.037	0.040	0.042	0.014
Carnarvon	Transmission (Morning)	0.085	0.100	0.115	0.078
Carnarvon	Transmission (Afternoon)	0.095	0.107	0.114	0.045
Carnarvon	Scatter (Morning)	0.119	0.106	0.119	0.068
Carnarvon	Scatter (Afternoon)	0.081	0.103	0.110	0.046
Narrabri	Transmission (Morning)	0.076	0.074	0.081	0.030
Narrabri	Transmission (Afternoon)	0.073	0.077	0.086	0.041
Narrabri	Scatter (Morning)	0.060	0.068	0.074	0.026
Narrabri	Scatter (Afternoon)	0.060	0.061	0.068	0.035

Note. In the Northern hemisphere, the measured summer months each year are the beginning of July to the end of August, and in the Southern hemisphere, the beginning of January until the end of the February. Only 1995 onwards has been considered in order to remove any exceptional atmospheric events. The data sets are sufficiently large that the standard error of the mean (i.e., $\sigma/\sqrt{N}$) is very small and therefore not presented.

extinction values derived from the scattered light are a slight underestimate as expected (see Appendix), but otherwise show the same trends. The distribution of extinction coefficients appears to show two combined trends. The first, is a set of normally distributed values centered on approximately 0.05, and a long positive tail that corresponds to the periods of

Table 4
Winter Extinction Coefficients from All Sites

Location	Detector	Mode	Median	Mean	Sigma
Mount Wilson	Transmission (Morning)	0.048	0.077	0.086	0.051
Mount Wilson	Transmission (Afternoon)	0.075	0.070	0.074	0.032
Mount Wilson	Scatter (Morning)	0.046	0.060	0.070	0.035
Mount Wilson	Scatter (Afternoon)	0.053	0.056	0.059	0.027
Las Campanas	Transmission (Morning)	...	...	...	...
Las Campanas	Transmission (Afternoon)	...	...	...	...
Las Campanas	Scatter (Morning)	0.028	0.028	0.031	0.015
Las Campanas	Scatter (Afternoon)	0.027	0.028	0.030	0.010
Izaña	Transmission (Morning)	0.045	0.048	0.059	0.035
Izaña	Transmission (Afternoon)	0.045	0.048	0.062	0.062
Izaña	Scatter (Morning)	0.042	0.043	0.051	0.030
Izaña	Scatter (Afternoon)	0.038	0.038	0.047	0.045
Sutherland	Transmission (Morning)	0.038	0.049	0.061	0.042
Sutherland	Transmission (Afternoon)	0.039	0.044	0.056	0.046
Sutherland	Scatter (Morning)	0.033	0.041	0.057	0.056
Sutherland	Scatter (Afternoon)	0.029	0.037	0.049	0.045
Carnarvon	Transmission (Morning)	0.051	0.056	0.065	0.028
Carnarvon	Transmission (Afternoon)	0.056	0.060	0.065	0.026
Carnarvon	Scatter (Morning)	0.051	0.060	0.068	0.032
Carnarvon	Scatter (Afternoon)	0.052	0.053	0.059	0.026
Narrabri	Transmission (Morning)	0.057	0.058	0.061	0.024
Narrabri	Transmission (Afternoon)	0.064	0.061	0.064	0.016
Narrabri	Scatter (Morning)	0.040	0.046	0.050	0.017
Narrabri	Scatter (Afternoon)	0.044	0.046	0.049	0.016

Note. In the Northern hemisphere, the measured winter months each year are the beginning of January until the end of the February, and in the Southern hemisphere, the beginning of July to the end of August. Only 1995 onwards has been considered in order to remove any exceptional atmospheric events. The data sets are sufficiently large that the standard error of the mean (i.e., $\sigma/\sqrt{N}$) is very small and therefore not presented.

mineral dust events. If the summer and winter periods are analyzed separately, then the winter does indeed show a mean of approximately 0.05–0.06, and the summer a much higher mean of 0.14–0.15 with correspondingly higher standard deviation. A similar atmospheric analysis on data from telescopes on the Canary Archipelago, which included our BiSON data from this instrument, was made by Laken et al. (2016), and the modal values are consistent given the uncertainties. Laken et al. (2016) provides a thorough investigation into the occurrence of dust events, and how they change over both short seasonal periods and longer timescales. Several authors have investigated this in more detail see, e.g., Guerrero et al. (1998), Jimenez et al. (1998), Siher et al. (2004), García-Gil et al. (2010), and Laken et al. (2014).

Siher et al. (2002) present extinction values for the International Research of the Interior of the Sun (IRIS) site based at Izaña. IRIS was a similar network to BiSON and used a similar observational technique but made use of the shorter-wavelength sodium absorption line at 589.6 nm. They quote an average extinction value of 0.111, which is slightly higher than the mean values found here. The higher value is expected due to their use of the shorter wavelength, as shorter wavelengths tend to suffer greater extinction. Jimenez et al. (1998) found the extinction to vary during 1984–1989 between 0.04 and 0.07 at 680 nm, which is in agreement with the values found here for Izaña. During dust storms, Jimenez et al. (1998) reported values up to 0.8, which again is in agreement with our findings. The Global Oscillation Network Group (GONG) site-survey (Hill et al. 1994) at Izaña measured an average extinction value of 0.1169 from 1985 September to 1993 July. GONG is another network similar to BiSON. Their initial site-survey used a normal incidence pyrheliometer manufactured by Eppley Laboratories (Fischer et al. 1986), which has a wide spectral sensitivity range of 250–3000 nm. Light from the Sun is broadly like that of a black body at a temperature of 6000 K, meaning that the intensity peaks at a wavelength of approximately 500 nm and decays quickly in the infrared. The value measured by the pyrheliometer will be strongly weighted toward the peak wavelength of the solar spectrum. King (1985) discusses the wavelength dependence of typical expected atmospheric extinction from 300 to 1100 nm at the Roque de los Muchachos Observatory, on the adjacent island of La Palma. At 500 nm, an extinction coefficient of 0.1244 can be expected, very close to the value determined during the GONG site-survey. At 300 nm, near the short end of the pyrheliometer sensitivity range, values well over 3 can be expected. The average extinction measured by the pyrheliometer will be much higher than the monochromatic values measured by BiSON at 769.9 nm, and unfortunately, this means that no comparison can be made between the results from BiSON and the GONG site-survey.

Probably the most striking feature in Figure 3 is the increase in extinction following the eruption of Mount Pinatubo in the Philippines on 1991 June 15. There is also a hint of the tail-end of effects from the El Chichón eruption in Mexico in 1982 April, where the start of the data in 1984 show extinction values around 0.1, double the typical value expected outside of a dust intrusion. Both Guerrero et al. (1998) and García-Gil et al. (2010) have previously observed these features from telescopes based at the Canary Islands.

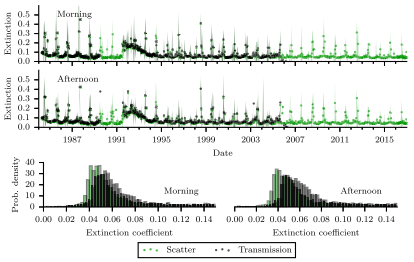

Figure 3. Extinction coefficients and statistical distribution from Izaña, Tenerife. Each dot represents the median value over 14 days. The colored banding represents ±3 times the standard error on each median. Dates indicate January 1 for each year.

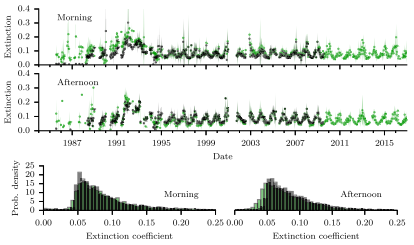

Figure 4. Extinction coefficients and statistical distribution from Carnarvon, Western Australia. Each dot represents the median value over 14 days. The colored banding represents ±3 times the standard error on each median. Dates indicate January 1 for each year.

4.2. Carnarvon

The BiSON node at Carnarvon is based at the historic Overseas Telecommunications Commission Satellite Earth Station, around 900 km north of Perth in Western Australia. The measured extinction coefficients over the operational lifetime of the site are shown in Figure 4.

The median morning extinction is 0.072, and the afternoon is 0.078. This increases to approximately 0.1 in the summer and decreases to around 0.06 in the winter. The standard deviation shows a similar increase in summer compared to winter. In the summer, the standard deviation is noticeably lower in the afternoon compared to the morning, at 0.05

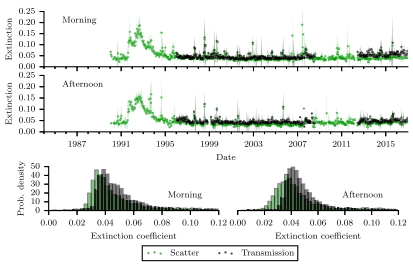

Figure 5. Extinction coefficients and statistical distribution from Sutherland, South Africa. Each dot represents the median value over 14 days. The colored banding represents ±3 times the standard error on each median. Dates indicate January 1 for each year.

and 0.07, respectively. Morning data are collected over the plains of Western Australia, and afternoon data are collected through air over the Indian ocean, and so it is expected that the sandy environment of Carnarvon would have a greater impact on extinction during morning observations. There are no comparison sites available near Carnarvon.

4.3. Sutherland

The BiSON node at Sutherland is situated 360 km north–east of Cape Town, at the SAAO. The measured extinction coefficients over the operational lifetime of the site are shown in Figure 5.

The median morning extinction is 0.046, and the afternoon is 0.045. There are no significant differences in the environment around Sutherland, and this is reflected by the stability of the extinction coefficients between morning and afternoon. The site shows stable performance with little change in the median values throughout the year. The atmosphere is particularly stable, with standard deviations of between 0.04 and 0.06 in the winter, and approximately 0.02 during the summer. Kilkenny (1995) state a mean extinction coefficient of 0.07 in the I_c band, which is slightly higher than is found here, but is within our measured standard deviation.

4.4. Las Campanas

The BiSON node at Las Campanas is situated 630 km north of Santiago, at the Las Campanas Observatory. The measured extinction coefficients over the operational lifetime of the site are shown in Figure 6.

Transmission data from this site are unreliable due to instrumentation issues, and so only the scattered-light data have been used to derive the extinction coefficients. The median extinction is 0.033 for both morning and afternoon, with a standard deviation of less than 0.02. There is also little variation between winter and summer periods, with the mean extinction at 0.028 and 0.042, respectively. Morning data are collected through air masses over South America, and afternoon data are collected through air over the South Pacific ocean. There appears to be no significant difference between the two zones, and Las Campanas has the most stable atmosphere of all BiSON sites, showing the lowest standard deviation.

Siher et al. (2002) performed a similar atmospheric analysis for the IRIS site based at La Silla, the adjacent mountain ridge to Las Campanas. They quote a value of 0.097, which is higher than is found here, but we again have to consider the shorter wavelength used by IRIS, and also that we expect the extinction from the scattered light to be a slight underestimate of the equivalent extinction from the transmitted light. The measured standard deviation is similar at 0.028.

4.5. Narrabri

The BiSON node at Narrabri is situated 525 km north–west of Sydney, at the Paul Wild Observatory. The measured extinction coefficients over the operational lifetime of the site are shown in Figure 7.

The median morning extinction is 0.074, and afternoon is 0.077. The standard deviation is approximately 0.04. There are no significant differences in the environment around Narrabri and again this is reflected in the stability of the extinction coefficients between morning and afternoon. There is, however, a strong seasonal variation with median values dropping to 0.06 in the winter and rising to almost 0.08 in the summer. Narrabri tends to have lower extinction and lower standard deviation than the coastal town of Carnarvon, although it does suffer from a higher percentage of cloudy days overall. No data sets could be found for comparison with Narrabri.

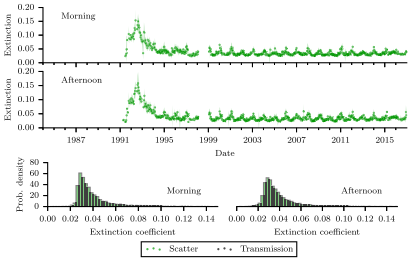

Figure 6. Extinction coefficients and statistical distribution from Las Campanas, Chile. Each dot represents the median value over 14 days. The colored banding represents ±3 times the standard error on each median. Dates indicate January 1 for each year. There is no transmitted light measured at Las Campanas, and so only the scattering values are presented.

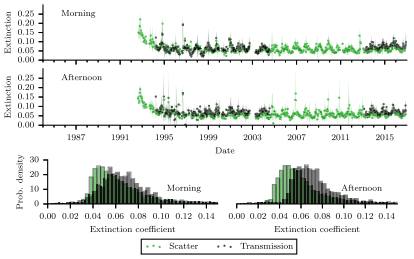

Figure 7. Extinction coefficients and statistical distribution from Narrabri, eastern Australia. Each dot represents the median value over 14 days. The colored banding represents ±3 times the standard error on each median. Dates indicate January 1 for each year.

4.6. Mount Wilson

The BiSON node at Mount Wilson is situated 52 km north–east of Los Angeles, at the Mount Wilson (Hale) Observatory. The measured extinction coefficients over the operational lifetime of the site are shown in Figure 8.

The median morning extinction is 0.085, and afternoon is 0.081. Mount Wilson is an atmospherically interesting site due to its location close to the major city of Los Angeles. Morning data are collected over the San Gabriel mountains, while afternoon data are observed through air over the city, and

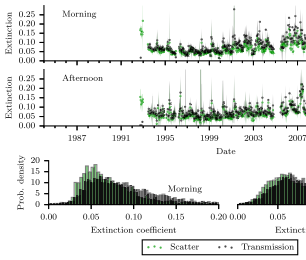

Figure 8. Extinction coefficients and statistical distribution from Mount Wilson, California USA. Each dot represents the median value over 14 days. The colored banding represents ±3 times the standard error on each median. Dates indicate January 1 for each year.

so it may be expected that afternoon extinction values would be worse. The median values found here show that morning and afternoon are generally similar; however, the modal values do show an increase from 0.048 in the morning to 0.070 in the afternoon. There is an upward trend in extinction levels from around 1999, becoming, from 2007, part of the large scatter in values that could be considered to be a result of the highly variable atmosphere near a large city. In clear sky conditions during our observations, the inversion boundary layer traps any pollution at a relatively low altitude—the well-known Los Angeles smog. Because the observatory elevation is well above the inversion layer, the pollution would not be seen except at very high zenith angles that have been specifically excluded from this analysis. At this site, light is collected via two mirrors, known as a coelostat, and it is possible that the increase in scatter is due to gradual reduction in performance due to deterioration of the mirrors; however, since the extinction is determined from a daily calibration, any variation in long-term performance is removed completely and so this is just speculation and the cause of the increased scatter is not clear. During the winter months, Mount Wilson does suffer a high proportion of cloudy days, but when the sky is clear, there appears to be little variation between seasons, both in terms of absolute extinction and standard deviation.

4.7. Site Summary

A summary of the extinction coefficients from 1995 over all seasons is given in Table 2. In order to compare seasonal differences, the values measured over two months for each mid-summer (July–August in the northern hemisphere, January–February in the southern hemisphere) are presented in Table 3, and during mid-winter (January–February in the northern hemisphere, July–August in the southern hemisphere) in Table 4.

5. Conclusion

Over 30 years of helioseismic data have been acquired by BiSON from several international observatories, the locations of which are summarized in Table 1. In this paper, we have made innovative use of these data to derive measurements of atmospheric opacity in the I_2 band, and we have presented the column atmospheric extinction coefficients from each site over the years 1984–2016. This is an important contribution to the literature, as there are limited data on AOD from other sources prior to the mid-1990s.

The median and standard deviation values for mid-summer periods from Table 3, and for mid-winter periods from Table 4, are shown graphically in Figure 9. The mean of the morning and afternoon values were taken for each site, and the extinction distribution standard deviation shown as error bars. We find that the best results are from the Las Campanas and Sutherland observatories with consistent year-round performance. Izaña offers comparable high-performance during the winter months, but becomes the worst site during mid-summer due to high concentrations of mineral dust in the atmosphere from the Western Sahara on the North African coast. Carnarvon similarly suffers degradation during the summer months, which again is most likely due to wind-borne sand in the dry environment of northern Western Australia. Narrabri and Mount Wilson both offer consistent performance.

The atmospheric extinction data presented here and the code to generate the figures are open access and can be downloaded from the University of Birmingham ePapers data archive (Hale 2017).

We would like to thank all those who have been associated with BiSON over the years. We particularly acknowledge the technical assistance at our remote network sites, with sincere apologies to anyone inadvertently missed. At Mount Wilson:

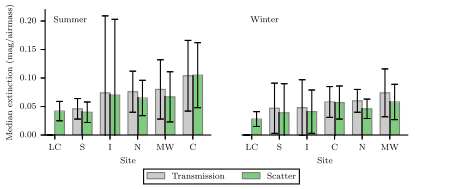

Figure 9. Summary of median extinction coefficients and distribution standard deviations. Only 1995 onwards has been considered in order to remove any exceptional atmospheric events. Left panel: summer periods from Table 3. Right panel: winter periods from Table 4. The mean of the morning and afternoon median coefficients have been taken for each site, and the extinction distribution standard deviation shown as error bars. The *x*-axis labels are the initials of each site name.

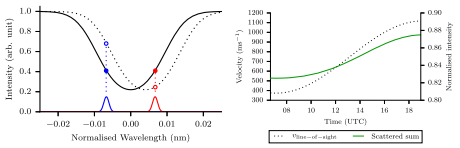

Figure 10. Left: the solid line represents an absorption line with zero line-of-sight velocity offset. The sum of the intensity measurement of the two wings (solid circles) is 0.82. The dashed line represents an offset where $v_{\text{line-of-sight}}$ is 1600 m s^{-1}. The sum of the intensity measurement of the two wings (open circles) rises to 0.93. Neither measurement returns the unity intensity measured outside the absorption line. Right: the expected variation in scattered-sum over typical line-of-sight Doppler velocities, in the absence of any atmospheric extinction. The amplitude of the daily change in effective-intensity varies throughout the year with changes in Earth's orbital velocity.

Stephen Pinkerton, the team of USC undergraduate observing assistants, former USC staff members Maynard Clark, Perry Rose, Natasha Johnson, Steve Padilla, and Shawn Irish, and former UCLA staff members Larry Webster and John Boyden. At Las Campanas: Patricio Pinto, Andres Fuentevilla, Emilio Cerda, Frank Perez, Marc Hellebaut, Patricio Jones, Gastón Gutierrez, Juan Navarro, Francesco Di Mille, Roberto Bermudez, and the staff of LCO. At Izaña: all staff at the IAC who have contributed to running the Mark I instrument over many years. At Sutherland: Pieter Fourie, Willie Koorts, Jaci Cloete, Reginald Klein, John Stoffels, and the staff of SAAO. At Carnarvon: Les Bateman, Les Schultz, Sabrina Dowling-Giudici, Inge Lauw of Williams and Hughes Lawyers, and NBN Co. Ltd. At Narrabri: Mike Hill and the staff of CSIRO. M.N.L. acknowledges the support of The Danish Council for Independent Research–Natural Sciences (Grant DFF-4181-00415). Funding for the Stellar Astrophysics Centre (SAC) is provided by The Danish National Research Foundation (Grant DNRF106). BiSON is funded by the Science and Technology Facilities Council (STFC). We thank the anonymous reviewers for their help and useful comments.

Appendix
Details on BiSON Observations and Line-of-sight Velocity Effects

The BiSON solar spectrometers precisely measure the line-of-sight velocity of the solar surface by looking at the Doppler shift of a Fraunhofer line using a potassium vapor reference cell.

The line-of-sight velocity is dominated by three components: the rotation of the Earth, the orbital velocity of the Earth, and the gravitational redshift. At high airmasses, the extended source of the Sun also suffers differential extinction, which causes the line-of-sight velocity due to solar rotation to become unequally weighted across the solar disc (Davies et al. 2014). In addition to these factors are small oscillations of the solar surface, which are the primary science output of the instrument, but for this analysis, can be considered to be insignificant. The lab-frame absorption line is Zeeman-split by placing the vapor cell in a longitudinal magnetic field, which, when combined with suitable polarization control, places the passbands of the spectrometer on the wings of the corresponding solar absorption line at the points that provide the highest sensitivity

to Doppler shift caused by the line-of-sight velocity. By taking the sum of the intensity measured on both wings of the absorption line, this can be used as a proxy for the total intensity. If the solar absorption line is modeled as a thermally and rotationally broadened Gaussian, then the left panel of Figure 10 shows the expected scattered-sum of the two wings when the Doppler-velocity is zero, and for when the Doppler-velocity is at the maximum expected redshift of 1600 m s^{-1}. Neither scattered-sum correctly recovers the unity-intensity continuum outside of the Gaussian that would be measured from the transmitted light.

The difference between the measurement of the transmitted light and the scattered-sum varies with Doppler-velocity. The right panel of Figure 10 shows the expected effect of this in the absence of atmospheric extinction over the range of line-of-sight velocities experienced during a typical day. The amplitude of the daily change in effective-intensity varies throughout the year with changes in Earth's orbital velocity, and this creates a very small seasonal effect in the scattering extinction coefficients. The increased effective-intensity is most pronounced at high redshifts experienced during the local-afternoon of each site. The gradual increase in effective-intensity will cause any extinction coefficients derived from the scattered-sum to have a slight over-estimate in the morning and a larger underestimate in the afternoon.

If the precise shape and depth of the solar absorption line were known, then it would be possible to correct for this effect and recover the equivalent transmitted intensity. Unfortunately, the solar disc-integrated line profile changes with magnetic activity, and so one can not correct based on a single line profile. Barreto et al. (2014) have resolved this problem for our data from Izaña by making use of a quasi-continuous Langley calibration technique, where they have achieved a *mean bias* (defined as the mean difference between the transmission and scattering extinction coefficients) of ≤0.01.

In practice, extinction data from the scattering detectors offers a slight underestimate of extinction in comparison to transmitted measurements for both morning and afternoon periods. However, this is simply a small systematic offset on the absolute value, and the extinction coefficients observed from real data are otherwise identical to those derived from the transmitted light and show the same trends.

ORCID iDs

S. J. Hale https://orcid.org/0000-0002-6402-8382
W. J. Chaplin https://orcid.org/0000-0002-5714-8618
G. R. Davies https://orcid.org/0000-0002-4290-7351
R. Howe https://orcid.org/0000-0002-3834-8585
M. N. Lund https://orcid.org/0000-0001-9214-5642

References

Barreto, A., Cuevas, E., Pallé, P., et al. 2014, AMT, 7, 4103
Davies, G. R., Chaplin, W. J., Elsworth, Y., & Hale, S. J. 2014, MNRAS, 441, 3009
Fischer, G., Hill, F., Jones, W., et al. 1986, SoPh, 103, 33
García-Gil, A., Muñoz-Tuñón, C., & Varela, A. M. 2010, PASP, 122, 1109
Guerrero, M. A., García-López, R. J., Corradi, R. L. M., et al. 1998, NewAR, 42, 529
Hale, S. J. 2017, BiSON—Atmospheric extinction coefficients in the I$_2$ band—1984 to 2016, University of Birmingham, UK: Birmingham Solar Oscillations Network, http://epapers.bham.ac.uk/3022/
Hale, S. J., Howe, R., Chaplin, W. J., Davies, G. R., & Elsworth, Y. P. 2016, SoPh, 291, 1
Hill, F., Fischer, G., Forgach, S., et al. 1994, SoPh, 152, 351
Jiménez, A., González Jorge, H., & Rabello-Soares, M. C. 1998, A&AS, 129, 413
Kasten, F., & Young, A. T. 1989, ApOpt, 28, 4735
Kilkenny, D. 1995, Obs, 115, 25
King, D. L. 1985, Atmospheric Extinction at the Roque de los Muchachos Observatory, La Palma ROO/La. Palma Technical Note 31, Instituto de Astrofísica de Canarias, La Laguna, Tenerife, Canary Islands, http://www.ing.iac.es/Astronomy/observing/manuals/ps/tech_notes/tn031.pdf
Laken, B. A., Parviainen, H., García-Gil, A., et al. 2016, JCli, 29, 227
Laken, B. A., Parviainen, H., Pallé, E., & Shahbaz, T. 2014, QJRMS, 140, 1058
Roca Cortés, T., & Pallé, P. L. 2014, MNRAS, 443, 1837
Siher, E., Benkhaldoun, Z., & Fossat, E. 2002, EuSA, 13, 159
Siher, E. A., Ortolani, S., Sarazin, M. S., & Benkhaldoun, Z. 2004, Proc. SPIE, 5489, 138

List of References

Agilent (**2005** June). Agilent 8614xB Series Optical Spectrum Analyzer User's Guide. *Datasheet 86140-90U03*, Agilent Technologies Deutschland GmbH, Herrenberger Str. 130, 71034 Böblingen, Germany. URL http://literature.cdn.keysight.com/litweb/pdf/86140-90U03.pdf?id=72624 [page 125]

Agilent (**2012** January). Agilent 86142B Optical Spectrum Analyzer Technical Overview. *Datasheet 5980-0177E*, Agilent Technologies Deutschland GmbH, Herrenberger Str. 130, 71034 Böblingen, Germany. URL http://literature.cdn.keysight.com/litweb/pdf/5980-0177E.pdf?id=1627506 [page 125]

Allen, J. F. (**1983** November). Maintaining Knowledge About Temporal Intervals. *Commun. ACM*, 26(11):832–843. ISSN 0001-0782. URL http://dx.doi.org/10.1145/182.358434 [page 38]

Antia, H. M. and Basu, S. (**2010** September). Solar Rotation Rate During the Cycle 24 Minimum in Activity. ApJ, 720:494–502. URL http://dx.doi.org/10.1088/0004-637X/720/1/494 [page 25]

Antia, H. M., Basu, S., and Chitre, S. M. (**2008** July). Solar Rotation Rate and Its Gradients During Cycle 23. ApJ, 681:680–692. URL http://dx.doi.org/10.1086/588523 [page 25]

Appourchaux, T., Belkacem, K., Broomhall, A.-M., et al. (**2010** February). The quest for the solar g modes. A&A Rev., 18:197–277. URL http://dx.doi.org/10.1007/s00159-009-0027-z [page 5]

ASTM (**2012**). G173-03(2012) Standard Tables for Reference Solar Spectral Irradiances: Direct Normal and Hemispherical on 37° Tilted Surface. *International Standard G173-03*, ASTM International, West Conshohocken, PA, USA. URL http://dx.doi.org/10.1520/G0173-03R12 [page 141]

Baker, B. (**2011a**). How delta-sigma ADCs work, Part 1. *Data Acquisition SLYT423A*, Texas Instruments Inc. URL http://www.ti.com/lit/an/slyt423a/slyt423a.pdf [page 204]

Baker, B. (**2011b**). How delta-sigma ADCs work, Part 2. *Data Acquisition SLYT438*, Texas Instruments Inc. URL http://www.ti.com/lit/an/slyt438/slyt438.pdf [page 205]

Barnes, I. (**2008** January). The Installation of the Sutherland temperature controllers in 2007 August. *BiSON Technical Report Series 294*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2036/ [page 65]

Barnes, I. (**2009a** January). The Sutherland temperature controllers. *BiSON Technical Report Series 316*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2040/ [page 65]

Barnes, I. (**2009b** September). The Jabba Temperature Controllers. *BiSON Technical Report Series 325*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK [page 63]

Barnes, I. (**2010** November). The Narrabri Temperature Controller. *BiSON Technical Report Series 333*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL `http://epapers.bham.ac.uk/2047/` [page 74]

Barnes, I. (**2011** May). The Jabba Detector (JDET). *BiSON Technical Report Series 342*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK [page 180]

Barnes, I. and Hale, S. J. (**2005** September). Carnarvon Trip Report — May 2005. *BiSON Technical Report Series 253*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL `http://epapers.bham.ac.uk/2026/` [page 60]

Barnes, I. and Hale, S. J. (**2010** November). The Installation of New Temperature Controllers at Narrabri in 2010 February. *BiSON Technical Report Series 332*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL `http://epapers.bham.ac.uk/2046/` [page 74]

Barnes, I., Jackson, B., and Miller, B. A. (**2003** November). BeebSub #56A and BeebSub #56B — BBC Computer Substitutes. *BiSON Technical Report Series 218*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL `http://epapers.bham.ac.uk/2014/` [page 50]

Barnes, I., Jackson, B., and Miller, B. A. (**2004** November). BeebSub #56C — An Improved BBC Computer Substitute. *BiSON Technical Report Series 243*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL `http://epapers.bham.ac.uk/2023/` [page 50]

Barnes, I. and Miller, B. A. (**2006** October). The Grand Opening of the Sutherland Zoo. *BiSON Technical Report Series 276*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL `http://epapers.bham.ac.uk/2032/` [page 65]

Barnes, I. and Miller, B. A. (**2009** September). Jabba is Returned to Carnarvon in 2009 July. *BiSON Technical Report Series 323*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL `http://epapers.bham.ac.uk/2043/` [page 60, 215, 243]

Barnes, I. and Miller, B. A. (**2011a** May). A Computer-Controlled Temperature Controller with Five Configurable Inputs. *BiSON Technical Report Series 335*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK [page 209, 210, 211]

Barnes, I. and Miller, B. A. (**2011b** May). The mount controller: A Digital Autoguider for Las Campanas. *BiSON Technical Report Series 344*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL `http://epapers.bham.ac.uk/2049/` [page 69]

Barnes, I. and Miller, B. A. (**2012** January). The Mount Controller: A Digital Autoguider for Carnarvon. *BiSON Technical Report Series 349*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK [page 63]

Barnes, I., Miller, B. A., and Jackson, B. (**2007** February). The Removal of Jabba from Carnarvon in 2006 November. *BiSON Technical Report Series 282*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL `http://epapers.bham.ac.uk/2034/` [page 60]

Barreto, A., Cuevas, E., Pallé, P. L., et al. (**2014** December). Recovering long-term aerosol optical depth series (1976-2012) from an astronomical potassium-based resonance scattering spectrometer. *Atmospheric Measurement Techniques*, 7:4103–4116. URL http://dx.doi.org/10.5194/amt-7-4103-2014 [page 94]

Basu, S. and Antia, H. M. (**2004** May). Constraining Solar Abundances Using Helioseismology. ApJ, 606:L85–L88. URL http://dx.doi.org/10.1086/421110 [page 24]

Basu, S. and Antia, H. M. (**2010** July). Characteristics of Solar Meridional Flows during Solar Cycle 23. ApJ, 717:488–495. URL http://dx.doi.org/10.1088/0004-637X/717/1/488 [page 25]

Basu, S., Broomhall, A.-M., Chaplin, W. J., and Elsworth, Y. P. (**2012** October). Thinning of the Sun's Magnetic Layer: The Peculiar Solar Minimum Could Have Been Predicted. ApJ, 758:43. URL http://dx.doi.org/10.1088/0004-637X/758/1/43 [page 25]

Basu, S., Chaplin, W. J., Elsworth, Y. P., New, R., and Serenelli, A. M. (**2009** July). Fresh Insights on the Structure of the Solar Core. ApJ, 699:1403–1417. URL http://dx.doi.org/10.1088/0004-637X/699/2/1403 [page 24]

Biermann, L. (**1946** August). Zur Deutung der chromosphärischen Turbulenz und des Exzesses der UV-Strahlung der Sonne. *Naturwissenschaften*, 33:118–119. URL http://dx.doi.org/10.1007/BF00738265 [page 11]

Brookes, J. R. (**1974** September). *An Accurate Measurement of the Difference Between the Solar and Laboratory Wavelengths of the Sodium Resonance Lines*. Ph.D. thesis, Department of Physics, University of Birmingham, UK [page 32]

Brookes, J. R., Isaak, G. R., and van der Raay, H. B. (**1976** January). Observation of free oscillations of the sun. Nature, 259:92–95. URL http://dx.doi.org/10.1038/259092a0 [page 18, 40, 41]

Brookes, J. R., Isaak, G. R., and van der Raay, H. B. (**1978a** October). A resonant-scattering solar spectrometer. MNRAS, 185:1–18. URL http://dx.doi.org/10.1093/mnras/185.1.1 [page 32]

Brookes, J. R., Isaak, G. R., and van der Raay, H. B. (**1978b** October). The observation of a rotating body using high-resolution spectroscopy. MNRAS, 185:19–22. URL http://dx.doi.org/10.1093/mnras/185.1.19 [page 114]

Broomhall, A.-M. (**2007** November). *The hunt for low-frequency modes of oscillation of the Sun: application of statistical techniques and instrumentation*. Ph.D. thesis, School of Physics and Astronomy, University of Birmingham, UK [page 114]

Broomhall, A. M., Chaplin, W. J., Elsworth, Y. P., and Appourchaux, T. (**2007** July). Needles in haystacks: how to use contemporaneous data in the search for low-frequency modes of oscillation of the Sun. MNRAS, 379:2–10. URL http://dx.doi.org/10.1111/j.1365-2966.2007.11905.x [page 4]

Burr-Brown (**1984**). Dual Low Noise Precision Difet OPERATIONAL AMPLIFIER. *Datasheet OPA2111*, Texas Instruments Inc., 12500 TI Boulevard Dallas, Texas, 75243, USA. URL http://www.ti.com/lit/ds/symlink/opa2111.pdf. Revised 2009 [page 181, 187, 196]

Carter, C. L., Snodgrass, H. B., and Bryja, C. (**1992** May). Telluric water vapor contamination of the Mount Wilson solar Doppler measurements. Sol. Phys., 139(1):13–24. URL http://dx.doi.org/10.1007/BF00147879 [page 31]

CC BY 4.0 (**2013** November). Creative Commons Attribution 4.0 International. URL https://creativecommons.org/licenses/by/4.0/ [page 295]

Centronic (**2017**). General Purpose Silicon Sensors (Series 5T). *Datasheet Series 5*, Centronic, Centronic House, King Henry's Drive, Croydon, CR9 0BG. URL http://www.centronic.co.uk/get/5a3fdbfe-ef10-4c8e-9447-f60e6023730b.pdf/general-purpose-detector-series-5t-data-sheet [page 184, 188, 195]

Chaplin, W. J. (**2006**). *Music of the Sun: The Story of Helioseismology*. Oneworld Publications. ISBN 9781851684519. URL https://books.google.co.uk/books?id=eJTvAAAAMAAJ [page 11]

Chaplin, W. J., Appourchaux, T., Baudin, F., et al. (**2004a** October). Flag Hare-And Exercise: on the Extraction of Sectoral Mode Splittings from Full-Disc Sun-As Data. In Danesy, D. (editor), *SOHO 14 Helio- and Asteroseismology: Towards a Golden Future*, volume 559 of *ESA Special Publication*, 356. URL http://adsabs.harvard.edu/abs/2004ESASP.559..356C [page 4]

Chaplin, W. J., Christensen-Dalsgaard, J., Elsworth, Y. P., et al. (**1999** September). Rotation of the solar core from BiSON and LOWL frequency observations. MNRAS, 308:405–414. URL http://dx.doi.org/10.1046/j.1365-8711.1999.02691.x [page 25]

Chaplin, W. J., Dumbill, A. M., Elsworth, Y. P., et al. (**2003** August). Studies of the solar mean magnetic field with the Birmingham Solar-Oscillations Network (BiSON). MNRAS, 343:813–818. URL http://dx.doi.org/10.1046/j.1365-8711.2003.06715.x [page 36]

Chaplin, W. J., Dumbill, A. M., Elsworth, Y. P., et al. (**2004b** April). Extended search for correlation between solar mean magnetic field BiSON data and coronal mass ejections. Sol. Phys., 220:307–316. URL http://dx.doi.org/10.1023/B:SOLA.0000031398.11428.f8 [page 36]

Chaplin, W. J., Dumbill, A. M., Elsworth, Y. P., et al. (**2005a** January). The Search for Correlation between BiSON SMMF Data and CME Events. *Highlights of Astronomy*, 13:141. URL http://adsabs.harvard.edu/abs/2005HiA....13..141C [page 36]

Chaplin, W. J., Elsworth, Y. P., Howe, R., et al. (**1996a** September). BiSON performance. Sol. Phys., 168:1–18. URL http://dx.doi.org/10.1007/BF00145821 [page 45]

Chaplin, W. J., Elsworth, Y. P., Howe, R., et al. (**1997** October). Techniques used in the analysis of data collected by the Birmingham Solar-Oscillations Network (BiSON). II. Frequency domain analysis & data merging. A&AS, 125:195–205. URL http://dx.doi.org/10.1051/aas:1997371 [page 38, 82]

Chaplin, W. J., Elsworth, Y. P., Isaak, G. R., et al. (**1996b** November). Solar core rotation from low-degree BiSON p-mode splittings: 1981-95. MNRAS, 283:L31–L35. URL http://dx.doi.org/10.1093/mnras/283.2.L31 [page 25]

Chaplin, W. J., Elsworth, Y. P., Isaak, G. R., et al. (**2002** November). Peak finding at low signal-to-noise ratio: low-ℓ solar acoustic eigenmodes at $n \leq 9$ from the analysis of BiSON data. MNRAS, 336:979–991. URL http://dx.doi.org/10.1046/j.1365-8711.2002.05834.x [page 4]

Chaplin, W. J., Elsworth, Y. P., Isaak, G. R., et al. (**2004c** October). Novel techniques for the identification of noise contributions to full-disc helioseismic power spectra. In Danesy, D. (editor), *SOHO 14 Helio- and Asteroseismology: Towards a Golden Future*, volume 559 of *ESA SP-*, 360. URL http://adsabs.harvard.edu/abs/2004ESASP.559..360C [page 85]

Chaplin, W. J., Elsworth, Y. P., Isaak, G. R., et al. (**2005b** May). Noise characteristics of full-disc helioseismic observations made by resonant scattering spectrometers. MNRAS, 359:607–614. URL http://dx.doi.org/10.1111/j.1365-2966.2005.08920.x [page 42, 85, 113, 114]

Chaves, J. (**2008**). *Introduction to Nonimaging Optics*. Optical Science and Engineering. CRC Press. ISBN 9781420054323. URL https://books.google.co.uk/books?id=ykk1xQc25c0C [page 145, 150]

Chou, D.-Y., Sun, M.-T., Huang, T.-Y., et al. (**1995** September). Taiwan Oscillation Network. Sol. Phys., 160:237–243. URL http://dx.doi.org/10.1007/BF00732806 [page 24]

Christensen-Dalsgaard, J., Dappen, W., Ajukov, S. V., et al. (**1996** May). The Current State of Solar Modeling. *Science*, 272:1286–1292. URL http://dx.doi.org/10.1126/science.272.5266.1286 [page 4]

Christensen-Dalsgaard, J., Gough, D. O., and Thompson, M. J. (**1991** September). The depth of the solar convection zone. ApJ, 378:413–437. URL http://dx.doi.org/10.1086/170441 [page 24]

Claverie, A., Isaak, G. R., McLeod, C. P., van der Raay, H. B., and Roca Cortés, T. (**1979** December). Solar structure from global studies of the 5-minute oscillation. Nature, 282:591–594. URL http://dx.doi.org/10.1038/282591a0 [page 18, 51]

Claverie, A., Isaak, G. R., McLeod, C. P., van der Raay, H. B., and Roca Cortés, T. (**1981** October). Rapid rotation of the solar interior. Nature, 293:443–445. URL http://dx.doi.org/10.1038/293443a0 [page 25]

Cooley, J. W. and Turkey, J. W. (**1965**). An algorithm for the machine calculation of complex Fourier series. *Math. Comput.*, 19:297–301. URL http://dx.doi.org/10.1090/S0025-5718-1965-0178586-1 [page 47, 218]

Corbard, T., Di Mauro, M. P., Sekii, T., and GOLF Team (**1998**). The Solar Internal Rotation from GOLF Splittings. In Korzennik, S. (editor), *Structure and Dynamics of the Interior of the Sun and Sun-like Stars*, volume 418 of *ESA Special Publication*, 741. URL http://adsabs.harvard.edu/abs/1998ESASP.418..741C [page 25]

Davidson, C. and Williams, H. K. (**2004** March). Progress Report on the Construction of the South African Dome — 1989 November/December. *BiSON Technical Report Series 226*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2017/ [page 65]

Davies, G. R. (**2011** July). *Investigating the low-frequency stability of BiSON's resonant scattering spectrometers*. Ph.D. thesis, School of Physics and Astronomy, University of Birmingham, UK. URL https://etheses.bham.ac.uk/id/eprint/1609/ [page 114]

Davies, G. R., Broomhall, A. M., Chaplin, W. J., Elsworth, Y. P., and Hale, S. J. (**2014** April). Low-frequency, low-degree solar p-mode properties from 22 years of Birmingham Solar Oscillations Network data. MNRAS, 439:2025–2032. URL http://dx.doi.org/10.1093/mnras/stu080 [page 4]

Davies, G. R., Chaplin, W. J., Elsworth, Y. P., and Hale, S. J. (**2014**). BiSON data preparation: a correction for differential extinction and the weighted averaging of contemporaneous data. *Monthly Notices of the Royal Astronomical Society*, 441(4):3009–3017. URL http://dx.doi.org/10.1093/mnras/stu803 [page 5, 38, 43, 82, 90, 93]

de Jager, C. (**1954** April). High-energy Microturbulence in the Solar Photosphere. Nature, 173:680–681. URL http://dx.doi.org/10.1038/173680b0 [page 119]

Dench, P. and Gregg, A. (**2010**). *Carnarvon and Apollo: One Giant Leap for a Small Australian Town*. Rosenberg Publishing Pty, Limited. ISBN 9781877058974. URL https://books.google.co.uk/books?id=EmD6QwAACAAJ [page 54]

Denny, B. (**1997**). Astronomy Common Object Model (ASCOM). https://ascom-standards.org/ [page 230]

Deubner, F.-L. (**1975** November). Observations of low wavenumber nonradial eigenmodes of the sun. A&A, 44:371–375. URL http://adsabs.harvard.edu/abs/1975A&A....44..371D [page 16, 17]

Di Mille, F., Boutsia, K., Infante, L., et al. (**2018**). Las Campanas Observatory. volume 10704, 10,704 – 10,704 – 9. URL http://dx.doi.org/10.1117/12.2312825 [page 66]

Domingo, V., Fleck, B., and Poland, A. I. (**1995** December). The SOHO Mission: an Overview. Sol. Phys., 162:1–37. URL http://dx.doi.org/10.1007/BF00733425 [page 21]

Duffett-Smith, P. (**1988**). *Practical Astronomy with Your Calculator*. Cambridge University Press. ISBN 9780521356992. URL https://books.google.co.uk/books?id=DwJfCtzaVvYC [page 285, 290, 291, 292, 293]

Dumusque, X., Glenday, A., Phillips, D. F., et al. (**2015** December). HARPS-N Observes the Sun as a Star. ApJ, 814:L21. URL http://dx.doi.org/10.1088/2041-8205/814/2/L21 [page 245]

Duvall, Jr., T. L., Dziembowski, W. A., Goode, P. R., et al. (**1984** July). Internal rotation of the sun. Nature, 310:22–25. URL http://dx.doi.org/10.1038/310022a0 [page 25]

Duvall, Jr., T. L. and Harvey, J. W. (**1984** July). Rotational frequency splitting of solar oscillations. Nature, 310:19–22. URL http://dx.doi.org/10.1038/310019a0 [page 25]

Duvall, Jr., T. L., Harvey, J. W., Libbrecht, K. G., Popp, B. D., and Pomerantz, M. A. (**1988** January). Frequencies of solar p-mode oscillations. ApJ, 324:1158–1171. URL http://dx.doi.org/10.1086/165971 [page 24]

Edmondson, W. and Egerton, A. (**1927**). The Vapour Pressures and Melting Points of Sodium and Potassium. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 113(765):520–533. ISSN 0950-1207. URL http://dx.doi.org/10.1098/rspa.1927.0005 [page 122, 123]

Elsworth, Y. (**1992** August). Update on My Report on the Trip to Mount Wilson (Hale) Observatory. *BiSON Technical Report Series 5*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2003/ [page 77]

Elsworth, Y. P., Broomhall, A.-M., Gosain, S., et al. (**2015** December). The Importance of Long-Term Synoptic Observations and Data Sets for Solar Physics and Helioseismology. *Space Science Reviews*, 196(1):137–166. ISSN 1572-9672. URL http://dx.doi.org/10.1007/s11214-015-0212-5 [page 5]

Elsworth, Y. P., Howe, R., Isaak, G. R., McLeod, C. P., and New, R. (**1990** October). Evidence from solar seismology against non-standard solar-core models. Nature, 347:536–539. URL http://dx.doi.org/10.1038/347536a0 [page 24]

Elsworth, Y. P., Howe, R., Isaak, G. R., et al. (**1995** October). Techniques used in the analysis of solar oscillations data from the BiSON (University of Birmingham) network. I. Daily calibration. A&AS, 113:379. URL http://adsabs.harvard.edu/abs/1995A&AS..113..379E [page 38, 42, 82]

Elsworth, Y. P., Jefferies, S. M., McLeod, C. P., et al. (**1989** March). The 160 minute solar oscillation - an artifact? ApJ, 338:557–562. URL http://dx.doi.org/10.1086/167217 [page 19]

Espagnet, O., Muller, R., Roudier, T., Mein, N., and Mein, P. (**1995** January). Penetration of the solar granulation into the photosphere: height dependence of intensity and velocity fluctuations. A&AS, 109:79–108. URL http://adsabs.harvard.edu/abs/1995A&AS..109..79E [page 6]

Evans, J. W. and Michard, R. (**1962a** September). Observational Study of Macroscopic Inhomogeneities in the Solar Atmosphere. II. Brightness Fluctuations in Fraunhofer Lines and the Continuum. ApJ, 136:487. URL http://dx.doi.org/10.1086/147402 [page 15]

Evans, J. W. and Michard, R. (**1962b** September). Observational Study of Macroscopic Inhomogeneities in the Solar Atmosphere. III. Vertical Oscillatory Motions in the Solar Photosphere. ApJ, 136:493. URL http://dx.doi.org/10.1086/147403 [page 15]

Fischer, G., Hill, F., Jones, W., et al. (**1986** January). A simple irradiance monitor for testing solar global oscillation network sites. Sol. Phys., 103:33–39. URL http://dx.doi.org/10.1007/BF00154856 [page 99]

Fossat, E. (**1991** May). The IRIS network for full disk helioseismology - Present status of the programme. Sol. Phys., 133:1–12. URL http://dx.doi.org/10.1007/BF00149818 [page 23]

Fossat, E., Boumier, P., Corbard, T., et al. (**2017** August). Asymptotic g modes: Evidence for a rapid rotation of the solar core. A&A, 604:A40. URL http://dx.doi.org/10.1051/0004-6361/201730460 [page 5]

Fossat, E. and IRIS Group (**2002** June). 11 years of IRIS network exploitation. In Combes, F. and Barret, D. (editors), *SF2A-2002: Semaine de l'Astrophysique Francaise*, 521. URL http://adsabs.harvard.edu/abs/2002sf2a.conf..521F [page 23]

Fröhlich, C., Romero, J., Roth, H., et al. (**1995** December). VIRGO: Experiment for Helioseismology and Solar Irradiance Monitoring. Sol. Phys., 162:101–128. URL http://dx.doi.org/10.1007/BF00733428 [page 21]

Fukugita, M., Ichikawa, T., Gunn, J. E., et al. (**1996** April). The Sloan Digital Sky Survey Photometric System. AJ, 111:1748. URL http://dx.doi.org/10.1086/117915 [page 154]

Gabriel, A. H., Grec, G., Charra, J., et al. (**1995** December). Global Oscillations at Low Frequency from the SOHO Mission (GOLF). Sol. Phys., 162:61–99. URL http://dx.doi.org/10.1007/BF00733427 [page 21]

García, R. A., Bertello, L., Turck-Chièze, S., et al. (**2001** January). Analysis of low frequency signal with the GOLF experiment: methodology and results. In Wilson, A. and Pallé, P. L. (editors), *SOHO 10/GONG 2000 Workshop: Helio- and Asteroseismology at the Dawn of the Millennium*, volume 464 of *ESA Special Publication*, 473–478. URL http://adsabs.harvard.edu/abs/2001ESASP.464..473G [page 4]

García, R. A., Jiménez-Reyes, S. J., Turck-Chièze, S., and Mathur, S. (**2004** October). Helioseismology from the Blue and Red Wings of the NA Profile as Seen by GOLF. In Danesy, D. (editor), *SOHO 14 Helio- and Asteroseismology: Towards a Golden Future*, volume 559 of *ESA Special Publication*, 432. URL http://adsabs.harvard.edu/abs/2004ESASP.559..432G [page 6]

García, R. A., Turck-Chièze, S., Jiménez-Reyes, S. J., et al. (**2007** June). Tracking Solar Gravity Modes: The Dynamics of the Solar Core. *Science*, 316:1591. URL http://dx.doi.org/10.1126/science.1140598 [page 5]

García-Gil, A., Muñoz-Tuñón, C., and Varela, A. M. (**2010**). Atmosphere Extinction at the ORM on La Palma: A 20 yr Statistical Database Gathered at the Carlsberg Meridian Telescope. *Publications of the Astronomical Society of the Pacific*, 122(895):1109. URL http://stacks.iop.org/1538-3873/122/i=895/a=1109 [page 99, 100]

Gough, D. O. (**1981** September). A new measure of the solar rotation. MNRAS, 196:731–745. URL http://dx.doi.org/10.1093/mnras/196.3.731 [page 25]

Gough, D. O. (editor) (**1986**). *Seismology of the sun and the distant stars; Proceedings of the NATO Advanced Research Workshop, Cambridge, England, June 17-21, 1985*, volume 169 of *NATO Advanced Science Institutes (ASI) Series C*. URL http://adsabs.harvard.edu/abs/1986ASIC..169.....G [page 24]

Graeme, J. (**1996**). *Photodiode Amplifiers: Op Amp Solutions*. Gain technology. McGraw Hill. ISBN 9780070242470. URL https://books.google.co.uk/books?id=sHV0c5hBW4QC [page 182]

Grec, G., Fossat, E., and Pomerantz, M. A. (**1980** December). Solar oscillations - Full disk observations from the geographic South Pole. Nature, 288:541–544. URL http://dx.doi.org/10.1038/288541a0 [page 20]

Guerrero, M. A., García-López, R. J., Corradi, R. L. M., et al. (**1998** November). Extinction over the Canarian observatories: the limited influence of Saharan dust. New A Rev., 42:529–532. URL http://dx.doi.org/10.1016/S1387-6473(98)00066-9 [page 99, 100]

Hale, G. E. (**1908a** April). The Tower Telescope of the Mount Wilson Solar Observatory. ApJ, 27:204. URL http://dx.doi.org/10.1086/141546 [page 12]

Hale, G. E. (**1908b** November). On the Probable Existence of a Magnetic Field in Sun-Spots. ApJ, 28:315. URL http://dx.doi.org/10.1086/141602 [page 12]

Hale, S. J. (**2003** December). *Scientific advancements in analysis of solar oscillation data*. Master's thesis, School of Physics and Astronomy, University of Birmingham, UK. URL https://etheses.bham.ac.uk/id/eprint/5952/ [page 38, 82]

Hale, S. J. (**2012** July). Sutherland 2012 June Installation of Tiger Counters. *BiSON Technical Report Series 358*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2053/ [page 65]

Hale, S. J. (**2013a** September). Replacement of counters, temperature monitor, and Peltiers in Carnarvon in 2013 September. *BiSON Technical Report Series 361*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2056/ [page 63]

Hale, S. J. (**2013b** December). The Installation of a Digital Autoguider in Sutherland in 2013 November. *BiSON Technical Report Series 362*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2057/ [page 65]

Hale, S. J. (**2014a** June). Autoguider repairs at Mount Wilson in 2014 April. *BiSON Technical Report Series 365*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2060/ [page 82]

Hale, S. J. (**2014b** July). Blind and mount controller repairs in Las Campanas in 2014 June. *BiSON Technical Report Series 366*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2061/ [page 70]

Hale, S. J. (**2014c** December). A check on Carnarvon in 2014 November following NBN construction. *BiSON Technical Report Series 367*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2062/ [page 60]

Hale, S. J. (**2015a** May). BiSON - All Sites - 1995 to 2014 - Performance Check. URL http://dx.doi.org/10.25500/eData.bham.00000059 [page 18, 45, 295]

Hale, S. J. (**2015b** May). The Hannibal Temperature Controllers. *BiSON Technical Report Series 372*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2067/ [page 69]

Hale, S. J. (**2015c** May). New temperature controller for Hannibal in Las Campanas in 2015 April. *BiSON Technical Report Series 373*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2068/ [page 69]

Hale, S. J. (**2015d** December). Temperature control repairs in Narrabri in 2015 October. *BiSON Technical Report Series, Number 375*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2084/ [page 74]

Hale, S. J. (**2015e** December). Aft-oven removed from Klaus at Mount Wilson in 2015 December. *BiSON Technical Report Series, Number 376*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2099/ [page 82]

Hale, S. J. (**2015f**). Performance of the Birmingham Solar-Oscillations Network (BiSON). In *National Astronomy Meeting 2015*. URL http://epapers.bham.ac.uk/3027/ [page 45]

Hale, S. J. (**2016a** March). Repairs to Jabba at Carnarvon in 2016 March. *BiSON Technical Report Series, Number 378*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/3130/ [page 63]

Hale, S. J. (**2016b** June). Autoguider repairs at Mount Wilson in 2016 April. *BiSON Technical Report Series, Number 380*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/3132/ [page 82]

Hale, S. J. (**2016c** August). Repairs to Jabba at Carnarvon in 2016 August. *BiSON Technical Report Series, Number 381*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/3133/ [page 63]

Hale, S. J. (**2016d** December). Fibre-feed tests at Mount Wilson in 2016 September. *BiSON Technical Report Series, Number 382*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/3134/ [page 82, 240]

Hale, S. J. (**2017a**). BiSON - Atmospheric extinction coefficients in the I_c band - 1984 to 2016. URL http://dx.doi.org/10.25500/eData.bham.00000025 [page 89, 303]

Hale, S. J. (**2017b** September). Fibre-feed tests at Izaña in 2017 May. *BiSON Technical Report Series, Number 384*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/3135/ [page 211, 218, 254]

Hale, S. J. (**2017c** September). Fibre-feed tests at Izaña in 2017 September. *BiSON Technical Report Series, Number 385*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/3136/ [page 211, 219, 224, 254]

Hale, S. J. (**2017d** December). Mount controller repairs at Las Campanas in 2017 November. *BiSON Technical Report Series, Number 386*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/3137/ [page 70]

Hale, S. J. (**2017e** May). Sunrise on Research. Gatsby Charitable Foundation Trust. URL http://epapers.bham.ac.uk/2987/. #TechniciansMakeitHappen [page 241]

Hale, S. J. (**2018a** February). Peltier repairs at Narrabri in 2018 February. *BiSON Technical Report Series, Number 388*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/3138/ [page 74]

Hale, S. J. (**2018b** August). Klaus fibre conversion at Mount Wilson in 2018 June. *BiSON Technical Report Series, Number 390*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/3140/ [page 82, 224, 248, 255]

Hale, S. J. (**2018c** November). Hannibal vapour cell replaced at Las Campanas in 2018 October. *BiSON Technical Report Series, Number 391*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/3174/ [page 73]

Hale, S. J., Chaplin, W. J., Davies, G. R., et al. (**2017**). Atmospheric Extinction Coefficients in the I_c Band for Several Major International Observatories: Results from the BiSON Telescopes, 1984-2016. *The Astronomical Journal*, 154(3):89. URL https://doi.org/10.3847/1538-3881/aa81d0 [page 89, 303]

Hale, S. J. and Davies, G. R. (**2013** May). The Installation of a Digital Autoguider and Tiger Counters in Narrabri in 2013 April. *BiSON Technical Report Series 360*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2055/ [page 74]

Hale, S. J., Howe, R., Chaplin, W. J., Davies, G. R., and Elsworth, Y. P. (**2016** January). Performance of the Birmingham Solar-Oscillations Network (BiSON). Sol. Phys., 291(1):1–28. ISSN 1573-093X. URL http://dx.doi.org/10.1007/s11207-015-0810-0 [page 11, 45, 295]

Hale, S. J. and Miller, B. A. (**2006** February). The Grand Opening of the Las Campanas Zoo. *BiSON Technical Report Series 261*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2029/ [page 69]

Hale, S. J. and Miller, B. A. (**2012** January). The Installation of a Digital Autoguider in Carnarvon in 2011 September. *BiSON Technical Report Series 348*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2051/ [page 63]

Hale, S. J. and Ross, E. (**2018** June). Jabba fibre conversion at Carnarvon in 2018 April. *BiSON Technical Report Series, Number 389*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/3139/ [page 63, 224, 255]

Harvey, J. W., Hill, F., Hubbard, R. P., et al. (**1996** May). The Global Oscillation Network Group (GONG) Project. *Science*, 272:1284–1286. URL http://dx.doi.org/10.1126/science.272.5266.1284 [page 23]

Hicks, W. T. (**1963**). Evaluation of Vapor-Pressure Data for Mercury, Lithium, Sodium, and Potassium. *The Journal of Chemical Physics*, 38(8):1873–1880. URL http://dx.doi.org/10.1063/1.1733889 [page 122, 123]

Hill, F., Fischer, G., Forgach, S., et al. (**1994** July). The Global Oscillation Network Group site survey, 2: Results. Sol. Phys., 152:351–379. URL http://dx.doi.org/10.1007/BF00680444 [page 99]

Horowitz, P. and Hill, W. (**2015**). *The Art of Electronics*. Cambridge University Press. ISBN 9780521809269. URL https://books.google.co.uk/books?id=LAiWPwAACAAJ [page 187, 191, 192, 197, 198, 199, 200, 201]

Houdek, G., Balmforth, N. J., Christensen-Dalsgaard, J., and Gough, D. O. (**1999** November). Amplitudes of stochastically excited oscillations in main-sequence stars. A&A, 351:582–596. URL http://adsabs.harvard.edu/abs/1999A&A...351..582H [page 4]

Houdek, G., Chaplin, B., and Elsworth, Y. P. (**1998** August). How are solar oscillations damped? *Astronomy and Geophysics*, 39:30. URL http://adsabs.harvard.edu/abs/1998A&G....39d..30H [page 4]

Houdek, G., Chaplin, W. J., Appourchaux, T., et al. (**2001** October). Changes in convective properties over the solar cycle: effect on p-mode damping rates. MNRAS, 327:483–487. URL http://dx.doi.org/10.1046/j.1365-8711.2001.04749.x [page 4]

Howard, R. (**1962** July). Preliminary Solar Magnetograph Observations with Small Apertures. ApJ, 136:211. URL http://dx.doi.org/10.1086/147365 [page 15]

Howe, R., Chaplin, W. J., Davies, G. R., et al. (**2018** October). Changes in the sensitivity of solar p-mode frequency shifts to activity over three solar cycles. MNRAS, 480:L79–L84. URL http://dx.doi.org/10.1093/mnrasl/sly124 [page 25]

Howe, R., Christensen-Dalsgaard, J., Hill, F., et al. (**2009** August). A Note on the Torsional Oscillation at Solar Minimum. ApJ, 701:L87–L90. URL http://dx.doi.org/10.1088/0004-637X/701/2/L87 [page 25]

Howe, R., Davies, G. R., Chaplin, W. J., Elsworth, Y. P., and Hale, S. J. (**2015** December). Validation of solar-cycle changes in low-degree helioseismic parameters from the Birmingham Solar-Oscillations Network. MNRAS, 454:4120–4141. URL http://dx.doi.org/10.1093/mnras/stv2210 [page 25]

Howe, R., Davies, G. R., Chaplin, W. J., et al. (**2017** September). The Sun in transition? Persistence of near-surface structural changes through Cycle 24. MNRAS, 470:1935–1942. URL http://dx.doi.org/10.1093/mnras/stx1318 [page 25]

Hoyng, P. (**1989** October). On the sensitivity of resonant scattering spectrometers for whole-disk solar velocity oscillation measurements. ApJ, 345:1088–1103. URL http://dx.doi.org/10.1086/167978 [page 7, 42, 87, 113]

Isaak, G. R. (**1961** February). An Atomic Beam Spectrophotometer. Nature, 189:373–374. URL http://dx.doi.org/10.1038/189373a0 [page 28]

Isaak, G. R. (**1992**). Global Solar Oscillations. *Carnegie Institution of Washington Year Book*, 92:149–156. URL https://archive.org/details/yearbookcarnegi199293carn [page 23]

Isaak, G. R., McLeod, C. P., Pallé, P. L., van der Raay, H. B., and Roca Cortés, T. (**1989** January). Solar oscillations as seen in the NaI and KI absorption lines. A&A, 208:297–302. URL http://adsabs.harvard.edu/abs/1989A&A...208..297I [page 32]

Jackson, B. (**2003** November). Potassium Vapour Cells. *BiSON Technical Report Series 220*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK [page 122]

Jackson, B. and Miller, B. A. (**2004** November). The Grand Opening of the Narrabri Zoo in 2004 July. *BiSON Technical Report Series 241*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2022/ [page 73]

James, J. and Sternberg, R. (**1969**). *The Design of Optical Spectrometers*. Chapman & Hall. URL https://books.google.co.uk/books?id=M15KAAAAIAAJ [page 263]

Jimenez, A. J., Gonzalez Jorge, H., and Rabello-Soares, M. C. (**1998** April). Diurnal atmospheric extinction over Teide Observatory (Tenerife, Canary Islands). A&AS, 129:413–423. URL http://dx.doi.org/10.1051/aas:1998193 [page 99]

Johnson, H. and Graham, M. (**1993**). *High-speed Digital Design: A Handbook of Black Magic*. Prentice Hall Modern Semicondu. Prentice Hall. ISBN 9780133957242. URL https://books.google.co.uk/books?id=H5SsQgAACAAJ [page 189]

Johnson, J. B. (**1928** July). Thermal Agitation of Electricity in Conductors. *Phys. Rev.*, 32:97–109. URL http://dx.doi.org/10.1103/PhysRev.32.97 [page 186]

Karak, B. B. and Nandy, D. (**2012** December). Turbulent Pumping of Magnetic Flux Reduces Solar Cycle Memory and thus Impacts Predictability of the Sun's Activity. ApJ, 761:L13. URL http://dx.doi.org/10.1088/2041-8205/761/1/L13 [page 254]

Kasten, F. and Young, A. T. (**1989** November). Revised optical air mass tables and approximation formula. *Appl. Opt.*, 28(22):4735–4738. URL http://dx.doi.org/10.1364/AO.28.004735 [page 91]

Kilkenny, D. (**1995** February). The effects of the Mt. Pinatubo eruption on the atmospheric extinction at SAAO, Sutherland. *The Observatory*, 115:25–28. URL http://adsabs.harvard.edu/abs/1995Obs...115...25K [page 102]

King, D. L. (**1985** September). Atmospheric Extinction at the Roque de los Muchachos Observatory, La Palma. *RGO/La Palma Technical Note 31*, Instituto de Astrofisica de Canarias, La Laguna, Tenerife, Canary Islands. URL http://www.ing.iac.es/Astronomy/observing/manuals/ps/tech_notes/tn031.pdf [page 100]

Koehler, R. L. and von Herzen, R. P. (**1986** January). A miniature deep sea temperature data recorder : design, construction, and use. *WHOI Technical Reports WHOI-86-3*, Woods Hole Oceanographic Institution. URL https://hdl.handle.net/1912/7828 [page 163]

Laken, B. A., Parviainen, H., García-Gil, A., et al. (**2016**). Thirty Years of Atmospheric Extinction from Telescopes of the North Atlantic Canary Archipelago. *Journal of Climate*, 29(1):227–240. URL http://dx.doi.org/10.1175/JCLI-D-14-00600.1 [page 99]

Laken, B. A., Parviainen, H., Pallé, E., and Shahbaz, T. (**2014**). Saharan mineral dust outbreaks observed over the North Atlantic island of La Palma in summertime between 1984 and 2012. *Quarterly Journal of the Royal Meteorological Society*, 140(680):1058–1068. ISSN 1477-870X. URL http://dx.doi.org/10.1002/qj.2170 [page 99]

Laskar, J. (**1986** March). Secular terms of classical planetary theories using the results of general theory. A&A, 157:59–70. URL http://adsabs.harvard.edu/abs/1986A&A...157...59L [page 293]

Lazrek, M., Pantel, A., Fossat, E., et al. (**1996** June). Is the Solar Core Rotating Faster of Slower Than the Envelope? Sol. Phys., 166:1–16. URL http://dx.doi.org/10.1007/BF00179353 [page 25]

Lefebvre, S., García, R. A., Jiménez-Reyes, S. J., Turck-Chièze, S., and Mathur, S. (**2008** November). Variations of the solar granulation motions with height using the GOLF/SoHO experiment. A&A, 490:1143–1149. URL http://dx.doi.org/10.1051/0004-6361:200810344 [page 6]

Leibacher, J. W. and Stein, R. F. (**1971**). A New Description of the Solar Five-Minute Oscillation. Astrophys. Lett., 7:191–192. URL http://adsabs.harvard.edu/abs/1971ApL.....7..191L [page 15]

Leighton, R. B., Noyes, R. W., and Simon, G. W. (**1962** March). Velocity Fields in the Solar Atmosphere. I. Preliminary Report. ApJ, 135:474. URL http://dx.doi.org/10.1086/147285 [page 12, 28, 77]

Lewis, D. J. (**1996** April). *A Dual Optical, Resonance-Scattering Spectrometer and the Probing of the Solar Atmosphere*. Ph.D. thesis, School of Physics and Space Research, University of Birmingham, UK [page 37]

Libbrecht, K. G. and Kaufman, J. M. (**1988** January). Frequencies of high-degree solar oscillations. ApJ, 324:1172–1183. URL http://dx.doi.org/10.1086/165972 [page 24]

Lieske, J. H., Lederle, T., Fricke, W., and Morando, B. (**1977** June). Expressions for the precession quantities based upon the IAU /1976/ system of astronomical constants. A&A, 58:1–16. URL http://adsabs.harvard.edu/abs/1977A&A....58....1L [page 293]

Linear (**2005**). LTC6240/LTC6241/LTC6242 Single/Dual/Quad 18MHz, Low Noise, Rail-to-Rail Output, CMOS Op Amps. *Datasheet 624012fe*, Linear Technology Corporation, 1630 McCarthy Blvd., Milpitas, CA, 95035-7417, USA. URL http://cds.linear.com/docs/en/datasheet/624012fe.pdf [page 196]

Lines, R. (**1998** February). A Visit to Las Campanas After the Lightning Strike that Destroyed our Step-Down Transformer. *BiSON Technical Report Series 71*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2005/ [page 70]

Lines, R. and Miller, B. A. (**1996** February). Scattered-Light Detector Manual. *BiSON Technical Report Series 47*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK [page 180]

Lingfors, D. and Volotinen, T. (**2013** July). Illumination performance and energy saving of a solar fiber optic lighting system. *Opt. Express*, 21(S4):A642–A655. URL http://dx.doi.org/10.1364/OE.21.00A642 [page 8]

Livingston, W. and Wallace, L. (**1985** February). Water vapor and Fe 5250.2. Sol. Phys., 95:251–252. URL http://dx.doi.org/10.1007/BF00152402 [page 31]

Lund, M. N., Chaplin, W. J., Hale, S. J., et al. (**2017**). Spatial incoherence of solar granulation: a global analysis using BiSON 2B data. *Monthly Notices of the Royal Astronomical Society*, 472(3):3256–3263. URL http://dx.doi.org/10.1093/mnras/stx2177 [page 5, 37]

Mangum, B. W. (**1983** December). Triple point of succinonitrile and its use in the calibration of thermistor thermometers. *Review of Scientific Instruments*, 54:1687–1692. URL http://dx.doi.org/10.1063/1.1137311 [page 163]

Mathur, S., Eff-Darwich, A., García, R. A., and Turck-Chièze, S. (**2008** June). Sensitivity of helioseismic gravity modes to the dynamics of the solar core. A&A, 484:517–522. URL http://dx.doi.org/10.1051/0004-6361:20078839 [page 5]

Mathur, S., Turck-Chièze, S., Couvidat, S., and García, R. A. (**2007** October). On the Characteristics of the Solar Gravity Mode Frequencies. ApJ, 668:594–602. URL http://dx.doi.org/10.1086/521187 [page 5]

McLeod, C. P. (**2002** July). Mark I Scaler System. *BiSON Technical Report Series 184*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2009/ [page 49]

Meeus, J. (**1991**). *Astronomical Algorithms*. Willmann-Bell. ISBN 9780943396354. URL https://books.google.co.uk/books?id=1smAQgAACAAJ [page 285]

Mészáros, L., Jaskó, A., Pál, A., and Csépány, G. (**2014** August). Accurate Telescope Mount Positioning with MEMS Accelerometers. PASP, 126:769–782. URL http://dx.doi.org/10.1086/677943 [page 235]

Miller, B. A. (**1997** November). The Trip to Las Campanas During the Big Snowstorm of 1997 August. *BiSON Technical Report Series 62*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2004/ [page 69]

Miller, B. A. (**1998a** February). BiSON Resonance-Scattering Spectrometers. *BiSON Technical Report Series 67*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK [page 48]

Miller, B. A. (**1998b** April). A Visit to Las Campanas in 1998 January. *BiSON Technical Report Series 74*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL `http://epapers.bham.ac.uk/2006/` [page 70]

Miller, B. A. (**2000** August). The Replacement of the Blind Motor and the Connection of the Station to the Internet in Narrabri in 2000 March. *BiSON Technical Report Series 138*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL `http://epapers.bham.ac.uk/2008/` [page 74]

Miller, B. A. (**2002** August). How Do You Get to the Zoo? *BiSON Technical Report Series 187*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL `http://epapers.bham.ac.uk/2010/` [page 59]

Miller, B. A. (**2003a** January). The Grand Opening of the Carnarvon Zoo in 2002 November. *BiSON Technical Report Series 193*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK [page 59]

Miller, B. A. (**2003b** February). RED v1.6 — A Cubic Fit. *BiSON Technical Report Series 198*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK [page 43]

Miller, B. A. (**2005** September). The Grand Opening of the Mount Wilson Zoo. *BiSON Technical Report Series 255*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL `http://epapers.bham.ac.uk/2027/` [page 81]

Miller, B. A. (**2011** May). The Installation of a Digital Autoguider in Las Campanas in 2011 March. *BiSON Technical Report Series 343*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL `http://epapers.bham.ac.uk/2048/` [page 69]

Miller, B. A. (**2012a** January). Stokes Vectors. *BiSON Technical Report Series, Number 350*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK [page 158]

Miller, B. A. (**2012b** January). Stokes Parameters and Mueller Matrices. *BiSON Technical Report Series, Number 351*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK [page 158]

Miller, B. A. (**2012c** January). Haidinger's Brush. *BiSON Technical Report Series, Number 352*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK [page 158]

Miller, B. A. (**2012d** January). The Optical Properties of Quartz. *BiSON Technical Report Series, Number 353*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK [page 158]

Miller, B. A. (**2012e** January). A Single-Element Pockels Cell. *BiSON Technical Report Series, Number 354*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK [page 158]

Miller, B. A. (**2012f** January). A Field-Widened Pockels Cell. *BiSON Technical Report Series, Number 355*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK [page 158, 159]

Miller, B. A. and New, R. (**1999** May). The Installation of Klaus at Mount Wilson in 1996 June. *BiSON Technical Report Series 106*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2007/ [page 78]

Mishra, B. (**1950** January). Lifetimes of Mercury and Potassium Atoms in Excited States. *Phys. Rev.*, 77:153–153. URL http://dx.doi.org/10.1103/PhysRev.77.153 [page 117]

Mitchell, A. and Zemansky, M. (**1961**). *Resonance Radiation and Excited Atoms*. The Cambridge series of physical chemistry. University Press. URL https://books.google.co.uk/books?id=Ic7vAAAAMAAJ [page 115, 116, 118, 120, 121]

Muñoz-Jaramillo, A., Dasi-Espuig, M., Balmaceda, L. A., and DeLuca, E. E. (**2013** April). Solar Cycle Propagation, Memory, and Prediction: Insights from a Century of Magnetic Proxies. ApJ, 767:L25. URL http://dx.doi.org/10.1088/2041-8205/767/2/L25 [page 254]

Nagel, L. W. (**1975**). *SPICE2: A Computer Program to Simulate Semiconductor Circuits*. Ph.D. thesis, EECS Department, University of California, Berkeley. URL http://www.eecs.berkeley.edu/Pubs/TechRpts/1975/9602.html [page 185, 271]

Nagel, L. W. and Pederson, D. (**1973** April). SPICE (Simulation Program with Integrated Circuit Emphasis). *Technical Report UCB/ERL M382*, EECS Department, University of California, Berkeley. URL http://www.eecs.berkeley.edu/Pubs/TechRpts/1973/22871.html [page 185, 271]

New, R. and Hale, S. J. (**2006** January). Carnarvon Trip Report — July/August 2005. *BiSON Technical Report Series 260*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2028/ [page 60]

New, R. and Isaak, G. R. (**2003** October). Work Carried Out at Narrabri in 2003 July. *BiSON Technical Report Series 213*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2013/ [page 74]

NOAA (**1976** October). U.S. Standard Atmosphere, 1976. *Technical Report NASA-TM-X-74335, NOAA-S/T-76-1562*, National Oceanic and Atmospheric Administration, Washington, DC, United States. URL https://ntrs.nasa.gov/search.jsp?R=19770009539 [page 142]

NREL (**2012**). G173-03 Standard Tables for Reference Solar Spectral Irradiances. URL http://rredc.nrel.gov/solar/spectra/am1.5/ [page 141]

Nyquist, H. (**1928** July). Thermal Agitation of Electric Charge in Conductors. *Phys. Rev.*, 32:110–113. URL http://dx.doi.org/10.1103/PhysRev.32.110 [page 186]

O'Gallagher, J. (**2008**). *Nonimaging Optics in Solar Energy*. Synthesis lectures on energy and the environment. Morgan & Claypool Publishers. ISBN 9781598293302. URL https://books.google.co.uk/books?id=hrKcNlXPu88C [page 150]

Ozyagcilar, T. (**2015a** November). Calibrating an eCompass in the Presence of Hard- and Soft-Iron Interference. *Application Note 4246*, Freescale Semiconductor, Inc. URL http://www.nxp.com/assets/documents/data/en/application-notes/AN4246.pdf [page 236, 237]

Ozyagcilar, T. (**2015b** November). Layout Recommendations for PCBs Using a Magnetometer Sensor. *Application Note 4247*, Freescale Semiconductor, Inc. URL http://www.nxp.com/assets/documents/data/en/application-notes/AN4247.pdf [page 236]

Ozyagcilar, T. (**2015c** November). Implementing a Tilt-Compensated eCompass using Accelerometer and Magnetometer Sensors. *Application Note 4248*, Freescale Semiconductor, Inc. URL http://www.nxp.com/assets/documents/data/en/application-notes/AN4248.pdf [page 236]

Ozyagcilar, T. (**2015d** November). Accuracy of Angle Estimation in eCompass and 3D Pointer Applications. *Application Note 4249*, Freescale Semiconductor, Inc. URL http://www.nxp.com/assets/documents/data/en/application-notes/AN4249.pdf [page 236]

Pedley, M. (**2013** March). Tilt Sensing Using a Three-Axis Accelerometer. *Application Note 3461*, Freescale Semiconductor, Inc. URL http://www.nxp.com/assets/documents/data/en/application-notes/AN3461.pdf [page 231, 233, 234]

Pedley, M. (**2015** October). High-Precision Calibration of a Three-Axis Accelerometer. *Application Note 4399*, Freescale Semiconductor, Inc. URL http://www.nxp.com/assets/documents/data/en/application-notes/AN4399.pdf [page 231, 232, 233]

Pesnell, W. D. (**2008** October). Predictions of Solar Cycle 24. *Solar Physics*, 252(1):209–220. ISSN 1573-093X. URL http://dx.doi.org/10.1007/s11207-008-9252-2 [page 5, 254]

Phillips, F. (**2011**). *An Introduction to Crystallography*. Read Books. ISBN 9781447417002. URL https://books.google.co.uk/books?id=zUhzkgEACAAJ [page 157]

Press, W., Teukolsky, S., Vetterling, W., and Flannery, B. (**2002**). *Numerical Recipes in C++: The Art of Scientific Computing*. Cambridge University Press. ISBN 9780521750332. URL https://books.google.co.uk/books?id=mSLhDt_XIUQC [page 218]

Rabl, A. (**1975** February). Optical and thermal properties of compound parabolic concentrators. *NASA STI/Recon Technical Report N*, 76:23,717. URL http://adsabs.harvard.edu/abs/1975STIN...7623717R [page 150]

Rabl, A. (**1976** January). Comparison of solar concentrators. *Solar Energy*, 18:93–111. URL http://dx.doi.org/10.1016/0038-092X(76)90043-8 [page 150]

Rabl, A. and Winston, R. (**1976** November). Ideal concentrators for finite sources and restricted exit angles. *Appl. Opt.*, 15(11):2880–2883. URL http://dx.doi.org/10.1364/AO.15.002880 [page 150]

Rhodes, Jr., E. J. and Ulrich, R. K. (**1977** March). The Use of p Mode Oscillations to Measure the Solar Rotation Rate. In *Bulletin of the American Astronomical Society*, volume 9 of BAAS, 336. URL http://adsabs.harvard.edu/abs/1977BAAS....9..336R [page 25]

Rhodes, Jr., E. J., Ulrich, R. K., and Simon, G. W. (**1977** December). Observations of nonradial p-mode oscillations on the sun. ApJ, 218:901–919. URL http://dx.doi.org/10.1086/155745 [page 16, 17]

Roca Cortés, T. and Pallé, P. L. (**2014** September). The Mark-I helioseismic experiment - I. Measurements of the solar gravitational redshift (1976-2013). MNRAS, 443:1837–1848. URL http://dx.doi.org/10.1093/mnras/stu1238 [page 48, 96]

Salabert, D., Fossat, E., Gelly, B., et al. (**2002a** August). IRIS^{++} database: Merging of IRIS + Mark-1 + LOWL. A&A, 390:717–723. URL http://dx.doi.org/10.1051/0004-6361:20020751 [page 23]

Salabert, D., Jiménez-Reyes, S. J., Fossat, E., Gelly, B., and Schmider, F. X. (**2002b** June). Variability of p-mode parameters in 11 years of IRIS^{++} data. In Wilson, A. (editor), *From Solar Min to Max: Half a Solar Cycle with SOHO*, volume 508 of *ESA Special Publication*, 95–98. URL http://adsabs.harvard.edu/abs/2002ESASP.508...95S [page 23]

Sallen, R. P. and Key, E. L. (**1955**). A practical method of designing RC active filters. *IRE Transactions on Circuit Theory*, 2(1):74–85. ISSN 0096-2007. URL http://dx.doi.org/10.1109/TCT.1955.6500159 [page 191]

Scherrer, P. H., Bogart, R. S., Bush, R. I., et al. (**1995** December). The Solar Oscillations Investigation - Michelson Doppler Imager. Sol. Phys., 162:129–188. URL http://dx.doi.org/10.1007/BF00733429 [page 21]

Scherrer, P. H. and Gough, D. O. (**2019** April). A Critical Evaluation of Recent Claims Concerning Solar Rotation. *arXiv e-prints*, arXiv:1904.02820. URL http://adsabs.harvard.edu/abs/2019arXiv190402820S [page 5]

Schou, J., Scherrer, P. H., Bush, R. I., et al. (**2012** January). Design and Ground Calibration of the Helioseismic and Magnetic Imager (HMI) Instrument on the Solar Dynamics Observatory (SDO). Sol. Phys., 275:229–259. URL http://dx.doi.org/10.1007/s11207-011-9842-2 [page 22]

Schunker, H., Schou, J., Gaulme, P., and Gizon, L. (**2018** June). Fragile Detection of Solar g-Modes by Fossat et al. Sol. Phys., 293:95. URL http://dx.doi.org/10.1007/s11207-018-1313-6 [page 5]

Schwarzschild, M. (**1948** January). On Noise Arising from the Solar Granulation. ApJ, 107:1. URL http://dx.doi.org/10.1086/144983 [page 11]

Seidel, L. (**1857**). *Ueber die Theorie der Fehler, mit welchen die durch optische Instrumente gesehenen Bilder behaftet sind, und über die mathematischen Bedingungen ihrer Aufhebung.* Abhandlungen der Naturwissenschaftlich-Technischen Commission bei der Königl. Bayerischen Akademie der Wissenschaften in München. Cotta. URL https://books.google.co.uk/books?id=f8rVPgAACAAJ [page 262]

Severnyi, A. B., Kotov, V. A., and Tsap, T. T. (**1976** January). Observations of solar pulsations. Nature, 259:87–89. URL http://dx.doi.org/10.1038/259087a0 [page 18]

SIGRIST Photometer (**2017**). VisGuard 2 In-situ Visibility Monitor. *Technical Report 14289E/4*, SIGRIST Photometer AG, Hofurlistrasse 1, 6373 Ennetbürgen, Switzerland. URL https://classic.photometer.com/svc/document.axd?id=14761&hl=E [page 96]

Siher, E. A., Benkhaldoun, Z., and Fossat, E. (**2002**). Diurnal Atmospheric Extinction Over Oukaïmeden Observatory using IRIS Database: 1989–1997. Comparison With The Other IRIS Sites. *Experimental Astronomy*, 13(3):159–170. ISSN 1572-9508. URL http://dx.doi.org/10.1023/A:1025535615069 [page 99, 103]

Siher, E. A., Ortolani, S., Sarazin, M. S., and Benkhaldoun, Z. (**2004** October). Correlation between TOMS aerosol index and astronomical extinction. In Oschmann, Jr., J. M. (editor), *Ground-based Telescopes*, volume 5489 of Proc. SPIE, 138–145. URL http://dx.doi.org/10.1117/12.549804 [page 99]

Stallman, R. M. (**1983** September). new Unix implementation. https://www.gnu.org/gnu/initial-announcement.html [page 59, 206]

Steinhart, J. S. and Hart, S. R. (**1968**). Calibration curves for thermistors. *Deep Sea Research and Oceanographic Abstracts*, 15:497–503. URL http://dx.doi.org/10.1016/0011-7471(68)90057-0 [page 163]

STELLA Robotic Observatory (**2017**). Environmental Status (Dust). URL http://stella.aip.de/stella/status/getdetail.php?typ=24 [page 96]

Stevens, W. and Rago, S. (**2013**). *Advanced Programming in the UNIX Environment*. Addison-Wesley Professional Computing Series. Pearson Education. ISBN 9780321638007. URL http://books.google.co.uk/books?id=kCTMFpEcIOwC [page 207]

Stokes, G. G. (**1851**). On the Composition and Resolution of Streams of Polarized Light from different Sources. *Transactions of the Cambridge Philosophical Society*, 9:399. URL http://adsabs.harvard.edu/abs/1851TCaPS...9..399S [page 158]

Taylor, T., Barnes, I., and Miller, B. A. (**2009** September). The Tiger Counters. *BiSON Technical Report Series 326*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK [page 63]

Texas Instruments (**1996** January). ADS1210/11 24-bit Analog-to-Digital Converter. *Datasheet SBAS034B*, Texas Instruments Inc., 12500 TI Boulevard Dallas, Texas, 75243, USA. URL http://www.ti.com/lit/ds/symlink/ads1211.pdf. Revised 2005 [page 204]

The GONG Project (**1985**). A Proposal to Study the Solar Interior by Measuring Global Oscillations with a World-Wide Network of Instruments. URL http://gong.nso.edu/science/papers/Report_Proposal_1.pdf [page 23]

Thompson, M. J. (**1990** January). A new inversion of solar rotational splitting data. Sol. Phys., 125:1–12. URL http://dx.doi.org/10.1007/BF00154773 [page 25]

Tomczyk, S., Schou, J., and Thompson, M. J. (**1995** July). Measurement of the Rotation Rate in the Deep Solar Interior. ApJ, 448:L57. URL http://dx.doi.org/10.1086/309598 [page 24]

Toptica (**2017** January). Tunable Diode Lasers. *Datasheet BR-199-68A*, Toptica Photonics AG, Lochhamer Schlag 19, D-82166 Graefelfing, Munich, Germany. URL https://www.toptica.com/fileadmin/Editors_English/11_brochures_datasheets/01_brochures/toptica_BR_Scientific_Lasers.pdf [page 124]

Torvalds, L. B. (**1991** October). Free minix-like kernel sources for 386-AT. https://groups.google.com/d/msg/comp.os.minix/4995SivOl9o/GwqLJlPSlCEJ [page 59, 206]

Turck-Chièze, S., Mathur, S., Ballot, J., et al. (**2008** October). The GOLF-NG prototype and the solar European perspective for cosmic vision 2015-2025. In *Journal of Physics Conference Series*, volume 118 of *Journal of Physics Conference Series*, 012044. URL http://dx.doi.org/10.1088/1742-6596/118/1/012044 [page 170, 172]

Ullah, I. and Shin, S.-Y. (**2012** September). Development of Optical Fiber-based Daylighting System with Uniform Illumination. *J. Opt. Soc. Korea*, 16(3):247–255. URL http://dx.doi.org/10.3807/JOSK.2012.16.3.247 [page 8]

Ulrich, R. K. (**1970** December). The Five-Minute Oscillations on the Solar Surface. ApJ, 162:993. URL http://dx.doi.org/10.1086/150731 [page 15]

Underhill, C. J. (**1993** May). *An Analysis Of Possible Line-Of-Sight Velocity Effects In The Power Spectra Of Solar Oscillations*. Master's thesis, School of Physics and Space Research, University of Birmingham, UK [page 34, 173]

Williams, H. K. (**2004** March). The Construction of the Narrabri Dome — 1992 August. *BiSON Technical Report Series 227*, High-Resolution Optical-Spectroscopy Group, University of Birmingham, UK. URL http://epapers.bham.ac.uk/2018/ [page 73]

Willson, R. C. (**1979** January). Active cavity radiometer type IV. Appl. Opt., 18:179–188. URL http://dx.doi.org/10.1364/AO.18.000179 [page 21, 28]

Winston, R. (**1974** October). Principles of solar concentrators of a novel design. *Solar Energy*, 16:89–95. URL http://dx.doi.org/10.1016/0038-092X(74)90004-8 [page 150]

Winston, R., Minano, J., Benitez, P., and Bortz, W. (**2005**). *Nonimaging Optics*. Optics and photonics. Elsevier Science. ISBN 9780080479736. URL https://books.google.co.uk/books?id=aa9iLJGbWJUC [page 150]

Wolff, C. L. (**1972** October). The Five-Minute Oscillations as Nonradial Pulsations of the Entire Sun. ApJ, 177:L87. URL http://dx.doi.org/10.1086/181058 [page 11, 17]

Wolff, C. L. (**1973** September). What is the Horizontal Scale of the 5-min Oscillations? Sol. Phys., 32:31–39. URL http://dx.doi.org/10.1007/BF00152727 [page 17]

Woodard, M. and Hudson, H. S. (**1983** October). Frequencies, amplitudes and linewidths of solar oscillations from total irradiance observations. Nature, 305:589–593. URL http://dx.doi.org/10.1038/305589a0 [page 21, 28]

Yeates, A. R., Nandy, D., and Mackay, D. H. (**2008** January). Exploring the Physical Basis of Solar Cycle Predictions: Flux Transport Dynamics and Persistence of Memory in Advection-versus Diffusion-dominated Solar Convection Zones. ApJ, 673:544–556. URL http://dx.doi.org/10.1086/524352 [page 254]

Index

C

D

E

F

G

H

N

O

P

Q

R

S

T

U

V

W

Y

Z