The hunt for low-frequency modes of oscillation of the Sun: application of statistical techniques and instrumentation

Anne-Marie Broomhall

A thesis submitted to
The University of Birmingham
for the degree of
DOCTOR OF PHILOSOPHY

School of Physics and Astronomy
The University of Birmingham
Submitted in November 2007

Abstract

The main focus of this thesis is to search for low-frequency solar acoustic and gravity modes of oscillation of the Sun. Low-frequency modes are sensitive to the structure of the solar core and so provide stringent tests for models of the internal structure of the Sun. In this thesis we develop and use statistical techniques that aim to detect low-frequency modes by searching for statistically significant features in frequency-power spectra. Frequency-amplitude spectra are compared for coincident prominent features. The amount of common noise shared by two data sets affects the probability that a coincident prominent feature is noise and so the statistical tests were non-trivial to derive. The optimum condition for comparing frequency-amplitude spectra is when no common noise is present. The power and source of the common noise is dependent on the data that are compared.

Statistical tests are used to search contemporaneous and non-contemporaneous observations made by the Birmingham Solar Oscillations Network (BiSON) and the Global Oscillations at Low Frequencies (GOLF) instrument. We also searched contemporaneous and non-contemporaneous sets of BiSON data. Many different combinations of BiSON data are compared to find contemporaneous BiSON timeseries that contain as little common noise as possible. Detection threshold levels are significantly reduced by searching for coincident prominent features, however, no new mode candidates are detected.

To investigate whether the gaps in the list of detected modes are expected, given the reduced threshold levels, predictions on the detectability of low-frequency modes are made, based on estimates of the modes' powers and damping times. The fraction of modes detected in Monte Carlo simulations is small at low frequencies. Significantly more detections would be made if the power signal to noise could be increased by a factor of 2, which could be achieved by lengthening timeseries and by reducing the amount of noise in the data.

The quality of solar observations is affected by instrument stability. We consider three effects that mean unresolved Doppler velocity observations are not homogenous across the solar disc; solar rotation, limb darkening and the position of a detector. The results indicate that observations made by BiSON instruments do not represent a uniform average across the solar disc. Different detectors are weighted towards different regions of the solar disc, and so observe different realizations of the solar noise. The calculated weightings are compared to BiSON observations. Agreement between the calculations and the observations is improved by altering various parameters, such as the observed size of the solar image. We find that the weighting varies with epoch and so can have consequences for the quality of BiSON data.

Acknowledgements

Many thanks go to my supervisors, Professor Yvonne Elsworth and Dr Bill Chaplin, for their much appreciated support, advise and enthusiasm throughout my time at Birmingham. They were inspirational and I feel honoured to have had the privilege of working with them. I would also like to thank my husband, Mark, for his encouragement over the last 3 years, for proof-reading this thesis and to listening to multiple repetitions of talks that I have given. My thanks also go to my parents, sister and grandparents who have always believed in my ability and have been very supportive and encouraging throughout my education.

I also offer my thanks to the members of the HiROS group, both past and present, for their help, their support and for making the last three years very enjoyable: John Allison, Ian Barnes, Guy Davies, Stephen Fletcher, Steven Hale, Barry Jackson, Bernard McCarty, Brek Miller, Mark Mobley, Sadie Oliver, Rosie Simionello, Neil Tarrant and Graham Verner

CONTENTS

LIST OF FIGURES

LIST OF TABLES

1. INTRODUCTION

1.1 Introduction to global helioseismology

The detection of solar oscillations has vastly increased our knowledge of the workings of the solar interior. Helioseismology is the study of the solar interior using observations of waves that propagate within the Sun. The development of helioseismology has dramatically improved our understanding of the internal properties of the Sun by testing the physical inputs that are used to model stellar interiors and by providing a detailed map of the Sun's structure and internal rotation.

Solar oscillations were first definitively observed by Leighton et al. (1962) when they detected oscillations of localized regions of the solar surface with periods of approximately 5 minutes. As the observations were made over a limited spatial extent it was initially thought that these 'five-minute modes' were a local phenomenon. However, it was not until Ulrich (1970) and Leibacher and Stein (1971) suggested that the five-minute modes might be acoustic modes, which are excited and trapped within the solar interior that a satisfactory theory to explain the observations was produced. The global nature of the modes was not revealed until 1979 when Claverie et al. detected low-degree modes with well-defined frequencies in unresolved solar data. It is now known that the Sun can support wave motions, which set up global resonant modes. Properties of the solar interior, such as temperature and composition, influence the periods of the oscillations. Accordingly, properties of the solar oscillations, such as their frequency, can yield valuable information about the internal structure of the Sun.

The first major discovery that was made because of the emergence of helioseismology was a determination of the depth of the solar convection zone. Observations by Deubner

(1975), Gough (1977) and Ulrich and Rhodes (1977) showed that the convection zone was actually 50% deeper than implied by conventional models. The depth of the convection zone can be determined by observing the discontinuity in the sound speed profile between the convection zone and the radiative core. Christensen-Dalsgaard et al. (1991) found the base of the convection zone to be 0.713 ± 0.003 solar radii from the centre of the Sun. This is in very good agreement with the more recent results of Basu and Antia (2004), who found the base of the convection zone is at a radius of 0.7133 ± 0.0005 solar radii. The value of the depth of the convection zone can only be constrained this tightly using helioseismology.

Helioseismology has also helped to uncover the solar rotation rate beneath the photosphere (see for example Rhodes and Ulrich 1977; Gough 1981; Thompson 1990). It has been determined that the differential rotation that is observed at the Sun's surface continues to the base of the convection zone. Below this the interior rotates like a solid body. This causes a region of intense shear at the base of the convection zone, known as the tachocline.

The composition of the solar interior, especially the abundance of helium, is important not only for solar modelling but also for modelling the solar system and the galaxy. The helium abundance cannot be inferred from spectroscopy as the necessary spectral lines are formed in the chromosphere and so are not in thermal equilibrium. Christensen-Dalsgaard and Gough (1980) used helioseismic data in comparison with the standard solar model to infer that the helium abundance is not vastly different from the model value assumed at the time. One recent estimate of the helium abundance that can be determined by helioseismology, which gives the proportion of the total mass of the Sun that is helium, is 0.2485 ± 0.0034 (Basu and Antia, 2004).

Until recently this value was in good agreement with the helium abundance determined by standard solar models. The values determined by the solar models are dependent on the heavy element abundance. However, new determinations of the heavy element abundance have recently been performed by Asplund et al. (2005). They deter-

mined the ratio of the mass fraction of heavy elements to hydrogen, Z/X, to be 0.0176. This can be compared to the value found previously by Grevesse and Sauval (1998) of $Z/X = 0.0229$. Recent studies, such as Basu and Antia (2004) and Zaatri et al. (2007), have compared seismically determined properties of the solar interior with those determined by solar models, constructed with both the old and the new abundances. The properties that were compared included the sound speed profile, the density profile, the helium abundance and the depth of the convection zone. The studies found that there was a large discrepancy between the seismic determinations and those found when the new abundances were used in solar models. Significantly better agreement was found when the old abundances were used. Basu et al. (2007) compared frequency properties of the modes themselves with those determined by solar models and concluded that lower abundance models have core structures that do not match the Sun's. Chaplin et al. (2007) obtained a seismic value for the metallicity of between 0.0252 to 0.03103, with errors of between ± 12 and 19%. These values are significantly larger than the metallicities derived by Asplund et al. (2005). The disparity between the seismically obtained results and those given by solar models, constructed with the new abundances, means that these new abundances have produced a real, and as yet unsolved problem.

Until recently one of the biggest unsolved conundrums was the solar neutrino problem. The neutron flux observed from the Sun was approximately $\frac{1}{3}$ of the value predicted by nuclear physics and solar models (Davis et al., 1968). However, now a possible solution has been found: the neutrinos change flavour as they travel from the Sun to Earth. Detectors are only sensitive to one flavour of neutrino but the neutrino itself could be any one of three different flavours by the time it reaches Earth, and so only $\frac{1}{3}$ of the neutrinos that reach Earth can be detected. Helioseismology provided very precise oscillation frequencies, which implied that the problem was with the neutrino physics and not the solar models that predicted the number of neutrinos that were expected to be observed on Earth (Elsworth et al., 1990).

Clearly observations of the solar oscillations have led to many important discoveries.

We now describe the solar oscillations in more detail.

1.2 Solar oscillations

Solar oscillations can be observed in 2 ways: by line-of-sight Doppler velocity measurements over the visible disc; or by measuring variations in the continuum intensity of radiation, which are caused by the compression of the radiating gas by the waves. Some instruments observe the integrated light from the whole of the solar disc as a single time-series. This is known as observing the 'Sun as a star'. Only wave motions with the largest horizontal scales are detectable in unresolved, Sun-as-a-star observations. Spatially resolved measurements of the Sun determine the simultaneous motion of the solar surface with a high spatial resolution. Using spatially resolved measurements, observations can be made of oscillations that are only visible if a portion of the solar disc is considered, as opposed to the whole disc. The amplitude of an individual mode observed in Sun-as-a-star data varies from $0.2\,\mathrm{m\,s^{-1}}$ for the strongest modes, to less than $0.01\,\mathrm{m\,s^{-1}}$ for the weaker modes. The total oscillatory signal that is observed is of the order of $1-2\,\mathrm{m\,s^{-1}}$. The modes with the largest amplitudes have frequencies of $\sim 3000\,\mu\mathrm{Hz}$, and periods of $\sim 5\,\mathrm{mins}$. Consequently these modes are known as the 'five-minute modes'. The solar oscillations that we observe in the Sun are acoustic waves that are sustained by a pressure-gradient restoring force and, consequently, are known as p modes.

Approximately the outer 30% of the Sun is known as the convection zone. It is generally believed that solar oscillations are stochastically excited by turbulent convective motions in the upper part of the convection zone, which generate acoustic noise. This noise is a broadband source, however, only waves that satisfy the appropriate resonant conditions are able to constructively interfere and become resonant modes. It is broadly accepted that the majority of mode excitation occurs in the downward plumes, where material, previously brought to the surface by the convection, has cooled and so flows back into the solar interior. Large excitation events, such as solar flares, may also contribute

to the modes' excitation. Once excited the modes are intrinsically damped by their interaction with the convection.

The simplest approximation of the Sun is a self-gravitating spherical ball of plasma that is in hydrostatic equilibrium. The horizontal structure of the modes can be expressed in terms of spherical harmonics, $Y_l^m(\theta, \phi)$, where θ is the colatitude and ϕ is the longitude. The degree of the mode, l, represents the total number of surface node circles on the sphere. The azimuthal order, m, gives the number of surface node circles that intersect with the poles. The azimuthal order takes values in the range $-l \leq m \leq l$. When $m = 0$ the angular degree is given by the number of surface nodes around the Sun's observed circumference. When $m = \pm l$ the angular degree gives the number of nodes around the Sun's equator. The structure of the eigenfunction in the radial direction is described by the radial order, n. The radial order is the number of nodes between the upper and the lower turning points of the mode. The upper turning point varies slightly with frequency but the lower turning point depends strongly on a mode's angular degree, l. At fixed l and m the mode frequency is a monotonic increasing function of n. If the Sun was wholly symmetric the frequency of a mode would be independent of m. However, any deviations from spherical symmetry break the degeneracy and cause the mode to become split in frequency, thereby introducing a frequency dependence on m. In the Sun the largest departure from the symmetrical state is due to the solar rotation, although the Sun's magnetic field is also important. Modes with the same l and n but different m are known as a multiplet. By convention, acoustic p modes have positive values of n (whereas solar gravity modes, which are introduced later in this section, have negative radial orders).

Only modes with $l \leq 3$ can be easily detected in Sun-as-a-star observations. Around the five-minute modes there is some evidence for $l = 4$ and 5 modes, but the power observed in these modes is very small. The perturbations of high-degree modes cancel when they are averaged over the solar disc as small regions of outward motion can be paired with small regions of inward motion. Therefore, high-degree modes cannot be observed in Sun-as-a-star data. Furthermore, spherical harmonics imply that unresolved

Figure 1.1: The propagation of sound waves through a cross-section of the solar interior. The path a wave takes through the interior is bent, by the increasing sound speed, until it reaches an inner turning point (indicated by the dotted circles). At the surface the waves are reflected by a rapid decrease in density. Image from Christensen-Dalsgaard (2003).

observations are only sensitive to modes where $l + m$ is even because of Earth's near-equatorial view of the Sun. On the other hand, spatially resolved observations have been able to detect modes with degrees of several thousands.

Acoustic waves travel from their excitation point near the surface to deeper in the Sun. As the pressure and sound speed in the Sun increase with depth, refraction causes the oscillations to follow a curved path which takes them back towards the surface (see Figure 1.1). Note that this is true for all waves except for those that propagate exactly vertically inward. The radius of the lower turning point increases with l and so only the low-l modes sample the deep interior. Furthermore, the $l = 0$ radial modes are the only modes that penetrate to the solar core. It is, therefore, very important to observe low-l modes as they allow us to infer conditions deep in the solar interior.

When the modes reach the photosphere they are reflected because of the sharp decrease in the density of the plasma. If the density scale height is smaller than the length scale of a mode the pressure changes required to make the wave cannot be maintained over a length of time that matches the wave period. Therefore, the compressions are smoothed as the gas can readjust on a timescale that is much shorter than the longer wave period and so the wave is reflected back into the interior. The maximum frequency at which modes are reflected back into the solar interior is known as the acoustic cut-off

frequency. For the Sun the acoustic cut-off frequency is approximately $5500\,\mu$Hz. As waves with frequencies higher than the acoustic cut-off are not reflected back into the solar interior they never become standing waves. Hence, there is a limit on the highest frequency modes that can be detected in solar observations. As the density scale height decreases rapidly with altitude in the outer regions of the Sun the depth at which modes with different periods are reflected varies. Low-frequency modes are reflected back into the solar interior at a lower altitude than high-frequency modes as the density scale height is greater at low altitudes. Therefore, the power in the low-frequency modes is more attenuated at the height in the photosphere at which the observations are made. This makes low-frequency modes harder to observe than high-frequency modes.

Another classification of mode that are thought to exist in the Sun are known as g-modes as their restoring force is gravity. These g modes rely on buoyancy and so they can only exist in regions of stable stratification. Therefore, they are trapped in the radiative interior, below the convection zone. As they are evanescent in the convection zone g modes are predicted to have very small amplitudes in the photosphere, where the observations are made. Consequently, at the time of writing, these modes have not been incontrovertibly observed. Solar g modes have very low frequencies and negative radial orders, n.

As mentioned already in this chapter solar observations can be either resolved, allowing the signal from high-l modes to be detected, or unresolved, enabling low-l modes to be observed. Furthermore, observations can be made of the Doppler velocity shifts or variations in the intensity of light can be measured. A brief description of some observational programs, designed to observe solar oscillations, will now be given.

1.3 Solar observational programs

To obtain a clean, well resolved frequency-power spectrum, in which modes can be observed, it is important to obtain long, preferably uninterrupted timeseries. One way of

doing this is to observe the Sun from space, outside of the Earth's atmosphere. The SOlar and Heliospheric Observatory (SOHO) is a satellite positioned at the L1 Lagrangian point, 1.6×10^6 km from Earth. The SOHO mission is jointly funded by the European Space Agency (ESA) and the National Aeronautics and Space Administration (NASA) and was launched in 1995. Onboard SOHO are many instruments, three of which were designed to observe solar oscillations. The Global Oscillations at Low Frequencies (GOLF) instrument makes unresolved Doppler velocity observations of the Sun using the neutral sodium doublet at 589.0 nm and 589.6 nm. The Michelson Doppler Imager (MDI) makes spatially resolved Doppler velocity observations of the Nickel I line. Finally, VIRGO (Variability of solar IRradiance and Gravity Oscillations) has two Sun PhotoMeters (SPM) that measure disc integrated intensities in three wavelength regions. Another part of VIRGO is the Luminosity Oscillations Imager (LOI), which measures solar intensities with limited spatial resolution.

Another way of obtaining continuous observations is to observe the Sun by means of a network of telescopes, that are strategically positioned around the globe in order to, potentially, obtain continuous observations of the Sun. This is the method employed by the Global Oscillations Network Group (GONG), which is a ground based network of 6 sites (Harvey et al., 1996). GONG makes resolved observations of the Sun and so can observe high-degree modes. The GONG instruments observe the Doppler velocity shifts in the Nickel I line at 676.8 nm. The Birmingham Solar Oscillations Network (BiSON) is a network of 6 sites that observe the Doppler shift in the potassium D1 line (Elsworth et al., 1995). BiSON makes Sun-as-a-star observations and so observes low-degree modes.

The majority of the work in this thesis uses BiSON data, although some GOLF data are used. The six BiSON stations are at Izaña, Tenerife; Carnarvon, Western Australia; Sutherland, South Africa; Las Campanas, Chile; Mount Wilson, California; and Narrabri, New South Wales. The data from the different sites are coherently combined, taking proper account of the relative quality of the data from each site. As mentioned it is important to create a timeseries with a high duty cycle. However, this requirement must

be balanced against the possibility that if the data were of poor quality it could drive up the overall level of noise by an amount greater than if the data were not included in the first place. Therefore, a 'figure of merit' is used to determine whether data should be included in the final timeseries or not (Chaplin et al., 1997). We will now briefly describe the basic principles behind the BiSON instrumentation.

1.4 BiSON observations

The stations in the BiSON network were established at different times using various instrument designs. However all of them are based on the same physical principles. We will now give a basic description of how the BiSON observations are made.

1.4.1 Choice of Fraunhofer line

The BiSON instruments use resonant scattering spectrometers to observe Doppler velocity shifts in the neutral potassium D1 spectral line, which has a wavelength of ~ 769.9 nm. We observe the solar potassium resonance line as it is narrow, deep and not blended with other lines, which are related to transitions in other elements. The solar potassium line has two components, however, the large separation in wavelength between the components (3.5nm) means they can be easily isolated. The D2 potassium component is blended with the telluric molecular oxygen line and so would confuse any solar observations. Thus the BiSON instruments observe the D1 potassium component. In comparison the sodium lines, which are observed by the GOLF instrument, are only separated by 0.6 nm. Furthermore one of the sodium components is blended by the atmospheric water vapour line. Obviously this is not a concern for the GOLF instrument as it is positioned above the Earth's atmosphere, however, it would be a problem for the ground-based observations. The solar potassium D1 line has steeper sides than both of the sodium lines, which means that observations of the potassium line are more sensitive to lineshifts. In fact Isaak et al. (1989) found that the potassium D1 line is ~ 2.3 times more sensitive to lineshifts than

Figure 1.2: A simple schematic representation of a resonant scattering spectrometer. Light from the Sun enters the instrument from the left. The vapour cell is positioned in a uniform longitudinal field of strength B.

the sodium D1 line. .

In a BiSON resonant scattering spectrometer potassium is heated to approximately 100°C, so that it becomes a vapour, and is enclosed in a small cell. Another reason for observing potassium is that atomic vapours are relatively easy to produce with all alkali metals such as potassium and sodium. However, a cell containing sodium must be heated to a higher temperature than a cell containing potassium to obtain the same optical depth.

A description of a resonance scattering spectrometer can be found in Brookes et al. (1978a). A brief but more up-to-date description of the BiSON instruments will now be given.

1.4.2 BiSON instrumentation

A simple schematic of a resonant scattering spectrometer is shown in Figure 1.2. Incident solar radiation passes through a coarse filter with a bandwidth of a few tens of nanometres. This is done to reduce the chances of overheating and the possibility of stray light being detected. The filtered radiation then passes through an interference filter, which has a much smaller passband, typically 1 nm centred on 769.898 nm. This narrow bandwidth removes any radiation from the D2 line at 766.4 nm. The light then enters a combination of a linear polarizer and a Pockels cell, which produce circularly polarized light that is allowed to enter the vapour cell. An alternating current is applied to the Pockels cell

which causes a rapid switching between left and right circularly polarized light. The handedness of the polarization is switched at a higher frequency than that associated with atmospheric fluctuations, thus minimizing the atmospheric noise observed. The light then passes into a vapour cell, which contains a heated potassium vapour and is enclosed in a longitudinal magnetic field.

Under the influence of an external magnetic field the single potassium absorption line is split into two and the separation between components is dependent on the strength of the field. These passbands lie at equal distances in wavelength from the unperturbed line centre. BiSON vapour cells are positioned within a uniform magnetic field of $\sim 0.18\,\mathrm{kG}$, which corresponds to a separation of $0.013\,\mathrm{nm}$. Each component is only sensitive to either left-handed or right-handed circularly polarized light. The right-handed circularly polarized component is shifted to a longer wavelength and so will now be referred to as the red instrumental component. The left-handed circularly polarized component is shifted to a shorter wavelength and so will be referred to as the blue instrumental component. As the light that enters the vapour cell is already circularly polarized, only one component of the split absorption line will be present. The handedness of the filtered light that is allowed to enter the vapour cell, therefore, determines which passband is observed.

The solar Fraunhofer line is broader than the potassium lines in the cell due to the Sun's rotation and because the Sun's turbulent atmosphere is significantly hotter than the vapour in the cell. Therefore the two shifted instrumental components are located on either side of the solar line. Recording the intensities of radiation scattered into each passband will give the relative solar flux at the wavelength of each passband on the solar resonance line. If the line-of-sight velocity between the cell and the Sun is zero the intensity observed by the two components will be the same (left panel of Figure 1.3). However, if the solar line is Doppler shifted the intensity of the scattered light in each component will differ (right panel of Figure 1.3). The scattered radiation is then detected using photodiodes that are orientated at 90° to the incident beam.

If the Doppler shifts in the wavelength of the solar line are sufficiently small the

Figure 1.3: The Fraunhofer line observed by a resonant scattering spectrometer. The black solid line represents the solar Fraunhofer line, the red dot-dashed line represents the red instrumental component, and the blue dashed line represents the blue instrumental component. In the left hand panel there is no line-of-sight velocity between the Sun and the instrument and so the red and blue components observe the same intensities. In the right-hand panel there is a line-of-sight velocity between the Sun and the instrument and so the intensities observed by the red and blue components are different.

portion of the solar line over which the left and right circularly polarized components observe can be considered to be linear. Assuming the line is symmetric the observed velocity, $v_{\rm obs}$, is given by

$$v_{\rm obs} \propto \frac{I_B - I_R}{I_B + I_R},$$ (1.1)

where I_B is the intensity observed in the blue wing of the solar line, which corresponds to the left circularly polarized light and I_R is the intensity observed in the red wing of the solar line, which corresponds to the right circularly polarized light. The constant of proportionality is obtained from a daily calibration, which makes use of the calculated velocity between the Sun and the Earth. The proportionality constant is known as the velocity sensitivity and is of the order of $\sim 3000\,{\rm m\,s^{-1}}$ per unit ratio. The observed velocity can be split into several different components, which will now be described.

1.5 Factors Contributing to the line of sight velocity

The velocity of the Sun's surface relative to a point on the Earth, $v_{\rm obs}$, can be broken up into several component parts and so can be written as (e.g. Brookes et al., 1976)

$$v_{\rm obs} = v_{\rm orb} + v_{\rm spin} + v_{\rm grs} + v_{\rm osc} + v_{\rm other},$$ (1.2)

where v_{orb} describes the line-of-sight velocity due to the Earth's orbit around the Sun, v_{spin} gives the line-of-sight velocity brought about because Earth spins on it's own axis, v_{grs} represents the gravitational redshift velocity and v_{osc} is the velocity due to the solar oscillations. Any line-of-sight velocities introduced by instrumental noise and atmospheric and solar effects are included in v_{other}. For all of these terms receding velocities are taken to be positive, as is the convention.

The velocity due to the oscillations, v_{osc}, is the cumulative velocity observed from a large number of modes. Over the entire Sun the superposition of high-order modes averages out and so only the lower-degree modes actually contribute to the observed velocity. The oscillation velocity signal only constitutes a small proportion of v_{obs}.

The gravitational redshift velocity, v_{grs} is not a true velocity signal but is indistinguishable from one. A photon of light detected by an instrument on Earth has to do work to escape the Sun's gravitational field before it can be observed. This manifests itself in a redshift of the light, and so by the time the photon is observed it has a longer wavelength than than when it left the Sun. The magnitude of the shift depends mostly on the solar mass and radius, but also on the wavelength of the emitted light. The offset observed by BiSON instruments corresponds to a velocity of $v_{\text{grs}} = +632\,\text{m}\,\text{s}^{-1}$. It should be noted that the size of the offset is wavelength dependent and so this value of the gravitational redshift velocity is specific to the observations of the solar potassium line, which has a wavelength of 769.9nm. The gravitational redshift velocity is always present and essentially constant.

The orbital velocity, v_{orb}, is due to the radial component of the Earth's orbit around the Sun. It arises primarily because of the slight eccentricity in the Earth's orbit. The magnitude of v_{orb} is also influenced by the effect on Earth's orbit of the moon and the other planets, especially Jupiter. The value of this component can be found using a standard Almanac. The BiSON group use look-up tables that are taken from the JPL ephemeris, which are available through the JPL Horizons system. The value of v_{orb} varies between $\sim +500\,\text{m}\,\text{s}^{-1}$ in April and $\sim -500\,\text{m}\,\text{s}^{-1}$ in October and is predicted to

a precision of less than $0.01\,\mathrm{m\,s^{-1}}$. The change in v_{orb} from day to day can be as large as $\sim 10\,\mathrm{m\,s^{-1}}$ but is generally much less than this.

The line-of-sight velocity component of Earth's spin, denoted by v_{spin}, varies throughout a day. In the morning an observer on Earth is approaching the Sun and so v_{spin} is negative; at local noon v_{spin} is zero; while in the afternoon an observer on Earth is moving away from the Sun and so v_{spin} is positive. The line-of-sight spin velocity is dependent on both the altitude and latitude of the observing site as well as the declination of the Sun (Brookes et al., 1976). It can be described by the equation (e.g. Elsworth et al., 1994)

$$v_{\mathrm{spin}} = \omega R_E(\lambda)\cos\lambda\cos\delta\sin\left[\frac{\pi}{12}\left(t - t_0\right)\right],\qquad(1.3)$$

where ω is the angular velocity of the Earth, $R_E(\lambda)$ is the observers distance from the Earth's centre, λ is the observers latitude, t_0 is the time of local noon in hours, t is the time of the observation in hours and δ is the declination of the Sun. The value of v_{spin} varies sinusoidally throughout a day. For a mid-latitude site the maximum change in velocity in any one day is about $800\,\mathrm{m\,s^{-1}}$.

The term v_{other} covers a wide range of effects. These include velocity noise introduced by solar, instrumental and atmospheric effects (see Chapter 2 for more details). Also, as discussed in Chapter 8, the Sun's rotation and limb darkening can alter the total line-of-sight velocity that is observed.

To gain information about the oscillation velocities many different analysis techniques are required. A brief description of some of the techniques that are used most frequently in this thesis will now be given.

1.6   Data analysis techniques

Much of the work for this thesis has involved analysing data in the frequency domain, rather than the time domain. The raw data are recorded in the form of a timeseries and so we begin this section by describing the technique used to convert the data between

time and frequency domains.

1.6.1   Fourier analysis

The Fourier transform defines a relationship between a signal in the time domain and its representation in the frequency domain. The standard Fourier transform, $F(\nu)$, of a function in the time domain, $f(t)$ is defined as

$$F(\nu) = \int_{-\infty}^{\infty} f(t)\mathrm{e}^{-i2\pi\nu t}\mathrm{d}t, \tag{1.4}$$

where ν is frequency in Hertz. Consider a signal, such as the one from the Sun, which is a superposition of many different oscillatory signals. The Fourier transform effectively allows this complicated signal to be separated out into individual frequency components, which appear as narrow peaks in the power spectrum. For discrete sets of data, such as those observed by the BiSON instruments, an approximate form of the Fourier transform is used, known as the fast Fourier transform (FFT). Using the FFT allows the transform to be calculated quickly and efficiently. An example of an FFT calculation is

$$G(\nu) = \frac{1}{N} \sum_{x=0}^{N-1} g(x)\mathrm{e}^{\frac{-i2\pi\nu x}{N}}, \tag{1.5}$$

where N is the number of points in the timeseries. $G(\nu)$ then provides the approximate Fourier transform up to the Nyquist frequency, ν_{nq}, which is defined as

$$\nu_{\mathrm{nq}} = \frac{1}{2\Delta t}, \tag{1.6}$$

where Δt is the cadence of the timeseries. BiSON data is normally stored with a cadence of $40\,\mathrm{s}$ and so the frequency domains extends to a frequency of $12,500\,\mu\mathrm{Hz}$. However, some of the data used later in this thesis (Chapter 5) has a cadence of $120\,\mathrm{s}$ and so the observations are only sensitive up to a frequency of $4167\,\mu\mathrm{Hz}$. $G(\nu)$ is the frequency-amplitude spectrum, which is complex. During this thesis we will also consider frequency-

Figure 1.4: Frequency-power spectrum obtained from 10yrs of BiSON data observed between 1st January 1996 and 1st January 2007.

power spectra, $P(\nu)$, which are given by

$$P(\nu) = |G(\nu)|^2. \tag{1.7}$$

The resolution in frequency, $\Delta\nu$, of a frequency-power or frequency-amplitude spectrum is determined by the length of the timeseries, T:

$$\Delta\nu = \frac{1}{T}. \tag{1.8}$$

Figure 1.4 shows an example of a frequency-power spectrum constructed from 10 yrs of BiSON data and so $\Delta\nu = 0.003\,\mu$Hz. The plotted spectrum has been restricted to the frequency range $1000 \leq \mu \leq 5000\,\mu$Hz to clearly demonstrate the individual resonant peaks that are visible in the spectrum. Notice that the strongest peaks are found in the five-minute range, at $\sim 3000\,\mu$Hz.

We now move on to describe how different sets of data can be compared, in the frequency domain, to obtain information about any data that may be common to both sets of observations.

1.6.2 Cross-spectral analysis

Cross-spectral analysis is often used when looking for low-frequency modes as it can improve the signal to noise of modes by combining two spectra and retaining any power that is detected by both instruments, while any contributions from independent noise sources are reduced. One of the main aims of this thesis is to develop the means by which contemporaneous sets of solar observations can be compared for coincident prominent features (Chapters 4, 5 and 7). To do this we must take proper account of the amount of noise that is common to the two sets of data. To determine the amount of common noise present we perform a cross-spectral analysis. Furthermore, in Chapter 7, detailed cross-spectral analyses are performed on different sets of BiSON data. These were done in an attempt to understand the common noise present, with the ultimate aim of producing two sets of BiSON data that contained a minimal amount of common noise. As cross-spectral analyses are a common theme throughout this thesis we now provide a brief introduction to the methodologies that are used.

The averaged cross-amplitude, $X_{A,B}$, of two sets of data, denoted by $a(t)$ and $b(t)$ in the time domain, is given by

$$X_{A,B} = \langle A(\nu) \cdot B^*(\nu) \rangle, \tag{1.9}$$

where $A(\nu)$ is the frequency-amplitude spectrum of timeseries $a(t)$ and $B^*(\nu)$ is the complex conjugate of the frequency-amplitude spectrum of timeseries $b(t)$. Any power that is present in both sets of data will be retained in the cross-amplitude spectrum. As we will see later in this thesis (Chapter 7) this power could be solar in origin, atmospheric noise and/or instrumental noise. The exact origin of the common noise depends on the data that are compared.

The cross-amplitude can then be normalized to give the coherency, which is a measure of the proportion of commonality shared by the two sets of data. The coherency is given

by

$$Y_{A,B} = \frac{X_{A,B}}{\sqrt{\langle A(\nu)A(\nu)^* \rangle \langle B(\nu)B(\nu)^* \rangle}}, \qquad (1.10)$$

and is zero when the frequency-amplitude spectra, $A(\nu)$ and $B(\nu)$, are completely independent and unity when the frequency-amplitude spectra are identical. Effectively, the coherency represents the ratio of the coherent power and the total power in the two spectra. In both equations the angled parentheses indicate that the mean value is determined. It is necessary to average when calculating the coherency. If no average is taken the numerator and denominator of the coherency will cancel and the determined coherency will be unity at all frequencies. This average can be calculated in two ways, which will now be described in turn:

- Method 1: Frequency-amplitude spectra can be found for short contiguous 2 hr segments, which are contained in the timeseries, $a(t)$ and $b(t)$. As the 2 hr timeseries are short it is important that they all have a 100% fill as otherwise severe noise will be introduced by the window function. The FFT of these short timeseries can be taken to produce low-resolution frequency-amplitude spectra. The average at each frequency bin in all of the 2 hr frequency-amplitude spectra can then be taken over all such spectra that are produced.

- Method 2: The FFT of both complete timeseries, $a(t)$ and $b(t)$, can be performed to produce two high-resolution frequency-amplitude spectra. The mean over a given range of frequencies can then be determined.

The phase spectrum, ϕ, of the two sets of data can be found using the cross-amplitude spectrum. We can see from equation 1.9 that the cross-amplitude spectrum has both real and imaginary parts. The phase, ϕ, is given by

$$\phi = \arctan \frac{\mathrm{Re}(X_{A,B})}{\mathrm{Im}(X_{A,B})} \qquad (1.11)$$

When comparing two sets of contemporaneous data the phase of any individual mode

should be approximately zero as the oscillations observed in two sets of contemporaneous data should be approximately in phase. However, when a time shift is introduced to make the data non-contemporaneous the observed oscillations will no longer be in phase, thus $\phi \neq 0$.

1.7 Motivation for searching for low-frequency p modes and g modes

The main aim of this thesis is to detect low-frequency p modes and g modes. Therefore a valid question to ask is why do we want to find evidence for these modes?

To date a multitude of p modes have been observed over a wide range of frequencies. However, no independently confirmed detections of low-degree p modes, with frequencies below $\sim 973\,\mu$Hz, have been made (e.g. García et al., 2001a; Chaplin et al., 2002a; Broomhall et al., 2007). As low-frequency p modes are expected to have very long lifetimes (see for example Houdek et al., 1998, 1999; Houdek et al., 2001; Chaplin et al., 2004a) their detection would allow their frequencies to be measured to very high accuracies and precisions. This in turn means that the physical quantities that can be inferred from the mode frequencies can also be determined more accurately and precisely. This is crucially important as the properties of low-l p modes are affected by conditions deep in the solar interior and so their frequencies act as a probe of these regions, allowing rigourous constraints to be placed on the structure of the solar core. However, as mentioned earlier, the depth beneath the photosphere of the upper turning point of a p mode decreases with mode frequency. Since the modes are evanescent in the photosphere the signal from low-frequency p modes is very weak. This results in a very small signal-to-noise ratio, meaning low-frequency p modes are very difficult to observe. Therefore, we have developed different analysis procedures in attempts to detect these modes.

As already mentioned, gravity (g) modes can only be sustained in regions that are stable against convection and so those of interest for studying the solar core are confined to the radiative interior. They would, therefore, provide a far more sensitive probe of the

deep solar interior than p modes. However, since the g modes are evanescent outside the radiative interior their signatures are significantly attenuated by the time they reach the photosphere. Furthermore, g modes have very low frequencies (below $\sim 450\,\mu$Hz) but the noise in a frequency power spectrum is largest at low frequencies. Hence g modes are predicted to have very small signal-to-noise ratios. It is, therefore, a major observational challenge to observe these modes. To date, there have been no independently confirmed observations of g modes.

Some of the low-l p-mode overtones, with frequencies below $\sim 450\,\mu$Hz, are thought to be mixed in character (Provost et al., 2000). This means that their sensitivity of the internal structure is similar to that of a g mode in the core, giving them excellent diagnostic potential. However, they behave like a p mode near the surface, which means, potentially, they should be easier to detect than g modes.

Therefore, the main focus of this thesis is the detection of low-frequency p modes and g modes. A brief description of the rest of this thesis will now be given.

1.8   Thesis synopsis

The main aim of this thesis is to develop and use techniques that search for low-frequency p modes and g modes. The techniques that are developed make use of the statistical structure of the noise in Doppler velocity observations. Therefore, they rely on a sound knowledge of the background noise. The structure of the observed noise can be split into three segments: solar noise, atmospheric noise and instrumental noise. Chapter 2 describes each of these factors in turn before considering the overall structure of the noise. The statistical structure of the background noise continuum is then used to develop various statistical tests, with the aim of detecting low-frequency p modes and g modes.

Chapter 3 develops statistical tests that can be used to search a single frequency-power spectrum. The simplest test searched for an individual statistically prominent spike in a frequency-power spectrum. However, other tests are developed that take advantage of

known properties of the modes to search for prominent structures in a frequency-power spectrum. To investigate the effectiveness of these statistical tests they were applied to data that were simulated to mimic Sun-as-a-star observations.

Chapter 4 moves on to consider how advantage can be taken of the numerous sets of Sun-as-as-star data that are available by comparing contemporaneous observations for coincidences. This involved deriving the probability of observing a prominent spike or structure of spikes in the same frequency bins in each of two frequency-amplitude spectra. When calculating the probability proper account must be taken of the amount of common noise present between the two sets of data, as this alters the probability that any detections are due to noise. The presence of this common noise means that the probability is non-trivial to derive.

In Chapter 5 the statistical tests derived in Chapters 3 and 4 are used to search BiSON and GOLF data for low-frequency p modes and g modes. The data are searched individually and they are compared for prominent coincidences. Both contemporaneous and non-contemporaneous data are searched.

Comparing the sets of data allows detection threshold levels to be reduced. However, there are still many gaps in the list of modes that are detected at low frequencies in Chapter 5. Chapter 6 investigates whether the number of recorded candidates is consistent with expectations based on estimates of the powers of the modes. These estimates were obtained by assuming the functional relationships observed at higher frequencies can be extrapolated down to low frequencies.

The positioning of the stations in the BiSON network means that there can be as many as 4 instruments making observations of the Sun at any one time. Also, many of these instruments take two sets of measurements. Chapter 7 aims to take advantage of the many sets of contemporaneous BiSON data that are available to search for low-frequency modes. When comparing two sets of data, detection threshold levels are lowest when the amount of common noise shared by the data is zero. The observed common noise could be solar, instrumental or atmospheric in origin. Therefore, in Chapter 7 we

strategically combine BiSON observations to create two sets of contemporaneous data that contain a relatively small amount of common noise. These data are then searched, using the statistical tests derived in Chapters 3 and 4, for low-frequency p modes and g modes.

In Chapter 7 we observe that the amount of common solar noise shared by two sets of data varies depending on the origin of the data that are being compared. This variation is due to Doppler imaging, which is investigated in Chapter 8. Unresolved measurements, such as those made by the BiSON instruments, are not homogenous across the solar disc and so the observed data do not represent a uniform average over the entire surface. Two effects that influence the inhomogeneity are due to the solar rotation and limb darkening. In Chapter 8 a further effect is considered that originates from the instrumentation itself. The intensity of light observed from a particular region on the solar disc is dependent on the distance between the detector and that region on the image of the solar disc formed in the instrument. As mentioned the majority of BiSON instruments have two detectors positioned on opposite sides of the image of the solar disc. The observations made by each detector are weighted towards differing regions of the disc. The effect of the line-of-sight velocity between the Sun and an instrument on the weighting of the solar disc is also investigated. We compare the calculated results with a velocity offset that can be observed in the BiSON data. This velocity offset varies with the line-of-sight velocity between the Sun and an observer. Therefore, the presence of the velocity offset in the BiSON data has consequences for the long-term stability of the instruments, and ultimately the quality of the BiSON data.

2. CHARACTERISTICS OF THE NOISE SPECTRUM

A frequency-power spectrum, constructed from Doppler velocity observations of the Sun, contains information about the solar modes, noise from the solar continuum, atmospheric noise and instrumental noise. At low frequencies the noise dominates a frequency-power spectrum making it extremely difficult to observe solar modes. This is because the power of the noise increases with decreasing frequency while the power of the modes decreases at low frequencies. The power of the total noise increases at low frequencies predominantly because the power in the solar noise increases at low frequencies (see Section 2.1). Characterizing the noise spectrum can, therefore, aid the search for low-frequency modes as it will enable simulations of the noise to be produced and may lead to the non-solar noise sources being reduced in power. Furthermore, the background solar noise can contain information about convection processes and wave propagation in the solar atmosphere. We will now describe the main factors that contribute to the total background noise. We start with the solar background continuum, which is the most dominant source of noise at low frequencies.

2.1 The solar background continuum

The noise from the solar continuum can be modelled, simply, as the superposition of noise due to the temporal evolution of active regions and the various types of granulation, namely normal granulation, mesogranulation and supergranulation (Harvey, 1985). The contribution to the velocity field of each component over time can be described by exponential decay functions where the characteristic times for active regions and the various scales of solar convection are different. The general form of the power contributed

Figure 2.1: An estimate of the solar background power spectrum for Sun-as-a-star Doppler velocity observations as predicted by the Harvey model (Harvey, 1985). The orange long-dashed curve, labelled 'SUM', is the sum of the spectra for the individual physical processes, which include granulation (solid black line, labelled 'G'), mesogranulation (short-dashed red line, labelled 'MG'), supergranulation (dot-dashed green line, labelled 'SG') and active regions (triple-dot-dashed blue line, labelled 'AR'). For reference the frequency of a day has been marked (purple, long-dashed line).

to the solar continuum, P_{cont}, by the various components can be modelled as

$$P_{\text{cont}} = \frac{2\sigma^2 \tau_d}{1 + (2\pi\nu\tau_d)^2}, \tag{2.1}$$

where τ_d is the characteristic decay timescale and σ is the root mean square velocity of the solar surface for full-disc observations (Elsworth et al., 1994). Equation 2.1 has been divided by a factor of 2 compared to the equation that appears in Harvey (1985). This was done to ensure the normalization is consistent as throughout this thesis the spectra are normalized by considering both the positive and negative frequencies, thus satisfying Parseval's theorem. In other words the data are normalized as double-sided spectra. An estimate of the power spectrum of each source of solar noise is depicted in Figure 2.1. Above $\sim 400\,\mu$Hz the normal granulation provides the greatest contribution to the total solar noise. Although the level of the noise produced by the normal granulation remains constant below $\sim 400\,\mu$Hz the amount of noise due to the supergranulation still increases and thus becomes the most dominant source of solar noise below $2 \times 10^{-4}\,$Hz.

The solar continuum noise is statistically random in nature and can be well represented with Gaussian statistics, taking a χ^2, two degrees of freedom (2 d.o.f.) distribution in the power spectrum. Between 400 and 1500μHz approximately half of the total observed noise in a frequency-amplitude spectrum is thought to be solar in origin. However, as yet the exact proportion of the total observed noise that can be attributed to the solar continuum is uncertain. Elsworth et al. (1994) performed a least squares fit to low-frequency regions of the cross-amplitude spectrum produced using data from the Izaña and Sutherland sites of BiSON to find the characteristic velocity (σ) and time constant (τ_d) for the granulation. Notice that it was possible to find the characteristic properties of the granulation alone as this dominates the solar noise frequency-power spectrum until very low frequencies. Elsworth et al. found the following values for the rms velocity and characteristic decay times:

$$\sigma = 0.319 \pm 0.009\,\mathrm{m\,s^{-1}}, \tag{2.2}$$

$$\tau_d = 260 \pm 23\,\mathrm{s}. \tag{2.3}$$

These can be compared to the values suggested by Harvey (1985), which were $\sigma = 0.7\,\mathrm{m\,s^{-1}}$ and $\tau_d = 372\,\mathrm{s}$. Elsworth et al. found that between 400 and $1500\,\mu$Hz the power of noise in BiSON data was a factor of 5 below the level predicted by Harvey's noise model. It should be noted, however, that the results found by Elsworth et al. are specific to the potassium line, which is used by BiSON instruments, whereas the values quoted by Harvey are for a more generalized 'strong photospheric line'. Harvey notes that the strength of the various components may vary with spectral line choice as different lines observe at different depths in the solar atmosphere.

The background solar noise is frequency dependent and will be very similar for observations made simultaneously with two different BiSON instruments. However, slight differences will occur in the solar noise observed by two instruments because of Doppler imaging. This will be dealt with in more detail in Chapter 8. Essentially, Doppler imag-

ing means that Sun-as-a-star data does not represent a uniform average over the whole solar disc (Brookes et al., 1978b). This weighting will vary from instrument to instrument as their position and motion relative to the Sun changes. This means that different instruments will observe at slightly different heights in the solar atmosphere and each instrument will be observing the granulation slightly differently. Both of these effects lead to slightly differing realizations of the solar noise in a frequency-power spectrum.

The realization of the granulation observed by an instrument depends on the depth in the photosphere at which the observations are made, which, in turn, depends on the wavelength observed by the instrument. A full calculation of the depth of formation of a solar Fraunhofer line is very complicated. However, an estimate can be made by assuming that all radiation with a given wavelength originates from the height at which unit optical depth, at that wavelength, is reached. Underhill and Speake (1996) estimated that from the centre to the wings of the potassium line, observed by BiSON instruments, the formation altitude varied by $\sim 400\,\mathrm{km}$. Each BiSON instrument analyses a certain range in wavelength on the the blue and red wings of the solar line profile. Given that the formation altitude varies with wavelength the blue wing will observe at a different depth to the red wing. Therefore, the granulation observed by the red-wing and blue-wing instrumental profiles will be slightly different. The depth of observation will also vary throughout the year as the observed solar line profile is Doppler shifted because of the line-of-sight velocity between the Sun and the Earth.

BiSON observations are made by examining Doppler shifts in the solar potassium line at 769.9 nm. However, other instruments, such as GOLF, make their observations using the neutral sodium doublet at 589.0 nm and 589.6 nm. The altitude at which the optical depth reaches unity is dependent on the wavelength at which the observations are made and the solar abundance of the element being observed. As sodium is far more abundant in the solar atmosphere than potassium (see for example Grevesse and Sauval, 1998; Grevesse et al., 2007) the solar sodium line will be formed at a higher altitude than the potassium line. Palle et al. (1992) estimated that the difference between the

formation heights of the sodium and potassium lines is $\sim 200\,\mathrm{km}$. However, a more recent estimate of the observation depths is given by Baudin et al. (2005), who calculated that the BiSON instruments observe $280\,\mathrm{km}$ above the photosphere, while GOLF observes $340\,\mathrm{km}$ above the photosphere. However, Baudin et al. comment that these values are only an estimation as it is very difficult to define the height of formation physically given the variation in formation altitude from the centre to the wings of the solar line. However, it is still true that the BiSON and GOLF data observe different depths in the solar photosphere and so the granulation noise detected by each instrument will not be identical.

Overall, some of the observed solar noise will be coherent between two instruments taking contemporaneous data but not all of it. The amount of coherent noise will depend on which instruments are being compared. For example, the Las Campanas A and B instruments that are separated physically on an astronomical mount by $\sim 60\,\mathrm{cm}$ will share a higher proportion of common solar noise than GOLF data will share with any of the BiSON instruments.

As we will see later the amount of common noise present between two sets of data affects the statistical mode detection techniques that are developed in this thesis. However, more fundamentally, the ability to detect low-frequency modes is limited by the amount of solar noise present in a set of data, as this is the dominant source of noise at low frequencies. One source of noise that will not be common between the ground-based BiSON instruments and the space-based GOLF instrument is atmospheric noise. Although not present in GOLF data, atmospheric noise does contribute towards the noise observed by the BiSON instruments and so the nature of the atmospheric noise will now be described.

2.2 Atmospheric noise

Atmospheric noise is evident in Earth-based solar observations in many different ways. The Sun's zenith angular distance will affect the amount of atmospheric extinction in a

non-random manner. If the Sun is close to the horizon the light from it will have to travel through more of the atmosphere than if the Sun is observed close to the zenith. This effect can be modelled and the modelled effect can be removed from the data. Furthermore, the optical path length through the atmosphere from different parts of the solar disc will vary, meaning that the amount of atmospheric extinction will vary across the solar disc.

Random atmospheric noise can also be caused by fluctuations in the atmosphere's transparency and seeing. Transparency fluctuations occur because of variations in the extinction coefficient of the atmosphere over the solar disc. These fluctuations can be caused by unevenly spread water vapour, dust or aerosols.

Transparency noise will not be correlated between two well-separated instruments. The separation required between two instruments so that the transparency fluctuations are not correlated is frequency dependent. Clette (1993) found that the frequency at which the noise from transparency fluctuations stops being coherent increases as the separation between the two instruments, L, decreases. More precisely, Clette found that the critical frequency, where the noise ceases to be coherent, is dependent on $L^{-\frac{1}{2}}$. Above $10,000\,\mu$Hz a separation of ~ 77m is required before the transparency fluctuations are no longer correlated. However, above $300\,\mu$Hz a separation of ~ 1400m is required. Therefore the noise from transparency fluctuations will be the same for the port and starboard observations made by a single BiSON instrument as they observe the same image of the Sun in the same vapour cell. The transparency noise will also be common to data taken by two instruments with a small separation. For example, the Las Campanas A and B instruments will observe coherent transparency noise as they are only separated by ~ 60 cm. However, the transparency noise will be independent in sets of data observed by the different BiSON sites. Obviously GOLF data does not suffer from atmospheric noise and so when the BiSON and GOLF data are compared, as they are in Chapter 5, this will be a source of independent noise that is only present in the BiSON data.

Another source of atmospheric noise is known as scintillation, which occurs because of turbulence in the Earth's atmosphere. The turbulence causes fluctuations in the refractive

index and airmass. The Kolmogorov model of turbulence (Kolmogorov, 1961) distributes the energy equally between atmospheric cells at all spatial scales. The power in this noise then varies as a function of frequency, ν. More precisely the power in the scintillation noise varies as $\nu^{-5/3}$ in a power spectrum. This model agrees well with the observed scintillation noise except at low frequencies (Young, 1974). Fossat et al. (1981) found that in the frequency range 150 to 3200 μHz the power in the airmass fluctuations varied as ν^{-1}.

Scintillation variations are coherent over smaller length scales than transparency noise. The noise created by scintillation is typically coherent over 5 to 15cm (Fried, 1965). This corresponds to an angular scale of a few seconds of arc, which is only a small fraction of the angular size of the solar disc. As a result Sun-as-a-star, unimaged, data are less susceptible to this noise source than resolved measurements. The only time this noise will be common to two sets of observations is when the port and starboard data from the same BiSON instrument are considered. It should be noted here that not all instruments are able to take port and starboard measurements. For example, in the BiSON network the Carnarvon A instrument and the Izaña instrument only have one detector. The scintillation noise will not be coherent between the data taken by Las Campanas A and B because they are separated by $\sim$ 60 cm, which is larger than the typical coherency length scale of the scintillation noise.

In Chapter 7 we compare port and starboard data from the same instrument to determine the level of noise that is common to both sets of data. In addition to solar and atmospheric noise some of the coherent data could be due to instrumental noise. Therefore, we go on to describe the instrumental noise that is present in BiSON data.

2.3 Instrumental noise

Instrumental noise comes from many sources including the finite counting rate of the detector, electronic noise and imperfections in the instrumental components. Hoyng

(1989) performed a detailed theoretical study on the sensitivity of resonance scattering detectors like the BiSON and GOLF instruments. He found that there are numerous causes of instrumental noise, each of which will now be described in turn.

1. **Guidance errors:** Guidance errors occur when light from the edge of the solar disc fails to pass through the spectrometer and into the vapour cell, which means that the instrument cannot observe the entire solar disc. Furthermore, the intensity of light transmitted by a Pockels cell is a function of the angle of incidence of light falling onto the cell. Therefore, if the direction of light received at the instrument, from the Sun, is not parallel to the alignment of the optical components of the instrument the intensity of light observed will be reduced. Any variation in the intensity of light observed can be misinterpreted as a Doppler velocity.

2. **Stability of the interference filter:** The interference filter is sensitive to changes in temperature. Any thermal fluctuations shift the wavelength response profile of the filter to higher/lower wavelengths. This in turn leads to a variation in the relative level of radiation let through the filter for each passband. This light can get scattered in the vapour cell causing a variation in the relative flux of scattered radiation, which can be misinterpreted as a Doppler velocity signal from the Sun.

3. **Efficacy of the Pockels cells:** A combination of a Pockels cell and a polariser can produce either clockwise or anti-clockwise circularly polarized light. However, the polarization of light will not be 100% efficient. If the voltage driving the Pockels cell is not constant a variable amount of radiation intended for the blue passband may be let through when the instrument is measuring the intensity of radiation scattered into the red passband and vice versa. This adds noise to the relative scattered intensity measured by the instrument.

4. **Temperature variations in the vapour cell:** Temperature variations in the vapour cell cause variations in the number density of atoms in the cell, which

changes the optical depth of the vapour. This can alter the amount of radiation scattered out of the cell to the detectors. If the temperature fluctuations occur at a faster rate than the passband switching interval, it could result in a spurious scaling in the relative intensities recorded for each sense of incident circularly polarized light. This, in turn, would lead to errors in the calculated solar velocity.

5. **Variations in the magnetic field enclosing the vapour cell:** An additional source of noise associated with the vapour cell is due to variations in the magnetic field enclosing the vapour cell. Variations in the magnetic field strength lead to variations in the sensitivity of the instrument as the wavelengths at which the red and blue passbands observe will be altered. However, the size of the variations and the level of noise introduced by the fluctuations will be small relative to the other noise sources described in this chapter.

6. **Temperature variations in the detectors:** The temperature of the detectors has to be carefully monitored as, in general, the detectors become noisier at higher temperatures. Temperature variations in the detectors generate variations in the dark current, which is the relatively small electric current that runs through a photosensitive device even when no photons are entering the device. The higher the temperature the larger the dark current that is produced. The magnitude of the dark current then becomes another variable in addition to the scattered intensities. It is possible that any temperature variations could alter the gain of the detectors. However, the magnitude of this variation due to small temperature fluctuations is likely to be minimal. Work is currently in progress within the BiSON group to investigate the effect of temperature changes on the gain of the detectors.

7. **Photon statistical noise:** The number of photons entering a BiSON resonant scattering spectrometer varies between $\sim 10^6 \, \text{s}^{-1}$, for the instrument at Izaña, and $\sim 10^{10} \, \text{s}^{-1}$ for the more recently installed instruments, such as Mount Wilson. Photon arrival times are governed by Poisson statistics and so there is an inherent variability

in this rate. However, a Poisson distribution approaches a normal distribution if the number of photons collected by the instrument is large. The standard deviation of the photon noise is given by the square root of the the number of photons collected, so that the signal-to-noise ratio of the photon noise is also given by the square root of the number of photons collected. Therefore, the low number of photons observed by the instrument at Izaña means that the shot noise is particularly evident in this data and in fact it dominates the power spectrum at high frequencies. However, as the counting rate in the newer instruments has been increased by a factor of 10,000 the shot noise is 100 times lower relative to the signal. The finite counting rate should produce a flat spectrum with random statistics (χ^2, 2 d.o.f. in the frequency-power domain). In a BiSON frequency-power spectrum this source of noise is not important at low frequencies as other noise sources, like the solar background, dominate. However, at high frequencies the solar noise has diminished significantly and so the photon noise contributes a larger portion of the total noise. García et al. (2005) reported that the GOLF detectors had lost a factor of ~ 4.5 in its counting rate over GOLF's lifetime, predominantly because of aging effects in the photomultiplier tubes. This means that the photon noise between 900 and 1500 μHz in GOLF data is only a factor of ~ 3.5 less than the solar background and so now constitutes an important part of the low-frequency noise power spectrum.

8. **Noise in amplifiers, voltage-to-frequency converters and other electronic components:** Scattered photons that are incident on the detectors produce a voltage that is amplified and passed to a voltage-to-frequency converter (VCF). The VCFs produce pulses that are counted by scalers. A scaler is an electronic circuit that records the aggregate of a specific number of signals that occur too rapidly to be recorded individually. The electronic components all have an inherent noise characteristic that is folded into the measured count rate.

Figure 2.2: The mean power spectrum of a set of BiSON data. The mean power was found for regions with widths of $100\,\mu$Hz. Each point is plotted at the minimum frequency of the $100\,\mu$Hz region. The size of the error bars were determined by the standard deviation of the power in the $100\,\mu$Hz regions. The peak in the mean level at $\sim 3100\,\mu$Hz is due to the presence of the five-minute modes, whose signals contain significant power. Below $\sim 1500\,\mu$Hz the mean power level increases with decreasing frequency. This is because the level of noise in the data increases and so the spectrum cannot be regarded as white at low frequencies.

We have now described the main features of the background noise individually and so we move on to discuss the more general features of the total noise background.

2.4 Total noise background

The central limit theorem states that any quantity produced by many independent variables will have a Gaussian distribution regardless of the distribution of the original variables, provided the independent variables have a finite variance. For example, a Lorentzian function does not have a finite variance and so if one of the independent variables had a Lorentzian distribution the central limit theorem can not be applied. However, that is not thought to be the case for any of the noise sources present in the background continuum. Therefore, although not all of the individual components that contribute to the noise background have a Gaussian distribution, because there are many independent contributing factors the resulting combination will tend towards a normal distribution.

When considered in its entirety a solar power spectrum is regarded as 'pink' because the mean level of the data increases as the frequency decreases. Figure 2.2 shows how the

mean power of noise in 8.5yrs of BiSON data varies with frequency. The bump in the mean level that peaks at $\sim 3100\,\mu$Hz is caused by the large signal-to-noise (S/N) ratio of the five-minute modes. If a small range of frequencies are considered the total background noise from solar observations has a Gaussian distribution in the frequency-amplitude domain. Therefore its power has a negative exponential distribution (χ^2 with 2 d.o.f.) in the frequency-power domain and can be described as flat and 'white'. Figure 2.3 shows the distribution of the power in different frequency ranges in a frequency-power spectrum of BiSON data. Each frequency range has a width of $100\,\mu$Hz. The data used was ~ 8.5 yrs in length, with a cadence of 120s, and will be used later when searching for low-frequency modes (Chapter 5). The distribution of the noise in the $500 - 600\,\mu$Hz, $1000 - 1100\,\mu$Hz and $1500 - 1600\,\mu$Hz frequency ranges are well represented by the linear fit until high powers. The deviation at high powers is particularly noticeable in the $1500 - 1600\,\mu$Hz frequency range where the departure from the linear fit is likely to be caused by the presence of the modes, which have reasonably large powers compared to the noise level in this frequency range. In general the deviations from the linear best fit line at high powers are due to the occasional randomly large spike in the frequency-power spectrum. The probability of observing the most prominent spike seen in the $500 - 600\,\mu$Hz range is ~ 0.03. This implies that this spike does not represent a large departure from the distribution. Similarly the probability of observing the most prominent spike that can be seen in the frequency range $1000 - 1100\,\mu$Hz is ~ 0.07. The probability of encountering the most prominent spike that is observed in the frequency range $1500 - 1600\,\mu$Hz is ~ 0.01. However, it should be remembered that in this frequency range the modes are reasonably prominent. Clearly in the frequency range $0 - 100\,\mu$Hz the noise cannot be regarded as 'white' even over this short frequency range. This is because the level of noise increases rapidly at low frequencies.

If we know the statistics of the solar noise background we can determine the probability that a prominent feature is due to noise. This allows power and amplitude spectra to be searched for prominent features that are statistically unlikely to be part of the noise
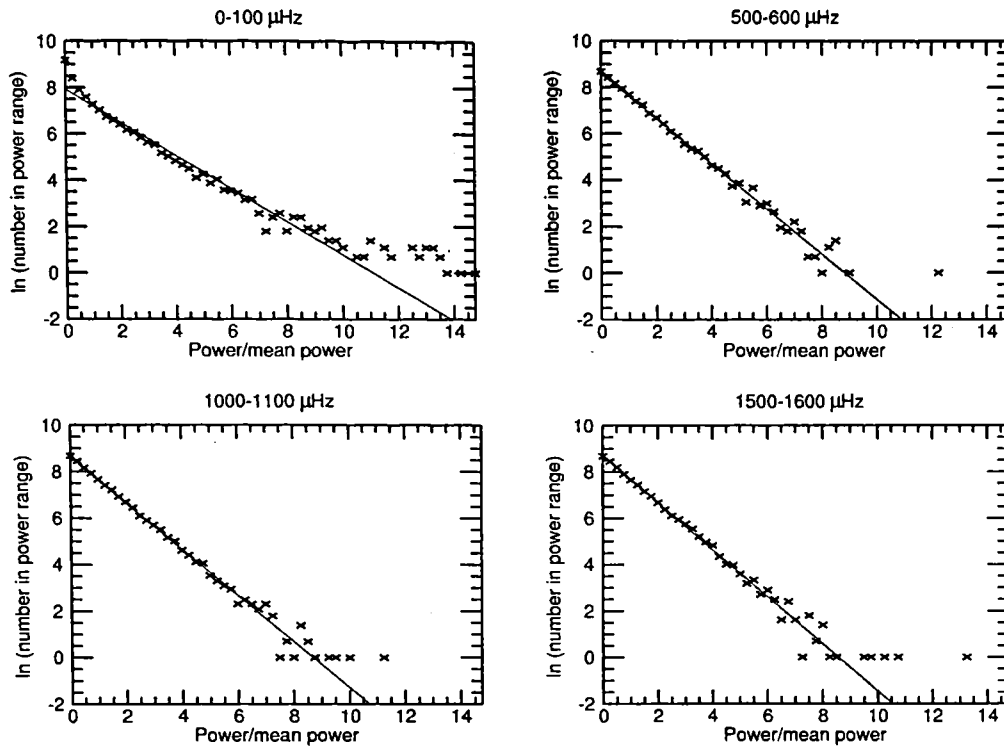
Figure 2.3: Distribution of the power in different frequency ranges of a power spectrum of BiSON data. The frequency ranges are all $100\,\mu$Hz in width. The black crosses represent the observed data, while the red solid line shows a linear fit to the distribution.

background. It is possible that a prominent feature provides evidence of a mode. We will now go on to describe how the distribution of the noise in frequency-amplitude and -power spectra can be used to search for prominent features. It is, therefore, vital to have a good understanding of the noise present in solar observations. Some of the statistical mode detection techniques described in this thesis involve comparing data from different instruments. When deriving these statistics proper account must be taken of the amount of noise that is common to the two sets of data. The amount of common noise varies depending on which instruments have measured the data. Searches for low-frequency modes are ultimately limited by the amount of solar noise present in the data as the level of solar noise increases in power at low frequencies while the power of the modes decreases.

3. DETECTION OF PROMINENT SPIKES IN ONE SPECTRUM

The signal-to-noise ratio (S/N) of low-frequency modes is small making them very difficult to observe. The depth beneath the photosphere at which the top of a mode's acoustic cavity occurs increases as the frequency of the mode decreases. Since observations of the modes are made in the photosphere, where the modes are evanescent, the observable signal is weaker for low-frequency modes than for high-frequency modes. Low-frequency modes are, therefore, observed with a power/amplitude that may not be much greater than the background level. Hence methods have been developed for revealing signatures that might be modes by assessing the abundance of statistically significant spikes that appear as part of the background noise. Using these methods we can identify prominent features that are statistically very unlikely to be part of the background noise and so can be considered as possible mode candidates.

3.1 The detection of individual prominent spikes in a pure noise spectrum

Consider a frequency-power spectrum created by taking the Fast Fourier Transform (FFT) of solar velocity data. As we have shown in Chapter 2 over a small range in frequency ($\sim 100 \, \mu$Hz) the background of this spectrum can be described as flat 'white' noise that has a χ^2 2 degrees of freedom (d.o.f.) distribution. It should be noted that this description is not accurate at very low frequencies (below $\sim 100 \, \mu$Hz). A spike is defined as the power that is contained in one frequency bin. Consider a spike at a frequency, ν, with a power, ξ_ν. Let the mean background level over a certain frequency range that

contains ξ_ν be given by $\langle \xi_\nu \rangle$. The relative height of the spike is then defined as

$$s_\nu = \frac{\xi_\nu}{\langle \xi_\nu \rangle}. \tag{3.1}$$

The probability of observing a spike with a relative height greater than or equal to s_ν in any given bin in a spectrum of noise that has a χ^2, 2 d.o.f. distribution is given by

$$p(s_\nu) = \exp(-s_\nu). \tag{3.2}$$

The probability of observing at least r spikes with a relative height greater than or equal to s_ν across a range of N bins is formally written as

$$P_s(s_\nu) = P[r; p(s_\nu), N], \tag{3.3}$$

where the $p(s_\nu)$ is defined as above and r is the minimum number of spikes that are required to be greater than or equal to s_ν in N bins. The probability of observing at least one spike ($r = 1$) with a relative height of s_ν or greater across N bins is equivalent to 1 minus the probability of not observing a spike greater than s_ν across a range of N bins and so can be written as

$$P_s(s_\nu) = P[1; p(s_\nu), N] = 1 - [1 - p(s_\nu)]^N. \tag{3.4}$$

An alternative way of writing this expression, which becomes more useful when $r > 1$, is in terms of the cumulative binomial (Bernoulli) distribution:

$$P_s(s_\nu) = P[r'; p(s_\nu), N] = \sum_{r=r'}^{N} p(s_\nu)^r [1 - p(s_\nu)]^{N-r} \frac{N!}{r!(N-r)!}. \tag{3.5}$$

For a spectrum of length T the number of bins, N, in a frequency range, Δ_p, is given by $N = \Delta_p T$. In Chaplin et al. (2002a), it was assumed that $p(s_\nu)$ is small, allowing both

Figure 3.1: The distribution of a pure noise frequency-power spectrum created by taking the negative natural logarithm of a uniform distribution. The straight line indicates that this produces a good χ^2 distribution with two degrees of freedom until high values of the normalized power.

equation 3.4 and equation 3.5 to be approximated as

$$P_s(s_\nu) \approx N p(s_\nu) \approx \Delta_p T \exp(-s_\nu).$$ (3.6)

Taking the natural logarithm of this expression allows the relative height for a given probability and frequency range to be calculated:

$$s_\nu = \ln \Delta_p T - \ln P_s(s_\nu).$$ (3.7)

This equation can then be used to set threshold heights in the spectrum for given probabilities and frequency ranges.

Pure noise frequency-power spectra were created to test the statistics outlined above in a set of simple Monte Carlo simulations. Anderson et al. (1990) suggest that a χ^2 2.d.o.f. distribution can be created by taking the negative natural logarithm of a set of uniformly distributed random numbers. Therefore 7 million uniformly distributed numbers were generated to test the distribution produced by this method. Once the numbers had been created it was possible to find the relative height, as defined by equation 3.1, of each number in the distribution that was created. The number of values between 2 different relative heights in the entire set of 7 million points was then determined. A graph of relative height against the natural logarithm of the number of occurrences in each range of relative heights should give a straight line if a χ^2 2d.o.f distribution has

Table 3.1: The average number of frequency ranges (Δ_p) containing at least 1 false detection in the whole spectrum.

Probability (P_s)	0.005	0.010	0.050	0.100
$\Delta_p(\mu\mathbf{Hz})$				
90	0.698	1.394	6.755	13.122
95	0.659	1.318	6.412	12.449
100	0.625	1.243	6.076	11.806
105	0.599	1.198	5.840	11.341
110	0.570	1.139	5.549	10.769

been produced. Figure 3.1 indicates a good χ^2 2d.o.f. distribution is produced until high values of the relative height. The gaps in the distribution mean that there are no spikes generated for certain high relative heights. Given a χ^2 2d.o.f. distribution the probability of observing a spike with a relative height between 12.825 and 12.900 is very small. Therefore statistically only 1 or 2 numbers out of the 7 million simulated points should lie in this range of relative heights. However statistical fluctuations could mean that no numbers are generated between the relative heights 12.825 and 12.900. The size and number of gaps in the distribution could be reduced by producing a larger sample of numbers. However, for relative powers greater than approximately thirteen the distribution is very hard to test without making the size of the sample significantly larger. Alternatively the relative height bands could be widened to increase the probability of producing a number in that range. However, the widths would have to be increased significantly before any effect would be seen and this would disrupt the smoothness of the straight line that is observed at lower normalized heights. Therefore, it is reasonable to believe that taking the negative natural logarithm of a uniform distribution does produce the required χ^2 2d.o.f. distribution and so this method will be used in the following simulations.

The method described above for producing a χ^2 2d.o.f. distribution was used to simulate 1000 independent spectra containing 7,076,160 random numbers that were assigned a cadence of 40 seconds. This number of points and value of the cadence were chosen

to be consistent with the length of the timeseries used in Chaplin et al. (2002a). We adopted values of $\Delta_p \approx 100\mu$Hz as this value will be used later in the analysis of real BiSON data. The number of frequency ranges, of width Δ_p, that contained at least one spike with a relative height greater than or equal to s_ν in the whole spectrum was counted for 1000 simulated spectra and the average number of frequency ranges that contained a detection in each individual spectrum was calculated. In later sections, a spike will be considered as a mode candidate if the probability that it is part of the background noise (P_s) is less than 0.1, and so the same threshold has been applied here. Additionally, in the later analysis of real BiSON and GOLF data thresholds will also be set at $P_s < 0.05$ and $P_s < 0.01$ and so the number of detections due to noise at these probabilities has also been investigated. The results are shown in Table 3.1 and Figure 3.2. A false detection can be defined as a spike with a relative height of at least s_ν that is due only to noise. In actual solar data a false detection could be mistakenly considered as a candidate for a mode. From Figure 3.2 we can see that the number of frequency ranges in the entire spectrum containing at least one false detection decreases as the probability of finding a false detection, P_s, decreases but increases as the frequency range, Δ_p, decreases. The number of frequency ranges of length Δ_p in the whole spectrum is

$$L = \frac{1}{2\Delta_p \text{cad}},\quad (3.8)$$

where cad is the cadence, taken to be 40s in these simulations. Therefore, if the cadence is constant, as Δ_p decreases, L increases. The number of frequency ranges that will contain at least one spike with a relative height $\geq s_\nu$ in the whole spectrum can then be predicted using the equation

$$\text{Number of } \Delta_p \text{ containing at least 1 false detection} = P_s L.\quad (3.9)$$

Hence the number of Δ_p containing at least one false detection is directly proportional to the number of Δ_p that are searched. Therefore, for a fixed value of P_s, as Δ_p decreases

Figure 3.2: The number of Δ_p in a simulated spectrum (7,076,160 points) that contained at least one spike with a relative height greater than s_ν. The simulations were performed for different values of Δ_p, ranging from 90 μHz to 110 μHz and these are represented by the different lines (see legend).

and L increases the number of Δ_p containing at least one false detection increases, thus explaining the observed results. The decision over the width of Δ_p is very important. As Δ_p increases the probability of detecting a prominent spike by chance in each Δ_p increases, therefore the size of Δ_p must be limited. However, as Δ_p decreases the number of frequency ranges of width Δ_p in the spectrum increases. Therefore the probability that one of these ranges will contain at least one prominent spike also increases and so it is not desirable to have a small value of Δ_p. Hence, when deciding on the optimum value of Δ_p a balance between these two requirements must be sought.

The total number of spikes in the entire spectrum with a relative height above the threshold level can be found by taking $N = \Delta_p T = 1$ bin. Equation 3.9 implies the proportion of Δ_p in the whole spectrum that contain at least 1 false detection is simply

$$\text{Proportion of } \Delta_p \text{ containing at least 1 false detection} = \frac{P_s L}{L} = P_s, \qquad (3.10)$$

as expected.

The predicted results agree moderately well with the simulated results as the difference

Table 3.2: The original probability is the value of P_s used in equation 3.7 to set the threshold, s_ν, for the simulations. The recalculated probability was found by substituting this value of s_ν into equation 3.4. These probabilities can be compared to the probability of observing at least one spike with a power greater than s_ν in 100μHz found by the simulations.

Original Probability	0.00500	0.01000	0.05000	0.10000
Recalculated Probability	0.00499	0.00995	0.04877	0.09516
Probability from Simulations	0.00500	0.00994	0.04861	0.09445

between the predicted and the simulated results is less than 6% of the predicted result for all values of P_s and Δ_p. Interestingly, for $P_s = 0.005$ and $P_s = 0.01$ the majority of the observed number of detections are higher than the predicted number. This is likely to be because only a small number of detections are expected, therefore the observed number is very susceptible to the effect of statistical fluctuations. However, for $P_s = 0.05$ and $P_s = 0.1$ the observed numbers of detections are consistently less than the predicted number. The numbers of expected detections are larger for these values of P_s and so the results will be less susceptible to statistical fluctuations. This could, therefore, imply inaccuracies in the equations that predict the expected number of detections (equations 3.6 to 3.9). Hence we go on to investigate the discrepancy between the predicted and simulated results further.

3.2   Improving the predictions

If the threshold relative height (s_ν) calculated from equation 3.7 is transferred back into equation 3.4 a slightly different probability (P_s) is found. Table 3.2 shows this effect becomes increasingly important as P_s increases. This recalculated probability is significantly closer to the probability revealed by the results of the simulations, i.e. the average number of Δ_p found containing at least one prominent spike (given in Table 3.1) divided by the total number of Δ_p searched (125 here). This implies that the value calculated for the threshold height for a given P_s is not accurate.

Since the predictions made in Section 3.1 rely heavily on the assumption that $P_s(s_\nu) \approx$

$Np(s_\nu)$ it is reasonable to investigate whether the calculated predictions will be more accurate if this approximation is either improved upon or, alternatively, dropped completely. The binomial expansion of equation 3.4 (and also the expansion for $r = 1$ in equation 3.5) gives

$$P_s(s_\nu) = 1 - \left[1 - Np(s_\nu) + \frac{N(N-1)}{2}p(s_\nu)^2 - \frac{N(N-1)(N-2)}{6}p(s_\nu)^3 \right.$$
$$\left. + \frac{N(N-1)(N-2)(N-3)}{24}p(s_\nu)^4 - ... \right]. \tag{3.11}$$

It must be noted that this particular expansion is only valid for the case when $r = 1$ as for higher values of r not all of the terms in the above equation will be present. Equation 3.11 can be simplified to give

$$P_s(s_\nu) \approx \sum_{q=1}^{M} \frac{(-1)^{q-1}N!}{q!(N-q)!}p^q, \tag{3.12}$$

where q is simply the integer over which the summation is made and M can be taken to be any value between 1 and N depending on how many terms of the expansion are required. For example, the approximation used by Chaplin et al. (2002a) only keeps the first term of equation 3.12 (approximation 1) and so $M = 1$. However, by including the next term in the series ($M = 2$, approximation 2) the equations for P_s and s_ν respectively become

$$P_s \approx Np(s_\nu) - \frac{N(N-1)}{2}p(s_\nu)^2, \tag{3.13}$$

$$s_\nu \approx -\ln\left\{ \frac{1}{N-1} \pm \left[\frac{N - 2P_s(N-1)}{N(N-1)^2} \right]^{\frac{1}{2}} \right\}. \tag{3.14}$$

In equation 3.14 we take the negative of the square root and not the positive square root. This is because, although both are solutions to equation 3.13, only the negative option produces the correct value of P when substituted into equations 3.2 and 3.4. As can be seen in Figure 3.3, for a given relative height the inclusion of the second term in equation 3.11 noticeably alters the probability. However, as more terms are added the

Figure 3.3: The probability of observing at least one spike in a frequency range Δ_p with a height greater than s_ν, as given by equation 3.12, is plotted for different values of M (where M is the number of terms in the expansion).

series converges to a value only slightly greater than the probability when just two terms are included. Approximation 2 should therefore predict the number of false detections expected more accurately than approximation 1.

Simulations the same as those performed in Section 3.1 were repeated here, except equation 3.14 was used to determine the threshold height. When approximation 2 is used to calculate the threshold height, s_ν, for certain values of P_s and Δ_p the difference between the simulated and predicted number of frequency ranges containing at least 1 false detection is less than when approximation 1 is used. For probabilities of 0.005 and 0.01 the observed numbers of detections are still greater than the predicted numbers of detections for the majority of Δ_p. Again this is likely to be because only a small number of detections are expected. Also the validity of approximation 1 improves as the probability that a single spike is above a threshold level, p, decreases. However, when $P_s = 0.05$ and 0.10 the observed numbers of detections are now greater than the predicted numbers for only half of the different Δ_p examined. The threshold height calculated using approximation 2 represents the data better than the original threshold height (see table 3.2) and so will be used in all that follows.

Until now we have been describing the statistics of observing a single prominent spike.

However, solar p modes are damped and so the width of a mode can be resolved in a frequency-power spectrum. Furthermore the rotation of the Sun means that some modes exhibit a multiplet structure in the frequency-power spectrum. It is possible to exploit these properties in order to develop more sophisticated detection techniques. We begin by taking advantage of the width in frequency of a mode that is observed in a frequency-power spectrum.

3.3 The detection of peaks in a single noise spectrum

Solar p modes are excited stochastically by turbulence in the outer regions of the convection zone and damped intrinsically by the convection. If a timeseries is of sufficient length this damping will mean that the resonant peaks of the modes will have an observed width in a frequency-power spectrum. The full width of a peak can then be used to provide a measure of the damping each mode undergos. The width of a mode, $\Delta\nu$, is related to the mode lifetime (e-folding time), τ, by

$$\Delta\nu = \frac{1}{\pi\tau}. \tag{3.15}$$

If the length of a timeseries is significantly greater than the lifetime of a mode the timeseries will extend over several realizations of the mode and so the modal peak will be resolved across several bins in the frequency domain. It is therefore desirable to search for a concentration of prominent spikes over a range of bins similar to the expected modal line width (Chaplin et al., 2002a). A peak is regarded as several spikes with relative heights $\geq s_\nu$ that cover a range of frequencies.

The lifetimes of modes increase as the frequency of the modes decreases. Below $\sim 2500\,\mu$Hz the widths of the mode peaks decrease rapidly with frequency. If the length of the timeseries is less than the lifetime of a mode the width of the mode will be narrow enough to confine the majority of the mode's power into one bin. The signal will behave like an undamped sine wave, however, the finite length of the observations will mean

that the frequency-power spectrum of the signal will actually be a sinc squared function centred on the frequency of the mode. If the acoustic signal from the mode is not commensurate with the time bins the power of the mode may lie in more than one bin even if the underlying width of the mode cannot be resolved. Furthermore, although any solar-cycle variation of low-n eigenfrequencies would be small it could be enough to split the signal between two bins. Therefore, even at low frequencies, where the lifetime of modes may be longer than the length of the timeseries, it is still possible to search for more than one spike that contains power from a mode.

In Section 3.5 artificial data will be used to assess how effective statistical tests are at detecting modes. The artificial spectrum will be 9.5 artificial years in length. In a spectrum of 9.5 years the width of the peaks is expected to be less than the bin width for modes with a frequency of less than $\sim 950\mu$Hz. Therefore in the artificial frequency-power spectrum modes cannot be resolved at frequencies of less than $\sim 950\mu$Hz. So a spectrum can be searched, at frequencies above $\sim 950\mu$Hz, for at least two spikes, both of which are greater than some relative height, $s_{\nu 1}$, and both of which are separated by less than twice the expected full width at half maximum (FWHM) of the mode.

Let N_1 be the number of bins that cover twice the FWHM of a mode. The probability, p_{width}, of finding at least two spikes that are less than N_1 bins apart anywhere in the spectrum is given by

$$p_{\text{width}} = p(s_{\nu 1}) \left[1 - (1 - p(s_{\nu 1}))^{N_1 - 1} \right]. \tag{3.16}$$

In other words p_{width} is given by the probability that an individual spike is greater than $s_{\nu 1}$ and that at least one spike in $N_1 - 1$ bins has a relative power of at least $s_{\nu 1}$. The probability of then observing at least one of these structures over N bins is

$$P_{\text{width}} = 1 - (1 - p_{\text{width}})^{N - N_1 + 1}. \tag{3.17}$$

A special case of this arrangement occurs when the prominent bins are forced to be in consecutive bins. As mentioned previously the power from an unresolved mode may

lie in two consecutive bins if the frequency of the mode is not commensurate with the window function of the timeseries. Therefore searching for two consecutive prominent spikes is relevant even when the power of a mode should be confined to one bin only. Furthermore this would provide a more stringent test for any mode whose power is spread across several bins. The probability of observing 2 spikes in consecutive bins is simply

$$P_{2\text{bins}}[2; p(s_{\nu 1}), 2] = p(s_{\nu 1})^2.$$ (3.18)

The number of bins in Δ_p is N and so the number of arrangements of pairs is $N - 1$. The probability that a range, Δ_p, contains at least 1 pair of such spikes is therefore

$$
\begin{aligned}
P_{\Delta\nu}[1; P_{2\text{bins}}, N - 1] &= 1 - [1 - P_{2\text{bins}}]^{N-1} \\
&\approx (N - 1)P_{2\text{bins}} - \frac{(N - 1)(N - 2)}{2}P_{2\text{bins}}^2 \\
&\approx (N - 1)p(s_{\nu 1})^2 - \frac{(N - 1)(N - 2)}{2}p(s_{\nu 1})^4.
\end{aligned}
$$ (3.19)

This can be rearranged to give

$$p(s_{\nu 1}) = \left[\frac{1}{N - 1} \pm \left\{ \frac{(N - 1) - 2P_{\Delta\nu}(N - 2)}{(N - 1)(N - 2)^2} \right\}^{\frac{1}{2}} \right]^{\frac{1}{2}},$$ (3.20)

and

$$s_{\nu 1} = -\frac{1}{2} \ln \left[\frac{1}{N - 1} \pm \left\{ \frac{(N - 1) - 2P_{\Delta\nu}(N - 2)}{(N - 1)(N - 2)^2} \right\}^{\frac{1}{2}} \right].$$ (3.21)

As with equation 3.14 the negative square root is taken. Incidently this is approximately half the height required when looking for a single spike at a given threshold probability ($P_{\Delta\nu}$ or P_s). These equations, along with equation 3.9, were then used to predict the number of Δ_p containing at least 1 pair of prominent consecutive spikes in a spectrum of noise. Once again simulations were performed to test the accuracy of these equations. The predictions agree with the simulations to an accuracy of better than 0.5%.

When searching the spectra for statistically prominent features a low threshold value

of P is set so that there is a very small probability of observing a prominent feature at least once in N bins. Using this threshold probability and equation 3.21 it is possible to determine threshold relative heights, such that if a feature is observed to be greater than this threshold height it is considered to be statistically significant. The width of a mode is not the only property that can be used to search for low-frequency modes. We now investigate how the multiplet structure exhibited by some modes can be used to construct a statistical test.

3.4 The detection of multiplet structures in a single pure noise spectrum

The rotation of the Sun splits modes with l greater than zero into $2l + 1$ components, each of which are described by a different azimuthal order, m. The azimuthal order can take any value between $\pm l$. However, in Sun-as-a-star data only $l + 1$ components are observed, instead of $2l + 1$, as the near-equatorial view of the Sun limits the visibility of certain modes. In fact modes are only clearly visible in Sun-as-a-star data when $l + m$ is even. Chaplin et al. (2001) found the separation between adjacent m for the $l = 1$, 2 and 3 modes, at frequencies above 1500μHz, is approximately 400nHz. At mode frequencies greater than 400μHz theoretical modelling implies that the magnitude of the frequency splitting of p-mode multiplets is relatively insensitive to the poorly understood rotation at less than 0.2 solar radii (Provost et al., 2000). It is therefore reasonable to assume that the synodic splitting in the searched-for low-frequency modes is similar to that observed at higher frequencies. However, magnetic fields introduce asymmetries in the splitting which become apparent in full-disc data for modes with $l \geq 2$. If the magnetic field originates near the surface the magnitude of the $l = 2$ and $l = 3$ mode asymmetry is only of the order of a few nHz at low frequencies. The effect of a deep-seated field on the asymmetry of the splittings is less straightforward to predict as the magnitude of the asymmetry varies in a quasi-oscillatory manner with frequency. Current data implies this variation is also small, but the evidence is by no means conclusive. The possibility

of an asymmetric arrangement needs to be taken into account when defining a multiplet structure. Therefore we take the uncertainty in the difference in frequency between any two components of a multiplet, $\delta\nu$, to be ±100nHz. The number of bins, N_δ, in a frequency range $\delta\nu$ is given by $N_\delta = \delta\nu T$, where T is the length of the timeseries. In the frequency range that will be searched this is always greater than the width of the mode, $\Delta\nu$. As $l + m$ must be even for the modes to be observed the difference in m between observable modes of a given l, δm, will always be a multiple of 2. The aim is therefore to search for prominent spikes that lie at multiples of $\sim 800\pm100$nHz from each other.

The probability of finding at least one spike by chance with a relative height above $s_{\nu 2}$ over the uncertainty range $2N_\delta$ is

$$P_\delta = P[1; p(s_{\nu 2}), 2N_\delta] = 1 - [1 - p(s_{\nu 2})]^{2N_\delta}. \tag{3.22}$$

Therefore, the probability of finding 2 spikes above $s_{\nu 2}$ that satisfy the requirements for a multiplet pair is

$$P_m = p(s_{\nu 2})P_\delta = p(s_{\nu 2})\{1 - [1 - p(s_{\nu 2})]^{2N_\delta}\}. \tag{3.23}$$

The number of possible arrangements of the structure over a range Δ_p is dependent on the synodic splitting, $N_{\text{syn}} = 400$nHz, and the uncertainty allowed (N_δ). In a well resolved spectrum N_{syn} and N_δ are both significantly less than the number of bins in Δ_p, N. Therefore, the number of arrangements can be taken to be approximately N. Hence the probability of observing at least one multiplet pair in a frequency range Δ_p is

$$P_{\text{multi}} \approx P[1 : P_m, N] = 1 - [1 - P_m]^N. \tag{3.24}$$

Simulations were performed in the same manner as in Section 3.1, except that multiplets containing two prominent spikes were searched for. In this case there is no simple, accurate expression for $s_{\nu 2}$ and so in the simulations $s_{\nu 2}$ was fixed and used to calculate

the corresponding probability. The simulated results are approximately 2% lower than the predicted results.

For $l = 2$ and 3 modes it is possible that three components of the multiplet could be observed. It is therefore viable to search for triplets, where a central component is flanked by two outer components that are 800±100nHz away. The probability of observing a triplet, P_t, is given by

$$P_t = p(s_{\nu 2})P_\delta^2 = p(s_{\nu 2})\{1 - [1 - p(s_{\nu 2})]^{2N_\delta}\}^2. \qquad (3.25)$$

The probability of observing at least one triplet structure in a frequency range of Δ_p is

$$P_{\text{triplet}} = 1 - [1 - P_t]^N. \qquad (3.26)$$

We have now developed the tools with which frequency-power spectra can be searched for both multiplet pairs and triplets that contain two or three prominent spikes respectively. We go on to determine how effective the statistical tests that we have developed in this chapter are at detecting simulated modes. We therefore search a set of simulated data for prominent single spikes, peaks of spikes and multiplet structures.

3.5 Detection of mode candidates in FLAG data

The Solar Fitting at Low Angular degree Group (FLAG) is an international collaboration that aims to improve analysis techniques for low-l solar p modes. FLAG data are artificial data that are designed to mimic unresolved Doppler velocity observations of the visible disc with a S/N per unit time characteristic of BiSON and GOLF measurements. The FLAG data, therefore, represent Sun-as-a-star observations. A timeseries containing low-frequency modes was investigated as a blind test of the statistics developed in this chapter. The data were made by W.J. Chaplin and the frequencies of the modes he included in the data were only revealed after the statistical tests were performed. Modes

with angular degrees between $0 \leq l \leq 3$ were included in the timeseries along with the appropriate level of broadband noise. The timeseries contained 7,464,960 points, with a cadence of 40s, and so spanned approximately 9.5 artificial years. No gaps were included in the timeseries and so the duty cycle was 100%. The data mimic observations that were made near the ecliptic plane, and so only modes where $l + m$ is even are visible. Since $\mid m \mid \leq l$ a mixture of different m will only be observed when $l \geq 1$. In all of the following searches the minimum threshold height for a spike to be considered a candidate corresponded to a probability of $P_s = 0.1$. This threshold was chosen as it means that there is less than a 10% chance that any detected mode candidates are part of the noise background and so the confidence level is 90%. However, more detailed discussions on whether this is an appropriate threshold level are included later in this section and in Chapter 5.

The statistical predictions of the distribution of noise that were described in Sections 3.2, 3.3 and 3.4 have been used to look for prominent spikes in the FLAG data. It should be noted that the statistical predictions give the probability that a prominent spike or structure of spikes is due to noise, which is the H0 hypothesis. Five tests were performed on the frequency-power spectrum produced by taking the FFT of the FLAG timeseries:

Test 1: The frequency-power spectrum was searched for prominent individual spikes. The probability that a spike is due to noise was calculated using equations 3.2 and 3.4.

Test 2: A test was performed to search for two prominent spikes in consecutive bins. The power of the least prominent of the two spikes was used in equation 3.2 to calculate p. This was then used with equation 3.19 to find the probability that a peak is due to noise.

Test 3: The spectra were searched for two prominent spikes that are separated by less than twice the predicted width of the mode. A power-law extrapolation, used in Chaplin et al. (2002a), was used to predict how the width of a mode varies with frequency. Equations 3.2, 3.16 and 3.17 were then used to determine the probability that any

such structures are due to noise.

Test 4: A test was performed to search for pairs of spikes in multiplets. Each pair of spikes must be separated by $800 \pm 100\text{nHz}$, $1600 \pm 100\text{nHz}$ or $2400 \pm 100\text{nHz}$ which correspond to separations in the azimuthal order of 2, 4 or 6 respectively. This ensures that $l + m$ is even for both components of the multiplet. Therefore, when the spikes are separated by $800 \pm 100\text{nHz}$ modes with $l = 1$, 2 or 3 can be revealed as each exhibit components that are separated by $\delta m = 2$. However, when the spikes are separated by $1600 \pm 100\text{nHz}$ only modes with $l = 2$ or 3 can be detected as the required difference in m is now $\delta m = 4$. Only $l = 3$ modes can be detected when the spikes are separated by $2400 \pm 100\text{nHz}$ as the required difference in m is $\delta m = 6$. The probability that each multiplet is due to noise was calculated using equations 3.2, 3.23 and 3.24.

Test 5: This test searched for multiplets containing three spikes. The central spike is separated from the outer components of the triplet by $800 \pm 100\text{nHz}$. Therefore this test can search for modes with $\ell = 2$ or 3 as these modes contain at least 3 visible components. Equations 3.2, 3.25 and 3.26 were then used to find the probability that each triplet would occur in pure noise.

Disregarding noise the power of a peak can be reduced to approximately 40% of its maximum value depending on how commensurate the phase of the signal is with the frequency spectrum bins. Noise can interfere both constructively and destructively with the signal from the mode and so can increase (or decrease) the severity of this effect. One possible solution to this is to oversample the spectrum by zero-padding the timeseries. Alternatively removing successively appropriate small numbers of data from the end of the main timeseries will shift the absolute frequency locations of the bins, thereby allowing the observer to scan for the frequency at which the signal is commensurate. Removing different numbers of bins from the end of a spectrum alters the location of the bins and therefore allows the spectrum to be evaluated at different locations. If a

sufficient number of shifts are made different realizations of the noise background will be given. Monte Carlo simulations have shown that shifting the spectrum 25 times will give the equivalent of 3 realizations of the noise (Chaplin et al., 2002a). However, the Monte Carlo simulations also showed that little is gained by performing more shifts. Therefore the length of the timeseries was altered 25 times and the same analysis was performed on each of the resulting 25 spectra. The length of the timeseries was altered by removing a different number of bins from the end of the original timeseries. The largest number of bins removed was 1460 and so a timeseries containing 7,463,500 bins was produced. The smallest number of bins removed was 85, creating a timeseries containing 7,464,875 bins. These numbers were chosen because they ensure the sum of the prime factors of the remaining number of points in the timeseries is small as this has repercussions for the time taken to perform the FFT. As the numbers of bins to be removed only constitute a small fraction of the base line in its entirety the 25 spectra should provide a fairly even sampling across a range of bin shifts.

The solar spectrum is, in fact, only locally regarded as being flat ('white'). A 'white' spectrum can be approximated by calculating a running mean for the data over regions of approximately 100μHz. However, this approximation falls short at very low frequencies as the level of noise across a 100μHz bin cannot be regarded as constant. Hence a running mean was taken to ensure that any prominent spikes or structures are central to the 100μHz region for which the mean level is calculated. By using a running mean we are effectively pre-whitening the spectrum. To test whether this has an effect on the underlying statistics the simulations described in Section 3.1 were repeated but using this method of searching for modes. The results of this simulation were the same as the results found in Section 3.1 and so taking a running mean does not affect the statistics.

The FLAG frequency-power spectrum was searched at frequencies between 1000 and 1500μHz for statistically significant spikes and structures. Although the exact frequencies of the modes that were included in the FLAG data were not known it had been acknowledged that they all had frequencies greater than 1000μHz. The frequency, ν,
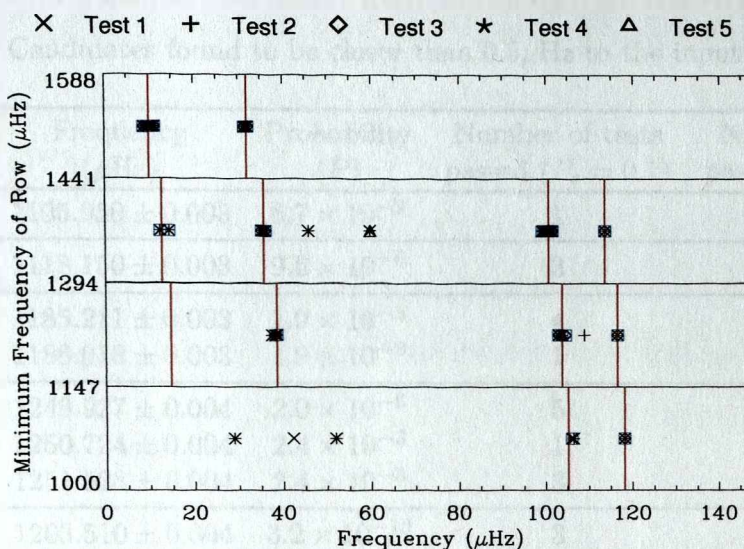
Figure 3.4: The possible candidates uncovered when searching for prominent spikes and structures in the FLAG data. The black symbols at the centre of each row represent the detections made for a 10% threshold. Different symbols are used to distinguish between the different tests. The red vertical lines indicate the frequencies of the input modes. The blue squares denote the detections that lie within 0.5μHz of an input mode frequency.

was recorded for any spike or structure that had less than a 10% chance of being due to noise. However, a more stringent threshold was also set at $P_s \leq 0.01$ to investigate the reliability of the tests. It is important to note that any spikes found by this method are not definitely evidence of a mode. There is simply a large probability they are not due to noise and therefore any spikes detected will be referred to as mode candidates.

Figure 3.4 shows all of the mode candidates that had a less than 10% chance of occurring in a pure noise spectrum in an echelle diagram. The entire frequency range that was searched has been split into strips of 147μHz and these strips have been placed one above another. The length of the baseline of the echelle diagram was chosen to be 147μHz so that the overtones of modes with a given l lie in near-vertical lines. The red vertical lines show the central frequency of the input modes. The different symbols represent the different tests that were performed (see the legend on Figure 3.4). Mode candidates that lie less than 0.5μHz from any components of the input modes have been highlighted with a blue square. In the case of the multiplet and triplet tests all of the components of the detected structure must correspond to input mode frequencies. Table

Table 3.3: Candidates found to be closer than 0.5μHz to the input frequencies.

l	n	m	Frequency (μHz)	Probability (P)	Number of tests passed ($P_s = 0.1$)	Number of tests passed ($P_s = 0.01$)
2	6	+2	1105.950 ± 0.003	6.7×10^{-2}	1	0
0	7	0	1118.150 ± 0.003	9.6×10^{-6}	3	2
1	7	-1	1185.211 ± 0.003	1.9×10^{-3}	4	2
1	7	+1	1186.018 ± 0.003	1.9×10^{-3}	1	1
2	7	-2	1249.927 ± 0.004	2.0×10^{-6}	5	4
2	7	0	1250.724 ± 0.004	2.4×10^{-3}	1	1
2	7	+2	1251.528 ± 0.004	2.4×10^{-3}	3	2
0	8	0	1263.510 ± 0.004	3.2×10^{-10}	3	3
3	7	-3	1305.590 ± 0.005	1.7×10^{-2}	1	0
3	7	+1	1307.200 ± 0.005	1.7×10^{-2}	1	0
1	8	-1	1329.211 ± 0.004	2.3×10^{-11}	4	4
1	8	+1	1330.018 ± 0.004	2.3×10^{-11}	4	4
2	8	-2	1393.891 ± 0.005	7.4×10^{-10}	5	3
2	8	0	1394.722 ± 0.005	7.4×10^{-10}	5	5
2	8	+2	1395.513 ± 0.005	7.4×10^{-10}	5	5
0	9	0	1407.616 ± 0.005	0	3	3
3	8	-3	1449.893 ± 0.006	9.8×10^{-9}	4	3
3	8	+1	1451.437 ± 0.007	1.1×10^{-2}	1	0
3	8	+3	1452.291 ± 0.006	9.8×10^{-9}	4	4
1	9	-1	1472.441 ± 0.005	0	4	4
1	9	+1	1473.213 ± 0.005	0	4	4

3.3 shows those candidate detections that are found within 0.5μHz of the frequencies of the input modes. Notice that there are six clear false detections made (not including the multiplet tests that contained one false detection and one candidate mode component). We have searched 5 frequency ranges of width $100\,\mu$Hz with a threshold of $P_s = 0.1$. By altering the length of the timeseries 25 times we have created 3 realisations of the noise. Furthermore, we have performed 5 different tests. Therefore the number of false detections we would expect is $0.1 \times 5 \times 3 \times 5 = 7.5$. Hence the number of false detections observed is in line with expectations.

This method is not very effective for finding the exact frequency of peaks that have

a large width as there are significant numbers of spikes that are greater than s_ν even when P_s is extremely small and s_ν is very large (e.g $s_\nu \geq 30$). At frequencies greater than 1500μHz the power of the modes is spread over many bins and the powers in the majority of these bins is greater than the threshold height. Therefore the search was limited to frequencies less than 1500μHz. Even so some of the higher-frequency modes in this range were detected in more than one bin, therefore, the probability that the most prominent spike (or structure of spikes) is (are) part of the background noise has been recorded.

In a timeseries of this length all of the input modes are oversampled and so have a resolved width in frequency. Therefore the errors on the observed frequency quoted in Table 3.3 have been calculated following the method described in Chaplin et al. (2002b). The error on the frequency, $\delta\nu_{nl}$ is given by

$$\delta\nu_{nl} = \left(\frac{f[\beta_{nl}]\Delta\nu_{nl}}{4\pi T}\right)^{\frac{1}{2}}, \tag{3.27}$$

where $\Delta\nu_{nl}$ is the width of the line, T is the length of the timeseries and $f[\beta_{nl}]$ is a function that is dependent on the noise-to-signal ratio, β_{nl}. The function, $f[\beta_{nl}]$ takes the form

$$f[\beta_{nl}] = (1 + \beta_{nl})^{\frac{1}{2}} \left[(1 + \beta_{nl})^{\frac{1}{2}} + \beta_{nl}^{\frac{1}{2}}\right]^3, \tag{3.28}$$

while $\beta_{nl} = 1/s_\nu$.

Only 1 of the 12 included modes was not detected in some capacity, meaning that at least one component of the mode passed at least one test. The completely undetected mode is the $l = 3$, $n = 6$ mode at 1161.69μHz. However, as can be seen from Figure 3.4 there are a lot of detections that do not correspond to the input modes. These are known as false detections. It is therefore useful to introduce the further requirement that the candidate modes pass at least two of the tests that were performed. When this constraint is introduced the number modes that are not detected increases to three as the $l = 2$, $n = 6$ mode at 1105.17μHz and the $l = 3$, $n = 7$ mode at 1306.80μHz do not pass this

requirement. However, the number of false detections is reduced from 6 to 1. This is still not ideal as the false detection could be mistaken for a real mode in actual solar data. It is therefore worth considering reducing the threshold limit to $P_s = 0.01$. This disqualifies all of the false detections and it does not reduce the number of modes that are detected. Therefore it is definitely advantageous to reduce the threshold levels to a 1% chance that any candidates could occur in noise. However, it should be noted that one of the false detections has a probability of being noise of less than 1% and so it is important to combine this more stringent threshold level with the requirement that each detection passes at least two of the tests.

Overall the statistical tests have proved to be reasonably effective at detecting the simulated modes. However, a significant number of false detections were also made. It is important to minimize the number of false detections as they may be mistaken for modes. Therefore in the future the threshold probability was set at $P_s = 0.01$ and each mode candidate was required to pass at least two tests.

3.6 Summary

The statistics that describe the noise background in a frequency-power spectrum have been used to develop various statistical tests. These tests can be used to search for prominent spikes or structures of spikes that are statistically unlikely to be part of the noise background. The tests are based on those developed in Chaplin et al. (2002a). However, here we have added an extra term in the approximation of P_s to produce a more accurate description of the background statistics. We have taken advantage of some of the physical properties of a mode observed in a frequency-power spectrum. These properties include the width of a mode, which is observed because the power of a mode is damped, and the multiplet structure of a mode, which is observed because of the solar rotation. Therefore, in addition to searching for individual prominent spikes tests can be performed that search for prominent peaks of spikes, multiplet pairs and triplets.

A blind test was performed on FLAG data to determine how effective these tests are at finding low-frequency modes. When the threshold probability was set at $P_s = 0.1$ all but one of the input modes were found but a substantial number of false detections were also made. However, the number of false detections was reduced, without significantly lowering the number of mode detections made, when the threshold probability was reduced to $P_s = 0.01$ and each candidate was required to pass at least two tests.

As the statistical tests were successful at finding the simulated modes they have been used to search real Sun-as-a-star data for low-frequency modes (Chapter 5). However, it is necessary to stress here that any mode candidates are still not definitely modes, there is just a large probability that they are not due to the background noise.

4. DETECTION OF PROMINENT SPIKES IN TWO SPECTRA

BiSON data are just one of many sets of solar observations that are currently available. Another example is the data recorded by the GOLF instrument on board the SOHO spacecraft. Like BiSON, GOLF measures spatially unresolved solar Doppler velocities. BiSON and GOLF data can, therefore, be searched for prominent spikes or structures of spikes using the statistics described in the previous chapter. The BiSON network and GOLF have both been making near-continuous observations of the Sun for more than 10 years. Therefore it is possible to compare contemporaneous observations made by the BiSON network and GOLF. Contemporaneous data should show a common signal from solar oscillations. This chapter will describe how the probability of observing prominent spikes (or patterns of spikes) that are found in the same bin (or bins) in each of two frequency spectra can be found. The statistics developed in Chapter 3 have been adapted so that they can be applied to two sets of contemporaneous observations. This problem is trivial when the data are completely independent (Section 4.1). However, the background noise in the two sets of observations will not be completely independent as there will be some common noise from the solar granulation. There will also be independent instrumental noise and any data taken by ground-based instruments, such as BiSON data, will contain atmospheric noise. Furthermore, some of the solar noise will be independent. The amount of noise common to the two sets of data will affect the probability that any prominent structures present at the same frequencies in both sets of data are due to noise. Section 4.2 will calculate this probability taking into account the amount of common noise present.

It should be noted that comparisons are not limited to BiSON and GOLF data.

Observations from other instruments, such as MDI, which is also onboard SOHO, can be used. The BiSON data are measured by a network of instruments that are positioned at different longitudes around the world. The positioning of the BiSON instruments allows near-continuous solar observations to be made, assuming that the instruments are fully operational and that the weather permits observations to be made. However, the separation in longitude between some of the BiSON sites means that it is possible that the Sun will be above the horizon, and so observable, at more than one site. In fact at certain times as many as four BiSON instruments could be taking contemporaneous observations of the Sun. It is, therefore, possible to compare contemporaneous data from different BiSON instruments and this will be done in Chapter 7. The remainder of this chapter will describe a method for comparing data in a generalized manner that is independent of the instruments from which the data being compared originate.

4.1 Detection of prominent spikes in two independent spectra

As an initial, simplistic example we will describe the statistics governing two timeseries that contain completely independent Gaussian noise. Assuming the noise in the power spectra created from the independent timeseries can be described by a χ^2 2 d.o.f. distribution the probability that a frequency range of width Δ_p will contain at least one false detection in the same bin, in each of the two independent spectra, can be found. The probability that two spikes in the same bin in two independent spectra will have relative heights greater than or equal to the relative height $s_{\nu 2}$ is

$$P_2 = p(s_\nu)^2 = \exp(-2s_{\nu 2}), \qquad (4.1)$$

where $s_{\nu 2}$ is defined as the ratio between the height of a spike and the mean level of noise surrounding that spike. Therefore, the probability that a frequency range, Δ_p, which consists of N bins, will contain at least 1 pair of spikes with a relative height greater

than $s_{\nu2}$ in the same frequency bin is

$$P_{s2}[1; P_2, N] = 1 - [1 - P_2]^N. \qquad (4.2)$$

Using an expansion similar to the one given in equation 3.13 it can be shown that

$$P_{s2}[1; P_2, N] = Np(s_{\nu2})^2 - \frac{N(N-1)}{2}p(s_{\nu2})^4, \qquad (4.3)$$

which can then be rearranged to give an expression for a threshold relative height, $s_{\nu2}$, for given values of N and P_{s2}:

$$s_{\nu2} = -\frac{1}{2}\ln\left\{\frac{1}{N-1} \pm \left[\frac{N - 2P_{s2}(N-1)}{N(N-1)^2}\right]^{\frac{1}{2}}\right\}. \qquad (4.4)$$

As with equation 3.14 the negative square root is taken. Comparing this equation with equation 3.14 shows that when looking for coincident spikes in two independent power spectra the threshold height is half that required when looking for prominent spikes in just one spectrum.

Simulations that counted the number of detections made in pairs of completely independent noise spectra were performed to check the validity of this equation. The simulations were performed for various frequency ranges, Δ_p, and probabilities, P_{s2}. A false detection was made if a frequency range of width Δ_p contained at least one pair of spikes that had a relative height greater than $s_{\nu2}$ positioned in the same frequency bin in each of the two simulated spectra. There was only a 1% difference between the predicted and the simulated number of false detections for all values of Δ_p and P_{s2} for which the simulations were performed.

As mentioned in Section 3.3 the modes of the Sun do not usually appear as a single spike in a single bin, but rather as a peak of spikes that cover several bins. For some of the modes this is because the energy of the mode is heavily damped over time. However, the power from modes with long lifetimes can also be spread over two consecutive bins if

the frequency of the mode is not commensurate with the frequency bins in the frequency-power spectrum. It is therefore usually necessary to look for spikes larger than the threshold height over several bins. The probability of observing a prominent spike in the first spectrum and a prominent spike in the second spectrum which lies either in the same bin as the spike in the first spectrum or 1 bin either side is given by

$$
\begin{aligned}
P_3 &= P[1; p(s_{\nu 2}), 1] \times P[1; p(s_{\nu 2}), 3] \\
&= p(s_{\nu 2}) \left[p(s_{\nu 2})^3 - 3p(s_{\nu 2})^2 + 3p(s_{\nu 2}) \right].
\end{aligned}
\tag{4.5}
$$

The value of P_3 can then be substituted into equation 3.4 to find the probability of such an arrangement of spikes occurring at least once in N bins.

Once again simulations were performed to ensure equation 4.5 represents the statistics accurately. A false detection was made if there were at least one pair of spikes, one spike from each spectrum, that both had relative heights greater than $s_{\nu 2}$. This time the spike in the second frequency-power spectrum had to be either in exactly the same frequency bin as the spike in the first frequency-power spectrum or one bin either side of that frequency bin. The simulated and predicted number of false detections were in good agreement as the difference between them was less than 0.6% for all Δ_p and $s_{\nu 2}$ that were used in the simulations.

The statistics described in this section only apply when the noise in two sets of data are completely independent. However, the assumption that two spectra are completely independent is very simplistic and, when considering contemporaneous data, often inaccurate. It is far more likely that two sets of contemporaneous data will contain some common noise and so we begin the next section by discussing the possible origins of the common noise. The presence of the common noise alters, in a non-trivial manner, the probability of observing a prominent feature at the same frequency in two frequency-power spectra. The remainder of this chapter is, therefore, dedicated to discussing the statistics governing two sets of data with a known amount of common noise.

4.2 Statistics of coherent noise spectra

Any two instruments that are making contemporaneous observations of the Sun will not only be observing the same oscillations but some of the solar background continuum observed by both instruments will be common to both sets of data. The granulation that can be continually observed on the Sun produces noise in Doppler velocity observations (see Chapter 2). This source of noise is particularly important when looking for low-frequency p modes and g modes as the level of solar noise increases as frequency decreases. Real data will contain some solar noise that is common to two sets of contemporaneous observations and some solar noise that is characteristic to each individual set of data and therefore not common. This is because the instruments making the observations may have slightly different views of the Sun meaning the observations will be weighted to different parts of the solar disc (see Chapter 8 for more details). Also the height in the photosphere at which the observations are made is likely to be different both because of Doppler imaging and because different instruments may observe different Fraunhofer lines. For example, BiSON uses the neutral K line at $\sim 770\,\mathrm{nm}$ but GOLF observations use the neutral sodium doublet at $\sim 590\,\mathrm{nm}$. When BiSON and GOLF data are compared there will also be noise due to instrumental effects and, for the BiSON data, atmospheric effects that will not be common to both sets of observations. In chapter 7 we compare many different sets of contemporaneous BiSON data. In some cases much of the instrumental noise and some of the atmospheric noise are also common to both sets of data.

The amount of common noise will vary depending on frequency and the instruments from which the data have been taken. It is, however, possible to determine the probability that a spike in the same frequency bin in each of two spectra, which contain a known amount of common noise, has a relative amplitude greater than a given threshold amplitude. Until now we have considered the statistics that are relevant to frequency-power spectra. However, in this section we will be considering frequency-amplitude spectra.

This is because the statistics that we are now going to develop assume the noise has a Gaussian distribution. The noise in frequency-amplitude spectra is normally distributed, however, this in not the case in frequency-power spectra, as we have already seen. To calculate the probability of observing a coincident prominent spike the amount of common noise must be determined. Therefore, we go on to describe the manner in which the common noise affects the probability and how the amount of common noise shared by two sets of data can be found.

Let the proportion of noise in two sets of data be represented by α such that when $\alpha = 0$ there is no coherent data but when $\alpha = 1$ the two sets are identical. The 'coherency' of the two data sets can be related to α in a manner that will be described later (see equation 4.70).

We will assume that the real and imaginary frequency-amplitude spectra contain normally distributed noise. We define x_1 and y_1 as the real and imaginary amplitudes of the first spectrum. They are given by the equations

$$x_1 = a + \alpha b, \tag{4.6}$$

$$y_1 = c + \alpha d, \tag{4.7}$$

where a, b, c and d are arrays of normally distributed random numbers. Then we define x_2 and y_2 as the real and imaginary amplitudes of the second spectrum. They are given by the equations

$$x_2 = \alpha a + b, \tag{4.8}$$

$$y_2 = \alpha c + d. \tag{4.9}$$

The equations for the real part of the observed noise, x_1 and x_2, can then be expressed

in the form

$$x_1 = \alpha a + \alpha b + (1 - \alpha)a, \tag{4.10}$$

$$x_2 = \alpha a + \alpha b + (1 - \alpha)b. \tag{4.11}$$

Clearly $\alpha a + \alpha b$ represents the common 'solar noise' and $(1 - \alpha)a$ and $(1 - \alpha)b$ in x_1 and x_2 respectively represent all 'other noise' that is not common to both sets of observations. The variance of the distributions, a, b, c and d are taken to be equal (i.e. $\sigma_a^2 = \sigma_b^2 = \sigma_c^2 = \sigma_d^2$). However, the two spectra may contain different amounts of 'other noise'. Therefore a multiplying factor of k is included to alter the variance of the 'other noise' in x_2. Hence

$$x_1 = \alpha a + \alpha b + (1 - \alpha)a,$$

$$x_2 = \alpha a + \alpha b + k(1 - \alpha)b. \tag{4.12}$$

The variance of the other noise in x_2 is, therefore, k^2 times the variance of the other noise in x_1. Note that it is reasonable to assume that the variance of the real and imaginary parts of each spectrum will be equal. Therefore a multiplying factor of k can also be included to alter the variance of the other noise in y_2. So we now have

$$y_1 = \alpha c + \alpha d + (1 - \alpha)c,$$

$$y_2 = \alpha c + \alpha d + k(1 - \alpha)d. \tag{4.13}$$

The variance of the other noise in y_2 is now k^2 times the variance of the other noise in y_1. It is also necessary to multiply the whole of x_2 and y_2 by a factor of j to allow for differences in the variance of the overall distributions. This gives

$$x_1 = \alpha a + \alpha b + (1 - \alpha)a = a + \alpha b, \tag{4.14}$$

$$y_1 = \alpha c + \alpha d + (1 - \alpha)c = c + \alpha d, \tag{4.15}$$

$$x_2 = j\left[\alpha a + \alpha b + k(1-\alpha)b\right], \tag{4.16}$$

$$y_2 = j\left[\alpha c + \alpha d + k(1-\alpha)d\right]. \tag{4.17}$$

The general expression for a probability density function is given by

$$P(X) = \frac{1}{(2\pi)^{\frac{z}{2}}}\frac{1}{\sqrt{|C|}}\exp\left(-\frac{1}{2}X^T C^{-1} X\right), \tag{4.18}$$

where z is the number of variables, C is a coherency vector and X is a matrix of the variables. In this case the variables are the real and imaginary amplitudes of both spectra and so $z = 4$. The column vector of the variables, X, is given by

$$X = \begin{pmatrix} x_1 \\ y_1 \\ x_2 \\ y_2 \end{pmatrix}, \tag{4.19}$$

and the coherency matrix, C, is given by the following expectation values:

$$C = \begin{pmatrix} E(x_1^2) & E(x_1 y_1) & E(x_1 x_2) & E(x_1 y_2) \\ E(y_1 x_1) & E(y_1^2) & E(y_1 x_2) & E(y_1 y_2) \\ E(x_2 x_1) & E(x_2 y_1) & E(x_2^2) & E(x_2 y_2) \\ E(y_2 x_1) & E(y_2 y_1) & E(y_2 x_2) & E(y_2^2) \end{pmatrix}. \tag{4.20}$$

Now the expectation value of x_1^2 can be determined as follows:

$$\begin{aligned} E(x_1^2) &= \langle a^2 + 2\alpha ab + \alpha^2 b^2 \rangle, \\ &= \langle a^2 \rangle + 2\alpha\langle ab \rangle + \alpha^2\langle b^2 \rangle. \end{aligned} \tag{4.21}$$

For any variable, x, with a given probability density distribution, $P(x)$, the function $f(x)$

has an expectation value given by

$$\langle f \rangle = \int_{-\infty}^{\infty} f(x) P(x) \mathrm{d}x.$$ (4.22)

Therefore, for a Gaussian distribution, such as a,

$$\langle a^2 \rangle = \int_{-\infty}^{\infty} a^2 \frac{1}{\sigma_a \sqrt{2\pi}} \mathrm{e}^{-a^2/2\sigma_a^2} \mathrm{d}a = \sigma_a^2.$$ (4.23)

Also, as a and b are both normally distributed $\langle ab \rangle = 0$. Therefore equation 4.21 reduces to

$$\mathrm{E}(x_1^2) = \sigma_a^2 + \alpha^2 \sigma_b^2 = (1 + \alpha^2)\sigma_a^2,$$ (4.24)

as $\sigma_a^2 = \sigma_b^2$. Using equation 4.22 to find the expectation value of x_2 gives

$$\begin{aligned} \mathrm{E}(x_2^2) &= j^2 \left(\alpha^2 \langle a^2 \rangle + 2 \left[\alpha^2 + k\alpha(1-\alpha) \right] \langle ab \rangle + \left[\alpha^2 + k^2(1-\alpha)^2 \right] \langle b^2 \rangle \right), \\ &= j^2 \left(\alpha^2 + [k(1-\alpha) + \alpha]^2 \right) \sigma_a^2. \end{aligned}$$ (4.25)

Similarly

$$\begin{aligned} \mathrm{E}(x_1 x_2) &= j \left(\alpha \langle a^2 \rangle + \left[\alpha^2 + \alpha + k(1-\alpha) \right] \langle ab \rangle + \alpha \left[k(1-\alpha) + \alpha \right] \langle b^2 \rangle \right), \\ &= j\alpha \left[1 + k(1-\alpha) + \alpha \right] \sigma_a^2. \end{aligned}$$ (4.26)

The symmetry of the real and imaginary parts of each spectrum means that

$$E(y_1^2) = (1 + \alpha)\sigma_a^2,$$ (4.27)

$$E(y_2^2) = j^2 \left(\alpha^2 + [k(1-\alpha) + \alpha]^2 \right) \sigma_a^2,$$ (4.28)

$$E(y_1 y_2) = j\alpha \left[1 + k(1-\alpha) + \alpha \right] \sigma_a^2,$$ (4.29)

$$E(x_1 y_1) = E(x_1 y_2) = E(x_2 y_1) = E(x_2 y_2) = 0.$$ (4.30)

Letting $A = (1 + \alpha^2)$, $B = j^2 \left(\alpha^2 + [k(1-\alpha) + \alpha]^2 \right)$ and $D = j\alpha \left[1 + k(1-\alpha) + \alpha \right]$ and

substituting equations 4.27 to 4.30 into equation 4.20 then gives

$$C = \sigma_a^2 \begin{pmatrix} A & 0 & D & 0 \\ 0 & A & 0 & D \\ D & 0 & B & 0 \\ 0 & D & 0 & B \end{pmatrix}. \tag{4.31}$$

It is now possible to find $|C|$ as follows:

$$
\begin{aligned}
|C| &= \sigma_a^8 \left[A \begin{vmatrix} A & 0 & D \\ 0 & B & 0 \\ D & 0 & B \end{vmatrix} + D \begin{vmatrix} 0 & A & D \\ D & 0 & 0 \\ 0 & D & B \end{vmatrix} \right], \\
&= \sigma_a^8 (A^2 B^2 - ABD^2 - ABD^2 + D^4), \\
&= \sigma_a^8 (A^2 B^2 - 2ABD^2 + D^4). \tag{4.32}
\end{aligned}
$$

Substituting back in to equation 4.32 for A, B, and D then gives

$$
\begin{aligned}
|C| &= \sigma_a^8 \left\{ j^2 (1 + \alpha^2) \left(\alpha^2 + [k(1-\alpha) + \alpha]^2 \right) - j^2 \alpha^2 [1 + k(1-\alpha) + \alpha]^2 \right\}^2, \tag{4.33} \\
&= j^4 \sigma_a^8 \left(\alpha^4 - 2\alpha^2 [k(1-\alpha) + \alpha] + [k(1-\alpha) + \alpha]^2 \right)^2, \\
&= j^4 \sigma_a^8 \left(\alpha^2 - [k(1-\alpha) + \alpha] \right)^4. \tag{4.34}
\end{aligned}
$$

Finally, this can be rearranged to give

$$|C| = j^4 \sigma_a^8 \left[\alpha^2 + \alpha(k-1) - k \right]^4. \tag{4.35}$$

Using equation 4.18 we can see that the probability that a spike at the same frequency in each of two frequency-amplitude spectra will have an amplitude of $x_1 + iy_1$ and $x_2 + iy_2$

in each spectrum, respectively, is then given by

$$p_1 = \frac{1}{(2\pi)^2 \sigma_a^4 j^2 \left[\alpha^2 + \alpha(k-1) - k\right]^2} \exp\left[-\frac{1}{2} X^T C^{-1} X\right]. \tag{4.36}$$

To allow us to evaluate equation 4.36 we now need to find an expression for C^{-1}. Once again using the expressions defined earlier for A, B and D and taking advantage of the fact that CC^{-1} is equal to the identity matrix, the inverse of C can be found using the equation

$$\begin{pmatrix} A & 0 & D & 0 \\ 0 & A & 0 & D \\ D & 0 & B & 0 \\ 0 & D & 0 & B \end{pmatrix} \begin{pmatrix} u & 0 & v & 0 \\ 0 & u & 0 & v \\ v & 0 & w & 0 \\ 0 & v & 0 & w \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \tag{4.37}$$

where u, v and w are the non-zero components of C^{-1} that need to be found. Equation 4.37 gives us four simultaneous equations,

$$Au + Dv = 1, \tag{4.38}$$

$$Av + Dw = 0, \tag{4.39}$$

$$Dv + Bw = 1, \tag{4.40}$$

$$Du + Bv = 0, \tag{4.41}$$

that can be solved to give

$$u = \frac{B}{AB - D^2}, \quad v = \frac{-D}{AB - D^2}, \quad w = \frac{A}{AB - D^2}. \tag{4.42}$$

Therefore

$$C^{-1} = \frac{-1}{\sigma_a^2(AB - D^2)} \begin{pmatrix} B & 0 & -D & 0 \\ 0 & B & 0 & -D \\ -D & 0 & A & 0 \\ 0 & -D & 0 & A \end{pmatrix}. \tag{4.43}$$

Hence using equations 4.19 and 4.43 an expression for $X^T C^{-1} X$ can be found as follows

$$X^T C^{-1} X = \frac{-1}{\sigma_a^2 (AB - D^2)} X^T \begin{pmatrix} Bx_1 - Dx_2 \\ By_1 - Dy_2 \\ Ax_2 - Dx_1 \\ Ay_2 - Dy_1 \end{pmatrix}$$

$$= \frac{1}{\sigma_a^2 (AB - D^2)} \left[B(x_1^2 + y_1^2) + A(x_2^2 + y_2^2) - 2D(x_1 x_2 + y_1 y_2) \right]. \quad (4.44)$$

We can now replace A with $(1 + \alpha^2)$, B with $j^2 (\alpha^2 + [k(1 - \alpha) + \alpha])$ and D with $j\alpha [1 + k(1 - \alpha) + \alpha]$. This allows equation 4.44 to be written as

$$X^T C^{-1} X = \frac{1}{\sigma_a^2 j^2 [\alpha^2 + \alpha(k - 1) - k]^2} \left\{ j^2 \left(\alpha^2 + [k(1 - \alpha) + \alpha]^2 \right) (x_1^2 + y_1^2) \right.$$
$$\left. + (1 + \alpha^2)(x_2^2 + y_2^2) - 2j\alpha [1 + k(1 - \alpha) + \alpha] (x_1 x_2 + y_1 y_2) \right\}. \quad (4.45)$$

We can now define p as the probability that, in a given frequency bin, the real amplitude x_i is greater than some threshold ξ_i and the imaginary amplitude y_i is greater than some threshold χ_i. In practice these amplitudes will be set to high values so that p is small. The above equations show that

$$p = \frac{1}{(2\pi)^2 \sigma_a^4 j^2 [\alpha^2 + \alpha(k - 1) - k]^2} \int_{y_2 = \chi_2}^{\infty} \int_{y_1 = \chi_1}^{\infty} \int_{x_2 = \xi_2}^{\infty} \int_{x_1 = \xi_1}^{\infty}$$
$$\exp \left\{ \frac{1}{2\sigma_a^2 j^2 [\alpha^2 + \alpha(k - 1) - k]^2} \left[2j\alpha [1 + k(1 - \alpha) + \alpha] (x_1 x_2 + y_1 y_2) \right. \right.$$
$$\left. \left. - j^2 \left(\alpha^2 + [k(1 - \alpha) + \alpha]^2 \right) (x_1^2 + y_1^2) - (1 + \alpha^2)(x_2^2 + y_2^2) \right] \right\} dx_1 dx_2 dy_1 dy_2. \quad (4.46)$$

Substitutions for $x_i = \rho_i \cos \theta_i$ and $y_i = \rho_i \sin \theta_i$ can be made to give

$$p = \frac{1}{(2\pi)^2 \sigma_a^4 j^2 [\alpha^2 + \alpha(k - 1) - k]^2} \int_{\theta_2 = 0}^{2\pi} \int_{\theta_1 = 0}^{2\pi} \int_{\rho_2 = r_2}^{\infty} \int_{\rho_1 = r_1}^{\infty}$$
$$\exp \left\{ \frac{1}{2\sigma_a^2 j^2 [\alpha^2 + \alpha(k - 1) - k]^2} \left[2j\alpha [1 + k(1 - \alpha) + \alpha] \rho_1 \rho_2 \cos(\theta_1 - \theta_2) \right. \right.$$

$$-j^2 \left(\alpha^2 + [k(1-\alpha) + \alpha]^2\right)\rho_1^2 - (1+\alpha^2)\rho_2^2)\bigg] \bigg\} \rho_1\rho_2 \mathrm{d}\rho_1 \mathrm{d}\rho_2 \mathrm{d}\theta_1 \mathrm{d}\theta_2. \qquad (4.47)$$

Now p gives the probability that, in a given frequency bin, ρ_i is greater than some threshold r_i for all θ_i between 0 and 2π radians. As we are interested in the difference between θ_1 and θ_2 only we can change these variables to the difference between the two angles, θ. Hence, we can substitute $\theta = \theta_1 - \theta_2$ into equation 4.47 and multiply equation 4.47 by 2π, to give

$$p = \frac{1}{2\pi\sigma_a^4 j^2 \left[\alpha^2 + \alpha(k-1) - k\right]^2} \int_{\theta=0}^{2\pi} \int_{\rho_2=r_2}^{\infty} \int_{\rho_1=r_1}^{\infty}$$

$$\exp\left\{ \frac{1}{2\sigma_a^2 j^2 \left[\alpha^2 + \alpha(k-1) - k\right]^2} \left[2j\alpha\left[1 + k(1-\alpha) + \alpha\right]\rho_1\rho_2\cos\theta \right.\right.$$

$$\left.\left. -j^2\left(\alpha^2 + [k(1-\alpha) + \alpha]^2\right)\rho_1^2 - (1+\alpha^2)\rho_2^2)\right]\right\}\rho_1\rho_2\mathrm{d}\rho_1\mathrm{d}\rho_2\mathrm{d}\theta. \qquad (4.48)$$

The Modified Bessel Function of the First Kind is defined, in the integral form, as (Abramowitz and Stegun, 1972)

$$I_0(z) = \frac{1}{\pi}\int_0^\pi \exp\{\pm z\cos\theta\}\mathrm{d}\theta. \qquad (4.49)$$

$I_0(z)$ may also be expressed as a summation:

$$I_0(z) = 1 + \frac{\frac{1}{4}z^2}{(1!)^2} + \frac{(\frac{1}{4}z^2)^2}{(2!)^2} + \frac{(\frac{1}{4}z^2)^3}{(3!)^2} + \cdots. \qquad (4.50)$$

Comparing equation 4.50 with equation 4.48 shows that

$$I_0 = 1 + \left(\frac{\alpha\left[1 + k(1-\alpha) + \alpha\right]\rho_1\rho_2}{2\sigma_a^2 j\left[\alpha^2 + \alpha(k-1) - k\right]^2}\right)^2 + \frac{1}{2^2}\left(\frac{\alpha\left[1 + k(1-\alpha) + \alpha\right]\rho_1\rho_2}{2\sigma_a^2 j\left[\alpha^2 + \alpha(k-1) - k\right]^2}\right)^4$$

$$+ \frac{1}{(3!)^2}\left(\frac{\alpha\left[1 + k(1-\alpha) + \alpha\right]\rho_1\rho_2}{2\sigma_a^2 j\left[\alpha^2 + \alpha(k-1) - k\right]^2}\right)^6 \cdots,$$

$$= \sum_{n=0}^{\infty}\frac{1}{(n!)^2}\left(\frac{\alpha\left[1 + k(1-\alpha) + \alpha\right]\rho_1\rho_2}{2\sigma_a^2 j\left[\alpha^2 + \alpha(k-1) - k\right]^2}\right)^{2n}. \qquad (4.51)$$

Substituting this equation into equation 4.48 gives

$$p = \frac{1}{\sigma_a^4 j^2 \left[\alpha^2 + \alpha(k-1) - k\right]^2} \int_{r_1}^{\infty} \int_{r_2}^{\infty} \sum_{n=0}^{\infty} \frac{1}{(n!)^2} \rho_1 \rho_2 \left(\frac{\alpha\left[1 + k(1-\alpha) + \alpha\right]\rho_1\rho_2}{2\sigma_a^2 j \left[\alpha^2 + \alpha(k-1) - k\right]^2}\right)^{2n}$$

$$\exp\left\{\frac{-1}{2\sigma_a^2 j^2 \left[\alpha^2 + \alpha(k-1) - k\right]^2}\left[j^2 \left(\alpha^2 + [k(1-\alpha) + \alpha]^2\right)\rho_1^2\right.\right.$$

$$\left.\left.+(1 + \alpha^2)\rho_2^2\right]\right\}\mathrm{d}\rho_1\mathrm{d}\rho_2. \tag{4.52}$$

It is now possible to integrate this expression term by term. Let

$$A_1 = \frac{\left(\alpha^2 + [k(1-\alpha) + \alpha]^2\right)}{2\sigma_a^2 \left[\alpha^2 + \alpha(k-1) - k\right]^2}, \tag{4.53}$$

and

$$A_2 = \frac{(1 + \alpha^2)}{2\sigma_a^2 j^2 \left[\alpha^2 + \alpha(k-1) - k\right]^2}. \tag{4.54}$$

The $n = 0$ term of equation 4.52 is given by

$$p_{n=0} = \frac{1}{\sigma_a^4 j^2 \left[\alpha^2 + \alpha(k-1) - k\right]^2} \int_{r_2}^{\infty} \exp\{-A_2\rho_2^2\}\rho_2\mathrm{d}\rho_2 \int_{r_1}^{\infty} \exp\{-A_1\rho_1^2\}\rho_1\mathrm{d}\rho_1. \tag{4.55}$$

Breaking this equation down, it can be shown that

$$\int_{r_i}^{\infty} \exp\{-A_i\rho_i^2\}\rho_i\mathrm{d}\rho_i = \frac{1}{2A_i}\exp\{-A_ir_i^2\}, \tag{4.56}$$

and so

$$p_{n=0} = \frac{1}{\sigma_a^4 j^2 \left[\alpha^2 + \alpha(k-1) - k\right]^2}\frac{1}{4A_1A_2}\exp\{-A_1r_1^2 - A_2r_2^2)\}. \tag{4.57}$$

The $n = 1$ term of equation 4.52 is given by

$$p_{n=1} = \frac{1}{\sigma_a^4 j^2 \left[\alpha^2 + \alpha(k-1) - k\right]^2}\left(\frac{\alpha\left[1 + k(1-\alpha) + \alpha\right]}{2\sigma_a^2 j \left[\alpha^2 + \alpha(k-1) - k\right]^2}\right)^2$$

$$\times \int_{r_2}^{\infty} \rho_2^3 \exp\{-A_2\rho_2^2\}\mathrm{d}\rho_2 \int_{r_1}^{\infty} \rho_1^3 \exp\{-A_1\rho_1^2\}\mathrm{d}\rho_1. \tag{4.58}$$

Now, with the help of equation 4.56, it can be shown that

$$
\begin{aligned}
\int_{r_i}^{\infty} \rho_i^3 \exp\{-A_i \rho_i^2\} \mathrm{d}\rho_i &= \frac{r_i^2}{2A_i} \exp\{-A_i r_i^2\} + \frac{1}{A_i} \int_{r_i}^{\infty} \rho_i \exp\{-A_i \rho_i^2\} \mathrm{d}\rho_i \\
&= \frac{1}{2A_i} \left(r_i^2 + \frac{1}{A_i} \right) \exp\{-A_i r_i^2\}.
\end{aligned}
\tag{4.59}
$$

Therefore

$$
\begin{aligned}
p_{n=1} &= \frac{1}{\sigma_a^4 j^2 \left[\alpha^2 + \alpha(k-1) - k\right]^2} \left(\frac{\alpha\left[1 + k(1-\alpha) + \alpha\right]}{2\sigma_a^2 j \left[\alpha^2 + \alpha(k-1) - k\right]^2} \right)^2 \\
&\quad \times \frac{1}{4A_1 A_2} \left(r_1^2 + \frac{1}{A_1} \right) \left(r_2^2 + \frac{1}{A_2} \right) \exp\{-A_1 r_1^2 - A_2 r_2^2\}.
\end{aligned}
\tag{4.60}
$$

Similarly

$$
\begin{aligned}
p_{n=2} &= \frac{1}{\sigma_a^4 j^2 \left[\alpha^2 + \alpha(k-1) - k\right]^2} \frac{1}{4} \left(\frac{\alpha\left[1 + k(1-\alpha) + \alpha\right]}{2\sigma_a^2 j \left[\alpha^2 + \alpha(k-1) - k\right]^2} \right)^4 \\
&\quad \times \int_{r_2}^{\infty} \int_{r_1}^{\infty} \rho_1^5 \rho_2^5 \exp\{-A_1 \rho_1^2 - A_2 \rho_2^2)\} \mathrm{d}\rho_1 \mathrm{d}\rho_2 \\
&= \frac{1}{4} \left(\frac{\alpha\left[1 + k(1-\alpha) + \alpha\right]}{2\sigma_a^2 j \left[\alpha^2 + \alpha(k-1) - k\right]^2} \right)^4 \frac{1}{4A_1 A_2} \left(\frac{r_1^4}{2} + \frac{r_1^2}{A_1} + \frac{1}{A_1^2} \right) \\
&\quad \times \left(\frac{r_2^4}{2} + \frac{r_2^2}{A_2} + \frac{1}{A_2^2} \right) \exp\{-A_1 r_1^2 - A_2 r_2^2\},
\end{aligned}
\tag{4.61}
$$

and

$$
\begin{aligned}
p_{n=3} &= \frac{1}{\sigma_a^4 j^2 \left[\alpha^2 + \alpha(k-1) - k\right]^2} \frac{1}{36} \left(\frac{\alpha\left[1 + k(1-\alpha) + \alpha\right]}{2\sigma_a^2 j \left[\alpha^2 + \alpha(k-1) - k\right]^2} \right)^6 \\
&\quad \times \int_{r_2}^{\infty} \int_{r_1}^{\infty} \rho_1^7 \rho_2^7 \exp\{-A_1 \rho_1^2 - A_2 \rho_2^2)\} \mathrm{d}\rho_1 \mathrm{d}\rho_2 \\
&= \frac{1}{36} \left(\frac{\alpha\left[1 + k(1-\alpha) + \alpha\right]}{2\sigma_a^2 j \left[\alpha^2 + \alpha(k-1) - k\right]^2} \right)^2 \frac{1}{4A_1 A_2} \left(\frac{r_1^6}{2} \frac{3r_1^4}{2A_1} + \frac{3r_1^2}{A_1^2} + \frac{3}{A_1^3} \right) \\
&\quad \times \left(\frac{r_2^6}{2} \frac{3r_2^4}{2A_2} + \frac{3r_2^2}{A_2^2} + \frac{3}{A_2^3} \right) \exp\{-A_1 r_1^2 + A_2 r_2^2\}.
\end{aligned}
\tag{4.62}
$$

More generally it can be shown that

$$p_{n=n} = \frac{1}{\sigma_a^4 j^2 \left[\alpha^2 + \alpha(k-1) - k\right]^2} \frac{1}{(n!)^2} \left(\frac{\alpha\left[1 + k(1-\alpha) + \alpha\right]}{2\sigma_a^2 j \left[\alpha^2 + \alpha(k-1) - k\right]^2}\right)^{2n}$$

$$\frac{1}{4A_1 A_2} \left[\sum_{m=0}^{n} \frac{n! r_1^{2(n-m)}}{(n-m)! A_1^m}\right] \left[\sum_{q=0}^{n} \frac{n! r_2^{2(n-q)}}{(n-q)! A_2^q}\right] \exp\left\{-A_1 r_1^2 - A_2 r_2^2\right\}. \qquad (4.63)$$

By combining equations 4.57, 4.60, 4.61, 4.62 and 4.63 an expression for the probability that a spike in the first set of data has an amplitude greater than r_1 and that a spike in the same frequency bin in the second set of data has an amplitude greater than r_2 can be found:

$$p = \sum_{n=0}^{\infty} \frac{1}{(n!)^2} \left(\frac{\alpha\left[1 + k(1-\alpha) + \alpha\right]}{2\sigma_a^2 j \left[\alpha^2 + \alpha(k-1) - k\right]^2}\right)^{2n} \frac{\left[\alpha^2 + \alpha(k-1) - k\right]^2}{(1+\alpha^2)\left(\alpha^2 + \left[k(1-\alpha) + \alpha\right]^2\right)}$$

$$\left[\sum_{m=0}^{n} \frac{n r_1^{2(n-m)} 2^m \sigma_a^{2m} \left[\alpha^2 + \alpha(k-1) - k\right]^{2m}}{(n-m)(\alpha^2 + \left[k(1-\alpha) + \alpha\right]^2)^m}\right] \left[\sum_{q=0}^{n} \frac{n r_2^{2(n-q)} 2^q \sigma_a^{2q} j^{2q} \left[\alpha^2 + \alpha(k-1) - k\right]^{2q}}{(n-q)(1+\alpha^2)^q}\right]$$

$$\exp\left[-\frac{\left(\alpha^2 + \left[k(1-\alpha) + \alpha\right]^2\right) r_1^2}{2\sigma_a^2 \left[\alpha^2 + \alpha(k-1) - k\right]^2} - \frac{(1+\alpha^2) r_2^2}{2\sigma_a^2 j^2 \left[\alpha^2 + \alpha(k-1) - k\right]^2}\right]. \qquad (4.64)$$

Equation 4.64 converges to a limit as $n \to \infty$. It was found that, at most, 20 terms (so $n = 19$) were needed to find an accurate value of p.

To ensure that this equation correctly predicts the probability some simulations for different values of α were performed. The results of one such simulation are shown here. This example takes $k = 2$, $j = 2$, $\sigma^2 = 1$, $r_1 = 2 < \rho_1 >$ and $r_2 = 2 < \rho_2 >$. Figure 4.1 shows that the predictions and the simulations for this example are in good agreement.

Equation 4.64 gives the probability that prominent spikes that are found in the same frequency bin in each of two frequency-amplitude spectra. The probability of observing at least 1 pair of prominent spikes in a frequency range Δ_p is

$$P = 1 - (1-p)^{\Delta_p T}, \qquad (4.65)$$

where T is the length of the timeseries. In equation 4.64 α accounts for the amount of

Figure 4.1: The probability that spikes in the same frequency bin in each of two spectra have amplitudes greater than twice the mean level for different values of α. In this simulation $k = 2$, $j = 2$ and $\sigma^2 = 1$. The crosses represent the probabilities found by simulations and the black line represents the probabilities found using Equation 4.64.

common noise in the two frequency-amplitude spectra. We will now relate α to the more commonly used measure, coherency.

4.2.1 Comparison between coherency and α

It is useful to compare α with the more commonly used coherency. The coherency of two sets of data is defined as (from equation 1.10, Chapter 1)

$$Y_{1,2} = \frac{\langle (x_1 + iy_1) \cdot (x_2 - iy_2) \rangle}{\sqrt{\langle (x_1 + iy_1) \cdot (x_1 - iy_1) \rangle \langle (x_2 + iy_2) \cdot (x_2 - iy_2) \rangle}}. \qquad (4.66)$$

Using equations 4.22 to 4.30 it can be shown that

$$\langle (x_1 + iy_1) \cdot (x_2 - iy_2) \rangle = 2j\alpha \left[1 + k(1 - \alpha) + \alpha \right] \sigma_a^2, \qquad (4.67)$$

$$\langle (x_1 + iy_1) \cdot (x_1 - iy_1) \rangle = 2(1 + \alpha^2)\sigma_a^2, \qquad (4.68)$$

and

$$\langle (x_2 + iy_2) \cdot (x_2 - iy_2) \rangle = 2j^2 \left(\alpha^2 + [k(1 - \alpha) + \alpha]^2 \right) \sigma_a^2. \qquad (4.69)$$

Therefore, the coherency is given by

$$Y_{1,2} = \frac{\alpha\left[1 + k(1-\alpha) + \alpha\right]}{\sqrt{(1+\alpha^2)(\alpha^2 + [k(1-\alpha)+\alpha]^2)}}, \tag{4.70}$$

which can be rearranged to give α for known values of k. Notice that j and σ_a are not in this equation because of the normalization.

We now have an equation for the probability of observing coincident prominent spikes, which is related to the coherency of the two frequency-amplitude spectra. The coherency of the spectra is related to α and k. Before we can determine the probability that a pair of prominent spikes are noise we must determine α, k, j and σ_a. We therefore move on to discuss how these parameters can be found from the data.

4.2.2   Calculating the variables in equation 4.64

Equation 4.64 contains four variables that need to be found from the data, namely α, k, j and σ_a^2. These parameters may not be constant throughout the frequency-amplitude spectra and so we will now show how these variables can be found for a frequency range of width Δ_p. Let V_1 be the variance of the real part of the first spectrum, defined by equation 4.14. Since x_1 is the sum of two Gaussian distributions it has a zero mean and the variance is given by the mean of x_1^2:

$$V_1 = \langle x_1^2 \rangle = (1 + \alpha^2)\sigma_a^2. \tag{4.71}$$

Let V_2 be the variance of the real part of the second spectrum, defined by equation 4.16. As x_2 is also the sum of two Gaussian distributions V_2 is given by

$$V_2 = \langle x_2^2 \rangle = j^2(\alpha^2 + [k(1-\alpha)+\alpha]^2)\sigma_a^2. \tag{4.72}$$

In practice the variances may be estimated by measurement of the mean power in each spectrum in the $N = \Delta_p T$ bins being searched. Additional information is provided by

calculation of the mean cross-amplitude (Elsworth et al., 1994), over the same N bin slice. The mean cross-amplitude is given by

$$
\begin{aligned}
X_{1,2} &= \langle (x_1 + iy_1) \cdot (x_2 - iy_2) \rangle, \\
&= 2j\alpha\sigma_a^2(1 + k[1 - \alpha] + \alpha).
\end{aligned}
\tag{4.73}
$$

We now have three equations but four variables that need to be found. However, if the data comes from observations made using a similar method it is reasonable to assume that the overall level of noise between the two spectra is of approximately the same order of magnitude. This implies that the magnitude of the common noise observed by the two instruments is the same. Note that the magnitude of the observed independent noise can still be different. If both sets of observations come from Doppler velocity measurements or if they both come from intensity measurements this is a reasonable assumption. Figure 4.2 shows the ratio of the mean absolute amplitude observed in contemporaneous BiSON and GOLF data. The ratio of the observed absolute amplitudes is always between 2.5 and 1.0 in the frequency range of interest in this thesis (below $\sim 1500\,\mu\mathrm{Hz}$), and so it is reasonable to assume that the overall noise level is of the same order of magnitude. Therefore, any variations in the total noise level between the two sets of observations occur because the amount of independent noise in the data is different i.e. $k \neq 1$. If this is the case we can assume that $j = 1$. Therefore equations 4.71, 4.72 and 4.73 may then be used to find values for α, σ_a^2 and k. Equation 4.71 may be rearranged to find σ_a^2 in terms of α and V_1:

$$
\sigma_a^2 = \frac{V_1}{1 + \alpha^2}.
\tag{4.74}
$$

Equation 4.74 may be substituted into equations 4.72 and 4.73 to give

$$
V_2 = \left(\alpha^2 + [k(1 - \alpha) + \alpha]^2 \right) \frac{V_1}{1 + \alpha^2},
\tag{4.75}
$$

Figure 4.2: Ratio of the mean absolute amplitude of 8.5years of BiSON and GOLF data.

and

$$X_{1,2} = \frac{2\alpha V_1}{1 + \alpha^2} \left(1 + k\left[1 - \alpha\right] + \alpha\right).$$ (4.76)

From equation 4.76 it can be shown that

$$k\left(1 - \alpha\right) = \frac{X_{1,2}\left(1 + \alpha^2\right)}{2\alpha V_1} - 1 - \alpha,$$ (4.77)

which may be substituted into equation 4.75 to give

$$V_2 = \frac{4\alpha V_1\left(1 + \alpha^2\right)\left(\alpha V_1 - X_{1,2}\right) + X_{1,2}^2(1 + \alpha^2)^2}{4\alpha^2 V_1\left(1 + \alpha^2\right)}.$$ (4.78)

Now the only unknown left in equation 4.78 is α. Equation 4.78 may be used to find a quartic in terms of α:

$$\left(4V_1^2 + X_{1,2}^2 - 4V_1 V_2\right)\alpha^4 - 4V_1 X_{1,2}\alpha^3 + \left(4V_1^2\right.$$

$$\left. + 2X_{1,2}^2 - 4V_1 V_2\right)\alpha^2 - 4V_1 X_{1,2}\alpha + X_{1,2}^2 = 0.$$ (4.79)

This equation has four roots, however, by definition α should take a value between 0 and 1. Therefore, when calculating α we have found the real root of equation 4.79 that lies between 0 and 1. Use of equations 4.74 and 4.77 then yields estimates of σ_a^2 and k respectively.

We finish by stressing that these statistics are relevant for searching frequency-amplitude

spectra but not frequency-power spectra, as was done when looking at the spectra individually. This is because the statistics have been derived assuming real and imaginary normally distributed noise, which is the case for frequency-amplitude spectra. However, the noise in a power spectrum has a χ^2 2 d.o.f distribution. These statistics can now be applied to the frequency-amplitude spectra of contemporaneous sets of data, such as BiSON and GOLF data, to search for low-frequency modes.

4.3 Summary

We began this chapter by discussing the probability of observing prominent spikes in the same frequency bin in each of two frequency-power spectra (equation 4.1). This equation is based on the assumption that the noise in the two frequency-power spectra is independent. However, this is unlikely to be the case if the data are contemporaneous. The amount and source of any common noise shared by the data vary depending on the instruments that are used to make the observations. However, there will at least be some solar noise, from the solar granulation, which is common to both sets of data.

Equation 4.64 gives the probability that a coincident, prominent spike in two frequency-amplitude spectra is due to noise. Derivation of the probability is non trivial and it takes proper account of the level of common noise in the frequency-amplitude spectra. The amount of common noise is given by α, which is related to the coherency (equation 4.70). The probability that a coincident prominent spike is due to noise (equation 4.64) is also related to k, which determines the difference in the amount of other noise, which is not coherent between the two frequency-amplitude spectra. The other parameters in equation 4.64 are σ_a, which determines the variance of the noise in the two frequency-amplitude spectra, and j, which allows the total amount of noise in the two frequency-amplitude spectra to be different.

To determine the probability that a coincident prominent spike is due to noise we need to know the values of α, k, σ_a and j. Assuming the observations that are being

compared are both Doppler velocity measurements or both intensity measurements we can take $j = 1$. We have then shown that α, k and σ_a can be found from the data using the variances of the real parts of the two frequency-amplitude spectra and the cross-amplitude spectrum.

So, in this chapter we have derived an equation that can be used to search for coincident prominent spikes in two frequency-amplitude spectra that contain a known amount of common noise. In the next chapter we use this probability to compare BiSON and GOLF data with the aim of detecting low-frequency p modes and possibly even g modes. We will also use equation 4.64 to compare two sets of contemporaneous BiSON data (Chapter 7) and to discuss the visibility of low-frequency p modes (Chapter 6).

5. RESULTS OF SEARCHING BISON AND GOLF DATA

The previous two chapters have discussed various statistical properties of the noise in solar power (Chapter 3) and amplitude spectra (Chapter 4). The statistics developed in Chapter 3 can be used to calculate the probability that a prominent feature in a frequency-power spectrum could occur by chance, while the statistics derived in Chapter 4 can be used to search for coincident prominent features in a pair of frequency-amplitude spectra. Therefore, statistical techniques can be used to search the frequency-power and frequency-amplitude spectra of Sun-as-a-star data for features that are statistically significant compared to the background noise, with the aim of uncovering potential mode candidates. Sun-as-a-star data are taken by the BiSON network and the GOLF instrument onboard SOHO. Various statistical tests have been applied to BiSON and GOLF data. This chapter goes on to describe how these statistical tests were applied and the results of the tests. However, we begin by describing the data used.

5.1 Observational data

BiSON and GOLF timeseries were created for the Phoebus collaboration, which is an international group of helioseismologists who aim to detect low-frequency solar g modes. Each timeseries contains 3071d of Sun-as-a-star Doppler velocity observations that were made between 20[th] April 1996 and 15[th] September 2004. BiSON data are normally stored with a cadence of 40s while GOLF data usually has a cadence of 20s. The data were re-binned to a common cadence of 120s. This cadence was chosen so that the BiSON and GOLF data were also compatible with other sets of data that have a cadence of 60s. For the BiSON data this was done by taking the average over three original 40s bins.

However, this average was found only when all three bins contained useful data. Take, for example, the first three time bins in the original BiSON data set. If at least one of the three bins contained null data the newly binned data point was set to zero. This only decreased the fill of the timeseries very slightly as, when creating timeseries from BiSON data, if a single point of data is missing this point is interpolated. Hence there were very few instances when data were lost because not all of the bins contained useful data. The re-binned timeseries contain 2,211,120 points and the bin width in the frequency spectrum is 3.77 nHz. The duty cycle of the BiSON timeseries is 78.6%, while the duty cycle of the GOLF timeseries is 93.4%. The frequency spectra of these timeseries were then searched for prominent spikes. The spectra were both compared with each other and searched individually. We will now describe the tests that were applied to search for low-frequency modes.

5.2 The tests applied to the data

Frequency-amplitude spectra were created from both the GOLF and the BiSON timeseries. Both frequency-amplitude spectra were searched in slices of 100μHz for the presence of prominent spikes or prominent patterns of spikes that occurred in the same frequency bin in each spectrum. BiSON and GOLF frequency-power spectra were also searched individually. When the BiSON and GOLF frequency-power spectra were searched individually the tests that were applied to the FLAG data in Section 3.5 were again applied here. The noise in a frequency-power spectrum exhibits a χ^2 2 d.o.f. distribution while the noise in a frequency-amplitude spectrum is distributed normally. The statistics governing the tests described in Section 3.5 assume a χ^2 2d.o.f. distribution. Hence it was the frequency-power spectra that were searched when the BiSON and GOLF data were considered separately.

When it comes to comparing the spectra we will use the statistics described in Chapter 4. These statistics assume that the noise is distributed normally and is complex, as is the

case in a frequency-amplitude spectrum. Therefore, when the BiSON and GOLF data are compared to look for coincident prominent features it was the frequency-amplitude spectra that were searched.

A brief summary of the tests applied when comparing the data will now be given. The tests that were applied when comparing the two spectra are similar to, but not identically the same as the tests described in Section 3.5.

Test 1: The frequency-amplitude spectra were searched for prominent, individual spikes that were found in the same frequency bin of each frequency-amplitude spectrum. We call this a 'pair of spikes'. Substituting the observed amplitudes into equation 4.64 gives the probability of observing such a pair of spikes by chance in the same frequency bin of each frequency-amplitude spectrum. The probability that this pair of spikes occurs at least once over a range of N bins is then given by equation 3.4.

Test 2: A test was performed to search for two prominent spikes in the same two consecutive bins in each frequency-amplitude spectrum. The amplitude of the least prominent of the two spikes in each frequency-amplitude spectrum was used in equation 4.64 to calculate p, the probability of a pair of spikes occurring by chance. This was then used with equations 3.18 and 3.19 to find the probability that a peak is due to noise.

Test 3: The spectra were searched for two or more pairs of prominent spikes that were separated by less than twice the predicted width of the mode. The width of a mode was predicted by extrapolating to low frequencies the widths of well-observed higher-frequency modes. Each prominent spike must, once again, lie in the same frequency bin in each frequency-amplitude spectrum. Equations 4.64, 3.16 and 3.17 can be used to determine the probability that any such structures are due to noise. It should be noted that this test is only different to Test 2 above $\sim 950\,\mu\mathrm{Hz}$, as below this frequency the extrapolation implies that the lifetime of a mode is greater than the length of the timeseries. Therefore the width of a mode with a frequency

below $\sim 950\,\mu$Hz is not resolved and all of the power should be contained in one bin.

Test 4: A test was performed to search for multiplets that contained two spikes. Each constituent spike in the multiplet must lie in the same frequency bin in each frequency-amplitude spectrum. It was assumed that the rotationally split components of a mode are separated by $400\pm100\,$nHz in frequency (see Section 3.4 for more details). However, as the only components visible in Sun-as-a-star data are components for which $l + m$ is even, for two spikes to constitute a multiplet they must be separated by multiples of $800 \pm 100\,$nHz. If the multiplet's spikes are separated by $800 \pm 100\,$nHz they could belong to an $l = 1$, 2 or 3 mode. If the spikes are separated by $1600 \pm 100\,$nHz they could belong to an $l = 2$ or 3 mode. If the spikes are separated by $2400 \pm 100\,$nHz they could be the outer components of an $l = 3$ mode. The probability that each multiplet is due to noise can be calculated using equations 4.64, 3.22, 3.23 and 3.24.

Test 5: This test searched for multiplets containing three spikes. Once again the spikes that formed the triplet must be in the same frequency bins in each frequency-amplitude spectrum. The components of the triplet must again be separated by $800\pm100\,$nHz. Therefore this test can only observe $l = 2$ and 3 modes. Equations 4.64, 3.25 and 3.26 can be used to find the probability that each triplet would occur at least once in $100\,\mu$Hz of pure noise.

The threshold levels were set so that the probability of any noise detections in a $100\,\mu$Hz range was less than 0.01. This lower threshold was used following the results in Section 3.5, which showed that reducing the threshold probability to 1% from 10% significantly reduced the number of false detections. Once again the bin shifting strategy described in Section 3.5 was used to scan the frequency bins in an attempt to force the frequency bins to be commensurate with the signal from a mode. The same number of points were removed from the BiSON and the GOLF data. As before, 25

differing numbers of points were removed from the end of both timeseries. This time the smallest number of points removed from the end of the timeseries was 15, while the largest number of points removed was 1458.

5.3   Finding α, k and σ_a^2

As stated in Section 4.2.2 there are three variables required to calculate the probability that any prominent spikes are noise. These variables need to be found from the data. The parameters are α, which provides a measure of the amount of noise that is common to the two sets of data, k, which allows one of the sets of data to contain more independent noise than the other, and σ_a^2, which provides a measure of the overall level of noise in the data. The procedure for finding these parameters is also described in Section 4.2.2. Let the first spectrum described in Section 4.2 represent the BiSON data and the second spectrum described represent the GOLF data. So x_1 and y_1 represent the real and imaginary parts of the BiSON frequency-amplitude spectrum respectively while x_2 and y_2 represent the real and imaginary parts of the GOLF frequency-amplitude spectrum. To find α, k and σ_a^2 it was necessary to find the variance of the real amplitude in the BiSON and the GOLF spectra and these are V_1 and V_2 respectively (see equations 4.71 and 4.72). It was also necessary to compute the BiSON-GOLF cross-amplitude spectrum (equation 4.73). Estimates of the variances of the two frequency-amplitude spectra and the cross-amplitude can be seen in Figure 5.1. The BiSON data have a much higher variance than the GOLF data at low frequencies, however, above $\sim 1000 \mu$Hz the variances become very similar. At low frequencies BiSON spectra become contaminated by artifacts related to the window-function. In particular, the daily harmonics are very prominent at low frequencies in a BiSON spectrum. Also there is a known artefact in BiSON data that is thought to occur because BiSON instruments are placed on an equatorial mount, which simplifies the automation of the instruments, but produces a quasi-oscillatory signal that is particularly visible below 1000μHz. A low-frequency function can be fitted to

the data to remove much of the unwanted low-frequency noise but some evidence of of the artefact still remains (Chaplin et al., 2002a). A cross-amplitude spectrum will suppress any uncommon noise and it gives an indication as to the amount of common noise present in two sets of data. The level of the cross-amplitude spectrum decreases with increasing frequency implying that the amount of common, solar noise is reduced at higher frequencies.

The values found for α, using equation 4.79, can be seen in the upper panel of Figure 5.2. As can be seen the shape of α is closely related to the shape of the coherency (red dashed line), although the values found for α are consistently smaller than the coherency. The middle panel of Figure 4.79 shows the values found for k. As the variance of the BiSON data is larger than the GOLF data at low frequencies k is significantly less than 1. However, as the variances of the noise in the two data sets converge k approaches a value of unity. If the amount of uncommon noise in the two data sets was exactly the same k would equal unity. The bottom panel of Figure 4.79 shows that the shape of σ_a follows closely the shape of V_1 and V_2 (shown in Figure 5.1). This is expected as σ_a is a measure of the level of noise in the two sets of data and we have already seen that the amount of noise in the BiSON and GOLF amplitude spectra decreases with increasing frequency.

As an aside, one assumption that is crucial to determining the probability that any coincident spikes are due to noise (equation 4.64, section 4.2) is that the variances of the real and imaginary parts are equal. Figure 5.3 shows the ratio of the variances of the real and imaginary amplitudes over the range of frequencies searched in both the BiSON data (top panel) and the GOLF data (bottom panel). The average ratio over regions of $100\,\mu$Hz has been plotted. The ratio of the real and imaginary variances is always between 0.98 and 1.02 for both the BiSON data and the GOLF data in the frequency range of interest here (i.e. below $\sim 1500\,\mu$Hz). This implies that it is valid to assume that the variances of the real and imaginary parts are equal.
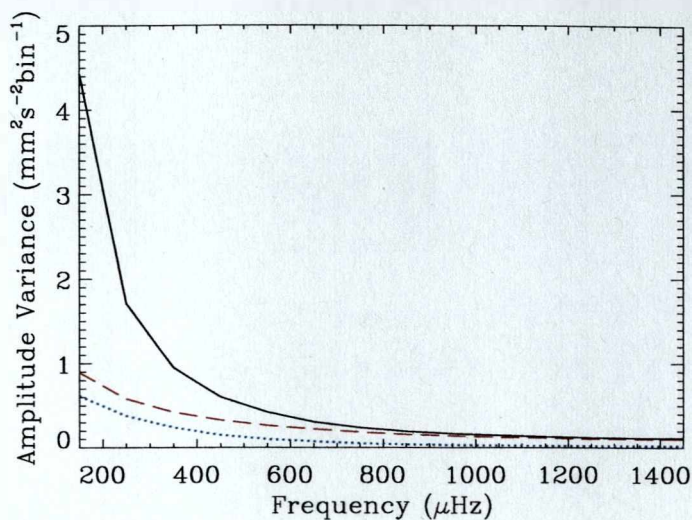
Figure 5.1: The variances of the real part of each frequency-amplitude spectrum and the cross-amplitude spectrum are plotted. The black solid line shows the variance of the BiSON data, which is equivalent to V_1 in Section 4.2. The red dashed line represents the variance of the GOLF data, which is equivalent to V_2. The blue dotted line shows the cross-amplitude of the two spectra, which is equivalent to $X_{1,2}$ in equation 4.73.

5.4   Results

The left panel of Figure 5.4 shows the absolute threshold levels (in units of mm s^{-1} per root bin) for the BiSON data, as a function of frequency, for each of the five tests. The threshold levels have been calculated for a 1% chance of passing the tests at least once anywhere in 100μHz. The levels are specific to the BiSON data even though the thresholds are calculated using the joint BiSON and GOLF probabilities, as the plotted thresholds are dependent on the mean amplitude level, which is different for the BiSON and GOLF data. The right panel of Figure 5.4 shows the percentage reduction in the threshold amplitudes with respect to the amplitudes given by the analysis of the BiSON data alone. The reduction in threshold levels is clearly frequency dependent. Nonetheless use of two data sets has allowed the threshold levels to be reduced by $\sim 28\%$. The shape of this graph can be compared to the results for α in Figure 5.2. When α is lower the decrease in the threshold levels is larger. Therefore, the amount of noise common to both sets of data is crucial in determining the effectiveness of comparing two sets of data.

The best-case scenario, leading to the largest reduction in threshold levels, would

Figure 5.2: The upper panel shows how α varies with frequency (the black line). To allow a comparison to be made the coherency for the same frequency region has been plotted as a red dashed line. The middle panel shows k for different frequency ranges. The lower panel shows the variation of σ_a with frequency.
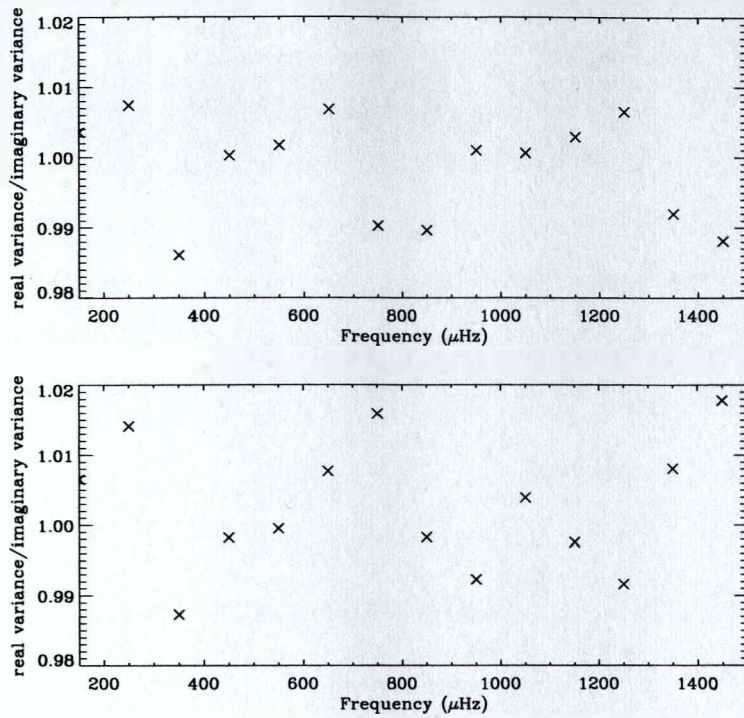
Figure 5.3: Top panel: The ratio of the variances of the real and imaginary parts of the BiSON amplitude spectrum for different frequencies. Bottom panel: The ratio of the variances of the real and imaginary parts of the GOLF amplitude spectrum for different frequencies.
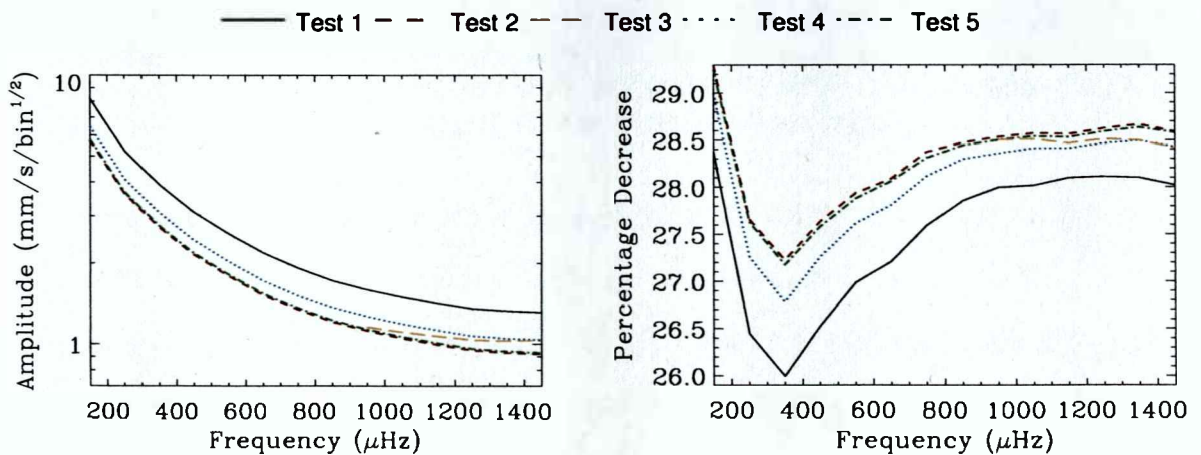
Figure 5.4: Left panel: threshold amplitudes (for BiSON) needed to record $P \leq 1\%$ in each of the five tests (see figure legend). Right panel: percentage reduction in the threshold amplitudes, with respect to the amplitudes given by analysis of the BiSON data alone.
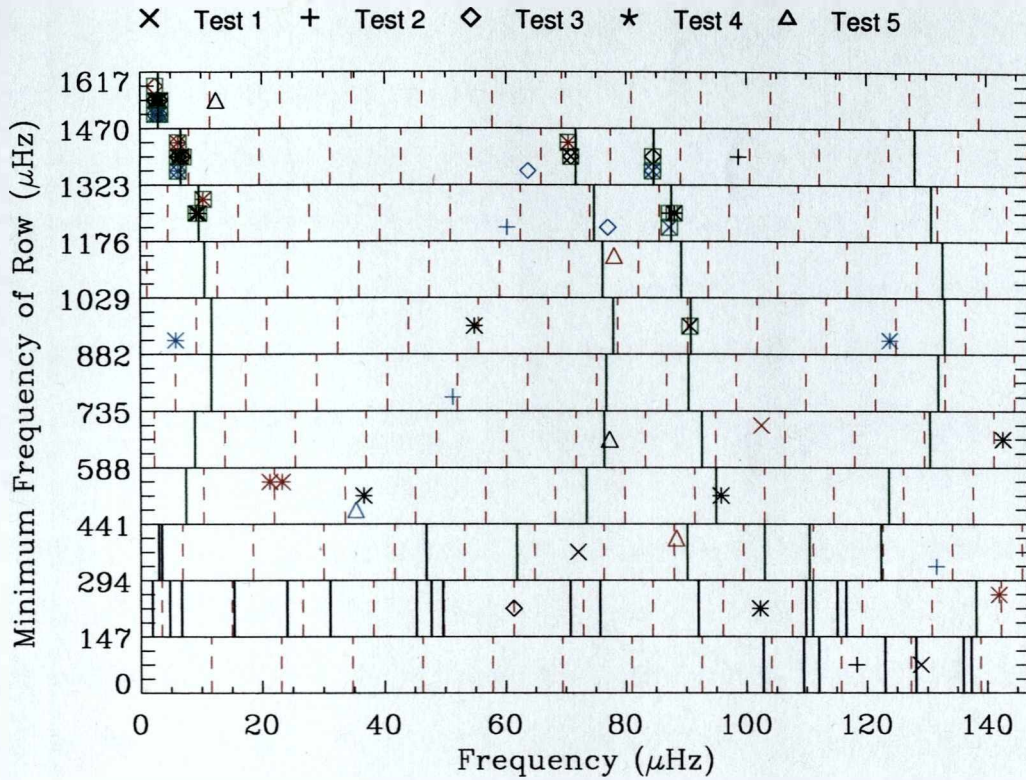
Figure 5.5: An Echelle plot, modulo 147 μHz, marking the locations in frequency of occurrences uncovered by the test searches. Locations in frequency where spikes, or patterns of spikes, were found in the same bin, or bins, of each spectrum at levels sufficient to record $P \leq 1\%$ are marked by the black symbols in the middle of each row. A different symbol has been used for each test (see figure legend). We have also recorded prominent spikes or patterns of spikes found by searching both the BiSON frequency-power spectrum (red symbols at top of each row) or the GOLF frequency-power spectrum (blue symbols at bottom of each row) alone. Symbols surrounded by a green square represent the prominent occurrences listed in Table 5.1. The green vertical lines mark locations of the p-mode frequencies predicted by the Saclay seismic model (Turck-Chièze et al., 2001). The navy blue vertical lines show the g-mode frequencies predicted by the M1 model from Provost et al. (2000). The vertical dashed red lines mark locations in frequency that are overtones of the 11.57 μHz diurnal frequency.

occur when $\alpha = 0$. In other words it is most advantageous to compare two completely independent frequency-amplitude spectra with no common background noise. At a particular frequency, ν, the threshold height in the BiSON frequency-amplitude spectrum is given by some value times the mean noise level in the BiSON data, s_ν, say. At the same frequency the threshold height in the GOLF data is given by the mean level of noise times this same value, s_ν. Therefore, when $\alpha = 0$ the threshold amplitudes would be reduced by a factor $\sqrt{2}$, implying a $\sim 29\%$ reduction in size.

Figure 5.5 is an Echelle diagram that shows the locations in frequency where promi-

nent spikes or patterns of spikes were found. Strips of length 147 μHz are placed one above another. The strips are 147μHz in length as consecutive overtones of low-l p modes are separated in frequency by approximately this amount. The green vertical lines indicate the frequencies of the modes predicted by the Saclay seismic model (Turck-Chièze et al., 2001). As can be seen the choice of the lengths of the strips has allowed these green lines to be arranged in four near-vertical strips, one strip for each l in the range $0 \leq l \leq 3$. The navy blue vertical lines represent the g-mode frequencies predicted by the M1 model (Provost et al., 2000).

The vertical red dashed lines indicate the frequencies of the daily harmonics that occur at multiples of 11.57μHz. These harmonics are particularly prominent in the low-frequency BiSON data. The daily harmonics are caused partially by periodic gaps in the data that are separated by 24 hrs and partially due to the fitting procedure that aims to remove the daily effect of the Earth's spin from the Doppler velocity residuals. There are also diurnal frequencies present in the GOLF data due to spacecraft operation signatures.

The black symbols in the middle of each row represent the frequencies of spikes or patterns of spikes that were found in the same bin in each of the frequency-amplitude spectra, which had less than a 1% chance of occurring by chance. Each of the different symbols show results for the different tests. Also shown are the frequencies of prominent spikes or patterns of spikes that passed the tests when the BiSON data were searched alone (represented by the red symbols at the top of each row) and the spikes or patterns of spikes that passed the tests when the GOLF data were searched alone (represented by the blue symbols at the bottom of each row). The symbols surrounded by a green square indicate that the frequency of the feature is less than 0.5 μHz from the location of one of the frequencies predicted by the Saclay seismic and M1 models (Turck-Chièze et al., 2001; Provost et al., 2000). The prominent spikes or patterns of spikes that lie within 0.5 μHz of a model frequency have been recorded in Table 5.1. The occurrences listed in Table 5.1 have been identified with particular p-mode components on the basis of their placement in frequency. The errors quoted in Table 5.1 have been calculated in

the manner described in Section 3.5.

When the outer components of a rotationally split mode have been found the mean of these components has been calculated to give an estimate of the frequency of the centroid of that mode. However, this can only be taken as an estimate because taking the straight mean ignores the possibility that the components are not an equal distance in frequency from the centre. During epochs of high solar magnetic activity it has been observed that the components of the $l = 2$ modes are not distributed symmetrically (Chaplin et al., 2003b). However, the data examined here covers epochs of both low and high activity and so this effect will be diminished. Furthermore the Moreno-Insertis and Solanki (2000) model implies that the observed asymmetry decreases with decreasing frequency. Moreno-Insertis and Solanki modelled the magnetic field near the surface. The altitude of the upper turning point of a mode's path through the solar interior decreases with mode frequency. In the Moreno-Insertis and Solanki model low-frequency modes sample less of the region of the solar interior that contains the model magnetic field. Therefore the presence of the magnetic field has a smaller effect on the observed frequencies of low-frequency modes than modes at higher frequencies. These predictions appear to be supported by observations of both BiSON and GOLF data (Chaplin et al., 2003b). As only frequencies less than 1500μHz were considered here it is unlikely that any multiplet asymmetry will be significant. Therefore, at these frequencies, and, given the length of the observations, it is reasonable to simply take the mean of the two components.

No new candidates were found but all of the listed candidates correspond well with previous claimed detections of p modes such as those made in Toutain et al. (1998), Bertello et al. (2000), García et al. (2001a), Chaplin et al. (2002a). The $l = 1$, $n = 7$ mode was not detected by Chaplin et al. (2002a). However, they did detect a component of the $l = 2$, $n = 7$ mode at $1250.564\,\mu$Hz in BiSON data, which is not seen here. However, closer inspection reveals Chaplin et al. found that the $m = 2$ component of the $l = 2$, $n = 7$ passed the single spike test with a probability of 0.04, which is above the threshold probability of 0.01 that was set here. Furthermore, the observed S/N of a mode

Table 5.1: Candidates found to be closer than 0.5μHz to predicted frequencies of modes from the Saclay seismic model and the M1 model.

l	n	m	Frequency (μHz)	Probability (P)	Number of tests passed	Distance from model in frequency (μHz)
0	6	0	972.613 ± 0.002	3.4×10^{-4}	2	0.132
1	7	-1	1185.196 ± 0.005	1.9×10^{-3}	4	
1	7	$+1$	1185.981 ± 0.005	3.0×10^{-3}	1	
		mean	1185.589 ± 0.004			0.027
0	8	0	1263.210 ± 0.007	6.0×10^{-4}	3	0.314
1	8	-1	1329.236 ± 0.005	0.0000	4	
1	8	$+1$	1330.037 ± 0.007	1.7×10^{-3}	1	
		mean	1329.637 ± 0.004			0.060
2	8	-2	1393.871 ± 0.007	4.0×10^{-4}	1	0.036†
0	9	0	1407.479 ± 0.007	4.0×10^{-8}	3	0.149
1	9	-1	1472.432 ± 0.008	3.8×10^{-11}	4	
1	9	$+1$	1473.269 ± 0.009	3.7×10^{-8}	4	
		mean	1472.851 ± 0.006			0.122

†: Difference with model frequency assumes $m = -2$ component lies $-0.8\,\mu$Hz from central frequency.

can be altered by beating with the solar noise background. The interaction between the mode and the noise can act both constructively, so as to increase the S/N of the mode, or destructively, so as to decrease the S/N of the mode. It is, therefore, possible that the S/N of a mode in two different sets of BiSON data will be different. This means that the mode can be detected in one set of data but not the other. It should be noted that the power in $l = 2$ modes is suppressed in GOLF observations because of the response of the instrument over the observed solar disc. This makes $l = 2$ modes difficult to detect in GOLF data.

Clearly some of the prominent features that are observed lie away from the model frequencies. The majority of these lie close to the diurnal frequencies demonstrating the possible impact of the daily harmonics on both sets of data. However, despite the reduced level of the thresholds allowed by comparing the data there are still a large number of modes that remain undetected. Chapter 6 goes on to investigate whether this is consistent with the number of detections expected from simulations based on predictions of powers and damping rates. These predictions will be made by extrapolating the powers and

widths present in well-observed modes to low frequencies.

Turck-Chièze et al. (2004) found a prominent structure in GOLF data with 5 components at $\sim 220\,\mu$Hz. They postulate that this structure could be a detection of a g mode, possibly the $l = 2$, $n = -3$ mode or the $l = 3$, $n = -5$ mode, or a combination of components from the two modes. These modes are predicted to have frequencies at $222.02\,\mu$Hz and $217.07\,\mu$Hz respectively, while the detected components have frequencies of $218.31\,\mu$Hz, $218.95\,\mu$Hz, $219.59\,\mu$Hz, $220.12\,\mu$Hz and $221.26\,\mu$Hz. No prominent structures were detected at or close to these frequencies here.

Although the statistical methods used by Turck-Chièze et al. (2004) are fundamentally the same as the statistics outlined here the actual application of these statistics to the data is different in many ways. Firstly no searches were performed for quintuplets here for reasons that follow. In order to observe a quintuplet whose components are signals from one mode only the rotation axis of the core must be different to the rotation axis of the rest of the Sun. This possibility was not considered here. BiSON and GOLF both take Sun-as-a-star observations, which means that the measurements are averaged over the whole disc. The visibility of a mode, therefore, depends on the average of its displacement pattern over the whole disc. As a result high-degree modes are averaged out as their positive and negative fluctuations cancel. In Doppler velocity observations made by a resonance scattering instrument, such as the BiSON instruments, this means that there is no net shift in the wavelength of the observed spectral line. Therefore, in full-disc observations only $l = 0$, 1, 2, and 3 modes are clearly visible. There is evidence for $l = 4$ and 5 modes at higher frequencies but the observed powers of these modes are very small. When observations are made close to the ecliptic plane, as is the case for the BiSON and GOLF instruments, only those components with $l + m$ even are visible. Therefore, the maximum number of components that are visible in a multiplet in Sun-as-a-star data is 4 (the $m = -3$, -1, +1, +3 components of an $l = 3$ mode). Furthermore the visibility of the different components is sensitive to azimuthal order, with the outermost components of any multiplet being the most prominent. In BiSON data the $l = 1$,

$m = \pm 2$ peaks are 1.8 times more prominent in power than their $m = 0$ counterpart. Similarly $m = \pm 3$ components are 2.6 times larger than the $m = \pm 1$ components in an $l = 3$ mode (Chaplin et al., 2004b). Therefore, even though it is possible that an $l = 4$ mode might just be visible it is likely that only the outer components will be prominent enough to be observed. Moreover, the separation of the components of the quintuplets that Turck-Chièze et al. (2004) search for is ~ 400 nHz. Therefore, they are searching for modes with consecutive m. In this thesis it was assumed that only modes with $l + m$ even would be visible and so the separation of visible components was taken to be a multiple of ~ 800 nHz.

The threshold level was set by Turck-Chièze et al. (2004) so there is a 10% chance of making a false detection in a frequency range of $\Delta_p = 20\,\mu\text{Hz}$. A more stringent threshold of a 1% chance of making a false detection was used here. Additionally we have looked at frequency ranges of $\Delta_p = 100\,\mu\text{Hz}$. While restricting the frequency range to $\Delta_p = 20\,\mu\text{Hz}$ does allow a lower threshold level to be set it increases the total number of Δ_p being searched. For example, in the frequency range 0 to $1500\,\mu\text{Hz}$ the number of $100\,\mu\text{Hz}$ bins is 15 but the number of $20\,\mu\text{Hz}$ bins in $1500\,\mu\text{Hz}$ is 75. Therefore, in this thesis, where the threshold was set for a 1% chance of a false detection in $100\,\mu\text{Hz}$, the number of false detections expected when searching a frequency range of $1500\,\mu\text{Hz}$ should be less than 1 (0.15 to be precise). However, in Turck-Chièze et al. (2004), where the threshold was set for a 10% chance of making a false detection in $20\,\mu\text{Hz}$, one would expect at least 7 false detections when searching a frequency range of $1500\,\mu\text{Hz}$ (7.5 to be precise).

Therefore the threshold levels set in this thesis require the power or amplitude of any modes to be larger than in Turck-Chièze et al. (2004) for a prominent structure to be considered as a candidate mode. However, it was decided that ensuring as few false detections as possible are made was paramount. In hypothesis testing a Type I error occurs when the null hypothesis is wrongly rejected. Here that would mean wrongly rejecting the hypothesis that a prominent structure is due to noise. In other words saying a prominent structure is not due to noise. A Type II error occurs when the null

hypothesis is not rejected when, in fact, it is false. In the statistics used in this thesis this would mean accepting that a feature is just noise when it is not. Generally, Type I errors are considered to be more serious than Type II errors. For this reason, in this thesis, we have opted to err on the side of caution and set very stringent requirements that need to be passed before a structure is considered to be a mode candidate.

The photon noise level in GOLF data has been steadily increasing over time as the instrument's components, mainly the photomultiplier tubes, age. García et al. (2005) reported that GOLF's photon counting rate had decreased by a factor of ~ 4.5. This means that the photon noise is becoming increasingly important at the low-frequency end of a spectrum of GOLF data. As the data used in Turck-Chièze et al. (2004) are from a slightly earlier epoch than the data used here the fact that the structure observed by Turck-Chièze et al. is not detected in this thesis could simply be due to the ever decreasing quality of the GOLF data. Therefore the structure could now be swamped by noise that was less prominent over the time span Turck-Chièze et al. searched.

Garcia et al. (2007) detect a significant periodic structure that is compatible with the periodic separation predicted by the theory for gravity modes. The separation in periods between g modes with the same l and consecutive n is approximately constant and Garcia et al. (2007) take advantage of this to look for periodic structures in the power spectrum. To do this they take the Fourier transform of the frequency-power spectrum of GOLF data between 25 and $140\,\mu$Hz to produce a periodogram. They find a broad structure in the region centred on a period of ~ 24mins that they show to be statistically significant using Monte Carlo simulations. The period of this feature is synonymous with the expected separation between dipolar g modes from $n = -4$ to -26. Figure 5.5 shows that no detections are made in the region 25 to $140\,\mu$Hz here. Simulations have been performed (by W.J. Chaplin) to determine how the detectability of modes varies as a function of lifetime and S/N. Artificial spectra containing $l = 1$ modes were created. The $l = 1$ modes included in the spectra were given a period spacing that matched the period separation described in Garcia et al.. These simulations determined that modes with the
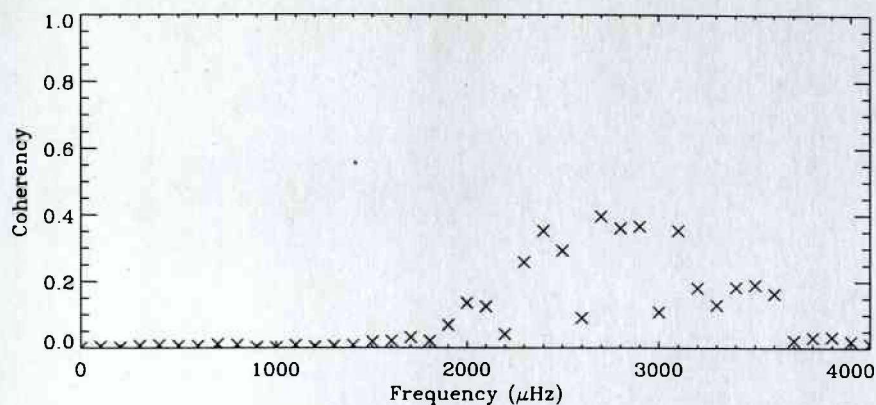
Figure 5.6: The coherency of the BiSON and GOLF data found when the GOLF observations start 24hrs after the BiSON observations.

lifetimes and S/N ratios similar to the modes detected by Garcia et al. should be seen in our spectra. Therefore these simulations imply that our results contradict those found by Garcia et al..

5.5 Use of non-contemporaneous data

Until now contemporaneous BiSON and GOLF data have been compared as this means that the signal from the modes should be common to both sets of data. However, as the data are contemporaneous the solar noise, from granulation, is also common to the two data sets. A small shift in the start time of one of the data sets would mean that the data are no longer contemporaneous and so the amount of low-frequency common noise could be reduced. As we have seen in Chapter 4 the advantage gained by using two sets of data is greatest when the level of common noise is zero. Furthermore at low frequencies the lifetimes of the modes are similar to, if not longer than, the length of the observations. Therefore, if only a small shift in the start times is implemented, the majority of the signal from low-frequency modes should remain commensurate. This is demonstrated in more detail in Section 7.2.1.

Figure 5.6 shows the coherency of the BiSON and GOLF data used in the previous section. However, the data from the first day in the GOLF timeseries has been removed. Therefore the GOLF observations commence 24 hrs after the BiSON data. As the length

Figure 5.7: An echelle diagram showing the results when the start of the GOLF observations is 24hrs after the start of the BiSON observations. The lines and symbols have the same meaning here as in Figure 5.5.

of the timeseries still needs to be the same for both timeseries the data from the last day of observations have been removed from the end of the BiSON data. The data now contain 3070 days of data or 2,210,400 time bins with a cadence of 120s. Figure 5.6 shows that separating the start times of the BiSON and GOLF data has reduced the low-frequency coherency to zero. Again, this is because the solar noise, which was the source of the low-frequency coherent noise in the contemporaneous data, will no longer be coherent. It is also interesting to notice that the coherency of the five-minute modes has also been reduced. The amount by which the coherency of a particular mode is reduced is related to the lifetime of that mode. This will be covered in more detail later, in Section 7.2.1. Suffice to say that as the low-frequency modes, which we are searching for in this thesis, have significantly longer lifetimes than the five-minute modes they should still be coherent even when the start times of the two timeseries are separated by 24hrs. The methods of searching for the modes described in Section 5.2 were applied to these new

Table 5.2: Candidates found to be closer than 0.5μHz to predicted frequencies of modes from the Saclay seismic model or the M1 model when the GOLF observations start 24hrs after the BiSON observations.

l	n	m	Frequency (μHz)	Probability (P)	Number of tests passed	Distance from model in frequency (μHz)
1	1	-1	284.665 ± 0.004	6.1×10^{-3}	1	0.065†
0	6	0	972.613 ± 0.002	1.4×10^{-4}	1	0.132
1	7	-1	1185.196 ± 0.005	1.4×10^{-5}	3	
1	7	$+1$	1185.981 ± 0.005	1.9×10^{-3}	1	
		mean	1185.589 ± 0.004			0.027
0	8	0	1263.210 ± 0.007	7.8^{-4}	1	0.314
1	8	-1	1329.236 ± 0.005	0.0000	4	
1	8	$+1$	1330.037 ± 0.007	1.0×10^{-3}	1	
		mean	1329.637 ± 0.004			0.060
2	8	-2	1393.871 ± 0.007	2.0×10^{-4}	1	0.036‡
0	9	0	1407.479 ± 0.007	9.4×10^{-4}	3	0.149
1	9	-1	1472.432 ± 0.008	0.0000	4	
1	9	$+1$	1473.269 ± 0.009	2.3×10^{-5}	4	
		mean	1472.851 ± 0.006			0.122

†: Difference with model frequency assumes $m = -1$ component lies $-0.4\,\mu$Hz from central frequency.

‡: Difference with model frequency assumes $m = -2$ component lies $-0.8\,\mu$Hz from central frequency.

timeseries, including the bin shifting strategy.

The results of this search can be seen in Figure 5.7. Once again the results are displayed in the form of an echelle diagram. The candidates shown have less than a 1% chance of being due to noise. The symbols and vertical lines have the same meaning as in Figure 5.5. Candidates that lie within $0.5\,\mu$Hz of the predicted frequencies are again highlighted by a green square. These candidates are also shown in Table 5.2. There is one new candidate found at $284.665 \pm 0.004\,\mu$Hz. This is close to the frequency predicted by the Saclay seismic model for the $m = -1$ component of the $l = 1$, $n = 1$ mode, which is expected to be found at $284.730\,\mu$Hz. This candidate passes Test 2 with a probability that it is due to noise of 6.1×10^{-3}. However, this does not necessarily mean that the candidate is a mode. For us to have more confidence it should pass more than one test because, as shown in Section 3.5, this requirement significantly reduces the number of false

detections. It should be pointed out that this mode is detected by García et al. (2001b) in GOLF data as a single peak, with a 10% confidence interval. However, this does not improve confidence in the detection as it just indicates that this spike is prominent in GOLF data, which were used in both analyses.

5.6 Summary

In this chapter we have used the statistical tests developed in Chapters 3 and 4 to search BiSON and GOLF data for low-frequency p modes and g modes. More precisely, we have searched the BiSON and GOLF frequency-power spectra individually for prominent spikes/structures of spikes that are statistically unlikely to be part of the background noise. We have also compared contemporaneous BiSON and GOLF frequency-amplitude spectra for statistically prominent features that occur in the same frequency bin in each spectrum.

When comparing the frequency-amplitude spectra three variables must be determined from the data before the statistics can be calculated. These variables are α, which provides a measure of the amount of common noise present in the data, k, which allows one set of data to contain more independent noise than the other and σ_a, which acts as a measure of the overall level of noise. The variables were found using the variances of the real parts of the BiSON and GOLF frequency-amplitude spectra and the cross-amplitude. The values of α, k and σ_a all vary with frequency and so are specific to the region of the frequency-amplitude spectra that are being searched.

The amplitude thresholds levels were calculated for a 1% chance of a false detection in $100\,\mu$Hz. When comparing the frequency-amplitude spectra the threshold levels for the five different tests that were performed were all less than $10\,\mathrm{mm\,s^{-1}\,bin^{1/2}}$ for the BiSON data. This constitutes a $\sim 28\%$ reduction in threshold levels compared to the threshold amplitudes required when searching one spectrum alone. The threshold levels can be reduced by the greatest amount if there is no common noise shared by the two

sets of data at low frequencies. In this case the threshold levels can be reduced by $\sim 29\%$. This is only slightly more than the reduction achieved here. All of the mode candidates that were found correspond to previously claimed detections of p modes (see for example García et al. 2001a, Chaplin et al. 2002a). However, despite the reduction in threshold levels no new mode candidates were revealed.

In addition to searching the contemporaneous data we have also shifted the start time of the GOLF timeseries so that it begins 24 hrs after the start of the BiSON timeseries, meaning that the data are no longer contemporaneous. This reduced the coherency of the noise at low frequencies to zero, which is the optimal condition as the thresholds are lowest. One new potential candidate is found at $\sim 285\,\mu$Hz. However, this candidate only passed one of the tests and so cannot be considered as a detection.

Clearly there are gaps in the list of candidate detections. We know that the S/N ratios of low-frequency modes are small. Although we cannot yet observe many low-frequency modes we can estimate their powers and widths based on extrapolations to lower frequencies of the powers and widths of observed modes with large S/N ratios. In the next chapter we go on to investigate whether the number of candidates observed here is consistent with these estimates of the powers and widths of modes or whether it is a puzzle that, given the decrease in threshold levels obtained by looking for coincidences, we have not observed more low-frequency modes.

6. THE VISIBILITY OF LOW-FREQUENCY P MODES

To date a multitude of solar acoustic (p) modes have been observed over a wide range of frequencies. However, no independently confirmed detections of low-degree (low-l) p modes, with frequencies below $\sim 973\,\mu$Hz, have been made (e.g. García et al. 2001a, Chaplin et al. 2002a, Broomhall et al. 2007 and Chapter 5 of this thesis).

In Chapter 5 we used statistical techniques to search for coincident prominent features in BiSON and GOLF frequency-amplitude spectra. We found that despite reducing amplitude detection thresholds to less than $10\,\mathrm{mm\,s^{-1}bin^{-1/2}}$ many low-frequency p modes remained undetected. However, we did find the $l = 0$, $n = 6$ mode at $\sim 973\,\mu$Hz to be extremely prominent. The $l = 0$, $n = 6$ mode has also been detected in other studies such as García et al. (2001a) and Chaplin et al. (2002a). The prominence of this mode is conspicuous because of the lack of evidence for higher-frequency modes such as the $l = 0$, $n = 7$ mode at $\sim 1118\,\mu$Hz. As yet the $l = 0$, $n = 7$ mode has only been observed in Sun-as-a-star data by García et al. (2001a). These higher-frequency modes should be easier to observe as theoretically they should exhibit larger amplitudes in the photosphere, where the observations are made. It is therefore of interest to investigate how prominent the $l = 0$, $n = 6$ mode, and the modes surrounding it in frequency, should be. This will provide an indication as to whether we should be observing modes that we are not; or whether some effect makes the $l = 0$, $n = 6$ mode more prominent than the rest.

How easy low-frequency modes are to detect depends upon their power and width in a frequency-power spectrum. As many low-frequency modes have not yet been detected their powers cannot be determined from solar data directly. Here, predictions of the power in very low-frequency modes have been made by extrapolating the power observed

in well-defined, higher-frequency modes. The height of a mode in a frequency-power (amplitude) spectrum also depends upon the width in frequency over which the power is spread. As the power in the signal from a mode is damped over time the signal from the mode may be spread over several bins of a frequency spectrum. The lifetime of a mode varies with frequency and so it is also possible to infer from extrapolation the width that a mode at a particular frequency is expected to exhibit from results for well-observed modes. We begin this chapter by discussing how the extrapolations were made. To make the extrapolations we have assumed that the simple functional relationships that turn out to describe the powers and heights of the well-observed modes are still valid at low frequencies. The validity of this assumption was tested using predictions of the powers and widths of low-frequency modes made by the Cambridge stochastic excitation and damping codes (e.g. Houdek et al., 1999).

6.1 Predicting the widths and powers of modes

Various properties of modes can be found by fitting frequency-power spectra. These properties include the width of a mode, and therefore its damping time, and the height of a mode, i.e. its maximum power spectral density. The product of the height and width is proportional to the total power in the mode. A BiSON spectrum, consisting of 3071 d of Sun-as-a-star Doppler velocity observations, was fitted using the methods described in Fletcher (2007). This procedure involved fitting a Lorentzian-like model to the various mode peaks in a frequency-power spectrum. The data were the same as the BiSON data used in Chapter 5 except that, instead of being rebinned to have a cadence of 120 s, the data were left with the nominal 40 s cadence on which the BiSON data are stored.

The visibility of a mode in a frequency-power spectrum will depend on the power and lifetime of the mode. We will now briefly discuss how these properties affect the visibility.

6.1.1 How the power and width of a mode affect its visibility

The width of a mode, $\Delta\nu$, is related to its lifetime, τ, by

$$\Delta\nu = \frac{1}{\pi\tau}. \tag{6.1}$$

In this thesis we refer to the maximum power spectral density per bin of a mode as its height, H. When the power spectrum is fitted it is actually the maximum power spectral density, H, which corresponds to the height of the Lorentzian, that is determined. The height of the most prominent spike of a resonant peak in a frequency-power spectrum will be greater than the height of the fitted Lorentzian. This is because of the random nature of the excitation of the modes, which means that the power will have a χ^2 2d.o.f. distribution about the underlying Lorentzian. Therefore, the power in some of the bins across the width of the mode will be greater than the height of the Lorentzian. These are the spikes that are most likely to be detected.

If the length of a timeseries, T, is significantly longer than the lifetime of a mode, τ, the timeseries will extend over several realizations of the mode and so the modal peak will be resolved across several bins in the frequency domain. However, if a mode's lifetime is significantly longer than the length of a timeseries all of the mode's power will be contained in a single bin[1]. It is, therefore, necessary to use two different equations to describe H that depend on whether or not the width of the mode can be resolved. Hence, H is given by

$$H = \begin{cases} 2V^2\,(\tau/T) & \text{for } T >> 2\tau \\ V^2 & \text{for } T << 2\tau \end{cases} \tag{6.2}$$

where V^2 is the total power of the mode. However, this does not describe H when $T \sim 2\tau$. Fletcher et al. (2006) used Monte Carlo simulations to investigate the height, H, when

---

[1] The power can be split between 2 bins if the frequency of the mode is not commensurate with the spectrum's frequency bins. Also the finite length of a timeseries means that the mode will appear in a frequency-power spectrum as a sinc squared function.
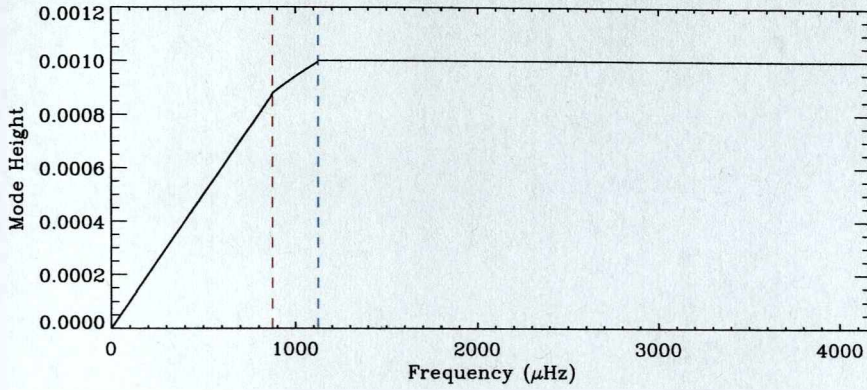
Figure 6.1: The height of a mode, as given by equations 6.2 and 6.3. Here we have taken $T = 2\tau$ at $1000\,\mu$Hz. The red and blue vertical lines highlight sharp changes in the gradient of the slope as we change which equation is used to calculate the height. Below the red vertical line $T \ll 2\tau$, between the two vertical lines $T \approx 2\tau$ and above the vertical blue line $T \gg 2\tau$.

$T \sim 2\tau$ and found that H is well modelled by the following equation:

$$H = \frac{2V^2}{T/\tau + 2}.$$

(6.3)

It should be noted that it is the power that we are going to extrapolate and not the peak height of the mode. This is because the height of the mode is not a smooth function of frequency, as it depends on whether the power of the mode is spread over more than one bin or whether the mode's power is confined to one bin only. On the other hand the power of a mode is a smooth function of frequency and so can be extrapolated more easily. This can be clarified by taking a simple example where the power of a mode, V^2, is approximately equal to the frequency of the mode. Figure 6.1 then shows how the height varies with frequency in this example. The height is given by equations 6.2 and 6.3. We have arbitrarily taken $T = 2\tau$ at $1000\,\mu$Hz and τ to be inversely proportional to the mode frequency. Clearly the height of a mode, H, is not a smooth function of frequency in spite of the fact that V^2 is. Therefore, when determining the properties of low-frequency modes it is better to extrapolate the powers, V^2, of modes rather than their height, H. The width, $\Delta\nu$, and height, H, of the mode can be used to find the power of the mode.

A well-defined mode is one with a power and a width that are sufficiently large for
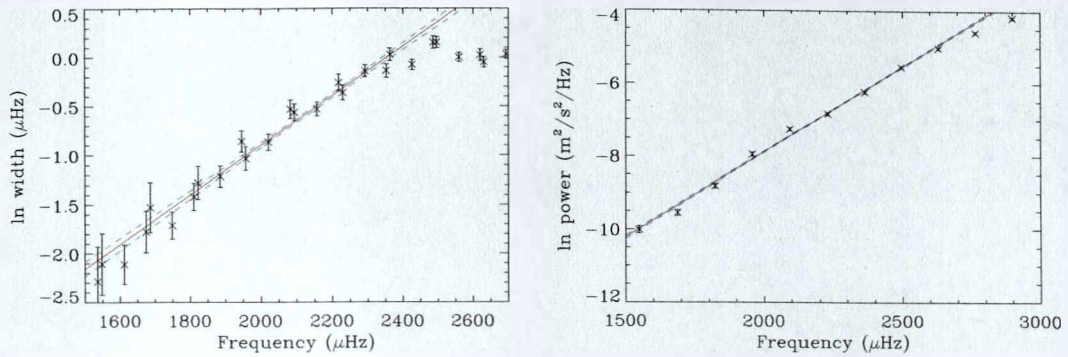
Figure 6.2: Results for method 1. Left panel: The black crosses show how the natural logarithm of the observed width of a mode varies with frequency. The errors of the widths are those associated with the fits. The red solid line shows the linear line of best fit for the widths at frequencies between $\sim 1530\,\mu$Hz and $\sim 2400\,\mu$Hz. The blue dashed lines indicate the errors on this linear fit. Right panel: The black crosses show how the natural logarithm of a mode's power varies with frequency. The red solid line shows the linear line of best fit for the powers at frequencies between $\sim 1540\,\mu$Hz and $\sim 2500\,\mu$Hz. The blue dashed lines indicate the errors on this fit.

the mode to be detected easily, allowing the shape of the mode in a frequency-power spectrum to be fitted accurately. In the case of the Sun, when fitting a set of data that spans ~ 8.5 yr, well-defined modes have frequencies greater than $\sim 1500\,\mu$Hz. Here we are concerned with modes at low frequencies that cannot be observed clearly. Therefore it was necessary to extrapolate the widths and powers of well-defined modes to obtain estimates of the widths and powers of low-frequency modes. We have assumed that simple functional relationships, which describe the variation of the parameters with frequency at higher frequencies, persist to lower frequencies. In section 6.1.5 we have used predictions from a stochastic excitation model to test the validity of this assumption. Extrapolations and simulations were performed to determine how often $l = 0$, 1 and 2 modes can be detected in the artificial data. However, the general trends in the results were very similar for each l and so in the rest of this chapter we will only show the results found for $l = 0$ modes. The extrapolation can be performed in two different ways, each of which will now be described in turn.

6.1.2 Method 1: The ln-linear relationship

An approximately linear relationship is observed between the natural logarithm of a mode's width and its frequency for well-defined modes with frequencies below $\sim 2400\,\mu\text{Hz}$ (see the left panel of Figure 6.2). The well-known plateau in the mode widths above $\sim 2400\,\mu\text{Hz}$ (see for example Chaplin et al. 2005) means that the linearity does not extend to higher frequencies. The observed widths have been fitted using data from modes with different degrees. In Sun-as-a-star observations only low-l modes can be clearly observed. Here widths from $l = 0$, 1 and 2 modes have been used. The width of a mode is not independent of degree as high-l modes have a lower inertia than low-l modes and so high-l modes have faster damping rates. However, the difference in damping rates is minimal over the confined range of l used here.

We now move on to discuss in more detail how we used the fitted values of the height, H, and width, $\Delta\nu$, of a mode to determine the power, V^2, of a mode. If the width of a mode can be resolved the power of that mode, V^2, can be obtained from its height, H, and width, $\Delta\nu$:

$$V^2 = \frac{\pi}{2} T \Delta\nu H, \tag{6.4}$$

where T is the length of the timeseries. Here H is in units of power per bin (rather than power per Hz) but the width, $\Delta\nu$, has units of Hz. Therefore, to give the width in terms of the number of bins, rather than Hz, we have multiplied $\Delta\nu$ by a factor of T. This equation has been obtained by rearranging the upper part of equation 6.2 (when $T >> 2\tau$). Modes with frequencies greater than $1500\,\mu\text{Hz}$ have a well-resolved width in the frequency-power spectrum that was fitted to determine the heights and widths of the modes. We can, therefore, use equation 6.4 to determine the powers of the modes from the fitted heights and widths. Equation 6.4 can be expressed in logarithmic form as

$$\ln V^2 = \ln \frac{\pi}{2} + \ln \Delta\nu + \ln H. \tag{6.5}$$

Let $p = \ln P$, $w = \ln \Delta \nu$ and $h = \ln H$. The variance of the natural logarithm of the power is given by

$$\sigma_p^2 = \sigma_h^2 + \sigma_w^2 - 2\rho(w, h)\sigma_h \sigma_w, \qquad (6.6)$$

where ρ is the correlation coefficient and is given by

$$\rho = \frac{cov(w, h)}{\sigma_w \sigma_h} = \frac{\overline{wh} - \overline{w}\,\overline{h}}{\sigma_w \sigma_h}. \qquad (6.7)$$

Notice that this term is subtracted in equation 6.6 as the coherency in the height and width should reduce the error of the power. The width and height of a mode have a negative correlation of $\rho \approx 0.95$ and so the resulting error bars on the calculated powers are small.

The relationship between the natural logarithm of a mode's power and its frequency is also approximately linear for well-defined modes with frequencies below $\sim 2500\,\mu$Hz (see the right panel of Figure 6.2). The power of a mode is dependent on its degree and so the values of the coefficients of the linear relationship are specific to each value of l. The right panel in Figure 6.2 shows the results for $l = 0$ modes only.

For both the power and the width a linear, least-squares fit was performed to determine the gradients and the zero-frequency intercepts. The gradients and zero-frequency intercepts of these fits and the errors associated with them can be seen in Table 6.1. This method of fitting the data will be referred to as method 1.

6.1.3 Method 2: The ln-ln relationship

An approximately linear relationship is also observed between the natural logarithm of the width of a mode and the natural logarithm of its frequency (see the left panel of Figure 6.3). Once again the linearity is only found at frequencies below $\sim 2400\,\mu$Hz. Likewise an approximately linear relationship is found when the natural logarithm of the power of a mode is plotted against the natural logarithm of the frequency of that mode (see the right panel of Figure 6.3). As with method 1 a linear least squares fit was

Figure 6.3: Results for method 2. Left panel: The black crosses show how the natural logarithm of the observed width of a mode varies with the natural logarithm of its frequency. The errors bars shown for the widths are those associated with the fits. Right panel: The black crosses show how the natural logarithm of the mode power varies with the natural logarithm of mode frequency. As in Figure 6.2 the red solid lines in both panels represent the linear best fits and the blue dashed lines represent the errors on the fits.

Table 6.1: Gradients and zero-frequency intercepts found by least squares fits for method 1 and method 2. The reduced χ_ν^2 value for each fit is also given.

	Gradient	Zero-frequency intercept	χ_ν^2
	Width		
Method 1	$(2.5 \pm 0.1) \times 10^{-3}$	-6.0 ± 0.2	0.91
Method 2	$(5.2 \pm 0.2) \times 10^{0}$	-40 ± 2	0.70
	Power		
Method 1	$(4.68 \pm 0.06) \times 10^{-3}$	-17.2 ± 0.1	6.5
Method 2	$(9.8 \pm 0.1) \times 10^{0}$	-82.3 ± 0.9	6.9

performed to find the gradients and the zero-frequency intercepts for both the width and the power. The gradients and zero-frequency intercepts resulting from fitting the powers and widths in this manner can be seen in Table 6.1. This approach to fitting the data will be referred to as method 2.

6.1.4 Extrapolating method 1 and method 2 to low frequencies

To determine the value of the width and power of a mode at very low frequencies the gradients and zero-frequency intercepts were used to extrapolate the linear relationships from method 1 and method 2 to lower frequencies. The widths found using both method

Table 6.2: Predicted widths of various $l = 0$ modes found by extrapolating both the relationship found in method 1 and the relationship found in method 2.

Frequency (μHz)	Width found using method 1 (μHz)	Width found using method 2 (μHz)
1407.627	0.091 ± 0.008	0.068 ± 0.011
1263.524	0.064 ± 0.006	0.039 ± 0.008
1118.15	0.044 ± 0.005	0.021 ± 0.005
972.745	0.031 ± 0.004	0.010 ± 0.003
825.365	0.021 ± 0.003	0.0044 ± 0.002

Table 6.3: Predicted powers of various $l = 0$ modes found by extrapolating the relationship found in method 1 and the relationship found in method 2.

Frequency (μHz)	Power found using method 1 (ms^{-2})	Power found using method 2 (ms^{-2})
1407.627	$(2.38 \pm 0.09) \times 10^{-5}$	$(1.32 \pm 0.08) \times 10^{-5}$
1263.524	$(1.21 \pm 0.05) \times 10^{-5}$	$(0.46 \pm 0.03) \times 10^{-5}$
1118.15	$(0.61 \pm 0.03) \times 10^{-5}$	$(0.14 \pm 0.01) \times 10^{-5}$
972.745	$(0.31 \pm 0.02) \times 10^{-5}$	$(0.035 \pm 0.002) \times 10^{-5}$
825.365	$(0.16 \pm 0.01) \times 10^{-5}$	$(0.0078 \pm 0.0008) \times 10^{-5}$

1 and method 2 for $l = 0$ modes, in the frequency range 800 to 1450 μHz, can be seen in Table 6.2; and the powers for the same modes, also found using both method 1 and method 2, can be seen in Table 6.3. Some of these modes have not yet been observed and so we have used the mode frequencies predicted by the Saclay seismic model (Turck-Chièze et al., 2001).

As can be seen the widths extrapolated using method 2 are consistently smaller than the widths extrapolated using method 1. In fact the widths inferred by the two methods do not agree to within their associated error bars at any frequency. The difference between the estimated width of a particular mode increases as the frequency of the mode decreases. Below 1000 μHz the widths predicted by method 1 are ~ 3 times larger than the widths predicted by method 2. Additionally method 2 consistently predicts lower powers than method 1. Again the difference between the two extrapolations increases as

the frequency of the mode decreases and at no time do the powers predicted by the two methods agree to within their respective error bars. At the higher frequencies the powers predicted by method 2 are just under half the powers predicted by method 1. However, below $1000\,\mu$Hz the powers predicted by method 1 are approximately a factor of 10 larger than the powers predicted by method 2. Clearly there is a significant discrepancy between the powers and widths predicted by the two methods.

Table 6.1 gives the reduced χ_ν^2 values of the linear fits. The lower the reduced χ_ν^2 value the better the fit is at representing the data. However, if the reduced χ_ν^2 value is significantly less than unity either the model used to fit the data is too complicated or the errors on the data have been overestimated. The χ_ν^2 values are very similar for each method with method 2 providing a slightly better fit for the widths, but with method 1 giving a better representation of the powers. Both methods appear to give very good fits to the width. The values of the reduced χ_ν^2 determined for the linear power fits are large because of the small error bars associated with the observed powers.

Monte Carlo simulations were performed using the results of both extrapolations. These simulations will be described in Section 6.2. However, before we describe the simulations it is interesting to compare the extrapolated parameters with theoretical predictions that can be found using a stochastic excitation model.

6.1.5 Comparing the results of the extrapolation with theoretical predictions

Model mode damping and excitation rates can be calculated using various analytical models, which describe the interactions between the convection and the oscillations. The power of a mode, V^2, is given by

$$V^2 = \frac{P}{2\pi\Delta\nu I}, \tag{6.8}$$

where P is the energy supply rate, I is the mode inertia and $\Delta\nu$ is, as defined previously, the width of the mode. The mode inertia determines how easily a mode is excited and
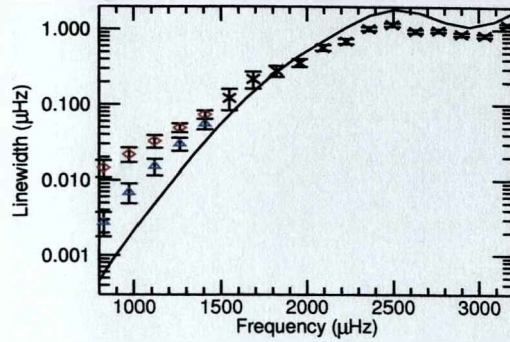
Figure 6.4: A comparison between observed, modelled and extrapolated linewidths. The solid line represents the modelled linewidths, the black crosses represent the linewidths found by fitting BiSON data, the red diamonds show the widths found using extrapolation method 1 and the blue triangles show the widths found using extrapolation method 2.

damped. The inertia can be defined in terms of the mode's total (kinetic + potential) energy, E, and the mean-square value of its surface velocity, V_s, i.e. $E = IV_s^2$ (Houdek et al. 1999, also see Christensen-Dalsgaard and Berthomieu 1991 and Christensen-Dalsgaard 2003 for more details). In what follows we will compare the powers and linewidths that have been predicted by one such model with the linewidths and powers that are observed in and can be extrapolated from BiSON data. Chaplin et al. (2005) calculated model energy supply rates and linewidths using the Cambridge stochastic excitation and damping codes. Figure 6.4 shows that the model widths drop off more rapidly than the observed widths at low frequencies. This is a known, and as yet unresolved problem, with the modelling of the linewidths (see for example Chaplin et al., 2005). It is therefore encouraging to note that both extrapolation methods predict larger widths than the solar excitation model.

The left panel of Figure 6.5 shows a comparison between the model, observed and extrapolated powers. As we are interested in the frequency dependence of the power the model powers have been scaled to ensure that the maximum model power and the maximum observed power are equal. The model powers decrease less rapidly than the powers observed in the BiSON data. Therefore the powers estimated by both method 1 and method 2 are smaller than the model powers at low frequencies. However, this is understandable as the model widths decrease more rapidly than the observed widths.

Chaplin et al. (2005) note that the model widths are too narrow at low frequencies, which implies that the extrapolations performed tend in the right direction.

It is also of interest to find the model powers using the observed widths. The right panel Figure 6.5 shows that using the observed widths produces a far better agreement with the powers observed in the BiSON data. This indicates that the majority of the discrepancy between the observed and model powers seen in the left panel of Figure 6.5 is due to the smaller model widths. The extrapolated powers also appear to agree better with the model values calculated using the observed widths (right panel of Figure 6.5). However, it is difficult to tell which extrapolation method produces the best agreement with the model powers. Also plotted in the right panel of Figure 6.5 are the model powers that are found when the widths estimated by both extrapolation methods are used. The agreement between the extrapolated powers and the model powers is poor when method 1 is used. However, the agreement between the extrapolated and model powers is better when method 2 is used. It is interesting that, when method 1 is used, the model powers are lower than the extrapolated powers. However, when method 2 is used the model values are greater than the extrapolated powers. This is likely to be because the widths extrapolated by method 1 are significantly larger than the widths estimated by method 2.

6.2 Simulating the modes

In this section we describe how we made artificial p-mode timeseries. These timeseries were analyzed, and results on detections of their artificial modes compared with the results of Chapter 5, on real p-mode data. In Chapter 5 BiSON and GOLF frequency-amplitude spectra were searched for coincident prominent structures that occurred at the same frequency in each spectrum. Here we needed to simulate pairs of timeseries, which both contained the signal from a mode and normally distributed noise. The statistical tests described in Chapter 5 were then applied to the simulated data to determine how

Figure 6.5: Left panel: A comparison between the observed, modelled and extrapolated powers. The solid line represents the modelled powers, the black crosses represent the powers found by fitting BiSON data, the red diamonds show the powers found using extrapolation method 1 and the blue triangles show the powers found using extrapolation method 2. Right panel: A comparison between the observed, modelled and extrapolated powers, however the model powers have been found using the observed widths. The symbols and lines have the same definitions as in the left panel. Also plotted are the model powers found using the widths estimated by extrapolation method 1 (red dotted line) and the model powers calculated using the widths extrapolated using method 2 (blue dashed line).

often the simulated modes could be detected.

Modes of a given frequency and width were simulated by randomly exciting an oscillator that was damped over the extrapolated lifetime. The mode's lifetime was predicted by rearranging equation 6.1 and using the extrapolations described in Section 6.1 for linewidth. The simulations produced timeseries that contained the signal from a simulated mode. The total power of the simulated mode was scaled to the power predicted by the simulations.

Two timeseries containing normally distributed random noise were created. The signal from two sets of contemporaneous real Sun-as-a-star data taken by different instruments, such as BiSON and GOLF, contains some coherent noise. The level of coherent noise is frequency dependent and at around $1000\,\mu$Hz the coherency between BiSON and GOLF data is ~ 0.1. Therefore 10% of the noise in the two simulated timeseries was set so it was common to both sets of simulated data. The timeseries containing the signal from the simulated mode was then added to each noise timeseries.

Solar modes are excited stochastically by turbulence in the convection zone, which is caused by the solar granulation. As the coherent noise in two sets of Sun-as-a-star data

is from the granulation it is possible that the coherent noise found between the BiSON and GOLF data is also correlated to the amplitude of excitation of the mode. In the simulations the amplitude of the excitation was determined by a normally distributed array. For some of the simulations the correlated noise that was added to the simulated data was taken to be the array that determined the amplitude of the excitation of the mode. Simulations were also performed when the noise that was coherent between the two simulated sets of data was independent of the amplitude of excitation of the mode.

The simulations have been performed for the $l = 0$ modes between ~ 800 and $\sim 1450\,\mu$Hz, the extrapolated properties of which are given in Tables 6.2 and 6.3. The timeseries were simulated to contain 2,211,120 points with a cadence of 120s to be consistent with the sets of BiSON and GOLF data used in Chapter 5. The simulated and BiSON timeseries used in this chapter cover the same length in time ($\sim 8.5\,$yrs), despite having different cadences and containing a different number of points. This means that the power of a mode in the simulated and observed timeseries is the same. Furthermore, the number of frequency bins that cover the width of a mode in the simulated frequency-power spectra is the same as in the fitted BiSON frequency-power spectrum. Therefore, the height of the mode in a simulated frequency-power spectrum will be the same as the height predicted by the extrapolations. The level of noise in the simulated spectra was scaled to mimic the mean level of noise observed in the BiSON frequency-power spectrum in the $100\,\mu$Hz surrounding the frequency of the mode that was simulated.

Each extrapolated power and width has an error associated with it. To obtain an upper limit on the fraction of times a simulated mode could be detected Monte Carlo simulations were performed using the $\pm 1\sigma$ values of the extrapolated widths and powers. As we were interested in determining an upper limit on the number of detections, simulations were performed using the maximum power predicted by each extrapolation method. The effect of varying a mode's width on its visibility is not obvious. Figure 6.6 shows various Lorentzian profiles that can be constructed using the properties, and their associated errors, which were estimated from method 1 for the $l = 0$, $n = 6$ mode.
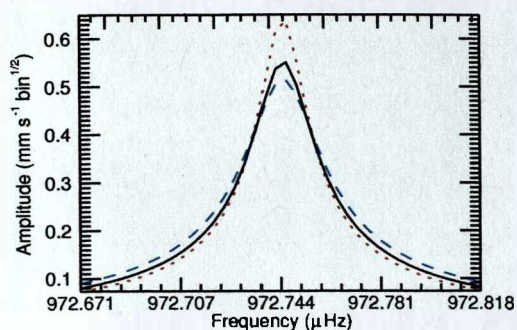
Figure 6.6: The Lorentzian profiles of the $l = 0$, $n = 6$ mode that can be produced using extrapolation method 1. The black solid line represents the Lorentzian profile produced using the power and width predicted by method 1. The red dashed line depicts the Lorentzian profile created using the maximum predicted power and the minimum predicted width. The blue dashed line shows the Lorentzian profile produced using the maximum predicted power and the maximum predicted width.

The general properties of the Lorentzian profiles that are observed in Figure 6.6 are the same for all of the modes that were simulated. Clearly the height of the Lorentzian is largest when the maximum power and minimum width are used. This is because the height, H, is proportional to the power, V^2, and inversely proportional to the width, $\Delta\nu$. As a larger height, H, should make a mode easier to detect simulations were performed using the maximum power and the minimum width allowed by the error bars on the extrapolations. Although the Lorentzian profile in Figure 6.6 that is plotted using the maximum power and the maximum width has a smaller height, H, there is more power in the wings of this profile. It is unclear how the increased power in the wings will affect a mode's visibility and so simulations were performed using the maximum power and maximum width allowed by the errors associated with the extrapolations.

As the Sun is continually being observed the length of timeseries available containing good quality data is ever increasing. Therefore, for comparison purposes, simulated timeseries were created that were twice as long as the original timeseries and so contained 6142 d of simulated data. In these simulations we have assumed that the total power in the timeseries, V^2, is twice the power predicted by the extrapolations, whilst the width, $\Delta\nu$, in units of frequency, was assumed to be independent of the length of the timeseries. As the width, $\Delta\nu$, of all of the simulated modes will be resolved in a timeseries of this

length, equation 6.2 implies that the height of the modes, H, will be the same. The power of the noise in a frequency-power spectrum is inversely proportional to the number of data points in the timeseries. As the timeseries is twice as long the mean power of the noise will be halved. Therefore, the mean S/N ratio of the simulated modes is potentially doubled.

Once timeseries containing a mode and noise had been created several tests were performed on the resulting frequency-amplitude spectra to determine whether the mode could be detected. For each mode 1000 pairs of timeseries were simulated. The number of pairs of frequency-amplitude spectra in which a mode was detected was then counted. We will now outline each of the statistical tests that were performed in turn.

6.3 Statistical tests

The tests performed are based on the statistics described in detail in Chapters 4 and 5. Here we summarize the statistical tests that were used to search the simulated spectra. The simplest test (Test 1 in Chapter 5) involves searching for a single prominent spike that is above a given threshold level in the same bin in each of the two simulated frequency-amplitude spectra. The threshold level for detection was set at a 1% chance of getting at least one false detection anywhere in $100\,\mu$Hz. This level and frequency range were chosen to maintain consistency with the detection methods employed in Chapter 5. A detection was considered to have been made if it was positioned within one linewidth of the input mode frequency. The linewidths were determined by the extrapolated values and so varied between modes. The second test involves searching for two prominent spikes in the same consecutive frequency bins of each frequency-amplitude spectrum (Test 2 in Chapter 5). Both spikes must lie within one peak width of the mode's input frequency for a detection to be counted. The third test performed also searched for two prominent spikes but this time the spikes did not need to be in consecutive bins (Test 3 in Chapter 5). These two spikes are known as a two-spike cluster. Each spike in the cluster had to

lie less than one linewidth from the mode's input frequency for a detection to be counted. The two prominent spikes in the cluster needed to be in the same bins in each frequency-amplitude spectrum. This test was then extended to search for clusters containing three, four and five prominent spikes. Furthermore, each prominent spike in the cluster had to be in the same frequency bin in each frequency-amplitude spectrum. All spikes in the clusters needed to lie within one linewidth of the mode's input frequency for a detection to be counted.

The simulated spectra were searched to determine how often these statistical tests were passed for each mode. The results of these tests will now be described.

6.4   Results of the simulations

Figure 6.7 shows the fraction of the 1000 simulations that were performed in which the mode being simulated was detected when looking for a single spike in the same bin in each frequency-amplitude spectrum. The results are shown for both types of extrapolation (method 1 and method 2). The displayed results are for the simulations that were performed where the noise that was correlated between the two datasets was coherent with the amplitude of the excitation of the mode for each extrapolation method. The results are also shown for the simulations in which the correlated noise was independent of the amplitude of excitation.

First let us consider Panel (a) in Figure 6.7. It shows three clusters of lines. The solid, black line in the cluster labelled sf=1 represents the results from simulations where the power and width of the input mode were predicted by method 1. The results show that the number of detections decreases rapidly as we move to lower frequencies. As each of the statistical tests are looking for prominent spikes it is the height the mode exhibits in the frequency-amplitude/power spectra that determines whether or not it can be detected. The height of a mode in a frequency-power spectrum is proportional to $V^2/\Delta\nu$ (see equation 6.4). The upper limit on the number of detections is given by the

(a) Method 1, noise correlated with excitation amplitude.

(b) Method 1, noise uncorrelated with excitation amplitude.

(c) Method 2, noise correlated with excitation amplitude.

(d) Method 2, noise uncorrelated with excitation amplitude.
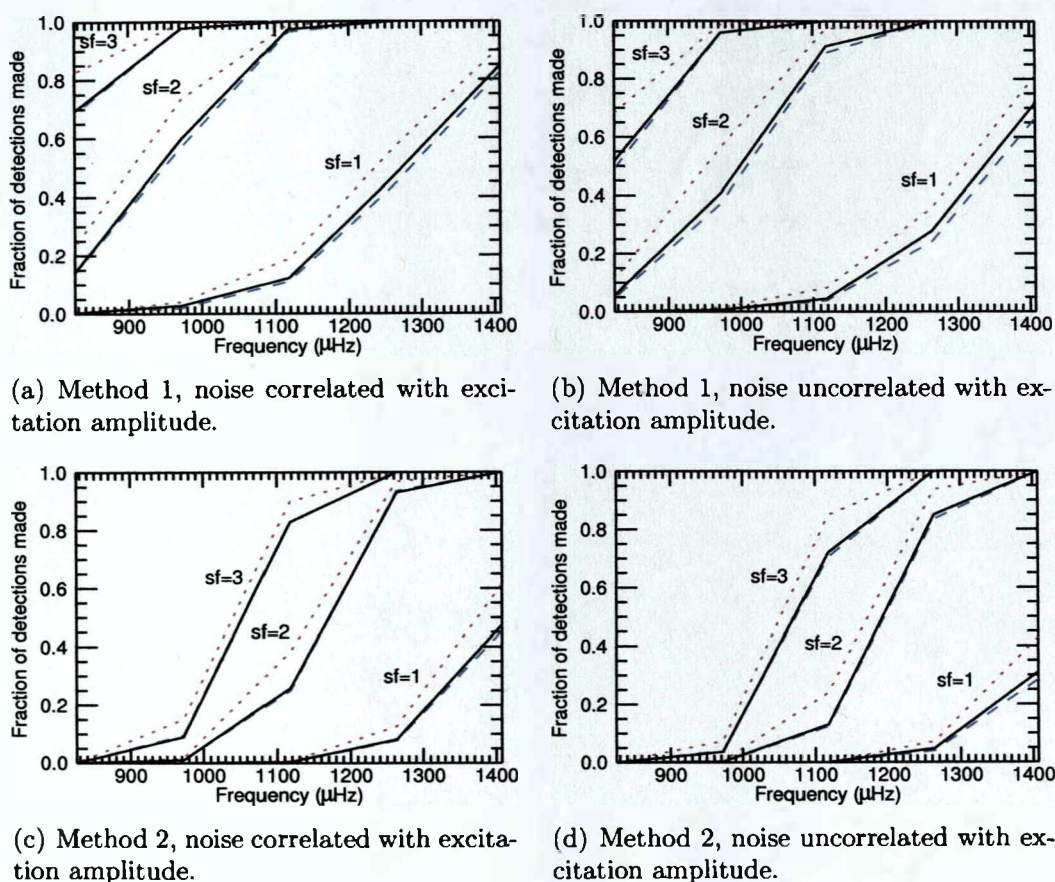
Figure 6.7: Results of searching for a single spike. The different panels represent different sets of simulations. The black solid lines show the results for the actual extrapolated values. The red dashed lines show the results when the maximum power and minimum widths are used in the simulations. The blue dashed line shows the results of the simulations that used the maximum powers and the maximum widths. The three clusters of lines are present as for each set of parameters the power was scaled by a different amount. The clusters labelled sf=1 represent the true predicted values. The clusters labelled sf=2 show the results when the power was increased by a factor of 2 and the clusters labelled sf=3 are the results when the power was increased by a factor of 3.

case when the maximum power and minimum width is used. Using the maximum power and minimum width means that the height of the mode, H, is, potentially, larger and so the signal-to-noise ratio of the mode should be increased. It is therefore understandable that using the maximum power and minimum width leads to more detections. When the simulations were performed using the maximum width the power is spread over a larger number of bins and so this negates the effect on the signal-to-noise of the mode of increasing its power.

Simulations were also performed to investigate the effect of increasing the observed signal-to-noise ratio. The power given to each simulated mode was increased by factors of 2 (labelled sf=2 in each panel of Figure 6.7) and 3 (labelled sf=3 in each panel of Figure 6.7). This significantly increases the number of detections made at low frequencies. The fraction of spectra in which the $l = 0$, $n = 6$ mode was detected when the scale factor was 1 was at most 0.05. For the $l = 0$, $n = 6$ mode the fraction of detections made when the scale factor was 2 was 0.60 and when the scale factor was 3 the fraction of simulated spectra in which the mode was detected was 0.98.

Panel (b) in Figure 6.7 also shows the results from the simulations using the mode properties predicted by method 1. However, this time the noise that was correlated between the two sets of data was independent of the amplitude of the excitations. Panels (c) and (d) in Figure 6.7 show the results of the simulations performed using the powers and widths predicted using method 2. Panel (c) shows the results for the simulations where the common noise is correlated to the mode excitation and panel (d) shows the results of the simulations where the common noise is independent of the excitation amplitude. The basic shape of the results of all four sets of simulations are similar, with the fraction of detections dropping off steeply at low frequencies.

Figure 6.8 shows a comparison between the four different cases shown in Figure 6.7. The results shown are for the simulations performed using the powers and widths predicted by the different extrapolation methods. The results show that using the parameters predicted by method 1 leads to more detections than the parameters predicted by method
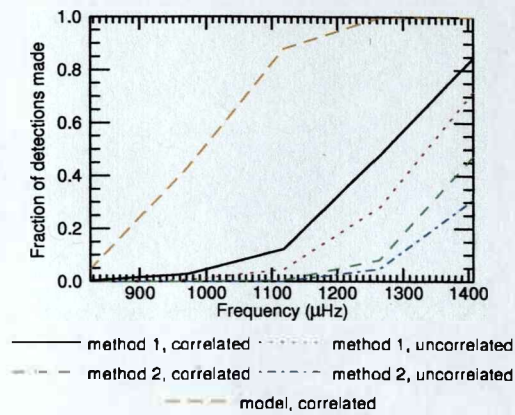
Figure 6.8: Comparison between the results of the simulations that are shown in Figure 6.7. The results shown are for the actual predicted values.

2. This is expected as the powers predicted by method 1 are significantly larger than the powers predicted by method 2. Also, more detections are made when the noise that is coherent between the two sets of data is also correlated to the amplitude of the excitation of the mode. This is also understandable. If, for example, at a point in time the level of solar noise is larger than average the amplitude of the excitation of the mode will also be larger than average enabling it to potentially remain distinguishable from the noise. However, even for the best case scenario (method 1, noise correlated to the excitation amplitude) the simulations show that the fraction of spectra in which the $l = 0$, $n = 6$ mode should be detected is less than 0.05. Also plotted on Figure 6.8 are the results of simulations performed using the theoretical widths and powers predicted by the stochastic excitation model described in Chaplin et al. (2005). The theoretical powers have been found using the model widths rather than the observed widths. Clearly more detections are made when the theoretical parameters are used.

Figure 6.9 shows a comparison between the fraction of detections made using the different statistical tests described in Section 6.3. The comparison shown here is for simulations performed using the results of method 1 when the correlated noise is coherent with the amplitude of the random excitation of the mode. This example has been chosen as this combination led to the most detections in the single spike test. However, the basic features are the same for all of the other cases shown in Figure 6.7. Above $\sim 1150\,\mu$Hz

Figure 6.9: Comparison between the different statistical tests described in Section 6.2. The results are shown for the simulations performed using the extrapolated values from method 1, where the noise that is common between the two sets of data is also correlated with the excitation of the modes.

the most successful test is the search for any 5 prominent spikes over twice the width of the mode. The test that searches for two consecutive prominent bins detects the fewest modes above $\sim 1150\,\mu$Hz. Below $\sim 1150\,\mu$Hz the single spike test detects the most modes while the 5 spike cluster test detects the least.

Figure 6.10 shows the results of the simulations performed using a timeseries containing 6142 d of simulated data. This is twice as long as the original timeseries. The results are shown for the most optimistic case, when the mode properties are found using method 1 and the common noise is correlated to the amplitude of the excitations. Lengthening the timeseries increases the number of frequency bins across the width of the mode, $\Delta\nu$, as the frequency bin width is halved, and the S/N of the mode is potentially doubled. Therefore Figure 6.10 shows the results for the 5-spike cluster test, as this gives an indication of how the results are affected by both the increased S/N and the larger number of frequency bins across the width of the mode.

6.5 Discussion

Clearly the visibility of the modes drops off significantly at low frequencies. The most optimistic results are achieved when the extrapolation is performed using method 1 and
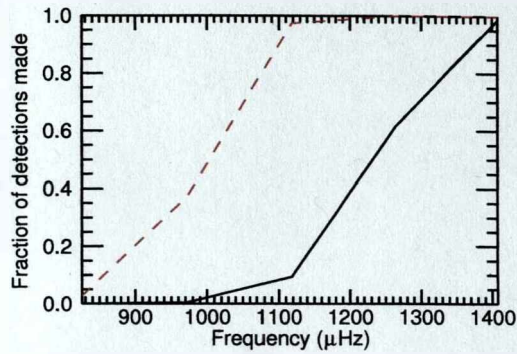
Figure 6.10: The fraction of 5-spike cluster detections made when the length of the timeseries is doubled. The power and widths were predicted using method 1 and the common noise is correlated to the excitation amplitude. The black solid line shows the results for the standard 3071 d timeseries while the red dashed line shows the results for the 6142 d timeseries.

the common noise is correlated to the amplitude of the excitation of the mode. Clearly more detections are made when the theoretical model parameters are used. However, these results can only be treated as an upper limit as we have already seen that the model underestimates the widths of the modes. This not only leads to the power of the modes being overestimated but also means that this power is spread over fewer bins in a frequency-power spectrum. Hence the height of a mode is increased, and so a mode is more likely to be detected.

For both extrapolation methods, irrespective of whether the noise is correlated or uncorrelated with the excitation amplitude, the fraction of detections at $\sim 1000\,\mu$Hz is less than 0.1. On the assumption that method 1 produces a more accurate prediction of the mode's powers and widths than method 2 the simulations imply that if the power S/N ratio can be increased by a factor of 2 a significantly higher proportion of modes could be detected. It would therefore be beneficial to continue efforts to try and improve the quality of the data further.

The total background continuum is a combination of instrumental noise, solar noise and, for Earth-based instruments such as the BiSON network, a small amount of atmospheric noise. If the atmospheric and instrumental noise can be reduced this would increase the S/N ratio of the mode making it easier to detect. The limiting factor to the improvement that can be made to the data quality is the amount of solar noise present.
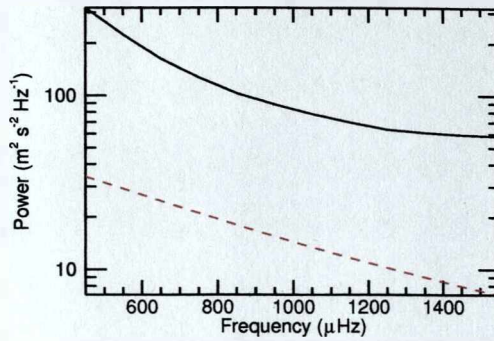
Figure 6.11: The power of the solar noise predicted by the Harvey model and the characteristic velocity and decay time found by Elsworth et al. (1994) (red dashed line) can be compared with the mean power of the BiSON data, which was used when fitting the well-defined modes.

Harvey (1985) modelled the solar continuum as a superposition of noise due to three scales of granulation (granulation, mesogranulation and supergranulation) and active regions (see Chapter 2 for details). Above $\sim 400\,\mu$Hz the dominant source of solar noise is the normal granulation. Elsworth et al. (1994) found values of the root mean square velocity of the solar surface, σ, and the characteristic decay time, τ_d, for the normal granulation by fitting equation 2.1 to the cross-amplitude spectrum found using BiSON data from two different sites. The red dashed line in Figure 6.11 shows the power of the solar noise as predicted by characteristic values found by Elsworth et al. (1994), namely $\sigma = 0.319\,\mathrm{m\,s^{-1}}$ and $\tau_d = 260\,\mathrm{s}$. Also plotted is the mean power of the BiSON data, which was used to predict the powers and widths of the low-frequency modes. The mean power of the BiSON data is always more than twice the predicted power of the solar noise. Therefore it is theoretically possible to significantly reduce the level of noise in this frequency range.

At frequencies above $\sim 1150\,\mu$Hz the test that detected the highest proportion of modes was the search for 5 prominent spikes across twice the width of the mode. However, at lower frequencies this test detects the least number of modes. The width of the lowest-frequency mode for which the simulations were performed, at a frequency of $\sim 825\,\mu$Hz, covers 4 bins if the width is found using method 1. As we allow the spikes in a cluster to be spread over twice the width of the mode the cluster test for five spikes can be used to search for this mode. If method 2 is used to extrapolate the widths this test cannot be

used to search for the $l = 0$, $n = 5$ mode as this method predicts that all of the power of the mode will be confined to one bin. The five-spike cluster test can be used to search for the $l = 0$, $n = 6$ mode although this would require five consecutive spikes to be higher than the threshold. At higher frequencies there is little difference between the results of the five-spike cluster test and the four-spike cluster test. This implies that little would be gained by performing statistical tests that look for clusters containing more than 5 spikes. Since low-frequency modes have such long lifetimes all of the power from the signal of a low-frequency mode is concentrated in a few frequency bins only. Therefore, they would be too narrow to be detected by a test which searched for clusters that contain a larger number of spikes.

Figure 6.10 shows that longer timeseries would be beneficial when searching for low-frequency modes, as the chances of detecting modes are increased. It should be noted that the results of these simulations are in good agreement with the results of the simulations when the power was increased by a factor of two. As mentioned earlier, this is expected as the S/N should also be increased by a factor of 2 in the simulations where the length of the timeseries is doubled.

It is possible to compare the results of these simulations to the observed candidates found in Chapter 5. According to the simulations the $l = 0$, $n = 9$ mode at $\sim 1407\,\mu$Hz should be detected in the vast majority of spectra and this mode is observed to be very prominent in the BiSON and GOLF data used in Chapter 5. In Chapter 5 we also detect the $l = 0$, $n = 8$ mode at $\sim 1263\,\mu$Hz while the simulations indicate that the fraction of time that this mode should be detected is at least 0.6. The $l = 0$, $n = 7$ mode at $\sim 1118\,\mu$Hz and the $l = 0$, $n = 5$ mode at $\sim 825\,\mu$Hz are not detected in the BiSON and GOLF data used in Chapter 5. The fraction of simulated spectra in which the $l = 0$, $n = 7$ mode was detected was, at most, 0.19, while the fraction of simulations in which the $l = 0$, $n = 5$ mode was detected was less than 0.01. Therefore it is reasonable to expect that the modes might not be detected in the BiSON and GOLF data.

It should of course be noted here that all of these results are based on the assumption

that the simple relationships for the mode powers and widths observed in well-defined modes are still valid at low frequencies. The solar excitation model suggests that this may not be an unreasonable assumption, although it is difficult to determine which extrapolation method predicts the lifetimes and powers more accurately.

The $l = 0$, $n = 6$ mode is significantly more prominent in the observations than the simulations. Therefore this mode will now be considered in more detail.

6.5.1   The $l = 0$, $n = 6$ mode

The $l = 0$, $n = 6$ mode is observed to be prominent in Sun-as-a-star data. However, some modes with higher frequencies remain undetected. The most optimistic simulations imply that the fraction of spectra in which the $l = 0$, $n = 6$ mode should be detected is less than 0.05. In the simulations performed here, for method 1, with noise correlated to the amplitude of the excitation, the average simulated S/N ratio in amplitude was ~ 2.1. However, the observed amplitude S/N ratio in the GOLF and BiSON data examined in Chapter 5 was greater than 3.4. In fact, the fraction of simulations in which the S/N ratio of the mode was greater than 3.4 was 0.014. Therefore the simulations poorly represent the observed properties of this mode. It should be noted that the amplitude S/N ratio of other modes that were detected in Chapter 5 are well represented by the simulations in this chapter. For example, the $l = 0$, $n = 9$ mode was observed to have a S/N ratio of 3.1 while the average S/N ratio produced by the simulations ranged from 2.9 to 3.7 depending on which extrapolation method was used and whether the noise was correlated to the excitation amplitude.

When the mean amplitude S/N ratio of the simulated $l = 0$, $n = 6$ mode was set as 3.4 the fraction of spectra in which the mode was detected increased to ~ 0.6. A S/N ratio of 3.4 is $\sim 3\sigma$ from the mean amplitude S/N found in the original simulations where extrapolation method 1 (ln-linear) was used and the noise was correlated to the excitation amplitude. Therefore, the simulations and the observations are not so different that new physics is required to explain the observations. It is more likely that, at some

point in time, the mode has been randomly excited to a larger amplitude because of the stochastic nature of mode excitation.

6.6 Summary

To investigate whether the gaps in the list of modes detected in Chapter 5 are reasonable we have made predictions of the detectability of low-frequency p modes. Estimates of the powers and damping times of these low-frequency modes were determined by extrapolating the observed powers and widths of higher-frequency modes with large observed S/N ratios. The extrapolations were made using two different methods. However, both methods predict that the low-frequency modes will have small S/N ratios and narrow widths in a frequency-power spectrum. Nonetheless, comparisons with the Cambridge stochastic excitation model indicate that the extrapolated trends are tending towards the right direction.

Monte Carlo simulations were then performed where timeseries containing mode signals and normally distributed noise were produced. The mode signals were simulated to have the powers and damping times predicted by the extrapolations. Various statistical tests were then performed on the frequency-amplitude spectra formed from these timeseries to investigate the fraction of frequency-amplitude spectra in which the modes could be detected.

The fraction of simulated frequency-amplitude spectra in which modes were detected decreases rapidly as the frequency of modes decreases and so the fraction of simulations in which the low-frequency modes were detected was small. However, increasing the S/N ratio of the low-frequency modes by a factor of 2 above the extrapolated values led to significantly more detections. Similarly doubling the length of the simulated timeseries led to more detections being made. Therefore efforts should continue to further improve the quality of the solar data and the length of timeseries that are currently available.

Using the properties predicted by extrapolation method 1 led to more detections

than when the properties predicted by method 2 were used in the simulations. An upper limit on the number of detections that could be made was obtained by using the properties predicted by the Cambridge stochastic excitation and damping codes. More detections were made when the noise was correlated to the excitation of the modes. In the simulations performed here it was assumed that the amount of data coherent with the excitation of the modes was the same as the amount of coherent data present in BiSON and GOLF timeseries, as this is solar in origin. However, as BiSON and GOLF observe different heights in the atmosphere, and because of the effects of Doppler imaging (see Chapter 8) not all of the granulation noise observed by the two instruments will be coherent. In other words, each set of data will contain some independent solar noise. Therefore, it is likely that a larger proportion of the noise in each timeseries is correlated to the mode excitation than is simulated here. It would, therefore, be interesting to investigate the effect of introducing some noise, which although not correlated between the two simulated timeseries, is coherent with the mode excitation.

The $l = 0$, $n = 6$ mode at $\sim 973\,\mu$Hz is observed to be more prominent in real data than in the simulations. However, we have shown here that the discrepancy is not significant. It is, therefore, possible that the $l = 0$, $n = 6$ mode has just randomly been excited to a larger amplitude because of the stochastic nature of the excitation. Similarly, it is possible that the mode has been randomly more lightly damped, thereby increasing its lifetime. Hence, a possible extension to this work is to study the evolution of the $l = 0$, $n = 6$ mode over time.

In this chapter we have shown that the gaps in the list of detected modes found in Chapter 5 are not entirely unexpected. In Chapter 5 we compared contemporaneous BiSON and GOLF data. We now move on to investigate whether any low-frequency modes can be detected in contemporaneous sets of BiSON data. To do this we need to find at least two sets of contemporaneous data with a relatively low coherency at low-frequencies as the less coherent the data are the lower the threshold levels.

7. THE COHERENCY OF BISON DATA

The main focus of this thesis is to search for low-frequency p modes and g modes. We have already outlined statistical techniques aimed at finding low-frequency modes and we have used the statistics to search BiSON and GOLF data. We have previously shown that detection thresholds levels can be reduced by comparing two frequency-amplitude spectra for coincident prominent spikes or structures of spikes. In this chapter we use the statistical tests to search pairs of contemporaneous and non-contemporaneous BiSON frequency-amplitude spectra.

Usually a BiSON timeseries will contain the best data accumulated from 6 BiSON sites, which are strategically positioned so that continuous observations of the Sun can, potentially, be made. In reality bad weather and technical difficulties mean that the fill of a BiSON timeseries is $\sim 80\%$. However, the latitudes at which the instruments are situated mean that much of the time at least two instruments are making contemporaneous observations of the Sun. When creating a single BiSON timeseries we can take advantage of time overlaps in the data measured by the different stations by including only the best quality data. However, we can also keep the overlapping data separate to create multiple sets of contemporaneous timeseries that can be compared.

Until recently two of the BiSON sites, Las Campanas and Carnarvon, had two instruments and these instruments were designated A and B. Furthermore, the instruments Las Campanas A, Las Campanas B, Carnarvon B, Mount Wilson, Narrabri and Sutherland each take two sets of measurements known as port and starboard observations. Usually these port and starboard measurements are averaged to produce a mean timeseries for each instrument. The data from the different stations can then be coherently added to

create a timeseries with a high fill. However, it is possible to keep the data from the individual instruments and the port and starboard data separate to create individual timeseries.

Therefore, there are many different sets of contemporaneous BiSON data that can be searched for low-frequency modes. For example, comparisons can be made between the port and starboard data and the data from different stations. Furthermore, data from the different sites can be coherently combined in numerous ways to create pairs of contemporaneous BiSON timeseries that can be searched at low frequencies for p and g modes.

As the behaviour of the data varies with frequency we have found it useful to split many of the spectra plotted in the frequency domain into the following three regions:

1. The 'low-frequency region' covers the range $0 - 1500\,\mu\text{Hz}$. This is the region where we are searching for low-frequency p modes and g modes. The amplitudes of the modes are very small at these frequencies and so the spectrum is dominated by noise. In many of the graphs that follow in the rest of this chapter (e.g. Figure 7.1) the upper boundary of the low-frequency region is given by the vertical dashed red line.

2. The 'five-minute region' covers the range $1500 - 5500\,\mu\text{Hz}$. This frequency band contains the majority of the power from the oscillatory modes of the Sun. The acoustic cut-off frequency is at $5500\,\mu\text{Hz}$, while below $1500\,\mu\text{Hz}$ the power of the modes are only just visible above the background noise. In this region the power observed from the modes is significantly greater than the power of the noise. In the majority of the figures in this chapter the 'five-minute region' is the frequency range between the vertical red dashed line and the vertical blue dot-dashed line.

3. The 'high-frequency region' covers the range $5500 - 12500\,\mu\text{Hz}$. This is above the acoustic cut-off frequency and so there should be very little mode power present. Therefore this region is dominated by noise. The upper limit is bounded by the

cadence of the timeseries, which is 40 s for the data used in this chapter. The lower limit of the high-frequency region is denoted by the vertical dot-dashed blue line.

When comparing two spectra with the aim of detecting modes with a low signal-to-noise ratio (S/N) the presence of common noise increases the probability that any prominent spike found in the same bin in each of the two spectra is due to noise. The optimum circumstances for detecting modes occur when the coherency between the two spectra in the region of interest is zero as the detection thresholds are lowest. In other words we are more likely to detect a mode if the two spectra are completely independent.

In this chapter we have aimed to make use of the multitude of BiSON timeseries available to search for low-frequency p modes and g modes. To do this we needed to find two sets of BiSON data with as little common noise as possible that we ultimately searched for low-frequency modes. With this in mind we begin this chapter by looking at the coherency of different combinations of BiSON data.

7.1 Comparison of port and starboard data

As mentioned above six BiSON instruments are able to take port and starboard measurements. The first pair of timeseries that were examined contained a combination of data from the six instruments with port and starboard data available. However, the port and starboard data were kept separate to form a port timeseries and a starboard timeseries. Each time overlapping data from different stations occurred the best quality data point was selected. To maximize the quality of the low-frequency data a rejection threshold of $150\,\mathrm{m\,s^{-1}\,Hz^{-1}}$ is set in the region of $800 - 1300\,\mu\mathrm{Hz}$. No moving mean has been taken as this would suppress the low-frequency modes. Timeseries lasting 10 years, commencing on the 1st January 1996, were used and so the data span the majority of activity cycle 23. The timeseries contained 7,890,481 data points and, as with most BiSON data, had a cadence of 40 s.

Figure 7.1 shows the ratios of the port and starboard mean absolute amplitudes,
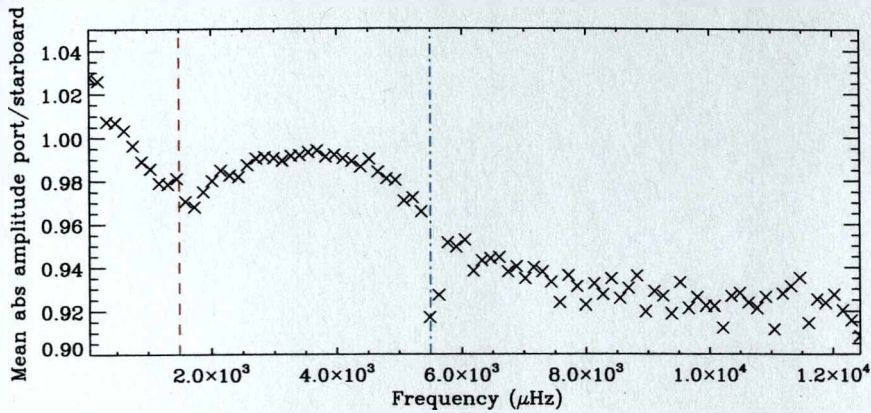
Figure 7.1: Ratio the port and starboard mean absolute amplitude smoothed over $139\,\mu$Hz.

which are given by $\left([\text{real amplitude}]^2 + [\text{imaginary amplitude}]^2\right)^{\frac{1}{2}}$. Later in this chapter the coherency will be found using short 2hr timeseries with a 100% fill (for a description of this method see Section 1.6.2). The bin width in the frequency-amplitude spectrum formed from a 2 hr timeseries is $\sim 139\,\mu$Hz and so, to maintain consistency, the plotted ratio has been smoothed over $139\,\mu$Hz. As the ratio is less than one at frequencies greater than $\sim 500\mu$Hz, Figure 7.1 shows that the port measurements are marginally less noisy than the starboard measurements. However, the amplitude of the modes in the five-minute peak are approximately the same.

The majority of the solar noise, atmospheric noise and some of the instrumental noise will be common to both the port and starboard observations. The common instrumental noise includes pointing noise, due to guidance errors, noise due to temperature variations in the instrument, noise due to variations in the magnetic field that encloses the vapour cell and noise due to the efficacy of the Pockels cell (see Section 2.3 for more details). Some of the instrumental noise will be independent such as some electronic noise and the photon statistical noise. The port and starboard detectors are positioned on opposite sides of the vapour cell, which means that the observations will be weighted towards different regions on the solar disc (see Chapter 8). Therefore, the solar background noise observed by port and starboard detectors are perceptibly different.

Figure 7.2 shows the cross-amplitude, phase and coherency spectra of the multi-
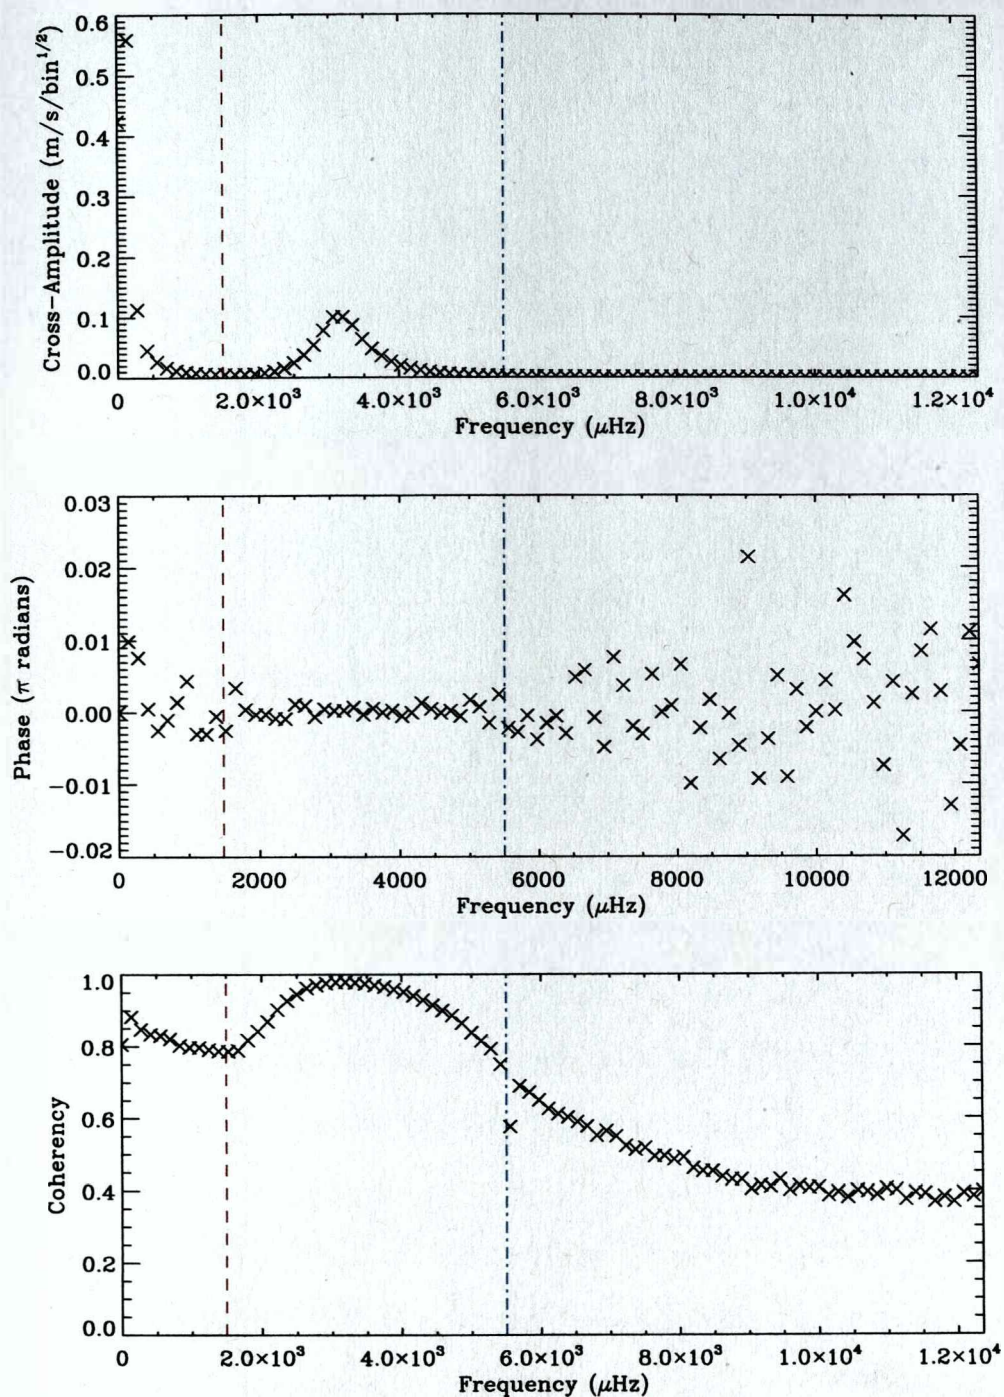
Figure 7.2: The cross-amplitude spectrum (top), phase (middle) and coherency (bottom) of the port and starboard datasets containing the combined data from all of the sites that take separate port and starboard measurements.

Figure 7.3: The port and starboard absolute amplitude spectra between 5000 and 6000 μHz. The data have been smoothed over 1 μHz bins.

site port and starboard data. Each was calculated by determining the average of all available 2 hr spectra with a 100% fill (see Section 1.6.2 for method). In the low-frequency region the cross amplitude increases with decreasing frequency and the coherency is ~ 0.8. As there is little mode power at low-frequencies this implies that the port and starboard frequency-amplitude spectra are dominated by common noise, which could be atmospheric, solar or instrumental. The phase spectrum shows that this noise is almost in phase. In the five-minute region the coherency is approximately unity, as the detectors observed almost the same signal from the modes, which are dominant here.

Notice the sharp dip in the coherency at $\sim 5500\mu$Hz in the bottom panel of Figure 7.2. Figure 7.3 shows the smoothed port and starboard frequency-amplitude spectra between 5000 and 6000 μHz. A definite peak is present at $\sim 5500\,\mu$Hz that is significantly larger in the starboard measurements than in the port observations. It is thought that this peak is due to a 'gear spike'. The gear on the motor that drives the direction in which an instrument is pointing has teeth on it. Every time the gear steps over one of these teeth

the alignment of the instrument is shifted slightly off centre. Another motor corrects for this effect but it nonetheless leads to a periodic noise source in the data that becomes evident when the power spectrum is examined. As the shift will alter the alignment in the same direction every time and the port and starboard detectors are situated on opposite sides of the vapour cell the misalignment will affect one detector more than the other. As can be seen from Figure 7.3 the starboard instrument is affected more severely than the port instrument. Furthermore, the noise created by this shift will be anti-correlated between the port and starboard data, thus explaining the dip in coherency observed in the bottom panel of Figure 7.2. The oscillatory signal observed below 5500 μHz in Figure 7.3 is from high-frequency p modes. The short lifetimes of modes at the high-frequency end of the spectrum mean that high-frequency modes are visible as broad peaks in a frequency-amplitude spectrum. This is why their signal has the appearance of a regular oscillation rather than the distinct features that are observed at low-frequencies.

We have seen that there is a large proportion of coherent data at low frequencies that is common to the port and starboard data. This coherent data can be explained in terms of solar, atmospheric and instrumental noise. However, it is also interesting to note the presence of coherent data above the acoustic cut-off frequency ($\sim 5500\mu$Hz) where there are no modes. The source of this high-frequency coherent noise is less obvious. Slightly below the acoustic cut-off frequency the shape of the modes in a frequency-power spectrum can be described by relatively wide Lorentzian profiles that have extended tails. It is possible that the high-frequency coherent data is due to the accumulation of these Lorentzian tails. We will now describe various simulations that were performed to determine whether the power in the Lorentzian tails of mode profiles are able to explain the observed high-frequency coherency.

7.1.1 *The high-frequency coherency produced by the wings of Lorentzian mode profiles*

To test whether the Lorentzian tails are a possible mechanism for generating the observed high-frequency coherency two simulated timeseries, which contained the same modes and

Figure 7.4: The coherency of two FLAG timeseries that contain identical modes but completely independent noise. The coherency was found by looking at 2hr spectra with 100% fill.

totally incoherent noise, were generated. The simulated data were originally constructed as frequency-amplitude spectra with both real and imaginary parts. Both the real and imaginary noise were simulated to have a Gaussian distribution but the variance of the noise was varied with frequency in an attempt to mimic the level of noise in real Sun-as-a-star data. Two completely independent frequency-amplitude spectra were simulated in this manner. The frequency-amplitude spectrum of FLAG modes was added to both noise spectra and the inverse FFT was performed on both simulated frequency-amplitude spectra to create two timeseries. The coherency was then found by taking the average over all 2hr spectra (see Figure 7.4). Although the coherency decreased above the acoustic cut-off frequency the minimum value reached was approximately 0.07. The only coherent data in the two simulated timeseries were the signals from the modes, none of which had frequencies above $\sim 5500\,\mu$Hz. Therefore the observed high-frequency coherency must be due to the Lorentzian tails of the mode profiles. However, the coherency observed in this simulation is still significantly less than the coherency observed in the real port and starboard data at high frequencies (see Figure 7.2). Hence there must be at least one other source of high-frequency noise that is coherent between the port and starboard BiSON data.

The high-frequency coherency found in these simulations (see Figure 7.4) is in good agreement with the high-frequency coherency found by Elsworth et al. (1994) between data from Izaña and Sutherland. This implies that the high-frequency coherent noise

shared by the port and starboard data that cannot be explained by the Lorentzian tails is not common between instruments from different sites. Therefore, the extra coherent noise shared by the port and starboard data could be instrumental and/or due to transparency fluctuations. Although the solar noise observed in the port and starboard data will not be identical it will still be more coherent than the solar noise observed by two different stations. The solar background continuum decreases rapidly in power as frequency increases and so should not be prominent in the high-frequency region. The cross-amplitude spectrum in the top panel of Figure 7.2 indicates that there is almost no coherent signal at high frequencies between the port and starboard data. Therefore it is possible that, although the magnitude of solar noise is small at high frequencies, it could make a significant contribution to the total high-frequency noise background as the magnitudes of other noise sources are also small at these frequencies. The common solar noise could, therefore, account for some of the observed high-frequency coherency.

We can, therefore, conclude that although some of the high-frequency coherency observed above the acoustic cut-off frequency in the real port and starboard data (see Figure 7.2) may be due to the Lorentzian tails the majority is not. Hence another form of noise must be responsible for the observed coherency and possible noise sources include instrumental, atmospheric and solar noise. One other possible source of high-frequency coherent data are pseudo-modes. These modes have been detected in various sets of solar observations (see for example Jefferies et al. 1988, Libbrecht 1988, Duvall et al. 1991 and García et al. 1998) including BiSON data (Chaplin et al., 2003a). Although there is some speculation as to the source of these modes it is thought that they are High-frequency Interference Peaks (HIPs). HIPs are thought to be the manifestation of constructive interference between waves that are emitted in an outward direction from their sources below the photosphere and waves that initially propagate inwards (García et al., 1998). However, as we are more interested in the low-frequency region here we will not investigate this further.

The simulations used to determine the contribution of the Lorentzian tails to the

high-frequency coherency also produced some unexpected results at low frequencies. We move on to describe these results and discuss the reasons they are observed.

7.1.2 The contribution to the low-frequency coherency of the wings of Lorentzian mode profiles

As we have seen Figure 7.4 shows the coherency observed between two simulated time-series that contained the same FLAG modes but independent noise. It is interesting to notice that below $\sim 1000\,\mu$Hz the coherency decreases very slowly with frequency. This is unexpected as all of the modes included in the timeseries had frequencies greater than $1000\,\mu$Hz. Also, low-frequency modes have relatively narrow Lorentzian profiles that cover a few bins only. Therefore the influence of the low-frequency modes' Lorentzian tails should be minimal. It is, however, possible that the low-frequency tails from higher-frequency modes are not negligible here. Before investigating this further it should be noted that it is possible that the overall effect of the Lorentzian tails are enhanced here as the level of noise included in the simulation has been designed to mimic the level of noise in real low-frequency data approximately only.

Figure 7.5 shows the Lorentzian profiles of the modes included in the simulated timeseries used in Section 7.1.1. As can be seen when the total power from the Lorentzian profiles is calculated power from the wings of the profiles is still present at low frequencies. Therefore it is understandable that some coherency was found at low frequencies in the simulation performed in Section 7.1.1. Figure 7.4 shows that at $\sim 1000\,\mu$Hz the coherency in the simulated data is approximately 0.1, while the ratio of the sum of the Lorentzian profiles and the simulated noise is ~ 0.1 at $\sim 1000\,\mu$Hz, implying that the Lorentzian tails can explain the low-frequency coherency of the simulated data observed in Figure 7.4.

It should be noted that the low-frequency coherency produced by the Lorentzian profiles is not large enough to explain the low-frequency coherency observed in the real data (see Figure 7.2). There are large amounts of data common to both the port and
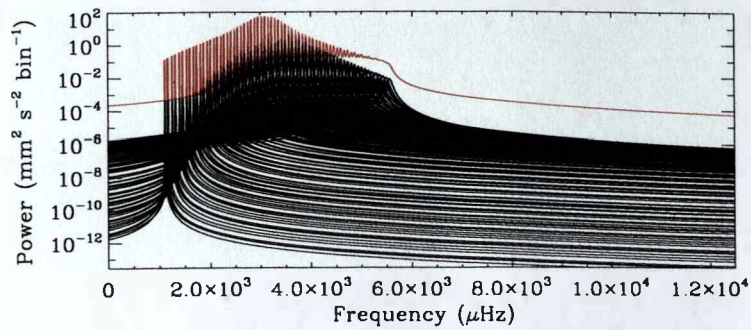
Figure 7.5: The Lorentzian profiles of FLAG modes. The black lines represent the individual Lorentzian profiles and the red line represents the sum of all of the Lorentzian profiles.

starboard observations at low frequencies. Therefore using these timeseries to search for low-frequency modes would not be ideal as the detection thresholds levels would be relatively large. However, it is possible that shifting the start time of the timeseries so they are no longer contemporaneous would remove the unwanted, coherent, low-frequency noise. We will now describe the coherency found when non-contemporaneous port and starboard data are compared.

7.2 Non-contemporaneous port and starboard data

In Section 7.1, the coherency was found using short 2hr timeseries. Therefore separating the start times of the port and starboard timeseries by more than 2hrs would mean the 2hr frequency-amplitude spectra that are compared when calculating the coherency are from independent blocks in time. If two timeseries come from independent blocks of time the coherency of the noise in the data should, theoretically, be reduced to ~ 0.

The coherency of any individual mode depends on the length of separation between the start times. In turn, the extent to which the length of the separation affects the coherency depends on the damping rate of the mode. For example, if a mode has a lifetime of the order of a few days, as is the case for modes with periods of about 5 minutes, the mode will experience significant damping over 24hrs. Therefore, if the start times of two sets of observations are separated by $\sim 1\,$d the coherency of the mode will be significantly attenuated. However, here we are concerned with low-frequency modes

whose the damping times are of the order of months. Hence low-frequency modes should still be reasonably coherent even if the start times of the two timeseries are separated by 24hrs. Separating the start times in such a manner will also provide a useful investigation into the temporal behaviour of the noise in a frequency-amplitude spectrum. However, before we investigate the coherency of non-contemporaneous port and starboard data we ensure in more detail that separating the start times of two timeseries by $\sim 1\,d$ will not destroy the coherency of the low-frequency modes that we are trying to detect.

7.2.1 The effect of a mode's lifetime on the coherency of shifted spectra

Simulations were performed to investigate how a modes's lifetime affects the coherency of the mode in non-contemporaneous data. The modes were simulated by creating a damped oscillator that was re-excited randomly in time. The damping time of a mode was calculated from the fitted width of the mode in an 8.5 year-long set of BiSON data. The power of each mode in the simulation was scaled to match the fitted power from the BiSON data. The results for each mode were simulated separately. Simulations were performed for $l = 1$ modes with $10 \leq n \leq 22$. These modes are at higher frequencies than the ones we are searching for in this thesis. However, they have a large S/N ratio and so their properties, such as lifetime and power, can be determined accurately. Initially simulations were performed using two identical timeseries that contained the signal from a simulated mode only.

For this investigation the average coherency over all 5.5 hr timeseries was calculated as increasing the length of the short timeseries reduces the frequency bin width in the resulting frequency-amplitude spectra to $50\,\mu\text{Hz}$. Timeseries that were 5.5 hrs in length were not used when examining real data as the number of 5.5 hr segments with 100% fill in a BiSON timeseries is small. The improved resolution in frequency-amplitude spectra made from 5.5 hr timeseries enabled a clearer evaluation of the effect of a time shift on the coherency of the modes to be made.

The top panel of Figure 7.6 shows the coherency of the modes when the start times

of the simulated timeseries were shifted so that they were separated by 12hrs (black crosses) and 24hrs (black diamonds). The coherency of two unshifted identical sets of data is always unity independent of the power and width of the simulated mode. Clearly the effect on the coherency is frequency dependent. Comparison with the lifetimes of the modes, which are shown in the middle panel of Figure 7.6, implies that the magnitude of the decrease in coherency is also lifetime dependent. The coherency of a low-frequency mode with a long lifetime is reduced by less than the coherency of a higher-frequency mode with a short lifetime. This is because the time shift represents a smaller proportion of the low-frequency modes' lifetime and so the low-frequency mode will experience less damping over the length of the shift. The magnitude of the decrease in coherency increases as the length of the shift in start times increases. Modes with a frequency of about 1500μHz have a lifetime of approximately 1 month and so 24 hrs is equivalent to $\sim \frac{1}{30}$th of the mode's lifetime. As a final test on the effect of shifting the start times of the timeseries on the coherency of the modes, the start times of the two timeseries were separated by $\frac{1}{30}$th of the lifetime of the mode being simulated. The top panel of Figure 7.6 shows that this only decreases the coherency by a very small amount at all frequencies. More careful examination shows that the coherency decreases slightly over the range of frequencies examined here. This could be related to the power and width of the mode in frequency-power spectra as they also decrease with frequency.

These simulations are not very realistic as the timeseries contain no noise. Therefore, two timeseries were simulated to contain the same simulated mode signal but independent noise. The start times of the timeseries were shifted by 12 hrs and 24 hrs and the coherency at the simulated mode's frequency was determined. The results of these simulations, which can be seen in the bottom panel of Figure 7.6, show that the coherency of different modes is still dependent on the length of the imposed shift in the start times. Notice that the unshifted coherency decreases with frequency. This is because the power of the mode decreases with frequency and so the S/N ratio of the mode decreases. Additionally the width of a mode decreases with frequency and so a low-frequency mode covers a

Figure 7.6: Top panel: The effect of separating the start time of timeseries that only contain the signal from a mode. The pairs of timeseries were created to contain one mode that has the same power and lifetime as is observed in ~ 8.5yrs of BiSON data. Middle panel: The lifetimes of the modes simulated in the top and bottom panels. Bottom panel: The coherency of simulated modes found when the start times of the timeseries are separated by various lengths in time. The timeseries contain the signal from a mode and independent noise.

Figure 7.7: The coherency when the start time of the port data has been shifted to begin 12hrs after the start of the starboard data.

smaller percentage of a 50μHz bin than a high-frequency mode. This means that at low-frequencies the mode has less influence on the calculated coherency of the bin and, as a result, the coherency of the noise, which is zero here, becomes more dominant. The bottom panel of Figure 7.6 indicates that shifting the start times still has a less severe effect on the coherency of lower-frequency modes than higher-frequency modes, even when noise is included. The red triangles in the bottom panel of Figure 7.6 show the result of separating the start times of the simulated timeseries by $\frac{1}{30}$th of the lifetime of the simulated mode. As can be seen, in each case the reduction in coherency is minimal implying that shifting the data by both 12 hrs and 24 hrs would have little effect on the coherency of low-frequency modes.

We have shown that separating the start times of two timeseries does not significantly affect the coherency of low-frequency modes. Therefore, we now investigate the effect of separating the start times of the port and starboard data by both 12 hrs and 24 hrs. This will not only provide useful information about the time evolution of noise in BiSON data but it may also lower the observed low-frequency coherency, thus producing the optimum conditions for searching for coincident low-frequency modes in two frequency-amplitude spectra.

Figure 7.8: The coherency when the start time of the port data has been shifted to begin 24hrs after the start of the starboard data.

7.2.2 *The coherency of non-contemporaneous port and starboard BiSON data*

The start time of the port timeseries was shifted to be both 12hrs and 24hrs after the start time of the starboard timeseries by removing 1080 and 2160 points, respectively, from the start of the port timeseries. The equivalent number of points were removed from the end of the starboard timeseries so that the port and starboard timeseries still covered the same length in time. As can be seen in Figure 7.7 introducing the 12 hr shift produced the desired effect of reducing the low-frequency coherency to approximately zero. However, it should also be noted that the coherency of the five-minute peak has also been significantly reduced. This means that these modes are less coherent, however, as shown in Section 7.2.1 the coherency of the low-frequency modes will only be reduced by a small amount. The coherency is still approximately zero at low frequencies because low-frequency modes have a small S/N and a narrow width and so each bin in the 2 hr frequency-amplitude spectra is dominated by noise that is, now, incoherent.

Interestingly when the data are shifted by 24 hrs the low-frequency coherency increases (see Figure 7.8). The coherency in the low-frequency range has a maximum of 0.2 at very low frequencies and a minimum of ~ 0.03 at $\sim 1500 \mu$Hz. Although not shown here this behaviour can also be seen if the start times of the data are 48hrs apart. This implies there is some source of low-frequency, coherent noise that repeats every 24hrs.

Simulations were run to investigate whether this effect was a real part of the data, a side effect of the analysis or due to the window function of the data, which contains

24hr periodicities. Two sets of Gaussian noise were created with a variance equal to the variance of the amplitude of the port data at $\sim 1400 \mu$Hz. The two noise timeseries were given a coherency of 50% and a frequency-amplitude spectrum of FLAG modes was then added to both sets of data. The frequency-amplitude spectra were then Fourier transformed back into timeseries. The window function of the port data was applied to one timeseries and the window function of the starboard data was applied to the other timeseries. The data are not plotted here but, as expected, the coherency was flat and equal to 0.5 both below and above the frequencies of the modes. Also as expected the coherency peaked at a value of 1 at $\sim 3300 \mu$Hz, where the power of the modes dominate. When the data were shifted by 12hrs the coherency either side of the modes was approximately zero. The five-minute peak had a coherency of just under 0.4, which is the same as the coherency at the five-minute peak observed when the port and starboard data are used. When the start time of the data was shifted by 24hrs the coherency of the five-minute peak was reduced to 0.15, which is only slightly less than the coherency of the port and starboard 24 hr shifted data. However, the low-frequency coherency of the simulated data was zero. This implies that the effect observed in the real port and starboard data was not a by-product of the analysis techniques or due to the data's window function and so must be present in the data themselves. It was possible that this low frequency coherency was observed because two data points separated by 24 hrs have probably been measured by the same instrument. To improve our understanding of the low-frequency coherent noise present in 24hr shifted port and starboard data further we investigated the autocorrelation of the port and starboard data. We now describe the results.

7.2.3 Autocorrelation of the port and starboard data

To further investigate the observed low-frequency coherency when the data are shifted by 24 hrs the autocorrelations of the port and starboard timeseries were determined. The autocorrelation is a measure of how well a signal matches a time-shifted version of itself

Figure 7.9: The upper image contains the autocorrelation of the port timeseries and the lower image shows the autocorrelation of the starboard timeseries.

and so it can be used for finding repeated patterns in a signal. The top graph in Figure 7.9 shows the autocorrelation of the port timeseries. The autocorrelation shows dominant peaks that are spaced at regular intervals of 24hrs. The strength of these peaks decreases as the length of the separation in time increases. This indicates that there is some sort of correlation in the noise every 24hrs but that the correlation decreases with time. The autocorrelation of the starboard timeseries (see the lower image in Figure 7.9) is very similar to the autocorrelation of the port timeseries. The main peaks all occur at the same positions in time but the strength of the peaks are very slightly different. To ensure that this effect is not due to the window function the autocorrelation of pure Gaussian noise with the same window functions as the port and starboard data was found. The resulting plot has not been included here but no features were visible implying that the spikes in Figure 7.9 are not due to the window function. These regular peaks could be due to the fact that two individual bins of data that are separated by 24hrs are likely to have been measured by the same instrument.

We have shown that the coherency of non-contemporaneous port and starboard multi-site data is less than the coherency of contemporaneous data. More precisely, the low-frequency coherency of the port and starboard data was approximately zero when the start times were shifted by 12 hrs. However, when the start times were shifted by 24 hrs some coherent data were present. The most probable explanation for this coherent data is that two data points separated by 24 hrs have probably been measured by the same instrument. We shall therefore go on to investigate the coherency of timeseries created from the data of each site individually.

7.3 Comparisons between port and starboard data from individual sites

BiSON timeseries can be created for individual sites while still keeping the port and starboard data separate. For example the port data from Las Campanas B can be compared to the starboard data from Las Campanas B. The cross-amplitude, phase and coherency spectra for Las Campanas B port and starboard data are shown in Figure 7.10. The results for the other five instruments that have port and starboard data available are not included as they are all very similar to this. The three images are very similar to the equivalent graphs for the port and starboard data from all of the sites combined (Figure 7.2). Although the low-frequency cross amplitude is slightly lower for the single site data the coherency is slightly higher than when the data are from multiple sites, which implies that the common noise makes up a higher proportion of the total noise in the single site data. Furthermore, in general, the high-frequency coherency of the individual site data is greater than the high-frequency coherency of the combined-site data. This may be because time overlaps in the data observed from different sites mean that, at certain points in time, the data in the port timeseries may have been observed by a different instrument to the data in the starboard timeseries. For example, assume there to be data available from both Las Campanas B and Sutherland at a given time. The Las Campanas B port data may be of better quality than the Sutherland port data at that time and so the Las

Campanas B data will be included in the final port timeseries. However, the reverse may be true for the starboard data, and so the Sutherland observations are included in the final starboard timeseries. This means that the atmospheric and intrumental noise will no longer be common at that time and the solar noise will be less coherent because of the effect of Doppler imaging, thereby reducing the coherency in the combined-site port and starboard data.

Since the coherency is high for all frequencies, comparing port and starboard data from one site would not be ideal for detecting low-frequency modes. When we compared non-contemporaneous port and starboard timeseries that contained data from different sites the low-frequency coherency was reduced. However, there was still some coherent low-frequency noise present when the start times were separated by 24 hrs, which could be observed because two points of data that are separated by 24 hrs are likely to have been observed by the same instrument. If this is the case then there should be some low-frequency coherent noise present in port and starboard data from an individual instrument that have start times separated by 24hrs. Therefore, we now describe the coherency found when the start times of port and starboard data from one site only were separated by 24 hrs.

7.3.1 *Non-contemporaneous data from the same site*

Figure 7.11 shows the coherency found when the Las Campanas B port data is shifted to start 24hrs after the Las Campanas B starboard data. Clearly there is a significant amount of low-frequency coherent noise. Although not plotted here the coherency found using data from the other sites also show significant amounts of low-frequency coherent noise. This supports the conclusion that the low-frequency noise in the 24hr shifted data is due to measurements being taken by the same instruments. In the combined site timeseries data separated by 12hrs will have been measured by two different instruments and so it is not possible to find the coherency for data from the same site that is shifted by 12hrs. However, this also explains why the low-frequency coherency is approximately

Figure 7.10: The top panel contains the cross-amplitude spectrum of the Las Campanas B port only and starboard only data. The middle graph shows the phase of the two datasets and the bottom graph gives the coherency.

Figure 7.11: The coherency of the port and starboard Las Campanas B data found when the start time of the port data was 24hrs after the start time of the starboard data.

zero when the data is shifted by 12hrs.

When the start of the timeseries are separated by 24 hrs any solar noise and atmospheric scintillation noise will no longer be common. It is possible that there will be a small amount of coherent noise due to atmospheric transparency fluctuations. Some of this noise is caused by the path length through the atmosphere of light arriving at the instrument. This will be approximately the same for two measurements that are taken 24hrs after one another. If there is a regular inaccuracy in the pointing this could also lead to coherent noise in port and starboard data observed from the same site but separated in time by 24 hrs.

The high level of coherency observed between the contemporaneous port and starboard data means that using these sets of data to search for low-frequency p modes would not be ideal. However, some of the noise sources that are coherent in data measured by the same instrument will not be common in data from two different instruments and so the low-frequency coherency may be reduced. Although some success has been achieved in reducing the coherency at low-frequencies by looking at non-contemporaneous data we would still like to find two contemporaneous BiSON timeseries with as little low-frequency common noise as possible. Therefore, we move on to compare data from different sites. This should reduce the amount of common data present as the atmospheric and instrumental noise will no longer be coherent.

7.4   Comparing the port and starboard data from different instruments

The longitudes at which the BiSON sites are positioned around the world means that often there are data available from more than one site at a particular time. When observations from the sites are combined to form a coherent timeseries the 'best quality' data are chosen. However, it is possible to compare data from two different sites, for example, a comparison can be made between Las Campanas B observations and Mount Wilson observations. This will provide further insight into the structure of the noise present in BiSON frequency-power spectra. We begin by considering a special case where data are available from two different instruments at the same observing site.

7.4.1   Las Campanas A and Las Campanas B

At Las Campanas there are two instruments, named A and B, on the same site. The two Las Campanas instruments share the same mount and are positioned parallel to each other with a separation of $\sim 60\,\mathrm{cm}$. Therefore, as long as both instruments are working, they will both be observing the Sun from approximately the same place at the same time and so there are large amounts of contemporaneous data that can be compared. The electronics controlling the two instruments are completely separate apart from a common ac mains power source and a common timing signal. This means that any noise from the temperature regulators, detectors and Pockels cell drivers are completely independent. The solar noise will be similar but, because of Doppler imaging, not identical. The small separation between the two instruments means that the atmospheric transparency noise will be coherent (Clette, 1993). However, noise from atmospheric scintillation is only coherent over a distance of $\sim 10\,\mathrm{cm}$ in average conditions (Young, 1974) and so is unlikely to be common between the A and B data. As the instruments share the same mount any pointing noise will be common to both instruments. Photon arrival times are totally independent between different detectors and so the photon noise will be incoherent between the A and B instruments.

Figure 7.12: Top panel: The coherency of the Las Campanas A port data and the Las Campanas B port data. Bottom panel: The coherency of the starboard only data for Las Campanas A and B.

The upper image in Figure 7.12 shows the coherency of the port only data from Las Campanas A and Las Campanas B. Above ~ 2500μHz the coherency is similar to the coherency found when looking at the port and starboard data from an individual site (Figure 7.10) and the coherency found when the data from all of the sites were used to create port and starboard timeseries (Figure 7.2). The coherency is approximately 1 at the five-minute peak and then it decreases steadily to a coherency of approximately 0.4 at the high-frequency end of the spectrum. This implies that the high-frequency common noise is still coherent. The observed high-frequency coherency implies that instrumental noise is not the main cause of the high-frequency coherent noise that is observed when the port and starboard data are compared. However, the high-frequency coherent noise could be due to solar granulation, Lorentzian tails, atmospheric transparency noise and/or pointing noise.

The lower image in Figure 7.12 shows the coherency of Las Campanas A and B star-

Figure 7.13: The black crosses represent the Las Campanas A-B starboard cross-amplitude spectrum and the red plus signs represent the Las Campanas A-B port cross-amplitude spectrum.

Figure 7.14: The black crosses show the ratio of the absolute amplitudes of the port and starboard data from Las Campanas A. The red plus signs represent the ratio of the absolute amplitude of the port and starboard data from Las Campanas B.

board only observations. It is interesting to notice that the coherency drops off at a faster rate at high frequencies. Above the acoustic cut-off frequency the coherency declines to less than 0.2 at $12,500\,\mu$Hz. Figure 7.13 shows the Las Campanas A-B starboard cross-amplitude spectrum (black crosses) and the Las Campanas A-B port cross-amplitude spectrum (red plus signs). Although the cross-amplitude of the five-minute modes is the same the two cross-amplitude spectra diverge at high frequencies, with the port cross-amplitude becoming increasingly larger than the starboard cross-amplitude. This, in conjunction with the different coherency levels observed at high frequencies between the Las Campanas A and B port data and the Las Campanas A and B starboard data, implies that the Las Campanas A-B port data contains a more coherent high-frequency noise than the Las Campanas A-B starboard data. Figure 7.14 shows the ratios of the absolute amplitudes of the port and starboard data for Las Campanas A (black crosses)

Figure 7.15: The ratio of the absolute amplitudes of the port and starboard data from Las Campanas A. The starboard data from 1998 and the first half of 1999 have been set to zero.

and B (red plus signs). Both ratios are less than 1 above the acoustic cut-off frequency, which implies that the starboard data is noisier than the port data for both instruments. However, the disparity is significantly larger in the Las Campanas A data. Although not shown here plotting the port-starboard ratios for individual years reveals that the problem is confined to 1998 and the first half of 1999. Figure 7.15 shows Las Campanas A port/starboard amplitude ratios, which were created by setting the Las Campanas A starboard data from 1998 and the first half of 1999 to zero. Clearly this figure is very different from Figure 7.14 and in fact implies that the port data now contains more noise than the starboard data. Therefore the Las Campanas A starboard data from 1998 and the first half of 1999 reduces the overall quality of the starboard data as it is very noisy. This noise is not common to both Las Campanas instruments. We can conclude, therefore, that the level of coherency between Las Campanas A-B starboard observations is reduced by the noisy Las Campanas A data in 1998-99. This analysis has shown that poor quality data in a timeseries can be found by determining the coherency spectrum of data from two different detectors.

Returning to Figure 7.12, there is a smaller proportion of common, low-frequency noise in both the port and the starboard Las Campanas A-B data than when the port and starboard data from the same instrument were compared. The atmospheric transparency and pointing noise observed by the A and B instruments should be the same as the instruments are on the same mount. This also means that the solar noise should be

similar as Doppler imaging will only have small effect at this separation. One source of noise that would no longer be coherent is the atmospheric scintillation noise. Also, as mentioned earlier, a lot of the instrumental noise will not be common in Las Campanas A and B data.

The low-frequency coherency between Las Campanas A and B is still ~ 0.5 and so comparing contemporaneous data from these two instruments would not be optimal for searching for low-frequency modes. However, as with the combined site and individual site data it is likely that looking at non-contemporaneous data will reduce the coherency. Therefore the coherency of non-contemporaneous Las Campanas A and B data was determined and we will now describe the results.

7.4.2 Non-contemporaneous data from Las Campanas A and Las Campanas B

In Section 7.2.2 an unexpected amount of low-frequency coherent data was found when the start times of the combined-site port and starboard timeseries were separated by 24hrs. This coherent data could be observed because two data points separated by 24hrs are likely to have been measured by the same instrument. This theory was supported by the results of Section 7.3.1, as a large amount of coherent data were observed when the Las Campanas B port and starboard timeseries had start times separated by 24hrs. The prime candidates for this coherent noise are instrumental noise and transparency fluctuations.

Figure 7.16 shows the coherency found between the mean Las Campanas A data (found by averaging the port and starboard data) and the mean Las Campanas B data when the start time of the Las Campanas A timeseries was 24hrs after the start time of the Las Campanas B timeseries. As can be seen, although there is very little low-frequency coherent data, the coherency is not zero. As the atmospheric scintillation noise and the solar noise will not be coherent with a 24hr time separation it is likely that this noise is due to guidance errors or transparency noise. This also indicates that the majority of the low-frequency coherency observed in the non-contemporaneous port and starboard Las

Figure 7.16: The coherency between non-contemporaneous mean Las Campanas A and B data. The Las Campanas A data has been shifted to start 24hrs after the Las Campanas B data.

Table 7.1: Combinations of sites and the number of overlapping 2hr timeseries with 100% fill available.

Site 1	Site 2	Number of Port Overlaps	Number of Starboard Overlaps
Las Campanas B	Carnarvon B	0	0
Las Campanas B	Mount Wilson	466	442
Las Campanas B	Narrabri	78	64
Las Campanas B	Sutherland	579	585
Carnarvon B	Mount Wilson	0	0
Carnarvon B	Narrabri	117	102
Carnarvon B	Sutherland	157	131
Mount Wilson	Narrabri	272	194
Mount Wilson	Sutherland	0	0
Narrabri	Sutherland	8	8

Campanas B data is instrumental.

The results for Las Campanas A and B are a special case as they share the same mount. We will now describe the results for two well-separated instruments that are at different observing sites. The only noise source common to contemporaneous data from two different sites is the solar noise and even this will not be identical because of Doppler imaging.

7.4.3 Comparisons between instruments at different sites

Table 7.1 shows the number of overlapping 2hr spectra with 100% fill in the port and starboard data for different combinations of sites. Las Campanas and Mount Wilson are

Figure 7.17: The top panel: The coherency of the Las Campanas B port only data and the Mount Wilson port only data. Bottom panel: The coherency of the starboard only data for Las Campanas B and Mount Wilson.

separated by only $\sim 47°$ in longitude and so there are significant amounts of overlapping data that can be compared for these sites. Las Campanas B was chosen in preference to Las Campanas A for this analysis because it has a higher fill and so there are more 2hr segments with 100% fill available.

Figure 7.17 shows the coherency of Las Campanas B and Mount Wilson data with the top panel showing the port only comparison and the lower panel showing the starboard coherency. Notice that the port and starboard graphs are almost identical, unlike the port and starboard coherency of the Las Campanas A and B data. The low-frequency coherency is smaller than the coherency of Las Campanas A and B. This is because the atmospheric and instrumental noise of these observations are not coherent, as the data are from different stations. Also Doppler imaging means that the observations from each site will be biased towards different parts of the solar disc (see Chapter 8) and so the solar noise observed by each instrument will be slightly less coherent.

Figure 7.18: Top panel: The coherency of the Las Campanas B port data and the Sutherland port data. Bottom panel: The coherency of the starboard only data for Las Campanas B and Sutherland.

Notice that the high-frequency coherency is lower than when port and starboard data from the same instruments were compared. The high-frequency coherency is also smaller than the high-frequency coherency observed in Las Campanas A and B data. This indicates that the majority of coherent noise at high frequencies was from a source that is specific to the individual sites. Therefore, the high-frequency coherency could be due to transparency fluctuations, guidance errors or solar noise, which will be more coherent for instruments at the same site, or an accumulation of these noise sources. The high-frequency coherency that is observed here is similar in value to the high-frequency coherency observed when the simulations were performed using the FLAG data (see Figure 7.4). Therefore this coherency could be caused by the Lorentzian tails of the mode profiles.

There are also a large number of overlapping 2hr spectra with 100% fill observed by Las Campanas B and Sutherland as the two sites are separated by only 91° in longitude.

Notice that even though these two sites are further apart than Las Campanas and Mount Wilson there are more 2hr spectra available. This is because the Sutherland data has a greater fill than the Mount Wilson data. Figure 7.18 shows the port only (top panel) and starboard only (bottom panel) coherency spectra from these two sites. Once again the port and starboard coherency spectra are very similar. The low-frequency coherency is even smaller for these two sites than in the Las Campanas B-Mount Wilson example. This can be explained in terms of Doppler imaging as the solar noise observed by Las Campanas and Sutherland will be less coherent than the solar noise observed by Las Campanas and Mount Wilson because Las Campanas and Sutherland are further apart. The high-frequency coherency is similar to that observed by Las Campanas B and Mount Wilson (Figure 7.17). Again the high-frequency coherency observed here could be caused by the Lorentzian tails. This also supports the theory that the larger high-frequency coherency observed in the noise of instruments from the same site is due to atmospheric transparency, instrumental noise and/or solar noise.

Data from two different sites would be better for searching for low-frequency p modes and g modes than data from the same station as the low-frequency coherency is smaller. As the port and starboard data from a particular instrument are not 100% coherent it is possible that comparing the Las Campanas B port data with the Sutherland starboard data may reduce the low-frequency coherency. However, as can be seen from the top image in Figure 7.19 the coherency is the same as the port-only and starboard-only coherency (Figure 7.18). The lower image in Figure 7.19 shows the coherency of the mean data (port and starboard averaged) from Las Campanas B and Sutherland. The coherency is very similar to that observed in the port-only and starboard-only cases. The advantage of using the mean data is that the fill and quality of the data is improved. However, the fill of the two mean sets of data is still relatively low ($\sim 25\%$) and this will reduce the quality of a frequency-amplitude spectrum formed using data from one site only. Therefore we now consider combining data in such a way as to improve the fill of the data while keeping the amount of common noise to a minimum.

Figure 7.19: Top panel: The coherency of the Las Campanas B port data and the Sutherland starboard data. Bottom panel: The coherency of the mean of the port and starboard data for Las Campanas B and Sutherland. There is little or no difference between these two graphs and the coherency shown in Figure 7.18.

7.5 Combining data from different sites

As we have already seen a lot of overlapping 2hr spectra are available between Las Campanas B and Mount Wilson and between Las Campanas B and Sutherland. Conversely there are no overlapping 2hr timeseries between Sutherland and Mount Wilson (see Table 7.1). A timeseries containing data from both Mount Wilson and Sutherland would provide even more overlapping data that can be compared to the Las Campanas B data. Such a timeseries would have an improved fill compared to a one-site set of data and therefore the quality of a frequency-amplitude spectrum produced by the combined-site timeseries would be improved. Although the fill of a timeseries is less important when calculating the coherency using the short 2hr spectra, as each short spectrum has 100% fill, it does become important when the data are searched for low-frequency modes as it is the frequency-amplitude spectrum created by taking the FFT of the complete timeseries that will be searched.

We have strategically and coherently combined data from all of the different BiSON sites to create two timeseries. Importantly the data from a particular instrument did not appear in more than one timeseries. The main aims of combining the data were to maximize the amount of overlapping data in the two timeseries and to share the fill as evenly as possible between the two sets of data. A timeseries was created using data from Las Campanas B, Narrabri, Izaña and Carnarvon B that had a fill of $\sim 56\%$. From now on we will refer to this timeseries as 'Series 1'. A timeseries consisting of data from Las Campanas A, Mount Wilson, Sutherland and Carnarvon A was also created and had a fill of $\sim 63\%$. We refer to this timeseries as 'Series 2'. For the sites that have both port and starboard data available the mean of the two was used. If there were any time overlaps in the observations from the stations in any one timeseries the highest quality data was included in the final timeseries. Over 8600 overlapping 2hr spectra with 100% fill were available between Series 1 and Series 2.

The cross-amplitude, phase and coherency spectra of Series 1 and Series 2 were found using short 2hr spectra (Figure 7.20). Not only is the low-frequency coherency smaller than the coherency when data from just two different sites are used but the improved fill means the signal-to-noise of the modes will be larger. Notice that the cross-amplitude of the five-minute modes is significantly larger than the low-frequency coherent noise. The only source of low-frequency coherent noise that should be present in Series 1 and 2 is the solar noise and this will not be entirely coherent because of Doppler imaging. The phase is slightly scattered at low frequencies compared to the five-minute mode region meaning the low-frequency noise is not entirely in phase. The coherency over the five-minute mode frequency range is very close to one implying that the modes are coherent.

Clearly this is the lowest coherency produced by a pair of contemporaneous BiSON data investigated in this chapter. The timeseries also have the added advantage of having a relatively high fill, compared to the single site data. Therefore these timeseries will be used to search for low-frequency modes. However, as we have seen many times in this chapter already comparing non-contemporaneous data reduces the coherency further and

Figure 7.20: The cross-amplitude, phase and coherency of Series 1 and Series 2.

Figure 7.21: Top panel: The coherency found between Series 1 and Series 2 when the start time of Series 1 is 12hrs after the start time of Series 2. Bottom panel: The coherency found between Series 1 and Series 2 when the start time of Series 1 is 24hrs after the start time of Series 2.

so we have also found the coherency of Series 1 and Series 2 when the start times are separated by both 12 hrs and 24 hrs.

7.5.1 Non-contemporaneous data from different sites

The data can be made non-contemporaneous by shifting the start time of one of the timeseries by a given amount. Here the start time of Series 1 was shifted by both 12 hrs and 24 hrs. The coherency of the resulting timeseries has been found for both cases and the results are shown in Figure 7.21. As can be seen the coherency of the low-frequency noise is approximately zero in both cases. This is expected as the only source of common noise in the contemporaneous data is the solar noise, which will not be coherent in non-contemporaneous data. The modes in the 24hr shifted data are less coherent than the modes in the 12hr shifted data. Even though this should have less of an effect on low-frequency modes, as they have longer lifetimes, the 12hr shifted timeseries will be used

when searching for low-frequency modes.

7.6 Results of searching BiSON data for low-frequency modes

When contemporaneous, Series 1 and Series 2 have a relatively small low-frequency coherency (compared to say port and starboard data taken by the same instrument). Series 1 and 2 also have a larger fill than when data from a single station is considered, which has repercussions for the quality of their frequency-amplitude spectra. Therefore, we have searched Series 1 and 2 for low-frequency p modes and g modes. We have shown that when the start time of one series is shifted by 12 hrs, so that the data are no longer contemporaneous, the low-frequency coherency is reduced to zero. Hence we have also searched Series 1 and 2 when the start·times were separated by 12 hrs. We have used the statistical tests described in Chapters 3, 4 and 5. We have compared Series 1 and Series 2 for coincident prominent features in the frequency-amplitude domain. We have also searched the frequency-power spectra of Series 1 and 2 individually for prominent features. For all the statistical tests the threshold was set so that there was less than a 1% chance that any detections in a $100\,\mu$Hz range were noise. We will now describe the results. It should be noted that phase changes in the modes when the data are non-contemporaneous are not important as the statistics are designed to search absolute frequency-amplitude spectra.

7.6.1 Searching contemporaneous BiSON data

Figure 7.22 is an Echelle diagram that shows the locations in frequency where prominent spikes or patterns of spikes were found when Series 1 and Series 2 were searched for low-frequency p modes and g modes. The methods for searching for the modes described in Section 5.2 were applied to the timeseries, including the bin shifting strategy. As in Chapter 5 the vertical red dashed lines indicate the frequencies of the daily harmonics that occur at multiples of 11.57μHz. As these harmonics are particularly prominent in

Figure 7.22: An Echelle plot, modulo $147\,\mu$Hz, marking the locations in frequency of occurrences uncovered by the test searches. Locations in frequency where spikes, or patterns of spikes, were found in the same bin, or bins, of each spectrum at levels sufficient to record $P \leq 1\%$ are marked by the black symbols in the middle of each row. A different symbol has been used for each test (see figure legend). We have also recorded prominent spikes or patterns of spikes found by searching both Series 1 (red symbols at bottom of each row) or Series 2 (blue symbols at top of each row) alone. Symbols surrounded by a green square represent the prominent occurrences listed in Table 7.2. The green vertical lines mark locations of the p-mode frequencies predicted by the Saclay seismic model (Turck-Chièze et al., 2001). The navy blue vertical lines show the g-mode frequencies predicted by the M1 model from Provost et al. (2000). The vertical dashed red lines mark locations in frequency that are overtones of the 11.57-μHz diurnal frequency.

the low-frequency BiSON data any detections that were found to lie closer than $0.3\,\mu$Hz to one of the diurnal frequencies were disregarded.

The black symbols in the middle of each row represent the frequencies of spikes or patterns of spikes that were found in the same bin in each of the spectra and which had a probability of less than 1% of occurring by chance. Also shown are the frequencies of prominent spikes or patterns of spikes that passed the tests when Series 1 was searched alone (represented by the red symbols at the bottom of each row) and the spikes or patterns of spikes that passed the tests when Series 2 was searched alone (represented by the blue symbols at the top of each row). The symbols surrounded by a green square correspond to features that lie less than $0.5\,\mu$Hz from the location of one of the model frequencies, predicted by the Saclay seismic model (Turck-Chièze et al., 2001) and the M1 model (Provost et al., 2000). The prominent spikes or patterns of spikes that lie within $0.5\,\mu$Hz of a model frequency have been recorded in Table 7.2. The occurrences listed in Table 7.2 have been identified with particular p-mode components on the basis of their placement in frequency. The errors quoted in Table 7.2 have been calculated in the manner described in Section 3.5. When the outer components of a rotationally split mode have been found the mean of these components has been calculated to give an estimate of the frequency of the centroid of that mode.

The frequencies of the candidates agree well with those found in Chapter 5. However, it is disappointing that there are a large number of candidates observed in Chapter 5 that are not detected here. The threshold levels used here will be larger than those used in Chapter 5 because the low-frequency coherency is larger when Series 1 and 2 are compared. However, it is more likely that the fact that fewer modes are detected here is highlighting the improvement in the overall quality of a timeseries obtained by combining data from all available sites. It is, nonetheless, interesting that all three potentially visible components of the $l = 2$, $n = 8$ mode at $\sim 1395\,\mu$Hz are detected in Series 1 and 2. These modes are actually detected as a triplet. However, only the $m = -2$ component is detected when the BiSON and GOLF data are compared. This could be

Table 7.2: Candidates found to be closer than 0.5μHz to predicted model frequencies of modes.

l	n	m	Frequency (μHz)	Probability (P)	Number of tests passed	Distance from model in frequency (μHz)
1	8	-1	1329.235 ± 0.006	9.5×10^{-7}	3	0.062†
2	8	-2	1393.871 ± 0.007	8.3×10^{-4}	1	
2	8	0	1394.702 ± 0.007	8.3×10^{-4}	1	0.007
2	8	2	1395.469 ± 0.007	8.3×10^{-4}	1	
0	9	0	1407.479 ± 0.007	1.7×10^{-4}	1	0.149
3	8	-3	1449.743 ± 0.008	7.3×10^{-3}	2‡	0.118♯
3	8	-1	1450.540 ± 0.008	7.3×10^{-3}	1‡	0.122†
1	9	-1	1472.434 ± 0.007	9.1×10^{-8}	4	
1	9	$+1$	1473.271 ± 0.007	4.5×10^{-6}	4	
		mean	1472.853 ± 0.005			0.121

†: Difference with model frequency assumes $m = -1$ component lies $-0.4\,\mu$Hz from central frequency.

♯: Difference with model frequency assumes $m = -3$ component lies $-1.2\,\mu$Hz from central frequency.

‡: These modes are detected as part of a triplet but the frequency of the third component does not correspond to the frequency of a model multiplet component.

because $l = 2$ modes are less prominent in GOLF data. Nonetheless, the detections must be regarded with suspicion as the components only passed one of the statistical tests (the triplet test). As we have seen in Chapter 3 ensuring that all candidates pass at least two of the statistical tests dramatically reduces the number of false detections that are made. It is encouraging to notice that the $l = 0$, $n = 9$ mode at $\sim 1407\,\mu$Hz is detected here. This mode is not detected in Chapter 5 when the BiSON data are searched alone. This implies that is it beneficial to split the BiSON data into contemporaneous timeseries to search for low-frequency modes.

7.6.2　Searching non-contemporaneous BiSON data

As we have seen in Chapter 4 the advantage gained by using two sets of data is greatest when the level of common noise is zero. Furthermore at low frequencies the lifetimes of the modes are similar to, if not longer than, the length of the observations. Therefore, if only a small shift in the start times is implemented, the majority of the signal from low-frequency modes should remain commensurate. Therefore, we have searched Series

Figure 7.23: An Echelle diagram showing the results when the start of Series 1 begins 12hrs after the start of Series 2. The lines and symbols have the same meaning here as in Figure 7.22.

1 and Series 2 for low-frequency p modes and g modes when the start time of Series 1 was 12 hrs after the start time of Series 2. The methods for searching for the modes described in Section 5.2 were applied to the non-contemporaneous timeseries, including the bin shifting strategy.

The results of this search can be seen in Figure 7.23. Once again the results are displayed in the form of an Echelle diagram. The candidates shown have less than a 1% chance of being due to noise. The symbols and vertical lines have the same meaning as in Figure 7.22. Candidates that lie within 0.5 μHz of the predicted frequencies are again highlighted by a green square. These candidates are also shown in Table 7.3. Notice in Figure 7.23 that there is a detection close to the $l = 1$, $n = 7$ mode at $\sim 1186\,\mu$Hz that is not highlighted by a green box. This is a multiplet detection and although one of the detected components was close to a predicted frequency, the other detected component was not. Also notice there is a detection close to the daily harmonic at $\sim 1354\,\mu$Hz.

Table 7.3: Candidates found to be closer than 0.5μHz to predicted frequencies of modes from the Saclay seismic model or the M1 model when Series 1 started 12hrs after Series 2.

l	n	m	Frequency (μHz)	Probability (P)	Number of tests passed	Distance from model in frequency (μHz)
1	8	-1	1329.237 ± 0.006	1.97741×10^{-7}	3	0.060
2	8	0	1394.701 ± 0.007	7.4×10^{-3}	1	0.006
2	8	$+2$	1395.472 ± 0.007	7.4×10^{-3}	1	
1	9	-1	1472.432 ± 0.007	3.5×10^{-9}	3	
1	9	$+1$	1473.254 ± 0.007	2.8×10^{-7}	2	
		mean	1472.843 ± 0.005			0.130

Again this is a multiplet detection, where the plotted frequency is the mean frequency of the two observed components. In this case the individual components lie on either side of the harmonic frequency. Disappointingly less modes are detected when the non-contemporaneous data are searched than when Series 1 and 2 are contemporaneous.

7.7 Summary

BiSON is a network of instruments that are positioned at strategic longitudes across the surface of the Earth so that often there is more than one instrument observing the Sun. Furthermore, many of the instruments have port and starboard detectors. The data observed by these detectors and instruments can be kept separate to form contemporaneous sets of BiSON data. The focus of this chapter was to find and compare contemporaneous BiSON data, with the aim of detecting low-frequency modes. A prerequisite for any pairs of data that we wished to search was that they had a low coherency at low-frequencies, to minimize detection threshold levels.

The first sets of BiSON data that were compared took advantage of the port and starboard detectors. Timeseries were created that contained data from all of the instruments with port and starboard detectors. However, the data observed by the port and starboard detectors were kept separate. We therefore had a combined-site port timeseries and a combined-site starboard timeseries. The two timeseries had a large coherency (above

0.8) at low-frequencies and so were not ideal for searching for low-frequency modes. This common noise is likely to be solar, atmospheric and instrumental in origin. Coherent data was also observed at high frequencies. The Lorentzian tails of the mode profiles can account for some, but not all, of the observed high-frequency coherency. Comparison with the high-frequency coherency found between data from different instruments (e.g. Figures 7.12, 7.17 and 7.18) implies that the majority of the high frequency coherency is due to transparency fluctuations, guidance errors and/or solar noise.

We have shown that shifting the start times of the data by 12 hrs and 24 hrs does not significantly reduce the coherency of low-frequency modes as they have such long lifetimes. When the combined-site port and starboard data were shifted so that the start times were separated by 12 hrs the low-frequency coherency was reduced to zero. Although not performed here this data could, therefore, be searched for low-frequency p modes and g modes. When the start times were separated by 24 hrs the low-frequency coherency was determined to be greater than zero. The low-frequency coherency is, most likely, observed because two points of data that are separated by 24 hrs will probably have been observed by the same instrument. This theory is supported by the fact that coherent data is also observed at low-frequencies when the port and starboard data, with the start times separated by 24 hrs, from one site only are compared. Additionally, the low-frequency coherency is zero when data from different sites are compared and the start times are separated by 24 hrs.

There are two instruments at the Las Campanas station, called A and B, and contemporaneous data from these two instruments were compared. Although the low-frequency coherency was smaller than when the port and starboard data from the same instrument were compared it was still reasonably large ($\sim$ 0.5). The majority of instrumental noise and atmospheric scintillation noise was independent but the solar noise, atmospheric transparency noise and pointing noise was still common. The high-frequency coherency observed between the port Las Campanas A and B data was different to the starboard Las Campanas A and B data. This was because the starboard Las Campanas A data in

1998 and the first half of 1999 was poor in quality.

The smallest low-frequency coherency (~ 0.4) observed between contemporaneous BiSON data was found when data from two different sites were compared. Here we compared data from Las Campanas B and Mount Wilson and Las Campanas B and Sutherland. Any coherent noise that is shared by data observed at different sites is solar in origin. To improve the quality of the timeseries the port and starboard data were averaged. The problem with comparing data from one site only was that there was only a small fill, which has repercussions for the quality of a frequency-amplitude spectrum created from the data. However, it was possible to coherently combine data from various different sites to create two separate timeseries, Series 1 and Series 2. The data from each instrument were only included in one of the timeseries. Series 1 contained data from Las Campanas B, Narrabri, Izaña and Carnarvon B, while Series 2 contained data from Las Campanas A, Mount Wilson, Sutherland and Carnarvon A.

When Series 1 and 2 were compared the low-frequency coherency was determined to be ~ 0.3, which is relatively low. Therefore, these data were searched for low-frequency modes. The list of candidates detected when Series 1 and Series 2 were searched agreed well with the list of candidates found when the BiSON and GOLF data were searched. However, a lot of the candidates observed in Chapter 5 were not observed here. The only modes that were detected here and not in Chapter 5 (the $l = 2$, $n = 8$, $m = 0$ and $+2$ modes at $\sim 1395\,\mu$Hz) only passed one test and so must be regarded with suspicion. Although we have combined data from different sites to improve the fill of the timeseries the duty cycles of Series 1 and Series 2 are still less than the fill of the BiSON data searched in Chapter 5. This could explain why less modes are detected here. However, it was encouraging that one mode (the $l = 0$, $n = 9$ mode at $\sim 1407\,\mu$Hz) was detected here that was not detected when only the BiSON data used in Chapter 5 were searched. When the start time of Series 1 was shifted to be 12 hrs after the start time of Series 2 the low-frequency coherency was zero. Therefore, we also searched Series 1 and 2 for low-frequency modes when they were not contemporaneous. However, no new mode

candidates were detected and the list of modes was again disappointingly short.

A thorough understanding of the noise in unresolved Doppler velocity observations is useful for detecting low-frequency modes. Throughout this thesis we have been commenting on the fact that the solar noise observed by two different instruments will not be identical because of Doppler imaging. In the next chapter we describe in detail how Doppler imaging affects unresolved solar Doppler velocity observations and we also discuss differences in the data observed by the port and starboard instruments.

8. WEIGHTING OF THE SOLAR IMAGE

8.1 Introduction

BiSON instruments make unresolved observations of the Sun, which means that the intensity recorded by the instrument is integrated over the entire solar disc. However, the response of a BiSON instrument to different areas of the solar disc is not homogenous and so the data obtained do not represent a uniform average over the entire surface. This uneven weighting is known as Doppler imaging and is caused by several different factors, which will be described in this chapter. The first two effects to be considered are properties of the Sun itself; the effect of the Sun's rotation and the weighting due to limb darkening. These effects have been considered previously by Brookes et al. (1978b) and the results are well known. Here we add a third factor, which is instrument specific as it is dependent on the positioning of the detector with respect to the solar image observed at the vapour cell. Once this effect is included various other observational parameters, such as the size of the observed solar disc and the angle of inclination of the rotation axis, can be altered to make the calculations more realistic.

As we have seen, in the previous chapter, the majority of BiSON instruments take two sets of observations that are known as port and starboard observations. The detectors responsible for making these observations lie on opposite sides of the vapour cell (see Figure 1.2). In this chapter we show that the port measurements are biased towards the limb of the image of the solar disc in the vapour cell that lies closest to the port detector. Similarly the starboard observations are biased towards the limb of the image of the solar disc that is closest to the starboard detector. We will now describe each of the three weighting processes mentioned above individually and then combine the effects
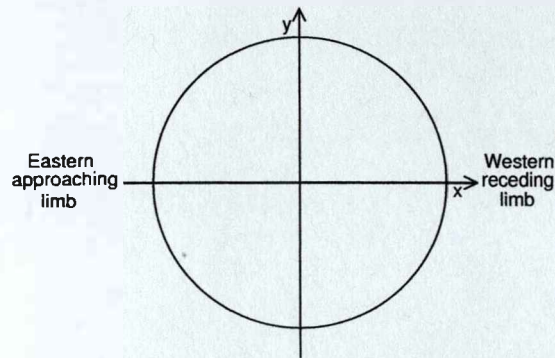
Figure 8.1: The image of the solar disc can be described by Cartesian coordinates that are constrained by the condition $x^2 + y^2 \leq R_\odot^2$, where $R_\odot$ is the radius of the Sun.

Figure 8.2: The Fraunhofer line observed when looking at the centre of the solar disc where the line-of-sight velocity due to the solar rotation is zero (black crosses). Also plotted is a Gaussian approximation to this line (red line).

together to produce a description of how the observations are weighted across the solar disc.

First, it is important to set up a coordinate system to describe the solar disc so that the results from each individual effect can be coherently combined. In what follows we describe the solar disc in terms of Cartesian coordinates with the origin at the centre of the solar disc (as shown in Figure 8.1). The edge of the observed Sun is constrained by the equation $x^2 + y^2 = R_\odot^2$, where $R_\odot$ is the observed radius of the solar disc.

Before we consider the weighting across the solar disc it is useful to discuss the solar Fraunhofer absorption line that is observed by the instrument.

8.2 Solar absorption line

If resolved observations are taken at the very centre of the solar disc, where $x = 0$ and $y = 0$ in Figure 8.1, the line-of-sight velocity due to the solar rotation will be zero. Doppler velocity observations of the centre of the solar disc have been made by Themis, which is a 90cm solar telescope located at Izaña, Tenerife. The observations used here were provided by R. Simionello in a private communication. The spectral line observed by Themis can be seen in Figure 8.2. Assuming the natural width of this line is caused by thermal Doppler broadening only, the shape of the line will follow that of a Gaussian distribution. Thermal Doppler broadening occurs because the emitting gas molecules in the Sun have random motions that mean the wavelengths at which photons from the Sun are emitted are Doppler shifted by a random amount. The spectral line shown in Figure 8.2 is well represented by a Gaussian centred on 769.897nm, with a FWHM, Γ, of 0.018nm and a 77% depth. The equation of the Gaussian representation of the absorption line is given by

$$I(\lambda) = 1 - 0.77 \exp\left[4\ln 2\left(\frac{(\lambda - \lambda_c)^2}{\Gamma^2}\right)\right]. \qquad (8.1)$$

Here $I(\lambda)$ is the intensity of light observed from the Sun, where an intensity of 1 means that no light has been absorbed and an intensity of zero means that all of the light at that wavelength is absorbed. Wavelength is represented by λ and λ_c is the central wavelength of the line. The red line in Figure 8.2 shows the Gaussian representation of the Fraunhofer line given by equation 8.1. As mentioned, in this example $\lambda_c = 769.897$nm but it is important to note that the wavelength of the line centre is dependent upon the date and time at which the observations are made for reasons that follow. Both Themis and the BiSON spectrometers observe the potassium line. If the line-of-sight velocity between the observer and the target, in this case the Sun, is zero the centre of the line λ_c has a wavelength of 769.898 nm. Any line-of-sight velocity between the observer and the Sun shifts the centre of the line away from this value by an amount that is dependent on the magnitude of the line-of-sight velocity. The line-of-site velocity between the observing

site and the Sun will change as the Earth orbits the Sun and as the Earth spins on it's own axis i.e. as v_{orb} and v_{spin} change. Therefore the central wavelength of an absorption line received on Earth from the Sun varies throughout the year and over the course of a day.

Now consider an observed spectral line centered on a wavelength, λ_c, with a Doppler FWHM, V_d, that is caused by the thermal motions of the absorbing gas. We take V_d to be the same as the width of the solar absorption line observed by Themis. The change in wavelength, $\Delta\lambda$, caused by a Doppler velocity, v, is given by

$$\Delta\lambda = \lambda\frac{v}{c}, \tag{8.2}$$

where c is the speed of light and λ is the original wavelength. We can, therefore, use the FWHM observed by Themis, $\Delta\lambda = \Gamma = 0.018$nm, to find the FWHM of the line in terms of velocity and so $V_d = 6892\,\text{m s}^{-1}$. It should be noted that this is, in fact, larger than the width expected due to Doppler thermal broadening only. This is because the width of the solar line is increased by the Doppler velocity of turbulence in the Sun's atmosphere. However, as we have already seen a Gaussian shape is still a good approximation for the observed solar Fraunhofer line.

We now go on to investigate the effect of rotation on the observed solar absorption line.

8.3   Weighting due to solar rotation

As the Sun rotates one limb of the solar disc moves towards an observer on Earth while the other limb moves away from an observer. In other words the rotation causes a line-of-sight velocity between an observer and the Sun that varies across the solar disc. This line-of-sight velocity will affect Doppler velocity observations, such as those made by the BiSON instruments. Resolved Doppler velocity observations made of the eastern limb of the solar disc will be blue shifted while observations of the western limb will be red

shifted. In what follows the effect of the variation in the line-of-sight velocity across the solar disc on unresolved Doppler velocity measurements will be investigated.

Consider the Sun to be a body that is circularly symmetric about its rotation axis and assume that it rotates with a constant angular velocity, ω. This is a rather simplistic assumption as the surface of the Sun actually exhibits differential rotation. However, it will still be able to give a general idea of how the rotation affects the Doppler velocity measurements. We take the rotation to be in an anti-clockwise direction. We also assume that the Sun is observed from a considerable distance along a line of sight that is perpendicular to the rotation axis. We have taken the rotation axis to be the y axis in figure 8.1. Again this is a simplistic assumption and later in this chapter we will investigate the effect of varying the orientation of the rotation axis. If the angular velocity is in an anti-clockwise direction it can be shown that the radial velocity along the line of sight, V_r, of a point on the surface with coordinates (x, y) is given by

$$V_r = \omega x. \tag{8.3}$$

The material at the solar equator rotates once approximately every 25 days and so we will take the Sun's angular velocity to be $2.9 \times 10^{-6}\,\mathrm{rad\,s^{-1}}$. Using equation 8.3 and taking $x = R_\odot$ we can see that the maximum line-of-sight velocity due to the solar rotation is approximately $2000\,\mathrm{m\,s^{-1}}$. Lines of constant velocity occur at constant values of x and so they are parallel to the rotation axis and independent of the y coordinate (see Figure 8.3).

As shown in Figure 8.3 the solar disc can be split into contours of constant rotation that are evenly spaced. These lines of constant rotation correspond to values of constant x in the Cartesian coordinate system described by Figure 8.1. Here, the solar disc was split into 501 evenly spaced values of x that range in position from one edge of the solar disc to the other, and so $-6.955 \times 10^8\mathrm{m} \leq x \leq 6.955 \times 10^8\mathrm{m}$. Let x_i represent a different value of the x coordinate for each $i = 0 \ldots 500$. We also divided the solar disc 501 times

Figure 8.3: The black vertical lines represent contours of constant velocity on the solar disc.

in the vertical direction so that $-6.955 \times 10^8 \text{m} \leq y \leq 6.955 \times 10^8 \text{m}$. Let y_j represent a different value of the y coordinate for every $j = 0 \ldots 500$. Let a line profile, centred on a wavelength, $\lambda_{i,j}$, be observed from every position on the solar disc where x and y coordinates cross. The line-of-sight velocity, brought about by solar rotation, means that the central wavelength of each line is Doppler shifted. Let λ_c be the central wavelength of the line observed at the centre of the solar disc. This means that the value of λ_c depends only on the line-of-sight velocity between the observer and the Sun and is independent of the solar rotation. The central wavelength of the line observed at (x_i, y_j), $\lambda_{i,j}$, is related to the wavelength of the line observed at the centre of the solar disc, λ_c, by

$$\lambda_{i,j} = \lambda_c \left(1 + \frac{\omega x_i}{c} \right). \tag{8.4}$$

The intensity of light that can be observed from each region of the Sun is related to the area of that region. Let each coordinate (x_i, y_j) correspond to a 'pixel' on the Sun. In both the x and y directions each pixel extends half way to its neighbouring point. In other words the pixel is a rectangle limited by the regions $x_i - 0.5(x_i - x_{i-1}) \leq x \leq x_i + 0.5(x_{i+1} - x_i)$ and $y_i - 0.5(y_i - y_{i-1}) \leq y \leq y_i + 0.5(y_{i+1} - y_i)$. Let us not forget that the pixels are limited by the edge of the solar disc $(x^2 + y^2 = R_\odot^2)$. Therefore some of the pixels will not form complete rectangles. Figure 8.4 shows a pixelated image of the Sun. The two black crosses are at the centres of pixels that are complete rectangles. The edges of each pixel are depicted by the black lines. The red triangle highlights the

Figure 8.4: An example of pixels on the Sun. The edges of the pixels are marked by the black lines. The two crosses show examples of complete pixels. The crosses are at the centre of the pixels and so show the (x_i, y_j) coordinates. The red triangle shows an (x_i, y_j) coordinate whose pixel is not a complete rectangle.

x and y coordinates of an incomplete pixel. The intensity of light observed on Earth from each pixel is proportional to the area of the pixel. Pixels whose shapes are limited by the edge of the solar disc will have a smaller area than the rectangular pixels and so less light will be observed from an incomplete pixel than a complete pixel. Furthermore, as all of the incomplete pixels are from the edge of the solar disc they will be more susceptible to the effects of limb darkening (see Section 8.4), which will reduce their intensity further. This means that the incomplete pixels contribute less to the total unresolved observations than the complete pixels. As the vast majority of the pixels are complete ($\sim 99\%$) it is possible to approximate the behaviour of unresolved solar observations by only considering complete pixels.

Let a line profile be observed from each of the complete pixels on the solar disc. Now, as the area of each pixel that we are considering is the same and as we are neglecting the effect of limb darkening at this stage, the solar absorption line profile observed by an instrument from each pixel has the same depth. Furthermore, as each pixel only constitutes a small area on the solar disc the difference in the line-of-sight velocity due to the solar rotation across the pixel will be small ($\sim 8\,\mathrm{m\,s^{-1}}$). Hence, the absorption line from each pixel will not be noticeably broadened by the solar rotation. Therefore the width of the absorption line observed from each pixel will be the same. We have, therefore, taken the line profile observed from each pixel to be the same as the Fraunhofer line observed by Themis, and so the FWHM of each line is 0.018 nm and the minimum

intensity of each line is 0.23. It is now necessary to consider how the instrument will observe these profiles.

8.3.1   Instrumental observations

We take the instrumental response in the vapour cell to have a Gaussian absorption profile with a FWHM of $700\,\mathrm{m\,s^{-1}}$. The height of the Gaussian is not important as we are only interested in how the observed intensity varies over the solar disc and not the value of the absolute intensity observed from each pixel.

Let the instrumental response of the vapour cell be centred on a wavelength λ_{inst} in a rest frame with no line-of-sight velocity. Some of the light received in the vapour cell from the Sun is absorbed by the gas in the vapour cell and some of the light is transmitted. The absorbed light is then re-emitted, allowing it to be detected. In other words the intensity of light that is detected by an instrument is proportional to the intensity of light that is absorbed by the vapour. The contribution of each region of the Sun to the intensity of the absorbed light will now be calculated.

As mentioned in Chapter 1, the light entering a BiSON instrument is filtered before it reaches the vapour cell. Let the light entering the vapour cell be restricted to the wavelength range $\lambda_1 \leq \lambda \leq \lambda_2$. Let $L_{i,j}(\lambda)$ be the line profile that corresponds to the coordinate (x_i, y_j) on the solar disc. The total intensity of light received at the vapour cell from a pixel (x_i, y_j) is given by

$$I_{i,j}^{\mathrm{r}} = \int_{\lambda_1}^{\lambda_2} L_{i,j}(\lambda)\mathrm{d}\lambda. \tag{8.5}$$

The line profiles, $L_{i,j}(\lambda)$, that are received at the vapour cell from two pixels with different locations on the solar disc can be seen in panels (a) and (b) of Figure 8.5 (the orange and black lines), as can the red instrumental absorption profile (red dot-dashed line). The wavelength of the instrumental profile is only affected by the influence of the magnetic field across the vapour cell. In other words it is not shifted by a line-of-sight velocity.

(a) Solar Fruanhofer line received at vapour cell from a pixel at $x = -0.8R_\odot$, $y = 0$

(b) Solar Fruanhofer line received at vapour cell from a pixel at $x = +0.8R_\odot$, $y = 0$

(c) Light transmitted from a pixel at $x = -0.8R_\odot$, $y = 0$

(d) Light transmitted from a pixel at $x = +0.8R_\odot$, $y = 0$

(e) Light absorbed from a pixel at $x = -0.8R_\odot$, $y = 0$

(f) Light absorbed from a pixel at $x = +0.8R_\odot$, $y = 0$

Figure 8.5: Panels (a) and (b) show the spectrum of light received from two different pixels on the Sun. In panel (a) the solar absorption line (the orange line) is blue shifted because the pixel is on the side of the Sun which is approaching an observer on Earth. In panel (b) the solar absorption line (the black line) is red shifted as the pixel is on the side of the Sun that is moving away from an observer on Earth. Panels (c) and (d) show a comparison between the light received and the light transmitted. In panels (a) to (d) the red dot-dashed line represents the red component of the instrumental absorption line, whose wavelength is only shifted by the influence of the magnetic field across the vapour cell. The difference between the light received and the light transmitted is the light that is absorbed. Panels (e) and (f) show the difference between the light received and the light transmitted, i.e. the light absorbed by the vapour.

The solar absorption line in Figure (a) (the orange line) has been blue-shifted because the pixel it comes from is moving towards an observer on Earth. The solar absorption line in Figure (b) (the black line) has been red-shifted because the pixel it comes from is moving away from an observer on Earth. The red instrumental line (the red dot-dashed line) in both panels is on the red limb of both absorption lines because is has been shifted to a longer wavelength by the magnetic field that is placed across the instrument vapour cell.

To determine the intensity of light that is absorbed by the instrument we must compare the intensity of light that is received from the pixel of the Sun with coordinates (x_i, y_j) with the intensity of the transmitted light. The total intensity of light that is transmitted by the instrument from this region of the solar disc between the wavelengths λ_1 and λ_2, $I_{i,j}^t$, is then given by

$$I_{i,j}^t = \int_{\lambda_1}^{\lambda_2} L_{\text{inst}}(\lambda) L_{i,j}(\lambda) d\lambda, \tag{8.6}$$

where $L_{\text{inst}}(\lambda)$ is the instrument profile. The spectra of the light that are transmitted by the vapour cell from the two chosen pixels can be seen in panels (c) and (d) of Figure 8.5 (the purple and green dashed lines).

The total intensity of light absorbed by the gas in the vapour cell from a pixel with coordinates (x_i, y_j) is given by the difference between the intensity of light received from the Sun and the intensity of light that is transmitted by the vapour:

$$I_{i,j}^a = I_{i,j}^r - I_{i,j}^t. \tag{8.7}$$

The intensity of light absorbed from a pixel with the coordinates $x = -0.8R_\odot$, $y = 0$, is equivalent to the area between the orange solid line and the purple dashed line in panel (c) of Figure 8.5. The spectrum of light absorbed from this pixel can be seen in panel (e) of Figure 8.5. Let the area under this sky blue curve be A_{I1}. Similarly the intensity of light absorbed from a pixel with the coordinates $x = +0.8R_\odot$, $y = 0$, is equivalent to

the area between the black solid line and the green dashed line in panel (d) of Figure 8.5. The spectrum of light absorbed from this pixel can be seen in panel (f) of Figure 8.5. Let the area under the pink curve be A_{I2}. Clearly A_{I2} is smaller than A_{I1}, meaning that the red wing of the instrument will observe more light from the pixel with the coordinates $x = -0.8R_\odot$, $y = 0$ than from the pixel with the coordinates $x = +0.8R_\odot$, $y = 0$.

8.3.2 The weighting of the solar disc due to the Sun's rotation

Figure 8.6 shows a contour map that displays the relative weighting from different areas of the solar disc. The weighting has been scaled relative to the intensities of the most and least prominent pixels and so does not represent the true intensity observed by the instrument from a particular pixel. The pixels that contribute the most to the total observed intensity have been given a weighting of 100. The pixels that contribute the least to the total intensity have been given a weighting of 0. The weighting of all other pixels are distributed within the range 0-100. In other words, the larger the number assigned to each contour, the greater the intensity of light that is observed from that region on the solar disc.

The central wavelength of the Fraunhofer line observed at the centre of the solar disc, λ_c, depends on the line-of-sight velocities due to the Earth's spin and orbit and the apparent gravitational redshift velocity. We define the line-of-sight velocity between the observer and an instrument, v_{los}, as

$$v_{los} = v_{orb} + v_{spin} + v_{grs}, \tag{8.8}$$

where v_{orb} is the line-of-sight velocity due to the Earth's orbit around the Sun, v_{spin} is the line-of-sight velocity because of Earth's spin on its own axis, and v_{grs} is the gravitational redshift velocity. A description of each of these velocities is given in Chapter 1. Here, for simplicity, we have taken the line-of-sight velocity between the observer and the Sun, v_{los}, to be zero. Therefore the wavelength of a solar line originating from the centre of
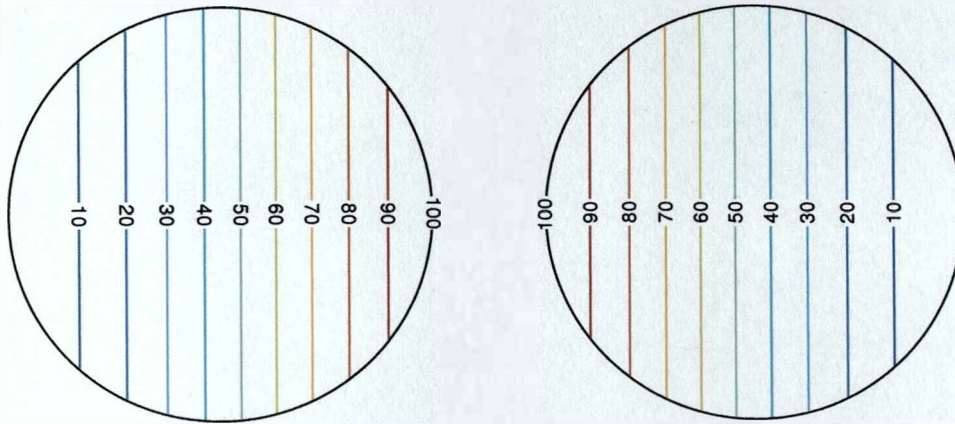
Figure 8.6: A contour map showing the effect of rotation on the contribution of different areas to the total observed intensity. The left hand panel shows the results for the blue wing instrumental response and the right hand panel shows the results for the red wing. Here $v_{\text{los}} = 0\,\text{m}\,\text{s}^{-1}$

the solar disc, λ_c, has the same value as the absorption wavelength of the potassium in the vapour cell ($\lambda_c = 769.898\,\text{nm}$). We take the magnetic field across the vapour cell to have a strength of 1.8kG and so the two Zeeman components of the potassium line are symmetrically split by $5200\,\text{m}\,\text{s}^{-1}$. Therefore the two instrumental components are separated in wavelength by $0.013\,\text{nm}$ and so the centre of the red wing component of the instrumental profile has a wavelength of $769.905\,\text{nm}$ and the centre of the blue wing component has a wavelength of $769.891\,\text{nm}$.

The left hand panel of Figure 8.6 shows the weighting due to the rotation observed by the instrumental blue wing and the right hand panel shows the weighting observed by the instrumental red wing. Lines of constant weighting are vertical straight lines, as expected, given that these also represent lines of constant line-of-sight velocity due to the Sun's rotation. The blue wing is weighted most strongly towards the limb that is moving away from an observer on Earth, while the red wing is weighted towards the limb that is approaching Earth. This is counterintuitive but can be explained when we look back at Figure 8.5. There we saw that the instrumental red wing observes more light from the pixel on the approaching side of the Sun ($x = -0.8R_\odot$) than the pixel on the receding side of the Sun ($x = +0.8R_\odot$). The reverse is true for the blue wing. As $v_{\text{los}} = 0\,\text{m}\,\text{s}^{-1}$ the red and blue wing instrumental components are symmetrically split about the observed

Figure 8.7: A contour map showing the effect of rotation on the contribution of different areas to the total observed intensity. The left hand panel shows the results for the blue wing instrumental response and the right hand panel shows the results for the red wing. Here the line-of-sight between the Sun and the observer is $1100\,\mathrm{m\,s^{-1}}$.

line centre, λ_c, and so the two panels in Figure 8.6 are mirror images.

As v_{los} varies with time so too does the observed weighting across the solar disc. The results in Figure 8.7 were calculated when the line-of-sight velocity between the Sun and an observer was $1100\,\mathrm{m\,s^{-1}}$. This is approximately the maximum orbital velocity between the Sun and the Earth in any one year. The variation in the blue wing observations across the solar disc is now less steep. The gradient of the weighting of the red wing observations is steeper, with the observations weighted even more heavily towards the approaching limb of the solar disc than when $v_{\mathrm{los}} = 0\,\mathrm{m\,s^{-1}}$.

In the next section we consider the effect of limb darkening on the weighting of Sun-as-a-star observations. We will then combine the effect of limb darkening with the results found in this section for the rotation.

8.4 Limb-darkening

The observed intensity of the limb of the Sun is diminished compared to the intensity at the centre of the solar disc. This effect is known as limb darkening. The limb darkening across the Sun can be calculated using the limb darkening function, which is given in

Figure 8.8: A birds-eye view of the geometric set up for calculating the effect of solar limb darkening. The observer is positioned at P, a distance d from the near edge of the Sun. The solar radius is $R_\odot$ and point O shows the centre of the Sun. If observations are made of a point S on the solar disc then ψ is the angle between the vector from the observer to this point and the radius vector between the centre of the Sun and S.

Allen (1955) as

$$\frac{I(\psi)}{I(0)} = 0.57 + 0.43 \cos \psi, \qquad (8.9)$$

where $I(\psi)$ is the intensity at a heliocentric angle ψ, $I(0)$ is the intensity if no limb darkening were present and ψ is the angle between the radius vector and the line of sight, which is shown in Figure 8.8. Figure 8.8 shows a birds-eye view of the Sun with the centre of the Sun at O. The observer is positioned at P and is observing a point S on the solar disc. The radius of the Sun is denoted by $R_\odot$ and the observer is a distance d from the edge of the Sun. The heliocentric angle, ψ, is then the outside angle between OS and PS. To calculate the intensity weighting across the solar disc due to limb darkening ψ must be related to the cartesian description of the solar disc. Projecting Figure 8.8 onto a flat image of the solar disc that would be seen by the observer at P (as in Figure 8.9), S will be observed as a point a distance r from the origin, O. Assuming the Sun is spherical, every position on the Sun that is a distance r from O will have the same heliocentric angle, ψ, and these positions are shown by the dashed circle in Figure 8.9. Looking back at Figure 8.8, r is given by the line connecting S with OP. Let this line

Figure 8.9: The view of the solar disc from P. All positions on the solar disc a distance r from the centre of the disc, O, such as S will have the same heliocentric angle, ψ.

meet OP at a point, T, where T is positioned so that the angle between ST and OP is a right angle. Let r_1 be the distance between T and the edge of the Sun on the line OP.

To evaluate equation 8.9 it is necessary to find an expression for $\cos\psi$ in terms of r, $R_\odot$, and d as these are all either known or measurable. Let θ be the angle between PO and PS. From Figure 8.8 it can be seen that

$$
\begin{aligned}
\sin\theta &= \frac{r}{\text{PS}}, \\
&= \frac{r}{\sqrt{r^2 + (d+r_1)^2}}, \\
&= \frac{r}{\sqrt{r^2 + (d + R_\odot - OT)^2}}, \\
&= \frac{r}{\sqrt{r^2 + \left(d + R_\odot - \sqrt{R_\odot^2 - r^2}\right)^2}}.
\end{aligned}
\tag{8.10}
$$

Now, using the sine rule on the triangle OPS and taking the angle at S to be $(\pi - \psi)$ rad we have

$$
\begin{aligned}
\frac{R_\odot}{\sin\theta} &= \frac{R_\odot + d}{\sin(\pi - \psi)}, \\
&= \frac{R_\odot + d}{\sin\psi}.
\end{aligned}
\tag{8.11}
$$

Then, using $\cos^2 \psi + \sin^2 \psi = 1$, it can be shown that

$$\cos \psi = \frac{\sqrt{R_\odot^2 - (R_\odot + d)^2 \sin^2 \theta}}{R_\odot}. \tag{8.12}$$

By using equation 8.10 to replace $\sin \theta$ in equation 8.12 an expression for $\cos \psi$ in terms of r, $R_\odot$ and d can be found:

$$\cos \psi = \frac{\sqrt{R_\odot^2 r^2 + R_\odot^2 \left(d + R_\odot - \sqrt{R_\odot^2 - r^2}\right)^2 - (R_\odot + d)^2 r^2}}{R_\odot \sqrt{r^2 + \left(d + R_\odot - \sqrt{R_\odot^2 - r^2}\right)^2}}. \tag{8.13}$$

Finally this can be related to the initial Cartesian set of coordinates, which were used in Section 8.3 to describe the solar disc, by taking

$$r^2 = x_i^2 + y_j^2. \tag{8.14}$$

Then, using equation 8.9, the limb darkening of a pixel with the coordinates (x_i, y_j) can be found. The intensity of a particular pixel, $I_{i,j}^l$, with coordinates (x_i, y_j) is given by

$$I_{i,j}^l = I(0) \left(0.57 + 0.43 \cos \psi\right), \tag{8.15}$$

where $\cos \psi$ is given by equation 8.13. When combining the effects of rotation and limb darkening we have taken $I(0) = I_{i,j}^a$, where $I_{i,j}^a$ is the intensity observed from the pixel when only the effect of the solar rotation is considered and is given by equation 8.7. The total intensity can be found by summing $I_{i,j}^l$ over all i and j.

We have, therefore, found an expression that takes into account both the solar rotation and the limb darkening. We go on to use equation 8.15 to produce a contour map of the solar disc that describes how a combination of the solar rotation and the limb darkening affects the weighting of the observations across the solar disc. However, first we consider the effect of limb darkening alone.

Figure 8.10: The weighting of observations of the solar disc when only limb darkening is considered.

8.4.1 *Weighting of the solar disc due to limb darkening*

Figure 8.10 shows the weighting of the solar disc when the effect of limb darkening is considered on its own. The relative weighting caused by limb darkening is independent of the line-of-sight velocity between the Sun and the observer, v_{los}. Furthermore, when considered on its own, the limb darkening affects blue and red wing observations in the same manner. Once again the weighting has been scaled so pixels with the largest observed intensity have a weighting of 100 and the pixels with the smallest observed intensity have a weighting of zero. All of the other weightings are distributed accordingly between these values. The weighting pattern consists of concentric circles with the observations being most heavily weighted towards the central regions. This is expected because, as seen in Figure 8.9, contours of constant ψ are also concentric circles.

8.4.2 *Weighting of the solar disc due to the Sun's rotation and limb darkening*

We now combine the weighting due to limb darkening with the effect of the solar rotation. Figure 8.11 shows the weighting of the solar disc when both the rotation of the solar surface and limb darkening are considered. In Figure 8.11 the line-of-sight velocity between the Sun and the observer, v_{los}, is zero. As expected, given the weighting observed when the rotation was considered alone, the contours are no longer centralized on the origin but are shifted towards the limbs. The red wing is weighted more strongly towards the

Figure 8.11: The weighting of the solar disc once the effect of limb darkening and the weighting due to the solar rotation are combined. The left panel shows the results for the instrumental blue wing, while the right panel shows the results for the instrumental red wing. The line of sight velocity between the Sun and the observer, v_{los}, is $0\,\text{m}\,\text{s}^{-1}$.

advancing limb of the solar disc and the blue wing is weighted towards the receding limb.

Figure 8.12 shows the weighting across the solar disc when $v_{\text{los}} = 1100\,\text{m}\,\text{s}^{-1}$. The introduction of a line-of-sight velocity between the Sun and the observer has shifted the contours towards the approaching limb of the solar disc. The effect of introducing a line-of-sight velocity is larger in the blue wing observations. This is because the red wing observations are already shifted towards the approaching limb of the solar disc.

The heliocentric angle, ψ, is defined in Figure 8.8. When $v_{\text{los}} = 0\,\text{m}\,\text{s}^{-1}$ the heliocentric angle of the pixel with the maximum weighting, ψ_{max}, is $41.0°$ in both the red and the blue wing observations. However, when $v_{\text{los}} = 1100\,\text{m}\,\text{s}^{-1}$ the heliocentric angle, ψ_{max}, of the pixel with the maximum weighting is $49.3°$ for the red wing observations. For the blue wing observations, the maximum weighted pixel has a heliocentric angle of $\psi_{\text{max}} = 32.0°$. This work can be compared with that of Underhill (1993), who performed a similar analysis. Underhill also found that, when $v_{\text{los}} = 0\,\text{m}\,\text{s}^{-1}$ the red and blue contour maps were mirror images of each other, with $\psi_{\text{max}} = 46.9°$. When $v_{\text{los}} = 1000\,\text{m}\,\text{s}^{-1}$ Underhill found that the heliocentric angle of the most heavily weighted bin, ψ_{max}, was $38.3°$ for the blue wing observations and $52.8°$ for the red wing observations. The heliocentric angle of the most heavily weighted pixels, ψ_{max}, determined here are in reasonable agreement with the values found by Underhill. Small differences may occur because some of the inputs in

Figure 8.12: The weighting of the solar disc once limb darkening has been included in addition to the weighting due to the solar rotation. The left panel shows the results for the instrumental blue wing, while the right panel shows the results for the instrumental red wing. This time v_{los} was $1100\,\text{m s}^{-1}$, which is approximately the maximum line-of-sight velocity between the Sun and the Earth in any one year.

the calculations may vary, such as the width and depth of the unshifted solar Fraunhofer line. However, it is encouraging that both Underhill and the calculations performed here find that as v_{los} increases ψ_{max} increases in the red wing observations and decreases in the blue wing observations.

In this section we have shown that the introduction of the limb darkening has turned the lines of constant weighting into ellipsoids. The exact positioning of the contours depends on v_{los}. However, the position of the detector with respect to the image of the Sun observed by an instrument also affects the bias in the observations. Therefore, we now model the effect of the position of a detector on the weighting of the observations across the solar disc.

8.5   Weighting of the solar image due to the position of the detector

As mentioned in Chapter 7 the majority of BiSON instruments each have two detectors, which are known as the port and starboard detectors. The port and starboard detectors are on opposite sides of the vapour chamber (see Figure 8.13). We consider a vertical slice through the vapour cell, which we take to be a square of side 15 mm. The intensity of light detected from a particular region on the solar disc by the port detector depends
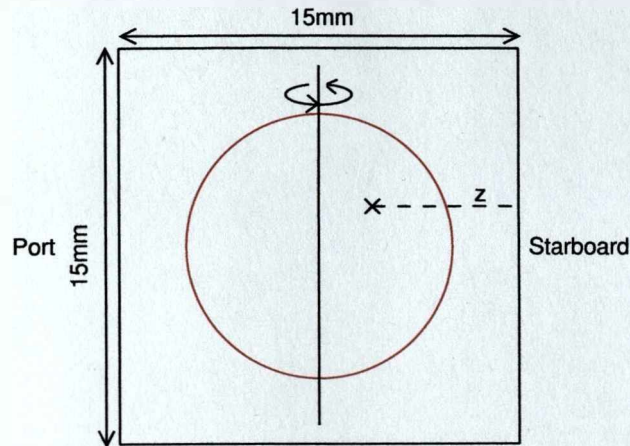
Figure 8.13: A schematic of a vapour cell. The vapour cell is 15 mm × 15 mm. At the centre of the vapour cell is the observed image of the Sun (red circle). The port and starboard detectors are positioned on either side of the vapour cell. Here we take the optical path of the light to the detector to be horizontal. In this figure the optical path is shown by the dashed line and z is the distance travelled by light through the vapour. Also shown are the axis and direction of rotation.

on the distance between that region on the image of the Sun in a vapour cell and the inside wall of the vapour cell on the side of the port detector (z in Figure 8.13). Here we have only considered the perpendicular distance between the region on the image and the wall of the vapour cell. This is not strictly true to a real instrument but it does provide a first order approximation.

The intensity observed by the detector from a given region on the Sun, I, is determined by the optical depth of the vapour in the cell, τ, and is given by

$$I = I^0 e^{-\tau}, \tag{8.16}$$

where I^0 is the intensity of light received from the Sun. The differential optical depth can be defined as

$$\mathrm{d}\tau = K\mathrm{d}z, \tag{8.17}$$

where $\mathrm{d}z$ is the optical path and K is the extinction coefficient. We have assumed that K is constant throughout the vapour cell and have defined z as the distance between a

pixel and the edge of the vapour cell along the optical path (see Figure 8.13). Therefore

$$\tau = \int_0^z K \mathrm{d}z = Kz. \tag{8.18}$$

We can then substitute this expression for the optical depth into equation 8.16 to give

$$I_{i,j}^s = I_{i,j}^0 \mathrm{e}^{-Kz}, \tag{8.19}$$

where $I_{i,j}^s$ is the intensity observed from a pixel with coordinates (x_i, y_j) and $I_{i,j}^0$ is the intensity of light absorbed by the vapour in the cell from that pixel. We originally assumed that the vapour has an optical depth of unity at the centre of the cell and so $K = 133.3\,\mathrm{m}^{-1}$. Initially we gave the image of the solar disc in the vapour cell a radius of 5 mm and we took the centre of the image of the solar disc to coincide with the centre of the vapour cell. We also took the image of the solar disc to be pixelated in the manner described in Section 8.3 and the position of each pixel on the image of the solar disc is determined by the positioning of the pixels in the real sized solar disc.

We now consider the effect on the weighting of unresolved solar observations of the detector position alone. We will then combine this effect with the previously considered results for limb darkening and solar rotation.

8.5.1 Weighting of the solar image due to the position of the detector

Figure 8.14 shows the results of considering just the effect of observing on different sides of the vapour cell. As before the most heavily weighted region is given a weighting of 100 and the region from which the smallest intensity of light is observed has been given a weighting of 0. As expected the observations are weighted most heavily towards the side of the solar disc that is closest to the detector. Thus, the starboard observations are weighted towards the receding limb of the Sun, while the port observations are weighted towards the approaching solar limb. It should be noted that when considered on its own the effect of the detector's position is the same for both the blue and red wing

(a) Port detector   (b) Starboard detector
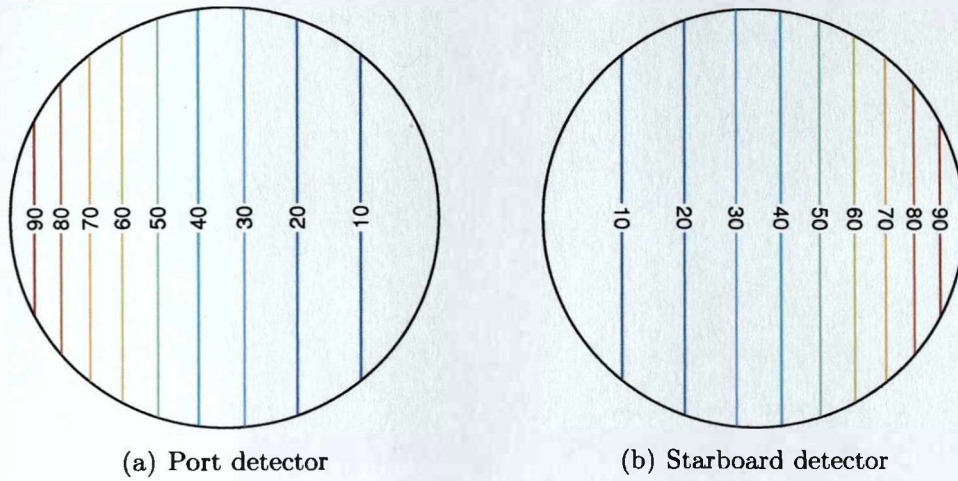
Figure 8.14: The weighting across the solar disc when the positioning of the detector is considered on its own.

instrumental observations. The results shown in Figure 8.14 are also independent of v_{los}. As we have only considered the horizontal distance between a pixel on the solar image and the edge of the vapour cell the contours of constant weighting are vertical lines, and so the weighting is independent of a pixel's y coordinate.

We now move on to amalgamate the effects of the Sun's rotation, limb darkening and the position of the detector. We begin by discussing how these three effects were combined.

8.5.2 Weighting of the solar disc due to the combined effect of the solar rotation, limb darkening and the position of the detector

The weighting due to all three effects was found by combining equations 8.7, 8.15 and 8.19 in the following manner. $I_{i,j}^{a}$ was found using equation 8.7 and then substituted for $I(0) = I_{i,j}^{a}$ in equation 8.15. Equation 8.15 was used to find $I_{i,j}^{l}$ and then $I_{i,j}^{l}$ was substituted for $I_{i,j}^{0}$ in equation 8.19. In other words

$$I_{i,j}^{a} = I_{i,j}^{r} - I_{i,j}^{t}, \tag{8.20}$$

$$I_{i,j}^{l} = I_{i,j}^{s}(0.57 + 0.43 \cos \psi), \tag{8.21}$$

$$I_{i,j}^{s} = I_{i,j}^{l} e^{-Kz}. \tag{8.22}$$

(a) Blue wing, port detector
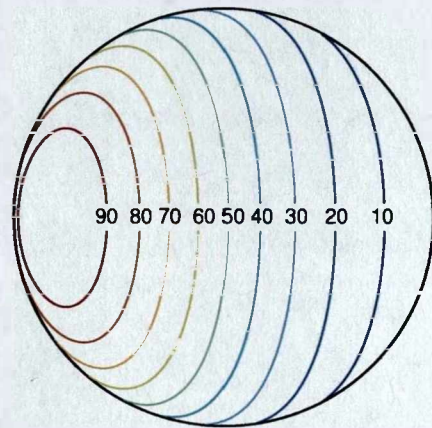
(b) Red wing, port detector

(c) Blue wing, starboard detector

(d) Red wing, starboard detector

Figure 8.15: Contour maps showing the weighting of the solar disc that is observed when the position of the detector is considered as well as the limb darkening and the effect of the solar rotation. The radius of the image in the vapour cell is 5 mm and $v_{los} = 0$.

In the following results the line-of-sight velocity between the Sun and the detectors was $v_{los} = 0\,\mathrm{m\,s}^{-1}$ and the image was positioned at the centre of the vapour cell. Figure 8.15 shows that, even when the other effects are introduced, the observed weighting still depends on the positioning of the detector with respect to the solar image. This means that the unresolved port and starboard data are not observing the same regions of the Sun. Therefore the solar noise observed by the two instruments will be slightly different. As $v_{los} = 0\,\mathrm{m\,s}^{-1}$ the red, starboard map is the mirror image of the blue, port map and the red, port map is the mirror image of the blue, starboard map. Before the instrumental effect was considered the blue wing observations were weighted towards the receding limb of the Sun (the western limb in Figure 8.15). When the Sun is observed by the port detector the results are shifted towards the approaching solar limb. The balance between these two effects results in the observations being weighted most heavily towards the centre of the disc. This reasoning can be reversed to explain why the red, starboard observations are also weighted towards the centre of the solar disc. Closer inspection reveals that the blue port observations are weighted very slightly more heavily towards the approaching limb of the Sun, while the red starboard observations are weighted slightly more heavily towards the receding solar limb. Therefore, the results observed in Figure 8.15 imply that the dominant effect is that of the position of the detector.

As mentioned previously the line-of-sight velocity between the Sun and an observer, v_{los}, varies throughout the year. Figure 8.16 shows contour maps of the weighting of the solar image when $v_{los} = 1100\,\mathrm{m\,s}^{-1}$. This line-of-sight velocity shifts the observations so that they are more strongly weighted towards the approaching limb of the solar disc than when $v_{los} = 0\,\mathrm{m\,s}^{-1}$. Notice the elongated contours that are observed in the red wing, starboard detector map. These occur because the effect of the line-of-sight velocity is to shift the observations towards the approaching solar limb. However, this effect is counterbalanced by the effect of the position of the starboard detector, which shifts the observations towards the receding solar limb.

We have now built up an image of the weighting of the Sun that takes account of the

(a) Blue wing, port detector

(b) Red wing, port detector

(c) Blue wing, starboard detector

(d) Red wing, starboard detector
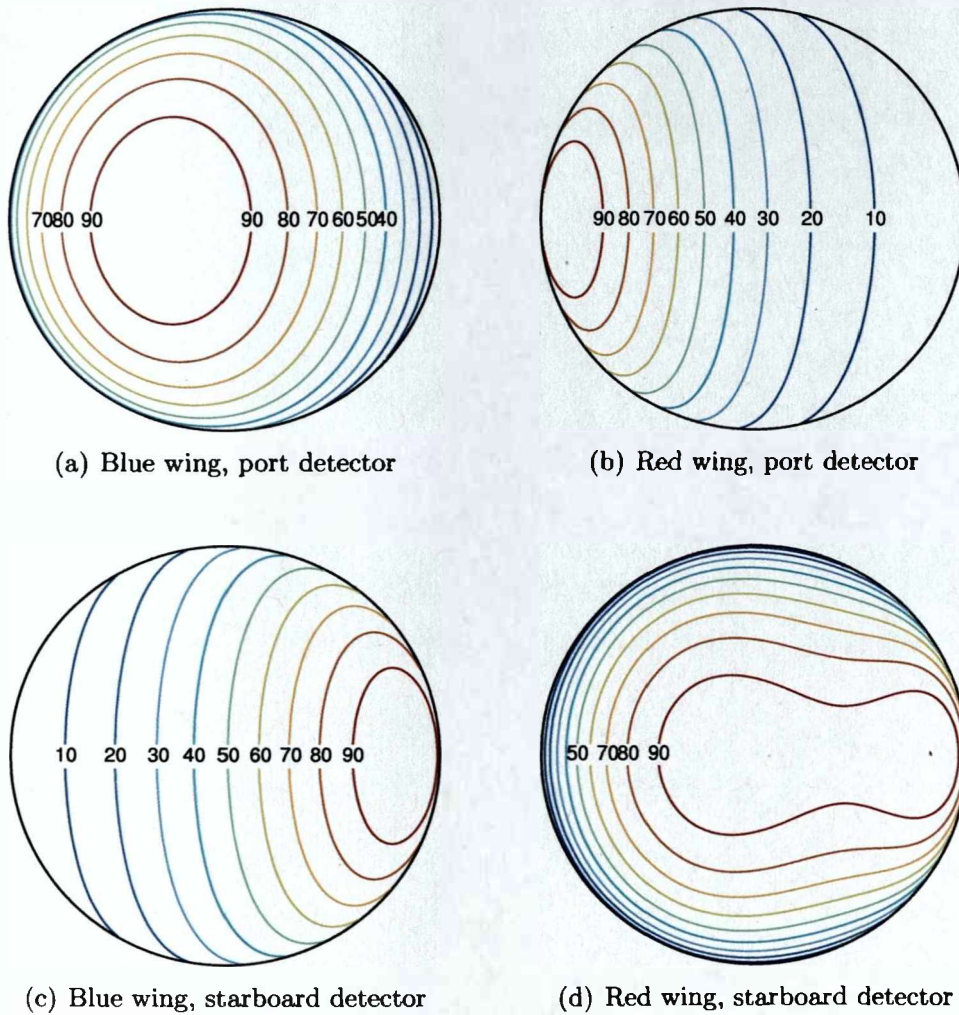
Figure 8.16: Contour maps showing the weighting of the solar disc that is observed when the position of the detector is considered as well as the limb darkening and the effect of the solar rotation. The radius of the image in the vapour cell is $5\,\text{mm}$ and $v_{\text{los}} = 1100\,\text{m}\,\text{s}^{-1}$.

solar rotation, limb darkening and the positioning of the detector. We have also shown that the exact structure of this weighting varies with the line-of-sight velocity between the Sun and the instrument. As the Earth orbits the Sun and spins on its own axis the line-of-sight velocity between the Sun and an Earth-based instrument will vary. Hence we move on to investigate in more detail how variations in v_{los} affect the observed weighting.

8.6 *Variation in the weighting with line of sight velocity*

As we have already seen, the line-of-sight velocity of the solar surface varies with distance from the rotation axis. Each pixel not only corresponds to an x and y coordinate but it also corresponds to a line-of-sight Doppler velocity, which is given by equation 8.3. As has been shown throughout this chapter the observations are weighted more heavily towards some areas of the solar disc than others and so the observed data are weighted towards some velocity in the line of sight, which occurs because of the solar rotation. Therefore, the weighting of the solar disc means that a Doppler velocity shift could be introduced into the Sun-as-a-star data. The weighted mean of the velocity observed over the solar disc then gives an indication of the mean line-of-sight velocity observed by the instrument. The weighting of a particular pixel is given by the ratio of the intensity of light from the pixel observed by the instrument and the total intensity observed from the whole Sun, which can be calculated by summing the observed intensity over all pixels. In other words

$$I_{\text{tot}} = \sum_{i=0}^{500} \sum_{j=0}^{500} I_{i,j}^s, \qquad (8.23)$$

where $I_{i,j}^s$ is given by equation 8.22. The weighted mean velocity was determined separately for the blue and red wing observations and for the port and starboard detectors. The results are shown in Figure 8.17. The calculated weighted mean velocity varies with v_{los}. The determined weighted mean velocity also depends on which detector is being considered and which instrumental profile is used. Table 8.1 shows the weighted mean velocity calculated for different combinations of the instrumental wing and detec-

(a) Port detector

(b) Starboard detector
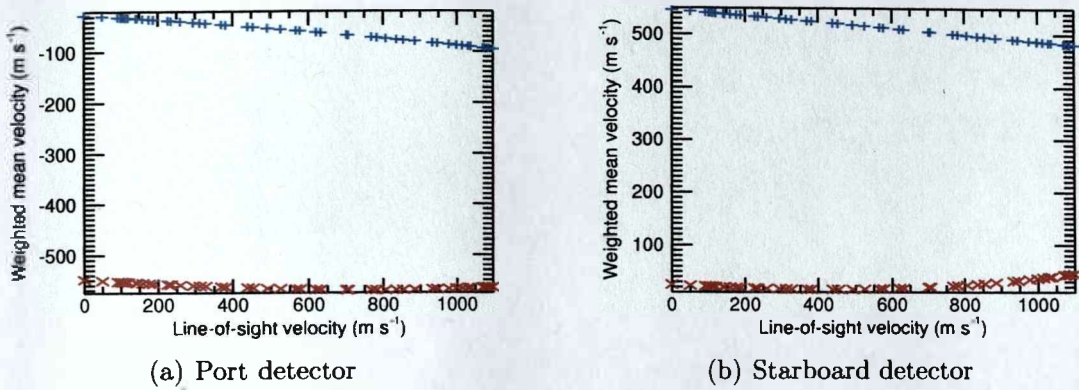
Figure 8.17: The weighted mean velocity of unresolved observations varies with the line-of-sight velocity between the Sun and an observer. In both panels the red crosses represent the red wing observations, while the blue plus signs represent the blue wing observations. The size of the image in the vapour cell was 5 mm.

Table 8.1: The weighted mean velocity observed at different values of v_{los}.

	Weighted mean velocity when $v_{los} = 0\,\mathrm{m\,s^{-1}}$ ($\mathrm{m\,s^{-1}}$)	Weighted mean velocity when $v_{los} = 1100\,\mathrm{m\,s^{-1}}$ ($\mathrm{m\,s^{-1}}$)
Blue wing, port detector	-26.8	-94.5
Red wing, port detector	-545.9	-565.2
Blue wing, starboard detector	+545.9	+481.6
Red wing, starboard detector	+26.8	+41.5

tor. The weighted mean velocity has been calculated when $v_{los} = 0\,\mathrm{m\,s^{-1}}$ and when $v_{los} = 1100\,\mathrm{m\,s^{-1}}$. The results in Table 8.1 indicate that when $v_{los} \neq 0\,\mathrm{m\,s^{-1}}$ the observations, apart from the red wing, starboard measurements, are shifted towards the approaching limb of the Sun. The increase in the weighted mean velocity with v_{los} observed in the red wing, starboard results probably occurs because the contours become elongated, as seen in Figure 8.16. The change in the calculated weighted mean velocity is larger for the blue wing observations as before the position of the detector was considered the red wing observations were weighted towards the approaching limb of the Sun.

Real BiSON observations combine the results for the blue and red wing observations in the manner described in Section 1.4.2. Therefore, the total observed weighted mean velocity is a combination of the weighted mean velocities observed by the blue and red wings. Hence, we have calculated the mean of the blue and red wing weighted mean

Figure 8.18: Mean weighted velocity as a function of line-of-sight velocity. The black diamonds represent the results when only the effect of rotation and limb darkening are considered. The green crosses represent the starboard observations and the pink plus signs represent the port observations. The size of the image in the vapour cell was 5 mm.

velocities for both the port and the starboard detectors. Figure 8.18 shows the mean weighted velocity when only the effect of the rotation and the limb darkening are included (black diamonds). As expected, when the line of sight velocity is zero the blue and red wing observations are symmetric and so the weighted mean velocity is $0\,\mathrm{m\,s^{-1}}$. However, as the line-of-sight velocity increases the weighted mean velocity becomes increasingly negative. Also plotted in Figure 8.18 are the weighted mean velocities calculated when the position of the detector is included in the calculations. In the calculations the radius of the image of the Sun in the vapour cell was 5 mm. The results show that the weighted mean velocity, which will be introduced into the observed data, depends on which detector is considered. The difference between the port and starboard weighted mean velocity is approximately constant at all values of v_{los}. We will compare these results with a velocity offset that is observed in the BiSON data.

8.6.1   Velocity offset in BiSON observations

In Chapter 1 we mentioned that the Doppler velocities observed by BiSON instruments are calculated using the ratio

$$v_{\mathrm{obs}} = k\frac{I_b - I_r}{I_b + I_r} = kR,$$

(8.24)

where I_b is the total intensity observed by the instrumental blue wing and I_r is the total intensity observed by the red wing. The value of k varies from day to day and is approximately $3000\,\mathrm{m\,s^{-1}}$. Determining the value of k is not a simple matter as the relationship between the intensity ratio, R, and the observed velocity, v_{obs}, is not linear because of the variation in the gradient of the solar Fraunhofer line.

The ratio, R, can be expanded in the form of a polynomial:

$$R(v_{\mathrm{obs}}) = \sum_{i=0}^{N} \left[a_i \left(v_{\mathrm{obs}} - v_{\mathrm{grs}} \right)^i \right], \qquad (8.25)$$

where v_{grs} is the line-of-sight velocity due to the gravitational redshift (see Section 1.5), N is the order of the polynomial and a_i are the coefficients of the polynomial. The order of the polynomial must be chosen carefully, as if there are not enough terms the resolution will be too poor to follow the change in sensitivity throughout the day. However, if too many terms are included in the polynomial the fit becomes unstable and the signal from the oscillations can be lost when the series is subtracted from the measured values. Currently, a cubic ($N = 3$) polynomial is used, as it is found to remove many unwanted low-frequency effects. It should be remembered that v_{obs} includes the orbital velocity, the spin velocity and the gravitational redshift velocity in addition to any oscillation velocities.

Once both v_{orb} and v_{spin} have been determined (see Section 1.5) a best fit algorithm is used to determine the coefficients, a_i, and v_{grs} on a day-by-day basis. The power series is then subtracted from the measured ratio, R_{meas}:

$$R_{\mathrm{osc}} = R_{\mathrm{meas}} - \sum_{i=0}^{N} \left(a_i \left(v_{\mathrm{obs}} - v_{\mathrm{grs}} \right)^i \right), \qquad (8.26)$$

where R_{osc} is the ratio given by the oscillations and any noise. The ratio, R_{osc}, is then divided by the the sensitivity, which is given by $(\delta R_{\mathrm{osc}}/\delta v)$.

Here we are mainly interested in equation 8.25. Usually the ratio is calculated separately for the port and starboard observations and then averaged. Here, the port and the

starboard observations are kept separate. The daily cubic polynomial, given by equation 8.25, pairs the observed daily ratio points with their instantaneous line-of-sight velocity. As equation 8.25 takes into account the varying values of v_{orb} and v_{spin}, each value of $R(v_{obs})$ should correspond to the same velocity from day to day. The change in the line-of-sight velocity between the Sun and the instrument because of the spin and orbital velocities should only alter the observed value of the ratio. We have plotted the value determined for $R(v_{obs})$, given by equation 8.25, as a function of v_{obs} using the calculated best fit coefficients. We have then determined the velocity intercept that corresponds to a value of $R(v_{obs}) = 0.15$. We have called this velocity the velocity offset. The value of the velocity offset was found when $R(v_{obs}) = 0.15$ rather than when $R(v_{obs}) = 0$ as the value of the ratio is often greater than 0. Therefore, to determine the velocity that corresponds to a ratio value of $R(v_{obs}) = 0$ we would have to extrapolate. This would introduce errors into the determined value of the offset velocity. On the other hand a ratio of $R(v_{obs}) = 0.15$ is observed on the majority of days and so no extrapolation is required.

Figure 8.19 shows the velocity offset observed by two BiSON instruments as a function of the known line-of-sight velocity between the Sun and the instrument. The analysis was performed for different instruments (and although not shown here, different years) to ensure the observed effects were not instrument (or year) specific. The data shown in Figure 8.19 were observed in 2006, but only days where at least 500 ratios were observed were used in the analysis. The results are very similar for both sites. As can be seen the value of the velocity offset not only varies with v_{los}, but it is also dependent on which detector is considered.

We can compare the observational results shown in Figure 8.19 with those calculated in Figure 8.18. Some similarities can be observed, such as the fact that the starboard observations are always positive and the port observations are always negative. Excluding outliers the range in the observed velocity offset is of the order of $50\,\mathrm{m\,s^{-1}}$ for both the port and the starboard observations, and for both instruments. The range in the

Figure 8.19: Left panel: The velocity offset observed in Las Campanas A data in 2006 as a function of line-of-sight velocity. Right panel: The velocity offset observed in Narrabri data in 2006 as a function of line-of-sight velocity. The red diamonds represent the results from the port detector, while the green crosses show the velocity offset observed by the starboard detector.

calculated weighted mean velocity is also of the order of $50 \, \text{m s}^{-1}$ for both the port and starboard observations. Additionally, both the observed and the calculated port velocities decrease with increasing line-of-sight velocity. However, there are also some clear differences between both panels of Figure 8.19 and Figure 8.18. For example, the gradient of the observed starboard velocity offset changes from positive to negative at a line-of-sight velocity of approximately $550 \, \text{m s}^{-1}$, whereas the gradient of the calculated velocity is always negative. Furthermore, the calculated velocities are approximately symmetrically split about a weighted mean velocity of $0 \, \text{m s}^{-1}$, whereas the magnitude of the observed port velocity offset is greater than the magnitude of the observed starboard velocity offset at all v_{los}. However, this is because the absolute value of the observed velocity offset is dependent on the value of the ratio, R, that is used to calculate the offset. Therefore, using a different value of the ratio, R, will shift the results to different velocities. More importantly, the difference between the observed port and starboard offset velocity ($\sim 250 \, \text{m s}^{-1}$) is significantly less than the calculated difference between the port and starboard weighted mean velocity ($\sim 600 \, \text{m s}^{-1}$).

It is, therefore, interesting to see whether we can replicate more accurately the main features of the observed velocity offset by altering various input parameters in our calculations. We therefore move on to investigate the sensitivity of the calculated weighted mean velocity to various input parameters, with the aim of improving the agreement

Figure 8.20: Sensitivity of the weighted mean velocity to the FWHM of the solar absorption line. The different coloured symbols represent the results calculated using different values of the FWHM of the solar absorption line.

between the observations and the calculations.

8.6.2 Sensitivity of the calculations to the width of the solar absorption line

One of the main input parameters in the calculations that determine the effect of the solar rotation on the observed weighting is the width of the solar absorption line. In Section 8.2 the width of the solar absorption line was determined to be 0.018 nm, using resolved observations of the centre of the solar disc. This value was then taken to be the width of the absorption line observed from all pixels on the Sun. To investigate the sensitivity of the weighted mean velocity to this width calculations were made where the width of the solar absorption line observed from each pixel was doubled. In other words the FWHM of the solar absorption line observed from each pixel on the solar disc was 0.036 nm. We also performed the calculations using a FWHM of 0.009 nm. The weighted mean velocity was calculated by only considering the effect of the rotation and limb darkening as the sensitivity should be independent of the position of the detector. The results are shown in Figure 8.20. The gradient of the slope is very sensitive to the assumed FWHM of the solar absorption line. It is therefore possible that the FWHM can be scaled slightly so that the range in the calculated weighted mean velocities matches the range in the observed velocity offset.

Figure 8.21: Sensitivity of the observed weighting to the solar rotation rate. The black diamonds show the results when the rotation rate was 25 days and the blue triangles represent the results when the rotation rate was 36 days.

8.6.3 Sensitivity of the calculations to the solar rotation rate

When calculating the effect of the solar rotation on the observed weighting we assumed that the Sun rotated as a solid body with a period of 25 days. However, the Sun actually exhibits differential rotation, which means that the rotation rate is latitude dependent. The material at the solar equator does rotate with a period of approximately 25 days, however, the material at the poles rotates with a period of approximately 36 days. To investigate how sensitive the calculations are to the assumed rotation rate we calculated the weighted mean velocity when the rotation rate was 36 days. The weighted mean velocity was calculated considering the effect of the solar rotation and limb darkening only. The results are shown in Figure 8.21. Although the sensitivity to the rotation rate increases as v_{los} increases changing the rotation rate alters the weighted mean velocity by, at most, $3.7\,\mathrm{m\,s^{-1}}$. This implies that including differential rotation in the calculations would make little difference to the observed results.

8.6.4 Sensitivity of the calculations to the width and separation of the instrumental absorption lines

When calculating the effect of the rotation on the observed weighting we had to include various parameters that described the instrumental absorption lines. We took the width

Figure 8.22: Sensitivity of the calculated weighted mean velocity to the width and separation of the instrumental absorption lines. The black diamonds show the results when the widths of the instrumental absorption lines were $700\,\mathrm{m\,s^{-1}}$ and the separation of the components was $5200\,\mathrm{m\,s^{-1}}$. The blue asterisks represent the results when the separation of the components was $5200\,\mathrm{m\,s^{-1}}$ the the FWHM of the instrumental absorption lines were $1400\,\mathrm{m\,s^{-1}}$. The green squares represent the results when the widths of the instrumental absorption lines were $700\,\mathrm{m\,s^{-1}}$ and the separation of the components was $2600\,\mathrm{m\,s^{-1}}$. The red crosses represent the results when the widths of the instrumental absorption lines were $700\,\mathrm{m\,s^{-1}}$ and the separation of the components was $7800\,\mathrm{m\,s^{-1}}$.

of these lines to be $700\,\mathrm{m\,s^{-1}}$ and we took the separation between the blue and red instrumental wings to be $5200\,\mathrm{m\,s^{-1}}$. However, these value have not been measured accurately and so it is important to determine how sensitive the calculations are to variations in the input values. Therefore, we calculated the weighted mean velocity, due to solar rotation and limb darkening, when the widths of the instrumental absorption lines were $1400\,\mathrm{m\,s^{-1}}$. Figure 8.22 indicates that doubling the width of the instrumental absorption lines has little effect on the calculated weighted mean velocity. Also plotted on Figure 8.22 are the results obtained by altering the separation of the instrumental components. The weighted mean velocity was calculated when the FWHM of the instrumental line was $700\,\mathrm{m\,s^{-1}}$ but the separation of the blue and red wing components was both $2600\,\mathrm{m\,s^{-1}}$ and $7800\,\mathrm{m\,s^{-1}}$. Altering the separation of the instrumental components affects the gradient of the slope. Therefore, this parameter could also be altered to improve the agreement between the range in the calculated weighted mean velocity and the observed velocity offset.

8.6.5   Varying of the size of the solar image

The observed radius of the Sun varies throughout the year. More precisely the angular radius of the Sun varies from $\sim 16'16"$ in early January to $\sim 15'44"$ in early July. This variation constitutes approximately 3% of the maximum observed radius. To investigate whether this variation would have a noticeable effect on the calculated weighted mean velocity throughout the year we varied the size of the image of the Sun observed at the vapour cell. The radius of the image was altered by a maximum of 3% and was varied in a sinusoidal manner so that the image had the largest radius at the start of January and the smallest radius at the start of July. However, changing the size of the image by this small amount only altered the weighted mean velocity by a maximum of $\sim 4\,\mathrm{m\,s^{-1}}$. Therefore the varying size of the solar image by this small amount does not explain the differences observed between the calculated weighted velocity and the velocity offset found in the BiSON data.

As varying the size of the image did alter the calculated weighted mean velocity slightly, we investigated the effect of altering the size of the image by a larger amount. Therefore, we calculated the weighted mean velocity observed at different values of v_{los} when the size of the image was 4 mm and 7.5 mm. Note that 7.5 mm is the largest possible radius that the image can have, assuming that the vapour cell observes the whole disc. Figure 8.23 shows that reducing the size of the observed image decreases the difference between the port and starboard weighted mean velocity. Therefore, the observed velocity offset is best replicated when the size of the image is less than 4 mm.

In all of the above calculations we have taken the optical depth to be unity at the centre of the vapour cell. We move on to investigate the effect of changing the optical depth on the weighting of the observations.

8.6.6   Varying the optical depth of the vapour

An optical depth of unity at the centre of the vapour cell is a reasonable but arbitrary value and is based on no direct observations. Therefore, we investigated the effect of

Figure 8.23: A comparison between the weighted mean velocity determined for different image radii. The weighted mean velocity is plotted as a function of line-of-sight velocity. The different colours show the results for the different detectors, while the different symbols represent the results for various image sizes.

varying the optical depth at the centre of the vapour cell, τ. This was done by altering the value of the extinction coefficient. Calculations were made when the optical depth at the centre of the cell, τ, was 0.8 and 1.2. The results are plotted in Figure 8.24. The radius of the image was again taken to be 5 mm. Altering the optical depth of the vapour has a similar effect on the calculated weighted mean velocity as altering the size of the image. In fact, changing the optical depth by a value of 0.2 is equivalent to altering the size of the image by 1 mm. The difference between the port and starboard velocities is reduced at low optical depths, and so the agreement between the calculated and the observed velocities is best when the optical depth at the centre of the vapour cell is less than unity.

8.6.7 *Closest agreement between the calculated and observed velocities*

The overall separation of the port and starboard observed velocity offsets is best replicated when the size of the image is small and the value of the optical depth at the centre of the cell is reduced. The actual size of the image and optical depth in the vapour cell are not yet known and thus need to be determined before accurate calculations can be made. However, one estimate of the size of the image, obtained simply by observing the image

Figure 8.24: A comparison between the weighted mean velocity calculated when the optical depth of the vapour is altered. The weighted mean velocity is plotted as a function of line-of-sight velocity. The different colours show the results for the different detectors, while the different symbols represent the results for different optical depths. The radius of the image was always 5 mm.

formed in a BiSON instrument, places the radius in the range of 4 – 5 mm. We have calculated the weighted mean velocity when the size of the image is 4 mm and we have altered the optical depth of the vapour in the cell until the separation in the calculated velocities is in agreement with the observed separation. We found that a separation between the port and the starboard calculated velocities of approximately $250 \, \text{m s}^{-1}$ is obtained when the optical depth at the centre of the vapour cell, τ, is 0.65. A comparison between the observed and calculated results can be seen in Figure 8.25. An optical depth of $\tau = 0.65$ is significantly less than the value of unity that was originally assumed and so efforts should be made to measure this value directly. In Figure 8.25 the shape of the observed starboard velocity offset is not well replicated by the calculations. Therefore, efforts should continue to improve accuracy of the calculations.

With this in mind we considered the effect of altering the orientation of the Sun's rotation axis. The Sun's rotation axis is tilted by an angle of $\sim 7.25°$ from vertical and so the observed orientation of the rotation axis varies as the Earth orbits the Sun. However, the calculations show that varying the orientation of the solar rotation axis makes very little difference to the calculated weighted mean velocity (less than $4.5 \, \text{m s}^{-1}$).

Figure 8.25: A direct comparison between the observed and calculated velocity offset. The calculations have been made for an image with a radius of 4 mm and the optical depth at the centre of the cell, τ, was 0.65.

The observed ratio is determined using observations of the total solar Fraunhofer line. Therefore, it is interesting to examine the total solar absorption line produced when the effects of the solar rotation, limb darkening and the position of the detector are combined.

8.7 Total observed line

The total observed line from the whole solar disc, which will be measured by Sun-as-a-star instruments, is given by the sum of the lines observed by all of the complete pixels on the Sun. In other words the total intensity, $L_{tot}(\lambda)$, which is a function of wavelength, is given by

$$L_{tot}(\lambda) = \sum_{i=0}^{500} \sum_{j=0}^{500} L_{i,j}(\lambda), \qquad (8.27)$$

where $L_{i,j}(\lambda)$ is the line profile observed from the pixel centred on (x_i, y_j). $L_{tot}(\lambda)$ was calculated when the effect of the rotation, limb darkening and the position of the detector were considered together. The line produced can be seen in the top panel of Figure 8.26. The resulting FWHM of this total observed line was calculated to be $\sim 7.3\,\mathrm{km\,s^{-1}}$. It should be noted that this process does not give any information about the actual observed intensity. Obviously this line was significantly larger than the individual lines, $L_{i,j}$, and so

Figure 8.26: Top panel: The line profile observed when the contributions from each pixel are summed and the line is scaled to have an 78% depth. Middle panel: The difference between the total observed line found here and a Gaussian profile. Bottom panel: The difference between the total observed line found here and the profile suggested by Underhill (1993).

the line was scaled to have a 78% depth. This value has been chosen to allow a comparison with the work of Underhill (1993), which was based on instrumental observations.

The shape of the total observed line can be compared to a Gaussian representation. The middle panel of Figure 8.26 shows the difference between the total observed line and a Gaussian profile with a 78% depth and the correct FWHM. The agreement between the total line and the Gaussian representation is good at all wavelengths. However, the total observed line does depart slightly from a Gaussian shape in the wings of the profile. As a Gaussian profile is symmetric but the difference between the two profiles is not we can

conclude that the total line profile calculated here is not symmetric. When calculating the observed velocity from the red and blue wing intensities (equation 8.25) it is assumed that the solar absorption line is symmetric. However, the departures from symmetry observed here are small and so are unlikely to affect the results greatly.

Underhill (1993) empirically determined a suitable profile for Sun-as-a-star data by comparing different model line shapes with the line profile observed by the BiSON instruments. Underhill found that the BiSON data was fitted well by a line profile with the form

$$L(\lambda) = \exp\left[\beta(\lambda)\right] \quad \text{with} \quad \beta(\lambda) = 1.514\exp\left[-\left(\frac{\lambda - \lambda_0}{92}\right)^2\right], \quad (8.28)$$

where they defined λ in mÅ and λ_0 was the line centre. The bottom panel of Figure 8.26 shows the difference between the total line profile determined here, $L_{\text{tot}}(\lambda)$, and the profile suggested by Underhill. The agreement between this profile and $L_{\text{tot}}(\lambda)$ is not as good as the agreement between the Gaussian profile and $L_{\text{tot}}(\lambda)$.

The Gaussian line profile is a better representation of the total line found in these calculations than the line profile determined by Underhill (1993). As the profile found by Underhill is based on observations of the line observed by the BiSON instruments it is possible that this indicates that these calculations do not fully replicate the real observations. Therefore, attempts should be made to make these calculations more realistic.

8.8 Summary

Unresolved Doppler velocity observations are not homogenous across the solar disc. In fact, Sun-as-a-star observations are weighted more heavily towards some areas of the solar disc than others. We began this chapter by considering the effect on unresolved observations of the solar rotation. To do this we calculated the spectrum of light that is absorbed by the potassium vapour in the instrument from pixels on the Sun. When the effect of the solar rotation was considered on its own the contours of constant weighting

were vertical lines. Counterintuitively the instrumental red wing observations were more heavily weighted towards the approaching limb of the Sun, while the instrumental blue wing observations were more heavily weighted towards the receding limb of the Sun. The positioning of the contours were dependent on v_{los}.

We then moved on to determine the effect of limb darkening on the weighting of the solar disc. When limb darkening was considered alone the contours of constant weighting were concentric circles centred on the middle of the solar disc. The contours were independent of the instrumental wing being considered. We combined the effect of the limb darkening with the effect of the Sun's rotation and found that, when $v_{los} = 0\,\mathrm{m\,s^{-1}}$, the contours of constant weighting become ellipse-like. The blue and red wing observations were mirror images, with the blue wing observations weighted most heavily towards the receding limb of the Sun and the red wing observations weighted most heavily towards the approaching limb of the Sun. When $v_{los} = 1100\,\mathrm{m\,s^{-1}}$ the contours were shifted towards the approaching limb of the solar disc.

Both the solar rotation and the limb darkening were considered by Brookes et al. (1978b). We added a third effect: the position of the detector with respect to the image of the solar disc in the vapour cell. This is important in BiSON instruments as many of them have two detectors, which are positioned on opposite sides of the vapour cell. When the position of the detector was considered on its own we found that the observations were most heavily weighted towards the side of the disc that was closest to the detector.

We then combined the effects of the solar rotation, limb darkening and the position of the detector. When $v_{los} = 0\,\mathrm{m\,s^{-1}}$ we found that the starboard, blue wing observations were a mirror image of the port, red wing observations and the contours of constant weighting were ellipse-like. The port, blue wing observations were a mirror image of the starboard, red wing observations and again the lines of constant weighting were ellipses. The results implied that the position of the detector is the most dominant effect. Again the results were affected by the line-of-sight velocity.

We investigated the effect of varying v_{los} on the weighting across the disc. The rota-

tion of the Sun and the inhomogeneous weighting indicate that observations are weighted towards a certain line-of-sight velocity. The calculated weighted mean velocity was dependent on v_{los}, which detector is considered and which instrumental wing is considered. The calculated weighted mean velocity was compared to a velocity offset observed in BiSON data. In an attempt to obtain better agreement between the calculated weighted mean velocity and the observed BiSON velocity offset we varied various input parameters.

The calculated weighted mean velocity is sensitive to the width of the solar absorption line observed from each pixel on the solar disc. Variations in the FWHM of the solar absorption line alter the gradient of the variation in the calculated weighted mean velocity with v_{los}. Comparison with the range in offset velocities observed in the BiSON data indicated that our choice of FWHM produces approximately the correct gradient. However, it is possible that a slight variation in the assumed FWHM of the solar absorption line would improve agreement further. We altered the assumed speed of solar rotation but found this had little effect on the calculated weighted mean velocity. Thus implying that including differential rotation in the calculations would only alter the results marginally. Additionally, the calculated weighted mean velocity is also relatively insensitive to variations in the width of the instrumental absorption lines. However, the calculated weighted mean velocity is sensitive to the separation of the instrumental components. This also alters the gradient of the variation in the calculated weighted mean velocity with v_{los}. Therefore, this is another parameter that could be altered to improve the agreement between the calculated and observed velocities.

When we decreased the size of the image by a relatively large amount and when the optical depth of the vapour in the cell was decreased better agreement between the calculated weighted mean velocity and the observed BiSON velocity offset was obtained. The best agreement between the observed and calculated velocities was obtained when the radius of the image was 4 mm and the optical depth at the centre of the vapour cell was 0.65.

Both the size of the image and the optical depth have not been measured accurately.

This needs to be done before rigourous calculations can be made. Some of the assumptions made in the calculations are simplistic. For example, when determining the effect of the position of the detector on the observed weighting, we have assumed that the path length of any detected photons is the horizontal distance between a pixel on the image of the Sun and the wall of the vapour cell. However, this is not necessarily true. Therefore, an extension to this work would be to make the calculations more realistic. The fact that the velocity offset varies throughout the year has an effect on the long-term stability of the BiSON observations, and thus affects the quality of the data. If the observed velocity offset can be replicated by these calculations it can be removed from the data, thereby reducing the noise in the data. As we have seen in Chapter 6, simulations imply that if the level of noise in the data is reduced we will significantly improve the chances of detecting low-frequency modes.

9. OVERVIEW

The main focus of this thesis was the development and application of techniques that aim to search for low-frequency p modes and g modes. Low-frequency modes are sensitive to the structure of the solar core and so provide stringent tests for models of the internal structure of the Sun. Crucial to the understanding and employment of these techniques is a sound knowledge of the background noise that is present in solar observations. At low frequencies the noise dominates a frequency-power spectrum making it extremely difficult to observe solar modes. In Chapter 2 we discussed the main sources of noise in BiSON data, which can be split into three main categories: solar noise, atmospheric noise and instrumental noise. As the total noise is an amalgamation of many different sources, the central limit theorem implies that the total background noise in the time domain tends towards a Gaussian distribution. This means that the total background noise in a frequency-amplitude spectrum should tend towards a normal distribution and the background noise in a frequency-power spectrum should tend towards a χ^2 2 degrees of freedom distribution. However, the noise in a frequency-power spectrum increases at low frequencies, and so a statistical description of the noise departs from a χ^2 2 degrees of freedom distribution at low frequencies. A flat frequency-power spectrum, with a χ^2 2 degrees of freedom distribution, can be approximated by calculating a running mean for the data over regions of approximately 100μHz. However, this approximation falls short at very low frequencies where the level of noise across a 100μHz range cannot be regarded as constant. Pre-whitening of a spectrum, to make the level of noise flat, is difficult as care must be taken not to alter the statistical distribution of the noise. We moved on to take advantage of the statistical distribution of the noise in both frequency-amplitude

and frequency-power spectra to search for low-frequency p modes and g modes.

In Chapter 3 we described statistical tests that can be applied to a frequency-power spectrum. These tests determine the probability that a prominent spike or structure of spikes is part of the noise background, which is the H0 hypothesis. Much of this work is based on that of Chaplin et al. (2002a), who use an approximation to determine threshold power levels. Here, we improved the accuracy of this approximation by including an extra term in the binomial expansion of P_s, the probability that a spike, larger than the threshold level, will be observed at least once in a frequency range, Δ_p. Various statistical tests were developed that not only searched for individual prominent spikes but also took advantage of known properties of the p modes to search for peaks and multiplets. The statistical tests were used to search artificial FLAG data. The statistical tests were very effective at detecting the modes that were included in the FLAG data. Initially the thresholds were set so there was less than a 10% chance that any detections were noise. However, it was found that the more stringent threshold level of a 1% chance that a detection is noise and the requirement that each mode candidate pass at least two tests significantly reduced the chances of making a false detection, whilst not dramatically reducing the number of modes that were detected. In other words, the statistical tests were able to detect the majority of the modes in the FLAG data, whilst making very few false detections.

In Chapter 4 the statistics that allow contemporaneous sets of data to be searched for prominent coincidences were considered. The probability of observing a prominent spike or structure of spikes in the same frequency bin in each of two frequency-amplitude spectra was found. Derivation of this probability was non-trivial, as proper account was taken of the amount of common noise present in the two sets of data. The origin and power of the common noise varies depending on the source of the data.

In Chapter 5 we used the statistics developed in the previous two chapters to search contemporaneous BiSON and GOLF data. The data were compared for prominent coincident features and searched individually. In BiSON and GOLF data the common noise

is solar in origin and, more specifically, is due to the solar granulation. Comparing the data allowed threshold levels to be reduced to less than $10\,\mathrm{mm\,s^{-1}\,bin^{-1/2}}$. By comparing two frequency-amplitude spectra we were able to reduce threshold levels by $\sim 28\%$ in amplitude, which is close to the maximum possible reduction of 29%. The exact magnitude of the reduction is dependent on the amount of common noise present and the optimal case is when no common noise is shared by the data. We also compared non-contemporaneous BiSON and GOLF data, which contained no common noise, and so the maximum reduction in threshold detection levels was obtained. The statistical tests detected low-frequency modes in both the contemporaneous and non-contemporaneous data, whose frequencies were in good agreement with those found in previous studies (for example Chaplin et al. 2002a). However, despite the reduction in threshold levels there were still gaps in the list of mode candidates.

We moved on to investigate whether or not we would expect to be able to detect more low-frequency modes than we detected in Chapter 5. In Chapter 6 we made predictions of how easily low-frequency modes could be detected. These predictions were based on estimates of the modes' powers and lifetimes, which were obtained using the simple function relationships that were found to describe the variation of these properties with frequency at higher frequencies. We extrapolated two relationships to low frequencies, which we called method 1 (ln-linear) and method 2 (ln-ln). A comparison was made between the extrapolated powers and widths and those predicted by the Cambridge stochastic excitation and damping codes. It is well known that the Cambridge model underestimates mode linewidths, and so it was encouraging that the widths predicted by both extrapolation methods were larger than the model widths. The underestimation of the model widths means that the model powers are overestimated. Therefore, it was again pleasing that the powers estimated by both extrapolation methods were smaller than the model powers. Monte Carlo simulations were performed to determine how often the modes could be detected.

The number of modes detected in the simulated data at low frequencies was small.

Most detections were made when the simulations used the parameters predicted by method 1 (the ln-linear relationship) and the noise was correlated to the mode excitation amplitude. Increasing the S/N of the mode by a factor of two led to significantly more detections. The S/N of a mode can be enhanced by increasing the length of the timeseries and by reducing the amount of noise in the data. Therefore, it is essential that high-quality observations of solar oscillations continue to be made, thus increasing the length of available timeseries. It is also important that efforts continue to reduce the amount of noise in the data. The simulations performed here imply that if longer, cleaner sets of data were available significantly more modes could be detected.

It is possible that this work has underestimated the amount of noise in the individual simulated timeseries that is correlated to the mode excitation. Therefore future work my include noise that is individual to the two timeseries but is still correlated to the mode excitation. Chapter 6 concentrated on determining how prominent individual mode peaks are. However, if $l \geq 1$ multiplet components are observed and so this work could be developed further by determining the visibility of multiplets.

In general the results of the simulations agreed well with those observed in Chapter 5. However, the $l = 0$, $n = 6$ mode, which is observed to be very prominent in the BiSON and GOLF data, was only detected in a small fraction of simulated spectra (≤ 0.05). One possible reason for the observed prominence of this mode in the solar data is that, because of the stochastic nature of the mode excitation, it has been randomly excited to a larger amplitude. This is encouraging as it is always possible that other modes could also experience a randomly large excitation, making them easier to detect if the appropriate epoch is considered. An extension to this work would be to examine the evolution with time of the amplitude of this and other modes. It is possible that such an investigation would enable information about the lifetime of this mode to be obtained. Furthermore, it would also be interesting to investigate the evolution in time of other modes to determine how their powers vary with time.

In Chapter 7 we concentrated on searching for low-frequency modes in BiSON data.

To ensure relatively low detection threshold levels it was necessary to find a pair of contemporaneous BiSON timeseries whose frequency-amplitude spectra contained as little common low-frequency noise as possible. The majority of BiSON instruments have two detectors, which are called port and starboard detectors. Usually the data measured by the port and starboard detectors are averaged to produce a mean timeseries for each instrument. The data from the different stations can then be coherently added to create a timeseries with a high fill. However, it is possible to keep the data from the individual instruments and the port and starboard data separate to create individual timeseries.

The first pair of BiSON timeseries that were compared in Chapter 7 kept the data from the port and starboard detectors separate. We created a timeseries that contained the port data from several different BiSON instruments. We also created a timeseries that contained the starboard data observed by different BiSON instruments. These combined-site port and starboard timeseries were found to contain a large proportion of low-frequency coherent noise. A comparison was also made between non-contemporaneous combined-site port and starboard data. When the start times were separated by 12 hrs the coherency was ∼ zero. However, when the start times were separated by 24 hrs the low-frequency coherency was greater than zero. It was determined that this low-frequency coherent noise was observed because two data points separated by 24 hrs were probably measured by the same instrument.

The positioning of the BiSON stations around the world means that at any one time as many as four stations could be observing the Sun. It is, therefore, possible to compare contemporaneous data from different BiSON sites. When data from different sites are compared the low-frequency coherency is relatively small, compared to the coherency of the port and starboard data. However, the fill of a timeseries containing data from one site only is limited and this has consequences for the quality of a frequency-amplitude spectrum. A timeseries was created using data from Las Campanas B, Narrabri, Izaña and Carnarvon B that had a fill of ∼ 56%. Additionally, a timeseries consisting of data from Las Campanas A, Mount Wilson, Sutherland and Carnarvon A was created

and had a fill of $\sim$ 63%. The low-frequency coherency of these two timeseries was relatively small and so these timeseries were compared, as both contemporaneous and non-contemporaneous data, to search for low-frequency modes. The list of observed candidates agreed well with those detected in Chapter 5. However, is it possible that more detections would be made if the timeseries that were compared had a higher fill.

The statistical tests that allow data to be compared are a powerful and versatile tool and they should continue to be used to help detect low-frequency p modes and g modes. The statistical techniques could be used to search other sets of data, such as MDI, VIRGO and GONG data. In Chapter 7 we showed that, because of their long lifetimes, the coherency of low-frequency modes is only reduced by a small amount when the start times of the timeseries are separated by $\sim$ 24 hrs. Therefore, further advantage could be taken of the large number of BiSON timeseries that are available by searching non-contemporaneous data for low-frequency p and g modes. For example, the combined-site port and starboard data were not searched here as, when contemporaneous, they contained a large amount of common noise. However, when the start times were separated by 12 hrs the coherency was zero at low frequencies, meaning that detection threshold levels would be relatively low. Furthermore, these timeseries have a reasonably high fill and so window function effects would be reduced. It is also possible that the statistical tests can be used in conjunction with other search techniques, such as collapsograms, as comparing different sets of data would allow the detection threshold levels that are relevant to such techniques to be reduced. As mentioned the tests used here were based on the H0 hypothesis. A further extension to this work would be to develop statistics based on the H1 hypothesis, which assumes there is a mode present in the data. Also the statistical tests described in this thesis are based on the observed properties of p modes. To improve the chances of detecting g modes the statistical tests could be tailored to take advantage of specific g mode properties. For example, the separation in periods between g modes with the same l and consecutive n is approximately constant. Overall the statistical techniques developed in this thesis are advantageous to the search for low-

frequency p modes and g modes because of the reduction in threshold levels that can be potentially achieved.

When the data from two different stations are compared the only source of common noise is the solar background. The amount of common noise present varies depending on the separation of the sites being compared because of Doppler imaging. In Chapter 8 we investigated the consequences of Doppler imaging and other effects on unresolved Doppler velocity observations, such as those made by BiSON. Unresolved Doppler velocity observations are more heavily weighted towards certain areas of the solar disc than others. Two reasons for this are solar rotation and limb darkening, which have been considered previously by Brookes et al. (1978b). When these two effects are combined the instrumental red wing observations are weighted most heavily towards the approaching limb of the solar disc and the instrumental blue wing observations are most heavily weighted towards the receding solar limb.

We considered the additional effect of the positioning of the detector with respect to the image of the solar disc seen by the vapour cell. The observations are weighted most heavily towards the side of the solar image that is closest to the detector. When all three effects are combined it is the position of the detector that is most dominant in determining the weighting of the solar image. As the line-of-sight velocity of the solar surface varies across the solar disc, the observations are weighted towards a certain line-of-sight velocity. This velocity was calculated by determining the weighted mean velocity over the solar disc. The weighting of the solar disc, and so the weighted mean velocity, is dependent on the line-of-sight velocity between the Sun and an observer, v_{los}. The manner in which the weighted mean velocity varies with v_{los} is dependent on which instrumental wing is considered and which detector is considered. The calculated weighted mean velocity was compared to a velocity offset that is observed in BiSON data. The best agreement between the observed velocity offset and the calculated weighted mean velocity was found when the size of the image and the optical depth were small. More precisely, the calculated weighted mean velocity is in good agreement with the observed velocity offset when the

radius of the image was 4 mm and the optical depth at the centre of the cell was 0.65.

The size of the image and the optical depth of the vapour have not been measured and this should be done to improve the accuracy of the calculations further. Work within the BiSON group is currently ongoing to determine the optical depth in the vapour cell. However, initial results imply that the optical depth is less than 1, in agreement with the results found here. The calculations are also sensitive to input parameters, such as the FWHM of the absorption line observed at the centre of the Sun and the separation of the blue and red instrumental components. Efforts should be made to measure the input parameters in the calculations carefully as an improvement in the accuracy of measurement would naturally lead to an improvement in the accuracy of the calculations. Furthermore, some of the assumptions made in these calculations could be improved. For example, we have assumed that the optical path between a pixel of the image of the Sun seen in a vapour cell and the detector is horizontal. Although any departures from this will be small making a more accurate assumption could improve the agreement between the calculations and the observations further. Nonetheless, the good agreement between the results and the calculations obtained here is encouraging and it is hoped that with further work this agreement can be improved further. A long term aim for this work would be to eventually remove this offset from the data, thereby improving the long-term stability of the results. This would improve the quality of the BiSON data and the simulations performed in Chapter 6 imply that if we can reduce the level of noise in the data we can significantly improve our chances of detecting low frequency modes.

Appendix A

Refereed Publication

The work on searching for statistically prominent, coincident features in frequency-amplitude spectra, which is described in Chapters 4 and 5 of this thesis, was accepted to be published by Monthly Notices of the Royal Astronomical Society. The paper is included here for reference.

Needles in haystacks: how to use contemporaneous data in the search for low-frequency modes of oscillation of the Sun

A. M. Broomhall,[1]* W. J. Chaplin,[1] Y. Elsworth[1] and T. Appourchaux[2]

[1]*School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT*
[2]*Institut d'Astrophysique Spatiale, CNRS–Université Paris XI UMR 8617, 91405 Orsay Cedex, France*

ABSTRACT

We show how to take advantage of contemporaneous data from two different instruments in the search for low-frequency modes of oscillation of the Sun. Contemporaneous data allow searches to be made for prominent, sharp concentrations of power which are coincident in frequency. Crucial to determining objective measures of the joint probability of the random occurrence of such features, which are potential candidates for modes, is a good understanding of the characteristics of the background noise. In this paper we show how to make proper allowance, in the calculation of the probability, for noise that is common to data from different instruments. This common noise is solar in origin, and comes from the solar granulation. Its presence makes calculation of the probability a non-trivial problem. We demonstrate application of the technique in searches for low-frequency p modes. The data we searched comprised 3071 d of contemporaneous Sun-as-a-star Doppler velocity observations made by the ground-based Birmingham Solar-Oscillations Network (BiSON), and the GOLF instrument onboard the ESA/NASA *SOHO* spacecraft.

Key words: methods: data analysis – methods: statistical – Sun: helioseismology – Sun: oscillations.

1 INTRODUCTION

A high priority for helioseismology is the detection of global acoustic and buoyancy, or gravity, modes of oscillation at very low frequencies. Observational data have revealed a rich spectrum of acoustic (p) modes. However, independently confirmed detections of core-penetrating, low-degree (low-*l*) p modes are absent at frequencies below ~970 μHz (e.g. García et al. 2001; Chaplin et al. 2002). Observation of the p modes anyway becomes difficult at frequencies below ≈1400 μHz. The frequency of the fundamental radial (*l* = 0) p mode is predicted to be ≈ 250 μHz. No indisputable detections of gravity (g) modes, which for given *l* lie at frequencies lower than the p modes, have been made to date (e.g. Appourchaux et al. 2000; Gabriel et al. 2002; Turck-Chièze et al. 2004).

Detection of the g modes remains something of a Holy Grail. The interior buoyancy waves, which form the g modes, are confined within the radiative interior and core. The high dwell time of the buoyancy waves in the core in principle provides an exquisite probe of the core conditions. But observations of the modes are made at the photospheric level, and the (evanescent) g-mode signatures are attenuated significantly by the time they reach the photosphere. Detection of the g modes is therefore a major observational challenge.

*E-mail: amb@bison.ph.bham.ac.uk

While the low-frequency p modes are not as sensitive to the core conditions as the g modes, detection of the p modes is still a high priority.

The very low-frequency p modes are expected to have extremely long lifetimes. Robust detection would therefore allow measurement of their frequencies to very high accuracy and precision. But, much more importantly, some of the low-*l* overtones show a *mixed* character (e.g. see Provost, Berthomieu & Morel 2000). They have a sensitivity to the internal structure which looks like that of a g mode in the core, giving excellent diagnostic potential, and like that of a p mode at the surface, making detection comparatively easier than for a g mode.

The low-frequency p- and g-mode signals are very weak. This places demands on two crucial parts of the detection problem: first, on the quality of the observations; and secondly, on the analysis procedures used to process the data and flag potential candidate detections. In this paper we deal with an important aspect of the second part of the problem: how to take advantage in the analysis of contemporaneous data collected by two instruments. Use of data from two instruments allows searches for *coincidences* to be made. In the context of searches for g and p modes, this means occurrences of prominent concentrations of power, which lie significantly above the local noise background, in the same frequency bin (or bins) of the frequency–amplitude spectra of the data sets. If the frequency spectrum of background noise is assumed to vary slowly in frequency,

and have no sharp features, occurrences of coincident, prominent concentrations of power will serve to provide potential candidate mode detections.

Crucial to determining objective measures of the joint probability of occurrence of such features is a good understanding of the characteristics of the background noise. In this paper we show how to make proper allowance in the calculation of the probability for noise that is common to data from two different instruments. This common noise is solar in origin, and comes from the solar granulation. Its presence makes calculation of the probability a tricky problem, and the derivations by necessity involve some non-trivial mathematics. The presentation of these derivations, and the formalism for calculating the probability, is the most important part of this paper. The derivations and formalism are presented in sufficient detail to allow the reader to use them in their own data analysis. This could be in searches for solar g or p modes; or in other astrophysical scenarios where the problem is similar, that is, where searches for narrow-band features need to make allowance for partially correlated background noise.

One could alternatively combine two data sets in an effort to form an optimal combination for analysis. It is, however, unclear what the optimal method for such a combination of helioseismic data would be. Potentially one could argue that the analysis presented in this paper could guide such a combination.

Here, we wish to keep discussion of other aspects of the solar mode detection problem to a reasonable minimum to allow us to concentrate on presentation of the new formalism. We therefore restrict practical application of the method to searches for low-frequency p modes in the frequency range from ~500 to ~1500 μHz, to illustrate use of the method. The data we searched comprised 3071 d of contemporaneous Sun-as-a-star Doppler velocity observations made by the ground-based Birmingham Solar-Oscillations Network (BiSON), and the GOLF instrument onboard the ESA/NASA *SOHO* spacecraft. Both data sets were prepared for the Phoebus collaboration. Detailed results on the lower frequency range will be the subject of follow-up, Phoebus-related, studies.

The layout of the rest of the paper is as follows. In Section 2 we give the reader the necessary background on the p-mode detection problem at hand, including the basic approach to calculation of the probabilities. Section 3 is the main part of the paper. In this section we give full derivations on how to allow for common background noise in calculation of the final probabilities. In Section 4 we discuss results on application of the method in searches for coincidences in the contemporaneous BiSON and GOLF data. Finally, we give a brief summary of the paper in Section 5.

2 THE PROBLEM: SEARCHES FOR LOW-FREQUENCY p MODES

2.1 The low-frequency spectrum

The p modes appear most prominently at frequencies of just over 3000 μHz. Here, the resonant peaks of the low-*l* modes show widths of ~1 μHz in the frequency spectrum. At successively lower frequencies the observed amplitudes and widths decrease in magnitude. The observed amplitudes decrease because the tops of the acoustic cavities move to greater depths below the photosphere. The p-mode signals are as a result more strongly attenuated by the time they reach the photosphere, making them harder to observe. The widths decrease because the modes are more lightly damped.

The decrease of observed amplitude and width makes detection of modes difficult at frequencies below about 1400 μHz. First, because

the observed signal-to-noise ratio (S/N) of mode peaks in the frequency spectrum is small; and secondly, because there will be very few bins in the frequency spectrum that contain signal due to the p modes. Detection is rendered more difficult still by the observed increase towards successively lower frequencies of the background amplitude levels. By background amplitude, we mean amplitude from sources other than the p modes. These sources include solar background noise, from active regions and the various scales of granulation, and instrumental noise. Photon shot noise contributes a flat background over the range in frequency occupied by the p modes (see García et al. 2005, for a discussion of the behaviour at extremely low frequencies). This combination of background sources gives a smoothly varying total background in the frequency spectrum. The distribution of amplitudes within any one bin is observed to follow a normal distribution (e.g. Appourchaux et al. 2000; García et al. 2001; Chaplin et al. 2002; Gabriel et al. 2002).

Isolation of candidate p modes involves searches for prominent spikes in the frequency spectrum, whose amplitudes exceed significantly the local mean background amplitude. By 'spike', we mean a prominent amplitude in a *single* bin of the frequency spectrum. A simple strategy based on searches for occurrences of individual prominent spikes is appealing in that it makes no a priori assumptions regarding the properties of the p modes (aside from the fact they are assumed to be sharp in frequency). Additional tests may nevertheless be formulated, which do make some assumptions on the expected structure (e.g. peak width giving spikes in consecutive bins; or patterns of nearby spikes given by rotational splitting). We discuss the various searches applied to the BiSON and GOLF data in Section 4.2 below. For now, let as assume searches are confined to the simplest test: searches for occurrences of individual prominent spikes.

2.2 Probability of observing a prominent spike by chance

Suppose we observe a prominent spike, at an amplitude several times the local mean background amplitude. Since the distribution of background amplitudes is known, it is a trivial matter to find the probability, p, of observing by chance a prominent spike of at least that amplitude in that bin. The probability of observing by chance at least one spike, of amplitude greater than or equal to the observed relative amplitude, over a region of N bins in the frequency spectrum is then

$$P = 1 - [1 - p]^N. \tag{1}$$

The value of P gives an objective measure of the probability of the spike being part of the broad-band, multicomponent background. This is because the probability is formulated on the assumption that the time-series contains only normally distributed noise. A low value for P flags a spike as being deserving of consideration as a possible candidate p mode. That said, it is important to stress that at no point does this *null* (H0) hypothesis test give the probability that the spike *is* a mode. Tests may be formulated that do. The p modes may, for example, be modelled as sine waves, and the probabilities calculated against the assumption the time-series contains normally distributed noise *and* an embedded sine wave (Gabriel et al. 2002), that is, the H1 hypothesis.

Crucial to fixing the value of P is the choice for N in equation (1). In what follows we shall assume frequency spectra are searched in slices of N independent bins at a time. The number K of independent N-bin slices that are searched will depend on the range of interest in the frequency spectrum. The total number of bins searched over the entire range of interest is then given by $M = KN$.

4   A. M. Broomhall et al.

If N is too large, the probability of finding at least one spike by chance within any one slice of N bins will increase. If N is too small the number K of independent slices over the entire region of interest will be large. The probability that at least one N-bin slice will contain at least one prominent noise spike by chance will then also increase. A trade-off must therefore be struck.

Here, we settled on a value for N that was equal to the number of bins in 100 μHz of the 3071-d BiSON and GOLF spectra, that is, $N = 26\,533$ bins. We performed searches in N-bin ranges between 550 and 1550 μHz, implying $K = 10$.

2.3 Use of contemporaneous data from two instruments

The discussion above gave some of the salient points on how to estimate whether or not a prominent spike is likely to be part of the background noise. It did so for searches of one spectrum. What if we have contemporaneous time-series from two different instruments to analyse? The advance we offer in this paper is a formalism for taking advantage of just such a situation, and the extra diagnostic power afforded by searches for coincidences that may then be performed.

We ask the question: if prominent spikes (or patterns of spikes) are uncovered in the same bin (or bins) of two frequency spectra, how does one calculate the joint probability of occurrence? At first glance this looks like a trivial problem. But it is not, because some proportion of the background noise will be correlated in the two (contemporaneous) time-series. This common noise, which is solar in origin, comes from the solar granulation. There will also be other independent sources of noise, as noted above, including some independent solar noise.

In the next section we come to the main part of the paper: how to calculate this joint probability, making allowance for the common noise.

3 FORMULATION OF JOINT PROBABILITY FOR SEARCHES IN TWO SPECTRA

We begin by defining a coefficient α to measure the proportion of noise which is common in two sets of contemporaneous noise data. The coefficient may be related formally to the coherence between the two sets (see equation 21 in Section 3 below). When $\alpha = 0$ there is no coherent data, but when $\alpha = 1$ the two sets are identical.

Next, we define x_1 and y_1 as the real and imaginary amplitudes of the first of two contemporaneous noise spectra. These amplitudes may be described by

$$x_1 = \alpha a + \alpha b + (1 - \alpha)a = a + \alpha b, \tag{2}$$

$$y_1 = \alpha c + \alpha d + (1 - \alpha)c = c + \alpha d, \tag{3}$$

where a, b, c and d are arrays of normally distributed random numbers. We also define x_2 and y_2 as the real and imaginary amplitudes of the second spectrum. They are given by the equations

$$x_2 = \alpha a + \alpha b + k(1 - \alpha)b, \tag{4}$$

$$y_2 = \alpha c + \alpha d + k(1 - \alpha)d. \tag{5}$$

Here $\alpha a + \alpha b$ represents the *common* solar noise, and $(1 - \alpha)a$ and $(1 - \alpha)b$, in x_1 and x_2 respectively, represent all 'other noise' that is *not* common to both sets of observations. Similarly $\alpha c + \alpha d$ represents the imaginary common solar noise and $(1 - \alpha)c$ and

$(1 - \alpha)d$, in y_1 and y_2, respectively, represent all 'other noise' in the imaginary part of the spectrum.

It is reasonable to assume that the variance of the real and imaginary parts of each spectrum will be equal. However, any individual noise that is not common to both spectra may have a different variance in the first spectrum compared to that in the second spectrum. Therefore, while the variances of the distributions, a, b, c and d have been taken to be equal (i.e. $\sigma_a^2 = \sigma_b^2 = \sigma_c^2 = \sigma_d^2$) we have also included a multiplying factor of k to alter the variance of the 'other noise' in x_2 and y_2. The variance of the 'other noise' in x_2 is therefore k^2 times the variance of the 'other noise' in x_1.

Next, we define p to be the probability that in a given bin the real amplitude x_i is greater than some threshold ξ_i, and that the imaginary amplitude y_i is greater than some threshold χ_i. The threshold amplitudes are in practice set to high levels so that prominent spikes whose amplitudes exceed the thresholds may be flagged as being of interest as potential candidate modes. Put more formally: should the amplitudes of a spike exceed ξ_i and χ_i, the probability of that spike being present by chance will be small (subject, of course, to the assumptions made on the distribution of the noise).

It can be shown that (see Appendix A)

$$
\begin{aligned}
p = {} & \frac{1}{4\pi^2 \sigma_a^4 [\alpha^2 + \alpha(k-1) - k]^2} \int_{y_2 = \chi_2}^{\infty} \int_{y_1 = \chi_1}^{\infty} \\
& \times \int_{x_1 = \xi_1}^{\infty} \int_{x_2 = \xi_2}^{\infty} \exp\left\{ \frac{-1}{2\sigma_a^2 [\alpha^2 + \alpha(k-1) - k]^2} [(\alpha^2 \right. \\
& + [k(1 - \alpha) + \alpha]^2)(x_1^2 + y_1^2) + (1 + \alpha^2)(x_2^2 + y_2^2) \\
& \left. - 2\alpha[1 + k(1 - \alpha) + \alpha](x_1 x_2 + y_1 y_2)] \right\} \mathrm{d}x_1 \, \mathrm{d}x_2 \, \mathrm{d}y_1 \mathrm{d}y_2.
\end{aligned}
\tag{6}
$$

Changing the variables, such that $x_i = \rho_i \cos\theta_i$ and $y_i = \rho_i \sin\theta_i$, and replacing θ_1 and θ_2 with the difference between them, θ, gives

$$
\begin{aligned}
p = {} & \frac{1}{2\pi\sigma_a^4 [\alpha^2 + \alpha(k-1) - k]^2} \int_{\theta=0}^{2\pi} \int_{\rho_2 = r_2}^{\infty} \int_{\rho_1 = r_1}^{\infty} \\
& \times \exp\left\{ \frac{1}{2\sigma_a^2 [\alpha^2 + \alpha(k-1) - k]^2} [2\alpha[1 + k(1 - \alpha) \right. \\
& + \alpha]\rho_1\rho_2 \cos\theta - (\alpha^2 + [k(1 - \alpha) + \alpha]^2)\rho_1^2 \\
& \left. - (1 + \alpha^2)\rho_2^2] \right\} \rho_1\rho_2 \mathrm{d}\rho_1 \mathrm{d}\rho_2 \mathrm{d}\theta.
\end{aligned}
\tag{7}
$$

Equation (7) may be integrated, with the help of the modified Bessel function (see Appendix B), to give the sum

$$
\begin{aligned}
p = {} & \frac{1}{\sigma_a^4 [\alpha^2 + \alpha(k-1) - k]^2} \int_{r_1}^{\infty} \int_{r_2}^{\infty} \sum_{n=0}^{\infty} \frac{1}{(n!)^2} \\
& \times \left(\frac{\alpha[1 + k(1 - \alpha) + \alpha]\rho_1\rho_2}{2\sigma_a^2 [\alpha^2 + \alpha(k-1) - k]^2} \right)^{2n} \\
& \times \exp\left\{ \frac{-1}{2\sigma_a^2 [\alpha^2 + \alpha(k-1) - k]^2} [(\alpha^2 \right. \\
& \left. + [k(1 - \alpha) + \alpha]^2)\rho_1^2 + (1 + \alpha^2)\rho_2^2] \right\} \rho_1\rho_2 \, \mathrm{d}\rho_1 \, \mathrm{d}\rho_2.
\end{aligned}
\tag{8}
$$

It is then possible to integrate equation (8) (see Appendix C) to give the sum

$$p = \sum_{n=0}^{\infty} \frac{1}{(n!)^2} \frac{\{\alpha[1 + k(1-\alpha) + \alpha]\}^{2n}}{2^{2n}\sigma_a^{4n+4}[\alpha^2 + \alpha(k-1) - k]^{4n+2}}$$
$$\times \frac{1}{4A_1 A_2}\left[\sum_{m=0}^{n} \frac{n! r_1^{2(n-m)}}{(n-m)! A_1^m}\right]\left[\sum_{q=0}^{n} \frac{n! r_2^{2(n-q)}}{(n-q)! A_2^q}\right]$$
$$\times \exp\left\{-A_1 r_1^2 - A_2 r_2^2\right\}, \tag{9}$$

where

$$A_1 = \frac{\alpha^2 + [k(1-\alpha) + \alpha]^2}{2\sigma_a^2[\alpha^2 + \alpha(k-1) - k]^2} \tag{10}$$

and

$$A_2 = \frac{(1+\alpha^2)}{2\sigma_a^2[\alpha^2 + \alpha(k-1) - k]^2}. \tag{11}$$

Equation (9) gives the probability that a spike in the first set of data has an amplitude greater than r_1 and that a spike which lies in the *same* frequency bin in the second set of data has an amplitude greater than r_2. Substitution of equation (9) into equation (1) yields the probability of at least one pair of such spikes occurring by chance in *N* bins of the spectrum.

Equation (9) converges as the number of terms, *n*, increases. We found that use of 20 terms gave accurate probability estimates. Simulations were performed, with pairs of randomly generated, complex Gaussian noise spectra having various α, which confirmed equation (9) gave the expected results.

Equation (9) contains three variables that need to be found from the data, namely α, *k* and σ_a^2. We now describe how to estimate the values. Let V_1 be the variance of the real part of the first spectrum defined by equation (2). Since x_1 is the sum of two Gaussian distributions it will have a zero mean, therefore the variance is given by the mean of x_1^2:

$$V_1 \equiv \langle x_1^2 \rangle = (1+\alpha^2)\sigma_a^2. \tag{12}$$

Let V_2 be the variance of the real part of the second spectrum defined by equation (4). This is then given by

$$V_2 \equiv \langle x_2^2 \rangle = \{\alpha^2 + [k(1-\alpha) + \alpha]^2\}\sigma_a^2. \tag{13}$$

The variances may be estimated in practice by measurement of the mean power in each spectrum in the 100-μHz slice being searched. Additional information is provided by calculation of the mean cross-amplitude (Elsworth et al. 1994), over the same 100-μHz slice. The mean cross-amplitude of the two spectra, $X_{1,2}$, is given by

$$X_{1,2} \equiv \langle (x_1 + y_1)(x_2 - y_2) \rangle$$
$$= 2\alpha\sigma_a^2(1 + k[1-\alpha] + \alpha). \tag{14}$$

Equations (12)–(14) may then be used to find values for α, σ_a^2 and *k*. Equation (12) may be rearranged to find σ_a^2 in terms of α and V_1:

$$\sigma_a^2 = \frac{V_1}{1+\alpha^2}. \tag{15}$$

Equation (15) may be substituted into equations (13) and (14) to give

$$V_2 = \{\alpha^2 + [k(1-\alpha) + \alpha]^2\}\frac{V_1}{1+\alpha^2} \tag{16}$$

and

$$X_{1,2} = \frac{2\alpha V_1}{1+\alpha^2}[1 + k(1-\alpha) + \alpha]. \tag{17}$$

From equation (17) it may be shown that

$$k = \frac{(1+\alpha)X_{1,2}}{2\alpha V_1} - 1, \tag{18}$$

which may be substituted into equation (16) to give

$$V_2 = \frac{4\alpha V_1(1+\alpha^2)(\alpha V_1 - X_{1,2}) + X_{1,2}^2(1+\alpha^2)^2}{4\alpha^2 V_1(1+\alpha^2)}. \tag{19}$$

The only unknown in equation (19) is now α. Equation (19) may be used to find a quartic in terms of α:

$$\left(4V_1^2 + X_{1,2}^2 - 4V_1 V_2\right)\alpha^4 - 4V_1 X_{1,2}\alpha^3$$
$$+ \left(4V_1^2 + 2X_{1,2}^2 - 4V_1 V_2\right)\alpha^2 - 4V_1 X_{1,2}\alpha + X_{1,2}^2 = 0. \tag{20}$$

Use of equations (15) and (18) then yields estimates of σ_a^2 and *k*, respectively. As an aside, we note that since α is found using an equation for the cross-amplitude it is possible to relate α to the coherency between the two spectra, $Y_{1,2}$:

$$Y_{1,2} = \frac{X_{1,2}}{\sqrt{\langle x_1^2 + y_1^2 \rangle \langle x_2^2 + y_2^2 \rangle}}$$
$$= \frac{\alpha(1 + k[1-\alpha] + \alpha)}{\sqrt{(1+\alpha^2)\{\alpha^2 + [k(1-\alpha) + \alpha]^2\}}}. \tag{21}$$

The coherency gives the proportion of the cross-amplitude spectrum that is due to coherent signal.

4 RESULTS ON CONTEMPORANEOUS BISON AND GOLF TIME-SERIES

4.1 Observational data

The BiSON and GOLF time-series each consisted of 3071 d of Sun-as-a-star Doppler velocity observations, made between 1996 April 20 and 2004 September 15, an epoch spanning most of so-lar activity cycle 23. Both data sets were made for the Phoebus collaboration. The BiSON data were processed in the manner described by Appourchaux et al. (2000) and Chaplin et al. (2002). The GOLF data were processed in the manner described by García et al. (2005). The processing gave two time-series sampled at different cadence. The nominal cadence on which the BiSON data are stored is 40 s. The corresponding cadence for the GOLF data is 20 s. Both time-series were rebinned to a common cadence of 120 s. The time-series therefore comprised 2211120 samples, giving a binwidth in the frequency domain of 3.77 nHz. The duty cycle of the BiSON time-series was 78.6 per cent; the duty cycle of the GOLF time-series was 93.4 per cent. The frequency spectra of the 120-s time-series were then searched for prominent spikes.

4.2 Searches for candidate modes: tests applied to data

We searched both spectra, in slices of 100 μHz, for the presence of prominent spikes, or prominent patterns of spikes, in the same frequency bin, or bins. The tests were based on those discussed in detail by Chaplin et al. (2002). The tests as outlined in that paper involved searches for prominent spikes, or prominent patterns of spikes, in a single frequency–power spectrum. The *power* spectrum was searched because Chaplin et al. (2002) formulated the probability *p* (Section 2.2) in terms of the negative exponential statistics displayed by the noise background in power. For this paper the probability *p* was instead given by equation (9). This probability is based on a description of the complex noise in amplitude, and gives the probability that a spike in the same frequency bin in *both* the

6   *A. M. Broomhall et al.*

Figure 1. Upper panel: variance of the amplitude in the BiSON (red, dashed line) and GOLF (black, solid line) data sets, equivalent to, respectively, V_1 and V_2 in Section 3; and the cross-amplitude of the BiSON and GOLF data sets (blue, dotted line), which is equivalent to $X_{1,2}$ in Section 3). Lower panel: value of α calculated using equation (20).

BiSON and GOLF frequency–amplitude spectra is greater than a certain threshold in *amplitude*. Hence, we searched the frequency–amplitude spectra.

The calculation of values of p required estimates of the parameters α, k and σ_a^2 locally in frequency. The parameters, and procedure, were described above, in Section 3. The procedure involved measurement of the variance of the amplitude in the BiSON and GOLF frequency spectra, and calculation of the BiSON–GOLF cross-amplitude spectrum. The upper panel of Fig. 1 shows estimates of the spectrum variances and the cross-amplitude, while the lower panel shows the coefficient α, which measures the proportion of common noise in the two sets.

Next, we give a brief description of each of the tests that were applied to the BiSON and GOLF data.

Test 1: This is a search for *a prominent spike* in the same frequency bin of each spectrum. We call this a 'pair of spikes'. Substitution of the observed amplitudes of the spikes into equation (9) yields an estimate for the joint probability p, which is the probability of both spikes appearing by chance in the same frequency bin of each spectrum. The probability, P, that this pair occurs by chance at least once over a range of N bins is then given by use of equation (1), which requires the calculated p as input.

Only those occurrences with $P \leqslant 1$ per cent were recorded for further scrutiny. We adopted the same probability threshold for the other tests.

The next two tests admit the possibility that power from a mode may be spread over two, or more, consecutive bins. Power-law extrapolation to lower frequencies of the widths of modes between ~1500 and ~2000 µHz suggests that confinement of the majority of power within a single bin of a 3071-d spectrum is expected at mode frequencies below about ~950 µHz (Chaplin et al. 2002). Even if the peak is not resolved in frequency, its maximum amplitude may be diminished, and the majority of its amplitude divided between two consecutive bins, if the mode signal is not commensurate with the window function of the observations. We therefore also applied the 'bin-shifting' strategy of Chaplin et al. (2002) in an attempt to circumvent this problem. [An alternative strategy involves zero-padding the time-series; see, for example, Gabriel et al. (2002).]

Test 2: This is a search for the presence of *two prominent spikes* in the same consecutive frequency bins of each spectrum. The amplitudes of the least prominent spikes in each spectrum are used in the calculation of p. The formula for P now differs from, and is more complicated than, the 'single-spike' equation (1). We used equation (13) of Chaplin et al. (2002).

Test 3: This is a search for *two or more pairs of spikes*. Spikes constituting each pair must of course lie in the same frequency bin of each spectrum. But the bins occupied by different pairs may lie up to twice the predicted peak width apart. The power-law extrapolation, mentioned briefly above, was used to derive a look-up curve of the predicted width as a function of frequency. We used equation (12) of Chaplin et al. (2002) to determine estimates of P.

The final two tests look for patterns of prominent spikes, where those patterns mimic the patterns expected from the rotational splitting (with allowance made for the influence of magnetic fields). Only every other component of the non-radial modes is in principle detectable in the BiSON and GOLF Sun-as-a-star data. That is because the Sun-as-a-star observations have very poor sensitivity to the $l + m$ odd components. The mean observed (synodic) separation between the $l + m$ even components of high-S/N modes is about 0.8 µHz.

Test 4: This test is a search for *two pairs of spikes*. Spikes constituting each pair must of course lie in the same frequency bin of each spectrum. And bins occupied by different pairs must be separated by 0.8 ± 0.1 µHz. We used equation (C6) of Chaplin et al. (2002) to determine estimates of P. This test can reveal patterns of spikes that may be candidates for $l = 1, 2$ or 3 modes.

Test 5: This test is a search for *three pairs of spikes*. Spikes constituting each pair must of course lie in the same frequency bin of each spectrum. And bins occupied by different pairs must be separated by multiples of 0.8 ± 0.1 µHz. We used equation (C7) of Chaplin et al. (2002) to determine estimates of P. This test can reveal patterns of spikes that may be candidates for $l = 2$ or 3 modes. (The $l = 1$ modes would only reveal two visible components in the Sun-as-a-star data.)

The choice of the mean separation, and the choice of the allowable range in that separation (± 0.1 µHz), was discussed at length in Chaplin et al. (2002). In summary, the choices were predicated on assumptions that: (i) the rotational frequency splitting of very low-frequency p modes would not differ significantly from the splitting measured in the high-S/N modes and (ii) departures from a uniform spacing (e.g. due to the influence of magnetic fields) would be no more severe than in the high-S/N modes.

Thresholds for BiSON data

Figure 2. Upper panel: threshold amplitudes (for BiSON) needed to record $P \leqslant 1$ per cent in each of the five tests (see figure legend). Lower panel: percentage reduction in the threshold amplitudes, with respect to the amplitudes given by analysis of the BiSON data alone.

4.3 Results

We begin our discussion of the results by showing threshold amplitudes for $P \leqslant 1$ per cent. The upper panel of Fig. 2 plots the absolute thresholds (in units of mm s^{-1} per root bin) for the BiSON data, as a function of frequency, for each of the five tests. By how much have these threshold amplitudes been reduced by making use of two contemporaneous data sets, as opposed to just one data set? The lower panel of Fig. 2 plots the percentage reduction in the threshold amplitudes, with respect to the amplitudes given by analysis of the BiSON data alone. Inspection of the plot reveals that use of two sets has reduced the threshold amplitudes by about 28 per cent on average. The percentage reduction is seen to be frequency dependent.

The percentage reduction in the threshold amplitudes is not far from the best case reduction, corresponding to $\alpha = 0$. Then, there would be no common background noise in the two data sets. At a particular frequency, v, the threshold amplitude in the BiSON spectrum is given by some value, s_v say, times the mean noise amplitude in the BiSON spectrum. At the same frequency the threshold amplitude in the GOLF spectrum is given by the mean noise ampli-

tude times this same value, s_v. Therefore, when $\alpha = 0$, the threshold amplitudes would both be reduced by a factor $\sqrt{2}$, implying a $\sim$29 per cent reduction in size.

Fig. 3 shows visually locations in frequency where prominent spikes, or patterns of spikes, were recorded. We also mark other frequencies of interest. The plot is an echelle plot. The entire range that was searched has been split into strips of 135 µHz. The strips have then been placed one above another. We chose a repeat frequency of 135 µHz, rather than the 100-µHz slice on which searches were performed, because consecutive overtones of the low-l p modes are separated in frequency by about this amount. The black vertical lines mark locations of the frequencies predicted by the Saclay seismic model (Turck-Chièze et al. 2001); the choice of repeat frequency causes the black lines to be arranged in four near-vertical strips, one strip for each l.

The vertical dashed red lines mark locations in frequency that are overtones of the 11.57-µHz diurnal frequency. Artefacts may occasionally be expected to show up in the frequency spectrum because of window function (for the ground-based BiSON) or spacecraft operation (for the GOLF) signatures.

Locations in frequency where spikes, or patterns of spikes, were found in the same bin, or bins, of each spectrum at levels sufficient to record $P \leqslant 1$ per cent are marked by the black symbols in the middle of each row. A different symbol has been used for each test (see figure legend). We have also recorded prominent spikes or patterns of spikes found by searching either the BiSON (red symbols at top of each row) or GOLF (blue symbols at bottom of each row) spectrum alone.

Symbols surrounded by a green square mark prominent spikes or patterns of spikes that are listed in Table 1. These occurrences all lie within ± 0.5 µHz of the mode frequencies predicted by the Saclay seismic model (Turck-Chièze et al. 2001). We have identified the occurrences with particular p-mode components on the basis of their placement in frequency, and by comparison with previous detections. In cases where the outer components of a rotationally split mode appear to have been uncovered, we have calculated the mean of those outer frequencies to give an estimate of the centroid frequency of the mode. Estimates of the uncertainties in frequency were calculated in the manner described by Chaplin et al. (2002).

We find no new candidate detections. All listed candidates correspond to previously claimed detections of p modes (e.g. see Toutain et al. 1998; Bertello et al. 2000; García et al. 2001; Chaplin et al. 2002). However, by searching for coincidences in contemporaneous data we have reduced the threshold amplitudes required to record occurrences with $P \leqslant 1$ per cent by $\sim$29 per cent, compared to the single-spectrum case. Or, put in terms of power, the demands placed on the S/N in power have been reduced by $\sim$58 per cent.

We make three additional important points. First, it is apparent from Fig. 2 that the analysis uncovers several occurrences of $P \leqslant 1$ per cent that lie well away from the predicted model frequencies. It is also interesting to note that while these occurrences do not sit precisely on the diurnal frequencies (vertical dashed red lines), the majority do lie fairly close to the diurnal frequencies.

Secondly, we demanded that prominent spikes, or prominent patterns of spikes, had $P \leqslant 1$ per cent to merit further consideration as potential candidates for modes. This was in contrast to previous studies in the literature (e.g. Appourchaux et al. 2000; García et al. 2001; Chaplin et al. 2002; Gabriel et al. 2002), which adopted a value of $P \leqslant 10$ per cent. The change to a probability threshold of 1 per cent was agreed recently by the Phoebus consortium. We note that use of a 10 per cent threshold did not result in any further candidates being flagged.

8   *A. M. Broomhall et al.*

Figure 3. Echelle plot, modulo 135 μHz, marking locations in frequency of occurrences uncovered by the test searches. Locations in frequency where spikes, or patterns of spikes, were found in the same bin, or bins, of each spectrum at levels sufficient to record $P \leqslant 1$ per cent are marked by the black symbols in the middle of each row. A different symbol has been used for each test (see figure legend). We have also recorded prominent spikes or patterns of spikes found by searching either the BiSON (red symbols at top of each row) or GOLF (blue symbols at bottom of each row) spectrum alone. Symbols surrounded by a green square represent the prominent occurrences listed in Table 1. The black vertical lines mark locations of the frequencies predicted by the Saclay seismic model (Turck-Chièze et al. 2001). The vertical dashed red lines mark locations in frequency that are overtones of the 11.57-μHz diurnal frequency.

Table 1. Candidates found to be closer than 0.5 μHz to predicted frequencies of modes from the Saclay seismic model.

l	n	m	Frequency (μHz)	Probability (P)	Number of tests passed	Distance from model in frequency (μHz)
0	6	0	972.613 ± 0.002	3.4×10^{-4}	2	0.132
1	7	-1	1185.196 ± 0.005	1.9×10^{-3}	4	
1	7	$+1$	1185.981 ± 0.005	3.0×10^{-3}	1	
		Mean	1185.589 ± 0.004			0.027
0	8	0	1263.210 ± 0.007	6.0×10^{-4}	3	0.314
1	8	-1	1329.236 ± 0.005	0.0000	4	
1	8	$+1$	1330.037 ± 0.007	1.7×10^{-3}	1	
		Mean	1329.637 ± 0.004			0.060
2	8	-2	1393.871 ± 0.007	4.0×10^{-4}	1	0.036^a
0	9	0	1407.479 ± 0.007	4.0×10^{-8}	3	0.149
1	9	-1	1472.432 ± 0.008	3.8×10^{-11}	4	
1	9	$+1$	1473.269 ± 0.009	3.7×10^{-8}	4	
		Mean	1472.851 ± 0.006			0.122

^{a}Difference with model frequency assumes $m = -2$ component lies -0.8 μHz from central frequency.

Thirdly, it is evident from Table 1 and Fig. 3 that there are gaps in the list of candidate detections. The S/N is evidently low in this part of the spectrum, and destructive interferences with noise would be expected to hinder searches for the modes; but is the number of recorded candidates consistent with expectations based on estimates of the powers of the modes (e.g. from extrapolation to lower frequencies of the observed powers of the high-S/N modes)? Even though use of contemporaneous data reduced the threshold

amplitudes needed to flag possible detections, we found no new modes. Is this a puzzle, or is the result expected given estimates of the likely power of these modes? A follow-up study is being conducted, from which we hope to obtain answers to these questions.

5 SUMMARY

We have shown how to take advantage of contemporaneous data from two different instruments in the search for low-frequency modes of oscillation of the Sun. Use of contemporaneous data allows searches to be made for prominent, sharp concentrations of power which are coincident in frequency, and which are significantly above the local background noise. Occurrences of this type are potential candidate modes.

Crucial to determining objective measures of the joint probability of the random occurrence of such features is a good understanding of the characteristics of the background noise. In this paper we showed how to make proper allowance in the calculation of the probability for noise that is common to data from different instruments. This common noise is solar in origin, and comes from the solar granulation. The main part of the paper gave a derivation of the formula that is needed to compute the probability, and a recipe for how to calculate from the two data sets the observational parameters required as input to the formula.

We demonstrated application of the technique by searching two time-series of Sun-as-a-star Doppler velocity data for low-frequency p modes. The data comprised 3071 d of contemporaneous Sun-as-a-star observations made by the ground-based BiSON, and the GOLF instrument onboard the ESA/NASA *SOHO* spacecraft. Both data sets were prepared for the Phoebus collaboration. Use of the two data sets, as opposed to just one, reduced by more than 50 per cent the S/N in power needed to flag potential mode candidates.

REFERENCES

Abramowitz M., Stegun I., 1965, Handbook of Mathematical Functions. Dover Books on Advanced Mathematics. Dover Press, New York
Appourchaux T. et al., 2000, ApJ, 538, 401
Bertello L. et al., 2000, ApJ, 535, 1066
Chaplin W., Elsworth Y., Isaak G., Marchenkov K., Miller B., New R., Pinter B., 2002, MNRAS, 336, 979

Elsworth Y., Howe R., Isaak G., McLeod C., Miller B., New R., Speake C., Wheeler S., 1994, MNRAS, 269
Gabriel A. H. et al., 1995, Sol. Phys., 162, 61
Gabriel A. H. et al., 2002, A&A, 390, 1119
García R. A. et al., 2001, Sol. Phys., 200, 361
García R. A. et al., 2005, A&A, 442, 385
Provost J., Berthomieu G., Morel P., 2000, A&A, 353, 775
Toutain T., Appourchaux T., Fröhlich C., Kosovichev A. G., Nigam R., Scherrer P. H., 1998, ApJ, 506, L147
Turck-Chièze S. et al., 2001, ApJ, 555, L69
Turck-Chièze S. et al., 2004, ApJ, 604, 455

APPENDIX A: DERIVATION OF EQUATION (5)

The general expression for a probability density function is

$$P(X) = \frac{1}{(2\pi)^{z/2}} \frac{1}{\sqrt{|C|}} \exp\left(-\frac{1}{2} X^{\mathrm{T}} C^{-1} X\right), \tag{A1}$$

where z is the number of variables. Our analysis uses two amplitude spectra that each has a real and imaginary part: so $z = 4$. X is the matrix of the variables given by

$$X = \begin{pmatrix} x_1 \\ y_1 \\ x_2 \\ y_2 \end{pmatrix}. \tag{A2}$$

C is the coherency matrix given by the expectation values

$$C = \begin{pmatrix} E(x_1^2) & E(x_1 y_1) & E(x_1 x_2) & E(x_1 y_2) \\ E(y_1 x_1) & E(y_1^2) & E(y_1 x_2) & E(y_1 y_2) \\ E(x_2 x_1) & E(x_2 y_1) & E(x_2^2) & E(x_2 y_2) \\ E(y_2 x_1) & E(y_2 y_1) & E(y_2 x_2) & E(y_2^2) \end{pmatrix}. \tag{A3}$$

Take the first element as an example. It is given by

$$\begin{aligned} E(x_1^2) &= \langle a^2 + 2aab + \alpha^2 b^2 \rangle \\ &= \langle a^2 \rangle + 2\alpha\langle ab \rangle + \alpha^2 \langle b^2 \rangle \\ &= \sigma_a^2 + \alpha^2 \sigma_b^2, \end{aligned} \tag{A4}$$

since a and b have Gaussian distributions. We find the other expectation values in a similar manner and substitute $K = \alpha[1 + k(1 - \alpha) + \alpha]$ and $J = \alpha^2 + [k(1 - \alpha) + \alpha]^2$ so that C can be written as

$$C = \sigma_a^2 \begin{pmatrix} 1 + \alpha^2 & 0 & K & 0 \\ 0 & 1 + \alpha^2 & 0 & K \\ K & 0 & J & 0 \\ 0 & K & 0 & J \end{pmatrix}. \tag{A5}$$

It is possible to show that

$$|C| = \sigma_a^8 [\alpha^2 + \alpha(k - 1) - k]^4 \tag{A6}$$

and

$$\begin{aligned} C^{-1} = &\frac{-1}{\sigma_a^2 [\alpha^2 + \alpha(k - 1) - k]^2} \\ &\times \begin{pmatrix} J & 0 & -K & 0 \\ 0 & J & 0 & -K \\ -K & 0 & 1 + \alpha^2 & 0 \\ 0 & -K & 0 & 1 + \alpha^2 \end{pmatrix}. \end{aligned} \tag{A7}$$

These equations can then be substituted back into equation (A1) to give equation (6).

10   *A. M. Broomhall et al.*

APPENDIX B: INTEGRATION OF EQUATION (6) WITH RESPECT TO θ

The modified Bessel function of the first kind is defined in integral form as (Abramowitz & Stegun 1965)

$$I_0(z) = \frac{1}{\pi} \int_0^\pi \exp\{\pm z \cos\theta\} \, d\theta. \tag{B1}$$

It may also be expressed as a summation:

$$I_0(z) = 1 + \frac{\frac{1}{4}z^2}{(1)^2} + \frac{\left(\frac{1}{4}z^2\right)^2}{(2)^2} + \frac{\left(\frac{1}{4}z^2\right)^3}{(3)^2} + \cdots. \tag{B2}$$

Comparison with equation (7) shows that

$$\begin{aligned}
I_0 &= 1 + \left(\frac{\alpha[1 + k(1-\alpha) + \alpha]\rho_1\rho_2}{2\sigma_a^2[\alpha^2 + \alpha(k-1) - k]^2} \right)^2 \\
&\quad + \frac{1}{2^2} \left(\frac{\alpha[1 + k(1-\alpha) + \alpha]\rho_1\rho_2}{2\sigma_a^2[\alpha^2 + \alpha(k-1) - k]^2} \right)^4 \\
&\quad + \frac{1}{(3!)^2} \left(\frac{\alpha[1 + k(1-\alpha) + \alpha]\rho_1\rho_2}{2\sigma_a^2[\alpha^2 + \alpha(k-1) - k]^2} \right)^6 + \cdots, \\
&= \sum_{n=0}^\infty \frac{1}{(n!)^2} \left(\frac{\alpha[1 + k(1-\alpha) + \alpha]\rho_1\rho_2}{2\sigma_a^2[\alpha^2 + \alpha(k-1) - k]^2} \right)^{2n}.
\end{aligned} \tag{B3}$$

Applying this to equation (7) gives equation (8).

APPENDIX C: INTEGRATION OF EQUATION (7) WITH RESPECT TO ρ_1 AND ρ_θ

By substituting equations (10) and (11) into equation (8), and letting

$$B_n = \frac{1}{n!^2} \left(\frac{\alpha[1 + k(1-\alpha) + \alpha]}{2\sigma^2[\alpha^2 + \alpha(k-1) - k]^2} \right)^{2n}, \tag{C1}$$

the nth term of the sum in equation (8) reduces to

$$\begin{aligned}
p_{n=n} = &\frac{1}{\sigma_a^4[\alpha^2 + \alpha(k-1) - k]^2} \\
&\times \int_{r_1}^\infty \int_{r_2}^\infty B_n \rho_1^{2n+1} \rho_2^{2n+1} \exp\left(-A_1\rho_1^2 - A_2\rho_2^2 \right) d\rho_1 \, d\rho_2.
\end{aligned} \tag{C2}$$

For $i = 1, 2$ let

$$\text{Int}_{i,n} = \int_{r_i}^\infty \rho_i^{2n+1} \exp\left(-A_i\rho_i^2 \right) d\rho_i. \tag{C3}$$

So for $n = 0$

$$\begin{aligned}
\text{Int}_{i,0} &= \int_{r_i}^\infty \rho_i \exp\left(-A_i\rho_i^2 \right) d\rho_i \\
&= \frac{1}{2A_i} \exp\left(-A_i r_i^2 \right),
\end{aligned} \tag{C4}$$

for $n = 1$

$$\begin{aligned}
\text{Int}_{i,1} &= \int_{r_i}^\infty \rho_i^3 \exp\left(-A_i\rho_i^2 \right) d\rho_i \\
&= \frac{1}{2A_i} r_i^2 \exp\left(-A_i r_i^2 \right) + \frac{1}{A_i} I_{i,0} \\
&= \frac{1}{2A_i} r_i^2 \exp\left(-A_i r_i^2 \right) + \frac{1}{2A_i^2} \exp\left(-A_i r_i^2 \right),
\end{aligned} \tag{C5}$$

and for $n = 2$

$$\begin{aligned}
\text{Int}_{i,2} &= \int_{r_i}^\infty \rho_i^5 \exp\left(-A_i r_i^2 \right) d\rho_i \\
&= \frac{1}{2A_i} r_i^4 \exp\left(-A_i r_i^2 \right) + \frac{2}{A_i} I_{i,1} \\
&= \frac{1}{2A_i} r_i^4 \exp\left(-A_i r_i^2 \right) \\
&\quad + \frac{1}{A_i^2} r_i^2 \exp\left(-A_i r_i^2 \right) + \frac{1}{A_i^3} \exp\left(-A_i r_i^2 \right).
\end{aligned} \tag{C6}$$

In fact, more generally,

$$\text{Int}_{i,n} = \frac{1}{2A_i} \sum_{m=0}^n \frac{n! r_i^{2(n-m)}}{(n-m)! A_i^m} \exp\left(-A_i r_i^2 \right). \tag{C7}$$

Substitution back into equation (C2) for both $i = 1$ and 2, and replacing B_n with equation (C1), gives equation (9).

Appendix B

Submitted Publication

The work on the visibility of low frequency modes, which is described in Chapters 6 of this thesis, was submitted to be published by Astronomische Nachrichten in October 2007. The paper is included here for reference.

The visibility of low-frequency solar acoustic modes

A. M. Broomhall[1], **W. J. Chaplin**[1], **Y. Elsworth**[1], **S.T. Fletcher**[2]

[1] School of Physics and Astronomy, University of Birmingham, Edgbaston, Birmingham B15 2TT
[2] Faculty of Arts, Computing, Engineering and Sciences, Sheffield Hallam University, Sheffield, S1 1WB

The dates of receipt and acceptance should be inserted later

Key words Sun: helioseismology - - Sun: oscillations - - Methods: statistical

We make predictions of the detectability of low-frequency p modes. Estimates of the powers and damping times of these low-frequency modes are found by extrapolating the observed powers and widths of higher-frequency modes with large observed signal-to-noise ratios. The extrapolations predict that the low-frequency modes will have small signal-to-noise ratios and narrow widths in a frequency-power spectrum. Monte Carlo simulations were then performed where timeseries containing mode signals and normally distributed Gaussian noise were produced. The mode signals were simulated to have the powers and damping times predicted by the extrapolations. Various statistical tests were then performed on the frequency-amplitude spectra formed from these timeseries to investigate the fraction of spectra in which the modes could be detected. The results of these simulations were then compared to the number of p-modes candidates observed in real Sun-as-a-star data at low frequencies. The fraction of simulated spectra in which modes were detected decreases rapidly as the frequency of modes decreases and so the fraction of simulations in which the low-frequency modes were detected was very small. However, increasing the signal-to-noise (S/N) ratio of the low-frequency modes by a factor of 2 above the extrapolated values led to significantly more detections. Therefore efforts should continue to further improve the quality of solar data that is currently available.

1 Introduction

To date a multitude of solar acoustic (p) modes have been observed over a wide range of frequencies. However, no independently confirmed detections of low-degree (low-l) p modes, with frequencies below $\sim 970\,\mu$Hz, have been made (e.g. García et al. 2001, Chaplin et al. 2002, Broomhall et al. 2007). As low-frequency p modes are expected to have very long lifetimes their detection would allow their frequencies to be measured to very high accuracies and precisions. This is crucially important as the properties of low-l p modes are affected by conditions deep in the solar interior and so their frequencies act as a probe of these regions. However, the signal from low-frequency p modes is very weak. Therefore, different analysis procedures have been developed in attempts to detect these modes.

Broomhall et al. (2007) used statistical techniques to look for coincident prominent features in BiSON and GOLF frequency-amplitude spectra. They found that many low-frequency p modes remain undetected despite reducing amplitude detection thresholds to less than $3\,\mathrm{mm\,s^{-1}\,bin^{-1/2}}$. However, Broomhall et al. found the $l = 0, n = 6$ mode at $\sim 973\,\mu$Hz to be extremely prominent. This mode has also been detected in other studies such as Garcia et al. (2001) and Chaplin et al. (2002). The prominence of this mode is conspicuous because of the lack of evidence for higher-frequency modes such as the $l = 0, n = 7$ mode at $\sim 1118\,\mu$Hz. As yet this mode has only been observed in Sun-as-a-star data by Garcia et al. (2001). These higher-frequency modes should be easier to observe as theoreti-

cally they should exhibit larger amplitudes in the photosphere, where the observations are made. It is, therefore, of interest to investigate how prominent the $l = 0, n = 6$ mode, and the modes surrounding it in frequency, should be. This will provide an indication as to whether we should be observing modes that we are not; or whether some effect makes the $l = 0, n = 6$ mode more prominent than the rest.

How easy low-frequency modes are to detect depends upon their power and width in a frequency-power spectrum. As many low-frequency modes have not yet been detected their powers cannot be determined from solar data directly. Here, predictions of the power in very low-frequency modes have been made by extrapolating the power observed in well-defined, higher-frequency modes. The height of a mode in a frequency spectrum also depends upon the width in frequency over which the power is spread. As the power in the signal from a mode is damped over time the signal from the mode may be spread over several bins of a frequency spectrum. The lifetime of a mode varies with frequency and so it is also possible to infer from extrapolation the width a mode at a particular frequency is expected to exhibit from results for well-observed modes. The details of how the extrapolations were made will be explained in Section 2. To make the extrapolations we have assumed that the simple functional relationships that turn out to describe the powers and heights of the well-observed modes are still valid at low frequencies. The validity of this assumption was tested using predictions of the powers and widths of low-frequency modes made by the Cambridge stochastic excitation and damping codes (e.g. Houdek et al. 1999). Monte Carlo sim-

ulations based on the results of the extrapolations were then performed to investigate how often simulated modes could be detected (Section 3). The simulations involved creating timeseries to mimic Sun-as-a-star data. The simulated data contained mode signals that have the properties predicted by the extrapolations. Various statistical tests, which are described in Section 4, were performed on the frequency-amplitude spectra that were created from the simulated timeseries. The statistical tests followed those described in Broomhall et al. (2007) and determined how often the simulated modes could be detected. The results of these tests are detailed in Section 5. A discussion of the results is presented in Section 6. In this discussion we pay particular attention to the unusual $l = 0$, $n = 6$ radial mode.

2 Predicting the Widths and Powers of modes

Various properties of modes can be found by fitting frequency-power spectra. These properties include the width of a mode, and therefore its damping time, and the height of a mode, i.e. its maximum power spectral density. The product of the height and width is proportional to the total power in the mode. A BiSON spectrum, consisting of 3071 d of Sun-as-a-star Doppler velocity observations, was fitted using the methods described in Fletcher (2007). The fitting procedure involved taking the Fourier transform of the timeseries to produce a frequency-power spectrum and then fitting a Lorentzian-like model to the various mode peaks in the resulting frequency-power spectrum. The observations were made between 1996 April 20 and 2004 September 15, an epoch which spans most of solar activity cycle 23. The data were processed in the manner described by Appourchaux et al. (2000) and Chaplin et al. (2002). The data were the same as the BiSON data used in Broomhall et al. (2007) except that, instead of being rebinned to have a cadence of 120 s, the data were left with the nominal 40 s cadence on which the BiSON data are stored.

The visibility of a mode in a frequency-power spectrum will depend on the power and lifetime of the mode. We will now briefly discuss how these properties affect the visibility.

2.1 How the power and width of a mode affects its visibility

The width of a mode, Δ, is related to its lifetime, τ, by

$$\Delta = \frac{1}{\pi\tau}. \tag{1}$$

If the length of a timeseries is significantly longer than the lifetime of a mode the timeseries will extend over several realizations of the mode and so the modal peak will be resolved across several bins in the frequency domain. However, if a mode's lifetime is significantly longer than the

length of a timeseries all of the mode's power will be contained in a single bin[1].

If the width of a mode can be resolved the power of that mode, V^2, is given by

$$V^2 = \frac{\pi}{2}T\Delta H, \tag{2}$$

where T is the length of the timeseries and H is the maximum power spectral density per bin. When the power spectrum is fitted it is actually the maximum power spectral density, which corresponds to the height, H, of the Lorentzian, that is determined. The width, Δ, and height, H, of the mode can then be used to find the power of the mode. In this paper the observed signal-to-noise ratio of a mode refers to the ratio of the height of the most prominent spike across the width of the mode and the mean level of the background around it. We define a spike as the power (or amplitude) contained in one bin in a frequency-power (or amplitude) spectrum. The height of the most prominent spike of a resonant peak in a frequency-power spectrum will be greater than the height of the fitted Lorentzian. This is because of the random nature of the excitation of the modes, which means that the power will have a χ^2 2 degrees of freedom distribution about the underlying Lorentzian. Therefore, the power in some of the bins across the width of the mode will be greater than the height of the Lorentzian. These are the spikes that are most likely to be detected.

Notice that it is the power, V^2, that we are going to extrapolate and not the peak height of the mode. This is because the height, H, of the mode is not a smooth function of frequency, as it depends on whether the power of the mode is spread over more than one bin or whether the mode's power is confined to one bin only. On the other hand the power of a mode is a smooth function of frequency and so can be extrapolated more easily.

A well-defined mode is one with a power and a width that are sufficiently large for the mode to be detected easily, allowing the shape of the mode in a frequency-power spectrum to be fitted accurately. In the case of the Sun, when fitting a set of data that spans $\sim$ 8.5 yr, well-defined modes have frequencies greater than $\sim$ 1500 μHz. Here we are concerned with modes at low frequencies that cannot be observed clearly. Therefore it was necessary to extrapolate the widths and powers of well-defined modes to obtain estimates of the widths and powers of low-frequency modes. We have, therefore, assumed that simple functional relationships, which describe the variation of the parameters with frequency at higher frequencies, persist to lower frequencies. In section 2.5 we have used predictions from a stochastic excitation model to test the validity of this assumption. The extrapolation can be performed in two different ways, each of which will now be described in turn.

---

[1] The power can be split between 2 bins if the frequency of the mode is not commensurate with the spectrum's frequency bins. Also the finite length of a timeseries means that the mode will appear in a frequency-power spectrum as a sinc squared function.

Fig. 1 Results for method 1. Top panel: The black crosses show how the natural logarithm of the observed width of a mode varies with frequency. The errors of the widths are those associated with the fits. The red solid line shows the linear line of best fit for the widths at frequencies between $\sim 1530\,\mu$Hz and $\sim 2400\,\mu$Hz. The blue dashed lines indicate the errors on this linear fit. Bottom panel: The black crosses show how the natural logarithm of a mode's power varies with frequency. The red solid line shows the linear line of best fit for the powers at frequencies between $\sim 1540\,\mu$Hz and $\sim 2500\,\mu$Hz. The blue dashed lines indicate the errors on this fit.

2.2 Method 1: The ln-linear relationship

An approximately linear relationship is observed between the natural logarithm of a mode's width and its frequency for well-observed modes with frequencies below $\sim 2400\,\mu$Hz (see the top panel of Figure 1). The well-defined plateau in the mode widths above $\sim 2400\,\mu$Hz (see for example Chaplin et al. 2005) means that the linearity does not extend to higher frequencies. The observed widths have been fitted using data from modes with different degrees. In Sun-as-a-star observations only low-l modes can be clearly observed. Here widths from $l = 0$, 1 and 2 modes have been used. The width of a mode is not independent of degree as high-l modes have a lower inertia than low-l modes and so high-l modes have faster damping rates. However, the difference in damping rates is minimal over the confined range of l used here.

The relationship between the natural logarithm of a mode's power and its frequency is also approximately linear for well-defined modes with frequencies below $\sim 2500\,\mu$Hz (see the bottom panel of Figure 1). The power of a mode is dependent on its degree and so the values of the coefficients of the linear relationship are specific to each value of

Table 2 Predicted widths of various $l = 0$ modes found by extrapolating both the relationship found in method 1 and the relationship found in method 2.

Frequency (μHz)	Width Found Using Method 1 (μHz)	Width Found Using Method 2 (μHz)
1407.627	0.091 ± 0.008	0.07 ± 0.01
1263.524	0.064 ± 0.006	0.039 ± 0.008
1118.15	0.044 ± 0.005	0.021 ± 0.005
972.745	0.031 ± 0.004	0.010 ± 0.003
825.365	0.021 ± 0.003	0.0044 ± 0.002

l. The bottom panel of Figure 1 shows the results for $l = 0$ modes only. The power is calculated from equation 2 using the width and the height of the best fit Lorentzian for the mode. The width and height of a mode have a negative correlation of ~ 0.95 and so the resulting error bars on the calculated powers are small.

For both the power and the width a linear, least-squares fit was performed to determine the gradients and the zero-frequency intercepts. The gradients and zero-frequency intercepts of these fits and the errors associated with them can be seen in Table 1. This method of fitting the data will be referred to as method 1.

2.3 Method 2: The ln-ln relationship

An approximately linear relationship is also observed between the natural logarithm of the width of a mode and the natural logarithm of its frequency (see the top panel of Figure 2). Once again the linearity is only found at frequencies below $\sim 2400\,\mu$Hz. Likewise an approximately linear relationship is found when the natural logarithm of the power of a mode is plotted against the natural logarithm of the frequency of that mode (see the bottom panel of Figure 2). As with method 1 a linear least squares fit was performed to find the gradients and the zero-frequency intercepts for both the width and the power. The gradients and zero-frequency intercepts resulting from fitting the powers and widths in this manner can be seen in Table 1. This approach to fitting the data will be referred to as method 2.

2.4 Extrapolating method 1 and method 2 to low frequencies

To determine the value of the width and power of a mode at very low frequencies the gradients and zero-frequency intercepts were used to extrapolate the linear relationships from method 1 and method 2 to lower frequencies. The widths found using both method 1 and method 2 for $l = 0$ modes, in the frequency range 800 to 1450 μHz, can be seen in Table 2; and the powers for the same modes, also found using both method 1 and method 2, can be seen in Table 3. Some of these modes have not yet been observed and so we have used the mode frequencies predicted by the Saclay Seismic Model (Turck-Chièze et al. 2001).

Table 1   Gradients and zero-frequency intercepts found by least squares fits for method 1 and method 2. The reduced χ_ν^2 value for each fit is also given.

	Gradient	Zero-frequency Intercept	χ_ν^2
Width			
Method 1	$(2.5 \pm 0.1) \times 10^{-3}$	-6.0 ± 0.2	0.91
Method 2	$(5.2 \pm 0.2) \times 10^{0}$	-40 ± 2	0.70
Power			
Method 1	$(4.68 \pm 0.06) \times 10^{-3}$	-17.2 ± 0.1	6.5
Method 2	$(9.8 \pm 0.1) \times 10^{0}$	-82.3 ± 0.9	6.9

Table 3   Predicted powers of various $l = 0$ modes found by extrapolating the relationship found in method 1 and the relationship found in method 2.

Frequency (μHz)	Power Found Using Method 1 (ms^{-2})	Power Found Using Method 2 (ms^{-2})
1407.627	$(2.38 \pm 0.09) \times 10^{-5}$	$(1.32 \pm 0.08) \times 10^{-5}$
1263.524	$(1.21 \pm 0.05) \times 10^{-5}$	$(0.46 \pm 0.03) \times 10^{-5}$
1118.15	$(0.61 \pm 0.03) \times 10^{-5}$	$(0.14 \pm 0.01) \times 10^{-5}$
972.745	$(0.31 \pm 0.02) \times 10^{-5}$	$(0.035 \pm 0.002) \times 10^{-5}$
825.365	$(0.16 \pm 0.01) \times 10^{-5}$	$(0.0078 \pm 0.0008) \times 10^{-5}$

As can be seen the widths extrapolated using method 2 are consistently smaller than the widths extrapolated using method 1. The difference between the estimated width of a particular mode increases as the frequency of the mode decreases. In fact the widths inferred by the two methods do not agree to within their associated error bars at any frequency. Below $1000\,\mu$Hz the widths predicted by method 1 are ~ 3 times larger than the widths predicted by method 2. Additionally method 2 consistently predicts lower powers than method 1. Again the difference between the two extrapolations increases as the frequency of the mode decreases and at no time do the powers predicted by the two methods agree to within their respective error bars. At the higher frequencies the powers predicted by method 2 are just under half the powers predicted by the method 1. However, below $1000\,\mu$Hz the powers predicted by method 1 are approximately a factor of 10 larger than the powers predicted by method 2. Clearly there is a significant discrepancy between the powers and widths predicted by the two methods.

Table 1 gives the reduced χ_ν^2 values of the linear fits. The lower the reduced χ_ν^2 value the better the fit is at representing the data. However, if the reduced χ_ν^2 value is significantly less than unity either the model used to fit the data is too complicated or the errors on the data have been overestimated. The χ_ν^2 values are very similar for each method with method 2 providing a slightly better fit for the widths but method 1 giving a better representation of the powers. Both methods appear to give very good fits to the width. The values of the reduced χ_ν^2 determined for the linear power fits are large because of the small error bars associated with the observed powers.

Monte Carlo simulations were performed using the results of both extrapolations. These simulations will be described in Section 3. However, before we describe the sim-

ulations it is interesting to compare the extrapolated parameters with theoretical predictions that can be found using a stochastic excitation model.

2.5 Comparing the results of the extrapolation with solar excitation model predictions

Theoretical model mode damping and excitation rates can be calculated using various analytical models, which describe the interaction of the convection and the oscillations. The power of a mode, V^2, is given by

$$V^2 = \frac{P}{2\pi \Delta I} \tag{3}$$

where P is the energy supply rate, I is the mode inertia and Δ is, as defined previously, the width of the mode in a frequency-power spectrum. In what follows we will compare the powers and linewidths that have been predicted by one such stochastic excitation model with the linewidths and powers that are observed in and can be extrapolated from BiSON data. Chaplin et al. (2005) calculated model energy supply rates and linewidths using the Cambridge stochastic excitation and damping codes. Figure 3 shows that the model widths drop off more rapidly than both the observed and extrapolated widths at low frequencies. This is a known, and as yet unresolved problem, with the modelling of the linewidths (see for example Chaplin et al. 2005).

The top panel of Figure 4 shows a comparison between the model, observed and extrapolated powers. As we are interested in the frequency dependence of the power the model powers have been scaled to ensure that the maximum model power and the maximum observed power are equal. The model powers decrease less rapidly than the powers observed in the BiSON data. Therefore the powers estimated by both method 1 and method 2 are smaller than the model powers at low frequencies. However, this is understandable

Fig. 2 Results for method 2. Top panel: The black crosses show how the natural logarithm of the observed width of a mode varies with the natural logarithm of its frequency. The errors bars shown for the widths are those associated with the fits. Bottom panel: The black crosses show how the natural logarithm of the mode power varies with the natural logarithm of mode frequency. As in Figure 1 the red solid lines in both panels represent the linear best fits and the blue dashed lines represent the errors on the fits.

Fig. 4 Top panel: A comparison between the observed, modelled and extrapolated powers. The solid line represents the modelled powers, the black crosses represent the powers found by fitting BiSON data, the red diamonds show the powers found using extrapolation method 1 and the blue triangles show the powers found using extrapolation method 2. Bottom panel: A comparison between the observed, modelled and extrapolated powers, however the model powers have been found using the observed widths. The symbols and lines have the same definitions as in the top panel. Also plotted are the model powers found using the widths estimated by extrapolation method 1 and the model powers calculated using the widths extrapolated using method 2.

Fig. 3 A comparison between observed, modelled and extrapolated linewidths. The solid line represents the modelled linewidths, the black crosses represent the linewidths found by fitting BiSON data, the red diamonds show the widths found using extrapolation method 1 and the blue triangles show the widths found using extrapolation method 2.

and model powers seen in the top panel of Figure 4 is due to the smaller model widths. The extrapolated powers also appear to agree better with the model values calculated using the observed widths (bottom panel of Figure 4). However, it is difficult to tell which extrapolation method produces the best agreement with the model powers. Also plotted in the bottom panel of Figure 4 are the model powers that are found when the widths estimated by both extrapolation methods are used. The agreement between the extrapolated powers and the model powers is poor when method 1 is used. However, the agreement between the extrapolated and model powers is better when method 2 is used.

3 Simulating the modes

In this section we describe how we made artificial p-mode timeseries. These timeseries were analyzed, and results on detections of their artificial modes compared with the results of Broomhall et al. (2007), on real p-mode data. In Broomhall et al. BiSON and GOLF frequency-amplitude spectra were searched for coincident prominent structures that occurred at the same frequency in each spectrum. Here

as the model widths decrease more rapidly than the observed widths. Chaplin et al. (2005) note that the model widths are too narrow at low frequencies, which implies that the extrapolations performed tend in the right direction.

It is also of interest to find the model powers using the observed widths. The bottom panel Figure 4 shows that using the observed widths produces a far better agreement with the powers observed in the BiSON data. This indicates that the majority of the discrepancy between the observed

239

we needed to simulate pairs of timeseries, which both contained the signal from a mode and normally distributed noise. The statistical tests described in Broomhall et al. (2007) were then applied to the simulated data to determine how often the simulated modes could be detected.

Modes of a given frequency and lifetime (width) were simulated by randomly exciting an oscillator that was damped over the correct timescale. The damping time was predicted using equation 1 and the extrapolations described in Section 2 for linewidth. The simulations produced timeseries that contained the signal from a simulated mode. The total power of the simulated mode was scaled to the power predicted by the simulations.

Two timeseries containing normally distributed random noise were created. The signal from two sets of contemporaneous real Sun-as-a-star data taken by different instruments, such as BiSON and GOLF, contains some coherent noise. This noise is solar in origin and comes from the solar granulation. The level of coherent noise is frequency dependent and at around $1000\,\mu$Hz the coherency between BiSON and GOLF data is ~ 0.1. Therefore 10% of the noise in the two simulated timeseries was set so it was common to both sets of simulated data. The timeseries containing the signal from the simulated mode was then added to each noise timeseries.

Solar modes are excited stochastically by turbulence in the convection zone, which is caused by the solar granulation. As the coherent noise in two sets of Sun-as-a-star data is from the granulation it is possible that the coherent noise found between the BiSON and GOLF data is also correlated to the amplitude of excitation of the mode. In the simulations the amplitude of the excitation was determined by a normally distributed array. For some of the simulations the correlated noise that was added to the simulated data was taken to be the array that determined the amplitude of the excitation of the mode. Simulations were also performed when the noise that was coherent between the two simulated sets of data was independent of the amplitude of excitation of the mode.

The simulations have been performed for the $l = 0$ modes between ~ 800 and $\sim 1450\,\mu$Hz, the extrapolated properties of which are given in Tables 2 and 3. The timeseries were simulated to contain 2,211,120 points with a cadence of 120s to be consistent with the sets of BiSON and GOLF data used in Broomhall et al. (2007). The simulated and BiSON timeseries cover the same length in time (~ 8.5 yrs), despite having different cadences and containing a different number of points. This means that the power of a mode in the simulated and observed timeseries is the same. Furthermore, the number of frequency bins that cover the width of a mode in the simulated frequency-power spectra is the same as in the fitted BiSON spectrum. Therefore, the height of the mode in a simulated frequency-power spectrum will be the same as the height predicted by the extrapolations. The level of noise in the simulated spectra was scaled to mimic the mean level of noise observed in the Bi-

SON frequency-power spectrum in the $100\,\mu$Hz surrounding the frequency of the mode that was simulated.

Each extrapolated power and width has an error associated with it. To obtain an upper limit on the fraction of times a simulated mode could be detected Monte Carlo simulations were performed using the $\pm 1\sigma$ error values of the extrapolated widths and powers.

Once timeseries containing a mode and noise had been created several tests were performed on the resulting frequency-amplitude spectra to determine whether the mode could be detected or not. For each mode 1000 pairs of timeseries were simulated. Each pair of simulated timeseries corresponds to a different realization. The number of pairs of spectra in which a mode was detected was then counted. Various statistical tests were used to detect the modes. We will now outline each of these tests in turn.

4 Statistical Tests

The tests performed are based on the statistics described in detail in Chaplin et al. (2002) and Broomhall et al. (2007). Here we summarize the statistical tests that were used to search the simulated spectra. The simplest test involves searching for a single prominent spike that is above a given threshold level in the same bin in each of the two simulated frequency-amplitude spectra. The threshold level for detection was set at a 1 per cent chance of getting at least one false detection anywhere in $100\,\mu$Hz. This level and frequency range were chosen to maintain consistency with the detection methods employed in Broomhall et al. (2007). The probability of observing such a 'pair of spikes' takes proper account of the level of common noise present in the two frequency-amplitude spectra as this affects the probability that any detection is due to noise. A detection was considered to have been made if it was positioned within one linewidth of the input frequency. The linewidths were determined by the extrapolated values and so varied between modes. As the modes are damped it is possible to take advantage of the width the mode may exhibit in frequency. Furthermore, if the frequency of a mode is not commensurate with the frequency bins of the spectrum the power of the mode may be spread over more than one bin. The second test involves searching for two prominent spikes in the same consecutive frequency bins of each frequency-amplitude spectrum. Both spikes must lie within one peak width of the mode's input frequency for a detection to be counted. The third test performed also searched for two prominent spikes but this time the spikes did not need to be in consecutive bins. These two spikes are known as a two-spike cluster. The bins occupied by the two spikes could lie up to twice the predicted peak width apart. However each spike had to lie less than one linewidth from the mode's input frequency for a detection to be counted. The two prominent spikes in the cluster needed to be in the same bins in each frequency-amplitude spectrum. This test was then extended to search for clusters containing three, four and five prominent spikes. The

Fig. 5 Comparison between the fraction of modes detected when different input parameters were used.

highest and lowest frequency spikes in the cluster had to be separated by less than twice the width of the mode. Furthermore, each prominent spike in the cluster had to be in the same frequency bin in each frequency-amplitude spectrum. All spikes in the clusters needed to lie within one linewidth of the mode's input frequency for a detection to be counted.

The simulated spectra were searched to determine how often these statistical tests were passed for each mode. The results of these tests will now be described.

5 Results of the Simulations

Figure 5 shows the fraction of the 1000 simulations that were performed in which the mode being simulated was detected when looking for a single spike in the same bin in each spectrum. The results show that using the parameters predicted by method 1 leads to more detections than the parameters estimated by method 2. This is expected as the powers predicted by method 1 are significantly larger than the powers predicted by method 2. More detections are made when the noise that is coherent between the two sets of data is also correlated to the amplitude of the excitation of the mode. This is also understandable. If, for example, at a point in time the level of solar noise is larger than average the amplitude of the excitation of the mode will also be larger than average enabling it to potentially remain distinguishable from the noise. However, all of the simulations that use the extrapolated values predict that the number of detections at low frequencies will be small. Also plotted on Figure 5 are the results of simulations performed using the model widths and powers predicted by the stochastic excitation model described in Chaplin et al. (2005). The model powers have been found using the model widths rather than the observed widths. Clearly more detections are made when the model parameters are used.

Figure 6 plots results from the spike tests where the coherent noise was correlated to the mode excitation. Each panel of Figure 6 shows three clusters of lines. The solid, black lines in the clusters labelled sf=1 represent the results from simulations where the power and width of the

input mode were predicted by the respective extrapolation methods. The results show that the number of detections decreases rapidly as we move to lower frequencies. As each of the statistical tests are looking for prominent spikes it is the height the mode exhibits in the frequency-amplitude/power spectra that determines whether or not it can be detected. The height of a mode in a frequency-power spectrum is proportional to V^2/Δ (see equation 2). The upper limit on the number of detections is given by the case when the maximum power and minimum width is used. Using the maximum power and minimum width means that the height of the mode, H, is, potentially, larger and so the signal-to-noise ratio of the mode should be increased. It is therefore understandable that using the maximum power and minimum width leads to more detections. When the simulations were performed using the maximum width the power is spread over a larger number of bins and so this negates the effect on the signal-to-noise ratio of the mode of increasing its power.

Simulations were also performed to investigate the effect of increasing the observed signal-to-noise ratio. The power given to each simulated mode was increased by factors of 2 (labelled sf=2 in each panel of Figure 6) and 3 (labelled sf=3 in each panel of Figure 6). This significantly increases the number of detections made at low frequencies. The fraction of spectra in which the $l = 0$, $n = 6$ mode was detected when the scale factor was 1 was at most 0.05. However, the fraction of detections made when the scale factor was 2 was 0.60 and when the scale factor was 3 the fraction of simulated spectra in which the $l = 0$, $n = 6$ mode was detected was 0.98.

6 Discussion

Clearly the visibility of the modes drops off at low frequencies. The most optimistic results are achieved when the extrapolation is performed using method 1 and the common noise is correlated to the amplitude of the excitation of the mode. Clearly more detections are made when the stochastic excitation model parameters are used. However, these results can only be treated as an upper limit as we have already seen that the model underestimates the widths of the modes. This not only leads to the power of the modes being overestimated but also means that the height of a mode is increased, and so a mode is more likely to be detected.

For both extrapolation methods, irrespective of whether the noise is correlated or uncorrelated with the excitation amplitude, the fraction of detections at $\sim 1000\,\mu$Hz is less than 0.1. On the assumption that method 1 produces a more accurate prediction of the mode's powers and widths than method 2 the simulations imply that if the power signal-to-noise ratio can be increased by a factor of 2 a significantly higher proportion of modes could be detected. It would therefore be beneficial to continue efforts to try and improve the quality of the data further.

Fig. 6 Results of searching for a single spike when the noise is correlated with the excitation amplitude. The black solid lines show the results for the actual extrapolated values. The red dashed lines show the results when the maximum power and minimum widths are used in the simulations. The blue dashed line shows the results of the simulations that used the maximum powers and the maximum widths. The three clusters of lines are present as for each set of parameters the power was scaled by a different amount. The clusters labelled sf=1 represent the true predicted values. The clusters labelled sf=2 show the results when the power was increased by a factor of 2 and the clusters labelled sf=3 are the results when the power was increased by a factor of 3.

The total background continuum is a combination of instrumental noise, solar noise and, for Earth-based instruments such as the BiSON network, a small amount of atmospheric noise. If the atmospheric and instrumental noise can be reduced this would increase the signal-to-noise ratio of a mode, making it easier to detect. The limiting factor to the improvement that can be made to the data quality is the amount of solar noise present. However, the mean power of the BiSON data is always more than twice the power of solar noise predicted by the Harvey model and the characteristic parameters for the granulation found by Elsworth et al. (1994). Therefore it is theoretically possible to significantly reduce the level of noise in this frequency range.

It is possible to compare the results of these simulations to the observed candidates found in Broomhall et al. (2007). According to the simulations the $l = 0$, $n = 9$ mode at $\sim 1407\,\mu$Hz should be detected in the vast majority of spectra and this mode is observed to be very prominent in the BiSON and GOLF data used in Broomhall et al.. The $l = 0$, $n = 8$ mode at $\sim 1263\,\mu$Hz is also detected by Broomhall et al.. The simulations indicate that the fraction of time that this mode should be detected is at least 0.6. The $l = 0$, $n = 7$ mode at $\sim 1118\,\mu$Hz and the $l = 0$, $n = 5$ mode at $\sim 825\,\mu$Hz are not detected in the BiSON and GOLF data used in Broomhall et al.. The fraction of simulated spectra in which the $l = 0$, $n = 7$ mode was detected was, at most, 0.35, while the fraction of simulations in which the $l = 0$, $n = 5$ mode was detected was less than 0.06. Therefore it is reasonable to expect that the modes might not be detected in the BiSON and GOLF data.

It should of course be noted here that all of these results are based on the assumption that the simple relationships for the mode powers and widths observed in well-defined modes are still valid at low frequencies. The solar excitation model suggests that this may not be an unreasonable assumption, although it is difficult to determine which ex-

trapolation method predicts the lifetimes and powers more accurately.

The $l = 0$, $n = 6$ mode is significantly more prominent in the observations than the simulations. Therefore this mode will now be considered in more detail.

6.1 The $l = 0$, $n = 6$ mode

The $l = 0$, $n = 6$ mode is observed to be prominent in Sun-as-a-star data. However, some modes with higher frequencies remain undetected. The most optimistic simulations imply that the fraction of spectra in which the $l = 0$, $n = 6$ mode should be detected is less than 0.05. In the simulations performed here, for method 1, with noise correlated to the amplitude of the excitation, the average simulated signal-to-noise ratio in amplitude was ~ 2.1. However, the observed amplitude signal-to-noise ratio in the GOLF and BiSON data examined in Broomhall et al. (2007) was greater than 3.4. In fact, the fraction of simulations in which the signal-to-noise ratio of the mode was greater than 3.4 was 0.014. Therefore the simulations poorly represent the observed properties of this mode. It should be noted that the amplitude signal-to-noise ratio of other modes that were detected in Broomhall et al. are well represented by the simulations in this paper. For example, the $l = 0$, $n = 9$ mode was observed to have a signal-to-noise ratio of 3.1 while the average signal-to-noise ratio produced by the simulations ranged from 2.9 to 3.7 depending on which extrapolation method was used and whether the noise was correlated to the excitation amplitude.

When the mean amplitude signal-to-noise ratio of the simulated $l = 0$, $n = 6$ mode was set as 3.4 the fraction of spectra in which the mode was detected increased to ~ 0.6. A signal-to-noise ratio of 3.4 is $\sim 3\sigma$ from the mean amplitude signal-to-noise found in the original simulations where extrapolation method 1 (ln-linear) is used and the noise is correlated to the excitation amplitude. Therefore, the simulations and the observations are not so different that new

physics is required to explain the observations. It is more likely that at some point in time the mode has been randomly excited to a larger amplitude because of the stochastic nature of mode excitation.

Acknowledgements

References

Appourchaux, T., Fröhlich, C., Andersen, B., Berthomieu, G., Chaplin, W. J., Elsworth, Y., Finsterle, W., Gough, D. O., et al.: 2001, ApJ 538, 401

Broomhall, A. M., Chaplin, W. J., Elsworth, Y., Appourchaux, T.: 2007, MNRAS 379, 2

Chaplin, W.J., Elsworth, Y., Isaak, G.R., Marchenkov, K.I., Miller, B.A., New, R., Pinter, B.: 2002, MNRAS 336, 979

Chaplin, W. J., Houdek, G., Elsworth, Y., Gough, D. O., Isaak, G. R., New, R.: 2005, MNRAS 360, 859

Elsworth, Y., Howe, R., Isaak, G.R., McLeod, C.P., Miller, B.A., New, R., Speake, C.C., Wheeler, S.J.: 1994, MNRAS 269, 529

Fletcher, S. T.: 2007, PhD thesis, University of Birmingham

García, R. A., Régulo, C., Turck-Chièze, S., Bertello, L., Kosovichev, A. G., Brun, A. S., Couvidat, S., Henney, C. J., et al.: 2001, Sol. Phys. 200, 361

Houdek, G., Balmforth, N. J., Christensen-Dalsgaard, J., Gough, D. O.: 1999, A&A 351, 582

Turck-Chièze, S., Couvidat, S., Kosovichev, A. G., Gabriel, A. H., Berthomieu, G., Brun, A. S., Christensen-Dalsgaard, J., García, R. A., et al.: 2001, ApJ 555, L69

BIBLIOGRAPHY

Abramowitz, M. and Stegun, I. A. (1972). *Handbook of Mathematical Functions*. Handbook of Mathematical Functions, New York: Dover, 1972.

Allen, C. W. (1955). *Astrophysical quantities*. [London] University of London, Athlone Press, 1955.

Anderson, E., Duvall, Jr, T., and Jefferies, S. (1990). Modeling of solar oscillation power spectra. *ApJ*, 364:699–705.

Asplund, M., Grevesse, N., and Sauval, A. J. (2005). The Solar Chemical Composition. In Barnes, III, T. G. and Bash, F. N., editors, *Cosmic Abundances as Records of Stellar Evolution and Nucleosynthesis*, volume 336 of *Astronomical Society of the Pacific Conference Series*, pages 25–+.

Basu, S. and Antia, H. M. (2004). Constraining Solar Abundances Using Helioseismology. *ApJL*, 606:L85–L88.

Basu, S., Chaplin, W. J., Elsworth, Y., New, R., Serenelli, A. M., and Verner, G. A. (2007). Solar Abundances and Helioseismology: Fine-Structure Spacings and Separation Ratios of Low-Degree p-Modes. *ApJ*, 655:660–671.

Baudin, F., Samadi, R., Goupil, M.-J., Appourchaux, T., Barban, C., Boumier, P., Chaplin, W. J., and Gouttebroze, P. (2005). Inferred acoustic rates of solar p modes from several helioseismic instruments. *A&A*, 433:349–356.

Bertello, L., Henney, C. J., Ulrich, R. K., Varadi, F., Kosovichev, A. G., Scherrer, P. H., Cortés, T. R., Thiery, S., Boumier, P., Gabriel, A. H., and Turck-Chièze, S. (2000).

Comparison of Frequencies and Rotational Splittings of Solar Acoustic Modes of Low Angular Degree from Simultaneous MDI and GOLF Observations. *ApJ*, 535:1066–1077.

Brookes, J., Isaak, G., Mcleod, C., van der Raay, H., and Roca Cortes, T. (1978a). The search for solar oscillations, 1974 to 1976. *MNRAS*, 184:759–767.

Brookes, J., Isaak, G., and van der Raay, H. (1976). Observations of free oscillations of the sun. *Natur*, 259:92.

Brookes, J., Isaak, G., and van der Raay, H. (1978b). The observation of a rotating body using high resolution spectroscopy. *MNRAS*, 185:19–22.

Broomhall, A. M., Chaplin, W. J., Elsworth, Y., and Appourchaux, T. (2007). Needles in haystacks: how to use contemporaneous data in the search for low-frequency modes of oscillation of the Sun. *MNRAS*, 379:2–10.

Chaplin, W., Elsworth, Y., Isaak, G., Marchenkov, K., Miller, B., and New, R. (2001). Rigid rotation of the solar core? on the reliable extraction of low-l rotational p-mode splittings from full-disc observations of the sun. *MNRAS*, 327:1127.

Chaplin, W., Elsworth, Y., Isaak, G., Marchenkov, K., Miller, B., New, R., and Pinter, B. (2002a). Peak findings at low signal-to-noise ratio: low-l solar acoustic eigenmodes at $n \le 9$ from the analysis of bison data. *MNRAS*, 336:979–991.

Chaplin, W., Elsworth, Y., Isaak, G., Miller, B., and New, R. (2002b). On the measurement precision of solar p-mode eigenfrequencies. *MNRAS*, 330:731–736.

Chaplin, W. J., Appourchaux, T., Baudin, F., Boumier, P., Elsworth, Y., Fletcher, S. T., Fossat, E., García, R. A., Isaak, G. R., Jiménez, A., Jiménez-Reyes, S. J., Lazrek, M., Lochard, J., New, R., Pallé, P., Régulo, C., Salabert, D., and Toutain, T. (2004a). Flag Hare-And Exercise: on the Extraction of Sectoral Mode Splittings from Full-Disc Sun-As Data. In *ESA SP-559: SOHO 14 Helio- and Asteroseismology: Towards a Golden Future*, page 356.

Chaplin, W. J., Appourchaux, T., Elsworth, Y., Isaak, G. R., Miller, B. A., and New, R. (2004b). On comparing estimates of low-l solar p-mode frequencies from Sun-as-a-star and resolved observations. *A&A*, 424:713–717.

Chaplin, W. J., Elsworth, Y., Howe, R., Isaak, G. R., McLeod, C. P., Miller, B. A., and New, R. (1997). 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.

Chaplin, W. J., Elsworth, Y., Isaak, G. R., Marchenkov, K. I., Miller, B. A., and New, R. (2003a). High-frequency interference peaks in BiSON data. In Sawaya-Lacoste, H., editor, *GONG+ 2002. Local and Global Helioseismology: the Present and Future*, volume 517 of *ESA Special Publication*, pages 247–250.

Chaplin, W. J., Elsworth, Y., Isaak, G. R., Miller, B. A., New, R., Thiery, S., Boumier, P., and Gabriel, A. H. (2003b). Observation of, and temporal variations in, solar p-mode multiplet frequency asymmetries at $l = 2$. *MNRAS*, 343:343–352.

Chaplin, W. J., Houdek, G., Elsworth, Y., Gough, D. O., Isaak, G. R., and New, R. (2005). On model predictions of the power spectral density of radial solar p modes. *MNRAS*, 360:859–868.

Chaplin, W. J., Serenelli, A. M., Basu, S., Elsworth, Y., New, R., and Verner, G. A. (2007). Solar heavy element abundance: constraints from frequency separation ratios of low-degree p modes. *ArXiv e-prints*, 705.

Christensen-Dalsgaard, J. (2003). *Lecture Notes of Stellar Oscillations*, chapter Excitation and Damping of the Oscillations. Aarhus University, 5th edition.

Christensen-Dalsgaard, J. and Berthomieu, G. (1991). *Theory of solar oscillations*, pages 401–478. Solar interior and atmosphere (A92-36201 14-92). Tucson, AZ, University of Arizona Press, 1991, p. 401-478. Research supported by SNFO and CNRS.

Christensen-Dalsgaard, J. and Gough, D. O. (1980). Is the sun helium-deficient. *Natur*, 288:544–547.

Christensen-Dalsgaard, J., Gough, D. O., and Thompson, M. J. (1991). The depth of the solar convection zone. *ApJ*, 378:413–437.

Claverie, A., Isaak, G. R., McLeod, C. P., van der Raay, H. B., and Cortes, T. R. (1979). Solar structure from global studies of the 5-minute oscillation. *Natur*, 282:591–594.

Clette, F. (1993). Properties of the atmospheric noise in full-disk photometric observations of solar oscillations - Implications for asteroseismology from the ground. *A&A*, 267:577–585.

Davis, Jr., R., Harmer, D., and Hoffman, K. (1968). *Frontiers in Physics, Volume 92: Solar Neutrinos - The First Thirty Years*, chapter Search for neutrinos from the Sun, pages 136–140. Addison-Wesley.

Deubner, F.-L. (1975). Observations of low wavenumber nonradial eigenmodes of the sun. *A&A*, 44:371–375.

Duvall, Jr., T. L., Harvey, J. W., Jefferies, S. M., and Pomerantz, M. A. (1991). Measurements of high-frequency solar oscillation modes. *ApJ*, 373:308–316.

Elsworth, Y., Howe, R., Isaak, G., McLeod, C., Miller, B., New, R., Speake, C., and Wheeler, S. (1994). Solar seismology - the velocity continuum spectrum. *MNRAS*, 269:529.

Elsworth, Y., Howe, R., Isaak, G., McLeod, C., Miller, B., New, R., and Wheeler, S. (1995). Techniques used in the analysis of solar oscillations data from the bison network. i. daily calibration. *A&AS*, 113.

Elsworth, Y., Howe, R., Isaak, G. R., McLeod, C. P., and New, R. (1990). Evidence from solar seismology against non-standard solar-core models. *Natur*, 347:536–539.

Fletcher, S. T. (2007). *New and novel techniques for determining the parameters of solar and stellar global oscillations.* PhD thesis, University of Birmingham.

Fletcher, S. T., Chaplin, W. J., Elsworth, Y., Schou, J., and Buzasi, D. (2006). Frequency, splitting, linewidth and amplitude estimates of low-l p modes of α Cen A: analysis of Wide-Field Infrared Explorer photometry. *MNRAS*, 371:935–944.

Fossat, E., Grec, G., and Harvey, J. W. (1981). Power Spectrum of Differential Refraction and Comparison with Solar Diameter Fluctuation Measurements. *A&A*, 94:95.

Fried, D. L. (1965). Statistics of a Geometric Representation of Wavefront Distortion. *Journal of the Optical Society of America (1917-1983)*, 55:1427–1435.

García, R. A., Bertello, L., Turck-Chièze, S., Couvidat, S., Gabriel, A. H., Henney, C. J., Régulo, C., Robillot, J. M., Roca Cortés, T., Ulrich, R. K., and Varadi, F. (2001a). Analysis of low frequency signal with the GOLF experiment: methodology and results. In Wilson, A. and Pallé, P. L., editors, *ESA SP-464: SOHO 10/GONG 2000 Workshop: Helio- and Asteroseismology at the Dawn of the Millennium*, volume 10, pages 473–478.

García, R. A., Palle, P. L., Turck-Chieze, S., Osaki, Y., Shibahashi, H., Jeffries, S. M., Boumier, P., Gabriel, A. H., Grec, G., Robillot, J. M., Roca Cortes, T., and Ulrich, R. K. (1998). High-Frequency Peaks in the Power Spectrum of Solar Velocity Observations from the GOLF Experiment. *ApJL*, 504:L51+.

García, R. A., Régulo, C., Turck-Chièze, S., Bertello, L., Kosovichev, A. G., Brun, A. S., Couvidat, S., Henney, C. J., Lazrek, M., Ulrich, R. K., and Varadi, F. (2001b). Low-Degree Low-Order Solar p Modes As Seen By GOLF On board SOHO. *SoPh*, 200:361–379.

García, R. A., Turck-Chièze, S., Boumier, P., Robillot, J. M., Bertello, L., Charra, J., Dzitko, H., Gabriel, A. H., Jiménez-Reyes, S. J., Pallé, P. L., Renaud, C., Roca Cortés, T., and Ulrich, R. K. (2005). Global solar Doppler velocity determination with the GOLF/SoHO instrument. *A&A*, 442:385–395.

Garcia, R. A., Turck-Chieze, S., Jimenez-Reyes, S. J., Ballot, J., Palle, P. L., Eff-Darwich, A., Mathur, S., and Provost, J. (2007). Tracking Solar Gravity Modes: The Dynamics of the Solar Core. *Science*, page 1140598.

Gough, D. O. (1977). Mixing-length theory for pulsating stars. *ApJ*, 214:196–213.

Gough, D. O. (1981). A new measure of the solar rotation. *MNRAS*, 196:731–745.

Grevesse, N., Asplund, M., and Sauval, A. J. (2007). The Solar Chemical Composition. *Space Science Reviews*, 130:105–114.

Grevesse, N. and Sauval, A. J. (1998). Standard Solar Composition. In Fröhlich, C., Huber, M. C. E., Solanki, S. K., and von Steiger, R., editors, *Solar Composition and Its Evolution – From Core to Corona*, pages 161–+.

Harvey, J. (1985). *Future Missions in Solar, Heliospheric and Space Plasma Physics*, chapter High Resolution Helioseismology, pages 199–208. ESA Publ. Division.

Harvey, J. W., Hill, F., Hubbard, R., 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., Patrón, J., Saá, O., and Yasukawa, E. (1996). The Global Oscillation Network Group (GONG) Project. *Science*, 272:1284–+.

Houdek, G., Balmforth, N. J., Christensen-Dalsgaard, J., and Gough, D. O. (1999). Amplitudes of stochastically excited oscillations in main-sequence stars. *A&A*, 351:582–596.

Houdek, G., Chaplin, B., and Elsworth, Y. (1998). How are solar oscillations damped? *Astronomy and Geophysics*, 39:30–+.

Houdek, G., Chaplin, W., Appourchaux, T., Christensen-Dalsgaard, J., Däppen, W., Elsworth, Y., Gough, D., Isaak, G., New, R., and Rabello-Soares, M. (2001). Changes in convective properties over the solar cycle: effect on p-mode damping rates. *MNRAS*, 327:483.

Hoyng, P. (1989). On the sensitivity of resonant scattering spectrometers for whole-disk solar velocity oscillation measurements. *ApJ*, 345:1088–1103.

Isaak, G. R., McLeod, C. P., Palle, P. L., van der Raay, H. B., and Roca Cortes, T. (1989). Solar oscillations as seen in the NaI and KI absorption lines. *A&A*, 208:297–302.

Jefferies, S. M., Pomerantz, M. A., Duvall, Jr., T. L., Harvey, J. W., and Jaksha, D. B. (1988). Helioseismology from the South Pole: Comparison of 1987 and 1981 results. In Rolfe, E. J., editor, *Seismology of the Sun and Sun-Like Stars*, volume 286 of *ESA Special Publication*, pages 279–284.

Kolmogorov, A. N. (1961). *Turbulence - Classis Papers on Statistical Theory*, page 151. Interscience Publishers, Inc., New York.

Leibacher, J. W. and Stein, R. F. (1971). A New Description of the Solar Five-Minute Oscillation. *ApL*, 7:191–192.

Leighton, R., Noyes, R., and Simon, G. (1962). Velocity fields in the solar atmosphere. i. preliminary report. *ApJ*, 135:474.

Libbrecht, K. G. (1988). Solar p-mode phenomenology. *ApJ*, 334:510–516.

Moreno-Insertis, F. and Solanki, S. K. (2000). Distribution of magnetic flux on the solar surface and low-degree p-modes. *MNRAS*, 313:411–422.

Palle, P. L., Regulo, C., Roca-Cortes, T., Sanchez-Duarte, L., and Schmider, F. X. (1992). Solar Radial Velocity and Oscillations as Measured by Sodium and Potassium Resonant Scattering Spectrometers. *A&A*, 254:348–+.

Provost, J., Berthomieu, G., and Morel, P. (2000). Low-frequency p- and g- mode solar oscillations. *A&A*, 353:775.

Rhodes, Jr., E. J. and Ulrich, R. K. (1977). The Use of p Mode Oscillations to Measure the Solar Rotation Rate. In *Bulletin of the American Astronomical Society*, volume 9 of *Bulletin of the American Astronomical Society*, pages 336–+.

Thompson, M. J. (1990). A new inversion of solar rotational splitting data. *SoPh*, 125:1-12.

Toutain, T., Appourchaux, T., Fröhlich, C., Kosovichev, A. G., Nigam, R., and Scherrer, P. H. (1998). Asymmetry and Frequencies of Low-Degree p-Modes and the Structure of the Sun's Core. *ApJ*, 506:L147-L150.

Turck-Chièze, S., Couvidat, S., Kosovichev, A. G., Gabriel, A. H., Berthomieu, G., Brun, A. S., Christensen-Dalsgaard, J., García, R. A., Gough, D. O., Provost, J., Roca-Cortes, T., Roxburgh, I. W., and Ulrich, R. K. (2001). Solar neutrino emission deduced from a seismic model. *ApJ*, 555:L69-L73.

Turck-Chièze, S., García, R. A., Couvidat, S., Ulrich, R. K., Bertello, L., Varadi, F., Kosovichev, A. G., Gabriel, A. H., Berthomieu, G., Brun, A. S., Lopes, I., Pallé, P., Provost, J., Robillot, J. M., and Roca Cortés, T. (2004). Looking for Gravity-Mode Multiplets with the GOLF Experiment aboard SOHO. *ApJ*, 604:455-468.

Ulrich, R. (1970). The five-minute oscillations on the solar surface. *ApJ*, 162:993.

Ulrich, R. and Rhodes, E. (1977). The sensitivity of nonradial p mode eigenfrequencies to solar envelope structure. *ApJ*, 218:521-529.

Underhill, C. (1993). An analysis of possible line-of-sight velocity effects in the power spectra of solar oscillations. Master's thesis, University of Birmingham.

Underhill, C. J. and Speake, C. C. (1996). Height Dependence of the Amplitudes of Solar Oscillations Viewed with a Potassium Resonance Scattering Spectrometer. *SoPh*, 163:231-240.

Young, A. T. (1974). Seeing: its Cause and Cure. *ApJ*, 189:587-604.

Zaatri, A., Provost, J., Berthomieu, G., Morel, P., and Corbard, T. (2007). Sensitivity of low degree oscillations to the change in solar abundances. *A&A*, 469:1145-1149.