CHAPTER 8

TECHNIQUES FOR OBSERVING SOLAR OSCILLATIONS

https://doi.org/10.1017/S007418090015853X Published online by Cambridge University Press

TECHNIQUES FOR OBSERVING SOLAR OSCILLATIONS

Timothy M. Brown
High Altitude Observatory/National Center for Atmospheric Research*
P.O. Box 3000
Boulder, CO 80307
U.S.A.

ABSTRACT. There is currently a wide range of methods for observing the properties of solar oscillations. It is generally true, however, that the means used to analyze the data are just as important to the final result as the instrumentation. I discuss first the instrumental techniques used to observe solar p-modes, with attention given both to low- and no-resolution systems, and to systems with spatial resolution. Then I describe the reduction techniques that are used to convert the raw observations into useful form.

1. INTRODUCTION

Solar p-mode oscillations were first observed by Leighton, *et al.* (1962), using photographically-generated Dopplergrams obtained with a spectroheliograph. There followed a long period of confusion, during which the physical origin of the oscillations was unknown, and, as a result, there was little agreement about which observing and data reduction methods were appropriate to study the oscillations. Deubner (1975) was the first to understand the importance of doing the best possible Fourier decomposition of the observed velocity field in space and time. This understanding led him to take observations over a substantial fraction of the area of the solar disk, for long time intervals -- observations that immediately validated the (by then 5 year old) idea that the oscillations resulted from sound waves trapped in a subphotospheric cavity (Ulrich 1970, Leibacher and Stein 1971). I resurrect this well-known tale simply to emphasize that p-mode observation does not rely on stable and sensitive instruments alone: useful results emerge only when the observing and analysis programs are properly attuned to the physical processes on the Sun.

The properties of solar p-modes are by now well enough understood that one can design observing programs to study them with a fair likelihood of success. There are three fundamental facets of the oscillations that are important in this regard. (1) The observable quantities (velocity, temperature, etc.) associated with a single oscillation mode vary at the solar surface according to

$$q = q_0 Y_l^m(\theta,\phi)\, e^{-i\omega t} \quad .$$

(2) The amplitude associated with any single oscillation mode is very small: a few cm s^{-1} for the velocity, or a few parts in 10^6 for the relative intensity. (3) There are very many (perhaps 10^7)

*The National Center for Atmospheric Research is sponsored by the National Science Foundation

J. Christensen-Dalsgaard and S. Frandsen (eds.), Advances in Helio- and Asteroseismology, 453–465.

distinct oscillation modes excited on the Sun, within the relatively narrow frequency band between roughly 2 and 5 mHz. Confusion between modes is therefore prone to be a problem.

In what follows, I will attempt to describe most of the instrumental and reduction techniques used in the study of solar p-modes, including both Doppler and intensity observations, and both spatially resolved and unresolved observations of the solar disk. I will discuss neither high spatial resolution spectrographic methods nor astrometric methods.

2. OBSERVATIONAL METHODS

2.1. No- and Low-Resolution Methods

No- and low-resolution techniques to study p-modes measure properties of light integrated over the entire solar disk, or measure the difference between Doppler shifts averaged over a few (usually two) large subsections of the disk. Because of the extreme spatial averaging involved in such observations, they are sensitive only to oscillations of low spherical harmonic degree ($l = 0$–5, typically, with observations that integrate over the full disk restricted to $l = 0$–3). Such observations contain relatively little information, yet they are very important, since they are currently the *only* practical means of examining oscillations with $l \leq 2$. These modes are crucial to understanding the Sun, because they are the only ones that probe the Sun's central regions. Moreover, the solar low-degree modes have provided both inspiration and potentially useful techniques for studies of solar-like oscillations on other stars.

No- and low-resolution instruments have so far been constructed in three distinct forms: atomic resonance devices, differential spectrographs, and irradiance sensors. All of these are illustrated (very schematically) in Fig. 1. Fig. 1a shows one form of *atomic resonance detector*. This device is similar to that described by Brookes *et al.* (1978), and to the ones used by Grec and Fossat (1977). It is essentially a Doppler detector, and works by alternately sampling the light scattered from the red and blue wings of a chosen resonance line in the solar spectrum. It allows sunlight to pass first through a prefilter to isolate the desired solar line, and then through a linear polarizer and a $\pm\lambda/4$ modulator. The effect of these optical components is to pass alternately right- and left-circularly polarized light to the resonant scattering cell, which is the heart of the apparatus. The scattering cell contains a small amount of sodium or potassium vapor, and is held in a longitudinal magnetic field. Zeeman splitting of the atomic energy levels then causes the two polarizations to scatter at different wavelengths. By properly choosing the magnitude of the magnetic field and the optical depth of the metallic vapor, one can arrange for the scattering to occur from two relatively narrow wavelength regions situated near the steepest part of the solar line profile. The difference in the amount of light scattered from the red and blue line wings (measured by the photomultiplier) then provides a measure of the Doppler shift of the solar line relative to the lab.

The *differential spectroscopic approach* is illustrated in Fig. 1b. This method has been used and described by Kotov *et al.* (1978), and by Scherrer (1979). In this technique, one images the Sun onto a mask, which encodes the light from the central region of the disk and that from the surrounding annulus into different forms (usually right and left circular polarizations). This encoding mask is followed by decoding optics, which allow one to chop between the two regions of the solar disk. The light then passes through a high-dispersion spectrograph, and the Doppler shift of the desired spectral line is determined in conventional fashion (perhaps using a Babcock magnetograph arrangement). By rapid chopping between the two regions of the disk, one can then measure the difference in their line of sight velocity with high precision, almost independent of slow drifts in the instrumental calibration.

Fig. 1c illustrates the method used for *irradiance measurements* with the SMM ACRIM instrument (Willson (1979), Woodard and Hudson (1983)). In this technique (an active cavity radiometer), one uses a shutter to alternately block and transmit the sunlight incident on a good

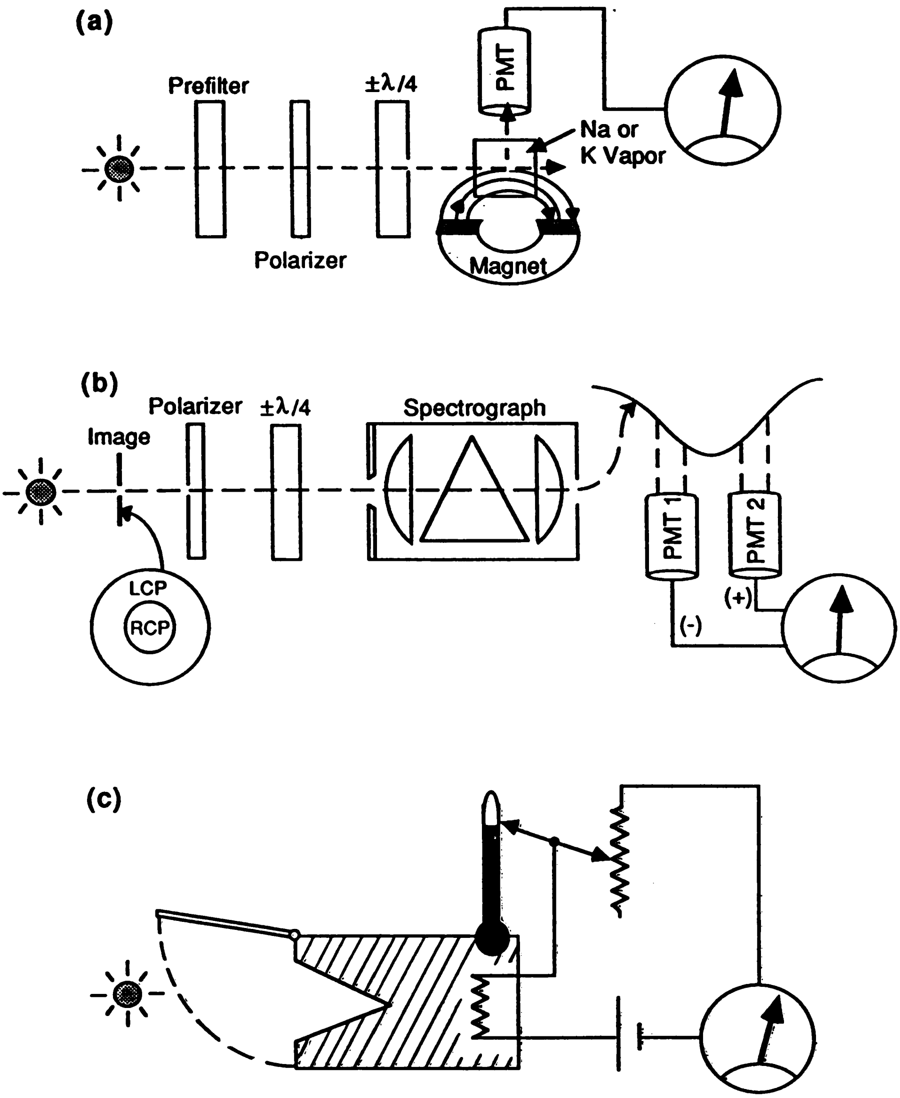

Figure 1. *Schematic illustrations of instruments for no- and low-resolution measurements of solar oscillations. (a) Atomic resonance method. (b) Differential spectroscopic method. (c) Active cavity radiometer.*

https://doi.org/10.1017/S007418090015853X Published online by Cambridge University Press

black-body absorber. When the shutter is open, the incident flux heats the cavity; when it is closed, an electric heater in a servo loop maintains the cavity at the same temperature as when the shutter is open. Measuring the current to the heater then provides an estimate of the flux incident on the absorber. This instrument is essentially a bolometer, measuring the temperature fluctuations resulting from the p-mode oscillations (as well as much larger variations from many other solar sources). Since the relative irradiance variations are only a few parts in 10^6, this kind of instrument can only be operated outside the Earth's atmosphere, where transparency variations do not occur.

The no- and low-resolution methods just described have a number of advantages and drawbacks, compared to other means for studying solar p-modes. On the positive side, it is worth repeating that, in spite of intensive development of various spatially-resolved methods in recent years, the integrated-light techniques remain the only ones that are effective for studying modes of low degree ($l < 3$, say). These methods are also very sensitive and very stable. For this reason they have been useful in examining the low- and high-frequency ends of the p-mode amplitude envelope, and they are well adapted to studies of slowly varying phenomena, such as the 160-minute oscillation. Finally, the atomic resonance devices are relatively small and rugged, which has allowed them to be used at the South Pole (Grec, *et al.* 1983) to obtain uninterrupted timestrings of several days' duration.

Unfortunately, all of the integrated light techniques are sensitive to variations in atmospheric transmission. This is obvious in the case of irradiance instruments. For the Doppler measurements, the sensitivity arises through the combined effects of solar rotation and spatial gradients in the atmospheric transmission. These methods measure a sort of intensity-weighted mean of the Doppler shift, averaged over the solar disk. Since the solar rotation imposes nearly 4000 ms^{-1} of shift between the east and west limbs, even a small transmission difference between different places on the solar disk can result in substantial apparent velocity signals. Another difficulty with these methods is that the measurements are ambiguous regarding the precise values of degree l and azimuthal order m that correspond to any particular power spectrum peak. Sometimes (as with the p-modes) the structure of the observed spectrum is clear enough that this ambiguity is not a problem. However, in other cases (notably reported detection of g-modes, and especially the 160-minute oscillation) lack of knowledge about the spatial structure responsible for the observed signal has significantly impeded progress.

2.2. Intermediate-l Methods

Measurements of intermediate-l oscillations ($3 \leq l \leq 200$, say) have relied principally on *Doppler measurements,* partly for instrumental reasons and partly because the intrinsic S/N ratio, set by solar processes, is better in the Doppler signal than in the intensity signal. To obtain useful information on the oscillations, one needs adequately resolved time series (typically one sample per minute) of spatially-resolved line position measurements. These have been obtained using instruments of two generic types: narrow band filters and Fourier tachometers.

There are several kinds of *narrow-band filter* currently in use to study solar oscillations: the atomic resonance filter (Cacciani and Fofi 1978, Rhodes *et al.* 1986), the birefringent filter (Libbrecht, 1986, Hill *et al.* 1986), and the Fabry-Perot (Rust *et al.* 1986). These employ different optical methods to obtain their narrow, stable bandpasses, but their application to solar oscillations is in all cases the same. First one chooses filter parameters so that the filter bandpass is narrower than the solar spectrum line to be observed. Then one tunes the filter so that the bandpass lies alternately on the red and blue wings of the line, and obtains spatially resolved images of the Sun in both positions. The Doppler shift of the solar line relative to the mean bandpass position may then be estimated from some measure of the difference between the two images, *e.g.*

$$\delta\lambda \approx \frac{I_{red} - I_{blue}}{I_{red} + I_{blue}} .$$

Since solar spectrum lines are typically rather narrow, even moderate lineshifts result in quite substantial intensity differences. The narrow-band filter methods are thus very sensitive. The way in which Doppler shifts are extracted from intensity measurements is both simple and easy to implement, so that narrow-band filters lend themselves to applications using large, rapidly-scanned detector arrays. On the other hand, the narrowness of the solar lines also means that the relation between intensity differences and velocities is not a linear one, and because of the large Doppler shift imposed by the solar rotation, one faces this difficulty in a rather severe form.

The *Fourier tachometer* (e.g. Brown, 1984, Evans 1980) accepts some degradation of S/N in order to obtain a device that is linear and stable. The device passes incoming sunlight through a prefilter, which isolates the solar spectrum line of interest. The light next passes through an interferometer, which has a transmission profile that is a sinusoidal function of frequency. The separation between transmission peaks is chosen to be roughly the width of the solar line; the interferometer path difference can be varied by a wavelength or so, in order to move the transmission peaks in wavelength by several fringes. With this optical arrangement and an appropriate demodulation scheme, one can measure the amplitude and phase of a chosen Fourier component of the line profile. The amplitude depends mainly on the line's strength, while the phase is a direct measure of the center frequency of the line.

This device is attractive because of its intrinsic linearity, and because it is fairly easy to calibrate, using a stable wavelength reference that need not be particularly near the wavelength of the solar line. In addition, its optical components can be made with large optical throughput, are fairly simple, and require no exotic materials. On the negative side, all of the light within the prefilter bandpass but outside the solar line profile contributes to the photon noise, but not to the signal. Since the prefilter bandpass must be wide compared to the solar line (or the instrument measures the prefilter wavelength and not the line wavelength), the instrumental noise tends to be larger for a Fourier tachometer than for instruments based on narrow-band filters. Also, the influence of the solar continuum appears in a number of systematic effects that bias the observed velocities, and that are difficult to calibrate using laboratory sources.

It is possible to detect p-modes using *intensity measurements* instead of measurements of line Doppler shifts. Such measurements tend to have higher noise than Doppler measurements, both because of competing solar processes with similar time and length scales (especially granulation), and because changes in atmospheric transmission appear directly in the intensity signal. The effects of atmospheric transmission have so far limited intensity measurements to degrees greater than 20. However, these disadvantages are balanced by the relative convenience cf measuring a simple intensity. Moreover, intensity fluctuations are nearly independent of distance from the solar limb, in contrast to the velocity signal; the latter goes to zero near the limb, because there the (mainly vertical) material motions associated with the p-modes are perpendicular to the line of sight. Thus, intensity measurements can be used over a larger fraction of the visible solar disk, improving one's ability to distinguish between modes with similar angular degree.

The most successful intensity observations have so far been obtained using comparatively narrow-band observations in temperature-sensitive parts of the solar spectrum (*e.g.* the Ca II K line and the CN bandhead) (Duvall *et al.* 1986, Kneer, Newkirk and Uexkull 1982). Collisionally-dominated lines formed near the solar temperature minimum seem particularly good for this purpose. In the best such spectral regions, one can obtain enhancement of about a factor of 5 in the relative intensity variation associated with the oscillations, compared to the variation seen in the continuum.

It has also proved possible to detect the p-modes in the continuum intensity, especially if the measurements are chiefly concerned with high-degree modes (Frazier 1968, Brown and Harrison 1980). Recent advances in detector technology have made such observations considerably more attractive (Nishikawa, *et al.* 1986), especially for spacecraft applications where atmospheric transmission variations are absent. All such observations remain limited by the solar convective background.

2.3. Observational Requirements

Assuming that one has an instrument that is adequate to observe solar p-modes, how must one use it to obtain useful results, and what limitations does one typically encounter? The principal requirements are that one must observe with time resolution that is adequate to sample the oscillations, and for as long a time span as possible. For a nominal period of 300 s, the signal would be critically sampled with images separated in time by 150 s. In practice this is inadequate, mainly because there is significant power in the solar p-modes up to frequencies of 5 mHz or more, corresponding to periods of 200 s or less. Some important work has been done with sampling intervals as long as 90 s (Duvall *et al.* 1986), but most workers prefer to sample no slower than once per 60 s.

Having sampled often enough, one must also sample long enough. The relevant frequency differences in p-mode studies tend to be very small: the differences between model mode frequencies and observed ones are about 10 μHz; the intrinsic linewidth of many modes appears to be less than 1 μHz; the frequency splitting between modes with adjacent m-values is about 0.4 μHz. Thus, the best frequency resolution obtainable is often none too good. Since the best possible frequency resolution is proportional to $1/T$, where T is the total length of the observing run, it is necessary to observe for long times, as continuously as possible. To get 1 μHz resolution, one must observe for about 12 days. Also, it is worth noting that improved resolution is advantageous even when the observing time is longer than mode lifetimes, since without such resolution one has difficulty estimating the width, shape, or central frequencies of the p-mode line envelopes.

Gaps in the observed data string (*e.g.* diurnal gaps) are objectionable because they introduce sidelobes in the frequency response. These sidelobes increase the effective noise, and lead to confusion between modes with different degrees (especially near certain degrees, where the frequency spacing between consecutive l-values is a multiple of the 11.6-μHz spacing due to diurnal sidelobes). There are several possible solutions to this problem. One may observe from places where diurnal gaps do not occur, such as the South Pole or from a satellite in an appropriate orbit. The former method has been practiced with success by Grec *et al.* (1983), and by Duvall *et al.* (1986). The latter has been the subject of much discussion (*e.g.* Noyes and Rhodes 1984), but has not yet produced a spaceborne instrument. Alternatively, one may construct a network of groundbased instruments, in hopes that at least one will be in sunlight at all (or almost all) times. At least two such networks for no-resolution observations are now in progress (see Fossat, these proceedings), and the GONG network promises similar advances for spatially resolved observations. Finally, one may attempt to overcome the effects of missing data by means of nontraditional analysis methods (*e.g.* Brown and Christensen-Dalsgaard 1986, Kuhn 1984). With or without such analysis methods, it remains unclear how large a duty cycle is required in order to avoid serious systematic errors in the power spectra. For example, current disagreements between South Pole data (Duvall *et al.* 1986) and low-latitude data (Brown 1985) regarding rotational splittings leaves some doubt as to the interpretation of data with a duty cycle of 0.3.

In addition to long data runs, one requires certain minimum seeing and transparency conditions. Transparency variations degrade data by two mechanisms. The first arises from the combination of solar velocity gradients with spatial gradients in the atmospheric transmission. Most Doppler instruments form, within each resolution element, what amounts to an intensity-

weighted mean of the line-of-sight velocity. When this velocity varies considerably over a resolution element (as is certainly the case for no- and low-resolution instruments), then small variations in the transmission from one side of a pixel to the other can result in noticeable spurious velocity signals. For example, in a no-resolution instrument, changing the transmission between the East and West limbs of the Sun by 0.1% would result in an apparent velocity change of about 4 ms^{-1}. Fortunately, noise from this source scales inversely as the square of the pixel width, and so is negligible for spatially resolved observations.

Transparency variations can also cause noise in Doppler instruments through their temporal fluctuations. All currently available Doppler analyzers derive their signals from differences of images obtained at different times. Transmission variations during this time interval can look like a velocity signal. Since the temporal spectrum of transmission variations is usually largest at low frequencies, it is advantageous to modulate rapidly.

Atmospheric image degradation (seeing) is also deleterious to p-mode observations. Seeing blurs the solar image, which lowers one's sensitivity to modes with high degree. This is bothersome both because it reduces the S/N ratio at high l, and because it interferes with attempts to measure mode amplitudes accurately. Seeing also moves the image about, on a wide range of time scales. These image motions combine with natural solar velocity gradients (especially supergranular velocity fields) to cause velocity noise. The motions can also combine with intensity gradients in the image to cause spurious modulation, in a fashion similar to the temporal transmission changes discussed above. Experience indicates that these seeing effects are largest near the solar limb, where both velocity and intensity gradients are large, but that their effects are probably not as dire as originally feared (Ulrich *et al.* 1984).

2.4. Observational Challenges

The instrumentation and observing techniques for studying p-modes have by now reached a fairly mature state, but a number of major challenges remain. The first of these is to lower the observational noise level still further, by improving the instrumentation, by devising better analysis methods, and by finding means to minimize the effect of solar noise (perhaps by better choices of spectrum lines). Driving the noise level as low as possible is important, both because it allows one to study modes in at higher and lower frequencies (in the tails of the p-mode amplitude envelope), and because low-noise observations are more susceptible to advanced data processing techniques, such as the gap-filling measures discussed earlier.

One can expect to see efforts to minimize the effects of the Earth's atmosphere on p-mode observations. This might be done, for instance, by correcting low-l observations for measured transmission variations, by increasing modulation speeds (or detecting all the modulation phases simultaneously, using different detectors), and by stabilizing images to compensate for seeing motion.

There will certainly be efforts to improve the temporal stability of spatially resolved observations, particularly for low degrees. This would be of great assistance in resolving questions about the visibility and properties of long-period oscillations, and particularly of g-modes. Making such improvements will require better calibration and stabilization techniques in order to eliminate diurnal drifts, or perhaps entirely new instruments with better intrinsic stability.

Finally, there will be more efforts to obtain longer, more nearly continuous observing runs, so that ambiguities resulting from diurnal gaps can be eliminated. This will require coordination between different observers, networks of observing stations, and the development of better and faster ways to reduce the very large amounts of data that will result.

3. DATA REDUCTION METHODS

As I emphasized in the Introduction, observing the Sun with the required precision, resolution, and spatio-temporal coverage is only half of the task required before one can make meaningful statements about solar oscillations. The other half is the reduction process, which converts the seemingly random observed velocity field into the ordered quantities that are needed for subsequent analysis. This process typically involves four major steps: preliminary reduction, spatial filtering, time series analysis, and spectrum interpretation.

3.1. Preliminary Reduction

By "preliminary reduction" I mean the process that begins with raw data, and produces standard maps of Doppler shift (or relative intensity variation), corrected for instrumental effects and with uninteresting solar phenomena removed. The details of this process tend to be very instrument-dependent, but they typically involve most of the following operations:

(1) Correction for gain and dark current variations in the instrument's detector. Different detectors and different demodulation schemes are more or less sensitive to these problems, but all suffer from them to some degree. Gain variations across the detector chip can be particularly intractable, since uniform flat-field sources are notoriously difficult to construct.

(2) Raw data to velocity conversion. This conversion takes different forms in different instruments. For example, narrow-band filter instruments usually define the velocity as a slightly nonlinear function of

$$\delta\nu = \frac{I_{red} - I_{blue}}{I_{red} + I_{blue}} \quad ,$$

whereas the Fourier Tachometer employs the arctangent of the ratio of intensity differences obtained at various modulation phases:

$$\delta\nu = V_0 \; arctan \frac{I_1 - I_3}{I_2 - I_4} \quad .$$

Whatever method is used, a moderate arithmetic capability is required to keep up with the flow of data.

(3) Image registration is necessary to compensate for image motions caused by telescope drive errors, pointing jitter, and atmospheric seeing. Such motions are detrimental for several reasons. Least serious (but still obnoxious) is the blurring of the image that results from superposing poorly-registered individual images. This blurring lowers the effective sensitivity to modes with high degree. A more serious effect of image motion is that it introduces noise into the measured velocities, via mechanisms that have already been discussed. Finally, if images are not properly registered and centered, then it becomes impossible to decompose them into spherical harmonics in a meaningful way. Image registration is usually done by determining an image center position from estimates of the limb position taken at many places around the solar circumference, and then using an interpolation technique to shift the digital image to a standard position.

(4) Once the images have been properly registered, it is desirable to remove any major sources of signal that are not relevant to the oscillations. Large, slowly changing signals put undue demands on the dynamic range of the various data processing steps, and tend to combine with minor defects in the analysis to produce large amounts of noise. For Doppler measurements, such signal sources include the Sun's differential rotation, the so-called limb red shift, and contributions to the line-of-sight velocity from the Earth's rotation and orbital motion. Fortunately, these phenomena are well understood (with the possible exception of the limb red shift), and can be subtracted from the data with good precision.

3.2. Spatial Filtering

Spatial filtering is the process of decomposing the observed velocity field into spherical harmonics. It is the most computation-intensive of the data reduction steps, so efficient algorithms are especially important. There are three techniques currently in use, or under consideration.

(1) The *direct multiplication* method is the most expensive of the methods, but also the most conceptually straightforward. To implement this technique, one first generates a number of masks, which are to be projected onto the observed data, one time step at a time. These masks are simply the desired Y_l^m functions (with due allowance for foreshortening of the geometrical pattern, and perhaps also for projection of the velocity field onto the line of sight), on the same spatial grid as is used in the observations. One then multiplies each mask onto the observed velocity or intensity image, and integrates the result over the area of the solar disk. The resulting numbers are taken to be the coefficients in a spherical harmonic expansion of the observed quantity. It is important to note that this identification is not entirely correct, since the Y_l^m are orthogonal on the sphere, but not on the distorted hemisphere of the Sun that is accessible to observation. This non-orthogonality tends to mix the coefficients of spherical harmonics with similar values of l and m – a fault that is common to all spatial filtering methods.

The direct multiplication method is easy to visualize and implement, but it is computationally expensive. If N is the number of pixels along one edge of a square image and M is the number of spherical harmonic coefficients desired, then the number of arithmetic operations required scales as N^2M. To preserve all of the information in the image, one must compute about as many coefficients as there are pixels, in which case the computational cost scales as N^4. Since this computation must be performed for each individual image in a time series, it becomes computationally prohibitive to compute all of the desired coefficients. For this reason, the direct multiplication method has been used mainly in studies involving only a restricted set of spherical harmonics of particular interest (*e.g.* Libbrecht and Zirin 1986, who quote results only for $l \leq 20$).

A more complicated (but also more efficient) technique is to decompose the spherical harmonic projection into two one-dimensional transforms. In this method, first described in Brown (1985), one first interpolates the observed velocity field onto a grid that is equally spaced in longitude and sine latitude. The $e^{im\phi}$ part of the spherical harmonic can then be separated from the Legendre transform, and implemented as an FFT. The Fourier transformed data array contains complex amplitudes as a function of m and of sine latitude. Finally, one performs 1-dimensional Legendre transforms (by performing a brute-force dot product with the appropriate Legendre functions) on each column individually. Efficient recursive techniques may be used to generate the Legendre functions as they are needed; this usually proves faster than computing all of the functions in advance and retrieving them from disk. Duvall *et al.* (1986) have improved this scheme by computing coefficients only for alternate values of m. This is permissible because the limited coverage in longitude (less than 180°) makes half of the coefficients redundant.

With the same notation as used above, the computational costs for this technique may be estimated as follows: The interpolation for each frame scales as N^2, the number of pixels in both the original and interpolated images. There are N Fourier transforms required, each at a cost that scales as $N \ln N$, so the Fourier transforms scale as $N^2 \ln N$. Finally, each Legendre coefficient requires a dot product of length N (with a similar cost, which scales in the same way, to generate the needed Legendre function), or NM operations for M coefficients. When M is large, the Legendre transforms dominate the cost, and to maintain all of the information ($M = N^2$), the total effort scales like N^3.

Even the 2 x 1-D transform method requires large amounts of computer time, so that an improved technique would be desirable. One possible improvement is a technique due to Dilts

(1985), which starts by performing a 2-dimensional FFT on the observed data, and then combines the resulting coefficients in such a way as to produce the desired spherical harmonic coefficients. Claims of substantial speed gains have been attached to this method, but caution is indicated, since early descriptions of its efficiency by Duvall (private communication) are very disappointing.

One problem afflicting the 2 x 1-D transform method of spatial filtering is that, in its simplest form, it assumes that the symmetry axis relative to which the spherical harmonics are computed lies in the plane of the sky, *i.e.*, that the B angle is zero. This is usually not the case when observing the Sun, and though the B angle is always small ($\pm 7^{\circ}$, at most), it is large enough to be significant, especially for high degree modes. There are two solutions to this problem. The simplest is to interpolate along curved lines in the observed image, giving interpolated data that truly lie on parallels of solar latitude. Another (perhaps slightly cheaper) technique is to compute Y_l^m coefficients with the rotation axis in the plane of the sky, and then transform the coefficients to a new coordinate system. Since the rotation transformation connects only modes with different m but the same l, the transformation may be done for each degree individually. Both of these techniques are currently in use, and both appear to work well.

In spite of all the work that has gone into spatial filtering in recent years, several important questions remain: (1) What spatial filters (*i.e.* functions to be multiplied onto the observed data) are really optimal, in the sense of minimizing the confusion between neighboring l and m values? Christensen-Dalsgaard (1984) has done some valuable work on this problem, but there is more to do, both on the theory and in the application of results to practical reductions. (2) What spatial filtering algorithms are most efficient, and what trade-offs exist between efficiency and accuracy? (3) How much of the solar disk should be used in the analysis, and how should it be apodized? How do the answers to these questions depend on the kind of data being processed (velocity or intensity), and on the S/N? (4) Are there other (perhaps nonlinear) techniques that can better discriminate between spherical harmonics with similar l and m?

3.3 Time Series Analysis

Once the individual images of varying quantities have been decomposed into their spherical harmonic components, one must compute temporal power spectra of the individual components. This process usually consists of several steps.

Data culling is usually required to identify and eliminate bad data values. Such values maybe due to weather, instrument failures, calibrations, moths flying through the light beam, or many other causes. Finding these bad values tends to be a tedious and labor-intensive job, and one that is difficult to automate except for very high-quality data.

After bad data are eliminated one must extract timestrings of individual spherical harmonic coefficients. This is essentially a trivial operation, requiring transposition of the 3-dimensional data matrix (amplitude *vs.* l, m, and time) so that its elements are available in a convenient order. For all its triviality, the large amounts of data involved (typically more than 100 Mbyte, for a 10-day run) make this another lengthy chore, and one that is substantially impeded by disk quotas, computer system managers, and others who would prevent one from monopolizing computer resources.

Filling gaps in the data strings can be accomplished by any of several techniques. These work fairly well, provided that the S/N ratio in the data is high, that the gaps are not too long, and that there are adequately long segments of good data adjoining the gaps. These restrictions have so far limited gap-filling to cases in which a few minutes of data are lost to clouds, instrument calibrations, or the like. Filling diurnal gaps has succeeded on artificial data, but has not yet been successfully done with real data.

Apodization involves multiplying the observed timestrings by a tapering function, intended to smooth the transition into data gaps. The object of this operation is to minimize

https://doi.org/10.1017/S007418090015853X Published online by Cambridge University Press

the far sidelobes of the observing window function. Some workers do this, and some do not. Spectra made from the same timestring with and without apodization are clearly different from one another, but it is difficult to convince oneself that apodization results in a significantly cleaner spectrum.

Having apodized the time series (or not), one is ready to compute power spectra. This must be done individually for each spherical harmonic timestring, but with FFTs the power spectrum computation is among the least difficult tasks to perform.

Finally, one must display the power spectra in some intelligible way. This is necessary for debugging, and very helpful for visualizing ways to perform later quantitative analyses. It is not a very easy thing to do, however, due once again to the large quantity of data involved and to its inherently 3-dimensional nature. 2-D image displays (gray-scale or color) do the best, allowing one to display slices at, *e.g.*, constant m or constant l (the by now traditional $l-\nu$ and $m-\nu$ diagrams). These are especially useful because the eye can detect linear or curved patterns in such displays that are clearly present, but in which no individual pixel is significantly different from its neighbors. However, for making quantitative estimates of S/N or line shapes, 1-dimensional plots of power spectra are probably the best.

3.4 Interpreting Power Spectra

Considering how difficult really good power spectra are to produce, it is a pity they they are not more directly useful. Unfortunately, raw power spectra are difficult to interpret, for several reasons: (1) There are very many oscillation modes with detectable amplitudes, and each excited mode generally appears in several spectra with similar nominal l and m values. This makes it difficult (though far from impossible) to identify particular power spectrum peaks with particular modes. (2) For data that suffer from diurnal gaps, each mode is accompanied by several sidelobes, adding to the confusion. (3) Raw power spectra are in any case in the wrong form for easy comparisons with theory. What one really desires is frequency, amplitude, and perhaps phase as functions of n, l, and m.

In the face of these difficulties, one retrieves quantitative information about the spectra in several different ways. To obtain structural information about the Sun (*e.g.* sound speed as a function of depth), it is useful to compute average spectra, averaged over m at a given nominal value of l (Duvall and Harvey, these proceedings). Such averages must be made with due allowance for the frequency variation with m that arises from the Sun's differential rotation, but when this is done they can be used to derive very precise frequency measurements. Convenient and accurate summaries of such frequency information can be obtained by fitting the individual frequencies to the Duvall law (Duvall 1982), especially if α in the law is allowed to be weakly dependent on frequency. It is perhaps a little surprising that two functions, each of a single variable, can represent the observed frequencies to within 1 μHz or so, but such appears to be the case. Indeed, there is good reason to believe that most of this error is in the observed frequencies themselves, and that better observations will reduce the scatter still further.

For studies of the solar differential rotation, most work has so far relied on cross-correlation analyses. These proceed by averaging all of the spectra at a given nominal l value over m (again allowing for some estimated differential rotation), and then cross-correlating this average spectrum with each of the individual spectra. The lag at which this cross-correlation reaches a maximum provides a measure of the frequency offset between each m spectrum and the average. The way in which these frequency offsets (or *splittings*) vary with m and l is then the basic information needed to estimate the nature of the Sun's internal rotation. Cross-correlation methods are convenient and appear to suffer from relatively low noise, because of all the averaging involved. This same averaging causes difficulties, however. The range of l and ν that goes into the average degrades the depth resolution that one may attain with such an analysis. Even worse phenomena may be seen at very low degrees, where leakage from neighboring l values probably causes one to estimate splittings that are systematically too large.

This form of analysis has been very useful up to now, but soon it should be supplanted by something better.

The best means for interpreting spectra would clearly be to make a thorough mode identification, *i.e.* to identify each peak on every spectrum with its corresponding p-mode. Such identifications provide the most natural and the most fine-grained correspondence between theory and observation. Ideally, one would desire the mode frequency, amplitude, phase, and lifetime for each n, l, and m. Something along these lines has been done by Libbrecht (1986) for modes with degrees between 5 and 20, by fitting Gaussians to his observed spectra. Other techniques (*e.g.* Brown, these proceedings, and Duvall, private communication) are currently under development. A great deal of work remains to be done in this area, both in developing new and efficient automated methods to identify modes, and in understanding the quirks and systematic failings of these methods.

In summary, it appears that the discipline of observing solar oscillations has passed through its childhood, and is at last able to pay serious attention to its adult responsibilities. It seems likely that in the next decade, these will be not so much to invent tools to study the Sun, but rather to refine and use the tools that we have. If results to date are any indication, this can only lead us to a deeper understanding of the Sun's internal structure and dynamics.

REFERENCES

Brookes, J.R., Isaak, G.R. and Van der Raay, H.B. 1978, *Mon. Not. Roy. Astron. Soc.* **185,** 1.

Brown, T.M. 1984, in *Solar Seismology from Space,* R.K. Ulrich (ed.), JPL Publication 84-84, p. 157.

Brown, T.M. 1985, *Nature* **317,** 591.

Brown, T.M. 1986 (these proceedings).

Brown, T.M. and Christensen-Dalsgaard, J. 1986, *Astrophys. J.* (submitted).

Brown, T.M. and Harrison, R.L. 1980, *Astrophys. J.* **236,** L169.

Cacciani, A. and Fofi, M. 1978, *Solar Phys.* **59,** 179.

Christensen-Dalsgaard, J. 1984, in *Solar Seismology from Space,* R.K. Ulrich (ed.), JPL Publication 84-84, p. 157.

Deubner, F.-L. 1975, *Astron. Astrophys.* **44,** 371.

Dilts, G.A. 1985, *J. Comp. Phys.* **57,** 439.

Duvall, T.L. 1982, *Nature* **300,** 242.

Duvall, T.L. and Harvey, J.W. 1986 (these proceedings).

Duvall, T.L., Harvey, J.W. and Pomerantz, M.L. 1986, *Nature* **321,** 500.

Evans, J.W. 1980, in *Solar Instrumentation -- What's Next?* R.B. Dunn (ed.), Sacramento Peak Observatory, Sunspot, p. 155.

Frazier, E.N. 1968, *Z. Astrophys.* **68,** 345.

Grec, G. and Fossat, E. 1977, *Astron. Astrophys.* **55,** 411.

Grec, G., Fossat, E. and Pomerantz, M. 1983, *Solar Phys.* **82,** 55.

Hill, F., Haber, D., Toomre, J. and November, L. 1986, in *Seismology of the Sun and the Distant Stars,* D.O. Gough (ed.), D. Reidel, p. 85.

Kneer, F., Newkirk, G. and von Uexkull, M. 1982, *Astron. Astrophys.* **113,** 129.

Kotov, V.A., Severny, A.B. and Tsap, T.T. 1978, *Mon. Not. Roy. Astron. Soc.* **183,** 61.

Kuhn, J.R. 1982, *Astron. J.* **87,** 196.

Leibacher, J.W. and Stein, R.F. 1971, *Astrophys. Lett.* **7,** 191.

Leighton, R.B., Noyes, R.W. and Simon, G.W. 1962, *Astrophys. J.* **135,** 474.

Libbrecht, K.G. 1986, *Nature* **319,** 753.

Libbrecht, K.G. and Zirin, H. 1986, *Astrophys. J.* **308,** 413.

Nishikawa, J., Hamana, S., Mizugaki, K. and Hirayama, T. 1986, *Publ. Astron. Soc. Japan* **38,**

277.
Noyes, R.W. and Rhodes, E.J. (eds.) 1984, *Probing the Depths of a Star: The Study of Solar Oscillations from Space*, JPL.
Rhodes, E.J., Cacciani, A., Tomczyk, S. and Ulrich, R.K. 1986, in *Seismology of the Sun and the Distant Stars*, D.O. Gough (ed.), D. Reidel, p. 309.
Rust, D.M., Kunski, R. and Cohen, R. 1986, *Johns Hopkins APL Tech. Digest* **7**, 209.
Scherrer, P.H. 1978, in *Proc. Symp. on Large-Scale Motion on the Sun*, Sacramento Peak Obs., Sunspot, p. 12.
Ulrich, R.K. 1970, *Astrophys. J.* **162**, 933.
Ulrich, R.K., Rhodes, E.J., Cacciani, A. and Tomczyk, S. 1984, in *Solar Seismology from Space*, R.K. Ulrich (ed.), JPL Publication 84-84, p. 157.
Willson, R.C. 1979, *Appl. Optics* **18**, 179.
Woodard, M. and Hudson, H. 1983, *Solar Phys.* **82**, 67.