Measurements of the Effects of Gravity Waves in the Middle Atmosphere Using Parametric Models of Density Fluctuations. Part I: Vertical Wavenumber and Temporal Spectra

R. J. SICA AND A. T. RUSSELL

Department of Physics and Astronomy, University of Western Ontario, London, Ontario, Canada

ABSTRACT

Parametric models of spectral analysis offer several distinct advantages over statistical methods such as the correlogram analysis. These advantages include higher spectral resolution and the ability, in principle, to separate correlated (i.e., wave) behavior from noise-driven (i.e., turbulent) behavior in the measurements. Here parametric models are used to highlight the spatial and temporal intermittency of the gravity wave spectrum. In Part II of this series the spatial and temporal spectrum are used to calculate energy dissipation and the eddy diffusion coefficient. The spectra are computed from measurements of density fluctuations obtained using a large power-aperture product lidar during a 6-h period on 30 August 1994. It is shown that parametric models provide an excellent representation of the temporal and spatial data series. One difficulty of parametric models is selecting the model order, an analogous situation to determining the proper lag in the correlogram procedure or the window length in the periodogram method. Extensive experimentation has shown that the ratio of the data matrix eigenvalues to the photon noise eigenvalues is an excellent indicator for the selection of the model order. The underlying spectral form found using the parametric models is similar to the standard correlogram method, that is, nominal underlying spatial and temporal spectral slopes between -2 and -4 and -1.25 and -2, respectively, with variability outside this range. The spatial–temporal behavior of the spectra is highly variable with numerous intermittent and intense features rising well above the photon noise floor. The vertical wavenumber spectra on this night may show a variation of spectral slope with height; however, the slope is both extremely sensitive to the noise level of the data, steepening as the signal-to-noise level increases, and highly variable in time. The temporal spectra also show considerable variation with height, both in magnitude and slope.

1. Introduction

The basic shape of the temporal and spatial spectrum of gravity waves and turbulence has been well established by many different techniques, for example, radar (Tsuda et al. 1989), lidar (Wilson et al. 1991), rocket soundings (Blix et al. 1990), and balloon soundings (Fritts et al. 1988). Often the available measurements lack either the temporal resolution to study the gravity wave spectrum at resolutions of tens of minutes or less, the capability to make long-term (e.g., hours of) measurements in the same air mass or the bandwidth to investigate more than two scale heights of atmospheric variations. The development of large power-aperture product lidars has, in part, addressed these issues. These lidars can have large bandwidths (70 km or more) and make relatively continuous measurements of density fluctuations in clear-sky conditions at high signal-to-noise ratios.

While the basic spectrum of gravity waves is well known, no first principle theory has been generally accepted to explain their spectrum. Explanations of the shape and magnitude of the spatial spectrum of gravity waves in the "saturated" region have evolved in complexity since the original linear saturation theories (e.g., Dewan and Good 1986; Smith et al. 1987) and now includes theories involving nonlinear interactions between waves, diffusion, and the nonseparability of the spatial and temporal spectra (e.g., Fritts 1985; Weinstock 1990; Hines 1991; Gardner 1994). An explanation of the temporal spectrum invoking turbulence arguments has been made by Gage (1979). One of the underlying themes of these various studies is the need to know how the spectrum evolves at high spatial–temporal resolution, in addition to knowing what geophysical conditions the breaking waves experience. The former topic is the motivation of this study, while the latter topic is motivating a search for unstable layers in the middle atmosphere (Sica and Thorsley 1996).

The discussion in this paper will be limited to the

Corresponding author address: Dr. Robert J. Sica, Department of Physics and Astronomy, University of Western Ontario, London, ON N6A 3K7, Canada.
E-mail: sica@uwo.ca

analysis of density fluctuations measurements from the Rayleigh-scatter beam of the University of Western Ontario's Purple Crow Lidar. An extremely important part of the analysis of the gravity wave spectrum is the effect of the background wind on the wave spectra. Temporal spectra reported in this work and in Part II use the observed frequency, rather than the intrinsic frequency. Gardner et al. (1993) have shown that data averaging can substantially reduce the spectral magnitude at high wavenumbers due to Doppler shifting. The spatial–temporal averaging chosen for this study should be sufficiently high enough that these effects are small. As argued by Manson (1990), the Doppler-shifting effects on the observed frequencies should be small for conditions with low winds as anticipated in August from general circulation models of the middle atmosphere (e.g., Hamilton et al. 1995). In the future the incorporation of wind information will be pursued using the U.K. Meteorological Office assimilation forecasts, using measurements in the troposphere, upper mesosphere, and lower thermosphere from the VHF and MF radar facilities at the University of Western Ontario and, eventually, wind measurements from the lidar itself.

2. The measurements

The Purple Crow lidar is a monostatic lidar system that can simultaneously measure both Rayleigh and sodium resonance fluorescence backscatter. The power spectral density (henceforth PSD) determinations in this paper use measurements from the Rayleigh-scatter system, which is described in detail by Sica et al. (1995). Briefly, the Purple Crow lidar has a large power-aperture product due to the use of a high power transmitter and large aperture receiver. The transmitter is a Nd:YAG laser operating at the second harmonic, with an output energy of nominally 600 mJ per pulse at 20 Hz. The receiver is a 2.65-m diameter liquid mercury mirror, which is coupled to a photon counting system by an optical fiber. Sica et al. (1995) have shown that the liquid mirror has a performance similar to a traditional glass telescope of the same area and also that the system efficiently detects the large backscatter returns. The large power-aperture product allows high signal-to-noise ratio density fluctuation measurements to be obtained.

Signal processing procedures to determine density fluctuations from the returned power have been given by Gardner et al. (1989b). In their scheme the relative density profile is detrended and the residuals corrected for bias due to atmospheric transmission variations and laser power fluctuations. The fluctuations are detrended by removing a third- or fifth-order polynomial fit to the average density profile depending on the measurement's bandwidth. The appropriate polynomial must fit the measurements sufficiently better (i.e., at least 10 times or smaller) than the size of the atmospheric fluctuations, and of course, must fit better than the photon noise floor

of the measurements. The mean is then removed from the residuals.

It is well known that to conserve kinetic energy density gravity waves will grow in amplitude with height. Hence, in this simple view, the vertical wavenumber spectrum will consist of growing sine waves. Spectral estimation methods typically assume sinusoidal signals of constant amplitude. To compensate for the effect of the background density profile each vertical wavenumber data series is scaled by the square root of the measured density as suggested by Hines (1960). Figure 1 shows the unscaled density perturbations as compared to the scaled perturbations for the measurement period. Both panels have a vertical resolution of 48 m and a temporal resolution of 9 min. The measurements have smoothed spatially with a filter with a Kaiser–Bessel filter with a 1008-m cutoff and temporally by 3's and 5's with a 63-min cutoff (Hamming 1977). It is clear that the square root of density scaling helps to normalize the amplitude of the major perturbations.

A representative height profile of density fluctuations is shown in Fig. 2a. The profile has 125 samples spaced 144 m apart and was acquired over 1 min. The ensemble of profiles is temporally smoothed using a finite-impulse response Kaiser–Bessel filter with a 21-min cutoff. Figure 2b shows the same residual scaled by the square root of density. In this particular measurement the effects are most noticeable above 45-km altitude. The difference in the PSD between the unscaled and scaled residuals is shown in Fig. 3a. Note both PSDs have been scaled to 0 dB at the peak to account for differences in units. Both PSDs show the same features, though the peak around 5×10^{-4} m^{-1} is at a bit higher wavenumber in the unscaled spectrum. The unscaled spectrum has a more pronounced low wavenumber peak with more low wavenumber rolloff than the scaled case.

The residuals may also be detrended by removing a least squares linear fit. The effect of detrending is shown for the scaled residual in Fig. 3b. Detrending improves the low wavenumber end of the spectrum with little effect above the initial peak. All residuals used for vertical wavenumber spectral estimations were scaled and linearly detrended.

3. Parametric models of the power spectral density

a. Introduction to the models

Spectral analysis of the density fluctuations is performed using parametric models. Statistical procedures such as correlogram analysis require the autocorrelation sequence of a process to be set to zero outside a specified spectral window (e.g., Blackman and Tukey 1958). The correlogram procedure is well known, and has been extensively used for studies of the atmospheric PSD (e.g., Tsuda et al. 1989). The autocorrelation function is determined out to a specified lag and a window function applied to the autocorrelation. The appropriately scaled

FIG. 1. Density fluctuation measurements from the night of 30 August 1994. Both data series have a resolution of 48 m in height and 9 min in time. The measurements are smoothed using a digital filter with a 1008-m spatial bandwidth and a 63-min temporal bandwidth. (a) The unweighted residuals expressed as a percentage. (b) The residuals weighted by the square root of the measured density, in units of $(kg\ m^{-3})^{1/2}$. A color version of this figure is available at http://PCL.physics.uwo.ca.

Fourier transform of the autocorrelation is then the PSD. The art of the correlogram is the trade-off of bias and spectral variance as the lag number is varied (''window carpentry'').

Parametric models can use information outside of this range imposed by a spectral window, which eliminates the sidelobes associated with window functions and increases spectral resolution. It has been well established for over 30 yr that parametric models are useful when applied to real measurements that contain noise, particularly if the noise can be independently estimated (as is possible for the Rayleigh-scatter lidar). Textbooks such as Bloomfield (1976), Jenkins and Watts (1968), and Marple (1987) show examples of using parametric

models on datasets including noise. Autoregressive (AR) models, moving average (MA) models, the combination of these two models, the autoregressive-moving average (ARMA) model, and the multiple signal classification (MUSIC) algorithm have been considered in detail here. As a general disclaimer, all of these methods (including the statistical methods, the correlogram, and its equivalent, the periodogram) assume the process is stationary. For a stationary process the mean does not vary with time and the process's autocorrelation function depends only on the difference of the time indices. This standard assumption is questionable, particularly for the temporal series, which extend over several hours, but is probably the best we can do at this time. Non-

FIG. 2. Individual density perturbation measurement at 0702 UTC. This measurement is shown at a resolution of 144 m in height and was obtained in 1 min. The measurement has been smoothed temporally using a filter with a 21-min bandwidth. (a) The unweighted residuals expressed as a percentage. (b) The residuals weighted by the square root of the measured density, in units of $(\text{kg m}^{-3})^{1/2}$. The linear trend of both residuals has been removed.
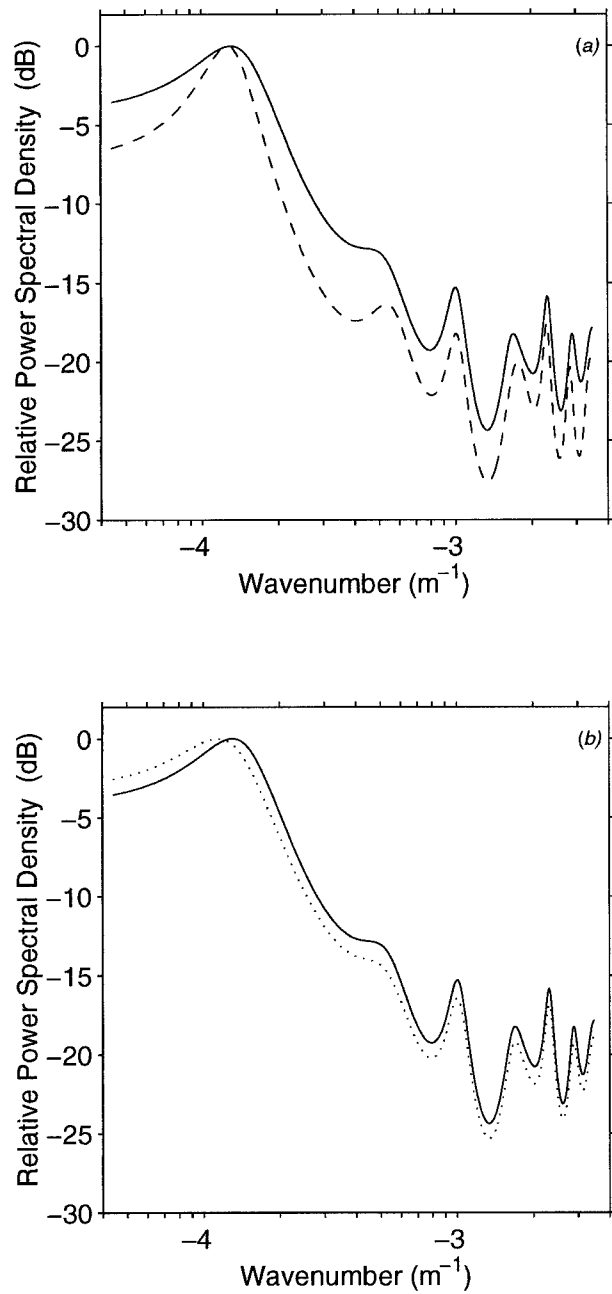

FIG. 3. Autoregressive (AR) model PSDs for the data series shown in Fig. 2. (a) PSD with (solid line) and without (dashed line) density scaling, without linear trend removal. Both PSDs have been normalized to 0 dB at the peak. (b) Normalized PSD for density-scaled data series without (dotted line) and with (solid line) the linear trend removed.

stationary techniques should (and will) be considered in the future.

An AR model of a data series $x[n]$ attempts to estimate the autocorrelation sequence of the data series from a set of linear equations. Here the noise-driven output of the model, u, is passed back through the p model coefficients in a manner similar to an infinite impulse response filter of the form

$$x[n] = -\sum_{k=1}^{p} a[k]x[n - k] + u[n], \qquad (1)$$

where the a's are the model coefficients. If the noise-driven output of the model is assumed to have zero mean and variance ρ_w, then $u[n]$ can be replaced by ρ_w in (1). Note the underlying process here is correlated, as each new prediction depends on past information. The atmospheric analog is correlations in the density fluctuations induced by the passage of waves. Solution of the model coefficients and the driving noise variance can be cast as a least squares problem of fitting the measurements to the model. The order of the process, p, is the order of the model fitted. The Fourier transform of the output sequence can be shown to be related to the Fourier transform of the input sequence, yielding the following form for the PSD as a function of frequency f:

$$P_{\mathrm{AR}}(f) = \frac{\tau \rho_w}{|A(f)|^2}, \tag{2}$$

where the quantity $A(f)$ is the Fourier transform of the model coefficients for measurements sampled at spacing τ (Marple 1987).

An infamous method for calculating the model coefficients, the maximum entropy method, attempts to extrapolate the autocorrelation sequence on the assumption of a uniform distribution of the autocorrelation sequence. The pitfalls of this technique are well known, in particular line splitting of spectral features and spurious peaks. Maximum entropy methods such as Berg's method were not investigated for this study. The model coefficients and driving noise variance can be estimated directly from the measurements rather than the autocorrelation function using least squares techniques (Marple 1987). The AR model coefficient determinations, which act directly on the measurements, include the covariance and modified covariance methods. Covariance methods have been shown to be resilient to line splitting.

In the covariance method a coefficients are chosen to minimize the difference between the model and measurements in the least squares sense,

$$\mathbf{e}_p^f = \mathbf{T}_p \cdot \begin{pmatrix} 1 \\ \mathbf{a}_p^f \end{pmatrix}, \tag{3}$$

where the differences between the measurement and model is the vector $\mathbf{e}_p^f$ and, in this case, the data matrix is a Toeplitz matrix, $\mathbf{T}_p$. The data matrix is given by

$$\mathbf{T}_p = \begin{bmatrix} x[p+1] & \cdots & x[1] \\ \vdots & \ddots & \vdots \\ x[K-p] & \cdots & x[p+1] \\ \vdots & \ddots & \vdots \\ x[K] & \cdots & x[K-p] \end{bmatrix}, \tag{4}$$

where K is the number of data samples.

The MA model has similar resolution to the correlogram (Jenkins and Watts 1968). The form of this model is similar to a finite impulse response filter. White

noise is passed through each of the q model coefficients at different delay times, then summed. The physical meaning of this model is that the atmosphere is a finite bandwidth black box that redistributes the gravity wave energy. The model equation is

$$x[n] = \sum_{k=1}^{q} b[k] u[n-k] + \rho_w, \tag{5}$$

where the b's are the model coefficients and we have replaced $u[n]$ by zero-mean white noise with variance ρ_w. The PSD is found from the relation

$$P_{\mathrm{MA}}(f) = \tau \rho_w |B(f)|^2, \tag{6}$$

where $B(f)$ is the Fourier transform of the model coefficients. The MA model coefficients can be found by solving a high-order AR model (Marple 1987).

A model equation to a process can be found that uses the AR model as a first pass to fit the measurement followed by an MA model to reduce the residual variations. Such a model allows for both correlations in the measurements in addition to any underlying statistical process. The appropriate model is the ARMA model, which is the combination of (1) and (5)

$$x[n] = -\sum_{k=1}^{p} a[k] x[n-k] + \sum_{k=0}^{q} b[k] u[n-k] \tag{7}$$

and has the associated PSD

$$P_{\mathrm{ARMA}}(f) = \tau \rho_w \left| \frac{B(f)}{A(f)} \right|^2. \tag{8}$$

In general, since this model has both the advantages of finite and infinite impulse responses it should both be versatile and use fewer model coefficients than pure AR or MA processes. If the AR model order and the MA model order are equal it can be shown that the ARMA model order is equal to the number of waves needed to fit the data. However, the sequential ARMA models investigated appear to be highly sensitive to noise (instrumental and geophysical) in the measurements as will be discussed later.

It is important to understand the difference between driving noise variance, ρ_w, the spectral variance, and the measurement noise. The driving noise variance of the model and the associated spectral variance of a PSD estimate are separate notions from any measurement noise that may be present. The AR, MA, and ARMA models calculate coefficients that can be considered as a ''black box'' whose input is a zero-mean noise signal, the driving noise variance, and whose output fits the measurements. The Fourier transform of the process's output is the PSD. The driving noise variance scales the magnitude of the PSD as is evident from (2), (6), and (8). A large spectral variance means the PSD has large fluctuations and appears noisy. A PSD can have a large spectral variance even for a data series with no measurement error (i.e., a correlogram of white noise using

a large number of lags, say half the data series). A trade-off must be made between spectral resolution and variance.

In the case of the lidar measurements, the dominant measurement noise is due to the counting statistics of the photomultiplier, the photon noise. The photon noise can be converted into a constant photon noise floor in the transform domain, which can then be compared to the PSD estimate (Gardner et al. 1989b). The photon noise is in addition to the spectral variance.

Eigenanalysis-based methods of frequency estimation are extremely good at removing white noise (e.g., temporal data series) and reasonably good with pink noise (e.g., the spatial data series). While these techniques allow high-resolution spectral estimation they do not give a true PSD (in the sense that the area under the spectrum is not equal to the power). The MUSIC method is based on determining the eigenvalues of the data matrix by single value decomposition. The p eigenvalues are sorted in descending order and divided into M signal and $(p - M)$ noise subspaces (details on this separation for the measurements will be discussed later). For white noise, the eigenvectors will be orthogonal to the signal eigenvectors allowing a ''PSD'' to be defined as

$$P_{\text{MUSIC}}(f) = \frac{1}{\mathbf{e}^H(f) \cdot \left(\sum_{k=M+1}^{p} \mathbf{v}_k \mathbf{v}_k^H \right) \cdot \mathbf{e}(f)}, \qquad (9)$$

where the $\mathbf{e}$ vectors are the noise eigenvectors, the $\mathbf{v}$ vectors are the signal eigenvectors, and the superscript H means the Hermitian transpose of the vector (Marple 1987). Hence, the denominator will be small (ideally 0) when $\mathbf{v}$ and $\mathbf{e}$ are orthogonal, corresponding to peaks in the spectrum. The quotes around PSD above (9) remind us this estimate is not a true PSD, in the sense the area under the MUSIC spectrum does not recover the power in the modeled process as is the case for a true PSD.

b. Validation of the models on a real data series

The data series shown in Fig. 2b (henceforth called the test series) will be used to illustrate the advantages and disadvantages of the parametric models compared to the correlogram (or its equivalent the periodogram) technique. First, a simplified representation of the series will be used to illustrate the techniques. Second, the actual test series will be analyzed in detail.

The test series is an extremely complex waveform. The test series was fit to 30 damped sinusoids using Prony's method, which will be discussed in detail in Paper 2 of this series. Figure 4a shows the test-case-data points, while the solid line is the Prony fit. The fit is an extremely good match to the data series, and in fact has an rms deviation from the data three times less than the photon noise floor of the measurements. From these 30 waves four representative waves are picked

with amplitudes greater than the photon noise floor (Fig. 4b and Table 1). The sum of these four waves is shown in Fig. 4c. The fit is reasonable except at the lowest heights. The bandwidth and relative energy density spectrum of these four waves is shown in Fig. 4d. The narrow bandwidths correspond to waves that have smaller damping. What appears in Fig. 4a to be a single wave of 8–10-km vertical wavelength is actually two waves of 9.2- and 5.6-km wavelength. Having created a simpler subset of the test series the different models and their trade-offs can be compared.

To further mimic the measurements normally distributed white noise of the same magnitude as the measurements was added to the 4-wave series. The correlogram of the 4-wave series is shown in Fig. 5a using 62 lags, the maximum resolution possible. Even at this resolution, the two low wavenumber waves (waves 1 and 2) cannot be resolved, while waves 3 and 4 are visible, but are surrounded by considerable spectral variance at high wavenumbers. The MA model for an order of 62 is, as expected, similar to the correlogram, though the MA model underestimates the magnitude of the low wavenumber waves (Fig. 5b). However, it does not show spurious high wavenumber components like the correlogram.

The covariance and modified covariance models are shown using an order of 16 (Figs. 5c and 5d). The order was chosen in a manner analogous to window closing for the correlogram, that is, the maximum order with reasonable spectral variance. Both methods show higher resolution than the correlogram. The covariance spectrum has better definition at the low wavenumber peaks, but cannot separate waves 1 and 2. Wave 3 is apparent in both spectra, but suppressed in magnitude due to its relatively large bandwidth.

The ARMA model, though in principle offering the best features of the AR and MA models, has excessive variance at high wavenumbers for an equivalent order compared to the AR models (here $p = 16$; Fig. 5e). However, the ARMA model is superior to the AR models for the location of wave 3. Decreasing the order, while decreasing the spectral variance, causes the model to fail in finding either of the low wavenumber waves detected at an order of 16.

The MUSIC spectral estimation at eighth order is able to resolve the low wavenumber peaks as well as waves 3 and 4. Due to the large bandwidths of waves 1 and 2 the relative magnitudes of these waves is not correctly retrieved. The MUSIC algorithm is extremely good at eliminating the noise from the signal as well as suppressing spurious peaks.

The full test series (Fig. 2b) was analyzed with the same suite of techniques as the 4-wave series (Fig. 6). Again, the correlogram and MA models have relatively high spectral variances associated with the high lag (order) used. Even at this high order the feature at 1×10^{-3} m^{-1} does not clearly stand out above the noise floor. The covariance method has the best trade-off be-

FIG. 4. Decomposition of the test series. (a) The test series (points) and the fit to the test series using Prony's method (solid line). (b) The four waves whose parameters are listed in Table 1. (c) The sum of the four waves (solid line) compared to the test series (points). (d) Relative energy spectral density and bandwidth of the four waves.

tween preserving the ratio of the peaks compared to the correlogram, MA, modified covariance and ARMA models, with reasonable spectral variance. The best wavenumber estimation is achieved by the MUSIC algorithm, which can resolve the low wavenumber peaks as well as the high-frequency peak at 1×10^{-3} m^{-1}.

Since each spectrum cannot be individually tuned for order an order selection criteria needed to be developed. Several criteria (e.g., final prediction error, Akaike information criteria, minimum description length, and criterion autoregressive transfer) exist and were tried, but were not found to be satisfactory as typically they underestimated the process order. Since the noise in the

signal is known a new criteria was developed based on the eigenvalues of the data matrix. The photon noise is used to generate a noise matrix of normally distributed random numbers with zero mean and variance equal to the photon noise. The ratio of the noise eigenvalues to the eigenvalues of the data matrix (e.g., a signal-to-noise ratio) is then computed. A threshold cutoff of 0.825 was found experimentally to give a good balance between resolution and spectral variance. The signal-to-noise ratio of the eigenvalues for the test series is shown in Fig. 7. It can be shown that the appropriate AR model order is twice the number of signal eigenstates (Marple 1987). Hence, for the test series where the number of signal eigenstates is 8 the appropriate AR model order is 16.

c. Mechanics of the procedure

For the purposes of this initial study measurements from the night of 30 August 1994 are used. The measurements commenced at 0326 UTC (2326 EDT) and ended at 0933 UTC. This night is chosen for two reasons. The first reason is that about 6 h of continuous measurements are available with no data gaps to com-

TABLE 1. Wave parameters for the 4-wave test series.

| Wave no.	Wavenumber (m^{-1})	Amplitude [(kg m^{-3})$^{1/2}$]	Growth rate* (m)	Phase (°)
1	1.09×10^{-4}	1.33×10^{-3}	-2.00×10^{-4}	-92.8
2	1.78×10^{-4}	8.03×10^{-4}	-1.48×10^{-4}	36.0
3	4.23×10^{-4}	6.61×10^{-4}	-4.56×10^{-4}	28.0
4	9.72×10^{-4}	8.10×10^{-5}	-6.29×10^{-5}	-30.4

* Negative growth rate means the wave is damped.

FIG. 5. Correlogram and parametric models of the 4-wave test sequence with added noise equivalent to the photon noise for the measurement at 0702 UTC. The noise floor of the PSDs are indicated by dotted lines on the figures (except for the MUSIC spectrum). (a) Correlogram. (b) Moving average model. (c) Covariance AR model. (d) Modified covariance autoregressive model. (e) Autoregressive-moving average (ARMA) model. (f) Multiple Signal Classification (MUSIC) algorithm.

plicate the analysis. The second reason is the temperature profile is relatively unperturbed and is close to that anticipated from empirical model atmospheres; hence, there is nothing obviously "exotic" about this night. For the vertical wavenumber spectra, the density fluctuations are divided up into three height ranges with bandwidths of about 20 km. As described in Gardner et al. (1989a), there is a trade-off between bandwidth and noise floor for the spatial spectrum. This bandwidth allows for an acceptable noise floor but sufficient bandwidth to locate the low wavenumber peaks anticipated in the gravity wave spectrum. An overlap of 5 km between regions was used to ensure continuity in the residuals.

Before determining the model coefficients the vertical (temporal) density fluctuations are temporally (spatially) filtered using finite impulse response filters. The ideal low-pass filter was not used due to its poor sidelobe behavior. The data series was smoothed using either Kaiser–Bessel filters with 40 dB or more of attenuation or by 3's and 5's (Hamming 1977). The length and hence the variance decrease of the filter is adjusted for the desired resolution. Infinite impulse response filters were tested but found to give unreasonable variance reduction estimates. The noise floor due to photon counting is

estimated from the measurements. The noise floor depends directly on the sampling interval and filter variance reduction and inversely on the number of counts in the lidar profiles. The noise floor can be decreased by coadding appropriately in space and time or increasing the bandwidth of the filter.

The PSDs calculated in this study have units of meters (spatial) or seconds (temporal); that is, the angular units have been removed. From these natural units the gravity wave polarization equations can be used to scale the spectrum to a kinetic energy density spectrum in units of $(J\ m^{-3})\ m$ or $(J\ m^{-3})\ s$ via

$$\kappa(z,\ t) = \frac{1}{2}\left(\frac{g(z,\ t)}{N(z,\ t)}\right)^2 \qquad (10)$$

(Gardner et al. 1989a). Here g is the acceleration due to gravity and N the angular buoyancy frequency determined from the lidar temperature measurements. The $(z,\ t)$ reminds us these quantities are height and time dependent.

4. Vertical wavenumber spectra

The density perturbations were divided in three regions, low (L_{sp}; 32–50 km), middle (M_{sp}; 45–65 km),

FIG. 6. Correlogram and parametric models for the actual test series shown in Fig. 2, plotted in the same format as Fig. 5.

and high (H_{sp}; 60–80 km). The measurements were coadded to height bins of 288 m and temporal records of 540 s. Smoothing temporally by 3's and 5's gives a bandwidth of 63 min and a variance decrease of 6. This choice of sampling gives acceptable spectral estimates at the upper height ranges where the signal-to-noise ratio
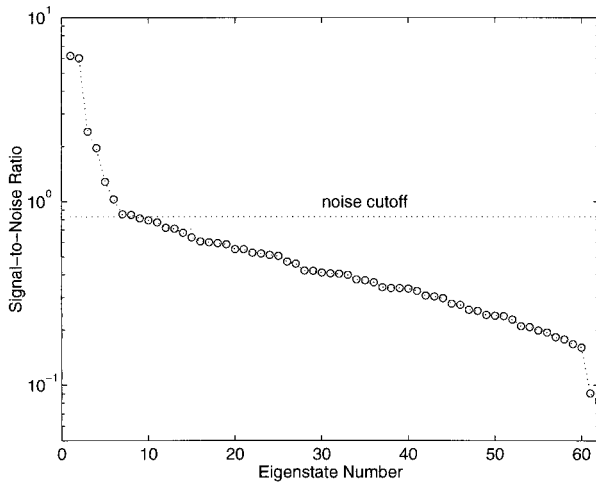

FIG. 7. Signal-to-noise ratio as a function of eigenstate number for the test series. The empirically derived cutoff between noise and signal eigenstates is shown in the figure.

is lower. However, the photocount rates in the low region are sufficiently high that higher spatial–temporal resolution can be employed. The resulting data series in the 32–50-km region is called the L'_{sp} series. For this series the measurements are coadded to 144-m bins and 61-s temporal records, then temporally smoothed with a 21-point Kaiser–Bessel filter with a 21-min bandwidth. The L'_{sp} series will be discussed first and used to illustrate the method of model-order choice, as well as the intermittency in the vertical wavenumber spectrum. The lower resolution series will then be used to illustrate the variations in the spectrum between height regions.

High spatial–temporal resolution vertical wavenumber spectra in the upper stratosphere

The covariance AR model order for the L'_{sp} series using the signal-to-noise ratio based on the eigenvalues of the data matrix is shown in Fig. 8. The order varies between 8 and 28 over the measurement period. Figure 9 shows the mean covariance AR PSD for this period with the noise floor removed. The mean correlogram at 62 lags is also shown (divided by 10 for ease of comparison). The arrow on the ordinate indicates the value of the Hines wavenumber for the AR PSD. This wavenumber, m_σ, was defined by Hines (1991) as

FIG. 8. Variation with time of the AR model order for the L'_{sp} series.

FIG. 9. Average covariance AR model PSD for the L'_{sp} data series (top curve) at the orders shown in Fig. 8 compared to the correlogram (lower curve) at 62 lags (i.e., half the data series). The correlogram has been divided by 10, and the noise floor has been removed from both spectra. The Hines wavenumber (described in the text) for the AR PSD is indicated by the arrow. The standard deviation of the mean Hines wavenumber over the measurement period is ±25%.

$$m_\sigma = \frac{N/(2\pi)}{V_{RMS}} = \frac{[\rho N/(2\pi)]^2}{g \int (PSD)\, dm}, \qquad (11)$$

where V_{RMS} is the root-mean-square wind speed. Using (10) yields the right-hand side of (11), where the integral of the power spectral density is over all wavenumbers available. The spectrum in the regions of and to the right of m_σ in Fig. 9 is considered to be nonlinear. The conversion from density fluctuations to kinetic energy density should be treated with caution in these regions. For the average spectrum shown in Fig. 9, the majority of the tail region is linear (that is to the left of the arrow).

Figure 10 shows contours of the individual kinetic energy density spectra as a function of time and wavenumber for the correlogram, covariance AR model, and MUSIC algorithm. The photon noise floor from each individual scan has been removed from the correlogram and the AR model. Isolated, coherent features appear particularly at low wavenumbers, where there is considerable power from the beginning of the measurements until about 0500 UTC, then another enhanced period from 0530 UTC until almost 0800 UTC, followed by a period of considerably lower power. All three spectral contours show this general trend, with increasingly higher resolution. The higher spectral variance of the correlogram is evident at wavenumbers higher than 1/(2000 m). The MUSIC PSDs show a narrow and nearly continuous band of power around 1/(10 000 m) from the beginning of the time series until about 0800 UTC. During the first period of intensification (0336–0500 UTC) there is also considerable power at higher wavenumbers, but during the second intensification between 0530 and 0730 UTC little power exists above 1/(4000 m) = 2.5×10^{-4} m^{-1}. This dynamic behavior of the PSDs characterizes the intermittency of the gravity wave spectrum.

Some individual AR PSDs for this period are shown in Fig. 11 to emphasize the variability of the total power with time. The peak magnitude of the spectrum is seen to vary from lows around 25 (J m^{-3}) m to almost 1000 (J m^{-3}) m with essentially constant noise floors (i.e., equivalent signal-to-noise ratios in the residuals). At times several distinct features are present, though the spectrum at 0856 UTC has essentially only the low wavenumber peak. These differences in spectral form are not merely due to the order of the model calculation. The orders of the spectra shown in Figs. 11a–d are 20, 16, 24, and 16, respectively. The variations in the spectra are geophysical variations of the density perturbation measurements.

Intermittency is also evident from the variations in the Hines wavenumber over this period (Fig. 12). The Hines wavenumber varies by a factor of 3, ranging from 7×10^{-4} m^{-1} to 2.15×10^{-3} m^{-1}. As one would anticipate, during the periods of highest power the Hines wavenumber is smallest and more of the spectrum is nonlinear. When indicating the Hines wavenumber on a spectrum averaged over a long period (i.e., hours, days, or weeks), it is important to consider the amount of variability in the average. The same comment would, of course, extend to the spectral averages themselves given the variability shown in Fig. 10.

An individual kinetic energy density spectrum can be integrated over wavenumber to produce a total kinetic energy density time series. From this time series the change of kinetic energy density with time, or power, can be calculated (Fig. 13). This calculation shows that the maximum power variations are around 0620 UTC when the low wavenumber feature around 1×10^{-4} m^{-1}
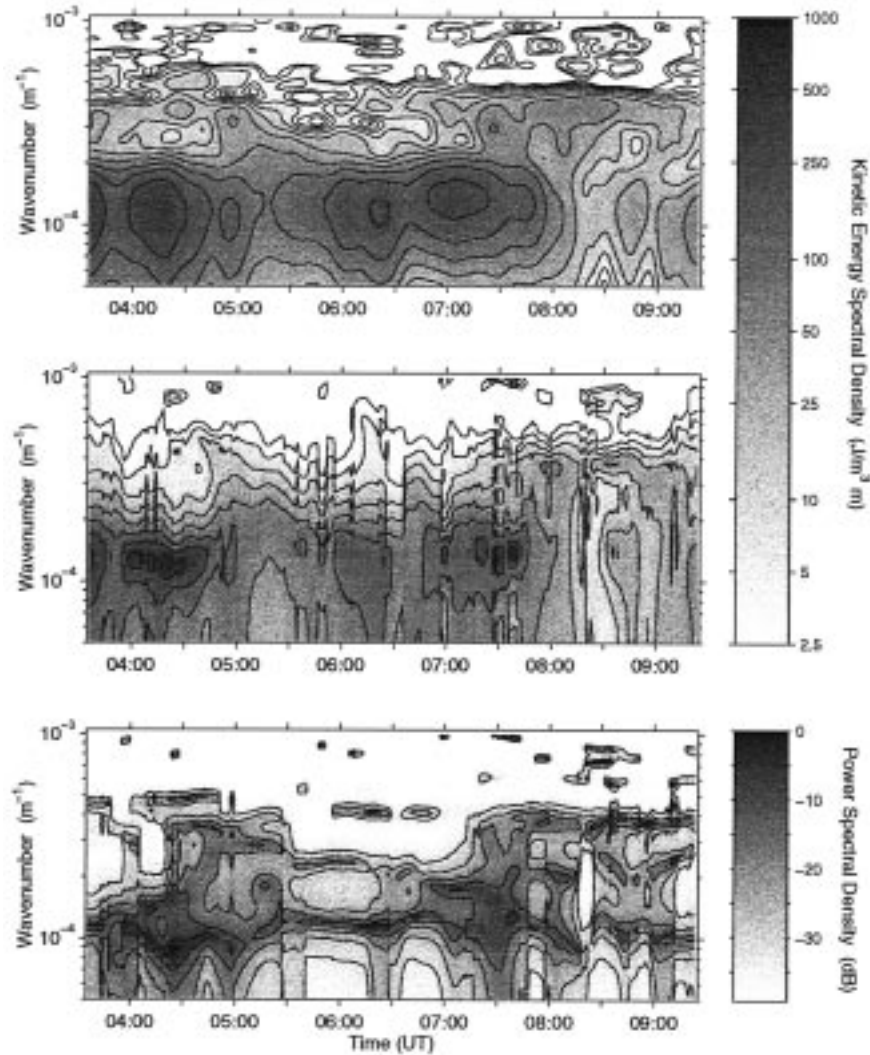
FIG. 10. Contour plots of the L'_{sp} kinetic energy density PSDs. (top) Correlogram. The peak value is 388 (J m^{-3}) m. (middle) The AR covariance model spectra. The peak value is 858 (J m^{-3}) m. (bottom) MUSIC algorithm. The correlograms and AR model spectra have the photon noise floor removed. The MUSIC spectra are in units of dBs. A color version of this figure is available at http://PCL.physics.uwo.ca.

is at its narrowest and highest intensity. The power "bursts" have magnitudes around 5–10 μW m^{-3}.

5. Height variations of the vertical wavenumber spectra

The order selection method discussed previously was used to pick the appropriate order for the L_{sp}, M_{sp}, and H_{sp} data series. For the L_{sp} data series the average order is 34, or about half the data series. The average order for the M_{sp} data series is 16, while for the H_{sp} data series it is 10.

a. Variations in kinetic energy density

The average AR kinetic energy density spectrum for the L_{sp}, M_{sp}, and H_{sp} data series is shown in Fig. 14. The spectra were calculated at the three heights at the same spatial and temporal resolution, hence the L_{sp} and M_{sp} spectra are at a higher than optimal signal-to-noise ratio and have excessive spectral variance evident as spikes above 1×10^{-3} m^{-1}. The L_{sp} spectrum has a low wavenumber peak at about 1.2×10^{-4} m^{-1} in agreement with the analysis of the L'_{sp} data series, with additional peaks at higher wavenumbers. At this high-order number the spectral variance is increased relative to the L'_{sp} series, with noise features present above 1×10^{-3}
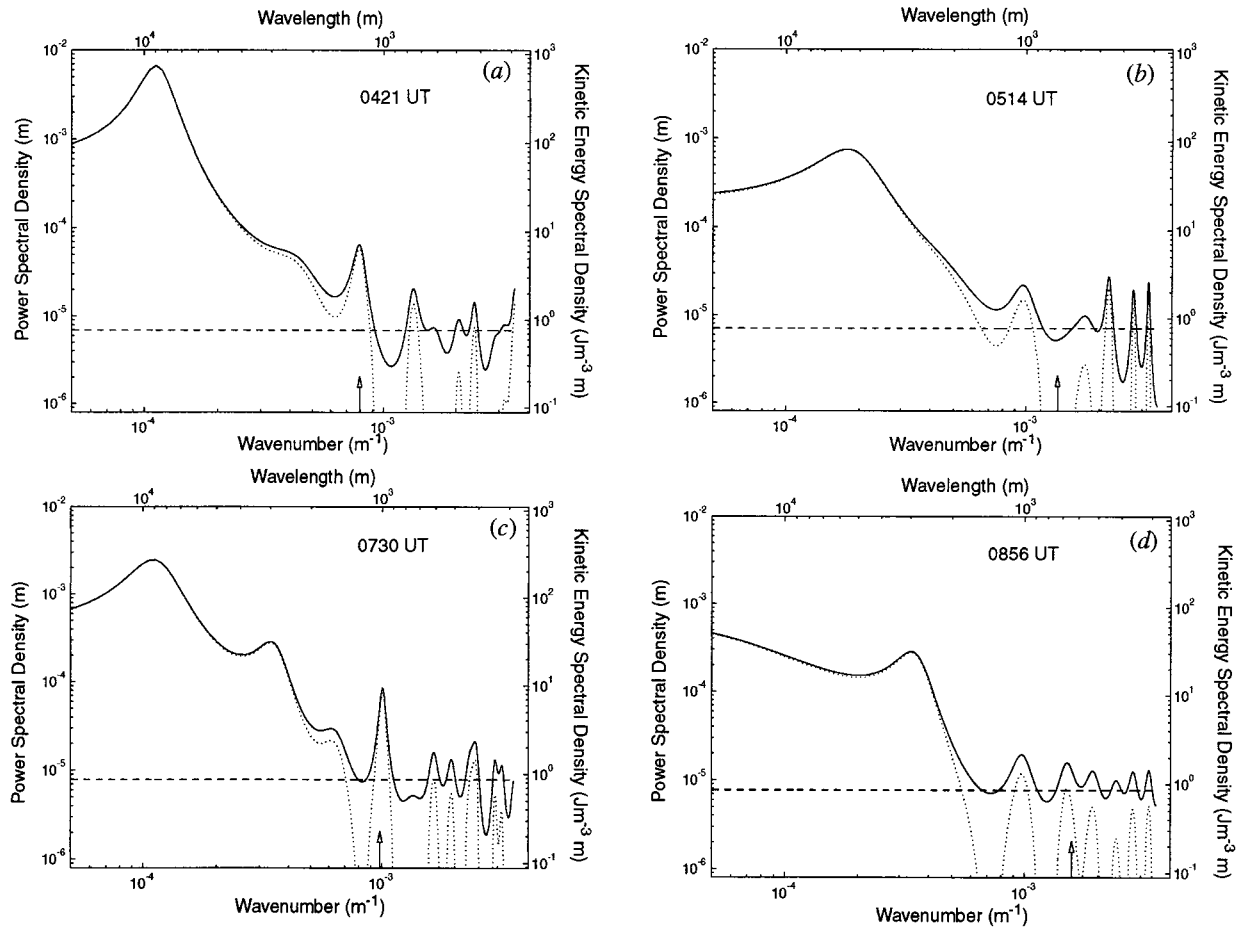
FIG. 11. Selected 1-min AR kinetic energy density PSDs from the spectra shown in Fig. 10 with and without the photon noise floor (indicated by the dashed line).


FIG. 12. The Hines wavenumber as a function of time for the upper-stratospheric L'_{sp} series. Smaller values of the Hines wavenumber mean more of the tail spectrum is nonlinear. Hence, the general trend over the measurement period is for the tail spectrum to become more linear.


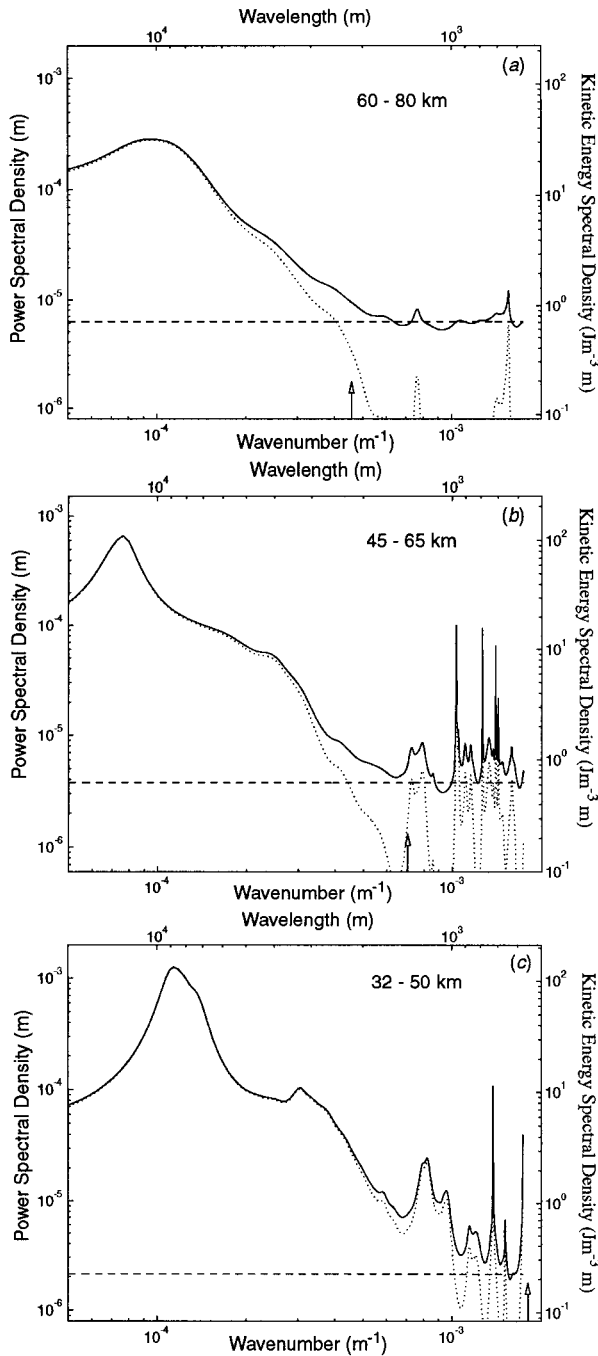FIG. 13. Total power in the L'_{sp} spectra as a function of time.

FIG. 14. The average kinetic energy density for the vertical wave-number spectrum in three height regions, H_{sp}, 60–80 km (a); M_{sp}, 45–65 km (b); and L_{sp}, 32–50 km (c). The mean photon noise floor is indicated on the figure by the dashed line, while the mean value of the Hines wavenumber is indicated by the arrow. The M_{sp} and L_{sp} spectra some spikes at high wavenumber (greater than 1.1×10^{-3} m^{-1} and 1.4×10^{-3} m^{-1} respectively) that are artifacts of the high model order used to allow direct comparison with the H_{sp} spectrum. At more modest model orders these noise spikes (which for any given spike only occur in a single spectrum) would not occur (see, e.g., Figs. 9 and 11 for the L'_{sp} series).


FIG. 15. Ratio of the average vertical wavenumber kinetic energy density PSD for the three height regions shown in Fig. 14: the ratio of the M_{sp} to L_{sp} spectra (solid line), the ratio of the H_{sp} to L_{sp} spectra (dotted line), and the ratio of the H_{sp} to M_{sp} spectra (dashed line). The solid line at unity indicates the level at which the magnitude of the kinetic energy density PSDs are equal.

m^{-1}. The M_{sp} spectrum is shifted to lower wavenumbers relative to the L_{sp} and H_{sp} spectrum, and also has noise features above 1×10^{-3} m^{-1}. The H_{sp} spectrum has a broad low wavenumber peak around 9.5×10^{-5} m^{-1}.

The changes in spectral shape with height are more readily seen by plotting the ratio of the averages (Fig. 15). The ratio of the M_{sp} to L_{sp} spectrum suggests that the spectral energy in near the low-wavenumber peak of the L_{sp} was transferred and shifted to lower wavenumber in the M_{sp} spectrum. Higher wavenumber waves (between about 2.8 and 6×10^{-4} m^{-1}) appear to be lost between the L_{sp} and M_{sp} regions. The ratio of the H_{sp} to M_{sp} spectra show a similar loss of kinetic energy density, between 1.5 and 3.5×10^{-4} m^{-1}. The H_{sp} to M_{sp} ratio also shows the M_{sp} spectrum has more power near the peak, which does not reach the greater heights, as opposed to the low wavenumber behavior of the M_{sp} to L_{sp} ratio. The ratio of the L_{sp} and H_{sp} series shows at wavenumbers higher than the L_{sp} series peak the kinetic energy density has decreased by about a factor of about 2.

Contours of the individual kinetic density spectra are shown for the three data series in Fig. 16 for the AR model. The L_{sp} spectra show the same general features as the L'_{sp} spectra studied earlier. Both low wavenumbers enhancements seen in the L'_{sp} spectra and the L_{sp} spectra are also seen at the greater heights, but a bit later in time. The first increase, from the start of measurements until about 0500 UTC in the L_{sp} spectra continues to 0530 UTC in the M_{sp} spectra, but is not evident in the H_{sp} spectra. The strong "island" in the L_{sp} spectra from 0430 to 0530 UTC at $1/(3200$ m$) = 3.12 \times 10^{-4}$ m^{-1} may be associated with the increase around $1/(6300$ m$) = 1.6 \times 10^{-4}$ m^{-1} in the M_{sp} spectra at approximately

FIG. 16. Contours of the AR kinetic energy density PSDs for the three height regions whose average is shown in Fig. 14. (top) H_{sp}, 60–80 km. The peak value is 95 (J m^{-3}) m. (middle) M_{sp}, 45–65 km. The peak value is 903 (J m^{-3}) m. (bottom) L_{sp}, 32–50 km. The peak value is 580 (J m^{-3}) m. A color version of this figure is available at http://PCL.physics.uwo.ca.

the same time. The second low wavenumber increase between about 0545 and 0800 UTC is seen in all three spectral series, though the onset of the increase seems almost simultaneous between the upper stratosphere and upper mesosphere, with an approximately 30-min delay before the increase in the lower mesosphere (around 0600 UTC).

The behavior of the Hines wavenumber for the three data series shows the L_{sp} has the most variability (Fig. 17). Initially when the kinetic energy density is relatively low for the H_{sp} series the Hines wavenumber is larger, and in fact slightly greater than the M_{sp} series value. After 0500 UTC when the magnitude of the PSD increases the H_{sp} spectra become more nonlinear and the Hines parameter decreases.

b. Power-law exponents (spectral slopes)

Both the spatial and temporal spectra of gravity waves, measured and modeled, are often characterized by a power law, whose exponent in the tail is nominally -3 for vertical wavenumber spectra and $-5/3$ for temporal spectra. As

FIG. 17. Variation of the Hines wavenumber with time for the L_{sp} (circles), M_{sp} (×'s), and H_{sp} (asterisks) data series.

is apparent from the spectra shown in the previous discussion there is a large amount of variation in the vertical wavenumber spectrum with time (and as will be discussed subsequently, also in the temporal spectrum). Where the concept of a power law is useful both heuristically as well as practically for parameterization schemes, the exponent of the power law for measured spectra must be treated with some caution and in fact may be useful only as an independent check on whether the measurements signal-to-noise ratio is adequate.

Since a power law plotted on a log–log scale is a straight line with slope equal to the exponent of the power law, the exponent of the power law will henceforth be called the slope of the spectrum. The slope of the average AR PSD for the four data series over various ranges are listed in Table 2. Slope determinations for all the data series could be reliably made from the peak of the PSD out to 6.00×10^{-4} m^{-1}. The slope increases with height from the L_{sp} to the series H_{sp}. The slope of the L_{sp} series is considerably less than the slope of the L'_{sp} series, primarily due to the better resolved large bandwidth peak around 3.25×10^{-4} m^{-1}. This difference is due to the higher signal-to-noise ratio of the L_{sp} compared to the L'_{sp} data series (Fig. 9 and Fig. 14). More prominent spectral peaks decrease the slope relative to the underlying decrease in the absence of any peaks (e.g., cf. Figs. 11c and 11d).


FIG. 18. Variability in the spectral slope of the L'_{sp} data series. (a) The absolute value of the slope as a function of time. (b) The absolute value of the regression coefficient for the least squares determination of the slope.

Is the increase of slope with height real or an artifact of the signal-to-noise of the residuals? Consider the following two exercises. The first exercise is to look at the variation of the slope of the L'_{sp} spectra as a function of time (Fig. 18a). The average slope of the individual spectra is -3.75 ± 1.64, which is considerably different from the slope of -2.75 for the average of the individual spectra between 1.26×10^{-4} to 6×10^{-4} m^{-1}. Furthermore, the slope does not seem to correlate with power, as during the two high power periods the slope is both large and small and is probably more dependent on the number and intensity of wave features in the spectrum. Neither the error of determination of the slope or the regression coefficient are particularly illuminating

TABLE 2. Power-law exponents and errors of the average AR PSDs for different data series and over different wavenumber ranges.

Data series	Low wavenumber (m^{-1})	High wavenumber (m^{-1})	Slope	Regression coefficient
L'_{sp}	1.26×10^{-4}	3.50×10^{-3}	-2.05 ± 0.03	-0.90
L'_{sp}	1.26×10^{-4}	1.00×10^{-3}	-2.60 ± 0.05	-0.98
L'_{sp}	1.26×10^{-4}	6.00×10^{-4}	-2.75 ± 0.06	-0.98
L_{sp}	1.14×10^{-4}	1.00×10^{-3}	-2.10 ± 0.05	-0.92
L_{sp}	1.14×10^{-4}	6.00×10^{-4}	-2.44 ± 0.06	-0.95
M_{sp}	7.64×10^{-5}	6.00×10^{-4}	-3.06 ± 0.06	-0.97
H_{sp}	9.34×10^{-5}	6.00×10^{-4}	-3.41 ± 0.05	-0.98

TABLE 3. Power-law exponents and errors of the average AR PSDs at fifth order for the L_{sp}, M_{sp}, and H_{sp} data series.

Data series	Low wavenumber (m^{-1})	High wavenumber (m^{-1})	Slope	Regression coefficient
L_{sp}	1.44×10^{-4}	5.00×10^{-4}	-2.57 ± 0.05	-0.98
M_{sp}	1.44×10^{-4}	5.00×10^{-4}	-2.87 ± 0.03	-1.00
H_{sp}	1.44×10^{-4}	5.00×10^{-4}	-2.20 ± 0.02	-1.00

for the fitting of vertical wavenumber spectrum. The larger slopes (i.e., > -3) are always associated with regression coefficients less than -0.9, while the smaller slopes sometimes have lower regression coefficients, as seen in Fig. 18b.

The second exercise involves recomputing the average PSD for the L_{sp}, M_{sp}, and H_{sp} data series using a low-order ($p = 5$) model. Using a low-order model inhibits the ability to find peaks and will thus enhance the ''background'' form of the spectrum. The resulting slopes over identical intervals are listed in Table 3. The M_{sp} spectrum still has a greater slope than the L_{sp} spectrum, but the H_{sp} spectrum now has the shallowest slope. These two exercises may suggest to the skeptic that slope determinations may be as much a matter of opinion as fact.

6. Temporal spectra

The parametric models, as well as statistical signal-processing techniques, were also applied to the time series measured in the stratosphere and mesosphere. In Part 2 of this series these spectra will be used to find energy dissipation and eddy diffusion coefficient profiles (Sica 1999). The generally accepted reason for the shape of the temporal spectrum is a statistical notion, in the sense that the temporal spectrum arises from turbulence (Gage 1979). One might anticipate that MA models, which have a finite bandwidth and can be re-


FIG. 19. Decrease with height (e.g., signal-to-noise ratio) of the AR covariance model order for the temporal spectrum from 31.7 to 39.0 km.

lated to statistical signal processing, would fit the measurements better than AR models (keeping in mind that high-order AR models can in fact be used to determine the MA model coefficients). However, the AR model orders used for both the spatial and temporal spectra are typically one-sixth of the data series length or less, which is not appropriate for MA modeling. Furthermore, if the underlying processes are purely turbulent the bandwidth of the model (related to the model order) would be short, for example, minutes rather than hours, which is not the case. The temporal spectra often show wave features, particularly at higher frequencies. Based on experiments similar to the 4-wave model discussed previously, the AR models are found to be the most appropriate choice for the temporal processing, contrary to our initial intuition.

a. High temporal resolution spectra in the stratosphere

As in the case of the spatial spectrum, the temporal spectrum will first be explored at higher resolution in the stratosphere. The high-resolution temporal series from 31.6- to 38.2-km altitude has 48-m height bins and a sampling interval of 1 min. This height range was chosen to give an adequate signal-to-noise ratio for the eigenvalues at the maximum temporal resolution of the measurements. The time series are spatially smoothed with a Kaiser–Bessel filter with a bandwidth of 1008 m, 33 dB or smaller sidelobes, and a variance decrease of 17.

The optimum order for the high-resolution time series is determined using the single value decomposition described earlier (Fig. 19). The order decreases with signal-to-noise ratio, for example, altitude. Since there are 360 points in each time series the optimum order is less than 5% of the time series length above 33 km, as opposed to traditional ad hoc order choice of one-fourth to one-third of the time series length. However, between 31.6 and 33 km the order increases rapidly, much more rapidly than the signal-to-noise ratio that decreases a factor of roughly 1.25 times per 2 km, suggesting the spectra at the lower heights are extremely structured, as will be evident shortly.

The average AR model spectrum shows features consistent with the correlogram, without the low-frequency rollover of the correlogram (Fig. 20). Both the MUSIC technique and the correlogram show a strong feature at 2.1×10^{-3} s^{-1} [1/(8 min)]. The feature has a narrow bandwidth, which means it corresponds to a relatively

FIG. 20. Average temporal spectrum in the stratosphere (31.7–39.0 km) with the noise floor removed. The AR covariance model at a mean order of 15 is the upper spectrum, which can be compared to the correlogram (the lower spectrum) computed at $N/5$ lags (i.e., 72 lags). The correlogram is divided by 10 for clarity. The arrow indicates the average measured buoyancy frequency.

undamped oscillation. A feature at twice this frequency is also present (and perhaps also at the third harmonic) and is discussed in detail in a study of the high-frequency tail of the temporal spectra (manuscript in preparation).

The underlying power-law behavior of the AR model spectrum continues past the measured buoyancy frequency (3.4×10^{-3} s^{-1}). As discussed previously the slope of the spectrum may be more qualitative than quantitative, but is calculated to demonstrate the consistency with previous long-term averages. In the "monotonic range" from 4.6×10^{-5} s^{-1} to 1.0×10^{-3} s^{-1} the slope is -1.49 ± 0.01, consistent with expectations for the temporal spectrum. As discussed for the slope of the vertical wavenumber spectrum, if the linear fit is made to higher frequencies (say to the buoyancy frequency) the peaks in the spectra lessen the calculated slope.

Figure 21 shows the average spectrum of Fig. 20 in approximately 1-km height bins (e.g., filter bandwidths). The average order of the AR model decreases from 42 at the lowest level to 12 at the top level. The buoyancy frequency is approximately constant with height and is indicated by a vertical average on the figure. Several interesting peaks are worth noting. The 32.2-km spectrum has a broadband peak around 4.0×10^{-4} s^{-1} [$1/(42$ min)]. A "hump" around 9×10^{-4} s^{-1} [$1/(19$ min)] is evident at the averages centered at 33.3 and 34.4 km. The intense feature at 2.1×10^{-3} s^{-1} decreases in intensity while increasing in bandwidth, until it disappears at the 37.7-km height bin. The increase in bandwidth of the feature means the oscillation is more strongly damped with height. All the spectra show power above the buoyancy frequency. The spectra at 33.3, 34.4, and 37.7 km show significant power decreases at the buoyancy frequency, but also some structure above. The signal-to-noise ratio at the latter two heights is sufficiently high that the power excess above the buoyancy frequency is real.

The MUSIC algorithm was used to attempt to localize individual peaks in the temporal spectrum. A height contour of this calculation using a conservative order number is shown in Fig. 22. At the lowest frequencies strong wave features appear around 35 km at $1/(298$ min) $= 5.6 \times 10^{-5}$ s^{-1} and $1/(139$ min) $= 1.2 \times 10^{-4}$ s^{-1}. Below 36 km the aforementioned peak at $1/(8$ min) $= 2.1 \times 10^{-3}$ s^{-1} is extremely clear, as is a less intense, broader-bandwidth feature at $1/(84$ min) $= 2.0 \times 10^{-4}$ s^{-1}. Possible harmonics of the feature at $1/(8$ min) are evident below 33 km. The MUSIC results provide independent support for a combination of a background "unstructured" spectrum with imposed wave features, since the MUSIC algorithm is most sensitive to wave features.

b. Temporal spectra from 32 to 80 km

The computed temporal spectra in the 31.6–81-km height region are averaged in five groups shown in Table 4. These groups were selected to give approximately the same noise floor for each average spectrum. The average spectrum for each group is shown in Fig. 23. The spectral average at 60 km has a sharp feature at 8.25×10^{-5} s^{-1} [$1/(202$ min)] and is a persistent feature in both the parametric models and the correlogram analysis at various noise floor levels. The sharp feature at 2.1×10^{-3} s^{-1} [$1/(8$ min)] in the 35-km average is also present as a broadband feature in the 42-km average. Since the broader bandwidth is associated with stronger wave damping perhaps this wave is becoming saturated in the upper stratosphere.

The amplitude of the average spectra are consistent with the spatial spectra shown in Fig. 14. At low frequencies the power is similar in the stratosphere, then decreases by about a factor of 5 in the mesosphere (not accounting for the low-frequency peak in the 60-km spectrum). The spectra slopes show no clear variation with height (Table 5). Rather the spectra at low frequencies {i.e., from the bandwidth of the measurements out to 3×10^{-4} s^{-1} [$1/(56$ min)]} appear to have a relatively steep slope (i.e., -2 to -2.5), while out to higher frequencies, structure in the temporal spectral causes the fitted slope to vary between -1 and -2. The 60-km spectrum is anomalously steep due to the strong low-frequency feature. What can be concluded from this exercise is that the underlying slope of the process appears to be about -2.

7. Discussion

Correlogram analysis of the spatial and temporal spectrum show a general agreement with previous re-
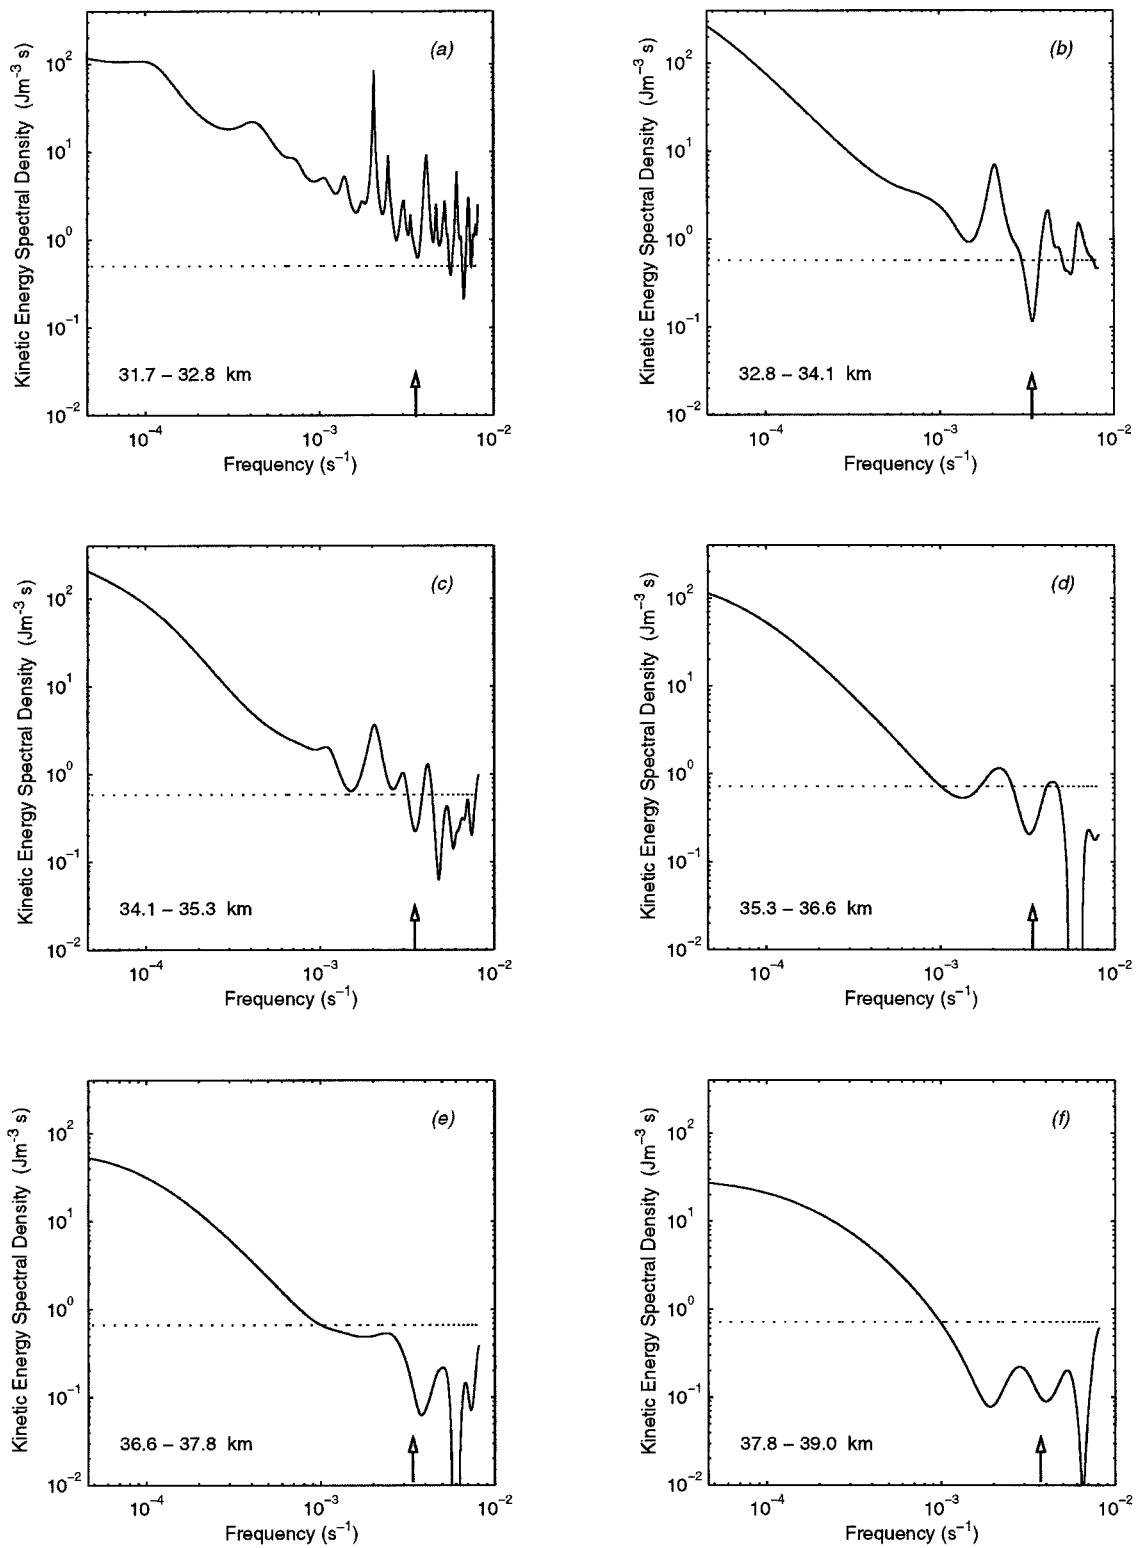
FIG. 21. Subsets of the average temporal spectrum shown in Fig. 20 in approximately 1-km height regions as indicated on the figure. The noise floor and average buoyancy frequency are also indicated in each panel.

FIG. 22. Contour plot of the variation of the temporal spectra with height using the MUSIC algorithm. The intense feature at $1/(8 \text{ min}) = 2.1 \times 10^{-3} \text{ s}^{-1}$ is clearly evident, as are possible second and third harmonics of this feature. A color version of this figure is available at http:// PCL.physics.uwo.ca.

ports of the gravity wave spectrum, that is, the appropriate slopes and magnitudes, as do the parametric models introduced and investigated. While statistical signal processing techniques, such as the correlogram, may be sufficient for average representations of the spectrum, the high-resolution techniques used in this study show the gravity wave spectrum to be extremely rich in variability and detail.

The combination of measurements of density fluctuations from a large power-aperture lidar system and high-resolution-parametric-modeling techniques unveil a wealth of features in the spatial and temporal spectra. As window closing is the "art" of the correlogram process, order selection is the art of parametric modeling. Single value decomposition of the eigenvalues of the data matrix is used to make the order selection as quantitative as possible. Hence, the variation of order as a function of time is an indication, in part, of the underlying process's complexity and intermittency.

The measurements are best represented by AR rather than MA or ARMA models. The order of the AR model is typically relatively low, that is less than 1/10 of the data series length, and argues against the representation of the measurements as a MA process, which can be related to high-order AR processes. Since the MA models can be related to classical statistical time series methods such as the correlogram and the periodogram, this result implies that these classic techniques are not best suited for the study of the gravity wave spectrum at high temporal–spatial resolution. Gravity waves introduce strong correlations in height and time, which are consistent with AR models rather than statistical models, which may be better suited for studying processes due to turbulence. The measurements on this night suggest that this result is also true for the temporal as well as the vertical wavenumber spectrum, at least out to the buoyancy frequency.

This variability in time for the vertical wavenumber

TABLE 4. Temporal time series specifics. A 21-point Kaiser–Bessel spatial filter was applied to each time series.

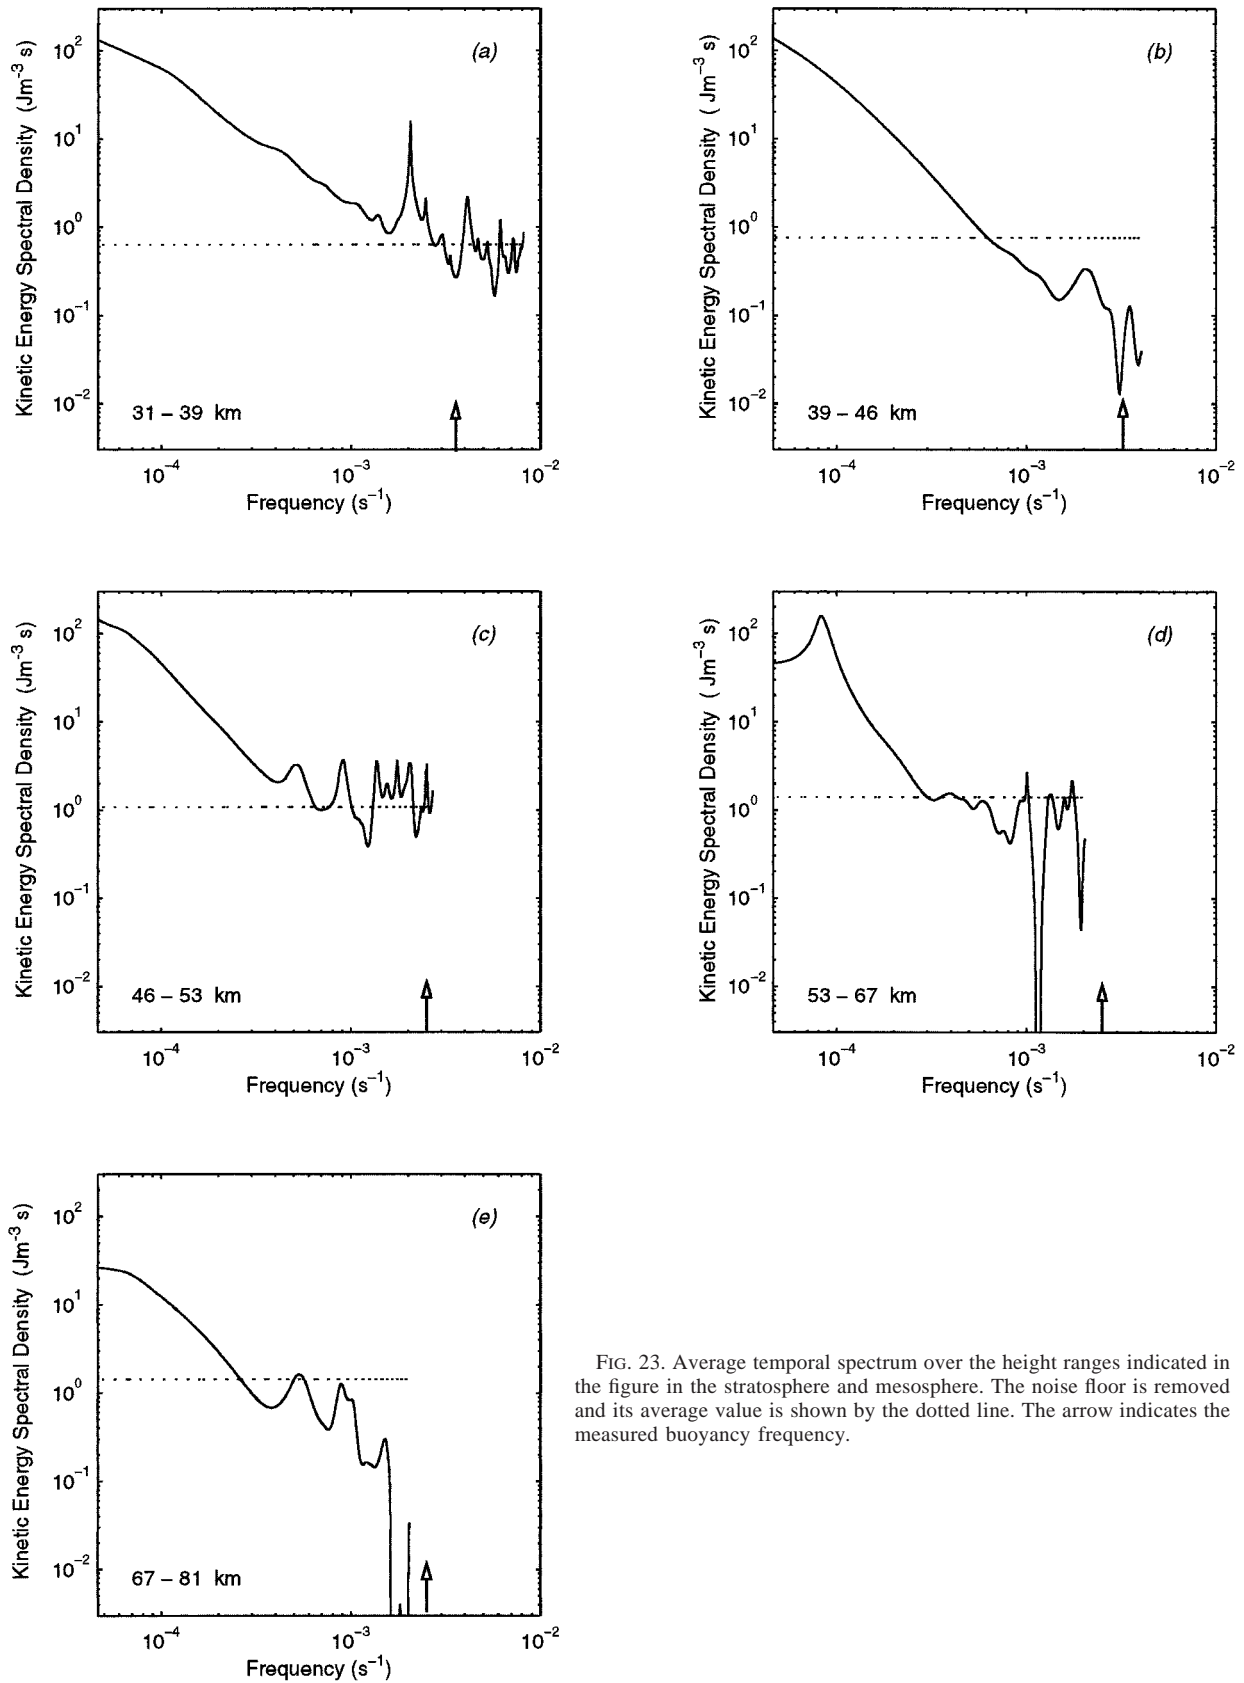| Time series label	Height range (km)	Height resolution (m)	Spatial filter bandwidth (m)	Temporal resolution (min)
35 km	31.6–39	48	1008	1
42 km	39–46	72	1512	2
50 km	46–53	96	2016	3
60 km	53–67	120	2520	4
74 km	67–81	120	2520	4

FIG. 23. Average temporal spectrum over the height ranges indicated in the figure in the stratosphere and mesosphere. The noise floor is removed and its average value is shown by the dotted line. The arrow indicates the measured buoyancy frequency.

TABLE 5. Slopes of the temporal time series. The fits begin at 4.62 $\times$ 10^{-5} s^{-1} and end as indicated in the table.

Time series label	Fit end (s^{-1})	Slope	Regression coefficient
35 km	1.00×10^{-3}	-1.49 ± 0.01	-1.00
42 km	1.00×10^{-3}	-2.12 ± 0.01	-1.00
50 km	1.00×10^{-3}	-1.44 ± 0.02	-0.91
50 km	3.00×10^{-4}	-2.18 ± 0.01	-1.00
60 km	1.00×10^{-3}	-1.60 ± 0.02	-0.90
60 km	3.00×10^{-4}	-2.69 ± 0.06	-0.95
74 km	1.00×10^{-3}	-1.18 ± 0.02	-0.86
74 km	3.00×10^{-4}	-2.00 ± 0.02	-0.99

spectrum is illustrated by the variations in the Hines wavenumber as well as the total power density changes of the spectrum, which are as large as a 10 μW m^{-3}. As a function of height the vertical wavenumber spectrum gained some power at wavenumbers greater than 1×10^{-4} m^{-1}, but lost significant amounts of power at larger wavenumbers between the upper stratosphere and the lower mesosphere. This loss is less between the lower and upper mesosphere. A similar energy loss with height occurs in the temporal spectrum.

Intermittency is the key, as the spectral shape can rapidly vary with time and height. Contour maps of the individual spectral determinations show the variability in the underlying processes. At times the spectrum is relatively smooth, at other times it is weak and variable, and at times features 20–30 dB above the background smooth spectrum develop and persist, sometimes for tens of minutes or more.

Parametric models of the spectrum yield slopes that are very similar to the correlogram method and consistent with previous observations. It was demonstrated that the concept of slope, while useful both heuristically and for practical gravity wave parameterization schemes, is a dynamic quantity that shows large variations in space and time. The large number of features present at most times preclude any definitive slope determination; rather the features ride on a background that has a slope of nominally -2 to -3 for the vertical wavenumber spectra and about -2 for the temporal spectra. At low wavenumbers and long-time integrations there may be changes of the slope with height for the vertical wavenumber and temporal spectrum. Such a variation of slope with height has been measured by Beatty et al. (1992) using Rayleigh-scatter and sodium resonance fluorescence lidar techniques from a low-latitude station (Arecibo, Puerto Rico) during late winter and early spring.

8. Conclusions and future directions

Applying parametric models to high temporal–spatial resolution Rayleigh-scatter lidar measurements of density fluctuations in the middle atmosphere have led to the following results.

1) The nature of the autoregressive models is one of correlations between neighboring measurements, as would be caused in the atmosphere by wave activity. The nature of the moving average process is one of noise driven through an atmospheric filter, akin to the process of turbulent dissipation of energy. Both the temporal and vertical wavenumber spectra are best represented by relatively low-order autoregressive models. The order of the parametric model can be selected by eigenvalue analysis of both the signal and the photon noise. The underlying spectral form is similar between the parametric models and the standard correlogram method.

2) The spatial–temporal behavior of the spectra is highly variable with numerous intermittent and intense features rising well above the photon noise floor of the measurements. The intermittency is apparent from the spectra at timescales from a few minutes to an hour or more of averaging. The parameter suggested by Hines to indicate the transition from a linear to nonlinear vertical wavenumber spectrum is also highly variable with time.

3) Both the total power and power spectral density of the spectra changed rapidly, with variations at low frequencies of over a factor of 100. The vertical wavenumber spectrum showed several strong intensifications, particularly at low wavenumbers over the observing period. The variations of the temporal spectra will be used in the following paper (Sica 1999) to estimate gravity wave energy dissipation and eddy diffusion.

4) When the power spectral densities are transformed to kinetic energy density spectra using the gravity wave polarization relations and the measured density and temperature, a decrease in kinetic energy with height is evident for both the vertical wavenumber and temporal spectra. On this night, most of this decrease occurs between the upper stratosphere and the lower mesosphere.

5) The slope of the spatial and temporal spectra could be determined with high statistical precision, though the slope obtained is extremely sensitive to averaging and, hence, the signal-to-noise ratio of the measurements. The vertical spectrum has a slope of between -2 and -3, while the temporal spectrum has a slope around -2. With appropriate averaging, the vertical wavenumber and temporal spectra have slopes that vary with height.

Now that a tool box has been developed and characterized in detail for a single night's measurements the next step is to apply these techniques to the over 200 nights of measurements already obtained by the Purple Crow Rayleigh-scatter lidar. In the second paper of this series the parametric modeling techniques will be used to compute profiles of energy dissipation and eddy diffusion due to gravity waves. Studies are under way to

investigate the seasonal variations of the spectra as well as variations of the spectra due to inversion layers.

REFERENCES

Beatty, T. J., C. A. Hostetler, and C. S. Gardner, 1992: Lidar observations of gravity waves and their spectra near the mesopause and stratopause at Arecibo. *J. Atmos. Sci.,* **49,** 477–495.

Blackman, R. B., and J. W. Tukey, 1958: *The Measurement of Power Spectra.* Dover Publications, 190 pp.

Blix, T. A., E. V. Thrane, and Ø. Andreassen, 1990: In situ measurements of the fine-scale structure and turbulence in the mesosphere and lower thermosphere by means of electrostatic positive ion probes. *J. Geophys. Res.,* **95,** 5533–5548.

Bloomfield, P., 1976: *Fourier Analysis of Time Series: An Introduction.* John Wiley and Sons, 258 pp.

Dewan, E. M., and R. E. Good, 1986: Saturation and the ''universal'' spectrum for vertical profiles of horizontal scalar winds in the atmosphere. *J. Geophys. Res.,* **91,** 2742–2748.

Fritts, D. C., 1985: A numerical study of gravity wave saturation: Nonlinear and multiple wave effects. *J. Atmos. Sci.,* **42,** 2043–2058.

——, T. Tsuda, T. Sato, S. Fukao, and S. Kato, 1988: Observational evidence of a saturated gravity wave spectrum in the troposphere and lower stratosphere. *J. Atmos. Sci.,* **45,** 1741–1759.

Gage, K. S., 1979: Evidence for a $k^{5/3}$ power law inertial range in mesoscale two-dimensional turbulence. *J. Atmos. Sci.,* **36,** 1950–1954.

Gardner, C. S., 1994: Diffusive filtering theory of gravity wave spectra in the atmosphere. *J. Geophys. Res.,* **99,** 20 601–20 622.

——, M. S. Miller, and C. H. Liu, 1989a: Rayleigh lidar observations of gravity wave activity in the upper stratosphere at Urbana, Illinois. *J. Atmos. Sci.,* **46,** 1838–1854.

——, D. C. Senft, T. J. Beatty, R. E. Bills, and C. A. Hostetler, 1989b: Rayleigh and sodium lidar techniques for measuring middle atmosphere density, temperature, and wind perturbations and their spectra. *World Ionosphere/Thermosphere Study Handbook,* C. H. Liu, Ed., Vol. 2, ICSU Scientific Committee on Solar Terrestrial Physics, 148–187.

——, C. A. Hostetler, and S. Lintelman, 1993: Influence of the mean wind field on the separability of atmospheric perturbation spectra. *J. Geophys. Res.,* **98,** 8859–8872.

Hamilton, K., R. J. Wilson, J. D. Mahlman, and L. J. Umscheid, 1995: Climatology of the SKYHI troposphere–stratosphere–mesosphere general circulation model. *J. Atmos. Sci.,* **52,** 5–43.

Hamming, R. W., 1977: Design of nonrecursive filters. *Digital Filters,* Prentice-Hall, 95–117.

Hines, C. O., 1960: Internal atmospheric gravity waves at ionospheric heights. *Can. J. Phys.,* **38,** 1441–1481.

——, 1991: The saturation of gravity waves in the middle atmosphere. Part II: Development of Doppler-spread theory. *J. Atmos. Sci.,* **48,** 1360–1379.

Jenkins, G. M., and D. G. Watts, 1968: Introduction to time series analysis. *Spectral Analysis and Its Applications,* Holden-Day, 140–208.

Manson, A. H., 1990: Gravity wave horizontal and vertical wavelengths: An update of measurements in the mesopause region (80–100 km). *J. Atmos. Sci.,* **47,** 2765–2773.

Marple, S. L., 1987: *Digital Spectral Analysis.* Prentice-Hall, Inc., 492 pp.

Sica, R. J., 1999: Measurements of the effects of gravity waves in the middle atmosphere using parametric models of density fluctuations. Part II: Energy dissipation and eddy diffusion. *J. Atmos. Sci.,* **56,** 1330–1343.

——, and M. D. Thorsley, 1996: Measurements of superadiabatic lapse rates in the middle atmosphere. *Geophys. Res. Lett.,* **23,** 2797–2800.

——, S. Sargoytchev, P. S. Argall, E. F. Borra, L. Girard, C. T. Sparrow, and S. Flatt, 1995: Lidar measurements taken with a large-aperture liquid mirror. 1. Rayleigh-scatter system. *Appl. Opt.,* **34,** 6925–6936.

Smith, S. A., D. C. Fritts, and T. E. Van Zandt, 1987: Evidence for a saturated spectrum of atmospheric gravity waves. *J. Atmos. Sci.,* **44,** 1404–1410.

Tsuda, T., T. Inoue, D. C. Fritts, T. E. van Zandt, S. Kato, T. Sato, and S. Fukao, 1989: MST radar observations of a saturated gravity wave spectrum. *J. Atmos. Sci.,* **46,** 2440–2447.

Weinstock, J., 1990: Saturated and unsaturated spectra of gravity waves and scale-dependent diffusion. *J. Atmos. Sci.,* **47,** 2211–2225.

Wilson, R., M. L. Chanin, and A. Hauchecorne, 1991: Gravity waves in the middle atmosphere observed by Rayleigh lidars. Part 2: Climatology. *J. Geophys. Res.,* **96,** 5169–5183.