7

ACQUISITION AND PROCESSING OF ATMOSPHERIC BOUNDARY LAYER DATA

Much of what we know about the structure of the boundary layer is empirical, the result of painstaking analysis of observational data. As our understanding of the boundary layer evolved, so did our ability to define more clearly the requirements for sensing atmospheric variables and for processing that information. Decisions regarding choice of sampling rates, averaging time, detrending, ways to minimize aliasing, and so on, became easier to make. We find we can even standardize most procedures for real-time processing. The smaller, faster computers, now within the reach of most boundary layer scientists, offer virtually unlimited possibilities for processing and displaying results even as an experiment is progressing.

The information we seek, for the most part, falls into two groups: (1) time-averaged statistics such as the mean, variance, covariance, skewness, and kurtosis and (2) spectra and cospectra of velocity components and scalars such as temperature and humidity. We discuss them separately because of different sampling and processing requirements for the two. A proper understanding of these requirements is essential for the successful planning of any experiment.

In this chapter we discuss these considerations in some detail with examples of methods used in earlier applications. We will assume that sensors collecting the data have adequate frequency response, precision, and long-term stability and that the sampling is performed digitally at equally spaced intervals. We also assume that the observation heights are chosen with due regard to sensor response and terrain roughness.

7.1 Time-averaged statistics

For calculations of means and higher order moments we need time series that are long enough to include all the relevant low-frequency contributions to the process, sampled at rates fast enough to capture all the high-frequency contributions the sensors are able to measure. Improper choices of averaging times and sampling rates can indeed compromise our statistics. We need to understand how those two

factors affect our measurements in order to make sensible decisions on how long and how fast to sample.

7.1.1 *Choice of averaging time*

Ideally, the statistical quantities we would use to describe the properties of turbulent flow would be those obtained by "ensemble" averaging (i.e., averaging over many realizations under identical conditions). In reality we are forced to describe them in terms of averages over time, making the "ergodic" hypothesis that time averages are, for all practical purposes, equivalent to ensemble averages. For this assumption to be true, the fluctuations must be statistically stationary over the period chosen for analysis. Only then can we justify the application of the Reynolds averaging rules (Appendix 7.1), which specify conditions the data must satisfy for proper separation of the fluctuating components from the mean.

Statistical stationarity of a time series $\alpha(t)$ demands that variances and covariances approach stable values as the averaging time T is increased, which implies that an integral time scale $\mathcal{T}_\alpha$ exists for $\alpha(t)$. The requirements for averaging time T with $T \gg \mathcal{T}_\alpha$ can then be expressed in terms of $\sigma^2_{\overline{\alpha}}$, the variance of the measured time mean $\overline{\alpha}$ about the expected ensemble mean, and σ^2_α, the ensemble variance of α. From Lumley and Panofsky (1964), we have

$$\sigma^2_{\overline{\alpha}} \simeq \frac{2\sigma^2_\alpha \mathcal{T}_\alpha}{T}. \tag{7.1}$$

For averaging times $T \gg \mathcal{T}_\alpha, \sigma^2_{\overline{\alpha}}$ becomes negligible. We can now specify an acceptable level of error $\sigma_{\overline{\alpha}}/\overline{\alpha}$ (say ϵ) and express the averaging time required to keep $\sigma_{\overline{\alpha}}/\overline{\alpha}$ within that level as

$$T \simeq \frac{2\sigma^2_\alpha \mathcal{T}_\alpha}{\overline{\alpha}^2 \epsilon^2}. \tag{7.2}$$

In order to use (7.2) to estimate suitable averaging times, we have to replace the ensemble variance σ^2_α by the more accessible time average variance—a further application of the ergodic hypothesis. For typical daytime values of $\sigma_u = 1$ m s^{-1}, $\mathcal{T}_u = 10$ s, and $\overline{u} = 5$ m s^{-1} and specifying $\sigma_{\overline{u}} = 0.1$ m s^{-1} (i.e., $\epsilon = 0.02$), we get $T = 2000$ s $\simeq 30$ min, a reasonable averaging time for mean horizontal winds, corresponding to the passage of two or three of the large convective cells that extend through the depth of the CBL. Lumley and Panofsky (1964) offer a clear discussion of ergodicity and averaging requirements. Using the expression (7.2), they show how T increases with the order of the moment, estimating that it takes roughly five times longer to measure the fourth moment to the same accuracy as the second moment, assuming $\mathcal{T}_\alpha$ is the same for both (here α represents the higher powers or products of measured variables). The reason for this is the tendency for $\sigma_\alpha/\overline{\alpha}$ to increase with the order of the moment. The

averaging time requirement can also be different for different second moments, depending on the magnitude of $\sigma_\alpha/\overline{\alpha}$.

If the averaging period can be increased indefinitely, ϵ can be kept to within negligible levels. $T \simeq 1$ h is, however, about as long as we can extend the averaging period without encountering nonstationarity in the form of diurnal variations in surface heat flux and boundary layer depth. With an averaging period this long we can expect errors on the order of 3.5%, 5%, and 10%–50% for near-surface σ_w^2, $\overline{w'\theta'}$, and $\overline{u'w'}$, respectively (Haugen, 1978).

7.1.2 *Choice of sampling rate*

Having established a suitable averaging time based on the accuracies we desire in the results, we turn to the next important consideration: how frequently to sample the signal. We assume the sampling process itself is virtually instantaneous, that is, the sensors are at least fast enough to respond to all the frequencies contributing to the process. (In the boundary layer, this means a high-frequency cutoff somewhere in the inertial subrange where spectral energy is falling off rapidly.) The conservative approach would be to sample at a rate commensurate with the frequency response of the sensors. But this is really not necessary because the accuracy in our statistics should, under stationary conditions, be a function only of the number of samples, N.

For equally spaced instantaneous samples with sampling interval $\Delta t > \mathcal{T}_\alpha$, the error variance can be expressed in a form similar to (7.1):

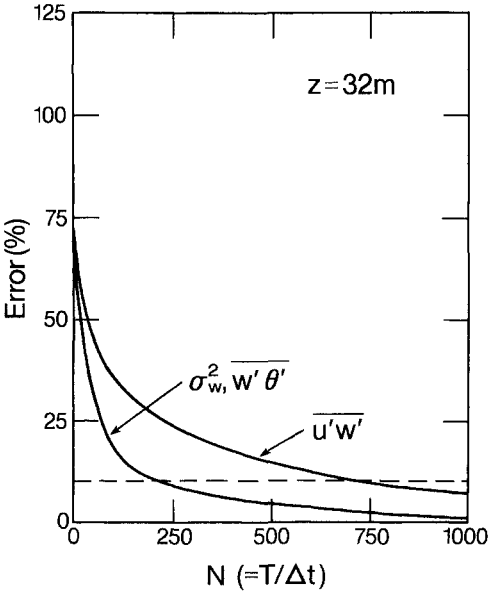

FIG. 7.1 Percentage error in calculated variances and fluxes as functions of the number of discrete samples used (after Haugen, 1978).

$$\sigma_{\overline{\alpha}}^2 = \frac{\sigma_\alpha^2}{N}, \tag{7.3}$$

where $N = T/\Delta t$. As with T in (7.2), we can estimate the number of samples needed to maintain a given level of confidence in $\overline{\alpha}$. For $\sigma_u = 1$ m s^{-1}, we need $N = 100$ to ensure that $\sigma_{\overline{u}} \leq 0.1$ m s^{-1}. For $T = 2000$ s, calculated from (7.2) for the same set of conditions, this value of N suggests a sampling interval $\Delta t = 20$ s. Here too, it can be shown that the number of samples needed to keep ϵ within desired limits increases with the order of the moment.

The results of Haugen's (1978) calculations, summarized in Fig. 7.1, confirm the dependence of ϵ on N. Using various combinations of T ($\gg \mathcal{T}_\alpha$) and Δt, he derived empirical curves from field data that show the number of samples for 10% accuracy in σ_w^2 and $\overline{w'\theta'}$ to be 200 and in $\overline{u'w'}$ to be 750. The implication is that Δt can be increased without loss of accuracy as long as T is increased proportionally to maintain N at the same level. Note that we still require the sampling to be near instantaneous. The accuracy in the moment calculations depends on our ability to capture, at least occasionally, the peaks and valleys in the signal, as illustrated in Fig. 7.2. In the frequency domain, the process is explained in terms of spectral folding (see Section 7.2).

7.2 Spectra and cospectra

Conversion of data collected in the form of equally spaced samples in the time domain to spectral and cospectral estimates in the frequency domain involves a series of calculations we refer to as spectrum analysis or Fourier analysis. Several methods for accomplishing this conversion exist, but the one most widely used today is the fast Fourier transform (FFT), which accepts data points in groups of powers of 2 (e.g., $N = 2^{10} = 1024$) and calculates the Fourier components in an elegant, time-efficient manner ideally suited to low-level computer assembly language. [For details on FFT processing see Bendat and Piersol (1971) and

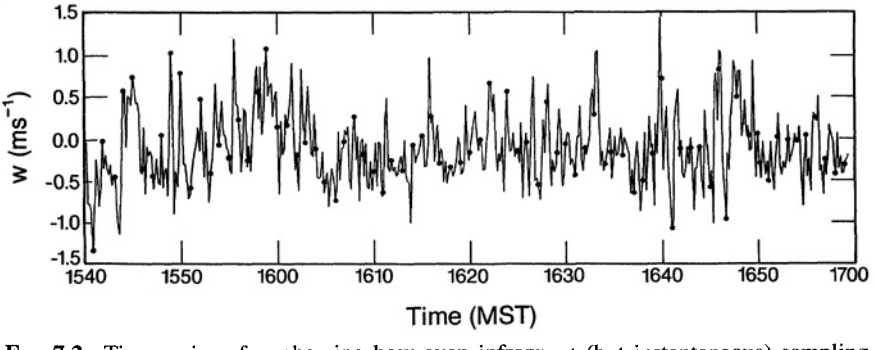

FIG. 7.2 Time series of w showing how even infrequent (but instantaneous) sampling (represented by the dots) can capture enough values at the peaks and valleys to approximate the true variance. The time indicated is Mountain Standard Time (MST).

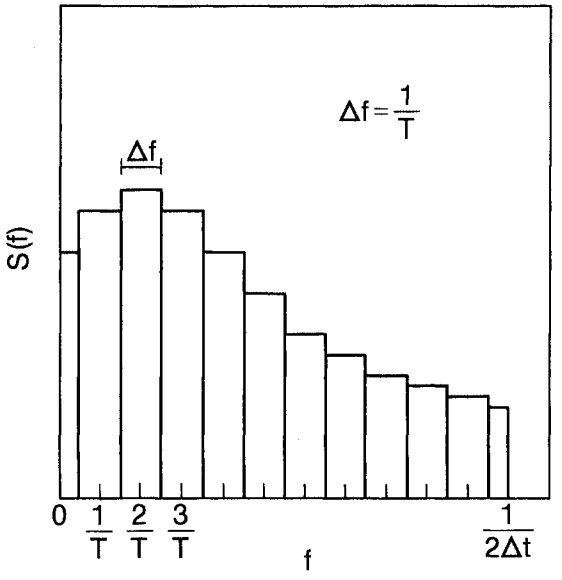

FIG. 7.3 Averaging period T and its relationship to spectral resolution and to the center frequency of spectral estimates.

Ramirez (1985).] The choices we make for T and Δt have particular relevance for FFT analysis because they directly influence the range and resolution of the spectral information the analysis provides. We discuss below the implications of those choices.

7.2.1 *Choice of record length*

For spectrum analysis we usually require record lengths somewhat longer than for variance calculations because of the clearer definition they promise in the energy peak and the rolloff on the low-frequency side of the spectrum. The selection of T automatically establishes the lowest frequency ($1/T$) we can resolve as well as the width of each elementary frequency band Δf in the spectrum output:

$$\Delta f = \frac{1}{T} = \frac{1}{N\Delta t}. \tag{7.4}$$

The elementary bands (excluding the first and last, which are only $\Delta f/2$ wide) are centered on frequencies $1/T, 2/T, 3/T \ldots, (1/2\Delta t - 1/T)$, as shown in Fig. 7.3. The number of spectral (or cospectral) estimates we obtain from an FFT analysis[1] is $N/2$, which gives us a total bandwidth of $1/2\Delta t$ in our spectral data.

[1]The FFT yields $N/2$ estimates that are complex numbers, so there is no loss of information. The squared magnitudes of the spectral estimates comprise the power spectrum, whereas the phase information is used in forming the cospectra and quadrature spectra (Appendix 2.1, Chapter 2).

With spectra (and cospectra) it is easy to determine whether the chosen T is long enough. A well-defined spectral peak in the $fS(f)$ [or $fC(f)$] plot and a rolloff over a decade or so on the low-frequency side are indications of an adequate T and a reasonably good approach to stationarity. To achieve this for wind and temperature spectra over most of the ABL requires averaging times of 60–90 min, which often conflict with nonstationarity caused by diurnal variations. If there is no evidence of a peak or even a leveling of the spectrum (or cospectrum) at the low end, the time series should be examined, even plotted out, to determine if T should be increased (to include a significant peak that was not originally anticipated) or if the data need to be filtered (to remove a long-term trend that is masking a real peak within the existing spectral range). These questions and more will be discussed in the next section dealing with the preparation of time series for spectrum analysis.

7.2.2 *Choice of sampling rate*

Here, the bandwidth desired in the calculated spectra and cospectra determines the rate at which the variables should be sampled. The decision is simple if the spectrum of the variable has a clearly defined cutoff frequency f_c. Shannon's sampling theorem tells us that at least two samples per cycle are needed to define a frequency component in the original signal completely; for this band-limited signal the sampling interval Δt should be

$$\Delta t = \frac{1}{2f_c}. \tag{7.5}$$

Sampling at this rate is referred to as "critical sampling." Sampling more frequently ($\Delta t < 1/2f_c$) yields no additional benefits[2], since no useful spectral information exists between f_c and $1/2\Delta t$. Sampling less frequently ($\Delta t > 1/2f_c$) yields a spectral bandwidth inadequate to cover the full range of frequencies contained in the signal. The spectral information at frequencies above $1/2\Delta t$ is not lost but folded back into the spectrum in an accordion fold, as shown in Fig. 7.4. This process is called "aliasing," since it results in high-frequency components emerging as low-frequency components within the resolved bandwidth. The frequency $f_0(= 1/2\Delta t)$ is referred to as the folding frequency or Nyquist frequency.

We rarely find sharp cutoffs in boundary layer turbulence. The inertial subrange, however, presents a logical place to establish the cutoff for data acquisition purposes. There are three benefits:

1. Being in a region of rapidly falling spectral energy, the folding is confined to no more than the first fold and its effects seldom extend below $0.5f_0$ (Fig. 7.5).

[2]Note that, in theory, (7.5) applies to an infinite record. With a finite record we must sample more frequently to properly resolve f_c. In practice, this would only be a problem for very short records.

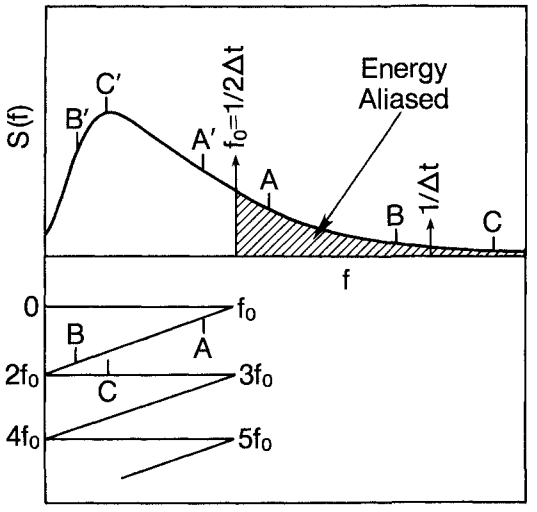

FIG. 7.4 Folding of spectral energy in an aliased spectrum. Frequencies A, B, and C appear at A', B', and C' in the resolved spectrum.

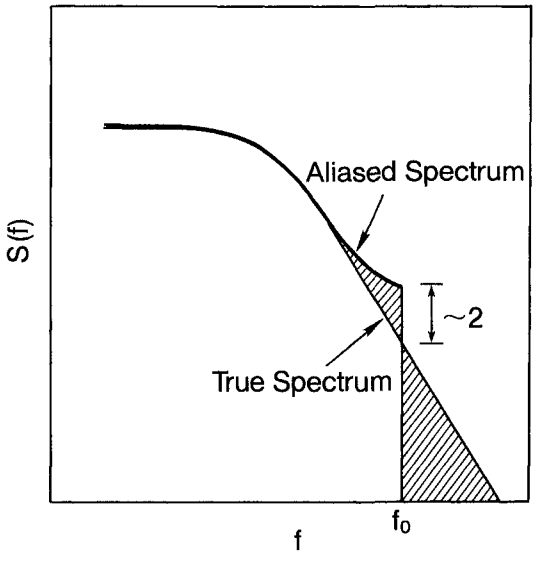

FIG. 7.5 Aliased energy shown raising spectral levels at the Nyquist frequency by a factor of 2. (Log scales on both axes.)

2. With the slope of the spectrum known, the true spectrum can be approximated by simple extrapolation. In the absence of any high-frequency noise, the spectrum at f_0 is raised by a known factor (~ 2).
3. No significant transport of momentum and heat occurs in this region of the spectrum. Therefore, a cutoff here is ideal from the point of view of flux measurements.

The ideal choice for f_0 is one that would reveal at least two octaves of the inertial subrange (allowing for aliasing errors) to provide a basis for extrapolation. This implies $f_0 > 4f_i$, where $f_i(\sim 2\overline{u}/z$, when M-O scaling applies) is the low-frequency limit of the inertial subrange and $\overline{u}$ is the mean wind speed at height z. This criterion is often stretched for measurements close to the ground. (In that case, particular care should be taken to minimize aliasing by using methods discussed in the following section.) The sampling frequency for turbulence data over reasonably flat terrain is usually set at 10 or 20 Hz depending on whether the observation height is above or below 5 m.

In selecting the sampling frequency for turbulence sensors, it is wise to consider where the power line frequency (50 or 60 Hz) would fall within the spectrum. Since line frequency contamination is present to some extent in most analog information transmitted over signal cables, it is important to place the aliased line frequency where it will do the least harm, that is, at $f = 0$. For $\Delta t = 1/10$ s, both 50 and 60 Hz frequencies should fold back to $f = 0$. The choice of $\Delta t = 1/20$ s, although ideal for the 60-Hz line frequency, would, however, place the 50-Hz line frequency at $f = f_0$ as seen in Fig. 7.6. Hence, the logical choice for Δt in countries with 50-Hz line frequency would be 1/25 s in applications where sampling at 1/10 s is not fast enough.

7.3 Preparing data for spectrum analysis

Measurements in the real atmosphere may not always satisfy our definitions of stationarity, or the sampling rate may not be fast enough to avoid serious aliasing

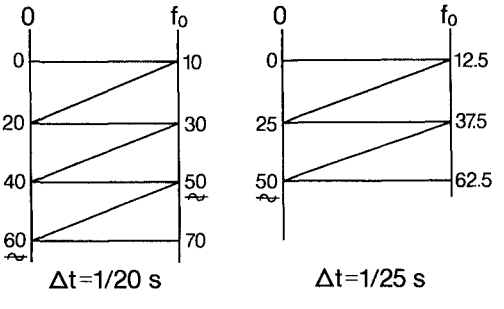

FIG. 7.6 Aliasing of power line frequency for 20- and 25-Hz sampling rates. Line frequencies are given in Hz.

in the part of the inertial subrange we are able to cover. These deficiencies can be met through proper intervention before the spectrum analysis. Some of the techniques most commonly used to minimize their effects are discussed below.

7.3.1 *Reduction of aliasing*

Since the very act of sampling at discrete intervals causes energy above frequency f_0 to be aliased back to lower frequencies, any scheme to minimize aliasing has to be applied very early in the data handling stage. There are two simple approaches; both depend on low-pass filtering to reduce the energy above f_0 that is available for aliasing. Both can be incorporated in the sensor electronics.

1. *Analog prefiltering.* Before digital conversion, the analog signal is low-pass filtered with an analog filter whose half-power point is set at $f_0 = 1/2\Delta t$. This will reduce the energy at the Nyquist frequency f_0 by 50%. When this filtered signal is sampled at interval Δt, aliasing restores the energy to nearly its true value at f_0 (assuming energy only in the first fold counts) and to slightly lower values in the range $0.5f_0 < f < f_0$ as shown in Fig. 7.7. The steeper the filter rolloff, the narrower the frequency band affected. The advantage of this method is that multiplexing and analog-to-digital (A/D) conversion speeds can be maintained at a lower rate. The disadvantage is the possibility of additional noise and drift from low-pass filters inserted ahead of the multiplexer contaminating the data. No filters are needed, however, if the sensor response drops off naturally to half-power at f_0.

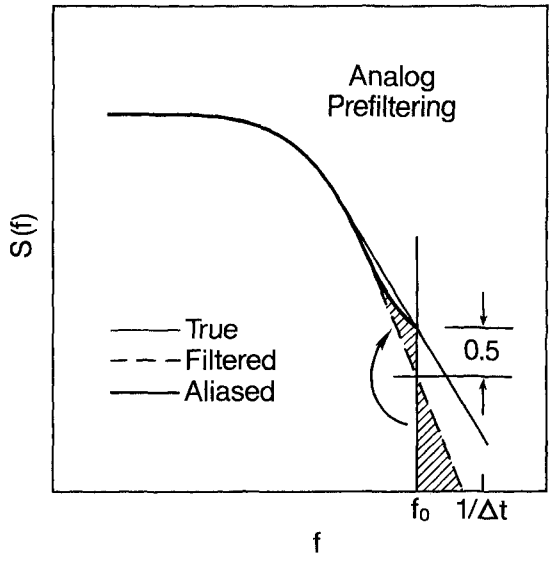

FIG. 7.7 Analog prefiltering aimed at reducing aliasing and restoring full energy at the Nyquist frequency. Filtering reduces energy at Nyquist frequency by a factor of 0.5. (Log scales on both axes.)

2. *Digital prefiltering.* When sampling speed is not a limiting factor (but the digital recording capacity is), it is more convenient to sample the variable at rates 10 to 20 times higher than f_0 and have the low-pass filtering performed digitally at the sensor output by constructing nonoverlapping block averages of the time series. Taking nonoverlapping block averages is equivalent to subjecting the time series to a moving average filter of width Δt but using only points Δt apart. With the first step, we impose the power transfer function $(\sin^2 \pi f \Delta t)/(\pi f \Delta t)^2$ on the original spectrum; with the second we alias back into the spectrum whatever energy remains above f_0 after filtering, as shown in Fig. 7.8.

The block averaging is accomplished by accumulating successive readings in the output buffer of the sensor (if it has a digital output) or of the local A/D converter sampling the signal (in the case of analog signals). Note that in this simple digital prefiltering scheme, aliasing restores only part of the energy lost through filtering (Fig. 7.8) because of the 60% drop in power at f_0. (Aliasing introduced by the original sampling at the higher rate will not normally affect the spectrum, unless large noise spikes exist at frequencies that could fold back into the frequency band between 0 and f_0.) The advantage of this approach is its simplicity. More sophisticated digital filters can be used if recording capacity is not a limiting factor; the low-pass filtering would then be performed at a later time, after the experiment.

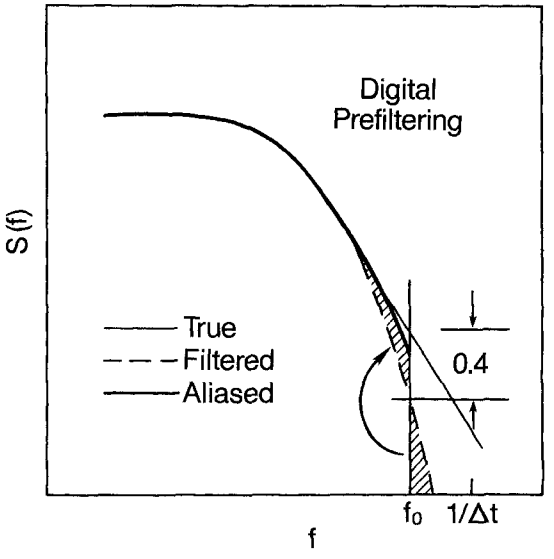

FIG. 7.8 Digital prefiltering achieved by computing nonoverlapping block averages with oversampled data. Energy at the Nyquist frequency is only partially restored because filtering reduces it by a factor of 0.4. (Log scales on both axes.)

7.3.2 *Trend removal*

The presence of a trend in the time series makes the data nonstationary and therefore less suitable for analysis. We define a trend as any frequency component with a period longer than the record length T. In the limit of very long periods (e.g., diurnal variation in temperature), the trend may appear to be linear within that period. Least-squares methods are often used to remove both linear and polynomial trends. For many applications, digital high-pass filtering is preferred because it is simpler and better understood. The same digital filter should be applied to all the variables processed, ensuring uniform low-frequency treatment of the signals.

Trends in the time series produce distortions at the low-frequency end of the spectrum (Fig. 7.9). This distortion, if large enough, could totally mask the true maximum in the $fS(f)$ spectrum and replace it with a spectrum that continues to rise with decreasing f (Fig. 7.10a). Detrending offers no guarantee that the true shape of the spectrum can be retrieved. The particular detrending method used often determines the shape of the detrended spectrum and the location of its maximum (Fig. 7.10b). Trend removal should be performed only if trends are physically expected or clearly apparent in the time series. Automatic detrending is not recommended, except for certain variance and flux calculations where the presence of trends can be highly detrimental (e.g., $\overline{u'\theta'}$ with trends likely in u and θ).

One indication of a linear trend in the time series is an f^{-1} slope at the low end of the $fS(f)$ versus f plot of the spectrum (Fig. 7.10a) best seen in temperature spectra. A linear trend would appear as a replicated ramp function in the Fourier analysis, producing harmonics that decrease in power as f^{-2}, hence the f^{-1} slope in the $fS(f)$ plot. Autocorrelation plots also announce the presence of trends (Fig. 7.9). The correlation function fails to drop to zero even at large lags; trend removal brings it down to zero very quickly.

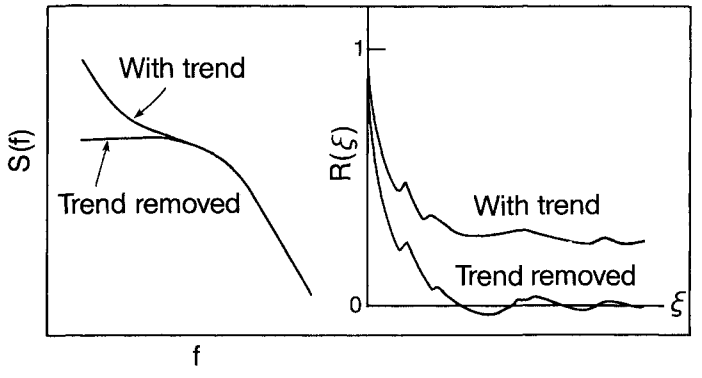

Fig. 7.9 Effects of trend on the spectrum (left) and on the autocorrelation function (right). (Log scales for spectrum axes.)

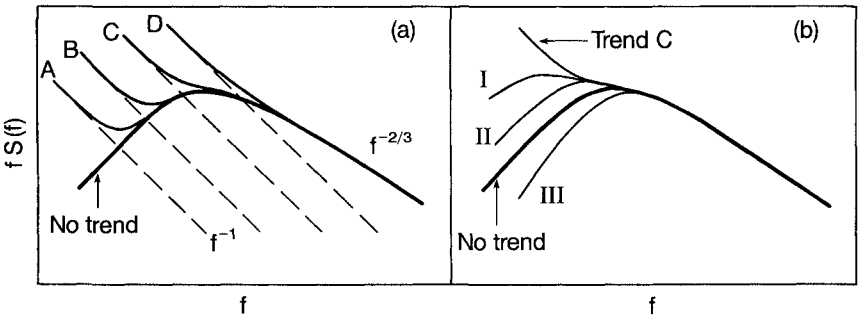

FIG. 7.10 (a) Spectrum distortion from trends of varying magnitudes (A-D). (b) Sensitivity of low-frequency spectrum shape to three hypothetical detrending procedures (I-III). (Log scales on all axes.)

There are arguments both for and against detrending. Scientists who work with data from aircraft find it essential to detrend. Others argue that trends, especially those in w associated with wave motions in stable air (curves A and B, Fig. 7.10a) are physically significant and should be included in the variance and flux calculations because they contribute to vertical transport. The undetrended spectrum should always be available for comparison with the detrended one to ensure that the detrending process removes only the suspected trend and no more. Autocorrelation functions and integral time scales calculated from such data can be particularly misleading because of their strong sensitivity to the detrending procedures used.

The simplest high-pass filter for detrending is one in which the original time series x_i is differenced from an equally weighted running mean of width τ_L. The process can be viewed as a sequence of low-pass, moving-average filtering and subtraction. The filtered time series x_i' becomes

$$x_i' = x_i - y_i,$$

where y_i represents the low-pass filtered time series. The power spectral transfer function $[K(f)]^2$ for this high-pass filter is well known:

$$[K(f)]^2 = [1 - H(f)]^2, \tag{7.6}$$

where

$$H(f) = \frac{\sin \pi f \tau_L}{\pi f \tau_L}. \tag{7.7}$$

$H(f)$ is the low-pass filter function. All frequencies with periods longer than τ_L are attenuated when the power spectrum of x_i is multiplied by $[K(f)]^2$.

Much sharper detrending filters, with less oscillation near $f = 1/\tau_L$, can be constructed using recursive techniques in which the results of successive filtering operations are fed back to the input terms. The operation involves fewer terms than the above "boxcar" approach, and the filter characteristics can be tailored to approximate those of electronic filters (e.g., R-C, Butterworth). McMillen (1988) describes one that simulates an ideal R-C filter with an easily specified time constant. With x_i and y_i representing, as before, the original and low-pass-filtered time series, respectively, a simple recursive filter can be constructed for which

$$y_i = ay_{i-1} + (1 - a)x_i, \tag{7.8}$$

where

$$a = e^{\Delta t/\tau_c}. \tag{7.9}$$

Δt is the time interval between data points and τ_c is the time constant of the desired low-pass filter. When $\Delta t/\tau_c \ll 1$, we have $a \simeq 1 - (\Delta t/\tau_c)$. The new time series is $x' = (x_i - y_i)$, as in the simple moving-average detrending filter.

It is important to point out in this context that departures from the running mean cannot be used directly for computing the fluxes as we would use departures from the conventional time average. Running means do not satisfy Reynolds averaging rules (see Appendix 7.1) so we cannot assume, as we do with simple deviations from the time average, that

$$\overline{\underline{w}'\underline{\theta}'} = \overline{\underline{w}\,\underline{\theta}} - \overline{\underline{w}}\,\overline{\underline{\theta}},$$

where the underbarred terms represent the running means and the deviations from them.

Integral time scales $\mathcal{T}_\alpha$ calculated from high-pass filtered time series also need to be treated with caution. This is because the integral to infinity of the autocorrelation $R(\xi)$ is zero for any high-pass filtered signal. The form of $R(\xi)$, over the range of time lags we normally use to compute the function, is very sensitive to τ_H, the time constant of the high-pass filter, for $\tau_H < 10\mathcal{T}_\alpha$ (see Appendix 7.2). For $\tau_H > 10\mathcal{T}_\alpha$, the time lag at the 1/e point on the correlation curve is usually a good estimate of $\mathcal{T}_\alpha$. At $\tau_H \simeq \mathcal{T}_\alpha$, the time lag corresponding to the first zero-crossing can be taken as an estimate of $\mathcal{T}_\alpha$. At $\tau_H < \mathcal{T}_\alpha$, the correlation curve will be too compromised for any calculation of $\mathcal{T}_\alpha$. All these points are treated in detail in Appendix 7.2.

7.3.3 *Tapering the time series*

When the sampling duration is too short to satisfy the condition $T \gg \mathcal{T}_\alpha$, a tapering window is applied to the time series to bring the values down to zero, or

close to it, at both ends of the sampling period. It minimizes the adverse effect of finite sampling[3] on the magnitude of the computed spectrum and improves our ability to resolve discrete contributions to the spectrum from waves in the signal.

The process of limiting the data to a finite period is equivalent to multiplying the time series with a rectangular time window of unit height and width T. In the frequency domain, this translates to a convolution of the true transform of the process with the function $(\sin \pi fT) / \pi fT$. Convolution, being a smoothing operation, smears out details present in the original transform and extends its range along the frequency axis. The effect on the power spectrum is loss of resolution and overestimation in regions where the power spectrum is dropping off rapidly, as in the inertial subrange. The smaller the T, the wider the lobes of the $(\sin \pi fT) / \pi fT$ function and the greater the smearing and overestimation. For a discussion of the effects of finite sampling on power spectra, see Kaimal et al. (1989).

By multiplying the time series (from which the mean has been removed) with a tapered window we are, in effect, replacing $(\sin \pi fT) / \pi fT$ with a different function, one with a slightly wider main lobe but with greatly suppressed side lobes (the negative and positive oscillations on either side of the main lobe). In the time domain, the tapering reduces the discontinuity at the boundaries of the data when viewed as a replicated sequence of the same time series. Kaimal and Kristensen (1991) tested a number of tapered windows and found the Hamming window brings the measured spectrum to within 1% of the spectral levels we might expect from a long enough record. The window has the form

$$w(n) = 0.54 + 0.46 \cos\left(\frac{2\pi n}{N}\right), n = -\frac{N}{2} \cdots 0 \cdots + \frac{N}{2}, \qquad (7.10)$$

where $w(n)$ is the window function and N is the number of equally spaced data points in the sample. The taper in the window reduces the variance in the time series, so the spectrum has to be compensated for that loss by multiplying it with the ratio of the squares of the areas of the two windows. For the Hamming window, this ratio is 2.52. (Note that, after windowing, the finite time series may have acquired a non-zero mean which must be removed before applying the FFT.)

A striking example of its application is the shipboard measurement of fluxes in the marine boundary layer through the "inertial-dissipation" method (Fairall et al., 1990), where it is necessary to compute the dissipation rates of turbulent kinetic energy, temperature variance, and humidity variance from very short samples of data ($T \approx \mathcal{T}_\alpha$). Use of the Hamming filter on samples of the order of seconds yields dependable inertial-subrange spectral intensities for the flux extimates.

[3]In earlier discussions of the effects of finite sampling (e.g., Pasquill and Smith, 1983) the process is treated as being analogous to high-pass filtering with a spectral transfer function $[1 - H(f)^2]$, where $H(f) = (\sin \pi fT)/\pi fT$; the actual transfer function is more complicated (Kaimal et al., 1989) and approaches the above function only as $T \to \infty$.

Other windows tested by Kaimal and Kristensen (1991) fared less well. The Hanning (cosine squared), Gaussian, and Bartlett (triangle) windows performed almost as well as the Hamming window, but the Poisson (exponential) and the 25% cosine taper windows were clearly unacceptable. [See Harris (1978) for definitions of above windows.] The last two caused overestimations approaching 50%, about the same as reported for the untapered (rectangular) window. Clearly, the commonly used cosine taper is not recommended for very short samples.

The other important benefit to be derived from tapering is the reduction in "side-band leakage" (the leakage of energy to neighboring frequencies through the side lobes in the convolving function), which is especially bad with the $(\sin \pi fT) / \pi fT$ function. Spurious spikes could appear in the spectrum that may be difficult or impossible to distinguish from spikes representing discrete waves. Finnigan et al. (1984) successfully used a Hamming window to identify gravity waves in their studies of wave-turbulence interactions.

7.3.4 *Addition of zeros*

When the time series available falls short of the length required for the spectrum analysis program on hand (power of 2, if an FFT is used), it is common practice to add zeros to the data sequence to make up the required number of points. This procedure, referred to as "padding", should be approached with caution because the consequences of improper application can be severe. The addition of zeros reduces the spectral estimates by a factor $(N - N_z)/N$, where N_z is the number of zeros added and N is the total number of points including the zeros. The spectral estimates should be corrected for this reduction. We can assume that the spectral estimates are diminished uniformly across the bandwidth as long as $N_z < N/3$. It is essential that means and trends be removed from the time series before adding zeros.

7.3.5 *Block averaging*

When the number of data points in the period selected for analysis is too long for a designated spectrum analysis program, it is customary to block average the time series in unweighted, nonoverlapping blocks to reduce the original time series to the number of points desired. Such compression of data enables us to examine the low-frequency behavior in spectra and cospectra quickly. (FFT programs that can handle an entire hour of data, 36,000 data points for instance, in one pass certainly exist, but the computing time needed to produce spectra and cospectra with all the points would be longer and the data handling more cumbersome.) The block averaging introduces some attenuation at the high-frequency end (see Section 7.1), but the loss is easily restored in the spectral output. A good reason for block averaging is to keep aliasing under control.

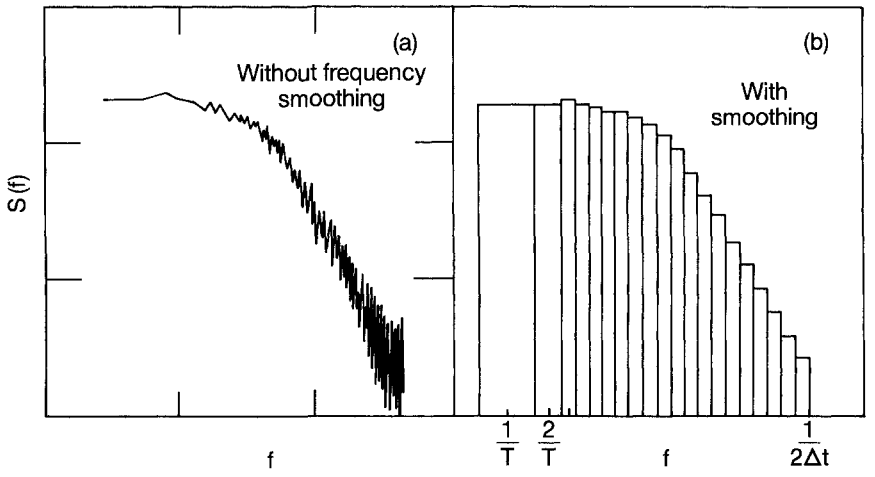

FIG. 7.11 (a) Scatter in the unsmoothed high-frequency spectral estimates. (b) Scatter reduced by smoothing with a frequency window that expands in width with frequency. (Log scales on all axes.)

7.4 Processing spectral data

The raw spectral density estimates derived from spectrum analysis programs, such as the FFT, are too numerous and too ragged to be of direct use to meteorologists.[4] Atmospheric turbulence covers a spectral range of five decades; this is best represented on a logarithmic scale not a linear one. The spectral density estimates generated by the digital programs, however, appear at equally spaced intervals ($\Delta f = 1/T$) on the frequency scale (Fig. 7.3). On a logarithmic scale this spacing results in excessive crowding and large scatter of spectral estimates at the high-frequency end, as illustrated in Fig. 7.11a. Some form of frequency smoothing is needed to extract a representative spectral curve from the estimates. Spectra derived through separate computations covering different segments of the desired spectral bandwidth are easier to splice if they are properly smoothed.

7.4.1 *Frequency smoothing*

An effective smoothing procedure for boundary layer work is one in which the averaging interval keeps expanding with frequency f. If m is the number of estimates in each nonoverlapping block, m is systematically increased as a function

[4]The power spectral estimates obtained directly from the FFT program (which when summed equal the variance of the original time series) have to be divided by $\Delta f(= 1/T = 1/N\Delta t)$ to get them in the form of spectral densities $S(f)$ used in the spectral plots. $S(f)$ has units of variance per Δf, and the area under that curve equals the variance. The frequency-weighted spectrum $fS(f)$, on the other hand, has units of variance, since multiplication by $f(= i\Delta f$, where $i = 1, 2, 3, \ldots \mathrm{N}/2)$ removes the Δf dependence.

of frequency (roughly exponentially) to yield about seven to eight estimates per decade, as in Fig. 7.11b. Each smoothed estimate (of width $m\Delta f$) is assigned to the center frequency of the band. In practice the first few estimates are accepted as they are; then m is increased in steps 3, 5, 7, and so on, until the density of smoothed estimates per decade reaches seven (or eight). The expected power law in the inertial subrange can only be tested dependably after such smoothing.

Other types of frequency smoothing are sometimes applied to spectral data. A three-point Hanning window (weighted 1/4, 1/2, 1/4) is one of many recommended in earlier treatises on the subject (Blackman and Tukey, 1958). Filtering the spectral estimates with such fixed bandwidth windows may provide the smoothing needed in many engineering applications where the spectral range of interest is often narrow, but is not particularly useful for boundary layer work because it still causes crowding of estimates at the high-frequency end.

7.4.2 *Spectral splicing*

When dealing with a very long time series ($N \gg 2^{10}$), it is advisable to split the record into two sets:

1. A set of r short records each with N/r data points
2. A series made up of N/r nonoverlapping block averages (r point averages) of the original time series

The former yields a bandwidth $r/N\Delta t$ to $1/2\Delta t$ and the latter $1/N\Delta t$ to $1/2r\Delta t$ (Fig. 7.12). They overlap over the range $r/N\Delta t$ to $1/2r\Delta t$ which, on a log scale, corresponds to a frequency ratio of $N/2r^2$. [For the Kansas spectra, Kaimal et al. (1972) had $\Delta t = 1/20$ s, $T = 3600$ s, $N = 72{,}000$, and $r = 16$ with an $N/2r^2 = 136$, which corresponds to an overlap of more than two decades. With

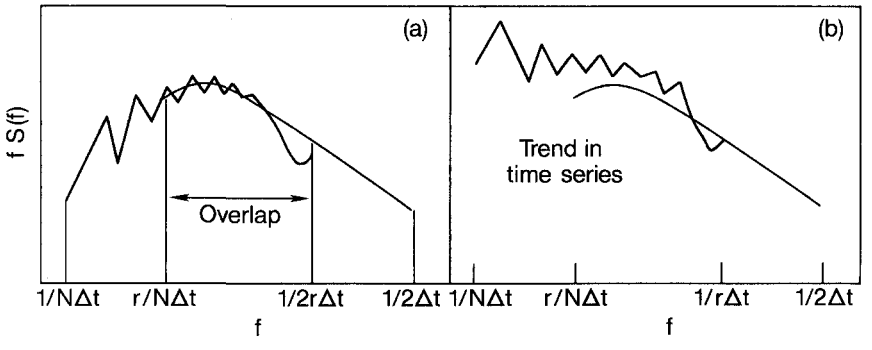

Fig. 7.12 (a) Splicing spectra from the shortened and the block-averaged time series in the absence of significant trends in the data. (b) Effect of trend in the time series that precludes matching in the regions of overlap. (Log scales on all axes.)

this choice of parameters, the same 4096-point FFT routine was used for both the high-end and the low-end spectral computations.]

The r successive spectra from the short records are first averaged to produce a single high-frequency spectrum. This "incoherent averaging" within each frequency band produces a spectrum that is very smooth. By comparison, the low-frequency spectrum will typically have more scatter. Figure 7.12a illustrates how the two usually combine. The match in the region of overlap is often very good so the analyst can decide where to drop the estimates from one spectrum and start with the other. If there is a mismatch, as in Fig. 7.12b, it is invariably the result of a long-term trend in the data. When that occurs, either the data should be detrended or the run discarded as unsatisfactory.

Unacceptably high noise levels at the high-frequency end may also be grounds for rejecting the spectrum. They usually appear as a large rise in the spectrum, approaching an f^{+1} slope, stretching over a decade or more at the high end, where the spectrum should be falling off (Fig. 7.13). This type of noise can be traced to one of the following:

- Sensor noise rising above signal levels as the signal drops below the sensor's noise threshold
- Sporadic spikes in the signal from radio frequency interference, a faulty cable, or mistriggering in the sensor (e.g., sonic anemometer or thermometer)

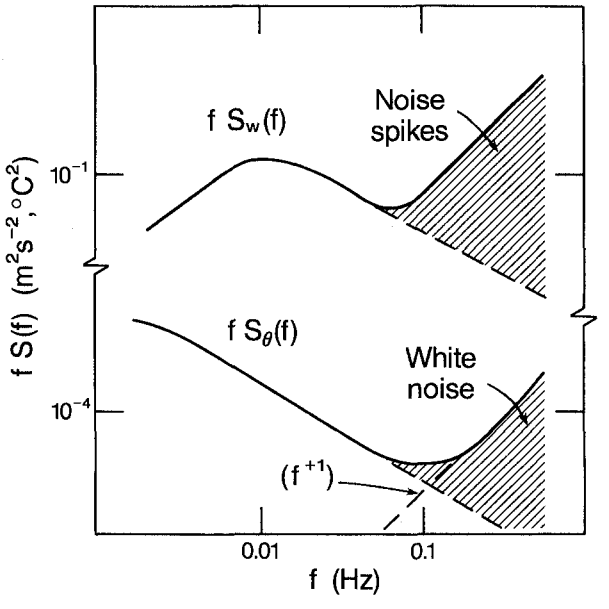

FIG. 7.13 High-frequency distortions arising typically from spikes in a sonic anemometer w signal and white noise in a platinum wire thermometer. (Log scales on both axes.)

Their effects on spectra are illustrated in Fig. 7.13. The noise threshold effect is usually observed in platinum-wire fast-response thermometer spectra when the temperature $fS(f)$ levels drop to 10^{-4} °C^2 or lower. The noise spectrum has the characteristics of white noise: flat in the $S(f)$ versus f plot, f^{+1} slope in the $fS(f)$ versus f plot (Fig. 7.13). This noise is often accepted as inevitable. Spikes, on the other hand, cannot be ignored, unless they are isolated and infrequent, as their effect can extend to all frequencies. Their contribution to the high end of the spectrum can cause it to exceed in magnitude the energy in the true turbulence peak. Corrective action (adjustments and repairs) should be taken to prevent the spikes from contaminating the signal.

In the context of spectral splicing we have to consider the possible implications of deriving the high-frequency portion of the spectrum by averaging spectra from relatively short segments of the time series. The overestimation predicted by Kaimal et al. (1989), discussed earlier, has negligible effect on the spectral estimates in most applications unless we choose the duration T of those segments to be smaller than the period of the spectral peak (τ_m). In the Kansas data analysis T was 3.75 min with τ_m typically 3–4 s for daytime w spectra. In Kaimal et al.'s (1989) study, the effects of finite sampling became apparent only when T dropped below about $10\tau_m$.

7.5 Archiving strategies

In field experiments of limited duration, raw data can be easily stored on digital magnetic tapes and played back as needed for analysis. Different strategies are needed for data collection over extended periods. In an operation similar to the Boulder Atmospheric Observatory where data collection is continuous, it is imperative that a system of archival and retrieval be worked out so the information is standardized and accessible to future users (Kaimal and Gaynor, 1983).

At the BAO, slow-response channels are sampled once a second and fast-response channels 10 times a second. A schematic of the operations at the BAO is shown in Fig. 7.14. In the fast response sonic anemometer channels, aliasing is minimized through digital prefiltering (20-point block average); in the fast-response platinum-wire temperature channels the natural rolloff from thermal conduction losses provides the needed high-frequency attenuation. Three types of data are prepared by the data acquisition computer at the BAO site for transmission to the central computer in Boulder, 25 km away:

- 10-s averaged values of readings on all channels, fast and slow
- 10-s grab samples (last data point in each 10-s block) of all fast-response channels
- Frequency-smoothed spectral estimates (35 estimates) over the frequency range 0.01–5 Hz, updated every 20 min

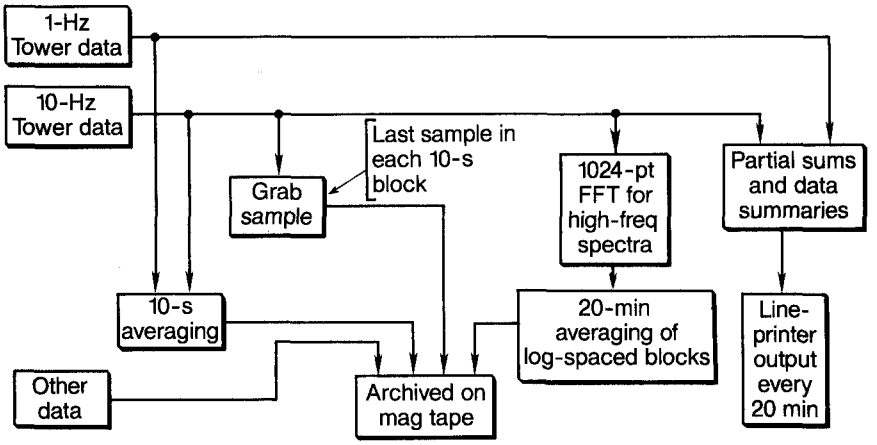

FIG. 7.14 Schematic of data acquisition and processing at the Boulder Atmospheric Observatory (BAO).

The on-site computer also generates partial sums of fast-response data and summaries of slow-response data for presentation of profiles, variances, fluxes, Obukhov lengths, and so on, at the end of each 20-min archiving period.

The 10-s averaged data are saved for future reconstruction of time plots and for filling in the low-frequency end of the spectra (0–0.05 Hz) as discussed in the section above. The 10-s grab samples are the "decimated" data points needed for recomputing fluxes, variances, and third moments over periods longer than 20 min. They are also useful for detecting the presence of noise in the original signal, which may not show up in plots of the 10-s averaged data.

Each 20-min spectrum is an average of ten 2-min (1024 point) FFT spectra. The choice of 20 min as the basic archiving block was arbitrary, selected as a compromise between the need for stability in the statistics and the need to track mesoscale variations in the boundary layer. Frequency smoothing is performed over blocks of increasing width, as discussed in Section 7.4. To fill in the low-frequency end of the spectrum, a 512-point FFT program is used on the 10-s averaged data with the option of adding zeros when the number of data points available falls short of 512. (Even combining four 20-min periods, the time series will be 32 points short.) How well this scheme fits in with spectral shapes in stable and unstable air can be seen in Fig. 7.15.

The partial sums generated in real time for the calculation of data summaries at the end of each 20-min period significantly reduces the processing time at the end. Individual variables and their products (e.g., $\overline{w}, \overline{\theta}, \overline{w^2}, \overline{\theta^2}, \overline{w\theta}$) are accumulated in separate registers as they are sampled. At the end of the 20-min period, variances and fluxes are generated by simple subtractions:

$$\overline{w'^2} = \overline{w^2} - (\overline{w})^2, \tag{7.11}$$

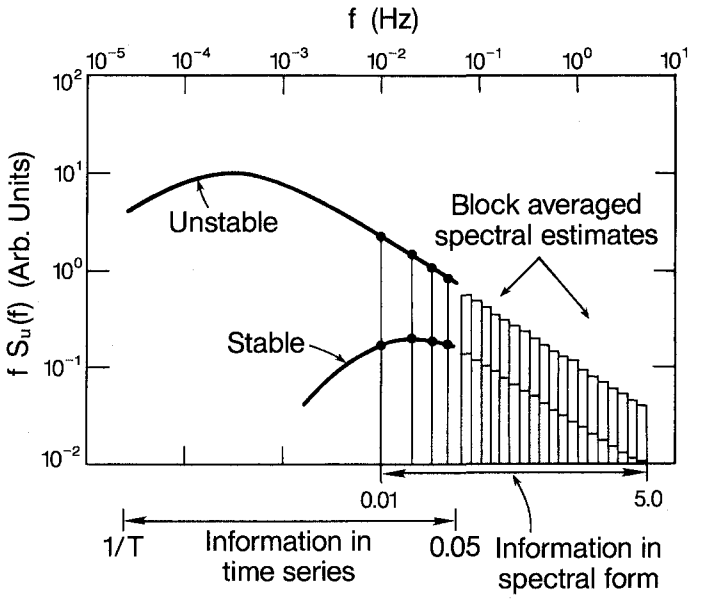

FIG. 7.15 Scheme for compressing turbulence data: high-frequency information stored as smoothed spectral estimates, low-frequency information as time series.

$$\overline{\theta'^2} = \overline{\theta^2} - (\overline{\theta})^2, \tag{7.12}$$

$$\overline{w'\theta'} = \overline{w\theta} - \overline{w}\,\overline{\theta}. \tag{7.13}$$

In deriving (7.11) and (7.12) we assume that $\overline{w'\overline{\theta}} = \overline{\overline{w}\theta'} = 0$, which is valid if $w = \overline{w} + w'$ and $\theta = \overline{\theta} + \theta'$. The same principle can be used for more complicated calculations involving coordinate transformation of the horizontal velocity measurements to the longitudinal (u) and lateral (v) components. If u_m and v_m are horizontal wind components (right-handed) measured along the probe axes and w_m is the measured vertical component, we should be able to compute the variances and fluxes in single operations immediately following ingest of the last data point in that averaging period in order to have them ready for display before the start of the next period. The algebraic operations involved are

$$\begin{aligned}\overline{u'^2} = &\left[\overline{u_m^2} - (\overline{u}_m)^2\right] \cos^2\theta_r + \left[\overline{v_m^2} - (\overline{v}_m)^2\right] \sin^2\theta_r \\ &+ 2\left[\overline{u_m v_m} - \overline{u}_m\overline{v}_m\right] \sin\theta_r \cos\theta_r,\end{aligned} \tag{7.14}$$

$$\overline{v'^2} = \overline{v_m^2}\cos^2\theta_r - 2\overline{u_m v_m}\sin\theta_r\cos\theta_r + \overline{u_m^2}\sin^2\theta_r, \tag{7.15}$$

$$\overline{u'w'} = [\overline{u_m w_m} - \overline{u}_m \overline{w}_m] \cos\theta_r + [\overline{v_m w_m} - \overline{v}_m \overline{w}_m] \sin\theta_r, \tag{7.16}$$

$$\overline{v'w'} = \overline{v_m w_m} \cos\theta_r - \overline{u_m w_m} \sin\theta_r, \tag{7.17}$$

$$\overline{u'v'} = \left[\overline{v_m^2} - \overline{u_m^2}\right] \sin\theta_r \cos\theta_r + \overline{u_m v_m} \cos 2\theta_r, \tag{7.18}$$

where the angle of rotation (counterclockwise) is

$$\theta_r = \tan^{-1}\left(\frac{\overline{v}_m}{\overline{u}_m}\right). \tag{7.19}$$

Similar expressions can be derived for a full three-dimensional coordinate rotation if a tilt correction for sloping terrain or for an incorrectly mounted sensor makes it necessary. In the above transformation we assumed $\overline{w}_m = 0$.

7.6 Special symbols

f_c	cutoff frequency
f_i	low-frequency limit of the inertial subrange
f_0	folding (Nyquist) frequency
n	number of data points
N	number of data points in period T
N_z	number of zeros added for padding
r	factor by which number of data points is reduced
Δt	sampling interval
T	averaging time
u_m, v_m, w_m	measured wind components along instrument axes
$w(n)$	window function, function of n
x_i, y_i	time series of variables x and y
$\alpha(t)$	any time series
θ_r	azimuth rotation angle
$\rho_\alpha(\xi)$	autocorrelation function of α
τ_c	time constant for recursive low-pass filter
τ_H	time constant for high-pass filter
τ_L	cutoff period for boxcar low-pass filter

Appendix 7.1 Reynolds averaging and running mean filters

In standard operations of turbulence analysis, such as the derivation of flow equations or the computation of eddy fluxes, the averaging operator (denoted by an overbar) is assumed

to separate the mean and fluctuating parts of a variable according to certain rules. These are known as the Reynolds averaging conditions and may be summarized as follows:

1. All primed (fluctuating) quantities must average to zero ($\overline{w'} = 0$).
2. The correlations between primed and averaged quantities must vanish ($\overline{w'\overline{\theta}} = 0$).
3. The average of an average must be equal to the same average ($\overline{\overline{w}} = \overline{w}$).
4. Differentiation in space and time commutes with the averaging operations

$$\overline{\frac{\partial w}{\partial x_i}} = \frac{\partial \overline{w}}{\partial x_i} \quad ; \quad \overline{\frac{\partial w}{\partial t}} = \frac{\partial \overline{w}}{\partial t}.$$

For example, when computing fluxes by the eddy correlation method, we usually write

$$\overline{w'\theta'} = \overline{w\theta} - \overline{w}\overline{\theta}, \tag{7.20}$$

neglecting the terms $\overline{w'\overline{\theta}}$ and $\overline{\overline{w}\theta'}$. These terms are, by definition, zero if the overbar represents a true Reynolds average. In fact, disappearance of these terms can be used as a test of its validity for the data being used.

Only ensemble averages can be expected to obey the Reynolds averaging rules precisely, but in practice we usually use time averages, making the ergodic hypothesis that they are equal to ensemble averages. (If the time series is not statistically stationary over relevant time scales, this assumption must be made cautiously.) In many situations, a running-mean low-pass filter is applied to the time series to approximate the "background" variations; the filtered time series is then subtracted from the original signal (i.e., the original is high-pass filtered) to derive the fluctuating components. We should point out that the terms $\underline{w'\overline{\theta}}$ and $\underline{\overline{w}\,\theta'}$ do not necessarily vanish (underbar denotes running-mean quantities) because $\underline{w'}$ and $\underline{\theta'}$ may not have zero means and $\underline{\overline{w}}$ and $\underline{\overline{\theta}}$ have spectral energy spread over finite bandwidths. Thus, the spectra of $\underline{w'}$ and $\underline{\overline{\theta}}$, for example, can be expected to overlap in the region where the high-pass and low-pass filters also overlap (Fig. 7.16). As a result, we can assume $\underline{w'\overline{\theta}} = 0$ only if a spectral gap exists in either w or θ in the region where the two filters overlap. The same comments apply to the computation of variances, since, by the same argument, $\underline{w'\overline{w}}$ and $\underline{\theta'\overline{\theta}}$ cannot be expected to vanish. These considerations are important if moments computed in this way are used in budget calculations that are based on equations derived by Reynolds averaging.

Appendix 7.2 High-pass filtering and integral time scales

In Section 7.4 we discussed the use of digital high-pass filtering to remove low-frequency trends in the data and its consequences for the estimation of integral time scales $\mathcal{T}_\alpha$. These consequences arise because the procedure has the effect of forcing the spectrum at zero frequency $\hat{S}_\alpha(0)$ to zero, causing $\mathcal{T}_\alpha$ to be zero, if the standard definition of the integral time scale is used (Appendix 2.1). Here we will examine how high-pass filtering affects the autocovariance function and through it the integral time scale.

For illustration, we use an example provided by Leif Kristensen (personal communication), which starts with a near-realistic spectrum for variable α

$$\hat{S}_\alpha(\omega) = \frac{\mathcal{T}_\alpha \sigma_\alpha^2}{\pi} \left(\frac{1}{1 + (\omega \mathcal{T}_\alpha)^2} \right). \tag{7.21}$$

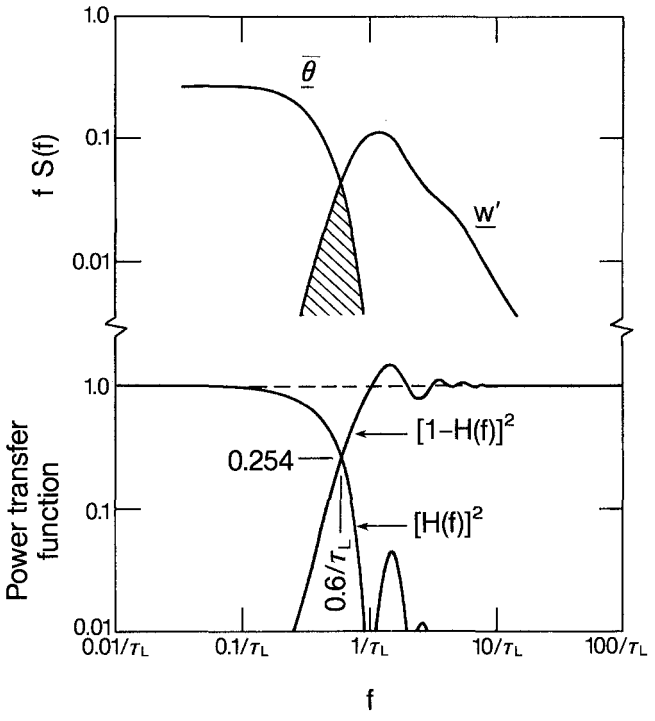

FIG. 7.16 Overlap of $\underline{w'}$ and $\overline{\theta}$ spectra resulting from overlap in the transfer functions for the running mean high-pass and low-pass filters.

(For convenience, we retain the two-sided spectral convention used in Appendix 2.1.) Its autocovariance $R_\alpha(\xi)$ has the simple exponential form

$$\begin{aligned} R_\alpha(\xi) &= \int_{-\infty}^{\infty} \hat{S}_\alpha(\omega) e^{i\omega\xi} d\xi \\ &= \sigma_\alpha^2 e^{-|\xi|/\mathcal{T}_\alpha} \end{aligned} \tag{7.22}$$

and an integral time scale, by our standard definition,

$$\frac{1}{\sigma_\alpha^2} \int_0^\infty R_\alpha(\xi) d\xi = \int_0^\infty e^{-|\xi|/\mathcal{T}_\alpha} d\xi = \mathcal{T}_\alpha. \tag{7.23}$$

At $\omega = 0$, the spectrum reduces to the form derived in Appendix 2.1:

$$\hat{S}_\alpha(0) = \frac{\mathcal{T}_\alpha \sigma_\alpha^2}{\pi}. \tag{7.24}$$

We now apply a first-order high-pass filter with time constant τ_H and a power transfer function $(\omega\tau_H)^2/[1 + (\omega\tau_H)^2]$ that ensures zero contribution at $\omega = 0$. The filtered spectrum will be

$$[\hat{S}_\alpha(\omega)]_{\text{filt.}} = \frac{\mathcal{T}_\alpha \sigma_\alpha^2}{\pi} \left[\frac{1}{1 + (\omega\mathcal{T}_\alpha)^2}\right] \cdot \left[\frac{(\omega\tau_H)^2}{1 + (\omega\tau_H)^2}\right], \tag{7.25}$$

and the corresponding autocovariance function becomes

$$\left[R_\alpha(\xi)\right]_{\text{filt.}} = \frac{\tau_H \sigma_\alpha^2}{\mathcal{T}_\alpha^2 - \tau_H^2} \left(\mathcal{T}_\alpha e^{-|\xi|/\tau_H} - \tau_H e^{-|\xi|/\mathcal{T}_\alpha}\right). \tag{7.26}$$

$$= \frac{\sigma_\alpha^2}{2}(1 - |\xi|/\mathcal{T}_\alpha)\, e^{-|\xi|/\mathcal{T}_\alpha} \quad \text{for} \quad \tau_H = \mathcal{T}_\alpha. \tag{7.27}$$

The effect of filtering on the autocorrelation function $\rho_\alpha(\xi)$ (the normalized autocovariance function) is shown in Fig. 7.17 for different ratios of $\tau_H/\mathcal{T}_\alpha$:0.1, 1.0, 10, and ∞. The effect of high-pass filtering is to introduce a negative lobe (visible, at least, in the first two curves), large enough to make the integral to infinity equal zero. For $\tau_H/\mathcal{T}_\alpha = 1$, the zero crossing occurs at $\xi = \mathcal{T}_\alpha$. For $\tau_H/\mathcal{T}_\alpha < 1$, the zero crossing is shifted closer to the origin and for $\tau_H/\mathcal{T}_\alpha > 1$, farther away from $\xi/\mathcal{T}_\alpha = 1$.

For $\tau_H/\mathcal{T}_\alpha > 10$, the zero crossing occurs outside the plot, at some distant ξ and $\rho_\alpha(\xi)$ is only slightly altered in the range we usually observe. Fortunately, this is the case in most applications, for otherwise the experimenter would be collecting highly compromised data.

The shapes of the correlation functions will, no doubt, vary with the type of filter chosen, but the above exercise does offer some useful insights on high-pass filtering.

1. It tells us roughly how far τ_H has to be separated from $\mathcal{T}_\alpha$ to ensure that important turbulence information is not lost along with the unwanted trend. Note that at $\tau_H/\mathcal{T}_\alpha = 1$,

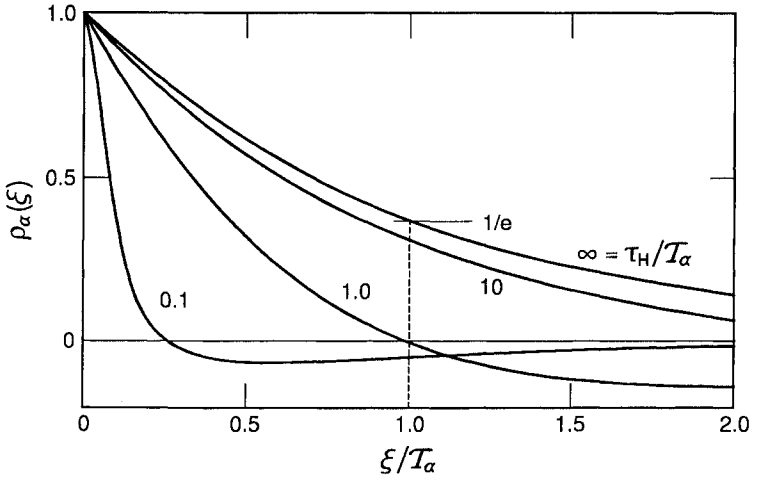

FIG. 7.17 Effect of high-pass filtering on the autocorrelation function for various ratios of filter time constant to integral time scale (from Leif Kristensen, personal communication).

the filter half-power point ($f = 1/2\pi\tau_H$) corresponds approximately to the frequency where the $fS(f)$ spectrum peaks [which roughly equals $1/2\pi\mathcal{T}_\alpha$ as shown in (2.82), Appendix 2.3]. So, a filter with $\tau_H/\mathcal{T}_\alpha > 10$ would preserve much of the energy at the spectral peak and the frequencies above.

2. Designing a high-pass filter for a given application requires prior knowledge of the magnitude of $\mathcal{T}_\alpha$, or where $fS(f)$ might peak. The relationships in Chapter 2 could serve as a guide because the collected data may be too contaminated by the trend to provide that information. The experimenter, by choosing τ_H in this manner, runs the risk of excluding from consideration conditions that deviate significantly from the norm. (The choice of record length T can be limiting in the same way.)

3. Practical alternatives to obtaining $\mathcal{T}_\alpha$ by integrating $\rho_\alpha(\xi)$ to infinity are needed because it is, first, impossible to implement with real data and, second, of questionable value, given the theoretical expectation of $\mathcal{T}_\alpha = 0$ for filtered data. One approach is to choose the value of ξ at which $\rho_\alpha(\xi) = 1/e \simeq 3.7$ ($e \simeq 2.72$ being the base for natural logarithms), which, for an exponential autocorrelation function, is precisely $\mathcal{T}_\alpha$, as pointed out in Chapter 2. Another alternative is to integrate to the first zero crossing or to the point where $\rho_\alpha(\xi)$ is very close to zero. Both these methods will systematically underestimate $\mathcal{T}_\alpha$, seriously so if $\tau_H/\mathcal{T}_\alpha$ is much less than 10.

The concept of the Eulerian integral time scale, or length scale, is intimately related to the question of ergodicity—whether time averages converge to constant values as the averaging time approaches infinity. As ξ increases, an ergodic variable not only becomes uncorrelated with itself, it also becomes statistically independent of itself. The integral time scale is a measure of the time for which $\alpha(t)$ "remembers" itself. At $\xi \gg \mathcal{T}_\alpha$, it approaches statistical independence. A well-designed high-pass filter would preserve the scales of motion that contribute to this decorrelation and remove only those that interfere with it. As a final caution, note that analyses such as those involved in the estimate of errors in moments from finite time averages (Section 7.3) are based on the standard definition of the integral time scale.

References

Bendat, J. S., and A. G. Piersol, 1971: *Random Data: Analysis and Measurement Procedures*. Wiley-Interscience, New York, 407 pp.

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

Fairall, C. W., J. B. Edson, S. E. Larsen, and P. G. Mestayer, 1990: Inertial-dissipation air-sea flux measurements: A prototype system using realtime spectral computations. *J. Atmos. Oceanic Tech., 7*, 425–453.

Finnigan, J. J., F. Einaudi, and D. Fua, 1984: The interaction between an internal gravity wave and turbulence in the stably stratified nocturnal boundary layer. *J. Atmos. Sci., 41*, 2409–2436.

Harris, F.J., 1978: On the use of windows for harmonic analysis with the discrete Fourier transform. *Proc. IEEE, 66*, 51–83.

Haugen, D.A., 1978: Effects of sampling rates and averaging periods on meteorological measurements. In *Proc. Fourth Symposium on Meteorological Observations and Instrumentation*, April 1978, Denver, CO, American Meteorological Society, Boston, MA, 15–18.

Kaimal, J.C., and J.E. Gaynor, 1983: The Boulder Atmospheric Observatory. *J. Clim. Appl. Meteor., 22*, 863–880.

Kaimal, J.C., and L. Kristensen, 1991: Time series tapering for short data samples. *Bound.-Layer Meteor., 57*, 187–194.

Kaimal, J.C., S.F. Clifford, and R.J. Lataitis, 1989: Effect of finite sampling on atmospheric spectra. *Bound.-Layer Meteor., 47*, 337–347.

Kaimal, J.C., J.C. Wyngaard, Y. Izumi, and O.R. Coté, 1972: Spectral characteristics of surface layer turbulence. *Quart. J. Roy. Meteor. Soc., 98*, 563–589.

Lumley, J.L., and H.A. Panofsky, 1964: *The Structure of Atmospheric Turbulence.* Wiley-Interscience, New York, 239 pp.

McMillen, R.T., 1988: An eddy correlation technique with extended applicability to non-simple terrain. *Bound.-Layer Meteor., 43*, 231–245.

Pasquill, F., and F.B. Smith, 1983: *Atmospheric Diffusion* (3rd ed.), Ellis Horwood Ltd., Chichester, England, 437 pp.

Ramirez, R. W., 1985: *The FFT: Fundamentals and Concepts*. Prentice-Hall, Englewood Cliffs, NJ, 178 pp.