USE OF COMPUTER FOR DATA ACQUISITION AND PROCESSING

Jimmy Tan-atichat, Scott Woodward, and William K. George
Turbulence Research Laboratory
Department of Mechanical and Aerospace Engineering
State University of New York at Buffalo
Buffalo, New York 14260

June 16, 1999

1 DIGITAL VERSUS ANALOG: INSTRUMENTATION AND PROCESSING

The design of any data acquisition and processing system eventually involves a decision as to whether the data collected will ultimately be analyzed by analog or digital techniques. Instrumentation for measurement systems can be broadly classified as analog, digital or hybrid systems. Hybrid systems incorporate both analog and digital devices and/or signal processing techniques in an effort to combine the more desirable features inherent in each. Signals that vary in a continuous manner and can assume an infinity of values in any given range are called analog signals. Devices which produce or directly process such signals are called analog devices. In contrast, signals or data which vary in discrete steps and can therefore take on only a finite number of different values in a given range are known as digital signals (or digital data). Thus, devices which produce such signals or are used to perform desired calculations by numerical operations are called digital devices or digital instruments. Some devices are inherently digital (e.g., switches and most types of mechanical or electronic counters), while others such as ordinary mercury-in-glass thermometers, strain gages, D'Arsonval voltmeters are analog in nature. However, in many applications the instrumentation for a particular measurement system can be implemented using devices which are analog, digital or both. In general, digital instrumentation and data processing techniques can offer higher accuracy and analysis speed than is available through analog processing techniques alone. It is also very difficult to design and build analog electronic equipment which is immune to the large number of interference factors caused by ambient temperature and other environmental condition variations (e.g., electrical and electromagnetic noise, component-aging, etc.) that tend to degrade the performance and limit the accuracy of such systems to typically 0.5 - 1.0% (equivalent to approximately 7 bits). While the early digital instruments and measurement systems offered higher accuracy than their analog counterparts, they were also much more expensive and difficult to operate, thus limiting their use to specialists in large research laboratories. Advances in integrated circuit technology in the 1960's and early 1970's paved the way for the extensive use of digital instrumentation, signal processing, and digital control because of increased functionality, ease of use, and frequent price reductions due to competition and high manufacturing yields. Undoubtedly the two most important contributing factors for the widespread use of digital data acquisition and processing systems were the availability of algorithms for the fast computation of discrete Fourier transformations on a digital computer (FFT routines) and low-cost general purpose digital computers, first brought about by the mini-computer in the mid and late 1960's and then by the invention and rapid development of the MOS microprocessor/microcomputer in the early 1970's. With the low cost of microprocessor-based systems and the advent of dedicated digital signal processing (DSP) chips, analog signal processors and hard-wired digital logic instrumentation systems of the late 1960's and early 1970's genre are all but gone. The newer systems are programmed to perform more functions, give results with higher accuracy and sell at lower prices. Although logic functions implemented with software or firmware and executed by a microprocessor are substaintially slower than the hard-wired counterpart, partly due to the different semi-conductor technologies involved, the speed is more than adequate in all but the most time-critical high-speed control applications. Digital systems designed with microprocessors are more versatile than hard-wired logic systems since the logic functions can be easily altered by changing the software being executed. Such changes can be accomplished by firmware (Read-Only-Memory or ROM) replacement for microprocessor-based systems whereas new circuits would have to be designed and built to modify hard-wired digital processing equipment. As more powerful and faster microprocessors are being developed, new applications in high-speed digital data acquisition and processing systems are being found. Despite the above-mentioned advantages of digital instrumentation and data analysis techniques, data presentation to human observers using analog displays are far superior to digital (numerical) displays. For example, it is much easier to visually comprehend the pertinent characteristics or to extract salient features of a time history of some signal (i.e., its waveform) on an oscilloscope or a graphic display than it is to comphrend the equivalent numerical time history of the same signal on a numerical readout device. Therefore, if the data has been processed digitally, a means of converting the results to an analog form for display purposes is very desirable. In control applications, it is often necessary to convert the digital information to an analog form for use as inputs. This is accomplished by digital-to-analog (D/A) converters.

1.1 DATA ACQUISITION

The applications of data acquisition systems can be roughly organized into three categories:

- Measurement - In a measurement application the process model is not known. In fact, the quantities being measured may not be well understood. Thus, measurement is the gathering of sufficient data in an attempt to construct a model of the unknown.

- Testing and Calibration - A testing or calibration approach acquisition represents a situation where a device is being checked against its design standards or measurements of certain parameters are made to establish the values of other calibration parameters. Variables to be measured and requirements for accuracy and precision have been established earlier and are known prior to the tests and/or calibrations.

- Control - In a control application, the system initiates a series of actions, measures them, and takes corrective action if the desired results are not achieved.

1.1.1 General Considerations in Data Acquisition System Selection

In selecting an optimum data acquisition system perhaps the most relevant question that needs to be answered is: 'What sort of information is ultimately going to be extracted from the data acquired?' or equivalently, 'What processing functions are going to be applied to the data?' These questions will be addressed in detail later in the text. Depending on the application, data acquisition systems can range from a simple low-speed data logger to systems which include simultaneous multichannel high-speed digitizing, real-time signal analysis and graphic displays that are controlled by dedicated fast computers. Some important points that need to be considered are:

- Overall Features and Capabilities - The selection process can be aided by providing answers to the following series of questions:

  1. How many channels are required for the proposed work?
  2. Can the number of channels be easily increased? If so, how will such expansions affect other system performance specifications

such as acquisition speed, power supply capability, portability, etc.?

3. What types of transducers will be used and what are their output ranges, output impedances, signal conditioning requirements, etc.?
4. What peripherals or options (if any) will be required to implement a complete, operational data acquisition system with the desired capabilities?
5. What performance boosting options such as enhanced graphics or floating-point processors might be desirable?

- Accuracy - An analysis of the application and the desired accuracy of the final result will reveal whether transducers, signal conditioners, A/D and D/A converters which may only be accurate to within a few percent are acceptable or whether state-of-the-art instrumentation modules must be used. Most requirements typically fall somewhere between the two extremes mentioned above. In others, a compromise between speed and accuracy must be made due to limitations in the present state-of-the-art. The compromise that must be made for high accuracy is the extra cost of the system components, and very often it comes at a reduction in the maximum conversion speed of the available A/D converters. High absolute accuracy (with calibrations traceable to the National Bureau of Standards) may be a prime concern for some applications while in others systems exhibiting only high relative accuracy (i.e., high precision) may be needed. In either case, temperature and long-term stability of the instrumentation system could be an important consideration.

- Acquisition Speed - The key to optimal acquisition of data which can be defined as technically correct for the data processing anticipated and the most economic in terms of data storage and processing time often lies in the proper selection of the acquisition or the sample rate. The optimal sampling rate will be discussed in detail later in this section.

- Data Storage Compatibility - Although many acquisition and processing systems can operate in a stand-alone mode, the complexity of certain data analysis procedures such as computation of discrete Fourier transforms with a large number of data points sometimes requires that the data be transported and analyzed on a more powerful computer system. On these occasions, it is important to ensure that the method (usually disks or network) used for transferring the data from one system to another be compatible, both with respect to the physical medium and to the format of the data created by the user or the operating system software. With the increasing use of computer networking, the problem of data transfer media compatibility has been partially alleviated because collected data can be transferred between computer systems through the network without using a physical transfer medium. One must still ascertain that the internal or binary representation of the data is compatible on both computers.

- System Portability - When data acquisition and control systems are required in the field, electrical generators or battery powered inverters can provide an AC power source. However, in many applications, lightweight, battery-powered computers and data acquisition systems can be used.

- Total System and Operating Costs - Depending on the technical requirements and system specifications, data acquisition and control systems can cost from a few hundred to several hundred thousand dollars or more to purchase. Besides the initial investment, other costs, direct and indirect, involved in the operation and maintenance of the system and especially in programming it should be considered as well. For highly specialized applications, system reliability and the level of technical support available could also play an important role in the selection of a data acquisition system.

1.1.2 General Considerations in Analog to Digital Conversion

An Analog-to-Digital converter provides a series of numbers, each of which corresponds to the weighted integral of the analog signal over some time period, Δ_s, called the 'aperture' which is less than or equal to the sampling period Δ. Several examples are shown in Figure 1. Thus if $u(t)$ is the signal and u_n is the digital sample,

$$u_n = (1/\Delta) \int_{t_n - \Delta}^{t_n} u(t)dt + \mu \quad (1)$$

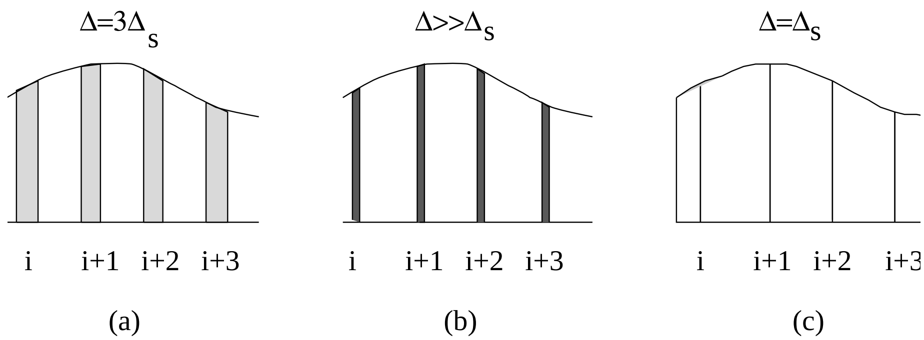

Figure 1: A/D Converter output corresponds to shaded area.

μ is the quantization noise discussed below.

If Δ_s/Δ is much less than unity, as is often the case, a convenient continuous approximation to the sampled signal is

$$u_o(t) = \Delta\delta(t - n\Delta)u(t); n = 1, 2, 3, \ldots \tag{2}$$

where the δ represents the impulse function. It is easy to see that the time average of $u_o(t)$ is exactly average of the numbers generated by the A/D, i.e.

$$U_{oT} = (1/T)\int_0^T u_0(t)dt = (1/N)\Sigma_{n=1}^N u_n \tag{3}$$

since the number of samples N is equal to T/Δ.

Although there are literally dozens of ADC architectures known and in use (see Table 1 for some representative types), two types represent over 90% of all ADC's in the market. They are: 1) the shift-programmed successive approximation ADC for high-speed applications, and 2) the dual-slope integrating ADC for low speed, low cost applications. The only inherent advantage of successive-approximation ADC's is speed. In all other respects, including cost and reliability, the integrating ADC is most likely to be superior. In recent years, the desirable characteristics of the integrating ADC have led to the development of multi-slope versions that extend its applications into the moderate speed range without major compromises.

1.1.3 Quantization Errors

Quantization noise arises from the approximation of a continuous signal by a finite set of numbers. In an ideal ADC, the quantizing error is the difference

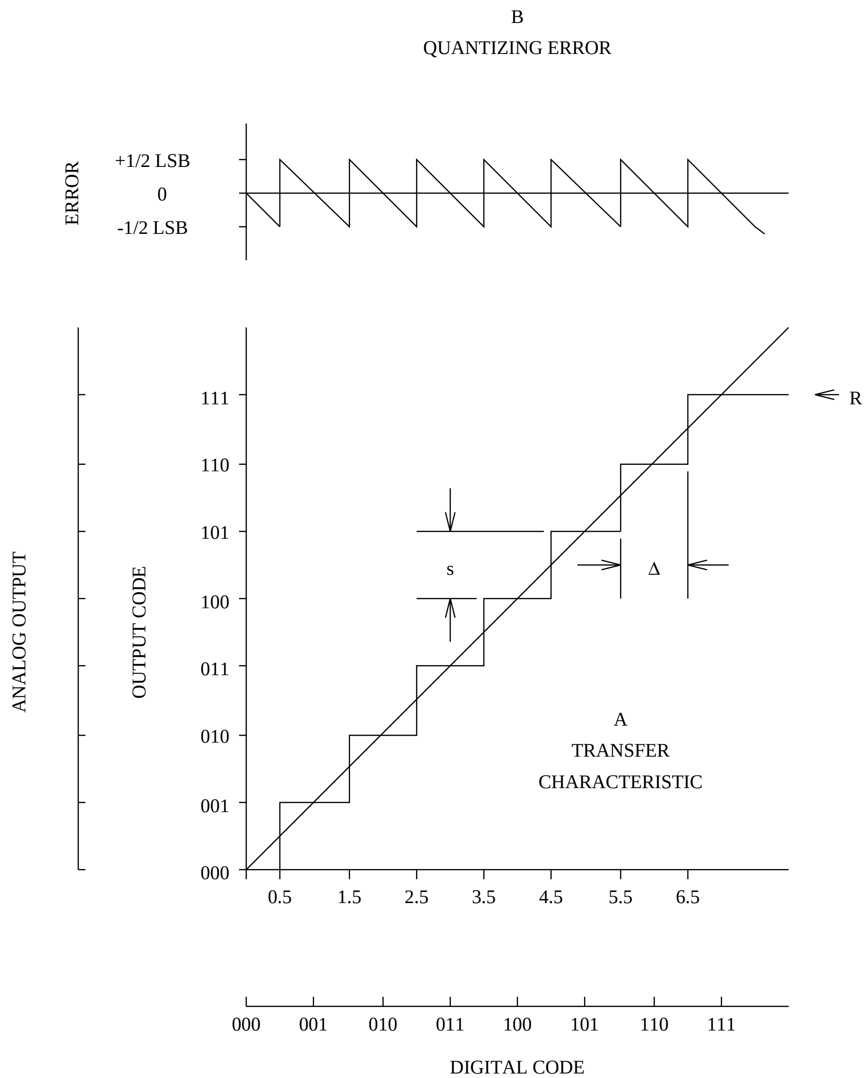

Figure 2: a) Theoretical Transfer Function of a 3-bit ADC and a 3-bit DAC (outer set of axes). b) ADC Quantizing Error

Type of A/D Converter	Application
Shift-Programmed Successive Approximation	General Purpose, High-speed
Dual-slope Integrating	Low Cost, Low Speed
Multi-slope Integrating	Moderate Speed
Parallel-Threshold (Flash)	Highest Possible Speed
Iterative ADC	Extremely High Speed, Less Expensive than Parrallel

Table 1: Representative types of Analog-to-Digital Converters

between the value of the analog input and the digital output due to the finite resolution of the ADC. Figure 2 shows the approximation of a simple ramp signal by a simple three-bit A/D converter. The quantization error (denoted as μ above) acts as a noise and is introduced by the sampling process.

There are a number of different models which can be chosen to compute the effect of the quantization noise on the statistics of the signal. In some cases, as for the ramp shown above and for periodic signals, the quantization noise is periodic. For many random signals a better approximation to the actual noise is obtained by assuming all values of μ between two adjacent levels to be equally probable. For this case the mean square noise is $s^2/12$, where s is the resolution of the A/D, and the noise spectrum is flat with respect to the sampling frequency.

The experimenter must provide enough gain to insure that the mean square value of his signal and its spectral levels are sufficiently above the quantization noise. Sometimes it is necessary to boost the spectral level of the analog signal before digitization by some form of prewhitening (eg., differentiation) to insure that interesting portions of the signal spectrum are not buried in the quantization noise. Alternately, a higher resolution A/D converter can be used to reduce the quantization noise to acceptable levels.

1.1.4 Clipping Errors

There exist a wide variety of ways in which an A/D converter can be configured. Important considerations include the input voltage range and polarity (unipolar/bipolar) required to generate the full scale output code, and whether the input lines are to be differential or single-ended. Regardless of the configuration, clipping errors can arise from the fact that the A/D con-

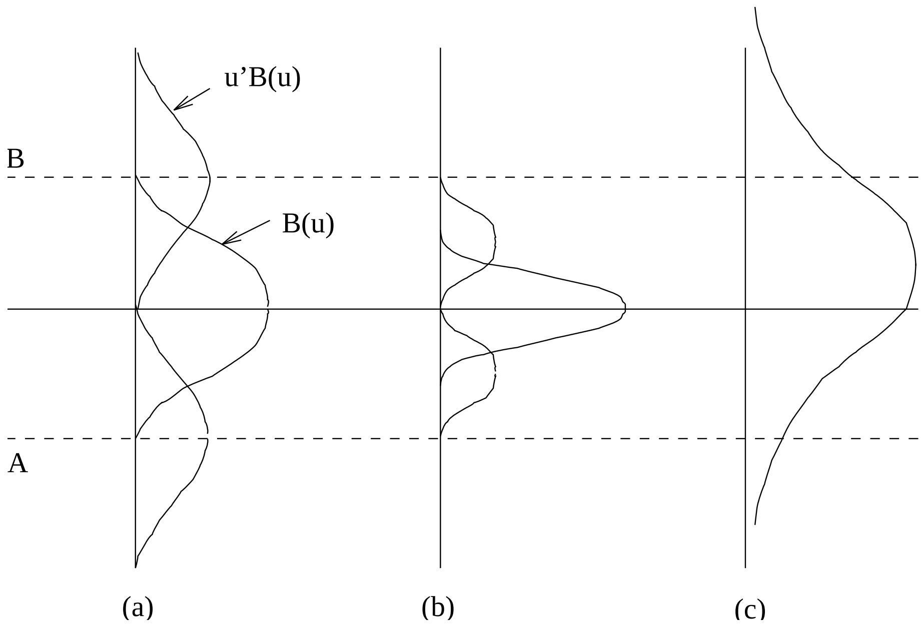

Figure 3: a) Measured mean square will be too low, even though most of the probability density is in the data window. b) Input signal is properly scaled for measurement of mean square. c) Gain and D.C. offset are large enough that clipping will substantially affect all moments, including the mean.

verter has a finite range. In Figure 2 the range is $[0, R]$ and the sampled signal is shown to be clipped at those values.

Clipping can substantially affect the measurement of averaged moments. The higher the moment, the greater the possibility of error. This can be illustrated as follows: If $B(u)$ is the probability density function of the signal $u(t)$, then the n^{th} moment of u, say $\overline{u^n}$, is given by

$$\overline{u^n} = \int_{-\infty}^{\infty} u^n B(u) du \tag{4}$$

But the value obtained from the sampled signal is only

$$\overline{u^n_{meas}} = \int_{A_2}^{A_1} u^n B(u) du \tag{5}$$

where A_1 and A_2 represent the upper and lower clipping levels. The higher the value of n, the more the tails of $B(u)$ are weighted. Since it is the tails of the probability density which are clipped, the more the tails are weighted, the greater the error. Figure 3 (a) illustrates this problem. Since the areas under the curves shown correspond to the first and second moments, it is clear that the latter will be seriously in error while the same clipping will have only a relatively small effect on the mean. Higher moments will be even more seriously affected than the second.

Since the clipping levels of the A/D converter are often fixed, it is usually desirable to precondition the analog signal before digitization to minimize the errors for the particular moment being measured. This can be accomplished by a judicious application of D.C. offset and gain. Figure 3 (b) illustrates a situation where the gain has been reduced to provide an improved estimate of the second moment at the expense of increased quantization noise. Note, however, that the estimate of the lower moments could be contaminated unnecessarily by the quantization noise. Figure 3 (c) illustrates the necessity for offsetting the mean to avoid the bias which can be introduced by clipping.

It is often not possible to optimize the A/D system for measuring all statistical quantities to a sufficient degree of accuracy from a single digitized record. Since it is impossible to correct for errors unless the answers are known in advance (in which case measurement would be unnecessary), the experimenter must often digitize the signal on several channels, each channel having the appropriate preconditioning for the statistical quantity desired.

1.1.5 The Sampling Rate

Table 2 shows typical values of A/D conversion rates and how these rates depend on the accuracy required. Since the requirements for high speed, high accuracy, and low cost are mutually exclusive, the sampling rate/accuracy trade-off can be very important in the selection of an A/D converter. Like the choices outlined above for quantization and clipping problems, the correct or optimal sampling rate is very much a function of the particular experiment and of the type of processing that will be carried out on the sampled data. If only probability analysis (i.e., mean values, moments, or histograms) are desired, the optimal sampling rate is determined by the need to have a sufficient number of statistically independent data points over a sufficiently long time interval to insure statistical convergence. Thus the problem is primarily 'how slowly' (as opposed to 'how rapidly') the data must be taken for a specified period (set by statistical convergence criteria) to insure that adequate storage is available and to minimize computational time. On the other hand, if spectral analysis or time series reconstruction is to be carried

Resolution (Bits)	Typical Conversion Speed (μs)
8	.5
10	2
12	5
14	8
15	10
16	12

Table 2: ADC Resolution versus Typical Conversion Speed

out with the data, both minimum sampling rate and minimum record length criteria must be met. These conflicting criteria are usually most efficiently resolved by digitizing at two or more different rates when both time series and statistical information are desired.

1.1.6 Probability Analysis

The most important factor determining the optimal sampling rate for probability analysis is the requirement to have enough statistically independent samples to insure statistical convergence. Convergence (or accuracy) can be measured by the variability which is defined as the ratio of the standard deviation of the estimate for the statistical quantity to the true value of that quantity. For the mean value, an estimate based on N statistically independent samples has variability

$$\varepsilon^2 = \frac{var\{U_N\}}{\overline{u}^2} = \frac{1}{N}\frac{var\{u\}}{\overline{u}^2} \tag{6}$$

where U_N is the estimate for the mean value, $\overline{u}$ is the true mean value, and $var\overline{u}$ is the true variance of the fluctuating signal. Thus the percent error of the measurement depends on the variance of the phenomenon and the inverse square root of the number of independent samples.

The N samples in equation 6 must be statistically independent. In general, samples taken in an arbitrary manner from a continuous process will not be statistically independent, but will be correlated. The time over which a signal is correlated with itself can be measured by an appropriately defined integral scale. It can be shown that samples separated by two or more integral scales affect convergence as though they were statistically independent.

Therefore, if Υ is the integral scale and Δ is the time between samples, the effective number of independent samples is

$$N_i = \begin{cases} N\Delta/(2\Upsilon) & , \Delta < 2 \\ N & , \Delta \geq 2 \end{cases} \tag{7}$$

Thus if the sampling rate is faster than one sample every two integral scales ($\Delta < 2\Upsilon$), the additional samples contribute nothing to convergence and only increase the memory and data storage requirements. On the other hand, if the sampling rate is slower than one sample every two integral scales, all samples are effectively independent, but the total length of time required to achieve a given accuracy is longer. The optimal sampling rate is clearly equal to one sample every two integral scales. This rate will always be lower than the rate required for time series analysis (see below), although the total length of record required will be similar. Thus, if only probability analysis is required, not only can there be a considerable savings in data storage requirements, but the A/D hardware complexity is reduced since the A/D conversion rate requirements are much less.

1.1.7 Time Series Analysis

An entirely different set of sampling constraints applies if spectral analysis or time reconstruction of the signals is to be carried out. The basic problem is that of spectral 'folding' or 'aliasing', which arises from the fact that the Fourier transform of a digitally sampled signal is a periodic version of the Fourier transform of the original signal. In brief, the transform of the original signal is repeated at intervals corresponding to the rate at which the data are sampled. Thus if $u(f)$ is the Fourier transform of $u(t)$, the Fourier transform of the digitized signal is

$$u_o(f) = \Sigma_{-\infty}^{\infty} u(f - n/\Delta) \tag{8}$$

This is illustrated in Figure 4.

If all signals were bandlimited (as in Figure 4 (a) & (b)) to frequencies less than half the sampling rate, problems would never arise. That serious problems can arise is illustrated in Figure 4 (d) where the highest frequency of the signal is greater than half the sampling rate. The spectra are folded into each other so that it is impossible to determine the original spectrum. If folding occurs, the data are said to be aliased. It is impossible to recover time series information once it is lost due to aliasing.

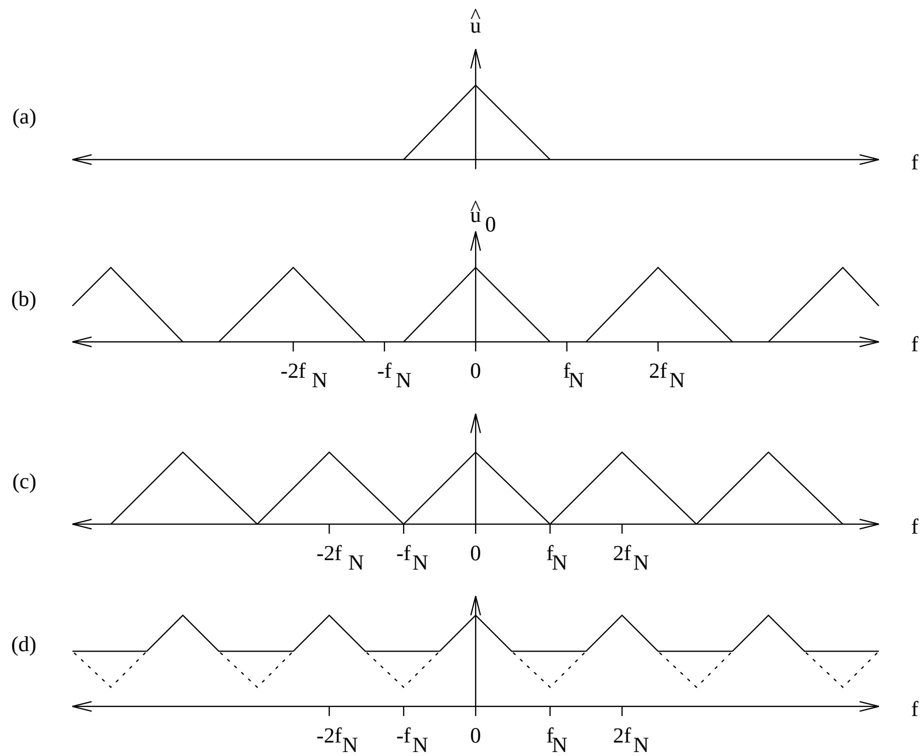

Figure 4: a) Transform of input analog signal b) Sample rate greater than twice the highest frequency. c) Sample rate equal to twice the highest frequency. d) Sample rate less than twice the highest frequency. Original transform no longer recoverable.

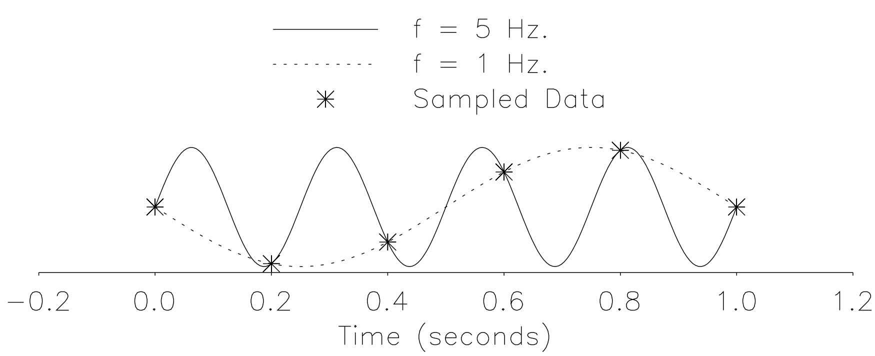

Figure 5: Aliasing of a 4 Hz. sine wave to a 1 Hz. sine wave by a 5 Hz. sampling rate.

The folding frequency, which is half the sampling rate, is referred to as the Nyquist frequency, thus $f_N = 1/2\Delta$. Aliasing can be avoided only if the Nyquist frequency is greater than any frequency present in the signal, the so-called Nyquist criterion. Figure 5 illustrates what happens when a sine wave is sampled at something less than twice its frequency. Instead of appearing at the proper frequency, say f_o, the peak is aliased to the frequency $2f_N - f_o$ which corresponds to the negative side of the first repeated version of the Fourier Transform of the original signal.

The phenomenon of aliasing does not occur if the rate at which the signal is sampled is at least twice the highest frequency. In fact, all of the information present in the original signal is also present in the digitized signal (but only if the record length is infinite!). This fact is the essence of the Shannon Sampling Theorem which states that the original time series can be reconstructed if the Nyquist criterion is satisfied. A corollary of this result is the fact that once a signal has been sampled in a manner which introduces aliasing, the original time series can not be recovered and the spectral information is aliased forever, regardless of how much digital processing is carried out.

It is not always convenient to sample at rates higher than twice the highest frequency because of the need to have sufficiently long records to insure that the lowest frequencies are captured. Also, in some situations the highest frequencies are difficult to determine in advance, or may be uninteresting. In such situations, it is necessary to precondition the analog signal by analog low-pass filtering prior to sampling. Since no analog filter completely removes the high frequencies, the signal must be digitally sampled at a rate higher than twice the cutoff frequency of the filter, the exact value being determined by the roll-off of the analog filter and the signal itself.

As a final note it should be pointed out that aliasing does not affect the probability measures of a signal, only its spectral character. Therefore the Nyquist criterion should not be applied if only probability measures are desired since the only effect will be to greatly increase the amount of data which must be stored with no improvement in statistical accuracy.

1.1.8 Simultaneous Sample and Hold

Almost all A/D converters employ a sample and hold circuit to latch the voltage which is to be converted. The sample-and-hold amplifier is a circuit that can be digitally commanded to operate in two sequential modes:

1) The sample mode, and 2) the hold mode. Ideally, the output is an exact reproduction of the input in the sample mode. In the hold mode, the output is 'frozen' and maintained without decay or drift, regardless of the length of time or the amount of external influences. To reduce hardware cost and complexity the various analog inputs are often multiplexed into a single sample-and-hold amplifier at a rate determined by the user and/or the device clock. (A multiplexer is a circuit designed to connect one of a number of input signal paths ('channel') to an output load, switching from channel to channel in an arbitrary, consecutive, or changing sequence, in accordance with a digitally coded 'channel address' instruction from the system program control source. Most MUX modules contain internal decoding circuitry to convert the binary address input to one-in-eight, one-in-sixteen, or whatever else is required to switch on one channel at a time.) The rate at which each channel of data is digitized is the rate at which the A/D converter is operating divided by the number of channels multiplexed into it. Thus, if the number of channels is high, the required conversion rate can be quite high to achieve even modest Nyquist frequencies on the individual channels.

A consequence of the multiplexing of analog data in the manner described above is that all channels of data are not digitized simultaneously, but rather in sequence. These time delays (or phase lags) between channels can have a serious effect on both probability and time series analyses. For example, if the sample interval is Δ, the time between the conversion of channel m and channel n is $(m-n)\Delta$. Thus the cross-covariance between the two channels is not $\overline{u^2}$, but rather $\overline{u(t)u(t+(m-n)\Delta)}$ which can be substantially different depending on the nature of the cross-correlation $\overline{u(t)u(t+\tau)}$ and the lag time $(m-n)\Delta$. As is shown below, the correct cross-moments can be found by first carrying out a spectral analysis of the data.

This phase lag also has an effect on spectral analysis. The cross spectrum can be shown to be given by

$$S^{mn}_{meas}(f) = S^{mn}_{true}(f)exp(-i(m-n)2\pi f\Delta) \tag{9}$$

From equation 9 it is clear that the correct cross-spectrum can be obtained from the measured cross-spectrum by simply multiplying the latter by the complex conjugate of the frequency dependent exponential. Since the cross-correlation is the inverse Fourier transform of the cross-spectrum, the true cross correlation can be obtained by making the phase correction before transforming. If this is done then the true cross moments can be evaluated, either from the cross-correlation at zero-lag or from the integral of the cross-spectrum.

An alternative to these correction procedures is to avoid the problem altogether. One method is to introduce multiple sample-and-hold amplifiers before the multiplexer so that all channels are latched simultaneously before being read sequentially by the A/D. Another is to simply provide each channel with its own A/D converter and synchronize their operation. When used in conjunction with dedicated data buffers, extremely high acquisition rates can be achieved in this manner since no multiplexing is required.

1.1.9 Digital-to-Analog Conversion

Sometimes it is desirable to reconstruct an analog signal from digital samples using a D/A converter. Digital-to-analog converters are devices that accept digital signals (coded sets of bits) at their input terminals, and generate an analog current or voltage output that is exactly representative of the input code, if each bit were perfectly weighted. Figure 2 shows the theoretical transfer function for the three least-significant bits of a binary input code. Note that these three bits yield all the output values for the eight possible binary input codes from 000 to 111. The transfer function for an 8-bit D/A converter would have 256 steps and a 12-bit converter would have 4,096.

An analog signal is best generated by using an approximation to the Whittaker interpolation formula on the digital data before D/A conversion. This formula is given by

$$w(t) = \sin(\pi t/\Delta)/(\pi t/\Delta) \tag{10}$$

Other digital low-pass filters can also be used if their cutoff frequency is sufficiently below the sampling frequency so that a reasonably smooth output is produced. Often the digital samples are fed directly to a D/A whose output is analog low-pass filtered to smooth it. If this is done, the number of samples per unit time fed to the D/A should be 5-10 times higher than the cutoff of the filter to insure a reasonable reproduction of the original analog signal.

1.2 PROCESSING OF DIGITAL DATA

1.2.1 Data Validation

There are a variety of errors which can be introduced into digitized data by random hardware and transmission errors. These may involve only a single bit or may represent a more serious loss of data. It is important before investing a substantial effort in data processing that the data be subjected to some preliminary screening. If possible the data should be plotted out as a time series so the user can inspect it visually since few computer algorithms are as effective at spotting data problems as the human eye-mind combination of a skilled investigator.

Probability analysis can also provide useful clues as to the validity of data. In particular, the probability density function can be especially useful in spotting wild points and clipping. Statistical criteria can often be utilized to remove or replace bad data. The exact treatment depends on the nature of the errors, the signal itself, and the type of analysis to be carried out.

1.2.2 Computation of Moments

Assuming that a sufficient number of data are available to insure statistical convergence, the moments can be computed in a straightforward manner. For example, the mean and variance are computed from

$$u_N = \frac{1}{N}\Sigma_{i=1}^{N} u_i \tag{11}$$

and

$$var\{\overline{u}_N\} = \frac{1}{N-1}\Sigma_{i=1}^{N}(u_i - u_N)^2 \tag{12}$$

The factor of $(N - 1)$ in the last expression is preferable to N since it can be shown to have less variability. The primary problem which can occur is that the summation becomes so large that succeeding data make no contribution since their scaled magnitude is less than the lowest significant bit carried by the machine in the summation. This problem is called underflow and can be very important when computing statistical quantities from large blocks of data. As a preventive measure, it is a good practice to use at least double precision variables for summation variables.

1.2.3 Computation of Probability Density Functions

The estimation of marginal or joint probability density functions from digital data requires only an algorithm to sort the data into appropriately defined boxes. The boxes must be sufficiently small that the shape of the measured density is not distorted by the curvature of $B(u)$ across the box. The variability of such estimators is inversely proportional to the expected number of independent occurrences in each box. Samples may be assumed independent if separated by at least two integral scales in time.

For example, if $B(u)$ is the marginal probability density and Δu is the box width at u, the expected number of occurrences in this box is $NB(u)\Delta u$ where N is the total number of samples in the ensemble. Similarly for the joint probability density function $B(u,v)$, the expected number of occurrences in a particular box is $NB(u,v)\Delta u \Delta v$. (If the sampling rate is faster than one sample every two integral scales, this number must be reduced by a factor $2\Upsilon/\Delta t$ so that it correspondss to the number of independent samples in the box.) Thus extremely long records are required to accurately estimate values of the probability density function in the far tails which can be very small. Obviously the situation can be relieved somewhat by increasing the size of the boxes in the tails, but at the risk of increased bias due to curvature effects.

1.2.4 The Autocorrelation and Cross-correlation

Direct computation of the auto- and cross-correlations from digitized signals can easily be implemented on a digital computer by simply cross-multiplying the data strings with the appropriate time lag and then averaging. Thus

$$C_N^{uv}(\tau) = \frac{1}{N}\Sigma_{n=1}^{N-m} u_n v_{n+m} \tag{13}$$

where $\tau = m\Delta t$. Note that the sum is divided by N instead of $N - m$ since this provides the least variability in the estimator. As always, care must be taken to avoid underflow.

The statistical error in the estimate for the correlation is

$$\varepsilon^2 = \frac{var\{C_N(\tau)\}}{C^2(\tau)} \tag{14}$$

$$= \frac{2\Upsilon}{T}\frac{C(0)}{C(\tau)}^2 \tag{15}$$

where $T = N\Delta t$ is the record length and Υ is the integral scale for the process. Note that since $C(\tau)/C(0) \rightarrow 0$ as $\tau \rightarrow 0$ because of stationarity, the variability becomes unbounded. Thus, the smaller the correlation, the larger the record length must be relative to the integral scale in order to achieve a specified accuracy. Note that the variability is not dependent on the A/D conversion rate $1/\Delta t$, therefore Δt can be chosen for a convenient estimation of the correlation at the desired lags, as long as the spectrum is not to be computed from the result. (If the spectrum is to be computed, the Nyquist criteria must be satisfied.) Alternately, computations can be minimized by computing only those time lags desired (as opposed to all those possible).

A considerable savings in computational time can be achieved if only pairs of points separated by two integral scales are utilized. This is because pairs closer together are not statistically independent and therefore do not contribute to reduce the variability.

An alternate means by which the correlation can be estimated is to first compute the spectrum by means of a Fast Fourier Transform and then inverse transform it to get the correlation. This method is sometimes referred to as the 'direct method' and will be discussed below under Spectral Analysis.

1.2.5 Fourier Analysis

The digital Fourier transform is defined as

$$\tilde{u}_k = \Sigma_{n=1}^{N} u_n e^{-i2\pi nk/N} \tag{16}$$

This can readily be shown to be the counterpart of the conventional continuous transform

$$\tilde{u}(f) = \int_{-\infty}^{\infty} u(t) e^{-i2\pi ft} dt \tag{17}$$

if the frequency is taken equal to the product of the index k and the bandwidth $B = 1/T$ where T is the record length. Note that only $N/2$ frequencies are generated since both real and imaginary parts must be determined from the N data points.

In many situations there is a considerable advantage to implementing equation 16 by the so-called Fast Fourier Transform (FFT). The FFT can be implemented in a variety of ways, all of which consist of shuffling the data into bit-reversed order and introducing the appropriate phase factors. The savings in computational time can be substantial since only $Nln(N)$ operations are required for an N-point transform as opposed to N^2 operations for the conventional implementation. Which particular FFT algorithm is chosen is very much a function of the particular machine on which it is implemented and the speed and accuracy desired.

Since the Fourier transform of even a real signal is in general complex, it would appear that complex computations must be carried out. This is not the case since a number of algorithms exist which treat the real and imaginary parts separately using only real variables. In fact, it is possible to utilize the complex nature of the transform to simultaneously transform two signals which have been interwoven; i.e., u_1, v_1, u_2, v_2, etc.

1.2.6 Spectral Analysis

There are two methods by which an estimator for the spectrum can be generated. The first of these requires computing the correlation function first, then Fourier transforming it to obtain the spectrum. This method follows directly from the fact that the correlation and spectrum are a Fourier transform pair; that is,

$$S(f) = \int_{-\infty}^{\infty} C(\tau) e^{-i2\pi f\tau} d\tau \tag{18}$$

and

$$C(\tau) = \int_{-\infty}^{\infty} S(f) e^{+i2\pi f\tau} df \tag{19}$$

It is usually more efficient on a digital computer to utilize a FFT algorithm directly on the time series data, and then compute the spectrum from the resulting Fourier coefficients. A 'direct' estimator for the spectrum is given by

$$S_T(f) = \frac{|\tilde{u}(f)|^2}{T} \tag{20}$$

where T is the record length. The frequencies are given in multiples of the bandwidth $B = 1/T$ (i.e. $f = B, 2B, \cdots, NB/2$), the index corresponding to the index in the transform vector. Note that even though the spectrum is real, it is determined by both the real and imaginary parts of the Fourier transform.

If the signal which is being analyzed is a realization of a random signal, its Fourier transform and the resulting spectrum will also be random. It can be shown that the variability of a spectral estimator is given by

$$\varepsilon^2 = \frac{var\{S_T(f)\}}{S^2(f)} = \frac{1}{BT} \tag{21}$$

where $S(f)$ is the true spectrum, B is the bandwidth and T is the length of the record from which the spectral estimate is made. Since the bandwidth is $B = 1/T$, the rms fluctuation of the spectral estimate is equal to the

Lag Window Name	Lag Window Formula
Rectangular	$D_R(\tau) = \begin{cases} 1 & , \lvert\tau\rvert \leq M \\ 0 & , \lvert\tau\rvert > M \end{cases}$
Bartlett	$D_B(\tau) = \begin{cases} 1 - \frac{\lvert\tau\rvert}{M} & , \lvert\tau\rvert \leq M \\ 0 & , \lvert\tau\rvert > M \end{cases}$
Tukey	$D_T(\tau) = \begin{cases} 1/2\left\{1 + \cos\frac{\pi\tau}{M}\right\} & , \lvert\tau\rvert \leq M \\ 0 & , \lvert\tau\rvert > M \end{cases}$
Parzen	$D_P(\tau) = \begin{cases} 1 - 6\left\{\frac{\tau}{M}\right\}^2 + 6\left\{\frac{\lvert\tau\rvert}{M}\right\}^3 & , \lvert\tau\rvert \leq \frac{M}{2} \\ 2\left\{1 - \frac{\lvert\tau\rvert}{M}\right\}^3 & , \frac{M}{2} \leq \lvert\tau\rvert \leq M \\ 0 & , \lvert\tau\rvert > M \end{cases}$

Table 3: Equations for Common Lag Windows

spectrum itself! This is not acceptable since every realization of the signal can be expected to yield spectra which are completely different. Therefore some means of averaging or smoothing the spectral data must be utilized.

The simplest means of reducing the fluctuations in the spectrum is to simply average the spectral estimates obtained from independent data records, the variability being reduced by the inverse of the number of estimates utilized. If these blocks of data are obtained successively by subdividing a record into m-subrecords, the reduction in the variability of the estimator can be shown to result from the increased bandwidth since $B = 1/T_m = m/T$. This procedure is usually referred to as block averaging.

An alternate method of reducing the spectral variability to acceptable levels is by averaging adjacent frequency estimates, thereby increasing the effective bandwidth of the estimate and reducing the number of independent frequencies for which estimates are made. This operation corresponds to convolving the original spectral estimate with a filter window, the window being determined by the weighting used when summing adjacent coefficients. In continuous variables,

$$S_{fil}(f) = \int_{-\infty}^{+\infty} S(f_1)W(f - f_1)df_1 \tag{22}$$

where $W(f)$ is the frequency space representation of the window.

Several commonly used filter windows and their Fourier transforms are shown in Figures 6 and 7 respectively. It is easy to show that the block

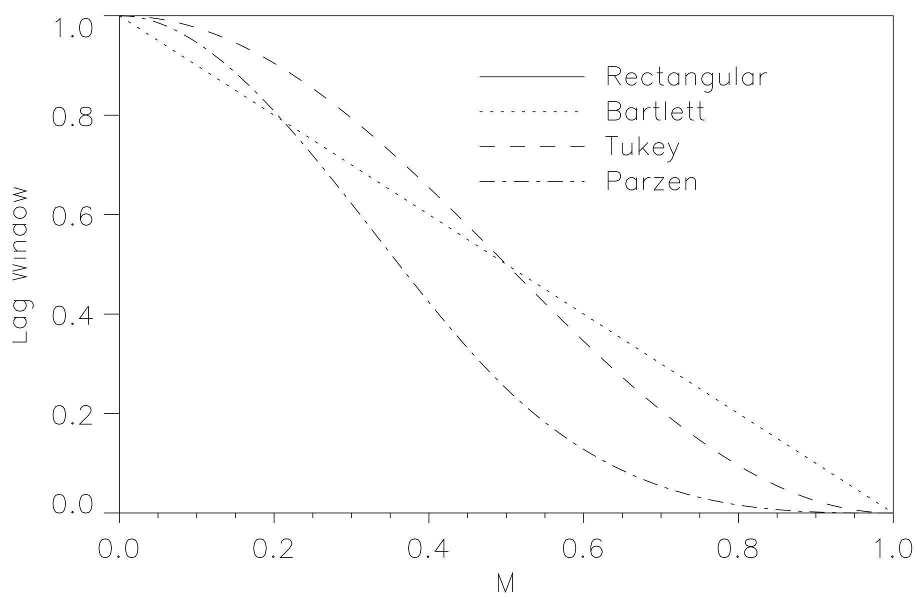

Figure 6: Lag windows commonly in use. (see Table 3 for equations)

Lag Window Name	Spectral Window Formula	Effective Bandwidth ($\Delta\omega$)
Rectangular	$W_R(\omega) = \frac{M}{\pi}\left\{\frac{\sin \omega M}{\omega M}\right\}$	$\frac{\pi}{M}$
Bartlett	$W_B(\omega) = \frac{M}{2\pi}\left\{\frac{\sin \omega M/2}{\omega M/2}\right\}^2$	$\frac{3\pi}{M}$
Tukey	$W_T(\omega) = \frac{M}{2\pi}\left\{\frac{\sin \omega M}{\omega M}\right\}\left[\frac{1}{1-(\omega M/\pi)^2}\right]$	$\frac{8\pi}{3M}$
Parzen	$W_P(\omega) = \frac{3M}{8\pi}\left\{\frac{\sin \omega M/4}{\omega M/4}\right\}^4$	$\frac{3.77\pi}{M}$

Table 4: Equations for Common Spectral Windows and Effective Bandwidth

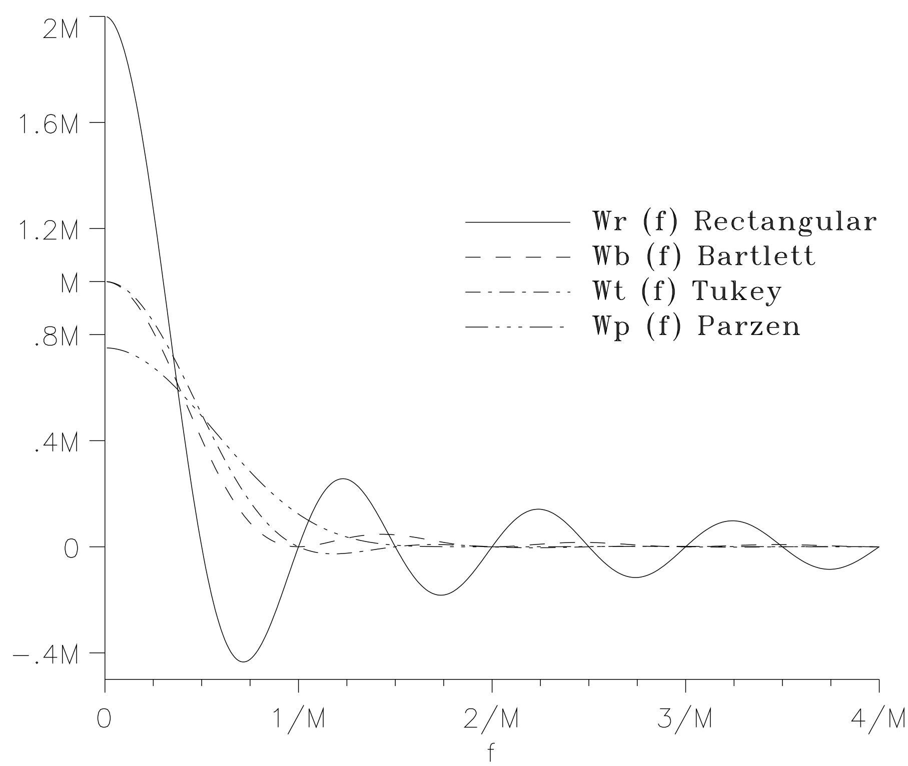

Figure 7: Spectal windows commonly in use. (see Table 4 for equations)

averaging method employed above is equivalent to convolving the spectrum with a Bartlett (or triangular) window. Note that several of the windows have negative side lobes which can yield physically unrealistic negative spectral estimates. These side lobes can lead to a problem called 'leakage' which occurs when the estimate at a given frequency is determined primarily by the spectrum at different frequencies which contain more energy and have been sampled by the side lobes. Leakage can be avoided by using windows which are sufficiently narrow or have side lobes which roll off more rapidly than the spectrum itself.

Convolution in frequency space corresponds to multiplication in time space and vice versa. This can be used to facilitate the smoothing of spectra. First the spectrum is inverse-transformed (by an FFT for example) to yield the correlation function. This correlation function is multiplied point-by-point by the appropriate time window, and then the result is transformed to yield the smoothed spectra.

It should be noted that even though low frequency data may not be desired, the record length must always be much greater than the integral scale for the process, unless the time the process is correlated with itself is reduced by high-pass filtering. This is because the expected value of the spectral estimate (for the direct case) is really

$$S_T(f) = \int_{-T/2}^{T/2} C(\tau) e^{-i2\pi f\tau} \{1 - \frac{|\tau|}{T}\} d\tau \tag{23}$$

whereas what we desire is given by equation 18. The finite limits in equation 22 are responsible for the $1/T$ bandwidth discussed earlier. It is clear that the estimate converges to the desired result only if the factor $(1 - |\tau|/T)$ is approximately unity over the range for which $C(\tau)$ is non-zero. If not, spectral leakage will be introduced by the 'default' window, the record length. Similar considerations can be shown to apply for the indirect estimator as well.

1.2.7 Cross Spectra

There are fundamentally no differences between computing spectra and cross-spectra by either method outlined above. If the direct method is used, the transform of each channel is first computed, then the cross-spectrum is obtained as

$$S_T^{uv}(f) = \frac{\tilde{u}(f)\tilde{v}^*(f)}{T} \tag{24}$$

where $*$ denotes the complex conjugate.

As before, these cross-spectra must be either block-averaged or smoothed to reduce the variability. The procedure differs only in that the negative frequencies must be accounted for separately since the cross-spectra are in general neither symmetric nor real.

The constraints on minimum record length can be more severe than for the spectrum and a significantly longer record may be required for cross-spectra than for spectra. This is easily seen from the estimator which can be shown to reduce (for the direct case) to

$$S_T^{uv}(f) = \int_{T/2}^{T/2} C^{uv}(\tau) e^{-i2\pi f\tau} \left[1 - \frac{|\tau|}{T}\right] d\tau \tag{25}$$

Unlike the autocorrelation, the cross-correlation does not in general peak at $\tau = 0$ and in fact may have its peak far from the origin. The record must still be long enough to insure that the factor $(1 - |\tau|/T)$ makes a negligible contribution over any region in which $C^{uv}(\tau)$ differs from zero.

To illustrate the kind of problems which can arise, consider two probes responding to the same signal but with one delayed with respect to the other. The record must be sufficiently long to insure that there is enough overlapping information (i.e., the same detail sensed by both probes although at different times) to permit averaging long enough to reduce the variability to acceptable levels. This problem often arises with hard-wired analyzers and in small computers where the record length is limited by the available memory. In such situations an external delay must be introduced into one input to move the peak of the cross-correlation to the origin, thereby reducing the record length requirement to approximately that for the spectrum alone. The phase introduced into the cross-spectrum by the external delay can be removed by complex multiplication since the spectrum computed in this manner is readily shown to be given by

$$S_{lag}^{uv}(f) = e^{-i2\pi f\tau_{lag}} S^{uv}(f) \tag{26}$$

where τ_{lag} is the externally introduced delay.

References

J.S. Bendat and A.G. Piersol. *Random Data: Analysis and Measurement Processes.* Wiley, New York, 1971.

Analogic Corp. 1982 design engineers handbook and selection guide: A/d and d/a converter modules signal conditioning and control modules. Technical Report BR1021, Rev. 1, Wakefield, MA, 1981.

Digital Equipment Corporation. PDP-11, LSI-11, and VAX hardware/software handbooks. Technical report, Billerica, MA, 1983.

L.D. Enochson and R.K. Otnes. *Programming and Analysis for Digital Time Series Data.* Number SVM-3 in The Shock and Vibration Monograph Series. Naval Publication and Printing Office, Washington, 1968.

W. K. George. Processing of random signals, in dynamic flow measurement in unsteady flows. In *Proceedings of the Dynamic Flow Conference 1978*, pages 757–793, P.O. Box 121, 02740, Skovlunde, Denmark, 1979.

B.M. Gordon. Linear electronic analog/digital conversion architectures, their origins, parameters, limitations, and applications. *IEEE Trans. on Circuits and Systems*, 25(7), July 1978.

J.L. Hilburn and P.M. Julich. *Microcomputers/Microprocessors: Hardware, Software, and Applications.* Prentice-Hall, Englewood Cliffs, NJ, 1976.

G.M. Jenkins and D.G. Watt. *Spectral Analysis and its Applications.* Holden-Day, San Francisco, 1968.

S. Leibson. Leibson on i/o. Technical report, Hewlett-Packard Desktop Computer Div., Fort Collins, CO, 1981.

J.L. Lumley and H.A. Panofsky. *The Problems of Measurement and the Interpretations of Data in Geophysics, in The Structure of Atmospheric Turbulence.* Interscience, New York, 1964.

D.R. McGlynn. *Microprocessors: Technology, Architecture, and Applications.* John Wiley, New York, 1976.

R.K. Otnes and L.D. Enochson. *Applied Time Series Analysis, Basic Techniques*, volume 1. Wiley, New York, 1978.

G.V. Rao. *Microprocessors and Microcomputer Systems.* Van Nostrand Reinhold, New York, 2 edition, 1982.

C. W. Van Atta. Sampling techniques in turbulence measurements. *Annual Review of Fluid Mechanics*, 6:75–91, 1974.

A.H. VanDoren. *Data Acquisition Systems.* Reston Publishing Co., Reston, VA, 1982.