



THE METEOROLOGICAL MAGAZINE

HER MAJESTY'S
STATIONERY
OFFICE

May 1981

Met.O. 942 No. 1306 Vol. 110

THE METEOROLOGICAL MAGAZINE

No. 1306, May 1981, Vol. 110

551.501.45: 519.24/.25

Smoothing and filtering of meteorological data

By A. C. L. Lee

(Meteorological Office, Bracknell)

Summary

The concept of filtering is discussed as a means of estimating a representative signal from data that have been corrupted by noise. Comparisons are made between the meteorologically common 'arithmetic average' filter and other filters. The Bessel filter emerges as a filter that gives a more representative estimate, can be readily implemented, and is particularly useful where considerable smoothing is required.

1. Introduction

In meteorology, in common with many other physical sciences, there is a need to make representative measurements of a relatively slowly moving 'signal' (such as wind direction), that is corrupted by 'noise' (such as turbulence) which has a wide spectral content.

A common technique is to smooth or filter the data over some period of time. This paper is not intended to consider the length of time over which any specific meteorological phenomenon should be smoothed, but rather to discuss the concept of filtering as a method of representing the underlying phenomenon of interest, often referred to as the 'mean' value. Several alternative filters, having the property of easy implementation, are considered, and the conflict between noise removal and phase distortion is highlighted. The Bessel, or Thompson, filter is shown to be a useful compromise. Its performance in estimating 'mean' quantities is shown to be superior to that of a simple arithmetic average in non-stationary situations, which are always those of interest. This is particularly true where derivative quantities, such as wind veer, are significant. Finally, a technical discussion is presented indicating how Bessel, or similar, filters can be implemented in hardware or in economical software. The latter is particularly significant in an era of cheap, although low-performance, microprocessor technology.

2. The purpose of filtering

It is helpful to pursue the specific example of wind direction further. We may wish to measure the synoptic wind direction, but a single sensor will measure the synoptic variable with turbulent noise superimposed. There will also be effects due to topography and the friction layer, but this is a separate issue which for our purpose will be ignored.

For the above example, we might ideally have a dense network of sensors and use spatial averaging over a 'synoptic' scale length to smooth out noise. Clearly this is impracticable; instead some form of 'time averaging' or filtering is often used, the assumption being made that the 'noise' has a zero mean, and that if we can average for long enough the residual noise will be small.

Unfortunately we cannot 'average' for an arbitrarily long time; the synoptic situation also changes and we are interested in following certain rapidly changing situations such as occur during the passage of a front.

There are two subtly different viewpoints when considering data from a single sensor:

(a) We are interested in temporal structures of time-scales longer than a certain 'characteristic time'. Longer structures are essentially synoptic, while shorter structures are essentially unwanted turbulent noise. However, we do wish to discriminate between separate events in time as far as possible in order to distinguish the sequence of events.

(b) We are interested in spectral frequencies below a certain 'characteristic frequency'. Lower frequencies are essentially synoptic, while higher frequencies are essentially unwanted turbulent noise. Thus we wish to be able to discriminate frequencies as far as possible so as to retain the wanted lower frequencies, and reject the unwanted high frequencies.

The two viewpoints are not mutually exclusive; they can both be satisfied using a 'low-pass' filter which attempts to exclude frequencies above the 'characteristic frequency', whose value is related to the scale length of the desired spatial averaging.

In general 'noise' will have a continuous spectrum, and for any characteristic frequency there will always be a contribution to the final signal from the low-frequency noise below the characteristic frequency in addition to any high-frequency noise which may not be adequately suppressed. At the same time we wish to be able to follow the more rapid synoptic variation without undue suppression or other (phase) distortion.

It must be realized that filtering is not the same as taking an arithmetic 'average' over some strictly defined period of time. If the latter is specifically required, it defines a specific type of filter. If, however, the problem is left more general, in the sense that the purpose is to monitor synoptic variations while suppressing 'noise' which has a zero average, this leaves more flexibility in the design of the filter. This extra flexibility can result in a filtered signal which more nearly approximates to the 'synoptic' signal than does the arithmetic average. In particular, different types of filter can place different emphases on distortion and noise removal, corresponding to the two viewpoints above.

In this paper, the purpose of low-pass filtering is taken to be the estimation of a representative non-stationary signal which has been corrupted by noise having a zero mean.

3. Approaches to filter implementation

Filtering is a subject well familiar to engineers, and there appears to be spasmodic interest in the meteorological literature. This interest has usually centred around low-pass filtering of data series consisting of samples equally spaced in time. Most authors have discussed various approaches to weighting the time series with a suitable finite set of weights, often in the form of a bell-shaped distribution.

An early paper in the geophysical literature is by Holloway (1958). He discusses the problem of assigning weights, without reaching any conclusion on 'optimum' methods. Craddock (1968) suggests an empirical scheme for determining a suitable distribution. Passi (1976) shows that 'smoother' data can

be obtained by making the distribution correspond to a mathematical function. Ruthroff and Bodtmann (1976) emphasize the 'smoothness' by showing how noise-reduced derivatives can be extracted from filtered data. Charnock (1957) uses a truncated approximation to a $(\sin x)/x$ impulse response in order to obtain a highly discriminating frequency response. Kaiser and Reed (1977) exhibit a more modern approach by finding alternative ways of approximating an ideally sharp cut-off frequency, while avoiding undesirable spreading into higher frequencies through 'ringing', by recognizing that only a finite number of weights will be used. A feature of their approach is to manipulate mathematical functions having known properties such as the maximally flat Butterworth polynomial, or the equi-ripple Chebyshev polynomial. By proper manipulation of the known functions, the desired type of approximation can be made to the ideally sharp cut-off low-pass filter.

The above papers share a common trait—they do their best to remove high frequencies, but pay little regard to the shape of the weighting function itself. To obtain the sharpest cut-off, the weighting function will have large ringing excursions, and so the filter response to an isolated 'impulse' will exhibit the same ringing. As sharper cut-off is achieved at some characteristic frequency, the necessary ringing weighting function extends in time. In this sense higher discrimination in the frequency domain leads to poorer temporal discrimination. If both time and frequency discrimination are considered important, the resulting set of weights closely resembles the smooth bell-shaped Gaussian or normal distribution, and this sort of distribution has been used by many authors. The necessary truncation of this smooth curve tends to throw up minor ringing in the frequency domain. In a recent review paper, Harris (1978) describes a number of useful alternative 'optimum' approaches to obtaining finite sets of weights which are discriminating in both domains.

From the viewpoint of smoothing time-series data, the finite weighting function approach has an important drawback. If the original continuous data contained high frequencies, then frequent samples must be taken if the data are to be fully represented. Specifically, a minimum of two samples must be taken per cycle at the highest frequency present in the data at any significant amplitude i.e. the Nyquist constraint (Shannon 1949). If the data are to be 'smoothed' over a relatively long period, then a large number of weights are involved. In a software implementation this implies a large number of multiplications per time-step, which is computationally expensive. In a hardware environment it is difficult to implement such a finite set of weights conveniently and accurately, if at all.

An alternative approach is to use a sampled data version of the modern analogue filter theory. The basic ideas of modern filter theory are discussed in textbooks such as Herrero and Willoner (1966), and Zverev (1967). Here the data are filtered by implementing a transfer function in the frequency domain which can be mathematically represented by numerator and denominator polynomials in the complex frequency s . In the analogue world, such transfer functions are easily implemented because most useful linear signal transmitting media approximate to this sort of transfer function.

In a sampled data system, it is possible to approximate such transfer functions closely by evaluating recursive relations between the data going into and out of a filter. Meyer and Isacker (1975) discuss recursive filters, although they emphasize high-frequency discrimination without regard to time discrimination. An exact equivalence cannot be obtained between the effective weighting function produced by these transfer functions and the finite set of weights described above, because the effective weighting function is strictly infinite in duration, although only a finite portion contributes significantly to the filter output. Thus the two approaches produce different types of filter. It will be shown in later sections that these recursive relationships can often be implemented with little more than one multiplication per time-step on average, especially if the data are smoothed over a relatively long period, and that this multiplication can often be degenerated into an accumulator shift. This has obvious advantages for microprocessor applications.

The nearly exact equivalence of the modern filter theory approach in both continuous analogue and sampled data systems means that the same filter can be implemented either way, or else by a hybrid, which can have advantages.

The simplest low-pass filters are obtained when the numerator becomes unity, leaving just a denominator polynomial. The resulting class of filters are the so-called 'polynomial filters'. The order of the filter corresponds to the highest power of s in the denominator polynomial. A good deal of attention has been paid over the years to obtaining the optimum performance (in some carefully defined manner) from low-order polynomial filters, which correspond to simple physical or recursive realizations. These will be emphasized in the next section. It will be seen that, for our purposes, little is lost by degenerating the numerator.

4. Polynomial filters and the arithmetic-average filter

Figs 1–4 describe the properties of several low-pass types of polynomial filter. Each has been normalized in frequency so that the 'characteristic frequency' (defined here as having 3 dB attenuation) occurs at $\omega = 1$ radian per second. The most widely known polynomial filter is the Butterworth, whose characteristics are shown in Fig. 1. Fig. 1(a) shows the filter attenuation in decibels against frequency on a logarithmic scale for several orders ($n = 1$ –10) of filter. Note the change in scale at $\omega = 1$. The polynomial coefficients in the Butterworth transfer function have been arranged so that all the derivatives of the absolute value of the transfer function with respect to frequency, except the n th, are zero at $\omega = 0$. In this sense the transfer function is maximally flat at low frequencies, while at high frequencies it is dominated by the highest power of s present.

The Butterworth is an excellent filter for separating sinusoidal signals at low and high frequencies. Unfortunately, the time by which the signal is delayed on passing through the filter (the group delay, shown in Fig. 1(b)) varies with ω , and this applies to frequencies below the characteristic frequency. The distortion which this produces on non-harmonic signals can be clearly seen in Figs 1(c) and (d) which show the filter response to a Dirac impulse (of vanishingly small width and unit area), and to a unit step. The filter 'rings' for a considerable time. The impulse response is of particular interest as, with a reversed time-scale, it represents the continuous weighting function which must be applied to the previous history of the raw data sequence to produce the filtered output. Evidently the Butterworth's ringing intermingles discrete events in time.

Fig. 2 shows the properties of a particular Chebyshev filter. In this filter a considerable increase in the rate of change of attenuation near the characteristic frequency can be achieved at the expense of allowing a defined amount of ripple (here 0.5 dB) in the passband (below the characteristic frequency). As might be expected, the group delay and impulse responses are even worse, from our point of view, than the Butterworth.

The Butterworth and Chebyshev filters are optimized in the frequency domain, and represent attempts to obtain a box-like discrimination with zero attenuation below the characteristic frequency, and infinite attenuation above. It can be shown using the Hilbert transform (Zverev 1967) that attempts to produce this 'ideal' frequency discrimination will always exhibit large variations in group delay and considerable spreading of the impulse response through ringing, independently of whether polynomial or other transfer functions are used. Thus box-like frequency discrimination produces poor time discrimination.

It is common practice in meteorology to form an arithmetic average of the measured signal for some finite time, e.g. a 10-minute average wind direction or wind speed. Some measuring devices will integrate a signal, present the integrated values, and restart a new integration. This form of widely spaced

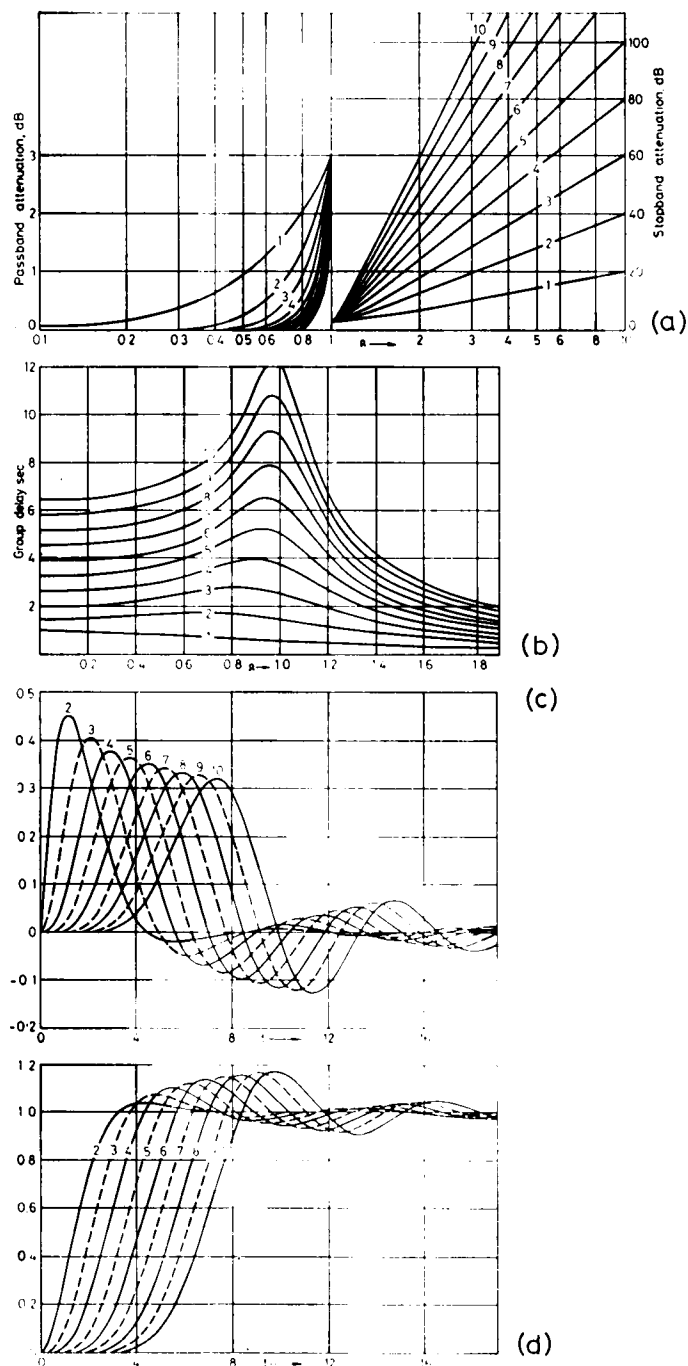


Figure 1. Butterworth filter characteristics. (a) Attenuation (decibels) v. frequency (radians per second). (b) Group delay (seconds) v. frequency. (c) Impulse response (for unit impulse) v. time (seconds). (d) Step response (for unit step) v. time.

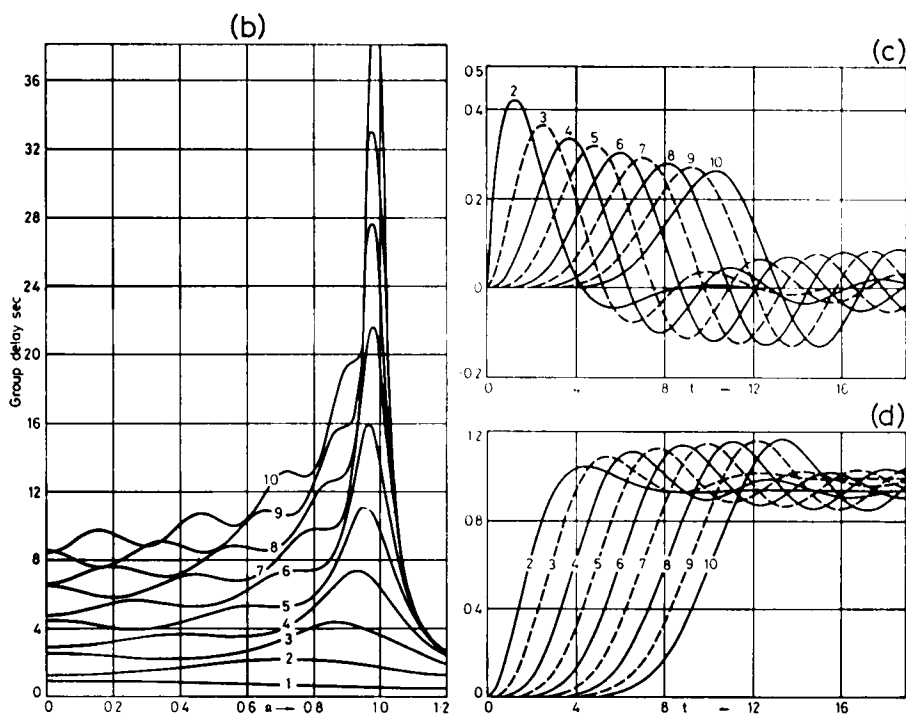



Figure 2. 0.5 dB Chebyshev filter characteristics. (a) Attenuation v. frequency. (b) Group delay v. frequency. (c) Impulse response. (d) Step response.

sampled output is difficult to compare with a continuous filter as the results from the two systems will in general be different, and there is no continuous trace to show how representative the samples are of the previous 10 minutes. Instead a filter will be considered which continuously presents the result of an arithmetic average over the previous T seconds, and whose response $R(\omega)$ to a harmonic signal $\sin(\omega T + \delta)$ is given by:

$$R(\omega) = (1/T) \int_t^{t+T} \sin(\omega t + \delta).dt = \{2/(\omega T)\} \sin(\omega T/2) \sin(\omega t + \delta + \omega T/2).$$

The factor $\sin(\omega t + \delta + \omega T/2)$ represents the original harmonic input with a phase delay of $\omega T/2$, or a group delay of $T/2$ seconds at all frequencies. The remaining factor, $\{2/(\omega T)\} \sin(\omega T/2)$, represents the amplitude response as a function of frequency. At low frequencies this tends to unity. For compatibility, choose $T = 2.778697$ seconds, so that we have 3 dB attenuation at $\omega = 1$. Fig. 3 shows the properties of this filter. The impulse response, Fig. 3(c), shows constant weighting over T seconds. It will be no surprise to find that we have achieved 'ideal' box-like discrimination in the time domain.

The transfer function of a filter and its impulse response form a Fourier transform pair. Moreover, the transform from frequency to time domain is identical to that from time to frequency domain, apart from a scaling factor. As we have already observed that a box-like discrimination in the frequency domain gives rise to poor discrimination through ringing in the time domain, we should not be surprised to find that a box-like discrimination in the time domain gives rise to poor discrimination through ringing in the frequency domain, as evidenced in Fig. 3(a). The attenuation peaks of the arithmetic-average filter, which go to infinity, occur at frequencies whose periods are exact submultiples of the integration period. Apart from these peaks, the general level of attenuation is rather low, increasing at a rate of about 12 dB per octave, or 20 dB per decade. This means that high-frequency noise, which usually contains most of the noise energy, is poorly suppressed compared to the performance of previous filters. Lewis (1960) describes the 'Slutzky-Yule' effect where a moving-average filter can appear to give apparent periodicities in the smoothed data. Peacock (1970) points out that the phase reversals in the arithmetic-average transfer function which occur after each infinite attenuation peak can have the effect of reversing the sign of quasi-periodic events which can still be seen in the smoothed time series.

If infinite attenuation peaks, which are necessarily associated with instantaneous phase reversals, are avoided, flat group delay is a reliable indicator of a distortion-free filter. The arithmetic-average filter has this desirable property in its passband, as can be seen by the constant group delay of $T/2$ shown in Fig. 3(b).

We have now reached a point where it appears that excessive zeal in producing the 'ideal' filter in one domain produces problems in the other, and we should consider filters that give nearly equal importance to both domains. A mathematical function of great interest in this context is the Gaussian function $g(t)$:

$$g(t) = \exp \{-t^2/(2\sigma^2)\}.$$

It will be seen that this is the bell-shaped normal distribution curve, familiar to statisticians. Its Fourier transform $G(\omega)$ is given by:

$$G(\omega) = \sigma \exp \{-(\sigma^2/2)\omega^2\}.$$

It can be seen that the time-domain Gaussian transforms to a Gaussian shape in the frequency domain. Thus we now have a frequency domain function which, if used as a low-pass filter, gives equal discrimination in both time and frequency domains. The characteristic frequency can be varied by changing σ .

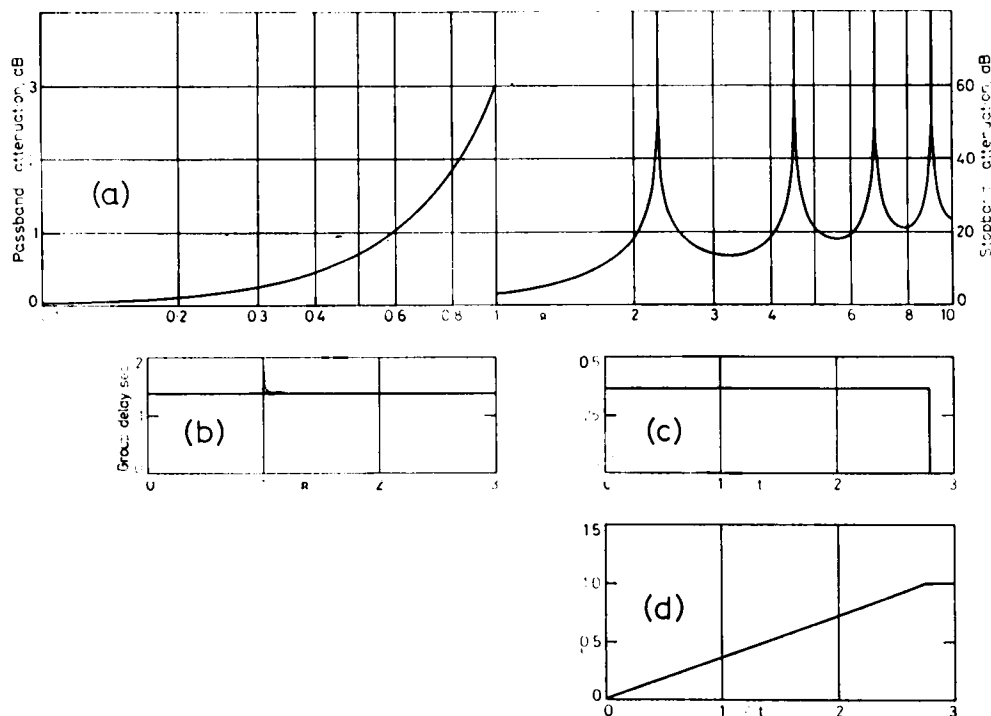


Figure 3. Arithmetic-average filter characteristics. (a) Attenuation v. frequency. (b) Group delay v. frequency. (c) Impulse response. (d) Step response.

Because $G(\omega)$ is purely real (having no imaginary part), the phase shift is zero at all frequencies, so the group delay has the highly desirable property of being constant—at zero. However, if we consider $g(t)$ as an equivalent time-reversed weighting function, we see that it is non-zero for negative values of t ; we require future values of data for the filter as well as past values. These will obviously not be available, so that a physical realization is not possible. We can begin to by-pass this problem, without losing the equal discrimination property, by considering the delayed Gaussian $g_{\tau}(t) = g(t-\tau)$. In the frequency domain, the attenuation of $G_{\tau}(\omega)$ is the same as for $G(\omega)$, but we now have phase shifts corresponding to a constant group delay of τ seconds. Negative values of t will still correspond to non-zero values of $g_{\tau}(t)$, but as τ is increased these values rapidly approach zero. Thus, although we may not be able to implement the Gaussian filter exactly, there is hope that it can be approximated arbitrarily closely provided that we can accept a group delay which increases with the closeness of the approximation.

In the weighting function approach discussed earlier, the finite extent of the weights limits the degree of approximation that can be achieved. The closest approach uses weights extending in time considerably beyond that period which contributes most significantly to the filter output, implying a large number of multiplications per time-step. It also implies that a filtered value representing any particular point in time must await data well beyond that time before calculation of the filter output can be made. In this way the closeness of approximation to $G_{\tau}(\omega)$ extracts its penalty of group delay, although much ingenuity has been expended in optimizing the approach to approximation (Harris 1978).

Thomson (1952) attempted to implement the highly desirable property of constant group delay in polynomial filters by arranging that the group delay was maximally flat (in the same sense as discussed earlier). The result is usually called a Bessel filter, and is shown in Fig. 4. As the order of the filter is increased, the flatness of the group delay (Fig. 4(b)), and the approximation of the attenuation characteristic to the Gaussian shape (Fig. 4(a)), extends to higher frequencies, until deviation from either Gaussian property is important only for those frequencies which are regarded as not significant, and thus are strongly attenuated by the filter. Thus the ideal of the Gaussian filter can be adequately approximated to without the need for a numerator polynomial in the transfer function.

The attenuation characteristics are much more gradual than either the Butterworth or Chebyshev filter. However, in comparison with the arithmetic-average filter we find:

(a) Within the passband, i.e. for the 'signal' frequencies, the attenuation characteristic of a Bessel filter of order greater than 3 is almost indistinguishable from that of an arithmetic-average filter; and for orders greater than about 5, the flatness of the group delay is adequate. Thus the signals of interest are essentially identical for Bessel or arithmetic-average filters. In this sense a Bessel filter of any relatively high order can be thought of as an ' n -second mean filter', where $n = 2.778697/\omega_c$, and ω_c is the 3 dB attenuation frequency in radians per second.

(b) Within the stop-band, the general level of attenuation by the Bessel filter is considerably greater than the arithmetic average, thus suppressing noise more effectively. For most purposes, extension of flat group delay to frequencies significantly higher than the 3 dB frequency is unnecessary. We are now seeing 'noise', and at the higher frequencies we have attenuation; in any case, the Bessel filter avoids the problems of phase reversal. For some applications, however, a high-order Bessel can be used to extend the flat group delay.

Thus the Bessel filter passes 'signals' in essentially the same way as the arithmetic-average filter which has the same 3 dB frequency, but attenuates high-frequency 'noise' better. In this sense the filter output is a better approximation of the underlying 'representative' value.

Some authors (e.g. Craddock and Grimmer 1960) have suggested that an 'exponential average' or time-constant filter might be applied. This corresponds to a first-order polynomial filter, whose characteristics are in Figs 1, 2 or 4 which give identical curves for first order. Attenuation is considerably less discriminating than with the higher-order Bessel or arithmetic-average filters, and the group delay is less flat in the significant frequency range.

There are other methods of approximating to a Gaussian filter using a polynomial implementation. Zverev (1967) discusses an equi-ripple (rather than maximally flat) approximation to constant group delay. For any order of filter, this gives a group delay which allows a defined amount of ripple about the mean value. In return, the 'flat' group delay extends to somewhat higher frequencies, and marginally higher frequency discrimination is available. The 'best' amount of ripple to allow depends on both 'signal' and 'noise' spectra, and the weighted significance of certain frequency bands in the context in which the data are to be used. The root-mean-square deviation between pure 'signal', and 'filtered signal plus noise' could be minimized to determine two parameters—the 3 dB frequency and the optimum amount of ripple allowed. This method could also be used to determine all the parameters available—the polynomial coefficients—without particular regard to known filter types as described above. In practice, a great deal of confidence in the representativeness of the available data, the weighting of the significance, and the penalties involved in unusual situations, would be required before such an abstract approach could replace the semi-intuitive approach described above, or be allowed to deviate too far from it. The results are not likely to be significantly different to those from a Bessel filter

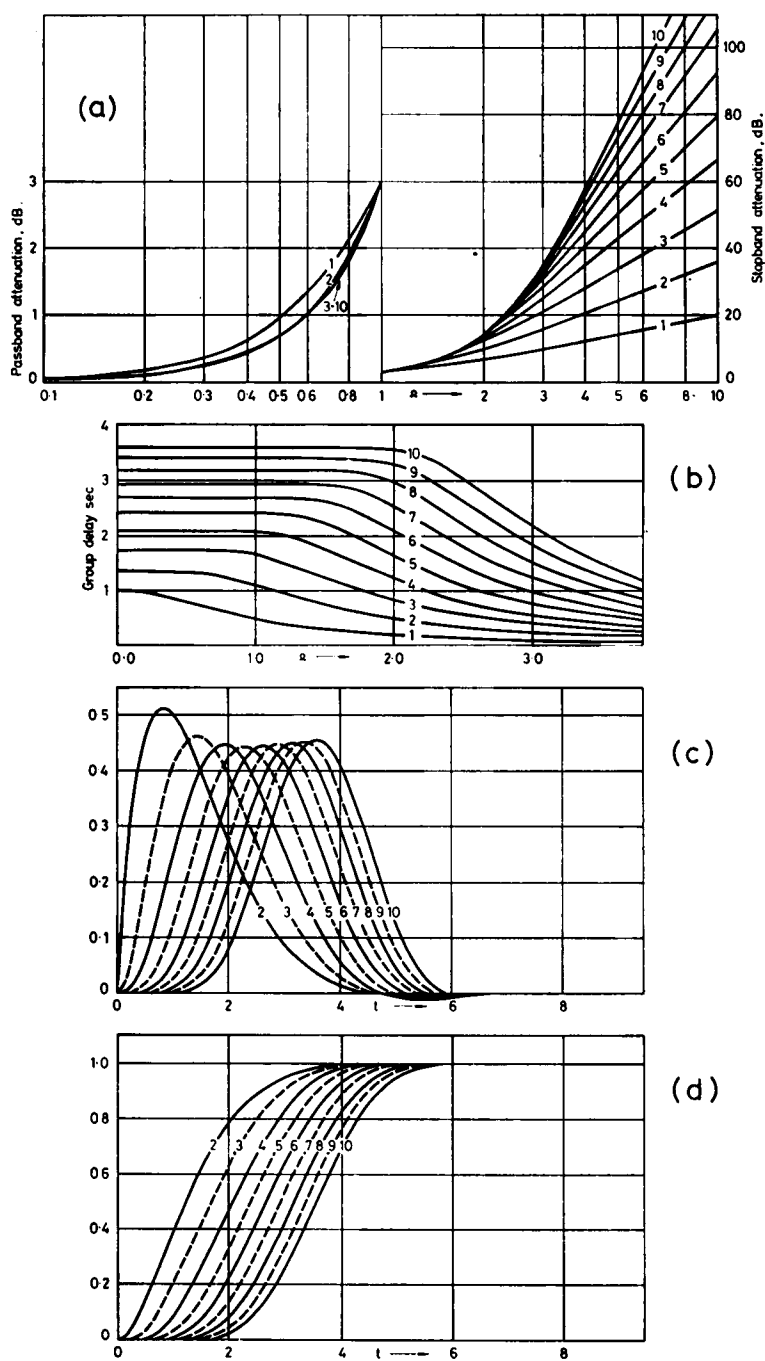


Figure 4. Bessel filter characteristics. (a) Attenuation v. frequency. (b) Group delay v. frequency. (c) Impulse response. (d) Step response.

at a suitable characteristic frequency for the application; the most important step has already been taken in moving away from a time-domain viewpoint (the arithmetic average) to a combined time- and frequency-domain viewpoint.

5. Some examples of filtering

Two examples of filtering of wind data are shown here, both based on the same sample of Cardington wind-vector data. Fig. 5 shows about 90 minutes of raw data (curve (a)), the results of a 10-minute arithmetic-average filter (curve (b)), and a 10-minute 6-pole Bessel filter (curve (c)). Curves (a) and (b) are shifted horizontally and vertically with respect to (c) to allow for group delay, and to avoid overlapping of the curves. Curve (c) commences at time zero, and the extent of the horizontal shift for (a) and (b) can be seen from their initial flat portions.




Figure 5. Wind-vector data and 10-minute averages. (a) Raw data. (b) 10-minute arithmetic-average filtering. (c) '10-minute' 6-pole Bessel filtering.

The effect of filter initialization and group delay can be seen in the first 15 minutes or so of data. Because data are delayed on filtering, the early filtered output does not correspond to real data, but rather depends on the initial values present in the filter stores. For cosmetic reasons these have been chosen to be near the expected raw data values, and this is the way to minimize the initial transient filter response. Where data are being measured continuously the transient response occurs only on first installation and does not cause problems.

Because of its poor high-frequency attenuation, the moving-mean filter exhibits structure on a time-scale significantly shorter than 10 minutes. As a consequence, values on the curve intended to give a measurement representative of the surrounding 10 minutes or so are sensitive to the exact time at which the value is chosen, i.e. the reported signal is noisy.

A highly subjective way of comparing the filters is to decide which is the closest to the curve one would

draw by eye through the data when 'smoothing' over a 10-minute interval. This process may more accurately approximate to a real human observer estimating 10-minute winds, than a machine implementing WMO rules for 'mean' wind.

Fig. 6 shows in more detail 35 minutes taken from the middle of Fig. 5. Arithmetic-average and Bessel-filtered data are also shown for 2-minute filters, suitably shifted as before. Although the raw data are the same, the different averaging time means that different structures are selected by the filter, so that a new comparison between filters is being made. Note that the arithmetic-average filter still reports noisy data.




Figure 6. Wind-vector data and 2-minute averages. (a) Raw data. (b) 2-minute arithmetic-average filtering. (c) '2-minute' 6-pole Bessel filtering.

The particular examples shown above are for the absolute value of a function. In meteorology one is often interested in the derivatives of functions. This is analogous to being interested in the shape of the contours on a chart rather than their absolute values. Examples of a potential interest in time derivatives might be:

(a) Looking for variations in wind direction or speed, or in temperature, to denote the passage of a front.

(b) Use of isallobars, being contours of derivatives of pressure. The shapes of the contours are functions of the second derivative of pressure.

(c) Estimating wind speed and direction from the values of radar range, azimuth and elevation of a balloon.

In all these examples the process of differentiation accentuates the higher frequencies and makes the filter rate of attenuation more important. In the time domain, e.g. in Figs 5 and 6, the arithmetic-average curves are rougher than the corresponding Bessel curves, so that their derivatives will be more noisy. Similar comments have even greater force for higher derivatives.

6. Poles, zeros and filter synthesis

The transfer function $F(s)$ of a continuous data filter represents the linear relationship in frequency space between the filter input $V_{IN}(s)$ and the output $V_{OUT}(s)$ as a function of the complex frequency s . For a more complete discussion of this type of description see Di Stefano *et al.* (1976). For a polynomial filter we have:

$$F(s) = V_{OUT}(s)/V_{IN}(s) = b_0/(b_0 + b_1s + b_2s^2 + \dots + b_ns^n),$$

where n is the order of the filter, and b_0 to b_n are the polynomial coefficients. This low-pass filter tends to unity gain as s tends to zero, and at high frequencies the rate of increase of attenuation depends primarily on n . The polynomial can be factorized into factors of the form $(\alpha_0 + s)$ or $(\alpha_r + i\omega_r + s)$, and the values of s obtained by putting each factor equal to zero are termed the poles of the filter—they are those complex frequencies for which $F(s)$ becomes infinite. Non-polynomial filters have a numerator polynomial in $F(s)$ and, by a similar method, those complex frequencies for which $F(s)$ becomes zero are termed zeros.

Values for poles (and zeros) which characterize various types of polynomial and other filters are to be found in books such as Zverev (1967) and Christian *et al.* (1966). For convenience, Table I shows the poles for Bessel filters up to order 8 having 3 dB attenuation at 1 radian per second. De-normalizing to give 3 dB attenuation at ω_c radians per second requires multiplying the poles by ω_c .

Table I. Bessel filter poles ($-\alpha, \omega$) normalized to 3 dB at 1 radian per second

Order	($-\alpha, \omega$)
2	(-1.10349, $\pm 0.63710i$);
3	(-1.32475, 0); (-1.04909, $\pm 1.00085i$)
4	(-1.37222, $\pm 0.41085i$); (-0.99681, $\pm 1.25913i$)
5	(-1.50470, 0); (-1.38301, $\pm 0.71904i$)
	(-0.95913, $\pm 1.47345i$);
6	(-1.57404, $\pm 0.32145i$); (-1.38405, $\pm 0.97304i$)
	(-0.93219, $\pm 1.66457i$);
7	(-1.68713, 0); (-1.61464, $\pm 0.59014i$)
	(-1.38113, $\pm 1.19353i$); (-0.91130, $\pm 1.83942i$)
8	(-1.76034, $\pm 0.27330i$); (-1.63966, $\pm 0.82417i$)
	(-1.37613, $\pm 1.39067i$); (-0.89433, $\pm 2.00166i$)

Realization of a polynomial filter corresponds to realization of a transfer function having the right poles. The complex factors of the polynomial are combined in complex conjugate pairs to form real coefficient quadratic factors in s . It is common to implement transfer functions using the first- and second-order real coefficient factors now generated into individual realizations or stages, and to realize the overall filter by passing data through each stage in turn (cascading).

It is possible to implement individual sections independently in the most convenient way. For example, the first stage of an odd-order polynomial filter might be the single real pole implemented as a simple electrical time-constant, as this is the least critical pole with respect to tolerances. This first stage does reduce the bandwidth of the raw data, so that if later stages are implemented by sampling and digitization followed by sampled data filtering, the Nyquist criterion allows the use of a lower sampling rate, and this requires less computation. Each successive sampled data filter stage reduces the bandwidth further, so that after each stage it may be possible to reduce the sampling rate further to keep the computational load light. In general, analogue techniques have the advantage of being able to handle wide-bandwidth data easily, while digital techniques have the advantage of being stable—especially when smoothing over long periods, when analogue methods lose accuracy. Thus the combination described above makes good use of both methods of implementation.

7. Hardware implementation

Sections 7 and 8 are intended to introduce the interested reader to some of the available technical literature, and provide an idea of the complexity of implementation. Without this, the advantages of polynomial implementation of Bessel-like filters are not apparent.

Implementation of filters using passive components is discussed in Zverev (1967). However, at the relatively low characteristic frequencies that are usually of interest in meteorology (say 100 Hz to 10^{-3} Hz), active electrical filters are likely to be more practical. Implementation of filters specified in terms of poles and zeros is discussed by Huelsman (1976) in a collection of significant papers together with an editorial. Implementation of a single real pole corresponds, of course, to a simple time-constant. For low-pass Bessel filters, which are relatively undemanding of component tolerances at meteorological frequencies, the simplest (Sallen and Key) filter implementation is likely to be adequate for many applications, and Shepard (1969) lists components for Bessel and other filters. Roughly speaking, one resistor, one capacitor and half an amplifier are required per pole. At the lowest characteristic frequencies, however, accuracy is likely to suffer, giving distorted signals.

8. Digital recursive implementation

The theory of sampled data systems is covered by Tou (1959). A more practical discussion of recursive implementation is presented by Gold and Rader (1969).

Standard z-transform theory (Tou 1959), shows how polynomial filter transfer functions can be expressed as recursive relations between sampled data points. Two important results from a modified theory (see Appendix) are:

(a) Implementation of a real pole:

$$F(s) = K\alpha/(s + \alpha).$$

$F(s)$ represents a time-constant of $1/\alpha$ seconds, with a low-frequency gain of K .

Using starred symbols to represent quantities sampled every T seconds, the integer n to indicate time and $:=$ to mean 'is replaced by', to implement $F(s)$ we have:

$$\left. \begin{aligned} v_{OUT}^*(nT) &:= C_1 v_{IN}^*(nT) + C_2 v_{OUT}^*\{(n-1)T\} \\ C_1 &= K\{1 - \exp(-\alpha T)\} \\ C_2 &= \exp(-\alpha T), \end{aligned} \right\} \dots \dots \dots (1)$$

where v_{IN}^* and v_{OUT}^* represent the filter input and output respectively as sampled functions of time.

(b) Implementation of a complex conjugate pole pair:

$$F(s) = K(\alpha^2 + \omega^2)/(s^2 + 2\alpha s + (\alpha^2 + \omega^2)).$$

$F(s)$ represents a low-pass second-order resonant structure (complex conjugate pole pair) with an undamped natural frequency of ω radians per second, a damping time-constant of $1/\alpha$ seconds and a low-frequency gain of K . For more information see Di Stefano *et al.* 1976.

Using starred symbols, K , T and n as before:

$$\left. \begin{aligned} v_{OUT}^*(nT) &:= C_3 v_{OUT}^*\{(n-1)T\} + C_4 v_{OUT}^*\{(n-2)T\} + C_5 v_{IN}^*\{(n-1)T\} \\ C_3 &= 2 \cos(\omega T) \exp(-\alpha T) \\ C_4 &= -\exp(-2\alpha T) \\ C_5 &= K\{1 - 2 \cos(\omega T) \exp(-\alpha T) + \exp(-2\alpha T)\}. \end{aligned} \right\} \dots \dots (2)$$

With these two recursive relations, any low-pass polynomial filter can be implemented by cascading stages.

When using any sampled system, the process of converting from continuous analogue data to sampled data (sampling) must be regarded as a significant process. Unfortunately, this is not always done in meteorology. If the sampling rate is $1/T$ samples per second, then frequencies above the Nyquist frequency $1/(2T)$ Hz in the analogue data cannot be represented (Shannon 1949, Tou 1959). The energy in these higher frequencies will appear in the sampled data under an alias frequency (aliasing). For example, a sine wave at just under $1/T$ Hz aliases as a low frequency, irreversibly corrupting the significant data. Thus, when using digital forms of filtering, the sampling rate must be high enough to ensure that only negligible energy is aliased. This is achieved either by making use of known frequency limitations of the sensor/environment (and ensuring that no other form of noise can intrude), or else by using an analogue filter before sampling. This can either implement a required real pole, or else implement a non-critical short time-constant pole which will have little effect on the overall filter.

The concept of bandwidth reduction by a filter stage, to allow operation at a lower sampling rate for the next stage, can be carried right through the recursive implementation. In fact, the low characteristic frequencies of many meteorological filters (when compared with the upper frequency limit of the raw data) makes it useful to precede the required filter by one or more spurious real-pole digital filters with suitably short time-constants, merely so that the following stages can be operated at a lower sampling rate to reduce the computational load. It is useful in this situation to use odd-order filters, and to ensure that all the real-pole time-constants involved accumulate to the required real-pole time-constant, to ensure the minimum distortion of the final filter performance.

As the actual values of the spurious time-constants are not critical, values can be chosen so that equation (1) degenerates to:

$$v_{OUT}^*(nT) := v_{IN}^*(nT) + [v_{OUT}^*\{(n-1)T\} - v_{IN}^*(nT)]/(2^m),$$

where m is an integer, so that the division can be implemented with an accumulator shift. Many low-performance microprocessors do not have a hardware multiply, so that this technique combined with reduced sampling rates can reduce the computational load to considerably below one multiplication per time-step. This option is not immediately applicable to the weighting function approach to filtering.

At some stage, multiplication by fixed non-binary numbers is required. Computational loading can be reduced by replacing each multiplication with several operations of the type 'shift right n bits, then add to (or subtract from) accumulator'. Resolution to i decimal places seems to require about i operations of the above type, considerably reducing loading over a general multiply routine. This technique is theoretically available to the weighting function approach, but the required increase in software, because of the large number of required weights, tends to preclude its use.

Detailed examination of equations (1) and (2) shows that when αT is small, the values of C_1 and C_6 tend to be small compared to unity. Thus, as v_{OUT}^* is comparable with v_{IN}^* , digital implementation using a system with a short word-length can reduce the effective resolution of the input data. This is not as catastrophic as it sounds. Provided that the noise level exceeds the effective resolution limit, the noise itself ensures that the signal is above and below the relevant digitization level for the correct proportion of time, and the output signal/noise ratio is nearly independent of the effective resolution. In fact, it is well known that high-resolution filtered signals can be extracted from suitable systems by filtering the output of a single bit digitizer (Davenport 1953).

Thus, only a small number of discrete digitization levels (although correctly spaced) are required in hardware, or early-stage software, for noisy signals such as wind direction or speed, provided that the result is to be filtered. This can lead to economies in hardware, and this effect is used in at least one

commercial wind vane which uses a small number of reed relays to sense wind direction. In practice, of course, any serious software resolution problems can be overcome by using spurious real-pole stages where αT is not too small.

9. Conclusions

In meteorology, data are often smoothed or filtered to reduce the effects of noise while measuring a representative signal. A common meteorological filter is the arithmetic average, or moving mean. A filter which gives a more accurately representative value, while retaining the ability to avoid excessive smoothing of wanted frequencies, is a Gaussian filter.

Close approximations to the unrealizable Gaussian filter can be made on sampled data using a weighting function approach. Due attention must be paid, as with all sampled data systems, to the Nyquist constraint. The weighting function approach implies a heavy computational load.

An alternative approach is to use polynomial transfer functions, which can be implemented in analogue or sampled data form. The analogue form can be implemented in simple hardware, but the sampled data (digital) form gives a more stable transfer function. A suitable approximation to the Gaussian ideal is the Bessel filter. By splitting the Bessel filter into a number of stages, each of which reduces the bandwidth required to represent the data presented to the next stage, the more critical stages can be implemented in recursive digital equations at a relatively low sampling rate. This method can be realized with a much lower computational loading on a processor, making it a practical proposition for applications such as automatic weather stations.

References

- | | | |
|---|------|--|
| Charnock, H. | 1957 | Notes on the specification of atmospheric turbulence. <i>J R Stat Soc, A</i> , 120, Part IV, 398-408. |
| Christian, E. and Eisenmann, E. | 1966 | Filter design tables and graphs. New York, John Wiley and Sons. |
| Craddock, J. M. | 1968 | Statistics in the computer age. London, English Universities Press. |
| Craddock, J. M. and Grimmer, M. | 1960 | The estimation of mean annual temperature from the temperatures of preceding years. <i>Weather</i> , 15, 340-348. |
| Davenport, W. B. | 1953 | Signal to noise ratios in band-pass limiters. <i>J Appl Phys</i> , 24, 720-727. |
| Di Stefano, J. J., Stubberud, A. R. and Williams, I. J. | 1976 | Theory and problems of feedback and control systems. New York, McGraw-Hill. |
| Gold, B. and Rader, C. M. | 1969 | Digital processing of signals. New York, McGraw-Hill. |
| Harris, F. J. | 1978 | On the use of windows for harmonic analysis with the discrete Fourier transform. <i>Proc Inst Electr Electron Eng</i> , 66, 51-83. |
| Herrero, J. L. and Willoner, G. | 1966 | Synthesis of filters. Englewood Cliffs, New Jersey, Prentice-Hall. |
| Holloway, J. L. | 1958 | Smoothing and filtering of time series and space fields. <i>Adv Geophys</i> , 4, 351-389. |
| Huelsman, L. P. | 1976 | Active RC filters: Theory and applications. New York, John Wiley and Sons. |
| Kaiser, J. F. and Reed, W. A. | 1977 | Data smoothing using low-pass digital filters. <i>Rev Sci Instrum</i> , 48, 1447-1457. |
| Lewis, P. | 1960 | The use of moving averages in the analysis of time-series. <i>Weather</i> , 15, 121-126. |
| Meyer, F. and Isacker, J. | 1975 | The synthesis of high resolution digital filters. <i>Institut Royal Météorologique de Belgique, Publications Série A</i> , No. 90. |

Passi, R. M.	1976	A weighting scheme for autoregressive time averages. <i>J Appl Meteorol</i> , 15 , 117-119.
Peacock, J. E.	1970	Some notes on the smoothing of time series. <i>Royal Observatory of Hong Kong, Technical Note (Local) No. 8</i> .
Ruthroff, C. L. and Bodtmann, W. F.	1976	Computing derivatives from equally spaced data. <i>J Appl Meteorol</i> , 15 , 1152-1159.
Shannon, C. E.	1949	Communication in the presence of noise. <i>Proc Inst Radio Eng</i> , 37 , 10-21.
Shepard, R. S.	1969	Active filters part 12: short cuts to network design. <i>Electronics International</i> , New York. 18 August, 82-91.
Thomson, W. E.	1952	Networks with maximally flat delay. <i>Wireless Eng</i> , October, 256-263.
Tou, J. T.	1959	Digital and sampled data control systems. New York, McGraw-Hill.
Zverev, A. I.	1967	Handbook of filter synthesis. New York, John Wiley and Sons.

Appendix

We wish to replace an analogue system with signals limited to frequencies below $1/(2T)$ Hz, and transfer function $F(s)$, by a sampler which samples the analogue data every T seconds, followed by a digital filter operating on these samples. Provided we are prepared to regard the filtered samples as being from a time series which is limited to frequencies below $1/(2T)$ Hz, the required z -transformed relationship between output and input sampled signals is given by:

$$V_{OUT}(z)/V_{IN}(z) = T G(z), \quad \dots \dots \dots (1)$$

where $G(z)$ is the z -transform of $F(s)$. The apparently spurious T is required to offset the gain of the sampler. A discussion of the relative gain of sampled to analogue signals and tables of z -transforms can be found in Tou (1959).

We can manipulate equation (1) and then take the inverse z -transform to obtain an equation for $v_{OUT}^*(nT)$ in terms of known quantities. Consider, for example, the real-pole transfer function:

$$F(s) = K\alpha/(s + \alpha).$$

Substituting its z -transform into equation (1) we have:

$$V_{OUT}(z)/V_{IN}(z) = K\alpha zT/\{z - \exp(-\alpha T)\},$$

and a little manipulation gives:

$$V_{OUT}(z) = K\alpha T V_{IN}(z) + \exp(-\alpha T) z^{-1} V_{OUT}(z).$$

Taking the inverse z -transform, and remembering that z^{-1} corresponds to a time-shift of T , we have:

$$v_{OUT}^*(nT) = K\alpha T v_{IN}^*(nT) + \exp(-\alpha T) v_{OUT}^*\{(n-1)T\}. \quad \dots (2)$$

In this way, standard z -transform theory can be made to give recursive equations for $v_{OUT}^*(nT)$ in terms of v_{IN}^* at various times, and previous values of v_{OUT}^* for any $F(s)$. The same results can be obtained from first principles without the use of the z -transform. The impulse response, or inverse Laplace transform, of $F(s)$ is sampled every T seconds (using the 'ideal sampler': Tou 1959) producing an infinite sampled time series. Taking the Laplace transform, using the quantity $\exp(-sT)$ as the transform of a delayed impulse, an infinite series in the frequency domain is obtained which can be summed and manipulated as above. The inverse Laplace transform gives equation (2).

This basic method of using recursive equations to represent a sampled impulse response does have an important limitation in the data smoothing context. Any analogue impulse response will contain frequencies up to infinity, but a sampled data stream cannot represent frequencies above the Nyquist frequency $1/(2T)$ Hz. Thus the sampled time-series cannot fully represent the impulse response, and the recursive equations cannot exactly implement $F(s)$. Where T is sufficiently short compared to $1/\alpha$ and $1/\omega$, there is little error; but there is an incentive to make T relatively long in the final filter stages to reduce computational 'loading'. The most important effect of this error is that if equation (2) is used directly, it gives the wrong gain at zero frequency. This can be seen by setting K equal to unity, and considering a constant signal which should give the same value for v_{IN}^* and v_{OUT}^* . To obtain this condition, the sum of the coefficients should be unity. Clearly they are not, although they approach unity for small αT . Replacing the first (gain) coefficient by $K(1 - \exp(-\alpha T))$ ensures the necessary unity sum of coefficients for unity K , and this form is given in Section 8 equation (1). The other small errors due to large T in the implementation of $F(s)$ can be tested for engineering significance using test signals, or predicted from the known complementary signal components produced by a sampler (Tou 1959).

Similarly, standard implementation of the complex conjugate pole-pair transfer function in Section 8 gives a similar recursive relationship to equation (2) of Section 8, except that C_5 appears as:

$$KT\{(\alpha^2 + \omega^2)/\omega\} \exp(-\alpha T) \sin(\omega T).$$

The form given in Section 8 is also modified to ensure unity sum of coefficients for unity K , and hence correct gain.

The necessity for ensuring unity sum of coefficients if unity gain is required should be taken right through to binary implementation of the coefficients. The limited resolution often available for C_2 or C_5 can mean that an error in the least significant bit has noticeable effects.

The accuracy of London Weather Centre forecasts of surface wind and total wave heights and their comparison with computer products

By R. M. Morris

(London Weather Centre)

Summary

Forecasts of wind speed and total significant wave height for spot locations on the Continental Shelf are part of the continuous service provided by London Weather Centre for the Offshore Industry. A check on the accuracy of these forecasts has been carried out over a three-year period for a number of sites where verifiable data are considered to be sufficiently reliable. Forecasters receive considerable numerical support from the operational 10-level model in the prediction of winds and waves and a check has also been made upon the accuracy of the computer products and results compared with the London Weather Centre 'man-machine mix' versions.

Introduction

London Weather Centre (LWC) is responsible for the preparation of weather and sea-state forecasts and their dissemination to the Offshore Industry for any part of the Continental Shelf. Most of the forecasts are for specific fixed sites, e.g. production platforms and drilling rigs, but forecasts are also prepared for tows of rigs and barges to and from the offshore sites.

The fixed-site forecasts usually consist of wind speed and direction (10-minute mean at 10 metres and also at deck level, with appropriate gusts), visibility, weather, cloud, total significant wave height* and swell direction, height and period, for forecast periods up to five days ahead. Undoubtedly the most important of these elements for on-site operations are wind speed, visibility, total significant wave height, and the range of significant wave periods. The forecasts are issued twice a day with almost all companies taking an early morning issue between 0600 and 0800 local time, but the second issue is split between those companies that prefer to have an afternoon forecast on their shore-base desks before 1700 and those companies that are able to take a more realistic and up-to-date issue between 1800 and 2000.

Until the middle of 1977 the forecasts were prepared using the routine 24-hour surface prognoses issued by the Central Forecasting Office (CFO) at Bracknell every 6 hours and the twice-daily issues of 48- and 72-hour surface prognoses. The sequence of wind evolution would be interpolated between current analyses and the winds deduced from the forecast (surface-pressure) charts. Routine conferences between the forecasters in CFO and the offshore forecasters at LWC facilitated the exchange of views on the synoptic developments and communicated the degree of confidence felt by CFO in the prognoses.

The prediction of sea state was based entirely on graphical methods derived by Darbyshire and Draper (1963) and later upon techniques contained in a comprehensive WMO publication (World Meteorological Organization 1976). Essentially, the use of these graphical techniques requires a knowledge of the wind speed, duration and fetch at specific points in order to predict the growth and decay of waves.

* Total significant wave height is defined as the mean of the highest one-third of all the waves observed during a period of 12-15 minutes and necessarily refers to the combined effects of wind-sea and swell; the expression 'total seas' is sometimes used in this paper as a convenient brief synonym.

In August 1977 LWC began receiving directly from COSMOS (the Meteorological Office computing system) computed values of surface wind speed and direction at each grid point of the fine-mesh version of the 10-level model (Benwell *et al.* 1971). These winds are derived from the 900 mb winds in the model, using the relationships described by Findlater *et al.* (1966). The forecast wind fields are in 3-hour steps up to 36 hours and are derived twice daily from the early operational fine-mesh model computations at about 0300 GMT and 1500 GMT.

Numerical computation of wind-waves and swell over the Continental Shelf also became available to LWC on an operational basis (once a day initially, but later twice daily) in spring 1978. The computation of sea state is based on the work of Golding (1979) and takes into account variations of depth. Values of wave height, direction and period are calculated for each of the fine-mesh grid points and forecast values are produced in 6-hour steps up to 36 hours.

More recently, in the spring of 1979, surface wind forecasts were derived from the octagon version of the 10-level model, covering the period from 36 to 72 hours ahead in 12-hour intervals, and these winds are sent to LWC twice daily some one or two hours after the fine-mesh data. All the numerical products are received at LWC in digital code and graphical chart formats by dedicated teleprinter and facsimile lines respectively.

The derivation of the 'man-machine mix' product is a rather complex process. Continental Shelf charts (1 in 3×10^6 scale) are plotted and analysed every 3 hours and sea-state analyses on a similar scale and for a similar area are prepared every 6 hours on the offshore bench at LWC. First of all the forecaster will compare the computer analysis field with his own hand-drawn analysis to look for errors in small detail. He will also compare the early stages of the computer prediction fields with his own latest updated analyses to evaluate the early trends. Similarly, the forecaster will check the computer sea-state analyses with his own hand-analysed charts and will also check the forecasts for certain specific locations using the empirical graphical techniques.

Since the forecaster is predicting conditions for individual rig or platform locations, there is no doubt that for the first 6–12 hours of any period the skilled forecaster should be able to improve on a numerical model output, if necessary by his superior appraisal of the fine detail. Beyond 12 hours the forecaster's modifications to the model depend upon how well he thinks the model has predicted the early stages and also how good he thinks the model is at predicting particular types of weather pattern. If the forecaster makes modifications to the wind forecasts then, clearly, he must modify the sea-state forecasts too.

It is clearly essential to know how useful the numerical guidance has become to the offshore forecasters at LWC. Unfortunately, staff availability did not permit verification of offshore forecasts to be carried out objectively before the summer of 1977 and hence comparisons of forecast performance before and after the advent of numerical guidance are not available. However, a comprehensive verification scheme has been fully maintained since September 1977 with the help and co-ordination of the offshore forecasters.

Verification is important for two fundamental reasons:

- (a) It reveals how useful forecasts are to the user and highlights certain strengths and weaknesses of both forecaster and forecast.
- (b) It brings out the qualities and deficiencies of the numerical model, which not only helps the forecaster but also maintains an effective feedback from the forecaster to the research modeller who needs the information to improve his product.

Probably the most important reason of all for verification is the fact that forecasts are sold under commercial contract. It is essential to know how good (or poor) the forecasts are in order to sell the forecast service to a market that is not obliged, of necessity, to take the products of the state weather service.

Method and scope

Offshore verification is carried out on five sets of data but, because of the continual development of the output program since August 1977, the verification is not uniform. Table I displays the amount of verification carried out on each set of data so far.

Table I. *Amount of verification carried out on five sets of data for offshore forecasts*

	Forecast period					
	T + 0	T + 12	T + 24	T + 36	T + 48	T + 72
Fine-mesh winds	Sept. 1978 to date	Sept. 1977 to date	Sept. 1977 to date	Sept. 1977 to date		
Octagon winds				May 1979 to date	May 1979 to date	May 1979 to date
Fine-mesh total seas	Sept. 1978 to date	May 1978 to date	May 1978 to date	May 1978 to date		
LWC winds		Sept. 1977 to date	Sept. 1977 to date		Sept. 1977 to date	Sept. 1977 to date
LWC total seas		Sept. 1977 to date	Sept. 1977 to date		Sept. 1977 to date	Sept. 1977 to date

Verification of the numerical products is carried out at six grid points (total seas verified at five points prior to September 1978), three of which are in the North Sea and a fourth corresponds to the national data buoy DB 1. The fifth and sixth points are different for winds and total seas. Since the wind verification started before that of the total seas, two points were chosen west of Ireland and north-west Scotland within the Continental Shelf to complete the circle of points. However, it seemed more prudent to verify the total seas at the Ocean Weather Stations 'L' and 'R' in view of the general lack of sea-state data immediately west of the British Isles. The three points in the North Sea were chosen to coincide with the main areas of offshore activity where it was also convenient to verify the LWC forecast issues.

The method of carrying out the verification needs to be carefully explained. Since the verification is for fixed times there is no difficulty in extracting the computer wind and total seas data from the digital code. The LWC forecasts are for 12- and 24-hour periods, which means that interpolation is required to extract the appropriate figures. In practice this is usually not difficult because, if the forecaster gives a trend through a period, invariably the following period will start with the values appropriate to the end of the previous period, e.g. forecast 0800–2000, 20 kn increasing to 30 kn and later to 35 kn; forecast 2000–0800, 35 kn, decreasing slowly to 25 kn. Interpolation would yield: 1200, 30 kn; 1800, 35 kn; 0001, 30 kn; 0600, 25 kn. The data used to verify the forecasts are derived from the routine surface and sea-state analyses prepared by the offshore forecasters. These analyses embrace the Continental Shelf ($1 \text{ in } 3 \times 10^6$ scale) at 3-hourly (weather) and 6-hourly (sea state) intervals. The quality of offshore observations is not a subject that can be dealt with in a few sentences, but it should be recognized that offshore forecasters acquire a considerable amount of skill in interpreting observations from familiar installations.

It is realized that verification for fixed times will tend to penalize forecasts which often contain useful and important details on the trend of wind and wave heights. Not only are these details ignored but timing errors tend to be more heavily penalized than is perhaps justified. Ideally, the verification should be carried out at 6-hourly intervals during the first 24 hours and perhaps at even 3-hourly intervals during the first 12 hours. This problem can also be tackled by plotting graphs of forecast values against observed values.

The three North Sea verification points were chosen in regions where certain observations are considered reliable and are also supported by good quality data nearby. In general the quality control of

wind is much easier than that of sea state, mainly because the latter tends to be a mixture of instrumental and visual observations. Every attempt has been made in this verification work to rely on known instrumental observations with the possible exception of sea state at the Ocean Weather Stations 'L' and 'R'.

Verification is confined to two parameters, wind speed and total significant wave height. Verification of wave-period forecasts would be extremely useful but this can be adequately carried out only by examination of the wave-trace charts. This sort of analysis should be possible with data from the buoy DB 1 when they become available.

Results

It will be convenient to discuss winds and waves separately whilst making comparisons between computer guidance and the LWC forecasts. These are not independent, of course, and one of the aims of this study is to ascertain how much improvement the forecaster can make to direct computer predictions.

(a) Winds

(1) *Computer forecasts.* For six months there were two types of verification, a qualitative assessment of the chart as a whole and a quantitative assessment at six grid points. The qualitative assessment was based upon a scale of six divisions, defined as follows:

- A: Good guidance.
- B: Feature movement too slow, good wind speeds.
- C—: Good feature movement, underestimated wind speeds.
- C+: Good feature movement, overestimated wind speeds.
- D: Poor feature movement, poor wind speeds.
- E: Major evolution errors.

The results of the six-month assessment are shown in Table II and clearly indicate why the wind fields are highly regarded by forecasters, though with a reservation that wind speeds tend to be underestimated. This method of assessment was discontinued when it became clear that results were constant.

Table II. *Chart assessment summary for the period from September 1977 to February 1978, expressed in percentage of total marks allocated for each forecast period*

Assessment	Forecast period					
	T+6	T+12	T+18	T+24	T+30	T+36
A	54	44	39	37	36	32
B	1	2	4	4	6	3
C—	41	46	47	45	41	42
C+	1	2	1	1	3	3
D	2	4	6	9	10	13
E	1	2	3	4	4	7

The quantitative verification is presented in two ways. Mean modulus errors over all wind speeds for each grid point are contained in Table III. The data are for up to and inclusive of August 1980 and it should be noted that the rectangle period started in September 1977 and the octagon period started in May 1979. Errors were calculated up to one decimal place before rounding off in Table III.

Table III. *Mean modulus wind speed errors in knots*

Position	Forecast period						
	T+0	Rectangle		T+36	T+36	Octagon	T+72
		T+12	T+24			T+48	
61°N 02°E	4	5	5	5	6	6	8
60.5°N 06°W	4	5	5	5	5	5	6
57.5°N 03°E	4	4	5	5	5	6	7
53.5°N 04°E	4	4	5	5	5	5	6
55.5°N 10°W	5	5	5	5	5	5	6
48.5°N 09°W	4	4	5	5	4	5	5

Table III reveals a marked uniformity in wind speed errors as regards location and irrespective of forecast period, except for T + 72 in the northern North Sea. Mean errors of this magnitude are extremely good. However, it is important to know how these errors relate to the actual wind speeds reported.

Table IV reveals the errors with respect to ranges of wind for two points. The upper section of each box refers to 61°N 2°E and the lower section refers to 48.5°N 9°W.

Considering the fine-mesh (rectangle) figures first, there are three main points to make:

- (i) The relatively small mean errors in Table III are due to the high frequency of observed wind speed below 20 kn.
- (ii) The errors do not increase significantly with increasing forecast period even at high wind speeds (excepting from T + 12 to T + 24, range 30–39 kn at DB 1).
- (iii) Mean errors are roughly 20–25% of observed wind speed at all wind speeds.

Table IV. *Errors (in knots) with respect to ranges of wind speeds for two points. For each forecast period the upper figures refer to 61°N 2°E and the lower figures refer to 48.5°N 9°W.*

Forecast period	Ranges of observed wind speeds							
	0–19		20–29		30–39		≥ 40	
	No. of occns	Mean error	No. of occns	Mean error	No. of occns	Mean error	No. of occns	Mean error
Rectangle—September 1979–August 1980								
T+0	475	4	179	5	59	9	12	13
	665	4	127	3	20	5	3	10
Rectangle—September 1978–August 1980								
T+12	1007	4	330	6	104	9	18	11
	1139	5	243	5	43	5	7	12
T+24	1005	5	326	6	102	8	20	9
	1142	5	254	5	35	7	7	12
T+36	1005	5	323	7	100	11	18	11
	1137	5	254	5	43	9	7	19
Octagon—May 1979–August 1980								
T+36	723	5	196	10	59	14	11	18
	832	4	130	6	21	9	2	21
T+48	722	5	195	9	59	15	12	20
	828	4	131	7	21	10	3	19
T+72	721	6	190	10	63	18	11	27
	833	5	127	8	21	12	2	29

Turning next to the octagon figures and noting that the number of cases is much less than that of the fine-mesh set, we find that there are also three main points to make:

(i) At $T + 36$ the octagon errors tend to become significantly greater than the rectangle errors at higher wind speeds at $61^{\circ}\text{N } 2^{\circ}\text{E}$ but not at $48.5^{\circ}\text{N } 9^{\circ}\text{W}$.

(ii) Octagon errors increase with increasing period of forecast and more substantially with increasing wind speeds, especially above 30 kn.

(iii) Mean octagon errors are generally 30% and perhaps 50% in the northern North Sea for observed wind speeds over 20 kn.

Undoubtedly the octagon wind errors are not simply due to a loss of wind speed in the model; phase and evolution errors must be increasingly important at extended forecast periods.

(2) *LWC forecasts.* Table V contains mean forecast wind speed errors for observed wind speed

Table V. Mean modulus wind speed errors (in knots) in forecasts issued by London Weather Centre from September 1977 to August 1980

Ranges of observed wind speeds <i>knots</i>													All speeds
		0-9		10-19		20-29		30-39		≥ 40			
Position	Forecast period	No. of occons	Mean error	No. of occons	Mean error	No. of occons	Mean error	No. of occons	Mean error	No. of occons	Mean error	Mean error	
Statfjord 61°N 2°E	T+12	299	7	873	5	670	6	280	6	55	7	6	
	T+24	303	9	873	6	674	6	282	6	55	9	6	
	T+48	302	11	870	7	676	6	282	7	55	14	7	
	T+72	302	12	870	7	676	6	282	9	55	14	8	
Montrose 57.5°N 1.5°E	T+12	237	6	894	5	680	5	215	5	41	5	5	
	T+24	237	7	894	6	683	5	217	6	41	7	6	
	T+48	237	9	893	6	683	6	217	8	41	15	7	
	T+72	237	11	893	7	681	7	217	9	41	15	8	
Placid 53.5°N 2°E	T+12	335	6	1107	5	589	4	140	5	22	5	5	
	T+24	341	7	1107	5	589	5	140	5	22	10	5	
	T+48	344	9	1107	5	589	6	140	8	22	16	6	
	T+72	344	10	1107	5	589	6	140	10	22	17	6	

ranges for forecasts issued by LWC to three locations in the North Sea. Table V contains an extra range below 20 kn compared with Table IV. There are some interesting features which may be listed as follows:

(i) Not surprisingly, errors are greatest at both low and high observed wind speeds and it is clearly almost as difficult to predict very light winds as to predict very strong winds for periods beyond 24 hours ahead.

(ii) For the broad range 10-39 kn the forecast errors are impressively small.

(iii) For all wind speeds mean errors increase slowly with increasing forecast period.

(3) *Comparison of computer forecasts with LWC forecasts.* The only comparison between computer forecasts and LWC forecasts is available for $61^{\circ}\text{N } 2^{\circ}\text{E}$ (Statfjord) and even here the length of verification period is different (cf. Tables IV and V). Even so, there are probably enough data to show whether

the forecaster is making a significant improvement to the computer guidance. An illustration is shown in Table VI which contains comparisons at the Statfjord Field (61°N 2°E) for two ranges of observed wind speed. The variations in verification period should be borne in mind, particularly for the octagon winds.

Table VI. Mean modulus errors (in knots) of wind speeds in the ranges 0–20 and 30–39 knots at 61°N 2°E with respect to period of forecast

Method	Range of wind speeds (kn)	Forecast period				
		T+12	T+24	T+36	T+48	T+72
Rectangle	0–20	4	5	5	—	—
	30–39	—	—	5	5	6
Octagon	0–20	9	8	11	—	—
	30–39	—	—	14	15	18
LWC forecaster	0–20	6	7	—	8	8
	30–39	6	6	—	7	9

Table VI reveals that for observed light winds the computer guidance is slightly superior to the LWC forecasts, but for observed strong winds the LWC forecasts are significantly better and substantially so at extended forecast periods. The computer products undoubtedly tend to underforecast wind speeds and this fact probably explains the computer superiority at low wind speeds. Table VI suggests, then, that the LWC forecasters are making a positive improvement to the rectangle wind forecasts (up to 36 hours ahead) in strong wind situations with some help from the CFO 24-hour surface prognostic chart. Since the LWC forecaster is dependent upon discussion with the CFO medium-range forecaster for forecast periods beyond 36 hours, then Table VI presents indirect evidence that the medium-range forecaster at CFO is producing a superior surface prognosis than is produced by the computer for 48 hours and 72 hours ahead.

(4) *Trends in accuracy.* The verification results examined so far do not contain any discrimination for time of year. Verification has been carried out in two-calendar-month periods from March 1978 (the periods from September 1977 to February 1978 were combined) and there are probably just enough data to identify any marked seasonal bias. Tables VII and VIII show the trend of wind forecast accuracy at 61°N 2°E for both light and strong winds.

Table VII. Trend of accuracy in forecast surface wind speeds issued by LWC forecasters for 61°N 2°E. Mean errors with respect to observed wind speed range 0–19 knots

Forecast period	Sept. '77– Feb. '78	Mar.– Apr. '78	May– June	July– Aug.	Sept.– Oct.	Nov.– Dec.	Jan.– Feb. '79	Mar.– Apr.	May– June	July– Aug.	Sept.– Oct.	Nov.– Dec.	Jan.– Feb. '80	Mar.– Apr.	May– June	July– Aug.
T+24	9	8	5	6	8	11	7	6	5	4	6	8	8	7	5	5
T+48	11	9	5	6	9	10	8	7	6	4	7	11	10	8	6	6
T+72	11	9	6	7	10	11	9	9	7	6	9	14	11	9	7	6

Table VIII. Trend of accuracy in forecast surface wind speeds issued by LWC forecasters for 61°N 2°E. Mean errors with respect to observed wind speed range 30–39 knots

Forecast period	Sept. '77– Feb. '78	Mar.– Apr. '78	May– June	July– Aug.	Sept.– Oct.	Nov.– Dec.	Jan.– Feb. '79	Mar.– Apr.	May– June	July– Aug.	Sept.– Oct.	Nov.– Dec.	Jan.– Feb. '80	Mar.– Apr.	May– June	July– Aug.
T+24	6	6	7	7	7	7	6	7	6	6	5	5	7	3	10	2
T+48	7	5	14	15	7	5	11	6	11	12	7	6	10	3	10	2
T+72	8	8	16	12	8	6	11	8	13	15	9	8	11	8	13	5

The forecasters' ability to predict light winds shows no significant trend other than that it is more difficult in winter than in summer (from May to August) with mean errors about double the values in other seasons.

(b) *Total significant wave height*

(1) *Computer forecasts.* Table IX contains mean errors of forecast total significant wave height at each of the six verification points. The most striking aspects of the results are as follows:

(i) Except locally at very high seas the forecast errors are no greater than analysis errors; neither do the errors tend to increase with increasing forecast periods.

(ii) At practically all locations, up to 9-metre seas, forecast errors do not increase greatly with increasing observed seas. In fact there is a tendency for errors to become proportionally smaller with higher observed seas.

The T + 0 total seas analyses are derived entirely from the wind fields and although some of the error must be due to error in the wind analysis itself, forecasters' experience with the computer output suggests that the modelling of swell may be the chief cause of these errors. The analysis of swell is a highly subjective process, although forecasters with long experience can usually detect the presence of swell from a series of observations at a fixed location. Impressive as the results are in Table IX it is generally

Table IX. *Rectangle total seas verification, March 1978–August 1980. Mean modular differences in half-metre units. Height ranges in actual significant waves.*

Position	Forecast period	Height range							
		0–3 m	Mean modular diff.	3·1–5·5 m	Mean modular diff.	5·6–9·0 m	Mean modular diff.	>9·0 m	Mean modular diff.
Grid 3–5 (Statfjord)	T+0	757	1·3	272	1·7	23	1·4	4	6·3
	T+12	938	1·3	316	2·0	30	2·2	4	7·6
	T+24	938	1·3	316	1·9	28	3·2	4	8·3
	T+36	938	1·3	318	2·2	28	3·5	4	7·9
Grid 4–6 (Forties)	T+0	834	1·2	188	1·3	14	1·7	0	—
	T+12	992	1·1	205	1·8	14	2·1	0	—
	T+24	1057	1·2	203	2·0	14	3·3	0	—
	T+36	1065	1·3	209	2·3	14	4·3	0	—
Grid 5–7 (Placid)	T+0	1027	1·0	63	1·3	0	—	0	—
	T+12	1169	1·0	75	1·5	0	—	0	—
	T+24	1259	1·1	69	1·6	0	—	0	—
	T+36	1248	1·2	74	1·6	0	—	0	—
Grid 6–2* (OWS 'L')	T+0	523	1·7	405	2·2	59	1·8	7	5·4
	T+12	600	1·7	468	2·0	78	2·5	8	7·0
	T+24	656	1·7	485	2·1	76	2·5	8	7·9
	T+36	660	1·7	483	1·9	70	2·9	9	7·5
Grid 9–4† (OWS 'R')	T+0	461	2·6	227	2·5	55	3·4	10	5·7
	T+12	626	2·0	290	2·4	58	4·4	14	6·2
	T+24	637	2·1	269	2·2	68	4·4	15	6·0
	T+36	631	2·0	283	2·4	58	4·7	12	6·2
Grid 8–5 (DB 1)	T+0	666	1·6	249	2·1	37	2·9	3	1·4
	T+12	642	1·7	263	2·1	31	2·9	3	3·7
	T+24	696	1·7	258	1·9	36	3·0	3	4·5
	T+36	692	1·6	266	1·7	31	2·8	3	4·9

* Excludes March–April 1980.

† March 1978–February 1980 only.

believed by forecasters that underforecasting of the wind seas by the model (associated with under-forecast wind fields) is compensated by overforecasting of swell. Unfortunately, only careful scrutiny of the data tapes from DB 1 will confirm or deny these views, since no other verification point can include measurements of direction in the wave height and period data.

Table X reveals the trend of forecast performance through the verification period. The $T + 0$ errors have a distinct minimum during May–June which is probably the period of least cyclonic activity over the North Atlantic and therefore the period of minimum swell generation. It is noticeable that the minimum errors are maintained through July–August at Statfjord but not at DB 1. This latter period was markedly cyclonic over the east Atlantic and the United Kingdom. There is no evidence from Table X that forecast performance has improved over the period, nor is there any suggestion that better analyses produce better forecasts.

Table X. *Trend of mean errors in computer forecasts of total significant wave heights in the range 3–5.5 metres (errors in half metres)*

Position	Forecast period	1978					1979					1980				
		Mar.–Apr.	May–June	July–Aug.	Sept.–Oct.	Nov.–Dec.	Jan.–Feb.	Mar.–Apr.	May–June	July–Aug.	Sept.–Oct.	Nov.–Dec.	Jan.–Feb.	Mar.–Apr.	May–June	July–Aug.
Statfjord	$T + 0$	—	—	—	—	2.2	1.4	2.4	1.0	1.0	1.6	1.8	1.5	1.6	1.1	1.4
	$T + 24$	2.5	1.2	1.3	1.4	1.9	1.4	2.2	1.4	2.5	1.9	2.1	1.7	1.9	2.3	2.4
DB 1	$T + 0$	—	—	—	—	2.0	1.7	2.1	1.1	2.8	1.7	2.9	2.3	2.0	0.8	2.1
	$T + 24$	—	—	—	2.7	1.9	1.5	1.4	1.2	1.0	1.8	2.7	2.2	1.6	1.4	2.0

(2) *LWC forecasts.* Table XI contains mean errors of forecast total significant wave height in forecasts issued by forecasters at LWC for three main locations in the North Sea. (Forecasts for periods beyond 36 hours are based upon graphical techniques using the forecast wind speed, fetch and duration, and taking into account the general development of the sea state during the first 36 hours of the period.)

Table XI. *Verification of LWC forecasts of total significant wave heights during the period September 1977–August 1980. Mean modular differences in half-metre units. Height ranges in actual significant waves.*

Position	Forecast period	0–3 m				Height range 3.1–5.5 m				>5.5 m			
		No. of cases	Mean modular diff.	PERS	CLI	No. of cases	Mean modular diff.	PERS	CLI	No. of cases	Mean modular diff.	PERS	CLI
Statfjord	$T + 12$	1415	1.0	1.0	1.9	652	1.5	1.6	2.3	124	2.6	2.6	7.2
	$T + 24$	1410	1.2	1.4		659	1.9	2.3		124	3.1	4.9	
	$T + 48$	1410	1.4	1.9		659	2.0	2.6		124	4.8	5.2	
	$T + 72$	1412	1.7	2.0		659	2.3	2.6		124	5.6	6.0	
Montrose	$T + 12$	1600	1.1	1.0	1.6	412	1.7	2.2	2.6	50	2.7	3.6	7.5
	$T + 24$	1600	1.3	1.4		412	1.9	2.7		50	3.9	6.1	
	$T + 48$	1600	1.5	1.5		412	2.2	3.1		50	5.7	6.8	
	$T + 72$	1600	1.6	1.9		412	2.8	3.3		50	6.5	7.4	
Placid	$T + 12$	1959	0.9	0.8	1.2	160	1.2	1.8	3.9	8	1.5	3.5	9.3
	$T + 24$	1959	1.0	1.3		160	1.4	2.5		9	2.5	5.1	
	$T + 48$	1959	1.2	1.4		160	2.3	2.9		9	5.5	6.5	
	$T + 72$	1959	1.4	1.6		160	2.9	3.3		9	6.0	6.8	

In Table XI, London Weather Centre mean forecast errors are compared with mean errors that would have occurred using persistence (PERS) and climatology (CLI); errors are in units of half-metres. The main features of Table XI are:

(i) Below 5.5 m waves, forecaster errors show only a small tendency to increase with increasing forecast period.

(ii) Below 5.5 m waves, forecaster errors do not markedly increase with increasing observed wave heights.

(iii) Above 5.5 m waves, forecaster errors are greater but still probably below 30% except for the longer-range forecasts.

(iv) Apart from the first 12-hour forecast period, LWC forecasters are on average superior to persistence forecasts, particularly when waves are high.

(v) Climatology forecasts are substantially inferior to LWC forecasters when waves are high.

Table XII shows the trend in LWC forecaster accuracy throughout the verification period, for one range of wave heights for one location. There is no evidence of any significant trend but it is interesting to note that for T + 72-hour forecasts the largest errors are on average occurring in early summer. This is the time of shortest wavelengths in the tropospheric circulation and small errors in forecast synoptic evolution may make relatively large errors in wind and associated sea state.

Table XII. *Trend of mean errors in forecasts of wave heights in the range 3-5.5 metres for Statfjord (errors in half-metres)*

Forecast period	1977			1978			1979			1979			1980			1980		
	Nov.- Dec.	Jan.- Feb.	Mar.- Apr.	May- June	July- Aug.	Sept.- Oct.	Nov.- Dec.	Jan.- Feb.	Mar.- Apr.	May- June	July- Aug.	Sept.- Oct.	Nov.- Dec.	Jan.- Feb.	Mar.- Apr.	May- June	July- Aug.	
T + 24	1.8	2.2	2.0	1.8	0.8	2.0	2.0	2.0	2.3	2.0	2.1	1.9	2.1	2.2	2.2	1.6	2.4	
T + 48	2.2	2.2	2.4	3.6	1.4	2.4	1.7	2.4	2.3	2.6	1.7	1.8	1.9	2.4	2.6	2.2	2.6	
T + 72	2.2	2.6	2.8	4.0	1.6	2.6	1.7	2.5	2.5	2.4	3.2	2.1	2.1	2.4	2.2	3.0	2.4	

Table XIII compares the computer forecast errors with the LWC forecaster errors for the Statfjord location. This Table reveals that forecasters are usually making a positive contribution by improving upon the computer guidance. However, the differences are not great.

Table XIII. *Computer forecast errors compared with LWC forecaster errors for the Statfjord location (errors in half-metres)*

Forecast period	T + 12			T + 24			T + 36			T + 48			T + 72		
	0-3	3-5.5	>5.5	0-3	3-5.5	>5.5	0-3	3-5.5	>5.5	0-3	3-5.5	>5.5	0-3	3-5.5	>5.5
Method															
Rectangle	1.3	2.0	2.7	1.3	1.9	3.8	1.3	2.2	3.7	—	—	—	—	—	—
LWC forecaster	1.0	1.5	2.6	1.2	1.9	3.1	—	—	—	1.4	2.0	4.8	1.7	2.3	5.6

(c) Concluding remarks

Although the verification results support the assertion that forecasters are making a positive contribution to the computer guidance in the prediction of wind and waves up to 72 hours ahead, the generally high performance of the rectangle (up to 36 hours ahead) cannot be overstressed. There are occasions when forecasters are under considerable pressure due either to bad or complicated weather conditions or to the volume of forecast enquiries and requests. Under these circumstances the need to have readily available detailed quantified objective predictions of wind and sea state around the Continental Shelf is absolutely essential and the computer guidance provides invaluable material for this purpose.

The verification of wind and sea-state forecasts is maintained continuously at LWC.

References

- | | | |
|---|------|---|
| Benwell, G. R. R., Gadd, A. J., Keers, J. F.,
Timpson, Margaret S., and White, P. W. | 1971 | The Bushby-Timpson 10-level model on a fine mesh. <i>Sci Pap, Meteorol Off</i> , No. 32. |
| Darbyshire, Mollie and Draper, L. | 1963 | Forecasting wind-generated sea waves. <i>Eng</i> , 195, 482-484. |
| Findlater, J., Harrower, T. N. S.,
Howkins, G. A. and Wright, H. L. | 1966 | Surface and 900 mb wind relationships. <i>Sci Pap, Meteorol Off</i> , No. 23. |
| Golding, B. W. | 1979 | Computer calculations of waves from wind fields. <i>Proceedings of Institute of Mathematics and its Applications Conference on Power from Sea Waves, Edinburgh, 28 June 1979</i> , 115-134. |
| World Meteorological Organization | 1976 | Handbook on wave analysis and forecasting. Geneva, WMO No. 446. |

Review

Weather modification: prospects and problems, by G. Breuer (translated from the German by H. Mörrth). 220 mm × 140 mm, pp. xiii + 178, *illus.* Cambridge University Press, London, 1980. Price £10.50.

This book reviews the science and practice of weather modification and discusses its potential impact upon society. The author has consulted a wide range of workers in this field and writes more in the style of a reporter than from the viewpoint of an active participant. According to the dust-jacket the presentation is aimed at informing the non-scientist.

Chapter 1 describes the scientific and technical background to the practice of weather modification. It opens with an extremely compressed account of the factors governing the general circulation and the predictability of the atmosphere. One of the principal difficulties in the evaluation of weather modification techniques is the separation of real effects from the natural variability of the atmosphere. The statistical nature of this problem is clearly brought out by the author. The chapter concludes with a brief but reasonably clear description of cloud microphysical processes.

Chapter 2 provides a comprehensive review of weather modification, both scientific experiments and commercial schemes. Many untested proposals are also discussed. An interesting account of the history of cloud seeding is followed by a discussion of 'warm' and 'cold' fog dispersal, rainfall enhancement, hail and lightning prevention, and the modification of tropical storms. The chapter concludes with a discussion of inadvertent and intentional climate modification. The presentation is often uncritical and, therefore, very dependent upon the quality of the source material. For example, despite the reasoned discussion on the need for statistical proof in Chapter 1, on page 61 the author states: 'Since farmers, who are known to be good with figures, year after year employ the services of weather modifiers for working on cumulus clouds, it must be assumed that they do so on the grounds of their experience, irrespective of the fact that a firm statistical proof of success—or failure—of these operations can generally not be produced'.

The third and concluding chapter is concerned with the legal, social, ecological, and military implications of weather modification. The quality of the presentation is again very variable. There is an interesting discussion on the difficulty of performing cost-benefit calculations and also on potential conflicts of interest—for example between the farming and tourist industries. Unfortunately, a few pages later one finds the fanciful suggestion that present technology is sufficiently advanced to make it

possible, at a reasonable expenditure, to clear fog from cities, in order to provide the inhabitants with a few hours of extra sunshine. The author makes the valid point that if weather modification is to be controlled effectively, international agreements should be established before it is widely used. However, the overall impression left by Chapter 3 is that, considering the limited proof of the efficacy of most weather modification techniques, it would be more profitable to worry about existing world problems than about hypothetical new ones.

Weather modification will be of interest to those who wish to acquaint themselves with a wide range of proposals for modifying the weather on all scales. They may also appreciate the copious references. However, anybody interested in a critical scientific discussion of weather modification techniques will be disappointed.

R. Brown

Mr Frederick James Parsons, M.B.E., M.A.

The death occurred at his home in Porthcawl, on 10 December 1980, of Mr Frederick James Parsons, M.B.E., M.A. He was 89 years of age.

Mr Parsons commenced weather observing at the County Observatory, Ross-on-Wye, in July 1914. Apart from absence on war service during the First World War, he maintained a program of meteorological observations until his retirement on 1 June 1975. Meteorological readings from Ross-on-Wye commenced in 1859 and Mr Parsons succeeded only two earlier observers. The station closed soon after his retirement.

In June 1965 he was awarded the M.B.E. in recognition of his long and valuable service.

Miss Dorothy Redfearn, B.E.M.

The death occurred on 18 December 1980 of Miss Dorothy Redfearn, B.E.M.

Miss Redfearn, with her sister Miss Mary Redfearn, had made regular weather observations from the post office at Forest-in-Teesdale, Co. Durham. Reports commenced on 13 December 1937 and terminated with the sisters' retirement on 31 December 1978. The reports were made hourly between 0700 and 1800.

Miss Redfearn and her sister were each awarded the B.E.M. in 1978 in recognition of their unflinching service.

Correction

In the article by Bond, Browning and Collier, *Meteorol Mag*, 110, 1981, 29–40, in the third line of the caption to Fig. 4, '10 m s⁻¹', should be '5 m s⁻¹'.

THE METEOROLOGICAL MAGAZINE

No. 1306

May 1981

Vol. 110

CONTENTS

	<i>Page</i>
Smoothing and filtering of meteorological data. A. C. L. Lee	115
The accuracy of London Weather Centre forecasts of surface wind and total wave heights and their comparison with computer products. R. M. Morris	133
Review	
Weather modification: prospects and problems. G. Breuer (translated from the German by H. Mörrth). <i>R. Brown</i>	143
Mr Frederick James Parsons, M.B.E., M.A.	144
Miss Dorothy Redfearn, B.E.M.	144
Correction	144

NOTICES

It is requested that all books for review and communications for the Editor be addressed to the Director-General, Meteorological Office, London Road, Bracknell, Berkshire RG12 2SZ, and marked 'For Meteorological Magazine'. The responsibility for facts and opinions expressed in the signed articles and letters published in this magazine rests with their respective authors.

Complete volumes of 'Meteorological Magazine' beginning with Volume 54 are now available in microfilm form from University Microfilms International, 18 Bedford Row, London WC1R 4EJ, England.

Full size reprints of out-of-print issues are obtainable from Johnson Reprint Co. Ltd., 24-28 Oval Road, London NW1 7DX, England.

Please write to Kraus Microform, Rte 100, Millwood, NY 10456, USA, for information concerning microfiche issues.

© Crown copyright 1981

Printed in England by Heffers Printers Ltd, Cambridge
and published by
HER MAJESTY'S STATIONERY OFFICE

£1.80 monthly

Dd. 716670 K15 5/81

Annual subscription £23.80 including postage

ISBN 0 11 726282 X

ISSN 0026-1149