Geophys. J. R. astr. Soc. (1975) **43**, 87–101

Complex Demodulation of Geomagnetic Data and the Estimation of Transfer Functions

R. J. Banks

(Received 1975 February 19)

Summary

Complex demodulation is a technique that allows the examination of the variation with time of the amplitude and phase of selected frequency components of a time series. Complex demodulates can form the basis for estimates of the power spectrum of the time series. They are computed most efficiently by the use of the Fast Fourier Transform.

Two examples of the application of complex demodulation to the analysis of geomagnetic time series are given. In the first example, the technique is used to demonstrate the modulation of a periodic phenomenon, the daily variation, by mechanisms with apparent periods of 6 months and 27 days. The second example, discussed in greater detail, is an application to the calculation of Geomagnetic Deep Sounding transfer functions. The spectra of simultaneous records of variations in the three components (H, D and Z) of the magnetic field recorded at a single station are divided into bands, and each band demodulated in turn. The polarization azimuth of the horizontal field at each instant of time can be computed, and, for the example considered, the azimuth tends to be constrained to the north–south direction.

The detailed response of the local conductive structure to different source field polarizations can be demonstrated. Everett & Hyndman's Unit Vector Method is used to investigate the effect of the observed bias in the azimuth of the source field on estimates of GDS transfer functions made by conventional methods. The demodulates can be used to devise criteria for selecting events according to their signal/noise ratio. The selection procedure is most useful when applied to records containing only sporadic activity in the frequency band of interest.

1. Introduction

In 1967, Bingham, Godfrey & Tukey published a paper summarizing the impact that the rediscovery of the Fast Fourier Transform (FFT) had had on time series analysis and spectral estimation. They also discussed likely developments in time series analysis which would be made possible by the advent of the FFT algorithm. Many references to this paper have been made by workers in many different fields, but nearly all of them refer to it as a source of information on the application of the FFT to conventional spectral analysis. In fact, approximately half of the paper

consists of a discussion of the FFT algorithm, and how it can be used to estimate power spectra, either directly or by way of the autocovariance function.

However, a very substantial fraction of the paper is devoted to a quite different though related technique of time series analysis and spectral estimation known as complex demodulation. As far as I have been able to determine, there have been very few references made to this part of the paper, and none that I have been able to find in the field of geophysics. It is the purpose of this paper to draw the attention of geophysicists to the technique, and to give examples that illustrate its use in the analysis of geomagnetic time series. Of the sections which follow, those which define the technique and describe its computational implementation are very much indebted to the paper by Bingham *et al.*

2. The definition and significance of complex demodulates

The process of complex demodulation as applied to a time series $x(t)$ (assumed to be sampled at a series of points spaced at equal intervals of time Δt) is defined in terms of two very simple mathematical operations. Each frequency band of interest in the spectrum is shifted to zero frequency by multiplying each term of the time series by the complex exponential function $\exp(-i\omega' t)$. ω' is the central frequency of the shifted band. A new series $X_s(\omega', t)$ is produced for each frequency band, i.e.

$$X_s(\omega', t) = x(t) \exp(-i\omega' t).$$

The frequency-shifted series is low-pass filtered using a set of weights $a_k (k = -m \text{ to } +m)$, and the result is the (complex) demodulated time series $X_d(\omega', t)$.

$$X_d(\omega', t) = \sum_{k=-m}^{k=+m} a_k X_s(\omega', t + k\Delta t).$$

The demodulates can be expressed most conveniently in the form

$$X_d(\omega', t) = |X_d(\omega', t)| \exp\left(-i\phi(\omega', t)\right)$$

in terms of the modulus and phase of $X_d(\omega', t)$.

2.1 *Relationship to power spectra*

If $S_{xx}(\omega)$ is the power spectrum density of $x(t)$, then by definition $S_{xx}(\omega)\,\Delta\omega$ is the variance in the frequency band $\omega - \Delta\omega/2, \omega + \Delta\omega/2$. The power spectrum is usually estimated as the Fourier Transform of the autocovariance function. However, as Bingham *et al.* show, and as is fairly clear from the definition, the spectral density can also be estimated from a suitable average of the complex demodulates:

$$S_{xx}(\omega)\,\Delta\omega = \langle X_d(\omega, t)\, X_d^*(\omega, t)\rangle$$

where the brackets represent an average over successive points in the demodulated time series, and $X_d(\omega, t)$ is the demodulate of the frequency band centred on ω and lying between $\omega - \Delta\omega/2$ and $\omega + \Delta\omega/2$. In a similar way, the cross-spectrum of two time series $x(t)$ and $y(t)$ can be estimated by

$$S_{xy}(\omega)\,\Delta\omega = \langle X_d(\omega, t)\, Y_d^*(\omega, t)\rangle.$$

Power and cross-spectra calculated by conventional means are averages over the whole length of the data series. Power spectra calculated by complex demodulation can also be averages over the whole series, but alternatively they can be taken over some specially selected subset of the data. The data with which we have to deal in geophysics are usually non-stationary, and therefore properties of the data determined

from averages over the whole of the time series (such as the conventionally calculated power spectrum) will not necessarily be the same as the 'instantaneous' (what Bingham *et al.* call the 'time-local') properties of the data. By taking suitable averages, over subsets of the data chosen on the basis of our special interest in it, we can use the complex demodulates to investigate the non-stationary properties of the time series.

2.2 *Demodulation of data containing periodic components*

Let us suppose that the data contain a periodic component that would produce a peak in the spectrum at frequency ω_0 i.e.

$$x(t) = A\cos(\omega_0 t+\gamma).$$

The simple complex demodulate centred on frequency $\omega' = \omega_0+\delta\omega$ will be

$$\begin{aligned} X_s(\omega', t) &= (A/2)\exp\left(-i(\omega_0+\delta\omega)t\right)\left\{\exp i(\omega_0 t+\gamma)+\exp\left(-i(\omega_0 t+\gamma)\right)\right\} \\ &= (A/2)\left\{\exp\left(-i(\delta\omega.t-\gamma)\right)+\exp\left(-i\left((2\omega_0+\delta\omega)t+\gamma\right)\right)\right\} \end{aligned}$$

which contains frequencies $-\delta\omega$ and $-(2\omega_0+\delta\omega)$.

The frequency-shifted series is low-pass filtered, and hopefully the component at $-(2\omega_0+\delta\omega)$ is removed completely, leaving

$$X_d(\omega_0+\delta\omega, t) = (A/2)\exp\left(-i(\delta\omega.t-\gamma)\right)$$

(I have assumed that the filter response is unity and introduces no phase shift at frequency $-\delta\omega$.)

If we know that the data contain a periodic variation with a frequency of ω_0, it is natural to choose $\omega' = \omega_0$ as the central frequency of the demodulate, and to shift down to zero the band of frequencies immediately around ω_0. In that case, $\delta\omega = 0$, and

$$X_d(\omega_0, t) = (A/2)\exp i\gamma,$$

i.e. the modulus of the demodulate is $A/2$ and its phase is equal to the phase of the periodic variation at ω_0. If A and γ change with time, there will be a corresponding change in the modulus and phase of the demodulate. An example of this application of complex demodulation is given in Section 4.

Another use of complex demodulation is to detect the presence of a hidden periodicity in a data series, and to determine its frequency and phase. If we compute $X_d(\omega', t)$, where $\omega' = \omega_0+\delta\omega$, the phase of the demodulate will be $-\delta\omega.t+\gamma$, which will vary linearly with t. By plotting ϕ as a function of t, we can estimate $\delta\omega$ and hence ω_0.

3. Computational procedures

The obvious technique is to use the defining formulae—to multiply each term of $x(t)$ by $\exp(-i\omega' t)$, and then to low-pass filter the resultant data series. Such a procedure is computationally expensive and leads to a redundancy of information in the demodulated series. Because X_s has been low-pass filtered, adjacent points in the demodulated series are strongly correlated, and the new series need only be sampled at every nth point, where the value of n will depend on the characteristics of the filter.

Much the fastest way of computing the demodulates is by means of the Fast Fourier Transform. This method has the further advantage that the oversampled character of the resultant time series is made quite explicit, and the number of

independent points required to fully specify the demodulated series is quite clear from the frequency response of the filter.

The raw data series is prepared for the FFT in the usual way—reduced to zero mean and zero linear trend, the ends tapered to zero by half cosine bells, and zeros added to bring the number of data points up to 2^K (where K is an integer). We then compute the FFT of the modified time series. The real and imaginary parts of the Fourier Transform are multiplied by a suitable discrete function of frequency centred on ω'. This function defines the filter which smoothes the series of raw demodulates X_s. The resultant band of frequencies is shifted to zero, care being taken to arrange the resulting negative frequency components in the correct storage locations. The new set of Fourier Transforms created by this sequence of operations can be truncated so as to produce a demodulated series consisting of independent data points. The new Nyquist Frequency (ω_N') is chosen so that there is no overlap of the positive and negative frequency components (aliasing), and so as to reduce zeros to as few as are compatible with computational efficiency (see Fig. 1(d)). Finally, an inverse FFT operation converts the transforms back into the time domain, and generates the series of complex demodulates centred on frequency ω'. The sampling interval of the demodulated series is $\Delta t'$, given by $\Delta t' = 1/2\omega_N'$.

The complete procedure is illustrated diagrammatically in Fig. 1. In the example shown, the spectrum of the data series has a peak at $\omega = \omega_0$, and the demodulated band is naturally centred on $\omega' = \omega_0$. Exactly the same procedure would apply to any other part of the spectrum not containing a periodic component.

Problems in the choice of filter

Ideally, we would wish to use a filter of the form

$$W(\omega) = 1 \qquad \omega' - \Delta\omega/2 \leqslant \omega \leqslant \omega' + \Delta\omega/2$$
$$= 0 \qquad \text{elsewhere}$$

which simply passes a band of frequencies of width $\Delta\omega$ centred on $\omega = \omega'$. However, multiplication in the frequency domain is equivalent to convolution in the time domain, and we need to bear in mind the form of the filter weights in the time domain corresponding to the frequency characteristics specified above. In addition, we have to remember that because of the cyclic nature of the Fourier Transform, the demodulated series and the set of weights repeat themselves at interval $T = N\Delta t'$ (N is the number of points in the demodulated series).

The set of weights in the time domain corresponding to the simple rectangular bandpass filter will vary with t as $\sin(\Delta\omega.t/2)/(\Delta\omega.t/2)$ which has substantial sidelobes that die away rather slowly. The result of the convolution of such a set of cyclic weights with the cyclic data series is leakage from one data interval T to the next. One way of getting round the difficulty is to modify $W(\omega)$ so as to make the corresponding set of weights die away more quickly. A possible choice of bandpass filter would be a cosine bell:

$$W(\omega) = \frac{1}{2}\left(1+\cos\frac{2\pi(\omega-\omega')}{\Delta\omega}\right) \qquad \omega' - \Delta\omega/2 \leqslant \omega \leqslant \omega' + \Delta\omega/2$$
$$= 0 \qquad \text{elsewhere.}$$

If such a function is felt to have insufficient definition of the required frequency band, a compromise would be a ' boxcar ' filter tapered at the low and high frequency cut-offs by half cosine bells. A second method of minimizing leakage is to add zeros at each end of the data series, thus spacing it further from its cyclic ' image '. Such a procedure is, in any case, normally followed for reasons of computational efficiency.

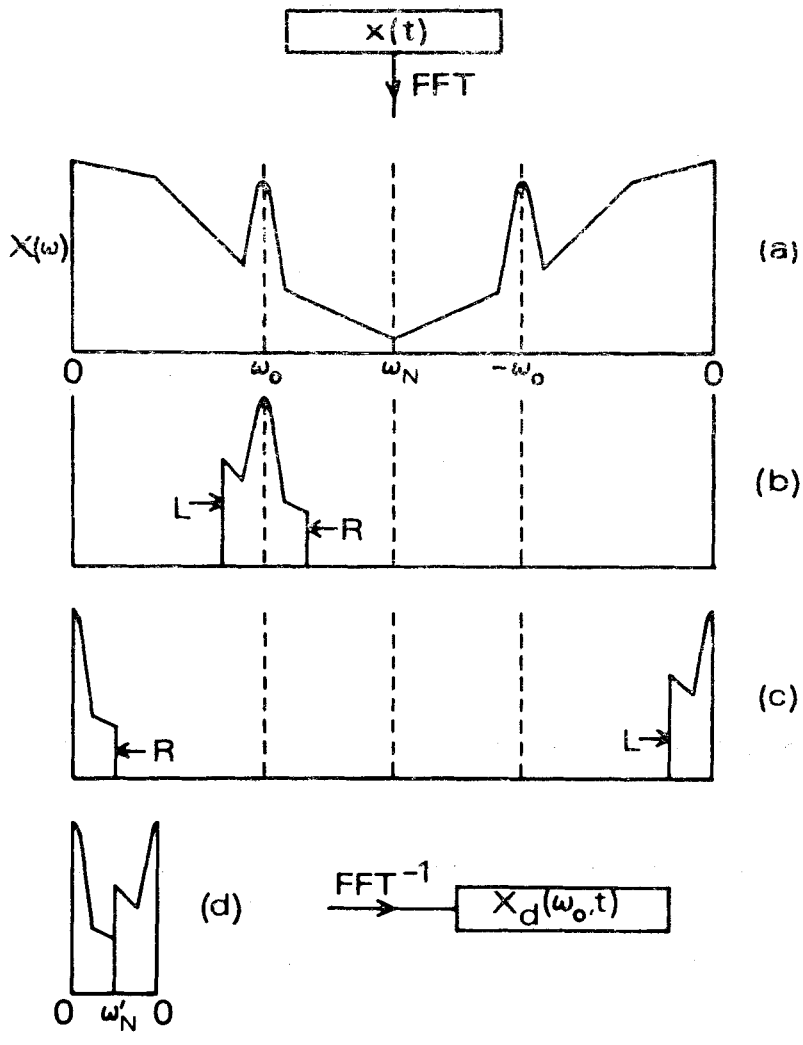

FIG. 1. Schematic diagram of the procedure for computing complex demodulates $X_d(\omega_0, t)$ from the time series $x(t)$. (a) Raw Fourier Spectrum; (b) Filtered Fourier Spectrum; (c) Frequency shifted Spectrum; (d) Truncated frequency shifted Spectrum. Note that no attempt has been made to indicate a realistic shape for the filter.

4. Complex demodulation of the daily variation

In order to illustrate the use of complex demodulation as a means of studying a periodic phenomenon, I have computed the demodulates of a series of hourly mean values of the horizontal north (H) component of the geomagnetic field recorded at Hartland. The original data series was 2208 points long, from October 1 to December 31, 1964. Its length has been increased to 3072 points by adding zeros at each end (after removal of the mean, etc.). The central frequency of the demodulated band is 1 c day^{-1}, at which there is known to be a peak in the spectrum produced by the daily (S) variation, and the band stretches from 0.8 to 1.2 c day^{-1}.

In Fig. 2 the modulus and phase of the demodulated time series are plotted, which

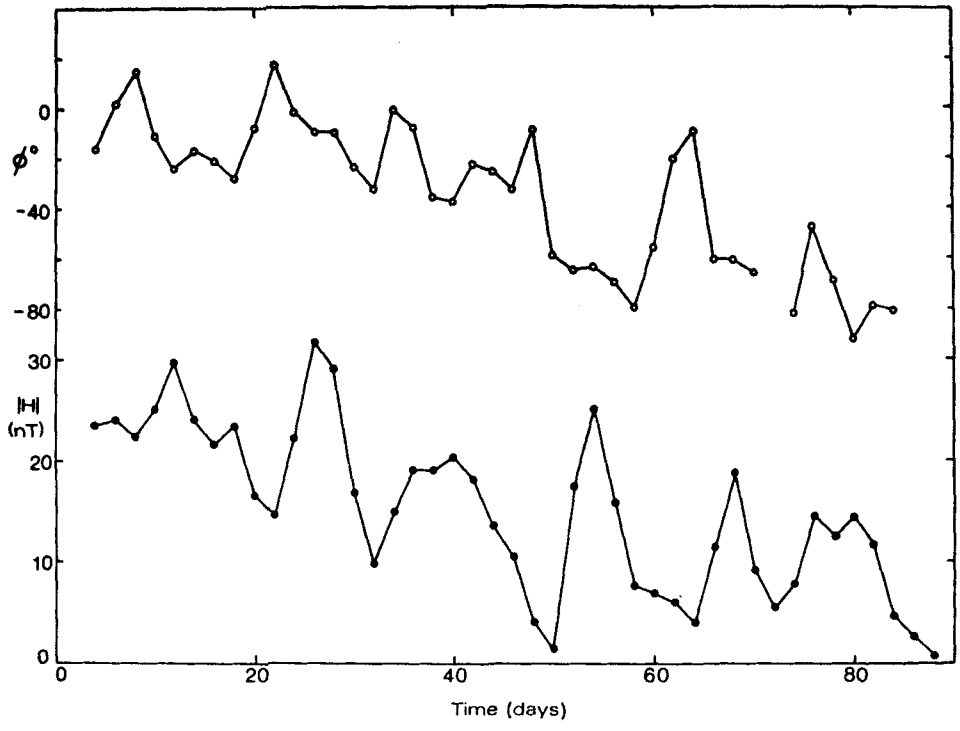

FIG. 2. Modulus and phase of Hartland H component daily variation demodulates. Time in days relative to October 1, 1964.

correspond to ' time-local ' estimates of the amplitude and phase of the daily variation. The sampling interval of the demodulated series is 2 days. The modulus shows very clearly the effects of an annual or semi-annual modulation, with the maximum near the beginning of the record and the minimum at the end. In addition, it shows what appears to be a rather irregular modulation with a period of about 13.5 days, presumably associated with the 27-day recurrence tendency of magnetic storms. The phase also shows interesting features, including a long-term trend and consistently high values on days when the amplitude of S is small.

It is not my purpose to discuss here the detailed temporal behaviour of the daily, variation, or the mechanisms responsible for the modulation, but to point out how the demodulates can help us to study that behaviour. For example, one line of enquiry made easy by the demodulates would be to look at the relationship between the amplitude of the 13.5-day modulation and the incidence of magnetic storms.

5. The application to geomagnetic deep sounding (GDS) data

5.1 *Estimation of GDS transfer functions*

The raw data derived from geomagnetic deep sounding (GDS) experiments consists of simultaneous magnetograms of the three components (probably H, D and Z, the horizontal north, declination, and vertical components respectively) of the magnetic field recorded at one or more locations. The first step in the analysis is the transformation of these time series into the frequency domain. Two different procedures have been applied to two different classes of magnetic disturbance:

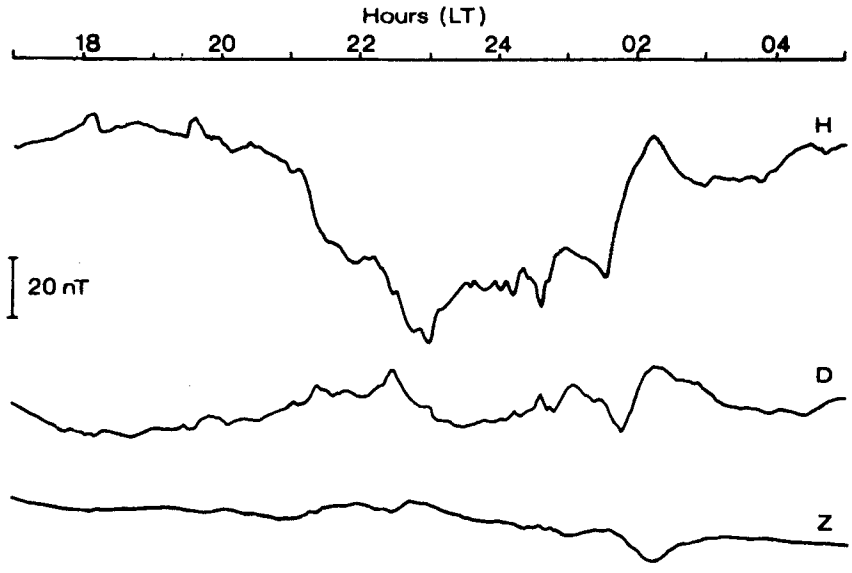

FIG. 3. Kabianga magnetogram for April 21–22, 1971.

(1) Bays and other quasi-cosinusoidal transient variations are picked out from the magnetograms, and the disturbed portion of the record Fourier Transformed directly. There are two advantages to this approach: a good signal/noise ratio is guaranteed at frequencies around $1/\Delta T$, where ΔT is the length of the oscillatory disturbance; and the polarization of the horizontal field, and probably also of the external current source, is likely to be relatively constant throughout the event, and therefore easy to define. The disadvantage of the method is that a large number of events is needed to sample the response of the conductive structure to all polarizations and frequencies of the external field.

(2) Magnetic storms and other continuously disturbed periods, containing superimposed variations generated by independent current systems, and having a wide range of source field polarizations and frequencies, are analysed by computing the power spectrum of the entire record. The power and cross-spectra may be calculated either from smoothed Fourier Transforms, or from modified autocovariance and cross-covariance functions (computed either naively or by an inverse Fast Fourier Transform of the raw power spectrum or cross-spectrum). The advantage of this approach is that the spectra are averages over a wide range of disturbances, and substantial power is present at a wide range of frequencies. There are a number of disadvantages: the form of the averaging in a given spectrum is unknown; it is not certain that the averages give equal weight to all polarizations of the external field, and it is likely that the field will tend to be constrained to some range of azimuths. Some frequencies may only be present in the magnetic record for a short length of time. Their contribution to the total power spectrum is diluted by the uncorrelated digital noise which is present throughout the record, and whose relative contribution to the spectrum increases as the length of the record increases.

On the whole, it is clear that method (1) is to be preferred. Unfortunately, very often it is just not possible to obtain a sufficient number of magnetograms of bays or similar disturbances that covers the whole frequency spectrum and samples a reasonable range of source polarizations. Stretches of continuous disturbance of type (2) tend to be much more common.

Complex demodulation provides a means of analysing type (2) data so as to obtain type (1) results. If we make use of complex demodulation, we have to surrender some of the frequency resolution we would obtain using conventional spectral analysis. In practice, this is not a serious limitation because the response of the Earth to magnetic field variations must vary smoothly and relatively slowly with frequency, and response estimates centred on half-a-dozen representative frequencies (based on a division of the spectrum into a corresponding number of frequency bands) will probably provide us with more information than we can comfortably handle. In return for this loss of resolution, we are put in a position to examine the way in which the amplitude, phase, and polarization of the disturbance in each frequency band change with time.

5.2 *Complex demodulation of Kenya GDS data*

Fig. 3 displays the three components of a magnetic disturbance recorded at Kabianga in western Kenya. The low amplitudes of all frequency components in the Z magnetogram are very characteristic of magnetic disturbances at low magnetic latitudes, as also is the strong tendency for the horizontal field variations to be polarized in a magnetic north–south direction. Kabianga is approximately 10° south of the magnetic equator, and the magnetic disturbances recorded there are predominantly generated by the return currents of high latitude events, which tend to flow east–west, parallel to the magnetic equator.

The sampling interval of the digitized magnetogram is 1 min, and there are 720 points in the raw record. It was high-pass filtered, reduced to zero mean, and tapered. High-pass filtering was found to produce less distortion of the lowest frequencies than removal of a linear trend. The use of a linear filtering operation preserves linear relationships between components over the whole frequency spectrum, though it is more time consuming than detrending. Zeros were added to bring the length of the series to 1024 points, giving a fundamental spectral resolution of $\delta f = 1/1024$ c min^{-1}. For the purposes of complex demodulation, the spectrum has been divided into six bands centred on frequencies (in units of δf) of 5, 12, 24, 48, 96, and 192. The bands contain Fourier components with periods of 512–128, 128–64, 64–32, 32–16, 16–8, and 8–4 min respectively. Fig. 4(a) displays the modulus of the demodulate of the Z component for band 4 (periods of 32–16 min). There is a continuous but very low level of activity in this band throughout most of the record, except for the event which occurred at 2240 h.

Previous studies of these data using conventional techniques of spectral analysis (Banks & Ottey 1974) have shown that, at Kabianga, Z variations in the period range 5–100 min are strongly correlated with horizontal field variations having azimuths of $-120°$ or $+60°$. It is now possible, using the demodulates, to see whether this correlation stands up in detail. For each instant of time, the demodulates of the H and D fields can be used to calculate the polarisation characteristics of the horizontal field: the azimuth ψ of the major axis of the polarization ellipse, the magnitude H_ψ of the major axis of the ellipse, and the ratio of the lengths of minor and major axes (Lilley & Bennett 1972). When investigating the dependence of anomaly response on the azimuth ψ, allowance has to be made for changes in the amplitude of the horizontal field variations, so rather than simply comparing Z with ψ, we compare Z/H_ψ and ψ.

Fig. 4(b) shows the azimuth ψ of the major axis of the polarization ellipse plotted as a function of time. It shows quite clearly the confinement of the horizontal field to azimuths between 0° and $-40°$ (no distinction is made in the diagram between azimuths of ψ and $\psi + 180°$). The conventional analysis indicates that such a polarisation is very unfavourable for the excitation of the conductivity anomaly, which is confirmed by the low level of the Z/H_ψ ratio ($<0{\cdot}1$) for such polarizations (Fig. 4(c)). However, at about 2030 and 2230, the polarization azimuth swings to values of 60°E

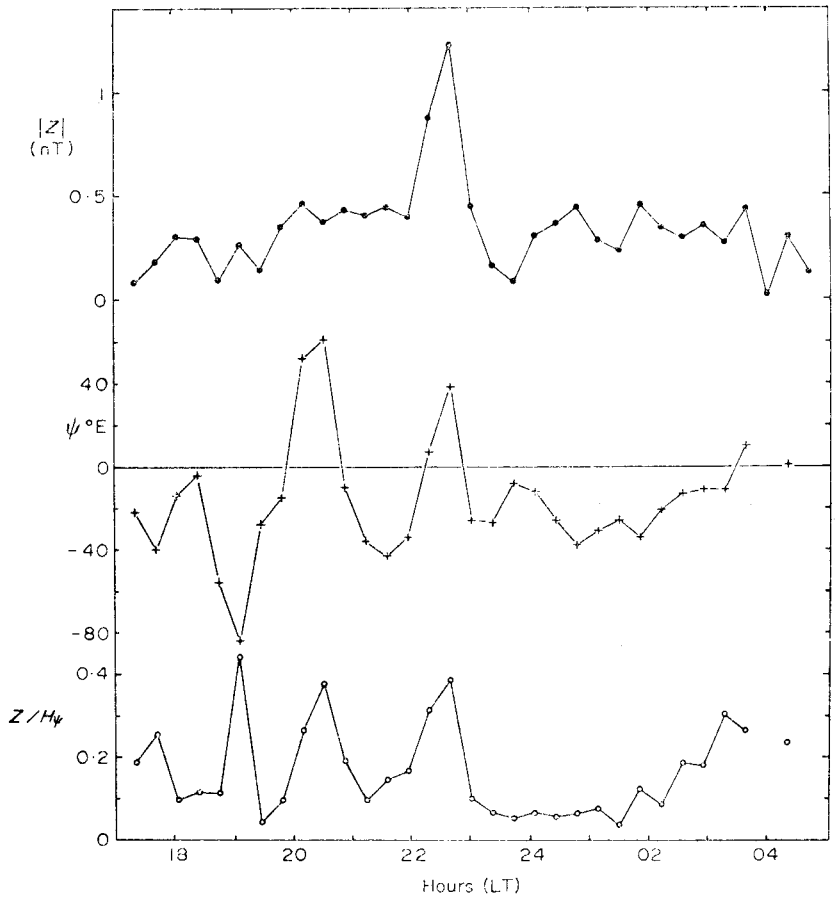

FIG. 4. Demodulates of the Kabianga magnetogram: (a) Modulus of Z component demodulate; (b) Azimuth ψ of major axis of horizontal field polarization ellipse; (c) Z/H_ψ. Periods 32–16 min.

and 40° E respectively. The favourable azimuth leads to the excitation of the Z anomaly, producing peaks of about 0·4 in the Z/H_ψ ratio. A similar peak is produced at 1900 h by an azimuth of $-85°$.

A diagram like Fig. 4 enables us to see in detail how the response of the anomaly changes with changing source azimuth. If we are to be able to investigate the functional form of the relationship, some means of plotting Z/H_ψ against ψ is required. For instance, the demodulates and polarization estimates could form the basis of a Parkinson plot. However, there are difficulties in adequately representing by graphical means the azimuthal dependence of the anomaly response. Conventional Parkinson plots take no account of the phase differences between variations in different magnetic field components. Similarly, a straightforward plot of Z/H_ψ as a function of ψ does not take account of the elliptical rather than linear polarization of most of the horizontal field disturbances. If an event is elliptically polarized, then although the azimuth of the major axis may be such as not to produce any Z anomaly, the azimuth of the minor axis will be most favourable, and the anomaly will respond accordingly, producing a Z variation whose amplitude will depend on the value of the ratio R_ψ of minor/major axis lengths.

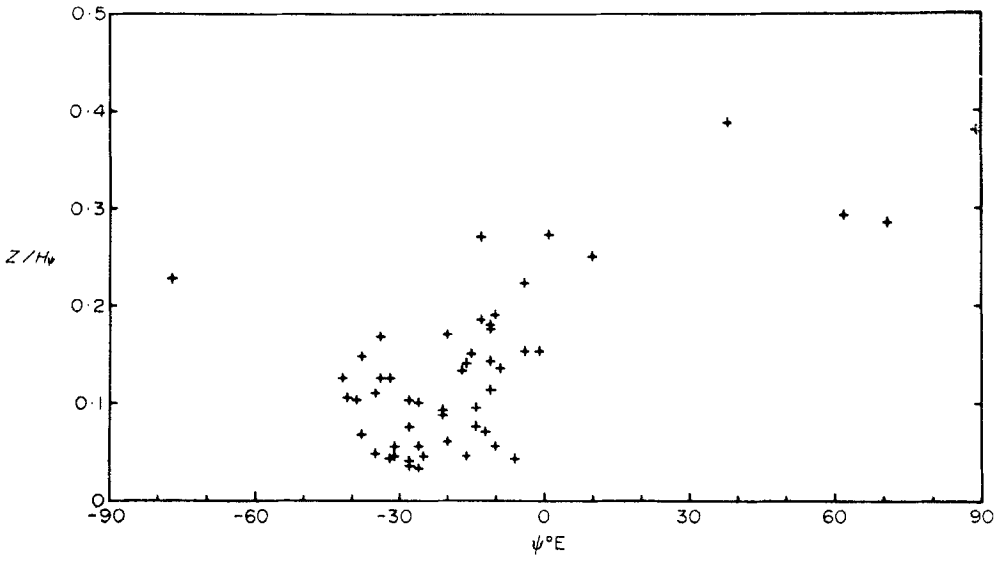

FIG. 5. Plot of Z/H_ψ against ψ for all Kabianga records, using events for which minor axis/major axis ratio was less than 0·2. Periods 32–16 min.

Problems of representation can be overcome if sufficient data are available. However, for present purposes, the problem can be partly overcome by plotting only those points for which the value of R_ψ is small, say, less than 0·2. Fig. 5 is an example of such a diagram; the results from the analyses of two other similar stretches of data have been included. Even so, the sampling of the response at a wide range of azimuths is scarcely adequate. The vast majority of the data points have azimuths which lie in the range $-40°$ to $0°$. The distribution of azimuths for the whole data set can be demonstrated more clearly by plotting a histogram showing the numbers of points in ten-degree sectors as a function of azimuth (Fig. 6).

In essence, the GDS single station transfer function technique involves fitting a relationship of the form

$$Z = GH_\psi \cos(\psi - \theta_p) \tag{1}$$

to the plot of Z/H_ψ as a function of ψ (θ_p is the azimuth of the horizontal field with which Z shows maximum correlation, G is the transfer function, and the disturbance is assumed to be linearly polarized at azimuth ψ). The plot of Z/H_ψ should show a cosinusoidal dependence on ψ, with $Z/H_\psi = 0$ at $\psi = \theta_p + 90°$. Fig. 5 does show the expected features, if a little imagination is used to interpolate the data. $Z/H_\psi = 0$ at $\psi = -25°$, and a roughly cosinusoidal ψ dependence can be detected.

However, the diagram does raise the question of whether the procedure used to fit the data to equation (1), or to a more sophisticated form such as $Z = A.H + B.D$, can cope with very unequal distributions of polarization. It is important to know whether the source bias leads to a corresponding bias in the transfer function estimates, and the effect on the accuracy of the estimates. This problem can be investigated by using Everett & Hyndman's (1967) 'Unit Vector' approach to the estimation of GDS transfer functions.

5.3 *The unit vector method*

In the unit vector method, as described by Everett & Hyndman, individual events are linearly combined in such a way as to form two horizontal vectors of unit amplitude,

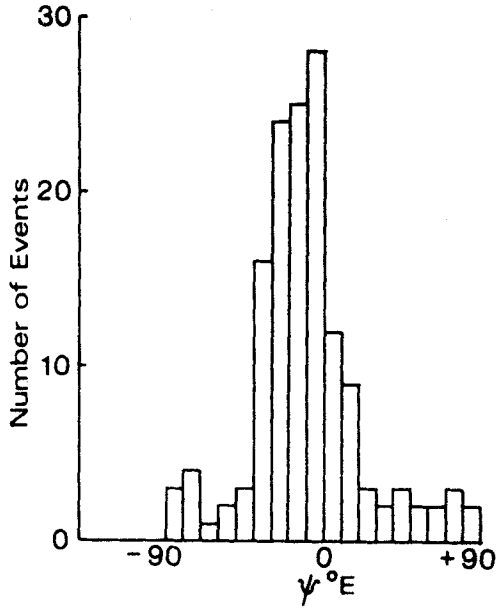

FIG. 6. Histogram showing distribution with azimuth of H field polarizations recorded at Kabianga. Periods 32–16 min.

polarised north and east, respectively. The vertical fields that correspond to the linear combinations are estimates, A and B, of the ideal transfer functions, A' and B', and they are unbiased estimates. Everett & Hyndman write

$$Z_r = AH_r + BD_r + e_r = A'H_r + B'D_r + e_r' \tag{2}$$

where the residual e_r is an unbiased estimate of e_r'. In Everett & Hyndman's notation, the subscript r refers to the result of the Fourier Transformation of the rth record, usually of a single event (method 1 of Section 5.1). However, it can equally well refer to a single complex demodulate, estimated at or around one particular instant of time. In what follows, an event can either be a selected record from which conventional Fourier Transforms are computed, or a single demodulate.

Complex weighting factors are found such that

$$\sum_{r=1}^{n} \gamma_r(H_r, D_r, Z_r) = (1, 0, A), \tag{3}$$

i.e. a linear combination of the events is formed to produce an event of unit magnitude in the horizontal north direction. The corresponding linear combination of vertical component events must then be an estimate of the transfer function A. The expectation of the square of the modulus of the error in A is:

$$E(\delta A . \delta A^*) = \sum_{r=1}^{n} \gamma_r e_r' . \gamma_r^* e_r'^* \tag{4}$$

The likely error $|\delta A|$ is made a minimum by minimizing $E(\delta A . \delta A^*)$ with respect to γ_r, subject to the constraints

$$\sum_{r=1}^{n} \gamma_r H_r = 1 \quad \text{and} \quad \sum_{r=1}^{n} \gamma_r D_r = 0$$

using the method of Lagrange multipliers.

Everett & Hyndman weight their events before the analysis by a factor $\sqrt{T_r}$, T_r being the length of the rth record. All the events should then have the same expected noise level, and equation (4) can be simplified:

$$E(\delta A.\delta A^*) = e^2 \sum_{r=1}^{n} \gamma_r \gamma_r^* \tag{5}$$

where $e^2 = E(e_r'.e_r'^*)$. The set of weights is found to be

$$\gamma_r = \frac{S_{DD} H_r^* - S_{HD}^* D_r^*}{S_{HH} S_{DD} - S_{HD} S_{HD}^*} \tag{6}$$

S_{DD} is the power spectrum of D, S_{HD} is the cross-spectrum of H and D, etc., and

$$A = \sum_{r=1}^{n} \gamma_r Z_r = \frac{S_{DD} S_{ZH} - S_{HD}^* S_{ZD}}{S_{HH} S_{DD} - S_{HD} S_{HD}^*}. \tag{7}$$

Equation (7) is the familiar expression usually obtained by a simple minimization of $E(e_r' e_r'^*)$. An equivalent set of weights δ_r, which determine B, can be found by minimizing $E(\delta B.\delta B^*)$, subject to the constraints

$$\sum_{r=1}^{n} \delta_r D_r = 1 \quad \text{and} \quad \sum_{r=1}^{n} \delta_r H_r = 0.$$

The equivalence of the results obtained by the straightforward least squares approach and by the unit vector method, shows that estimates of the transfer functions using formulae such as equation (7) do involve a quite subtle weighting of individual events. The nature of the relative weighting can be seen more clearly if we restrict our attention to a group of events for which the cross spectrum of H and D, S_{HD}, is zero. For such a set of data, the weight to be applied to the rth event would be

$$\gamma_r = \frac{H_r^*}{S_{HH}} = H_r^* \sum_{r=1}^{n} H_r H_r^*.$$

In other words, the contribution which the rth event makes to the estimate of A is directly proportional to the H amplitude of the event. If there are only a few events in the group of digitized records that have significant variations in the H component, they will be the ones to effectively determine the estimate of A. The remaining events, though they may form the majority, will make a much smaller contribution, determined by their weighting. We can see that the effect of the weighting procedure is to prevent the source bias from biasing the transfer function estimates.

However, we can expect source bias to reduce the accuracy of one or other of the transfer function estimates. The minimum error in A can be estimated from equation (5) by substituting for γ_r from equation (6). The error is then (approximately) given by

$$|\delta A|^2 = \frac{e^2}{S_{HH}(1 - R_{DH}^2)} \tag{8}$$

where R_{DH}^2 is the coherence of D and H. Assuming again a group of events for which S_{DH} is zero (i.e. $R_{DH}^2 = 0$)

$$|\delta A|^2 = \frac{e^2}{S_{HH}}$$

and the error in A is inversely proportional to the square root of the total power of the H component records. If a group of events happens to have a predominantly east–west polarization, S_{HH} will be small, and the resulting error in A will be large. By analogous reasoning, the corresponding error in B will be smaller.

5.4 *Weighting and selection of events*

The investigation of the East Africa data in Section 5.2, using complex demodulation, demonstrated the unequal distribution of polarizations of the events. The unit vector approach to transfer function estimation shows that the standard least squares method weights events in such a way as to minimize the effects of source bias, and explains the effectiveness of the method in coping with inadequate data. However, the unit vector method also enables the estimation of the errors in A and B, and these should be computed because they will show the effects of source bias.

It might be thought that the weights γ_r, δ_r represent the optimum set, and that for a given set of data, no further weighting or selection is required. Further weighting or selection must reduce S_{HH}, and according to equation (8) increase the likely errors in A and B. However, a form of weighting has already been applied to the data, since the records being analysed have themselves been selected from the totality of available magnetic data. Records used for transfer function analysis are selected by visual inspection, involving some assessment of the degree of disturbance and of the range of frequencies present. We would not consider analysing a record that contained nothing in H or D but digital noise; we would intuitively expect to obtain superior results from the analysis of records with large amounts of power in each frequency band.

The discrepancy between equation (8) and what we intuitively expect is explained by the assumption made in the unit vector analysis (on which equation (8) is based) that H and D are noise-free. This assumption is implicit in equation (2). If the assumption were strictly correct, then every event, however small the values of H_r and D_r, would contribute to the estimates of A and B, and would help to reduce the errors δA and δB. Digital noise in H and D makes no contribution to the estimation of A and B, but does lead to an apparent (though spurious) reduction in the errors.

The intuitive approach is consequently a reasonable one, and we may expect to achieve more reliable transfer function estimates if we apply a process of selection or weighting to the demodulates, based on the amplitude of the disturbance, as measured by H_ψ for instance. However, if we were to weight the demodulates in proportion to $(H_\psi)_r$, one of the requirements of the unit vector analysis would no longer be met—that the events, or demodulates, should have the same expected noise level. Instead, the noise level would be proportional to $(H_\psi)_r$, and we would have to minimize the full expression in equation (4), rather than the simpler form in equation (5).

In order to avoid this complication, we can apply a simple selection procedure to the data, by assigning weight 1 to all events for which $(H_\psi)_r$ exceeds a specified level, and weight 0 to all events for which the reverse is true. The determination of the critical value of H_ψ can be based on a visual inspection of the complete set of demodulates from all available records. It should be possible to estimate the likely maximum noise level, and to use this as a basis for the choice of the critical value of H_ψ. More complicated automatic schemes could be devised if required.

Selection criteria of this kind have been applied to the East African data, and the estimates of A, B, δA, δB, and other transfer functions compared with those obtained from the complete data set. Fig. 7 is a comparison of the coherence R^2 of the observed vertical component with that predicted by $Z = AH + BD$. R^2 is equal to the fraction of the total vertical field power S_{ZZ} that can be fitted to the linear combination of horizontal field components. It is clear that the effect of selection is to increase the proportion of coherent Z power, and that the criterion used is discriminating against

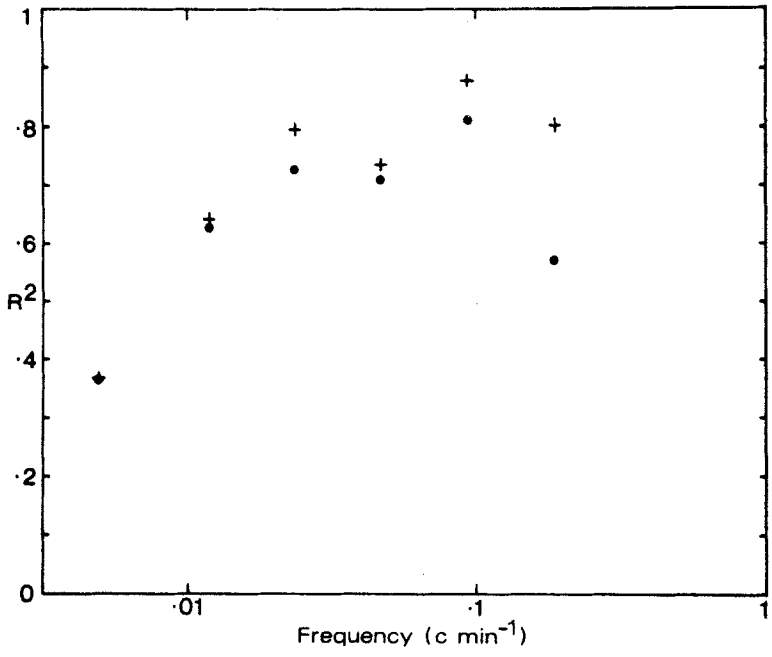

FIG. 7. Coherent fraction of vertical field power in the six frequency bands, at Kabianga: all data (circles), and selected data (crosses).

uncorrelated noise. The difference is most marked at the higher frequencies, where the discrimination is most severe—only only 28 per cent of the data were passed in band 5, and 15 per cent in band 6, compared with percentages of 83, 67, 45, and 43 in bands 1 to 4 respectively. The transfer function estimates themselves (not shown) were very little affected by the discrimination; the estimated errors in the transfer functions increased slightly, presumably because of the reduction in the total amount of data. The slightness of the effects produced might be seen as an indication that there is no need under any circumstances for any refinement of the usual methods of analysis. However, my own view is that the insensitivity of the transfer functions to the removal of substantial portions of the data is rather an indication that the original visual selection procedure was carried out efficiently and correctly. I can readily envisage other types of data to which some such selection procedure could be applied with great benefit. For instance, in my (1969) analysis of long period magnetic storm data, I selected records from periods of sunspot maximum in order to increase the real power levels of Z variation data. Using the demodulates, a much more refined selection procedure could be applied to the data to enhance signal/noise ratios. The greatest benefits will accrue when we are dealing with events for which the level of the horizontal field disturbance is generally low, and only infrequently rises substantially above the noise level.

6. Conclusions

The experience that has been gained so far in the calculation of GDS transfer functions by way of complex demodulation indicates that the approach carries with it substantial advantages. The spectrum can be divided up into bands whose widths can be related to the power levels. The straightforward Fourier spectrum is really tailored to processes generating essentially white spectra. Where power levels vary

logarithmically with period (as is often the case for geomagnetic data) a division of the spectrum into logarithmic frequency ranges seems appropriate.

The computed transfer functions are automatically averaged over the chosen frequency bands, thereby reducing the independent frequency estimates to a manageable number, and producing much more stable estimates in the process. The estimation of errors for the transfer functions is a simple extension of the basic calculations if Everett & Hyndman's unit vector method is used. However, similar results could be achieved by averaging the raw spectral estimates in the frequency domain. The major bonus that complex demodulation has to offer is information about temporal changes in the characteristics of the source field, and the possibility of using the information as a basis for selecting those parts of the data that are likely to give the most satisfactory transfer function estimates.

Results similar to those achieved by complex demodulation could, up to a point, be obtained by dividing the time series into a number of segments, and estimating the spectrum of each portion in turn by the usual methods. However, complex demodulation, besides being very much more efficient in computational terms, allows a very much greater degree of flexibility. The information can be divided in any way we choose between the time and frequency domains; between the two extremes represented by the original time series of N data points and the raw Fourier spectrum with N Fourier coefficients. Complex demodulation seems to be an excellent way of having the cake and eating it. Geophysicists should look hard at complex demodulation as a method of spectral estimation, and as a way of simultaneously examining the non-stationary properties of time series.

Department of Environmental Sciences,
University of Lancaster

References

Banks, R. J., 1969. *Geophys. J. R. astr. Soc.*, **17**, 457.
Banks, R. J. & Ottey, P., 1974. *Geophys. J. R. astr. Soc.*, **36**, 321
Bingham, C., Godfrey, M. D. & Tukey, J. W., 1967. *IEEE Trans. Audio Electroacoustics*, AU-15, 56.
Everett, J. E. & Hyndman, R. D., 1967. *Phys. Earth Planet. Int.*, **1**, 24.
Lilley, F. E. M. & Bennett, D. J., 1972. *Geophys. J. R. astr. Soc.*, **29**, 49.