Godfrey

COMMONWEALTH SCIENTIFIC AND INDUSTRIAL RESEARCH ORGANIZATION

DIVISION of FISHERIES and OCEANOGRAPHY

Report No. 80

THE SPECTRUM OF SEA LEVEL AT SYDNEY

By B. V. Hamon

Reprint No. 959

Marine Laboratory
Cronulla, Sydney
1977

ISBN 0 643 02048 9

Printed by CSIRO, Melbourne

77.85-600

THE SPECTRUM OF SEA LEVEL AT SYDNEY

B. V. Hamon

Abstract

The spectrum of sea level at Fort Denison (Sydney Harbour, 33°53'S 151°10'E) has been computed in the range 0-4 cycles per day (cpd) from three years of filtered 3-hourly tide heights. Separate computations for 0-0.8 cpd, from 19 years of 12-hourly mean sea levels, are also reported, together with the spectrum of mean sea level adjusted to constant atmospheric pressure. The main results are the low values of spectral energy density in the intertidal and transtidal bands, and a slight but significant peak in the spectrum at 0.35-0.4 cpd.

INTRODUCTION

Tidal records are among the few relatively long continuous geophysical records available for study. They also contain much information at frequencies outside the tidal bands.

Tidal recording at Fort Denison goes back to at least 1897, but no routine extraction of hourly heights has been made. Following a suggestion from W.H. Munk some years ago, this Division, with the cooperation of the Maritime Services Board of N.S.W., began digitising the accumulated tide records from Fort Denison. We are working backwards from 1975, and at present (1976) the work has been completed back to 1950. It is planned to continue the work as time permits. We are also digitizing the more recent records, with about one year delay.

This paper is an interim report on the spectra that have been computed, using up to 19 years of data. Its main purpose is to present the spectra and discuss methods used. Little attempt is made at interpretation. A much more extensive study, including other variables, would be needed for this.

DATA AND METHODS

The basic data is the set of hourly tide heights for Fort Denison (Sydney Harbour) for approximately 19 years (August 1954-December 1973). These hourly heights had been read in this Division from the original chart records, which were made available to us by the Maritime Services Board of N.S.W. and the Archives Office, N.S.W. The charts were read on a (manually operated) chart reader, either "D-Mac" (D-Mac Ltd., Glasgow) or Summagraphics (Summagraphics Corp., Fairfield, Conn., U.S.A.). (Further details of the use of a chart reader for tide records is given by Greig (In press).)

The chart reader gives a digital x,y coordinate pair, on punched paper tape or directly to a computer, for any input data point. The sets of co-ordinates are computer processed, using the appropriate chart scales and time and height checks, to give the required tide heights at hourly intervals.

These hourly tide heights are then checked (Lennon (1965), Greig (In press)) for smoothness. Suspected values are checked by hand against the original records. After correction, and a repeat check, the hourly tide heights are stored on magnetic tape. The stored values are in units of feet, rounded to 0.01 ft (≈0.3 cm), and are relative to the zero of the tide gauge, which is approximately Indian Spring Low Water.

The spectra were computed by the FOS (Faded Overlapping Segments) method (Groves and Hannan, 1968), using a set of FORTRAN subroutines (ARAND) made available to us by the University of Oregon, U.S.A.

In the FOS method, the data (length N=number of observations) are divided into overlapping segments, each of length 2m where usually $2m \ll N$, and m+1 is the number of frequency bands desired. Each segment is "tapered" of "faded" by term-by-term multiplication by a triangular or bell-shaped "data window". The squared moduli of the Fourier components of each tapered segment are then computed, and the values in each frequency band averaged over all available segments to give the spectrum. (The method can also be used for cross spectra, coherence and phase, but only the spectra of single series are of concern here).

Two data windows were used, although there is little reason to choose between them for the present applications. The two windows are (ARAND subroutine WINDOW 1):

$$w_1(j) = 1-\{2(j-(2m-1)/2)/(2m+1)\}^2$$

$$w_2(j) = 1-|2(j-(2m-1)/2)/(2m+1)|$$

where j = 0,1,2...2m-1.

These windows result in half-power spectral bandwidths (in units of cycle/time) of $b_1 \approx 1.2/(2m+1)$ and $b_2 \approx 1.33/(2m+1)$, respectively. (Normalisation of the data windows is provided for in the subroutines that compute the spectrum).

Ninety-five percent confidence intervals were estimated from Fig. 3.10 of Jenkins and Watts (1968), using 1.25N/m as an estimate of degrees of freedom (ν), for either data window. (For $\nu>100$, the confidence limits were estimated from an approximation to the χ^2 distribution, which led to $1\pm2.765/\nu^{\frac{1}{2}}$ for the 95 percent confidence band).

THE SYDNEY SPECTRUM, 0-4 cycles/day

This spectrum was computed mainly to look for evidence of spectral peaks between and above the main tidal frequency bands (~1 and ~2 cycles/day (cpd), for diurnal and semi-diurnal tides). The computation was also regarded as a useful test of the spectral techniques used.

Munk and Bullard (1963) were the first to compute spectral energy densities close to and between the two main tidal bands. They used specially designed band pass filters to discriminate against the unwanted tidal energy. More recently, Munk and Cartwright (1966), (referred to herein as MC) and Cartwright (1968) have successfully computed spectra in the frequency range 0-4 cpd without first filtering out the tides. We have used essentially the same methods as given in Cartwright (1968).

The basic data set, with time interval Δt of 1h, should in theory permit the spectrum to be computed up to a Nyquist frequency ($f_N=1/2\Delta t$) of 12 cpd. The upper half to one-third of such a spectrum might be contaminated by aliasing, and the computation would be lengthy if good frequency resolution is to be retained, and a long enough record used so that confidence limits are not too large. We have not attempted such a computation, although it would not be impractically large for the CDC7600. (The computation reported below took only 15 sec, including the initial low-pass filtering).

Following Cartwright (1968), we smoothed the hourly tide heights with a 61-term low pass filter with half-amplitude cut-off at 4.0 cpd, computing only every third of the possible output values, so that the series used for spectral analysis had Δt=3 h.

Only three of the 19 years of available data were used, since only an experimental "first look" was required. The years 1965-1967 were used. The number of (3-hourly) observations was N=8717.

When a series contains a number of strong spectral "lines" of known frequencies, as in the present case (the major diurnal and semi-diurnal tidal constituents), it is an advantage to choose the segment length for spectral analysis (2m) so that each segment contains nearly an integer number of wave lengths of the strong lines. This reduces the leakage of spectral energy from strong to weak parts of the spectrum. In the present case we used 59 day segments, or 2m=472. The frequency resolution is then Δf=4/236=.016949 cpd, and each segment contains 114.005 cycles of the M_2 component (amplitude at Sydney ~51 cm) and 118.001 cycles of the S_2 component (amplitude (12.8 cm)). Note however that this segment length is not as suitable for the other three of the five components of maximum amplitude at Sydney. These are N_2 (11.4 cm, 111.863 cycles/segment), K_1(14.8 cm, 59.162 cycles/segment) and O_1 (9.5 cm, 54.843 cycles/segment). No attempt was made to find a segment length more suited to Sydney's tidal constituents.

The spectrum is shown in Fig. 1. The most obvious features are the bands containing the diurnal and semi-diurnal tides. Significant energy also appears in the two higher-frequency tidal bands, near 3 and 4 cpd. There is the usual steep rise towards zero frequency. The annual and semi-annual "tides" are not resolved, and the monthly and fortnightly tides do not stand out against background. The broad peak 0.3-0.6 cpd will be discussed later.

The contribution to the spectrum from rounding-off the original hourly heights to the nearest 0.01 ft was estimated to be only 0.6×10^{-3} $cm^2(cpd)^{-1}$. This is negligible, compared to even the lowest spectral energy densities in Fig. 1.

The variances in the two main tidal bands obtained by numerically integrating the spectral energy densities were compared with the variances computed from the published amplitudes of the relevant tidal constituents. The comparisons were made for the bands 0.847-1.102 cpd and 1.763-2.119 cpd for the diurnal and semi-diurnal bands respectively. Two sets of published amplitudes were available, one based on analysis of data for the year 1934, and one for 1968. It was necessary to adjust the published amplitudes for node factors appropriate to mid 1966 (middle of the three year period used for spectrum analysis).

The results are given in Table 1. The agreement between all three values for the dirunal tides is excellent. The maximum difference is of the same order as the variance estimated by integrating the non-tidal background ($\sim 1.2 cm^2 (cpd)^{-1}$) over the same band.

In the semi-diurnal band, the spectral estimate agrees very well with the variance from the 1934 analysis, the difference being smaller than the difference between the variances from the 1934 and 1968 analyses. This latter difference is due almost wholly to a difference in amplitude of M_2 (1934 analysis: 51.75 cm, 1968 analysis: 51.18 cm). Amplitude differences of this order, obtained from analyses for the same port but using data from different years, have been noted before (e.g. Rossiter, 1969). It would be interesting, though probably of little practical value, to see how the M_2 amplitude and phase for Sydney vary in analyses based on each of the 19 years for which data are now available.

The background energy density in Fig. 1 at frequencies above about 1.2 cpd is in the range $0.2\text{-}1 cm^2 (cpd)^{-1}$. This is of the same order as at Honolulu (MC, Fig. 7), and less than an order of magnitude below the corresponding values at Newlyn (England) (MC, Fig. 12).

The background energy density just outside the semi-diurnal tidal band in Fig. 1 (i.e. at ~1.7 and ~2.2 cpd) might be regarded as being slightly higher than would be expected from a gross smoothing of the background over the whole spectrum. This could be due to some form of (non-linear) interaction between the strong semi-diurnal tides and the "background", to non-linearity in the tide recorder, or to computation effects (see e.g. Rossiter and Lennon (1968), Amin (1976)). As a rough estimate, the rise might be taken to average $0.4 cm^2 (cpd)^{-1}$ over the range 1.7 to 2.2 cpd, giving a total variance of only $0.2 cm^2$. Further work on the point does not seem to be warranted, since the effect, whatever its cause, is so small.

THE SUB-TIDAL SPECTRUM

Several computations of the spectrum in the range 0-0.8 cpd were made, using mean sea levels at 12 hour intervals. The mean sea levels were computed from the basic set of hourly heights by using Munk's "Tidekiller" filter (Munk, personal communication). This is a symmetric low pass filter whose weights have been chosen so that the filter amplitude response is very small in both the diurnal and semi-diurnal tidal bands (Fig. 2). Table 2 gives the normalised filter coefficients. (In the inter-tidal band, the amplitude response rises to ~0.042. Above 2.1 cpd, there are several extrema in amplitude response, with absolute values less than .02. The very low energy densities in these parts of the spectrum (Fig. 1) make such extrema unimportant). At sub-tidal frequencies, the filter behaves as a low-pass filter, with half-power cut-off at 0.31 cpd. The spectra in Fig. 3 have been corrected for the effect of the filter.

The solid line in Fig. 3 shows the spectrum 0-0.8 cpd, with resolution 0.02 cpd, computed from the whole 19 years of data. The spectrum as drawn is a weighted mean of three separate spectra, each of which showed the same features. The separate spectra were computed from three portions of the 19 year data, chosen to avoid two small gaps in the data. Separate calculations (not plotted) were made by choosing data segments in the four months beginning each 15 November to represent summer, and in the four months beginning each 15 May to represent winter, but the summer and winter spectra did not appear to differ significantly from one another, or from the solid curve in Fig. 3. A further calculation covered the range 0-0.1 cpd, with 0.002 cpd resolution. This showed a peak about 1.7 times the background level at the two adjacent frequencies of 0.044 and 0.046 cpd, but this is probably not significant. (Degrees of freedom=35, so 95 percent confidence band is 0.66/1/71).

The only significant feature of the solid curve in Fig. 3 is the broad peak around 0.35-0.4 cpd. This peak was evident in the summer and winter spectra, and in each of the three spectra that were combined to give Fig. 3. It has already been noted in discussing Fig. 1.

Although the 0.4 cpd peak is statistically significant, it represents little energy. The area under the peak, and above some assumed continuum level, was estimated to represent a variance of only ~0.5 cm^2 (0.7 cm rms), although such an estimate is difficult where the spectrum is so steep.

More work needs to be done before the 0.4 cpd peak can be explained. It might represent a local accummulation of continental shelf wave energy, due to a zero group velocity of one of the shelf wave modes (Buchwald and Adams (1968), Cutchin and Smith (1973)). With presently available estimates of shelf parameters, however, it would have to correspond to at least the second mode, since the estimated zero group velocity for the first mode is ~0.8 cpd. The presence of a second mode in this area has not yet been demonstrated, although no special search for it has been tried. Other mechanisms involving shelf waves have been suggested (Buchwald, 1976).

The root-mean-square amplitude of ~0.7 cm, estimated above, is perhaps not so small, when compared with the observed amplitude of shelf waves on this coast (~5 cm). Presumably second mode shelf waves, if present would have an rms amplitude appreciably less than 5 cm.

The dashed curve of Fig. 3 is the spectrum of "adjusted" sea level at Sydney, for the period November 1961-December 1969. Three-hourly atmospheric pressures for Mascot were filtered, using a low-pass filter with a frequency response similar to that of Munk's filter, to give mean pressures at 12-hour intervals. These were then added to the corresponding mean sea levels, to form the adjusted sea level series. With units of centimeters and millibars respectively, this process is equivalent to assuming that the ocean responds to pressure changes as an inverse barometer. The spectrum of adjusted mean sea level shows a peak at a lower frequency - about 0.25-0.3 cpd, instead of 0.35-0.4 cpd found from the spectrum of observed mean sea level. Interpretation of the shift is difficult, due to the complex relation between sea level, atmospheric pressure and shelf waves (Hamon, 1962, 1966).

Both curves in Fig. 3 show a weak maximum at 0.72-0.74 cpd. The rise in the spectra at the top end (0.76-0.8 cpd) is probably due to the near-by diurnal tide band, though it is not clear why the rise is more pronounced for adjusted sea levels. (Note that correction for Munk's filter involves dividing the computed spectrum by 0.0018 at 0.8 cpd!).

The two curves in Fig. 3 compare reasonably with the corresponding curves in Fig. 2 of Hamon (1962), if allowance is made for the smaller degrees of freedom due to use of only 1.5 years of data.

DISCUSSION

The results show only one feature of interest - the 0.35-0.4 cpd peak in sea level spectrum. Otherwise, the spectrum for Sydney already given over three decades of frequency (0.0005-0.5 cpd) (Hamon, 1968) has been confirmed at least above about 0.01 cpd, and extended to 4 cpd. With the scales used in 1968 however, the extension shows a very small "variance/octave" (see Appendix) at the inter-tidal and trans-tidal frequencies studied. From the background levels estimated from Fig. 1, the variance/octave is 1 cm^2 at 1.5 cpd and 0.7 cm^2 at 3.5 cpd (cf. levels in the range 10-20 cm^2 for frequencies between 0.01 and 0.2 cpd, Fig. 1 of Hamon (1968)).

ACKNOWLEDGEMENT

I thank the Maritime Services Board of N.S.W. and the Archives Office of N.S.W. for making the valuable original Fort Denison records available. Mr M. A. Greig of this Division was responsible for supervising the digitising, and has also given valuable comments on this paper.

REFERENCES

Amin, M. 1976. The fine resolution of tidal harmonics. *Geophys. J.R. Astr. Soc.*, 44, 293-310.

Buchwald, V.T., and Adams, J.K. 1968. The propagation of continental shelf waves. *Proc. Roy. Soc., A.*, 305, 235-250.

Buchwald, V.T. 1976. Diffraction of shelf waves by an irregular coastline. Proc. I.U.T.A.M. Conf. on Waves on Water of Varying Depth. *Aust. Acad. Sci.*, Canberra. (In press).

Cartwright, D.E. 1968. A unified analysis of tides and surges round north and east Britain. *Phil. Trans. Roy. Soc. Lond.*, Ser. A, 263, 1-55.

Cutchin, D.L., and Smith, R.L. 1973. Continental shelf waves: low frequency variations in sea level and currents over the Oregon continental shelf. *J. Phys. Oceanogr.*, 3, 73-82.

Greig, M.A. (In press). Digitising of Torres Strait tidal records. *CSIRO Div. Fish. Oceanogr. Rep.* No. 74.

Groves, G.W., and Hannan, E.J. 1968. Time series regression of sea level on weather. *Rev. Geophys.*, 6, 129-174.

Hamon, B.V. 1962. The spectrum of mean sea level at Sydney, Coff's Harbour and Lord Howe Island. *J. Geophys. Res.*, 67, 5147-5155. (Correction: ibid, 68, 4635, 1963).

Hamon, B.V. 1966. Continental shelf waves and the effects of atmospheric pressure and wind stress on sea level. *J. Geophys. Res.*, 71, 2883-2893.

Hamon, B.V. 1968. Spectrum of mean sea level at Lord Howe Island in relation to circulation. *J. Geophys. Res.*, 73, 6925-6927.

Jenkins, G.M., and Watts, D.G. 1968. Spectral analysis and its applications. Holden Day, California, 525 pp.

Lennon, G.W. 1965. The treatment of hourly elevations of the tide, using an IBM 1620. *Int. Hyd. Rev.*, 42, 125-148.

Munk, W.H., and Bullard, E.C. 1963. Patching the long-wave spectrum across the tides. *J. Geophys. Res.*, 68, 3627-3634.

Munk, W.H., and Cartwright, D.E. 1966. Tidal spectroscopy and prediction. *Phil. Trans. Roy. Soc. Lond.* Ser. A, 259, 533-581.

Rossiter, J.R. 1969. A note on some troublesome harmonic constituents. *Inst. Coastal Oceanog. and Tides, Int. Rep.* No. 14, 8 pp.

Rossiter, J.R., and Lennon, G.W. 1968. An intensive analysis of shallow water tides. *Geophys. J. R. Astr. Soc.*, 16, 275-293.

TABLE 1

Comparison of variances computed from spectrum and from published analyses

Band	Variances (cm^2)		
	1934 Analysis	1968 Analysis	Spectrum (Fig. 1)
Diurnal	200.6	200.4	200.4
Semi-diurnal	1448.7	1414.2	1442.9

TABLE 2

Munk's "Tidekiller" Filter
Normalised Weights

Index	Weight	Index	Weight	Index	Weight
1	.0013307	11	.0180727	21	.0338603
2	.0030073	12	.0195518	22	.0354118
3	.0047028	13	.0208050	23	.0370094
4	.0060772	14	.0219260	24	.0386839
5	.0072261	15	.0234033	25	.0395287*
6	.0085349	16	.0251492		
7	.0101603	17	.0278167		
8	.0122665	18	.0300054		
9	.0146225	19	.0314959		
10	.0165525	20	.0325633		

*The filter is symmetric about this central weight

APPENDIX

THE INTERPRETATION AND PLOTTING OF SPECTRA

Spectrum analysis is still confusing to many people, and the results are regarded as difficult to interpret. I have been guilty of adding to the confusion by using (in earlier papers) a method of presentation common in meteorology, but not in oceanography, and (worse) of inventing the term "variance/octave". A few words of explanation might help.

The spectral energy density, which we will denote x(f), is usually computed at equal intervals of frequency, Δf, and in units of (data unit)^2 (cycle per time unit)$^{-1}$. The word "energy" is historical; "variance" would be a better general term, and will be used here.

The values of x(f) may be interpreted as an "analysis of variance", based on frequency. Thus the part of the total variance of the original record that is due to frequencies between f_1 and f_2 is

$$\int_{f_1}^{f_2} x(f).df(\text{data unit})^2,$$

and the *total* variance of the original record σ^2, should equal

$$\int_{0}^{f_N} x(f).df,$$

where f_N is the Nyquist frequency (1/2 Δt, where Δt is the time interval between observations). (This is a very useful check on units and scaling - and more necessary than one would think, since some available computer programs have been found to be in error by a factor of 2π!).

If x(f) is plotted as ordinate and f as abscissa, both on *linear* scales, the plot has the desirable property that the area under the curve, between any two frequencies, is proportional to the variance of the original record in that frequency band. Unfortunately, most geophysical time series have spectra that rise so steeply with decreasing frequency that this simple plotting cannot be used.

A plot of f.x(f) against log f also has the above desirable property, and is the method of plotting favoured in meteorology. It is the method used in Hamon (1962, 1968). It has two main advantages:

(i) a logarithmic frequency scale is more "natural" than a linear one (and of course the only choice when a wide frequency range is being studied), and
(ii) it tends to flatten the usually steep geophysical spectra. Other advantages are that poor (fractional) frequency resolution at the lowest few frequencies of any given analysis is emphasised, and that the ordinate scale has simpler units.

The ordinate f.x(f) has units (data unit)^2; i.e. frequency no longer appears. What does the quantity mean? If one goes back to the meaning of x(f), then f.x(f) can be thought of as the "variance in a frequency band of width f, centered at frequency f", of course this is not literally true unless the spectrum is uniform, but it conveys the idea. Since a band of width f *starting* at f extends to 2f, i.e. through one octave, the notation "variance/octave" seemed appropriate.

If f.x(f) is plotted against $\log_{10}f$ with scale factors β(data unit)^{2}cm^{-1} and α cm/decade respectively, the variance between two frequencies can be estimated from the area (cm^2) under the curve, and bounded by the two frequencies, by multiplying by $\ell n(10)\beta/\alpha$, i.e. by $2.30\beta/\alpha$.

If the process being studied contains a component of fixed frequency with amplitude a, (e.g. a tidal "constituent"), the spectrum will contain a peak centered at the appropriate frequency. The height and width of the peak depend on the frequency resolution used, and the band-width of the analysis method. The *area* under the peak (with due allowance for background, if any) will be a measure of the variance $\frac{1}{2}a^2$ of the component. This applies also when plotting in the form f.x(f) *vs* log f. The point needs stressing, since one might be tempted to think that, because f.x(f) has units of variance, the height (alone) of the peak should be a measure of the variance of the component.

In practice, as in the present paper, one is often forced to use plotting methods that do not preserve the equivalence between variance and area under the curve. (It is worth noting that log(x(f)) or log (f.x(f)) *vs* log f is usually used in turbulence studies, since one often expects a power law, of the form $x(f) \propto f^n$).

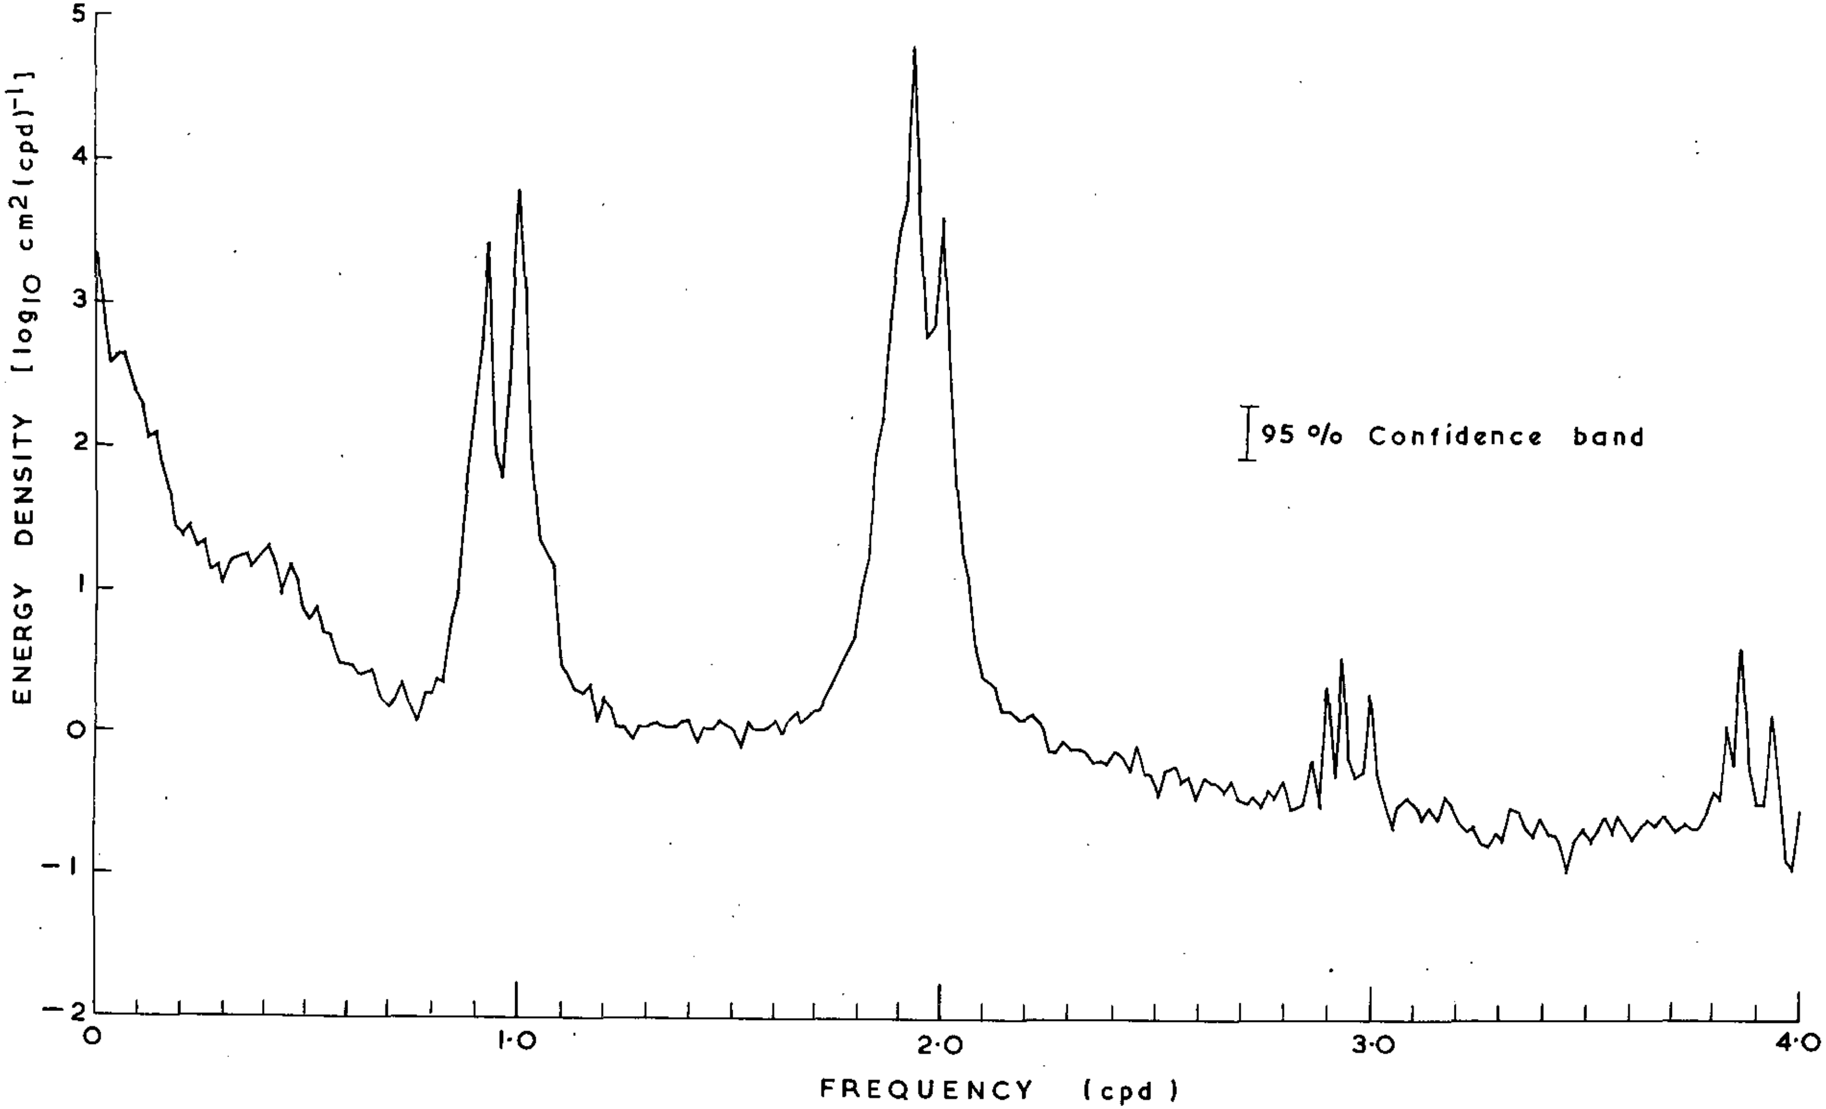

Fig. 1. The spectrum of sea level at Fort Denison (Sydney), 0.4 cycles per day (cpd).

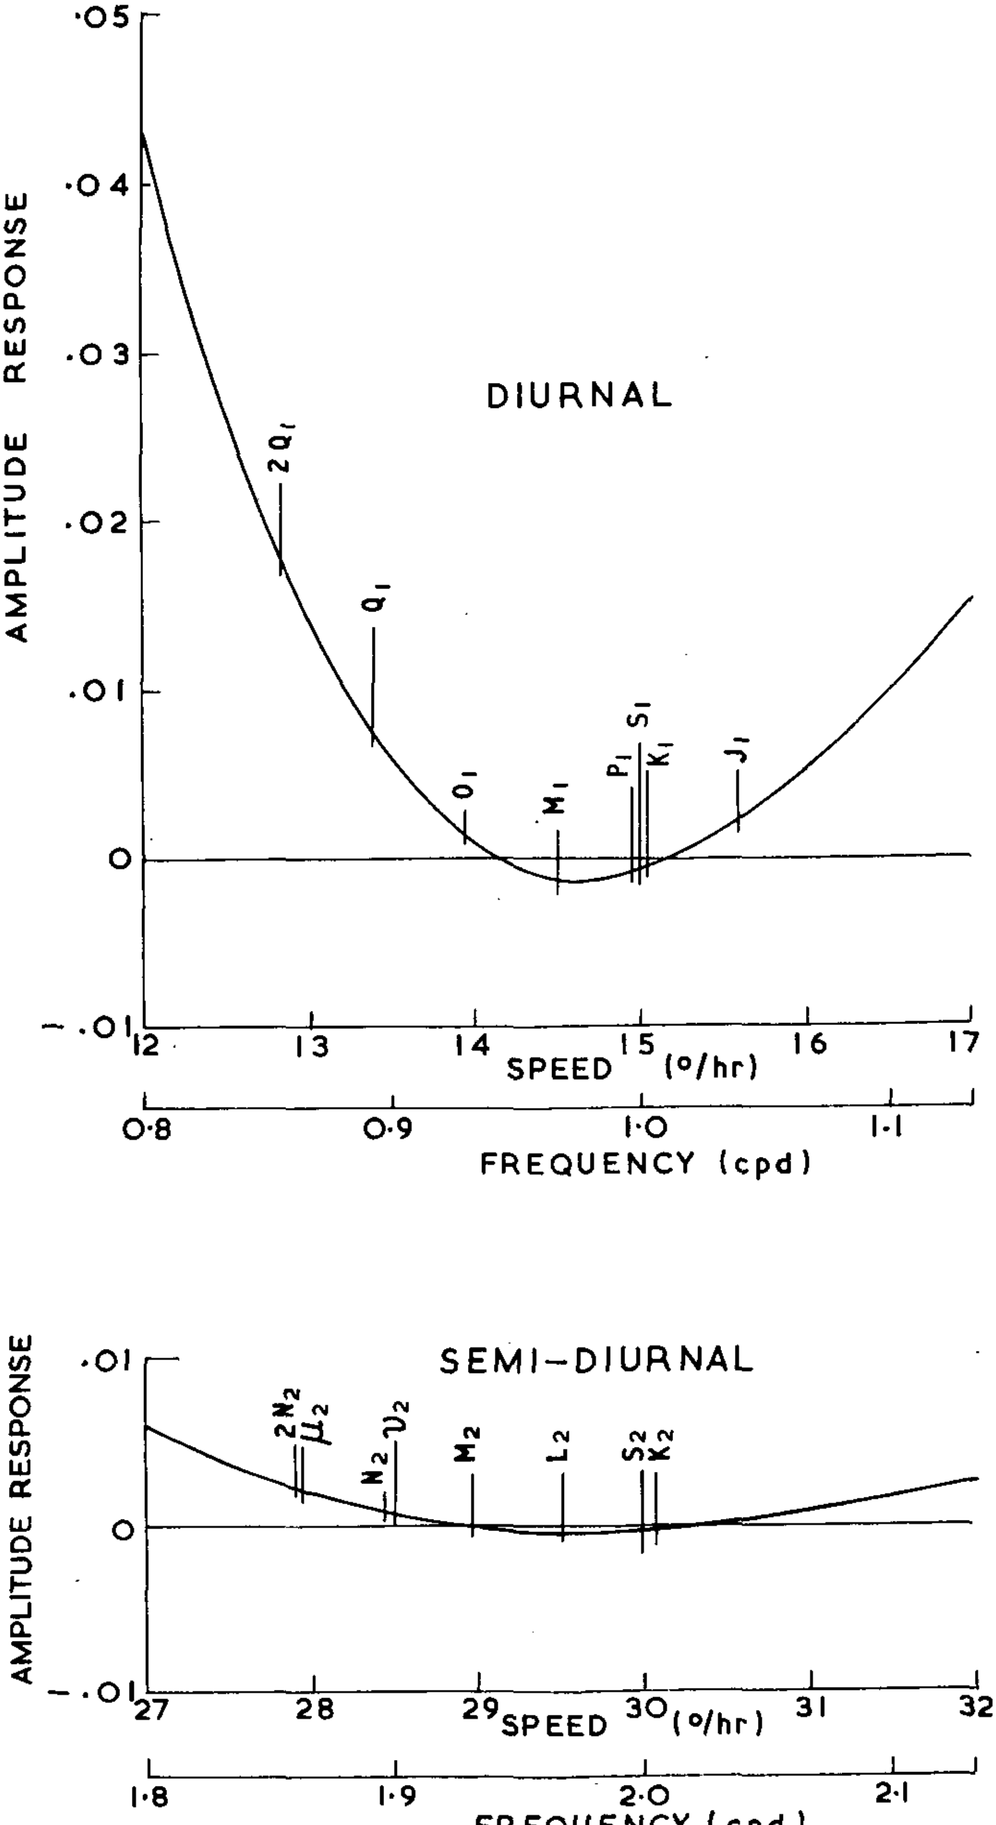

Fig. 2. The amplitude response of Munk's "Tidekiller" filter in the diurnal and semi-diurnal tidal bands. The frequencies of the main tidal lines in each band are indicated.

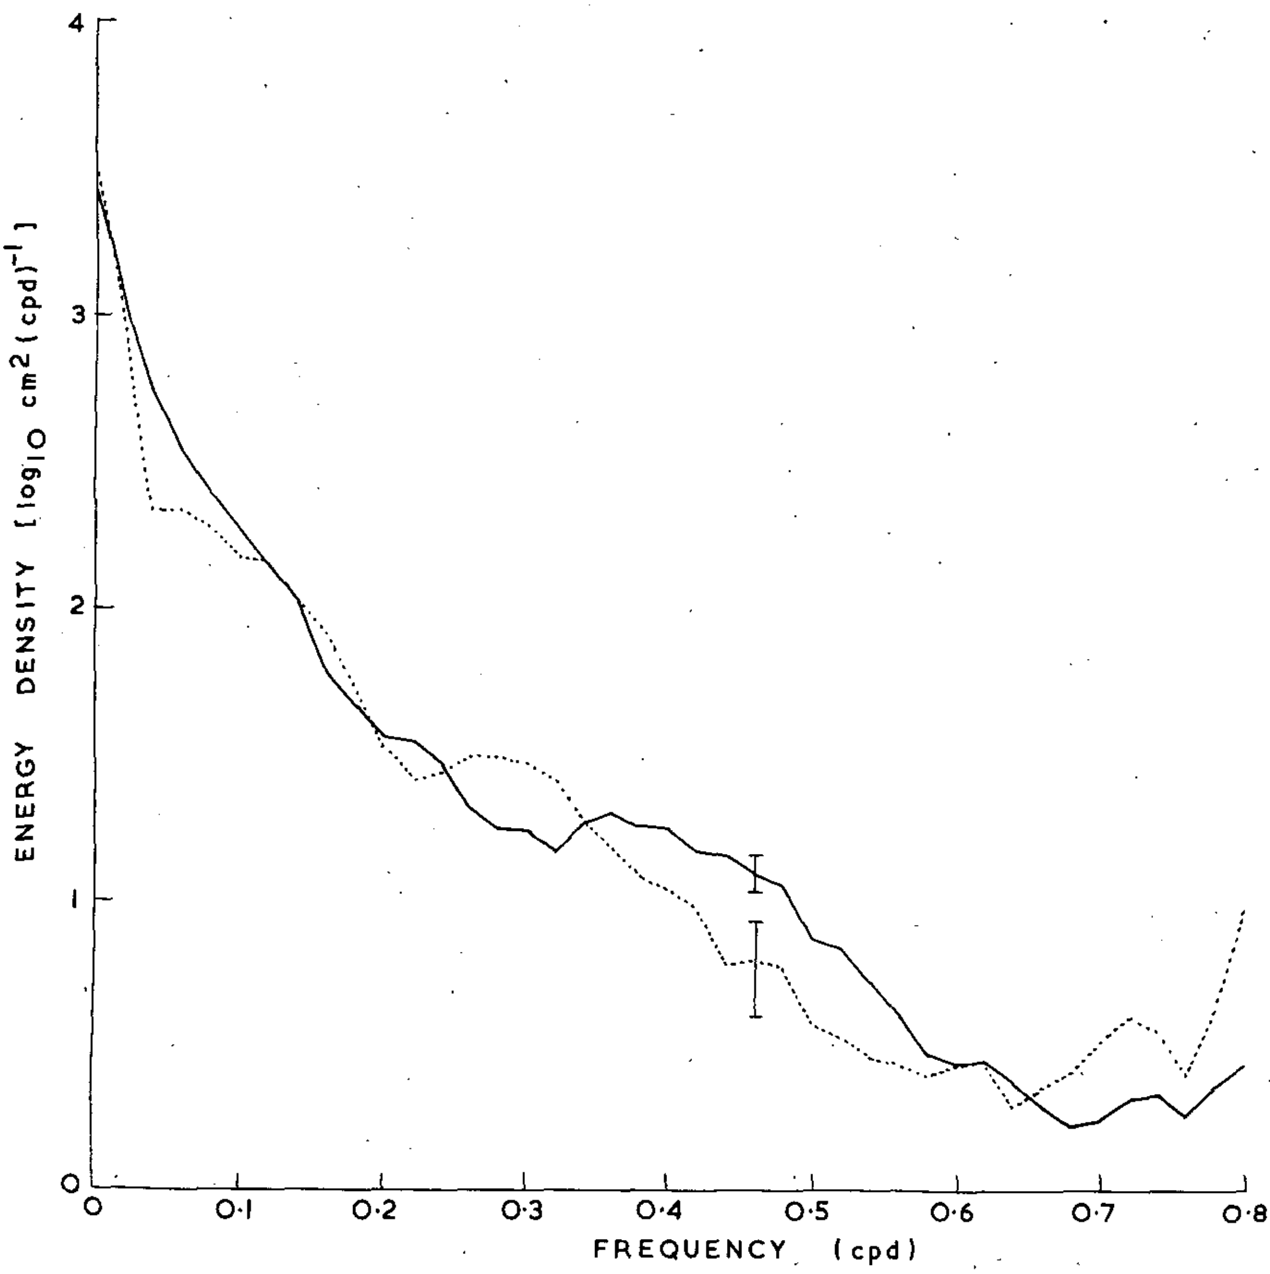

Fig. 3. Solid line: spectrum of observed mean sea level at Fort Denison, 0-0.8 cpd, based on 19 years of data, and corrected for the filter used to compute mean sea level. Dotted line: spectrum of adjusted mean sea level (adjusted to fixed atmospheric pressure) at Fort Denison, from 8 years of data. The vertical bars indicate the 95 percent confidence intervals.