Chapter 3: The analysis of tidal data

3.1 Introduction

Chapter 2, above, is devoted to a discussion of the tidal potential and its immediate consequences – the solid and ocean tides. These tidal signals can be considered to be the inputs to a physical system – a hydrothermal convection cell – whose outputs are the observed variations in effluent temperature, chemistry and flow rate. The input signals are easily predicted, and their characteristic features are discussed in detail in Chapter 2. It would be advantageous to have an equally precise knowledge of the measured output signals, but this is rarely possible. Collecting data on the seafloor is difficult and expensive, so the datasets are often limited in length to a few days or weeks. Furthermore, the resulting time-series tend to be rather noisy. Consequently, the aim of this chapter is to compare the different methods of extracting information from a limited, imperfect hydrothermal time-series.

The analysis of tidal signals has been a subject of great interest for many years. Originally, the motivation was to establish the nature of the ocean tide at a particular port, to ensure the safety of shipping. The sea level in the port was measured with a tide-gauge over a period of time and a harmonic decomposition of the resulting time-series was used to extract the harmonic constants of the port. The number of harmonic constants which can be determined increases with the length of the time-series (Doodson & Warburg, 1941). Given suitable data, this method produces an accurate and theoretically sound description of oceanic tidal signals. Furthermore, it has successfully been applied to the study of earth tides, where the data to be analysed include time-series of gravity, strain and tilt (Melchior, 1983).

It must be stressed, however, that traditional harmonic methods are ill suited to the analysis of hydrothermal time-series. The first reason for this is that data from the seafloor are constrained to span limited short time intervals. In contrast, land-based equipment such as tide gauges can obtain very long time-series at minimal cost. For example, Munk & Cartwright (1966) were able to analyse 19 years of tide observations in order to establish the harmonic constants of the ocean tide at Newlyn and Honolulu. The second obstacle to the application of harmonic methods is that hydrothermal time-series tend to be much noisier than traditional tidal signals - such as those from tide-gauges or strain meters - which appear smooth in the time-domain. To a good approximation, these smooth tidal signals contain spectral power only at the tidal frequencies. A hydrothermal time-series, on the other hand, often has a significant amount of power in a smooth background spectrum, in addition to power in line components at the tidal frequencies.

In summary, the task of extracting tidal information from hydrothermal time-series is made difficult by the fact that the datasets tend to be (1) short and (2) noisy. The implications of these problems are discussed in this chapter by dividing the techniques of time-series analysis into two classes – parametric and non-parametric.

Among non-parametric techniques, the periodogram spectral estimator has been widely used in the interpretation of hydrothermal time-series despite the fact that it is demonstrably ill suited to short and noisy data. The multiple window method (Thomson, 1982) is much better suited to the problem, and is recommended as the best non-parametric technique for hydrothermal time-series. The MWPS code which implements this technique (A. Chave, *pers. comm.*, 1999) is applied to hydrothermal time-series data in Chapter 4.

Among parametric techniques, the traditional Harmonic Method is of limited use with a short time-series, and it is recommended that the Admiralty Method (Admiralty Tidal Handbook, No. 3) be used in its place. Furthermore, to combat the problems of noise, the Admiralty Method can be made more robust (Chave *et al.,* 1987), and combined with a technique which removes drift from the data signal (Tamura *et al.,* 1991). The subsequent hybrid technique is recommended as the optimal parametric technique for extracting information from hydrothermal time-series. This technique has been implemented in the new computer code HYBRID. It is applied to hydrothermal time-series data in Chapter 4.

3.2 Non-parametric analysis

Many hydrothermal time-series clearly have the appearance of tidal signals when viewed in the time-domain because they display diurnal and semi-diurnal periodicity (Chapter 4). However, it should not be assumed that a signal is tidal without further supporting evidence. Consequently, the first step in the examination of a time-series should be an impartial, non-parametric analysis of the data, in which the prejudice of the analyst plays no part. An attempt to fit the data with a parametric model should only be made if evidence of tidal modulation has previously been uncovered by a non-parametric analysis.

A hydrothermal time-series can be considered to be a finite sample from a stationary process. The problem is then to make an estimate of the power spectrum of this process, using the information from the finite dataset. It must be stressed that the true power spectrum of the process can never be deduced from a finite amount of data (Thomson, 1982). However, the data can be used to construct several different *estimates* of the true power spectrum (Section 3.2.1).

Once an estimate of the power spectrum has been made, it can be examined to see if there is evidence of significant power at the tidal frequencies. The principal components of the tidal potential are *M2* (period 12.42 hours), *S2* (period 12 hours), *K1* (period 23.94 hours) and *O1* (period 25.82 hours) (Table 2.1). It is likely that *any* tidal signal has significant spectral power at these frequencies. In most tidal signals the *M2* component dominates (Schwiderski, 1980). Consequently, an initial test of whether a signal is influenced by the tides is simply to see if it has significant spectral power at ~1.93 cpd (equivalent to a period of ~12.42 hours).

It must be stressed that the presence of spectral power at ~1.93 cpd does not *prove* that a time-series is tidally modulated. It may be that there is some other physical process – entirely independent of the moon - which causes modulation at this frequency. For the purposes of this dissertation, however, it is supposed that this is not the case. Consequently it is assumed that all signals whose spectra are *correlated* with tidal processes are in fact *caused* by tidal processes.

3.2.1 Simple Fourier transform methods

3.2.1.1 The periodogram

The simplest spectral estimate is the periodogram, which is based on the discrete Fourier transform (Press *et al.,* 1986). Given a time-series dataset $\{x_0,\ldots,x_{N-1}\}$, the periodogram estimate $\hat{S}_P(f)$ of the power at frequency f is:

$$\hat{S}_P(f) = \left| \sum_{n=0}^{N-1} x_n e^{2\pi i f n / N} \right|^2 \tag{3.1}$$

Normally, this spectral estimate is made at the discrete set of frequencies $\{f_j = j/(N\Delta t)\}$, where Δt is the sampling interval between successive data-points in the time-domain, and j lies in the range $\{-N/2,\ldots,0,\ldots,N/2\}$. (It is assumed, for convenience, that N is even.) Unfortunately, the periodogram $\hat{S}_P(f)$ is not a good estimator of the true spectrum $S(f)$.

The Fejer kernel $K_F(f)$ is defined by:

$$K_F(f) = \frac{1}{N^2}\left[\frac{\sin(\pi f)}{\sin(\pi f/N)}\right]^2 \quad (3.2)$$

It can be shown that the periodogram estimate $\hat{S}_P(f)$ is equal to the convolution of the true spectrum $S(f)$ with the Fejer kernel $K_F(f)$ as follows:

$$\hat{S}_P(f) = \int_{s=-\infty}^{\infty} K_F(s) S(f-s) ds \quad (3.3)$$

Consequently, the periodogram estimate $\hat{S}_P(f_j)$ of the power at frequency f_j is a weighted average of the true power at *all* other frequencies. For this reason, the periodogram is said to be biased by spectral leakage.

3.2.1.2 Local and broadband bias

The bias can be divided into two parts by defining a bandwidth W in the frequency domain. Leakage from frequencies in the range (f_j-W, f_j+W) is termed local bias, while leakage from the other more distant frequencies is termed broadband bias. The bandwidth can be expressed in non-dimensional form, as the time-bandwidth product $NW\Delta t$. This is often chosen to be an integer, in which case it represents the number of frequency bins, on either side of f_j, which contribute to the local bias in the estimate at frequency f_j. Consequently the local bias in the estimate at f_j consists of the spectral leakage from the frequencies $\{f_{j-NW\Delta t}, \ldots, f_{j+NW\Delta t}\}$ (Thomson, 1982).

3.2.1.3 The windowed periodogram

The periodogram estimate of equation (3.3) suffers very badly from broad band bias, because the side-lobes of the Fejer kernel fall away slowly (equation (3.2)). The problem can be addressed by multiplying the dataset $\{x_0, \ldots, x_{N-1}\}$ by a data window $\{w_0, \ldots, w_{N-1}\}$ before performing the Fourier transformation (Press *et al.,* 1986). This produces a spectral estimate known as the windowed periodogram estimate $\hat{S}_{WP}(f)$, defined by:

$$\hat{S}_{WP}(f) = \left|\sum_{n=0}^{N-1} w_n x_n e^{2\pi i f n/N}\right|^2 \quad (3.4)$$

In this case, the estimate $\hat{S}_{WP}(f)$ is the convolution of the true spectrum $S(f)$ with a spectral window $W_S(f)$:

$$\hat{S}_{WP}(f) = \int_{s=-\infty}^{\infty} W_S(s) S(f-s) ds \quad (3.5)$$

The data window $\{w_0, \ldots, w_{N-1}\}$ is usually normalised so that:

$$\sum_{n=0}^{N-1} w_n^2 = N \tag{3.6}$$

It then follows that the spectral window $W_S(f)$ is given by:

$$W_S(f) = \left| \sum_{n=0}^{N-1} w_n e^{2\pi i f n / N} \right|^2 \tag{3.7}$$

Many data windows are in common use, and the choice of a particular data-window is made according to the desired bias properties of its spectral window $W_S(f)$.

3.2.1.4 Discrete prolate spheroidal sequences

A family of data windows known as the discrete prolate spheroidal sequences (or DPSS's) is of particular interest. Given a choice of the time-bandwidth product ($NW\Delta t$) it is possible to define an ordered set of N DPSS's $\{\{v_n^{(0)}\},\ldots,\{v_n^{(N-1)}\}\}$ (Thomson, 1982). For a particular value of the time-bandwidth product ($NW\Delta t$), the first DPSS ($\{v_n^{(0)}\}$) can be shown to have the smallest local bias of all possible data windows (Slepian, 1978).

This optimality of the DPSS's - in particular the first such window $\{v_n^{(0)}\}$ - makes them a natural choice as the data windows for use in windowed periodograms. The choice of time-bandwidth product should be made individually for each problem, by weighing up the likely effects of local and broad-band bias and the desired resolution of the spectral estimate.

The preceding remarks have important consequences for the analysis of tidal signals from the noisy, short datasets which are obtained at hydrothermal systems. The true spectrum of the tidal potential is known to consist of a large number of closely-spaced harmonic lines, and it is expected that all tidal signals will have similar line components in their spectra, in addition to the background noise from other processes. When spectral estimates are made, the large concentrations of power in the harmonic lines will tend to leak into adjacent frequencies, making it difficult to distinguish individual lines. The problem gets worse as the time-interval spanned by the dataset ($(N-1)\Delta t$) decreases, because the separation between adjacent frequency bins is $1/(N\Delta t)$. For example, a dataset spanning a time-interval of 14 days gives a frequency bin separation of ~0.07 cpd. The set of tidal harmonic lines in the semi-diurnal frequency band covers a bandwidth of about ~0.3 cpd and would therefore contain only 4 frequency bins. This is clearly insufficient to resolve the location of a useful number of harmonic lines (Figure 2.7).

The bias problems of the unwindowed periodogram (equation (3.1)) are severe, and it should not be used in tidal analysis. Nonetheless, many studies of hydrothermal tidal signals contain little or no discussion of the spectral estimators which are employed, and it is quite possible that they rely on unwindowed periodograms (Johnson & Tunnicliffe, 1985; Little *et al.,* 1988; Chevaldonné *et al.,* 1991; Johnson *et al.,* 1994; Copley *et al.,* 1999). Windowed periodograms are an improvement, and a DPSS should be used as the data window to minimise bias, with the time-bandwidth product chosen to suit the particular problem.

Even with a DPSS as the data window, the windowed periodogram estimate (equation (3.4)) is not optimal for application to hydrothermal time-series. Periodogram estimates are not statistically efficient, in the sense that the variance of the estimate does not decrease as the number of data-points (N) increases. The two traditional methods of countering this problem are band-averaging and section-averaging (Press *et al.,* 1986; Chave *et al.,* 1987), but both have the undesirable side effect of reducing resolution in the frequency domain. This can ill be afforded in the analysis of short tidal signals.

3.2.1.5 Band-averaging

In band-averaging, a windowed periodogram estimate $\hat{S}_{WP}(f)$ is smoothed in the frequency domain by taking its convolution with a suitable function $K(f)$. Therefore the band-averaged spectral estimate $\hat{S}_{BA}(f)$ is:

$$\hat{S}_{BA}(f) = \int_{s=-\infty}^{\infty} K(s)\hat{S}_{WP}(f-s)ds \tag{3.8}$$

The convolution in equation (3.8) has the effect of smearing the power from tidal harmonic lines into neighbouring frequencies and therefore leads to a loss of resolution in the frequency domain. Consequently, band-averaging is not recommended for the analysis of hydrothermal time-series.

3.2.1.6 Section-averaging

In section averaging, the time-domain dataset $\{x_0,\ldots,x_{N-1}\}$ is split into M subsections, which may overlap with each other. For each section m, a windowed periodogram is used to produce a spectral estimate $\hat{S}_m(f)$. The section-averaged spectrum $\hat{S}_{SA}(f)$ is then calculated by taking an average of the spectral estimates obtained from each section. The averaging process can be done in several ways (Chave *et al.,* 1987), but the most obvious method is simply to take the arithmetic mean:

$$\hat{S}_{SA}(f) = \frac{1}{M}\sum_{m=1}^{M}\hat{S}_m(f) \tag{3.9}$$

Regardless of the details of the averaging process, this approach inevitably reduces resolution in the frequency domain, because the individual spectra ($\hat{S}_m(f)$) are obtained from subsections which are shorter than the original time-series. Consequently, section-averaging is not recommended for the spectral analysis of hydrothermal time-series.

3.2.1.7 Conclusion

In summary, methods based on the periodogram suffer from severe problems as means of extracting tidal information from hydrothermal time-series, although they have been used extensively for this purpose. The unwindowed periodogram of equation (3.3) is badly biased, while the windowed periodogram of equation (3.4) must be smoothed, by band- or section-averaging, to make it reasonably efficient. However, both of these methods lead to reduced resolution in the frequency domain, making it difficult to separate individual line components. It is strongly recommended that the periodogram methods of Section 3.2.1 should *not* be used to analyse hydrothermal time-series.

3.2.2 Multiple Window Power Spectra

The optimal non-parametric technique for analysing hydrothermal time-series is arguably the multiple window (or multiple taper) method of Thomson (1982), which makes use of the DPSS data windows discussed in Section 3.2.1. A brief synopsis of the method is provided here.

The first step is to choose a value for the time-bandwidth product $NW\Delta t$. The choice is influenced by the length of the time-series $(N\text{-}1)\Delta t$ and by the desired accuracy of the spectral estimate in the frequency domain. Inherent in this choice is a trade-off between efficiency and spectral resolution. A small time-bandwidth product allows closely spaced lines to be separated, but at the cost of low efficiency. Conversely, a large time-bandwidth product increases efficiency at the cost of smearing adjacent line components together. In the case of hydrothermal time-series, it is helpful that the spacing of the tidal harmonic lines in the frequency domain is known *a priori*.

As discussed in Section 3.2.1, a particular choice of time-bandwidth product $NW\Delta t$ defines a set of DPSS's $\{\{v_n^{(0)}\},\ldots,\{v_n^{(N-1)}\}\}$. In the multiple window method, each DPSS is used as the data window in a separate windowed periodogram of the entire time-series. This creates a set of N separate spectral estimates $\{\hat{S}_0(f),\ldots,\hat{S}_{N-1}(f)\}$ which are known as the eigenspectra. The multiple window spectral estimate $\hat{S}_{MWPS}(f)$ is then obtained by taking a weighted average of the eigenspectra at each frequency. It can be shown that it is acceptable to restrict attention solely to the first $K=2NW\Delta t$ eigenspectra (Thomson, 1982). Consequently, $\hat{S}_{MWPS}(f)$ is given by an expression of the form:

$$\hat{S}_{MWPS}(f)=\sum_{k=0}^{K-1} d_k(f)\hat{S}_k(f) \tag{3.10}$$

The weights $d_k(f)$ are calculated from the data in a manner described by Thomson (1982). It is important to note that this method relies on windowed periodograms of the *entire* data sequence, and does not perform any convolutional smoothing in the frequency domain. Thus, unlike band- or section-averaging, it achieves statistical efficiency without sacrificing any spectral resolution. This is clearly desirable when seeking to identify closely spaced tidal harmonic lines in short time-series.

3.2.2.1 Identification of line components

The multiple window method is particularly well suited to tidal analysis as it incorporates an explicit procedure for the identification of significant line components in the spectrum. Firstly, the complex harmonic constants for the time-series are estimated, by regression over the eigenspectra, for each frequency bin (f_j). It must then be decided whether the estimated power of the harmonic constants represents a statistically significant line component. To this end, an F-statistic is constructed for the estimated harmonic constants in each frequency bin. It can be shown that this statistic is distributed as $F_{2,2K-2}$ for Gaussian data. Consequently, an unusually large value of the F-statistic is interpreted as evidence that there is a significant line component at the frequency under consideration. For a time-series of N data-points, Thomson (1982) suggests that the F-statistic should be considered significant if its value exceeds the 100*(N-1)/N percentage point of the $F_{2,2K-2}$ distribution. In Chapter 4, this criterion is used to decide if line components in the spectrum are significant.

Once the significant line components have been identified, a residual time-series is constructed by subtracting the signal generated by the significant line components from the original time-series. The power spectrum of the residual time-series is then calculated using the multiple window method. Of course, with the power in the significant line components removed from the residual signal, its spectrum is much less prone to distortion by spectral leakage. Finally, the power in the significant line components is added to the spectrum of the residual time-series to give the overall spectral estimate.

A computer code to perform these calculations, MWPS (A. Chave, *pers. comm.*, 1999) is used for the non-parametric analysis of seafloor data in Chapter 4.

From the point of view of hydrothermal time-series, the multiple window method offers reasonable solutions to the two key problems which afflict seafloor data - shortness and noisiness. Firstly, by retaining the highest possible resolution in the frequency domain, it makes maximum use of the limited length of the dataset, unlike the band- and section-averaged periodogram methods. Secondly, by identifying significant line components, and removing them before calculating the background spectrum, it reduces the effect of spectral leakage from the powerful tidal harmonic lines into the surrounding spectrum.

The multiple window method is very well suited to the analysis of data with a tidal component. It is therefore surprising, and disappointing, that it seems to have been used only once for this purpose (Thomson *et al.,* 1986). Nonetheless, it is proposed here that it should be the non-parametric method of choice for all subsequent analysis of seafloor hydrothermal time-series.

3.3 Parametric analysis

Non-parametric analysis (Section 3.2) can be used to examine a time-series without prejudice as to its nature or cause. If non-parametric analysis suggests that the time-series may have a tidal component, the next step is to postulate a parametric model for the time-series and estimate the values of the parameters. The underlying objective is to extract as much useful information from the data as possible, given the parametric form assumed for the data. The choice of a particular parametric model is made according to the expected structure of the time-series. It must be remembered, therefore, that each parametric model contains an implicit set of assumptions, or prejudices, about the nature of the time-series.

The Harmonic Method (Section 3.3.1) relies on fewer assumptions than the Admiralty Method (Section 3.3.2). In a Harmonic Method decomposition, the sole assumption is that all tidal signals have spectral lines at the same discrete set of frequencies as the tidal potential. The disadvantage of the Harmonic Method is that a large number of harmonic constants must be estimated to yield a reasonably accurate description of the tidal signal. This requires a long time-series (~1 year or more) and so the Harmonic Method is unsuitable for the short time-series obtained at the seafloor.

In contrast, the Admiralty Method of Section 3.3.2 requires only 8 harmonic constants to describe a tidal signal, and the values of these parameters can be estimated from relatively short time-series (~7 days). However, these advantages come at the cost of making extra assumptions about the nature of tidal signals. In common with the Harmonic Method, it is assumed in the Admiralty Method that all tidal signals have harmonic lines at the same discrete set of tidal frequencies. In addition it is supposed that certain magnitude and phase relationships which hold between selected harmonic components of the tidal potential hold good for all other tidal signals. The Admiralty Method is very successful in describing the ocean tides in the world's ports, which suggests that these assumptions are justified.

The Bayesian method of drift removal, discussed in Section 3.3.3, relies on the assumption that the drift component of a time-series is smooth. It yields parameter estimates which are maximum likelihood estimates when the drift has a particular stochastic structure.

Finally, the Hybrid Method of Section 3.3.4 combines the ideas of the earlier sections in a new computer code (HYBRID) designed for the parametric analysis of tidal signals. The code decomposes a time-domain dataset into subsections and analyses each subsection separately. The code removes the non-tidal part of the data within each subsection in one of two ways, with the choice being made by the user. In the first method, the mean value of the subsection is removed. The Admiralty Method harmonic constants are then estimated by solution of the appropriate normal equations. In the second method, a Bayesian estimate of the non-tidal drift in the signal is made concurrently with the estimates of the Admiralty Method harmonic constants. The estimated Admiralty Method harmonic constants for each subsection are then combined in a robust manner (Chave *et al.,* 1987) to give an overall estimate of the tidal component of the original data signal.

3.3.1 The Harmonic Method

It is assumed in the Harmonic Method (Section 2.3.4) that an arbitrary tidal signal $\zeta(t)$ can be written as a sum of harmonic terms at the N tidal frequencies $\{\omega_1,\ldots,\omega_N\}$, plus an error term $\varepsilon(t)$:

$$\zeta(t) = \sum_{j=1}^{N} \hat{H}_j \cos\left(\omega_j t - \hat{g}_j\right) + \varepsilon(t) \tag{3.11}$$

Thus, given the set $\{\omega_1,\ldots,\omega_N\}$ *a priori*, equation (3.11) can be postulated as a parametric model for a time-series exhibiting tidal modulation. The set of harmonic constants $\{\hat{H}_1,\ldots,\hat{H}_N$

, $\hat{g}_1,\ldots,\hat{g}_N\}$ is then chosen to provide the best-fitting model for a given time-series. Usually, the harmonic constants are chosen to minimise the least-squares expression:

$$\int [\varepsilon(t)]^2\, dt = \int \left[\zeta(t) - \sum_{j=1}^{N} \hat{H}_j \cos\left(\omega_j t - \hat{g}_j\right) \right]^2 dt \tag{3.12}$$

The harmonic constants which minimise the error of equation (3.12) can be found directly by solving the appropriate normal equations (Press *et al.,* 1986).

Unfortunately, the Harmonic Method is not a good means of extracting tidal information from a short time-series, for the reasons outlined in Chapter 2. It is only possible to separate harmonic components at angular frequencies ω_1 and ω_2 with a time-series whose length exceeds $2\pi/(\omega_1-\omega_2)$. Consequently, a prohibitively long time-series is required to estimate the harmonic constants at a useful number of tidal frequencies.

3.3.2 The Admiralty Method

The Admiralty Method of tidal analysis represents a distinct improvement over the Harmonic Method for seafloor hydrothermal systems because it can extract reasonably accurate harmonic constants from much shorter time-series (Doodson & Warburg, 1941; Admiralty Tidal Handbook, No. 1; Admiralty Tidal Handbook, No. 3). There are, in fact, two related versions of the Admiralty Method, but only one of them is considered in detail here. The first version is the Annual Grouping Method (Admiralty Tidal Handbook, No.1) which is designed for use with time-series longer than one year. Hydrothermal time-series are rarely as long as this, so the Annual Grouping Method will not be considered further. The second version of the Admiralty Method (Admiralty Tidal Handbook, No.3) is designed for the analysis of time-series of one month's duration or less. It is this version of the method which is considered here.

Given a tidally modulated time-series $\zeta(t)$ the task is to derive a finite set of parameters which summarises the tidal information contained within the signal. In a Harmonic Method decomposition (Section 3.3.1), the time-series $\zeta(t)$ is written as the sum of N harmonic components, plus an error term $\varepsilon(t)$, as follows:

$$\zeta(t) = \sum_{j=1}^{N} \hat{H}_j \cos\left(\omega_j t - \hat{g}_j\right) + \varepsilon(t) \tag{3.13}$$

It is conceptually useful to consider the sum in equation (3.13) to be a sum of N basis functions:

$$\left\{\hat{b}_j(t) = \cos(\omega_j t)\right\} \tag{3.14}$$

In equation (3.13), the basis functions of equation (3.14) are time-shifted by the phase-lags $\{\hat{g}_j\}$ and scaled by the amplitudes $\{\hat{H}_j\}$ to create a synthetic tidal signal. Consequently, it is the set of $2N$ numbers $\{\hat{H}_j,\hat{g}_j\}$ which describes the tidal component of the time-series $\zeta(t)$.

In the Admiralty Method, the tidal signal is written in a manner analogous to equation (3.13), as a sum of components, plus an error term $\varepsilon(t)$:

$$\zeta(t) = \sum_{j=1}^{4} H_j F_j(t)\cos(\phi_j(t) - g_j) + \varepsilon(t) \tag{3.15}$$

The advantage of the Admiralty Method is that it has only four basis functions, defined by:

$$\left\{b_j(t) = F_j(t)\cos(\phi_j(t))\right\} \tag{3.16}$$

The functions $F_j(t)$ and $\varphi_j(t)$ are of great importance as they endow each Admiralty Method basis function with a degree of non-stationarity. The presence of $F_j(t)$ and $\varphi_j(t)$ in equation (3.16) shows that the Admiralty Method basis functions are not pure sinusoidal functions of time like the basis functions of the Harmonic Method (equation (3.14)). Rather, they might be described as *pseudo*-sinusoidal functions whose amplitude is modulated by the function $F_j(t)$ and whose phase is controlled by the function $\varphi_j(t)$. It is shown in Sections 2.3.1 – 2.3.3 that the tidal potential can be viewed as a non-stationary mixture of diurnal and semi-diurnal oscillations whose magnitude and phase vary according to the positions of the sun and moon. The functions $F_j(t)$ and $\varphi_j(t)$ are calculated from astronomical theory at any time t and impose the same astronomical non-stationarity on all functions which are linear combinations of Admiralty Method basis functions (equation (3.16)). Consequently, the use of the Admiralty Method basis functions ensures automatically that the expected astronomical non-stationarity is present in all tidal signals. Daily values of the functions $F_j(t)$ and $\varphi_j(t)$ – valid at 00h00 GMT - are published annually in Table VII of the Admiralty Tide Tables under the heading 'Tidal Angles and Factors'. The values at intermediate times are found by interpolation. The four Admiralty Method basis functions $\{b_1(t),b_2(t),b_3(t),b_4(t)\}$ are shown over a period of 20 days in 1994 in Figure 3.1

The parameters in an Admiralty Method decomposition are the amplitudes $\{H_1,\ldots,H_4\}$ and phase-lags $\{g_1,\ldots,g_4\}$ which must be applied to the four basis functions of equation (3.16). The parameter set $\{H_1,\ldots,H_4,g_1,\ldots,g_4\}$ describing the signal $\zeta(t)$ is usually chosen - according to a least-squares criterion - to minimise:

$$\int [\varepsilon(t)]^2\, dt = \int \left[\zeta(t) - \sum_{j=1}^{4} H_j f_j(t)\cos(\phi_j(t) - g_j) \right]^2 dt \tag{3.17}$$

It is important to note the limits which are placed on the amplitudes and phase-lags. The amplitudes must be non-negative

$$H_j \in [0,\infty) \tag{3.18}$$

while the phase-lags are constrained as follows:

$$g_j \in [0,2\pi) \tag{3.19}$$

This means that although the basis function $b_j(t)$ can be scaled arbitrarily (equation 3.18), it can only be shifted in phase by one cycle (equation 3.19).

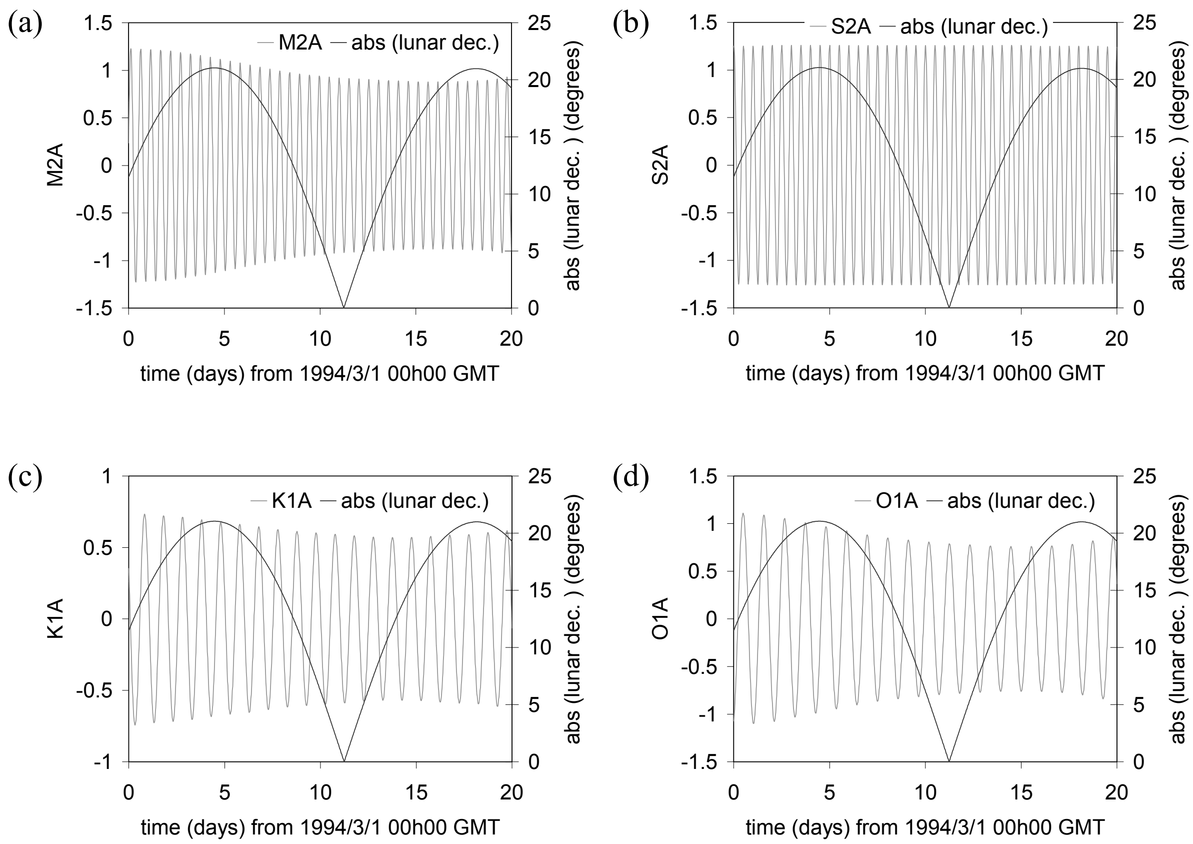

Figure 3.1: The four basis functions of the Admiralty Method (equation 2.16) as functions of time over a period of 20 days in 1994. The non-stationarity of these basis functions is controlled by the positions of the sun and the moon. The absolute lunar declination is shown for comparison. (a) The *M2A* basis function $b_1(t)$. (b) The *S2A* basis function $b_2(t)$. (c) The *K1A* basis function $b_3(t)$. (d) The *O1A* basis function $b_4(t)$.

It should be noted that in equation (3.19) the phase lag g_j is expressed in radians. The usual convention in the tidal literature (Doodson & Warburg, 1941) is to express angles in degrees. Both conventions are used in this dissertation.

3.3.2.1 Derivation of the Admiralty Method basis functions $\{b_j(t)\}$

The basis functions of the Admiralty Method (equation (3.16)) are derived from a subset of the basis functions of the Harmonic Method (equation (3.14)). Since the Admiralty Method is designed for use with time-series shorter than one month, it cannot extract harmonic constants equivalent to the long-period species of the Harmonic Method. The Admiralty Method is only capable of extracting information from the diurnal and semi-diurnal frequency bands. Under a Harmonic Method decomposition of the tidal potential, the most significant diurnal components are *K1* and *O1*, and the most significant semi-diurnal components are *M2* and *S2*. Each of these four principal components {*M2*,*S2*,*K1*,*O1*} is associated with a group of harmonic components with similar frequencies (Table 3.1). Consequently, the 20 most important Harmonic Method components are organised into four groups.

Admiralty Method component	Harmonic Method component	period (hours)	relative magnitude within group
M2A or $b_1(t)$	*2N2*	12.90	0.02534
	$\mu 2$	12.87	0.03057
	N2	12.66	0.19146
	v2	12.63	0.03636
	M2	12.42	1.00000
	$\lambda 2$	12.22	0.00738
	L2	12.19	0.02827
S2A or $b_2(t)$	*T2*	12.02	0.05861
	S2	12.00	1.00000
	K2	11.97	0.27215
K1A or $b_3(t)$	$\pi 1$	24.13	0.01939
	P1	24.06	0.33093
	K1	23.94	1.00000
	$\varphi 1$	23.81	0.01424
	J1	23.09	0.05591
O1A or $b_4(t)$	*2Q1*	28.02	0.02534
	$\sigma 1$	27.84	0.03056
	Q1	26.87	0.19146
	$\rho 1$	26.73	0.03637
	O1	25.83	1.00000

Table 3.1: The derivation and naming of the four Admiralty Method basis functions is based on four groups of Admiralty Method basis functions. Table from Admiralty Tidal Handbook, No.3.

The four Admiralty Method basis functions represent the overall, combined behaviour of the Harmonic Method components within each group in Table 3.1. For this reason, it is conceptually useful to refer to the four Admiralty Method basis functions $\{b_1,b_2,b_3,b_4\}$ by using the notation $\{M2A,S2A,K1A,O1A\}$. However, it is very important to stress the difference between the set $\{M2,S2,K1,O1\}$ of Harmonic Method basis functions and their Admiralty Method counterparts $\{M2A,S2A,K1A,O1A\}$. The former set consists of stationary, sinusoidal basis functions of the form $\{\hat{b}_j(t)=\cos(\omega t)\}$ (equation (3.14)). The latter set consists of the non-stationary, *pseudo*-sinusoidal basis functions $\{b_j(t)=F_j(t)\cos(\varphi_j(t))\}$.

In summary, the Admiralty Method of decomposing a tidal signal makes use of the fact that the astronomical motions of the sun and moon are well known and are expected to affect all tidal signals in a similar manner. This assumption allows tidal signals to be described by a set of 8 harmonic constants which can be extracted from reasonably short (~7 days) time-series

3.3.2.2 The information content of a tidal signal

Consider a (synthetic) tidal signal created using the 1200 components of the Tamura tidal potential catalogue (Tamura, 1987) by the ETGTAB code (Figure 3.2a). Admiralty Method harmonic constants can be extracted from the original data series (sampling interval Δt = 15 min, series length $N\Delta t$ = 7 days) according to the least-squares criterion of equation (3.17). The set of 8 Admiralty Method harmonic constants $\{H_1,\ldots,H_4,g_1,\ldots,g_4\}$ which are obtained can be expressed in complex number form by the set of 4 complex numbers $\{A_1,\ldots,A_4\}$, defined by:

$$\begin{cases}\mathrm{Re}(A_j)=H_j\cos(g_j)\\ \mathrm{Im}(A_j)=-H_j\sin(g_j)\end{cases} \tag{3.20}$$

The set of complex harmonic constants $\{A_1,\ldots,A_4\}$ can then be graphed in the complex plane and labelled according to the corresponding basis functions $\{M2A, S2A, K1A, O1A\}$ (Figures 3.2b,c). The magnitudes $\{H_1,\ldots,H_4\}$ of the components are represented graphically by the length of the phasors. The phase lags $\{g_1,\ldots,g_4\}$ are represented by the angles, measured positive *anticlockwise*, between the real axis and the phasor. It is important to stress the physical meaning of these phase lags. They represent the phase lag of each component relative to the corresponding component of the tidal potential on the Greenwich meridian. The time-series analysed in Figure 3.2 is an estimate of the tidal potential on the 30°W meridian. Semi-diurnal components of the tidal potential travel westwards round the globe at 30° per hour, while diurnal components of the tidal potential travel westwards at 15° per hour. This explains why the semi-diurnal components have phase lags of 60° (Figure 3.2b) while the diurnal components have phase lags of 30° (Figure 3.2c). This illustrates an important feature of Admiralty Method harmonic constants.

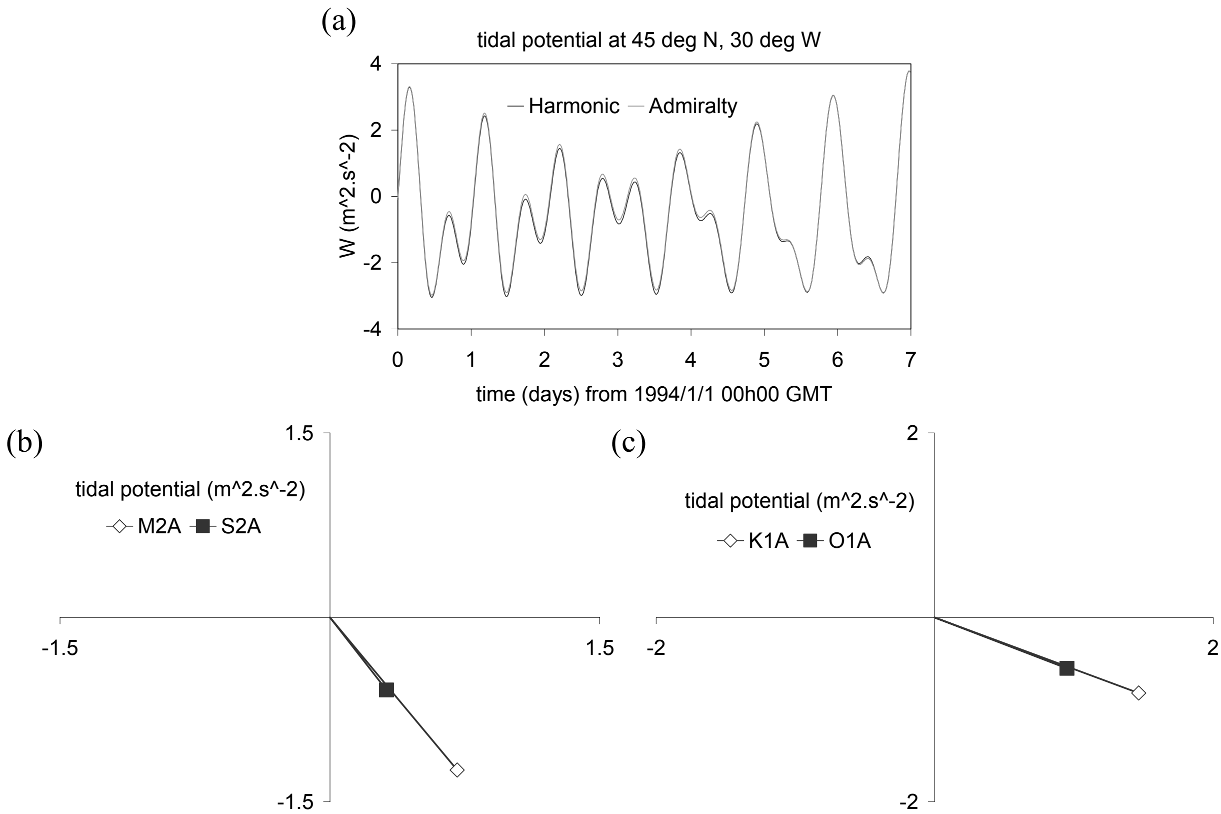

Figure 3.2: The extraction of Admiralty Method harmonic constants from a tidal potential generated using 1200 components of a Harmonic Method decomposition. (a) The tidal potentials generated by the Harmonic Method and the Admiralty Method in the time domain. (b) The semi-diurnal Admiralty Method harmonic constants in the complex plane. (c) The diurnal Admiralty Method harmonic constants in the complex plane. Original time-series generated using the ETGTAB code.

For the tidal potential and the solid tide, the Admiralty Method harmonic constants have distinctive phase lags. The phase lags for the diurnal components (g_3 and g_4) equal the westward longitude, while those for semi-diurnal components (g_1 and g_2) equal twice the westward longitude. As a consequence, it is a striking feature of the diagrams that the two semi-diurnal components have the same phase lag (i.e. $g_1 = g_2$ in Figure 3.2b), and the two diurnal components have the same phase lag (i.e. $g_3 = g_4$ in Figure 3.2c). In general these phase lag equalities do not occur for the ocean tide (Figure 3.3). Consequently, if plots of the Admiralty Method harmonic constants for a time-series show that $g_1 \approx g_2$ and $g_3 \approx g_4$ then this constitutes evidence that the time-series is more closely correlated to the solid tide than the ocean tide.

It is of interest to compare the original time-series generated by the Harmonic Method harmonic constants $\{\hat{H}_1,\ldots,\hat{H}_{1200},\hat{g}_1,\ldots,\hat{g}_{1200}\}$ with the time-series generated by the Admiralty Method harmonic constants $\{H_1,\ldots,H_4,g_1,\ldots,g_4\}$ (Figure 3.2a). The time-series generated by the 8 parameters of the Admiralty Method does not match the original perfectly, but the fit is nonetheless remarkably good. It is tempting to conclude that a parametric description with 2400 degrees of freedom has been replaced by a parametric description with 8 degrees of freedom, but this is not so. The harmonic decomposition of a *tidal* signal into N Harmonic Method components does not have $2N$ degrees of freedom, because certain magnitude and phase relationships always hold between some of the components for all tidal signals. This should not be too surprising. The working definition of a tidal signal is 'a time-series whose fundamental physical cause is the tidal potential'. It is shown in Section 2.3 that the harmonic decomposition of the tidal potential is derived from the simpler long-period/diurnal/semi-diurnal decomposition. The essence of the Admiralty Method is the assumption that the astronomical motions of the sun and moon impose the same constraints on all tidal signals. It must be stressed that this is only a hypothesis, but it is one that has been shown to be justified in the many applications of the Admiralty Method to the description of ocean tides in the world's ports.

Figure 3.3 shows the Admiralty Method harmonic constants extracted from an estimate of the ocean tide at 45°N, 30°W. The estimate of the ocean tide is made using the CSR code which is based on the results of satellite altimetry (Schrama & Ray, 1994). Figure 3.3a shows that the 8 Admiralty Method harmonic constants reproduce the original signal remarkably well. Figures 3.3b,c demonstrate that the phase lags for the ocean tide are not generally as simple as those for the solid tide (Figures 3.2b,c). For the ocean tide, it is not generally true that $g_1 \approx g_2$ or $g_3 \approx g_4$.

3.3.2.3 Astronomical and intrinsic non-stationarity

By definition, every tidal signal is the output from some physical system whose input is the tidal potential. The tidal potential exhibits astronomical non-stationarity due to the motions of the sun and moon (Sections 2.3.1 – 2.3.3). Consequently, if the physical mechanism linking the tidal potential to the observable time-series does not change over time, it is to be expected that the observable time-series will exhibit the same astronomical non-stationarity.

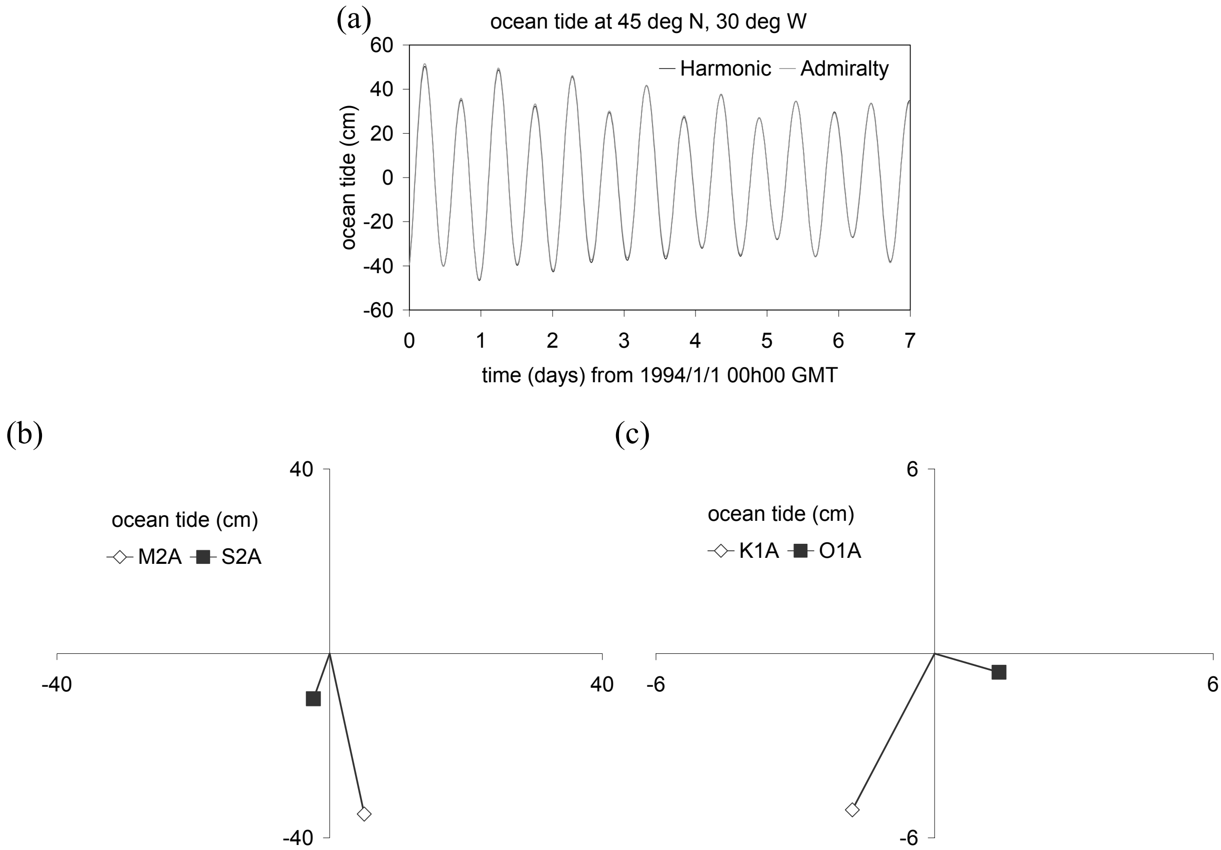

Figure 3.3: The extraction of Admiralty Method harmonic constants from an estimate of the ocean tide generated from a Harmonic Method decomposition. (a) The ocean tide signals generated by the Harmonic Method and the Admiralty Method in the time domain. (b) The semi-diurnal Admiralty Method harmonic constants in the complex plane. (c) The diurnal Admiralty Method harmonic constants in the complex plane. Harmonic Method ocean tide time-series generated using the CSR code.

Astronomical non-stationarity is easily predicted from astronomical formulae and is explicitly accommodated in the basis functions of the Admiralty Method (equation (3.16)). As an example, consider the parametric analysis of a tidal signal of two weeks duration. The time-series could be split into disjoint sections, each a week long, which are to be analysed separately. The set of $2N$ parameters extracted in a Harmonic Method decomposition $\{\hat{H}_1,\ldots,\hat{H}_N , \hat{g}_1,\ldots, \hat{g}_N\}$ changes from one week to the next because of the astronomical non-stationarity induced by declinational changes (Figure 2.3). However, an Admiralty Method decomposition yields the *same* set of 8 parameters $\{H_1,\ldots,H_4,g_1,\ldots,g_4\}$ for both weeks. For this reason, the Admiralty Method is immune to astronomical non-stationarity.

However, there is a second type of non-stationarity which might be present in a time-series, to which the Admiralty Method is not immune. This can be labelled intrinsic non-stationarity, and occurs if the physical mechanism linking the tidal potential to the observable time-series is *itself* subject to change over the observation interval. For example, if the observable time-series is the temperature of hydrothermal effluent, a change in seafloor hydrology could lead to a change in the spectral character of the temperature signal. Other examples would be drift caused by an erratic instrumental power supply, or by hydrothermal precipitation on seafloor sensors.

If a time-series dataset is split into subsections, each can be analysed separately using the Admiralty Method. If the harmonic constants obtained from these analyses change significantly between the subsections, it can be surmised that there is an intrinsic non-stationarity in the time-series.

3.3.2.4 Robustification of the Admiralty Method

Since hydrothermal time-series are noisy, it is of interest to consider how the estimation of Admiralty Method harmonic constants could be made more robust. In general a data section of ~7 days length is sufficient to yield accurate estimates of the Admiralty Method harmonic constants. Consequently, a time-series longer than ~7 days can be split into shorter subsections for which the Admiralty Method harmonic constants are estimated separately. Chave *et al.* (1987) suggest a manner in which power spectrum estimates from the separate subsections can be combined in a robust manner to yield an overall estimate. Their method is here adapted to the case of estimating Admiralty Method harmonic constants. As an example, consider the estimation of the (complex) harmonic constant A_1 associated with the *M2A* component (equation (3.20)). Suppose that the original time-series is split into M

subsections, and that the estimate of A_1 obtained from the m^{th} subsection is labelled $A_{1,m}$. The technique of Huber weighting can be used to derive an overall estimate from the set $\{A_{1,m}\}$ of estimates obtained from the separate subsections. The real and imaginary parts of A_1 are estimated in the same manner, and so estimation of the real part only is described here. Firstly, the median μ and interquartile distance σ of the set $\{\mathrm{Re}(A_{1,m})\}$ are calculated. (The median and interquartile distance are robust equivalents of the mean and standard deviation). The scaled distance of the m^{th} estimate from the median is given by:

$$x_m = \frac{\mathrm{Re}(A_{1,m}) - \mu}{\sigma} \tag{3.21}$$

The purpose of the robustification is to downweight the influence on the overall estimate of outlier subsection estimates for which $|x_m|$ is large. Consequently a set of weights $\{w_m\}$ is defined by:

$$\begin{aligned} w_m &= 1 & \left|x_m^2\right| \le k \\ w_m &= \sqrt{\frac{2k}{\left|x_m^2\right|} - \frac{k^2}{x_m^2}} & \left|x_m^2\right| > k \end{aligned} \tag{3.22}$$

Chave *et al.* (1987) suggest that k = 1.5 gives good results. The overall estimate of the real part of the harmonic constant is then given by:

$$\mathrm{Re}(A_1) = \sum_{m=1}^{M} w_m \,\mathrm{Re}(A_{1,m}) \tag{3.23}$$

Furthermore, this robust section averaging technique can be used to estimate jack-knife confidence limits for the estimate $\mathrm{Re}(A_1)$ on the assumption that the real and imaginary parts of A_1 are normally distributed.

Robust estimates of Admiralty Method harmonic constants are made using this technique for the data presented in Chapter 4. For sufficiently long time-series, robust section averaging is used and 95% confidence limits can be placed on the harmonic constants in the complex plane.

3.3.3 Drift Removal by Bayesian methods

In many cases, the time-series obtained from a hydrothermal system appears to be the sum of three separate signals - a tidal signal, a background drift signal and residual high frequency noise. For example, the temperature measurements obtained at TAG on the Mid-Atlantic Ridge (Schultz *et al.,* 1996) show tidal variations of ~0.5 °C superimposed on a slow drift whose amplitude is ~10 °C. (Figure 4.11). It is desirable to have a well-defined method of decomposing a time-series into a drift signal, a parametrised tidal signal and residual error. An interesting method, whose heuristic justification relies on techniques from Bayesian statistics, is proposed by Tamura *et al.* (1991). In their original method, the tidal signal is represented parametrically by a Harmonic Method decomposition. In order to apply their method to short hydrothermal time-series, it is adapted here to incorporate an Admiralty Method decomposition. Accordingly, what follows is a summary of the theory presented by Tamura *et al.* (1991) and Akaike (1980), modified where necessary to incorporate the Admiralty Method.

Suppose that the time-series to be analysed consists of the data $\{x_1,\ldots,x_N\}$, collected at times $\{t_1,\ldots,t_N\}$. The objective is then to estimate a drift signal $\{d_1,\ldots,d_N\}$ and a set of Admiralty Method harmonic constants $\{H_1,\ldots,H_4,g_1,\ldots,g_4\}$ which generate a tidal signal $\{\zeta_1,\ldots,\zeta_N\}$ according to:

$$\zeta_i = \sum_{j=1}^{4} H_j F_j(t_i)\cos\left(\phi_j(t_i) - g_j\right) \tag{3.24}$$

Consequently, the original time-series $\{x_1,\ldots,x_N\}$ can be decomposed into a tidal signal (ζ_i), a drift signal (d_i), and a residual error signal (ε_i) as follows:

$$x_i = \zeta_i + d_i + \varepsilon_i \tag{3.25}$$

It is reasonable to demand two properties of the extended parameter set $\{d_1,\ldots,d_N,H_1,\ldots,H_4,g_1,\ldots,g_4\}$ which is to be estimated from the data $\{x_1,\ldots,x_N\}$. Firstly, the drift $\{d_i\}$ should be smooth, and secondly the resulting error $\{\varepsilon_i\}$ should be small. The smoothness of the drift is quantified by the expression:

$$\sum_{i=3}^{N} \left| d_i - 2d_{i-1} + d_{i-2} \right|^2 \tag{3.26}$$

A small value of this quantity is desirable to ensure that the drift is smooth.

The size of the error can be quantified by the expression:

$$\sum_{i=1}^{N} \left|\varepsilon_i\right|^2 = \sum_{i=1}^{N} \left| x_i - d_i - \zeta_i \right|^2 \tag{3.27}$$

A small value of this quantity is desirable to ensure that the error is small.

In order to achieve the simultaneous requirements of smooth drift and small error, the extended parameter set $\{d_1,\ldots,d_N,H_1,\ldots,H_4,g_1,\ldots,g_4\}$ is obtained by finding the values which minimise the penalised least-squares expression:

$$\sum_{i=1}^{N} \left| x_i - d_i - \zeta_i \right|^2 + \nu^2 \sum_{i=3}^{N} \left| d_i - 2d_{i-1} + d_{i-2} \right|^2 \tag{3.28}$$

Equation (3.28) shows that the relative importance of smooth drift and small error is controlled by a hyperparameter ν whose value must be chosen before the least-squares solution is found. (Later, an analogy with Bayesian statistics will provide an explicit criterion for selecting the optimal value of v.)

The penalised least squares minimisation of equation (3.28) can be written in matrix form by defining an extended data vector $\underline{X}$ (of dimension $2N$-2), and an extended parameter vector $\underline{P}$ (of dimension N+8) as follows:

$$\underline{X} = \begin{pmatrix} x_1 \\ \vdots \\ x_N \\ 0 \\ \vdots \\ 0 \end{pmatrix} \quad , \quad \underline{P} = \begin{pmatrix} d_1 \\ \vdots \\ d_N \\ H_1 \cos(g_1) \\ H_1 \sin(g_1) \\ \vdots \\ H_4 \cos(g_4) \\ H_4 \sin(g_4) \end{pmatrix} = \begin{pmatrix} \underline{d} \\ \underline{\theta} \end{pmatrix} \tag{3.29}$$

For later convenience the vector $\underline{P}$ is split into two parts: $\underline{d}=(d_1,\ldots,d_N)$ and $\underline{\theta}=(H_1\cos(g_1),\ldots,H_4\sin(g_4))$. It should be noted that the set of Admiralty Method harmonic constants $\{H_1,\ldots,H_4,g_1,\ldots,g_4\}$ is easily recovered from the elements of $\underline{\theta}$. Therefore the vector $\underline{d}$ represents the drift signal while the vector $\underline{\theta}$ represents the parameters which generate the estimated tidal part of the signal.

The N-by-8 matrix T is defined by:

$$T = \begin{bmatrix} f_1(t_1)\cos(\phi_1(t_1)) & f_1(t_1)\sin(\phi_1(t_1)) & \cdots & f_4(t_1)\cos(\phi_4(t_1)) & f_4(t_1)\sin(\phi_4(t_1)) \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ f_1(t_N)\cos(\phi_1(t_N)) & f_1(t_N)\sin(\phi_1(t_N)) & \cdots & f_4(t_N)\cos(\phi_4(t_N)) & f_4(t_N)\sin(\phi_4(t_N)) \end{bmatrix} \tag{3.30}$$

The (N-2)-by-N matrix D is defined by:

$$D = \begin{bmatrix} 1 & -2 & 1 & & & 0 \\ & 1 & -2 & 1 & & \\ & & \ddots & \ddots & \ddots & \\ 0 & & & 1 & -2 & 1 \end{bmatrix} \tag{3.31}$$

The ($2N$-2)-by-(N+8) matrix A is defined by:

$$A = \begin{bmatrix} I_N & T \\ vD & O \end{bmatrix} \tag{3.32}$$

where I_N is the N-by-N identity matrix and O is an (N-2)-by-8 matrix of zeroes.

It then follows that the solution $\underline{P}^*=(\underline{d}^*,\underline{\theta}^*)$ of the penalised least squares minimisation of equation (3.28) is the value of $\underline{P}$ which minimises the matrix norm:

$$\left\| A\underline{P} - \underline{X} \right\|^2 \tag{3.33}$$

For a given value of the hyperparameter *v*, the solution $\underline{P}^*$ which minimises this norm can be found explicitly by constructing and solving the appropriate normal equations (Press *et al.,* 1986). The solution is:

$$\underline{P}^* = \left(A^T A\right)^{-1} A^T \underline{X} \qquad (3.34)$$

A criterion for deciding the optimal value of the hyperparameter *v* is now constructed by analogy with Bayesian statistics.

Consider the idealised case where the drift can be represented stochastically by an integrated random walk. Suppose that the $\{E_I\}$ is a set of random variables, drawn independently from a zero-mean normal distribution of variance σ^2/v^2. The drift is then given by:

$$d_i = 2d_{i-1} + d_{i-2} + E_i \qquad (3.35)$$

In the language of Bayesian statistics, equation (3.32) defines a prior distribution for the drift signal, given by:

$$\pi\left(\underline{d}\middle|\sigma^2, v\right) = \left(\frac{1}{2\pi\sigma^2}\right)^{N/2} \left(\det\left(v^2 D^T D\right)\right) \exp\left(\frac{-v^2}{2\sigma^2} \underline{d}^T \left(D^T D\right)\underline{d}\right) \qquad (3.36)$$

It is now supposed that the error terms $\{\varepsilon_i\}$ of equation (3.25) are drawn independently from a zero-mean normal distribution of variance σ^2. It follows that the data distribution is proportional to:

$$f\left(\underline{x}\middle|\underline{d}, \sigma^2, v, \underline{\theta}\right) = \exp\left(\frac{-1}{2\sigma^2}\left\|\underline{x} - \underline{d} - T\underline{\theta}\right\|^2\right) \qquad (3.37)$$

From equations (3.36) and (3.37), the posterior distribution for the data $\underline{x}$ is proportional to:

$$f\left(\underline{x}\middle|\underline{d}, \sigma^2, v, \underline{\theta}\right)\pi\left(\underline{d}\middle|\sigma^2, v\right) = \left(\frac{1}{2\pi\sigma^2}\right)^{N/2} \left(\det\left(v^2 D^T D\right)\right) \exp\left(\frac{-1}{2\sigma^2}\left(\left\|\underline{x} - \underline{d} - T\underline{\theta}\right\|^2 + \left\|vD\underline{d}\right\|^2\right)\right) \qquad (3.38)$$

Given the stated assumptions, it is reasonable to define estimates of σ^2 and $\underline{\theta}$ which maximise the posterior likelihood (equation (3.38)) to be optimal. It can be shown that equation (3.38) is maximised with respect to $\underline{\theta}$ and σ^2 by setting $\underline{\theta} = \underline{\theta}^*$ (the solution from equation (3.34)) and by setting $\sigma^2 = \sigma_B^2$, where σ_B^2 is defined by:

$$\sigma_B^2 = \frac{1}{N}\left\|A\underline{P}^* - \underline{X}\right\|^2 \qquad (3.39)$$

Consequently:

$$f\left(\underline{x}\middle|\underline{d}, \sigma_B^2, v, \underline{\theta}^*\right)\pi\left(\underline{d}\middle|\sigma_B^2, v\right) = \left(\frac{1}{2\pi\sigma_B^2}\right)^{N/2} \left(\det\left(v^2 D^T D\right)\right)^{1/2} \exp\left(\frac{-N}{2}\right) \exp\left(\frac{-1}{2\sigma_B^2}\left(\underline{d} - \underline{d}^*\right)^T \left(I + v^2 D^T D\right)\left(\underline{d} - \underline{d}^*\right)\right) \qquad (3.40)$$

Finally, the posterior marginal likelihood $L(v)$ is defined by integrating equation (3.40) with respect to $\underline{d}$:

$$L(v) = \int f\left(\underline{x}\middle|\underline{d}, \sigma_B^2, v, \underline{\theta}^*\right)\pi\left(\underline{d}\middle|\sigma_B^2, v\right)d\underline{d} \tag{3.41}$$

It can be shown that:

$$L(v) = \left(\frac{1}{2\pi\sigma_B^2}\right)^{N/2}\left(\det\left(v^2 D^T D\right)\right)^{1/2}\left(\det\left(I + v^2 D^T D\right)\right)^{-1/2}\exp\left(\frac{-N}{2}\right) \tag{3.42}$$

The optimal value of the hyperparameter v is now to defined to be the value which *maximises* the posterior marginal likelihood $L(v)$ of equation (3.42). Equivalently, the optimal value of v is the value which *minimises* the quantity $-2\log(L(v))$, defined by:

$$-2\log(L(v)) = N + N\log\left(2\pi\sigma_B^2\right) + \log\left(\det\left(I + v^2 D^T D\right)\right) - \log\left(\det\left(v^2 D^T D\right)\right) \tag{3.43}$$

When comparing different values of the hyperparameter v in a particular problem, equation (3.31) shows that the matrix D remains unchanged. It follows that $\log(\det(v^2 D^T D)) = N\log(v^2) + constant$ in equation (3.43). Consequently, the optimal value of v in a particular problem is defined to be the value which minimises the quantity:

$$ABIC(v) = N + N\log\left(2\pi\sigma_B^2\right) + \log\left(\det\left(I + v^2 D^T D\right)\right) - N\log\left(v^2\right) \tag{3.44}$$

The quantity $ABIC(v)$ is known as 'Akaike's Bayesian Information Criterion' (Tamura *et al.,* 1991).

In practice, a computer can be programmed to vary the value of v, obtaining a solution from equation (3.34) each time, until the least value of $ABIC(v)$ is found. The final estimate of the drift $\underline{d}$ and the Admiralty Method harmonic constants $\underline{\theta}$ are the solutions obtained from equation (3.34) when v minimises $ABIC(v)$.

3.3.4 The HYBRID code – a parametric method for analysing tidal data

The ideas of the previous sections are combined in a new method of parametric analysis, which is specifically designed to extract tidal information from hydrothermal time-series. This method is implemented by the new HYBRID code which is used for parametric data analysis in Chapter 4.

The HYBRID code is designed to estimate Admiralty Method harmonic constants from hydrothermal time-series. For sufficiently long time-series, the time-series is split into separate subsections which are analysed independently. For the analysis of tidal signals, a subsection length of 1 week is often used as this is sufficiently long to allow accurate estimates to be made of the four complex harmonic constants $\{A_1, A_2, A_3, A_4\}$ (equation

(3.20)). In order to make clear their meaning, these harmonic constants are labelled according to the Admiralty Method basis function to which they correspond {*M2A*, *S2A*, *K1A*, *O1A*}. The HYBRID code analyses the individual data subsections in one of two ways, with the choice being made by the user. The first method is labelled 'mean removal', and involves the subtraction of the mean value of the subsection before the harmonic constants are estimated. This technique is simple and works well when the time-series does not suffer from excessive drift. The second method is labelled 'Bayesian drift removal'. This technique (Section 3.3.3) is much more computationally intensive, and is designed to decompose the time-series into (1) a smooth drift signal, (2) a tidal signal generated by the Admiralty Method harmonic constants, and (3) residual noise. Use of this method in Chapter 4 suggests that it works very well when the time-series follows closely the parametric form assumed in Section 3.3.3. However, for very noisy time-series, better results are obtained by using 'mean removal' in each subsection.

Finally, the harmonic constants estimated for each subsection are combined using Huber weighting to yield an overall robust estimate for the Admiralty Method harmonic constants. 95% jackknife confidence limits are provided for these estimates on the assumption that the real and imaginary parts of the harmonic constants are independently normally distributed.

3.4 Conclusions

The aim of this chapter is to explain the techniques for analysing hydrothermal time-series which are implemented in Chapter 4.

The first test of a time-series is to examine it in the time domain. In many cases diurnal and semi-diurnal periodicities are immediately apparent and it can be suspected that the time-series is tidally modulated. It is of interest to decide whether the tidal modulations are caused by the ocean tide or the solid tide and so it may be useful to graph these 'input signals' alongside the data signal in the time domain.

The next stage in the investigation of a time-series is to obtain a non-parametric estimate of its frequency spectrum. In the past, many authors have used the periodogram to estimate the power spectrum of seafloor time-series in order to establish whether it is tidally modulated. It is argued here, in Section 3.2.1, that the periodogram is an extremely bad estimator of the true power spectrum and should not be used. It is suggested instead that the multiple window method of Thomson (1982) should be used in the non-parametric analysis of tidal signals in seafloor hydrothermal systems. Consequently, the MWPS code (A. Chave, *pers. comm.*, 1999) is used in Chapter 4 to obtain spectral estimates for hydrothermal time-series. This code incorporates an explicit method for the identification of significant line components in the spectrum. If significant line components are identified at the known tidal frequencies (in particular at ~1.93 cpd corresponding to the *M2* component) then it can be concluded with confidence that the time-series is tidally modulated.

In the cases where non-parametric analysis reveals tidal modulation of the time-series, a parametric analysis can be attempted in order to summarise the nature of the tidal signal. It is suggested that the Admiralty Method is the optimal parametric description of tidal signals since it requires only 8 harmonic constants to describe an arbitrary tidal signal. Furthermore, the harmonic constants for any time-series can be extracted from a fairly short data section of ~7 days. The Admiralty Method is able to describe tidal signals using so few parameters because much of the information content of a tidal signal is due to the well known astronomical motions of the sun and moon and is the same for *all* tidal signals.

A new computer code – HYBRID – for the parametric analysis of tidally modulated signals is described (Section 3.3.4). The HYBRID code is designed to estimate the Admiralty Method harmonic constants of a tidally modulated signal. It incorporates a Bayesian technique for drift removal (Section 3.3.3) and robust section averaging (Section 3.3.2.4). The HYBRID code is used extensively in Chapter 4 to quantify the tidal part of a number of seafloor time-series.

The (complex) harmonic constants extracted from a time-series using the Admiralty Method can be plotted in the complex plane for ease of interpretation. In Chapter 4, the harmonic constants of a seafloor time-series are compared with the harmonic constants of the solid and ocean tides in order to establish whether the ocean tide or the solid tide is responsible for the tidal modulation of the time-series.