Data Processing and Analysis

Rick Aster and Brian Borchers

September 24, 2008

Energy and Power Spectra

It is frequently useful to study the distribution and power of a signal in the frequency domain. The simplest such measure is the energy spectral density, which is just the amplitude of the spectrum squared

$$|\Phi(f)|^2 = \Phi(f)\Phi^*(f) \ . \tag{1}$$

Using the convolution theorem, and noting that phase conjugation in the frequency domain corresponds to time reversal in the time domain, this can be recognized as the Fourier transform of the autocorrelation

$$\Phi(f)\Phi^*(f) = F[\phi(t) * \phi^*(-t)] = F[\phi(t) \text{ cor } \phi^*(t)] \ . \tag{2}$$

We can thus observe that function which has a sharp and narrow autocorrelation function will have a broad energy spectral density, while a function which has a broad autocorrelation function will have a narrow energy spectral density. This can perhaps be understood better by considering what is in fact required of a time-domain function for it to have narrow (in the limit, delta-like) autocorrelation function; the function must change rapidly, so that it does not resemble itself very much for a small shift from zero lag. For a function to change rapidly, it must have high frequency energy in its spectrum. Note that, because the units of a spectrum are $u \cdot \mathrm{s} = u/\mathrm{Hz}$, the units of (1) are u^2/Hz^2, where u denotes the physical units of $\phi(t)$ (e.g., Volts, Amperes, meters/s, etc.).

Many interesting signals, such as those arising an incessant excitation, are, practically speaking, unbounded in time or *continuous* (as opposed to limited in time, or *transient*). If the statistical behavior of the signal (we will look at statistical aspects much more later on in the class) doesn't change with time, so that the spectral and other properties of the signal are time-invariant, it is referred to as *stationary*. Examples where signals can often be considered to be stationary include seismic, thermal, or electromagnetic noise, tides, winds, temperatures, and currents. Some signals of interest exhibit strong periodicities (tides, for example) because they are associated with astronomical or other periodic effects. Because of their continuous nature, such signals have infinite total energy

$$E_T = \lim_{T->\infty} \int_{-T/2}^{T/2} |\phi(t)|^2 \ dt = \ \infty, \tag{3}$$

so that the Fourier transform of the autocorrelation (2) won't converge. The frequency content of such signals may instead be examined as a *power spectral density*, or simply *PSD*.

The signal power averaged over some interval T is simply the energy (3) normalized by the length of the observation

$$P_T = \frac{1}{T}\int_{-T/2}^{T/2} |\phi(t)|^2 \, dt = \frac{1}{T}\int_{-\infty}^{\infty} |\phi(t)\cdot\Pi(t/T)|^2 \, dt \; . \tag{4}$$

As the observation interval T becomes long, this converges to the true signal power

$$P = \lim_{T\to\infty} P_T \; . \tag{5}$$

The PSD is defined as

$$PSD(\phi(t)) = \lim_{T\to\infty} \frac{1}{T}\Phi_T(f)\cdot\Phi_T^*(f) \tag{6}$$

where

$$\Phi_T(f) = F[\phi(t)\cdot\Pi(t/T)] \; . \tag{7}$$

Note that (6) has units of the power spectral density (6) are u^2/Hz. Further note that that PSDs will be real, symmetric functions over f for the common case where $\phi(t)$ real (and thus has a Hermitian spectrum). For this reason, power spectra of real functions are typically plotted only for positive frequencies. Because we can never do calculations on an infinite-length signal, all PSDs in practice are estimates of P.

The simplest (but not best!) way to estimate a PSD is to simply truncate the data with a T-length rectangular time window extending in time from $-T/2$ to $T/2$. This estimate has a long history, and it is sometimes referred to as a *periodogram*. To understand the relationship between the periodogram estimate and the true PSD (6), note that for a rectangular window of width T and a real-valued time series (which has a Hermitian spectrum)

$$PSD_{periodogram} = \frac{1}{T}|\Phi_T(f)|^2 = \frac{1}{T}|F[\phi(t)\Pi(t/T)]|^2 \tag{8}$$

Using the convolution theorem, this gives

$$PSD_{periodogram} = \frac{1}{T}|\Phi(f) * \mathrm{sinc}(Tf)|^2 \tag{9}$$

where, recall, the Fourier transform of $\Pi(t/T)$ is

$$\mathrm{sinc}(Tf) = \frac{\sin \pi T f}{\pi T f} \; . \tag{10}$$

Thus, what we actually obtain in a periodogram estimate is the true PSD of the process convolved in the frequency domain with the $\mathrm{sinc}(Tf)$ function. Figure 1 shows a periodogram estimate for a sinusoidal process. The underlying

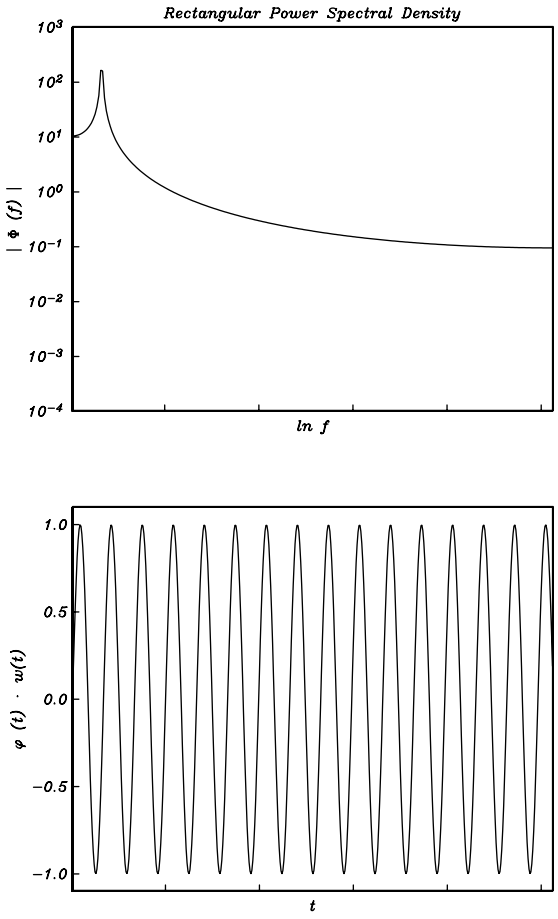

Figure 1: A Periodogram Estimate of a Pure Sinusoidal Process

process has a delta function spectrum, (9) results in a broader peak in the PSD estimate obtained using (8). The spearing effect of the convolution is an indication of limited *spectral resolution*.

The loss of resolution caused by the convolution in (9) can be highly undesirable. As convolution is essentially a smoothing operation, our windowed estimate in (9) is a blurred picture of the true spectrum. In the Periodogram case, this blurring takes the form of convolution with a sinc function because we chose an (abrupt) boxcar data truncation on the $\pm T/2$ interval. Because of its slow ($(Tf)^{-1}$) fall-off and oscillatory behavior, the periodogram should never be used in practice except for quick and dirty estimates. The smearing of spectral resolution due to the convolution of the true spectrum with the Fourier transform of the windowing function is called *spectral leakage*, as the frequency domain convolution in (9) causes power from surrounding frequencies to "leak" into the estimate at any particular frequency. In its simplest form, spectral leakage in the periodogram will make the PSD estimate for a function that is really a sinusoid of frequency f have the appearance of sinc functions centered on the true frequencies ($\pm f$) of the continuous signal, rather than the true delta functions.

Spectral leakage can be reduced by increasing T, so that the Fourier transform of the windowing function becomes reciprocally (by a factor of $1/T$) narrower. However, for statistical reasons involving the variance of the estimate that we will not elaborate on here, this is still a poor way to estimate the PSD. A better way to reduce spectral leakage, at the cost of eliminating the statistical contributions of data near the endpoints of the data series, is to window with a smoother time function than the boxcar that has a Fourier transform that is more delta-like by some measure. For example, consider the *Bartlett* or *Parzen* window

$$\Lambda(2t/T) = 4/T^2 \left(\Pi(2t/T) * \Pi(2t/T)\right) \tag{11}$$

which is a unit height triangle function spanning the interval $-T/2$ to $T/2$. $\Lambda(2t/T)$ is easily seen by the convolution theorem to have a Fourier transform given by the sinc function squared

$$F[\Lambda(2t/T)] = 4/T^2 F[\Pi(2t/T) * \Pi(2t/T)] = \mathrm{sinc}^2(fT/2) \tag{12}$$

which falls off asymptotically as $(Tf)^{-2}$ and is positive everywhere (although it is still oscillatory; Figure 2).

The formulation of various data windows such as the Parzen window has historically formed a veritable cottage industry in signal processing, and numerous function are in usage (may of which can be readily generated using various MATLAB functions in the signal processing toolbox). The general tradeoff in window selection is between the width of the main lobe of the leakage function and the rate of falloff away from the center frequency. A few examples of commonly used windows and their corresponding spectral leakage properties when they are applied to a true sinusoidal signal (which has a delta function PSD), are shown in the following figures.

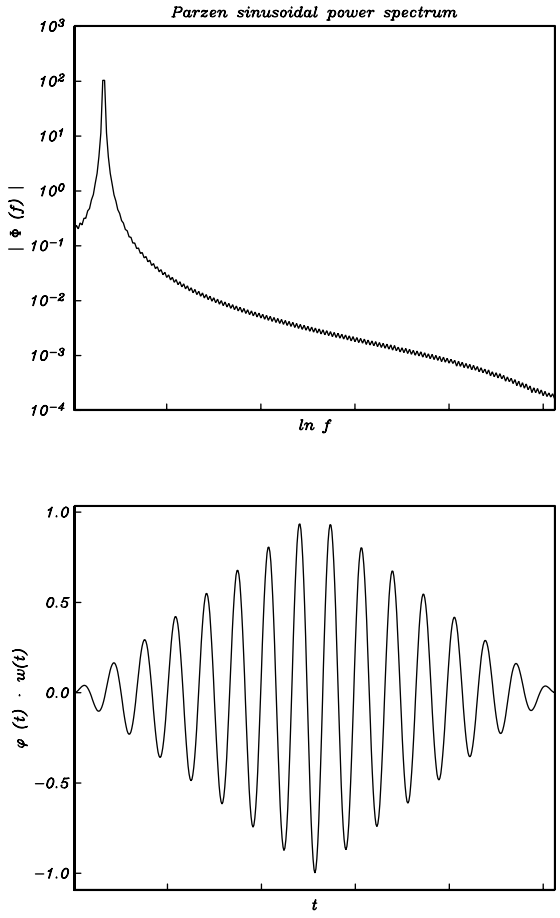

Figure 2: A Bartlett or Parzen window estimate of a pure sinusoidal process spectrum.

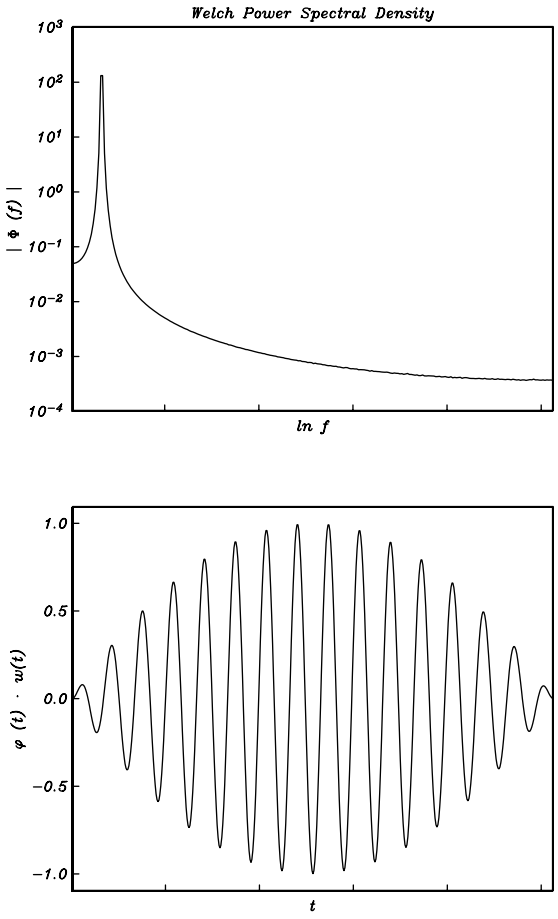

Figure 3: A Welch window estimate of a pure sinusoidal process spectrum.

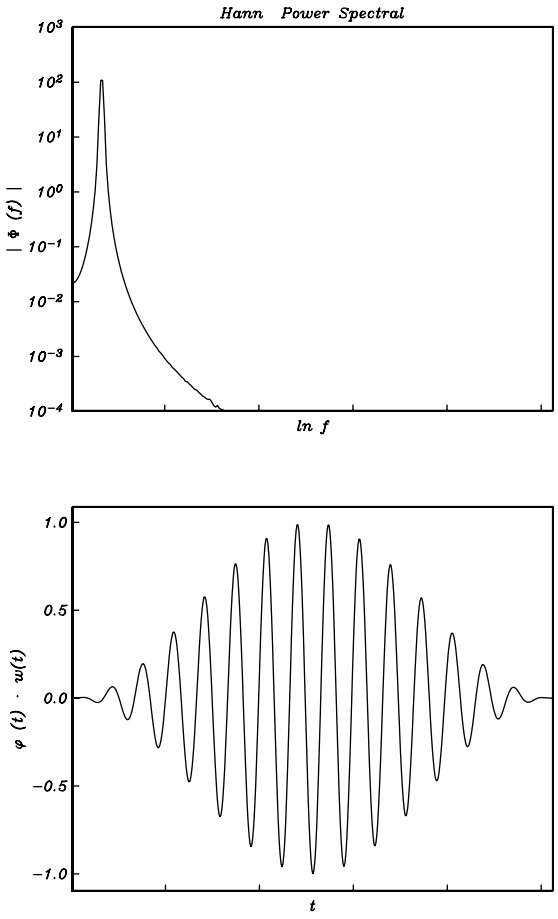

Figure 4: A Hann window estimate of a pure sinusoidal process spectrum.

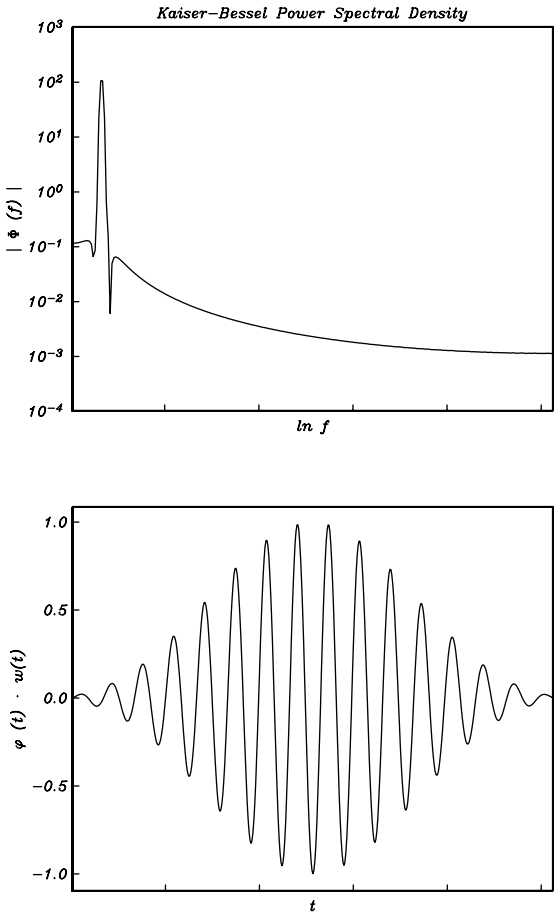

Figure 5: A Kaiser-Bessel window estimate of a pure sinusoidal process.

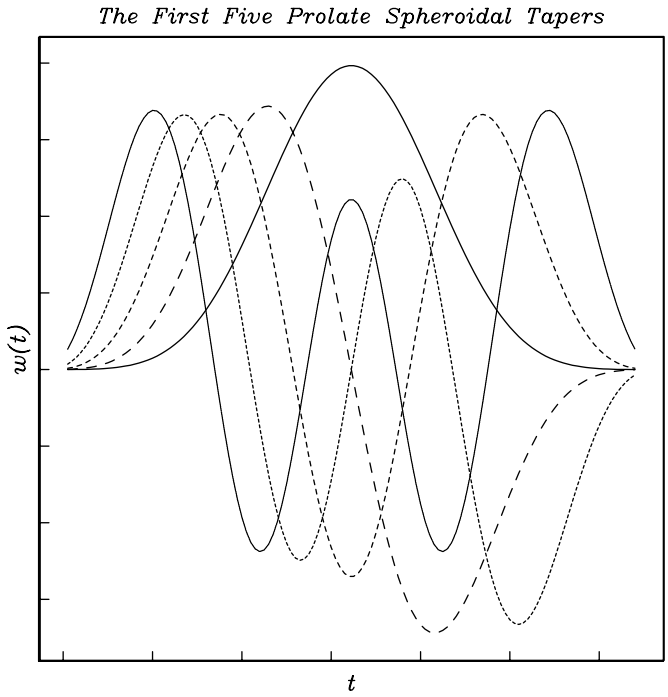

Figure 6: Prolate Spheroidal Taper Functions ($0 \leq k \leq 4; NW = 4$).

An interesting issue in spectral estimation that arises from the use of windows is the "throwing out" of data resulting from tapering near the data segment endpoints. The result is that we are downweighting information and thus increasing the statistical uncertainty of the PSD estimate. For long, stationary time series, one simple and widely-applied method of addressing this issue is to evaluate a suite of either overlapping or nonoverlapping spectral estimates for a host of window locations, and to subsequently average them and calculate statistical bounds on this mean. The most commonly used technique along these lines is called *Welch's Method* (see the *welch* function in the MATLAB signal processing toolbox).

An elegant, computationally-intensive, and increasingly widely utilized method of estimating spectra (see the *pmtm* function in MATLAB's signal processing toolbox) is *multitaper spectral estimation* (e.g., Thomson, proc. IEEE V. 70, No 9, September 1982). In multitaper spectral estimation, a family of statistically independent spectral estimates is obtained from a signal using an orthogonal set of windows on the estimation interval that are referred to as *prolate spheroidal tapers* (Figure 6).

In multitaper spectral estimation individual spectra obtained from the prolate spheroidal tapers are combined in a weighted sum to produce a spectral estimate with leakage that is approximately limited to some specified frequency

band, $\pm W$. Specifically, for a specified time-bandwidth product, NW, the multitapers are the Fourier Transforms of solutions, U_k, to the frequency-domain eigenvalue-eigenfunction equation

$$\int_{-W}^{W} \frac{\sin N\pi(f-f')}{\sin \pi(f-f')} U_k(N,W;f')\, df' = \lambda_k(N,W) \cdot U_k(N,W;f)\ . \qquad (13)$$

where the λ_k are eigenvalues (the first $2NW$ of which are close to one), and N is the discrete length of the taper sequence (this is a discrete formulation for spectral estimation on sampled time series, which we shall discuss next shortly. The integral in (13) is a convolution in the frequency domain between the U_k and the Dirichlet kernel, a function that arises frequently in discrete Fourier analysis because it is the Fourier transform of the sampled counterpart of the boxcar function (more on this later). Solutions to (13) form an orthogonal family of functions which have the greatest fractional energy concentration in the frequency interval $(-W, W)$. The eigenvalues in (13) are measures of the degree to which spectral leakage is confined to $(-W, W)$. Spectral leakage becomes increasingly worse for higher-order tapers, with the energy leakage being given approximately as

$$1 - \lambda_k \approx \frac{\sqrt{2\pi}}{k!}(8c)^{k+1/2} e^{-2c} \qquad (14)$$

where $c = \pi NW$. Figure 7 shows the fractional leakage for the first five multitapers.

Because of the appreciable leakage of the higher order tapers the lowest few (typically six or so, depending on the values of N and W) are typically used in practice. Figures 8 through 12 show the five lowest order multitaper estimates for $NW = 4$ for the example sine wave signal used in the earlier figures. Figure 13 shows the multitaper estimate obtained by averaging them. The leakage function displayed in Figure 13 approximates a frequency boxcar of width $2W$.

An example geophysical application of the PSD is to quantify the background noise characteristics of seismic stations, so as to gauge, for example, how they compare to known very quiet sites, and to assess what frequency bands good or bad for signal detection. This is of considerable importance both for earthquake and Earth structure studies and for estimating detection thresholds for clandestine events (e.g., nuclear tests). Figure 14 shows PSD estimates for a fairly quiet IRIS broadband seismic station in the Tien Shan mountains near Ala Archa, Kyrgyzstan, at periods ranging from 0.1 to 10^3 s (about 17 minutes). The bounding curves are empirically-based high- and low-noise extremal models for broadband stations. Noise at short periods is dominated by cultural (man-made), wind, and other rapidly varying environmental effects. The prominent noise peaks near 7 and 14 seconds are globally observed and are generated by ocean waves. The long-period power is higher on the horizontal sensors as opposed to the vertical sensors because they are sensitive to tilt caused by barometric, thermal, or other long-period noise sources. The peak near 1.6 s is unusual and may represent microseismic wave noise from the nearby Issyk Kul, one of the largest high-altitude alpine lakes in the world.

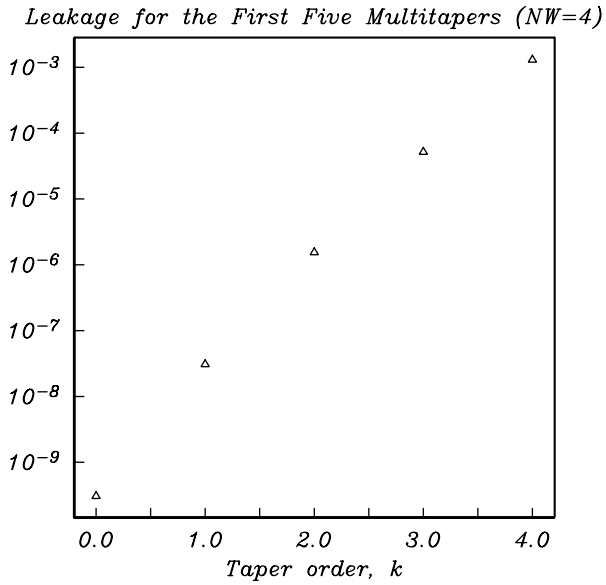

Figure 7: Fractional energy leakage outside of $f = (-W, W)$ for the first five multitapers ($NW = 4$).

As a final indication of the great utility of the PSD, the Figure (15) shows processed PSDs from a broadband seismometer (Guralp CMG-3Tb) located in a 255-m deep borehole in the polar icecap near the South Pole. A great many of 1-hour data length, 50% overlap, PSDs using a hamming taper, were calculated from the month of May, 2003, and the resulting individual PSDs were used to assemble an empirical probability density function for the signal characteristics at he station The bifurcation of the high frequency noise is caused by intermittent periods where tractors are moving snow near the station. Pink misty areas concentrated around 1 and 20 seconds are PSDs that include teleseimic earthquake signals. For the most part, the high frequency noise at this site reestablishes the Peterson Low Noise model (lower curve) and is this among the quietest stations on Earth.

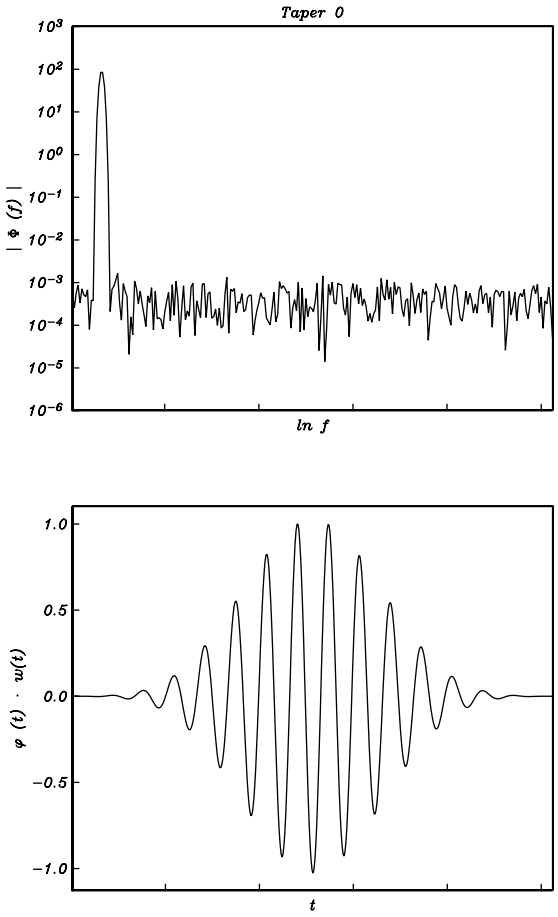

Figure 8: Prolate spheroidal taper spectral estimate ($k = 0$; $NW = 4$).

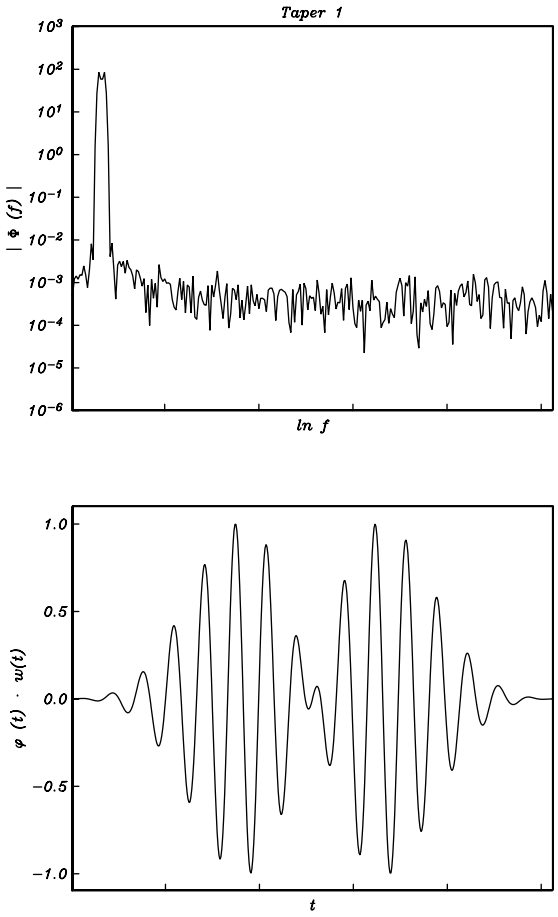

Figure 9: Prolate spheroidal taper spectral estimate ($k = 1$; $NW = 4$).

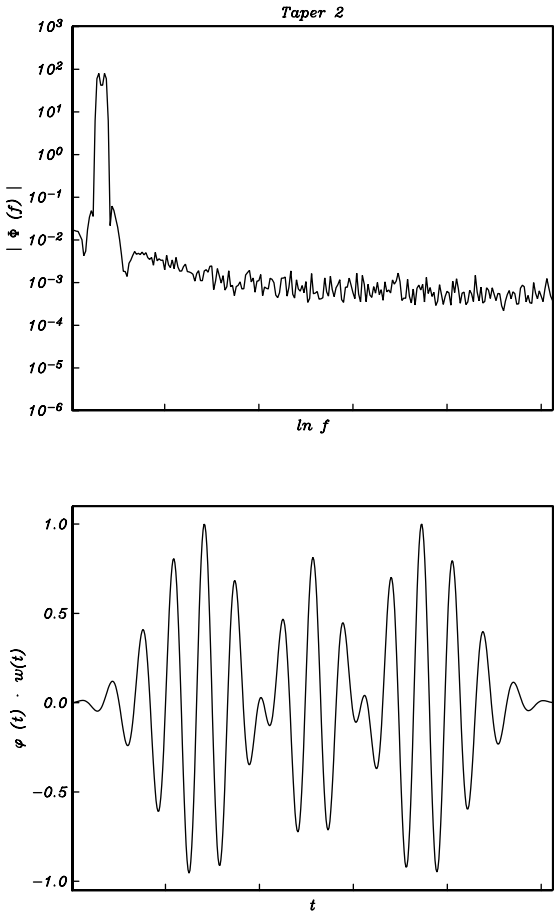

Figure 10: Prolate spheroidal taper spectral estimate ($k = 2$; $NW = 4$).

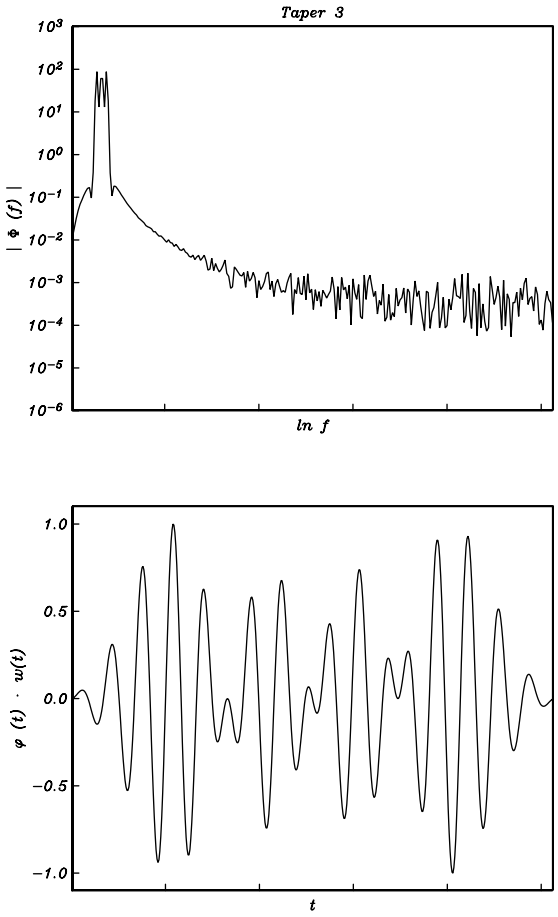

Figure 11: Prolate spheroidal taper spectral estimate ($k = 3$; $NW = 4$).

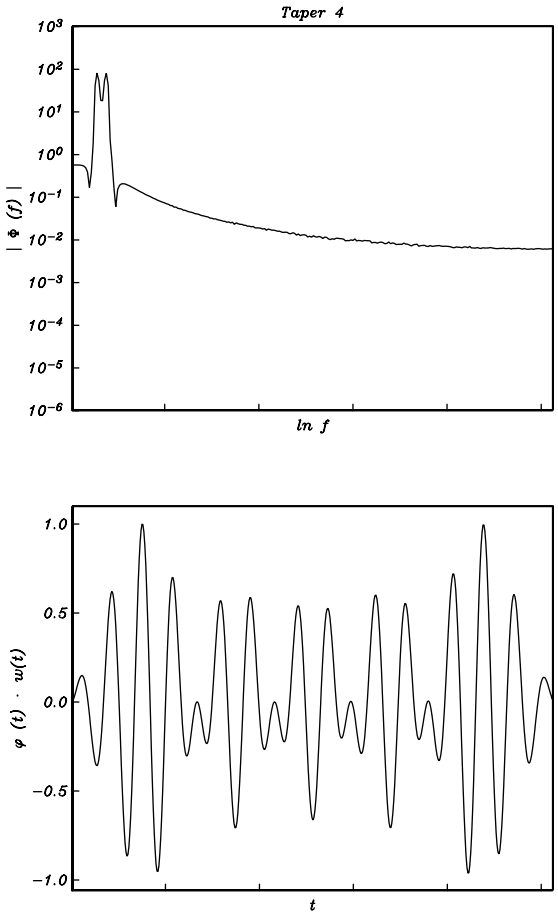

Figure 12: Prolate spheroidal taper spectral estimate ($k = 4$; $NW = 4$).

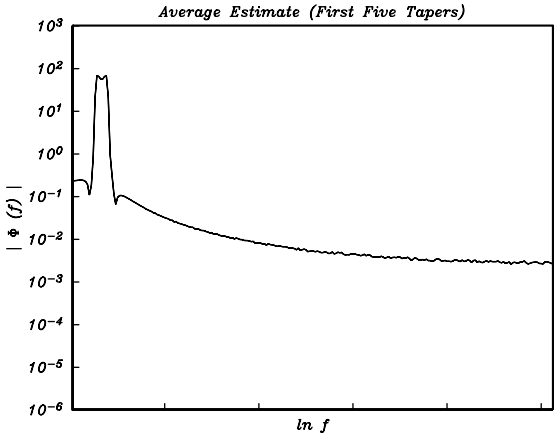

Figure 13: Prolate spheroidal taper average spectral estimate ($0 \leq k \leq 4$; $NW = 4$).

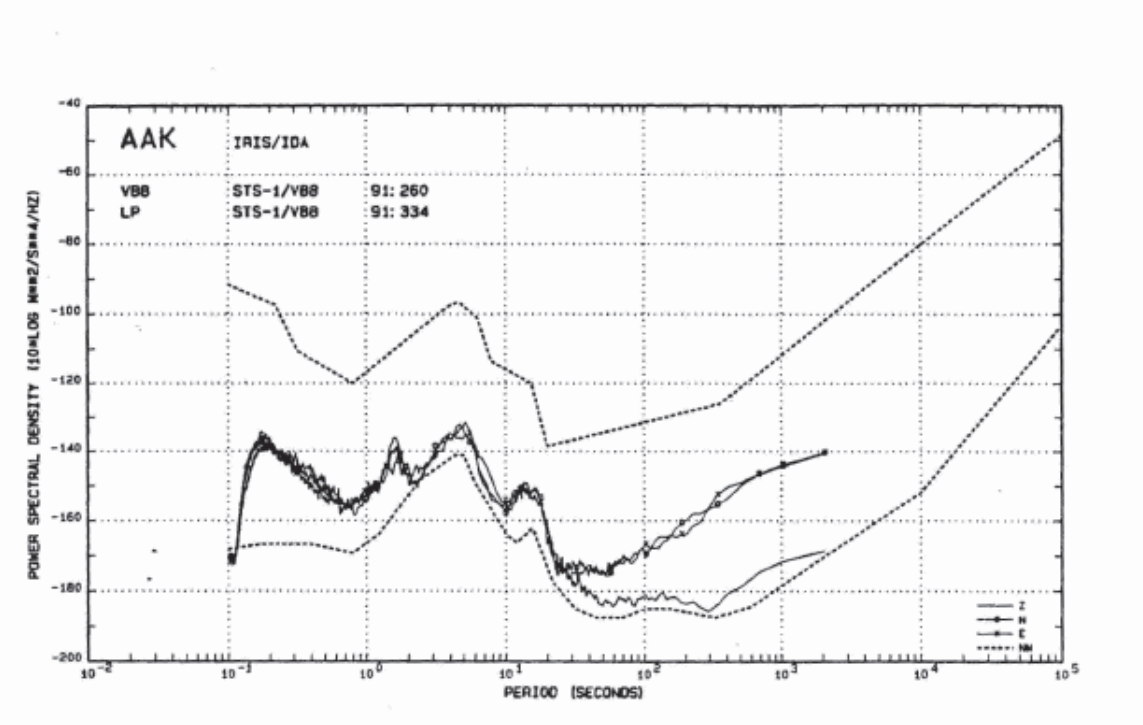

Figure 14: Earth Acceleration Power Spectral Density for background noise at the Ala Archa IRIS/IDA station as a function of period. Z, N, E refer to vertical, north, and east seismometer components. Curves labeled NM are the empirical noise model bounds of Peterson (1994) denoting to extremal PSD values from stations installed around the world. The reference (0 db) level is $(1\ \mathrm{m/s^2})^2/\mathrm{Hz}$. PSD estimates were obtained using Welch's method.

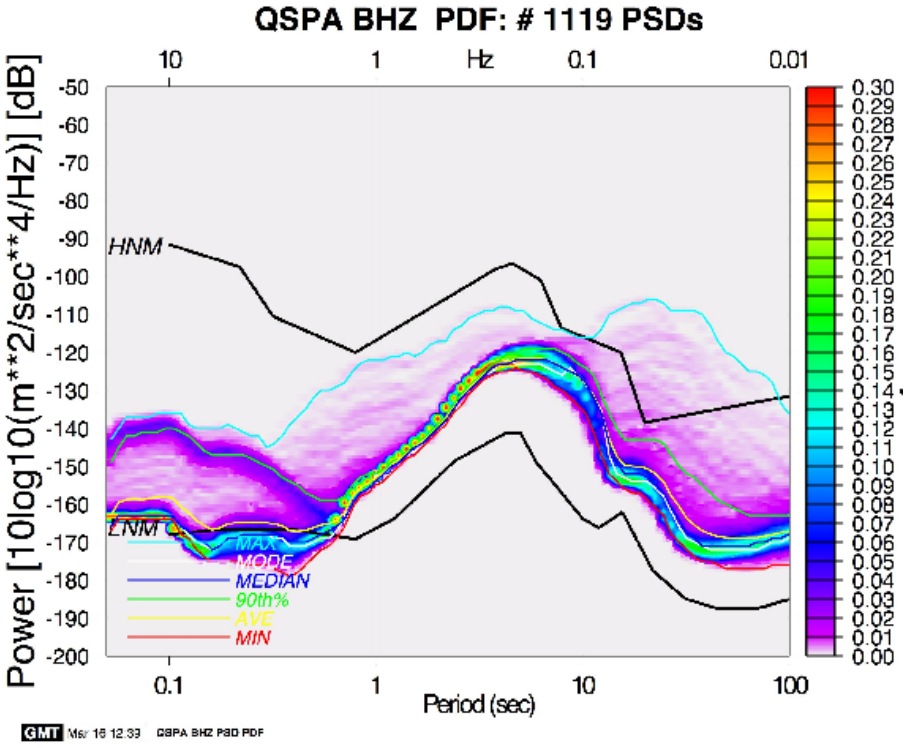

Figure 15: Quiet South Pole (QSPA) Global Seismic Network Station, power spectral density probability density plot