Investigation of Signal Characteristics Using the Continuous Wavelet Transform

John Sadowsky

The continuous wavelet transform (CWT) displays the scale-dependent structure of a signal as it varies in time. This scale-dependent structure, in turn, is essentially the instantaneous frequency, so that the CWT provides a view of the frequency versus time behavior of the signal and therefore has great potential as a preliminary tool for investigating wideband, nonstationary, or other types of signals having time-dependent spectral characteristics. In this article, the CWT is considered as a qualitative tool that can be used to analyze such signals. The results of this analysis would then be used to construct appropriate signal processing algorithms to detect, characterize, and classify the signals. A particular complex-valued wavelet, proposed by Morlet, has good properties for use as the kernel in a qualitative CWT. These properties are presented, and the use of the CWT for studying wideband communications signals and for finding features that might be of phenomenological significance in a seismic signal is discussed.

INTRODUCTION

The continuous wavelet transform (CWT) provides a method for displaying and analyzing characteristics of signals that are dependent on time and scale. It therefore is potentially a useful tool for detecting and identifying signals with exotic spectral features, transient information content, or other nonstationary properties. The CWT is an operator that takes a signal and produces a function of two variables: time and scale. As a two-variable function, it can be considered as a surface or image. The idea to be explored in this article is that features in this CWT surface result from parameters of the signal that could prove useful for its detection, characterization, classification, or conditioning.

As a first step toward a process of signal characterization, we review the definition of the CWT and why this definition would have any relationship to the time-varying scale and frequency structure of a signal. A particular wavelet well-suited to qualitative analysis of time series is the Morlet wavelet. An analysis of this wavelet is presented and its basic properties are described. Analytic calculations of the Morlet-based CWT on fundamental signals such as sinusoids, impulses, and linear frequency modulation are then derived, and the resulting CWT surfaces are shown and discussed.

Finally, a set of nonstationary and wideband signals is presented for which the CWT is an especially good tool for analysis and for the design of signal processing algorithms. These signals are of special interest in advanced communications and radar systems and indicate a particular advantage to considering the CWT as part of a signal intelligence toolkit. A particular empirical example of the use of the CWT for qualitative investigation of signal characteristics is also presented. In this example, a sampled time series of an underground explosion intended to emulate a seismic event is used as input to a CWT surface-generation algorithm. This time series is a good example of a transient waveform with certain additional characteristics. The CWT surface has features that can apparently be attributed to the shock and aftershock features in the event, and these features are not readily apparent for the time series or a spectral estimation from the time series.

REVIEW OF THE CONTINUOUS WAVELET TRANSFORM

In this section, we review the definition and basic properties of the CWT. These properties were described in a previous article,[1] and more complete and mathematical treatments can be found in the classical references in the field (see, for example, Daubechies[2] and Meyer[3]). We also introduce the particular mother wavelet used for the investigation in this article.

The CWT is defined with respect to a particular function, called a mother wavelet, that satisfies some particular properties. Not every function can qualify to be a mother wavelet. As the kernel function of a signal transform, it is important that the mother wavelet be designed so that the transform can be inverted—there must be some related transform that permits one to recover the original signal from its CWT. Even if the application of the CWT does not require such transform inversion, the invertibility of the CWT is necessary to assure that no signal information is lost in the CWT. Signal information may be restructured or rearranged, but it must still be present in the CWT for the original signal to be reconstructed.

The most important property that must be satisfied by the mother wavelet is the "admissibility condition," which is required for an inverse wavelet transform to exist. This admissibility condition is discussed formally in the boxed insert (Admissibility Condition for the Mother Wavelet). It implies that the Fourier transform of the mother wavelet is 0 at frequency 0. Thus, the mother wavelet has no DC bias and, therefore, must have oscillations to cause it to act as a bandpass filter. Figures 1a and 1b illustrate this property for two example real-valued mother wavelets, the Mexican hat function $\Psi(t) = c(1-t^2)e^{-t^2/2}$ (where c is a normalizing constant and t is time), and the Daubechies D5 function,[2] a wavelet with nice smoothness and compactness characteristics but whose definition is somewhat complicated and is indirectly given through the coefficients of a dilation equation satisfied by the wavelet.

Suppose that $\psi(t)$ denotes the mother wavelet, and that $\hat{\psi}(\omega)$ denotes its Fourier transform. We use the definition (in which j denotes $\sqrt{-1}$)

$$\hat{\psi}(\omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \psi(t) e^{-j\omega t} dt \,. \tag{1}$$

Then the admissibility condition implies that $\hat{\psi}(0) = 0$, that is, $\int \psi(t) dt = 0$. Figure 2 illustrates this effect for the Morlet wavelet $\psi(t)$, defined by the formula

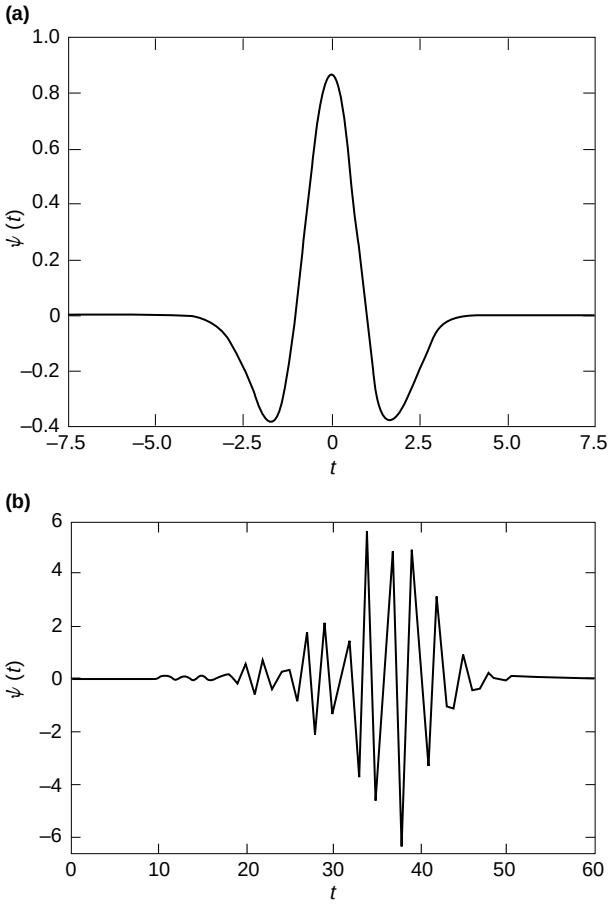

Figure 1. Examples of real-valued wavelets satisfying the admissibility condition. (a) The Mexican hat wavelet, $\Psi(t) = (2/\sqrt[4]{\pi}\sqrt{3})(1-t^2)e^{-t^2/2}$. (b) The Daubechies D5 wavelet, derived from dilation coefficients (see Ref. 2, pp. 194–202).

$$\psi(t) = \sqrt{2} e^{-t^2/\alpha^2} \left(e^{j\pi t} - e^{-\pi^2\alpha^2/4} \right). \tag{2}$$

The parameter α in Eq. 2 can be set to obtain desired time–frequency shaping in the CWT. Note that in the graph of $\psi(t)$ in Fig. 2a, there appear to be no oscillations, but this apparent absence is the result of the graph being that of the amplitude of the complex waveform—the oscillations are contributed by the phase. One can compute directly the Fourier transform of this wavelet. It is

$$\hat{\psi}(\omega) = \alpha e^{-\alpha^2(\pi^2+\omega^2)/4} \left(e^{\pi\alpha^2\omega/2} - 1 \right). \tag{3}$$

From Eq. 3, one immediately sees that $\hat{\psi}(0) = 0$ and that the admissibility condition holds. A more explicit

ADMISSIBILITY CONDITION FOR THE MOTHER WAVELET

For the continuous wavelet transform to be invertible, the mother wavelet $\Psi(t)$ must satisfy the admissibility condition:

$$\int_{-\infty}^{\infty} \frac{\left|\hat{\Psi}(\omega)\right|^2}{|\omega|}\, d\omega < \infty ,$$

where $\hat{\Psi}(\omega)$ is the Fourier transform of the wavelet. We suppose that $\Psi(t)$ is continuous with continuous Fourier transform. If $\hat{\Psi}(0) \neq 0$, then from continuity, there is a small interval I containing 0, and an $\epsilon > 0$ such that $\left|\hat{\Psi}(\omega)\right| > \epsilon$ for all ω in I. But it would then follow that

$$\int_{-\infty}^{\infty} \frac{\left|\hat{\Psi}(\omega)\right|^2}{|\omega|}\, d\omega \geq \int_{I} \frac{\left|\hat{\Psi}(\omega)\right|^2}{|\omega|}\, d\omega \geq \int_{I} \frac{\epsilon^2}{|\omega|}\, d\omega = \infty .$$

The admissibility condition therefore implies that the mother wavelet has no DC component, that is, $\hat{\Psi}(0) = 0$. One computes directly that

$$\hat{\Psi}(0) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \Psi(t)\, dt ,$$

and so the admissibility condition implies that the integral of the mother wavelet is zero. For this to occur, the mother wavelet must contain oscillations; it must have sufficient negative area to cancel out the positive area. Of course, this is what it means to have no DC component.

We will explain in what sense the admissibility condition implies that the continuous wavelet transform has an inverse. For this, we use the notation, $\Psi^{a,b}(t)$ to denote the scaled, time-shifted wavelet, $\Psi^{a,b}(t) = \frac{1}{\sqrt{|a|}} \Psi\left(\frac{t-b}{a}\right)$. We can therefore express the continuous wavelet transform of a signal $s(t)$ as $(T^{wav}s)(a,b) = \int_{-\infty}^{\infty} s(t)\overline{\Psi^{a,b}}(t)dt$. Also, we will use the notation $\langle f, g \rangle$ for the scalar product in the Hilbert Space of finite energy signals, $\langle f, g \rangle = \int f(t)\overline{g(t)}\, dt$.

The admissibility condition implies that we can define a finite value C_Ψ via

$$C_\Psi = 2\pi \int_{-\infty}^{\infty} \frac{\left|\hat{\Psi}(\omega)\right|^2}{|\omega|}\, d\omega .$$

If $f(t)$ and $g(t)$ are two finite energy signals, then one can compute the following using the inverse Fourier transform and the Plancherel theorem:

$$\iint (T^{wav}f)(a,b)\overline{\left\langle g, \Psi^{a,b} \right\rangle} \frac{da\, db}{a^2} = C_\Psi \langle f, g \rangle .$$

Thus, considered as operators on the signal $g(t)$, we see that the linear functionals $\iint (T^{wav}f)(a,b)\left\langle \Psi^{a,b}, \_ \right\rangle \frac{da\, db}{a^2}$ and $C_\Psi \langle f, \_ \rangle$ both perform the same operation. It is in this weak convergence sense that we therefore can conclude that

$$f(t) = \iint (T^{wav}f)(a,b)\Psi^{a,b}\, \frac{da\, db}{a^2} ,$$

i.e., that the difference between the two finite energy signals acts as a zero linear functional on the space of finite energy signals. We note that this inversion is as one would expect—the continuous wavelet transform of the signal provides the coefficients of a decomposition of the signal into a superposition of scaled and translated mother wavelets, integrated with respect to the scaling metric $\frac{da\, db}{a^2}$.

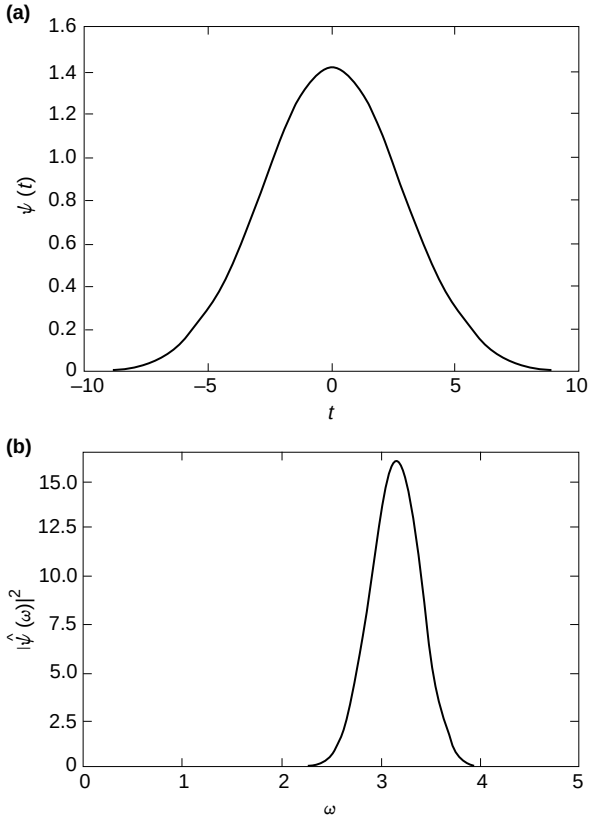

Figure 2. The Morlet wavelet $\psi(t)$ and its power spectrum. For simplicity, the parameter a is set to 4. (a) Amplitude of the complex Morlet wavelet. (b) Power spectrum of the Morlet wavelet.

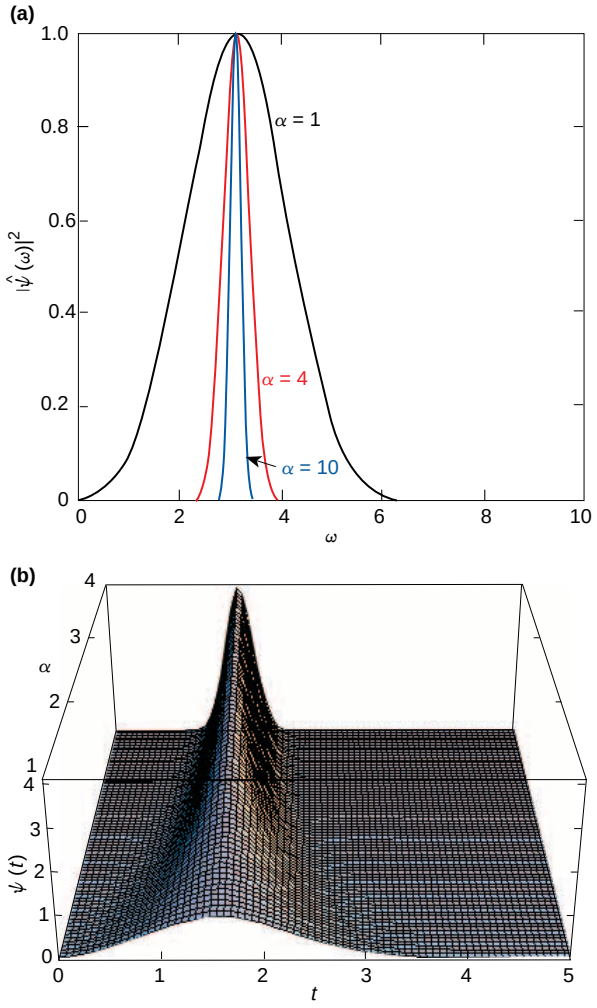

Figure 3. The power spectrum (normalized) of the Morlet wavelet for several values of the shaping parameter α. (a) Power spectrum at α = 1, 4, and 10. (b) Fourier transform for α varying continuously from 1 to 4.

verification of the admissibility condition is given in the boxed insert (Proof of the Admissibility of the Morlet Wavelet).

In Fig. 2, we see that the Morlet wavelet, in the frequency domain, is a complex bandpass filter. Its effect as a filter would be to limit a signal to a band centered about the frequency of approximately π rad/s, with the center point approaching π as the parameter α gets large. One can also calculate that the 3-dB bandwidth of the Morlet wavelet is approximately $3.33022/\alpha$ (where $4\sqrt{\ln(2)} \approx 3.33022$) and is thus inversely proportional to the parameter α. Figure 3 illustrates this for several values of the parameter α.

The CWT of a signal $s(t)$ with respect to the wavelet $\Psi(t)$ is a function of the two variables $\alpha > 0$ and b, and is defined by the expression

$$(\mathfrak{C}s)(a,b) + \frac{1}{\sqrt{a}} \int_{-\infty}^{\infty} s(t) \Psi^* \left(\frac{t-b}{a} \right) dt \,, \tag{4}$$

where * denotes complex conjugation. Suppose that we denote by subscript a,b the rescaling of a function by a and the translation of a function by b; that is, let $\Psi_{a,b}(t) = (1/\sqrt{a})\Psi[(t-b)/a]$. Then $(\mathfrak{C}s)(a,b)$ is just the correlation of the signal $s(t)$ by $\Psi_{a,b}(t)$. Now, one can easily compute the Fourier transform of $\Psi_{a,b}(t)$ as

$$\hat{\Psi}_{a,b}(\omega) = \sqrt{a}\hat{\Psi}(a\omega)e^{-j\omega b} \,. \tag{5}$$

A rescaling by a in the time domain becomes a rescaling by $1/a$ in the frequency domain, which has the effect of moving the center frequency of the passband from π to π/a, with a similar rescaling of the 3-dB bandwidth.

Effectively through a Parseval's identity, one can compute a frequency domain formulation of the CWT and obtain

PROOF OF THE ADMISSIBILITY OF THE MORLET WAVELET

To prove that the Morlet wavelet, $\psi(t) = \sqrt{2}e^{-t^2/\alpha^2}\left(e^{j\pi t} - e^{-\pi^2\alpha^2/4}\right)$, satisfies the admissibility condition, we must look at upper bounds on $\left|\hat{\psi}(\omega)\right|^2 \Big/ |\omega|$ for ω near 0 and approaching $+\infty$ and $-\infty$.

Now, as was noted in the text, $\hat{\psi}(\omega) = e^{-\alpha^2(\pi^2+\omega^2)/4}\left(e^{\pi\alpha^2\omega/2} - 1\right)$. Near $\omega = 0$, the term $e^{-\alpha^2(\pi^2+\omega^2)/4}$ is approximately $e^{-\alpha^2\pi^2/4}$. Thus, we need to consider the behavior of $\left(e^{\pi\alpha^2\omega/2} - 1\right)^2 \Big/ |\omega|$ as ω approaches 0 from the left and right. But, by L'Hospital's rule, for example, it is clear that this ratio approaches 0 in these cases, and so the integral $\int_{-1}^{1} \frac{\left|\hat{\psi}(\omega)\right|^2}{|\omega|}\, d\omega$ is finite.

As ω approaches $+\infty$, one has

$$\begin{aligned}\hat{\psi}(\omega) &= e^{-\alpha^2(\pi^2+\omega^2)/4}\left(e^{\pi\alpha^2\omega/2} - 1\right) \le e^{-\alpha^2\pi^2/4} e^{-\alpha^2(\omega^2 - 2\pi\omega)/4} \\ &= e^{-\alpha^2(\omega-\pi)^2/4},\end{aligned} \tag{I}$$

and so it follows that for some positive B,

$$\int_B^{\infty} \frac{\left|\hat{\psi}(\omega)\right|^2}{|\omega|}\, d\omega \le \int_B^{\infty} \frac{1}{\omega} e^{-\alpha^2(\omega-\pi)^2/2} d\omega \le \int_B^{\infty} \frac{d\omega}{\omega^2} < \infty. \tag{II}$$

The second inequality of Expression II is true for all ω large enough so that $e^{-\alpha^2(\omega-\pi)^2/2} \le \frac{1}{\omega}$, i.e., greater than some B so that $(\omega-\pi)^2 > \frac{2}{\alpha^2}\log(\omega)$.

As ω approaches $-\infty$, one has $e^{\pi\alpha^2\omega/2}$ approaching 0 and so, in particular, for ω less than some negative C, $\left(e^{\pi\alpha^2\omega/2} - 1\right) \le 1$. Thus, we have

$$\int_{-\infty}^{C} \frac{\left|\hat{\psi}(\omega)\right|^2}{|\omega|}\, d\omega \le \int_B^{\infty} \frac{1}{\omega} e^{-\alpha^2(\omega^2+\pi^2)/2} d\omega < \infty, \tag{III}$$

using the same reasoning used in Expression II.

$$(\mathfrak{C}s)(a,b) = \sqrt{a}\int_{-\infty}^{\infty} \hat{s}(\omega)\hat{\Psi}^*(a\omega)e^{j\omega b}d\omega\,. \qquad (6)$$

Both the time domain and frequency domain formulations of the CWT can be used to compute the Morlet wavelet transform of the prototypical basic signals, the complex sine wave, and the impulse. For the complex sine wave of frequency f_0 Hz, $s(t) = e^{j2\pi f_0 t}$, it is simpler to use the frequency domain formulation, as $\hat{s}(\omega) = \sqrt{2\pi}\delta(\omega - 2\pi f_0)$, with $\delta(t)$ being the Dirac delta distribution. Thus,

$$(\mathfrak{C}s)(a,b) = \sqrt{2\pi a}\hat{\psi}^*(2\pi a f_0)e^{j2\pi f_0 b}\,. \qquad (7)$$

The equation for the Morlet CWT of the sine wave (Eq. 7) therefore separates into a product of a function only of a and a function only of b. The b function is strictly a periodic phase function whose period is proportional to the frequency of the sine wave. The a function is a scaled version of the Fourier transform of the Morlet wavelet and thus has a peak at scale value $a = 1/2f_0$. The CWT surface should, therefore, show a ridge at this location, the width of which is proportional to a. Figure 4 contains CWTs for three different complex sine waves to illustrate this surface feature.

The CWTs illustrated in Fig. 4 and the subsequent figures are complex-valued surfaces. Each point of the surface is therefore a complex number of a given amplitude and phase. Amplitude is represented in the figures by pixel density (color saturation), from low density (white background) for near-zero amplitude, to saturated color for high density. The phase is represented by the color spectrum from phase 0° (purple) to phase 360° (red). The horizontal axis represents time in the CWT (variable b), and the vertical axis is logarithmic and represents scale [variable $\log(a)$]. In the vertical axis, $a = 0$ is at the top of the figure, so that small scale (high frequency) is rendered above larger scale (low frequency), as is customary for time–frequency plots.

The CWT surfaces in Fig. 4, therefore, contain horizontal ridges at the appropriate scales. The ridge widths are constant for the three frequencies because the scale axis is logarithmic and bandwidth in a wavelet transform is constant over logarithmic scales. In Fig. 5, we illustrate a CWT for an impulse, $s(t) = \delta(t - t_0)$, in which the impulse occurrence t_0 is placed 1/3 of the way across the b axis. For this signal, it is simpler to use the time-domain formulation and compute

$$(\mathfrak{C}s)(a,b) = \frac{1}{\sqrt{a}}\psi^*\left(\frac{t_0 - b}{a}\right). \qquad (8)$$

Ideally, an impulse would appear as a vertical ridge in the CWT surface. The fanning effect at lower frequencies in Fig. 5 is a result of the specific form of the Morlet wavelet.

In many applications, a modified version of the Morlet wavelet is used, which we shall call the Morlet

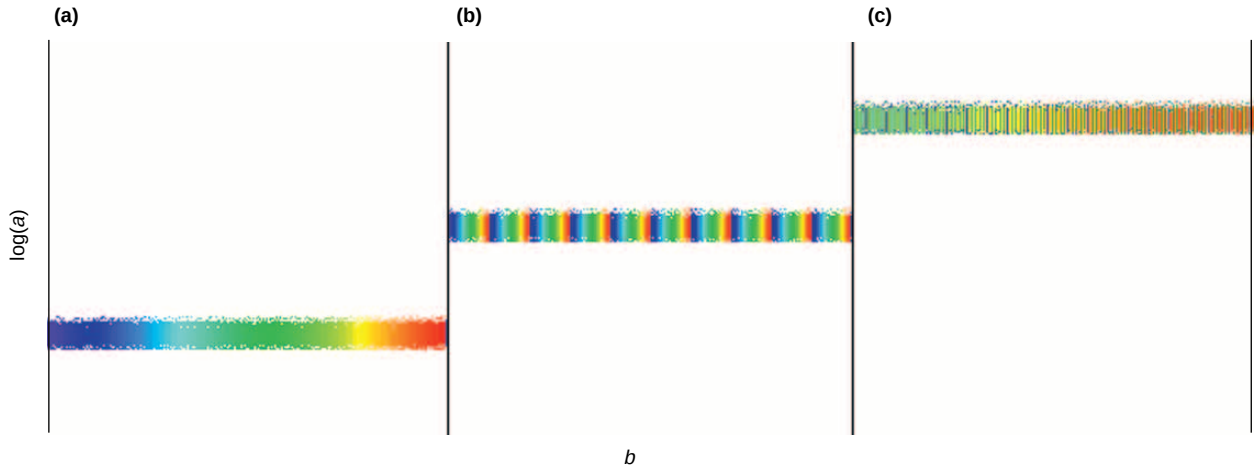

Figure 4. The continuous wavelet transform of a complex sine wave generated using the Morlet wavelet. The horizontal axis is time translation (variable b) and the vertical axis is log scale [variable $\log(a)$], with small scale (high frequency) above larger scale (low frequency.) (a) Low-frequency sine wave. (b) Medium-frequency sine wave. (c) High-frequency sine wave.

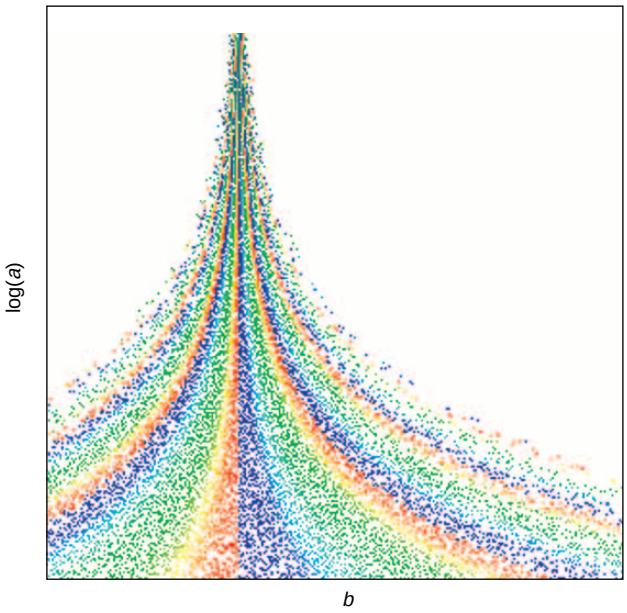

Figure 5. The continuous wavelet transform of an impulse generated using the Morlet wavelet. The horizontal axis is time translation (variable *b*) and the vertical axis is log scale [variable log(*a*)], with small scale (high frequency) above larger scale (low frequency.)

pseudowavelet. Kronland-Martinet et al.[4] have used it in the analysis of sound, for example, and Dutilleux[5] has used it in the analysis of musical instruments. Its general form is

$$\xi(t) = \frac{1}{\sqrt{2\pi}} e^{-(t^2/2) + j\omega_0 t}, \tag{9}$$

where ω_0 is a parameter that can be set to center wavelet over the appropriate range of frequencies. In analogy with the Morlet wavelet, for example, one has the normalized value of $\omega_0 = \pi$. In Ref. 5, a suggested reasonable value for ω_0 for the analysis of speech or music is between 5 and 6.

Strictly speaking, the Morlet pseudowavelet is not a mother wavelet because it can be shown that it does not satisfy the admissibility condition. Indeed, one can easily compute

$$\int_{-\infty}^{\infty} \xi(t) d\xi = e^{-\omega_0^2/2} \neq 0. \tag{10}$$

Equation 10 can be made arbitrarily close to 0 by choosing ω_0 large enough. For $\omega_0 = 5$, for example, this integral is less than 4×10^{-6}. Although the pseudowavelet is not necessarily useful for reconstructing a signal from its CWT, as the resolution of the identity no longer holds, it is quite useful for time–frequency display of signals, because relevant features of the signal appear as patterns in the surface and the Morlet pseudowavelet is much simpler than the Morlet wavelet to use in computations. For example, the Fourier transform of the Morlet pseudowavelet takes the much simpler form

$$\hat{\xi}(\omega) = \frac{1}{\sqrt{2\pi}} e^{-(\omega - \omega_0)^2/2}. \tag{11}$$

An example of a signal with finite duration and with a time-varying spectrum is a chirp signal, that is, a signal with quadratic phase. Because the phase is quadratic, its instantaneous frequency (the derivative of its phase) is a linear function of time, and so the signal is often called a linear frequency modulation (LFM) signal. Thus, the frequency spectrum is shifting either up or down (depending on the quadratic coefficient in the phase) over time. For our example, we have restricted the chirp signal to a finite duration with a Gaussian envelope. The form for the signal is

$$s(t) = e^{-\alpha t^2/2 + j\beta t^2/2 + j\gamma t}. \tag{12}$$

Here, the parameter α controls the Gaussian envelope time duration, β is the linear frequency change rate, and γ is the initial frequency. The frequency spectrum for this chirp signal is centered about the frequency γ rad/s and has a bandwidth proportional to $\sqrt{\alpha^2 + \beta^2}$. Figure 6 shows the chirp signal and its power spectrum. Although the parameters of the chirp can be estimated from the spectrum, if the signal of interest is known to be a chirp and if the duration parameter of the Gaussian window is known, there is nothing implicit in the shape of its power spectrum that specifically characterizes LFM.

Figure 7 illustrates the CWT of a chirp signal. Here we have used the Morlet pseudowavelet to simplify the computation somewhat. A Morlet wavelet would decrease the amount of lower-amplitude spreading but would not qualitatively differ significantly in appearance. Through a tedious but not difficult calculation, the CWT for the chirp whose form is given in Eq. 12, with a Morlet pseudowavelet with parameter *k*, can be expressed in closed form. The expression is messy but straightforward to graph.

THE CWT OF A SEISMIC EVENT

Figure 8 contains the time series of measurements of a seismic event. A section of the CWT generated from this time series is shown in Fig. 9. This example

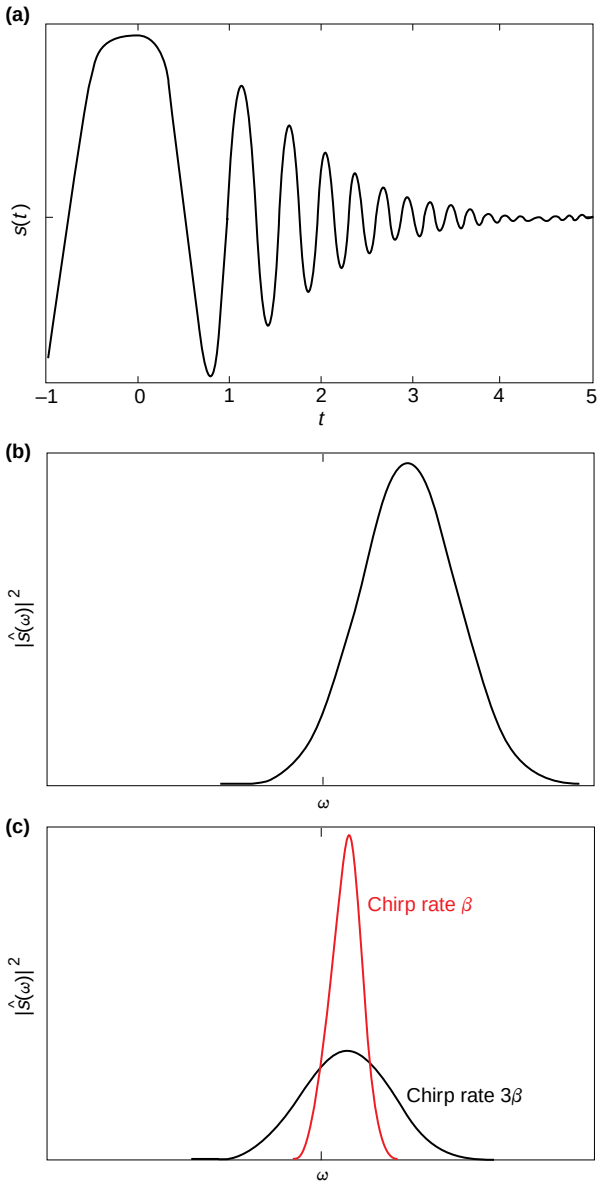

Figure 6. Time- and frequency-domain plots of the chirp (linear frequency modulated, LFM) signal with Gaussian window. (a) Time-domain plot of a chirp signal with Gaussian window. (b) Power spectrum of chirp signal. (c) Power spectrum of two chirp signals with the same initial frequency γ but different chirp rates β.

was generated from a data set of a sampled time series of shock waves at a specific location from a seismic event simulated by an underground explosion. As in the previous figures, the CWT surface is shown with the time delay (variable *b*) as the horizontal axis and the logarithm of the scale [log(*a*), for variable *a*] as the vertical axis with small scale above large scale. For clarity of the image, the background zero-amplitude regions are displayed as black rather than white. Figure 9 shows a section of the overall CWT surface. Again, amplitude is indicated with color density and phase with color.

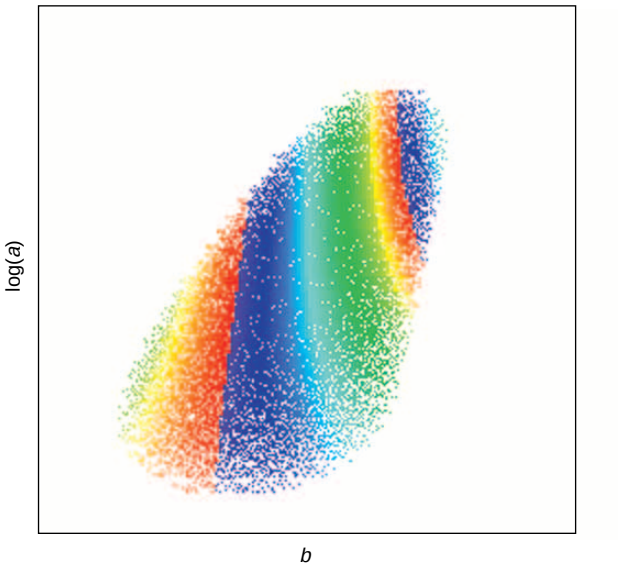

Figure 7. The CWT of a chirp signal showing a feature generated in a CWT surface by a chirp signal of sufficient chirp rate to span 5 octaves of scale change over its duration. The feature was isolated via a 3-dB threshold applied to the surface. The horizontal axis is time delay *b*, and the vertical axis is log scale log(*a*), with small scale above large scale.

Except for the striking patterns of phase changes across increasing scale at various time delays, it is difficult to see much structure in the amplitude variations. Thus, a threshold was applied to the image, presenting only those amplitudes within 5 dB of the maximum amplitude in the surface. The resulting surface is shown in Fig. 10.

In Fig. 10, we have divided the CWT surface into two sections to show a fairly long duration (range of variable *b*), presented as Figs. 10a and 10b. The initial shock is shown as a moderate-bandwidth, high-frequency island of relatively long duration. One can see in the figure that this shock actually consists of

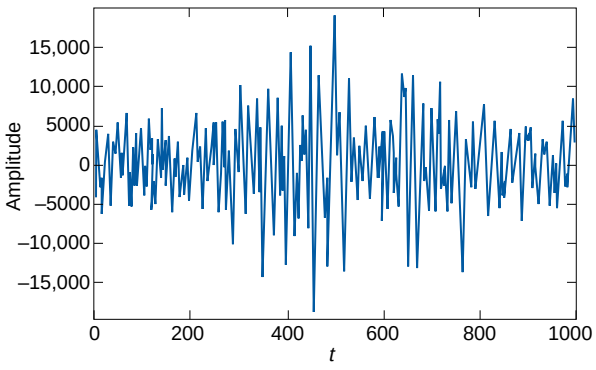

Figure 8. Time series of a seismic event simulated with an underground explosion.

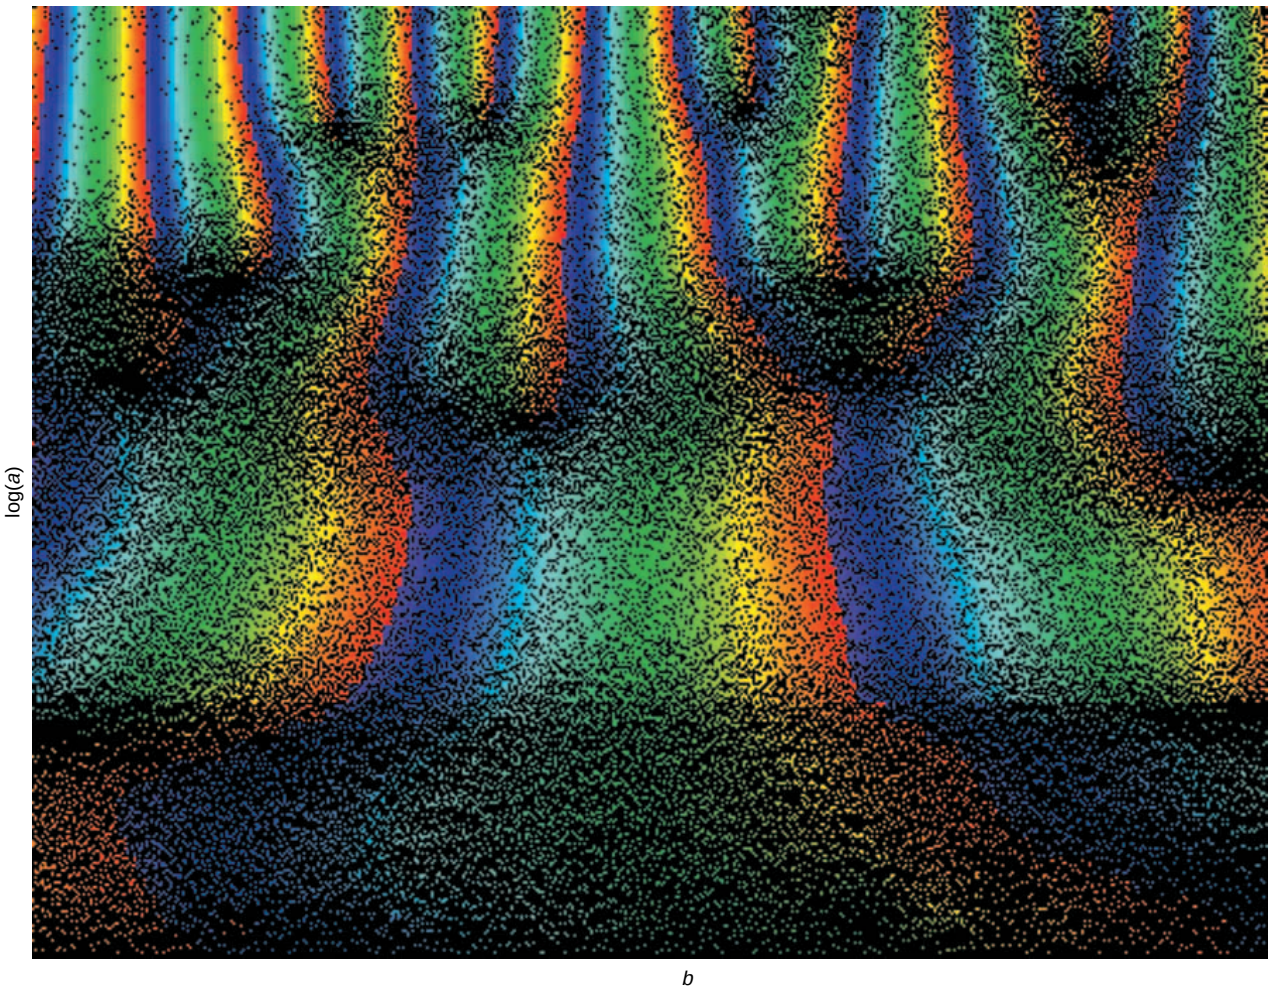

Figure 9. The CWT of a seismic event generated using the Morlet pseudowavelet. The background (zero-amplitude region) is black. The horizontal axis is time delay *b*, and the vertical axis is log scale log(*a*), with small scale above large scale.

components at two scales (frequencies) that undulate and blend together at certain times over their duration. An immediate aftershock appears a short time later and has two major components. One component is a high-frequency piece with characteristics similar to the original shock but with the larger of the two scale components somewhat reduced and disconnected.

The aftershock displays another component at a lower frequency and much smaller bandwidth, especially considering that the scale axis is logarithmic, so that bandwidth would increase in the CWT with increasing scale. This lower-frequency component appears on the right side of the image in Fig. 10a and is delayed in time, with respect to the higher frequency aftershock, by an amount approximately equal to the time delay between the shock and aftershock. Much smaller aftershocks continue to appear at spectral frequencies between these dominant frequencies for a period of time, as illustrated in Fig. 10b.

Although there are small nonzero slopes in the ridges in the CWT surface, the ridges are basically horizontal, indicating very little to no linear or nonlinear frequency modulation in the components. This is confirmed as well by the colored phase lines in the components, which maintain a fairly even crest-to-crest spacing, indicating constant frequency over the time duration. Moreover, the phase branching patterns shown in the nonthresholded surface of Fig. 9 appear in the thresholded surface to occur below the threshold of significant features. They reside in the noise and are essentially artifacts because the threshold-crossing ridges are at different scales and therefore have different phase periods, but the CWT generation algorithm strongly attempts phase continuity where possible.

To provide a physical interpretation for these features, perhaps as underground echo returns or other structural modes of vibration, would require seismological expertise. The CWT analysis is not intended to

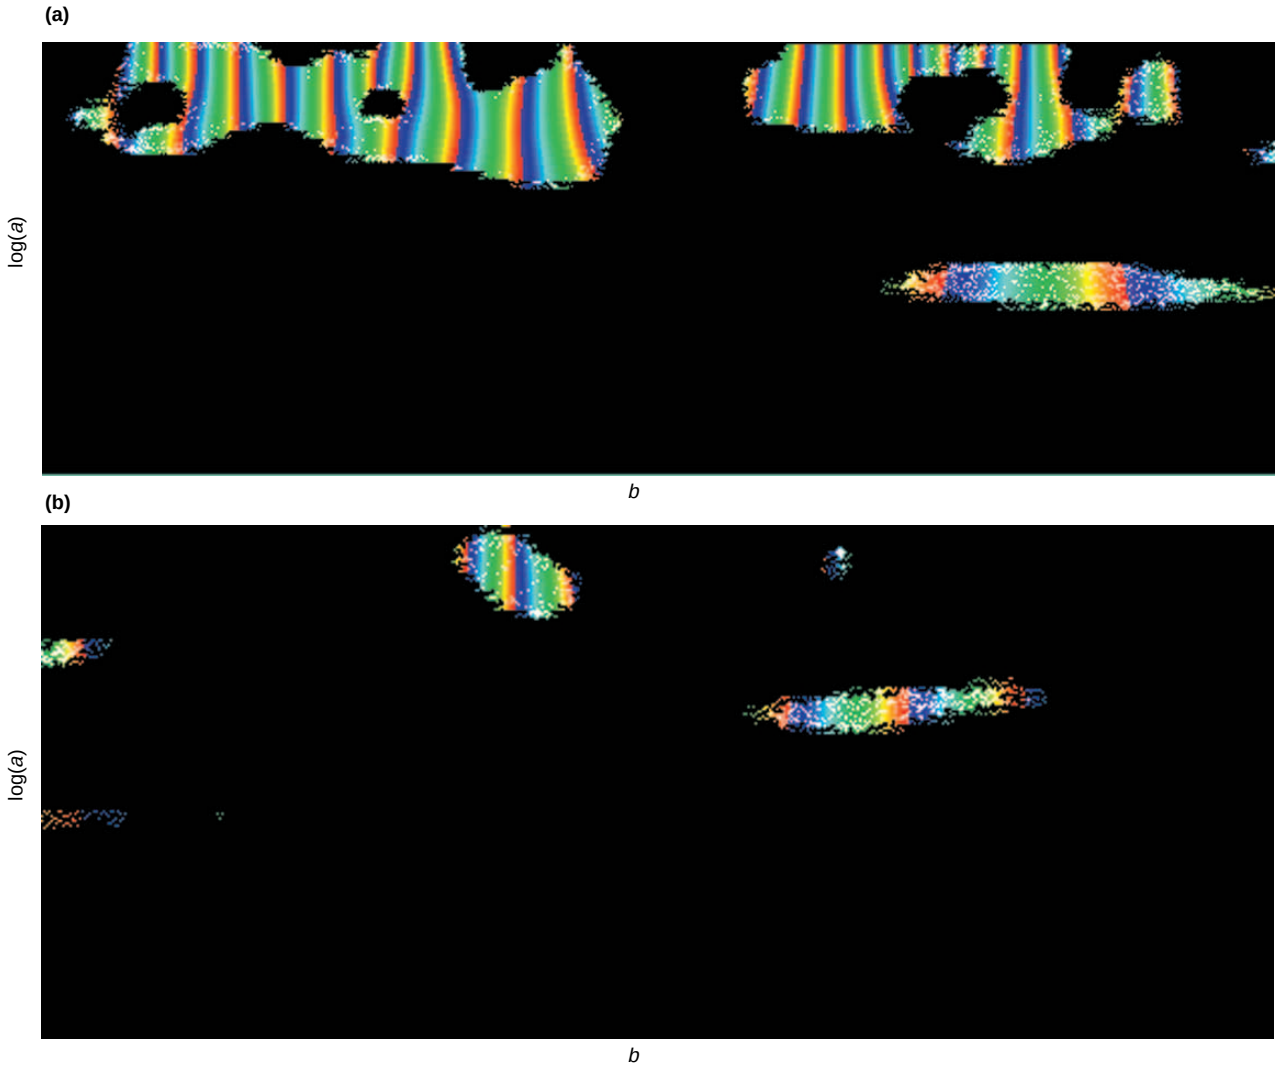

Figure 10. The CWT of the seismic event of Fig. 8, except with a threshold set at 5 dB below the surface maximum amplitude. The surface is a rather long time series and so is displayed in two rows: (a) onset of a seismic event, and (b) a continuation in time. The horizontal axis is time delay *b*, and the vertical axis is log scale log(*a*), with small scale above large scale.

provide such physical modeling. Rather, the CWT surface displays in a basically qualitative fashion those patterns that are likely to have physical interpretation and contribution to the phenomenological modeling. The CWT is therefore useful as an early data analysis tool to provide a different view and insight into time-scale and time–frequency behavior.

SIGNAL PROCESSING WITH THE CWT

The CWT is thus a natural tool to be used early in the investigation of the properties of signals to develop processing algorithms and concepts. It produces a representation of the time–frequency features of the signal, with the additional benefit of a so-called "zooming" effect. This zooming effect modifies the spectral resolution to be a function of scale—small-scale structure has higher resolution (frequency bandwidth) than does large-scale structure. In other words, the CWT zooms in to display detailed, fine features (high frequency) and zooms out to display large, coarse trends (low frequency). Thus, for signals for which spectral characteristics or statistical properties are likely to change over time, the CWT could be used to identify those signal features that are potentially exploitable by signal processing. Moreover, a processing approach to such signals might be to generate a CWT with the appropriate fast algorithm and process the resulting image with pattern recognition algorithms to detect and characterize objects in the CWT surface.

Signal intelligence (SIGINT) is an area of signal processing application that would benefit greatly from qualitative and quantitative analysis using the CWT. An important goal of SIGINT processing is to detect

and characterize communications, radar, and electronic countermeasures signals that are deliberately designed to be difficult to detect, locate, and characterize. Methods for design generally are based on the structure of such signals in the frequency domain and in the time domains. By presenting the same information contained in the signal in a different manner, such as the time-scale analysis of the CWT, it is possible that otherwise hidden structure can be emphasized.

CWT analysis has great potential for wideband communications signals, such as frequency-hopped and spread-spectrum waveforms. In all such signals, information for detection and characterization is found in the instantaneous spectrum, and the property that makes these signals most useful is that this spectrum changes over time. The signals are important for secure or covert communications specifically because these characteristics can be difficult to estimate by traditional signal processing methods. The CWT can be a useful qualitative tool for early analysis of signals of interest as preparation for subsequent design and development of signal-specific processing algorithms.

Frequency-hopped signals, for example, have narrowband instantaneous frequency that is centered at specific values for short durations, but over time these values can vary over a wide range of frequency. Thus, time-averaged spectral estimators would compute a much wider band spectrum. These signals are therefore quite useful for wideband jamming, for example. Notched filters in the jamming band are of little use, because the hopped signal's instantaneous frequency remains in the passband of the notched filter for a very small portion of the time. Attempts to cover more of the components of the total jamming spectrum with several notches or a wider stopband filter will seriously degrade a processor's performance on the signal being jammed because the filters suppress too much of the signal energy.

A better approach would be to track the jammer's frequency hopping and remove only that energy. The frequency hopper appears in the CWT as short ridges of limited time duration and at different scales corresponding to the hopping pattern. The locations of these ridges contain valuable signal intelligence for the design of an antijam processor, such as the locations of the instantaneous frequency centers, the durations of each frequency dwell, duty cycles, possible hopping patterns, and hopping sequences for deriving more sophisticated information about the hopping pattern (e.g., an initial register for linear shift register generator estimation).

A second use of the CWT would be as a direct jammer filter. Sampled signals can be used to compute the coefficients of a discrete wavelet transform, that is, the set of coefficients that represent the signal with respect to a basis consisting of scalings and translations of a given mother wavelet. From these coefficients, the original sampled signal can be reconstructed. Both the computation of the coefficients and the reconstruction of the signal from the coefficients can be performed very efficiently with the linear-time Mallat algorithm (see, for example, Ref. 6). By zeroing out those coefficients corresponding to the instantaneous spectrum locations of the jammer, as determined by the thresholding of a CWT surface, for example, and reconstructing the signal from the resulting coefficients, one would effectively have excised the jammer with minimal degradation in the energy of the signal of interest.

Spread-spectrum signals are usually created through a technique known as direct sequence spread spectrum (DSSS). In this method, a narrowband communications signal of any type (digital or analog) is multiplied by the "DSSS chip rate," a phase-shift keyed (PSK) signal having a significantly higher keying rate than the signal bandwidth. (The data stream used to spread the signal is called the sequence of "chips.") The data sequence of the PSK signal is pseudorandom, usually generated with a linear recurrence sequence or some other pseudorandom number generator, to assure a balanced spectrum. The resulting effect on the spectrum of the DSSS signal is the convolution of the spectra from the original signal and the PSK signal. Because the PSK signal has a high chip rate, its spectrum is significantly wider than the original signal and so the resulting signal is spread to a much wider bandwidth. By conservation of energy, the energy can be spread over a wider spectral band only if the instantaneous frequency energy is reduced. Therefore, DSSS signals can be spread to lie below background noise for secure communication. By matching the PSK signal on the receiver end of a communication system, the DSSS signal is returned to its original bandwidth, completing the communications link. If the DSSS signal is multiplied by some PSK signal that is not matched to the keying sequence, the signal is not returned to its original bandwidth and remains spread in frequency. In this way, multiple DSSS signals of different keying sequences can share the same part of the spectrum simultaneously, thus allowing for code division multiple access multiplexing. The detection and characterization of such signals is a difficult problem because DSSS signals often lie below the noise.

One method for their detection uses a signal processing algorithm called a chip rate detector. This algorithm multiplies a DSSS signal with a square wave of frequency equal to the chip rate, offset by one-half wavelength from the chip boundaries of the DSSS's PSK component. The resulting signal has a strong sinusoidal trend at the chip rate, which can then be detected and estimated with a discrete Fourier transform. The problem with this chip rate detector is that it requires knowledge of the chip rate and testing of

various offsets to obtain an offset approximately halfway between DSSS chip boundaries. The CWT surface effectively displays a range of possible chip rates and offsets. By beginning with a mother wavelet that is square-wave-like (the admissibility condition forces the wavelet to have many oscillations), such as the Haar wavelet or its modifications (see Ref. 2), the range of values of the variable a rescales the wavelet and, hence, produces a continuously changing hypothetical chip rate. Near the true chip rate of the DSSS signal, the two signals (the DSSS and the scaled mother wavelet) slowly slide past each other, so that the two signals are offset by one-half wavelength for a time. By computing a one-dimensional Fourier transform of each row of a CWT surface (i.e., at constant values of the scale variable a), one should detect a peak in the transformed surface at the chip rate.

Other signals of interest for intelligence and electronic warfare applications of the CWT include radar signals with pulse modulation and modulation of the pulse repetition interval. Again, such signals exhibit changes in instantaneous spectrum over time that are important to their function as military radar signals and that are displayed as features in a CWT surface. Researchers are currently pursuing these applications of the CWT, which may lead to new and exciting signal processing techniques and algorithms based on this powerful tool from applied mathematics.

REFERENCES

1. Sadowsky, J., "The Continuous Wavelet Transform: A Tool for Signal Analysis and Understanding," *Johns Hopkins APL Tech. Dig.* **15**(4), 306–318 (1994).
2. Daubechies, I., *Ten Lectures on Wavelets*, SIAM Press, Philadelphia (1992).
3. Meyer, Y., *Ondelettes et Opérateurs I: Ondelettes*, Hermann Editeurs des Sciences et des Arts, Paris (1990).
4. Kronland-Martinet, R., Morlet, J., and Grossman, A., "Analysis of Sound Patterns Through Wavelet Transforms," *Int. J. Pattern Recognition and Artificial Intelligence*, Special Issue on Expert Systems and Pattern Analysis, **1**(2), 97–126 (1989).
5. Dutilleux, P., "An Implementation of the 'Algorithme à Trous' to Compute the Wavelet Transform," in *Wavelets: Time-Frequency Methods and Phase Space*, J. Combes, A. Grossmann, and Ph. Tchamitchian (eds.), Springer-Verlag, New York, pp. 298–304 (1989).
6. Mallat, S., "Multiresolution Approximation and Wavelets," *Trans. Am. Math. Soc.* **315**, 69–88 (1989).

ACKNOWLEDGMENT: The author wishes to thank Kishin Moorjani, whose suggestions and encouragement were important factors in the production of this article, and Stan Favin, who provided the seismic data used in the analysis.

THE AUTHOR

JOHN SADOWSKY is the assistant supervisor of the Advanced Signal and Information Processing Group of the M. S. Eisenhower Research Center at APL. He earned a B.A. in mathematics from The Johns Hopkins University in 1971 and a Ph.D. in mathematics from the University of Maryland in 1980. In the Research Center, his work includes advanced transforms and algorithms for signal processing and, in collaboration with researchers in The Johns Hopkins University Schools of Engineering, Arts and Sciences, and Medicine, he coordinates the sensory engineering program. Before joining APL in August 1989, he was a principal scientist and supervisor at the Systems Engineering and Development Corporation (now a division of the Essex Corporation), where he led programs to develop advanced transforms, algorithms, and architectures for processing systems for electronic and communication signals intelligence. His e-mail address is John.Sadowsky@jhuapl.edu.