WAVELETS

Wavelets have been found to be very useful in many scientific and engineering applications, including signal processing, communication, video and image compression, medical imaging, and scientific visualization. The concept of wavelets can be viewed as a synthesis of ideas that originated during the last several decades in engineering, physics, and pure mathematics. Although wavelets are a rather simple mathematical tool with a great variety of possible applications, the subject of wavelets is often introduced at a high level of mathematical sophistication. The goal of this article is to develop a basic understanding of wavelets, their origin, and their relation to scaling functions, using the theory of multiresolution analysis.

HISTORICAL PERSPECTIVE (1–14)

Prior to the 1930s, the main tools of mathematics for solving scientific and engineering problems traced back to Joseph Fourier (1807) with his theory of frequency analysis. He proposed that any 2π-periodic function $f(t)$ can be represented by a linear combination of cosines and sines:

$$f(t) = a_0 + \sum_{k=1}^{\infty} (a_k \cos kt + b_k \sin kt) \tag{1}$$

The coefficients a_0, a_k, b_k are the *Fourier coefficients* of the series and are given by

$$a_0 = \frac{1}{2\pi} \int_0^{2\pi} f(t)\,dt \tag{2a}$$

$$a_k = \frac{1}{2\pi} \int_0^{2\pi} f(t) \cos kt\,dt \tag{2b}$$

$$b_k = \frac{1}{2\pi} \int_0^{2\pi} f(t) \sin kt\,dt \tag{2c}$$

After 1807, mathematicians gradually were led from the notion of frequency analysis to the notion of scale analysis—that is, analyzing $f(t)$ by creating a mathematical structure that varies in scale. A. Haar, in his thesis (1909), was the first to mention using wavelets. An important property of the wavelets he used is that they have compact support, which means that the function vanishes outside a finite interval. Unfortunately, Haar wavelets are not continuously differentiable, which limits their application.

From the 1930s to the 1960s, several groups, working independently, researched the representation of functions using scale-varying basis functions. By using one such function, the *Haar basis function,* Paul Levy investigated Brownian motion and thereby laid the foundation for the modern theory of random processes. He found that the Haar basis function is superior to the Fourier basis functions for studying small and complicated details in Brownian motion.

Also during the 1930s, research was done by Littlewood, Paley, and Stein on computing the *energy* of a function $f(t)$:

$$\text{energy} = \frac{1}{2}\int_0^{2\pi} |f(t)|^2\,dt \tag{3}$$

Their computation produced different results when the energy was concentrated around a few points and when it was distributed over a larger interval. This observation disturbed many scientists, because it indicated that energy might not be conserved. Later on, they discovered a function that can both vary in scale and conserve energy at the same time, when computing the functional energy. David Marr applied this work in developing an efficient algorithm for numerical image processing using wavelets in the early 1980s.

Between 1960 and 1980, the mathematicians Guido Weiss and Ronald Coifman studied the simplest elements of a function space, called *atoms,* with the goals of finding the atoms for a common function and finding the *construction rules* that allow the reconstruction of all the elements of the function space using these atoms. In 1980, Grossman and Morlet recast the study of quantum physics in the context of wavelets using the concept of frames. Morlet introduced the term "wavelet" as an abbreviation of "wavelet of constant shape." These new insights into using wavelets provided a totally new way of thinking about physical reality.

In the summer of 1985, Stephane Mallat applied wavelets to his work in digital signal processing. He discovered a relationship between quadrature mirror filters, the pyramid algorithm, and orthonormal wavelet bases. Inspired by these results, Y. Meyer constructed the first nontrivial wavelets. Unlike the Haar wavelets, the Meyer wavelets are continuously differentiable; however, they do not have compact support. In the early 1990s, Ingrid Daubechies used Mallat's work to construct a set of orthonormal wavelet basis functions that are perhaps the most elegant, and have become the cornerstone of wavelet applications today.

The development of wavelets is an emerging field comprising ideas from many different fields. The foundations of wavelet theory have been completed, and current research is in the refinement stage. The refinement involves generalizations and extensions of wavelets, such as extending wavelet packet techniques.

The future of wavelets depends on the possibility of applications. Wavelets have so far been limited in practical applications by their lack of compact support.

FOURIER ANALYSIS

Time-series data have traditionally been analyzed in either the time or the frequency domain. Fourier analysis is quite useful in identifying frequency components of a signal, but it cannot describe when those frequency components occurred, since it lacks time resolution. This is particularly important for signals with time-varying frequency content, as in human speech and video images.

The Fourier transform is characterized by the ability to analyze a signal in the time domain for its frequency content. The transform works by first translating a function in the time-domain into a function in the frequency domain. The signal can then be analyzed for its frequency content, because the Fourier coefficients of the transformed function represent the contribution of each sine and cosine function at each frequency. An Inverse transform does the opposite by transforming data from the frequency domain into the time domain. Although the time-series data can have infinitely many sample points, in practice one deals with a finite time interval using a sampling mechanism. The discrete Fourier transform (DFT) estimates the Fourier transform of a function from a finite number of its sampled points. The sampled points are supposed to be typical of what the signal looks like at all other times. The DFT has symmetry properties almost exactly the same as the continuous Fourier transform. To approximate a function by samples, and to approximate the Fourier integral by the DFT, requires multiplication by a matrix which involves on the order of n^2 arithmetic operations. However, if the samples are uniformly spaced, then the Fourier matrix can be factored into a product of just a few sparse matrices, and the resulting factors can be applied to a vector in a total on the order of $n \log n$ arithmetic operations. This technique is the so-called fast Fourier transform (FFT).

WAVELET VERSUS FOURIER TRANSFORM

The FFT and the discrete wavelet transform (DWT) are both linear operations that generate a data structure containing $\log_2 n$ segments of various lengths, usually filling it and transforming it into a different data vector of length 2^n.

The mathematical properties of the matrices involved in the transforms are similar as well. The inverse transform matrix for both the FFT and the DWT is the transpose of the original. As a result, both transforms can be viewed as a rotation in function space to a different domain (1). For the FFT, this new domain contains basis functions that are sines and cosines. For the wavelet transform, this new domain contains more complicated basis functions called wavelets, mother wavelets, or analyzing wavelets.

The two transforms have another similarity. The basis functions are localized in frequency, making mathematical tools such as power spectra (how much power is contained in a frequency interval) and scalegrams useful at picking out frequencies and calculating power distributions.

The most interesting dissimilarity between these two kinds of transforms is that individual wavelet functions are localized in space. Fourier sine and cosine functions are not. This localization in space, along with wavelets' localization in frequency, makes many functions and operators using Wavelets *sparse* when transformed into the wavelet domain. This sparseness, in turn, makes wavelets useful for a number of applications such as data compression, feature detection in images, and noise removal from time series.

One way to see the time–frequency resolution difference between the two transforms is to look at the basis–function coverage of the time–frequency plane (7,14). Figure 1 shows a windowed Fourier transform, where the window is simply a square wave. The square-wave window truncates the sine or cosine function to particular width. Because a single window is used for all frequencies in the WFT, the resolution of the analysis is the same at all locations in the time–frequency plane. An advantage of wavelet transforms is that the windows vary. In order to isolate signal discontinuities, one would like to have some very short basis functions. At the same time, in order to obtain detailed frequency analysis, one would like to have some very long basis functions. A way to achieve this is to have short high-frequency basis functions and long low-frequency ones. This happy medium is exactly what you get with wavelet transforms. Figure 2 shows the coverage in the time–frequency plane with one wavelet function, the Daubechies wavelet.

WAVELET ANALYSIS

Wavelets are a class of functions used to localize a given function in both space and scaling. The basic construction of wavelets is based on the family of *mother wavelets* ϕ, consisting of almost any function defined in a finite interval. *Daughter wavelets* are then formed by translation (b) and scaled contraction (a). An individual wavelet can be defined by

$$\phi_{[a,b]}(t) = |a|^{-1/2}\phi\left(\frac{t-b}{a}\right) \tag{4}$$

An example of a typical Wavelet is $a = 2^{-j}$ and $b = ak$:

$$\phi_{[j,k]}(t) = 2^{j/2}\phi(2^j t - k) \tag{5}$$

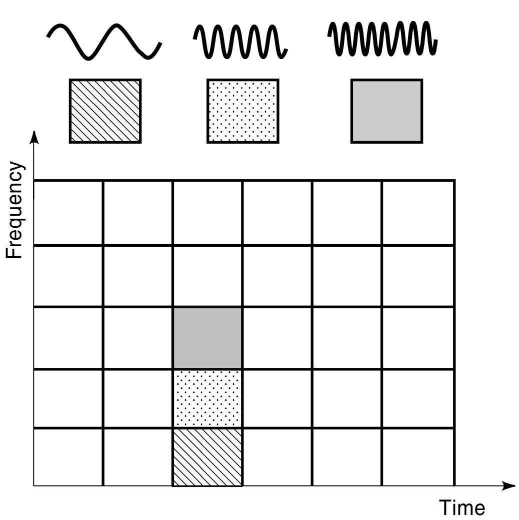

Figure 1. Fourier basis functions, time–frequency tiles, and coverage of the frequency plane (1).

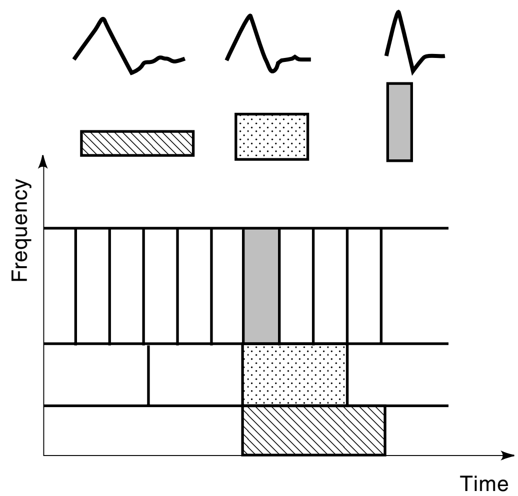

Figure 2. Daubechies wavelet basis functions, time–frequency tiles, and coverage of the frequency plane (1).

More precisely, it can be written as a linear combination of scaling functions $\phi(2^j t - k)$, which are 2^{-j}-scaled and $k/2^j$-translated version of $\phi(t)$, as follows:

$$\phi_{[j,k]}(t) = \sum_{k=-\infty}^{\infty} p_k \phi(2^j t - k) \tag{6}$$

This is often referred to as the two-scale function for the scaling functions, and the sequence $\{p_k\}$ is called the two-scale sequence of ϕ.

Given a scaling function ϕ, the basic assumption of multiresolution analysis is that there exists another function ψ, called a wavelet, such that ϕ forms a basis for the reconstruction of ψ, analogously to the relation for a scaling function. The reconstruction of the wavelet can be expressed as follows:

$$\psi_{[j,k]}(t) = \sum_{k=-\infty}^{\infty} q_k \phi(2^j t - k) \tag{7}$$

The two-scale relations in Eqs. (6) and (7) together are called the reconstruction relations. Since both $\phi(2x)$ and $\phi(2x - 1)$ are in the subspace of the analyzing function, Eqs. (6) and (7) can be combined to form the *reconstruction relations:*

$$\phi(2t - l) = \sum_{k=-\infty}^{\infty} [a_{l-2k}\phi(t-k) + b_{l-2k}\psi(t-k)] \tag{8}$$

Although many classes of wavelets exist, there are two main typical classes:

1. The wavelets defined on the real line, such as the Haar wavelet, the Daubechies Wavelet, and B-spline wavelets in general. Linear, quadratic, and cubic wavelets have been studied.
2. Multiwavelets, or waveletlike functions defined on finite intervals, such as Legendre wavelets and flatlet wavelets.

The interested reader is referred to Refs. 11–14.

Haar Wavelet

The Haar scaling function and Haar wavelet are a very simple example to illustrated many nice properties of scaling functions and wavelets, and are of practical use as well. The Haar scaling function is defined by

$$\phi(t) = \begin{cases} 1 & \text{for} \quad 0 \le x \le 1 \\ 0 & \text{otherwise} \end{cases} \tag{9}$$

The two-scale relation can be expressed in a summation as shown below:

$$\phi_2^H(t) = \sum_{k=0}^{2} p_k \phi(2t-k) \tag{10a}$$

or

$$\phi(t) = \phi(2t) + \phi(2t-1) \tag{10b}$$

The Haar wavelet corresponding to the Haar scaling function is given by

$$\psi(t) = \begin{cases} 1 & \text{for} \quad 0 \le x \le \frac{1}{2} \\ -1 & \text{for} \quad \frac{1}{2} \le x \le 1 \\ 0 & \text{otherwise} \end{cases} \tag{11}$$

We can easily construct the two-scale relation for the Haar wavelet as

$$\psi(2t) = \phi(2t) - \phi(2t-1) \tag{12}$$

The two-scale relations in Eq. (10b) express $\phi(t)$ in terms of $\phi(2t)$ and $\phi(2t) - 1$), while the two-scale relations in Eq. (12) for Haar wavelets express $\psi(t)$ in terms of $\phi(2t)$ and $\phi(2t) - 1$). The reconstruction relations can be written in the matrix form

$$\begin{bmatrix} \phi(t) \\ \psi(t) \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \begin{bmatrix} \phi(2t) \\ \phi(2t-1) \end{bmatrix} \tag{13}$$

The decomposition relations are easily derived by just inverting the reconstruction relations as follows:

$$\begin{bmatrix} \phi(2t) \\ \pi(2t-1) \end{bmatrix} = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & -\frac{1}{2} \end{bmatrix} \begin{bmatrix} \phi(t) \\ \psi(t) \end{bmatrix} \tag{14}$$

Daubechies Wavelets

Another example of wavelets defined on the real line is Daubechies wavelets. The Daubechies scaling function ϕ_3^{D} is defined by the following relation:

$$\phi_3^{\mathrm{D}}(t) = \sum_{k=0}^{3} p_k \phi(2t-k) \tag{15}$$

where two-scale sequence $\{p_k\}$ are

$$\{p_1, p_2, p_3, p_4\} = \left\{ \frac{1+\sqrt{3}}{4}, \frac{3+\sqrt{3}}{4}, \frac{3-\sqrt{3}}{4}, \frac{1-\sqrt{3}}{4} \right\} \tag{16}$$

In general, the two-scale sequence for any scaling functions has the property

$$\sum_k p_{2k} = \sum_k p_{2k+1} = 1 \tag{17}$$

There is no closed form for ϕ_3^{D}; however, one can use numerical computation to draw the graph of ϕ_3^{D} (1,2).

The Haar wavelet is the simplest one. It has found many applications. However, it has the drawback of discontinuity. It consists entirely of rectangular functions and cannot reproduce even linear functions smoothly in finite series for practical use. On the other hand, B-spline wavelets have higher continuity than Haar wavelets. They are more suitable for representing any continuous function. However, the complications of calculating its wavelet decomposition and reconstruction relation coefficients have limited its usefulness.

SUBBAND CODING (7,12–15)

One of the main applications of subband coding is compression. A key concept in signal analysis is that of *localization* in time and frequency. Another important intuitive concept is that of *multiresolution,* or the idea that one can consider a signal at different levels of resolution. These notions are particularly evident in image processing and computer vision, where coarse versions of images are often used as a first approximation in computational algorithms. In signal processing, a low-pass and subsampled version is often a good coarse approximation for many real-life signals. This intuitive paradigm leads to the mathematical framework for wavelet constructions (12). The wavelet decomposition is a successive approximation method that adds more and more projections onto *detail* spaces, or spaces spanned by wavelets and their shifts at different scales.

In addition, this multiresolution approximation is well suited to many applications. That is true in cases where successive approximation is useful, for example, in browsing through image databases, as is done for instance on the World-Wide Web. Rather than downloading each full image, which would be time-consuming, one only needs to download a coarse version, which can be done relatively fast. Then, one can fetch the rest, if the image seems of interest. Similarly, for communication applications, multiresolution approximation leads to transmission methods where a coarse version of a signal is better protected against transmission errors than the detailed information. The assumption is that the coarse version is probably more useful than the detail.

There are many techniques for image coding. Subband coding is the most successful today. Pyramid coding is effective for high-bit-rate compression, and transform coding based on the discrete cosine transform has become the JPEG standard. Subband coding using wavelets (the tree-structured filterbank approach) avoids blocking at medium bit rates, because its basis functions have variable length. It uses an adapted basis (the transformation depends on the signal). Long basis functions represent flat background (low frequency), and short basis functions represent regions with texture. This feature is good for image enhancement, image edge detection, image classification, videoconferencing, video on demand, tissue and cancer cell detection (16), and so on. And due to its adapted basis functions, one can also develop a set of algorithms for adaptive filtering systems (7).

WAVELET APPLICATIONS

This section discusses the applications of wavelets in digital communication systems and in image and video compression. We will present a brief description of how wavelets can be used (1) to improve communication efficiency in digital communication systems, and (2) to remove redundant information (e.g., spatial redundancy, spectral redundancy, and temporal redundancy) in image and video compression for an FBI fingerprint application. Other applications in computer and human vision and for denoising noisy data are also described.

Wavelets in Digital Communications

The channel coding that follows source coding is designed to reintroduce—in a controlled manner—a prescribed level of redundancy into the source-coded streams to mitigate the anticipated effects of the channel, and best performance is achieved when the redundancy is tailored to the specific characteristics of the channel. Although channel coding and modulation applications have received comparatively little attention to date, wavelet theory has an important complementary role to play in this aspect of the communications problem. Wavelet representation can be exploited to develop systems for reliable transmission over specific channels such as satellite or mobile wireless channels or a combination of both. The problem of multiple-user communication over a wireless channel is of greatest interest in practice. A wireless communication channel is subject to mutual interference among users and fading due to time-varying multipath propagation. To combat the interference and fading effects, a popular multirate modulation technique, referred to as *spread-spectrum code division multiplexing access* (SS CDMA) has been used widely in the industry (17–22). This subsection focuses on the use of wavelets for SS CDMA digital communication systems. In addition, a brief description of fractal modulation using wavelets is described.

Single-Scale Covert Communication Waveform. The idea of wavelet-based SS CDMA was proposed in Ref. 18. The proposed algorithm is summarized below:

- k information bits are grouped together to form an M-ary set of sequences, where

$$M = 2^k \tag{18}$$

- A pseudonoise (PN) code selects a sequence (out of 2^k sequences) for each input information bit generated by the inverse discrete wavelet transform (IDWT).
- The coded data are sent as binary phase-shift keying (BPSK).
- The receiver reverses the above steps, that is, it performs BPSK demodulation, DWT, and data extraction.

Note that only one of 2^J wavelet coefficients is used for BPSK data and the remaining coefficients are unused, where J denotes the number of stages in the wavelet transform. If we let L be the length of the filter band of the DWT, then the processing gain equals the composite impulse response length, that is,

$$G = (2^J - 1)(L - 1) - 1 \tag{19}$$

This wavelet-based SS technique is equivalent to a single-scale covert communication waveform.

Multiple-Scale Covert Communications Waveform. A multiple-scale covert waveform employing a wavelet-transform-domain SS was proposed in Ref. 19. This technique is summarized as follows:

- Encode the data into an M-ary alphabet.
- Apply the PN code to the encoded data.
- Serial-to-parallel demultiplex the resultant into scale streams.
- Represent the SS chips as wavelet coefficients.
- Create a waveform by IDWT.
- Receiver reverses the above steps.

Figure 3 (20, p. RO-72) illustrates the multiple-scale SS system.

Wavelets for Code Division Multiplexing Access. The above waveforms are suitable for multiple users where each user has the same waveform algorithm with fixed system parameters and independent PN drivers for the pseudorandom parameters. The design of a CDMA using wavelets has been addressed in Ref. 21, which concerns itself with the three key issues:

- *Channel Capacity.* The number of supportable users per unit bandwidth at a fixed bit error rate (BER) and bit signal-to-noise ratio (SNR), E_b/N_0
- *Near–Far Problem.* The effect of large interference due to a nearby user (this problem is discussed in another article in this encyclopedia)
- *Implementation.* Computational complexity

Regarding the channel capacity, the use of wavelets will lower the required bit SNR. This means that the multiple-access (MA) noise decreases on using a proper wavelet transform. Since the wavelet transform is chosen to allow the user to operate at smaller bandwidth, one can impose tighter

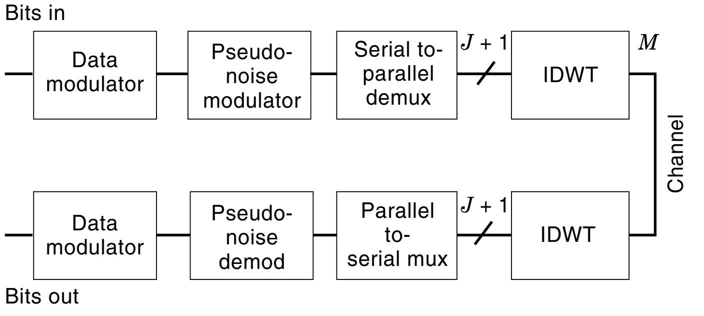

Figure 3. An illustration of a multiple-scale covert communications system using wavelet transforms (4, p. RO-72).

bandpass filtering at the transmitter. This means improved bandwidth efficiency. In addition to these features, modulation using wavelets occupies all available degrees of freedom in amplitude and phase.

Concerning the near–far problem, the Gaussian character of wavelet-based waveform makes the interference noiselike at all levels, and the near–far effects are perfectly predictable based on power levels.

Finally, the construction of a wavelet-based signal uses a finite number of subsequences, which allows the transmitter to be implemented with only load operations.

This sub-subsection briefly explains the use of the single-scale and multiple-scale waveforms in CDMA.

Single-Scale CDMA. According to the sub-subsection "Single-Scale Covert Communication Waveform" above, the single-scale CDMA waveform is generated using the following algorithm:

- For each input k information bits we send one of the M complex-valued sequences a_n, where the index n denotes the nth of the M sequences, and M is given in Eq. (18).
- M complex-valued sequences are derived from the M-ary wavelet coefficient matrix. Note that these sequences have approximately Gaussian distribution.
- The set of sequences changes pseudorandomly for each transmission.

If we let the transform length N be k times the alphabet size M,

$$N = kM \tag{20}$$

and the chip pulse with duration T be $P(t)$, then the signal model for the single-scale CDMA can be written as

$$S(t) = \sum_{n=1}^{N(M-1)} a_n P(t - nT) \tag{21}$$

For unfiltered BPSK or quaternary PSK (QPSK), the pulse shape is square. For filter BPSK/QPSK or wavelet use, the pulse shape is equivalent to the truncated approximation of a sinc pulse, defined as follows:

$$\mathrm{sinc}(t) = \frac{\sin t}{t} \tag{22}$$

It should be noted here that the multiple user waveforms described above possess the following properties:

- Common structure with independent pseudorandom drivers
- Statistical orthogonality

Let us define U as the number of supportable users, G as the processing gain given by Eq. (19), $(E_b/N_0)_{nom}$ as the nominal bit SNR without MA noise, $(E_b/N_0)_{aci}$ as the actual bit SNR with MA noise and intersymbol interference, and

$$\alpha = \frac{(E_b/N_0)_{act}}{(E_b/N_0)_{nom}} \tag{23}$$

The total number of supportable users for unfiltered QPSK is found to be (21)

$$U = \frac{3(1-\alpha^{-1})G}{4(E_b/N_0)_{nom}} + 1 \tag{24}$$

whereas for filtered QPSK and a wavelet-based waveform it is

$$U = \frac{(1-\alpha^{-1})G}{(E_b/N_0)_{nom}} + 1 \tag{25}$$

A plot of Eqs. (24) and (25) is shown in Fig. 4 (20, p. RO-90) for a QPSK signal with 16-ary simplex. This figure shows that the wavelet SS can support more users than the filtered QPSK.

Multiple-Scale CDMA. According to the sub-subsection "Multiple-Scale Covert Communications Waveform" above, the multiple-scale CDMA waveform uses the wavelet transform domain structure shown in Fig. 3. In this figure, each user signal is driven by an independent PN modulator. The resulting signals will be statistically orthogonal (as good as the PN sequence orthogonality). The performance of the multiple-scale CDMA is as good as that of the single-scale, that is, it has the same channel capacity, BER, interference tolerance, and so on. However, the computational complexity for multiple-scale CDMA is less than for single-scale CDMA. Multiple-scale systems can acquire the signal without a training sequence (20,21), and they can be used for low probability of interception (LPI) and low probability of detection (LPD) networks.

Fractal Modulation. There are some noisy channels whose key characteristic is that the channel is open for some finite but unknown time interval, during which it has some finite but unknown bandwidth. Such models are useful for a range of wireless and secure communications applications, as well as for broadcast applications in which information is being

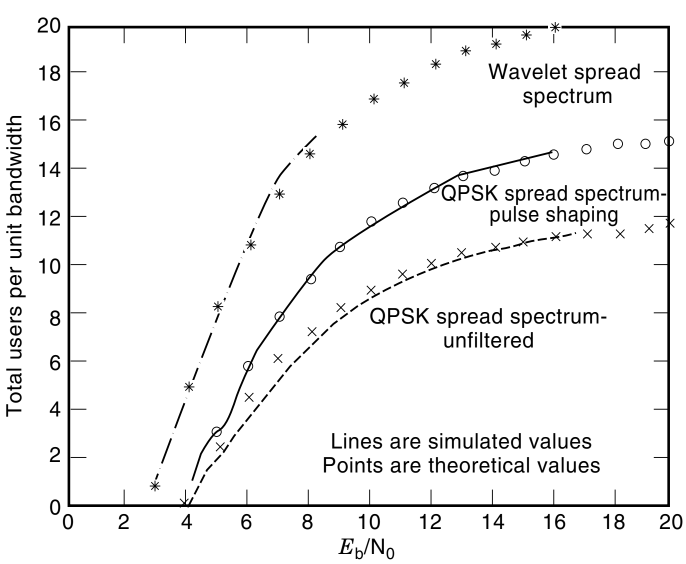

Figure 4. Supportable number of users per unit bandwidth as a function of bit SNR for various QPSK waveforms (4, p. RO-90).

transmitted to receivers whose front ends have different bandwidths and processing capabilities. A wavelet-based modulation model, referred to as *fractal modulation,* which makes efficient use of iterated multirate filter banks, has been shown to provide a novel diversity strategy for communication over such channels (23). The essence of this scale diversity involves dividing the available transmit spectrum into multiple adjacent octave-spaced bands, and modulating periodic extensions of the symbol stream into these bands at corresponding rates. That is, the information stream is present on all time scales, providing a novel and efficient form of diversity for such applications.

FBI Fingerprint Compression

Between 1924 and today, the U.S. Federal Bureau of Investigation has collected about 30 million sets of fingerprints (24). The archive consists mainly of inked impressions on paper cards. Facsimile scans of the impressions are distributed among law enforcement agencies, but the digitization quality is often low. Because a number of jurisdictions are experimenting with digital storage of the prints, incompatibilities between data formats have recently become a problem. This problem led to a demand in the criminal-justice community for a digitization and compression standard. In 1993, the FBI's Criminal Justice Information Services Division developed standards for fingerprint digitization and compression in cooperation with the National Institute of Standards and Technology, Los Alamos National Laboratory, commercial vendors, and criminal-justice agencies (25).

Let us put the problem of storing the data of digital fingerprints in perspective. Fingerprint images are digitized at a resolution of 500 pixels/in. (200 pixels/cm) with 256 levels of gray-scale information per pixel. A single fingerprint is about 700,000 pixels and needs about 0.6 Mbyte of storage. A pair of hands, then, requires about 6 Mbytes of storage. So digitizing the FBI's current archive would result in about 200 terabytes of data. Obviously, data compression is important to bring these numbers down. The data compression standard *wavelet / scalar quantization* (WSQ) implements a hand-tuned custom wavelet basis developed after extensive testing on a collection of fingerprints. The best compression ratio achieved with these wavelets is 26 : 1.

Computer and Human Vision

In the early 1980s, David Marr began work at MIT's Artificial Intelligence Laboratory on artificial vision for robots. He is an expert on the human visual system, and his goal was to learn why the first attempts to construct a robot capable of understanding its surroundings were unsuccessful. Marr believed that it was important to establish scientific foundations for vision, and that while doing so, one must limit the scope of investigation by excluding everything that depends on training, culture, and so on, and focus on the mechanical or involuntary aspects of vision. This low-level vision is the part that enables us to recreate the three-dimensional organization of the physical world around us from the excitations that stimulate the retina. Marr asked these questions:

- How is it possible to determine the contours of objects from variations in their light intensity?
- How is it possible to sense depth?
- How is motion sensed?

He then developed working algorithmic solutions to answer each of these questions. Marr's theory was that the human visual system has a complicated hierarchical structure that involves several layers of processing. At each processing level, the retinal system provides a visual representation that scales progressively in a geometrical manner. His arguments hinged on the detection of intensity changes. He theorized that intensity changes occur at different scales in an image, so that their optimal detection requires the use of operators of different sizes. He also theorized that sudden intensity changes produce a peak or trough in the first derivative of the image. These two hypotheses require that a vision filter have two characteristics: it should be a differential operator, and it should be capable of being tuned to act at any desired scale. Marr's operator is referred to today as a Marr wavelet.

Denoising Noisy Data

In diverse fields, from planetary science to molecular spectroscopy, scientists are faced with the problem of recovering a true signal from incomplete, indirect, or noisy data. Can wavelets help solve this problem? The answer is certainly yes, through a technique, called *wavelet shrinkage and thresholding,* that David Donoho of Stanford University has worked on for a number of years (26). The technique works in the following way. When you decompose a data set using wavelets, you use filters that act as *averaging* filters, and others that produce *details.* Some of the resulting wavelet coefficients correspond to details in the data set. If the details are small, they might be omitted without substantially affecting the main features of the data set. The idea of *thresholding,* then, is to set to zero all coefficients that are less than a particular threshold. The remaining coefficients are used in an inverse wavelet transformation to reconstruct the data set.

REMARKS

There are many applications waiting for wavelet techniques to improve their usefulness beside those mentioned above. Examples are speech compression in mobile communication and digital answering machines; audio compression in digital broadcasting, HDTV, VSAT, storage devices, multimedia, high-fidelity audio, and music; and ECG heart waveform monitoring systems and archives for cardiologists.

In signal processing, wavelets make it possible to recover weak signals from noise. This has proven useful especially in the processing of X-ray and magnetic resonance images in medical applications. Images processed in this way can be cleaned up without blurring or muddling the details (16).

In Internet communications, wavelets have been used to compress images to a greater extent than is generally possible with other methods. In some cases, a wavelet-compressed image can be as small as about 25% of the size of similar-quality images using the more familiar JPEG method. Thus, for example, a photograph that requires 200 kbyte and takes a minute to download in JPEG format might require only 50 kbyte and take 15 s to download in wavelet-compressed format.

Wavelet compression works by analyzing an image and converting it into a set of mathematical expressions that can then be decoded by the receiver. Wavelet compression is not yet widely used on the Web. The most common compressed image formats remain the GIF, used mainly for drawings, and JPEG, used mainly for photographs.

BIBLIOGRAPHY

1. A. Graps, An introduction to wavelets, in *IEEE Computational Science and Engineering, Summer 1995,* Vol. 2, No. 2, Los Alamitos, CA: IEEE Computer Society Press, 1995.
2. C. K. Chui, *An Introduction to Wavelets,* New York: Academic Press, 1992.
3. C. K. Chui (ed.), *Wavelets: A Tutorial in Theory and Application,* New York: Academic Press, 1992.
4. I. Daubechies, Orthonormal bases of compactly supported wavelets, *Commun. Pure Appl. Math.,* **41**: 909–996, 1988.
5. I. Daubechies, *Ten Lectures on Wavelets,* CBMS–NSF Regional Series in Applied Mathematics, Philadelphia, PA: SIAM, 1992.
6. I. Daubechies, Where do wavelets come from?—A personal point of view, *Proc. IEEE,* **84**: 510–513, 1996.
7. G. Strang and T. Nguyen, *Wavelets and Filter Banks,* Wellesley-Cambridge Press, 1996.
8. R. K. Young, *Wavelet Theory and its Applications,* Norwell, MA: Kluwer, 1993.
9. G. Beylkin, R. Coifman, and V. Rokhlin, Fast wavelet transforms and numerical algorithms I, *Commun. Pure Appl. Math.,* **44**: 141–183, 1991.
10. G. Beylkin, R. Coifman, and V. Rokhlin, Wavelets in numerical analysis, in M. B. Ruskaied (ed.), *Wavelets and Their Applications,* Boston: Jones and Bartlett, 1992.
11. M. Ueda and S. Lodha, Wavelets: An elementary introduction and examples, UCSC-CRL 94-47, 1995.
12. S. G. Mallat, A theory for multiresolution signal decomposition: The wavelet representation, *Commun. Pure Appl. Math.,* **41**: 674–693, 1988.
13. K. Ramchandran, M. Vetterli, and C. Herley, Wavelets, subband coding, and best bases, *Proc. IEEE,* **84**: 541–560, 1996.
14. M. Vetterli and C. Herley, Wavelets and filter banks: Theory and design, *IEEE Trans. Signal Process.,* **40**: 2207–2232, 1992.
15. K. Cinkler, Very low bit-rate wavelet video coding, *IEEE J. Selected Areas Commun.,* **16**: 4–11, 1998.
16. M. Unser and A. Aldroubi, A review of wavelets in biomedical applications, *Proc. IEEE,* **84**: 1996, 626–638.
17. C. E. Cook et al. (eds.), *Spread-Spectrum Communications,* New York: IEEE Press, Wiley, 1983.
18. M. J. Medley, G. Saulnier, and P. Das, Applications of the wavelet transform in spread-spectrum communications systems, *Proc. NJIT Symp. Appl. Subbands and Wavelets,* March 1994, in *Proc. SPIE,* **2242**: 54–68, 1994.
19. R. S. Orr, C. Pike, and M. J. Lyall, Wavelet transform domain communications, *Wavelet Applications II, Proc. SPIE,* **2491**: 271–282, 1995.
20. R. S. Orr and P. P. Gandhi, Wavelets in communication systems, T06, *ICC'96 Tutorials and Workshop,* June 23, 1996.
21. R. S. Orr, C. Pike, and M. Bates, Multiple access communications using wavelet technology, *Wavelet Applications II, Proc. SPIE* **2491**: 283–294, 1995.
22. G. W. Wornell, Emerging applications of multirate signal processing and wavelets in digital communications, *Proc. IEEE,* **84**: 586–603, 1996.
23. G. W. Wornell and A. V. Oppenheim, Wavelet-based representation for a class of self-similar signals with application to fractal modulation, *IEEE Trans. Inf. Theory,* **38**: 785–800, 1992.
24. V. Wickerhauser, *Adapted Wavelet Analysis from Theory to Software,* Boston: A. K. Peters, 1994, pp. 213–214.
25. J. Bradley, C. Brislawn, and T. Hopper, The FBI wavelet/scalar quantization standard for gray-scale fingerprint image compression, Tech. Rep. LA-UR-93-1659, Los Alamos, NM: Los Alamos Ntl. Lab., 1993.
26. D. L. Donoho, Nonlinear wavelet methods for recovery of signals, densities, and spectra from indirect and noisy data, in I. Daubechies (ed.), *Different Perspective on Wavelets,* Proc. Symp. Appl. Math. 47, Providence: Amer Math. Soc., 1993, pp. 173–205.

BOJ N. TRAN
The Boeing Company
TIEN M. NGUYEN
The Aerospace Corporation
MAZEN M. SHIHABI
Hughes Space and Communications Company

WAVELETS AND FILTER BANKS. See WAVELET TRANSFORMS.