6

Fractal Modulation

6.1 INTRODUCTION

There are a number of interesting potential applications for the homogeneous signal theory developed in the Chapter 5. In this chapter, we focus on a particular example as an indication of the direction that some applications may take. In particular, we explore the use of homogeneous signals as modulating waveforms in a communication system [91]. Beginning with an idealized but general channel model, we demonstrate that the use of homogeneous waveforms in such channels is both natural and efficient, and leads to a novel multirate diversity strategy in which data is transmitted simultaneously at multiple rates.

Our problem involves the design of a communication system for transmitting a continuous- or discrete-valued data sequence over a noisy and unreliable continuous-amplitude, continuous-time channel. We must therefore design a modulator at the transmitter that embeds the data sequence $q[n]$ into a signal $x(t)$ to be sent over the channel. At the receiver, a demodulator must be designed for processing the distorted signal $r(t)$ from the channel to extract an optimal estimate of the data sequence $\hat{q}[n]$. The overall system is depicted in Fig. 6.1.

The particular channel we consider has the characteristic that it is "open" for some time interval T, during which it has a particular bandwidth W and signal-to-noise ratio (SNR). This rather basic channel model is a useful one for a variety of settings, and in particular it can be used to capture both characteristics of the transmission medium and constraints inherent in

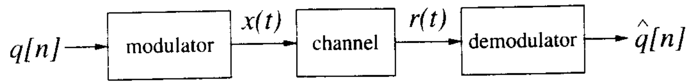

Figure 6.1. A communication system for transmitting a continuous- or discrete-amplitude data sequence $q[n]$ over a noisy and unreliable continuous-amplitude, continuous-time channel.

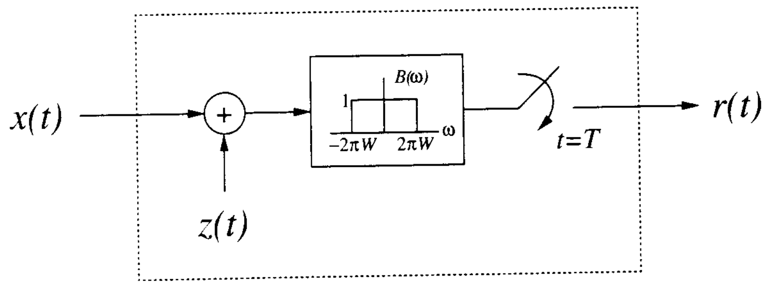

Figure 6.2. The channel model for a typical communications scenario.

one or more receivers. When the noise characteristics are additive, the overall channel model is as depicted in Fig. 6.2, where $z(t)$ represents the noise process.

When either the bandwidth or duration parameters of the channel are known *a priori*, there are many well-established methodologies for designing an efficient and reliable communication system. However, we restrict our attention to the case in which *both* the bandwidth and duration parameters are either unknown or not available to the transmitter. This case, by contrast, has received comparatively less attention in the communications literature, although it arises rather naturally in a range of both point-to-point and multiuser communication scenarios involving, for example, jammed and fading channels, multiple access channels, covert and low probability of intercept (LPI) communication, and broadcast communication to disparate receivers.

In designing a suitable communication system for such channels, we require that the following key performance characteristics be satisfied:

1. Given a duration-bandwidth product $T \times W$ that exceeds some threshold, we must be able to transmit $q[n]$ without error in the absence of noise, i.e., $z(t) = 0$.

2. Given increasing duration-bandwidth product in excess of this threshold, we must be able to transmit $q[n]$ with increasing fidelity in the presence of noise. Furthermore, in the limit of infinite duration-bandwidth product, perfect transmission should be achievable at any finite SNR.

The first of these requirements implies that, at least in principle, we ought to be able to recover $q[n]$ using arbitrarily narrow receiver bandwidth given sufficient duration, or, alternatively, from an arbitrarily short duration segment given sufficient bandwidth. The second requirement implies that we ought to be able to obtain better estimates of $q[n]$ the longer a receiver is able to listen, or the greater the bandwidth it has available. Consequently, the modulation must contain redundancy or diversity of a type that can be exploited for the purposes of improving the reliability of the transmission. As we demonstrate, the use of homogeneous signals for transmission appears to be rather naturally suited to fulfilling both these system requirements.

The minimum achievable duration-bandwidth threshold in such a system is a measure of the efficiency of the modulation. Actually, because the duration-bandwidth threshold $T \times W$ is a function of the length L of the data sequence, it is more convenient to transform the duration constraint T into a symbol rate constraint $R = L/T$ and phrase the discussion in terms of a rate-bandwidth threshold R/W that is independent of sequence length. Then the maximum achievable rate-bandwidth threshold constitutes the *spectral efficiency* of the modulation, which we denote by η. The spectral efficiency of a transmission scheme using bandwidth W is, in fact, defined as

$$\eta = R_{\max}/W$$

where $R_{\max}$ is the maximum rate at which perfect communication is possible in the absence of noise. Hence, the higher the spectral efficiency of a scheme, the higher the rate that can be achieved for a given bandwidth, or, equivalently, the smaller the bandwidth that is required to support a given rate.

When the available channel bandwidth is known *a priori*, a reasonably spectrally efficient, if impractical, modulation of a data sequence $q[n]$ involves expanding the sequence in terms of an ideally bandlimited orthonormal basis. Specifically, with W_0 denoting the channel bandwidth, a transmitter produces

$$x(t) = \sum_n q[n] \sqrt{W_0} \operatorname{sinc}(W_0 t - n)$$

where

$$\operatorname{sinc}(t) = \begin{cases} 1 & t = 0 \\ \dfrac{\sin \pi t}{\pi t} & \text{otherwise} \end{cases}.$$

In the absence of noise, a (coherent) receiver can, in principle, recover $q[n]$ from the projections

$$q[n] = \int_{-\infty}^{\infty} x(t) \sqrt{W_0} \operatorname{sinc}(W_0 t - n)\, dt$$

which can be implemented as a sequence of filter-and-sample operations. Since this scheme achieves a rate of $R = W_0$ symbols/sec using the double-sided bandwidth of $W = W_0$ Hz, it is characterized by a spectral efficiency of

$$\eta_0 = 1 \text{ symbol/sec/Hz}. \tag{6.1}$$

Because the transmitter is assumed to have perfect knowledge of the rate-bandwidth characteristics of the channel, this approach is viable only for those point-to-point channels in which there exists a feedback path from receiver to transmitter. However, we consider a more general case in which we have either a broadcast channel (i.e., a scenario with a single transmitter and multiple receivers), or a point-to-point channel in which no such feedback path is available. In these cases, the approach outlined above does not constitute a viable solution to our communications problem. Indeed, in order to accommodate a decrease in available channel bandwidth, the transmitter would have to be accordingly reconfigured by decreasing the parameter W_0. Similarly, for the system to maintain a spectral efficiency of $\eta_0 = 1$ when the available channel bandwidth increases, the transmitter must be reconfigured by correspondingly increasing the parameter W_0. Nevertheless, while not a solution to the problem of communication without a feedback path, the perfect-feedback solution provides a useful performance baseline in evaluating the solution we ultimately develop for this problem. In the sequel, we therefore refer to this as our benchmark modulation scheme.

A viable solution to the problem of interest requires a modulation strategy that maintains its spectral efficiency over a broad range of rate-bandwidth combinations using a fixed transmitter configuration. A rather natural strategy of this type arises out of the concept of embedding the data to be transmitted into a homogeneous signal. Due to the fractal properties of the transmitted signals, we refer to the resulting scheme as "fractal modulation."

6.2 TRANSMITTER DESIGN: MODULATION

To embed a finite-power sequence $q[n]$ into a bihomogeneous waveform $x(t)$ of degree H, it suffices to consider using $q[n]$ as the coefficients of an expansion in terms of a wavelet-based orthonormal self-similar basis of degree H, i.e.,

$$x(t) = \sum_n q[n]\, \theta_n^H(t)$$

where the basis functions $\theta_n^H(t)$ are constructed according to (5.21). When the basis is derived from the ideal bandpass wavelet, as we will generally assume in our analysis, the resulting waveform $x(t)$ is a power-dominated homogeneous signal whose idealized time-frequency portrait has the form depicted in Fig. 5.3. Consequently, we may view this as a *multirate modulation* of $q[n]$ where in the mth frequency band $q[n]$ is modulated at rate 2^m using a double-sided bandwidth of 2^m Hz. Furthermore, the energy per symbol used in successively higher bands scales by $\beta = 2^{2H+1}$. Using a suitably designed receiver, $q[n]$ can, in principle, be recovered from $x(t)$ at an arbitrary rate 2^m using a baseband bandwidth of 2^{m+1} Hz. Consequently, this modulation has a spectral efficiency of

$$\eta_F = \frac{1}{2} \text{ symbol/sec/Hz.}$$

We emphasize that in accordance with our channel model of Fig. 6.2, it is the baseband bandwidth that is important in defining the spectral efficiency since it defines the highest frequency available at the receiver.

While the spectral efficiency of this modulation is half that of the benchmark scheme (6.1), this loss in efficiency is, in effect, the price paid to enable a receiver to use any of a range of rate-bandwidth combinations in demodulating the data. Fig. 6.3 illustrates the rate-bandwidth tradeoffs available to the receiver. In the absence of noise the receiver can, in principle, perfectly recover $q[n]$ using rate-bandwidth combinations lying on or below the solid curve. The stepped character of this curve reflects the fact that only rates of the form 2^m can be accommodated, and that full octave increases in bandwidth are required to enable $q[n]$ to be demodulated at successively higher rates. For reference, the performance of the benchmark modulation is superimposed on this plot using a dashed line. We emphasize that in contrast to fractal modulation, the transmitter in the benchmark scheme requires perfect knowledge of the rate-bandwidth characteristics of the channel.

Although it considerably simplifies our analysis, the use of the ideal bandpass wavelet to synthesize the orthonormal self-similar basis in our modulation strategy is impractical due to the poor temporal localization in this wavelet. However, we may, in practice, replace the ideal bandpass wavelet with one having not only comparable frequency domain characteristics and better temporal localization, but sufficiently many vanishing moments to ensure that the transmitted waveform is power-dominated as well. Fortunately, there are many suitable wavelets from which to choose, among which are those due to Daubechies [12]. When such wavelets are used, the exact spectral efficiency of the modulation depends on the particular definition of bandwidth employed. Nevertheless, using any reasonable definition of bandwidth, we would expect to be able to achieve, in practice, a spectral efficiency close to $(1/2)$ symbols/sec/Hz with this modulation, and, as a result, we assume $\eta_F \approx 1/2$ in subsequent analysis.

Another apparent problem with fractal modulation as initially proposed is that it requires infinite transmitter power. Indeed, as Fig. 5.3 illustrates,

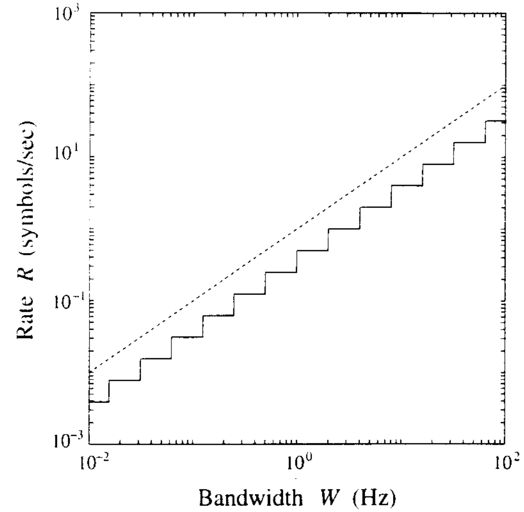

Figure 6.3. Spectral efficiency of fractal modulation. At each bandwidth B, the solid curve indicates the maximum rate at which transmitted data can be recovered in the absence of noise. The dashed curve indicates the corresponding performance of benchmark scheme.

$q[n]$ is modulated into an infinite number of octave-width frequency bands. However, it should be appreciated that in a practical implementation only a finite collection of contiguous bands $\mathcal{M}$ would, in fact, be used by the transmitter. As a result, the transmitted waveform

$$x(t) = \sum_n q[n] \sum_{m \in \mathcal{M}} \beta^{-m/2} \psi_n^m(t) \tag{6.2}$$

would exhibit self-similarity only over a range of scales, and demodulation of the data would be possible at one of only a finite number of rates. In terms of Fig. 6.3, the rate-bandwidth characteristic of the modulation would extend over a finite range of bandwidths chosen to cover extremes anticipated for the system.

The fractal modulation transmitter can be implemented in a computationally highly efficient manner, since much of the processing can be performed using the discrete-time algorithms of Section 5.4. For example, synthesizing the waveform $x(t)$ given by (6.2) for $\mathcal{M} = \{0, 1, \ldots, M-1\}$ involves two stages. In the first stage, which involves only discrete-time processing, $q[n]$ is mapped into M consecutive octave-width frequency bands to obtain the sequence $p^{[M]}[n]$. This sequence is obtained using M iterations of the synthesis algorithm (5.32) with the QMF filter pair $h[n]$, $g[n]$ appropriate to the wavelet basis. The second stage then consists of a discrete- to continuous-time transformation in which $p^{[M]}[n]$ is modulated into the continuous-time frequency spectrum via the appropriate scaling function according to

$$x(t) = \sum_n p^{[M]}[n]\,\phi_n^M(t) = \sum_n p^{[M]}[n]\,2^M \phi(2^M t - n).$$

It is important to point out that because a batch-iterative algorithm is employed, potentially large amounts of data buffering may be required. Hence, while the algorithm may be computationally efficient, it may be considerably less so in terms of storage requirements. However, in the event that $q[n]$ is *finite length*, it is possible that memory-efficient implementations may be constructed as well.

The transmission of finite length sequences using fractal modulation requires some basic modifications to the scheme. In fact, as initially proposed, fractal modulation is rather inefficient in this case, in essence because successively higher frequency bands are increasingly underutilized. In particular, we note from the time-frequency portrait in Fig. 5.3 that if $q[n]$ has finite length, e.g.,

$$q[n] = 0, \qquad n < 0,\ n > L - 1,$$

then the mth band completes its transmission of $q[n]$ and goes idle in half the time it takes the $(m-1)$st band, and so forth. However, finite length messages may be accommodated rather naturally and efficiently by modulating their periodic extensions $q[n \bmod L]$ thereby generating a transmitted waveform

$$x(t) = \sum_n q[n \bmod L]\,\theta_n^H(t)$$

which constitutes a periodicity-dominated homogeneous signal of the type discussed in Section 5.3. If we let

$$\mathbf{q} = \{q[0]\ q[1]\ \cdots\ q[L-1]\}$$

denote the data vector, then the time-frequency portrait associated with this signal is shown in Fig. 6.4. Using this enhancement of fractal modulation, we not only maintain our ability to make various rate-bandwidth tradeoffs at the receiver, but we acquire a certain flexibility in our choice of time origin as well. Specifically, as is apparent from Fig. 6.4, the receiver need not begin demodulating the data at $t = 0$, but may more generally choose a time-origin that is some multiple of LR when operating at rate R. Additionally, this strategy can, in principle, be extended to accommodate data transmission on a block-by-block basis.

The final aspect of fractal modulation that remains to be considered in this section concerns the specification of the parameter H. While H has no effect on the spectral efficiency of fractal modulation, it does affect the power efficiency of the scheme. Indeed, it controls the relative power distribution between frequency bands and, hence, the overall transmitted power

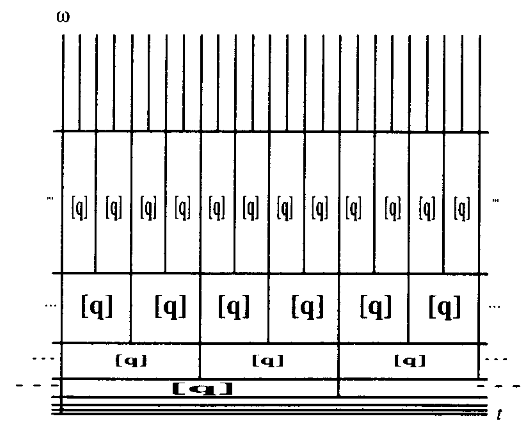

Figure 6.4. A portion of the time-frequency portrait of the transmitted signal for fractal modulation of a finite-length data vector **q**. The case $H = -1/2$ is shown for convenience.

spectrum, which takes the form (5.25) where $\gamma = 2H + 1$. Consequently, the selection of H is important when we consider the presence of additive noise in the channel.

For traditional additive stationary Gaussian noise channels of known bandwidth, the appropriate spectral shaping of the transmitted signal is governed by a "water-filling" procedure [24] [94], which is also the method by which the capacity of such channels is computed [95]. Using this procedure, the available signal power is distributed in such a way that proportionally more power is located at frequencies where the noise power is smaller.

When there is uncertainty in the available bandwidth, the water-filling approach leads to poor worst-case performance. As an example, for a channel in which the noise power is very small only in some fixed frequency band $\omega_L < \omega \le \omega_U$, where $0 < \omega_L < \omega_U < \infty$, the water-filling recipe would locate the signal power predominantly within this band. As a result, the overall SNR in the channel would strongly depend on whether the channel bandwidth is such that these frequencies are passed. By contrast, the distribution of power according to a spectral-matching rule that maintains an SNR that is independent of frequency leads to a system whose performance is uniform with variations in bandwidth and, in addition, is potentially well-suited for LPI communication. Since power-dominated homogeneous signals have a power spectrum of the form of (5.25), the spectral-matching rule suggests that fractal modulation may be naturally suited to channels with additive $1/f$ noise whose degree H is the same as that of the transmitted signal.

As we discussed in Chapter 3, the class of $1/f$ processes includes not only classical white Gaussian noise ($H = -1/2$) and Brownian motion ($H = 1/2$), but, more generally, a range of rather prevalent nonstationary noises that exhibit strong long-term statistical dependence.

In this section, we have developed a modulation strategy that satisfies the first of the two system requirements described at the outset of Chapter 6. We now turn our attention to the problem of designing optimal receivers for fractal modulation, and, in the process, we will see that fractal modulation also satisfies the second of our key system requirements.

6.3 RECEIVER DESIGN: DEMODULATION

Consider the problem of recovering a finite length message $q[n]$ from band-limited, time-limited, and noisy observations $r(t)$ of the transmitted wave-form $x(t)$ consistent with our channel model of Fig. 6.2. We assume that the noise $z(t)$ is a Gaussian $1/f$ process of degree $H_z = H$, and that the degree H_x of the homogeneous signal $x(t)$ has been chosen according to our spectral-matching rule, i.e.,

$$H_x = H_z = H. \tag{6.3}$$

We remark at the outset that if it is necessary that the transmitter measure H_z in order to perform this spectral matching, the robust and efficient parameter estimation algorithms for $1/f$ processes developed in Chapter 4 may be conveniently exploited.

Depending on the nature of the message being transmitted, there are a variety of different optimization criteria from which to choose in designing a suitable receiver. As representative examples, we consider two cases. In the first, the transmitted message is a digital data stream and we focus on the design of minimum probability-of-error receivers. In the second, the transmitted message is an analog data sequence for which we design minimum mean-square error receivers.

6.3.1 Demodulation of Digital Data

In this case, our transmitted message is a random bit stream of length L represented by a binary-valued sequence

$$q[n] \in \{+\sqrt{E_0}, -\sqrt{E_0}\}$$

where E_0 is the energy per bit. For this data, we develop in this section a receiver that demodulates $q[n]$ so as to minimize the bit-error probability. We remark in advance that efficient demodulation of more general real- or complex-valued M-ary sequences for $M > 2$ is possible using straightforward extensions of the results we develop in this section.

We begin by noting that an efficient implementation of the optimum receiver processes the observations $r(t)$ in the wavelet domain by first extracting the wavelet coefficients r_n^m using the DWT (2.21). These coefficients take the form

$$r_n^m = \beta^{-m/2}\, q[n \bmod L] + z_n^m, \tag{6.4}$$

where the z_n^m are the wavelet coefficients of the noise process, and where we have assumed that in accordance with our discussion in Section 6.2 the periodic replication of the finite length sequence $q[n]$ has been modulated. To simplify our analysis, we further assume that the ideal bandpass wavelet is used in the transmitter and receiver, although we reiterate that comparable performance can be achieved when more practical wavelets are used.

The duration-bandwidth characteristics of the channel in general affect which observation coefficients r_n^m may be accessed. In particular, if the channel is bandlimited to 2^{M_U} Hz for some integer M_U, this precludes access to the coefficients at scales corresponding to $m > M_U$. Simultaneously, the duration-constraint in the channel results in a minimum allowable decoding rate of 2^{M_L} symbols/sec for some integer M_L, which precludes access to the coefficients at scales corresponding to $m < M_L$. As a result, the collection of coefficients available at the receiver is

$$\mathbf{r} = \{r_n^m, m \in \mathcal{M}, n \in \mathcal{N}(m)\}$$

where

$$\mathcal{M} = \{M_L, M_L + 1, \ldots, M_U\} \tag{6.5a}$$

$$\mathcal{N}(m) = \{0, 1, \ldots, L\,2^{m-M_L} - 1\}. \tag{6.5b}$$

This means that we have available

$$K = \sum_{m=M_L}^{M_U} 2^{m-M_L} = 2^{M_U - M_L + 1} - 1 \tag{6.6}$$

noisy measurements of each of the L non-zero samples of the sequence $q[n]$. The specific relationship between decoding rate R, bandwidth W, and redundancy K can, therefore, be expressed in terms of the spectral efficiency of the modulation η_F as

$$\frac{R}{W} = \frac{2\eta_F}{K+1}, \tag{6.7}$$

where, as discussed earlier, $\eta_F \approx 1/2$. Note that when $M_U = M_L$ we have $K = 1$, and (6.7) attains its maximum value, η_F.

The optimal decoding of each bit can be described in terms of a binary hypothesis test on the set of available observation coefficients $\mathbf{r}$. Denoting by H_1 the hypothesis in which $q[n] = +\sqrt{E_0}$, and by H_0 the hypothesis in which

$q[n] = -\sqrt{E_0}$, we may construct the likelihood ratio test for the optimal decoding of each symbol $q[n]$. The derivation is particularly straightforward because of the fact that, in accordance with the wavelet-based models for $1/f$ processes developed in Chapter 3, under each hypothesis the z_n^m in (6.4) may be modeled as independent zero-mean Gaussian random variables with variances

$$\text{var}\, z_n^m = \sigma_z^2 \beta^{-m} \tag{6.8}$$

for some variance parameter $\sigma_z^2 > 0$. Consequently, given equally likely hypotheses (i.e., a random bit stream) the likelihood ratio test readily reduces to the test

$$\ell \underset{H_0}{\overset{H_1}{\gtrless}} 0,$$

where

$$\ell = \sum_{m=M_L}^{M_U} \beta^{m/2} \sum_{l=0}^{2^{m-M_L}-1} r_{n+lK}^m \tag{6.9}$$

is a sufficient statistic.

Before turning to a discussion of the resulting performance, it should be emphasized that, as in the case of the transmitter, the receiver has a convenient, computationally efficient, hierarchical implementation based on the DWT. Specifically, assuming $r(t)$ to be bandlimited to resolution 2^{M_U}, it may be sampled at rate 2^{M_U}, then successively filtered and downsampled to level $m = M_L$ according to the wavelet decomposition tree of Fig. 2.6(a). To produce the sufficient statistic ℓ, at each level m the terms from the detail sequence r_n^m corresponding to the same value of the $q[n]$ are collected together, weighted by the factor $\beta^{m/2}$, and accumulated with the weighted r_n^m from previous stages. Again, however, this is a batch algorithm, and while computationally efficient, this implementation may be less efficient in terms of storage requirements.

Performance

Since the statistic ℓ is Gaussian under each hypothesis, the performance associated with this optimal receiver is straightforward to derive. In particular, since

$$\begin{aligned} E[\ell|H_1] = -E[\ell|H_0] &= \sqrt{E_0}K \\ \text{var}[\ell|H_1] = \text{var}[\ell|H_0] &= \sigma_z^2 K \end{aligned}$$

and since

$$\Pr(\ell > 0|H_0) = \Pr(\ell < 0|H_1),$$

the bit-error probability can be expressed as

$$\Pr(\varepsilon) = \Pr(\ell > 0|H_0) = Q\left(\sqrt{K\sigma_c^2}\right), \tag{6.10}$$

where $Q(\cdot)$ is again defined by (4.32), and where σ_c^2 is the SNR in the channel, i.e.,

$$\sigma_c^2 = \frac{E_0}{\sigma_z^2}.$$

Substituting for K in (6.10) via (6.7) we can rewrite this error probability in terms of the channel rate-bandwidth ratio as

$$\Pr(\varepsilon) = Q\left(\sqrt{\sigma_c^2\left[\frac{2\eta_F}{R/W} - 1\right]}\right), \tag{6.11}$$

where, again, $\eta_F \approx 1/2$. Note that, as we would anticipate with the strategy, the performance of fractal modulation is independent of the spectral exponent of the noise process when we use our spectral matching procedure.

To establish a performance baseline, we also evaluate a modified version of our benchmark modulation in which we incorporate repetition coding, i.e., in which we add redundancy by transmitting each sample of the message sequence K times in succession. This comparison scheme is not particularly power efficient both because signal power is distributed uniformly over the available bandwidth irrespective of the noise spectrum, and because much more effective coding schemes can be used with channels of known bandwidth [96] [97]. Nevertheless, with these caveats in mind, such comparisons do lend some insight into the *relative* power efficiency of fractal modulation.

In our modified benchmark modulation, incorporating redundancy reduces the effective decoding rate per unit bandwidth by a factor of K, i.e.,

$$\frac{R}{W} = \frac{\eta_0}{K}, \tag{6.12}$$

where η_0 is the efficiency of the modulation without coding, i.e., unity. When the channel adds stationary white Gaussian noise, for which $H = -1/2$, the optimum receiver for this scheme demodulates the received data and averages together the K symbols associated with the transmitted bit, thereby generating a sufficient statistic. When this statistic is positive, the receiver decodes a 1-bit, and a 0-bit otherwise. The corresponding performance is, therefore, given by [97]

$$\Pr(\varepsilon) = Q\left(\sqrt{K\sigma_c^2}\right) = Q\left(\sqrt{\sigma_c^2\left[\frac{\eta_0}{R/W}\right]}\right), \tag{6.13}$$

where the last equality results from substituting for K via (6.12).

Comparing (6.13) with (6.11), we note that since $\eta_0 \approx 2\eta_F$, the asymptotic bit-error performances of fractal modulation and the benchmark scheme are effectively equivalent for $R/W \ll \eta_F$, as is illustrated in Fig. 6.5. In Fig. 6.5(a), $\Pr(\varepsilon)$ is shown as a function of R/W at a fixed SNR of 0 dB ($\sigma_c^2 = 1$), while in Fig. 6.5(b), $\Pr(\varepsilon)$ is shown as a function of SNR at a fixed $R/W = 0.125$ bits/sec/Hz. Both these plots reveal strong thresholding behavior whereby the error probability falls off dramatically at high SNR and low R/W. It is important to emphasize that comparisons between the two schemes are meaningful only for the case in which the noise has parameter $H = -1/2$, corresponding to the case of stationary white Gaussian noise. For other values of H, the performance of the benchmark modulation is not only difficult to evaluate, but necessarily poor as well because of inefficient distribution of power among frequencies.

6.3.2 Demodulation of Analog Data

In this section, we assume that $q[n]$ is a continuous-valued sequence of independent, identically distributed, zero-mean Gaussian random variables, each with variance

$$\operatorname{var} q[n] = \sigma_x^2,$$

and develop a receiver yielding the minimum mean-square error (MSE) estimate of $q[n]$ based on our corrupted observations $r(t)$.

We proceed in a manner analogous to that described in Section 6.3.1 for the case of digital data. In particular, we first project our observations onto the ideal bandpass wavelet basis from which $x(t)$ was synthesized, so that our observations may again be expressed in the form (6.4). Given the set of K accessible observation coefficients r_n^m specified by (6.5) with (6.6), we readily obtain that the optimum estimates of $q[n]$ are of the form

$$\hat{q}[n] = E\left[q[n] | \mathbf{r}\right] = \frac{\ell}{K + (1/\sigma_c^2)} \qquad (6.14)$$

where ℓ is the sufficient statistic (6.9) and where σ_c^2 is the SNR in the channel, defined by

$$\sigma_c^2 = \frac{\sigma_x^2}{\sigma_z^2}.$$

Evidently, the sufficient statistic ℓ defined in (6.9) plays a key role in the demodulation of both digital and analog data, and in fact its calculation dominates the computational complexity of the receiver. However, we emphasize that, as was discussed in Section 6.3.1, that ℓ can be obtained via a computationally efficient algorithm by exploiting the DWT.

In general, and as we would expect, the optimum estimate represents a blend of *a priori* information about $q[n]$, and information obtained from the observations. At high SNR ($\sigma_c^2 \gg 1/K$), the *a priori* information is essentially ignored, and the resulting estimator specializes to the maximum likelihood

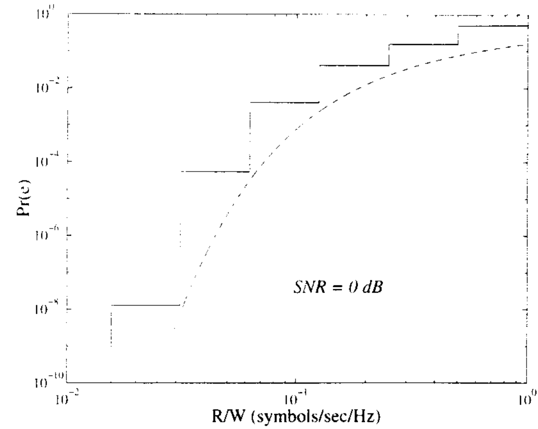

(a)

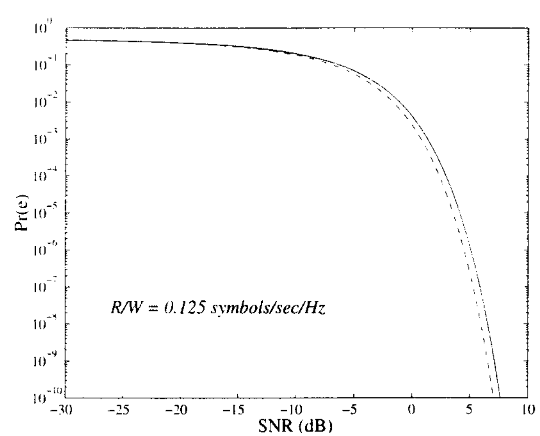

(b)

Figure 6.5. Bit-error performance of fractal modulation with digital data. Solid lines indicate the performance of fractal modulation, while dashed lines indicate the performance of the benchmark modulation with repetition coding. (a) Bit-error probability $\Pr(e)$ as a function of Rate/Bandwidth ratio R/W at 0 dB SNR. (b) Bit-error probability $\Pr(e)$ as a function of SNR at $R/W = 0.125$ symbols/sec/Hz.

estimator. At low SNR ($\sigma_c^2 \ll 1/K$), the observations are essentially ignored, and the estimator approaches the *a priori* estimate

$$\hat{q}[n] = E\left[q[n]\right] = 0.$$

Finally, we remark that the optimum receiver (6.14) with (6.9) is a linear data processor, as would be anticipated since we have restricted the discussion to Gaussian sequences and Gaussian noise. In non-Gaussian scenarios, the receivers we have developed are the best *linear* data processors; i.e., no other linear data processor is capable of generating an estimate of $q[n]$ with a smaller mean-square error.

Performance

The normalized MSE associated with the optimum receiver (6.14) can be readily derived as

$$\epsilon^2 = \frac{E\left[(q[n] - \hat{q}[n])^2\right]}{E\left[(q[n])^2\right]} = \frac{E\left[\mathrm{var}(q[n]|\mathbf{r})\right]}{\mathrm{var}\, q[n]} = \frac{1}{1 + K\sigma_c^2}. \tag{6.15}$$

Generally, it is convenient to substitute for K in (6.15) via (6.7) to get

$$\epsilon^2 = \frac{1}{1 + \sigma_c^2 \left[\dfrac{2\eta_F}{R/W} - 1\right]}, \tag{6.16}$$

where $\eta_F \approx 1/2$, and where $R/W \leq \eta_F$ by virtue of our definition of η_F. From (6.16) we see, then, that for $R/W \ll \eta_F$, the MSE is given asymptotically by

$$\epsilon^2 \sim \frac{1}{\sigma_c^2 \dfrac{2\eta_F}{R/W}}. \tag{6.17}$$

Note that, as in the case of digital data, the performance (6.16) is independent of the parameter H when we use spectral matching.

For the purposes of comparison, consider the MSE performance of our benchmark modulation with repetition coding in the presence of stationary white Gaussian noise. As in the case of digital data, incorporating redundancy reduces the effective rate-bandwidth ratio by a factor of K, yielding (6.12). The optimum Bayesian receiver for this scheme, using a minimum MSE criterion, demodulates the repeated sequence, and averages the terms corresponding to the same value of $q[n]$ to generate $\hat{q}[n]$. Hence this K-fold redundancy leads to a normalized MSE of

$$\epsilon^2 = \frac{E\left[(q[n] - \hat{q}[n])^2\right]}{E\left[(q[n])^2\right]} = \frac{1}{1 + \sigma_c^2 K} \tag{6.18}$$

where σ_c^2 is the SNR, i.e., the ratio of the power in $q[n]$ to the density of the white noise power spectrum. Combining (6.18) with (6.12) we get

$$\epsilon^2 = \frac{1}{1 + \sigma_c^2 \dfrac{\eta_0}{R/W}} \tag{6.19}$$

whenever $R/W \leq \eta_0$. By comparison with (6.17) we see that when $R/W \ll \eta_0$,

$$\epsilon^2 \sim \frac{1}{\sigma_c^2 \dfrac{\eta_0}{R/W}}, \tag{6.20}$$

which is essentially (6.17) since the achievable η_0 is unity and $\eta_F \approx 1/2$ as discussed earlier. This means that, at least asymptotically, the performance between the two schemes is comparable in the presence of white noise.

This behavior is reflected in the performance curves for both fractal modulation and the benchmark modulation with repetition coding of Fig. 6.6. In Fig. 6.6(a), MSE is shown as a function of R/W at a fixed SNR of 0 dB ($\sigma_c^2 = 1$), while in Fig. 6.6(b), MSE is shown as a function of SNR at a fixed $R/W = 0.125$ symbols/sec/Hz. As we expect, the longer the channel is open, or the greater the available bandwidth in the channel, the better the performance of fractal modulation. Although comparisons between the two modulation schemes are appropriate only for the special case of additive white Gaussian noise channels, we reiterate that the performance of fractal modulation (6.16) is independent of the spectral exponent of the $1/f$ noise. By contrast, we would not, in general, expect (6.19) to describe the performance of the benchmark modulation with repetition coding in the presence $1/f$ noise.

6.4 SUMMARY

In this chapter, we developed a novel and powerful diversity strategy that we referred to as fractal modulation. As developed, fractal modulation constitutes a compelling paradigm for communication over noisy channels of simultaneously uncertain duration and bandwidth. As we discussed, this channel is a rather realistic model for a variety of scenarios encountered in point-to-point, broadcast, and multiple-access communications.

With fractal modulation, we used the orthonormal self-similar basis expansions derived in Chapter 5 to develop an approach for modulating discrete- or continuous-valued information sequences onto homogeneous signals. The result was a modulation scheme in which information was embedded in the transmitted waveform on all time scales.

We then considered the problem of optimum demodulation of such transmissions. Two representative cases were considered. The first involved

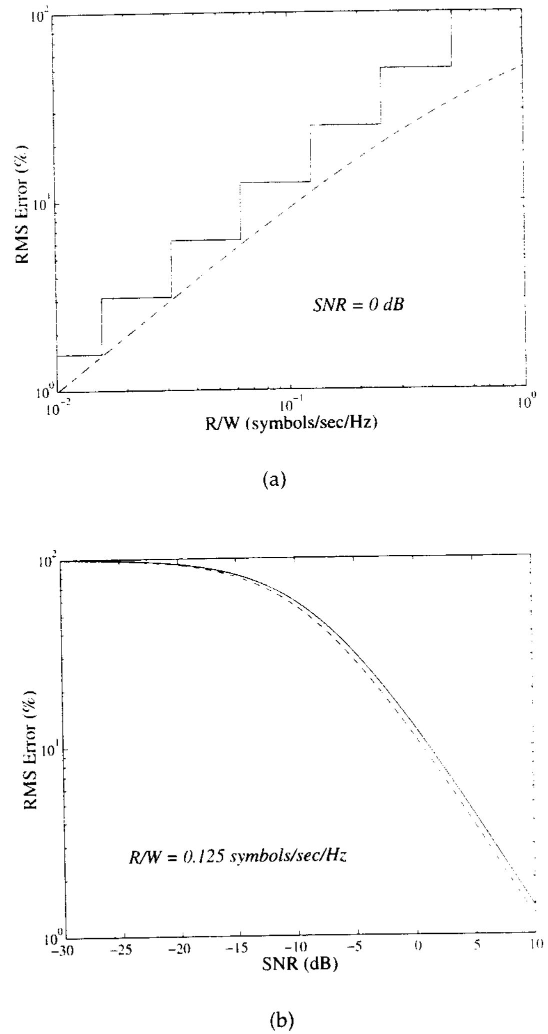

Figure 6.6. Tradeoffs between error, rate, and bandwidth for fractal modulation with the optimum receiver for noisy analog data. The solid lines represent the performance of fractal modulation, while the dashed lines correspond to the performance of the benchmark modulation with repetition coding. (a) MSE ϵ^2 as a function of Rate/Bandwidth ratio R/W at 0 dB SNR. (b) MSE ϵ^2 as a function of SNR at $R/W = 0.125$ symbols/sec/Hz.

minimum probability-of-error demodulation of a digital data stream, while the second involved minimum mean-square error demodulation of an analog data stream.

Of considerable practical interest, we also showed that the use of self-similar bases derived from wavelet bases as developed in Chapter 5 leads directly to computationally efficient discrete-time implementations of the resulting transmitters and receivers for fractal modulation systems.

Our development included an evaluation of several aspects of the performance of this diversity strategy on the channel of interest. In order to provide a performance baseline, our analysis included comparisons to more traditional forms of modulation and diversity.

A number of interesting open questions were not addressed in our development. For example, the extent to which fractal modulation is optimum for the unknown duration-bandwidth channel remains to be explored. Closely related is the issue of how to define a useful notion of capacity for such channels, since the usual Shannon capacity degenerates in this case. It is also possible, though not obvious, that efficient coding techniques [98] can be used in conjunction with fractal modulation to provide a more effective diversity benefit on such channels.

As a final remark, we point out that fractal modulation and its generalizations also have a potentially important role to play in secure communications. In such applications, considerations such as vulnerability to detection, interception, and exploitation by various strategies naturally become important [99]. To meet the particular needs that arise in these scenarios, a variety of basic extensions to fractal modulation can be developed. For example, traditional direct-sequence spread-spectrum type techniques [97] can be combined with fractal modulation in a relatively straightforward manner to further enhance its suitability to such applications.