1

Introduction and Overview

Signal processing plays a central role of a truly enormous range of modern electronic systems. These include, for example, voice, data, and video communication and storage systems; medical imaging and diagnostic systems; radar, sonar and satellite remote sensing systems; and automated manufacturing systems. Signal processing algorithms required for these kinds of applications take many forms. Efficient speech compression and recognition systems, for example, require sophisticated signal *analysis* algorithms. By contrast, high-speed modems for data communication require not only powerful signal detection algorithms at the receiver, but efficient signal *synthesis* algorithms at the transmitter. And several other kinds of algorithms—including signal restoration, enhancement, and manipulation—are also important in diverse applications.

Whenever we construct algorithms for signal synthesis, analysis or processing, we invariably exploit some *model* for the signals of interest. In some instances this model is implicit. For example, spline-based strategies are often used in signal interpolation problems even when there is no quantitative model for the signal. In such cases, there is an implicit assumption that the signal is smooth at least in some qualitative sense. In other instances the model is explicit. For example, bandlimited interpolation algorithms are designed for a specific and well-defined class of signals whose frequency content is zero except within some known bandwidth.

The signal model may arise out of a deterministic formulation (e.g., a sinusoid of possibly unknown frequency and/or phase), or a stochastic one (e.g., a first-order autoregressive process). It may encompass a particularly

broad class of signals, or a relatively narrow one. The class of stationary random processes, for instance, is much larger than the class of first-order autoregressive processes. Not surprisingly, better performance can be expected of algorithms tuned to more narrowly defined classes of signals. This of course inevitably comes at the expense of robustness and flexibility: we cannot in general expect systems to behave well on signals for which they have not been designed.

Choosing a signal model for a particular application involves many factors in addition to those cited above. For example, naturally we seek models that capture the important characteristics of the physical signals of interest as closely as possible. On the other hand, an overly complex model can lead to an essentially intractable framework for algorithm development. Hence, the art and science of signal model selection invariably involves a compromise.

Much of traditional signal processing has relied upon a relatively small class of models for use in wide-ranging applications. These models include, for example, periodic and bandlimited signals. They also include important classes of stationary random signals derived from filtered white noise processes such as the autoregressive moving-average (ARMA) models. And certain cyclostationary signals turn out to be important models for many signal processing applications involving digital communications.

It is worth noting that as a whole these models have tended to arise out of signal processing's deep traditional foundation in linear time-invariant (LTI) system theory [1] [2] [3]. For example, periodic signals are characterized by a form of translational invariance: the associated waveforms are invariant to time translations by multiples of the period. Likewise, stationary and cyclostationary signals are also characterized by a form of translational invariance: the *statistics* of these processes are invariant to appropriate time translations of the underlying waveforms. With the close connection to LTI system theory, it is not surprising, then, that the Fourier transform plays such a key role in the analysis and manipulation of these models, nor that fast Fourier transform (FFT) plays such a key role in the implementation of algorithms based on such models.

Interestingly, even the most sophisticated models that are used in signal processing applications typically have evolved out of useful extensions of an LTI framework. For example, the time-varying signal models required for many applications often developed out of such extensions. As a consequence windowed Fourier transform generalizations such as the short-time Fourier transform (STFT) frequently play an important role in such applications. Even state-space formulations, which can in principle accommodate a broad range of behaviors, are firmly rooted in a time-invariant perspective.

However, a strong theme that runs throughout this book is that there are many physical signals whose key characteristics are fundamentally different

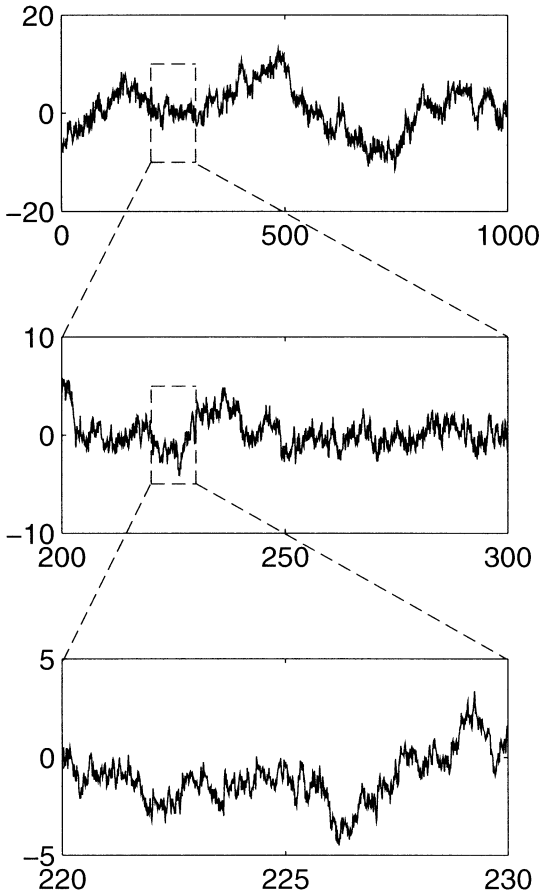

Figure 1.1. A sample waveform from a statistically scale-invariant random process, depicted on three different scales.

than can be produced by many of these traditional models and such generalizations. In particular, there are many signals whose defining characteristic is their invariance not to translation but rather to *scale*. For example, in the stochastic case, we mean that the statistics of the process do not change when we stretch or shrink the time axis. Both qualitatively and quantitatively the process lacks a characteristic scale: behavior of the process on short time scales is the same as it is on long time scales. An example of a process with this behavior is depicted in Fig. 1.1.

Typically, an important consequence of this scale-invariance is that the resulting waveforms are *fractal* and highly irregular in appearance [4]. In general, fractals are geometric objects having nondegenerate structure on all temporal or spatial scales. For the fractals most frequently of interest, the structure on different scales is related. The fractal processes on which we focus in particular have the property that this structure at different scales is *similar.* This, in turn, has the important consequence that even well-separated samples of the process end up being strongly correlated.

Collectively, these features are, of course, extremely common in many types of physical signals, imagery, and other natural phenomena. In fact, as we discuss in Chapter 3, fractal geometry abounds in nature. Fractal structure can be found, for example, in natural landscapes, in the distribution of earthquakes, in ocean waves, in turbulent flow, in the pattern of errors and data traffic on communication channels, in the bronchi of the human lung, and even in fluctuations of the stock market.

It is worth emphasizing that random process models traditionally used in signal processing, by contrast, typically do not exhibit this kind of behavior. Fig. 1.2, for example, depicts the behavior of a simple, first-order autoregressive process on different time scales. Hence, in order to use fractal signal models with statistical scale-invariance characteristics in applications, we require techniques for detecting, identifying, classifying, and estimating such signals in the presence of both other signals and various forms of additional distortion. In Chapter 4, we therefore develop a variety of such signal processing algorithms. As an example, we develop robust and efficient algorithms for estimating the fractal dimension of signals from noisy measurements, and for optimally restoring such signals.

In the deterministic case, scale-invariance means that the actual waveform is repeatedly embedded within itself, so that temporal dilations and compressions simply reproduce the signal. An example of a "self-similar" waveform of this type is depicted in Fig. 1.3 on different time scales. As we develop in Chapter 6, such signals have a potentially important role to play as modulating waveforms in, for example, secure communication applications, where they provide a novel and powerful form of transmission diversity. In particular, the resulting waveforms have the property that the transmitted information can be reliably recovered given either severe time- or band-limiting in the communication channel. In light of the geometric properties of the resulting transmitted waveforms, we refer to this diversity strategy as "fractal modulation."

At one time it was believed that relatively complicated signal models were necessary to produce the kind of special behavior we have described. However, as will become apparent in subsequent chapters, while fundamentally new model structures are required, they are in fact no more complex than (and in many respects at least as tractable as) many traditional signal

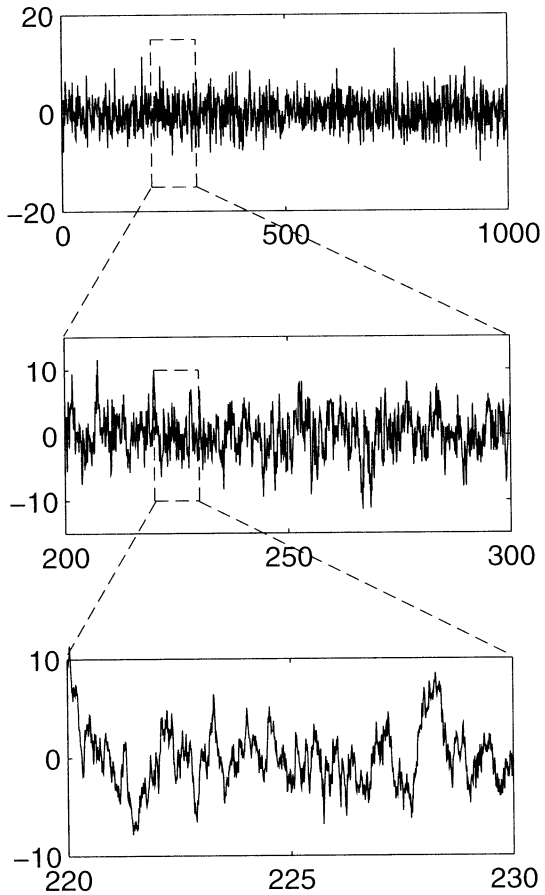

Figure 1.2. A sample waveform from a simple first-order autoregressive process, depicted on three different scales.

processing models. In particular, remarkably simple characterizations that are useful both conceptually and analytically are developed in Chapter 3 for the case of statistically scale-invariant signals, and in Chapter 5 for the case of deterministically scale-invariant signals.

In addition, while the Fourier transform plays a central role in the analysis and manipulation of both statistically and deterministically *translation*-invariant signals, as we will see it is the *wavelet* transform that plays an analogous role for the kinds of scale-invariant signal models that are the focus of this book. Wavelet representations are expansions in which the basis

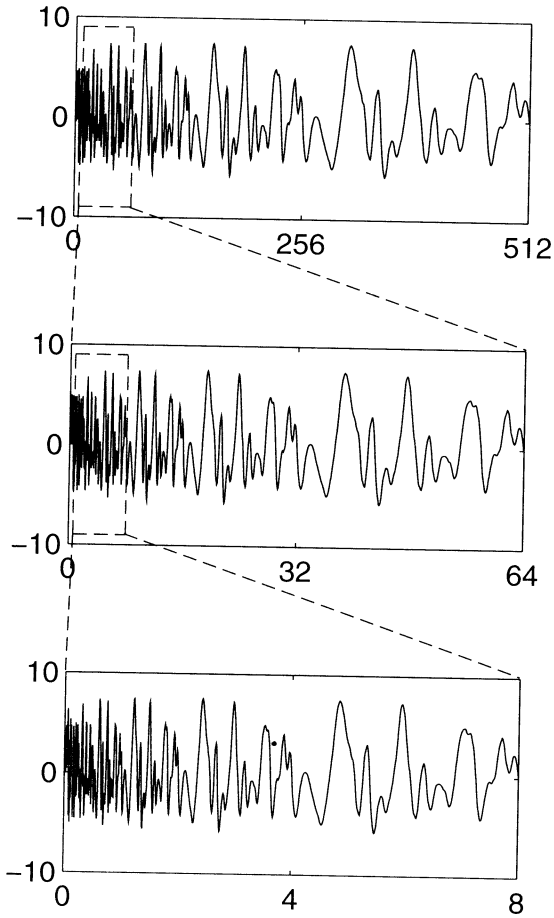

Figure 1.3. Dilated self-similar waveform.

signals are all dilations and translations of a suitable prototype function. In Chapter 3, for example, we develop the role of such wavelet representations as *Karhunen-Loève* type expansions for stochastically scale-invariant random processes: when projected onto wavelet bases, the resulting expansion coefficients are effectively uncorrelated. In Chapter 5, we show how wavelet bases can be used to construct orthonormal bases of self-similar waveforms that provide efficient representations for deterministically scale-invariant waveforms.

Yet wavelet representations are not merely a conceptual tool in the manipulation of fractal signals. In fact, just as the FFT plays an important role in practical algorithms for processing translation-invariant signals, it is the discrete wavelet transform (DWT) that is frequently a key component of practical algorithms for processing scale-invariant signals. In Chapter 4, for example, we see how fast algorithms for estimating the fractal dimension of signals rely heavily on the DWT. Similarly, in Chapter 6, we see how the DWT algorithm is critical to efficient implementations of fractal modulation transmitters and receivers.

Finally, in Chapter 7 we explore some system theoretic foundations for the classes of signals developed in earlier chapters. We discuss the duality between time-invariant and scale-invariant systems, and show how Mellin, Laplace, and wavelet transforms arise rather naturally in the representation of the latter class of systems. In the process, some interesting and unifying themes are introduced.

Given the central role that the wavelet transform plays in the efficient representation of scale-invariant signals, we begin with a self-contained overview of the relevant aspects of wavelet theory in Chapter 2.