HST.582J / 6.555J / 16.456J Biomedical Signal and Image Processing
Spring 2007

Chapter 1 - DATA ACQUISITION

Introduction

The goal of *data acquisition* is to capture a signal and encode in a form suitable for computer processing with minimum loss of information. Data acquisition typically consists of three stages: transduction, analog conditioning, and analog-to-digital conversion. *Transduction* is the conversion from one form of energy to another. In present technology, the only form of energy suitable for encoding into a computer is electrical energy, therefore signals need to be converted to *analog* voltages whose waveforms are ideally the same as those of the original signals. For example, we use a microphone to transduce an acoustic signal, or an electric thermometer to measure temperatures. Transducers are specific to each type of signal, and the study of such devices is beyond the scope of these notes.

The second stage of data acquisition, *analog signal conditioning,* usually consists of amplifying and filtering the analog signal measured with a transducer. Because the purpose of this stage is to provide a good match between the typically low-amplitude, wide-bandwidth transducer signals and the analog-to-digital (A/D) converter, conditioning is best understood after studying A/D conversion.

An *analog-to-digital converter* is a device that transforms a continuous-time signal measured with a transducer into a *digital* signal that can be represented in a computer. Conceptually, it can be divided into a series of two operations (which are realized simultaneously in actual devices): *sampling,* in which the continuous-time, analog signal is converted into one that is only defined for discrete times, but whose amplitude can take arbitrary values, and *quantization,* in which a continuous-amplitude signal is converted into a digital signal that can only take a finite set of values. The sampling operation is particularly critical if we want to avoid loss of information in the conversion.

1.1 Continuous-time and discrete-time signals

Many signals are *continuous-time* in the sense that they are defined at arbitrarily-close points in time. Sine waves are important examples of continuous-time signals because Fourier's theorem states that most signals of practical interest can be decomposed into an infinite sum of sine waves. A continuous-time sine wave is defined by:

$$x(t) = a\cos(\Omega t + \phi) = a\cos(2\pi F t + \phi) \tag{1.1}$$

where F is the frequency, Ω is the angular frequency, a the amplitude, and ϕ the phase.

Discrete-time signals (also called *time series)* are defined over the set of integers, that is, they are indexed sequences. A discrete-time sine wave is defined by

$$x[n] = a\cos(\omega n + \phi) = a\cos(2\pi f n + \phi) \tag{1.2}$$

It is important to understand that the frequency of a discrete-time sinusoid is not uniquely defined. This fundamental ambiguity is a consequence of the basic trigonometric result that the value of a sinusoid does not change if an integer multiple of 2π is added to its argument:

$$\cos(\theta) = \cos(\theta + 2\pi k), \quad \text{where } k \text{ is an integer} \tag{1.3}$$

Adding the $2\pi kn$ to the argument of (1.2), and applying (1.3), we get

$$x[n] = a\cos(2\pi f n + \phi) = a\cos(2\pi f n + 2\pi k n + \phi)$$

Two cases must be distinguished. If $k \geq -f$, this is equivalent to a sinusoid with frequency $f + k$ with no change in phase.

$$x[n] = a\cos(2\pi f n + \phi) = a\cos(2\pi(f + k)n + \phi), \quad k \geq -f \tag{1.4a}$$

On the other hand, if $k < -f$, (1.4a) leads to a negative frequency. To avoid this, we introduce $l \triangleq -k$, and make use of $\cos(\theta) = \cos(-\theta)$, to obtain a sinusoid of frequency $l - f$ with a reversal in phase:

$$x[n] = a\cos(2\pi f n + \phi) = a\cos(2\pi(l - f)n - \phi), \quad l > f \tag{1.4b}$$

Thus, a discrete-time sinusoid with frequency f is identical to a same-phase sinusoid of frequency $f + k$, where k is any integer greater than $-f$, or to a phase-reversed sinusoid of frequency $l - f$ with $l > f$.

The above relations are more concisely expressed using complex exponential notation. Specifically, (1.2) can be written as:

$$x[n] = Re\left[e^{j(2\pi f n + \phi)}\right] = Re\left[e^{-j(2\pi f n - \phi)}\right] \tag{1.5}$$

Because value of a complex exponential does not change if a multiple of 2π is added to its argument, we get:

$$x[n] = Re\left[e^{j(2\pi(f+k)n + \phi)}\right] = Re\left[e^{-j(2\pi(k-f)n - \phi)}\right] \tag{1.6}$$

which is equivalent to (1.4ab). Because of this fundamental frequency ambiguity, we will often implicitly assume that the angular frequency of a discrete-time sinusoid is restricted to the range $-\pi \leq \omega \leq \pi$, or, equivalently, that $-\frac{1}{2} \leq f \leq \frac{1}{2}$.

1.1.1 Sampling a continuous-time signal

Certain signals are inherently discrete, for example the periodic measurements of a patient's temperature, or a stock-market index at daily closing times. More often, discrete-time signals are obtained by sampling a continuous-time signal at regular intervals:

$$x[n] \triangleq x(nT_s), \quad -\infty < n < \infty \tag{1.7}$$

T_s is the *sampling interval,* and $F_s \triangleq \frac{1}{T_s}$ is the *sampling frequency* or *sampling rate.* Because we use square brackets for discrete-time signals, and parentheses for continuous-time signals, it is clear that the two "x" in (1.7) represent different signals.

1.1.2 Sampling a sinusoid - Aliasing

For example, assume that we sample a continous-time sinusoid as in (1.1). The corresponding discrete-time signal is:

$$x[n] = x(nT_s) = a\cos(2\pi F nT_s + \phi) = a\cos(2\pi nF/F_s + \phi) \tag{1.8}$$

This is a discrete-time sinusoid whose frequency is related to that of the original signal by:

$$\omega = \Omega T_s \quad \text{and} \quad f = F/F_s \tag{1.9}$$

Equation (1.9) of great practical importance when processing signals with digital systems because it relates the dimensionless frequency f used by the computer to the physical frequency F.

Despite the simplicity of (1.8), it hides a difficulty arising from the ambiguity of frequency for discrete-time sinusoids. Specifically, from (1.4), the same discrete-time signal $x[n]$ in (1.9) could have been obtained by sampling any of the following continous-time sinusoid at the same rate F_s:

$$x_k(t) = a\cos(2\pi(F + kF_s)t + \phi), \quad k \geq -F/F_s \tag{1.10a}$$

or

$$x_l(t) = a\cos(2\pi(lF_s - F)t - \phi), \quad l > F/F_s \tag{1.10a}$$

In other words, once the signal is sampled, it is not possible to know if the frequency of the original continuous time signal was F or $F + F_s$ or $F + 2F_s$, etc, or, with phase reversal, $F_s - F$ or $2F_s - F$, etc. This phenomenon is known as *aliasing* because frequencies may not be what they appear to be once a continuous-time signal is sampled. Figure 1.5 illustrates how two continuous-time sinusoids with different frequencies can lead to identical discrete-time signals. Figure 1.14 shows, in the frequency domain, some of the continuous-time sinusoids that are aliased onto the same discrete-time signals.

In practice, aliasing is not an issue for single sinusoids so long that the frequency of the analog signal is known withing a factor of $F_s/2$. Specifically, assume that the discrete-time signal (1.2) is obtained by sampling a continous-time signal $x(t)$. Further assume that we know that the frequency of the original signal was between $kF_s/2$ and $(k+1)F_s/2$, where k is a positive integer. Then the analog signal can be uniquely determined as:

$$x(t) = \cos(2\pi(f + k/2)F_s t + \phi), \quad \text{if } k \text{ is even} \tag{1.11a}$$

$$x(t) = \cos(2\pi((k + 1)/2 - f)F_s t - \phi), \quad \text{if } k \text{ is odd} \tag{1.11b}$$

1.2 The Nyquist sampling theorem

The purpose of sampling continuous-time signals is to allow their processing by digital computers. This is useful only if the information carried by the original signal is retained in the sampled version. Intuitively, it seems that, if the continuous-time signal is sufficiently smooth, sufficiently close samples will provide a good approximation to the original signal. The *Nyquist sampling theorem* formalizes this intuition: It states that, if a signal $x(t)$ contains no frequency components higher than W, it can be exactly reconstructed from samples taken at a frequency $F_s > 2W$. Further, the Nyquist theorem gives an explicit *interpolation formula* for reconstructing $x(t)$ from the discrete-time signal $x[n]$:

$$x(t) = \sum_{n=-\infty}^{\infty} x[n]\phi(t - nT_s) \qquad (1.12a)$$

with

$$\phi(t) = \frac{\sin(\pi F_s t)}{\pi F_s t} \qquad (1.12b)$$

This interpolation formula expresses $x(t)$ as a linear function (a weighted sum) of the samples $x[n]$. The time-dependent weights $\phi(t - nT_s)$ are obtained by delaying a basic function $\phi(t)$ shown in Fig. 1.4. This function verifies the property

$$\phi(nT_s) = \begin{cases} 1 & \text{if } n = 0 \\ 0 & \text{otherwise} \end{cases} \qquad (1.13)$$

This property implies that $x(t) = x[n]$ for $t = nT_s$, which is reassuring since this is the very definition of sampling.

A signal is said to be sampled at the *Nyquist frequency* if F_s is exactly $2W$. Signals sampled at a rate higher than the Nyquist frequency are said to be *oversampled.*

1.2.1 Sampling theorem seen in the frequency domain

A general proof of the Nyquist sampling theorem will be given in Chapter 5 using Fourier transforms. Appendix 1.A.2 gives a derivation of the sampling theorem in the special case of periodic signals. This case is easily handled because a bandlimited periodic signal is entirely specified by a finite set of Fourier series coefficients. Fourier series are introduced in Appendix 1.A.1. Here, we give a qualitative argument for the Nyquist sampling theorem based on an informal version of the Fourier theorem.

The Fourier theorem states that any reasonably well behaved signal can be expressed as a "sum" of sinusoids or, alternatively, complex exponentials. Informally, one has:

$$x(t) = \sum_F A(F)\cos(2\pi F t + \phi(F)), \quad 0 \le F < W \qquad (1.14)$$

where the "sum" is from $F = 0$ to a maximum frequency W if $x(t)$ is bandlimited. If this signal is sampled at a rate $F_s > 2W$, each continous sinusoid with frequency F inside the sum is converted into a unique discrete-time sinusoid of frequency $f = F/F_s$:

$$x[n] = \sum_f A(fF_s)\cos(2\pi f n + \phi(fF_s)), \quad 0 \le f < 1/2 \qquad (1.15)$$

Because the discrete-time frequencies are all distinct, there is no loss of information, and it is possible to reconstruct each component of the original signal, and therefore the signal $x(t)$, from $x[n]$. In fact, using (1.11), this would work even if the original signal had been limited to any frequency range from $kF_s/2$ to $(k+1)F_s/2$, where k is any positive integer, so long that the range of frequencies is known. This situation is illustrated in Fig. 1.15.

1.2.2 Aliasing for complex signals

What happen when the Nyquist criterion is not met, i.e. when the sampling rate F_s is less than twice the frequency range W of the analog signal? In that case, there is no longer a one-to-one relationship between frequency components of the analog signal $x(t)$ and those of the discrte-time signal $x[n]$. Specifically, frequency components of $x(t)$ higher than $F_s/2$ are aliased into low-frequency components that interact with low-frequency components of $x(t)$. These interactions can be additive if the co-aliased components are in phase, or destructive if they are out of phase. In such cases, there is irreversible loss of information, i.e. it is no longer possible to reconstruct the anlog signal from the discrete-time signal.

To illustrate these effects of aliasing, consider the two-component signal shown in Fig. 1.6:

$$x(t) = \cos 2\pi F_1 t + \cos(2\pi F_2 t + \phi), \tag{1.16}$$

with $F_1 < F_2$. If this signal is sampled at frequency $F_s = F_1 + F_2$, the low-frequency component F_1 does not alias because $F_1 < F_s/2$, but the high-frequency component F_2 is aliased to $F_s - F_2 = F_1$, with a phase reversal:

$$x[n] = \cos 2\pi f_1 n + \cos(2\pi f_1 n - \phi), \quad f_1 = F_1/F_s \tag{1.17}$$

If the phase ϕ is zero, the two components of the discrete-time signal add in phase, so that the two-component discrete-time signal is undistinguishible from a single sinusoid that has double the amplitude of either component. On the other hand, if $\phi = \pi$, the two components have opposite phases, so that the aliased component exactly cancels the f_1 component, and the sampled signal $x[n]$ is identically zero. This case is shown in Fig. 1.6.

In practice, one can always avoid aliasing by lowpass filtering the continuous-time signal $x(t)$ before sampling. An analog filter with sharp cutoff is desirable to avoid having to oversample the signal. Actual analog filters always require a finite transition band between the passband and the stopband. **Even if the signal is known to be band-limited, it is still important to use an antialiasing lowpass filter because the signal might be contaminated by broadband noise whose high-frequency components would be aliased into the low-frequency range of the signal.**

1.3 Quantization

The second conceptual step in analog-to-digital conversion consists in quantizing the sampled signal. A quantizer takes as input a discrete-time, continuous-valued signal $x[n]$, and produces a signal $x_Q[n]$ that can only take a finite number of values. In one common quantization scheme, called *rounding*, the quantizer output $x_Q[n]$ is equal to kQ, where Q is the *quantization step,* and k is the integer closest to $x[n]/Q$. In most quantizers, the number of quantization steps is a power of two of the form 2^B, where B is called the *number of bits* of the quantizer. With a B-bit quantizer, one can only encode signals whose values lie in the range

$$-V_{\max} \leq x[n] < V_{\max} \tag{1.18}$$

where $V_{\max}$ is related to the quantization step by:

$$V_{\max} = 2^{B-1}Q \tag{1.19}$$

Figure 1.8 shows $x_Q[n]$ as a function of $x[n]$ for $B = 4$ and $Q = 1$, corresponding to $V_{\max} = 8$.

One way to describe the effects of quantization is to write the quantized signal $x_Q[n]$ as the sum of the input signal $x[n]$ and a *quantization error signal* $q[n]$:

$$x_Q[n] \triangleq x[n] + q[n] \tag{1.20}$$

A convenient characterization of the degradation caused by quantizing is the *signal-to-noise ratio (SNR):*

$$\text{SNR} \triangleq \frac{\text{mean power in } x[n]}{\text{mean power in } q[n]} \triangleq \frac{P_x}{P_q} \tag{1.21}$$

For the rounding scheme of quantization, one always has:

$$-Q/2 < q[n] \leq Q/2 \tag{1.22}$$

If the quantization step is sufficiently small relatively to the signal amplitude, and the signal is sufficiently irregular, the quantization error signal will be equally likely to take any values between $-Q/2$ and $Q/2$. We will show in Chapter 11 that, under these conditions, the mean power in the error signal is:

$$P_q = \frac{Q^2}{12} \tag{1.23}$$

The mean signal power P_x can be written as a product of three terms:

$$P_x = \frac{P_x}{X_{\max}^2} \left(\frac{X_{\max}}{V_{\max}}\right)^2 V_{\max}^2 \tag{1.24}$$

where $X_{\max}$ is the maximum value of $x[n]$.

The first term on the right side of (1.24) is the reciprocal of the *peak factor* K_P. This is a dimensionless number that characterizes the shape of the signal, and does not change if the signal is multiplied by a constant gain. For example, for a sinusoid, the peak factor is always 2. For signals that contain sharp pulses, with little energy between pulses, the peak factor is larger.

The second term on the right side of (1.24) represents the ratio of the peak amplitude in the signal to the maximum amplitude that can be encoded by the quantizer. Because this ratio must always be smaller than one to avoid clipping, we will call it the *clipping factor* K_C. Making use of (1.23), (1.24), and (1.21), the signal-to-noise ratio becomes:

$$\text{SNR} = \frac{12 K_C^2 V_{\max}^2}{Q^2 K_P} = \frac{3 K_C^2 2^{2B}}{K_P} \tag{1.25}$$

Because the signal-to-noise ratio can vary over a wide range, it is convenient to express it in decibels (dB):

$$\text{SNR(dB)} \triangleq 10 \log_{10} \frac{P_x}{P_q} = 6B + 4.8 + 20 \log_{10} K_C - 10 \log_{10} K_P \tag{1.26}$$

This formula shows that the signal-to-noise ratio (in decibels) increases linearly with the number of bits in the quantizer at a rate of 6 dB/bit. Note also that, for a given quantizer and type of signal, the best signal-to-noise ratio is achieved if the clipping factor is 1, i.e. if the maximum signal amplitude matches the maximum quantization level. It is also apparent that signals that have low peak factors can be quantized with better signal-to-noise ratios than signals with high peak factors. For example, for 12-bit quantization of a sine wave, and optimal clipping factor, the signal-to-noise ratio will be 73.8 dB. Under most circumstances quantization error for a 12-bit quantizer is small relative to other sources of noise.

1.4 Oversampling analog-to-digital converters

In the conventional analog-to-digital techniques described above, the signal is sampled at a rate close to the Nyquist limit using a high-resolution A/D converter (typically 12 bits or more). The last decade has seen the development of new A/D conversion techniques that are characterized by very high sampling rates (many times the Nyquist rates), and low-resolution quantizers, as low as 1 bit in some cases. In effect, these schemes are based on a trade off between sampling rate and quantization noise. Very high sampling rates provide the additional flexibility of shaping the spectrum of the quantization noise so that it lies mostly in frequency regions well above the the signal, leading to an increase in signal-to-noise ratio in the frequency range of interest.

1.4.1 Oversampling

Assume that we want to sample a signal bandlimited to W Hz. From Nyquist's theorem, the sampling rate F_s should be higher than $2W$, and, to ensure efficient processing and storage, it is best that F_s not be much greater than the Nyquist limit. For example, audio signals, which are bandlimited to 20 kHz, are typically sampled at rates of 44-48 kHz. In oversampling A/D systems, this operation is accomplished in two steps (Fig. 1.9). First, the signal is sampled at a rate KF_s much higher than the Nyquist limit, i.e. $K >> 1$. Second, this high-rate signal is digitally downsampled or decimated to the desired rate F_s. In order to prevent aliasing, the downsampling operation must be preceded by an antialiasing filter with a cutoff frequency $F_c < F_s/2$. This antialiasing filter plays the same role as in conventional A/D conversion, with the important difference that it is implemented digitally rather than using analog circuitry. Thus, in oversampling systems, there are two antialiasing filters, an analog one preceding the high-rate A/D converter, and a digital filter preceding the decimation stage.

The oversampling technique has two major advantages over conventional A/D systems.

1. The requirements on the *analog* antialiasing filter are much less severe than in conventional systems. In conventional systems, there is only a narrow frequency band between the upper frequency limit W of the signal and the Nyquist frequency $F_s/2$, so that a "brickwall" filter with sharp cutoffs is required. In oversampling systems, the Nyquist frequency $KF_s/2$ is much greater than W, so that a filter with gentle roll-off suffices. In many cases, it is possible to use a simple RC circuit for antialiasing in oversampling systems. On the other

hand, the requirements on the *digital* antialiasing filter preceding the decimation stage are quite stringent. This, however, is much less of a problem because it is usually easier and less costly to design digital filters with sharp cutoffs than analog filters. Thus, oversampling systems trade complexity in analog design for complexity of digital design.

2. Oversampling systems achieve moderate improvements in signal-to-noise ratio over conventional systems. This can be understood from a frequency-domain analysis of the quantization noise. If the input signal is sufficiently irregular, the quantization noise can be modeled as white noise, meaning that it occupies the entire frequency band of the A/D converter. Because the total noise power depends only on the quantization step, the wider the bandwidth, the smaller the quantization noise *per unit bandwidth.* In conventional systems, the bandwidth of the quantization noise is $F_s/2$, while in oversampling systems, it is $KF_s/2$. Thus, the signal-to-noise ratio *in the frequency band of interest* $(F < W)$ will be K times greater in oversampling systems than in conventional systems, even though the number of bits in the quantizer remains the same. The additional quantization noise in frequency regions above $F_s/2$ is effectively removed by the digital antialiasing filter. Thus, the digital filter has a dual function in oversampling systems, that of antialiasing, and that of eliminating high-frequency quantization noise.

Despite these advantages, oversampling systems as described above are not practical for applications such as audio signal processing. Consider specifically an audio signal sampled at 48 kHz using a 16-bit conventional A/D converter. With an oversampling system using a sampling rate of 3 MHz (64 times oversampling), one would need a 13-bit A/D converter to obtain the same signal-to-noise ratio. Such combination of high rates and high resolution are very costly, if not beyond the limit of current technology. Thus, it is necessary to reduce the resolution of the A/D converter if oversampling is to be used. This can be accomplished with no loss in signal-to-noise ratio by moving most the quantization noise to frequencies above the region of interest, a technique known as *noise shaping.*

1.4.2 Noise shaping

The most widely used technique for noise shaping is *sigma-delta modulation.* As shown in Fig. 1.10, sigma-delta modulation consists of A/D converting the integral (sigma) of the difference (delta) between the original signal and the quantized signal. The A/D converter has typically low resolution, as low as 1 bit, but its sampling rate is very high, much as in the oversampling systems described above. If the converter has 1 bit, its output will consist of a sequence of $+1$ and -1 (Fig. 1.11). If the input signal is near zero, there will be an equal number of plus and minus output samples. If the input amplitude is increased, the percentage of positive output samples will increase. By averaging over many samples, the digital antialiasing filter that follows the sigma-delta converter is able to reconstruct the input signal.

In the diagram of Fig. 1.10, the A/D and D/A converters are inverse of one another (except for the addition of quantization noise by the A/D converter), so that they can be eliminated from the diagram. Specifically, the sigma-delta modulator can be approximated by the linear model shown in Fig. 1.12. In this model, if $x(t)$ represents the input signal, $y(t)$ the output, and $q(t)$ the quantization noise, one has:

$$y(t) = \int_0^t (x(\tau) - y(\tau))\, d\tau + q(t) \tag{1.27}$$

Converting to the frequency domain, and rearranging terms, we obtain:

$$Y(F) = \frac{X(F)}{1 + j2\pi F} + \frac{j2\pi F Q(F)}{1 + j2\pi F} \tag{1.28}$$

Thus, the output is the sum of a term due to the signal and a term representing the quantization noise. At low frequencies ($F << 1$), the signal term approaches $X(f)$, while the noise term approaches 0. On the other hand, at high frequencies, the noise term becomes much higher than the signal term. Thus, the sigma-delta converter acts as a lowpass filter with respect to the signal, and as a highpass filter with respect to quantization noise. Fig. 1.13 compares the frequency distribution of the quantization noise in oversampling systems with and without noise shaping. The sigma-delta converter effectively moves most the quantization noise beyond the frequency band of the signal, yielding high signal-to-noise ratio despite coarse quantization. Sigma-delta converters are widely used in digital audio applications.

1.5 Summary

The conversion of a continuous-time signal into a digital signal can be decomposed into two conceptual stages, sampling and quantization. For band-limited signals, sampling can be achieved without loss of information providing that the sampling frequency F_s is greater than twice the highest frequency in the signal. In this case, the continuous-time signal can be exactly reconstructed from the samples by means of the interpolation formula:

$$x(t) = \sum_{n=-\infty}^{\infty} x[n]\phi(t - nT_s)$$

where the interpolating function is given by

$$\phi(t) = \frac{\sin \pi F_s t}{\pi F_s t}$$

If a signal is sampled at a rate lower than the Nyquist frequency, frequency components above half the Nyquist frequency get aliased into low-frequency components, resulting in an unrecoverable loss of information. Aliasing can (and should) always be avoided by lowpass filtering the signal before sampling.

Quantization consists in converting a continuous-amplitude signal into a discrete-valued signal. A measure of the degradation produced by quantization is the ratio of the signal power to the power in the quantization error signal. This signal-to-noise ratio increases linearly with the number of bits in the quantizer at a rate of 6 dB per bit, and is largest when the peak signal amplitude matches the voltage range of the quantizer. In many circumstances, commercial quantizers provide signal-to-noise ratios that exceed the signal-to-noise of the measurement, so that quantization is of little consequence in practice. Oversampling combined with noise shaping make it possible to achieve high signal-to-noise ratio with low-resolution quantizers.

1.6 Further reading

Siebert, Chapter 12, Sections 3-5; Chapter 14, Section 3
Oppenheim and Schafer, Chapter 4, Sections 1, 3, 8, and 9; Chapter 6, Section 6
Oppenheim, Willsky and Nawab, Chapter 7, Sections 1-3
Karu, Chapters 19 and 20

1.A.1 Fourier series

1.A.1.1 Sine-cosine form

The basic idea behind Fourier series is that a periodic signal with period T can be expressed as a sum of sine waves at frequencies $1/T, 2/T, 3/T, ..., k/T, ...$ Specifically, consider how we might approximate a periodic signal $x(t)$ by the sum of sinusoids $\hat{x}(t)$:

$$\hat{x}(t) \triangleq \sum_{k=0}^{K} A_k \cos(2\pi kt/T) + \sum_{k=1}^{K} B_k \sin(2\pi kt/T) \tag{1.A.1}$$

To obtain the best approximation, we minimize the average power P_e of the error signal $e(t) \triangleq x(t) - \hat{x}(t)$ over one period of $x(t)$:

$$P_e \triangleq \frac{1}{T}\int_0^T e(t)^2 dt = \frac{1}{T}\int_0^T \left(x(t) - \sum_{k=0}^{K} A_k \cos(2\pi kt/T) - \sum_{k=1}^{K} B_k \sin(2\pi kt/T)\right)^2 dt \tag{1.A.2}$$

This technique of minimizing the power in an error signal representing the difference between the desired signal and an approximation (or estimate) is very widely used in signal processing, and is called *least-squares estimation.* To minimize P_e we set to zero its partial derivatives with respect to each of the A_k and B_k coefficients.

$$\frac{\partial P_e}{\partial A_l} = \frac{2}{T}\int_0^T \cos(2\pi lt/T)\left(x(t) - \sum_{k=0}^{K} A_k \cos(2\pi kt/T) - \sum_{k=1}^{K} B_k \sin(2\pi kt/T)\right) dt = 0 \tag{1.A.3a}$$

$$\frac{\partial P_e}{\partial B_l} = \frac{2}{T}\int_0^T \sin(2\pi lt/T)\left(x(t) - \sum_{k=0}^{K} A_k \cos(2\pi kt/T) - \sum_{k=1}^{K} B_k \sin(2\pi kt/T)\right) dt = 0 \tag{1.A.3b}$$

Simplifying and rearranging terms, this becomes

$$\int_0^T x(t)\cos(2\pi lt/T)dt = \sum_{k=0}^{K} A_k \int_0^T \cos(2\pi kt/T)\cos(2\pi lt/T)dt + \sum_{k=1}^{K} B_k \int_0^T \sin(2\pi kt/T)\cos(2\pi lt/T)dt \tag{1.A.4a}$$

$$\int_0^T x(t)\sin(2\pi lt/T)dt = \sum_{k=0}^{K} A_k \int_0^T \cos(2\pi kt/T)\sin(2\pi lt/T)dt + \sum_{k=1}^{K} B_k \int_0^T \sin(2\pi kt/T)\sin(2\pi lt/T)dt \tag{1.A.4b}$$

It can be shown that

$$\int_0^T \cos(2\pi kt/T)\cos(2\pi lt/T)dt = \begin{bmatrix} T & \text{if } k = l = 0 \\ T/2 & \text{if } k = l > 0 \\ 0 & \text{if } k \neq l \end{bmatrix} \quad (1.A.5a)$$

$$\int_0^T \sin(2\pi kt/T)\cos(2\pi lt/T)dt = 0 \quad 0 \leq k,\ l \leq Ko \quad (1.A.5b)$$

$$\int_0^T \sin(2\pi kt/T)\sin(2\pi lt/T)dt = \begin{bmatrix} 0 & \text{if } k = l = 0 \\ T/2 & \text{if } k = l > 0 \\ 0 & \text{if } k \neq l \end{bmatrix} \quad (1.A.5c)$$

Because these integrals are zero except when the two signals within the integral are the same, it is said that the sine and cosine functions for frequencies that are multiple of $1/T$ are *orthogonal* over the interval $[0, T]$. This term is justified by analogy with orthogonal vectors whose dot product is equal to zero. Replacing the integrals in (1.A.4) by their values from (1.A.5), we obtain:

$$A_0 = \frac{1}{T}\int_0^T x(t)dt \quad (1.A.6a)$$

$$A_l = \frac{2}{T}\int_0^T x(t)\cos(2\pi lt/T)dt \quad 1 \leq l \leq K \quad (1.A.6b)$$

$$B_l = \frac{2}{T}\int_0^T x(t)\sin(2\pi lt/T)dt \quad 1 \leq l \leq K \quad (1.A.6c)$$

These formulas for computing coefficients that provide the best approximation to $x(t)$ do not change when the number of coefficients in the approximation increases. This is a consequence of the orthogonality of the sine and cosine functions.

Fourier's theorem states that, if the function $x(t)$ is reasonably well-behaved (i.e. if it does not have "too many" discontinuities), the average power P_e in the error signal tends to zero when the number of coefficients K increases. This means that the estimate $\hat{x}(t)$ becomes arbitrarily close to $x(t)$, and in the limit one can write

$$x(t) = \sum_{k=0}^{\infty} A_k \cos(2\pi kt/T) + \sum_{k=1}^{\infty} B_k \sin(2\pi kt/T) \quad (1.A.7)$$

1.A.1.2 Exponential form

Equations (1.A.6) and (1.A.7) are called the *sine-cosine form* of Fourier series. Simpler equations can be obtained if we introduce the *exponential form*, which is based on the complex Fourier series coefficients X_k:

$$X_0 \triangleq A_0 \quad (1.A.8a)$$

$$X_k \triangleq (A_k - jB_k)/2 \quad k \geq 1 \quad (1.A.8b)$$

$$X_{-k} \triangleq (A_k + jB_k)/2 \quad k \geq 1 \quad (1.A.8c)$$

Conversely, the sine and cosine coefficients can be computed from the complex Fourier coefficients:

$$A_0 = X_0 \quad (1.A.9a)$$

$$A_k = X_k + X_{-k} \qquad k \geq 1 \tag{1.A.9b}$$

$$B_k = j(X_k - X_{-k}) \qquad k \geq 1 \tag{1.A.9b}$$

Replacing the A_k and B_k in (1.A.6) and (1.A.7) by their expressions as function of X_k, we obtain the exponential form of the Fourier series:

$$x(t) = \sum_{k=-\infty}^{\infty} X_k e^{j2\pi kt/T} \tag{1.A.10}$$

with

$$X_k = \frac{1}{T}\int_0^T x(t)e^{-j2\pi kt/T}dt \tag{1.A.11}$$

Note that the summation in (1.A.10) is from $-\infty$ to ∞, and that there is a minus sign in the argument of the complex exponential in (1.A.11).

1.A.1.3 Properties

Fourier series have useful symmetry properties that are easily verified:

$$x(t) \text{ is real} \quad \text{iff} \quad X_{-k} \ = \ X_k^*$$

X_k^* denotes the complex conjugate of X_k. Assuming now that $x(t)$ is real,

$$x(-t) \ = \ x(t) \quad \text{iff} \quad X_k \text{ is real and } X_{-k} \ = \ X_k$$

$$x(-t) \ = \ -x(t) \quad \text{iff} \quad X_k \text{ is imaginary and } X_{-k} \ = \ -X_k$$

$$x(t - T/2) \ = \ -x(t) \quad \text{iff} \quad X_k \ = \ 0, \text{ for } k \text{ even}$$

Another property, known as *Parseval's theorem,* is that the mean power in one period of $x(t)$ is equal to the sum of the powers of the Fourier components:

$$P_x \triangleq \frac{1}{T}\int_0^T x(t)^2 dt = \sum_{k=-\infty}^{\infty} |X_k|^2 \tag{1.A.12}$$

where $|X|$ denotes the magnitude of X.

1.A.2 Sampling theorem and interpolation formula for periodic signals

In this section, we derive the Nyquist sampling theorem in the special case of periodic signals. This case is particularly simple because a bandlimited periodic signal is completely characterized by a finite number of Fourier series coefficient. On this other hand, the periodicity assumption is not as restrictive as it seems because any finite-duration signal can be considered as one period of a periodic signal. The derivation makes use of the Fourier series intriduces in Section 1.A.1.

1.A.2.1 Discrete Fourier series

Let $x(t)$ be a periodic signal that contains no frequency components at and above W, and let K be the largest integer such that $K/T < W$. All Fourier coefficients X_k must be zero for $|k| > K$, otherwise the signal would not be bandlimited to W. Therefore, $x(t)$ can be decomposed into a sum of a finite number N of complex exponentials, where $N \triangleq 2K + 1$ in the general case, and $N \triangleq 2K$ in the special case when the highest-frequency component X_K is purely cosinusoidal, and therefore real:

$$x(t) = \sum_{k=-[N/2]}^{[N/2]} X_k e^{j2\pi kt/T} \tag{1.A.13}$$

where $[N/2]$ denotes the integer part of $N/2$, i.e. $(N-1)/2$ if N is odd, and $N/2$ if N is even. The sampling theorem for periodic signals states that $x(t)$ can be exactly reconstructed from N samples at intervals of $T_s = T/N$. In other words, the sampling frequency $F_s = N/T$ must be more than twice the highest frequency component K/T in the signal.

To prove this theorem, we will first show that the Fourier coefficients X_k can be computed from N samples of $x(t)$. The signal at any time t will then be reconstructed by replacing the X_k in (1.A.13) by their expression as a function of the N samples. Figure 1.2a illustrates the relations between the continuous-time signal $x(t)$, the sampled signal $x[n]$ and the Fourier coefficients X_k.

Equation (1.A.13) holds for any time t, including the sampling times nT/N:

$$x[n] = x(nT/N) = \sum_{k=-[N/2]}^{[N/2]} X_k e^{j2\pi kn/N} \qquad 0 \leq n \leq N-1 \tag{1.A.14}$$

This provides a set of N linear equations for the N Fourier coefficients X_k as a function of the N samples $x[n]$*. [1] We will show presently that the solution to this set of equations is:

$$X_k = \frac{\alpha_k}{N} \sum_{n=0}^{N-1} x[n] e^{-j2\pi kn/N} \tag{1.A.15a}$$

where

$$\alpha_k = \begin{bmatrix} 1/2 & \text{if } |k| = N/2 \ \text{(requiring N even)} \\ 1 & \text{if } |k| < N/2 \end{bmatrix} \tag{1.A.15b}$$

The relations (1.A.14) and (1.A.15) are called *discrete Fourier series.* They are the equivalent for discrete, periodic signals of relations (1.A.10) and (1.A.11) for continuous-time signals. Note the symmetry between (1.A.14) and (1.A.15), the only differences being the signs of the arguments of complex exponentials and the $1/N$ factor in (1.A.15). These relations are the basis for the *discrete Fourier transform,* which is one of the most useful tools in digital signal processing, and will be studied in detail in Chapter 4.

[1] *Because the Fourier coefficients are complex, strictly speaking, one should write equations separately for the real and imaginary parts. However, due to symmetry relations between coefficients (specifically, $X_{-k} = X_k^*$), this gives only N independent real equations.

1.A.2.2 Periodic interpolation formula

We have not shown that (1.A.15) is the solution to the set of equations (1.A.14). A direct proof is simple (it involves the orthogonality of discrete-time complex exponentials), but we will not do it. Rather, we will assume that (1.A.15) holds, and use it to express the original signal $x(t)$ as a function of its samples $x[n]$. In the special case when $t = nT_s$, this interpolation formula will verify that (1.A.15) is the solution to (1.A.14). Specifically, let us replace the Fourier coefficients X_k in (1.A.13) by their values as a function of the samples $x[n]$ in (1.A.15):

$$x(t) = \sum_{k=-[N/2]}^{[N/2]} \left[\frac{\alpha_k}{N} \sum_{n=0}^{N-1} x[n] e^{-j2\pi kn/N} \right] e^{j2\pi kt/T} \tag{1.A.16}$$

Interchanging the order of summations over k and n, and noting that $1/N = T_s/T$, this becomes:

$$x(t) = \sum_{n=0}^{N-1} x[n] \phi_N(t - nT_s) \tag{1.A.17}$$

with

$$\phi_N(t) \triangleq \frac{1}{N} \sum_{k=-[N/2]}^{[N/2]} \alpha_k e^{j2\pi kt/T} \tag{1.A.18}$$

Formula (1.A.17) expresses the signal $x(t)$ as a weighted sum of delayed interpolating functions $\phi_N(t)$. The weights are the samples $x[n]$, and the delays multiples of the sampling interval nT_s. A closed-form expression for $\phi_N(t)$ is easily obtained by making use of the formula for the partial sum of a geometric series

$$\sum_{k=0}^{N-1} z^k = \frac{z^N - 1}{z - 1}$$

where $z \triangleq e^{j2\pi t/T}$.

$$N\phi_N(t) = \sum_{k=-[N/2]}^{[N/2]} \alpha_k z^k = \left[\begin{array}{ll} \frac{z^{N/2} - z^{-N/2}}{z^{\frac{1}{2}} - z^{-\frac{1}{2}}} & \text{if } N \text{ is odd} \\ \frac{(z^{N/2} - z^{-N/2})(z^{\frac{1}{2}} + z^{-\frac{1}{2}})}{2(z^{\frac{1}{2}} - z^{-\frac{1}{2}})} & \text{if } N \text{ is even} \end{array} \right.$$

Replacing z by its value, and making use of the fact that $e^{j\theta} - e^{-j\theta} = 2j \sin\theta$, we obtain the final formula for the interpolating function:

$$\phi_N(t) = \left[\begin{array}{ll} \frac{\sin(\pi Nt/T)}{N \sin(\pi t/T)} = \frac{1}{N} \frac{\sin(\pi F_s t)}{\sin(\pi F_s t/N)} & \text{if } N \text{ is odd} \\ \frac{\sin(\pi Nt/T)\cos(\pi t/T)}{N \sin(\pi t/T)} = \frac{1}{N} \sin(\pi F_s t) \text{cotan}(\pi F_s t/N) & \text{if } N \text{ is even} \end{array} \right. \tag{1.A.19}$$

The function $\phi_N(t)$ is illustrated in Fig. 1.3 for $N = 9$. It is periodic with the same period $T = NT_s$ as $x(t)$, and verifies the property

$$\phi_N(nT_s) = \begin{cases} 1 & \text{if } n \text{ is a multiple of } N \\ 0 & \text{otherwise} \end{cases} \tag{1.A.20}$$

These properties imply that $x(t) = x[n]$ for $t = nT_s$, thereby confirming that (1.A.15) is indeed the solution to the set of equations (1.A.14).

1.A.2.3 Aliasing of a periodic signal

We will now examine what happend when the periodic signal is no longer band-limited so that the Nyquist criterion is not met. Specifically, suppose we take samples at frequency $F_s = N/T$ of the wide-bandwidth signal:

$$x(t) \triangleq \sum_{k=-\infty}^{\infty} X_k e^{j2\pi kt/T} \qquad (1.A.21)$$

By making use of the relation (1.A.15), one can as before define N discrete Fourier coefficients from N samples $x[n]$. However, we anticipate that, due to aliasing, the discrete Fourier coefficients of $x[n]$ will no longer be the same as the coefficients X_k of the continuous-time signal $x(t)$, so they will be denoted $\tilde{X}_k, -[N/2] \leq k \leq [N/2]$. The relations between the continuous-time signal $x(t)$, the samples $x[n]$, the continuous Fourier coefficients X_k and the discrete Fourier coefficients $\tilde{X}_k$ is illustrated in Fig. 1.2b. To relate the Fourier coefficients of the discrete and the continuous signals, we express the samples both as a finite sum, as in (1.A.14), and as an infinite sum as in (1.A.21). From (1.A.14):

$$x[n] = \sum_{k=-[N/2]}^{[N/2]} \tilde{X}_k e^{j2\pi kn/N} \qquad (1.A.22)$$

From (1.A.21),

$$x[n] = x(nT/N) = \sum_{l=-\infty}^{\infty} X_l e^{j2\pi ln/N} \qquad (1.A.23)$$

The infinite sum in (1.A.23) can be decomposed into finite sums over N coefficients. Specifically, l can be expressed uniquely in the form $k + rN$, where r is an integer, and $-[N/2] \leq k \leq [N/2]$*. [2] Eq. (1.A.23) becomes:

$$x[n] = \sum_{r=-\infty}^{\infty} \sum_{k=-[N/2]}^{[N/2]} X_{k+rN} e^{j2\pi(k+rN)n/N} \qquad (1.A.24)$$

Interchanging the order of summations, and making use of the 2π periodicity of the complex exponential, (1.A.24) yields:

$$x[n] = \sum_{k=-[N/2]}^{[N/2]} \left[\sum_{r=-\infty}^{\infty} X_{k+rN} \right] e^{j2\pi kn/N} \qquad (1.A.25)$$

Comparing (1.A.25) with (1.A.24) gives the desired relation between the Fourier coefficients of the discrete and continuous-time signals:

$$\tilde{X}_k = \sum_{r=-\infty}^{\infty} X_{k+rN} \qquad (1.A.26)$$

That is, the k^{th} Fourier coefficient of the sampled signal $x[n]$ is equal to the k^{th} coefficient of the original signal $x(t)$ plus the sum of all the aliased components that are of order k plus a multiple of N. The relation between $\tilde{X}_k$ and X_k is illustrated in Fig 1.7 for different values of N. As the number of samples decreases, the discrete Fourier coefficients $\tilde{X}_k$ increasingly deviate from the continuous coefficients X_k. In the special case when Nyquist's criterion is met, $X_k = 0$ for $|k| > [N/2]$, so that $\tilde{X}_k$ is equal to X_k, i.e. the Fourier coefficients are not affected by aliasing.

[2] *For simplicity, we assume that N is odd in this derivation.

1.A.2.4 Generalization of the interpolation formula to arbitrary signals

The interpolation formulas (1.A.17) and (1.A.19) are valid only for periodic signals. Similar formulas can be obtained for arbitrary band-limited signals. For deriving these more general formulas, we will let the period of the signal tend to infinity. Then, the number of samples N in (1.A.17) must also go to infinity because, if W is the largest frequency component in the signal, one must have $N = F_s T > 2WT$ in order to verify Nyquist's condition. Because the signal is assumed to be periodic, (1.A.17) can be written as:

$$x(t) = \sum_{n=-[N/2]}^{[N/2]} x[n]\phi_N(t - nT_s)$$

As the period tends to infinity, this formula becomes

$$x(t) = \sum_{n=-\infty}^{\infty} x[n]\phi(t - nT_s) \qquad (1.A.27)$$

with

$$\phi(t) \triangleq \lim_{N\to\infty} \phi_N(t) = \lim_{N\to\infty} \frac{1}{N}\frac{\sin(\pi F_s t)}{\sin(\pi F_s t/N)} \qquad (1.A.28)$$

(1.A.29) As N becomes very large, while t remains fixed, one has

$$N \sin(\pi F_s t/N) \approx \pi F_s t$$

so that

$$\phi(t) = \frac{\sin(\pi F_s t)}{\pi F_s t} \qquad (1.A.30)$$

Together, Equations (1.A.27) and (1.A.30) give the general interpolation formula (1.12).

Formula (1.A.27) holds for arbitrary band-limited signals, including periodic ones. If (1.A.17) and (1.A.27) are to be both true for periodic signals, one must have:

$$\phi_N(t) = \sum_{r=-\infty}^{\infty} \phi(t - rNT_s) \qquad (1.A.31)$$

It can be shown that this is indeed the case.

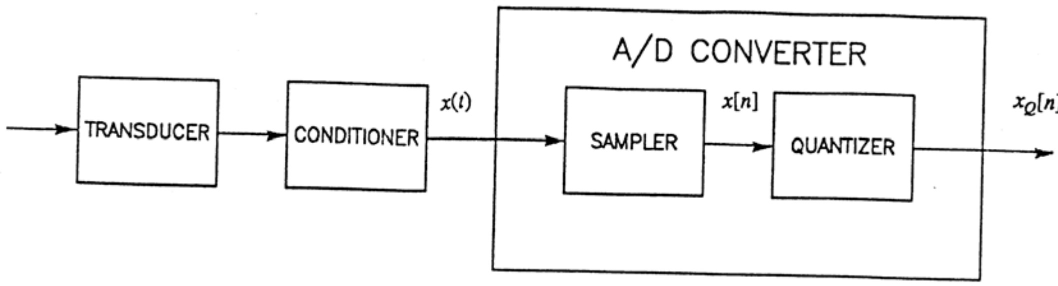

Figure 1.1:

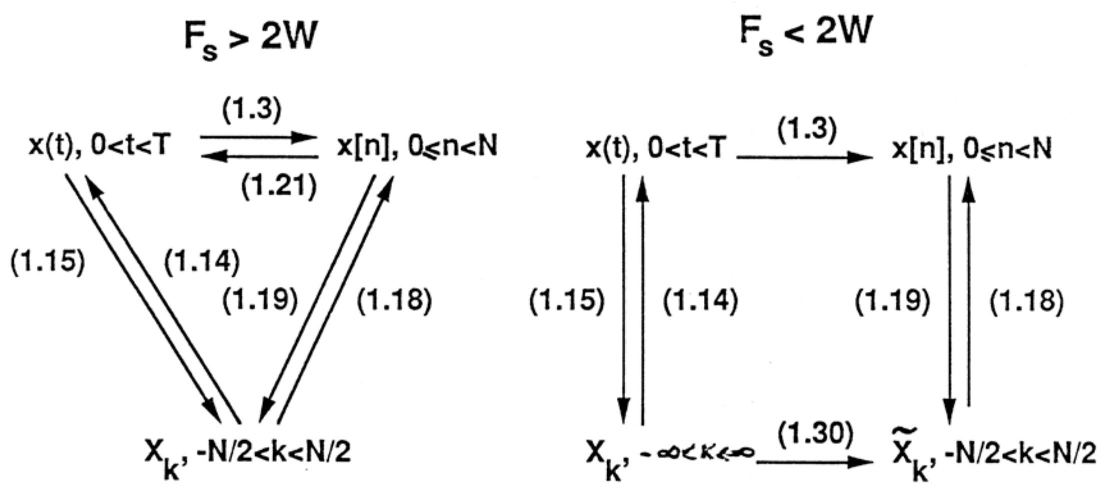

Figure 1.2: A. Relation between the continuous time signal $x(t)$, the sampled signal $x[n]$ and the Fourier coefficients X_k when F_s is above Nyquist's rate. Numbers refer to formulas in the text. B. Same as A. when F_s is below Nyquist's rate. $\tilde{X}_k$ refers to the Fourier coefficients of the sampled signal which now differ from those of $x(t)$.

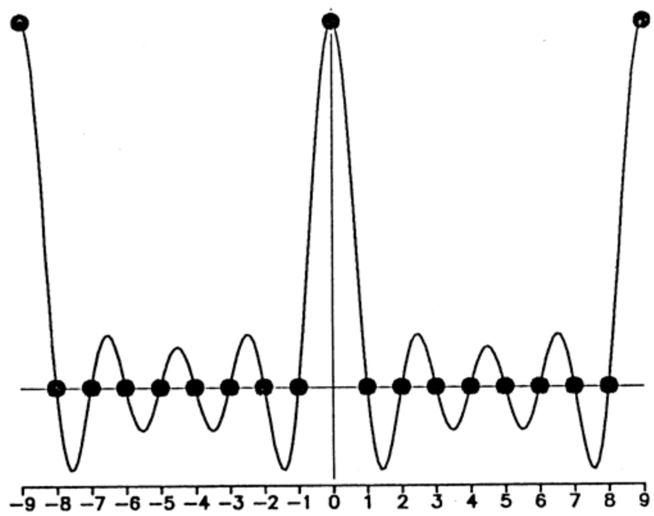

Figure 1.3: The periodic interpolating function $\phi_N(t)$ for $N = 9$ and $F_s = 1$.

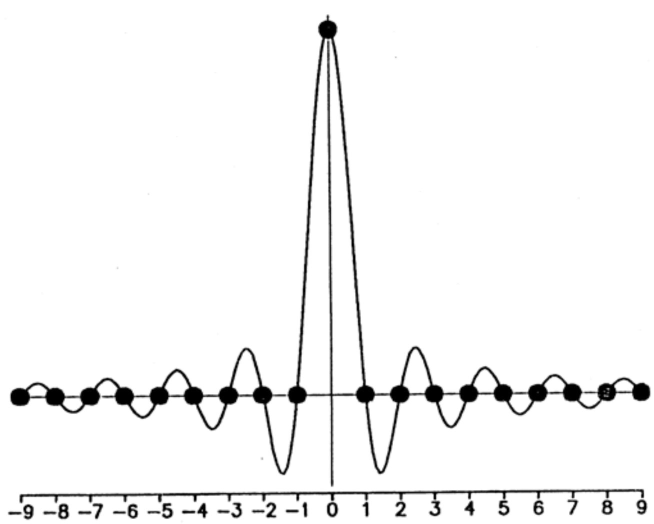

Figure 1.4: The general interpolating function $\phi(t)$ for the same sampling frequency as in Fig 1.2.

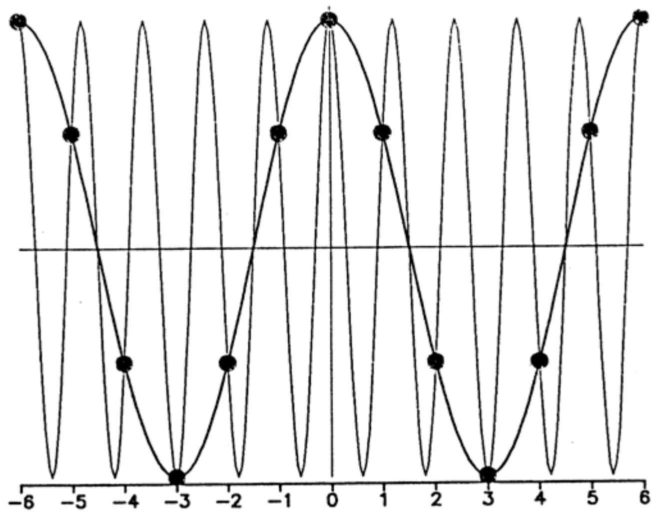

Figure 1.5: Aliasing: Two sinewaves with different frequencies that have the same samples.

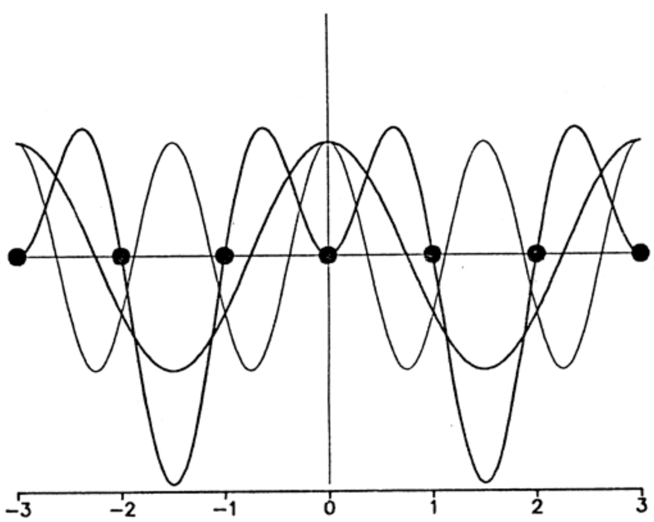

Figure 1.6: A two-component signal whose samples are identically zero due to aliasing of the high-frequency component.

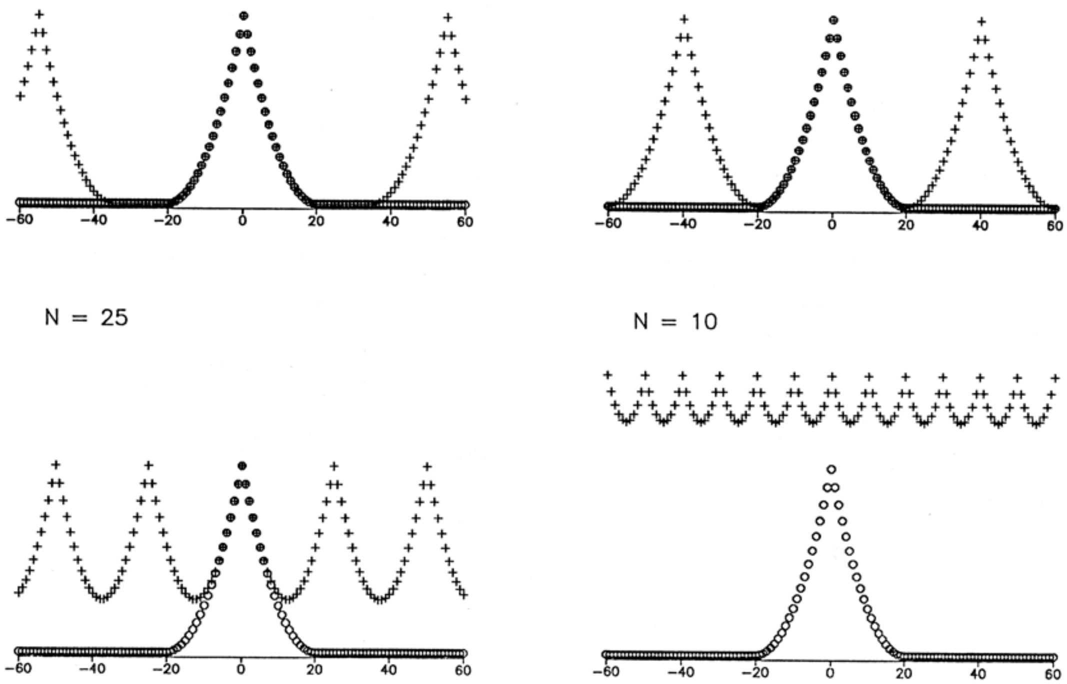

Figure 1.7: Effects of sampling at 4 different frequencies on the spectrum of a periodic signal with 39 components. Crosses show the Fourier coefficients of the sampled signal, circles those of the continuous-time signal.

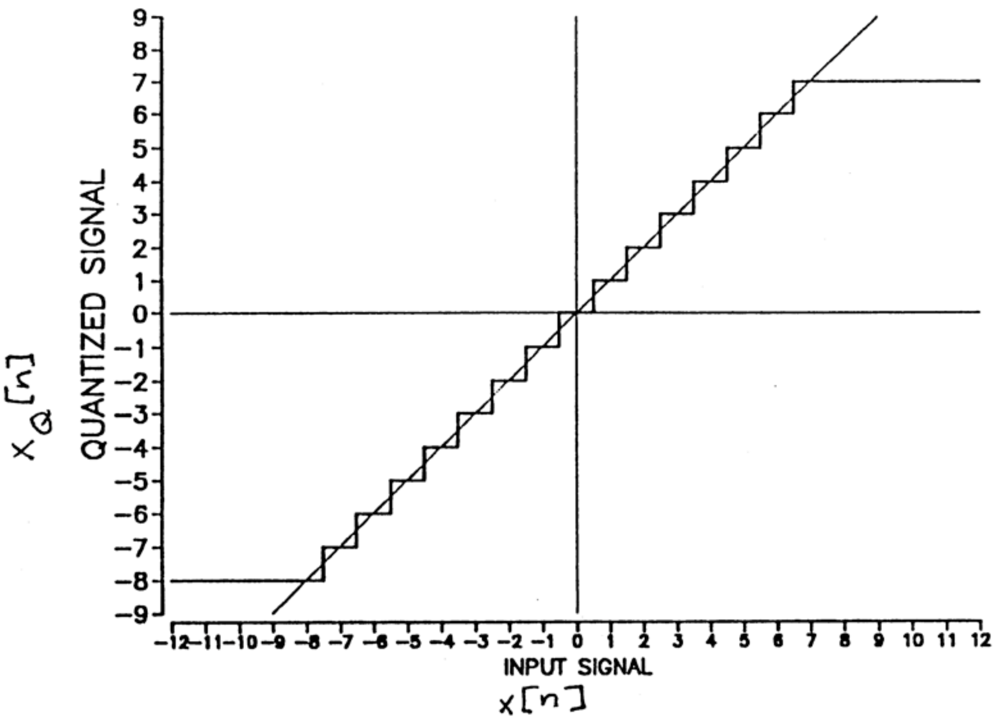

Figure 1.8: Input-output function for a 4-bit quantizer.

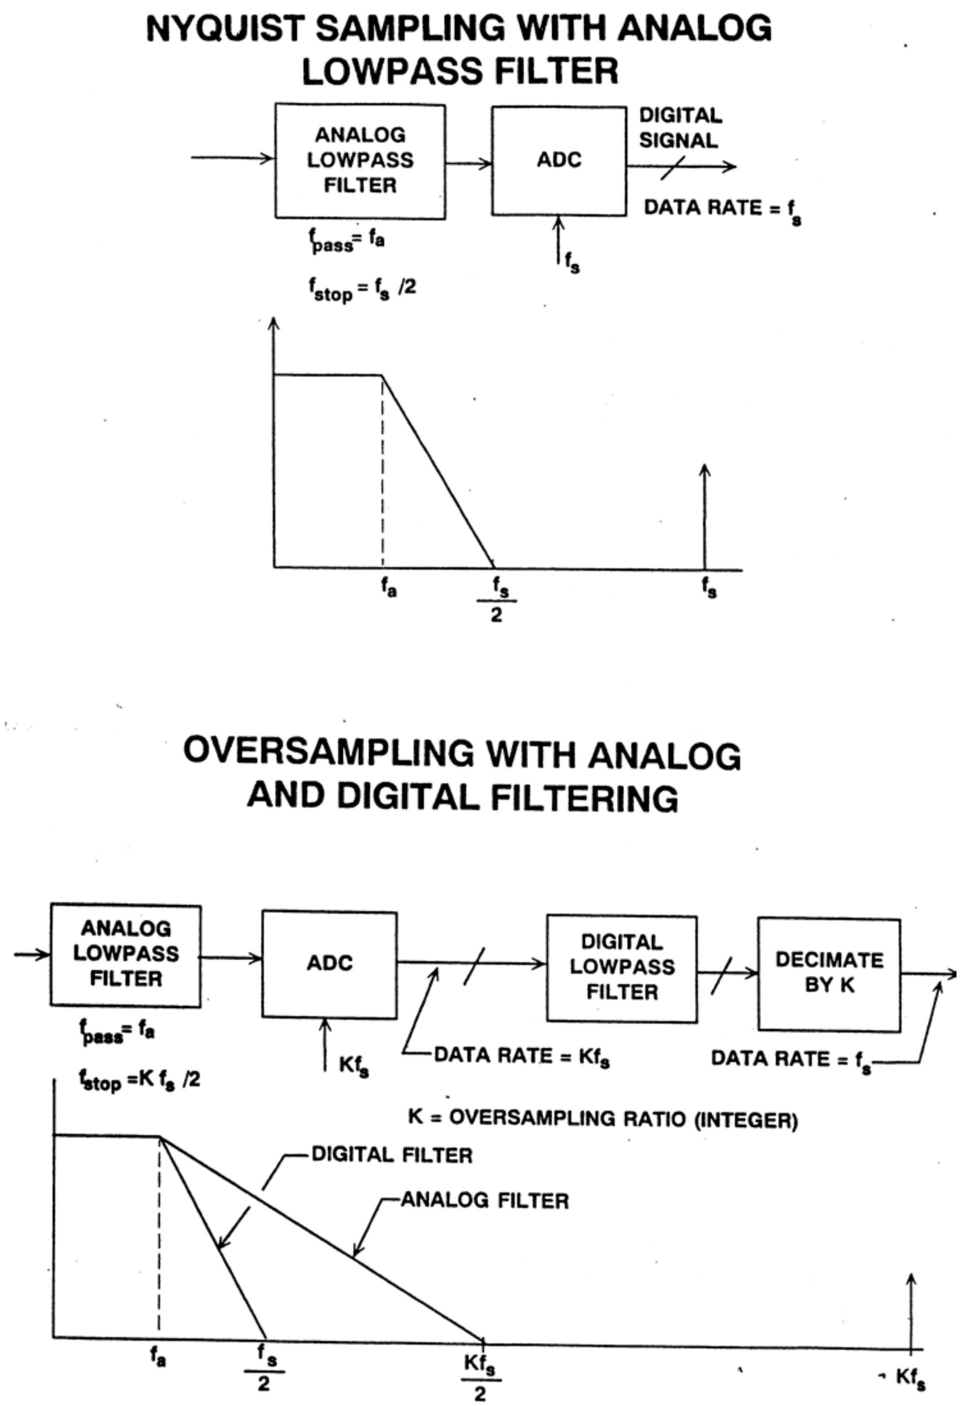

Figure 1.9: (from *Analog Devices*) Block diagrams of conventional and oversampling A/D systems.

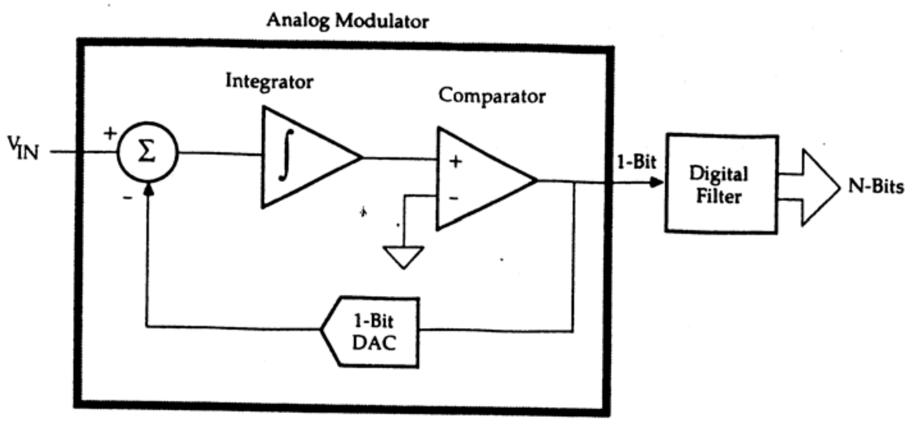

Figure 1.10: (from *Analog Devices*) Block diagram of sigma-delta modulation A/D system.

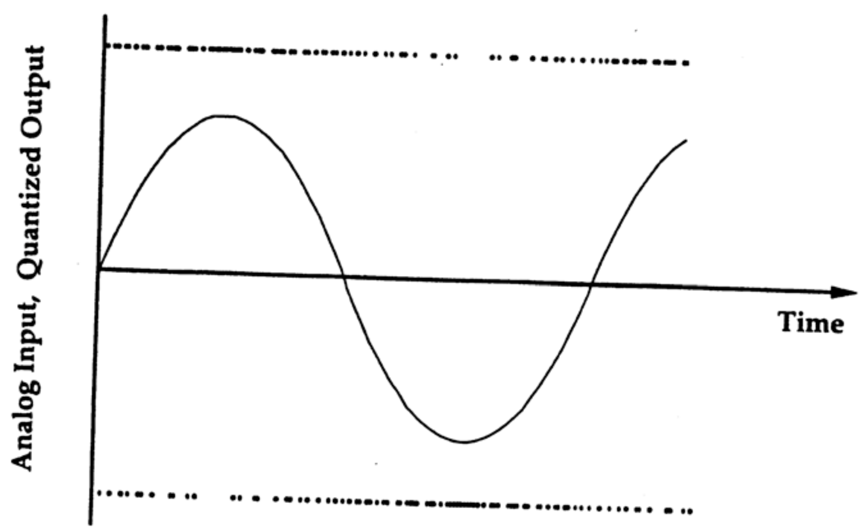

Figure 1.11: (from *Analog Devices*) Input and output signals for 1-bit sigma-delta modulator.

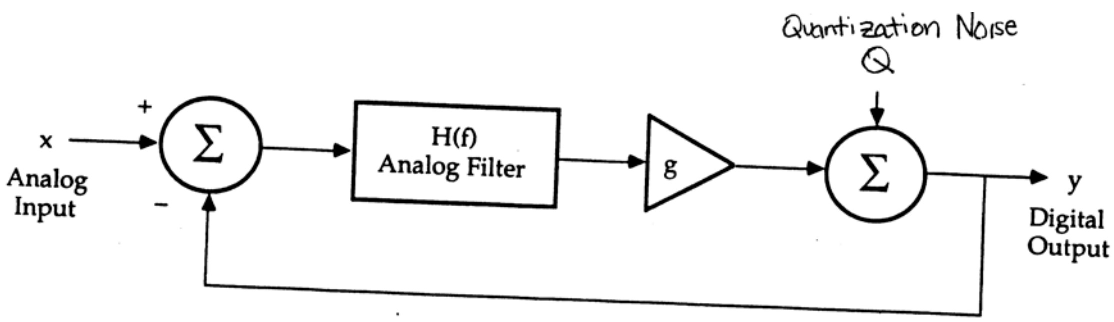

Figure 1.12: (from *Analog Devices*) Linear-system model for sigma-delta modulator.

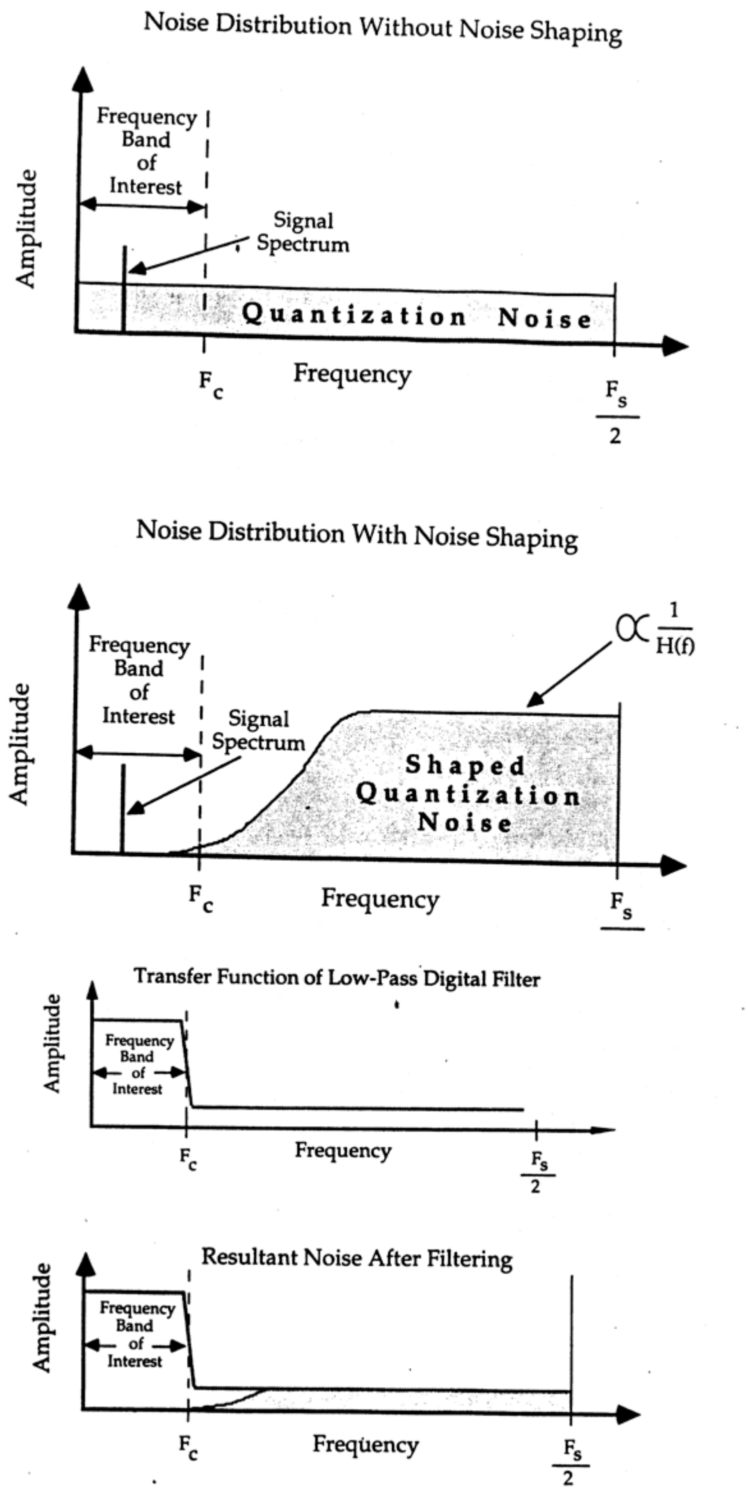

Figure 1.13: (from *Analog Devices*) Spectrum of quantization noise with and without noise shaping.

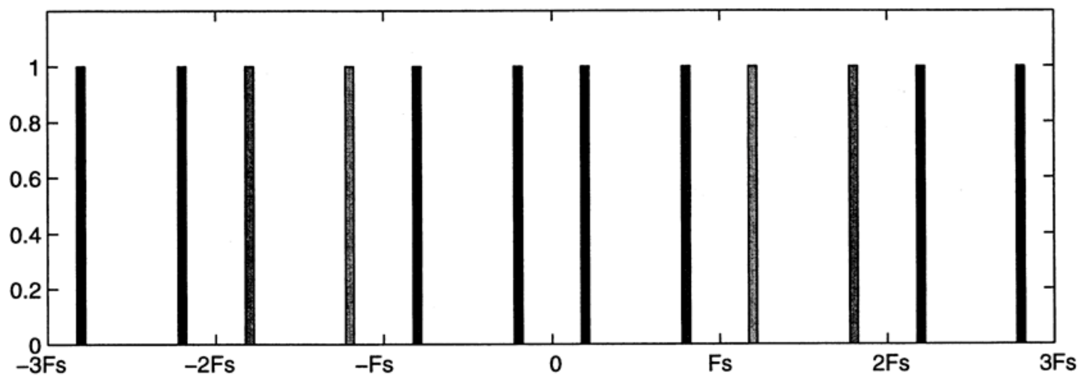

Figure 1.14: Sampling of an analog sine wave, $F_s = 5F$

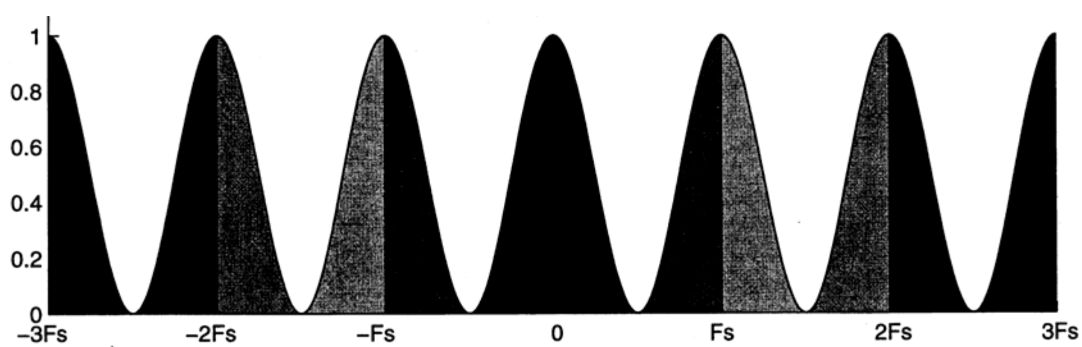

Figure 1.15: Sampling of a bandlimited signal, $F_s = 2W$