Chapter 8

Sampling, Aliasing, and Data Conversion

8.1. Introduction

The process of converting a continuous-time signal to a sequence of numbers is called sampling. Sampling is the fundamental operation of DSP, and avoiding or at least minimising aliasing is the most important aspect of sampling.

8.2. Nyquist Sampling Theorem

A continuous-time signal $x(t)$ is called *band limited* if its Fourier transform is zero outside a certain frequency range. Suppose we sample a band-limited signal and choose the sampling frequency such that $f_{sam} \geq 2f_m$. The spectra of the continuous-time signal and of the sampled signal are exemplified in Figure 8.1. Notice that in this case the replicas in the sampled signal do not overlap. This is the principle of the Nyquist rate of sampling.

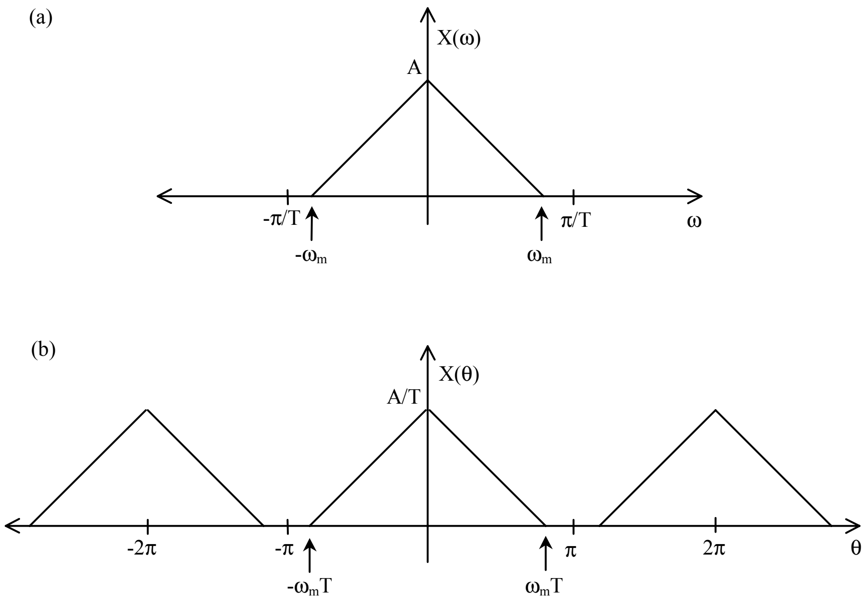

Figure 8.1: Sampling of a band limited signal above the Nyquist rate: (a) Fourier transform of the continuous-time signal; (b) Fourier transform of the sampled signal.

The *Nyquist sampling rate*, or the critical sampling rate for a band limited signal is mathematically defined by Equation 8.1, where f_m is the frequency of the signal.

$$f_{sam} = \frac{1}{T} = 2f_m \tag{8.1}$$

By definition, if a band limited signal is sampled at a rate equal to or greater than the Nyquist sampling frequency $2f_m$. Then the shape of the Fourier transform of the sampled signal in the range [-π, π] is identical to the shape of the Fourier transform of the given signal, except for multiplication of the frequency axis by a factor T, and multiplication of the amplitude axis by a factor $1/T$, as shown in Figure 8.1(b).

8.3. Aliasing

Aliasing is caused by sampling at a rate lower than that of the Nyquist frequency for a given signal. Now let us consider the case of a band limited signal when it is sampled at a lower rate than the Nyquist frequency. As an example, Figure 8.2 illustrates the case of sampling at $f_{sam} = 3f_m/2$. Here, the shape of the Fourier transform of the sampled signal in the range [-π, π] becomes distorted. Distortion occurs in this frequency range because two adjacent replicas overlap and their superposition give rise to the shape illustrated by Figure 8.2 (b).

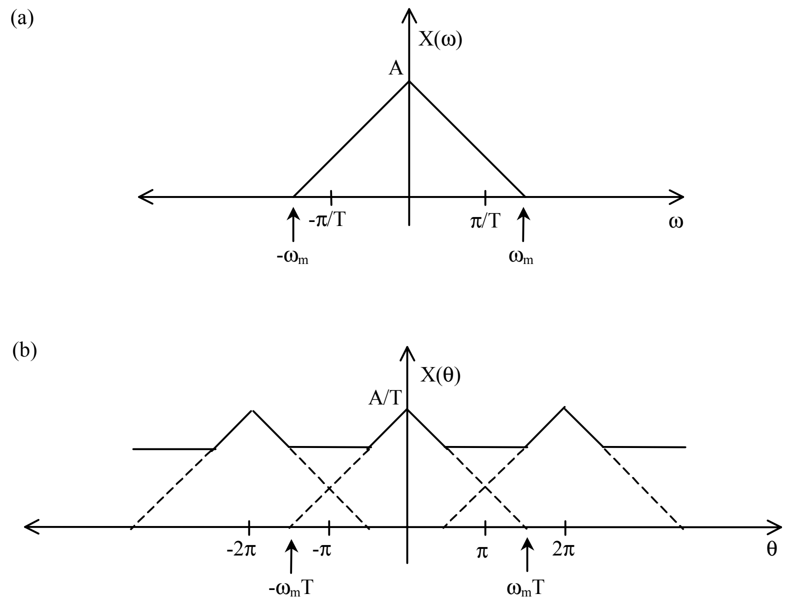

Figure 8.2: Sampling of a band limited signal below the Nyquist rate: (a) Fourier transform of the continuous-time signal; (b) Fourier transform of the sampled signal.

By definition, if a band limited signal is sampled at a rate lower than that of the Nyquist sampling frequency $2f_m$. Then the shape of the Fourier transform of the sampled signal in the range [-π, π] is distorted relative to the Fourier transform of the given signal. This distortion which is more commonly referred to as aliasing, results from overlapping of the replicas as illustrated by Figure 8.2 (b).

8.4. Anti-alias Filter Requirements

A common practise in the sampling of a continuous–time signal is to filter the signal before it is passed to the sampler. The filter used for this purpose is usually an analogue low-pass filter whose cut-off frequency is not larger than half the sampling rate. Such a filter is called an *anti-aliasing* filter. In addition, a further consideration in the design of the anti-alias filter is to implement a so-called *Guard Band*, as illustrated in Figure 8.3.

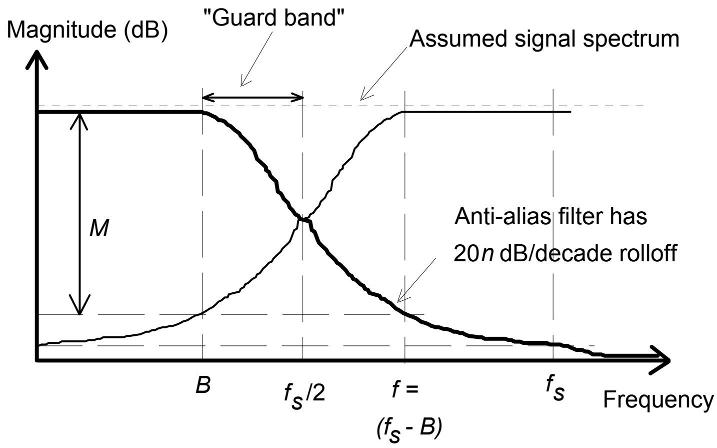

Figure 8.3: Anti-alias filter guard band

Using simple geometry on Figure 8.3, it can be mathematically shown that:

$$M = 20n\log_{10}\left(\frac{f}{B}\right) = 20n\log_{10}\left(\frac{f_{sam} - B}{B}\right) = 20n\log_{10}\left(\frac{f_{sam}}{B} - 1\right) \qquad (8.2)$$

Where B is the signal bandwidth of interest.
M is the maximum allowable size of alias components within B.
n is the order of the anti-alias roll-off filter.
f_{sam} is the sampling frequency.

8.5. Modelling of Quantisation Noise

The sampling of an analogue continuous-time signal is normally implemented using a device called an *analogue-to-digital converter* (A/D). The continuous-time signal is first passed through a device called a *sample-and-hold* (S/H), whose function is to measure the input signal value at the clock instant and hold it fixed for a time interval long enough for the A/D operation to complete. Analogue-to-digital conversion is potentially a slow operation, and a variation of the input voltage during the conversion may disrupt the operation of the converter. The S/H prevents such disruption by keeping the input voltage constant during the conversion. This is schematically illustrated by Figure 8.4.

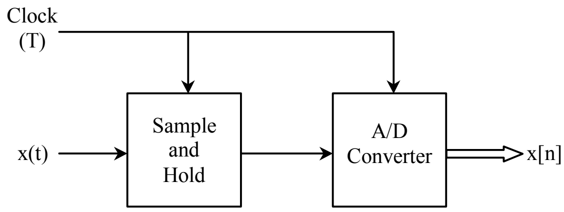

Figure 8.4: Schematic diagram of sampling by an A/D converter.

After a continuous-time signal has been through the A/D converter, the quantised output may differ from the input value. With reference to Figure 8.5, it illustrates that the maximum possible output value after the quantisation process could be up to half the quantisation level q above or q below the ideal output value. This deviation from the ideal output value is called the *quantisation error*. Similarly to A/D converters, D/A converters also exhibit saturation and quantisation errors.

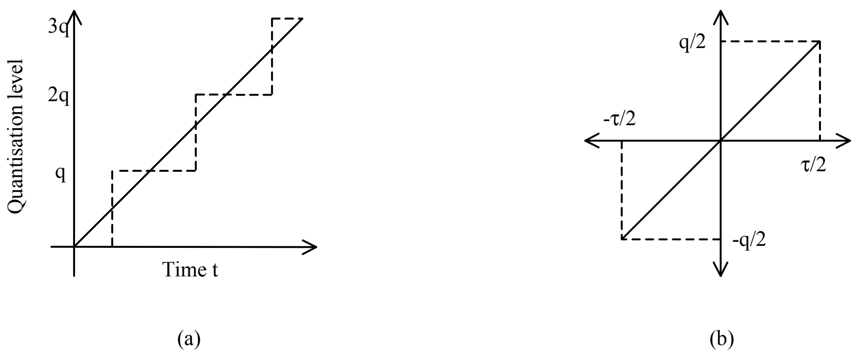

Figure 8.5: (a) Quantisation of the input signal. (b) Quantisation error.

8.6. Derivation of the Mean Square Error of the Quantisation Process

The mean square error (MSE) of the quantisation process can be derived as follows:

$$Error = m \cdot t \qquad -\frac{\tau}{2} \le t \le \frac{\tau}{2}$$

$$MSE = \frac{1}{\tau} \cdot 2 \cdot \int_0^{\tau/2} m^2 \cdot t^2 dt = \frac{2m^2}{\tau} \int_0^{\tau/2} t^2 dt$$

$$MSE = \frac{2m^2}{\tau} \cdot \left[\frac{t^3}{3} \right]_0^{\tau/2} = \frac{m^2 \tau^2}{12} \tag{8.3}$$

With respect to Figure 8.5 (b), the quantisation error at $\tau/2$ is given by $q/2$. Hence,

$$\frac{q}{2} = m \cdot \frac{\tau}{2} \tag{8.4}$$

Re-arranging Equation 8.4 to make m the subject gives:

$$m = \frac{q}{\tau}$$

Now substituting *m* into Equation 8.3 gives the MSE of the quantisation process for a signal in terms of the quantisation step level *q*.

$$MSE = \frac{q^2 \tau^2}{12\tau^2} = \frac{q^2}{12} \tag{8.5}$$

This undesirable noise is sometimes referred to as quantisation *noise power* and is modelled as white noise on top of the output signal, as illustrated by Figure 8.6.

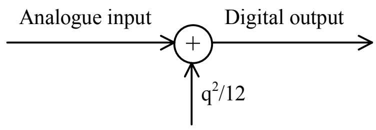

Figure 8.6: Modelling of the quantisation process.

8.7. SNR for a Sine Wave

The number of quantisation bits *N* that are required to produce a given signal-to-noise-ratio (SNR) on a sine wave acquisition can be described as follows. With reference to Figure 8.7, the magnitude of the sine wave is 2A and is quantised with 2^N quantisation levels, with a quantisation step size of *q*.

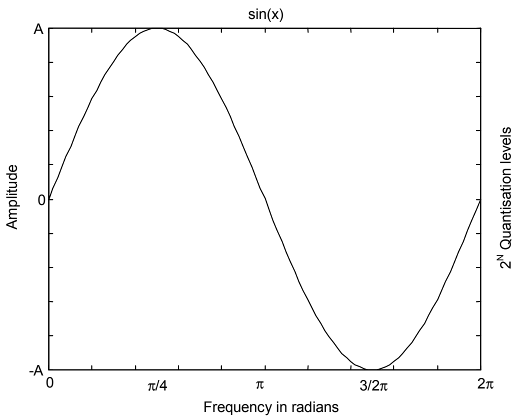

Figure 8.7: Sine wave with amplitude A, quantised with 2^N quantisation levels.

It can be shown from the diagram that the magnitude of the sine wave is equal to the number of quantisation levels multiplied by the step size, represented by Equation 8.6.

$$2A = 2^N \cdot q \tag{8.6}$$

Re-arranging Equation 8.6 to make *A* the subject gives the following:

$$A = \frac{2^N \cdot q}{2}$$

The *signal power* of a sine wave is mathematically defined by Equation 8.7.

$$signal \quad power = \frac{A^2}{2} \tag{8.7}$$

Now substituting for *A* in Equation 8.7, the *signal power* of a sine wave can be represented in terms of the quantisation step size and the number of quantisation bits. This is defined by Equation 8.8 below.

$$signal \quad power = \frac{1}{2}\left[\frac{2^N \cdot q}{2}\right]^2 = 2^{2N-3} \cdot q^2 \tag{8.8}$$

The SNR can now be deduced by inserting Equation 8.5 and 8.8 into Equation 8.9.

$$SNR = \frac{signal \quad power}{noise \quad power} \tag{8.9}$$

$$SNR = \frac{2^{2N-3} \cdot q^2}{q^2 / 12} = 2^{2N} \cdot \frac{3}{2}$$

It is common practise to express the SNR of a signal in terms of decibels (dB). This is achieved by taking the $10log_{10}$ of the SNR. Hence, the SNR of a sine wave quantised with N bits is given by Equation 8.10 below.

$$SNR = 10\log_{10}\left[2^{2N} \cdot \frac{3}{2}\right] = [6.02N + 1.76]dB \tag{8.10}$$

8.8. Dynamic Range

It is often the case that a guaranteed SNR is required for a signal which may be any amplitude within a given dynamic range. In this case, it is necessary to determine how many ADC bits are required to achieve the SNR for the smallest signal, and then to have additional bits to ensure that the ADC does not saturate for the largest signal. For example, if the amplitude of a sine wave is known to lie between 1V and 8V, then the number of additional bits to accommodate this dynamic range would be 3 (2^3=8).

Example: Calculate the number of bits required to produce a SNR of 30dB for a sine wave acquisition that varies between 0.5V to 6V.

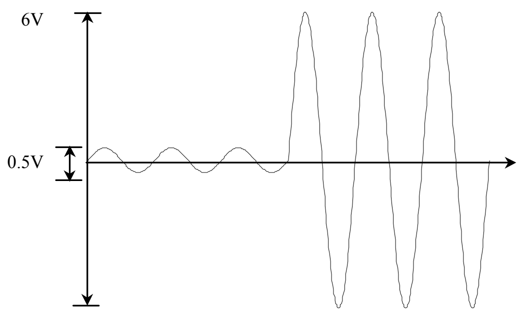

Figure 8.8: A sine wave acquisition varying between 0.5V and 6V.

The number of bits that are required to maintain a SNR of 30 dB for the 0.5V part is:

$$SNR = [6.02N + 1.76]dB$$

$30 = 6.02N_1 + 1.76$

$N_1 = 4.69$ bits

An additional N_2 bits are also required to maintain a SNR of 30dB for the 6V part of the waveform. This is calculated by re-arranging Equation 8.6 to making N_2 the subject and taking the logarithm.

$$2^{N_2} = \frac{6}{0.5}$$

$$N_2 = \log_2(12) = \frac{\log_{10}(12)}{\log_{10}(2)}$$

$N_2 = 3.58$ bits

Hence, the total number of bits required to produce a SNR of 30dB for a sine wave acquisition that varies between 0.5V and 6V is given by:

Total bits = $N_1 + N_2$

Total bits = 4.69 + 3.58 = 8.27 = 9 (rounded up)

It is important to realise that the bits must be rounded up to become a real integer as we are using binary. Since the number of quantisation bits is 9, the quantisation step size q can now be calculated.

$$q = \frac{6}{2^9} = 0.0117\text{V}$$

8.9. Signal Reconstruction

If a sampled signal $x[n]$ has been obtained from a band-limited signal $x(t)$ by sampling at the Nyquist rate (or higher), the signal $x(t)$ can be perfectly reconstructed using the formula:

$$x(t) = \sum_{n=-\infty}^{\infty} x(nT)\,\mathrm{sinc}\left(\frac{t-nT}{T}\right) \qquad (8.11)$$

However, this represents a filter that has an infinite impulse response, which is therefore non-causal.

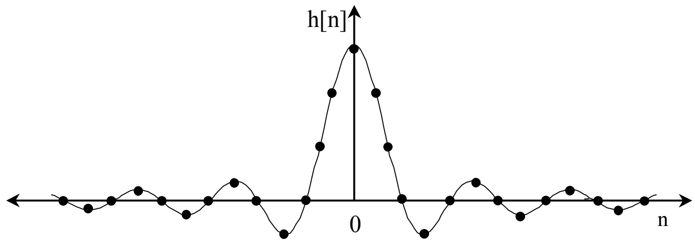

Figure 8.9: Impulse response of perfect reconstruction filter

In practice, therefore, it is usual to use the "zero-order hold" reconstruction filter which has the formula:

$$h_{zoh}(t) = \begin{cases} 1, & 0 \le t < T \\ 0, & otherwise \end{cases} \qquad (8.12)$$

which has an impulse response and overall response as shown: