Chapter 11

Sampling and Reconstruction

Digital hardware, including computers, take actions in discrete steps. So they can deal with discrete-time signals, but they cannot directly handle the continuous-time signals that are prevalent in the physical world. This chapter is about the interface between these two worlds, one continuous, the other discrete. A discrete-time signal is constructed by **sampling** a continuous-time signal, and a continuous-time signal is **reconstructed** by **interpolating** a discrete-time signal.

11.1   Sampling

A sampler for complex-valued signals is a system

$$Sampler_T : [Reals \rightarrow Complex] \rightarrow [Integers \rightarrow Complex], \qquad (11.1)$$

where T is the **sampling interval** (it has units of seconds/sample). The system is depicted in figure 11.1. The **sampling frequency** or **sample rate** is $f_s = 1/T$, in units of samples/second (or sometimes, Hertz), or $\omega_s = 2\pi/T$, in units radians/second. If $y = Sampler_T(x)$ then y is defined by

$$\forall\, n \in Integers, \quad y(n) = x(nT). \qquad (11.2)$$

11.1.1   Sampling a sinusoid

Let $x : Reals \rightarrow Reals$ be the sinusoidal signal

$$\forall\, t \in Reals, \quad x(t) = \cos(2\pi f t), \qquad (11.3)$$

x: Reals → Complex y: Integers → Complex
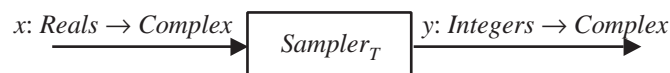
$Sampler_T$

Figure 11.1: Sampler.

> **Basics: Units**
>
> Recall that frequency can be given with any of various units. The units of the f in (11.3) and (11.4) are Hertz, or cycles/second. In (11.3), it is sensible to give the frequency as $\omega = 2\pi f$, which has units of radians/second. The constant 2π has units of radians/cycle, so the units work out. Moreover, the time argument t has units of seconds, so the argument to the cosine function, $2\pi ft$, has units of radians, as expected.
>
> In the discrete time case (11.4), it is sensible to give the frequency as $2\pi fT$, which has units of radians/sample. The sampling interval T has units of seconds/sample, so again the units work out. Moreover, the integer n has units of samples, so again the argument to the cosine function, $2\pi fnT$, has units of radians, as expected.
>
> In general, when discussing continuous-time signals and their sampled discrete-time signals, it is important to be careful and consistent in the units used, or considerable confusion can result. Many texts talk about **normalized frequency** when discussing discrete-time signals, by which they simply mean frequency in units of radians/sample. This is **normalized** in the sense that it does not depend on the sampling interval.

where f is the frequency of the sinewave in Hertz. Let $y = Sampler_T(x)$. Then

$$\forall\, n \in Integers, \quad y(n) = \cos(2\pi fnT). \tag{11.4}$$

Although this looks similar to the continuous-time sinusoid, there is a fundamental difference. Because the index n is discrete, it turns out that the frequency f is indistinguishable from frequency $f + f_s$ when looking at the discrete-time signal. This phenomenon is called **aliasing**.

11.1.2   Aliasing

Consider another sinusoidal signal u given by

$$\forall\, t \in Reals, \quad u(t) = \cos(2\pi(f + Nf_s)t),$$

where N is some integer and $f_s = 1/T$. If $N \neq 0$, then this signal is clearly different from x in (11.3). Let

$$w = Sampler_T(u).$$

Then for all $n \in Integers$,

$$w(n) = \cos(2\pi(f + Nf_s)nT) = \cos(2\pi fnT + 2\pi Nn) = \cos(2\pi fnT) = y(n),$$

because Nn is an integer. Thus, even though $u \neq x$, $Sampler_T(u) = Sampler_T(x)$. Thus, after being sampled, the signals x and u are indistinguishable. This phenomenon is called **aliasing**, presumably because it implies that any discrete-time sinusoidal signal has many continuous-time identities (its "identity" is presumably its frequency).

Example 11.1: A typical sample rate for voice signals is $f_s = 8000$ samples/second, so the sampling interval is $T = 0.125$ msec/sample. A continuous-time sinusoid with frequency 440 Hz, when sampled at this rate, is indistinguishable from a continuous-time sinusoid with frequency 8,440 Hz, when sampled at this same rate.

Example 11.2: Compact discs are created by sampling audio signals at $f_s = 44,100$ Hz, so the sampling interval is about $T = 22.7$ μsec/sample. A continuous-time sinusoid with frequency 440 Hz, when sampled at this rate, is indistinguishable from a continuous-time sinusoid with frequency 44,540 Hz, when sampled at this same rate.

The frequency domain analysis of the previous chapters relied heavily on complex exponential signals. Recall that a cosine can be given as a sum of two complex exponentials, using Euler's relation,

$$\cos(2\pi f t) = 0.5(e^{i2\pi ft} + e^{-i2\pi ft}).$$

One of the complex exponentials is at frequency f, an the other is at frequency $-f$. Complex exponential exhibit the same aliasing behavior that we have illustrated for sinusoids.

Let $x:$*Reals* $\to$ *Complex* be

$$\forall\, t \in \textit{Reals}, \quad x(t) = e^{i2\pi ft}$$

where f is the frequency in Hertz. Let $y = $*Sampler*$_T(x)$. Then for all n in *Integers*,

$$y(n) = e^{i2\pi fnT}$$

Consider another complex exponential signal u,

$$u(t) = e^{i2\pi(f+Nf_s)t}$$

where N is some integer. Let

$$w = \textit{Sampler}_T(u).$$

Then for all $n \in$ *Integers*,

$$w(n) = e^{i2\pi(f+Nf_s)nT} = e^{i2\pi fnT}e^{i2\pi Nf_snT} = e^{i2\pi fnT} = y(n),$$

because $e^{i2\pi Nf_snT} = 1$. Thus, as with sinusoids, when we sample a complex exponential signal with frequency f at sample rate f_s, it is indistinguishable from one at frequency $f + f_s$ (or $f + Nf_s$ for any integer N).

There is considerably more to this story. Mathematically, aliasing relates to the periodicity of the frequency domain representation (the DTFT) of a discrete-time signal. We will also see that the effects of aliasing on real-valued signals (like the cosine, but unlike the complex exponential) depend strongly on the conjugate symmetry of the DTFT as well.
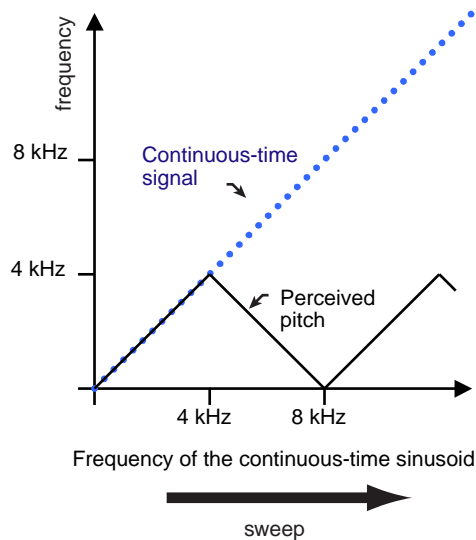
Figure 11.2: As the frequency of a continuous signal increases beyond the Nyquist frequency, the perceived pitch starts to drop.

11.1.3   Perceived pitch experiment

Consider the following experiment.[1] Generate a discrete-time audio signal with an 8,000 samples/second sample rate according to the formula (11.4). Let the frequency f begin at 0 Hz and sweep upwards through 4 kHz to (at least) 8 kHz. Use the audio output of a computer to listen to the resulting sound. The result is illustrated in figure 11.2. As the frequency of the continuous-time sinusoid rises, so does the perceived pitch, until the frequency reaches 4 kHz. At that point, the perceived pitch begins to fall rather than rise, even as the frequency of the continuous-time sinusoid continues to rise. It will fall until the frequency reaches 8 kHz, at which point no sound is heard at all (the perceived pitch is 0 Hz). Then the perceived pitch begins to rise again.

That the perceived pitch rises from 0 after the frequency f rises above 8000 Hz is not surprising. We have already determined that in a discrete-time signal, a frequency of f is indistinguishable from a frequency $f + 8000$, assuming the sample rate is 8,000 samples/second. But why does the perceived pitch drop when f rises above 4 kHz?

The frequency 4 kHz, $f_s/2$, is called the **Nyquist frequency**, after Harry Nyquist, an engineer at Bell Labs who, in the 1920s and 1930s, laid much of the groundwork for digital transmission of information. The Nyquist frequency turns out to be a key threshold in the relationship between discrete-time and continuous-time signals, more important even than the sampling frequency. Intuitively, this is because if we sample a sinusoid with a frequency below the Nyquist frequency (below half the sampling frequency), then we take at least two samples per cycle of the sinusoid. It should be intuitively appealing that taking at least two samples per cycle of a sinusoid has some key

---

[1]This experiment can be performed at http://www.eecs.berkeley.edu/ẽal/eecs20/week13/aliasing.html. Similar experiments are carried out in lab C.11.
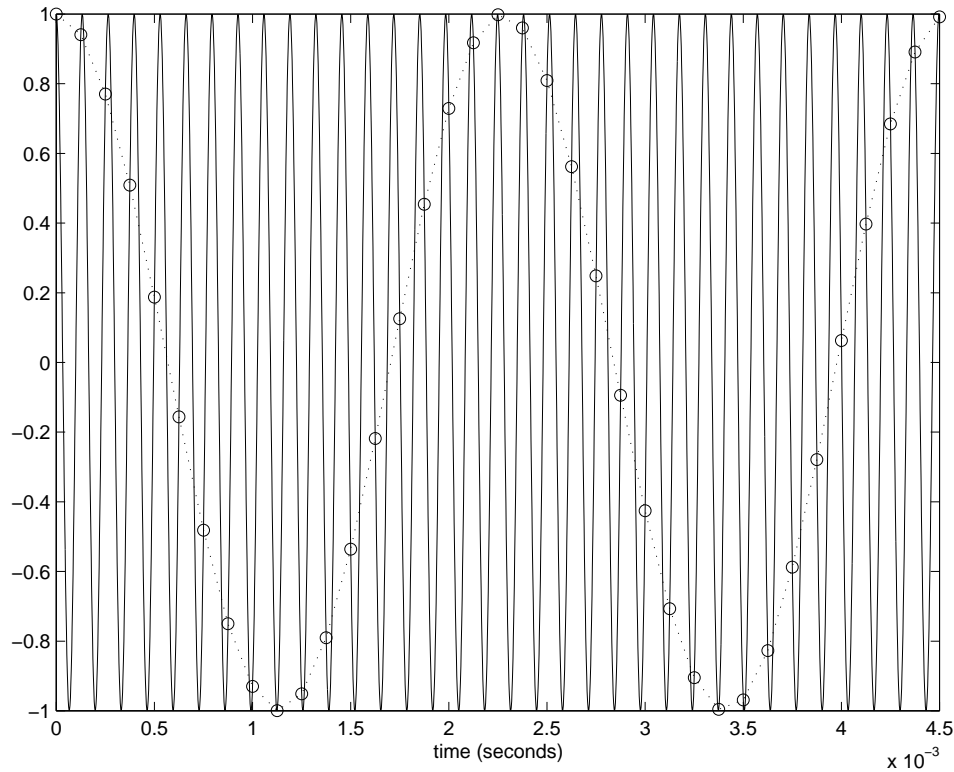
Figure 11.3: A sinusoid at 7.56 kHz and samples taken at 8 kHz.

significance. The two sample minimum allows the samples to capture the oscillatory nature of the sinusoid. Fewer than two samples would not do this. However, what happens when fewer than two samples are taken per cycle is not necessarily intuitive. It turns out that the sinusoid masquerades as one of another frequency.

Consider the situation when the frequency f of a continuous-time sinusoid is 7,560 Hz. Figure 11.3 shows 4.5 msec of the continuous-time waveform, together with samples taken at 8 kHz. Notice that the samples trace out another sinusoid. We can determine the frequency of that sinusoid with the help of figure 11.2, which suggests that the perceived pitch will be $8000 - 7560 = 440$ Hz (the slope of the perceived pitch line is -1 in this region). Indeed, if we listen to the sampled sinusoid, it will be an A-440.

Recall that a cosine can be given as a sum of complex exponentials with frequencies that are negatives of one another. Recall further that a complex exponential with frequency f is indistinguishable from one with frequency $f + Nf_s$, for any integer N. A variant of figure 11.2 that leverages this representation is given in figure 11.4.

In figure 11.4, as we sweep the frequency of the continuous-time signal from 0 to 8 kHz, we move from left to right in the figure. The sinusoid consists not only of the rising frequency shown by the dotted line in figure 11.2, but also of a corresponding falling (negative) frequency as shown in figure
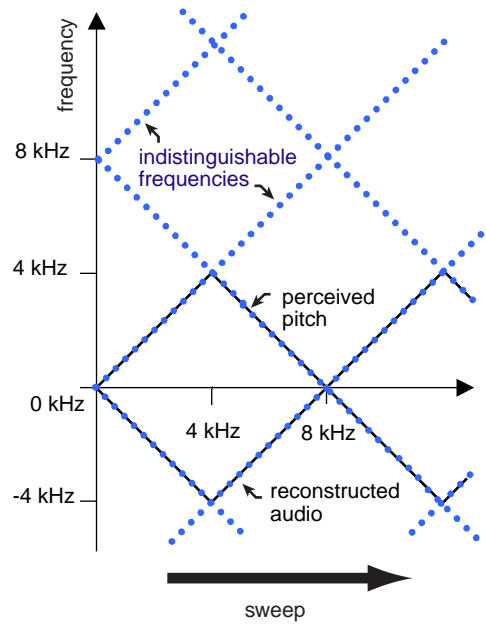
Figure 11.4: As the frequency of a continuous signal increases beyond the Nyquist frequency, the perceived pitch starts to drop because the frequency of the reconstructed continuous-time audio signal stays in the range $-f_s/2$ to $f_s/2$.

11.4. Moreover, these two frequencies are indistinguishable, after sampling, from frequencies that are 8 kHz higher or lower, also shown by dotted lines in figure 11.4.

When the discrete-time signal is converted to a continuous-time audio signal, the hardware performing this conversion can choose any matching pair of positive and negative frequencies. By far the most common choice is to select the matching pair with lowest frequency, shown in figure 11.4 by the solid lines behind dotted lines. These result in a sinusoid with frequency between 0 and the Nyquist frequency, $f_s/2$. This is why the perceived pitch falls after sweeping past $f_s/2 = 4$ kHz.

Recall that the frequency-domain representation (i.e. the DTFT) of a discrete-time signal is periodic with period 2π radians/sample. That is, if X is a DTFT, then

$$\forall\, \omega \in Reals, \quad X(\omega) = X(\omega + 2\pi).$$

In radians per second, it is periodic with period $2\pi f_s$. In Hertz, it is periodic with period f_s, the sampling frequency. Thus, in figure 11.4, the dotted lines represent this periodicity. This periodicity is another way of stating that frequencies separated by f_s are indistinguishable.

11.1.4 Avoiding aliasing ambiguities

Figure 11.4 suggests that even though a discrete-time signal has ambiguous frequency content, it is possible to construct a uniquely defined continuous-time signal from the discrete-time waveform by choosing the one unique frequency for each component that is closest to zero. This will always result in a reconstructed signal that contains only frequencies between zero and the Nyquist frequency.

Correspondingly, this suggests that when sampling a continuous-time signal, if that signal contains only frequencies below the Nyquist frequency, then this reconstruction strategy will perfectly recover the signal. This is an intuitive statement of the **Nyquist-Shannon sampling theorem**.

If a continuous-time signal contains only frequencies below the Nyquist frequency $f_s/2$, then it can be perfectly reconstructed from samples taken at sampling frequency f_s. This suggests that prior to sampling, it is reasonable to filter a signal to remove components with frequencies above $f_s/2$. A filter that realizes this is called an **anti-aliasing filter**.

> **Example 11.3:** In the telephone network, speech is sampled at 8000 samples per second before being digitized. Prior to this sampling, the speech signal is lowpass filtered to remove frequency components above 4000 Hz. This lowpass filtered speech can then be perfectly reconstructed at the far end of the telephone connection, which receives a stream of samples at 8000 sample per second.

Before probing this further, let us examine in more detail what we mean by reconstruction.

11.2 Reconstruction

Consider a system that constructs a continuous-time signal x from a discrete-time signal y,

$$DiscToCont_T : DiscSignals \rightarrow ContSignals.$$

Probing further: Anti-Aliasing for Fonts

When rendering characters on a computer screen, it is common to use **anti-aliasing** to make the characters look better. Consider the two figures below:

At the left is an image of the Greek letter omega. At the right is the result of sampling that rendition by taking only one pixel out of every 100 pixels in the original (every 10-th pixel horizontally and vertically), and then rescaling the image so it has the same size as the one on the left. The original image is discrete, and the resulting image is a smaller discrete image (this process is known as **subsampling**). Rendered with normal-sized pixels the character on the right looks like this:

To the discerning eye, this can be improved considerably. The problem is that the character at the upper left above has hard edges, and hence high (spatial) frequencies. Those high frequencies result in aliasing distortion when subsampling. To improve the result, we first lowpass filter the character (blurring it), and then subsample, as shown below:

The result looks better to the discerning eye:

$$y\text{: } Integers \to Complex \quad \boxed{DiscToCont_T} \quad x\text{: } Reals \to Complex$$

Figure 11.5: Discrete to continuous converter.

This is illustrated in figure 11.5. Systems that carry out such 'discrete-to-continuous' conversion can be realized in any number of ways. Some common examples are illustrated in figure 11.6, and defined below:

- **zero-order hold**: This means simply that the value of the each sample $y(n)$ is held constant for duration T, so that $x(t) = y(n)$ for the time interval from $t = nT$ to $t = (n+1)T$, as illustrated in figure 11.6(b). Let this system be denoted

$$ZeroOrderHold_T\text{: } DiscSignals \to ContSignals.$$

- **linear interpolation**: Intuitively, this means simply that we connect the dots with straight lines. Specifically, in the time interval from $t = nT$ to $t = (n+1)T$, $x(t)$ has values that vary along a straight line from $y(n)$ to $y(n+1)$, as illustrated in figure 11.6(c). Linear interpolation is sometimes called **first-order hold**. Let this system be denoted

$$LinearInterpolator_T\text{: } DiscSignals \to ContSignals.$$

- **ideal interpolation**: It is not yet clear what this should mean, but intuitively, it should result in a smooth curve that passes through the samples, as illustrated in figure 11.6(d). We will give a precise meaning below. Let this system be denoted

$$IdealInterpolator_T\text{: } DiscSignals \to ContSignals.$$

11.2.1 A model for reconstruction

A convenient mathematical model for reconstruction divides the reconstruction process into a cascade of two systems, as shown in figure 11.7. Thus

$$x = S(ImpulseGen_T(y)),$$

where S is an LTI system to be determined. The first of these two subsystems,

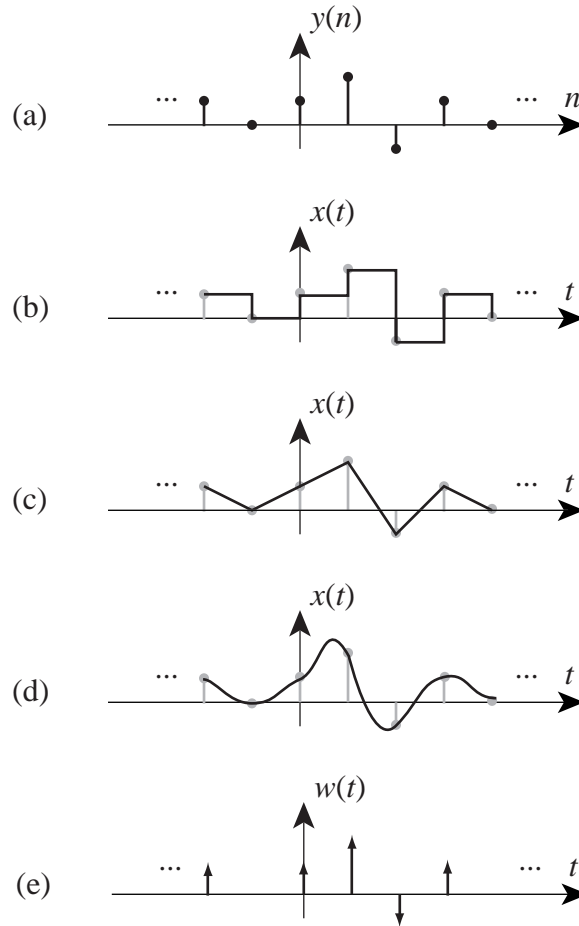$$ImpulseGen_T\text{: } DiscSignals \to ContSignals,$$

Figure 11.6: A discrete-time signal (a), a continuous-time reconstruction using zero-order hold (b), a reconstruction using linear interpolation (c), a reconstruction using ideal interpolation (d), and a reconstruction using weighted Dirac delta functions (e).
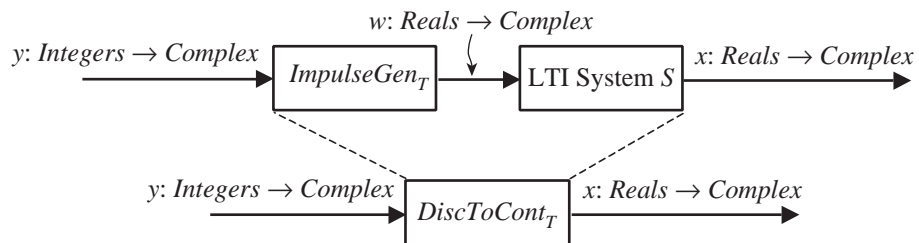

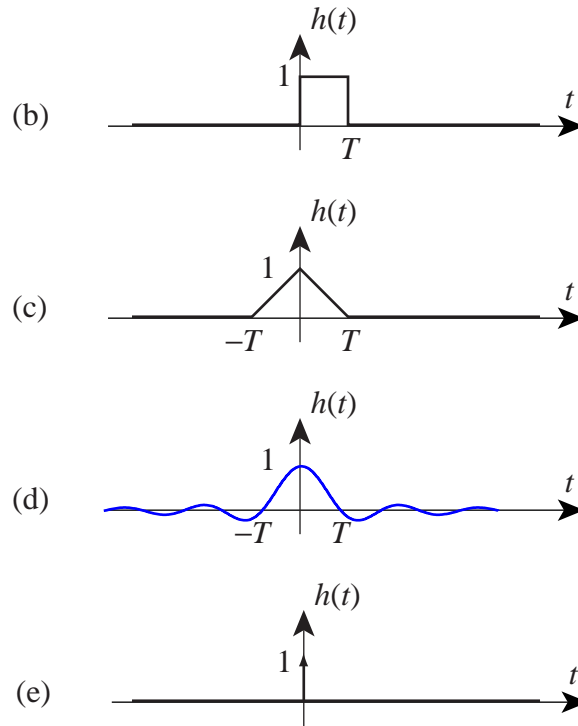Figure 11.7: A model for reconstruction divides it into two stages.

Figure 11.8: The impulse responses for the LTI system S in figure 11.7 that yield the interpolation methods in figure 11.6(b-e).

constructs a continuous-time signal, where for all $t \in Reals$,

$$w(t) = \sum_{k=-\infty}^{\infty} y(k)\delta(t - kT).$$

This is a continuous-time signal that at each sampling instant kT produces a Dirac delta function with weight equal to the sample value, $y(k)$. This signal is illustrated in figure 11.6(e). It is a mathematical abstraction, since everyday engineering systems do not exhibit the singularity-like behavior of the Dirac delta function. Nonetheless, it is a useful mathematical abstraction.

The second system in figure 11.7, S, is a continuous-time LTI filter with an impulse response that determines the interpolation method. The impulse responses that yield the interpolation methods in figure 11.6(b-e) are shown in figure 11.8(b-e). If

$$\forall t \in Reals, \quad h(t) = \begin{cases} 1 & 0 \leq t < T \\ 0 & otherwise \end{cases}$$

then the interpolation method is zero-order hold. If

$$\forall t \in Reals, \quad h(t) = \begin{cases} 1 + t/T & -T < t < 0 \\ 1 - t/T & 0 \leq t < T \\ 0 & otherwise \end{cases}$$

then the interpolation method is linear. If the impulse response is

$$\forall\, t \in Reals, \quad h(t) = \frac{\sin(\pi t / T)}{\pi t / T}$$

then the interpolation method is ideal. The above impulse response is called a **sinc function**, and its Fourier transform, from table 10.4, is given by

$$\forall\, \omega \in Reals, \quad X(\omega) = \begin{cases} T & \text{if } |\omega| \leq \pi/T \\ 0 & \text{otherwise} \end{cases}$$

Notice that the Fourier transform is zero at all frequencies above π/T radians/second, or $f_s/2$ Hz, the Nyquist frequency. It is this characteristic that makes it ideal. It precisely performs the strategy illustrated in figure 11.4, where among all indistinguishable frequencies we select the ones between $-f_s/2$ and $f_s/2$.

If we let $Sinc_T$ denote the LTI system S when the impulse response is a sinc function, then

$$IdealInterpolator_T = Sinc_T \circ ImpulseGen_T.$$

In practice, ideal interpolation is difficult to accomplish. From the expression for the sinc function we can understand why. First, this impulse response is not causal. Second, it is infinite in extent. More importantly, its magnitude decreases rather slowly as t increases or decreases (proportional to $1/t$ only). Thus, truncating it at finite length leads to substantial errors.

If the impulse response of S is

$$h(t) = \delta(t),$$

where δ is the Dirac delta function, then the system S is a pass-through system, and the reconstruction consists of weighted delta functions.

11.3   The Nyquist-Shannon sampling theorem

We can now give a precise statement of the **Nyquist-Shannon sampling theorem**:

If x is a continuous-time signal with Fourier transform X and if $X(\omega)$ is zero outside the range $-\pi/T < \omega < \pi/T$ radians/second, then

$$x = IdealInterpolator_T(Sampler_T(x)).$$

We can state this theorem slightly differently. Suppose x is a continuous-time signal with no frequency larger than some f_0 Hertz. Then x can be recovered from its samples if $f_0 < f_s/2$, the Nyquist frequency.

Probing further: Sampling

We can construct a mathematical model for sampling by using Dirac delta functions. Define a pulse stream by

$$\forall\, t \in Reals, \quad p(t) = \sum_{k=-\infty}^{\infty} \delta(t - kT).$$

Consider a continuous-time signal x that we wish to sample with sampling period T. That is, we define $y(n) = x(nT)$. Construct first an intermediate continuous-time signal $w(t) = x(t)p(t)$. We can show that the CTFT of w is equal to the DTFT of y. This gives us a way to relate the CTFT of x to the DTFT of its samples y. Recall that multiplication in the time domain results in convolution in the frequency domain (see table 10.9), so

$$W(\omega) = \frac{1}{2\pi}X(\omega) * P(\omega) = \frac{1}{2\pi}\int_{-\infty}^{\infty} X(\Omega)P(\omega - \Omega)d\Omega.$$

It can be shown (see box on page 386 that the CTFT of $p(t)$ is

$$P(\omega) = \frac{2\pi}{T} \sum_{k=-\infty}^{\infty} \delta(\omega - k\frac{2\pi}{T}),$$

so

$$
\begin{aligned}
W(\omega) &= \frac{1}{2\pi}\int_{-\infty}^{\infty} X(\Omega)\frac{2\pi}{T} \sum_{k=-\infty}^{\infty} \delta(\omega - \Omega - k\frac{2\pi}{T})d\Omega \\
&= \frac{1}{T} \sum_{k=-\infty}^{\infty} \int_{-\infty}^{\infty} X(\Omega)\delta(\omega - \Omega - k\frac{2\pi}{T})d\Omega \\
&= \frac{1}{T} \sum_{k=-\infty}^{\infty} X(\omega - k\frac{2\pi}{T})
\end{aligned}
$$

where the last equality follows from the sifting property (9.11). The next step is to show that

$$Y(\omega) = W(\omega/T).$$

We leave this as an exercise. From this, the basic Nyquist-Shannon result follows,

$$Y(\omega) = \frac{1}{T} \sum_{k=-\infty}^{\infty} X\left(\frac{\omega - 2\pi k}{T}\right).$$

This relates the CTFT X of the signal being sampled x to the DTFT Y of the discrete-time result y.

Probing further: Impulse Trains

Consider a signal p consisting of periodically repeated Dirac delta functions with period T,

$$\forall\, t \in \textit{Reals}, \quad p(t) = \sum_{k=-\infty}^{\infty} \delta(t - kT).$$

This signal has the Fourier series expansion

$$\forall\, t \in \textit{Reals}, \quad p(t) = \sum_{m=-\infty}^{\infty} \frac{1}{T} e^{i\omega_0 m t},$$

where the fundamental frequency is $\omega_0 = 2\pi/T$. This can be verified by applying the formula from table 10.5. That formula, however, gives an integration range of 0 to the period, which in this case is T. This integral covers one period of the periodic signal, but starts and ends on a delta function in p. To avoid the resultant mathematical subtleties, we can integrate from $-T/2$ to $T/2$, getting Fourier series coefficients

$$\forall\, m \in \textit{Integers}, \quad P_m = \frac{1}{T} \int_{-T/2}^{T/2} \left[\sum_{k=-\infty}^{\infty} \delta(t - kT) \right] e^{i\omega_0 m t}\, dt.$$

The integral is now over a range that includes only one of the delta functions. The kernel of the integral is zero except when $t = 0$, so by the sifting rule, the integral evaluates to 1. Thus, all Fourier series coefficients are $P_m = 1/T$. Using the relationship between the Fourier series and the Fourier Transform of a periodic signal (from section 10.6.3), we can write the continuous-time Fourier transform of p as

$$\forall\, \omega \in \textit{Reals}, \quad P(\omega) = \frac{2\pi}{T} \sum_{k=-\infty}^{\infty} \delta\left(\omega - \frac{2\pi}{T} k \right).$$
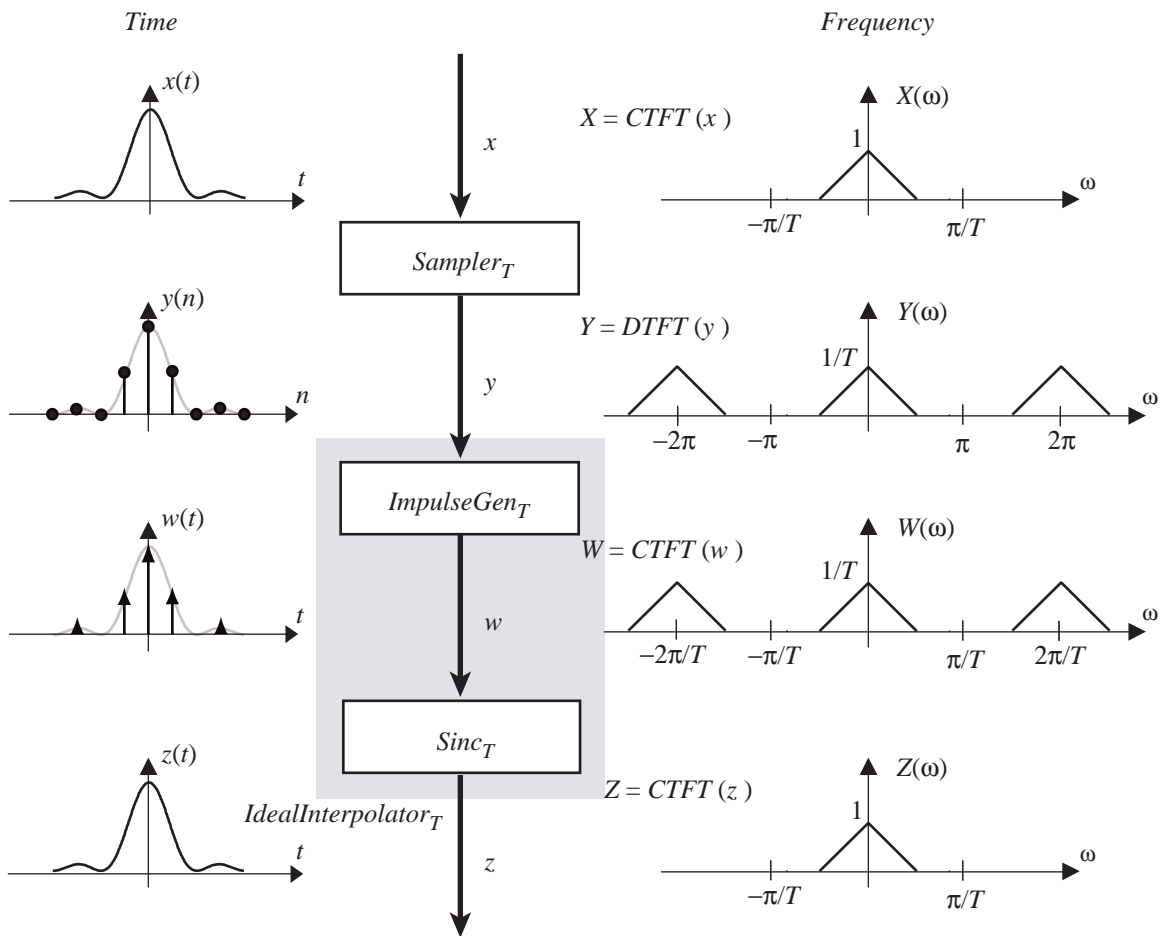
Figure 11.9: Steps in the justification of the Nyquist-Shannon sampling theorem.
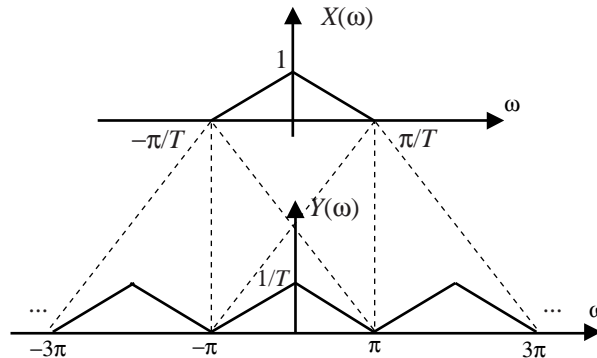
Figure 11.10: Relationship between the CTFT of a continuous-time signal and the DTFT of its discrete-time samples. The DTFT is the sum of the CTFT and its copies shifted by multiples of $2\pi/T$, the sampling frequency in radians per second. The frequency axis is also normalized.

A formal proof of this theorem involves some technical difficulties (it was first given by Claude Shannon of Bell Labs in the late 1940s). But we can get the idea from the following three-step argument (see figure 11.9).

Step 1. Let x be a continuous-time signal with Fourier transform X. At this point we do not require that $X(\omega)$ be zero outside the range $-\pi/T < \omega < \pi/T$. We sample x with sampling interval T to get the discrete-time signal

$$y = Sampler_T(x).$$

It can be shown (see box on page 385) that the DTFT of y is related to the CTFT of x by

$$\boxed{Y(\omega) = \frac{1}{T} \sum_{k=-\infty}^{\infty} X\left(\frac{\omega}{T} - \frac{2\pi k}{T}\right).}$$

This important relation says that the DTFT Y of y is the sum of the CTFT X with copies of it shifted by multiples of $2\pi/T$. Also, the frequency axis is normalized by dividing ω by T. There are two cases to consider, depending on whether the shifted copies overlap.

First, if $X(\omega) = 0$ outside the range $-\pi/T < \omega < \pi/T$, then the copies will not overlap, and in the range $-\pi < \omega < \pi$,

$$Y(\omega) = \frac{1}{T} X\left(\frac{\omega}{T}\right). \tag{11.5}$$

In this range of frequencies, Y has the same shape as X, scaled by $1/T$. This relationship between X and Y is illustrated in figure 11.10, where X is drawn with a triangular shape.

In the second case, illustrated in figure 11.11, X does have non-zero frequency components higher than π/T. Notice that in the sampled signal, the frequencies in the vicinity of π are distorted by the overlapping of frequency components above and below π/T in the original signal. This distortion is called **aliasing distortion**.
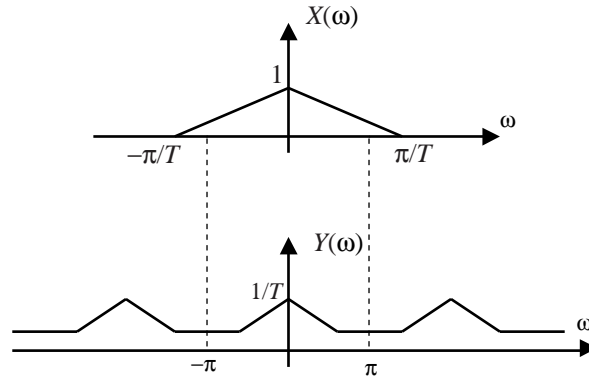
Figure 11.11: Relationship between the CTFT of a continuous-time signal and the DTFT of its discrete-time samples when the continuous-time signal has a broad enough bandwidth to introduce aliasing distortion.

We continue with the remaining steps, following the signals in figure 11.9.

Step 2. Let w be the signal produced by the impulse generator,

$$\forall\, t \in Reals, \quad w(t) = \sum_{n=-\infty}^{\infty} y(n)\delta(t-nT).$$

The Fourier Transform of w is $W(\omega) = Y(\omega T)$ (see box on page 385).

Step 3. Let z be the output of the *IdealInterpolator*$_T$. Its Fourier transform is simply

$$\begin{aligned} Z(\omega) &= W(\omega)S(\omega) \\ &= Y(\omega T)S(\omega), \end{aligned}$$

where $S(\omega)$ is the frequency response of the reconstruction filter *IdealInterpolator*$_T$. As seen in exercise 21 of chapter 10,

$$S(\omega) = \begin{cases} T & -\pi/T < \omega < \pi/T \\ 0 & \text{otherwise} \end{cases} \tag{11.6}$$

Substituting for S and Y, we get

$$\begin{aligned} Z(\omega) &= \begin{cases} TY(\omega T) & -\pi/T < \omega < \pi/T \\ 0 & \text{otherwise} \end{cases} \\ &= \begin{cases} \sum\limits_{k=-\infty}^{\infty} X(\omega - 2\pi k/T) & -\pi/T < \omega < \pi/T \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

If $X(\omega)$ is zero for $|\omega|$ larger than the Nyquist frequency π/T, then we conclude that

$$\forall\, \omega \in Reals, \quad Z(\omega) = X(\omega).$$

That is, w is identical to x. This proves the Nyquist-Shannon result.

However, if $X(\omega)$ does have non-zero values for some $|\omega|$ larger than the Nyquist frequency, then z will be different from x, as illustrated in figure 11.11.

11.4   Summary

The acts of sampling and reconstructing a continuous-time signal bridge the continuous-time world with the discrete computational world. The periodicity of frequencies in the discrete world implies that for each discrete-time sinusoidal signal, there are multiple corresponding discrete-time frequencies. These frequencies are aliases of one another. When a signal is sampled, these frequencies become indistinguishable, and aliasing distortion may result. The Nyquist-Shannon sampling theorem gives a simple condition under which aliasing distortion is avoided. Specifically, if the signal contains no sinusoidal components with frequencies higher than half the sampling frequency, then there will be no aliasing distortion. Half the sampling frequency is called the Nyquist frequency because of this key result.

Exercises

Each problem is annotated with the letter **E, T, C** which stands for exercise, requires some thought, requires some conceptualization. Problems labeled **E** are usually mechanical, those labeled **T** require a plan of attack, those labeled **C** usually have more than one defensible answer.

1. **E** Consider the continuous-time signal

$$x(t) = \cos(10\pi t) + \cos(20\pi t) + \cos(30\pi t).$$

   (a) Find the fundamental frequency. Give the units.

   (b) Find the Fourier series coefficients $A_0, A_1, \cdots$ and $\phi_1, \phi_2, \cdots$.

   (c) Let y be the result of sampling this signal with sampling frequency 10 Hz. Find the fundamental frequency for y, and give the units.

   (d) For the same y, find the discrete-time Fourier series coefficients, $A_0, A_1, \cdots$ and $\phi_1, \cdots$.

   (e) Find

$$w = \textit{IdealInterpolator}_T(\textit{Sampler}_T(x))$$

   for $T = 0.1$ seconds.

   (f) Is there any aliasing distortion caused by sampling at 10 Hz? If there is, describe the aliasing distortion in words.

   (g) Give the smallest sampling frequency that avoids aliasing distortion.

2. **E** Verify that $\textit{Sampler}_T$ defined by (11.1) and (11.2) is linear but not time invariant.

3. **E** A real-valued sinusoidal signal with a negative frequency is always exactly equal to another sinusoid with positive frequency. Consider a real-valued sinusoid with a negative frequency -440 Hz,

$$y(n) = \cos(-2\pi 440 nT + \phi).$$

Find a positive frequency f and phase θ such that

$$y(n) = \cos(2\pi f nT + \theta).$$

4. **T** Consider a continuous-time signal x where for all $t \in Reals$,

$$x(t) = \sum_{k=-\infty}^{\infty} r(t-k).$$

where

$$r(t) = \begin{cases} 1 & 0 \le t < 0.5 \\ 0 & otherwise \end{cases}.$$

   (a) Is $x(t)$ periodic? If so, what is the period?

   (b) Suppose that $T = 1$. Give a simple expression for $y = Sampler_T(x)$.

   (c) Suppose that $T = 0.5$. Give a simple expression for $y = Sampler_T(x)$ and $z = IdealInterpolator_T(Sampler_T(x))$

   (d) Find an upper bound for T (in seconds) such that $x = IdealInterpolator_T(Sampler_T(x))$, or argue that no value of T makes this assertion true.

5. **T** Consider a continuous-time signal x with the following finite Fourier series expansion,

$$\forall\, t \in Reals, \quad x(t) = \sum_{k=0}^{4} \cos(k\omega_0 t)$$

where $\omega_0 = \pi/4$ radians/second.

   (a) Give an upper bound on T (in seconds) such that $x = IdealInterpolator_T(Sampler_T(x))$.

   (b) Suppose that $T = 4$ seconds. Give a simple expression for $y = Sampler_T(x)$.

   (c) For the same $T = 4$ seconds, give a simple expression for

$$w = IdealInterpolator_T(Sampler_T(x)).$$

6. **T** Consider a continuous-time audio signal x with CTFT shown in figure 11.12. Note that it contains no frequencies beyond 10 kHz. Suppose it is sampled at 40 kHz to yield a signal that we will call x_{40}. Let X_{40} be the DTFT of x_{40}.

   (a) Sketch $|X_{40}(\omega)|$ and carefully mark the magnitudes and frequencies.

   (b) Suppose x is sampled at 20,000 samples/second. Let x_{20} be the resulting sampled signal and X_{20} its DTFT. Sketch and compare x_{20} and x_{40}.

   (c) Now suppose x is sampled at 15,000 samples/second. Let x_{15} be the resulting sampled signal and X_{15} its DTFT. Sketch and compare X_{20} and X_{15}. Make sure that your sketch shows aliasing distortion.
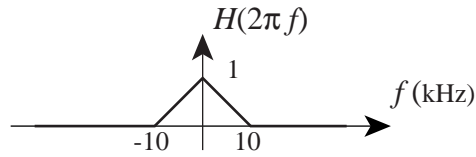
Figure 11.12: CTFT of an audio signal considered in exercise 6.

7. **C** Consider two continuous-time sinusoidal signals given by

$$x_1(t) = \cos(\omega_1 t)$$

$$x_2(t) = \cos(\omega_2 t),$$

with frequencies ω_1 and ω_2 radians/second such that

$$0 \leq \omega_1 \leq \pi/T \quad \text{and} \quad 0 \leq \omega_2 \leq \pi/T.$$

Show that if $\omega_1 \neq \omega_2$ then

$$Sampler_T(x_1) \neq Sampler_T(x_2).$$

I.e., the two distinct sinusoids cannot be aliases of one another if they both have frequencies below the Nyquist frequency. **Hint**: Try evaluating the sampled signals at $n = 1$.