Introduction to Embedded Systems
Notes for EECS 124
Spring 2008
Edward A. Lee, Sanjit A. Seshia, and Claire Tomlin
eal@eecs.berkeley.edu, sseshia@eecs.berkeley.edu, tomlin@eecs.berkeley.edu
Electrical Engineering & Computer Sciences
University of California, Berkeley
February 3, 2008

Chapter 2

Sensors

This chapter, yet to be written, gives an overview of sensor technology with emphasis on how to model sensors.

2.1 Signal Conditioning

Sensors convert physical measurements into data. Invariably, they are far from perfect, in that the data they yield gives information about the physical phenomenon that we wish to observe and other phenomena that we do not wish to observe. Removing or attenuating the effects of the phenomena we do not wish to observe is called **signal conditioning**.

Suppose that a sensor yields a continuous-time signal x. We model it as a sum of a **desired part** x_d and an **undesired part** x_n,

$$x(t) = x_d(t) + x_n(t). \tag{2.1}$$

The undesired part is called **noise**. To condition this signal, we would like to remove or reduce x_n without affecting x_d. In order to do this, of course, there has to be some meaningful difference between x_n and x_d. Often, the two parts differ considerably in their frequency content.

> **Example 2.1:** Consider using an accelerometer to measure the orientation of a slowly moving object. The accelerometer reacts to changes in orientation because they change the direction of the gravitational field with respect to its axis. But it will also report acceleration due to vibration. Let x_d be the signal due to orientation and x_n be the signal due to vibration. We will assume that x_n has higher frequency content than x_d. Thus, by frequency-selective filtering, we can reduce the effects of vibration.

Analysis of the degree to which frequency selective filtering helps requires having a model of both the desired signal x_d and the noise x_n. Reasonable models are usually statistical, and analysis of the signals requires using the techniques of random processes. Although such analysis is beyond the scope of this text, we can gain insight that is useful in many practical circumstances through a purely deterministic analysis.

Our approach will be to condition the signal $x = x_d + x_n$ by filtering it with an LTI system S called a **conditioning filter**. Let the output of the conditioning filter be given by

$$y = S(x) = S(x_d + x_n) = S(x_d) + S(x_n),$$

where we have used the linearity assumption on S. Let the error signal be defined to be

$$r = y - x_d.$$

This signal tells us how far off the filtered output is from the desired signal. The **energy** in the signal r is defined to be

$$||r||^2 = \int_{-\infty}^{\infty} r^2(t)dt.$$

We define the **signal to noise ratio** (**SNR**) to be

$$SNR = \frac{||x_d||^2}{||r||^2}.$$

Combining the above definitions, we can write this as

$$SNR = \frac{||x_d||^2}{||S(x_d) - x_d + S(x_n)||^2}. \tag{2.2}$$

It is customary to give SNR in **decibels**, written **dB**, defined as follows,

$$SNR_{dB} = 10\log_{10}(SNR).$$

Note that for typical signals in the real world, the energy is effectively infinite if the signal goes on forever. A statistical model, therefore, would use the **power**, defined as the expected energy per unit time. But since we are avoiding using statistical methods here, we will stick to energy as the criterion.

A reasonable design objective for a conditioning filter is to maximize the SNR. Of course, it will not be adequate to use a filter that maximizes the SNR only for particular signals x_d and x_n. We cannot know when we design the filter what these signals are, so the SNR needs to be maximized in expectation. That is, over the ensemble of signals we might see when operating the system, weighted by their likelihood, the expected SNR should be maximized.

Although determination of this filter requires statistical methods beyond the scope of this text, we can draw some intuitively appealing conclusions by examining (2.2). The numerator is not affected by S, so we can ignore it and minimize the denominator. It is easy to show that the denominator is bounded as follows,

$$||r||^2 \leq ||S(x_d) - x_d||^2 + ||S(x_n)||^2 \tag{2.3}$$

which suggests that we may be able to minimize the denominator by making $S(x_d)$ close to x_d (i.e. make $||S(x_d) - x_d||^2$ small) while making $||S(x_n)||^2$ small. That is, the filter S should do minimal damage to the desired signal x_d while filtering out as much as much as possible of the noise. This, of course, is obvious.

As illustrated in Example 2.1, x_d and x_n often differ in frequency content. We can get further insight using **Parseval's theorem**, which relates the energy to the Fourier transform,

$$||r||^2 = \int_{-\infty}^{\infty} (r(t))^2 dt = \frac{1}{2\pi} \int_{-\infty}^{\infty} |R(\omega)|^2 d\omega = \frac{1}{2\pi} ||R||^2$$

where R is the Fourier transform of r.

The filter S is an LTI system. It is defined equally well by the function $S\colon (\mathbb{R} \to \mathbb{R}) \to (\mathbb{R} \to \mathbb{R})$, by its impulse response $h\colon \mathbb{R} \to \mathbb{R}$, a continuous-time signal, or by its **transfer function** $H\colon \mathbb{R} \to \mathbb{C}$, the Fourier transform of the impulse response. Using the transfer function and Parseval's theorem, we can write

$$SNR = \frac{||X_d||^2}{||HX_d - X_d + HX_n||^2}, \tag{2.4}$$

where X_d is the Fourier transform of x_d and X_n is the Fourier transform of x_n. In Problem 1, we explore a very simple strategy that chooses the transfer function so that $H(\omega) = 1$ in the frequency range where x_d is present, and $H(\omega) = 0$ otherwise. This strategy is not exactly realizable in practice, but an approximation of it will work well for the problem described in Example 2.1.

Note that it is easy to adapt the above analysis to discrete-time signals. If $r\colon \mathbb{Z} \to \mathbb{R}$ is a discrete-time signal, its energy is

$$||r||^2 = \sum_{n=-\infty}^{\infty} (r(n))^2.$$

If its discrete-time Fourier transform is R, then Parseval's relation becomes

$$||r||^2 = \sum_{n=-\infty}^{\infty} (r(n))^2 = \frac{1}{2\pi} \int_{-\pi}^{\pi} |R(\omega)|^2 d\omega.$$

Note that the limits on the integral are different, covering one cycle of the periodic DTFT. All other observations above carry over unchanged.

2.2 Sampling

Almost every embedded system will sample and digitize sensor data. In this section, we review the phenomenon of aliasing. We use a mathematical model for sampling by using the Dirac delta function δ. Define a pulse stream by

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

Consider a continuous-time signal x that we wish to sample with sampling period T. That is, we define a discrete-time signal $y\colon \mathbb{Z} \to \mathbb{R}$ by $y(n) = x(nT)$. Construct first an intermediate continuous-time signal $w(t) = x(t)p(t)$. We can show that the Fourier transform of w is equal to the DTFT of y. This gives us a way to relate the Fourier transform of x to the DTFT of its samples y.

Recall that multiplication in the time domain results in convolution in the frequency domain, 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 19) that the Fourier transform 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 of Dirac delta functions. The next step is to show that

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

which follows easily from the definition of the DTFT Y and the Fourier transform W. 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 Fourier transform X of the signal being sampled x to the DTFT Y of the discrete-time result y.

This important relation says that the DTFT Y of y is the sum of the Fourier transform 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 $-\pi < \omega < \pi$,

$$Y(\omega) = \frac{1}{T} X\left(\frac{\omega}{T}\right). \tag{2.5}$$

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 2.1, where X is drawn with a triangular shape.

In the second case, illustrated in figure 2.2, 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**.

From these figures, we get the guideline that we should sample continuous time signals at rates at least twice as high as the largest frequency component. This avoids aliasing distortion.

Probing further: Impulse Trains

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

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

This signal has the Fourier series expansion

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

where the fundamental frequency is $\omega_0 = 2\pi/T$. The Fourier series coefficients can be given by

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

The integral is 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 of the Dirac delta function, 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, we can write the continuous-time Fourier transform of p as

$$\forall\, \omega \in \mathbb{R}, \quad P(\omega) = \frac{2\pi}{T} \sum_{k=-\infty}^{\infty} \delta\left(\omega - \frac{2\pi}{T} k\right).$$

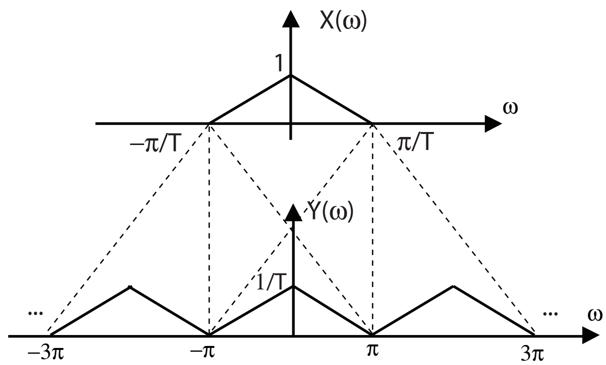

Figure 2.1: Relationship between the Fourier transform of a continuous-time signal and the DTFT of its discrete-time samples. The DTFT is the sum of the Fourier transform and its copies shifted by multiples of $2\pi/T$, the sampling frequency in radians per second. The frequency axis is also normalized.

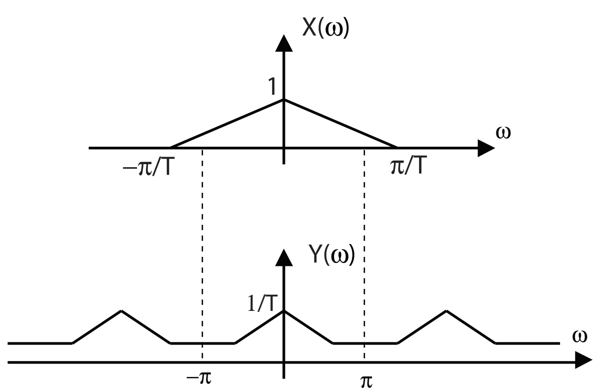

Figure 2.2: Relationship between the Fourier transform of a continuous-time signal and the DTFT of its discrete-time samples when the continuous-time signal has a broad enough bandwidth to introduce aliasing distortion.

Exercises

1. Consider the accelerometer problem described in Example 2.1. Suppose that the change in orientation x_d is a low frequency signal with Fourier transform given by

$$X_d(\omega) = \begin{cases} 2 \text{ for } |\omega| < \pi \\ 0 \text{ otherwise} \end{cases}$$

This is an ideally bandlimited signal with no frequency content higher than π radians/second, or 0.5 Hertz. Suppose further that the vibration x_n has higher frequency components, having Fourier transform given by

$$X_d(\omega) = \begin{cases} 1 \text{ for } |\omega| < 10\pi \\ 0 \text{ otherwise} \end{cases}$$

This is again an ideally bandlimited signal with frequency content up to 5 Hertz.

(a) Assume there is no frequency conditioning at all, or equivalently, the conditioning filter has transfer function

$$\forall\, \omega \in \mathbb{R}, \quad H(\omega) = 1.$$

Find the SNR in decibels.

(b) Assume the conditioning filter is an ideal lowpass filter with transfer function

$$H(\omega) = \begin{cases} 1 \text{ for } |\omega| < \pi \\ 0 \text{ otherwise} \end{cases}$$

Find the SNR in decibels. Is this better or worse than the result in part (a)? By how much?

(c) Find a conditioning filter that makes the error signal identically zero (or equivalently makes the SNR infinite). Clearly, this conditioning filter is optimal for these signals. Explain why this isn't necessarily the optimal filter in general.

(d) Suppose that as in part (a), there is no signal conditioning. Sample the signal x at 1 Hz and find the SNR of the resulting discrete-time signal.

(e) Describe a strategy that minimizes the amount of signal conditioning that is done in continuous time in favor of doing signal conditioning in discrete time. The motivation for doing this is analog circuitry can be much more expensive than digital filters.