Introduction to DSP

This lecture introduces DSP systems and reviews the sampling theorem, the effects of quantization and discrete-time convolution.

After this lecture you should be able to: explain some of the advantages and disadvantages of DSP, predict effects of changing the sampling frequency, anti-aliasing and reconstruction filters. You should also be able to compute (by hand) the linear convolution of two discrete-time sequences.

This material is covered in Chapter 1 of the text.

A General DSP System

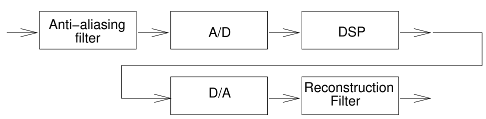

The figure above shows a general DSP system. An analog input is filtered to band-limit it. It is then sampled and quantized by an analog-to-digital converter (ADC). This "digital" signal is the processed by a digital circuit, often a computer. The processed digital signal is then converted back to an analog signal by a digital-to-analog converter (DAC). The resulting step waveform is converted to a smooth signal by a reconstruction (or "anti-image") filter.

Advantages of DSP

A DSP implementation:

- allows more complex processing than is possible with analog circuitry
- provides better signal quality and repeatable performance because of the digital processing
- is more flexible since the digital processing can often be easily modified (e.g. software can be upgraded)
- usually results in a lower cost for equivalent performance

Disadvantages of DSP

The main disadvantage of DSP techniques is that they are limited to signals with relatively low bandwidths. The point at which DSP becomes too expensive will depend on the application and the current state of conversion and digital processing technology. Currently DSP systems are used for signals up to video (about 5 MHz) bandwidths. The cost of high-speed ADCs and DACs and the amount of digital circuitry required to implement very high-speed designs (> 100 MHz) makes them impractical for many applications. As conversion and digital technology improve the bandwidths for which DSP is economical continue to increase.

There are some other disadvantages to a DSP implementation which can be significant in some applications:

- the need for an ADC and DAC makes DSP uneconomical for simple applications (e.g. a simple filters)
- higher power consumption and size of a DSP implementation can make it unsuitable for simple very low-power or small-size applications

Applications

Some examples of applications for DSP include:

- digital sound recording such as CD and DAT
- speech and compression for telecommunication and storage
- implementation of wireline and radio modems (including digital filtering, modulation, echo cancelation and other functions)
- image image enhancement and compression

- speech synthesis and speech recognition

You may be interested in reading the descriptions of several DSP applications in section 1.5.

Sampling Theorem

A basic parameter of any DSP system is the rate at which the input is sampled. The sampling theorem gives the minimum sampling rate required.

If a signal of bandwidth $f_{\max}$ is sampled at a rate F_s which is more than $2f_{\max}$ (the Nyquist rate) then it is possible to recover the original (continuous) signal by passing the samples through an ideal low-pass "reconstruction" filter.

Aliasing

However, if the signal is not strictly band-limited to a bandwidth of $f_{\max}$ then the frequency components above $F_s/2$ are "folded-back" and appear in the sampled signal at a different frequency. This causes a noise-like effect.

Since practical filters cannot implement the ideal "brick-wall" amplitude response there will always be some aliasing.

Anti-Aliasing Filtering

The undesired frequency components can be attenuated by using an analog filter before sampling. The design of this filter will depend on the amount of aliasing noise that can be tolerated.

The textbook shows examples of two criteria: using a filter whose response at the maximum desired frequency is 2% of the passband signal level, and an filter that will attenuate aliased components to a level equal to the quantization noise level.

As an example the textbook uses a filter with a Butterworth response as an example. The amplitude response of a Butterworth filter of order n is:

$$|H(f)| = \frac{1}{\sqrt{1 + (f/f_c)^{2n}}}$$

Exercise: What is the attenuation of this filter at the frequency f_c?

Dynamic Range and Quantization Noise

Dynamic Range

The dynamic range of the ADC is the ratio of the maximum to minimum signal levels. In dB, this is:

$$D = 20 \log_{10} 2^B$$

Quantization Noise

Quantization noise is the noise-like distortion that appears to be added to the signal because the analog signal must be rounded (or truncated) to the nearest quantized value.

For a sinusoidal input signal at the maximum ADC input level the quantization signal-to-noise ratio in dB for a B-bit ADC is given by:

$$\text{SQNR} = 6B + 1.8$$

Exercise: What is the dynamic range of a 10-bit ADC? What is the quantization SNR for a full-scale sinusoidal input?

Effect of DAC Zero-Order Hold

The spectrum of the sampled signal contains replicas of the original signal's spectrum at multiples of the sampling frequency. By passing this sampled signal through an ideal low-pass filter with a bandwidth equal to that of the original signal we can extract only the spectrum of the original signal and thus reconstruct the original (continuous) signal.

However, typical D/A converters output a step waveform instead of the theoretical impulse train. This "zero order hold" has the same effect as passing the signal through a filter whose impulse response is a pulse whose width is the sampling period. The amplitude response of this hypothetical filter's transfer function has a $sin(x)/x$ shape and this effect must be taken into account.

This effect must be compensated for by adding a filter with the inverse ($x/sin(x)$) amplitude response in either the DSP processor or the reconstruction filter.

Reconstruction Filter

The process involved in the selection of a reconstruction filter is similar to that involved in the selection of the anti-aliasing filter. The filter must be designed so that the amount of signal from aliased frequency components that "leaks" through the filter is reduced to an acceptable level. The amount which is acceptable will depend on the application.

Discrete-Time Convolution

Convolution is the most common DSP operation. As in a continuous-time system the output of a digital filter is the convolution of the input sequence and the filter's impulse response.

The linear convolution of two discrete sequences, $x(n)$ and $h(n)$, is given by the sum:

$$y(n) = h(n) * x(n) = \sum_{k=-\infty}^{\infty} h(k)x(n-k)$$

As in the continuous-time case to compute the convolution we reverse and shift the sequence $x(k)$, multiply it by $h(k)$ and then find the sum of the products.

Exercise: Compute the convolution of the sequences $h(k) = 1, 2, 3$ and $x(k) = 2, 1$ for $n = 0, 1, 2$.

Note that the symbol$\circledast$ is normally reserved for circular convolution where the convolution sum is computed over N points and the indices of h and x are taken to be modulo-N:

$$y(n) = h(n) \circledast x(n)$$

$$= \sum_{k=0}^{N} h((k)_N)x((n-k)_N)$$