FIR Filter Design using Windows

This lecture describes a technique that can be used to obtain the coefficients for FIR filters.
After this lecture you should be able to: (1) based on its impulse response determine whether a filter has a linear phase response, and (2) design an FIR filter using the Window method.
This material is covered in Sections 6.1 to 6.5 of the textbook.

Linear Phase

Filters with linear phase characteristics avoid distorting the signal waveform. This is required for many applications such as video and data transmission and filtering of music.

Delay causes phase shifts that are proportional to frequency (i.e. linear phase). We can add or subtract delays to any system without affecting its linear-phase characteristic.

In addition to delay, there are two conditions that can cause the phase response to be constant independent of frequency. The first is when the transfer function, $H(\omega)$, is purely real. In this case the phase is zero. From the properties of the Fourier Transform if $H(\omega)$ is real then the impulse response must be even. Similarly, if the transfer function is purely imaginary (phase fixed at $\pm\pi/2$) then the impulse response must be odd.

Therefore, a condition for a system to have linear phase is that its impulse response have even or odd symmetry about some point.

Exercise: Can an FIR filter be designed to *not* have linear phase? Can an IIR filter be designed to have linear phase? If so, under what conditions?

Group Delay

When describing systems that are approximately linear phase it's often more convenient to describe how the systems deviates from the ideal (linear) phase response. For this purpose a useful measure is the *group delay* which is the derivative of the phase with respect to frequency. Using the derivative of the phase instead of the phase itself makes it easier to see differences from the ideal linear phase characteristic. For a system with linear phase the group delay will be a constant.

Exercise: What are the units of phase? What are the units of group delay?

FIR Filter Design with Windows

In theory we could just take the inverse Fourier Transform of the desired frequency response and use the values of impulse response as the filter coefficients.

Although this would work, in many cases of interest the desired frequency response is zero in the stop-band and there is no defined transition region. This finite-"duration" frequency response results (e.g. a pulse) in an infinitely long impulse response (e.g. $sinc(f)$). It would require infinitely many filter coefficients to implement a filter with such an impulse response.

We thus approximate the desired impulse response by truncating it. The effect of truncating the impulse response to a length of N samples is the same as multiplying it by a "window" function of width N. Viewed in the frequency domain the effect of the windowing is to convolve the desired transfer function with the Fourier Transform of the window function. The basis of the window filter design method is to select a window function that will modify the desired transfer function in such a way that the resulting filter transfer function is "close enough" to the desired one.

The table below shows some of the different window functions that are available. Each window function is a compromise between the width of the main lobe and the sidelobe levels.

Window Name	Main Lobe Width	Sidelobe Level	Stopband Attenuation
Rectangular	1	13	21
Triangular	2	26	25
Hanning	3	31	44
Hamming	3	41	53
Blackman	5	57	74

The convolution of the Fourier Transform of the window function and the desired transfer function will "smear" the desired transfer function. This will have two important effects. The first is that the transition region will now have a finite width. This width will be approximately the width of the main lobe of the Fourier Transform of window function. The second effect is that the response will now will be "leakage" into the passband from the stopband causing a non-zero stopband attenuation. The level of the stopband attenuation will depend on the sidelobe levels of the Fourier Transform of the window function.

The width of the main lobe for a rectangular window will be approximately the inverse of the window duration. The table above shows the relative width of the main lobe for the other windows.

Exercise: What is the approximate width of the transition region (in Hz) for a 80-coefficient filter designed using a Hamming window if the sampling rate is 8 kHz.

Another window function, the Kaiser window, has a parameter that allows for tradeoffs between main lobe width and the stopband attenuation. It is optimum in the sense of having the most energy in the main lobe for a given peak sidelobe level.

The method for designing a filter using the window method can be summarized as follows:

- select a window function (e.g. Hamming) that meets the required stopband attenuation and (possibly) passband ripple

- compute the approximate number of filter coefficients, N, from the required transition region bandwidth

- compute the N values of the window function $w(n)$

- compute the impulse response of the ideal filter for a filter whose cutoff frequencies are centered in the transition regions. For simple filters closed form solutions can be found (see table 6.2 in the text). In other cases the DFT can be used if enough zero-padding is used to reduce the effects of periodicity.

- multiply the window and impulse response values to obtain the filter coefficients. Typically these will be symmetrical.

- compute the filter's frequency response to check that the result meets the requirements

Exercise: A 400 Hz to 3.4 kHz band-pass filter with a minimum attenuation at DC and 4 kHz of 50 dB is required. The sampling rate is 8 kHz. Which of the above windows would be suitable for use in designing such a filter? Which would require the least number of filter coefficients to meet the specifications? Approximately how many filter coefficients will be required?

The main advantage of the window FIR filter design method is that it is relatively easy to compute the coefficient values. However, with the widespread availability of PCs the window design method has largely been replaced by the optimum equiripple ("Parks-McClellan") method that will be described in the next lecture.