Lecture 11

Windowing

Introduction

We have learned how to design FIR filters to perform lowpass, highpass, bandpass and bandstop operations. We can control the "steepness" of the transition from pass band to stop band with our choice of the filter length. However, we've seen that the filter response is always plagued by the Gibbs phenomena; there are "ripples" in both the pass and stop bands. The solution to the ripple problem is *windowing* which involves a trade-off between ripple magnitude and width of the transition band. We will explore this trade-off by examining a few specific windows.

Effect of windowing on frequency response

The output y of a discrete LTI system is a convolution of the input x with an impulse response h

$$y(n)=\sum_{k=-\infty}^{\infty} h(k)x(n-k) \tag{1}$$

This is expressed in the frequency domain as

$$Y(\omega)=H(\omega)X(\omega)$$

where

$$H(\omega)=\sum_{n=-\infty}^{\infty} h(n)e^{-j\omega n}$$

is the system frequency response.

Suppose we multiply the impulse response by a *window* function $w(n)$ to produce a new impulse response

$$h_w(n)=w(n)h(n)$$

The corresponding frequency response is

$$H_w(\omega)=\sum_{n=-\infty}^{\infty} h_w(n)e^{-j\omega n}=\sum_{n=-\infty}^{\infty} w(n)h(n)e^{-j\omega n}$$

Express $h(n)$ as an inverse Fourier transform

$$H_w(\omega)=\sum_{n=-\infty}^{\infty} w(n)\left[\frac{1}{2\pi}\int_{-\pi}^{\pi} H(\alpha)e^{j\alpha n}d\alpha\right]e^{-j\omega n}$$

Change the order of summation and integration to get

$$\begin{aligned}H_w(\omega)&=\frac{1}{2\pi}\int_{-\pi}^{\pi} H(\alpha)\left[\sum_{n=-\infty}^{\infty} w(n)e^{-j(\omega-\alpha)n}\right]d\alpha\\&=\frac{1}{2\pi}\int_{-\pi}^{\pi} H(\alpha)W(\omega-\alpha)d\alpha\end{aligned} \tag{2}$$

where

$$W(\omega)=\sum_{n=-\infty}^{\infty} w(n)e^{-j\omega n} \tag{3}$$

is the Fourier transform of the window. The frequency response of the windowed impulse response is the convolution of the non-windowed (ideal) response with the response of the window.

Selected windows

The topic of windows is extensive. Reference [1] gives a good overview of the subject. Here we examine a few including the *Hamming window* which we will adopt as our default window for the remainder of the course.

Rectangular window

The rectangular window corresponds to simple truncation. It is the window we have implicitly used in designing FIR filters so far. It's definition is

$$w(n)=\begin{cases}1 & 0\le n\le M\\ 0 & \text{otherwise}\end{cases} \tag{4}$$

This window and its frequency response are graphed in Fig. 1. The frequency response displays a relatively narrow *mainlobe* centered on $f=0$. If this was the only component in the frequency response, its effect would be a smoothing or "blurring" of the ideal filter response. However, there are also significant *sidelobes*, the first having amplitudes of about $-13\,\text{dB}$. These are what give rise to the Gibbs phenomena and make the rectangular window largely unacceptable for practical filter design.

Triangular window

Assuming the discontinuity of the rectangular window at $n=0,M$ is the cause of the large sidelobes, an obviously strategy is to adopt a window which, in some sense, "gradually" tapers to zero at its left and right edges.

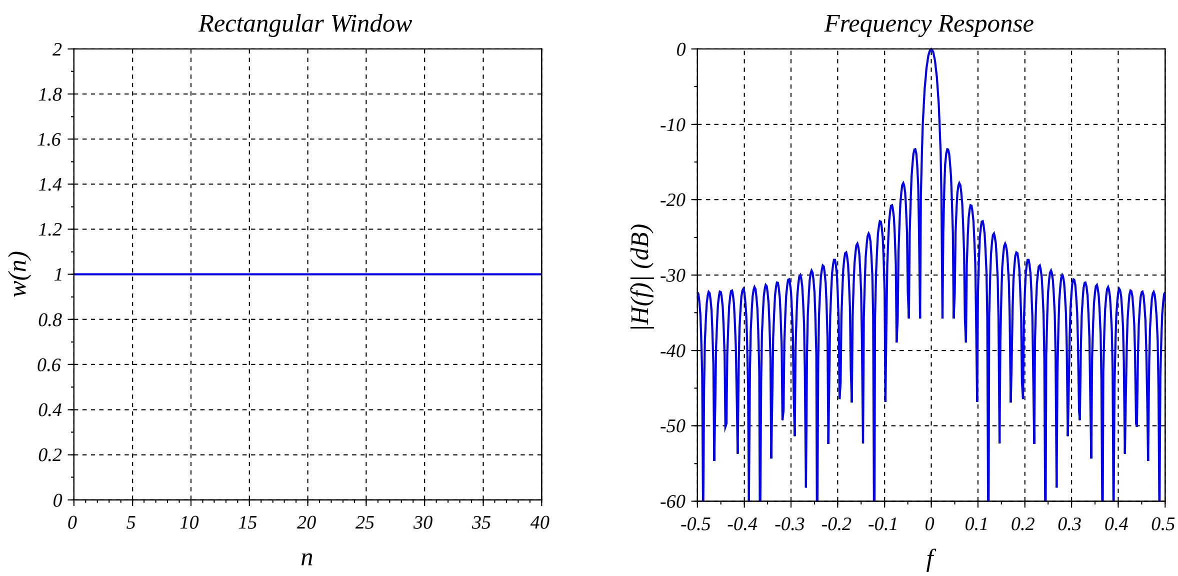

Fig. 1: (Left) rectangular window for $M=40$. (Right) magnitude frequency response (normalized).

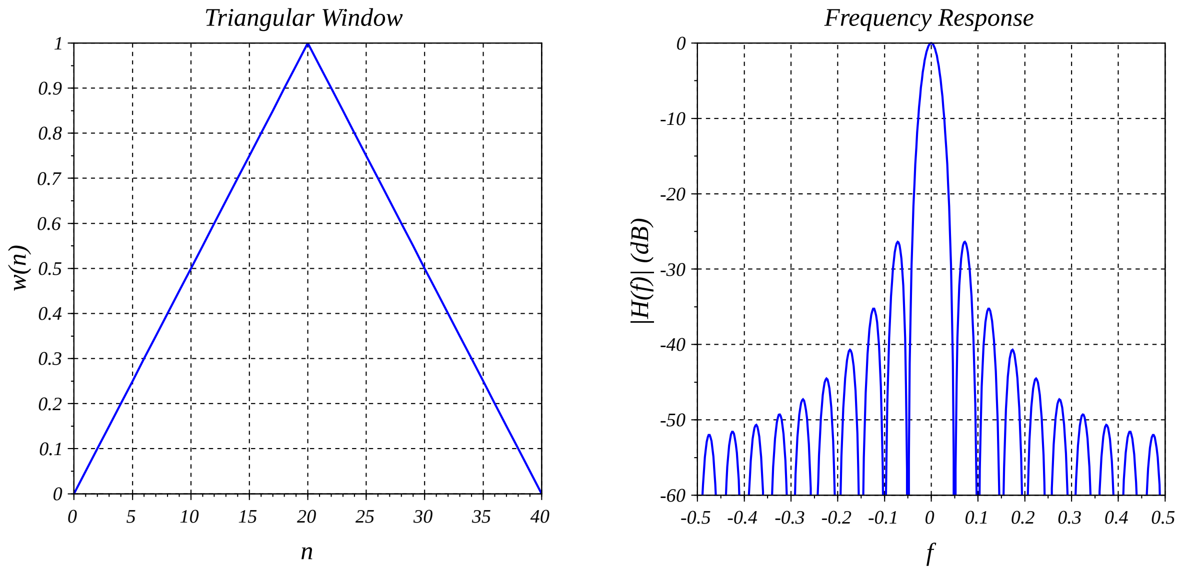

Fig. 2: (Left) Triangular window for $M=40$ *. (Right) Magnitude frequency response (normalized).*

The triangular window

$$w(n)=\begin{cases} 1-\left|\dfrac{n-M/2}{M/2}\right| & 0\le n\le M \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

has a linear taper. This window and its frequency response are graphed in Fig. 2.

Two things stand out in the frequency response of this window relative to the rectangular window. First, the side lobes are much reduced in amplitude. Second, the main lobe is wider. We seem to have traded off resolution for side-lobe suppression.

Hann window

The Hann window

$$w(n)=\begin{cases} \dfrac{1}{2}\left[1-\cos\left(\dfrac{2\pi n}{M}\right)\right] & 0\le n\le M \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

is type of “raised-cosine” window (Fig. 3). It is sometimes mistakenly called a “Hanning” window, possibly due to confusion with the “Hamming” window we discuss below. It smoothly transitions, in both amplitude and slope, to zero at its boundaries. Its response has a slightly wider main lobe, but smaller sidelobes than the triangular window response. The first sidelobe amplitude is about $-32\,\text{dB}$. Subsequent sidelobe amplitudes are even lower.

Hamming window

The Hann window has a constant term and a cosine term, each with amplitude ½. The Hamming window (Fig. 4) is essentially a Hann window in which these amplitudes have been adjusted to suppress the first sidelobe of the frequency response. The required amplitudes are $25/46$ and $21/46$, respectively

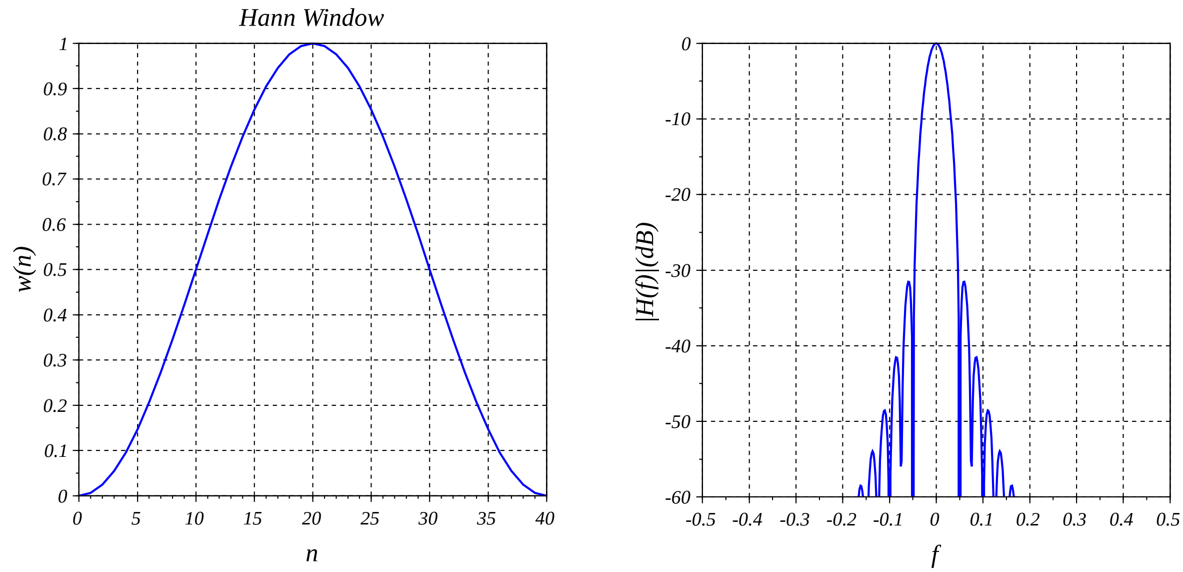

Fig. 3: (Left) Hann window for $M=40$. (Right) Magnitude frequency response (normalized).

$$w(n)=\begin{cases}\frac{25}{46}-\frac{21}{46}\cos\left(\frac{2\pi n}{M}\right) & 0\le n\le M\\ 0 & \text{otherwise}\end{cases} \tag{7}$$

As a result the Hamming window is not zero at its edges, but rather has the value

$$w(0)=w(M)=\frac{4}{46}>0 \tag{8}$$

Although this does null the first sidelobe, subsequent sidelobes do not fall off as fast as they do in the Hann window frequency response. This is generally considered an advantageous trade off. Using the approximations $25/46\approx 0.54$ and $21/46\approx 0.46$

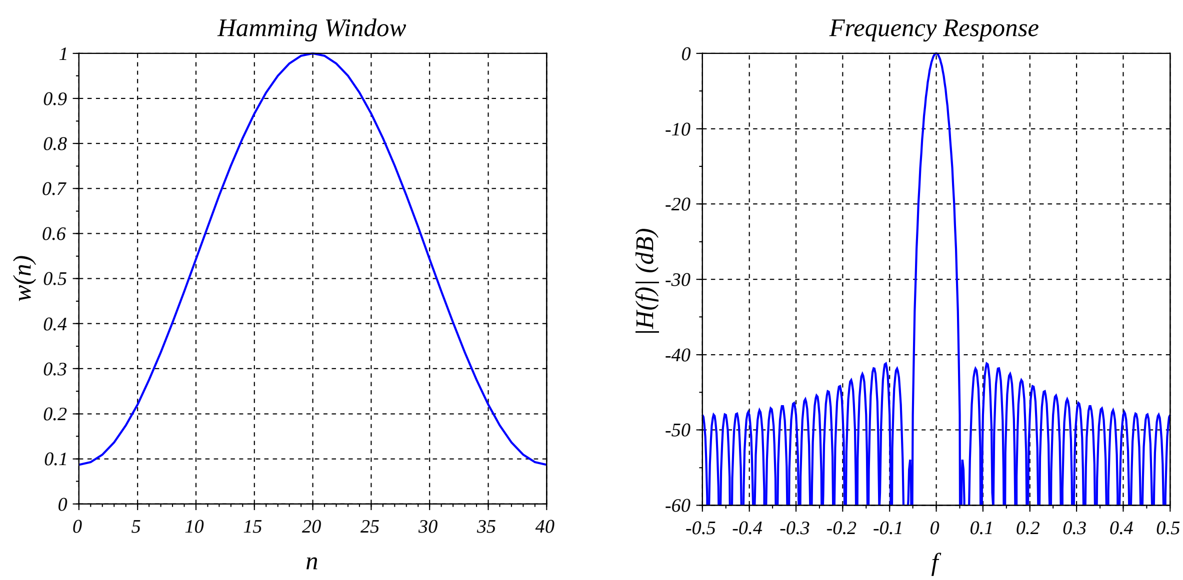

Fig. 4: (Left) Hamming window for $M=40$. (Right) Magnitude frequency response (normalized).

$$w(n)=\begin{cases} 0.54-0.46\cos\left(\frac{2\pi n}{M}\right) & 0\le n\le M \\ 0 & \text{otherwise}\end{cases} \tag{9}$$

Actually brings the second sidelobe down about 1 dB more to about $-43\,\text{dB}$ (Fig. 5). Form (9), that we might call a "modified Hamming window" is what is usually implemented as a "Hamming window" as opposed to Hamming's original version (7). In this course our "go-to window" will be the Hamming window (form (9)) as it provides a good compromise between main-lobe width and side-lobe amplitudes.

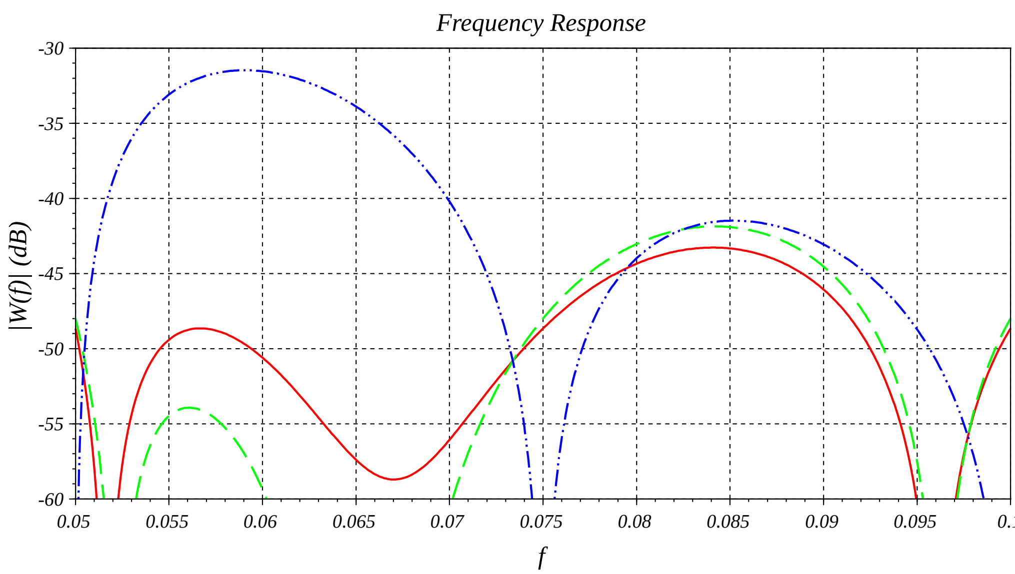

Fig. 5: Response of $M=40$ Hann (blue dot-dash), Hamming (green dashed), and modified Hamming (solid red).

Blackman-Nuttall Window

A window that trades off even more main lobe width for smaller sidelobe ripples is the Blackman-Nuttall window. Its definition is

$$w(n)=\begin{cases} a_0-a_1\cos\left(\frac{2\pi n}{M}\right)+a_2\cos\left(\frac{4\pi n}{M}\right)-a_3\cos\left(\frac{6\pi n}{M}\right) & 0\le n\le M \\ 0 & \text{otherwise}\end{cases} \tag{10}$$

with

$$a_0=0.3635819\,,\ a_1=0.4891775\,,\ a_2=0.1365995\,,\ a_3=0.0106411$$

The window and its frequency response are shown in Fig. 6. The mainlobe is considerably wider than those of the windows we have previously considered, but the sidelobes are also considerably lower, a bit below $-90\,\text{dB}$. This window goes even further in the the mainlobe-width vs. sidelobe-amplitude trade off.

Scilab code to apply Hann, Hamming and Blackman-Nuttall windows to an impulse response is given in the Appendix.

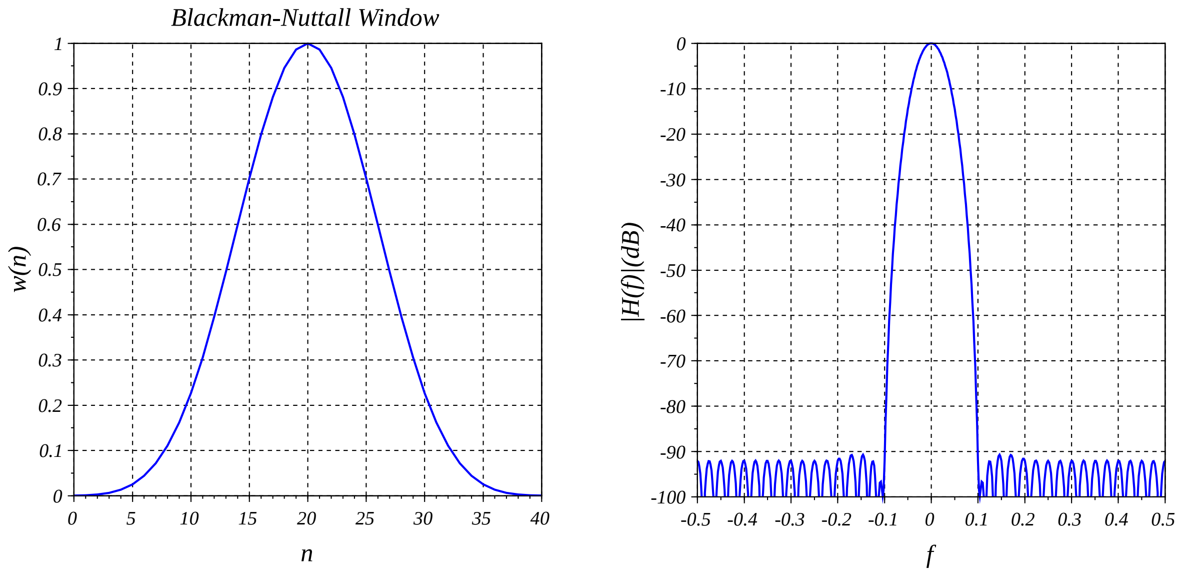

Fig. 6: (Left) Blackman-Nuttall Window for $M=40$. (Right) Magnitude frequency response (normalized).

Other window types

All the windows we have considered so far have a mainlobe-plus-sidelobe structure. Although very common, this is not a necessary property of a window. For example, the Hann-Poisson window (Fig. 7) adds an extra exponential taper to the Hann window with the result that there are no sidelobes. Instead there is a single mainlobe that decreases monotonically and never reaches zero. This makes it useful in algorithms where it is critical that the frequency response have no ripples due to sidelobes. For example, where maxima in the frequency response need to be found using numerical optimization methods. There is no "best" window. There are only "best" windows for given applications.

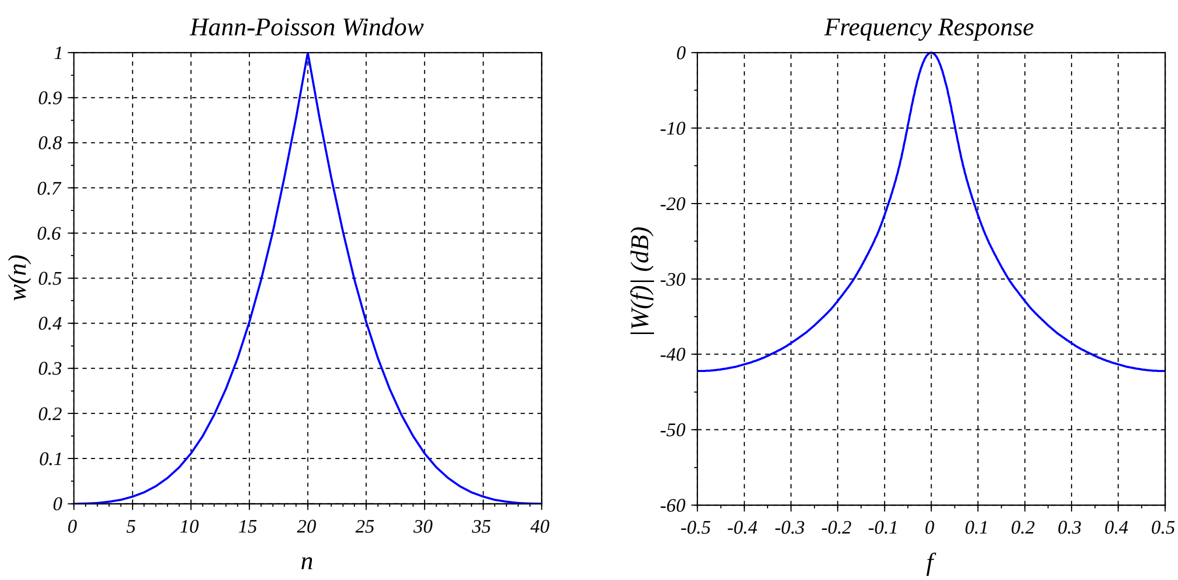

Fig. 7: (Left) Hann-Poisson Window for $M=40$. (Right) Magnitude frequency response (normalized). This window has the property that its frequency response decreases monotonically, so there are "no sidelobes."

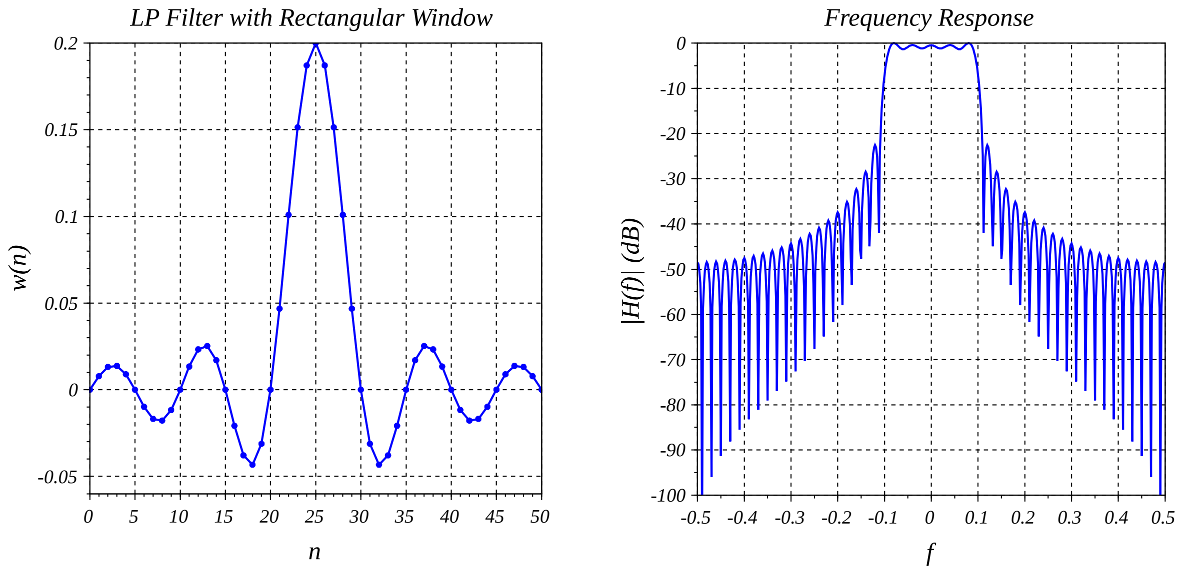

Fig. 8: Lowpass filter with rectangular window. (Normalized to 0 dB peak frequency response.)

Frequency response of windowed FIR filters

Now we let the rubber hit the proverbial road and see what effects these various windows have on filter responses. In Fig. 8, Fig. 9 and Fig. 10 we show the impulse and frequency response of a lowpass filter with $f_c = 0.1$, $M = 50$ and rectangular, Hamming and Blackmann-Nuttall windows, respectively.

The progressively larger mainlobe widths of the window's frequency responses translate into less rapid transition from passband to stopband. Smaller sidelobes translate into smaller ripples in the passband and stopband.

What is the practical implication of a frequency response having ripples? In `sounds.zip` there are four files named `chirp.mp3`, `chirpLPrectangle.mp3`,

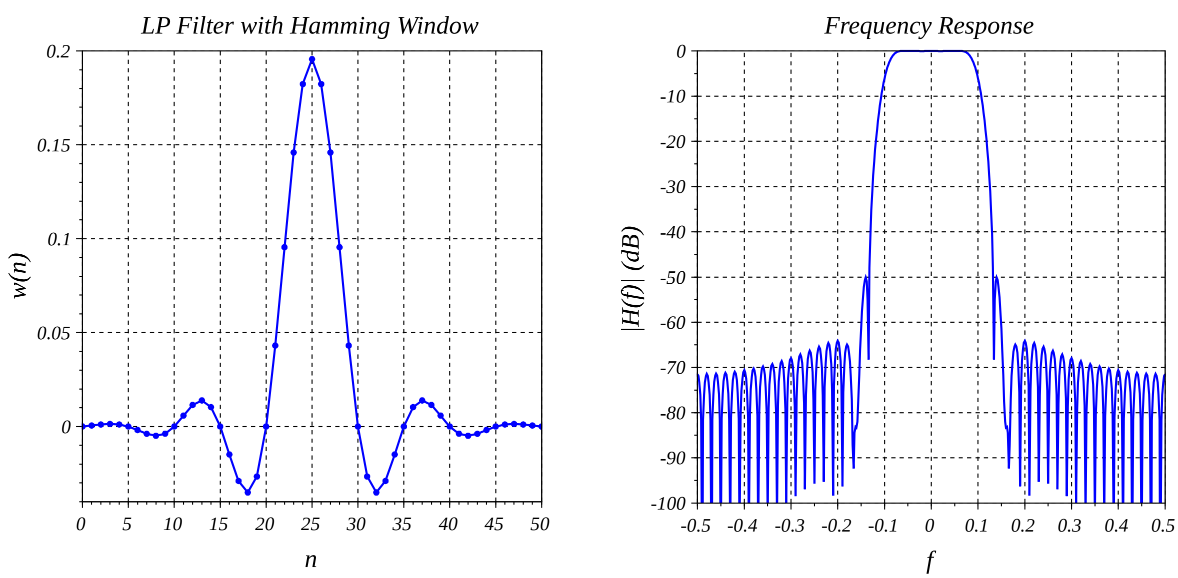

Fig. 9: Lowpass filter with Hamming window.

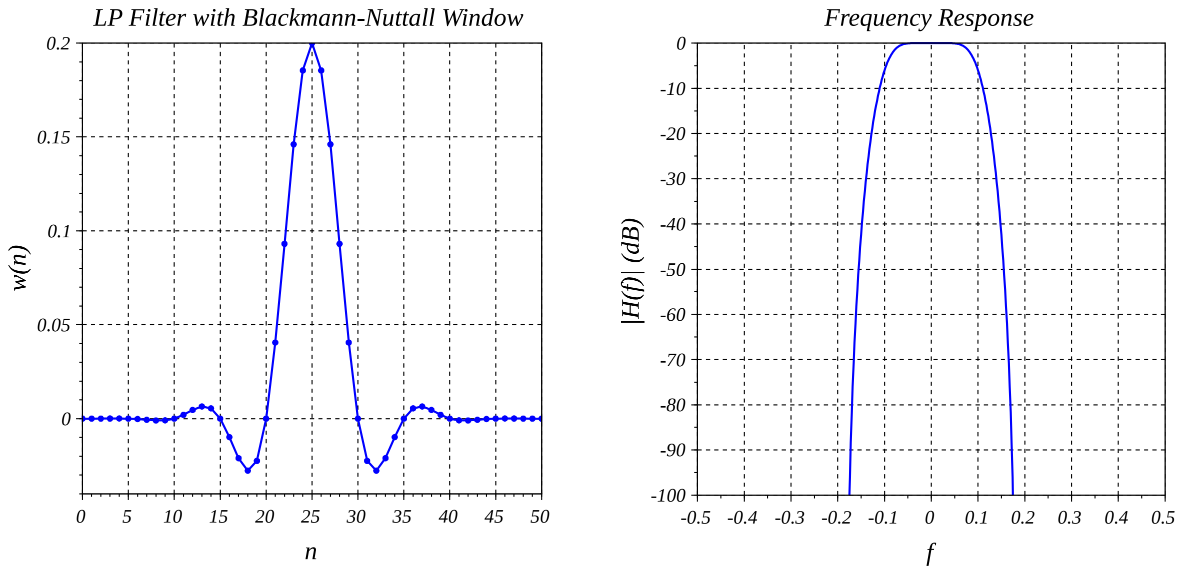

Fig. 10: Lowpass filter with Blackmann-Nuttall window.

`chirpLPhamming.mp3`, and `chirpLPnuttall.mp3`. The first is a "chirp" signal (linearly increasing frequency) covering frequencies $0 \le f \le 0.2$ with $F_s = 8\,\mathrm{kHz}$. The other three are lowpass filtered versions with cutoff frequency $f_c = 0.1$ and $M = 100$ using rectangular, Hamming, and Blackman-Nuttall windows. You can clearly hear the sidelobes of the rectangular window as a series of pulsations in the stopband. These are barely audible with the Hamming window and not audible with the Blackman-Nuttall window.

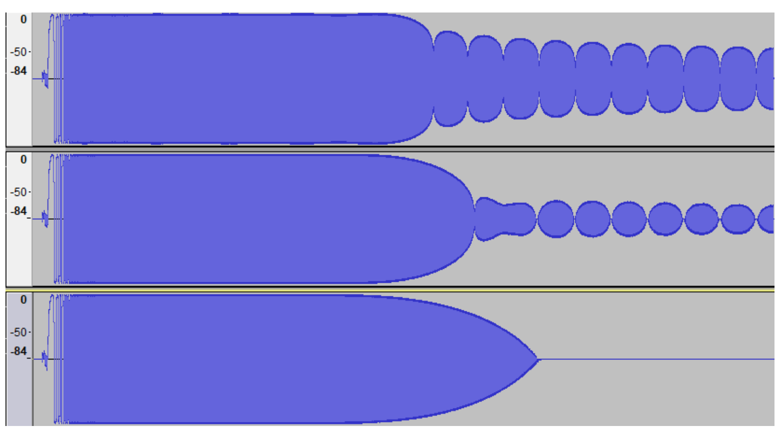

Fig. 11: Lowpass-filtered chirp signal using (top) rectangular, (middle) Hamming, and (bottom) Blackman-Nuttall windows. Amplitudes are in dB ranging from 0 to −84 dB. (Audacity screenshot.)

The filtered audio waveforms are shown in Fig. 11. The transition from passband to stopband (amplitude reduction in filtered signal) is most rapid for the rectangular window, less rapid for the Hamming window and least rapid for the the Blackman-Nuttall window. This is the transition-band vs. side-lobe trade off.

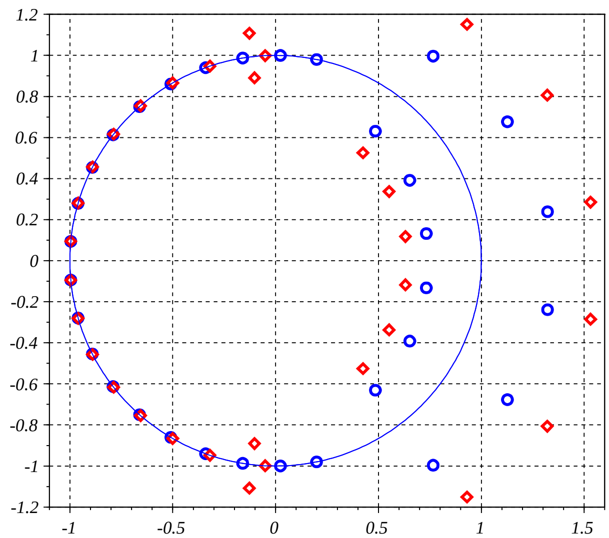

Fig. 12: Lowpass filter z-plane zeros with (circles) rectangular window and (diamonds) Hamming window.

Effect of windowing on transfer function zeros

In the previous lecture we pointed out that any FIR filter transfer function can be written in factored form as

$$H(z)=b_0(1-\zeta_1 z^{-1})(1-\zeta_2 z^{-1})(1-\zeta_3 z^{-1})\cdots(1-\zeta_M z^{-1}) \tag{11}$$

where the ζ_k are the zeros of $H(z)$. To within the multiplicative constant b_0, $H(z)$ is fully specified by the locations of these zeros in the *z* plane. The effect of a particular window can be thought of as simply moving the transfer function zeroes. For example, Fig. 12 shows the zeros of a lowpass filter with rectangular and Hamming windows.

Perfect Reconstruction

The desired (ideal) impulse response of a lowpass filter is

$$h_d^{LP}(n)=\frac{\sin(\omega_c n)}{\pi n} \ , \ h_d^{LP}(0)=\frac{\omega_c}{\pi} \tag{12}$$

As we have noted previously, the corresponding highpass filter impulse response is $\delta(n)$ minus the lowpass filter impulse response. Therefore

$$h_d^{LP}(n)+h_d^{HP}(n)=\delta(n) \tag{13}$$

and the sum of lowpass- and highpass-filtered signals is the original signal. This is the *perfect reconstruction* property of FIR filters. Let’s see if this still applies when we apply an arbitrary window function?

For the lowpass and highpass filters with impulse response of length $M+1$ let’s write

$$h_w^{LP,HP}(n)=h_d^{LP,HP}(n-M/2)w(n) \tag{14}$$

Here $h_d^{LP,HP}(n-M/2)$ is the ideal impulse response shifted to produce a causal impulse response and $w(n)$ is the windowing function. Summing the windowed lowpass and highpass impulse responses we find

$$\begin{aligned} h_w^{LP}(n)+h_w^{HP}(n) &= \left[h_d^{LP}(n-M/2)+h_d^{HP}(n-M/2)\right]w(n) \\ &= \delta(n-M/2)\,w(n) \\ &= \begin{cases} w(M/2) & n=M/2 \\ 0 & n\neq M/2 \end{cases} \end{aligned} \tag{15}$$

The only condition for a windowed FIR filter to provide perfect reconstruction is $w(M/2)=1$. All of the windows we have considered have this property. So, provided they have the same length, same cutoff frequency, and use the same window, the two-band equalizer shown in Fig. 13 will achieve $y(n)=x(n)$ when $g_{\text{bass}}=g_{\text{treble}}=1=0\,\text{dB}$ regardless of the window used.

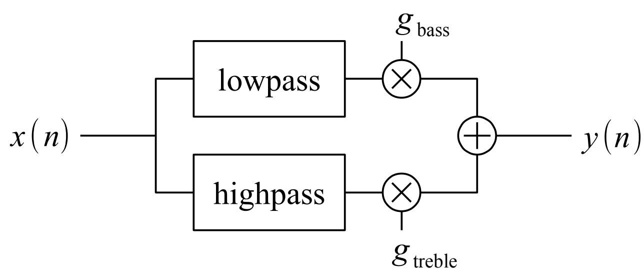

Fig. 13: Two-band equalizer.

Likewise for the *N*-band equalizer shown in Fig. 14. Provided the filters have the same length, matching band edges, and use the same window, we obtain $y(n)=x(n)$ when $g_1=g_2=\cdots=g_N=1=0\,\text{dB}$.

References

1. https://en.wikipedia.org/wiki/Window_function

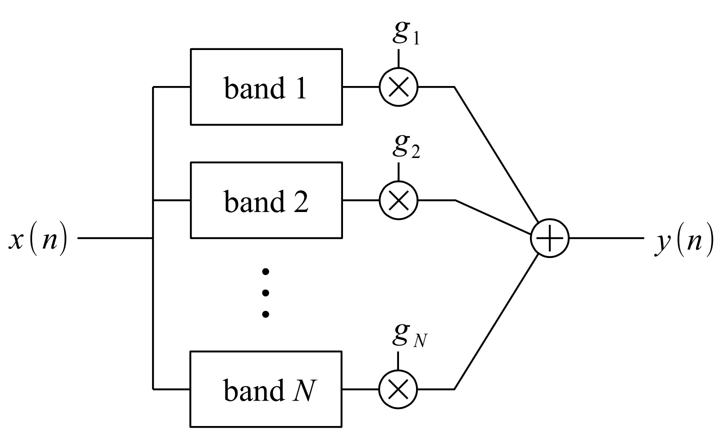

Fig. 14: N-band equalizer.

Appendix

```
//applies a Hann window to signal h(n)
function hw = winHann(h)
  N = length(h);
  hw = h.*(1-cos(2*%pi*[0:N-1]/(N-1)))/2;
endfunction

//applies a Hamming window to signal h(n)
function hw = winHamming(h)
  M = length(h)-1;
  hw = h.*(0.54-0.46*cos(2*%pi*[0:M]/M));
endfunction

//applies a Blackman-Nuttall Window to signal h(n)
function hw = winBlackmanNuttall(h)
  N = length(h);
  p = 2*%pi*[0:N-1]/(N-1);
  hw = h.*(0.3635819-0.4891775*cos(p)..
    +0.1365995*cos(2*p)-0.0106411*cos(3*p));
endfunction
```