IIR Filter Design

Topics: IIR filter design, pole-zero placement, allpass, notch filters

An IIR filter is computed by the basic formula that combines inputs x[n] and outputs y[n] as follows:

$$y[n]=\sum_{k=0}^{M} b_k x[n-k]+\sum_{k=1}^{N} a_k y[n-k]$$

The coefficients b_k are the "feedforward"coefficients, because they only act on the input signal x[n], and the coefficients a_k are the "feedback" coefficients, because they act on past outputs y[n]. We have already seen that by using the Z-transform that we obtain a simple input-output relationship:

$$\boldsymbol{Y}(z)=\boldsymbol{H}(z)\boldsymbol{X}(z)$$

The transfer function of the IIR filter is:

$$\boldsymbol{H}(z)=\frac{b_0+b_1 z^{-1}+\cdots+b_M z^{-M}}{1-a_1 z^{-1}-\cdots-a_N z^{-N}}$$

Notice the numerator and denominator are both polynomials, and therefore, like all polynomials, have roots. The roots of the numerator are called the "zeros" of the filter, and the roots of the denominator, are called the "poles". We can factor both the numerator and denominator into their roots, and write:

$$\boldsymbol{H}(z)=\frac{b_0 z^{-M}(z-z_1)(z-z_2)\cdots(z-z_M)}{z^{-N}(z-p_1)\cdots(z-p_N)}$$

Notice that there are M zeros, $z_{1,}\ldots, z_M$, and N poles, $p_{1,}\ldots, p_N$. At each of the zeros, $\boldsymbol{H}(z)=0$, i.e., $\boldsymbol{H}(z_k)=0$ for $k=1,\ldots, M$. In contrast, at each of the poles, $\boldsymbol{H}(z)$ is not defined since the denominator is zero there. Furthermore, there are z^{N-M} zeros (or poles, if M>N) at the origin $z=0$.

Example:

$$\boldsymbol{H}(z)=\frac{1+2z^{-1}+1z^{-2}}{1-z^{-1}+0.25z^{-2}}=\frac{z^{-2}(z+1)(z+1)}{z^{-2}(z-0.5)(z-0.5)}=\frac{(z+1)(z+1)}{(z-0.5)(z-0.5)}$$

This filter has two zeros at $z=-1$ and two poles at $z=0.5$

Frequency response in terms of poles and zeros

The frequency response can be determined from the transfer function by substituting $z=e^{j2\pi F}$. This gives

$$H(F)=\frac{b_0+b_1 e^{-j2\pi F}+\cdots+b_M e^{-j2\pi FM}}{1-a_1 e^{-j2\pi F}-\cdots-a_N e^{-j2\pi FN}}=\frac{b_0 e^{-j2\pi F M}(e^{j2pi F}-z_1)(e^{j2pi F}-z_2)\cdots(e^{j2pi F}-z_M)}{e^{-j2\pi F N}(e^{j2\pi F}-p_1)\cdots(e^{j2\pi F}-p_N)}$$

What is important is to understand how each pole or zero affects the frequency response. Consider a filter with a single pole at $z=p_1$ and a single zero at $z=z_1$, i.e.,

$$H(F)=\frac{e^{j2\pi F}-z_1}{e^{j2\pi F}-p_1}\Rightarrow|H(F)|=\frac{|e^{j2\pi F}-z_1|}{|e^{j2\pi F}-p_1|}$$

For any frequency F, the point $z=e^{j2\pi F}$ is a point on the unit circle, and therefore the magnitude response is the ratio of lengths from $z=e^{j2\pi F}$ to the zero $z=z_1$ and to the pole $z=p_1$. This is illustrated in Figure 1 below.

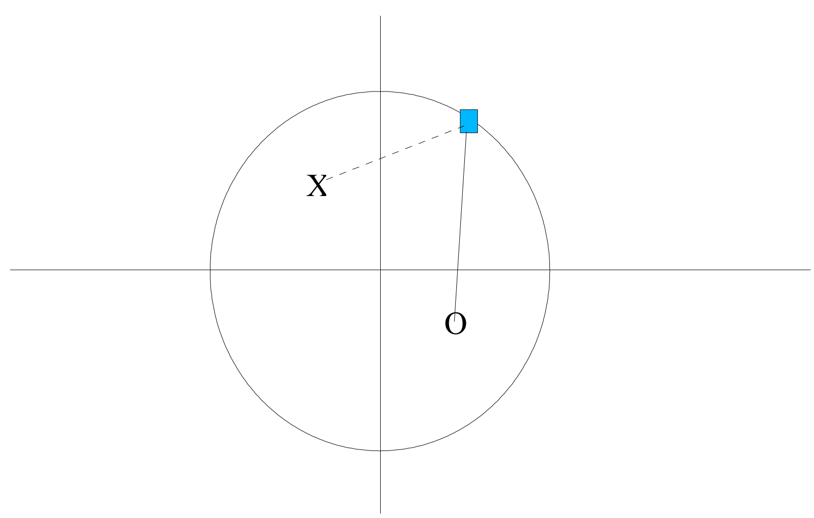

Figure 1: The magnitude response at any frequency (indicated by the point on the unit circle) is the ratio of lengths to the zero (solid) and to the pole (dashed).

Consequently, when $z = e^{j2\pi F}$ approaches the location of the zero, the numerator decreases, and therefore so too does $|H(F)|$. Conversely, when $z = e^{j2\pi F}$ approaches a pole, the denominator decreases, and therefore, the magnitude *increases*.

Example: a notch filter

A notch filter is a filter that rejects only one frequency, called the "notch" frequency. Notch filters are frequently used to remove 60 Hz power line interference in sensor data. A 60 Hz notch filter is easily created by placing a zero at $F_n = 60/f_s$, yielding

$$\boldsymbol{H}_1(z) = (1 - e^{j2\pi F_n} z^{-1})(1 - e^{-j2\pi F_n} z^{-1}) = 1 - 2\cos(2\pi F_n) z^{-1} + z^{-2}$$

Notice that we need two zeros, at complex-conjugate locations, in order to obtain a filter with real coefficients. This is a simple FIR filter with 3 coefficients.

A sharper notch results from placing a pole at $z = r e^{j2\pi F_n}$, with $0 \le r < 1$, so that it is just inside the unit circle, at the same orientation as each of the zeros. This gives

$$\boldsymbol{H}_2(z) = \frac{(1 - e^{j2\pi F_n} z^{-1})(1 - e^{-j2\pi F_n} z^{-1})}{(1 - r e^{j2\pi F_n} z^{-1})(1 - r e^{-j2\pi F_n} z^{-1})} = \frac{1 - 2\cos(2\pi F_n) z^{-1} + z^{-2}}{1 - 2r\cos(2\pi F_n) z^{-1} + r^2 z^{-2}}$$

The frequency response of both types of notch filters is plotted below[1].

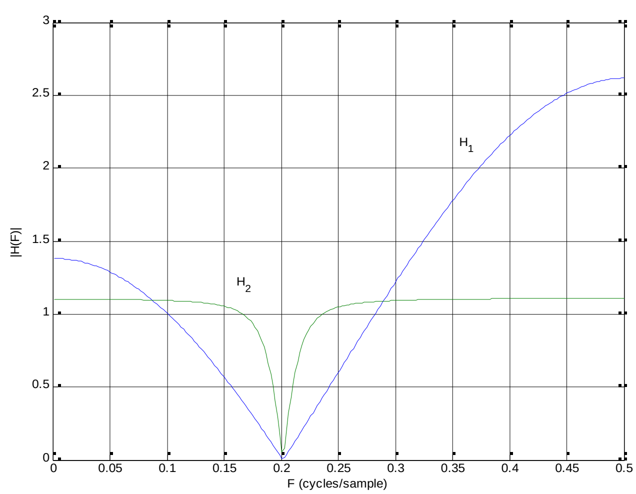

Figure 2: The magnitude response of the two notch filters $\boldsymbol{H}_1$ *and* $\boldsymbol{H}_2$ *(with r=0.9).*

1 The MATLAB script "iirfreqresp.m" was used to plot this response.

Example: tunable filters

These are examples of filters whose frequency characteristics are easily adjusted.

Lowpass filter

$$\boldsymbol{H}_{LP}(z)=\frac{1-\alpha}{2}\frac{1+z^{-1}}{1-\alpha z^{-1}}$$

Properties: $|H_{LP}(0)|=1$, $|H_{LP}(1/2)|=0$, i.e., the DC response is 1, and the filter has a null at Nyquist. Futhermore, it can be shown that $|H(F_c)|=\frac{1}{\sqrt{2}}$ at the frequency

$$F_c=\frac{1}{2\pi}\cos^{-1}(\frac{2\alpha}{1+\alpha^2})\leftarrow\rightarrow\alpha=\frac{1-\sin(2\pi F_c)}{\cos(2\pi F_c)}$$

The frequency F_c is called the "3-dB frequency" of the filter, since the magnitude response is down by 3 dB from the value at DC, i.e.,

$$20\log_{10}|H(F_c)|=20\log_{10}\left|\frac{1}{\sqrt{2}}\right|=-3\,dB$$

For example, if $\alpha=\frac{1}{2}$, then $F_c\approx 0.10$ cycles/sample. Note that $0\leq\alpha<1$ is required.

Highpass filter

$$\boldsymbol{H}_{HP}(z)=\frac{1-\alpha}{2}\frac{1-z^{-1}}{1+\alpha z^{-1}}$$

Properties: $|H_{LP}(0)|=0$ and $|H_{LP}(1/2)|=1$, i..e, a null at DC and unity gain at Nyquist. Its 3-dB frequency is given by $F_c=\frac{1}{2}-\frac{1}{2\pi}\cos^{-1}(\frac{2\alpha}{1+\alpha^2})$

Bandpass (BP) and bandstop (BS) filters

$$\boldsymbol{H}_{BP}(z)=\frac{1-\alpha}{2}\frac{1-z^{-2}}{1-\beta(1+\alpha)z^{-1}+\alpha z^{-2}} \quad and \quad \boldsymbol{H}_{BS}(z)=\frac{1+\alpha}{2}\frac{1-2\beta z^{-1}+z^{-2}}{1-\beta(1+\alpha)z^{-1}+\alpha z^{-2}}$$

Properties: The bandpass filter has nulls at DC and Nyquist, while the bandstop has unity gain at those frequencis. The center of the passband for the bandpass filter (where it has unity gain) is

$$F_{center}=\frac{1}{2\pi}\cos^{-1}(\beta)\leftarrow\rightarrow\beta=\cos(2\pi F_{center})$$

The center of the stopband for the bandstop filter is given by exactly the same formula. The 3-dB bandwidth, which the width of band in between the 3-dB points, is given by the formula

$$W_3=\frac{1}{2\pi}\cos^{-1}\left(\frac{2\alpha}{1+\alpha^2}\right)\leftarrow\rightarrow\alpha=\frac{1-\sin(2\pi W_3)}{\cos(2\pi W_3)}$$

An example response for this type of bandpass filter is shown in Figure 3.

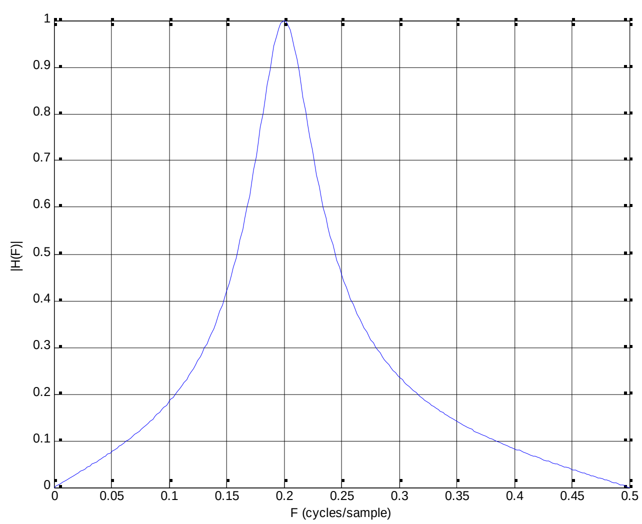

Figure 3: Magnitude response for the bandpass filter with $F_{center}=0.2$ *and* $W_3=0.05$ *. This requires* $\alpha=0.7265$ *and* $\beta=0.3090$.

The pole zero plot for this filter, obtained using MATLAB's “zplane” function is shown in Figure 3.

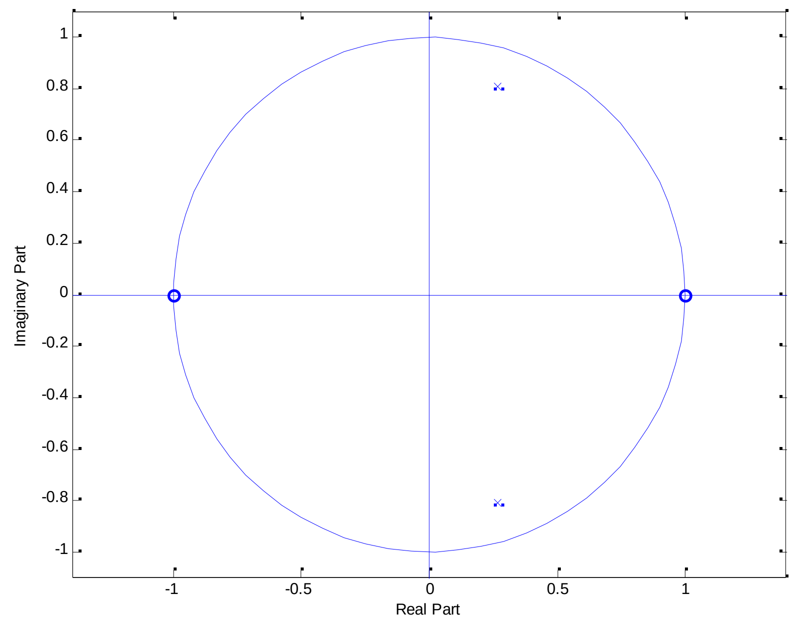

Figure 4: Pole-zero plot for the filter in Figure 3.

Cascading filters

The responses of filters connected in series multiply together, which provides a method to produce sharper filter responses.

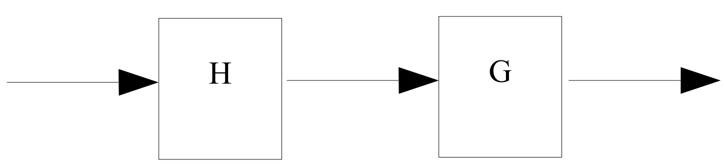

The frequency response of the series is $H(F)G(F)$, i.e., the product of the two.

Comb filters

A filter's response can be repeated in the same band of frequencies by replacing every instance of z in the transfer function $\boldsymbol{H}(z)$ by z^L. For example, suppose that we replace z for the notch filter $\boldsymbol{H}_2(z)$ in Figure 3 by z^3. We obtain the response shown in Figure 5. Note that the response repeats L times in the band [0, 1]

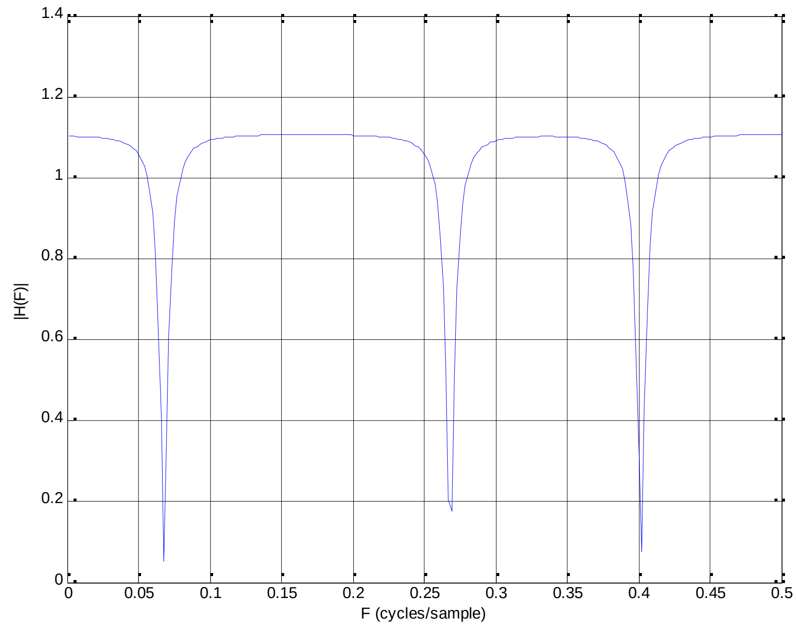

Figure 5: Comb version of the notch filter created by replacing $\boldsymbol{H}_2(z)$ from Figure 3 with $\boldsymbol{H}_2(z^3)$.

Allpass filters

These filters alter only the phase of a signal, and not its magnitude. An allpass is created by using denominator whose coefficients are in the reverse order as the numerator. Mathematically, an allpass has the form

$$\boldsymbol{H}_{AP}(z)=\frac{b_0+b_1 z^{-1}+\cdots+b_{M-1} z^{-(M-1)}+z^{-M}}{1+b_{M-1} z^{-1}+\cdots+b_0 z^{-M}}$$

Therefore, the frequency response is

$$H_{AP}(F)=e^{-j2\pi F M}\frac{b_0 e^{j2\pi F M}+b_1 e^{j2\pi F(M-1)}+\cdots+b_{M-1} e^{j2\pi F}+1}{1+b_{M-1} e^{-j2\pi F}+\cdots+b_0 z^{-j2\pi F M}}$$

Notice that aside from a factor of $e^{-j2\pi FM}$, the numerator is the complex conjugate of the denominator. Since for any complex number z , the ratio $e^{j\theta}\frac{z}{z^*}$ has magnitude one, we see that $|H_{AP}(F)|=1$ for all frequencies. Hence, an allpass filter changes only the phase of the input.

One application of an allpass filter is to create reverberation. The filter

$$\boldsymbol{H}_{rev}(z)=\frac{\alpha+z^{-R}}{1+\alpha z^{-R}}$$

is often used, where R is an integer specifying the number of samples before the initial echo. The group delay[2] of an all pass with $R=110$ is shown in Figure 6.

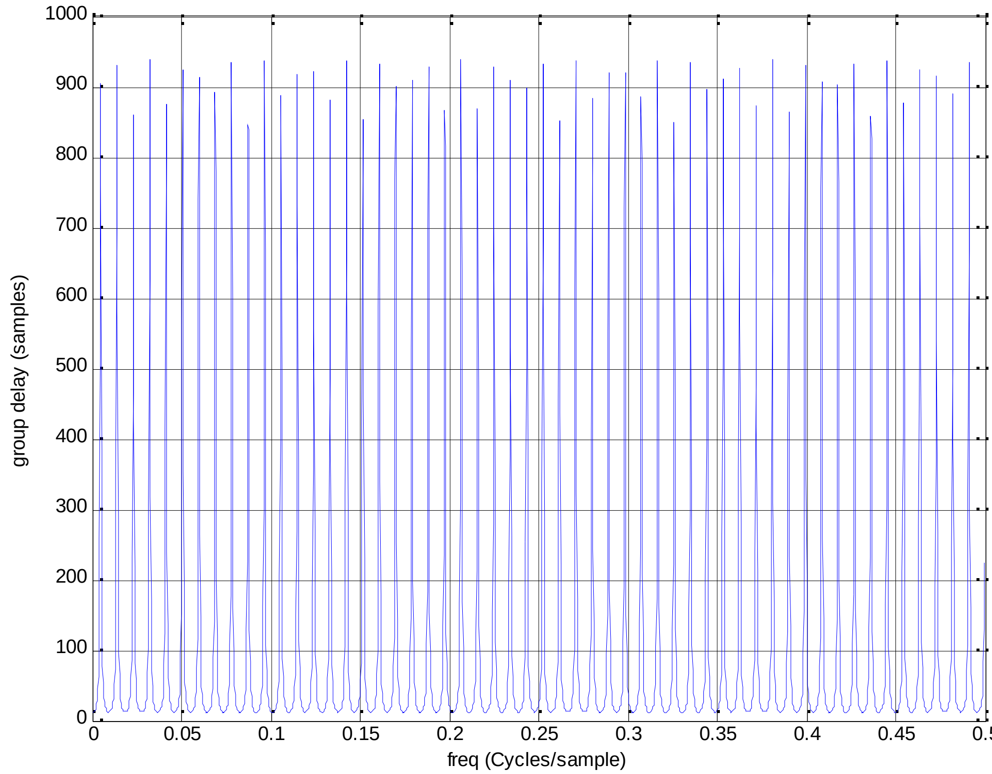

Figure 6: Group delay of the allpass filter $\boldsymbol{H}_{rev}(z)$ with $R=110$.

2 The MATLAB script "reverbdemo.m" was used to create this plot.

Classical prototypes

The study of analog filters is well-established, and many design techniques are known. There are several well-known analog prototypes, including: Butterworth, Chebyshev (types I and II), Elliptic (or Cauer). Of these, the Butterworth has the smoothest profile but the slowest transition from passband to stopband. The Elliptic has the most ripple in the passband and stopband, but the sharpest transition from passband to stopband. Design of these filters is best accomplished by using a filter design tool such as "fdatool" in MATLAB.

IIR vs FIR

Compared to IIR filter, FIR filters have two advantages: they are always stable, they are less sensitive to finite precision implementation, and they can have exactly constant group delay. IIR filters are rarely used in image processing due to the need for linear phase response (constant group delay).

The advantage of IIR filters is that they can achieve a desired sharpness of response with many fewer coefficients, and they can implement unusual characteristics such as being an allpass.