

Digital Signal Processing Laboratory (EE39203)

Autumn, 2022-23

Experiment 5	Digital Filter Design	
Slot:	Date:	
Student Name:	Roll No.:	

Grading Rubric

	Tick the best applicable per row			
	Below	Lacking	Meets all	Points
	Expectation	in Some	Expectation	
Completeness of the report				
Organization of the report (5 pts)				
With cover sheet, answers are in the same				
order as questions in the lab, copies of the				
questions are included in report, prepared				
Quality of figures (5 pts)				
and name(s)				
Understanding and				
implementation of simple FIR filter				
(35 pts)				
Difference eq., flow diagram, impulse				
response, plots of magnitude response,				
plots of original and filtered signals and				
their DTFT, matlab code, questions				
Understanding and				
implementation of simple IIR filter				
(35 pts)				
Difference eq., flow diagram, impulse				
response, piors of magnitude response,				
their DTFT, matlab code, questions				
Understanding parameters of				
lowpass filter design (20 pts)				
Magnitude response plots with marked				
regions, questions.				
TOTAL (100 pts)				

Total Points (100):

TA Initials:

1. Introduction

The first week of the laboratory covers some basic examples of FIR and IIR filters, and then introduces the concepts of FIR filter design. Then the second week covers systematic methods of designing both FIR and IIR filters.

2. Background

In digital signal processing applications, it is often necessary to change the relative amplitudes of frequency components or remove undesired frequencies of a signal. This process is called filtering. Digital filters are used in a variety of applications. In Laboratory 4, we saw that digital filters may be used in systems that perform interpolation and decimation on discrete-time signals. Digital filters are also used in audio systems that allow the listener to adjust the bass (low-frequency energy) and the treble (high frequency energy) of audio signals.

Digital filter design requires the use of both frequency domain and time domain techniques. This is because filter design specifications are often given in the frequency domain, but filters are usually implemented in the time domain with a difference equation. Typically, frequency domain analysis is done using the Z-transform and the discrete-time Fourier Transform (DTFT).

In general, a linear and time-invariant causal digital filter with input x[n] and output y[n] may be specified by its difference equation

$$y[n] = \sum_{i=0}^{N-1} b_i x[n-i] - \sum_{k=1}^{M} a_k y[n-k]$$
(1)

where b_i and a_k are coefficients which parameterize the filter. This filter is said to have N zeros and M poles. Each new value of the output signal, y[n], is determined by past values of the output, and by present and past values of the input. The impulse response, h[n], is the response of the filter to an input of [n], and is therefore the solution to the recursive difference equation

$$h[n] = \sum_{i=0}^{N-1} b_i \delta[n-i] - \sum_{k=1}^{M} a_k h[n-k]$$
(2)

There are two general classes of digital filters: infinite impulse response (IIR) and finite impulse response (FIR). The FIR case occurs when $a_k = 0$, for all k. Such a filter is said to have no poles, only zeros. In this case, the difference equation (2) becomes

$$h[n] = \sum_{i=0}^{N-1} b_i \delta[n-i]$$
(3)

Since (3) is no longer recursive, the impulse response has finite duration N.

In the case where $a_k \neq 0$, the difference equation usually represents an IIR filter. In this case, (2) will usually generate an impulse response which has non-zero values as $n \to \infty$.

However, later we will see that for certain values of $a_k \neq 0$ and b_i , it is possible to generate an FIR filter response.

The Z-transform is the major tool used for analyzing the frequency response of filters and their difference equations. The Z-transform of a discrete-time signal, x[n], is given by

$$X(z) = \sum_{n=-\infty}^{\infty} x[n] z^{-n}$$

where z is a complex variable. The DTFT may be thought of as a special case of the Z-transform where z is evaluated on the unit circle in the complex plane.

$$X(e^{j\omega}) = X(z)|_{z=e^{j\omega}} = \sum_{n=-\infty}^{\infty} x[n]e^{-j\omega n}$$

From the definition of the Z-transform, a change of variable m = n - K shows that a delay of *K* samples in the time domain is equivalent to multiplication by z^{-K} in the Z-transform domain.

$$x[n-K] \stackrel{z}{\leftrightarrow} z^{-K}X(z)$$

We may use this fact to re-write Eq. (1) in the Z-transform domain, by taking Z-transforms of both sides of the equation:

$$Y(z) = \sum_{i=0}^{N-1} b_i z^{-i} X(z) - \sum_{k=1}^{M} a_k z^{-k} Y(z)$$
$$Y(z) \left(1 + \sum_{k=1}^{M} a_k z^{-k} \right) = X(z) \sum_{i=0}^{N-1} b_i z^{-i}$$
$$H(z) \triangleq \frac{Y(z)}{X(z)} = \frac{\sum_{i=0}^{N-1} b_i z^{-i}}{(1 + \sum_{k=1}^{M} a_k z^{-k})}$$

From this formula, we see that any filter which can be represented by a linear difference equation with constant coefficients has a rational transfer function (i.e. a transfer function which is a ratio of polynomials). From this result, we may compute the frequency response of the filter by evaluating H(z) on the unit circle:

$$H(e^{j\omega}) = \frac{\sum_{i=0}^{N-1} b_i e^{-j\omega i}}{(1 + \sum_{k=1}^{M} a_k e^{-j\omega k})}$$

There are many different methods for implementing a general recursive difference equation of the form (1). Depending on the application, some methods may be more robust to quantization error, require fewer multiplies or adds, or require less memory. Fig. 1 shows a system diagram known as the direct form implementation; it works for any discrete-time filter described by difference equation (1). Note that the boxes containing the symbol z^{-1} represent unit delays, while a parameter written next to a signal path represents multiplication by that parameter.



Figure 1: Direct form implementation for a discrete-time filter described by a general difference equation of the form in equation (1).

3. Design of a Simple FIR Filter

Download nspeech1.mat from https://engineering.purdue.edu/VISE/ee438L/lab5/data/nspeech1.zip

Download DTFT.m from

https://engineering.purdue.edu/VISE/ee438L/lab5/data/DTFT.zip



Figure 2: Location of two zeros for a simple FIR filter.

To illustrate the use of zeros in filter design, you will design a simple second order FIR filter with the two zeros on the unit circle as shown in Fig. 2. In order for the filter's impulse response to be real-valued, the two zeros must be complex conjugates of one another:

$$z_1 = e^{j\theta}$$
, $z_2 = e^{-j\theta}$

where θ is the angle of z_1 relative to the positive real axis. We will see later that $\theta \in [0, \pi]$ may be interpreted as the location of the zeros in the frequency response.

The transfer function for this filter is given by

$$H_f(z) = (1 - z_1 z^{-1})(1 - z_2 z^{-1}) = 1 - 2\cos\theta \, z^{-1} + z^{-2}$$

Use this transfer function to determine the difference equation for this filter. Then draw the corresponding system diagram and compute the filter's impulse response h[n].

This filter is an FIR filter because it has impulse response h[n] of finite duration. Any filter with only zeros and no poles other than those at 0 and $\pm \infty$ is an FIR filter. Zeros in the transfer function represent frequencies that are not passed through the filter. This can be useful for removing unwanted frequencies in a signal. The fact that $H_f(z)$ has zeros at $e^{\pm j\theta}$ implies that $H_f(e^{\pm j\theta}) = 0$. This means that the filter will not pass pure sine waves at a frequency of $\omega = \theta$.

Use Matlab to compute and plot the magnitude of the filter's frequency response $|H_f(e^{j\omega})|$ as a function of ω on the interval $-\pi < \omega < \pi$, for the following three values of θ :

i)
$$\theta = \pi/6$$

ii)
$$\theta = \pi/3$$

iii)
$$\theta = \pi/2$$

Put all three plots on the same figure using the **subplot** command.

Lab Report:

Submit the difference equation, system diagram, and the analytical expression of the impulse response for the filter $H_f(z)$. Also submit the plot of the magnitude response for the three values of θ . Explain how the value of θ affects the magnitude of the filter's frequency response.

In the next experiment, we will use the filter $H_f(z)$ to remove an undesirable sinusoidal interference from a speech signal. To run the experiment, first download the audio signal *nspeech1.mat*, and the M-file DTFT.m Load *nspeech1.mat* into Matlab using the command load nspeech1. This will load the signal *nspeech1* into the workspace. Play *nspeech1* using the sound command, and then plot 101 samples of the signal for the time indices (100:200).

We will next use the DTFT command to compute samples of the DTFT of the audio signal. The DTFT command has the syntax [X,w]=DTFT(X,M), where x is a signal which is assumed to start at time n = 0, and M specifies the number of output points of the DTFT. The command [X,w]=DTFT(x,0) will generate a DTFT that is the same duration as the input; if this is not sufficient, it may be increased by specifying M. The outputs of the function are a vector X containing the samples of the DTFT, and a vector w containing the corresponding frequencies of these samples.

Compute the magnitude of the DTFT of 1001 samples of the audio signal for the time indices (100:1100). Plot the magnitude of the DTFT samples versus frequency for $|\omega| < \pi$. Notice that there are two large peaks corresponding to the sinusoidal interference signal. Use the Matlab max command to determine the exact frequency of the peaks. This will be the value of θ that we will use for filtering with $H_f(z)$. Hint: Use the command

[Xmax,Imax]=max(abs(X)) to find the value and index of the maximum element in X. θ can be derived using this index.

Write a Matlab function FIRfilter(x) that implements the filter $H_f(z)$ with the measured value of θ and outputs the filtered signal (Hint: Use convolution). Apply the new function FIRfilter to the nspeech1 vector to attenuate the sinusoidal interference. Listen to the filtered signal to hear the effects of the filter. Plot 101 samples of the signal for the time indices (100:200), and plot the magnitude of the DTFT of 1001 samples of the filtered signal for the time indices (100:1100).

Lab Report:

For both the original audio signal and the filtered output, hand in the following:

- The time domain plot of 101 samples.
- The plot of the magnitude of the DTFT for 1001 samples.

Also hand in the code for the FIRfilter filtering function. Comment on how the frequency content of the signal changed after filtering. Is the filter we used a lowpass, highpass, bandpass, or a bandstop filter? Comment on how the filtering changed the quality of the audio signal.

4. Design of a Simple IIR Filter

Download pcm.mat https://engineering.purdue.edu/VISE/ee438L/lab5/data/pcm.zip

While zeros attenuate a filtered signal, poles amplify signals that are near their frequency. In this section, we will illustrate how poles can be used to separate a narrowband signal from adjacent noise. Such filters are commonly used to separate modulated signals from background noise in applications such as radio-frequency demodulation.

Consider the following transfer function for a second order IIR filter with complexconjugate poles:



Figure 3: Location of two poles for a simple IIR filter.

Figure 3 shows the locations of the two poles of this filter. The poles have the form

$$p_1 = re^{j\theta}$$
, $p_2 = re^{-j\theta}$

where r is the distance from the origin, and θ is the angle of p_1 relative to the positive real axis. From the theory of Z-transforms, we know that a causal filter is stable if and only if its poles are located within the unit circle. This implies that this filter is stable if and only if |r| < 1. However, we will see that by locating the poles close to the unit circle, the filter's bandwidth may be made extremely narrow around θ .

This two-pole system is an example of an IIR filter because its impulse response has infinite duration. Any filter with nontrivial poles (not at z = 0 or $\pm \infty$) is an IIR filter unless the poles are cancelled by zeros.

Calculate the magnitude of the filter's frequency response $|H_i(e^{j\omega})|$ on $|\omega| < \pi$ for $\theta = \pi/3$ and the following three values of r.

- *r* = 0.99
- *r* = 0.9
- *r* = 0.7

Put all three plots on the same figure using the subplot command.

Lab Report:

Submit the difference equation, system diagram and the analytical expression of the impulse response for $H_i(z)$. (Hint: The frequency response of the system can be obtained by restricting the z-transform to the unit circle. So the DTFT of $h_i[n]$ is $H_i(e^{j\omega})$. Therefore, to get $h_i[n]$, you can take the inverse Fourier transform of $H_i(e^{j\omega})$.) Also submit the plot of the magnitude of the frequency response for each value of r. Explain how the value of r affects this magnitude.

In the following experiment, we will use the filter $H_i(z)$ to separate a modulated sinusoid from background noise. To run the experiment, first download the file pcm.mat and load it into the Matlab workspace using the command load pcm. Play pcm using the sound command. Plot 101 samples of the signal for indices (100:200), and then compute the magnitude of the DTFT of 1001 samples of pcm using the time indices (100:1100). Plot the magnitude of the DTFT samples versus radial frequency for $|\omega| < \pi$. The two peaks in the spectrum correspond to the center frequency of the modulated signal. The low amplitude wideband content is the background noise. In this exercise, you will use the IIR filter described above to amplify the desired signal, relative to the background noise.

The pcm signal is modulated at 3146Hz and sampled at 8kHz. Use these values to calculate the value of θ for the filter $H_i(z)$. Remember from the sampling theorem that a radial frequency of 2π corresponds to the sampling frequency.

Write a Matlab function IIRfilter(x) that implements the filter $H_i(z)$. In this case, you need to use a for loop to implement the recursive difference equation. Use your calculated value of θ and r = 0.995. You can assume that y(n) is equal to 0 for negative values of n. Apply the new command IIRfilter to the signal pcm to separate the desired signal from the background noise, and listen to the filtered signal to hear the effects. Plot the filtered signal for indices (100:200), and then compute the DTFT of 1001 samples of the filtered signal using the time indices (100:1100). Plot the magnitude of this DTFT. In order to see the DTFT around $\omega = \theta$ more clearly, plot also the portion of this DTFT for the values of ω in the range[$\theta - 0.02, \theta + 0.02$]. Use your calculated value of θ .

Lab Report:

For both the pcm signal and the filtered output, submit the following:

- The time domain plot of the signal for 101 points.
- The plot of the magnitude of the DTFT computed from 1001 samples of the signal.
- The plot of the magnitude of the DTFT for ω in the range $[\theta 0.02, \theta + 0.02]$.

5. Lowpass Filter Design Parameters

Download nspeech2.mat

https://engineering.purdue.edu/VISE/ee438L/lab5/data/nspeech2.zip

Oftentimes it is necessary to design a good approximation to an ideal lowpass, highpass or bandpass filter. Figure 4 illustrates the typical characteristics of a real low-pass filter. The frequencies $|\omega| < \omega_p$ are known as the passband, and the frequencies $\omega_s < |\omega| \le \pi$ are the stopband. For any real filter, $\omega_p < \omega_s$. The range of frequencies $\omega_p \le \omega \le \omega_s$ is known as the transition band. The magnitude of the filter response, $H(e^{j\omega})$, is constrained in the passband and stopband by the following two equations

 $|H(e^{j\omega}) - 1| \le \delta_p \text{ for } |\omega| < \omega_p$

$$|H(e^{j\omega})| \leq \delta_s$$
 for $\omega_s < |\omega| \leq \pi$

where δ_p and δ_s are known as the passband and stopband ripple respectively. Most lowpass filter design techniques depend on the specification of these four parameters: ω_p , ω_s , δ_p and δ_s .



Figure 5: DTFT of a section of noisy speech.

To illustrate the selection of these parameters consider the problem of filtering out background noise from a speech signal. Figure 5 shows the magnitude of the DTFT over a window of such a signal, called nspeech2. Notice that there are two main components in nspeech2: one at the low frequencies and one at the high. The high frequency signal is noise, and it is band limited to $|\omega| > 2.2$. The low frequency signal is speech and it is band limited to $|\omega| < 1.8$. Download the file nspeech2.mat and load it into the Matlab workspace. It contains the signal nspeech2 from Fig. 5. Play the nspeech2 using the sound command and note the quality of the speech and background noise.

In the following sections, we will compute low-pass filters for separating the speech and noise using a number of different methods.

5.1. Filter Design Using Truncation

Ideally, a low-pass filter with cutoff frequency ω_c should have a frequency response of

$$H_{ideal}(e^{j\omega}) = \begin{cases} 1 & |\omega| \le \omega_c \\ 0 & \omega_c < |\omega| \le \pi \end{cases}$$

and a corresponding impulse response of

$$h_{ideal}(n) = \frac{\omega_c}{\pi} \operatorname{sinc}\left(\frac{\omega_c n}{\pi}\right) \text{ for } -\infty < n < \infty$$
(4)

However, no real filter can have this frequency response because $h_{ideal}(n)$ is infinite in duration.

One method for creating a realizable approximation to an ideal filter is to truncate this impulse response outside of $n \in [-M, M]$.

$$h_{trunc}(n) = \begin{cases} \frac{\omega_c}{\pi} \operatorname{sinc}\left(\frac{\omega_c}{\pi}n\right) & n = -M, \dots, -1, 0, 1, \dots, M\\ 0 & otherwise \end{cases}$$



Figure 6: Frequency response of low-pass filter designed using the truncation method.

Figure 6 shows the magnitude response of the lowpass filter with cutoff frequency $\omega_c = 2.0$, with the impulse response truncated to $n \in [-10,10]$. Notice the oscillatory behaviour of the magnitude response near the cutoff frequency and the large amount of ripple in the stopband.

Due to the modulation property of the DTFT, the frequency response of the truncated filter is the result of convolving the magnitude response of the ideal filter (a rect) with the DTFT of the truncating window. The DTFT of the truncating window, shown in Fig.

7, is similar to a sinc function since it is the DTFT of a sampled rectangular window. Notice that this DTFT has very large sidelobes, which lead to large stopband ripple in the final filter.



Figure 7: DTFT of a rectangular window of size 21.

A truncated impulse response is of finite duration, yet the filter is still noncausal. In order to make the FIR filter causal, it must be shifted to the right by M units. For a filter of size N = 2M + 1 this shifted and truncated filter is given by

$$h(n) = \begin{cases} \frac{\omega_c}{\pi} \operatorname{sinc}\left(\frac{\omega_c}{\pi}\left(n - \frac{N-1}{2}\right)\right) & n = 0, 1, \dots, N-1 \\ 0 & otherwise \end{cases}$$
(5)

This time shift of (N - 1)/2 units to the right corresponds to multiplying the frequency response by $e^{-j\omega(N-1)/2}$. It does not affect the magnitude response of the filter, but adds a factor of $-j\omega(N-1)/2$ to the phase response. Such a filter is called linear phase because the phase is a linear function of ω .

It is interesting to see that the filter formula of (5) is valid for N both even and odd. While both of these filters are linear phase, they have different characteristics in the time domain. When N is odd, then the value at n = (N - 1)/2 is sampled at the peak of the sinc function, but when N is even, then the two values at n = N/2 and n = (N/2) - 1 straddle the peak.

To examine the effect of filter size on the frequency characteristics of the filter, write a Matlab function LPFtrunc(N) that computes the truncated and shifted impulse response of size N for a low pass filter with a cutoff frequency of $\omega_c = 2.0$. For each of the following filter sizes, plot the magnitude of the filter's DTFT in decibels. Hints: The magnitude of the response in decibels is given by $|H_{dB}(e^{j\omega})| = 20 \log_{10} |H(e^{j\omega})|$. Note that the log command in Matlab computes the natural logarithm. Therefore, use the log10 command to compute decibels. To get an accurate representation of the DTFT make sure compute least sample that you at 512points using the command [X,w]=DTFT(filter_response,512).

- *N* = 21
- N = 101

Now download the noisy speech signal nspeech2.mat, and load it into the Matlab workspace. Apply the two filters with the above sizes to this signal. Since these are FIR filters, you can simply convolve them with the audio signal. Listen carefully to the unfiltered and filtered signals, and note the result. Can you hear a difference between the two filtered signals? In order to hear the filtered signals better, you may want to multiply each of them by 2 or 3 before using sound.

Lab Report:

- Submit the plots of the magnitude response for the two filters (not in decibels). On each of the plots, mark the passband, the transition band and the stopband.
- Submit the plots of the magnitude response in decibels for the two filters.
- Explain how the filter size effects the stopband ripple. Why does it have this effect?
- Comment on the quality of the filtered signals. Does the filter size have a noticeable effect on the audio quality?

Acknowledgement: This manual is based on "Purdue University: ECE438 - Digital Signal Processing with Applications (2016)" by Prof. Charles Bouman and Prof. Mireille Boutin.