

**ECE 2025 Spring 2005**  
**Lab #8: Filter Design: Bandpass Filters**

Date: 9–15 March 2005

---

**You should read the Pre-Lab section of the lab and do all the exercises in the Pre-Lab section before your assigned lab time.**

The Warm-up section of each lab must be completed **during your assigned Lab time** and the steps marked *Instructor Verification* must also be signed off **during the lab time**. After completing the warm-up section, turn in the verification sheet to your TA.

*Forgeries and plagiarism are a violation of the honor code and will be referred to the Dean of Students for disciplinary action. You are allowed to discuss lab exercises with other students and you are allowed to consult old lab reports, but you cannot give or receive written material or electronic files. Your submitted work should be original and it should be your own work.*

It is only necessary to turn in Section 4 as this week’s lab report; the lab report format is **Informal**. The report will be **due on either 16, 17-March or 28, 29-March (after Spring Break)**.

---

## 1 Introduction

The goal of this lab is to study the response of FIR filters to inputs such as complex exponentials and sinusoids. In the experiments of this lab, you will use `firfilt()`, or `conv()`, to implement filters and `freqz()` to obtain the filter’s frequency response.<sup>1</sup> As a result, you should learn how to characterize a filter by knowing how it reacts to different frequency components in the input.

## 2 Pre-Lab

This lab also introduces the process of “filter design” and then uses that skill to design practical filters: *lowpass*, *highpass* and *bandpass* filters. Bandpass filters can be used to detect and extract information from sinusoidal signals, e.g., individual notes in a musical passage or tones in a touch-tone telephone dialer.

### 2.1 Frequency Response of FIR Filters

The output or *response* of a filter for a complex sinusoid input,  $e^{j\hat{\omega}n}$ , depends on the frequency,  $\hat{\omega}$ . Often a filter is described solely by how it affects different input frequencies—this is called the *frequency response*. The frequency response of a general FIR linear time-invariant system is

$$H(e^{j\hat{\omega}}) = \sum_{k=0}^M b_k e^{-j\hat{\omega}k} \quad (1)$$

For the process of filter design, it is important to recognize that the choice of filter coefficients  $\{b_k\}$  determines the frequency response. The most useful filter design methods provide a formula for the  $\{b_k\}$  coefficients that will give a particular behavior for  $H(e^{j\hat{\omega}})$ , e.g., the formula in (2).

---

<sup>1</sup>If you are working at home and do not have the function `freqz.m`, there is a substitute available called `freqz.m`. You can find it in the *SP-First Toolbox*, or get it from the ECE-2025 WebCT page.

### 2.1.1 MATLAB Frequency Response for Lowpass Filter

MATLAB has a built-in function for computing the frequency response of a discrete-time LTI system. When the filter coefficients are given by the formula:

$$h[n] = \frac{\sin(\hat{\omega}_c(n - M/2))}{\pi(n - M/2)} \frac{\sin(2\pi(n - M/2)/M)}{2\pi(n - M/2)/M} \quad \text{for } n = 0, 1, 2, \dots, M \quad (2)$$

where  $M$  is the filter order, which should be an **even integer**. The first term in  $h[n]$  is a *sinc function*; second term in  $h[n]$  is called a Lanczos window.<sup>2</sup> The design parameter  $\hat{\omega}_c$  is called the *cutoff frequency* of the filter because it determines the passband and stopband regions of the frequency response (below). The following MATLAB statements show how to use `freqz` to compute and plot both the magnitude (absolute value) and the phase of the frequency response of the filter in (2) (when  $\hat{\omega}_c = 0.32\pi$ ) as a function of  $\hat{\omega}$  in the range  $-\pi \leq \hat{\omega} \leq \pi$ :

```
M = 50;    nn = 0:M;
LW = sin(2*pi/M*(nn-M/2)) ./ (2*pi/M*(nn-M/2));    %-- Lanczos window
LW(M/2 + 1) = 1;                                     %-- fix divide by zero
wc = 0.32*pi;
sincwc = sin(wc*(nn-M/2)) ./ (pi*(nn-M/2));    %-- sinc function
sincwc(M/2 + 1) = wc/pi;                             %-- fix divide by zero
bb = sincwc.*LW;                                       %-- Filter Coefficients
ww = -pi:(pi/400):pi;                                  %-- omega hat
HH = freqz(bb, 1, ww);                                %<--freakz.m is an alternative
subplot(2,1,1);
plot(ww, abs(HH)), zoom on, grid on
subplot(2,1,2);
plot(ww, angle(HH)), zoom on, grid on, shg
xlabel('Normalized Radian Frequency')
```

For FIR filters, the second argument of `freqz( -, 1, - )` must always be equal to 1.<sup>3</sup> The frequency vector `ww` should cover an interval of length  $2\pi$  for  $\hat{\omega}$ , and its spacing must be fine enough to give a smooth curve for  $H(e^{j\hat{\omega}})$ .

## 2.2 Passband Defined for the Frequency Response

Certain types of digital filters have a frequency response (magnitude) that is close to one in some frequency regions, and close to zero in others. For example, the plot in Fig. 1 is a lowpass filter whose magnitude is close to one when the frequency  $\hat{\omega}$  is near zero. This region is called the *passband* of the filter. It will be useful to have a precise measurement of the passband width so that we can compare different filters.

- From the plot of the magnitude response in Section 2.1.1 determine the set of frequencies where the magnitude is very close to one, as defined by  $||H(e^{j\hat{\omega}})| - 1|$  being less than 0.0125. This should be a region of the form  $-\hat{\omega}_p \leq \hat{\omega} \leq \hat{\omega}_p$ . Determine  $\hat{\omega}_p$  for the case where  $M = 50$ .
- The parameter  $\hat{\omega}_p$  is called the *passband edge*. Compare the value of  $\hat{\omega}_p$  found in the previous part to the design parameter  $\hat{\omega}_c$  in (2).

## 2.3 Stopband Defined for the Frequency Response

When the frequency response (magnitude) of the digital filter is close to zero, we have the stopband region of the filter. In the lowpass filter example of Section 2.1.1, the magnitude is close to zero when the frequency

<sup>2</sup>The Lanczos window is formed by sampling the mainlobe of a sinc function.

<sup>3</sup>If the output of the `freqz` function is not assigned, then plots are generated automatically; however, the magnitude is given in decibels which is a logarithmic scale. For linear magnitude plots a separate call to `plot` is necessary as in the code above.

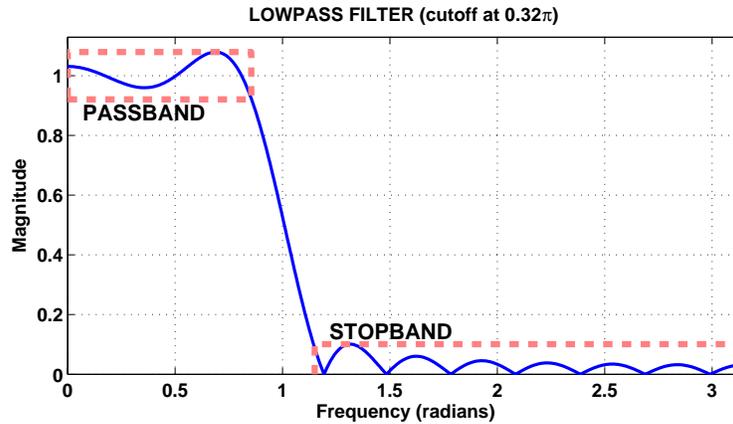


Figure 1: Passband and stopband defined for a typical lowpass filter. This particular filter is a length-21 FIR filter designed with rectangular window and a sinc function with a cutoff frequency of  $\hat{\omega}_c = 0.32\pi$ . The approximate value of the passband edge is  $\hat{\omega}_p = 0.27\pi = 0.85$ ; the stopband edge,  $\hat{\omega}_s = 0.366\pi = 1.15$ .

$\hat{\omega}$  is near  $\pi$  (a high frequency). This region is called the *stopband* of the filter. We can make a precise measurement of the stopband edge as follows:

- From the plot of the magnitude response in Section 2.1.1 determine the set of frequencies where the magnitude is nearly zero, as defined by  $|H(e^{j\hat{\omega}})|$  being less than 0.0125. This should be two regions:  $\hat{\omega}_s \leq \hat{\omega} \leq \pi$  in positive frequencies, and  $-\pi \leq \hat{\omega} \leq -\hat{\omega}_s$  for negative frequencies. Determine  $\hat{\omega}_s$  for the case where  $M = 50$ .
- The parameter  $\hat{\omega}_s$  is called the *stopband edge*. Compare the value of  $\hat{\omega}_s$  found in the previous part to the design parameter  $\hat{\omega}_c$  in (2).

## 2.4 Linear Phase in the Frequency Response

The phase of the frequency response can be related to time delay. In the lowpass filter example of Section 2.1.1, the phase plot appears to be jagged, but it is actually linear.

- Determine the slope of the linear segments of the frequency response for the  $M = 50$  filter in Section 2.1.1. Your answer should be an integer.
- Plot the impulse response of the digital filter in Section 2.1.1 for the range  $n = 0, 1, \dots, M$ . Then determine the symmetry point for  $h[n]$ , i.e., find the integer  $n_s$  such that  $h[n_s + n] = h[n_s - n]$ .
- Compare the value of the slope (from part (a)) to the “symmetry point” of  $h[n]$ . Verify that the phase slope is equal to  $-n_s$ .

## 2.5 Use the MATLAB FIND Function

In MATLAB the `find` function returns a list of indices where a logical condition is true. See `help relop` for information. For example, if you have the frequency response from Section 2.1.1 you could use the `find` command to determine the indices where `abs(HH)` is close to one, and then use those indices to display the list of frequencies in the passband.

### 3 Warm-up

The objective of this warm-up is to use the MATLAB GUI `dltidemo`<sup>4</sup>. A second objective is to demonstrate the frequency responses of lowpass, highpass and bandpass filters designed by the *Lanczos-window* method.

#### 3.1 LTI Frequency Response Demo

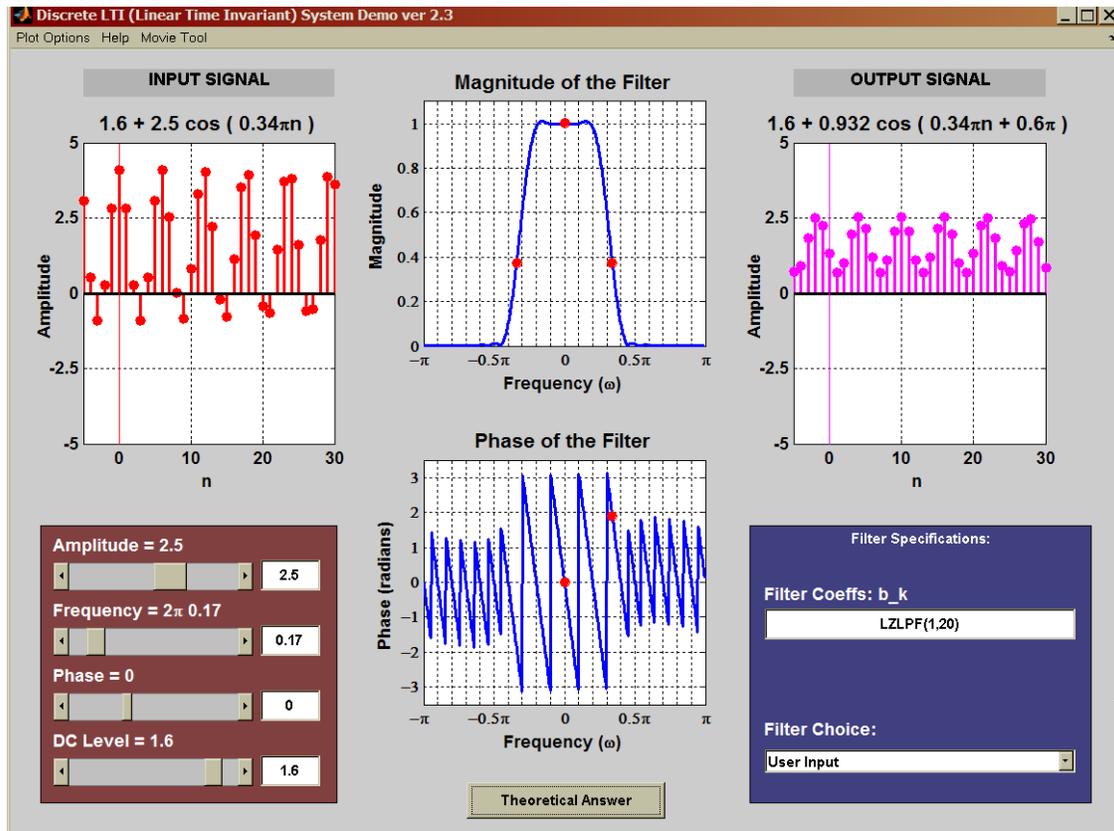


Figure 2: DLTI demo interface. The frequency label is  $\omega$  because MATLAB won't display  $\hat{\omega}$ . When the Filter Choice is set to `User Input`, MATLAB code can be entered in the text box for the filter coefficients.

The `dltidemo` GUI illustrates the “sinusoid-IN gives sinusoid-OUT” property of LTI systems. In this demo, you can change the amplitude, phase and frequency of an input sinusoid,  $x[n]$ , and you can change the digital filter that processes the signal. Then the GUI will show the output signal,  $y[n]$ , which is also a sinusoid (*at the same frequency*). Figure 2 shows the interface for the `dltidemo` GUI. It is possible to see the formula for the output signal, if you click on the `Theoretical Answer` button located at the bottom-middle part of the window. The digital filter can be changed by choosing different options in the `Filter Specifications` box in the lower right-hand corner. Right-clicking on the red dots in the frequency response panels (middle) will give the numerical value of the magnitude or phase at that point; likewise, right-clicking on a signal point in the left or right panels will give the numerical value of the signal at one index.

In the Warm-up, you should perform the following steps with the `dltidemo` GUI:

- (a) Set the input to  $x[n] = 1.5 \cos(0.1\pi(n - 4))$
- (b) Set the digital filter to be a 9-point averager.

<sup>4</sup>The `dltidemo` GUI is part of the *SP-First* toolbox, and is already installed in the ECE lab. The latest versions of all the *SP-First* GUIs can also be found at: <http://users.ece.gatech.edu/mcclella/matlabGUIs/index.html>

- (c) Determine the formula for the output signal and write it in the form:  $y[n] = A \cos(\hat{\omega}_0(n - n_d))$ .
- (d) Using  $n_d$  for  $y[n]$  and the fact that the input signal had a peak at  $n = 4$ , determine the amount of delay through the filter. In other words, how much has the peak of the cosine wave shifted?
- (e) Now, change the length of the averaging filter so that the output will be zero, i.e.,  $y[n] = 0$ . Use the GUI to show that you have the correct filter to zero the output. If the filter length is more than 15, you will have to enter the “Filter Specifications” with the  option.
- (f) When the output is zero, the filter acts as a *Nulling Filter*, because it eliminates the input at  $\hat{\omega} = 0.1\pi$ . Which other frequencies  $\hat{\omega}$  are also nulled? Find at least one.

**Instructor Verification** (separate page)

### 3.2 Lowpass, Bandpass and Highpass Filters

The **dltidemo** can be used to show different filter types: lowpass filters (LPF), bandpass filters (BPF) and highpass (HPF) filters. You should perform the following steps with the **dltidemo** GUI:

- (a) Set the input to  $x[n] = 2 + 1 \cos(0.9\pi n)$ , which is a signal consisting of one low frequency component and one high frequency component.
- (b) Set the digital filter to be the filter **Lowpass, L=15**. Observe the output and determine the exact mathematical formula for the output signal.
- (c) Set the digital filter to be the filter **Highpass, L=15**. Observe the output and determine the exact mathematical formula for the output signal.
- (d) Set the digital filter to be the filter **Bandpass, L=21**. Observe the output and determine the exact mathematical formula for the output signal.
- (e) Use the frequency response,  $H(e^{j\hat{\omega}})$ , shown in the **dltidemo** GUI to explain why the outputs are different in these three cases.

**Instructor Verification** (separate page)

### 3.3 MATLAB Function for LPF Design

In this section you must write an M-file that will design a lowpass filter according to the formula given in (2). The M-file should return the filter coefficients given the order ( $M$ ) and the frequency  $\hat{\omega}_c$  as inputs.

- (a) Use the following comments as a template for writing the M-file.

```
function hh = LZLPF( wc, M )
%LZLPF  LPF design function with Lanczos window
%
%  wc = design parameter giving the approximate location
%       of the passband and stopband edges
%  M = filter order
%  hh = impulse response of the FIR lowpass filter
```

- (b) Demonstrate that your M-file works by using it in the **dltidemo** GUI. There is a  option as one of the filter choices. You should be able to write a call to LZLPF in the text box of the user input. Make a length-37 FIR LPF whose cutoff frequency is approximately  $0.37\pi$ . Recall that the length of an FIR filter is  $L = M + 1$ . If you look carefully at  $h[n]$  for the “length-37” Lanczos LPF, what is the true length of the impulse response?

**Instructor Verification** (separate page)

## 4 Lab Exercises

Among the many uses of lowpass filters (LPFs) is their role as *anti-aliasing* filters when downsampling signals or images. One everyday use of anti-aliasing filters is the display of fonts on a computer monitor. “Anti-aliased fonts” solve the following problem: how to display text on a low-resolution device. Since high quality printing (as seen in books) requires the equivalent of about 1200 or 2400 dots per inch (dpi), while a computer screen generally has 72 dots per inch, significant downsampling is needed to render text on a computer monitor. In reality, the problem is not that severe because laser printers at 600 dpi can produce printed pages that are indistinguishable from typeset book pages.

From the theory of sampling and aliasing, we know that high-frequency signal components will be aliased if the Sampling Theorem is not obeyed. This brings up the question of how to describe displayed fonts as having high-frequency or low-frequency components. The purpose of this lab exercise is to show that LPFs used prior to downsampling are beneficial. The LPF removes high-frequency components prior to downsampling and, as a result, the signal being downsampled is forced to obey the sampling theorem by the “anti-aliasing LPF.”

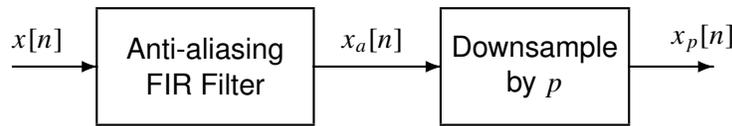


Figure 3: Cascade of anti-aliasing LPF and a downsampler.

The preceding discussion raises the following question: “if the signal  $x[n]$  is downsampled by a factor of 2, how do we guarantee that there will be no aliasing?” Furthermore, how do we generalize to the case where the downsampling is by a factor of  $p$ ? The answer comes from the following ad-hoc reasoning: Suppose that  $x[n]$  came from the sampling of the analog signal  $x(t)$  at a sampling rate of  $f_s$ . If we downsample  $x[n]$  by a factor of two, that is equivalent to sampling the analog signal  $x(t)$  at a rate of  $\frac{1}{2}f_s$ . At this lower sampling rate, the signal must have no frequency components above  $0.25f_s$ . Now, when  $x[n]$  was created by sampling at the rate of  $f_s$ , the frequency  $f = 0.25f_s$  mapped to  $\hat{\omega} = 2\pi f/f_s = 0.5\pi$ . Therefore, we conclude that downsampling  $x[n]$  by a factor of  $p = 2$  will not alias if the spectrum of  $x[n]$  has no frequency components in the region  $0.5\pi \leq |\hat{\omega}| \leq \pi$ . One way to enforce this condition is to apply a LPF to  $x[n]$ , and design the LPF to have its stopband cover the region  $0.5\pi \leq |\hat{\omega}| \leq \pi$ . On the other hand, its passband should cover the region  $|\hat{\omega}| < 0.5\pi$  because we want to preserve all the frequency components in the original  $x(t)$  below  $0.25f_s$ . A lowpass filter that performs this task is called an *anti-aliasing* filter.

### 4.1 Downsampling Music

Changing the sampling rate of an audio signal makes it easy to see how the LPF performs its anti-aliasing function. In this section, you must design an anti-aliasing filter and then downsample a music signal from 8000 Hz to 2000 Hz. The music signal was produced with the clarinet FM synthesis function, so the notes have high-frequency harmonics that might alias at lower sampling rates.

- Design a length-31 LPF that will perform the anti-aliasing. Explain how the cutoff frequency is chosen, and where stopband and passband edges actually end up.
- Load the signal from `ClarinetScale.wav`, and then change its sampling rate from 8000 Hz to 2000 Hz by first filtering with the anti-aliasing filter and then downsampling.
- Display three spectrograms: the original signal, the original downsampled (without an anti-alias LPF), and the signal from the previous part. Explain how these spectrograms confirm the useful function of the anti-aliasing filter.

## 4.2 Anti-Aliased Fonts

In this section, you must change the DPI (dots per inch) of a text image from 600 dpi to 100 dpi. The original image is too large to fit on the screen because its dimensions are  $4153 \times 3470$ , but when reduced by a factor of six the size is less than  $700 \times 600$ , which should fit on all computer monitors. In the test below, the effects of aliasing versus anti-aliasing can be seen visually.

- (a) Design a length-41 LPF that will perform the anti-aliasing. Explain how the cutoff frequency is chosen, and where stopband and passband edges actually end up.
- (b) Load the image from `text8.png` using `imread`, and then change its sampling rate from 600 dpi to 100 dpi by first filtering with the anti-aliasing filter and then downsampling. In the case of 2-D signals such as images, the FIR filter must be applied to all the rows and all the columns of the image.
- (c) Display two images: the original downsampled (without an anti-alias LPF), and the signal from the previous part. Explain how these images confirm the benefits of the anti-aliasing filter.

**Image Display:** The original image is black and white with numerical values of either 0 or 255. After the LPF is applied the numerical values of the image will actually be outside of the  $[0,255]$  range. In addition, the start-up and ending regions of the FIR Filter produce values that do not translate to reasonable gray-scale values. Since we want an image display where black is black, and white is white, the numerical values greater than 255 should be set equal to 255, and the values less than zero should be set equal to zero. This can be accomplished with `show_img` by using the “noscale” option.

- (d) Print three images for your lab report: the original, the original downsampled (without an anti-alias LPF), and the signal from the previous part. Assess the quality of the correctly downsampled image versus the original.
- (e) Determine whether changing the length of the FIR filter (greater than 41 or less than 41) would change the quality of the processed image very much.

## 4.3 Efficiency of 2-D Filtering (Optional)

This section is primarily for information. The original image is extremely large, so the time to run the FIR filter over all the rows and columns can be very long. Depending on how you call the `conv2` function, the time can vary quite a bit. To quantify the running time, compare the time to filter all the columns versus filtering all the rows.

```
bb = LZLPF(???);    %-- design the filter
tic, yy = conv2(bb', xx); toc    %-- xx holds the image
tic, yy = conv2(bb, xx); toc
```

Armed with this information, you can devise different strategies for doing the filtering efficiently. For example, you can filter the rows by doing three steps: transpose the image, filter the columns, and transpose it back. Which way is faster?

## Lab #8

ECE-2025

Spring-2005

### INSTRUCTOR VERIFICATION PAGE

*For each verification, be prepared to explain your answer and respond to other related questions that the lab TA's or professors might ask. Turn this page in at the end of your lab period.*

Name: \_\_\_\_\_

Date of Lab: \_\_\_\_\_

Part 3.1(d) and (f) Use the **dltidemo** GUI to illustrate the operation of a 9-point averaging filter. Determine the amount of delay through the filter and the frequency components that are nulled, and write your answers in the space below.

Verified: \_\_\_\_\_

Date/Time: \_\_\_\_\_

Part 3.2(b,c,d,e) Use the **dltidemo** GUI to demonstrate lowpass, highpass and bandpass filters. For each filter, write the expressions for the output  $y[n]$  in the space below.

$y_L[n] =$

$y_H[n] =$

$y_B[n] =$

Verified: \_\_\_\_\_

Date/Time: \_\_\_\_\_

Part 3.3(b) Demonstrate that your M-file for LPF design works correctly by using it in the **dltidemo** GUI to obtain a length-37 LPF.

Verified: \_\_\_\_\_

Date/Time: \_\_\_\_\_