
DSP First

Lab 10: Octave Band Filtering

¹Pre-Lab and Warm-Up: You should read at least the Pre-Lab and Warm-up sections of this lab assignment and go over all exercises in the Pre-Lab section before going to your assigned lab session.

Verification: 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**. One of the laboratory instructors must verify the appropriate steps by signing on the **Instructor Verification** line. When you have completed a step that requires verification, simply demonstrate the step to the TA or instructor. Turn in the completed verification sheet to your TA when you leave the lab.

Lab Report: It is only necessary to turn in a report on Section 4 with graphs and explanations. You are asked to **label** the axes of your plots and include a title for every plot. In order to keep track of plots, include your plot *inlined* within your report. If you are unsure about what is expected, ask the TA who will grade your report.

1 Introduction

This lab introduces a practical application where we attempt to extract information from sinusoidal signals—in this case, piano notes. Bandpass FIR filters can be used to extract the information encoded in the waveforms. The goal of this lab is to design and implement several bandpass FIR filters in MATLAB, and use the filtered outputs to determine automatically which note is being played. However, since there are 88 keys on the piano, we will only require the system to figure out which octave the note is in, not the exact note. In the experiments of this lab, you will use FIR filter, or convolution, VIs to implement filters and “freqz.vi” to obtain the filter’s frequency response. As a result, you should learn how to characterize a filter by knowing how it reacts to different frequency components in the input.

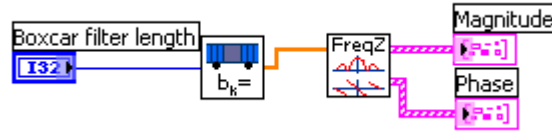
1.1 Frequency Response of FIR Filters

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 frequencies—this is called the frequency response. The frequency response of a general FIR linear time-invariant system is

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

The DSP First toolkit has a function for computing the frequency response of a discrete-time LTI system. The following LabVIEW code shows how to use freqz.vi to compute and plot the magnitude (absolute value) of the frequency response of an L -point averaging system² as a function of $\hat{\omega}$ in the range $-\pi \leq \hat{\omega} \leq \pi$:

² The filter length L is equal to $M + 1$.



Inside MathScript nodes, we will always use capital HH for the frequency response. When using the freqz VI, the frequency vector will cover the interval $-\pi \leq \hat{\omega} \leq \pi$ for $\hat{\omega}$, and its spacing should be fine enough to give a smooth curve for $H(e^{j\hat{\omega}})$.

2 Background and Pre-Lab

2.1 Piano Notes

A piano keyboard consists of 88 keys grouped into octaves. Each octave contains 12 notes, the notes in one octave being twice the frequency of the notes in the next lower octave.

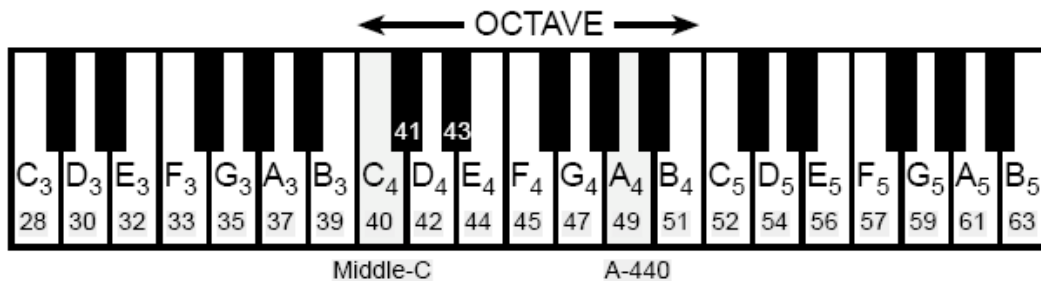


Figure 1: Layout of a piano keyboard. Key numbers are shaded. The notation C4 means the C-key in the fourth octave.

Frequencies of the notes are defined by setting the frequency of one note and referring all other frequencies to that note. The reference note is the A above middle-C, which is usually called A-440 (or A4) because its frequency is 440 Hz. Each octave contains 12 notes (5 black keys and 7 white) and the ratio between the frequencies of the notes is constant between successive notes. As a result, this ratio must be $2^{1/12}$. Since middle C is 9 keys below A-440, its frequency is approximately 261 Hz. Consult Chapter 3 for more details.

If we want to produce a system capable of writing music directly from a recorded signal $x(t)$, we need to analyze the frequency content of the signal. One way to do this analysis is to use a set of bandpass FIR filters, each one having its passband designed for one note on the keyboard. This would require a very large number of filters. Another way to do the analysis would be to use a two stage approach. First, we would use a set of bandpass FIR filters where each passband would pass the frequencies in one octave. Then these “octave filters” would be followed by more precise bandpass filters (BPFs) that would determine which key inside the octave is being played. The work in this lab will be to produce a working set of “octave filters.”

3.1 Create a Bandpass Filter (BPF)

There are many ways to get the filter coefficients for a bandpass filter. One easy choice is to define the impulse response of an L -point FIR filter as:

$$h[n] = \frac{2}{L} \cos(\hat{\omega}_c n), \quad 0 \leq n < L$$

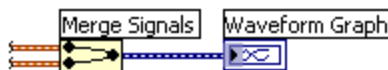
where L is the filter length, and $\hat{\omega}_c$ is the center frequency that defines the frequency location of the passband. For example, we would pick $\hat{\omega}_c = 0.2\pi$ if we want the peak of the filter's passband to be centered at 0.2π . The bandwidth of the bandpass filter is controlled by L ; the larger the value of L , the narrower the bandwidth. This particular filter is also discussed in the section on useful filters in Chapter 7.

- Generate the impulse response of a length-25 bandpass filter with $\hat{\omega}_c = 0.2\pi$, and plot $h[n]$ with a stem plot.
- Compute the frequency response of the length-25 BPF from the previous part and plot its magnitude and phase versus $\hat{\omega}$ over the range $-2\pi \leq \hat{\omega} \leq 2\pi$.
- Use the magnitude response plot to describe the passband of the BPF. For example, measure the bandwidth at a convenient point such as 50% of the peak.

Instructor Verification (separate page)

3.2 Overlay Plotting

Sometimes it is convenient to overlay multiple signals in one LabVIEW graph. The LabVIEW function "Merge Signals" will allow you to perform these sorts of overlays.



Test out this functionality by following these directions:

- Plot the magnitude response of the 25-point BPF filter defined in (2). For this plot, it is sufficient to use a horizontal frequency axis that extends from $-\pi$ to $+\pi$.
- Use the stem function to place vertical markers at several points on the frequency response at frequencies $\{0, \pm 0.2\pi, \pm 0.5\pi\}$.

```
hold on, stem(pi*[-0.5,-0.2,0,0.2,0.5],0.9*ones(1,5),'r. '), hold off
```

3.3 Signal Concatenation

In a previous lab, a very long music signal was created by joining together many sinusoids. When two signals were played one after the other, the composite signal was created by the operation of concatenation. In MathScript, this can be done by making each signal a row vector, and then using the matrix building notation as follows:

$$xx = [xx, xxnew];$$

where $xxnew$ is the sub-signal being appended. The length of the new signal is equal to the sum of the lengths of the two signals xx and $xxnew$. A third signal could be added later on by concatenating it to xx . In LabVIEW, we can use the “Append Signals” express VI to attach waveforms to each other. If we are only dealing with arrays, we can build an array using two preexisting arrays.

3.3.1 Comment on Efficiency

In LabVIEW, the efficiency of concatenation depends on the type of loop used. However, in general, when we append a waveform in a loop ($xx = [xx, xxnew]$), this becomes an inefficient procedure if the signal length gets to be very large. The reason is that LabVIEW must re-allocate the memory space for xx every time a new sub-signal is appended via concatenation. If the length xx were being extended from 400,000 to 401,000, then a clean section of memory consisting of 401,000 elements would have to be allocated followed by a copy of the existing 400,000 signal elements and finally the append would be done. This is clearly inefficient, but would not be noticed for short signals.

An alternative is to pre-allocate storage for the complete signal vector, but this can only be done if the final signal length is known ahead of time. Additionally, if a for loop is used with auto-indexing enabled, LabVIEW will do this automatically.

3.4 Filtering a Signal

Use the filter from a previous section (Sec. 3.1) to process an input signal composed of several sinusoids:

$$x[n] = \begin{cases} \cos(0.5\pi n) & 0 \leq n < 200 \\ 2 & 200 \leq n < 400 \\ 0.5 \cos(0.2\pi n) & 400 \leq n < 600 \end{cases}$$

- Generate $x[n]$.
- Filter $x[n]$.
- Make stem plots of $x[n]$ and the filtered signal.
- Use the frequency response to validate that the output signal has the correct magnitude and phase in each of the three regions where the input has different frequencies.
- Comment: observe the transient effect at the transitions ($n = 200$ and $n = 400$). This is due to the start-up of the FIR filter as it encounters a new sinusoid. How long does the transient last (in samples)?

4 Lab Exercises: Bandpass Filter Design

Section 4.1 should not be included in the lab report. It is provided as an extension of the Warm-up. Section 4.2 must be included.

4.1 Simple Bandpass Filter Design

The L -point averaging filter is a lowpass filter. Its passband width is controlled by L , being inversely proportional to L . It is also possible to create a filter whose passband is centered around some frequency other than zero. One simple way to do this is to define the impulse response of an L -point FIR as:

$$h[n] = \frac{2}{L} \cos(\hat{\omega}_c n), \quad 0 \leq n < L$$

where L is the filter length, and $\hat{\omega}_c$ is the center frequency that defines the frequency location of the center of the passband.

- (a) Generate a bandpass filter that will pass a frequency component at $\hat{\omega} = 0.4\pi$. Make the filter length (L) equal to 40. Make a plot of the frequency response magnitude and phase.
- (b) The passband of the BPF filter is defined by the region of the frequency response where $|H(e^{j\hat{\omega}})|$ is close to its maximum value of one. Typically, the passband width is defined as the length of the frequency region where $|H(e^{j\hat{\omega}})|$ is greater than some level such as 0.5, or $0.707 = 1/\sqrt{2}$. Note: you can use the find function or “search all in 1D array.vi” to locate those frequencies where the magnitude satisfies $|H(e^{j\hat{\omega}})| \geq 0.5$ (similar to Fig. 2). Use the plot of the frequency response for the length-40 bandpass filter from part (a), and determine the passband width using the 0.5 level to define the pass band.
- (c) Make two other plots of BPFs for $L = 20$ and $L = 80$ with the same $\hat{\omega}_c$; and measure the passband width for both. Then explain how the width of the passband is related to filter length L , i.e., what happens when L is doubled or halved.

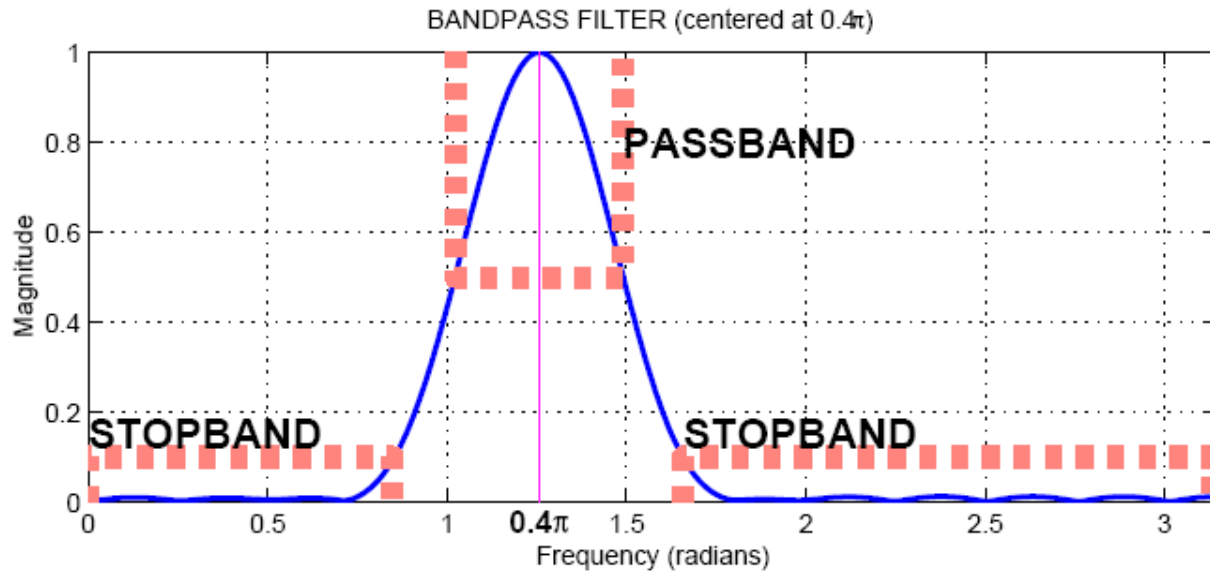


Figure 2: The frequency response of an FIR bandpass is shown with its passband and stopband regions. In this case, the passband is defined to be the region where $|H(e^{j\hat{\omega}})|$ is greater than 0.5; and the stopband is the region where the magnitude is less than 0.1.

4.2 A Better BPF

It is possible to get better performance in the frequency response by modifying the filter coefficients slightly. One easy way is to use a “Hamming window.” In this case, the impulse response for a length- L bandpass filter would be given as:

$$h[n] = (0.54 - 0.46 \cos(2\pi n / (L-1))) \cos(\hat{\omega}_c (n - (L-1)/2)) \quad \text{for } n = 0, 1, 2, \dots, L-1$$

where $\hat{\omega}_c$ is the desired center frequency for the BPF. The first term is the Hamming window. As before, the filter length L determines the passband width, although the Hamming BPF tends to be a little wider than the BPF in the previous section. The big advantage of the Hamming BPF is that its stopband has ripples that are very small (usually less than 0.01). Use zooming to figure out the height of the ripples in the stopband when you plot the frequency response (below).

- (a) Generate a Hamming bandpass filter that will pass a frequency component at $\hat{\omega} = 0.25\pi$. Make the filter length (L) equal to 41. Make a plot of the frequency response magnitude and phase. Measure the response of the filter (magnitude and phase) at the following frequencies of interest:

$$\hat{\omega} = \{0, 0.1\pi, 0.25\pi, 0.4\pi, 0.5\pi, 0.75\pi\}. \text{ Summarize the values in a table.}$$

Hint: use the freqz VI to calculate these values, or the find functionality to extract this information from the vector that produced the plot.

- (b) The passband of the BPF filter is defined by the region of the frequency response where $|H(e^{j\hat{\omega}})|$ is close to its maximum value of one. In this case, the passband width is defined as the length of the frequency region where $|H(e^{j\hat{\omega}})|$ is greater than 50% of the peak magnitude value. The stopband of the BPF filter is defined by the region of the frequency response where $|H(e^{j\hat{\omega}})|$ is close to zero.

Use the plot of the frequency response for the length-41 bandpass filter from part (a), and determine the passband width using the 50% level to define the pass band. Make two other plots of BPFs for $L = 21$ and $L = 81$, and measure the passband width in both. Then explain how the width of the passband is related to filter length L , i.e., what happens when L is (approximately) doubled or halved.

- (c) If the input signal to the length-41 FIR BPF is:

$$x[n] = 2 + 2 \cos(0.1\pi n + \pi/3) + \cos(0.25\pi n - \pi/3)$$

Determine (by hand) the formula for the output signal. (Hint: use the magnitude and phase measurements from part (a).) Comment on the relative amplitudes of the three signal components in the output signal, observing whether or not they were in the passband or stopband of the filter.

- (d) Use the frequency response (and passband width) of the length-41 bandpass filter to explain how the filter is able to pass the components at $\hat{\omega} = \pm 0.25\pi$, while reducing or rejecting others.

5 Lab Exercises: Piano Note Decoding

It is possible to decode piano signals into octaves using a simple FIR filter bank. The filter bank in Fig. 3 consists of seven bandpass filters which each pass only one of the seven possible octaves. The input signal for all the filters is the same.

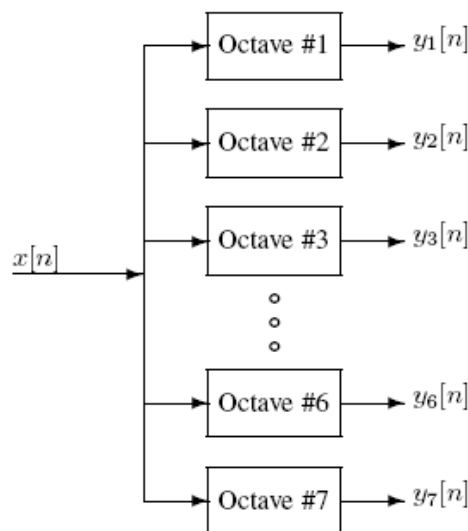


Figure 3: Filter bank consisting of bandpass filters which pass frequencies corresponding to the seven octaves on a piano. (Each octave contains 12 keys, but we will ignore the four extra keys outside the seven complete octaves.)

The octave filtering system needs two pieces: a set of bandpass filters (BPFs) to isolate individual frequency bands (i.e., the octaves), and a detector to determine whether or not a given octave is active. The detector must “score” each BPF output and determine whether or not the octave contains one or more notes. In a practical system where noise and interference are also present, this scoring process is a crucial part of the system design, but we will only work with noise-free signals to illustrate the basic functionality of the system.

To make the whole system work, you will have to write at least three VIs:

1. A filter design function to produce the filter coefficients for the BPFs in Fig. 3.
2. A filtering function that will filter the input signal through each channel.
3. A scoring function that will evaluate the output of each channel for “activity.”

The following sections discuss how to create or complete these functions.

5.1 Piano Octaves

The bandpass filter specs are determined by the frequencies of different octaves on the piano. The standard notation is to start an octave at the “C” key. For example, octave #4 starts at “middle-C” which is key #40 and extends up to key #51; the fifth octave starts at key #52 and extends to key #63, and so on. See Fig. 1 for the layout of a piano keyboard.

In this implementation, we want to design filters that will isolate five different octaves: octaves #2 through #6. Octave #2 starts with key #16, octave #3 starts at key #28, and octave #6 starts at key #64 (the last key in octave #6 is key #75).

In order to design the BPFs, we need the upper and lower frequencies for each octave in $\hat{\omega}$. We can convert key number to frequency in hertz, and then use the sampling frequency to convert analog frequency to $\hat{\omega}$. For this lab, assume that the sampling frequency is $f_s = 8000$ Hz. Make a table defining the five BPFs that you must design.

(a) For each octave, determine the lower and upper frequencies which will become the lower and upper edges of the passband of the BPF. Give these numbers in hertz and also in $\hat{\omega}$.

(b) Then compute the center frequencies of the BPFs as the midpoint between the lower and upper band edges of the filter. Save this information for use in the Filter Design Section (next).

5.2 Bandpass Filter Bank Design

The FIR filters that will be used in the filter bank (Fig. 3) should be constructed according to the Hamming impulse responses of (3). These “Hamming” BPFs require two parameters: the length L and the center frequency, f_c . Use the previous table of piano octave frequencies to define the passbands.

Furthermore, the filters should be scaled so that their maximum magnitude at the center frequency of the

passband is equal to one. This can be done by introducing a third parameter β that will scale all the filter coefficients.

$$h[n] = \beta(0.54 - 0.46 \cos(2\pi n / (L-1))) \cos(\hat{\omega}_c (n - (L-1)/2)) \quad \text{for } n = 0, 1, 2, \dots, L-1$$

The constant β gives flexibility for scaling the filter's gain to meet a constraint such as making the maximum value of the frequency response equal to one. The bandwidth of the bandpass filter is controlled by L ; the larger the value of L , the narrower the bandwidth.

1. Devise a strategy for picking the constant β so that the maximum value of the frequency response will be equal to one. Write the method that will do this scaling operation in general. There are two approaches here:
 - a. *Mathematical:* derive a formula for β from the formula for the frequency response of the BPF. Then use LabVIEW to evaluate this closed-form expression for β . In this case, such analysis is nearly impossible.
 - b. *Numerical:* let LabVIEW measure the peak value of the unscaled frequency response, and then compute β to scale the peak to be one.
- (b) You must design five separate bandpass filters for the filter bank system. Each filter has a different length and a different center frequency. The lengths have to be different because the bandwidths of the octaves are different. In fact, the bandwidth of each octave is twice that of the next lower octave.

For each filter, determine the length L that is required to get the correct bandwidth. Use the bandedges determine in Section 5.1. This filter design process will be trial-and-error, so each time you change L you will have to make a frequency response plot (magnitude only) to see if the filter is correct.

- (c) Generate the five (scaled) bandpass filters. Plot the magnitude of the frequency responses all together on one plot (the range $0 \leq \hat{\omega} \leq \pi$ is sufficient because $|H(e^{j\hat{\omega}})|$ is symmetric). Indicate the locations of each of the center frequencies of the five octaves on this plot and illustrate that the passbands cover the separate octaves.
- (d) As help for the previous parts, here are some comments: The passband of the BPF filter is defined by the region of $\hat{\omega}$ where $|H(e^{j\hat{\omega}})|$ is close to one. In this lab, the passband width is defined as the length of the frequency region where $|H(e^{j\hat{\omega}})|$ is greater than 0.5.

Filter Design Specifications: For each octave, choose L so that the entire octave of frequencies lies within the passband of the BPF.

The *stopband* of the BPF filter is defined by the region of $\hat{\omega}$ where $|H(e^{j\hat{\omega}})|$ is close to zero. In this case, we can define the stopband as the region where $|H(e^{j\hat{\omega}})|$ is less than 0.01 because the Hamming filters have excellent stopbands. Notice, however, that the stopbands do not eliminate all the other octaves because the stopband does not start where the passband ends. There is a

small “transition region” between the pass and stop bands where the frequency response is small, but not as small as in the stopband.

Zoom in to show the frequency response over the frequency domain where the DTMF frequencies lie. You can zoom in on a chart or graph by right-clicking on the indicator and navigating to **Visible Items » Graph Palette** and enabling this feature. This will allow you to zoom in using the now enabled magnifying glass tool.

5.3 Piano Octave Decoding

There are several steps to decoding a piano signal to figure out which note is being played:

1. Filter the individual segments to extract the possible frequency components using carefully designed BPFs.
2. Determine a set of time segments over which to examine the output signals. In general, this could be hard to implement so we will restrict the notes to occur at regular intervals based on the timing of the song.
3. Determine which frequency components are present in each time segment by measuring the size of the output signal from all of the bandpass filters. This tells you which octaves are active.

Here is how the system should work: When the input to the filter bank is a piano signal, the outputs from a few of the bandpass filters (BPFs) should be larger than the rest. If we can detect (or measure) the large outputs, then we know which octaves contain notes. These octaves are then used to label the notes partially.

A good measure of the output levels is the *peak value* of the filter outputs, because when the BPF is working properly it should pass only the sinusoidal signals within one octave and the peak value would represent the combined amplitude of those sinusoids. If there is only one sinusoid within the octave then the peak value is exactly equal to the amplitude of the one sinusoid present.

- (a) In order to test your design so far, generate the following signal composed of several sinusoids:

$$x(t) = \begin{cases} \cos(2\pi(220)t) & 0 \leq t < 0.25 \text{ sec} \\ \cos(2\pi(880)t) & 0.3 \leq t < 0.55 \text{ sec} \\ \cos(2\pi(440)t) + \cos(2\pi(1760)t) & 0.6 \leq t < 0.85 \text{ sec} \end{cases}$$

Use a sampling rate of $f_s = 8000$ Hz, and generate $x(t)$ in a vector.

- (b) Filter the vector through the five filters of the filter bank.
- (c) Make plots of the five outputs all together in one figure.
- (d) Use the frequency responses to validate that the output signals have the correct magnitude and phase in each of the three regions where the input has different frequencies. You must determine

which octave the signals are in.

- (e) Comment: observe the *transient* effect at the transitions. This is due to the start-up of each FIR filter as it encounters a new sinusoid.

5.4 Scoring the Activity

The final objective is detecting which octaves contain notes. In order to automate this detection process, we need a score function that rates each octave for possible activity.

Complete the octave scoring function based on the skeleton given in Fig. 4. The input signal to the function must be the input signal containing all the notes. Rather than trying to detect the beginning and end of individual notes, we will have the score function calculate a score every 50 milliseconds.

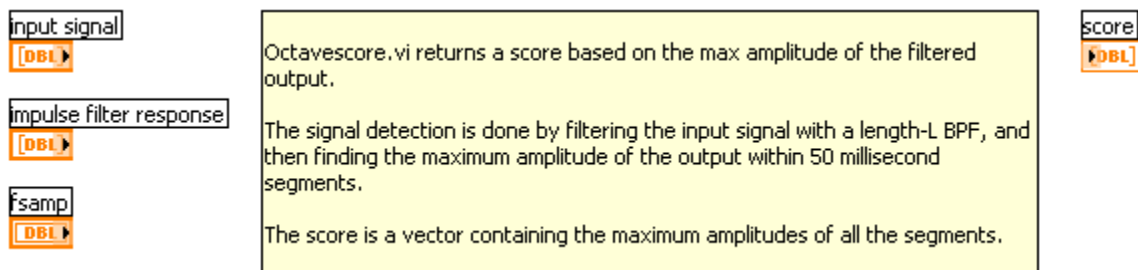


Figure 4: Skeleton of the octavescore.vi function. Complete this code with additional functions.

- (a) Use the following rule for scoring: the score for the i -th channel equals $\max_n |y_i[n]|$, where the maximum is taken over a 50 millisecond interval. The signal $y_i[n]$ is the output of the i -th BPF. Since score must be computed every 50 milliseconds, each channel has a score vector with 20 scores per second.
- (b) Note: After filtering and prior to scoring, adjust the output signal for the delay through the filter. Since each filter has a different length, the delays are slightly different. For the Hamming BPFs, the delay is $(L-1)/2$. This is a minor issue, but it might cause some misalignment of the output scores.
- (c) Finally, use all the score vectors together to implement an automatic detection system, based on the fact that all filter passbands are above 50%. If all the input signals have an amplitude of one, then comparing the scores to a threshold of 0.5 would detect the presence of a signal. This is a big assumption on the inputs, but it will be true in the test data (below).

Comparing the scores to a threshold will give an output that is either zero (score less than threshold) or one (score greater than threshold). So you will end up with a $5 \times N$ matrix consisting of ones and zeros, where N is the number of 50-millisecond intervals. Use your experience with plotting the output signals in the previous section to judge how well your automatic system is working, i.e., are the ones in the correct locations and correct octaves?

5.4.1 Testing

Once you get your system working you could test it in three ways (only #1 needs to be included with the lab report):

1. Use the test signal provided with the lab in 'MATLAB Files' link which contains the MAT file called labtest.mat. It contains a signal sampled at 8000 samples/sec whose duration is 3.77 seconds. The signal has either 0, 1 or 2 sinusoidal notes present at one instant of time. The amplitude of each sinusoid is one. You should run your system to determine which octaves are active in each 50-millisecond time slice. The MAT file also contains a matrix called notes which has the correct notes in the input (as key numbers). The durations vary between 0.16 and 0.32 secs.
2. Generate a music signal that is a progression of notes starting at key #16, ending at key #87, and taking every other note, or every third note. To keep the overall signal reasonably short, make each note about 0.2 seconds long, and make the amplitude of every sinusoid equal to one. This test signal is useful because you know the correct answer.
3. Generate your own musical tones, as you did in Lab 04. Make sure to use $f_s = 8000$ Hz. In this case, you will have several notes playing simultaneously, so the system will have a hard time making 100% correct decisions.

When you run your system, there might be errors some of the time, especially when a note frequency is near the boundary between two octaves. Summarize the behavior of your test cases above, and comment on how well your system works. For example, what is the percentage of errors (i.e., incorrect octave designations)? What happens if two simultaneous notes are in the same octave?

5.4.2 Demonstration

When you turn in your lab report, you must demonstrate to the Lab TA that your system works correctly. We will have a "mystery" signal similar to the signal in labtest.mat.

Lab 10

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: Create a bandpass filter; plot the magnitude response, and then determine the width of the passband at 50% of the peak.

Verified: _____

Date/Time: _____

Part 3.4: Filter the three-sinusoid signal through the BPF, and explain the amplitudes of the output signal in the three sections.

Verified: _____

Date/Time: _____