

Laboratory Assignment 3

Digital Music Synthesis: Beethoven's Fifth Symphony Using MATLAB

PURPOSE

In this laboratory assignment, you will use MATLAB to synthesize the audio tones that make up a well-known sequence from Beethoven's Fifth Symphony by using basic mathematical functions as fundamental building blocks. You will explore how mathematical operations on these signals yield improvements in perceived quality and can be used to generate special effects. By generating a spectrogram of the signal, you will learn how to represent changes in the frequency content of signals over time and how to interpret such spectrograms.

3.1 OBJECTIVES

By the end of this assignment, you should be able to:

Time Domain

1. Construct physically meaningful signals mathematically in MATLAB using the basic signal building blocks and operations on these signals.
2. Gain physical understanding of constructed signals by viewing them on the oscilloscope and using audio playback.
3. Improve the perceived audio quality, shift pitch and generate special effects using mathematical operations on signals.

Frequency Domain

4. Generate a time/frequency representation of a signal and understand the information contained therein.
5. Explain how the spectrograph and musical score represent the same information.
6. Predict how time-domain changes to the signal that produce special effects will alter the spectrograph.

3.2 REFERENCE

—
Review Topics

1. Signal representations, particularly sinusoids and tones
2. Sampling, interpolation and decimation from Lab 2
3. Fourier series

Exploratory Topics

1. Special effects
2. Fourier analysis
3. Time I frequency representations
4. Musical notes and practices

3.3 LABORATORY PREPARATION

Problems

Question 1: Using the information in the background section, write a mathematical expression for a CT signal $c(t)$ that represents your composition. Use shifted CT unit step functions to express your signal analytically as a single function. Be sure to pay attention to the short pauses after the notes (otherwise, the three G notes will sound like one long tone), the frequencies of the notes, and the durations of the notes and rests. For simplicity, you can assume that the amplitude (volume) for each note is constant over its duration.

Question 2: Using your answer to Question 1, write a mathematical expression for the DT signal $c[n]$ obtained by sampling the CT signal $c(t)$ at 8 kHz. Be sure to convert your CT unit step functions to valid DT unit step functions.

Question 3: Determine analytical expressions for the operations performed by your functions **half** and **double** and the MATLAB function **flipud**. Assume that your original signal $x(t)$ is sampled and stored as a column vector \mathbf{x} and that the operations result in a column vector y , which results in the signal $y(t)$ for audio playback. Express $y[n]$ using time scaling and shifting of $x[n]$, and express $y(t)$ using time scaling and shifting of $x(t)$.

Question 4: Determine how you would use MATLAB to generate the signal in Question 2. Also indicate how you would need to modify this MATLAB function to cause the volume of each note to decay exponentially with time.

Question 5: Find mathematical expressions for some of the effects outlined in the section on volume variations on page 29 that you intend to use. Use simple examples to illustrate your answers.

Question 6: How do you expect the volume variations you discussed in Questions 4 and 5 to impact the signal spectrum? What will be the impact on the spectrograph?

Question 7: Given the signal $x(t)$ consisting of two time-varying tones, sketch the spectrograph

$$x(t) = [\cos(100t) + 2 \sin(\omega_1(t)t + \pi/4)]u(t)$$

$$\omega_1(t) = 100(1 - 2e^{-10t})$$

3.4 BACKGROUND


In this section, we explore how to use simple tones to compose a segment of music. By using tones of various frequencies, you will construct the first few bars of Beethoven's famous piece Symphony No. 5 in C-Minor. Each musical note can be represented by a sinusoid whose frequency depends on the note pitch.

Frequencies in Music

Musical notes are arranged in groups of twelve, called octaves. The notes that we'll be using are in the octave containing frequencies ranging from 220Hz to 440Hz. The twelve notes in each octave are logarithmically spaced in frequency, with each note frequency being $2^{1/12}$ times the frequency of the next lowest note. Thus, a 1-octave pitch shift upwards corresponds to a doubling of the frequencies of the notes in the original octave. Table 2 shows the ordering of notes in the octave to be used to synthesize the music for this part of the experiment as well as the fundamental frequencies for those notes.

Table 2: Notes in the 220-440Hz octave

Note	Frequency (Hz)
A	220
A [#] , B ^b	$220 * 2^{1/12}$
B	$220 * 2^{2/12}$
C	$220 * 2^{3/12}$
C [#] , D ^b	.
D [#] , E ^b	
E	
F	
F [#] , G ^b	
G	.
G [#] , A ^b	



A musical score is essentially a program of sorts—a plot of frequencies (notes, on the vertical scale) versus time (measures, on the horizontal scale). The musical sequence of notes to the piece you will synthesize is shown in Figure 3.4.1.




Figure 3.4.1 Opening Notes of Beethoven's Fifth Symphony (1st movement, Op. 67)

Musical Notation

The horizontal lines on the staff in Figure 3.4.1 represent the notes E, G, B, D, and F from bottom to top. The spaces between the lines are used to represent the notes F, A, C, and E, again from bottom to top. Note that A-G only yields seven notes.

The additional changes in pitch are denoted by adding the symbols \sharp (sharp: increase pitch by $2^{1/12}$) or \flat (flat: decrease pitch by $2^{1/12}$) to a given note.

In the musical score in Figure 3.4.1, the first three eighth notes are all G. While it appears that the first half note should be an E, it is actually an E^b due to the inclusion of the three flat symbols at the left of the score. The next three eighth notes are all F, and the final half note is D. You can get the fundamental frequencies for these notes by following the pattern given in Table 2.

Note Durations

In the simplest case, each note may be represented by a burst of a sinusoidal tone followed by a shorter period of silence (a pause). The pause allows us to distinguish between separate notes at the same pitch.

The duration of each tone burst is determined by whether the note is a whole note, half note, quarter note, eighth note, etc. Obviously, a quarter note has twice the duration of an eighth note and so on. The short pause following each note should be of the same length regardless of the note's duration. In composed music, longer periods of silence that are part of the musical score are indicated by one or more rest symbols. The two "4" symbols at the left of the score indicate that there are 4 beats per measure in the score and that a quarter note lasts one beat. For this particular piece of music, the duration of a beat should be about 1/2 second.

Time / Frequency Representations of Signals

Creating music from a musical score is a straightforward process. What about creating a musical score from music? There has to be a way to determine the notes in a sampled signal and to find out when they begin and end.

Such a process is called a time / frequency representation of the signal, or a spectrograph. Time is plotted on one axis, frequency on another, and color is used to represent the presence or absence of frequencies in the signal during that particular time. This is easily done in MATLAB by breaking the sampled signal into small time segments, of approximately 50 msec in duration, and plotting the energy present at each frequency for that time segment. This process is performed by the file **P_3_9 .m**; check the help on this function for specific details of its operation.

Note that the information displayed by this function is almost identical to a musical score. There are some errors in what is displayed, since evaluating the frequency information for a small piece of data is inaccurate.

Improving Perceived Quality

Volume variations: Typically, when a note is played, the volume rises quickly from zero and then decays over time, depending on how hard the key is struck and how long it is depressed. The variation of the volume over time is divided into four segments: Attack, Decay, Sustain, and Release (ADSR). For a given tone, volume changes can be achieved by multiplying a sinusoid by another function called a windowing function. A decaying exponential is the simplest way to modulate the tone volume. You can try concatenating different functions to model ADSR.




Figure 3.4.2 An ADSR Envelope

Overlapping tones: Consider a beginning piano player's attempts at playing a simple tune. This student generally will quit pressing down one key before pressing down the next. A more advanced player would have some notes overlapping (either by using floor pedals or by pressing keys simultaneously). As the volume of one note is decaying, another note is played. Mathematically this can be accomplished by allowing the time regions occupied by different sinusoids to overlap. This will yield a much smoother, less staccato-sounding piece.

Harmonics: While tones are represented by sinusoids, in real instruments the vibrations that generate a given note also include harmonic components. The frequency of the note is called the fundamental frequency f_0 the n th harmonic is a sinusoid having frequency equal to n times the fundamental frequency: nf_0 . The amplitude, or power, of the harmonic components is generally less than that of the fundamental and decreases with increasing frequency. For each note, adding in some lower-power harmonics makes the note sound "richer," whereas a single tone sounds very crisp. A word of caution: adding in harmonics of equal or greater power will have undesired effects, such as increasing the perceived note fundamental, and may make the note sound harsh.

To understand harmonics, recall the Fourier series from mathematics; evaluating a series for a certain number of harmonic components is the desired operation. The magnitude of a Fourier series coefficient indicates the energy in a harmonic frequency. Consider the Fourier series for a square wave - it has only odd harmonics and the magnitude of the n th harmonic is proportional to $1/n$.

Reverb: Instruments are usually played in rooms that generate reverberations. Reverberations are reflections off of room surfaces, including walls, ceiling, floors, and obstructions. In DT mathematics, you can view this as generating echoes of the original signal, with decreasing amplitude as the delay time increases. However, a major surface (e.g., a back wall), may generate more power (amplitude) at specific delays. You can determine the delay time by dividing the distance sound travels by the speed of sound (-- 300 m/s at sea level). A model for generating an echo is given in Laboratory Assignment 2.

3.5 LABORATORY ASSIGNMENT

Part I: Song Creation

Problem 1: Using MATLAB, construct a row vector that is a discrete-time representation of the notes and rests that make up the musical piece in Figure 3.4.1. Let the sampling rate be 8 kHz for this particular composition. Be sure to pay attention to the short pauses after the notes (otherwise, the three G notes will sound like one long tone), the frequencies of the notes, and the durations of the notes and rests. Once you have a piece that you think accurately represents the music, play the composition. Remember to scale your signal before saving it as a sound file.

Verify that the timing and pitch of your composition is an accurate representation of the musical score. Describe how it sounds and how it can be improved. Think about terms like richness of sound, clarity of pitch, flow of composition, musical expression, instrumental versus vocal sounds, etc.

Problem 2: Using what you know about digital pitch shifting and the logarithmic frequency relation between the notes in Table 2, modify your composition using **half** and **double** so that it is played back an octave higher and one octave lower. Note that in the process of digital pitch shifting, you also changed the duration of the signals.

Discuss what happens to the pitch and duration of a signal when you use your functions **half** and **double** to modify it. Be specific: indicate new durations in seconds and new pitches by both frequency and musical pitch.

Problem 3: Try to have the volume of each note decay over time to make your music sound more interesting and realistic.

How did changing the signal amplitude affect your perception of the volume? If you double the signal amplitude, does the volume double as well?

Part II: Note Modifications

Problem 4: Modify your pitch-shifted composition so that each note is of the same duration as in the original composition.

Problem 5: Can you pitch shift your composition up by a single half-step? Consider interpolating first and then resampling.

Problem 6: Allow decaying notes to overlap slightly in time.

Problem 7: Try adding harmonics to create a richer sound.

Problem 8: Use the M-file you created in Lab 2 to generate an echo to create reverb effects.

For the sequence of the possible improvements above, determine the mathematical expressions for the sequence of operations needed to achieve the desired result. Use a simple mathematical example to illustrate your reasoning. Are you able to achieve the desired result? Discuss what worked well and what needs improvement. Explain why.

Part III: Frequency Analysis

Problem 9: Compute the time-frequency representation of your signal from Problem 1 with the function **P_3_9 m**. Compare the plot with the musical score in Figure 3.4.1, and identify each note from the score on the spectrogram.

Problem 10: Generate spectrograms of your pitch-shifted composition from Problem 2. You will need to change the frequency range displayed by **P_3_9.m** to see any results.

How do these differ from the spectrogram of the original composition? Verify that your results are as expected by identifying the pitch and duration of each note on the spectrogram.

Problem 11: Generate spectrograms for some of the special effects you generated in Problems 4 - 8.

How do these spectral effects and improvements alter the spectrogram? Based on the operations you performed and properties of the Fourier Transform, do these changes make sense?