

ECE 2025 Spring 2000
Lab #3: Harmonic Signals & Musical Notes

Date: 31Jan–3Feb 2000

A short Lab Quiz will be held during the first 20 minutes of this lab. It will cover basic MATLAB issues.

This is *the official* Lab #3 description; it consists of some material from the two labs in Appendix C.2 and C.3 of the text.

The Warm-up section of each lab must be completed in Lab and the steps marked *Instructor Verification* must also be signed off **during your scheduled lab time**. One of the laboratory instructors must verify the appropriate steps by initialing on the **Instructor Verification** line. When you have completed a step that requires verification, simply raise your hand and demonstrate the step to the instructor.

Lab Report: Please turn in Sections 3, 4 and 5 with explanations as this week's lab report. Staple the **Instructor Verification** sheet to the end of your lab report as evidence that the appropriate steps were witnessed by the instructor.

The report will **due during the week of 7–10 Feb. at the start of your lab**.

1 Introduction and Overview

The goal of this laboratory is to gain familiarity with sums of complex exponentials and their use in representing more complicated signals such as speech and music. One additional objective is to establish the connection between musical notes, piano keys, and their frequencies. The general form can be expressed in two ways:

$$x(t) = \Re \left\{ \sum_{k=1}^N X_k e^{j2\pi f_k t} \right\} \quad (1)$$

$$x(t) = \sum_{k=1}^N A_k \cos(2\pi f_k t + \phi_k) \quad (2)$$

where the complex amplitude X_k is $X_k = A_k e^{j\phi_k}$.

In this lab we will synthesize waveforms composed of sums of sinusoids, sample them, and then reconstruct them for listening. We will use the sum in equation (1) to synthesize the following signals:

1. Sine waves at a specific frequency played through a D/A converter.
2. Periodic signals that are sums of harmonically related sinusoids.
3. Music signals that match the frequency of specific notes and chords.

1.1 Harmonic Sinusoids

There is an important special case where $x(t)$ is the sum of N cosine waves whose frequencies (f_k) are *different*. If we concentrate on the case where the frequencies (f_k) are all multiples of one basic frequency f_0 , i.e.,

$$f_k = k f_0 \quad (\text{HARMONIC FREQUENCIES})$$

then the sum of N cosine waves given by (1) becomes

$$x_h(t) = \sum_{k=1}^N A_k \cos(2\pi k f_0 t + \phi_k) = \Re \left\{ \sum_{k=1}^N X_k e^{j2\pi k f_0 t} \right\} \quad (3)$$

This particular signal $x_h(t)$ has the property that it is guaranteed to be periodic. Its period is $T_0 = 1/f_0$, because each of the cosines in the sum repeats with period T_0 . The frequency f_0 is called the *fundamental frequency*, and T_0 is called the *fundamental period*. Notice that T_0 is also the shortest possible period.

2 Warm-up

The instructor verification sheet is included at the end of this lab.

In this lab, the periodic waveforms and music signals will be created with the intention of playing them out through a speaker. Therefore, it is necessary to take into account the fact that a conversion is needed from the digital samples, which are numbers stored in the computer memory to the actual voltage waveform that will be amplified for the speakers.

2.1 D-to-A Conversion

The digital-to-analog conversion process has a number of aspects, but in its simplest form the only thing we need to worry about at this point is that the time spacing (T_s) between the signal samples must correspond to the rate of the D-to-A hardware that is being used. From MATLAB, the sound output is done by the `soundsc(xx, fs)` function¹ which does support a variable D-to-A sampling rate if the hardware on the machine has such capability. A convenient choice for the D-to-A conversion rate is 11025 samples per second,² so $T_s = 1/11025$ seconds; another common choice is 8000 samples/sec. Both of these rates satisfy the requirement of *sampling fast enough* as explained in the next section. In fact, most piano notes have relatively low frequencies, so an even lower sampling rate could be used.

2.2 Theory of Sampling

Chapter 4 treats sampling in detail, but this lab is usually done prior to lectures on sampling, so we provide a quick summary of essential facts here. The idealized process of sampling a signal and the subsequent reconstruction of the signal from its samples is depicted in Fig. 1. This figure shows a continuous-time in-




Figure 1: Sampling and reconstruction of a continuous-time signal.

put signal $x(t)$, which is sampled by the continuous-to-discrete (C-to-D) converter to produce a sequence of samples $x[n] = x(nT_s)$, where n is the integer sample index and T_s is the sampling period. The sampling rate is $f_s = 1/T_s$ where the units are samples per second. As described in Chapter 4 of the text, the ideal discrete-to-continuous (D-to-C) converter takes the input samples and interpolates a smooth curve between them. The *Sampling Theorem* tells us that if the input signal $x(t)$ is a sum of sine waves, then the output $y(t)$

¹In MATLAB version 5, the function `soundsc(xx, fs)` performs automatic scaling to avoid saturating the D-to-A converter which would give a distorted sound. If the `sound()` function were used instead of `soundsc()`, it would be necessary to scale the vector `xx` so that it lies between ± 1 . Consult `help sound`.

²This sampling rate is one quarter of the rate (44,100 Hz) used in audio CD players.

will be equal to the input $x(t)$ if the sampling rate is more than twice the highest frequency f_{\max} in the input, i.e., $f_s > 2f_{\max}$. In other words, if we *sample fast enough* then there will be no problems synthesizing the continuous audio signals from $x[n]$.

Most computers have a built-in analog-to-digital (A-to-D) converter and a digital-to-analog (D-to-A) converter (usually on the sound card). These hardware systems are physical realizations of the idealized concepts of C-to-D and D-to-C converters respectively, but for purposes of this lab we will assume that the hardware A/D and D/A are perfect realizations.

- (a) The ideal C-to-D converter is, in effect, being implemented whenever we take samples of a continuous-time formula, e.g., $x(t)$ at $t = t_n$. We do this in MATLAB by first making a vector of times, and then evaluating the formula for the continuous-time signal at the sample times, i.e., $x[n] = x(nT_s)$ if $t_n = nT_s$. This assumes perfect knowledge of the input signal, but we have already been doing it this way in Lab #2.

To begin, compute a vector `x1` of samples of a sinusoidal signal with $A_1 = 100$, $\omega_1 = 2\pi(800)$, and $\phi_1 = -\pi/3$. Use a sampling rate of 11025 samples/second, and compute a total number of samples equivalent to 2 seconds time duration. You may find it helpful to recall that the MATLAB statement `tt = (0 : 0.01 : 3)`; would create a vector of numbers from 0 through 3 with increments of 0.01. Therefore, it is only necessary to determine the time increment needed to obtain 11025 samples in one second. You should use the `add_cos()` function from the previous lab for this part.

Use `soundsc()` to play the resulting vector through the D-to-A converter of the your computer, assuming that the hardware can support the $f_s = 11025$ Hz rate. Listen to the output.

- (b) Now compute another vector `x2` of samples of a second sinusoidal signal (0.8 secs. in duration) for the case $A_2 = 80$, $\omega_2 = 2\pi(1200)$, and $\phi_2 = +\pi/4$. Listen to the signal reconstructed from these samples. How does its sound compare to the signal in part (a)?
- (c) **Insert** the signal `x2` in the middle of `x1` by using the colon notation. Do the insertion so that `x2` starts at time $t = 0.75$ secs. You should be able to use a statement something like:

$$\mathbf{xx} = \mathbf{x1}; \quad \mathbf{xx}(n1:n2) = \mathbf{x2};$$

assuming that both `x1` and `x2` are row vectors. Determine the values of `n1` and `n2` so that the second signal is in the correct place. Listen to this signal. Explain what you heard.

- (d) To verify that the insert operation was done correctly in the previous part, make the following plot:

$$\mathbf{tt} = 0:(1/11025):2; \quad \text{plot}(\mathbf{tt}, \mathbf{xx});$$

This will show the “envelope” of the signal and verify that the amplitude changes from 100 to 80 at $t = 0.75$ secs. Make sure that the `tt` vector is the same as you used in part (a).

- (e) Now send the vector `xx` to the D-to-A converter again, but double the sampling rate in `soundsc()` to 22050 samples/second. *Do not recompute the samples in `xx`*, just tell the D-to-A converter that the sampling rate is 22050 samples/second. Describe how the *duration* and *pitch* of the signal were affected. Explain.

Instructor Verification (separate page)

2.3 Vectorizing in MATLAB

In MATLAB loops are usually very inefficient, so the `for` loop should be avoided as much as possible. This is especially true when the function is going to process a very long signal, e.g., an audio signal lasting many seconds or several minutes.

Learning to write “vectorized” code is easy *if you understand matrix and vector notation*. For example, a sum of products can be rewritten as an *inner product* of vectors:

$$c = \sum_{k=1}^N z_k^* y_k \quad \Longrightarrow \quad c = \mathbf{z}^* \mathbf{y}^T$$

if we define the *row* vectors \mathbf{z} and \mathbf{y} in the following manner:

$$\mathbf{z} = (z_1 \quad z_2 \quad \cdots \quad z_N) \quad \text{and} \quad \mathbf{y} = (y_1 \quad y_2 \quad \cdots \quad y_N)$$

The superscript $*$ as in \mathbf{z}^* denotes the conjugate operator, and the superscript T as in \mathbf{y}^T denotes transpose. Be careful in MATLAB because the prime operator does both the transpose and the conjugate together (see `help transpose`).

How does this impact coding in MATLAB? The answer is that most summations should not be done as `for` loops in MATLAB, but rather should be implemented as some sort of matrix (or vector) multiplication. It turns out that this strategy is not limited to arithmetic operations. It is also possible to vectorize logical operations as discussed in Appendix B, pp. 412–413.

2.4 Vectorizing Arithmetic Operations

Vectorize the following function; in other words, replace the loop with *one* MATLAB statement that does the same thing. Recall mag-squared identity for complex numbers: $|z|^2 = z^* z$.

```
function ss = sumsqd(z)
%SUMSQD do the sum of the magnitude squared of all the elements in z.
% (The input vector contains complex elements)
%
ss = 0;
for k=1:length(z)
    magsqd = conj(z(k)) * z(k);
    ss = ss + magsqd;
end
```

Demonstrate that your vectorized function works by executing: `sumsqd(exp(j*pi*(1:7)/7))`. It is easy to figure out what the answer should be, by using properties of the complex exponential.

Instructor Verification (separate page)

2.5 Vectorizing a Copy Operation

- (a) The *outer product* is a special matrix product that multiplies a column vector by a row vector. It can be used to do some clever operations. Execute the following lines of MATLAB code so that you can explain what operation they perform:

```
yy = ones(5,1) * randn(1,4);
size(yy)
```

```

%--The next 2 statements do the same operation using a different trick
%----that is much faster in execution, although harder to understand
xx = randn(1,4);
yy = xx([1,1,1,1,1],:);

```

- (b) Write a function that performs the same task as the following without using a for loop.

```

function Z = expand(x,ncol)
%EXPAND Function to generate a matrix Z with identical columns equal
%       to an input vector x, i.e., make ncol copies of x.
% usage:
%       Z = expand(x,ncol)
%       x = the input vector containing one column for Z
%       ncol = the number of columns needed in Z
%
x = x(:); %--this turns the input vector x into a column vector
Z = zeros(length(x),ncol);
for i=1:ncol
    Z(:,i) = x;
end

```

2.6 Vectorizing Logical Operations

- (a) Execute the following MATLAB code. Explain what operations are performed by the last 2 lines.

```

A = randn(6,3);
A = A .* (A>0);
A = A + pi*(A==0);

```

- (b) Write a new function that performs the same task as the following function without using the for loops. Use the idea in part (a) and also consult Section B.7.3 on vector logicals in Appendix B: *Using MATLAB*. In addition, the MATLAB logical operators are summarized via `help relop`.

```

function Z = replacez(A,afix)
%REPLACEZ replace negative elements of a matrix with the number afix
% usage:
%       Z = replacez(A)
%       A = input matrix whose negative elements are to be replaced
%       afix = replacement value for the negative elements
%
[M,N] = size(A);
for i=1:M
    for j=1:N
        if A(i,j) < 0
            Z(i,j) = afix;
        else
            Z(i,j) = A(i,j);
        end
    end
end
end

```

3 Exercises: Complex Exponentials

3.1 Sinusoidal Synthesis with an M-file

Since we will generate many functions that are a “sum of sinusoids,” it will be convenient to have a function for this operation. To be general, we will allow the frequency of each component (f_k) to be different. In the MATLAB function, we will use the form given in (4) with the complex amplitudes X_k defined as $X_k = A_k e^{j\phi_k}$.

$$x(t) = \Re \left\{ \sum_{k=1}^N X_k e^{j2\pi f_k t} \right\} \quad (4)$$

Consult equation (2) for the equivalent sinusoidal form.

3.2 Write a Vectorized Sum of Cosines

Write an M-file called `vsum_cos` that will synthesize a waveform in the form of (4). In the previous lab, this function was written with `for` loops, but these are rather inefficient in MATLAB. *In this lab, you must rewrite the function with NO loops.* The first few statements of the M-file are the comment lines—they should look like:

```
function [xx,tt] = vsum_cos(fk, Xk, fs, dur, tstart)
%VSUM_COS VECTORIZED Function to synthesize a sum of cosine waves
% usage:
% [xx,tt] = vsum_cos(fk, Xk, fs, dur, tstart)
% fk = vector of frequencies
% (these could be negative or positive)
% Xk = vector of complex amplitudes: Amp*e^(j*phase)
% fs = the number of samples per second for the time axis
% dur = total time duration of the signal
% tstart = starting time (default is zero, if you make this input optional)
% xx = vector of sinusoidal values
% tt = vector of times, for the time axis
%
% Note: fk and Xk must be the same length.
% Xk(1) corresponds to frequency fk(1),
% Xk(2) corresponds to frequency fk(2), etc.
```

The MATLAB syntax `length(fk)` returns the number of elements in the vector `fk`, so we do not need a separate input argument for the number of frequencies. On the other hand, the programmer (that’s you) should provide error checking to make sure that the lengths of `fk` and `Xk` are the same. See `help error`. *Include a copy of the MATLAB code with your lab report.*

3.3 Testing

In order to demonstrate that your `vsum_cos()` function is correct, synthesize the “vowel” signal shown in lecture and described in Chapter 3 of the book. These signals are shown in Figs. 3.9, 3.10 and 3.11 in the *DSP First* textbook. Make plots for two different cases: the sum of two components which is the signal $x_4(t)$ in Fig. 3.9, and the sum of all five components which is shown in Fig. 3.11. Plot these two signals in the same window with a two-panel subplot. *Include these plots in your lab report with an explanation of why the period is different for the two cases.*

4 Periodic Waveforms: Fourier Analysis and Synthesis

The Fourier Series method provides both *analysis* and *synthesis* capability.

1. If we start with a periodic signal $x(t) = x(T+T_0)$, then we can compute the Fourier Series coefficients by using the following analysis integral:

$$a_k = \frac{1}{T_0} \int_0^{T_0} x(t) e^{-j2\pi kt/T_0} dt$$

2. If we are given the Fourier Series coefficients, $\{a_k\}$, then we can synthesize a periodic signal using:

$$x_N(t) = \sum_{k=-N}^N a_k e^{j2\pi kt/T_0}$$

where $2N+1$ is the number of terms used to form the signal and T_0 is the period.

3. In order to use the `vsum_cos()` function for Fourier synthesis, we must include all the $\{a_k\}$ coefficients for both the positive and negative k . Likewise the frequency vector, `fk` must contain both positive and negative harmonics.

In this exercise, you must synthesize a set of waveforms using different numbers of Fourier coefficients. The formula for the coefficients is:

$$a_k = \begin{cases} \frac{j^k}{(\pi k)^2} (e^{jk\pi} - 1) & k \neq 0 \\ \frac{1}{2} & k = 0 \end{cases} \quad (5)$$

- (a) Write a short MATLAB function that will generate the Fourier coefficients $\{a_k\}$ over the range $-N \leq k \leq N$. Input is N , output is a vector containing the a_k coefficients.
- (b) Use your `vsum_cos` M-file with the fundamental frequency $f_0 = 10$ Hz, to generate synthesized signals with a finite number of coefficients. For the time interval of the synthesis, use $-0.1 \leq t \leq 0.4$ secs.

Make plots for three different cases: $N = 3, 5,$ and 9 (where N is the largest subscript for a_k). Use a three-panel subplot to show all three signals together.

- (c) Explain what happens as N gets larger and also draw a sketch of the waveshape obtained as $N \rightarrow \infty$. If necessary, run your program for $N = 49$ to get the “converged” answer.

Note: You have to set up a time grid for the time interval $-0.1 \leq t \leq 0.4$ secs. The spacing between grid points must be small enough so that the sampling rate (grid points per second) is at least *twice as high as the highest frequency component in your signal*. This fact will be explained in great detail when sampling is discussed in Chapter 4. However, for this section you should just choose a very small spacing between grid points to get a smooth plot. A rule of thumb is to let the sampling rate f_s be at least 50 times the fundamental frequency, i.e., at least 50 samples per period.

5 Piano Keyboard

In the next lab, you will have to build a large program that will synthesize the notes of a well known piece of music.³ Since these signals require sinusoidal tones to represent piano notes, a quick introduction to the frequency layout of the piano keyboard is needed. The result will be that we can develop a formula that gives the frequency for each key. On a piano, the keyboard is divided into octaves—the notes in each octave being twice the frequency of the notes in the next lower octave. For example, the reference note is the A above




Figure 2: Layout of a piano keyboard. Key numbers are shaded. The notation C_4 means the C-key in the fourth octave. The key A_4 is A-440.

middle-C which is usually called A-440 (or A_4) because its frequency is 440 Hz. One octave below is A_3 which has a frequency of 220 Hz; one octave above is A_5 at 880 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. Thus this ratio must be $2^{1/12}$. Since middle C is 9 keys below A-440, its frequency is approximately 261 Hz. Consult chapter 9 for even more details.

Another interesting relationship is the ratio of fifths and fourths as used in a chord. Strictly speaking the fifth note should be 1.5 times the frequency of the base note. For middle-C the fifth is G, but the frequency of G is about 392 Hz which is not exactly 1.5 times 261.6. It is very close, but the slight detuning introduced by the ratio $2^{1/12}$ gives a better sound to the piano overall. This innovation in tuning is called “equally-tempered” and was introduced in Germany in the 1760’s and made famous by J. S. Bach in the “Well Tempered Clavichord.”

You can use the ratio $2^{1/12}$ to calculate the frequency of notes anywhere on the piano keyboard. For example, the E-flat above middle-C (black key number 43) is 6 keys below A-440, so its frequency should be $f = 440 \times 2^{-6/12} = 440/\sqrt{2} \approx 311$ Hz. Now write an M-file to produce a desired note for a given duration. Your M-file should be in the form of a function called `key2note.m`. You may want to call the `vsum_cos` function that you wrote previously. Your function should have the following form:

```
function xx = key2note(X, keynum, dur)
% KEY2NOTE Produce a sinusoidal waveform corresponding to a
% given piano key number
%
% usage: xx = key2note (A, phi, keynum, dur)
%
% xx = the output sinusoidal waveform
% X = complex amplitude for the sinusoid, X = A*exp(j*phi).
% keynum = the piano keyboard number of the desired note
```

³If you have little or no experience reading music, don’t be intimidated. Only a little knowledge is needed to carry out this upcoming lab. On the other hand, the experience of working in an application area where you must quickly acquire new domain knowledge is a valuable one. Many real-world engineering problems have this flavor, especially in signal processing which has been applied in diverse areas such as geophysics, medicine, radar, speech, financial markets, etc.


```

%         dur = the duration (in seconds) of the output note
%
fs = 11025;      %-- or use 8000 Hz
tt = 0:(1/fs):dur;
freq =
xx = real( );

```

For the `freq =` line, use the formulas given above to determine the frequency for a sinusoid in terms of its key number. You should start from a reference note (middle-C or A-440 is recommended) and solve for the frequency based on this reference. For the `xx = real()` line, generate the actual sinusoid as the real part of a complex exponential at the proper frequency. *For your lab report, include the MATLAB code for `key2note()`.*

5.1 Testing `vsum_cos()` and `key2note()` with Sounds

Summation: Use your `vsum_cos()` function to generate the **sum** of four sinusoids with frequencies that correspond to notes in the A-minor piano chord which consists of the keys $\{C_4, E_4, A_4, C_5\}$. If these four notes are played at the same time, the result should be an A-minor chord. Make all the amplitudes equal, pick the phases to be $\frac{1}{2}\pi$, and make the length of the signal 0.8 seconds. Use a sampling rate of 11025 Hz. Play the signal with `soundsc()` to verify that it makes a pleasant sounding chord.

Concatenation: Utilize `key2note()` to generate a **sequence** of notes for $\{C_4^\sharp, E_4^\flat, F_4^\sharp, G_4^\sharp, B_4^\flat\}$. These notes should be played individually, one after another. For the five notes, choose the amplitudes so that the loudness increases by 10% for each successive note; pick the phases to be zero. Make the duration of each note equal to 0.20 seconds, and also put a very short pause (`zeros`) of 0.05 secs. in between each note. **Note:** in MATLAB row vectors can be concatenated by writing `xxx = [xx1, zz, xx2, zz, xx3, zz, xx4]`.

For your lab report, include the MATLAB code used to generate the sounds in both cases. In the summation case, draw a sketch of the spectrum; for the concatenation case, determine the total length of the signal in seconds, and also in number of samples.

Optional: In the summation exercise, try different phases. Pick all the phases to be different and pick the values at random (in MATLAB use the `rand()` function). By listening to several cases, you should be able to convince yourself that “the human ear cannot detect phase.”

Lab #3
ECE-2025
Spring-2000
INSTRUCTOR VERIFICATION SHEET

Staple this page to the end of your Lab Report.

Name: _____

Date of Lab: _____

Part 2.2(e) Synthesizing two sinusoids, combining them and playing them at two different D-to-A rates. Explain the plot and how you placed the shorter sinusoid at the correct starting time.

Verified: _____

Date/Time: _____

Part 2.4

Vectorize the function `sumsqd()` and demonstrate that it works by executing

`sumsqd(exp(j*pi*(1:7)/7))`.

Tell the TA/instructor what you expect the answer to be by doing a mathematical derivation. Write your derivation in the space below.

Verified: _____

Date/Time: _____