

1 Abstract

This lab begins the study of music signal processing by determining the frequencies of musical notes. The goals of this lab are: (1) To be able to use a simple signal processing algorithm to compute frequencies of sampled sinusoids; (2) To use this algorithm to compute the frequencies of musical tones in a pure tonal version of “The Victors”; (3) To use simple data visualization methods to determine the relations between these frequencies. The result will be the 12-tone chromatic scale used for most Western music. You will need to know these frequencies to build simple music synthesizers and transcribers in Projects 1 and 3.

2 Background

We now make a first attempt to analyze musical signals. It makes sense to start with the simplest musical signals: pure sinusoidal tones. Even an untrained ear can sense that these are very simple signals.

We will analyze a basic tonal version of “The Victors.” (If you do not know “The Victors,” then your admission to the University of Michigan may be rescinded). This version contains sinusoids of different frequencies, so we need a way to compute the frequencies of sinusoids from their samples. Then we need a way of interpreting these frequencies: Why them? What are the relations between them? How can we compute them from the tonal signals?

3 Frequency computation

This lab uses extensively the trigonometric angle sum and difference identities:

$$\begin{aligned}\cos(a + b) &= \cos(a) \cos(b) - \sin(a) \sin(b), \\ \cos(a - b) &= \cos(a) \cos(b) + \sin(a) \sin(b).\end{aligned}$$

Adding and subtracting these gives this equality, called a product-to-sum identity:

$$2 \cos(a) \cos(b) = \cos(a + b) + \cos(a - b). \tag{1}$$

This simple formula leads to two different methods for determining the frequency of a sinusoidal signal.

3.1 Tuning fork method

Musicians can use a “tuning fork” to tune instruments. Understanding how this method works mathematically is best illustrated using an example. Substituting $a = 2\pi 443t$ and $b = 2\pi 3t$ into (1) yields:

$$\underbrace{\cos(2\pi 440t)}_{\text{play 440}} + \underbrace{\cos(2\pi 446t)}_{\text{play 446}} = \underbrace{2 \cos(2\pi 3t) \cos(2\pi 443t)}_{\text{play}}.$$

Generalizing the preceding example, one can see that the sum of two sinusoids with the same amplitude and with approximately equal frequencies (that differ by ν Hz) is equal to a sinusoid having frequency midway between these frequencies but with “amplitude” that *varies sinusoidally with frequency $\nu/2$* . You heard an example like this in Lab 1.

Hence one way of measuring the frequency of a sinusoid is as follows:

- Listen to the sum of the sinusoid and another sinusoid at a known and similar frequency;
- This sum will sound like this: “loud-soft-loud-soft,” etc. In the example above, even though the formula for the “amplitude” is $2 \cos(2\pi 3t)$, which is a 3 Hz variation mathematically, our ears hear the sound get louder and softer 6 times per second (6 Hz) because positive and negative amplitudes sound the same to our ears. This amplitude variation is called “beating” and in this example the “beat” frequency is 6 Hz.

- Vary the known frequency until this “beat” disappears. Then the two sinusoids have the same frequency.

Musicians will recognize this method is similar to how pianos are tuned using a tuning fork¹. Guitarists use this technique extensively (at least when their electronic tuner has dead batteries). You do not even have to know the two frequencies, as long as the tuning fork vibrates at the correct frequency.

3.2 Signal processing method using arccos

We prefer an automatic procedure that operates directly on samples of the sinusoid and uses a computer to determine its frequency, like in a smart phone application for instrument tuning. We can do this as follows. First recall that sinusoidal signal is:

$$x(t) = A \cos(2\pi ft + \theta).$$

Let $x[n]$ denote samples of the sinusoid at sampling rate S . Recall from Lab 1 that the samples of a sinusoidal signal are:

$$x[n] = x(t) \Big|_{t=n/S} = A \cos(2\pi ft + \theta) \Big|_{t=n/S} = A \cos(2\pi(f/S)n + \theta). \quad (2)$$

Substituting $a = 2\pi(f/S)n + \theta$ and $b = 2\pi f/S$ into (1) and multiplying by A gives:

$$\underbrace{A \cos(2\pi f/S(n+1) + \theta)}_{x[n+1]} + \underbrace{A \cos(2\pi f/S(n-1) + \theta)}_{x[n-1]} = 2 \underbrace{A \cos((2\pi f/S)n + \theta)}_{x[n]} \cos(2\pi f/S).$$

We rewrite this equality using (2) as follows:

$$x[n+1] + x[n-1] = 2 \cos(2\pi f/S) x[n].$$

Now we rearrange to solve for the frequency f :

$$f = \frac{S}{2\pi} \arccos\left(\frac{x[n+1] + x[n-1]}{2x[n]}\right). \quad (3)$$

(We must choose a time sample n for which $x[n]$ is nonzero.) This is a simple formula for computing the frequency f of a pure sinusoidal signal from (almost) any three of its consecutive samples $\{x[n-1], x[n], x[n+1]\}$. We can also use the arccos formula (3) for a sequence of sinusoidal signals having different frequencies (like in music) because it requires only a couple of time samples. This is called *pitch tracking*. Keep in mind that this arccos method works only for pure sinusoidal signals. Real music is more complicated than (2).

RQ Lab2.1. A sinusoidal signal is sampled at rate $S = 16000 \frac{\text{Sample}}{\text{Second}}$ and four of its consecutive samples are $(x[11], x[12], x[13], x[14]) = (-1.951, 0, 1.951, 3.827)$. Determine the frequency of the sinusoidal signal. (Round to nearest integer.)

¹ “You Can Tune a Piano, But You Can’t Tuna Fish” was an early album by REO Speedwagon, back in the vinyl LP era.

4 Data visualization

Before proceeding with the lab instructions we also must discuss *data visualization*. Recall the steeply dropping curve that depicts radioactive decay of a radioisotope. Also recall the steeply dropping curve that depicts a cable supporting a suspension bridge, from a tower to the midpoint of a bridge. They look similar, but are they? This is a question about modeling and many science and engineering problems require mathematical models for quantities of interest. Common models include linear, affine, polynomial, power, and exponential. We focus on the power and exponential models here.

4.1 Power model: $y = ax^p$

Suppose we believe two sets of data x and y are related by the model

$$y = bx^p \tag{4}$$

for two constants b and p (where p need not be an integer). How can we test this hypothesis and determine b and p ? Taking logarithms² gives

$$y = bx^p \implies \log(y) = \underbrace{p}_{\text{slope}} \log(x) + \underbrace{\log(b)}_{\text{intercept}}. \tag{5}$$

So a plot of $\log(y)$ versus $\log(x)$ is a straight line with slope = p and y -intercept = $\log(b)$. Making such a plot is a quick way of testing this hypothesis and for determining p and b . This kind of plot is called a *log-log plot* because we plot the log of both the x and the y values. `Matlab`'s `loglog` command makes such plots directly, or we could combine the `plot` command and the `log10` command by typing something like `plot(log10(x), log10(y), '*')`.

4.2 Exponential model: $y = ba^x$

On the other hand, if we think that x and y are related by the model

$$y = ba^x \tag{6}$$

for some constants a and b , then we can plot $\log(y)$ versus x because

$$y = ba^x \implies \log(y) = \underbrace{\log(a)}_{\text{slope}} x + \underbrace{\log(b)}_{\text{intercept}}, \tag{7}$$

which is a straight line with slope = $\log(a)$ and y -intercept = $\log(b)$. This is called a *semi-log plot* because we apply the logarithm only to the vertical axis. `Matlab`'s `semilogy` command makes such plots directly, or we could type something like `plot(x, log10(y), '*')`.

Back before computers, people made such plots by plotting data directly on *log-log* or *semi-log* graph paper that has non-uniform scaling that “takes the logarithms” for you. This paper avoided the need to use a calculator to compute logarithms, and also had the advantage that the data values themselves, not their logarithms, are on the axes. The `loglog` and `semilogy` commands share this advantage.

Plotting the two curves mentioned above as semi-log and log-log plots would reveal that radioactive decay follows the exponential model $y = ba^x$ where x is time and $-1/\log_2(a)$ is the half-life, whereas the suspension bridge cable curve is a parabola $y = bx^2$ which is a power model. (This parabola should not be confused with a freely-hanging cable curve, called a catenary.)

Note that both (4) and (6) involve exponentiation, but in the power model (4) we have the independent variable x raised to some power, whereas in the exponential model (6) the independent variable x is in the exponent.

² In this class, and in `Matlab`, “log” always means *natural log* (even though that might be “ln” on your calculator). In this class if we need a base-10 logarithm we will write \log_{10} and in `Matlab` we use the command `log10`.

4.3 Example 1: Semi-log plots



We observe the x, y pair data in the table below left. We wish to determine a model $y = g(x)$ relating x and y . The log-log plot of the data is curved (below middle), whereas a semi-log plot of the data gives a straight line (below right). This tells us that an exponential model is appropriate.

The line's slope is $(2.4 - 1.7)/(5 - 1) = 0.175 = \log_{10}(a) \implies a = 10^{0.175} = 1.5$.

Its y -intercept (where $x = 0$) is $1.505 = \log_{10}(b) \implies b = 10^{1.505} = 32$.

So a reasonable model that fits the given data perfectly is the "exponential" model: $y = g(x) = 32(1.5)^x$.

x	$y = g(x)$
1	48
2	72
3	108
4	162
5	243



4.4 Example 2: Log-log plots

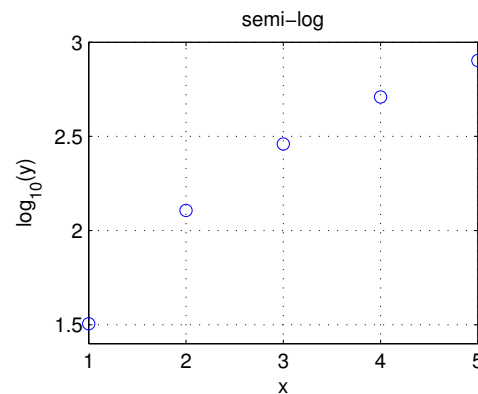

We observe the data in the table below left. We wish to determine a model for the function $y = g(x)$.

The semi-log plot of the data is curved (below right), but a log-log plot gives a straight line (below middle).

Its slope is $(2.9 - 1.5)/(0.7 - 0) = 2 = p$ (power). Its y -intercept is $1.505 = \log_{10}(b) \implies b = 10^{1.505} = 32$.

So a reasonable model that fits the given data perfectly is the "power" model: $y = g(x) = 32x^2$.

x	$y = g(x)$
1	32
2	128
3	288
4	512
5	800

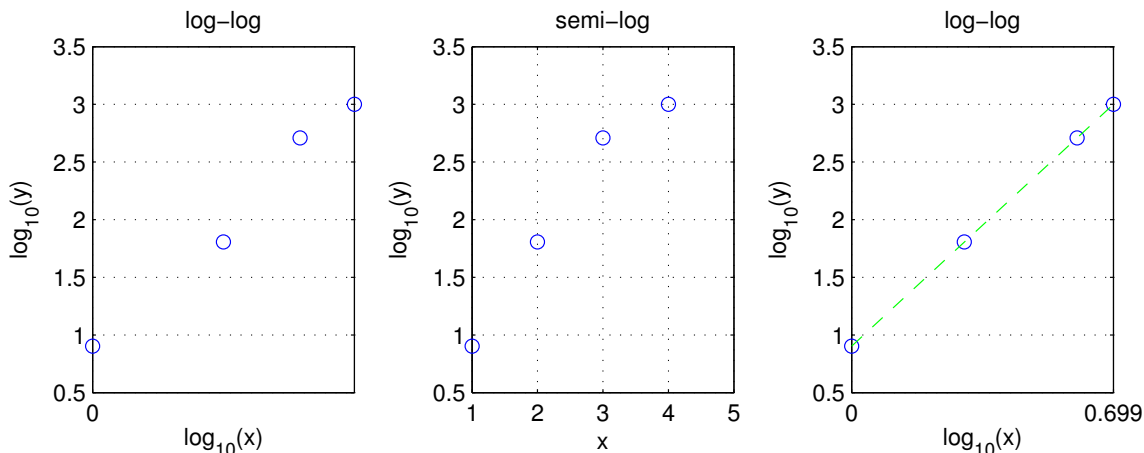


RQ Lab2.2. A particular cubic function has the formula $y = 5x^3$. If we have several x, y pairs and we plot $\log(y)$ versus $\log(x)$ we will get a line. What is the slope of that line?

4.5 Example 3: Determination of missing data

Now suppose that we know $y = g(x) = [8, 64, 512, 1000]$ for *some subset* of $x = \{1, 2, 3, 4, 5\}$. (You will see why this situation is relevant soon.) Can we still find a simple model for $y = g(x)$?

First try plotting y versus $x = 1:4$ on both a log-log and semi-log scale (left and middle plots below).



The semi-log plot gives a curve, whereas the log-log plot gives a straight line with a break, suggesting that $x = 3$ is missing. To verify, plot y again versus $x = [1\ 2\ 4\ 5]$ on a log-log scale. This yields the plot on the right above, a straight line. So a reasonable model is $y = g(x) = bx^p$. The slope is $(3.0 - 0.9)/0.7 = 3 = p$, and the intercept is about $0.9 = \log_{10}(b) \implies b = 10^{0.9} \approx 8$, so a reasonable model is $y = g(x) = 8x^3$. To verify that we computed everything correctly, substitute $x = 1, 2, 4, 5$ into the model $y = g(x) = 8x^3$ and you will get the values 8, 64, 512, 1000 that match y where we started.


Next we apply all of these tools to music signals.

4.6 Matlab tips

- In Matlab, `log` means *natural log*, *i.e.*, log base e , often denoted “ln” on calculators. For base-10 log, use the function `log10()` in Matlab.
- In Matlab, `exp(x)` computes e^x . Functions like `log` and `log10` and `exp` all apply to every element of a vector. To see this, type `log10([100 1000 0.1])`
- In Matlab, `10^x` computes 10^x if x is a scalar. If x is a vector $[a\ b\ \dots\ z]$ then `10.^x` computes the vector $[10^a\ 10^b\ \dots\ 10^z]$ To see this, type `10.^[-1 2 1]` Try typing the following to see what happens: `10.^[-1 2 1]` (That one extra dot `.` makes a big difference!)

4.7 Missing frequencies in “The Victors”

The following diagram may help explain part 5.2d below.



Missing: 1 3 6 8 10 11 12 13 ... -1 -2 -3 ...

5 Lab 2: What you must do

5.0 Frequency measurements for a single sinusoid [8]

- (a) [1] Make a sinusoidal signal and plot it as follows:
`S = 8192; t = [0:999]/S; x = 3 * cos(2*pi*64*t); plot(t, x, '-');
xlabel 't', ylabel 'x(t)'`
- (b) [1] Now use the arccos formula (3) to determine the frequency (in Hz) of this sinusoidal signal as follows:
`n = 7; f0 = (S/2/pi) * acos((x(n+1) + x(n-1)) ./ (2*x(n)))`
- (c) [2] Explain the result of part (b) in light of part (a).
- (d) [1] Repeat (b) with `n = 33` instead of `n = 7`.
- (e) [2] Explain what went wrong in (c). Hint: examine `x(32:34)`
- (f) [1] Repeat (b) with `n = 7:11` instead of `n = 7`.

5.1 Frequency measurements for a “song” [14]

- (a) Download “`vectors_tone.mat`” from the Data folder on Canvas.
Type `load vectors_tone.mat; S = 8192; sound(x, S)` Recognize this?
- (b) [1] Type `n = 3000 + (-200:200); clf, subplot(321), plot(n, x(n), '-')`
- (c) [1] Is it a single pure sinusoid? (Look carefully.)
- (d) [1] Type `n = 2:numel(x)-1; f1 = (S/2/pi) * acos((x(n+1) + x(n-1)) ./ (2*x(n)));
subplot(322), plot(n, real(f1), 'b-', n, imag(f1), 'g-')`
This implements the frequency estimation formula (3) for $S = 8192 \frac{\text{Sample}}{\text{Second}}$.
- (e) [1] Did it work?
- (f) [1] Type `c1 = (x(n+1) + x(n-1)) ./ (2*x(n)); subplot(323), plot(n, c1)`
- (g) [2] Now explain why the first plot of `f1` was not entirely successful.
- (h) [1] Describe how this behavior could actually be helpful for *segmenting* the signal into separate tones.
- (i) [1] Type `c2 = reshape([0, c1, 0], 3000, 26);
f2 = S/2/pi*acos(mean(c2(2:end-1,:))); f2=sort(f2); subplot(324), plot(f2, 'o')`
- (j) [1] This eliminates the first and last numbers in each segment that are *outliers* because of the transition between tones, takes the mean (average) of the numbers in each segment, and sorts the result.
Are these reasonable estimates for a frequency in Hertz?
- (k) Type `audiowrite('vectors_test', x, S); [z S2] = audioread('vectors_test');`
This writes the signal samples $x[n]$ in `x` to the `.wav` file, and then reads the values from that file into the array `z`. A `.wav` file also stores the sampling rate `S`.
- (l) [1] Type `c3 = (z(n+1) + z(n-1)) ./ (2*z(n)); subplot(325), plot(n, c3)`
- (m) [1] Type
`f3 = S/2/pi*acos(c3); subplot(326), plot(n, real(f3), 'c-', n, imag(f3), 'm-')`
- (n) [2] Now we are trying to determine the frequency, using the same formula as before, but from the signal $z[n]$ that is the version of `x` that was written to the `.wav` file. What changed? (You might need to zoom in and look closely.)

5.2 Data visualization [18]

Now we are going to use the note frequencies found in the previous part to determine the relationship between between those frequencies and piano keys. To do this we must somehow number the keys (in ascending order). For now we will take the lowest note in this version of the Victors and label it number 0, and then increment by one for each key to the right.

(a) Type `f = [? ? ? ? ? ?]`

where each “?” is a frequency (in `f2`) that you found, in ascending order. Hint: type `disp(f2)` and `unique(f2)`.
Hint: the first “?” above will be around 392 Hz.

(b) [1] Type `clf, subplot(321), plot(f, '*'), grid`

(c) [1] Type `subplot(322), plot(log10(f), '*'), grid`

(d) Neither of these plots look linear because this song does not use all possible notes so there are “missing frequencies.”

The person who played this song knows which notes were used, so here we give you that data.

(Although you could figure it out from Section 4.5.)

Type `key = [0 2 4 5 7 9]`

(e) [1] Type `subplot(323), plot(key, f, 'o'), grid`

(f) [1] Type `subplot(324), plot(key, log10(f), 'o'), grid`

(g) [6] One (and only one!) of this second pair of plots has points that lie along a line. Determine which one, and then use the ideas from Section 4 to determine a mathematical formula that relates key k to frequency f . Your formula should look something like $f = 3k^2$ but of course not that particular formula.

(h) [2] Your formula will be valid for all keys and notes, not just the 6 notes used in this particular song. Give numerical values for 12 of the frequencies (one octave’s worth), whether they appear or not.

Hint, if your formula were $f = 3k^2$ then in `Matlab` you could see an octave’s (12 notes) worth by typing:

`k=0:11; f = 3 * k.^2`

(i) [3] Use `subplot(325)` and make a scatter plot of the 12 frequencies you found. (Hint: the values will not be on a straight line.) Label axes!

(j) [3] Use `subplot(326)` and make a semi-log plot of the 12 frequencies you found. (Hint: the values should be on a straight line.) Label axes!

5.3 Lab report

Make a brief report that includes your 3 pages of plots and your concise answers to the questions above, including your list of frequencies (both measured and inferred). The audience will be your lab instructor, not your discussion instructor. Save your report as a pdf file and upload it to Canvas. Although we recommend using Google Docs, you may use any editor as long as you export it pdf format for uploading to Canvas. For full credit, the file must be uploaded by the start of next week’s lab. If you were unable to finish it during this week’s lab, then seek help in office hours (if needed) before the due date.

5.4 Next week

Now that you know what pure musical tones are, and their frequencies, you will write a `Matlab` program for a simple musical tone synthesizer with a `Matlab` graphical user interface (*i.e.*, an on-screen keyboard) that generates musical tones. Then you will write a `Matlab` program for a simple musical note analyzer that, when given a pure tonal musical signal, produces a `Matlab` stem plot resembling musical staff notation.