PHYS 391 – Lab 4: Fourier Transforms

Key Concepts

- Fast Fourier Transforms
- Aliasing
- Windowing
- Harmonic Overtones

4.1 Introduction

This lab will explore the experimental aspects of using the Fast Fourier Transform (FFT). This is going to be a much more qualitative lab than the past few labs in the sense that the idea here is to let you explore some of these concepts "by hand" rather than just relying on a mathematical descriptions. In your lab writeup, you should be careful to number each section as you do it and keep descriptive notes of what you are doing at each step. There really isn't much data to collect in this lab, so you will be graded mostly on how well you can explain what you are doing, what you observe, and how you can relate this to the theory of the Fourier Transform.

4.2 DAQ Setup

We will use a program on the PC to emulate a digital oscilloscope and acquire sampled data of voltage versus time. This time-domain data is also sometimes called a "waveform".

Start by finding a PC with an NI USB-6009 or Vernier SensorDAQ data acquisition box (DAQ stands for Data AQuisition). This small white box with screw terminals should be plugged into the USB port of the PC. Attach the output from a Pasco function generator (use the HI impedance output) into one of the *Analog Input* channels on the USB-6009/SensorDAQ. For the SensorDAQ the analog inputs in differential mode are pins 10 & 11. For the USB-6009 pick one of the pairs of terminal blocks labelled AI (for example AI1) and attach both the positive and negative leads from the function generator into the terminal block *and tighten the terminal screws down.* Turn on the function generator and select a 1 kHz sine wave output.

If you are feeling lucky, you can now start up the LabVIEW program ScopeTest (which will either be on the Desktop or can be downloaded from the lab website) and search for this signal. There are two strings in the upper left hand corner of this panel which controls the input data source. The first is the device name, which follows a pattern like `DEV1`. The second is a channel name which follows the a pattern like `AI1`. Case does not matter here, so enter the proper channel name and see if you see something which looks like a sine wave. If there is a red message on the bottom of the panel, you have selected an invalid device name. Try changing the device number starting at zero and see if you can find something which seems to look like a sinusoidal signal. The box is likely either `DEV1` or `DEV2`.

If you would like to be more systematic, you can use the National Instruments *Measurement and Automation Explorer* program to find the USB-6009 device directly and make sure it is working properly. Open the *NI MAX* application (round blue icon on the desktop), click on the Devices and Interfaces icon, then select NI-DAQmx Devices. You should now see the 6009/SensorDAQ in the list with the device name specified. If you select the 6009/SensorDAQ colored green in the list and click on the icon at the top which says "Test Panel" you should be able to check directly that this device is providing good data (hint, set to continuous mode before starting). If the MAX application can't see data from the box, there is no hope for the LabVIEW application to work properly. If you really thing everything is set up properly but you still can't see a signal, try a different input channel on the 6009/SensorDAQ I have seen channels stop working for mysterious reasons. Ask your TA for help if you are having problems.

4.3 Frequency Domain Setup

The ScopeTest application starts with a waveform display looking at Voltage vs. Time (time domain). Set up a sine signal (using the Pasco Function Generator) with frequency $\nu = 1$ kHz and amplitude ± 1 V. The data acquisition parameters should be set to an initial sample frequency $\nu_s = 20$ kHz and a total number of samples $N = 100$. With these settings, what is the minimal frequency resolution $\Delta\nu$ expected from the

FFT? What is the Nyquist limit in terms of the maximum frequency that you expect to be able to resolve? Verify in the time domain that the signal indeed appears to have the correct period (although this may be a bit difficult as we don't have a trigger to synchronize the signal to the display). The 'Hold Data' button can be used to temporarily freeze the display.

Select the tab at the top of the display to change to the frequency domain. Make sure all of the switches on this panel are off (down) so that you have a linear Y scale with averaging turned off. The window parameter should be set to "Rectangular" which is the same as no specific window. You can flip the Autoscale switch on/off to get the signal in range, or else just select the number at the top of the Y scale and type in the maximum amplitude you would like to see. Briefly sketch what you see in your log book and note whether the frequency and amplitude match what you expect. Is the amplitude here the peak amplitude, or the RMS amplitude? The RMS amplitude of a sine wave is defined as $V_{\mathrm{RMS}} = \sqrt{\overline{V(t)^2}} = V_0/\sqrt{2}$, where V_0 is the peak amplitude $V(t) = V_0 \sin \omega t$. In electricity, this is also called the AC voltage.

Turn on the cursor display and measure (by dragging the cursor line back and forth) the frequency spacing between neighboring points in the frequency domain plot. Does this agree with what you expect for $\Delta\nu$ from your sampling frequency ν_s and sample length N? What is the maximum frequency displayed on the plot? Why doesn't this go up to ν_s?

What is the meaning of the point at $\nu = 0$? Adjust the DC offset on the function generator and watch what happens here. (note, there appears to be some sort of ground loop issue between the Lenovo PCs and the Pasco function generators that applies an arbitrary DC voltage (shift up or down) to your sine wave. Experiment using a second Pasco generator as a variable DC source (lower frequency push-button range switch until the display shows all zeros with no decimal point), with the Amplitude button now setting the DC voltage output. See if you can figure out how to wire this up.

4.4 Leakage and Windowing

4.4.1 Frequency Leakage

Slowly turn the frequency on the function generator up from 1 kHz to a frequency exactly $\Delta\nu$ higher (or lower). Describe in words what you see. To take a better look, set the function generator frequency exactly in between two neighboring primary frequencies defined by the points on the frequency domain plot (such as 1 kHz + $\Delta\nu/2$). Sketch in your log book (or print out from the screen) what the FFT looks like and include a rough estimate of the full width at half maximum of the peak. You may want to use the "X zoom" feature which will allow you to zero in on one part of the frequency spectrum. You also may want to momentarily use the "Averaging" if you have excessive noise. Describe in words why you expect to see this "frequency leakage". How does the peak height (amplitude) now compare to what you had before? Estimate whether you would recover the expected amplitude if you were to integrate the area under the peak (estimate this by taking the two largest amplitude points and adding them together).

4.4.2 Frequency Resolution

Calculate how you would expect the frequency resolution ($\Delta\nu$) to change if you were to double the sampling frequency ν_s or double the sample size N. First try doubling ν_s and explore what happens to the frequency peak. Either adjust the slider, or just type the desired frequency (in kHz) into the text box below the slider. Can you understand what happened? Note that the horizontal frequency scale should have changed. Why did this happen? Change the function generator frequency to put the signal right in between two primary frequency points and again estimate the FWHM. Has this improved (reduced) the width of the peak or made it worse? Describe in words what benefit you do gain from a higher sampling frequency ν_s.

Now put the sampling frequency back to 20 kHz and instead double the number of points N. Now describe the effect on $\Delta\nu$ and on the shape of the peak in the case of maximum leakage. Compare the FWHM for this condition to the other two. Increase N to 1000. Does this look better?

Can you think of a case where increasing N alone without increasing ν_s wouldn't actually lead to any improvement? What would happen if you had a transient, non-periodic signal like a short single pulse? What would the effective "sample time" be in this case?

4.4.3 Windowing

Set all of the acquisition parameters back to where they were with $\nu_s = 20$ kHz and $N = 100$. Make sure the averaging is turned off or else you will see really strange things. Slowly scan the function generator frequency and watch how the width and amplitude of the peak changes as the frequency passes through integer multiples of $\Delta\nu$. Now change the FFT window parameter to "Hanning" and repeat this exercise. Describe the differences. How does the peak amplitude now compare to what you expect from the time domain? Does the peak maximum height or integrated peak area best match the time domain amplitude? How did the width of the peak change? Change to a log scale on the vertical axis and compare the effects of applying a Rectangle, Hanning, Hamming, or Blackman-Harris window. Which gives the narrowest central peak? Which gives the smallest tails? You may need to either set the Y axis autoscale or set the Y axis minimum value by hand to see the entire shape of the curve.

If you change the sample length to something large (like $N = 1000$), does the central width and tails from the window go down, or remain about the same? Looking on both the log and linear Y scale is probably a good idea here.

Leave the Hanning window selected for the rest of the lab.

4.5 Aliasing

Change the sampling frequency down to $\nu_s = 5$ kHz to make this section less tedious, but for now set the sample size to something small ($N = 100$ should work). Use a linear scale in Y and make sure that averaging is turned off. Slowly scan the function generator frequency up until you reach the Nyquist critical frequency. What is this frequency? Look at the time domain briefly to see what this looks like. Is there much of a sine wave visible? If you can't see the individual points, you probably need to reduce N. Can you understand what you are observing from what you expect to be happening?

Switch back to the frequency domain and keep increasing the function generator frequency, looking at the spectra as you increase the frequency. Explain what you see. When you reach a frequency of close to 5 kHz, again look at the time domain. What do you see? What is the apparent frequency? Does this agree with the FFT? Does this seem correct? Explain what is going on. If you keep increasing the frequency past the sampling frequency, now what happens in the frequency domain? Change the sampling frequency up to something very large like 40 kHz and prove to yourself that you really aren't looking at a low frequency signal. In fact, if you slowly change the sampling frequency by dragging the slider, note that you can tell the difference between an aliased frequency and a "true" frequency. If the apparent frequency changes with ν_s, it is an aliased artifact!

4.6 Frequency Overtones

Set the function generator back to 1 kHz. Set the sampling frequency to something large (like $\nu_s = 40$ kHz) and the sample length to something large as well (like $N = 2000$). Look in the frequency domain and make sure you have a reasonable window selected. Enable the cursor and move the cursor to select the observed primary frequency peak.

Now, change the function generator to produce a square wave rather than a sine wave. Explain what you see. Find the next highest peak in the spectrum and measure the frequency and amplitude compared to the primary. Using the "X zoom" will probably make this easier. If you set the cursor first, the zoom will magnify the current cursor position. Note that the amplitude at the cursor position is given at the bottom of the panel next to the cursor switch. Turning the averaging on might be helpful as well. To change the frequency range displayed, you can use the scroll bar, drag the cursor to one edge of the visible chart, or just type in directly an upper and lower frequency range.

From the handout of notes in class, what would you expect the frequency and amplitude of this next peak to be? Remember, we worked out the Fourier Series for a step function. Is this an even or odd harmonic of the primary frequency? Select a log scale and look at the multitude of frequency lines. Are all of these consistent with being higher harmonics of the primary frequency as you would expect, or is there some evidence of aliasing? (Turn off the "X zoom" to see this more clearly.) Very slowly raise the function generator frequency (with the averaging off). Do all of the lines more up, or do some move down? What happens if you set the frequency to something like 4 kHz? Where are the aliased high-frequency components in the FFT frequency band? Play around a bit and convince yourself that you could greatly confuse yourself this way. Does it seem like a good idea to have the Nyquist frequency an integer multiple of your primary frequency? Ideally, we would include a lo-pass filter between the function generator and the ADC to directly filter out any frequency above our Nyquist limit.

Go back to a linear scale and set the frequency back to something like 1 kHz. Change from a square wave to a triangle wave. How do the higher frequency harmonics change? Adjust the function generator frequency (or the sampling frequency) to convince yourself that you are not looking at aliased frequencies.

4.6.1 Averaging

Return to some low frequency like 1 kHz and set the function generator back to a sine wave. With the log scale on, you can see the "noise floor" as the constant fuzz at the very bottom of the FFT graph. If you can't see this, select the Y-axis autoscale, or set the minimum Y-axis value to something very small. The variation comes from the randomly changing frequency components present in the noise.

Now, turn on the averaging switch and describe what you observe happen. The averaging algorithm averages the RMS amplitude in each bin of frequency over some number of recent acquisition cycles. It isn't quite a straight average, but rather the most recent data are weighted higher so that you can slowly see changes develop. How does this help identify very small amplitude signals? What is the approximate magnitude of the noise in mV? Does the frequency generator really produce a pure sine wave, or are there other harmonics now visible?

4.7 Musical Instruments

Replace the Pasco frequency generator with the line output of a Casio keyboard. Before plugging keyboard in, understand how to select different instruments (once it is plugged in through the line out the speaker will stop making sound). Start with the flute (which is the closest to a pure sinusoid) and note that if you hold a key down it will continue to play at a roughly uniform volume. Choose suitable acquisition parameters (and describe them) so that you can readily see note played on the keyboard in the FFT panel. Measure the frequency of the highest and lowest note on the Casio keyboard. Adjust your acquisition parameters if necessary.

4.7.1 Beat Frequencies

Play two notes at the same time a half-step apart (such as a white key and a neighboring black key). Use the "X Zoom" feature to see the two frequencies on the frequency diagram. Check the various windows to see how useful each is to separating these frequencies. Adjust your acquisition parameters if necessary to cleanly separate the two frequencies. The 6009 has a maximum buffer size which will be substituted for N if you specify a number which is too large. Measure the frequency difference and from this calculate the expected beat frequency. Check in the time domain and try to verify that the "envelope" follows the expected beat frequency. Compare the frequency difference for two notes low on the keyboard to two notes high on the keyboard. Is the frequency difference constant across the instrument?

4.7.2 Frequency Overtones

Again using the flute, pick a particular note and measure the frequency. Move the cursor to this point to keep track of where the primary frequency for this note is located. Switch instruments to the trumpet and play the note again. You may need to adjust the frequency range so that you can see at least a few overtones above the primary frequency. A range from 0 to 5 kHz is probably good. Sketch roughly what you see from the trumpet in both the time and frequency domain. Measure the relative amplitudes of the first few overtones above the primary frequency. Does this seem more like a square wave or a triangle wave?

Try this again using a different instrument, either a pipe organ or a jazz organ. How do the overtones here differ? Are the amplitudes of the higher harmonics greater or less than the primary? Is there any evidence of a sub-harmonic frequency (below the primary as defined by the flute)? Feel free to play around here a bit and try to compare what your ear hears to what the FFT shows you in terms of frequency components.

Do some research and explain, based on the physics of musical instruments, why a trumpet, jazz organ, or another instrument (whose spectra you examined) have their own particular spectra. See if the spectra you examined using the Casio to synthesize these sounds matches what you find out about a particular instrument.

4.7.3 Melody

Go back to the flute, and compare the individual frequencies for a standard C chord (C, E, and G). If you know nothing about music, these are helpfully labelled for you on the keyboard. Is there any obvious pattern for how the frequencies of the E and G notes match to the fundamental C? Repeat this for the trumpet which may make the relationship more clear. How do the harmonics of particularly the C and G note relate to each other? Does there seem to be an explanation from the frequency domain measurements for why certain combinations of notes might sound “pleasing” rather than dissonant?

4.8 Final Comments

This lab has only scratched the surface in demonstrating some of the useful things that can be done with the FFT experimentally, but it has illustrated several of the pitfalls which can easily be avoided. The issues of aliasing, leakage, and windowing are always present in any experimental frequency-domain measurement. Without an understanding of where these issues arise from and bit of care to mitigate their effects, one can be severely fooled by discrete sampled data. Hopefully this lab will save you an afternoon of grief some day when you have to do this for real.