How to Analyze the Frequency Content of a Signal

Author: John M. Cimbala, Penn State University
Latest revision: 10 October 2013

Introduction

- When the frequency content of a signal is *known*, or at least *approximately known*, it is easy to choose the sampling frequency – we choose a sampling frequency larger than twice the maximum expected frequency of the signal (using the Nyquist criterion).
- In this learning module, we discuss what to do in situations in which the frequency content of a signal is *unknown* (as in many typical laboratory situations). In particular, *how then can we avoid aliasing errors*?

Apply a low-pass (anti-aliasing) filter

- The best way to eliminate aliasing errors entirely is to ***filter*** the signal *before* taking the digital data.
  - In particular, a ***low-pass filter*** is used; thus, a low-pass filter is often called an ***anti-aliasing filter***.
  - Here is how this anti-aliasing-filter technique works:
    - A low-pass filter lets low frequency components of the signal *pass* through (that's why it is called a "low-pass" filter), but cuts off frequency components above some specified ***cutoff frequency***.
    - An ***ideal low-pass filter*** *abruptly* cuts off the high frequency content of the signal – all content of the signal at frequencies greater than the cutoff frequency is removed.
    - Ideally then, the cutoff frequency of the ideal low-pass filter is set to one-half of the sampling frequency to avoid aliasing, $f_{\text{cutoff}} = f_s/2$ (following the Nyquist criterion).
    - That way, *all frequencies that would have produced aliasing errors are removed from the signal before the digital data are sampled, and therefore no aliasing is possible*.
  - Unfortunately, real-life low-pass filters do not cut off high frequencies abruptly; instead, the attenuation of higher frequencies falls off rather slowly with frequency.
  - Because of this, most real-life data acquisition systems employ a cutoff frequency *several times smaller* than the sampling frequency to ensure proper anti-aliasing.
- Low-pass filters are discussed in more detail in another learning module.

Sample at different sampling frequencies

- Even without a low-pass anti-aliasing filter, we are often able to analyze the frequency content of the signal by *sampling at several different sampling frequencies, and comparing the corresponding frequency spectra*.
  - Principle: If we sample at two different sampling frequencies, and the peaks in the frequency spectra appear at *different* frequencies, we can be sure that aliasing errors are occurring.
  - In this case, the sampling frequency must be continually increased until the peaks in the spectra *do not change*, and *the aliased peaks can be consistently explained*. An example is provided below.
- In many practical applications, such as measurement of turbulent fluctuations in a fluid flow, there are sometimes no distinct peaks in the frequency spectrum; instead there is a broad range of frequencies that contribute to the spectrum.
  - In this type of situation, instead of a sharp peak, we see a broad bump in the frequency spectrum.
  - If the sampling frequency is high enough, the frequency spectrum should drop off toward zero amplitude near the folding frequency of the spectrum. Otherwise, we are not sampling at a high enough frequency, and the high-frequency components of the signal may contribute to aliasing errors.
- These techniques are best illustrated by examples, as follows.

Example

- To illustrate how to correctly find the frequency content of a signal, a voltage signal with more than one frequency component, along with some random noise, is measured using digital data acquisition.
- We *do not know* the frequency content of the signal, but we expect a frequency component around 15 Hz.
- We sample at more than one frequency, as discussed above, to illustrate the technique.
- The first case is sampled at f_s = 30 Hz, with N = 256. Note that 30 Hz is chosen to illustrate the technique. In general, we start with a sampling frequency greater than twice the maximum expected frequency of the signal, in order to avoid potential aliasing problems. Since we expect a frequency component around 15 Hz in this example, we sample at twice that frequency.
- The time trace and frequency spectrum are shown below for the first case, f_s = 30 Hz and N = 256:

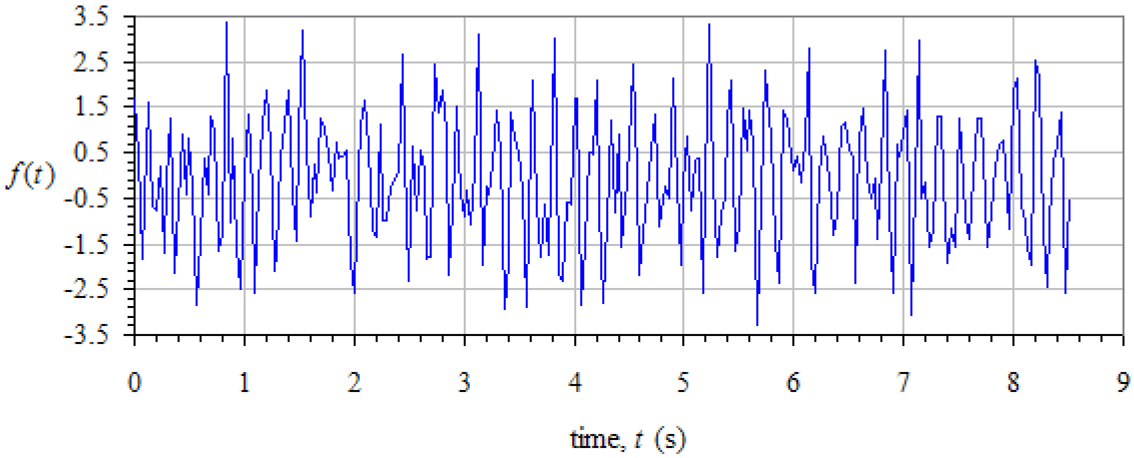

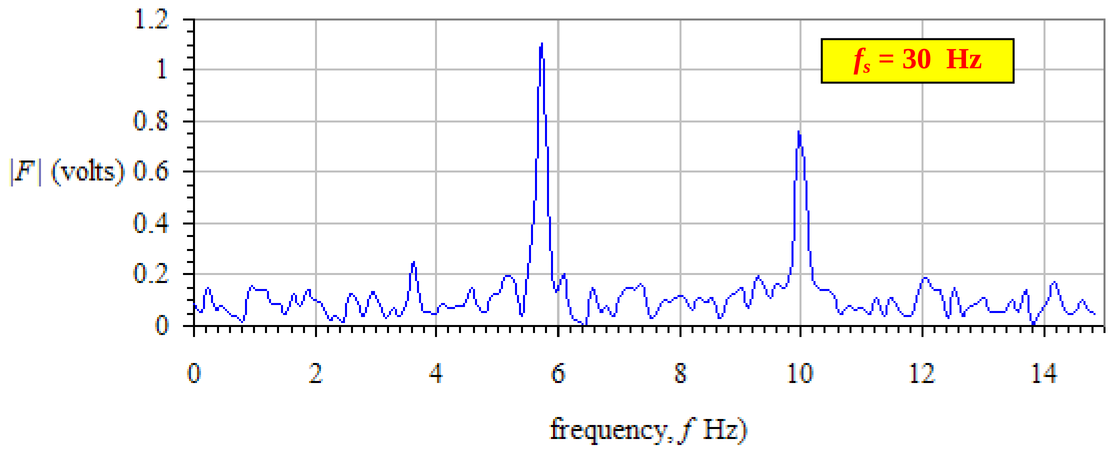

- There appears to be a strong 10 Hz component of this signal (amplitude around 0.8 V), along with an even stronger component at around 5.7 Hz (amplitude around 1.2 V).
- Since the sampling time is arbitrary, there is some leakage, as discussed in a previous learning module.
- The rest of the frequency spectrum shows random noise.
- Can this spectrum be trusted? In other words, is there possibly some aliasing error? One quick way to tell is to *sample at a higher frequency*.
- *Note*: It is important to sample at a new sampling frequency that is *not* an integer multiple of the previous sampling frequency; otherwise it is possible for the aliasing frequency to appear the same in both spectra. [This would potentially give us a false confidence in our frequency analysis.]
- For example, 60 Hz (30 × 2), 90 Hz (30 × 3), etc. would *not* be good choices for our new f_s.
- We choose f_s = 66 Hz. (N is kept at 256 for all of these examples.) The time trace and the corresponding frequency spectrum are shown below. (*Note*: The sampling period T has been reduced from about 8.5 s to about 3.9 s, since we are sampling at a higher frequency but with the same number of data points.)

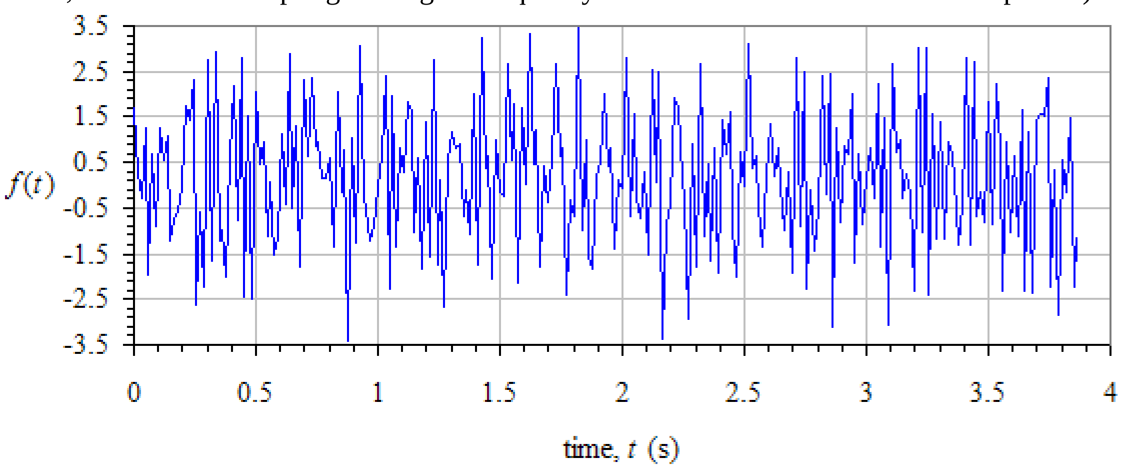

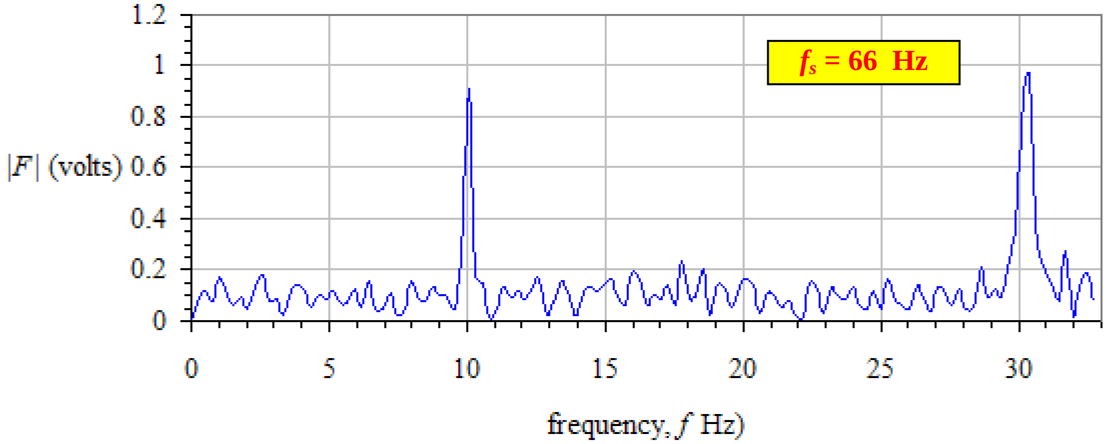

- The peak at 10 Hz is still present, but the one at 5.7 Hz is gone; in its place is a peak at 30.3 Hz! Clearly, *aliasing errors occurred in the first case*, and the sampling rate may *still* not be high enough.
- We examine the first two cases to see if they are *consistent*:
  - Sampling at 66 Hz yields a frequency component at 30.3 Hz.
  - Since f_s = 66 is greater than 2(30.3) = 60.6, it is possible that this represents a legitimate frequency component of the signal.
  - If this *is* legitimate, we should be able to *predict* the aliased frequency from the first case. Let's check:
    - Sampling at f_s = 30 Hz yields a frequency component at 5.7 Hz.
    - If the actual frequency component of the signal were 30.3 Hz, we can predict the aliasing frequency, $f_a = |f - f_s \cdot \text{NINT}(f/f_s)| = |30.3 - 30.0 \cdot \text{NINT}(30.3/30.0)| = 0.3$ Hz when sampling at 30 Hz.
    - However, this is *not* consistent with the observed frequency of 5.7 Hz when sampling at 30 Hz.
  - So, because of the inconsistencies, we conclude that the sampling rate of 66 Hz is *still* not high enough.
  - Sampling must be increased to an even *higher* frequency to resolve the problem.
- We must sample at an even higher sampling frequency than 66 Hz. We need to choose a frequency that is not an integer multiple of either of the first two (30 and 66). So, for example, 90 Hz would not be a good choice (30 × 3). 96 would not be a good choice either, since it is the sum of 30 and 66. 100 Hz is a good choice.
- We sample a third case at 100 Hz. We do not show the time traces from here on. The frequency spectrum is shown below:

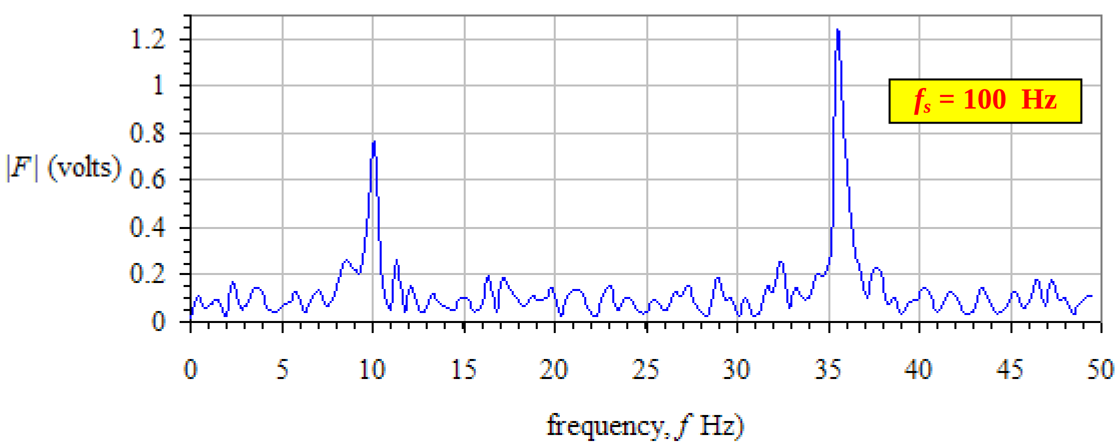

- This time, the peak at 10 Hz remains – we are confident now that the peak at 10 Hz is real. But the other peak has again moved – this time to about 35.7 Hz. Apparently, the previous peaks at 5.7 Hz and 30.3 Hz are *not real*, but are *falsely perceived due to aliasing*.
- This indicates that f_s = 66 Hz is *still* not a high enough sampling frequency to avoid aliasing.
- Is 100 Hz a sufficient sampling rate for this signal? Is aliasing still a problem? To check, we examine the first three cases to see if they are consistent:
  - When f_s = 100 Hz, the peak frequency is at 35.7 Hz, but when f_s = 66 Hz, the peak is at 30.3 Hz.
  - If the actual frequency component of the signal were 35.7 Hz, we could predict an aliasing frequency of $f_a = |f - f_s \cdot \text{NINT}(f/f_s)| = |35.7 - 66.0 \cdot \text{NINT}(35.7/66.0)| = 30.3$ Hz when sampling at 66 Hz. This is consistent with the frequency observed when sampling at 66 Hz.

- When sampling at 30 Hz, if the actual frequency component of the signal were 35.7 Hz, we could predict an aliasing frequency of $f_a = |f - f_s \cdot \mathrm{NINT}(f/f_s)| = |35.7 - 30.0 \cdot \mathrm{NINT}(35.7/30.0)| = 5.7$ Hz. This is consistent with the frequency observed when sampling at 30 Hz.
- Thus, since everything is now consistent, we conclude that the sampling rate of 100 Hz *is* high enough.
- The bottom line – the signal contains:
  - A component at **10 Hz**, with an amplitude greater than **0.75 volts** (the exact amplitude cannot be known because leakage has reduced the height of the peak at 10 Hz).
  - A component at **35.7 Hz**, with an amplitude greater than **1.2 volts** (the exact amplitude cannot be known because leakage has reduced the height of the peak at 35.7 Hz).
  - Some low amplitude noise at all frequencies (random noise).

- As a final check, another case is run at a sampling frequency of 125 Hz, the spectrum of which is shown:

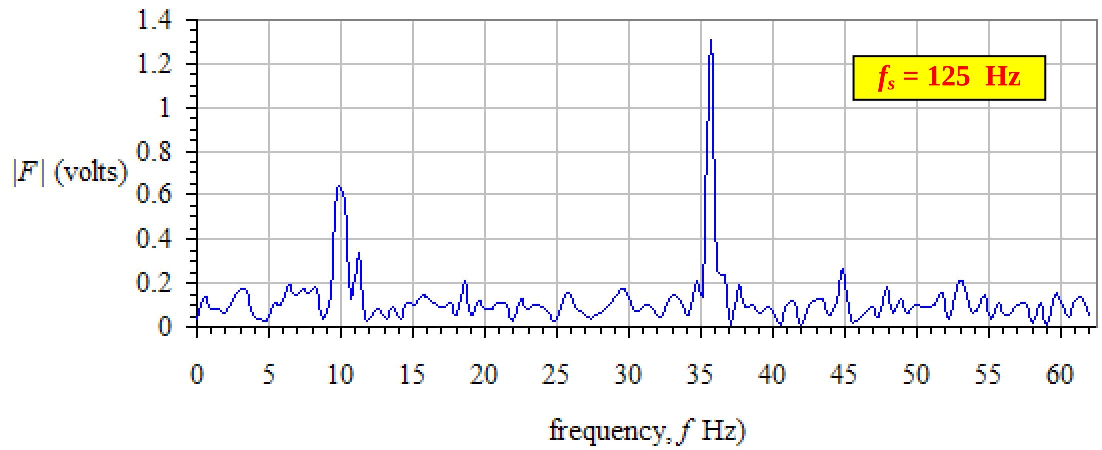

- Now, finally, the peaks have remained at the same frequencies (10 Hz and 35.7 Hz) as previously. We are very confident now that the signal is being sampled at high enough frequency to avoid aliasing. The two major peaks in the frequency spectrum have been correctly captured.

Why not sample at a really high frequency to begin with?

- Why didn't we just sample at a very high frequency right at the beginning, and avoid all this trouble?
- To answer this question, we do one more case, at a sampling frequency of f_s = 500 Hz, while keeping the number of data points at N = 256. The frequency spectrum is plotted below:

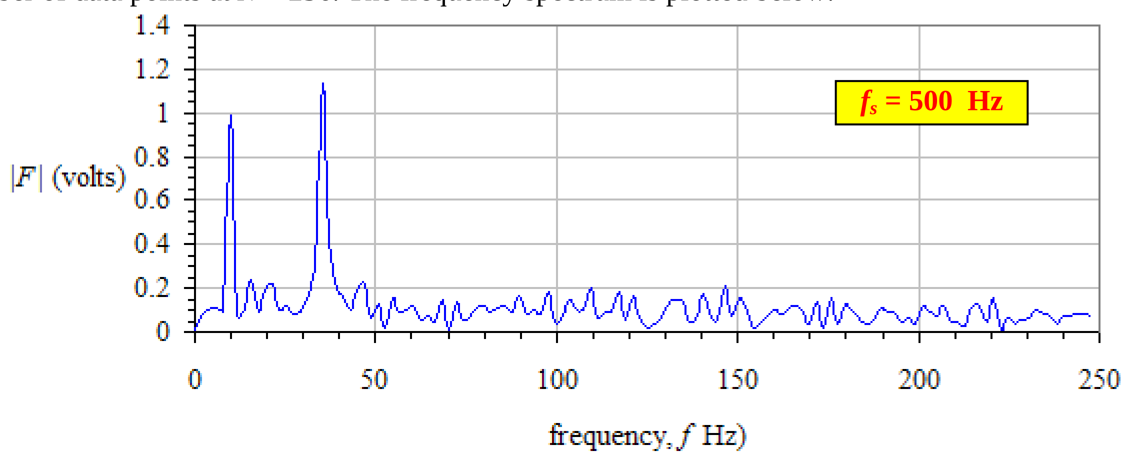

- If this were the first spectrum we created, there would be several problems:
  - With this one spectrum, we cannot be sure that the observed peaks are real or aliased, unless we know for certain beforehand that the maximum possible frequency in the signal is below 250 Hz.
  - The frequency resolution has been reduced significantly. For example, at a sampling frequency of 100 Hz, and with 256 data points, $\Delta f = 1/T = f_s/N$ = (100 Hz)/256 = 0.391 Hz. By comparison, at a sampling frequency of 500 Hz, and with 256 data points, $\Delta f = 1/T = f_s/N$ = (500 Hz)/256 = 1.953 Hz.
  - In fact, the peaks in the above spectrum occur at 9.76 and 35.2 Hz instead of 10.0 and 35.7 Hz – we lose frequency resolution at high sampling frequencies.
  - Data sampled at frequencies greater than twice the frequency component at 35.7 Hz are "wasted."

Effect of adding an anti-aliasing filter

- So far, this signal was not filtered in any way. There is still non-zero amplitude due to noise in the signal, apparently at all frequencies, even those near the folding frequency of 250 Hz for the case with f_s = 500 Hz.
- However, the high frequency energy is merely *noise*, and is really of no interest in this particular problem.
- One way to remove the high frequency noise in this example is to *low-pass filter the signal, using a cutoff frequency somewhat above the maximum frequency of interest in the problem*, which is 35.7 Hz here, and then *sample at a frequency several times higher than this cutoff frequency*, depending on the type of low-pass filter used.
- To illustrate, the signal is filtered with an eighth-order low-pass Butterworth filter, with a cutoff frequency of 50 Hz. Data are sampled at f_s = 150 Hz and N = 256. The frequency spectrum is shown below:

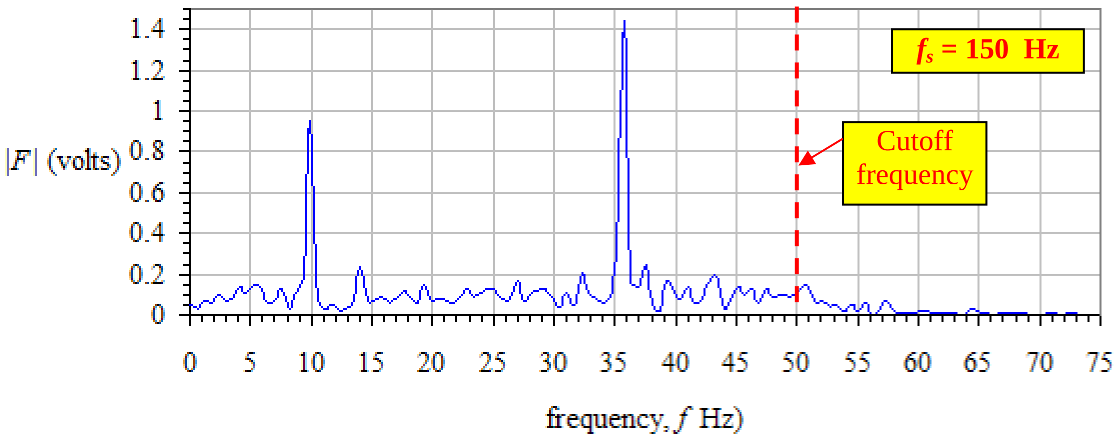

- Notice how the magnitude of the frequency spectrum drops off beyond the cutoff frequency, as desired.
- *Note*: Due to the random nature of the signal, the signal sampled for this FFT is not the identical signal that had been sampled for the previous FFTs.
- Only now can we completely trust the results of the frequency spectrum. In this example, we state confidently the following:
  - There is a frequency component of the signal at 10 Hz, with an amplitude greater than 0.9 V.
  - There is an even stronger frequency component at 35.7 Hz, with an amplitude greater than 1.4 V.
  - There is random noise over the rest of the frequency range.
  - There *may* be some other minor peaks at frequencies above 75 Hz, but these cannot be identified, and are not of interest; they are buried in the background noise.
- This example illustrates some of the frustrations and problems associated with FFT spectral analysis. You must be very careful when applying FFTs so that the frequency spectrum yields the correct frequency content of the signal!