Pitfalls in the Analysis of Machinery Vibration Measurements

John L. Frarey, Schenectady, New York

Strictly speaking, the FFT analyzer is intended for use with harmonically related signals that are stationary in time. In other words, all signals are in the center of the FFT bin and never change with the analysis time. We have learned how to live with non-bin centered data by the use of windows. We still have a way to go in handling very noisy data. That path is even longer in being able to extract useful data from transient vibration data. The secret to obtaining useful data in these cases is an understanding of the FFT process and how it may be used to provide good approximations of transient data. This article will discuss several cases where great care must be taken in setting up the analyzer or very misleading data will be produced.

The knowledge and understanding of the FFT process by vibration specialists have greatly expanded since I gave my first talk on the FFT at the 1980 Annual Meeting of the Vibration Institute in New Orleans. Then, data acquisition times, windows and the relationship between the A/D sample rate and the F_{max} of the 400 line spectrum were not widely understood. Today, anyone certified as a Level 1 Vibration Specialist has a pretty good idea of how the FFT vibration analyzer works. In fact they probably use one daily in their machinery condition monitoring program.

In spite of this improved understanding of the FFT process, there are still some areas where the unsuspecting can get into serious trouble. In some monitoring programs the number of averages to take for accelerometer data seems to be a constant of nature. We always take 8 averages for example. The number of averages should be determined after a review of the character of the data. For extremely noisy data, the combination of too few averages coupled with very tight vibration limit specs may lead to an unnecessarily high number of alarms.

Strictly speaking, the FFT process should be used only on stationary data. In other words, the data must not be varying in frequency or statistical characteristics. Additionally, it was intended to only treat harmonic data where each signal is precisely a harmonic of the center of the first bin. If these rules were strictly followed, then the FFT would not be the popular processing tool it has become. One must realize however, that when one uses the process in nonstationary data situations, one has to be very careful in setting up the analyzer or serious errors can occur. The FFT has even been used to process transient data. The idea that the peak hold feature can produce the amplitude portion of a Bode plot during startups or shutdowns can result in getting bad data for the response at the critical speed, depending on the setup of the analyzer. Some of these situations will be analyzed in this article.

The Number of Averages is a Constant of Nature

The statement that "When looking at accelerometer data from a new machine, a 400 line spectrum and 8 averages is a reasonable first choice" has been interpreted in some plants as "Use a 400 line spectrum and 8 averages for all accelerometer data." The first statement assumed that one would produce a spectrum and then depending on the character of the data, modify the number of lines and the number of averages to produce a spectrum where the noise level will impact the discrete frequency signal by an acceptable amount. The random noise level in the same bin as a discrete signal may add or subtract from the amplitude of the signal. (This actually takes place in the real and imaginary spectra before combining them into the magnitude spectrum.) The number of averages smooths the noise level to its average value but averages the magnitude of the discrete signal to a value nearer its actual amplitude. Therefore, the variability in the discrete signal amplitude will approach zero if enough averages are taken.

Figure 1 shows the time history of a noise signal including one discrete frequency component. The spectrum shows a 400 line analysis that has been averaged by summation eight times. The spectrum appears quite clean. Figure 2 shows the amplitude of the discrete signal after repeating the test ten times. The actual level of the signal is shown, as are the ten amplitudes measured, the average of the ten amplitudes and the $\pm 3\sigma$ values. About 99% of the data will fall between these 3σ values. One can see that in two successive tests, the data for the constant amplitude discrete signal varied from 8.5 to 9.7. This is a 14% increase in the measured amplitude when none existed in the actual values. Figures 3 and 4 show the same test except that the number of averages was 64. Of the ten trials, the minimum amplitude was 9.15 and the maximum amplitude was 9.41. This is an increase in the measured amplitude of 2.8% when none existed in the actual value of the discrete signal.

Suppose the above variations in reading the amplitude of the discrete signal were coupled with a machine speed that could vary. Assuming a Hanning window is used, an additional 14% increase could occur if a speed change moved the discrete signal frequency from the edge of the bin to the center of the bin.

The lesson from the above data is that there is a relationship between averaging time or data acquisition time for higher line spectra and how tight one can set the alert limits for a changed vibration reading. Eight averages of a 400 line spectrum for all acceleration data is not a wise rule for data collection. Rather we must make the decision based on consideration of the time to take the data and the amount of variability we will accept in the amplitude of discrete signals.

Beware of Fat Peaks

In an ideal situation, all signals would reside in the middle of a given FFT bin and not vary from this location during the time to acquire data. Sometimes in real life this condition cannot be fulfilled. Figure 5 is a comparison of two time histories and their respective spectra. Note that the peak amplitude of both traces is approximately 250 mV. The top spectrum, associated with the top time trace only shows a peak amplitude of 180 mV. The bottom spectrum shows the correct amplitude of 250 mV. Besides the different amplitude readings, the major difference between the two spectra is the width of the peak. Very close inspection of the top trace would show that the time per cycle is varying. In other words, the signal is frequency modulated and thus moving through several bins during the data acquisition time. In a real case this means that the machine speed has varied during the data acquisition time or in the case of an instability such as oil whip or steam whirl, the frequency of the signal has varied during the data acquisition period.

The true amplitude of a fat peak can be partially corrected by taking the square root of the sum of the squares of the data in each bin. While this correction will partially correct the amplitude of the varying signal, it may not give us the true value of 250 mV. The signal resides for some time at the edge

---

Based on a paper presented at the National Technical Training Symposium and 26th Annual Meeting of the Vibration Institute, Pittsburgh, PA, June 2002.

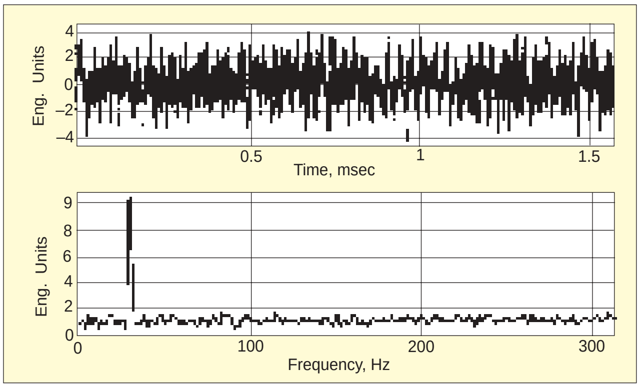

Figure 1. The top trace is the time history of a signal containing one discrete frequency component and a high level of noise. The bottom trace is the spectrum of these data after 8 averages.

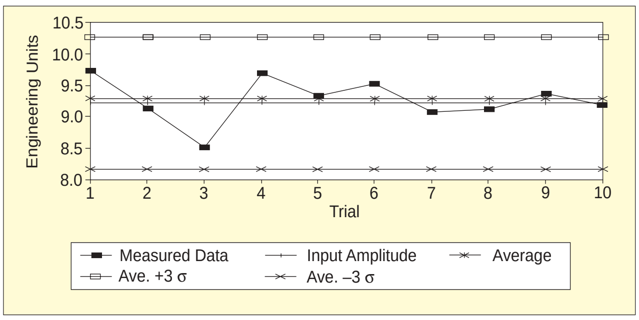

Figure 2. The amplitude of the discrete signal as measured after 8 averages. The test is repeated ten times. The average of the peak readings, the actual input data and the ±3σ levels are also shown.

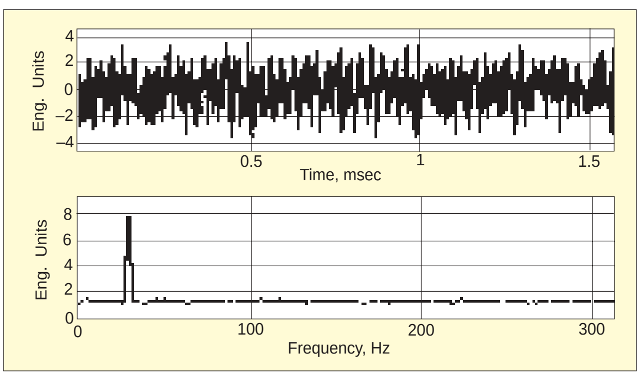

Figure 3. The top trace is the time history of a signal containing one discrete frequency component and a high level of noise. The bottom trace is the spectrum of these data after 64 averages.

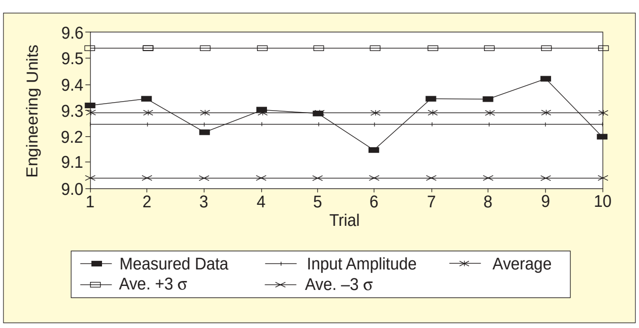

Figure 4. The amplitude of the discrete signal as measured after 64 averages. The test is repeated ten times. The average of the peak readings, the actual input data and the ±3σ levels are also shown.

of the bin, which introduces a picket fence error (reduces the amplitude by 14% for a Hanning window) in the signal reconstruction. A solution to the problem would be to reduce the resolution of the spectrum by either selecting a higher f_{max} or reducing the number of lines of resolution. This will increase the bin width, allowing for better containment of the signal

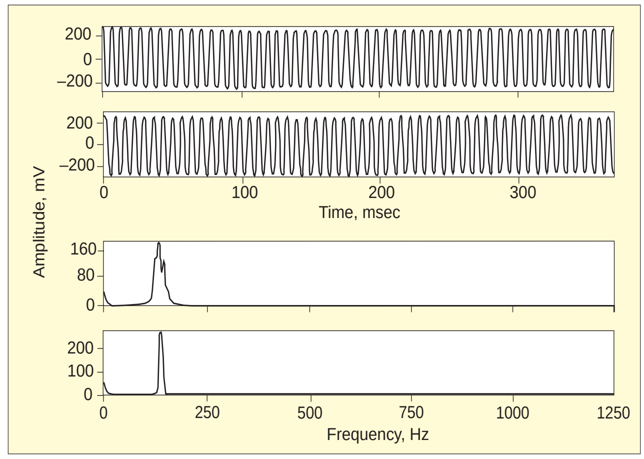

Figure 5. Two signals of equal amplitude are shown in the top two traces. The spectra for these two traces are shown below. Note that the spectra of the data in the top trace show a reduced amplitude and a widened peak. The bottom spectrum reads the amplitude correctly. Careful examination of the top time trace shows that the frequency is not constant but varying.

within its bandwidth.

We Can Make the Vibration Disappear

This example is taken from "Case Histories – Power Industry" by Kevin Guy and published by the Vibration Institute. Unstable vibrations began occurring during the operation of a 650 MW turbine generator. Operation would continue for some time with the vibration levels on the turbine supervisory system showing 1 to 2 mils but then growing to 8.5 mils for about one sec. As time went on, the levels became higher and in order to prevent an alarm, the power level was decreased by 20 MW which caused the vibration variations to be temporarily eliminated. After a period of time the high vibration excursions would return and the power would again be reduced.

Vibration spectra were produced from the shaft riders that were installed on the unit. When using only one average, the vibration could be higher on a few spectra while the majority of the single spectra showed low vibration. A sub synchronous vibration at 1575 CPM was observed. The peak was very wide and was exciting the shaft critical at 1575 CPM. If several averages of the spectral data were taken, the vibration levels were low and indicated a smooth running machine. In order to investigate the spectral width, a zoom spectrum was produced. This showed even lower levels and when several zoom spectra were averaged, the vibration levels were even lower.

The utility management decided that since this vibration dropped below the alert level when averaged and was even lower with a zoom spectrum, this was not the problem. Their solution was to filter the data below 3000 CPM, which produced uniformly low vibration readings.

Understanding the FFT process can explain the disappearance of the vibration averaging and using a zoom spectrum. First, the high vibration would occur on only a few of the non-averaged spectra. This is consistent with the turbine supervisory system. Averaging the data simply averaged one high vibration spectrum with several low vibration spectra resulting in a low level of vibration in the averaged spectrum. The excitation occurred over a wide band so that when the bandwidth was narrowed in the zoom spectrum, the levels became even lower. The referenced case history did not describe what happened next but one can be sure this was not the end of that story.

What Answer Do You Want?

The peak hold option of an FFT analyzer will store the maximum amplitude that was achieved in a bin during the input time. It is touted as the way to get the amplitude portion of the bode plot during coast-downs or startups. This situation is the epitome of non-stationary data. Using the FFT analyzer in this

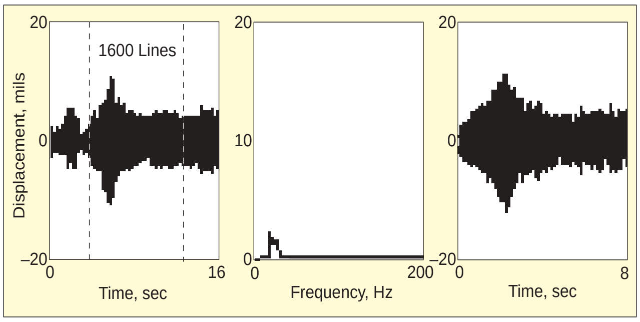

Figure 6. The complete time trace of a coast down is shown in the left panel. The amount of this transient that is required to fill the input buffer for a 1600 line spectrum is shown in the right panel. The middle panel shows the peak hold representation of the transient using 90% overlap.

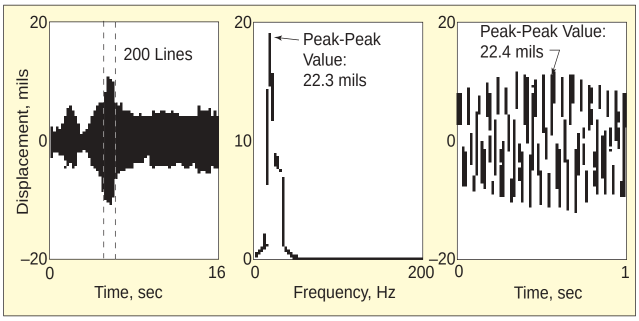

Figure 7. The complete time trace of a coast down is shown in the left panel. The amount of this transient that is required to fill the input buffer for a 200 line spectrum is shown in the right panel. The middle panel shows the peak hold representation of the transient using 90% overlap.

case violates the rule that the data must be stationary and leads to the question, "Is it possible to get a good amplitude versus speed plot during a coast down?" The answer is yes, if one is careful to set up the analyzer in a manner that will give it a chance to get something close to accurate data. If the analyzer is not set up properly, inaccurate or misleading results can be obtained. Unfortunately, the results of an expensive FFT analyzer are usually trusted regardless of how it was set up.

The data presented here are from a classroom demonstration given by Nelson Baxter in the Level 3 Vibration course, and is used here with his permission. The first advantage we have today is that many analyzers allow us to capture the complete coast-down in the input buffer and then decide how best to set up the analyzer to handle the data. Not so long ago, one had to set up the analyzer correctly before the data were seen or use an analog tape recorder to replay the data until the result was closer to the actual case.

One of the advantages of modern analyzers is that they allow us to select the number of lines of resolution in the spectrum. We can select from 200 to 1600 lines of resolution in most analyzers and have even more choices in some analyzers. One would think that the highest resolution would be the proper choice to exactly determine the critical speed.

Figure 6 shows a complete coast-down and peak-hold spectrum. The right panel shows the data needed to fill the input buffer for a 1600 line spectrum. The spectrum shows that the peak to peak response at the critical speed is about 2.3 mils. If we did not observe the data that filled the input buffer and remained confident with our firm belief that the analyzer can not lie to us, we might report this as the response of the machine while it coasted down through the critical speed. We look at the data in the input buffer that was used to produce the spectrum and we see two things. For 1600 lines it takes 8 sec to acquire the data needed to produce the spectrum. During this time both the amplitude and the frequency have changed dramatically. The peak hold value saved in the spectrum is low because the data were within the bandwidth of the bin for a very short time. The peak to peak amplitude of the data in the input buffer also changed during the time taken to fill the buffer. Clearly we have done something wrong.

We fell into the trap of thinking that higher resolution is better. To obtain high resolution requires a lot of data to fill the input buffer, which takes time. We should have minimized the time to fill the buffer and not maximized the resolution.

Figure 7 shows the same coast down with a 200 line spectrum instead of the 1600 line spectrum. Instead of 8 sec to fill the input buffer it now takes only 1 sec. The panel at the right shows the data in the input buffer that is used to compute the FFT. Now looking at the peak hold coast-down shown in the middle panel, one can see that the peak to peak amplitude as the machine coasted down through the critical is 22.3 mils. Next, looking at the time trace of the data as it traverses the critical, one can measure the actual peak to peak value of the maximum amplitude which is 22.4 mils. Notice that there is almost a ten to one difference in the measured response between the 1600 line and the 200 line spectrum. The 200 line spectrum is very close to the actual peak to peak value as measured from the time trace at the critical speed. The actual RPM may now be determined by measuring the period of the data at the critical speed.

I would guess that there has been a lot of incorrect data reported on a machine's response as it passed through the critical by analysts that only looked at the peak hold spectrum of the data. This was particularly true before analyzers allowed one to store the complete coast down in the input buffer.

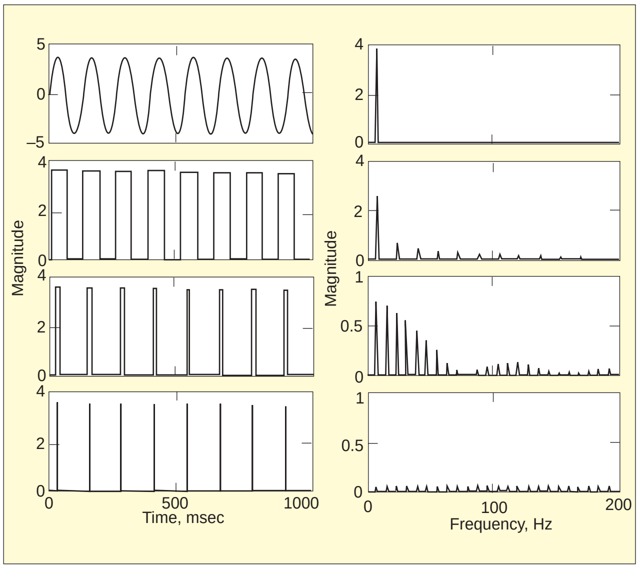

Figure 8. Spectrum energy distribution as a function of pulse shape.

Bad Bearing Detection? Now it is Much Easier

Back in the dark ages of vibration analysis in the 1960s and early 1970s, detecting a defective rolling element bearing sounded simple but unfortunately turned out to be very difficult. At the time, equations for the frequency of signals generated by a defective bearing depended on the location of the defect, the geometry of the bearing and the speed. Conceptually the bearing defect frequencies would be calculated and then the spectrum would be examined for a signal at these frequencies. If a signal were there, then the bearing would be defective. If no signal were detected, then the bearing must be ok. The problem with this approach was that in many situations where no signal was detected, the bearing was known to be defective or worse yet, the bearing soon failed. In cases where defective bearings were detected using this procedure, the background noise in the spectrum was extremely low or the bearing defect was very large.

The reason for the poor success of this method lies in the nature of the signal generated by the defective bearing. Figure 8 shows the spectra generated for several types of signals. The top pair shows the spectrum generated by a sine wave, a single

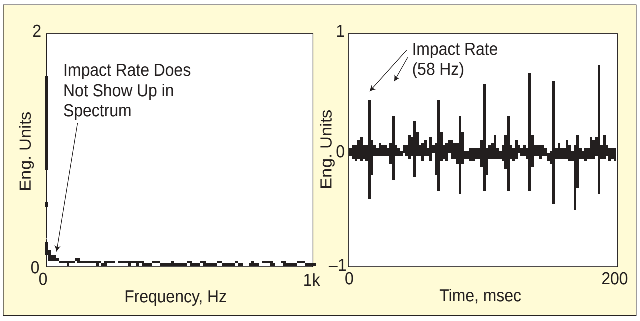

Figure 9. A simulation of the time trace of a defective bearing and the resultant spectrum of this signal.

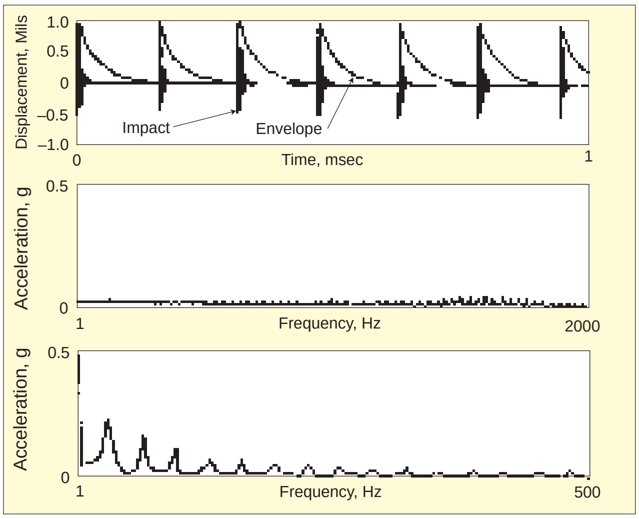

Figure 10. The top trace shows the simulated impact and ring down of a defective bearing. The middle trace is the spectrum of this impact and ring down. The second signal in the top trace is the envelope of the impact and ring down. The bottom trace is the spectrum of the envelope.

peak. The next pair of signals shows the spectrum generated by a square wave. The energy of the signal is shared by all the odd harmonics of the repetition rate of the square wave. The fundamental is still the largest signal in the spectrum. The third pair shows the spectrum for a much shorter duration pulse than the square wave. The fundamental is still the largest signal in the spectrum but much more energy is shared with all the higher harmonics. The fourth pair of signals shows a very short impact in the time trace. This spectrum shows that there are many harmonics of the fundamental and all are very low in amplitude compared with the time trace amplitude. This is the typical spectrum generated by a hammer in modal analysis testing.

If the shape of the defect time trace were a sine wave, then one should have very little trouble in detecting the defect signal in the spectrum. On the other hand, if the time trace is more like a hammer impact, then we can see that the amplitude of the fundamental is very low and shared almost equally with all harmonics. We can imagine that a ball impacting a pit in the raceway would produce a signal more like the impact hammer than the sine wave. Figure 9 shows the spectrum and the time trace for a defective bearing. This figure was furnished by Nelson Baxter from the Level 3 Vibration course. Note that the time trace shows a very sharp impact, meaning that the energy will be spread among many harmonics, as the spectrum shows. The fundamental defect frequency and its harmonics are barely detectable. If a little machinery background noise were added to the time trace, one can see that the evidence of a defect could easily be obscured. Some early reports on the use of this technique claimed that the defect showed up in the third harmonic while others reported seeing evidence in the seventh harmonic. What was happening in these cases was much like the impact hammer modal test – the harmonic was

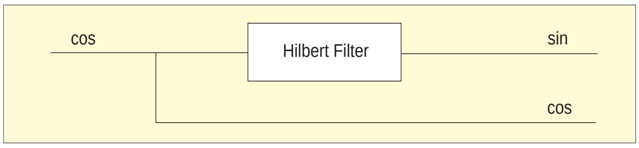

Figure 11. Hilbert filter and analytic signal.

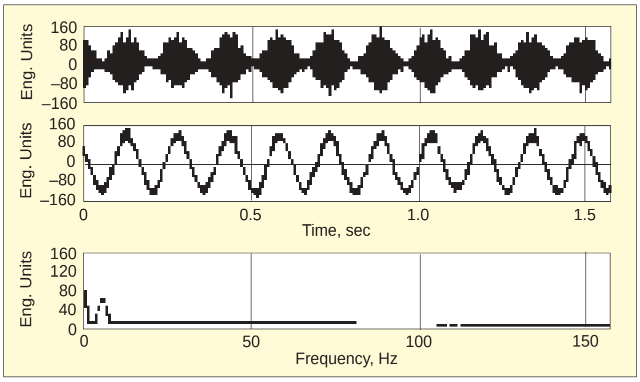

Figure 12. The top trace is a high frequency modulated by a sine wave. The second trace is the envelope of the top trace using the Hilbert transform. The bottom trace is the spectrum of the envelope.

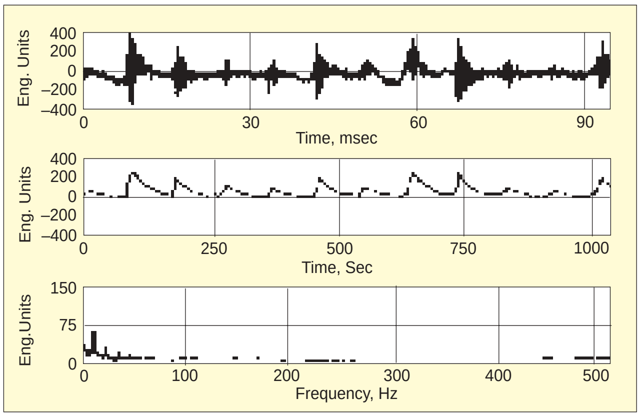

Figure 13. The top trace is the signal from a defective bearing. The second trace is the envelope of the top trace using the Hilbert transform. The bottom trace is the spectrum of the defective bearing.

amplified by a resonance. It is quite common to see a bearing defect pulse excite a structural resonance but this usually happens at much higher frequencies where very small motions can produce very high g levels.

Figure 10 shows three traces. In the top trace a simulated bearing defect impact train is shown. In the middle trace the spectrum of the impact train is shown. Note that there are very low amplitude signals at the defect frequency in the lower frequency portion of the spectrum. The data at the high end of the spectrum are the actual ring down frequency of the train of impacts shown at the top. This resonance is very strongly modulated by the defect signal. One could determine the presence of the defect signal from the sideband spacing of the high frequency data. In real life cases, machinery noise obscures this sideband structure or possibly more than one resonance is excited and overlaps the first resonance. In any event, detecting the low frequency defect frequency or detecting the defect frequency by the resonant sidebands is very difficult. One could measure the time between the impacts in the top trace and convert this period to a frequency to show that it is the defect frequency of the bearing. This may often be done when viewing a bearing defect in the time domain. Again, sometimes machine noise and other resonances preclude making an accurate measurement of the impact spacing.

Suppose one could beef up the thickness of the impact pulse and also remove the ring down frequency. This is done in the second trace in the top chart. The line drawn is called the envelope of the train of pulses and may be produced by a simple diode, capacitor and resistor network. The circuit must have a very fast rise time while the decay time is set by the time constant of the *RC* circuit. The spectrum of the envelope is shown in the bottom trace. The defect signal is now clearly identified. One might wonder what the harmonics of the defect frequency mean. The answer is that they have no physical significance but only define the decaying pulse shape which is controlled by the *R* and *C*, not by the condition of the bearing. This technique has dramatically improved the ability to detect defective bearings and has become one of the standard methods.

Software Envelope Detector

We all know that a function such as an envelope detector, implemented in hardware, costs money to originally design and test. In addition, it costs money for parts and labor every time we make one. If we could duplicate the function in software, we would still have to design and test it, but it would not cost anything to 'build.' We can charge enough for the software function to pay the nonrecurring costs. Since there are no recurring costs, we should increase our profit. In addition, a software solution does not soak up any power, weighs nothing and uses no space. Therefore, all other things being equal, a software solution is preferred. The manufacturers make more money and the user has a smaller, lighter unit.

It is possible to make a software envelope detector utilizing a technique called the Hilbert transform. The Hilbert transform shifts the phase of every signal in the input by 90 degrees, shown schematically in Figure 11. By utilizing the original and shifted data one can see that we have a situation analogous to real and imaginary signals. The amplitude of this signal may be obtained by taking the square root of the sum of the squares at each sample of the data. For example, if the signal had a fixed amplitude, this technique would produce a fixed amplitude DC value. If the original signal varies with time or has an envelope varying with time, the output of the combination will be the envelope of the input signal.

Figure 12 shows a high frequency signal modulated by a low frequency sine wave. The second trace is the result of shifting the phase of the signal by the Hilbert transform and then combining the 'real' and 'imaginary' signals to get the amplitude. This amplitude is the envelope of the signal or the modulating signal. The bottom trace is the spectrum of the envelope.

Figure 13 shows the same sequence of signals. The top trace shows raw data from a defective bearing. The second is the envelope of the defect signal and the bottom is the spectrum of the envelope. The f_{max} of the spectrum is high since the resonance excited by the defect impacts is a high frequency. To examine the spectrum of the defect signal, a zoom Hilbert transform could be performed. This has the effect of moving the defect signal to the right in the spectrum. The disadvantage of the zoom spectrum is that it destroys the typical defect type signals in the top two traces.

Time Series Averaging – Extracting Signals from Noise

What is required for time series averaging is both the noisy signal and a pulse synchronized to the data. One possibility is to use a key phasor marker on the input shaft of a gear box and a signal from an accelerometer mounted on the gearbox. Although most analyzers perform time series averaging in a more sophisticated manner, it is easy to picture if one envisions continually adding the synchronized signal to itself and then stopping the process and dividing the sum by the number of additions made to arrive at the average value. One can also see that adding a random noise signal to itself (random phase and sign) and eventually dividing the sum by the number of additions will result in a noise signal that will approach zero as the number of averages becomes large. Figure 14 shows this process for 1 average (raw signal), 48 averages and 256 averages. The synchronized signal emerges from the noisy signal as more averages are taken.

Figure 15 dramatically shows the value of time series averaging. The top trace is the original signal. The second trace is the result of 128 synchronous averages. This trace clearly shows that one pulse is different from the others, such as one tooth mesh in the signal generated in a gearbox. The bottom trace is the spectrum of the averaged signal. It is almost impossible from this spectrum to determine that one tooth mesh is different from the others.

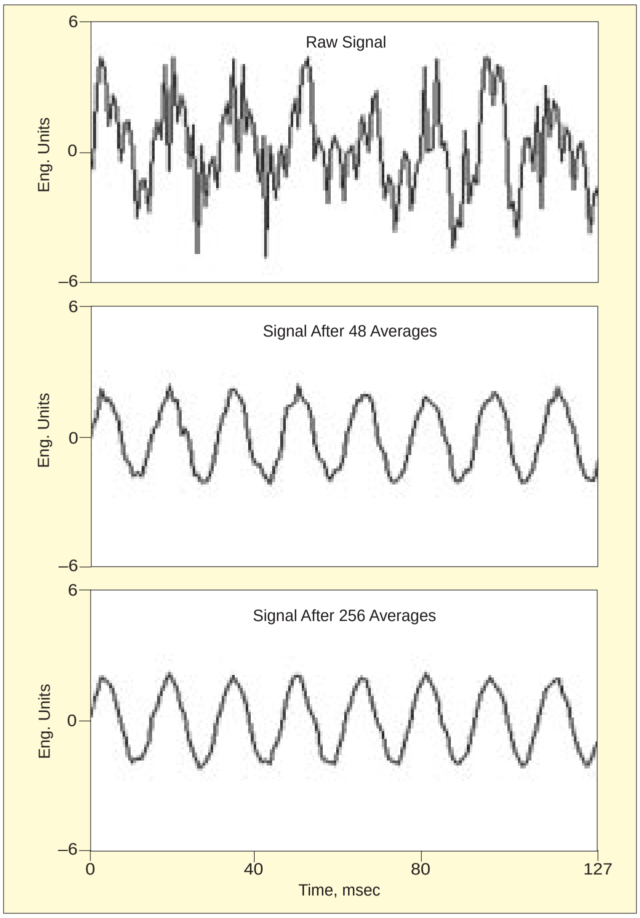

Figure 14. Shows the effect of time series averaging for 1, 48 and 256 averages.

Figure 16 is a plot of the averaged noise amplitude versus the number of averages. Notice that this is a smooth curve where the noise level is inversely proportional to the square root of the number of averages. For example, the fourth average is one over the square root of 4 times the first average or a ratio of 1:0.5 It is thus easy to predict the signal to noise ratio from the number of averages.

Figure 17 is the same type of plot but here the signal is an unsynchronized sine wave instead of random noise. For the first few averages, the amplitude drops dramatically as the unsynchronized signal drifts out of phase with the first signal. Then the amplitude increases as the signals drift back in phase. To ensure a good ratio between a synchronized and a non-synchronized signal one must either take a great number of averages or observe the running average and stop it when the desired ratio is achieved or the non-synchronized signal has disappeared. Figure 18 shows the results of 16 trials of 8 averages of a non-synchronized signal. Note that the resulting amplitude could vary from 30 to 160.

Time Series Averaging Can Also Fool You

To illustrate that there is no one magic data processing technique, Figure 19 shows data from a single mesh gear box. The input shaft is bent so that the contact point of teeth varies around the pitch line circle in a sinusoidal manner. This modulation is clearly shown in both the top time trace and the top spectrum. Note that if the time series averaging were synchronized to the output shaft, all evidence of the modulation disappears. This illustrates the fact that one should not get locked

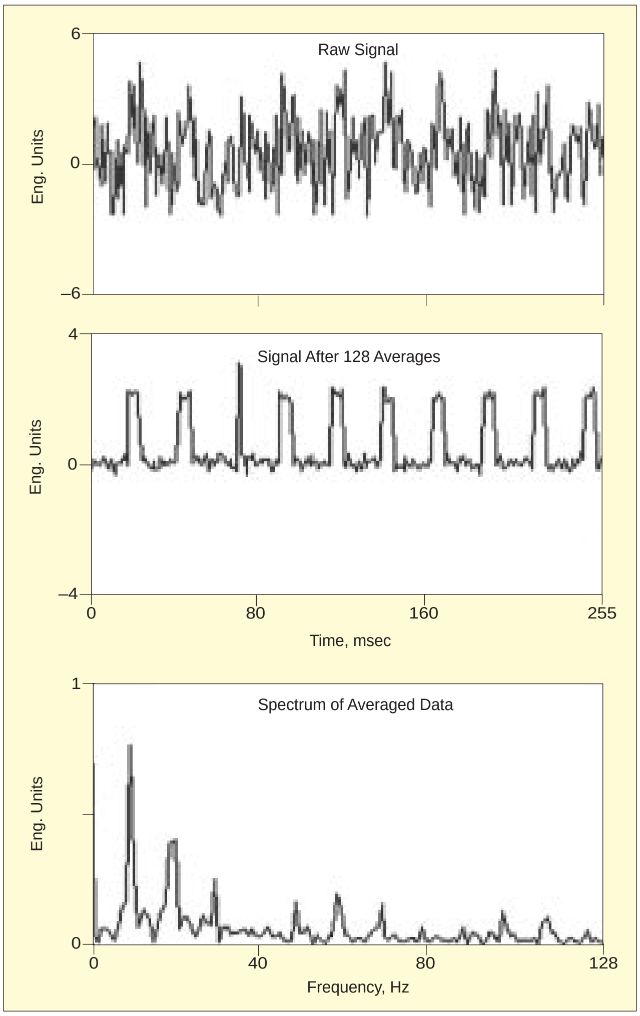

Figure 15. Time series average is used to recover a signal from noise and display it in the time domain.

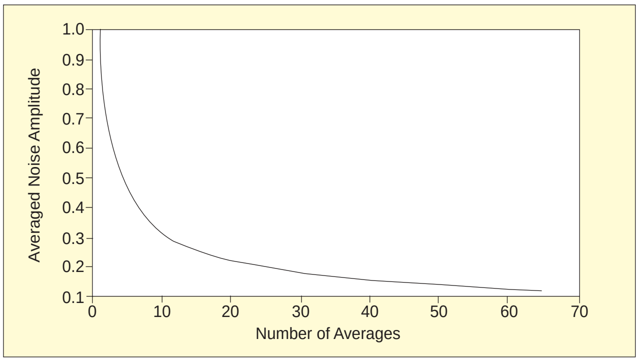

Figure 16. Averaged amplitude of noise versus number of averages.

into one data processing technique. One must look at time trace data and spectral plots. If time series averaging is used, it is helpful to obtain plots where the data are synchronized once to the input shaft and once to the output shaft.

The Inverse FFT

The inverse FFT (IFFT) is a very powerful tool that is not often used. For example, consider a Fourier transform process starting with 1024 input samples. The FFT computes a 512 line spectrum, but to reduce the danger of aliasing, generally only 400 lines are kept and displayed. If all 512 lines and the 512 phase angles are kept, then the inverse FFT may be performed. The process is as follows:

FFT the 1024 time domain data samples to get 512 real spectral values and 512 imaginary spectral values. Then process that data to get 512 spectral amplitudes and 512 phase angles.

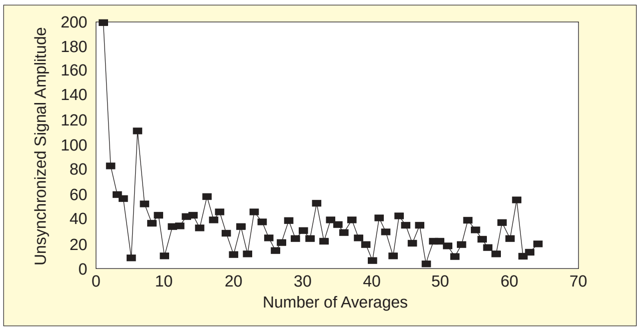

Figure 17. Averaged amplitude of unsynchronized signal vs the number of averages

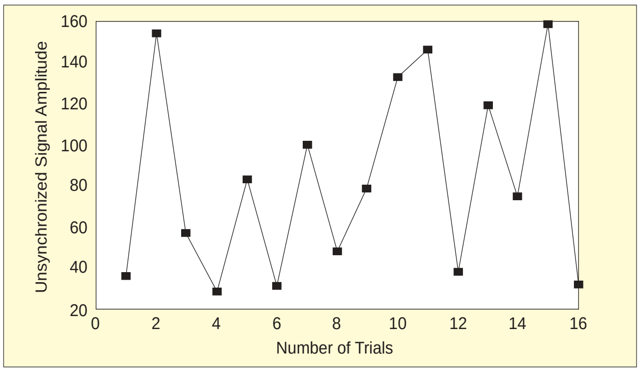

Figure 18. Sixteen trials of 8 averages of unsynchronized signal.

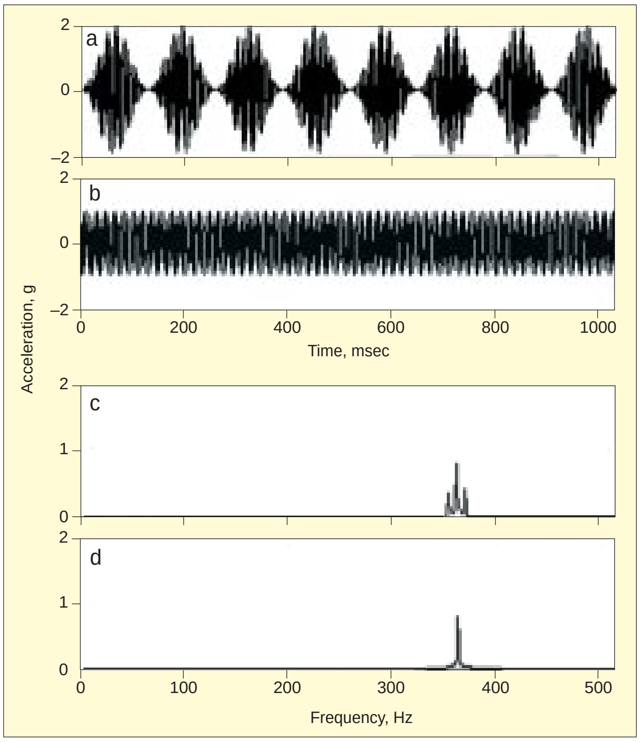

Figure 19. 128 time series averages of a modulated signal when synchronized to the input and output shaft.

This is the displayed spectrum. IFFT the 512 real spectral values and 512 imaginary spectral values to get the 1024 time data samples from the original. This would just be an interesting process except that the spectral data may be modified before the IFFT is performed.

Figure 20 (top trace) shows the original time trace of the combination of three discrete signals and noise. The next trace is the 512 line spectrum. In the third trace, all data are set to zero except in the immediate vicinity of the first discrete signal. The final trace is the IFFT of the modified signal and shows a relatively clean sine wave.

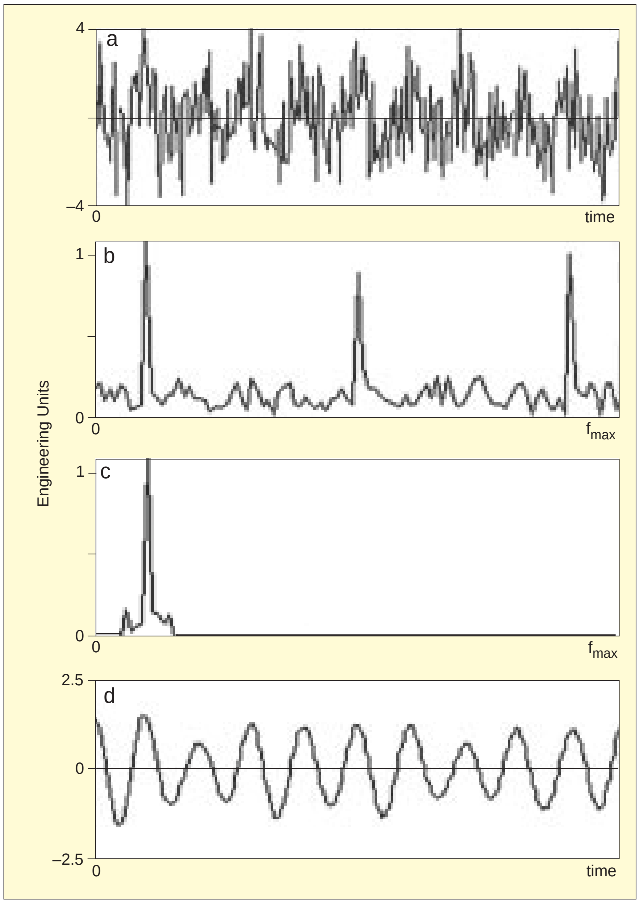

Figure 20. Using the FFT and IFFT as a filter: (a) original time trace; (b) 512 line spectrum; (c) all data set to zero except near the first discrete signal; (d) IFFT of the modified signal.

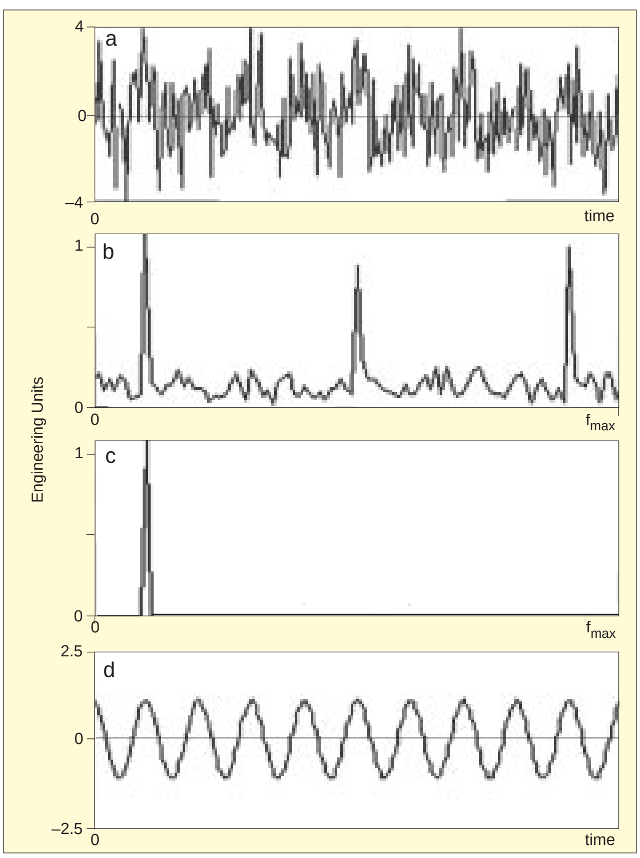

Figure 21. Using the FFT and IFFT as a filter: (a) original time trace; (b) 512 line spectrum; (c) all bins set to zero except for discrete signal; (d) IFFT of the modified signal.

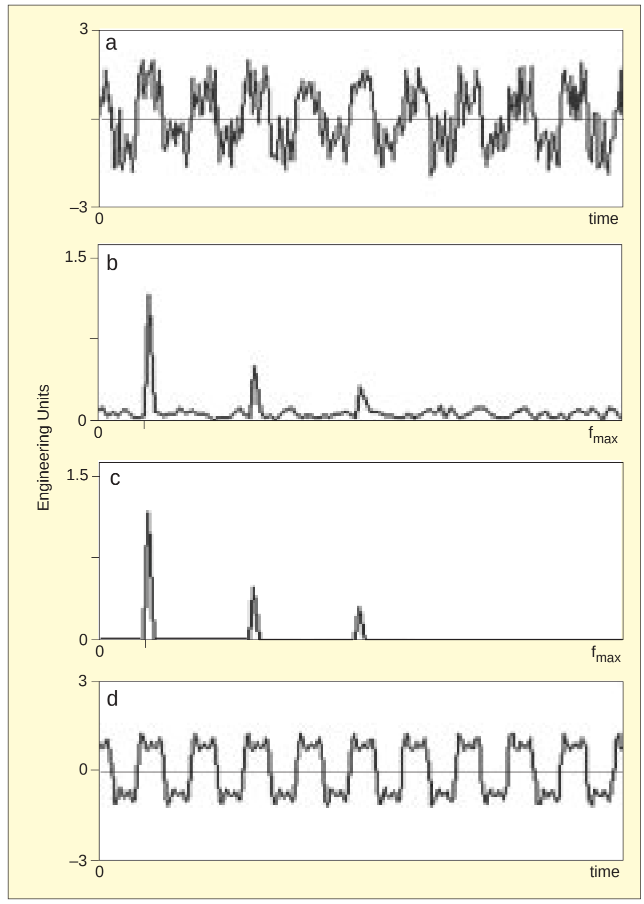

Figure 22. Removing noise from a square wave signal: (a) mixture of a square wave and noise; (b) spectrum of the top trace; (c) all bins set to zero except the three harmonics; (d) IFFT.

Figure 21 is the same process except that all FFT bins have been set to zero except for the single bin containing the discrete signal. The IFFT now produces a pure sine wave. This process may be used to filter the data using very narrow band pass filters. Realize however that if an original signal is modulated and/or is a shape other than a sine wave, then additional bins must be included to obtain the original signal. One must include sideband bins for modulated data and harmonics for non-sinusoidal wave shapes.

Figure 22 illustrates this point for a non-sinusoidal wave shape. The top trace is a mixture of a square wave (harmonics 1,3 and 5) and noise. The next trace is the spectrum of the top time trace. In the next spectrum, all bins have been set to zero except the three harmonics. The final trace is the IFFT, which now clearly shows square wave character.

Conclusion

The Fourier transform was originally conceived to determine the harmonic content of a given repetitive wave shape. Noisy signals and transient phenomena were far from Fourier's mind in his original formulation. The Fast Fourier Transform algorithm is now in common use in the vibration analysis of rotating machinery vibration. We commonly encounter signals that are not bin centered, very noisy or exhibit transient phenomena. The FFT may still be used in these cases to gain insight into the condition of the machinery, but it must be realized that one is skating on very thin ice and serious errors can easily occur for the careless. The cases discussed in this article show both sides of the question. Many errors can be encountered by carelessly following rules of thumb and not understanding the FFT process. On the other hand, careful, knowledgeable use of the FFT can provide very valuable insight about the condition of machinery in noisy and transient conditions. S V

The author can be contacted at: jlfanalys@aol.com.