



Fourier Theory & Practice, Part I: Theory (HP Product Note 54600-4)

By: Robert Witte
Hewlett-Packard Co.

Introduction:

This product note provides a brief review of Fourier theory, especially the unique behavior of the FFT. The note also describes some typical applications and provides some tips on how to get the most out of the FFT capability of the HP 54600 series scopes with an HP 54657A or HP 54658A FFT module.

Fourier Theory

Normally, when a signal is measured with an oscilloscope, it is viewed in the time domain (Figure 1a). That is, the vertical axis is voltage and the horizontal axis is time. For many signals, this is the most logical and intuitive way to view them. But when the frequency content of the signal is of interest, it makes sense to view the signal in the frequency domain. In the frequency domain, the vertical axis is still voltage but the horizontal axis is frequency (Figure 1b). The frequency domain display shows how much of the signal's energy is present at each frequency. For a simple signal such as a sine wave, the frequency domain representation does not usually show us much additional information. However, with more complex signals, the frequency content is difficult to uncover in the time domain and the frequency domain gives a more useful view of the signal.




Figure 1

- (a) A signal shown as a function of time.
- (b) A signal shown as a function of frequency.

Fourier theory (including both the Fourier Series and the Fourier Transform) mathematically relates the time domain and the frequency domain. The Fourier transform is given by:

$$V(f) = \int_{-\infty}^{\infty} v(t) e^{-j2\pi ft} dt$$

We won't go into the details of the mathematics here, since there are numerous books which cover the theory extensively (see references). Some typical signals represented in the time domain and the frequency domain are shown in Figure 2.




Figure 2: Frequency spectrum examples.

The Fast Fourier Transform

The discrete (or digitized) version of the Fourier transform is called the Discrete Fourier Transform (DFT). This transform takes digitized time domain data and computes the frequency domain representation. While normal Fourier theory is useful for understanding how the time and frequency domain relate, the DFT allows us to compute the frequency domain representation of real-world time domain signals. This brings the power of Fourier theory out of the world of mathematical analysis and into the realm of practical measurements. The HP 54600 scope with Measurement/Storage Module uses a particular algorithm, called the Fast Fourier Transform (FFT), for computing the DFT. The FFT and DFT produce the same result and the feature is commonly referred to as simply the FFT.

The HP 54600 series scopes normally digitize the time domain waveform and store it as a 4000 point record. The FFT function uses 1000 of these points (every fourth point) to produce a 500 point frequency domain display. This frequency domain display extends in frequency from 0 to $f_{eff}/2$, where f_{eff} is the effective sample rate of the time record (Figure 3a).




Figure 3

(a) The sampled time domain waveform.

(b) The resulting frequency domain plot using the FFT.

The effective sample rate is the reciprocal of the time between samples and depends on the time/div setting of the scope. For the HP 54600 series, the effective sample rate is given by:

$$f_{\text{eff}} = \frac{\text{record length}}{10 * \text{time/div}} = \frac{1000}{10 * \text{time/div}} = \frac{100}{\text{time/div}}$$

So for any particular time/div setting, the FFT produces a frequency domain representation that extends from 0 to $f_{\text{eff}}/2$ (Figure 3b). When the FFT function is active, the effective sample rate is displayed when the time/div knob is turned or the $\pm$ key is pressed. Note that the effective sample rate for the FFT can be much higher than the maximum sample rate of the scope. The maximum sample rate of the scope is 20 MHz, but the random-repetitive sampling technique places samples so precisely in time that the sample rate seen by the FFT can be as high as 20 GHz.

The default frequency domain display covers the normal frequency range of 0 to $f_{\text{eff}}/2$. The Center Frequency and Frequency Span controls can be used to zoom in on narrower frequency spans within the basic 0 to $f_{\text{eff}}/2$ range of the FFT. These controls do not affect the FFT computation, but instead cause the frequency domain points to be replotted in expanded form.

Aliasing

The frequency $f_{\text{eff}}/2$ is also known as the folding frequency. Frequencies that would normally appear above $f_{\text{eff}}/2$ (and, therefore, outside the useful range of the FFT) are folded back into the frequency domain display. These unwanted frequency components are called aliases, since they erroneously appear under the alias of another frequency. Aliasing is avoided if the effective sample rate is greater than twice the bandwidth of the signal being measured.

The frequency content of a triangle wave includes the fundamental frequency and a large number of odd harmonics with each harmonic smaller in amplitude than the previous one. In Figure 4a, a 26 kHz triangle wave is shown in the time domain and the frequency one. Figure 4b shows only the frequency domain representation. The leftmost large spectral line is the fundamental. The next significant spectral line is the third harmonic. The next significant spectral line is the fifth harmonic and so forth. Note that the higher harmonics are small in amplitude with the 17th harmonic just visible above the FFT noise floor. The frequency of the 17th harmonic is $17 \times 26 \text{ kHz} = 442 \text{ kHz}$, which is within the folding frequency of $f_{\text{eff}}/2$, (500 kSa/sec) in Figure 4b. Therefore, no significant aliasing is occurring.




Figure 4a:
The time domain and frequency domain displays of a 26 kHz triangle wave.




Figure 4b:
Frequency spectrum of a triangle wave.




Figure 4c: With a lower effective sample rate, the upper harmonics appear as aliases.




Figure 4d: With an even lower effective sample rate, only the fundamental and third harmonic are not aliased.

Figure 4c shows the FFT of the same waveform with the time/div control turned one click to the left, resulting in an effective sample rate of 500 kSa/sec and a folding frequency of 250 kSa/sec. Now the upper harmonics of the triangle wave exceed the folding frequency and appear as aliases in the display. Figure 4d shows the FFT of the same triangle wave, but with an even lower effective sample rate (200 kSa/sec) and folding frequency (100 kSa/sec). This frequency plot is severely aliased.

Often the effects of aliasing are obvious, especially if the user has some idea as to the frequency content of the signal. Spectral lines may appear in places where no frequency components exist. A more subtle effect of aliasing occurs when low level aliased frequencies appear near the noise floor of the measurement. In this case the baseline can bounce around from acquisition to acquisition as the aliases fall slightly differently in the frequency domain.

Aliased frequency components can be misleading and are undesirable in a measurement. Signals that are bandlimited (that is, have no frequency components above a certain frequency) can be viewed alias-free by making sure that the effective sample rate is high enough. The effective sample rate is kept as high as possible by choosing a fast time/div setting. While fast time/div settings produce high effective sample rates, they also cause the frequency resolution of the FFT display to degrade.

If a signal is not inherently bandlimited, a lowpass filter can be applied to the signal to limit its frequency content (Figure 5). This is especially appropriate in situations where the same type of signal is measured often and a special, dedicated lowpass filter can be kept with the scope.




Figure 5: A lowpass filter can be used to band limit the signal, avoiding aliasing.

Leakage*

The FFT operates on a finite length time record in an attempt to estimate the Fourier Transform, which integrates over all time. The FFT operates on the finite length time record, but has the effect of replicating the finite length time record over all time (Figure 6). With the waveform shown in Figure 6a, the finite length time record represents the actual waveform quite well, so the FFT result will approximate the Fourier integral very well.




Figure 6

(a) A waveform that exactly fits one time record.

(b) When replicated, no transients are introduced.

However, the shape and phase of a waveform may be such that a transient is introduced when the waveform is replicated for all time, as shown in Figure 7. In this case, the FFT spectrum is not a good approximation for the Fourier Transform.




Figure 7

(a) A waveform that does not exactly fit into one time record.

(b) When replicated, severe transients are introduced, causing leakage in the frequency domain.

Since the scope user often does not have control over how the waveform fits into the time record, in general, it must be assumed that a discontinuity may exist. This effect, known as LEAKAGE, is very apparent in the frequency domain. The transient causes the spectral line (which should appear thin and slender) to spread out as shown in Figure 8.

* For a related experiment, see "Teaching Math on an Oscilloscope, Part 3".




Figure 8

Leakage occurs when the normally thin spectral line spreads out in the frequency domain.

The solution to the problem of leakage is to force the waveform to zero at the ends of the time record so that no transient will exist when the time record is replicated. This is accomplished by multiplying the time record by a WINDOW function. Of course, the window modifies the time record and will produce its own effect in the frequency domain. For a properly designed window, the effect in the frequency domain is a vast improvement over using no window at all.¹ Four window functions are available in the HP 54600 scopes: Hanning, Flattop, Rectangular and Exponential.

The Hanning window provides a smooth transition to zero as either end of the time record is approached. Figure 9a shows a sinusoid in the time domain while Figure 9b shows the Hanning window which will be applied to the time domain data. The windowed time

domain record is shown in Figure 9c. Even though the overall shape of the time domain signal has changed, the frequency content remains basically the same. The spectral line associated with the sinusoid spreads out a small amount in the frequency domain as shown in Figure 10.2 (Figure 10 is expanded in the frequency axis to show clearly the shape of the window in the frequency domain.)




Figure 9

- (a) The original time record.
- (b) The Hanning Window.
- (c) The windowed time record.

The shape of a window is a compromise between amplitude accuracy and frequency resolution. The Hanning window, compared to other common windows, provides good frequency resolution at the expense of somewhat less amplitude accuracy.

The FLATTOP window has fatter (and flatter) characteristic in the frequency domain, as shown in Figure 11. (Again, the figure is expanded in the frequency axis to show clearly the effect of the window.) The flatter top on the spectral line in the frequency domain produces improved amplitude accuracy, but at the expense of poorer frequency resolution (when compared with the Hanning window).




Figure 10

The Hanning Window has a relatively narrow shape in the frequency domain.




Fig. 11

The flattop window has a wider, flat-topped shape in the frequency domain.

The Rectangular window (also referred to as the Uniform window) is really no window at all; all of the samples are left unchanged. Although the uniform window has the potential for severe leakage problems, in some cases the waveform in the time record has the same value at both ends of the record, thereby eliminating the transient introduced by the FFT. Such waveforms are called SELF-WINDOWING. Waveforms such as sine bursts, impulses and decaying sinusoids can all be self-windowing.

A typical transient response is shown in Figure 12a. As shown, the waveform is self-windowing because it dies out within the length of the time record, reducing the leakage problem.




Figure 12

- (a) A transient response that is self-windowing.**
- (b) A transient response which requires windowing.**
- (c) The exponential window.**
- (d) The windowed transient response.**

If the waveform does not dissipate within the time record (as shown in Figure 12b), then some form of window should be used. If a window such as the Hanning window were applied to this waveform, the beginning portion of the time record would be forced to zero. This is precisely where most of the transient's energy is, so such a window would be inappropriate.

A window with a decaying exponential response is useful in such a situation. The beginning portion of the waveform is not disturbed, but the end of the time record is forced to zero. There still may be a transient at the beginning of the time record, but this transient is not introduced by the FFT. It is, in fact, the transient being measured. Figure 12c shows the exponential window and Figure 12d shows the resulting time domain function when the exponential window is applied to Figure 12b. The exponential window is inappropriate for measuring anything but transient waveforms.

¹ *The effect of a time domain window in the frequency domain is analogous to the shape of the resolution bandwidth filter in a swept spectrum analyzer.*

² *The shape of a perfect sinusoid in the frequency domain with a window function applied is the Fourier transform of the window function.*

Selecting a Window

Most measurements will require the use of a window such as the Hanning or Flattop windows. These are the appropriate windows for typical spectrum analysis measurements. Choosing between these two windows involves a tradeoff between frequency resolution and amplitude accuracy. Having used the time domain to explain why leakage occurs, now the user should switch into frequency domain thinking. The narrower the passband of the window's frequency domain filter, the better the analyzer can discern between two closely spaced spectral lines. At the same time, the amplitude of the spectral line will be less certain. Conversely, the wider and flatter the window's frequency domain filter is, the more accurate the amplitude measurement will be and, of course, the frequency resolution will be reduced. Choosing between two such window functions is really just choosing the filter shape in the frequency domain.

The rectangular and exponential windows should be considered windows for special situations. The rectangular window is used where it can be guaranteed that there will be no leakage effects. The exponential window is for use when the input signal is a transient.³

³ *For more information on windows, see references 2 and 6.*

References:

1. Brigham, E. Oran, The Fast Fourier Transform and Its Applications, Englewood Cliffs, NJ: Prentice-Hall, Inc., 1988.
2. Hewlett-Packard Company. "Fundamentals of Signal Analysis", Application Note 243, Publication Number 5952-8898, Palo Alto, CA, 1981.
3. McGille Clare D. and George R. Cooper. Continuous and Discrete Signal and System Analysis. New York: Holt, Rhinehart and Winston, Inc., 1974.
4. Ramirez, Robert W. The FFT Fundamentals and Concepts. Englewood Cliffs, NJ: Prentice-Hall, Inc., 1985.
5. Stanley, William D., Gary R. Dougherty and Ray Dougherty. Digital Signal Processing, 2nd ed. Reston, VA: Reston Publishing Company, Inc., 1984.
6. Witte, Robert A. Spectrum and Network Measurements. Englewood Cliffs, NJ: Prentice-Hall, Inc., 1991.