


FFT Applications for TDS Oscilloscopes



Introduction

The ability to analyze the frequency domain with the Fast Fourier Transform (FFT) has opened up a whole new set of opportunities in digital oscilloscopes. However, the user needs to understand the capabilities and constraints of both the measurement techniques and the instruments that perform them. With the aid of practical examples, this paper helps the user understand the theory behind the techniques, to identify the measurements that the FFT can perform, and to set up the oscilloscope for optimum results.

1.1 FFT Terminology in the TDS Oscilloscope Environment

Each acquisition of a waveform by the oscilloscope is referred to as an acquired **record** of data. The **FFT time record** is the portion of the acquired time domain record that will be input to the FFT. The FFT time record always starts at the beginning of the acquired record of data. The **zero phase reference point** for the FFT is always at the middle of the FFT time record.

Figure 1.1 shows the FFT time record for each record length selected in the TDS oscilloscope. The FFT length equals the record length for record lengths up to 5k. For record lengths greater than or equal to 15k, the FFT is 10k in length.




Figure 1.1 Relationship of the FFT time record to the acquired data record.

Figure 1.2 shows how the FFT time record is transformed into a frequency domain record of 1/2 the length. Each point in the FFT frequency domain record may be referred to as a bin.

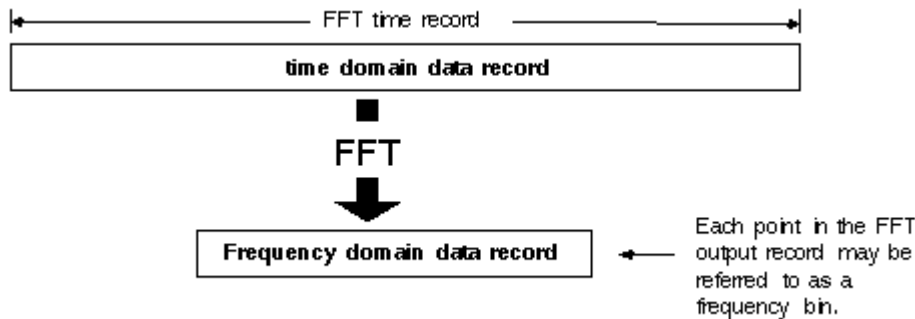


Figure 1.2 The frequency domain record is 1/2 the length of the FFT time record.

The change in frequency from one bin to the next is shown by equation 1.1

$$\Delta F = \frac{\text{sample rate}}{\text{FFT length}}$$

where:

ΔF is the frequency difference between bins

1.2 FFT Frequency Range

The range of frequencies covered in the output record from the FFT is 0 to 1/2 the sample rate of the acquired data record. For example, a sample rate of 20 MS/s (megasamples per second) would give an FFT range of 0 to 10 MHz. The value of 1/2 MS/s sample rate is often referred to as the Nyquist point.

Although the non-interleaved sample rate of the TDS 540 ADC (analog-to-digital converter) is always 250 MHz, the sample rate of the acquired data record may vary from 5 S/s to 100 GS/s depending on the sec/div setting of the time base.

1.3 DC Correction in the FFT Output Record

The output of a standard FFT computation yields a DC value which is twice what it should be relative to the other frequencies. When **vertical position** and **offset** are set to zero, the displayed output of the FFT on TDS oscilloscopes is corrected to show the true value for the DC component of the input signal. The DC level is also corrected according to the window used. (See Section 1.20 for information about windows.) Keep in mind that 1 VDS equals 1 Vrms, and the display is in dB with respect to 1 Vrms or the display is in linear Vrms.

1.4 Impulse Testing to Obtain Magnitude and Phase Response

Impulse testing can be used for rapid determination of the magnitude and phase response of a system as a function of frequency. In loudspeaker testing, this method allows the phase and magnitude response to be computed before the sound from the speaker has a chance to reflect from the walls, eliminating the need for an expensive anechoic test chamber. An ideal impulse has infinite amplitude, infinitely small width, and an area of one. The Fourier transform of an impulse is 1. This means that the impulse contains all frequencies at equal amplitude with a zero phase reference point starting at the time position of the impulse. In other words, each harmonic starts a cycle at the impulse time position. An impulse may be applied to the input of a linear time invariant system, and the FFT of the impulse response of that system will yield an approximation of the magnitude and phase response. This is an advantage over Bode plots because the magnitude and phase are displayed in real-time, allowing for real-time adjustments to the system under test while observing effects on phase and magnitude. Figure 1.3 shows a test setup for impulse testing.

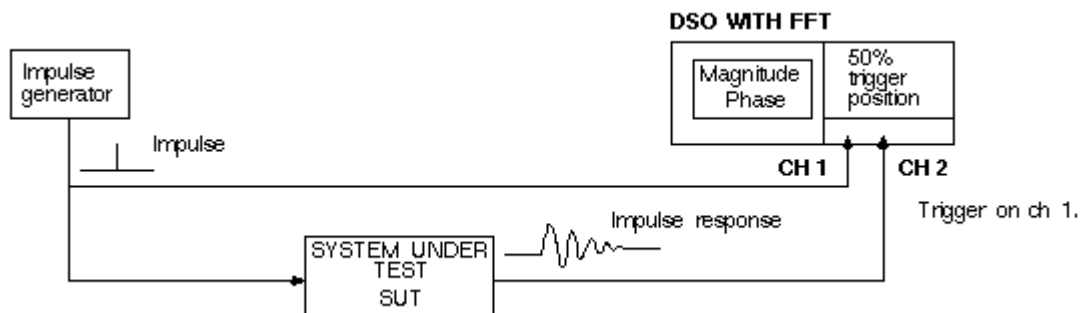


Figure 1.3 Test configuration for impulse response testing of a system.

1.5 Setting up the Pulse Generator to Approximate an Impulse

Because the ideal impulse has infinite amplitude, infinitely small width, and an area of one, the pulse generator must be set to provide a usable approximation of an ideal impulse. The generator should be adjusted to output a pulse with as much amplitude as possible without causing the system under test to clip. If the amplitude is not large enough, or the width is too narrow, it will not contain sufficient energy.

If the pulse amplitude is too large, it will cause the system under test to clip. The width of the pulse should be less than one sample interval at the sample rate of the oscilloscope's current time base setting. If it is too wide, the response will be distorted. These details suggest that the sample rate on the oscilloscope should be as low as possible in order to cover the frequency range of the system under test. (The sample

rate is displayed in the upper left-hand corner of the screen on TDS oscilloscopes.) The pulse repetition rate should be high enough to provide a fast update rate to the display if real-time evaluation of the phase and magnitude is necessary. Also, it should be low enough that only one pulse occurs in the acquired record and the impulse response fully decays to zero before the next oscilloscope acquisition begins.

1.6 Setting up the Scope for Impulse Testing

Setting up the Vertical Sensitivity:

- The oscilloscope should be set to trigger on Channel 1 (Ch 1), which is connected to the pulse generator.
- The impulse response from the system under test goes into Ch 2.
- Adjust the vertical sensitivity of Ch 2 (make sure **Zoom** is off) so that the response covers as many divisions as possible without clipping. This will give maximum vertical resolution to the signal. The signal may go one division above and one below the viewable area on the screen before it clips.
- To determine if clipping occurs, turn on **Peak to Peak** measurement on Ch 2. A message is displayed on the screen if the signal clips. Turn off the **Peak to Peak** measurement after setting the vertical gain.
- Now turn on **Zoom** -- with Ch 2 selected -- and adjust the vertical scale knob to reduce the size of the impulse response waveform on the display. (**Zoom** affects only the display scaling.)

Setting the Horizontal Trigger Position:

- Turn Zoom off. Press the HORIZONTAL MENU button and select Trigger Position.
- Press the Set to 50% selection in the vertical menu, thus aligning the impulse and the FFT zero phase reference.

The TDS oscilloscopes have an FFT zero phase reference at 50% of the FFT time record width. The FFT time record is 10k long and always starts at the beginning of the record. (See Figure 1.1). To obtain the correct phase response for the system under test, the impulse must occur at the same position as the FFT zero phase reference in the time record. Since the FFT length is equal to the record length, this reference is at a 50% pretrigger setting.

Note: There are two exceptions to this when using the TDS 540: for a 15k record length, set the Trigger Position to 33.33%; for a 50k record length, set it to 10%. (The Trigger Position control does not have sufficient resolution to be set exactly to 33.33% for the 15k record length. To set this, turn Zoom back on.)

Setting the Horizontal Time/Div:

The frequency range covered by the FFT will be from 0 Hz to 1/2 of the sample rate shown in the upper left part of the display.

- With Zoom turned off and Ch 2 selected, adjust the Horizontal Scale until the optimum sample rate is obtained. The sample rate needs to be high enough that high-frequency components in the impulse response are not aliased. (Refer to Section 1.13 for an explanation of aliasing.)
- After selecting the desired sample rate, choose a record length which will allow the entire impulse response to be captured in the acquired data record.

Selecting an Acquisition Mode:

Start in Sample mode

- After the FFT has been turned on, Hi Res mode may be needed to remove noise which is not synchronized with the trigger. (Hi Res mode affects the frequency response of the signal being acquired. Refer to Section 1.18 for details on the Hi Res mode frequency response.)
- If this is not desired and the impulse is repetitive, Average mode can be used to remove noise from the signal at the cost of slower display response.

Note: Peak Detect and Envelope modes can add significant distortion to the FFT results and are not recommended for this application.

Turning on the FFT Magnitude Waveform:

- Press the MORE button and select FFT from the bottom menu. If FFT is not in the menu, press the Change Math waveform definition selection in the Vertical Menu and select FFT.
- Use the Vertical Menu buttons to assign the FFT to Ch 2.
- Select a Rectangular Window function and set the vertical option to dB with respect to 1 Vrms.
- Press the OK button to turn on the FFT waveform.

Turning on the Phase Waveform:

To turn on the phase waveform, a second FFT must be turned on.

- Press the MORE button and select one of the three Math Waveform buttons in the bottom row. Do not select the FFT math waveform which was turned on in the previous step.
- Press the Change Math definition button in the vertical menu and set up the definition: FFT, vertical phase degrees, rectangular window, phase suppression to -100 dB. The phase suppression may be adjusted later to remove noise phase from the display. (Refer to Section 1.14)




Figure 1.4 Oscilloscope display of impulse testing of a parametric equalizer for a sound system.

Figure 1.4 shows the results of impulse testing configured as in Figure 1.3. In this example, the system under test was a parametric equalizer for professional sound systems. Its filter response can be adjusted for center frequency, bandwidth, and passband gain or attenuation.

A key point to remember for impulse testing is that, in order to obtain the desired phase display, the impulse must occur at the zero phase reference point in the FFT time record.

Because the sample clock in the oscilloscope is not synchronous with the trigger point, the samples can be positioned relative to the trigger by plus or minus one sample interval. At the higher frequency in the FFT output display, there are only two samples per cycle. This means that the phase error at this frequency could be off by 0 to plus 180 degrees. At the beginning of the record, the phase error will be very small, since there are many samples per cycle at these frequencies. Therefore, **the phase display is more accurate at the beginning of the record than at the end of the record.**

1.7 An Alternative to Impulse Response Testing: Step Response

The previous section has demonstrated how the phase and magnitude response of a system may be obtained by taking an FFT of the system's impulse response. If the testing is for a very high-frequency system, a good impulse generator may not be available. An alternative way to obtain the phase and magnitude response is to apply a step function to the input of the system: take the derivative of the step response at the output of the system and save it to **ref1** memory in the TDS oscilloscope (the derivative of the step response equals the impulse response); then set up the oscilloscope to do an FFT of the waveform in **ref1**.

Make sure the rising edge of the step is positioned at the zero phase reference point in the FFT time record. This is at a 50% pretrigger setting for record lengths up to and including 5k. (Sections 1.8 and 2.3 describe more about the zero phase reference point.)

This procedure does not produce a live display. However, if the derivative waveform is selected in the **MORE** menu and if the **SAVE/RECALL WAVEFORM** menu is displayed on the screen, each update of the display requires pressing only a single button.




Figure 1.5 Obtaining the magnitude and phase response of a filter by taking the FFT of the derivative of the filter's step response.

1.8 Measuring the Harmonic Content of Signals

Another major use of the FFT is to observe the frequency content of a signal. Figure 1.6 illustrates an FFT of a sine wave. The phase is shown in the bottom trace. Since the zero phase reference point for the FFT is in the center of the FFT time domain record, the value of the displayed phase will depend on where this phase reference is located with respect to the positive peak of the harmonic of interest. Its position may be set by adjusting the amount of pretrigger in the Horizontal Menu or by changing the trigger level. The hand-drawn dotted sine cycle in Figure 1.6 shows how the phase reference for a given frequency bin always starts with its peak at the center of the FFT time record and is the same frequency as the bin of interest.

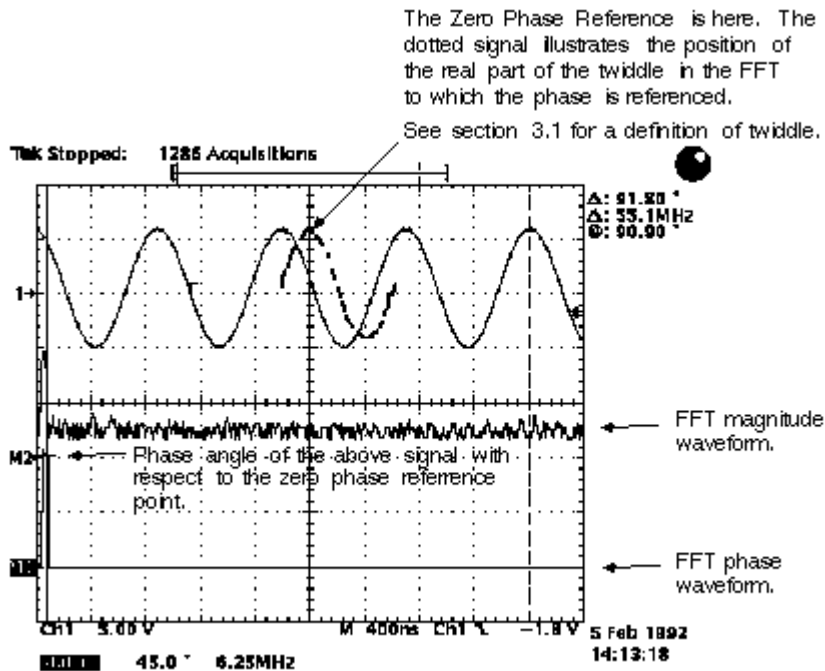


Figure 1.6 Measuring the frequency and phase of a sine wave.

For the above example, the entire 500 point data record is shown on the screen. Notice that the acquired signal leads the zero phase reference point by approximately 90 degrees. The paired cursor is attached to the phase waveform and reads out the phase angle. Since the FFT time record is the same as the acquired data record, we could move the FFT time record with respect to the acquired data by adjusting the amount of pretrigger in the Horizontal Menu of the oscilloscope. This position can be adjusted in finer detail by adjusting the trigger level, when the slope of the signal is not too large. **This positioning allows the user to set the phase of the fundamental of the signal to zero degrees.** (See Figure 1.6 and Figure 1.7.)

Figure 1.7 illustrates how the FFT time record and zero phase reference point may be moved with respect to the trigger position.

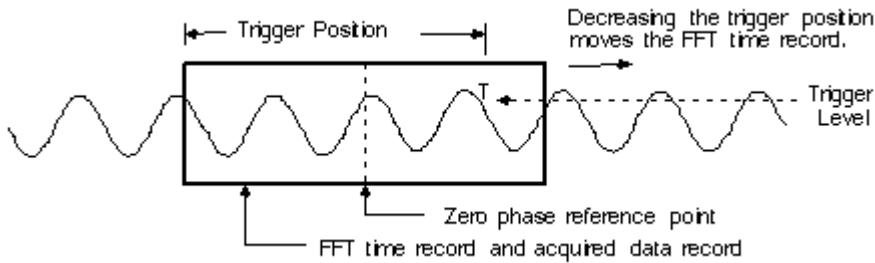


Figure 1.7 Positioning the FFT time record with respect to the trigger.

Zero phase for the fundamental of the above waveform occurs when the FFT time record aligns with the acquired signal such that the zero phase reference point aligns with the positive peak of one cycle.

Figure 1.8 shows the harmonic content of a square wave where the FFT time record was positioned to obtain a fundamental at approximately zero degrees. Notice that the positive half cycle is centered about the zero phase reference point.

Note that aliases folded back into the magnitude display are visible in Figure 1.8. More details on how to identify aliases are given in [section 1.13](#).

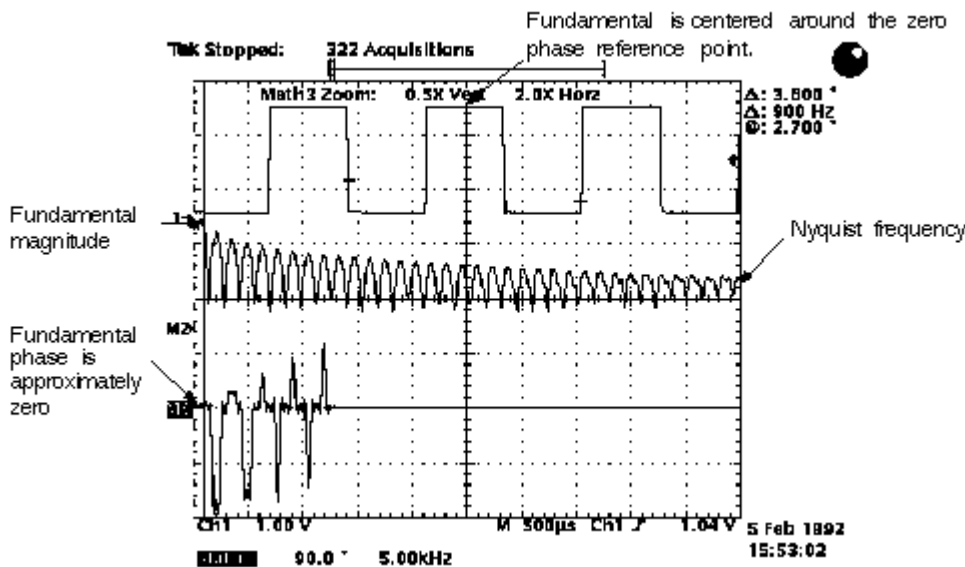


Figure 1.8 Measuring the harmonic content of a square wave.

1.9 Analyzing Line Current Harmonics

The International Standard CEI/IEC 555 deals with disturbances in power supply systems caused by household appliances and similar electrical equipment. [Section 1.11](#) deals with harmonics. A test setup for observing these harmonics is shown below in Figure 1.9 (Test setup courtesy of Carl Matson, Group Engineering, Tektronix.)

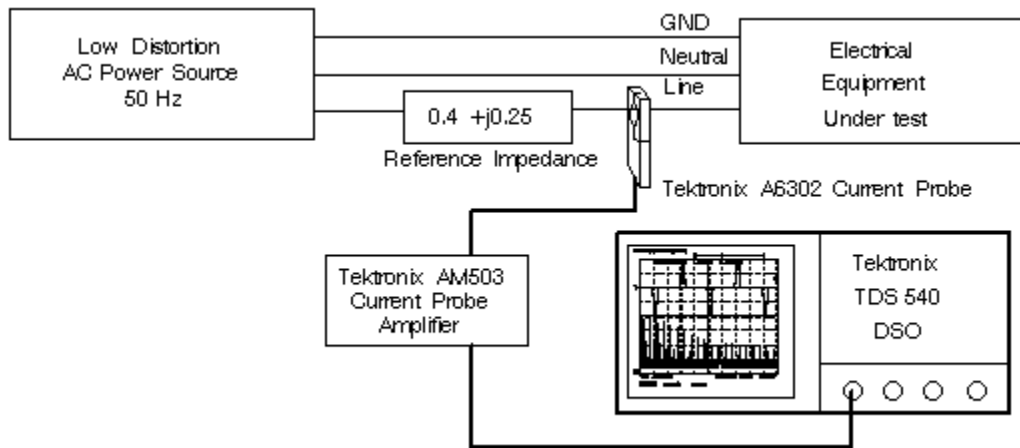


Figure 1.9 Test setup for analyzing line current harmonics.

For this example a test was made on a piece of electrical equipment containing a switching power supply. An FFT of the line current waveform was taken using a sample rate of 10 kS/s, a record length of 5000, dBVrms on the vertical scale, and a Hanning window. The results are shown below in Figure 1.10. The frequency range covered on the display is 0.0 to 2.0 kHz.




Figure 1.10 Analyzing line current disturbance of electrical equipment.

1.10 FFT Boosts Productivity for Power Supply Designer

(Thanks for this application to Carl Matson, Design Engineer at Group Engineering, Tektronix. When Carl was designing power supplies, he found the TDS with FFT to be a time-saving tool for solving noise problems.)

Circuit and power supply designers, working with low frequencies, constantly encounter signals that show up as chaotic time domain displays (see lower display in Figure 1.12.) In spite of using sophisticated time domain test equipment to view these signals, very little can be learned about them. Designers often make trial-and-error changes to their circuits, spending hours trying to locate the noise sources within their designs. This example tells how a TDS oscilloscope with FFT capability can help power supply designers quickly locate the source of a problem.




Figure 1.11 The voltage reference circuit and test setup used in Section 2.

Figure 1.11 illustrates a simple precision voltage reference circuit. A 15 VDC supply provides bias through a 10 kilohm bias resistor. A shunt voltage regulating device gave a stable DC output at V_{out} . A $0.01 \mu\text{F}$ capacitor, in parallel across the voltage reference device, filtered high-frequency components.

Ch 3 of the TDS oscilloscope was connected to the V_{out} test point to monitor the output voltage. **Sample** mode (the default acquisition mode) was used at a sample rate of 100 kS/s. Coupling was AC input.

Figure 1.12 shows that the Ch 3 time domain display was chaotic and far from being a stable ripple-free output. Several undesirable non-DC components are visible at V_{out} . Faced with this evidence, a designer couple spend hours trying to locate the cause of the noise. An engineer's first response typically would be to add more capacitance to improve filtering.




Figure 1.12

The lower waveform in Figure 1.13 shows the time domain signal at V_{out} after changing the capacitance value to $0.22 \mu\text{F}$ to increase the $C_{effective}$. The new response at V_{out} showed a small reduction in high-frequency components, but not enough to qualify as a good design. V_{p-p} at Ch 3 was marginally reduced from 4.96 mV to 3.28 mV.

For a better look at the signal, a real-time FFT was activated on Ch 3 and the signal was examined in the frequency domain. (TDS setup: 20 MHz bandwidth limit, Blackman-Harris Window, 5k points/div.)

The large bump at around the 5 kHz region of the spectrum told the designer that non-DC components, albeit at very low frequencies, were making their way to V_{out} . This information about the frequency content was possible only in the frequency domain. Once the presence of low-frequency components were identified using an FFT, the designer recognized the futility of trying to compensate by making changes in capacitance, because capacitive filters reduce only high frequencies.

The simple FFT spectrum in Figure 1.13 prompted the designer to make changes to other elements of the circuit, rather than continue with the trial-and-error method of tweaking various parameters. To solve this problem, Carl decided to change the voltage reference device itself.




Figure 1.13

Figure 1.14 shows time domain and FFT plots with the new voltage reference device in the circuit. Ch 3 Vp-p dropped from 3.28 mV to a desirable value of 640 μ V. The FFT spectral amplitude exhibits no peak as it did in Figure 1.13.




Figure 1.14 Time domain and FFT plots, with the changed voltage reference device.

Some Benefits of a TDS Scope with FFT: Time and real-time frequency domains can be shown in one display. Time display highlights the problem; the FFT allows insight into the nature of the problem.

Real-time FFT shows changes in circuit response in tandem with changing circuit parameters. A common heterodyne swept spectrum analyzer would take much longer to display the amplitude peaks in the low-frequency region of the FFT spectrum.

1.11 Using the FFT to Observe Distortion in Signals

A pure sine wave with no harmonics can be fed into a system under test. Any measurable distortion in the system will add harmonics that will show up in the FFT of the system's output.

Figure 1.11 shows a 20 MHz sine wave. Distortion in this particular signal shows as a second harmonic at 40 MHz, which is 45 dB below the fundamental, and a third harmonic at 60 MHz, which is 60 dB below the fundamental.

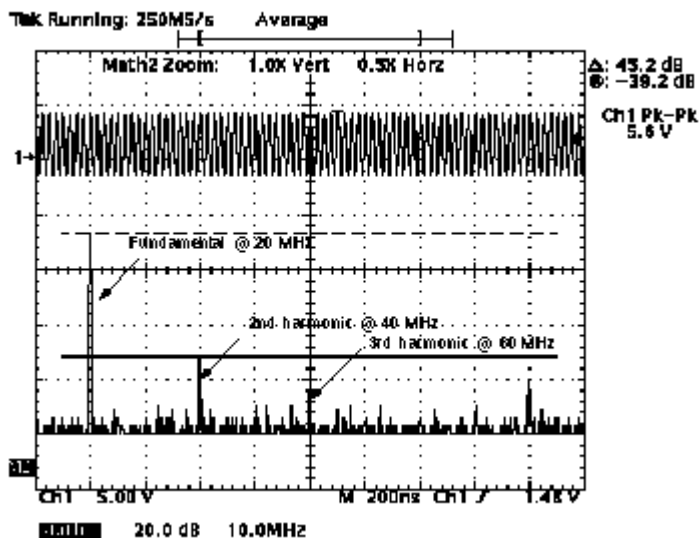


Figure 1.15 Distortion shows up as undesired harmonics.

1.12 Using the FFT to Identify Noise Sources

Noise in mixed digital/analog systems can be easily observed with an oscilloscope in the time domain. Identifying the sources of the observed noise in the time domain can be difficult, if not impossible. An FFT of the noise will produce a display of its frequency content. These frequencies can then be associated with known frequencies of clocks, oscillators, strobos, and other devices with the system, and the necessary steps can be taken to eliminate them.

Figure 1.12 demonstrates how a 31 MHz memory strobe signal shows up in an amplifier's noise floor. (The solid vertical cursor is next to the fundamental.) Notice that the 62 MHz second harmonic of this strobe also shows up.




Figure 1.16 Identifying noise sources using the FFT

1.13 Recognizing Aliasing in the FFT

The greatest frequency that can be input into a sampler without aliasing is 1/2 of the sample frequency. It is common to have acquired signals with a fundamental frequency less than 1/2 the sample rate, but the harmonics of that signal may be greater than 1/2 the sample rate and they will alias. This shows up in the FFT as frequencies that fold back into the display. Here are three methods for identifying aliases if they occur.

1. Fast-rising edges in a waveform create many high-frequency harmonics. These harmonics typically decrease in amplitude as their frequency increases. Figure 1.17 shows how these harmonics fold back into the display at the Nyquist point and are easily identified.

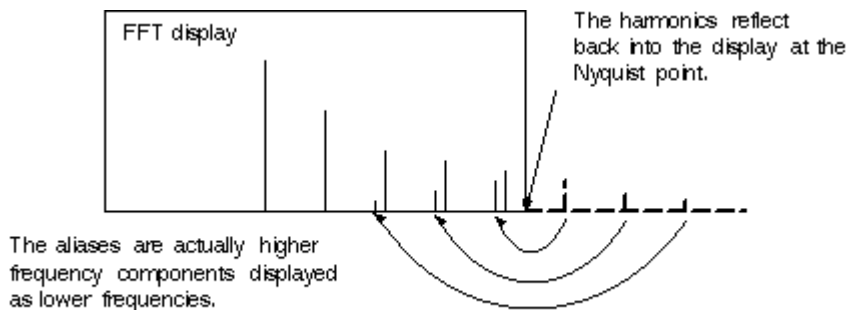


Figure 1.17 Recognizing aliases.

2. A second way to identify the aliases in the display is to select the channel on which the FFT is being applied and increase the sample rate by turning the Horizontal SCALE knob. This will increase the Nyquist frequency point and cause the aliased signals to unfold and no longer be aliased.

Figure 1.18, below, illustrates how the signal in Figure 1.17 would unfold as a result of adjusting the sample rate.




Figure 1.18 Increasing the sample rate unfolds the aliases.

3. A third method of identifying aliases can be used if the input signal's frequency can be adjusted. As the input frequency is increased, the non-aliased harmonics will move toward the right-hand side of the screen. However, the aliased harmonics will move toward the left. It is also possible for the aliases to move toward the right. As the input generator frequency is increased, the aliases move toward the left of the screen. When they reach the edge, they will reflect back into the display and begin moving toward the right again.




Figure 1.19 Change the generator frequency to identify aliases.

Although the TDS oscilloscopes do not provide anti-alias filters, the bandwidth limit filters can provide some alias attenuation in some situations. Hi Res mode also can provide some degree of alias attenuation. Also, Average and Hi Res modes can attenuate aliased noise.

1.14 Adjusting the FFT Phase Suppression

The noise signal in the FFT can have phase angles that randomly vary from $-\pi$ to π . This could make the phase display unusable. The **phase suppression control** in the **Change Waveform menu** has been provided to suppress noise phase to zero on the display. This control allows the user to select a threshold level in dB with respect to 1 Vrms. If the magnitude of the FFT is greater than this threshold, its phase angle will be displayed. If it is less than this threshold, the phase angle will be set to zero.

The threshold can be viewed easily with a cursor by turning on an FFT with dBVrms for the vertical scale. Make sure the FFT waveform is selected, then press the **CURSOR** button and choose the **horizontal bar cursor**. Move this cursor to the level at which phase suppression is desired. The dB readout of the cursor is the value which should be entered into the **phase suppression** selection on the second FFT turned on. This FFT should have **phase degrees** or **phase rad** selected for the vertical. (Make sure that both FFTs have the same window function turned on.)

1.15 Improving Oscilloscope Response When Using a Longer Record Length FFT

The longer record length FFTs slow down the oscilloscope's response. A way to improve response is to store the time domain signal in a reference memory and then select that reference memory as the FFT source. The FFT will compute only when the reference memory is changed by saving a new waveform.

1.16 Longer Record Length FFT Gives a Lower Noise Floor

The longer FFTs provide a lower noise floor. Figure 1.20 (below left) shows the 500 point FFT of a sine wave from a signal generator. Figure 1.20 (below right) shows the 5000 point FFT of the same signal, demonstrating the lower noise floor. The cursor setting and the sample rate are the same for both figures.

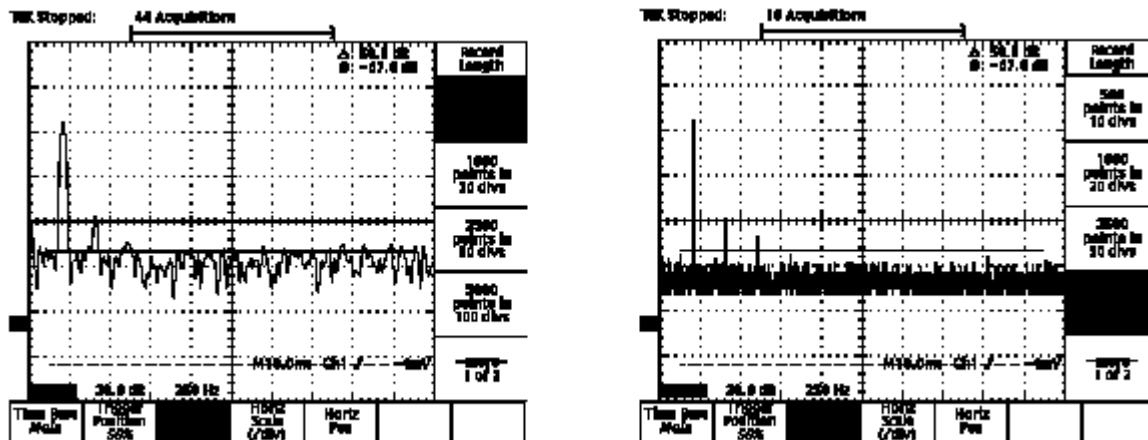


Figure 1.20 At left, a 500 point FFT. At right, a 5000 point FFT showing the lower noise floor.

1.17 Increasing FFT Length Increases the Frequency Resolution

When analyzing AM modulated signals, it may be desirable to observe the sidebands and resolve them from the fundamental. The choice of the FFT window function, the sample rate, and the FFT length are important parameters that must be adjusted to obtain the desired results. (See Sections 1.20, 1.21, 1.22, 1.23, 1.24 for an explanation of the various window functions.) The rectangular window provides the greatest frequency resolution because it has the narrowest main lobe. At a given sample rate, longer FFTs will give smaller frequency steps per FFT bin for greater frequency resolution. A higher sample rate will give more samples per cycle for better trigger stability. Figure 1.21 shows a 1 MHz carrier AM modulation by a 13 kHz sine wave. The Blackman-Harris window was used and the sample rate was 5 MS/s. The only difference between the two setups is that the record length at left is 500 points and the record length at right is 15k points with a 10k FFT. The sidebands are undetectable in the 500 point example, while they show quite clearly with the higher-resolution 15k point FFT.

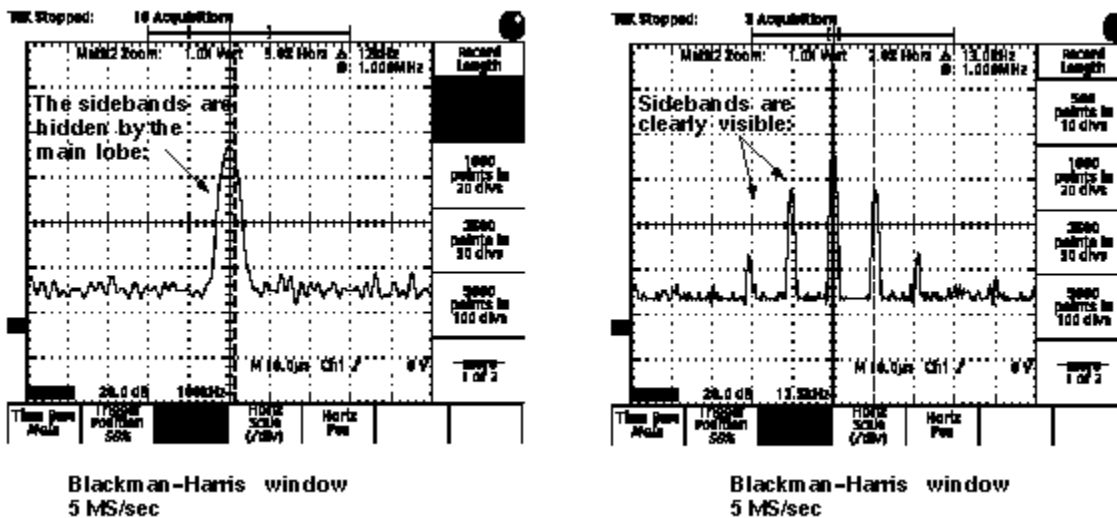


Figure 1.21 At left, AM modulated signal. 500 point FFT does not show the sidebands. At right, the same signal with a 10k FFT clearly shows the sidebands.

The sidebands are high enough in amplitude that switching to a rectangular or Hamming window would allow them to become slightly visible in the 500 point example above. As stated earlier, the disadvantage of using the longer record lengths is the longer processing time. The 500 point acquisitions and FFT provide very nice real-time update rate which is helpful when adjustments are being made to the system under test.

1.18 Hi Res Acquisition Mode Removes Noise

Hi Res mode (not available on all TDS oscilloscopes) can be used to average out random noise in repetitive or single-shot events. It also improves the vertical resolution. Custom hardware allows Hi Res mode to operate with little effect on the acquisition rate. Hi Res mode has no effect on the signal when the sample rate is 250 MS/s or greater.

The Hi Res frequency response the acquired record's sample rate is essentially the same for every time base setting for the range of frequencies from 0 to 0.5 of the sample frequency. The frequency response gradually rolls off to approximately 0.63 at the Nyquist frequency (see Figure 1.22c on the next page). Aliases are attenuated by varying amounts, depending on the sample rate, because the response for frequencies greater than 0.5 of the sample frequency varies with the time base setting. When analyzing the FFT display, this frequency response must be taken into account only if Hi Res mode is on.




Figure 1.22 Hi Res mode frequency response from 0 to 0.5 of the sample frequency.

The response shown in Figure 1.22a is for a time basing setting of 2 $\mu\text{s}/\text{div}$ and in Figure 1.22b is for 5 $\mu\text{s}/\text{div}$. On the TDS 520 and TDS 540 oscilloscopes, the ADCs are always running at 250 MS/s. For some time base settings, this rate is decimated to a lower sample rate for the acquired record. When Hi Res mode is turned on, all the samples in the decimation interval are averaged to obtain each output point. Because this is done by a dedicated circuit (a custom ASIC), only a small decrease in the acquisition rate occurs, even when the decimation interval is very large. The average is the summation of all the points in the decimation interval divided by the number of points in the interval. This is equivalent to a FIR filter with all the coefficients equal to $1/L$, where L is the number of points in the decimation interval. Since the value of L is different for each time base setting in which the oscilloscope is decimating, the frequency response of the filter changes with the value of L . This frequency response looks similar to Figure 1.22a, with the number of lobes changing with each value of L . Since there are both an input sample rate and a decimated sample rate, the response, which varies at the ADC sample rate, affects the amplitude of all signals including aliases (Figure 1.22a). Signals which are not aliased will always show the same frequency response according to Figure 1.22c.

Hi Res mode tends to average to zero signals in the time domain which are not synchronous with respect to the trigger. If it is desirable to see these random signals in the FFT's frequency domain, use Sample mode acquisition.

Here is the equation for Hi Res frequency response:

$$F_n = \frac{1}{L} \sum_{i=0}^{L-1} e^{-kn}$$

where:

$$k = \frac{2}{N} \pi \sqrt{-1}$$

$$n = 0 \dots N-1$$

N = number of points to compute

L = decimation ration = 250 MS/s sample rate.

eq(1.2)

The first $N/2$ points computed by Equation 1.2 will yield the frequency response according to Figure 1.22a. The first $N/(2L)$ points will yield the Hi Res frequency response according to Figure 1.22c.

Figure 1.23 shows an FFT taken with Sample mode acquisitions and the same with Hi Res mode acquisitions. To obtain both traces on the screen, the oscilloscope was first set to Hi Res mode and the FFT magnitude waveform was saved into **ref1**. Because different acquisition modes cannot run simultaneously, the oscilloscope was then switched to Sample mode acquisition. This figure shows how effective Hi Res mode is at removing random noise.

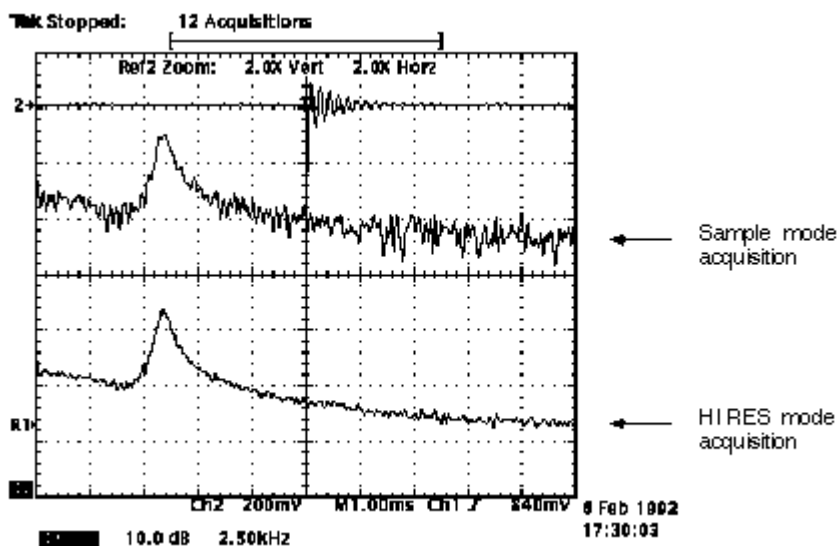


Figure 1.23 Hi Res mode acquisition for the FFT.

1.19 Average Acquisition Mode and the FFT

Both Hi Res mode and Average mode can be useful for removing aliased noise from the FFT display. If the signal is repetitive and triggered, Average mode can be used to remove unwanted noise from the display.

Average mode will attenuate signals not synchronized with the trigger. To see these signals in the FFT output, use Sample mode instead of Average mode.

1.20 Windowing -- What Does it Accomplish?

Windowing is a technique that compensates for some of the limitations of FFT analysis. The windowing operation consists of multiplying the time record by various weighting curves often named for their originators. The next few sections of this paper describe the four windows offered in the TDS oscilloscopes.

The FFT operation is equivalent to curve-fitting a sum of complex sinusoids at discrete frequencies, to the observed data. The amplitudes and phases of these sinusoids are adjusted (computed by the FFT) to fit the observed data. If the number of sinusoids selected is equal to the data length (as in the FFT), the data can be fit exactly and *uniquely*. One limitation of using the FFT as a spectral estimator is that it cannot arbitrarily choose the frequencies of the sinusoids. The sinusoids are constrained to be at integer multiples of the sampling frequency divided by the record length (Sampling Frequency/Record Length).

When a frequency component in a signal does not lie exactly on a bin, the FFT algorithm must assign a non-zero amplitude to many unrelated sinusoids in order to fit the time data exactly. The effect is often referred to as *spectral leakage*. This condition occurs more often than not when analyzing real-world signals.

Since the algorithm assigns energy to these other components, the amplitude of the main component (the closest line to the actual signal frequency) tends to be underestimated. In addition, if there are smaller amplitude signals present in the data, they can be totally obscured by the leakage components.

There are three parameters of concern when using the FFT for power spectrum estimation: Amplitude Accuracy, Frequency Resolution, and Spectral Leakage. Windowing reduces the effect of spectral leakage.

1.21 Four FFT Windows

TDS oscilloscopes offer four windows for the FFT. These are the rectangular, Hamming, Hanning, and Blackman/Harris windows. The process of *windowing* the data is shown below in Figure 1.24. The FFT time domain record is multiplied point by point with the FFT window. When the Hanning or Blackman/Harris window is used, the data point amplitudes taper to zero at the end of the record.

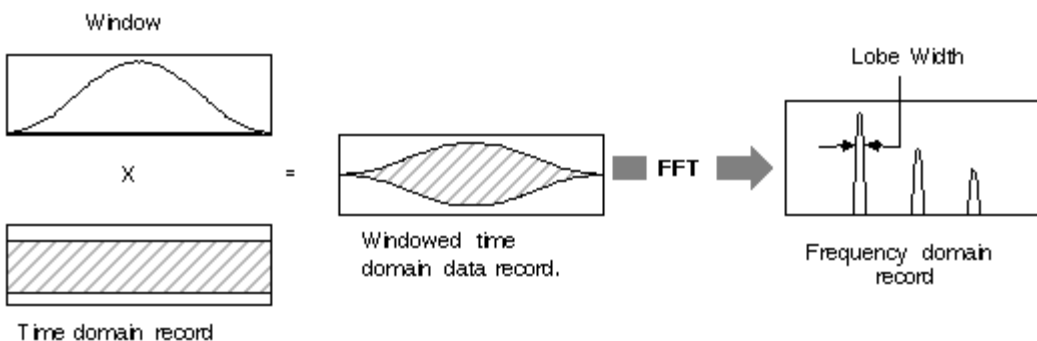


Figure 1.24 Multiplying the time domain data record by a window.

The rectangular window is equivalent to multiplying all data record points by one. It generally gives the best frequency resolution because it results in the narrowest lobe width in the FFT output record. It gains frequency resolution at the expense of amplitude accuracy if the frequency of the signal being observed has a non-integer number of cycles in the FFT time record.

The other three windows are bell-shaped, as shown in Figure 1.26. The Hanning and Blackman/Harris windows taper the data at the ends of the record to zero, while the Hamming window tapers the data to smaller values, but not to zero. The bell-shaped windows decrease the amount of energy spillover into adjacent frequency bins, increasing the amount of amplitude accuracy at the expense of decreasing the frequency resolution because of wider lobe widths.

The windows, listed in order of decreasing frequency resolution and increasing lobe width, are:

1. Rectangular
2. Hamming
3. Hanning
4. Blackman/Harris

When using windows, be sure that the most interesting parts of the signal in the time domain record are positioned in the center of the window so the tapering does not cause significant errors.

The best window for a given application depends on various factors. Figure 1.25 shows the frequency domain characteristics for each of the four windows. The window data was quantized to 16 bits. The frequency domain diagrams illustrate the way energy spillover affects the bins adjacent to the main lobe. The Hamming window has lower side lobes adjacent to the main lobe than the Hanning window. However, side lobes farther from the main lobe are lower for the Hanning. The characteristics of the signal under observation will dictate whether a Hamming or Hanning window is the better choice.

Refer to Section 1.24 for the equations used to compute each window.

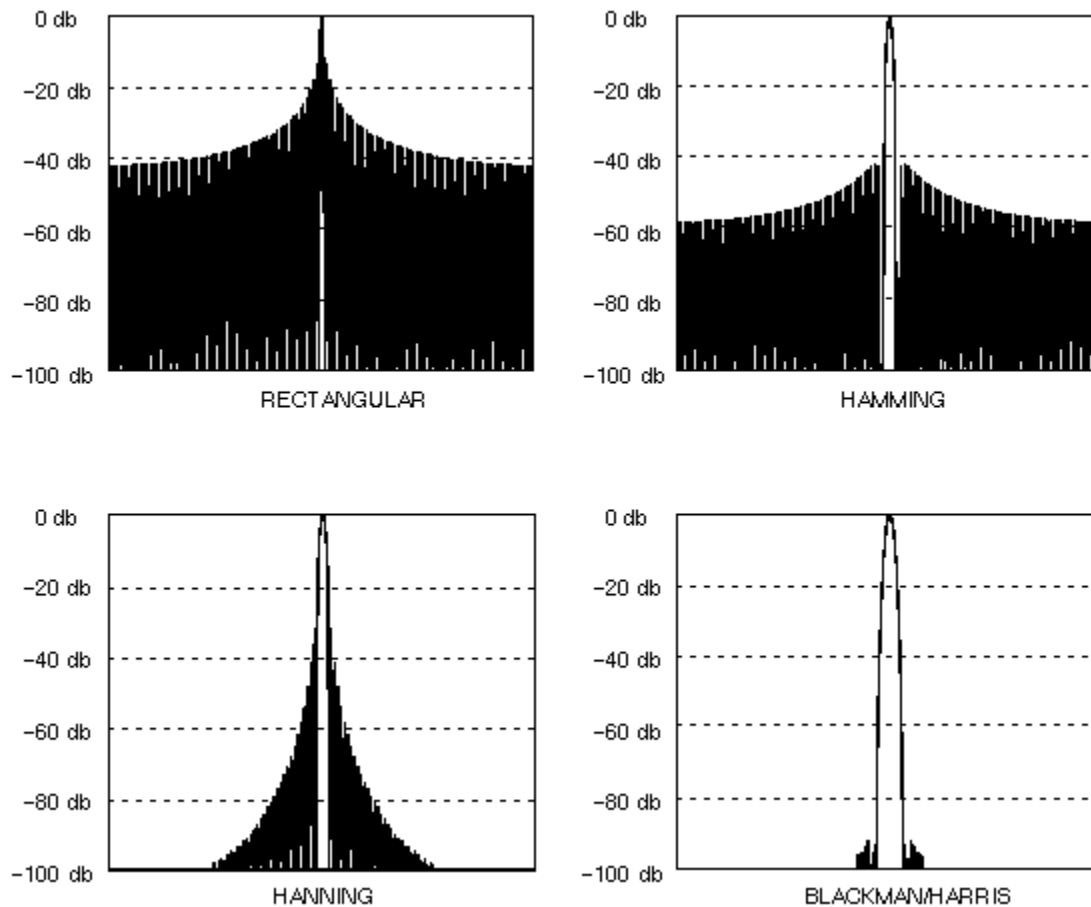


Figure 1.25 Frequency domain shapes for the various windows.

The time domain shape for each of the four windows is shown below in Figure 1.26.

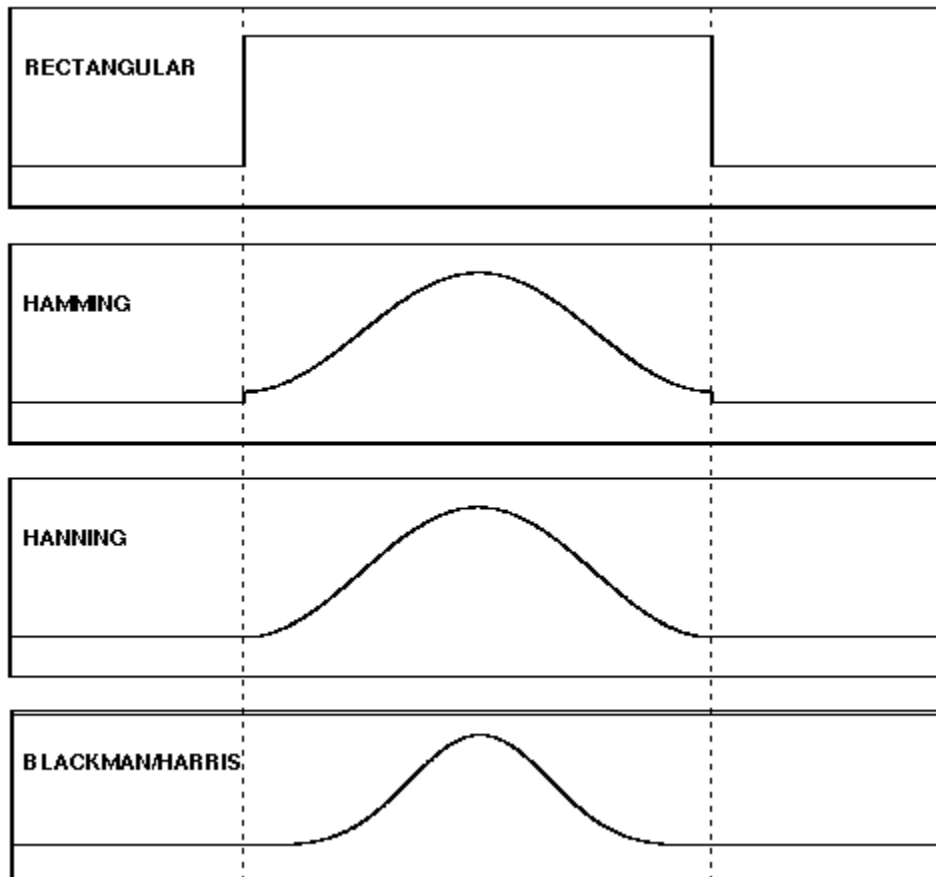


Figure 1.26 Time domain shapes of the various windows.

A rectangular window is typically used for impulse response testing since the beginning points are usually zero and the data tapers to zero at the end of the record. The beginning points are zero because the impulse is normally placed in the center of the time domain record at the zero phase reference point. Other windows can be used if desired. If phase is not important, the impulse can be placed at the beginning of the record. For this case, the window must be rectangular.

For harmonic analysis of continuous time signals, the best window choice depends on the signal characteristics and on the particular characteristics that are of most interest. The use of Blackman/Harris, Hanning or Hamming windows typically makes harmonics observation easier.

1.22 Scallop Loss of Windows

Consider the peak of the main lobe to be zero dB at frequency bin f_1 . As a constant amplitude input frequency is swept between the two frequency bins, f_1 and f_2 , the maximum magnitude deviation from zero dB is referred to as *scallop loss* or the picket fence effect. Figure 1.27 illustrates the concept. When the input frequency falls exactly halfway between two frequency bins, the scallop loss is at a maximum. Curves 1, 2, and 3 depict the main lobe at three different positions as the input frequency is swept. Figure 1.28 shows the maximum scallop loss for each window.




Figure 1.27 Scallop loss or picket fence effect.

	Rectangular	Hamming	Hanning	Blackman/Harris
Highest Side Lobe	-13	-43	-32	-94
3 dB BW (bins)	0.89	1.3	1.44	~1.92
6 dB BW (bins)	1.21	1.81	2.00	~2.73
Max Scallop Loss dB	3.92	1.78	1.42	~0.81

Figure 1.28 Filter characteristics of the FFT windows.

1.23 Coherent Gain Correction of Windows

Except for the rectangular window, each window introduces an energy loss due to the tapering of amplitude at the ends of the input time record. Because this would cause a magnitude error, the TDS oscilloscopes correct the FFT magnitude output data by a factor of $1/CG$, where CG is the coherent gain of the window. The net result of this correction is that main lobes *centered on a bin* will show no magnitude difference as the user changes from one window function to another. When the lobe is between two bins, magnitude differences between windows will be observable due to different scallop losses for

1.24 Equations for the FFT Windows

Except for the rectangular window, the windows used in the TDS oscilloscopes are derived from a cosine series. The window type is obtained by substituting the correct coefficients for the cosine terms into eq(1.3).

$$W_n = \sum_{i=0}^m \left(a_i \cdot \cos \left(\frac{2\pi i}{N} \cdot \left(n - \frac{N-1}{2} \right) \right) \right)$$

$$n = 0 \text{ to } N-1$$

where:

- w is the window data array
- a is the array of window coefficients
- N is the window length
- m is the window order
- i is the index to the window coefficients array
- n is the index to the window data array.

eq(1.3)

The following table (Figure 1.29) indicates the values to substitute into eq(1.3) to obtain the various windows.

	Hamming	Hanning	Blackman/Harris
m	1.0	1.0	3.0
a_0	0.543478	0.50000	0.35875
a_1	0.456522	0.50000	.48829
a_2	--	--	0.14128
a_3	--	--	0.01168
Highest side lobe	-43 dB	-32 dB	-94 dB

Figure 1.29 Window coefficients for eq(1.3)

For the rectangular window:

$$W_n = 1.0$$

1.25 Effects of Display Interpolation on the FFT

The user can horizontally zoom a waveform on the display with either Sin(x)/x or linear interpolation. (The choice is made from the **DISPLAY -> Filter** menu selection.) When the record length is set to 500 points, the FFT waveforms are always displayed zoomed by a minimum of 2X. Other length FFTs have a minimum zoom of 1X or less. The effects of Sin(x)/x interpolation are dependent on the window used. The user can check the effects of the interpolation by toggling between Sin(x)/x and linear interpolation, observing the results on the display.

Another way to view the effects of interpolation is to use *real point intensification*, which is turned on by pressing **DISPLAY -> Style -> Intensified Sample**, a setting in which FFT output points, before zooming, are displayed at brighter intensity than interpolated points.

Sin(x)/x interpolation can give good results on the FFT magnitude display. Linear interpolation should be used for phase displays because Sin(x)/x will cause an undesirable ringing in the results.

2.1 Mixed Radix FFT Derivation

The standard Discrete Fourier Transform (DFT) is the starting point for the mathematical derivation of the mixed radix FFT. The DFT is shown below as eq(2.10):

$$X(k) = \frac{1}{N} \sum_{n=0}^{N-1} x(n) e^{-j2\pi nk/N}$$

for: $k = 0$ to $N-1$

where:

$x(n)$ = a point in the input time domain data array
 $X(k)$ = a point in the output frequency domain data array
 N = Length of the FFT input data block
 k = index to the current output data point
 n = index to the input data array
 j = square root of -1

eq(2.1)

The value of **Wnk** is commonly referred to as the **twiddle** or phase factor. By using Euler's formula the expression for **Wnk** may be written as eq(2.3):

ALSO:

$$\text{Twiddle factor : } Wnk = e^{-j2\pi nk/N}$$

eq(2.2)

$$Wnk = \cos\left(\frac{2\pi nk}{N}\right) - j \sin\left(\frac{2\pi nk}{N}\right)$$

eq(2.3)

Thus the twiddle factors are obtained either by computing sine and cosine values or by looking them up in a table. For faster execution, table lookup is used in the TDS oscilloscopes.

2.2 A Look at the Mixed Radix FFT for 1000 Points

In the TDS oscilloscopes, FFT record lengths were made to be products of powers of 2 and 5 to give a more convenient human interface; therefore, the data blocks can be broken down by factors of 2 and 5. Since the sample rates in the oscilloscopes also are products of powers of 2 and 5, the delta frequency (Δf) per output point of the FFT results in nice round values. This concept is illustrated in Figure 2.1 using the 100 GS/s equivalent-time sample rate. The change in frequency between points is 100 MHz.




Figure 2.1 Illustration of delta frequency per FFT bin.

If computed as presented, the DFT, eq(2.1), contains many redundant calculations. This equation can be factored into a form called a *mixed radix FFT*, making the computation faster by removing much of the computational redundancy. As an example, a 1000 point FFT is shown below.

A 1000 point data block can be split into two 500 point data blocks. Each data block is obtained by taking every second point out of the original 1000 point sequence, as shown in eq(2.4). This equation is obtained by factoring eq(2.1).

$$X(k) = \frac{1}{N} \sum_{m=0}^1 \sum_{n=0}^{499} x(2n + m)e^{-i2\pi(2n+m)k/N}$$

for: $N = 1000$ and $k = 0$ to 999

eq(2.4)

From eq(2.4), each of the 500 point data blocks can be broken down into two blocks of 250 point length, as shown in eq(2.5).

$$X(k) = \frac{1}{N} \sum_{p=0}^1 \sum_{m=0}^1 \sum_{n=0}^{249} x(4n + 2m + p)e^{-i2\pi(4n+2m+p)k/N}$$

eq(2.5)

The 250 point data blocks are again divided by 2 to obtain:

$$X(k) = \frac{1}{N} \sum_{q=0}^1 \sum_{p=0}^1 \sum_{m=0}^1 \sum_{n=0}^{124} x(8n + 4m + 2p + q)e^{-i2\pi(8n+4m+2p+q)k/N}$$

eq(2.6)

The 125 point data blocks from eq(2.6) are not evenly divisible by 2; however, they may be broken down by factors of 5, in three stages, to obtain the following:

$$X(k) = \frac{1}{N} \sum_{s=0}^1 \sum_{r=0}^1 \sum_{q=0}^1 \sum_{p=0}^4 \sum_{m=0}^4 \sum_{n=0}^4 x(200n + 40m + 8p + 4q + 2r + s)e^{-i2\pi(200n+40m+8p+4q+2r+s)k/N}$$

eq(2.7)

At this point, eq(2.7) is factored to obtain a form that more closely represents how it will be computed in the final algorithm. For the final implementation, the $1/N$ factor is distributed throughout each stage to prevent accumulator overflow in the fixed point processor.

$$T = -j2\pi \frac{k}{N}$$

eq(2.8)

$$X(k) = \frac{1}{N} \sum_{s=0}^1 e^{Ts} \sum_{r=0}^1 e^{2Tr} \sum_{q=0}^1 e^{4Tq} \sum_{p=0}^4 e^{8Tp} \sum_{m=0}^4 e^{40Tm} \sum_{n=0}^4 x(200n + 40m + 8p + 4q + 2r + s)e^{-200Tn}$$

eq(2.9)

The 1000 point mixed radix FFT described by eq(2.9) has three stages containing 5 point DFTs, followed by three stages containing 2 point DFTs. The 5 point DFT is known as radix 5 butterfly, and the 2 point DFT is a radix 2 butterfly. Figure 2.2 shows a data flow diagram that illustrates how the computations take place for a radix 5 butterfly. Each triangle represents a multiply by a complex twiddle value.




Figure 2.2 Flow diagram for a radix 5 butterfly.

Figure 2.3 illustrates a complete data flow diagram for a 50 point mixed radix FFT. It helps visualize how the input data array is scrambled and how the radix 5 and radix 2 butterflies are applied to the data. Each line represents a multiply by a complex twiddle value.




Figure 2.3 Computational flow diagram for a 50 point mixed radix FFT

2.3 Understanding FFT Phase

This section presents mathematical detail explaining how to interpret the phase data from the FFT. The expression for the DFT is written again in eq (2.10) below. The complex magnitude for a given frequency bin indexed by k is simply the summation of all points in the data block, $x(n)$, multiplied by a complex sinusoid block of data of a frequency determined by k . The variable k is the index to the FFT output frequency bin and equals the number of cycles in the twiddle table.

$$X(k) = \frac{1}{N} \sum_{n=0}^{N-1} x(n)e^{-i2\pi nk/N}$$

$$\text{for: } k = 0 \text{ to } N - 1$$

eq(2.10)

$$X(k) = \frac{1}{N} \sum_{n=N/2}^{N/2-1} x(n)e^{-i2\pi nk/N}$$

$$\text{for: } k = 0 \text{ to } N-1$$

eq (2.11)

The *real* part of the complex sinusoid twiddle factor is show in Figure 2.4 as described by eq (2.10) for values of $k=1$, $k=2$, and $k=3$. Notice they all start at the beginning of a cycle peak at the start of the record. This is the zero phase reference point as implemented by eq(2.10).




Figure 2.4 The real part of the twiddle factors for $k=1$, $k=2$, and $k=3$. This is according to eq(2.10).

By adding to the phase angle of every odd FFT bin, the zero phase reference point is moved to the center of the FFT time record in the oscilloscope. This is described by eq (2.11) and is shown in Figure 2.5.




Figure 2.5 The real part of the twiddle factors according to eq(2.11).

Because the FFT windows in the TDS oscilloscopes are symmetrical about the centers of the records, it is desirable that the zero phase reference point be in the center. This makes the phase display much easier to read and interpret when a window other than the rectangular is in use.

The output data from the FFT computation is a set of complex numbers of the form:

$$F = a + jb$$

where:

F is a complex FFT output data point
 a is the real part and b is the imaginary part
 j is the square root of negative one.

eq(2.12)

The magnitude, **M**, of **F** is given by:

$$M = \sqrt{a^2 + b^2}$$

eq(2.13)

The phase, **T**, of **F** is given by:

eq(2.14)

The phase angle, obtained from eq(2.14), is with respect to the zero phase reference point in the FFT time record (as shown in Figure 2.5). In other words, **the phase of a frequency bin of interest is the phase with respect to the twiddle factor of that same frequency.**

2.4 FFT Phase and Magnitude Accuracy

The accuracy of the FFT phase and magnitude is difficult to generalize in that large variations are possible depending on factors such as: FFT window in use, sample rate, trigger position with respect to the sample clock from one acquisition to the next, and type of display interpolation.

The actual Digital Signal Processing implementation of the TDS 540 FFT uses 16-bit data storage for intermediate stages, and 16-bit twiddle storage. Theoretically, the best dynamic range is 96 dB; round-off and truncation noise, however, will limit accuracy at the levels around 70 dB below the maximum output values from the ADC. These errors are insignificant compared to errors possible in the generic use of the FFT.

Magnitude Error: The largest magnitude errors that occur are due to scallop loss when the signal frequency is halfway between two bins (see Section 1.22). A second source of magnitude error is due to such FFT window characteristics as main lobe width and side lobe attenuation (see Section 1.21). Another source of magnitude error could be the interpolation when using Zoom mode to expand the signal. There is a choice of linear or Sin(x)/x interpolation for horizontal expansion of the signal display.

Phase Error: Errors in the phase display can range up to 180 degrees. Because the sample clock is asynchronous with respect to the trigger, the phase can vary from 0 to -180 degrees from one acquisition to the next.

The FFT frequency range is from 0 Hz to 1/2 the sample rate. For input frequencies close to 1/2 the sample rate, there are only 2 or 3 samples per cycle. The least error in the phase display is obtained when there are many samples per cycle in the signal of interest. For example, the trigger jitter is 0 to +1

sample interval; if there are 100 samples per cycle, the phase error would be approximately 1/100 or 1%. Also, if the signal contains many harmonics, the lower frequency components will have better phase accuracy than the higher-frequency components.

Remember that the phase reference point is at 50% of the input time record to the FFT. Adjusting the pretrigger position or the trigger level can change this reference point's position with respect to the input waveform, changing the phase angle difference between adjacent bins. If the difference between FFT bin 1 and bin 2 is 45 degrees and then the pretrigger position is adjusted, the phase angle difference will change to another value.

Finally, note that the phase computation is not unwrapped. In other words, if the phase angle is really $3[\pi]$, it would be displayed as $[\pi]$. In this context, the phase error could be even larger than just described.

Summary

Key features of FFT implementation in TDS oscilloscopes

1. The FFT input length equals the acquired record length. Exceptions: for record lengths of 15k and 50k, the FFT is 10k long, starting at the beginning of the acquired record.
2. Rectangular, Hamming, Hanning, and Blackman/Harris windows are provided.
3. The zero phase reference point is in the center of the FFT time domain record.
4. A phase suppression control is provided to remove noise phase.
5. The vertical scale can be displayed as linear RMS, dBV_{rms} , phase degrees, or phase radians.

Primary uses for the FFT

1. Impulse response testing to obtain the phase and magnitude as a function of frequency.
2. Step response testing as an alternative to impulse testing.
3. Harmonic analysis and interpretation of the phase display.
4. Analyzing line current harmonics of electrical equipment.
5. Observing distortion.
6. Characterizing noise sources in a system.

Characteristics and cautions for FFT use

1. Aliased signals can be wrapped into the display. They can be identified by recognizing the characteristics of the signal under observation, or by changing the sample rate, or by changing the frequency of the input signal when possible. Bandwidth limit filters or Hi Res acquisition mode can be used to attenuate alias signals in some situations. Average or Hi Res mode also can attenuate aliased random noise.
2. Longer record lengths provide better frequency resolution and a lower noise floor.

3. Performing an FFT on a waveform stored in a ref memory required less computing power from the oscilloscope's signal processor since it computes only when a new waveform is saved into the ref.
 4. Hi Res mode can be used to reduce noise that is not synchronized with the trigger, but the frequency response characteristic of Hi Res mode must be taken into account when interpreting the FFT results. Hi Res mode is good for single-shot or repetitive acquisitions.
 5. For repetitive triggered acquisitions, Average mode can be used to reduce noise not synchronized with the trigger.
 6. If the signal contains components of interest that are not synchronized to the trigger, Hi Res and Average acquisition modes are not appropriate.
 7. Be aware of the effect that $\text{Sin}(x)/x$ or linear interpolation may have on the phase or magnitude display of the FFT output. Use linear interpolation for phase displays.
 8. For impulse response testing, position the impulse at the zero phase position in the center of the FFT time record.
-

FFT Glossary

Alias

Nyquist theorem says that a signal must be sampled at a minimum rate equal to twice the input signal frequency for accurate representation. Signals sampled at a rate lower than this minimum (the Nyquist Rate) will show up as low-frequency signals due to insufficient sample points. Such a signal is called an alias. The phenomenon is called aliasing.

Bin

Each frequency point represented in the frequency domain display of an FFT is called a bin.

Bode plots

The amplitude and phase plots as a function of frequency are known as Bode plots. The logarithmic amplitude thus defined is measured in terms of dB.

Decimation interval

The time interval between adjacent sample points.

FFT

Fast Fourier Transform -- a procedure for calculating discrete frequency components from sampled time data. Special digital signal processing chips perform FFTs in today's oscilloscopes.

FFT Phase

The time comparison in degrees or radians of a signal harmonic and the twiddle factor of the same frequency.

FIR filter

Finite Impulse Response filter -- the impulse response for this system has a finite duration; it is non-zero only over a finite time interval. After a finite time (based on filter design), the system response becomes zero. An FIR filter is sometimes known as a *convolver*. An FIR filter has all zeros in its transfer function, cannot oscillate, and is always stable. It is typically designed to have a linear phase characteristic.

Frequency domain

Signals displayed with frequency along the X-axis and magnitude along the Y-axis.

Linear time invariant

Systems whose parameters do not change with time invariant (also constant parameter, fixed, or stationary). If a system itself does not change with time, then for any given initial state, the shape of the output will depend only on the shape of the input, and not on the instant at which the input is applied.

Noise floor

That amplitude threshold in a frequency domain instrument below which the presence of a signal will not be measurable (discernible). For example, if the noise floor of an instrument is -75 dB, a signal at -78 dB will be buried in the noise, and not be measurable.

Nyquist point

The minimum sampling frequency as defined by the Nyquist Theorem which requires that a signal be sampled at a minimum rate equal to twice the input signal frequency for faithful reproduction. Sampling below this frequency will lead to aliasing.

Record length

The number of data points in a time record. For every point in a sampled time record, there is a corresponding point in the frequency record. Half of the points in the frequency record get discarded as these are image frequencies (reflections).

Time domain

Signals displayed with time on the X-axis and amplitude on the Y-axis. Instruments which display signals in this format are called time domain instruments.

Twiddle factor

The complex sinusoid coefficients which are multiplied by the input data when computing an FFT.

Zero phase reference point

The point in the time domain record to which the phase values of the FFT are referenced.

Contact Tektronix:

ASEAN / Australasia / Pakistan (65) 6356 3900
Austria +41 52 675 3777
Balkan, Israel, South Africa and other ISE Countries +41 52 675 3777
Belgium 07 81 60166
Brazil & South America 55 (11) 3741-8360
Canada 1 (800) 661-5625
Central Europe & Greece +41 52 675 3777
Central East Europe, Ukraine and Baltics +41 52 675 3777
Denmark +45 80 88 1401
Finland +41 52 675 3777
France & North Africa +33 (0) 1 69 81 81
Germany +49 (221) 94 77 400
Hong Kong (852) 2585-6688
India (91) 80-22275577
Italy +39 (02) 25086 1
Japan 81 (3) 6714-3010
Luxembourg +44 (0) 1344 392400
Mexico, Central America & Caribbean 52 (55) 56666-333
Middle East, Asia and North Africa +41 52 675 3777
The Netherlands 090 02 021797
Norway 800 16098
People's Republic of China 86 (10) 6235 1230
Poland +41 52 675 3777
Portugal 80 08 12370
Republic of Korea 82 (2) 528-5299
Russia, & CIS 7 095 775 1064
South Africa +27 11 254 8360
Spain (+34) 901 988 054
Sweden 020 08 80371
Switzerland +41 52 675 3777
Taiwan 886 (2) 2722-9622
United Kingdom & Eire +44 (0) 1344 392400
USA 1 (800) 426-2200
USA (Export Sales) 1 (503) 627-1916
For other areas contact Tektronix, Inc. at: 1 (503) 627-7111
Updated April 6, 2005

Copyright 2005, Tektronix, Inc. All rights reserved. Tektronix products are covered by U.S. and foreign patents, issued and pending. Information in this publication supersedes that in all previously published material. Specification and price change privileges reserved. TEKTRONIX and TEK are registered trademarks of Tektronix, Inc. All other trade names referenced are the service marks, trademarks or registered trademarks of their respective companies.