# Fast Fourier Transforms explained

In this white paper Pico Technology discusses how Fast Fourier Transforms (FFTs) can be used to analyze signals in the frequency domain, as well as which window to use improve your understanding of specific signals.

Frequency-domain fundamentals

When using an oscilloscope to analyze a signal, the instrument displays how the signal amplitude varies versus time. This is known as time-domain analysis. However, any continuous signal can be represented by multiple harmonically related sinusoidal waves of varying amplitude and frequency known as a Fourier series.

In the diagram below, we can see how a square wave can be approximately represented by summing the odd harmonic weighted values to the fundamental frequency. The amplitude of the 3[rd] harmonic is one third of the amplitude of the fundamental frequency, the 5[th] harmonic is one fifth of the amplitude of the fundamental and so on. In the approximation shown below, the Fourier series constructed from the fundamental frequency up to the 11[th] harmonic is:

$$y = \sin x + \frac{1}{3}\sin 3x + \frac{1}{5}\sin 5x + \frac{1}{7}\sin 7x + \frac{1}{9}\sin 9x + \frac{1}{11}\sin 11x$$
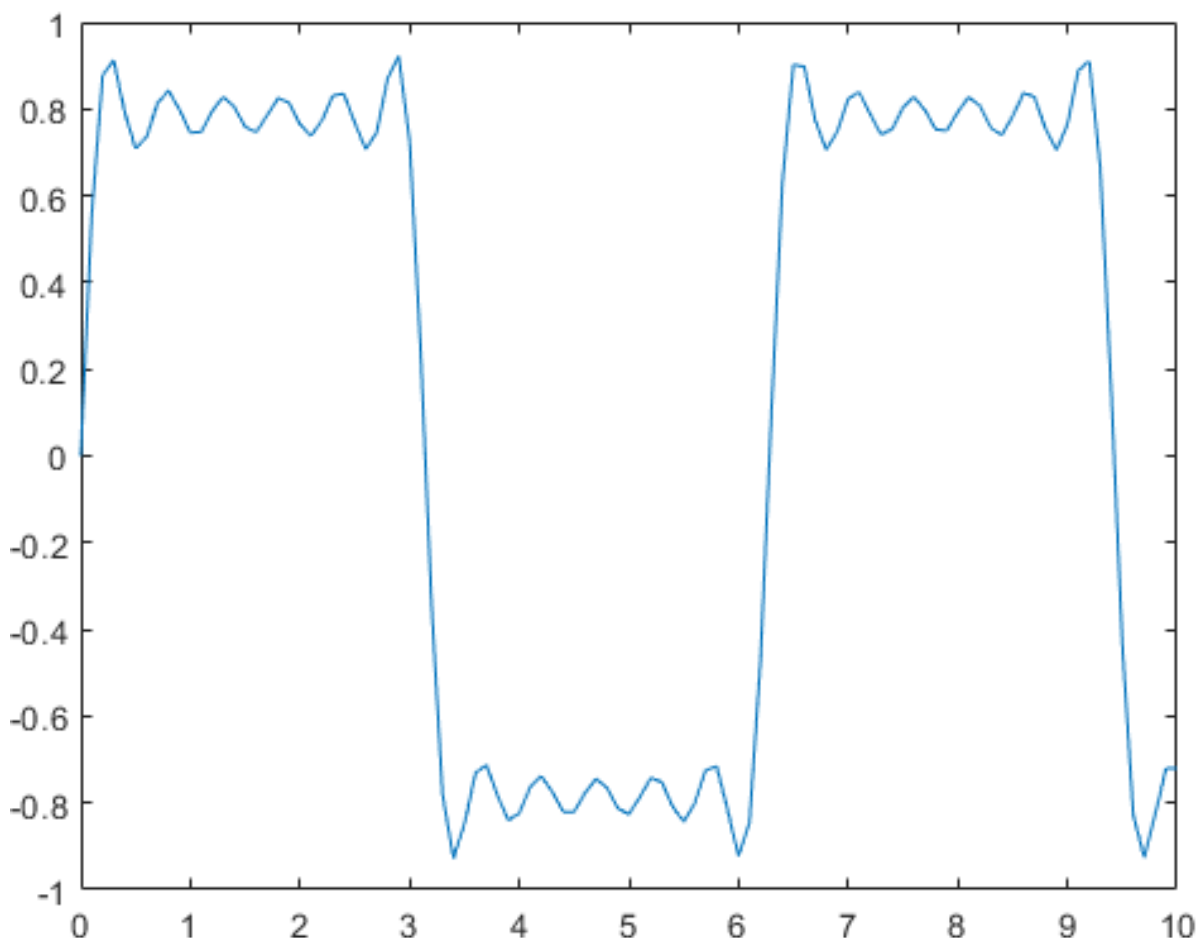
Approximation of a square wave (fundamental to the 11[th] harmonic)

If we then use the same function up to the 21st harmonic, we get a much more accurate approximation of the square wave:
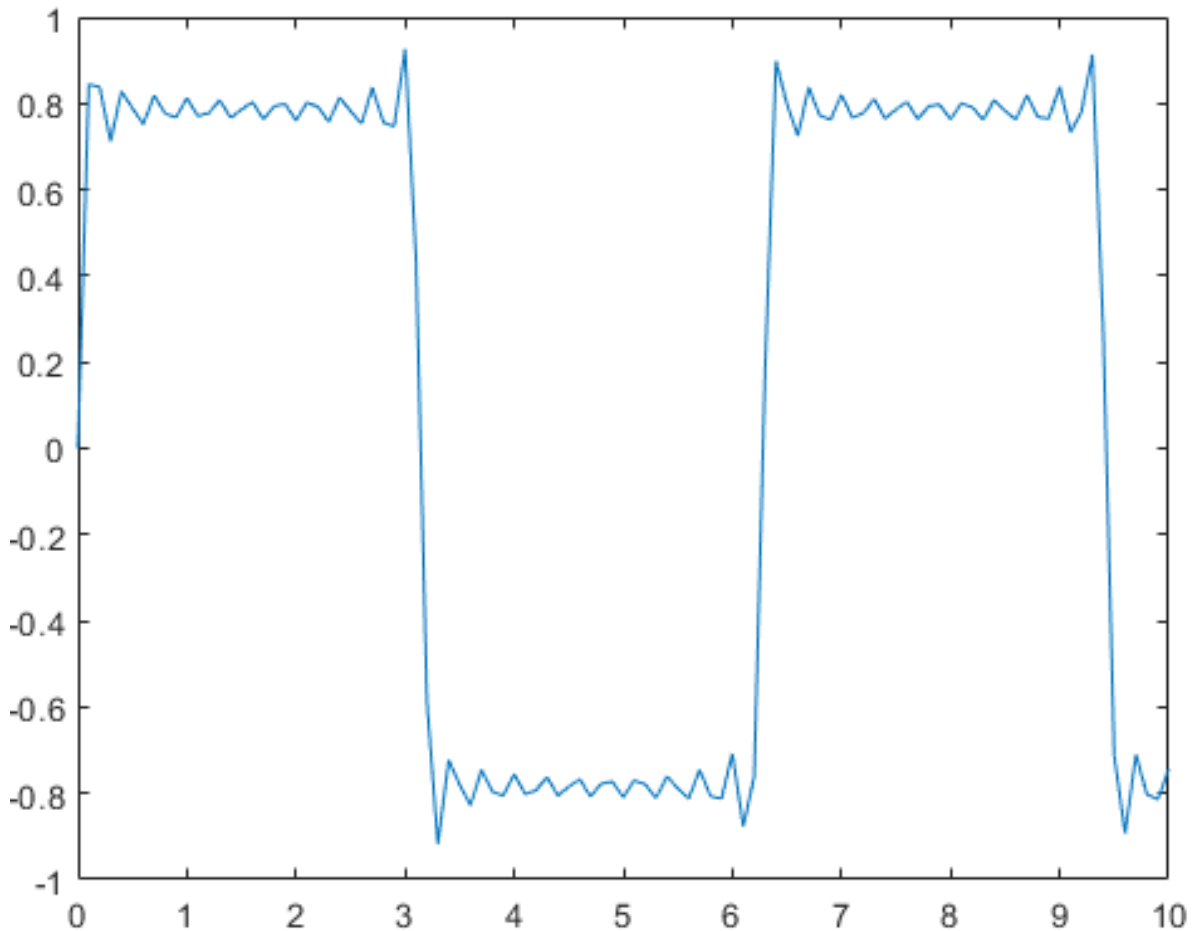
Approximation of a square wave (fundamental to the 21st harmonic)

If we use the reverse process to deconstruct a waveform into a Fourier series of sinusoids, the instrument can display how the signal amplitude varies versus frequency (spectrum). This is known as frequency-domain analysis.

Gibbs phenomenon

When introducing the Fourier series above, we described the input as "any continuous signal". Discontinuous signals – those with infinitely sharp edges, such as ideal square waves – contain infinitely high frequencies and cannot be exactly represented by a finite Fourier series. This limitation is known as the Gibbs phenomenon.

The Gibbs phenomenon can be seen as the ringing effects in the above waveform. Adding more terms to the series, representing more harmonics, increases the frequency of the ringing while the overshoot converges to a predictable value of just under 9% of the step size.

Why is frequency-domain analysis useful?

To use an everyday example, if the output signal of the music player in a mobile phone is deconstructed using digital signal processing (DSP) into its component frequencies, the amplitude of each band of frequencies can be boosted or attenuated to modify the equalization of the original music.

The Fast Fourier Transform

The examples shown above demonstrate how a signal can be constructed from a Fourier series of multiple sinusoidal waves. In order to analyze the signal in the frequency domain we need a method to deconstruct the original time-domain signal into a Fourier series of sinusoids of varying amplitudes. To implement this, we need to use a Discrete Fourier Transform (DFT), which deconstructs samples of a time-domain signal into its frequency components as discrete values also known as frequency or spectrum bins. An optimized and computationally more efficient version of the DFT is called the Fast Fourier Transform (FFT).

The oscilloscope spectrum display below shows a 10 kHz square wave deconstructed using an FFT, displaying up to the 9th harmonic at 90 kHz.
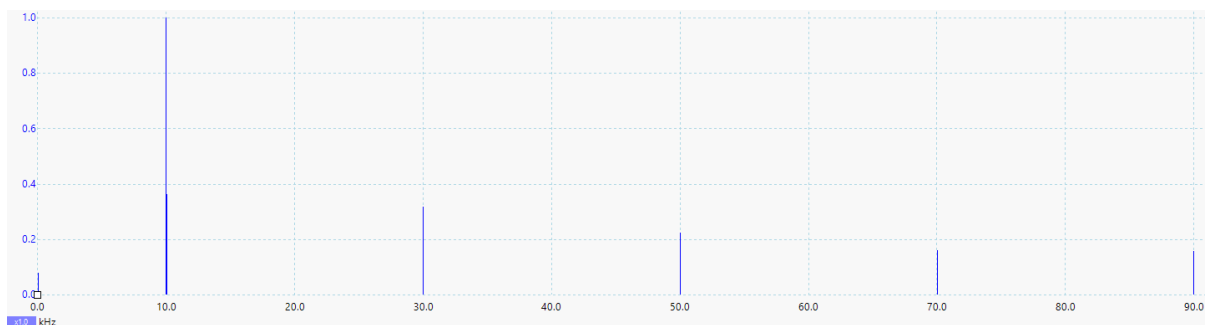
Spectral display of a 10 kHz square wave

FFT spectrum bins

A DFT or FFT can be expressed as $X(k) = \sum_{n=0}^{N-1} x(n * T_s) e^{-\frac{i 2 \pi k n}{N}}$

Where

$X(k)$	Complex discrete frequency spectrum
$x(n * T_s)$	Sample at the time $n * T_s$
k	Index of each discrete spectrum bin k = 0,1,2,3 etc.
n	Index of each time-domain sample n = 0,1,2,3 etc.
T_s	Sample period or interval
N	Number of samples (Window Size)

Each spectrum bin can be represented as $f(k) = k * 1/(N * T_s)$

This means that the measurement period, also known as the time gate ($N * T_s$), determines the frequency resolution bandwidth or bin width of the transformed signal. Finer frequency resolution is achieved by selecting a longer measurement period, either by increasing the number of samples or the sample period. However, as we shall see shortly, the sample period is derived

from the sample rate, which is indirectly derived from the frequency span. It may seem counterintuitive that in the frequency domain *lowering* the sample rate to increase the sample period improves resolution, because of course, in the time domain *increasing* the sample rate to reduce the sample period improves resolution.

The number of spectrum bins (FFT Size) should not be more than half of the number of samples of the original time-domain signal ($N/2$) and in PicoScope this is automated. Setting the number of spectrum bins to 131,071 will automatically set the number of acquired samples to 262,140.
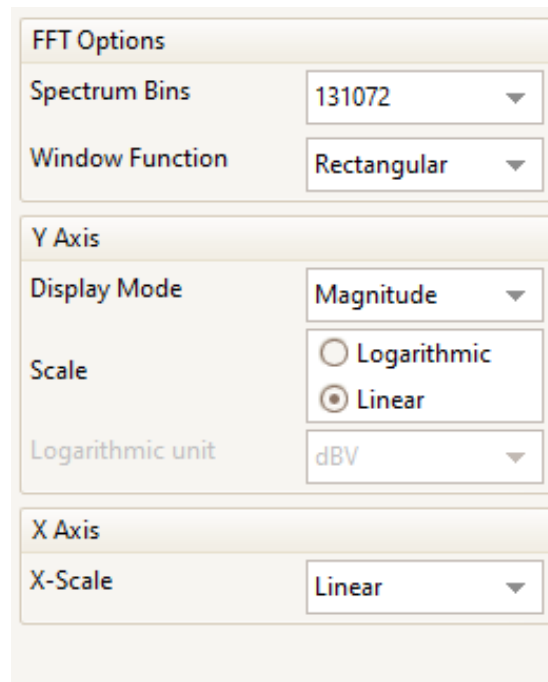
PicoScope FFT options dialog

The frequency span (or maximum observable frequency) is the Nyquist frequency, which is half the sample rate. Therefore, to observe the 9[th] harmonic of a 10 kHz square would require a frequency span of 90 kHz with a minimum sample rate of 180 kS/second. Setting the frequency span to 100 kHz will in PicoScope also set the sample rate to 200 kS/second, giving a sample period of 5 μs. The best frequency resolution is achieved by setting the frequency span to the minimum necessary, thus increasing the sample period, and the number of bins as large as possible whilst giving acceptable acquisition performance.

Properties ✕

|	
Sample interval	5 µs
Sample rate	200 kS/s
No. samples	262,140
Window	Rectangular
No. bins	131072
Bin width	762.9 mHz
Time gate	1.311 s

PicoScope FFT measurement properties

In this case we have a frequency span of 100 kHz spread over 131072 spectrum bins, so the minimum frequency resolution (bin width) is 100000/131072 = 0.7629 Hz and the measurement period or time gate ($N*T_s$) is 262140 * 5 µs = 1.311 seconds.

Reducing the resolution bandwidth of a swept spectrum analyzer improves noise performance. In the same way, increasing the FFT measurement period (time gate) and hence frequency resolution also reduces the FFT noise floor, which in turn improves the dynamic range of the measurement. The quantization noise of the analog to digital converter determines the ultimate measurement noise floor.

For an accurate representation of the frequency components to be calculated, the signal should ideally be periodic, and when using a rectangular measurement window, the measurement period ($N*T_s$) should be an integer multiple of the signal period. This is because if ($N*T_s$) is not an integer multiple of the signal period, the start and end of the measurement period will not be at zero-crossing points of the time domain signal, leading to discontinuities in the FFT. This results in both amplitude errors and frequency components that are not present in the original signal known. This "smearing" of the frequency spectrum is known as spectral leakage. Using a longer measurement period reduces the effect of the discontinuities and hence artificial frequency components but does not resolve the amplitude errors. To reduce the discontinuities themselves, it is necessary to gradually fade in and fade out the measurement window using a non-rectangular window function.

Window functions

The window function is multiplied by the sampled time-domain data to give a windowed waveform. Note that using a rectangular window does not modify the original data and has the same effect as not using any window function.
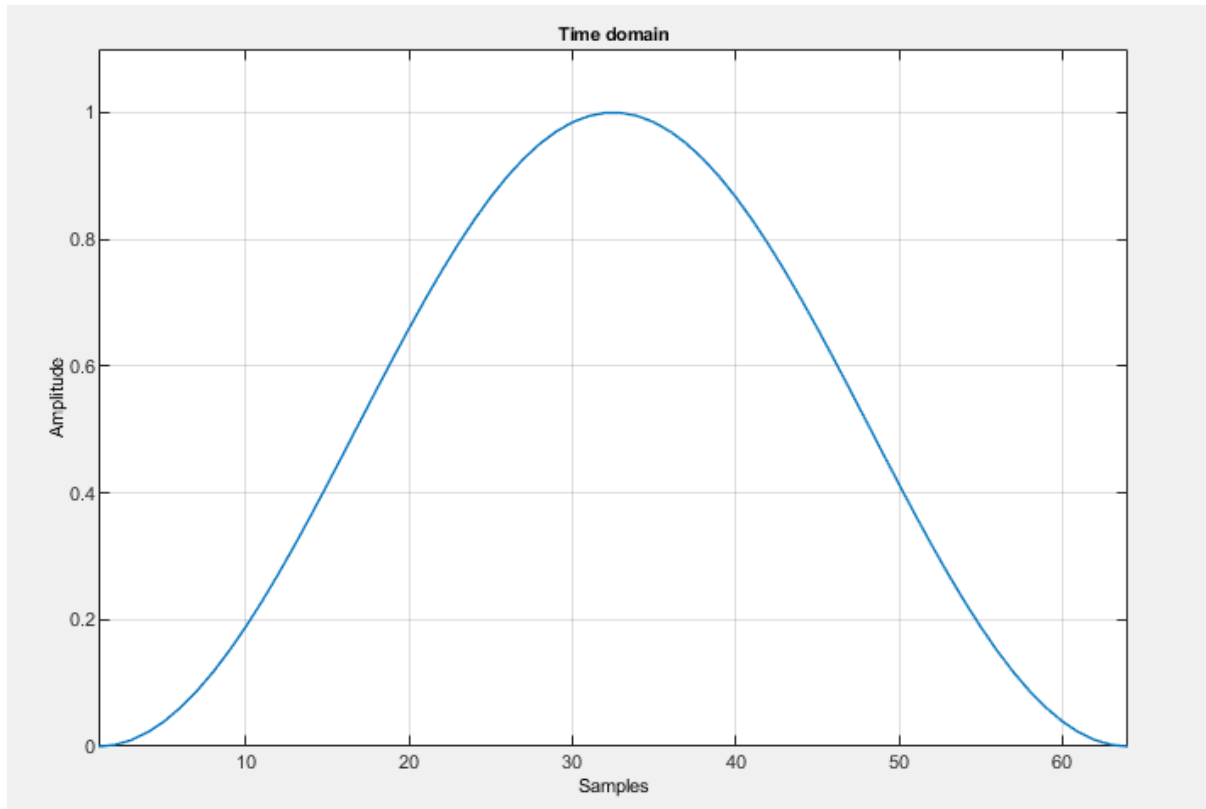

Window function amplitude versus sample count

Many window functions use a characteristic bell-shaped curve that smoothly fades the signal amplitude in and out from the start of the window to the end. This reduces the discontinuities that cause leakage and amplitude errors.

There are trade-offs when choosing an appropriate FFT window function. The FFT consists of multiple sinusoidal waveforms at different frequencies consisting of a main lobe (the peak centred on DC) and several side lobes (all the other peaks).

Spectral leakage is caused by sinusoidal side lobes overwhelming nearby sinusoidal main lobes. To minimize leakage the side lobe amplitude (dB) should be minimized, but the trade-off is that the main lobe bandwidth (main peak width in bins at −3 dB roll-off) is increased, thus reducing frequency resolution. It is also possible to improve reduce leakage without having a severe impact on the main lobe bandwidth, by selecting a window function that increases the side lobe roll-off rate (dB/octave).
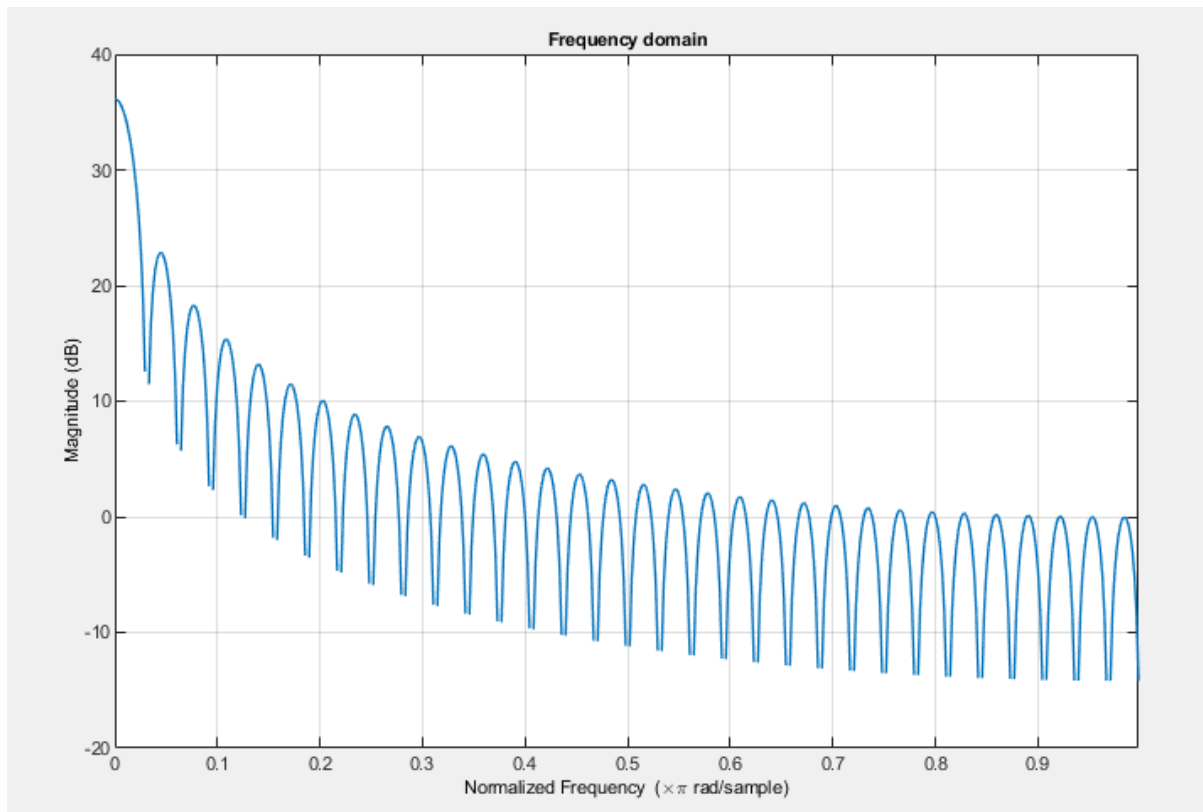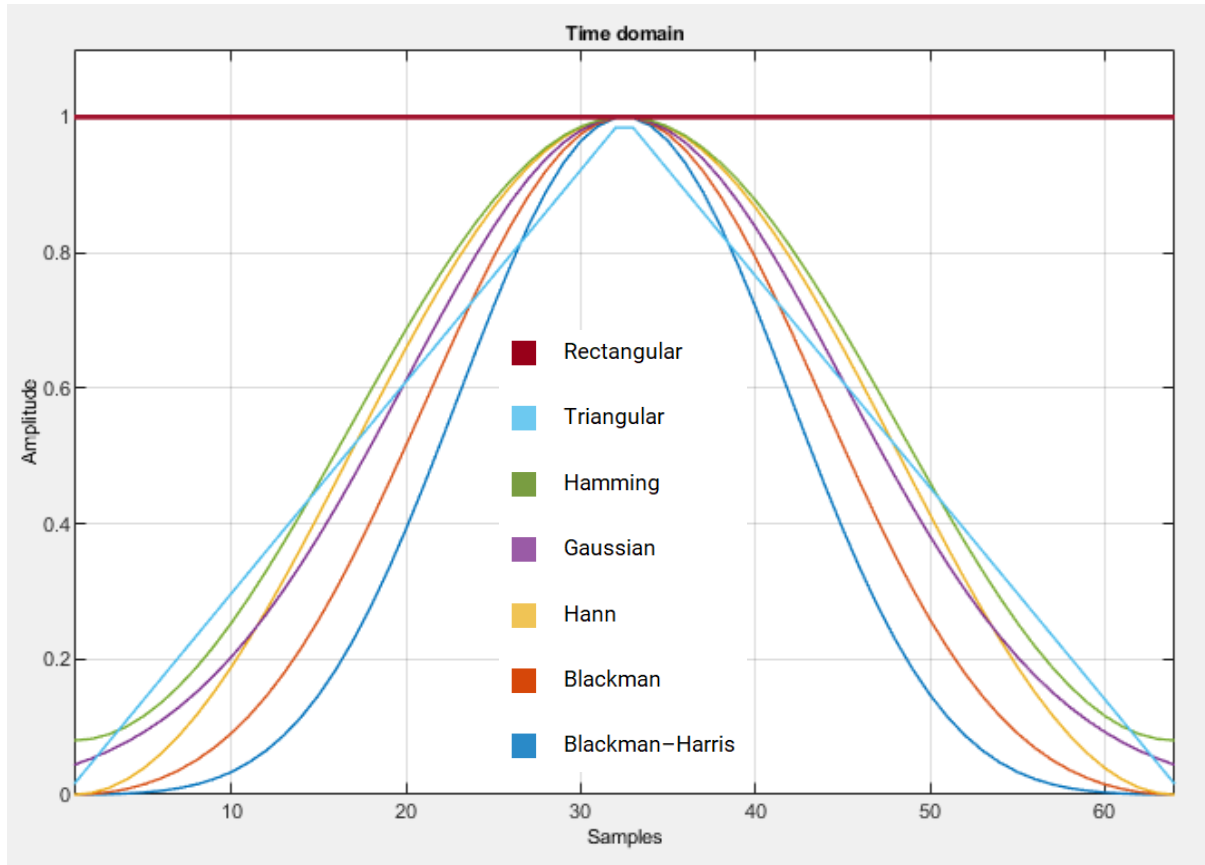
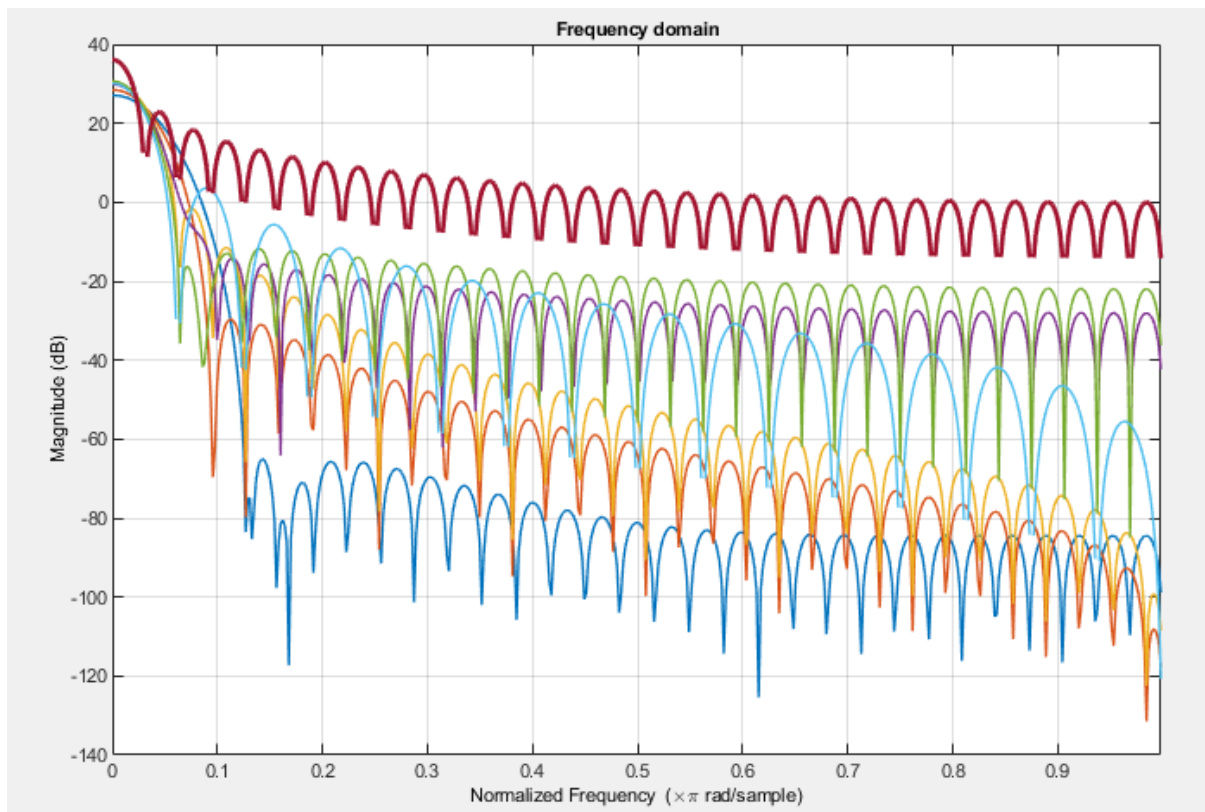Rectangular window function main lobe and side lobes

The table below shows the available window functions in PicosScope sorted in order of main lobe 3 dB bandwidth (peak width). As can be seen, the rectangular window gives the best frequency resolution at the expense of very poor side lobe performance. The Hann window offers better side lobe performance along with good side lobe roll-off at the expense of a ~60% increase in main lobe bandwidth. Overall it is a good compromise, which is why it is suitable in most applications and makes an excellent starting point when undertaking FFT analysis. Window functions that have the best (lowest) side lobe performance and highest side lobe roll-off will have the lowest spectral leakage, albeit with poorer frequency resolution.

| Window	Window Main peak width (bins at 3dB roll-off)	Highest side lobe (dB)	Side lobe roll-off (dB/octave)
Rectangular	0.89	-13	6
Triangular	1.28	-27	12
Hamming	1.30	-42	6
Gaussian	1.33	-42	6
Hann	1.44	-32	18
Blackman	1.68	-58	18
Blackman-Harris	1.90	-92	6
Flat-Top	2.94	-44	6

PicoScope window functions

PicoScope window functions displayed in the time domain

PicoScope window functions displayed in the frequency domain

Rectangular window
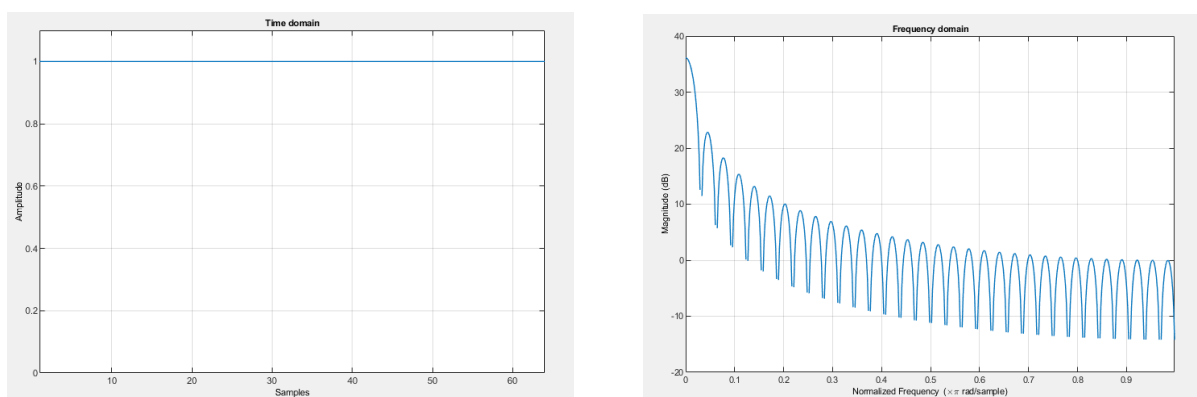
Rectangular window

The rectangular window (sometimes referred to as a *uniform window*) offers the best frequency resolution at the expense of poor side lobe performance. This leads to high levels of spectral leakage and poor amplitude accuracy on repetitive waveforms, where the start and end of the window do not correspond to zero crossing points. This window is best suited to making impulse response or transient measurements where the signal starts at zero amplitude and decays back to zero amplitude, thus avoiding any discontinuities at the start and end of the window. Another application well suited to the rectangular window is where the signal contains many closely spaced carriers of a similar power, such as broad-bandwidth constant-amplitude signals. With the best frequency resolution of any window function, the rectangular window enables the closely spaced carriers to be resolved.
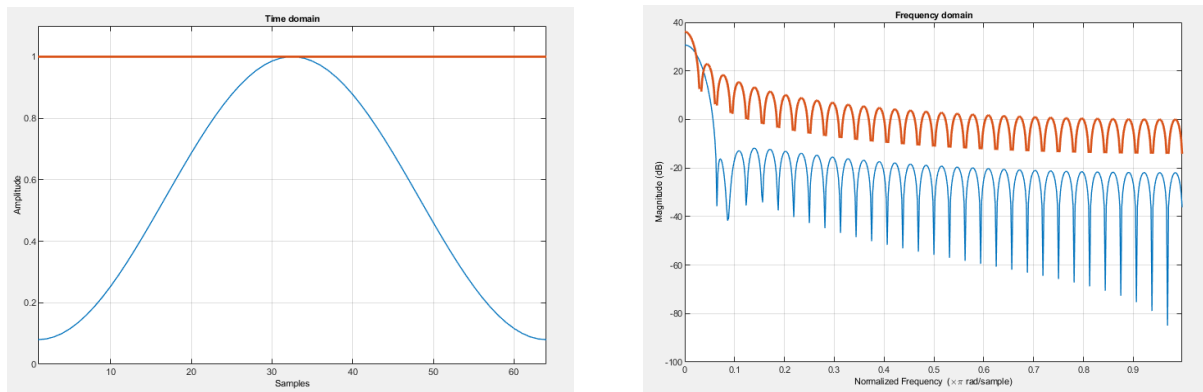
Triangular window
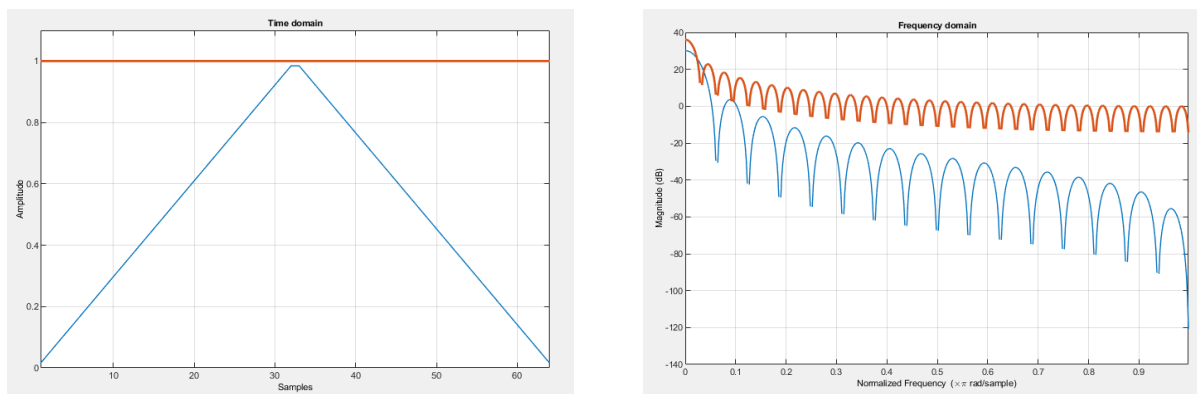
Triangular window compared with rectangular window

The triangular window (sometimes referred to as a Bartlett window) is a simple window function that offers the narrowest main lobe bandwidth performance after the rectangular function. It has poor side lobe performance and as such the worst spectral leakage after the rectangular function. In most applications this makes the triangular window less useful than the general-purpose Hann window and other more specialized window functions. Its biggest advantage over other window functions is that it has low *processing loss*, which is the degradation in signal-to-noise ratio caused by the window function. This does make the triangular window useful in applications requiring resolution of signals close to the noise floor as well. The triangular window is also good at resolving closely spaced carriers.

Hamming window

Hamming window compared with rectangular window

The Hamming window function does not reach zero amplitude at the start and end of the function. As a result, there are residual discontinuities which result in good first side lobe performance but poorer performance for the outer side lobes than some other windows. The Hamming function maintains good main lobe bandwidth and hence frequency resolution. Amplitude accuracy is adequate. Because of the good near side lobe performance and good frequency resolution, the Hamming window is useful in applications with narrowly spaced carriers, whilst being less susceptible to amplitude inaccuracy than the rectangular or triangular windows.

Gaussian window
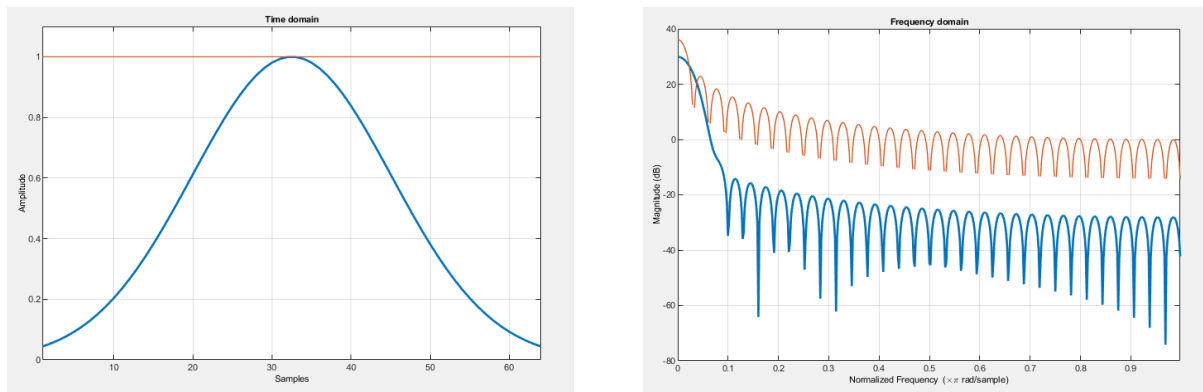
Gaussian window compared with rectangular window

The Gaussian window has adequate main lobe performance. Although the 3 dB bandwidth is good, the lower skirt is wider than for the previous windows discussed. It has good first side lobe performance and the outer side lobes roll off at the same rate as the Hamming window, giving very low spectral leakage. Its outstanding feature is that it offers the best time-bandwidth product of any of the window functions discussed in this white paper. This leads to very good frequency and time accuracy performance. The Gaussian window also offers good amplitude accuracy, making it well suited to calibration applications and other applications where measurement accuracy is of primary importance.

Hann window
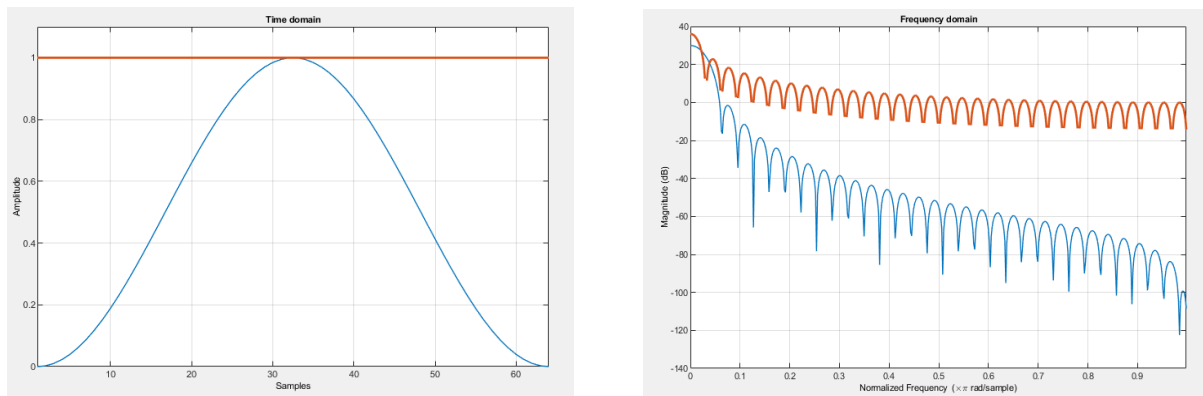
Hann window compared with rectangular window

The Hann window (sometimes referred to incorrectly as a *Hanning window*) is a good general-purpose window function for many applications. It has good side lobe roll-off and hence much reduced spectral leakage. The Hann window also maintains good main lobe bandwidth and offers adequate amplitude accuracy. It is the recommended FFT window function to use when the spectral composition of the signal is not well understood or does not require the specific characteristics of other functions. As a result, the Hann window can be used in most applications with satisfactory results.
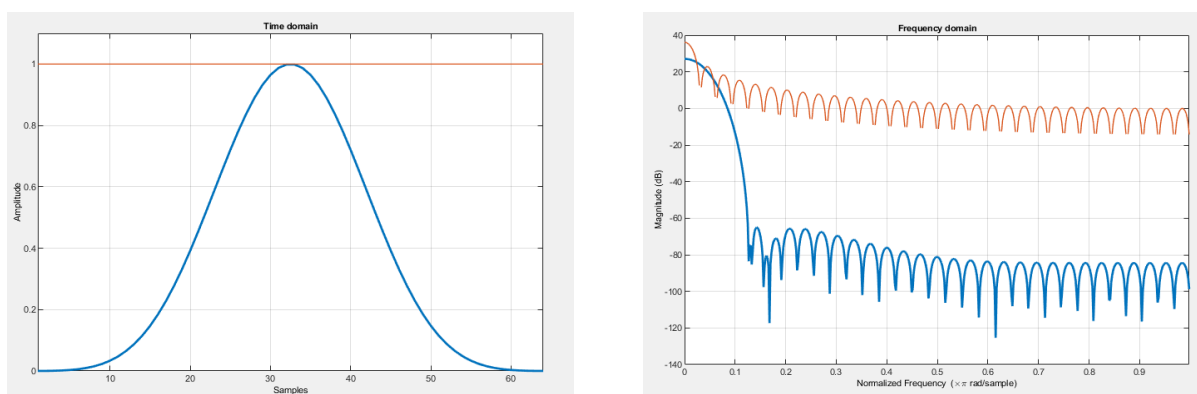
Blackman window
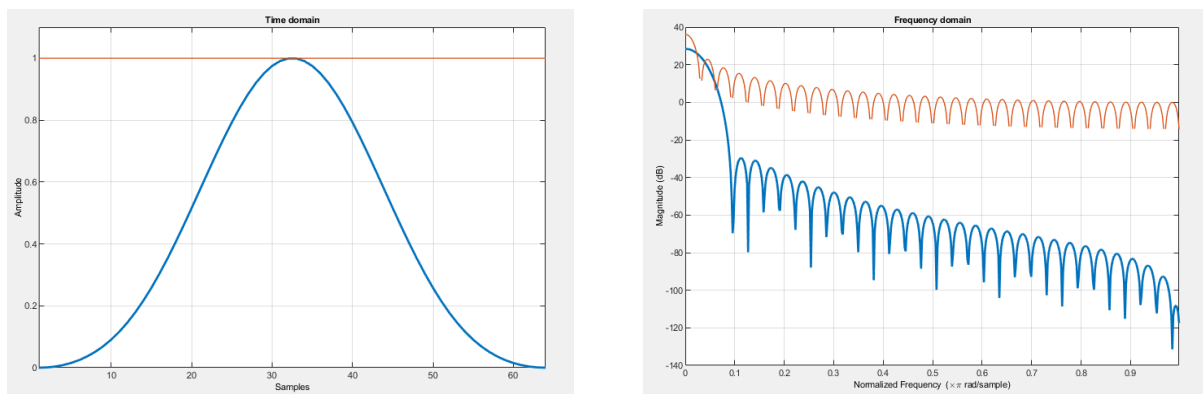
Blackman window compared with rectangular window

The Blackman window was designed to have a minimal spectral leakage, which it achieves at the expense of poorer main lobe bandwidth performance. The Blackman window offers good amplitude accuracy. The combination of extremely low spectral leakage, with a corresponding lack of artificial frequency artefacts as well as good amplitude accuracy, make the Blackman window well suited to audio analysis applications.

Blackman-Harris window

Blackman−Harris window compared with rectangular window

The Blackman-Harris window was again designed to offer minimal spectral leakage, which it achieves at the expense of a wider main lobe bandwidth than the Blackman window. First side lobe performance is outstanding at −92 dB with a gentle roll-off of the outer side lobes. The intention is to offer a general-purpose window function with a much better dynamic range than other functions such as the Hann window. The combination of high dynamic range, low levels of artificial frequency artefacts and good amplitude accuracy make the Blackman-Harris window very useful in many general-purpose applications. The limitation is the wide main lobe bandwidth, making it unsuitable for applications with narrowly spaced carriers.

Flat-top window
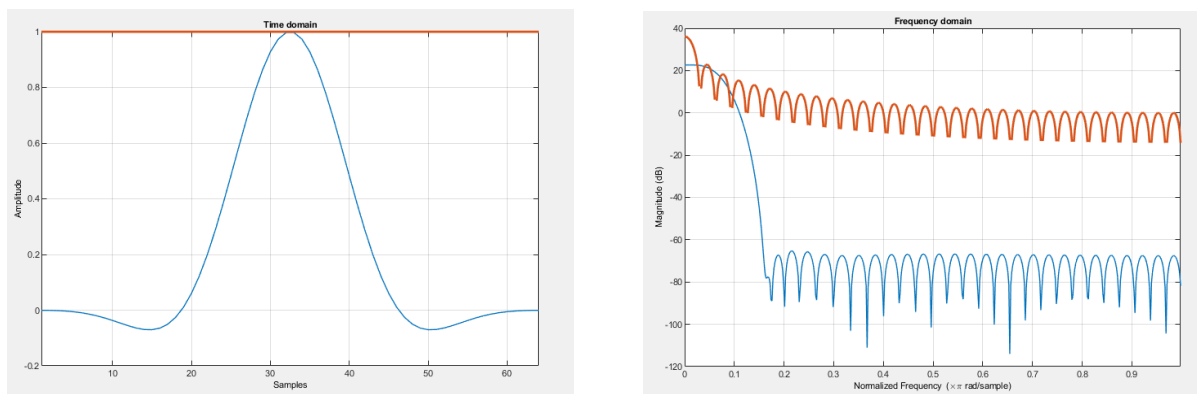
Flat-top window compared with rectangular window

In some ways the flat-top window has some undesirable characteristics. The main lobe bandwidth is very wide, and it has the highest processing loss of all the windows we have considered. This leads to a poor signal to noise ratio and makes the window not as useful for analyzing signals near the noise floor as other windows. As can be seen from the above diagrams, the flat-top window crosses below the zero amplitude line, which causes the very wide main lobe. However, the amplitude of the main lobe is the most accurate of any of the windows we have considered. The flat-top window has the best amplitude accuracy of any of the windows considered, also resulting in very low levels of pass band ripple. This amplitude accuracy makes the flat-top window ideally suited to calibration applications where amplitude accuracy is of primary importance, such as the calibration of transducers.

Summary

When analyzing a signal using an FFT, the best frequency resolution is achieved by setting the frequency span to the minimum necessary, thus increasing the sample period. The number of spectrum or frequency bins should be as large as possible, whilst maintaining acceptable acquisition performance.

When selecting which FFT window function to use, it is important to consider both the harmonic content of the time domain signal and which characteristics of the signal are of primary importance.

If the signal contains narrowly spaced carriers, choose a window with a narrow main lobe bandwidth such as rectangular, triangular or Hamming. However, if the signal of interest is in danger of being overwhelmed by nearby interferers, then choose a window with a low first side lobe level such as Blackman or Blackman–Harris. If overall measurement accuracy is most important, a Gaussian window is a good choice. If dynamic range is important then the Blackman-Harris window is the best choice. For the best amplitude accuracy, the flat-top window offers the best performance. If the signal content is unknown, then the Hann window's balance of characteristics makes it an excellent starting point.

All the spectrum windows discussed in this article are available in PicoScope 6, shown below.
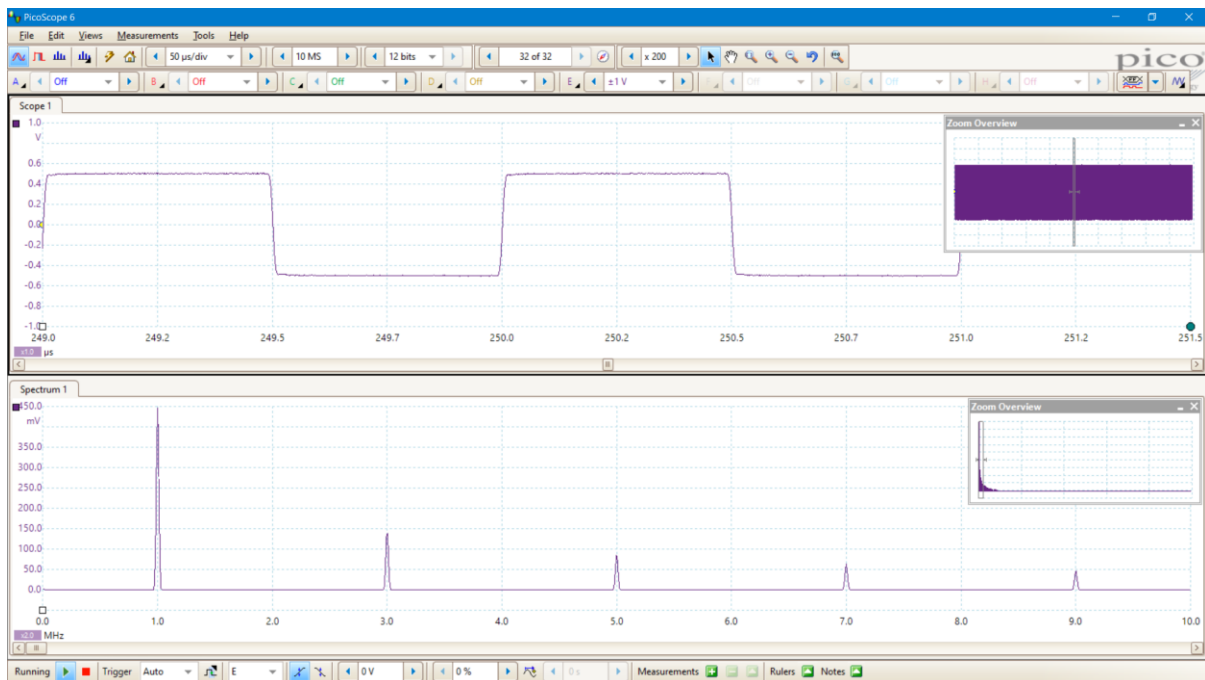
PicoScope FFT spectrum plot of a square wave, showing declining harmonics