SPECTRUM

▶ Application Note

Signal Processing for Digitizers

Modular digitizers allow accurate, high resolution data acquisition that can be quickly transferred to a host computer. Signal processing functions, applied in the digitizer or in the host computer, permit the enhancement of the acquired data or the extraction of extremely useful information from a simple measurement.

Modern digitizer support software, like Spectrum's SBench 6 and many third party programs incorporate many signal processing features. These include waveform arithmetic, ensemble and boxcar averaging, Fast Fourier Transform (FFT), advanced filtering functions and histograms. This application note will investigate all these functions and provide typical examples of common applications for these tools..

Analog Calculation (Waveform Arithmetic)

Analog calculation includes addition, subtraction, multiplication, and division of acquired waveforms. These functions are applied to data in order to improve signal quality or to derive alternative functions. One example is the use of subtraction to combine differential components into a differential waveform with reduced levels of common mode noise and pickup. Another is taking the product of current and voltage waveforms to compute instantaneous power.

Each of these arithmetic functions is applied to waveforms on a sample by sample basis. This assumes that the waveforms being combined have the same record length. Figure 1 shows the SBench 6 setup path for selecting analog calculation.

Right clicking on the desired source channel brings up a selection box. Selecting Calculation opens additional choices for measurements, signal calculation, signal conversion, and averaging. A choice of Signal Calculations provides access to Fast Fourier Transform (FFT), histograms,

Figure 1:Signal processing selections in SBench 6 are associated with each signal.Here A1-Ch0 is chosen and analog calculations selected.The Calculation dialog box is setup to add signals A1-Ch0 and A1Ch1.

filtering and several other functions. If Analog Calculation is chosen the Calculation dialog box pops up allowing the setup of the desired arithmetic operations. In this example the two source signals are to be added. Other choices are subtraction (SUB), multiplication (MULTI) and division (DIV). Similar selection paths can bring up all the other signal processing functions to be discussed.
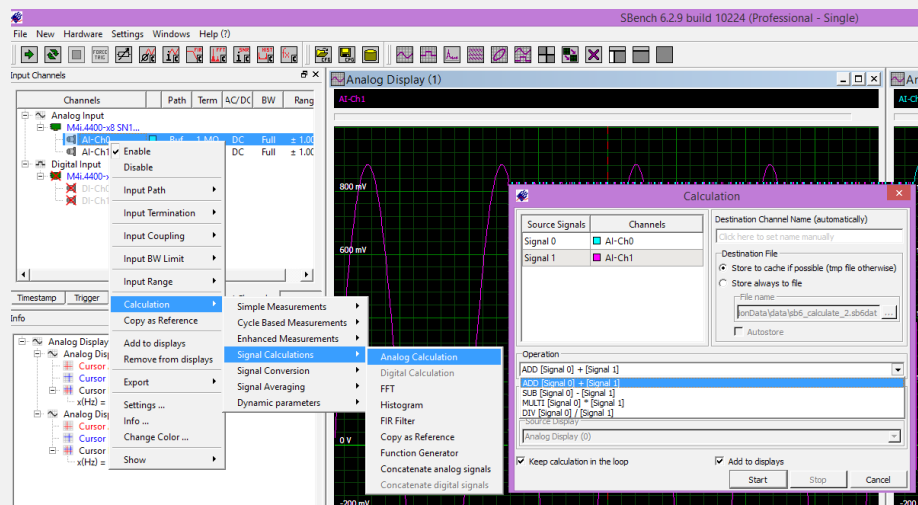
▶ Application Note

The first example of applying waveform arithmetic to practical problems is to subtract one signal component from another to compute the differential signal. This is shown in Figure 2.

Differential signals are commonly used to improve signal integrity. In the example in Figure 2 the 'P' and 'N' components of a 1 MHz clock (shown in the two right hand panels) are combined using the subtraction operation. The resultant differential signal is shown in the left hand grid. The Info pane at the left center uses parameters to measure the peak to peak and average value of each waveform. Note that the differential signal has twice the peak to peak amplitude and a near zero mean value. Note also that the common mode noise on the differential components has been eliminated.

Figure 2:Using the subtraction function to combine two differential components into a differential signal.

The second example multiplies a voltage waveform by a current waveform to obtain the instantaneous power as shown in Figure 3.

The source waveforms are the voltage across the power field effect transistor (FET) and the FET channel current in a flyback mode switching power supply. The product of those waveforms represents the instantaneous power dissipated by the FET. The current waveform (upper right grid) shows a linearly increasing ramp during the FET conduction peaking at 600 mA. The voltage across the FET is at a minimum during conduction but rises to a peak value of 260 V when the device is off. The product of those two waveforms is shown in the left grid. This is the instantaneous power waveform which shows significant peaks occur during the transitions between the on (conduction) and off states. The average (5.111W) and peak power (44.25W) are determined using parameters and appear in the info pane at the left center.
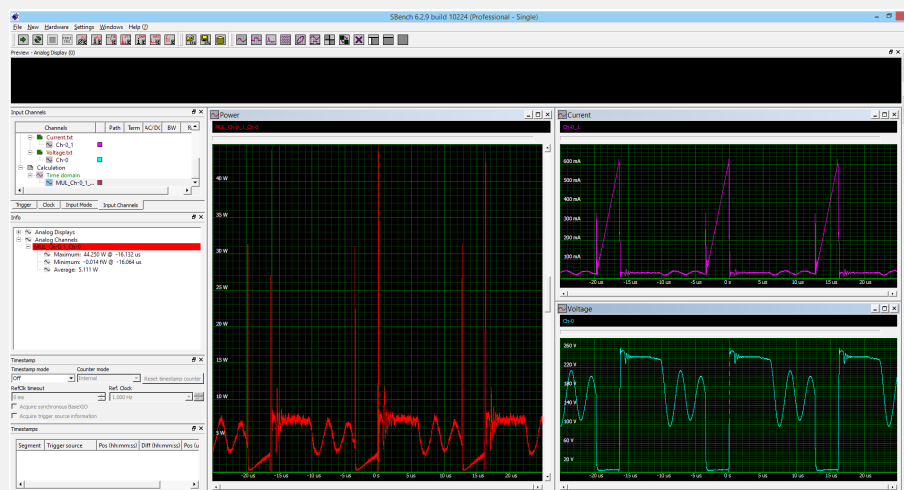

Figure 3: Using the multiplication function to compute instantaneous power from the current and voltage waveforms of a switch mode power supply.

These examples show how the analog calculations can be used to derive other important waveforms from those that are initially acquired.

Averaging

Averaging is a signal processing tool used to reduce the effects of noise and non-synchronous periodic waveforms on acquired signals. It requires multiple acquisitions and a stable trigger. Signal components that are not synchronous with the trigger timing, including random noise, are reduced in amplitude. The degree of reduction is dependent on the waveform characteristics and the number of acquisitions added to the average.

Spectrum's SBench 6 software, used in this application note, and most oscilloscopes perform ensemble averaging, meaning that the same sample location in multiple acquisitions are averaged together. If a stable trigger is available, the resulting average has a random noise component lower than that of a single-shot record.

Summed Averaging

Summed Averaging uses a fixed number of acquisitions and is the repeated addition, with equal weight, of the same sample locations from successive waveform acquisitions. Whenever the maximum number of sweeps is reached, the averaging process either stops or is reset to begin again.

Figure 4 shows the concept of a summed ensemble average:

In Figure 4 the arrows indicate the nth point. The amplitude value of the nth point of each acquisition is summed with those of the other acquisitions. The sum is then divided by the number of acquisitions to determine the nth value of the average. This takes place for all sample points in the acquisition group. The resultant averaged waveform has the same number of points as each acquired waveform.

Averaging is supported for both normal acquisition and for multiple (segmented) acquisitions. Multi averaging calculations permit the average of consecutive segments of the multiple recording acquisition.
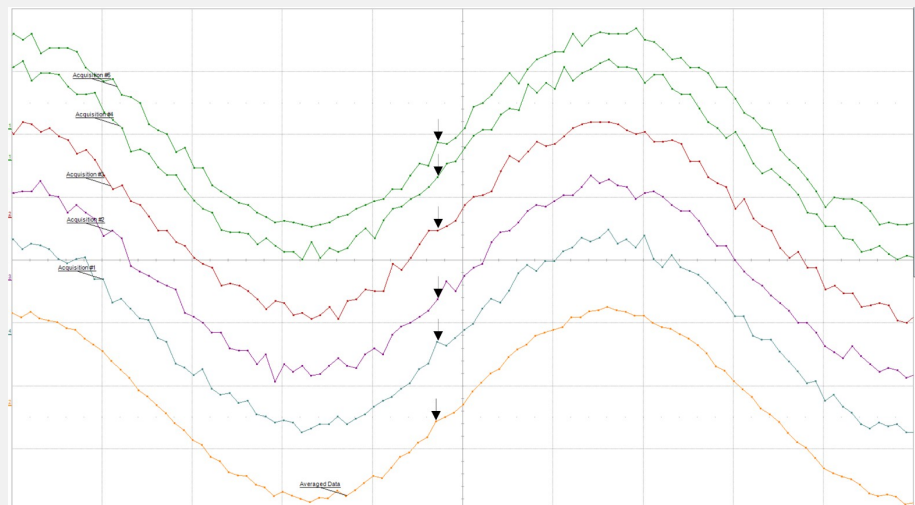

Figure 4: Summed ensemble averaging adds the nth point of multiple acquisitions and then divides the sum by the number of acquisitions to determine the averaged value for the nth point.

What Improvement Can You Expect?

When a signal is averaged additive broadband Gaussian noise will be reduced by the square root of the number of averages. So averaging four acquisitions can improve the signal to noise ratio by two to one. Similarly, non-synchronous periodic signals will be reduced in the average. The degree of reduction depends on the phase variation of the interfering signal from acquisition to acquisition. Signals synchronous to the trigger, such as distortion products, will not be reduced in amplitude by averaging.

Example of Averaging

Figure 5 shows a typical example where averaging is useful. The acquired signal (left grid) is a linearly damped sine wave with additive vertical noise. Note that as the sine amplitude decreases in the presence of fixed amplitude noise it disappears into the noise.

Averaging 1024 acquisitions increases the signal to noise ratio to a point where the sine wave can be discerned through the whole waveform.

The principle limitation of the summed ensemble average is that it requires multiple repetitive waveforms with a stable trigger.
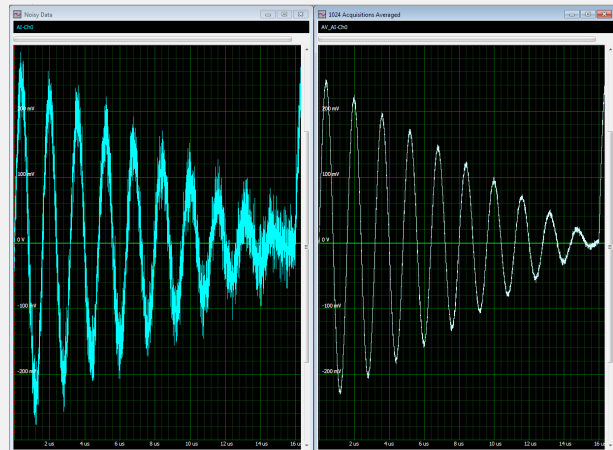

Figure 5: An example of using summation averaging to improve signal to noise ratio.With 1024 acquisitions averaged, the sine wave is visible above the noise.

Moving Average

The moving average, sometimes called a 'boxcar' average or smoothing, takes an average of a user defined number of symmetrically placed adjacent samples. For a sample size of five the process is defined mathematically by the following equation:

Averaged Sample = [sample (x-2) + sample (x-1) + Sample (x) +sample (x+1) + sample (x=2)] / 5

The number of samples used in the average must be matched to the period of variations in the waveform otherwise the moving average can reduce the amplitude of narrow features.

Figure 6 provides an example of using a moving average of 50 adjacent samples, shown in the left hand grid. Note the smoothing and elimination of noise compared to the acquired waveform shown in the grid on the right.

The samples are uniformly weighted and the average is taken running along the samples of the acquisition. The advantage to a moving average is that the signal need not be repetitive. The tradeoff is that in creating a smoothed waveform there is a corresponding loss of high-frequency information. Care must be exercised in setting the number of samples averaged.
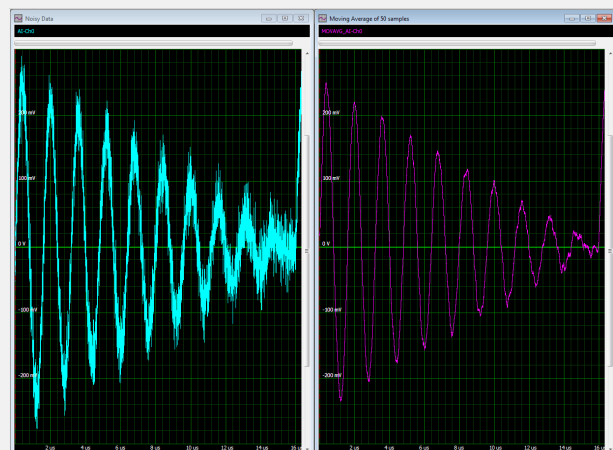

Figure 6: An example of using a moving average using 50 adjacent samples of the acquired signal (left hand grid) showing the resultant smoothing in the right hand grid.

Fast Fourier Transform

The Fast Fourier Transform (FFT) maps the acquired waveform from the time domain (amplitude versus time) into the frequency domain spectrum (amplitude versus frequency). This allows observation of the frequency components which make up the signal. The FFT does not improve signal quality directly but it shows the structure of the signal and provides information on how to remove undesirable spectral components.

▶ **Application Note**

The frequency spectrum resulting from an FFT has a discrete time axis just as the time domain signal has discrete time samples. Samples in the spectrum, often referred to as bins or cells, are spaced by the resolution bandwidth (∆f) which is inversely proportional to the acquired signals record length. So to increase the frequency resolution of the FFT spectrum you have to increase the acquired signals record length.

The frequency range or span of the spectrum display is one half of the sample rate at which the signal was acquired. So to increase the span you must increase the sampling rate.

Vertical scaling of the FFT in SBench 6 can be in linear units of Volts or in logarithmic units expressed in decibels (dB). The decibel scale can be referenced to full scale of the digitizer range (dBFS), one milliwatt (dBm), 1 µV (dbµV), or to the largest peak in the spectrum which is assumed to be the modulated carrier (dBc).

Weighting Functions

The theoretical Fourier transform assumes that the input record is of infinite length. A finite record length can introduce discontinuities at its edges. This introduces pseudo-frequencies in the spectral domain, distorting the real spectrum. When the start and end phase of the signal differ, the signal frequency falls within two frequency bins, broadening the spectrum.

The broadening of the spectral base, stretching out in many neighboring bins, is termed leakage. Cures for this are to ensure that an integral number of periods is contained within the display grid or that no discontinuities appear at the edges. Both need a very precise synchronization between waveform signal frequency and digitizer sampling rate and an exact setup of the acquisition length, what is normally only possible in lab and not with real world signals. Another is to use a window function (weighting) to smooth the edges of the signal.

In an effort to minimize these effects a weighting function is applied to the acquired signal which forces the end points of the record to zero. The FFT in SBench 6 offers the user a choice of any of eight weighting functions. Weighting functions have the effect of changing the shape of the spectral lines. One way of thinking about the FFT is that is synthesizes a parallel bank of band pass filters spaced by the resolution bandwidth. The weighting function affects the shape of the filter frequency response. Figure 7 compares the spectral responses for four of the most commonly used weighting functions.
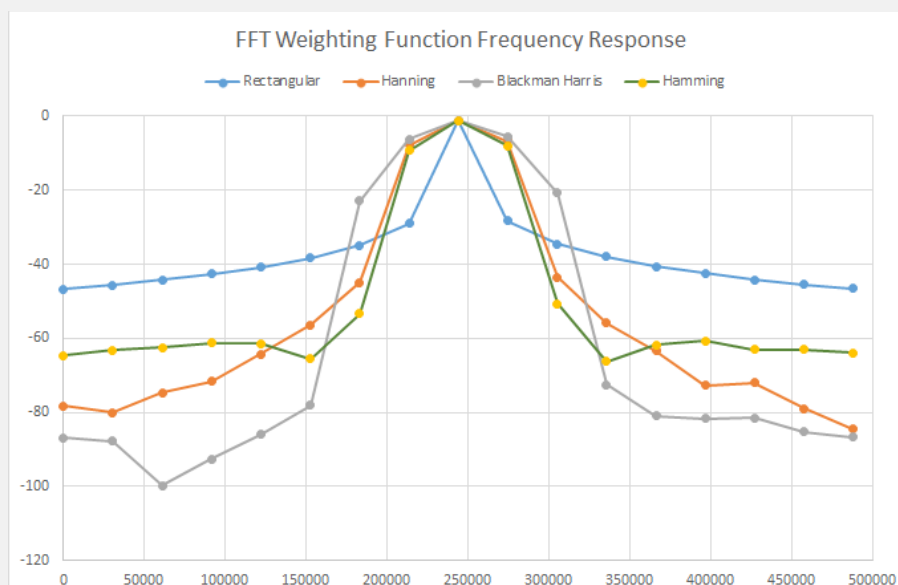

Figure 7: A comparison of the spectral shape for four most commonly used weighting functions including rectangular (no weighting), Hanning, Hamming, and Blackman-Harris.

Table 1 shows the key characteristics of each weighting function.

▶ Application Note

Ideally, the main lobe should be as narrow and flat as possible to represent the true spectral components, while all side lobes should be infinitely attenuated. The window type defines the bandwidth and shape of the equivalent filter to be used in the FFT processing. The maximum side lobe amplitudes of the spectral response are shown in table 1. Minimum side lobe levels help discriminate between closely spaced spectral elements.

As mentioned previously, the FFT frequency axis is discrete having bins spaced in

FFT Window Filter Parameters				
Window Type	Highest Side Lobe (dB)	Scallop Loss (dB)	ENBW (bins)	Coherent Gain (dB)
Rectangular	-13	3.92	1.0	0.0
Hanning	-32	1.42	1.5	-6.02
Hamming	-43	1.78	1.37	-5.35
Blackman-Harris	-67	1.13	1.71	-7.53

Table 1: Characteristics of the four most commonly used weighting functions

multiples of the resolution bandwidth. If the input signal frequency falls between two adjacent bins the energy is split between the bins and the peak amplitude is reduced. This is called 'picket fence' effect or scalloping. Broadening the spectral response decreases the amplitude variation. The scallop loss column in table 1 specifies the amplitude variation for each weighting function.

Weighting functions effect the bandwidth of the spectral response. Effective noise bandwidth (ENBW) specifies the relative change in bandwidth relative to that of rectangular weighting. Normalizing the power spectrum to the measurement bandwidth (power spectral density) requires dividing the power spectrum by the ENBW times the resolution bandwidth ($\Delta f*ENBW$).

Coherent Gain specifies the change in spectral amplitude for a given weighting function relative to rectangular weighting. This is a fixed gain over all frequencies and can be easily normalized.

The rectangular weighting function is the response of the acquired signal without any weighting at all. It has the narrowest bandwidth but exhibits rather high side lobe levels. Because the amplitude response is uniform over all points in the acquired time domain record it is used for signals that are transient in nature (are much shorter than the record length). It is also used when the best frequency accuracy is required.

The Hanning and Hamming weighting functions have good, general purpose responses providing good frequency resolution along with reasonable side lobe response.

Blackman-Harris is intended for the best amplitude accuracy and excellent side lobe suppression.

FFT Application Example

Figure 8 shows a typical example where the FFT is useful. The signal from an ultrasonic range finder is acquired using a broadband instrumentation microphone and a Spectrum M4i series 14 bit digitizer.

The acquired time domain signal is in the left grid. The time domain record includes 16,384 samples taken at a sample rate of 3.90625 MHz. Its duration is 4.2 ms. The resultant FFT (right grid) has 8,192 bins spaced at a resolution bandwidth of 238 Hz (the reciprocal of the 4.2 ms record length) for a span of 1.95 MHz (half of the sampling rate). The spectrum in the lower right is the full span. The zoom view in the upper right shows only the first 100 kHz, allowing a better view of the main spectral components.
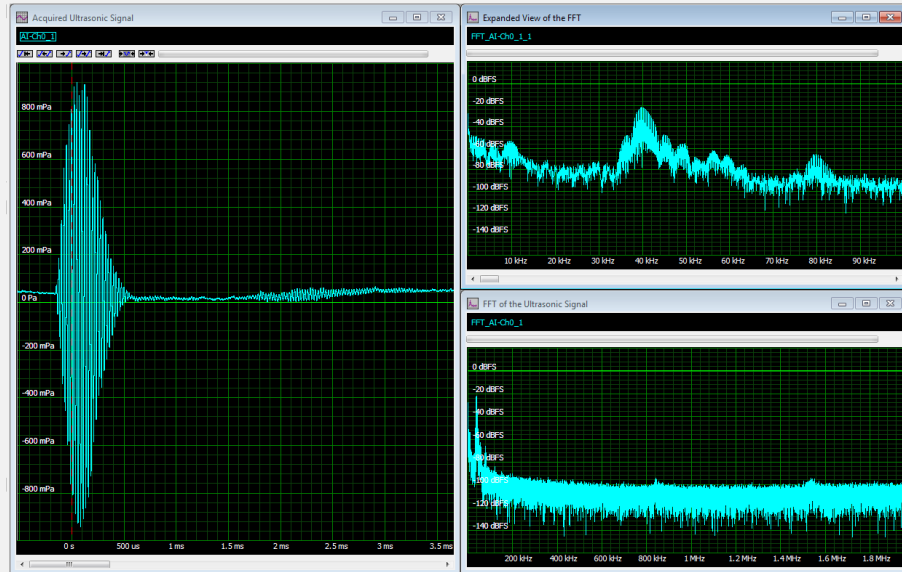

Figure 8: A 40kHz ultrasonic pulse (left)and its associated FFTs (lower right full spectrum, upper right zoom view) showing the principal spectral response at 40 kHz along with undesired lower and higher frequency components.

The FFT allows us a better understanding of the elements that make up this signal. It is a transient signal whose duration is less than the acquired record length. In this case rectangular weighting has been used. The primary signal is the 40 kHz burst which is clearly the frequency component with the highest amplitude. There is an 80 kHz signal which is the second harmonic of the 40 kHz component. Its amplitude is about 45 dB below the 40 kHz signal component. There are also a lot of low frequency components between 0 Hz and 10 kHz. The highest components, those near DC, are ambient noise found in the room where the device was used.

The goal here is to be able to measure the time delay between the transmitted burst and the 40 kHz reflection. To improve this measurement we can remove the signal frequency components outside of the range of the

40 kHz components. This spectral view will be our guide in setting up a filter to remove the unwanted frequency components.

Filtering

The SBench 6 professional software includes finite impulse response (FIR) digital filters in low pass, band pass, or high pass configuration. Filters can be created based directly in SBench 6 by entering the desired filter type, cutoff frequency or frequencies, and filter order. SBench 6 advises if the filter is not realizable and suggests solutions to correct the problem. Alternatively, you can enter filter coefficients derived from another source. These filters can be applied to the acquired signal and we can compare the results with the raw or averaged acquisitions. In Figure 9 an FIR band pass filter with cutoff frequencies of 30 and 50 kHz has been applied to the acquired signal.

The upper left grid contains the raw waveform. Below that

Figure 9: : Showing the effects of applying an FIR band pass filter with cutoff frequencies of 30 to 50 kHz to the 40 kHz ultrasonic signal.The raw waveform and its FFT are on the left side of the display.The filtered signal and its FFT are on the right side. Note the flatness in the filtered baseline, the result of eliminating the low frequency pickup.

is the FFT of the raw signal which we have seen before. The upper right grid contains the band pass filtered waveform. The FFT of the filtered signal is in the grid on the lower right. Note that the band pass filter has eliminated the low frequency pickup and the 80 kHz second harmonic. The time domain view of the filtered signal now has a flat baseline. The reflections are more clearly discernible, which is the goal of the processing. Again the FFT provides great insight into the filtering process.
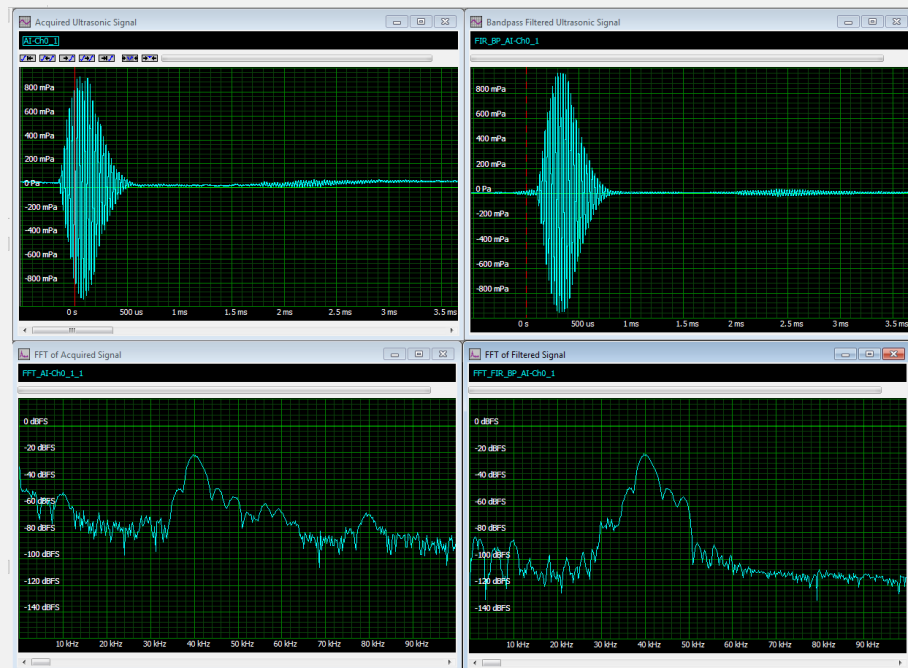
Histograms

So far we have looked at data in both the time and frequency domain. Each of these views adds something to our understanding of the data we are acquiring. We can also view the data in the statistical domain which deals with the probability that certain amplitude values occur. This is conveyed by the histogram which plots the frequency of occurrence versus amplitude value. The histogram is a finite record length estimate of the signals probability distribution. SBench 6 offers the ability to create histograms of acquired waveforms. Some examples are shown in Figure 10 including sine, triangle, and noise waveforms and their associated histogram distributions.

The top row of waveforms show a sine, triangle and noise waveforms. Below them are the associated histograms.

The horizontal axis of the histogram represents the amplitude of the signal. The vertical axis

shows the number of values in a small range of values (binning).

Each histogram distribution is distinct and the difference is related to the signal characteristics. The distribution of the sine wave shows high peaks on either extreme and a saddle shaped mid-region. The reason for this shaping is that the sinewave's rate of change varies through each cycle. It is highest at the zero crossing and slowest at the peaks. If the sine is sampled at a uniform sampling rate there will be more samples at the positive (rightmost peak in the histograms) and negative (left most peak) peaks and the fewest samples at the zero crossing (in the center of the histogram horizontally).

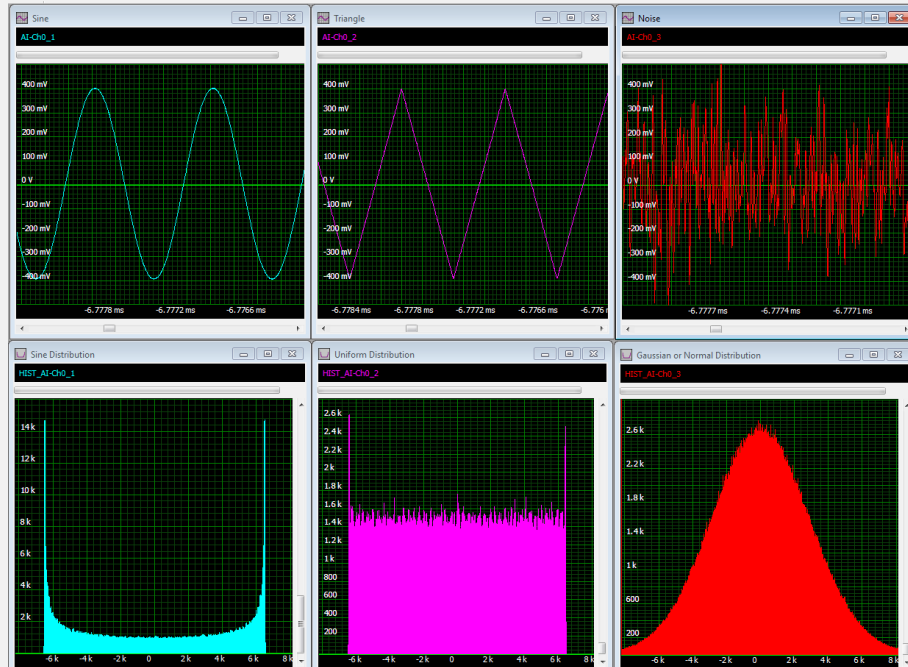The triangle wave has a constant slope, either positive or negative. The resulting

Figure 10: Some example of common waveforms including sine, triangle, and noise and their associated histograms

histogram has a uniform distribution except at the extremes. The peaks are there because the signal generator has limited bandwidth which rounds the peaks and a greater number of samples are acquired at those points.

The histogram of the noise signal results in a Gaussian or normal distribution. I probably don't have to spend too much time on that distribution for those who have taken a statistics course. The unique characteristic of the Gaussian distribution is that it is not bounded. The other distributions have amplitude limits, the horizontal range is fixed. The Gaussian distribution has 'tails' which extend theoretically to infinity (in practical instruments the tails will be limited by clipping in the analog to digital converter).

So histograms tell their own stories about acquired signals. They are good at showing up asymmetries (distortion) and low probability 'glitches' in waveforms. Figure 11 shows a histogram of a sine wave with a glitch at the zero crossing.
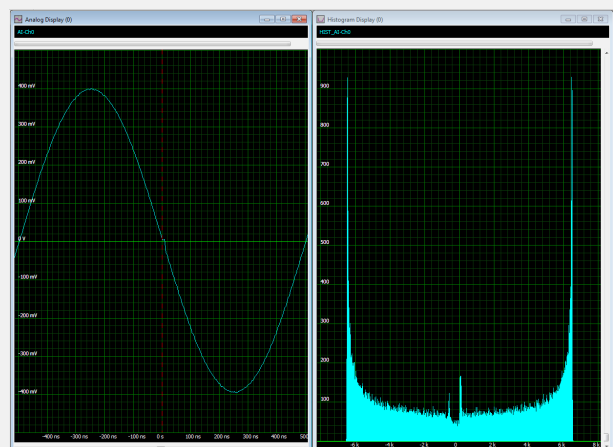

Figure 11:A sine wave exhibiting a zero crossing glitch and its associated histogram.Note the peaks near the center of the histogram

The histogram clearly shows significant peaks near the zero crossing that are not present in

the sine histogram in Figure 10.

Conclusion

Signal processing tools like analog calculations, averaging, FFT, filtering and histograms help interpret acquired data and derive secondary signals yielding greater insight into you data.

Authors:

Arthur Pini   Independent Consultant

Oliver Rovini   Technical Director, Spectrum GmbH

Greg Tate   Asian Business Manager, Spectrum GmbH