November 30, 2011
Dr. Michael Lauterbach

Summary

The digital oscilloscope has become the most powerful and commonly used instrument for design engineers to use when characterizing and troubleshooting electronic devices. As the characteristics of the devices become better with each new generation demands are placed on the oscilloscope to measure these improved signals – often with tighter specs that require higher accuracy measurements. This paper will examine methods which may be used to capture, display and measure signals while achieving maximum accuracy.

The Basic Problem

All real time oscilloscopes incorporate a wide band amplifier and ADC in the front end of the scope. As signal speeds have become faster, the demand from the market has been for higher speed amplifiers and faster sampling ADCs. Yet the inherent noise of a wideband amplifier tends to increase with the square root of the bandwidth. No engineer wants to spend his time characterizing the signal quality of his device and then discover later that he was actually characterizing the noise added to his signal by the front end of an oscilloscope. In the worst case the engineer may spend his time troubleshooting noise sources he believes are in the device he is designing and perhaps discover – after a competitor has beaten him to market – that he had an excellent design and the observed noise was added in the acquisition process.

Solutions

Using a low noise, high resolution oscilloscope

The most obvious solution is that vendors of oscilloscopes should create a new generation of wide bandwidth, low noise, and fast sampling rate instruments. The "sweet spot" of the oscilloscope market is in the 500 MHz range. There are many people who buy/use scopes above and below this bandwidth, but if a manufacturer could offer a scope with 500 MHz bandwidth (or a bit better) then the majority of applications would be addressed by the new technology. Of course such an instrument should also have options for adding digital channels for mixed signal/embedded applications, the ability to view signals in the time, frequency and statistical domains and some optional packages to trigger and decode the ubiquitous serial data signals (USB, RS232, UARTs, I^2C, SPI, etc.). This is the approach taken by the LeCroy HRO High Resolution Oscilloscopes. They have very low noise front end amplifiers and signal path. They incorporate 2GS/s ADCs which have 12 bit resolution. Overall the scopes can achieve a signal-to-noise ratio of 55 dB compared to 35 dB commonly found in 8 bit oscilloscopes. This has dramatic positive effects in viewing and measuring signals in both the time and frequency domains. It also adds accuracy to statistical domain characterization by reducing the effects of acquisition noise when looking for worst cast signal parameters or accumulating histograms of parameter values. The user can characterize his actual signal with very small impact from the front end of the oscilloscope.

Signal Averaging

A second solution to the problem of characterizing signals in the presence of noise is to use signal averaging. This is a well known technique available in most oscilloscopes. The user sets up a trigger and acquires the noisy signal many times. A math buffer adds up N acquisitions and divides by N. This reduces the level of noise. Using this technique the oscilloscope can display the latest acquisition (with the noise) as well as the averaged signal, thereby observing the effect of the averaging. When the desired level of noise reduction is achieved the user can view the shape of the averaged signal and make measurements. There are a few caveats to keep in mind. First, the signal must be repetitive. This technique does not work for observing or characterizing unstable signals or intermittents. Also, the signal feature the user wants to view/measure must be stable in time with respect to the trigger event that initiates capture of the waveform. Thirdly, the property the user is measuring is one that should be unrelated to signal quality because the averaging process dramatically changes the quality of the signal by removing noise – both the acquisition noise of the oscilloscope front end and the "true" noise that is part of the signal. It is clear that measurements of signal quality such as Signal to Noise Ratio, eye patterns and jitter (where a prime source of jitter is the vertical noise in the signal) should not be done using signal averaging. But even measurements of parameter values and parameter stability will also be affected by removing the noise which is part of the signal. Measurements of rise time, pulse width, overshoot and other signal parameters will be changed by the removal of part of the signal. A final factor to keep in mind is that the signal is never truly repetitive nor is the trigger time completely stable. The noise in the signal will cause the trigger time to shift from acquisition to acquisition thereby distorting the averaged waveform.

Filtering

A third solution to the problem is to apply filtering to remove noise. Many oscilloscopes have this capability. Both analog and digital filters are built into many models from a variety of vendors. This method does not rely on having a repetitive waveform or on trigger stability. A single acquisition is used and the scope applies a filter to the ADC samples. Figures 1a and 1b show an example comparing the use of averaging and filtering.

Figure 1a: *The upper trace is the raw samples from the ADC. The lower trace is 1000 averages. The noise (from the signal and from the acquisition process) are greatly reduced by averaging.*

Figure 1b: *The same signal on the upper trace as in Figure 1a. But the lower trace is a single acquisition – no averaging. A digital filter (defined in the lower right corner) has been applied to the raw data. The filter removes noise above 8 MHz – both "real" noise that was part of the signal and noise added by the acquisition process. Lower frequency noise is not removed.*

The signal is a damped sine wave with substantial noise. Some of the noise is due to the fact that the signal is noisy and some of the noise comes from the acquisition process (oscilloscope amplifier noise). A quick glance reveals the power of filtering – the averaged signal in Figure 1a looks quite similar to the filtered signal in Figure 1b. A closer examination of Figures 1a/b reveals the drawback of filtering. You can see the details of the damped sine wave look different, especially toward the center of the screen. The filtered waveform shows some interesting "bumps" riding on the sine wave. Are those real? Does the underlying signal have such structure? Or are the "bumps" artifacts added by the filter? The user doesn't know. In this particular example the underlying damped sine wave is smooth. But when seeing this waveform on his oscilloscope maybe the user will believe the underlying signal should be a smooth, damped sine wave – but maybe he will believe there is crosstalk that is adding unexpected signal content (the "bumps"). If the unexpected content is not correlated to the trigger time (which is very likely) then this important signal content would be hidden by signal averaging but "revealed" by the process of filtering a single shot acquisition. In this example the filter has a -3dB bandwidth of 8 MHz as shown in the lower right corner of Figure 1b. This still lets through some noise – which is causing the bumpy structure. If the filter was changed to a lower bandwidth then the user would see the filtered signal in Figure 1b more closely match Figure 1a. How does the user know which filter to use? The answer is that he doesn't – unless he already knows the correct shape of his waveform before he acquires/filters it. You can use filters to make your waveform conform to the shape you imagine it should be – whether that shape is the true shape of the actual waveform or not. Filters remove acquisition noise but they also remove real signal content. The filter reshapes – some people would say it distorts – the signal.

Aside from the quandary of not knowing which filter is the “right” one to use when viewing your signal, the process of filtering also has some of the same drawbacks as averaging. Because filtering removes part of the signal (the noisy part) an oscilloscope user cannot make meaningful signal quality measurements. Signal to noise ratio, eye diagrams and jitter measurements will obviously be incorrect if a filtered signal is used to make the measurement. More subtly, if the user is testing a device to meet a worst case specification for rise time, overshoot or other signal parameters a filtered signal should not be used because the filter may be removing exactly the piece of the signal content that the worst case spec is supposed to catch/prohibit.

Some examples of Real Signals

Testing a Power Supply

The most basic signal is DC. A majority of engineers do not design power devices but they often have to check that their power devices do not have problems such as too much noise and ripple. An example of a power supply that has substantial problems is shown in Figure 2. Any device, power supply or otherwise, that has large, obvious faults can be viewed, characterized and debugged using 8 bit oscilloscopes.

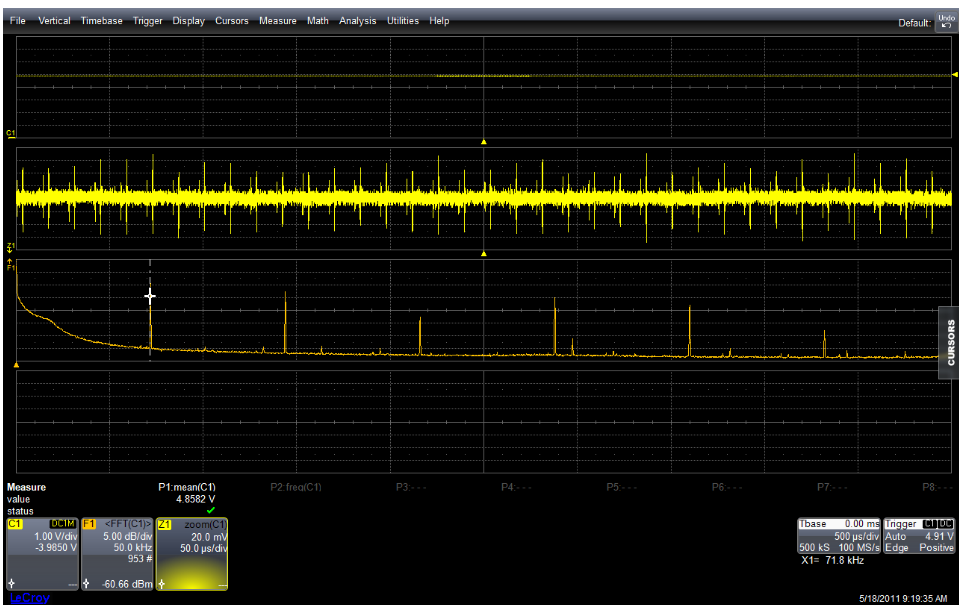

Figure 2: *The top trace is a nearly flat line from a DC power supply. But the zoom (middle trace) reveals substantial noise. The lower trace is an FFT which will help the user identify the source of the noise. Signals with large, obvious problems such as this can be characterized and debugged using 8 bit oscilloscopes.*

The next example is shown in Figures 3a and 3b which were both acquired using 8 bit oscilloscopes. Again we are looking at a DC power supply. The engineer will come to the same conclusion no matter which brand of 8 bit oscilloscope is used. The power supply has 8 divisions of peak to peak noise at 20 mV per division – a total of 160 mV of peak to peak noise.

Figure 3a: *The flat line in the upper trace is from a DC power supply. The middle trace is a zoom, showing 8 divisions of noise at 20 mV/div (red box). The lower trace is an FFT.*

Figure 3b: *The same signal as Figure 3a but acquired on a different brand of 8 bit oscilloscope. The zoom trace shows the same 8 divisions of noise at 20 mV/div (red box).*

Perhaps an engineer who tested the power supply twice, using two brands of oscilloscopes would call up his power supply vendor and berate him for delivering a device that had such poor performance. He might ship the power supply back to the vendor and ask for a replacement. Or he might look at the FFT of the noise in the power supply and try to troubleshoot the source. In either case he would be wasting his time. A hint at the true nature of the problem would come if he compared the FFT of the power supply noise when using the first scope to the FFT of the power supply noise when using the second scope. The FFT's would not match. One scope would claim the noise had different frequency characteristics than the other scope. So maybe the next argument would be with the oscilloscope vendors? Which scope is correct? This would be another waste of time.

Figure 4 resolves the problems discussed in the previous paragraph. If the power supply is tested using a low noise high resolution oscilloscope the user sees only 5 divisions of noise at 10 mV/div.   A very reasonable (for this application) 50 mV of peak to peak noise. The 160 mV reported by both of the 8 bit scopes was noise added to the signal by the 8 bit oscilloscopes. This resolves the argument with the power supply vendor. It also resolves the argument with the scope vendors. Since the FFT was primarily showing the frequency content of noise in the front end amplifier both scopes were showing accurate FFTs – but the two front end amplifiers have different designs and therefore have different noise characteristics.

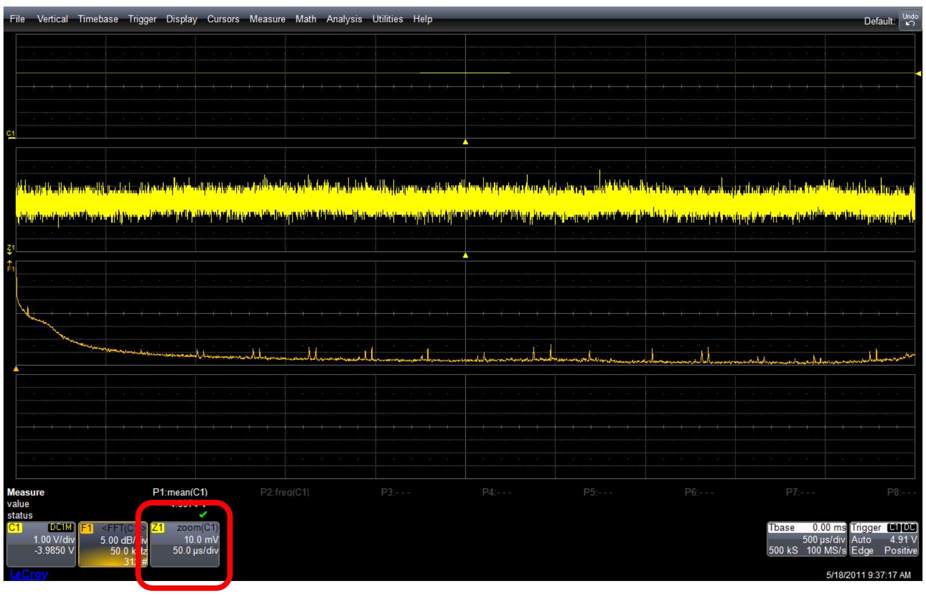

Figure 4: *The same DC power supply as shown in Figure 3a/b but this time acquired with a low noise, high resolution oscilloscope. Note the zoom (center trace) has 5 divisions of noise at 10 mV/div (red box). This is an accurate measurement of the power supply noise- the front end noise from the oscilloscope is much smaller than 50 mV.*

Could the engineer have achieved the correct measurement of 50 mV peak to peak noise using an 8 bit oscilloscope by using either averaging or filtering? The answer is, yes-but you have to know the right answer before you acquire the data. If you know the true signal has 50 mV of noise than you could set up averaging and stop the averaging process when the peak noise reached 50 mV. Similarly you could take a single shot of the signal with an 8 bit oscilloscope and try a variety of filters until you get the answer you know is correct. Of course real engineering doesn't happen this way.

Testing a Signal with low SNR

The example shown in Figure 5a/b is a medical signal. In such applications the signal often has low SNR. One can certainly see the basic shape of the signal but the interesting portion, which allows a trained observer to detect problems, is hidden in the noise. Applying a filter removes the noise (Figure 5b) revealing a signal shape that has some interesting structure. Is it the true shape of the signal? It is hard to know (unless you already know the right answer before you acquire the signal).

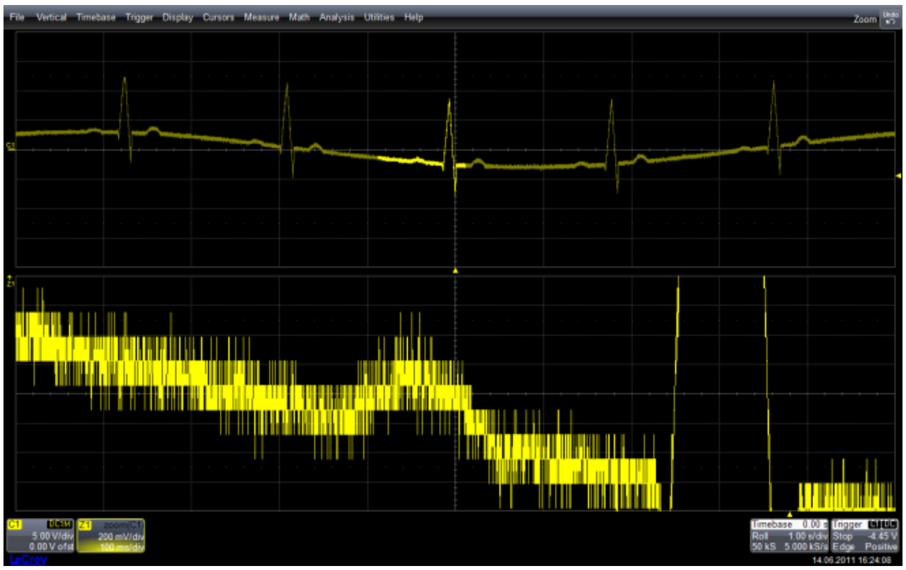

Figure 5a: *A signal that contains substantial noise – some of which is part of the fundamental signal and some from the acquisition process. The basic acquisition is the upper trace. The lower trace is a zoomed portion.*

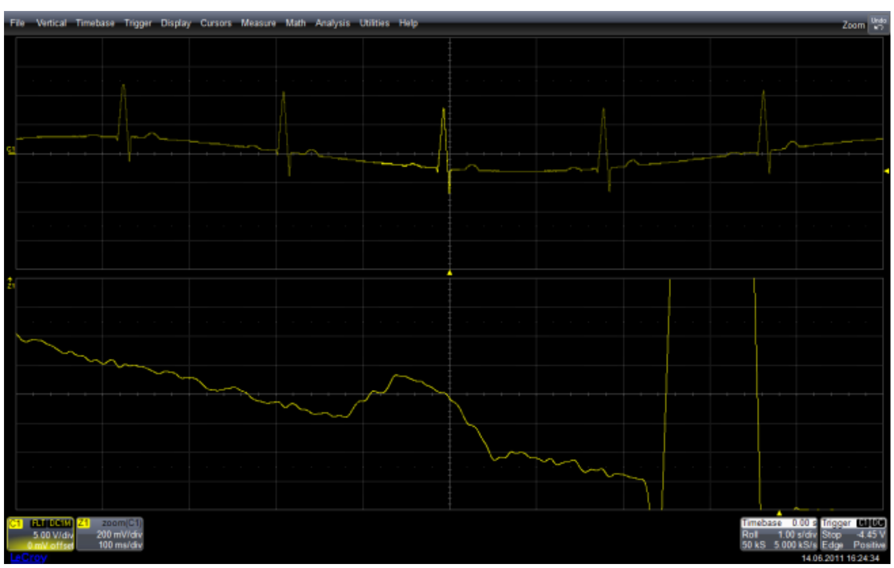

Figure 5b: *The same signal as Figure 5a but the signal has been filtered to remove high frequency noise (both noise in the signal and noise from the acquisition process). Are the bumps shown in the zoom the real shape of the signal? Or are they artifacts from the filter?*

The same signal is acquired by a low noise, high resolution HRO in Figure 6a/b. The zoom of the unprocessed ADC samples in Figure 6a shows substantially less noise than in Figure 5a. Now we know the true signal shape has lower noise. And of course the user can apply averaging or filtering to a waveform acquired by an HRO in the same fashion as an 8 bit scope (Figure 6b). The difference is that the processing is done starting with a waveform that has much less acquisition noise. In this case, a trained observer might understand something interesting and useful by seeing the overshoot shown in the middle of the screen.

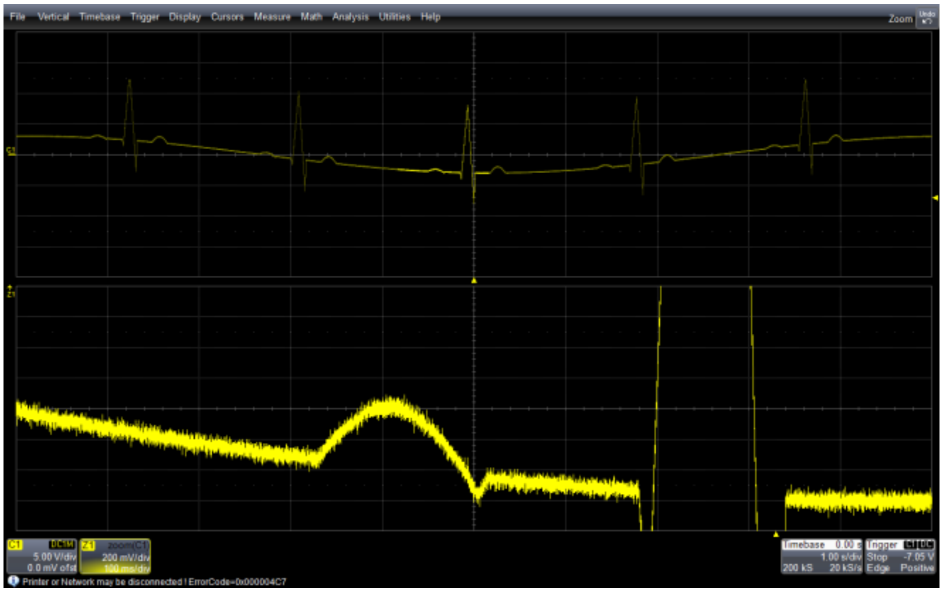

Figure 6a: *The same signal as in Figure 5a/b but this time it is acquired with a low noise, high resolution HRO oscilloscope. The noise level is much lower. There is a small contribution of noise from the acquisition process but primarily the scope is showing a clear view of the real signal.*

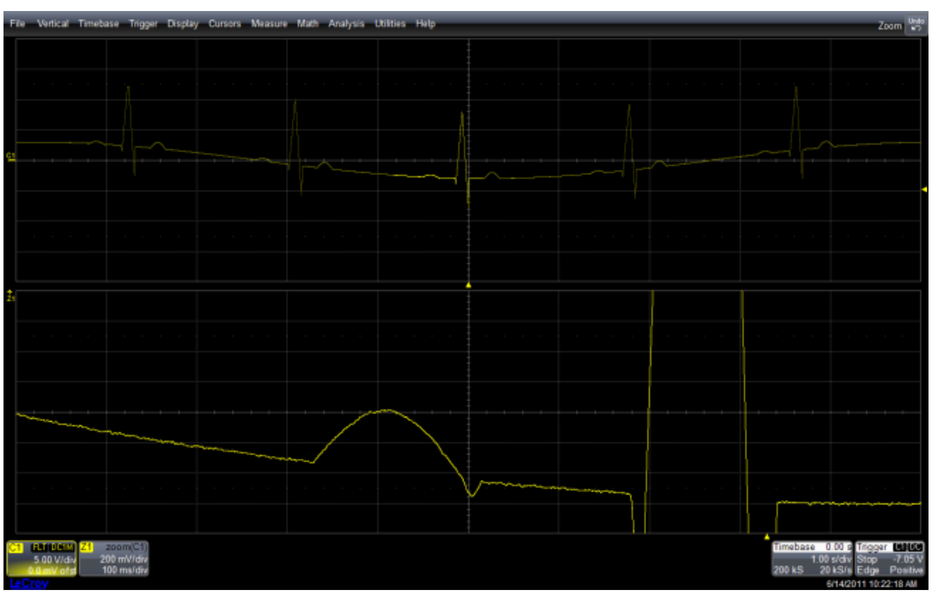

Figure 6b: *The same signal, acquired on the same HRO as in Figure 6a. But this time a filter has been used to remove noise (mostly noise that was part of the signal). An interesting overshoot can be observed in the center of the zoom just after the large bump.*

The basic lesson is one engineers have known for several decades: Garbage In – Garbage Out. If an oscilloscope adds a substantial amount of noise to a signal in the acquisition process then the view and measurements of the signal are likely to be incorrect—even if you try to apply tools such as averaging and filtering. On the other hand, if the scope acquires a clean acquisition of the signal, then it is much more likely the user can apply the analysis capabilities of the scope to get a useful conclusion.

ERES vs. HiRes

The discussion above focused on the general characteristics of acquiring, viewing and measuring signals using low noise HRO oscilloscopes, signal averaging and filtering. It also showed examples of real world situations and methods to obtain more accurate views and measurements. Since many engineers already own oscilloscopes – and often have access to oscilloscopes from more than one company, the final section of this document will focus on the use of two very specific types of signal filtering methods that are commonly available – the Enhanced Resolution (ERES) math function available in many LeCroy oscilloscopes and the High Resolution (HiRes) acquisition mode available in Tektronix scopes. Both features offer Finite Impulse Response (FIR) filters to remove high bandwidth noise from a signal and obtain additional vertical resolution. Aside from that very basic similarity the two techniques are very different.

ERES is a math function invoked by the user who can control the filter length, the number of bits of added vertical resolution and the resulting new, reduced bandwidth of the oscilloscope (as shown in the lower right hand corner of Figure 1b). The math function can operate on any waveform but is usually applied to the ADC samples of a single shot digitized waveform. If desired, the user can view both the true signal shape (with full bandwidth) and the filtered signal shape (using ERES). On the other hand, HiRes is not a math function. It is an acquisition mode. If the user is operating a Tektronix oscilloscope under conditions where the number of points the user requests on the screen is smaller than the number of points generated by the ADC running at its full sample rate then HiRes will use the oversampling of the ADC to generate a sequence of measurements with reduced bandwidth, lower noise and higher vertical resolution. The raw ADC samples are lost and replaced by the new HiRes measurements which are displayed as a waveform for the user.

If the two filters were the same it might be a small difference that one method allowed viewing of the full bandwidth true signal shape and the other one did not. After all, the point of the filters is to remove noise and allow the user to view a modified version of the signal at lower bandwidth and better vertical resolution. However the ERES and HiRes filters are dramatically different.

HiRes is a simple “boxcar” shaped filter. If the user requests 10,000 sample points on the screen for a timebase where the ADC running at full sample rate generates 100,000 points then HiRes will take each set of 10 ADC samples “on the fly”, average them and put the resulting averaged number into memory. The average is a simple summation of N points divided by N. The only user control is the ability to turn HiRes on/off and to request a number of samples on the screen. The scope does not report the number of extra bits of vertical resolution, the filter length or the reduced bandwidth of the waveform. A benefit of this filtering method is that it is very fast and does not require a powerful processor.

The ERES filter is bell shaped. It is implemented on a point-by-point basis to a waveform. The user controls the length of the filter and the scope reports the new bandwidth of the ERES waveform and the number of added bits of vertical resolution. The mathematical process of the calculation is to replace each individual point of a raw waveform with one new point of the filtered waveform. At each point in time, the raw ADC sample taken at that time is given the most weight in calculating the value of the ERES point. The ADC points directly adjacent (one sample period earlier and one sample period later) are weighted somewhat less than that. The weighted summation continues with less weighting for “farther out” samples for the length of the filter as specified by the user. The weighting curve is bell shaped. The only points “lost” in this process are at the very beginning and end of the original waveform – because there are not enough points to perform the weighted average. The reduction in the number of points in the ERES waveform is equal to the filter length. If the original waveform is 100,000 points and the filter length is 10 then the resulting ERES waveform is 99,990 points in length. This contrasts to HiRes where a filter length of 10 reduces a waveform from 100,000 points to 10,000. The shape of the ERES filter is set to maintain exactly linear phase response. This means two events in a waveform will always be separated by the same distance in time before and after the filtering. The method of implementation also means that a waveform feature digitized with a certain number of samples will continue to have that number of samples in the ERES waveform. This is particularly important in clocks, data and any type of digital signal. If any brand of scope acquires 10 samples on an edge, the shape of the edge is accurately defined and displayed to the user. But if a new waveform is formed by decimating the original waveform by a factor of 10 (or more) then the new waveform has only one sample on the edge (or less). This leads to a mathematical artifact called Gibbs Ears. The undersampled edge does not have enough information for the scope to know its shape. If a Sin(x)/x interpolation method is used to estimate the signal shape then fictitious pre-shoot and overshoot are added to the true signal shape. Figure 7a shows this type of problem.

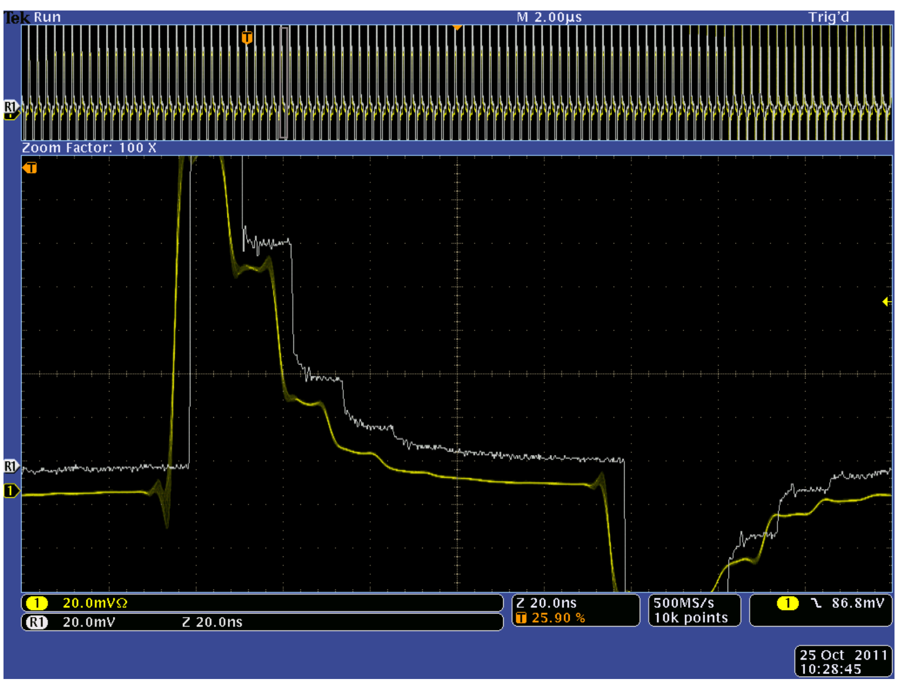

Figure 7a: *The grey waveform is 100k samples acquired using Normal acquisition mode at 5 GS/s and saved as a reference. The lower trace in yellow is the same signal acquired on the same timebase using HiRes mode. The user has asked for a display of 10k points at 500 MS/s. Note the introduction of preshoot prior to the edges.*

Figure 7b demonstrates that ERES filtering does not introduce Gibbs Ears. If a scope user is aware of this possibility then he can ignore the overshoot and undershoot introduced by HiRes on edges. But a less expert scope user might believe the use of HiRes is showing a more accurate version of the waveform and think the overshoot is real.

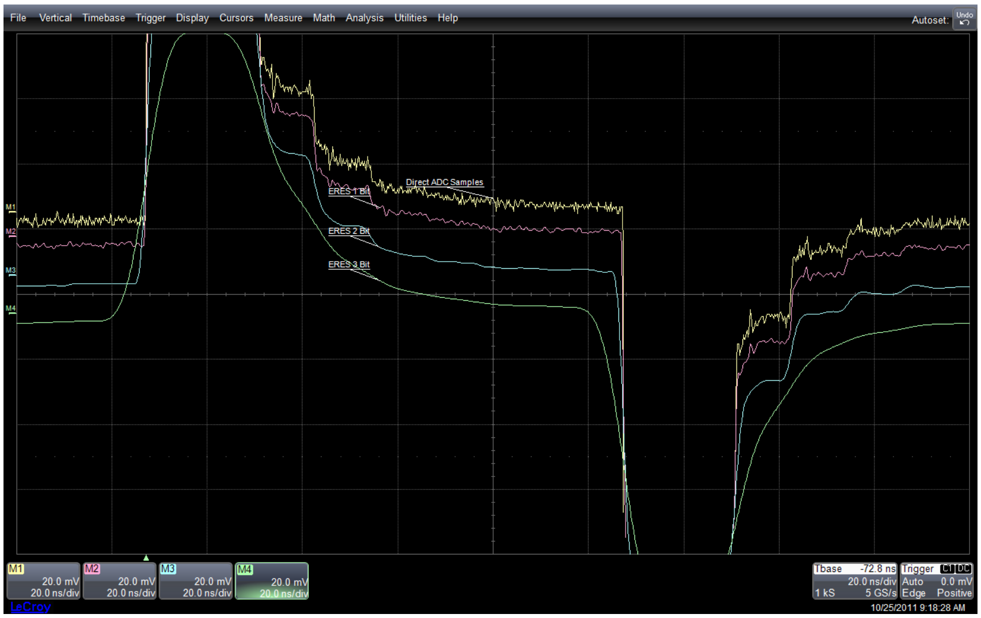

Figure 7b: *The same waveform as Figure 7a. In this case the acquisition at 5 GS/s is the top trace shown in yellow. The same acquisition is also displayed using ERES filters of different lengths to obtain 1 (2nd trace in red), 2 (3rd trace in blue) and 3 (4th trace in green) extra bits of resolution No preshoot is introduced.*

Advanced Application of ERES

In many cases a filter such as ERES is used to remove noise – both acquisition noise and signal noise – which obscures a lower frequency portion of the signal. But it is also possible the user simply wants to use the extra bits of vertical accuracy available with ERES to view/measure more accurately a waveform that has very little noise. If such a signal is repetitive, the oscilloscope can acquire the signal using RIS (Random Interleaved Sampling- sometimes called "equivalent time" acquisition mode) which greatly boosts the sampling rate. The acquired waveform might have an equivalent time sampling rate of 200 GS/s rather than 5GS/s. This means the application of ERES will have much less effect on the usable bandwidth of the oscilloscope. If the waveform is not repetitive, the user can apply interpolation to upsample a single shot acquisition of the waveform by up to a factor of 50 and then apply ERES to the interpolated waveform. An example is shown in Figure 8 using Sin(x) /x interpolation to upsample the waveform by a factor of ten. It is clear that applying ERES to an upsampled waveform preserves more of the signal content than applying the same filter to the raw waveform.

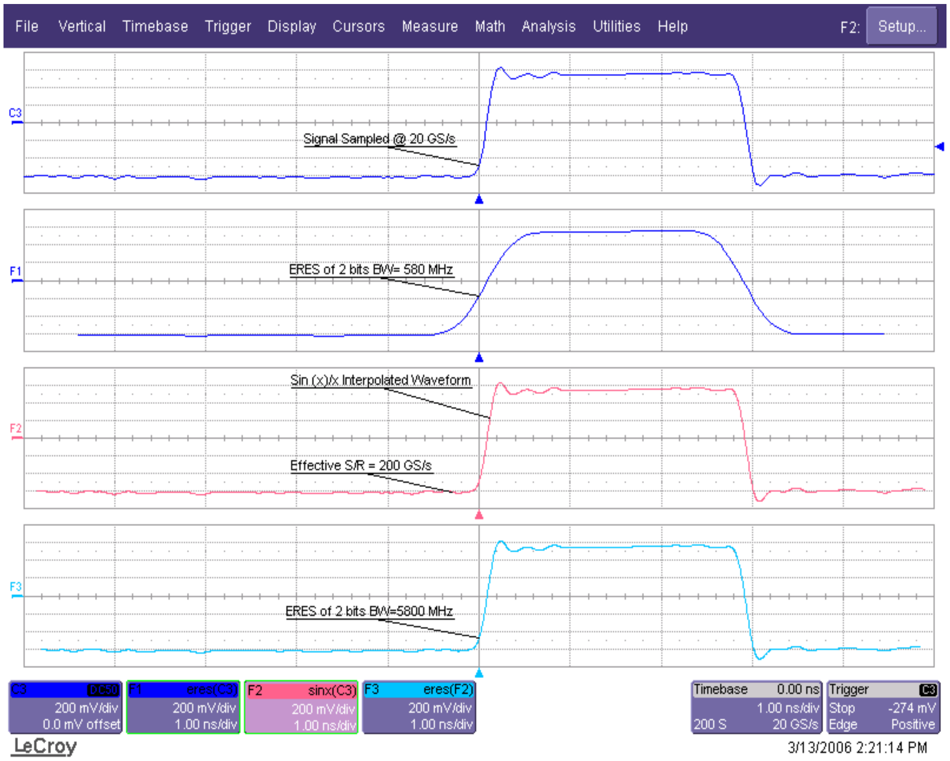

Figure 8: *The top trace shows a waveform acquired at 20 GS/s. In the 2nd trace ERES filters that waveform. The 3rd trace is the original waveform with Sin(x)/x interpolation applied to upsample by a factor of 10. The final waveform is the result of applying ERES to the upsampled waveform.*

Summary

If an engineer needs to make more accurate measurements the most straightforward way to accomplish this is to acquire the signal using a low noise, high resolution oscilloscope. Getting the cleanest possible acquisition is the best way to achieve more detailed views and more precise measurements. It also avoids loss of time in “troubleshooting” noise or artifacts introduced by the acquisition process. In cases where the user already knows the correct shape of a signal or already knows the answer he wants from a measurement filtering can be used to change the signal shape and obtain the desired result. It is also possible to using averaging or filtering to reduce noise in cases where the final answer is not known but the scope user needs to be aware these processes remove part of the real signal and therefore he must be cautious about interpreting the results.