

Enhanced Resolution Filtering

The available sampling rate of LeCroy oscilloscopes is often higher than that required for the analyzed signal's bandwidth. Oversampling, particularly pronounced in the long-memory models, can be used to increase the displayed trace's effective resolution: the ability of the instrument to distinguish closely spaced voltage levels. This is done by filtering the digitized signal using Enhanced Resolution, available in the optional WP02 Advanced Math Package.

Although similar to signal smoothing using a simple moving-average filter, enhanced resolution filtering is more efficient both in terms of bandwidth and the superior passband characteristics that result. And on waveforms with single-shot characteristics, it can be used instead of successive trace averaging.

Advantages

Enhanced resolution filtering improves two important characteristics of the oscilloscope.

Resolution is improved by a fixed amount for each filter. This true increase in resolution occurs whether or not the signal is noisy, or whether a single-shot or repetitive signal.

Signal-to-noise ratio (SNR) is improved in such a way as to be dependent on the form of the noise in the original signal. This is because the enhanced resolution filtering decreases the bandwidth of the signal, therefore filtering out some of the noise.

Implementation

The oscilloscope's set of constant-phase, FIR (Finite Impulse-Response) filters provide fast computation, excellent step response in 0.5 bit steps, and minimum bandwidth reduction for resolution improvements of between 0.5 and 3 bits. Each step corresponds to a bandwidth reduction of a factor of two, allowing easy control of the bandwidth/resolution trade-off. The parameters of the six filters are given in the following table:



FIR Enhanced-Resolution-Filter Parameters		
Resolution Increase (Enhancement)	-3 dB Bandwidth (\times Nyquist)	Filter Length (Samples)
0.5	0.5	2
1.0	0.241	5
1.5	0.121	10
2.0	0.058	24
2.5	0.029	51
3.0	0.016	117

With low-pass filters, the actual SNR increase obtained in any particular situation depends on the power spectral density of the noise on the signal. The improvement in SNR corresponds to the improvement in resolution if the noise in the signal is white — that is, if it is evenly distributed across the frequency spectrum. If the noise power is biased towards high frequencies, the SNR improvement will be better than the resolution improvement. Whereas the opposite may be true if the noise is mostly at lower frequencies.

SNR improvement due to the removal of coherent noise signals — feed-through of clock signals, for example — is decided by the fall of the dominant frequency components of the signal in the passband. This is easily ascertained using Spectral Analysis.

The filters have a precisely constant zero phase response. This has two desirable properties. First, the filters do not distort the relative position of different events in the waveform, even if the events' frequency content is different. And second, because the waveforms are stored, the delay normally associated with filtering (between the input and output waveforms) can be exactly compensated during the computation of the filtered waveform.

All the filters have been given exact unity gain at low frequency. Enhanced resolution should thus not cause overflow if the source data is not overflowed. If part of the source trace were to overflow, filtering would be allowed, but the results in the vicinity of the overflowed data — the filter impulse response length —

would be incorrect. This is because in some circumstances an overflow may be a spike of only one or two samples, and the energy in this spike may not be enough to significantly affect the results. It would then not be desirable to disallow the whole trace.

When to Use It?

In general, enhanced resolution is used to replace the averaging function in situations where the data record has a single-shot or slowly repetitive nature and averaging cannot be used.

There are two particular situations in which enhanced resolution is especially useful. One is when the signal is noticeably noisy and measurements of the noise are not required. The signal can be “cleaned up” by using the enhanced resolution function. The other is when — even if the signal is not particularly noisy — high-precision measurements of the waveform are required (when using Expand with high vertical gain, for example). Enhanced resolution will then increase the resolution of the measurements.


The examples on the following pages illustrate how enhanced resolution can be used.



Examples

Low-pass Filtering


This screen shows the spectrum of a square signal before (top grid) and after (bottom grid) enhanced resolution processing. The result clearly illustrates how the filter rejects high-frequency components from the signal. The higher the bit enhancement, the lower the resulting bandwidth.



Enhanced Resolution

Increasing Vertical Resolution

In this example the bottom trace has been significantly enhanced by a three-bit enhanced resolution function.




Note: The original signal being highly oversampled, the resulting bandwidth is still high enough for the signal not to be distorted.



Reducing Noise

The following illustration shows the effect of enhanced resolution on a noisy signal. The original trace (top grid) has been processed by a two-bit enhanced resolution filter. The result (bottom grid) shows a “smooth” trace, where most of the noise has been eliminated.



Enhanced Resolution

Cautionary Notes

- The enhanced resolution function can only improve the resolution of a trace; it cannot improve the accuracy or linearity of the original quantization by the eight-bit ADC.
- The constraint of good temporal response excludes the use of maximally-flat filters. The pass-band will therefore cause signal attenuation for signals near the cut-off frequency. The highest frequencies passed may be slightly attenuated. The frequency response of a typical enhanced resolution filter (the 2-bit enhancement filter) is illustrated below, indicating the -3 dB cut-off frequency of 5.8% of the Nyquist frequency.
- The filtering must be performed on finite record lengths: data will be lost at the start and end of the waveform, so that the trace becomes slightly shorter after filtering.
- The number of samples lost is exactly equal to the length of the impulse response of the filter used, and thus varies between two and 117 samples. Owing to the oscilloscope's very long waveform memories, this loss — just 0.2% of a 50 000 point trace — is not normally noticed. However, it is possible to demand filtering on a record so short, there would be no data output, and in this case the scope will not allow filtering.

