

Interpolation in Your DSO

Peter J. Pupalaiakis

Product Marketing Manager

LeCroy Corporation, New York

Introduction

Interpolation is an important feature in today's digital oscilloscopes. The main purpose of a DSO is to analyze and view analog waveforms. To do this, the DSO samples a waveform at some finite sample rate, generating a vector of voltages with respect to time. Since this vector represents a set of points (not the actual smooth analog waveform), it is often desirable to modify the acquired waveform by generating samples that are predicted between the actual acquired points. The generation of samples that occur in between actual waveform samples is called *interpolation*. When done properly, this results in a waveform with a higher resultant sample rate that is a closer approximation of the analog waveform under analysis.

This paper will address the two most popular interpolation methods and explain basically how they work. It then addresses how to ensure good interpolation results and how to determine when interpolation is appropriate. The interpolation performance of three high-end scopes are compared with some simple experiments. Finally, differences in interpolation operation are contrasted.

Linear Interpolation

The simplest form of interpolation is linear interpolation. Linear interpolation is performed by assuming that a straight line joins each waveform sample. This is a very simple, but naïve method that provides limited results.

Linear interpolation can be viewed as convolution of an upsampled waveform with a triangular window. The triangular window is formed by generating a triangle that has a height of 1.0 and a width that is twice the sample period. As the window slides to the right, interpolated points are found by calculating the sum of the values of the window multiplied by the value of actual samples at the times where the window and the sample points intersect. The interpolated sample is placed at the time of the window apex. The width of the window defines its memory - the time over which actual samples affect the interpolated samples. Since the width of the window is twice the sample period, only samples that bracket the time of the interpolated sample affect the interpolated value. This convolution can also be realized with a digital filter in an upsampling arrangement. The arrangement shown is a 5 point upsampler. In this arrangement, every new sample produces 5 new ones. The filter output is a function only of the new input and the last one. The filter coefficients are generated by sampling the window. By examining the diagrams, it's easy to see the pattern for developing the filter coefficients.

SinX/X Interpolation

A popular and more complicated form of interpolation is called SinX/X (also referred to as Sync or simply SinX interpolation). SinX interpolation gets its name from the well known shape of the window function used for the convolution. Unlike the narrow pointed triangle of linear interpolation, the window for SinX interpolation is a theoretically never ending damped sinewave.

This window shape derives from an important assumption that Nyquist's criterion has been obeyed in the sampling of the original waveform. In other words, it assumes that all of the frequency content in the analog waveform sampled lies below one-half the sample rate that the waveform was sampled at - a reasonable assumption. When this assumption is made, and the inverse Fourier transform of this assumed spectrum is calculated, the result is this well-known function.

In the general case, this assumption is the best assumption that can be made - but it is not always a correct assumption, as we shall see. As a result, SinX interpolation is truly the most valid interpolation method. This is understood by examining the meaning of Nyquist's criterion. Nyquist said that when all of the frequency content of a signal lies below one-half the sample rate, then the continuous analog signal can be completely determined from the sampled points. SinX interpolation is merely the mechanics for obtaining the continuous analog signal.

SinX interpolation suffers from some mathematical and practical technicalities that make it impossible for this method to be perfect. First of all, the Sinc function goes forever and must be truncated at a point where the truncation error dips acceptably low. This is due to the fact that a truly bandwidth limited signal must have infinite length implying that all sample points must be known for all time. As it turns out, the influence of points further and further from the point being interpolated diminishes rapidly and the truncation provides highly acceptable results. Another drawback is that in a sampled system, noise and artifacts due to DSO architecture like channel interleaving creep in, causing noise and distortion above the Nyquist limit. Again, the errors caused by this can be kept acceptably low.

When Is Interpolation Valid

In pure mathematical terms, neither interpolation method is valid. For example, linear interpolation implies that waveform points are joined by straight lines - a technical impossibility due to bandwidth limitations alone. SinX interpolation, as mentioned, is valid only when Nyquist's criteria is met - which is never fully the case - and when the waveform is infinitely long.

Suffice it to say, interpolation can be valid *to a large extent*. If you are uncomfortable with this concept, consider the fact that a digital oscilloscope is used to view, analyze, measure and otherwise make judgements on an analog signal. The validity of interpolation is philosophically related to the concept that statements can be made about an analog waveform utilizing only an imperfect digital representation. Since we know that depending on bandwidth, sample rate, signal fidelity etc. that we can make good assumptions about an analog waveform with the digital scope, we can also say that interpolation is a generally good method.

The two most important things to know when using interpolation are:

- How to set up the DSO for acquisition such that the assumptions needed for interpolation are as valid as possible.
- How valid the interpolation method is under the circumstances.

It cannot be overstated that these two things are actually the knowledge needed to use the DSO properly and effectively anyway and cannot be avoided.

Setting up the DSO to Enhance Validity of SinX Interpolation

All interpolation methods gain in validity as the ratio of the sample rate to the bandwidth grows. Interpolation will always improve as the sample rate is made higher. Some rules of thumb are in order. Linear interpolation works very well only when the ratio of the sample rate to the highest frequency

component is at least 10 to 1. SinX interpolation works very well only when this ratio is greater than 2:1 - 3:1 is a good ratio with 4:1 usually working almost perfectly.

Using the LeCroy WaveMaster 8620A as an example, SinX interpolation is almost perfectly valid at the highest channel sample rate of 20 GS/s. This is because the bandwidth of the scope is 6 GHz, with such a sharp dropoff in response that the signals are greatly attenuated at and above 7 GHz. Since the Nyquist rate at 20 GS/s is 10 GHz, Nyquist's criterion is met and SinX interpolation is highly effective. In effect, the bandwidth limitation of the scope ensures that interpolation is always valid at 20 GS/s. At lower sample rates, you must do some thinking if you want to use SinX interpolation - you must determine the highest frequency in your input signal. Another useful rule of thumb is the bandwidth x risetime multiplier. On the LeCroy WaveMaster DSO, the bandwidth and risetime are related by:

$$\text{Bandwidth} \times \text{Risetime} = 0.45$$

This means that a signal with greater than 90 ps risetime will have a bandwidth requirement of less than 5 GHz. While the bandwidth point is not the highest frequency content, only the point where the content drops 3 dB below the DC content, it provides some practical ability to estimate the frequency content, and therefore the sample rate required where interpolation works well. For example, a 150 ps risetime signal has a bandwidth of 3 GHz, making SinX interpolation work very well at 10 GS/s.

Determining the Validity of SinX Interpolation

The best piece of good news is that the validity of SinX interpolation can be determined. It is possible to determine this not only qualitatively but also quantitatively. Not only that, it can be determined without the slightest bit of Fourier transforms and without pen, paper or computer. All you need is the scope and a repetitive signal. The scope is easy, for the analysis will only be good on the particular scope you use. This is because so many idiosyncratic differences exist between scopes made by different vendors that the test provides scope specific results only.

The repetitive waveform may not be so easy, for the real-time DSO is usually utilized to analyze time varying waveform characteristics, but you can usually find a way. For example, if all you have is a random data pattern, try to arrange the input signal so you can trigger on a particular pattern repetitively, or try to make the system generate a repeating constant pattern. The repetitive waveform is only needed to examine the interpolation validity.

Simply set up the scope to trigger repetitively on the waveform and build a persistence map of the waveform to ensure that the waveform is repetitive. Make sure that interpolation is turned off at this step. If the waveform is repetitive, it will build a tight persistence map. Then, once the repetitiveness is determined, you need a high effective sample rate rendition of the waveform as a reference. Most high-end DSOs provide an equivalent time mode (called RIS on LeCroy scopes) and this mode can be used to acquire the reference trace. LeCroy DSOs also provide a processing function called *Persistence Trace Mean* that extracts the mean waveform from a persistence map. Save the reference waveform for later compare. If you cannot figure out how to save a reference waveform, simply print out the persistence view for comparison. Then, turn on SinX interpolation and repetitively trigger on the waveform building up a new persistence map. The test criterion is simple - the degree to which the new persistence map matches the reference waveform (or original persistence map) determines the degree of validity of the interpolation method.

Here is an example of this test applied on a WaveMaster 8620A:

This picture demonstrates how to test SinX interpolation validity with a WaveMaster scope. The repetitive signal is applied in this case to channel 2 with persistence turned on and SinX interpolation disabled. F1 is defined as the Persistence Trace Mean of channel 2. Over time, the persistence plot is filled in and the Ptrace mean function generates a clean picture of the input analog waveform. F1 is saved to memory M1 and becomes the reference trace.

M1 is left on the screen as the reference, F1 is turned off, and SinX interpolation is enabled on channel 2. The scope is triggered repeatedly, generating a new persistence plot of the repeatedly acquired waveform and its interpolated samples. Since the new persistence map is essentially an exact replica of the reference (barely visible and located in the center of the persistence plot), the test concludes that under these conditions, SinX interpolation is good to use. By the way, since the applied edge is extraordinarily fast, you can conclude that the WaveMaster 8620A will always provide proper SinX interpolation when operated at its maximum sample rate of 20 GS/s regardless of the input signal applied.