Waveform Reconstruction

(Interpolation)

Agilent Technologies

Notes to presenter: Slides 1-23 are a generic, fundamental view of interpolation. If your audience is interested in specifics about the interpolation and filtering implemented in the Agilent 54850 series scopes, slides 24 through 30 discuss this topic in more detail.

In this presentation we will consider waveform reconstruction, otherwise called interpolation, in the context of digitizing oscilloscopes. Interpolation is key to measurement accuracy in digitizing oscilloscopes, so this is an important topic to understand if you care about measurement accuracy.

Why Interpolation?

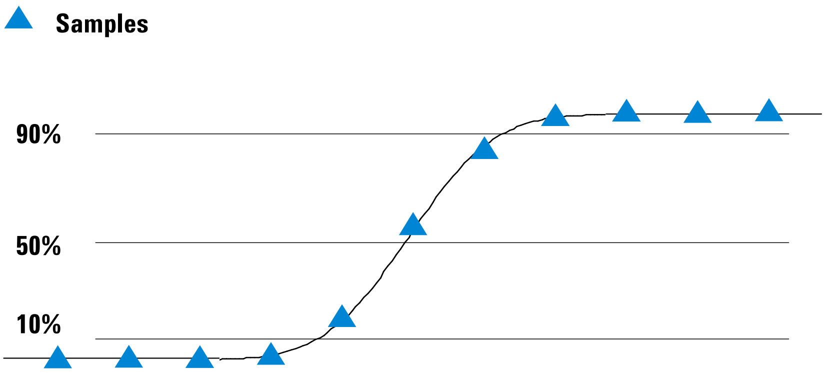

Waveform Reconstruction
December 18, 2002

Agilent Technologies

Page 2

Interpolation is necessary to achieve the greatest accuracy in measurements on sampled waveform data. In this example, none of the acquired samples coincides exactly with the 10%, 50%, or 90% points on the rising transition. Interpolating between the acquired samples provides a much more accurate measure of the corresponding time intervals than simply using the time of the nearest sample as an estimate.

What Is Interpolation?

- *Interpolation* increases the effective sampling density of a waveform consistent with the *bandwidth* allowed by the original sample rate. This process is *required* for optimal waveform *viewing* as well as for making accurate *measurements.*

- *Interpolation* cannot add *information* about the signal which was not present in the original samples.

- *Interpolation* is (for Agilent oscilloscopes) a process which follows the rules of linear systems theory, and can be observed in the time and frequency domains.

Some key points about interpolation are discussed on this slide.

Note that interpolation cannot increase the bandwidth. The bandwidth is set by the original sampling rate.

If the Shannon criterion was not violated, all the information about the signal is present in the original set of samples. Interpolation is not adding any new information about the signal.

Interpolation as practiced in Agilent oscilloscopes is a linear process, so everything you learned about linear systems theory can be applied to interpolation, in both the time and frequency domains.

The Sampling Process

When a signal is *sampled*, its frequency spectrum is *copied* at multiples of the sample rate, Fs.

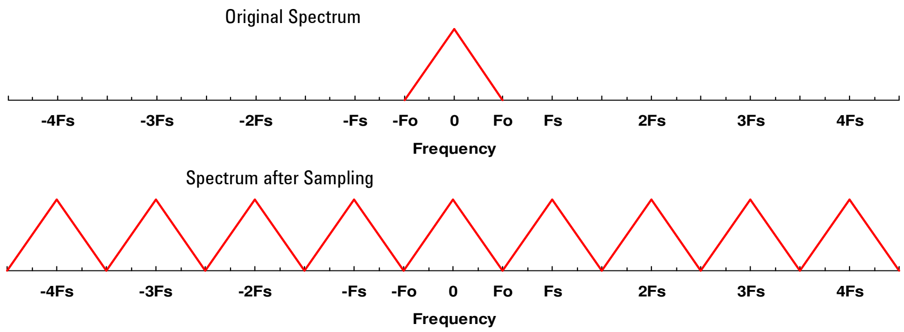

Mathematically, *any one* (or an infinite number of combinations) of these spectrums *can account for* the sampled data.

When a signal is sampled, its frequency spectrum is copied at multiples of the sample rate, Fs. For this illustration, we chose a simple triangular spectrum for the original signal. Note that in this example, the original spectrum only includes frequencies up to the foldover frequency, which is half the sampling rate. After sampling, the spectrum is copied an infinite number of times at multiples of the sampling rate.

What Is Aliasing?

Aliasing occurs when the *maximum* frequency component in the signal, the Nyquist Frequency (Fn), is greater than the Foldover Frequency (Fo = Fs/2).

When aliasing occurs, the *information bandwidth* has been *exceeded* and the spectrum is *corrupted.* Where the spectrums *overlap*, the *resulting* spectrum depends on the *phases* and *magnitudes* of the overlapping components.

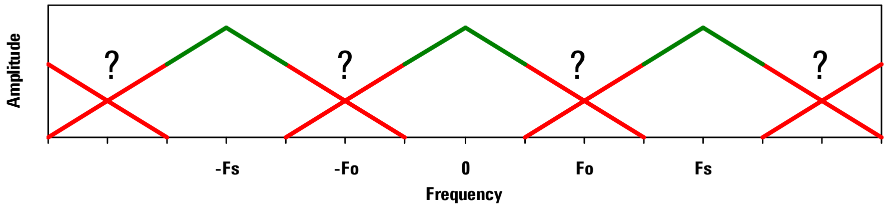

If the original spectrum had been wider, so it included components higher than the foldover frequency, the spectrum after sampling contains ambiguous components in the overlapping regions. This is commonly referred to as aliasing. The information about the original signal cannot be reliably determined by examining the sampled spectrum.

The Interpolation Process

The first step in Interpolation is to *add* zero-valued samples to achieve the desired sample *density*. For 4:1 interpolation, 3 zeros are added for each actual sample.

Note the presence of *high frequency* components (abrupt slope changes) in the zero-filled signal.

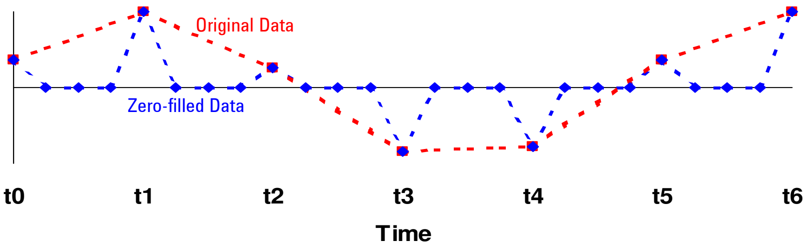

In this slide, the red squares correspond to the original samples of a signal. The first step in interpolation is to add zero-valued samples between the original samples. For 4:1 interpolation, we would add 3 zeros between every pair of actual samples. As you can intuit, this new set of samples contains some higher-frequency components, as implied by the abrupt changes in the zero-filled signal.

The Interpolation Process

Adding zeros has the effect of *rescaling* the frequency axis of the spectrum to the *new* sample rate. For 4:1 interpolation, the new sample rate is Fs' = 4Fs and the new foldover frequency is Fo' = 4Fo = 2Fs.

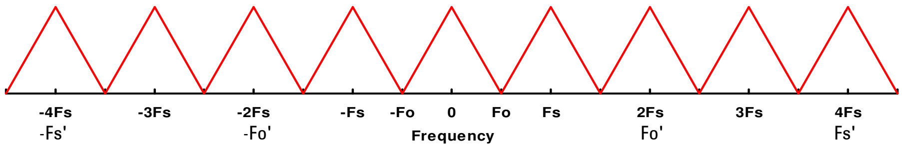

A *combination* of the replicated spectrums is *required* to create the zero-filled spectrum.

This operation effectively rescales the frequency axis of the spectrum to the new (4X higher) sampling rate. The new foldover frequency is now 2Fs.

The Interpolation Process

The *lowest-bandwidth* continuous signal which could *produce* the zero-filled signal has the following spectrum:

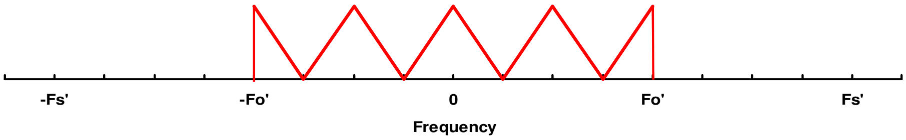

Note that there are N-1 = 3 *undesired* images of the original spectrum, and all N = 4 images have the *same* amount of *energy.* The *SNR* at this point in the process is 1/(N-1) = 1/3 = -10*log(N-1)dB = -4.8dB!

If we just stopped there, we would have 3 undesired images of the original spectrum, and a signal-to-noise ratio of -4.8 dB, in other words, more noise than signal. This should be obvious from visual examination of the zero-filled signal on page 6.

The Interpolation Process

The final step of interpolation is to *low-pass filter* the zero-filled signal to *attenuate* the *undesired* spectral images.

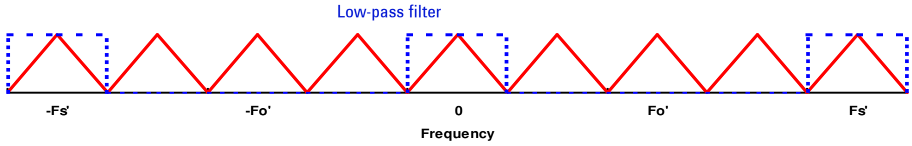

The *resulting* spectrum is the same as if the original *band-limited* signal had been sampled at *N = 4 times* the *original* sample rate.

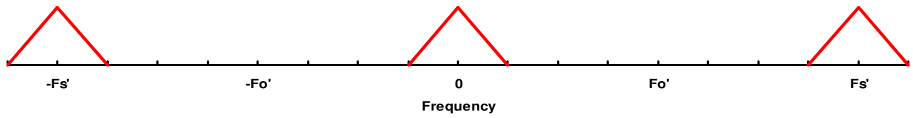

The next step is to low-pass filter the zero-filled signal to attenuate the undesired spectral images. The low-pass filter shown in the illustration would do that.

The Interpolation Process

The reconstructed waveform has *smooth* edges relative to simple *linear* interpolation.

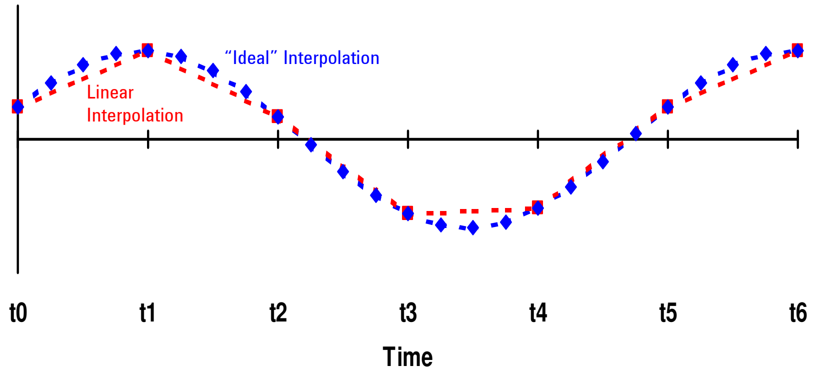

After low-pass filtering, the reconstructed waveform has smooth edges relative to simple linear interpolation (red dashed line).

The Interpolation Process

If *aliasing* occurred when the signal was *originally* sampled, then the interpolated *waveform* is still aliased.

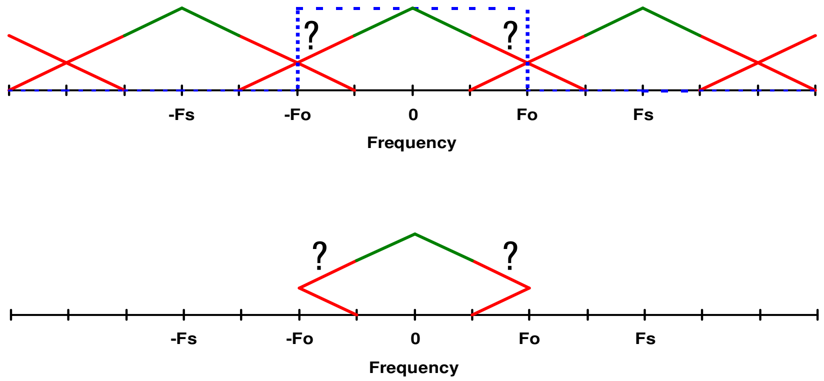

Note that if aliasing occurred when the original signal was sampled, then the interpolated waveform is still aliased. The portion in red is ambiguous as to the magnitude and phase of the original signal that it represents.

Desirable Characteristics For a Low-Pass Filter

Eight bits of resolution is equivalent to a SNR of 49.8dB. On average, the filter must *reject* the undesired spectral images by 49.8dB + 10*log(N-1)dB = *61.6dB* for 16:1 interpolation or *66.7dB* for 50:1 interpolation. (Where N is the interpolation ratio)

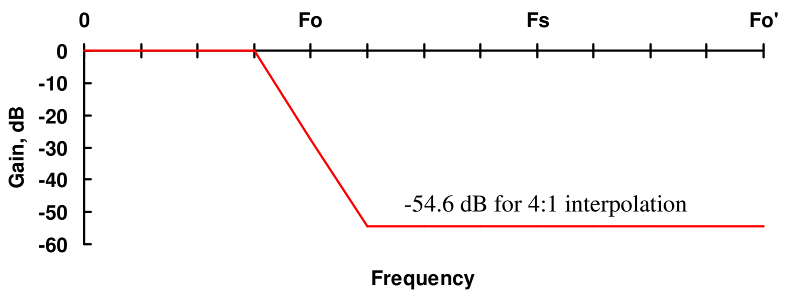

Waveform Reconstruction
December 18, 2002

Agilent Technologies

Page 12

What characteristics are desirable for the low-pass filter? Let's assume for this example that our digitizing scope has an ADC with 8-bit resolution. That's equivalent to a SNR of almost 50 dB. In order not to degrade the performance of the scope, the filter must reject the undesired spectral images by 61.6 dB for 16:1 interpolation, or 54.6 dB for 4:1 interpolation.

Desirable Characteristics For a Low-Pass Filter

Linear phase provides a *constant group delay*; all frequencies are delayed by a fixed amount of time. *Timing measurements are not distorted* due to phase effects! Linear phase is the source of the preshoot in the step response.

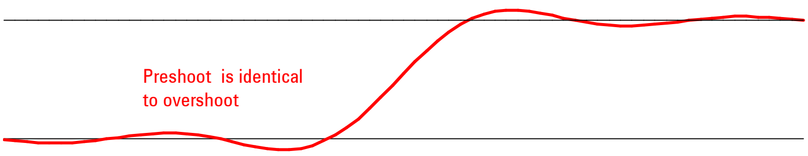

To make the most accurate time interval measurements (rise time, pulse width, etc.), it is important that all the components of the original signal be delayed by the same amount of time. This implies constant group delay. If you're using your scope to measure digital data signals that are not simple square waves, nonlinear phase will also contribute to excessive intersymbol interference.

One characteristic of a linear-phase filter with constant group delay is that the preshoot and pre-transition ripple and the overshoot and post-transition ripple are symmetrical. If you rotate the picture about the center point of the rising edge in both the vertical and horizontal axes, the two halves should overlay exactly if the group delay is constant.

Note that what we're looking at here is the response of the scope, i.e. what you would see if you stimulated the scope with a step whose rise time is much faster than the scope's rise time. For an optimally-designed filter response, if you were to use the scope to observe a band-limited signal, i.e. a step with a rise time slower than the scope's own rise time, the preshoot and overshoot would not be evident.

Desirable Characteristics For a Low-Pass Filter

The filter's response should *complement* the analog response to optimize the overall system response. A maximally flat *system* provides the most accurate measurements.

The filter response should complement the scope's analog hardware response to provide the optimal overall system response. The resulting maximally-flat system will provide the most accurate measurements.

Desirable Characteristics For a Low-Pass Filter

A low filter bandwidth provides a level of *anti-aliasing protection*. Any signal which aliases into the stop band is attenuated. In many cases, the -3dB *system* bandwidth remains the same -- the system merely rolls off faster.

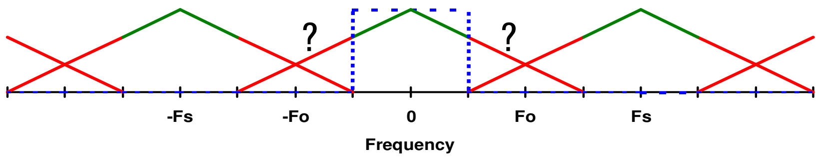

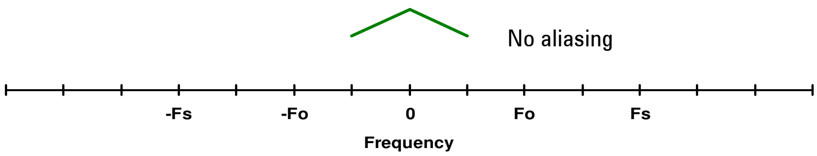

Finally, the filter should provide some level of anti-aliasing protection. This is a good time to point out that what we're really interested in is the overall response of the scope, which includes the analog hardware. As we mentioned in the last slide, a desirable characteristic of the filter we choose is to complement the scope's analog hardware. The analog hardware usually rolls off significantly before the foldover frequency is reached, so it is the primary anti-aliasing filter. For example, the Agilent 54855A scope has a bandwidth of 6 GHz and a sampling rate of 20 GS/s. Its analog response has very strong attenuation at 7 GHz.

The "Brickwall" Filter

This is the *"ideal"* reconstruction filter. It's main advantages are:

1) the original samples are not modified, and
2) the bandwidth is as high as possible, or Fo.

Its impulse response is $\sin(\pi x)/\pi x = \mathrm{sinc}(x)$.

Its impulse response is *infinite* in extent and it is, therefore, *unrealizable.* Implementations involve a complex set of trade-offs.

Now that we've established some criteria for a filter, let's examine some possible responses.

The brickwall filter would appear at first glance to be the ideal reconstruction filter, because it yields the highest possible bandwidth and it does not modify the original samples. However, it is unrealizable, as its impulse response is infinite in extent.

The "Brickwall" Filter

Impulse Response

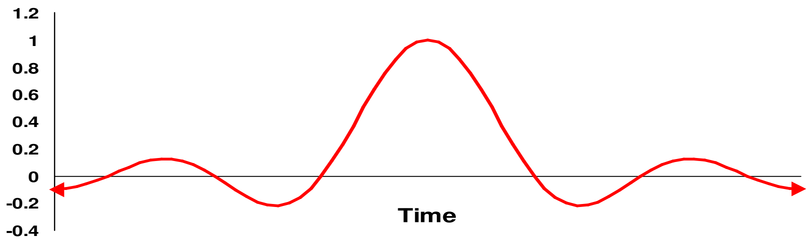

Frequency Response

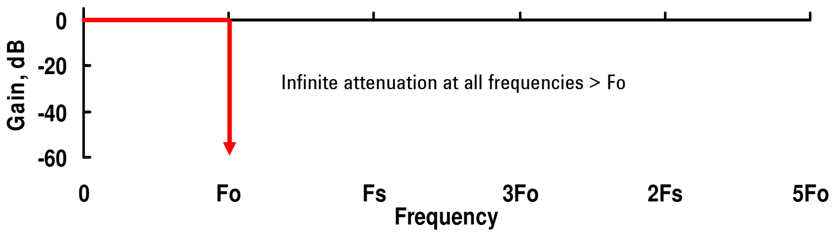

The brickwall filter has no attenuation up to the cutoff frequency (in this case we pick the foldover frequency, half the sample rate, for illustration). Above the cutoff frequency, it has infinite attenuation.

The Gaussian Filter

A filter with a *gaussian* impulse response has the *fastest* settling possible with *no overshoot.* Although ideally infinite in extent, it is simple to implement adequately.

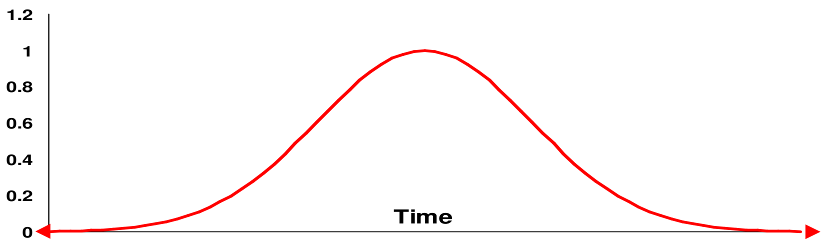

A Gaussian filter has the advantage of having no overshoot or ripple. In the “good old days” this was considered by many to be the ideal scope response. The Gaussian filter also has the advantage of being relatively simple to implement.

The Gaussian Filter

Its frequency response is also gaussian (when plotted on a linear y axis). In order to *adequately reject* undesired images, it has an *unacceptably* low bandwidth.

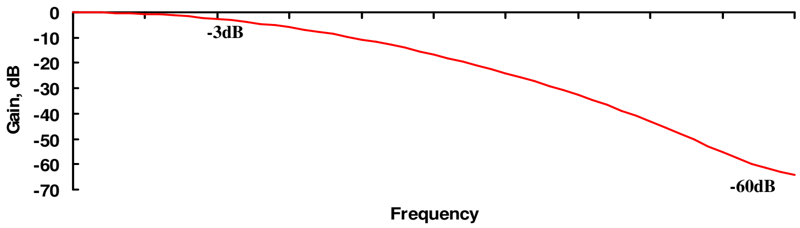

Waveform Reconstruction
December 18, 2002

Agilent Technologies

Page 19

The big problem with the Gaussian response is the tradeoff it requires between attenuation in the stopband and attenuation in the passband. The Gaussian filter rolls off very smoothly and slowly. If it has adequate attenuation above the foldover frequency to reject the spurious images, then it attenuates too much at lower frequencies. If you designed a Gaussian filter to be down 55 dB at the foldover frequency, it would have a 3-dB bandwidth that is approximately 22% of the foldover frequency. For a 20 GSa/s scope, that would make the 3-dB bandwidth only 2.2 GHz, which is unacceptably and unnecessarily low.

The Gaussian filter also causes unacceptable measurement inaccuracy. Even below the 3-dB point, it is throwing away valuable information about the signal. This was the source of the old “rule of thumb” that your scope had to have a bandwidth that was 3 times the bandwidth of your signal to make accurate measurements.

The Linear Interpolation Filter

Linear interpolation can also be described by this process. Its filter has a simple *triangular* impulse response.

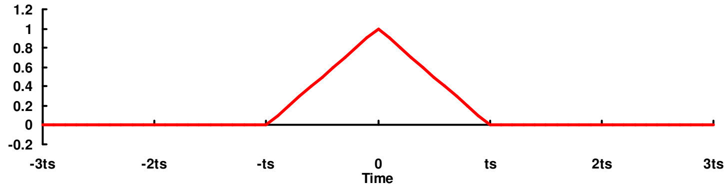

The *discontinuities* in the slope indicate that it has *high frequency* content -- it does not reject the spectral images very well.

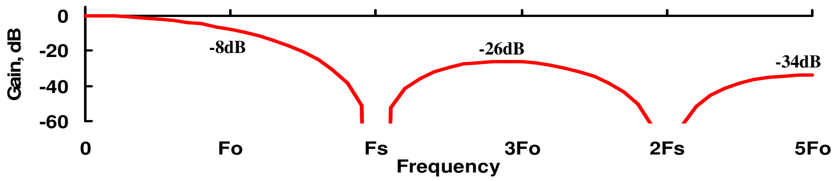

Linear interpolation is also a possible choice. As you can see in this slide, its response is not well behaved. It is not successful in rejecting spurious images, as its response keeps bouncing back up in the stopband.

Other Characteristics To Consider

- The filter should not be overly long (too many taps):
  - The longer the filter the longer it takes to run - processing time is directly proportional to filter length
  - If the filter has a length of N, you lose N-1 samples out of the record.

Filters in Agilent Oscilloscopes

A *brickwall* filter is one *extreme;* a *gaussian* filter is another *extreme.* The filters we use in our real-time scopes are somewhere in the *middle.* The specific filter varies with the oscilloscope's model number as well as the sample rate and time/div settings.

These filters are designed to *complement* the analog front end designs of our products. When combined with the products' analog front ends, they provide *exceptional* signal *fidelity* and *measurement accuracy*.

The interpolation filters have a *gain of 1* and *no phase shift* out to the system bandwidth. Therefore, the filters *do not modify* any samples of a signal which is *within* the system bandwidth.

As you can deduce, neither the brickwall filter nor the Gaussian filter would be an optimum choice. The filters implemented in Agilent's real-time scopes are a tradeoff between the two extremes. The filters are designed to to complement the analog hardware to provide the greatest overall measurement accuracy, which includes exceptionally flat response and constant group delay out to the scope's bandwidth. The filters have no phase shift out to the scope's bandwidth, therefore they do not modify any samples of a signal which is within the scope's bandwidth.

Summary of Interpolation

Interpolation is required for viewing or making measurements on high bandwidth real time signals.

Interpolation is a *mathematically linear* process which consists of adding zero-valued samples, then low-pass filtering.

The choice of a low pass filter is *critical* for overall *waveform fidelity* and *measurement accuracy*. It should be chosen to *optimize* the *measurement system*.

Once aliasing occurs, it cannot be corrected. However, limiting the filter bandwidth can reduce aliased signal components (along with some non-aliased components).

To summarize:

Interpolation is a requirement for viewing and for making accurate measurements on high-bandwidth real time signals.

Interpolation is a mathematically linear process, consisting of adding zero-value samples then low-pass filtering.

The choice of filter response is critical. The filter should be designed to optimize the overall measurement system.

Once aliasing occurs, it cannot be eliminated by filtering. However, limiting the filter bandwidth can help to reject aliased signal components – but also some non-aliased components.

Implementation Details in Agilent 54850A

- As a function of sample rate
- With (Sin X/X) interpolation turned on
- By model number
- On exported waveform data
- On automatic measurements
- On jitter measurements (EZJIT)

The type of interpolation and the type of filter used depends on a number of variables. This next set of slides describes the current implementation (rev. A.03.19 or below)

As a Function of Sampling Rate

- At 20 GSa/s, response optimizing filter is used with a system bandwidth >6 GHz (54855A), very flat response out to 6 GHz, and linear phase. This filter is always used, even if (Sin X)/X interpolation is turned off.
- At 5 GSa/s and 10 GSa/s, a response optimization filter is used with a bandwidth of Fs/4, if (Sin X)/X is turned on.
- At 2 GSa/s, no optimization filtering is done. The user can still select Sin (X)/X interpolation.

When the sampling rate is set to 20 GSa/s, the oscilloscope uses a filter characteristic described in the previous portion of this presentation, i.e., maximally flat response and linear phase out to the scope's specified bandwidth. On the 54853A, the optimization filter also operates at 10 GSa/s sampling rate.

When the sampling rate is set to 5 GSa/s or 10 GSa/s, an optimization filter with a lower cutoff frequency is used to help reduce aliasing while preserving very flat frequency response and linear phase response.

The analog hardware roll-off still provides the primary anti-aliasing filter. Using the interpolation filter for anti-aliasing only works if this is the case. For example, consider the Agilent 54855A sampling at 10 GSa/s. The analog response has good attenuation (anti-aliasing) by 7 GHz. At 10 GSa/s, the aliased band from 5 to 7 GHz folds into (aliases into) the 5 to 3 GHz band respectively. With the Fs/4 filter (2.5 GHz), aliasing is essentially eliminated, given the response of the front end. At 3 GHz, the digital filter is moderately attenuating the signal (both non-aliased 3 GHz and aliased 7 GHz components). At 4 GHz the digital filter provides full attenuation, making up for the lack of attenuation in the analog path. Between 3 and 4 GHz, the combination of the two roll-offs attenuate signals from 7 to 6 Ghz to prevent aliasing. Signals from 3 to 4 GHz are only attenuated by the digital filter and, consequently, aren't attenuated as much. From 4 to 5 GHz, the digital filter is providing all of the attenuation of aliased components.

Note that if the analog path didn't roll off, then signals from 7.5 GHz to 12.5 GHz would alias into the 0 to 2.5 GHz band and would not be eliminated.

With (Sin X)/X Selected

- With Sin X/X interpolation turned on, the amount of interpolation ranges from 4:1 to 16:1, depending on the sampling rate and time base settings.
- We use at least 4:1 interpolation on full sample rate data to improve edge measurements, eliminate visual aliasing, improve Delayed and function magnification performance
  - The user can turn off sin X/X interpolation if desired to improve throughput

The degree of interpolation varies depending on the time base setting and sampling rate selected. The tradeoff is made to enhance throughput to the screen.

By Model Number

- **On all three models, the response optimization filter is optimized for overall flat system response and linear phase out to the scope's bandwidth.**
- **On the 54853A, an optimization filter is not used if the sampling rate is set to 10 GSa/s**
- **On the 54852A, no optimization filter is used. The hardware is sufficiently flat and has linear phase out to the 2 Ghz bandwidth of the 54852A, so no response optimization filter is needed.**

On all three models, 54853A, 54854A, and 54855A, response optimization filter is used to provide optimal flatness and linear phase out to the scope's bandwidth when operating at full sampling rate.

The 54853A oscilloscope has a bandwidth of 2.5 GHz. Therefore the analog front end provides sufficient antialiasing at 10 GSa/s. Therefore the Fs/4 filter is not implemented at 10 GSa/s.

On Exported Waveform Data

Waveform data saved to a file (.wfm or .csv) for analysis by an external application such as a spreadsheet or the Agilent 89601A Vector Signal Analysis always has the enhancement filter applied if it was sampled at 20 GSa/s (also at 10 GSa/s on the 54853A).

If the scope sampling rate is set to 20 GSa/s, waveform data saved to a file for analysis by an external application always has the optimization filter applied.

Automatic Measurements

When automatic measurements are performed, linear interpolation between the nearest points is always used to enhance accuracy. This is true whether Sin X/X interpolation is on or off.

When you use an automatic measurement, the scope always uses linear interpolation between the nearest points to find, for example, the 10%, 50%, 90% point. If Sin X/X interpolation is turned on, linear interpolation is applied between the two nearest interpolated points, not limited to the original samples. This yields the optimal accuracy on measurements.

Jitter Measurements (EZJIT)

When EZJIT (option 002 or E2681A) is turned on:

•If there are >=10 samples on an edge, no additional interpolation is done.

•If there are <10 samples on an edge, then 16:1 Sin X/X interpolation is invoked.

In all cases, linear interpolation between nearest points is used to optimize accuracy.

If EZJIT is turned on to make jitter measurements, sin X/X interpolation is used to ensure at least 10 points on any edge used in a measurement. In addition to this, linear interpolation is used between the nearest points to optimize the accuracy.