Voltage and Time Resolution in Digitizing Oscilloscopes

Application Note 348

November 1986

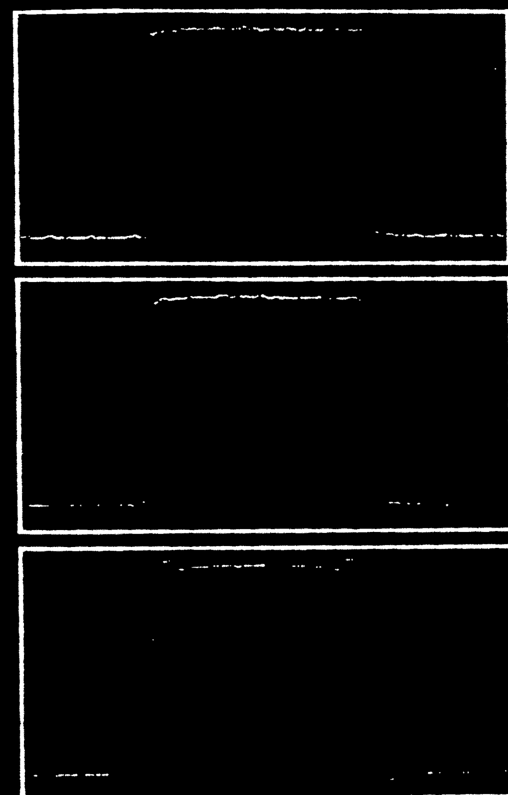

Voltage and Time Resolution in Digitizing Oscilloscopes

Application Note 348

Voltage Resolution

Introduction

Bandwidth, sensitivity and writing rate are familiar terms to most oscilloscope users. With the rapid growth of digitizing oscilloscopes, users need to know more about specifications like sampling rate, bits of resolution and differential non-linearity to determine what kind of instruments they need to make a measurement. A large part of this new vocabulary concerns voltage and time resolution, and here people often make misleading over-simplifications.

This application note explains resolution, errors in digitization, and their measurement and improvement. It discusses the relationship between real-time sampling rate and analog bandwidth, and the effects of these considerations on automatic parametric measurements. Although the oscilloscope user usually doesn't need to know these internal design considerations, occasionally they do affect measurements.

To decide whether you can make an overshoot measurement on an oscilloscope you must know its resolution. Analog oscilloscope resolution is generally specified in lines per division; digitizing oscilloscope resolution in bits. To compare the different measurements, you must understand differences in the front-ends of both kinds of oscilloscope as well as a figure of merit known as "effective bits." Figures 1a and 1b show an example of this.

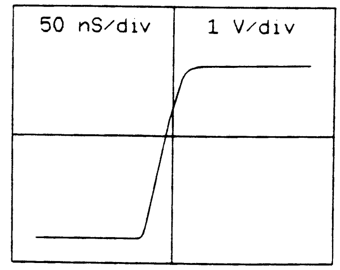

Figure 1a. *On 1V/div scale. The analog oscilloscope's resolution is 33mV. Since the signal must not overdrive the vertical amplifier, for a TTL signal this is the most sensitive scale useable. The smallest discernible overshoot is 0.6%. Uncalibration of the vertical helps but doesn't improve measurement capability.*

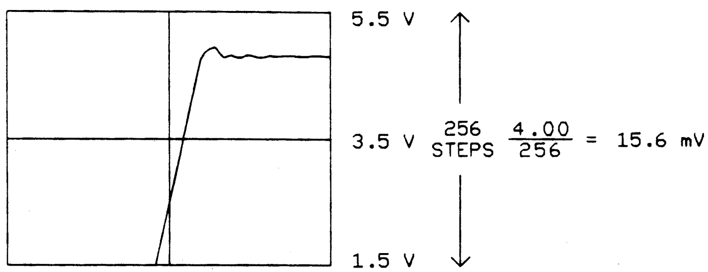

Figure 1b. *Since digitizing oscilloscopes process numerical information, the digitizing limits can be set on a much smaller range. Here 8 bits of resolution would produce a voltage resolution of 15.6mV, allowing a 0.3% overshoot measurement.*

Analog oscilloscope resolution is defined [1] as the total number of trace lines discernible along the coordinate axes, bounded by the extremities of the graticule or other specific limits. For a typical analog oscilloscope, this means that vertical resolution depends on the vertical attenuation setting. Thus, a different voltage resolution is associated with each vertical scale. A typical value is 30 lines per division, which on the 1 V/div scale represents a voltage resolution of 33 millivolts.

Now enter digitizing oscilloscopes, and a host of other digitizing instruments, whose manufacturers specify resolution in bits. The voltage related to the least significant bit out of the digitizer is the resolution of the instrument at a particular setup. How many bits of resolution does an analog oscilloscope have? Assuming 30 lines per division and a full scale signal, this is 240 lines on screen, or almost eight bits. Eight bits would be 256 lines on screen, or two raised to the eighth power.

The Digitization Process

Knowing the number of effective bits of resolution for analog or digitizing oscilloscopes does not allow you to compare their measurement capabilities directly. Equating analog oscilloscope resolution to digitizing oscilloscope resolution assumes that all aspects of a digitizing oscilloscope's vertical subsystem are functionally similar to those of the typical analog oscilloscope. With most analog oscilloscopes, you must keep the signal on screen to avoid overdriving the input amplifiers. However, a digitizing oscilloscope need not have these offset limitations, so its resolution must be considered more carefully. HP's digitizing oscilloscopes provide magnifying or scaling capabilities that allow the full range of the display to be focused on a portion of the signal. A few analog oscilloscopes have had this capability in the past, but it is being applied more frequently today in digitizing oscilloscopes. Because of its offset capability, the digitizing oscilloscope outperforms the standard analog oscilloscope for the overshoot measurement shown in figure 1.

Because the digitizing oscilloscope of figure 1b doesn't have the offset limitations of the analog oscilloscope of figure 1a, it makes a better overshoot measurement with the same number of bits of resolution because a more sensitive vertical scale can be used. However, the real comparison is not this easy. This is because a converter's resolution can be different from that implied by the number of bits it outputs. An eight-bit converter often doesn't produce eight useful or "effective" bits without sacrificing bandwidth by running at a slow sampling rate, or by passing the signal through analog or digital filters. HP's digitizing oscilloscopes have approximately the same amount of vertical resolution as analog oscilloscopes, balancing the "effective" resolution of the converters and front end offset capabilities.

Digitizing oscilloscope measurement capability is not determined by the number of A/D output bits, nor by the number of bits that can be spanned (as in the above example) by the converter. It is governed instead by the effective amount of resolution that can be applied to the waveform. This resolution can be measured in "effective bits." To understand the concept of effective bits, it's necessary to understand various problems arising in A/D conversion.

The ideal A/D converter has a transfer function that looks like a set of stairs (fig. 2a). Quantization error is the only error to which it is subject because it produces the same output code for a range of input voltages. Real converters are subject to additional distortions that are examined in this section.

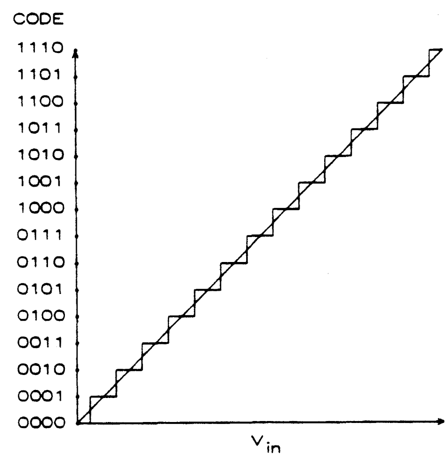

Figure 2a. *An ideal A/D converter. Note that thresholds are all equally spaced.*

Obtaining numbers that represent a time-varying voltage introduces errors. [3] These errors degrade the performance of A/D converters and limit resolution. The digitizing oscilloscope designer's problem is to control these errors to produce useful information, which involves controlling noise, non-linearities, missing codes and aperture uncertainty.

Noise

Primary noise sources are thermal processes and in the digitizing oscilloscope, the quantization process itself. The signal actually digitized may be thought of as having a random signal added to it:

$$V(t) = s(t) + N(t) \qquad (1)$$

where s(t) is the underlying or "true" signal, and N(t) represents noise. N(t) is a random process [2]. Consider for each instant in time t_i, $N(t_i)$ is a random variable. For thermal noise, which stems from the molecular motion in a conductor, each $N(t_i)$ is a rmally-distributed random variable with mean zero and a certain variance. Since N(t) is a random process, V(t) is a random process, and each $V(t_i)$ is a random variable with mean $s(t_i)$ and variance identical to that of $N(t_i)$. We make an additional requirement that the random process N(t) be stationary*, that is, for each instant in time, the associated random variable $N(t_i)$ has the same mean and variance. Notice that V(t) is not stationary, since its mean is s(t). When the real signal V(t) is digitized, the error N(t) may exceed the voltage represented by the smallest quantization step. This results in a different code than the one representing the voltage s(t).

Because the value of the nearest threshold is "unknown," another random process acts. Another random number is added to the noise-corrupted signal at every sampling instant, so that the output represents a discrete voltage associated with a particular output code. This can be thought of as a quantization noise process, Q(t). Each $Q(t_i)$ has mean zero and a uniform distribution from -q/2 to q/2. So now:

$$\begin{aligned} CODE(t) &= V(t) + Q(t) \qquad (2) \\ &= s(t) + N(t) \end{aligned}$$

This is an expression for a noisy, quantized signal, which will later be used to explain the effect of averaging and filtering.

Non-linearities and Missing Codes

Non-linearities in the A/D converter are another source of error in digitized data. Missing codes are a form of differential non-linearity. Differential non-linearity results when one or a few of the codes of the A/D converter is not the right size. Figure 2a shows an ideal A/D's transfer function, and figure 2b shows the transfer function of an A/D with a

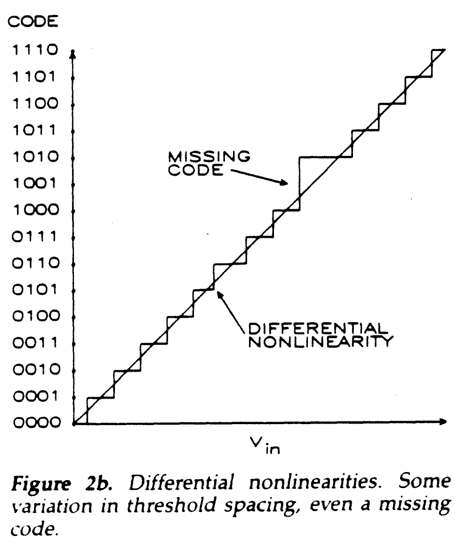

Figure 2b. *Differential nonlinearities. Some variation in threshold spacing, even a missing code.*

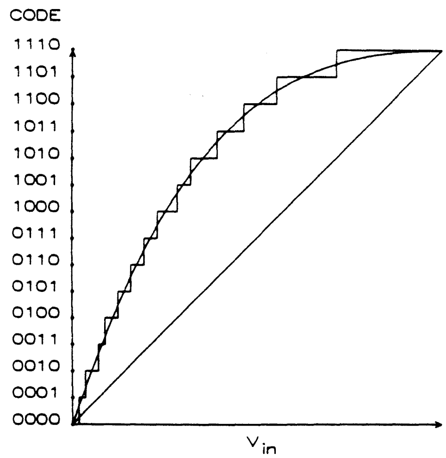

Figure 2c. *Integral nonlinearity. Small-scale signals show only gain and offset errors. Full-scale signals are severely distorted. Integral nonlinearity can also be caused by other circuit elements in the signal path.*

differential non-linearity and missing code. Notice that in the ideal A/D converter, all the codes are the same size—that is, v(n+1)-v(n) is a constant for all n, each n representing the nth threshold. The effects of missing codes or differential non-linearities on full-scale signals are usually not great. However, they may severely distort small scale signals that cross the non-linearity.

Integral non-linearity distorts the overall waveform (fig 3c). It can result from non-linear conditioning of the signal either before it reaches the A/D converter or within the converter. For small signals, distortion is minimal and is reflected in gain or offset errors. On the other hand, for full-scale signals, significant distortion occurs. Note that integral non-linear distortion is not limited to A/D converters, but that it affects overall signal fidelity.

*It must also be ergodic. For more explanation see [2].

Performance Measurement of Digitizers

Aperture Jitter

Sampling aperture is defined as the interval of time that the sampler is 'turned on.' Errors in the time placement of the sampling aperture are known as aperture jitter, and they are caused by noise in the time base. When the placement of the sampling aperture is at the wrong time, an incorrect voltage sampled is reported as having occurred at the proper time. Clearly, no amount of A/D converter resolution remedies this kind of error. Noise in the time base can result either from a noisy trigger, which causes uncertain placement of waveforms with respect to the trigger point--particularly serious when repeated acquisitions are being compared against each other--or from non-linearities in the timebase circuitry itself.

Noise, non-linearities, and missing codes cause a converter to have less resolution than indicated by the number of its output bits. If a converter has missing codes or other differential non-linearities in certain frequency ranges, then its ability to resolve signals at those frequencies is clearly less than an ideal converter with the same number of output bits. The same holds true for non-linear distortion, whether it is introduced by non-linearities in amplifiers or other devices such as charge-coupled devices (CCDs), which are often used to sample the analog signal for later digitization at a slower rate. Every A/D converter is subject to these errors. To compare between converters, several tests have been devised [3]. Effective bits, a simple and easy to understand figure of merit, is discussed in the next section.

Vertical resolution can be estimated by measuring the fidelity with which a signal is digitized, and then stating this measure understandably. This is the method of the sine wave curve fit test, or as it is sometimes called, the effective bits test. A real digitizer has N effective bits if it performs as well as an ideal N-bit digitizer at a given frequency. An ideal N-bit digitizer is subject only to quantization error.

Here "as well as" is usually taken to mean in a root-mean-square error (RMS) sense. The voltages represented by the A/D converter output codes are compared to the input voltages at the sample times. The RMS error between the output of an A/D converter V_o, and its input V_i, is defined as:

$$e_{RMS} = \sqrt{\frac{1}{N-1} \sum_{i=1}^{N} (v_0 - v_i)^2} \qquad (3)$$

All A/D converters, even ideal converters, introduce errors. The RMS error introduced by an ideal converter corresponds to its quantization error, which declines as the number of bits increases. The greater the resolution of the converter, the less its RMS error.

Relating effective bits to RMS error:

$$\text{Effective bits} = N - \log_2 \left[\frac{RMS_{MEASURED}}{RMS_{IDEAL}} \right] \qquad (4)$$

As the measured RMS error increases, the number of effective bits decreases. If an A/D converter is measured and has RMS equal to the ideal RMS for a given frequency, then it has N effective bits at that frequency. Note also that if the measured RMS is less than the ideal RMS, a converter has more than N effective bits. This is possible when the signal is processed either by a smoothing filter (for single-shot samplers) or by averaging (for repeating samplers). Assuming a small quantization step size, the quantization error density function over individual codes is uniform. Then RMS(ideal) = $Q/\sqrt{12}$ where Q is the quantization step size. A more detailed description of the test procedure for determining the effective bits can be found in [4].

The effective bits figure of merit depends on frequency. A digitizing oscilloscope with a seven-bit A/D converter may have more resolution at some frequencies than one with an eight-bit converter if it acquires signals with less RMS error.

Improving Voltage Resolution

Two techniques are particularly useful in improving the information quality from A/D converters—signal averaging for repetitive signals, and digital filtering for single-shot signals. Signals can be both filtered and averaged as well. In practice it's possible to increase resolution by two to four bits using these techniques, even to achieve better resolution than possible with an ideal converter of the same number of output bits.

Signal Averaging

Consider a signal corrupted by random noise, as discussed above:

$$V(t) = s(t) + N(t).$$

we let each $N(t_i)$ be normally distributed with mean zero and variance σ^2 then the RMS error associated with a digitization of 1000 points of this waveform is:

$$e_{RMS} = \sqrt{\frac{1}{999} \sum_{i=1}^{1000} [v(t_i)-s(t_i)]^2}$$

$$= \sqrt{\frac{1}{999} \sum_{i=1}^{1000} N^2(t_i)} \approx \sigma \qquad (5)$$

Averaging the waveform, reduces the RMS error by a factor of $1/\sqrt{A}$, where A is the number of averages:

$$\overline{N}(t_i) = \sum_{k=1}^{A} N_k(t_i)$$

$$\text{so} \quad \text{var}\ \overline{N}(t_i) = \frac{\sigma^2}{A} \qquad (6)$$

$$e_{RMS,A} = \sqrt{\frac{1}{999} \sum_{i=1}^{1000} \overline{N}^2(t_i)} \approx \frac{\sigma}{\sqrt{A}} \qquad (7)$$

This analysis assumes no quantization error. However it remains almost the same if quantization error is added. If the noise on a digitized signal has very small variance compared with the quantization noise, averaging the signal only clarifies the quantization steps. To be averaged effectively, a signal must have enough noise on it to be quantized to a step adjacent to the one representing the underlying signal for a fair fraction of the samples.

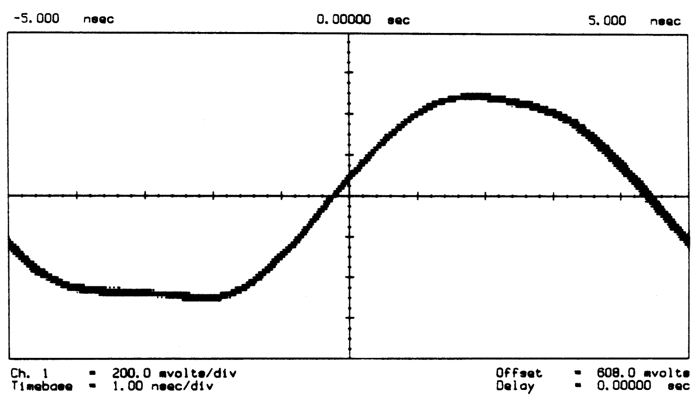

Figure 3a. Signal before averaging

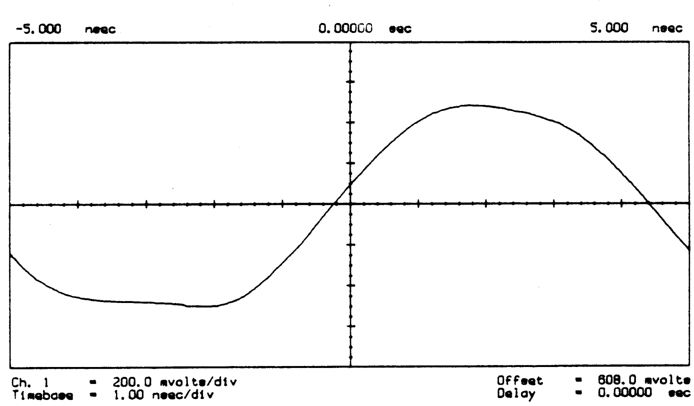

Figure 3b. Signal after averaging

To demonstrate the effect of averaging, a noisy signal was input to the HP 54100D digitizing oscilloscope. Figures 3a and 3b show the effect of averaging. Unlike bandwidth limiting, another noise-reducing technique, averaging does not distort the underlying waveform. Through averaging, it is possible to improve voltage resolution from seven bits to ten.

Digital Filtering Increases Vertical Resolution.

Digital filtering is a more sophisticated means of increasing the resolution of digitizing oscilloscopes. In essence, digital filtering trades vertical bandwidth for vertical resolution, and may be used for both single-shot and repetitive measurements. It is most effective for single-shot data.

Just as an analog filter can be thought of as the implementation of a differential equation, a digital filter is the implementation of a difference equation.[5] Processing the data using a digital filter can increase the effective bits figure of merit for a signal.

In the s-domain, a filter transfer function relates the input signal X(s) and output Y(s) by Y(s)=H(s)X(s). H(s) is normally a rational function that can be written as H(s)= N(s)/D(s). Then,

$$s^k Y(s) \Leftrightarrow \frac{d^k}{dt^k} y(t) \qquad (8)$$

Transforming into the time domain, noting that

$D(s)\ Y(s) = N(s)\ X(s)$,

$$\sum_{n=0}^{M} b_n \frac{d^n}{dt^n} y(t) = \sum_{n=0}^{L} a_n \frac{d^n}{dt^n} x(t) \qquad (9)$$

A completely analogous pair of expressions exist for sampled signals. If X(z) and Y(z) are the z-transforms of sampled sequences x_k and y_k, and H(z) is the transfer function of a digital filter, then Y(z) = H(z)X(z). Again, H(z) normally is a rational function. In this case, however, a time-difference rather than a differentiation relation applies:

$$z^i Y(z) \Leftrightarrow y_{k+i} \qquad (10)$$

and the result is:

$$\sum_{n=0}^{M} b_n y_{k-n} = \sum_{n=0}^{L} a_n x_{k-n} \qquad (11)$$

This is just a weighted average of inputs and outputs. Digitizing oscilloscopes employing filters generally use weighted averages of just the inputs; these are known as Finite Impulse Response (FIR) filters.

How does filtering improve resolution? Showing this mathematically is beyond the scope of this note, but one can see how this works as follows: a simple filter, $y_n = 0.25x_{n-1} + 0.5\ x_n + 0.25x_{n+1}$ is applied to acquired data. If the input sequence is corrupted by noise, the averaging process of the filter reduces the noise variance, just as averaging repeated occurrences did earlier. Now, however, x_n cannot be changing so rapidly anywhere that its weighted averages do not give a reasonably good representation of the original sequence. Another way to consider this is that the signal part of a sample x_n is correlated to nearby samples, but the noise part is not. The noise component is filtered out.

Filtering improves the effective bits figure of merit because it reduces the amount of noise. A sinewave quantized to four bits is filtered using the above filter (figures 4a and 4b). The output

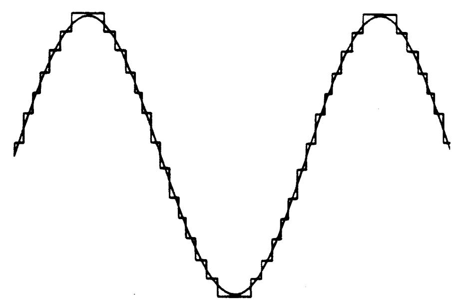

Figure 4a. *Sine wave digitized with 4 bits of resolution, no filter.*

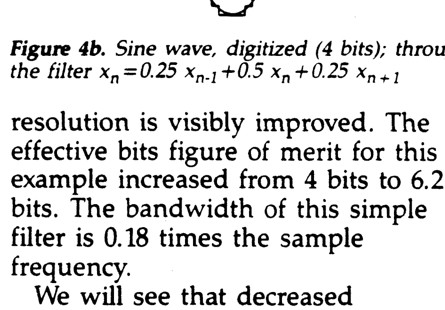

Figure 4b. *Sine wave, digitized (4 bits); through the filter* $x_n = 0.25\ x_{n-1} + 0.5\ x_n + 0.25\ x_{n+1}$

resolution is visibly improved. The effective bits figure of merit for this example increased from 4 bits to 6.2 bits. The bandwidth of this simple filter is 0.18 times the sample frequency.

We will see that decreased bandwidth is the cost of additional horizontal resolution as well. Thus, signal fidelity in the time domain and analog bandwidth are traded off at a given sampling rate.

Single-shot Time Resolution

For analog oscilloscopes, time resolution is defined as the number of lines visible per screen. The pixel resolution of a digitizing oscilloscope may also limit time resolution, but the usual limiting factor in single-shot instruments is the sampling interval. Many instruments use repeating sampling techniques such as random repetitive [7] or sequential sampling to increase time resolution. Here the emphasis will be on reconstruction filters to increase time resolution in single-shot (sometimes called "real-time") digitizing oscilloscopes.

For a digitizing oscilloscope, the best resolution obtainable without signal processing is one sample period. We can measure the time of the trigger to a much finer resolution than one sampling period (figure 5), but the resolution of single-shot measurements involving the trigger remains one sampling period.

However, the horizontal resolution can be improved using a reconstruction filter. According to Nyquist's Sampling Theorem, if a signal is sampled at a frequency twice its bandwidth, then a reconstruction algorithm exists to completely recover its value at every instant. [5]

Approximating this reconstruction formula allows the digitizing oscilloscope to reconstruct signals and achieve greater horizontal resolution. The simplest reconstruction (or interpolation) algorithm involves simply drawing a line between points and computing approximate waveform values on the line. If you do this to a sine wave (figure 6), it's possible to reconstruct the waveform to an arbitrary horizontal resolution, but now RMS error has been added. You have passed the sine wave through a reconstruction filter waveform to an arbitrarily fine horizontal resolution. One problem: this increases the RMS error associated with the signal.

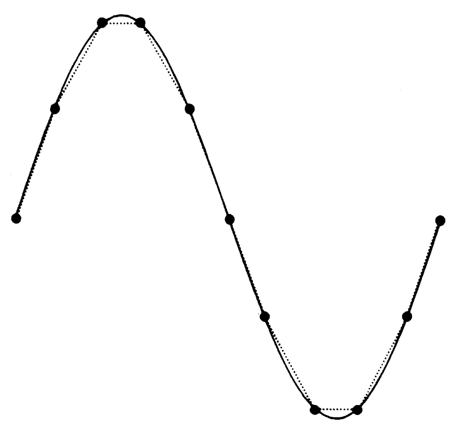

Figure 6. Linear reconstruction improves horizontal resolution, but also increases RMS error.

Reconstruction has a transfer function of its own, and may introduce distortion into signals. Conservatively applied, even linear reconstruction increases the horizontal resolution without significantly degrading signal fidelity. However, more sophisticated signal processing techniques can be used to reconstruct signals with horizontal resolution of one hundredth of a sampling period (figure 7). This can be accomplished without significantly degrading vertical resolution by conservatively limiting the analog bandwidth of input signals.

If the bandwidth of the reconstruction filter is increased, two effects result. First, single-shot waveform fidelity deteriorates (figure 8). Second, if a slow rolloff filter is used to minimize pulse distortion, more signal above $f_s/2$ passes through. This induces jitter on edges and, ultimately, aliasing.

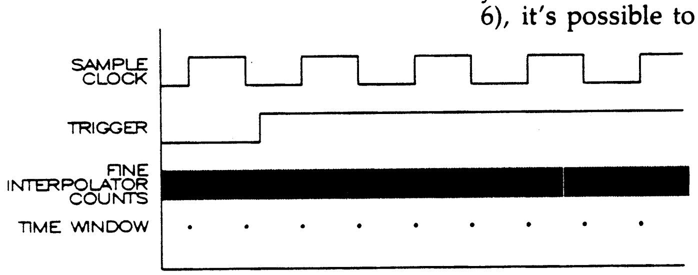

Figure 5. On repeated sweeps, time interval resolution can be improved by use of a "Fine interpolator" which measures the phase of the sample clock at the trigger instant. The first point after the trigger is displayed according to the number of fine interpolator counts from the trigger to the next sample clock. Upon repeated sweeps, it is possible to place the displayed points using the phase of the sample clock and "fill in" the intervals between points. This technique, while reducing jitter, does not necessarily improve time resolution.

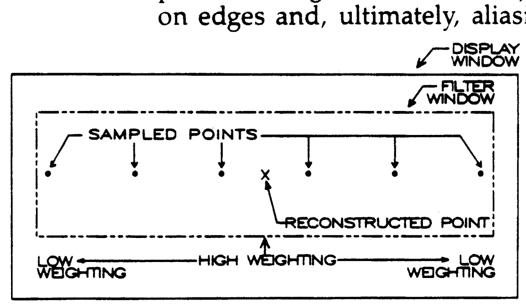

Figure 7. Depiction of how a reconstructed point is obtained. The sampled points are used in a weighted average to calculate the value of the reconstructed point. For other points, the window is moved.

Sampling Rate and Analog Bandwidth

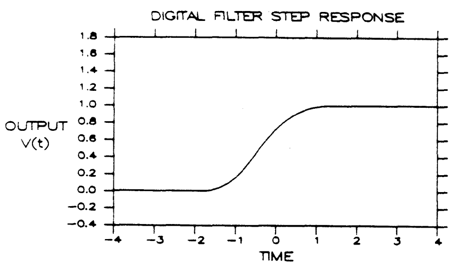

Figure 8a. Step response of digital filter with bandwidth 0.1f$_s$.

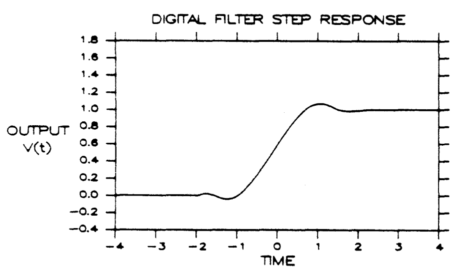

Figure 8b. Step response of lowpass digital filter with band width 0.25f$_s$.

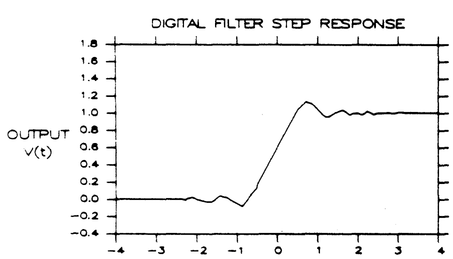

Figure 8c. Lowpass digital filter step response bandwidth 0.5f$_s$.

Oscilloscopes are designed to display time-varying signals without inordinate distortion. Analog bandwidth is defined as the frequency at which an instrument attenuates a signal by 3 dB. Real signals, such as square waves and step inputs cause problems because they are not bandwidth-limited. To minimize distortion of these kinds of signals, oscilloscope designers use a transfer function between the input and display that rolls oft slowly and preserves phase linearity. Typical designs use Gaussian or maximally-flat time delay (MFTD) filters.

To prevent aliasing, the digitizing oscilloscope must significantly attenuate signal spectral content above one-half of the sampling frequency. In digitizing oscilloscopes, several analog and digital elements control the bandwidth (figure 9). We are interested in the overall transfer function of these elements.

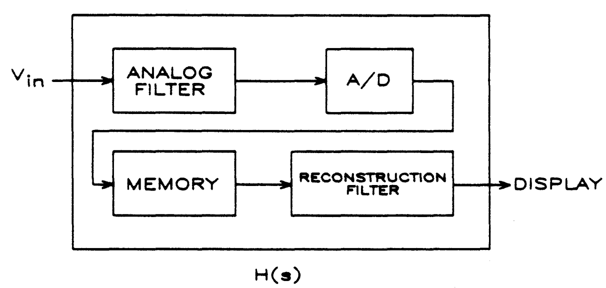

***Figure 9.** Both analog and digital circuit elements control a digitizing oscilloscope's analog bandwidth.*

The significant attenuation at half the sampling frequency requirement forces the digitizing oscilloscope designer to make another compromise. Essentially, bandwidth and signal fidelity must be traded off at a given sampling frequency. Choosing a limited bandwidth can yield exceptional pulse waveform fidelity; tolerating distortion yields higher bandwidth and better instrument rise time.

Looking at several lowpass digital filters over a range of bandwidths demonstrates this. Their step responses are shown in figure 8. To achieve good signal fidelity for a variety of common pulse waveforms, it is necessary to limit the analog bandwidth to sampling rate ratio to about 0.25. Beyond this, even though continuous-wave (CW) signals are not attenuated, pulse waveforms are unacceptably distorted.

Conclusion—Parametric Measurements are What Really Matter

Also, oscilloscope users frequently wonder about the voltage of the input signal between sample points. Many suspect that there are wide variations or glitches in the signal that are not being digitized, and in fact are being ''missed'' by the sampling process. However, digitizing and analog oscilloscopes perform similarly in this regard because of the analog front-end of digitizing oscilloscopes. If the analog bandwidth is to be limited, spectral content above this bandwidth must be eliminated by analog filtering to prevent aliasing. But the bandwidth limiting feature also spreads and attenuates high frequency glitches, making many wide enough to appear on one or two nearby sample points. Of course, if the glitch is short enough that its frequency components are extremely high when compared to the instrument's bandwidth, the sampler detects nothing. This is identical to the analog oscilloscope's response to wildly out-of-band inputs.

Voltage Resolution

To measure overshoot, you must know the limitations of the input circuitry of your oscilloscope. Digitizing oscilloscopes have either magnification capability or the ability to set gain and offset to allow a small vertical range of an input signal to be digitized. In this way, equal or superior resolution to analog oscilloscopes can be achieved on most measurements. Digitizing oscilloscopes process data numerically using averaging or filtering to improve resolution. With improved vertical resolution, it's possible to place voltage markers (percentages of the top and base of a signal) on a signal. Precise voltage determination means both precise vertical and horizontal measurements.

Time Resolution

Most pulse parameters are time-interval measurements. The measurement of a single-shot pulse width can have high variance if a signal is undersampled, or reconstructed using a filter designed to maximize CW bandwidth. Figure 9 shows an undersampled pulse. If no reconstruction is done, the pulse width may vary from measurement to measurement by as much as a full sample period. Even with linear interpolation, this situation may be as bad. Undersampling introduces edge jitter, and pulse width is uncertain. However, by limiting system bandwidth to approximately one fourth the sampling rate, it is possible to build reconstruction filters with minimal edge jitter. With these filters, precise time-interval measurements may be made on such narrow pulses.

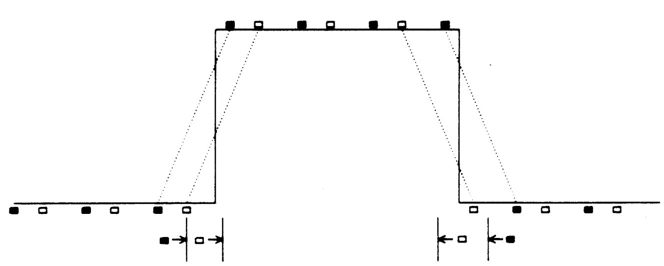

Figure 10. *Pulse width measurement varies on repetition by as much as 25% for this short undersampled pulse. Reconstruction filtering minimizes the measurement variation.*

References

[1] **Bandwidth and Sampling Rate in Digitizing Oscilloscopes,** HP Application Note 344, April 1986

[2] **Dynamic Performance Testing of A to D Converters,** Hewlett-Packard Product Note 5180A-2. HP part number 5952-7629.

[3] Jay, Frank, ed., **IEEE Standard Dictionary of Electrical and Electronics Terms,** 3rd Ed., IEEE, Inc., New York, NY. 1984.

[4] Liu, B., and Peled, A., **Digital Signal Processing,** John Wiley & Sons, New York, 1976.

[5] Muto, A. S., Neil, J., and Peetz, B. E., **Measuring Waveform Recorder Performance,** Hewlett-Packard Journal, Nov. 1982. P21.

[6] Oppenheim, A. V. and Schaeffer, R. W., **Digital Signal Processing,** Prentice-Hall, New York. 1975.

[7] Papoulis, A., **Probability, Random Variables, and Stochastic Processes,** McGraw Hill, New York, NY. 1965.

Printed in U.S.A.
10/86

5954-2652

HEWLETT
PACKARD