
Feeling Comfortable With Digitizing Oscilloscopes



© Copyright Hewlett Packard Company 1987

Manual Part Number 9320-5776

Printed in U.S.A. March 1987

Printing History

First Edition - October 1986
Second Edition - March 1987

Printed in U.S.A.

Table of Contents

Section 1: 1-1 Digitizing or Analog?

Section 2: 2-1 Analog Scopes: How Do They Work?
 2-1 What's Inside?
 2-2 The Vertical Channels
 2-3 Triggering
 2-5 The Horizontal Section
 2-6 The CRT
 2-7 Summary

Section 3: 3-1 Why a Digitizing Scope?

Section 4: 4-1 Digitizing Scopes: How Do They Work?
 4-1 Why The Long Wait?
 4-2 Some Digitizing Concepts
 4-3 What Shape Is It?
 4-5 How Do We Sample?
 4-5 Real Time Sampling
 4-6 Repetitive Sampling
 4-6 Random Repetitive Sampling
 4-7 Sequential Sampling
 4-8 The Vertical Channels
 4-9 Converting to Digital
 4-11 Memory
 4-11 The Processor
 4-12 The CRT
 4-13 Summary

Table of Contents Con't

Section 5:

- 5-1 Specifications and Measurements
- 5-1 Vertical Resolution
- 5-3 Effective Bits of Resolution
- 5-4 Trace Noise
- 5-6 Time Interval Measurements
- 5-7 Jitter
- 5-9 Multi-Channel Display
- 5-12 Storing Images on Screen
- 5-13 Hardcopy Output
- 5-14 Summary

1

Digitizing or Analog?

You've finally captured an intermittent glitch on your analog storage scope and you're trying to get a picture of it with your scope camera. Unfortunately, before you can get a good picture, the trace fades away. Or, you're trying to determine the rise time of an edge with your scope, so you expand the trace to five vertical divisions, center the trace with the vertical position control, adjust the vertical expansion, readjust the position, and then count divisions. Unfortunately, the edge doesn't have enough horizontal expansion and only covers two time divisions, so you change the time/division knob and reposition the waveform.

While the analog scope has been a mainstay in the electronics industry, it has some shortcomings that we have all accepted because there wasn't anything better. Until recently, that is. The analog scope is being challenged by one of its offspring, the digitizing scope. Because of its newness and different means of displaying and capturing data, there is often some confusion about how digitizing scopes work and what they are good for. Learning about a new type of scope also involves a time investment, which may not be easy to justify without a view of the advantages.

The goal of this book is to point out and put into perspective the advantages of digitizing storage scopes. We're not going to tell you that a digitizing scope is better for every application; both analog and digitizing scopes have their shortcomings. We would, however, like to help sort out some of the misconceptions and confusion about where digitizing and analog scopes fit, while giving some basic information about how each works. We will also discuss some digitizing scope specifications to look for, as well as what they mean. All in all we hope this will help you to see that the advantages of digitizing scopes make the transition from analog well worthwhile.

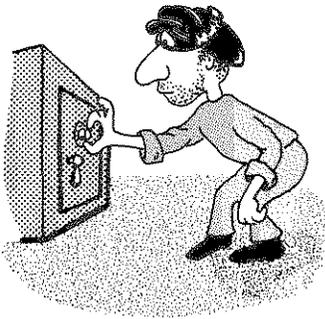


2

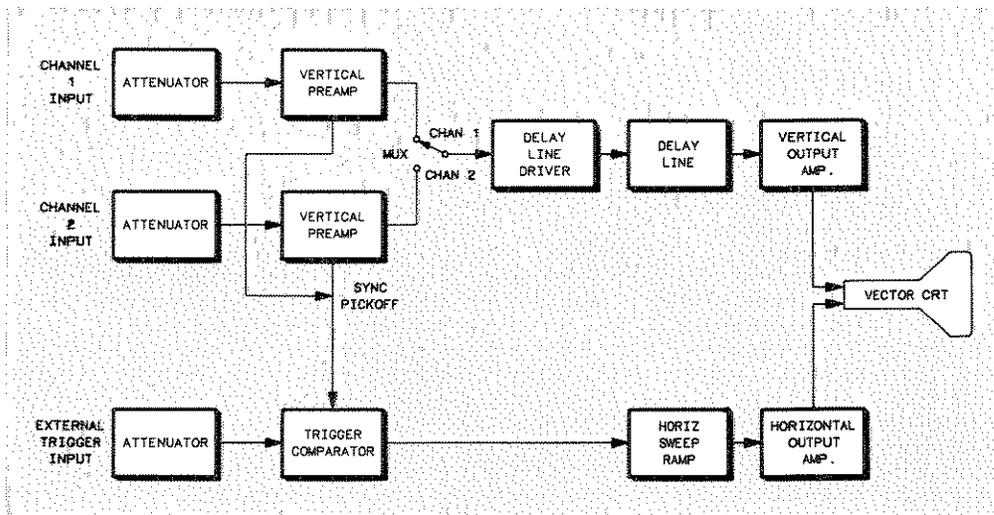
Analog Scopes: How Do They Work?

Before we talk about the advantages of analog scopes, let's review how they work for a few moments.

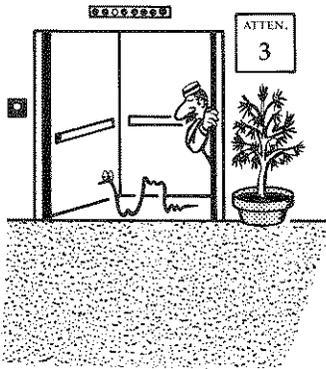
What's Inside?



As you can see in the diagram below, an analog scope has two major signal paths. The first is the vertical signal path, which ultimately is responsible for deflecting the CRT beam vertically in response to the input signal. The second path is the horizontal. It triggers the scope and moves the beam from left to right across the screen. In a typical display, time is represented horizontally and voltage is represented by the vertical axis.



The Vertical Channels

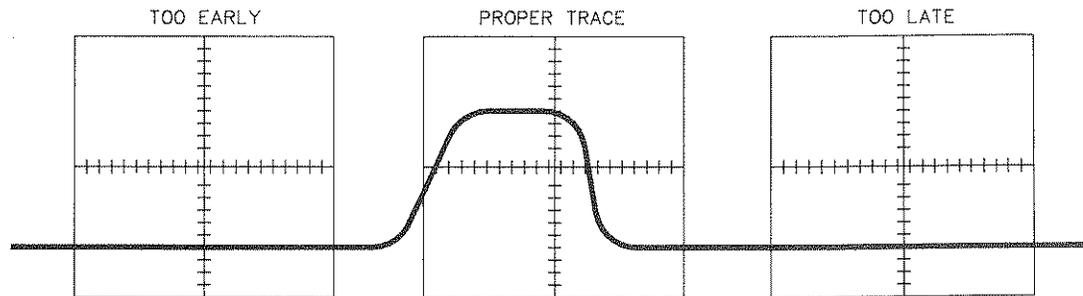


When a signal comes into an analog scope, the first thing it sees are the attenuators. Attenuators match the high impedance of the scope probes (typically $1\text{ M}\Omega$ or $10\text{ M}\Omega$) to the low impedance of the vertical preamplifiers. The attenuators also scale the input signals to a level the vertical preamps can handle. The amount of attenuation and preamp gain is set by the front panel vertical sensitivity knob.



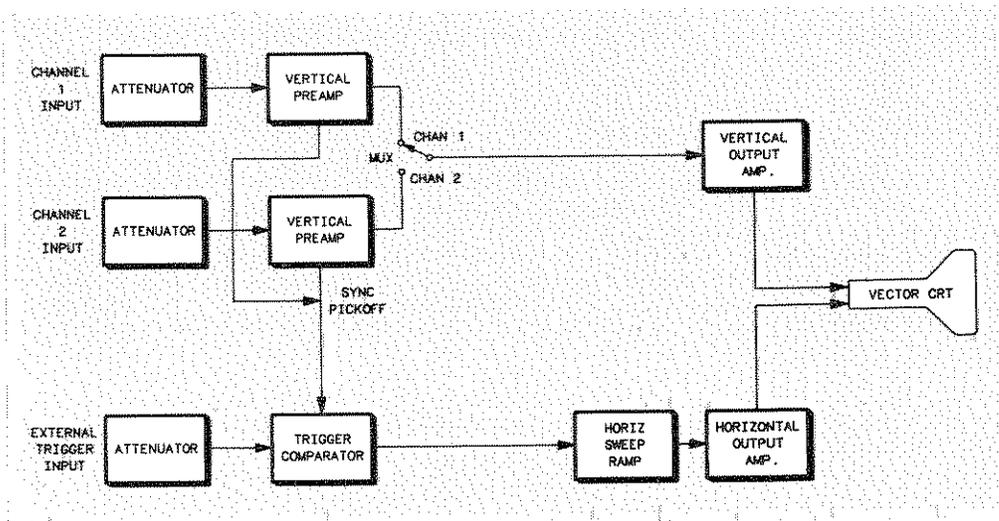
Triggering

The triggering portion plays a very important part in the operation of a scope--it determines where (in time) the trace starts. In essence, the triggering circuits tell the horizontal section when to start moving the beam from the left side of the CRT to the right. If the trace starts too early, the part of interest on the signal won't be seen. The same is true if it starts too late. The figure below gives you an idea of what happens.



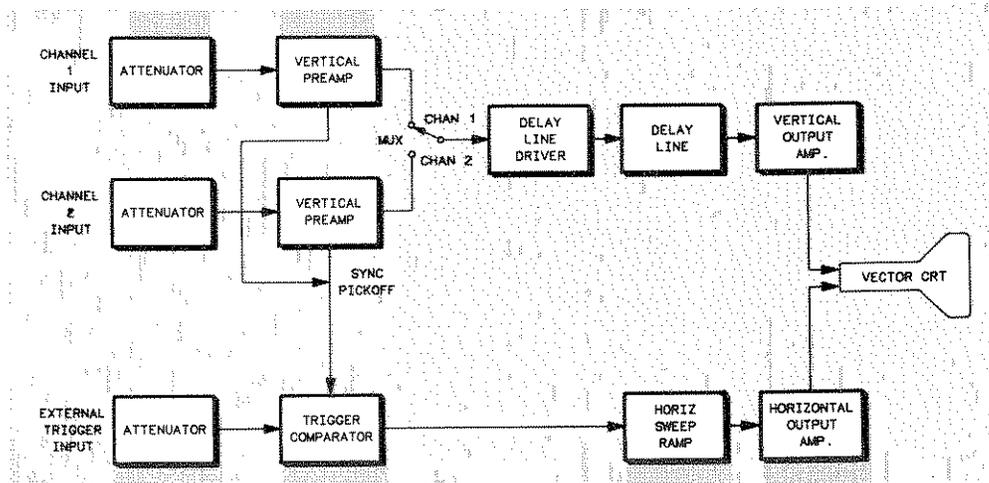
How does the trigger circuit know when to trigger? It gets a replica of the signal, called the sync pickoff, from the selected trigger source. This sync pickoff is compared to a pre-set trigger voltage that is set with the front panel trigger level knob. Most analog scopes let you specify a slope as well as a trigger voltage. This allows you to trigger at a specific point on a rising or falling transition.

When the trigger circuit finds a voltage and transition from the source that matches those set with the trigger controls, it tells the horizontal sweep circuits to start moving the beam from left to right. The speed of the beam is determined by the seconds/division knob on the front panel. As the beam is moved horizontally across the screen, the vertical amplifiers move the beam up and down, relative to the input voltage.



Both the horizontal sweep and vertical deflection information have to arrive at the CRT at the same time. If they don't, the scope won't be able to display the voltage information properly.

Look at the block diagram at the top of the page. Since the delays in the horizontal path are longer, vertical information will reach the CRT before the horizontal information. The solution to the problem is to put a calibrated delay into the vertical path so both horizontal and vertical signals will get to the CRT at the same time.



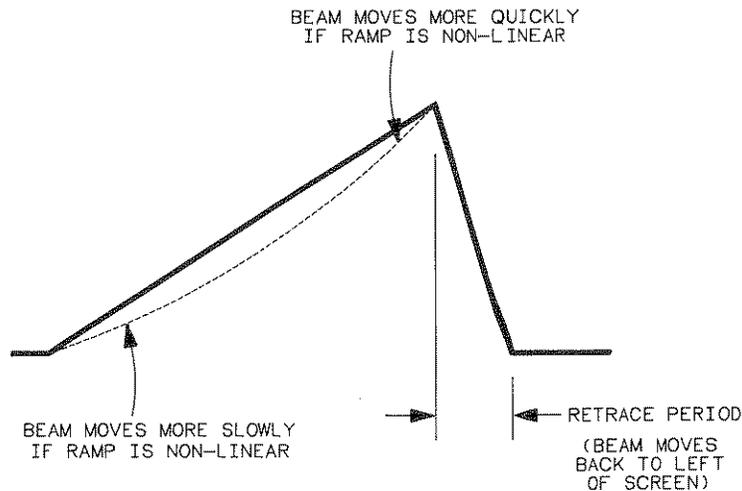
Analog Scopes: How Do They Work?

The Horizontal Section

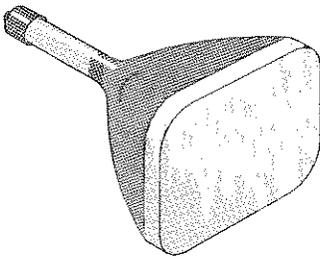
For external triggering, the horizontal path of an analog scope has an attenuator like the vertical channels. This attenuator serves the same purpose as those in the vertical channels, i.e., impedance matching and scaling the external trigger signal. However, the horizontal attenuator is followed by trigger comparison circuits, instead of a preamp, as in the vertical channels.

The horizontal portion of the scope, which is responsible for moving the trace along the time or horizontal axis, directly affects the time accuracy of an analog scope. The horizontal beam movement is controlled by a voltage ramp (called the sweep ramp); the time interval accuracy of the scope depends heavily on this ramp.

Once the trigger comparator has found a valid trigger, it tells the horizontal sweep ramp generator to start. As the ramp rises, it causes the beam to move from left to right across the CRT. Since the left to right movement represents time on the CRT, the ramp must be very linear. If the ramp has non-linearities, the beam moves at different rates across the screen. Typically, ramp linearity controls time interval accuracy of an analog scope within $\pm 3\%$.



The CRT



The last major portion of an analog scope is the display or CRT. Analog CRTs are vector displays that can move the beam to any point directly. A signal from the vertical amplifier moves the beam in the vertical direction. This may seem obvious, but it brings up a very important point. The CRT and its drivers must be able to deflect the beam vertically as fast as the signal rises. What this means is that the CRT bandwidth must be the same as the input bandwidth of the scope! High bandwidth CRTs pose several problems. As CRT bandwidth goes up, the following happens:

Cost of the CRT goes up;

Accuracy of the CRT goes down;

Reliability of the CRT goes down.

To keep the cost of the CRT down while keeping the accuracy and reliability up, the scope must use as low a bandwidth CRT as possible. However, since the CRT must have the same bandwidth as the scope, high bandwidth analog scopes demand high bandwidth CRTs. The only real solution is to move to a new architecture.

Summary

In this first section we have talked about how an analog scope works. We also pointed out some of the shortcomings of current analog scope architecture. Here are some of the key points to remember about analog scopes:

- There are two major signal paths--horizontal and vertical;
- Everything (including the CRT) must work at the same speed as the input signal;
- All input channels are usually multiplexed through a single vertical path to the CRT;
- The horizontal path is responsible for triggering;
- The scope triggers on a voltage level and rising or falling slope;
- As input bandwidth goes up, cost of the CRT also goes up, while reliability and accuracy of the CRT go down.

In the next section we will look at digital architectures that eliminate the need for all parts of the scope to work at the same speed as the input signal. These digital architectures give some measurement capabilities not available with analog scopes.



3

Why a Digitizing Scope?

The digitizing storage oscilloscope is a natural evolution of the analog storage scope. Advances in technology have allowed analog scopes to offer some impressive features over the years. However, because of the analog scope's architecture, it will be difficult or impossible to implement many new features that users have asked for. One such feature is the ability to capture data and send it to a computer in an automated test system. Another is the capability to store a captured waveform on screen indefinitely.

On the other hand, these features are easily implemented on a digitizing scope. In the next two sections we'll investigate how a digitizing scope works. We'll also see how this differs from analog scopes. As we talk about these subjects, we'll see how a digitizing architecture allows a whole new set of additional measurements that previously were either time-consuming or simply couldn't be done.

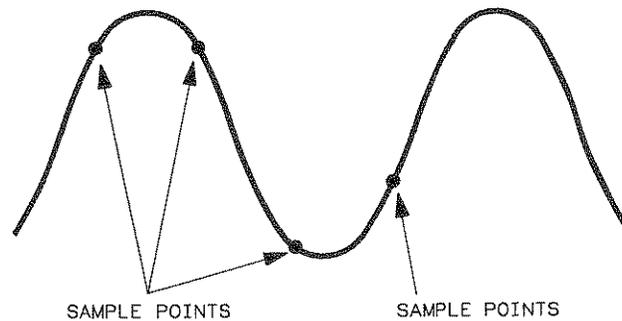
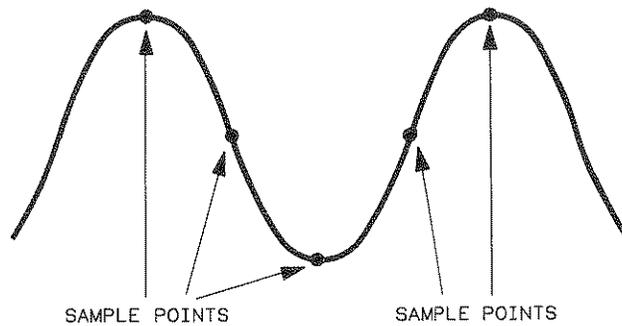


Some Digitizing Concepts

Before we talk about exactly how a digitizing scope works, let's look at some of the concepts behind one. There are three issues at hand when capturing an unknown signal: frequency, phase, and fidelity (for our purposes, fidelity includes shape and amplitude of the signal). In the realm of frequency we get some help from Nyquist. In his original article dated 1924, Nyquist formulated a theorem that says if we sample a signal at a rate of $2F$, the best information we can get is of frequency F . Said another way, in order to discern a frequency F from a signal, we must sample at a rate of at least $2F$.

Does this mean that a digitizing scope must sample at twice the bandwidth frequency? Well, no. That's really a misinterpretation of Nyquist's Theorem. One of the boundaries of Nyquist's Theorem assumes that you sample a signal infinitely in both directions if you are to completely reconstruct that signal from two points per period (with single shot sampling). Since digitizing scopes can't take an infinite number of samples in a single shot system and still display the waveform, manufacturers must limit the size of the sample window to something more reasonable. However, this also limits the accuracy of the reconstruction. The other alternative is to sample more than two times per period.

Since technology limits how fast we can sample, to get more samples per period we must lower the bandwidth limit. In other words, if we sample only twice per period for a limited amount of time, we limit the accuracy of the reconstructed signal. By sampling at a higher rate per period, we can make a more accurate reconstruction. Hewlett-Packard bandwidth limits their single-shot digitizing scopes so there are four points per period at the bandwidth.



What Shape Is It?

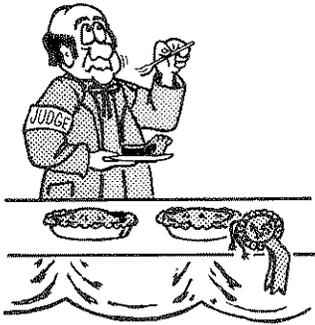


Suppose you have a 100 MHz analog scope and you put a 100 MHz square wave into it. What would you expect to see on screen? Well, unless your scope is broken you'd see a distorted 100 MHz sine wave! Why? Because a 100 MHz square wave is made up of a 100 MHz sine wave and an ideally infinite number of odd harmonics. At 100 MHz, the input signal is attenuated to -3dB and the higher the frequency, the greater the attenuation. At 100 MHz the scope has bandwidth limited most of the higher harmonics. What you have left are the fundamental and lower harmonics.

With many points per period, it is easy to reconstruct an input waveform from the digitized points. Such is the case at frequencies well below the bandwidth of the scope. As the bandwidth of the scope is approached, however, the number of points available for reconstruction drops (as we mentioned earlier, HP single-shot digitizing scopes have four points per period at bandwidth). Four points per period is adequate to characterize a sine wave. But what if the input signal is a square wave? Can a digitizing scope adequately reconstruct a square wave from four points?

To explain the point, let's use the previous example of a 100 MHz square wave. If we put a 100 MHz square wave into a 100 MHz digitizing scope, the scope will pass the square wave through a 100 MHz bandwidth limit filter. The output of this filter, which is a slightly distorted 100 MHz sine wave, will then go to the A/D. That means that the A/D is not digitizing a square wave but a sine wave with some small amount of harmonic information (the same information that the 100 MHz analog scope displays on the CRT). Since we can reconstruct a sine wave well from four points, this isn't too much of a problem. In other words, at or near the bandwidth frequency, the A/D will be digitizing sine waves (or nearly so) and reconstructing them as such. That doesn't mean that a digitizing scope sees everything as a sine wave. At frequencies less than the bandwidth we will have more sample points per period and can make a better determination of the actual shape—just as the analog oscilloscope will display a proper waveshape if the input frequency is considerably less than the specified -3dB bandwidth.

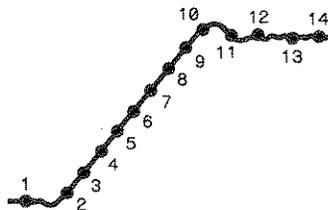
How Do We Sample?



There are two basic methods we can use to acquire these points on a waveform. We can capture them as they occur in real time by sampling at a very high rate, or by making several passes on the waveform and getting some points on each pass. We call the first method real time or single-shot sampling, as illustrated below.

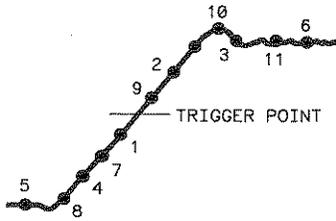
Real-Time Sampling

In this case, we sequentially sample the waveform on its first occurrence. The higher the sample rate in comparison to the frequency of the signal, the better our signal picture. Sometimes, additional points are interpolated mathematically between the actual data points. Real-time sampling allows you to capture and display events that occurred before the trigger (what we call negative time viewing).



Repetitive Sampling

The second method, called repetitive sampling, captures data on the waveform by acquiring points on more than one occurrence of the signal. This, of course, requires that the waveform itself is repetitive and not a single-shot event. On each occurrence of the trigger event a few more points are acquired until these points are put together into a very accurate composite reconstruction. Since the data is not all captured in real time on one pass, Nyquist criteria does not apply to repetitive sampling. As such, there is not necessarily any relationship between sample rate in a repetitive sampling scope and the bandwidth. As a matter of fact, the sample rate will usually be lower than the bandwidth. And, since many points are acquired on the waveform, simply connecting the data points with lines provides a very good picture of the waveform.



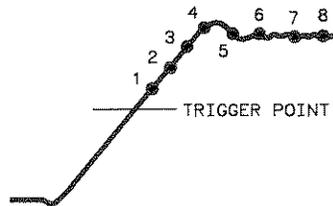
Random Repetitive Sampling

One type of repetitive sampling is random repetitive. The sampling is done constantly, not waiting for a trigger event. On each occurrence of the signal, more points are acquired. Then, all the sampled points are put together into one composite picture of the waveform. Each point is put into its proper place by measuring the amount of time that elapsed between it and the trigger point. Since all points are acquired asynchronous to the trigger point, the sampling is random in relation to the trigger point.

Because the sampled data is captured before and after the trigger point, we can actually see what happened before the trigger. We call this capability "negative time."

Sequential Sampling

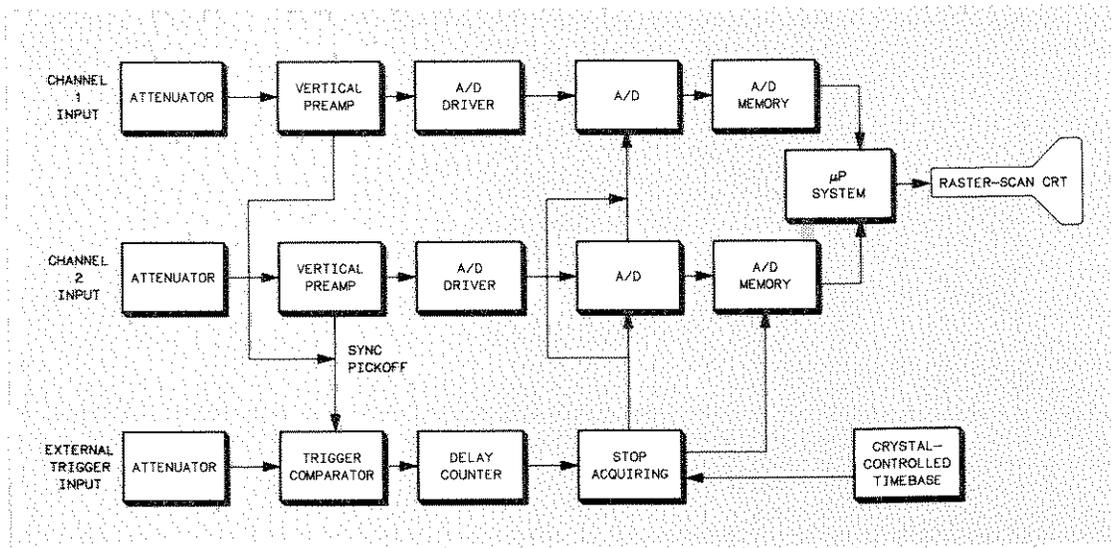
A second method of repetitive sampling looks for the trigger point on each pass, waits a predetermined amount of time, then takes a sample. On the next pass, the predetermined time interval is incremented and a new sample is taken. After a number of passes, the waveform is reconstructed just as in the random repetitive method. However, there is one major difference—with sequential sampling all the samples are taken after the trigger. This means that you can't have negative time



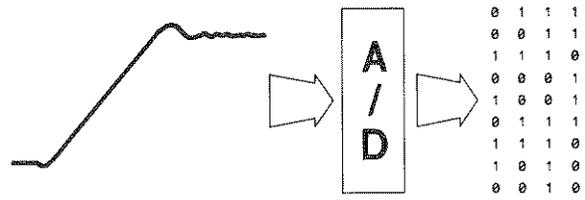
(i.e., data before the trigger) viewing. Sequential sampling does provide very accurate waveform reconstruction since it can use a slower, higher resolution A/D.

The Vertical Channels

If you look at the figure below, you'll notice that this block diagram of a digitizing scope has separate paths for each input channel. This is in contrast to most analog scopes that have only one vertical path to the CRT. By having separate paths for each channel, a digitizing scope can capture data on all channels simultaneously.



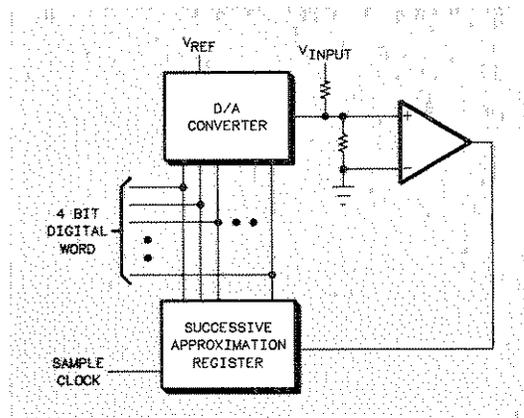
Like an analog scope, a digitizing scope has attenuators to scale and match impedances to the preamp. They also have vertical preamps to amplify the input signals to the A/Ds. The part of the block diagram following the vertical preamp is where a digitizing scope takes on a different look than its analog predecessor. Instead of amplifying the input signal and using it to drive the vertical deflection plates of the CRT, a digitizing scope changes the input signal into a digital word through the A/D convertor.



Converting to Digital

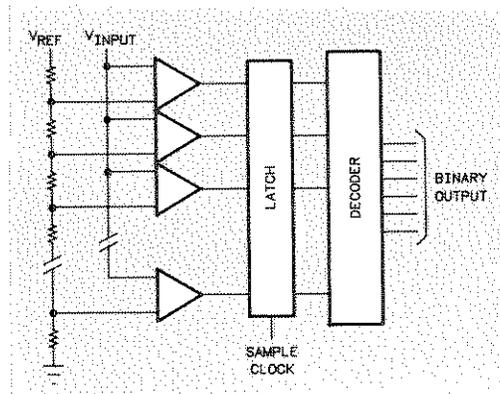
There are several ways to convert an analog input to a digital word in a scope. Digitizing scopes typically use either successive approximation converters or flash A/D converters.

Successive approximation converters are relatively simple from a hardware standpoint, since they require only one voltage comparator, as you can see below. This is essentially a serial step-by-step process.

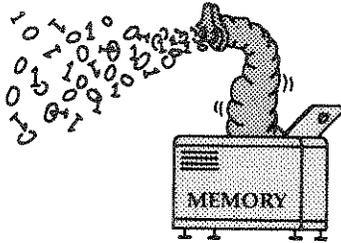


It takes N passes (where N is the number of bits of resolution) and N clocks for a successive approximation converter to convert an analog voltage to a digital output. While this scheme provides high resolution, it is not fast enough for use in a single-shot digitizing scope. Since a repetitive digitizing scope doesn't have to sample as quickly as a single-shot one, the successive approximation A/D may be used in repetitive architectures.

Single-shot digitizing scopes need a fast conversion A/D, since they digitize input voltages "on-the-fly." In such cases the flash A/D converter, similar to the one shown below, is used. A flash A/D converter is essentially a parallel process. Although fast, flash A/D converters are more complex hardware-wise. Its architecture requires $2^N - 1$ voltage comparators, again where N is the number of bits of resolution. To put that in perspective, a 6-bit flash A/D requires 63 comparators, a 10-bit needs 1023, while a 12-bit has to have 4095 comparators!



Memory

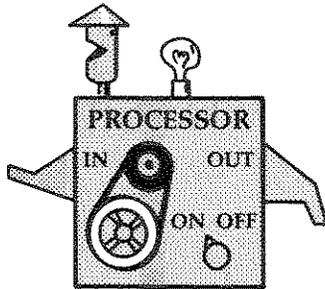


Once the data has been put in digital form, it is stored in memory. The memory must be able to store at the same rate as the A/D samples for most digital architectures. That means that if the sample rate is 200 MHz, the memory must have a write cycle time of 5 ns. Most digitizing scopes use special FISO (Fast In, Slow Out) memories for storage.

To allow the use of slower digitizers and storage memories, some manufacturers store the incoming waveform in CCDs (Charge Coupled Devices). Since the waveform can be read out of the CCD to the A/D at a slower rate, slow, high resolution A/Ds can be used. Shortcomings of this technique include noise, cell size and cell leakage, which limit horizontal and vertical resolution.

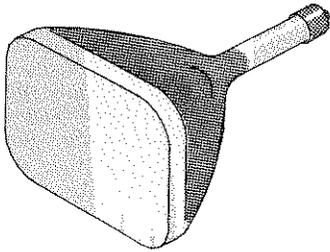
Once the digitized waveform is stored in semiconductor memory, the image can be held indefinitely without deteriorating.

The Processor



All digitizing scopes incorporate a microprocessor of some sort. The power of a processor in a digitizing scope profoundly affects its capabilities. Because the data is in digital form, it's easy to have the processor perform some specific tasks on the data. For instance, we can automatically measure such parameters as time intervals, rise and fall times, frequency and more. The processor can also allow such functions as waveform math. And surely not least of all, it's a simple matter to format it for output to a printer or plotter. This eliminates the need for scope cameras.

The CRT



The last fundamental difference between digitizing and analog scopes is the CRT. In Section 1 we talked about CRTs in analog scopes needing to be as fast as the input bandwidth. A digitizing scope, however, divorces the CRT from the input bandwidth by storing everything in memory. Data is stored in memory at the same rate it is digitized, but is taken out at a slower rate and displayed. This means that the CRT needn't run any faster than the processor can write data to it. Since the data is not written at the same speed as the input signal, a slower, less expensive raster CRT can be used. This brings some additional benefits besides cost savings, like higher reliability, lower cost, and longer life than high frequency analog CRTs. Other capabilities, like color, can be more easily implemented than is possible with high-speed analog CRTs.

Summary

In this section we talked about the concepts behind digitizing scopes while exploring how they work. We saw how digitizing scopes capture data by sampling it, storing it in memory, and then reconstructing it on screen. The major conceptual difference between this scheme and the way analog scopes work is that everything past the A/D convertor doesn't need to work as fast as the signal. In a digitizing scope everything following the memory works at the same speed as the microprocessor, instead of having to work at the same rate as the incoming signal. This makes it possible for the scope to process the waveform data and automatically measure parameters such as rise time, frequency, time intervals, and the like.

- There are two basic types of digitizing scopes—real time and repetitive. Real time digitizing scopes capture a waveform on a single pass, while repetitive digitizing scopes take a number of passes.
- Repetitive scopes can be further categorized into sequential and random repetitive. Sequential digitizing does not provide negative time viewing.
- Sample rate and bandwidth are related only in real time digitizing scopes.

In the next section we will look at how analog and digitizing scopes differ when being used to make the same measurements.



5

Specifications and Measurements

Instrument specifications have meaning only when they help you solve a problem. We will talk about some important specs here; more importantly, we'll show you how these specs relate to actual measurements. We'll also contrast how analog and digitizing scopes attack common measurement problems.

Vertical Resolution

The number of bits in the A/D of a digitizing scope helps set the vertical resolution. For example, a one-bit digitizer can measure two levels of a signal—one and zero, or high and low. A two-bit digitizer can discern four levels, a three bit gives eight levels, and so on. In general, a digitizer can discern 2^n discrete levels of an input signal. If you divide the input signal by the number of discrete levels of the A/D, you'll get an idea of what bits of resolution means in terms of voltage. For example, given a 1 V P-P input displayed full screen, a six-bit digitizing scope can measure voltage increments of

$$\frac{1 \text{ V (input V)}}{2^6 - 1} = 15.9 \text{ mV.}$$

This 15.9 mV is the resolution of the least significant bit (LSB) of the digitizing scope's A/D. In more general form we can say:

$$\frac{\text{Input Voltage}}{2^{(\text{Number of A/D Ranges})} - 1} = 1 \text{ LSB or Minimum Discernable Level}$$

Manufacturers may even specify the resolution in terms of the LSB.

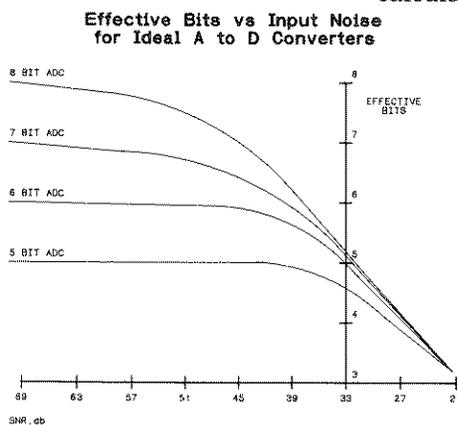
What does vertical resolution mean in terms of measurements? The issue that really concerns most users is what may be happening between those digitized levels. The maximum ideal error is $\pm 1/2$ LSB. To use the example on the previous page, if there is noise on the signal that's less than 15 mV, the digitizing scope may not display it. The only way to overcome this problem in a digitizing scope is to have more vertical resolution. However, having more bits in the A/D doesn't necessarily mean more vertical resolution. There are factors other than the number of A/D bits that affect the effective vertical resolution of a digitizing scope.

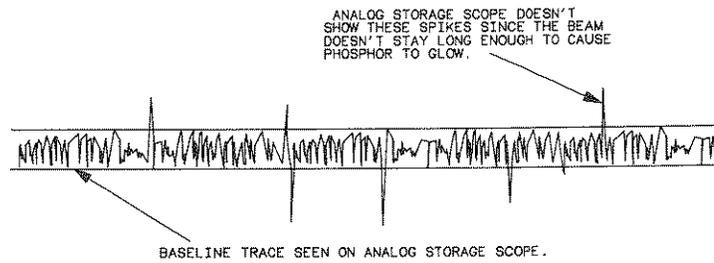
In Section 4, while talking about how to sample (page 4-11), we mentioned that some digitizing scopes use CCDs (Charge Coupled Devices) to capture the input waveform before it is digitized. This way a slower, higher resolution A/D may be used. While this allows more bits in the A/D, it also introduces more noise into the system before the waveform is digitized. So, while a higher resolution A/D can be used, there is more uncertainty given to the A/D in the form of noise. If you think of noise as high frequency components that are above the input bandwidth, you'll begin to understand why such noise is particularly undesirable. The input signal is bandwidth limited because the scope can't deal with frequencies higher than the bandwidth. However, CCDs add high frequency components (noise) that are not attenuated by the input filters because they are introduced from within the scope. Consequently, some of the additional resolution is negated by the effects of the internally generated noise. So, to get higher vertical resolution you must not only increase the number of bits in the A/D, you must also keep the internal system noise to a minimum.

Effective Bits of Resolution

How do you measure the vertical resolution of a digitizing scope if it involves more than just the number of bits in the A/D? We believe the best specification is one called "effective bits" of resolution. Effective bits is the resolution the user will see on the screen of the scope. To say it another way, effective bits of resolution quantizes internal system noise and its effect on the vertical resolution of a scope. An easy way to define effective bits is to think of it as an expression of a signal to noise ratio. In the last paragraph we talked about noise and how it lowers the effectiveness of the A/D. This lowered effectiveness is expressed by the quantity of effective bits. Effective bits takes a less than ideal A/D convertor system and expresses it in terms of an ideal A/D convertor system, i.e., one without noise and distortion. Effective bits will always be lower than the number of A/D bits and generally decreases as frequency goes up. How much lower depends on noise sources and nonlinearities within the scope (amplifiers, CCDs, etc.). The diagram below gives an idea of how noise can affect the A/D. Notice that as noise in the scope goes up, the effectiveness of the A/D goes down quite rapidly. So, when looking at digitizing scope specs, you need to know more than just the number of bits in the A/D.

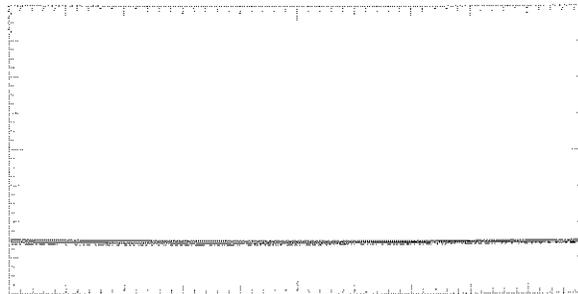
How does a manufacturer measure effective bits of resolution? One of the more common methods is called the "sine curve fit." With this method, a very stable sine wave of known frequency and amplitude is put into the scope. The digitized points are then read out to a computer. The computer calculates where the points should be for a sine function and then compares the digitized scope points to them. The difference between the calculated points and actual digitized points are used to calculate the number of effective bits.





Trace Noise

Because of the limited number of levels in the A/D of a digitizing scope, analog scopes are generally better than digitizing scopes for seeing some low levels of noise on a signal. However, digitizing scopes are often quieter internally than analog storage scopes. If you look at a baseline trace on an analog scope line, like the one above, it looks narrower than that of a digitizing scope. The digitizing scope trace, shown below, would seem to be much noisier. However, if you could turn up the intensity to maximum on an analog scope, without



it blooming, you'd notice that the trace is much fatter. Even that, however, doesn't give an accurate appraisal of how noisy the scope is internally. The beam of the scope is actually being deflected to higher levels, but since the beam doesn't stay at those levels for very long the phosphor doesn't receive enough energy to phosphoresce. Consequently, the trace looks much narrower than it actually is, indicating a lower level of noise.

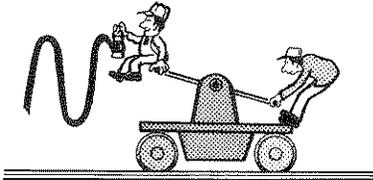
What this all means is that an analog scope may display a much lower noise level than is actually there. In addition, you have very little idea how much is scope noise and how much is signal noise. The governing principle is

$$\text{Noise}_{\text{Displayed}} = \text{Noise}_{\text{System}} + \text{Noise}_{\text{Measuring Instrument.}}$$

On the other hand, if you leave a digitizing scope in storage mode with no input, it will display every level it sees. In the picture below, a digitizing scope left for some time reaches a steady state of three digitized levels (zero, one level above zero, and one level below zero). You have a way of quantifying the system noise in a digitizing scope that you don't have with an analog scope.

If you don't know how much noise the measuring instrument (scope) is adding, you can't accurately determine how much of the displayed noise is from your system. With a digitizing scope you have a method of quantifying how much noise the instrument is adding to the measurement.

Time Interval Measurements

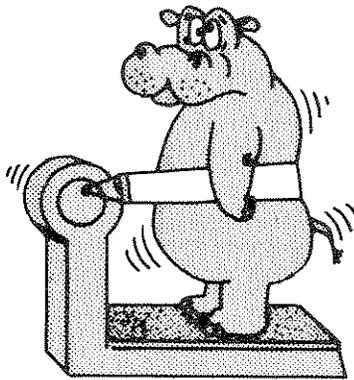


Some of the most common scope measurements are those involving a time interval. These measurements include finding a rise or fall time, determining a time interval, or measuring the frequency of a waveform. Traditionally, making such measurements with an analog scope meant displaying a part of the waveform and counting time divisions on the scope screen. Although these were time consuming measurements, there wasn't an easier way to make them with a scope. That is, until the digitizing scope.

For example, a rise or fall time measurement on an analog scope involves displaying an edge, expanding it to five vertical divisions, moving the edge up or down on screen, then counting the time divisions between the 10% and 90% lines.

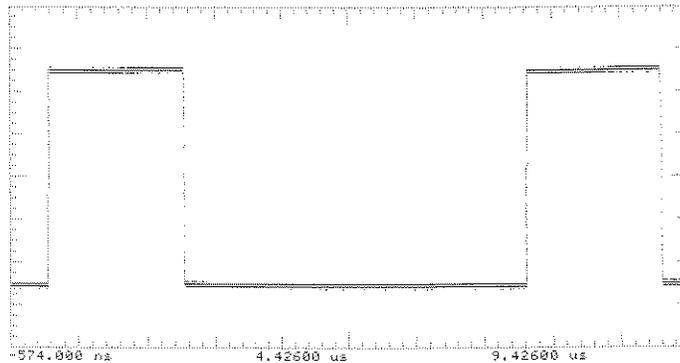
There is an easier way to make such measurements with a digitizing scope. For instance, in order to measure the rise time of an input waveform on an HP 54110D digitizing scope, you press the RISE TIME key. The scope displays the rise time at the bottom of the screen. As a matter of fact, any waveform parameter is just as easy to measure. If you want all of the parameters, you press the ALL key, and the scope measures the parameters shown at the bottom of the screen illustrated below.

Jitter



Another important time interval measurement is jitter. With an analog scope you display an edge, and then try to determine where the greatest excursions occur on screen before they fade. A common practice is to mark those excursions on the face of the CRT with a grease pencil.

The same measurement can be made more easily with a digitizing scope. You can put the scope in INFINITE PERSISTENCE mode, which lets all the data accumulate indefinitely. As you can see in the picture on the next page, every movement of the edge is accumulated on screen, making

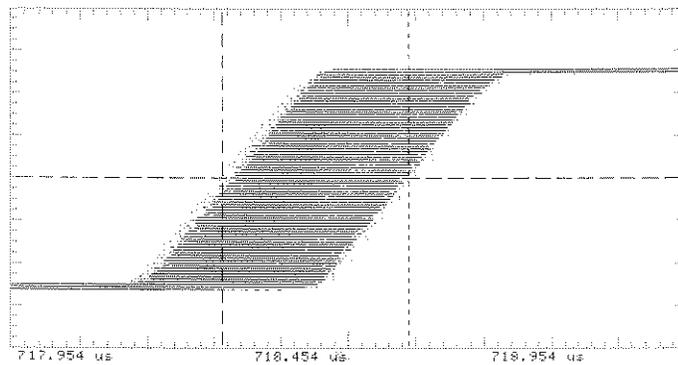


Ch. 1	=	1.000	volts/div	Offset	=	2.640	volts
Timebase	=	1.00	us/div	Delay	=	-574.000	ns
Ch. 1 Parameters				P-P Volts	=	5.125	volts
Rise Time	=	28.610	ns	Fall Time	=	4.420	ns
Freq.	=	141.699	kHz	Period	=	7.05722	us
+ Width	=	1.96379	us	- Width	=	5.09343	us
Overshoot	=	62.50	avolts	Preshoot	=	8.000	volts
RMS Volts	=	2.710	volts	Dutycycle	=	27.82	%

```

Trigger mode : Edge
On Pos. Edge on Chan1
Trigger Levels
Chan1 = 2.640 volts
Holdoff = 70.000 ns
    
```

it very easy to determine the total amount of jitter. Some digitizing scopes provide an ENVELOPE mode, which shows only the minimum and maximum excursions, making it even easier to measure the total amount of edge movement.

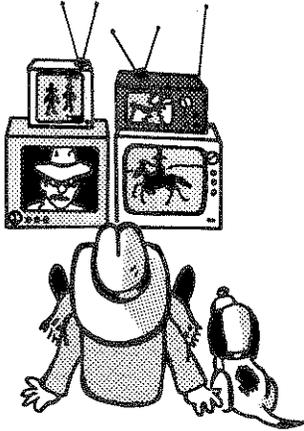


Ch. 1	=	1.000 volts/div	Offset	=	2.648 volts
Timebase	=	100 ns/div	Delay	=	717.954 us
Delta T	=	278.000 ns			
Start	=	718.266 us	Stop	=	718.544 us
Delta V	=	0.000 volts			
Vmarker1	=	2.648 volts	Vmarker2	=	2.648 volts

Trigger mode : Edge
On Pos. Edge on Chan1
Trigger Levels
Chan1 = 2.648 volts
Holdoff = 70.000 ns

Multi-Channel Display

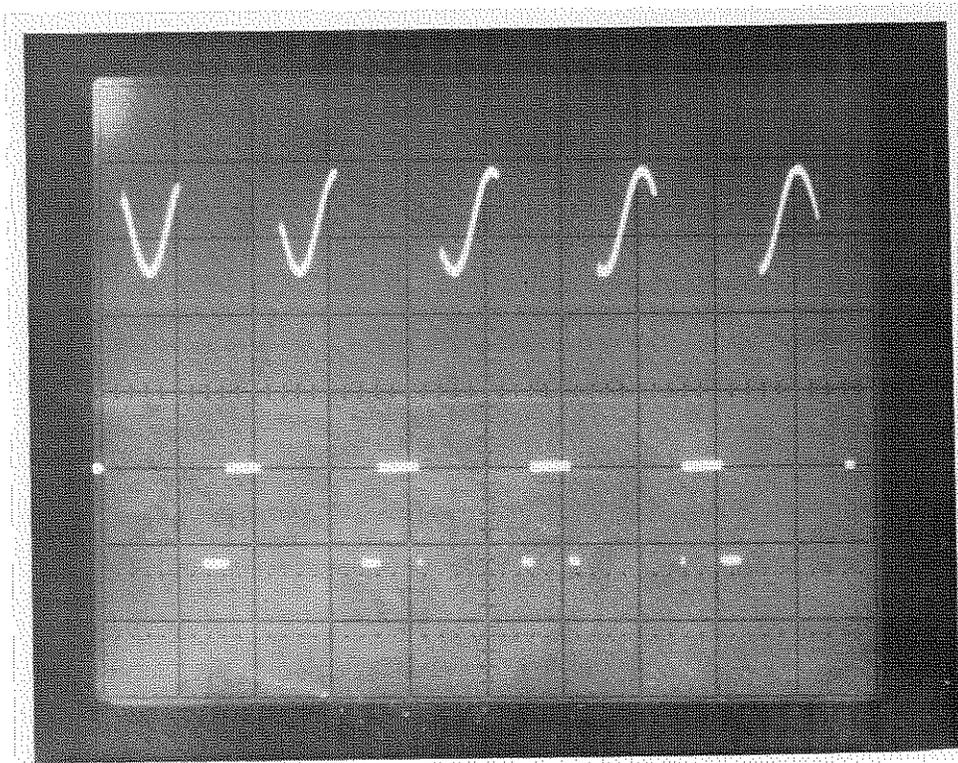
Many applications require a scope to display more than one channel. Most of us take that for granted. But let's look at some differences in bandwidth and display capability between analog and digitizing scopes when it comes to multi-channel displays.



An analog scope has several methods of displaying multiple channels. The best is a dual-beam CRT. A dual beam scope, as the name implies, has two electron beams. Since the CRT has two electron beams, one can be dedicated to each channel. However, their expense keeps many users from buying dual beam scopes for general purpose use.



The second method is called CHOP mode. CHOP mode uses a single CRT beam to draw both channels on a single sweep. The beam switches between channels at a rate between 100 kHz and 1 MHz. CHOP is used when you are trying to capture two asynchronous events and preserve their time relationship to one another. However, there is one drawback to CHOP that we can illustrate with an example. Suppose we have two 2 MHz signals we want to display. If we have the scope sweep speed set to $1 \mu\text{s}/\text{div}$ and the CHOP switching rate is 1 MHz, the beam spends $1 \mu\text{s}$ on one channel before switching to the other. That means the display shows one division of one channel and then spends the next time division ($1 \mu\text{s}$) on the other trace. When it switches back to the first trace, there is a $1 \mu\text{s}$ "hole" left in the trace. This continues throughout the trace on both channels as you can see below.



The truth is that in CHOP mode these “holes” are always there, but at slower sweep speeds, they are less noticeable and may even seem to disappear if the switching rate is fast in comparison to the sweep speed. At higher sweep speeds you may be missing critical information that occurs during the holes.

The third method of displaying multiple channels on an analog scope is ALTERNATE mode. In ALTERNATE the scope draws one complete trace and then draws the other. This method eliminates the “holes” in the traces produced by CHOP mode, but it introduces another problem. Since the scope triggers, draws one trace, then triggers again to draw the second trace, both triggers are shown at the left of the screen. How much time actually occurred between the triggers? We really have no idea—the time correlation between channels has not been preserved.

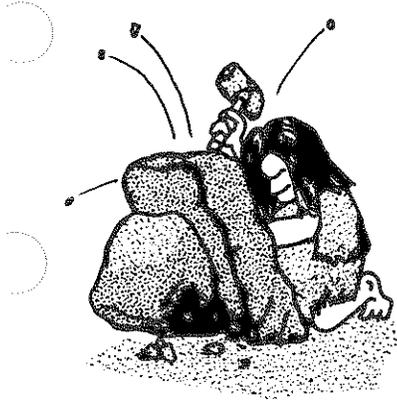
How does a digitizing scope overcome these multi-channel display problems? By having separate A/Ds. Information is captured and digitized simultaneously. By capturing data on all channels simultaneously, the time relationship among them is preserved. Secondly, since both channels are captured in their entirety, there are none of the holes associated with an analog scope's CHOP mode.

Storing Images On Screen

A major frustration with analog storage scopes is the limited amount of time images can be kept on screen. Within a few seconds the stored image fades positive (blooms), which means that the trace gets cloudy until it is just a bright cloud on the CRT. This is particularly true with single-shot events. This tendency to fade positive is related to the fact that the trace is stored electrically on a mesh inside the CRT (called the storage mesh). Because of a number of interactions within the CRT, the trace stored on the mesh tends to "spread" or become less distinct electrically, causing the bright cloud on screen. The end result is that the trace is obscured.

A digitizing scope stores traces in RAM rather than on an electrically-charged mesh. Once stored, the image can't change. Consequently, the trace can be kept on screen for an infinite amount of time without change.

Hardcopy Output



While developing a lab notebook or test procedure, it is desirable to have a picture of what we see on the screen of an oscilloscope. One frustration when keeping such records is using a scope camera. Anyone who has used a scope camera can probably remember taking several pictures before getting the proper exposure.

Again, because of their processor-based architecture, digitizing scopes simplify the task greatly. With scopes like the HP 54100 and HP 54200 series, you can use printers and/or plotters for hardcopy output. Once you have what you want on screen, you can transfer the waveforms along with setup and measurement information to the printer or plotter with the push of one button.

Summary

While adding a whole new set of measurement capability not possible with analog scopes, digitizing scopes have simplified most of the classic scope measurements. Because of their microprocessor-based architecture, digitizing scopes can easily interface with desktop and minicomputers in automated test environments. Here are some of the key points to remember about digitizing scopes.

- More A/D bits doesn't necessarily mean more resolution.
- Sample rate and bandwidth are not always related.
- Data is captured on all channels simultaneously, eliminating the need for CHOP and ALTERNATE modes found on analog scopes.
- The image can be kept on screen indefinitely without degradation.
- Because of their microprocessor architecture, automatic measurement of parameters like frequency, rise time, period, etc., is possible.
- You can easily copy what's on screen to a printer or plotter without having to use a scope camera.
- Connection to a desktop computer or minicomputer is easy for use in computer automated test systems.
- Waveforms can be stored in internal memory or on mass storage for future use and comparison.