

Design of a Mixed-Signal Oscilloscope

This combination of a digital oscilloscope and a logic timing analyzer provides powerful cross-domain triggering capabilities for capturing signals in mixed-signal environments. MegaZoom technology, consisting of advanced acquisition techniques and dedicated signal processing, maintains display responsiveness while making optimal use of deep sample memory.

by Matthew S. Holcomb, Stuart O. Hall, Warren S. Tustin, Patrick J. Burkart, and Steven D. Roach

The design of the HP 54645A/D oscilloscopes introduced in **Article 1** began with the HP 54645A, envisioned as a performance upgrade to the HP 54600 Series oscilloscopes.¹ These oscilloscopes, introduced in 1991, have an almost analog look and feel. Their three-processor design yields unprecedented display update rate and responsiveness at an affordable price. The major design goal for the HP 54645A was to improve the sample rate performance by an order of magnitude while maintaining the responsiveness and display update rate of the existing HP 54600 Series products.

Ultimately, two new products were created: the HP 54645A two-channel oscilloscope and the HP 54645D mixed-signal oscilloscope. The mixed-signal oscilloscope is a new product category that adds 16 logic timing analyzer inputs to the two channels of oscilloscope input. In addition to displaying all 16 logic channels, the HP 54645D provides advanced logic triggering functions on patterns that can span all 18 channels.

The HP 54645A and 54645D were designed concurrently. We made every effort to have the oscilloscope-only product (the HP54645A) be simply the combination product (the HP54645D), with an external trigger circuit substituted for the digital channel subsystem. Even the firmware ROM in the two products is identical.

Architecture

We started the design by modifying the architecture. A simplified block diagram of an HP 54600 Series oscilloscope is shown in Fig. 1. Two ICs form the core of the system: the acquisition processor and the display processor. The display system was left alone and only the acquisition circuitry was redesigned. We kept the same display, the same package, the same analog front end, and the same host 68000 processor.

In the original HP 54600 design, the acquisition processor IC was responsible for taking the digitized samples from the analog-to-digital converter (ADC) and placing them into an external memory as time-ordered pairs, which the display processor IC placed on the CRT display. More specifically, the acquisition tasks include:

- Generation of the sample clocks for the ADCs and controlling the decimation of the sample clock and dither algorithms for use at slower time base settings.
- Peak detection of the ADC data. This block calculates the minimum and maximum values of the sampled data.
- Intermediate storage of the ADC data (or minimum/maximum pairs if the peak-detector is used) into an internal circular 2K-sample memory. This memory, known as the *capture RAM*, holds the data until the trigger point is calculated.
- Accepting the analog trigger from one of the channels and measuring the time between the trigger and one of the sample clock edges.
- After the trigger is found, unloading the data from the capture RAM into an external RAM known as the *waveform RAM*. For each sample value, a corresponding x-axis (time) value is calculated.

All of these tasks were integrated into one chip for the HP 54600 oscilloscopes. For the new products, we divided the above functions into separate, discrete components, as shown in Fig. 2.

Clock Generation. Much of the difficulty of the original one-IC design stemmed from unwanted coupling between the sample clocks and the trigger clocks. In a digital oscilloscope, such coupling can severely corrupt the time base fidelity, causing time-axis nonlinearities (“bowing in time”), time-axis discontinuities, and sample bunching. In a higher-frequency instrument, the design of the clocking and trigger systems would have been all the more difficult. Consequently, the new products have a separate, dedicated bipolar IC for handling the clock generation and triggering.

Peak Detection. As before, we needed a handful of digital circuitry that stands between the ADC and the intermediate capture memory. This circuitry takes the 200-MSa/s 8-bit ADC data, calculates the appropriate minimum and maximum values, and stores the results sequentially into memory.

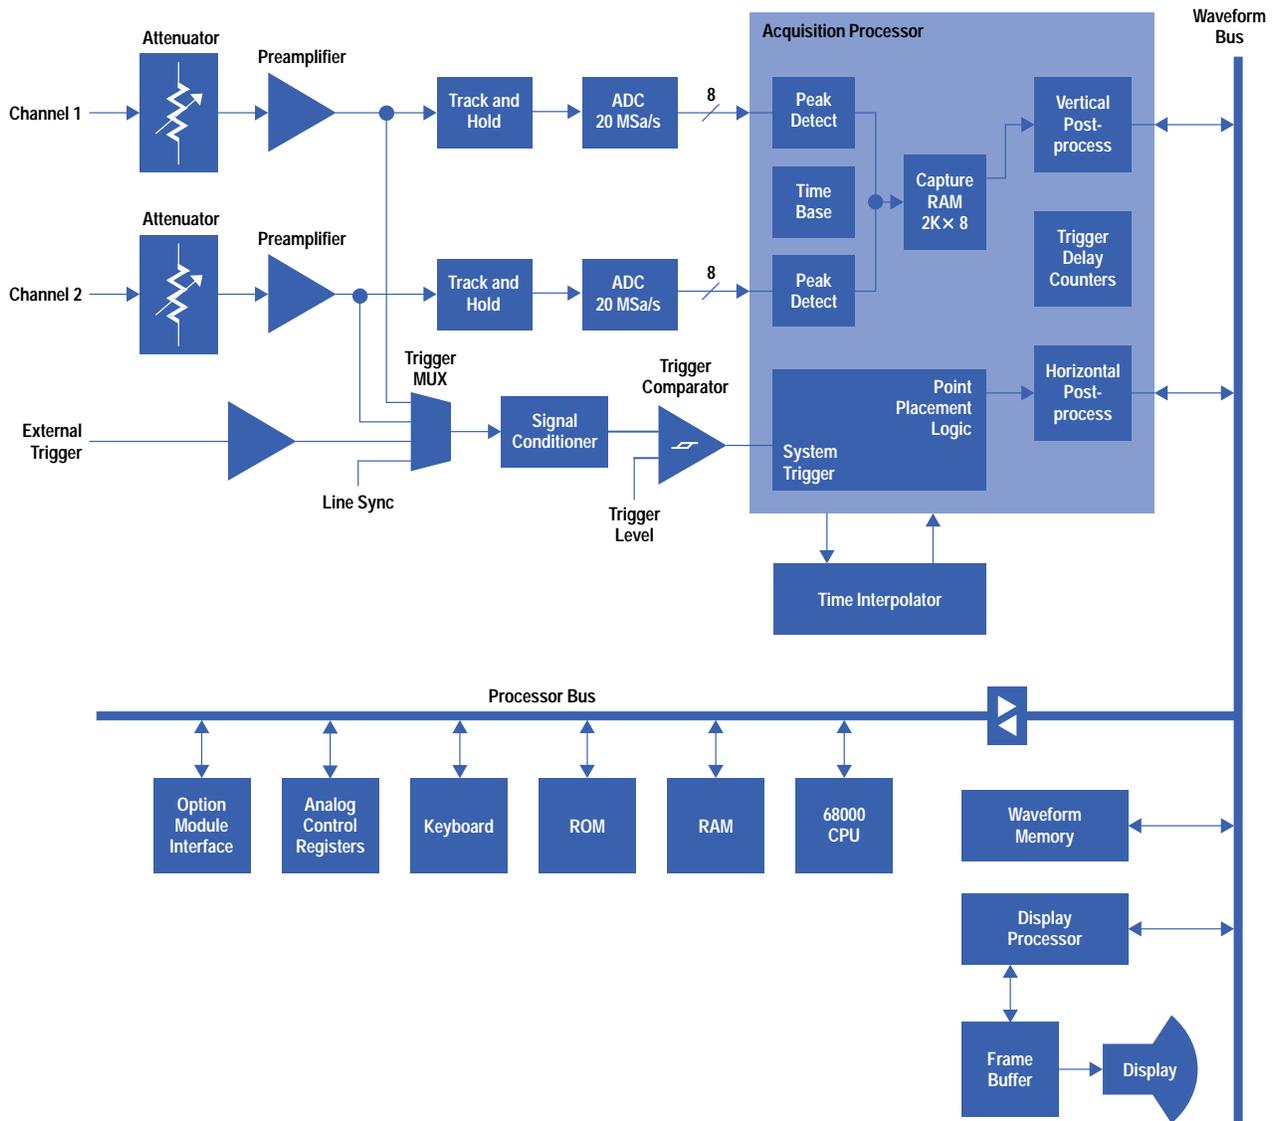


Fig. 1. Original HP 54600 Series oscilloscope block diagram.

Additionally, we decided to improve the averaging performance by digitally low-pass filtering the incoming sampled data before storing it into memory. This technique (later named *smoothing*) requires summing 32-bit values at a 200-MHz rate.

We chose to realize these functions in a CMOS gate array known as the *oscilloscope signal processor IC*. In addition to the functions described above, this IC decelerates (or fans out) the ADC data, steering and decimating the sampled data into the capture memory as described next.

Capture Memory. Recall that this memory is used to hold the incoming data until a suitable trigger point has been found and measured. In the HP 54600, the acquisition processor, being a fully custom design, had an internal 2K-byte (16K-bit) memory. Such memories are less available in gate array technology, so combining this function with the oscilloscope signal processor IC wasn't a particularly attractive option.

At 200-MHz rates, or 400-MHz rates for the timing channels, high-speed memories aren't off-the-shelf plug-in items. A traditional ECL design would be power hungry and costly. Our solution was to use commercially available video memories (VRAMs). Since these parts have a 50-MHz 16-bit interface, we could store data at a 100-Mbyte/s rate per memory. Each channel then requires only two devices. Since these memories are dual-ported, we can access the data (for plotting) simultaneously with the acquisition process, an advantage that we capitalized on further, as explained later.

Each VRAM has a capacity of 4M bits, resulting an 8M-bit (or 1M-byte) memory depth per channel. This increased memory depth is a huge improvement over previous designs and delivers value to the customer through increased time capture at a given sample rate (see [Article 4](#)).

Time Base and VRAM Controller. Along with the advantages of all of this deep memory came the problems of deep memory. Another CMOS gate array IC was designed for managing the storage and retrieval of data into the VRAM array. Its primary

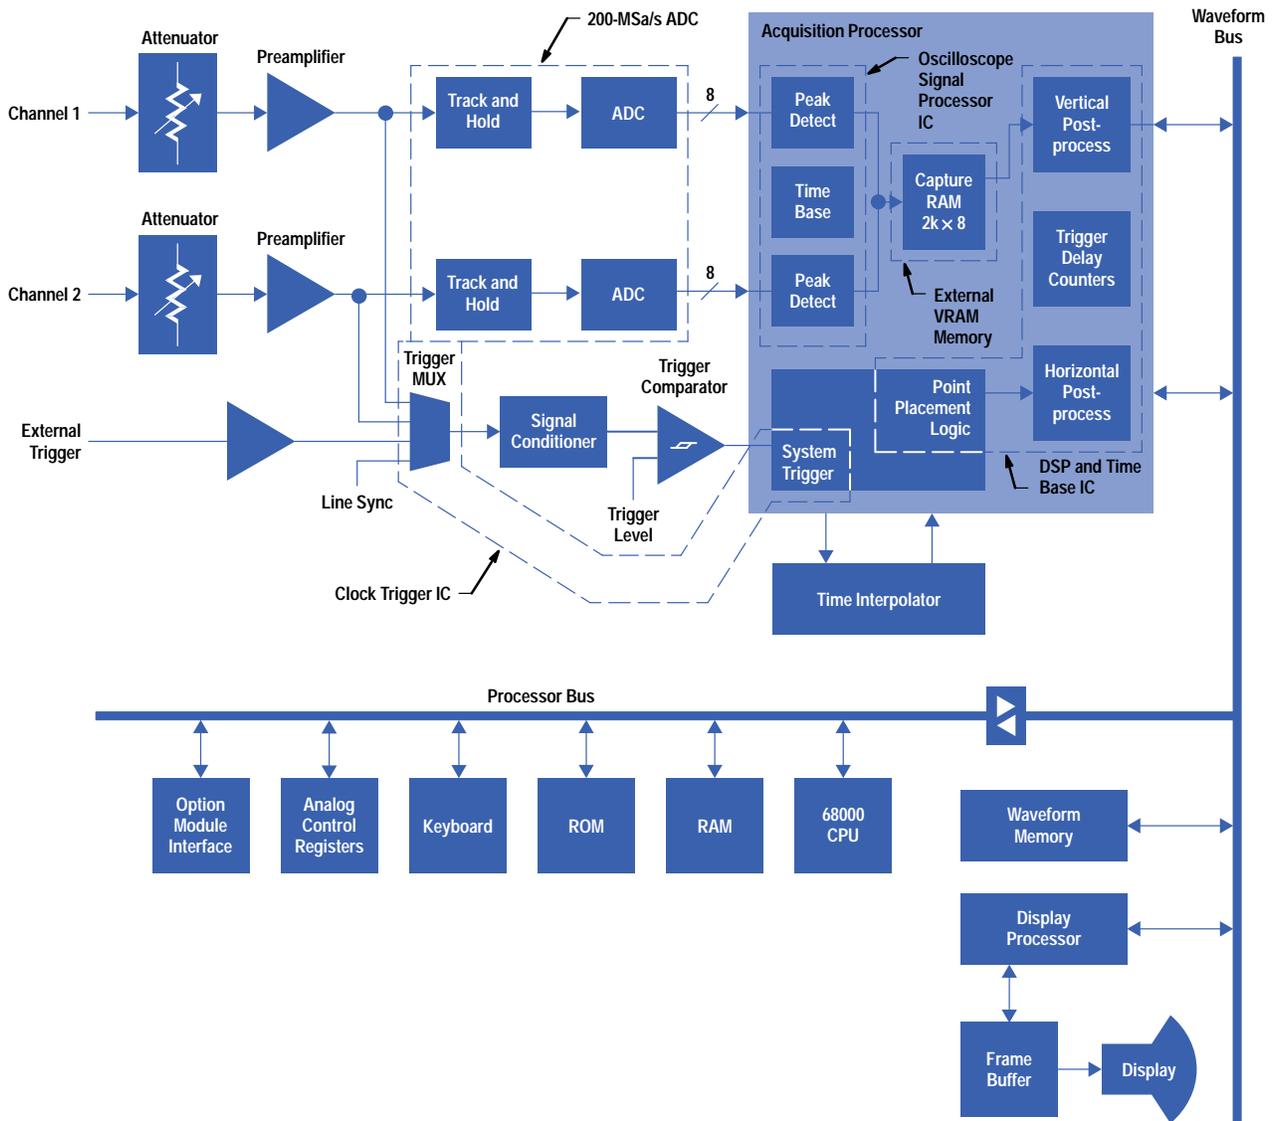


Fig. 2. HP 54645A oscilloscope architecture compared to the HP 54600. The topology remains relatively unchanged. Only the logic partitioning is different.

job is to provide the necessary commands to the VRAM to manage the data flowing into and out of its two ports. It also provides a high-speed interface to this memory for the rest of the system at a sustained reading or writing rate of over 10 million bytes per second. It controls the main time base algorithms, the trigger delay and holdoff counters, and a variety of interacquisition and intra-acquisition sample dithering circuits.

Data Processing. In previous products all of the time-axis acquisition algorithms were hard-coded into hardware. Such a system was fast and inexpensive, but more difficult to change. Therefore, for the new oscilloscopes, instead of hard-coding the mathematics into the time base and VRAM controller IC, we opted to use an off-the-shelf digital signal processor, or DSP: Texas Instruments' 25-MHz TMS320C50.

Like its predecessor, the DSP's job is to use the trigger information to determine where each of the samples stored in the VRAMs should be placed on the screen. Additionally, it handles all of the pan-and-zoom operations, allowing the user to redisplay the acquired data at other time base or delay settings after the acquisition is complete.

Mixed-Signal Oscilloscope

The idea behind the HP 54645D mixed-signal oscilloscope is to provide a truly seamless analog and digital waveform viewing instrument. A key goal for the project was to provide a product that was greater than the sum of just the two instruments. By combining an oscilloscope and a timing analyzer, we hoped to provide functionality and ease of use never before available in a laboratory quality oscilloscope.

Careful attention was paid to knitting together the analog and digital halves of the mixed-signal oscilloscope. Anything you can do with the analog channels, you can do with the digital channels: viewing, automatic measurements, triggering, and so on. Both halves share the same time base and the triggering system can operate on all 18 channels.

Where the oscilloscope channels have a *oscilloscope signal processor* CMOS gate array to decelerate and postprocess the ADC data, the digital channels have a similar *logic analyzer signal processor* CMOS gate array. However, the techniques for storing and plotting timing analyzer traces are radically different from oscilloscope traces. More specifically, on the faster time base settings, oscilloscopes use a technique known as *equivalent time sampling* to build up a display record over multiple triggers, resulting in an effective sample period of 50 ps or less (2.5 ps for the HP 54645A). Timing analyzer traces, however, are always single-shot—an entire waveform is reconstructed from each and every trigger, with an effective sample period exactly equal to the actual sample period. We therefore chose to double the maximum sample rate on the logic channels to 400 MSa/s to provide better single-shot time resolution. Using the same 50-MHz 16-bit VRAM memory as the HP 54645A, the HP 54645D has a total memory capacity of two million samples for each of eight digital channels (sampled at 400 MSa/s) or 1 million samples across all 16 digital channels (sampled at 200 MSa/s).

Use of Deep Memory

It has long been recognized that one of the factors slowing the acceptance of the digital storage oscilloscope over the traditional analog oscilloscope has been the superior update rate of the analog oscilloscope. Update rate describes the number of waveforms that the instrument can display per unit time. Since update rate is representative of the percentage of the input waveform that actually gets displayed on the screen, it translates to the amount of information relayed to the user. In the past, digital storage oscilloscope manufacturers used a single-processor architecture to handle all user interface, acquisition, and display tasks. This has resulted in digital storage oscilloscopes gaining a reputation for having far inferior update rate capability compared to analog oscilloscopes.

One of the advantages that digital oscilloscopes hold over analog oscilloscopes is their ability to store the data for closer examination after the acquisition is stopped. Digital oscilloscopes with deep memory have been gaining popularity because of their ability to store a given amount of time with more timing resolution than oscilloscopes with less memory. The deeper an oscilloscope's acquisition memory, the more samples per unit time it can capture. To make full use of the deep memory, digital oscilloscopes provide the capability of *panning* (moving the display window earlier or later in time) and *zooming* (changing the display window to display more or less time). This allows the user to display the portion of the acquisition that is of interest, with the desired amount of detail.

Unfortunately, the update rate of the oscilloscope normally suffers with deep-memory acquisition architectures. The update rate is inversely proportional to the amount of memory captured. Since every data point processed has to be read out of an ADC converter and placed in capture memory by the acquisition system, more points translates to more time to process those points. The deep-memory digital storage oscilloscopes on the market before the HP 54645A/D always place the update rate/memory depth trade-off in the hands of the user by means of a human interface control for the acquisition memory depth. If the user is interested in panning and zooming the acquired data after stopping the acquisition, the memory depth can be increased to gain more timing resolution for this task. If the user is only interested in viewing repeated acquisitions of the signal, the memory depth can be decreased to improve the update rate. This requires that the user have an understanding of the trade-offs involved and a knowledge of the user interface so that the changes can be made.

The HP 54645A/D oscilloscope architecture is designed to optimize the memory depth for the user's requirements. If the user is viewing repeated acquisitions of the signal, the memory depth is decreased if necessary to maintain the maximum update rate. If the user stops the acquisition, the memory depth of the last acquisition is changed to use all of the available acquisition memory. This feature is made possible by the same architectural feature that was designed to maximize the oscilloscope's update rate. During most continuously acquiring configurations, half of the acquisition memory is used to write newly acquired data while the other half is being read into the display. Since the architecture is fundamentally based on a 1-million-point memory system, 500,000 points are dedicated to the next trigger while 500,000 points are being read to the display.

If memory depth has been traded off for update rate, a subset of 500,000 points is used for each trigger during continuous acquisition mode. For example, at 200 $\mu\text{s}/\text{div}$, the acquisition time required to fill the screen with an acquired waveform is $200 \mu\text{s}/\text{div} \times 10 \text{ divisions} = 2 \text{ ms}$. The time required to acquire 500,000 points at 200 MSa/s is the sample period times the number of points, or $5.0 \times 10^{-9} \times 500,000 = 2.5 \text{ ms}$. Since the time required to capture 500,000 points is larger than the time required to fill the screen, we choose to reduce the number of acquired points to maintain maximum update rate. In this case, 400,000 points can be acquired in 2 ms ($400,000 \times 5.0 \times 10^{-9} = 2 \text{ ms}$), so we acquire 400,000 points. This maintains the theoretical update limitation of 500 waveforms per second ($1/0.002 = 500$). A more extreme trade-off occurs at 5 $\mu\text{s}/\text{div}$. Only 10,000 points can be captured in the 50- μs acquisition time ($10,000 \times 5.0 \times 10^{-9} = 50.0 \times 10^{-6}$).

If the user presses the Run/Stop key while an acquisition is in progress, the assumption is made that this may have been done to examine the last acquisition in more detail. The last acquisition that occurred before the user's action is stored in one half of the memory system. Since it cannot be assumed that more triggers will follow the user's action, this half of the memory may not be written to any more after the Run/Stop key is pressed. However, since the other half of the memory system now contains defunct data, the acquisition system is reconfigured to use the entire 500,000 points available in this other memory half, and keeps acquiring data, looking for a trigger. Since the primary task to be executed when the user presses the Run/Stop

key is to stop acquisition in what appears to be an instantaneous fashion, the system cannot wait forever for an additional trigger. Therefore, a timeout is initiated that will cause the system to cease looking for additional triggers after 100 ms. For the oscilloscope's acquisition system, this is a large window of opportunity to capture a deep-memory record. However, if no additional triggers occur, the acquisition system is brought to a stopped state, and the shorter acquisition obtained on the last trigger is used for postacquisition pan and zoom. Since most applications involve relatively frequently occurring triggers, this methodology will result in most users viewing deep-memory acquisitions every time they stop their oscilloscopes. For applications in which the rate at which triggers occur is uncertain, or in which there is only one trigger event, the user can press the Single key and guarantee a 1-million-sample acquisition for every trigger.

The Decimation Problem

Even with a million points of sample memory, at the slower time base settings, there will still be more samples across the screen than can be stored. We need to prune (*decimate*) the available data down to the memory size, then further thin the stored data down into 4000-point records for rapid display. In these situations, when the ADCs are sampling faster than the data can be used, the oscilloscope is said to be *oversampling*.

There are a variety of techniques available to perform this decimation process, each portraying a different representation of the input signal. Some methods will tend to highlight the subtle variations in the signal, others show the signal extremes, and others hide the signal extremes. As an example, suppose that the sample rate of the ADC is a factor of 100 greater than the desired store rate. Fig. 3 shows various decimation techniques for 100:1 oversampling.

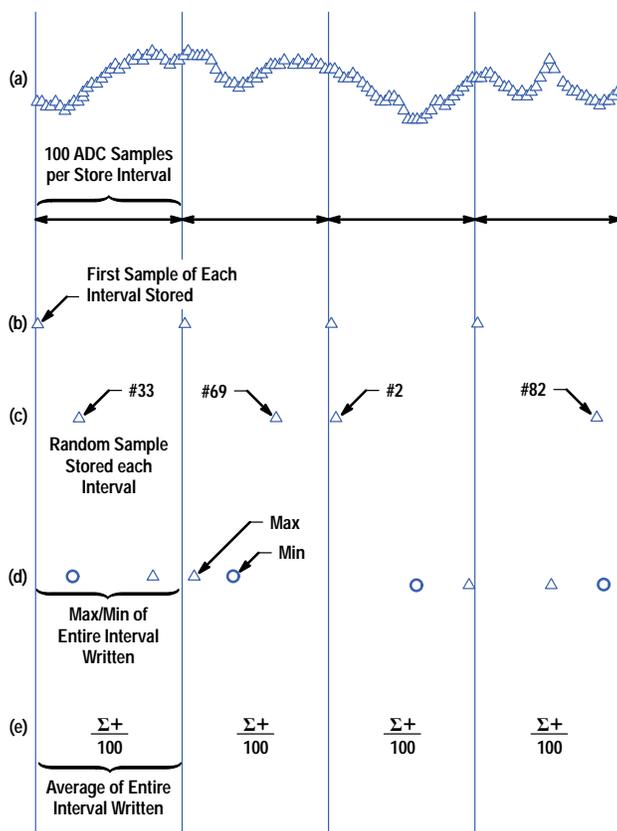


Fig. 3. Various decimation techniques shown at 100:1 oversampling. (a) ADC samples. (b) Simple decimation. (c) Intra-acquisition dither. (d) Peak detection. (e) Smoothing.

Simple Decimation. One sample is stored and the next 99 ignored; then the next sample is stored and the next 99 ignored, and so on. This approach is illustrated in Fig. 3b. Since it is the simplest to implement, it is the most common decimation technique in digital oscilloscopes. Unfortunately, because of the very regular spacing of the stored samples, it is exceptionally prone to aliasing. Therefore, the HP 54645A/D rarely use this technique (only when calculating FFTs, when exact sample spacing is called for).

Intra-Acquisition Dithering. Rather than store the first sample during the sample interval, this patented technique stores a randomly selected sample for each interval. In Fig. 3c, sample #33 is stored during the first interval, then #69 during the second, then #2 in the third, and so on. This technique, used in all of the 546xx oscilloscopes, is remarkably effective against aliasing. The stored samples are not evenly spaced, so it becomes much more difficult for the sampled signal to “lock” to the samples. The HP 54645A/D oscilloscopes use this technique with an improved folding pseudorandom number generator.

It is used whenever the instrument is in the normal display mode. For more information on this sampling technique, see [Article 5](#).

Peak Detection. Another common data compression technique is to store the minimum and maximum values of all of the samples in the sample interval, as shown in Fig. 3d. Widely used in digital oscilloscopes, this technique is called *peak detection*. When this technique is in use, rare, infrequent excursions are never ignored (lost), as they might be in the preceding two methods. This approach, however, tends to over-emphasize noise on the displayed record. Peaks are exaggerated and baselines become fatter at the expense of signal details. For example, an AM signal, peak detected, would show the modulation envelope quite clearly, but would lose the shape of the carrier wave. Statistical information about where the signal spends most of its time between the minimum and maximum values is lost.

The HP 54645A/D oscilloscopes combine traditional peak detection with intra-acquisition dithering when the instrument is in peak detection mode. They simultaneously acquire and plot both sampled versions of the incoming signal. The peak detected version highlights the signal extremes (one minimum/maximum pair per pixel column), while the denser conventionally sampled record (six points per pixel column) provides statistical information about the signal.

Smoothing. Yet another approach for decimating the incoming samples is shown in Fig. 3e. Here, rather than store the signal extremes for each interval, the average value is stored. This is logically equivalent to a simple low-pass boxcar filter cascaded with a decimating filter. At the slowest sweep speeds, millions of samples are averaged together for every point drawn on the screen, even on single traces. This display mode is useful for pulling the signal out of the noise.

While smoothing reduces the noise in the signal in a manner similar to conventional averaging, there are some differences. Smoothing works on signals acquired with a single trigger, while averaging requires multiple acquisitions to be effective. Smoothing functions as a low-pass filter with the cutoff frequency depending on the time base setting of the oscilloscope. When possible, the HP 54645A/D oscilloscopes use smoothing decimation when conventional averaging is turned on to provide additional noise reduction.

Logic Channel Decimation (Glitch Detection). The acquisition system of the HP 54645D mixed-signal oscilloscope has 16 logic channels. Decimating the series of 1s and 0s for these logic channels provides its own challenges. Simple decimation techniques would lose narrow glitches, so a peak detection technique (known as *glitch detection* in the logic analyzer domain) is always used. Two bits are sufficient to encode any sample interval, one value for each state (high and low). Not suffering from the drawbacks of peak detection on oscilloscope traces, logic waveforms are always acquired and plotted using glitch detection encoding regardless of the display mode selected for the oscilloscope channels.

Trigger System Features

Perhaps the most difficult task of the design was the trigger architecture for the mixed-signal oscilloscope. It needed to be a mixture of both analog and digital trigger systems. The traditional oscilloscope triggering modes (rising or falling edges on any channel, with a variety of coupling and noise-reject selections) needed to be coupled with the logic triggering modes (pattern triggering, sequencing, Boolean AND/OR, and so on). But more significant, all of the cross-domain triggering circuits needed to be defined and developed.

Development of this cross-domain triggering system architecture and choices about which trigger modes to support were guided by two key design objectives. The first of these was to seamlessly integrate the trigger configuration for all channels, whether analog or digital. This allows any channel to be used as a trigger source in any trigger setup. A second objective was to extend this seamless trigger configuration to a full set of trigger features. This feature set extends far beyond traditional single-channel, edge-mode triggering to include functionality essential in a mixed-signal test environment.

In mixed-signal testing, digital signals are often grouped together to form a common data or address word. Consequently, pattern triggering using terms including all channels was an obvious need. Inclusion of analog channels in patterns also provides the unique capability to gate a digital address trigger specification with analog-level conditions. Use of a pattern-qualified edge term (pattern AND edge) is a second key capability, since it allows a typical read or write to be captured by specifying the address in the pattern specification and the edge as the read or write strobe. Specification of trigger sequences in the mixed-signal oscilloscope is a further capability that is especially useful in capturing cause-and-effect relationships between analog and digital signals. Finally, pattern duration and Boolean combinations of trigger terms are included.

Trigger Types

Design of a cross-domain trigger system to meet these objectives was influenced by trigger methods specific to timing analyzers and digital oscilloscopes. Typically, timing (and logic) analyzers triggering on pattern or complex logical terms assume that all trigger events are simultaneous with clock transitions (in our case the ADC sample clock). This is referred to as *synchronous trigger generation*. In the analog world of the oscilloscope, all trigger events are assumed to be asynchronous to clock edges. Since each of these methods has advantages, the final implementation of the HP 54645D trigger hardware uses two independent trigger systems, one for synchronous and one for asynchronous triggering.

Trigger features typical of timing analyzer capability include pattern triggering, sequences of pattern terms, pattern duration, and logical combinations of pattern and edge terms. In a timing analyzer, digital channels are normally grouped together and

trigger conditions are evaluated only at sample clock intervals. There are several advantages to this type of synchronous sample clock triggering. One is that trigger events are always visible on the display because events that occur between sample clocks cannot cause triggers. Also, the signal transition causing the trigger event to occur will not display any jitter at the trigger time because that time occurs at a sample clock edge. Furthermore, complex trigger sequences can be easily implemented in hardware using sequential logic. Finally, it is easy to reject spurious patterns by qualifying pattern detection over multiple samples.

In a mixed-signal oscilloscope, triggering synchronously with the sample clock also has some significant drawbacks. Above all, it degrades equivalent time sampling operation for analog waveforms. This occurs because equivalent time sampling relies on random distribution of trigger events with respect to the sample clock and uses multiple triggered waveforms to enhance the analog display quality. The analog waveform display, when synchronously triggered at fast time base speeds, appears as a broken trace, with bunching of the signal at sample clock intervals (see Fig. 4). Since the HP 54645D samples at 200 MSa/s, degradation of waveforms is only severe at time base speeds of 20 ns/div or less. Furthermore, single-shot trigger captures are unaffected by synchronous triggering of the waveform.

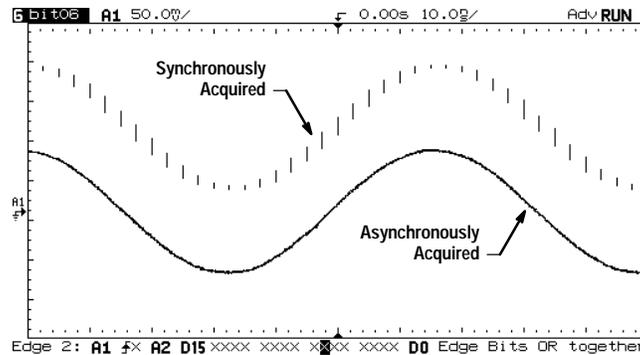


Fig. 4. Effect of synchronous triggering on the waveform display.

A second drawback of synchronous triggering is that if the sample rate is too low, trigger events can be totally missed. This is not a problem in the HP 54645D mixed-signal oscilloscope because the sampling interval (5 ns) is smaller than the minimum width of any pulse passed by the 100-MHz analog bandwidth (~7 ns).

Asynchronous trigger generation, traditional in digital oscilloscopes, provides precise trigger timing information to the waveform display. This allows the use of equivalent time waveform reconstruction to combine multiple trigger events and optimize signal display for repetitive signals at fast time base speeds. In addition, trigger time resolution for single-shot captures is limited only by analog channel bandwidth, rather than by sample time granularity. However, asynchronous trigger generation does have one significant drawback. It is possible that the event that generates the trigger may not be visible in the waveform display. This possibility exists because an analog event may cross a trigger threshold and then fall below the threshold again between successive sample clocks. In the HP 54645D, this effect is minimized because the analog bandwidth of the trigger path limits detectable input pulse widths to about 7 ns (depending on threshold level) while the minimum sampling interval is 5 ns.

Mixed-Signal Oscilloscope Trigger System Architecture

To achieve optimal signal display quality and maximize trigger capabilities, the HP 54645D mixed-signal oscilloscope uses both synchronous and asynchronous triggering. The hardware architecture, shown in Fig. 5, includes dual trigger paths that implement this trigger capability. The synchronous trigger path includes a full set of pattern and edge detectors, which allow independent detection of all trigger terms. A pattern duration block is also included to allow triggers based on pattern duration less than target time, greater than target time, and in a target time range. An important feature of this path is the inclusion of high-frequency pattern rejection blocks. These blocks require that patterns be present for a fixed number of sample periods. This enables rejection of patterns that are only present at bus transitions because of channel-to-channel skew. An example of this skew is shown in Fig. 6, where a transition is made between address values of D7 and A4 (hexadecimal). Since the sample clock is not synchronized with the address transition, the sample clock will occasionally occur during transition states (in this case the address B6) and trigger on these patterns unexpectedly. With the high-frequency reject enabled, these triggers are rejected because they are only present for very short times (typically less than 1 ns). The pattern rejection block is automatically programmed by the CPU to require a minimum of two samples for all pattern definitions with more than one defined level. Outputs from the trigger terms detected are then routed through combinatorial logic blocks, sequencing, and holdoff. The design of this path allows for a completely orthogonal set of trigger logic terms.

The asynchronous trigger path provides similar trigger functionality, including edge trigger, pattern detection, pattern-qualified edge detection, pattern duration, and Boolean logic combinations of trigger terms. Since this path lacks

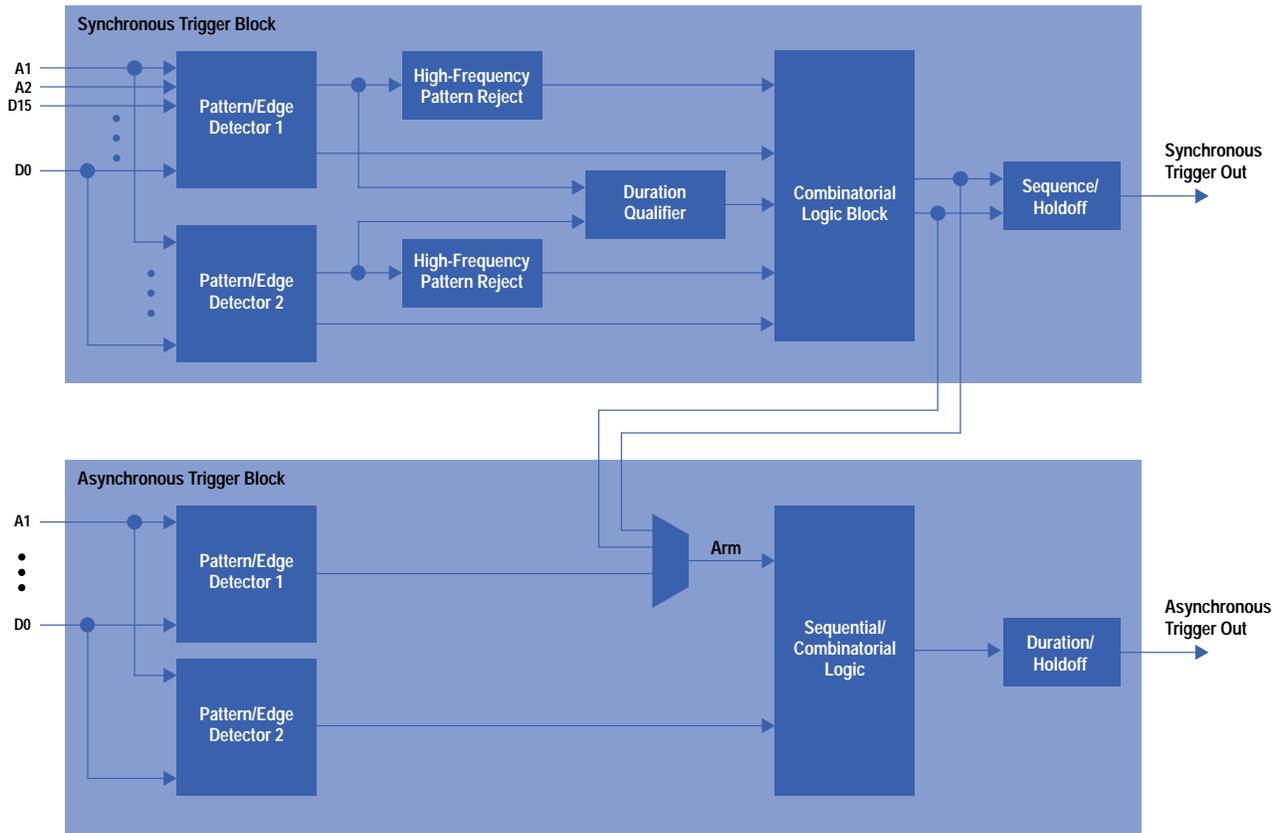


Fig. 5. HP 54645D trigger system block diagram.

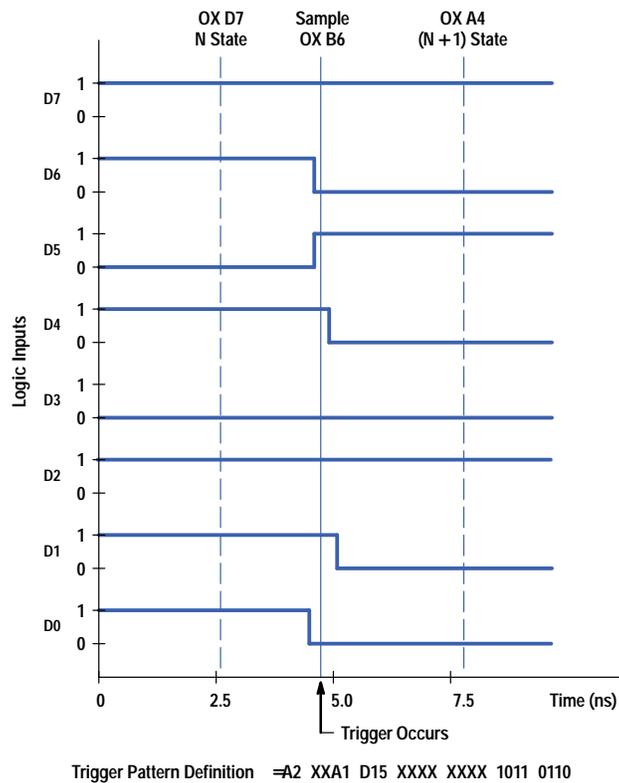


Fig. 6. Triggering on a channel-to-channel skew state (B6 hexadecimal).

multiple edge detection blocks, only single-edge trigger terms can be configured. Also, only one pattern-qualified edge detection can be configured. The ability to do cross-domain triggering with synchronously configured pattern terms arming the asynchronous trigger path significantly broadens the configuration choices. For example, an advanced pattern definition of (P1 AND E1) AND (NOT = P2) can be implemented by synchronously detecting P1 AND (NOT = P2) and using this to arm the asynchronous block, which then detects E1 (see Fig. 7).

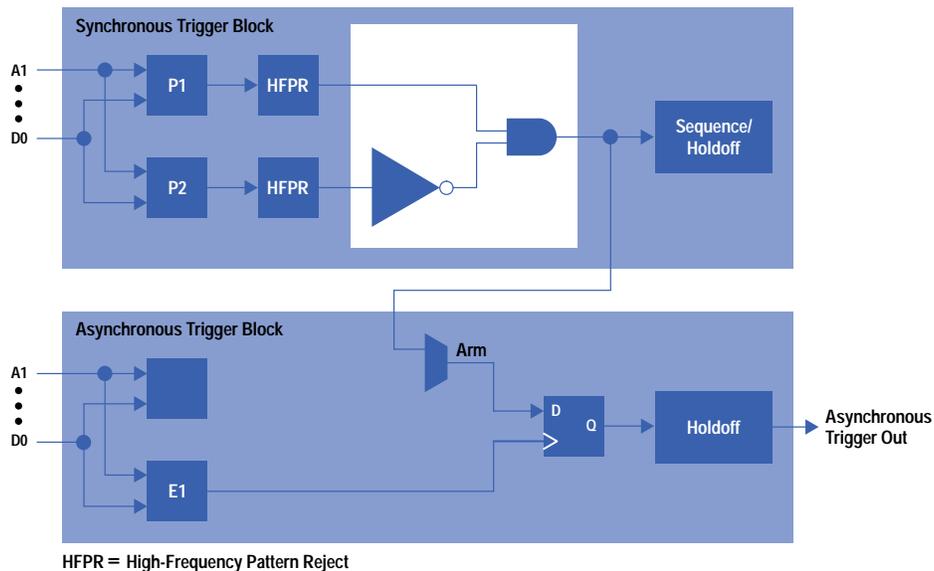


Fig. 7. Cross-domain trigger configuration.

Based on the current trigger term definitions and channel display selections, the mixed-signal oscilloscope CPU configures either the synchronous or the asynchronous trigger path as the trigger source for the instrument. To optimize the analog channel display quality, triggers are set up asynchronously, if possible, for all cases in which analog channels are on or analog edges are used in trigger terms. This context-sensitive implementation allows the use of equivalent time sampling to enhance analog waveform display in all possible cases. Consequently, waveform display is optimized without the need for the user to know the details of the trigger configuration.

Mixed-Signal Oscilloscope Trigger Applications

Typical trigger applications for the HP 54645D mixed-signal oscilloscope use the rich set of triggering functions to capture events combining analog and digital criteria. The following paragraphs present specific examples that show the unique capabilities of the mixed-signal oscilloscope trigger system.

ADC Performance Evaluation. Dynamic ADC performance can be evaluated using either pattern triggering to trigger on specific digital outputs or edge triggering with the analog threshold set to the level of interest. By providing an input sawtooth waveform, slew rate effects, missing codes, and nonlinearities can easily be evaluated.

Gated Processor Address Read/Write. An edge-qualified pattern detection can be used to detect a processor read or write to a specific address. By using analog channels in the pattern definition, it is possible to gate the trigger so that only processor reads and writes are displayed that occur when an analog signal is in a specific state. For example, in an alarm system in which a processor normally polls critical sensors to determine if alarm conditions exist, the use of a gated pattern and edge trigger could allow a display of all processor activity starting with (triggering on) the first processor read for which the sensor state indicates an alarm.

Glitch/Pattern Duration. Glitch triggering allows triggering on the duration between the positive and negative edges of a single signal. This allows capture of pulses less than a specific pulse width (glitches), greater than a specific pulse width, and within a specific pulse width range. This is very useful for capturing events that are in or out of specific timing requirements. Pattern duration allows the same duration-triggering choices but combines up to eighteen channels in the pattern definition so that triggering on the timing of combined channels is obtained. A specific example of using pattern duration is in capturing timing violations for a data handshake (see Fig. 8). The time for data acknowledgment after data valid must be less than a maximum time T. Violation of this timing could be found by setting pattern terms DAV low and DTACK high and looking for pattern duration greater than T.

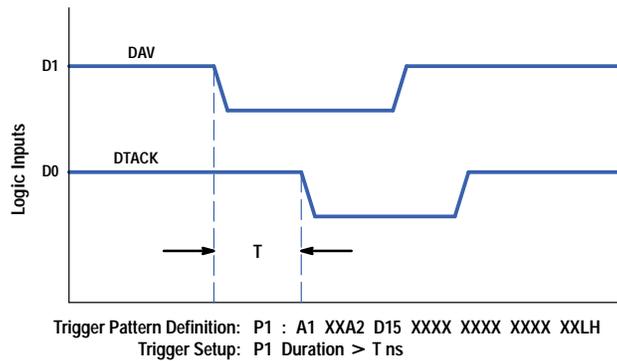


Fig. 8. Pattern duration triggering to find a timing violation.

Interrupt Response/Latency. Using a trigger sequence such as E1 THEN P1 AND E2 allows triggering on a processor response to an interrupt by first triggering on a hardware interrupt to a processor and then on a write to a register that is the ultimate response of the processor to the interrupt. The value written out can be read and subsequent events can be traced. Since the mixed-signal oscilloscope captures negative-time (pretrigger) events, this trigger can also be used to measure interrupt processing latency.

Register Write/Analog Response. In control applications, it is useful to be able to trigger on an analog signal that is controlled by the setting of a control register. This can also be achieved by a trigger sequence such as pattern AND edge THEN edge. For example in a motor control application, this trigger sequence can be used to trigger on a write to a control register followed by a motor commutation.

Display Enhancements

In the HP 546xx family of products, waveform update and display have always been high priorities in the design. Traditionally, digital oscilloscopes have had a slow update rate of data to the screen. In the HP 54645A/D the update rate has been increased to approximately three million points per second. This gives a much closer representation of the way an analog oscilloscope displays waveforms. For instance, on an amplitude modulated signal, an analog display would show the upper and lower boundaries with a filled-in area between them. On a traditional digital oscilloscope you would only see a few waveforms inside those boundaries. However, because of the increased update rate on the HP 546xx family, you see the boundaries and the area inside filled with data, a display that matches the traditional analog oscilloscope much more closely. This helps by displaying waveforms more closely to their true nature, but by itself does not address other features of an analog display such as its inherent intensity variation as a function of slew rate.

Analog oscilloscopes generally use a cathode-ray tube (CRT) in which an electron beam is swept across a display at a fixed speed while being deflected in the vertical direction by an analog signal. Because of this, areas of the waveform that have a high slew rate, such as edges, are displayed at a lower intensity than areas such as the top of a waveform, where the slew rate is low. This intensity variation provides information to the user that is not available with a traditional digital oscilloscope. For example, square waves displayed on an analog display are brighter at the top and bottom where the slew rate is low, and quickly changing edges are at a much lower intensity. For another example, consider a waveform that has a lot of baseline noise. On an analog display, the persistence of the phosphor causes the baseline to be very bright and infrequent noise to be dim. On a digital display, the noise is accentuated because all of the data points are displayed with the same relative intensity. This presents the digital oscilloscope designer using a raster display with a real challenge. How to display the waveforms to give a more accurate representation of their time-varying nature?

One thing to note is that the analog oscilloscope doesn't display infrequent signals very well. Because the phosphor has a limited persistence, the waveform will fade after a short time and if not redrawn will not be visible. One attempt to solve this problem is the analog storage oscilloscope, which uses a special plate in the CRT that has a much longer persistence. This is not a perfect solution because the waveform tends to bloom on this type of display and will still fade over time. It is also fairly expensive. A solution in the digital realm has been to use color, displaying data points that are sampled frequently in a bright color that is dimmed at a computed fade rate to simulate a long-persistence phosphor. However, color displays are more expensive, and more memory is needed to keep the counts for each data point. This can also be done with a monochrome display by varying the brightness of a data point in time, but more memory is still required, and in both cases an increased computational burden is placed on the display software. None of these methods addresses the problem of displaying different slew rates at different intensities.

The HP 54645A/D oscilloscopes use a proprietary variable-intensity algorithm to control the intensity of each data point. Adjacent points in the same raster row are used to set the intensity of the current point. Therefore, a point with horizontal neighbors will be brighter than if it has none. In the HP 54645A/D oscilloscopes, two intensity levels are used: full-bright and half-bright. When a point has a neighbor on both sides it is displayed full-bright. When it has zero or one neighbor, it is displayed half-bright. Thus, a high-slew-rate edge is displayed with less intensity because it has no horizontal neighbors.

Looking again at the square wave, the edges will be displayed at a lower intensity than the top and bottom of the signal, which results in a display that looks much more like the analog oscilloscope's representation of the waveform. For the signal with a noisy baseline, the noise away from the baseline is displayed at a lower intensity because the data points do not have horizontal neighbors. The cost to implement this enhancement is greatly reduced because inexpensive printed circuit board hardware is used. The update rate of the waveforms is not affected because this method can be used on the fly as data points are being displayed.

One other enhancement is used when the HP 54645D displays logic waveforms. These traditionally have a high value and a low value with edges drawn between them when the waveform changes. On a raster display, we have a fixed horizontal resolution in which to place the data points. The display is divided into rows and columns of addressable points called pixels. Suppose a point falls at the boundary between two horizontal pixels. On some acquisitions it will be displayed in one pixel, on others in the adjacent pixel. This results in a double-wide full-bright edge. Compared with a point that falls in the middle of a pixel, it will seem brighter. On a display with many logic waveforms this adds a perceptible flicker. The HP 54645D solves this by displaying double-wide pixels half-bright and single-wide pixels full-bright. This has the effect of equalizing the intensity and greatly reducing the perceived flicker.

In all of these methods an attempt was made to give the HP 54645A/D oscilloscopes a display that does its best to represent the waveforms the user might encounter.

Acknowledgments

We would like to acknowledge the technical contributions of Mike Beyers, Stu Hall, Dana Johnson, Pat Burkart, Warren Tustin, Derek Toepfen, Ken Gee, Frank Leri, Marty Grove, and Dan Timm. Jerry Murphy led the marketing effort. Bob Witte and Steve Warntjes managed the whole affair.

Reference

1. R.A. Witte, "Low-Cost, 100-MHz Digitizing Oscilloscopes," *Hewlett-Packard Journal*, Vol. 43, no. 1, February 1992, pp. 6-11.

-
-
- ▶ [Go to Subarticle 3a](#)
 - ▶ [Go to Next Article](#)
 - ▶ [Go to Journal Home Page](#)