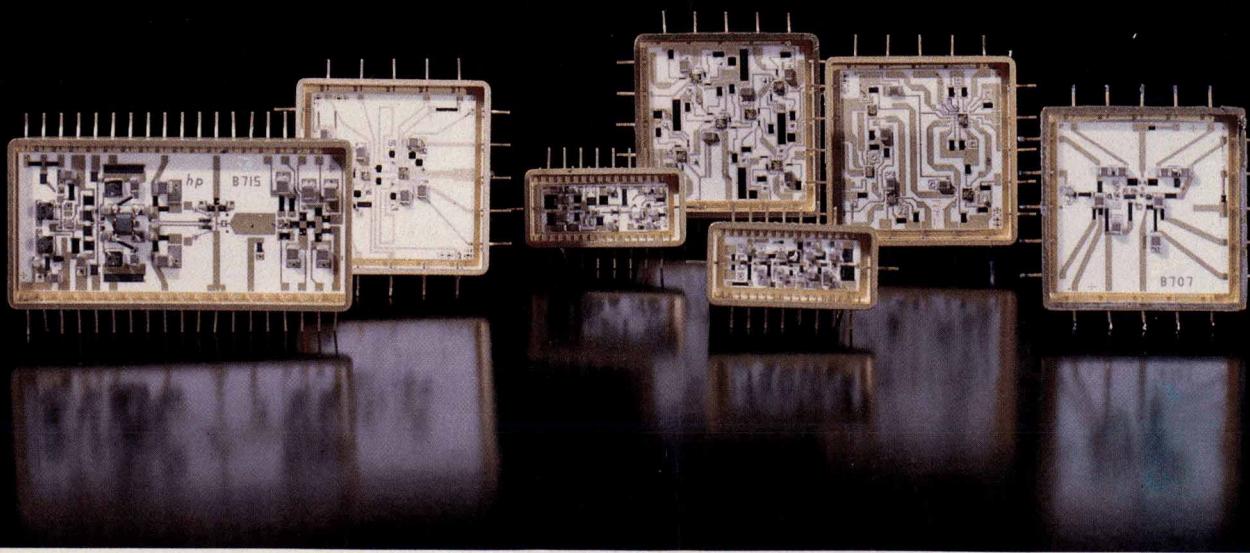
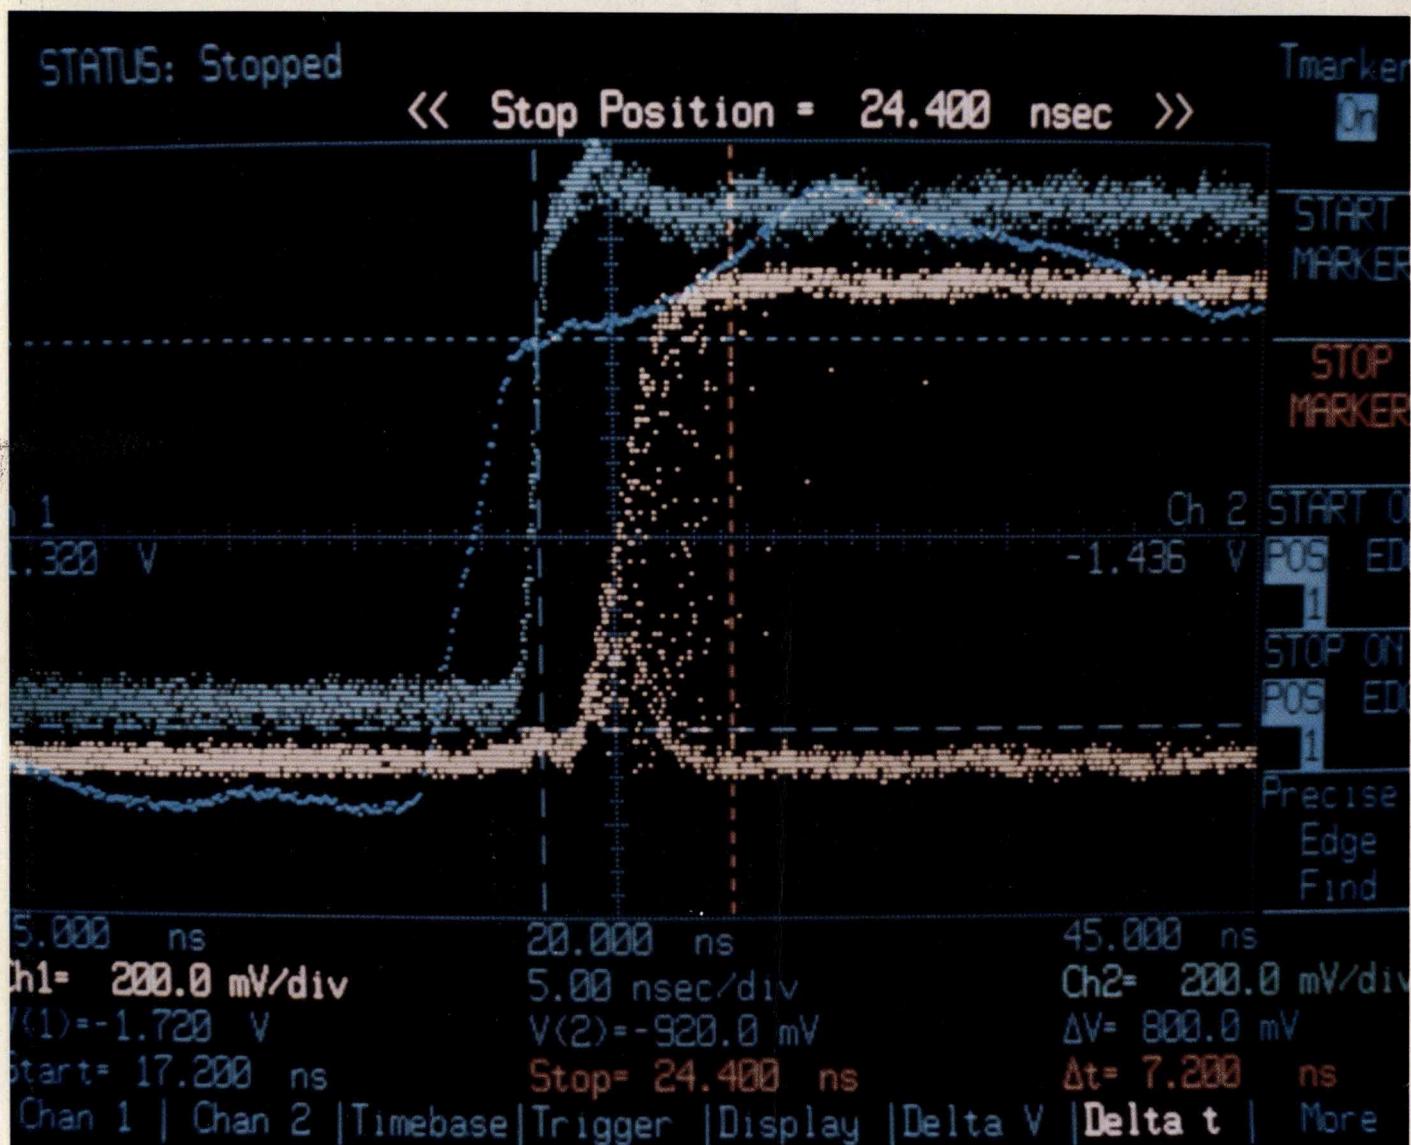


HEWLETT-PACKARD JOURNAL

APRIL 1986



HEWLETT-PACKARD JOURNAL

April 1986 Volume 37 • Number 4

Articles

-
- 4** **A Data Acquisition System for a 1-GHz Digitizing Oscilloscope**, by Kenneth Rush and Danny J. Oldfield Random repetitive sampling yields an equivalent digitizing rate of 100 GHz for repetitive signals.

5 General-Purpose 1-GHz Digitizing Oscilloscopes

- 11** **High-Performance Probe System for a 1-GHz Digitizing Oscilloscope**, by Kenneth Rush, William H. Escovitz, and Arnold S. Berger Because a typical system includes more than one class of signal, a versatile probe system is a must.
-

- 20** **Waveform Graphics for a 1-GHz Digitizing Oscilloscope**, by Rodney T. Schlater Fast plotting algorithms approximate the feel of an analog oscilloscope.
-

- 26** **Hardware Implementation of a High-Performance Trigger System**, by Scott A. Genther and Eddie A. Evel The trigger system resembles the front end of a logic analyzer.
-

- 33** **1-GHz Digitizing Oscilloscope Uses Thick-Film Hybrid Technology**, by Derek E. Toeppen Some new processes were developed and some old ones were improved.
-

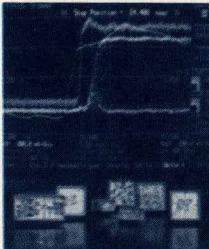
- 37** **A Modular Power Supply**, by Jimmie D. Felps A primary assembly drives two switching mode dc-to-dc converters.
-

- 40** **Program Helps Teach Digital Microwave Radio Fundamentals**, by Christen K. Pedersen Even someone entirely unfamiliar with electronics and communications can learn something.
-

- 47** Authors

Editor, Richard P. Dolan • Associate Editor, Business Manager, Kenneth A. Shaw • Assistant Editor, Nancy R. Teater • Art Director, Photographer, Arvid A. Danielson • Support Supervisor, Susan E. Wright
Illustrator, Nancy S. Contreras, • Administrative Services, Typography, Anne S. LoPresti • European Production Supervisor, Michael Zandwijken • Publisher, Russell M. H. Berg

In this Issue



One of my memories of my first engineering job is of my efforts to measure 0.1-microsecond pulse rise times with an oscilloscope—state-of-the-art at the time—that had a 0.35-microsecond rise time. How wonderful it would have been to have one of the scopes whose design is the subject of the first six papers in this issue. These instruments have 350-picosecond rise times, a thousand times shorter than my old scope's. On the other hand, the designers they're intended for are working with high-speed logic families and high-speed data communications, where rise times are typically less than one nanosecond. That's how it goes—better instruments help create better circuits, which create a need for better instruments, and so on.

The oscilloscopes discussed in this issue are the HP 54100A/D and HP 54110D Digitizing Oscilloscopes. All three models have 1-GHz bandwidth, 100-picosecond time base accuracy, 10-picosecond time interval resolution, and a digitizing rate of 40 megasamples per second. Using a technique called random repetitive sampling (page 4), they sample the input signal, convert the samples to digital form, store the resulting numbers, and display them on demand. Their storage capability lets them show input events that occurred before the trigger event and provide infinite-persistence displays. The HP 54100A has three input channels—two vertical and one trigger. The HP 54110D has an extra trigger input and provides complex pattern-recognition triggering like a logic analyzer (page 26). The HP 54110D adds a high-resolution full-color display, waveform math functions, semiautomatic channel-to-channel time interval measurements, and automatic measurements on stored waveforms. HP expects that up to 80 percent of its customers for 1-GHz oscilloscopes will prefer the full-color HP 54110D. All three models are fully programmable for use in automatic manufacturing test systems.

The box on page 5 introduces these three scopes and discusses the need for them. The paper on page 33 describes the thick-film hybrid technology that's responsible for the scopes' performance. Shown on the cover, in front of a display from the HP 54110D, are the seven hybrid circuits in each instrument: a 1-GHz preamplifier, a 3-GHz sampler, a 1-GHz probe receiver, a 300-MHz probe receiver, and three trigger hybrids making up a 500-MHz trigger system. Other papers describe the scopes' data acquisition system (page 4), probe system (page 11), waveform graphics (page 20), trigger system (page 26), and power supply (page 37).

Until recently, the most cost-effective method for long-distance voice communications over terrestrial and satellite microwave links was analog transmission using frequency division multiplexing and mechanical switching. Now, thanks to advances in integrated circuits and high-speed logic, the balance is beginning to tip in favor of the all-digital communications system, using pulse code modulation, time division multiplexing, and digital switches. The article on page 40 describes a software package that teaches the fundamentals of digital microwave radio by simulating a real system in action. It's intended for technical people who are new to digital communications, and for senior engineers and engineering managers who need to increase their familiarity with the subject.

-R.P. Dolan

What's Ahead

The design and features of a line of instrument modules for use with personal computers is the primary subject of the May issue. These members of HP's PC Instruments family make use of a personal computer's display as a soft front panel to reduce their cost. They come in small packages that can be stacked and racked.

Also included in the May issue is a research report about a new VLSI metallization technology using titanium silicide.

The HP Journal encourages technical discussion of the topics presented in recent articles and will publish letters expected to be of interest to our readers. Letters must be brief and are subject to editing. Letters should be addressed to: Editor, Hewlett-Packard Journal, 3000 Hanover Street, Palo Alto, CA 94304, U.S.A.

A Data Acquisition System for a 1-GHz Digitizing Oscilloscope

by Kenneth Rush and Danny J. Oldfield

THE DREAM OF MANY oscilloscope users is to own an instrument that can store for all time a record of a single event to high resolution in both amplitude and time. Towards this end, many people around the world are designing very fast analog-to-digital converters (ADCs) and producing digitizing oscilloscopes. The approach has typically been to digitize to a modest resolution at very high rates. Currently, several manufacturers offer 8-bit, 100-MHz ADCs. The reason for the emphasis on faster ADCs is the need of end users for greater time resolution in their measurements.

If we recognize that many signals of interest can be made to repeat, then alternate methods can be used to achieve high time resolution in a digitizing oscilloscope. This article describes a system used to realize one of these methods, called random repetitive sampling. The system achieves an equivalent sampling rate of 100 GHz and a bandwidth of 1 GHz for a time resolution of 10 ps.

Random Repetitive Sampling

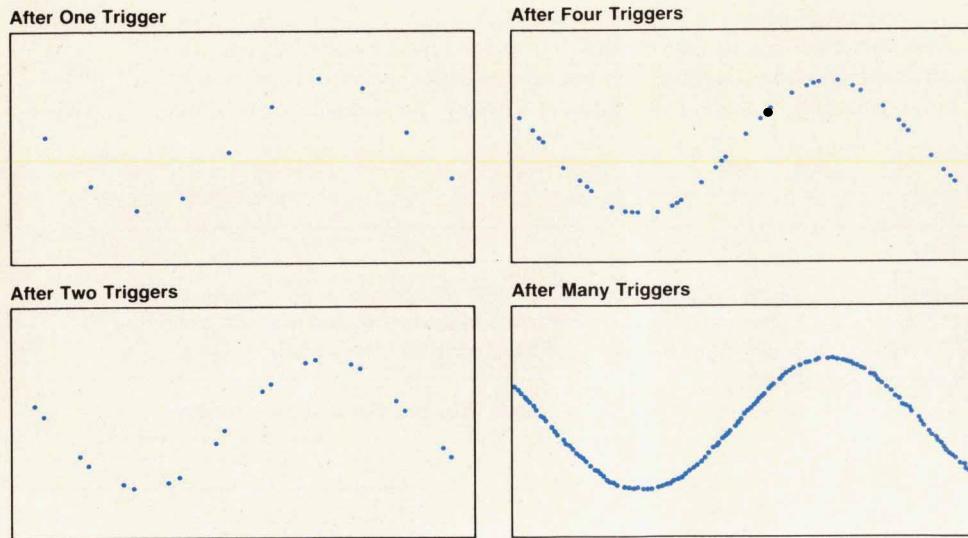
The HP 54100A/D Digitizing Oscilloscope makes random observations of input signal voltages and stores these observations in memory. The true timing of these observations is determined by measuring the time relationship between a trigger event related to the input signal and the sample clock in the instrument. Every 25 ns the input signals are sampled, digitized, and stored into memory. While each sample is time correlated with every other sample because the samples are spaced 25 ns apart, the samples are not correlated in time with the input signal. When a trigger event is detected, the timing relationship between the trigger event and the sample clock (40 MHz) is measured to a

resolution of 10 ps by a circuit called the trigger interpolator, to be described later.

If we are able to keep the information gained from this first measurement, that is, the sampled voltage values and the time relationship value, then the measurement could be repeated to obtain a new set of voltage values and a new random time relationship value. Repetitively making this measurement on a signal that doesn't change much from measurement to measurement will yield an accurate image of the signal waveform with 10-ps time resolution by reconstructing the waveform from the information stored over many measurements. Fig. 1 illustrates how data accumulates in a typical case.

If we have an acquisition memory depth of 1024 samples, and we do not arm the trigger until we have stored 1024 points, then the memory is holding a history of the signal that spans about 25 microseconds. If we stop storing samples as soon as we detect a trigger event, then we can accumulate over many triggers up to 25 microseconds of information before the trigger, even at high time resolution. If the sample rate is reduced by a factor of ten, then the amount of pretrigger viewing can be extended to 250 microseconds. This method of reducing the sampling rate is used to achieve very long pretrigger viewing in the HP 54100A/D. This is one of the contributions of random repetitive sampling, that is, pretrigger viewing with high time resolution without the necessity of delay lines to time shift the signals.

The HP 54100A/D has a posttrigger delay counter to program what portion of the 1024-point acquisition memory is stored before and after the trigger. When measurements are to be made a long time after the trigger, it is not necessary to change the sample rate. The posttrigger counter is simply



(continued on page 6)

Fig. 1. In random repetitive sampling, the accumulation of information over many triggers or measurements produces an accurate image of the signal waveform.

General-Purpose 1-GHz Digitizing Oscilloscopes

The six articles on pages 4 through 39 of this issue report on various aspects of the design of the HP 54100A/D and HP 54110D Digitizing Oscilloscopes. These general-purpose oscilloscopes take advantage of a completely digital oscilloscope architecture, random repetitive sampling, and state-of-the-art technologies to meet the new demands faced by engineers involved with digital design and high-speed data communications. The HP 54100A/D and HP 54110D simplify analog time-domain measurements on high-speed logic circuits. They can make the new measurements needed when working with the most recently developed logic families, and they can make many measurements faster and more accurately than conventional oscilloscopes.

Introduced in October 1984, the HP 54100A/D was HP's first completely digital oscilloscope, and industry's first 1-GHz-bandwidth digital oscilloscope. The A model has two vertical channels and one trigger channel, and provides three-bit pattern recognition triggering. The D model has an additional trigger channel and provides four-bit pattern recognition triggering and other complex triggering features. The HP 54110D (Fig. 1), introduced in October 1985, adds a color display and several additional measurement capabilities to the feature set of the HP 54100D.

In addition to their 1-GHz bandwidth (350-ps rise time), the HP 54100A/D and HP 54110D also provide a 100-ps/div time base and 10-ps time resolution. These capabilities allow the user to measure subnanosecond rise times, propagation delays, and glitches.

A 40-megasample/second digitizing rate and infinite persistence mode allow the user to capture fast, low-repetition-rate signals, infrequent error conditions, or variations in a signal. The infinite persistence mode is also useful when looking at worst-case conditions such as noise, jitter, and drift.

Three types of interchangeable input pods allow the user to measure signals from a variety of source impedances. The HP 54300A Probe Multiplexer accepts all three input pods and allows up to 16 probes to be multiplexed at one time. Up to two HP 54300As can be cascaded in series, increasing the number of inputs that can be multiplexed.

Design Trends

The oscilloscope has always been the workhorse of the digital circuit designer. This is still true today, especially for designers working to advance the state of the art in processes and technologies. Three trends pose a particular challenge to oscilloscope manufacturers if they are to continue to provide a useful tool for the digital designer: the increasing speeds of digital circuits, their increasing complexity, and the rate of turnover in technologies and products.

Not only are designers turning to faster logic families such as ECL 10KH and 100K, but semiconductor technologies are constantly undergoing improvement. With TTL and even CMOS now commercially available with edge transition times and gate delays approaching one nanosecond, an oscilloscope must have a 1-GHz bandwidth to be truly useful when working with these logic families.

As system speeds increase, bus cycle and memory cycle times shrink. With more going on in less time, the designer must pay stricter attention to precise timing. Transmission line characteristics are becoming more important. With a shorter cycle time, solutions to timing problems that may have worked in the past will no longer suffice. For example, with no extra time in the bus-cycle-time budget, the designer may not be able to slow down the clock rate or reshuffle events within the allotted time. With this increased emphasis on precise timing, good analog



Fig. 1. HP 54110D Digitizing Oscilloscope adds a color display and other functions to the features of the HP 54100D.

design techniques are essential to the success of digital designs. The oscilloscope is the most useful tool for characterizing the analog performance of digital circuits.

Examples of increasing logic speed are not limited to semiconductor gate delays and transition times. Disc drive manufacturers strive to pack more bits per square inch on the surface, which translates to higher data rates in both read and write circuits. Fiber optics and satellite communications offer the attraction of getting more information to more places faster and with fewer errors, pushing data rates up to gigabits per second. Local area networks, data multiplexers, and the computers' internal buses must all go faster to keep up with the logic circuits' ability to perform mathematical operations on numbers faster.

Oscilloscope Requirements

A 350-ps rise time, by itself, is not sufficient for an oscilloscope to make the required measurements. The precision, resolution, and range of time interval measurements are probably the most important attributes of an oscilloscope. The most basic measurement expectation of an oscilloscope is to measure the time between events on waveforms. How long after the driver started its transition did the input to the last RAM pass through logic low? How far from the center of the data window was the peak of a pulse on a disc drive? By how much time do the edges of the clock vary in different parts of a system?

In complex digital systems, it is important to be able to view events leading up to the trigger. Otherwise, it is difficult if not impossible to determine the cause of events. In an asynchronous system, where cycle times are not always the same, reliably measuring such basic parameters as setup time requires the ability to display events before the trigger.

One of the most frustrating behaviors of complex digital circuits for digital designers is the tendency for catastrophic errors, which occur very infrequently. Digital systems have an almost infinite number of states they can occupy, and an erroneous operation may be related to only one of these states (which the system has a very low probability of being in at any time) or, worse yet, to a unique sequence of states. Even normal periodic events, such as an individual pulse on a disc drive, may have a very low duty cycle in a complex digital system. It is difficult for the digital designer to see such events using a conventional oscilloscope. The infinite persistence mode on the HP 54100A/D retains information on the display indefinitely, so the designer can see

those very infrequent error conditions, as well as examine signals having a very low duty cycle.

Faster means smaller and denser, since distance is proportional to time in electric circuits. As circuits become denser, they become harder to probe. Frequently, the hardest part of making an oscilloscope measurement is probing the desired signal. At the same time, the effect of the probe on the circuit under test, and the distortions introduced by the probe and its grounding path on the signal going to the oscilloscope, are often the dominant source of error in high-bandwidth oscilloscope measurements; this effect becomes more pronounced at high speeds. Therefore, a truly general-purpose oscilloscope must offer solutions to the conflicting requirements of probing: small physical dimensions, low mass, high bandwidth, and minimum loading (particularly capacitive loading).

Chips, the simplest and lowest-level circuit building blocks, are rapidly becoming more complex, with many more functions and connections. Examples range from gate arrays to fiber optic receivers. Today's board is tomorrow's chip. As this happens, the number of variables in their transfer function increases exponentially, and the number of opportunities for errors that have an undesired impact on system operation increases proportionally. To characterize such a circuit—and especially to troubleshoot it if it isn't working properly—may require thousands of measurements.

Thus, the designer needs instruments—including oscilloscopes—that can be programmed to make measurements automatically. A truly general-purpose oscilloscope must be programmable, but shouldn't require a doctorate in computer science to program effectively. This makes it possible to program computers to take over the repetitive tasks of setting up measurements and of gathering, sorting, and analyzing the data. A new-generation oscilloscope must also be easy to operate manually. All the measurement features in the world are of little benefit if they are obscured by an arcane control architecture, or hidden in a bewildering array of buttons and knobs.

Meeting all of these requirements was the basis of the design objectives for the oscilloscopes discussed in this issue.

Arthur W. Porter
Product Marketing Engineer
Colorado Springs Division

(continued from page 4)

programmed to count the correct number of sample periods (greater than 1024) to reach the time window of interest. Using a crystal to control the sample clock guarantees high time base stability. Very long delays can be achieved without the long time base ramps required in conventional oscilloscopes. This leads to low overall time base jitter and

counter-like time base accuracy.

The availability of very high time resolution leads to the possibility of high bandwidth. By designing a high-bandwidth sample-and-hold circuit and using random repetitive techniques, we can achieve very high bandwidth. In the case of the HP 54100A/D we achieve bandwidth

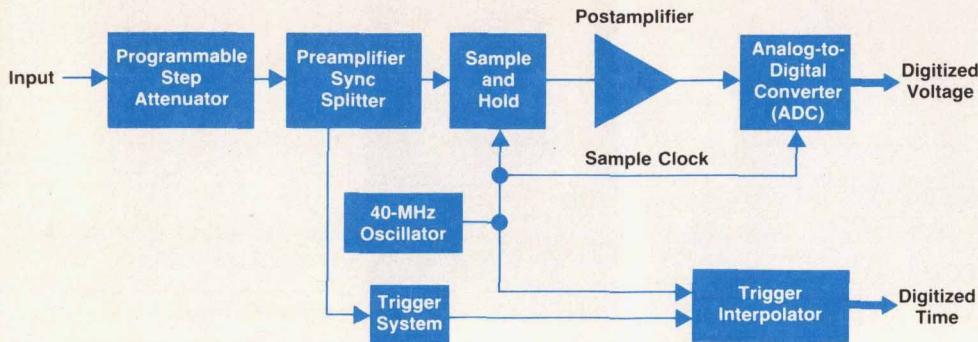


Fig. 2. Simplified block diagram of the HP 54100A/D data acquisition system.

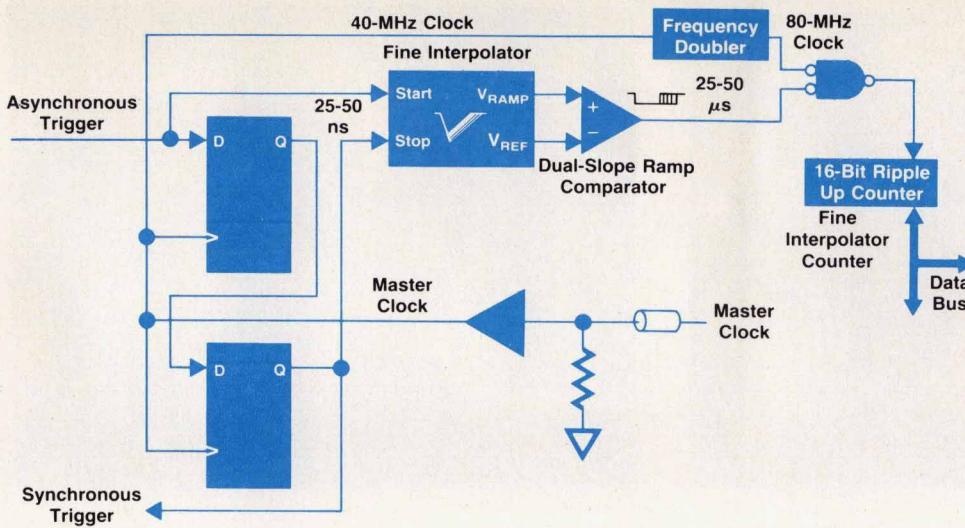


Fig. 3. Fine interpolator block diagram.

orders of magnitude higher than the sample rate. The rise time is limited by the bandwidth of a hardware preamplifier and sampler and not by the sample rate.

The advantages of accurate, stable, high-resolution timing and high bandwidth do not come free, however. There is one significant disadvantage of random repetitive sampling. If the sample rate is limited to 40 MHz as it is in the HP 54100A/D, then there exists a 25-ns hole between two data points taken on the same trigger. There is a finite probability that a time window narrower than 25 ns might not receive any data from a given acquisition cycle or trigger event. A 10-ns window will have a probability of only 10/25 of receiving one data point given one acquisition cycle, and a 1-ns window will have a probability of 1/25. Therefore, for a 1-ns window of interest, the signal must repeat 25 times on the average for each data point plotted on the screen. Because the data points are stored digitally and do not fade away, this effect is minimized. For a 1-ns time window (100 ps/div) and a signal repetition rate of 1000 Hz, the HP 54100A/D can still accumulate 40 digitally stored points per second, and the data rate goes up proportionally with signal repetition rate and time window width, reaching rates as high as 50,000 data points per second.

The realization of random repetitive sampling required the development of several new circuits. Fig. 2 illustrates

in greatly simplified form the basic components required. The circuits to be described in this article are the preamplifier/sync splitter, the sample-and-hold circuit, and the trigger interpolator. The programmable step attenuator is similar to the HP 33321G Attenuator. The trigger system is discussed in detail in the article on page 26. Two Siemens SDA 6020, 6-bit, 50-MHz flash ADCs are stacked to get seven-bit resolution at 40 MHz in the HP 54100A/D.

Fine Interpolator

The role of the fine interpolator is to provide the timing information that is necessary to reconstruct waveforms from multiple acquisitions of digitized data. The basic time base accuracy of the instrument is also set by the operation of this circuit. Fig. 3 shows a block diagram of the fine interpolator circuit. Its function is to measure very accurately the time between the trigger event and the next occurrence of the system master clock. Knowing the individual times allows reconstruction of random repetitively sampled waveforms from raw data.

Circuit operation begins with a block called the synchronous trigger generator. It has two inputs, the master clock and the asynchronous trigger coming from the trigger system discussed elsewhere in this issue. The synchronous

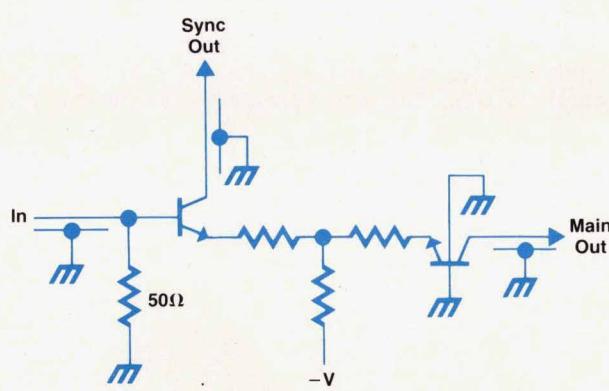


Fig. 4. Basic elements of the preamplifier and sync splitter.

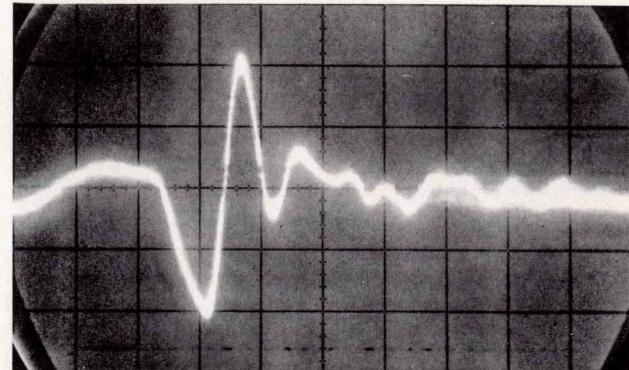


Fig. 5. Sampler kickout. Vertical: 5 mV/div. Horizontal: 200 ps/div.

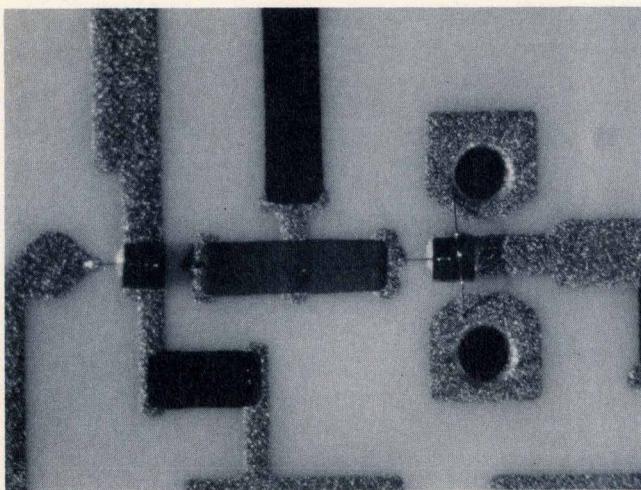


Fig. 6. Preamplifier thick-film hybrid substrate.

trigger generator consists of two high-speed ECL D flip-flops. The output of the circuit is another trigger signal that is coincident with the second occurrence of the master clock after the initial trigger is generated. The flip-flops are necessary to ensure that no false synchronous triggers are generated because of operation in a metastable region. Because of the random nature of the two inputs, the setup and hold specifications are by design often violated.

This output signal is referred to as the synchronous trigger. Because the asynchronous and synchronous triggers are related only by the master clock, the actual time difference is a random time between 25 ns and 50 ns, or one period of the 40-MHz master clock. This time is stretched by a factor of 1250 by the dual-slope analog-to-digital converter—the heart of the fine interpolator. This ADC consists of a high-speed differential comparator that switches a 10-mA reference current. The reference current discharges a capacitor for 25 to 50 ns. The capacitor is then charged from a current source 1250 times smaller than the discharge current. The ratio of the currents is the expansion ratio of the interpolator. The discharge is the fast ramp and the charge is the slow ramp of the dual-ramp circuit. The ramp is an input to a comparator that produces a logic output that is equal in time to the stretched or slow ramp time. The logic output enables a 16-bit counter that counts an 80-MHz clock derived from the master clock. This doubling in frequency is necessary to achieve 10-ps resolution for the system. The resultant digital word is then grouped with the digitized sampled data from the vertical system for reconstruction of the measured waveform.

Preamplifier

Typically the preamplifier in a conventional oscilloscope terminates the input signal and splits it into the main and sync paths. Isolation between the two paths is important, but is not a major constraint because the main path typically drives a terminated transmission line. For a random repetitive sampling system, however, the design constraint is different. The sampling process, by its very nature, generates a charge disturbance at the sampling bridge that propagates towards the preamplifier in the opposite direction

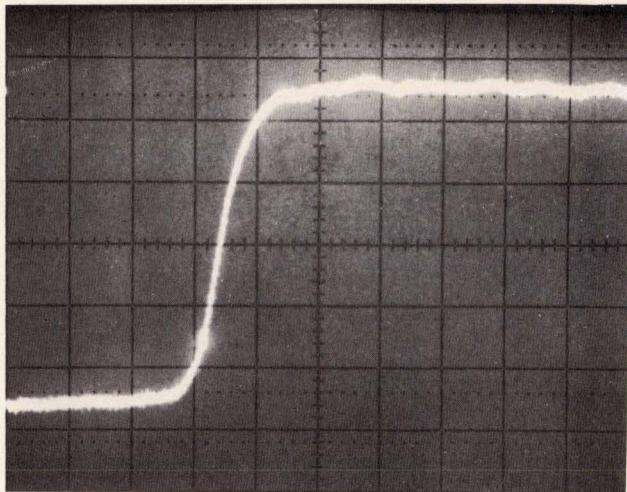


Fig. 7. Preamplifier step response. Vertical scale: 0.05V/div. Time scale: 200 ps/div.

from the input signal. Should the disturbance not be blocked by the preamplifier and couple into the sync path, it could cause a trigger event synchronous with the system master clock to occur. This would cause an erroneous trigger, or more likely, one that is statistically related to the master clock. The more the signals are coupled, the less random the acquisition of data.

Fig. 4 shows a schematic representation of the basic elements of the preamplifier and sync splitter. Because the circuit must pass the input signal with a bandwidth well in excess of 1 GHz, it is fabricated using a custom thick-film hybrid on a ceramic substrate. This is necessary to achieve fine geometries and low parasitic capacitance. The circuit is essentially a two-transistor differential amplifier with emitter resistors that reduce the gain and help stabilize the circuit. The first transistor splits the input signal into two parts—the main and sync signals. The second transistor is a common-base stage with its main function being to block any unwanted signal from the sampling bridge from coupling into the sync signal.

Fig. 5 shows a measurement made of "kickout" from the sampling bridge; this must be significantly attenuated before coupling into the sync signal. The design goal was to achieve 35-dB isolation between the preamplifier output and the sync output. The ingredients necessary to achieve this with the common-base stage were very low collector-to-emitter capacitance and near-zero base impedance to ground at 1 GHz. The latter was achieved with dual base

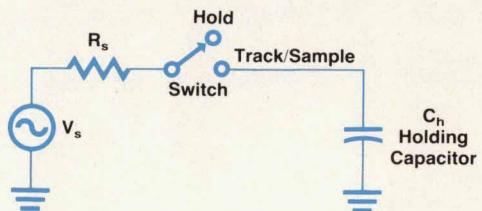


Fig. 8. Equivalent circuit of a track-and-hold or sample-and-hold circuit.

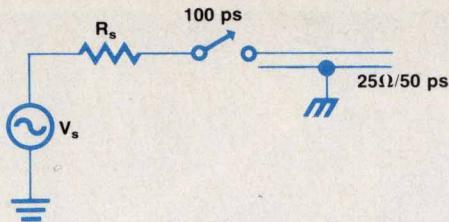


Fig. 9. Transmission line loading on a sampler.

bond wires that go directly to conductive vias (holes) near the die location. Total inductance at the base has been measured at 0.5 nH, which is sufficiently low to achieve the desired low impedance at that terminal. Fig. 6 is a photograph of the implementation on the preamplifier substrate. Fig. 7 shows the output response of the circuit measured at the main output with a 40-ps rise time stimulus. Note that the response has no overshoot and has a rise time of 200 ps.

Random Sampler

The key to extending the bandwidth of a digitizing oscilloscope without having to increase the digitizing rate is narrow-aperture sampling. Currently, there are two techniques used to achieve the narrow-aperture effect. The track-and-hold method is most commonly used at low frequencies where the required switch can be put into a feedback loop to reduce the nonlinear effects. The HP 54100A/D uses sample-and-hold methods because of the many advantages offered at high frequencies. The equivalent circuit used to illustrate these two methods is shown in Fig. 8. The circuit consists of a source voltage (V_s) to be measured and the associated back impedance (R_s) of the source, a switch, and a holding capacitor loading the switch. In a track-and-hold circuit, the switch is closed for relatively long times, and in this mode the voltage on the capacitor must track the source voltage. Then at a prescribed instant, the switch is opened. If the switch, being an electronic device (typically a diode bridge), changes state from low resistance to high resistance quickly, then a narrow aperture effect can be realized, that is, the voltage on the capacitor is held and is not affected much by the turn-off

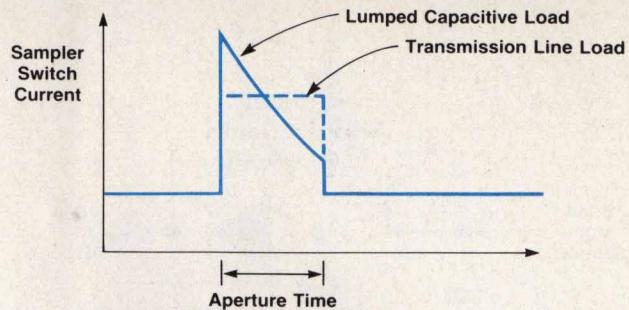


Fig. 10. Sampler switch current for lumped C or transmission line loading.

transient. However, in the track mode, the input signal is filtered through the source resistance and holding capacitance, limiting the bandwidth of the output signal. This filtering effect and the excess time used in tracking the signal limit the usefulness of track-and-hold methods at high frequencies.

Operating the same circuit in the sample-and-hold mode alleviates these two problems, but introduces more. A sample-and-hold circuit closes the switch only briefly (about 100 ps in the case of the HP 54100A/D), removing a small amount of charge from the source. This charge, when integrated on the hold capacitor, generates a voltage proportional to the average value of the input signal over the sampling aperture time (100 ps). There is typically a net voltage transfer loss with a sample-and-hold circuit because the short sampling aperture offers insufficient time for the voltage on the capacitor to achieve a steady-state value. The loss in voltage is called sampling efficiency. For example, if R_s is 25Ω, C_h is 2 pF, and the sampling aperture T_s is 100 ps, the sampling efficiency N_s , which is given by:

$$N_s = 1 - e^{-T_s/R_s C_h}$$

is about 86%.

In designing the HP 54100A/D, we found that the sampling efficiency could be improved by loading the sampling switch with a distributed capacitance or transmission line

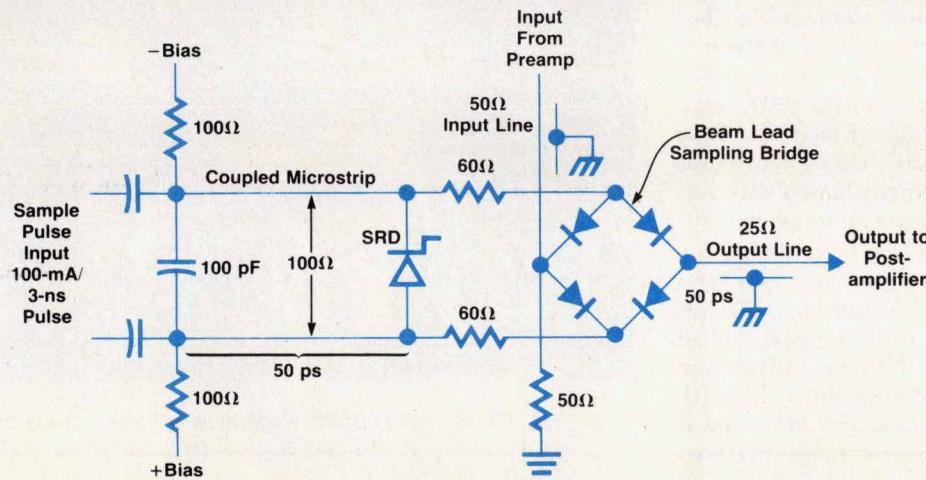


Fig. 11. Simplified schematic diagram of the sampling circuit.

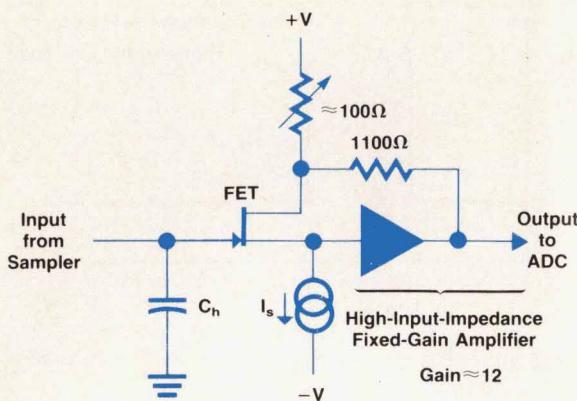


Fig. 12. Simplified postamplifier schematic diagram.

rather than a lumped capacitor. If, as shown in Fig. 9, we load the sampler with a 25Ω transmission line 50 ps long, then with a sampler aperture of 100 ps, the initial wavefront moving down the transmission line reflects off the end of the open line, returns to the switch just as the switch opens, and is trapped on the line. For this case, the sampling efficiency is exactly 100%. The efficiency can be made greater or less than 100% by making the transmission line load greater or less than 25Ω . The total capacitance of the transmission line load is $50 \text{ ps} \div 25\Omega$, or 2 pF , the same as the previous example, but we have increased the sampling efficiency by 14%. Another advantage of transmission line loading is greater uniformity of current flowing through the sampling switch during the aperture time. The initial current flowing into a lumped capacitor load is twice the current flowing into the matched 25Ω line, and decreases exponentially over the aperture duration while the current into the transmission line is constant (Fig. 10). The lower peak current leads to improved linearity for large signals.

Fig. 11 illustrates the relationship between the input circuitry, the output circuitry, and the sample strobe generation circuitry used in the HP 54100A/D. This circuit is realized as a thick-film microcircuit. The input travels down a 50Ω microstrip transmission line past the sampling bridge and is terminated in a 50Ω load. The sampling bridge taps the input line and connects to the output 25Ω microstrip line on the right. To the left of the input line is the sample strobe generation circuit. The primary source of the fast edge pulse used to turn the sampling bridge on is the step recovery diode (SRD). This diode generates a 10V, 50-ps-rise-time edge. The circuitry to the left of the SRD serves the dual purpose of matching the signal from the sample pulse generator to the SRD as well as shaping the 100-ps sample strobe pulse. The SRD is normally biased with current flowing from the positive bias supply through the SRD and out the negative bias supply. If the bias supplies are balanced, the positive and negative reverse bias of the sampling bridge is equal to half the SRD forward bias. To initiate sampling, a 100-mA, 3-ns-wide current pulse is applied from the sample pulse generator into the matching network with such a polarity as to reverse-bias the SRD.

When forward-biased, the SRD acts as a short circuit. It continues to act as a short circuit under reversed-bias cur-

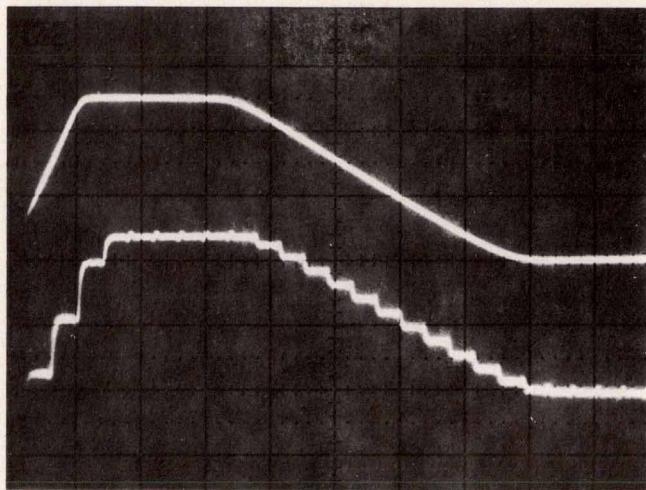


Fig. 13. A ramp-type input to the sampler and the output of the postamplifier held in steps 25 ns apart.

rent conditions until the stored charge in the diode goes to zero, at which time the diode quickly changes to an open circuit. In this case, there is sufficient bias current to allow the reverse current to reach about 100 mA before the diode opens. When the diode charge goes to zero, the 100-mA current is suddenly (<50 ps) interrupted, and a portion of the current is forced to flow through the 60Ω resistors into the sampling bridge. The rest of the current flows back into the 100Ω coupled microstrip structure where it is shorted by the 100-pF capacitor and is reflected back toward the SRD. When the wavefront reaches the SRD, it again forward-biases the SRD, causing the current flowing into the sampling bridge to turn off. The strobe time and thus the aperture time of the sampler are fixed at 100 ps by the two-pass length of the coupled microstrip. During this aperture time, the sampling bridge shorts the input line to the output line, allowing current to flow from one to the other. The amount of charge transferred into or out of the load capacitor is proportional to the difference between the

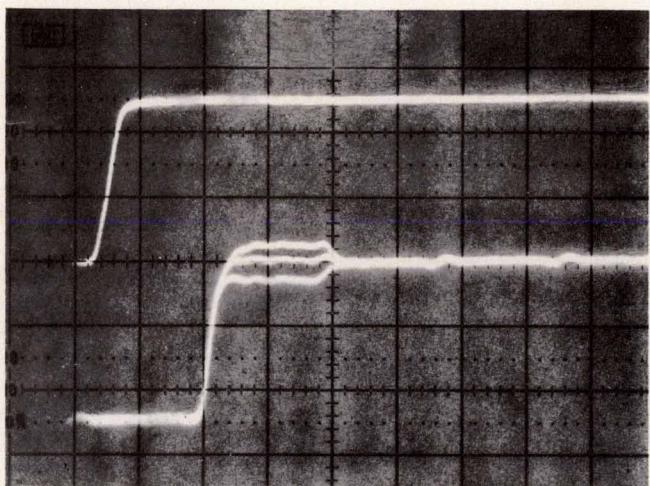


Fig. 14. Sampling efficiency adjustment. A step input to the sampler and the postamplifier output for three settings of the FET drain resistance.

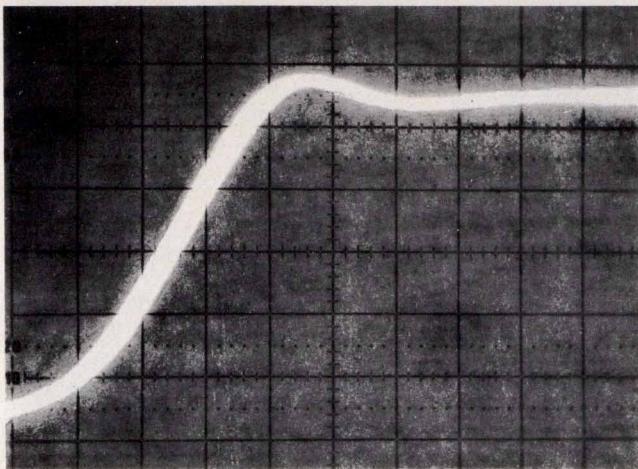


Fig. 15. Reconstructed sample step response has a rise time of 130 ps. Time scale: 50 ps/div.

input and output voltages. If the sampling efficiency is not exactly equal to 100%, then the first sample taken after a step change in the input signal will not be correct, but subsequent samples will measure the difference between the input and the output and reduce the error to zero. To minimize this error on the first sample, the sampling efficiency must be as close to 100% as possible.

The sampling efficiency depends on the load capacitance. If there were other parasitic capacitances in parallel with the load capacitor, such as the input capacitance of an amplifier needed to boost the output signal to a sufficient level to be digitized, then the sampling efficiency would not be 100%. Therefore, the postsampler amplifier must have zero input capacitance. In fact, if there are any vari-

ances in the load capacitance because of thick-film process variations, the input capacitance of the postamplifier might have to be negative to compensate for other parasitics. The design of the postsampler amplifier is therefore not a trivial matter, especially since its output must settle and be digitized at high rates (40 MHz in the HP 54100A/D).

Fig. 12 illustrates the postamplifier used in the HP 54100A/D. A source follower input stage is used with bootstrapping to reduce the effect of gate drain capacitance. In a normal case, the positive feedback to the drain is 100% (i.e., the drain voltage change is forced to be equal to the gate voltage change). Then the effective load capacitance of the FET is zero. The 100 Ω resistor in the drain circuit can be trimmed, allowing the $\pm 0.2\text{-pF}$ input capacitance to compensate for variances in the total holding capacitance according to:

$$C_{in} = C_{gd} (1 - A_{fb}),$$

where A_{fb} is the positive feedback gain and C_{gd} is the gate-to-drain capacitance of the FET.

If A_{fb} is greater than unity, the input capacitance of the amplifier can be negative. This does not cause a stability problem, however, because the holding capacitor makes the total capacitance of the node to ground positive.

Altogether the sampler serves to measure the instantaneous value of an input signal and hold this value to be digitized at a 40-MHz rate. Fig. 13 shows a ramp type input to the sampler, and the output of the postamplifier held in steps 25 ns apart.

Fig. 14 illustrates a step-like input to the sampler and the output of the postamplifier for three settings of the FET drain resistance. Fig. 15 is a reconstructed image of the sampler step response to an HP tunnel diode step generator. The rise time of this sampler is about 130 ps, indicating a bandwidth of about 2.7 GHz.

High-Performance Probe System for a 1-GHz Digitizing Oscilloscope

by Kenneth Rush, William H. Escovitz, and Arnold S. Berger

ATYPICAL SYSTEM to be measured by an oscilloscope usually contains not just one class of signal, but several classes. Making signals from these different classes available to the measuring oscilloscope is the goal of the probing system for the HP 54100A/D Digitizing Oscilloscope.

The probing system consists of the following products:

- The HP 54001A 1-GHz Miniature Active Probe
- The HP 54002A 50 Ω Input Pod
- The HP 54003A 1-M Ω Probe
- The HP 54300A Probe Multiplexer.

Circuits that are sensitive to capacitive loading are usually fast circuits having low resistances. For this class, the probe capacitance dominates the design trade-offs. Fast circuits, say ECL 100K, also demand a high-bandwidth probe to preserve the shape and timing of the signal sent to the oscilloscope. The HP 54001A Miniature Active Probe is the optimum probe for this class of signals. This probe's pole-zero cancellation technique results in an extremely small probe tip with low capacitance (2 pF), medium resistance (10 k Ω), and extremely high bandwidth (1 GHz). In size and bandwidth, the HP 54001A is an improvement

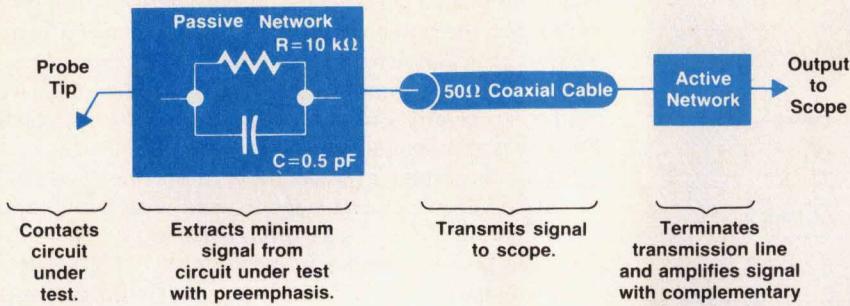


Fig. 1. Block diagram of the HP 54001A Miniature Active Probe.

over previous active probes.

A large class of signals exists in a 50Ω environment of coaxial connectors. The requirements of 50Ω terminated coaxial signals cannot be met easily by other probing solutions, so the HP 54002A was designed to cover this class. The HP 54002A offers low insertion loss and good termination.

Circuits that are sensitive to resistive loading (i.e., having resistances above a kilohm) are usually slow circuits and are not so sensitive to capacitive loading. A probe can be optimized for this class of circuits by keeping the probe resistance high at the expense of probe capacitance. Such is the case with the HP 54003A. The HP 54003A probe offers equivalent performance to conventional oscilloscope probes, high resistance ($1\text{ M}\Omega$), moderate capacitance (8 pF), and relatively high bandwidth (300 MHz) for requirements where resistive loading is more important than capacitive loading.

The idea that allows these products to make a significant contribution is a system of interchangeable pods and a method of multiplexing the outputs of the pods into an oscilloscope. The HP 54300A Probe Multiplexer allows these probes to form a measurement solution. Now signals from several classes of circuits can be brought together and measured in one system. Low-speed, high-resistance oper-

ational-amplifier signals can be measured by the HP 54003A and multiplexed with high-speed ECL signals measured by the HP 54001A. Amplifier outputs in 50Ω coax can now be measured in the same system as the low-power CMOS used to program the amplifier gain.

Miniature Active Probe

Careful consideration of noise, bandwidth, and loading effects led to some unusual design concepts in the HP 54001A. The single most significant problem in designing an oscilloscope probe is dealing with the cable capacitance. The cable capacitance in a conventional probe is part of a capacitive voltage divider. Hence, the higher the cable capacitance, the higher the probe tip capacitance must be. Historically, designers have used several methods to minimize cable capacitance. One common method is to make the cable short, since capacitance is proportional to length. Unfortunately, this leads to probes that are often too short to bridge the gap between the circuit under test and the oscilloscope. Another method is to make the cable have very high impedance. If we could realize a 150Ω cable, it would have one-third the capacitance of a 50Ω cable. Unfortunately, high-impedance coaxial cable tends to have a large outer diameter. Carrying both of these techniques to their extremes—maintaining sufficient length to be useful, and maintaining a small diameter cable—yields the HP 54003A, which is similar to the HP 10017A passive probe, having a tip capacitance of about 8 pF .

Other techniques have been used to reduce tip capacitance, but they usually have severe drawbacks. First, plac-

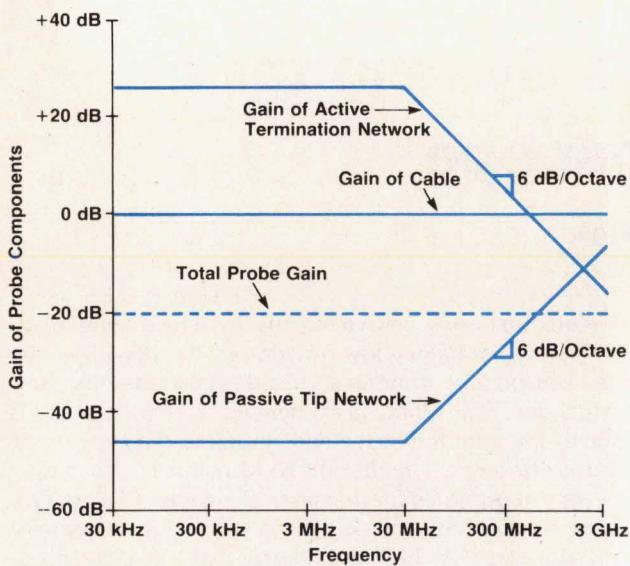


Fig. 2. Gain of HP 54001A probe components versus frequency.

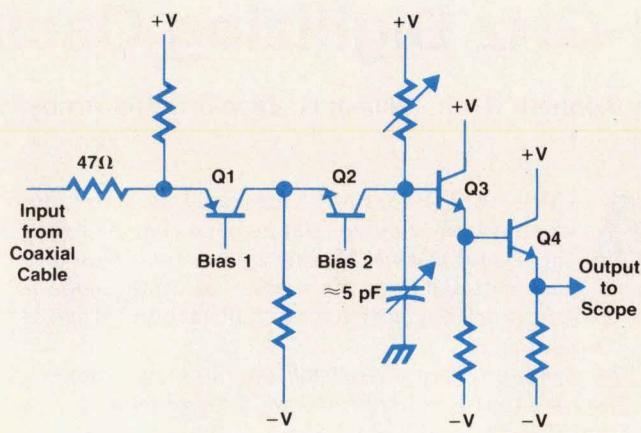


Fig. 3. Simplified schematic of the active network in the HP 54001A Miniature Active Probe.

ing an amplifier near the probe tip to drive the cable capacitance eliminates the tip impedance's dependence on cable properties, but power must be supplied to the amplifier, and it takes up space, making the probe tip large. Second, the cable could be considered as a transmission line and terminated as such as in resistive divider probes. However, this leads to probes with very low resistance at the probe tip— 500Ω for a 10:1, 50Ω probe such as the HP 10020A.

Obviously a new approach was needed to produce the ideal probe for high-speed logic. To achieve higher bandwidth in a longer probe requires low-loss, high-quality coaxial cable terminated in a 50Ω impedance. The small tip size restraint requires a passive tip. The requirement of high resistance forces us to introduce gain into the system to terminate the cable in 50Ω . One approach might be to use a resistive divider probe with a high division ratio, followed by an amplifier to make up the gain loss. This poses two problems. First, gain at very high frequencies (>300 MHz) is not easy to get, and if we could get it, the gain would add noise to the system. However, if we are willing to remove more signal from the circuit under test using slight capacitive loading, we will not be attenuating the signal as much at high frequencies. Thus, less gain is required. This is the method used in the HP 54001A probe.

Fig. 1 illustrates the concept in a block diagram. In the probe tip, a thick-film microcircuit is used to realize an RC passive network. The resistance is a laser-trimmed, thick-film resistor. The capacitance is the front-to-back capacitance of a couple of metal traces on the ceramic substrate and has a value of 0.5 pF .

Fig. 2 illustrates the transmission gain of the passive tip network and the active network gain required to get a total probe gain of -20 dB (10:1). Note the transmission zero caused by the parallel action of the resistance and capaci-

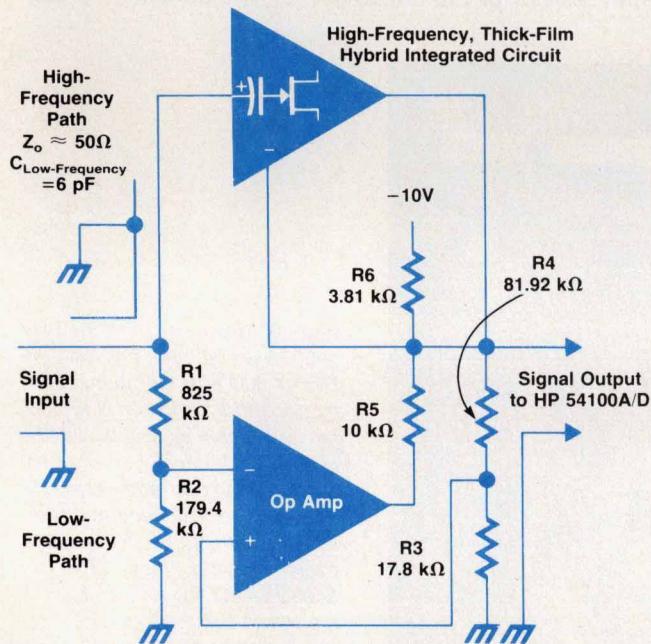


Fig. 4. Block diagram of the HP 54003A 1- $M\Omega$ Probe input pod.

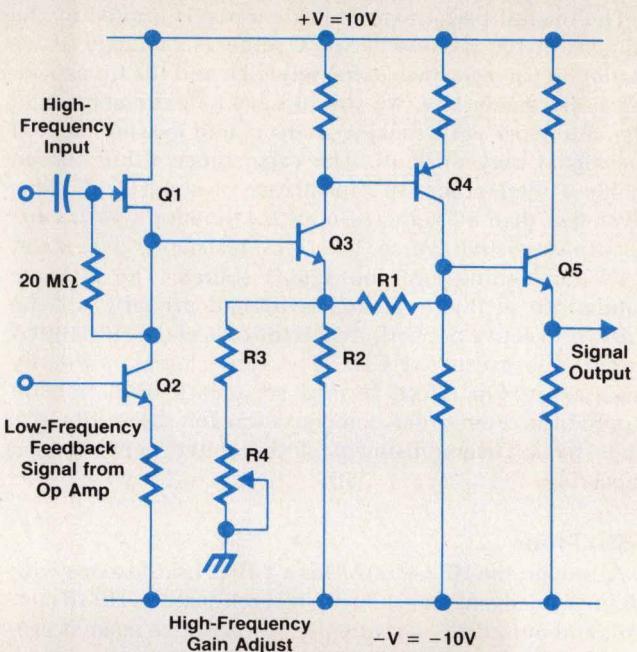


Fig. 5. Basic circuit of the high-frequency thick-film hybrid for the HP 54003A 1- $M\Omega$ Probe input pod.

tance in the probe tip. This transmission zero must be cancelled in the active network with a pole at 30 MHz. Below 30 MHz, the passive network has a loss of 46 dB, and the active network must have a gain of 26 dB while the coaxial cable has no loss or gain. Above 30 MHz, the gain of the passive network increases at 6 dB/octave, while the active network gain decreases at 6 dB/octave, yielding a constant -20 -dB gain over the frequency range of the probe. The required gain of the active network goes through 0 dB at about 630 MHz.

The requirements on the active network are rather severe. It must terminate the coaxial cable in its characteristic impedance of 50Ω , boost the signal by 26 dB below 30 MHz, roll off its gain at -6 dB/octave to well past the desired bandwidth of the probe, and drive the 50Ω input of the oscilloscope to a reasonable dynamic range of ± 2 V, giving a 20 V dynamic range at the probe tip (10:1 probe). Thick-film hybrid microcircuit technology is used to realize the circuit.

Four discrete microwave transistors are used in cascade. Fig. 3 illustrates the circuit. Transistor Q1 and the 47Ω resistor in series with its emitter terminate the coaxial cable in its characteristic impedance. The base bias of Q1 is automatically adjusted with an operational-amplifier feedback circuit to guarantee that the input offset is zero volts. Q2 acts as a buffer to isolate the high impedance at its collector from the low impedance of the active network input. At the collector of Q2 is an RC network that generates the required current-to-voltage conversion of the signal current passing through Q1 and Q2. The resistive and capacitive values are trimmed at calibration to guarantee a match with the passive network to achieve the proper gain at all frequencies. Q3 and Q4 form a cascade of emitter followers to isolate the high-impedance node at the base of Q3 from the low impedance of the output node.

The overall performance of this probe is limited by the choices of the passive network values. Obviously, if we made the tip resistance extremely high and the tip capacitance extremely low, we would have an extremely small signal to work with, making recovery and measurement of the signal very difficult. The capacitance value chosen yields a total probe tip capacitance of about 2 pF. This gives less than a 3% increase in rise time for a 50Ω, 1-ns-rise-time resistive source. The 10-kΩ resistance gives about 0.5% dc loading error on a 50Ω source. The ultimate bandwidth of this approach is limited primarily by the thick-film active network, and in the case of the HP 54001A probe, this exceeds 1 GHz. The cable losses, primarily caused by skin effect, form a secondary limitation on bandwidth. First-order compensation for these losses is done by an overadjustment of the active network trim capacitor.

1-MΩ Probe

Although the HP 54100A/D is a 1-GHz oscilloscope with 50 internal signal paths, its digital architecture, HP-IB control, and automatic measurement capabilities make it useful at much lower frequencies. To take advantage of this, the HP 54003A 1-MΩ Probe was designed. It is a unity-gain, general-purpose, 300-MHz input buffer amplifier with a 1-MΩ, 10-pF input. It requires an external HP oscilloscope probe. Probes for it include the HP 10017A 1-MΩ, 8-pF miniprobe, the HP 10014A 10-MΩ, 10-pF standard probe, and the HP 10032A 100:1 Probe. Its bandwidth of 300 MHz and its rise time of 1.2 ns are specified both at the BNC input and at the probe tip. It is useful for circuits ranging from high-impedance operational amplifiers to ECL logic circuits to power supplies up to 200V.

Most oscilloscope amplifiers are preceded by attenuators, which limit the dynamic range that the amplifiers must cover. The HP 54100A/D has 50Ω, 1-GHz attenuators

already built into it, so the HP 54003A buffer amplifier must have a large dynamic range. The dynamic range is $\pm 2V$, which covers most digital circuit and operational-amplifier applications with a $\times 10$ probe. This is about twice the dynamic range of buffer circuits in active probes and oscilloscope inputs.

As shown in Fig. 4, the HP 54003A is a two-path amplifier. Signal components less than about 1 MHz pass through the low-frequency path, an operational-amplifier circuit, then to a low-frequency input of the high-frequency amplifier. The operational amplifier also controls the operating point of the high-frequency amplifier.

Signal components higher than 20 Hz also follow the high-frequency path. There is sufficient overlap of the two circuits' passbands to ensure a smooth total response. Because of the closed-loop design and modern operational-amplifier stability, there is almost no visible drift. Also, the high open-loop gain of operational amplifiers means that 0.1% resistors R1, R2, R3, and R4 accurately and repeatably set the low-frequency gain to:

$$A = \{R2/(R1 + R2)\} \{(R3 + R4)/R3\}.$$

One adjustment in the high-frequency amplifier sets the high-frequency gain equal to the low-frequency gain so that the overall response is flat. The high-frequency circuit (Fig. 5) is a unity-gain amplifier with input, gain, and output sections. Because the input and output sections are followers with gain less than 1, a gain section is necessary. Transistors Q1 and Q2 make up the input section. FET source follower Q1 is operated at gate-to-source voltage $V_{gs} = 0$, maintained by the 20-MΩ gate-to-source resistor. However, the Miller effect between the gate and source causes the input impedance to be much greater than that resistance. The input resistance does not significantly load the 1-MΩ input resistor of the low-frequency operational-amplifier

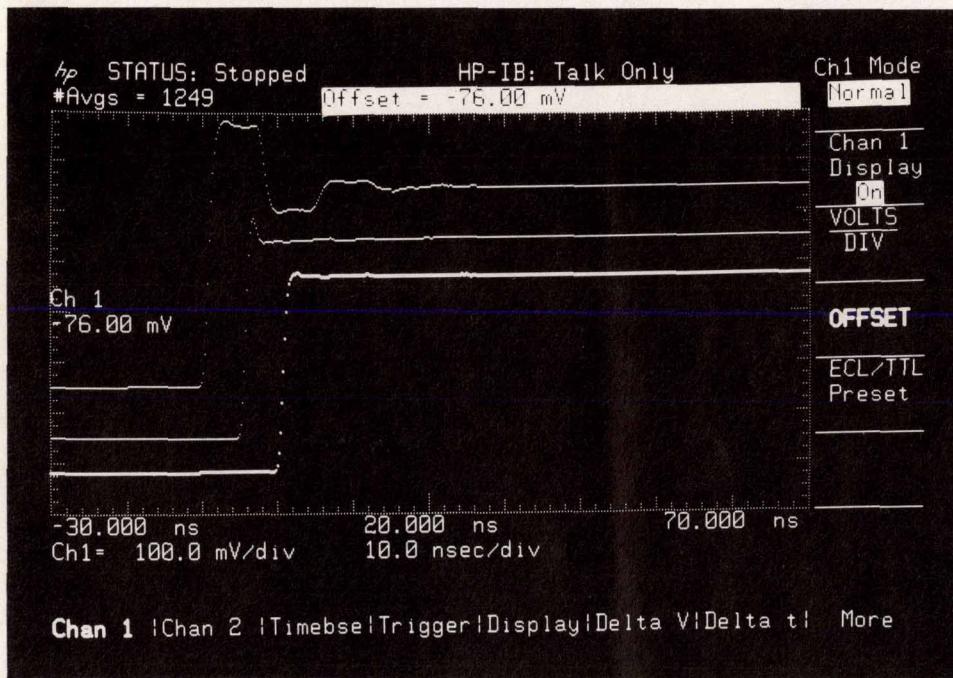


Fig. 6. Response of the HP 54003A probe input pod with the HP 10100C 50Ω feedthrough terminator and artificially long 600-mm 50Ω cable as a stub, with the HP 10100C directly on the HP 54003A input pod, and with no HP 10100C. The 50Ω feedthrough terminator with stub creates 33% overshoot (top trace). The HP 54003A input line as a stub (middle trace) causes less overshoot because it is fairly short for this bandwidth. Proper use of the HP 54003A is without feedthrough terminator (bottom trace).

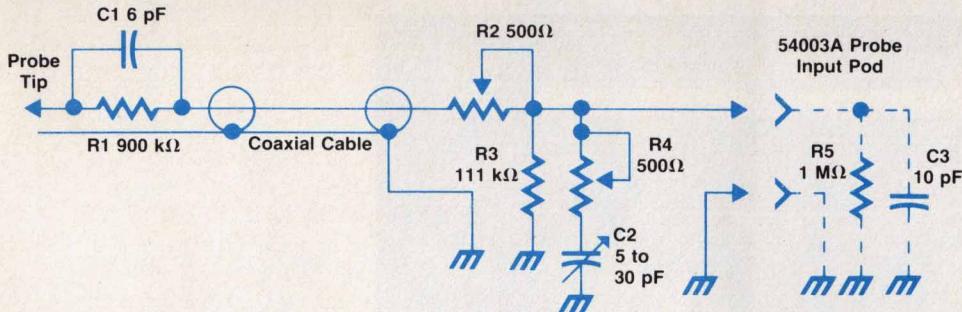


Fig. 7. Schematic diagram of the HP 54003-61617 1-MΩ, 8-pF, 10:1 oscilloscope probe (same as the HP 10017A).

circuit. The feedback input from the operational-amplifier circuit maintains the proper current in Q1 and Q2 for FET operation at $V_{gs} = 0$ regardless of the input bias level within the $\pm 2V$ dynamic range and regardless of individual FET variations.

The gain stage is the common-emitter/common-emitter circuit of Q3 and Q4 with voltage series feedback. Resistors R1 through R4 make up the feedback network. R4 is adjusted so that the high-frequency amplifier gain is equal to the low-frequency operational-amplifier circuit gain.

The high-frequency circuit is implemented as a thick-film hybrid. The technology is described in the article on page 33. The short length of chip-to-substrate wire bonds and repeatable geometry of traces minimize variable parasitic impedances at 300 MHz. This produces consistently high-performance circuits.

The HP 54003A input should be driven from a 50Ω source or oscilloscope probe such as the HP 54003-61617 (HP 10017A). Although the input capacitance is about 10 pF, the BNC input connector and the input line look like a 50Ω transmission line terminating at the input capacitance of the high-frequency amplifier, which is only about 2 or 3 pF. Thus, a 50Ω source can drive the input well beyond 300 MHz. It is common to terminate high-impedance inputs with 50Ω feedthrough terminators, such as the HP 10100C, to reduce reflections in high-frequency measurements. However, this causes overshoot with fast rise time signals because of the inevitable stub of BNC connector and buffer amplifier input line. For a long stub, the overshoot is 33%. Instead of a feedthrough termination, it is better to use a $\times 5$ or $\times 10 50\Omega$ attenuator at the BNC input if the signals are large enough. If the source impedance is close enough to 50Ω , no termination or attenuator is necessary. Fig. 6 shows how the waveform varies for three different termina-

tions. Only analysis of individual cases reveals the best termination, if any, to use with high-impedance inputs. Of course, the HP 54002A 50Ω Input Pod is best for dedicated use in a 50Ω system.

Most high-impedance inputs are used with probes to characterize or troubleshoot circuits. Probes represent the highest-resistance, lowest-capacitance, lowest-disturbance connection to circuits that do not have buffered 50Ω test points. The main probe for the HP 54003A is the HP 54003-61617, which is identical to the HP 10017A Miniature Divider Probe (Fig. 7). It is designed to drive 1-MΩ inputs with approximately 10 pF input capacitance. Up to about 10 MHz, it is a capacitor-compensated 10:1 resistive voltage divider. The divider consists of the parallel combination of R1 and C1 in series with the parallel combination of R3, R4, R5, C2, C3, and the cable capacitance. The R2 and R4 network terminates the characteristic impedance of the cable at very high frequencies, where transmission line effects in the cable become important. The center conductor of the cable is resistive to damp out remaining reflections.

1-MΩ Probe Performance

Fig. 8 is a block diagram of some of the equipment used to characterize the high-frequency performance of the HP 54003A. It produces flat pulses of approximately 1-ns rise time, typical of the fastest pulses the HP 54003A can usefully measure. The pulse flattener is a Schottky diode series clipper. When the HP 8082A output is negative, it is connected to the pulse flattener output through the diode, which is turned on. When the HP 8082A output returns to 0V, it is connected to the flattener output only through the diode capacitance, about 1 pF. The output of the flattener returns rapidly to 0V through the 50Ω output resistor. The small capacitance of the turned-off diode filters the pulse-

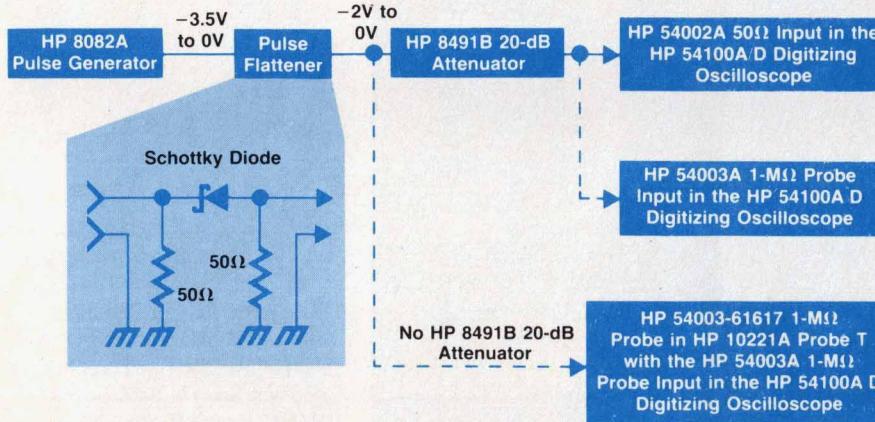
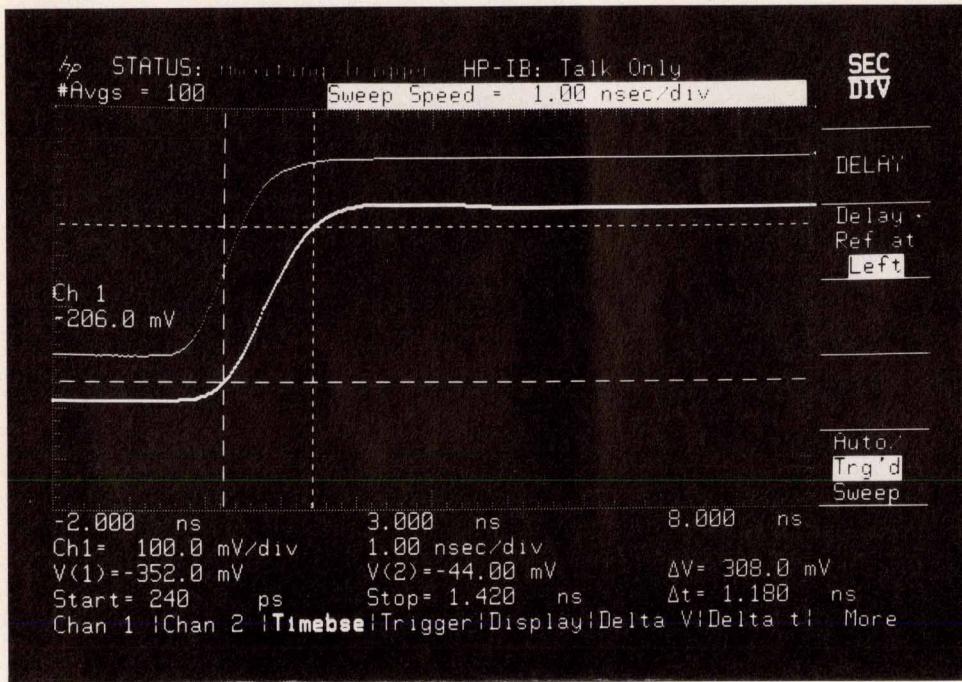


Fig. 8. Equipment to produce flat pulses for characterizing oscilloscopes and probes.



top aberrations of the pulse generator, which are already small, and the 50Ω output resistor reduces reflections when the diode is off. The 20-dB attenuator is a nearly 50Ω load for the pulse flattener regardless of the device under test. The HP 8082A's amplitude and offset are adjusted to produce the flattest reference pulse near the pulse transition.

Fig. 9 shows the response of the HP 54100A/D with the HP 54002A 50Ω input (top trace, the reference pulse) and the HP 54003A $1-M\Omega$ input (bottom trace). The reference trace has a 1.02-ns rise time and aberrations of about 1% peak-to-peak for the first 10 ns, which require the HP 54100A/D's magnify mode to see. The HP 54003A trace

has a rise time of 1.18 ns and about 2% overshoot compared to the reference pulse.

Fig. 10 compares the performance of the HP 54003A and the HP 54003-61617 $\times 10$ miniprobe (bottom trace) with the reference trace. No attenuator is used because the probing tee is already terminated in 50Ω . The rise time with the probe is 1.09 ns and the perturbations are about 3% peak-to-peak. The 1.09-ns rise time of the probe compared to the 1.02-ns rise time of the reference waveform or the 1.18-ns rise time of the HP 54003A itself does not mean that the $1-M\Omega$ probe is exceptionally fast, but only that it can be tuned using its adjustable termination resistors.

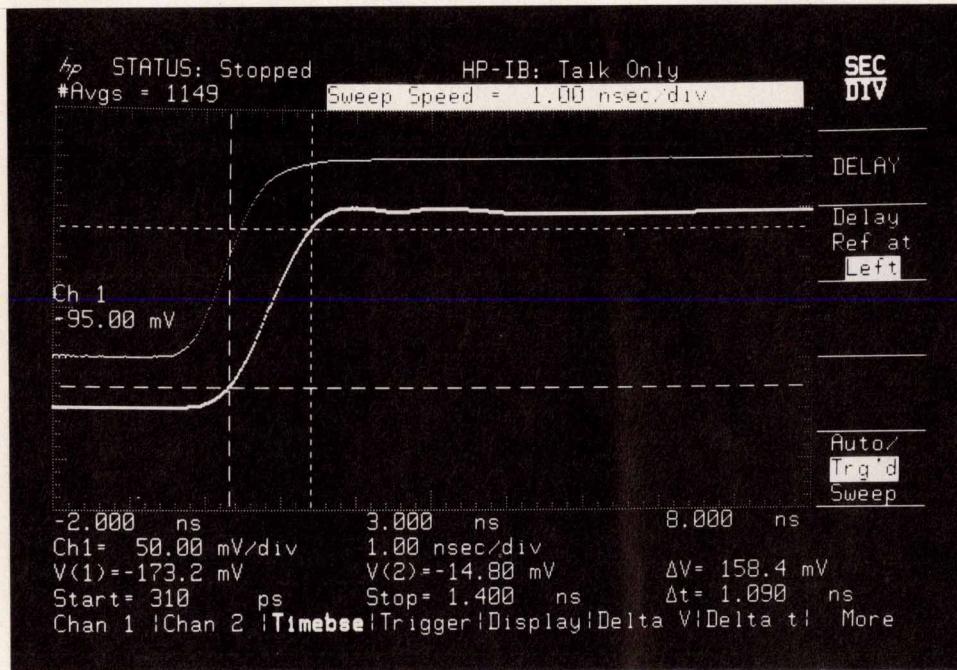


Fig. 9. Response of the HP 54100A/D with the HP 54002A 50Ω probe to a 1-ns rise time reference pulse (top trace), compared with the response of the HP 54003A $1-M\Omega$ input pod (bottom trace).

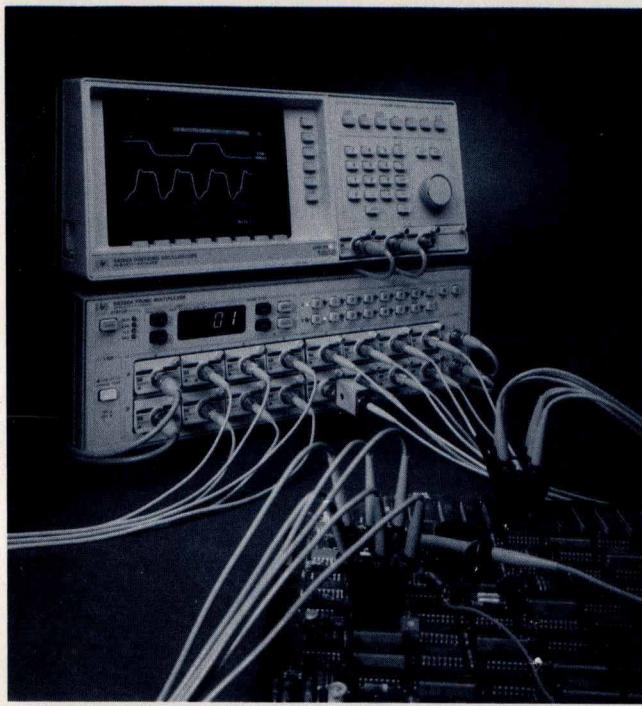


Fig. 11. The HP 54300A Probe Multiplexer, shown with the HP 54100A Digitizing Oscilloscope.

Probe Multiplexer

As previously discussed in this article, the difficulties associated with probing a circuit increase as the circuit's density and frequency limits increase. One manifestation of this problem is probing multiple nodes in a complex system. Quite often, the mechanical fixturing needed for probing points of interest in a circuit must be designed into the circuit itself. Thus, part of an overall probing strategy must address the need for a greater number of probes than can be accommodated by an instrument's input.

The HP 54100A/D can accommodate four HP 54000-series probes as inputs. Two of its inputs are dedicated to the vertical inputs, as in a traditional oscilloscope, and two are devoted to the time base trigger channel. The HP 54300A Probe Multiplexer is designed to expand the inputs of the HP 54100A/D or other 50Ω instruments. It is configured as two independent, eight-to-one, analog input channels. Each channel can be thought of as an eight-position rotary switch, capable of switching 50Ω transmission line signals with a signal bandwidth greater than 1.5 GHz.

Fig. 11 shows the HP 54300A with a full complement of probes. Any of the HP 54000-series probes can be interchanged in any slot position, thus allowing up to 16 points in the circuit to be probed without having to remove and reconnect the probe tips. Once the signal exits the probe, it is in a well-controlled 50Ω environment. The active channel is routed through an eight-channel RF coaxial switch to a BNC connector on the front panel.

Multiplexer System Electronics

Fig. 12 is a block diagram of the HP 54300A. The heart

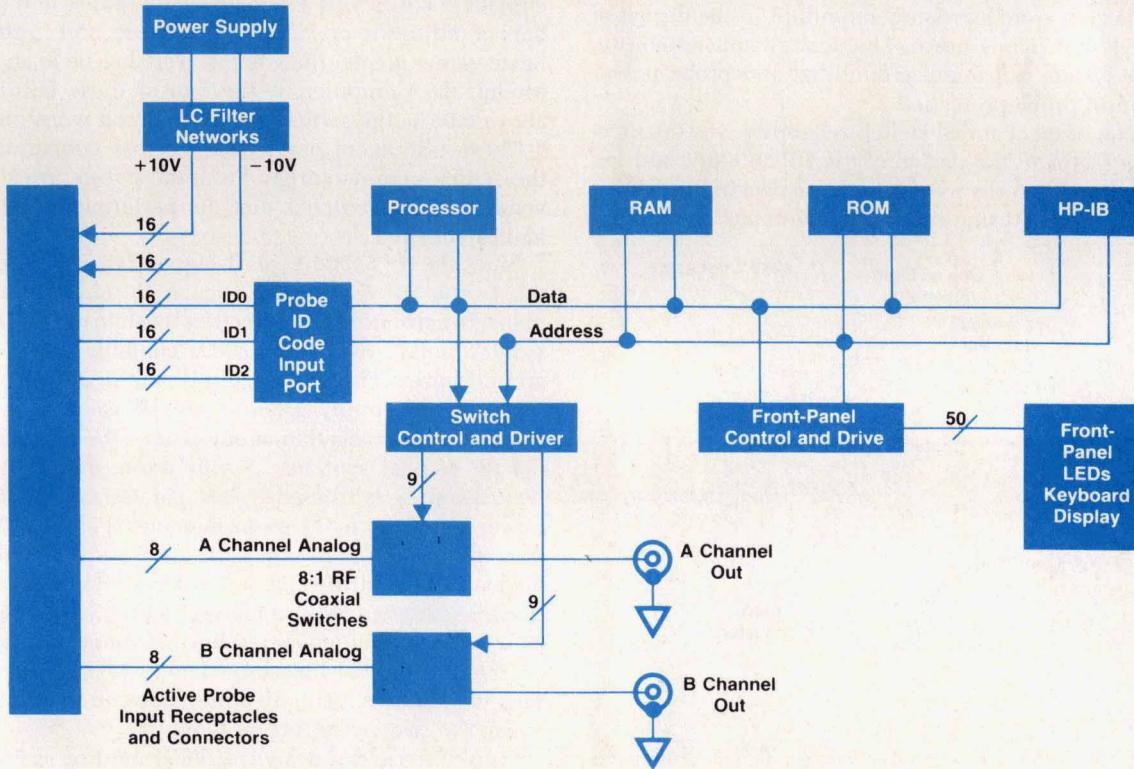


Fig. 12. Block diagram of the HP 54300A Probe Multiplexer control circuitry.

of the instrument is two eight-channel, RF coaxial switches that have bandwidths from dc to greater than 2 GHz. The switches are solenoid-actuated. Each solenoid is connected to a 50Ω stripline cable that can make a connection between the central contact of the switch and a contact around the perimeter. Actuating the solenoid with 24Vdc connects the associated 50Ω peripheral connector to the center one. Only one solenoid may be actuated at a time. The RF switch design is shown in Fig. 13.

The instrument is controlled by a 6809A microprocessor operating with a 1-MHz clock frequency. The entire operating firmware is contained in single $8K \times 8$ EPROM circuit. The remainder of the memory consists of 2K bytes of scratchpad memory, used for power-on system storage, and 4K bytes of battery-backed-up nonvolatile memory, which is used to store sequences of switch closures (lists).

The HP-IB interface, front-panel switches and display, and RF coaxial switches are controlled by LSI controller circuits.

The power supply is a linear design, capable of supplying 5V at 2A for the digital circuitry, $\pm 10V$ at 2A for active probe power, and 28Vdc at 500 mA for the solenoids. The 5V and $\pm 10V$ supplies are regulated by self-contained IC regulators. The 28V supply is essentially unregulated. However, to limit power dissipation within the solenoids, an LM317 regulator is used as a 30V clamp. At voltages less than 30Vdc, the LM317 looks like a series diode with a forward voltage drop of approximately 1V. When the input voltage rises above 30Vdc, the regulator begins to regulate, effectively clamping the voltage seen by the solenoids at 30Vdc.

Multiplexer Design Considerations

Three factors were extremely important in the design of the HP 54300A. They are mechanical switch reliability, electrical system noise and grounding, and probe power integrity and probe protection.

The issue of mechanical switch reliability was the most important factor in the design of the HP 54300A, and the most development time went into this aspect of the design. Although manufacturing RF coaxial switches is a mature

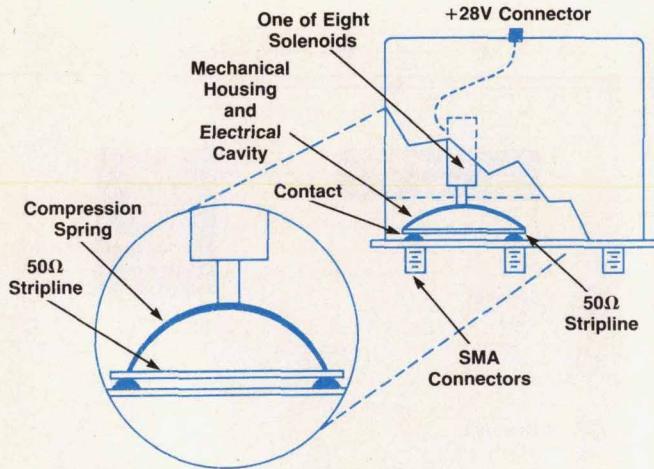


Fig. 13. Cutaway view of the multiplexer RF coaxial switches. The mechanical housing and 50Ω stripline form a wide-bandwidth, low-loss transmission environment.

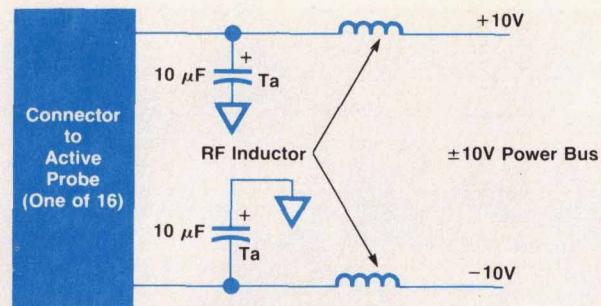


Fig. 14. The tantalum capacitors and RF inductors provide high-frequency power supply filtering and prevent cross talk between active probes.

technology, we had to be certain that the switches could achieve the performance level that our customers expect. With the help of marketing, we surveyed potential HP 54300A users to determine the limits of acceptable performance. This survey provided us with two performance benchmarks. First, the switches had to be able to perform reliably for 700,000 to 1,000,000 operations per switch contact, and second, the rate of soft failures, that is, partial closures or incomplete openings, should be less than 0.05% per contact.

We instituted an extensive test program to evaluate RF switch performance. The switches were tested by cycling each switch through all eight contact positions while ramping the temperature continuously from -20°C through $+60^\circ\text{C}$. After 200,000 cycles, the switch was removed from the environmental chamber and tested for RF performance and dc resistance. The dc test was the most stringent. Each contact position was exercised 10,000 times and the numbers of failures to open, failures to close, and contact resistance values greater than $50\text{ m}\Omega$ were logged to an HP 9000 Model 226 Computer. If the rate of these failures went above 0.05%, the switch was considered worn out.

These tests were performed in close cooperation with the switch manufacturers. From these tests, we selected a vendor whose switches met the performance criteria we had established.

Since the HP 54300A was designed as a companion product to the HP 54100A/D, a major design challenge was managing grounding and electrical system noise. To control system noise, the HP 54300A contains three separate ground buses. The analog supply, the digital supply, and the solenoid supply grounds are all isolated from each other and come together at one point—the flange of one of the RF coaxial switches. Furthermore, the $\pm 10V$ analog power supply is filtered by a separate LC circuit at each power connector to the probe housings (Fig. 14). Whenever a signal must cross between the analog and digital ground planes, it is sent through a $100\text{-k}\Omega$ series resistance that straddles the two ground planes on the inner layers of the main printed circuit board in the instrument. These signals are then buffered by high-impedance CMOS logic circuitry. This technique effectively eliminates currents flowing between the two grounding systems.

Even with careful design of the grounding system, early tests revealed noise pulses that were correlated with the

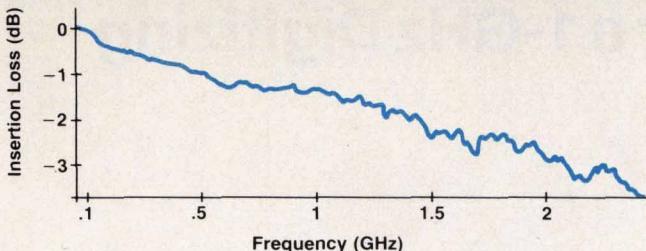


Fig. 15. Plot of insertion loss versus frequency for the HP 54300A Probe Multiplexer. The overall shape of the curve is the result of skin-effect losses. The fine structure is caused by small impedance discontinuities at internal connections.

1-MHz system clock emanating from the 6809A microprocessor. Since all external inputs to the microprocessor are interrupt driven, once an external (front panel or HP-IB) command is executed, the microprocessor need only wait until the next interrupt occurs. The SYNC instruction was used to put the 6809A in a wait mode. In this mode, all address and data bus activity ceases, and the microprocessor halts all instruction processing until an interrupt occurs. Thus, the digital system is only active when executing switch closure or system instructions. This effectively eliminates all coupled noise in the system.

As mentioned earlier, each connector to the probe housing has a separate LC filter at the $\pm 10V$ inputs. This helps eliminate cross talk between probes and provides clean power. The power supply is a linear design, which at worst case is only 40% efficient. The use of a linear rather than a switching power supply design was decided upon since the potential for noise coupling problems was much higher with a switching power supply. Since the HP 54300A chassis is rather roomy and a fan is used to cool the active probe circuitry, the heat developed by the power supply is not a problem in the system design.

The potential for a power supply failure destroying all the active probes in the instrument is handled by two separate circuits. The IC regulators are internally current-limited to a safe level to prevent catastrophic failure of the supply. In addition, overvoltage protection ICs, located downstream from the voltage regulator ICs, separately monitor the plus and minus 10V lines. Should either of the supplies go above 12Vdc, the circuits activate an SCR, which short-circuits the line and blows a protection fuse.

The analog voltages are monitored by a comparator circuit that can generate a processor interrupt if a power supply failure occurs.

Multiplexer Operational Features

Normal operation of the probe multiplexer is quite straightforward. The pushbuttons **A0** through **A7** and **B0** through **B7** toggle the corresponding switch inputs on and off. An LED lights to show the active channel. In addition, scanning sequences (lists) can be created and stored in nonvolatile memory. These lists are similar to those used in the HP 3488A Switch Control unit. A list can be up to 99 steps long and up to 100 lists may be created. The total number of stored steps may not exceed 3761, the capacity of the nonvolatile RAM array. A list may be recalled from memory and edited and/or executed. **NEXT** and **PREV** keys allow

sequencing through a list. If these keys are held down, the instrument will move rapidly through the list, but will not activate a channel until the key is released. Where practical, the HP-IB command set mimics that of the HP 3488A. The front panel also contains a four-digit LED display. The two left-hand digits display the currently active list, and the two right-hand digits display the current step in a list. The display is also used to display error messages in the format: **Er nn**, where **nn** is an error message number.

Lists can also be scanned by pulsing an external, rear-panel BNC input connector. A positive TTL pulse causes the current list to be scanned forward one entry. The instrument also issues a pulse from another rear-panel BNC output connector when the switch contact has settled and data can be taken. Thus, a data logger scheme can be implemented without the intervention of a computer/controller. In this mode of operation, the end of a list rolls around to the beginning as the next step.

Using the Multiplexer

The HP 54300A is easily connected to other test instruments. Its BNC outputs are 50Ω transmission lines and can be connected to any instrument that has 50Ω input terminations. When used with an HP 54002A 50Ω Input Pod, the HP 54300A has an input-to-output bandwidth greater than 1.5 GHz (3 dB). When connected to an HP 54100A/D, the system bandwidth is still greater than 1 GHz.

Since the analog signal path is essentially that of a 50Ω transmission line, the signal losses that occur through the instrument are those that would naturally occur in any cabling and connection system of equal length. Fig. 15 is a graph of the swept frequency response of the instrument. The losses are not the usual 6 dB per octave of a single pole, but rather skin-effect losses normally associated with transmission lines.

In theory, very large multiplexing systems could be built by cascading multiple HP 54300A instruments. However, because of the increased cabling and connections, the signal losses would become excessive. Therefore, we recommend that no more than one additional level of HP 54300As be cascaded, giving a total of 128 input channels. Cascading one additional level reduces the multiplexer's bandwidth to 1 GHz (3 dB), and reduces the system bandwidth of the HP 54100A/D to less than 1 GHz, but this may be a reasonable compromise in applications requiring many input channels.

Acknowledgments

The authors would like to acknowledge the contributions of the following people: Rick James and Frank Leri developed the mechanical design of the pods and pod housing. Martin Guth designed the thick-film microcircuits used in the HP 54001A probe tip and amplifier. Lewis Dove designed the thick-film microcircuits for the HP 54003A amplifier, and Ken Miller aided in the design and test of the amplifier circuit. Tom Cheski did the quality assurance testing of the RF coaxial switches. Bill G. Smith did the final mechanical design of the HP 54300A. Jonathan Mahaffey put the entire firmware for the HP 54300A into one 8K ROM. Don Smith guided the development of a number of disjointed projects with aggressive schedules and kept them all on track.

Waveform Graphics for a 1-GHz Digitizing Oscilloscope

by Rodney T. Schlater

WAVEFORM GRAPHICS are an important part of an oscilloscope. By looking at a displayed waveform, a user can understand more about a signal than if it is described by several pages of data. For example, at a glance a user can tell whether the rise time is too slow, whether the signal has ringing on it, whether the lower level is floating, or whether the signal is misbehaving.

Recognizing the importance of the display, special attention was paid to developing a good graphics interface for the HP 54100A/D Digitizing Oscilloscope. Several new methods of displaying waveforms were developed that provide users greater insight into their circuits. In addition, new techniques were created for plotting waveforms on the screen very quickly, making the HP 54100A/D very responsive and giving it the feel of an analog oscilloscope.

The three display modes of the HP 54100A/D are variable persistence, infinite persistence, and averaging. Variable persistence and infinite persistence are new to digitizing oscilloscopes. The averaging used is a continuous average. Most digitizing oscilloscopes now use a terminating type of average that stops when the selected number of repetitions have been averaged.

Variable Persistence

Variable persistence is the normal display mode of the HP 54100A/D. With variable persistence, the user can con-

trol how long waveforms remain on the screen. The waveform persistence can be set from 0.2 second to 10 seconds.

For short persistence times, variable persistence behaves like the normal display mode of an analog oscilloscope. The screen is continuously updated with new waveforms, and the picture on the screen changes as the input signal changes. For longer persistence times, variable persistence can convey more information than a conventional oscilloscope.

In addition to displaying amplitude versus time, variable persistence gives the user a view of what the waveform is doing over a long period of time (e.g., jittering, drifting, or changing in amplitude). In analog oscilloscopes, variable persistence has proved valuable for capturing single-shot and infrequently occurring signals. Because it is implemented using different techniques, the HP 54100A/D's variable persistence solves many of the problems that are inherent in analog techniques: the persistence time can be set precisely, the picture is always clear and easy to adjust, and the display does not fade or bloom.

Variable persistence overcomes one of the major shortcomings of digitizing oscilloscopes— inability to handle dynamic waveforms. Most digitizing oscilloscopes work well when the waveform is static and well-behaved, but have problems with waveforms that are changing, such as multiple-trace waveforms, envelope waveforms, and

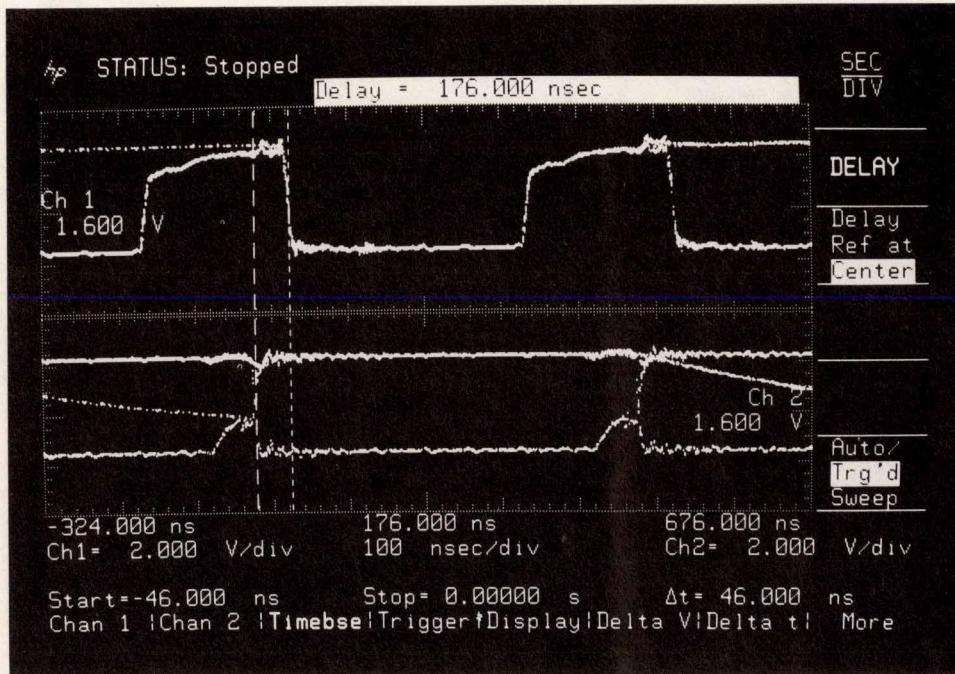


Fig. 1. With variable persistence, digitizing oscilloscopes can display waveforms that are multi-valued in time. The two vertical markers in the photo are positioned to measure the setup time between the data strobe (top waveform) and an address line (bottom waveform) of a 68000 microprocessor.

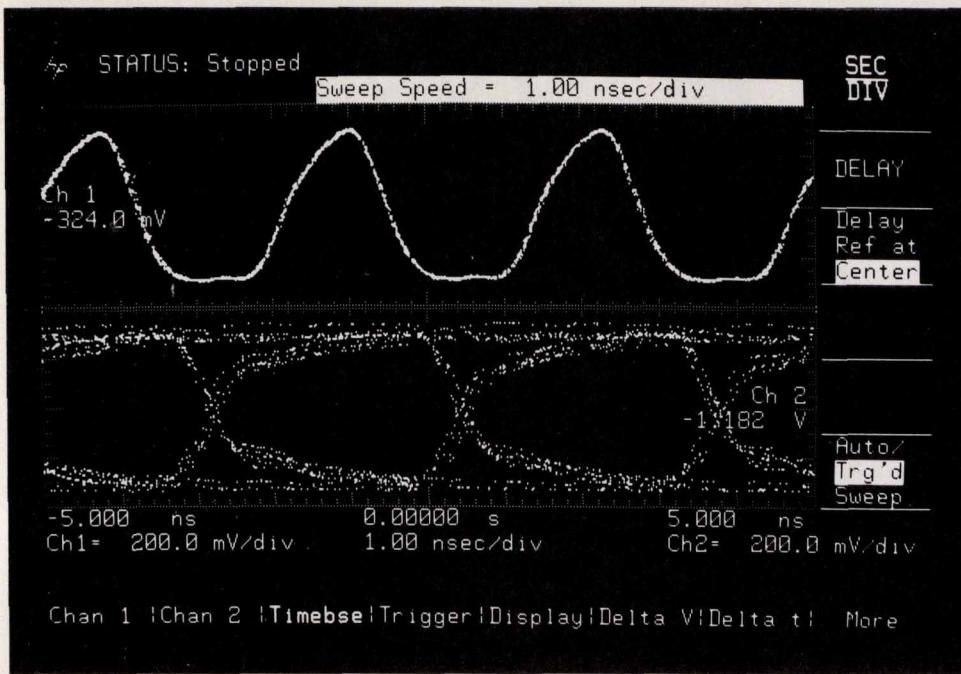


Fig. 2. Variable persistence shows the different waveform modes in an eye diagram. The quality of the transmission line can be determined by measuring the width and height of the eye opening.

waveforms that are aliased, jittering, or noisy. This is because they treat all waveforms as single-valued functions (i.e., having one vertical point for every time point), whereas many waveforms are multivalued over time.

Variable persistence handles changing waveforms very easily. Fig. 1 shows an example of a multiple-trace waveform displayed with the variable persistence set to 0.5 second. A multiple-trace waveform occurs when a signal has several modes depending on the circuit conditions. Sometimes it is possible to display only one of the waveform modes by selecting an appropriate trigger signal, but such a signal is not always available. In many cases, it is desirable to display the signal as a multiple-trace waveform. In the example, the multiple-trace waveform is used to measure the setup time between a 68000 microprocessor data strobe and an address line.

Fig. 2 shows another example of a multiple-trace signal, one that is both noisy and jittering in time. The display shows the eye diagram of an output signal from a long transmission. The quality of the transmission line can be determined by measuring the height and width of the eye pattern.

Infinite Persistence

The maximum setting on the variable persistence control knob selects the infinite persistence mode. In this mode, waveforms accumulate on the display indefinitely or until the user presses the **CLEAR DISPLAY** key. This erases the screen and restarts the process.

Infinite persistence is a feature that is new to both digital and analog oscilloscopes. The maximum persistence time in analog storage oscilloscopes has been limited to between 10 and 60 seconds in the viewing mode.

Infinite persistence is valuable because it can catch all the variations of a signal. It is useful for making worst-case measurements, such as jitter, noise, or metastable measure-

ments, and for trapping circuit problems that occur infrequently, even those that occur every few days.

Fig. 3 shows an example of an infrequently occurring problem that was solved using infinite persistence. The problem involved a two-stage ECL counter circuit. The problem only occurred about once every two seconds and turned out to be a 2-ns negative-going glitch on the carry output of the counters. In Fig. 3, the top waveform is the trigger signal and the bottom waveform is the carry output from the counter. The problem also illustrates the HP 54100A/D's ability to look back in time at events that occur before the trigger.

Averaging

The averaging display mode is useful for removing uncorrelated noise from signals and for increasing the resolution of displayed waveforms. Averaging is performed on the samples that fall within a time slot. Features that distinguish the HP 54100A/D's method are that the averaging is performed very quickly and that it is done using a continuous averaging algorithm, which causes the screen to be constantly updated with the newest averaged waveform. Other digitizing oscilloscopes that perform averaging generally use a terminating method in which the averaging stops after a selected number of sweeps.

Two front-panel controls affect averaging. The **CLEAR DISPLAY** key erases the screen and restarts the averaging process, and there is a control for selecting the number of averages, which can be a value between 1 and 2048. There is also a display readout that shows the current number of averages.

Fig. 4 shows an example of how averaging can be used to recover a signal buried in noise. The top waveform shows the original signal, and the bottom waveform shows the signal after averaging 128 sweeps. The top waveform was acquired in the variable persistence mode and then was

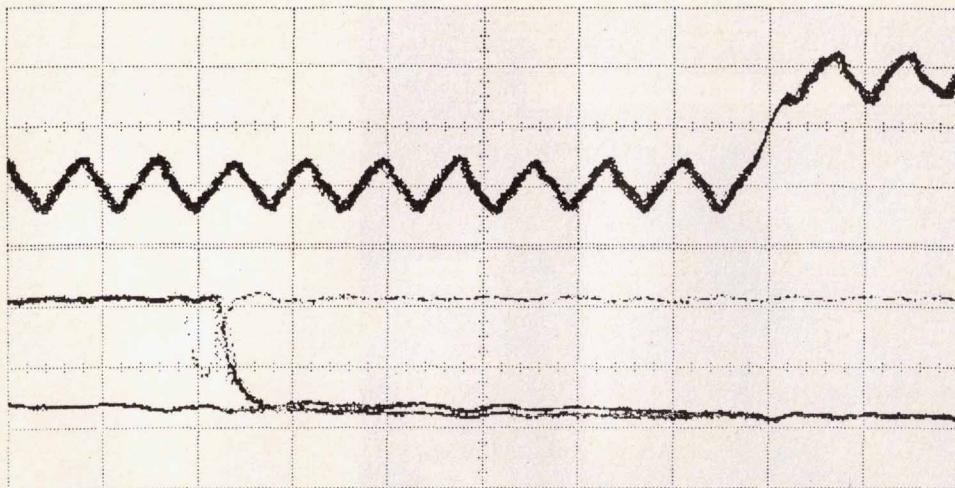


Fig. 3. An ECL counter was not reaching full count. The cause of the problem, a glitch on the carry output, was discovered using infinite persistence. The picture was obtained using the hard-copy output capability of the HP 54100A/D.

saved in the waveform memory.

The HP 54100A/D uses a seven-bit analog-to-digital converter (ADC). Averaging extends the resolution to ten bits. Normally, not all of the extra resolution is required, since the resolution of the screen is only eight bits. However, there are several places where the extra resolution is of value. One is for time measurements; more accurate results can be obtained using the extra precision. A second is for HP-IB applications; HP-IB users can read out full-precision values over the bus. A third is for display magnification.

The HP 54100A/D has a vertical magnification feature that is useful for looking at distortions on waveforms such as ringing and reflections. This feature works by placing a window over the part of the waveform that is to be magnified, and then turning the magnification on. Fig. 5 shows an example of the top of a pulse that is magnified by a factor of 16. The top waveform was acquired in the variable persistence mode. The ADC quantization levels show up

as lines across the screen. The bottom waveform was acquired in the averaging mode. With averaging, enough additional bits of resolution are provided to fill in the gaps between quantization levels.

Graphics System

Fig. 6 shows a simplified block diagram of the HP 54100A/D. The first block is the input circuitry and the sampler. The input circuitry contains attenuation and offset circuits to bring the signal within the range of the sampler and ADC. The sampler performs a sample-and-hold function. The sample window is narrow enough to ensure a 1-GHz bandwidth.

The sampled data is then passed to the ADC, which converts each sample to a seven-bit digital number. These values are stored in a high-speed memory; 1024 points are stored on each acquisition cycle. However, the number of points that reach the display screen depends on the time-

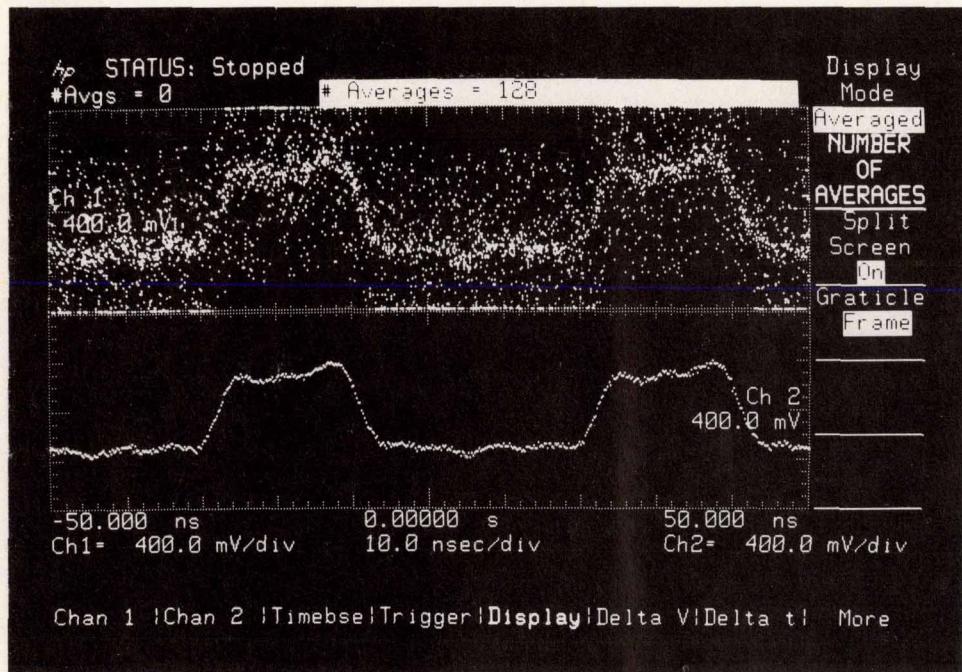


Fig. 4. The top waveform shows a signal that is buried in noise. The bottom waveform shows the signal after averaging.

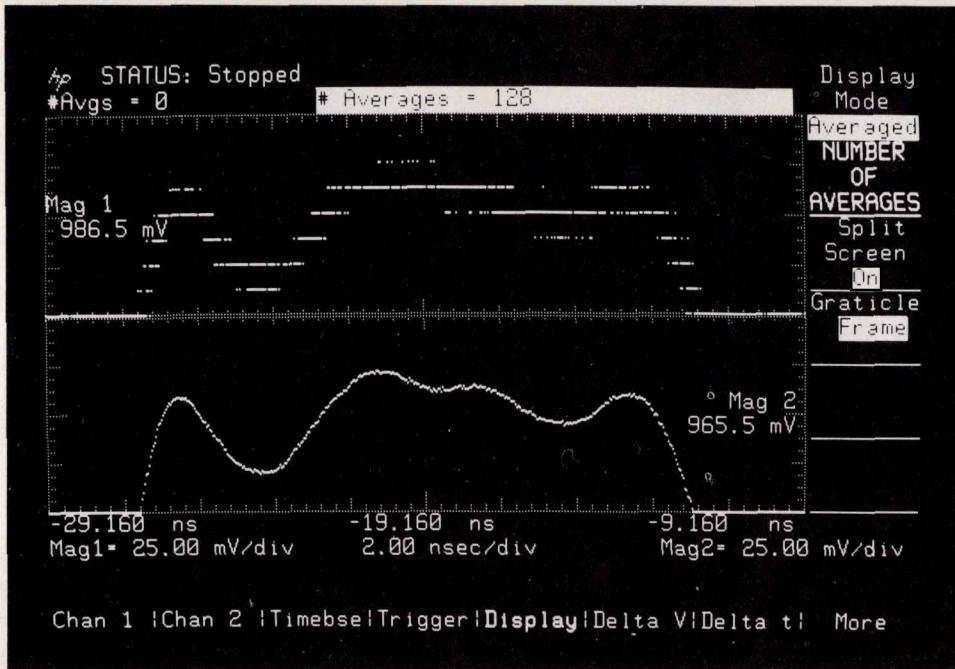


Fig. 5. The photo shows the top of a pulse magnified by a factor of 16. In the variable persistence mode (top waveform), the analog-to-digital quantization values show up as levels. In the averaging mode (bottom waveform), the points between the levels are filled in.

per-division setting.

When an acquisition cycle is complete, the 68000 microprocessor reads the digitized values from the high-speed memory, processes the points, and stores the points as bits in the waveform pixel memory. Every 1/60 second, the display circuitry reads the waveform memory and displays the points on the screen.

The HP 54100A/D uses two memory planes, one for the graticule and one for the live waveforms. Using two memory planes has several advantages. First of all, the graticule or static display is separated from the waveform or live display. This method results in a faster screen update rate, because the software does not have to take the graticule into account when the waveforms are displayed and erased. Second, the waveforms and graticule can be displayed at two intensities, half-bright and full-bright. This makes it easy for the user to distinguish the waveforms from the graticule.

The instrument uses a raster display. The screen is 576 pixels wide by 368 pixels high. The waveforms and graticule are drawn in the graphics area, which is 501 pixels wide by 256 pixels high. The area outside the graphics area is used for text.

The waveform and graticule memories are seen as RAM by the 68000 microprocessor. Each bit in the graphics memory corresponds to a pixel on the screen. When the microprocessor sets a bit in the memory, it shows up as a dot on the screen.

One of the early decisions in the project concerned the choice of a display. Two alternatives were considered; one was a raster display and the other was a vector display, which is traditionally used in an oscilloscope. The raster display was chosen primarily because a raster display's memory is a natural image memory for retaining infinite persistence waveforms. Once a dot is set in the display memory, it stays set until the microprocessor clears it.

Importance of Screen Update Rate

One of the design goals for the HP 54100A/D was to put waveforms on the screen at a very fast rate. The screen update rate has an important effect on the performance of a digital oscilloscope for several reasons.

First, it is necessary to update waveforms at least ten times per second to give the user the feeling that the oscilloscope is operating in real time and is responsive to changing conditions on the input signal. Generally, the higher the update rate, the more responsive the oscilloscope will be.

Second, the screen update rate affects the dead time, which is the time the oscilloscope spends in processing the waveforms rather than waiting for triggers. Reducing the dead time reduces the chances of missing a trigger. Normally, missing a trigger is not important unless it occurs on a waveform that contains an infrequent error for which the user is looking.

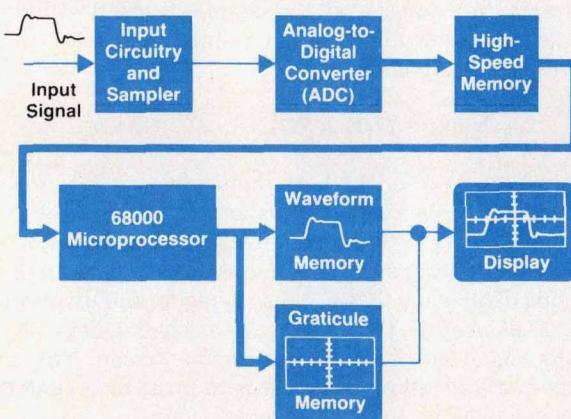


Fig. 6. Simplified block diagram of the HP 54100A/D Digitizing Oscilloscope.

Third, increasing the update rate increases the perceived bandwidth of the signal being displayed. This is because a random repetitive oscilloscope achieves a high bandwidth by sampling the same signal many times. The more times a signal is sampled within a given period, the higher the bandwidth will be. An example is a narrow glitch or vertical spike. The probability of catching the glitch is related to the width of the glitch versus the sample time. Sampling the glitch more often increases the chances that it will be detected.

Last, the screen update rate is important if a user is looking for an infrequent problem. Increasing the update rate decreases the time it takes for the problem to be displayed.

One of the features that sets the HP 54100A/D apart from all other digitizing oscilloscopes is its very high update rate. Up to 20,000 points are put on the screen each second in the variable persistence mode, and up to 50,000 points are put on the screen each second in the infinite persistence mode. The next section describes the plotting algorithm that is used to achieve such a high screen update rate.

High-Speed Plotting Algorithm

The design of the high-speed plotting routine eliminates floating-point calculations by two means:

- A lookup table is used to convert the digital sample values to Y values on the screen. The table is recomputed only when one of the magnification settings is changed.
- A 32-bit integer is used for the X coordinate. The upper 16 bits represent the X value in pixels and the lower 16 bits are used for accuracy.

A Pascal version of the high-speed plotting routine is shown below. The routine is actually coded in assembly language and consists of 17 statements, five of which are for plotting the point directly on the screen.

Integer Variables (Pascal Program)

```
X = X value in pixels, number between 0 and 500
Y = Y value in pixels, number between 0 and 255
A_to_d_data = sample value, 0 to 127
Number_of_points = number of points to be plotted
Long_x = 32-bit integer
16 MSBs = X in pixels
```

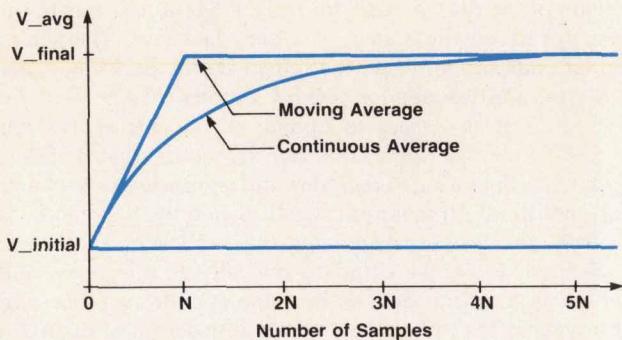


Fig. 7. Graph showing the behavior of a moving average and a continuous average. At time zero the input signal is changed from $V_{initial}$ to V_{final} . The graph shows the variation of the voltage with the number of samples.

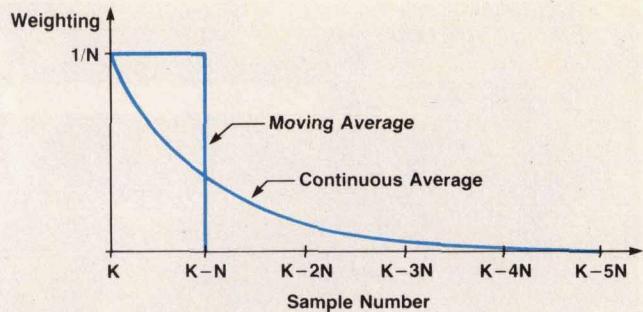


Fig. 8. Graph showing the weighting of samples over time for a moving average and a continuous average.

16 LSBs = accuracy bits
Long_delta_x = step size (same format as Long_x)

Plotting Routine (in Pascal)

```
For I := 1 TO Number_of_points DO
BEGIN
  X := Long_x/65536; {Compute X in pixels}
  Long_x := Long_x + Long_delta_x; {Add delta for next point}
  Y := Vertical_table(A_to_d_data[I]); {Look up Y value in table}
  Plot (X,Y); {Plot point on screen}
END;
```

The display routine takes 18 μ s to plot one point. Up to fifty 1000-point waveforms can be acquired and plotted in one second in the infinite-persistence mode.

Averaging Algorithm

The HP 54100A/D uses a continuous averaging algorithm, which is also referred to as an exponential average. The algorithm has been used in other applications, but there has been very little written about it, especially as it applies to oscilloscopes. The algorithm divides the waveform record into 501 time slots. The samples that fall within a time slot are averaged together.

Several other methods of averaging were considered. One method is the terminating average. With the terminating average, the averaging stops when the required number of samples has been acquired. The average for each time slot is simply the average of all the samples that fell within the time slot. Mathematically the terminating average can be expressed as follows:

$$V_{avg} = (V(1) + V(2) + \dots + V(N)) / N \quad (1)$$

where $V(1), V(2), \dots, V(N)$ are the sampled values for the time slot, N is the selected number of samples to be averaged, and V_{avg} is the average value for the time slot.

The terminating average is not appropriate for an oscilloscope display mode, since the averaging and display update stops after the final average is reached. Later changes in the signal are not reflected on the screen. The only method of updating the display is to press the **CLEAR DISPLAY** key, which restarts the averaging process. It is inconvenient for users to have to press the **CLEAR DISPLAY** key when signals are changing or when probing different sig-

nals. (Note: the HP 54100A/D does use a terminating average for data acquisitions over the HP-IB.)

The ideal continuous averaging method is the moving average. With the moving average, new samples are continuously acquired. The average for each slot is performed on the last N samples that fall in the time slot. The equation for the moving average can be expressed as follows:

$$V_{avg}(K) = (V(K+1-N) + V(K+2-N) + \dots + V(K))/N \quad (2)$$

where $V(K+1-N)$, $V(K+2-N)$, ..., $V(K)$ are the last N samples for the time slot, and $V_{avg}(K)$ is the Kth average.

The disadvantage of the moving average is the amount of memory required to implement it. The values for the last N samples have to be saved for each time slot. For example, for two channels, 501 time slots, and one byte of storage for each time slot, the amount of memory required to perform a moving average of eight samples is 8016 bytes ($2 \times 501 \times 1 \times 8$). For a moving average of 256 samples, the amount of memory required is 256,512 bytes.

The HP 54100A/D continuous averaging algorithm uses a total of 8,016 bytes for the two channels. Each average is computed by first multiplying the previous average by $(N-1)/N$ and then adding $1/N$ times the new sample. The equation is:

$$\begin{aligned} V_{avg}(K) &= ((N-1)/N)V_{avg}(K-1) + (1/N)V(K) \quad K \geq N \\ V_{avg}(K) &= ((K-1)/K)V_{avg}(K-1) + (1/K)V(K) \quad K < N \end{aligned} \quad (3)$$

where $V(K)$ is the current sample, N is the selected number of averages, $V_{avg}(K-1)$ is the previous average for the time slot, and $V_{avg}(K)$ is the current average for the time slot.

To maintain accuracy and perform the algorithm quickly, an alternate version of the equation is used. For every sample, a running sum, $Sum(K)$, is also computed.

$$Sum(K) = Sum(K-1) - V_{avg}(K-1) + 256V(K) \quad (4)$$

$$V_{avg}(K) = Sum(K) / N \quad (5)$$

The value $V(K)$ is multiplied by a factor of 256 (i.e., it is shifted eight bits) to maintain accuracy. This makes the lower eight bits in $Sum(K)$ and $V_{avg}(K)$ accuracy bits. Eight bytes are saved for every time slot: $Sum(K)$, which is four bytes, $V_{avg}(K)$, which is two bytes, and the average count, N, which is also two bytes.

Analysis of the Continuous Averaging Algorithm

Looking at the equation (3) for the continuous average, it is not obvious how the continuous average compares to a true average or moving average. For values of K less than the selected number of averages, N, the equation reduces to a true average. Substituting for $V_{avg}(K-1)$ gives the following results:

$$V_{avg}(K) = \frac{((K-1)/K)(V(1) + V(2) + \dots + V(K-1))}{(K-1)} + (1/K)V(K) \quad (6)$$

$$= (V(1) + V(2) + \dots + V(K))/K \quad (7)$$

The equation can be further analyzed by determining the effects of a change in signal level on the average. Assume that the input signal is at a steady-state level of $V_{initial}$ and is changed to a level V_{final} . Also assume that K is greater than or equal to N, so the number of averages is constant. After one sample, the average is:

$$V_{avg}(1) = ((N-1)/N)V_{initial} + (V_{final}/N) \quad (8)$$

After two samples, the average is:

$$V_{avg}(2) = V_{initial}((N-1)/N)^2 + (V_{final}/N)((N-1)/N + 1) \quad (9)$$

The first term grows as a power of n, where n is the number of samples, and the second term is a geometric series times (V_{final}/N) . After n samples the equation reduces to:

$$\begin{aligned} V_{avg}(n) &= V_{final} - (V_{final} - V_{initial})((N-1)/N)^n \\ &\approx V_{final} - (V_{final} - V_{initial})\exp(-n/(N-0.5)) \end{aligned} \quad (10)$$

A graph of this equation is shown in Fig. 7, along with a graph for the moving average. Initially, when the input signal is changed from $V_{initial}$ to V_{final} , both the moving average and the continuous average change at about the same rate. However, the moving average takes only N samples to reach the final voltage whereas the continuous average approaches V_{final} exponentially and takes between 3N and 5N samples to reach the final voltage.

For noisy steady-state signals, the continuous average and the moving average behave almost identically. The main difference is in how successive samples are weighted (Fig. 8). For the moving average, the last N samples contribute $1/N$ to the average. For the continuous average, the most recent sample contributes $1/N$ to the average. The effect of each sample lasts longer since the contributions decrease exponentially.

Another way of looking at the continuous average is to think of the averaging as a filtering process; the continuous average is, in fact, a low-pass digital filter. For example, compare equation (10) with the equation for the step response of an RC network.

$$V(t) = V_{final} - (V_{final} - V_{initial})\exp(-t/RC) \quad (11)$$

The time constant for the above equation is RC, and the time constant for the averaging equation is approximately $N-0.5$ (i.e., in $N-0.5$ samples, the average reaches 63% of its final value).

In using the low-pass filter analogy, it is important to keep in mind that averaging is done on time slots and that adjacent time slots are not averaged together. The filter, therefore, does not affect the signal itself, but rather filters out rapid deviations of the signal from its steady-state condition.

Reference

1. R. Schlater, "Digital Scopes Gain Persistence," *Electronics Week*, February 11, 1985.

Hardware Implementation of a High-Performance Trigger System

by Scott A. Genther and Eddie A. Evel

IN ADDITION TO BANDWIDTH and time base accuracy, the usefulness of an oscilloscope is limited by the performance and features of its trigger system. The wide vertical bandwidth and precise timing accuracy of the HP 54100A/D Digitizing Oscilloscope require a very high-quality trigger system. Therefore, much effort was put into the design of the basic trigger path to achieve high performance, and a radical departure was taken from the traditional oscilloscope feature set to provide triggering functions more useful to designers of digital hardware.

The front end of the system is designed for very fast response to provide repeatable and accurate timing between input transitions and the time base reference. At the same time, it is structured like the front end of a logic analyzer, where each input has its own comparator. This allows simultaneous and sequential processing of input signals so that sophisticated functions resembling those found in dual-time-base oscilloscopes and logic analyzers can be provided.

In addition to the traditional single-source triggering mode, the new architecture provides the ability to trigger on transitions of a parallel pattern of several inputs, like a logic timing analyzer. The digital nature of the instrument makes possible a holdoff function that is much improved over the corresponding delay feature on conventional oscilloscopes. The structure of the high-speed hardware also

makes it possible to implement the sophisticated features found in the D model. The first of these adds a time qualifier to the basic parallel pattern trigger to provide two modes of time qualification, another feature found in logic timing analyzers. Also included is a three-channel state trigger, which is a simple subset of the state analyzer function, and there are two delayed modes, one by time and one by events. These, combined with the windowing capability of the time base, fully duplicate the functionality of a dual-time-base oscilloscope.

User Interface

HP 54100A/D triggering capabilities are structured into modes. Only one mode is active at a time and each is associated with its own softkey menu. In an HP 54100D, the modes are: Edge, Pattern, State, Time-Delay, and Event-Delay. The softkey menu indicates the trigger mode.

Softkey menus are the means by which parameters are entered into the trigger system. Except for the Edge mode menu, descriptive softkey labels are employed that when read from top to bottom briefly summarize the function that is performed. For example, a typical Event-Delay function display might read After Pos Edge On Chan 1 Trig On 213 Events Of Neg Edge On Chan 2.

Except for trigger levels, all parameters in one menu are completely independent of the parameters in any other

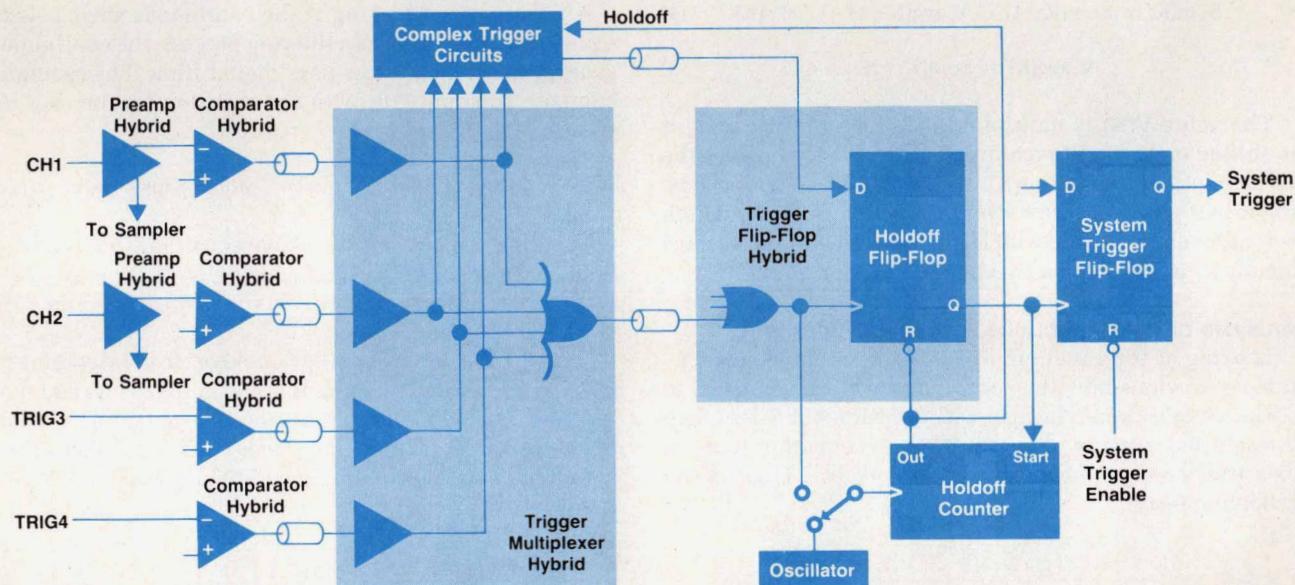


Fig. 1. Block diagram of the HP 54100A/D triggering system. Each triggering channel, two internal and two external, has its own comparator, unlike conventional oscilloscopes that multiplex analog sync signals to a single comparator.

menu. Trigger levels are excluded from the rule because it would be cumbersome to enter four trigger levels for each of the five modes. Instead, levels can only be entered in Edge and remain in effect when the oscilloscope is switched into another trigger mode. Another purpose of Edge mode is to avoid the necessity of having to learn a complex feature to do simple single-source triggering. This straightforward function is actually redundant, because it is included in Pattern mode. It is provided so that a complex feature set can be offered without inconveniencing users who do not need the full power of the instrument.

Fast-Path Hardware

Trigger performance consistent with the vertical bandwidth and the timing accuracy of the HP 54100A/D requires high sensitivity, high-frequency capability, low jitter, and minimum timing error for a wide range of signal conditions. In addition, the trigger must be programmable, and once programmed, must work in the expected fashion without requiring an operator to tweak it to get a stable waveform display on screen. This is of particular importance when the instrument is used in an automatic test system. These characteristics become more difficult to achieve as frequency increases up to the 500-MHz trigger specification of the HP 54100A/D. To achieve these characteristics, high-speed circuitry and fast edge speeds are required along with good interconnect techniques to minimize noise and timing signal degradation. The system achieves rms jitter of better than 50 ps and detectable timing errors of only a few parts per million under worst-case conditions.

The architecture of the HP 54100A/D trigger system is shown in Fig. 1. The most radical departure from the conventional oscilloscope trigger system is the use of separate high-speed trigger comparators for each of the four trigger source channels. This eliminates the need to multiplex low-level analog signals to a single comparator, thus elimi-

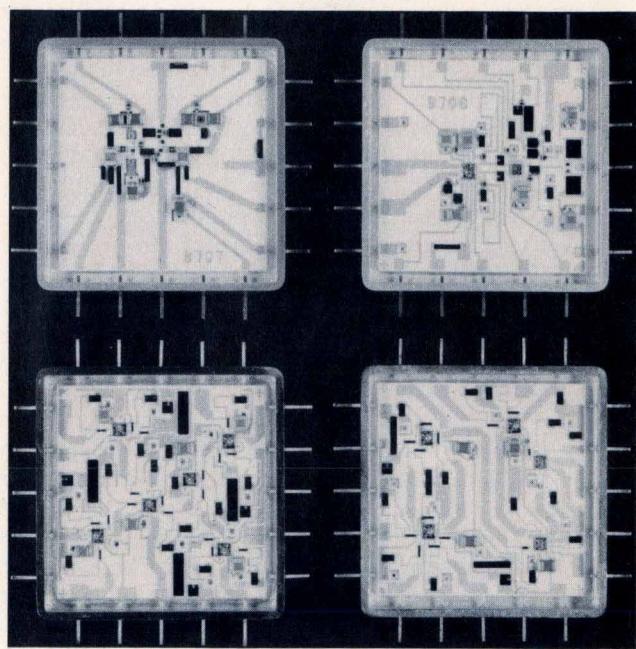


Fig. 2. The four thick-film hybrids that constitute the heart of the HP 54100A/D trigger system are the preamplifier (top left), the sync comparator (top right), the trigger multiplexer (bottom left), and the trigger flip-flop (bottom right).

nating a major source of noise coupling. Since the comparator outputs are logic signals, multiplexing is much easier. The other advantage is that now the trigger architecture is similar to that of a logic analyzer and logic triggering can now be performed.

A preamplifier in each vertical channel provides a convenient pickoff for the sync signal, along with necessary isolation from the vertical sampling circuit to prevent sampling transients from degrading the triggering.

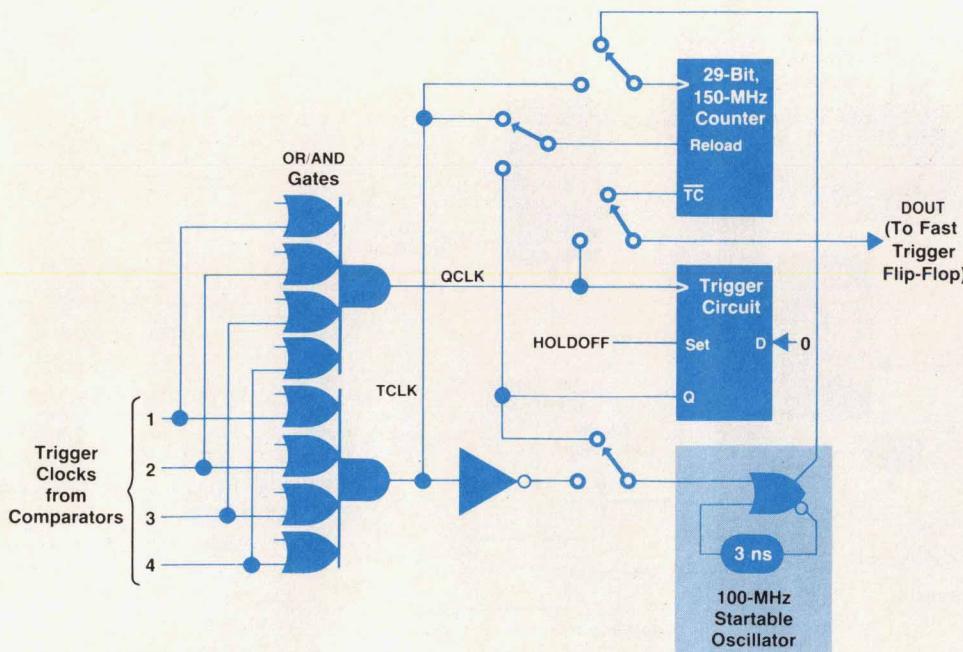


Fig. 3. Block diagram of the hardware that implements the complex triggering functions of the HP 54100D. A trigger circuit, counter, and startable oscillator operate in parallel with the fast hardware to integrate a variety of triggering functions with conventional oscilloscope triggering.

The comparator outputs go to the trigger multiplexer circuit, where one or more of the signals are used for triggering. The comparator signals are also routed to the complex trigger circuits. The multiplexer output goes to the input OR gate of the trigger flip-flop circuit and is used to clock the holdoff flip-flop. In the HP 54100A, which does not have the more complex features, the trigger multiplexer and complex trigger circuits are omitted and the comparators feed the trigger flip-flop hybrid directly. One or more of the signals may be used for simple logic pattern triggering, depending on which comparators are enabled. The trigger multiplexer and the complex trigger circuits are on a plug-in printed circuit board with coaxial signal connections. Thus, converting from the A model to the D model merely involves adding one plug-in board.

The holdoff flip-flop works with the holdoff counter to set up a fixed trigger cycle. Once a trigger occurs, further triggering is inhibited until holdoff times out. This feature is used to establish a unique trigger point on signals that cross the trigger threshold several times per cycle, such as a pulse burst. As an option, holdoff can be set to inhibit triggering for a selected number of transitions instead of a selected time.

There is a timing problem that occurs because the trigger signal must be gated by a system signal that is asynchronous—in this case, the end-of-holdoff signal from the holdoff counter. Since the signals are asynchronous, there is a possibility of violating the setup time requirements of the holdoff flip-flop if the timing of the signals is within a certain range. This could cause the flip-flop output to be time-skewed from the clock edge by more than normal, thus causing the displayed signal on screen to be time-displaced from normal—a trigger timing error. The probability of this happening increases as the trigger signal frequency increases and can be quite high at 500 MHz, the HP 54100A/D's bandwidth specification. To minimize this problem, the holdoff flip-flop is implemented as two flip-flops in

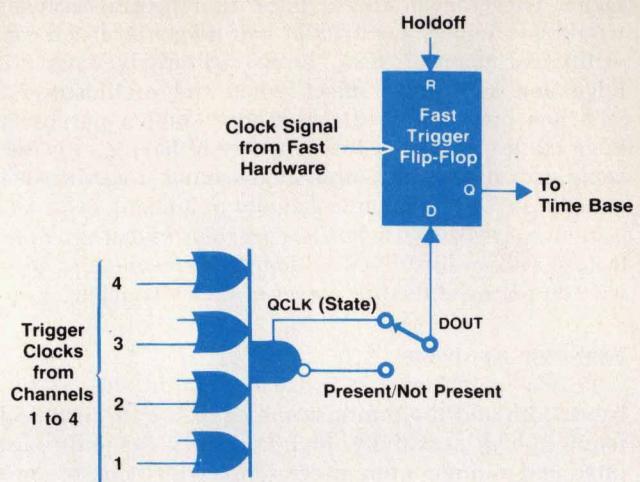


Fig. 4. Block diagram of the HP 54100D triggering system in State mode. The qualifying pattern signal generated at QCLK allows the fast trigger flip-flop to recognize edges only when the pattern (state) is true.

series, with the Q output of the first driving the D input of the second and the clock of the second delayed to allow time for the first to settle. The same problem exists at the system trigger flip-flop, but at a lower clock frequency because its clock rate is reduced by the holdoff system.

500-MHz Logic and Microwave Design

To achieve the desired performance, high-frequency design techniques are used throughout the trigger system. Four high-frequency thick-film hybrids were designed for the preamplifier, the comparator, the trigger multiplexer, and the trigger flip-flop circuits. These are the circuits that must handle signals up to the maximum trigger frequencies, above 500 MHz. Bipolar chips from HP's 5-GHz process

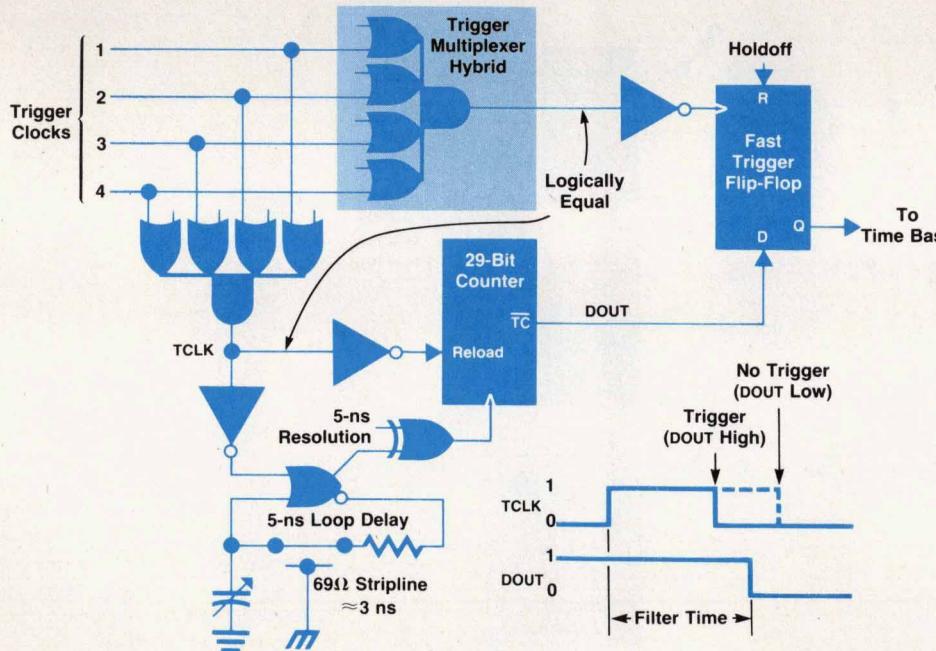


Fig. 5. Block diagram of the HP 54100D triggering system in Pattern Present < mode. The startable oscillator and the counter are configured as a timer to filter out triggers on patterns of duration greater than the programmed delay.

are mounted on the hybrids. Coaxial cables connect one printed circuit board to another. Signal connections on the printed circuit board are via 50Ω microstrip lines using a four-layer board structure with inner-layer ground planes. Special care was taken to separate the ground planes of various circuits to minimize ground coupling of noise into the trigger system. The 50Ω signal interconnect system helps minimize noise pickup and aids in maintaining the fast edge speeds necessary for low trigger jitter. Good line terminations are particularly important, because at certain frequencies, reflections can occur superimposed on pulse edges that are used for timing. This causes distorted timing edges that result in delay changes and increased jitter at those frequencies.

Fig. 2 is a picture of the four thick-film hybrids used in the trigger system. The hybrid approach allows good control over parasitic reactances. 50Ω signal lines can be run very close to the chips they connect. The package allows easy grounding and heat sinking to the back package surface. This approach permitted a fast design cycle, because the chips were already in production for other instruments; only the thick-film substrates required a new design effort.

The trigger flip-flop hybrid has several unusual features. First, ground feedthrough holes are used to improve chip grounding to a ground plane on the back of the substrate. This also greatly improves power supply bypassing through the chip capacitors mounted on the substrate. The wide traces are 50Ω microstrip lines used for signal connections. When the lines drive a chip input, the line is necked down to form a T-section network to compensate for the chip input capacitance and thus form a better line termination. There also is a small untrimmed resistor in series with some inputs to provide damping to prevent ringing on signal transitions. Ringing can cause jitter and timing variations similar to the effects of reflections mentioned above. The microstrip lines are expanded near where they connect to a package pin to provide some compensation for the inductance of the bond wire between the hybrid and the package pin.

Good grounding and control of layout parasitics are also important in the sync comparator hybrid design to mini-

mize coupling from the output to the comparator input, which could cause oscillations if not well-controlled.

Outside of the hybrids, the remainder of the trigger system is implemented with standard ECL and two HP LSI counter chips. Programming data is handled serially to minimize the number of connections from the trigger system to other digital circuits, thus reducing the noise coupled from those sources.

Dual-Path Triggering

The basic Pattern Entered/Exited and Edge triggering functions are implemented in the fast trigger path hardware. This is done in the multiplexer hybrid shown in Fig. 1. Hewlett-Packard EECL logic chips are connected to form an OR/AND gate that can selectively sum any combination of the four trigger signals from the front-end comparators.

The more complex functions of the HP 54100D, however, are not implemented with EECL because the quantity of logic required is prohibitive. Used instead are SSI and MSI functions from the 10KH ECL family. While being respectably fast, with 1-ns typical propagation delays, the 10KH logic is not capable of the speeds of EECL, so it could not be used between the comparators and the trigger flip-flop hybrid without sacrificing performance.

To resolve this dilemma, the logic for the complex functions is implemented in a path parallel to the fast hardware so as not to impede the latter's performance. This is accomplished by carefully tapping the four trigger clock signals without degrading the fidelity of their transmission to the trigger flip-flop. The 10KH logic operates on the tapped-off signals and generates a single output, DOUT, which goes to the D input of the fast flip-flop. Here the slower path interfaces back into the fast one. The 10KH logic can enable and disable the fast path, in real time, to select which edges from the multiplexer hybrid will be recognized. Thus, the complex trigger logic uses the same trigger signal as do the basic functions, but selects only certain transitions to generate triggers as determined by the active mode. Because the processing is done in parallel and path interfacing is carefully controlled, the full dynamic

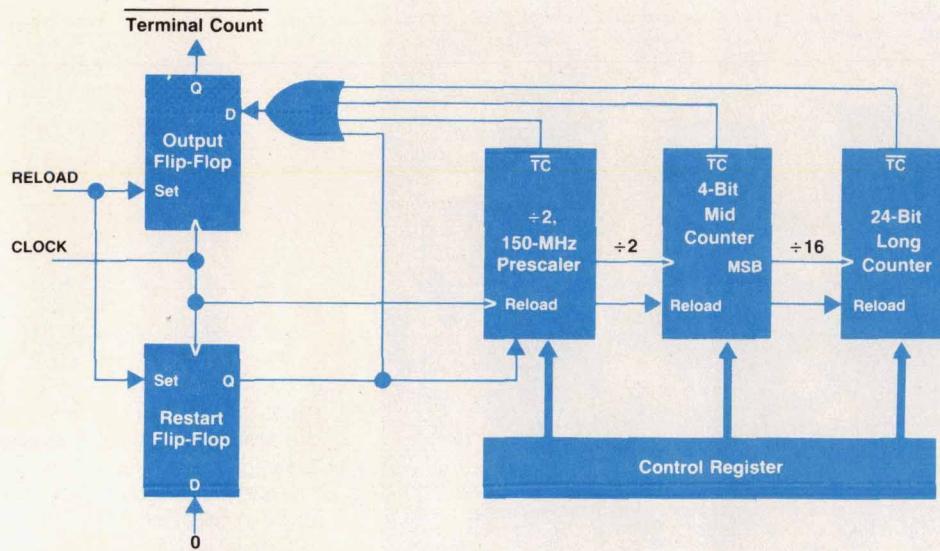


Fig. 6. Simplified block diagram of the HP 54100D's event counter. A ripple-clock, ripple-load scheme is used to achieve a 4-ns reload time and greater than 150-MHz operation.

performance of the fast hardware is preserved. Not only is nothing lost, but as will be seen later, the sophisticated features can take advantage of the performance of the fast hardware, and signal-to-trigger delay is consistent between modes. This is because the signal that is used to generate triggers always comes through the same path regardless of the mode that is active.

Fig. 3 (page 27) is a block diagram of the complex trigger hardware. The four buffered signals from the multiplexer hybrid are sent to two ECL OR/AND gates that can be independently programmed to pass any of the four clock signals or a logical product of them that constitutes a pattern signal. One gate produces the signal TCLK, which is always programmed to be the same as the signal passed to the fast trigger flip-flop. The other OR/AND gate produces a signal called QCLK, which is independent of TCLK. TCLK and QCLK can then be routed under microprocessor control to various hardware resources that can be interconnected in different ways to produce the function required by the chosen trigger mode. These hardware resources include a 29-bit counter with a 150-MHz clock rate and 5-ns reload time, a 100-MHz delay-line-type startable oscillator, and a 5-ns pulse-response trigger circuit. In general, QCLK is the signal that must occur before TCLK is processed, and TCLK is logically the same as the signal used to trigger the oscilloscope.

The key to offering a variety of triggering functions with this hardware set is its configurability. A control register, serially programmed by the microprocessor, interfaces to the hardware blocks to connect them together as called for by the user-selected trigger mode. The most straightforward example of this is State mode, in which the triggering system is equivalent to the block diagram shown in Fig. 4. State mode allows the user to define one input as the edge or clock source and then define a parallel pattern across the remaining three inputs. The oscilloscope will trigger on the edge of the clock source, but only if the parallel

pattern (state) is simultaneously true or false as desired. To implement this, the fast-path hardware is programmed to pass only the signal originating from the selected clock source. QCLK is then programmed to generate the pattern and is routed directly to the D input on the trigger flip-flop. Thus, triggers occur in response to the input selected as the clock, but only when QCLK, and hence the pattern, is true.

Time Qualification

The HP 54100D provides the ability to qualify triggers based on the time duration of a parallel pattern. In this case, the trigger can be set to occur on the true-to-false transitions of a pattern, but only if the pattern has been true for less than, or greater than, a programmed filter time. For a pattern defined over one channel, say Chan 1 - High, this function triggers on pulses qualified by their width. This is a feature found on many logic analyzers, but in the HP 54100D a different approach was taken to the design of the filter timer. Filter time accuracy and resolution are much improved over those of logic analyzers and, most important, recovery time is significantly reduced.

Fig. 5 shows a block diagram of the triggering system in Pattern Present < mode, which triggers on pattern durations shorter than the filter time. The user-defined pattern is generated at TCLK, which simultaneously goes to the trigger flip-flop through the fast hardware and to the complex trigger logic. The logic in this mode is configured as a timer. The signal being timed, TCLK, drives a startable oscillator that is used to clock an ECL counter. The counter output is routed to the D input of the fast trigger flip-flop, which is programmed to respond to negative edges (pattern exit edges) at TCLK. To follow the operation of this circuit, assume that the pattern is false, so TCLK is low. This low level forces the oscillator to reinitialize asynchronously, and also resets the counter, sending its output high. The

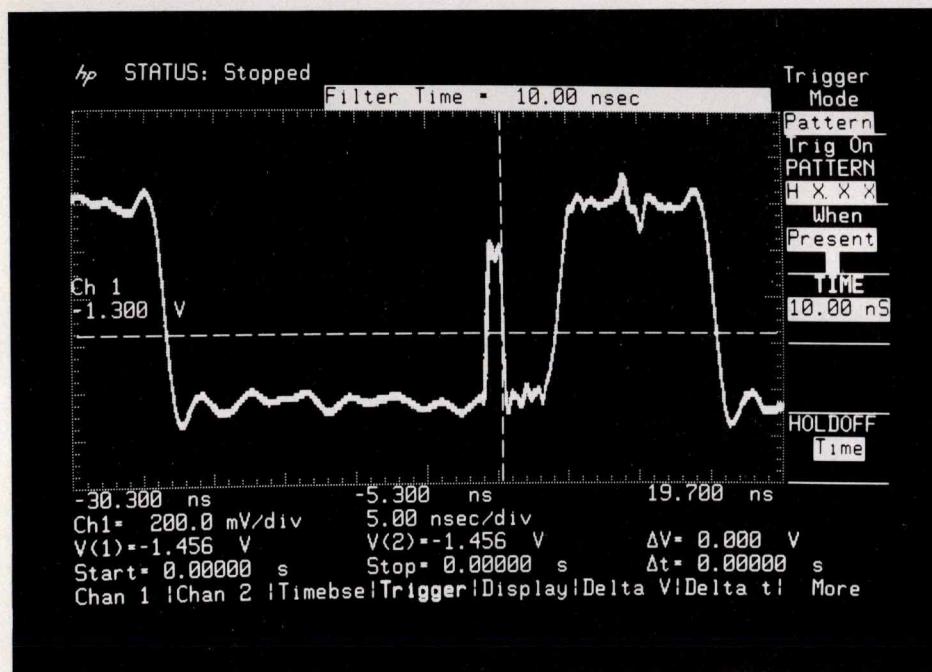


Fig. 7. Trace of a 1-ns glitch occurring within a complex data stream. The trigger is solidly locked onto the 1-ns glitch while rejecting pulses of longer duration, including the one that occurs only 5 ns later.

timer is now in its initialized state, awaiting a low-to-high transition at TCLK. When TCLK goes high, the oscillator starts sending the counter a 100-MHz square-wave clock that is synchronized to the input transition. When the counter overflows, the D input of the trigger circuit goes low and disarms the trigger flip-flop. The D input is initially high, but since the flip-flop is set to respond to negative edges, the first positive transition will not generate a trigger. The next negative transition will, however, but only if the timer has not yet run out. In this way, pulses at TCLK, which represent pattern-true durations, generate triggers only if they are shorter than the delay from TCLK to DOUT. This delay is varied by loading different values into the counter. Hence the circuit is a digitally programmable pulse width filter.

This timer scheme has several advantages. Before the occurrence of a pulse, the D input to the trigger flip-flop is high. This means that the minimum pulse width that can be recognized is limited not by the speed of the slower ECL logic, but only by the speed of the EECL logic in the fast path. So in the Pattern Present < mode, the pulse response is as good as in Edge mode, which is better than 1 ns. A second advantage is that the decision as to whether a pulse is short enough is made at the fast trigger flip-flop, so this device's coincidence-resolving ability is used to prevent jitter on pulses with widths close to the filter time. Other advantages include 5-ns resolution on the filter time setting and low timing jitter since the oscillator is synchronized to input transitions. Varying the counter load value provides 10-ns steps on the timer setting, and inverting the oscillator output under microprocessor control increases the resolution to one half the period, or 5 ns.

Pattern Present > mode is identical to the less-than mode except that the counter output is inverted before it is sent to the D input of the fast trigger flip-flop. With this adjustment the trailing edge of a TCLK pulse is recognized only if the width of the pulse is longer than the programmed delay from TCLK to DOUT, instead of shorter.

Recovery Time

Perhaps the most significant advantage of using a startable oscillator and counter combination to implement a timer is that a very short recovery time can be achieved. After the end of one pulse, the timer must be reinitialized in preparation for the next. Since this requires a finite amount of time, there is a limit to how closely two pulses can be separated if the second pulse is to be timed correctly. This timing constraint, known as recovery time, is less than 5 ns in the HP 54100D.

For the timer to initialize on a low transition of TCLK, three things must occur. First, the DOUT signal must go high before the next pulse occurs. Second, the oscillator must be reset to its nonoscillating static state. Third, the counter must be reloaded with the count value. Getting DOUT to return high quickly is simply a matter of using fast enough logic, and the type of startable oscillator used inherently resets quickly. Referring to Fig. 5, the oscillator consists of a two-input NOR gate with a delay line wrapped around from output to input. The line is simply a long series-terminated printed circuit board stripline, and the variable capacitor is used to set the frequency of oscillation.

The other input to the NOR gate is the control input for the oscillator.

A high-to-low transition at this point sends an edge through the NOR gate, which propagates around the loop at a 100-MHz rate. When the control input goes high, the NOR gate output is forced low. When this transition arrives at the other input to the gate, the oscillator is reinitialized. Thus, it takes an amount of time equal to the loop delay (only 5 ns in this circuit) to reset the oscillator and make it ready for the next pulse. It may seem necessary to wait for reflections from the capacitor load to settle out before restarting the oscillator, but because the line is terminated at one end, these reflections are absorbed. Also, they do not interfere with the oscillating mode should the oscillator be restarted soon after initialization, because the reflections and the normal transitions are propagating in opposite directions. Second-order effects, such as a slight impedance mismatch at the terminated end and nonlinear gate output impedance, can enhance the effect of reflections, but these were found to be insignificant for the purposes at hand.

Fast-Reloading Counter

With an oscillator that resets in 5 ns, a counter that reloads in 5 ns or less is needed if it is not to limit recovery time. Fig. 6 shows a simplified block diagram of this counter. It is 29 bits wide, can be clocked at 150 MHz, and reloads in about 4 ns. Because of the number of bits in the counter, clocking and reloading it in parallel was not practical. Instead, the counter is constructed as three cascaded stages: a divide-by-two prescaler, a 4-bit mid counter, and a 24-bit long counter. The prescaler clocks the mid counter, and the long counter counts the most significant bit of the mid counter.

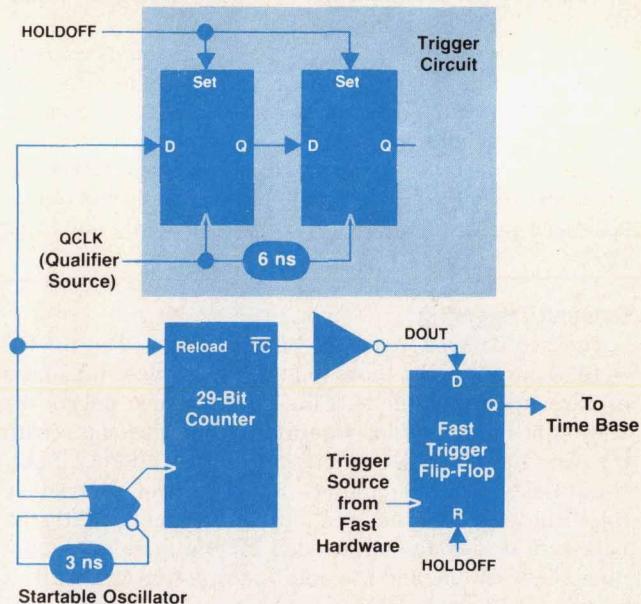


Fig. 8. Block diagram of the HP 54100AD triggering system in Time-Delay mode. The startable oscillator and counter are configured as a timer that is started by a trigger circuit. This allows time to be counted from an edge at QCLK before arming the fast trigger flip-flop.

Only the front-end prescaler has to operate at the full 150-MHz repetition rate. The mid counter operates at half this rate, and the long counter at 1/32 of this rate. In this way, successively slower logic can be used for the upper stages and the amount of logic that has to work at 150 MHz is minimal. The prescaler is implemented with a few 10KH SSI logic functions, the mid counter with a 10KH up-counter, and the long counter with an HP LSI ripple-down counter chip.

To achieve the 4-ns reload time, the counter is not only clocked, but also loaded in a ripple fashion. Some additional logic is included in the two earlier stages to generate initialization signals for the later stages, but this occurs only the first time through a counting operation. During the first few clocks of a count cycle, the prescaler keeps track of the edges, and at the same time causes the mid counter to be initialized. By the time the prescaler needs to generate a carry, the mid counter is ready to accept it. Similarly, the first time through its sequence, the mid counter initializes and loads the long counter, which is ready by the time the mid counter generates a carry. All it takes, then, to reload the entire counter is to take the prescaler back to its initial state, which is done by setting the output and restart flip-flops shown in Fig. 6. Hence the reload time for the counter is determined by the amount of time needed to set a pair of flip-flops and wait for their outputs to propagate through any adjacent logic. The counting operation then proceeds correctly, since regardless of the state of the rest of the logic, the counter is in its initial state when these two flip-flops are set. For smaller load values, the counting cycle is not long enough for all of this initialization to take place. In this case, parts of the counter are turned off through bits in the control register, which also holds the load value. When the load value is less than 37, the long counter is turned off and values less than 5 result in the mid counter's being turned off also.

Fig. 7 shows the trace of a signal consisting of a glitch occurring within a complex data stream. The oscilloscope is set to trigger on pulses less than 10 ns wide. This signal demonstrates both the excellent pulse response in this mode and the short recovery time. The trigger system solidly locks onto the glitch that is only 1 ns wide and rejects all longer pulses, including the one that occurs only 5 ns later.

Delayed Triggering

The two delayed modes of triggering offered in the HP 54100D are similar to those found on some dual-time-base oscilloscopes. One mode, called Time-Delay, delays the arming of the trigger for a period of time that starts with the occurrence of some qualifying event at another input. Event-Delay mode is similar, but instead of delayed by time, the trigger is delayed by signal events. With the hardware that is already needed for the State mode and time qualification, and the addition of a trigger circuit, it is straightforward to implement these delayed modes. Fig. 8 shows the block diagram of the trigger system in Time-Delay mode. Fig. 9 shows the same for the Event-Delay mode.

In Time-Delay mode, the counter and oscillator are connected to implement a timer. However, instead of being

driven by TCLK, the timer is driven by the trigger circuit, which is driven by QCLK. QCLK is programmed to pass the signal from the input the user specifies as the qualifying source and the fast hardware is programmed to pass the input selected as the trigger source. When an edge occurs at QCLK, the trigger circuit latches and starts the timer by releasing the startable oscillator. When the counter clocks out, DOUT is sent high and the trigger flip-flop is armed to trigger on the next edge coming through the fast hardware. When that edge occurs, a trigger is generated and the hold-off signal (HOLDOFF) pulses high. This resets the trigger circuit and the timer to repeat the cycle. The trigger circuit is similar to the fast trigger flip-flop, but is implemented with 10KH logic. Two flip-flops are connected in series with the clock of the second delayed. This arrangement filters out transients that might occur should QCLK and HOLDOFF transitions occur simultaneously, thus providing a timer-start signal well synchronized to QCLK.

The Event-Delay mode configuration is similar to Time-Delay except that the counter is driven by TCLK instead of the startable oscillator, and the trigger circuit consists of a single flip-flop. After an edge at QCLK, the count-start flip-flop latches. This releases the counter to begin counting TCLK edges. The counter sends DOUT high to arm the fast trigger on the second-to-last edge at TCLK so the fast trigger flip-flop counts the last event and generates a trigger. In this way, the actual trigger is always generated through the fast hardware, which results in consistent signal-to-trigger delay between modes. When the trigger occurs, HOLDOFF goes high and resets the circuit in preparation for another cycle.

If the counter is released near the time it receives a clock, it is possible that it can go to some incorrect state because of the ambiguous inputs. To avoid this, the counter has been designed such that the difference between its initial state and the next state involves only one flip-flop. In this way, the counter can go only one of two ways: either it recognizes the first edge, or it ignores it. Either way is valid. If the first state transition depended on the action of several

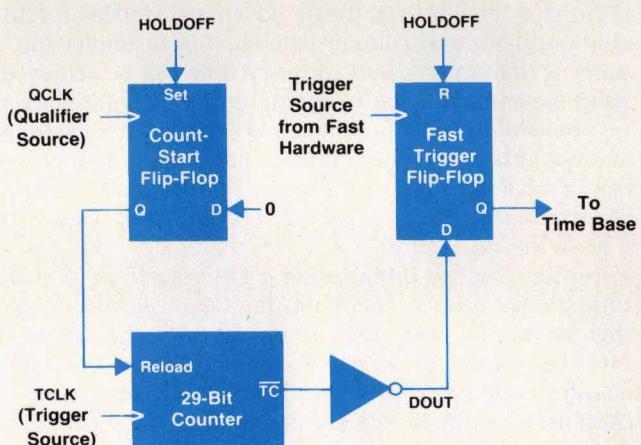


Fig. 9. Block diagram of the HP 54100D triggering system in Event-Delay mode. After an edge at QCLK, the counter is gated to begin counting TCLK edges. On the second-to-last event, the counter generates an output and arms the trigger flip-flop, which counts the last event and generates the trigger.

flip-flops, there would be no guarantee that they would all resolve the coincidence in the same way. This could result in a jump to an incorrect state, which could disrupt the counting sequence and make operation seem erratic.

1-GHz Digitizing Oscilloscope Uses Thick-Film Hybrid Technology

by Derek E. Toeppen

HYBRID MANUFACTURING TECHNOLOGY is widely used for high-frequency and microwave circuitry because of its ability to provide low-parasitic connections to devices and controlled-impedance connections between devices. Typically, the devices found on the hybrids are discrete rather than integrated because the components (step-recovery diodes, high-frequency pnp transistors, etc.) or the circuit structures (delay lines, inductors, etc.) do not lend themselves to integration. In the case of microwave circuitry, thin-film technology has been the usual choice because of its very controlled geometries. The drawback to using thin-film hybrids has always been the cost.

When the designers of the HP 54100A/D Digitizing Oscilloscope started looking at technologies for manufacturing their 1-GHz-bandwidth circuitry, they were interested in hybrids

for all the reasons just mentioned. However, because of the relatively low frequency of the circuits, the need for a wide range of resistor values required by the broadband nature of the circuitry, and the promise of lower cost, they became interested in using thick-film hybrid technologies.

Early investigation confirmed that the HP 54100A/D hybrids could be built using thick-film technology with the development of some new processes and some improvements of existing ones. New and improved processes were necessary because the facility chosen to build the hybrids had no previous experience building hybrids beyond 300 MHz. The primary goal of the process work was to allow the design and production of the hybrids required for the HP 54100A/D. However, a broader goal was to provide the capability to use thick-film hybrid technology for circuitry with bandwidths

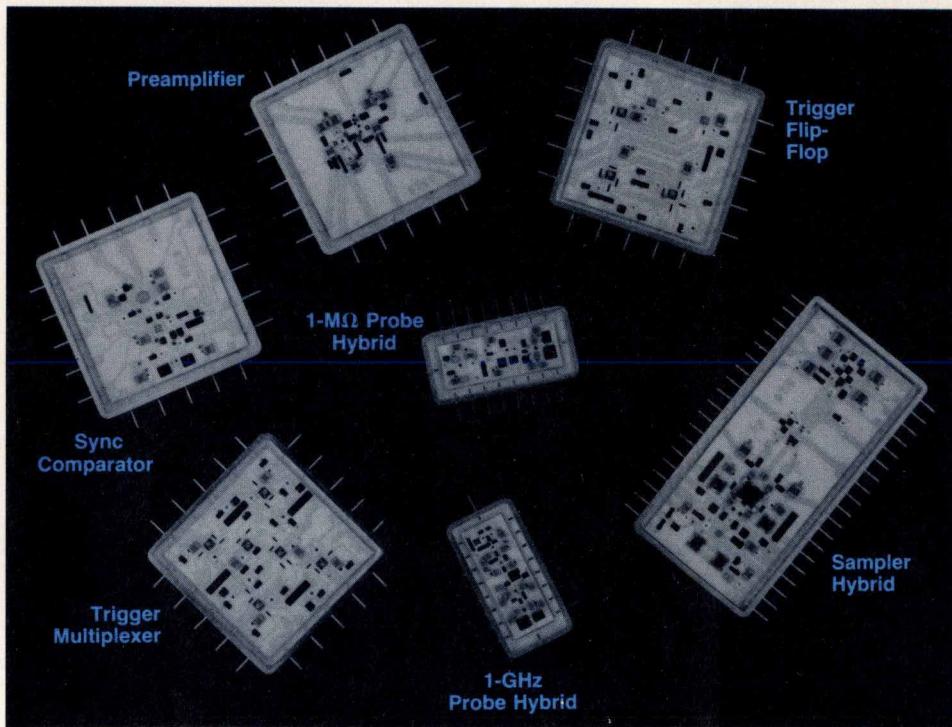


Fig. 1. The seven HP 54100A/D thick-film hybrid circuits without lids.

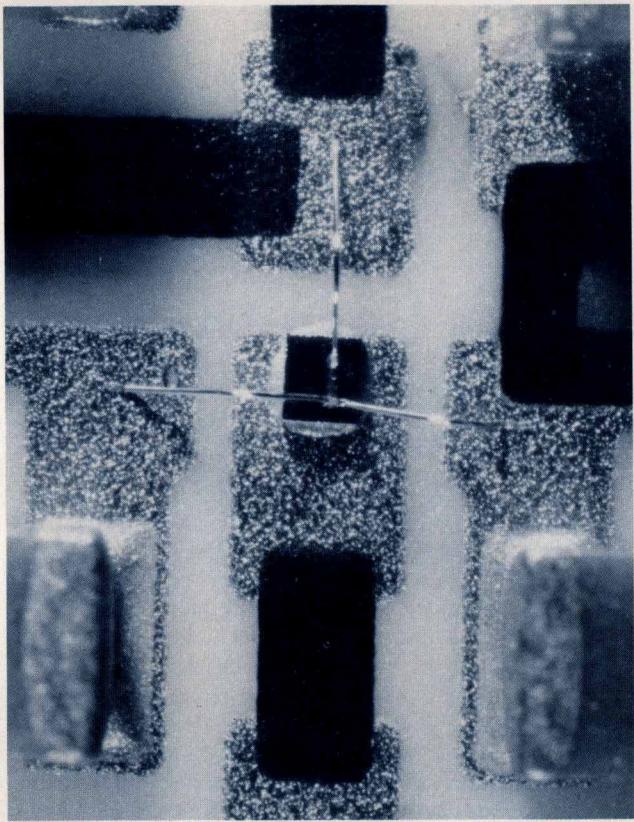


Fig. 2. Wedge and stitch bonds.

into the 3-to-4-GHz range. An integral part of this goal was the desire to provide circuit models for thick-film components and assembly techniques. Such capability would allow a designer to reduce hybrid development time and cost by permitting a complete and accurate evaluation of a hybrid's performance before one is ever built.

Seven hybrids were designed and are in production using the new processes: a 1-GHz preamplifier, a 3-GHz sampler, a 1-GHz probe receiver, a 300-MHz probe receiver, and three trigger hybrids making up a 500-MHz trigger system (Fig. 1).

Wedge Bonding

The wedge bond (Fig. 2) has been the most commonly used wire bonding technique for microwave hybrids. It is attractive for two reasons, first for its ability to bond to the small pads on microwave devices, and second for the direct path the wire takes to the substrate surface. The short path means a reduction in the wire inductance.

A further reduction in wire bond inductance can be made with what is called a stitch bond (Fig. 2). The stitch bond is a continuous wedge bond that starts on the substrate, goes to the pad on a device, then travels across the device and back down to the substrate surface on the other side. The result is two wire bonds on a single device pad. Running the bonds in opposite directions reduces mutual coupling between the bonds, allowing nearly a factor of two reduction in inductance compared to a single wire bond.

Both the single wedge bond and the stitch bond are used on the HP 54100A/D hybrids.

The actual inductance of a wire bond loop is difficult to

calculate theoretically. Table I shows the theoretical inductance per unit length of a wire over a ground plane. Measurements indicate that over the distance of a typical wire bond loop these numbers are reasonably accurate. Thus, in modeling a hybrid assembly, the values in Table I are used with no compensation for the wire bond shape. In all models discussed in this paper it is assumed that the physical size of the component being modeled is much smaller than the wavelength of the signal seen on it. This allows a lumped parameter model to be used. For circuit designs in the 3-to-4-GHz range, care should be taken to see that this condition is met.

Beam Lead Bonding

A beam lead device is an unpackaged monolithic device with small leads extending directly from the device surface (Fig. 3). Such devices can be mounted flat on a substrate with the leads bonded directly to the substrate's conductors. The result is a short and low-inductance connection between the device and the hybrid. A thermosonic bonding technique, using a modified wire bonder, was developed for bonding the beam lead diodes in the HP 54100A/D hybrids.

Models for beam lead devices supplied by device manufacturers usually include lead inductance, so it was unnecessary to measure this parameter.

Printed-Through Holes

One critical requirement for the successful use of any thick-film technology is the ability of the technology to

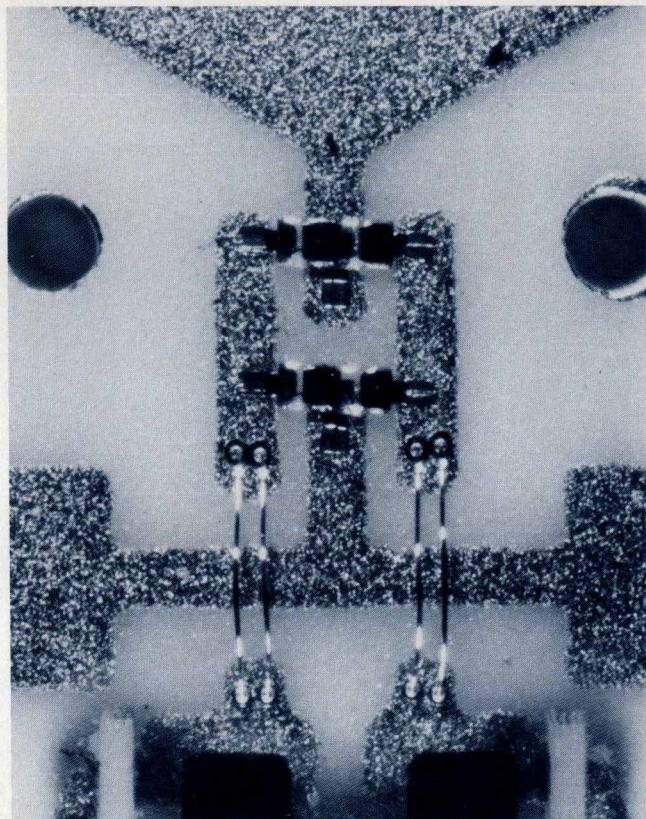


Fig. 3. Beam-lead device mounting.

Table I

Inductance of a wire bond lying flat over a ground plane (nH/mm)

	Substrate thickness (mm)	
	0.635	1.27
Wire bond diameter (μm)	25.4	1.0
	17.8	1.09
		1.27

provide a low-inductance path from anywhere on the substrate to ground. To ensure this, a process to coat holes in the substrate with a standard thick-film conductor was developed (Fig. 4). This provides the shortest possible path to the substrate ground plane.

The inductance of the printed-through holes was determined experimentally. The values are shown in Table II.

Reduced Thick-Film Resistor Geometry

A second critical requirement for the use of thick-film technology is the ability to minimize thick-film resistor geometries to reduce resistor parasitics. The reduction was achieved in two ways. First, work was done on the resistor printing and trimming processes. This allowed the

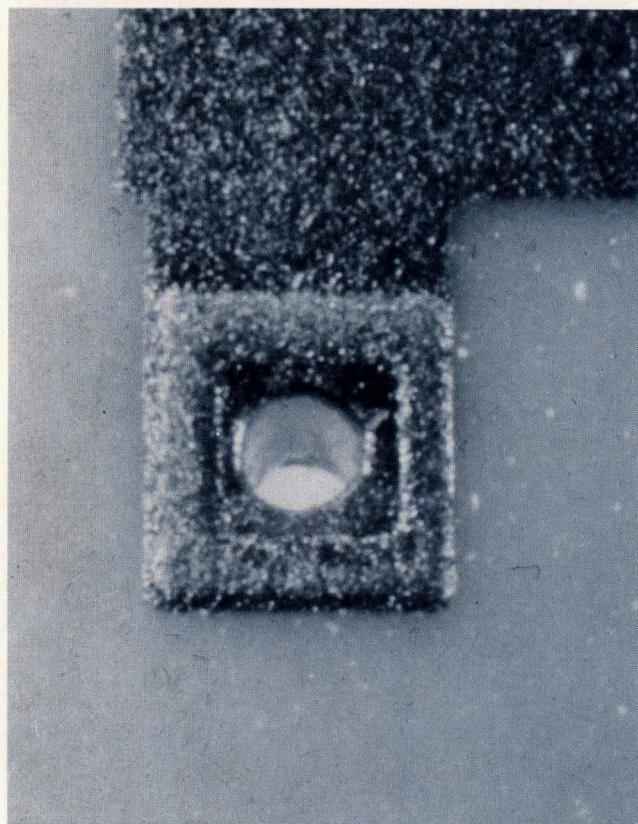


Fig. 4. Printed-through hole.

Table II

Inductance of thick-film printed-through holes (nH)

	Substrate thickness (mm)	
	0.635	1.016
Hole diameter (mm)	0.381	0.148
	0.508	0.114
		0.206

minimum resistor dimension specification to be reduced to 0.508 mm on a $\pm 1\%$ trimmed resistor and 0.254 mm on a $\pm 30\%$ untrimmed resistor. The untrimmed resistor is widely used for series damping where tolerance is not important, but where low parasitics are critical. For a $\pm 1\%$ 50Ω resistor, this first step brought about a 33% reduction in resistor area. Second, new resistor paste values were provided for designing resistors. This means a narrower range of resistor aspect ratios is required to make resistors. The result is an additional 53% reduction in the area of a 50Ω resistor.

The model for thick-film resistor parasitic inductance and capacitance is probably the most interesting result of the modeling effort because of its simplicity (Fig. 5). The parasitics are calculated by simply thinking of the physical resistor as a length of lossy stripline. Using stripline equations¹ and the geometry of the resistor, the characteristic impedance, wave velocity, and inductance and capacitance

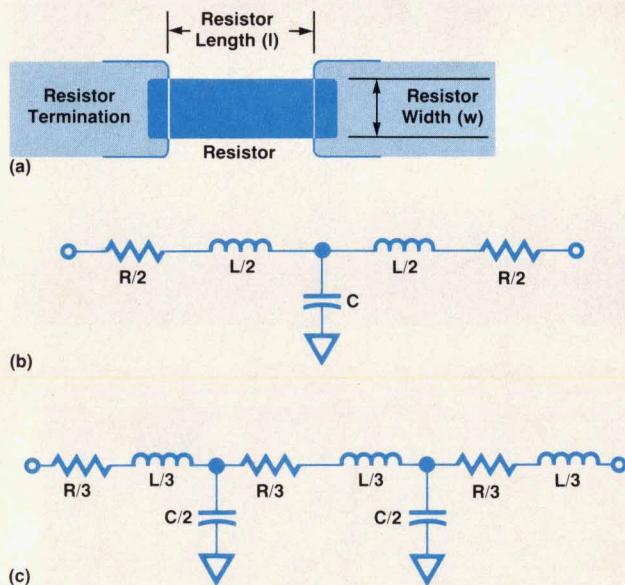


Fig. 5. The thick-film resistor R in (a) can be considered to be a length l of stripline of width w to calculate its parasitic inductance L and capacitance C . (b) shows a simple model of the resistor, and (c) shows an expanded model for higher-frequency analysis.

per unit length can be calculated. The inductance, capacitance, and resistance are then used to create a lumped parameter model of the resistor. The typical resistor model used for the HP 54100A/D hybrids is shown in Fig. 5b. If the circuit analysis is to extend into the several-GHz range the model can be expanded as shown in Fig. 5c. Note that the model shown in Fig. 5 does not include the capacitance of the resistor terminations.

Striplines

One of the essential features of the hybrid technology is the ability to produce controlled-impedance lines. Therefore, an effort was made to evaluate and improve the print resolution and quality of thick-film conductors. This effort was not so much to make a quantitative improvement in the specified print resolution but rather to improve print yield. The conductor width resolution specification is ± 0.0254 mm. This provides a 50Ω line with $\pm 0.25\%$ tolerance on a 1.016-mm-thick substrate and $\pm 2\%$ tolerance on a 0.635-mm substrate.

An accurate set of equations for stripline calculations has been published by H. A. Wheeler.¹ These equations have been programmed for the HP-41C Calculator by L. Mar and M. W. Gunn.² The stripline equations also proved useful in determining the parasitics of conductor runs, other than intentional controlled-impedance lines, and conductor areas such as resistor terminations. Values obtained from these equations were used to model resistor terminations and to place intentional parasitics for tuning a circuit's response. Response tuning was used widely in

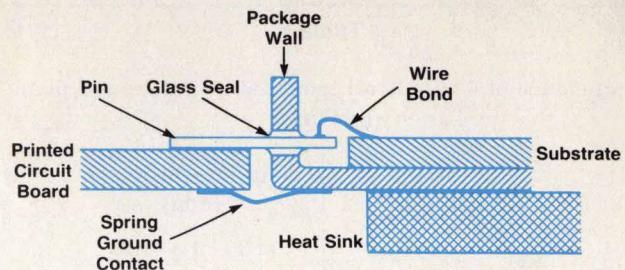


Fig. 7. Cross section of hybrid circuit package assembly.

the HP 54100A/D trigger hybrids to compensate for the input capacitance of the digital integrated circuits. Fig. 6a is a photograph of the compensation on one such hybrid. The equivalent circuit is shown in Fig. 6b. R1 is the untrimmed series damping resistor mentioned earlier.

Packaging Process

The typical microwave package is a machined metal cavity with coaxial connectors in the walls for high-frequency connections and pins for dc connections. In keeping with the desire to hold down hybrid cost, a number of alternative packages were investigated for the HP 54100A/D hybrids. A stamped metal package with horizontal pins out the cavity walls was ultimately chosen (Fig. 7). This package provides relatively low cost, easy hybrid mounting, an easy high-frequency interface, and metal construction for thermal concerns and noise protection. The hybrids are mounted in the package using conductive epoxy. The horizontal package pins are connected to the substrate using standard ball wire bonds, and the lid is sealed using a conductive epoxy preform.

Two things were done so that the horizontal pin interface would provide the required $<5\%$ reflection with a 300-ps-rise-time signal. First, the dielectric constant of the glass holding the pin in the package wall was specified so that the pin in the wall has close to a 50Ω impedance. Second, capacitance is added on either side of the package wall to compensate for high pin inductance outside the wall. The extra capacitance on the hybrid side is created using three wire bonds to connect pins to the substrate (instead of just one) and by adding a small area of thick-film conductor where the bonds connect to the substrate.

Acknowledgments

I would like to acknowledge the efforts of the Colorado Springs Technology Center's processes engineering group and especially Mark Michaels for their timely process development work, Mike McTigue, Martin Guth, and Lewis Dove for their work developing the hybrids, the Microwave Technology Division and the Network Measurements Division for freely sharing their microwave processes and experience, and Ken Rush for his insight in modeling high-frequency circuitry.

References

1. H. A. Wheeler, "Transmission-Line Properties of a Strip on a Dielectric Sheet on a Plane," *IEEE Transactions on Microwave Theory and Techniques*, Vol. MTT-25, no. 8, 1977, pp. 631-647.
2. L. Mar and M.W. Gunn, "HP-41 Solves Widest Range of Microstrip Design Problems," *Electronics*, February 10, 1981, pp. 156-157.

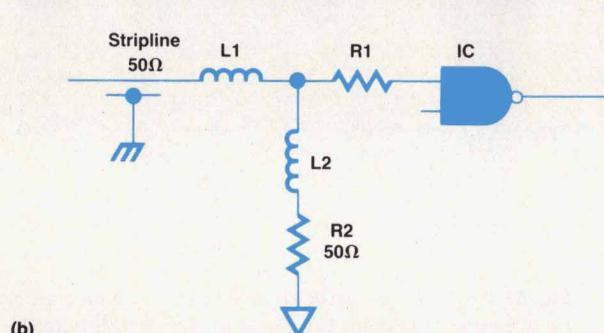
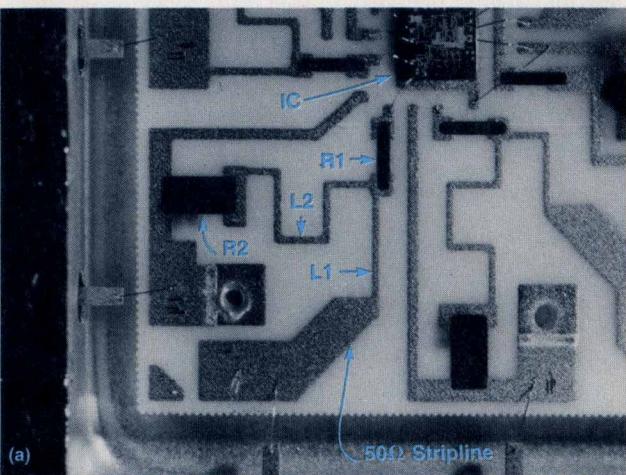


Fig. 6. (a) Photo showing response tuning of the trigger hybrid. (b) Equivalent circuit.

A Modular Power Supply

by Jimmie D. Felps

THE POWER SUPPLY for the HP 54100A/D Digitizing Oscilloscope is divided into three different modules (Fig. 1).

The primary assembly converts the input ac voltage to a dc voltage that is distributed to two dc-to-dc converters: the digital power supply and the analog power supply. These are switching mode power supplies using pulse width modulation (PWM) techniques at a frequency of about 70 kHz. Each power supply assembly is a plug-in card that plugs directly into a motherboard. The primary assembly mounts directly to the HP 54100A/D chassis. Two three-wire cables between the primary assembly and the digital and analog power supplies are the only chassis wiring required. Instrument cooling is accomplished by a single fan that varies in speed with the ambient temperature. The fan drive circuitry is located on the analog power supply. Each of the three modules can be tested and serviced individually.

Operation

The primary assembly rectifies and filters the input ac voltage. When the voltage select switch is in the 230V position, the ac input is rectified to yield ~300 Vdc. The ac input is doubled in the 115V position to yield the same dc voltage.

The primary assembly has surge protection circuitry that protects against ac line voltage transients and inrush current. Overvoltage crowbar devices throw the circuit breaker for sustained overvoltage conditions (e.g., 230Vac line voltage when in the 115V position). The main EMI filter is located in the dc current path to reduce component size and increase the effectiveness of the filter. The PWM control circuitry on each dc-to-dc converter is powered by bleeder resistors across the bulk energy storage capacitors.

The digital power supply (Fig. 2) has two output voltages, +5V and -5.2V, and can deliver 15A on either or both outputs. A single-ended forward converter topology was chosen for both of the dc-to-dc converters because of its simplicity, reliability, and performance at power levels under 200 watts. Postregulators provide voltage regulation for both outputs. Voltage regulation and current limiting are accomplished by V-I loop control of the PWM IC. Fault conditions use the shutdown line to latch the PWM IC in

the off position. Line isolation is accomplished via optoisolators and the main power transformer in the power stage.

An efficiency of >70% is achieved at full load by differentially regulating the outputs as shown in Fig. 3 (the outputs are regulated to a differential voltage of 10.2 volts while ignoring the voltage of each with respect to digital ground) and allowing only one postregulator to be in a linear mode (regulating) at a time. Load conditions can be such that both postregulators are saturated, but usually the lightly loaded output has its regulator in a linear mode of operation. A more detailed description of the voltage regulation is given later in this article.

Transformer-Coupled Series Regulators

High-efficiency transformer-coupled series regulators are used as postregulators.¹ Since both postregulators function in exactly the same manner, only the +5V postregulator (Fig. 4) will be discussed.

A conventional three-terminal linear regulator has a dropout voltage (minimum $V_{in} - V_{out}$ required for regulation) of ~2.5V. The output current passes through the pass transistor in the regulator, resulting in a significant power loss in high-current low-voltage outputs. In the HP 54100A/D power supply, the postregulators are inserted in the ac portion of the secondary so transformers can be used to overcome some of the deficiencies of conventional regulators.

Notice in Fig. 4 that the +5V output circuitry looks like a typical single-ended forward converter output except that T2 has been inserted. Capacitor C1 is required to block dc from the T2 secondary. Bridge rectifier CR1 converts the ac (bidirectional) current from T2 to a dc (unidirectional) current for the pass transistor Q1. A turns ratio of 50:1 was chosen for T2 because it allows Q1 to operate at one fiftieth of the output current and the effective pass voltage is one fiftieth of the voltage that appears across Q1. When regulating, the pulse from power transformer T1 is bucked by the regulator, reducing the output level during the charge portion of the cycle, t_{chg} , and the flywheel clamp voltage is boosted by the regulator during the flywheel portion of the cycle, t_{fly} (Fig. 5). The result is a level shifting of the waveform at point A by a voltage equal to V_{reg} , where $V_{reg} = V_{Q1}/50$. When Q1 is saturated, $V_{reg} < 0.1V$. To prevent excessive voltage on Q1, VR1 sets the limit to 75V. This results in a maximum value of 1.5V for V_{reg} . The output filter inductor L1 was chosen to be a coupled inductor to improve the cross regulation of the outputs.² L1 also limits the amount of peak detection that occurs on an unloaded output to ~1V because of its transformer action. Therefore, adequate regulation can be accomplished with a V_{reg} range of 0.1 to 1.5V. Magnetization current for T2 is supplied by bleeder resistor R1 so that regulation occurs for 0 to 100% load variations.

Refer to Fig. 2 to see how the outputs are balanced relative

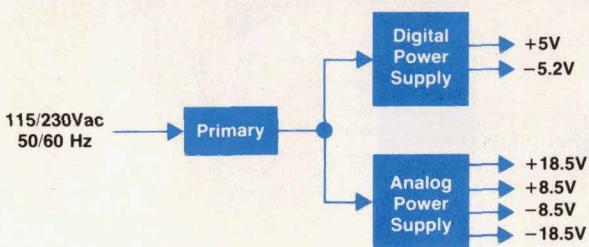


Fig. 1. Basic power supply block diagram.

to digital ground to give the desired output voltages (+ 5.0V and - 5.2V). The balance amplifier A5 compares the two outputs to digital ground to determine which output voltage is too high. (Only one output can attempt to go too high because the differential control loop, which is the fast control loop, maintains a 10.2V output differential.) The output voltage of A5 then increases or decreases to the threshold voltage (+ 2.2V for the + 5V regulator or + 2.5V for the - 5.2V regulator) of the error amplifier whose output is too high. That regulator then becomes linear and brings the high output back within limits. The gain of A5 is such that when the + 5V output is ~7 mV high, the + 5V regulator comes out of saturation and the - 5.2V output will be ~7 mV low. The opposite condition occurs when the - 5.2V output goes ~7 mV high.

The worst-case power dissipation for either pass transistor for all combinations of loads was found to be 4 watts.

Analog Power Supply

The primary circuitry of the analog power supply is identical to that of the digital power supply. The secondary circuitry differs in that there are no postregulators. The + 18.5V output is regulated via the V-I loop, and the three other outputs are "tag along" outputs. Cross regulation is improved by the use of a coupled inductor in the output filter. Local three-terminal regulators are located in the system to provide circuit isolation.

PWM IC

A semicustom PWM IC was designed to be operated on

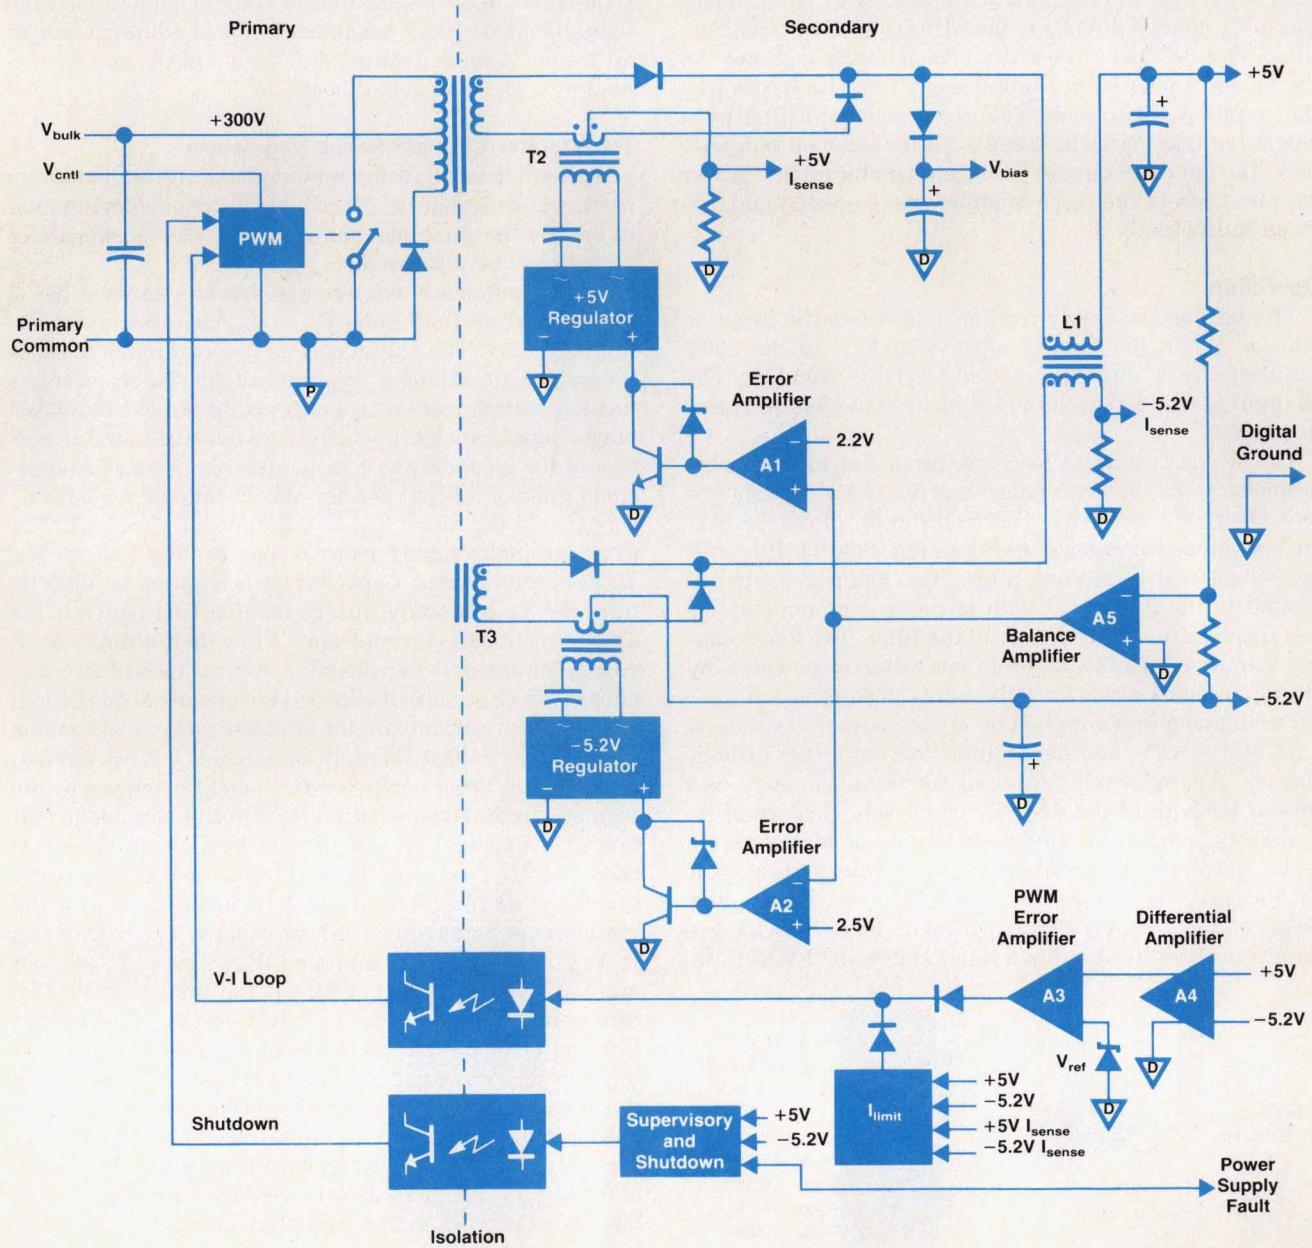


Fig. 2. Digital power supply schematic diagram.

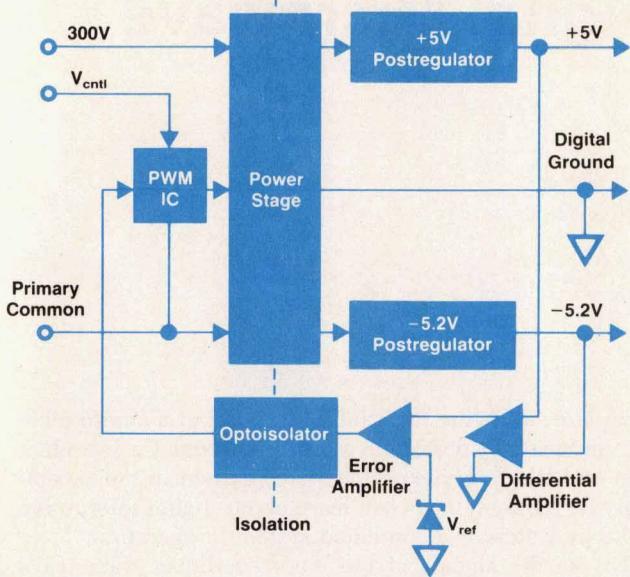


Fig. 3. Pulse width modulation (PWM) voltage control.

the primary side of the power line. This IC is used in both the digital power supply and the analog power supply. A feature set that could not be found in an off-the-shelf IC and was deemed necessary was:

- Feedforward capability
- Undervoltage sensing with hysteresis (V_{bulk} and V_{ctrl})
- Overvoltage sensing (V_{bulk})
- Low power (programmable)
- MOSFET drive output
- Demagnetization
- Dual-threshold current limit (cycle-by-cycle and hiccup mode)

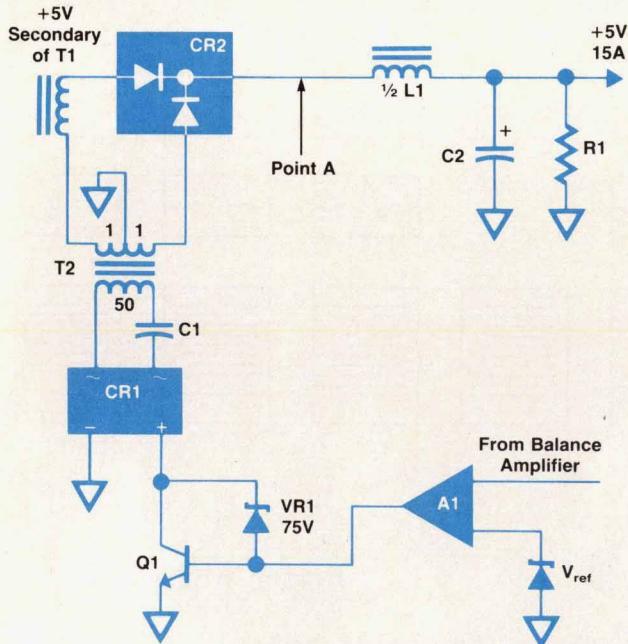


Fig. 4. +5V postregulator.

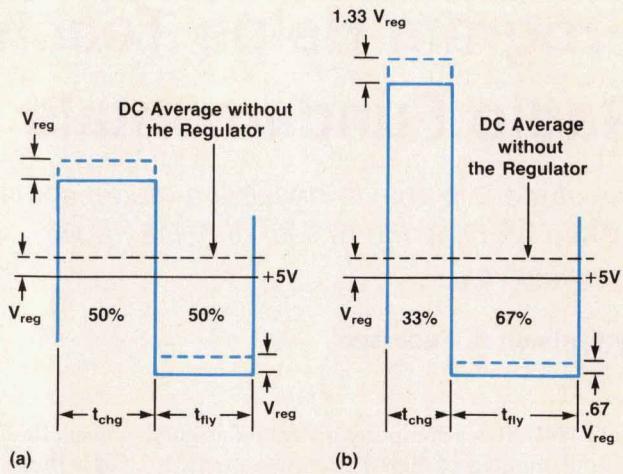


Fig. 5. Postregulator operation. This is the voltage at point A in Fig. 4. (a) 50% duty cycle. (b) 67% duty cycle.

- Overvoltage and overtemperature latch
- Internal Zener diode (for current-fed applications).

References

1. S. Pepper, "A New High-Efficiency Post-Regulation Technique for Multiple Output Converters," Proceedings of POWERCON 9, 9th International Solid-State Power Electronics Conference, 1982, Section D-3, pp 1-9.
2. W. Hirschberg, "Improving Multiple-Output Converter Performance with Coupled Output Inductors," Proceedings of POWERCON 9, 9th International Solid-State Power Electronics Conference, 1982, Section G-4, pp 1-5.

Program Helps Teach Digital Microwave Radio Fundamentals

The student is able to observe a model communications system and see the results of noise, fading, and nonlinearities.

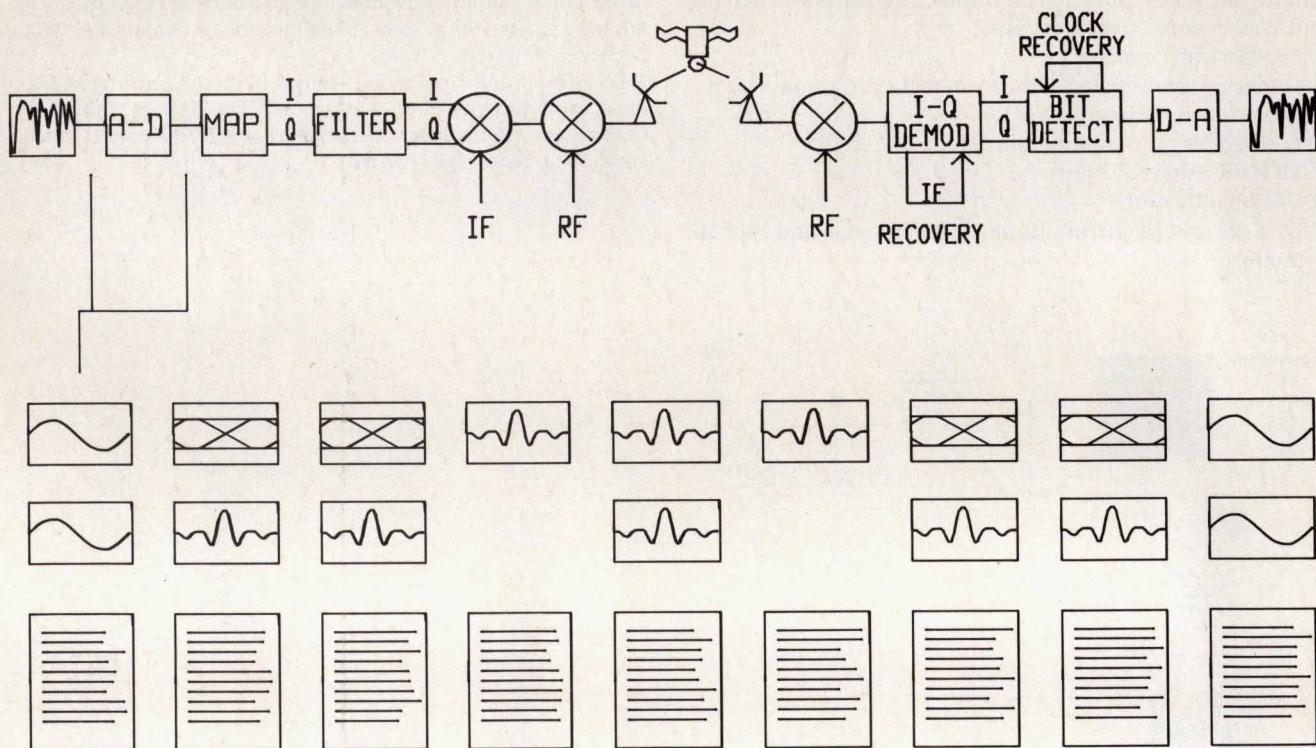
by Christen K. Pedersen

I-Q TUTOR is a computer program designed to teach the fundamentals of digital microwave radio to people from a wide variety of technical backgrounds. I-Q Tutor (HP 11736A) teaches fundamentals by modeling a digital radio. Displays are used to make the program more intuitive, and digital signal processing methods are used to model the digital radio signals and typical degradations like noise and multipath fades.

The I-Q Tutor computer program is graphically intensive and interactive. Instead of using a sequential teaching structure, I-Q Tutor acts as a real-time digital microwave radio

simulator, allowing the student to work at a comfortable rate and learn by observing. Working through the introduction and laboratory experiments contained in the accompanying manual, users can learn about digital microwave radio by watching a simulated system in operation.

To see the signals of the modeled digital microwave radio, the user positions a "probe pointer" with the computer's right and left cursor keys while looking at a system overview page (Fig. 1). This selects the point in the system block diagram where the signals will be "probed." Then a down cursor key activates the instrument displays avail-



Modulation Type: QPSK

SNR: 40 dB

Filter Alpha: .3

Fig. 1. This system overview screen shows the workbench structure of the I-Q Tutor computer program. The digital radio system is represented by the block diagram at the top with the signal displays and laboratory notebook pages below.

able at that point in the system. Drawings of the signals on the screen give a sense of reality to the program, making it feel more like a real system than a model. In addition to the instrument displays, there are also "lab notebook pages" containing descriptions of the block diagram elements and the displays. By simulating the complex signals of a real digital microwave radio, it is hoped that many of the intuitive leaps that normally come from working on real equipment in the lab might be achieved in the computer-aided instruction.

In addition to being highly visual, the I-Q Tutor program is easy to use. One of the significant contributions of the computer program is that it doesn't require any previous knowledge to use it. In fact, someone entirely unfamiliar with electronics and communications can operate it and still learn something.

Graphic Display Techniques

Part of I-Q Tutor's user friendliness comes from its emphasis on visual displays. Since the primary difference between analog and digital microwave communication is in the modulation schemes used, it was considered important that the program be able to communicate these differences and, if possible, show them. To do this, the signals of digital microwave radios had to be represented in an easy-to-understand way.

To provide a point of reference, traditional spectrum analyzer and oscilloscope instrument displays are mimicked (Fig. 2). The oscilloscope displays plot the time-domain data arrays in a linear horizontal fashion, and the spectrum displays plot the data obtained with the FFT (fast Fourier transform) on a log scale. These show the baseband modulating signals, their frequency spectra, the frequency translation processes, and the noisy baseband signals at the receiver. These are the displays that the engineer or technician familiar with traditional microwave communications is used to, so they serve as a point of reference.

Unfortunately, none of these displays completely shows the digitally modulated signals of digital microwave radio systems. To display digital microwave radio signals in a

meaningful way, those who work with digital microwave radios commonly use high-bandwidth X-Y oscilloscopes with the in-phase and quadrature (I and Q) channels of their radios displayed on the X and Y axes, respectively. When displayed like this, the phase, amplitude, and I and Q components can all be seen easily.

However, this is not the only way that digital microwave radio people look at their signals. Another standard way of looking at digital microwave radio signals is the eye diagram (Fig. 3). With one or two clock periods of I or Q shown on an oscilloscope, the eye diagram is formed by successive traces displayed because of the CRT's persistence. Of course, the scope must be triggered with the data clock for this display to have any meaning. The eye diagram is valuable because it shows how the data is represented by the signal displayed. Ideally, at the widest opening of the "eyes" there are only distinct levels representing the data levels to be transmitted or received. Degradations like noise on the signal cause the display trace to spread out, making the eye opening seem to close. A partially closed diagram indicates that the signal will be more difficult to detect and interpret correctly.

The eye diagram and Q-vs-I methods of displaying digital microwave radio signals are much more related than their methods of generation might indicate. In reality, there are three variables, all interrelated, that are important to understanding the signals in digital microwave radio: they are I, Q, and time. Since mastering this concept represents a significant level of understanding about how digital microwave radio systems work, it was considered important that the program make this accessible to the user.

To display the I, Q, and time signals in their eye and Q-vs-I orientations while still allowing the interaction between all three parameters, a technique called a vector display was developed. Since there are three parameters associated with the signals and only two dimensions on the computer screen, a two-dimensional projection of three-dimensional data is used. The projection looks like a view of a three-dimensional object that can be rotated in three axes. By performing this general transform rather than

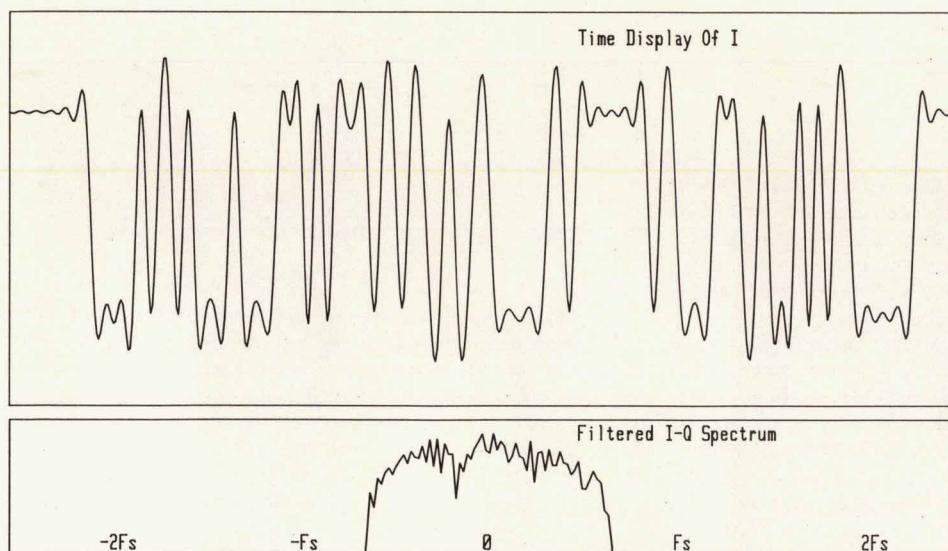


Fig. 2. Traditional methods of displaying signals, like oscilloscopes and spectrum analysis, are mimicked by the program, allowing the user to see displays of simulated digital microwave radio signals like these. Although these provide a useful point of reference, there are better ways to look at the baseband digital microwave radio signals.

just showing the eye and Q-vs-I orthogonal projections, it is possible to see the relationship of all three parameters and bring a subtle theoretical concept to life. The vector display allows the user to select the orientation of the I, Q, and time axes to view the Q-vs-I and eye diagrams as well as other arbitrary projections of the three orthogonal axes (see Fig. 3).

With the vector display, it is also possible to look at the I-Q signal states just at the clock instant on the vector display. This is the time in the data cycle when the signal should be representative of a valid data state. When observed with the time axis going into the screen, the digital microwave radio signal sampled at the clock instant should

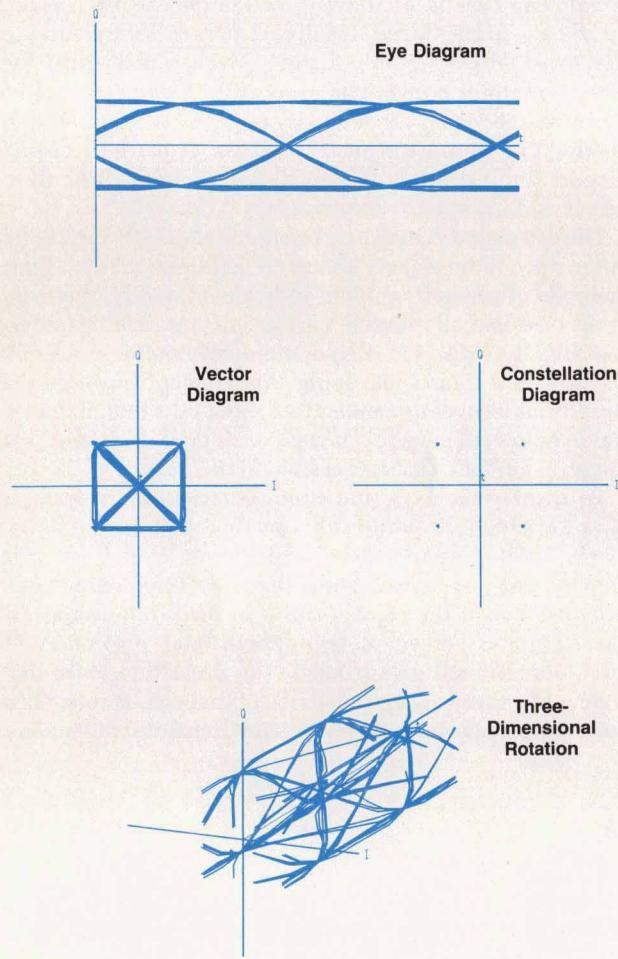


Fig. 3. In the field of digital microwave radio, the capabilities of normal measurement devices are often augmented to help understand the complete significance of digital microwave radio signals. The simplest of these displays is the eye diagram. By triggering an oscilloscope with the data clock and displaying either the I or Q baseband signals versus time, several data sequences can be seen at once with their valid data points appearing at the same horizontal position. The vector diagram results from displaying Q versus I and provides a view of the phase plane of the signal. The constellation diagram is a vector diagram without the transitions between symbol states. Finally, it is possible to display I, Q, and time as a three-dimensional figure. All of these techniques provide insightful views of otherwise meaningless signals, and for that reason the I•Q Tutor program is capable of all of them.

consist of points at each of the transmission states for the modulation format chosen. This display's dot structure led to its being called a constellation diagram. The constellation diagram is valuable in observing the effects of various degradations on the system.

Representing the Digital Microwave Radio Signal

Although the graphic displays of I•Q Tutor show a great deal of information, they would be pointless if the signals they show were not realistic. To obtain representative information, the I•Q Tutor program relies heavily on signal processing techniques to create data for the displays by mathematically modeling a real digital microwave radio. The program can be used to model an ideal system with any one of many filters or it can show more realistic signals with the addition of noise and some pertinent degradations, such as multipath fade and high-power-amplifier nonlinearities. To understand the modeling techniques, it is important to know a little about how digital microwave radios transmit information.

The complex signals used in digital microwave radio are modulated and demodulated using quadrature signals in I-Q modulators and demodulators. As shown in Fig. 4, the modulator and demodulator have two input and output signals: the I, or in-phase, and the Q, or quadrature-phase, signals. In the modulation process, the I and Q signals generate complex modulations in the I-Q plane. The product of this modulation process is a signal that is modulated in both phase and amplitude, an unusual concept to those accustomed to only amplitude or frequency modulation. One consequence of this process is that a complex array is required to represent the modulating and demodulated signals of the radio. The I•Q Tutor program models the baseband signals of the digital microwave radio using complex arrays made up of real and imaginary parts to represent the I and Q components. This data structure is also convenient for simulating some of the degradations, since complex FFTs are relatively common and greatly simplify the modeling algorithms.

The frequency spectra obtained with the FFT are instrumental in understanding the effects of filters and the addi-

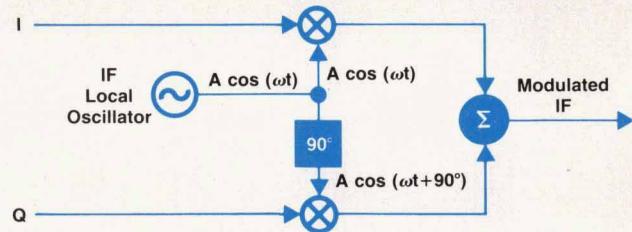


Fig. 4. Digital microwave radios differ from traditional microwave radios primarily in the way that they modulate information onto the radio signal. The quadrature or I-Q modulation process common to most digital microwave radios is illustrated above. By selectively adding predetermined levels of quadrature and in-phase local oscillator components, the modulator can produce any phase or amplitude of the carrier. The specific set of I and Q driving voltages that are used at the modulator's input determines the modulation type and information transmitted by the radio.

tion of noise. With a five-sample-per-clock-period resolution, the main lobe and one and a half sidelobes are visible in the frequency displays. The frequency spectra make many of the fundamental degradations of digital microwave radio clearly visible. For example, the noise floor added during transmission can be seen, as well as subtler effects like the regeneration of filtered sidebands when high-power-amplifier nonlinearities are present.

Filtering

Perhaps the most fundamental element of the digital microwave radio model is the filter function. Digital signals have filtering requirements that are best satisfied by Nyquist filters. The I•Q Tutor program allows the user to select Nyquist filters of various bandwidths to observe the effects they have on the system's signals. The Nyquist filter is defined in the frequency domain by the parametric equations shown in Fig. 5. The characteristic that makes Nyquist filters the choice for digital signal filtering is that they filter pulse code modulated signals without degrading the information levels. The filter doesn't introduce any interference because its step response in the time domain has a maximum response at zero time and a zero response at all other multiples of the clock period. This response ensures that the ringing of the filter because of one symbol state will not affect the data level of any other symbols.

The effects of various data filters are modeled by multiplying the appropriate Nyquist filter coefficients by the frequency spectra obtained using the FFT. When an inverse FFT translates the frequency spectra back to the time do-

main, the resulting arrays show the effects of filtering. Since filters are used at both the transmitter and the receiver of a digital microwave radio system, only the square root of the Nyquist filter's response is used when modeling the transmitter or receiver. This distributed filter design ensures that the back-to-back combination of transmit and receive filters will provide perfect data filtering and not introduce any distortions. The split filter also causes the filtered signals at the transmitter to exhibit some eye closure or intersymbol interference, since they have only been filtered by half of the distributed data filter.

The ability to change the Nyquist filter's bandwidth quickly and observe the effect it has on the system signals provides a powerful teaching tool. It would be impossible for a real laboratory to provide the diversity of filter designs available here, and even if it were possible, it's unlikely that the effects of the filter on the entire system could be as conveniently observed.

Simulating Noise Degradations

While filtering is the most critical digital microwave radio element, noise is the most fundamental degradation experienced by digital microwave radios. Indeed, much can be learned about how a digital microwave radio works in the real world by simply analyzing how it responds to various signal-to-noise ratios.

In any communication process, there is some noise present. The effect of noise on the transmission integrity varies greatly depending on the method used to transmit information. In digital microwave radio, the transmission path is

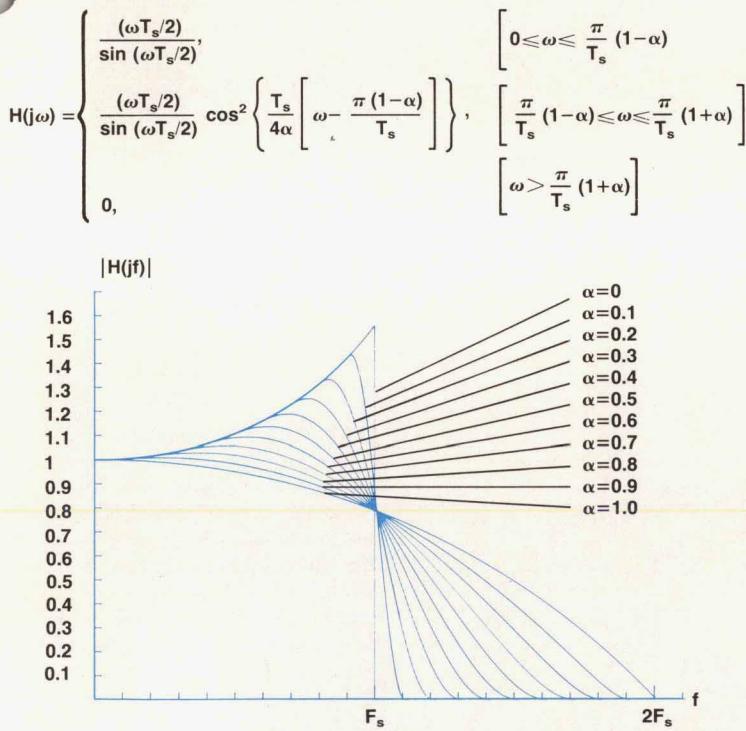
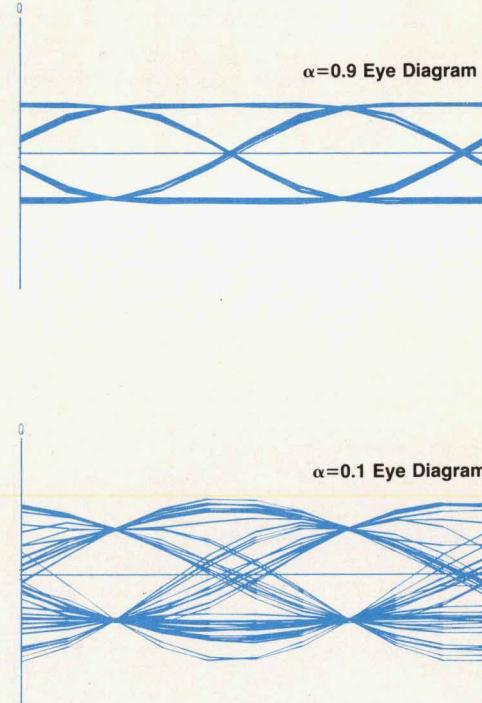


Fig. 5. The filtering requirements of digital microwave radio are very different from those of traditional analog systems. The filters must be able to filter pulse code modulated signals without introducing any interference to the symbol values. The I•Q Tutor program uses ideal Nyquist filter transfer functions to simulate real digital microwave radio filters. By selecting various filter rolloff factors (α), the user can see the effects of narrow and wide filters on the system, as shown in these two example eye diagrams from the program.



usually between two microwave towers or between an earth station and a satellite. For either case, some noise will be introduced in the transmission process, degrading the signal.

Since noise is one of the most fundamental degradations experienced by digital microwave radios, it was considered important that the I-Q Tutor program model it realistically. Conceptually, this is not a difficult task. Noise is familiar to many people as a flat noise floor on a spectrum analyzer. To model this, a flat noise floor can be added in the frequency domain representation of the filtered baseband signals given by the FFT. What phase should the flat noise floor have?

The nondeterministic nature of noise dictates that its frequency components should have random phase. The frequency spectrum obtained using the FFT is made up of real and imaginary components at every frequency, however, and the calculation of real and imaginary frequency components with constant amplitude and random phase involves transcendental functions. To avoid the additional computation time required to calculate the real and imaginary components of random-phase noise, a simplified approach is taken. By using one of four random phases (0, 90, 180, and 270 degrees) the effects of random noise can be approximated without incurring the additional computation time required by the transcendental functions. The level of the noise components to be added is calculated from the user-selected signal-to-noise ratio, and then the scaled noise vector is added to or subtracted from the real and imaginary frequency components of the spectrum.

This four-phase noise model can be used because its time-domain equivalent is approximately correct. When noise is present in a digital microwave radio signal, it causes a greater ambiguity in the signal. The effect of the noise on the signal can be seen easily in the I-Q phase plane. The noise is described by a Gaussian amplitude distribution and a linear random phase distribution. Since the four-phase model assigns each frequency component its own independent phase, the vector addition of all of the signals yields a random phase distribution and a good approximation to a Gaussian amplitude distribution. The effects of noise shown by the program include the spreading of states, eye diagram closure, and bit error rates (BERs), and are all very accurately modeled by the four-phase noise

addition (see Fig. 6).

To verify that the addition of noise is truly representative of theoretical systems to the extent that it is observable in the working program, many BER measurements were made with a variety of signal-to-noise ratios. Fig. 6 shows the results, which agree quite well with the theoretical curves for the same values.

Although much can be learned about digital microwave radio signals from looking at the relatively idealistic signals modeled by the simple addition of noise and filtering of data filters, the real world is much different. Two more complex degradations have particularly drastic effects on digital microwave radio systems: multipath fades and high-power-amplifier nonlinearities.

Simulating Multipath Impairments

Multipath fade distortions of the frequency spectrum can severely degrade digital microwave radio operations, so it is important that I-Q Tutor be able to show some of the distortions introduced by multipath fades. Before we discuss the algorithm it is important to understand the causes of multipath fade and the effects it has on digital microwave radio communication integrity (see Fig. 7).

Multipath fades occur in real telecommunications links when there is more than one transmission path and the paths are of different lengths. In spite of the time spent making sure that microwave antennas are correctly aligned, some multipath fading is unavoidable. Multiple transmission paths arise from a variety of causes, and are dependent on many environmental effects. The two most common causes of multipath fades are reflections and diffractive paths. Reflections arise from the microwave signal's reflecting off objects within the path, including clouds, trees, and the ground. Normally transmitters and receivers are located to avoid reflective objects in the transmission path; however, reflections from clouds are inevitable. Diffractive paths are caused by the bending of radio waves when they are transmitted through the air. The amount that the transmitted signals bend during transmission is dependent on the refractive index of air at microwave frequencies. As the air heats, cools, and moves, the refractive index is constantly changing, creating a dynamic fade that has dramatic effects on digital microwave radio communications.

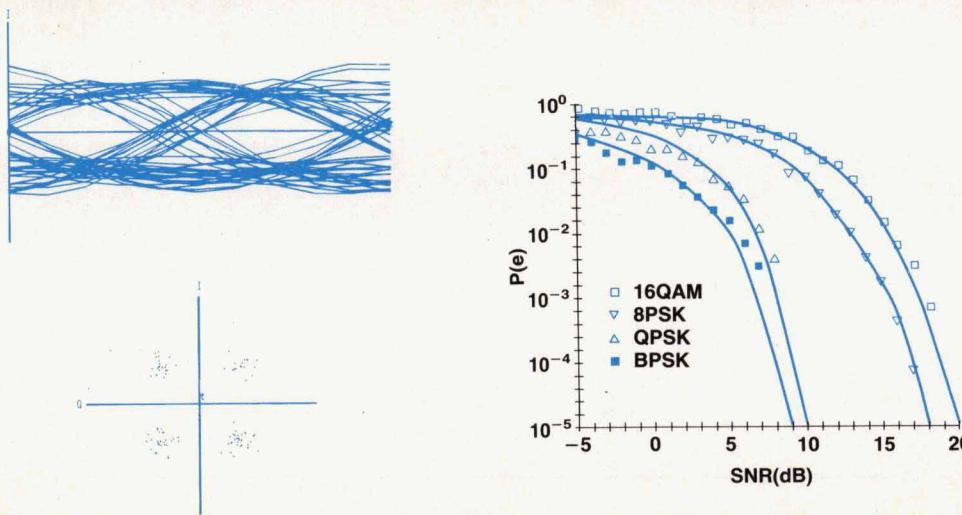


Fig. 6. Noise is introduced to digital microwave radio signals during transmission and has a profound affect on the quality of digital communications achieved. The effect that noise has on digital microwave radio signals is accurately modeled by I-Q Tutor and can be seen best in the program's eye and constellation diagram. The relationship of the $P(e)$ (probability of error) to the signal-to-noise ratio (SNR) is well understood and often plotted as shown above. The I-Q Tutor program's performance (dots) accurately matches the curves, tracking them well for large $P(e)$.

Although it is tempting to try to explain multipath fade in the time domain, the problem is easier to understand in the frequency domain. A time delay in the frequency domain has a flat amplitude response, but exhibits a positive sloping ramp in phase. This is because for a given time delay, the higher-frequency components are delayed more in phase than the low-frequency components. The vector addition of the individual frequency components for the path with the short delay and the path with the longer delay is the frequency spectrum incident at a receiver experiencing a multipath fade.

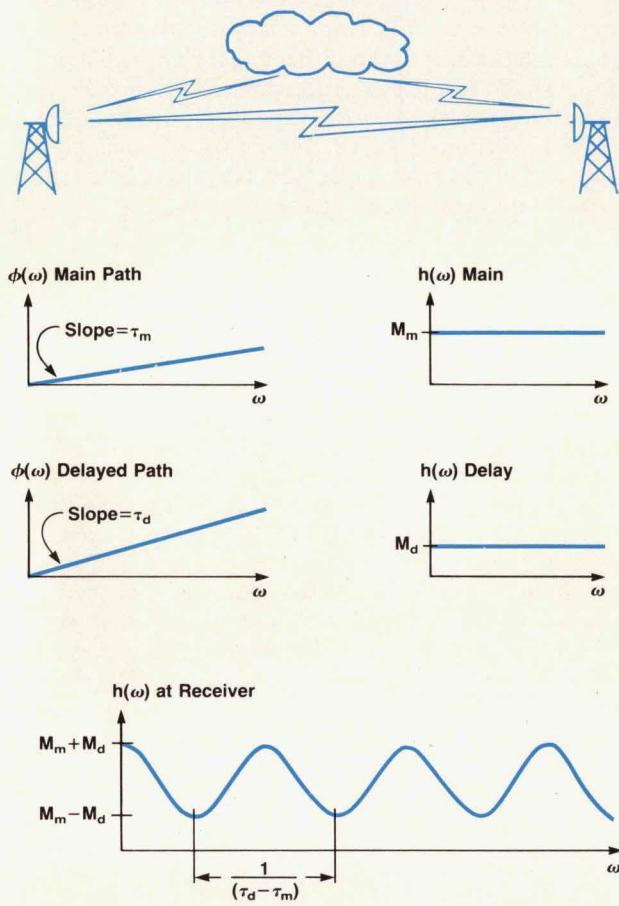


Fig. 7. The three most prevalent transmission paths for microwave signals are line-of-sight, diffractive, and reflective paths. The effect that multiple transmission paths have on the received digital microwave radio spectrum can be dramatic. Since a time delay results in a phase ramp in the frequency spectrum, the relative delay of one path with respect to another means that the phase of the delayed signal is shifted as a function of frequency with respect to the direct signal. When these add at the receiver, they add constructively or destructively depending on the specific frequency and delays. This results in a sinusoidal ripple in the received frequency spectrum, which when plotted on a log scale, appears as notches. The depth of the notches is a function of the amplitudes of the main and delayed signals, and their separation in the frequency domain is inversely proportional to the differential delay. By modeling the multipath fade with a frequency dependent phase shift, I•Q Tutor accurately simulates many of the effects of a real multipath fade.

As can be seen in Fig. 7, at some frequencies the phase of both signals adds constructively and at others it adds destructively. This leads to a sinusoidal variation with frequency in the amplitude response of the channel seen by the receiver. When plotted on a log scale this response appears as notches in the frequency domain. The frequency difference between the notches is the reciprocal of the differential delay, and their depth is a function of the magnitude of the two paths relative to one another.

From this description of multipath fading, a modeling technique follows directly. Although it is true that multipath fades in real systems are the product of an infinite number of transmission paths, many of the effects of multipath fade are represented well by the two-path model. The I•Q Tutor program allows the user to specify two-path fades in terms of three parameters: fade depth, notch position, and fade delay. The fade is modeled by adding a phase ramp to a copy of the spectrum of the transmitted signal before the addition of noise. This is done by first calculating the phase of each frequency component from its real and imaginary parts, then adding the appropriate phase shift. The real and imaginary components are then returned to the array. Once the phase ramp has been added to achieve the delay, the array is scaled to achieve the specified notch depth and then added to the original transmitted spectrum.

To verify the model's accuracy, the multipath effects have been verified by comparison with theoretical M curves, which show multipath fade degradation as a function of notch position. The close agreement found between the model and theoretical curves can be verified by the user through a lab exercise in the back of the I•Q Tutor manual.

Simulating High-Power-Amplifier Nonlinearities

High-power-amplifier nonlinearities are another major degradation experienced by digital microwave radio systems. Nonlinearities are particularly important to satellite systems where amplifiers have to be operated near compression to satisfy link power budget requirements. Although the algorithm is fairly straightforward, the following background describes why this method was chosen, and more important, why high-power-amplifier nonlinearities are important to digital microwave radios.

When an amplifier is operated near its compression point, its performance is degraded as its transfer function becomes more and more nonlinear. Two measures of the amplifier's nonlinearity are of particular value in evaluating the effects of the nonlinearities on digital microwave radio signals. These are amplitude modulation to amplitude modulation conversion (AM to AM) and amplitude modulation to phase modulation conversion (AM to PM). AM to AM describes how the amplifier's output amplitude varies as a function of input amplitude. Ideally, the amplifier's output amplitude is directly proportional to the input amplitude, but real amplifiers have limited maximum output amplitudes and usually behave nonlinearly before they reach these levels. AM to PM, on the other hand, is the amount of phase shift introduced in the output signal as a function of the input signal's amplitude. Once again, the ideal amplifier introduces no incremental phase shift, but real amplifiers often have delay characteristics that are a

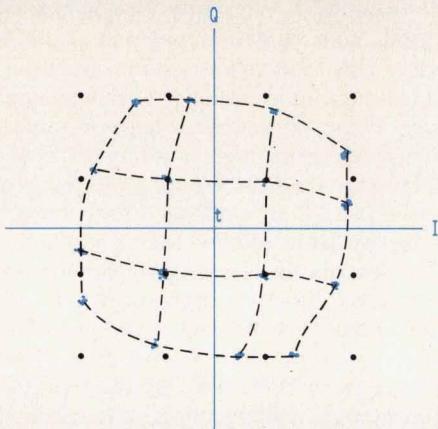


Fig. 8. The nonlinear characteristics of high-power amplifiers are highly visible on displays of digital microwave radio signal states like this constellation diagram. The compression of the amplifier at high amplitudes causes the outermost states to be drawn in. The amplitude dependent phase shifts also cause the outermost states to be rotated. A 16QAM modulation was used for this diagram because the amplitude variations help to excite the nonlinear channel. I•Q Tutor allows varying degrees of nonlinearities to be selected to observe the increase in degradation with increased amplifier drive levels.

function of the input amplitude, causing greater phase shifts in the output with higher input amplitudes.

The effect that these AM to AM and AM to PM characteristics have on digital microwave radio signals is easy to observe, since the signals are customarily displayed in the I-Q plane. Digitally modulated microwave signals that have been degraded by high-power-amplifier nonlinearities look distorted. In Fig. 8, the set of ideal 16QAM* states is shown with those of a 16QAM signal that has been degraded by high-power-amplifier nonlinearities. Notice that the outermost states, the states with the greatest amplitude, have been compressed in magnitude and rotated in phase while

*16QAM stands for 16-state quadrature amplitude modulation.

the inner states remain relatively unchanged.

Although the high-power-amplifier nonlinearities pose problems for digital microwave radio designers, they may be at least partially compensated for using predistorters and modulations that have smaller amplitude variations.

Modeling the nonlinear elements of the high-power amplifier must be done in a different way from the degradations previously discussed. The other impairments could all be modeled in the frequency domain, greatly simplifying the algorithms required in some cases. Nonlinear elements, however, cannot be modeled easily in the frequency domain, since the FFT only works for linear systems. Therefore, I•Q Tutor models high-power-amplifier nonlinearity in the time domain. I•Q Tutor allows the user to specify the operating level of the amplifier and then use predefined AM to AM and AM to PM performance curves to calculate the amplifier response for every time-domain sample. Since the phase rotations involve performing a rectangular to polar transform and then polar to rectangular, the modeling is somewhat time-consuming, but effective.

Surprising Results

The simulation of realistic degradations not only helps neophytes to understand the primary effects they have on digital microwave radios, but also provides insights to the experienced engineer. In some cases it is simply rewarding to have a tool that visually confirms principles that are difficult to believe. It is rewarding to see the sidebands of the digital microwave radio frequency spectrum regenerate when the high-power-amplifier nonlinearity is selected. Although it seems intuitive that by clipping the filtered signal the amplifier nonlinearities should do this, it is reassuring to observe it. One can also see that highly filtered signals are far less tolerant to high-power-amplifier nonlinearities than less filtered signals, and can appreciate the significance of such principles without having to get bogged down in the reasons that they are so. It is exactly this kind of accessibility that I•Q Tutor was designed to provide.

Authors

April 1986

4 — Data Acquisition System

Danny J. Oldfield



Born in Huntsville, Alabama, Dan Oldfield served in the U.S. Army and then studied electrical engineering at the University of Tennessee (BSEE 1978). With HP since 1978, he has contributed to the development of the HP 54100A Digitizing Oscilloscope and the HP 1727A Oscilloscope. He's now a process engineering manager for printed circuit board assembly at the Colorado Springs Division. He is the author of a paper on measuring metastable-state D flip-flops and coauthor of an *HP Journal* article on the HP 1727A Oscilloscope. Dan lives in Colorado Springs, Colorado with his wife and daughter. He's a member of Big Brothers and enjoys woodworking and antique cars. He currently owns a 1957 Corvette that is in mint condition.

Kenneth Rush



With HP since 1979, Ken Rush holds a BS degree in aerospace engineering (1973) and an MS degree in electronic engineering (1975), both awarded by the University of Tennessee. He is an expert in the design of high-speed ICs and thick-film circuits and was the chief analog circuit architect for the HP 54100A/D Digitizing Oscilloscope. His previous professional experience was with Oak Ridge National Laboratory, where he designed specialized electronic instruments. Ken was born in Winchester, Tennessee and now lives in Colorado Springs, Colorado with his wife and two children. He is interested in the natural history and the early exploration of the Rocky Mountains and in native American culture. He builds muzzle-loading guns, tans animal skins, and makes buckskin and fur clothing.

11 — Probe System

Arnold S. Berger



Born in Brooklyn, New York, Arnie Berger is a project manager at HP's Logic Systems Division. He holds a 1966 BS degree and 1971 PhD degree in materials science. Both degrees were awarded by Cornell University. With HP since 1979, he was project leader and hardware designer for the HP 54300A Probe Multiplexer and has worked on storage CRTs, including the HP 1727A Oscilloscope. Before joining HP, he did postdoctoral research at Cornell University and held research positions at Ford Motor Company and Argonne National Laboratory. He is the author of 11 scientific papers. Arnie lives in Colorado Springs, Colorado, is married, and has one daughter. His wife is an HP software engineer. He is a part-time instructor at a local college and enjoys skiing, jogging, bicycling, and home computers.

William H. Escovitz



Bill Escovitz came to HP in 1979 and has worked on CRTs for oscilloscopes and displays. More recently he designed the HP 54003A 1-MΩ Probe and the HP 10032A 100:1 Probe for the HP 54100A/D Digitizing Oscilloscope. He studied physics at Dartmouth College (AB 1968) and at the University of Chicago (SM 1973 and PhD 1979) and was an instructor at Thetford Academy, Vermont, before joining HP. He is the coauthor of 14 papers or conference presentations on field ionization microscopes, imaging biological specimens, and other technical subjects. He's a member of the American Physical Society and the IEEE. Born in Pittsburgh, Pennsylvania, Bill lives in Colorado Springs, Colorado. He is married and has two sons.

Kenneth Rush

Author's biography appears elsewhere in this section.

20 — Waveform Graphics

Rodney T. Schlater

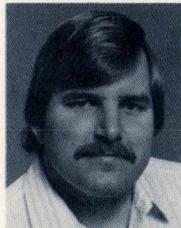


Rod Schlater received undergraduate degrees in physics in 1971 and in electrical engineering in 1973 from the University of Colorado. He completed work on his MSEE degree from the University of Colorado in 1983. Before coming to HP, he did signal process-

ing work and digital design for radar systems. Since joining HP in 1978 Rod has done hardware and software design for digital oscilloscopes. His work on variable and infinite persistence for the HP 54100A/D is the subject of a patent application. He is married and enjoys history, bridge, tennis, and travel.

26 — Trigger System

Scott A. Genther



With HP since 1981, Scott Genther was responsible for the high-speed digital design of the HP 54100D Digitizing Oscilloscope. Earlier he contributed to the design of the HP 1980A/B Oscilloscope. He is a graduate of the University of Tennessee, completing work for his BSEE degree in 1979 and for his MSEE degree in 1982. Scott was born in Brecksville, Ohio and now lives in Colorado Springs, Colorado. He coaches youth soccer and likes participating in team sports. A former All-American track athlete at the University of Tennessee, he enjoys playing soccer, softball, volleyball, and basketball.

Eddie A. Evel

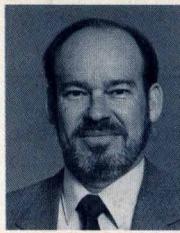


Since coming to HP in 1965, Ed Evel has designed a number of oscilloscope circuits, including the basic trigger system and two of the trigger hybrids for the HP 54100A. He also designed the digitizer for the HP 1980B Oscilloscope and is named as inventor or co-inventor on four patents on circuits for oscilloscopes and operational amplifiers. Born in Ransom, Kansas, Ed earned a BSEE degree from Kansas State University in 1962 and worked on missile guidance systems at Martin Marietta Corporation before coming to HP. He lives in Colorado Springs, Colorado with his wife and four daughters and is interested in outdoor activities, including camping, river rafting, snowmobiling, and motorcycling.

33 — Thick-Film Hybrids

Derek Toeppen

Derek Toeppen is a production engineer at HP's Colorado Springs Division.

Jimmie D. Felps

Born in Johnson City, Texas, Jim Felps studied electrical engineering at the University of Texas. He served in the U.S. Air Force and did electronic design at Texas Instruments before joining HP in 1979. He designed the semicustom pulse width modulator IC

for the power supplies in the HP 54100A/D and HP 54110D Oscilloscopes. He has also contributed to the development of the HP 1980B, the HP 1727A, and the HP 1726A Oscilloscopes. He is named inventor on a patent application related to the power supply in the HP 54100A/D. Jim and his wife live in Colorado Springs, Colorado and have four children. His outside interests include hunting, fishing, and camping.

Christen K. Pedersen

Born in Salt Lake City, Utah, Chris Pedersen attended Brigham Young University and received his BSEE degree in 1984. He joined HP's Stanford Park Division the same year as a product marketing engineer and is responsible for understanding and developing applications such as I-Q Tutor for digital microwave communications products. He is a member of the IEEE and is interested in digital image processing and in computer graphics and programming. Chris and his wife live in Campbell, California. He's an avid photographer who especially likes taking large-format black-and-white pictures in the deserts of southern Utah.

Address Correction Requested
Hewlett-Packard Company, 3000 Hanover
Street, Palo Alto, California 94304

HEWLETT-PACKARD JOURNAL

April 1986 Volume 37 • Number 4

**Technical Information from the Laboratories of
Hewlett-Packard Company**

Hewlett-Packard Company, 3000 Hanover Street

Palo Alto, California 94304 U.S.A.

Hewlett-Packard Central Mailing Department

P.O. Box 529, Startbaan 16

1180 AM Amstelveen, The Netherlands

Yokogawa-Hewlett-Packard Ltd., Suginami-Ku Tokyo 168 Japan

Hewlett-Packard (Canada) Ltd.

6877 Goreway Drive, Mississauga, Ontario L4V 1M8 Canada

Bulk Rate
U.S. Postage
Paid
Hewlett-Packard
Company

0200035216&&COLLERHO
MR R H COLLINS
2732 CHEROKEE DR
BIRMINGHAM AL 35216

CHANGE OF ADDRESS:

To subscribe, change your address, or delete your name from our mailing list, send your request to Hewlett-Packard Journal, 3000 Hanover Street, Palo Alto, CA 94304 U.S.A. Include your old address label, if any. Allow 60 days.