

THE CONCEPT IS SIMPLE: CONSTANTLY LOOK FOR FAULTS BEFORE DIGITIZING THE SIGNAL; STOP ACQUISITION WHEN SOMETHING GOES WRONG. THE SCOPE'S MEMORY THEN REVEALS WHAT LED TO THE FAILURE. IT TOOK SCOPE DESIGNERS MANY YEARS TO PERFECT THIS INVALUABLE APPROACH.

DSO fault triggers reveal what went wrong

ALMOST EVERY DIGITAL-CIRCUIT designer can tell a hair-raising story about spending a moment (or a week or a month) in troubleshooting/debugging hell. Troubleshooting an intermittent failure is one of the most difficult and frustrating problems a digital-system designer faces. Elusive, intermittent glitches; “runt” pulses; noise; and slow transitions cause project overruns, delayed product introductions, field failures, and even product recalls.

And signal-fault problems such as these will only get worse as processing speeds increase, circuits shrink, and project schedules tighten. Designers no longer have the luxury of generous design margins or the time to rework circuits to make them insensitive to signal faults. Dense, high-speed designs must have high signal integrity from the outset.


Thanks to recent developments in digital-oscilloscope-triggering technology, engineers who are responsible for digital-hardware debugging and system integration have new tools for tracking down elusive signal-integrity problems. You can use these tools in directed searches for the cause of prototype failures. But before looking at faults and new digital scopes' trigger capabilities, it's worthwhile to review conventional means of troubleshooting digital circuits with suspected signal faults.

USING AN ANALOG SCOPE TO LOOK FOR GHOST TRACES

Once upon a time, the accepted way of looking for intermittent signals buried in a good signal was to use a high-speed analog scope in a darkened room. Sometimes, you even had to improvise viewing hoods or turn off the lab lights to get a look at a faint ghost of a transient signal before the main trace bloomed and burned the phosphor.

A bigger problem, however, is the scope's blind period during retrace. During the retrace interval, you can entirely miss transients. Nobody talked much about that problem until recently, however. For typi-

Figure 1



Although an ideal digital signal assumes only two voltage levels and makes transitions between those levels in zero time, real signals include narrow spikes, low-amplitude noise, undersized “runt” pulses, and signals that make slow transitions between states.

cal analog scopes, retrace time can be 1 to 10 msec, so assuming a 20-MHz clock and a fault once every second, the probability of the scope input's being active when a fault occurs is 5 to 25%. You might see the fault once every minute or so if the scope's writing speed is high enough. But if the fault occurs every 10 minutes, every hour, or intermittently, your chances of seeing the fault with an analog scope are negligible.

Over the past three decades, the remaining analog-scope vendors have improved their products with higher writing speeds, storage, and image-enhancing techniques. So now, when you look for intermittent signals, perhaps the room need not be so dark. But capturing a signal fault with an analog scope is still a matter of watching, waiting—and luck.

When faced with intermittent failures, embedded-system developers often write short software loops to cycle the code through the failure mode. A scope can then display signal activity and—maybe—the fault. This approach works only if the small loop makes the circuit behave in the same way as the actual code does and if the developers have correctly identified the of-

fending loop. An in-circuit emulator (ICE) often helps in this trial-and-error process. Some ICEs include a trigger input to allow synchronizing the emulator and the scope.

FINDING FAULTS USING OLDER DIGITAL SCOPES

Digital scopes brought many new capabilities, but there was a price to pay in troubleshooting performance. Early DSOs, with their low sampling rates, worked well on repetitive signals using equivalent-time sampling but were limited when data was intermittent or nonrepetitive. In addition, early DSOs had problems with aliasing and were difficult to use. But as sampling rates increased, digital scopes' single-shot or transient-capture capability promised to turn the instruments into useful troubleshooting tools.

Transient capture stems from digital scopes' ability to look backward in time by stopping signal acquisition after an event occurs. This capability lets you trigger on the failure and look backward through the acquisition memory for the fault that caused the failure. With early digital scopes, a deep memory captured a long-time sample of a suspect signal just before the circuit failure. You could then search for the fault by examining the captured signal in detail. To help with the search, some DSOs allow you to define an envelope of a good waveform and display anything that doesn't fit into that envelope. Other scopes have a scroll mode in which captured data scrolls by. You can stop the scrolling when you see an anomaly. Even with these capabilities, the amount of time required to find a fault in a captured signal can be enormous. And the DSO input is inactive while you are looking through the captured data.

Another advantage of digital scopes is that they can display captured signals that remain in memory. By mapping captured signals on top of each other, the scope develops a time-lapse picture of the signal. If the scope captures and plots an intermittent or abnormal signal along with a number of valid signals, the abnormal signal stands out in the display. The abnormal waveform is especially visible if the scope uses color or gray-scale grading to differ-

entiate waveforms that recur at different duty cycles.

But even with a high sampling rate, a deep memory, and a persistent display, the probability of seeing an infrequent or transient signal is quite low when you use conventional triggering. The probability is low because the acquisition system (ADC and memory) is inactive during the display-processing period. This period can last 10 msec to 1 sec in older scopes. (Most scope manufacturers specify a display-update or waveform-capture rate that is the inverse of the processing dead time.) With conventional triggering, you have only a 0.01 to 0.0001% probability on any given sweep of capturing a fault that occurs once per second in a 10-MHz signal. Even with the lat-

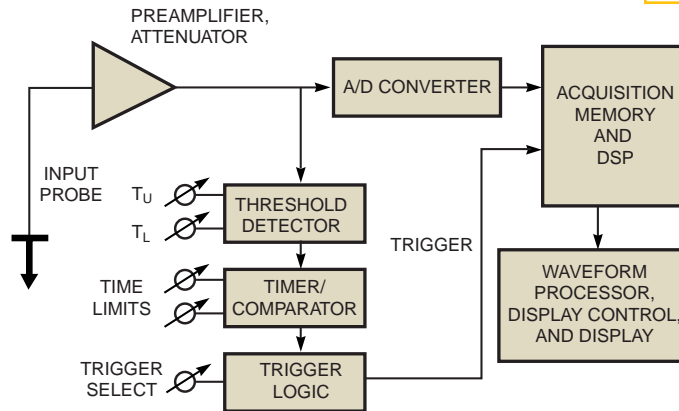
Over the past decade, DSO designers have developed powerful trigger schemes with which you can do amazing things. You can now selectively trigger on several types of fault, such as too short a pulse (a spike or glitch), too small a pulse (a runt or noise), or a floating signal (a transition-time or slew-rate fault). These triggers operate on the input signal before it reaches the sampler and ADC, so there are no effects from ADC speed, memory length, display technique, and processing dead time. The trigger circuits continuously monitor the signal and stop the ADC or freeze the acquisition memory so that the scope can capture and display transients.

Other technical innovations were more glamorous, however, and received most of the attention in technical articles, advertising,

and product literature. Tektronix focused on finding faster ways to get waveforms onto the screen. LeCroy (www.lecroy.com) pursued high-speed sampling and deep, flexible memory. HP (www.hp.com) emphasized ease of use and analog-like response. Gould (www.gouldis.com) concentrated on simulating intensity-modulated analog displays and enhancing ADC resolution beyond the 8 bits that are common in general-purpose DSOs. Triggering innovations seemed lost in the hype.

But for digital troubleshooting, these new trigger functions define a new breed of scopes that offers exceptional combinations of high-speed sampling, long memory, and advanced display technology along with fault-trigger capabilities. With these scopes, you can continuously monitor a suspect signal and trigger on and capture a signal fault in real time. Prices range from \$7000 to \$40,000 (Table 1).

You can classify fault triggers according to three general fault definitions that cover virtually all common types of digital signal faults. But, before looking at the triggers themselves, consider the ideal digital signal. An ideal digital signal instantly changes from one state to another and stays in the



The new generation of fault trigger continuously monitors signals. When a fault condition that you specify arises, the trigger circuit stops the flow of new information into the DSO's acquisition memory. A history of the signals that preceded the fault remains in memory.

ent display-update technology, the probability of the fault's occurring during any sweep is less than 25% under the most favorable conditions.

With all of these problems, few engineers tried to use early digital scopes to troubleshoot intermittent failures. Several years ago, a Tektronix (www.tek.com/Measurement) DSO designer wrote: "Troubleshooting digital circuits poses difficult triggering challenges. If an oscilloscope cannot trigger, the waveform cannot be examined closely. And the signal display most needed to identify a problem often is the one on which the oscilloscope cannot trigger." With fault-trigger schemes now available, troubleshooters or system integrators have a variety of tools to trigger on the exact problem for which they are looking.

new state for at least one clock cycle. The difference between states depends on the nominal circuit threshold, TN, which is a function of the type of logic (TTL, CMOS, or ECL). TN, however, is not a single absolute value; it can vary from component to component and from area to area on the circuit board. Each logic family has a defined level above which the state is clearly high, TU, and another level below which the state is clearly low, TL. Because its value never falls between these levels, the ideal digital signal never assumes a state that is open to question. Therefore, the ambiguity of TN is unimportant.

No real “digital signal” meets these ideal criteria, however. A digital signal is actually an analog signal that obeys three simple rules (Figure 1 and Table 2).

- The signal must spend a minimum amount of time in either the high or the low state before making another transition. A violation of this rule is called a “spike,” a “glitch,” or a “pulse-width fault.”
- The signal must move from the high state through the middle state to the low state or vice versa; it should not cross a threshold from the high or low state into the middle state and then return to its original state. A signal that violates this rule is called “noise” or a runt pulse.
- The signal must spend most of its time in either the high or the low state. If the signal spends too much time between states, it is said to “float.” A float fault is the most general form of “transition-time” or “slew-rate” fault.

These three rules coincide with the three continuous-time fault triggers now available on high-end DSOs. Runt and float-fault triggers require dual thresholds, although a spike or pulse-width trigger

requires only a single threshold. (Dual thresholds make the spike trigger even more powerful, however.) In all three cases, fault detection precedes the digitizer (Figure 2) and is independent of the scope’s dead time and ADC sampling rate. These real or continuous-time fault detectors depend only on the order and timing of the threshold crossings. The triggering is true continuous-time triggering because, once you activate the fault detectors, the detectors monitor the signal and wait for a trigger condition. The trigger event stops the storage of digitized data, and the

FAULT-TRIGGER CONDITIONS ARE MUTUALLY EXCLUSIVE, AND YOU MUST SELECT A TRIGGER MODE BEFORE YOU INVOKE THE FEATURE.

scope’s memory retains the signals that immediately preceded and followed the event. In fact, with these trigger circuits, the scope can trigger on faults that are too brief to display because the sampling rate is too low.

As currently implemented, fault-trigger conditions are mutually exclusive, and you must select a trigger mode before you invoke the feature. Thus, you must have some idea of the type of fault you are looking for before you proceed. Each scope manufacturer suggests ways to identify the fault before you select the trigger type. If, however, all three triggers are active simultaneously, the scope would function as a real-time fault monitor, which would trig-

ger on all types of fault. The user would need only to specify a well-behaved signal’s thresholds, allowed pulse width, and rise/fall time.

SPIKE AND GLITCH TRIGGERS

The most common type of digital-signal fault is a spike—a narrow pulse that goes from one state to another and back again in less than a specified time period. A spike is usually the first fault you suspect as the cause of a circuit failure. Circuits in the system under test can mistake a spike for a clock edge. Such an error can throw off timing. A spike can also cause an unwanted state change in a counting or control circuit.

Spikes and glitches have several causes, such as malfunctioning components, metastable conditions, and switching transients coupled from other areas of the circuit or even from other equipment. Intermittent connections or bad solder joints can generate spurious signals. A failed bypass capacitor that does not filter transients as intended can allow spikes to affect the circuit. A spike can result from a race condition in which two signals simultaneously vie for control of a signal line; the slightly different timing on the signal edges produces a spike. Crosstalk can also cause spikes, but a noise- or runt-pulse trigger detects crosstalk more effectively than a glitch detector does.

Varying environmental conditions and differences in component values among nominally identical circuits can cause spikes of the same amplitude and width to produce different effects at different times. This variability gives spikes and glitches their reputations for intractability.

Early-generation DSOs and logic analyzers employed a simple “glitch trigger” that used a peak-hold circuit to detect spikes that occurred between sample peri-

TABLE 1—DIGITAL SCOPES WITH THIRD-GENERATION FAULT TRIGGERS

Company	Tektronix	LeCroy	HP	Gould
Model series	TDS 500/600/700	LC500	54800	Classic 6000/9500
Bandwidth (MHz)	500 to 1000*	1000	500 to 1500*	200 to 400*
Maximum sample rate (G samples/sec)	2 to 5	4 to 8	1 to 8*	0.1 to 2*
Maximum memory (k points)	15 to 500	250 to 16,000*	32 to 64*	50 to 1000*
Fault triggers	Yes*	Yes	Yes	Yes
Minimum trigger event (nsec)	0.6 to 2	0.6	0.5 to 1	2.5
Variable persistence	Yes	Yes	Yes	Yes
Display update rate (waveforms/sec)	150 to 200,000*	150 to 40,000*	300 to 1950*	No spec
Dead time (μsec)	6700 to 5*	6700 to 25*	3300 to 513*	No spec
Price (1000)	\$10 to \$35	\$20 to \$40	\$10 to \$30	\$7 to \$12

*Available on selected models, as an option, or in selected modes only.

ods. These glitch detectors first appeared in logic analyzers to make up for the loss of signal-integrity information. Because the critical information is the state and timing relationship among signals, logic analyzers discarded analog information.

THE ADC NEVER SLOWS DOWN

In more recent digital scopes with higher sample rates, the ADC runs at full speed in all modes. To conserve memory at slow timebase settings, these scopes store every nth sample. Between stored samples, however, some scopes record the maximum and minimum values. This function, called peak- or min/max detection, is most useful when you are looking for narrow pulses in slowly changing waveforms. This mode is less effective on faster timebase settings, especially when glitches are infrequent, random, or intermittent.

Because they do not operate during the dead time when the scope processes the acquired signal for display, neither of these early glitch detectors is a continuous or real-time trigger. The glitch- or peak-detect function does not operate when the ADC is inactive; therefore, the scope does not catch glitches that occur at that time. Furthermore, these peak detectors often miss spikes that occur near valid transitions.

In the new spike- and pulse-width-trigger modes, the DSO's ADC runs continuously until the scope detects a trigger. At that point, the scope ceases to allow new samples into the acquisition memory. In this way, the glitch and other related signals remain in memory. Fault detection starts when the signal crosses a threshold and a timer or counter starts. The occurrence of a second crossing within a specified time initiates a spike or pulse-width trigger. For a spike fault, the trigger occurs on the last threshold crossing that defines the fault.

To set up the spike or pulse-width trigger, you must specify a threshold level, a minimum-acceptable pulse width, and a positive or negative spike condition (or both in some cases). The variable threshold and the choice of a positive or negative spike let you selectively look for high or low spikes, marginal spikes, and overshoot. In some cases, you can specify minimum and maximum pulse durations so that you can

look for pulses that are too long and too short. Some scope manufacturers refer to this mode as an "exclusion trigger" because the scope triggers on anything except a pulse of normal or acceptable length.

Some of the new scopes have both glitch and pulse-width triggers. Although these functions are essentially the same, you set them up in different ways.

NOISE AND OTHER TRIGGERS

High-speed signals coupled from one line to another can cause low-level noise, static, or crosstalk. Although the spike trigger with appropriate threshold settings can capture low-level spikes, the noise- or runt-

sions of the same circuit may respond differently to runt pulses, or the same circuit may respond differently at different times.

Usually, the circuit under test ignores noise or runt pulses that almost cross the threshold. Although most circuits tolerate a small amount of noise as long as the signal remains outside the indeterminate zone, the noise can become large enough to cause a transition under some conditions. Moreover, noise that doesn't normally cause problems can become a problem as conditions change. For example, supply variations can intermittently cause or prevent a runt pulse from crossing the logic threshold. The intermittent nature of noise

and runt pulses places them among the most difficult faults to deal with.

Two transition levels completely define the noise or runt trigger, although some DSOs let you specify a minimum or maximum time so that you can differentiate between noise and a longer diminished-amplitude pulse. To set up this trigger mode, you specify two threshold levels, the pulse sense (positive or negative) and, optionally, the minimum or maximum

time period. As with the spike fault, the noise or runt-fault trigger occurs on the last threshold crossing that defines the fault.

OVERSHOOT AND RINGING

An improperly terminated or uncompensated line is the usual cause of overshoot and ringing. Although overshoot itself is seldom much of a problem, a severe down cycle following the overshoot can appear as noise or a spike. Generally, overshoot and ringing result from hardware or design faults that are not intermittent. Because overshoot and ringing are repetitive, conventional scopes usually have little difficulty detecting them. The problems are visible as you probe the circuit after you first apply power.

Nevertheless, pulse-width and noise triggers have their uses in finding overshoot and ringing. You can use a spike trigger to detect overshoot: Set the threshold above the nominal high-signal level (or below the nominal low-signal level) and set the pulse width at a multiple of the rise time. You can detect the second half of the

TABLE 2—SIGNAL FAULTS AND THEIR TRIGGERS

Fault type	Trigger type		Float
	Pulse width	Noise/runt	
Large spike	XXX		
Pulse from race condition	XXX	xx	
Overshoot and ringing	XXX	xx	
Small spike	xx	XXX	
Crosstalk	xx	XXX	
Runt pulse	xx	XXX	
Ground bounce	xx	XXX	
Floating level		xx	xx
Slow rise/fall time			XXX
Droop			XXX

XXX=best trigger for fault.

xx=works for fault capture under some conditions.

pulse trigger is more convenient. And, because two thresholds and the crossing rule completely define the trigger conditions, a timer need not add uncertainty.

Runt or dwarf pulses are those pulses whose widths are similar to those of normal pulses but whose amplitudes fail to cross the nominal circuit threshold. Such pulses do not cause the circuit under test to change states. An occasional runt pulse can cause frustrating problems by failing to clock a flip-flop as intended. A runt pulse can occur when one gate tries to pull a line down, while another gate tries to pull it up. Or perhaps too large a fan-out or a faulty component causes excessive loading under certain conditions. An improper line termination can cause a reflection that results in a reduced-level pulse at a third node. Closely coupled crosstalk and ground bounce can also couple a signal to another line, albeit at a lower level. The data signal coming from a disk- or tape- storage device could develop runt pulses as the head or storage medium ages or becomes contaminated. Because of varying parameters, ver-

overshoot cycle if you set the runt detector's upper threshold just below the nominal pulse level or if you set the lower threshold just above the nominal level.

A slow transition can cause timing-uncertainty errors in many systems. A bad connection, a missing pull-up resistor or bypass capacitor, a high-impedance trace, or a marginal supply voltage can cause these problems. In some cases, changing a component value changes the circuit parameters so that faster or slower transitions cause the circuit to fail.

Some designs allow unenabled data lines to float. Floating lines can assume undefined states, but such states are generally not a problem when you use proper data-qualification techniques. If, however, the float condition occurs when environmental conditions or an inadequate design margin allow a control signal to drift into the undefined zone between TL and TU, even a slight perturbation can cause an unwanted state change. Such a condition may first show up as a glitch in a logic circuit, but the root cause is the floating signal line.

The most general type of transition fault is a float condition that occurs when the signal remains between states longer than a defined time interval, regardless of what occurs next. The high and low thresholds and a maximum allowable transition time define the time interval. A good rule of thumb for the maximum transition time is 10% of the clock-pulse width, although this may be too long for some circuits. The float- or transition-time trigger circuit consists of two threshold sensors and a timer. The first threshold crossing (from a valid state to the floating state) activates the timer. The second threshold crossing (from the floating state to the other valid state) stops the timer. The float-fault trigger occurs when the signal remains in the floating state for a longer time than specified. Some transition-time triggers are independent of the second threshold crossing. You can use them to find a long runt pulse as well as a slow rise/fall time. Other trigger circuits wait until the second threshold crossing to trigger.

To set up a float-trigger condition, you specify the high and low thresholds and an allowable time period. Some scopes let you specify either too long or too short a transition time. The short-interval trigger can help you identify the fast-changing signals that sometimes underlie crosstalk problems. Also, some scopes indicate the effec-

tive slope or slew-rate that the selected differential voltage and transition time imply.

Some scopes offer additional trigger functions that are useful when you face certain problems. A dropout or time-out trigger is useful when a line goes completely dead or never becomes active after you enable it. Although it is most useful in data-communications and network troubleshooting, this trigger mode can help in prototype debugging. You might also be able to use this trigger mode to detect a supply-voltage dropout, a processor crash, or a local- or wide-area-network hang-up.

Pattern triggers were first available in

IN SELECTING A SCOPE, CONSIDER THE RESOLUTION AND ACCURACY OF THE TRIGGER-CIRCUIT THRESHOLDS AND TIMING FUNCTIONS AND THE PERFORMANCE OF THE ARMING, QUALIFYING AND HOLD-OFF FUNCTIONS.

logic analyzers and have appeared in some multichannel scopes. Obviously, the number of channels that the scope provides limits the pattern width. The best implementations of a pattern trigger are in Tektronix's 16-channel logic scope and HP's mixed-signal scope, which has 16 logic-analyzer channels and two scope channels. Neither of these scopes has the fault-trigger capabilities that this article discusses, however.

Although a setup-and-hold violation is not strictly a signal fault, it can cause problems in your prototype. A scope with a setup-and-hold-violation trigger can help you to find such faults when you know which signal lines are suspect.

PRACTICAL CONSIDERATIONS

Although the new fault triggers are continuous-time monitors, the DSO's rearm and display-processing times may prevent the scope from capturing the second of two faults that occur in quick succession. After detecting the first fault, the trigger circuit is insensitive to additional faults for the

small amount of time it requires to rearm itself.

Even more significant, however, are the scope's memory length, memory segmentability, and display-processing time. If you can segment the scope's memory so that it stores multiple triggered events before displaying them, the scope can capture several fault events in quick succession (allowing for the trigger dead time). On the other hand, if the scope must prepare an image for display after each fault, the display-processing time must elapse before the scope is ready to acquire the next fault.

Typically, these dead times aren't much of a problem, because most troubleshooters look for faults one at a time, fixing each one before proceeding to the next. However, you should be aware of these dead times; if you want to know how long they are, you need to measure them, because scope manufacturers do not completely specify dead-time durations. Moreover, dead time can vary with setup conditions.

Interaction between the scope's sample rate and timebase can cause another type of problem: The new fault triggers have a typical time sensitivity or threshold of 500 psec to 3 nsec, so they can detect very fast faults. If the sample rate is too slow, however, the scope may be unable to display the fault that caused the trigger. Keep in mind that the effective sampling rate of most DSOs is a function of the timebase setting, so you don't always get the benefit of the instrument's advertised maximum sampling rate. For example, a 50M-sample/sec real-time sampling rate yields a 20-nsec sampling resolution. At this sampling rate, a scope can miss or distort the details of a transient shorter than 50 to 100 nsec. On slow timebase settings, you may have to use peak detection to get an indication of the fault.

THE IMPORTANCE OF SUBTLE DIFFERENCES

In selecting a scope, consider the resolution and accuracy of the trigger-circuit thresholds and timing functions and the performance of the arming, qualifying, and hold-off functions. Ask the scope manufacturer whether you can use the trigger circuits together in an AND or OR fashion. Data sheets rarely provide operational details of such features, despite the features' importance in some applications.

In all scope measurements, keep in mind the effect of your probe on the circuit and signal that you are testing. Too low a probe

impedance can load your circuit, cause a fast spike to drop in amplitude, or slow a fast rise time. An improperly compensated probe can distort waveforms and can cause your scope to miss faults that your circuit sees or cause a rise-time or runt-fault trigger that your circuit doesn't see.

MAKE NO MISTAKE

You can mistake some types of software errors for signal faults. For example, a software bug—the failure to set an enable flag—can cause a failure to read data into memory. However, a runt pulse or a spike, which might result from simultaneous requests for control of a bus line, can cause the same thing. In these cases, you might find it useful to have the DSO trigger a logic analyzer or to activate an ICE breakpoint. With the additional instruments, you can observe the hardware states or the code that preceded the fault. Be sure to account for the delay time from the displayed waveform to the trigger output; to measure the delay, connect the trigger output to a second scope channel.

Many of the factors that you need to

consider in choosing a scope depend on your preferences and your application. You can make a wise choice only after a period of using several competitive instruments in your lab. In particular, you can assess some of the nuances of the new trigger schemes only after an extended trial. Similarly, choosing between the high-speed capture and display techniques of some scopes and the deep memory and memory segmentation of others requires careful evaluation in your application.

CAREFUL EXAMINATION

When you use an oscilloscope without an appropriate fault trigger, you must carefully examine the signal, hoping to spot a rule violation. Although a long-persistence display and deep memory can help to capture the signal for later review, locating faults in this way consumes a surprising amount of engineering time. Although the capabilities you need for efficient signal-fault troubleshooting exist only in new, high-end DSOs, the time such scopes save soon repays the units' \$7000 to \$40,000 price. □

Authors' biographies

The authors are principals in Farnbach and Swift Technology LLC (FST), a consulting and technology licensing company. FST holds US patent 4,965,800, which covers a portion of the technology discussed in this article. In addition, the authors are employed independently.

William A Farnbach is a senior member of the technical staff at Rockwell Semiconductor Systems (San Diego) where his responsibilities include digital-circuit design and the development of custom ASICs. Previously, he was chief technical officer of Wavetek (San Diego), and he was a design engineer at Hewlett-Packard (Colorado Springs, CO). He holds six patents.

William E Swift is president and principal consultant of TMI (Thousands Oaks, CA), a strategic marketing company that serves the electronics industry. Previously, he worked in engineering and management at Fluke Corp (Everett, WA) and several other major electronics companies.