 By Colin Shepard, Tektronix Inc

READ BETWEEN THE LINES OF BANNER SPECIFICATIONS–BANDWIDTH, SAMPLE RATE, AND RECORD LENGTH–TO DRILL DOWN TO THE NUANCES AND LESS GLAMOROUS FEATURES THAT AFFECT EFFICIENCY AND EVEN THE VALIDITY OF YOUR DESIGN.

Evaluating oscilloscopes: Dig deeper

For years, electronic-system performance increased in step with steadily increasing semiconductor integration and functions, simulation-model improvements, architectural changes, and more. Through all these increases, interdevice-signaling speeds and techniques did not change significantly, because the I/O-signaling architectures were doing an adequate job, and the underlying technologies for change were not in place.

In the past half-decade or so, engineers have focused more on LVDS (low-voltage differential signaling) to achieve dramatic increases in system performance. Data rates have increased by an order of magnitude, driving wider adoption of complex serial protocols for communication among devices. These protocols include PCI Express, InfiniBand, and XAUI (10-Gigabit attachment-unit interface). These environments encompass a variety of data rates and transmission architectures, but all of them share a need for rigorous design and verification practices.

This need has made test equipment, such as oscilloscopes, more important than ever. Engineers rely on oscilloscopes to analyze the performance of their serial-device designs and to support verification and debugging work. Their tasks include precise parametric measurements, troubleshooting, and signal-integrity analysis. Later in the development process, they turn to oscilloscopes to produce eye-diagram displays for compliance testing.

Too often, an engineer choosing an oscilloscope considers only the specifications that appear in the headlines of product brochures and magazine ads. The best known of these banner specs are bandwidth, sample rate, and record length. Though important, these dimensions of oscilloscope performance provide an incomplete picture of the instrument's effectiveness in daily real-world use. The bandwidth figure, for example, states only the oscilloscope's gross frequency range; it says little about the instrument's ability to reliably detect and capture fast anomalies.

Consequently, it is important to read between the lines of the banner specifications when evaluating an oscilloscope. This advice actually has two contexts. First, it is a caveat to drill down to the nuances that make up heavily promoted specifications. Second, it is a reminder to study the features that, though less glamorous than vendors most often tout in the market, can substantially impact the efficiency and even the validity of the designer's work.

DEFINE AND REDEFINE BANDWIDTH

Bandwidth specifications are legitimately important. To a designer pushing the limits of high-speed serial-bus architecture, pure bandwidth is at the top

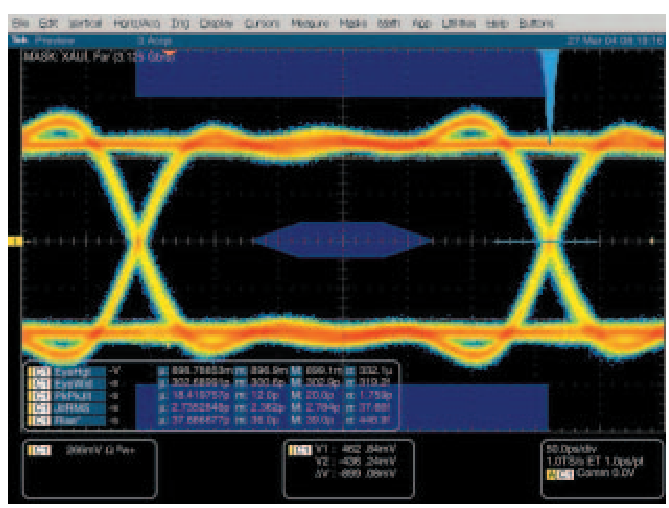

Figure 1 Eye-diagram analysis depends on oscilloscope banner specs, such as bandwidth. But it also relies on clock-recovery features, probing, triggering, and more. Looking beyond the banner specs is the best way to ensure efficient, accurate measurements.

of the shopping list for oscilloscope features. But bandwidth itself is simply a specification describing the frequency response of the instrument, the frequency at which a sine wave rolls off by −3 dB. Two oscilloscopes having the same bandwidth rating can have different rise times and disparate responses to complex wave shapes. Is there a between-the-lines specification or feature that might better inform the purchaser's decision? There are two facets to the answer to this question. One is the oscilloscope's true rise-time performance. The other is the instrument's behavior in its DSP modes.

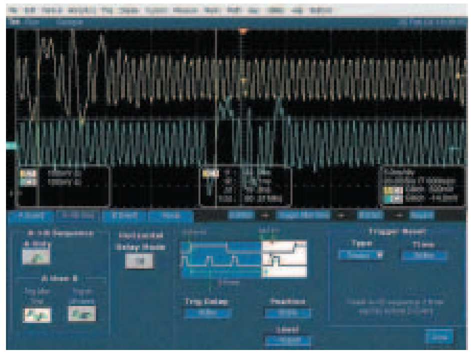

Figure 2 **A dual triggering screen shot shows a high-speed serial-lane skew violation. Note the trigger-reset conditions; the sequence restarts itself if a time limit expires before the second event occurs.**

The analog rise time is a function of the oscilloscope's bandwidth. It is tempting to simply calculate the rise time from the bandwidth using a textbook formula, and this calculation is the basis of some published rise-time specifications. An objectively measured rise time is a better foundation for measurements, either with or without DSP enhancement. Every engineer understands the importance of rise-time response. Weighing the difference between measured and calculated rise times is an example of reading between the lines.

You can use DSP filtering to extend the oscilloscope's net bandwidth, flatten its frequency response, and provide better matching between channels. These capabilities are key when the device under test employs a high-speed, multilane, serial-transmission environment. But DSPs can introduce errors that tend to increase in proportion to the frequency extension beyond the true analog bandwidth.

You should use DSP when measuring less-than-200-psec rise times or eye diagrams (**Figure 1**). In these cases, it's essential to extract the maximum bandwidth from the oscilloscope. Clearly, this need favors the DSP approach. The fastest measurements almost always demand the highest bandwidth. But it is sometimes useful to have a way to bypass the DSP extensions and make measurements using only the instrument's innate analog bandwidth and rise time. For example, some researchers use application-specific DSP algorithms and need to work with raw data from the oscilloscope. In such cases, a DSP-bypass feature is important. This kind of specification may not make headlines, but it is a factor in choosing a high-performance oscilloscope.

SIGNAL COMPLEXITY, NOT JUST SPEED

The phrase "high-speed measurement" has all kinds of implications about less-than-1-nsec edges and blindingly fast clock rates. Designers sometimes overlook the fact that these high-speed measurements are often high-complexity measurements. Capturing one symbol in a data stream can involve judgment, luck, estimation, guesswork, or the right choice of triggering features.

Oscilloscope triggering—just as much as the bandwidth and sample rate—determines what you can capture, view, and measure with the instrument, just as much as the bandwidth and sample rate do. The trigger system has its own set of specifications. The trigger path is typically a tributary of the main input-signal path and should embody many of the same environmental attributes: sensitivity, jitter, and more. Another dimension of trigger performance is the range of trigger types—the conditions that you can impose to define when a trigger occurs.

Of course, the trigger system has its own banner specs. A designer choosing an oscilloscope to measure fast serial signals might assume the trigger path has the same bandwidth as the instrument's specified bandwidth. In truth, the term of interest is "trigger sensitivity." This specification embodies one simple question: What are the amplitude requirements when capturing signals near the top of the frequency range? In many oscilloscopes, the trigger sensitivity does not match the analog-acquisition bandwidth.

Even though the normal content of a signal may be comfortably within the trigger's performance specifications, you might not detect narrow glitches or truncated pulses if the trigger sensitivity at high speed is inadequate. Fortunately, innovations such as SiGe (silicon-germanium) trigger-circuit topologies have begun to overcome this limitation. Engineers commonly regard the oscilloscope's triggering features as a given and assume that the edge and glitch triggers they've always used will suffice. In truth, triggering flexibility stands alongside the instrument's top banner specs when it comes to doing real work.

Every oscilloscope has edge triggering, and most higher end instruments have "advanced" triggers, as well. Edge triggering detects a simple voltage-threshold crossing, and advanced triggering applies many more qualifications pertaining to voltage, timing or logic conditions, and more. Advanced triggers are becoming increasingly important in serially transmitted digital signals.

In some cases, an advanced trigger setup may be the only way to trigger on the signal of interest. A designer working with a multilane InfiniBand device, for example, must ensure that the lanes are time-aligned within specific tolerances—that is, they must not only comply with the standard, but also function properly.

The normal way to approach this measurement challenge is to trigger on a single character in one data stream and measure the skew, or time shift, among the lanes. The measurement results summarize the skew value at one instant. The results are often good for the time being but insufficient for predicting stable operation in the long run.

Recently, oscilloscopes with full-featured dual triggering have dramatically simplified the complex task of observing these skew changes over time. You can define two advanced triggers from the full menu of trigger conditions. After the data character fires the first trigger, a secondary trigger can look for skew errors for a period of time or rearm the first trigger to start the search again (**Figure**

2). The setup can wait for days, if necessary, for an error combination to occur.

Trigger specifications are rarely the first order of business when evaluating an oscilloscope. But the trigger system is a full partner in detecting and capturing complex or intermittent events. The time you save making an unattended long-term skew measurement more than makes up for the time you spend reading between the lines of the trigger specifications.

MORE "SECONDARY" SPECIFICATIONS

The banner specs of bandwidth, sample rate, and the like have overshadowed the specifications this article has so far discussed. These "secondary" specs are not alone in this respect. Designers often treat many other parameters as secondary issues during the oscilloscope evaluation, but these issues can potentially cause you to make or break an urgent engineering deadline.

Embedded clock recovery is the foundation of oscilloscope-eye-diagram analysis for many serial standards and also supports measurements such as the CDR (clock-data recovery, **Figure 3**). The designer working with embedded clock signals needs to look beyond banner specs and ponder what the oscilloscope can do to make clock recovery faster, easier, more flexible, and more repeatable.

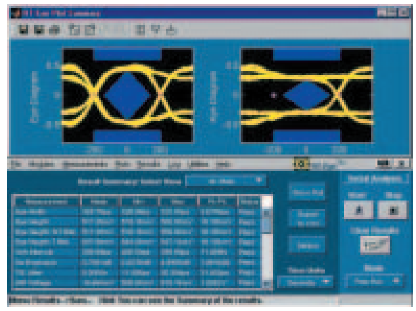

Figure 3 Eye-diagram-analysis applications simplify complex measurements but rarely appear in the list of oscilloscope banner specs.

As always, the needs of the application guide the designer's choice. You need to ask whether you will use the oscilloscope for troubleshooting or for compliance measurements, what clock-recovery mechanisms are available, and whether the oscilloscope can recover the clock in real time, enabling it to display dynamic eye-diagram behavior.

Most high-end oscilloscopes offer either software- or hardware-based clock recovery. You produce software clock recovery from stored acquisition data. The software approach is a well-accepted tool for conformance testing using applications such as Tektronix's TDSRT (real-time)-Eye automated compliance and analysis software.

You can use PLL-based clock recovery to acquire real-time eye diagrams, but herein lies another between-the-lines question: Does the PLL, which may be software- or hardware-recovered, adapt to the evolving clock frequencies in today's serial standards? Some do; some don't. It pays—literally—to understand the difference.

Eye-diagram measurements exemplify some of the most complex procedures that a design engineer needs to perform with an oscilloscope; another example is jitter measurements. In both cases, the designer can benefit from the domain expertise of application software running on the oscilloscope. The software tools minimize the learning curve and dramatically reduce setup, measurement, and analysis time (**Figure 3**). Yet, these tools never appear on a list of banner specs. You must do your homework, look beyond the glamorous specs at the top of the list, and make sure the right tools are available.

BEYOND PROBE BANNERS

Probing opens yet another spec discussion. All of the acquisition and analysis features that this article has so far explained depend on faithful transmission of the signal between the device under test and the oscilloscope itself. Many new high-speed-interface standards use differential signaling rather than the more familiar single-ended communication.

Although probing approaches have banner specs of their own, especially in bandwidth and loading, it is also important to understand the implications of the oscilloscope and the probe working together as a system. Does the oscilloscope system offer true differential-probing tools? If it does not use these tools, you need to use two single-ended probes and onboard math, precluding certain types of measurements. Moreover, issues such as common-mode rejection, sensitivity, response accuracy, and noise floor all affect the probe's impact on the signal. Small differences in these parameters can cause large probe-load distortions when you're measuring today's high-speed signals.

Probe-attachment methods rarely get top-tier visibility among an oscilloscope's banner specs, but they are critically important in every measurement. Some vendors fit devices under test with SMA test points, whereas others require access to individual pins on tiny surface-mount devices. Determine whether the oscilloscope's family of probes embraces all these needs.

In short, banner specs remain the accepted standard of comparison among oscilloscopes. But savvy engineers read between the lines to scrutinize the underlying performance that affects everyday tasks. If you take them as a whole, some less prominent specs in acquisition, triggering, analysis tools, and probing can be just as important as the big numbers describing bandwidth, sample rate, and record length. It's all about doing the job, and even the humblest specification can sometimes make the difference between success and failure.□

AUTHOR'S BIOGRAPHY

Colin Shepard is product-line manager of the performance-oscilloscopes product line at Tektronix Inc.

TALK TO US

Post comments via TalkBack at the online version of this article at www.edn.com.