

BY DAN STRASSBERG • CONTRIBUTING TECHNICAL EDITOR

SCOPES:

more than
meets the eye



MODERN DIGITAL SCOPES NOW DO MUCH OF THE HEAVY LIFTING IN MEASUREMENT AND ANALYSIS. BUT SUCCESSFUL USE OF THESE ADVANCED CAPABILITIES REQUIRES DOING YOUR HOMEWORK.

People often say that EEs are almost obscenely fortunate to have a tool that provides as much insight into fundamentally invisible processes as do oscilloscopes into the internal workings of electronic circuits and systems; no other profession has a tool that reveals as much. Despite the embarrassment of riches that scopes afford their users, manufacturers continue to find ways to make the instruments more valuable. Unquestionably, the old cries of “faster” (referring both to bandwidth and sampling rate), “deeper” (referring to depth of acquisition memory), and “less costly” continue to motivate scope designers. But the ways to make scopes even more useful are

growing—seemingly just as fast as are bandwidth, sampling rate, and memory depth.

Over the past few years, scopes’ analytical and computational prowess has shown no signs of slowing its ascent. Adding analytical capabilities is, however, only part of the challenge of designing computationally intensive oscilloscopes. Another important part is ensuring that new capabilities of mind-boggling sophistication don’t actually boggle the minds of the target users. A scope is probably better off without features that are so difficult to operate that users give up trying to make them work. Scope designers often liken their progeny to motor vehicles and refer to the panoply of usability issues under the heading of “How an instrument ‘drives.’”

As important as scopes are in EEs’ jobs, most engineers still regard the instruments as mere tools—adjuncts to accom-

plishing the task at hand, *not* the objects of the work. Greater ease of use both responds to and encourages this attitude; when you can make a measurement without giving the technique much thought, it is comforting to believe that the procedure merits little thought. Moreover, in this era of constricted schedules and budgets, there is rarely time to think about problems that seem peripheral to completing a job. Alas, such thinking can be dangerous (see **sidebar**, “Calibrating scopes’ high-frequency amplitude accuracy: more difficult than you might think”). Modern scopes make inherently difficult measurements seem easy, but, all too often, the measurements are less straightforward than they appear. Failure to recognize this fact and to understand the instrument and the technique can lead to erroneous or meaningless results—whose lack of validity can go unrecognized until the consequences

become painfully obvious and corrective action is prohibitively expensive.

A FOOL’S ERRAND

Becoming enough of a scope expert to select the best unit for your application and to use the instrument in the most advantageous possible way requires effort. Indeed, some say that attempts to find the best scope or to most effectively use it are fools’ errands. To begin with, no two engineers are likely to agree on definitions of “best” and “most advantageous” in the context of selecting and using scopes. Second, data sheets, the principal presale documents by which engineers select scopes, have become voluminous, sometimes exceeding 30 pages packed with footnotes and fine print. Third, many midmarket scopes and nearly all high-end units are now PC-based, which usually means based on standard versions of Windows. In such instruments, a Windows-based software application determines how you access the multitude of scope features.

The complexity of scope applications is certainly at least comparable with that of common shrink-wrapped office-software applications, such as Word and Excel from Microsoft (www.microsoft.com). Most office-application users avail themselves of only a small fraction of the software features. So it is with scope users. Moreover, a common problem for many scope users is that they don’t use the instruments every day, yet, when they slide into the driver’s seat, they need to quickly get answers to their questions about the unit or device under test. In other words, the methods of accessing

AT A GLANCE

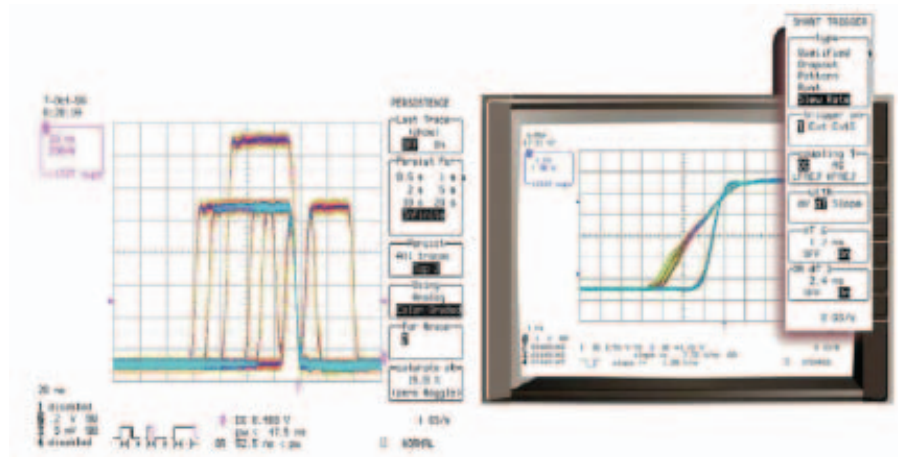
Especially when working with high-speed serial buses, engineers need active differential probes. Scope manufacturers differ on how best to design such probes. Although probing wideband circuits always affects the measured signals, well-designed probes minimize loading effects.

In response to users' demands to view waveforms in greater detail, some newer scopes sport screens as large as 12.1 in. diagonal.

For the widest-bandwidth measurements, the new NRO (near-real-time oscilloscope) minimizes the drawbacks of sequential-sampling instruments and provides rapid waveform acquisition and deep memory.

and using scope features should be intuitive—conforming, wherever possible, to the conventions with which the users are familiar.

Scope manufacturers point out that—at least with high-end instruments—your most valuable ally in selecting and effectively using the right instrument can be the field engineer who sold you the scope—or is trying to sell it to you. He can help you set up side-by-side compar-



Today's scopes can find anomalous waveforms for you if you tell them what to look for. LeCroy calls the feature exclusion triggering. Other manufacturers offer similar features but use different names.

isons with competitive units before your purchase and can supply advice and accessories to help you effectively use the scope. Representatives of distributors that sell scopes may offer similar services. Also, don't assume that factory support is unavailable to you because you purchased your instrument from a distributor. Depending on the manufacturer and the scope model you purchased, the factory may offer support. And remember that most scope vendors' Web sites offer a wealth of application notes containing information on effective use of the companies' products. **Tables 1 and 2**, at the Web version of this article at www.edn.com/060216cs, summarize key specifications of real-time-sampling scopes from four major manufacturers.


BEGINS WITH PROBE

An appropriate place to begin a discussion of modern scopes is with the probe. The probe tip is where the instrument meets the device under test. Time was that engineers considered only a few megahertz to be a high frequency. Now, probing gigahertz signals is commonplace, and familiar serial buses transmit signals at rates in excess of 3 Gbps. Scope manufacturers recommend that your scope and probe together have a -3 -dB bandwidth at least 1.8 times the bit rate. So, if you are working with a bus whose raw bit rate is 3.125 Gbps, your scope and probe should have a combined bandwidth

of at least 5.625 GHz. (A bus with a raw bit rate of 3.125 Gbps usually carries information at 2.5 Gbps; 8-bit/10-bit clocking embedded within the data stream limits the information rate to 80% of the raw bit rate.) The bandwidth closest to 5.625 GHz that scope manufacturers advertise is 6 GHz. The 6.67% margin above 5.625 GHz can help to compensate for bandwidth reduction attributable to the probe.

Several points are important. The first is that probing such high-speed serial buses is a job for differential active probes. At these speeds, nearly all buses are differential, and the signal swings are small for a variety of sound reasons: Unlike single-ended circuits, differential receivers tend to reject common-mode "noise," enabling the use of smaller signal swings; differential circuits also radiate less noise and subject power-supply rails to less transient loading than do single-ended circuits. But the smaller signal swings militate against passive probes, which, to reduce capacitive loading, generally attenuate their input signals. In addition, using two scope inputs to view one differential signal is out of the question. That approach effectively not only halves the number of channels on your scope, but also provides input-terminal pairs that are inadequately matched at the frequencies involved. The result can be the appearance on the screen of waveform artifacts that don't exist.

Multigigahertz-bandwidth differential active probes are amazingly clever, and



The three members of Tektronix's DPO7000 series sport 12.1-in.-diagonal, XGA-resolution, 1024×768-pixel screens. The top-of-the-line, 2.5-GHz-bandwidth unit accommodates 400M samples of acquisition memory, all of which you can assign to one channel.

their sophistication is likely to grow in the next few years. Manufacturers disagree about the best way to design and characterize these devices, but all manufacturers seem to agree on one fact: If you are trying to acquire multigigahertz signals, it is impossible to connect a probe to a unit under test without imposing *some* load on the signal you are trying to measure.

Manufacturers disagree, however, on whether that loading *always* has a *meaningful* effect on the waveforms you wish to view. Nevertheless, it is difficult to refute that, unless a probe is designed with the utmost care, the loading effects not only

can be meaningful, but also can make unacceptable waveforms appear perfectly fine or vice versa. For example, probe-induced errors can cause what is, in fact, a good waveform to appear to violate an eye-diagram mask or can make a waveform that violates the mask appear to comply.

That probes impose capacitive loads on units under test is well-known. However, a probe's series inductance is also important in determining the probe's response at several gigahertz. Moreover, the resonance between the probe's shunt capacitance and series inductance can have

even more dramatic effects both on the loading of the unit under test and on the probe's frequency and transient response.

PROBES GET SMARTER

Modern probing systems from all of the major scope manufacturers include facilities for bidirectional communication between the scope and the probe. Modern active probes do more than merely send the scope an amplified or buffered replica of the waveform at the probe tips, and the scope does more than just supply power to such probes. For example, LeCroy's newest probes store dynamic

CALIBRATING SCOPES' HIGH-FREQUENCY AMPLITUDE ACCURACY: MORE DIFFICULT THAN YOU MIGHT THINK

By Steve Sekel, LeCroy Corp

Customer questions and complaints about scope amplitude accuracy are fairly common. Customers try to measure the accuracy with a swept sine wave from a signal generator. Users shouldn't try this procedure themselves. Although the measurement sounds legitimate, the results are almost always wrong when the frequencies are higher than a couple of gigahertz.

The first problem is that you need to level the generator output at the output end of the cable. Even the best cables—those that cost more than \$1000—have some amplitude loss when you get to the several-gigahertz range. The *only* way to use a signal generator to measure amplitude accuracy is to use a high-quality, calibrated power divider at the end of the cable that connects to the oscilloscope.

One output of the power divider is connect-

ed directly to the power head of an RF-power meter that is calibrated for the frequency range and power levels you are testing. If you are testing all of the volts/division ranges, this measurement often requires using more than one power head. The power-meter readings normalize the output level at each frequency step. In an automated-calibration system, you perform this procedure under computer control. It is possible but tedious to manually perform the procedure.

REFLECTIONS

The second problem, which undoubtedly occurs in many cases, is dealing with the reflections from the scope input. In reality, the user is measuring the signal with the reflections superimposed. Scope inputs are not perfect 50 Ω terminations. Different attenuators switch using relays or electronic switching. Inevitably, the paths

are imperfect; they introduce some reflections at different frequencies.

Scope vendors work to minimize these reflections, but they all achieve about the same performance: a VSWR (voltage-standing-wave ratio) that, over the passband, can go from a perfect 1-to-1 to about 1.35-to-1. Whenever the termination reflects energy back into the line, the reflection creates standing waves at some frequency that relates to the length of the cable. Because they exhibit reflections at different frequencies, different models of oscilloscopes measure different amplitudes from the same generator-and-cable combination.

A user can reduce this effect by installing a high-quality, 6-dB attenuator at the scope's input and attaching the power-divider output to the attenuator. The attenuator improves the return loss by 6 dB, reducing the

effect of the reflection in the cable.

As you can see, the metrology required to accurately measure a scope's amplitude accuracy over frequency is complex. All scope manufacturers put considerable effort into designing and verifying the complex systems that designers use to calibrate the instruments. Attempting to manually replicate this measurement by using only a signal generator and cable can't produce results of the desired accuracy.


AUTHOR'S BIOGRAPHY

Steve Sekel is a product-marketing manager at LeCroy Corp. He has worked in the test-instrumentation industry for 28 years, serving in marketing, product-development management, and design engineering. He has a bachelor's degree in electrical engineering from Tri-State University (Angola, IN).

probe-calibration data. This data includes more than just the probe's offset voltage and dc gain; it includes high-frequency-gain- and phase (delay)-characterization data. According to Mike Lauterbach, PhD, LeCroy's director of product management, all ultrawideband scopes from all manufacturers use DSP-based techniques to correct the vertical-amplifier high-frequency-gain and -phase characteristics. The corrections improve the response so that it more closely resembles the desired response—often that of a fourth-order Bessel lowpass filter—than does the amplifier's uncorrected response.

As far as Lauterbach knows, however, only LeCroy's WaveLink probe family currently includes the probe response in the correction algorithm. Within seconds of your connecting a WaveLink probe to a compatible LeCroy scope, the correction routine uploads the calibration data from the probe and compensates the channel's vertical response for the probe's ac characteristics (as measured at the factory—or the last time you used a LeCroy-supplied fixture to characterize the probe). Including the probe in the calibration enables LeCroy, whose 11-GHz scopes offer narrower -3 -dB bandwidth than that of the nearest competitive models from Agilent or Tektronix, to nevertheless claim the most accurate high-frequency ac and transient response among real-time scopes in the more-than-10-GHz class. LeCroy also points out that, unlike at least one competitor, it does not currently use DSP to extend the bandwidth of its scopes.

In case you haven't noticed, modern wideband scopes do not have frequency response related to the 10 to 90% rise time by the time-honored formula $T_R = 0.35/BW$, where $T_R = 10$ to 90% rise time and $BW = -3$ -dB bandwidth. And you can't determine the combined rise time of the scope and probe from $\sqrt{(T_{R(SCOPE)}^2 + T_{R(PROBE)}^2)}$. For one thing, you must carefully check the data sheet's notes to determine whether each rise-time spec applies to the time the signal takes to traverse 10 to 90% or 20 to 80% of the input-step amplitude. Manufacturers sometimes specify both rise times.



LeCroy's WaveExpert, NRO (near-real-time oscilloscope), and SDA100G sampling scopes form a series that you can equip with sampling heads that provide 100-GHz bandwidth on four channels. Although they don't sample in real time, they sample 50 times as fast and store records many times as long as those of sequential-sampling scopes, the only scopes whose bandwidth is nearly as wide.

Some standards for bus physical layers use only the 20 to 80% values; using 10 to 90% values in such cases would only cause confusion. In addition to the "which-rise-time?" issue, however, the old formulas don't apply to new scopes and probes because the newer units' high-frequency-roll-off characteristics differ from those of the analog scopes whose behavior formed the basis for the old rules. To learn more about deep memory and finding ephemeral anomalies in long-waveform records, see **sidebar** "Acquisition memory: a deep subject" at the Web version of this article at www.edn.com/060216cs.

IT TAKES PERSISTENCE

Persistence mode doesn't work quite the way many people think it does (**Figure 1**, pg 52). To dispel the confusion, here is a brief explanation that generally applies to all scope brands. Note that persistence mode can often correctly acquire waveforms that—because of a limited real-time sampling rate—contain frequencies too high for the scope to capture in real time. Many scope users erroneously believe that capturing such waveforms requires using random equivalent-time sampling, a mode you must use with caution to avoid little-understood pitfalls (**Reference 1**).

To use persistence mode, the trigger must be stable in time with respect to the waveform that you want to capture. You can trigger on a waveform feature or use

another trigger source. Each time it triggers, the scope acquires waveform samples and places the corresponding dots on the screen with respect to the trigger time. It draws no line between the dots, though. By default, some scopes add sine x/x -interpolated dots, whereas others add none. The scope simply places the dots on the screen—or, to be more exact, it places the dots in an array in the display-processor IC, which draws the dots on the screen. The scope draws no line through the dots, however, and makes no attempt to re-create the shape of the incoming signal; such an attempt could violate the Nyquist criterion.

The scope then triggers repeatedly. Typically, it triggers several hundreds—or even thousands—of times. Each time, it acquires samples and places the dots on the screen, but it never attempts to "draw the trace." The scope simply displays the acquired samples with respect to the trigger time. If the trigger and the incoming waveform are stable, the set of dots is closely packed onto a line shaped like the signal and strongly resembling a waveform. If the trigger time or the waveform is unstable because of vertical noise or timing jitter, the persistence display shows a cloudier set of dots. If the signal shape exhibits occasional large, intermittent aberrations, you may see a large number of dots that follow the normal signal shape and a fainter number that show the abnormal shape.

SLOW REFRESH

Scope manufacturers make much of their instruments' fast screen-update rates and responsiveness to changes in control settings. Some companies refer to such attributes as "analog-scope feel." These claims are valid as well as important to the way in which you use a scope, but, if you think about the claims for a few moments, you can easily wonder how they can possibly not be exaggerations. Nearly all digital scopes refresh their screens only 30 or 60 times per second, yet many display many thousands of waveforms per second. They achieve this responsiveness by aggregating multiple changes to their screen bit maps between refreshes and displaying the

aggregated result at the next refresh.

This aspect of scope operation is philosophically similar to the way in which scopes whose displays have, say, 1024 pixels horizontally display million-point-deep records without forcing you to scroll endlessly through the long records. However, you can zoom to that mode, as well, if you choose. The simplest way to compress a million samples into 1000 pixel columns, each representing 1000 samples, is to find the minimum and maximum signal values in each 1000-sample group and illuminate all pixels in the column from the one that corresponds to the lowest value to the one that corresponds to the highest. This approach produces a “fat” trace, whose illumination is constant over its width. To show greater signal detail, a scope can determine how many times since the last screen update the signal level corresponded to each point in the screen’s pixel map and relate each pixel’s brightness or color to the number of “hits” at the associated point.

Scope manufacturers are also discovering the value of the big screen—not the living-room-dominating size of an HDTV and not even the wide aspect ratio of the screens on some laptop PCs but considerably larger in area than has been customary in scopes. Bigger screens on scopes make it easier to see waveform details. LeCroy started the trend a couple of years ago—at least in small-footprint scopes—with its WaveSurfer family, whose members sport 10.4-in.-diagonal, SVGA, 800×600-pixel screens in a 6-in.-deep package that occupies no more benchtop area than does a Tektronix TDS3000B, whose screen measures only 6.4-in. The WaveSurfer’s screen area is more than 2.5 times as great as that of the Tek unit. Now, LeCroy has added higher performance units to its stable of large-screen, small-footprint scopes. The three members of the WaveRunner Xi series, whose prices start at \$7500, are the same size as the WaveSurfers and also have 10.4-in. SVGA screens.

Not to be outdone, Tek, with its new DPO7000 series, has one-upped LeCroy on screen size and resolution. The DPO7000 screens measure 12.1 in. diagonal. Their area is approximately 3.6 times that of a 6.4-in screen, and they provide XGA, 1024×768-pixel resolution. The approx-

imately 12-in. package depth is roughly twice as great as that of LeCroy’s small-package units but is much shallower than most scopes. The DPO7000s, whose top-of-the-line unit can accommodate memory as deep as 400M samples—all of which is assignable to one channel—also attack LeCroy’s long-held dominance in memory depth.

Although welcoming the large screens and small benchtop footprints, engineers who incorporate scopes into larger systems—for example, for production test—may be less than thrilled with the new package geometries. For these engineers, selecting system components that occupy a minimum of rack space is of key importance. The new packages are taller than those of most traditional scopes. It seems likely that the solution to the

height problem will lie in LXI (LAN extensions for instrumentation), a new standard for system-component instruments. You can imagine low-profile LXI scopes whose screens lie flat atop them until an operator pulls them forward on their rack slides and hinges the screen into a vertical position.

BEYOND 20 GHz

A survey of the current state of digital-scope technology would be incomplete without some discussion of the widest bandwidth scopes—the class of instruments that engineers used to call sequential-sampling scopes. Until the advent, a year ago, of LeCroy’s WaveExpert and SDA100G series, the phrase “sequential sampling” was appropriate, and there were only two vendors, Agilent and Tektronix.





Figure 1 Analog-persistence mode maps frequency of occurrence into intensity or color variations on each pixel of the display simulating the phosphor response of an analog oscilloscope (courtesy LeCroy).

The LeCroy units essentially rewrote the book on how engineers design these ultrawide-bandwidth instruments (70 to 100 GHz, depending on the manufacturer). At the product introduction, LeCroy referred to its instruments simply as sampling scopes, because “sequential” didn’t really apply. But the problem with not qualifying “sampling” is that all digital scopes are sampling scopes. During the year, LeCroy solved its terminology problem by inventing a new term, “NRO” (near-real-time oscilloscope) and adding an NRO series to its line.

All scopes in this category—including the LeCroy units—depend on the signal’s occurring repetitively. It need not recur at a constant rate, but it must follow a trigger signal by an essentially fixed delay. Classic sequential-sampling scopes capture only one sample during each iteration of the input waveform, advancing the sampling point incrementally with each new trigger. Thus, despite their extremely wide bandwidth, these scopes acquire waveforms slowly. This low speed rules out instruments of this type in many common scope applications.

In some of these scopes, the analog sampler is separate from the scope mainframe. The sampler is a so-called zero-order hold circuit, which captures the input signal with femtosecond aperture uncertainty and maintains the captured voltage for tens of microseconds. The sampler output is thus a relatively low-frequency replica of a multigigahertz signal. From the sampler output onward, the analog signals that the scope deals with are relatively low in frequency. The ADCs in such scopes are usually high-resolution (14 bits or more) successive-approximation devices with conversion rates no higher than a few hundred kilo-



Acquiring multigigahertz signals without introducing loading effects that completely invalidate the measurement requires great care, an understanding of the physics of probing, and highly specialized probes. These 10-GHz Infiniimax units from Agilent work with one of the company’s DSO 80000-series scopes.

hertz. Memory depths in classic sequential-sampling scopes rarely exceed 100k samples.

BANDWIDTH TO 100 GHz

Advances in sampling technology enable the LeCroy units, with the appropriate sampling plug-ins, to achieve industry-leading 100-GHz bandwidth, whereas advances in ADC and memory technology make possible an architecture that differs considerably from that of sequential-sampling instruments. Instead of taking only one sample during each iteration of the input waveform, the LeCroy units take many. The company says that the sampling rate is 50 times that of the fastest competitive instrument. In addition, memory depths of hundreds of millions of samples are possible, and a built-in clock-recovery facility allows the scopes to operate, in many cases, without an external trigger. The scopes also accommodate built-in analysis features that you would probably expect to find only in real-time-sampling scopes. Thus, these scopes can handle many applications in which competitive instruments would acquire data too slowly, could not capture records of the necessary length, would require external equipment to trigger from the available signals, or would present a more complex interface to less-extensive analysis facilities.

As do Agilent and Tek, LeCroy offers optical-to-electrical converters to permit use of its ultrawideband scopes for fiber-

optic communication-system measurements. Unlike its competitors, though, LeCroy does not currently offer differential-input plug-ins for these scopes. As a result, you need two of the LeCroy mainframes to simultaneously view four more-than-20-GHz differential signals—a task the competitive units can perform with one mainframe. **EDN**

REFERENCE

■ Pupalaiakis, Peter J, *Random Interleaved Sampling*, November 2005, www.lecroy.com/tm/library/registerPDF.asp?wp=577.

AUTHOR’S BIOGRAPHY

Contributing Technical Editor Dan Strassberg has been covering test and measurement for EDN for more than 18 years. He holds a bachelor’s degree in electrical engineering from Rensselaer Polytechnic Institute (Troy, NY) and a master’s degree in electrical engineering from Massachusetts Institute of Technology (Cambridge). You can reach him at strassbergedn@att.net.

FOR MORE INFORMATION

MANUFACTURERS OF HIGH-PERFORMANCE BENCHTOP OSCILLOSCOPES

Agilent Technologies
www.agilent.com

LeCroy Corp
www.lecroy.com

Tektronix Inc
www.tektronix.com

Yokogawa Corp of America
www.yokogawa.com/us/

MANUFACTURERS OF VERY-HIGH-SPEED, HIGH-RESOLUTION MODULAR DIGITIZERS

Acqiris USA
www.acqiris.com

Gage Applied Technologies
www.gage-applied.com

National Instruments
www.ni.com

Ztec Inc
www.ztec-inc.com

MORE AT EDN.COM

✚ We encourage your comments! Go to www.edn.com/060216cs and click on Feedback Loop to post a comment on this article.

✚ This author writes about high-performance DSOs at www.edn.com/article/CA608891.

WHEN YOU NEED HIGH SAMPLING RATES BUT 8 BITS AREN'T ENOUGH

By Kaustubh Wagle, National Instruments

Today, many applications involving communications, medical instrumentation, ultrasonic NDT (nondestructive testing), audio, and video—to name a few—demand higher resolution digitization than is available in mainstream oscilloscopes but do not need modern digital scopes' versatile triggering capabilities. High-resolution-digitalizer modules in package formats such as PXI are ideal for such applications, especially when—as is often the case—applications require many channels. IC manufacturers have created ADCs with resolutions to 24 bits.

Manufacturers of instrument modules often face the trade-off of resolution for sampling speed and, hence, must produce creative designs and schemes to meet the increasing demand for high-resolution digitization at higher sampling rates. Two of these techniques are *time-interleaved sampling* and *linearization of multibit delta-sigma ADCs*.

Time-interleaved sampling, or *pingpong*, interleaves two or more

ADCs to achieve higher effective sampling rates. Time interleaving of n ADCs, each sampling at a rate of f_s , results in an effective sampling rate of $n \times f_s$. Many modern scopes use this technique to good advantage with minimal impact on accuracy. They can do so because of their modest ADC resolution—usually, 8 bits.

In higher resolution applications, there's a catch: Without great care, interleaving can result in unwanted frequency content. For example, ping-ponging two 12-bit, 100M-sample/sec ADCs achieves an effective sampling rate of 200M samples/sec with 12-bit resolution. However, the interleaved ADCs inevitably exhibit gain, offset, and phase mismatches (Figure A). In the frequency domain, gain and phase mismatches cause an *image spur* (a spectral line at $f_s/2$ minus the fundamental frequency) in the FFT, whereas offset mismatch causes an *offset spur* (a spectral line at $f_s/2$). Interleaving larger numbers of ADCs leads to more spurs. These spurious-frequency

components degrade such dynamic-performance specifications as dynamic range, SINAD (signal, noise, and distortion), and ENOB (effective number of bits), defeating the original purpose of achieving high resolution at higher sampling rates.

REDUCING SPURS

Designers of 8-bit systems, such as scopes, have learned how to control such problems, but, in higher resolution systems, reducing pingpong spurs requires more elaborate measures. For example, you can implement classic analog-matching techniques involving the use of common reference voltages, matched physical layouts, and equal-length traces. At resolutions of 12 bits and more, however, many of these techniques require additional circuits that add error sources of their own (Reference A).

The other technique for reducing pingponging effects is digital post-processing, either in software on the host PC or in a powerful FPGA on the digitizer itself. In any case, it is critical that you reduce, if not eliminate, the image and offset spurs that pingponging causes. Otherwise, you defeat the purpose of pingponging, and you achieve only the higher sampling speed—not the higher resolution. One caveat for scope and digitizer users is to carefully note the sampling speed for the dynamic specifications, which they can usually find in fine print in the footnotes of the data sheet.

Another creative technique for achieving high resolution at higher sampling speeds is the *linearization of multibit delta-sigma ADCs*. Single-bit delta-sigma ADCs provide high resolution and high dynamic range for low-frequency applications. However, because of

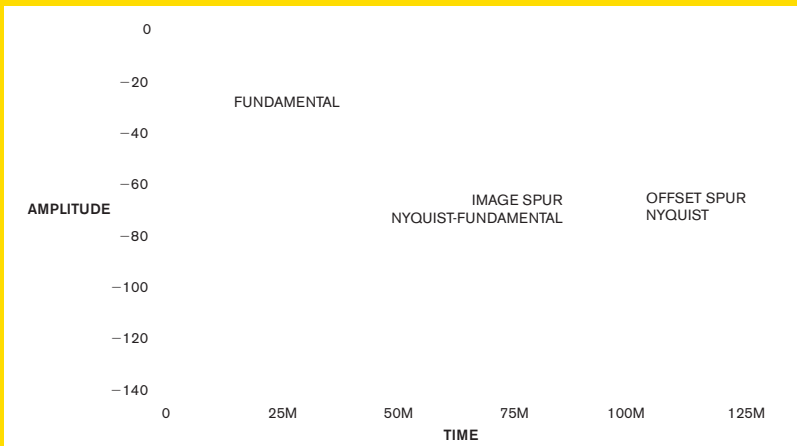


Figure A Time-interleaved sampling of two 12-bit, 100M-sample/sec ADCs results in an effective sampling rate of 200M samples/sec. Gain and phase mismatch between the ADCs results in the image spur, and the offset mismatch between the ADCs results in the offset spur.

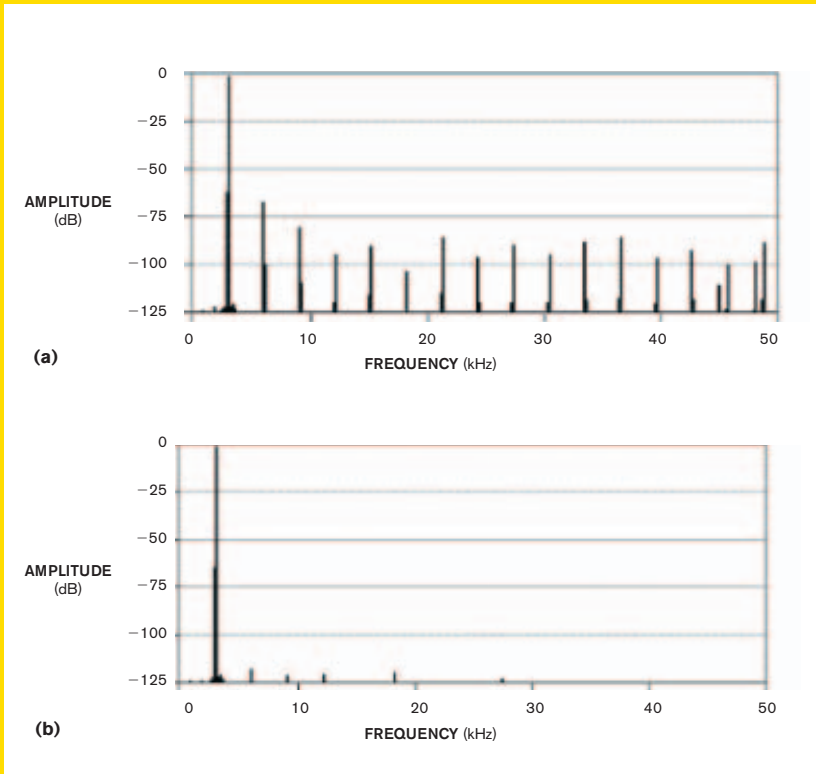


Figure B These FFT plots show what happens when you apply a pure sine wave to a 6-bit delta-sigma ADC before linearization (a) and afterward (b).




Figure C At less than 500k samples/sec, the PXI-5922 flexible-resolution digitizer can resolve 24 bits (0.06 ppm of full-scale). At 15M samples/sec, the resolution is still 16 bits (15 ppm).

the limited sampling speeds, single-bit delta-sigma ADCs are not available for applications involving dynamic signals at frequencies greater than a few hundred kilohertz. Multibit delta-sigma ADCs

can provide high dynamic range at high frequencies if you can linearize the ADC to remove the inherent nonlinearities.

Figure Ba illustrates how nonlinearities in the ADC show up as har-

monics in the frequency domain. The Flex II ADC (Reference B) from National Instruments uses a powerful FPGA and patented linearization techniques to digitally remove these nonlinearities and to provide unprecedented dynamic range at higher sampling rates (Figure Bb).

The increased dynamic range enables engineers to analyze a signal that the noise floor of traditional instrumentation would lose.

The PXI-5922 flexible-resolution digitizer incorporates the Flex II ADC to provide high resolution to 15M samples/sec. You can use this single digitizer to sample 24 bits at 500k samples/sec to 16 bits at 15M samples/sec (Figure C).

REFERENCES

- A** Looney, Mark, *Advanced Digital Post-Processing Techniques Enhance Performance in Time-Interleaved ADC Systems*, *Analog Dialogue*, August 2003.
- B** Wagle, Kaustubh and N Knudsen, *Flex II ADC Technology White Paper*, National Instruments 2005.

AUTHOR'S BIOGRAPHY

Kaustubh Wagle is product manager of the modular-digitizer line at National Instruments, which he joined in 2003. He holds bachelor's degrees in computer science and in economics from The University of Texas-Austin.

ACQUISITION MEMORY: A DEEP SUBJECT

One of the big deals in the scope world is the ever-increasing need for more acquisition memory. Among major manufacturers, only Tektronix continues to do a brisk business in benchtop scopes with relatively shallow memories—2500 to 10,000 samples/channel. These are lightweight, lunch-box-sized units in the attractively priced TDS3000B series and below. These scopes' FISO (fast-in/slow-out) architecture rapidly stores real-time samples in analog form in CCD (charge-coupled-device)-like ICs, sometimes referred to as analog shift registers. The samples shift from these ICs into moderate-speed ADCs. Tek may well have built more FISO scopes than it and all of its competitors have built scopes that use other architectures. FISO is cost-effective—as long as you can live with its limitations. Tek will likely continue to sell these or similar units in good volume for years to come, but, relative to all scopes sold, fewer and fewer FISO units are likely to appear atop design engineers' lab benches. Instead, such scopes will become even more pervasive in maintenance and repair work, where they are already popular.

Aside from the FISO corner of Tek's product line, deep memory is hot and is not destined to cool off any time soon. The reason

is that, as a scope's bandwidth increases, the instrument's acquisition rate must increase proportionally, and, as the sampling rate increases, a record of a given duration comprises proportionally more samples. To capture a 15-GHz signal, you must, as a practical matter, acquire at least 40G samples/sec. A 20-msec record (whose duration equals that of one cycle at 50 Hz) thus comprises at least 800M samples ($20 \times 10^{-3} \text{ sec} - 40 \times 10^9 \text{ samples/sec} = 800 \times 10^6 \text{ samples}$). To achieve this record length on four channels, you need a scope with a total acquisition memory of 3.2G samples. Bear in mind that this memory must be much faster and, hence, much more costly than that in a PC. At 40G samples/sec, the scope must store a new sample on each channel every 25 psec! No company currently sells such a scope. If such an instrument existed today, its selling price could easily approach \$250,000. Nevertheless, it would be wrong to bet against the introduction of such a product within the next couple of years. Moreover, when it does appear, the instrument is likely to cost much less than \$250,000 in 2006 dollars.

On Jan 4, 2006, Tektronix introduced its DPO7000 series, which includes four-channel, 500-MHz-, 1-GHz-, and 2.5-GHz-bandwidth units having base

US list prices of \$14,000 to \$24,900. This choice of bandwidth indicates that Tek focused its initial efforts in this round of scope introductions on the market area in which the money is. Lower performance scopes sell in larger unit volumes but at lower prices. Higher performance scopes can cost more than \$100,000 apiece, but the unit volumes are modest. When all channels are active, the series' highest performance unit, the 2.5-GHz-bandwidth DPO7254, captures a maximum of 10G samples/sec/channel. For a price premium of \$15,000, you can buy the scope with memory of 100M samples/channel. Moreover, if you use only one channel, you can assign all 400 Mbytes to the active channel. However, contrary to what has for the past few years been a de facto industry standard, assigning the unused channels' memory to the active channel doesn't force you to interleave those channels' ADCs with that of the active channel. Hence, the DPO7254 can capture a 40-msec record of a 2.5-GHz-bandwidth signal at 10G samples/sec. (See sidebar "When you need high sampling rates, but 8 bits aren't enough.")

EPIHEMERAL ANOMALIES

Being able to capture huge records can be a mixed blessing. Often, long records contain just a

few precious nuggets—instances when the unit under test misbehaved in important ways. Those waveforms are the ones you are almost surely looking for. But if the records are tens or hundreds of millions of points deep and contain just a few—or even a few dozen—anomalous waveforms, the probability of your missing something important in a visual scan is nearly 100%. One of these days—soon, with any luck—your scope may be able to take at least a first cut at what you ought to look at and what you can probably safely ignore. Today, though, you are likely to have to tell the scope what to find as opposed to what to ignore (Picture). If you think about the problem, however, telling the scope what to ignore—or even letting the instrument decide what to ignore—sounds like the better deal for the user.

To some extent, the mask-test feature of modern scopes can help you in your search. Scopes with this feature allow you to define a pass/fail mask. You capture a waveform that you consider normal for the measurement of interest. You then display the waveform and set limits—bands of time variations for the portions of the waveform that rise and fall rapidly and bands of voltage variations for the horizontal (or nearly horizontal) portions—in effect

creating a thickened version of the normal waveform. The scope may even further simplify this job by specifying default values for the width of the mask bands. When you run the scope in the mask-test mode, you can choose among several alternative actions the instrument can take if it detects a waveform that falls outside the mask. For example, the scope can stop capturing additional waveforms, it can sound an alarm, or it can e-mail you a message.

Performing mask testing while running the scope in what Agilent and LeCroy call the segmented-memory mode and what Tek calls the FastFrame mode might be more beneficial. By allowing you to divide a long acquisition memory into smaller sections, these modes enable the scope to rapidly capture multiple instances of waveforms that satisfy the trigger conditions. If, for example, your scope's acquisition memory held 2^{20} (1M) samples, you might divide it into 64 segments of 2^{14} (16k) samples. Scopes of some makes capture their greatest possible number of waveforms per second in this mode. It might be useful if you could set up the scope to capture a waveform, test it against the mask, and, if it fell within the mask—that is, if

it met the criteria you established for a normal waveform—write the next-acquired waveform over the just-acquired normal waveform. That way, when the scope's memory filled or when you manually stopped the test, you would have a library of aberrant waveforms, each tagged with its position in the acquisition sequence. When you had determined how long the scope took to test one waveform against the mask and how long it took to rearm for the next trigger, you would know not only what the various failing waveforms looked like, but also approximately what fraction of waveforms failed; if there were multiple failure modes, you would know the relative likelihood of each occurrence. Such a "capture-failing-waveforms-only" mode could greatly facilitate searches for failures that occur infrequently.

Nearly a decade ago, Tektronix developed a different way of performing a similar task. Tek now calls this mode FastAcq and the underlying technology DPX (digital phosphor) because it enables digital scopes to emulate the behavior of analog-storage scopes that are based on phosphors with long decay times. Other manufacturers have developed digital-scope

modes that also emulate analog-storage scopes. Although these persistence modes present displays that can resemble FastAcq's, the way in which these scopes create the displays differs significantly.

FastAcq is limited to relatively short records (1000 samples in the latest implementation). It bypasses the normal capture of waveforms as records of time-ordered samples and instead immediately builds a 3-D pixel map in which color grading normally represents the dimension perpendicular to the screen. By so doing, FastAcq, in its latest implementation, captures as many as 250,000 waveforms/sec, more than other manufacturers' persistence modes. However, persistence modes can handle much longer waveform records. Moreover, some scopes retain the original time-ordered records, enabling subsequent analysis of the captured data. If FastAcq displays an anomaly that you want to study, you must use one of the scope's standard modes to capture another instance of the anomaly. If the anomaly occurs infrequently, that approach can take a while.