Eight years later, MAKING

UNDERSTANDING WAVEFORM-GENERATOR OPERATION AND SPECIFICATIONS BEFORE YOU BUY IS AS IMPORTANT TODAY AS IT WAS EIGHT YEARS AGO.

BY DAN STRASSBERG • CONTRIBUTING TECHNICAL EDITOR

It would be untrue and grossly unfair to the test-and-measurement industry to say that, in the eight years since *EDN* published an article on waveform generators (**Reference 1**), the technology and the products have remained unchanged. The most obvious change is the disappearance of some of the manufacturers mentioned in '98—although those companies' product lines are mostly still available from the companies that acquired the businesses. Strikingly, though, despite the appearance of many improvements and higher performance products and the fact that a 2006 dollar buys more waveform-generation capability than did a 1998 dollar, the underlying technology has remained surprisingly stable. Except for some new features, the block diagrams of most 2006 waveform generators are remarkably similar to those of 1998 instruments. Moreover, the need to understand the specifications and operation of the waveform-generation products that you buy today is just as important as was that need eight years ago.

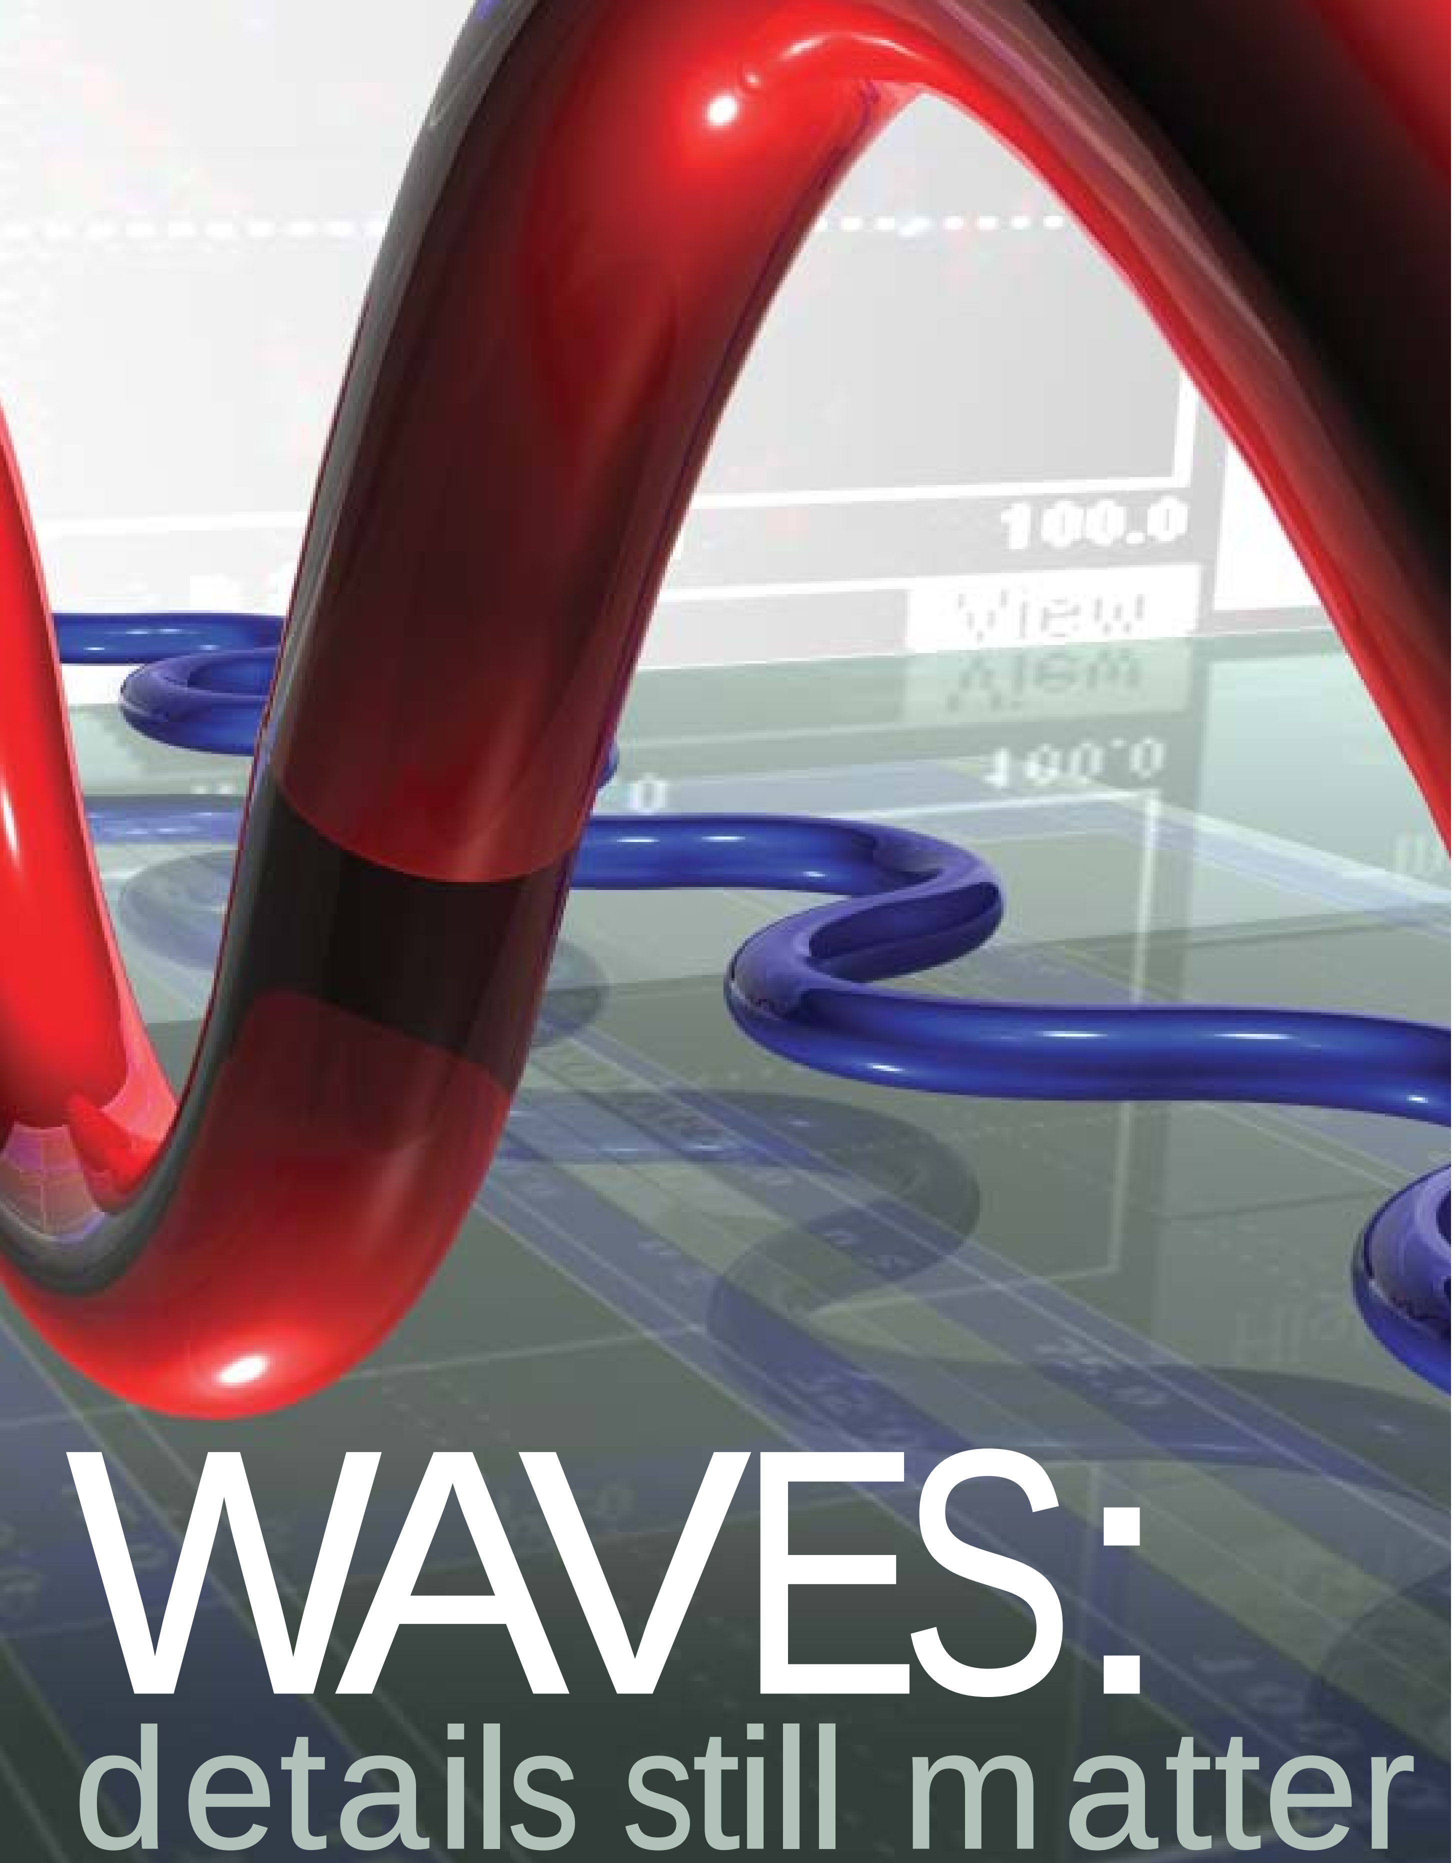

WAVES:

details still matter

AT A GLANCE

- Selecting a waveform generator can be confusing. Vendors use different terminology to describe their instruments' architecture. You need to carefully study the data sheet to determine an instrument's important characteristics. Failure to do so is likely to result in unpleasant surprises.

- Generators that use DDS (direct-digital synthesis) provide impressive features at attractive prices, but there are two basic ways to build DDS-based generators of user-defined arbitrary waveforms. For such waveforms, DDS-based true AWGs (arbitrary-waveform generators) are more flexible and generally more expensive than SFGs (synthesized-function generators).

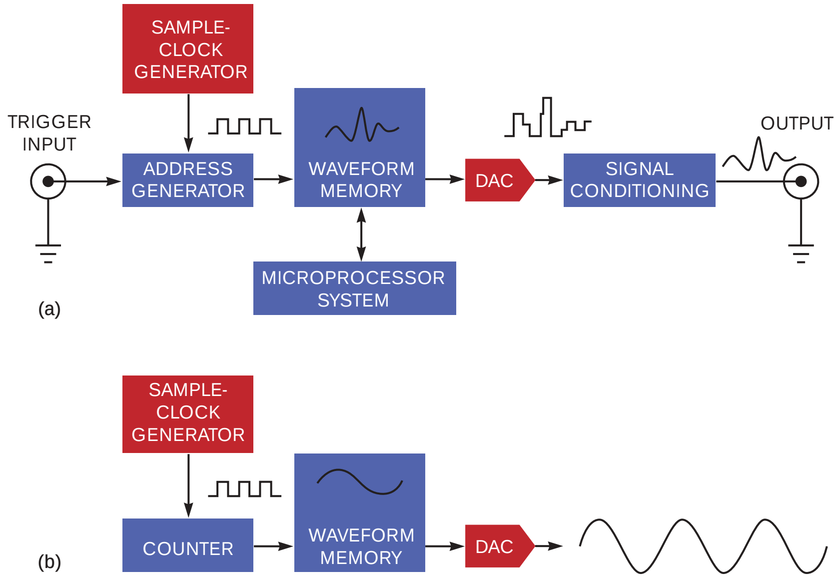

Figure 1 Whereas the block diagram of a true arbitrary-waveform generator (a) and that of an SFG (synthesized-function generator) are similar, the SFG usually lacks the microprocessor system and uses relatively simple filters in the signal-conditioning block. In a generator of fixed functions (b), the sample-clock generator may be DDS-based (courtesy Wavetek).

Even though their underlying technology is digital, the waveform generators that this article covers all produce analog waveforms. Generators that are analog throughout are still being manufactured—just as are analog oscilloscopes, but, like analog scopes, analog waveform generators are now restricted almost entirely to low-cost units whose primary market is in education. Internally digital generators range from units that produce sinusoidal signals whose output frequencies max out at a few megahertz to units that can deliver sine waves at frequencies as high as 500 MHz. Upper frequency limits for sinusoidal outputs are more commonly only 50 or 80 MHz, however. When producing nonsinusoidal waveforms other than square waves, these generators' maximum output frequencies are usually lower, although the maximum square-wave frequency is often equal to the maximum sine-wave frequency.

Prices for high-quality waveform generators begin at a little more than $1000 and range upward to more than $50,000

TABLE 1 WAVEFORM-GENERATOR ARCHITECTURES

Parameter	DDS-based synthesized-function generator	Arbitrary-waveform generator (ARB, AWG)	Tabor Wonder Wave series*
Sampling rate	Fixed	Variable	Variable to 1.2G samples/sec
Waveform type	Standard plus limited arbitrary	Arbitrary	All standard, arbitrary, and captured waveforms
Vertical resolution	Typically 14 bits	To 16 bits	To 16 bits
Memory depth	To 14k points	To 4M points	To 16M points
Memory management (sequencer)	No	Yes	Yes, as many as 16,000 segments, four advance modes
Signal integrity	Can skip or add points	Precise	Precise
Jitter and phase noise	High	Low	Low
Modulation	All frequency modulation	Amplitude modulation or none	All frequency, amplitude, and phase modulation, including digital
Frequency sweep	Yes	Limited	Yes
Trigger modes	Some	Yes	Continuous, trigger, burst, gated, retrigger
Amplitude range	10V p-p into 50Ω maximum	10V p-p into 50Ω maximum	To 32V p-p into open circuit
Digital outputs	No	Some	Yes
Ability to synchronize multiple instruments	No	Yes	Yes with extended facilities
Cost	Affordable	Can be expensive	Affordable

*True AWG with DDS clock and additional proprietary features (courtesy Tabor)

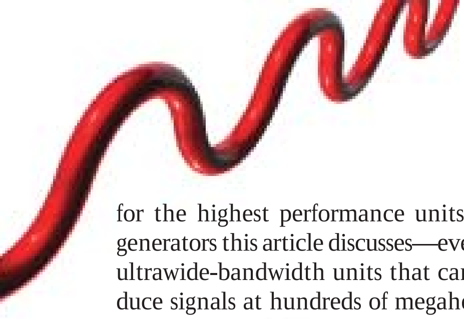

for the highest performance units. The generators this article discusses—even the ultrawide-bandwidth units that can produce signals at hundreds of megahertz—are usually thought of as baseband devices. Nevertheless, most generators that produce signals at frequencies beyond 80 MHz are categorized as RF-signal sources. Unlike baseband generators, however, RF generators can usually produce both modulated and unmodulated carriers. In most cases, RF generators include modulation sources. But these instruments usually also accept modulation signals from external sources, such as the baseband generators discussed here. Often—especially when the modulation is digital—the baseband generator must produce two independent outputs, I (in-phase) and Q (quadrature), because many forms of digital modulation depend on separately controlling the RF signal's I and Q components (see sidebar "Using waveform generators in IQ-modulator characterization" at the Web version of this article at www.edn.com/060914cs).

Several architectures dominate waveform generation (Figure 1). Probably the most common, albeit not the most familiar, is the SFG (synthesized-function generator), or function synthesizer. The SFG architecture is a clever elaboration on DDS (direct-digital (frequency) synthesis). The more familiar architecture is that of the ARB or AWG (arbitrary-waveform generator). In all likelihood, the reason that the ARB architecture is more familiar is that designers more easily understand it at the block-diagram level—though not necessarily at the detailed implementation level. Some generators—even relatively inexpensive ones—employ both architectures, using each architecture to perform those functions the instrument designers believe it performs best.

According to the company, Tabor's Wonder Wave series (Table 1) delivers, without compromise, the best features of AWGs and SFGs at prices approaching those of SFGs, which are generally less expensive than AWGs. Although this impressive claim suggests an architectural breakthrough, the generators are true AWGs that use DDS-based clocks.

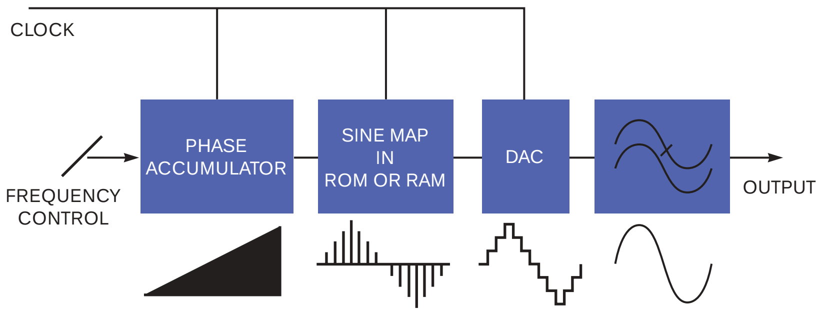

Figure 2 SFGs and most AWGs that have swept-frequency capabilities use clock generators based on the DDS technique known as phase accumulation.

STRENGTHS, WEAKNESSES

It is important to understand the strengths and weaknesses of the architectures and how these characteristics can either suit a generator for a certain job or provide a reason for you to select a different model for your application (Table 2). A lack of understanding can lead you to discover too late unexpected performance quirks that make using the generator considerably more difficult than you expected.

One complication in evaluating waveform generators is that manufacturers are

TABLE 2 GUIDE TO SELECTING WAVEFORM GENERATORS BY APPLICATION

	Generator category			
Criteria	Function/ sweep	Arbitrary/ function (SFG)	Arbitrary waveform	Vector signal
Frequency				
To 20 MHz	X	X		
20 to 240 MHz		X	X	
Greater than 240 MHz			X	X
Waveform characteristics				
Standard (sine, square, triange, sawtooth, ramp, noise, etc)	X	X		
Long, complex, fast transitions			X	X
IQ-modulation signals		X	X	X
Application				
High-speed serial data (PCI Express, SATA, HDMI)			X	
Low-speed serial data (I^2C, CAN, and others)		X		
RF test (wireless communication, RF, baseband, UWB, defense)		X	X	X
Optical- and magnetic-storage testing			X	
Digital test (mixed-device components, DAC, ADC, memory)			X	
Image-device test (capture and display)		X	X	
Other (automotive, education, medical)	X	X	X	
Available budget				
To $1500	X			
$1500 to $10,000		X		
More than $10,000			X	X

(Table courtesy Tektronix)

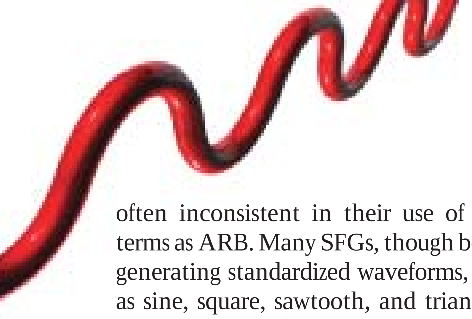

Tabor's WW1281 boasts 400-MHz bandwidth and true-AWG capabilities, including deep memory. The manufacturer claims that the generator family offers all of the advantages and none of the drawbacks of SFGs and true AWGs at prices close to those of SFGs.

often inconsistent in their use of such terms as ARB. Many SFGs, though best at generating standardized waveforms, such as sine, square, sawtooth, and triangular waves, can, within limits, produce arbitrary waveforms. Adding such capabilities to an SFG does not add tremendous cost, and users who believe that they may someday need to generate arbitrary waveforms find the feature an attractive security blanket. Unfortunately, the manufacturers sometimes describe these SFGs as ARBs or AWGs, even though the instruments can't match the arbitrary-waveform-generation capabilities of generally more expensive, true ARBs. Some manufacturers use the term "AFG" (arbitrary-function generator) to denote instruments that are basically SFGs but also offer limited arbitrary-waveform-generation capabilities.

Further complicating product selection is the fact that many true ARBs use DDS to generate their swept-frequency clocks (Figure 2). Whereas DDS is well-suited to generating frequency sweeps and a sweep capability is an important attribute of most SFGs, not all AWGs that use DDS to produce swept-frequency clocks deliver the best combination of AWG and SFG attributes.

THE PHASE ACCUMULATOR

The heart of an SFG is the phase accumulator. This elegant functional block converts a reference clock at a fixed-frequency, f_C, into a stable, lower frequency output clock whose frequency, f_O, you can set with great precision over a wide range and change essentially instantaneously in tiny, huge, or in-between steps. In addition, the output signal's phase is continuous when the frequency changes. That is, if you command the frequency to change from, say, 10 MHz to 100 kHz at the 227° point of an output cycle, the 100-kHz waveform begins at its 227° point. Moreover, if the phase accumulator drives a DAC that produces an analog output, no discontinuity occurs in the output amplitude—only a change in slope if the output is a sine wave or another continuous waveform.

DDS is well-suited to producing swept frequencies because you command the phase accumulator to produce a particular value of f_O by supplying it with a binary number, M. (Implementations also exist in which M is in binary-coded-decimal format.) The greater the value of M, the higher the output frequency. If you continuously vary M, f_O continuously varies. By supplying values of M that change over time in appropriate ways, you can create a limitless variety of sweep types, of which linear and logarithmic ramps, sawteeth, and sinusoidally varying frequencies are probably the most common.

You can think of this approach as dividing an output cycle into 2^N intervals or steps. Suppose $N=14$; then, $2^N=16{,}384$ and each step represents 360°/16,384 or 0.02197°. If $N=14$ and $M=2048$, the sample points occur at 45° intervals, and $f_O=f_C/8$. If $M=1$, the sample points occur at 0.02197° intervals and $f_O=2^{-14}\times f_C$. In this example, the values of M are powers of 2, but in a real system in which $N=14$ (actually a somewhat simplified version of a real system), M would be a 14-bit binary number and could therefore assume any integer value from 1 to 16,384. An output cycle need not contain an exact integer number of sample periods, however. Indeed, the lack of such a requirement makes possible the synthesizer's superb frequency resolution—that is, its ability to let you adjust f_O in such tiny increments.

BK Precision's DDS-based, 21-MHz-bandwidth 4070A offers many attractive features at a moderate price.

This digital-frequency synthesizer is a sampled-data system and is subject to the limitations of the sampling theorem. In the preceding example, the theoretical maximum output frequency is incrementally less than $0.5\times f_C$. The practical maximum is somewhat lower—approximately $0.4\times f_C$.

OTHER FUNCTIONAL BLOCKS

Although the phase accumulator is a key element in an SFG, it isn't the whole instrument. To produce analog waveforms, the instrument must convert the frequency synthesizer's swept- or variable-frequency output into an analog waveform. This conversion requires a waveform memory, a DAC, and output-signal conditioning—which often includes deglitching of the DAC output, always includes filtering, and usually includes amplification. With the frequent exception of square waves, which designers usually think of as trivial, the data sets that define the standard waveforms reside in

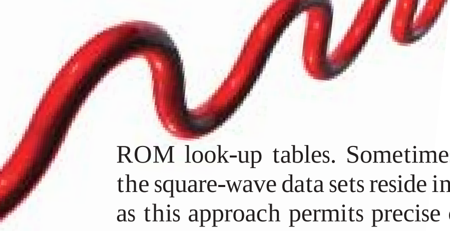

ROM look-up tables. Sometimes, even the square-wave data sets reside in ROM, as this approach permits precise control of rising- and falling-edge slew rates. For user-defined—that is, arbitrary—waveforms, the look-up tables reside in RAM instead of ROM. As f_O increases, the number of points in the waveform definition decreases. In other words, as f_O increases, the generator skips sending more and more of the values from the look-up table to the DAC. Conversely, at low values of f_O, the generator may repeatedly send some or all of the values from the look-up table to the DAC before moving to the next value in the data set.

Although it is not intuitively obvious, an SFG that uses a single-frequency clock to provide the input to its frequency synthesizer, as do most SFGs, does not need output lowpass filters whose corner frequency depends on the output signal's repetition rate. The ability to use relatively simple fixed-frequency output filters is a key reason that SFGs can be highly cost-effective.

For builders of test systems, several companies offer waveform generators as modular instruments. LeCroy's PXI-based PXA125 re-creates data sets as deep as 2M points at 125M samples/sec.

Though they offer a large number of highly desirable characteristics, SFGs are, nevertheless, not perfect. For example, the fact that a generator may skip points in the waveform-definition look-up table at high values of f_O and may repeat points at low values of f_O suggests that the output waveforms are not completely predictable. Moreover, because M, the number that determines f_O, need not be an integer submultiple of the look-up-table depth, cycle-to-cycle variations can easily occur in the waveforms. That is, on successive iterations of the output waveform, different look-up-table points will repeat or be absent. Thus, SFG outputs, though extremely stable in average frequency over the long term, can exhibit troublesome cycle-to-cycle variations in the short term. You can correctly characterize these variations as jitter.

In addition, whereas waveform-memory depths of 2^{14} (16,384) points—the value in several popular SFGs—may sound generous, user-defined waveforms are often as much as 16M points deep. Moreover, users may want to effectively modify the waveform depth on the fly by repeating waveform segments, often through the use of conditional branching and looping. In the waveform-generator context, conditional branching is the ability to jump to a specified point in the waveform-definition table after satisfying a test condition. Looping is the ability to repeat a range of waveform-definition points either a specified

number of times or until some test condition is satisfied. Although such capabilities are common in ARBs, no SFG currently on the market offers these features, and it is unclear that such capabilities would be consistent with the SFG architecture.

AWGs: SIMPLE IN CONCEPT

It would be convenient to say that AWGs developed as a result of SFGs' shortcomings, but AWGs predate SFGs and have evolved into the dominant high-end waveform-generation products as SFGs captured an increasing share of the high-volume applications. Conceptually, AWGs are simple: A clock, which can use DDS, drives a counter, which supplies addresses to the waveform memory. For arbitrary waveforms, the memory is always RAM. The sequencing logic, which makes possible the looping and branching capabilities, stands between the counter output and the memory-address lines. The data-set values from the memory become the input to a DAC, which drives analog-signal-conditioning circuits that supply the output signal. From a block-diagram perspective, the AWG's signal conditioning is much the same as that of the SFG, although, in practice, the filters can differ significantly.

Strictly speaking, although you can use a DAC on a PC-plug-in data-acquisition board to generate arbitrary waveforms, such a board does not fit the definition of an AWG. For example, data-acquisition boards incorporate no looping-and-branching memory-addressing logic. And, whereas a data-acquisition board may provide a means of adjusting the DAC's offset and full-scale-output voltages over narrow ranges, if you want to reduce the full-scale output from, say, 10.24 to 5.12V, either you must furnish your own output attenuator, thereby increasing the output-source resistance, or you must divide the DAC's digital inputs by two, thereby reducing the unit's resolution. On the other hand, many instrument-level AWGs have wide-range gain controls that require neither of these compromises.

Although data-acquisition boards aren't AWGs or SFGs, true AWGs and SFGs are available not only in familiar benchtop-instrument formats, but also as modules in such formats as CompactPCI and PXI, as well as in standard PC-plug-in—that is, PCI—formats. Among the suppliers are LeCroy, National Instruments, and VXI Technology.

MORE AT EDN.COM

+ We encourage your comments! Go to www.edn.com/060914cs and click on Feedback Loop to post a comment on this article.

Fluke's 291 is part of a series of one-, two-, and four-channel DDS-based generators that reproduce data at speeds to 100M samples/sec. Despite the high speed, the unit's price is close to those of some generators with much lower performance.

Output filtering is an area in which AWGs and SFGs can differ significantly. As stated, SFGs work well with relatively simple fixed-frequency lowpass filters. In contrast, AWGs can sometimes benefit from output lowpass filters whose −3-dB frequency, $f_{-3\text{ dB}}$, is a constant fraction of the clock frequency. With a DDS-based clock, $f_{-3\text{ dB}}$ would thus be a constant fraction of the DDS-output frequency. The problems with this approach, however, are its complexity and cost. As a result, practical generators sometimes settle for a few switchable, fixed-frequency filters.

In an AWG, there is merit to making $f_{-3\text{ dB}}$ and perhaps the shape of the filter's stopband response user-selectable and not linking $f_{-3\text{ dB}}$ to the clock frequency. The generator's designer has, at best, only a vague idea of the complexity of the waveforms the instrument users will want to synthesize. At a given repetition rate, more complex waveforms have greater high-frequency content than do simpler ones and may require a more complex filter topology, a higher $f_{-3\text{ dB}}$, or both. Using DSP to shape the values of the data-set points that define the waveform before applying the values to the DAC might be a way to optimize the generator's frequency response and minimize the filter requirements. The multimedia world uses upsampling, in which DSP techniques fill in intermediate values in sparsely sampled, bandwidth-limited data sets, thus simplifying the reconstruction of analog signals from sampled data. Such techniques might also work in waveform generators. Nevertheless, no commercial baseband waveform generator currently uses either approach.

DON'T GET CARRIED AWAY

Companies that market waveform generators emphasize that the biggest problems their customers experience with the products involve learning to use the more arcane features, especially defining arbitrary waveforms. Although many generators that produce arbitrary waveforms require a separate computer running appropriate software for creation of waveform-definition files, some generators include built-in waveform-definition applications. Users who lack experience with these applications and who believe that they can learn by creating a complex waveform that they need *now* are likely to suffer disappointment and frustration. Even for experienced users of waveform-definition applications, the creation of complex waveforms takes time. You should gain facility with such applications by first learning how to define relatively simple waveforms, and, to do so, you should set aside adequate learning time when you are not under pressure. You should try defining complex signals only after you have mastered the basic operations, and, again, you should avoid working on complex definitions when you are under time pressure.

Generally speaking, waveform-generator user interfaces are only moderately complex. Many newer instruments have front-panel waveform displays that give the generators the look of oscilloscopes. To the uninitiated, this appearance can be deceiving. Few of these instruments provide displays of actual output waveforms. Rather, the displays give you an idea of such waveform characteristics as rise time when you are specifying them.

Although the number of bits of vertical resolution in waveform-definition data sets is important, don't assume that a generator's vertical accuracy necessarily equals the full-scale output divided by 2^N, where N is the number of bits. Look instead for a specification of vertical resolution as a percentage of full scale; it is likely to correspond to a slightly smaller effective number of bits because of static and dynamic errors in the DAC and signal-conditioning circuits. Remember, too, that both the actual and the effective word length tend to be lower in generators with higher bandwidth. If you can find an instrument that simultaneously delivers 16 effective bits of resolution and 500-MHz bandwidth, it is likely to cost you dearly. EDN

REFERENCES

1 Strassberg, Dan, "Choosing a waveform generator: The devil is in the details," *EDN*, Sept 1, 1998, pg 75, www.edn.com/archives/1998/090198/pdfs/18df2.pdf.

2 "Direct-Digital Frequency Synthesis: A Basic Tutorial," Osicom Technologies Inc, www.ehb.itu.edu.tr/~eepazarc/ddstutor.html.

3 "The ABCs of Arbitrary Waveform Generation," Agilent Technologies, Publication No. 5989-4138, http://cp.literature.agilent.com/litweb/pdf/5989-4138EN.pdf, November 2005.

4 "XYZs of signal sources," Tektronix Inc, Publication No. 76W_16672_3, www2.tek.com/cmswpt/tidetails.lotr?ct=TI&cs=Primer&ci=2323&lc=EN, June 2005.

FOR MORE INFORMATION

Aeroflex
www.aeroflex.com

Agilent
www.agilent.com

Berkeley Nucleonics
www.berkeleynucleonics.com

BK Precision
www.bkprecision.com

Elan
www.elandigitalsystems.com

Fluke
www.fluke.com

Krohn-Hite Corp
www.krohn-hite.com

LeCroy
www.lecroy.com

National Instruments
www.ni.com

Rohde & Schwarz Company
www.rohde-schwarz.com

Stanford Research Systems
www.thinksrs.com

Tabor
www.taborelec.com

Tegam
www.tegam.com

Tektronix
www.tektronix.com

VXI Technology
www.vxitech.com

Yokogawa
www.yokogawa.com

AUTHOR'S BIOGRAPHY

Contributing Technical Editor Dan Strassberg has covered test and measurement for EDN *for nearly 19 years. He holds two degrees in electrical engineering—a bachelor's from Rensselaer Polytechnic Institute (Troy, NY) and a master's from the Massachusetts Institute of Technology (Cambridge).*

USING WAVEFORM GENERATORS IN IQ-MODULATOR CHARACTERIZATION

By Mark Albert, Tektronix

Modern mobile-telecommunications standards rely on IQ (in-phase/quadrature) modulation to impress digital data onto RF carriers. IQ modulation encodes the digital signal into amplitude variations of two analog sine waves that are phase-shifted by 90° relative to each other—the I and Q signals. Designers of IQ modulators must ensure that the devices meet certain performance criteria, such as IQ gain balance, phase skew, local-oscillator feedthrough, intermodulation distortion, and common-mode-rejection ratio.

To characterize IQ modulators with respect to these parameters, you need a dual-channel signal generator to simulate the I- and the Q-input signals. Whereas a number of dedicated vector-signal generators are available, most offer only limited flexibility for adjusting the signal parameters. Built-in modulation generators commonly have a bandwidth of only 40 MHz, which can be limiting, especially when you need to test devices for modern wideband standards such as WiMax. General-purpose AFGs (arbitrary-function generators), on the other hand, now offer bandwidths to 240 MHz and allow direct, fine-resolution control of all signal parameters. The following example illustrates the benefits of using a modern AFG for the physical characterization of IQ modulators.

On the input side, IQ modulators typically accept differential signals, matching the output-signal characteristics of modern high-speed DACs that, in normal use, provide the I and the Q signals. If the signal generator you use for the tests delivers only single-ended signals, you need a separate single-ended-to-differential converter. During characterization, you typically observe the IQ-modulator output on a spectrum analyzer.

To ensure high signal quality, the IQ modulator must, within certain tolerances, symmetrically apply the modulation with respect to gain and phase shift. Otherwise, the receiver cannot properly decode the demodulated IQ signal. Therefore, IQ-imbalance measurements are critical.

To measure the modulator's imbalance, you observe the modulator output on a spectrum analyzer. IQ imbalance results in imperfections in the desired RF spectrum. By adjusting amplitude and phase of the simulated I and Q signals, you can minimize the imperfections. The negative of the generator-adjustment values then reflects the mismatch inherent in the modulator.

Although some vector-signal generators allow manual adjustments of gain and phase skew, the settings often provide only coarse resolution. With other instruments, adjustments can require tedious reloading of vector signals. In comparison, modern AFGs allow direct adjustment of phase and amplitude in every channel with fine setting resolution of 0.01° for phase and four digits for amplitude.

While you make the adjustments, a detailed graphical-waveform representation on the instrument's display confirms the selected output waveforms, as well as their amplitude and phase offset, providing confidence that all relevant settings are correct (Figure A).

As this measurement example illustrates, modern AFGs' performance and flexibility enable the instruments to successfully address applications that only more expensive, dedicated test gear previously addressed. Nevertheless, for measurements that require stimuli at RF or IF, digital modulation, the simulation of complex protocols, or replication of real-world behaviors, high-end AWGs (arbitrary-waveform generators) with high bandwidths and sequencing capabilities remain the tools of choice.

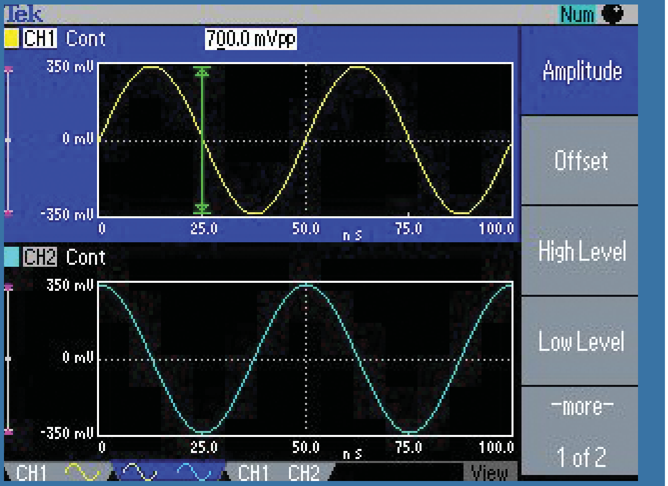

Figure A Though waveform-generator displays are not true scope displays of the output waveforms, some instruments' displays are realistic enough to indicate the 90° phase difference between the I and the Q outputs.

AUTHOR'S BIOGRAPHY

Mark Albert is a market specialist with Tektronix Inc's signal-sources-product line. He has been with the company for nine years and is currently responsible for function and arbitrary-waveform generators. Before joining Tektronix, he worked for more than six years in international product marketing with several high-technology companies. He holds a master's degree in electrical engineering from the Technical University of Darmstadt (Germany) and a master's in business administration from the University of North Carolina–Chapel Hill.