


AN OVERVIEW OF SIGNAL SOURCE TECHNOLOGY AND APPLICATIONS



Tektronix®





AN OVERVIEW OF SIGNAL SOURCE TECHNOLOGY AND APPLICATIONS

Who Should Read This Primer

This publication is aimed at engineers and technicians involved in the development of just about any kind of product with electronic content. If your job responsibilities include characterization, debug, verification, or testing of electronic circuits that carry AC signals (waveforms), then you are very likely to need a waveform source that can adapt to the unique needs of your design.

Often, the first breakthrough in a verification or troubleshooting task is finding a way to replicate the symptoms of a problem in your design. Usually this involves a signal source of some kind. Understanding the many types of signal sources, their features and functions, and their applications can help you choose a signal source that will deliver the waveforms you need. As always, selecting the right tool will make your job easier and will help you produce fast, reliable results.

AN OVERVIEW OF SIGNAL SOURCE TECHNOLOGY AND APPLICATIONS

Contents

Who Should Read This Primer	i
What is a Signal Source?	1
Types of Digital Signal Sources	1
Signal Source Hardware Architecture: The Arbitrary Waveform Generator	3
Doing the Math: Clock Frequency and Memory Depth Calculations	4
Modifying the Waveform: Filtering and Sequencing	5
Creating AWG Waveforms: Tools and Methods	6
Summary	7
Signal Source Hardware Architecture: The Arbitrary Function Generator	9
Understanding Direct Digital Synthesis	9
Summary	10
Signal Source Hardware Architecture: The Data Generator	11
Digital Patterns Differ from Analog Waveforms	12
Sequencer Expands Pattern Length Indefinitely	12
A Sharp Tool for Digital Faults	13
Analog Parameters in a Digital World	14
Making the Connection	15
Summary	15
Signal Source Applications	17
AWG Application: Simulating a Disk Drive Read-Channel Signal	17
AWG Application: Simulating Real-World Aberrations in 100Base-T Physical Layer Signals	19
Data Generator Application: Rambus® Design and Troubleshooting	21
Design margin testing	21
Clock Substitution	21
Testing Complementary Signal Paths	21
Dispelling the Effects of Ground Bounce and Transients	21
Some Signal Source Characteristics and Capabilities to Watch For	22
Characteristics	22
Sample Rate	22
Vertical Resolution	22
Memory Depth	22
Number of Channels	22
Capabilities	22
Modifying AWG Waveforms in Real Time: Integrated Editors	22
Creating Application-Specific Waveforms: Dedicated Functions	22
Expanding the Waveform Length: Sequencing	23
Synchronizing Analog and Digital: Markers	23
Using Waveforms from Other Resources: Importing Data	23
Signal Source Selection Guide	24
The Integrated Tool Set for Superior Measurement and Analysis	back cover

WHAT IS A SIGNAL SOURCE?

A signal source is nothing less than the cornerstone of almost any instrumentation setup used in hardware design, debug, or evaluation projects. It is a key engineering tool. It is an essential troubleshooting aid for the technician. It is a surrogate for an automotive ignition pulse, a heart pacemaker, or a guided missile's gyro output. Second only to the ubiquitous DMM, signal sources are perhaps the most universal class of electronic test instruments.

Unless you're working with a purely DC circuit (and when was the last time you designed a flashlight?), your circuit is likely to require some kind of AC stimulus signal as you evaluate components, functional blocks, and subsystems. The waveform from the signal source emulates a signal coming in from the outside world, such as a sensor output. Similarly, it can be used as a stand-in for waveforms that will appear in as-yet-unavailable parts of the circuit design.

Interestingly, the signal source's job is not simply to provide an "ideal" waveform. Often the instrument must add known, repeatable amounts and types of distortion (or errors) to the signal it delivers. This characteristic is one of the signal source's strongest virtues, since it is often impossible to create predictable distortion exactly when and where it's needed using only the circuit itself. The response of the unit-under-test (UUT) in the presence of these distorted signals reveals its ability to handle stresses that fall outside the normal performance envelope.

A stimulus signal may take the form of a low-distortion sine wave, a stream of logic pulses, a high-frequency radio carrier wave, a mobile telephone transmission, and many other formats. Traditionally, the task of producing these diverse waveforms has been filled by separate, dedicated signal sources, from ultra-pure audio sine-wave generators to multi-GHz RF signal generators. While there are many commercial solutions, the user must often custom-design or modify a signal source for the project at hand. It can be very difficult to design an instrumentation-quality signal generator, and of course, the time spent designing ancillary test equipment is a costly distraction from the project itself.

Fortunately, digital sampling technology and signal processing techniques have brought us a solution that answers almost any kind of signal generation need with just one type of instrument – the arbitrary generator.

Types of Digital Signal Sources

Broadly divided into arbitrary waveform generators (AWG), arbitrary function generators (AFG), and data or pattern generators (DG), digital signal sources span the whole range of signal-producing needs. Each of these types has its unique strengths:

- **AWG:** Whether you want a data stream shaped by a precise Lorentzian pulse for disk-drive characterization, or a complex modulated RF signal to test a GSM- or CDMA-based telephone handset, the AWG can produce any waveform you can imagine. You can use a variety of methods – from mathematical formulae to "drawing" the waveform – to create the needed output.
- **AFG:** Typically this instrument offers fewer waveform variations, but with excellent stability and fast response to frequency changes. If the UUT requires the classic "sine and square" waveforms (to name a few) and the ability to switch almost instantly between two frequencies, the AFG is the right tool. An additional virtue is the AFG's low cost, which makes it very attractive for applications that do not require an AWG's versatility.
- **DG:** This third type of signal source meets the special stimulus needs of digital devices that require long, continuous streams of binary data, with specific information content and timing characteristics. The rationale for this specialized type of signal source will become evident in the architectural discussions that follow.

Earlier, we described signal sources as a class of universal test instruments. With the advent of AWGs and AFGs in particular, the term "universal" takes on an even broader meaning, since many different signal generation tasks conceivably can be handled not by one class of instruments, but by just one instrument.

SIGNAL SOURCE HARDWARE ARCHITECTURE: THE ARBITRARY WAVEFORM GENERATOR

Fundamentally, an Arbitrary Waveform Generator (AWG) is a sophisticated “playback” system that delivers waveforms based on stored digital data that describes the constantly changing voltage levels of an AC signal.^{*1} It is a tool whose block diagram is deceptively simple. In this section we will look at the basic functional blocks, then go on to discuss some of the additional elements that support functions such as sequencing, modulation, etc.

To understand the AWG, first it’s necessary to grasp the broad concepts of digital sampling. Digital sampling is exactly what its name implies: defining a signal using samples, or data points, that represent a series of voltage measurements along the slope of the waveform. These samples may be determined by actually measuring a waveform with an instrument such as an oscilloscope, or by using graphical or mathematical techniques. Figure 1a depicts a series of sampled points. All of the points are sampled at uniform time intervals, even though the curve makes their spacing appear to vary. In an AWG, the sampled values are stored in binary form in a fast Random Access Memory (RAM).

With the stored information, the signal can be reconstructed at any time by reading back the memory locations and feeding the data points through a digital-to-analog converter (DAC). Figure 1b depicts the result. Note that the AWG’s output circuitry filters between the points to connect the dots and create a clean, uninterrupted waveform shape. The UUT does not “see” these dots as discrete points, but as a continuous analog waveform.

Figure 2 is a simplified block diagram of an AWG that implements these operations.

As always, there are some terms and conditions that affect the fidelity of the signal reproduction. Chief among these is the **Nyquist Sampling Theorem**, which states that the sampling frequency must be at least twice that of the highest frequency component

of the sampled signal. To sample a 1 MHz signal, for instance, it is necessary to acquire points at a frequency of at least 2 megasamples per second (MS/s). Although the theorem is usually cited as a guideline for acquisition (as with an oscilloscope), its pertinence to AWGs is clear: stored waveforms must have enough points to faithfully retrace the details of the desired signal.




Figure 1: AWGs use stored digital samples to construct a waveform.




Figure 2: AWG block diagram (simplified).

^{*1} Normally the term “AC” denotes a signal that goes positive and negative about a 0 volt (ground) reference and therefore reverses the direction of current flow once in every cycle. For the purposes of this discussion, however, AC is defined as any varying signal, irrespective of its relationship to ground. For example, a signal that oscillates between +1 V and +3 V, even though it always draws current in the same direction, is construed as an AC waveform. Most signal sources can produce either ground-centered (true AC) or offset waveforms.

An AWG can take these points and read them out of memory at any frequency within its specified limits. If a set of stored points conforms to the Nyquist Theorem and describes a sine wave, then the AWG's output will be a sine wave. However, there is a finite maximum frequency, or sample rate, at which the instrument can operate. This is usually specified in terms of megasamples or gigasamples per second. Today's fastest AWGs can achieve 2.6 GS/s.

Other AWG hardware characteristics, particularly vertical resolution and memory depth, are just as important as sample rate.

Vertical resolution (amplitude) expresses the voltage-measuring precision of the sample points described above. Resolution pertains to the binary word width, in bits, of the instrument's DAC, with more bits equating to higher resolution. While "more is better," higher-frequency AWGs usually have lower resolution – 8 or 10 bits – than general-purpose AWGs offering 12 or 14 bits.

An AWG with 10-bit resolution provides 1024 sample levels spread across the full voltage range of the instrument. If, for example, this 10-bit AWG has a total voltage range of $2 V_{p-p}$, each sample represents a step of approximately 2 mV – the smallest increment the instrument could deliver, assuming it is not constrained by other factors in its architecture.

Memory depth in an AWG plays a key role in the instrument's flexibility. More (deeper) memory provides either of two benefits:

- 1) More cycles of the desired waveform can be stored. This is useful because it reduces the number of "endpoints." An endpoint is the last memory location occupied by the waveform, after which the AWG must wrap around and return to the beginning in order to continue to produce the output signal. There are unavoidable errors that occur at this transition, so it is desirable to minimize the number of endpoints.
- 2) More waveform detail can be stored. Complex waveforms have high-frequency information in their pulse edges and transients. It is difficult to interpolate these fast transitions as we did with the simple, predictable sine wave. To faithfully

reproduce a complex signal, the available waveform memory capacity must be used to store more transitions and fluctuations rather than more cycles of the signal.

Today's state-of-the-art AWGs offer up to 8-Msample memory depth. When combined with the high sample rate that some top models deliver, these instruments can store and reproduce complex RF waveforms, even including pseudo-random bit streams for use in physical-layer testing of network equipment. Similarly, these fast AWGs with deep memory can generate very brief digital pulses and transients.

Doing the Math: Clock Frequency and Memory Depth Calculations

Calculating the frequency of the waveform that an AWG will produce is a matter of solving a few simple equations. Consider the example of an instrument with one waveform cycle stored in memory:

Given a 100 MS/s clock frequency and a memory depth of 4000 samples – an actual sample RAM would have 4096 samples, but we will round that figure to 4000 for the sake of simplicity – then:

$$F_{output} = \text{Clock Frequency} \div \text{Memory Depth}$$

$$F_{output} = 100,000,000 \div 4000$$

$$F_{output} = 25,000 \text{ Hz (or 25 kHz)}$$

Figure 3 illustrates this concept.




Figure 3: At a clock frequency of 100 MHz, the single 4000-point waveform is delivered as a 25 kHz output signal.

At the stated clock frequency, the samples are about 10 ns apart. This is the time resolution (horizontal) of the waveform. Be sure not to confuse this with the amplitude resolution (vertical) described earlier.

Carrying this process a step further, assume that the sample RAM contains not one, but four cycles of the waveform:

$$F_{output} = (Clock\ Frequency \div Memory\ Depth) \times (cycles\ in\ memory)$$

$$F_{output} = (100,000,000 \div 4000) \times (4)$$

$$F_{output} = (25,000\ Hz) \times (4)$$

$$F_{output} = 100,000\ Hz$$

The new frequency is 100 kHz, and we have reduced the number of endpoint transitions. Figure 4 depicts this concept.

In this instance, the time resolution of each waveform cycle is lower than that of the single-waveform example – in fact, it is exactly four times lower. Each sample now represents 40 ns in time. The increase in frequency (and the reduction in endpoint aberrations) comes at the cost of some horizontal resolution.




Figure 4: Using four stored waveforms and a 100 MHz clock, a 100 kHz signal is produced.

Modifying the Waveform: Filtering and Sequencing

Once the basic waveform is defined, other operations can be applied to modify or extend it.

Filtering allows you to remove selected bands of frequency content from the signal. For example, when testing an analog-to-digital converter (ADC), it is necessary to ensure that the analog input signal (which comes from the AWG) is free of frequencies higher than the converter’s clock frequency. This prevents unwanted aliasing distortion in the UUT output, which would otherwise compromise the test results.

One reliable way to eliminate these frequencies is to apply a steep low-pass filter to the waveform. This allows frequencies below a specified point to pass through, and drastically attenuates those above the cutoff.

Filters also can be used to “roll off” pure waveforms such as square and triangle waves. Sometimes it’s simpler to modify an existing waveform in this way than to create a new one.

In the past, it was necessary to use a signal generator and an external filter to achieve results such as these. Fortunately, the better present-day laboratory AWGs feature accurate built-in filters.

Waveform sequencing is another valuable enhancement to the basic AWG architecture. In effect, it allows you to store a huge number of “virtual” waveform cycles in the instrument’s memory. The waveform sequencer borrows instructions from the computer world: loops, jumps, and so forth. These instructions, which reside in a sequence memory separate from the waveform memory, cause specified segments of the waveform memory to repeat.

To cite a very simple example, imagine that the 4000-point memory described earlier is loaded with a clean pulse that takes up half the memory (2000 points), and a distorted pulse that uses the remaining half. If we were limited to basic repetition of the memory content, then the AWG would always repeat the two pulses, in order, until commanded to stop. But waveform sequencing changes all that.

Suppose you wanted the distorted pulse to appear twice in succession after every 511 cycles. You could write a sequence that repeats the clean pulse 511 times, then jumps to the distorted pulse, repeats it twice, and goes back to the beginning to loop through the steps again... and again. Figure 5 explains this premise. Loop repetitions can go into the hundreds of millions. Given what we have already discussed about the tradeoff between the number of cycles stored and the resulting horizontal resolution, it's clear that sequencing provides much-improved flexibility without compromising the resolution of individual waveforms.

Note here that any sequenced waveform segment must continue from the same amplitude point as the segment preceding it. In other words, if a sine wave segment's last sample value was 1.2 volts, the starting value of the next segment in the sequence must be 1.2 volts as well. Otherwise, an undesirable glitch can occur when the DAC attempts to abruptly change to the new value.

Although this example is very basic, it represents the kind of capability that is needed to detect irregular pattern-dependent errors and co-symbol interference in communications circuits. With waveform sequencing, it's possible to run long-term stress tests – extending to days or even weeks – with the AWG as the stimulus source.

Creating AWG Waveforms: Tools and Methods

We have discussed the AWG's ability to produce long, complex waveforms. But this would be almost useless if the engineer had to key in point-by-point values for every location in the waveform memory. Fortunately, there are many labor-saving approaches to waveform creation.

Most lab-quality AWGs are furnished with certain standard waveforms, either stored in the instrument's local nonvolatile memory, or provided on disk media. These waveforms make an excellent starting point for modification via the AWG's editing tools.

One of the most efficient ways to derive a new waveform is to “learn” it from a digital sampling oscilloscope (DSO). If, for example, stimulus signals are needed for testing a newly-developed product in manufacturing, the DSO can be used to capture the waveforms from a known-good engineering prototype. The DSO's acquisition memory contents can then be copied into the AWG's sample memory via a GPIB or Ethernet connection between the two instruments. Assuming the AWG has sufficient bandwidth and resolution, it can duplicate the waveform exactly. Perhaps just as important, the AWG's waveform can be modified with known impairments, if necessary.

A second waveform creation method is the AWG's built-in math tool. All “standard” waveforms (sine, square, triangle, etc.) can be derived from relatively simple mathematical formulae. Some AWGs have math editors that apply convolution, integration, and other operations. As an adjunct to this on-board math capability, many AWGs accept data from PC-based waveform simulation tools.




Figure 5: AWG waveform memory capacity can be “expanded” with loops and repetitions.

The third, and perhaps the most intriguing, waveform editing method is the built-in graphical editor. Some AWGs offer a means of actually drawing (on the instrument's own display screen) the waveform features you need. Of course there are finite limitations – you still can't make an edge transition in 0 ps! – but the instrument will reproduce all of the waveform details within its performance range. Figure 6 is a typical graphical waveform editing display.

Some AWGs can actually deliver an analog waveform output and a correlated digital signal at the same time. These instruments have “marker” outputs that can provide either a trigger pulse or a programmed binary expression when a specific sample value is clocked out of the memory.

Summary

The arbitrary waveform generator (AWG) offers a degree of versatility that few instruments can match. With its ability to produce any waveform you can imagine, the AWG embraces applications ranging from automotive ABS simulation to wireless network stress testing. The AWG is a tool that no design or development lab should be without.




Figure 6: AWG graphical waveform editor.

SIGNAL SOURCE HARDWARE ARCHITECTURE: THE ARBITRARY FUNCTION GENERATOR

The arbitrary function generator (AFG) shares many features with the AWG, although the AFG is, by design, a more specialized instrument with a narrower range of applications. The AFG offers unique strengths: it produces stable waveforms in standard shapes – particularly the all-important sine and square waves – that are both accurate and agile. Agility is the ability to change quickly and cleanly from one frequency to another.

Most AFGs offer some subset of the following familiar wave shapes:

- Sine
- Square
- Triangle
- Sweep
- Pulse
- Ramp
- Modulation
- Noise
- Haversine

While AWGs can certainly provide these same waveforms, today's AFGs are designed to provide improved phase, frequency, and amplitude control of the output signal. Moreover, many AFGs offer a way to modulate the signal from internal or external sources, which is essential for some types of standards compliance testing.

In the past, AFGs created their output signals using analog oscillators and signal conditioning. More recent AFGs rely on a completely different architec-

ture that is digital in nature and has much in common with the AWG. The same basic building blocks are there: waveform memory, DAC, and output signal conditioning. And like the best AWGs, the AFG uses a technique known as Direct Digital Synthesis (DDS) to determine the rate at which samples are clocked out of the memory. Since we didn't cover this architecture in the AWG section, we'll focus on it now.

Understanding Direct Digital Synthesis

DDS technology synthesizes waveforms by using a single clock frequency to spawn any frequency within the instrument's range. Figure 7 summarizes the DDS-based AFG architecture in simplified form.

In the phase accumulator circuit, the Δ (Delta) phase register receives instructions from a frequency controller, expressing the phase increments by which the output signal will advance in each successive cycle. In a modern high-performance AFG, the phase resolution may be as small as one part in 2^{30} , that is, approximately 1/1,000,000,000.

The output of the phase accumulator serves as the clock for the waveform memory portion of the AFG. The instrument's operation is almost the same as that of the AWG, with the notable exception that the waveform memory typically contains just a few basic signals such as sine and square waves. The analog output circuit is basically a fixed-frequency low-pass filter which ensures that only the programmed




Figure 7: AFG block diagram (simplified).

frequency of interest (and no clock artifacts) leaves the AFG output.

To understand how the phase accumulator creates a frequency, imagine that the controller sends a value of “1” to the 30-bit Δ phase register. The phase accumulator’s Δ output register will advance by $360 \div 2^{30}$ in each cycle, since 360 degrees represents a full cycle of the instrument’s output waveform.

Therefore, a Δ phase register value of “1” produces the lowest-frequency waveform in the AFG’s range, requiring the full 2^{30} increments to create one cycle. The circuit will remain at this frequency until a new value is loaded into the Δ phase register.

Values greater than one will advance through the 360 degrees more quickly, producing a higher output frequency (some AFGs use a different approach: they increase the output frequency by skipping some samples, thereby reading the memory contents faster). Note that the only thing that changes is the phase value supplied by the frequency controller. The main clock frequency does not need to change at all. In addition, it allows a waveform to commence from any point in the waveform cycle.

For example, assume it is necessary to produce a sine wave that begins at the peak of the positive-going part of the cycle. Basic math tells us that this peak occurs at 90 degrees. Therefore:

$$2^{30} \text{ increments} = 360^\circ; \text{ and}$$

$$90^\circ = 360^\circ \div 4; \text{ then,}$$

$$90^\circ = 2^{30} \div 4$$

When the phase accumulator receives a value equivalent to $(2^{30} \div 4)$, it will prompt the waveform

memory to start from a location containing the positive peak voltage of the sine wave.

As explained earlier, the typical AFG has just a few types of waveforms stored in its memory. In general, sine and square waves are the most widely used for many test applications. Some AFGs are “hybrid” units that deliver accurate, agile DDS-based sine and square waves, while other wave shapes are created using more conventional AWG techniques. This is a cost-effective solution that puts the highest performance behind the most critical functions.

DDS architecture provides exceptional frequency agility, making it easy to program both frequency and phase changes on the fly. Where is this important? Certainly it’s useful for testing any type of FM (frequency-modulated) UUT device – radio and satellite system components, for example. And if the specific AFG’s frequency range is sufficient, it’s an ideal signal source for test on FSK (frequency shift keying) and frequency-hopping telephony technologies such as GSM.

Summary

The arbitrary function generator (AFG) serves a wide range of stimulus needs; in fact, it is the prevailing signal-source architecture in the industry today. Although it cannot equal the AWG’s ability to create virtually any imaginable waveform, the AFG produces the most common test signals used in labs, repair facilities, and design departments around the world. Moreover, it delivers excellent frequency agility. And perhaps most important to many users, the AFG is available at a fraction of the price of top AWG models.

SIGNAL SOURCE HARDWARE ARCHITECTURE: THE DATA GENERATOR

A third type of signal source, the data generator (DG), is a more specialized tool for those with specific digital test requirements. Where the AWG and AFG are primarily designed to produce waveforms with “analog” shapes and characteristics, the DG’s mission is to generate volumes of binary information. Also known as a pattern generator, the DG produces the streams of 1s and 0s needed for testing computer busses, microprocessor IC devices, and other digital systems.

Because it is dedicated to digital testing applications, there might be a tendency to infer that the DG is a “limited” tool. After all, an AWG can generate pattern data but the DG cannot generate analog waveforms! The DG was never meant to be a universal tool; instead, its features give it unique strengths that neither AWGs nor AFGs can match.

Among these features are:

- **Sequencing:** An absolute necessity in the world of data and pattern generation. No internal memory can be deep enough to store the many millions of pattern words (also known as “vectors”) required for a thorough digital-device test. Consequently, DGs are equipped with sophisticated sequencers, far more so than those of other types of signal sources.
- **Multiple outputs:** Where the AWG or AFG may have two or four outputs, the DG may have hundreds of output channels. This is to support the numerous data and/or address inputs of the typical digital device.
- **Pattern data sources:** The modern DG must accept data from logic analyzers, DSOs, simulators, and even spreadsheets. Why? Because a complex digital pattern would be impossibly tedious and error-prone if entered by hand.

Moreover, digital data is usually available from various simulation and verifications steps in the design process.

- **Display:** The DG display must emphasize the details of many channels of pattern data simultaneously rather than the details of signal amplitude vs. time (as with the AWG display). It should offer markers, scrolling, and other time-saving features to help the user focus on the data of interest. Figure 8 is an example of a multi-channel DG display.




Figure 8: Multi-channel DG bus timing display.

Digital Patterns Differ from Analog Waveforms

Some aspects of the DG architecture will look familiar after reading the earlier sections on AWGs and AFGs. Figure 9 depicts the DG block diagram. Once again there is an address generator, a waveform (or pattern) memory, a shift register, etc. But notice one significant omission: the DAC is absent!

The DAC is not necessary because the DG doesn't need to trace out the constantly shifting levels of an analog waveform. Although the DG has an analog output circuit, this circuit is used to set voltage and edge parameters that apply to the whole pattern. For example, most DGs provide a way to program the logic "1" and "0" voltage values for the pattern.

Sequencer Expands Pattern Length Indefinitely

The DG's sequencer bears some similarity to that of the AWG described earlier. The DG must provide tremendously long and complex patterns, and must respond to external events – usually a UUT output condition that prompts a branch execution in the DG sequencer.

The DG's pattern memory capacity, typically about 256 Kbits maximum, may seem modest in comparison to the largest AWG waveform memory. But like the AWG sequencer, the DG sequencer can loop on short pattern segments to produce a data stream of much greater length. It can wait for an external event or trigger, then execute a series of repeat counts or




Figure 9: DG block diagram (simplified).

conditional jumps. In addition, the DG sequencer offers many levels of loop nesting and branch conditions. It is a tool that can be controlled using normal programming conventions, and produces address, data, clock, and control signals for almost any imaginable digital device. Follow the flow in Figure 10 to see how a few short instructions and pattern segments can unfold into millions of lines of stimulus data.

A Sharp Tool for Digital Faults

Although the DG lacks a DAC, it has some digital features not found in most AWGs or AFGs. The most critical use of these features resides in the delay circuit that precedes the analog output in Figure 9. The delay circuit is responsible for implementing the

small changes in edge positioning that support both jitter and timing tests.

The delay circuit can deliver tiny changes (on the order of picoseconds) in edge placement. For example, some state-of-the-art DGs provide simple front-panel controls that allow you to move all edges or selected edges in 5 ps steps within a range of ± 100 ps. These small timing changes model the classic jitter phenomenon in which the placement in time of a pulse edge moves erratically about a nominal center point. You can test jitter tolerance by changing and observing the effects of edge timing in relation to the clock.

In today's best DGs, it's possible to apply this jitter throughout the pattern, or on isolated pulses via a "masking" function that pinpoints specific edges.




Figure 10: DG sequence flow, showing loops and conditional events.

Figure 11 shows a digital phosphor oscilloscope (DPO) capture of a DG output signal with the addition of the jitter effect. The inset illustration provides a simplified and enlarged view of the same events. Other features give the modern DG even more flexibility for critical jitter testing. Some instruments have an external analog modulation input that controls both the amount of edge displacement (in picoseconds) and the rate at which it occurs. With so many jitter variables at your disposal, it's possible to subject the UUT to a wide range of real-world stresses.




Figure 11: The DG uses small timing shifts to simulate jitter.




Figure 12: Programmed voltage variations on a DG output signal.

The delay circuit plays a second, equally important role in testing for timing problems such as setup-and-hold violations. Most clocked devices require the data signal to be present for a few nanoseconds before the clock pulse appears (setup time), and to remain valid for a few nanoseconds (hold time) after the clock edge. The delay circuit makes it easy to implement this set of conditions. Just as it can move a signal edge a few picoseconds at a time, it can move that edge in hundreds of picoseconds, or in nanoseconds. This is exactly what is needed to evaluate setup and hold time. The test involves moving the input data signal's leading and trailing edges, respectively, a fraction of a nanosecond at a time while holding the clock edge steady. The resulting UUT output signal is acquired by a DSO or logic analyzer. When the UUT begins to put out valid data consistent with the input condition, the location of the leading data edge is the setup time. This approach can also be used to detect metastable conditions in which the UUT output is unpredictable.

Analog Parameters in a Digital World

Although the DG's repertoire does not include common signal conditioning operations such as filtering and modulation, it nevertheless offers some tools to manipulate the output signal. These features are needed because digital design problems are not limited to purely digital issues such as jitter and timing violations. Some design faults are the result of analog phenomena such as erratic voltage levels or slow edge rise times. The DG must be able to simulate both.

Voltage variations in the stimulus signal are a key stress-testing tool. By exercising a digital UUT with varying voltage level, especially levels immediately below the device's logic threshold, it is possible to predict the device's performance and reliability as a whole. A UUT with intermittent (and difficult to trace) failures will almost certainly turn into a "hard" failure when the voltage is reduced.

Figure 12 depicts the effect of programming a DG to produce several discrete logic levels. Here the results of several instructions are shown cumulatively, but in reality, the instrument applies a single voltage level throughout the pattern.

Edge transition times, or rise times, are another frequent cause of problems in digital designs. For example, a pulse with a slow edge transition may not trigger the next device in line in time to clock in data. Slow edges are notorious for causing race conditions, another cause of intermittent failures. A host of cumulative design factors, notably distributed reactances, can degrade the rise time of a pulse as it travels from source to destination. Therefore, engineers try to ensure that their circuits can handle a range of rise times. Like the voltage variations described earlier, slowing down the pulse edge rate is part of every stress and margin testing plan. Since the DG is a common tool in digital design environments, it's often called upon to simulate these transition-time problems.

Looking again at the DG block diagram in Figure 9, notice the input labeled “Transient Control” that feeds the analog output circuit. It is here that the user can program a broad range of edge rates for the instrument’s output signal.

Figure 13 illustrates the effect of the programmable edge-rate feature.

Making the Connection

As we have discussed, DGs are often used in tests that require critical pulse-edge characteristics, including voltage accuracy, rise-time performance, and edge placement. Unfortunately, simply providing a high-quality signal at the instrument’s front-panel connector is not enough. Often the signal must travel to a test fixture a meter or more away from the instrument, through cables and connectors that can seriously degrade the signal’s timing and edge details. Some modern DGs solve this problem with an external signal interface that buffers the signal and brings the performance of the instrument all the way out to the UUT.

Figure 14 shows a typical DG equipped with a signal interface. The interface minimizes rise-time degradations due to cable capacitance and provides ample local current to drive a UUT input without “loading down.” Interestingly, the external interface is the connection point for signals from the UUT, notably “Inhibit” instructions that cause the DG output to switch to a high-impedance state.

Summary

Data Generator Applications Span Many Technologies. In the design department, the DG is an indispensable stimulus source for almost every class of digital device. In broader terms, the DG is useful for functional testing, debug of new designs, and failure analysis of existing ones. It’s also an expedient tool to support timing and amplitude margin characterization.

The DG can be used early in the product development cycle to substitute for system components that are not yet available. For example, it might be programmed to send interrupts and data to a newly developed bus circuit when the processor that would normally provide the signals doesn’t yet exist. Similarly, the DG might provide addresses to a memory bus, or even the digital equivalent of a sine wave to a DAC under test.

With its extraordinarily long patterns and its ability to implant occasional errors in the data stream, the DG can support long-term reliability tests to ensure compliance with military or aerospace standards. In addition, its ability to respond to external events from the UUT as part of the pattern sequence provides even more flexibility in demanding characterization applications.

The DG is equally at home testing semiconductor devices such as ASICs and FPGAs, or rotating media – hard-disk drive write circuits and DVDs. Likewise, it’s useful for testing CCD image sensors and LCD display drivers/controllers. The DG is an effective solution just about anywhere a complex digital bit stream is needed to stimulate a UUT.




Figure 13: Programmed rise-time variations on a DG output signal.



Figure 14: DGs and external signal interface “pod.”

AWG Application: Simulating a Disk Drive Read-Channel Signal

With their ever-increasing speed and bit density, hard-disk drives (HDD) pose one of the most formidable test challenges existent today. The AWG plays an enabling role in the design and evaluation of new generations of HDDs. It provides signals that model the response of a read channel, permitting engineers to move ahead with the design of all the ancillary circuitry that delivers the disk data to the outside world.

The first step is to create an emulation of the transition response that occurs when a read head acquires data from the magnetic medium. This signal takes on the classic shape of a Lorentzian pulse.

To create the Lorentzian pulse, we can use an AWG with an Equation Editor. The rationale for the specific Lorentzian pulse formula is beyond the scope of this discussion^{*2}, but suffice it to say that the following formula is required:

$$v(t) = \frac{1}{1 + \left(\frac{2(t-10)/T}{PW50/T} \right)^2}$$

Where: t = time

T is a parameter representing the bit interval

PW50 is the pulse width at which the amplitude is at 50% of the peak voltage

Figure 15 shows this same formula as it appears on the screen display, along with an inset image that presents the resulting Lorentzian pulse wave shape.

While entering the above formula will certainly produce the desired waveform, there may be a more expedient way to get the job done. Some AWGs are supplied with a library of waveforms, either stored within a ROM in the instrument or on a floppy disk. Usually, it's much simpler to modify one of these factory-supplied waveforms than to design your own.

The next step is to create the actual binary data content using the AWG's Pattern Editor tools. The pattern has the fundamental period of the bit interval T (from the formula above) and should reflect any coding algorithms, such as RLL, that need to be applied. In this example, we will use a 62-bit NRZ (Non-Return to Zero) pattern. Figure 16 shows the Pattern Editor screen

After the pattern is entered via the editor, it's necessary to scale the signal to match the target bit interval. For the sake of simplicity, assume the data rate is 650 Mbps, so that each pattern bit represents 1.5 ns. Using an AWG with a 2.6 GS/s clock rate,

each signal point represents six data points in the AWG memory. It's necessary at this point to expand the pattern to a 248-point record.

The image inset in Figure 16 shows the resulting data stream.

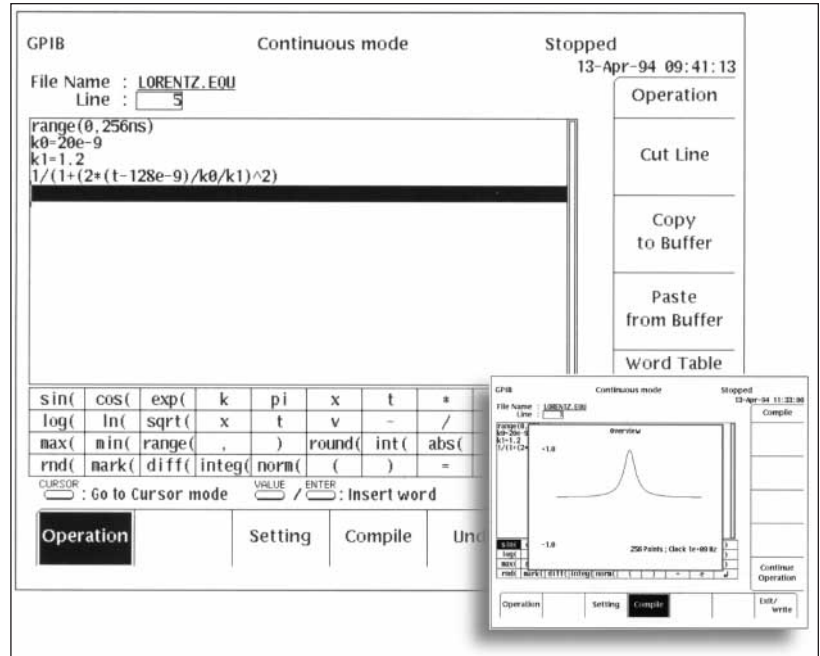


Figure 15: The AWG equation editor simplifies creation of the Lorentzian pulse.

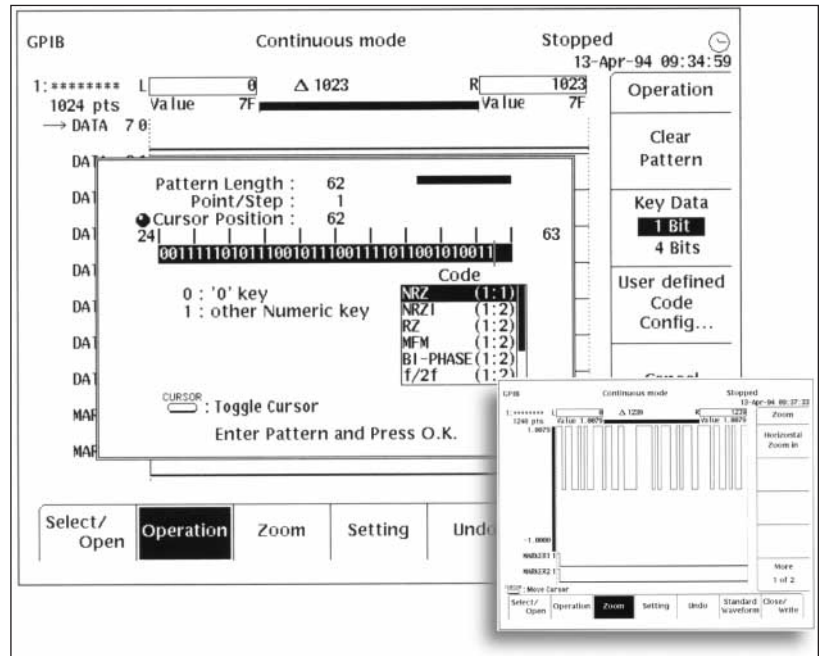


Figure 16: The Pattern Editor's tools create the needed binary information.

^{*2} For more information, see *Signals and Measurements for Disk Drive Design* at www.tektronix.com/Masurement/App_Notes/.




Figure 17: The hard disk Read channel signal.

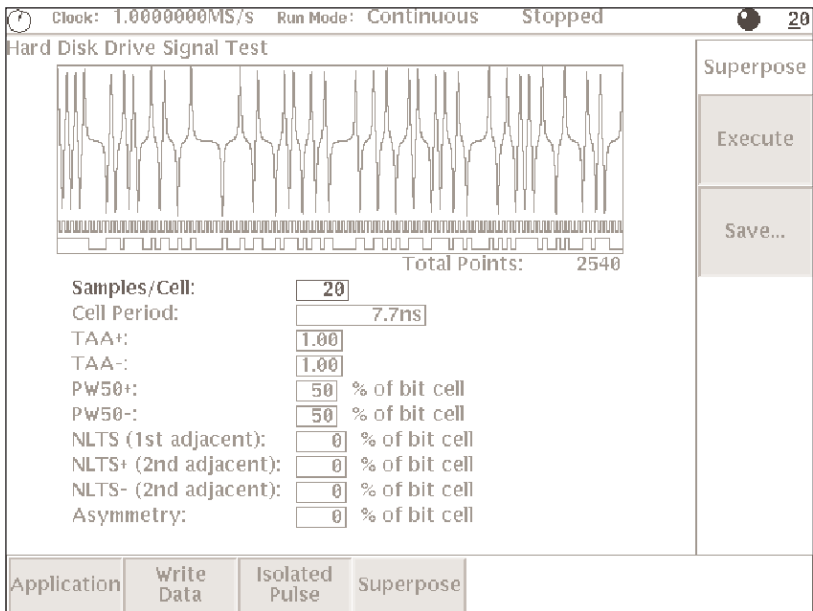


Figure 18: Advanced disk-drive application editor offered by the AWG610.

The final step in this process is to convolve the two signals – the Lorentzian pulse and the data stream – to produce a detailed model of a genuine read-channel signal. This could be a complex mathematical process, but once again, the AWG's built-in tools come to the rescue. State-of-the-art AWGs incorporate a convolution tool that includes options such as differential math. This is exactly what we need to produce the final waveform shape for the HDD read channel.

The differential math is needed because the binary waveform is a typical stream of 1 and 0 states, while the read channel's information content is imparted by transitions, not states. The differentiation ensures that the final data reflects this reality. The result is shown in Figure 17.

There is an alternative approach to that described above. The AWG's equation editor provides the tools to enter the waveform in mathematical terms. The editor supports polynomial formulas made up of many operators, functions, and variables. Although it sounds complex, this approach can actually simplify the design of complex waveforms for certain applications.

In addition to the steps outlined here, others may be required to "fine tune" the waveform for its ultimate use in a read channel. But the foregoing explanation has shown you how an AWG can generate and manipulate detailed read-channel waveforms.

Figure 18 shows the advanced disk-drive application editor provided by the AWG610 and AWG500 series. The disk-drive application editor uses a fill-in-the-blanks approach to quickly and easily create read-channel signals for simulation and test of read-channel amplifiers and other devices.

AWG Application: Simulating Real-World Aberrations in 100Base-T Physical Layer Signals

To simulate physical layer test signals for 100Base-T transceivers, you must consider a wealth of analog parameters: undershoot and overshoot, rise and fall time, ringing, amplitude variations, and specific timing variations such as jitter. AWGs provide an efficient method for generating signal impairments like these and many more. Using an AWG to replicate 100Base-T Ethernet signal impairments is the best way to stress components to the limits (margins) of their performance under real-world conditions.

The starting point for such a margin test is the 100Base-T signal shown in Figure 19, which was created by a modern AWG with a built-in Network Signal Application Editor. The basis of this signal is a PN9 pseudo-random data stream. The screen shown here is from an oscilloscope measurement, not the AWG itself.

The AWG display screen in Figure 20 highlights some of the built-in features of the application editor used to create a 100Base-T signal. Notice that you may choose any of several protocol standards, from ITU-T to SDH/SONET, as well as the bit rate, line code (MLT-3 in this example), and other parameters.




Figure 19: An application-specific network signal editor on an AWG was used to create this 100Base-T signal.

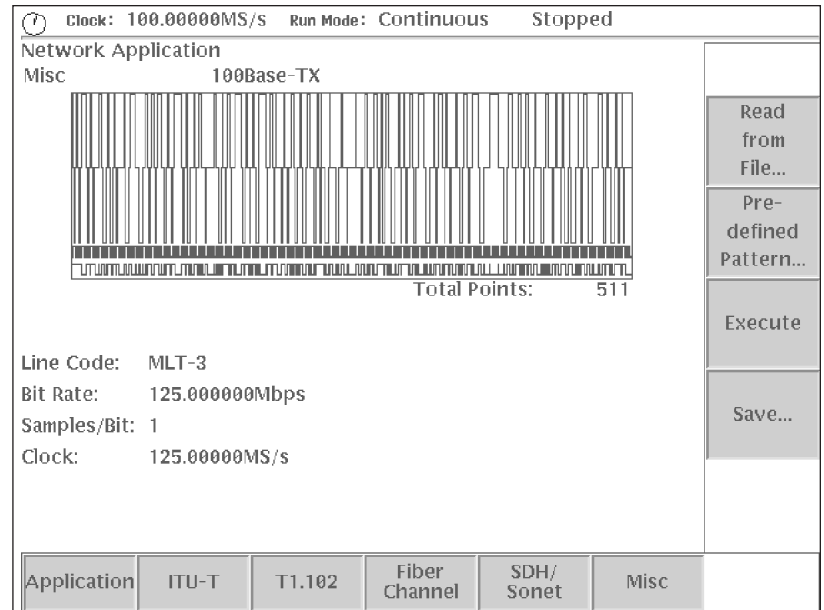
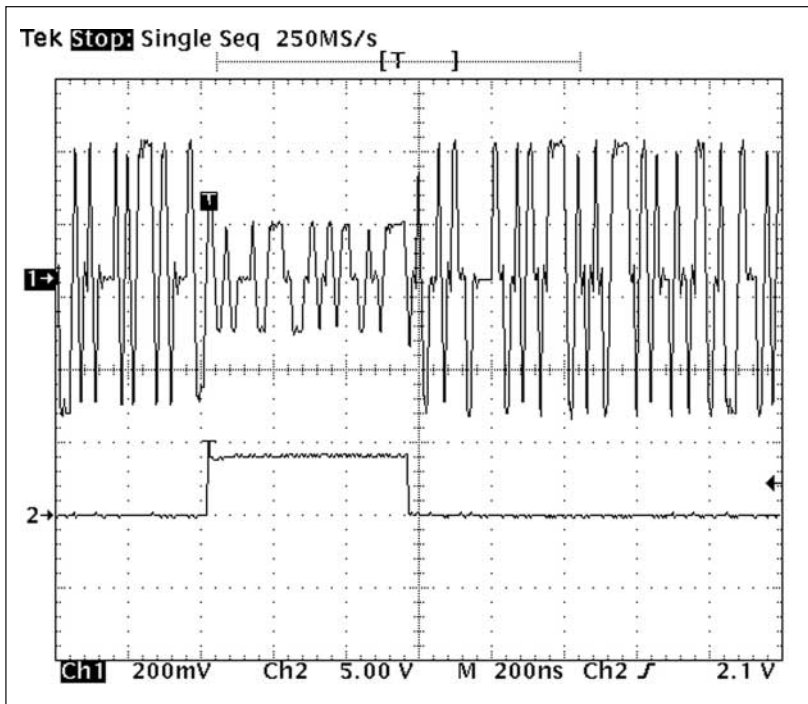


Figure 20: The network signal application editor screen.



In Figure 21, the 100Base-T signal shows the effects of amplitude variations applied with a real-time editing tool built into the AWG.

By adding momentary amplitude attenuations like this, it is possible to characterize the UUT's response to varying degrees of signal loss due to harsh real-world conditions, and adapt the design accordingly. Note also the lower trace, which is the AWG's marker output. It was used here to trigger the oscilloscope which captured this display.

Figure 21: Real-time editing tools can be used to apply brief signal variations.

Data Generator Application: Rambus® Design and Troubleshooting

Demand for video and audio web applications on personal computers has spawned a requirement for high-bandwidth memory that can keep pace with other system elements, especially processors and buses. In answer to this need, Rambus®, Inc. has developed a new memory architecture known as the (RDRAM™) which can outperform current SDRAM technology by a factor of ten.

Rambus' approach to providing increased memory system performance is based on a high-speed, chip-to-chip interface that transfers data over a relatively narrow bus known as the Rambus Channel. RDRAM devices support a transfer rate of 800 Mbps on each of the data pins.

At these tremendous data rates, tolerances can build up, ground bounce phenomena can emerge, and signals can deteriorate. It may be difficult to pin the blame on any one part of the circuit.

We have talked about measuring setup/hold times and jitter, for example. But what is the recourse when there is too much jitter? Jitter is a symptom – what is causing the symptom? Other potential problems range from improper clock and data waveform symmetry to excessive noise on the transmission lines.

The modern DG can augment measurement tools such as oscilloscopes and logic analyzers to resolve Rambus design problems. They can help find the root cause of problems that seem to have no clear origin. The process is one of signal substitution and fault isolation.

To be effective in the Rambus environment, the DG must deliver exceptional performance at high frequencies. Data rates in the range of 1.1 Gb/s are essential. In addition, the DG must provide an ultra-low-jitter signal (preferably in the range of 3 ps) with edge risetimes of 150 ps or less. To carry out certain stress tests, the DG must deliver a wide range of signal amplitudes – approximately $0.25 V_{p-p}$ to $2.5 V_{p-p}$. Programmable channel-to-channel delay and complementary outputs are also useful for the differential signal lines common in Rambus architecture. Following are some applications that make good use of a DG:

Design margin testing

Even a design that “works” has its limits. What are those limits? It's mandatory to confirm that the design offers sufficient margin to manage the vagaries of the manufacturing process. Tolerance buildup in clock signal path terminations and device AC parameters is another common problem in high-speed circuits.

A DG can stress system elements with signals that are too early, too late, too “small” (in amplitude), and so on.

Clock Substitution

Memory-access failures can have many causes. One such cause is setup and hold timing violations due to excessive clock jitter. By substituting an external low-jitter source for the clock that feeds the RDRAMs, it's possible to “divide and conquer” the circuit problem.

Figure 22 is a jitter characterization of a DG output, measuring only 12.170 ps jitter. The measurement was taken with a high-bandwidth oscilloscope. Since actual Rambus Channel jitter is in the 70 ps range (or higher), it's obvious that this DG is capable of handling Rambus clock substitution chores.

Testing Complementary Signal Paths

A DG with programmable complementary outputs can provide paired signals to evaluate the effects of amplitude variations, skew between Rambus Channel Clock To Master (CTM) and Clock From Master (CFM) signals, and symmetry referenced to V_{ref} .

Dispelling the Effects of Ground Bounce and Transients

Simultaneous bus access across a multi-channel Rambus system draws far more current than usual, causing ground bounce and possibly voltage transients that are echoed in the CTM signal line.

An external DG can usually provide much greater current source and sink capacity than the on-board clock generator. This extra output capacity offers greater resistance to waveform deformation, permitting the designer to evaluate the effects of a more perfect clock in the circuit.

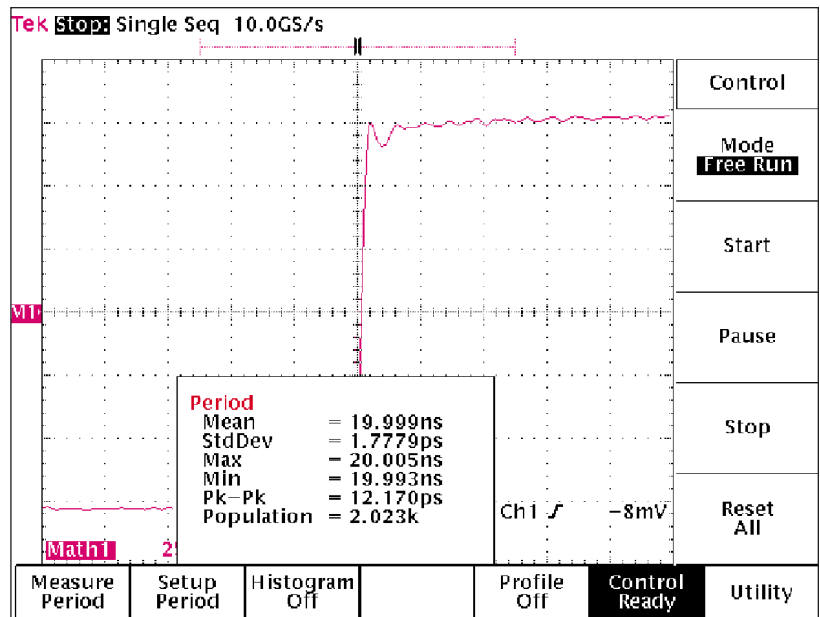


Figure 22: A low-jitter signal from an external DG can be used for Rambus clock substitution during troubleshooting.

SOME SIGNAL SOURCE CHARACTERISTICS AND CAPABILITIES TO WATCH FOR

Characteristics

A host of factors must be weighed when choosing a signal source. Of course, the “banner” specifications are important. However, you need to also look at the instrument’s adaptability to your particular application. It’s very easy to be swayed by sample rates and resolution specifications when the tool you really need is the one that can accept your EDA simulator data.

Most published specifications dwell on the key measures of signal source performance – the “banner” specifications – leaving the determination of adaptability up to your discernment.

The key specifications, in order of importance, are as follows:

Sample Rate

The maximum clock or sample rate at which the instrument can operate. The sample rate controls the frequency and fidelity of the main output signal. The Nyquist Theorem states that the clock rate must be at least twice the frequency of the desired signal’s highest component frequency to ensure accurate reproduction. Four, eight, or even 100 times over-sampling may be necessary in some applications.

Vertical Resolution

The vertical resolution of the DAC defines the amplitude accuracy and distortion of the reproduced waveform. A DAC with inadequate resolution contributes to quantization errors, causing imperfect waveform generation.

Memory Depth

Memory depth limits the length of a non-repeating waveform. It contributes to signal fidelity because using a low-frequency signal (relative to the sample clock rate) permits more data points to be stored per cycle. Without sufficient memory depth, you cannot take advantage of the high relative sample rate. Modern AWGs offer 2, 4, even 8 Msamples of memory.

Number of Channels

Many applications require more than one output channel. AWGs with up to four independent channels are available. This feature is indispensable for simulating sensors in automotive (anti-lock brakes)

or bio-electrical applications, and more. Similarly, dual-channel instruments can generate complex digital modulation signals such as I&Q modulated waveforms. Some AWGs have one or two analog outputs and up to 14 high-speed digital outputs for mixed-signal testing.

Capabilities

After you’ve determined that a signal source meets your needs with regard to the basic specifications, consider some of the other features that can make the instrument easier to use, more productive, or more cost-effective for your application:

Modifying AWG Waveforms in Real Time: Integrated Editors

Suppose you need a series of waveform segments having the same shape but different amplitudes as the series proceeds. To create these amplitude variations, recalculate the waveform or redraw it using an off-line waveform editor. But both approaches are unnecessarily time-consuming and error-prone. A better method is to use an integrated editing tool that can modify the waveform memory in both the time and amplitude domains. These editors automatically update the sample memory to deliver near real-time changes in the output signal.

Creating Application-Specific Waveforms: Dedicated Functions

AWGs are ideal for generating complex test waveforms for high-speed disk drive and network communication applications. Some AWGs provide application-specific signal creation functions for PRML signals with NLTS characteristics for the disk-drive industry. Others support network data signal generation with built-in 100baseT or Gigabit Ethernet signals.

Predictable jitter simulation is an application that cuts across many industries. The ability to inject jitter into the waveform file can be a time-consuming, math-intensive effort. Some of today’s state-of-the-art AWGs offer special “jitter editors” to assist in adding jitter to any appropriate signal. These editors make a complex process simple: define the digital data stream, select the jitter profile, and enter the jitter deviation and frequency. The instrument calculates the changes and saves them to the waveform file.

Expanding the Waveform Length: Sequencing

Often, it's necessary to create long waveform files to fully exercise the UUT. Where portions of the waveform are repeated, a waveform sequencing function can save you a lot of tedious, memory-intensive waveform programming. Sequencing allows you to fill in a table listing the waveforms in order of appearance. Programmable repeat counters, branching on external events, and other control mechanisms determine the number of operational cycles and the order in which they occur. With a sequence controller, you can generate waveforms of almost unlimited length.

Synchronizing Analog and Digital: Markers

Marker outputs provide a binary signal that is synchronous with the main analog output signal of

an AWG. In general, markers allow you to output a pulse (or pulses) synchronized with a specific waveform memory location. Marker pulses are usually used to synchronize the digital portions of a UUT while that unit is simultaneously stimulated by the analog signal.

**Using Waveforms from Other Resources:
Importing Data**

Data import functions allow you to use waveform files created outside the AWG. For example, a waveform captured by a modern digital storage oscilloscope can be easily transferred via GPIB or Ethernet to the AWG. All of the AWG's editing tools are available to manipulate the signal, just like any other stored waveform. Simulators and other EDA tools are another expedient source of waveforms.

SIGNAL SOURCE SELECTION GUIDE

AWG & AFG SELECTION GUIDE

Model	Maximum Clock Rate	Output Channels	Maximum Amplitude*	Region Shift	Vertical Direct Resolution	DSO or DPO Transfer	Parallel Digital Output
AWG610	2.6 GS/s	1 **	2	400 fs	8 Bits	Yes	2 ***
AWG520	1 GS/s	2	2	2 ps	10 bits	Yes	10 bits & 4 Markers
AWG510	1 GS/s	1	2	2 ps	10 bits	Yes	10 bits & 2 Markers
AWG2021	250 MS/s	1 Std, 2 opt	5	4 ps	12 bits	Yes	12 ECL or 12/24 TTL & 2 markers
AWG2005	20 MS/s	2 Std, 4 opt	10	50 ps	12 bits	Yes	24 TTL & 2 Markers
AFG320	16 MS/s	2	10	Phase Shift	12 bits	Yes	No
AFG310	16 MS/s	1	10	Phase Shift	12 bits	Yes	No

* Maximum Amplitude – V_{p-p} into 50 ohms

** Complementary

*** Independently programmable digital output markers

DATA GENERATOR SELECTION GUIDE

Model	Maximum Data Rate	Output Channels	Maximum Amplitude	Period	Timing Resolution	Rise/Fall Time	Pattern Depth	Looping
HFS 9DG1 Data Time Generator	630 Mbps	4 per module, 36 maximum	$3 V_{p-p}$ *	1.6 ns to ~20 ms	1 ps	<250 ps	64 Kbits	None
HFS 9DG2 Data Time Generator	350 Mbps	4 per module, 36 maximum	$5.5 V_{p-p}$ *	2.8 ns to ~20 ms	1 ps	800 ps to ~6 ns	64 Kbits	None
DG2020A	200 Mbps	12/24/36	$9 V_{p-p}$ *	5 ns	100 ps	<4 ns @ $5 V_{p-p}$ when used with P3420 probe	64 Kbits	2000 step/256 block looping nest
DG2030	400 Mbps	4/8	$5 V_{p-p}$ (50 ohm)	2.5 ns	20 ps	800 ps to 8 ns	256 Kbits	4000 step/256 block looping nest
DG2040	1.1 Gbps	2**	$2.5 V_{p-p}$ (50 ohm)	0.909 ns	10 ps	<150 ps	256 Kbits	4000 step/256 block looping nest

* Open circuit

** Complementary

THE INTEGRATED TOOL SET FOR SUPERIOR MEASUREMENT AND ANALYSIS

The Tektronix signal sources family is a key component of the powerful Tektronix instrument ensemble for testing the most challenging mixed-signal design and manufacturing test applications.

This integrated tool set includes sampling oscilloscopes, digital storage oscilloscopes, digital phosphor oscilloscopes, logic analyzers, and a host of complementary connection devices.



Measure Circuit Board Impedance

The TDS8000 Digital Sampling Oscilloscope with the 80E04 TDR/Sampling module is an outstanding solution for circuit board impedance measurements, featuring a reflected rise time of <35 ps and 20 GHz bandwidth.



Characterize High-speed Digital Signals

Measuring high-speed signal characteristics requires tools with uncompromised performance. The TDS694C Digital Storage Oscilloscope (DSO) has 3 GHz bandwidth and 10 GS/s sample rate across four channels simultaneously to ensure the most accurate single-shot rise-time and timing measurements available. A special cross-triggering capability enables correlation with TLA logic analyzers.



Analyze Complex Signals

An affordable solution for many applications is the TDS3000 Series Digital Phosphor Oscilloscope (DPO). This highly portable oscilloscope delivers up to 500 MHz bandwidth and sample rates up to 5 GS/s. Its intensity-graded color display helps you locate and characterize anomalies that are often elusive on traditional Digital Storage Oscilloscopes.



Examine Digital Systems in Real-Time

For design teams who need to debug and verify their product designs, Tektronix Logic Analyzers provide breakthrough features that capture, analyze, and display the real-time behavior of digital systems operation. The TLA700 and TLA600 Series feature powerful state-based triggering plus the ability to simultaneously acquire 2 GHz timing data and 200 MHz state data through the same probe. The TLA family operates with an open Windows® platform, allowing easy integration into the design environment.



Secure the Critical Connection

Even the most advanced instrument can only be as precise as the data that goes into it. Tektronix is the leader in technologies that enable access to the device-under-test. Our broad selection of probes and interconnect accessories includes the world's most advanced high-bandwidth active and differential probes which provide access to dense, high-speed circuitry while maintaining maximum signal fidelity.

For further information, contact Tektronix:



Worldwide Web: for the most up-to-date product information visit our web site at: www.tektronix.com/measurement/signal_sources/

ASEAN Countries (65) 356-3900; Australia & New Zealand 61 (2) 9888-0100; Austria, Central Eastern Europe, Greece, Turkey, Malta, & Cyprus +43 2236 8092 0; Belgium +32 (2) 715 89 70; Brazil and South America 55 (11) 3741-8360; Canada 1 (800) 661-5625; Denmark +45 (44) 850 700; Finland +358 (9) 4783 400; France & North Africa +33 1 69 86 81 81; Germany + 49 (221) 94 77 400; Hong Kong (852) 2585-6688; India (91) 80-2275577; Italy +39 (2) 25086 501; Japan (Sony/Tektronix Corporation) 81 (3) 3448-3111; Mexico, Central America, & Caribbean 52 (5) 666-6333; The Netherlands +31 23 56 95555; Norway +47 22 07 07 00; People's Republic of China 86 (10) 6235 1230; Republic of Korea 82 (2) 528-5299; South Africa (27 11)651-5222; Spain & Portugal +34 91 372 6000; Sweden +46 8 477 65 00; Switzerland +41 (41) 729 36 40; Taiwan 886 (2) 2722-9622; United Kingdom & Eire +44 (0)1344 392000; USA 1 (800) 426-2200.



From other areas, contact: Tektronix, Inc. Export Sales, P.O. Box 500, M/S 50-255, Beaverton, Oregon 97077-0001, USA 1 (503) 627-6877.

Copyright © 2000, Tektronix, Inc. All rights reserved. Tektronix products are covered by U.S. and foreign patents, issued and pending. Information in this publication supersedes that in all previously published material. Specification and price change privileges reserved. TEKTRONIX and TEK are registered trademarks of Tektronix, Inc. All other trade names referenced are the service marks, trademarks or registered trademarks of their respective companies.

0300 TD/XBS 76W-13740-0

Tektronix