Faster Troubleshooting with the TDS2000 Series Oscilloscopes

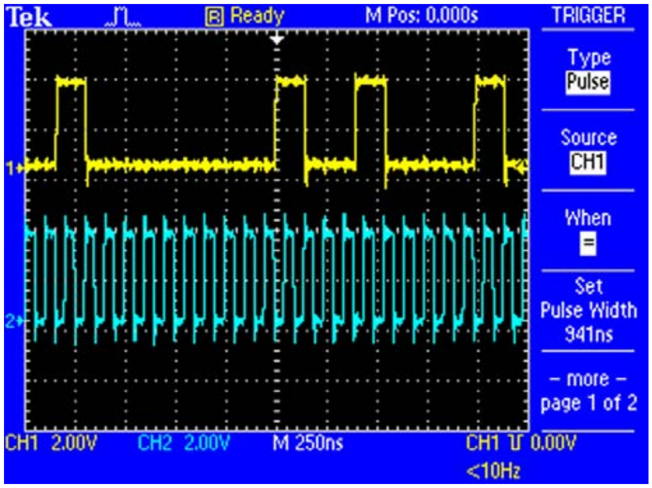

Introduction

While the fast, advanced processors in today's servers and PCs get most of the attention, humble 20- and 30-MHz processors are doing just as much work in the real world—if not more. Workhorse microprocessors, decades old in their design, are still earning sockets in machines, consumer electronics, and vehicles everywhere.

What do these embedded processors, and the applications they serve, have in common? Several characteristics stand out. The processors are well proven, well understood, widely supported, and easy to design with. Clock rates are relatively slow, and bus rates slower yet, compared to current high-end alternatives. The systems in which they are used—ranging from soft-drink machines to avionics--must have rock-solid reliability. And costs, including design, manufacturing, and maintenance expenses, must be held as low as possible.

One more characteristic is worth noting: there is a steady trend toward increasing the clock rates of these embedded components and buses. Not to catch up with the fastest servers, but toward "short-cycle" devices with clock rates five or six times faster than that of their predecessors. These devices have the same pinout and the same functionality, combined with the ability to do more work in a given span of time. This enables them to execute more instruction cycles and carry out more complex operations without slowing down overall system operation. For the software developer in particular, this is a great benefit. Time-consuming code optimization is less critical. New products can be brought to market more quickly and at lower cost.

All of these technical realities and market pressures result in the need for cost-effective, high-performance measurement instruments to support engineers' design and debug work. One class of tools—digital storage oscilloscopes (DSO)—continues to evolve in response to changing market requirements, such as the near-universal demand to embed digital subsystems in products of every kind, and the increasing "speeds and feeds" of the processors at the heart of these products.

Tektronix®
Enabling Innovation

The bandwidth of the low-cost oscilloscopes that support basic digital troubleshooting has doubled recently, reaching 200 MHz. Other valuable "high-end" measurement features, including advanced triggering, FFT analysis, and color displays, have migrated into entry-level instruments. Now designers have an affordable digital troubleshooting solution, such as the TDS2000 Series, for their work on commodity embedded processors.

The balance of this application note will focus on the application of some digital troubleshooting aids that are appearing in the emerging generation of low-cost DSOs.

MORE BANDWIDTH MEANS MORE VISIBILITY

A hidden "feature" of a processor manufactured a month ago, as opposed to the same device manufactured a decade ago, is the increase in the speed of its signal edges (transitions). Based on comparing the CMOS processes used 15 years ago, when many of these designs were created, to 5 years ago, when the fastest 5V capable processes were developed, speeds increased about three fold. Most new designs are either on these fastest 5V processes, or go to lengths to reduce voltages in the core areas, using 5V only on the periphery. With this later approach, even faster clock rates are possible. These faster speeds are a by-product of shrinking feature sizes on the silicon chip.

Faster edge performance is generally a good thing, reducing skew problems, setup times, and race conditions within the system. Faster propagation delays (driven largely by faster edge rates in CMOS) also create issues. Margins in address decoding that often depend on delays between address line logic and the bus control lines become more challenged when those delays become uniformly less. Regardless, designers need to see and understand these edges, as well as the narrower transients and other pulse characteristics that can arise in the presence of high-speed transitions.

When selecting a DSO for digital design work including a 20 MHz embedded processors, one might assume that an instrument with, say, 50 MHz or 100 MHz bandwidth is more than enough for the job. Certainly that is true for answering basic troubleshooting questions such as "Is the signal present? Is it correctly timed and synchronized?" But other details may be less obvious.

A DSO with high bandwidth provides far more insight into signal behavior than a lower-bandwidth instrument. That is simply because the oscilloscope's rise time becomes part of the equation that determines the quality of the viewed signal, as described by this formula:

$$\text{Measured Rise Time} = \sqrt{(\text{Oscilloscope Rise Time})^2 + (\text{Signal Rise Time})^2}$$

A pulse that appears to be "correct" when viewed at a lower bandwidth may have an amplitude aberration in its leading edge, causing it to act like two pulses. Or a narrow transient on a bus output might escape notice altogether, causing erratic operation on subsequent device inputs. As the formula implies, a 200 MHz DSO, such as the TDS2000 Series, can capture details that are invisible to a 100 MHz instrument.

The benefits of greater DSO bandwidth are not limited to viewing signal edges. Ground bounce, noise, crosstalk, and many other aberrations can be observed more readily—and overlooked less often—when a high-bandwidth instrument is used.

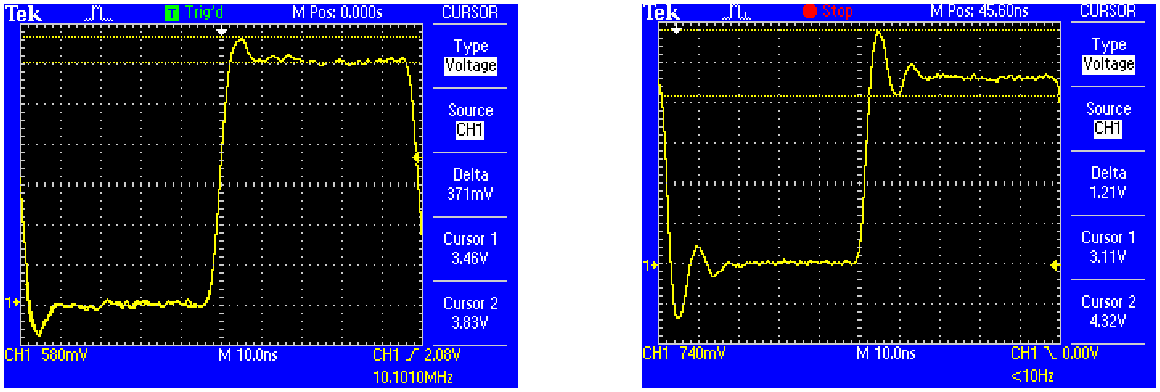

The higher the bandwidth, the more accurate the reproduction of your signal, as illustrated with a signal captured by the TDS2000 Series at 60 MHz and 200 MHz bandwidth levels.

CONDITIONAL TRIGGERING PINPOINTS TIMING PROBLEMS

An important but sometimes misunderstood laborsaving tool in every DSO is its selection of triggering conditions. When a DSO triggers, it is proof that the conditions you specified have been met. This simple fact, as much as the resulting waveform display, makes conditional triggering an essential tool for embedded system troubleshooting. Many people already use noise reject (which often increases trigger hysteresis) to limit runt pulses, and the various bandwidth limits to select out desired signals.

One of the most versatile triggering features, the pulse width trigger, has recently migrated from high-end lab instruments into low-cost DSOs, such as the TDS2000 Series. With this setting, the oscilloscope triggers when the incoming signal's pulse width is:

< (LESS THAN) the specified time

> (GREATER THAN) the specified time

= (EQUAL TO) the specified time (within a nominal tolerance)

≠ (NOT EQUAL TO) the specified time (within a nominal tolerance)

The LESS THAN pulse width trigger is one of the fastest ways to find suspected transient pulses on a bus or device output. Brief full-amplitude transients caused by crosstalk or timing issues can cause intermittent problems when they momentarily strobe a device's Output Enable or Chip Select input, for example. This may cause the device to send its data to the bus "out of turn," with unpredictable results. The LESS THAN trigger detects pulses shorter than a user-specified width and causes the oscilloscope to acquire the signals present at all connected probe inputs. By this means, it is possible to capture not only the transient itself, but also (if four oscilloscope channels are used) any resulting activity on the Output Enable line, the data bus, and so on.

Tektronix®
Enabling Innovation

The GREATER THAN trigger can help find "stuck bits" and other signals that fail to return to their default state after some transaction. This setting causes the oscilloscope to trigger itself if a trailing edge fails to occur within a specified time. For example, a bus data output signal may switch to a "1" in response to an Output Enable active, and not transition to a new state thereafter. This could be caused by the Output Enable signal itself being incorrect; by the transition time to tri-state being too long on the part being driven; or by the next value to be put on the data bus not being activated. The GREATER THAN trigger detects this error, again revealing contributing circumstances on all connected oscilloscope channels. With some probing, it should be possible to detect which of the above conditions actually created the difficulty. The time range here, as with the other pulse width trigger settings, extends from the tens of nanoseconds to multiple seconds. This provides ample time to ensure that the signal is truly "stuck" and not just delayed.

The EQUAL TO trigger offers an alternative to voltage threshold triggering approaches when the available trigger signals (such as Output Enables) are plagued by transients or noise, causing false triggering in the oscilloscope.

The basic embedded microprocessor illustrates this situation. Most of these devices include an external bus that allows for expansion of the processor's self-contained memory or peripheral interfaces. Often this bus allows peripherals to control the timing of their own transactions with the processor. The processor issues an address, and then sends out an address strobe. The affected peripheral eventually issues an "acknowledge" signal to confirm receipt of the processor's command. This process takes a known (and often user definable) number of delay cycles often unique to the specific peripheral or set of peripherals.

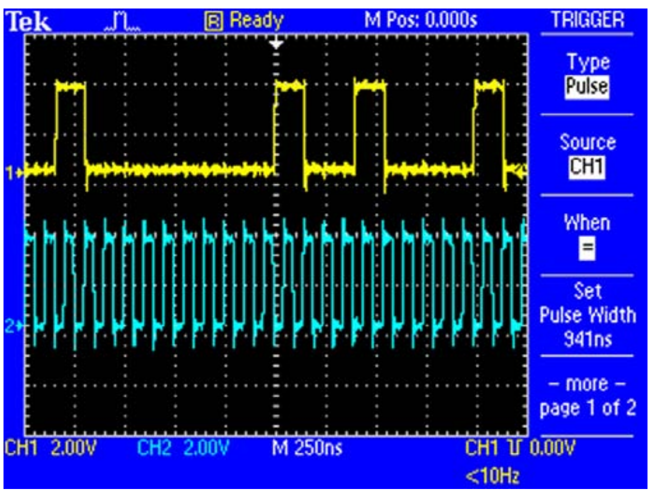

Setting the pulse width to = 941 ns enables you to capture signals with widths equal to 941 ns wide on channel 1 using the TDS2000 Series.

The knowledge of this delay time is the key to isolating any peripheral and examining the activity on various test points when it responds. The solution is simple: use the Address Strobe as the trigger line and set the pulse width trigger time EQUAL TO the number of clock delays associated with the specific peripheral. The leading edge of the Address Strobe will initiate the time countdown; then the trigger circuit will wait for the preset amount of time (say, three cycles' worth of time expressed in nanoseconds); then the oscilloscope will trigger and capture the activity on the test points. By definition, this is when the peripheral will be active on the bus.

Tektronix®
Enabling Innovation

Thus the EQUAL TO pulse width trigger allows the oscilloscope do a logic analyzer's job to some extent.

BUILT-IN COUNTER TAKES FREQUENCY MEASUREMENTS TO THE NEXT LEVEL

Automated frequency measurements have been part of the DSO feature set almost from the beginning. Normally these derive the period by examining the first full cycle of the captured waveform. This is a powerful tool to measure the characteristics of one-time events, but lacks the ability to extract a high accuracy average frequency value for a continuous waveform.

A different approach to frequency measurements is embodied in the common frequency counter, generally available and often fairly inexpensive (though never free). This approach can be implemented in a DSO by using the trigger signal as a source for average frequency readout, a new feature now finding its way into some low-cost DSOs. Frequency counters use many different means to make their measurements. The most common, simple solutions are the fixed frequency counter, where the number of input cycles are counted over a fixed interval of time (displaying the count) and the fixed period counter where the number of clock cycles are counted during one period (displaying the reciprocal of the count). Both solutions offer excellent accuracy when the number of accumulated counts is high, and poor accuracy when low. A related solution divides the measurement interval in half. For the first half, it counts both time base and stimulus counts. Once the half way point is reached, the next terminating transition (same polarity as the starting transition) on either input completes the measurement. The result is then calculated. This solution does not offer the maximum possible accuracy at the extreme cases, but consistently provides approximately ½ the maximum possible accuracy, providing a stable, at-a-glance indication (to six-digit precision) of the frequency of valid trigger events. Since virtually any event (within reasonable amplitude limits) can be a trigger event, the "readout" is in effect actually a useful general-purpose frequency counter.

In embedded system troubleshooting, it is often necessary to check the frequency of various local clock signals, including the master crystal oscillator. Here, the oscilloscope's trigger counter offers a quick built-in solution for the problem, as in the TDS2000 Series. The measurement is more accurate than that of the waveform-based automated frequency measurement, and it eliminates the need to set up a separate instrument for frequency counting.

The counter can also help in the search for crosstalk and noise sources. For example a noise signal on a bus trace, identified by the counter as having frequency of 100 kHz, might point toward crosstalk or grounding problems in the area of the switching power supply. Similarly, a noise signal with a frequency ½ that of the master clock might indicate crosstalk from an adjacent bus.

Because it is in fact derived from the trigger signal, the counter can detect the frequency of any conditional trigger event, not just the once-per-cycle voltage triggers. Combined with the pulse width trigger, for instance, the counter can determine the frequency with which a specific pulse width occurred within a varying stream of pulses.

WAVEFORMS IN LIVING COLOR

Formerly an exclusive feature in costly lab instruments, color liquid crystal displays are now a reality for some entry-level DSOs, such as the TDS2000 Series. Color brings an added dimension of information to the display, and makes troubleshooting with this class of instruments easier than ever.

A waveform is just a trace on a screen. How can it benefit by being presented in color? The value of color becomes clear when viewing multiple traces. Each of four traces has its own unique color. This color-coding scheme continues on the front panel, where the yellow-coded knobs control the yellow waveform, which is coming in via the yellow probe connector. The scheme can be extended all the way out to the probe tip or even to the device-under-test, marking test points with colored tags. The color oscilloscope display is a productivity solution that is foolproof in the best sense of the word—it minimizes the small human errors that can cost hours of troubleshooting time.

Color is also helpful when superimposing two traces to compare their differences; in addition, certain colors may be more legible in compromised lighting conditions.

CONCLUSION

Low-cost DSOs, such as the TDS2000 Series, bring important new capabilities to bear against the challenges of troubleshooting embedded digital systems. With increased bandwidth, significant enhancements in triggering, and other new features, these instruments are keeping pace with the market's demand for tools that will get the job done quickly and easily.