A Jitter Education: First-Year Lab

by

Gene L. Harding
Purdue University
glhardin@purdue.edu

Abstract: Timing jitter is a significant issue in high-speed digital design. Systems that interface with standard buses, like USB, IEEE1394, SPI4.2, Fibre Channel, and PCI-Express must meet the jitter specification required by each standard. Engineers and engineering technologists performing the design work need to understand what jitter is, how to characterize it, and how to find and mitigate its root causes so that their systems work properly.

This is the fifth in a series of papers addressing "A Jitter Education" (specifically, timing jitter in digital systems) in the context of the Electrical Engineering Technology (EET) curriculum. The first paper presented a number of jitter topics and proposed how to weave them into a four-year EET curriculum using a spiral approach. Subsequent papers explained how to introduce the subject of timing jitter to a first-year EET student, assessed those results, and discussed a more in-depth treatment appropriate for second-year EET students.

This paper proposes a lab experiment to follow and reinforce the first-year lecture material. In the first two procedures of the lab, students use a low-end digital storage oscilloscope (DSO) to investigate some of the display options of a DSO, including vector format, dot format, and persistence. In the last three procedures of the lab, they investigate the basic types of deterministic jitter and make measurements using a DSO that is not equipped to perform automatic jitter characterization. Alternative lab procedures are discussed, so instructors can modify the experiments, if needed. Options include a comparative analysis of jitter in circuits using different transistor technologies, a similar analysis between old and new function generators, and measurement of jitter in a low-cost fixture both with and without an electromagnetically-coupled interfering signal.

Keywords: Jitter, timing jitter, laboratories, education, educational technology.

I. Introduction

For the purposes of this paper, timing jitter will be defined as the deviation of a digital signal's transition from its expected time. If the deviation is large enough to move the transition into an adjacent clock cycle, the result is an error on the bus. Electrical engineers (EEs) and electrical

engineering technologists (EETs) who work with digital systems need to have a basic understanding of jitter for two reasons. First, it enhances troubleshooting to effect proper system performance. Second, since virtually all of today's high-speed digital bus standards include a jitter specification [1, 2, 3, 4, 5], it allows engineers to verify compliance with appropriate specs.

There are two broad categories of timing jitter: random (RJ) and deterministic (DJ). DJ is generally caused by issues over which the system designer has some control. Periodic jitter (PJ) is caused by electromagnetic interference (EMI), intersymbol interference (ISI) is caused by limited bandwidth in the transmitting medium, and duty cycle distortion (DCD) is normally caused by an incorrect threshold voltage setting. A more in-depth treatment of the types of jitter and their underlying causes is contained in [6].

There are three basic types of jitter measurements. The most basic is period jitter, which is the difference between the minimum and maximum periods of a waveform. Note the distinction between periodic jitter, which is a type of *jitter*, and period jitter, which is a jitter *measurement*. The abbreviation PJ, at least in this article, is used to denote *periodic jitter*. Cycle-to-cycle jitter is period jitter measured between two adjacent clock cycles of a waveform; a variation of this measurement is called *n*-cycle jitter, which is the same except that the periods measured are *n* cycles apart. The most complex measurement is time interval error (TIE). TIE requires reconstruction of the signal's ideal clock, and is useful for characterizing jitter's total effect over time [6].

Jitter measurements can be displayed in a number of different ways, each with its own advantages and disadvantages [6]. The histogram is a probability density function that illustrates the likelihood of occurrence for each jitter magnitude. The trend display maps jitter magnitude against "wall clock time" and is useful for seeing trends in the jitter, such as with spread spectrum clocking, and for correlating periodic jitter with the signal inducing it. The spectrum display, since it plots jitter against frequency, can also be helpful for determining sources of EMI. Perhaps the most intuitive display is the eye diagram. It plots repetitive waveform captures in an overlay fashion, showing both time and voltage information very clearly. A less common display is the bathtub plot, which shows the size of the data valid window for different bit error ratios.

A number of modern test instruments can measure timing jitter [6]. These include real-time DSOs, sampling oscilloscopes, logic analyzers, spectrum analyzers, bit error ratio testers (BERTs), and time interval analyzers (TIAs). The real-time DSO is by far the most common of these, and it is the instrument of choice for the lab described in this paper.

Although in-depth jitter analysis is a complex subject suitable for upperclassmen or graduate students, jitter concepts are very straightforward and well within the grasp of first- or second-year students. As such, the subject lends itself well to a spiral educational approach, revisiting jitter topics at increasing levels of complexity throughout an EE or EET curriculum. This is the approach described in [6], and the lab described herein is a part of it.

The remainder of this paper discusses the introductory jitter topics along with the corresponding prerequisite knowledge, a detailed description of the lab in its updated form, results of the first trial of the lab in its initial form, and some alternative approaches for the jitter measurements.

For the reader's convenience, Table 1 lists definitions for the acronyms used in this article.

Table 1: Acronym List

Acronym	Definition
BERT	Bit error ratio (or rate) tester
DCD	Duty cycle distortion
DJ	Deterministic jitter
DSO	Digital storage oscilloscope
EE	Electrical engineer(ing)
EET	Electrical engineering technology(ist)
EMI	Electromagnetic interference
FFT	Fast Fourier Transform
ISI	Intersymbol interference
LPF	Low-pass filter
LSB	Least significant bit
MSB	Most significant bit
PCB	Printed circuit board
PJ	Periodic jitter
PLD	Programmable logic device
RJ	Random jitter
TIA	Time interval analyzer
TIE	Time interval error

II. The Freshman Material

The beginning material, which I present to my second-semester freshman EET students at Purdue, is detailed in [7]. It begins with the basic definition of jitter and the difference between RJ and DJ. The different types of DJ are introduced (periodic jitter, ISI, and DCD) and their underlying causes discussed. Next, the three basic measurements are presented: period jitter, cycle-to-cycle and *n*-cycle jitter, and time interval error. The last topic is the importance of the DSO, which is primarily its ability to capture large amounts of data for statistical analysis.

There is some foundational knowledge required before the above material is covered [7]. The student must have a basic understanding of the DSO, analog-to-digital and digital-to-analog conversion, and the concepts of sampling and sampling rate. At this point in the curriculum, we have covered electromagnetic fields and transformers, so the concept of EMI (and its effects) is a natural progression. Likewise, RC time constants are covered in the first semester analog course along with rise and fall times. Threshold voltages are used in both of the first two semesters, and various voltage specifications are explained in the second semester. One more specific topic is needed to perform this lab: synchronous counters. By the time students do the jitter lab,

synchronous and asynchronous counters have been discussed in lecture and used in the lab multiple times, so they should be well prepared to build the 74LS160-based synchronous counter they need to perform this lab.

The title of the lab is "DSO Displays and Timing Jitter" because it covers both subjects. There are some pre-lab tasks that should take an hour or so to perform, five in-lab activities, and several analysis questions. Students have 2:50 (two hours and fifty minutes) to complete the lab. I expect some students to complete both the lab activities and analysis questions during lab time, while others will have to do some or all of the analysis questions outside of the lab. The next section describes the pre-lab activities.

III. Pre-Lab

There are five steps in the pre-lab. The first is a calculation of the time between samples of the oscilloscope, which is simply the reciprocal of the sampling rate:

$$T_{samples} = \frac{1}{SampleRate} \tag{1}$$

The purpose of this step is to review the relationship between sample rate and sample period.

Question two has students calculate the actual sample rate for different sweep rates. The sweep rates used in this portion of the lab will be 2.5 ms/div, 250 µs/div, and 25 µs/div. At these sweep rates the DSO's memory depth, not the maximum sample rate, determines the sample period:

$$T_{2.5ms/div\_sample\_rate} = \frac{\left(\frac{2.5ms}{div}\right)(10div)}{2500pts} = 10\ \mu s \tag{2}$$

$$T_{250\mu s/div\_sample\_rate} = \frac{\left(\frac{250\mu s}{div}\right)(10div)}{2500pts} = 1\ \mu s \tag{3}$$

$$T_{25\mu s/div\_sample\_rate} = \frac{\left(\frac{25\mu s}{div}\right)(10div)}{2500pts} = 100\ ns \tag{4}$$

This idea will be important in the second part of the lab, dealing with infinite persistence, when students need to determine why more samples appear on the signal transitions at faster sweep rates.

The third pre-lab question asks the students to predict whether the counter output edge rate will get faster, slower, or stay the same when the clock input increases. The purpose of this question is to get them to realize that the output edge speed is dependent on the *technology* used to make the chip, *not* the speed of the clock input.

The fourth step is construction of the counter circuit. Theoretically, this step should require a few minutes to look up the datasheet, a few minutes to determine how to wire the circuit, and a few more minutes to build the circuit. These steps can all be done outside of the lab, however, to prevent wasting valuable lab time.

For the final pre-lab step, students must draw the expected output waveforms of the most and least significant bits (MSB and LSB, respectively) of the four-bit counter, as shown in Figure 1. The reason for this step is to get students to think about what the output will be before lab to save more in-lab time. During lab, they will use their drawings to verify the correctness of the counter's outputs. These two waveforms will be used to illustrate intersymbol interference. The in-lab activities and associated analysis questions are discussed next.

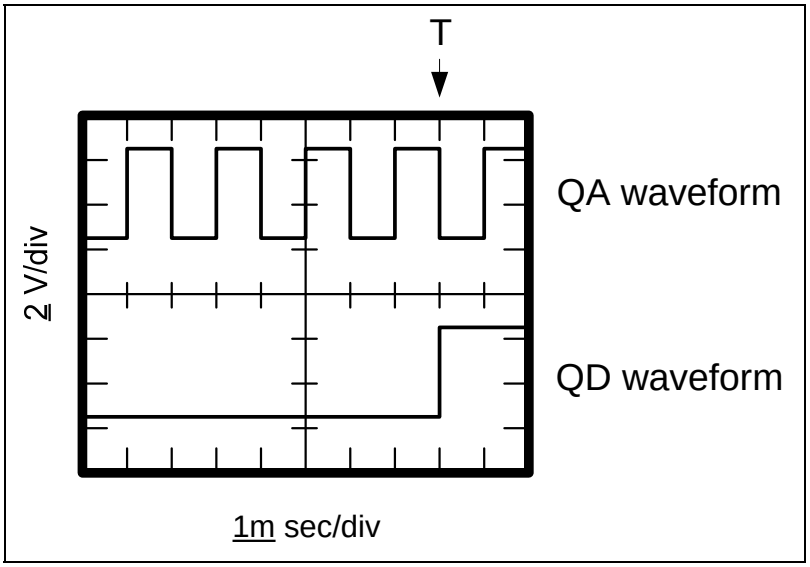

Figure 1: Counter LSB (top) and MSB (bottom) expected waveforms

IV. In-Lab Activities and Analysis Questions

The lab itself is divided into five tasks: two exercises to investigate DSO display options, then measurements of three different types of jitter.

The first lab exercise compares vector and dot display options in the DSO. A function generator is used to create a 5-V_{pp} square wave at frequencies of 100 Hz, 1 kHz, and 10 kHz. Using vector display format, these waveforms are viewed on the DSO at sweep rates of 2.5 ms/div, 250 µs/div, and 25 µs/div, respectively. Questions embedded within the lab steps prompt students to observe that the waveforms appear identical; the only noticeable difference is the time base of the oscilloscope. A screen capture of the 100-Hz square wave is shown in Figure 2.

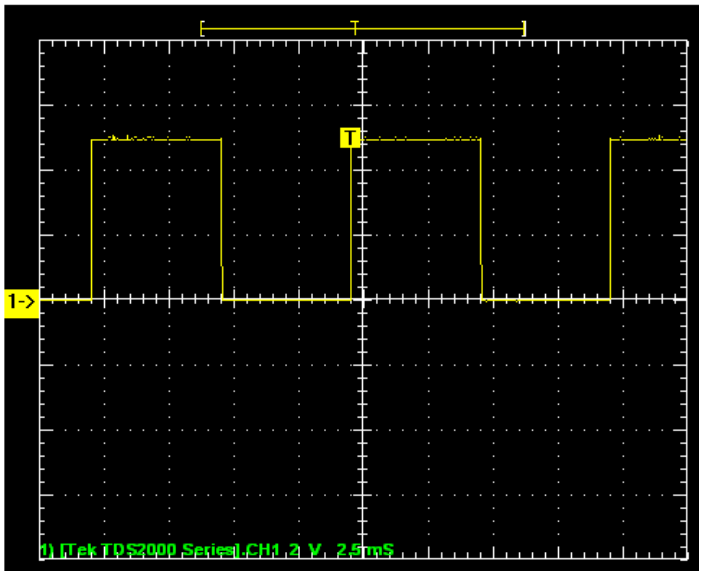

Figure 2: 100-Hz square wave, 50% duty cycle, 0V to 5V; vector display mode

These steps are then repeated with the oscilloscope in dot display mode. A difference is immediately apparent: the waveform transitions "disappear" because the dots are no longer connected. The display looks very much like the display of an old analog oscilloscope. Unfortunately, the screen capture software for our oscilloscopes "connects the dots," so I cannot include a representative figure to illustrate it, but it looks the same as Figure 2 without the vertical portions of the waveform.

Three of the analysis questions focus on this portion of the lab. The first question is

> "What is the difference between the waveforms, … and *why* are they different?"

The goal of this question is to have students observe that the waveform transitions look different because the sampling points (dots) are connected in vector mode.

The second question is

> "Why do the horizontal and vertical portions of the waveforms look the same in 'vector' display format, but different in 'dot' display format?"

The idea here is that vector display mode masks sampling density.

Question three is:

> "In step 1.g., why did the signal transitions disappear while the horizontal portions of the waveform were still clear?"

The horizontal portions of the waveforms are clear because many samples are taken in those areas, creating the appearance of a line even though the "dots" are not connected. Few, if any, samples are taken on the edges because of the fast edge rates, making the transitions "disappear."

The overall idea of the first exercise is that, although "vector" display mode provides a pleasing and intuitive display, some care must be taken when using it. Its output can sometimes be deceptive, especially when dealing with signal speeds near the sampling rate limits.

The second exercise explores display persistence. The DSO remains in "dot" display mode and is set for infinite persistence. Students take a 10-second capture of the same square waves used in part one. The noticeable difference in the waveforms is the number of samples on the waveform transitions. With a sampling period of 10 µs for the 100-Hz waveform, few if any samples occur at the edges, even after 10 seconds' worth of captures. For the 1-kHz waveform the sampling period is 1 µs, and a few samples appear at the transitions within 10 seconds. Increasing the waveform frequency and sweep rate by another factor of ten decreases the sample period to 100 ns, and the waveform edges become clearly visible.

There is one analysis question for this portion of the lab:

> "Why were there more dots at the signal transitions of the waveform in step 2.e. than in the waveforms of steps 2.c. and 2.d.? (Hint: Review pre-lab question 2.)"

The key takeaways here are: 1) sampling rate is limited by memory depth at all but the fastest sweep rates; 2) faster sweep rates mean faster sampling rates; and 3) a faster sampling rate means more samples at each edge.

The remaining parts of the lab involve measuring different types of jitter. In the third exercise, students measure jitter caused by simulated EMI, then by simulated noise. The oscilloscope remains in "dot" display mode with infinite persistence, the sweep rate is set to 100 µs/div, and the trigger point is moved near to the left side of the screen. The function generator is set to provide a 1-kHz square wave with a sinusoidal modulation (i.e., simulated EMI) of up to 200 µs. This places the rising edge of the waveform near the left side of the display and the unmodulated falling edge just past center screen. Students take 30 seconds worth of captures and use the DSO's cursors to measure the peak-to-peak jitter, which is at or very near 400 µs. The process is repeated using noise modulation, which results in about 25% less peak-to-peak jitter.

The primary purposes of this exercise are to give students their first experience of seeing what jitter looks like and to reinforce the correlation between noise and RJ. There are three very straightforward analysis questions (the second one is actually two questions in one):

> "Which type of modulation resulted in a larger magnitude of jitter?
>
> □Sine □ Noise"
>
> "Which type of modulation would correlate to deterministic jitter (DJ), and which type to random jitter (RJ)?
>
> Sine → DJ RJ
> Noise → DJ RJ"

The answers are Sine, DJ, and RJ, respectively.

The fourth and fifth lab procedures make use of the four-bit decade counter circuit constructed in pre-lab. The circuit, shown in Figure 3, includes low-pass filters (LPFs) built from discrete components attached to the counter's QA (LSB) and QD (MSB) outputs. QA is a square wave; QD is an asymmetric waveform, low for counts 0-7, and high for 8-9. As such, QD has a relatively long time to discharge to the zero state before its rising edge, which should correspond to QA's falling edge, as shown in Figure 1. The asymmetric nature of the MSB waveform is the property that can cause intersymbol interference (ISI) in bandwidth-limited media.

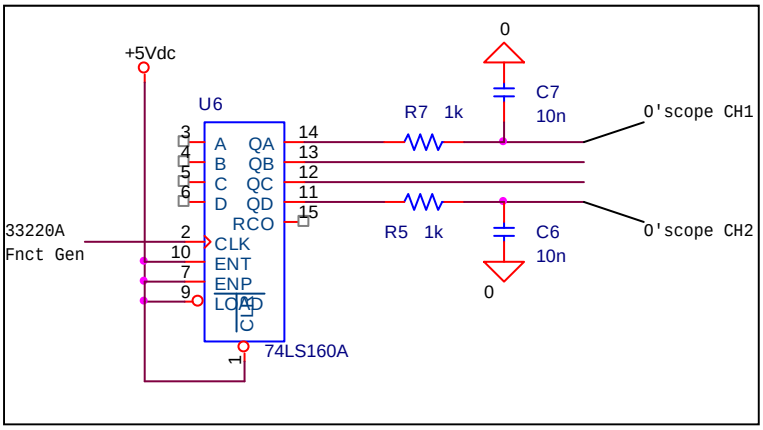

Figure 3: 74LS160 synchronous 4-bit decade counter circuit

To make the ISI measurement, students first align the vertical midpoints of the two waveforms with one of the display graticules, then use the cursors to measure the time difference (at the waveforms' mid-point) between the falling edge of QA and the rising edge of QD. This measurement is done at 1 kHz, 50 kHz, and 100 kHz. The jitter magnitude increases with the frequency. After the measurement at 100 kHz, the sweep rate is slowed down somewhat to show two full periods of both waveforms, and students are required to draw them. At this frequency, the bandwidth limitation is very obvious. Neither signal ever reaches the full voltage value of the high state, and QA also never fully reaches the low voltage value. This is illustrated in Figure 4.

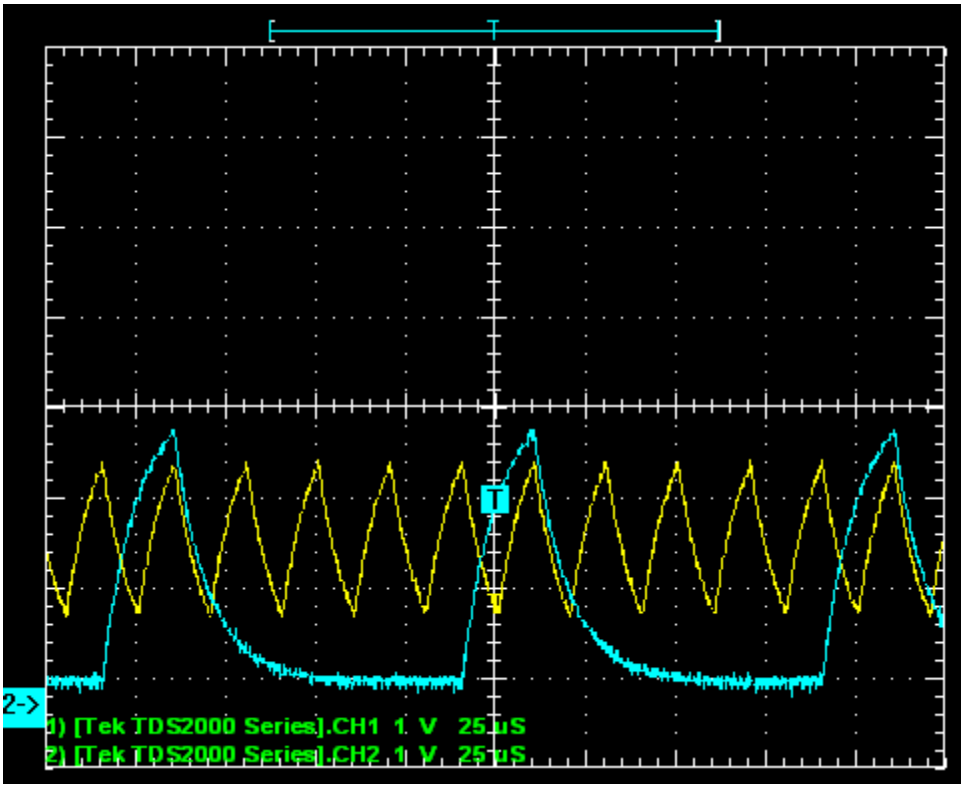

Figure 4: QA and QD waveforms at 100 kHz

There are two analysis questions relating to ISI. The first one asks for an explanation of the answer to pre-lab question number three. Hopefully, students will have figured out the answer correctly in pre-lab, and just need to explain in writing that the 74LS160's output edge rates are based on the technology used to create the chip, not the input clock rate. The second ISI question is (note that Table 1 in the *lab instructions* contains the jitter measurements at 1 kHz, 50 kHz, and 100 kHz):

> "If the ISI jitter measurements in Table 1 were exactly the same, how would their proportion of the clock period change as the frequency increased?
> □ Increase □ Stay the same □ Decrease"

The purpose of this question is to drive home the point that there are two things going on at the same time. As the frequency increases, the clock period decreases, which magnifies the impact of the increasing jitter.

The final in-lab procedure is to measure the jitter caused by an incorrect threshold voltage setting. All that is needed is a signal with an edge rate that is measurable with the DSO. Since the counter is already running and connected to the DSO, the rising edge of the low-pass-filtered QD signal is used. First, the clock rate is reduced back to 1 kHz to ensure complete high and low voltage swings. Next, the waveform's vertical midpoint is aligned with one of the display graticule lines, and cursors are used to measure the jitter that would be caused by a 0.5-V (10%) error in threshold voltage. This is illustrated in Figure 5.

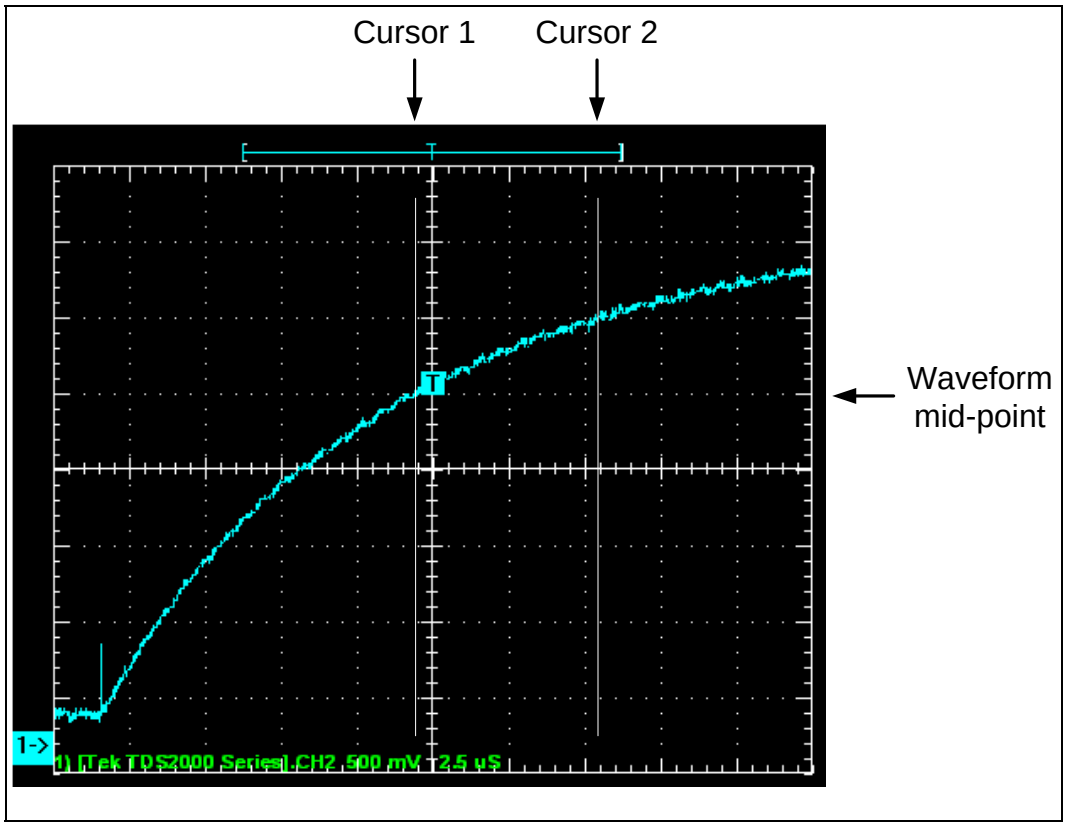

Figure 5: Cursor measurement of jitter caused by 0.5-V threshold voltage error

The last analysis question addresses duty cycle distortion:

> "Given your DCD jitter measurement, how much DCD jitter would be caused by a 250-mV error in the voltage threshold setting? (You may assume a linear waveform transition.) Show your calculation below. "

The point of this question is that the effect of various threshold settings can be approximated using linear interpolation (and to have students apply their algebra skills).

V. The First Trial

I administered an early version of this lab for the first time during the spring 2008 semester. It only had a single pre-lab question (the first question listed in the previous section), no ISI or DCD measurements, and half of the analysis questions. Although the average grade (95%) implies great success, my observation of the actual lab indicated otherwise. The students were not really engaged in the lab, their comprehension of the concepts seemed poor, and the lab was too short. (Even the slow students finished in a little over an hour.)

Part of the problem is certainly that this was the least motivated cohort of students I have seen to date, but there were other issues. First, the pre-lab activity was not nearly enough to get them

thinking about the lab beforehand. The additional pre-lab tasks are intended to remedy that issue. Second, the lab was clearly shorter and easier than appropriate, which indicated plenty of room for adding the ISI and DCD measurements, plus the additional analysis questions. Hopefully, these changes will also motivate the class to engage better with the lab and improve their overall comprehension of the material. The lab as outlined in this paper is planned to be administered in the spring 2009 semester.

VI. Alternative Lab Techniques

This section presents a few alternatives for investigating periodic jitter, intersymbol interference, and duty cycle distortion.

If one could easily produce printed circuit boards (PCBs), or had appropriate PCBs already available, it would be interesting to explore periodic jitter using a PCB with two traces running parallel and close together for several inches. The signal of interest could be sent down one trace, while an interfering signal of a specific frequency, not a harmonic of the signal of interest's frequency, could be transmitted down the other trace. The DSO's Fast Fourier Transform (FFT) function would be applied to the signal of interest, and a spike at the interfering signal's frequency could be observed on the FFT display. Another possibility that would obviate the need for a special PCB would be to operate a relatively high-powered transmitter (a cell phone, perhaps?) close to a regular proto-board circuit. This approach might work if the DSO had enough bandwidth to capture the cell phone's frequency with its FFT. Either way, observing the FFT spike at the interfering signal's frequency would likely be an eye-opening experience for most students, even if the magnitude of the periodic jitter were not measured directly in the time domain. On the other hand, this technique might be a little overwhelming for our freshmen, because they are not introduced to the frequency domain and Fourier series until their sophomore year.

Although it could be significantly more involved, it might be instructive to compare the ISI of two longer data streams. If the data streams were generated from a Programmable Logic Device (PLD), microprocessor, or even a programmable function generator, they could be created with a variety of patterns (1s and 0s) and the jitter compared.

Finally, it could be useful to compare DCD produced by different technologies. For instance, since their edge speeds are different, a 74LS-series chip could be compared to its faster equivalent in the 74AS-series. Likewise, a similar comparison could be done between old and new function generators if the DSO has adequate bandwidth to properly characterize the signal transitions of the new function generator.

VII. Conclusion

Timing jitter can be a major issue in some high-speed digital systems. It is important for EEs and EETs who work with such systems to understand jitter for two reasons: so they can effectively

troubleshoot their systems and mitigate its underlying causes, and so they can verify compliance with standard bus specifications.

Although detailed jitter analysis is a fairly complex subject, the concepts are very straightforward and intuitive. Thus, jitter analysis is a good topic to be taught with a "spiral" approach, introducing basic concepts early in the curriculum and revisiting the subject periodically with increasing levels of depth and complexity.

This article built on previous papers by describing a lab to introduce the three basic types of deterministic jitter—periodic jitter, intersymbol interference, and duty cycle distortion—while also delving into some of the display features of the digital storage oscilloscope. Interested instructors can obtain a copy of the lab instructions, in Microsoft Word format, by e-mailing the author (see e-mail address at the top of the paper). Constructive feedback is also welcome.

References

[1] Agilent Technologies, Inc., "Automated USB 2.0 Receiver Compliance Test and Characterization with the Agilent N5990A Software Platform," *application note*, document #5989-6232EN, 2007.

[2] The Institute of Electrical and Electronics Engineers, Inc., "1394b™ IEEE Standard for a High-Performance Serial Bus—Amendment 2," *IEEE Std 1394b™-2002*, Dec 2002.

[3] Optical Internetworking Forum, "System Packet Interface Level 4 (SPI-4) Phase 2 Revision 1: OC-192 System Interface for Physical and Link Layer Devices," Implementation Agreement OIF-SPI4-02.10, Oct 2003.

[4] National Committee for Information Technology Standardization (NCITS), "Fibre Channel – Methodologies for Jitter Specification, draft technical report," *NCITS T11.2 / Project 1230 / Rev 10*, Jun 1999.

[5] Merritt, R., "Next PCI Express to Deliver 5Gbps Data Rate", *Electronic Engineering Times*, Jul 2006.

[6] Harding, G. L., "A Spiral Approach to Teaching Jitter Analysis in the Undergraduate Curriculum," *the Technology Interface*, Spring 2007, Vol. 7, No. 2, ISSN # 1523-9926, http://technologyinterface.nmsu.edu/Spring07/.

[7] Harding, G. L., "A Jitter Education: An Introduction to Jitter for the Freshman" [CD-ROM], *2006 Annual Conference Proceedings, American Society for Engineering Education*.

Gene L. Harding is an associate professor of Electrical and Computer Engineering Technology at Purdue University, where he has been for 7 years. He also worked in industry for 3 years with Agilent Technologies, and has over 23 years of combined active and reserve service with the United States Air Force.