


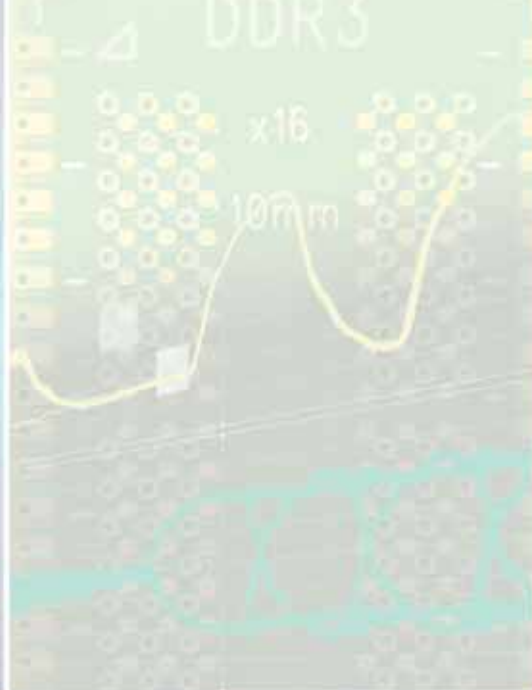


ISSUE FIVE • 2008

Agilent Measurement Journal



Agilent Technologies



Finding a New Path when Assumptions Break Down



Darlene J.S. Solomon
Chief Technology Officer, Agilent Technologies
Vice President, Agilent Laboratories

Every type of measurement is based on fundamental assumptions about natural laws and physical behavior. As we push the technological envelope, however, real physical phenomena often cause breakdowns in our assumptions. When this happens, we need new tools, techniques and measurements to characterize the unexpected behaviors.

In digital electronics, for example, the square waves that provide clock signals and represent data bits become distorted

at gigahertz frequencies and beyond. One explanation is rooted in Fourier theory, which states that any waveform can be represented by a combination of sine waves of specific frequencies, amplitudes and phase relationships. As high-frequency digital signals travel through physical media, delays to the constituent frequencies produce phase shifts and waveform distortion. These can cause unexpected or unreliable operation in digital circuits designed for sharply defined ones and zeroes.

For digital signals that exhibit such “analog-like” behavior, signal integrity measurements and jitter measurements are two ways to identify, estimate, quantify, and troubleshoot problems in digital systems. Tools such as the simulators within advanced design applications also help developers predict and minimize these problems earlier in the design cycle.

In chemical analysis, it’s easy to imagine today’s “microfluidic” instruments as containing miniaturized beakers, tubes and columns that handle volumes of less than one microliter. When channels and chambers are less than 100 micrometers wide, however, fluids may behave in unexpected ways — and a process as seemingly simple as the mixing of fluids

can become very difficult, or very easy. At the microscale, molecules diffuse quickly over small distances, and heating or cooling occurs very rapidly. If those molecules interact with the walls of a microfluidic chip, they tend to behave in ways not seen in the macroscopic world: Different rules take over and different physical properties dominate.

Researchers from Agilent and elsewhere have gone far beyond macroscale assumptions to understand fluid flow, diffusion, heat transfer, surface interactions, and fabrication at the micro- and nanoscale. From that knowledge, we’ve created instruments such as the high-pressure nanoflow liquid chromatograph/mass spectrometer (LC/MS). These advanced systems provide sensitive and repeatable quantitative or qualitative analyses of nanoliter and picoliter samples.

Flawed assumptions don’t have to be a dead end. As you push the envelope in your work, I encourage you to step back, check for “assumption breakdowns” — and look for new paths that lead to fresh, successful approaches. As we’ve learned at Agilent, sometimes it’s the fundamental stumbling block that leads to the most intriguing breakthrough.

Insight

Department

1 Finding a New Path when Assumptions Break Down

Physical phenomena often cause break-downs in our fundamental assumptions about natural laws and physical behavior, requiring new tools, techniques and measurements to characterize the unexpected behaviors.

Visit the *Journal's* online version to listen to Darlene's podcast interview at www.agilent.com/go/journal

Emerging Innovations

Department


- 4 • Testing MIPI D-PHY protocols
- Oscilloscope firmware update
- DNA replication
- Joining the LTE/SAE trial initiative
- Testing proteins
- Praise for DC power analyzer
- Scripting features for AMDS
- Testing hybrid PCBA systems
- Improving the gene expression system workflow

Solutions

Department

6 Turning a "Good Enough" Test Strategy into one that's Reliable, Repeatable and Traceable

Tweaking a "good enough" test strategy for a design under test yields one that is reliable, repeatable and traceable.



10 Understanding the Effects of Limited-Bandwidth Channels on Digital Data Signals

The effect of limited-bandwidth channels on digital data signals can be overcome through compensation for the bandwidth limitation by emphasis and equalization.

16 Introducing the 3GPP LTE Downlink


While LTE has the potential to enhance current deployments of 3GPP networks and enable significant new service opportunities, the new techniques used in the downlink add substantial complexity.

26 Validating the Physical and Protocol Layers in DDR Memory Interfaces

Addressing the challenges of validating DDR memory at the physical and protocol layers requires that special attention be paid to the probing methods of today's oscilloscopes and logic analyzers.

32 Understanding Total Jitter Measurements at Low Probabilities

A new bracketing approach coupled with a bit error rate tester can be utilized to make accurate peak-to-peak total jitter measurements.



44 Using Behavioral-Model Simulation to Accurately Predict First-Order PLL Performance

Increasing data rates for high-speed serial links and RF wireless frequencies are among the factors driving the need to characterize PLL circuit blocks and accurately assess their noise in a closed-loop condition.

www.agilent.com/go/journal

50 Creating Synchronous High-Frequency Sampling across Multiple Digitizers

Measuring sampling-clock delay using the acquired signal as a time reference offers a viable means of simultaneously measuring high-frequency signals over many channels.

56 Overcoming the Challenges of Testing FlexRay Networks

A protocol analyzer designed to handle issues specific to FlexRay networks helps create the quality metric developers need to obtain a clear indication of their work's quality.

Visit the *Journal's* online version to listen to a podcast concerning this article.

www.agilent.com/go/journal

Subscribe to *Agilent Measurement Journal*

Give yourself an edge in today's dynamic world of technology: Subscribe to the *Journal* and receive your own copy of the print edition in the mail.

To activate your free subscription, go to www.agilent.com/go/journal and look for the link "Manage your *Agilent Measurement Journal* subscription."

While visiting the site, check out the latest digital edition of *Agilent Measurement Journal*. Also, please help us improve and enhance future issues of the *Journal* by completing the survey located at the end of the digital edition.

64 Understanding Operating Life and Repeatability of Electromechanical Switches and their Effect on Total Cost of Ownership

When selecting an EM switch, careful consideration must be given to its operating life; repeatability; the switch design, mechanism and materials; and the process control used during mass production.

Campus Connection Department



70 Implementing Micro and Nano LC/MS Techniques for High Sensitivity Lipidomics Analysis

The University of California, Irvine performs leading-edge research into biologically active lipids — and is another excellent example of successful academic collaboration with Agilent.

Visit the *Journal's* online version to listen to a podcast concerning this article.

www.agilent.com/go/journal

AGILENT MEASUREMENT JOURNAL

William P. Sullivan | President and Chief Executive Officer

Darlene J.S. Solomon | Agilent Chief Technology Officer and Vice President, Agilent Laboratories

Heidi Garcia | Editor-in-Chief

ADVISORY BOARD

David Badtorff | San Diego, California, USA

Lee Barford | Santa Clara, California, USA

Bert Esser | Amstelveen, Netherlands

Pablo Estrada | Santa Rosa, California, USA

Mike Griffin | Hong Kong, Hong Kong

Johnnie Hancock | Colorado Springs Colorado, USA

Jean-Claude Krynicki | Palaiseau, Essonne, France

Linda Roettger | Santa Clara, California, USA

Craig Schmidt | Loveland, Colorado, USA

Roger Stancliff | Santa Rosa, California, USA

Kazuyuki Tamaru | Tokyo, Japan

Boon-Khim Tan | Penang, Malaysia

Giampaolo Tardioli | Santa Rosa, California, USA

Dominic Testa | Wilmington, Delaware, USA

Daniel Thomasson | Santa Rosa, California, USA

Kenn Wildnauer | Santa Rosa, California, USA

Jerry Yeoh | Penang, Malaysia

Martha Zemedé | Santa Rosa, California, USA

EDITORIAL

Please e-mail inquiries and requests to journal@agilent.com

Agilent Measurement Journal

Emerging Innovations

- **Industry's first MIPI D-PHY test solution facilitates debugging**

Agilent's new mobile industry processor interface (MIPI) D-PHY protocol test solution features the Agilent N4851A analysis probe and N4861A stimulus probe. This first-to-market system, based on the Agilent 16900 logic analyzer platform, supports camera serial interface (CSI-2) and display serial interface (DSI) and facilitates hardware and software debugging, reducing the need for interoperability testing. Real-time analysis and stimulus hardware provide bit-level-to-image-level test capabilities, hierarchical protocol display, real-time error detection, and automated tools for test-vector generation, allowing users to simulate, troubleshoot and verify designs.

- **Oscilloscope firmware upgrade runs MATLAB® scripts**

Researchers and designers can now run MATLAB scripts and directly calculate transmitter waveform dispersion penalty (TWDP) and other performance parameters without external processing courtesy of a firmware update for the Agilent 86100C Infiniium DCA-J wideband oscilloscope. A compliance methodology adopted by IEEE 802.3aq 10G Ethernet (LRM), TWDP quantifies transmitter performance for systems using dispersive channels and equalizing receivers. The updated firmware (revision 8.0) sends live oscilloscope waveforms to MATLAB where they can be analyzed directly. In addition, engineers can use any MATLAB script from any standard, or they can write their own scripts for custom outputs. The results are faster development cycles and greater insights toward design improvements.

- **DNA replication findings may explain disease origins**

Researchers at the Baylor College of Medicine have observed a new way that DNA additions or deletions are introduced in genes during cell division. The replications have been associated with a wide range of diseases such as Alzheimer's, Parkinson's, Potocki-Lupski Syndrome, and others.

The findings — aided by Agilent's custom oligonucleotide comparative genomic hybridization microarrays — were published in the journal *Cell* with details about how this newly discovered method, called replication fork stalling and template switching (FoSTes), adds or deletes segments of DNA in previously unexpected locations during replication. Baylor researchers studying Pelizaeus-Merzbacher disease (PMD) with the new FoSTes method found genomic changes, such as extra genetic material in the middle of another duplication, that previous DNA theories could not explain.

- **Scripting features highlight new AMDS offering**

The latest release in Agilent's Antenna Modeling Design System (AMDS) is now available. The new full-wave, 3-D electromagnetic modeling and simulation software contains a scripting feature that automates complex designs such as patch-array antennas. The enhancements enable designers to fine tune antennas for better performance within electronic devices such as cell phones. They also streamline and introduce greater accuracy into post-processing computations by allowing designers to write their own programs to automate element placement and incorporate mathematical functions that can perform antenna design analysis.

● High-sensitivity protein kit trumps silver-stained gel methods

Pharmaceutical and life-science researchers now have a more accurate way to analyze proteins with Agilent's new High Sensitivity Protein 250 Kit for the Agilent 2100 bioanalyzer. The kit delivers greater sensitivity than silver-stained SDS-polyacrylamide gel electrophoresis (SDS-PAGE) by detecting proteins to as little as 1 pg/ul and a 0.05 percent impurity. It also covers a sizing range from 10-250 kDa and boasts a quantitation range of four orders of magnitude. The new kit features a direct labeling reaction that is highly reproducible, and separation, quantitation and purity measurements can be performed in a single step with a 10-sample-per-hour throughput for QA/QC.

● And the award for best DC power analyzer goes to...

Agilent's N6705A DC power analyzer has taken home a bevy of awards from industry publications and organizations since its introduction in May 2007. *Test & Measurement World* named it 2008 Best in Test, while *Electronic Products* dubbed it 2007 Product of the Year. In addition, the power analyzer was honored with *Design News'* 2007 Golden Mousetrap Award and was a finalist for the International Engineering Consortium's 2008 Design Vision Award.

The N6705A measures a device under test's power consumption and does not require engineers to write any code for operation. It provides a variety of sourcing and measuring capabilities by combining up to four DC power supplies with digital multimeter (DMM), oscilloscope, arbitrary waveform generator, and data logger functionality.

● Hybrid PCBA test system reduces costs and resources

By combining two different electronics manufacturing test methodologies, Agilent has devised a new hybrid boundary scan and vectorless test extended performance (VTEP) printed circuit board assembly (PCBA) test solution. Based on Agilent's VTEP in-circuit test (ICT) software and bead probe technology, the system allows users to test PCBAs in a limited-access environment without sacrificing test coverage. As a result, designers can develop circuit boards with less electrical test access, ultimately saving fixture costs and reducing test resources. Fewer probes in the test fixtures also results in reduced PCBA strain.

● Agilent joins LTE/SAE trial initiative

Agilent has joined the 3GPP Long Term Evolution/System Architecture Evolution (LTE/SAE) Trial Initiative for high-speed wireless broadband technology validation. The initiative verifies LTE's ability to achieve true broadband performance in mobile devices. It is divided into three main phases: proof of concept, interoperability and trial. Agilent will contribute LTE interoperability testing and field trials for physical layer air interface testing through protocol layer test and network diagnostics.

● Improved workflow features added to gene expression system

Agilent's gene expression bioinformatics system now features guided workflow features designed to improve the user experience. The GeneSpring GX 9.0 visualization and analysis system provides step-by-step guidance for major microarray platforms, including Agilent SurePrint, Affymetrix GeneChips and Illumina Beadchips. New workflow features enable users to ask detailed questions about complex data sets such as t-tests, two-way and three-way ANOVA tests and one-way post-hoc tests, injecting reliability into the gene identification process.


MATLAB is a U.S. registered trademark of The Math Works, Inc.

Turning a “Good Enough” Test Strategy into one that’s Reliable, Repeatable and Traceable

Haze Hutmacher

Application Engineer, Agilent Technologies

haze_hutmacher@agilent.com



Often times, there's a surprisingly short distance between a test strategy that is "good enough" and one that is reliable, repeatable and traceable. Consider the case of a design under test (DUT) with three Ethernet ports, which must pass bi-directional UDP and TCP traffic. While this may seem like a fairly trivial exercise for any high quality Ethernet test tool, when addressed with a test methodology involving re-purposed PCs, it can quickly become problematic.

This article examines a real-world test strategy for addressing the DUT described above. It will illustrate how minor changes in this strategy can result in a higher quality test methodology.

Sketching the existing solution

Figure 1 shows the test methodology currently employed by a manufacturer to test the DUT. Using one test technician, two PCs and a shareware software program, the test technician would perform the following steps:

1. Send a port A enable command
2. Send a traffic enable command
3. Send UDP bi-directional traffic
4. Send TCP bi-directional traffic
5. Re-configure cables from port A to port B
6. Send a port B enable command
7. Send a traffic enable command
8. Send UDP bi-directional traffic
9. Send TCP bi-directional traffic
10. Manually document pass/fail results
11. Reconfigure for the next DUT

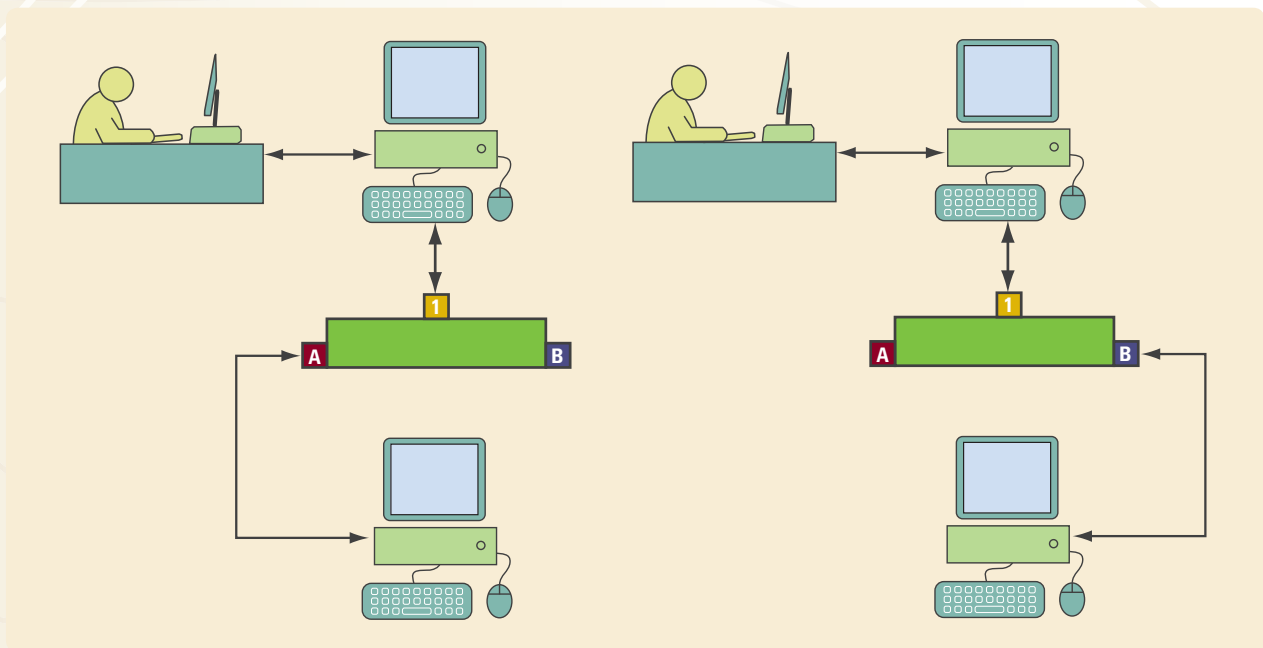


Figure 1. The initial test methodology required several manual steps to test ports A and B.

Identifying the major issues

A cursory glance at this methodology shows potential problems ranging from unknown pedigree of the shareware authors, extreme human intervention in the test process, an unknown ability to vary test parameters, lack of repeatability, no firm audit trail to test results, and no traceability to a fixed test methodology. Further complicating matters, the actual DUT must be configured through Port 1 to send bi-directional traffic from Port 1 to and from Port A. It must then be reconfigured to send traffic from Port 1 to and from Port B.

Creating a better solution

Improving this test strategy required a method for sending control data over the traffic link so that the ports could be reconfigured programmatically. Agilent's N2X system, a comprehensive multi-services test solution for converging network infrastructures, proved to be ideal for this task because of its ability to create proprietary or custom payload data units (PDUs), and capture and decode Ethernet traffic.

Before implementing the N2X system into the test strategy, it was first necessary to determine what was needed to configure the DUT ports. Data from a standard test cycle was captured using the free Wireshark software package. A traffic pattern emerged which could be duplicated using the N2X PDU builder functionality (Figure 2).

A test PDU was built on the N2X system and transmitted to a separate N2X port where the data was captured, decoded and verified as correct. The test PDU was then sent to the DUT where physical inspection verified that it had caused the desired response. Crucial to this test was the N2X system's ability to send a single unique PDU followed by another unique PDU, creating the necessary command string for DUT configuration. This iterative process was used to build, test and verify the three specific control signals needed to configure the DUT.

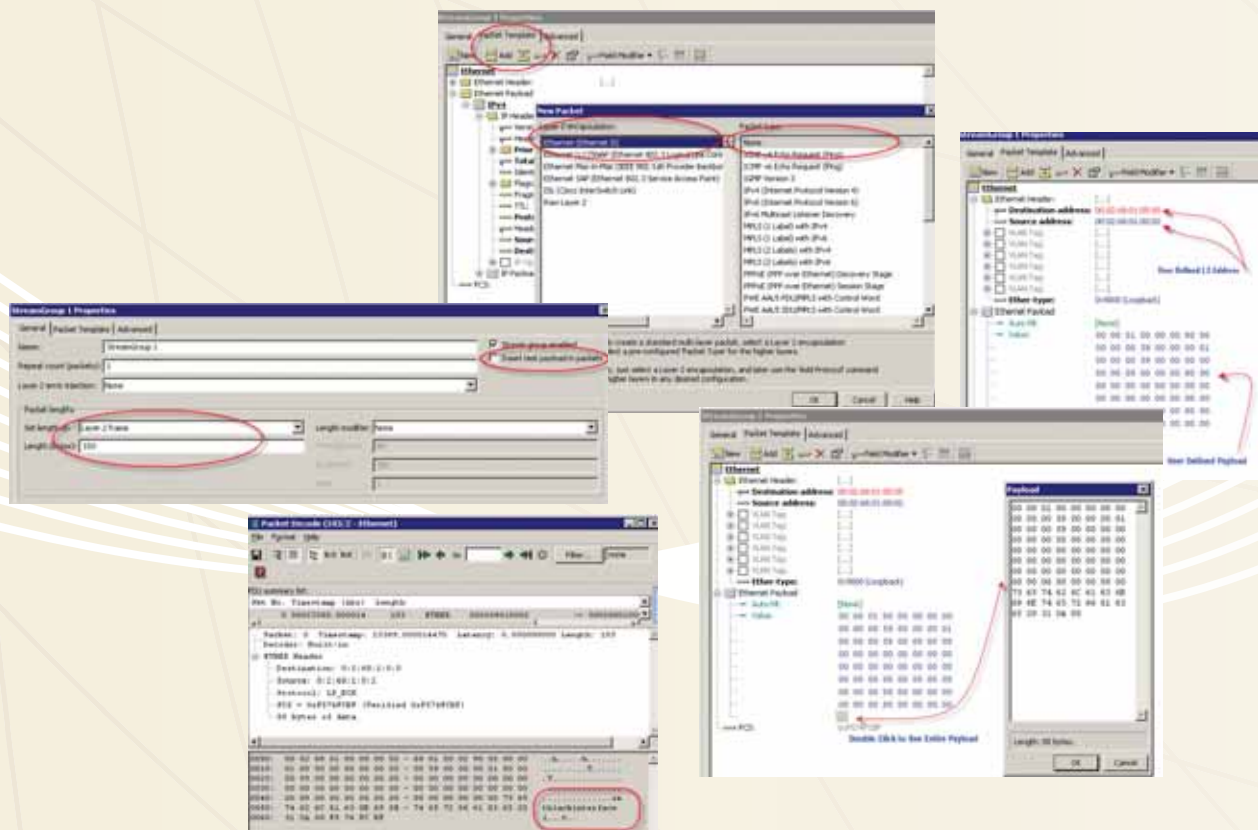


Figure 2. Working clockwise from the left, the N2X system walks the user through the process of building, testing and verifying a specific PDU to be sent to a DUT.

Outlining the benefits

Implementing the N2X system in the original test methodology resolved the problems previously identified. It automated the test process and removed the human intervention for anything other than inserting and removing the DUT into the test fixture and recording test results (Figures 3 and 4). It also allowed test parameters (e.g., frame size, data rate and payload size) to be varied using either a Quick Test (pre-written scripts designed to accomplish routine test tasks on the N2X) or custom script to quantify the operation of each DUT. Additionally, the N2X system's ability to test five DUT's at one time, using the 16-port Ethernet test module, provided a 500 percent increase in test throughput.

Name	Packet	L3 Source	L3 Destination	L3 Priority	Stream
Port 1037 (50.00%)					
Profile 1 (10.00%)	Enable Port A	Ethernet	*	*	*
Profile 2 (10.00%)	Enable Port B	Ethernet	*	*	*
Profile 3 (10.00%)	Enable Manufacture Test	Ethernet	*	*	*
Profile 4 (10.00%)	UDP Traffic	UDP (For 3H+3)/3H+4/Ethernet	192.1.1.2	192.9.1.2	0x00
Profile 5 (10.00%)	TCP Traffic	TCP (For 3H+3)/3H+4/Ethernet	192.1.1.2	192.9.1.2	0x00
Profile 6					
Profile 7					

Figure 3. A key feature of the test scenario, as illustrated in this profile configuration, was the ability to send a single frame of Ethernet data, select and send another single frame of data, and then blast traffic to verify the DUT's ability to pass traffic at a specified rate. Agilent's N2X system allows up to 15 individual traffic profiles on each port.

By scripting the test methodology into a single package that could be used for every test run, the issues of measurement repeatability, audit trail and traceability were resolved. Questionable metrics can now be traced back to the test script and evaluated in light of what has occurred during every test. Test results are logged to a data file on the N2X controller and may be accessed using FTP or TELNET from other PCs within the network.

Conclusion

While using re-purposed PCs in an in-house test system might seem like a good idea, often times a quality test strategy can be achieved with just a few minor tweaks in the process. In this case, use of the N2X system resulted in a repeatable, reliable and traceable test methodology for realizing cost savings (in terms of the test technician's time), reduced test errors and fixture wear and tear, and improved throughput.

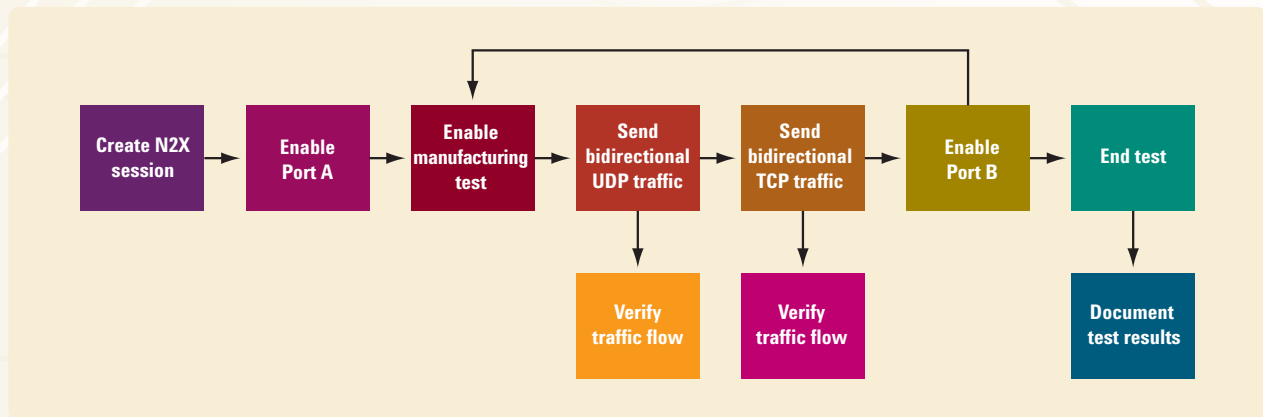
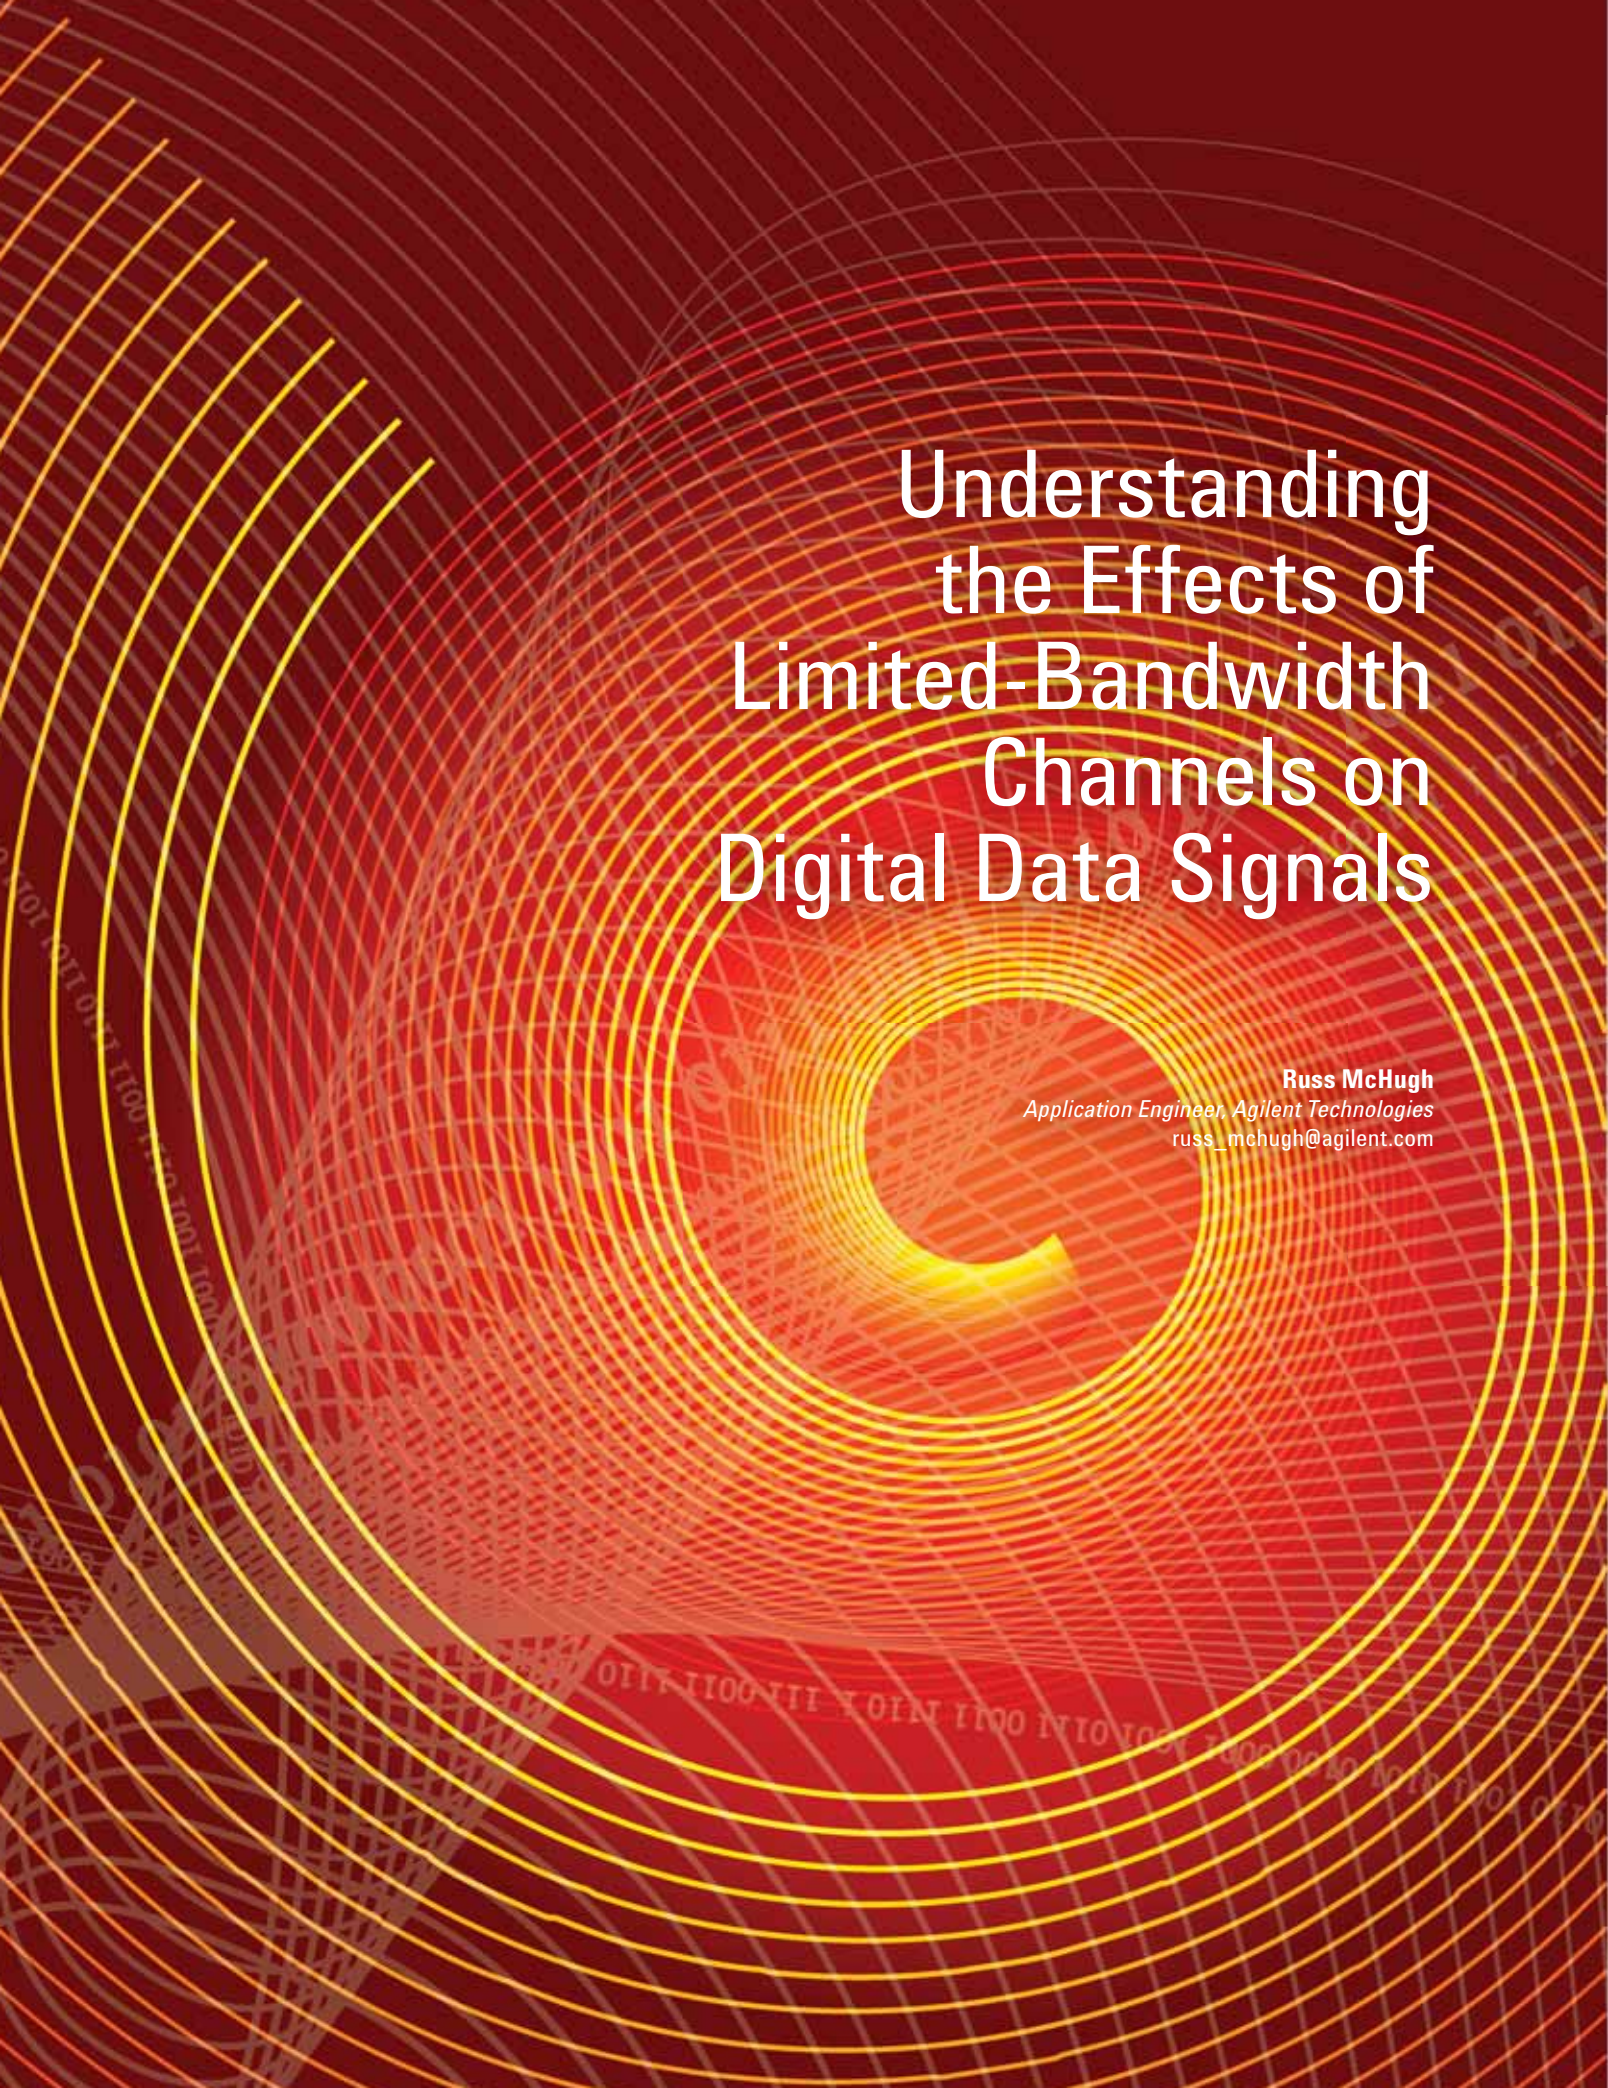


Figure 4. Using the profile configuration in Figure 3, an initial flowchart was created that required no user intervention. Once the test bed is cabled, the program sends a single PDU from Profile 1 and then Profile 3. Finally it sends bi-directional traffic using Profiles 4 and 5 to verify the DUT's throughput capability. The entire process is repeated using Profile 2 and Profile 3 to configure the Port B traffic.

The background features a series of concentric, glowing yellow and orange circles that create a tunnel-like effect. The circles are set against a dark red background. Faint binary code (0s and 1s) is scattered throughout the scene, particularly along the edges of the circles.

Understanding the Effects of Limited-Bandwidth Channels on Digital Data Signals

Russ McHugh

Application Engineer, Agilent Technologies
russ_mchugh@agilent.com

Data rates in digital systems are on the rise. These increased data rates need to be maintained in deployed systems and in new systems where cost is a concern. These constraints produce systems in which the bandwidth of the channel (e.g., backplane, circuit board) can not be increased as the data rate increases. At the same time, the channel's finite bandwidth heavily limits its high-speed performance. Consequently, this limited bandwidth constrains the maximum allowable data rate that can be transmitted through the channel. One of the limiting factors in increasing the data rate is the intersymbol interference (ISI) created by the limited-bandwidth data channel.

Limiting factors such as implementation costs and in-place systems hinder a designer's ability to increase channel bandwidths. In deployed systems, for example, it is not practical to change out the backplane to increase the channel bandwidth. It is therefore necessary to determine how much ISI the existing channels produce at the new higher data rate. For new designs, cost constrains the developer's ability to create a higher-bandwidth channel. Higher-bandwidth materials exist and can be used in place of standard FR4 but these materials come at a greatly increased cost.

This article examines the effect of limited-bandwidth channels on digital data signals. It also explains how to overcome these effects through compensation for the bandwidth limitation by emphasis and equalization.

Examining data signals in the time domain

Digital designers traditionally work in the time domain. Using the primary tool of choice — the oscilloscope — they look at how a signal's voltage levels vary with time. Figure 1 offers an example of a captured signal displayed in the form of a single-valued waveform. For each point in time, or position on the oscilloscope display, the waveform has one value. The rise and fall times of the signal, as well as the zero and one levels, can all be easily identified on the display; however, due to a typical screen's limited size and resolution, only a limited duration of the waveform can be seen. Because an overall view of the signal is not provided, it is easy to miss a waveform's extreme values. The designer can scroll the oscilloscope display to view different sections of the waveform, but still it is impossible to get an overall picture.




Figure 1. Display of a captured signal as a single-valued waveform

The technique used to obtain an overall view of the digital signal is to display an eye diagram, which is an overlay of the different bit transitions plotted one on top of the other (Figure 2). In the eye diagram, maximum and minimum values of the zero and one levels can be seen along with the extremes of rise and fall times.




Figure 2. Data signal eye diagram

Viewing data in the frequency domain

When digital signals propagate through a bandwidth-limited channel, they are low-pass filtered and the higher frequency content is attenuated. To understand this effect, look at the data signal in the frequency domain (e.g., the signal's spectrum). In the simplest case of one bit, or impulse, the spectrum takes the form of a sinc function (e.g., $\sin(x)/x$) (Figure 3).




Figure 3. Spectrum of an impulse, the sinc function

On a spectrum analyzer, the measured power spectrum shows the absolute value of the $\text{sinc}^2(f)$ function. If the digital signal is a repeating pulse, the resulting power spectrum will be spectral lines rather than a continuous envelope. The spectral nulls fall at the pulse frequency and its multiples. The spectral lines follow the pulse repetition rate. If the digital signal is more complicated than a repeating pulse such as a data signal, its spectrum is more complex. In this case, the spectral nulls will fall at the data rate and its multiples, similar to the pulse train. For a repeating pattern, such as a PRBS7, the spectral lines fall at the inverse of the pattern length in time. For a 5-Gb/s data rate, PRBS7 data signal, the spectral lines will be spaced at $5\text{e}9/127$ or 39.4 MHz (Figure 4). The longer the pattern, the closer the spectral lines fall and the more closely the spectrum envelope follows the $\text{sinc}^2(f)$ function.




Figure 4. Spectral line spacing of the PRBS7 pattern

The other significant effect on the spectrum is the transition time of the signal (e.g., its rise and fall times). The faster the transition time, the more energy there is at higher frequencies, which is indicated by increases in the power of the higher-frequency spectral lines. With the channel acting as a low-pass filter, it has a greater effect on these higher-frequency components. Increased data rates translate into a short bit time or period, necessitating faster transition times. The faster transition times mean that more of the signal's energy is at higher frequencies and therefore, the bandwidth limitation of the channel has a larger effect on the signal spectrum.

Characterizing channel bandwidth

An ideal channel has infinite bandwidth and does not affect the spectrum of a data signal. In the past, with data rates in the 100 to 200 Mb/s range, this assertion was basically true. With data rates now in the multi-Gb/s range, this assertion is no longer valid.

The bandwidth of common circuit-board materials, such as FR4, is limited in relation to these new, higher data rates. This bandwidth can be displayed in the frequency domain as S (scattering) parameters. The S-parameter for transmission is S_{21} . It displays the gain or loss (as in this case) of the signal propagating through the channel. The ideal channel has zero gain (loss) across the signal's frequency range, although a real channel will have increasing loss at higher frequencies. A sample channel derived using the 20-inch trace on a demo board was measured and is shown in Figure 5. The figure depicts a -3 dB point (e.g., 3 dB below the level at the minimum frequency) of 945 MHz. The channel bandwidth measurement was made up to 20 GHz. As will be seen, multi-Gb/s data signals have spectral



Figure 5. S_{21} of the 20-inch trace on a demo board

Observing the bandwidth-limited spectrum

To observe the bandwidth-limited spectrum, it is necessary to examine the spectrum of the data signal after it has passed through the channel. Figure 6 shows the original data signal spectrum (in yellow) and the spectrum after propagation through the demo board (in blue). The drop in amplitude starts almost at 0 Hz and is particularly large above 5 GHz.



Figure 6. Data signal spectrum before (yellow) and after (blue) propagating through the demo board

The same data signals can be viewed in the more familiar time domain (Figure 7). Notice that the limited bandwidth of the channel has slowed the transition time of the signal such that the edges appear more rounded and the amplitude of the isolated bits is reduced (lower trace in Figure 7).




Figure 7. Limited channel bandwidth effect displayed in the time domain (lower trace)

This change is easier to see when viewing this signal as an eye diagram (Figure 8). The eye, though still open, has been “pinched” and is more closed (lower trace in Figure 8). For some bit transitions, the amplitude does not rise to the “one” level at all. Reducing the channel bandwidth by having a longer trace, or increasing the data rate would potentially cause the eye to completely close.




Figure 8. Limited channel bandwidth effect displayed as an eye diagram (lower trace)

Dealing with the limited-bandwidth effect

To achieve higher data rates through existing or low-cost channels, the designer must compensate for the bandwidth limitation. Because the channel attenuates the higher frequency components, they must be amplified. Two common techniques exist for addressing this task.

One option is to apply emphasis at the transmitter, increasing the transmitted energy in the higher frequencies. This is done by increasing the amplitude of the data signal on the transition bits relative to the amplitude on the non-transition bits (Figure 9). A transition bit occurs where the data signal transitions from a one to a zero or from a zero to a one. Non-transition bits occur where the signal remains at the same level.




Figure 9. Transition bits showing emphasis

Because it is the transitions that contain higher frequency energy, emphasis increases the energy in the higher frequencies that are low-pass filtered by the channel. Emphasis can be created by boosting the transition level above the nominal signal level (pre-emphasis) or by reducing the non-transition level below the nominal signal level (de-emphasis). Since most digital standards constrain the maximum signal level, de-emphasis is the more common implementation. The effect that emphasis can have on a bandwidth-limited signal can be dramatic. As an example, compare the bandwidth-limited eye diagram in the lower trace in Figure 8 with the emphasized bandwidth-limited eye diagram in the lower trace of Figure 10.




Figure 10. Eye diagrams with emphasis, before and after the bandwidth limited channel

The other option is to apply equalization at the receiver to amplify the received high-frequency components. This compensates for the energy lost through the bandwidth-limited channel (Figure 11). The main limitation to equalization is noise. When amplifying the high frequencies of the signal, the noise in the signal is also amplified. This limits the maximum amount of equalization that can be applied. Some systems employ both emphasis and equalization to achieve higher data-transmission rates. The level of emphasis and equalization may even be varied during operation to accommodate variations over time in the channel-frequency response.

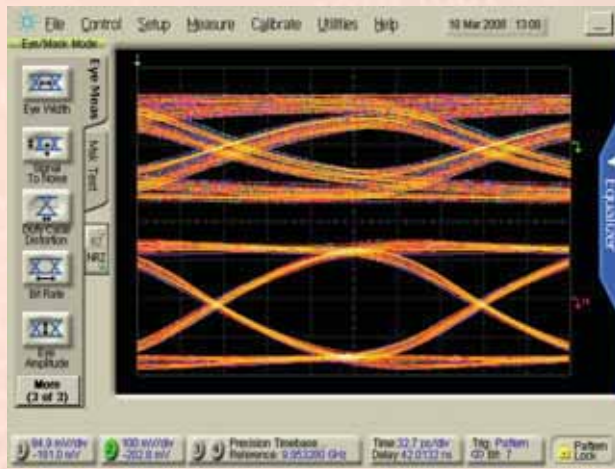


Figure 11. Eye diagram with equalization after bandwidth-limited channel (lower trace)

Conclusion

Digital signals, traditionally viewed in the time domain, can also be viewed in the frequency domain. Here the designer can see that data signals have frequency content above the nominal data rate that is important for signal quality. When these data signals propagate through a real channel (e.g., one with limited bandwidth), they are low-pass filtered and the higher-frequency content is attenuated. Looking at this band-limited data signal in the time domain, the effects on rise time, fall time and eye-diagram shape become visible. This limited-bandwidth effect is one of the ultimate limits on the data rate of a signal that will propagate through the channel. To deal with this effect, the designer can choose to employ emphasis at the transmitter, equalization at the receiver, or both.

The background features a city skyline at night with illuminated buildings. Overlaid on this is a blue hexagonal grid pattern. A network diagram with nodes and connecting lines is also visible, particularly in the lower right quadrant. The title text is positioned in the upper right area.

Introducing the 3GPP LTE Downlink

Moray Rumney BSc, C. Eng, MIET

Lead Technologist, Agilent Technologies

moray_rumney@agilent.com

Mitsuru Yokoyama

Architecture Lead Engineer, Agilent Technologies

mitsuru_yokoyama@agilent.com

Long-Term Evolution (LTE) is a project name of the Third Generation Partnership Project (3GPP) which aims to define a new air interface for mobile communications. LTE is an evolution of 3GPP's Universal Mobile Telecommunication System (UMTS) towards an all-IP network. It will provide a framework for increasing data rates and overall system capacity; reducing latency; and improving spectral efficiency and cell-edge performance (Figure 1).¹

Unlike UMTS, which is based on wideband code division multiple access (W-CDMA) technology, LTE is based on orthogonal frequency-division multiple access (OFDMA). In this regard, LTE is similar in concept to Mobile WiMAX™, another emerging technology for wireless broadband access, although the systems operate with different frame structures, subcarrier spacing and channel bandwidths.

In Issue Four of *Agilent Measurement Journal* we covered the LTE uplink in the article "3GPP LTE: Introducing Single-Carrier FDMA." This article will focus on the OFDMA LTE downlink, briefly describing the physical layer characteristics as defined in the 3GPP Release 8 specifications.

Overview of OFDM

To better understand the LTE downlink, let's look first at how the downlink resources are accessed by the users. The LTE downlink uses OFDMA, which is a variant of orthogonal frequency-division multiplexing (OFDM). For an overview of OFDM technology, please see "Understanding the Use of OFDM in IEEE 802.16 (WiMAX™)" in Issue Two of *Agilent Measurement Journal*. Although the article explains the basics of OFDM with reference to WiMAX, the general principles apply to LTE as well.

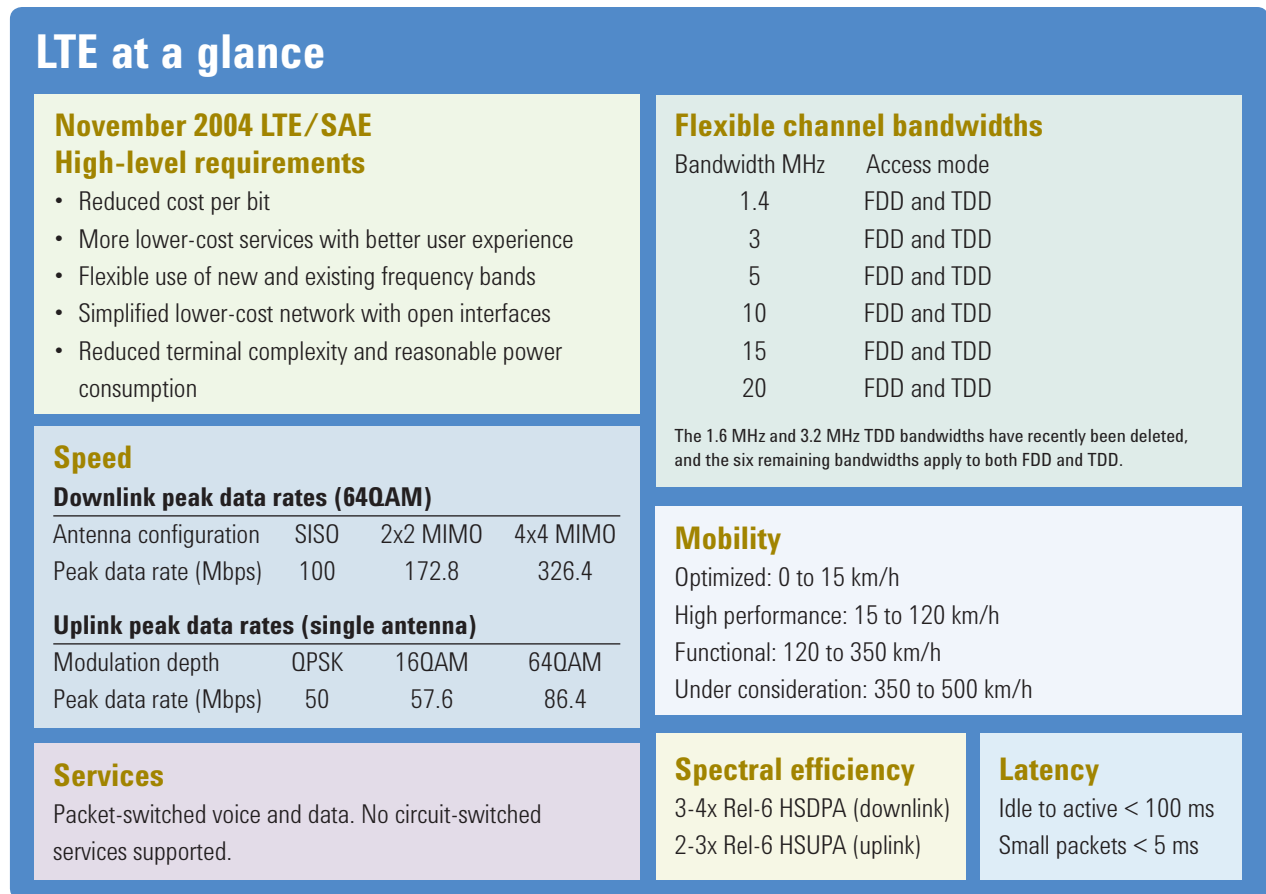


Figure 1. A look at some of the key characteristics of LTE

Rather than transmit a high-rate stream of data with a single carrier, OFDM uses a large number of closely spaced, orthogonal subcarriers that are transmitted in parallel. Each subcarrier is modulated with a conventional modulation scheme (e.g., QPSK, 16QAM or 64QAM), at a low symbol rate. The combination of hundreds or thousands of subcarriers enables data rates similar to conventional single-carrier modulation schemes in the same bandwidth.

Figure 2 illustrates several key features of an OFDM signal. In the frequency domain, multiple adjacent subcarriers are each independently modulated with data. Guard intervals are inserted between each of the symbols in the time domain to prevent intersymbol interference caused by multipath in the radio environment. The fast Fourier transform (FFT) and its inverse (IFFT) are used to generate and detect OFDM signals.

OFDM offers a number of distinct advantages in comparison to CDMA:

- OFDM easily scales up to wide channels which are more resistant to fading.
- OFDM channel equalizers are much simpler to implement because the OFDM signal is represented in the frequency domain rather than the time domain.
- OFDM is resistant to multipath because its long symbols can be separated by a guard interval also known as the cyclic prefix (CP). By sampling the received signal at the optimum time, the receiver can remove the time-domain interference between adjacent symbols caused by multi-path delay spread in the radio channel.

- OFDM is better suited to multiple antenna techniques such as MIMO. The frequency domain representation of the signal enables easy pre-coding to match the signal to frequency and phase characteristics of the multi-path radio channel.

OFDM also has some disadvantages. Its subcarriers are closely spaced, rendering OFDM sensitive to frequency errors and phase noise. This also makes OFDM sensitive to Doppler shift, which causes interference between the subcarriers. In addition, pure OFDM creates high peak-to-average ratio (PAR) signals, which cause difficulty with power amplifier design and power consumption.

In comparison to CDMA, OFDM is more difficult to operate at the network's cell edges. Whereas CDMA relies on different scrambling codes to protect against inter-cell interference at the cell edge, OFDM has no such feature. Some form of frequency planning at the cell edge is therefore required. A comparison of CDMA and OFDM is given in Table 1.

Table 1. Comparison of CDMA and OFDM

Attribute	CDMA	OFDM
Transmission bandwidth	Full system bandwidth	Variable up to full system bandwidth
Symbol period	Very short – inverse of the system bandwidth	Very long – defined by subcarrier spacing and independent of system bandwidth
Separation of users	Orthogonal spreading and scrambling codes	Primarily frequency and time, although scrambling and spreading can be layered on top




Figure 2. OFDM signal represented in frequency and time, taken from 3GPP Technical Specification TS 25.892²

Extending OFDM with OFDMA

With standard OFDM, low-rate user equipment (UE) transmissions occupy narrow frequency allocations and can suffer from narrowband fading and interference. Because of this, the LTE downlink uses OFDMA, which adds time-division multiple access (TDMA) to basic OFDM. Figure 3 shows that with basic OFDM each user is allocated a fixed set of subcarriers. With OFDMA, subcarriers are allocated dynamically among different users on the channel. The result is a more robust system with increased capacity, attributable to the efficiency gained from multiplexing low-rate users and dynamically scheduling users by frequency (according to each user's instantaneous channel conditions).

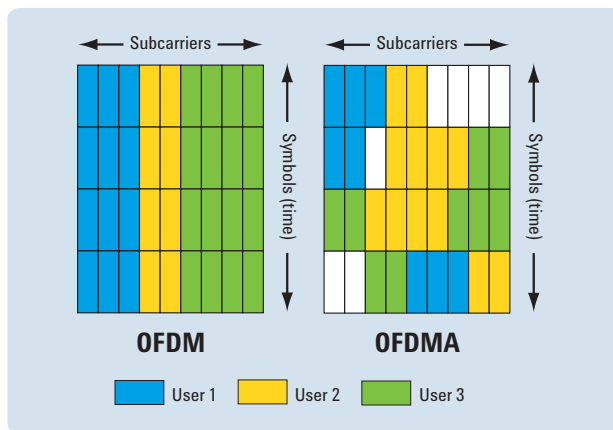


Figure 3. OFDM and OFDMA subcarrier allocation

A simplified OFDMA data transmission is summarized in Figure 4. The example uses four (M) subcarriers over two symbol periods with the payload data represented by QPSK modulation. Real LTE signals are allocated in units of 12 adjacent subcarriers and will be described in more detail shortly.

The M adjacent subcarriers are spaced 15 kHz apart and positioned at the desired place in the channel bandwidth. Each subcarrier is modulated for the OFDMA symbol period of 66.7 μ s by one QPSK data symbol. In this four subcarrier example, four data symbols are taken in parallel. Since these are QPSK data symbols, only the phase of each subcarrier is modulated, therefore subcarrier power remains constant between symbols. After one OFDMA symbol period has elapsed, a CP is inserted and the next four symbols are transmitted in parallel. The CP is shown as a gap for visual clarity; however, it is actually filled with a copy of the end of the next symbol. This means

that the transmission power is continuous but has a phase discontinuity at the symbol boundary. To create the transmitted signal, an IFFT is performed on each subcarrier to create M time-domain signals. In turn, these are vector-summed to create the final time-domain waveform used for transmission.

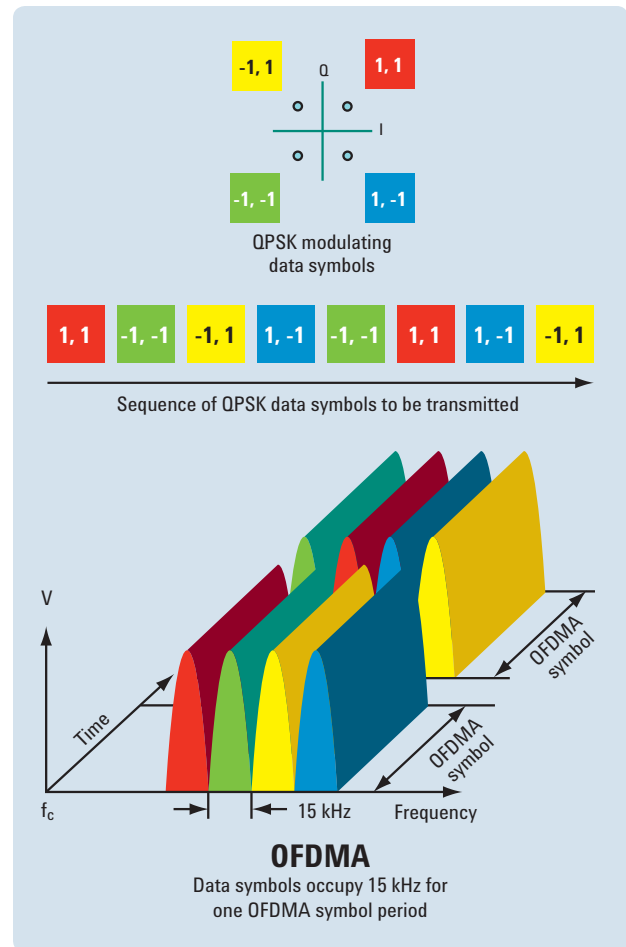


Figure 4. OFDMA transmitting a series of QPSK data symbols

The OFDMA signal in Figure 4 is clearly multi-carrier and it is this parallel transmission of multiple symbols that creates the undesirably high PAR of OFDM. As the number of subcarriers increases, the PAR of the OFDMA signal (with random modulating data) approaches Gaussian noise statistics. High PAR creates problems for power amplifier design and is a main reason why 3GPP developed SC-FDMA with its lower PAR for the LTE uplink.

Table 2. LTE downlink physical signals and physical channels

Downlink signals	Full name	Purpose
P-SCH*	Primary synchronization signal	Used for cell search and identification by the UE. Carries part of the cell ID (one of three orthogonal sequences).
S-SCH*	Secondary synchronization signal	Used for cell search and identification by the UE. Carries the remainder of the cell ID (one of 168 binary sequences).
RS	Reference signal (pilot)	Used for downlink channel estimation. Exact sequence derived from cell ID (one of $3 \times 168 = 504$).

Downlink channels	Full name	Purpose
PBCH	Physical broadcast channel	Carries cell-specific information
PMCH	Physical multicast channel	Carries the multicast (MCH) transport channel
PDCCH	Physical downlink control channel	Scheduling, ACK/NACK
PDSCH	Physical downlink shared channel	Payload
PCFICH	Physical control format indicator channel	Defines number of PDCCH OFDMA symbols per sub-frame (one, two or three)
PHICH	Physical hybrid ARQ indicator channel	Carries HARQ ACK/NACK

* Note: There are no formal acronyms to describe the primary and secondary synchronization signals. The terms P-SCH and S-SCH come from earlier 3GPP technical reports and are still used informally despite their suggestion of “channel” rather than “signal.”

LTE downlink structure

The LTE physical layer supports the use of paired and unpaired spectrum with frequency-division duplex (FDD) and time-division duplex (TDD) modes, respectively. Each of these modes has its own frame structure, described shortly. First, however, we will consider the composition of the downlink which comprises physical signals and physical channels as described in the 3GPP specifications.³

Downlink physical signals are generated in Layer 1 and used for system synchronization, cell identification and radio channel estimation. The primary synchronization signal (P-SCH) and secondary synchronization signal (S-SCH) encode the cell identification data, allowing the UE to identify and synchronize with the network. Downlink reference signals (RS), known as pilot signals in other standards, are used by the UE receiver to estimate the phase and flatness of the received signal. Errors in the received signal are the combination of errors in the transmitted signal and further imperfections in the radio channel. Without the use of reference signals spaced across the channel bandwidth, phase and amplitude shifts in the received signal would make demodulation unreliable, particularly at high modulation depths such as 16QAM or 64QAM. With these high modulation depths, even a small error in the received signal amplitude or phase can cause demodulation errors.

Alongside the physical signals are physical channels, which carry data such as control, scheduling and user payload to and from the higher layers. Since LTE is a packet-only system there is no need to define dedicated channels since all data is carried on the shared channel. The function of each LTE physical signal and channel is summarized in Table 2. Table 3 shows the modulation schemes allowed for the downlink signals and channels.

Table 3. Modulation schemes for the LTE downlink

Downlink signals	Modulation scheme
P-SCH	One of three Zadoff-Chu sequences
S-SCH	Two 31-bit BPSK M-sequence
RS	Complex $I + jQ$ pseudo random sequence (length-31 gold sequence) derived from cell ID

Downlink channels	Modulation scheme
PBCH	QPSK
PMCH	QPSK, 16QAM, 64QAM
PDCCH	QPSK
PDSCH	QPSK, 16QAM, 64QAM
PCFICH	QPSK
PHICH	BPSK modulated on I and Q with spreading factor 2 or 4 Walsh codes

Configuring bandwidth

LTE is being designed to support the international mobile-wireless market, regional spectrum regulations and spectrum availability. To accomplish this, variable channel bandwidths from 1.4 MHz to 20 MHz are specified. Standard subcarrier spacing is 15 kHz but subcarrier spacing of 7.5 kHz is also possible for the new LTE evolved multimedia broadcast multi-cast system (eMBMS), which is to be specified in Release 9 of the 3GPP specifications. Subcarrier spacing is independent of the channel bandwidth.

The smallest amount of allocated resource in both the downlink and uplink is called a resource block (RB). An RB is 180 kHz wide and lasts for 0.5 ms. For standard LTE, an RB consists of 12 subcarriers at a 15-kHz spacing, while for eMBMS using the optional 7.5-kHz spacing, the RB is 24 subcarriers wide. The maximum number of RBs supported by each channel bandwidth is given in Table 4.

Table 4. Channel bandwidth configurations

Channel bandwidth (MHz)	1.4	3.0	5	10	15	20
Nominal transmission bandwidth configuration (MHz)	1.08	2.7	4.5	9	13.5	18
Nominal transmission bandwidth configuration (resource blocks)	6	15	25	50	75	100

Frame structure

Two radio-frame structures are specified for the LTE downlink: frame structure type 1 (FS1) for full-duplex and half-duplex FDD, and frame structure type 2 (FS2) for TDD. These frame structures are shown in Figures 5 and 6. FS1 is optimized to co-exist with existing FDD UMTS systems and consists of ten 1 ms sub-frames, each composed of two 0.5 ms slots for a total duration of 10 ms. FS1 is the same in the uplink and downlink in terms of frame, sub-frame and slot duration, although the composition of the physical signals and channels is different. Uplink and downlink transmissions use different spectra.




Figure 5. Frame structure type 1 (TS 36.211 Figure 4.1-1)³

The structure of FS2 is a lot more flexible than FS1. An example of an FS2 structure is shown in Figure 6. This example is for a 5 ms switch-point periodicity and consists of two 5 ms half-frames for a total duration of 10 ms. Subframes consist of either an uplink or downlink transmission, or a special subframe




Figure 6. Frame structure type 2 (for 5 ms switch-point periodicity), TS 36.211 Figure 4-2.1³

containing the downlink and uplink pilot timeslots (DwPTS and UpPTS) separated by a transmission gap (GP). The allocation of the subframes for the uplink, downlink and special subframes is determined by one of seven different configurations. Subframes 0 and 5 are always downlink transmissions, while subframe 1 is always a special subframe. The composition of the other subframes varies depending on the configuration. For a 5 ms switch-point configuration, subframe 6 is always a special sub-frame as shown in Figure 6. With 10 ms switch-point periodicity, there is only one special subframe per 10-ms frame. The remainder of the article will focus on FS1 for the FDD downlink.

Figure 7 shows the downlink resource grid for a 0.5-ms timeslot which incorporates the concepts of a resource element and a resource block. A resource element is the smallest identifiable unit of transmission and consists of one subcarrier for one symbol period. However, transmissions are scheduled in larger units called resource blocks which comprise 12 adjacent subcarriers for a period of one 0.5-ms timeslot.




Figure 7. Downlink resource grid (TS 36.211. V8.2.0 Figure 6.2.2-1)

Mapping the downlink

Figure 8 presents a more detailed view of FS1 for the downlink, showing the downlink slot structure color coded for the different signals and channels. As the diagram shows, an entire 10-ms frame is required for the control channels to repeat. The frame structure is defined in units of T_s , which is the shortest time interval of the system defined as $1/(15000 \times 2048)$ seconds or 32.552 ns.

Figure 8 shows how the timeslot is divided up into seven symbols. Each symbol is extended by the length of the CP by copying the end of the symbol to the beginning. This process does not add new information to the signal. Rather, the CP adds redundancy to counteract the inter-symbol interference caused by multipath delay spread. Using the normal CP length of $144 \times T_s$ (4.69 μs), it is necessary to make the first CP slightly longer at $160 \times T_s$ so that the timeslot adds up to 0.5 ms. The CP is chosen to be slightly longer than the longest expected delay spread in the radio channel. For LTE, the normal CP length enables the system to cope with path delay variations up to about 1.4 km. Note that this figure represents the difference in path length due to reflections, not the size of the cell.

The mapping of the downlink physical signals and channels for the example in Figure 8 is as follows:

- RS are transmitted at OFDMA symbol 0 of the first subcarrier and symbol 4 of the fourth subcarrier of each slot. This is the simplest case for single-antenna use. The position of the RS varies with the antenna port number and the CP length.
- P-SCH is transmitted on symbol 6 of slots 0 and 10 of each radio frame, and occupies 62 subcarriers centered on the DC subcarrier.
- S-SCH is transmitted on symbol 5 of slots 0 and 10 of each radio frame, and occupies 62 subcarriers centered on the DC subcarrier.
- PBCH is transmitted on symbol 0 to 3 of slot 1 of each radio frame, and occupies 72 subcarriers centered on the DC subcarrier.

For simplicity, the PMCH, PCFICH and PHICH are not shown in this example. Note that the control channels are contained within the central 1.08 MHz of the signal so that system operation can be independent of the channel bandwidth. The length

72 for the P-SCH and S-SCH gives high correlation when using an allocation of 6 RB (72 subcarriers). The length 62 for the PBCH means that it can be detected using an FFT of length 64, thereby minimizing UE complexity.

Figure 9 shows the downlink mapping across frequency and time. The central DC subcarrier of the downlink channel is not used for transmission, but is reserved for energy generated due to local-oscillator feedthrough in the signal-generation process.

Table 5 summarizes the options for CP length and number of symbols per timeslot. The extended CP of $512 \times T_s$ (16.67 μs) is available for use in larger cells and provides protection for up to a 5-km delay spread. The price for this increased protection is a reduction in system capacity since the extended CP allows for only six symbols per timeslot. The longest protection from delay spread is achieved when using the extended CP of $1024 \times T_s$ (33.33 μs) with the optional 7.5-kHz subcarrier spacing for eMBMS. This enables transmissions from multiple cells to be




Figure 8. Frame Structure 1 for downlink

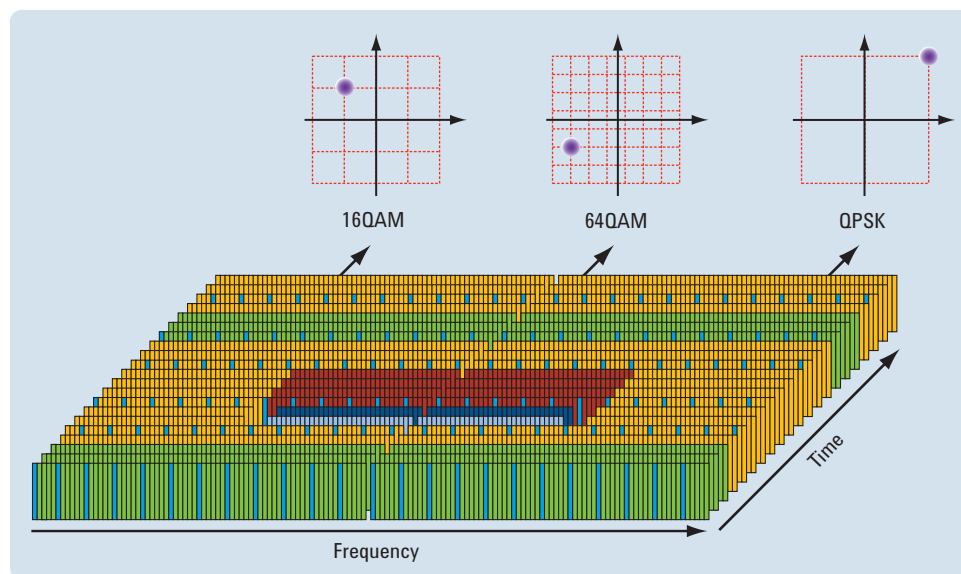


Figure 9. Frame Structure 1 for downlink showing one subframe versus frequency

combined in a Multicast/Broadcast over Single Frequency Network (MBSFN) with protection from delay spread of up to 10 km. This very long CP means there are only three symbols per timeslot, but this capacity loss is counteracted by the doubling up of the subcarriers.

Table 5. Cyclic prefix configurations for downlink FS1

		CP in Ts by symbol number						
		0	1	2	3	4	5	6
Normal	$\Delta f=15$ kHz	160	144	144	144	144	144	144
Extended	$\Delta f=15$ kHz	512	512	512	512	512	512	–
	$\Delta f=7.5$ kHz	1024	1024	1024	–	–	–	–

Verifying the downlink

Verifying the early functionality and performance of the LTE downlink requires flexible signal analysis. Figure 10 shows an outline of the downlink signal-generation process. First the incoming bit stream is scrambled using pseudo-random sequence generation. Next, the bits are mapped to a modulation symbol format. For example, the PDSCH can be mapped to QPSK, 16QAM or 64QAM containing two bits, four bits and six bits per symbol, respectively. Layer mapping is then applied to support various antenna configurations and precoding can be applied to adjust the phase and amplitude of each layer for each antenna. Next, resource mapping is applied and a downlink OFDMA signal is generated for each antenna. When the signal is received, it is demodulated using the inverse of this process.

In the real world, the base station and UE communicate with each other via signaling, exchanging essential information about the uplink and downlink signal composition. However, in the early phases of testing, this information exchange does not take place. Nevertheless, parameters for such things as resource mapping, modulation scheme, sequence generation, OFDM

signal generation, and power boosting are needed to allow the signal to be measured. Many of these parameters can be estimated or extracted from the signal under test, but for the remainder of the testing they must be entered manually.

Figure 11 shows one of the many signal configuration menus of the Agilent 89601A vector signal analysis software which allows users to specify numerous parameters for measuring the LTE downlink. Using this software, many key parametric and demodulation measurements can be performed to determine the downlink performance of LTE devices.

Figure 12 shows six different views of an LTE downlink signal. The top middle trace shows the auto-detected signals (P-SCH, S-SCH, PBCH, PCFICH, PDCCH, and RS) and one shared channel — PDSCH3 at 64QAM. Each channel type is color-coded to enable the different elements of the downlink to be identified on the other traces and independently measured. As well as auto-detecting the channels, this trace also measures the channel-specific EVM, the channel power and the modulation type. The top right trace gives further information about the modulation quality including the frequency error and the IQ offset and gain imbalance. The top left trace shows the IQ constellation of the signal, which shows everything from the Zadoff-Chu RS through the QPSK control channels to the 64QAM of the PDSCH3. The bottom left trace is a traditional power-versus-frequency spectrum, which shows this is a 5-MHz allocation. This is confirmed in the bottom middle trace which shows the EVM by subcarrier. Note that the X scale shows a range of 300 subcarriers which is an occupied bandwidth of 4.5 MHz at 15 kHz per subcarrier. EVM per subcarrier is an important measure since the EVM is likely to degrade at the channel edges due to the transmit filter. The final trace at the bottom right shows how the EVM varies in time across the 10-ms capture interval. This can be useful in determining downlink performance after transient events such as a power change.

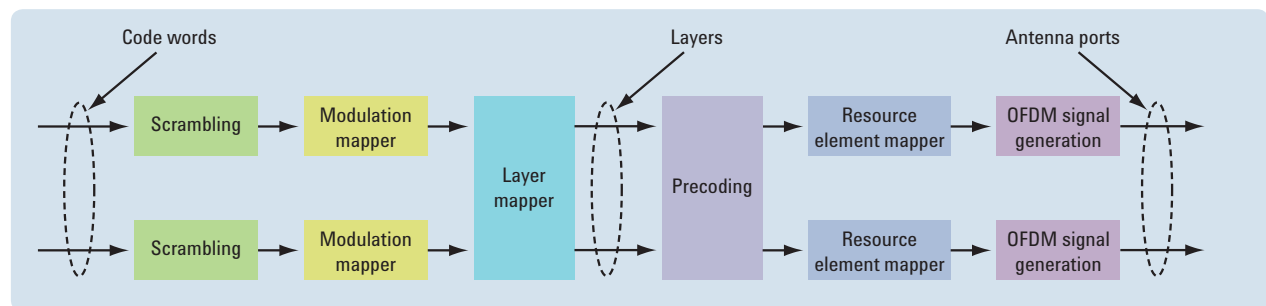


Figure 10. Downlink signal generation

Looking ahead

LTE has the potential to enhance current deployments of 3GPP networks and enable significant new service opportunities. Initially, though, LTE is expected to give equipment designers some difficulty because it is an evolving standard and, as such, is open to change and interpretation. From the technology

perspective, the number of new techniques used in the downlink (and the uplink as well) add substantial complexity. For example, the use of multiple antenna configurations to support high data rates makes the design of UE quite complicated, as does the introduction of the new downlink and uplink multiple-access schemes. It may be some time before the real-world behavior of these enhancements is well understood and products are optimized accordingly prior to system deployment.



Figure 11. Example of a setup menu for downlink signal analysis

References

1. Long Term Evolution of the 3GPP Radio Technology, www.3gpp.org/Highlights/LTE/LTE.htm
2. 3GPP TR 25.892 V6.0.0 (2004-06) www.3gpp.org/ftp/Specs/html-info/25892.htm
3. 3GPP TS 36.211 V8.2.0 (2008-03) www.3gpp.org/ftp/Specs/html-info/36211.htm

"WiMAX," "Fixed WiMAX," "Mobile WiMAX," "WiMAX Forum," the WiMAX Forum logo, "WiMAX Forum Certified," and the WiMAX Forum Certified logo are trademarks of the WiMAX Forum. All other trademarks are the properties of their respective owners.

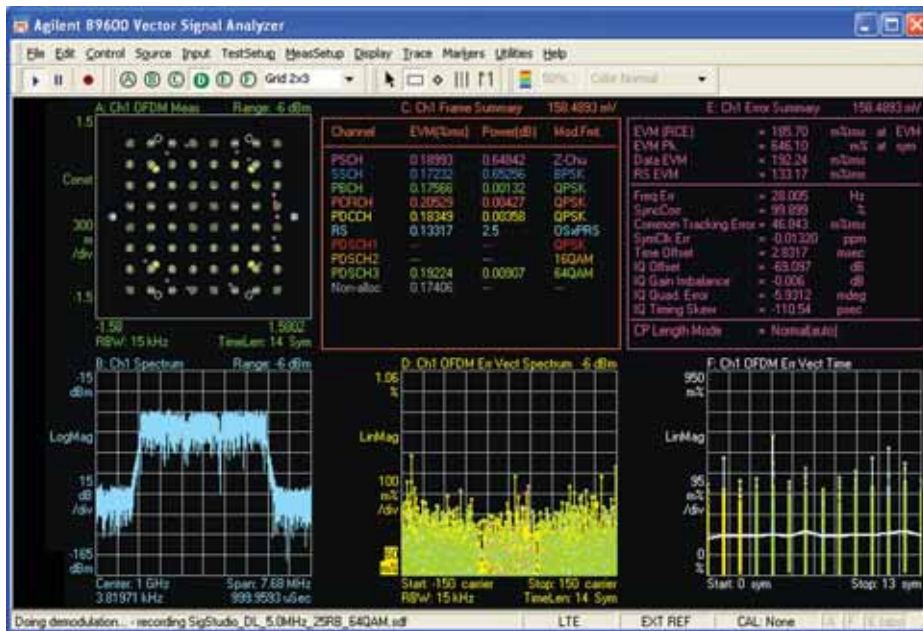


Figure 12. In-channel EVM measurements using 89601A LTE application

DDR3

x16

10mm

Validating the Physical and Protocol Layers in DDR Memory Interfaces

Min Jie Chong

Infiniium Scope Product Manager, Agilent Technologies

min-jie_chong@agilent.com

M

Memory devices are found almost everywhere, from computers and phones to HDTV sets and automobiles. One of today's most commonly used technologies is double data rate (DDR) synchronous dynamic random access memory (SDRAM). DDR technology is now in its third generation (DDR3) and offers faster transfer rates and lower energy consumption than previous versions.

The internal DDR architecture is special because it relies on parallel, single-ended signals that operate at the speed of today's new serial technologies. From a purely electrical perspective, parallel and serial buses generally don't work well together due to the interference generated between adjacent bus lines. Also, most of today's common high-speed architectures consist of differential signals that inherently reject common-mode noise. Thus, problems such as crosstalk, impedance mismatch, electromagnetic interference (EMI), jitter, and noise become highly likely in and around DDR devices. When characterizing DDR devices, probe loading also can affect the performance of measured signals.

As transfer rates increase, performance within the physical and protocol layers becomes increasingly important due to reductions in signal amplitude and timing margin — and performance in those layers is the key to system interoperability. Failures in the physical layer are correlated with protocol failures such as marginal timing relationships, protocol violations, clock jitter issues, and errors from other buses. After physical-layer validation has been performed, timing relationships and protocol violations can then be verified.

After outlining the challenges in validating DDR at the physical and protocol layers, this article suggests a few viable measurement methods that can overcome those challenges. These techniques are applicable to all DDR technologies and to the DDR side of fully buffered dual in-line memory modules (FB-DIMMs).

Overview: Probing methods

The Joint Electronic Device Engineering Council (JEDEC) specifies the DDR physical layer specifications in DDR devices for both the memory controller and the DDR DRAM. For true physical-layer compliance with the specification, the recommended approach is to probe at the DDR DRAM ball-grid-array (BGA) package ballout. This will provide greater insight into signal performance when performing physical layer testing. Because all signals are hidden beneath the package, however, it is difficult to probe directly at the DDR BGA ballout. The next best probing location is at the signal vias on the back of the circuit board because these are closest to the DRAM ballout.

Unfortunately, signal vias may not be present if components are densely arrayed on both sides of a circuit board. This suggests probing elsewhere, perhaps at signal traces or surface-mount resistors and capacitors. Although this might seem straightforward and easy, it often compromises the measurement

results because probing at these locations often causes signal reflections that produce non-monotonic edges. Rather than presenting true signal performance, the resulting measurements show a signal that has been affected by reflections from the components. This undesirable phenomenon can cause errors in slew-rate and setup/hold-time measurements since the non-monotonic edge shifts the timing at a specific voltage threshold. This problem can be avoided by probing near the DRAM ballout.

“Interposer probes” can be used for protocol validation by inserting the probe between the DIMM and connector.¹ Signals then are routed through the probe to a logic analyzer. Because this memory design is most commonly used in computer systems, this method does not apply to embedded systems in which the DDR DRAM is attached directly to a circuit board (e.g., without a connector). In such cases, the circuit board may have probe points that can be accessed by a logic analyzer; however, these might not be sufficient to enable probing of all

signals simultaneously due to the high density — and wide buses — of clock, strobe, data, address, and control signals. One alternative is to lay out a probing footprint on the boards prior to placing components for protocol validation. While this approach will solve the issues described above, it probably will require more design effort and may increase the manufacturing cost of larger boards.

BGA probe solution

A DDR BGA probe adapter can address the probing problems described above. The adapter is a thin fixture that can be attached between a DRAM chip and a circuit board because it has a compatible footprint on its top and bottom sides (Figure 1). The signals at the DRAM ballout are routed to the top side of the BGA probe adapter so they can be accessed by oscilloscope or logic-analyzer probes. This method provides a direct access point to the DRAM ballout for true physical and protocol validation versus the DDR specification.

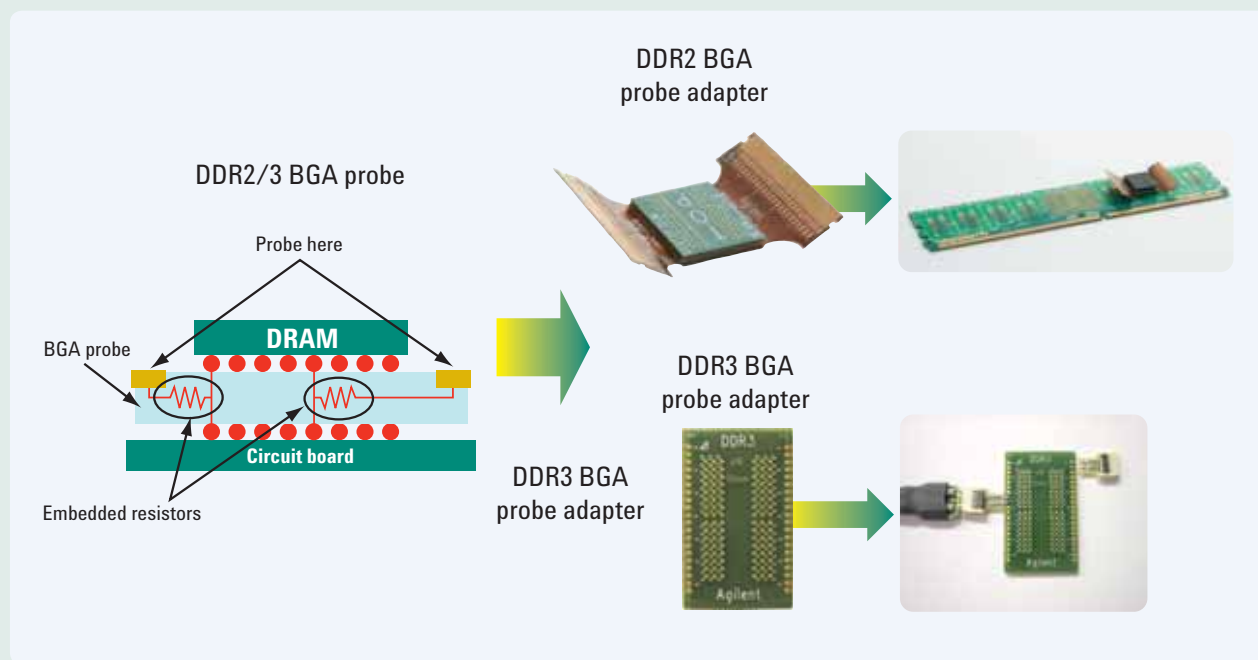


Figure 1. DDR2 and DDR3 BGA probe adapters are compatible with scope and logic-analyzer probes. Embedded resistors ensure highly accurate results by minimizing probe loading on incoming signals.

1. Interposer probes provide an electrically and mechanically nonintrusive connection between a logic analyzer and the DIMM, enabling the capture and observation of data traffic crossing the connector interface.

If the traces on the BGA probe adapter are designed with the same length, there is no skew between the signals. Inside the adapter, embedded resistors placed near the DDR signals prevent probe loading from interfering with the DDR signals. This design minimizes the capacitive loading of the stubs and probe, preserving operation of the high-speed DDR interface without impacting the actual signals. Waveforms obtained this way more closely represent actual signal performance and thereby provide highly accurate results for physical- and protocol-layer validation.

Examining the challenges in physical-layer validation

During physical validation, bi-directional transmission of strobe and data signals on the bus makes it difficult to separate the traffic moving between the memory controller and the DDR device. Examining these signals separately is required to enable independent analysis of their electrical and timing characteristics. If the signals cannot be isolated, it will be impossible to characterize the performance of the memory controller and the DDR device.

Three methods may be used to separate the read and write waveforms: triggering on the preamble bit width, triggering on the amplitude of the read or write signal and triggering on the command signals using a mixed-signal oscilloscope (MSO). Unfortunately, each of these separation methods has important limitations.

Width triggering: Because the JEDEC definition of the write preamble width is somewhat vague, it is best to assume a large variation in the width. In most cases, however, the write preamble width is similar to that of either the read preamble or the regular data-bit period. Due to these issues, there is a high likelihood that this separation method will not be effective.

Amplitude triggering: In some cases, it is possible to isolate the read and write signals if one is larger than the other; however, the read and write signals sometimes have similar amplitudes. As a result, predicting when this will be an effective way to separate the waveforms can be difficult.

Command-signal triggering: An MSO includes both analog and digital (logic) channels and can measure them simultaneously. As the DDR command signals are asserted during different operations, they can be used to trigger measurements of strobe and data signals. This is done by connecting the command signals to the MSO logic channels and the strobe and data signals to the MSO analog channels. There is one caveat: The bandwidth of a typical MSO is best suited to DDR devices with transfer rates of less than 400 MT/s.

Addressing the challenges

The latest oscilloscopes offer new capabilities that overcome the challenges of separating read and write signals, and also reduce the time and effort required to fully characterize a design. Examples include zone triggering and automated compliance measurements. With these functions, it is possible to exhaustively characterize and validate a memory interface in much less time than is required with manual methods.

Read and write separation with zone triggering

Setting the scope display to infinite persistence makes it possible to observe distinctive differences between read and write signal patterns. This is possible for two reasons: The data signals have different phase relative to the strobe signals, and no two pieces of silicon have identical electrical characteristics. By observing and understanding the meaning of these distinctive differences, it is possible to identify and isolate “zones” associated with the patterns of read and write signals. In Figure 2, the DDR signal shows distinctive patterns that can be isolated with a zone-triggering function. Once the read and write signals are separated, each can be easily measured.

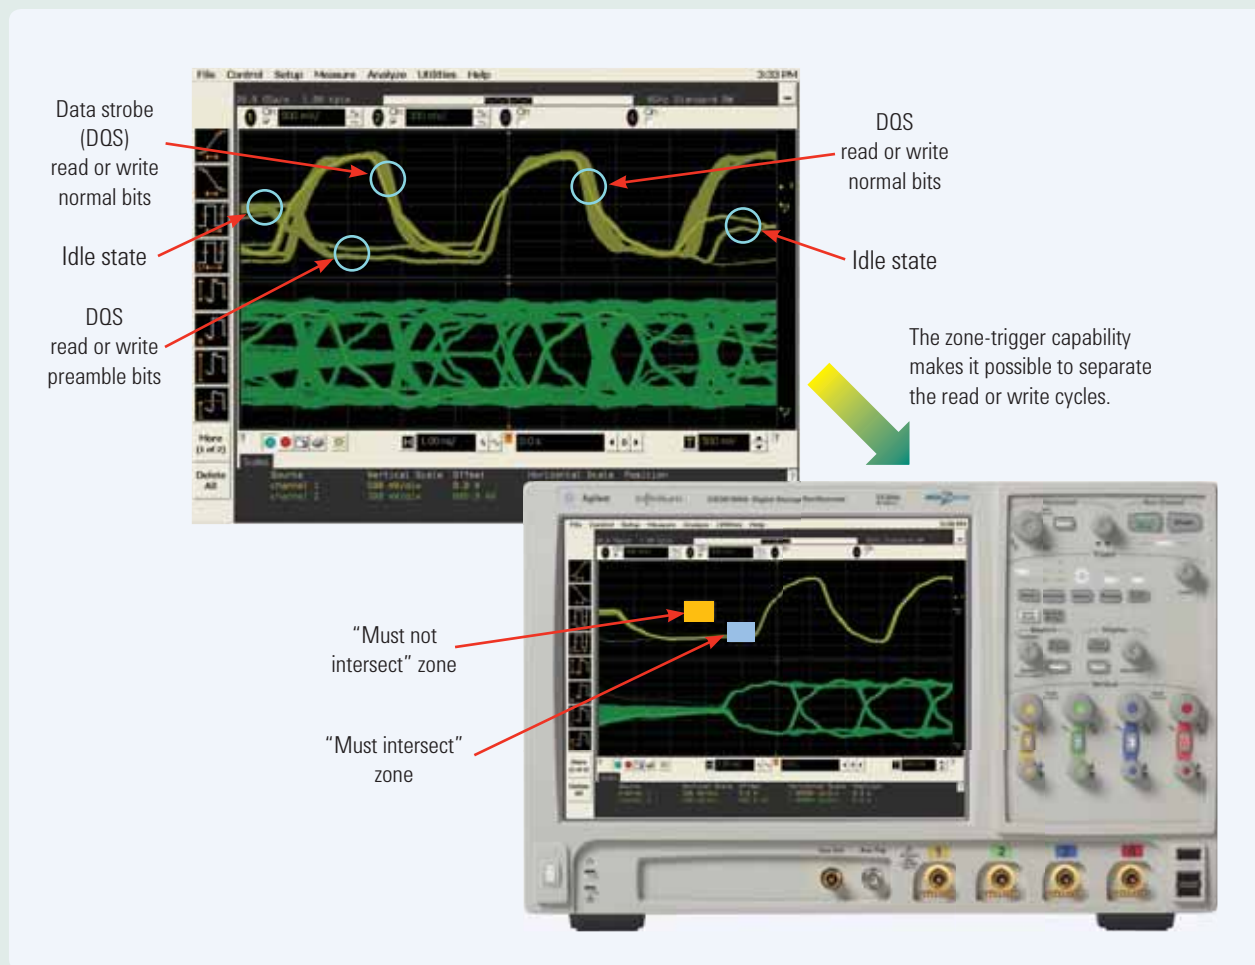


Figure 2. The zone-trigger capability makes it possible to separate the read or write cycles.

Automated measurements with a DDR application

Separating the read and write signals is just the first step of the validation task. It is still necessary to spend time with an oscilloscope manually validating each test parameter defined in the JEDEC specification. Because this is a very long list, however, it is often difficult to exhaustively characterize every test parameter. Even more problematic, the results must be manually recorded and formatted in a test report.

To save time and effort, many of the required steps can be automated with dedicated “applications” built into an oscilloscope. Using these automated routines, measurements of every test parameter can be repeated multiple times to thoroughly analyze a signal with complete statistical results, along with screen captures of worst-case results. Many applications also automatically generate comprehensive test reports for archiving or sharing.

Overcoming challenges in protocol-layer validation

As the DDR transfer rate gets faster, two common challenges crop up: signal-sampling position and signal correlation. Some logic analyzers provide capabilities that address these issues.

Signal-sampling position adjustment

The DDR technology transfer rate is quickly catching up with the sampling rate of the logic analyzer. As the data-valid window becomes smaller, a logic analyzer has to sample at the right position on the signal bit to capture the correct data. If the data is sampled outside of the data-valid window, the wrong data will be captured because the signal is sampled at a quasi-state or an invalid transition state. A “high” state may then be misinterpreted as a “low” state, or vice versa, and the logic analyzer may interpret the protocol incorrectly and consequently detect a too-high bit error rate.

To avoid quasi-state errors, some logic analyzers provide high-resolution data sampling capabilities that allow accurate sampling positioning of read and write data (Figure 3). The eye diagram provides comprehensive signal integrity information on all of the memory buses. The sample position can be adjusted to the center of the eye opening with 10 ps horizontal and 10 mV vertical resolution, enabling precise acquisition of accurate read and write data.

Signal correlation with protocol-aware sampling

The DDR3 and DDR2 DIMM architectures are markedly different, with DDR3 using a fly-by topology and DDR2 using a T-branch topology. In DDR2, the T-branch balances the signal delays to each memory device; however, the topology makes it difficult to manage signal reflections. In contrast, the fly-by topology used in DDR3 improves signal integrity on the command and address buses. The downside, however, is that signals from the memory controller arrive at each DDR DRAM at different times, causing skew. If no timing compensation is made, some signals will not be correlated. For example, if command or address samples are not correlated with the associated strobe and data information, protocol violations will occur.

A logic analyzer equipped with protocol-aware sampling and a time-adjustable delay circuit can compensate for the skew. With this capability, read and write cycle signals are properly aligned with the center of the eye opening and all signals will be time correlated. This ensures proper association of command and address signals with strobe and data, as well as accurate protocol-layer validation.

Conclusion

Rapidly evolving DDR memory technology is driving the need for new design and validation methods. With DDR speeds climbing higher to match those of the latest high-speed serial technologies, it is more likely that probe loading will distort the signal waveforms. For that reason alone, it is essential to pay special attention to probing — and every signal integrity measurement begins at the probe tip and reaches all the way back into the oscilloscope and logic analyzer. The ability to access all memory signals is important for full functional debug and testing with the logic analyzer. New capabilities now being built into oscilloscopes and logic analyzers greatly simplify the physical and protocol validation of DDR devices and interfaces.

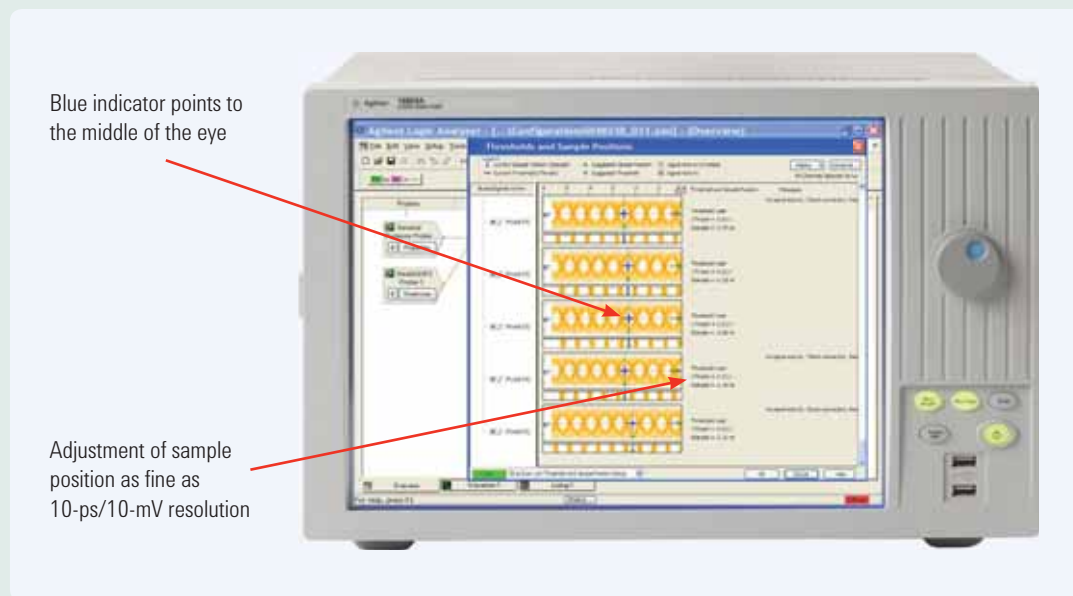
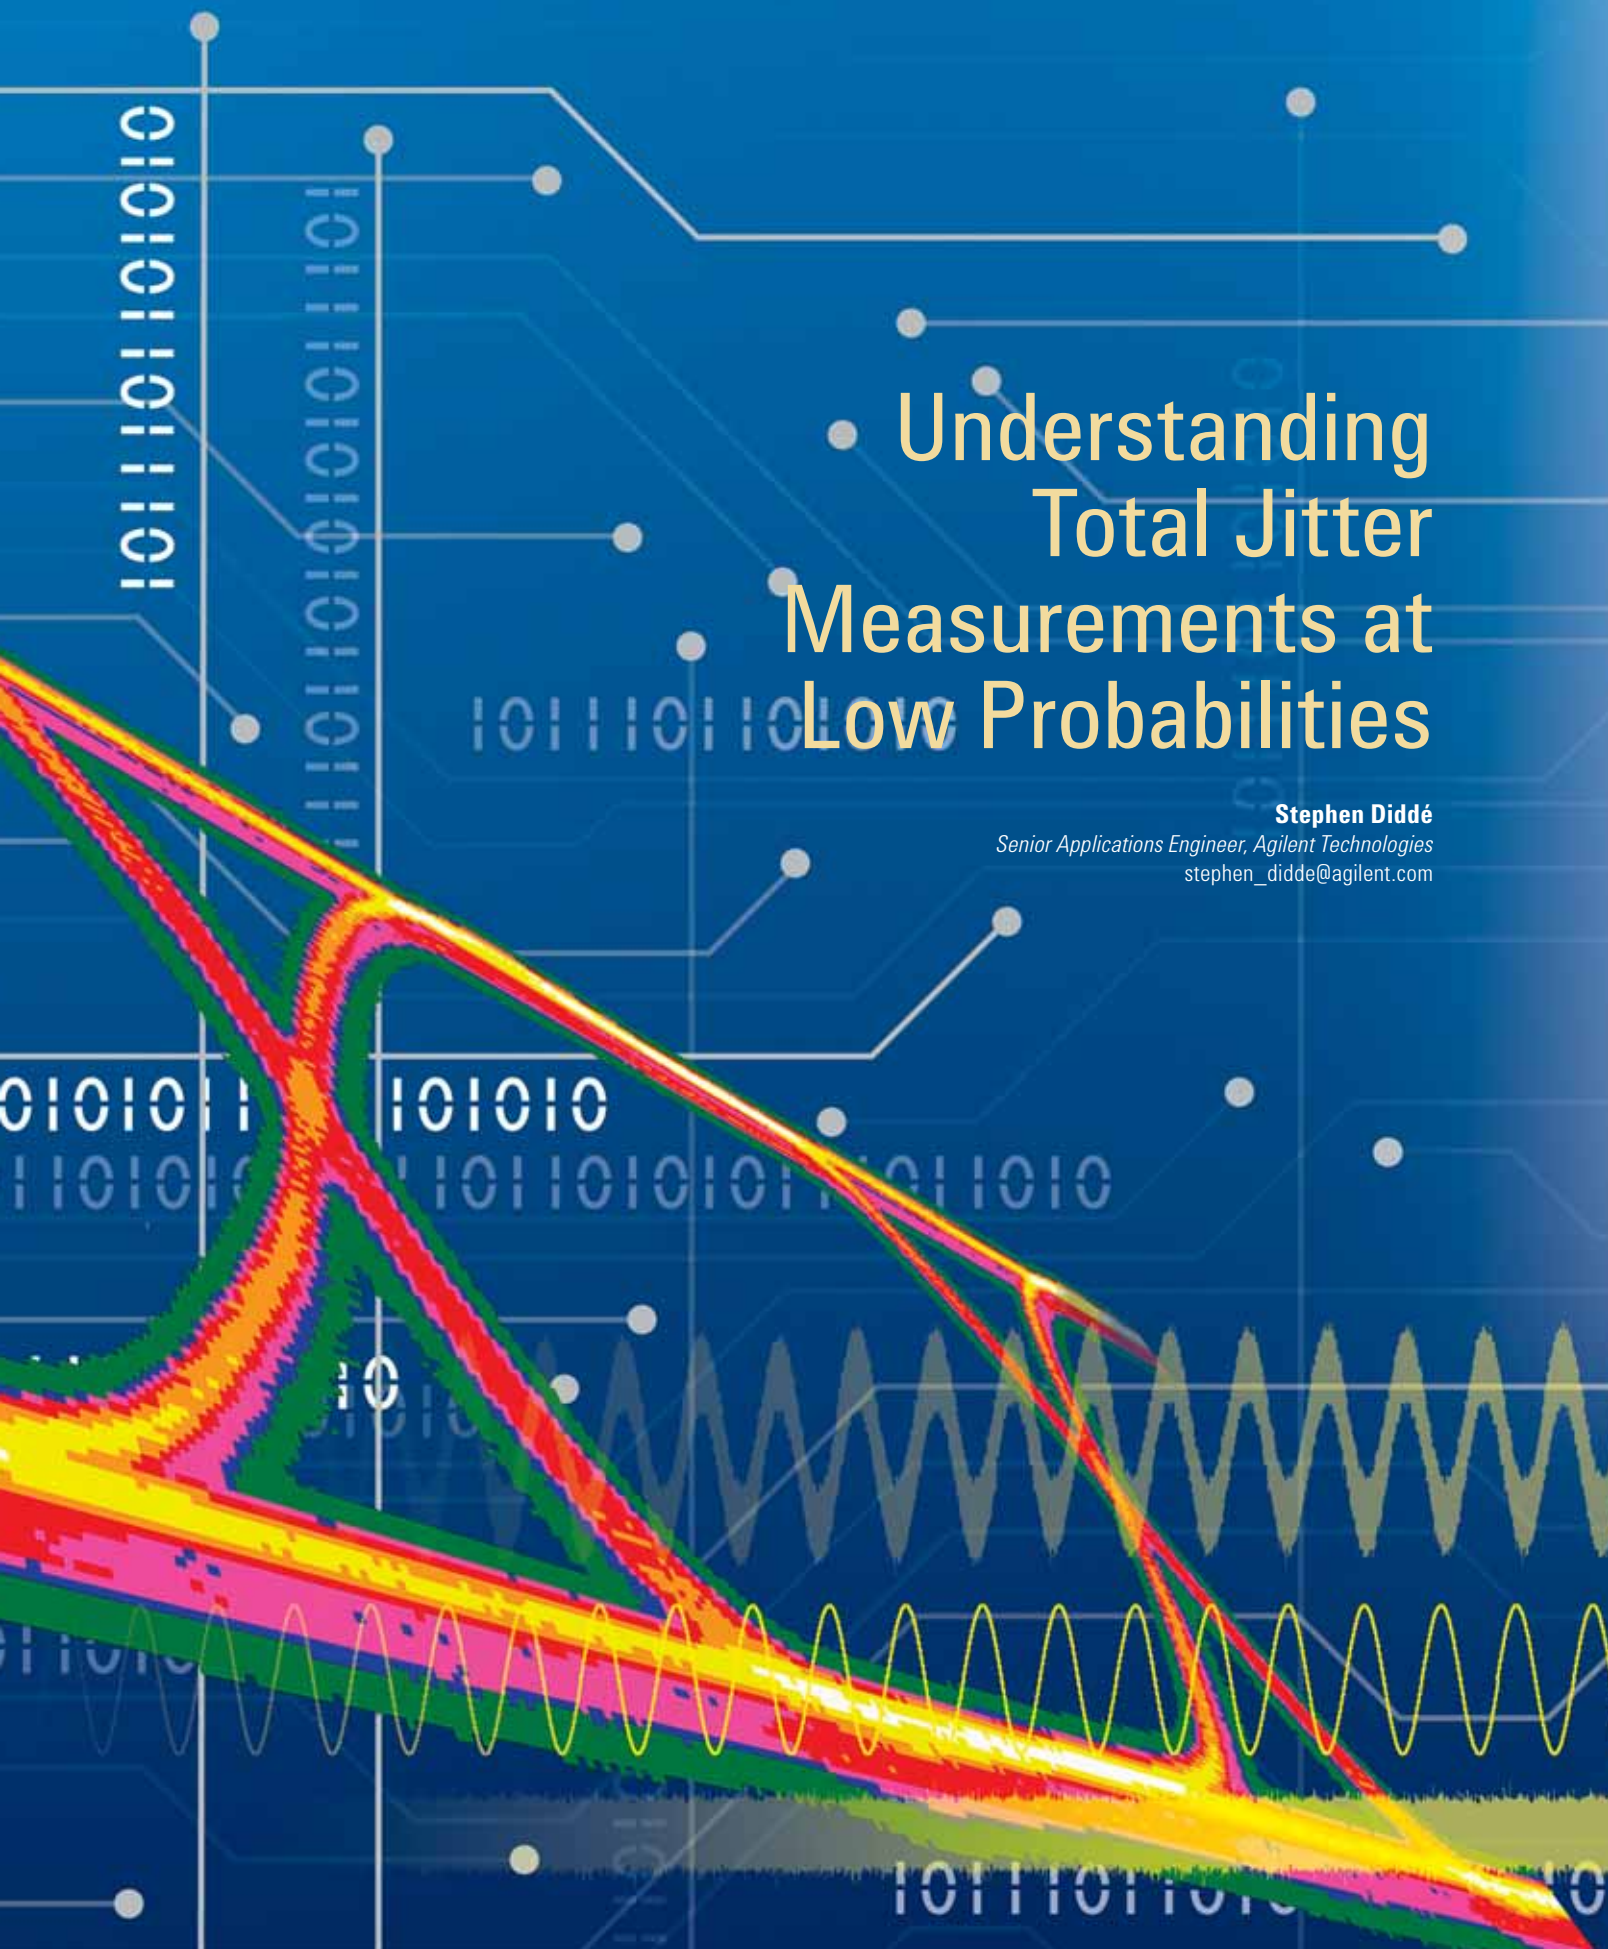


Figure 3. In protocol measurements, adjusting the logic analyzer sampling position to the center of the eye opening will maximize sampling accuracy.

Understanding Total Jitter Measurements at Low Probabilities

Stephen Diddé

Senior Applications Engineer, Agilent Technologies
stephen_didde@agilent.com



T

Total jitter (TJ) is defined by many high-speed, serial-data standards to ensure relatively low bit error rates (BERs) in the presence of timing misalignment. As the data rates are increased, this parameter becomes more critical. For a high-speed serial data link operating at rates beyond 2.5 Gb/s, the required error rate cannot exceed 10^{-12} . Timing misalignment, or jitter, contributes more to the overall BER than any other effect in a typical high-speed serial link and is caused by limited-channel bandwidth, system noise, phase noise, crosstalk, and power supply leakage — among other things (see sidebar, *Taxonomy of Jitter*). Total jitter, the sum of all the timing errors of the signal, can be measured with either an oscilloscope or bit error rate tester (BERT) by observing the timing error of the crossover point or the BER.

The BER measurement is based on the statistical properties of the underlying physical mechanisms. As the sample size is increased, the estimate of the parameter is improved. A confidence level can be assigned to this estimate, especially when the underlying mechanism is stochastic. Total jitter is derived by making multiple BER measurements while adjusting the sampling point. By defining the upper and lower limits on the confidence level, it is possible to identify the range of values for TJ. This article explains this bracketing approach in detail and outlines the search algorithms used to find the range and the confidence level for the estimate.

Representing jitter

Jitter quantifies the allowed uncertainty of the sampling instance within the timing window of the bit. It is formally defined as the expected deviation of a signal's timing event from its intended or ideal occurrence in time. Ideal occurrences are marked by reference-clock sources that specify when the ideal events should occur. In general, this ideal clock is realized by either a master network clock or a clock that is recovered from the data stream, depending on the communication scheme.

Mathematically, jitter can be represented in either the analog or digital domain. In analog communications, jitter is also known as phase noise and is defined as a phase offset that continuously changes the timing of a signal:

Equation 1.

$$S(t) = P[t + \phi(t)]$$

Where $S(t)$ is the jittered signal waveform, $P[t]$ is the undistorted waveform, and $\phi(t)$ is the phase offset, or phase noise. This definition is most useful in the analysis of analog waveforms such as clock signals and is frequently used to express the quality of oscillators.




Figure 1. Total jitter is the sum of the individual components, PJ and RJ, in this example

If the total jitter is denoted by $J(t)$ and the periodic, random and data-dependent jitters are denoted as $PJ(t)$, $RJ(t)$ and $DDJ(t)$, respectively, then the total jitter can be defined as:

Equation 2.

$$J(t) = PJ(t) + RJ(t) + DDJ(t) + \dots$$

Here, the waveform in time of the total jitter time is the sum of the individual components. As an example, consider a periodic jitter with a frequency of 10 MHz and an amplitude of 10 ps that is added to a random jitter with a standard deviation of 1.5 ps. This data is simulated, with Gaussian statistics for the random jitter signal and a pure sine wave for the periodic jitter. The results are shown in Figure 1.

To assemble the timing-jitter budget for a design, we must obtain the total jitter at a specified time to yield a single number and hence, a single figure of merit. Here a peak-to-peak value (e.g., the difference between the maximum value and the minimum value) is typically used. It is defined as:

Equation 3.

$$TJ_{pp} = J(t)_{max} - J(t)_{min}$$

The total peak-to-peak jitter for the example in Figure 1 is about 31 ps, but this result is not useful because the RJ term describes an unbounded random process. As a consequence, the observed minimum and maximum values become larger without limit — and consequently so does the TJ_{pp} . The usual way to deal with this problem is to use the individual terms. We can build histograms or calculate the probability density function (PDF) for the individual jitter components and then use a convolution to calculate the PDF:

Equation 4.

$$J(x) = PJ(x) \times RJ(x) \times DDJ(x) \times \dots$$

The TJ_{pp} value is then the maximum non-zero probability PDF value minus the minimum non-zero PDF value. Figure 2 shows the PDFs for our example: The TJ_{pp} is 31 ps, exactly the same value that was obtained from the time-domain waveform.




Figure 2. The total jitter PDF is the convolution of the individual component's PDFs

The PDF has two advantages over the time-domain waveform. First, it can be measured directly on many different types of test equipment such as sampling oscilloscopes, real-time oscilloscopes and time-interval analyzers. Secondly, the PDF of a Gaussian process is well known.

If we know the RJ RMS value and the PDFs of all the other jitter components, it is possible to calculate the PDF of the total jitter. Expressing TJ_{pp} as a function of a probability level can be done once we construct a cumulative probability distribution function (CDF) by integrating the PDF:

Equation 5.

$$CDF(t) = \int_{-\infty}^t PDF(x).dx$$

At each point in time, the CDF gives the probability that the transition happened earlier. TJ_{pp} for a probability level of γ is then the time value where $CDF=1-\gamma/2$, minus the time value where $CDF=\gamma/2$. Figure 3 shows the TJ CDF in the example in Figure 1.

The CDF is directly related to the BER, which is caused only by timing jitter. However, BER as measured by test equipment includes the effects of timing jitter as well as amplitude noise. Usually, the bit is sampled at the center of its period, where the signal is likely to have attained its maximum signal-to-noise power ratio. At this sampling instant, the contribution from timing jitter is usually small, compared to the contribution from amplitude noise. When the sampling instant deviates from this ideal location, due to timing misalignment, the contribution from timing jitter grows. As the sampling edge is brought closer to the edge of the bit, the contribution from timing jitter often dominates that from the amplitude noise.

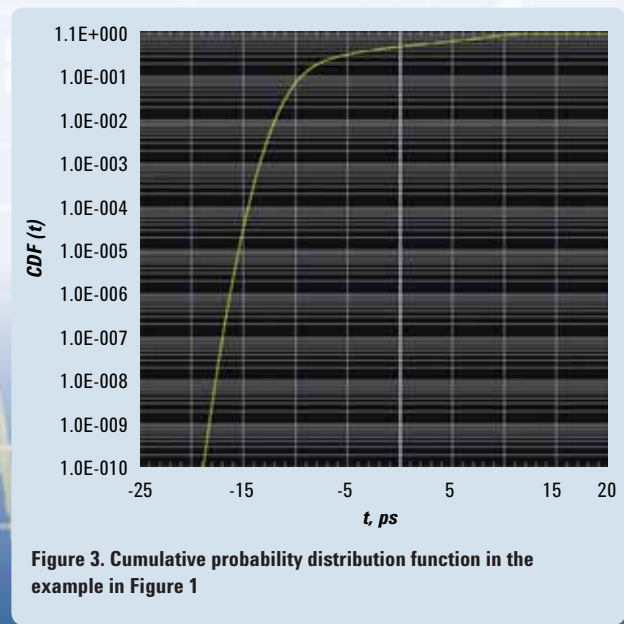


Figure 3. Cumulative probability distribution function in the example in Figure 1

Understanding bathtub curves

BERTSCAN is a technique first formalized by the ANSI for Fibre Channel and later adopted into the T11.2 Jitter standard.^{2,3} It defines a method for estimating total jitter by measuring the BER. The T11.2 document, treated as a reference standard for many datacom standards, models the timing error using a PDF. This timing error gives rise to BER, when the bits are sampled at instants away from the ideal sampling point, as illustrated in Figure 4.

From a qualitative perspective, the BER measured with the ideal sampling point yields the best possible value. The value worsens as the sampling point is moved across the eye in either direction, towards the eye crossing points (e.g., the left crossing point or the right crossing point in Figure 4). If the BER is plotted as a function of sampling time (or sampling delay offset), the resulting plot resembles a bathtub. The minimum BER occurs at

the optimum sample point, which is usually near the center of the eye. The maximum BER occurs at the left and right crossing points. Figure 5 depicts a typical bathtub plot for the case of large levels of random jitter. The total jitter is defined as the difference between the eye opening at 10^{-12} BER and the bit period.

In practice, generating this plot requires measurement of BER levels to very low probabilities. As this is a very time-consuming process, the usual practice is to measure only a part of the bathtub curve, usually down to a BER of 10^{-9} , and extrapolate to the required levels (e.g., 10^{-12}). The extrapolation process assumes certain models and the accuracy depends on the validity of these assumptions and models. For example, the T11.2 document suggests that the $RJ(x)$ in Equation 4 can be modeled as a Gaussian process. This enables the extrapolation of lower BER values from the data collected at higher BER levels.




Figure 4. Eye diagram showing the effect of sampling on the measured BER

Measuring BER

From the above description, it is clear that jitter measurements can be derived by making accurate BER measurements. Conceptually, the BER measurement is modeled by a binomial process. The error detector in a BERT detects the incoming bits (ones and zeros) and compares them with the expected bits. The errors occur as a result of stochastic noise processes, as well as deterministic effects in the link. The BERT is only able to estimate the probability of error, P_e , and asymptotically approaches the correct value as the sample size is increased.

Theoretical analysis

The fact that the bits transmitted are either ones or zeros allows us to model the process analogous to the tossing of a coin, which also has only two outcomes. If p is the probability of finding heads when a coin is tossed, the probability of finding the heads k times when the coin is tossed n times is given by the well-known binomial probability function,⁴

Equation 6.

$$P(k; n, p) = \binom{n}{k} p^k (1-p)^{n-k}$$

Extending this analogy to the BERT case, and noting that p corresponds to P_e (the probability of error) and k is the number of errors when n bits are received, we get:

Equation 7.

$$P(k; n, p) = \binom{n}{k} P_e^k (1-P_e)^{n-k}$$

Equation 7 provides likelihood of finding k errors when n bits are received and when the system has an underlying probability of error, P_e . The assumption made here is that each error event is independent of the next. Our goal is to find an accurate estimate of P_e in a reasonable amount of time. The estimate of P_e , as measured by the BERT, is given by:

Equation 8.

$$BER = k/n$$

In the actual measurement of bit errors, the estimate for the true P_e is exact only when the sample size or the number of observed bits n approaches infinity.

Equation 9.

$$BER \xrightarrow{n \rightarrow \infty} P_e$$

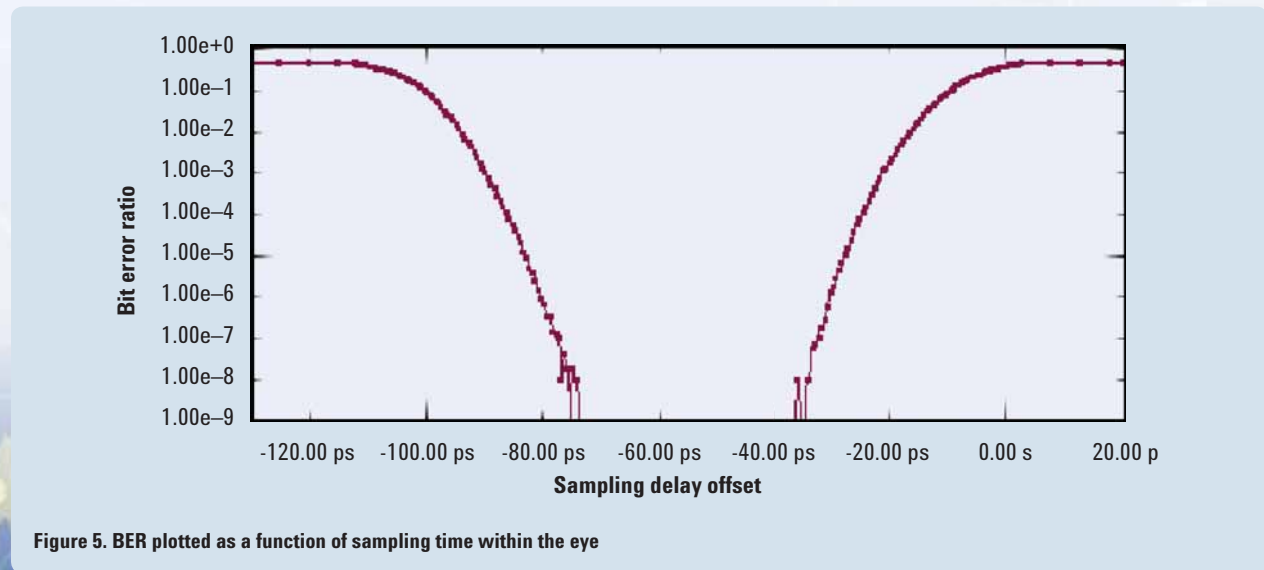


Figure 5. BER plotted as a function of sampling time within the eye

Since it is impossible to wait for an infinite number of bits, it is desirable to collect as many samples as possible. If we measure a large number of errors, say 1,000, then the estimate is much more accurate. However, this poses a serious constraint on the required test time. At a 2.5-Gb/s PCI Express data rate, for example, the time required to generate 1,000 errors could easily reach 400,000 seconds, which is over 111 days!

Returning to Equation 7, we observe that the value of P_e is much smaller than 1, typically in the range of 10^{-3} to 10^{-12} . When P_e is small, Equation 7 can be approximated to a Poisson distribution:⁴

Equation 10a.

$$P(k, \mu) = \frac{1}{k!} \cdot \mu^k e^{-\mu}$$

Equation 10b.

$$\text{where } \mu = n \cdot P_e$$

This expression is much more useful for analytical purposes.

In order to minimize the measurement time, we only want to measure a minimum number of bits that still guarantees a certain confidence level. If we desire a higher confidence level the measurement time is accordingly higher. This confidence level defines a probability measure for the given estimate. Specifically, the confidence level is the probability that P_e is less than a specified threshold, γ , given that k errors are observed when n bits were received.

Equation 11.

$$CL = \text{Prob}[P_e < \gamma \mid k, n]$$

Because this confidence level gives the probability that P_e is less than γ , we can define it as the lower limit of confidence, CL . We

can also define another useful quantity, the upper limit of the confidence level which defines the probability that P_e is greater than γ (usually referred to as the target BER).

Equation 12.

$$CL_u = \text{Prob}[P_e > \gamma \mid k, n]$$

Equation 10a gives the probability of finding k errors. Using this expression we can derive the probability of finding at most k errors when we vary the number of errors from zero to k . By summing all the probabilities we obtain the total probability of finding k or fewer errors.

Equation 13.

$$P[\text{errors} \leq k] = \sum_{i=0}^k P(i, \mu)$$

This quantity defines the confidence level for finding no greater than k errors.⁵ This expression is not very useful because it gives the confidence level if P_e is known. If we replace P_e with γ , then the resulting confidence level can be used as a guideline.

Equation 14a.

$$P[\text{errors} \leq k \mid \gamma] = \sum_{i=0}^k P(i, \nu),$$

Equation 14b.

$$\nu = n \cdot \gamma$$

Let's say we compute the confidence level for a given γ , n and k from Equation 14a to be high. And in the experiment, we measure a higher k . This means one of two things. Either we got lucky in picking a very unlikely event or the error mechanism is worse than anticipated. If we repeat the experiment and consistently measure higher value for γ , then we conclude that the latter is indeed true. This implies that the true value, P_e , is greater

than γ with a probability given by Equation 14a. As indicated by Equation 12, this is the confidence level for the upper limit.

Equation 15.

$$CL_u = \sum_{i=0}^k P(i, \nu) = \sum_{i=0}^k \frac{1}{i!} \cdot \nu^i e^{-\nu}$$

As an example, let's consider a target BER, γ , of 10^{-8} . If n is 10^9 then we expect to see 10 errors for k . But in the actual experiment, we observe 15 errors. CL_u then works out to 0.95 or 95 percent from Equation 14a. This is obtained by using 15 for k , 10^9 for n and 10^{-8} for γ . From this we can infer that P_e is higher than 10^{-8} with a 95 percent confidence level.

Its complement gives the probability that the errors are k or greater. The confidence level, CL_l , is then given by,

Equation 16.

$$CL_l = 1 - CL_u = 1 - \sum_{i=0}^k \frac{1}{i!} \cdot \nu^i e^{-\nu}$$

The confidence level here defines the probability that P_e is less than γ . From Equation 16 it is possible to find the minimum number of bits that are required for the estimate of the probability of error as a function of the confidence level, CL_l . When there are no errors received in an observation window, during which n bits are analyzed or when $k=0$, Equation 16 reduces to:

Equation 17.

$$CL_l = 1 - e^{-\nu}$$


Rearranging terms, and using Equation 14b, we get:

Equation 18.

$$n \cdot \gamma = -1n(1 - CL_l)$$

From Equation 18, the minimum n for a given CL_l can be calculated. When the confidence level is specified, a lower limit can be defined on the number of bits for a given target BER, γ . For example, if we want to verify that a communication link has a BER of less than 10^{-9} , and the confidence level desired is 95 percent, then a minimum n of approximately 3×10^9 (or $(-\ln(1-0.95)/10^{-9})$) bits are needed. In other words, when we receive 3×10^9 bits and there are no errors observed, then P_e is less than 10^{-9} with 95 percent confidence.

When k is non zero, the solution to Equation 16, as well as Equation 15, is difficult to obtain even with Poisson approximation. Previous work in this area attempts to solve it by further approximating the Poisson distribution to a Gaussian distribution.⁶ This turns out to be not so efficient, as it requires much higher n than would be required without the approximation.



Another method used to solve Equations 15 and 16 is the Poisson regression.⁷ While this method provides better results than the Gaussian approximation, the most accurate results come from empirically evaluating Equations 15 and 16.

In order to find the required confidence levels, the values of k and n are varied systematically and for a given γ , Equations 15 and 16 are computed. Figure 6 is generated by picking those values of k and n for which CL_U or CL_L is one of the three values, 90, 95 or 99 percent. From this graph it is clear that the measured BER approaches the target BER, γ , asymptotically from either direction.

If we examine the red curve corresponding to the 95 percent level, in Figure 6, we notice that at $n = 3 \times 10^9$, the observed value of k is zero, giving the measured BER of 0. But from the graph we can interpret this to mean that the actual probability of error P_e is less than the target value of 10^{-8} with a confidence level of 95 percent. The same result can be obtained with the help of Equation 18.

Finding total jitter

Using the BERTSCAN technique to find 10^{-12} P_e points, and therefore TJ, would be very time consuming. To increase the speed of measurement, without any approximations, we can use the concepts developed in the previous section. Since we are purely interested in the TJ result, it is sufficient to find the sample delay offsets on the left and right slope of the bathtub curve where the P_e is exactly 10^{-12} . We call these points x_L and x_R . The peak-to-peak total jitter, TJ_{pp} , is then simply the bit interval minus the difference between x_L and x_R .

Unfortunately, since the BERT has a finite delay resolution, it is virtually impossible to locate the sampling delay that corresponds exactly to these points. Even if there were sufficient timing resolution, an infinite number of bits would still need to be observed to prove that the P_e is exactly 10^{-12} .

Bracketing approach

Because we are unable to locate a single point on the slope where the P_e is exactly 10^{-12} , we instead target an interval that brackets it. The point where P_e is equal to 10^{-12} lies within this interval with a high confidence level. Figure 7 depicts this process for the right slope. We search for an interval $[x-, x+]$ that brackets the x_L point, where $x+$ and $x-$ are separated by no more than the desired delay-step resolution of the TJ measurement, Δx .

We don't need to know the exact P_e values at $x+$ and $x-$. It is sufficient to assert that $P_e(x-)$ is greater than 10^{-12} and $P_e(x+)$ is less than 10^{-12} at a desired confidence level. If we choose 95 percent, we have determined that x_L is within the interval $[x-, x+]$ with a confidence level better than 90 percent. For lack of better knowledge, we assume that x_L is in the middle of the bracketing interval. Now the distance between $x-$ and $x+$ is Δx , and hence x_L is only accurate to $\pm 0.5 \Delta x$. Replicating the same procedure for the right half of the bathtub curve yields x_R with the same accuracy. TJ_{pp} can then be calculated as before with an accuracy of $\pm \Delta x$.

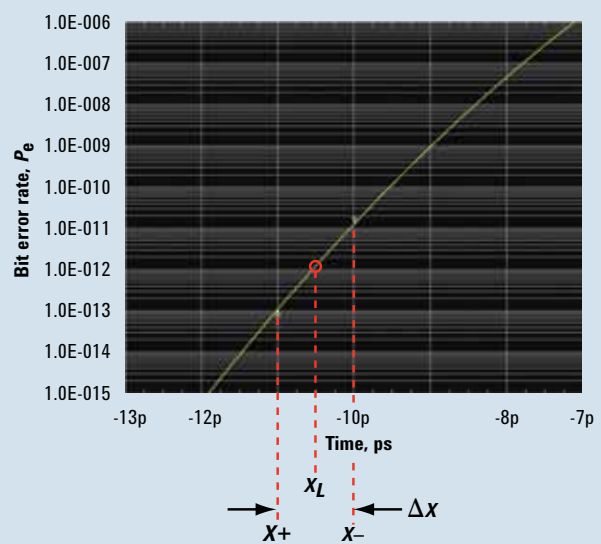


Figure 7. Definitions for the bracketing approach, on the right slope lower BER region

Strategies for searching

There are many possible algorithms to search for the interval $[x-, x+]$. There are several items to consider during the implementation of a fast TJ measurement using the bracketing approach:

- Because measurement times increase with decreasing P_e , the search should be performed from left to right for the left edge, and from right to left for the right edge.
- The main goal of the search algorithm is to minimize the number of failed attempts to find $x+$. At 10 Gb/s it takes about five minutes to compare $3e^{12}$ bits, the longest measurement time during each iteration.
- Once $x-$ has been determined, $x+$ is often within Δx because of resolution limitations.
- The search can be optimized if it is realized as an iterative process. The resolution can be fine tuned continuously from a coarse step to a fine step.
- The measurement times can be optimized for a given measurement accuracy.
- To get better initial values for the search, it is a good idea to perform a relatively fast complete bathtub scan. From the data, we can get reasonable first guesses for $x-$ and $x+$, either directly or by fitting an inverse error function to the data.
- If the device under test has a P_e floor, the search may be stuck since the P_e never gets below 10^{-12} . This can be accounted for with a robust implementation.

Specific implementation of the search algorithm was carried out using a linear search algorithm for various combinations of RJ and DDJ.⁸ Average measurement times for this implementation

are given in Tables 1 and 2. Note that measurements times with the bracketing approach are not strictly repeatable, since at lower P_e values the time until the first error is observed is randomly distributed.

Table 1. Measurement times for 10 Gb/s signals; time resolution set to 1 ps

	DJ = 0	DJ = 15 ps	DJ = 30 ps
RJ = 0 ps	–	10 min	10 min
RJ = 2 ps	16 min	12 min	10 min
RJ = 4 ps	17 min	16 min	12 min
RJ = 6 ps	17 min	26 min	–

Table 2. Measurement times for 10 Gb/s signals; time resolution set to 5 ps

	DJ = 0	DJ = 15 ps	DJ = 30 ps
RJ = 0 ps	–	10 min	10 min
RJ = 2 ps	10 min	10 min	10 min
RJ = 4 ps	10 min	14 min	10 min
RJ = 6 ps	11 min	8 min	–

The data in the tables above are obtained through simulations, which were carried out using the average time between errors.⁸ When the RJ values were low, the measurement was completed in 10 minutes, independent of DJ. This is because the slope of the bathtub curve is directly related to the value of RJ. In the low P_e region, it takes only one data point to bracket the 10^{-12} point. The measurement time is then dominated by the time required to compare $2.996e^{12}$ bits (five minutes), once per slope. This is independent of the delay resolution used. The minimum test time at 10 Gb/s is always 10 minutes, no matter how coarse the resolution.

For increasing RJ values, measurement time goes up because more points are located on the slope of the bathtub curve. The sawtooth shape in this region is really an indication of the random variability of the measurement time. It entirely depends on how many points are located on the slope and where. The lower resolution setting hits fewer points on the slope, so the measurement completes earlier with decreasing resolution.

From Tables 1 and 2, we see that an average measurement time of about 15 to 20 minutes is achieved, at 10 Gb/s and with a 1-ps delay-step resolution. Making the same measurement without using the bracketing approach and with 1-ps resolution generates a plot similar to that of Figure 5. It requires about 41.67 hours. The bracketing approach therefore reduces measurement times by a factor of about 40 (and as much as 100), depending on RJ and DJ values.

Conclusion

BER measurement is critical to making an accurate peak-to-peak total jitter measurement. A BERT can be used to make measurements with high confidence levels, while the bracketing approach can be used to control the test time based on the required confidence level in the measurement. Experimental data produced using this approach shows a 40x to 100x improvement in measurement time compared to a conservative bathtub measurement. For a TJ measurement that was done at the 10^{-12} BER level, with a confidence level of better than 90 percent, typical test times are approximately 20 minutes at 10 Gb/s, and a little more than one hour at 2.5 Gb/s.

Due to this direct measurement approach, accuracy of the results is independent of the TJ PDF. Consequently, the bracketing approach presents a significant advantage over other methods based on oscilloscopes or time-interval analyzers, which fail if the jitter distribution doesn't fit the extrapolation model.

References

1. Stephens, R. 2004. Analyzing Jitter at High Data Rate. *IEEE Optical Communications*, February 2004.
2. *Jitter test methodology for Fibre Channel*. ANSI X3.230-1994 Annex A, sub clause A.4.2.
3. Working draft, Fibre Channel—Methodologies for Jitter Specification-MJSQ. 2003, Rev 10. *National Committee for Information Technology Standardization (NCITS)*, 2003.
4. Papoulis, A. 1984. *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill, New York.
5. Reed, J. 2000. Calculating Statistical Confidence Levels for Error-Probability Estimates. *Lightwave Magazine*, April 2000: 110–114.
6. Jeruchim, M. C. Techniques for Estimating the Bit Error Rate in the Simulation of Digital Communication Systems, *IEEE J. Select. Areas Commun.*, vol. SAC-2, no. 1, 1984: 153–170.
7. Macdonald, P.D.M. *Foundations of Statistics - Confidence Intervals for a Poisson Mean*, available online at: www.math.mcmaster.ca/peter/s743/poissonalpha.html.
8. Mueller, M., Stephens, R., and McHugh, R. 2005. Total Jitter Measurement at Low Probability Levels, using Optimized BERT Scan Method. *Designcon 2005*, Worcester, Mass., September 19–21, 2005: TA-4.

Taxonomy of Jitter

Jitter is dependent on the type of communication system and can be caused by a number of factors such as crosstalk and noise. In system design, it is important to understand what specific mechanism is causing timing errors.

As shown in Figure 1, jitter is classified according to generic categories, the most common of which are bounded and unbounded, correlated and uncorrelated, and data-dependant and non-data-dependant, random and deterministic, periodic and non-periodic.¹ Total jitter is broadly divided into deterministic jitter and random jitter. Random jitter is due to noise effects that alter the bit arrival times. This jitter is unbounded, in the sense that the expected jitter grows with time. Deterministic jitter, on the other hand, is bounded and is produced by data-pattern dependencies and crosstalk from other signals.

Data-dependent jitter is subdivided into intersymbol interference (ISI) and duty-cycle distortion (DCD). ISI is caused by dispersion, while DCD is generated by the non-uniform response to the rise and fall times of the data signal. Crosstalk and periodic jitter are the result of interference and AM-PM conversion due to nonlinearity at the transmitter, the medium and the receiver.

References

1. Stephens, R. 2004. Analyzing Jitter at High Data Rate. *IEEE Optical Communications*, February 2204.

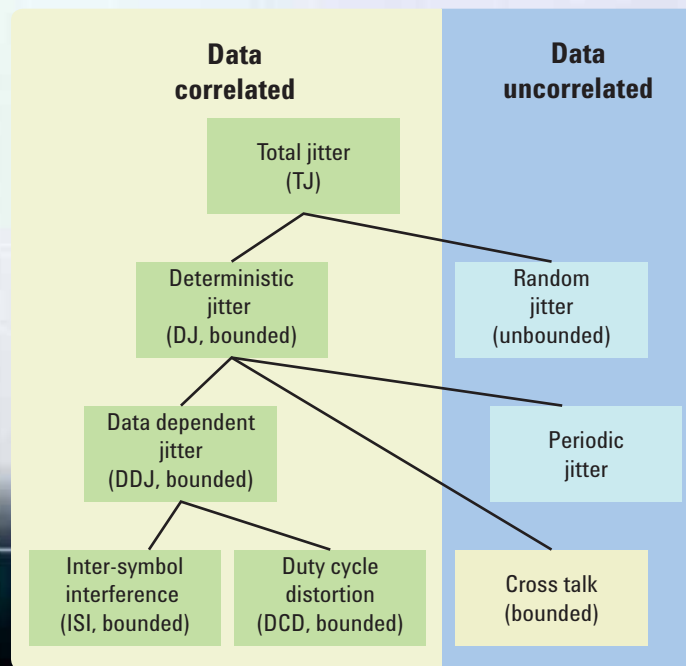
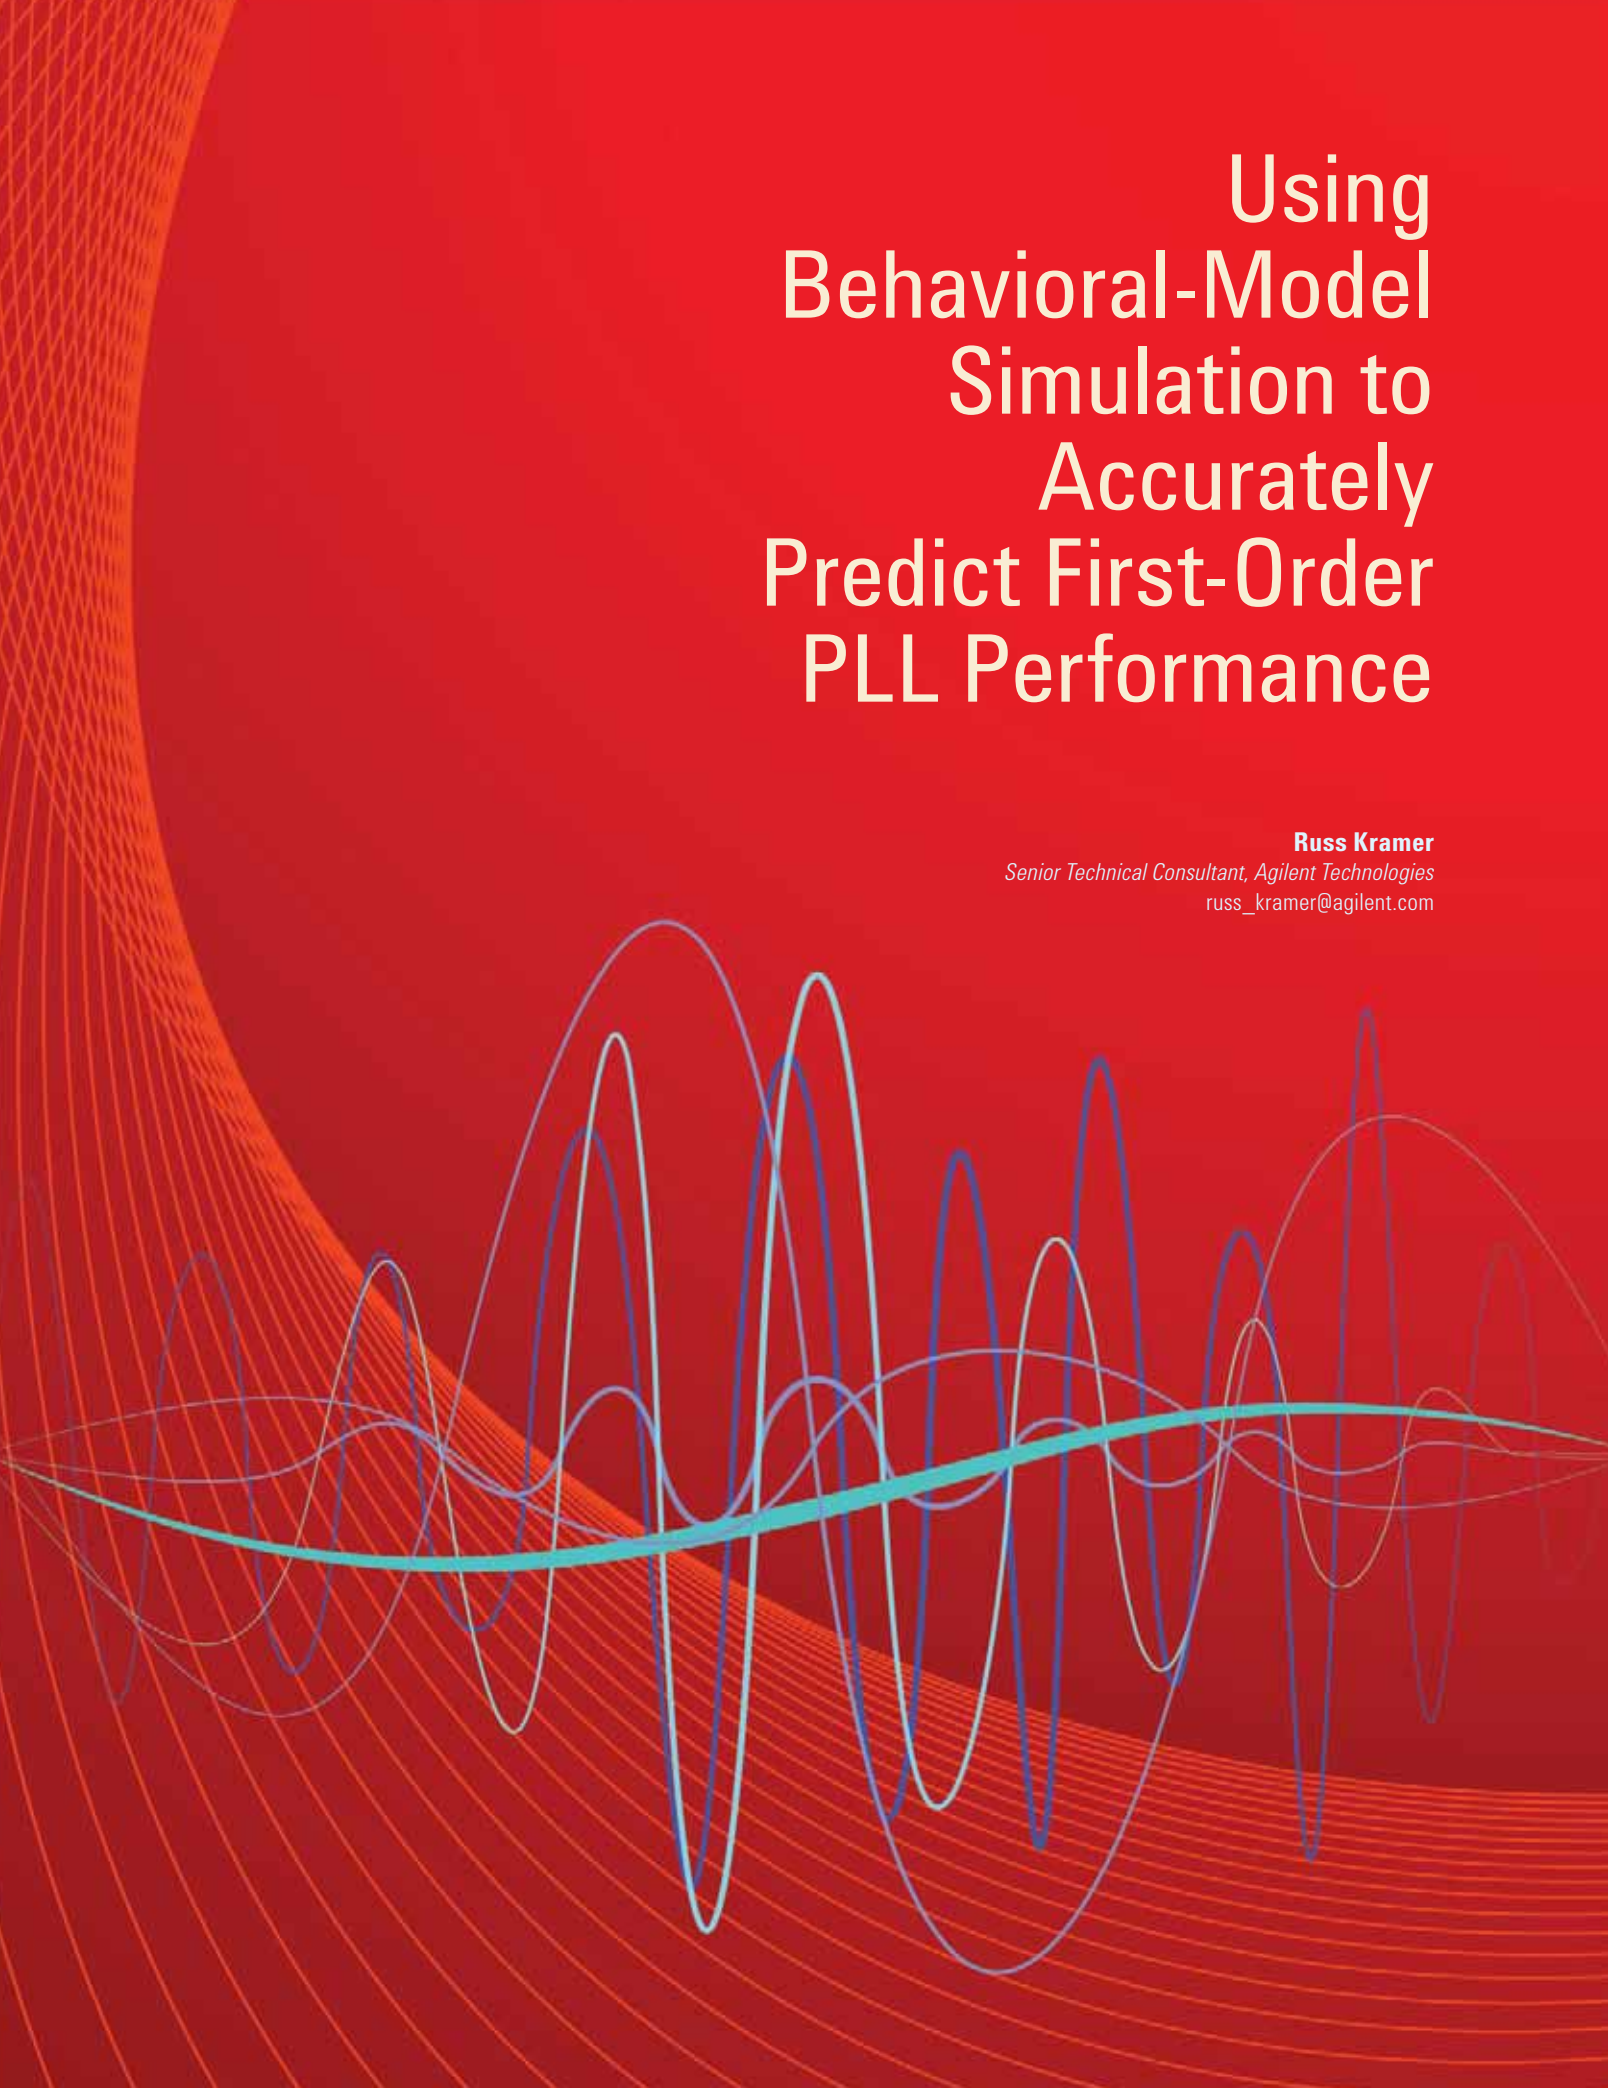


Figure 1. Classification of jitter components



Using Behavioral-Model Simulation to Accurately Predict First-Order PLL Performance

Russ Kramer

Senior Technical Consultant, Agilent Technologies

russ_kramer@agilent.com

P

Phase-locked loops (PLLs) are widely used as on-chip clock generators to synthesize and reshape a high-frequency internal signal that is derived from a lower-frequency external signal. In data communications, for example, PLLs are used as clock-recovery systems. In broadband optical communication networks, they serve as clock-and-data recovery (CDR) systems that generate the clock and resynchronize the data from the received signal. PLLs are also used as frequency synthesizers in wireless communications to synthesize an accurate frequency that can then be used to modulate or demodulate the incoming signals.

The random temporal variation of phase in the synthesized frequency is an essential measure of PLL performance. This “phase noise” or “jitter” is an undesired variation in the timing of events at the output of the PLL — and it is difficult to predict with the small-signal analysis capabilities of traditional circuit simulators. Because a PLL generates repetitive switching events as an essential part of its operation, the noise performance must be evaluated in the presence of this large-signal behavior.

Some design tools are well-suited to the simulation and characterization of noise performance in small circuits such as voltage-controlled oscillators (VCOs). However, a closed-loop PLL synthesizer may have a device gate count of 30,000 to 50,000, making it hard for some simulators to converge on a result. In contrast, the Agilent RF Design Environment (RFDE) and Advanced Design System (ADS) tools feature the capability to characterize noise performance and create behavioral models for every subcircuit of a PLL synthesizer. The phase noise or jitter performance of each subcircuit can be simulated separately with RFDE while noise parameters can be extracted for behavior modeling. Using the Agilent EEsof circuit-envelope capability, a closed-loop simulation can then be performed in a matter of several seconds.

Understanding PLL noise

Unlike a single tone in a frequency spectrum, the synthesized output frequency of a PLL is subject to all sorts of noise, occurring in both amplitude and phase. Common sources such as flick (1/f), thermal and shot noise are associated with devices. Comprehensive noise models also include additional sources such as power/ground, substrate noise coupling, signal intermodulation, and sub-optimum biasing.

Due to the nonlinear nature of oscillator circuits, amplitude fluctuation is inherently limited and phase variation is of key importance. Variations in phase — both short-term nonrandom and long-term — are usually due to an external reference source and discrete spurious signals, which can be removed or suppressed to a reasonable level using appropriate techniques. Short-term random variation in phase, which is the phase noise or jitter mentioned earlier, directly impacts the accuracy and stability of PLL performance. If the ideal output signal of an oscillator is sinusoidal (e.g., phasor), then the noise is a small perturbation added to this trajectory as shown:

Equation 1.

$$S(t) = [A_1 + n_1(t)]\cos[\omega_0 t + \theta(t)]$$

In the equation, A_1 is the constant amplitude, $n_1(t)$ is the amplitude (AM) noise (negligible in a well-designed oscillator), ω_0 is the waveform center frequency, and $\theta(t)$ the waveform phase perturbation.

Modulation theory shows that phase noise can be represented as a sideband with symmetrically smaller amplitude on both sides of the carrier frequency. Intuitively, phase noise in the frequency domain can be viewed as the cycle-to-cycle jitter in the time domain which changes the instantaneous zero-crossing of an otherwise perfect sinusoidal signal.

There are five significant sources of phase noise within a PLL synthesizer:

- VCO phase noise
- Reference-oscillator phase noise
- Thermal noise and device noise from components in the loop filter
- Noise from the digital dividers and phase detector
- Noise injected by the supplies and bias circuits

The first three sources are well understood and can usually be accurately modeled using measured phase-noise data for the VCO and reference, and conventional noise models from circuit theory for the loop filter. In contrast, noise from digital devices, dividers and phase/frequency detectors is difficult to model and constrains top-level simulation of all analog and digital functional blocks.

Modeling phase noise

Over the years, researchers have studied many ways to characterize the phase noise in electrical oscillators. Leeson’s model, published in 1966, was the first attempt to predict phase noise in oscillators.¹ It is described with the following equation:

Equation 2.

$$\ell(f) = \frac{2FkT}{P_s} \left(\frac{f_{osc}}{2Qf} \right)^2$$

F is the noise factor of the oscillator’s gain element, k is the Boltzmann constant, T is the absolute temperature, P_s is the oscillator signal power, f_{osc} is the carrier frequency, Q is the quality factor (formed from the resonating inductance (L) and capacitance (C) portion of the oscillator), and f is the offset frequency from the carrier.

These days, highly detailed models are being used to characterize oscillator performance. For example, Razavi uses a linear time-invariant (LTI) model to describe the behavior of phase noise in oscillators, while Hajimiri relies on a more accurate linear time-variant (LTV) model.^{2,3} Demir derived a nonlinear stochastic differential equation for phase error, and solves this equation in the presence of random perturbations.⁴

To better understand PLL noise, consider the generic PLL shown in Figure 1. It includes five key elements: phase frequency detector (PFD), charge pump (CP), loop filter (LF), VCO, and frequency divider (FD). The noise transfer function of the closed-loop PLL is derived as follows:

Equation 3.

$$\frac{\Theta_{out}}{\Theta_i} = \frac{NK_pK_{vco}F(s)}{Ns + K_pK_{vco}F(s)} = \frac{NG_{OL}(s)}{1 + G_{OL}(s)}$$

Equation 4.

$$\frac{\Theta_{out}}{\Theta_{PFD}} = F(s) \frac{K_{vco}}{s} \left(\frac{1}{1 + G_{OL}(s)} \right)$$

Equation 5.

$$\frac{\Theta_{out}}{\Theta_{LF}} = \frac{K_{vco}}{s} \left(\frac{1}{1 + G_{OL}(s)} \right)$$

Equation 6.

$$\frac{\Theta_{out}}{\Theta_{VCO}} = \left(\frac{1}{1 + G_{OL}(s)} \right)$$

In this derivation, Θ_i , Θ_{out} , Θ_{PFD} , Θ_{LF} , and Θ_{VCO} represent the noise signal at different stages; N is the divide ratio from the PLL divider circuit; K_p and K_{vco} are gain attributed to the PFD and VCO, respectively; $G_{OL}(s)$ is the open-loop gain; and s is time in seconds. Figure 2 shows the typical noise contribution from individual blocks in the example PLL. In the figure, P refers to the charge pump stage of the PLL and τ is the period of one clock cycle.

Note that within the PLL loop bandwidth (ω_c) the phase noise is typically dominated by contributions from the frequency dividers (blue) and phase detector (green). For frequencies well below the loop bandwidth, the phase-noise plot typically flattens out due to the cumulative noise being dominated by the phase frequency detector; thus, the resultant in-band noise is essentially flat below the loop bandwidth. Outside of the loop bandwidth the major noise contributor is the VCO.

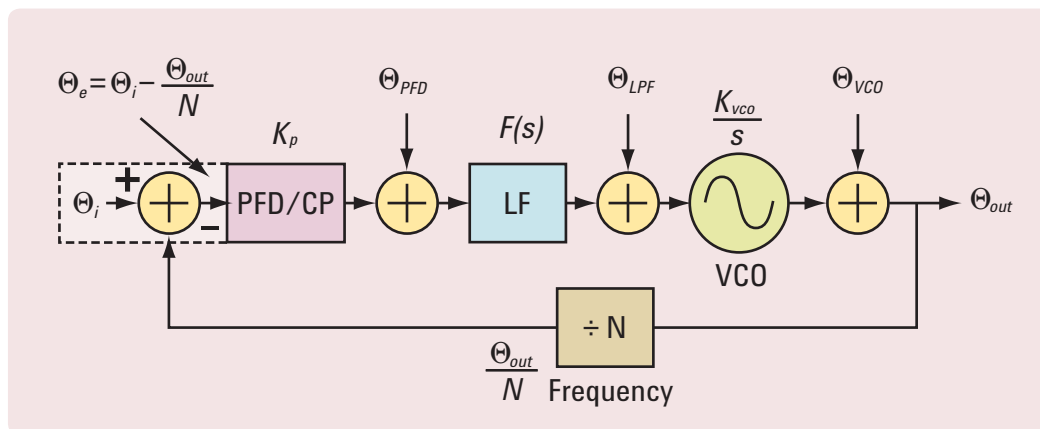


Figure 1. Block diagram of a generic PLL with noise signals at different nodes

Referring back to Equations 3 and 4, any noise from the input source and CP-based PFD is low-pass filtered. The noise due to static phase offsets, contributed by the CP leakage current, can be kept low using good PFD and CP design techniques. The noise from the LF is shaped by a bandpass transfer function whereas the noise contribution from the VCO is high-pass filtered.

Exploring analysis capabilities and methodologies

Tools such as RFDE and ADS provide an accurate, well-defined methodology for predicting first-order phase-noise performance in PLLs. These tools in particular provide five important noise-simulation capabilities:

- **Linear noise analysis:** S-parameter and small-signal AC (noise voltage).
- **Budget noise analysis:** System-level design at the block diagram level.
- **Nonlinear noise analysis:** Harmonic balance (HB) mixer noise figure, HB phase noise and noise voltage.
- **Transient nonlinear noise analysis:** Noise generated by noise sources, nonlinear devices and passive devices (jitter).
- **Circuit envelope noise analysis:** A Monte Carlo technique to simulate noise.




Figure 2. Noise contributions from individual blocks affect closed-loop PLL phase-noise performance

Using these tools, characterization and simulation of the PLL closed-loop phase noise can be accomplished using a five-step process:

1. Characterize the VCO/divide-by-N chain
2. Model the VCO/divide-by-N chain
3. Characterize the CP and PFD
4. Model the CP and PFD
5. Create the PLL functionality model

A closer look at these five steps will further illustrate what is possible with these tools.

Steps 1 and 2: Characterize and model the VCO/divide-by-N chain

A transistor-level VCO/divide-by-N chain can be simulated using the HB oscillator and noise analysis. The phase-noise modulator component from the ADS/RFDE model library can then be used to emulate VCO/divide-by-N phase noise. Figure 3 compares phase noise results from a model and an HB simulation.

Steps 3 and 4: Characterize and model the CP and PFD

Characterization of the transistor-level CP and PFD is done using a two-step simulation. First, transient analysis is performed to find PFD/CP sensitivity (e.g., CP current versus input phase difference at the PFD). Next is transient nonlinear noise analysis, which is performed with time-domain noise off and on.

CP noise current is then calculated by taking the difference between currents with noise on and off. Jitter can be calculated using the transient nonlinear noise and the sensitivity calculation of the PFD/CP. Jitter can then be used for a current-noise source or for the PFD/CP model.




Figure 3. Overlaid phase noise plots from a model (blue) and an HB simulation (red) show close agreement

Step 5: Model PLL functional behavior

Figure 4 shows an ADS schematic of a behavioral model of the PLL. It includes behavioral blocks for the VCO/divide-by-N chain and PFD/CP along with the circuit-level loop filter in ADS (Figure 5).




Figure 5. An example loop filter from a PLL

Circuit-envelope noise analysis can be run with the reference and oscillator frequencies. Post-processing of envelope-noise data in the ADS/RFDE data display can be used for calculation of phase noise from the nodal noise voltages.

Comparing simulated and measured results

Figure 6 shows the results of a simulation that followed the five-step process detailed in the previous section. Figure 7 shows the measurement of the modeled PLL as implemented with the IBM CMOS7RF process. This PLL has every block shown in Figure 1 except the LF. The key difference: The PLL on the IC had a variable-gain amplifier (VGA) following the VCO, a feature not included in the previous simulations.




Figure 6. Simulated phase noise from the PLL frequency synthesizer

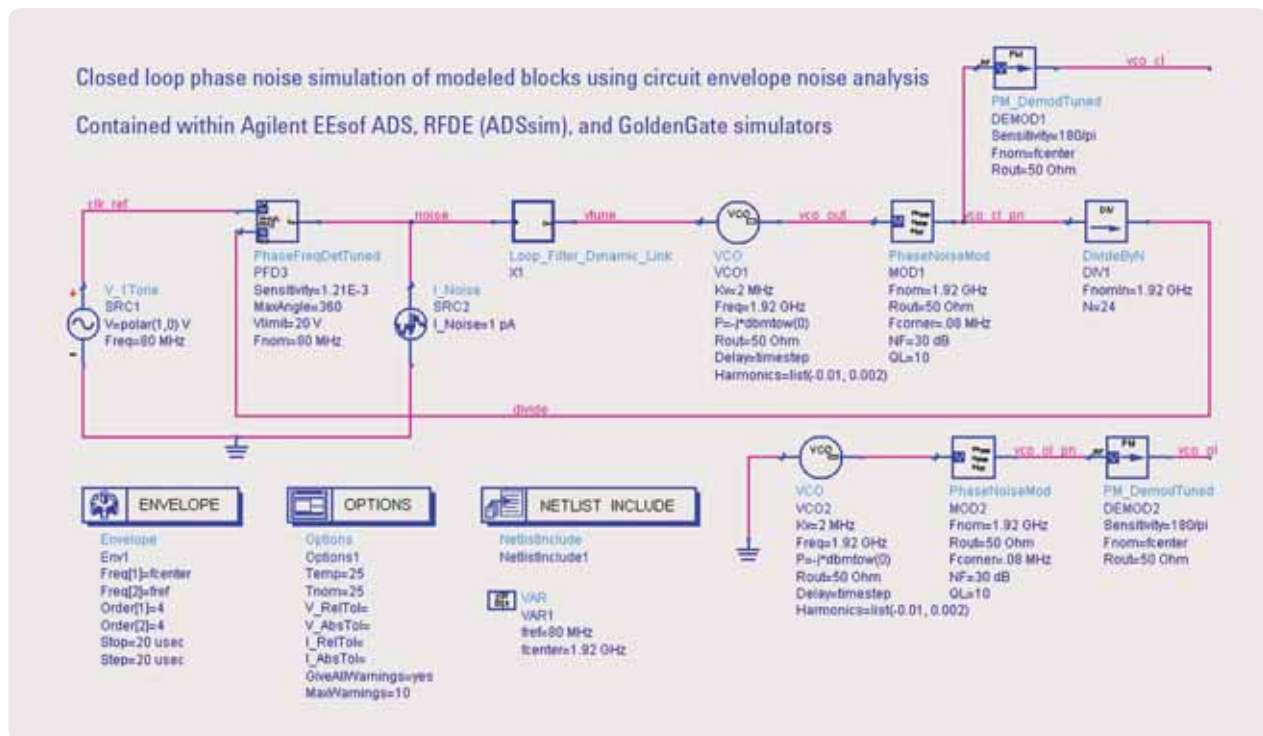


Figure 4. An ADS schematic of a behavioral PLL

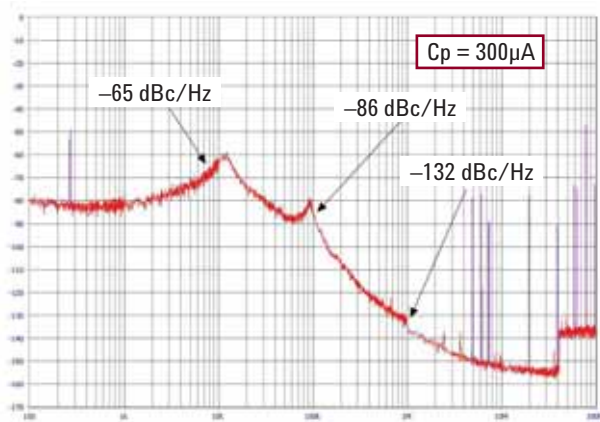


Figure 7. The measured phase noise of the PLL frequency synthesizer

Table 1 compares simulated and measured PLL closed-loop phase noise at key frequencies. Taking into account the circuitry differences between the simulated and measured PLL, the simulated results can be used as a good first-order prediction of the real PLL on the IC.

Table 1. Simulated versus measured PLL closed-loop phase noise at two key frequencies

Frequency	Simulated	Measured
10 kHz	-60.9 dBc/Hz	-65 dBc/Hz
100 kHz	-86.97 dBc/Hz	-86 dBc/Hz

Conclusion

PLLs or frequency-locked loops (FLLs) are common in many types of communication circuitry including serial/deserialize (SERDES), clock recovery, data recovery, and LO frequency synthesis. As data rates for high-speed serial links rise to 40 GB/s and as RF wireless frequencies increase, the noise (phase or jitter) of a locked PLL/FLL and first-order locked-loop performance must be designed to meet narrower phase-noise margins related to bit error rate (BER), packet error rate (PER), error vector magnitude (EVM), and dynamic range. When designing to these tighter specifications, the ability to characterize locked-loop circuit-level blocks and accurately assess the noise being introduced by these blocks in a closed-loop condition is extremely valuable. Meeting these specifications is critical to the creation of successful first-pass product designs that provide lower engineering and manufacturing costs. This is especially true at the IC level because the initial manufacturing costs tend to be very high.

As shown in this article, simulated results based on a PLL behavior model were consistent with measured results from an IC (allowing for the circuitry differences between the modeled and measured PLL). This work supports the idea that actual PLL performance can be predicted at the first-order by using a robust behavioral model simulation.

Acknowledgements

The author wishes to thank Yi Fang and Wei Zhang, formerly of M/A-COM, for their contributions to this article. He would also like to acknowledge the assistance of Andy Howard from Agilent Technologies on RFDE/ADS simulations and the support of Carlos Cardenas from M/A-COM on simulation tools.

In December 2007, an earlier version of this article was published in two parts on the Planet Analog website, www.planetanalog.com.

References

1. Leeson, D.B. 1966. A Simple Model of Feedback Oscillator Noise Spectrum. *Proc. IEEE*, Volume 54: 329-330.
2. Razavi, B. 1996. A Study of Phase Noise in CMOS Oscillators. *IEEE J Solid-State Circuits*, Volume 31, March 1996: 331-343.
3. Hajimiri, A. and Lee, T.H. 1998. A General Theory of Phase Noise in Electrical Oscillators. *IEEE J Solid-State Circuits*, Volume 33, February 1998: 179-194.
4. Demir, A., Mehtotha, A. and Roychowdhury, J. 1999. Phase Noise in Oscillators: A Unifying Theory and Numerical Methods for Characterization. *IEEE Trans. Circuits Syst. II*, Volume 46, Jan. 1999: 56-62.



Creating Synchronous High-Frequency Sampling Across Multiple Digitizers

Yves Maumary, Ph.D.

Software Design Manager, Agilent Technologies
yves_maumary@agilent.com

T

Today, high-speed digitizer systems operating at well above 100 MSa/s are being used in a diverse range of applications: operation of single-pulse linear induction accelerators for flash radiographic facilities, hypervelocity ballistic-range experiments, propulsion research, and more. A growing number of such applications require simultaneous measurement of high-frequency signals over many channels.

Most of today's high-speed digitizers or oscilloscopes feature a maximum of only four channels. For applications requiring more than four channels — and needing very precise time correlation between channels or accurate phase of continuous signals — it is necessary to synchronize the sampling clocks of multiple instruments.

With Agilent Acqiris digitizers, synchronous sampling can be achieved across several modules with ASBus, a bus system that distributes trigger and clock signals.¹ Up to seven modules can be connected with ASBus; however, by distributing a common, high-stability 10-MHz clock reference to all digitizers, it is possible to easily achieve synchronous sampling across a greater number of acquisition channels. One important challenge remains: measuring sub-nanosecond time delays between the synchronous samples of different channels. This article presents a method for measuring sampling-clock delay using the acquired signal as a time reference.

1. ASBus is a proprietary high-bandwidth auto-synchronous bus system that allows distribution of all necessary trigger and clock signals across up to seven digitizer modules.

2. Throughout this article, synchronous sampling is defined as follows: For any two channels A and B acquiring data in synchronous sampling mode, and where the i^{th} voltage sample u_{Ai} is the measured signal voltage on channel A at time t_i , there is a corresponding sample u_{Bi} measured on channel B at time $t_i + \Delta_{AB}$. Δ_{AB} is the sampling-clock delay of B with respect to A and is constant over all i values.

Achieving synchronous sampling

Accurate time correlation requires synchronous sampling across multiple digitizers and multiple channels.² This can be required when analyzing multi-channel single-shot events, for example, or when digital signal processing (DSP) operations combine samples from different signal channels before processing the data.

Although it is possible to achieve multi-channel synchronous sampling by distributing a common sampling clock to the various modules, this presents a major technical challenge at high frequencies. As one example, the backplane busses and connectors used in CompactPCI/PXI chassis are not well suited to high-frequency signals, and above about 100 MHz, clock-pulse edges deteriorate significantly and induce jitter. Using coaxial cables and proper connectors requires costly high-frequency fan-outs.

Another way to achieve synchronous sampling is to lock each digitizer's sample-clock generator to a common high-stability 10-MHz reference signal. By feeding the frequency reference to every module, the "sampling instants" on all channels will be synchronous (e.g., the sampling-clock delay between any two channels will be constant).

The sampling-clock delay includes all delays due to factors such as delay lines, signal path lengths and cable lengths (for the 10-MHz frequency reference). To verify that the criteria for synchronous sampling are satisfied, this delay must be shown to be constant. The sampling instants t_i are equally spaced (within the clock jitter) and have an interval equal to the inverse of the sampling frequency. With constant sampling-clock delay, the waveform data can be resampled using interpolation to yield a waveform with samples taken at exactly the same instants as the chosen reference channel. In cases that require data from several channels to be combined in DSP operations, it may be necessary to measure the sampling-clock delay to allow data resampling.

Measuring the sampling instants

Accurate measurements of the sampling-clock delay require a precise time reference common to both instruments. The trigger instant constitutes the only time reference for the waveform data from the digitizer itself. Unfortunately, many factors affect the precision of this time reference. For example, the path of the input signal — and of the analog-to-digital converter (ADC) data once the signal is converted — is different from the path of the trigger signal on the digitizer board, implying different propagation delays. Components on these paths have propagation delays that vary from component to component. This is the *trigger-vs.-channel delay*.

The trigger comparator also can affect trigger-time precision. The comparator threshold is calibrated to a finite resolution, and any noise on a signal entering the trigger comparator causes trigger-time jitter.

Consequently, another time reference is needed. One possible solution is to use the input signal. If the signal delivered to the channel input connectors is identical, the samples from each waveform can be positioned with respect to a reference instant in the signal itself.

Outlining the measurement method

It is possible to show the constant sampling-clock delay between two channels of different digitizer modules (when using a common clock reference) by repeatedly measuring the sampling instants of channels relative to each other, two-by-two.

The starting point is a time reference, which is needed to accurately position the samples on an absolute time scale. Because the trigger instant is known to a precision much larger than that required for a measurement, the measured signal can be used as the absolute time reference. In this case, it must be exactly the same signal delivered to the channel inputs. By choosing a sine signal and fitting the measured data to a sine function, it is possible to position the samples in time relative to the signal (e.g., relative to the first positive-going zero crossing). The difference between the fitted phases of the waveforms from each channel is the sought-after sampling-clock delay.

This method uses the trigger-time information from the digitizers for one purpose: to determine the starting samples in each waveform. This yields a measurement of the sampling-clock delay that is free of the errors and jitter caused by digitizer trigger systems.

The measured samples from each channel must be taken during the same period of the input signal. Therefore, the digitizers must be triggered at exactly the same time using an external trigger pulse. Note that this trigger pulse is completely asynchronous to the signal and the clocks. Even without using the trigger time from the digitizer, it is known that the first sample from each waveform is taken within *sampling interval + trigger-vs.-channel delay + trigger jitter* of the trigger instant.

The frequency of the input sine wave should be chosen appropriately. A higher frequency will provide better timing accuracy, but the period must be long enough compared to *sampling interval + trigger-vs.-channel delay + trigger jitter* to resolve the ambiguity due to the period folding (see next section). Also, more samples per period will allow the sine fit to converge without having to specify the frequency.³

Finally, the sine wave frequency is chosen such that the phase of the digitized samples is different for each period over the complete acquisition time window. Nonlinearity errors such as those due to the ADC eventually translate into time errors and can be eliminated by the averaging effect of the sine fit over many periods.

The principle of using a sine fit over many periods will average out digitizer imperfections such as nonlinearity in ADCs and high-frequency phase noise in the sampling-clock generator. The only remaining errors come from low-frequency phase noise in the clock generators.

The stability of the sampling-clock delay can be verified by repeating the measurement over many acquisitions.

Defining the folding ambiguities

The method described above contains two inherent folding ambiguities, which are explained in the diagrams of Figure 1.⁴ In the figure, t_i and t_j are the phases from the sine fit converted into time. D is the sampling-clock delay and s_i the sampling interval. Red crosses represent the samples from the reference

3. In experiments at Agilent, 16 samples per sine period (25 MHz sine wave sampled at 400 MSa/s) yielded the best results.

4. When considering two repetitive events relative to each other (e.g., the sampling instants on two different channels), one is chosen as the reference. Because it is not known which one occurs first, the hypothesis that the reference event occurs *before* the other event may, for example, subtract a full period to the calculation of the delay between the two events in the case the hypothesis turns out to be false. This method-induced effect (a full-period shift in the example) is called a *folding*.

channel, with A marking the first sample. B marks the first sample of the measured channel, represented as blue circles. T is the period of the sine wave.

The first ambiguity, called *trigger folding*, is due to the trigger instant falling between B and the subsequent red cross (one sampling interval after A), or between A and the subsequent blue circle (one sampling interval after B). With Agilent Acqiris digitizers, by definition, the first sample in an acquisition always precedes the trigger instant by less than a sampling interval. Therefore, if the trigger instant (which is common to both digitizers) arrives between A and the subsequent blue circle (second diagram), it means the hypothesis of A preceding B is false and it is necessary to subtract one sampling interval from the difference $t_i - t_j$ to obtain the sampling-clock delay D.

The second ambiguity, called the *sine-fit folding*, is due to the sine fit, which returns a phase of $\pm\pi$. The trigger can fall at a

time such that channel A has a negative phase and channel B a positive phase (Figure 1, bottom diagram). In such cases, the period T of the sine wave must be added to the $t_i - t_j$ difference to obtain the sampling clock delay D.

Examining measurement results

A series of actual measurements shows the precise synchronization that is possible with the methods described above. The three key parameters were the synchronization states, the sampling-clock delay and the stability of the sampling-clock delay.

Configuring the measurement

The test setup used for the measurements reported below was composed of two Agilent U1066A Acqiris DC440 12-bit digitizers (Figure 2).

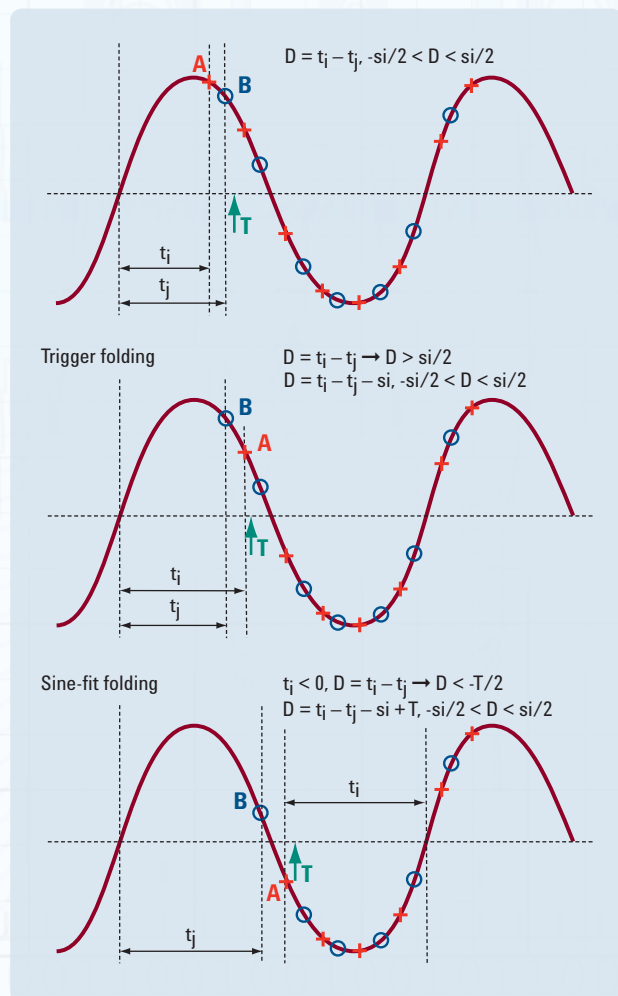


Figure 1. Two types of folding ambiguities — trigger folding and sine-fit folding — are present in the described measurement method.




Figure 2. The test setup included two digitizers, one signal generator and one function/arbitrary waveform generator for the trigger pulse.

The sine wave signal providing the time reference was produced by an Agilent synthesized signal generator. The signal amplitude was adjusted to be about 80 percent of full-scale input. To ensure the same signal was being fed into both channels, a 50 Ohm passive splitter was used to connect two cables of identical length. A similar splitter and identical length cables were used with an Agilent 33220A function/arbitrary waveform generator to provide a trigger pulse.

The digitizers were controlled through a MATLAB® script. Care must be taken to ensure that both digitizers are armed and ready before a trigger pulse is sent, preferably by controlling the 33220A with the MATLAB script.

Results: Synchronization states

When synchronizing the sampling clock generators with a common 10-MHz reference, several effects come into play that produce multiple possible synchronization states. The first is linked to the actual hardware implementation of the sampling-clock generator. In the case of Agilent Acqiris digitizers, six possible states are especially interesting: two are one-half of a sampling interval apart at the maximum sampling rate of 400 MSa/s and the other four are one-quarter of a sampling interval apart at 200 MSa/s. When decimation is used to achieve lower sampling frequencies, additional sampling-clock delay states may be created through the initial state of the decimation counter.⁵ These must be taken into account when comparing the results of sampling-clock delay measurements.

When the sampling frequency is set, the digitizer clock generator settles into one of the possible states due to the first effect described above. It remains locked in that state for all further acquisitions, until the registers controlling the clock generator are reloaded. The latter occurs either when a full self-calibration is performed or when the sampling frequency is changed. The state multiplication effect due to decimation, however, is always effective (the decimation counter is reset at the beginning of each acquisition).

A series of measurements at 400 MSa/s, 200 MSa/s and 100 MSa/s verified that the state was indeed conserved within a calibration epoch (unless decimation was active).

5. Decimation is the process of retaining only one of n samples to yield a lower effective sampling frequency than the sampling-clock oscillator can provide.

Results: Sampling-clock delay

A series of one million measurements was performed at 400 MSa/s. Each acquisition was 2,048 samples long, and the sine-wave signal frequency was set to 25.048828125 MHz. Figure 3 shows a histogram of the measured sampling-clock delay: The mean value is 1.115 ns and the standard deviation is 3.475 ps.




Figure 3. This histogram of the measured sampling-clock delay has a mean value of 1.115 ns and a standard deviation of 3.475 ps.

It is worth noting that the standard deviation includes delay drift due to temperature variations that occurred during a 35-hour measurement in an uncontrolled environment. This is illustrated in Figure 4, which plots the measured delay for each of the one million measurements. The delay variation due to temperature was measured to be 10 ps/°C. The sampling-clock delay variation with temperature was recorded over a period of about 14 hours (Figure 5). Over a short time period (e.g., 10,000 measurements) the typical standard deviation is reduced to about 2 ps.




Figure 4. The standard deviation of the sampling-clock delay includes delay drift caused by temperature variation during the 35-hour measurement.




Figure 5. Over a period of about 14 hours, the sampling-clock delay varied with temperature.

Results: Sampling-clock delay stability

To verify that the measured sampling-clock delay does not vary during an acquisition, the measurement method was modified slightly to use the sine fit for only small chunks of data from a very long acquisition. Figure 6 shows the delays obtained from 500 fits of 8,000 samples each along a 4 MSa acquisition. The red lines show the standard deviation, while the two external lines depict the maximum and minimum values. The standard deviation is about 1 ps.

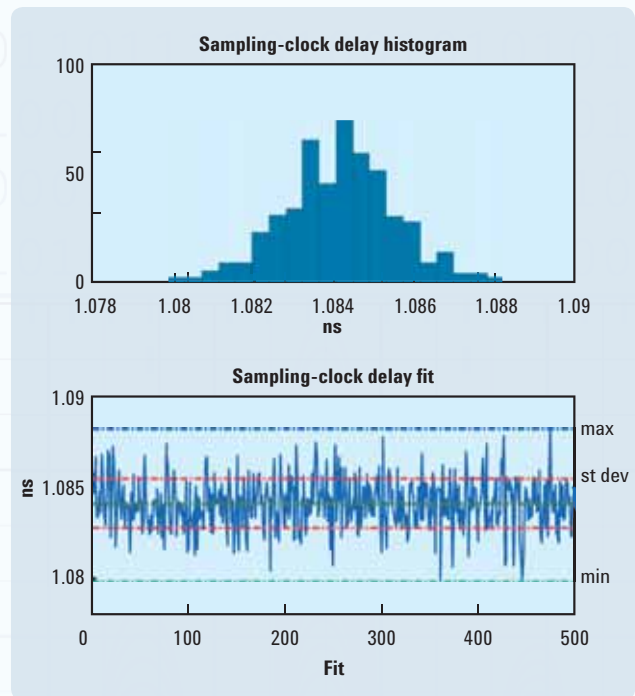



Figure 6. Delays obtained from 500 fits of 8,000 samples had a standard deviation of about 1 ps.

Conclusion

For highly demanding applications, synchronous sampling across multiple instruments can be achieved using a common 10-MHz clock reference, provided by either a master instrument or an external high-stability source (when better precision is required). When the sampling-clock delay between channels must be precisely measured, the sine-fit method allows effective suppression of most error sources.

The described method was used to verify the sampling synchronicity between Agilent U1066A Acqiris DC440 12-bit digitizers. The results show that it can be successfully used to determine the sampling-clock delay between channels.

MATLAB is a U.S. registered trademark of The Math Works, Inc.






Overcoming the Challenges of Testing FlexRay Networks

Cary Brown

Technical Marketing Manager, Agilent Technologies
cary_brown@agilent.com

Narcis Ilisei

R&D Software Lead, Agilent Technologies
narcis_ilisei@agilent.com



Today, the amount of in-vehicle electronics is clearly on the rise. What's more, with automakers continuing to introduce further advances in comfort, reliability and safety, this trend shows no sign of abating any time soon. As a result, the number of electronic control units (ECUs) required to control the associated automotive components will increase — as will overall system complexity. Unfortunately, that translates into a higher chance for errors, which, in turn, can severely erode a consumer's confidence in his or her vehicle.

Likely culprits of such errors are irreproducible electronic problems that leave no traces other than entries in error logs. Examples of such problems are vehicle-integrated cell phones that temporarily cease to function, and radios that can be switched on after starting the ignition but only produce sound after rebooting.

To help deal with the increasing complexity of car-internal networks and the number of ECUs, automakers and automotive electronic manufacturers have turned to FlexRay — a highly reliable protocol capable of high-speed communications — as the next-generation automotive control network. Offering a combination of high-speed static and dynamic transmission, it will likely become the *de facto* standard for in-vehicle communications.

As the FlexRay protocol becomes widely adopted as the network of choice for future in-vehicle applications, appropriate testing will become critical to ensuring proper operation of in-vehicle electronics — and ultimately ensuring consumer confidence. Accomplishing these goals requires a quality metric that gives developers a clear indication of the quality of their work. That metric should also provide the Quality Assurance (QA) department with a tool that enables a clear, reproducible and descriptive quality metric for every network design. The tool of choice for helping create this quality metric is a protocol analyzer designed to handle issues specific to FlexRay networks.

Taking a closer look at FlexRay

The FlexRay protocol is a new element within the increasingly complex automotive system. As a high-speed, deterministic communications technology, it has a maximum data rate of 10 Mb/s on two channels (gross data rate of up to 20 Mb/s), and is capable of guaranteeing real-time operation. Other prominent FlexRay features include time-triggered behavior, redundancy, safety and fault-tolerance. With these capabilities, an in-vehicle network with FlexRay at its core is capable of providing determinism for engine control and fault tolerance for steer-by-wire and brake-by-wire — which require 100 percent reliability — as well as other advanced safety applications (Figure 1).

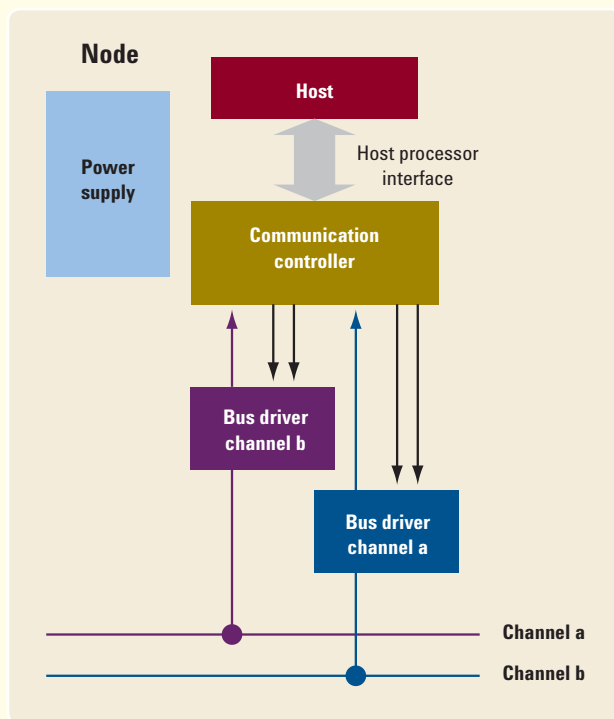


Figure 1. The FlexRay block diagram.

The FlexRay protocol is composed of both static (ST) and dynamic (DYN) segments arranged to form a periodically repeated bus cycle. The ST segment employs a generalized time-division multiple-access (GTDMA) scheme. In contrast, the DYN segment uses a flexible TDMA (FTDMA) bus-access scheme (Figure 2).

FlexRay provides scalable (e.g., single- or dual-channel) static and dynamic message transmission and incorporates the advantages of familiar synchronous and asynchronous protocols. It supports collision-free bus access, fault-tolerant clock synchronization through a global time base, and guaranteed message latency with message-oriented addressing occurring through identifiers. FlexRay has an error-management service that provides error handling and error signaling, a wakeup service that addresses the automotive system’s power-management needs, and a diagnosis service that tests the bus guardian on the physical layer.

While the FlexRay protocol offers definite benefits and holds great promise for advanced automotive communications systems, it also presents some difficulties. Namely, in a FlexRay network, sporadic errors are very hard to find and are often traced back to network communication issues. Recent studies have shown that errors in network communication are the number two cause of failures in automotive electronics and are among the highest-cost automobile repairs. The problem is further complicated by a design trend toward even more networking, which will result in higher complexity. The ability to find network communication errors depends on the engineer ensuring the signal integrity of FlexRay signals as well as proper timing of the time-triggered communication bus. It also requires a quality metric specific to the FlexRay protocol.

Examining the FlexRay quality metric

The quality metric for networks based on complex protocols such as FlexRay must meet a few requirements. In the case of the FlexRay protocol, the quality metric has two individually observable measurements: physical layer quality and protocol layer quality. Metrics for each of these measurements are defined as follows:

Physical layer quality metric: This metric covers the quality of the network cabling and identifies which impedance measurements exist at the various connectors, stubs and termination points. The FlexRay physical layer measurement of robustness is based on the collection of all measured electrical characteristics shown to be in accordance with design specifications. This quality measurement demonstrates if design parameters are obeyed and, subsequently, the achieved degree of quality.

Protocol layer quality metric: This metric covers the higher levels of the protocol, highlighting such problems as semantics, incorrect packet data and incorrect cyclic redundancy check (CRC) values. Timing dependencies are also observed, for instance, to check the update intervals of signals or sequences of events. Such timing errors or race conditions can produce errors that are very hard to reproduce: From an electrical point of view, the system may work perfectly; however, communication errors may still occur and the system will not behave as planned.

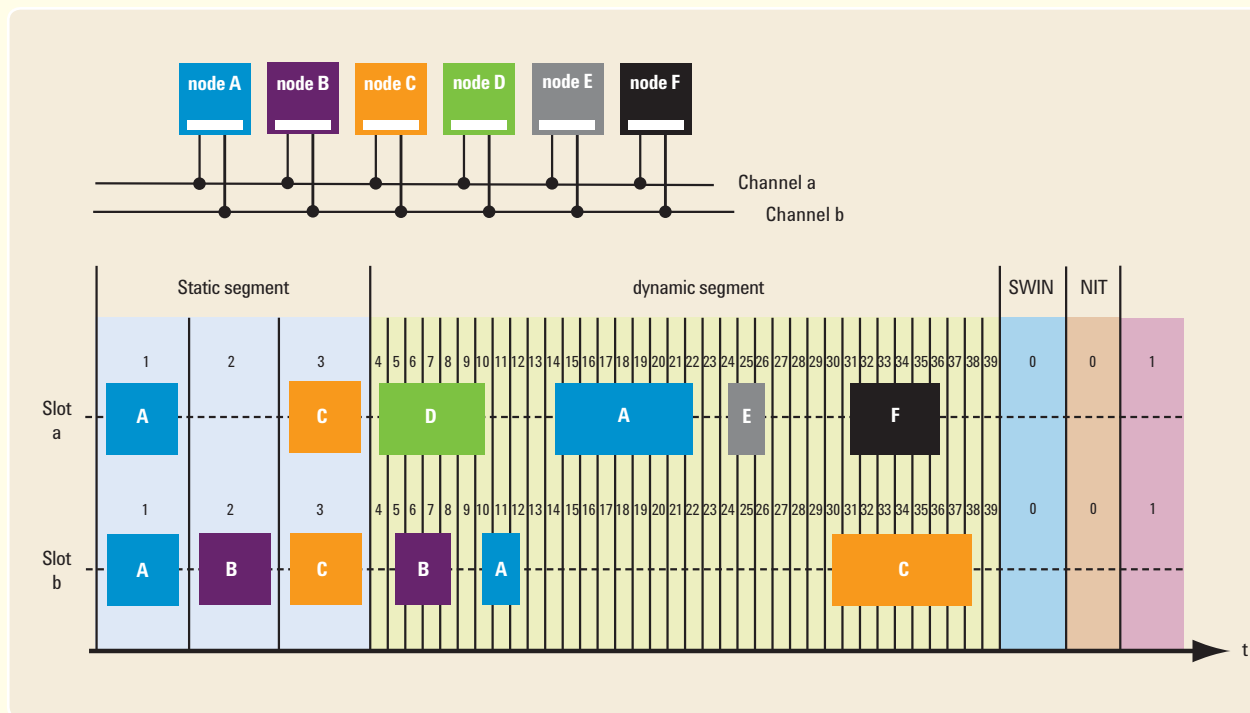


Figure 2. FlexRay segment configuration consists of the static segment with the fixed time trigger method and the dynamic segment with the flexible time trigger method.

Specific measurement tools are required to create quality metrics for FlexRay’s physical and protocol layers. Measurement tools are currently available to help create the protocol layer quality metric; however, this is not yet the case for the physical layer. No complete test solution is currently available to address this need. Instead, it is incumbent on the measurement industry to provide — as soon as possible — a measurement device that can fulfill this need. This is the only way to help FlexRay realize its promise of reliability, real-time capability and robustness for the physical layer.

Assessing the physical layer

Creating a quality metric for FlexRay’s physical layer requires the use of an appropriate measurement approach and solution. There are several approaches that are not well-suited to this task. They include:

- The classical method of using an oscilloscope, capable of producing an eye diagram, to measure electrical signal quality. This approach is not suitable for FlexRay networks for two reasons. First, each individual protocol specification defines clear masks for eye shapes that indicate compliance. In the case of the FlexRay protocol, this approach is viable only when using a point-to-point connection topology (Figure 3). That’s because the FlexRay

specification only defines a mask for this type of connection topology. Second, due to the internal architecture of many oscilloscopes, a period of dead time occurs after each trigger and before the next trigger event can be captured. During this dead time, the scope is essentially inactive, preventing a continuous measurement and making it virtually impossible to cover 100 percent of all bit transitions — a prerequisite for safety-critical applications such as X-by-wire systems.

- Use of a bit error rate tester (BERT) to validate the quality of a network and its physical connections. For a number of reasons, this method is not viable for the FlexRay network. To begin with, in order to calculate the bit error rate, the BERT must know what sort of data is coming so that it can compare the incoming data stream to the expected data. In a dynamic, non-deterministic system such as a car, this is simply not possible. Additionally, in a FlexRay network, each ECU has its own clock, which is used for data transmissions. Because BERT measurements require synchronization with a common unique clock, this makes it difficult to use a BERT for FlexRay. Further, by default, bit error rates show no errors in lab and test usage. Robustness, while not needed in test versions, is required in the final product. In order to make a valid and statistically sound statement at a 10^{-12} bit error rate, test times simply become impractical.

Rather than the approaches outlined above, what's required is a measurement tool that combines a number of key capabilities. The tool must, for example, use thresholds from the FlexRay specification to detect all bit transitions and have the ability to synchronize constantly on the byte-start sequence (BSS) of the FlexRay signal. This latter capability will help avoid problems associated with different clocks. The optimal tool should also be able to detect all bit transitions without any dead time. Additionally, the tool must use the appropriate multiple of the bit time — including the tolerance — to check for correctness at consecutive sequences of ones and zeros.

As previously mentioned, such a measurement tool is currently not yet available. However, corresponding developments and tests are currently underway that should enable the industry to soon create this type of measurement device.

Working at the protocol layer

There is no commonly agreed-to solution for the FlexRay protocol layer. This is due primarily to the layer's complexity. Defining a clear quality measure depends heavily on the application as well as the nature of the transmitted signals. Further complicating matters, errors in the protocol layer are often very broad and diverse. In general, they relate to the following questions:

- How often are signal updates expected?
- What is the maximum change rate of this signal?
- How well can missing packets (e.g., ZeroFrames) be tolerated?
- How are timing violations (e.g., slot violations in the static segment) checked?
- Is it possible to use, as an example, running counters as part of the payload to detect packet loss? If so, how are such errors detected and caught?

The description of the quality requirements for the protocol layer of a specific FlexRay network must be created by the network developer or in accordance with the specifications and requirements document. Because the typical controller used for FlexRay networks suppresses or filters out most error situations, developers need different measurement tools to identify these errors. The common tool of choice for such measurements is the protocol analyzer.

For this task, the ideal protocol analyzer must be capable of decoding and displaying the signals within a FlexRay network. Powerful triggering capabilities are also critical because these enable the developer to check for protocol errors and relevant events. To extract the necessary signal decoding, the protocol analyzer must support the latest FlexRay specification as well as common file formats such as Fibex files.

The ideal protocol analyzer must also be capable of handling issues specific to FlexRay networks. For example, to support a higher update rate for specific signals, the FlexRay specification allows multiplexing of signals in several static slots. This alleviates the need to reduce the entire cycle time just to fulfill the need of one signal. As a result, the protocol analyzer must be able to detect such a situation and correctly display dependent signal values which have been evaluated for consistency across multiple slot boundaries.

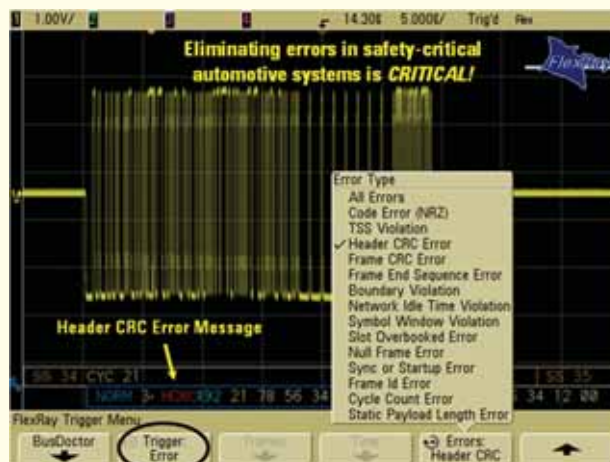


Figure 3. An example of an eye diagram in which a random timing violation occurred.

Also unique to the FlexRay specification is the creation of only generic error messages by standard FlexRay controllers (e.g., Bosch E-Ray controller). Only correct packets are handed to the upper layers; further details, such as why and when communication was terminated, are not shown. Therefore, the ideal protocol analyzer for FlexRay networks should include an additional measurement controller capable of detecting various protocol errors and timing issues. These errors can then be used as specific triggers for associated errors.

It is also important to note that there is no synchronous communication at the start and during the initial synchronization of the FlexRay protocol. Consequently, the ideal protocol analyzer must be able to capture asynchronous communication until synchronicity is reached. Switching between the asynchronous and synchronous modes should be automatic. This is crucial in cases where no Fibex data files are available. The ideal protocol analyzer must also be capable of displaying the FlexRay packets going across the bus — without knowing the cycle or slot timing.

Defining a viable solution

The Agilent J8130A Vehicle Protocol Tester Series 1000 (VPT1000) is a prime example of a FlexRay protocol analyzer that meets the requirements identified above for the protocol layer quality metric (Figure 4). It provides powerful hardware processing and the deep memory required to meet the analysis requirements linked to FlexRay networks running at speeds of up to 10 Mbps. Additionally, it provides the fastest time to insight by enabling car makers to increase productivity through powerful data capture, in-depth analysis with visualization capabilities and standalone data logging.



Figure 4. Agilent FlexRay VPT1000 supports protocol layer testing.

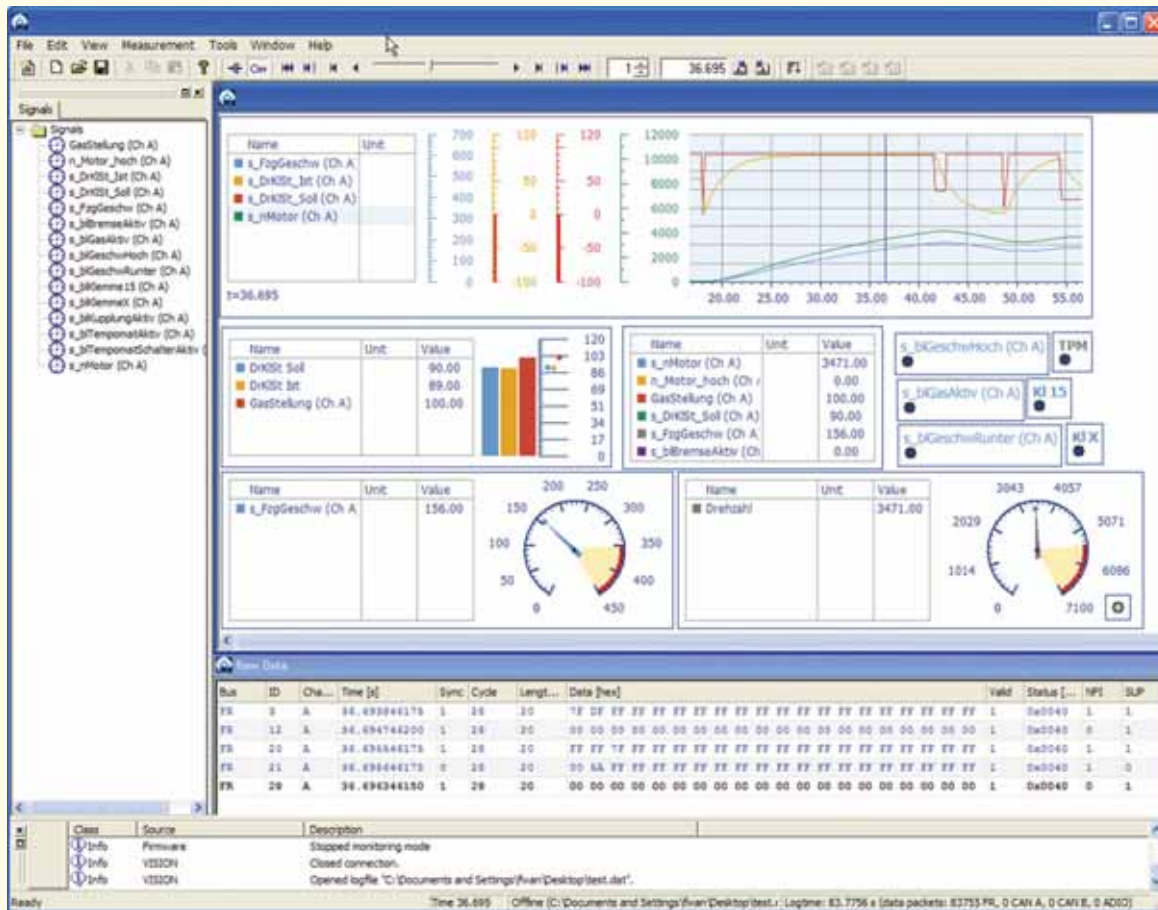


Figure 5. A sample screenshot of the VPT1000 FlexRay analyzer software.

The VPT1000 hardware is extremely flexible, supporting a wide range of use cases. It can be used with a PC for online logging and visualization, mounted in a car as a standalone data logger (without a PC), or as a protocol decode/trigger device together with the Agilent 6000/7000 Series mixed-signal oscilloscope (MSO) for combined physical layer and protocol analysis.

Several key features make the VPT1000 FlexRay protocol analyzer well suited for use with FlexRay networks:

- The VPT1000 is equipped with two FlexRay connections and two Controller Area Network (CAN) channels. CAN is a broadcast shared serial bus for microcontrollers used mainly in automotive applications. Both FlexRay and CAN can be displayed on the analyzer, on the same time base and with a resolution of 25 ns — a capability that is critical for integration tests and timing measurements. Also, up to four digital input/output (I/O) signals can be integrated into the measurement.
- The VPT1000 supports the correct display of multiplexed signals during a FlexRay cycle and can work in either synchronous or asynchronous modes. Switching between these modes occurs automatically, without any loss of data. This feature allows the analyzer to be used for start-up analysis or for measurements of disturbed networks, enabling fast and effective identification of the root cause of communication problems.
- The VPT1000 features an additional measurement controller to receive and detect more data in case of error situations. Timing and protocol error situations are clearly captured and can be used as trigger. Signals are decodable and a powerful trigger engine covers different protocol errors. The software supports direct import from Fibex files (Figure 5).


In addition to these features, the VPT1000 comes with a specific FPGA design to process full bus load based on the latest FlexRay Specification 2.1a. It also includes a FlexRay controller (Bosch E-Ray) and an additional controller implementation with extended protocol-error and frame-analysis capabilities.

Conclusion

Finding errors in a FlexRay network can be a difficult task, especially if those errors are sporadic in nature. A quality metric is vital to testing FlexRay networks and finding such errors, thereby ensuring consumer confidence as well as proper operation of in-vehicle electronics. Proper measurement tools are critical for accomplishing this goal. A FlexRay protocol analyzer such as the Agilent VPT1000 is well suited to helping create a quality metric for the FlexRay protocol layer. Currently, no such solution exists for the FlexRay physical layer. Obtaining such a tool will be vital to ensuring that the FlexRay protocol lives up to its true potential and enables the next generation of advanced automotive communications applications.

For more information, go to www.agilent.com/find/vpt1000

The authors would like to thank Tilmann Wendel, former Technical Marketing Manager, Agilent Technologies, for his contributions to this article.



Additional FlexRay Protocol Information

The FlexRay protocol was developed and proposed by a large consortium of automotive manufacturers and suppliers known as the FlexRay Consortium. Its core member companies include BMW, Bosch GmbH, DaimlerChrysler, Freescale, General Motors, NXP Semiconductors and Volkswagen. These seven companies have brought together their respective areas of expertise to define a communication system that is targeted to support the needs of future in-car control applications. There are also premium associate and associate members of the FlexRay Consortium. As of November 2007, there were 24 premium associate members and more than 70 associate members.

To date, the FlexRay Consortium has released the FlexRay Communication Systems Specification Set and it is now available to the general public. The first production vehicle with FlexRay was the 2007 BMW X5, although it was used only for the pneumatic damping system. Full use of FlexRay is expected in 2008.


The FlexRay protocol is viewed as the industry's answer for a high-speed, deterministic and fault-tolerant communications technology for advanced automotive control applications. While it was designed specifically for in-vehicle networking, FlexRay is not expected to replace existing networks such as CAN, Local Interconnect Network (LIN) or Media Oriented Systems Transport (MOST). Instead, it will work in conjunction with these systems.

FlexRay's advantage over other networks is predictable and reliable delivery of the very high data rates required by today's emerging systems such as X-by-wire, collision avoidance and driver assistance. MOST can support high data rates, but was designed specifically to connect in-vehicle multimedia. CAN is capable of connecting several ECUs, but at a typical data rate of only 500 Kbps to 1 Mbps. Also, because CAN employs priority arbitration for message delivery, only the highest priority message is assured delivery. What's more, its relatively low data rate and lack of fault-tolerance make it inadequate for advanced applications such as X-by-wire. LIN is essentially an inexpensive and comparatively slow sub-network to CAN, and as such cannot handle the data rates required by the advanced safety systems expected in next-generation vehicles.

For more information, please visit www.flexray.com.



Understanding Operating Life and Repeatability of Electromechanical Switches and their Effect on Total Cost of Ownership



Teh Chee-Leong

R&D Engineer, Agilent Technologies

chee-leong_teh@agilent.com

Ivan K.C. Hor

R&D Project Manager, Agilent Technologies

ivan-kc_hor@agilent.com

Optimizing the cost of test in a high-volume production environment has always been a formidable challenge for manufacturers, especially because test costs can be a significant percentage of overall product cost. One example is in mobile handset testing. According to IDC's *Worldwide Quarterly Mobile Phone Tracker*, global shipments of mobile handsets reached the one billion unit milestone in 2006. The actual shipment number — 1.019 billion — represented a 22.5 percent growth over the 832.8 million units shipped in 2005.

This impressive growth intensifies the importance of reliability and performance in the automated test systems (ATS) used for mobile-phone testing. Consequently, high-volume component manufacturers are always looking for test solutions that are not only cost effective in terms of throughput, but that also provide the highest accuracy and performance. Furthermore, the explosive growth in the very large scale integration (VLSI) arena, coupled with the need to squeeze more features and functionality into a single component requires expanded test coverage to cover both DC and the RF spectrum. Test systems and solutions have become highly complicated switching boxes or modules that guarantee the high level of accuracy, performance and reliability needed to meet the requirements of automated testing. Ultimately, this translates into a higher premium being placed on lower cost of ownership for these types of systems.

Consequently, there is a strong need for multiple switches capable of providing the various routings of test paths between DUTs and test and measurement systems with a guaranteed level of precision (switching repeatability) and accuracy (RF

specifications). The emphasis is also on switch operating life and reliability. Long life, repeatability and reliability directly lower the cost of ownership by reducing calibration cycles, increasing test-system uptime and ensuring ATS measurement integrity over time. As a result, operating life and repeatability of generic electromechanical (EM) switches have now become the two most critical properties that determine the total cost of ownership of ATS (see the sidebar, *Switch types*). However, the mechanical nature of the EM switching mechanism means that their life, precision and reliability are dependent on the switch design, mechanism and materials, as well as the process control used during mass production. As a result, careful consideration must be given to these factors when selecting an EM switch.

Switch types

There are basically two major types of connectorized RF and microwave switch modules: electromechanical (EM) switches, which rely on mechanical contacts as their primary switching mechanism, and solid-state switches. The latter category is comprised of field-effect transistors (FETs) and PIN diodes. FET switches create a channel (depletion layer) that allows current to flow from the drain to the source of the FET. The PIN diode consists of a high-resistivity intrinsic (I) layer that is sandwiched between highly-doped positively (P) charged material and negatively (N) charged material. Selecting which type of switch to use depends on the application. EM switches are often preferred due their very low insertion loss, high power-handling capability and excellent immunity to electrostatic discharge (ESD).

Operating life of an EM switch

The operating life of an EM switch is defined as the number of cycles the switch will complete while meeting all RF and repeatability specifications. Operating life also refers to the electrical life of the switch, as opposed to the mechanical life, which is significantly longer. One life cycle is defined as one closing and opening of the jumper contact — sometimes referred to as switch blade — or one on/off triggering of the electromagnetic coils in the switch. The operating life is very dependent on the jumper contact mechanism, contact resistance, and the material and plating used in all the key RF components.

Examining the conventional switch-contact mechanism

Conventional switches function by moving the jumper contact — a thick rectangular element — inside the RF housing. The jumper contact is joined by a push rod, generally made of a dielectric material such as polystyrene (PS), which moves inside an access hole in the RF housing. The tip of the jumper contact is pressed directly onto the flat surface on the tip of the connector's center conductors by a mechanical spring force from the actuator.

Figure 1 depicts an open RF line with the jumper contact retracted. Figure 2 illustrates a closed RF line in which the jumper contact forms a bridge between the input and output ports, thereby allowing propagation of the RF signal between these ports.




Figure 1. RF line opened




Figure 2. RF line closed

The jumper contact is usually thick and inflexible, as illustrated in Figure 3. The vertical motion of the jumper contact and push rod during the opening and closing sequences results in what is sometimes called “frictionless switching” because there is no friction produced between the jumper contact and the center conductor.




Figure 3. Conventional EM switch mating architecture

Although this design allows mechanical actuations on the order of tens of millions of cycles, there are some significant drawbacks. The continuous impact between the jumper contact and center conductor will gradually result in increasing wear and tear, producing some debris. This debris, along with dirt and contamination accumulated over time, remains on the tip. As a result, contact resistance increases over time, leading to increased insertion loss. This may or may not result in the switch

failing its RF specifications, but will have a significant effect on insertion-loss repeatability over time. The random nature of this particle buildup also means that such failure can be intermittent — and it may not be detectable, as these particles remain trapped on the surface of the center conductor throughout the lifetime of the switch. This adverse effect is the result of an inflexible jumper-contact design. Consequently, it is not uncommon to find switches with designs of this nature having loose repeatability specifications — or none at all.




Figure 4. The particle buildup phenomenon

Assessing contact-surface materials

An important factor in determining the performance and operating life of an EM switch lies in the materials, plating and surface profile employed at the contacting surfaces. The combination of contact finish and plating materials is also crucial to the handling of high-power signals. Contact finish affects the series resistance of a pair of closed jumper contacts, whereas plating material affects both the contact resistance and thermal conductivity of the assembly.

The jumper contacts used in EM switches are often fabricated from beryllium copper alloy, followed by a thin layer of a good conductivity metal (e.g., gold) on the contact areas. A thin layer of gold finish promotes excellent corrosion resistance, low contact resistance, good RF characteristics, and acceptable wear characteristics. Gold is the preferred plating material for the contact due to its intrinsically low resistance and its capability

to withstand oxidation and environmental corrosion. However, there is an issue with this choice of plating material. Because the gold layer is plated directly on the beryllium copper surface, the copper will eventually migrate into the gold layer, and the gold will likewise diffuse into the copper layer. This migration is further accelerated by the presence of oxygen, heat and humidity. To minimize the migration and prevent diffusion of gold and copper atoms, some form of barrier material is normally used between the beryllium copper and gold layers.

Beryllium copper is a metal alloy that typically contains copper with 1.8 to 2 percent beryllium. Sometimes additional alloying elements are added. In addition to possessing significant metalworking advantages, this amalgam exhibits good electrical and thermal conductivity characteristics. For example, its good thermal conductivity and very high tensile strength allow the jumper contacts to be exposed to high temperatures without any risk of melting or deterioration — key factors that ensure consistently good pressure contact and prolong the operating life of the switch.

Adding a wiping mechanism

Increasing the repeatability and operating life of a switch requires a design that essentially cleans off the center conductor tip during each switching cycle. This process eliminates the particle buildup prevalent in conventional EM switch designs.

The principle of contact wiping action is widely known and has been applied in relays and keyboard switches to break through surface corrosion and debris on contacting surfaces. In the context of EM switches, a properly designed wiping mechanism plays a very important role in prolonging electrical life and maintaining repeatability. In addition to wiping action, Agilent EM switches also use suitable lubricants and smooth finishing at the contacts. This unique design produces excellent repeatability of less than 0.03 dB across the operating life of the switches, while maintaining all RF specifications.

Essentially, the wiping action pushes any small particles or debris out of the contact zone, allowing the switch to self-clean. However, it is important to note that excessive wiping action combined with high contact pressures can generate debris due to excessive rubbing between the two surfaces. Therefore, the key is to generate the optimum amount of wiping action between the contacting surfaces.




Figure 5. A "microscopic wiping" switch-mating architecture

The wiping action mechanism employed within the Agilent EM switch family is illustrated in Figure 5. Here, the center-conductor profiles of the connectors are designed with a spherical mating surface. This mating surface is slightly curved to create a minor downward force and a small movement between the jumper contact and the mating surface. This movement is made possible with a thin, flexible jumper-contact design that promotes a microscopic wiping effect between the two surfaces. The wiping action continuously cleans the contact area by breaking through the surface films and moving debris away. Figures 6 and 7 show the wiping mechanism in action.




Figure 6. A piece of trapped debris




Figure 7. Debris being pushed away

Avoiding particle buildup

It has been well documented that particle buildup or debris inside the switch cavity can cause random failures and switching repeatability issues. These problems often occur when foreign particles migrate to the contact surfaces. The particles usually originate from two common sources:

- Contamination during manufacturing: Contamination debris that occurs during manufacturing can be minimized with a thorough cleaning process (e.g., ultrasonic) prior to the assembly on the RF microwave components.
- Material wear and tear: Particles are generated as two surfaces meet during movement of the jumper contacts. The resulting volume of particles depends on the amount of frictional force and the contact material's tendency to shed. Hence, during the design stage it is critical to consider the amount of surface contact at the jumper interfaces, the proper combination of materials and the plating profile.

A good example is Agilent's guide-rod design. Two thin push rods are used to mobilize each contact jumper. This design essentially inhibits any rotation or sliding occurrences. The elimination of additional centering rods coupled with the small surface areas of the two push rods further limits the possibility of particle buildup, in comparison to conventional designs.

Repeatability on measurement uncertainty

The repeatability of a switch has a direct effect on the measurement uncertainty of a test setup. Figure 8 shows an Agilent PNA network analyzer connected to a multi-port test set, which is used to test multiple devices. In this example, a total of three two-port devices can be tested simultaneously with any port. Because these errors are random rather than systematic, a root sum square (RSS) calculation is used to determine the total measurement uncertainty. Here, two scenarios are presented:

Scenario 1.

PNA repeatability = 0.01 dB, EM Switch repeatability = 0.03 dB
Total measurement uncertainty = $\sqrt{0.01^2 + 0.03^2 + 0.03^2} = 0.044$ dB

Scenario 2.

PNA repeatability = 0.01 dB, EM Switch repeatability = 0.1 dB
Total measurement uncertainty = $\sqrt{0.01^2 + 0.1^2 + 0.1^2} = 0.142$ dB




Figure 8. PNA network analyzer with a multi-port test set

The repeatability of the EM switch has a significant effect on the total measurement uncertainty of the system, in turn affecting the accuracy of all measurements. The wiping-action design of Agilent EM switches removes particle buildup to maintain a repeatability specification of 0.03 dB. This is crucial for ensuring a lower cost of test.

Conclusion

Operating life and switching repeatability are important criteria that must be considered when selecting an EM switch. Operating life is highly dependent on the jumper-contact mechanism, contact resistance, materials and plating used in all key RF components. As a result, these factors must be carefully considered during the design phase, and steps should be taken to ensure that they do not adversely affect operating life and switching repeatability. Doing so delivers the lower cost of test that today's high-volume component manufacturers demand.





Implementing Micro and Nano LC/MS Techniques for High Sensitivity Lipidomics Analysis

Giuseppe Astarita

*Project Scientist, University of California, Irvine
gastarit@uci.edu*

Hua Xiao

*Postdoctoral Fellow, University of California, Irvine
hxiao@uci.edu*

Faizy Ahmed

*Scientist, Agilent Technologies
faizy_ahmed@agilent.com*

Daniele Piomelli, Ph.D.

*Louise Turner Arnold Chair in the Neurosciences
Professor of Pharmacology and Biological Chemistry
Director, Center for Drug Discovery,
University of California, Irvine
piomelli@uci.edu*

After many years behind the scenes, lipids are making a spectacular comeback by being thrust into the limelight of biomedicine. This renewed excitement in lipids has been sparked by a series of convergent discoveries in the fields of biochemistry, cell biology and receptor pharmacology. Each of these discoveries is revealing the various roles played by lipid messengers and their receptors in health and disease. Most of these discoveries have been fueled by progress in techniques, such as mass spectrometry and atomic-force microscopy, which have opened experimental opportunities that were virtually unthinkable just ten years ago. Indeed, we are at a point today where a major goal is within our reach: profiling large-scale changes in lipid composition or determining the topographical distribution of individual lipid species in just a few cells.

The ability to profile the lipid composition of biological samples is important in disease diagnosis and drug discovery. However, constraints imposed both by the topological localization and small quantities of most biologically active lipids now demand the development of novel analytical approaches, and in particular, liquid chromatography mass spectrometry (LC/MS) techniques for high-sensitivity lipidomics.

The research work described in this article was performed in collaboration with the Department of Pharmacology, University of California Irvine (UCI). It was supported by an Agilent Technologies Foundation grant.

Novel analytical approaches to lipidomic analysis

In biological tissues, many bioactive lipids are present only in trace amounts. As a result, their measurement represents a major analytical challenge. Furthermore, most analytical approaches do not take into account the topological distribution of lipids in tissues and therefore often miss critical information. High-spatial resolution analyses of lipid metabolites would enable scientists to fully characterize the state and functionality of biological tissues. Topographical analysis by MS imaging, as detailed in the sidebar, has shown great promise in this area. However, the technique presents some limitations that must be addressed.

One of the challenges, for example, involves lipid metabolites of low abundance — which include many bioactive lipid messengers. These metabolites may remain undetected using the MS imaging technique due to suppression by more abundant constituents (e.g., phospholipids). Also, absolute analyte quantification is often difficult due to the lack of appropriate internal standardization. Consequently, the main analytical challenge facing MS imaging is its current lack of sensitivity and accurate quantification, both of which are important to the application of this technology to the understanding of neurological and psychiatric diseases.

We have begun to address these challenges working at the forefront of measurement technologies through Agilent's Program for University Research and with financial support from the Agilent Technologies Foundation. As described in the following paragraphs, the bioanalytical solution that we developed involves the use of laser microdissection (LMD), in combination with microfluidics LC coupled to MS, for high-spatial resolution and high-sensitivity quantification of trace amounts of lipid molecules in brain microstructures.

LC/MS system: Enabling lipidomic analyses of brain macrostructures

The hippocampus is a brain structure that plays a central role in memory and spatial navigation. Hippocampal degeneration is a key element in the pathogenesis of Alzheimer's disease. For the analysis of lipids in the hippocampus, a general workflow consists of slicing a frozen rat brain with a cryostat and then manually punching out hippocampal subregions (Figure 1a). The punches are homogenized in a small volume of cold methanol (0.3 ml) containing appropriate internal standards. The lipids are extracted using a mixture of chloroform and water (2:1; vol:vol). The organic layer is then collected and analyzed by LC/MS. In the example provided in Figure 1b, we separated multiple lipid classes using a small-particle-size column (1.8 μm). This allowed both high-throughput and high-resolution analysis.

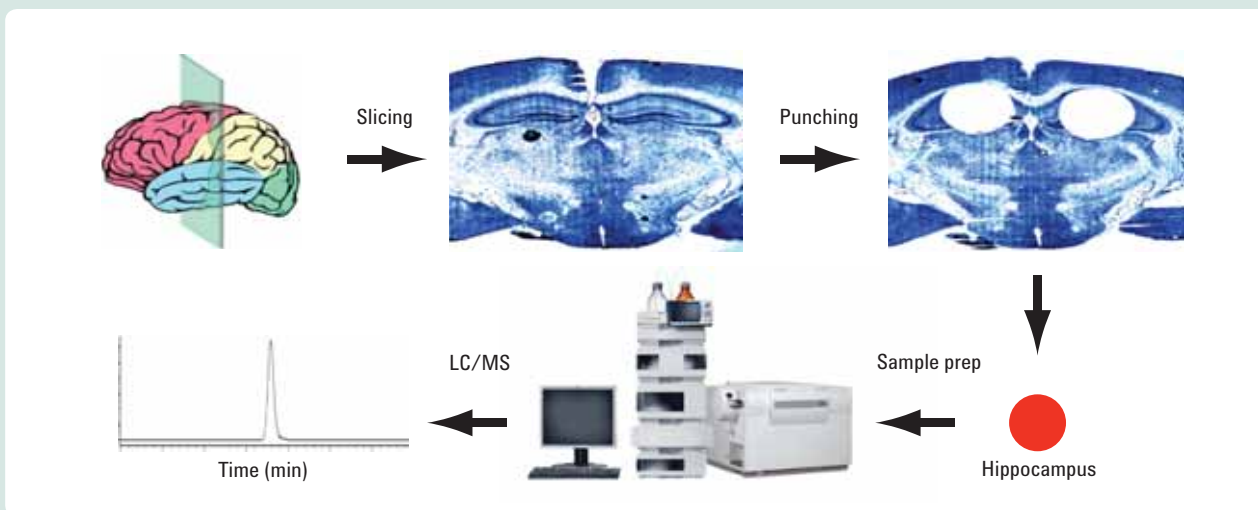


Figure 1a. The analysis of a brain macrostructure follows this workflow schematic.




Figure 1b. This graphic illustrates a representative LC/MS chromatogram for the analysis of lipids in brain macrostructure (left) and linearity of the instrument response to different concentration of fatty-acid ethanolamides (right). The method was linear from 0.2 pmol up to 75 pmol of fatty-acid derivatives.

LC conditions (Agilent 1100):

Column: Zorbax XDB Eclipse C18 (50 x 4.6 mm i.d., 1.8 μ m)

Mobile phase:

A= water

B= methanol

Gradient:

85% B at 0 min

90% B at 2 min

100% B at 3 min

Column temperature: 40° C

Flow rate: 1.5 ml/min

Injection volume (10 μ l)

MS conditions (Agilent 1946D):

Ionization mode: electrospray

Ionization polarity: positive

Capillary voltage: 3000 V

Fragmentor voltage: 120 V

Drying gas (N₂): 13 liters/min

Drying gas temperature: 350° C

Nebulizer pressure: 80 PSI

Developing a microflow LC/MS system

A physiological analysis of the hippocampus shows that neural information flows through a series of anatomical subregions — from the dentate gyrus to cornu ammonis 3 (CA3) and finally to CA1. Alterations in the lipid composition of each of these structures may contribute to hippocampal function in health and disease.

To define in greater detail the lipid composition of select hippocampal structures, we developed a high-sensitivity methodology for lipid analysis. Using an optimized-microflow LC/MS system, we were able to quantify as low as 20 fmol of fatty-acid derivatives (Figure 2).

Applying the microflow LC/MS system

To determine the endogenous lipid distribution within selected hippocampus microstructures, we laser-microdissected discrete regions of interest from 30- μm thick rat brain slices fixed in 4 percent paraformaldehyde. Microdissection was accomplished using the Leica LMD6000 laser microdissection system.

Under direct microscopic visualization, minute areas were selected on the computer screen and isolated with a laser beam as shown in Figure 3. The regions of interest were collected in the cap of a microcentrifuge tube. The cap was filled with acetonitrile (10 μl) containing appropriate internal standards to both inactivate metabolic enzymes during sample collection and

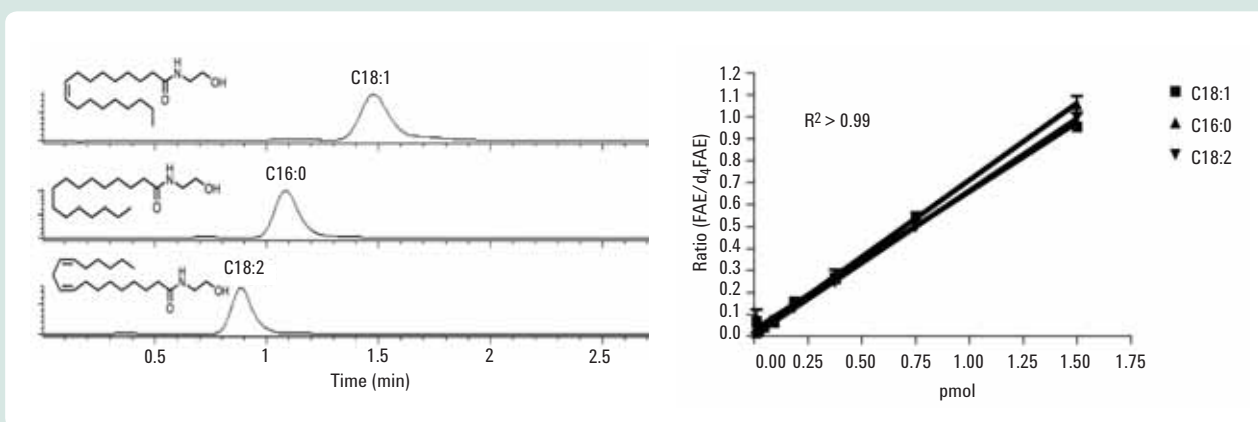


Figure 2. These representative LC/MS chromatograms of fatty-acid ethanolamides were obtained using a microflow LC/MS system (left). The method allowed for linear quantification from 0.02 pmol up to 4 pmol (right).

LC conditions (Agilent 1100):

Column: Zorbax XDB Eclipse C18 (50 x 0.5 mm i.d., 1.8 μm)

Mobile phase:

A= water

B= methanol

Isocratic for 10 min :

10% A

90% B

Column temperature: 30° C

Flow rate: 20 $\mu\text{l}/\text{min}$

Injection volume (0.2 μl)

MS conditions (Agilent 1946D):

Ionization mode: capillary-electrospray

Nebulizer: Microflow (50 μm i.d. inner needle).

Nospacer installed.

Ionization polarity: positive

Capillary voltage: 3500 V

Fragmentor voltage: 120 V

Drying gas (N_2): 13 liters/min

Drying gas temperature: 325° C

Nebulizer pressure: 30 PSI

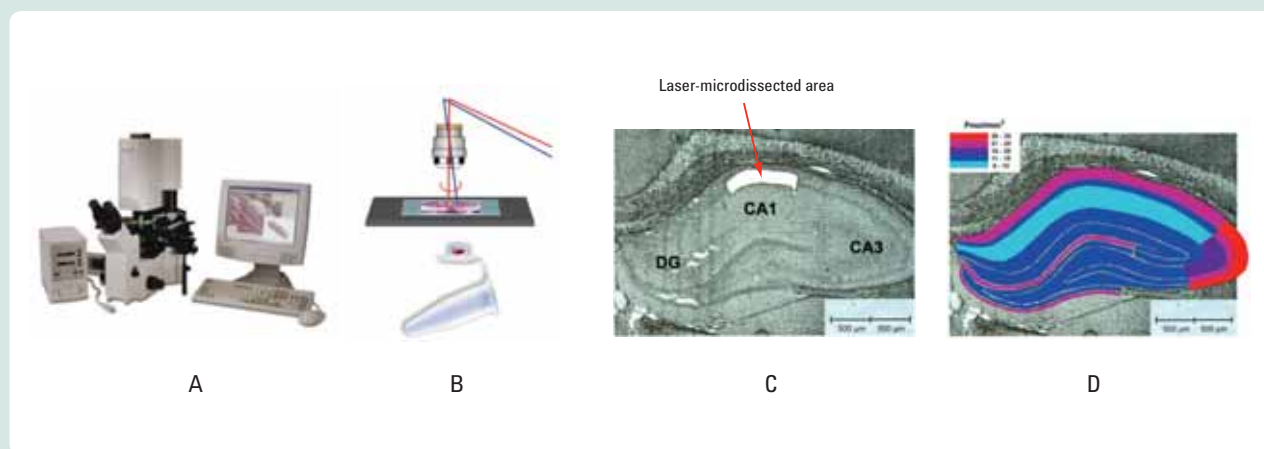


Figure 3. This topological analysis of 2-AG distribution in different layers of rat hippocampus was created using a laser microdissection system (A). Selected microstructures from hippocampus were laser-cut under the microscope (B and C) and the microstructure was then analyzed using the newly developed microflow LC/MS system (D).

LC conditions (Agilent 1100):

Enrichment column: Zorbax XDB Eclipse C8 (35 x 0.3 mm i.d., 3.5 μ m)

Analytical column: Zorbax XDB Eclipse C18 (50 x 0.075 mm i.d., 3.5 μ m)

Pumps mobile phases:

A= 5 mM NH_4Ac and 0.25% HAc in water

B= 5 mM NH_4Ac and 0.25% HAc in ACN

Capillary-pump:

Isocratic (2% B), flow rate: 4 μ l/min

Gradient nano-pump:

2% B at 4 min

100% B at 10 min

Flow rate: 1 μ l/min

Columns temperature: N/A

Enrichment column switch: 4 min

Injection volume (1 μ l)

MS conditions (Agilent 1946D):

Ionization mode: nanoelectrospray

Ionization polarity: positive

Capillary voltage: 1900 V

Fragmentor voltage: 70 V

Drying gas (N_2): 4 liters/min

Drying gas temperature: 320 $^\circ$ C

Nebulizer pressure: N/A

accurately quantify the lipid analytes. The samples were sonicated and concentrated in a final volume of 2 μ l before being analyzed using the newly developed microflow LC/MS system. As an example, the spatial distribution within the hippocampus of the bioactive lipid 2-arachidonoyl-sn-glycerol (2-AG) — an important ligand for cannabinoid receptors — is represented in Figure 3 as a bidimensional pseudocolor map.

Developing a nanoflow LC/MS platform

The lower limit of the tissue size that can be collected by LMD depends on the magnification and minimal width of the laser beam, which is in the range of 1 μ m. In principle, LMD can be used to isolate a few and even single cells. Therefore, in order to analyze trace amount of lipids in increasingly small brain nanostructures we developed an ultrahigh-sensitivity nanospray LC/MS method. This type of platform is often used in proteomics and can be adapted to the analysis of lipid metabolites.

We started our method development using a single custom-made column (Zorbax XDB C18, 50 x 0.075 mm i.d., 3.5 μ m) and interfaced the LC/MS system with a nanospray source. Because of the small column volume, we were required to inject no more than 0.05 μ l of sample to maintain a good peak resolution.

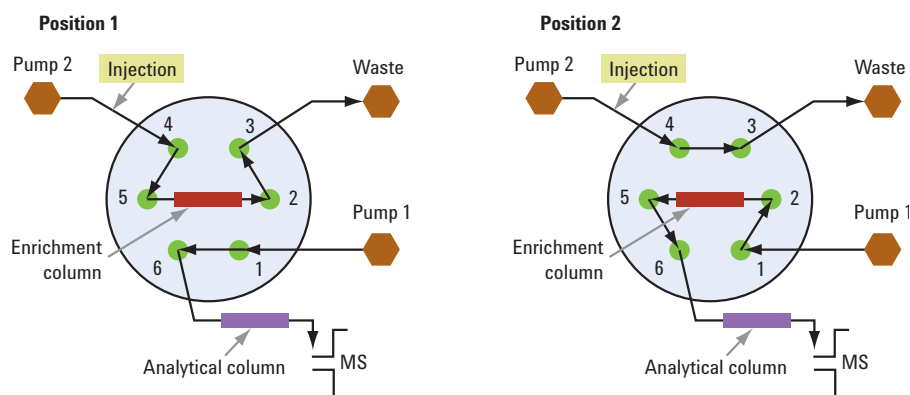


Figure 4a. This schematic diagram illustrates the flow path of the Agilent nanoflow lipidomics solution. In the first step of separation, the sample is loaded onto a C8 reversed-phase column and lipids that bind to the column are washed from impurities and concentrated (Position 1, left). In the second step, the enrichment column is switched into the solvent path of the nano-pump and back flushed onto the C18 nano-column on the right (Position 2, right). An increasing concentration of organic solvent elutes the concentrated sample and further separation is achieved onto the analytical reversed-phase column. The analytical column effluent sprays into the nano-electrospray source of the MS detector.

Lipids: The basics

Lipids are absolutely essential for life, playing diverse and important roles in nutrition and health. In the brain, lipids are primarily constituted of fatty acids, fatty-acid derivatives (e.g., glycerolipids, phospholipids and sphingolipids), and substances that are biosynthetically related to fatty acids (e.g., cholesterol and its derivatives). Lipids constitute the building blocks of neuronal and glial membranes, insulate nerve fibers and act as a source of energy. They also serve as a biosynthetic precursor to cellular messengers that control a diversity of brain functions. Changes in the levels of structural and signaling lipid species have been linked to the pathogenesis of several neuropsychiatric diseases, including Alzheimer's disease and schizophrenia.

The process of measuring lipids is called lipidomics. More specifically, lipidomics is defined as the large-scale analysis of lipid profiles in cells and tissues. This type of analysis provides an invaluable snapshot of brain-function physiology and can be used to identify and characterize both physiological and pathological states of neuronal and glial cells.

The complex nature of lipid signaling and its strict dependence on specific membrane localization imposes structural and topological constraints that make lipidomic analysis a challenging prospect. Currently this analysis can be tackled with one of three main techniques:

- In electrospray ionization mass spectrometry, a solution containing the analyte is sprayed through a needle, forming charged droplets. After desolvation, these droplets explode, forming ions that are detected by MS. Single-stage MS separates the ions according to their mass-to-charge ratio (m/z). Tandem MS and multi-stage MS (MS^n) allow multiple fragmentations to occur, generating fragments that are used to identify and quantify specific analytes in complex mixtures.
- In desorption ionization MS techniques (e.g., matrix-assisted laser desorption ionization MS (MALDI), cluster-secondary ion MS (SIMS) and desorption electrospray ionization (DESI)) the process is initiated by irradiating a defined spot on the solid-state sample with a focused excitatory beam (e.g., laser, ions and charged droplets of solvents). Upon impact, the sample surface releases a vapor of ionized molecules, which are detected by MS and used to create a topographic map of a specific analyte. Hence, these techniques are named MS imaging.
- In atomic-force microscopy, a tip connected to a cantilever is brought into close proximity to a membrane. The force between the tip and the membrane causes a deflection of the cantilever, which is measured by a laser detector (photodetector). As the tip scans across the membrane, system software creates a topographic map of the surface.




Figure 4b. The nanoflow LC/MS system, configured to perform LC/MS chromatogram for 2-AG, is able to reach a limit of quantification of 0.005 pmol.

To inject a larger volume and concentrate the sample injected, we added an enrichment column (Zorbax XDB C8, 35 x 0.3 mm i.d., 3.5 μ m) to the system before our analytical column. A valve was also added to switch the flows between columns as shown in Figure 4a. As configured, the system is able to provide sensitivity as low as 5 fmol for fatty-acid derivatives (Figure 4b), which represents a remarkable improvement over previous LC/MS analytical methodologies.

The future of lipidomic analysis

The initial results of our research and development should help future development of microfluidic-based LC-Chip-MS platforms for the analysis of lipid metabolites in complex biological matrices. Such platforms will allow researchers to integrate most of the nanoflow LC/MS components into a single microfluidic chip, thus avoiding the problems associated with capillary connections and the need to keep the system free of leaks, blockage and excessive dead volume. This technology will facilitate high-throughput and large-scale analyses of lipids in biological samples, adding a higher degree of sensitivity. Integrated LC-Chip-MS platforms may also be applied to comparative lipidomic analyses of very small amounts of neural tissue or other biological materials (e.g., skin, hair, saliva and sweat) that may be used for diagnostic purposes.

The authors would like to thank Jennifer Geaga, graduate student, University of California, Irvine, and Jason Clapper, graduate student, University of California, Irvine, for their contributions to this article.

A dream fulfilled

Faizy Ahmed

I joined Agilent in 2001 as a product specialist for consumable products, focused on LC and GC column products. My journey to UCI began in the spring of 2005, when I pursued an R&D position within Agilent. For family reasons, however, I could not relocate to Wilmington, Delaware, where the position was located.

Fortunately, the hiring manager, Bill Barber, suggested that I investigate collaborative opportunities near my home. Luckily, in January 2005 I had presented a technical seminar that was well attended by academics and industrial scientists. One attendee was Dr. Giuseppe Astarita, a project scientist at UCI. Through Dr. Astarita, I met Professor Daniele Piomelli and was able to find many productive collaborations in lipidomics research. Since 2006, I have recommended UCI for two Agilent Technologies Foundation grants and have made academic contributions to the research by serving as a mentor.

Thanks to the flexibility of Agilent management, today I enjoy the best of both worlds — contribution to academic and industrial research, as well as the opportunity to develop products for Agilent.

Contact Agilent

Email Update Service

Get the latest information and newest test skills

Subscribe to our free Email Update Service for the latest Agilent product, support and application information — customized for your interests and preferences.

Only the information you want, when you want it

Easily identify your specific interests from a selection of more than 100 product and application areas and specify the types of information you would like to receive:

- **Support information** such as new-product information and specifications; application notes; tradeshow, events and live or virtual seminars.
- **What's New**, including education/training and consulting services; firmware and software upgrades; and warranty, calibration and repair.

Total commitment to privacy

Agilent is proud to participate in the Better Business Bureau On-Line Privacy Program, and we adhere to all of the program's standards.



Agilent Email Updates

www.agilent.com/find/emailupdates



Agilent Test and Measurement RSS Feeds

Subscribing to Agilent's RSS feeds lets you receive Agilent news as soon as it is published. With over 50 topics — ranging from "AC Power Sources/Power Analyzers" to "Wireless Test Systems" — you may find several feeds of interest.



www.agilent.com/find/rss

Additional Agilent Resources

Agilent Direct Store

- Quickly build and place your order, 24x7
- Easily select and configure your products
- Share carts within your company (e.g., technical team and purchasing)
- Flexible payment options
- Secure transactions
- Order basic instruments, general purpose test accessories, microwave test accessories and replacement parts online today (available to U.S. customers only)



Agilent Direct

www.agilent.com/find/onlinestore

Infoline

Your online source for Agilent repair and calibration information



Infoline

www.agilent.com/find/service

Agilent CertiPrime

Your source for Agilent CertiPrime premium used equipment



www.agilent.com/find/used



Agilent Technologies