Vol. 7 • Issue 2

2025

Signal Integrity Journal™

Signal Integrity • Power Integrity • EMC/EMI

Cover Feature: p. 10

Optimization of IBIS-AMI Model Parameters with Machine Learning Algorithms

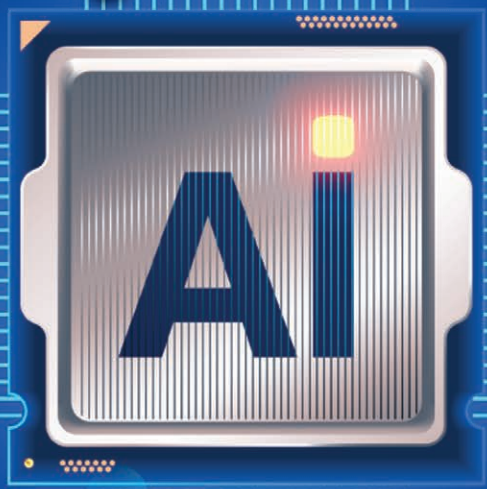

horizon house®

Founded in 1958

signalintegrityjournal.com

Si-FLY® HD

CPX CO-PACKAGED 224G SOLUTIONS

Samtec's Si-Fly® HD co-packaged cable system (CPX) is the highest density 224 Gbps PAM4 solution on the market. Co-packaged connectivity eliminates insertion losses in the BGA breakout and PCB, and enables a complete passive DAC 224G channel that is both extreme performance and low cost.

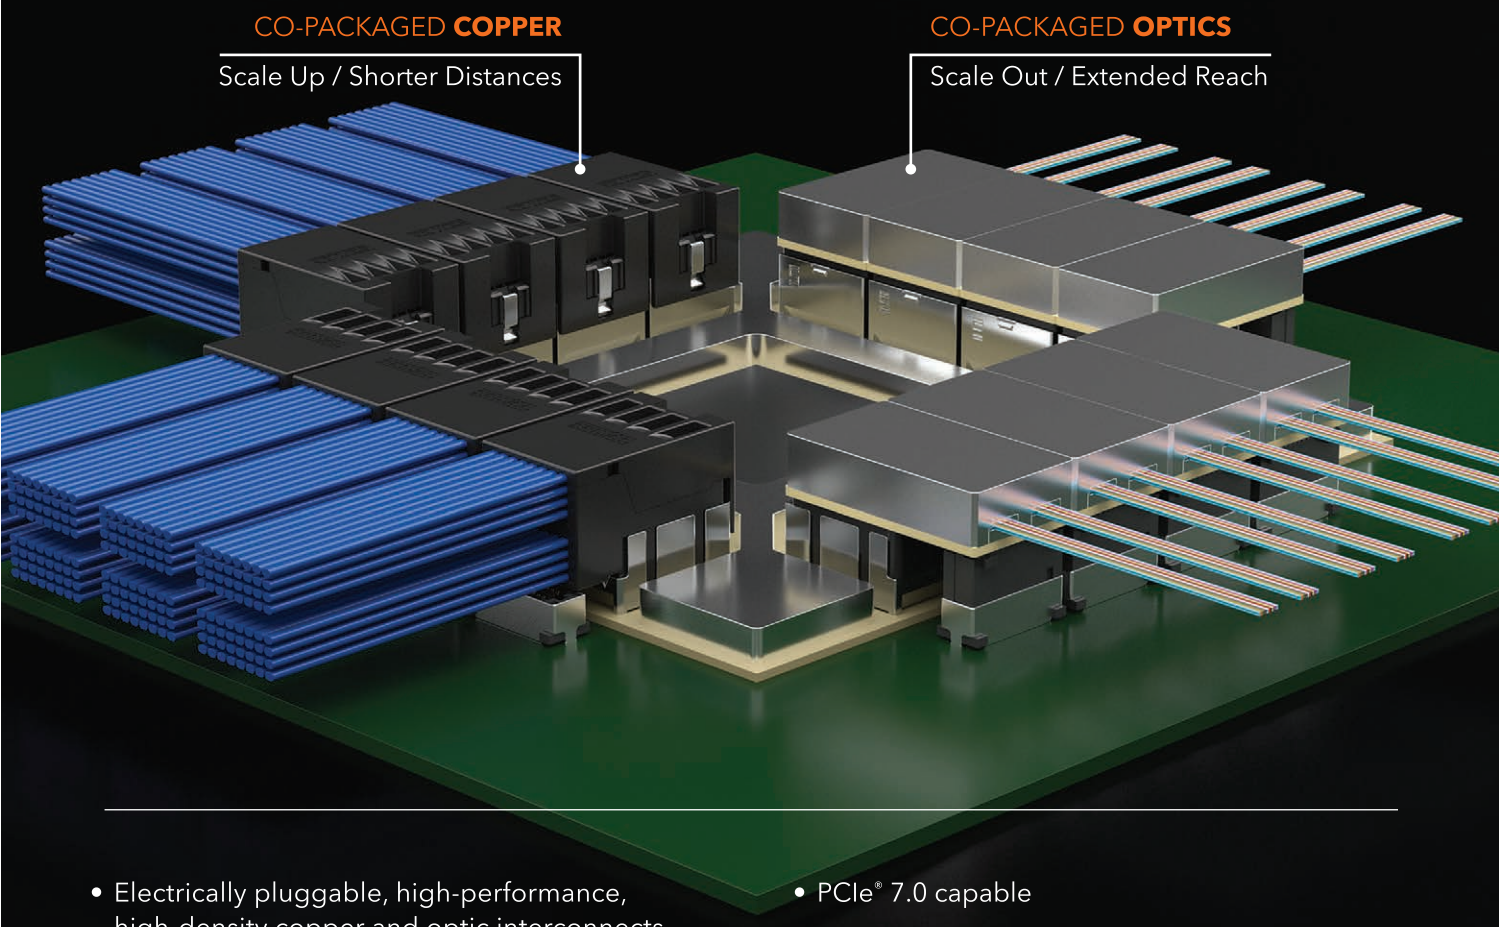

- Electrically pluggable, high-performance, high-density copper and optic interconnects on a 95 mm x 95 mm or smaller substrate
- Extreme density:170 differential pairs/sq. inch
- Samtec Eye Speed® hyper low skew twinax cable with 1.75 ps/m maximum intra-pair skew
- PCIe® 7.0 capable
- Surface mount reflow technology to the substrate
- Near-chip (ASIC adjacent) system also available
- To learn more, please visit **samtec.com/si-fly-hd** or contact **HDR@samtec.com**

samtec

www.samtec.com

TEST & MEASUREMENT PRODUCT GUIDE

Available for download and print

SCAN ME

Get our new T&M product guide for a complete overview of our portfolio, specifications and product models from Rohde & Schwarz. We offer a broad range of instruments to meet both your everyday needs and your most challenging requirements. Learn more about what's possible when faced with tomorrow's challenges today and prepare for what's next!

rohde-schwarz.com/discover

ROHDE&SCHWARZ
Make ideas real

TABLE OF CONTENTS

Editor's Note

6 The Impact of AI at the Singularity

Eric Bogatin, Technical Editor, Signal Integrity Journal

Cover Feature

10 Optimization of IBIS-AMI Model Parameters with Machine Learning Algorithms

Jared James and Ambrish Varma, Ph.D., Cadence Design Systems

Technical Features

16 Seeing Through the Noise: Reliable Power Rail Measurements in High-Current AI Systems

Seamus Brokaw, Tektronix, and Steve Sandler, Picotest

20 ManyPoint Networks: A System Co-Design Framework for 448 Gbps AI Fabrics and Beyond

Andrew Josephson, Samtec

26 High-Speed Digital Interface Characterization Requires New Test Approach

Hiroshi Goto, Anritsu Company

30 Using Ultra-Broadband Baluns to Perform Differential S-Parameter Measurements Using Single-Ended 2-Port VNA

Madrone Coopwood, Jason Yoho Ph.D., and James R. Andrews Ph.D., Hyperlabs

36 Material-Induced Skew in High-Speed Multilayer PCBs: Influences and Mitigation Strategies

Bob Nurmi, John Coonrod, and Vitali Judin, Rogers Corporation

42 The Imperfect Via: The Rough Truth Lurks Beneath the Surface

Bert Simonovich, Lamsim Enterprises

Follow Us @SIJupdates

Become a Fan at facebook.com/signalintegrityjournal/

Join the Signal & Power Integrity Community

Staff

TELEDYNE LECROY
Everywhereyoulook™

Test in Both Frequency and Time Domains

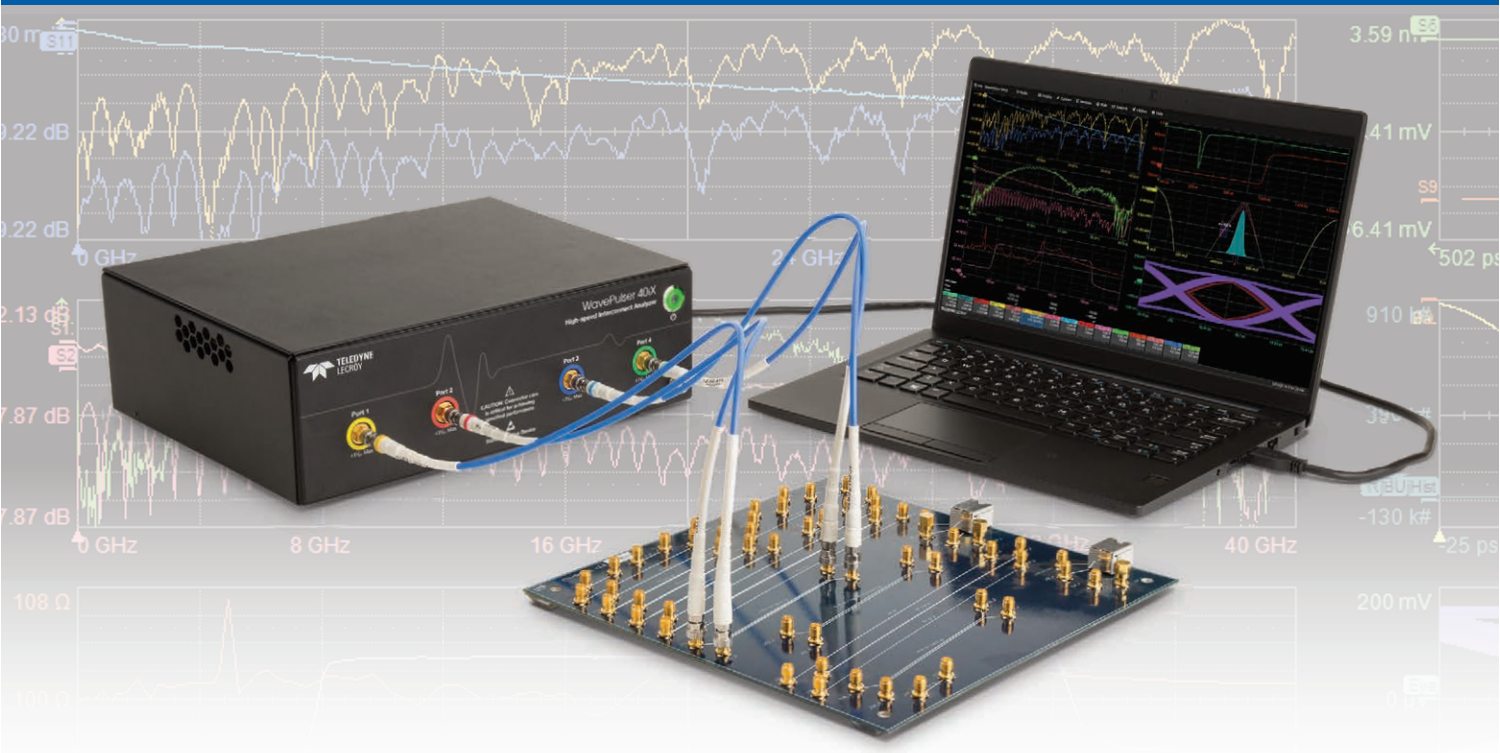

WavePulser 40iX

High-speed Interconnect Analyzer

- Measures S-parameters (like a VNA)
- Measures impedance profile (like a TDR)
- De-embedding, simulation, emulation, and time gating of results
- No calibration required
- Costs a fraction of the price of a VNA

teledynelecroy.com/wavepulser

The Impact of AI at the Singularity

Eric Bogatin, Technical Editor
Signal Integrity Journal

In 2005, noted futurist Ray Kurzweil published his book *The Singularity is Near* in which he applied his Law of Accelerating Returns to extrapolate computational performance.

The Law of Accelerating Returns applies Moore's law to all technological advances. By whichever metric selected, the performance of technology increases exponentially over time rather than linearly. This is a consequence of the differential equation that technology advances obey. The rate of advancement increase in each generation is proportional to the previous generation of technology. This is a linear first-order differential equation, and the solution is an exponential. Kurzweil offers numerous examples of the exponential growth of technologies.

In particular, he plotted the computation ability that $1000 buys you at its current value. He demonstrated that over the last 120 years, it has had a doubling time of about 1-2 years. His plot from 2005 is shown in ***Figure 1***.[1]

In his book, Kurzweil writes that "by around 2020," $1000 will buy computer power equal to a single brain. He also states that by 2045, the onset of the singularity, the same amount of money will buy one billion times more power than all human brains combined today.[2]

He calls the singularity the merger of human and machine intelligence.

Fast-forward to present day, 2025. An LLM running on a $1000 desktop computer is not so far off from Kurzweil's prediction of human brain-level computing performance. Take, for example, a typical hyperscale data center with 10,000 GPUs on which the LLMs are developed; we are nearing Kurzweil's singularity.

What does this mean? What is the consequence of the combination of AI software running on hardware optimized for AI processing?

As engineers, we interact with AI systems in three ways: as developers of the hardware and software platforms, as users of the new paradigm of AI-driven EDA design tools, and as consumers of AI tools in our daily lives.

Twenty years ago, the killer app that drove the need for speed was the transmission of video information over the internet. Netflix and YouTube created the internet. Today, AI data centers drive the need for speed and interconnect density. Heterogenous packaging is finally in production. Optical interconnects from board to board are also in production. Power distribution and thermal management for 50 kW processor boards are driving the emergence of new technology revolutions.

Agentic AI models, autonomous agents that learn hardware design on their own using EM simulation outputs, are being offered commercially. There seems to be a consensus that these tools can currently replace the junior engineer. The debate is not whether these tools will replace the engineering judgement of senior engineers, but when they will be capable of replacing senior engineers. If they do replace the junior engineer, where will the next generation of senior engineers come from?

All aspects of business and work life are currently impacted by the proliferation of AI tools. Education is not exempt from the revolutionary change brought about by AI. For example, Khanmigo displays the benefit of free personalized AI tutors in any subject for childern of any age, anywhere in the world.

We are in the early chaotic revolution of AI tools; it is used by faculty to grade content and by students to create content. Many institutions have yet to regulate the use of AI in academics, meaning anything goes.

Where lies the future of AI in education? I am reminded of an old joke I heard from another professor 50 years ago. It goes like this:

A professor realizes one day that he gives the same lectures year after year. He decides to record his lectures. The following year, he comes to class and turns on his tape recorder, and the class listens to his recorded lecture. It goes so well that he starts leaving class after five

iNRCORE®
FAMILY OF BRANDS

You CAN'T Build Without iNRCORE

SIGNAL
Ethernet, Transceiver Modules, MIL-STD-1553, Data Bus Couplers, Baluns, Delay Lines

CAPACITORS
High-Q/Low ESR Multi-Layer Ceramic, Broadband, Single Layer, Thin Film, Variable

FILTERS
Bandpass, High Pass, Low Pass, Notch, EMI/RFI, Bias Tee, Amplitude Equalizer, Diplexers/Multiplexers

INDUCTORS
Conical, Ceramic Chip, High Frequency, High Q, Planar, Toroidal, CMC, THT, SMT, MIL-STD-981, MIL-PRF-27

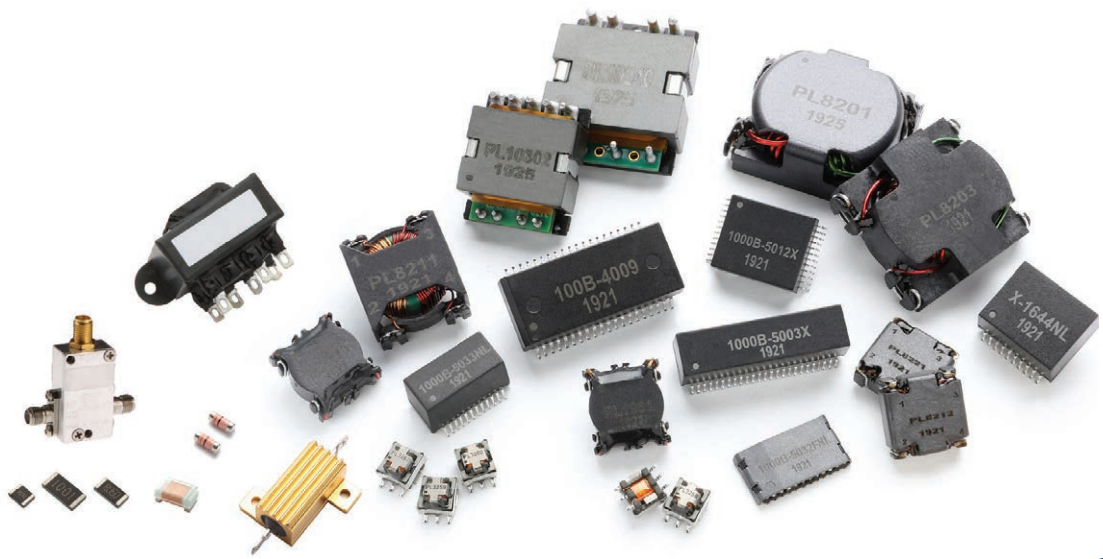

For more information
SCAN HERE!

ISO 9001 CERTIFIED
AS9100 CERTIFIED
ITAR Registered

www.iNRCORE.com +1-215-781-6400

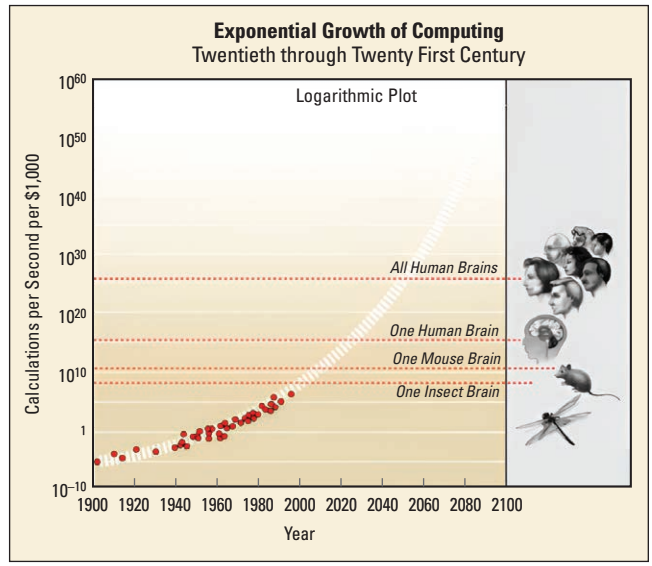

Fig. 1 The information processing ability that $1000 buys in current dollars increases exponentially.

minutes and just lets the recorder play.

One day, he decides to check on his class and comes back ten minutes before the end of the lecture. When he walks into the classroom, he sees a tape recorder on every desk, left by each student to record his recording.

We are now witnessing one outcome of the law of unintended consequences. Students create content using AI, which is then graded by professors using AI. Emails written and sent by AI are read and responded to by AI. Soon, an AI processor board for an AI-generated application will be designed by AI, verified by AI, and fabricated by an AI-driven assembly line. Where is the human in these loops?

In 2024, Kurzweil published *The Singularity is Nearer,* a sequel in which he paints a beautiful picture of all the new opportunities we can expect from the AI revolution: longer lifespans using designer genes, higher quality of life from new materials, instant communication, and safer travel.[3]

Before cars replaced horse-drawn carriages, who could have anticipated the problems of car accidents, rush-hour traffic, air pollution, and parking? Kurzweil has a great track record for his predictions coming true, but there are alternative visions of the future. Just look at some of the movies from the last 60 years, such as "I, Robot," "The Terminator" series, "The Matrix" series, "WarGames," or "Colossus: The Forbin Project." In all of these narratives, when humans were taken out of the loop, it did not end well for them.

As developers of the next generation of technology, we should always keep in mind that technology amplifies human nature. Let's pay attention to the unintended consequences of what we create. ■

REFERENCES

1. R. Kurzweil and Kurzweil Technologies, Inc., "PPTExponentialGrowthof-Computing.jpg," Wikimedia Commons, January 2008, commons.wikimedia.org/w/index.php?curid=3324354.
2. R. Kurzweil, *The Singularity is Near: When Humans Transcend Biology*, Viking, 2005.
3. R. Kurzweil, *The Singularity is Nearer: When We Merge with AI, Viking*, 2024.

KEEP YOUR EYES WIDE OPEN™

WITH

HYPERLABS

HYPERLABS HAS BEEN INNOVATING ULTRA-BROADBAND PRODUCTS TO 110GHz AND BEYOND, SINCE 1992.

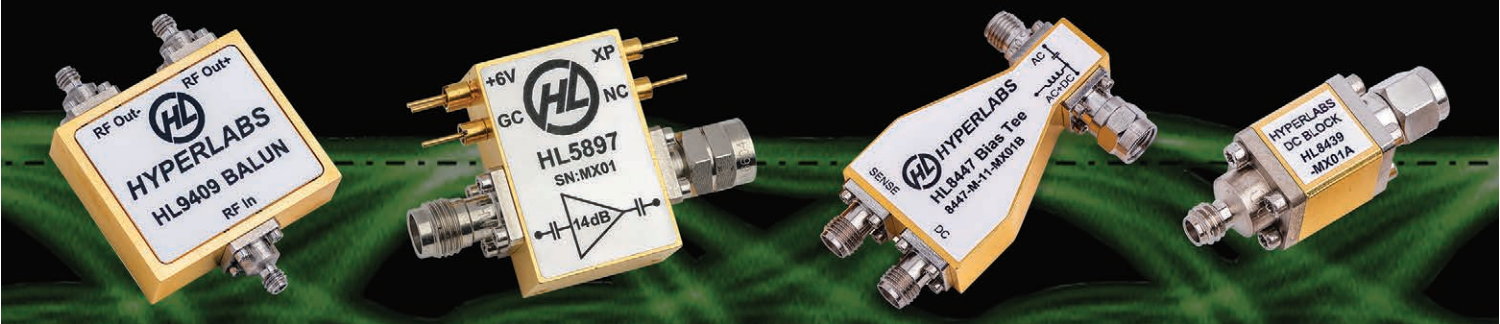

AMPLIFIERS, BALUNS, POWER DIVIDERS, DC BLOCKS, BIAS TEES, PICKOFF TEES, AND SO MUCH MORE...

HYPERLABS FOCUSES ON OPTIMAL SIGNAL INTEGRITY TO HELP YOU MAINTAIN THOSE CLEAN EYES THAT YOU'VE WORKED SO HARD TO CREATE WHEN EVERY BIT COUNTS!

HYPERLABS

www.HYPERLABS.com
+1.720.497.6538

Optimization of IBIS-AMI Model Parameters with Machine Learning Algorithms

Jared James and Ambrish Varma, Ph.D.
Cadence Design Systems, Santa Clara, Calif.

Serial link speeds have increased 25X in under 20 years, thus increasing the complexity of the IBIS-Algorithmic Modeling Interface (AMI) models used in simulating these links. With the increased speed and complexity of designs, it is crucial to analyze channels to ensure sufficient margin for error-free data transmission. An exhaustive manual search method is typically used to find the best set of parameters for a given channel, but given the increased number of model parameters and ranges, this approach can quickly become computationally expensive, even with parallel execution.

Machine learning (ML) techniques[1] have proven effective in modeling complex systems with numerous interacting components and nonlinear relationships. Some of these techniques can be employed to optimize the parameters of a complex system more efficiently than the exhaustive search method typically used in serial link simulations.

This article describes the use of Cadence's Sigrity™ signal and power integrity solution ML optimization algorithm to quickly and efficiently converge on the best set of parameters in a set of IBIS-AMI models. The application of Sigrity was investigated for refining IBIS-AMI parameters to find the optimal set of values to maximize a specific metric.

Reference Designs

As the speeds of serial link standards have increased, so has the complexity of the designs needed to transmit and receive data. This is reflected in the amount of transmitter (Tx) and receiver (Rx) equalization specified in the standards. This is needed to overcome the increased channel losses that come with the faster data rates.

In the past, IBIS-AMI models based on the Electronics Industries Association serial standards, called "reference designs," have swept all possible Tx and Rx parameter combinations to find the best solution for a given channel. While this was a reasonable approach in the past, this method has become more difficult to execute. Even with many cores and tools to manage the sweeping of these parameters, it has become computationally expensive and very time consuming.

Simulations can be run more efficiently using statistical methodology. However, these types of models lack nonlinear effects such as noise and signal clipping that can give less than optimal results. As a result, time domain channel simulation was used for this study, and a possibly more computationally efficient approach was used that applied an ML optimi-

zation algorithm to find the best set of IBIS-AMI model parameters for a given channel.

ML Optimization

The ML optimization process or methodology can be done in many ways. One method is to execute an exhaustive search that involves a brute force algorithm that systematically enumerates all possible solutions to a problem and checks each one to see if it is a valid solution. This algorithm is typically used for problems that have a small and well-defined search space, where it is feasible to check all possible solutions.[2] However, for cases where the search space is extremely large, a different approach is needed. This is where ML can be used to more efficiently find an optimal solution.

There are areas in the ML field that specifically deal with the optimization of parameters of a function. This function is often called an objective function, which is an unknown function (sometimes called a "black box"). The goal is to find the set of parameters that will provide the maximum (or minimum) value of the objective function. A model of this unknown objective function (sometimes called a surrogate model) is built and updated based on evaluations of the objective function. The update of the surrogate model is based on an acquisition function. While there are several types of surrogate models and acquisition functions from which to choose, the most common surrogate function is a Gaussian process model, and the most common acquisition function is the expected improvement.

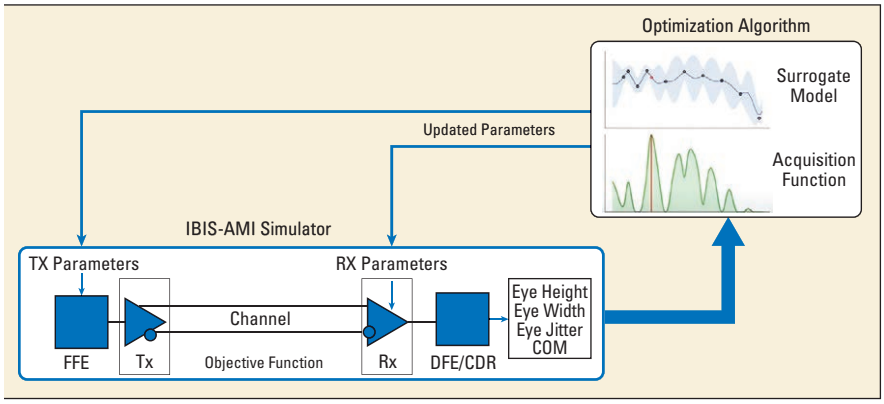

Fig. 1 **Machine learning optimization process.**

This optimization process (see ***Figure 1***) can be detailed as follows:[3]

1. An initial random sampling is collected by applying random parameter values to the objective functions. The size of this sampling can be scaled based on the number of parameters.
2. The surrogate model is trained on the initial random sampling.
3. Create updated sets of parameters that would:
   a. Move closer to the maximum (or minimum) of the acquisition function
   b. Further random sampling of the acquisition function.
4. Evaluate the objective function based on the updated parameters from the surrogate model.
5. Update the surrogate model based on the latest samples of the objective function.
6. Repeat steps 3-5 until a stopping criterion has been met.

When creating the updated set of parameters in Step 3, there is a tradeoff between the number of simulations that will exploit the acquisition function, called exploitation, and the number of simulations that will explore the acquisition functions, called exploration. This is an important tradeoff because too much exploitation will result in possibly not finding the best set of parameters, and too much exploration will not spend enough time working towards the best set of parameters. This tradeoff is sometimes called the "exploration-exploitation dilemma." Typically, more exploration is done early in the optimization process before it switches to more exploitation towards the end of the process.

For this article, the objective function was the simulation of a specific

TABLE 1

112G-LR COM PARAMETERS

Transmitter equalizer, minimum cursor coefficient	***c(0)***	***0.50***	—
Transmitter equalizer, third pre-cursor coefficient Minimum value Maximum value Step size	c(-3)	 -0.06 0 0.02	 —
Transmitter equalizer, second pre-cursor coefficient Minimum value Maximum value Step size	c(-2)	 0 0.12 0.02	 — — —
Transmitter equalizer, first pre-cursor coefficient Minimum value Maximum value Step size	c(-1)	 -0.34 0 0.02	 — — —
Transmitter equalizer, post-cursor coefficient Minimum value Maximum value Step size	c(1)	 -0.2 0 0.02	 — — —
Continuous time filter, DC gain Minimum value Maximum value Step size	g_{DC}	 -20 -2 1	 dB dB dB
Continuous time filter, DC gain2 Minimum value Maximum value Step size	g_{DC2}	 -6 0 1	 dB dB dB
Continuous time filter, scaled zero frequency	f_z	f_b/2.5	GHz
Continuous time filter, pole frequencies	f_{p1} f_{p2}	f_b/2.5 f_b	GHz GHz
Continuous time filter, low frequency pole/scaled zero	f_{LF}	f_b/80	GHz

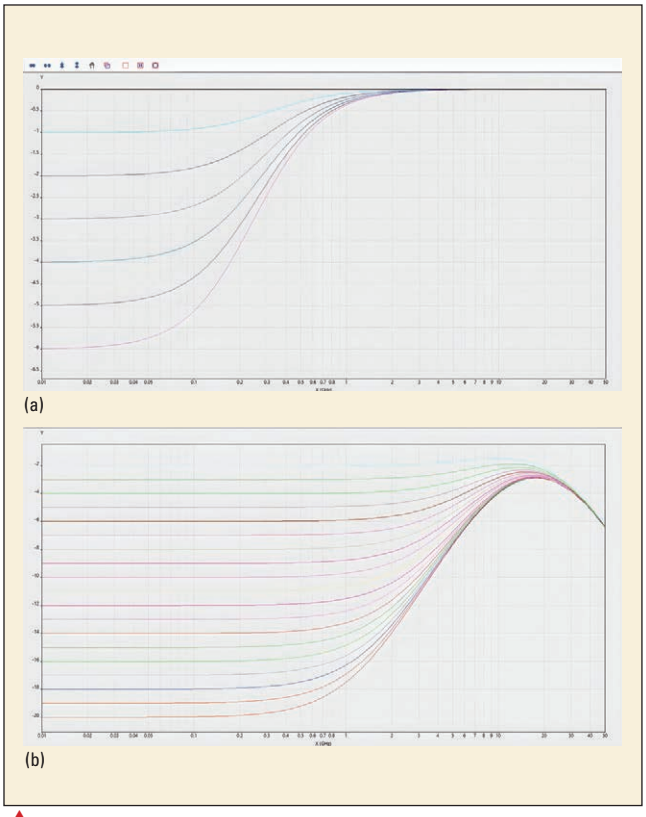

Fig. 2 **Bode plots of the two stages of CTLE.**

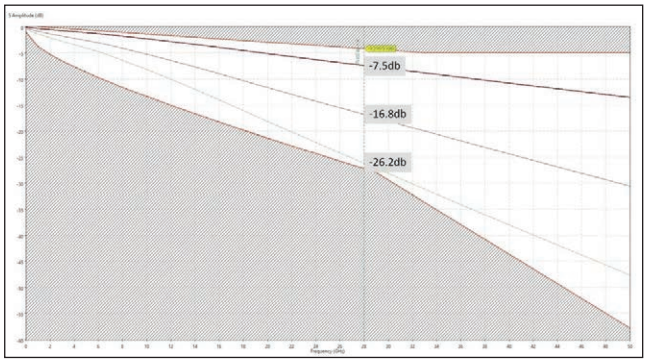

Fig. 4 **Insertion loss mask.**

channel with IBIS-AMI models at the Tx and Rx that, together, would have many equalization parameters. From the simulation, an output needed to be chosen for the surrogate model. Typical outputs for this type of simulation would include eye height, eye width, eye jitter, and channel operation margin (COM). A single output, or a combination of several outputs, could be used as an output to maximize or minimize. For most cases, this would be reliable but, as will be discussed in a later section, there are situations where these outputs would not work, and a different measurement would be needed.

IBIS-AMI Model Generation

For this study, heavily parametrized IBIS-AMI models were needed to leverage the Sigrity ML optimization. A model based on the OIF-CEI-112G-LR standard was selected due to its large number of parameter combinations, shown in ***Table 1***.[4]

The Tx AMI model has three taps of precursor and

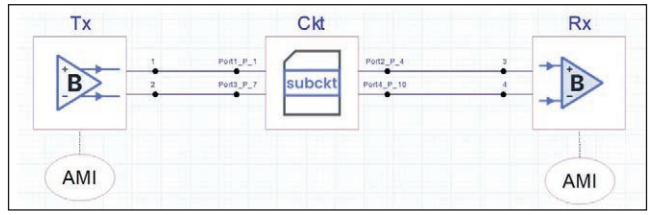

Fig. 3 **Testbench schematic.**

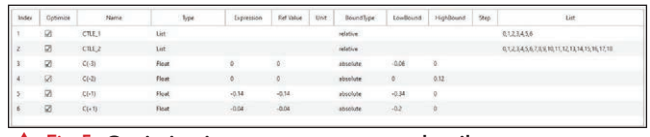

Fig. 5 **Optimization parameter setup details.**

one tap of post-cursor, with the ranges shown in Table 1. Note, not all combinations of values are valid for the Tx, as the main cursor cannot be lower than 0.5. This issue, the impact to the optimization algorithm, and the work-around are discussed in a later section.

The Rx AMI model is a two-stage continuous time linear equalization (CTLE) receiver with a 12-tap decision feedback equalizer (DFE). The two stages of CTLE, shown in ***Figure 2***, were based on the pole-zero pairs shown in Table 1, along with the various gain steps. The taps of the DFE were adapted and limited based on the command parameters. Had the brute force method been used, there were a total of six different parameters that could be optimized between the Rx and Tx models, translating to 636,804 simulations for complete coverage of the solution space.

Simulation Setup

A testbench was built in an IBIS-AMI simulator, as shown in ***Figure 3***. The simulation was run long enough to allow the DFE to adapt to the incoming data (~70 kbits), and to accumulate ~100 kbits for making measurements. The simulation was run with 64 steps per unit interval (UI) and a vertical resolution of 2048, allowing sufficient detail for taking accurate measurements.

To ensure that the optimization algorithm would perform for different setups, three different channels were selected based on the insertion loss mask from the OIF-CEI-112G-LR spec[4] shown in ***Figure 4***. The chan-nels were selected for a short, medium, and long channel, based on the loss at the fundamental frequency (28 GHz), as labeled in the figure.

For the optimization to run correctly, the algorithm needed to know which parameters could be adjusted and which outputs could be maximized for the best result. The CTLE control was set up as a list of numbers, while the Tx pre- and post-cursor were set up as a range of float numbers, with the upper and lower bound based on values from Table 1. ***Figure 5*** shows the parameter setup details. It was assumed the impact of the parameters would have a continuous impact on the simulation output when increased or decreased. Placing the parameter values in a random order could create discontinuities in the output, causing the algorithm to possibly not converge on a solution.

As stated in the previous section, the sum of the Tx pre- and post-cursor needed be 0.5 or less. The main cursor was calculated by summing the pre- and post-cursor

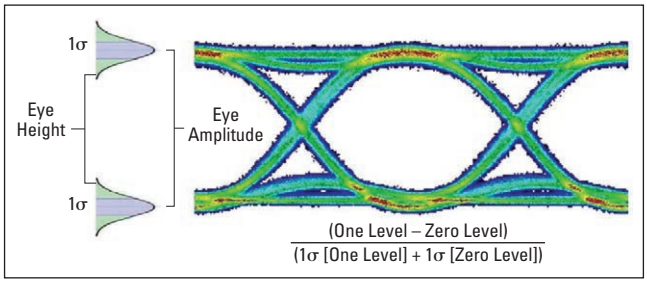

Fig. 6 **SNR measurement details.**

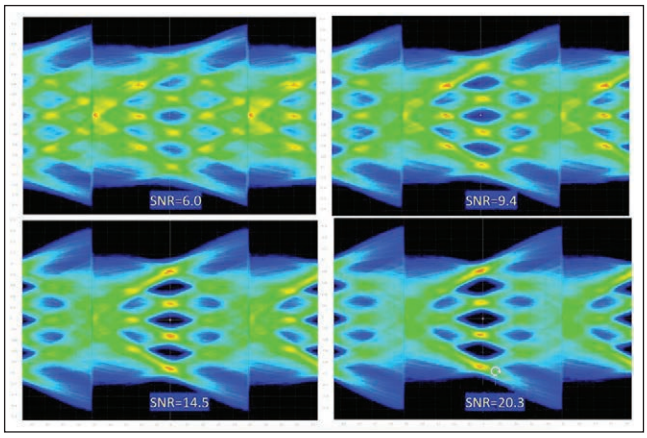

Fig. 7 **SNR comparison for various PAM4 eye density plots.**

values together and subtracting that number from 1.0. During the optimization process, it was not possible to allow the ML algorithm to skip these simulations, as they were not valid. This issue was solved by creating a custom code for the Tx AMI model to set all pre- and post-cursor coefficients to zero and set the main cursor to a very small value (0.01) whenever the sum of the coefficients was more than 0.5. This had the result of presenting a very poor result to the ML algorithm, forcing the adaptation of the parameters away from the invalid set of parameters. The output to be maximized should have been an output that would increase as the simulation result improved. This could be a combination of an eye height and eye width, or a relative measurement like COM. However, these types of measurements suffer from the fact that they are always zero for a closed eye simulation result, which should not be an issue if, across the global solution space, there are few closed eye simulation results. It has been found that for non-return-to-zero (NRZ) signaling simulations, this turns out to be true for most channels except for the most difficult. However, for four-level pulse amplitude modulated (PAM4) signaling simulations, the opposite is true; there are many more closed eye simulation results than open eye results.

This is an important detail, as the ML algorithm needs a gradient of results to help find a maximum/minimum value; known as gradient descent,[5] a common optimizing algorithm in the ML field. If the algorithm observes mostly zeros from the simulation results, the algorithm will have a gradient of zero and will be unable to find the best solution. Therefore, a measurement that will provide a non-zero result for all simulations is needed, including closed eye simulation results. This measurement should work for any signaling type (NRZ, PAM4, etc.). The signal-to-noise ratio

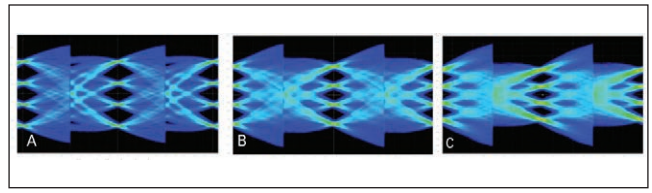

Fig. 8 **A: short channel results, B: medium channel results, and C: long channel results.**

(SNR) measurement would meet these requirements. A SNR measurement is defined as a ratio of the desired signal level to the level of background noise, plus any distortion.[6] For serial link simulations, the desired signal is defined as the difference between the high and low level of an eye diagram. The background noise and any distortions are measured by the sum of the one sigma of the high level and one sigma of the low level. This definition will work for an NRZ eye diagram, as well as a PAM4 eye diagram. The details of this measurement can be found in ***Figure 6***.

Figure 7 shows SNR measurement for various PAM4 eye density plots. While the lower plots had some eye opening, the upper plots had no eye opening and returned a zero for typical eye measurements, providing useless information for the ML algorithm. Also, the SNR measurement increased as the eye plot improved, even for the closed eye measurements.

As stated earlier, the optimization process can find a minimum or maximum value of the evaluation function. For this study, the ML algorithm was set up to find a minimum value. To accompany this, the SNR value was inverted because the ML algorithm needed to know what the stopping criterion should be. This could be either an output value below a target value, in which case the simulations would stop once this value was reached or a maximum number of simulations were run, or the finding of the best result after a specific number of simulations were run. For this study, the best SNR value was to be found within 100 simulation runs.

Parallel execution of the simulations was used, and due to the way in which the ML model is updated, the number of parallel simulations was limited to four. This may seem to be a disadvantage when compared to the almost unlimited number of cores that could be used for an exhaustive search; however, the results reflected that this limitation had a minimal impact on the time to convergence when measured by the number of simulations needed.

The optimization was run on the three different channels, as shown in Figure 6. For each channel, the following steps were performed:

1. Based on the number of parameters to be optimized (six), 30 initial simulations were run with random parameter settings. The SNR of each simulation is recorded.
2. The surrogate model is trained on the SNR output from the initial set of simulations.
3. The surrogate model is then queried for the next set of simulations.
4. The results from the next set of simulations are used to update the surrogate model.
5. Steps 3 and 4 are repeated until no parameter sets

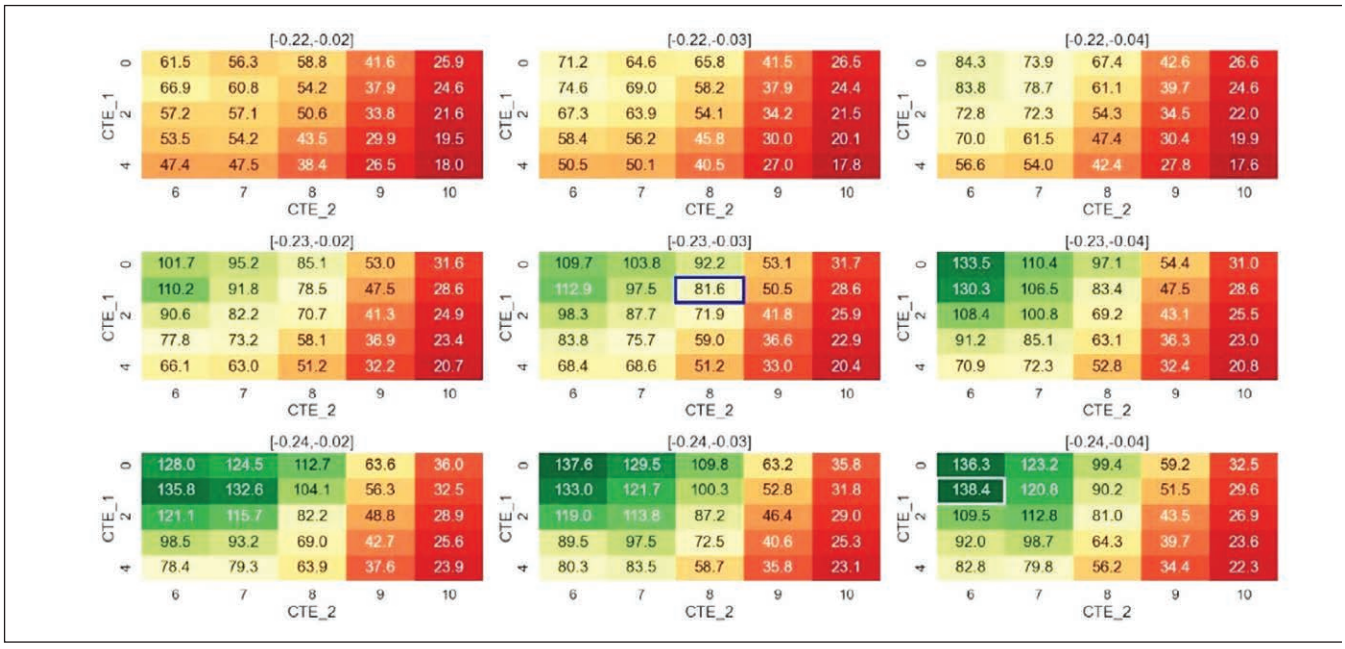

▲ Fig. 9 **Heatmap grid of the parameter sweep.**

are available, or if the total number of simulations has been reached (100 simulations).

Simulation Results

Figure 8 shows the optimization convergence for the three different channels detailed in Figure 6.

Based on the flow described in the previous section, it can be observed that after the initial 30 simulations, the algorithm quickly converged to a good result, with the following simulations used to refine the result to a better answer. In addition, the algorithm was able to converge for three different lossy channels, showing the robustness of the process.

The results clearly show that the ML optimization found parameters that resulted in good open eye results at the Rx. While these results were good, it was necessary to verify that they were the best results, or at a minimum, close to the best results. To confirm that the algorithm converged on the best set of parameters, the parameters from the best result of the medium channel were locally swept around to see if a better result could be found. Due to the difficulty of showing more than four parameter sweeps in a single graph, it was decided to keep the values for C(-3) and C(-2) constant, while C(-1), C(+1), and the two CTLE stage settings were swept around the best result found by the ML algorithm. Assuming the algorithm came close to the best result, this sweeping methodology should confirm the algorithm's result.

Figure 9 shows a grid of heatmaps of this manual sweep. The title for each heatmap shows the C(-1), C(+1) setting. Each heatmap shows the sweep values of the two CTLE stages. The best result from the ML algorithm is highlighted in blue, while the best result from the manual sweeping is highlighted in white. While the maximum SNR value found by the manual sweeping is much higher than that found by the algorithm, the parameters for the maximum SNR are only one step away from those found by the algorithm.

Conclusion

The application of Cadence's Sigrity SI/PI ML technology was investigated for refining IBIS-AMI parameters to quickly and efficiently converge on the best set of parameters in a set of IBIS-AMI models. The ML algorithm was applied to the optimization of AMI parameters in a serial link simulation. The results show that the algorithm was able to successfully find good results for three different channels. This method found a good set of parameters in fewer simulations than if a traditional manual method had been deployed, conserving the use of limited human and computing resources. In most test cases, it was found that only 100 simulations were needed to find the best set of parameters. ■

REFERENCES

1. M. Kashyap, K. Keshavan, and A. Varma, "A novel use of deep learning to optimize solution space exploration for signal integrity analysis," *IEEE*, April 5, 2018, https://ieeexplore.ieee.org/document/8329701.
2. "Implementation of Exhaustive Search Algorithm for Set Packing," https://www.geeksforgeeks.org/dsa/implementation-of-exhaustive-search-algorithm-for-set-packing/.
3. W. Wang, "Bayesian Optimization Concept Explained in Layman Terms," March 18, 2020, https://medium.com/data-science/bayesian-optimization-concept-explained-in-layman-terms-1d2bcdeaf12f.
4. "Common Electrical I/O (CEI) - Electrical and Jitter Interoperability agreements for 6G+ bps, 11G+ bps, 25G+ bps, 56G+ bps and 112G+ bps I/O," *Optical Internetworking Forum*, May 5, 2022, https://www.oiforum.com/wp-content/uploads/OIF-CEI-5.0.pdf.
5. J. Brownlee, "Gradient Descent for Machine Learning," *Machine Learning Algorithms*, August 12, 2019, https://machinelearningmastery.com/gradient-descent-for-machine-learning.
6. "Understanding Eye Pattern Measurements," *Anritsu*, https://rintintin.colorado.edu/~gifford/5830-AWL/Anritsu%20Eye%20Diagram.pdf.

Eliminate unwanted EMI with our MAGPro® Power Choke Impedance Finder

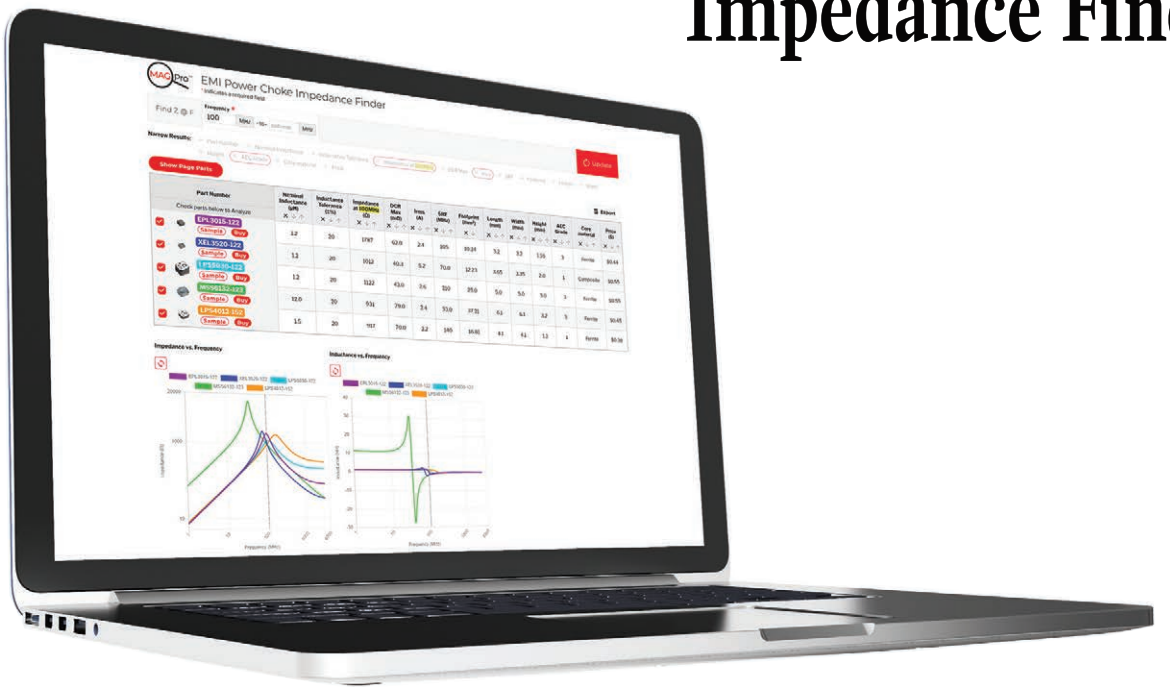

Find off-the-shelf, inductors/chokes for your EMI noise-filtering applications with this powerful new tool

Coilcraft's patented MAGPro EMI Power Choke Impedance Finder is a powerful search and analysis tool developed by engineers for engineers to help you identify the optimal inductors/chokes for various types of EMI filters.

It helps you optimize the design of an input EMI filter or a second stage output filter for switching power supplies and other applications requiring noise reduction.

The Z @ Frequency search option locates inductors that perform at your desired impedance and frequency, allowing you to analyze and compare multiple part numbers of your choice quickly and easily, and then order free evaluation samples for further review.

So take the first step toward suppressing unwanted EMI noise. Start your search today at **www.coilcraft.com/tools**.

Coilcraft®

WWW.COILCRAFT.COM

Seeing Through the Noise: Reliable Power Rail Measurements in High-Current AI Systems

Seamus Brokaw, Tektronix, Beaverton, Ore., and Steve Sandler, Picotest, Phoenix, Ariz.

The Challenge: Measurement Uncertainty Obscures VRM Performance Gains

Measuring power rail transient performance in response to high current dynamic loading is essential for power integrity validation. Modern AI GPU boards can experience dynamic currents exceeding 1000 Amps (A), creating demanding validation requirements. However, these evaluations have revealed that these critical power rail measurements vary wildly based on measurement location and equipment used, creating uncertainty that obscures validation of potential gains from next-generation Voltage Regulator Module (VRM) designs and can lead to incorrect engineering decisions.

In this article, power rail voltage measurement uncertainty is examined using several different probe configurations to monitor VCore measurements on a Picotest S2000 load stepper board. The results reveal measurement variations up to 27 mV — a level of uncertainty that can completely mask the performance improvements engineers are seeking from advanced VRM technologies.

How can engineers trust their measurements when the uncertainty exceeds the performance gains they are trying to validate?

TLVR Technology's 10 mV Promise

Recent research presented at DesignCon 2025 by Idan Ben Ezra of Broadcom, along with Steve Sandler, Heidi Barnes, and Benjamin Dannan, demonstrated the potential of Trans-inductor Voltage Regulator (TLVR) technology, a specialized type of VRM that enhances the transient response of power delivery systems, particularly for high-current, low-voltage applications Transformer-Less Voltage Regulator.[1] Their simulation results showed that TLVR VRMs with non-linear control could improve transient response significantly by approximately 10 mV compared to traditional VRM designs as shown in ***Figure 1***.

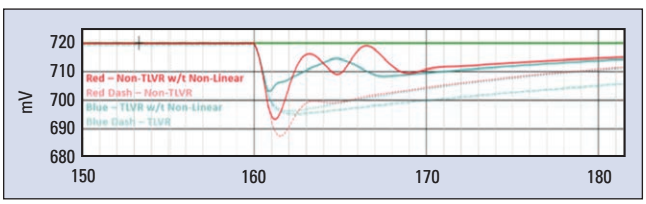

Fig. 1 Simulation results showed that TLVR VRMs with non-linear control could improve transient response significantly by approximately 10 mV compared to traditional VRM designs.[1] ***Source:*** **Idan Ben Ezra, Broadcom.**

In measurement this might not be as clear, but what are the ramifications of TLVR VRMs? The TLVR VRM requires a 2-winding inductor rather than a single-winding inductor. The secondary inductor windings from each VRM phase are series connected. If the series connection opens, the benefit is lost. Many applications monitor this winding loop and shut down the board if a connection is lost. In the case of a 16-module VRM, 32 switching phases are generally included, so 32 additional windings and one of the 32 connections opening can result in the board being disabled.

The impacts are summarized as:

1. Added complexity: TLVR requires dual-winding inductors instead of single-winding designs
2. Increased cost: Additional windings and monitoring circuitry
3. Reliability concerns: In a 16-module VRM with 32 switching phases, any of the 32 secondary winding connections could disable the entire board
4. Space constraints: Extra inductor windings consume valuable PCB real estate.

The engineering decision to adopt TLVR technology hinges on reliably measuring this 10 mV improvement. Yet measurement uncertainty of 27 mV — nearly three times the expected gain — makes this validation impossible with conventional measurement approaches.

Experimental Setup

All of the measurements were performed on a Picotest 2000 A demo board as shown in ***Figure 2***. This system uses a matrix of 256 individual GaN load cells that emulate the dynamic 2000 A, sub-nanosecond transients that today's VRM designs must withstand. A 16-bit high-speed microcontroller provides 11-bit load control (up to 2047 A with 1 A resolution) at sample rates exceeding 50 MSPS. A single probe test pad was used for the voltage monitoring; a single channel of a single oscilloscope was used for the testing to eliminate oscilloscope termination and setting variations from the results; and a pre-programmed demo transient pattern was used to create an identical dynamic load transient for each evaluation.

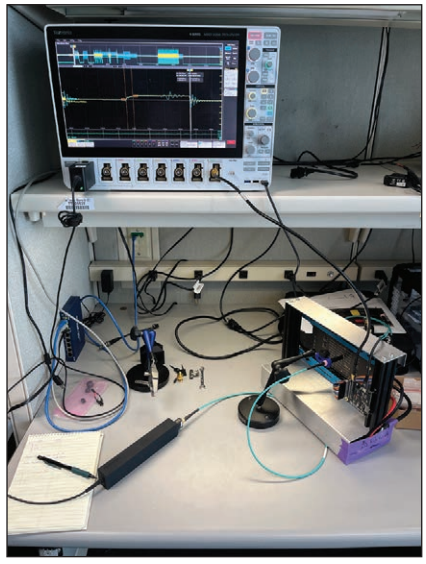

Fig. 2 **Experimental test setup.**

The comprehensive load pattern stresses all aspects of power rail design by including:

- Fast current steps and bursts with nanosecond edges
- Linear and exponential current ramps
- Sine wave patterns at various frequencies
- Gaussian noise patterns
- Current steps of varying amplitudes.

This pattern exposes non-linear, time-variant, and large-signal effects that traditional impedance measurements and Bode plot analysis cannot reveal.

Probe Configurations

Seven different Tektronix and Picotest probe configurations were evaluated to quantify measurement uncertainty (see ***Figure 3***):

1. TICP with J2115A isolator (P2104A 2X tip) - 25 acquisitions
2. TICP without isolator (P2104A 2X tip) - 15 acquisitions
3. TICP with 10X tip (baseline comparison) - 25 acquisitions
4. TPR with J2115A isolator (P2104A 1X tip) - 20 acquisitions
5. TPR without isolator (P2104A 1X tip) - 21 acquisitions
6. BNC with J2115A isolator (P2104A 1X tip) - 20 acquisitions
7. BNC without isolator (P2104A 1X tip) - 20 acquisitions

These configurations test the effects of isolation from both galvanically isolated TICP probes and transformer-coupled isolation via the J2115A. TICP provides differential measurements with DC-to-DC galvanic isolation, while the J2115A offers isolation above 1 kHz. In addition, TPR and BNC measurements are ground-referenced, making them more susceptible to ground loop errors.

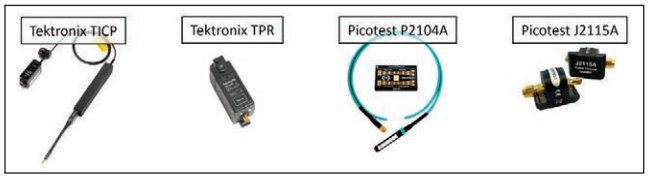

Fig. 3 **Examples of probe configurations tested.**

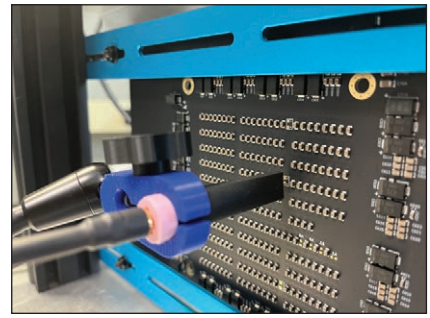

Fig. 4 **VCore test point location.**

Measurement Protocol

For each probe configuration, multiple acquisitions were performed (15-25 samples) to enable statistical analysis. All VCore voltage measurements, (see ***Figure 4***) were taken during identical load patterns on the same test pad to ensure fair comparison.

Statistical Analysis Approach:

- Mean, maximum, and minimum voltage values recorded for each acquisition
- Standard deviation calculated across multiple acquisitions
- Normalization using Mean VCore voltage (843.02 mV) for consistent comparison
- Uncertainty quantified as the range between maximum and minimum values.

Key Measurements:

- Mean voltage stability: baseline voltage consistency
- Maximum voltage deviation: peak voltage during load transients
- Minimum voltage deviation: voltage droop during load steps
- Standard deviation: measurement repeatability across acquisitions.

The Root Cause: Ground Loop Error and Probe Noise

The measurement results (see ***Figures 5*** and ***6***) revealed two critical factors affecting measurement quality: ground loop error and probe noise.

Ground Loop Error: The Primary Culprit

The measured data provides compelling evidence that ground loop error is the primary source of measurement uncertainty. While measurements varied significantly across configurations, solutions incorporating the J2115A coaxial isolator showed much closer agreement. TICP probe results were similar with and without the J2115A isolator, confirming that the large variance stems from ground loop error.

Ground loop error occurs when multiple ground paths create circulating currents that induce voltage errors in measurements. The J2115A coaxial isolator corrects this error at frequencies between approximately 1 kHz and 10 MHz. Above 10 MHz, coaxial probe cables often provide sufficient ground loop isolation on their own.

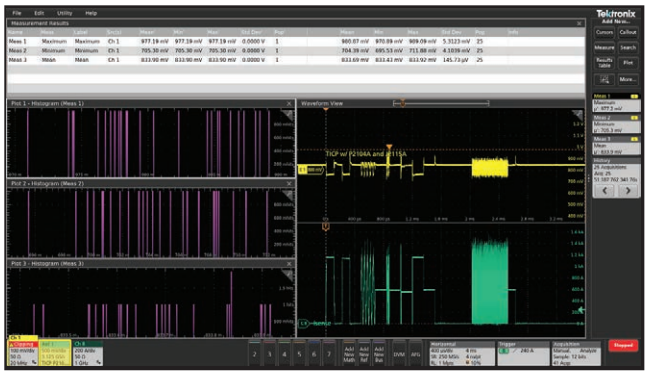

Fig. 5 VCore measurement statistics for one probe configuration.

	Uncertainty	Max (mV)	Min (mV)
highest		29.14	27.37
lowest		1.06	0.49

	Vdd Avg. (mV)	Normalized Max (mV)	Normalized Min (mV)	Max - Min
TICP P2104A 2X w/ J2115A	843.02	990.2	713.72	
TICP P2104A 2X		986.35	710.53	
TICP 10X tip		980.23	720.63	
TPR P2104A 1X w/ J2115A		985.29	714.21	
TPR P2104A 1X		962.03	737.9	
P2104A 1X w/ J2115A		989.18	711.94	
P2104A 1X		991.17	712.95	

Fig. 6 VCore measurement results.

Probe Noise: The Secondary Factor

Probe noise also significantly impacts measurement accuracy. When a 10X attenuating tip to the TICP probe was attached, the results varied dramatically, demonstrating how probe noise can sway measurements. This highlights why simply using a differential probe like a Tektronix TDP1500 is not always the solution — high-impedance differential probes are particularly susceptible to noise pickup, and many differential probes introduce their own noise.

Measurement Recommendations

Based on the findings, a multi-layered approach to minimize measurement uncertainty is recommended:

1. Isolation is Essential
   - TICP probes provide galvanic isolation down to DC, eliminating ground loops
   - J2115A coaxial isolator offers transformer-coupled isolation above 1 kHz
   - Combined approach: While TICP's high CMRR provides significant isolation, adding the J2115A isolator provides additional protection without degradation.
2. Low-Noise Probe Selection
   - Low-impedance probes reduce susceptibility to noise pickup
   - High shield attenuation cables minimize electromagnetic interference
   - Short probe pins reduce antenna effects
   - Ferrite-shielded probe tips block high-frequency noise.
3. Best Practices for Power Rail Measurements
   - Use differential probes, when possible, but ensure they have low noise characteristics
   - Implement coaxial isolation for ground-referenced measurements
   - Minimize probe lead length to reduce inductance and noise pickup
   - Verify measurements by shorting probes on a ground pad — TICP should show minimal difference with and without isolator, while TPR and P2104A/5A should show significant improvement with isolator.

Conclusions

The investigation reveals that measurement uncertainty in power rail validation stems primarily from ground loop error and probe noise. With measurement variations up to 27 mV — nearly three times the 10 mV improvement promised by TLVR technology — conventional measurement approaches cannot reliably validate next-generation VRM performance.

The solution requires a comprehensive approach:

1. Galvanic isolation through TICP probes eliminates ground loops down to DC
2. Coaxial isolation via J2115A isolators provides additional protection above 1 kHz
3. Low-noise probe design with high shield attenuation and ferrite protection
4. Proper measurement practices including short probe leads and ground pad validation.

While Picotest PDN cables provide higher shield attenuation than most and the Tektronix TICP probe offers much higher CMRR than conventional probes, neither provides infinite protection. Engineers must implement all available error minimization techniques to achieve the measurement certainty required for confident decisions.

The bottom line: Do not let measurement error exceed VRM performance gains. With proper probe selection and isolation techniques, engineers can achieve the measurement accuracy needed to validate the 10 mV improvements promised by next-generation VRM technologies.

Next Steps: Advanced Validation Techniques

Future work will explore additional validation techniques, including:

- Additional probe trials with Tektronix TDP1500 and Picotest P2105A browsers
- Better normalization techniques to add in offset gain and gain compensation to the existing offset normalization used in this article
- Ground to ground measurement and shorted ground measurements to quantify the ground bounce
- Frequency-domain analysis of measurement uncertainty.

By addressing measurement uncertainty at its source, engineers can move beyond "blindfolded engineering" and make confident decisions about power integrity optimization.

REFERENCE

1. S. Sandler, B. Dannan, H. Barnes, and I. B. Ezra. "Power Delivery Network Master Class on 2000A: How to Design, Simulate, & Validate," *DesignCon 2025*.

Siglent Delivers 12-Bit Oscilloscopes to address Signal Fidelity Challenges on every bench

12-bit

SIGLENT®

Every Bench. Every Engineer. Every Day.

support@siglentna.com www.siglentna.com

ManyPoint Networks: A System Co-Design Framework for 448 Gbps AI Fabrics and Beyond

Andrew Josephson
Samtec, New Albany, Ind.

Achieving 448 Gbps signaling may require system architects and standards bodies, as well as PHY and interconnect designers, to rethink bandwidth provisioning. AI-scale computing provides an opportunity to rethink system architectures, and it may be time to bridge the interconnect, PHY, and compute cluster domains with a single actionable performance metric.

This work introduces a hardware-centric definition of compute cluster bisection bandwidth as a performance metric for AI-scale 448 Gbps systems. Unlike traditional abstractions, this metric is grounded in physical interconnect layout and IO port availability, enabling system architects to evaluate bandwidth provisioning through real, bidirectional link paths. Because it focuses on the minimal physical bisection of a topology, this approach supports clean, cost-aware comparisons across architectures and avoids assumptions about traffic patterns or runtime behavior.[1]

This article presents a top-down co-design methodology that begins with cluster topology and works downward into the electromagnetic and physical layers. This inversion of the conventional signal integrity (SI) workflow allows one to engage directly with system-level performance metrics, bridging the gap between PHY feasibility and workload-driven cluster efficiency. The framework leverages practical hardware configuration item breakdowns to support scalable accelerator fabrics, enabling link-by-link and pad-by-pad tradeoff analysis across compute domains.

The methodology is demonstrated in the context of 448 Gbps pathfinding and 400G serial interface development, where mmWave test fixtures and parallel PHY validation tools now operate at 100 GHz. These capabilities invite a shift in engineering roles where PHY-level SI engineers take on traditional roles of hardware systems engineering to support the scale, efficiency, and reuse demands of modern GPU cluster networks. The cost and complexity of integration at 224 and 448 Gbps bandwidths will require systemic adaptability and reuse of lab infrastructure.

Not all mesh topologies are appropriate for accelerator mesh fabrics, and not all accelerator fabrics qualify as efficient manypoint networks. This framework distinguishes between them by evaluating how efficiently a topology uses its available IO to reach across the cluster. By aligning physical interconnect behavior with compute cluster performance, the methodology enables topology-aware, IO-efficient, and PHY-agnostic co-design of high bandwidth AI systems.

Accelerator Mesh Co-Design Challenges

In conventional system design, bisection bandwidth is often defined through abstract or runtime-centric lenses. Graph-theory models treat it as the capacity of the smallest edge cut, assuming unit link capacity and ignoring physical feasibility, while MPI benchmarks measure throughput across a cut during execution, entangled with routing, congestion, and software stack effects. Vendor marketing can

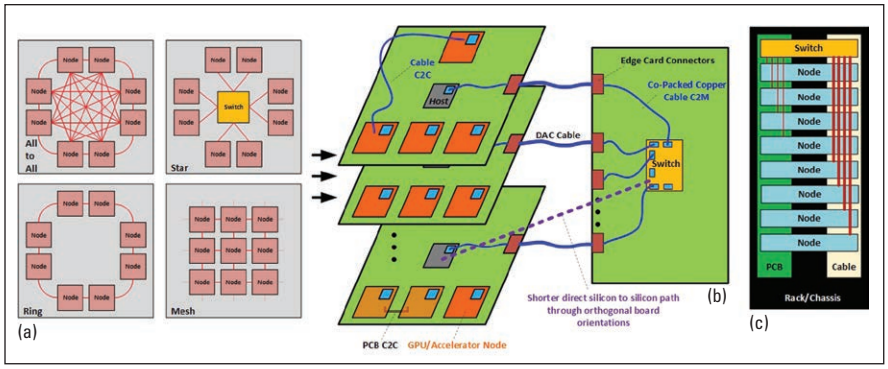

Fig. 1 Backplane topologies lead to node-to-node meshing.

cause further confusion by presenting inflated aggregate bandwidth numbers based on non-minimal cuts, often detached from real-world constraints. These approaches, while useful in their respective domains, fail to provide a clean, actionable metric for hardware-centric co-design. They obscure the physical realities of interconnect layout, IO port availability, and packaging constraints — precisely the factors that dominate system-level performance in scaling up accelerator mesh fabrics. Unlike graph-theory edge cut models or MPI runtime metrics that OEM architects rely on, this definition enforces a minimal physical bisection normalized by IO, creating a common language that bridges system-level abstractions and SI realities.

Bisection bandwidth is here defined as the count of real, bidirectional hardware links crossing a minimal physical bisection of the compute node cluster, normalized by the number of IO ports per compute node. This definition is flat, minimal, and grounded in actual hardware, making it cost-aware and topology-sensitive. It favors designs that use IO to reach more destinations, not merely to push faster links to the same endpoints with linear data rate scaling. By avoiding assumptions about traffic patterns, software stack behavior, or promotional bias, this definition offers a layer-aligned view through interconnect, PHY, and cluster topology to be co-designed.

Topology Analysis

Figure 1 introduces the concept of direct node-to-node mesh topology in the context of a SI channel design challenge. It originates from the 2023 SIJ article[2] comparing PCB versus cabled backplane implementations for 112 Gbps links. It highlights how cabled backplanes can extend physical reach beyond the limits of traditional PCB routing, enabling more nodes to be interconnected directly in a mesh fabric. This is especially relevant for scaling beyond blade-local accelerator node meshes, where the ability to span blades with high-speed links becomes critical. By showing how topology and physical channel design intersect, Figure 1 provides the beginnings of the many-points network concept, where interconnect architecture is co-designed with physical layer constraints to support larger, more efficient accelerator clusters.

Additionally, Figure 1 introduces co-packaged connectors and the physical layer challenges associated with scaling 400 Gbps links to the PCB baseboard. It highlights how co-packaged connector technology addresses reach limitations by enabling high-density, low-loss interconnects directly at the package level.[3] This is critical for implementing blade-level meshes, where local interconnects must support high bandwidth with minimal signal degradation.

For extended reach beyond the blade, co-packaged connector meshes can be paired with cabled backplanes or alternate board orientations within the cabinet. For example, a pizza-box switch oriented vertically (see Figure 1b) allows every direct-attach copper cable to remain short and directly connect to a horizontal compute blade stack, which is similar to how liquid cooling manifolds or 48 V bus bars are routed. Standard 19-in. rack mount form factors favor rapid development and high-volume manufacturing cost benefits, but they also make it challenging for short path lengths between nodes.

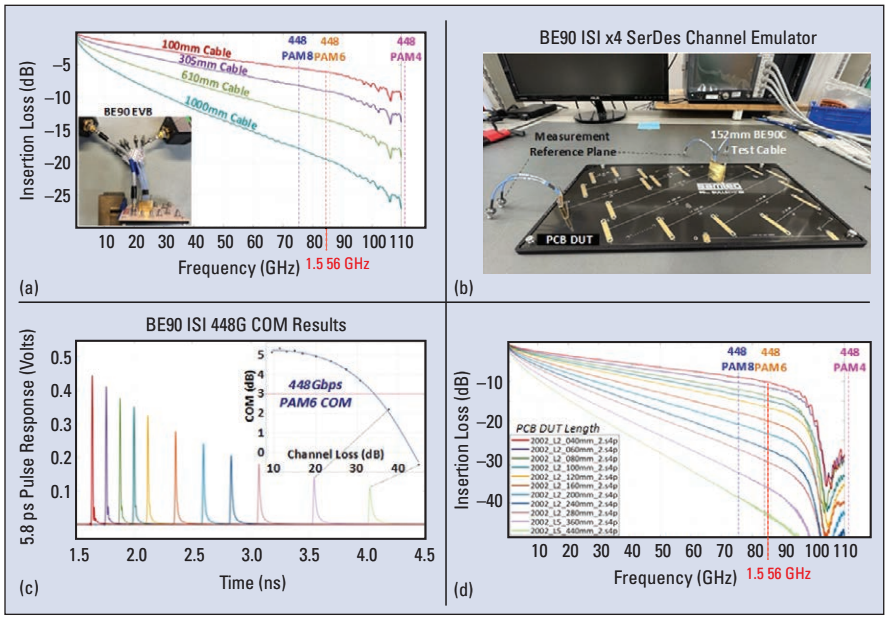

Fig. 2 Millimeter wave test fixtures, such as the Bulls Eye® ISI Evaluation Board,[4] find use in 400G industry pathfinding. (BE90 EVB: Bulls Eye 90 GHz evaluation board. COM: channel operating margin.)

Measurement Data for 448 Gbps Pathfinding

It is impossible to perform system-level analysis (and co-design) without a means for physical layer characterization. Currently, industry standards groups are convening to consider ways to achieve 448 Gbps signaling. Some components originally built for 224 Gbps analysis are proving helpful in this area. ***Figure 2*** presents four plots from Samtec's mmWave test fixture characterization work, each demonstrating aspects of SI performance at or beyond 100 GHz. These plots represent real, deployable measurement capabilities that directly support 400G PHY IEEE and OIF standardization efforts in the industry.[5-7]

The data in Figure 2 shows that the mmWave test platform demonstrates

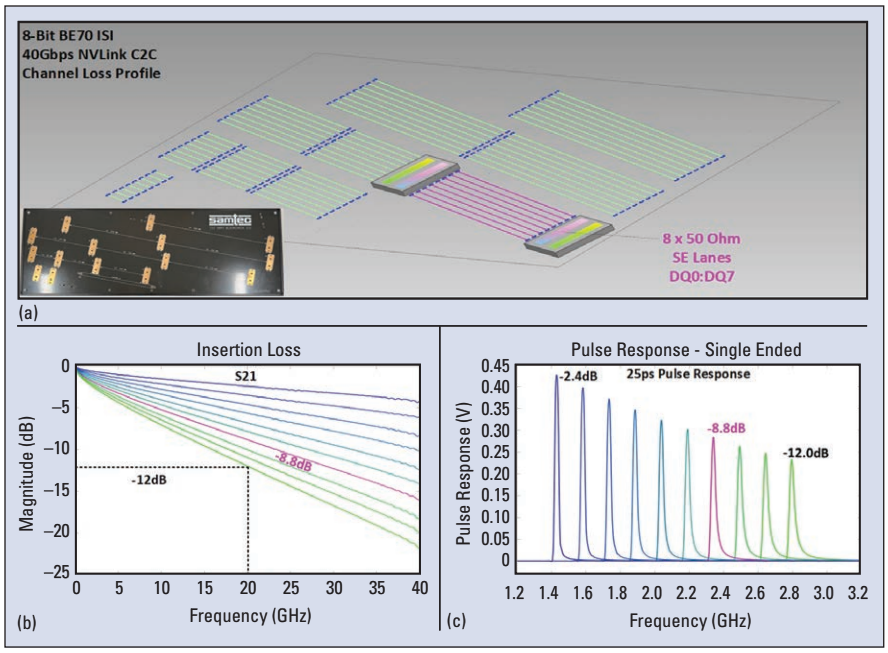

Fig. 3 When paired with low-skew, single-ended coaxial cabling and precision RF fixture design techniques implemented in HVM PCB processes and materials, the ISI evaluation board demonstrates C2C scale electrical budgets at blade-scale physical reaches.

insertion loss from approximately 6.1 dB to 42.5 dB at 82 GHz Nyquist across 40 to 442 mm paths, maintaining a near-linear response with ~1 dB ILD bandwidth to 90 GHz. Differential return loss remains around 15 dB through 90 GHz. Unlike fixtures that rely on tightly coupled differential lines, this design achieves ~55 dB adjacent single-ended lane isolation and leverages low-skew true/complement pairs (~1.5 ps P/N skew) to suppress skew-induced inter-symbol interference (ISI) during coupled-line propagation.

The fixture's ability to maintain linear insertion loss and avoid resonant roll-off up to 100 GHz (see Figure 2c) is critical for validating next generation SerDes channels, especially for PAM4 and PAM6 signaling. When designing interconnects for large-scale accelerator meshes, this level of measurement fidelity is essential because every millimeter of reach matters and every dB of margin is precious.

The mmWave test fixtures used in these next generation serial standard pathfinding studies, such as the Samtec Bulls Eye® ISI Evaluation Board (see Figure 2b), function as 100 GHz single-ended buses that are 8 bits wide for validating x4 SerDes quads. The well isolated single-ended design prevents skew from being converted to ISI during propagation in a coupled transmission. While originally designed for multi-lane SerDes characterization, these fixtures are broadly useful for parallel PHY validation as well due to the connector test points supporting wider coaxial count arrays. This flexibility enables high-fidelity signal capture across multiple lanes simultaneously.

Figure 3 presents measurements from a Samtec BE70 ISI Evaluation Board (a 2×8 lane, 224Gbps channel emulator originally designed for validating 4-lane differential SerDes quads). The board functions as an 8-bit single-ended bus, and its parallel lane structure, channel loss slope, and dynamic range align closely with the Nvidia NVLink chip-to-chip (C2C) die-to-die SI budget.[8] The measured reach of approximately 250 mm of PCB stripline plus 300 mm of cable significantly exceeds the 60 mm reach typically associated with NVLink C2C off-package links, suggesting that the electrical feasibility of NVLink C2C may extend well beyond its original design envelope. When paired with low-skew, single-ended coaxial cabling and precision RF fixture design techniques implemented in HVM PCB processes and materials, the evaluation board demonstrates C2C scale electrical budgets at blade-scale physical reaches.

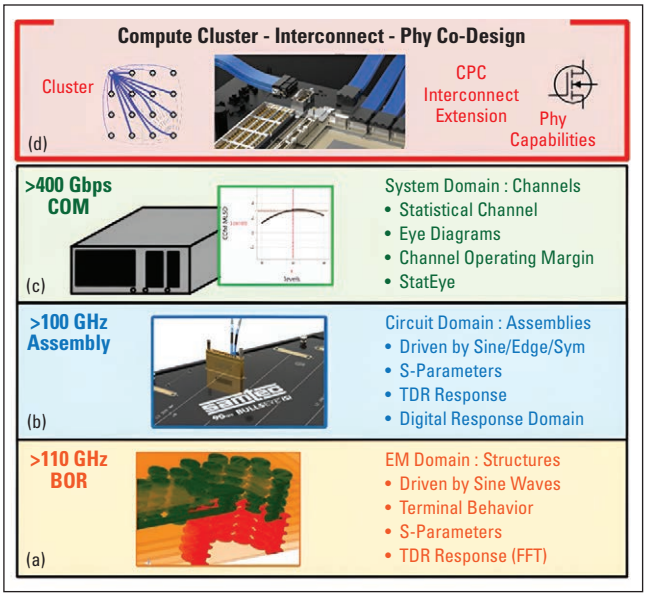

Fig. 4 Four levels of potential co-design encompass the accelerator mesh to the EM domain.

Accelerator Mesh Co-Design

Figure 4 illustrates the layered bandwidth stack from the electromagnetic (EM) domain to compute cluster topology. The bottom three layers (electromagnetic domain, circuit assemblies, and statistical channel modeling) were established in previous work as foundational to PHY and interconnect characterization.[1] The top layer added here represents the compute cluster performance domain, specifically for IO-bound workloads. This addition is enabled by the hardware-centric bisection bandwidth definition, which connects physical interconnect behavior to system-level performance through the stack. This alignment allows the construction of a closed-form analytic framework by inserting a realistic hardware configuration breakdown. Node counts, blade layouts, and IO port mappings can be thought of as configurations that translate abstract bandwidth metrics into actionable system architecture decisions.

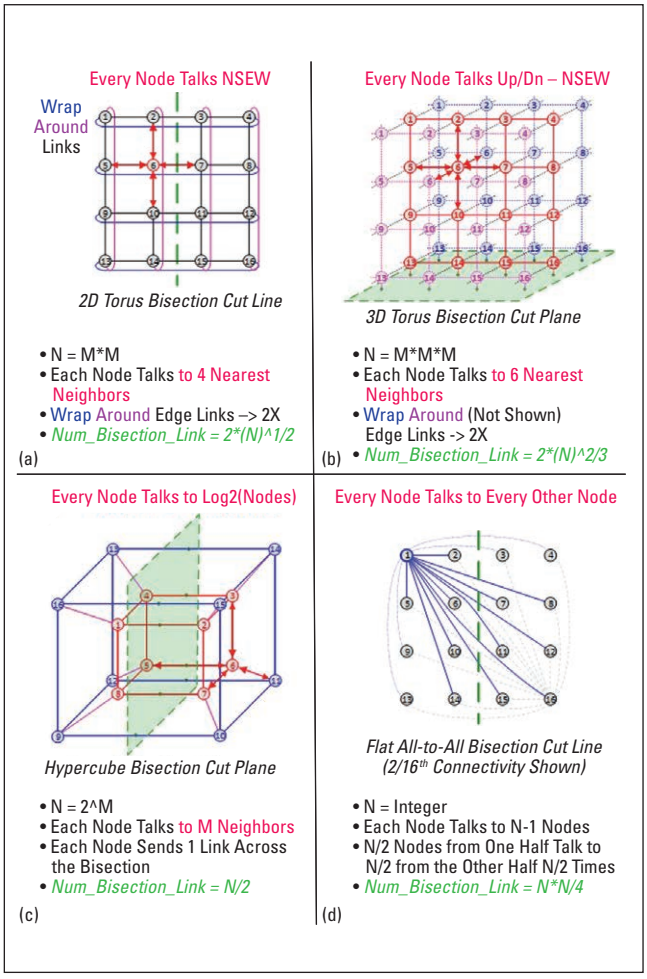

Fig. 5 Bisection links showing node mesh topologies and scaling. The bisection cut intersects a line, a surface, or a dimensional axis, depending on the topology.

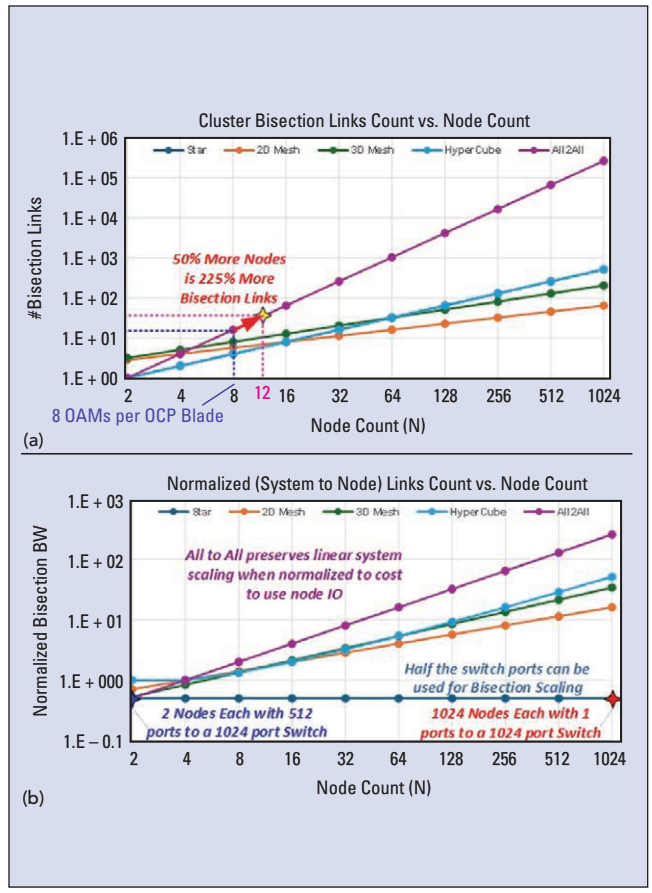

Fig. 6 Analyzing HW-centric bisection BW (links) vs. node count (normalized to the number of IO used) helps designers evaluate how efficiently a topology uses IO.

Bisection BW Detailed Explanations and Scaling

Understanding how bisection bandwidth scales across different mesh topologies is central to evaluating interconnect efficiency in accelerator clusters. ***Figure 5*** visualizes node arrangements, and the corresponding bisection cut lines for several canonical topologies: 2D torus, 3D torus, hypercube, and all-to-all. Each diagram shows how many links cross the bisection plane, which directly determines the system's ability to move data between two halves of the cluster. These link counts are derived from the physical layout and connectivity rules of each topology. For example, a 2D torus with wraparound links (see Figure 5a) yields a bisection link count of $2 \times \sqrt{N}$, while a 3D torus scales as $2 \times N^{(2/3)}$, and a hypercube scales linearly as $N/2$.

The approach shown in Figure 5 for defining bisection bandwidth is intentionally PHY-agnostic. It does not assume a specific signaling rate, encoding scheme, or protocol. Instead, it focuses on the number of bidirectional links crossing the minimal cut, which can then be multiplied by the per-link bus width and data rates later. This makes it a clean and minimal abstraction for system-level bandwidth provisioning. By avoiding assumptions about traffic patterns or runtime behavior, it enables comparisons across topologies and hardware implementations.

Figure 6 reinforces this concept by plotting the total bisection link count and the normalized bisection bandwidth per node IO. The latter metric (system bisection bandwidth divided by total outbound IO bandwidth) reveals how efficiently a mesh topology uses its available IO to reach across the cluster. It highlights that some topology types (such as all-to-all) scale bandwidth faster than IO port count, while others (such as star) remain fixed regardless of port count.

The normalized view in Figure 6 is especially powerful for co-design. It allows system architects to evaluate not just how much bandwidth a topology provides, but how efficiently it uses IO to scatter data away from each node. In AI-scale workloads, where collective operations and many-to-many exchanges dominate, this efficiency can matter more than raw data rate. A topology that doubles its bisection bandwidth without doubling its IO port count is inherently more scalable while also being more difficult to implement for interconnect. The plots in Figure 6 show that mesh topologies can achieve this, especially when paired with co-packaged interconnects and optimized PHYs. Here is the essence of the co-design framework: to enable topology-aware, IO-efficient, and PHY-agnostic co-design of accelerator mesh fabrics.

Scalable Accelerator Reference Architecture for Mesh Interconnects (SARAMI)

Figure 7 presents a tabular breakdown of node counts

# Nodes	X # Nodes per Long Blade Edge	Y # Nodes per Short Blade Edge	Z # Blades per Cabinet
2	2	1	1
4	2	2	1
8	4	2	1
16	4	4	1
16	8	2	1
16	4	2	2
32	4	4	2
32	4	2	4
64	4	4	4
64	4	2	8
128	4	4	8
128	4	2	16
256	4	4	16
256	4	2	32
432	4	3	36
512	4	4	32
512	4	2	64
528	4	3	44
1024	4	4	64
1024	4	2	128

(a)

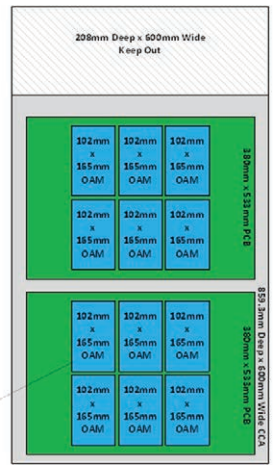

(b)

12 OAMs per 1U Blade to ~500kW per Cabinet

Fig. 7 SARAMI Reference Architecture table overlays actual dimensions, power envelopes, and packaging densities.

mapped to realistic OCP OAM blade form factors, serving as an example SARAMI. This figure translates the abstract concept of scalable accelerator mesh fabrics into tangible hardware configurations through a simple assumption of hardware configuration items breakdown, showing how compute nodes scale within blade and cabinet constraints. It overlays actual dimensions, power envelopes, and packaging densities. By grounding the possible mesh topologies to real-world form factors, system architects can quickly evaluate how many nodes can be supported per blade, how blades stack within a cabinet, and how interconnect provisioning aligns with physical layout and scale-up network goals. This then allows closely coupled PHY/interconnect solutions to be considered as candidate technologies for the links.

Conclusion and Next Steps

This work presents a reference framework for evaluating PHY and interconnect scalability in accelerator mesh fabrics intended for GPU and accelerator scale up networks. It is grounded in a hardware-centric definition of bisection bandwidth that enables PHY-agnostic interconnect provisioning. By aligning electromagnetic, circuit, and statistical channel domains with compute cluster performance, this article presents a blueprint for a simple closed-form analytic framework that spans from SI to system-level bandwidth provisioning tailored for AI parallel compute workloads.

Future generations of AI hardware systems may need to consider blade orientation in the cabinet as a degree of freedom. For instance, the same traditional 19-in. racks that enable rapid integration and HVM cost containment of circuit card assemblies also favor highly localized blade-level compute cluster mesh interconnect in implementation. In systems where liquid cooling is required for every 1kW node, and airflow constraints are no longer dictated by board orientation, orthogonal board orientations in cabinets can help reach more destinations with IO by reducing geometrical path length between compute nodes.

This work offers a practical lens for future system design, where interconnect topology, PHY performance, and packaging geometry are co-optimized. The importance of high system bisection bandwidth becomes especially clear when viewed through the lens of modern AI problem classes. Workloads such as large-scale model parallelism, deep reinforcement learning, graph neural networks, physics-informed simulations, and large-batch data parallelism all rely on frequent, high-volume communication across distributed nodes. Ranging from all-to-all exchanges of activations and gradients to structured mesh updates in scientific machine learning, these patterns demand interconnect architectures and mesh fabrics that can sustain bandwidth across the entire cluster.[9,10] The proposed blueprint framework directly addresses this need by overlaying topology-aware co-design that scales bandwidth with node count and IO efficiency as well as link data rate, ensuring that the interconnect does not become the bottleneck in compute-intensive, communication-heavy AI systems. ■

REFERENCES

1. A. Mittal, "Rethinking Distributed Computing for the AI Era. Communications of the ACM," (2025). https://cacm.acm.org/blogcacm/rethinking-distributed-computing-for-the-ai-era/ .
2. A. Josephson, B. Gore, and J. Sprigler. "Selecting a Backplane: PCB vs. Cable for High-Speed Designs" | Signal Integrity Journal.
3. B. Gore, et al. "Beyond 200G: Brick Walls of 400G links per Lane," DesignCon 2025.
4. Bulls Eye® ISI Evaluation Boards.
5. A. Josephson and L. Mei. "Developing a Test Fixture for 200G with Pathway to 400G," Korea Test Conference, July 14, 2025.
6. Josephson, Andrew et al. "Measured 400 Gb/s per lane channel files," IEEE 802.3 New Ethernet Applications Ad Hoc 2025.
7. Josephson, Andrew. "Measurement Results for a 448 Gbps Physical Channel" OIF 448G AI Workshop, April 2025.
8. Y. Wei et al., "9.3 NVLink-C2C: A Coherent Off Package Chip-to-Chip Interconnect with 40Gbps/pin Single-ended Signaling," 2023 IEEE International Solid-State Circuits Conference (ISSCC), San Francisco, CA, USA, 2023, pp. 160-162, doi: 10.1109/ISSCC42615.2023.10067395.
9. Y. Zhou, Y. Li, and Y. Wang, "Bruck Algorithm Performance Analysis for Multi-GPU All-to-All Communication," ACM Transactions on Architecture and Code Optimization, 2024. https://dl.acm.org/doi/fullHtml/10.1145/3635035.3635047
10. Zhang, Y.; Li, Y.; Wang, Y.; et al. "WSC-LLM: Efficient LLM Service and Architecture Co-exploration for Wafer-scale Chips," *Proceedings of the 51st Annual International Symposium on Computer Architecture (ISCA)*, 2024. https://dl.acm.org/doi/epdf/10.1145/3695053.3731101.

Trusted Next-Generation Testing

See more. When you look at things differently.

Signal Quality Analyzer-R MP1900A

Now conduct highly accurate intra-pair skew measurements with the detail necessary to meet modern design requirements.

Ultra High Speed Meets Ultra High Precision

Anritsu's Signal Quality Analyzer-R MP1900A is an all-in-one high-performance multi-channel BERT with 8 slots for expansion modules such as 128Gbps PAM4 Pulse Pattern Generator, 118Gbps PAM4 Error Detector, and Jitter/Noise Generator. The MP1900A expands measurement applications supporting signal integrity analyses of high-speed network interfaces such as with 200GbE/400GbE/800GbE standards.

MP1900A has built-in Link Training/ LTSSM analysis function supporting high-speed serial bus interfaces such as PCI Express and USB.

Learn more about our **Next Generation Signal Integrity Testing Solutions** and explore the Anritsu Digital and Optical Testing Solutions Resource Center.

Scan the QR code

www.anritsu.com 1-800-ANRITSU

Anritsu
Advancing beyond

High-Speed Digital Interface Characterization Requires New Test Approach

Hiroshi Goto
Anritsu Company, Morgan Hill, Calif.

The amount of data being processed is increasing exponentially and there's no end in sight. In 2024, 147 zettabytes (ZB) of data were processed. Compare this to 181 ZB of data expected to be generated in 2025. That's 2.5 quintillion bytes each day, or 29 terabytes per second.

Data growth is due to many factors. At the top of the list is artificial intelligence — unsurprisingly. Other reasons include high-definition video streaming and the continued rollout of IoT use cases, from autonomous driving and telehealth to smart manufacturing.

To keep pace, high-speed digital interfaces used in computing, server, storage, and other data systems are evolving rapidly. ***Table 1*** lists the performance of each recent generation of PCI Express®, USB, DDR, and Ethernet, displaying how they all have increased in speed and baud rate to meet the exploding data transmission demand.

Given the high bandwidth that these digital interfaces must support, design verification becomes more complex for engineers. Traditional characterization methods have limited capability or are not granular enough to confidently identify errors

TABLE 1

HIGH-SPEED DIGITAL INTERFACE PERFORMANCE CONTINUES TO RISE TO MEET DATA DEMAND

Standard		Baud-rate [Gbaud]	Modulation	Unit Interval [ps]
PCIe	Gen5	32	NRZ	31.25
	Gen6	32	PAM4	31.25
	Gen7	64	PAM4	15.63
USB	USB 3.2	10	NRZ	100
	USB4 v1	20	NRZ	50
	USB4 v2	25.6	PAM3	39.06
DDR	DDR5	6.4	NRZ	156.25
800GbE	800GAUI-8-C2M	53.125	PAM4	18.82
	CR8			
	KR8			
	800GAUI-4, CR4	106.25	PAM4	9.41
1.6TbE	1.6TAUI-8, CRB	106.25	PAM4	9.41

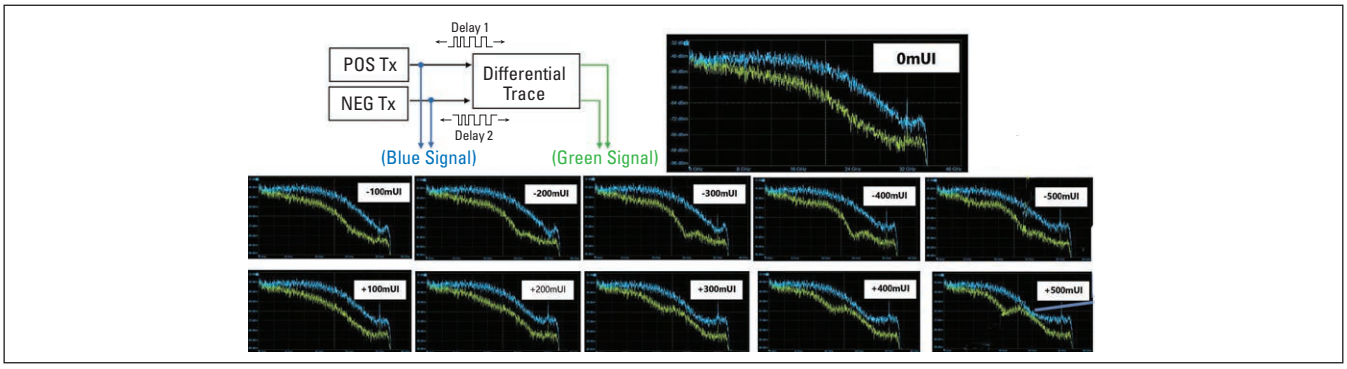

Fig. 1 **Frequency domain effect of intra-pair skew of a PCIe Gen6 MCP interface.**

Fig. 2 **The addition of skew degrades BER, making it difficult for receivers to interpret data.**

and anomalies in addition to ensuring standards compliance. A new methodology introduces skew delay generation using delay blocks in two Pulse Pattern Generators (PPGs) to conduct highly accurate intra-pair skew measurements with the detail necessary to meet modern design require-ments.

Intra-Pair Skew's Impact on Designs

Intra-pair skew refers to the time delay between two signals in a differential interface. High-speed interconnects require precise timing synchronization that can be disrupted by even the slightest bit of skew. That is because, as transmission speeds increase, the unit interval (UI) decreases. The result is that digital interfaces supporting current high-speed standards are more susceptible to bit errors. Components such as vias, design decisions like intentionally inserted gaps, and vestigial signals can introduce delays and alter the timing consistency across different signal paths.

The introduction of intra-pair skew in high-speed designs can cause reflections, crosstalk, or bit errors. Overall, signal integrity, and consequently, system reliability, will suffer. Understanding and minimizing interrupted skew helps engineers optimize signal path layouts, improve bandwidth, and ensure robust operation.

Let's look at PCIe interfaces as an example. PCIe Gen6 has a UI of 31.25 ps. A printed circuit board (PCB) with a dielectric constant of 3.5 and 5 mm trace length stripline for a PCIe Gen6 interface highlights the importance of skew testing. ***Figure 1*** shows the frequency domain effect of intra-pair skew of a PCIe Gen6 MCP interface at 32 Gbaud rate with a 1 UI unit time. As shown, there is a steep dip at the Nyquist frequency of 16 GHz with a maximum difference of about -8 dB when compared to 0 mUI.

PAM4 Modulation Impact

Intra-pair skew is especially significant on PAM4 sig-nals.

It can reduce the eye size by more than 3X, thereby increasing the signal's sensitivity to margin. This behavior is even more pronounced when using a differential trace.

PAM4 is the preferred modulation scheme for high-speed interfaces because it significantly improves data speed, yet only requires bandwidth similar to NRZ. Because PAM4 transmits two bits per symbol by using four distinct amplitude levels, it requires precise timing to accurately distinguish between the different levels, especially when signals are close together in time. Therefore, PAM4 signals are highly sensitive to intra-pair skew.

In PAM4, the differential signal when the skew is 0 UI maintains the shape of POS (or the inversion of NEG). When the skew increases, the differential signal is degraded, and it becomes difficult to distinguish 1 or 0 levels as a digital signal.

The Impact of Intra-Pair Skew on BER

Excessive intra-pair skew can lead to a sharp rise in bit error rate (BER), especially at data rates of several gigabits per second or higher. Once skew is added, further BER degradation occurs. Generally, the higher the order of Pseudo Random Bit Sequence (PRBS) the higher the BER (see ***Figure 2***), as timing mismatches compromise the receiver's ability to correctly interpret the data.

Accounting for pattern sensitivity, the BER impact from intra-pair skew can range from between 0.5 to 1 order of magnitude at 0.2 UI, so it can be a significant contributor to overall system margins. Test solutions and protocols must have tight control over intra-pair skew to ensure signal integrity and optimal data transmission quality as a result.

It is clear that intra-pair skew is an important factor to consider in system design. To overcome intra-pair skew challenges associated with high-speed digital interfaces, advanced and innovative evaluation solutions must be implemented to ensure highly accurate measurements.

Fig. 3 A high-quality BERT is needed for verifying high-speed digital interfaces currently in development.

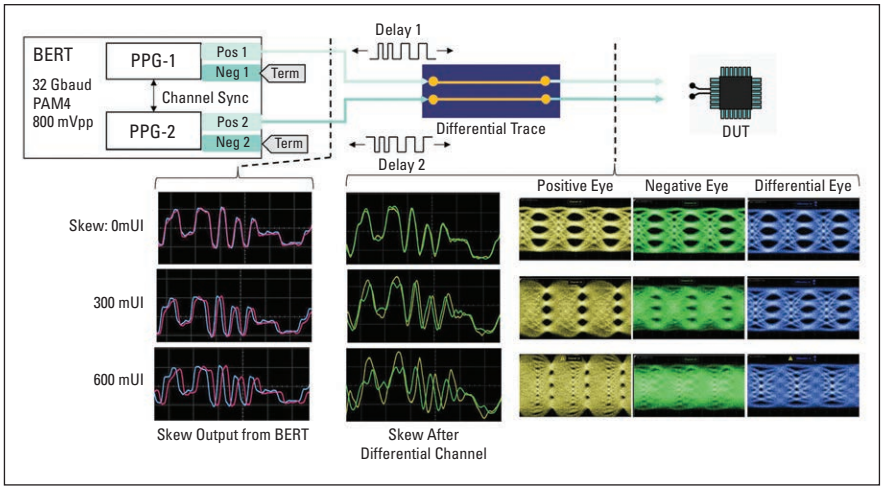

Fig. 4 Test setup and results of the effect of skew on PAM4 eye at 64 GT/s.

Inadequate Traditional Testing Methods

The traditional arbitrary intra-pair skew approach uses various coaxial adapters from multiple vendors to introduce an artificial mechanical delay or skew. It has become the conventional method because it allows engineers to step up skew in discrete amounts and allows skew to be introduced anywhere in the channel where there are coaxial connections.

The problem with this method is that it generates skew by creating what is called a trombone trace. In effect, a trombone trace is a routing pattern on a PCB designed to increase the length of one of the traces in a differential pair, thereby compensating for intra-pair skew. It adds a series of bends or loops, which can create significant challenges to accurately characterize designs. The bends and turns often create variations in impedance, leading to signal reflections.

Secondly, a trombone trace on one leg often has a different delay per unit length compared to the straight trace on the other leg. The adverse effect is mode conversion that may negatively impact the signal.

Another major issue with the arbitrary intra-pair skew approach is that the adapters that often introduce the skew can vary in length by as much as ~0.2 in. This makes it difficult to quantify the skew. Further, the method simply cannot truly support higher data rates, in large part because the skew steps are not very granular.

New Test System Configuration

Because the traditional coaxial adapter approach is unsatisfactory for modern high-speed interface designs, an improved test approach is necessary. The new methodology incorporates a high-speed Bit Error Rate Tester (BERT) with integrated PPGs for measuring performance under varying skew conditions.

Enabling this capability requires the use of dual transmitters on the BERT (see ***Figure 3***). Dual transmitters are used for single-ended control of the phase of the signals within the differential pair. The transmitters need to be clock- and jitter-synchronized, yet are independently controlled by phase from the second PPG module. The setup is not currently in the standard CEM Compliance Test configuration, but it is suggested for accurate skew testing.

The recommended test setup has a skew on a 64 GT/s differential channel. Differential signals are generated from the outputs of each PPG, allowing for precise control of skew. Both PPGs drive the same pattern, but the second PPG is a logic inversion. Data NEG on both PPGs are terminated at 50 Ω at the source or fed to an oscilloscope. The skewed signal also can be passed through a noise generator, if desired.

The synthesizer serves as the clock source. Optionally, the clock can be fed into a synthesizer from a 100 MHz source. The jitter module is used to produce multiple synchronized clocks to feed each PPG and add jitter. Skew is introduced by a delay block inside each PPG. The delay can be advanced or regressed by dialing skew in POS or NEG direction as needed.

Need for Calibration

Given the precise measurements required to verify today's high-speed designs, calibration is important for accurate characterization. It ensures the best match between POS and NEG. Parameters that must be matched between the two PPGs during calibration include amplitude, edge rate, duty cycle, jitter, and equalization.

Figure 4 shows results from the test setup. The display reveals the effect of skew on a PAM4 eye at 64 GT/s of a PCIe 6.0 interface with 2 mUI resolution, which is 0.0625 ps for PCIe 6.0. Based on the results, the intra-pair skew can have quantifiable impact to BER margins on PCIe at 64 GT/s, even without any additional impairments.

Summary

The increasing speeds of digital interfaces to meet the ever-growing data demands of modern society places stress on design engineers. Conventional test methods cannot support today's high-speed designs. For this reason, a new test methodology for measuring intra-pair skew is necessary to verify high-speed interfaces such as PCI Express.

The BERT-based setup uses a new method with dual transmitters to control the phase of signals within a differential pair, enabling granular measurement of intra-pair skew. This capability is crucial for understanding and mitigating the impact of skew on BER margins at high data rates such as 64 GT/s.■

Beat Skew. Reduce Loss. Boost Speed.

Rogers High Frequency Laminates. Trusted & Essential.

Material-induced skew doesn't have to slow your design. Rogers high-frequency laminates are engineered to minimize glass weave effects, ensure signal integrity, and support your next-gen systems at **224 Gbps** and beyond.

VISIT US TO EXPLORE MORE >>

ROGERS CORPORATION

ROGERSCORP.COM

Using Ultra-Broadband Baluns to Perform Differential S-Parameter Measurements Using Single-Ended 2-Port VNA

Madrone Coopwood, Jason Yoho Ph.D., and James R. Andrews Ph.D.
Hyperlabs, Louisville, Colo.

Most high-performance test instruments, including sampling oscilloscopes and vector network analyzers (VNAs), are single-ended, ground-referenced, 50 Ω instruments. There are now a growing number of electronic applications, both digital and analog, that are using differential circuit techniques and also balanced transmission lines. As a result, many engineers are being tasked to design new differential circuits without the benefit of having proper test instruments.

Most VNAs can be upgraded by the manufacturers by purchasing an optional, expensive, multi-port switch matrix to accompany 2-port VNAs. This article will offer a significantly lower-cost alternative. Instead of using a switch matrix, baluns can be used with conventional 2-port VNAs to perform differential measurements. Several companies offer ultra-wideband baluns that permit differential measurements to be made over extremely wide frequencies extending from the kHz range to well into the GHz band.

Baluns

Historically, when engineers needed to do differential measurements with conventional test instruments, they used a balanced to unbalanced transformer, or "balun" for short. ***Figure 1*** shows the typical schematic diagram for a balun. It consists of a simple transformer with one wire of the primary winding being the ground terminal for the unbalanced side. The balanced secondary winding is not connected to the ground terminal and is thus considered to be "floating" with respect to ground. Impedance transformation is also possible if the number of wire turns on the primary and secondary are unequal. The impedance transformation is equal to the square of the turns ratio (N). A classic example of a balun, which some may be familiar with, is the antenna transformer that used to be supplied with every TV receiver. It was used to match 300 Ω flat ribbon lead to 75 Ω coax cable. For a 300 Ω /75 Ω transformation, a 2:1 turns ratio is required.

Figure 2 shows another example of a balun. In this case, the balanced secondary consists of two identical windings that are connected as a center-tapped secondary. The

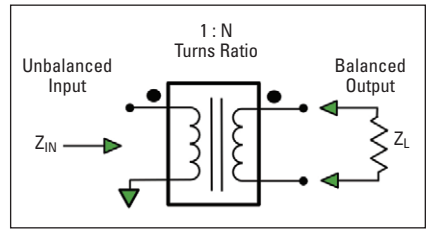

Fig. 1 Balun transformer schematic diagram.

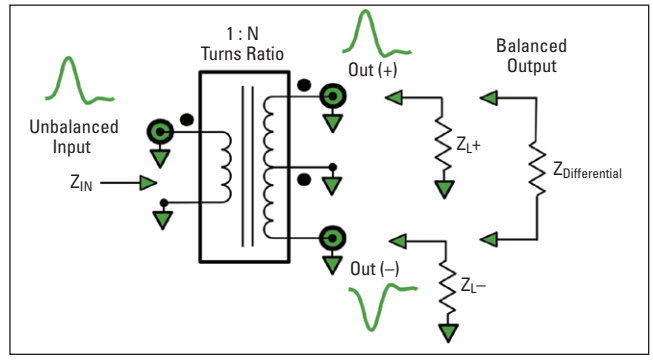

Fig. 2 Balun transformer with center-tapped output schematic diagram.

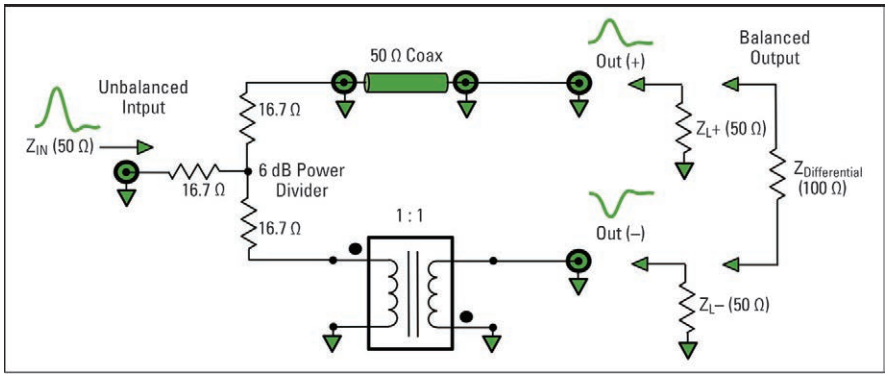

Fig. 3 Ultra-wideband balun transformer schematic diagram.

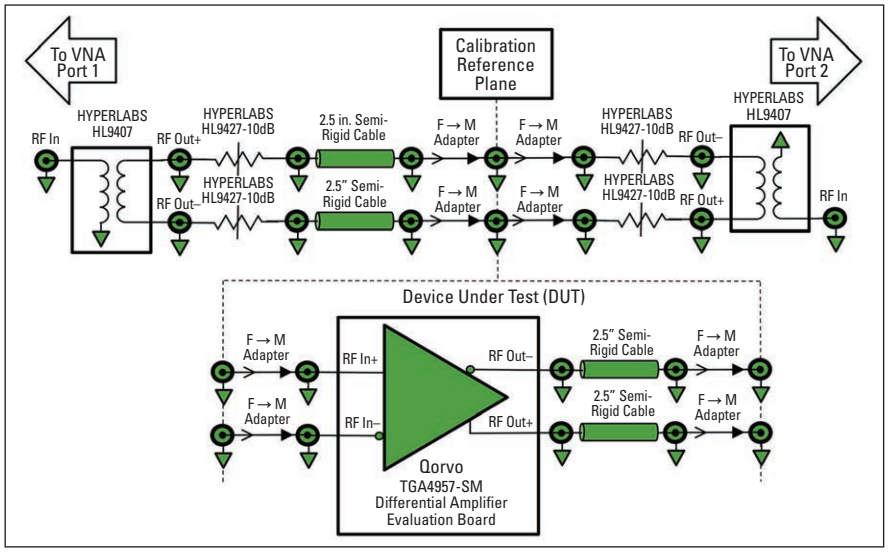

Fig. 4 Block diagram of 2-port VNA measurement system utilizing two baluns.

center tap is usually then connected to the common ground. Note that the black dots are polarity indicators for the various transformer windings. With the arrangement shown in Figure 2, one of the secondary outputs is "in-phase" with the input and is thus labeled as the (+), or "non-inverting" output. The other secondary output is "out-of-phase" with the input and is thus labeled as the (-), or "inverting" output. There is a 180-degree phase difference between the non-inverting and inverting outputs. There are now three output impedances to be considered. Z_{L^+} and Z_{L^-} are the impedances referenced to ground seen looking into the non-inverting and inverting outputs, respectively. There is also a differential impedance, $Z_{DIFFERENTIAL}$, which is the impedance seen between the two center pins of the (+) and (-) output coax connectors. $Z_{DIFFERENTIAL} = Z_{L^+} + Z_{L^-}$. The impedance transformation is still determined by the turns ratio, N, of the secondary and primary windings.

The balun designs shown in Figure 1 and Figure 2 are very widely used and are available from many manufacturers. They can be used for balanced to unbalanced transformations and to shift impedance levels by altering the turns ratio, N. The major limitation in these designs is bandwidth. They are built using conventional transformer designs and techniques. The transformer core material, number of wire turns, etc. are dictated by the desired operating frequency. It is difficult to design transformers, including baluns, to operate over more than one or two decades of bandwidth. For typical "wireless" RF applications, ultrawide bandwidth is not a requirement. For example, the TV antenna transformer mentioned earlier only needs to work from 50 to 800 MHz. However, for digital data, ultrawide bandwidth is a mandatory requirement. Digital data systems, such as 5G, PCIe 6.0, and USB 3.0 require bandwidths extending from applications, the balun designs of Figure 1 and Figure 2 are unsatisfactory.

There are several methodologies for designing an ultra-wideband balun design that works over many decades of bandwidth. One of these designs is shown in ***Figure 3***. It consists of a 50 Ω, impedance-matched, 6 dB power divider, a 50 Ω coaxial inverting 1:1 transformer, and a length of 50 Ω coax cable.

In this version of the design, the input signal is split into two identical signals by the 6 dB power divider. One of these signals is then inverted (180-degree phase shift) by the 1:1 inverting transformer. The other signal is sent through a coax cable whose length is chosen to match the propagation delay time of the 1:1 inverting transformer. The input impedance is 50 Ω. The output impedances of both the (+) and (-) coaxial outputs are also 50 Ω. The differential output impedance is thus 100 Ω. The 1:1 inverting transformer is a specialty design (sometimes referred to as a pulse inverter) that is a hybrid of coax cable and conventional transformer designs. This design concept results in 1:1 inverting transformers with more than six decades of bandwidth. The major limitation in this balun design is the 3 dB of loss suffered in the 6 dB power divider and that the impedance transformation is limited to 2:1 (i.e., 100 Ω differential output to 50 Ω single-ended input).

Differential S-parameter measurements using a single-ended 2-port VNA and a pair of ultra-broadband baluns are compared against measurements taken on a 4-port VNA. The device under test (DUT) is a broadband differential amplifier. The following portion of this article builds on prior work reported in an application note AN-21.[1] The referenced application note demonstrated differential S-parameter measurements to 10 GHz using broadband baluns. Utilizing the industry's higher bandwidth baluns, the same technique is shown to yield accurate differential S-parameter measurements to 40 GHz and beyond from a single-ended 2-port network analyzer.

2-Port VNA Measurement System

Starting with a 2-port VNA, begin with a typical setup external to the VNA to create a new calibration reference plane. A block diagram of the 2-port to 4-port measurement system is shown in ***Figure 4*** (the VNA is an Anritsu MS4644B and the baluns are HYPERLABS HL9407). These baluns feature -3 dB bandwidth from 500 kHz to

Fig. 5 Open reflection calibration of port 1.

67 GHz and have excellent phase and amplitude matching at the balanced port. To improve the impedance match of the differential test ports, 10 dB attenuators (HYPERLABS HL9427-10) are added to each of the four differential test ports.

As shown in Figure 4, the 2-port VNA measurement system was constructed with a male differential test port 1 and a female differential test port 2. This arrangement facilitates zero-length thru calibration of the VNA. The ports of the DUT were configured female-male accordingly for easy insertion. The connector spacing of the differential amplifier evaluation board does not match the connector spacing of the balun, so it was necessary to use interface cables. Short semi-rigid VNA loops were employed to interface the incompatible connector spacings. One set of cables was incorporated into test port 1, and the other set of cables became part of the DUT.

2-Port VNA Calibration

A 12-Term SOLT calibration was performed including isolation using two Anritsu 3652 "K" calibration kits. Each kit contains one female open, one female short, one male open, and one male short calibration standard. Using two kits, it was possible to simultaneously connect two female calibration standards to differential port 1 for each set of Short, Open, and Load reflection calibration.

Figure 5 shows two female and two male open calibration standards, connected to differential test port 1 and differential port 2, respectively. The calibration of port 1 and port 2 was completed simultaneously by cycling through the Short, Open, and Load reflection calibration standards. This approach eliminated the need to swap out equi-phase test port adapters during calibration.

Only one calibration kit's coefficients could be loaded into the MS4644B VNA. A small calibration error results from minor differences between the two calibration kits.

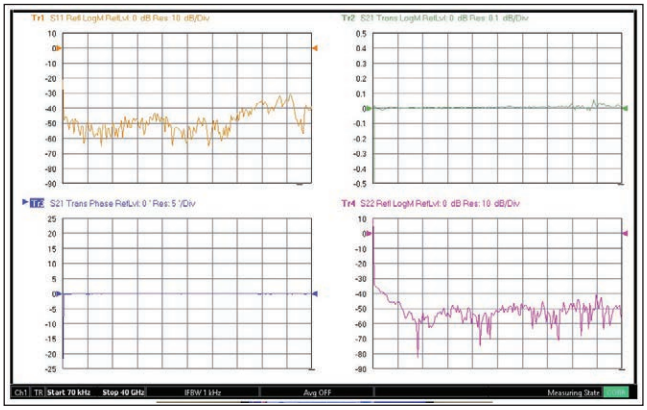

Fig. 6 Thru verification of differential calibration.

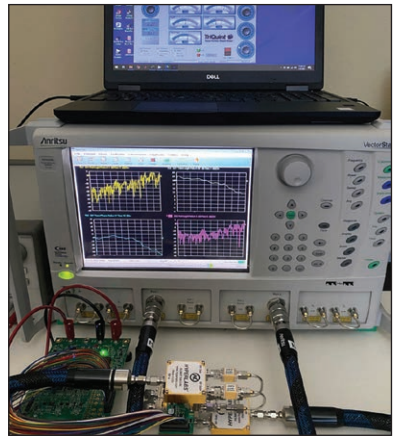

Fig. 7 Active 2-port differential measurement.

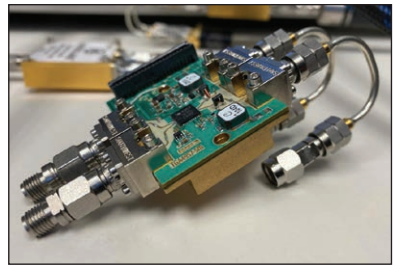

Fig. 8 Defined DUT.

The quality of the resulting calibration is shown in ***Figure 6***, the active 2-port differential measurement is shown in ***Figure 7***, and the defined DUT is shown in ***Figure 8***.

4-Port VNA Calibration and Measurement System

A block diagram of the 4-port measurement system is shown in ***Figure 9***. The VNA is Anritsu MS4647B with MN4697C Multi-Port Test Set. Calibrations were performed using Anritsu 36585V-2F Precision AutoCal.

After completing AutoCal, a Thru Update was performed for all port combinations shown in Figure 9 using Anritsu 33VFVF50C female-female adapter (23.62 mm). The delay of the 23.62 mm adapter was entered as a calibration coefficient and corrected by the Thru Update operation. However, the slight attenuation of the adapter was not corrected. As a result, insertion gain measurements taken on the 4-port VNA measurement system are slightly optimistic.

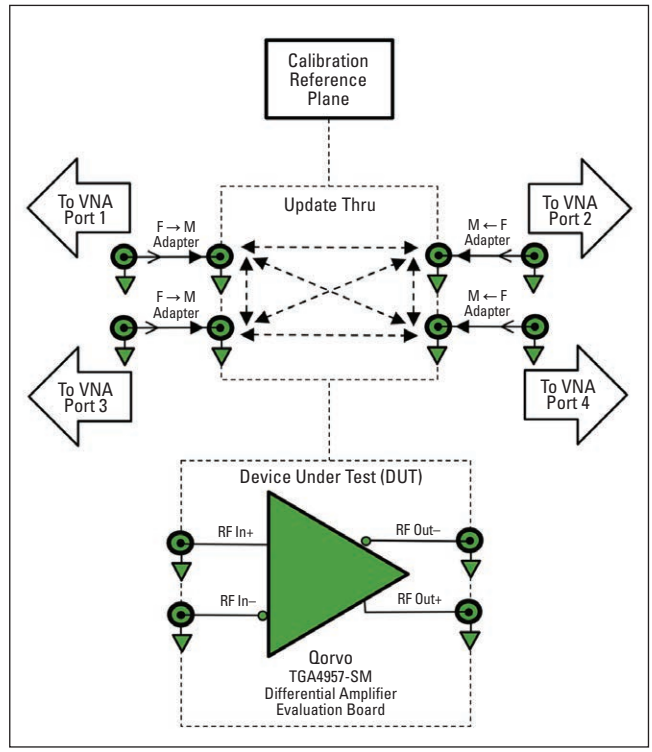

Fig. 9 Block diagram of 4-port VNA measurement system.

Leaders in High-fidelity Signal Solutions

Maury Microwave now unites under one identity, one website, and one vision, with a comprehensive portfolio of high-integrity signal solutions that drives innovation from concept to field operation. Noise sources, signal generators, synthesizers, and modules provide advanced, flexible signal generation, from low phase noise to wide bandwidth. Solid-state amplifiers deliver broadband, high-power performance with excellent linearity, while StabilityPlus™ cables maintain amplitude and phase stability with flexure. Together, these solutions deliver the highest signal fidelity to empower the future of connectivity.

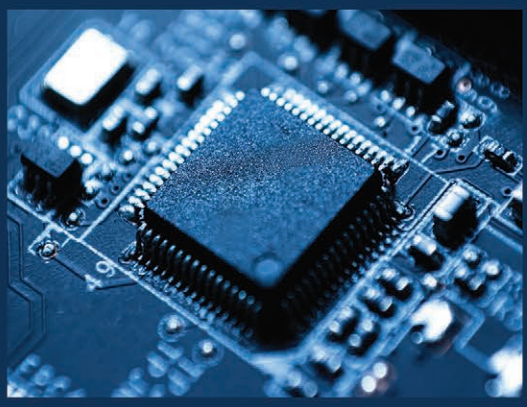

From R&D...

SGX1000 RF Signal Generators, UFX Noise Generators and MPA Series Amplifiers

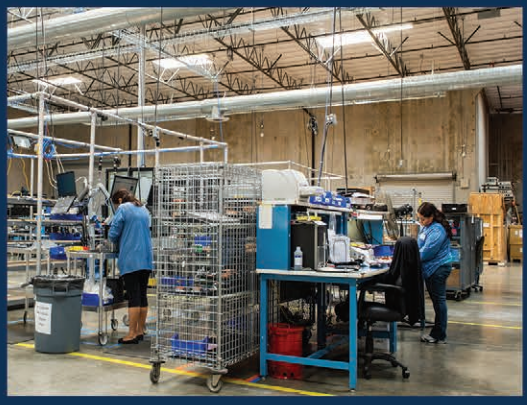

...Through Manufacturing...

HS, HSX, and HSY Series Ultra Low Phase Noise Multi-channel RF Synthesizers

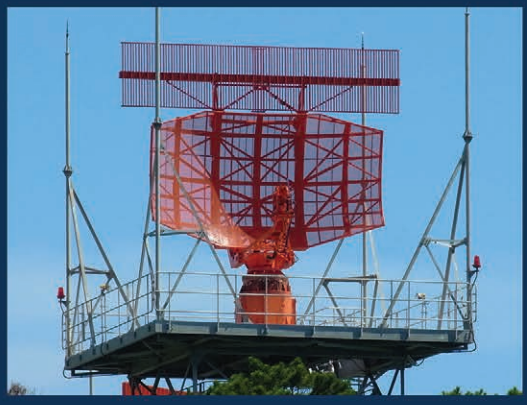

...To Deployment

HSM Series RF Synthesizer Modules and Interconnects

Learn More at maurymw.com

maury

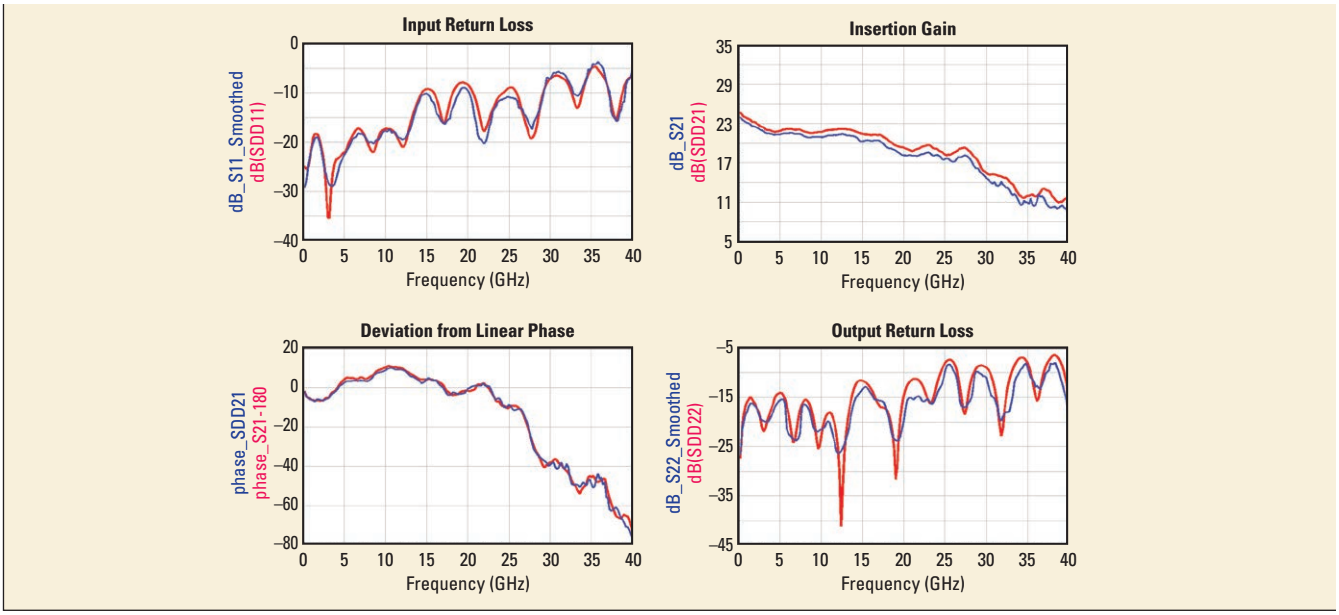

Fig. 10 Measured differential S-parameters (4-port VNA in red; 2-port VNA, baluns, and attenuators in blue).

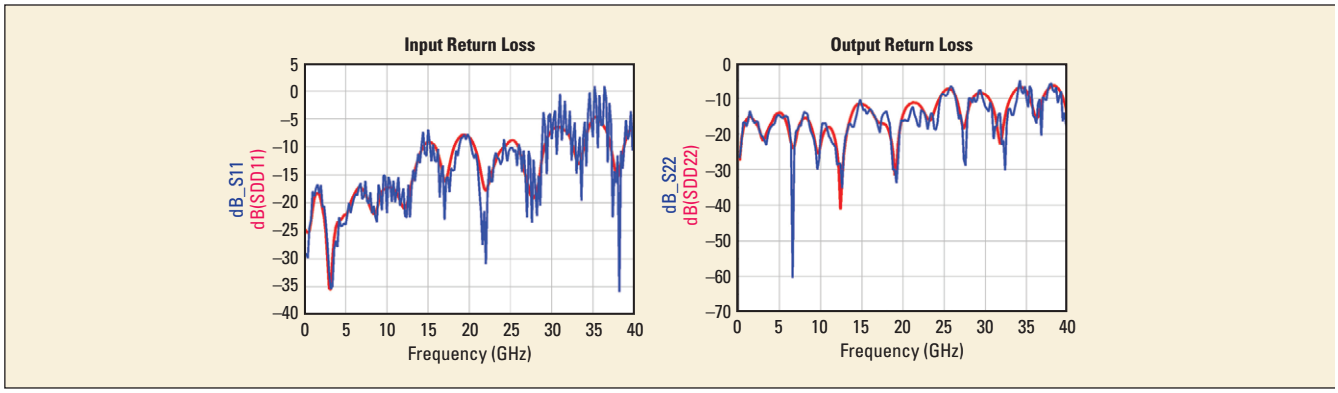

Fig. 11 Comparison of raw return loss data.

S-Parameter Test Results

Measurements obtained from both test systems are compared in ***Figure 10***. The red traces represent data collected on the 4-port VNA measurement system, and the blue traces represent data collected on the 2-port VNA system using HL9407 baluns.

The return loss data from the 2-port system is noisy and was smoothed with a 4.5% moving average in the plots of Figure 10. The raw data is presented in ***Figure 11***. Port 1 return loss data is noisier than port 2, which may be related to the semi-rigid cables that were integrated into port 1 of the calibrated measurement system. The cables may have deformed due to torquing and flexing after calibration.

The S21 gain response measured with the 4-port measurements system is notably higher than S21 measured on the 2-port measurement system. This can be attributed to two factors. First, the 4-port measurement system was used to characterize a DUT that did not include any coaxial adapters nor cables. In contrast, the 2-port measurements system was used to characterize a DUT that included adapters and cables. Second, as previously mentioned, the Thru Update used during calibration of the 4-port measurement system did not account for insertion loss of the Anritsu 33VFVF50C female-female adapter. These two factors explain the discrepancy in measured S21 response.

Conclusion

This article demonstrates accurate differential S-parameter measurements obtained from a single-ended 2-port VNA using ultra-broadband baluns and attenuators. This measurement system is a cost-effective alternative to purchasing a multi-port test set for a VNA. It is worthy of note that the differential S-parameters reported in this article are only a subset of the full mixed mode S-parameters that were obtained from the 4-port measurement system. The 4-port measurements system does yield common mode and mode conversion S-parameters in addition to differential S-parameters. ■

REFERENCE

1. J. R. Andrews, "Differential VNA Measurements Using Single-Ended, Two Port Instruments and BALUNs," *Picosecond Pulse Labs Application Note* AN-21, Dec. 2008. Currently available at: www.kh6htv.com/pspl-app-notes.

YOU CAN'T HAVE ACCURATE DATA WITHOUT ACCURATE CALIBRATION

AVAILABLE NOW

"We're developing the tools that define how engineers validate power integrity in the next decade."

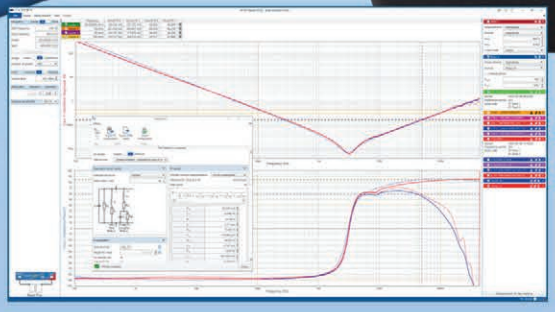

Picotest's latest calibration solutions yield the most precise results, necessary for today's high-speed, low-noise designs.

-We Supply-

Calibration Tools: Precision kits, ground isolators, fixtures
Measurement Solutions: VNAs, Injectors, Probes, TDR
Power Integrity Tools: Transient step load testers, impedance measurement, stability/PSRR validation
Support: Application notes, test boards

PICOTEST

picotest.com

Material-Induced Skew in High-Speed Multilayer PCBs: Influences and Mitigation Strategies

Bob Nurmi, John Coonrod, and Vitali Judin
Rogers Corporation, Chandler, Ariz.

In the relentless push toward faster digital communication and higher data throughput, signal integrity (SI) has emerged as a fundamental design challenge. Among various forms of signal degradation, skew — defined as the differential timing delay between two logically paired signals, often in differential pairs — is a key contributor to timing errors in systems operating at multi-gigabit data rates.

While conductor length mismatch is a known cause, an increasingly important factor at high speeds is material-induced skew due to local differences in dielectric constant arising from the PCB's woven glass fabric reinforcement. This article aims to unpack the root causes of material-induced skew, particularly focusing on the glass weave effect, and explores mitigation techniques ranging from laminate selection to signal routing strategies.[1]

Skew in High-Speed Circuits

Skew is the time delay difference between two signals that should ideally arrive simultaneously. In differential pairs, skew disrupts common-mode noise rejection and can significantly degrade eye diagrams, jitter margins, and bit-error rates. It becomes particularly critical in systems running above 10 Gbps, where picosecond-level timing mismatches can impact signal interpretation.

Origins of the Glass Weave Effect

Most rigid multilayer PCBs use woven glass fiber reinforcement in the dielectric to provide mechanical strength and thermal stability. This woven glass has a Dk of approximately 6 for standard e-glass to 4.5 for newer generation low Dk glass, while the resin matrix (e.g., Thermoset (PPE, Hydrocarbon) or PTFE-based systems) ranges from 2.1 to 3.8 Dk.

When a trace overlays different portions of this weave — such as a glass bundle (knuckle) versus an open resin pocket — the effective Dk seen by the trace can vary, leading to differential propagation delays between conductors in a pair.[2] Over lengths of several inches, even small local variations can integrate into measurable skew, especially for thin dielectric layers or small-width conductors. ***Figure 1*** shows how the placement of the conductor can overlap different portions of the PCB glass and resin, resulting in variations in dielectric constant, causing skew.

There are many different glass styles, with some common standard glass styles listed in ***Figure 2*** alongside their dimensions. These three glass styles were used in a recent study to quantify the glass weave effect.

Quantifying the Impact of the Glass Weave Effect

After the circuits were fabricated on these three styles of glass PCBs, thorough inspections were done to choose the appropriate circuits for testing (see ***Figure 3***). Chosen for testing were circuits with the signal conductor aligned to the glass knuckle-bundle run (high Dk), and circuits with the signal conduc-

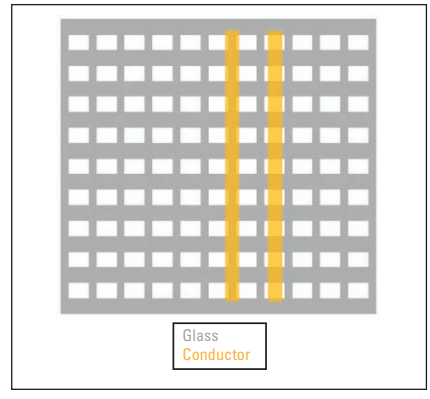

Fig. 1 **Diagram showing how the placement of the conductor can overlap different portions of the PCB glass and resin, resulting in variations in dielectric constant and therefore causing skew.**

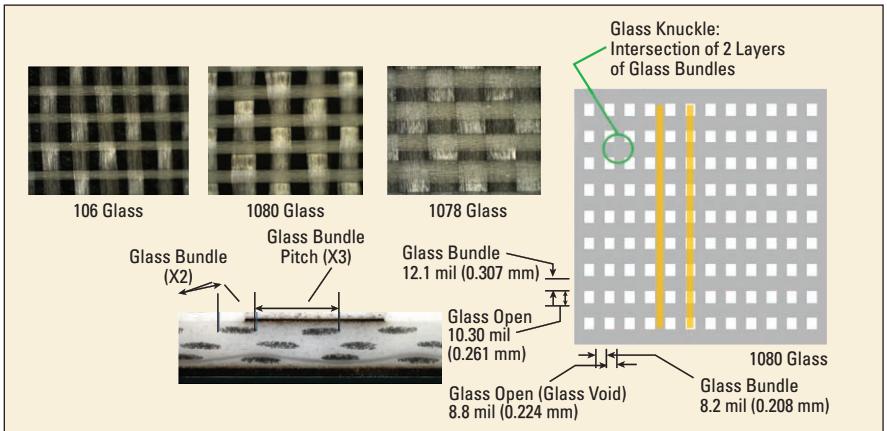

Glass Style	Type	Dimensions (mils)					
		X1	X2	X3	Y1	Y2	Y3
106	Open Weave, Balanced	1.00	4.80	18.50	0.60	10.20	20.60
1078	Spread, Balanced	1.60	8.20	18.50	1.10	12.10	18.50
1080	Open Weave, Unbalanced	1.60	8.20	17.00	1.10	12.10	22.40

(X1) is Glass Bundle Thickness

Fig. 2 **Some common standard glass styles with dimensions of glass bundles vs. resin.**

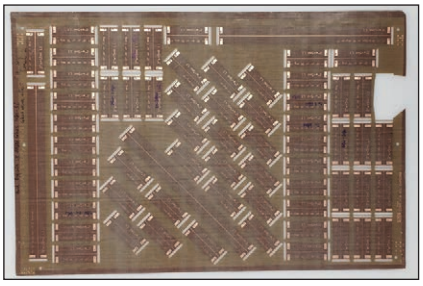

Fig. 3 **Test circuits fabricated on different glass styles of PCBs for testing.**

40 to 60 circuits.

Empirical Measurement Techniques

Time domain reflectometry and vector network analyzers were used with rise times down to 3.2 ps and bandwidths up to 110 GHz. Fast rise time is essential to resolve fine impedance variations from weave transitions. Impedance anomalies as small as 0.5 mils in trace width were detected — sufficient to produce a 1 Ω impedance swing. A network analyzer was used to measure the following properties: phase angle (unwrapped), group delay (based on phase angle which varies with frequency), propagation delay, effective Dk measurements based on phase angle, and impedance from reflected S11 and S22. Comparisons of these properties were done with circuits using 4 mil thick PTFE-woven-glass with 106 glass, 4 mil thick PTFE-woven-glass with 1080 glass, 4 mil thick PTFE-woven-glass with 1078 glass, and 4 mil thick RO4835™ LoPro® laminate with 1080 glass. Rogers Corporation conducted extensive testing of various laminates using microstrip test vehicles at frequencies up to 110 GHz.

Across three glass styles — 106 (standard), 1080 (open weave), and 1078 (spread weave) — data showed significant variation in performance due to the interaction between signal traces and localized glass weave features.

TABLE 1

GROUP DELAY MEASUREMENTS FOR EACH TYPE OF GLASS STYLE

Glass Type	*Δ Group Delay*	*Equivalent ΔDk*
106 (open weave)	4.7 ps	0.11
1080 (open weave)	7.3 ps	0.17
1078 (spread weave)	1.0 ps	0.02

Measured Results

The average group delay data differences between traces aligned with knuckle-bundle versus bundle-open regions are shown in ***Table 1***. Example measurements for 4 mil PTFE 106 glass style are shown in ***Figure 4***.

Using measured phase angle with the microstrip phase response formula finds effective Dk. Using Eff Dk, propagation velocity can be found, and from that, propagation delay can be found. A higher Dk will have a slower wave, which is increased propagation delay. For propagation delay, the average difference between 40 GHz and 80 GHz is 6.9 ps, which is equivalent to a Dk difference of 0.15.

The test results for phase angle show the difference at 77 GHz of 100°, which is equivalent to a Dk difference of 0.09 (see ***Figure 5***). This is the most accurate measurement given that it uses raw phase angle measurements from the network analyzer. A higher Dk will have an increased phase angle (more negative value, as formatted below). Test results for phase angle show the equivalence of 0.09 Dk at 77 GHz.

The average impedance difference is 3.1 Ω, which is equivalent to a difference in Dk of 0.40. The general trends are correct, but the impedance value is not correct for extracting Dk only. The circuits had a difference in conductor width and copper plated thickness that alters the impedance values much more than the glass-weave effect on Dk. The time axis for the impedance curve is

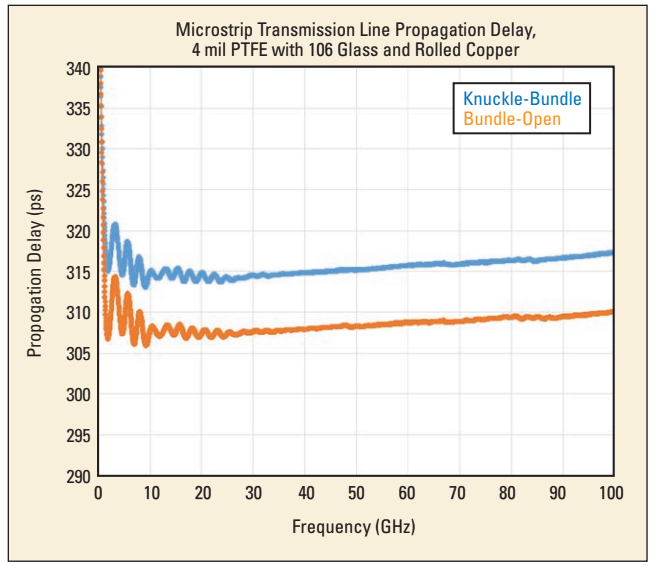

Fig. 4 **Microstrip transmission line propagation delay for 4 mil PTFE with 106 glass and rolled copper.**

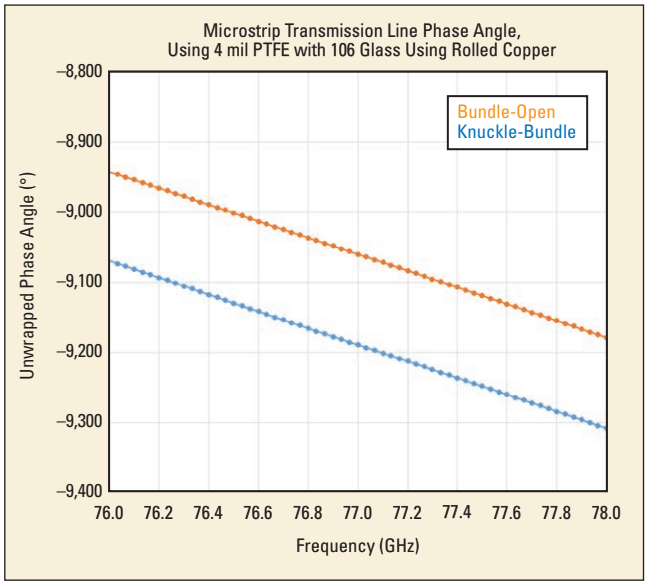

Fig. 5 **Microstrip transmission line phase angle for 4 mil PTFE with 106 glass and rolled copper.**

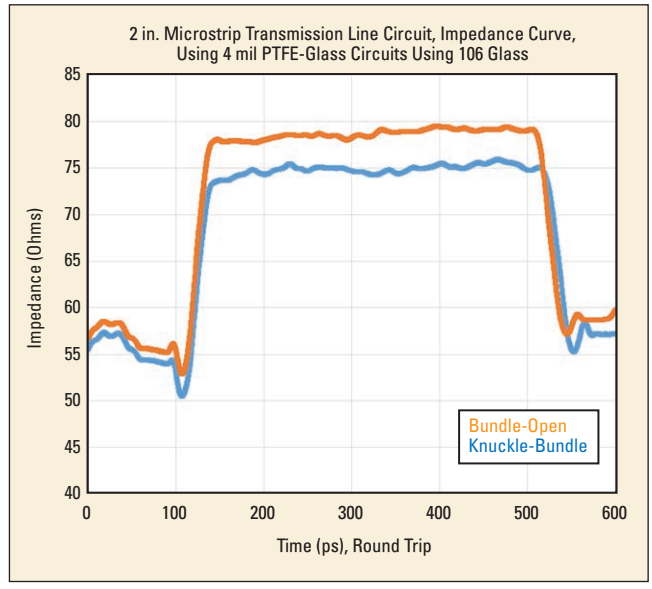

Fig. 6 **2 in. microstrip transmission line circuit impedance measurement for 4 mil PTFE with 106 glass and rolled copper.**

round trip, or the time is half, as shown in ***Figure 6***. A higher Dk will have a lower impedance and longer (slower) propagation time.

This summary of the data in ***Table 2*** highlights how spread weave glass styles significantly reduce the skew-inducing effects by smoothing Dk variation.

Influences Beyond Glass Weave: Laminate Composition

Ceramic-filled laminates, such as RO4835™ LoPro, shown in the test results, mitigate skew by smoothing Dk transitions between glass and resin. Compared to unfilled PTFE-glass laminates, ceramic-filled systems demonstrated lower Dk variation and minimal skew, even when using traditional 1080 glass styles.

A laminate with filler, which is usually a different Dk than the glass and the resin system, will help to minimize the Dk transitions between the glass fabric and the resin system (see ***Figure 7***). Smoothing the Dk transitions is helpful for reducing the glass-weave effect; however, the effect can still be observed, although to a lesser extent than unfilled glass reinforced laminates.

Table 3 shows test results for 1080 glass with laminate unfilled vs. filled, proving that a ceramic filled laminate significantly reduces the glass-weave effect on skew. Further to this improvement, RO4835T™ (a ceramic loaded thermoset similar to 4835 LoPro) and XtremeSpeed GB series© (a ceramic loaded PTFE) use advanced glass material specifically for high-speed, low-loss digital applications. These materials are constructed with spread glass 1078 and 1035 to further even the Dk distribution in the XY plane.

Copper Foil Surface Roughness

Rolled copper, being smoother than electro-deposited copper, contributes less impedance variation and delay uncertainty. In thin laminates, even small variations in copper surface texture can alter propagation velocity and phase delay, indirectly contributing to skew.

The results from this study used laminates with electro-deposited (ED) copper. Most ED copper does not have directionality. Rolled copper is extremely smooth and has a natural directionality, as one axis of the copper is slightly rougher than the other axis. The copper directionality typically has a small influence on circuit performance due to

TABLE 2

OVERVIEW OF THE TEST RESULTS MEASURED BETWEEN 40 TO 80 GHZ, FOR GROUP DELAY, PROPAGATION DELAY, PHASE ANGLE, AND THE EFFECTIVE CHANGE IN DK FOR EACH GLASS STYLE (106, 1078, AND 1080)

		Average Differences Between Knuckle-Bundle and Bundle-Open		
Glass Style	*Type*	*40 to 80 GHz*		*77 GHz Phase Angle (°)*
		Δ Group Delay (ps)	*Prop Delay (ps)*	
106	open weave, balanced	4.7	6.9	100
1078	spread, balanced	1	1.3	20
1080	open weave, unbalanced	7.3	10.1	149
		Equivalent Difference in Dk (ΔDk) Between Knuckle-Bundle and Bundle-Open		
Glass Style	*Type*	*40 to 80 GHz*		*77 GHz Δ Dk from Phase Angle*
		Δ Dk from Group Delay	*Δ Dk from Propagation Delay*	
106	open weave, balanced	0.11	0.15	0.09
1078	spread, balanced	0.02	0.03	0.02
1080	open weave, unbalanced	0.17	0.22	0.14

Unlock High-Speed Insights with the **QCA Series**

QCA Series digital sampling oscilloscopes deliver powerful, cost-effective performance and exceptional measurement accuracy demanded by today's leading engineers.

- High-quality **precision timebase** and ultra-low jitter, comparable to industry-leading performance.
- Supports **optical and electrical signals** up to 58 GBaud (NRZ and PAM4).
- Easy-to-use **VISEYE™** jitter measurement and eye diagram analysis software.
- Pair with the **QCR Series clock recovery instrument** for optimal measurement accuracy.
- Optimized for **design verification testing (DVT)** and **high-volume manufacturing (HVM)** of high-speed interconnects, ASICs, GPUs, PICs, CPO, and transceiver modules.
- Contact us for a demonstration today.

www.quantifiphotonics.com/QCA

QUANTIFI PHOTONICS®
A Teradyne Company

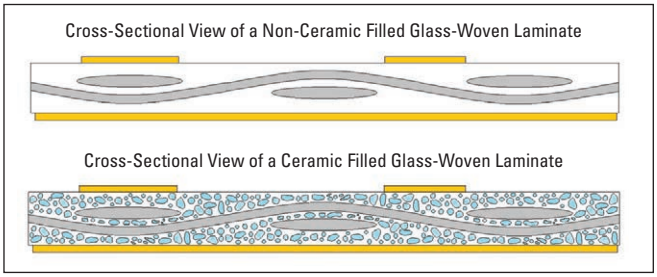

Fig. 7 **Cross section view of a non-ceramic filled glass-woven laminate versus ceramic filled.**

TABLE 3

CERAMIC-FILLED LAMINATES, SUCH AS RO4835™ LOPRO, COMPARED TO UNFILLED PTFE-GLASS LAMINATES, MITIGATE SKEW BY SMOOTHING DK TRANSITIONS BETWEEN GLASS AND RESIN EVEN WHEN USING TRADITIONAL 1080 GLASS STYLES

	Average Differences Between Knuckle-Bundle and Bundle-Open		
Laminate	*40 to 80 GHz*		*77 GHz Phase Angle (°)*
	Group Delay (ps)	*Prop Delay (ps)*	
PTFE with 1080 glass	7.3	10.1	149
RO4835 LoPro with 1080 glass	0.3	0.6	10
	Equivalent Difference in Dk (ΔDk) Between Knuckle-Bundle and Bundle-Open		
Laminate	*40 to 80 GHz*		*77 GHz Δ Dk from Phase Angle*
	Δ Dk from Group Delay	*Δ Dk from Propagation Delay*	
PTFE with 1080 glass	0.17	0.22	0.14
RO4835 LoPro with 1080 glass	0.01	0.02	0.01

panel orientation, but for thin circuits, it can have some impact. The surface of rolled copper is shown in ***Figure 8*** at high magnification from a scanning electron microscope photograph.

Surface roughness is slightly different for rolled copper along the grain direction vs. cross-grain direction. Rougher copper will slow the phase velocity and alter the phase angle.

Skew Mitigation Strategies: Material Construction

Using materials that have no glass fabric, such as RO3003™ and RO3003G2™, will yield the best results for minimizing skew. However, this is not practical in constructing high layer count, tight pitch, dense PCBs, as registration of inner layer circuits typically requires the use of woven glass for dimensional stability.

Minimizing the amount of glass for a given dielectric thickness to maximize the spacing between conductor and glass, using laminate and prepreg bond ply materials such as RO4835T™/RO4450T™, and the newest material soon to be released for Gen 9 designs,

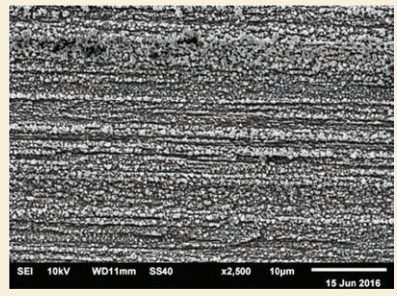

Grain Direction is from Left to Right

Fig. 8 **Surface texture of rolled copper as seen on a scanning electron microscope.**

XtremeSpeed RO1201™ laminate constructed with readily available low Dk spread glass and RO1101B prepreg bond-ply without any glass fabric combined with resin formulations that contains ceramic fillers for extremely low-loss multilayer stripline, will minimize or eliminate skew effects, differential pair 224Gb/s PAM4 performance, and avoid the use of limited supply glass and quartz.

Circuit Orientation

Skew mitigation has also included PCB rotation of ~12° relative to the glass cloth directional weave of the laminate and prepreg supplied panels.[3] Zig-zag routing also helps to average out local Dk transitions by traversing multiple wave regions.[3] These are intended to route a differential pair such that both traces experience the same dielectric zones.

With PCB rotation, the downside is that efficient utilization of the materials is lost. For example, if two PCBs each occupy 12 × 18 in. on a standard 18 × 24 in. panel, a rotation of 12° will yield only one PCB per 18 ×24 in., resulting in twice the material cost per PCB.

Conclusion

As circuit technology continues to progress to higher frequencies and data rates, PCB materials have an increasing influence on the performance of the system. Skew is an important bit error rate property influenced by the micro-local area Dk differences caused by woven glass fabric used in the manufacturing of PCB laminate and prepreg bonding materials. There has been significant progress in smoothing out the Dk differences in glass fabrics through balance square weave and spread glass designs as well as through the lowering of Dk in glass composition from 6.6 to 4.5 to reduce the difference between glass and the coated resin. Additionally, ceramic loading in resin formulations by the laminate manufacturers also helps to smooth out the Dk in XY planes, and the construction of laminates with higher resin content/thickness between the glass fabric and surface conductor (as with RO1201 laminate) have minimized or eliminated skew. Ultimately, no glass fabric (as with prepreg bond-ply RO1101B) eliminates skew.■

REFERENCES

1. Paper adapted from J. Coonrod Presentation, "An Overview of Glass-Weave Impact on Millimeter-wave PCB Performance," *Microwave Journal* webinar, October 2018.
2. A. F. Horn, III, J. W. Reynolds, and J. C. Rauto, "Conductor Profile Effects on the Propagation Constant of Microstrip Transmission Lines," *IEEE Tran. Microw. Trans Theory Tech. Symposium*, 2010.
3. J. Loyer, R. Kunze, and X. Ye, "Fiber Weave Effect: Practical Impact Analysis and Mitigation Strategies," *DesignCon*, Santa Clara, 2007.

SIGNAL EDGE SOLUTIONS®

SI/PI Modeling & Simulation Experts

info@signaledgesolutions.com

301.887.3371

signaledgesolutions.com

CUSTOM MODELING SOLUTIONS

We specialize in advanced SI/PI and EMC modeling services. If the model needed for your simulation is not available, we can develop one.

- Interconnect Modeling: Die-to-Die and High-Speed SI/PI
- Electromagnetic Model Extractions: Touchstone (SnP)
- Custom Measurement-Based Modeling
- IBIS Model Development
- Power Delivery Network Modeling: State-Space Average VRM Models
- Passive components (capacitors, inductors, resistors, embedded components)

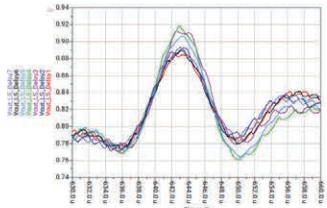

Our SSAM simulation

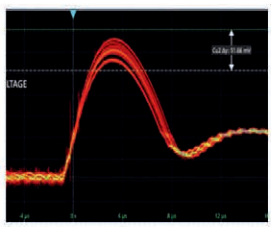

The measurement

SI/PI MODEL LIBRARY

We have a wide range of full state-space average VRM and regulator models available for products from brands including Texas Instruments, Analog Devices, Infineon, Renesas, Microchip, Qorvo, EMpower Semiconductor, Monolithic Power Systems, and more. Our SSAM models are ideal for applications including:

- Fast simulation
- Switching noise analysis
- Large and small signal analysis (time, frequency, harmonic balance)
- All PDN analysis
- EMI Simulation
- Crosstalk analysis

Models are available individually or with a subscription to our SI/PI Model Library.

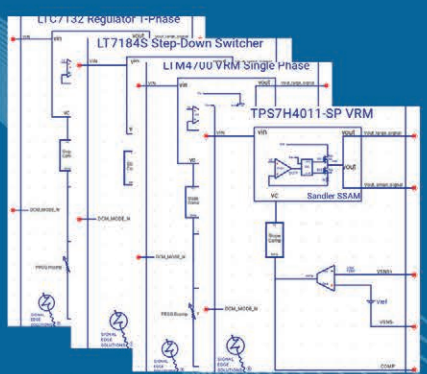

Full state-space average VRM models

SCAN HERE

for our SI/PI model solutions

www.signaledgesolutions.com/models

Keysight CERTIFIED EXPERT EDA SOFTWARE

DEPARTMENT OF DEFENSE DFARS Compliant UNITED STATES OF AMERICA

Veteran Owned Small Business VOSB CVE

The Imperfect Via: The Rough Truth Lurks Beneath the Surface

Bert Simonovich
Lamsim Enterprises, Ottawa, Canada

During DesignCon 2025, I had several side discussions about the findings presented in my DesignCon 2024 paper on dielectric anisotropy.[1] A key concern raised was the discrepancy between measured results and simulations when converting the out-of-plane dielectric constant (Dkz) to in-plane dielectric constant (Dkxy) using my heuristic method. While Isola's Tachyon 100G showed an average material anisotropy of approximately 4 to 6% across different glass styles, other researchers claimed that an anisotropy of 10 to 12% was necessary for accurate via simulation correlation.[5]

All glass-reinforced laminates are anisotropic, meaning dielectric properties vary depending on the orientation of the electric field within the structure. The Dkxy applies when the electric field is parallel to the fiberglass cloth, whereas the Dkz is when the field is perpendicular to it. Determining material anisotropy is strongly influenced by the specific test method used to extract dielectric properties and knowing the glass to resin volume ratios.

In my DesignCon 2024 paper, I defined percent anisotropy (Λ) as:

$$\Lambda \cong \left(\frac{Dkxy}{Dkz} - 1\right) \times 100 \quad (1)$$

Studies were done in an attempt to determine dielectric anisotropy using a quarter-wave resonant via stub structure.[4,5,6] This approach depends on the time delay (TD), which is influenced by the stub length and is equivalent to one-quarter of the period (T) of the resonant frequency. In theory, this approach seems like a sound method, but in reality, the as-fabricated product can skew the results.

Any quarter-wave resonant structure generates frequency nulls in the S21 insertion loss (IL), as illustrated in ***Figure 1***. The first resonant null at 13 GHz corresponds to the fundamental frequency (f_0), with additional nulls appearing at each odd-harmonic.

Given the speed of light (c), the length of the stub, and the effective dielectric constant (Dkeff) surrounding the via hole structure, the resonant frequency is predicted by:

$$f_0 \frac{c}{4 \times StubLen \times \sqrt{Dkeff}} \quad (2)$$

Adjusting Dk values within a 3D field solver to fit measured results based on as-fabricated PCB cross-section (x-section) dimensions only provides an effective anisotropy (Λeff) specific to a similar via structure utilizing the same dielectric material. It does not represent the true anisotropy of the bulk dielectric.

While material anisotropy con-

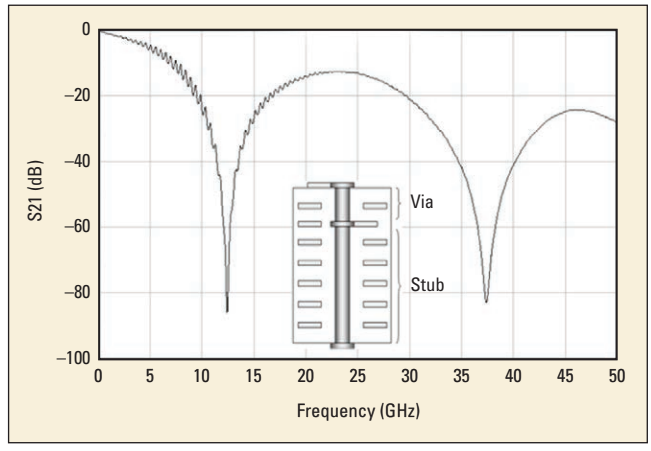

Fig. 1 **S21 IL plot showing resonant nulls due to quarter-wave stub resonances.**

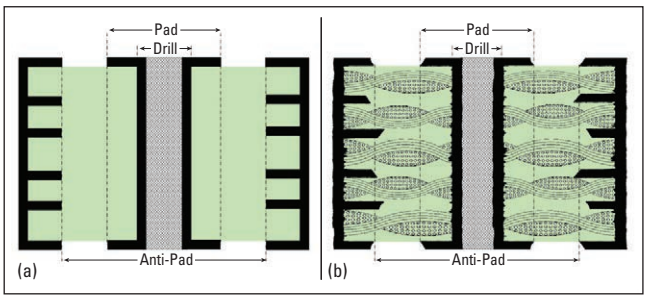

Fig. 2 **Cross-section illustration example of an ideal (a) as-designed via structure and (b) as-fabricated via structure.**

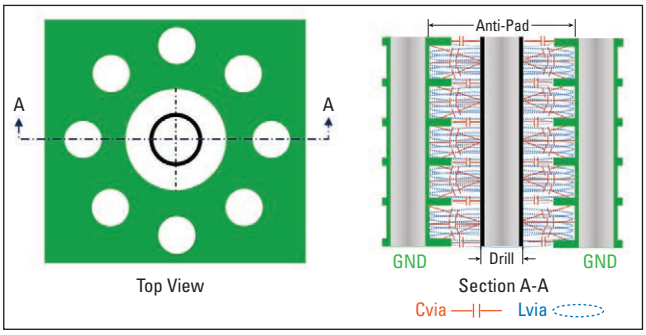

Fig. 3 **Anatomy of a single via structure surrounded by GND reference vias.**

tributes to Dkeff surrounding a via hole structure, several other factors must also be considered. One key factor is resin content of the dielectric. During the lamination process, the prepreg layers are pressed, leading to a decrease in resin volume. Since the glass volume remains unchanged, the overall Dk of the pressed laminate increases. This should be accounted for before applying my heuristic method to calculate Dkxy.

Another important consideration is drilled hole size. The actual dimensions of the via hole structure often differ from the specifications in the computer aided design (CAD) database, which can impact simulation accuracy.

Lastly, via barrel roughness plays a significant role. Just as foil roughness influences Dkeff and TD in transmission lines, via barrel roughness affects the surrounding dielectric properties as well. Increased via barrel roughness leads to higher TD and lowers the resonant frequency. Since quarter-wave stub resonance is used to determine Λeff, an increase in Dkeff and TD results in higher Λeff values.

To illustrate the impact of manufacturing tolerances on dielectric anisotropy, an ideal via structure can be compared with an actual fabricated version. An ideal via structure is depicted in ***Figure 2a***. The via barrels are perfectly smooth and antipads align symmetrically across all layers. The dielectric surrounding the via is assumed to be homogeneous. Many signal integrity (SI) engineers rely solely on the bulk Dk values provided in laminate suppliers' Dk/Df construction tables without accounting for material anisotropy. Additionally, they often assume that the final pressed dielectric thickness matches the stackup design specifications and that the specified drill size aligns with the actual drill bit dimensions.

In reality, an as-fabricated x-section reveals deviations from ideal conditions, as illustrated in ***Figure 2b***. Manufacturing tolerances result in misalignment of antipads across layers, and via barrels often exhibit rough surfaces with protruding whiskers which will affect dielectric properties. Moreover, since vias pass through a mixture of resin and fiberglass cloth, using bulk Dk values may not accurately represent material anisotropy. The Dkeff surrounding the via depends on the glass resin volume ratios of the pressed dielectric thickness and actual drill size used.

Drill Size

CAD software defines finished hole size (FHS) in the PCB layout. To add to the confusion, some CAD software also call this drill size. Fabrication notes will specify actual drill diameter tolerances, and the board shop will adjust these to meet plating hole thickness depending on the PCB class the design has to meet. The actual drill diameter is at least 2 mils larger than the FHS, but may be 3 to 4 mils larger depending on the plating requirements specified. When engineering design automation tools import the design database for SI analysis, it is the FHS that gets imported. This is a common trap SI engineers fall into when modeling vias; using the FHS instead of actual drill size will underestimate via capacitance and thus Dkeff.

Via Capacitance

In a coaxial structure, electromagnetic (EM) fields are fully contained within a grounded shield surrounding a central conductor, separated by a dielectric material. The electric field (E-field) dictates capacitance, while the magnetic H-field defines inductance, leading to transverse electromagnetic (TEM) wave propagation.

Although a via structure resembles coaxial design, it lacks a continuous shield. Instead, ground (GND) vias and antipad clearance holes confine EM fields within the dielectric cavity between reference planes, resulting in quasi-TEM propagation. As illustrated in ***Figure 3***, the anatomy of the via structure includes localized EM fields, but does not fully contain them.

In Figure 3, Section A-A, via capacitance (Cvia) is influenced by drill diameter (Drillϕ), antipad size (Antipadϕ), and nearby GND vias. Increasing Drillϕ or decreasing Antipadϕ raises via capacitance by reducing the space between the via barrel and antipad. The approximation for via capacitance is given as:

$$\text{Cvia} \cong \frac{2\pi\varepsilon_0}{\ln\left(\frac{\text{Antipad}\phi}{\text{Drill}\phi}\right)} \times \text{Dkxy} \qquad (3)$$

where: ε_0 is the permittivity of free space, Dkxy is the in-plane dielectric constant, and Antipadϕ and Drillϕ represent the antipad and drill diameters, respectively.

Via Roughness

Via barrel roughness mainly results from copper plating wicking into voids created by drill bit crazing of the glass reinforcement weave, often due to a dull drill bit. Though typically overlooked in via modeling, it affects SI correlation. ***Figure 4*** illustrates copper plating wicking into glass bundles.

Conductor roughness increases via capacitance and Dkeff, similar to how copper surface roughness raises self-capacitance (C_{11}) in transmission lines.[2] Wicking extends beyond the drill diameter, concentrating electric field

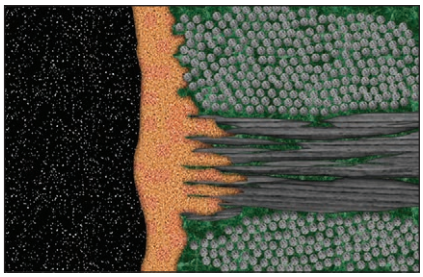

Fig. 4 **Copper plating wicking into glass crazing caused by drill bit.**

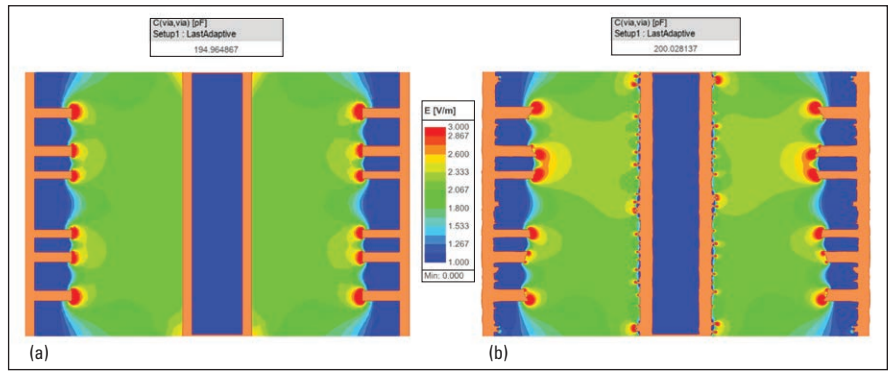

Fig. 5 **Electric field strength color map and capacitance of (a) smooth vias and (b) rough vias.** ***Source:*** **Juliano Mologni, Ansys.**

strength and further increasing capacitance.

HFSS simulations in ***Figure 5*** validate this effect. ***Figures 5a*** and ***5b*** show E-field strength in smooth and rough vias, respectively. The E-field is mostly contained within the antipad opening, like a coaxial geometry. Increased E-field strength along the roughness profile in Figure 5b leads to a 2.6% capacitance rise.

Achieving model correlation is difficult due to the randomness of wicking and its interaction with glass and resin. A single x-section only captures one slice of the 360° hole, where wicking varies around the circumference. ***Figure 6*** provides a microscopic top-down view of a plated-through hole showing copper wicking into the faintly visible glass weave.

Dkeff Compensation Due to Conductor Roughness

As shown in Figure 6, the measured inner ring diameter of 14.4 mils represents the FHS. The middle ring drill diameter is 18.43 mils. By inspection, the outer ring diameter of 18.80 mils represents the drill diameter plus the average roughness.

Heuristically, additional capacitance and Dkeff correction due to roughness can be estimated for the via example in Figure 6. If the ratio of $Dkeff_{rough}$ to $Dkeff_{smooth}$ is defined as:

$$\frac{Dkeff_{rough}}{Dkeff_{smooth}} = \frac{Cvia_{rough}}{Cvia_{smooth}} \quad (4)$$

If $Drill\phi smooth$ = 18.43 mils; $Drill\phi rough$ = 18.80 mils, and assuming a typical $Antipad\phi$ of 40 mils, then by combining Equation 3 with Equation 4, $Dkeff_{rough}$ can be expressed by Equation 5. When plugging in the numbers, one sees $Dkeff_{smooth}$ increases by 2.6% as compared to $Dkeff_{smooth}$.

$$Dkeff_{rough} \cong \frac{\ln\left(\frac{Antipad\phi}{Drill\phi_{smooth}}\right)}{\ln\left(\frac{Antipad\phi}{Drill\phi_{rough}}\right)} \times Dkeff_{smooth} \cong$$

$$\frac{\ln\left(\frac{40}{18.43}\right)}{\ln\left(\frac{40}{18.80}\right)} \times Dkeff_{smooth} \cong 1.026 \times Dkeff_{smooth} \quad (5)$$

Roughness Effect on TD

Extracting Dkeff from the first quarter-wave resonant null in an S21 IL plot follows Equation 6, which assumes Dkeff is purely capacitance-driven.

$$Dkeff = \left(\frac{c}{4 \times StubLen \times f_0}\right)^2 \quad (6)$$

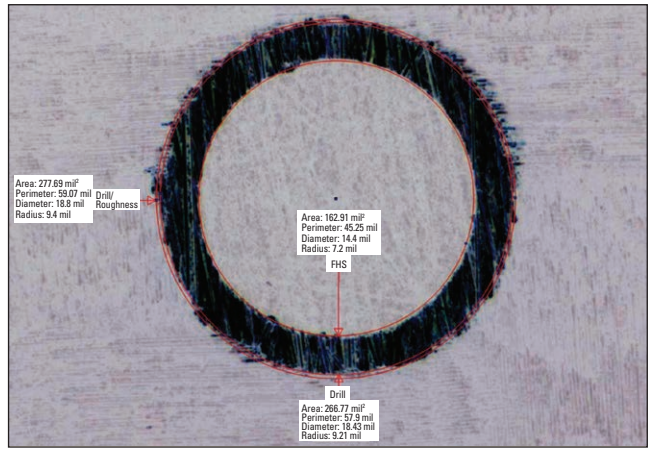

Fig. 6 **Microscopic top-down view of a slice of an actual plated through hole showing copper wicking into the glass weave, faintly visible running horizontally and vertically.**

However, for time-variant EM fields, inductance also affects TD. Via barrel roughness impacts self-inductance (L_{11}) similarly to copper surface roughness in transmission lines. In my previous paper,[3] Dkeff for time-variant fields is expressed as:

$$Dkeff = c^2(L_{11}C_{11}) \quad (7)$$

Since L_{11} increases Dkeff proportionally, failing to use a causal metal roughness model, such as Bracken's model,[7] can lead to misinterpretation of extracted values and anisotropy effects.

To validate this, I collaborated with Juliano Mologni from Ansys to introduce roughness into a via quarter-wave stub structure and quantify its effect using Equation 6.

Experiment Setup

A six-layer via stub structure, modeled in HFSS, featured a 10 mil drill, a 50 mil antipad, and six 10 mil stitching vias surrounding main via at 60 mil diameter. Microstrip traces on the top layer extended the stub length to 150 mils. A value of 3.97 was used for Dk.

Applying the Huray roughness model to all vias, parameterized from 0 to 2 µm in 0.1 µm increments (0 to 33 µm Rz equivalent roughness), and an HHSR of 4.9 based on the Simonovich-Cannonball roughness model,[8] we simulated 20 Huray NR values.

Results

Figure 7 plots S21 IL resonant nulls across all NR

OPHIR RF

High Power RF Systems

CLASS A LINEAR • 5-YEAR WARRANTY

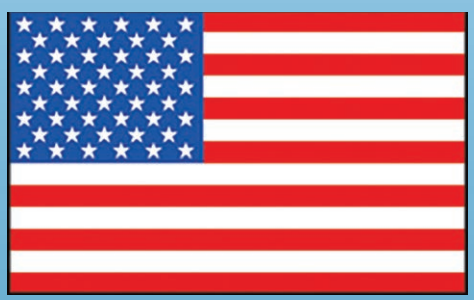

DESIGNED AND MANUFACTURED IN THE USA
SINCE 1992!!

0.01-225MHz 5kW

80-1000MHz 2kW

1.0-6.0GHz, 500W

6.0-18GHz, 250W

www.OphirRF.com
(310) 306-5556

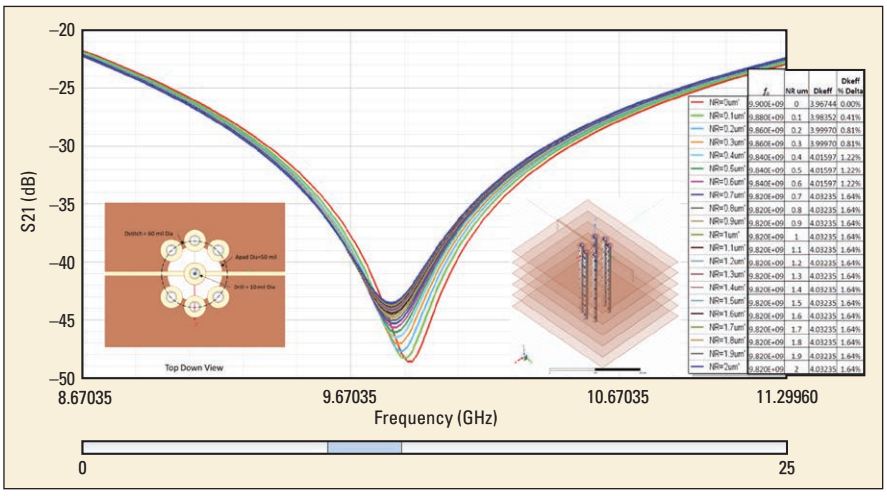

Fig. 7 **S21 IL showing the quarter-wave stub resonant nulls for 0-2 µm Huray nodule radius (NR) roughness parameters.** ***Source:*** **Juliano Mologni, Ansys.**

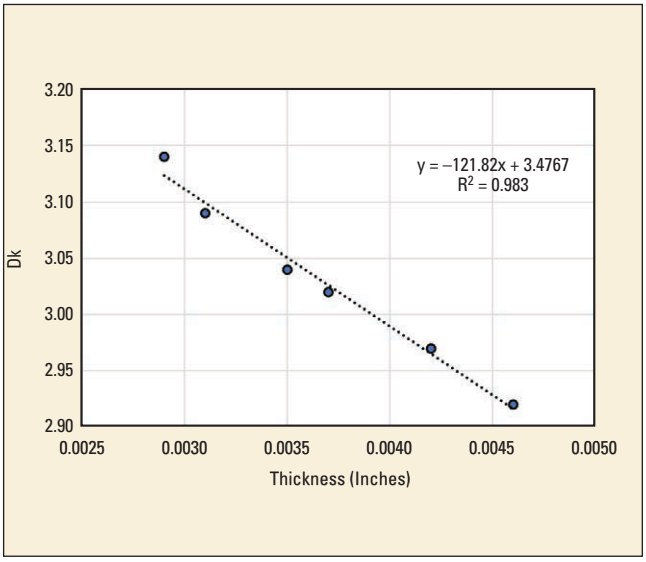

Fig. 8 **Linear fit of Dk vs. prepreg thickness for published values of Tachyon 100G 1078 glass style.[9]**

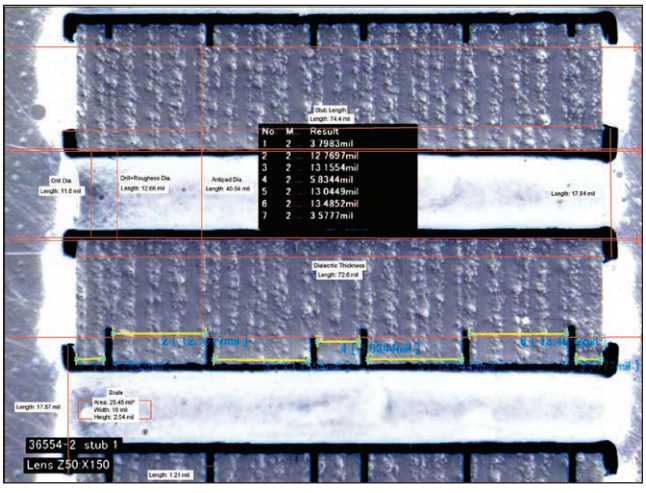

Fig. 9 **Negative image from original cross-section photo of Stub_1 showing pressed dielectric thickness measurements in yellow. Additional measurements for this case study are shown in red.** ***Source:*** **Scott McMorrow.[5]**

values. Each frequency was measured and converted to Dkeff using Equation 6. Dkeff increased from 3.98 for NR = 0 µm to 4.03 for NR = 0.7 µm and did not change for NR increase past that value. Since Bracken corrects only the imaginary impedance component of rough metal, not capacitance, the observed Dkeff increase is solely attributed to L_{11} validating the expectation.

Dkeff Due to Pressed Thickness

Dkeff of individual cores and prepreg layers varies based on final pressed thickness. Micro-section measurements of the tested board are important for accurate simulation correlation. Since stackups are designed using published Dk/Df values, they differ from actual pressed thickness post-fabrication.

During the pressing process, heat and pressure reduce resin content, thereby increasing Dk since published values are based on pre-pressed resin content. ***Figure 8*** illustrates the relationship between bulk Dk and thickness for Tachyon 100G 1078 glass. With thickness changes altering resin volume, a linear fit equation is used to adjust Dk accordingly before converting from Dkz to Dkxy.

Previous Study

In a DesignCon 2015 paper,[4] anisotropic dielectric models were proposed to align via simulations with measurements. Since final results were not available by publication, I reached out to Scott McMorrow, who graciously shared as-fabricated and simulation results.[5]

Two via stub test structures, Stub_1 and Stub_4, were cross-sectioned for analysis. Using measured via lengths and Dk values from the as-designed stackup for the simulation, there was a difference in stub resonance frequency of 1.054% for Stub_1 and 1.057% for Stub_4. These corresponded to as-fabricated effective anisotropy (Λeff) values of 11% and 12%, respectively, based on empirical data.

Figure 9 shows a negative image of the original x-section photo of Stub_1. The dielectric thickness measurements, shown in yellow and summarized in the black box at the center of the via, are from the original picture. Additionally, I performed further measurements, indicated by the red dimension lines.

Dkeff Due to Pressed Thickness

Compared to the as-designed stackup, the as-fabricated dielectric was thinner, increasing the average bulk Dkz from 3.00 to 3.07 (2.8%). Using 3.07 as the baseline and applying my heuristic conversion method, the average bulk Dkxy is 3.22, corresponding to 4.8% anisotropy.

Dkeff Compensation Due to Conductor Roughness

From the x-section measurements in Figure 9: Drillϕsmooth = 11.80 mils, Drillϕrough = 12.66 mils, and Antipadϕ = 40.04 mils. With $Dkeff_{smooth}$ = 3.22, Equation 5 yields the effective Dkxy due to roughness (Dkef-

ARTECH HOUSE
BOSTON | LONDON

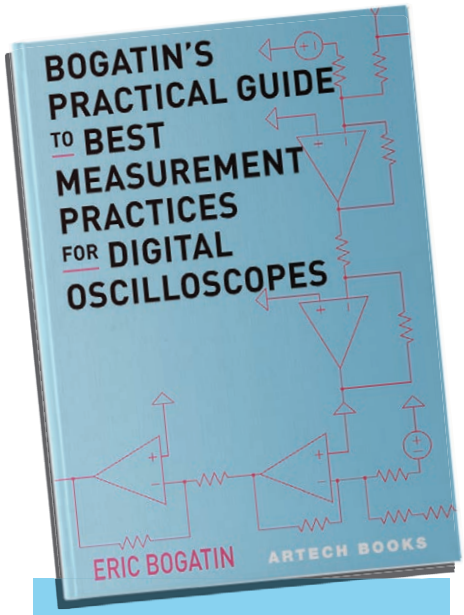

USE THE QR CODES BELOW TO LEARN MORE!

US, CANADA
SOUTH AMERICA, AUSTRALIA
US.ARTECHHOUSE.COM

UK, EUROPE
EMEA, ASIA
UK.ARTECHHOUSE.COM

MICROWAVE, DIGITAL OSCILLOSCOPES

BOGATIN'S PRACTICAL GUIDE TO BEST MEASUREMENT PRACTICES FOR DIGITAL OSCILLOSCOPES

Eric Bogatin

Bogatin's Practical Guide to Best Measurement Practices for Digital Oscilloscopes is a hands-on learning resource designed to teach you how to use digital oscilloscopes the right way. It focuses on real-world techniques and best practices that every engineer, technician, and student should know. This book includes 15 easy-to-follow chapters, each paired with video tutorials created by Eric Bogatin, demonstrating how to make accurate and reliable measurements using modern digital scopes.

YOU'LL LEARN HOW TO:

- Avoid common mistakes that lead to bad measurements - like incorrect probe use, poor cable setups, and scope bandwidth issues
- Use key oscilloscope features like FFTs and math functions effectively
- Understand what you're really measuring and why it matters

The book starts with the basics - like signal measurements, 10x probes, and transmission lines - and builds up to more advanced topics like debugging, working with differential signals, measuring noise and emissions, and even connecting your measurements to circuit simulations. To help you practice, the book comes with a free virtual oscilloscope and waveform emulator. This means you can run simulations on your computer and try out what you're learning, even without physical lab equipment. While the main focus is on 10x probes and scopes up to 200 MHz, there's also guidance for higher-frequency applications over 200 MHz.

Whether you're a student just starting out, a new engineer, or a technician looking to improve your measurement skills, this guide will help you build a strong foundation in oscilloscope use.

ISBN: 978-1-63081-963-7 · **E-Book:** 978-1-63081-964-4 · 2025
Hardcover · 430pp

Price: $124 /£94 · **E-Book Price:** $102 /£72

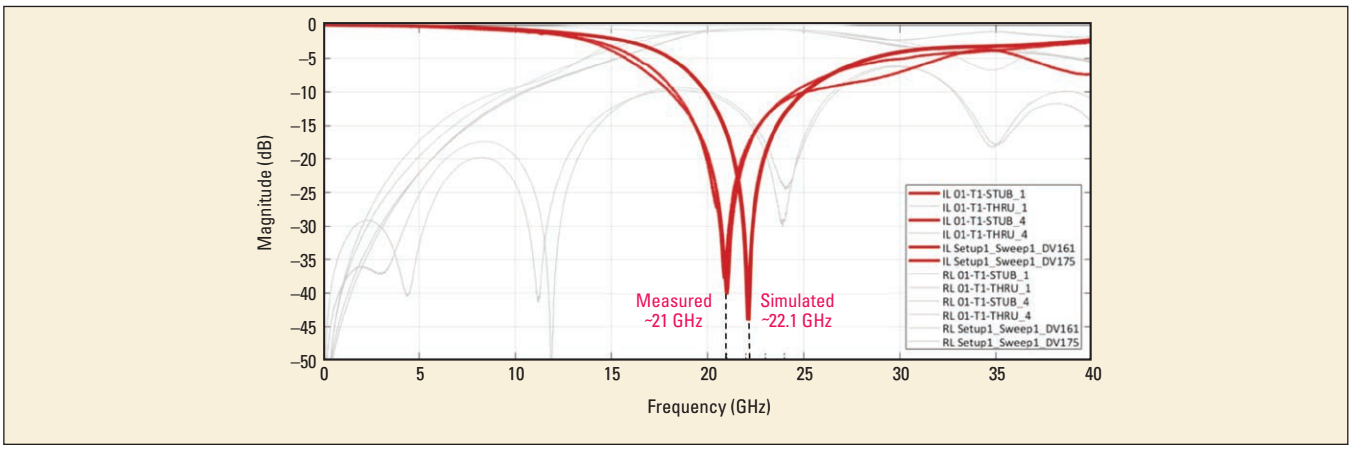

▲ Fig. 10 **Measured vs. simulated IL results.[5]**

fxy_{rough}) as:

$$Dkeffxy_{rough} \cong \frac{\ln\left(\frac{Antipad\phi}{Drill\phi_{smooth}}\right)}{\ln\left(\frac{Antipad\phi}{Drill\phi_{rough}}\right)} \times Dkeff_{smooth} \cong \frac{\ln\left(\frac{40.04}{11.80}\right)}{\ln\left(\frac{40.04}{12.66}\right)} \times 3.22 \cong 1.06 \times 3.22 \cong 3.42 \quad (8)$$

which increases Dkeffxy by 6.2%.

Excluding Dkeff correction for roughness-induced inductance, the modeled anisotropy of as-fabricated via Stub_1 is:

$$\Lambda eff_{rough} \cong \frac{Dkeffxy_{rough}}{Dkeffz_{pressed}} - 1 \cong \frac{3.42}{3.07} - 1 \cong 0.114$$

or 11.4% (9)

Taking half the difference between Drillϕrough and Drillϕsmooth, the surface roughness of the via barrel is 0.43 mils or approximately 10.9 µm. Assuming a similar Dkeff contribution as in Reference 5, the polynomial fit shown in Figure 9 suggests this roughness adds 1.36%.

This brings the total effective anisotropy to:

$$\Lambda eff_{total} \cong 11.4\% + 1.4\% \cong 12.8\% \quad (10)$$

Resulting in a final Dkeff of:

$$Dkeff_{final} \cong 3.22 \times 1.12 \cong 3.61 \quad (11)$$

Comparison of Simulated vs. Measured Results

Figure 10 shows simulated vs. measured IL. [5] Using measured stub lengths and Dk values in the HFSS model, the simulated quarter-wave resonant frequency is ~22.1 GHz, while the measured frequency was ~21 GHz.

With a measured stub length of 74.4 mils:

$$Dkeff_{meas} = \left(\frac{c}{4 \times fo \times StubLen}\right)^2 = \left(\frac{1.18E10}{4 \times 21.0E9 \times 74.4E-3}\right)^2 \cong 3.6 \quad (12)$$

$$Dkeff_{sim} = \left(\frac{c}{4 \times fo \times StubLen}\right)^2 = \left(\frac{1.18E10}{4 \times 22.1E9 \times 74.4E-3}\right)^2 \cong 3.2 \quad (13)$$

$$\Lambda eff = \frac{3.6}{3.2} - 1 = 12.5 \quad (14)$$

Summary and Conclusion

- Calculated Λeff from Equation 10 = 12.8%
- Measured Λeff from Equation 14 = 12.5%
- Calculated Dkeff from Equation 11 = 3.61
- Measured Dkeff from Equation 12 = 3.60

These results confirm excellent correlation and validate the hypothesis.

The extended case study revealed an effective anisotropy of 12.5%, compared to 4.8% from bulk Dk values. Adjusting Dkeff for pressed thickness and via roughness added ~8%, supporting the need to account for these factors in simulations.■

ACKNOWLEDGEMENTS

Thanks to Juliano Mologni and Scott McMorrow for simulation support and measured data.

REFERENCES:

1. B. Simonovich, "A Heuristic Approach to Assess Anisotropic Properties of Glass-reinforced PCB Substrates," *DesignCon 2024 proceedings*, Santa Clara, Calif.
2. B. Simonovich, "A Practical Method to Model Effective Permittivity and Phase Delay Due to Conductor Surface Roughness," *DesignCon 2017 proceedings*, Santa Clara, Calif.
3. B. Simonovich, "A Tale of Two Data Sheets and How Foil Roughness Affects Dk," White Paper, *Lamsim Enterprises Inc.*, Issue 03, March 7, 2022.
4. S. McMorrow et al., "Anisotropic Design Considerations for 28 Gbps Via to Stripline Transitions," *DesignCon 2015 proceedings*, Santa Clara, Calif.
5. S. McMorrow, "Anisotropy Test Vehicle Tachyon Results," *Samtec presentation*, 2022.
6. A. Neves et al., "Free Signal Integrity? How Understanding Anisotropic Materials & Tolerances Could Increase Performance at 112/224Gbps & Beyond," *DesignCon 2025 proceedings*, Santa Clara, Calif.
7. J. E. Bracken, "A Causal Huray Model for Surface Roughness.", *DesignCon 2012 proceedings*, Santa Clara, CA.
8. B. Simonovich, "Practical Method for Modeling Conductor Surface Roughness Using The Cannonball Stack Principle," *White Paper, Issue 2*, April 25, 2015.
9. Isola Tachyon 100G Dk/Df construction table.
10. B. Simonovich, "From Smooth to Imperfect Vias: The Rough Truth Impacting Simulation Model Accuracy," *Signal Integrity Journal*, June 3, 2015.

APE Asia Photonics Expo

Held in: SINGAPORE | Brought to you by: informa markets | Supported by: SINGAPORE EXHIBITION & CONVENTION BUREAU

APE 2026

ASIA PHOTONICS EXPO

Connecting Photonics Opportunities and Innovation in Asia

4-6 FEBRUARY 2026

SANDS EXPO AND CONVENTION CNETRE L1, SINGAPORE

Exhibit Profile:

- Optical Communication
- Optics
- Laser
- Infrared
- Sensing
- Display
- Quantum
- Semiconductor

Contact Us

Email: visitors-ape@informa.com

Website: www.asiaphotonicsexpo.com

Follow Us: #APE2026

Book a stand

Sign up to visit

AD INDEX

Anritsu 25
APE 2026 49
Artech House 47
Coilcraft 15
COMSOL, Inc. ..... COV 4
DesignCon 2026 ..... COV 3
emv 2026 8
HYPERLABS INC. 9
iNRCORE 7
Maury Microwave Corporation 33
Ophir RF 45
Picotest 35
Quantifi Photonics 39
Rogers Corporation 29
Rohde & Schwarz USA, Inc. 3
Samtec USA ..... COV 2
Siglent Technologies 19
Signal Edge Solutions 41
Teledyne Lecroy 5

USA

Eastern and Central Time Zones
Carl Sheffres
Group Director
(New England, New York, Eastern Canada)
Tel: (781) 619-1949
csheffres@horizonhouse.com

Michael Hallman
Associate Publisher
(NJ, Mid-Atlantic, Southeast, Midwest, TX)
Tel: (301) 371-8830
Cell: (781) 363-0338
mhallman@horizonhouse.com

Pacific and Mountain Time Zones
Brian Landy
Western Reg. Sales Mgr.
(CA, AZ, OR, WA, ID, NV, UT, NM, CO, WY, MT, ND, SD, NE & Western Canada)
Tel: (831) 426-4143
Cell: (831) 713-9085
blandy@horizonhouse.com

International

International Sales
Michael O'Kane
Tel: +44 (0) 1875 825 700
Cell: +44 (0) 7961 038 245
mokane@impressmedia.com

Germany, Austria, and Switzerland (German-speaking)
Victoria and Norbert Hufmann
victoria@hufmann.info
norbert@hufmann.info

SIJ

685 Canton Street
Norwood, MA 02062
www.signalintegrityjournal.com

Tel: (800) 966-8526
(781) 769-9750
Fax: (781) 769-5037

DesignCon® 2026

WHERE THE CHIP MEETS THE BOARD

Host Sponsor: **Amphenol**

February 24 – 26, 2026
Santa Clara Convention Center
Santa Clara, CA

The Premier High-Speed Communications and Systems Design Conference and Expo

- 170 Exhibitors
- 15 Track Conference
- Daily Keynotes and Open-to-All Education
- Networking Events

Use code **SIJ26** for **15% off** any education package or a complimentary expo pass

DesignCon.com

informa markets

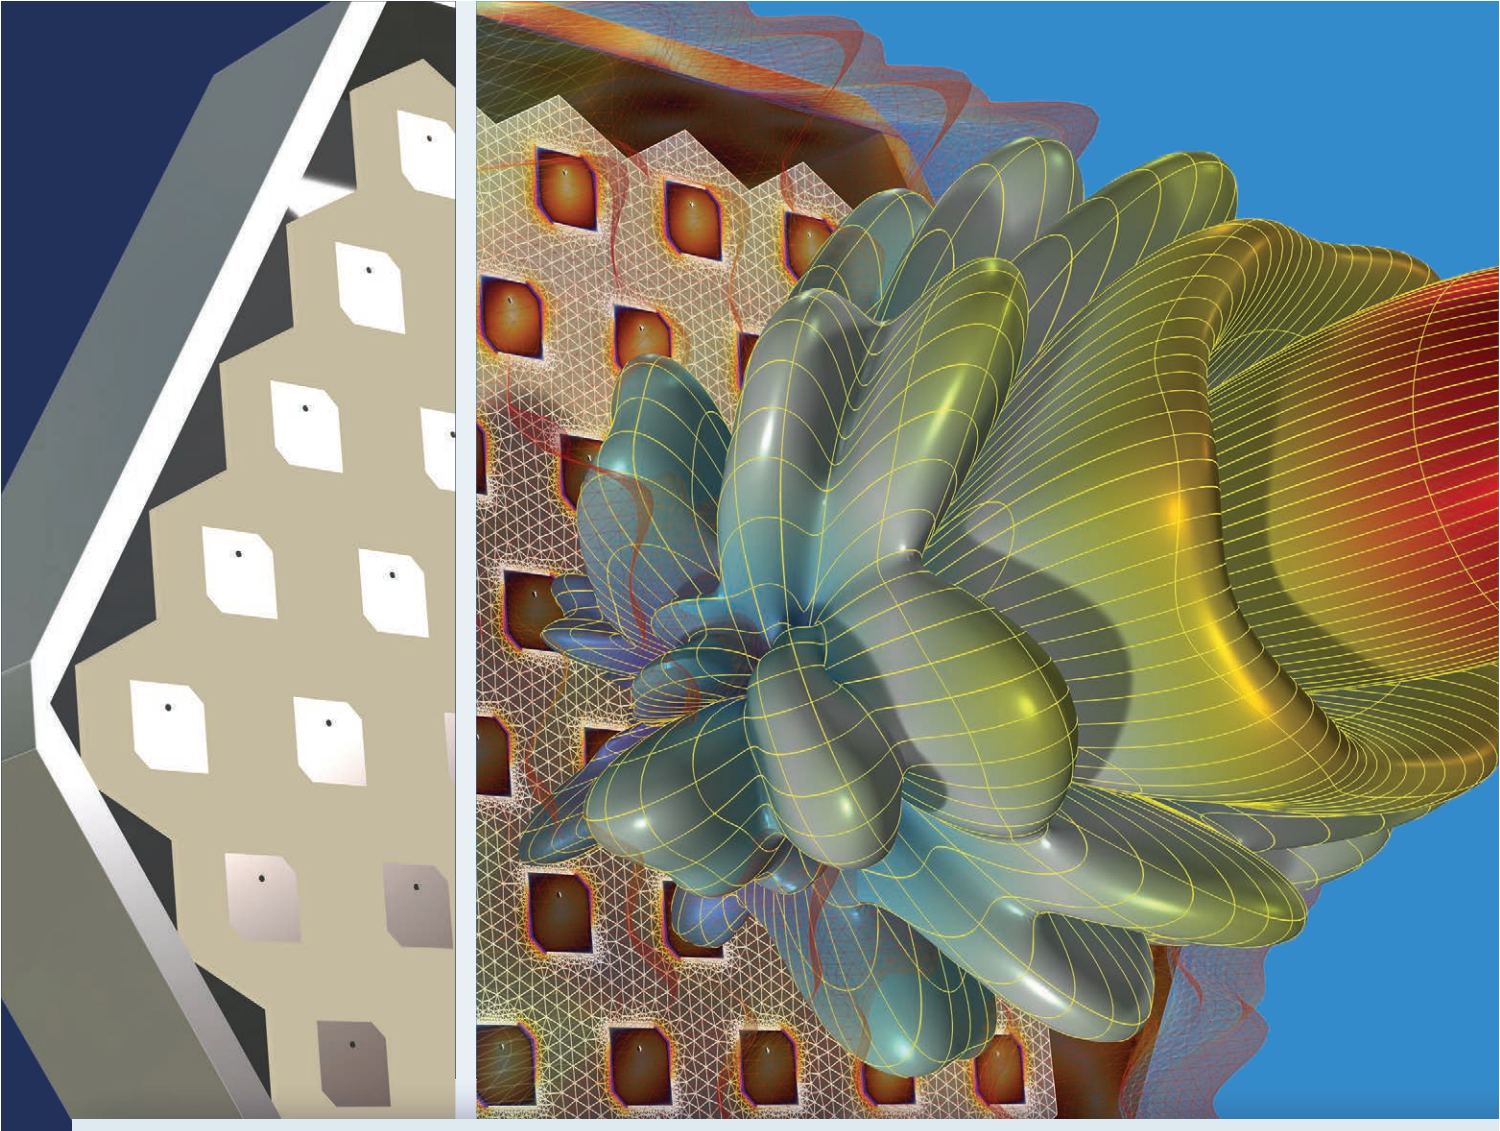

Take the Lead in RF Design

with COMSOL Multiphysics®

Multiphysics simulation is expanding the scope of RF analysis to higher frequencies and data rates. Accurate models of microwave, mmWave, and photonic designs are obtained by accounting for coupled physics effects, material property variation, and geometry deformation. Ultimately, this helps you more quickly see how a design will perform in the real world.

» *comsol.com/feature/rf-innovation*

COMSOL