Vol. 8 • Issue 1

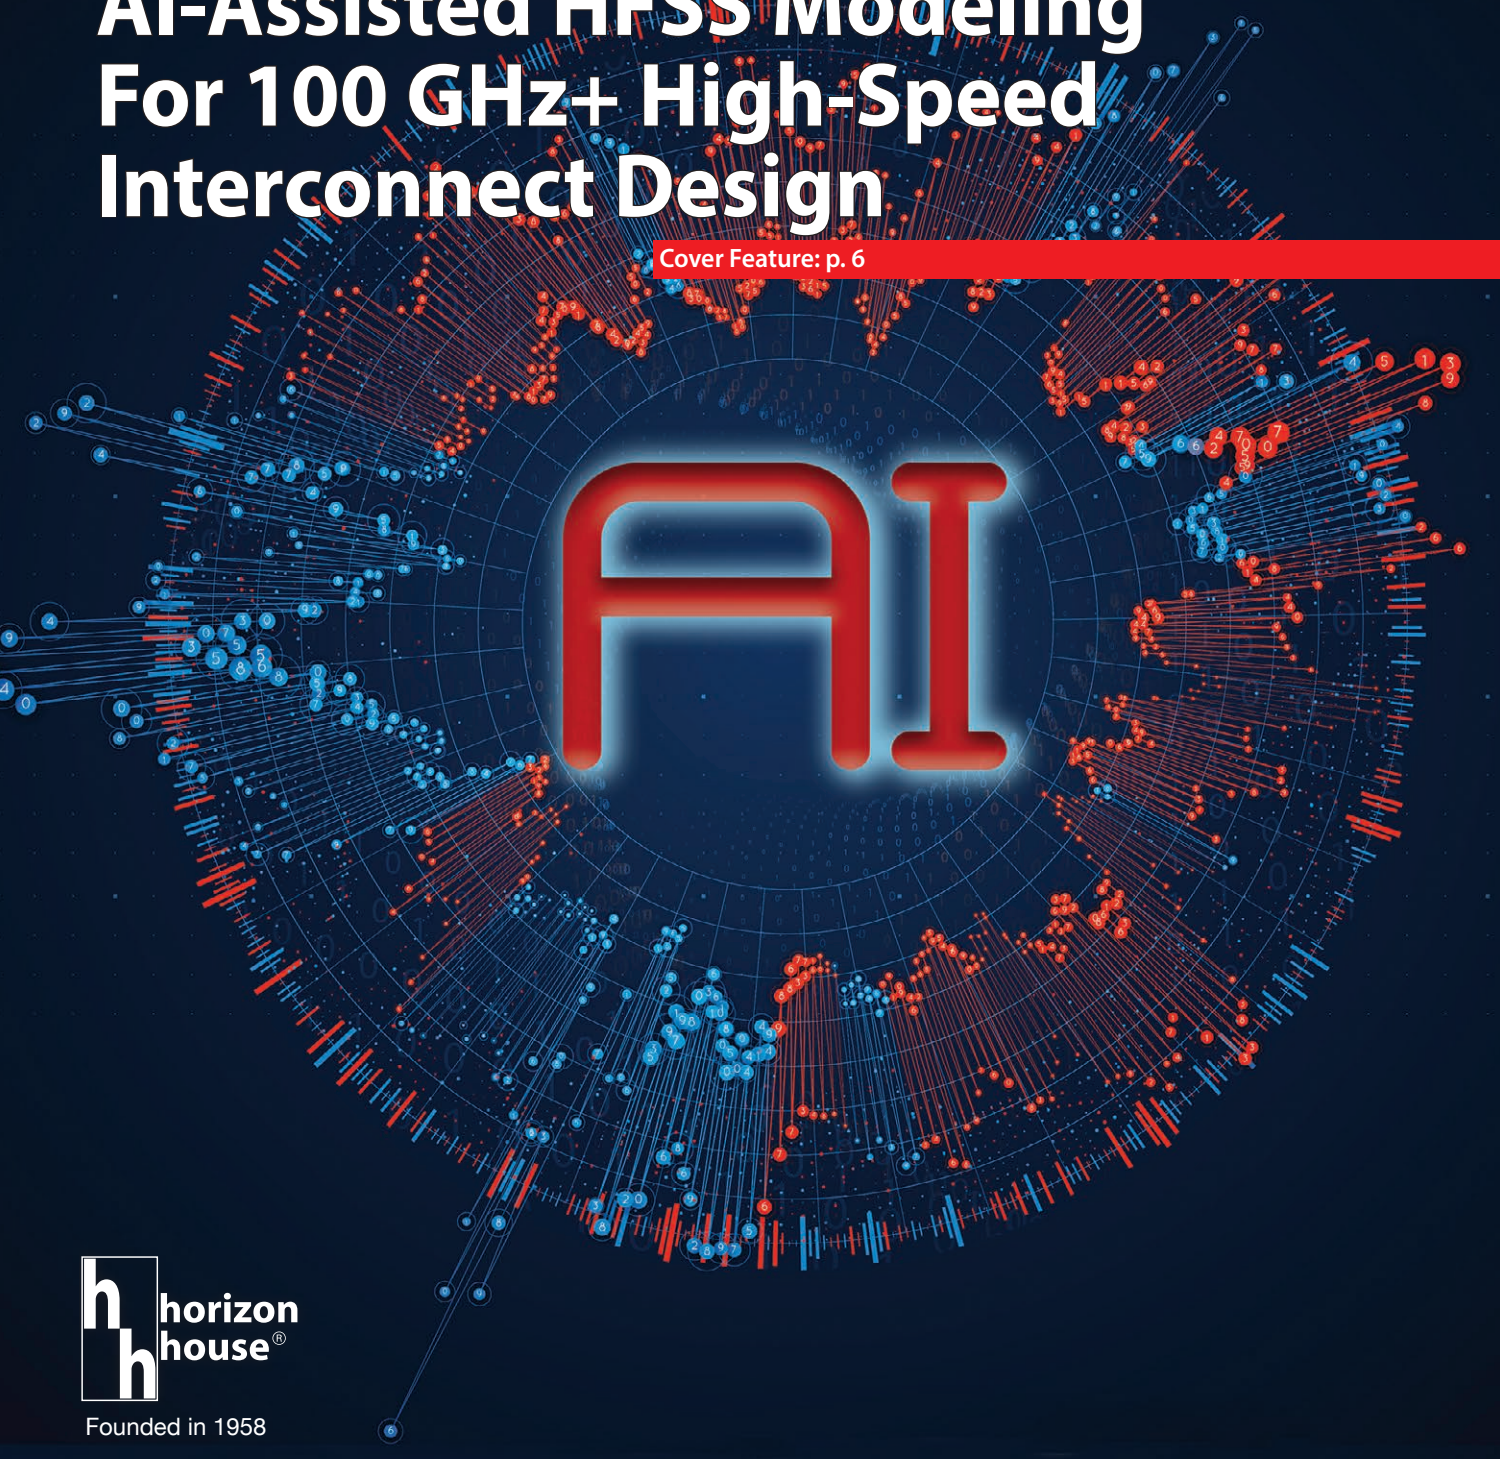
2026

Signal Integrity Journal™

Signal Integrity • Power Integrity • EMC/EMI

AI-Assisted HFSS Modeling For 100 GHz+ High-Speed Interconnect Design

Cover Feature: p. 6



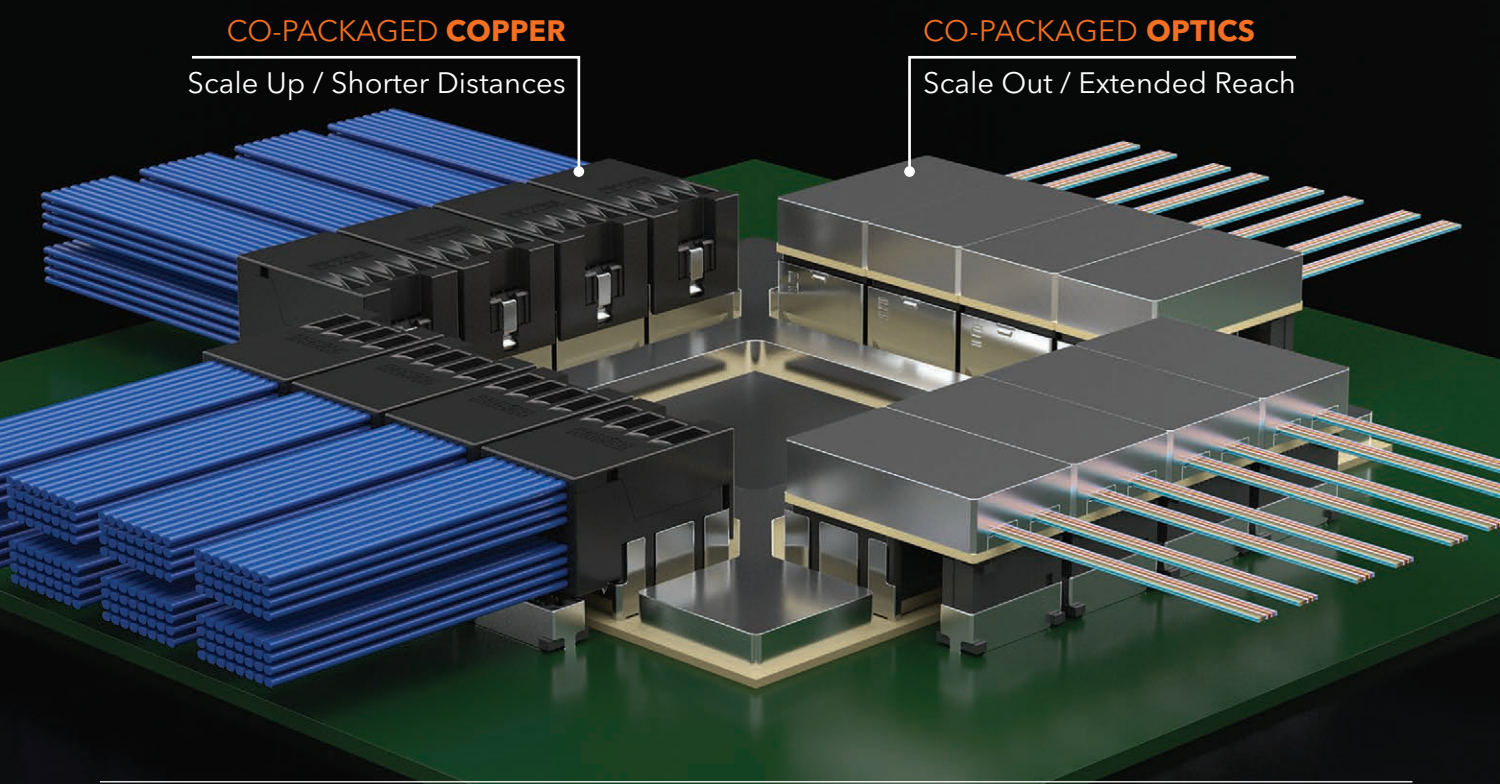
horizon house®

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.



KEY FEATURES & BENEFITS

- 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 per square 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

TECHNOLOGY ROADMAP


Si-Fly® HD
High-Density
Backplane Cable System


Si-Fly® HD
Mezzanine Arrays

samtec
www.samtec.com

For additional details, visit **samtec.com/si-fly-hd** or contact **HDR@samtec.com**.



R&S®ZNB3000

FAST FORWARD TO RESULTS

Best-in-class dynamic range

ROHDE&SCHWARZ

Make ideas real

TABLE OF CONTENTS

Cover Feature

6 AI-Assisted HFSS Modeling For 100 GHz+ High-Speed Interconnect Design

Jason Yoho, Ameya Ramadurgakar, and Madrone Coopwood, HYPERLABS

Technical Features

12 Why Power Architectures Constrain New Space AI Missions

Ken Coffman, Vicor

16 Machine Learning Models for SI/PI Analysis with Meshed Planes

Jiyue Zhu, Regina Thahir, Xiaoyan Xiong, Gang Kang, and Jian Liu, Cadence

22 Stop Breaking the Loop: Modern Closed–Loop Stability in ADS

Benjamin Dannan, Steve Sandler, and Heidi Barnes

28 Evolution of the PDN Design Process: Explaining the Second Paradigm Shift

Istvan Novak, Samtec

34 When Frequency Shapes Time: How S-Parameter Properties Shape Simulated TDR Behavior

Tim Wang Lee, Keysight Technologies

40 PAM2 vs. PAM4 Signaling: Simply Explained

Bert Simonovich, LAMSIM Enterprises Inc.

Follow Us @SIJupdates

Become a Fan at facebook.com/signalintegrityjournal/

Join the Signal & Power Integrity Community

Staff

Group Director: Carl Sheffres
Technical Editor: Eric Bogatin
Media Director: Pat Hindle
Brand & Content Director: Jennifer DiMarco
Editorial & Media Specialist: Kelley Roche
Associate Editor: Kaitlyn Joyner
Director of Production & Distribution: Edward Kiessling
Art Director: Janice Levenson
Graphic Designer: Ann Pierce

Editorial Advisory Board

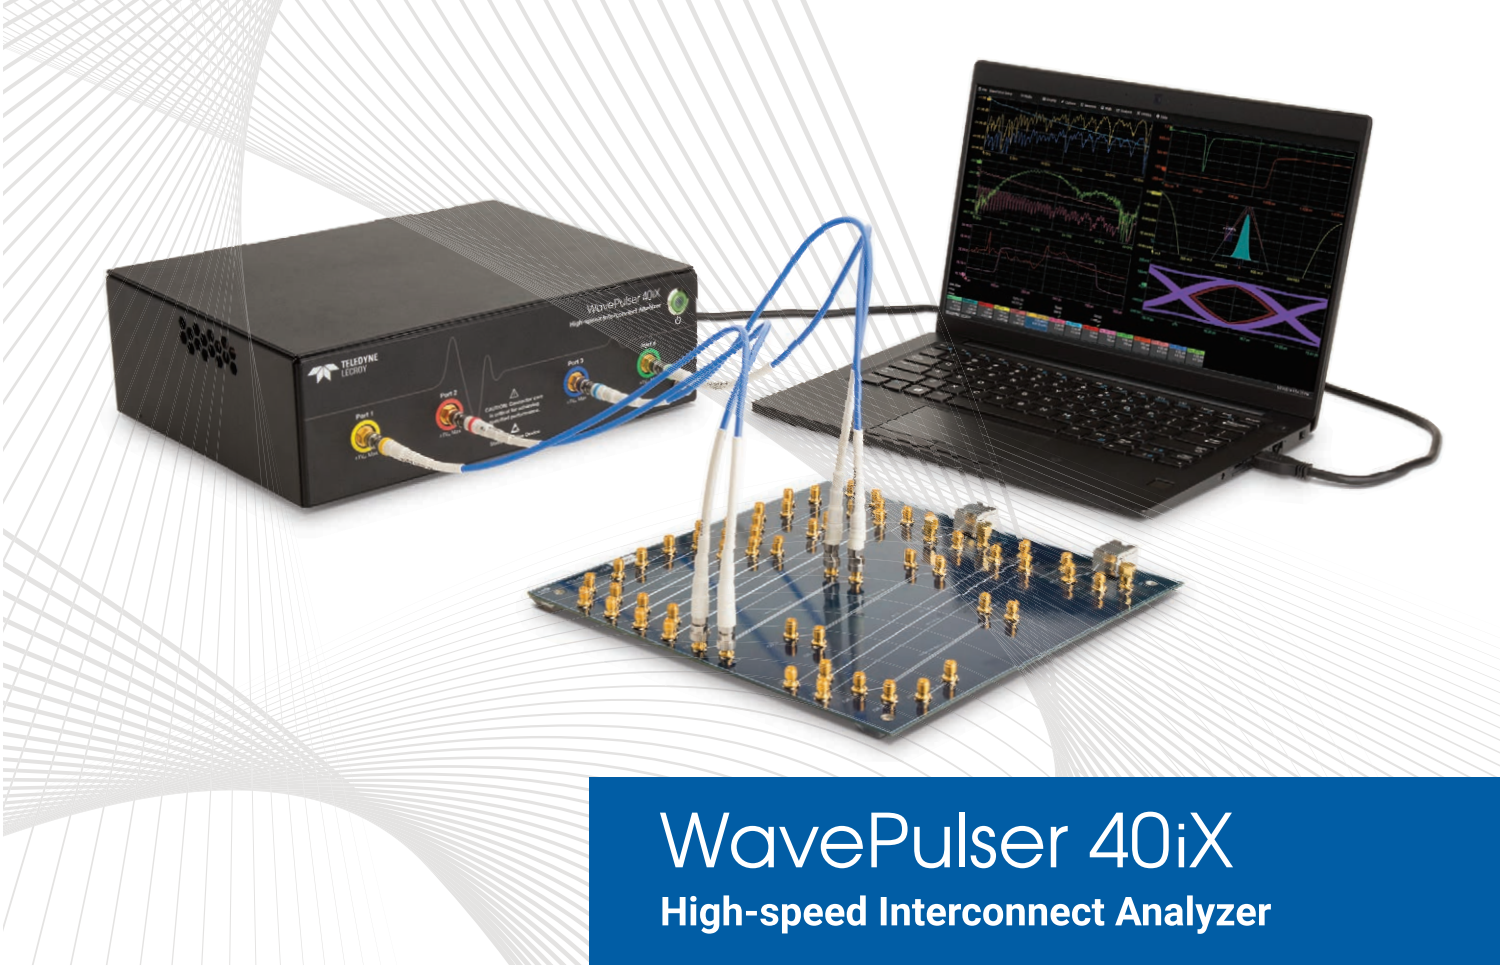
Rula Bakleh
Eric Bogatin
Karen Burnham
Benjamin Dannan
Benoit Derat
Jay Diepenbrock
Vladimir Dmitriev-Zdorov
Jason Ellison
Barry Katz
HeeSoo Lee
Cathy Ye Liu
Alfred Neves
Istvan Novak
Prashant Pappu
C. J. Reddy
Steve Sandler
Yuriy Shlepnev
Bert Simonovich
Larry Smith

CORPORATE STAFF

CEO: William M. Bazzy
President: Ivar Bazzy
Vice President: Jared Bazzy

TELEDYNE | LECROY

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)
- Advanced de-embedding to IEEE 370-2020 Standard
- Emulation and simulation of high-speed serial data channels
- No calibration required
- Costs a fraction of the price of a VNA

Learn more:

teledynelecroy.com/wavepulser

AI-Assisted HFSS Modeling For 100 GHz+ High-Speed Interconnect Design

Jason Yoho, Ameya Ramadurgakar, and Madrone Coopwood
HYPERLABS, Louisville, Colo.

As serial data rates continue to increase, signal integrity challenges increasingly resemble traditional mmWave design problems. At frequencies approaching and exceeding 100 GHz, geometry, fabrication tolerance, and electromagnetic detail govern performance. Launch structures, connectors, and transitions often become the dominant contributors to reflection, loss, and channel margin degradation.

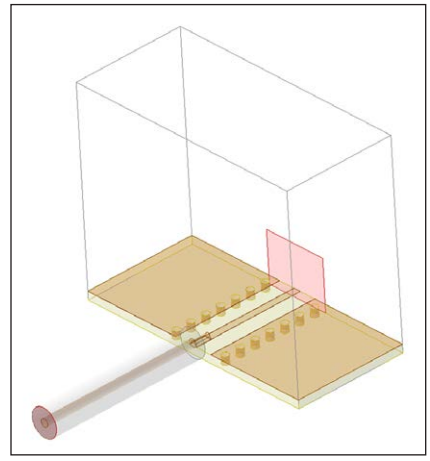
To this point, full wave electromagnetic simulations using commercially available solvers such as ANSYS's HFSS are no longer optional refinements applied late in the design cycle. They become central design tools used to understand, predict, and control high-speed interconnect behavior. However, engineers quickly discover that productivity can drop sharply. Small geometric changes can produce significant performance swings, many candidate geometries are not physically realizable, and traditional parametric sweeps can become expensive.

At the same time, artificial intelligence (AI) techniques have begun to be applied in electronic design workflows. For signal integrity applications above 100 GHz, AI does not replace electromagnetic expertise or engineering judgment. When applied carefully and within strict constraints, however, AI can serve as a practical assistant that potentially reduces iteration count, improves convergence reliability, and shortens the path to a manufacturable interconnect solution.

This technical feature describes how AI-assisted methods can be applied within a full wave electromagnetic solver workflow for high-speed interconnect design, with an emphasis on practical connector and launch structures relevant to signal integrity applications.

Where AI Fits in Optimizing High-Speed Interconnects

For high-speed interconnects operating in the mmWave frequencies, unconstrained or freeform AI-generated geometry is neither credible nor useful. Instead, AI provides value when it operates within boundaries defined by



▲ Fig. 1 Simulation test case structure.

known interconnect topologies and fabrication limits.

In practice, AI assistance is most effective in three areas:

- Proposing geometry variations within a fixed, well-understood topology
- Exploring tightly constrained design spaces more efficiently than brute-force sweeps
- Capturing and reusing design results related to geometry and solver setup.

Equally important are the limitations. Current publicly available AI does not reliably identify solver artifacts, cannot determine when a numerically converged result is physically incorrect, and does not account for manufacturing variability. For these reasons, experienced signal integrity engineers must remain firmly in the loop.

Test Case: Coaxial Feedthrough to CPW Launch Transition

The transition from a coaxial feedthrough to a coplanar waveguide (CPW) represents one of the most critical impedance discontinuities in high-speed interconnect design. At frequencies extending into the mmWave region, this transition is often dominated by high return loss and can therefore limit overall channel performance, even when the downstream transmission line is well controlled.

In the structure examined in ***Figure 1***, a UT-020 coaxial feedthrough provides excitation normal to the substrate surface and transitions directly into a ground-backed GCPW. The inner conductor of the coaxial feedthrough is soldered to the CPW center trace, while the outer conductor is tied to the surrounding ground planes through a combination of direct contact and via fencing. This geometry creates an inherently three-dimensional electromagnetic discontinuity in which field symmetry, return-current continuity, and uniform impedance must all be managed simultaneously for efficient propagation.

For these reasons, the coaxial feedthrough to CPW launch serves as an effective test case for evaluating AI-assisted electromagnetic design workflows. The topology is well-understood, fabrication constraints are well defined, and performance sensitivity is high, making it representative of the types of interconnect challenges encountered in modern high-speed signaling applications.

Parameterizing the Test Case

The UT-020 coaxial feedthrough geometry was held fixed to represent a realistic connector interface while the grounded CPW launch region was fully parameterized to enable controlled investigation of impedance transformation and return-current behavior.

Three geometric parameters were identified as primary contributors to broadband matching behavior and selected for systematic exploration, as shown in ***Figure 2***.

1. CPW Center Trace Width

The center conductor width directly influences the characteristic impedance of the CPW, and the capacitive loading presented to the coaxial feedthrough. Small changes in width can significantly alter the local impedance transition at the launch, particularly at the upper end of the frequency band.

2. CPW Ground Spacing (Gap)

The lateral spacing between the center trace and the ground planes affects both impedance and field confinement. At high frequencies, excessive gap spacing can allow return currents to spread, increasing inductance and degrading return loss, while overly narrow gaps can introduce excess capacitance.

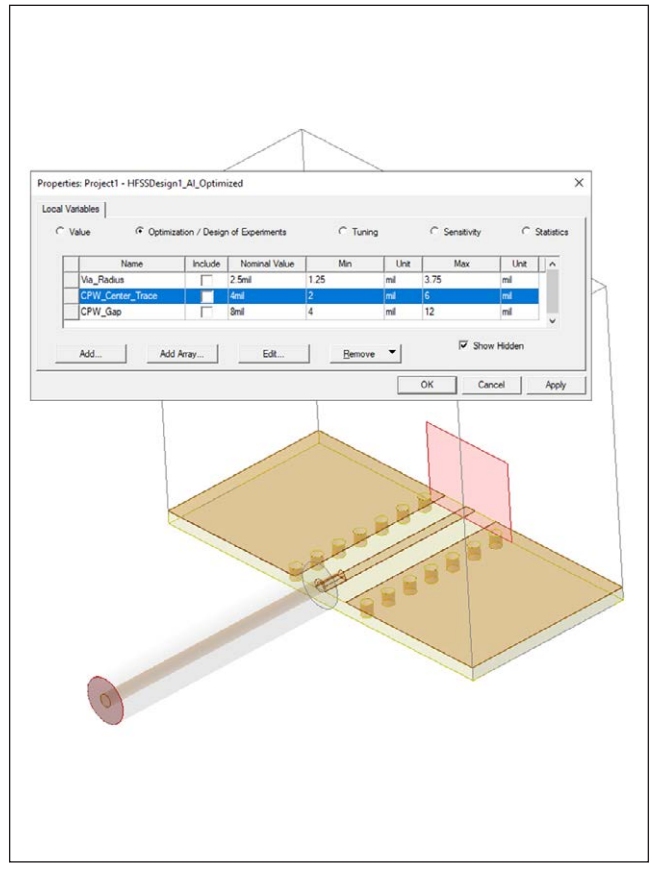


Fig. 2 Parameterized structure to create a dataset for the AI analysis.

3. Ground-Stitch Via Radius

The via radius influences the inductive and resistive properties of the return-current path between the coaxial outer conductor and the CPW ground planes. Larger vias generally reduce inductance and improve high frequency grounding, but are constrained by layout density and fabrication rules.

These parameters were swept within fabrication-realistic bounds to generate a discrete set of full wave simulations forming the basis for interpolation and response surface analysis. The overall topology, materials, and excitation method were held constant throughout the study to isolate the effects of geometric variation within the launch region.

This test case provides a challenging but well-understood benchmark for assessing AI-assisted design workflows, as small geometric changes can produce measurable broadband effects while remaining directly relevant to practical signal integrity applications.

AI Analysis Methodology

All parametric simulations were executed in the full wave solver using parametric analysis, with results exported in comma-separated value (.csv) format for post-processing by AI. Separate datasets were generated for each S-

parameter (S_{11}, S_{21}, S_{12}, and S_{22}), with each file containing frequency-dependent responses corresponding to a unique combination of geometric parameters. These CSV files served as the sole input to the AI analysis workflow, ensuring that all conclusions were derived directly from actual simulation results.

The first stage of analysis consisted of aligning the individual CSV files by geometric parameters and frequency. Each simulation point was uniquely identified by its CPW center trace width, CPW gap spacing, and ground-stitch via radius. Frequency grids were verified to be consistent across all runs to allow direct comparison and aggregation of results. Only converged simulation data were retained for further analysis.

For each geometry, the worst-case return loss was computed by the AI engine independently for both ports as the least negative value of S_{11} and S_{22} across the specified frequency range. A two-port worst-case return loss metric was then defined as the maximum of these two values, ensuring that optimization was not biased toward a single port or a narrow frequency region. Forward transmission parameters (S_{21} and S_{12}) were evaluated only as secondary checks to confirm that improvements in return loss were not achieved through nonphysical or impractical geometry changes.

The resulting scalar metric provided a single, broadband performance score for each geometric configuration. These scores were then used by the AI tool to construct a continuous response surface as a function of the three swept parameters. A second-order polynomial model, including linear, quadratic, and cross-coupling terms, was determined by the AI tool to fit to the dataset to interpolate behavior within the explored parameter bounds without extrapolation beyond the simulated space.

The interpolated response surface was evaluated by AI on a dense grid spanning the original sweep limits to identify the parameter combination minimizing the two-port worst-case return loss metric. This interpolated optimum represents a predicted best-performing geometry within the explored design space. Importantly, the interpolated result was not treated as final; instead, it was used to guide selection of one or more verification points for subsequent HFSS simulation.

This data-driven reduction and interpolation approach allows systematic extraction of broadband trends from a limited number of full wave simulations while maintaining full reliance on validated electromagnetic results for final design decisions.

Baseline Geometry Selection and AI-Estimated Refinement

An initial baseline geometry was selected by the authors based on conventional signal integrity practice and prior experience with ground-backed CPW launch structures. The baseline dimensions consisted of a 5 mil center trace width, 8 mil gap, and a 2.5 mil ground-stitch via radius. This configuration represents a reasonable starting point for a coaxial-to-GCPW transition, balancing impedance control, manufacturability, and layout density. Importantly, the baseline geometry was not selected to be optimal, but rather to be electrically plausible and fabrication-realistic without detailed tuning.

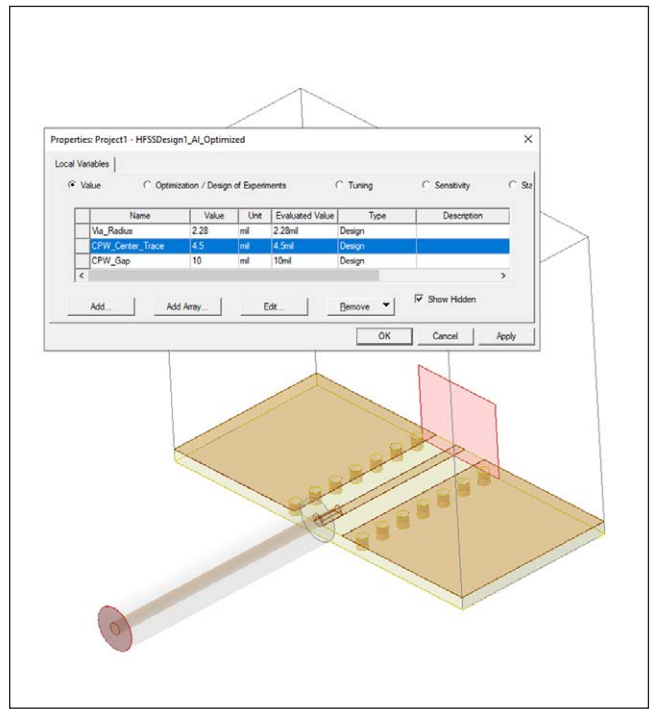


Fig. 3 The structure re-simulated with the AI optimized parameters.

Using this baseline as a reference, a structured parametric sweep was performed over fabrication-bounded ranges of the center trace width, gap spacing, and ground-stitch via radius, focusing on the 80 to 100 GHz region. These sweep results formed the input dataset for the AI-assisted analysis described previously. The AI tool was not permitted to modify the overall topology, materials, or excitation method; its role was limited to analyzing trends within the discrete parameter space defined by the parameter sweep data.

Based on interpolation of the sweep results and evaluation of broadband return loss metrics, the AI-assisted analysis identified an estimated optimum geometry of approximately 4.5 mil center trace width, 10 mil gap, and 2.28 mil ground-stitch via radius, shown in ***Figure 3***.

This geometry differs from the original baseline in two important ways. First, the wider CPW gap reduces capacitive loading and improves broadband field confinement near the launch region. Second, the reduced center trace width compensates for this increased spacing to maintain a favorable impedance transition. The slightly smaller via radius reflects a balance between return-current inductance and local parasitic capacitance within the explored parameter bounds.

Notably, the AI-estimated geometry lies within the parameter ranges explicitly provided by the parametric sweep data and does not introduce non-manufacturable features or extrapolated dimensions. The result represents a refinement derived from the available simulation data rather than a freeform or unconstrained design suggestion. This comparison highlights how AI-assisted analysis can guide designers away from reasonable but non-optimal initial guesses toward improved broadband performance using information already embedded in a limited set of full wave simulations.

Breaking Opto-Electronic Barriers with VectorStar™ ONA and O/E Calibration to 145 GHz

Upgradable to 170 GHz

O/E Calibration Module

The MN4765B O/E Calibration Module is a precision characterized, unamplified high-speed photodiode. By combining rigorous calibration, thermal stability, and broadband performance, the module delivers repeatable, traceable measurements critical for next-generation optical and photonic device development.

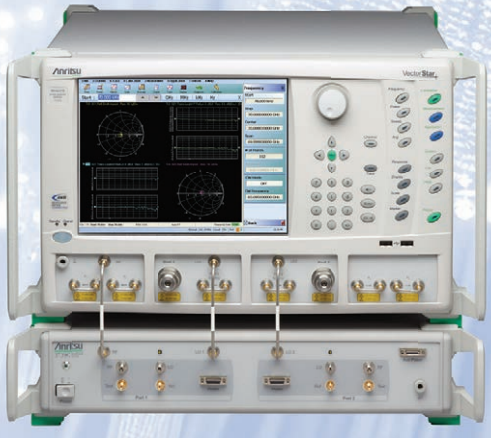


MN4765B O/E Calibration Module

- **Fast, Accurate Opto-Electronic Measurements**
- **Standards-Based Characterization**
- **Exceptional Temperature Stability**
- **Integrated Internal Biasing**
- **High Linearity Performance**
- **High Responsivity**

Opto-Electronic Network Analyzer (ONA)

The VectorStar ME7848A Opto-Electronic Network Analyzer is a VNA based measurement solution purpose-built for characterizing opto-electronic components used in high-speed photonic data transmission systems while eliminating the need for multiple analyzers when operating across different wavelengths.



VectorStar ME7838-Series VNA

Measurement Flexibility and Scalability

- Supports multiple wavelength bands including **850 nm, 1310 nm, and 1550 nm**
- Modular configuration allows rapid reconfiguration as measurement requirements evolve
- Leverages the industry leading performance of VectorStar VNAs for consistent RF-to-optical correlation



MN4775A E/O Converter

Anritsu
Advancing beyond

Discover how you can achieve measurement confidence with Anritsu.
Learn more at **www.anritsu.com**

©2026 Anritsu Company

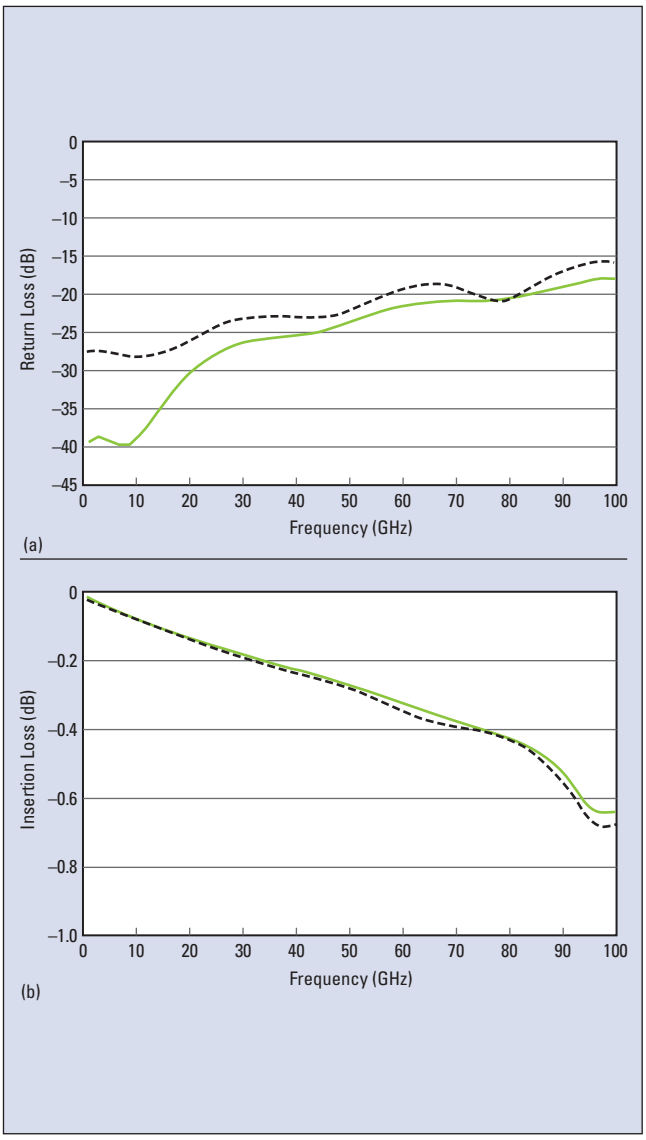


Fig. 4 Comparison of return loss (a) and insertion loss (b) between baseline geometry (black dashed curve) vs. AI-assisted geometry (green curve).

Comparison of Baseline and AI-Estimated Launch Performance

Figure 4 compares the simulated performance of the initial baseline launch geometry and the AI-estimated geometry derived from interpolation of the parametric sweep data. The baseline design (5 mil center trace width, 8 mil gap, 2.5 mil ground-stitch via radius) is shown in dashed black, while the AI-estimated geometry (4.5 mil center trace width, 10 mil gap, 2.28 mil ground-stitch via radius) is shown in green.

Figure 4a shows the broadband return loss of the launch structure from 1 to 100 GHz. Across the entire band, the AI-estimated geometry exhibits a consistent improvement in return loss relative to the initial guess. The improvement is most pronounced at lower and midband frequencies, where the baseline design shows elevated reflection levels, but it remains evident through the upper end of the band. Importantly, the AI-estimated result does not rely on a narrowband resonance or localized notch; instead, the improvement is smooth and broadband in nature, indicating a more balanced impedance transition through the launch region.

Figure 4b shows the corresponding insertion loss comparison. The two designs track closely across most of the band, indicating that the improved matching achieved by the AI-estimated geometry is not the result of excessive conductor widening or nonphysical impedance distortion. At the highest frequencies, the AI-estimated geometry shows slightly reduced insertion loss relative to the baseline, consistent with improved field confinement and reduced parasitic effects in the launch region.

These results demonstrate that the AI-assisted interpolation approach led to a geometry that improves broadband matching without introducing penalties in forward transmission. From a signal integrity perspective, this represents a more robust and predictable launch transition, achieved through the refinement of physically meaningful geometric parameters rather than changes to topology or materials.

Conclusion

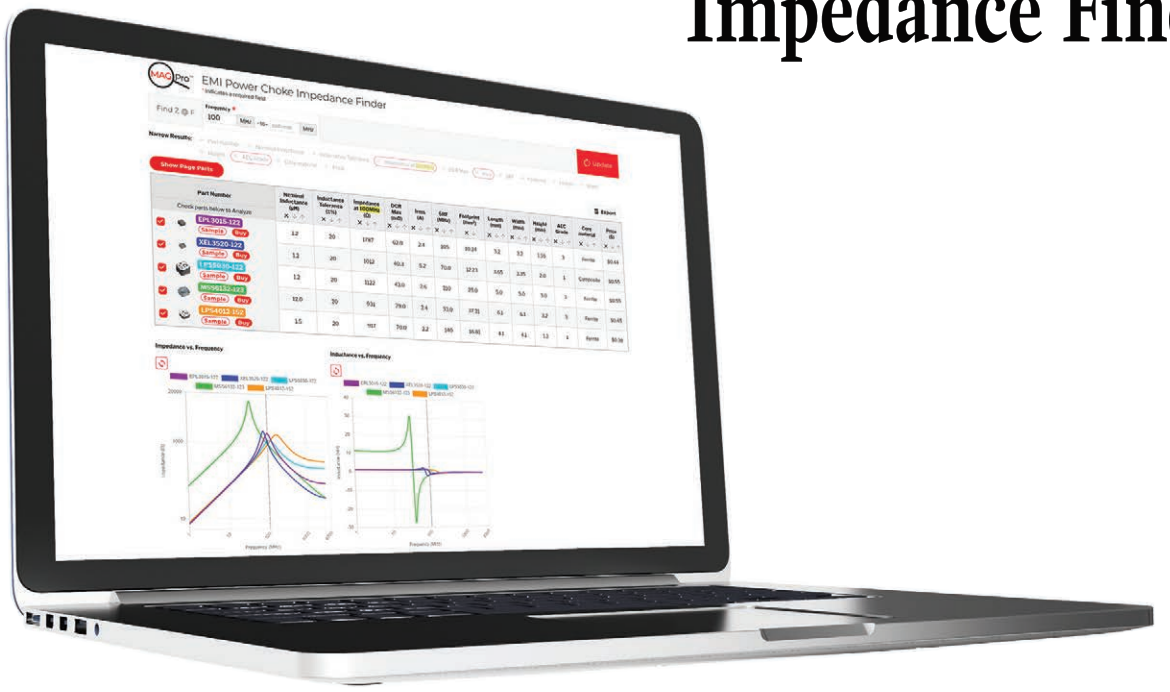
Based on the results of this study, the primary benefit of the AI-assisted workflow was not a drastic reduction in the time required to identify an acceptable launch geometry. A reasonable design target could have been achieved through conventional manual tuning within a comparable number of full wave simulations. However, manual processes can leave the hardware designer blind to the broader parametric landscape, offering little insight into whether a solution is near-optimal or if significant improvements remain unexplored.

In contrast, the AI-assisted approach enabled a systematic exploration of the design space, identifying an interpolated optimum that eludes intuitive manual adjustment. By analyzing trends across the full dataset rather than focusing on isolated tuning steps, the method revealed a geometry that delivered improved broadband performance while remaining consistent with fabrication constraints.

From a signal integrity perspective, the value of this approach lies in mitigating uncertainty rather than increasing raw design speed. However, this path is not without its own model uncertainties. The "optimum" parametric values identified by the AI tool are inherently tied to the specific algorithm and training data used; other mathematical techniques or specialized models may exist that could converge on similar or superior results more efficiently. Furthermore, there remains the risk that an AI-optimized solution may be an artifact of the mathematical model's interpolations rather than a reflection of physical reality.

As operating frequencies increase and interconnects become more sensitive to minute geometric details, the ability to identify robust optima — rather than simply satisfactory — designs will become increasingly critical. While commercially available full wave solver tools will likely incorporate specialized AI, the engineer must still weigh the trade-offs: Whether the marginal gains in performance justify the additional compute resources and time.■

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

Why Power Architectures Constrain New Space AI Missions

Ken Coffman
Vicor, Andover, Mass.

Satellite designers have crossed a threshold. We are no longer putting "a bit more DSP" in orbit; we are fielding on-board compute that rivals terrestrial edge servers, running AI inference, adaptive beamforming, dynamic spectrum management and real-time analytics at hundreds of TOPS. At that point, the hard limit on mission capability is not the processor data sheet — it is whether the power architecture can deliver sub-volt rails at hundreds of amperes, with microsecond-class transient performance, inside a small, thermally constrained, radiation-exposed enclosure.

Traditional space power topologies were not designed for this regime. They assumed modest current density, slow load steps and static operating profiles. AI workloads violate all three assumptions at once. The result is straightforward: with conventional bus and point-of-load schemes, power delivery — not compute capability — often becomes the gating function for new space missions targeting AI-class processing.

The New Reality of On-orbit Compute

Modern LEO payloads are expected to perform on-board inference, compression, routing, and autonomous control that were once strictly ground functions. Doing that efficiently pushes designers toward ultra-deep-submicron FPGAs, adaptive compute acceleration platforms, and custom ASICs whose core voltages continue to drop even as current demand grows.

The Spacechips AI1 transponder provides a concrete stress test for any power architecture. It is a radiation-tolerant, AI-enabled processor card delivering up to 133 TOPS of in-orbit AI throughput, with signal processing and machine learning tasks performed directly on-orbit rather than on the ground. Its main AI rail operates around 0.8 V with time-domain current on the order of 130 to 150 amps, precisely the regime where conventional bus-plus-POL infrastructures become inefficient and difficult to cool.

The load is not benign. AI accelerators swing between idle, inference, and burst modes, generating fast, discontinuous current steps. Voltage droop, ringing, or noise in that environment translates directly into de-rated clocks, reduced margins, and thrown-away compute capacity. With local decoupling limited by volume, mass, and MLCC reliability, the upstream power delivery network must do the heavy lifting on transient response.

System and Environmental Constraints

In new space, the usual design knobs are locked together. To control launch costs, platforms get smaller, smaller satellites harvest less power, even as workloads demand more. Every watt wasted in the power chain is a watt lost to payload function, and operators are increasingly unwilling to donate that margin.

Thermal constraints impose high current density on the power delivery network (PDN). I^2R losses in planes, vias, and wire harnesses rise with the square of current; copper that was acceptable at a few tens of amps becomes a significant heat source at hundreds. There is no forced air to save you. All waste heat must conduct through the structure and radiate away, so PDN loss appears as mechanical and mass penalties with larger radiator panels, heavier conduction paths, and more elaborate packaging.

Radiation tightens the constraints. Total ionizing dose-induced parameter shifts eat into already tight voltage margins on sub-1 V rails, while single-event effects can coincide with load transients to produce worst-case conditions exactly when the system is most stressed. Power components must tolerate radiation exposure typical of LEO and MEO missions without dragging in the cost, mass, and schedule overhead associated with fully rad-hard custom solutions.

Why Conventional Space Power Architectures are Inefficient at High Currents

Most heritage designs follow a familiar pattern: an isolated DC-DC converter steps down the spacecraft bus; a low-voltage distribution plane fans out; local PWM point-of-load regulators generate final core rails near the load. At modest current and voltage, this is workable. At 0.8 V and >100 A it becomes self-defeating.

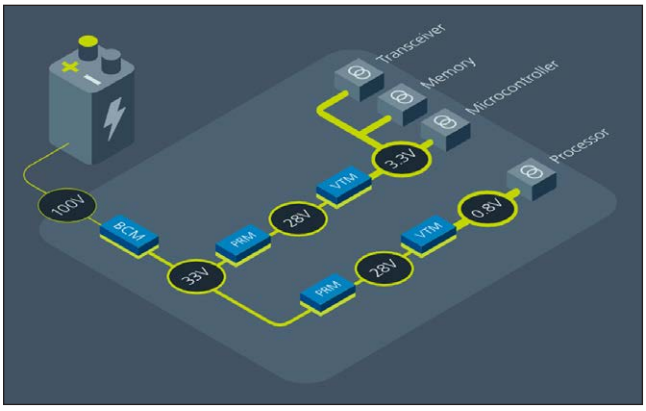


Fig. 1 Vicor Factorized Power Architecture.

- Distribution losses scale with the square of current. Pushing hundreds of amperes across even carefully optimized planes produces non-trivial I^2R heating and several tens of millivolts of droop that designers cannot afford to spend.
- Hard-switched POLs inject rich harmonic content that demands aggressive filtering — more inductors, more capacitance, and more copper — which, in turn, raises impedance and lengthens current loops.
- Classical control loops struggle with deep, fast load steps. Designers are forced to blunt the problem with large, low-ESL capacitor banks, trading away size, weight, and power while still falling short of ideal transient behavior.

At some point, trying to fix this architecture is ineffective. The structure itself is poorly suited to the load.

Current Multiplication and Factorized Power Architecture

Spacechips decided to replace the old topology rather than try to optimize it. They partnered with Vicor to implement a current-multiplication approach using the Vicor Factorized Power Architecture (FPA™) and radiation-tolerant modules specifically targeted at LEO and MEO missions, as shown in ***Figure 1***.

Current multiplication separates the functions of DC-DC conversion into independent power modules. The Vicor rad-tolerant topology splits the power conversion task into three sections: voltage transformation and isolation (BCM), load voltage regulation (PRM) and power delivery (VTM). This configuration reduces the last-centimeter copper losses and allows layout engineers more flexibility in placing modules where space is available.

The Vicor FPA deliberately splits the power conversion task into two distinct functions:

- A PRM™ regulator that generates a well-controlled intermediate voltage from an isolated BCM® bus converter. In the AI1 case, a rad-tolerant BCM steps 100 V down by one-third to 33 V and delivers several hundred watts at high efficiency; a PRM then regulates that rail to an optimal intermediate bus for the load.
- A VTM™ current multiplier located physically close to the processor that converts the intermediate rail directly to the 0.8 V AI core while multiplying current by 32.

For a given power level, increasing the distribution voltage reduces current and therefore I^2R loss by the square of the voltage ratio. As a simple example, delivering the same power on a 100 V bus instead of a 12 V bus reduces current by about 8.3× and idealized copper losses by roughly 69×, assuming the same conductor resistance.

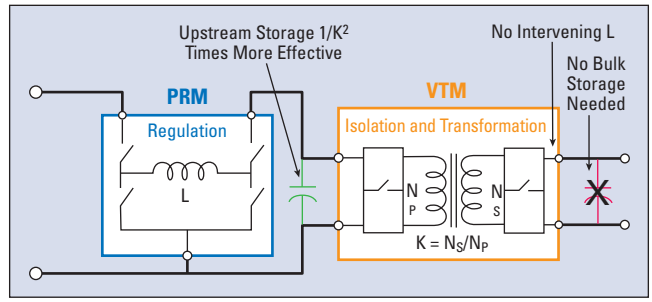


Fig. 2 Energy storage in a Factorized Power Architecture.

In practice, spacecraft power systems commonly use a 28 V bus, and higher power new space designs increasingly adopt elevated bus voltages such as 100 V to further reduce distribution losses. Even moving from 12 to 28 V cuts distribution current by a bit more than 2× and I^2R loss by about 4 to 5×, while the step from 28 to 100 V yields a further ~3.6× reduction in current and ~13× reduction in resistive loss for the same power. That is the difference between a hot marginal backplane and a tractable thermal design.

100-to-33-to-0.8 V Example

To make the architectural impact tangible, consider a representative case consistent with the AI1 power path: a 100 V spacecraft bus feeding two AI rails: one at 0.8 V, the other at 3.3 V, each delivered by a 150 W VTM for a total of 300 W of on-orbit AI compute (see ***Figure 2***). With voltage transformation designed in a very compact form factor and placed close to the load, designers have added flexibility in component placement and values.

The 150 W VTM at 0.8 V supplies about 187 A; two such rails represent 300 W of AI-class power near the die, even before accounting for headroom and conversion loss. One cannot route that current around a backplane; the point of the architecture is to confine those hundreds of amperes to a very short, explicitly engineered region near the processors.

In a more conventional bus-plus-PoL approach with, say, a 12 V intermediate rail feeding two local multiphase regulators, that same 300 W of AI implies about 25 A on the 12 V plane. With realistic copper and via resistance on a dense RF-plus-digital card, seeing 10 mΩ of effective distribution resistance is not optimistic. At 25 A, that path dissipates about 6.25 W in I^2R loss and drops 0.25 V, more than 2% of the 12 V rail, before the current even arrives at the PoLs. Scaling to 400 W or beyond for future variants makes the loss and droop climb quadratically.

The AI1 chain looks different. A rad-tolerant BCM takes the 100 V bus and delivers up to 400 W at an intermediate voltage at one-third the input (33 V), with enough headroom to drive two PRM: VTM rails at 150 W each. Each PRM regulates its own 33 V-class intermediate node; each VTM then converts to the appropriate point of load voltage. At 300 W total, the 33 V distribution bus carries only about 9 A instead of 25 A. With the same 10 mΩ

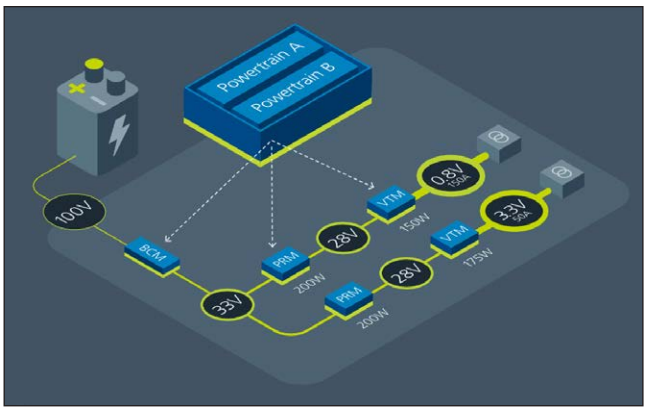


Fig. 3 Internally redundant power trains integrate and automate monitoring and control of the power rails.

effective path resistance, I^2R loss falls to roughly 0.8 W and the DC drop to about 0.09 V — less than 0.3% of the intermediate voltage. That is roughly a 5 to 8× reduction in distribution loss and a meaningful recovery of voltage margin, gained purely by moving the heavy lifting to a higher voltage, lower current domain.

Downstream, each VTM trades that 33 V back down to 0.8 V and multiplies current by 32× or 8×, so the high-current loops (approximately 187 A in the 32× rail) are reduced to a few centimeters of copper, via fields and decoupling specifically designed for that density.

The net result is that the same 300 W AI workload either taxes a hot, not well-regulated 12 V plane in a legacy architecture or rides on a relatively cool 33 V intermediate bus feeding two localized current-multiplier stages in the FPA implementation. The AI1 transponder takes the latter route and, in doing so, recovers both thermal and voltage headroom that would otherwise constrain on-orbit compute.

Electrical Advantages for AI-class Loads

Placing the current-multiplying stage centimeters from the FPGA or ACAP collapses the high-current path. The intervening copper is short and wide, and the effective output impedance at the die is dramatically lower than any equivalent 12 V distribution plus PoL scheme. Internally redundant power trains for single-event effects mitigation. For single-event effect mitigation, internally redundant power trains integrate and automate monitoring and control of the power rails — transparently relieving the design engineer from implementing this logic in external circuitry, as shown in ***Figure 3***.

Resonant, soft-switched conversion in the FPA modules further reduces conducted and radiated EMI compared to hard-switched alternatives. Lower harmonic content means lighter filters, shorter loops and better design margins against DO-160-style conducted emissions constraints when applicable to New Space platforms.

Crucially, the transformation stage does not rely on a slow global feedback loop to manage transient behavior. Its fixed-ratio operation and very low intrinsic impedance allow it to source abrupt current steps in microseconds instead of tens or hundreds of microseconds, keeping a 0.8 V rail inside a narrow tolerance window even under worst-case AI bursts.

There is a second-order benefit: capacitance multiplication. A current-multiplying transformer effectively scales the apparent load-side capacitance seen by the source by approximately the square of the turns ratio. In practice, that means a modest bank of carefully placed MLCCs at the AI rail behaves like a much larger reservoir as seen by the upstream PRM, allowing designers to meet transient requirements without covering the card in fragile, mechanically risky large-case ceramics.

Radiation Tolerance at the Architecture Level

Current multiplication also changes how designers approach radiation in the PDN. Rad-tolerant Vicor FPA modules for new space are qualified to levels typical of LEO and MEO missions (tens of krad total ionizing dose) and use redundant powertrains and SEE-tolerant control to maintain operation under single-event stress. But the more interesting aspect is architectural.

By keeping tight voltage regulation at a higher, more forgiving intermediate voltage and handling high current via a fixed-ratio transformer near the load, the design reduces the number of nodes where a radiation-induced parameter shift can immediately threaten a 0.8 V margin. The regulation stage sits at tens of volts with ample headroom; its drift and SEE behavior can be managed with conventional de-rating and system-level mitigation. The transformation stage, lacking a complex fast feedback loop, is inherently less susceptible to control-loop upset.

The AI1 implementation adds another layer: radiation-tolerant modules incorporate dual powertrains, allowing the load to be driven at full power from either side, with autonomous fault detection and switchover. Instead of hardening every discrete component in a sprawling PDN, Spacechips can treat radiation as a system-level problem with well-characterized, modular building blocks.

For many LEO and MEO missions, this rad-tolerant, modular approach avoids the cost, mass, and procurement complexity of fully rad-hard parts while still satisfying mission reliability objectives. It is a pragmatic middle path between consumer-grade COTS and bespoke rad-hard ASICs.

Raising the Ceiling on On-orbit AI

As satellites adopt higher levels of on-board autonomy, the practical limit on useful AI is increasingly set by the ability of the power architecture to deliver clean, tightly regulated, high-current rails under worst-case transients in a radiation environment and within tight size, weight, and power budgets. By moving from legacy low-voltage distribution to current multiplication with Factorized Power Architecture, Spacechips and Vicor effectively raised that ceiling. The AI1 transponder is not just another RF board. It demonstrates that the bottleneck has shifted from computing capability to optimized power delivery in new space.

If you design the power delivery network as if the processor was the scarce resource, you will leave compute capability on the table. If you design the processor board as if power density, current distribution, and radiation behavior are the scarce resources, you can put server-class AI in a small LEO satellite and still close the thermal and reliability budgets constraints. That is the architectural shift this generation of missions is forcing, and current multiplication is one way to achieve it. ■

SDS7000A Digital Storage Oscilloscope

Up to 8 GHz bandwidth, 12-bit ADC, with Real-Time Spectrum Analysis (RTSA).



Real-Time Spectrum Analysis (RTSA)

www.siglentna.com
quote@siglentna.com

SIGLENT®

Every Bench. Every Engineer. Every Day.

Machine Learning Models for SI/PI Analysis with Meshed Planes

Jiyue Zhu, Regina Thahir, Xiaoyan Xiong, Gang Kang, and Jian Liu
Cadence, San Jose, Calif.

As data rates continue to scale into the multi-tens of gigabits per second, the tolerance for uncertainty in interconnect behavior has significantly diminished. At the same time, packaging and board-level technologies are evolving toward higher density, heterogeneous integration, and greater compliance with standards. These trends have driven widespread adoption of meshed reference planes, including cross-hatch ground planes in flexible and rigid-flex designs (see ***Figure 1***), and perforated planes with degassing holes in 3D-IC packages (see ***Figure 2***).

From an electromagnetic (EM) perspective, meshed planes are a departure from the continuous reference planes assumed in most analytical and quasi-static models. The discontinuities introduced by the mesh alter current return paths, increase effective inductance, and can lead to non-uniform field distributions. These effects impact both signal integrity (SI), through changes in characteristic impedance, delay, coupling, and power integrity (PI), through degradation of power distribution network (PDN) impedance and resonance behavior.

Conventional approaches to modeling meshed planes rely on 3D full-wave EM solvers. However, a single metal layer may contain thousands to millions of apertures, making direct simulation computationally expensive and often impractical for iterative design flows. Furthermore, SI and PI effects must often be evaluated simultaneously, further increasing model complexity. This article presents machine learning (ML)-based modeling approaches that efficiently characterize SI and PI behavior in systems with meshed planes.

Artificial neural network (ANN) models are developed for pre-layout SI analysis of traces referenced to meshed ground planes, with Bayesian optimization employed for systematic hyperparameter tuning. The proposed methods demonstrate high accuracy compared with 3D full-wave EM simulations while achieving orders-of-magnitude reduction in computation time. The results suggest that ML-enabled modeling can serve as a practical and scalable solution for SI/PI analysis of complex meshed-plane structures.

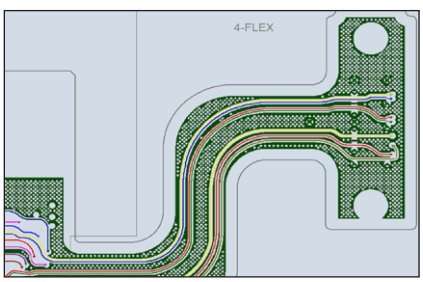


Fig. 1 Flexible and rigid-flex meshed plane design with cross-hatch planes.

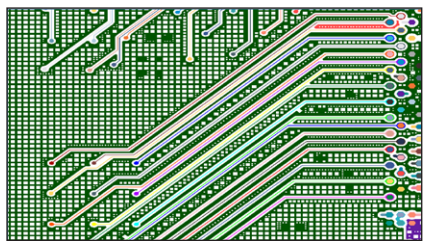


Fig. 2 3D-IC package design planes with degassing holes.

Challenges in SI/PI Analysis of Meshed Planes

The modeling of meshed planes poses several unique challenges that distinguish it from conventional solid reference plane modeling:

- Geometric complexity: The sheer number of holes or mesh elements creates a large-scale 3D problem that is difficult to simplify without loss of fidelity.

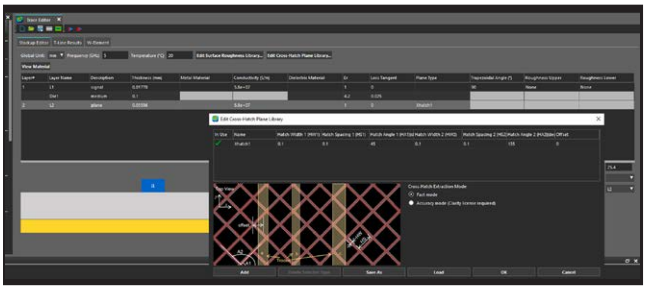


Fig. 3 Pre-layout SI analysis trace editor referenced to meshed ground (cross-hatch) planes.

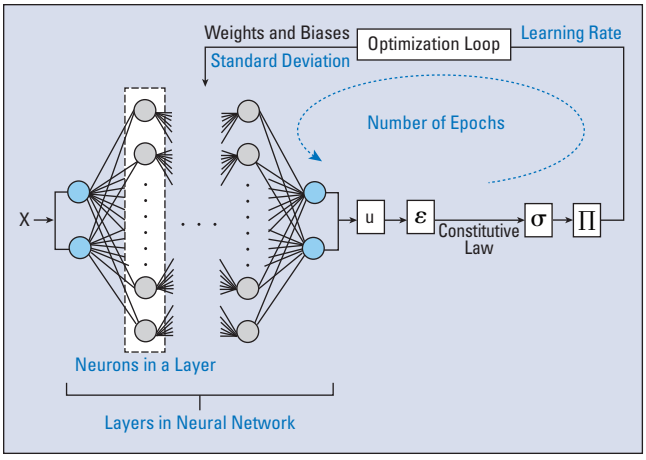


Fig. 4 Bayesian optimization for ANN hyperparameter tuning.

- Full-wave requirements: The EM behavior of meshed planes is inherently 3D, invalidating many quasi-static or 2D assumptions commonly used in interconnect modeling.
- Computational cost: Applying 3D full-wave methods to such structures requires substantial memory and runtime, limiting their use in early design stages.
- Coupled SI and PI effects: Meshed planes simultaneously affect signal return currents and power delivery paths, necessitating a unified modeling approach.

In practical design environments, these factors severely constrain the ability to perform rapid design-space exploration or optimization when meshed planes are involved.

ML Models for Pre-Layout SI Analysis

To address these challenges, ML models are being introduced as surrogate models for EM-based SI analysis. They focus on traces referenced to meshed ground planes, particularly cross-hatch structures commonly used in flexible and rigid-flex boards (see ***Figure 3***).

When a signal trace is referenced to a meshed ground plane, the return current distribution becomes fragmented and non-uniform. This behavior results in increased effective inductance and modified capacitance compared to solid-plane references, leading to deviations in characteristic impedance and signal delay. Accurately capturing these effects is essential for impedance control and timing closure.

The proposed ANN model takes a comprehensive set of design parameters as inputs, including cross-hatch parameters (hatch pitch, hatch width, and hatch angle), trace geometry parameters (trace width and trace-to-trace spacing for coupled lines), position of traces with respect to meshed-plane (trace offset and trace rotation angle),

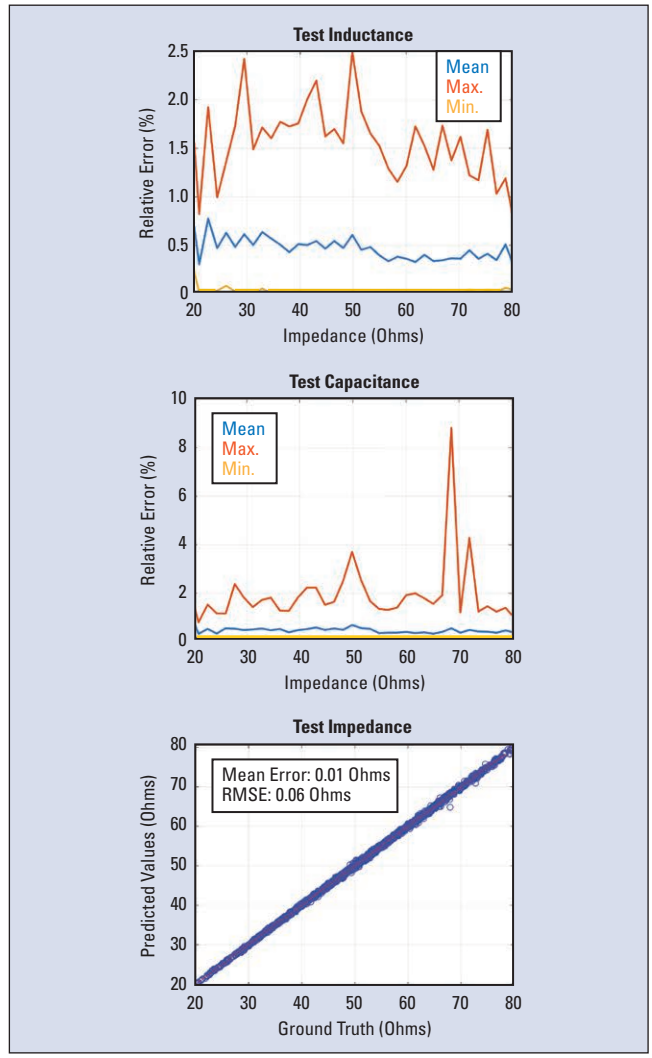


Fig. 5 ANN validation of a single trace microstrip with cross-hatch planes.

stackup parameters (thicknesses of metal layers, dielectric layers, and trace layers), and material parameters (dielectric constant and conductor properties).

The outputs of the model include per-unit-length inductance and capacitance, single-ended trace impedance, propagation delay, and velocity, differential and common-mode impedance for coupled traces, and differential delay and delay mismatch. These outputs directly support pre-layout SI analysis and constraint definition.

ANN Architecture and Hyperparameter Optimization

ANNs offer sufficient flexibility to approximate the nonlinear relationships between meshed-plane geometry and EM response. However, model performance is strongly influenced by hyperparameter selection, including hidden layer count, hidden dimension, learning rate, and training epochs.

Rather than relying on manual tuning or grid search, Gaussian process–based Bayesian optimization is employed to identify optimal ANN hyperparameters (see ***Figure 4***).

This approach provides several advantages:

- Systematic exploration of the hyperparameter space
- Flexible search domains, allowing continuous parameter

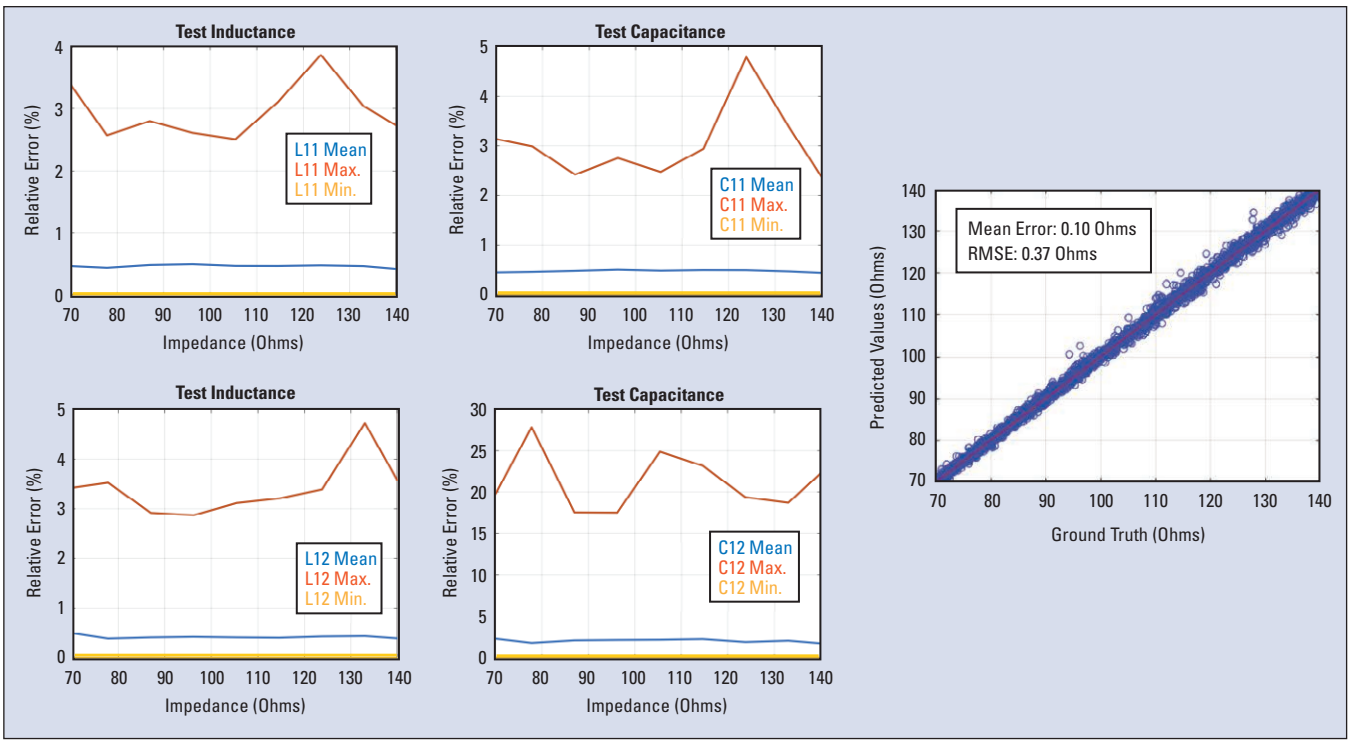


Fig. 6 ANN validation of coupled trace microstrip with cross-hatch planes.

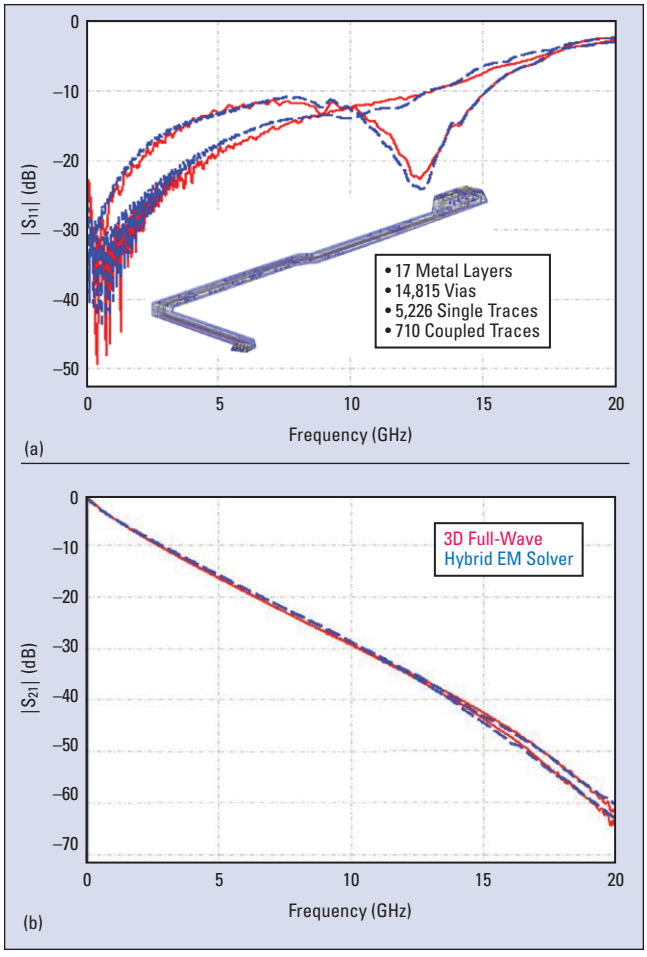


Fig. 7 The hybrid EM solver leverages ML techniques to extract part of a layered interconnector with high efficiency and reliability: (a) return loss and (b) insertion loss.

ranges rather than discretized values
- Efficient convergence, achieving improved accuracy with fewer training iterations.

Notably, the hidden layer dimensions are not constrained to traditional powers-of-two conventions, enabling more efficient network configurations.

Model Validation and Results

The trained ANN models were validated against reference data generated using 3D full-wave EM simulations. For single-ended microstrip traces referenced to cross-hatch ground planes, the ANN model demonstrates strong predictive accuracy (see ***Figure 5***). Across a characteristic impedance range of 40 to 60 Ω, the mean prediction error is below 0.6%, with a maximum error below 4%. The predicted impedance values closely track the full-wave simulation results, confirming the model's ability to capture the impact of meshed return paths.

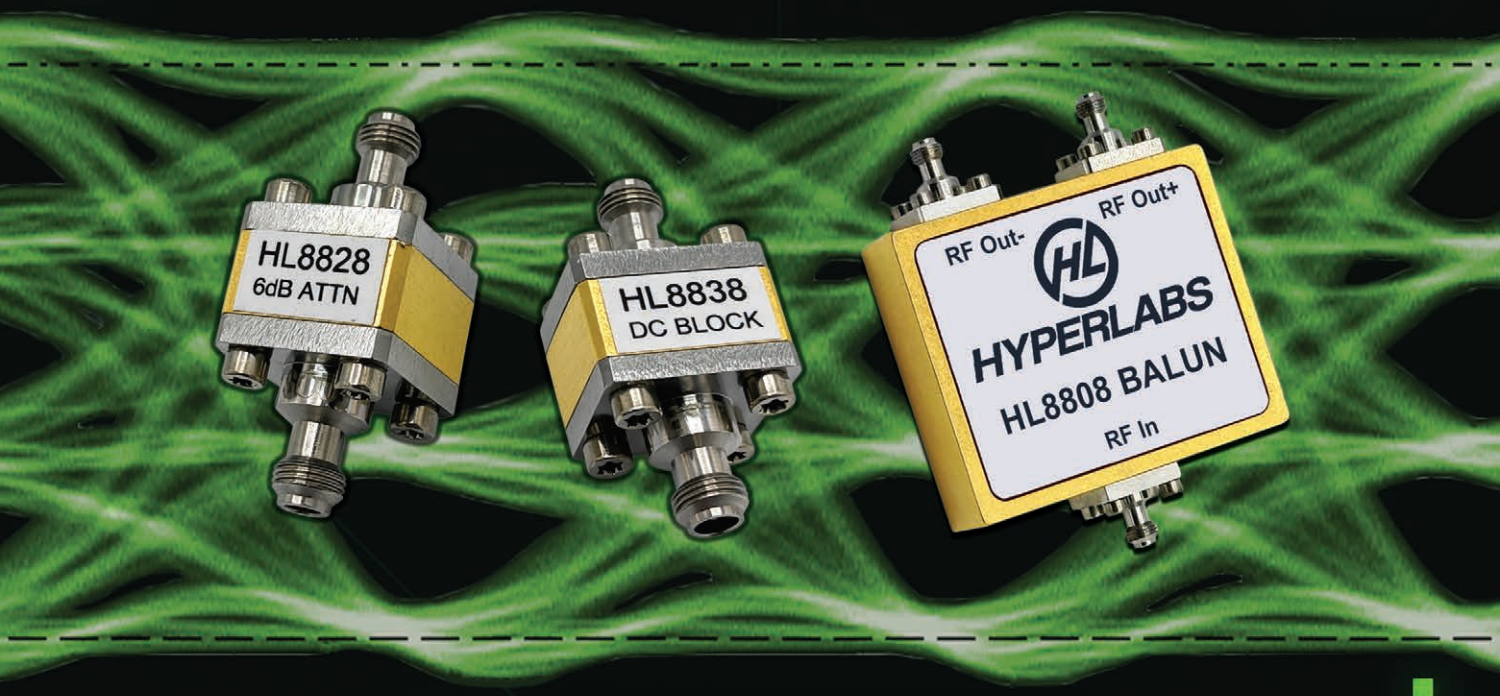
For coupled microstrip traces, both self- and mutual-coupling terms are evaluated (see ***Figure 6***). For differential impedance in the 100 to120 Ω range, the self-term predictions achieved a mean error below 0.6% and a maximum error below 4%. The mutual capacitance and inductance terms exhibited small mean error; however, larger maximum errors were observed in weakly coupled cases. These deviations are attributed to the increased sensitivity of coupling parameters when absolute coupling levels are very small. Overall, the validation results indicated that the ML models provide sufficient accuracy for practical SI analysis, particularly in early design and optimization stages.

Extension to Hybrid EM Solver-Based SI/PI Characterization

In addition to pre-layout SI modeling, ML models can be extended to SI/PI analysis of meshed planes by integrating

LEADING THE INDUSTRY TO 145 GHZ

KEEP YOUR EYES WIDE OPEN.™



- 0.8 MM CONNECTORS
- DC – 145 GHz BANDWIDTH
- ULTRA-LOW INSERTION LOSS
- PRECISION AMPLITUDE & PHASE BALANCE
- DESIGNED & MANUFACTURED IN THE USA

HYPERLABS

COLORADO | OREGON | USA
WWW.HYPERLABS.COM

them into hybrid EM solver workflows. This approach uses ML models to rapidly characterize the local EM behavior of meshed planes, while the hybrid EM solver captures global interactions across the package or board (see ***Figure 7***). The trained ANN models for parallel-meshed planes were used to correct the finite elements of the parallel-plane field domain, while the ML-based models for traces referenced to meshed ground were used to correct the transmission-line parameters of traces.

The hybrid EM solver leveraging ML models was used to extract the S-parameters of a small part of a layered interconnector with 17 metal layers, 14,815 vias, 5,226 single-ended traces, and 710 coupled traces (the complete structure is out of the capability of 3D full-wave simulation). The extraction was conducted on a Windows server with an Intel Xeon processor and 1.0 TB of RAM. The extracted return loss and insertion loss agreed well with the 3D full-wave simulation in a broad frequency region, as shown in Figure 7. The hybrid EM solver is much faster than 3D full-wave simulation and has obtained up to 50x speedup with consistent accuracy. It also requires less memory, as compared in ***Table 1***.

TABLE 1
PERFORMANCE COMPARISON

EM Solver	CPU Number	Simulation Time	Memory
3D full-wave solver	72	47 hrs. 11 mins.	374 GB
Hybrid EM solver	40	53 mins.	148 GB

Key Takeaways

- Meshed reference planes significantly alter current return paths and cannot be accurately modeled using solid-plane assumptions
- Direct 3D full-wave simulation of meshed planes is often impractical for iterative design flows
- ML models can accurately predict SI metrics such as impedance and delay when trained on high-quality EM data
- Bayesian optimization provides an efficient and robust method for ANN hyperparameter tuning
- ML-based surrogate models are well-suited for pre-layout analysis and early-stage design exploration involving meshed planes.

Conclusion

Meshed planes are becoming indispensable in modern electronic systems, yet they pose significant challenges for conventional SI and PI analysis methodologies. This article has discussed how ML-based modeling, combined with systematic hyperparameter optimization, offers a practical and accurate alternative to brute-force EM simulation. By enabling fast and reliable prediction of key SI metrics for traces referenced to meshed planes, the proposed approach supports efficient design-space exploration and informed engineering decision-making.■

SIJ
Stay Tuned for our Next Signal Integrity Journal Coming in October 2026
Cover Feature: High Speed Interconnect Technologies
Ad Close: 9/8
Ad Materials Due: 9/10
Get the Best of Both Worlds:
PRINT + LEADS...
ASK HOW!!
Signalintegrityjournal.com/advertise



StepLoad Pro™ with SEPIA

VRM & Power Supply Test Automation Software

Precise Measurement in Less Than 1 Minute

Transform VRM and power supply testing and get time back with StepLoad Pro with SEPIA. This software drives test equipment to measure step load response across a wide range of frequencies, and SEPIA integration allows the software to work **even on non-LTI systems**.

StepLoad Pro with SEPIA simulates a SPICE model of the VRM, automatically analyzes the impedance, Q, and phase margin in a time domain analysis, then identifies the necessary design fixes to reach stability.

ANALYZES	PROVIDES
• PDN Impedance	• Design Fixes
• Phase Margin	• SPICE Model
• Q	



SCAN HERE

to learn more

www.signaledgesolutions.com

SIGNAL EDGE SOLUTIONS

SI/PI Modeling & Simulation Experts

Stop Breaking the Loop: Modern Closed-Loop Stability in ADS

Benjamin Dannan, Steve Sandler, and Heidi Barnes

Power integrity engineers are running into a new reality: most modern voltage regulator modules (VRMs) control loops are no longer linear, time-invariant systems, and the old "break the loop and run a Bode plot" workflow does not tell the whole story anymore. Keysight's ADS offers a complete approach by letting you simulate the entire power-delivery path end-to-end — from the VRM and PCB to the decoupling network and all the way into the ASIC. With the Sandler State-Space Average Model at the core, ADS becomes a practical environment for closed-loop stability assessment that reflects how today's converters actually behave.

Two measurement-driven methods fall naturally into this workflow. Non-invasive stability measurement (NISM) provides a frequency-domain view of stability directly from output-impedance data, while scope embedded power integrity analysis (SEPIA) turns step-load responses into a time-domain stability model you can drop straight into ADS. Both approaches deliver quantitative stability metrics without breaking the loop, and both work cleanly with non-LTI systems. Bode plots are still available in ADS for traditional analysis, but they don't apply to most modern VRMs. These newer methods give engineers a realistic, simulation, and measurement-ready way to evaluate and optimize stability across the full power distribution network (PDN).

Control loop stability has always been of interest in power electronics because the control loop ultimately determines the circuit's end performance. Poor stability shows up immediately as elevated noise. Steve Sandler first presented on this topic at CMSE 2011, where he showed the relationship between stability and a power supply's closed-loop functions. He continued presenting on this in the years that followed, including his APEC 2012 paper presentation titled, "High-Fidelity Power Supply Measurement." A key takeaway from that paper was that all closed-loop performance is affected by control loop stability, as summarized on slide four of that presentation, see in ***Figure 1***.

We do not always have the option of traditional Bode plot measurements, since in many cases open-loop Bode plots are simply not possible. Voltage references and three-terminal regulators do not expose the feedback loop. Op-amps and RF amplifiers have bandwidths too high for practical signal injection. PMICs make this even more challenging by integrating multiple regulators into a single chip, with only the outputs available.

To overcome these limitations, Steve Sandler developed a closed-loop stability tool, NISM, which mathematically relates the measured output impedance to the stability of the control loop. NISM provided a simple, quantitative assessment of stability without access to the loop. Adoption was slow at first as engineers were comfortable with Bode plots, and NISM did not always agree with traditional phase-margin/gain-margin results. Sandler repeatedly showed that NISM was correct and that the Bode plot, limited to phase and gain margins, does not always indicate relative stability. Engineers seemed to know this already: after running a Bode plot, they would generally check the step-load response. If the step load was the trusted truth, why not just use that in the first place?

That led Sandler to develop SEPIA in 2013, which extracts quantitative stability directly from the step-load response. Engineers trusted step loads, SEPIA agreed with NISM, and both methods worked non-invasively. Still, adoption took time. It was not until the last few years — when measurement access became even more limited, noise limits tightened, dynamic currents increased, and VRMs shifted to non-linear, time-variant control schemes — that interest accelerated. Suddenly, the industry needed methods that do not rely on assumptions of linearity or access to the control loop.

Simulation has always been part of the workflow, and this is where the Sandler State-Space Average Model fits in. First published in 1995 and updated several times since, it provides a practical way to simulate closed-loop behavior, including large-signal effects, EMI, and other real-world behaviors. Integrating NISM and SEPIA into Keysight ADS was the obvious next step. Ben Dannan and Heidi Barnes played a major role in getting these methods into ADS, and Masashi Nogawa — through his publications and continued advocacy — has been instrumental in driving renewed interest and adoption. These methods are now becoming even more accessible as they are integrated directly into measurement instruments. NISM has been available in the OMICRON Lab Bode Suite since 2012 and is now supported in a range of oscilloscopes, VNAs, and simulators. SEPIA is now being integrated into oscilloscopes as well.

The introduction of Picotest high-speed step loads allows clean, high-bandwidth data to be acquired, furthering the benefits of SEPIA.

This article demonstrates a rigorous simulation-based approach to power integrity within Keysight ADS. We explore why traditional stability metrics often fail to capture the complexities of modern VRMs, particularly in non-linear or time-varying scenarios. By leveraging NISM and SEPIA directly within the ADS environment, designers can achieve a more realistic and accurate view of system stability. This unified simulation workflow ensures that the virtual model reflects the true behavioral dynamics of modern converters — principles that translate directly from the digital workspace to physical hardware validation.

What Can We Get from the Step Response?

Before we dive into the mechanics of SEPIA, let's revisit the fundamentals. Consider a standard PDN model shown in ***Figure 2***: if you were to apply a single step response, whether in a simulation or on a physical circuit, what data could you actually extract from this step response?

Most engineers would expect a basic look at transient behavior. But what if a single step response could reveal the entire DNA of your circuit? Imagine extracting every individual component value, alongside the loaded Q, phase margin, output impedance, and both natural and forced resonance frequencies. Furthermore, what if that same measurement told you exactly how to fix the design? SEPIA makes this possible, turning a simple voltage measurement into a complete diagnostic roadmap.

By utilizing the built-in SEPIA functions within Keysight ADS, we can transform a standard time-domain transient waveform into a wealth of design data. By simply placing two markers on the waveform, SEPIA extracts a complete, high-fidelity SPICE model of the circuit.

This goes far beyond basic visualization; it allows us to instantly determine the quality factor (Q) and phase margin — the two most critical metrics for assessing stability (see ***Figure 3***). Whether you are validating a physical board on an oscilloscope or iterating in simulation, SEPIA provides the automated support necessary to optimize your PDN and ensure your design meets its target impedance goals in under a minute.

What is SEPIA?

Developed by Steve Sandler of Picotest, SEPIA is a sophisticated tool designed to simplify power integrity design, analysis, and simulation. It allows engineers to extract a high-fidelity electrical SPICE model of a power supply's output impedance from a simple voltage measurement, making it highly effective for both LTI and non-LTI systems.

By accurately determining critical parameters, including loaded Q, phase margin, capacitance, excess inductance, and both natural and forced resonance frequencies, SEPIA provides a comprehensive view of system behavior. It specifically uses Q and phase margin to assess stability performance, and it can suggest the minimum adjustments required

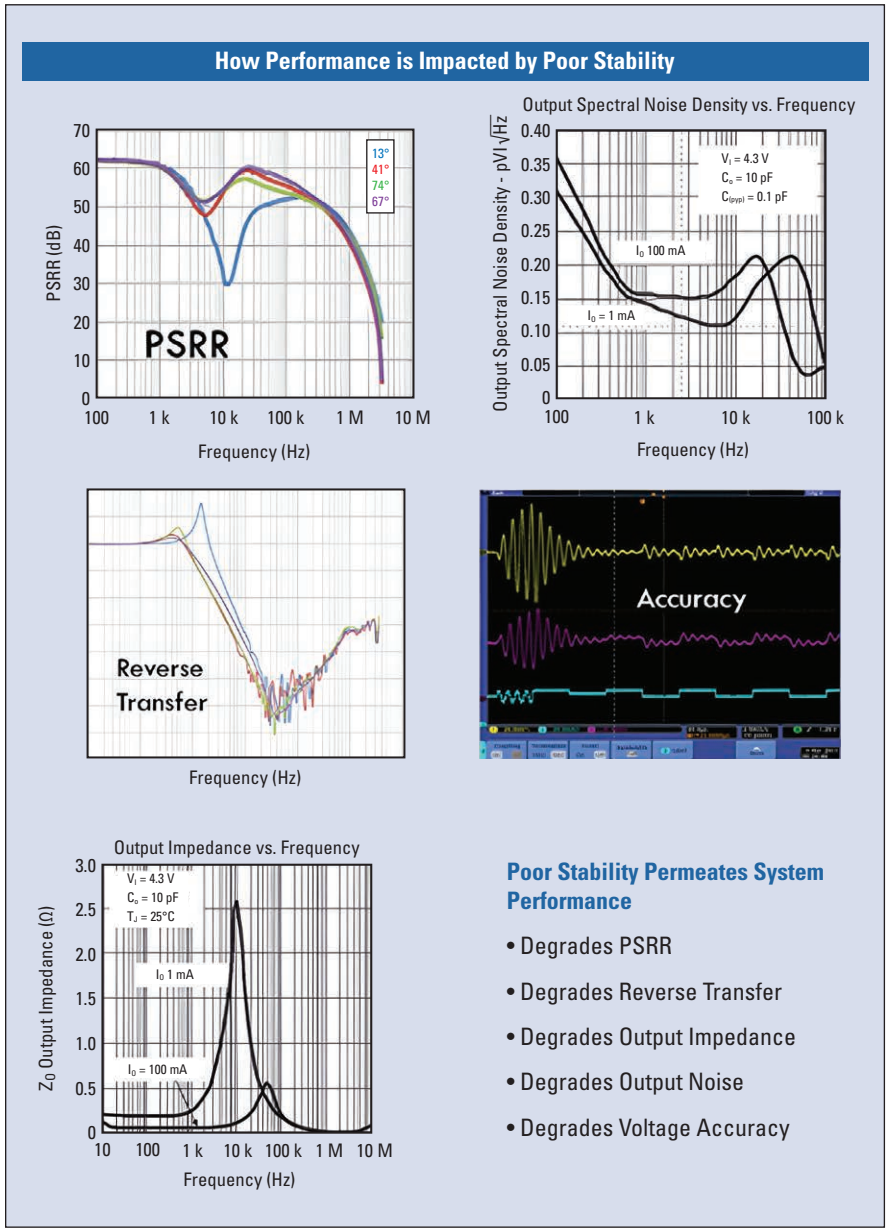


Fig. 1 A summary of all closed-loop power supply performance metrics and how they are impacted by control loop stability.[1]

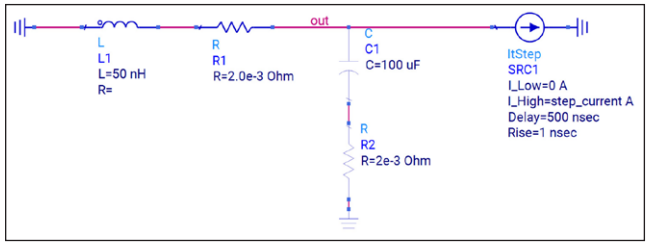


Fig. 2 Simple PDN with a step response.

to achieve a stable solution or identify the ideal capacitor values needed to meet a specific target impedance.

Key Capabilities of SEPIA

- Stability Metrics: Directly calculates phase margin and Q to ensure power supply reliability.
- Modeling: Converts physical measurements into actionable SPICE models.
- Optimization: Recommends specific capacitor solutions to hit your "Target Impedance" goals.
- Versatility: Functions across various system types, even those that are non-linear or time-varying.

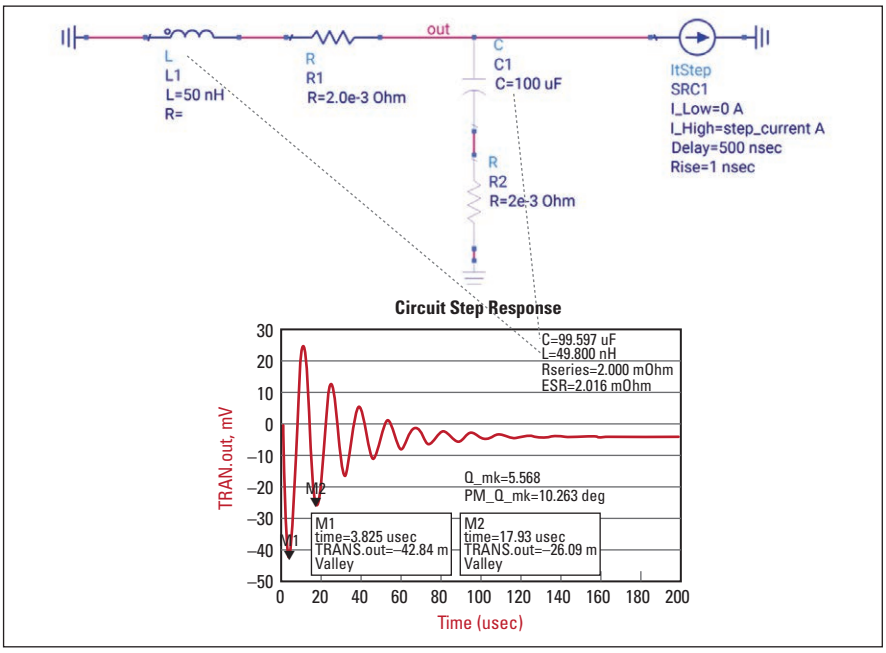


Fig. 3 Simple PDN with a step response showing SEPIA.

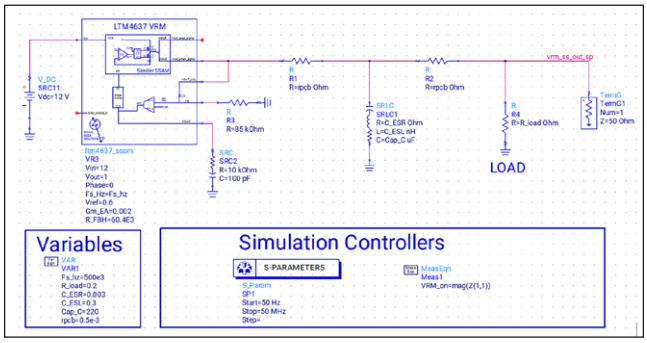


Fig. 4 LTM4637 stability frequency domain simulation example using NISM in ADS.

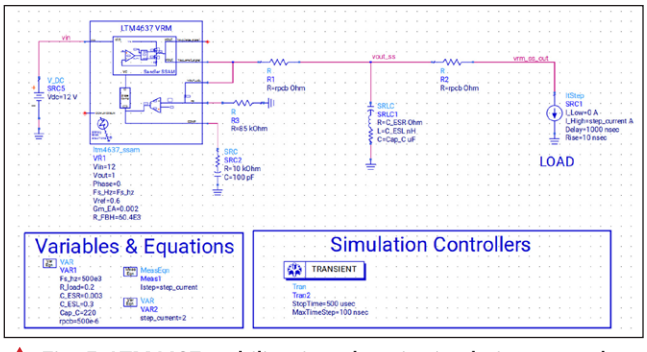


Fig. 5 LTM4637 stability time domain simulation example using SEPIA in ADS.

SEPIA Simulation Comparison to NISM Using State-Space Average Models

To validate the accuracy of our stability metrics, we transition from theory to a practical simulation using a Sandler State-Space Average VRM model of the LTM4637 in Keysight ADS. This setup allows us to directly compare results from NISM and SEPIA to observe how each handles the same circuit dynamics.

The test circuits (shown in ***Figures 4*** and ***5***) convert a 12 V input to a 1 V output, with R3 setting the output voltage and SRC2 providing control loop compensation. To simulate a realistic environment, we've included SRLC1, a lumped-capacitor model representing the PDN, and R1 and R2 to account for some PCB copper resistance. By applying a transient step response, we can now evaluate the SEPIA extraction against the NISM baseline to verify the precision of our phase margin and Q factor calculations.

Figure 6 shows the NISM results on the left and the SEPIA results on the right. If we compare the two results highlighted in green, we see that NISM calculated Q = 3.841, whereas SEPIA calculated Q = 4.051, or basically the same answer. Additionally, when we compare the phase margin results highlighted in purple, we see NISM calculated PM = 14.504 degrees, and SEPIA calculated PM = 14.072 degrees, or again, the same answer. This is very powerful because we have shown two different methods that correlated across the frequency and time domains. Additionally, this validates our SEPIA result.

Beyond providing stability metrics, SEPIA successfully extracted a complete, high-fidelity SPICE model from the transient data. The extracted parameters — capacitance (207.192 µF), ESR (0.245 mΩ), inductance (5.325 nH), and series resistance (1.006 mΩ) — allow us to reconstruct the physical PDN with surgical precision.

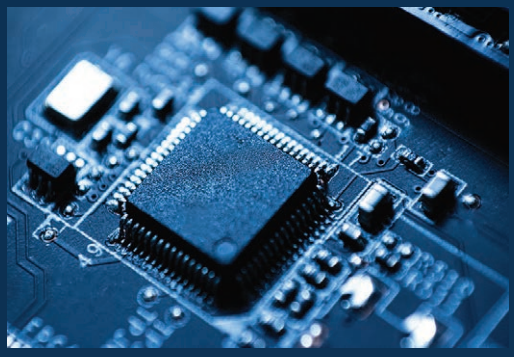
As shown in ***Figure 7***, we can generate a passive SPICE model using these values and re-run the time-domain simulation for validation. The comparison in ***Figure 8*** is striking: the step response of the original LTM4637 Sandler State-Space Average VRM model is virtually identical to the response of our extracted SEPIA model.

Conclusion

SEPIA represents a paradigm shift in power integrity by bridging the gap between measurement and simulation. Unlike traditional Bode plots, SEPIA excels where others fail — providing accurate stability metrics like phase margin and loaded Q for both LTI and non-LTI systems. By allowing engineers to extract a SPICE model from a simple voltage measurement, it facilitates a "design-by-analysis" approach that ensures the PDN remains stable and optimized.

Leaders in High-fidelity Signal Solutions

Maury Microwave has 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 Programmable Noise Generators and MPA Series Amplifiers



...Through Manufacturing...



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

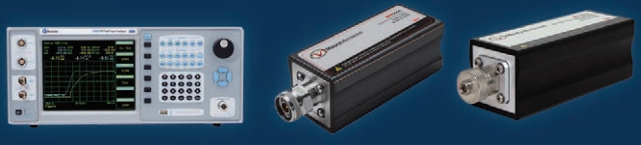
...To Deployment...



HSM Series RF Synthesizer Modules and Interconnects



...And Maintenance...



4500C Power Measurement Analyzer and RTP4000 and RTP5000 USB Peak Power Sensors

Learn More at maurymw.com

maury

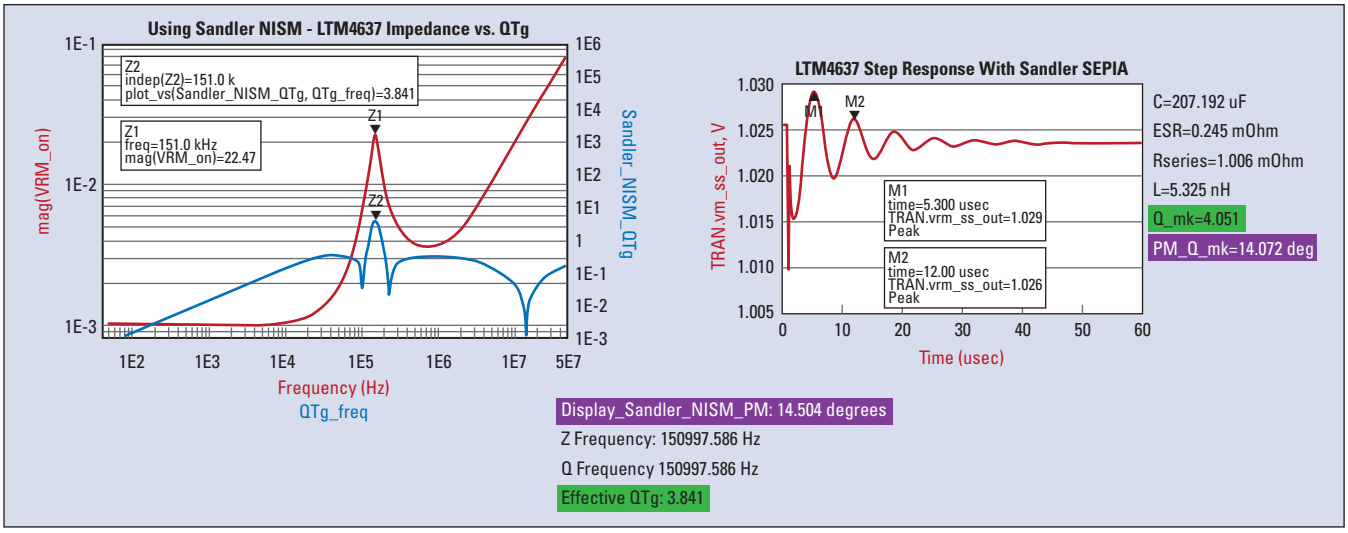


Fig. 6 LTM4637 stability simulation results — NISM (left) vs. SEPIA (right).

Key Takeaways for Power Integrity Success

- The Golden Rule: Always remember that noise follows impedance. To manage noise, you must manage your impedance profile.
- Design Strategy: Utilize SEPIA or NISM to aim for a flat impedance PDN. If a flat response is not achievable, ensure the system is designed so that Q < 2 to maintain stability.
- Comprehensive Modeling: Power integrity spans from the power supply to the load; achieving accuracy requires end-to-end modeling in ADS and validation through both large-signal and small-signal simulations.
- Efficiency: When paired with StepLoad Pro on a scope, SEPIA transforms complex characterization into a 60-second automated reality, offering a ubiquitous solution across the entire design cycle.

Ultimately, SEPIA is not just a measurement tool — it is an essential support system for the modern engineer. It provides the specific capacitor solutions and minimal adjustments needed to hit target impedances, ensuring your designs are robust, stable, and high performing from the lab bench to the final product.■

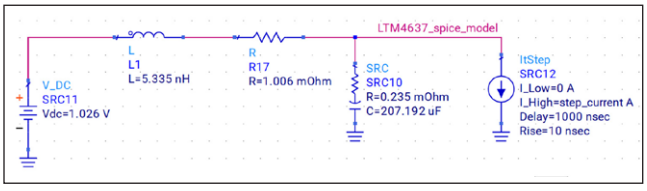


Fig. 7 Simulation setup for extracted SPICE model for LTM4637 PDN.

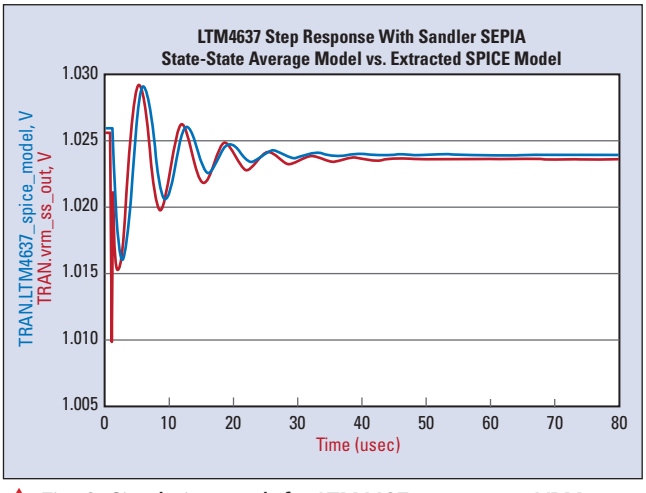


Fig. 8 Simulation result for LTM4637 state-space VRM model vs. extracted SPICE model for LTM4637 PDN using SEPIA.

REFERENCES

1. S. M. Sandler, "High Fidelity Power Supply Measurement," *APEC 2012*, Feb. 2012, Web: https://www.dropbox.com/scl/fi/ktz21ggfy5yejmdqds016/High-Fidelity-Power-Supply-Measurement-1011-Apec-2012.pptx?rlkey=kn9g5xgazf77bpzdhjmlaomkp&dl=0.
2. "Extracting Bode Plots from Output Impedance," *Picotest,* Web: https://www.picotest.com/wp-content/uploads/2024/05/Extracting_Bode-Plots_from_Output_Impedance_203PETSandler.pdf.
3. M. Nogawa, "Loop Analysis Directly from Time-Domain Waveform with SEPIA,"*Microwave Journal*, Vol. 67, No. 6, June 2024.
4. M. Nogawa, "ZOUT from Transient Response: The Link between Time Domain and Frequency Domain," *Microwave Journal*, Vol. 67, No. 4, April 2024.
5. "Non-Invasive Stability Measurement," *Picotest*, Web: https://www.picotest.com/non-invasive-stability-measurement.html.
6. B. Dannan, "Beyond the LTI Barrier: Quantitative VRM Stability and PDN Optimization from Time-Domain Step Load Measurement," *DesignCon 2026*.
7. S. Sandler, B. Dannan, H. Barnes, I. B. Ezra, and Y. Ni, "Design, Simulation, and Validation of a 2000-Amp Core Power Rail," *DesignCon 2024*.
8. S. Sandler, B. Dannan, H. Barnes, and C. Yots, "VRM Modeling and Stability Analysis for the Power Integrity Engineer," *DesignCon 2023*.
9. "SI/PI Model Library," Signal Edge Solutions, Web: https://www.signaledgesolutions.com/models.
10. S. M. Sandler and A.K. Davis, "Power Integrity Using ADS," *Faraday Press,* 2019.
11. "The Inductive Nature of Voltage-control Loops," EDN.
12. "Signal Integrity Power Integrity and Electromagnetic (EM) Modeling Publications," Signal Edge Solutions.
13. Keysight ADS, Web: https://www.keysight.com/us/en/products/software/pathwave-design-software/pathwave-advanced-design-system.html.
14. H. Barnes, "Optimizing Power Distribution Networks for Flat Impedance," *Signal Integrity Journal,* May 2020.
15. H. Barnes, "Power Integrity Fundamentals: Impedance vs. Frequency," *Signal Integrity Journal,* May 2021.
16. SEPIA ADS Example Workspace Download, Web: https://www.signaledgesolutions.com/downloads.
17. "StepLoad Pro VRM & Power Supply Test Automation Software," Signal Edge Solutions, Web: https://www.signaledgesolutions.com/product-page/stepload-pro-vrm-power-supply-test-automation-software.
18. "StepLoad Pro Product Bundle," Signal Edge Solutions, https://www.signaledgesolutions.com/product-page/stepload-pro-product-bundle.

ADVANCING OPTICAL CONNECTIVITY FOR AI-SCALE INFRASTRUCTURE

End-to-end optical, copper, and thermal innovations enabling next-generation AI data centers.

As AI workloads push bandwidth, density and power efficiency to new limits, system architects are moving optical interfaces closer to compute and switching silicon. At OFC 2026, TE Connectivity (TE) showcased innovations enabling scalable, high-performance architectures—from near chip connectivity to rack level optical infrastructure.

TE'S EXPANDING OPTICAL PORTFOLIO FOR AI DATA CENTERS INCLUDES:

High-Density Fiber Array Units (FAUs)

Enabled by TE's recent acquisition of RAM Photonics, these FAUs leverage advanced active alignment and automated optical fiber fusion splicing to deliver exceptional density, scalable manufacturing and lower insertion loss.

1.6T Linear Receive Optics (LRO)

High-performance connectivity supporting next-generation switch and accelerator architectures.

Co-Packaged Optics (CPO) Architectures

Enabling scalable near-chip optical connectivity for higher bandwidth and improved system efficiency.

End-to-End Optical Infrastructure

From near-chip optical interfaces to rack-level connectivity including optical backplanes, blind-mate connectors and high-density cabling.

DISCOVER TE'S INNOVATIONS FOR AI DATA CENTER CONNECTIVITY

te.com/AI

EVERY CONNECTION COUNTS

© 2026 TE Connectivity. All Rights Reserved.
TE Connectivity, TE, TE connectivity (logo) and EVERY CONNECTION COUNTS are trademarks owned or licensed by the TE Connectivity plc family of companies.

Evolution of the PDN Design Process: Explaining the Second Paradigm Shift

Istvan Novak
Samtec, Burlington, Mass.

Artificial intelligence (AI), together with machine learning (ML), created a huge surge of computing and networking infrastructure needs. This, in turn, has dramatically increased the power consumption of computing and networking chips. Traditional design and validation methods that served us well for decades are no longer able to cope with the new challenges: a paradigm shift is under way of how we design, optimize, and validate our systems and their power distribution network (PDN). In this article, we walk through the evolution of the PDN design and validation process and show what causes this second paradigm shift.

As a child hobbyist, I was introduced to electronics when vacuum-tube radios were still quite common. Small portable radios had already entered the market, with the first ones all using germanium transistors. The simpler units had the so-called "reflex" topology with only two transistors and no bypass capacitor on the supply rail.

Very early computer boards did not need any specific power distribution design. This is illustrated in ***Figure 1a***, which shows a diode-transistor logic card from computing equipment made in the 1970s. It was on a two-sided board, with no ground plane and no bypass capacitor. And many of us remember the prevailing rule of thumb from the 1980s: place a 0.1 uF ceramic capacitor across the power-ground pins of each logic chip, which, back then, usually was a dual-in-line (DIL) package (see Figure 1b).

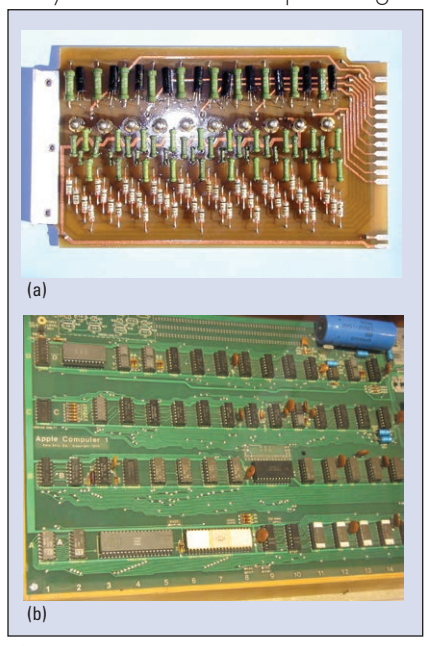


Fig. 1 Logic card from an old computer with no bypass capacitor (a), and Apple-1 motherboard with a few ceramic disc bypass capacitors (b).

First Paradigm Shift

During my career in professional electronics design, the first paradigm shift happened in the mid-to-late 1990s, when the central processing units (CPUs) in computers grew in complexity to the point that they required 10s of amperes of core current. The current transients of CPU core rails required a new approach: power ground plane pairs replaced the power-ground traces, the number and values of bypass capacitors increased, and a systematic frequency-domain design approach was developed. It was based on a calculated impedance target for the power distribution network.[1]

The impedance target in the 1990s was usually tens or hundreds of milliohms (mΩ), which immediately created a challenge for validation: instrument setups that were to measure impedance were not well suited to measure mΩ on working computer boards. To measure mΩ, the two-port

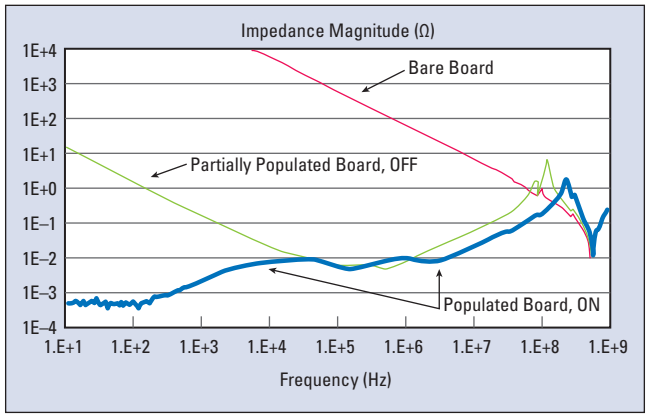


Fig. 2 PDN validation result on a computer board from the mid 2000s.

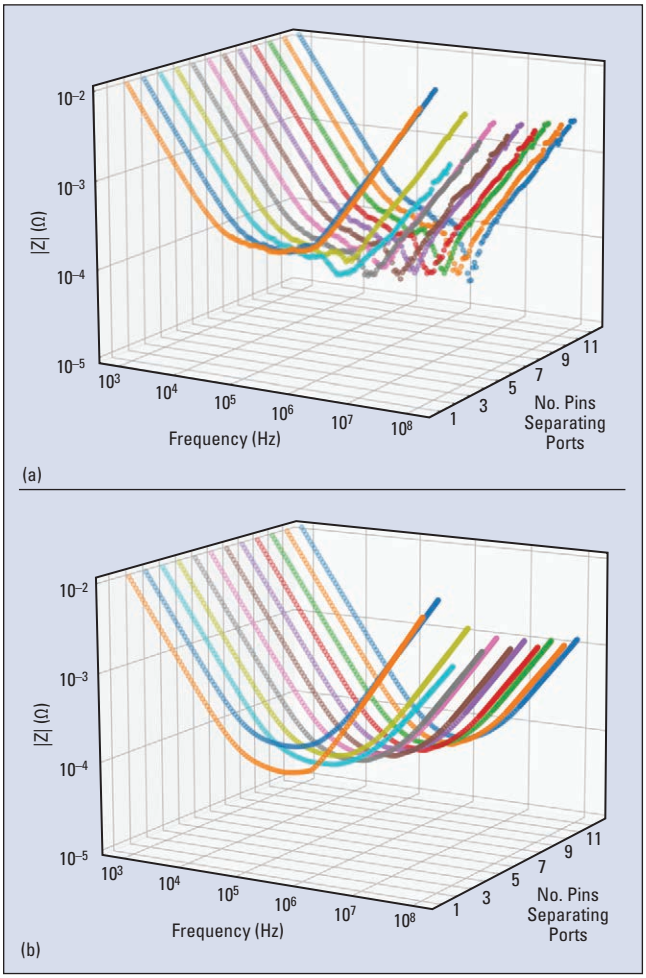


Fig. 3 Measured (top chart) impedance and simulated (bottom chart) as a function of pin location under a large computing chip (reproduced from Reference 4).

shunt-through measurement setup was developed in 1997, and it has been in use ever since.[2]

In multi-layer printed circuit boards (PCBs), using one or two plane pairs to serve as the low-impedance supply rails provided low enough series resistance that the horizontal power distribution lumped the entire PDN with tolerably low spatial variations. (As a reminder, the sheet resistance of a one-ounce copper plane is around 0.7 mΩ, which is just a small fraction of a 10 mΩ parallel target impedance.) This was also the time before the widespread use of digital power converters. As a result, such PDNs could be treated as linear and time invariant (LTI) networks, allowing us to use either the time or frequency domain to measure or simulate our circuits, because getting the other domain was straightforward and easy with Fourier transform. An impedance validation example is shown in ***Figure 2***.[3] Measuring PDN impedance at the mΩ level was readily doable with the two-port shunt-through scheme by probing the PDN at dedicated power-ground plated through hole pairs with probes connected to the vias from the top and bottom.

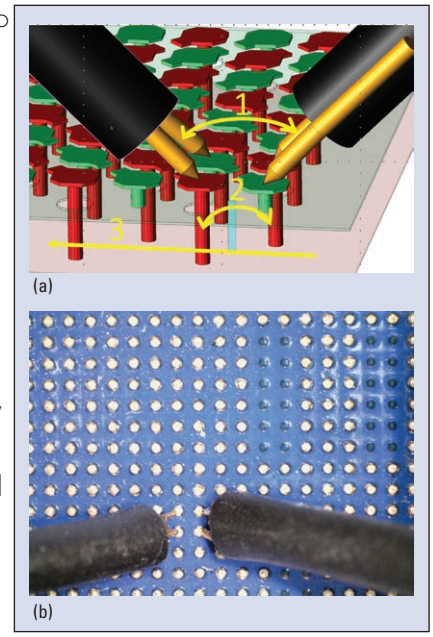


Fig. 4 The three dominant close-field 3D interactions are marked by yellow arrows and numbers: (1) probe-tip coupling, (2) via-barrel coupling, (3) finite plane impedance coupling (a). 1-mm PacketMicro PDN probes landed on a 1-mm via field (b).

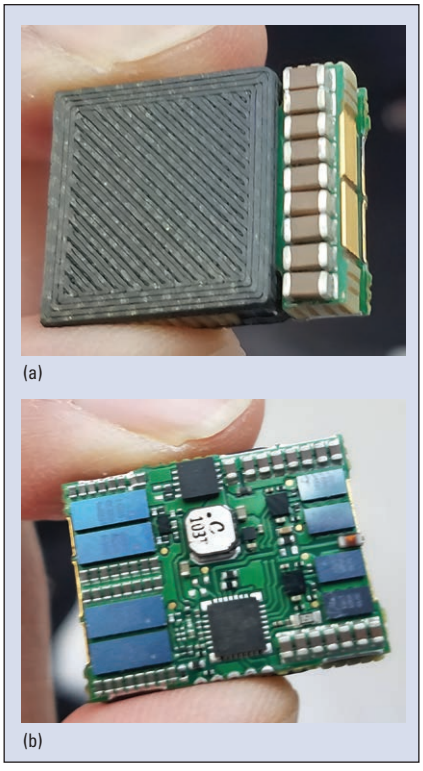
As the years passed, chips continued to become more power hungry: clock frequencies kept going up, supply voltage continued to drop, and the maximum current draw and current transients were also on the rise. Lower supply voltages require us to set tighter absolute noise budgets, which, when combined with rising current values, resulted in a steeply dropping impedance requirement.

Second Paradigm Shift

The boom of AI and ML accelerated the trend of power-hungry chips, so we have arrived at another paradigm shift. The largest chips now need kiloamperes and require tens of micro-ohms (μΩ) PDN impedance to feed them, creating significant new challenges.

Validating μΩ AC impedance is very complicated with available instrumentation. However, there are even bigger challenges. As illustrated in ***Figure 3***, PDNs cannot be considered lumped any more even at low frequencies, because the series plane impedance becomes comparable to or even bigger than the required parallel PDN impedance. This can result in strong spatial and/or frequency variation of impedance and noise, regardless of whether we look at it in the frequency or in the time domain. In turn, this complicates how we define, design for, and validate PDN requirements.

Very low impedance exposes 3D interactions between probes and the PDN that were previously masked by with higher PDN impedance.[5] The 3D rendering in ***Figure 4a*** (from Reference 4) shows the probing of the PDN in the footprint of a high current chip. PCBs that make use of blind vias to connect the surface pads to the power planes may



(a)

(b)

Fig. 5 Vicor high-density power converter circuit top and bottom views. ***Source:*** **Steve Sandler, Picotest.**

lack through-holes, so both probes for the two-port shunt-through testing must be placed on the pads on the same side of the board. Having wafer probes next to each other creates additional inductive loop coupling between the exposed probe tips, which may completely mask out the low DUT impedance we need to measure. Adding to the challenge, this probe tip coupling cannot be calibrated out with today's vector network analyzer (VNA) calibration methods. 3D rendering of a checker-board via field connecting the chip's footprint pads to the ground plane (green) and power plane (red) is shown in Figure 4a.

As if these challenges were not enough, the component density of our systems and subsystems puts potential noise sources and noise victims closer and closer together, giving rise to increased in-system interference. ***Figure 5*** (reproduced from Reference 6) shows a multi-kW very compact power converter. To reduce distribution loss through the connecting planes, we need to place the power converters close to the chips, where high-speed signals are also concentrated, further increasing the chances for in-system interference.

There are additional challenges as well: the nonlinear functions of digital power converters make PDN validation more complex because the easy conversion between time and frequency domain is valid only in the range of linear operation. The functionality under very large current transients has to be validated differently, such as by a custom fixture, capable of creating large current transients.[7]

It is still unclear which solutions the industry will settle on to solve these challenges. On the hardware side, an emerging possibility is using a vertical PDN, feeding the high-current supply rails of chips from below the main board while not relying on horizontal power planes in PCBs. However, vertical power distribution requires a significant redesign of the mechanical infrastructure of the system, and it carries its own challenges, including backward compatibility of racks and chassis elements as well as access to test points for validations. In addition, vertical power distribution does not eliminate the spatial variation within the large area required to carry kiloamperes of current. With the spatial variation, the question also remains: Where do we define the test points, and how does that relate to the overall noise specifications?

Evolution of PDN Design

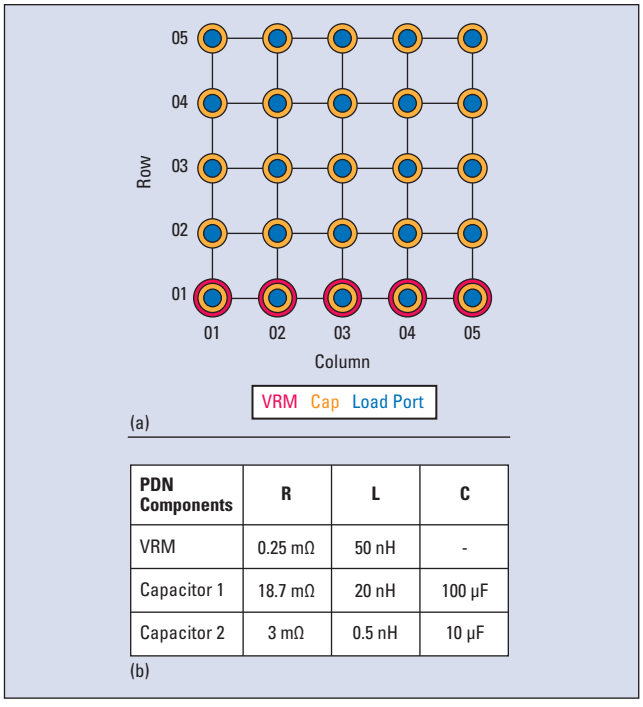
On-chip power distribution design processes have long focused on short-span transient time-domain noise simulations, but one of the earliest systematic off-chip PDN design processes (introduced in the early 1990s) is based on frequency-domain impedance targets. Due to the inherent lowpass filtering of the board-to-silicon path in off-chip PDN designs, the silicon must handle the high end of the power noise spectrum. Because of cost and technological constraints on the chips, lower-frequency bypassing has to be handled in the package and board design. Fortunately, the package and board have successively more space available, but the filtering in the PDN path means we cannot influence the high-frequency noise with components on the user board.

Off-chip impedance-based design and validation approaches focus on linear and time invariant (LTI) systems, and the impedance profile is typically simulated with frequency independent lumped RLC approximation of the off-chip power distribution. Using the frequency domain allows designers to decouple time domain excitation (in most cases highly statistical and able to represent only a few selected working conditions) from the rest of the PDN design and validation workflow.

Off-chip worst-case transient noise estimation, the Reverse Pulse Technique (RPT), was introduced in 2002.[8] This process is based on the step response of the PDN. It calculates the absolute worst-case transient noise of a minimal phase LTI system for any arbitrary sequence of current steps with a given transition time. In the early 2000s, the typical design had a combination of clock frequencies in the tens or hundreds of MHz, current consumptions in the 10s of amperes, and one or more 1 oz PCB power-ground planes in the PCB. These system parameters were sufficiently supported with PDN impedances in the tens of mΩ range and allowed designers to largely ignore nonlinearities in the power sources and chips as well as spatial variations on larger server boards.

When current consumption kept rising, DC drop calculations on power ground planes became important and had to be predicted and then validated. As the required DC and AC distribution impedance kept falling, spatial impedance variations became very noticeable. As a result, chip static and dynamic loading effects were added to simulation models. In larger systems, the RC delay caused by the R series plane resistance and the C bulk capacitance could cause enough phase shift between the output of a voltage regulator module (VRM) and its remotely located sense point that it can noticeably reduce the regulator's phase margin.[9]

Though the original target impedance approach and the RPT are still valid within the stated boundaries of their applications, applying them to demanding new applications faced multiple challenges: sometimes strong nonlinearities at the power source (voltage clipping), lack of including the silicon and its nonlinearities, and, most importantly, the lack of insight of the inter-relation between the spatially varying



PDN Components	R	L	C
VRM	0.25 mΩ	50 nH	-
Capacitor 1	18.7 mΩ	20 nH	100 µF
Capacitor 2	3 mΩ	0.5 nH	10 µF

(b)

Fig. 6 Node and component placement sketch (a), component values (b).

PDN transfer functions and spatially varying worst-case transient noise.

Addressing Today's Challenges

Multiple publications address different aspects of the challenges related to the second paradigm shift. On the system design side, changing from horizontal to vertical power distribution reduces, though does not eliminate, the increasing variability due to the horizontal PDN impedance. Moreover, unless the partitioning and component placement of the bottom-side power source module is optimized for vertical power distribution, the power module may inherit some of the challenges related to horizontal power distribution.

In current design methodologies, optimizing the response of the entire PDN, including the package and chip, offers better overall results.[10] However, designers must have access to sufficiently detailed information on all components, including the chip and its package; this is unfortunately not the case for many system and board designers.

Validating very low PDN impedance becomes easier if the test signal power is increased. For example, with two-port VNA measurements, boosting the source power to +20 dBm and using a low noise preamplifier, the equivalent-resistance noise floor of the measurement setup can be reduced to below 1 µΩ.[11] Although this addresses the challenge of measuring very low values of PDN impedance, characterizing the spatial variations within big chip footprints still requires a large number of measurements. The limitation stemming from the single-point connection of the VNA measurement setup can be removed by "area illumination," which has long been available from some of the CPU manufacturers as dedicated test gear. It illuminates the entire footprint (or an arbitrary portion of it) with a large-signal transient source with fast transient capability. Newer

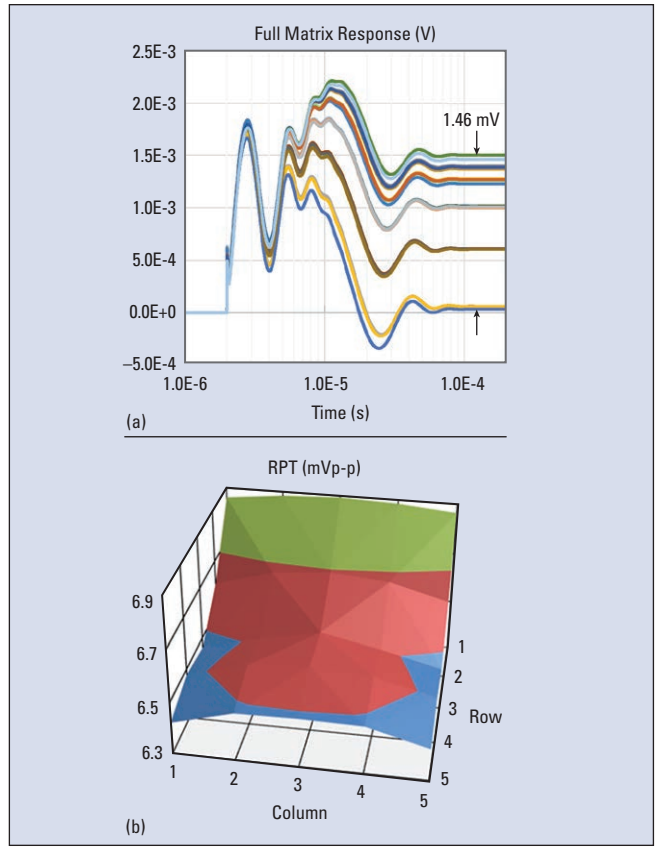


Fig. 7 Time-domain step response across all 25 nodes (a) and RPT transient-response results simulated on the circuit shown in Figure 6 (b).

implementations are described in References 12 and 13.

While with full-area excitation we can easily measure µΩ impedances, this is restricted to applications where a proper interposer can be made available that matches the exact footprint of the chip and mechanically and thermally fits into the system. Since the large-signal transient exercisers available today can only sink current, they can be used for testing active PDNs only. Moreover, they also do not answer the question about spatial distribution over the illuminated area.

Example PDN Analysis

To illustrate the potentially strong frequency and spatial variation of PDN noise, we take a simplified simulated case from Reference 14. We assume a hypothetical 4 × 4 in. power-ground plane pair with connection/observation nodes on a 1 in. grid. The sketch in ***Figure 6*** shows the geometry, node numbering, and component placement. The frequency and step responses are then simulated with SPICE. Time-domain results are shown in ***Figure 7***.

In this simple illustration (Figures 6 and 7), circuit values are intentionally exaggerated to clearly show the effect of spatial and frequency variations. Each square inch of the plane pair is represented by a SPICE equivalent circuit with 2.6 mΩ resistance, 142 pH inductance, and 184 pF capacitance. There are two capacitor banks at every node with the cumulative parameters shown in the table. The VRM phases are on the nodes in the first row.

The excitation itself has a uniform distribution of a 1A

eLearning
Signal Integrity • Power Integrity • EMC/EMI
Join Us in 2026...
June 24
SI/PI Summit
August 19
High Speed Interconnect Technologies Panel
September 23
AI in SI/PI Design Summit
October 14
AI in Design Panel
Organized by SIJ
www.signalintegrityjournal.com

total value, with equal current sharing and zero skew from node to node. The transient source is a piece-wise-linear step waveform, shaped by a fifth-order Butterworth filter to control the excitation bandwidth, which is set to a 100 MHz cut-off frequency. As explained in Reference 14, when we do AC simulations, the frequency response plot is called AC noise with voltage unit. Even though the total excitation current is 1 A, we don't call this impedance because the current is distributed across all nodes. The transient noise is calculated with the RPT which is, strictly speaking, applicable to single-node lumped PDNs. For this reason, we don't call this worst-case noise; we just calculate the peak-to-peak noise with the RPT process.

Summary

With the rapid drop of PDN impedance that is required to feed our biggest chips, design and validation processes have multiple new challenges. The main characteristics of the second paradigm shift is that spatial variance and 3D interactions of close-by PDN features must be considered. These challenges call for new design methodologies and validation processes. End-to-end design flow including VRM and silicon load, vertical power distribution, large-signal transient array loads, and VNA measurements with enhanced accessories are some of the key developments that can help. ■

Acknowledgement

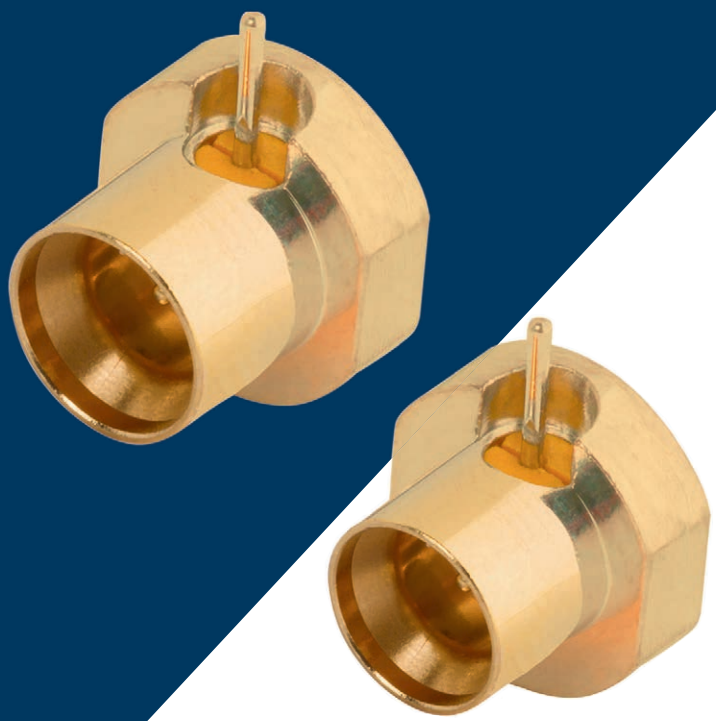
The author wishes to express his sincere thanks and appreciation to all of his past and current colleagues who shaped his understanding of electronics. A special thanks goes to Janine Love, Samtec, for her expert editing support.

REFERENCES

1. L. Smith, et. al, "Power Distribution System Design Methodology and Capacitor Selection," *IEEE Trans. Adv. Packag.*, Vol. 22, No. 3, Aug. 1999, pp. 284–291.
2. I. Novak, "Why 2-Port Low-Impedance Measurements Still Matter," Web: http://www.electrical-integrity.com/Quietpower_files/QuietPower-53.pdf.
3. I. Novak, "A New Dawn in R&D," *Microwave Magazine*, Vol. 12, No. 5, 2011.
4. J. Phillips, et al., "Determining the Requirements, Die vs. Package vs. Board: Multi-level Power Distribution Network Design," *DesignCon 2025*.
5. "3D Connection Artifacts in Modern PDN Measurements," *CadenceLive Boston*, September 12, 2023.
6. I. Novak, "The Future of Power Integrity through the Eyes of Experience," *Design Con 2023*, Web: http://www.electrical-integrity.com/Paper_download_files/DC23_TheFutureOfPowerIntegrity_final.pdf.
7. S. Sandler, et al., "Design, Simulation, and Validation Challenges of a Scalable 2000 Amp Core Power Rail," *DesignCon 2024*.
8. V. Drabkin, C. Houghton, I. Kantorovich, and M. Tsuk, "Aperiodic Resonant Excitation of Microprocessor Power Distribution Systems and the Reverse Pulse Technique," *IEEE 11th Topical Meeting on Electrical Performance of Electronic Packaging*, 2002.
9. J. Hartman, et al., "Impact of Regulator Sense-point Location on PDN Response," *DesignCon 2015*.
10. I. Zamek, "Reducing Chips Respin Rate by Fixing Common Decoupling Methods Flaws," *DesignCon 2025*.
11. K. Willis, "Removing Communication Barriers Between CAD and Instrumentation Companies with Open Source PCB," *IEEE EMC-SIPI Symposium*, 2025.
12. E. Koether and I. Novak, "Transient Load Tester for Time Domain PDN Validation," *EDI CON 2017*.
13. S. Sandler, B. Dannan, H. Barnes, I. Ben Ezara, and Y. Ni, "Design, Simulation, and Validation Challenges of a Scalable 2000 Amp Core Power Rail," *DesignCon 2024*.
14. E. Koether, et al., "Bridging the Time-Frequency Chasm in PDN Design: Leveraging Cumulative Power-Rail Noise and Reverse Pulse Techniques for Spatial-Frequency Insight," *DesignCon 2026*.

REVERSE MOUNT THRU-HOLE RF PCB CONNECTORS

Offers Superior Signal Integrity by Eliminating via Transitions. Ideal for High-Frequency Designs.



Features & Benefits

- Direct signal path reduces insertion loss and reflections
- Space-saving, low-profile mechanical integration
- Eliminates cable from the top of the PCB
- Clean, predictable impedance transitions
- SMP (26.5 GHz), SMPM (40 GHz)

Applications

- RADAR and Satellite Communications
- Compact modules
- mmWave systems
- 5G

Amphenol SV MICROWAVE

Learn More

www.svmicrowave.com

When Frequency Shapes Time: How S-Parameter Properties Shape Simulated TDR Behavior

Tim Wang Lee, Ph.D., Keysight Technologies
Santa Rosa, Calif.

As modern high-speed designs push into multi-gigabit and higher data rates, engineers rely on multiple signal integrity (SI) metrics to evaluate interconnect performance. Among these, time-domain reflectometry (TDR) remains one of the most direct methods for identifying controlled impedance and discontinuities.

A TDR plot shows the impedance profile the signal sees as it travels down the interconnect. ***Figure 1*** shows an example of the impedance of a differential pair, and the impedance a signal sees as it travels from the pins of the chip outward.

The impedance of the differential pair is initially higher due to the wider separation of the two traces. The TDR presents a direct method for identifying the impedance and length of the 85 Ω interconnect.

When Two TDRs Do Not Match

There are two primary approaches for obtaining a TDR response:

1. Measurement in the time domain using a sampling scope
2. Simulation from measured frequency-domain S-parameters.

In principle, both methods should yield identical impedance profiles. In practice, differences in S-parameter quality, bandwidth, and simulation setup can lead to noticeable discrepancies, as shown in ***Figure 2***. The poor quality of the S-parameter data and improper simulation parameters caused the two traces to differ.

This article first establishes a structured setup for converting measured frequency-domain data into a time-domain TDR response. Using an appropriate setup, various S-parameter properties and their impact on the TDR results are examined. Finally, it is concluded with a practical methodology that facilitates a reliable frequency-to-time TDR conversion process.

This article presents a structured

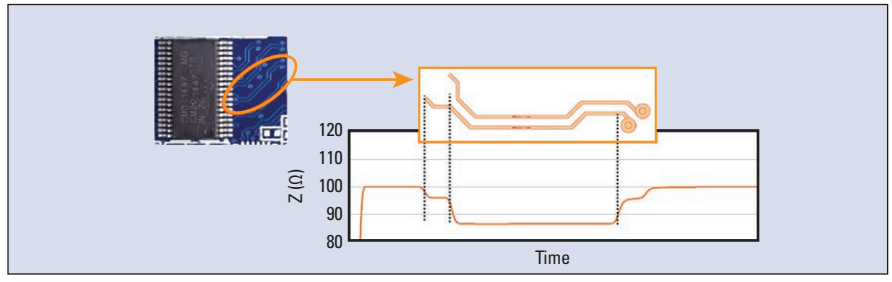


Fig. 1 An illustration of a differential pair and its corresponding TDR plot. The picture of the board is from an open-source board.[1]

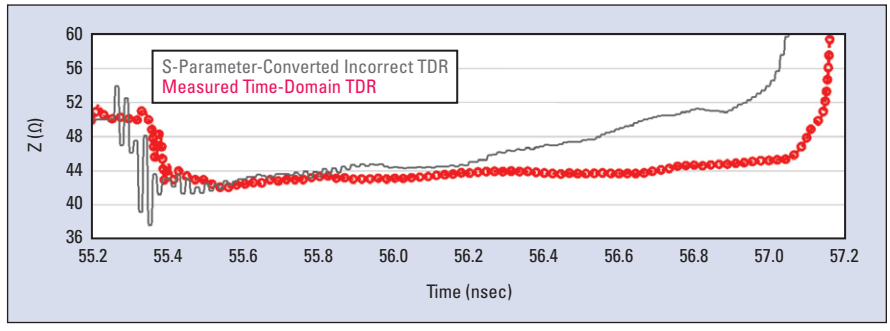


Fig. 2 The converted TDR (gray) appears different from the measured, time-domain TDR (red).

workflow for correlating time- and frequency-domain TDR analysis. Using this methodology, a simulated TDR generated from measured S-parameter data correlates with time-domain measurement within ±2.6%, as shown in ***Figure 3***. The high quality of the S-parameter data and the appropriate transient simulation setup enable a strong correlation between the two. Correlation was evaluated using the maximum impedance deviation between the simulated and the measured trace.

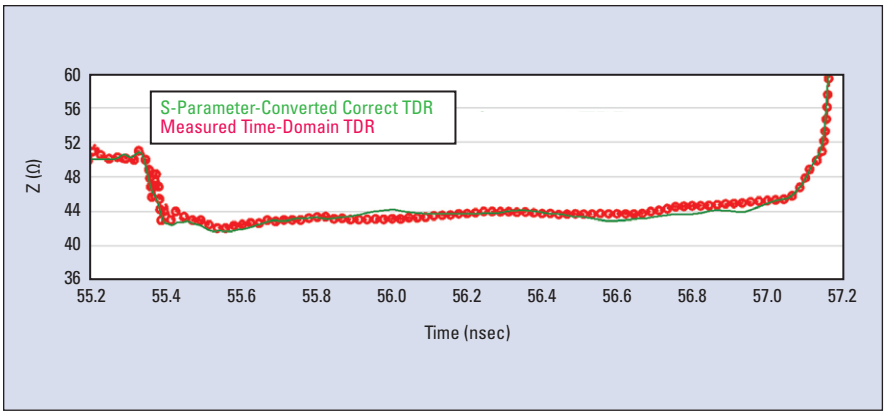


Fig. 3 The converted TDR (green) is within ±2.6% of the measured time-domain TDR (red).

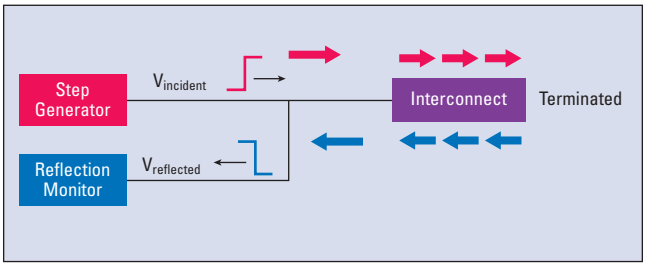


Fig. 4 An illustration of the step edge generation and the reflection monitoring in a TDR measurement instrument.

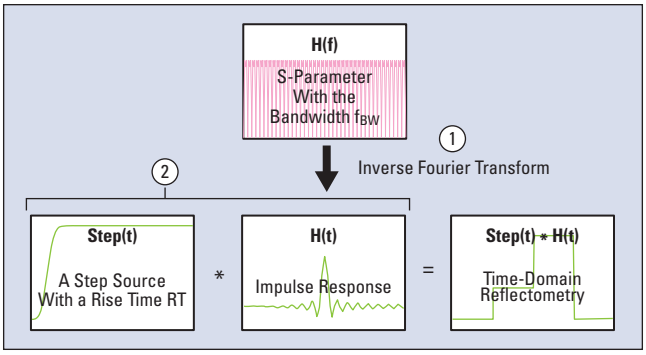


Fig. 5 An illustration of the general process to produce a TDR with a given S-parameter data.

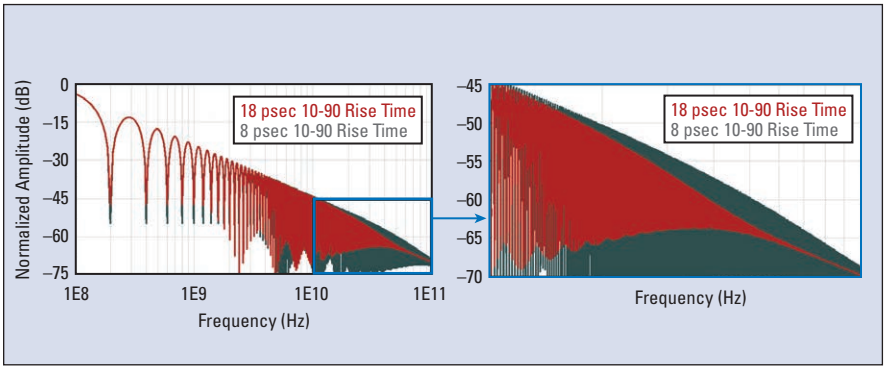


Fig. 6 The simulated and normalized spectrum of two steps with different rise times, 8 psec and 18 psec.

Bridging Two Domains: From Frequency to Time

Whether using a sampling scope or a vector network analyzer (VNA), the objective remains the same: to reveal the impedance profile of the interconnect. The sampling scope provides a direct time-domain snapshot, while the VNA captures frequency-domain behavior that can be converted to time using a transient simulation.

In a typical TDR measurement with a sampling scope, as shown in ***Figure 4***, a voltage source with a 10-90 rise time launches an incident step into the device under test (DUT). Discontinuities along the interconnect reflect part of this signal, and the reflection monitor at the source records the resulting voltage. From these voltages, the reflection coefficient and corresponding impedance profile are calculated.[2]

Not every SI lab is equipped with a TDR instrument, but most have a VNA. The VNA's measured S-parameters can reproduce the same insight by applying an inverse Fourier transform to obtain the impulse response. During a transient simulation, this impulse response is convolved with the source step waveform to generate the TDR response, as shown in ***Figure 5***.

Shorter Rise Time Translates to a Higher Bandwidth

In a TDR measurement setup with a step edge, the rise time is an important parameter. With a shorter rise time, the TDR can resolve smaller physical features that introduce impedance discontinuities. A shorter rise time corresponds directly to a broader frequency spectrum and therefore higher effective bandwidth.

According to the specification of the TDR instrument,[3] an 18 psec 10-90 rise time corresponds to approximately 35 GHz of bandwidth, while a step edge with an 8 psec rise time has a bandwidth of 50 GHz. The simulated spectrum comparison of the two different rise times in ***Figure 6*** confirms that a shorter rise time corresponds to a wider bandwidth.

Figure 6 further illustrates that, in a transient simulation, the step rise time must be constrained by the available S-parameter bandwidth. If the step source has a rise time corresponding to a 50 GHz bandwidth, S-parameter data with a 40 GHz bandwidth would cause the simulation result to have an unexpected simulation artifact. The graph shows that there is more high-frequency content when the step has a shorter rise time. As a result, the step edge with shorter rise time has a higher bandwidth.

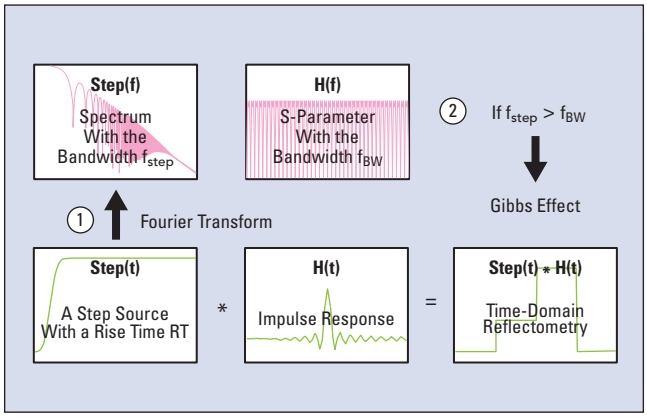


Fig. 7 When the corresponding spectrum of the step source has a higher bandwidth than the S-parameter data, the resulting TDR is expected to have the Gibb effect, a frequency-to-time domain conversion ringing artifact.

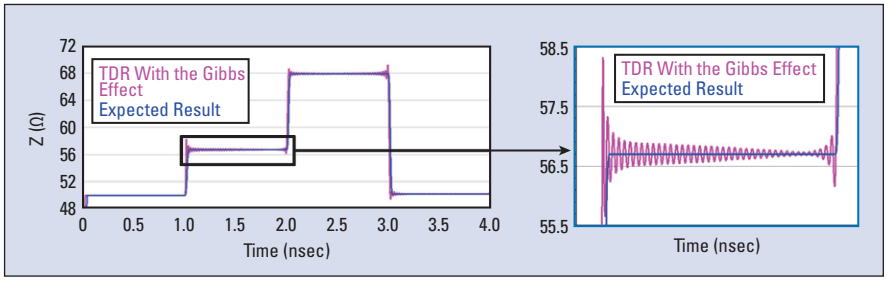


Fig. 8 A comparison between an expected TDR result and one with the Gibbs effect.

The Hidden Influencers: Bandwidth, Step Size, Causality, and Passivity

Because the measured S-parameter data is the focus of the transient simulation, careful examination of the properties of the S-parameter is critical in producing quality TDR plots.

Simulated TDR impedance correctness is fundamentally limited by the integrity of its input S-parameters. The following properties critically influence the outcome:

- Bandwidth: determines the smallest impedance features the TDR can resolve
- Frequency step: affects the impedance accuracy
- Causality: ensures correct timing of the TDR
- Passivity: ensures no energy generation; violations yield higher impedance
- Reciprocity: less critical for TDR since reflections, not transmissions, dominate the measurement.

Table 1 consolidates the S-parameter properties, their resulting time-domain artifacts, and the practical mitigation strategies required for accurate frequency-to-time conversion.

Frequency Bandwidth Determines the Shortest Rise Time

Assume there is measured S-parameter data with a 40 GHz bandwidth. This data is passive, causal, and reciprocal. If one simulates this measured data with a 40 GHz bandwidth with a step that has a higher bandwidth (short rise time), one expects the Gibbs phenomenon/Gibbs effect to be present. ***Figure 7*** gives a pictorial representation.

The difference between the source bandwidth and the bandwidth of the measured data causes the Gibbs effect. Although the Gibbs effect originally describes the ringing when one approximates a discontinuous function with jumps (a square wave) with a finite Fourier series,[4] the term has now become a descriptor for the ringing artifact when working with the frequency and time domains.

The Gibbs effect shows up in the TDR plot as ringing in the impedance plot, especially when there are impedance discontinuities, as shown in ***Figure 8***. The difference between the higher source bandwidth and the bandwidth of the data causes the Gibbs phenomenon, which appears as ringing.

To avoid the Gibbs effect in the TDR result, the step rise time must be short enough to resolve small features, yet not so short that it exceeds the S-parameter bandwidth. ***Figure 9*** shows the TDR result after simulating a 40 GHz bandwidth S-parameter with a step source that has a correct rise time configuration. A shorter rise time produces better TDR resolu-

TABLE 1

MAPPING BETWEEN S-PARAMETER PROPERTY VIOLATIONS AND THEIR OBSERVABLE EFFECTS IN THE CONVERTED TDR RESPONSE

S-parameter Property	*Time-Domain Symptom*	*Engineering Interpretation*	*Corrective Action*
Frequency step too large	Incorrect impedance result	Poor approximation of frequency response	≤ 10 MHz step for typical high-speed links
Non-causal data	Early arrival, pre-ringing	Violates physical timing constraint	Ensure causality in the calibrated measurement
Non-passive data	Artificially higher impedance result	Artificial energy generation	Ensure passivity in the calibrated measurement
Non-reciprocal data	Minimal impact on TDR	TDR dominated by reflection	Usually negligible

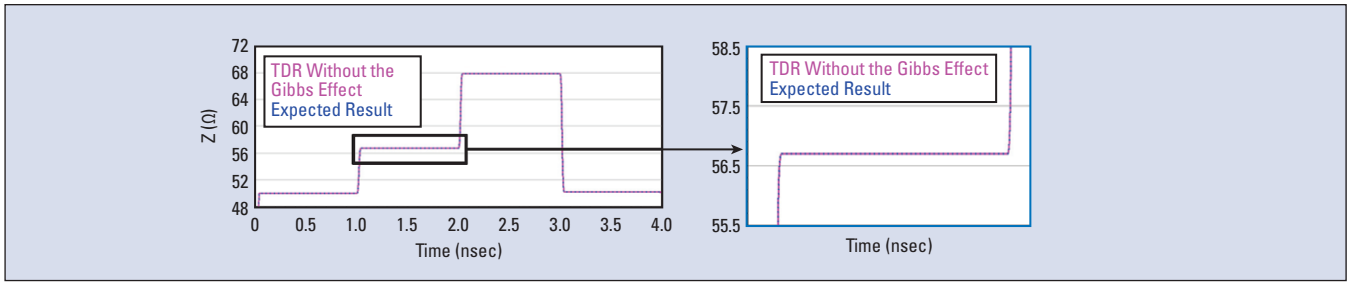


Fig. 9 By choosing the optimal rise time, one achieves the best impedance and length resolution while avoiding the Gibbs effect.

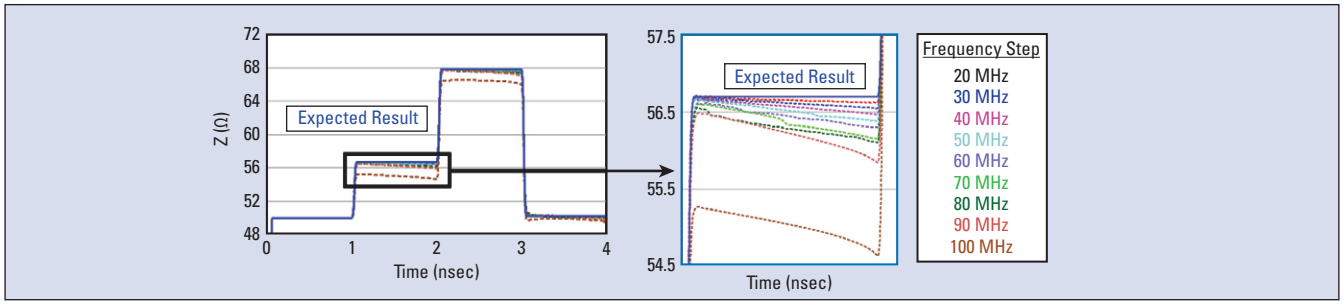


Fig. 10 If the frequency step becomes sufficiently large, the S-parameter data no longer correctly represents the device under test.

tion, but it cannot be so short that the step source bandwidth exceeds the bandwidth of the S-parameter data.

In practice, begin by setting the rise time according to the equation,

$RT_{opt}=k/f_{BW}$,

where RT_{opt} is the optimal rise time, f_{BW} is the S-parameter bandwidth, and k is the constant that provides the shortest rise time without producing the Gibbs effect. The constant, k, depends on the transient solver implementation and numerical windowing behavior. Verification that the resulting TDR waveform is free from artificial ringing is important.

In addition to the rise time, when defining the step edge for simulation, it is important to use an edge shape that resembles the actual TDR instrument. A real instrument's step edge follows a smooth transition rather than a linear slope. If possible, one should avoid using a linear edge to approximate the transition, as it can introduce unrealistic high-frequency components.

Fine Frequency Step Produces the Correct TDR Representation

The S-parameter data represents the frequency response of the DUT. The theoretically correct representation is one with infinite bandwidth and with an infinitesimal frequency step. In practice, S-parameter measurement produces an approximation of the frequency response because the bandwidth of the data is finite, and so is the frequency step.

Because measured S-parameters are discrete approximations of a continuous frequency response, one expects that as the frequency step becomes sufficiently large, the approximated frequency response of the DUT no longer accurately reflects its behavior. For the same reason, one expects this wrong approximation to also appear in the TDR response (see ***Figure 10***). This incorrect representation in the frequency domain translates to an incorrect TDR result in the time domain. For typical high-speed interconnects, a frequency step of 10 MHz is considered acceptable when capturing S-parameter data.

Non-Causality Introduces Incorrect TDR Impedance and Timing

A causal system is a system whose output at any given time depends only on the present and past inputs, not on future inputs. In a causal system, the effect cannot precede the cause. Interconnects are causal systems, and the S-parameter data representing them should also be expected to be causal.

TDR simulation using non-causal S-parameter data will exhibit incorrect timing behavior. ***Figure 11*** demonstrates the timing impact of the non-causality in the S-parameter data. The non-causality also leads to erroneous impedance

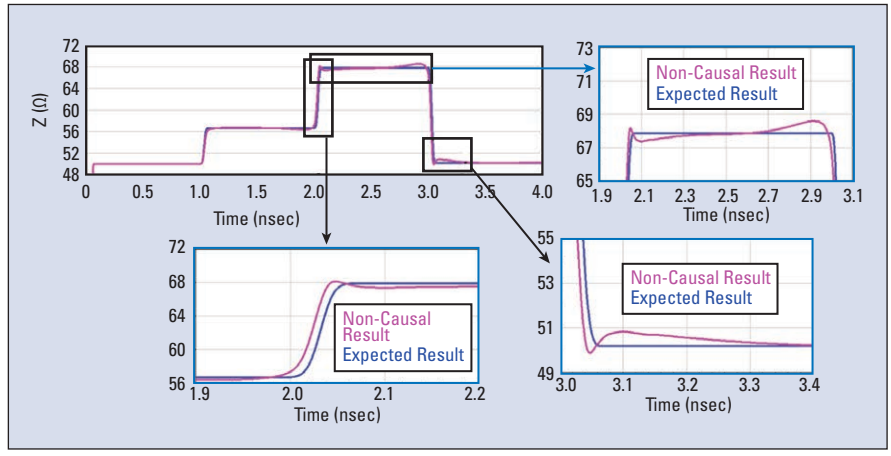


Fig. 11 A transient TDR simulation indicates that non-causal S-parameter data not only results in early timing, but it can also produce erroneous impedance dips and bumps in the TDR.

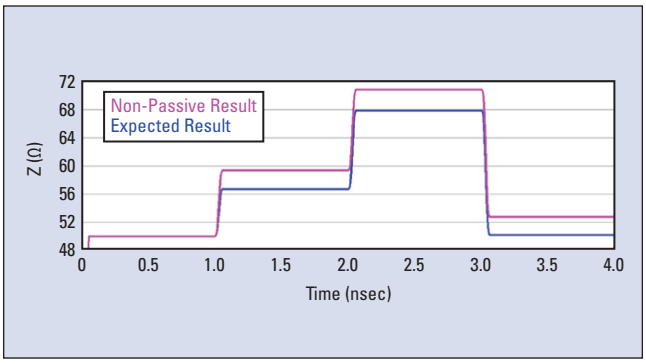


Fig. 12 A transient TDR simulation confirms that non-passive S-parameter data results in an erroneous impedance increase in the TDR.

dips and bumps.

Non-Passivity Creates Higher Impedance Readings

When a system is passive, it only consumes energy and does not produce any energy. Interconnects are passive and should not produce energy. If the S-parameter data violates passivity, the resulting TDR will exhibit artificially higher impedance due to non-physical energy gain.

Figure 12 confirms the higher impedance when non-passive S-parameter data was used in a TDR simulation. The non-passivity increases the impedance of the TDR.

TDR is Not Sensitive to Non-Reciprocity

A reciprocal system is one where the forward transmission and the reverse transmission are identical. Interconnects are generally reciprocal. In TDR analysis, reciprocity plays a limited role because the response is dominated by reflections rather than transmission paths.

Because the focus of TDR is on the reflected voltages, rather than the transmitted ones, the TDR result should not be sensitive to whether the forward and reverse transmissions match each other. ***Figure 13*** confirms that the TDR simulation with non-reciprocal S-parameter data remains consistent with the expected results.

Shaping Time with Care: Best Practices for Reliable Conversion

Producing a reliable TDR from frequency-domain data requires attention to both the integrity of S-parameter data and the simulation setup. The following practices enable reliable frequency-to-time conversion:

1. **Validate S-parameter integrity**
   Confirm fine frequency steps, adequate bandwidth, and verified passivity and causality. These ensure that the S-parameter model is physically accurate and minimizes simulation artifacts.
2. **Use an appropriate step edge**
   Apply a smooth step that reflects the real TDR instrument's transition shape. Avoid using a linear ramp, which introduces unrealistic high-frequency components.
3. **Match the rise time to the available bandwidth**
   Select a realistic rise time for the step source based on the S-parameter bandwidth. A practical starting point is $RT_{opt} = 3/f_{BW}$, followed by waveform inspection to confirm the absence of Gibbs-related ringing.

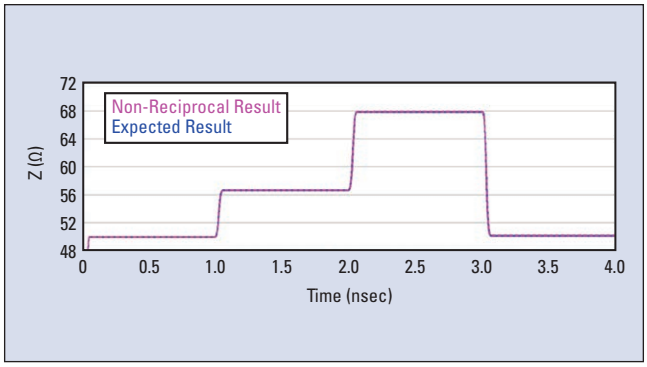


Fig. 13 A transient TDR simulation confirms that the TDR result is not sensitive to the reciprocity of S-parameter data.

By following these practices, engineers can confidently generate simulated TDRs that align with measured results, gaining insight into interconnect impedance without needing a dedicated TDR instrument.

Conclusion: Seeing Time Clearly

Whether using a sampling scope or a VNA to obtain the TDR plot, the goal remains the same: to reveal the impedance profile of the interconnect and evaluate interconnect SI. The sampling scope provides a direct time-domain view, while the VNA provides the frequency-domain data that, when handled carefully, can reconstruct the same information through simulation.

Simulated TDR correctness is fundamentally limited by S-parameter integrity. Each property — bandwidth, frequency step, causality, and passivity — shapes the impedance profile of the interconnect. When these factors are properly controlled, the TDR becomes a powerful diagnostic tool rather than a source of confusion.

TDR, however, is one of many tools in the SI toolbox. Consistency checks with S-parameters, such as comparing return loss and insertion loss, should always accompany TDR analysis to ensure a holistic signal integrity analysis is achieved.

Understanding and validating S-parameter fidelity enables engineers to correctly produce TDRs that correlate well with measurements. Applying these principles transforms TDR correlation from a trial-and-error exercise into a controlled engineering methodology.■

Acknowledgements

The author would like to thank several industry engineers for insightful discussions that helped shape this work.

REFERENCES

1. OpenRex Project, "OpenRex – Open Source Hardware Project," *iMX6 Rex Projects*, Web: https://www.imx6rex.com/open-rex/.
2. Keysight Technologies, *Time Domain Reflectometry Theory*, Application Note 5966-4855E, 2004.
3. Keysight Technologies, "N1055A 2-/4-Port TDR/TDT Remote Sampling Head," *Product Datasheet*, Web: https://www.keysight.com/us/en/product/N1055A/2-4-port-tdr-tdt-remote-sampling-head.html.
4. R. Baraniuk, et al., "Gibbs Phenomena," *Signals and Systems*, LibreTexts Engineering, 2022. Web: https://eng.libretexts.org/Bookshelves/Electrical_Engineering/Signal_Processing_and_Modeling/Signals_and_Systems_(Baraniuk_et_al.)/06%3A_Continuous_Time_Fourier_Series_(CTFS)/6.07%3A_Gibbs_Phenomena.

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

© 2025 Anritsu Company

Anritsu
Advancing beyond

PAM2 vs. PAM4 Signaling: Simply Explained

Bert Simonovich
LAMSIM Enterprises Inc., Ottawa, Canada

To keep up with the bandwidth demands of the industry, non-return-to-zero (NRZ) signaling has been replaced by pulse amplitude modulation level-4 (PAM4) signaling. To maintain the same naming convention, NRZ signaling is now referred to as PAM2. Both terms are acceptable.

PAM2 has been widely used in various serial communication interfaces over the years due to its simplicity, robustness, and lower cost to implement. The signal level remains constant during the bit interval, which helps in maintaining signal integrity with good margin. For decades, it has been the workhorse for standards like PCIe and Ethernet 802.3, but as bit rates surpassed 32 Gbps, PAM4 signaling has become the standard.

PAM4 signaling is used in high-speed data communication to effectively double the serial data rate without requiring an increase in channel bandwidth. It is commonly used in high-speed standards, such as 100 G, 200 G, and 400 G Ethernet. It has also been adopted for PCIe Gen 6 and future Gen 7 standard. But unlike PAM2, PAM4 uses four distinct voltage levels to represent data.

As shown in ***Figure 1a***, each bit time, or unit interval (UI) of a PAM2 data channel, contains one symbol, which can represent either a logic zero or a logic one. In this case, the bit rate is equal to the baud rate. In PAM4 signaling, one symbol from each of the PAM2 data channels is encoded into one bit stream using the Gray coding algorithm. The result is four voltage levels, with two symbols transmitted per UI, effectively doubling the bit rate for the same baud rate. For instance, if the bit rate is 56 Gbps, the baud rate would be 28 GBaud.

The PAM4 eye diagram in ***Figure 1b*** shows three distinct eye openings, with one-third the voltage amplitude VA, compared to the single eye opening in PAM2. These levels are binary coded as 00, 01, 10, and 11. This difference in eye opening, however, leads to a 9.6 dB reduction in signal to noise ratio. This means it is more susceptible to noise, requiring more advanced signal processing and error correction. Fortunately, forward error correction coding is used to compensate.

Throughout the course of my career, I have coined the phrase, "Never say never with copper. Silicon always finds a way!" The shift from NRZ (PAM2) to PAM4 signaling has enabled higher data rates while maintaining bandwidth efficiency, despite its increased complexity and noise sensitivity. As insatiable data needs continue to grow due to AI demands, the adoption of PAM4 standard has paved the way for future PAM-level encoding schemes to address future bandwidth limitations, prolonging the life of copper interconnect. ■

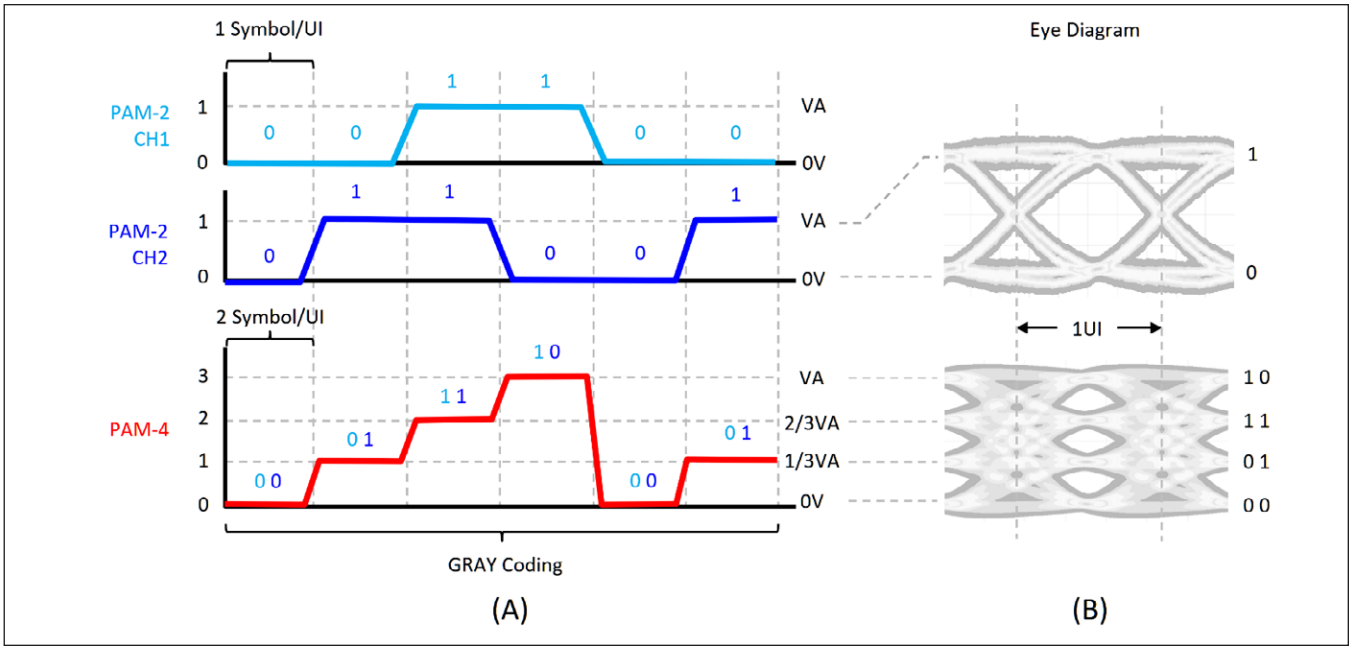


Fig. 1 Comparison of (a) NRZ (PAM-2) encoding versus PAM-4 encoding and (b) respective eye diagrams.

CHINA INTERNATIONAL
OPTOELECTRONIC EXPO

CIOE 2026

The 27th China International Optoelectronic Exposition

Enabling the Next Generation of High-Speed Interconnects



September 9-11, 2026

Shenzhen World Exhibition & Convention Center

Shenzhen China

Exhibition Areas	Exhibitors	Visitors
240,000 ㎡	3,800 +	150,000 +

8 Expos in One Show

- INFORMATION AND COMMUNICATION EXPO
- PRECISION OPTICS EXPO & CAMERA EXPO
- LASERS TECHNOLOGY & INTELLIGENT MANUFACTURING EXPO
- AR&VR EXPO
- INTELLIGENT SENSING EXPO
- INFRARED APPLICATIONS EXPO
- PHOTONICS INNOVATION EXPO
- DISPLAY TECHNOLOGY EXPO

Co-located with | IICIE INTERNATIONAL INTEGRATED CIRCUIT INNOVATION EXPO | elexcon

WWW.CIOE.CN/EN

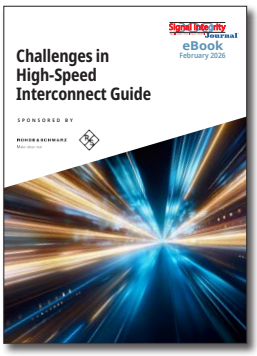
Booth Application

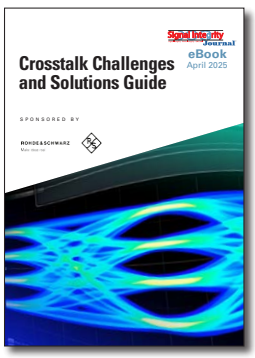
Register Now

AD INDEX

Anritsu 9, 39
CIOE 2026 41
Coilcraft 11
COMSOL, Inc. ..... COV 4
HYPERLABS INC. 19
IEEE EMC + SIPI 2026 ..... COV 3
Maury Microwave Corporation 25
Rohde & Schwarz USA, Inc. 3
Samtec USA ..... COV 2
Siglent Technologies 15
Signal Edge Solutions 21
Signal Integrity Journal 20, 32
SV Microwave, Inc. 33
TE Connectivity 27
Teledyne Lecroy 5

E-BOOKS







SIJ

The information you need, from industry experts.

signalintegrityjournal.com

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
Scott Blackburn
Western Reg. Sales Mgr.
(CA, AZ, OR, WA, ID, NV, UT, NM, CO, WY, MT, ND, SD, NE & Western Canada)
Cell: (303) 819-3958
sblackburn@mwjournal.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: (781) 769-9750

AUGUST 3-7

EMC+SIPI 2026

DALLAS, TEXAS

2026 IEEE INTERNATIONAL SYMPOSIUM ON ELECTROMAGNETIC COMPATIBILITY, SIGNAL & POWER INTEGRITY



3RD GLOBAL SIGNAL INTEGRITY AND POWER INTEGRITY (SIPI) UNIVERSITY

TUESDAY-THURSDAY, AUGUST 4-6, 2026

The Global SIPI University offers two full days of lectures carefully curated to cover basic and advanced concepts of SI & PI during the symposium week. SI & PI are gaining ever-growing attention due to the higher data rates and larger currents in modern high-speed digital systems. Industry requires skilled engineers with a background in these two disciplines to address the increasing complexity and challenges of electronic system design.

The program offers both lectures and practical demonstrations. The intent is to expand the discussion on the background concepts and to provide a more detailed presentation of advanced methods for addressing the current design challenges. The demos will encompass the use of instruments typically employed for analysis, measurement and troubleshooting for signal and power integrity.



CO-CHAIR
Francesco de Paulis
Univ. L'Aquila, L'Aquila, Italy

CO-CHAIR
Majid Ahadi Dolatsara
Keysight Technologies



TOPICS

OPENING SESSION:	Overview of High-Speed Signal Integrity (SI)
INTRODUCTION TO SI AND PI:	Evolution from the Basics to the Current Technology
SIGNAL INTEGRITY I:	Transmission Line Effects, Lumped Effects, Passive Interconnect Design, Terminations and Reflections
SIGNAL INTEGRITY II:	How Interconnects Work: The Physics of Losses, Reflections, and Coupling
SIGNAL INTEGRITY III:	S-params In the Eyes of the Beholder: Interconnect Design Considerations at 224 Gb/s and 448 Gb/s
SIGNAL INTEGRITY IV:	Measurements for Signal Integrity (VNA and TDR)
KEYNOTE:	Powering AI Performance from Silicon to System: The Role of SIPI
SIGNAL INTEGRITY DEMO I:	Time Domain – Scope Measurements of the Impact of Transmission Line Effects and Reflections Demonstrating Ringing Noise, Time of Flight and Eye Diagrams
SIGNAL INTEGRITY V:	Simulation and Modeling of Signal Integrity for High-Speed Wired Communications
SIGNAL INTEGRITY DEMO I:	Frequency Domain - VNA (Calibration and Channel Analysis)
SIGNAL INTEGRITY VI:	Unmasking The DUT: Fixture De-Embedding Essentials
POWER INTEGRITY I:	Power Integrity Fundamentals
POWER INTEGRITY II:	VRM Basics, VRM Modeling, and IR Drop
POWER INTEGRITY III:	Package and IC PDN, On-Chip VRM
POWER INTEGRITY DEMO I:	Low Impedance PDN Measurements
CLOSING SESSION:	Understanding Vias for Better SI and PI – 50+ Years of Physics-Based Modeling

Please view the presentation abstracts and instructor bios available on the Symposium website: 2026.emcsipi.org/global-sipi-university

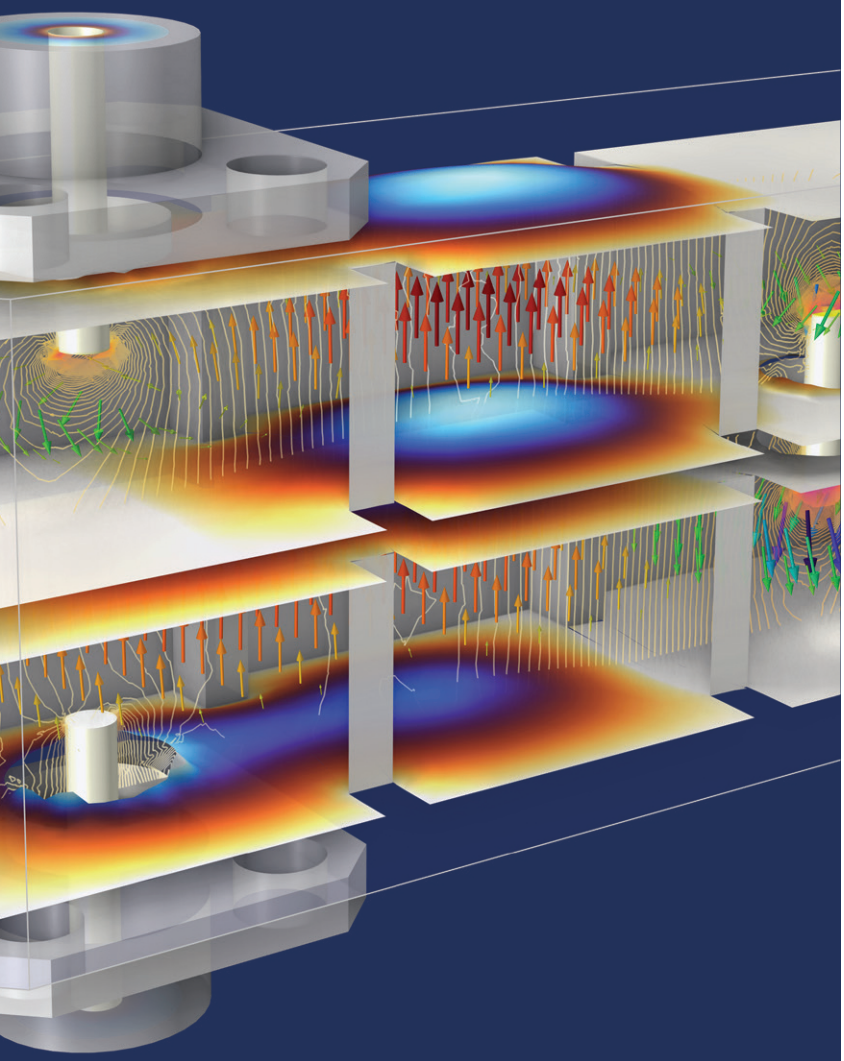
#IEEE_ESP26

2026.emcsipi.org

IEEE | EMC SOCIETY

COMSOL

Simulate real-world designs, devices, and processes with COMSOL Multiphysics®



Innovate faster.

Test more design iterations before prototyping.

Innovate smarter.

Analyze virtual prototypes and develop a physical prototype only from the best design.

Innovate with multiphysics simulation.

Base your design decisions on accurate results with software that lets you study unlimited multiple physical effects on one model.

SCAN ME TO LEARN MORE

comsol.com/feature/multiphysics-innovation